

Introducing the D2iQ Kubernetes Platform (DKP)

D2iQ Kubernetes Platform (DKP)

Exported on 03/29/2024

Table of Contents

1	Architecture.....	66
1.1	Learn the key concepts and architectural components of a DKP cluster .	66
1.2	Components for the Kubernetes Control Plane.....	66
1.3	Related Information.....	67
1.4	Next Steps.....	67
1.5	DKP Ports.....	67
1.5.1	Control-plane nodes.....	67
1.5.2	Worker nodes.....	68
1.5.3	Next Topic:.....	69
1.6	Supported Infrastructure Operating Systems.....	69
1.6.1	Amazon Web Services (AWS).....	70
1.6.2	Microsoft Azure.....	71
1.6.3	Google Cloud Platform (GCP).....	72
1.6.4	Pre-Provisioned.....	72
1.6.5	Pre-provisioned/Azure.....	74
1.6.6	vSphere.....	74
1.6.7	VMware Cloud Director (VCD) - Beta Product.....	76
1.6.8	Amazon Elastic Kubernetes (EKS).....	76
1.6.9	Azure Kubernetes Service (AKS).....	76
1.6.10	Next Step.....	77
2	Download DKP.....	78
2.1	Download from the Support Website.....	78
2.1.1	AWS Marketplace Only Download.....	78
2.1.2	Next Step:.....	78
3	Day 0 - Get Started with DKP.....	79
3.1	DKP Concepts and Terms.....	80
3.1.1	Next Topic:.....	81

3.1.2	Cluster Types and Concepts	81
3.1.2.1	Multi-cluster Environment	82
3.1.2.2	Single-cluster Environment	83
3.1.2.3	Other important concepts	83
3.1.2.4	Next Topic:	84
3.1.3	CAPI Concepts and Terms	84
3.1.3.1	Next Topic:	85
3.1.4	Air-gapped or Non-air-gapped Environment?	85
3.1.4.1	Air-gapped Environments	85
3.1.4.2	Non-air-gapped Environments	86
3.1.4.3	Next Step:	86
3.1.5	What is Pre-provisioned Infrastructure.....	86
3.1.5.1	Where is Pre-provisioned used?.....	86
3.1.5.2	Next Topic:	87
3.2	Provide Context for Commands with a kubeconfig File.....	87
3.2.1	Single-cluster Environment	88
3.2.1.1	Set the environment variable for all your operations.....	88
3.2.2	Multi-cluster Environment	88
3.2.2.1	Use a flag to reference the target cluster.....	88
3.2.3	Next Topic:	89
3.3	Storage	89
3.3.1	Ephemeral Storage	90
3.3.2	Persistent Volume.....	90
3.3.2.1	Persistent Volume Claim.....	91
3.3.2.2	Related Information	91
3.3.3	Next Step:	91
3.3.4	Default Storage Providers in DKP	91
3.3.4.1	Change or Manage Multiple StorageClasses.....	93
3.3.4.2	Driver Information	93

3.3.4.3	Related Information	96
3.3.5	Provision a Static Local Volume	97
3.3.5.1	Before you Begin	97
3.3.5.2	Provision the Cluster and a Volume	98
3.4	Resource Requirements	100
3.4.1	General Resource Requirements	100
3.4.2	Infrastructure Provider Specific	101
3.4.3	Kommander Component Requirements:.....	101
3.4.3.1	Next Topic:	101
3.4.4	Managed Cluster Requirements	101
3.4.4.1	Minimum Recommendation for Managed Clusters:	101
3.4.5	Management Cluster Application Requirements.....	103
3.4.5.1	Management cluster application minimum resources and storage requirements	103
3.4.6	Workspace Platform Application Defaults and Resource Requirements	105
3.4.6.1	Next Topic:	108
3.5	Install Overview	108
3.5.1	Prepare Your Environment for Install:.....	108
3.5.2	Next Step:	110
3.5.3	Install OS Version Compatibility	110
3.5.3.1	Amazon Web Services (AWS)	110
3.5.3.2	Microsoft Azure	112
3.5.3.3	Google Cloud Platform (GCP)	112
3.5.3.4	Pre-Provisioned.....	113
3.5.3.5	Pre-provisioned/Azure.....	114
3.5.3.6	vSphere.....	115
3.5.3.7	VMware Cloud Director (VCD) - Beta Product.....	116
3.5.3.8	Amazon Elastic Kubernetes (EKS).....	116
3.5.3.9	Azure Kubernetes Service (AKS).....	117

3.5.4	Install Related Version Information	117
3.5.4.1	Supported Kubernetes Version	117
3.5.4.2	Konvoy Image Builder	118
3.5.4.3	Storage Providers	118
3.5.4.4	Deploying a Cluster in FIPS mode.....	119
3.6	Prerequisites for Install	121
3.6.1	Prerequisites for the Konvoy Component	121
3.6.1.1	Non-air-gapped (all environments)	121
3.6.1.2	Air-gapped Only (additional requirements)	122
3.6.2	Prerequisites for Kommander Component	122
3.6.2.1	Non-air-gapped (all environments)	122
3.6.2.2	Air-gapped Only (additional requirements)	124
3.6.2.3	Next Step:	124
3.6.3	Container Runtimes	125
3.6.3.1	TIPS.....	125
4	Licenses.....	125
4.1	Buy a License	128
4.2	DKP Enterprise	130
4.2.1	Compatible Infrastructure	130
4.2.2	Platform Applications.....	130
4.2.3	Catalog Applications	130
4.2.4	Cluster Manager.....	130
4.2.5	Built-in GitOps.....	130
4.2.6	DKP Enterprise Multi-cluster UI.....	131
4.3	DKP Essential	132
4.3.1	Compatible Infrastructure	132
4.3.2	Platform Applications.....	132
4.3.3	Cluster Manager.....	132
4.3.4	Built-in GitOps.....	133

4.3.5	DKP Essential Single Cluster UI	133
4.4	DKP Government Advanced	134
4.4.1	Compatible Infrastructure	134
4.4.2	Platform Applications	134
4.4.3	Catalog Applications	134
4.4.4	Cluster Manager	134
4.4.5	Built-in GitOps	135
4.4.6	DKP Government Advanced Multi-cluster UI	135
4.4.7	Confirmed Stateside Support (CSS)	136
4.4.8	Military-Grade Security	136
4.5	DKP Government Essential	136
4.5.1	Compatible Infrastructure	137
4.5.2	Platform Applications	137
4.5.3	Cluster Manager	137
4.5.4	Built-in GitOps	138
4.5.5	DKP Government Essential Single Cluster UI	138
4.5.6	Confirmed Stateside Support (CSS)	138
4.5.7	Military-Grade Security	138
4.6	Add a DKP license	139
4.6.1	Prerequisites	139
4.6.2	Download the Container Image and Extract the Binaries	139
4.6.3	Obtain the Amazon Resource Number (ARN) from the AWS Marketplace UI	140
4.6.4	Enter License Information in the DKP UI	141
4.6.5	Enter a DKP License using kubectl	141
4.6.6	Activate a DKP Insights License Key	142
4.6.6.1	Prerequisites	142
4.6.6.2	Obtain a License Key	142
4.6.6.3	Enter License Information	142

4.6.6.4	Remove an Insights License	143
4.7	Remove a DKP license.....	143
4.7.1	Remove a License via the UI	143
4.7.2	Manually Remove a License using kubectl	143
4.8	MSP Program.....	145
4.8.1	Benefits.....	145
4.8.2	How does it work?	145
4.8.3	How do I become an MSP partner?	146
5	Day 1 - Basic Installs by Infrastructure.....	147
5.1	Scenario Based Installation Instructions.....	147
5.2	Pre-provisioned Install Options.....	147
5.3	AWS Install Options	149
5.4	EKS Install Options	149
5.5	vSphere Install Options.....	149
5.6	VMware Cloud Director Install Options	149
5.7	Azure Install Options.....	149
5.8	AKS Install Options	149
5.9	GCP Install Options.....	149
5.10	Pre-provisioned Install Options.....	149
5.10.1	Resources.....	149
5.10.2	Pre-provisioned Non-air-gapped Install	151
5.10.2.1	Next Step:	151
5.10.2.2	Pre-provisioned: Define Infrastructure	151
5.10.2.3	Pre-provisioned: Define Control Plane Endpoint.....	152
5.10.2.4	Pre-provisioned: Create a Management Cluster	155
5.10.2.5	Pre-provisioned: Configure MetalLB.....	158
5.10.2.6	Pre-provisioned: Install Kommander	161
5.10.2.7	Pre-provisioned: Verify Install and Log in to UI	164
5.10.2.8	Pre-provisioned: Create Managed Clusters Using the DKP CLI	166

5.10.3	Pre-provisioned Air-gapped Install.....	170
5.10.3.1	Next Step:	170
5.10.3.2	Pre-provisioned Air-gapped: Configure Environment	170
5.10.3.3	Pre-provisioned Air-gapped: Load the Registry.....	173
5.10.3.4	Pre-provisioned Air-gapped: Define Infrastructure	176
5.10.3.5	Pre-provisioned Air-gapped: Define Control Plane Endpoint.....	177
5.10.3.6	Pre-provisioned Air-gapped: Create a Management Cluster.....	179
5.10.3.7	Pre-provisioned Air-gapped: Configure MetalLB.....	183
5.10.3.8	Pre-provisioned Air-gapped: Install Kommander	185
5.10.3.9	Pre-provisioned Air-gapped: Verify Install and Log in to UI.....	188
5.10.3.10	Pre-provisioned Air-gapped: Create Managed Clusters Using the DKP CLI	191
5.10.4	Pre-provisioned FIPS Install	195
5.10.4.1	Next Step:	195
5.10.4.2	Pre-provisioned FIPS: Define Infrastructure.....	195
5.10.4.3	Pre-provisioned FIPS: Define Control Plane Endpoint	198
5.10.4.4	Pre-provisioned FIPS: Create a Management Cluster	199
5.10.4.5	Pre-provisioned FIPS: Configure MetalLB	204
5.10.4.6	Pre-provisioned FIPS: Install Kommander.....	205
5.10.4.7	Pre-provisioned FIPS: Verify Install and Log in to UI	209
5.10.4.8	Pre-provisioned FIPS: Create Managed Clusters Using the DKP CLI	211
5.10.5	Pre-provisioned FIPS Air-gapped Install.....	215
5.10.5.1	Next Step:	215
5.10.5.2	Pre-provisioned Air-gapped FIPS: Configure Environment.....	215
5.10.5.3	Pre-provisioned Air-gapped FIPS: Load the Registry	220
5.10.5.4	Pre-provisioned Air-gapped FIPS: Define Infrastructure	221
5.10.5.5	Pre-provisioned Air-gapped FIPS: Define Control Plane Endpoint.....	222
5.10.5.6	Pre-provisioned Air-gapped FIPS: Create a Management Cluster	225
5.10.5.7	Pre-provisioned Air-gapped FIPS: Configure MetalLB.....	229

5.10.5.8	Pre-provisioned Air-gapped FIPS: Install Kommander	230
5.10.5.9	Pre-provisioned Air-gapped FIPS: Verify Install and Log in to UI	233
5.10.5.10	Pre-provisioned Air-gapped FIPS: Create Managed Clusters Using the DKP CLI	234
5.10.6	Pre-provisioned with GPU Install.....	239
5.10.6.1	Next Step:	239
5.10.6.2	Pre-provisioned GPU: Nodepool Secrets and Overrides.....	240
5.10.6.3	Pre-provisioned GPU: Define Infrastructure	243
5.10.6.4	Pre-provisioned GPU: Define Control Plane Endpoint	243
5.10.6.5	Pre-provisioned GPU: Create a Management Cluster.....	246
5.10.6.6	Pre-provisioned GPU: Configure MetalLB	249
5.10.6.7	Pre-provisioned GPU: Install Kommander	251
5.10.6.8	Pre-provisioned GPU: Verify Install and Log in to UI.....	255
5.10.6.9	Pre-provisioned GPU: Create Managed Clusters Using the DKP CLI.....	258
5.10.7	Pre-provisioned Air-gapped with GPU Install	262
5.10.7.1	Next Step:	263
5.10.7.2	Pre-provisioned Air-gapped GPU: Configure Environment.....	263
5.10.7.3	Pre-provisioned Air-gapped GPU: Load the Registry	266
5.10.7.4	Pre-provisioned Air-gapped GPU: Nodepool Secrets and Overrides	268
5.10.7.5	Pre-provisioned Air-gapped GPU: Define Infrastructure.....	270
5.10.7.6	Pre-provisioned Air-gapped GPU: Define Control Plane Endpoint	272
5.10.7.7	Pre-provisioned Air-gapped GPU: Create a Management Cluster	274
5.10.7.8	Pre-provisioned Air-gapped GPU: Configure MetalLB	280
5.10.7.9	Pre-provisioned Air-gapped GPU: Install Kommander.....	282
5.10.7.10	Pre-provisioned Air-gapped GPU: Verify Install and Log in to UI	286
5.10.7.11	Pre-provisioned Air-gapped with GPU: Create Managed Clusters Using the DKP CLI.....	289
5.11	AWS Install Options	293
5.11.1	AWS Install	294
5.11.1.1	AWS Prerequisites	294

5.11.1.2	Next Step:	296
5.11.1.3	AWS: Create an Image.....	295
5.11.1.4	AWS: Create the Management Cluster	298
5.11.1.5	AWS: Install Kommander.....	301
5.11.1.6	AWS: Verify Install and Log in the UI	303
5.11.1.7	AWS: Create a Managed Cluster Using the DKP CLI	306
5.11.2	AWS Air-gapped Install	310
5.11.2.1	AWS Prerequisites	310
5.11.2.2	Next Step:	313
5.11.2.3	AWS Air-gapped: Create an Image.....	311
5.11.2.4	AWS Air-gapped: Load the Registry	313
5.11.2.5	AWS Air-gapped: Create the Management Cluster.....	314
5.11.2.6	AWS Air-gapped: Install Kommander	319
5.11.2.7	AWS Air-gapped: Verify Install and Log in the UI	321
5.11.2.8	AWS Air-gapped: Create a Managed Cluster Using the DKP CLI	323
5.11.3	AWS with FIPS Install	328
5.11.3.1	AWS Prerequisites	328
5.11.3.2	Next Step:	329
5.11.3.3	AWS FIPS: Create an Image	329
5.11.3.4	AWS FIPS: Create the Management Cluster	331
5.11.3.5	AWS FIPS: Install Kommander	335
5.11.3.6	AWS FIPS: Verify Install and Log in the UI.....	337
5.11.3.7	AWS FIPS: Create a Managed Cluster Using the DKP CLI.....	340
5.11.4	AWS Air-gapped FIPS Install	344
5.11.4.1	AWS Prerequisites	344
5.11.4.2	Next Step:	345
5.11.4.3	AWS Air-gapped FIPS: Create an Image.....	346
5.11.4.4	AWS Air-gapped FIPS: Load the Registry	347
5.11.4.5	AWS Air-gapped FIPS: Create the Management Cluster.....	349

5.11.4.6	AWS Air-gapped FIPS: Install Kommander.....	353
5.11.4.7	AWS Air-gapped FIPS: Verify Install and Log in the UI	355
5.11.4.8	AWS Air-gapped FIPS: Create a Managed Cluster Using the DKP CLI ...	358
5.11.5	AWS with GPU Install.....	363
5.11.5.1	AWS Prerequisites	363
5.11.5.2	GPU Prerequisites.....	364
5.11.5.3	Next Step:	364
5.11.5.4	AWS GPU: Use Node Label Automatic Configuration	364
5.11.5.5	AWS GPU: Create an Image	365
5.11.5.6	AWS GPU: Create the Management Cluster.....	368
5.11.5.7	AWS GPU: Install Kommander	371
5.11.5.8	AWS GPU: Verify Install and Log in the UI	375
5.11.5.9	AWS GPU: Create a Managed Cluster Using the DKP CLI.....	377
5.11.6	AWS Air-gapped with GPU Install.....	382
5.11.6.1	AWS Prerequisites	382
5.11.6.2	GPU Prerequisites.....	383
5.11.6.3	Next Step:	383
5.11.6.4	AWS Air-gapped GPU: Use Node Label Automatic Configuration	383
5.11.6.5	AWS Air-gapped GPU: Create an Image	384
5.11.6.6	AWS Air-gapped GPU: Load the Registry.....	386
5.11.6.7	AWS Air-gapped GPU: Create the Management Cluster	388
5.11.6.8	AWS Air-gapped GPU: Install Kommander.....	393
5.11.6.9	AWS Air-gapped GPU: Verify Install and Log in the UI.....	396
5.11.6.10	AWS Air-gapped GPU: Create a Managed Cluster Using the DKP CLI....	399
5.12	EKS Install Options	404
5.12.1	EKS Install.....	405
5.12.1.1	AWS prerequisites	405
5.12.1.2	Next Step:	406
5.12.1.3	EKS: Minimal User Permission for Cluster Creation.....	406

5.12.1.4	EKS: Cluster IAM Policies and Roles	409
5.12.1.5	EKS: Create an Image	414
5.12.1.6	EKS: Create an EKS Cluster	414
5.12.1.7	EKS: Grant Cluster Access	423
5.12.1.8	EKS: Retrieve kubeconfig for EKS Cluster	424
5.12.1.9	EKS: Attach a Cluster	427
5.13	vSphere Install Options.....	433
5.13.1	Overview	434
5.13.2	Next Step:	434
5.13.3	vSphere Prerequisites: All Installation Types.....	434
5.13.3.1	vSphere: Minimum User Permissions	435
5.13.3.2	vSphere: Storage Options.....	437
5.13.3.3	vSphere: VMware Prerequisites	437
5.13.4	vSphere Install.....	438
5.13.4.1	Further vSphere Prerequisites.....	439
5.13.4.2	Next Step:	439
5.13.4.3	vSphere: Image Creation Overview	439
5.13.4.4	vSphere: Create an Image	440
5.13.4.5	vSphere: Create a CAPI VM Template	440
5.13.4.6	vSphere: Create the Management Cluster	444
5.13.4.7	vSphere: Configure Metal LB.....	447
5.13.4.8	vSphere: Install Kommander	449
5.13.4.9	vSphere: Verify Install and Log in to the UI.....	451
5.13.4.10	vSphere: Create a Managed Cluster Using the DKP CLI.....	454
5.13.5	vSphere Air-gapped Install	459
5.13.5.1	Further vSphere Prerequisites.....	459
5.13.5.2	Next Step:	459
5.13.5.3	vSphere Air-gapped: Image Creation Overview.....	459
5.13.5.4	vSphere Air-gapped: Create an Image	460

5.13.5.5	vSphere Air-gapped: Load the Registry	461
5.13.5.6	vSphere Air-gapped: Create a CAPI VM Template.....	463
5.13.5.7	vSphere Air-gapped: Create the Management Cluster	467
5.13.5.8	vSphere Air-gapped: Configure Metal LB	471
5.13.5.9	vSphere Air-gapped: Install Kommander.....	473
5.13.5.10	vSphere Air-gapped: Verify Install and Log in to the UI	475
5.13.5.11	vSphere Air-gapped: Create a Managed Cluster Using the DKP CLI	477
5.13.6	vSphere FIPS Install.....	483
5.13.6.1	Further vSphere Prerequisites.....	483
5.13.6.2	Next Step:	483
5.13.6.3	vSphere FIPS: Image Creation Overview	483
5.13.6.4	vSphere FIPS: Create an Image.....	484
5.13.6.5	vSphere FIPS: Create a CAPI VM Template	485
5.13.6.6	vSphere FIPS: Create the Management Cluster.....	490
5.13.6.7	vSphere FIPS: Configure Metal LB.....	493
5.13.6.8	vSphere FIPS: Install Kommander	495
5.13.6.9	vSphere FIPS: Verify Install and Log in to the UI.....	497
5.13.6.10	vSphere FIPS: Create a Managed Cluster Using the DKP CLI	499
5.13.7	vSphere FIPS Air-gapped Install.....	505
5.13.7.1	Further vSphere Prerequisites.....	505
5.13.7.2	Next Step:	505
5.13.7.3	vSphere FIPS Air-gapped: Image Creation Overview	505
5.13.7.4	vSphere FIPS Air-gapped: Create an Image	506
5.13.7.5	vSphere FIPS Air-gapped: Load the Registry.....	507
5.13.7.6	vSphere FIPS Air-gapped: Create a CAPI VM Template	509
5.13.7.7	vSphere FIPS Air-gapped: Create the Management Cluster	513
5.13.7.8	vSphere Air-gapped FIPS: Configure Metal LB.....	517
5.13.7.9	vSphere FIPS Air-gapped: Install Kommander	519
5.13.7.10	vSphere FIPS Air-gapped: Verify Install and Log in to the UI.....	521

5.13.7.11	vSphere FIPS Air-gapped: Create a Managed Cluster Using the DKP CLI	523
5.14	VMware Cloud Director Install Options	529
5.14.1	Special Resource Requirements for VMware Cloud Director:	529
5.14.1.1	CPU and Memory	529
5.14.1.2	Storage	529
5.14.1.3	Next Step:	530
5.14.2	Cloud Director Service Provider	530
5.14.2.1	If Required - Further vSphere Prerequisites	530
5.14.2.2	Overview	530
5.14.2.3	Next Step:	531
5.15	Azure Install Options.....	531
5.15.1	Azure Install.....	532
5.15.1.1	Azure Prerequisites.....	532
5.15.1.2	Next Step:	533
5.15.1.3	Azure: Create Image	534
5.15.1.4	Azure: Create the Management Cluster	536
5.15.1.5	Azure: Install Kommander	538
5.15.1.6	Azure: Verify Install and Log in to the UI	540
5.15.1.7	Azure: Create a Managed Cluster Using the DKP CLI.....	543
5.16	AKS Install Options	547
5.16.1	AKS Install	547
5.16.1.1	DKP Prerequisites	548
5.16.1.2	AKS Prerequisites	548
5.16.1.3	Next Step:	550
5.16.1.4	AKS: Create Image	550
5.16.1.5	AKS: Create an AKS Cluster	551
5.16.1.6	AKS: Retrieve kubeconfig for AKS Cluster	555
5.16.1.7	AKS: Attach Cluster.....	558

5.17	GCP Install Options.....	563
5.17.1	GCP Install.....	563
5.17.1.1	GCP Prerequisites:.....	564
5.17.1.2	Next Step:.....	564
5.17.1.3	GCP: Create Image.....	564
5.17.1.4	GCP: Create the Management Cluster.....	567
5.17.1.5	GCP: Install Kommander.....	572
5.17.1.6	GCP: Verify Install and Log in the UI.....	574
5.17.1.7	GCP: Create a Managed Cluster Using the DKP CLI.....	576
6	Day 2 - Cluster Operations Management.....	582
6.1	Operations.....	582
6.1.1	Access Control.....	583
6.1.1.1	Centrally Manage Access Across Clusters.....	583
6.1.1.2	Special Limitation for Kommander Roles.....	583
6.1.1.3	Types of Access Control Objects.....	584
6.1.1.4	Related Information.....	586
6.1.1.5	Granting Access to Kubernetes and Kommander Resources.....	586
6.1.1.6	Onboarding Users onto a DKP Cluster.....	594
6.1.2	Identity Providers.....	597
6.1.2.1	Grant access to users in your organization.....	597
6.1.2.2	Benefits of Using an External Identity Provider.....	597
6.1.2.3	Prerequisites.....	597
6.1.2.4	Groups.....	598
6.1.2.5	Authorize a Group in Github.....	598
6.1.2.6	External LDAP directory.....	599
6.1.3	Infrastructure Providers.....	603
6.1.3.1	View and Modify Infrastructure Providers.....	604
6.1.3.2	Delete an infrastructure provider.....	604
6.1.3.3	Configure an AWS Infrastructure Provider with a User Role.....	604

6.1.3.4	Configure an AWS Infrastructure Provider with Static Credentials	611
6.1.3.5	Create an Azure Infrastructure Provider in the DKP UI	618
6.1.3.6	Create a vSphere Infrastructure Provider in the DKP UI	619
6.1.4	Add a Header, Footer, and Logo to Your DKP Implementation.....	619
6.1.4.1	Creating a Header Banner	620
6.1.4.2	Creating a Footer Banner.....	620
6.1.4.3	Adding Your Organization’s Logo to the Header	620
6.2	Applications	621
6.2.1	AppDeployment resources.....	622
6.2.1.1	Customization	622
6.2.1.2	Prerequisites	622
6.2.1.3	Customize Your Application.....	623
6.2.1.4	Print and Review the Current State of an AppDeployment Resource.....	624
6.2.1.5	Deployment Scope.....	625
6.2.1.6	More Information	625
6.2.2	Logging Stack Application Sizing Recommendations.....	625
6.2.3	Rook Ceph Cluster Sizing Recommendations	629
6.2.4	Manage an Application using the UI.....	633
6.2.4.1	Enterprise - Manage an Application using the UI.....	634
6.2.4.2	Essential - Manage an Application using the UI.....	637
6.2.5	Platform Applications.....	639
6.2.5.1	Related Topics	639
6.2.5.2	Deploy Platform Applications via CLI	639
6.2.5.3	Platform Application Dependencies	641
6.2.5.4	Workspace Platform Application Resource Requirements.....	647
6.2.6	Setting Priority Classes in DKP Applications	647
6.2.6.1	DKP Priority Classes.....	647
6.2.6.2	Prerequisites	648
6.2.6.3	Set the Priority Class	648

6.3	Workspaces.....	650
6.3.1	Global / Workspace UI.....	650
6.3.2	Default Workspace.....	650
6.3.3	Create a Workspace.....	650
6.3.4	Add, Edit, and Delete Workspace Annotations and Labels	650
6.3.5	Delete a Workspace	651
6.3.6	Workspace Applications.....	651
6.3.6.1	Cluster-scoped Application Configuration from the DKP UI	652
6.3.6.2	Cluster-scoped Configuration for Existing AppDeployments	654
6.3.7	Workspace Catalog Applications.....	663
6.3.7.1	Workspace DKP Catalog Applications	663
6.3.7.2	Deployment of Catalog Applications in Workspaces	671
6.3.7.3	Workspace Catalog Application Upgrades	674
6.3.7.4	Custom Applications	676
6.3.8	Workspace Role Bindings.....	685
6.3.8.1	Configure Workspace Role Bindings	685
6.4	Projects.....	686
6.4.1	Project Namespace	686
6.4.2	Create a Project	686
6.4.3	Create a Project - UI Method	687
6.4.4	Create a Project - CLI Method	687
6.4.5	Project Applications.....	688
6.4.5.1	Project Platform Applications	688
6.4.5.2	Project Catalog Applications.....	694
6.4.5.3	Project AppDeployments.....	713
6.4.6	Project Deployments	713
6.4.6.1	What is GitOps?.....	714
6.4.6.2	Continuous Delivery with GitOps.....	714
6.4.6.3	Continuous Deployment	724

6.4.6.4	Project Deployments Troubleshooting	727
6.4.6.5	View Helm Releases	728
6.4.7	Project Role Bindings.....	728
6.4.7.1	Configure Project Role Bindings - UI Method.....	728
6.4.7.2	Configure Project Role Bindings - CLI Method.....	728
6.4.7.3	Configure Project Role Bindings to Bind to WorkspaceRoles - CLI Method.....	729
6.4.7.4	Role Binding with VirtualGroup	731
6.4.8	Project Roles	732
6.4.8.1	Configure Project Role - UI Method	732
6.4.8.2	Configure Project Role - CLI Method	732
6.4.9	Project ConfigMaps	735
6.4.9.1	Configuring Project ConfigMaps - UI Method	735
6.4.9.2	Configuring Project ConfigMaps - CLI Method	735
6.4.10	Project Secrets.....	736
6.4.10.1	Configure Project Secrets - UI Method	736
6.4.10.2	Configure Project Secrets - CLI Method.....	737
6.4.11	Project Quotas & Limit Ranges	738
6.4.11.1	Creating Project Quotas & Limit Ranges- UI Method.....	738
6.4.11.2	Create Project Quotas & Limit Ranges - CLI Method.....	738
6.4.12	Project Network Policies	740
6.4.12.1	About Network Plugins.....	740
6.4.12.2	About Network Policies	741
6.4.12.3	Creating Network Policies	741
6.4.12.4	Network Policy Examples.....	742
6.5	Managing Clusters.....	745
6.5.1	Create a Managed Azure Cluster from the DKP UI	746
6.5.1.1	Prerequisites	746
6.5.1.2	Provision an Azure Cluster	746

6.5.2	Create a Managed Cluster on vSphere from the DKP UI	747
6.5.2.1	Prerequisites	747
6.5.2.2	Provision a vSphere Cluster	747
6.5.2.3	Advanced Configuration Parameters	751
6.5.2.4	Next Actions	752
6.5.3	Attach a Kubernetes Cluster	752
6.5.3.1	Attach Kubernetes Cluster	752
6.5.3.2	Requirements for Attaching an Existing Cluster	753
6.5.3.3	Create a kubeconfig File for Attachment	755
6.5.3.4	Attach a Cluster with no Networking Restrictions (UI)	758
6.5.3.5	Attach a Cluster with Networking Restrictions	772
6.5.3.6	Attach a DKP-created Kubernetes Cluster	824
6.5.3.7	Access a Managed or Attached Cluster	826
6.5.3.8	Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster	827
6.5.4	Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster	827
6.5.4.1	Prerequisites: General Prerequisites for your Cluster Conversion	828
6.5.4.2	Prerequisites: Prepare your Cluster's Existing Configurations	830
6.5.4.3	Prerequisites: Clone Git Repository from Gitea	831
6.5.4.4	Back up your Cluster's Applications and Persistent Volumes	832
6.5.4.5	UI: Convert an Essential Cluster into a Managed Cluster	842
6.5.4.6	CLI: Convert an Essential Cluster into a Managed Cluster	844
6.5.4.7	Post Conversion Cleanup: Cluster Autoscaler Configuration	845
6.5.4.8	Post Conversion Cleanup: Clusters run on Different Cloud Platforms	848
6.5.4.9	Troubleshooting	851
6.5.5	Advanced Creation of CLI Clusters	858
6.5.5.1	Generate Cluster Objects	858
6.5.5.2	Use the Upload YAML Form	858
6.5.6	Custom Domains and Certificates Configuration for All Cluster Types	858

6.5.6.1	Why Should You Set Up a Custom Domain or Certificate?	859
6.5.6.2	KommanderCluster and Certificate Issuer Concepts.....	860
6.5.6.3	Configure Custom Domains or Custom Certificates post Kommander Installation.....	861
6.5.7	Disconnect or Delete Clusters.....	867
6.5.7.1	Disconnect or delete a cluster	867
6.5.7.2	Disconnect vs. Delete	867
6.5.7.3	Troubleshooting.....	868
6.5.8	Management Cluster	868
6.5.8.1	Editing.....	869
6.5.8.2	Disconnecting	869
6.5.9	Cluster Statuses.....	869
6.5.10	Cluster Resources.....	870
6.5.11	DKP Platform Applications.....	871
6.5.12	Cluster Applications and Statuses.....	871
6.5.12.1	Custom Cluster Application Dashboard Cards	872
6.5.13	Kubernetes Cluster Federation (KubeFed).....	874
6.6	Backup and Restore.....	874
6.6.1	Configure Velero	874
6.6.1.1	Cloud Provider Storage.....	875
6.6.1.2	Use Velero with AWS	875
6.6.1.3	Use Velero with Azure.....	881
6.6.1.4	Use Velero with GCP	887
6.6.2	Back up with Velero	893
6.6.2.1	Install the Velero CLI.....	894
6.6.2.2	Regular Backup Operations.....	894
6.6.2.3	Restore a cluster from backup.....	897
6.6.2.4	Backup Service Diagnostics.....	899
6.7	Logging.....	899

6.7.1	Logging Operator	901
6.7.1.1	Fluent Bit Buffer Information.....	901
6.7.2	Admin-level logs.....	902
6.7.3	Enable Workspace-level Logging	902
6.7.3.1	How to enable Workspace-level Logging for use with DKP.	902
6.7.3.2	Logging Prerequisites and Architecture.....	903
6.7.3.3	Enable Logging Applications through the UI.....	904
6.7.3.4	Create AppDeployments to Enable Workspace Logging	905
6.7.3.5	Override ConfigMap to Restrict Logging to Specific Namespaces	907
6.7.3.6	Override ConfigMap to Modify the Storage Retention Period in Workspace Grafana Loki	909
6.7.3.7	Verify Cluster Logging Stack Installation	910
6.7.3.8	View Cluster Log Data	911
6.7.4	Multi-Tenant Logging Overview	913
6.7.4.1	Enable Multi-tenant Logging	915
6.7.4.2	Create a Project for Logging	916
6.7.4.3	Create Project-level Logging AppDeployments	916
6.7.4.4	Override ConfigMap to Modify the Storage Retention Period in Project Grafana Loki	917
6.7.4.5	Verify Project Logging Stack Installation	919
6.7.4.6	View Project Log Data	920
6.7.5	Fluent Bit.....	923
6.7.5.1	Audit Log Collection.....	923
6.7.5.2	Related Information	924
6.7.5.3	Collecting systemd Logs from a Non-default Path	924
6.7.6	Scaling the Logging Stack.....	929
6.7.6.1	Introduction.....	929
6.7.6.2	Logging Operator	929
6.7.6.3	Grafana Loki	930
6.7.6.4	Rook Ceph	930

6.7.6.5	Audit Log	931
6.7.7	Configuring Loki to use AWS S3 Storage in DKP.....	932
6.7.7.1	Configuring Loki.....	932
6.8	Security	934
6.8.1	Authentication and authorization architecture	934
6.8.1.1	Details on distributed authentication and authorization between clusters	934
6.8.1.2	Authentication	934
6.8.1.3	Authorization	935
6.8.2	OpenID Connect (OIDC) Introduction.....	935
6.8.2.1	An introduction to OpenID Connect (OIDC) Authentication in Kubernetes	935
6.8.2.2	Identity Provider	936
6.8.2.3	Add Login Connectors	937
6.8.2.4	Change the Access Token Lifetime	938
6.8.2.5	Authentication	938
6.8.3	SAML Connector	939
6.8.3.1	Connect your Kommander cluster to an IdP using SAML.....	939
6.8.3.2	Connect Kommander to an IdP Using SAML	939
6.8.4	Policy Controls	942
6.8.4.1	Workload Policy Controls	942
6.8.4.2	Enforce Policies Using Gatekeeper.....	942
6.8.5	Traefik-Forward-Authentication in DKP (TFA).....	946
6.8.5.1	TFA Overview	946
6.8.5.2	TFA Authentication Workflow	947
6.8.5.3	Default TFA Configuration in DKP.....	947
6.8.5.4	Cluster Storage Option	948
6.8.6	Create Local Access	948
6.8.6.1	Next Step:	949
6.8.6.2	Create Local Users during the Kommander Installation	949

6.8.6.3	Create Local Users after Installing Kommander	950
6.8.6.4	Add RBAC Roles to Local Users.....	952
6.8.6.5	Modify or Delete Local Users	953
6.9	Networking	954
6.9.1	Configure networking for Konvoy cluster.....	954
6.9.2	Networking Service.....	954
6.9.2.1	Service Topology	955
6.9.2.2	EndpointSlices	956
6.9.2.3	DNS for Services and Pods	956
6.9.2.4	Ingress and Networking.....	957
6.9.2.5	Network Policies	958
6.9.3	Required Domains.....	960
6.9.3.1	This section describes the required domains for DKP	960
6.9.4	Configure Ingress for load balancing	961
6.9.4.1	Learn how to configure Ingress settings for load balancing (layer-7)....	961
6.9.4.2	Prerequisites	961
6.9.4.3	Expose a pod using an Ingress (L7)	961
6.9.5	Ingress	963
6.9.5.1	Traefik Ingress Controller.....	963
6.9.5.2	Traefik v2.4.....	963
6.9.5.3	Related Information	964
6.9.6	Load Balancing	964
6.9.6.1	Load Balancing for Internal Traffic	964
6.9.6.2	Load Balancing for External Traffic	964
6.9.7	Use Istio as a Microservice	965
6.9.7.1	Prerequisites	965
6.9.7.2	Deploy Istio Using DKP	966
6.10	GPUs	968
6.10.1	Kommander GPU Overview	968

6.10.2	GPU Prerequisites	969
6.10.2.1	Configure GPU for Kommander clusters.....	969
6.10.2.2	Prerequisites	969
6.10.3	NVIDIA Platform Application Management Cluster	970
6.10.3.1	Enable NVIDIA Platform Application on Kommander for Management Cluster.....	970
6.10.3.2	Select the Correct Toolkit Version for your NVIDIA GPU Operator.....	971
6.10.4	NVIDIA Platform Application Attached or Managed Cluster	972
6.10.4.1	Enable NVIDIA Platform Application on Attached or Managed Clusters.....	972
6.10.4.2	Select the Correct Toolkit Version for your NVIDIA GPU Operator.....	973
6.10.5	Kommander GPU Settings and Troubleshoot.....	975
6.10.5.1	Validate and troubleshoot GPU.....	975
6.10.5.2	Validate that the Application has Started Correctly.....	975
6.10.5.3	NVIDIA GPU Monitoring.....	975
6.10.5.4	NVIDIA MIG Settings.....	976
6.10.5.5	Troubleshooting NVIDIA GPU Operator on Kommander.....	977
6.10.5.6	Disable NVIDIA GPU Operator Platform Application on Kommander	980
6.10.6	GPU Toolkit Versions.....	980
6.10.7	Enable GPU Usage After Installing DKP	981
6.11	Monitoring and Alerts	982
6.11.1	Recommendations.....	983
6.11.1.1	Prometheus.....	983
6.11.2	Grafana Dashboards.....	985
6.11.2.1	Add Custom Dashboards	986
6.11.3	Cluster Metrics.....	988
6.11.4	Configure Alerts Using AlertManager.....	988
6.11.4.1	Prerequisites	989
6.11.5	Centralized Monitoring	994
6.11.5.1	Centralized Metrics.....	995

6.11.5.2	Centralized Alerts.....	997
6.11.6	Centralized Cost Monitoring	998
6.11.6.1	Adding a Kubecost License Key to your Clusters	999
6.11.6.2	Centralized Costs.....	1000
6.11.6.3	Related Information	1001
6.11.7	Monitor Applications using Prometheus.....	1001
6.11.8	Set Storage Capacity for Prometheus	1003
6.12	Storage for Applications.....	1004
6.12.1	Rook Ceph in DKP	1004
6.12.1.1	Rook Ceph in DKP - Prerequisites	1005
6.12.1.2	Rook Ceph Configuration	1006
6.12.1.3	Rook Ceph Dashboard	1010
6.12.2	BYOS (Bring Your Own Storage) to DKP Clusters.....	1011
6.12.2.1	Disable DKP Managed Ceph	1011
6.12.2.2	Creating DKP Compatible Ceph Resources	1012
6.12.2.3	Configure Loki to use S3 Compatible Storage	1016
6.12.2.4	Overriding velero and grafana-loki Configuration	1017
6.12.2.5	Overriding project-grafana-loki Configuration.....	1018
6.13	DKP Troubleshooting.....	1019
6.13.1	Generate a Support Bundle	1020
6.13.1.1	Prerequisites	1020
6.13.1.2	Diagnostic Bundle	1020
6.13.1.3	Generate a Support Bundle	1021
6.13.1.4	SSH Fallback	1022
6.13.2	Custom Collectors	1024
6.13.2.1	Customizations	1024
6.14	Telemetry.....	1028
6.14.1	Information collected by D2iQ	1028
6.14.2	Data Storage Location.....	1029

6.14.3	Size Limitations.....	1029
6.14.4	Used Collectors.....	1029
6.14.4.1	Extended Bundle Collectors	1029
6.14.4.2	Telemetry Bundle Collectors	1030
6.14.5	Sending Telemetry Data for Analysis	1030
6.14.5.1	Prerequisites	1030
6.14.5.2	Generate and Send a Bundle	1031
6.14.5.3	Recommended Telemetry Upload Frequency.....	1032
6.14.6	Scheduling Telemetry Upload	1032
6.14.6.1	Prerequisites	1032
6.14.6.2	Create and Enable the CronJob	1032
6.14.6.3	GitHub Actions.....	1034
7	Custom Installation and Additional Infrastructure Tools	1036
7.1	Universal Configurations for all Infrastructure Providers.....	1036
7.2	AWS Infrastructure	1036
7.3	EKS Infrastructure.....	1036
7.4	Azure Infrastructure.....	1037
7.5	AKS Infrastructure	1037
7.6	Pre-provisioned Infrastructure	1037
7.7	vSphere Infrastructure.....	1038
7.8	VMware Cloud Director Infrastructure.....	1038
7.9	GCP Infrastructure	1038
7.10	Universal Configurations for all Infrastructure Providers.....	1038
7.10.1	Additional Configurations.....	1039
7.10.2	Container Runtime Engine (CRE)	1039
7.10.3	Configuring an HTTP/HTTPS Proxy	1039
7.10.3.1	Configure the Bootstrap Cluster HTTP Proxy Settings	1040
7.10.3.2	Create CAPI Components with HTTP/HTTPS Proxy Settings	1041
7.10.3.3	Create a Cluster with HTTP/HTTPS Proxy.....	1042

7.10.3.4	Configure an HTTP/HTTPS Proxy for the DKP Kommander Component.....	1044
7.10.4	Working with Load Balancers.....	1045
7.10.4.1	External Load Balancers.....	1045
7.10.4.2	Select your Connection Mechanism.....	1045
7.10.4.3	External Load Balancer for the DKP Kommander Component	1046
7.10.5	Registry and Registry Mirrors.....	1046
7.10.6	Subnets and Pods.....	1046
7.10.6.1	Next Step:	1048
7.11	AWS Infrastructure	1048
7.11.1	Configuration Types	1048
7.11.2	AWS Diagrams	1049
7.11.3	AWS Pricing Considerations	1050
7.11.4	AWS Service Limits.....	1050
7.11.5	AWS Prerequisites	1051
7.11.5.1	Konvoy Prerequisites	1051
7.11.5.2	Control Plane Nodes.....	1051
7.11.5.3	Worker Nodes.....	1052
7.11.5.4	AWS Prerequisites	1052
7.11.6	AWS Konvoy Image Builder.....	1053
7.11.6.1	Learn how to build a custom AMI for use with DKP	1053
7.11.6.2	Prerequisites	1053
7.11.6.3	Extract KIB Bundle	1054
7.11.6.4	1. Create a Custom AMI	1054
7.11.6.5	2. Air-gapped AMI	1056
7.11.6.6	Prerequisites	1056
7.11.6.7	Next Step:	1057
7.11.7	Minimal Permissions and Role to Create Clusters	1057
7.11.7.1	Prerequisites	1057

7.11.7.2	Minimal Permissions	1057
7.11.7.3	Create Resources in CloudFormation Stack	1057
7.11.7.4	Leverage the Role	1061
7.11.7.5	Use EC2 Instance Profiles	1062
7.11.7.6	Use Access Keys.....	1062
7.11.7.7	Next Step:	1063
7.11.8	Cluster IAM Policies, Roles, and Artifacts	1063
7.11.8.1	Prerequisites	1063
7.11.8.2	Instance Profiles	1064
7.11.8.3	Next Steps:	1071
7.11.9	AWS Cluster Creation Choices.....	1071
7.11.9.1	Custom AMI in Cluster Creation	1071
7.11.9.2	AWS Non-air-gapped Create a New Cluster	1073
7.11.9.3	AWS Air-gapped Create a New Cluster.....	1089
7.11.9.4	Explore New AWS Cluster	1097
7.11.9.5	Make the AWS Cluster Self-Managed.....	1098
7.11.9.6	Creating DKP Clusters on AWS from the DKP UI.....	1101
7.11.10	Delete an AWS Cluster.....	1102
7.11.10.1	Create a Bootstrap Cluster and Move CAPI Resources	1103
7.11.10.2	Delete the Workload Cluster.....	1105
7.11.10.3	Delete the Bootstrap Cluster	1106
7.11.10.4	Known Limitations	1106
7.11.10.5	Next Topic:	1107
7.11.11	Multiple AWS Accounts	1107
7.11.11.1	Objective.....	1107
7.11.11.2	Assumptions	1107
7.11.11.3	Glossary.....	1107
7.11.11.4	Prerequisites	1107
7.11.11.5	Deploy DKP on AWS.....	1108

7.11.11.6 Next Steps:	1110
7.11.12 GPUs in an AWS environment	1110
7.11.12.1 Understanding GPUs.....	1110
7.11.12.2 Install GPU Support for Supported Distributions on AWS.....	1110
7.11.12.3 Configure Konvoy Automatic GPU Node Labels.....	1111
7.11.12.4 Update NVIDIA GPU Clusters	1112
7.11.13 Manage AWS Node Pools	1113
7.11.13.1 Create AWS Node Pools.....	1114
7.11.13.2 List AWS Node Pools.....	1115
7.11.13.3 Scale AWS Node Pools.....	1116
7.11.13.4 Delete AWS Node Pools	1118
7.11.13.5 AWS Cluster Autoscaler	1119
7.11.14 AWS Certificate Renewal.....	1121
7.11.14.1 Prerequisites	1122
7.11.14.2 Create a Cluster with Automated Certificate Renewal	1122
7.11.14.3 Technical Details.....	1122
7.11.14.4 Debug.....	1123
7.11.15 Configure Infrastructure in UI.....	1123
7.11.15.1 Create Infrastructure Provider.....	1124
7.11.15.2 Fill out the Add Infrastructure Provider Form	1124
7.11.16 Replace an AWS Node	1124
7.11.16.1 Prerequisites	1124
7.11.16.2 Replace a Worker Node	1125
7.12 EKS Infrastructure.....	1127
7.12.1 EKS Introduction	1128
7.12.2 Minimal User Permission for EKS Cluster Creation.....	1129
7.12.2.1 Next Step:	1132
7.12.3 EKS Cluster IAM Permissions and Roles.....	1132
7.12.3.1 Prerequisites from AWS:	1132

7.12.3.2	EKS IAM Artifacts.....	1132
7.12.4	Create an EKS Cluster from the CLI.....	1137
7.12.4.1	Create a New EKS Kubernetes Cluster	1137
7.12.4.2	Known Limitations	1144
7.12.5	Create an EKS Cluster from the DKP UI.....	1144
7.12.5.1	Create an AWS Infrastructure Provider	1144
7.12.5.2	Provision an EKS Cluster	1145
7.12.5.3	Access EKS Cluster	1145
7.12.5.4	IAM User and Role Access for EKS Clusters.....	1146
7.12.6	Grant Cluster Access	1147
7.12.6.1	How to Grant EKS Cluster Access	1147
7.12.7	Explore EKS Cluster	1148
7.12.7.1	Explore the new Kubernetes cluster	1148
7.12.8	Manage EKS Nodepools.....	1151
7.12.8.1	Create a node pool.....	1151
7.12.8.2	Scaling Up Node Pools	1152
7.12.8.3	Delete EKS Node Pools.....	1152
7.12.9	Delete EKS Cluster from CLI.....	1153
7.12.9.1	Delete the EKS cluster	1153
7.12.9.2	Next Step:	1154
7.12.9.3	Known Limitations	1154
7.12.10	Delete EKS Cluster from the DKP UI	1154
7.13	Azure Infrastructure.....	1156
7.13.1	Azure Prerequisites.....	1156
7.13.1.1	DKP Prerequisites	1156
7.13.1.2	Azure Prerequisites.....	1157
7.13.2	Azure using Konvoy Image Builder	1159
7.13.2.1	Prerequisites	1159
7.13.2.2	Extract the KIB Bundle	1159

7.13.2.3	Configure Azure Prerequisites	1160
7.13.2.4	Next Step:	1161
7.13.3	Azure Bootstrap	1161
7.13.3.1	Prepare to deploy Kubernetes clusters	1161
7.13.3.2	Prerequisites	1161
7.13.3.3	Bootstrap Cluster Lifecycle Services	1162
7.13.3.4	(Optional) Create identity secret for Azure	1163
7.13.3.5	Next Step:	1163
7.13.4	Create a New Custom Azure Cluster	1163
7.13.4.1	Prerequisites	1163
7.13.4.2	Name your cluster	1164
7.13.4.3	Tips and Tricks	1164
7.13.4.4	Create a new Azure Kubernetes cluster	1165
7.13.4.5	Known Limitations	1172
7.13.5	Explore new Azure Cluster	1173
7.13.5.1	Learn to interact with your Kubernetes cluster	1173
7.13.5.2	Explore the new Kubernetes cluster	1173
7.13.6	Azure Make new Cluster Self-Managed	1177
7.13.6.1	Make the new Kubernetes cluster manage itself	1177
7.13.6.2	Known Limitations	1179
7.13.7	Azure Certificate Renewal	1179
7.13.7.1	Configure Automated Renewal for Managed Kubernetes PKI Certificates	1179
7.13.7.2	Requirements	1180
7.13.7.3	Prerequisites	1180
7.13.8	Azure Replace a Node	1182
7.13.8.1	Replace a worker node	1182
7.13.8.2	Prerequisites	1182
7.13.8.3	Replace a worker node	1182

7.13.9	Azure Delete Cluster	1185
7.13.9.1	Delete the Kubernetes cluster and clean up your environment	1185
7.13.9.2	Prepare to Delete a Workload Cluster	1185
7.13.9.3	Delete the Workload Cluster.....	1187
7.13.9.4	Delete the Bootstrap Cluster	1188
7.13.9.5	Next Step:	1188
7.13.9.6	Known Limitations	1188
7.14	AKS Infrastructure	1189
7.14.1	Create a New AKS Cluster	1189
7.14.1.1	Use DKP to create a new AKS cluster.....	1189
7.14.1.2	Name Your Cluster.....	1190
7.14.1.3	Create a New AKS Kubernetes Cluster.....	1190
7.14.1.4	Inspecting or Editing the Cluster Objects	1191
7.14.1.5	Known Limitations	1193
7.14.2	Explore New AKS Cluster	1193
7.14.2.1	Learn to interact with your AKS Kubernetes cluster	1193
7.14.2.2	Explore the New AKS Cluster	1193
7.14.3	Delete AKS Cluster	1197
7.14.3.1	Delete the AKS cluster and clean up your environment	1197
7.14.3.2	Delete the Workload Cluster.....	1197
7.14.3.3	Next Step:	1198
7.14.3.4	Known Limitations	1198
7.14.4	Create a new AKS Cluster from the DKP UI.....	1198
7.14.4.1	Prerequisites	1198
7.14.4.2	Provision an AKS Cluster.....	1199
7.14.4.3	Access an AKS Cluster	1200
7.15	Pre-provisioned Infrastructure	1200
7.15.1	Create a Kubernetes cluster on pre-provisioned nodes in a bare metal infrastructure.....	1200

7.15.2	Advanced Workflow	1201
7.15.3	Pre-provisioned Prerequisite Configuration	1201
7.15.3.1	Prerequisites for a Pre-provisioned infrastructure	1201
7.15.3.2	Machine Specifications	1202
7.15.3.3	Next Step:	1203
7.15.4	Pre-provisioned Air-gapped Define Environment	1204
7.15.4.1	Fulfill the prerequisites for using a pre-provisioned infrastructure when Air-Gapped	1204
7.15.4.2	Air-Gapped Registry Prerequisites	1204
7.15.4.3	Load the bootstrap image	1204
7.15.4.4	Copy air-gapped artifacts onto cluster hosts	1205
7.15.5	Pre-provisioned Set Infrastructure	1209
7.15.5.1	Define your infrastructure	1209
7.15.6	Pre-provisioned Define Control Plane Endpoint	1211
7.15.6.1	Define the Control Plane Endpoint for your cluster	1211
7.15.6.2	External Load Balancer	1212
7.15.6.3	Built-in virtual IP	1212
7.15.6.4	Single-Node control plane	1212
7.15.6.5	Known limitations	1212
7.15.7	Pre-provisioned Create Secrets and Overrides	1213
7.15.7.1	Create necessary secrets and overrides for pre-provisioned clusters	1213
7.15.7.2	Create a unique cluster name	1213
7.15.7.3	Create a secret	1213
7.15.7.4	Create Overrides	1214
7.15.7.5	Next Step:	1215
7.15.7.6	Pre-provisioned FIPS Create Secrets and Overrides	1215
7.15.8	Pre-provisioned Bootstrap Cluster	1219
7.15.8.1	Bootstrap a kind cluster and CAPI controllers	1219
7.15.8.2	Next Step:	1219

7.15.9	Pre-provisioned Create a New Cluster.....	1220
7.15.9.1	Create a Kubernetes cluster using the infrastructure definition.....	1220
7.15.9.2	Audit logs.....	1222
7.15.9.3	Pre-provisioned Air-gapped New Cluster	1223
7.15.10	Pre-provisioned Azure only Configurations.....	1227
7.15.10.1	Prerequisites:	1228
7.15.10.2	Set Environment Variables with Credentials:	1228
7.15.10.3	Next Step:	1231
7.15.11	Pre-provisioned Modify the Calico Installation	1231
7.15.11.1	Set the Interface.....	1231
7.15.11.2	Change the Encapsulation Type	1234
7.15.11.3	Provider Specific Settings	1234
7.15.12	Pre-provisioned Built-in Virtual IP	1236
7.15.12.1	Virtual IP example.....	1236
7.15.13	Provision on the Flatcar Linux OS.....	1236
7.15.13.1	Flatcar Linux example.....	1236
7.15.14	Pre-provisioned Use HTTP Proxy.....	1237
7.15.14.1	HTTP proxy example	1238
7.15.15	Pre-provisioned Use Alternate Pod or Service Subnets	1238
7.15.16	Pre-provisioned: Make Cluster Self-managed.....	1240
7.15.16.1	Make the new Kubernetes cluster manage itself	1240
7.15.17	Pre-provisioned Configure MetalLB.....	1243
7.15.17.1	Create a MetalLB configmap for your pre-provisioned infrastructure..	1243
7.15.17.2	Layer 2 Configuration	1243
7.15.17.3	BGP Configuration	1244
7.15.18	Pre-provisioned Create and Delete Node Pools.....	1245
7.15.18.1	Create a Pre-provisioned node pool	1245
7.15.18.2	Delete a node pool	1246
7.15.19	Pre-provisioned Add Nodes to Existing Node Pool	1247

7.15.19.1 Prerequisites	1247
7.15.19.2 Scale up a Cluster Node	1247
7.15.19.3 Scale Down a Cluster Node.....	1248
7.15.20 GPU Nodepools in a Pre-provisioned Environment	1249
7.15.21 Pre-provisioned Delete Cluster.....	1251
7.15.21.1 Prepare to Delete the Pre-provisioned Cluster.....	1251
7.15.21.2 Delete the Workload Cluster.....	1253
7.16 vSphere Infrastructure	1254
7.16.1 Creating DKP clusters in a VMware vSphere environment	1254
7.16.2 Next Step:	1255
7.16.3 vSphere Prerequisites.....	1256
7.16.3.1 Prepare your environment to run DKP with VMware vSphere.....	1256
7.16.3.2 DKP Prerequisites	1256
7.16.3.3 VMware vSphere Prerequisites.....	1257
7.16.3.4 Next Steps:	1258
7.16.3.5 Minimum User Permissions.....	1258
7.16.3.6 vSphere Storage Options.....	1260
7.16.4 Create a Base OS image in vSphere	1260
7.16.4.1 Create the Base OS Image	1261
7.16.4.2 Next Step:	1262
7.16.5 Create a VM Template.....	1262
7.16.5.1 Create a vSphere Template for Your Cluster from a Base OS Image...	1262
7.16.5.2 Specific to VMware Cloud Director.....	1264
7.16.6 vSphere Bootstrap	1266
7.16.6.1 Prepare to deploy Kubernetes clusters	1266
7.16.6.2 Prerequisites	1266
7.16.6.3 Bootstrap cluster lifecycle services.....	1266
7.16.6.4 Next Step:	1268
7.16.7 Create new vSphere Cluster	1268

7.16.7.1	Prerequisites	1268
7.16.7.2	Name your cluster.....	1268
7.16.7.3	Create a New vSphere Kubernetes Cluster	1268
7.16.7.4	Known Limitations	1275
7.16.7.5	Next Steps:	1275
7.16.8	Explore a vSphere Cluster	1275
7.16.8.1	Get the kubeconfig file for the new Kubernetes cluster	1275
7.16.8.2	Create a StorageClass with a vSphere datastore	1276
7.16.8.3	Explore nodes and pods in the new cluster	1276
7.16.9	Make vSphere Cluster Self-Managed	1279
7.16.9.1	Make the new Kubernetes cluster manage itself	1280
7.16.9.2	Known limitations	1282
7.16.10	Configure MetalLB for a vSphere infrastructure	1282
7.16.10.1	Create a MetalLB configmap for your vSphere infrastructure.....	1282
7.16.10.2	Layer 2 configuration.....	1282
7.16.10.3	BGP configuration.....	1283
7.16.11	Install vSphere Air-Gapped	1285
7.16.11.1	Create a Kubernetes vSphere cluster in a private network with no access to the Internet (air-gapped)	1285
7.16.11.2	Bastion Host.....	1285
7.16.11.3	Create a Base Air-gapped OS VM Image.....	1288
7.16.11.4	Create an Air-gapped CAPI VM Template	1289
7.16.11.5	vSphere Air-gapped Seed the Registry	1292
7.16.11.6	vSphere Air-gapped Bootstrap	1293
7.16.11.7	Create a new Air-gapped vSphere Cluster	1294
7.16.11.8	Make vSphere Air-gapped Cluster Self-Managed	1297
7.16.11.9	Explore vSphere Air-gapped Cluster	1301
7.16.12	vSphere Certificate Renewal	1305
7.16.12.1	Configure Automated Renewal for Managed Kubernetes PKI Certificates.....	1305

7.16.12.2 Certificate Renewal.....	1305
7.16.12.3 Prerequisites	1305
7.16.13 Delete vSphere Cluster	1307
7.16.13.1 Prepare to delete a self-managed workload cluster.....	1307
7.16.13.2 Delete the Workload Cluster.....	1310
7.16.13.3 Delete the Bootstrap Cluster	1311
7.16.13.4 Next Step:	1311
7.16.13.5 Known limitations	1311
7.16.14 Manage vSphere Node Pools.....	1311
7.16.14.1 Create vSphere Node Pools	1312
7.16.14.2 List vSphere Node Pools	1313
7.16.14.3 Scale vSphere Node Pools	1313
7.16.14.4 vSphere Cluster Autoscaler	1316
7.16.14.5 Replace a vSphere Node	1319
7.16.14.6 Delete vSphere Node Pools.....	1321
7.17 VMware Cloud Director Infrastructure.....	1322
7.17.1 Special Resource Requirements for VCD.....	1322
7.17.1.1 VM Sizing Policy - CPU and Memory	1322
7.17.1.2 VM Placement Policy.....	1322
7.17.1.3 Storage	1323
7.17.2 Next Step:	1323
7.17.3 Cloud Director for Service Providers	1323
7.17.3.1 Related Topics:	1325
7.17.3.2 Helpful VMWare Resources:	1325
7.17.3.3 Next Step:	1325
7.17.3.4 Cloud Director Prerequisites	1325
7.17.3.5 Cloud Director Configure the Organization	1329
7.17.3.6 Cloud Director Roles, Rights and Rights Bundles	1332
7.17.3.7 Cloud Director CAPVCD User Rights	1335

7.17.3.8	Cloud Director Create Image and Template.....	1339
7.17.3.9	Cloud Director Bootstrap the Cluster.....	1340
7.17.3.10	Cloud Director Create a New Cluster.....	1342
7.17.3.11	Cloud Director Explore the Cluster.....	1346
7.17.3.12	Cloud Director Make the New Cluster Self-Managed.....	1349
7.17.3.13	Cloud Director Install Kommander.....	1352
7.17.3.14	Cloud Director Manage Node Pools.....	1354
7.17.3.15	Cloud Director Delete a Cluster.....	1361
7.18	GCP Infrastructure.....	1365
7.18.1	GCP Prerequisites.....	1365
7.18.1.1	Prerequisites.....	1365
7.18.1.2	Control plane nodes.....	1365
7.18.1.3	Worker nodes.....	1366
7.18.1.4	GCP Prerequisite Roles.....	1366
7.18.2	GCP Konvoy Image Builder.....	1368
7.18.2.1	Prerequisites.....	1368
7.18.2.2	GCP Prerequisite Roles.....	1369
7.18.3	Bootstrap GCP.....	1371
7.18.3.1	Prerequisites.....	1371
7.18.3.2	Bootstrap Cluster Lifecycle Services.....	1372
7.18.4	Create a New GCP Cluster.....	1373
7.18.4.1	Name your cluster.....	1373
7.18.4.2	Create a new GCP cluster.....	1374
7.18.5	Explore the GCP Cluster.....	1376
7.18.5.1	Explore the new Kubernetes cluster.....	1377
7.18.6	Make the New GCP Cluster Self-Managed.....	1380
7.18.6.1	Make the new Kubernetes cluster manage itself.....	1380
7.18.6.2	Known Limitations.....	1382
7.18.7	Manage GCP Node Pools.....	1382

7.18.7.1	Create GCP Node Pools.....	1383
7.18.7.2	List GCP Node Pools	1384
7.18.7.3	Scale GCP Node Pools	1384
7.18.7.4	Delete GCP Node Pools.....	1387
7.18.7.5	GCP Cluster Autoscaler.....	1387
7.18.8	Delete a GCP Cluster	1390
7.18.8.1	Prepare to Delete a Workload Cluster	1390
7.18.8.2	Delete the Workload Cluster.....	1392
7.18.8.3	Delete the Bootstrap Cluster	1392
7.18.8.4	Known Limitations	1393
8	Additional Kommander Configuration.....	1394
8.1	Installing Kommander by Environment	1394
8.2	Kommander Customizations	1394
8.3	Installing Kommander by Environment	1394
8.3.1	1. Is your Environment Air-gapped or Non-Air-gapped?	1394
8.3.2	2. What License do you Have?	1394
8.3.3	3. Do you Want to Enable DKP Insights?	1395
8.3.4	Next Step:	1395
8.3.5	Identify or Change your Default StorageClass	1395
8.3.5.1	Identify your StorageClass	1396
8.3.5.2	Next Step:	1396
8.3.6	Install Kommander According to your Environment.....	1396
8.3.6.1	Install Kommander in an Air-gapped Environment	1397
8.3.6.2	Install Kommander in a Non-air-gapped Environment.....	1402
8.3.6.3	Install Kommander in a Pre-provisioned, Air-gapped Environment	1405
8.3.6.4	Install Kommander in a Pre-provisioned, Non-air-gapped Environment	1412
8.3.6.5	Install Kommander in a Small Environment.....	1416
8.3.7	Verify Kommander Installation	1419

8.3.7.1	Failed HelmReleases	1420
8.3.7.2	Next Step:	1420
8.3.8	Log in to the UI with Kommander	1420
8.3.8.1	Dashboard UI Functions	1421
8.3.8.2	Next Step	1421
8.4	Kommander Customizations	1422
8.4.1	Prerequisites	1422
8.4.2	Initialize a Kommander Installer Configuration File	1422
8.4.3	Configure Applications	1422
8.4.3.1	Inline configuration (using values)	1422
8.4.3.2	Reference another YAML file (using valuesFrom)	1423
8.4.4	Minimal Kommander Installation.....	1423
8.4.5	Install with Configuration File	1424
8.4.6	Verify Installation	1424
8.4.7	DKP Kommander Configuration Reference.....	1424
8.4.7.1	Configuration Parameters	1424
8.4.7.2	AppConfig.....	1427
8.4.7.3	IngressCertificate.....	1428
8.4.7.4	Airgapped	1428
8.4.8	Custom Domains and Certificates for Management and Essential Clusters.....	1428
8.4.8.1	Why to set up a Custom Domain or Certificate?	1429
8.4.8.2	Certificate Issuer and KommanderCluster Concepts.....	1430
8.4.8.3	Certificate Authority (CA) Specifics.....	1431
8.4.8.4	Configure the Kommander Installation with a Custom Domain and Certificate	1432
8.4.8.5	Verify and Troubleshoot the Domain and Certificate Customization... ..	1441
8.4.9	DNS Record Creation with External DNS	1442
8.4.9.1	Configure External DNS with the CLI: Management or Essential Cluster.....	1443

8.4.9.2	Configure External DNS with the UI: All Clusters	1445
8.4.9.3	Verify your External DNS Configuration.....	1447
8.4.10	External Load Balancer	1450
8.4.10.1	Load Balancing for External Traffic in DKP	1450
8.4.10.2	Configure Kommander to use an External Load Balancer	1451
8.4.10.3	Configure the External Load Balancer to Target the Specified Ports...	1452
8.4.11	Installing Kommander with an HTTP Proxy	1452
8.4.11.1	Prerequisites	1453
8.4.11.2	Enable Gatekeeper	1453
8.4.11.3	Create Gatekeeper ConfigMap in the Kommander Namespace	1454
8.4.11.4	HTTP Proxy Configuration Considerations	1455
8.4.11.5	Install Kommander.....	1456
8.4.11.6	Configure Workspace or Project.....	1456
8.4.11.7	Configure HTTP Proxy in Attached Clusters	1456
8.4.11.8	Create Gatekeeper ConfigMap in the Workspace Namespace	1457
8.4.11.9	Configure Your Applications	1458
8.4.11.10	Manually Configure Your Application	1458
8.4.12	Enable DKP Catalog Applications after Installing DKP.....	1458
8.4.12.1	Configure a Default Enterprise Catalog after Installing DKP	1459
8.4.12.2	DKP Catalog Application labels	1460
8.4.12.3	Next Step:	1460
9	Additional Konvoy Configurations	1461
9.1	Section Contents.....	1461
9.2	Konvoy Image Builder	1461
9.2.1	How KIB Works	1462
9.2.1.1	Using Override files.....	1462
9.2.1.2	Using HTTP/S Proxies	1462
9.2.2	Prerequisites	1463
9.2.2.1	Compatible DKP to KIB Versions	1463

9.2.2.2	Next Steps:	1464
9.2.3	KIB with AWS.....	1465
9.2.3.1	Minimal IAM Permissions for KIB.....	1465
9.2.3.2	AWS Image Integrated with DKP CLI.....	1468
9.2.3.3	Create a Custom AMI	1469
9.2.3.4	AWS Air-gapped AMI	1474
9.2.3.5	Add Custom Tags to Image	1475
9.2.3.6	KIB for EKS	1476
9.2.4	KIB for Azure	1477
9.2.4.1	Learn how to build a custom Azure Image for use with DKP.....	1477
9.2.4.2	Prerequisites	1477
9.2.4.3	Extract the KIB Bundle.....	1477
9.2.4.4	Configure Azure Prerequisites	1478
9.2.4.5	Build the Image	1479
9.2.4.6	Image Gallery	1480
9.2.4.7	Marketplace Images for Rocky Linux	1480
9.2.4.8	KIB for AKS.....	1481
9.2.5	KIB with GCP	1481
9.2.5.1	Prerequisites	1481
9.2.5.2	GCP Prerequisite Roles	1482
9.2.6	KIB for GPU.....	1485
9.2.6.1	Verification	1486
9.2.7	KIB with vSphere	1487
9.2.7.1	Create a vSphere Base OS Image	1487
9.2.7.2	Create a vSphere Virtual Machine Template.....	1489
9.2.8	KIB with Pre-provisioned Environments	1493
9.2.9	Konvoy Image Builder CLI	1494
9.2.9.1	Other resources:.....	1494
9.2.9.2	konvoy-image build.....	1494

9.2.9.3	konvoy-image completion	1501
9.2.9.4	konvoy-image generate-docs	1506
9.2.9.5	konvoy-image generate	1506
9.2.9.6	konvoy-image provision.....	1510
9.2.9.7	konvoy-image upload.....	1511
9.2.9.8	konvoy-image validate	1512
9.2.9.9	konvoy-image version.....	1513
9.2.10	Use Override Files with Konvoy Image Builder.....	1513
9.2.10.1	Learn how to use override files with Konvoy Image builder	1513
9.2.10.2	Default Override Files	1514
9.2.10.3	Custom Override Files	1520
9.3	FIPS 140-2 Compliance	1526
9.3.1	FIPS Support in DKP	1526
9.3.2	Infrastructure Requirements for FIPS-140-2 Mode	1527
9.3.2.1	Supported Operating Systems	1527
9.3.3	Deploying a Cluster in FIPS mode.....	1527
9.3.3.1	Supported FIPS Builds	1527
9.3.4	Create FIPS 140 Images: Non-air-gapped Environments	1528
9.3.4.1	Use Konvoy Image Builder to create images with FIPS-compliant binaries	1528
9.3.4.2	Non-air-gapped Environment Create FIPS-140 images	1528
9.3.5	Create FIPS 140 Images: Air-gapped Environment.....	1529
9.3.5.1	Examples:	1530
9.3.5.2	Pre-provisioned FIPS Infrastructure	1530
9.3.6	Validate FIPS in Cluster	1531
9.3.6.1	Run FIPS validation.....	1531
9.3.6.2	Signature Files.....	1532
9.3.7	FIPS 140 Mode Performance Impact	1533
9.3.7.1	Understand the performance impact from operating your cluster in FIPS 140 mode.....	1533

9.4	Registry Mirror Tools	1533
9.4.1	Air-Gapped Registry Prerequisites	1533
9.4.2	AWS ECR	1534
9.4.2.1	Prerequisites	1534
9.4.2.2	Upload the Air-gapped Image Bundle to the Local ECR Registry:.....	1534
9.4.2.3	Air-gapped Environment Information regarding your AWS ECR Account	1534
9.4.3	JFrog Artifactory	1535
9.4.4	Nexus Registry	1535
9.4.5	Harbor Registry	1535
9.4.6	Bastion Host.....	1535
9.4.7	Registry Mirror.....	1535
9.4.8	Next Topic:	1536
9.4.9	Use a Registry Mirror	1536
9.4.9.1	Alternative mirror example	1537
9.5	Air-gapped Seed the Registry	1537
9.6	Configure the Control Plane	1538
9.6.1	Prerequisites	1538
9.6.1.1	Modifying Audit Logs.....	1538
9.6.1.2	Viewing the Audit Logs.....	1545
9.7	Update Cluster Nodepools	1545
9.7.1	Prerequisites:	1546
9.7.2	Steps:	1546
9.7.3	Related Topics:	1546
9.8	Verify your Cluster and DKP Installation	1547
9.8.1	Check the Cluster Infrastructure and Nodes.....	1547
9.8.2	Monitor the CAPI resources	1548
9.8.3	Verify all Pods	1548
9.8.4	Troubleshooting.....	1549

9.9	GPU for Konvoy	1549
9.10	Delete a DKP Cluster with One Command	1549
10	Upgrade DKP	1551
10.1	Prerequisite	1551
10.1.1	Supported Upgrade Paths:	1551
10.1.2	Upgrade Order	1552
10.1.3	Next Step	1552
10.2	Upgrade Compatibility Tables	1552
10.2.1	Supported Operating Systems	1552
10.2.2	Konvoy Image Builder	1552
10.2.3	Kubernetes Version Supported	1552
10.2.3.1	Deploying Clusters Versions	1553
10.2.3.2	Attaching Clusters Versions.....	1553
10.2.4	Next Step:	1554
10.2.5	Upgrade Cluster Node Pools	1554
10.2.5.1	Prerequisites:	1554
10.2.5.2	Steps:	1554
10.2.5.3	Related Topics:	1555
10.3	Upgrade Prerequisites	1555
10.3.1	Prerequisites for the Kommander Component.....	1555
10.3.1.1	All Kommander Environments	1555
10.3.2	Prerequisites for the Konvoy Component	1557
10.3.2.1	All Konvoy Environments.....	1557
10.3.3	Third-Party and Open Source Attributions	1558
10.3.3.1	Next Steps:	1558
10.4	Upgrade: For Air-gapped Environments Only	1558
10.4.1	Overview of Seeding the Registry for Air-gapped Environment	1558
10.4.2	Download all Images for Air-gapped Deployments	1559
10.4.3	Extract Air-gapped Images and Set Variables	1559

10.4.4	Only Pre-provisioned: Load Images for Deployments - Konvoy	1559
10.4.5	Load Images to your Private Registry - Konvoy	1561
10.4.6	Load Images to your Private Registry - Kommander	1561
10.4.7	Load Images to your Private Registry - DKP Catalog Applications.....	1561
10.4.8	Next Step:	1562
10.5	Upgrade DKP Enterprise	1562
10.5.1	Steps to upgrade Konvoy via CLI	1562
10.5.2	Overview of Steps in Order	1563
10.5.3	Next Step:	1563
10.5.4	Enterprise: For Air-gapped Environments Only	1563
10.5.4.1	Extract Air-gapped Images and Set Variables	1564
10.5.4.2	Load Images to your Private Registry - Kommander	1564
10.5.4.3	Load Images to your Private Registry - DKP Catalog Applications.....	1564
10.5.4.4	Load Images to your Private Registry - Konvoy	1564
10.5.4.5	Next Step:	1565
10.5.5	Enterprise: Upgrade the Management Cluster and Platform Applications	1565
10.5.5.1	Upgrade Kommander.....	1566
10.5.6	Enterprise: Upgrade Platform Applications on Managed and Attached Clusters.....	1569
10.5.6.1	Prerequisites	1570
10.5.6.2	Upgrade Platform Applications from the CLI	1570
10.5.6.3	Next Step:	1571
10.5.7	Enterprise: Upgrade Workspace Catalog Applications.....	1571
10.5.7.1	Upgrade Catalog Applications	1572
10.5.7.2	Next Step:	1574
10.5.8	Enterprise: Upgrade Project Catalog Applications	1574
10.5.8.1	Upgrade with the UI.....	1574
10.5.8.2	Upgrade with the CLI	1575
10.5.8.3	Next Step:	1576

10.5.9	Enterprise: Upgrade Custom Applications	1576
10.5.9.1	Verify the compatibility of Custom Applications with the current Kubernetes version	1576
10.5.9.2	Next Step:	1576
10.5.10	Enterprise: Upgrade the Management Cluster CAPI Components	1576
10.5.10.1	Upgrade the CAPI Components	1577
10.5.10.2	Next Step:	1578
10.5.11	Enterprise: Upgrade the Management Cluster Core Addons	1578
10.5.11.1	See Also:	1580
10.5.11.2	Next Step:	1580
10.5.12	Enterprise: Upgrade the Management Cluster Kubernetes Version	1580
10.5.12.1	Upgrade the Kubernetes Version	1580
10.5.12.2	Next Step:	1585
10.5.13	Enterprise: Upgrade Managed Clusters	1585
10.5.13.1	Upgrade Managed Clusters	1585
10.5.13.2	Upgrade Kubernetes Version on a Managed Cluster	1585
10.5.13.3	Upgrade Attached Cluster	1587
10.6	Upgrade DKP Essential	1588
10.6.1	Steps to upgrade DKP via CLI	1588
10.6.2	Overview of Steps in Order	1588
10.6.3	Next Step:	1588
10.6.4	Essential: For Air-gapped Environments Only	1589
10.6.4.1	Extract Air-gapped Images and Set Variables	1589
10.6.4.2	Load Images to your Private Registry - Kommander	1589
10.6.4.3	Load Images to your Private Registry - Konvoy	1589
10.6.4.4	Next Step:	1590
10.6.5	Essential: Upgrade the Essential Cluster and Platform Applications	1590
10.6.5.1	Upgrade Kommander	1591
10.6.6	Essential: Upgrade the Cluster CAPI Components	1592

10.6.6.1	Upgrade the CAPI Components	1593
10.6.6.2	Next Step:	1594
10.6.7	Essential: Upgrade the Cluster Core Addons	1594
10.6.7.1	Upgrade the Core Addons	1594
10.6.7.2	Next Step:	1596
10.6.8	Essential Upgrade Kubernetes Version	1596
10.6.8.1	Upgrade the Kubernetes Version	1596
11	CLI and API Tools	1601
11.1	API Documentation.....	1601
11.1.1	apps.kommander.mesosphere.io/v1alpha2.....	1601
11.1.1.1	App.....	1601
11.1.1.2	AppDeployment.....	1601
11.1.1.3	AppDeploymentList	1602
11.1.1.4	AppDeploymentSpec	1602
11.1.1.5	AppList.....	1603
11.1.1.6	ClusterApp.....	1603
11.1.1.7	ClusterAppList.....	1603
11.1.1.8	CrossNamespaceGitRepositoryReference.....	1604
11.1.1.9	GenericAppSpec	1604
11.1.1.10	TypedLocalObjectReference	1605
11.1.2	apps.kommander.mesosphere.io/v1alpha3.....	1605
11.1.2.1	App.....	1605
11.1.2.2	AppDeployment.....	1605
11.1.2.3	AppDeploymentClusterCondition	1606
11.1.2.4	AppDeploymentList	1606
11.1.2.5	AppDeploymentSpec	1607
11.1.2.6	AppDeploymentStatus.....	1607
11.1.2.7	AppList.....	1608
11.1.2.8	ClusterApp.....	1609

11.1.2.9	ClusterAppList.....	1609
11.1.2.10	ClusterConfigOverrides	1609
11.1.2.11	ClusterDeploymentStatus	1610
11.1.2.12	CrossNamespaceGitRepositoryReference.....	1610
11.1.2.13	GenericAppSpec	1611
11.1.2.14	TypedLocalObjectReference	1611
11.1.3	dispatch.d2iq.io/v1alpha2.....	1612
11.1.3.1	GitopsRepository	1612
11.1.3.2	GitopsRepositoryList	1612
11.1.3.3	GitopsRepositorySpec	1612
11.1.3.4	GitopsRolloutTemplate.....	1613
11.1.4	kommander.mesosphere.io/v1beta1	1614
11.1.4.1	CAPIClusterReference	1614
11.1.4.2	ClusterReference	1614
11.1.4.3	GenericClusterReference.....	1614
11.1.4.4	IngressSpec.....	1614
11.1.4.5	IngressStatus	1615
11.1.4.6	IssuerReference	1616
11.1.4.7	KommanderCluster.....	1617
11.1.4.8	KommanderClusterList.....	1617
11.1.4.9	KommanderClusterSpec	1617
11.1.4.10	KommanderClusterStatus.....	1619
11.1.4.11	License	1620
11.1.4.12	LicenseCondition	1621
11.1.4.13	LicenseExternalAWS.....	1621
11.1.4.14	LicenseExternalReference.....	1622
11.1.4.15	LicenseList	1622
11.1.4.16	LicenseSpec	1622
11.1.4.17	LicenseStatus.....	1623

11.1.4.18 PlacementSelector.....	1624
11.1.4.19 VirtualGroup	1625
11.1.4.20 VirtualGroupClusterRoleBinding	1625
11.1.4.21 VirtualGroupClusterRoleBindingList	1625
11.1.4.22 VirtualGroupClusterRoleBindingSpec	1626
11.1.4.23 VirtualGroupList	1626
11.1.4.24 VirtualGroupSpec	1626
11.1.5 workspaces.kommander.mesosphere.io/v1alpha1	1627
11.1.5.1 KommanderProjectRole	1627
11.1.5.2 KommanderProjectRoleList	1627
11.1.5.3 KommanderProjectRoleSpec	1627
11.1.5.4 KommanderProjectRoleStatus	1628
11.1.5.5 KommanderWorkspaceRole	1628
11.1.5.6 KommanderWorkspaceRoleList	1628
11.1.5.7 KommanderWorkspaceRoleSpec	1629
11.1.5.8 KommanderWorkspaceRoleStatus.....	1629
11.1.5.9 Project.....	1629
11.1.5.10 ProjectCondition	1630
11.1.5.11 ProjectList	1630
11.1.5.12 ProjectRole	1631
11.1.5.13 ProjectRoleList.....	1631
11.1.5.14 ProjectRoleSpec.....	1631
11.1.5.15 ProjectRoleStatus	1632
11.1.5.16 ProjectSpec	1632
11.1.5.17 ProjectStatus.....	1632
11.1.5.18 VirtualGroupKommanderClusterRoleBinding	1633
11.1.5.19 VirtualGroupKommanderClusterRoleBindingList	1633
11.1.5.20 VirtualGroupKommanderClusterRoleBindingSpec	1634
11.1.5.21 VirtualGroupKommanderClusterRoleBindingStatus.....	1634

11.1.5.22	VirtualGroupKommanderProjectRoleBinding	1634
11.1.5.23	VirtualGroupKommanderProjectRoleBindingList	1635
11.1.5.24	VirtualGroupKommanderProjectRoleBindingSpec	1635
11.1.5.25	VirtualGroupKommanderProjectRoleBindingStatus.....	1636
11.1.5.26	VirtualGroupKommanderWorkspaceRoleBinding	1636
11.1.5.27	VirtualGroupKommanderWorkspaceRoleBindingList	1636
11.1.5.28	VirtualGroupKommanderWorkspaceRoleBindingSpec	1637
11.1.5.29	VirtualGroupKommanderWorkspaceRoleBindingStatus.....	1637
11.1.5.30	VirtualGroupProjectRoleBinding	1637
11.1.5.31	VirtualGroupProjectRoleBindingList	1638
11.1.5.32	VirtualGroupProjectRoleBindingSpec	1638
11.1.5.33	VirtualGroupProjectRoleBindingStatus	1639
11.1.5.34	VirtualGroupWorkspaceRoleBinding	1639
11.1.5.35	VirtualGroupWorkspaceRoleBindingList	1640
11.1.5.36	VirtualGroupWorkspaceRoleBindingSpec	1640
11.1.5.37	VirtualGroupWorkspaceRoleBindingStatus	1641
11.1.5.38	Workspace.....	1641
11.1.5.39	WorkspaceCondition	1641
11.1.5.40	WorkspaceList	1642
11.1.5.41	WorkspaceRole	1642
11.1.5.42	WorkspaceRoleList.....	1643
11.1.5.43	WorkspaceRoleSpec.....	1643
11.1.5.44	WorkspaceRoleStatus	1643
11.1.5.45	WorkspaceSpec	1644
11.1.5.46	WorkspaceStatus.....	1644
11.2	CLI Commands	1644
11.2.1	CLI Commands for DKP.....	1644
11.2.2	dkp attach.....	1645
11.2.2.1	Options	1645

11.2.2.2	SEE ALSO.....	1645
11.2.2.3	dkp attach cluster	1645
11.2.3	dkp check	1646
11.2.3.1	Options	1646
11.2.3.2	SEE ALSO.....	1646
11.2.3.3	dkp check cluster	1647
11.2.4	dkp cluster.....	1648
11.2.4.1	Options	1648
11.2.4.2	SEE ALSO.....	1649
11.2.4.3	dkp cluster type.....	1649
11.2.5	dkp completion	1649
11.2.5.1	Synopsis	1650
11.2.5.2	Options	1650
11.2.5.3	Options inherited from parent commands	1650
11.2.5.4	SEE ALSO.....	1650
11.2.5.5	dkp completion fish	1650
11.2.5.6	dkp completion bash	1651
11.2.5.7	dkp completion zsh.....	1652
11.2.5.8	dkp completion powershell	1653
11.2.6	dkp config.....	1654
11.2.6.1	Options	1654
11.2.6.2	SEE ALSO.....	1654
11.2.6.3	dkp config get	1655
11.2.6.4	dkp config set.....	1656
11.2.7	dkp create.....	1657
11.2.7.1	Options	1657
11.2.7.2	Options inherited from parent commands	1658
11.2.7.3	SEE ALSO.....	1658
11.2.7.4	dkp create workspace	1658

11.2.7.5	dkp create appdeployment.....	1659
11.2.7.6	dkp create capi-components	1660
11.2.7.7	dkp create image-bundle	1661
11.2.7.8	dkp create bootstrap	1661
11.2.7.9	dkp create cluster	1662
11.2.7.10	dkp create chart-bundle.....	1683
11.2.7.11	dkp create nodepool	1684
11.2.8	dkp delete	1697
11.2.8.1	Options	1697
11.2.8.2	Options inherited from parent commands.....	1698
11.2.8.3	SEE ALSO.....	1698
11.2.8.4	dkp delete bootstrap.....	1698
11.2.8.5	dkp delete capi-components.....	1698
11.2.8.6	dkp delete cluster.....	1699
11.2.8.7	dkp delete chart	1700
11.2.8.8	dkp delete nodepool	1701
11.2.9	dkp describe	1701
11.2.9.1	Options	1701
11.2.9.2	SEE ALSO.....	1701
11.2.9.3	dkp describe cluster	1701
11.2.10	dkp detach.....	1702
11.2.10.1	Options	1702
11.2.10.2	SEE ALSO.....	1702
11.2.10.3	dkp detach cluster	1703
11.2.11	dkp diagnose.....	1703
11.2.11.1	Options	1704
11.2.11.2	SEE ALSO.....	1705
11.2.11.3	dkp diagnose default-config	1705
11.2.11.4	dkp diagnose ssh.....	1705

11.2.12	dkp edit	1706
11.2.12.1	Synopsis	1706
11.2.12.2	Options	1706
11.2.13	dkp get	1707
11.2.13.1	Options	1708
11.2.13.2	Options inherited from parent commands	1708
11.2.13.3	SEE ALSO	1708
11.2.13.4	dkp get kubeconfig	1708
11.2.13.5	dkp get appdeployments	1709
11.2.13.6	dkp get workspaces	1710
11.2.13.7	dkp get nodepools	1710
11.2.13.8	dkp get chart	1711
11.2.13.9	dkp get clusters	1712
11.2.14	dkp import	1712
11.2.14.1	Options	1712
11.2.14.2	SEE ALSO	1712
11.2.14.3	dkp import image-bundle	1713
11.2.15	dkp install	1713
11.2.15.1	Options	1713
11.2.15.2	SEE ALSO	1714
11.2.15.3	dkp install kommander	1714
11.2.16	dkp move	1715
11.2.16.1	Synopsis	1715
11.2.16.2	Options	1715
11.2.16.3	SEE ALSO	1715
11.2.16.4	dkp move capi-resources	1716
11.2.17	dkp open	1716
11.2.17.1	Options	1716
11.2.17.2	SEE ALSO	1717

11.2.17.3 dkp open dashboard	1717
11.2.18 dkp push	1717
11.2.18.1 Options	1718
11.2.18.2 Options inherited from parent commands	1718
11.2.18.3 SEE ALSO	1718
11.2.18.4 dkp push chart-bundle	1718
11.2.18.5 dkp push chart	1719
11.2.18.6 dkp push bundle	1719
11.2.19 dkp scale	1720
11.2.19.1 Options	1720
11.2.19.2 SEE ALSO	1720
11.2.19.3 dkp scale nodepool	1720
11.2.20 dkp serve	1721
11.2.20.1 Options	1721
11.2.20.2 SEE ALSO	1721
11.2.20.3 dkp serve bundle	1721
11.2.21 dkp update	1722
11.2.21.1 Options	1722
11.2.21.2 SEE ALSO	1722
11.2.21.3 dkp update nodepool	1722
11.2.21.4 dkp update controlplane	1728
11.2.21.5 dkp update bootstrap	1734
11.2.22 dkp upgrade	1738
11.2.22.1 Options	1738
11.2.22.2 Options inherited from parent commands	1738
11.2.22.3 SEE ALSO	1738
11.2.22.4 dkp upgrade kommander	1738
11.2.22.5 dkp upgrade capi-components	1740
11.2.22.6 dkp upgrade catalogapp	1740

11.2.22.7	dkp upgrade workspace	1741
11.2.22.8	dkp upgrade addons	1742
11.2.23	dkp version	1749
11.2.23.1	Synopsis	1749
11.2.23.2	Options	1749
11.2.23.3	Options inherited from parent commands	1749
11.2.24	dkp upload	1749
11.2.24.1	Options	1749
11.2.24.2	SEE ALSO	1749
11.2.24.3	dkp upload telemetry	1750
12	Release Notes	1752
12.1	DKP 2.6.0 Release Notes	1752
12.1.1	Release Summary	1752
12.1.2	Additional resources	1753
12.1.3	DKP 2.6.0 Features and Enhancements	1753
12.1.3.1	Podman Support	1753
12.1.3.2	Kubernetes 1.26.6 Support	1753
12.1.3.3	Telemetry	1753
12.1.3.4	DKP AI Navigator	1754
12.1.3.5	High Availability of DKP Workloads	1754
12.1.3.6	Auto-attach CAPI Clusters	1754
12.1.3.7	AWS ECR Supported	1754
12.1.3.8	Customizable Banners in the Dashboard	1755
12.1.3.9	Support for Provisioning in vCloud Director	1755
12.1.3.10	Observability and Monitoring of Network-restricted Clusters at Edge	1755
12.1.3.11	Provision VMware vSphere Clusters from the DKP User Interface	1755
12.1.3.12	Cluster Autoscaler can Manage Multiple Clusters	1756
12.1.3.13	Advanced Scenario-Based Kommander Installations	1756
12.1.3.14	Scenario Based DKP Upgrades	1756

12.1.3.15 Updated Links to Documentation in the DKP UI	1756
12.1.3.16 Documentation Page Rating and Feedback.....	1757
12.1.3.17 Konvoy Image Builder (KIB) Release v2.5.0	1757
12.1.3.18 Display Application Dependencies in the UI.....	1757
12.1.3.19 DKP Insights Release Notes.....	1758
12.1.4 DKP 2.6.0 Supported Kubernetes Versions.....	1758
12.1.4.1 Kubernetes Version Supported	1758
12.1.5 DKP 2.6.0 Kubernetes Major Updates and Deprecations.....	1759
12.1.5.1 Deprecated API Services	1760
12.1.5.2 Removed OCI Image Registry	1760
12.1.5.3 Migration to External Cloud-Controller-Managers	1760
12.1.6 DKP 2.6.0 Components and Applications	1760
12.1.6.1 Components	1760
12.1.6.2 Applications	1762
12.1.7 DKP 2.6.0 Known Issues and Limitations	1769
12.1.7.1 AWS additionalTags cannot contain spaces	1769
12.1.7.2 Increased Time to Attach Several Clusters.....	1769
12.1.7.3 Use Static Credentials to Provision an Azure Cluster.....	1769
12.1.7.4 Containerd File Limit Issue.....	1770
12.1.7.5 Intermittent Error Status when Creating EKS Clusters in the UI	1770
12.1.7.6 Limitations to Disk Resizing in vSphere	1770
12.1.7.7 Error Status in Grafana Logging Dashboard with EKS Clusters.....	1770
12.1.7.8 Rook Ceph Install Error	1771
12.1.7.9 Kommander Installation Configuration File Changes	1772
12.1.7.10 FluentD Logging Operator Error	1772
12.1.7.11 Generate a Token or Enable the Konvoy Credentials Plugin for Proxied Clusters.....	1774
12.1.7.12 Potential Logging Interruption during Upgrade.....	1779
12.1.7.13 Gatekeeper Uninstallation Error	1780

12.1.8	DKP 2.6.0 Customer Incidents	1780
12.1.8.1	kubernetes/crictl Cannot Pull Images from Harbor Registry.....	1780
12.1.8.2	Mis-configuration of OIDC Identity Provider from CLI or Kommander UI Crashes Dex and Locks User out of Kommander UI	1781
12.1.8.3	Kommander Failed to Install on Microsoft Azure.....	1781
12.1.8.4	Failing Kommander Upgrade from DKP 2.2.3 to 2.3.2	1781
12.1.8.5	Add Note about Deploying with a Lower Kubernetes Version	1781
12.1.8.6	DKP Missing Required Images for kubetunnel in Air-gapped Environments	1781
12.1.8.7	Missing registry_2.tar File in 2.5 Beta Air-gapped Bundle.....	1781
12.1.8.8	Scaling Nodes for Pre-provisioned Clusters	1781
12.1.8.9	Velero 'backup logs' Command Doesn't Work	1782
12.1.8.10	Unable to Build Ubuntu KIB Images in GCP	1782
12.1.8.11	Unable to Deploy Kommander with HTTP_PROXY	1782
12.1.8.12	Traefik entry point TCP 9090 is tagged Velero-minio.....	1782
12.1.8.13	Kommander Install on Small Environment Needs Reference to Platform Application Dependencies.....	1782
12.1.8.14	Microsoft Azure Upgrade from DKP 2.3.0 to DKP 2.4.0 Breaks Volume Attachments for Pods	1782
12.1.8.15	License Counts the Control Plane Cores after the 2.4 Upgrade	1782
12.1.8.16	EKS Cluster Creation Fails when AWS User Credentials do not have Full IAM Access	1783
12.1.8.17	Increase Disk size for vSphere Cluster does not Increase ContainerD Space	1783
12.1.8.18	Erroneous Configuration of OIDC Identity Provider Prevents User from Accessing the UI	1783
12.1.8.19	Calico Not Updated During DKP Upgrade from 2.1.x to 2.2.2 on Flatcar.....	1783
12.1.8.20	Attaching a DKP Essential Cluster via Kubetunnel not Supported	1783
12.1.8.21	Instructions on Configuring a Custom Domain and Certificate After Installing Kommander are Missing.....	1783


12.1.8.22	dkp create cluster Fails when the Kubeconfig file from a Previous Cluster is set	1784
12.1.8.23	Sizing Guidelines for Logging Stack and Rook Ceph Applications.....	1784
12.1.9	DKP 2.6.0 Deprecations.....	1784
12.1.9.1	Deprecated dkp push Subcommand and Flag	1784
12.1.9.2	Catalog Applications Compatibility	1784
12.1.9.3	Spark Operator to be Removed in DKP 2.7.....	1785
12.2	DKP 2.6.1 Release Notes.....	1785
12.2.1	Release Summary	1785
12.2.2	Additional resources.....	1786
12.2.3	DKP 2.6.1 Features and Enhancements	1786
12.2.3.1	AI Navigator Optional Install	1786
12.2.4	DKP 2.6.1 Supported Kubernetes Versions.....	1786
12.2.4.1	Kubernetes Version Supported	1786
12.2.5	DKP 2.6.1 Kubernetes Major Updates and Deprecations.....	1787
12.2.5.1	Deprecated API Services	1787
12.2.5.2	Removed OCI Image Registry	1788
12.2.5.3	Migration to External Cloud-Controller-Managers	1788
12.2.6	DKP 2.6.1 Components and Applications	1788
12.2.6.1	Components	1788
12.2.6.2	Applications	1789
12.2.7	DKP 2.6.1 Known Issues and Limitations	1796
12.2.7.1	Known Issues	1797
12.2.8	DKP 2.6.1 Customer Incidents and Resolved Issues	1808
12.2.8.1	Prometheus has duplicate targets for the apiserver	1808
12.2.8.2	Running the DKP upgrade command twice causes a regression in the associated containerd secret.....	1808
12.2.8.3	Pre-provisioned Upgrade failed with Flatcar	1808
12.2.8.4	Editing Credentials of a vSphere cluster in the DKP UI were not updated.....	1808

12.2.8.5	Limits on Capi-controller needed raised.....	1809
12.2.8.6	Insight alerts from an Essentials cluster converted into Managed land into the management cluster workspace	1809
12.2.8.7	Grafana-Loki telemetry sending data back to Grafana.....	1809
12.2.8.8	Increased Time to Attach Several Clusters in 2.6.0.....	1809
12.2.9	DKP 2.6.1 Deprecations.....	1809
12.2.9.1	Deprecated dkp push Subcommand and Flag	1809
12.2.9.2	Catalog Applications Compatibility	1810
12.2.9.3	Spark Operator to be Removed in DKP 2.7.....	1810
12.3	DKP 2.6.2 Release Notes.....	1810
12.3.1	Release Summary	1810
12.3.2	Additional resources.....	1811
12.3.3	DKP 2.6.2 Features and Enhancements.....	1811
12.3.3.1	Kubernetes 1.26.14 Support	1811
12.3.3.2	AWS Custom AMI Required.....	1811
12.3.3.3	Konvoy Image Builder (KIB) Release v2.5.5	1812
12.3.3.4	AI Navigator Optional Install	1812
12.3.4	DKP 2.6.2 Supported Kubernetes Versions.....	1812
12.3.4.1	Deploying Clusters Versions	1812
12.3.4.2	Attaching Clusters Versions.....	1813
12.3.5	DKP 2.6.2 Kubernetes Major Updates and Deprecations.....	1813
12.3.5.1	Deprecated API Services	1814
12.3.5.2	Removed OCI Image Registry	1814
12.3.5.3	Migration to External Cloud-Controller-Managers	1814
12.3.6	DKP 2.6.2 Components and Applications	1814
12.3.6.1	Components	1814
12.3.6.2	Applications	1816
12.3.7	DKP 2.6.2 Known Issues and Limitations	1823
12.3.7.1	Known Issues	1823

12.3.8	DKP 2.6.2 Customer Incidents and Resolved Issues	1835
12.3.8.1	CAPI Move Causes Worker Kubelet Issue (CAPPP).....	1835
12.3.8.2	Upgrade Could Fail for EKS or AKS Clusters.....	1835
12.3.8.3	CAPA AMIs are Not Available for the Kubernetes Version used in DKP 2.6.2	1835
12.3.9	DKP 2.6.2 Deprecations.....	1835
12.3.9.1	Deprecated dkp push Subcommand and Flag.....	1835
12.3.9.2	Catalog Applications Compatibility	1836
12.3.9.3	Spark Operator to be Removed in DKP 2.7.....	1836
13	AI Navigator.....	1839
13.1	AI Navigator User Agreement and Guidelines.....	1839
13.1.1	Accept the User Agreement	1839
13.1.2	Guidelines for Using AI Navigator.....	1840
13.1.2.1	Next Topic	1841
13.2	AI Navigator in Installations and Upgrades.....	1841
13.2.1	Install DKP 2.6.1 and Later Releases.....	1841
13.2.2	Upgrade DKP 2.6.0 to 2.6.1	1841
13.2.2.1	For Upgrade in DKP Enterprise Environments	1841
13.2.2.2	For Upgrade in DKP Essential Environments	1842
13.2.2.3	Next Topic	1842
13.3	Access the AI Navigator	1842
13.3.1	Open the AI Navigator.....	1842
13.3.2	Close the AI Navigator	1843
13.3.2.1	Next Topic	1843
13.4	How to Create Good Queries.....	1843
13.4.1	What Goes in a Prompt.....	1843
13.4.1.1	Inline Commands or Code Snippets	1844
13.4.1.2	Code Blocks	1845
13.4.2	Selected Prompt Examples	1845

13.4.2.1	Evaluate a YAML Manifest	1845
13.4.2.2	Updating a TLS Thumbprint	1847
13.4.2.3	Describe Clock Skew in Ceph Cluster Status	1847
14	Access Documentation	1849
14.1	Supported Documentation	1849
14.2	Archived Documentation	1849
14.3	Download Documentation PDF.....	1849
14.4	Download Documentation.....	1849
15	D2iQ Security Updates.....	1851
15.1	CVE Policy	1851
15.1.1	CVE Management:	1852
15.1.1.1	Scanning Policy:.....	1852
15.1.1.2	Shipping Policy:.....	1852
15.1.1.3	Known Issues:	1852
16	Version Support Policy	1853
16.1	Supported Kubernetes Versions	1853
16.2	Features in Patches	1853
16.3	Experimental Status.....	1853
16.4	Beta Feature	1854
16.5	Technical Preview	1854
16.6	End-of-Life (EoL)	1854
16.7	Support Definitions	1854
16.7.1	Secondary Support	1854
16.8	Standard Level & Severity Definitions.....	1856
16.9	Operating System (OS) Version Support Policy	1856
16.9.1	Red Hat Enterprise Linux (RHEL) Version Support Policy	1856
16.9.1.1	Example:	1857
17	Legal Notices	1858

As the leading independent Kubernetes Management Platform in production, the D2iQ Kubernetes Platform (DKP) provides a holistic approach and a complete set of enterprise-grade technologies, services, training, and support to build and run applications in production at scale. Built around the open-source Cluster API, the new version of DKP becomes the single, centralized point of control for an organization’s application infrastructure, empowering organizations to more easily deploy, manage, and scale Kubernetes workloads in Day 2 production environments.



Features	Benefits
Container Orchestration	Leverage an industry standard distribution of open-source Kubernetes for cluster and container management.
Application Management and Deployment	Deploy applications and services within Kubernetes clusters with Helm ¹ .
Observability	Gain deep insight into your Kubernetes clusters and applications with open source metrics leveraging Telegraf, Prometheus, and Grafana ² .

1 <https://helm.sh/>

2 <https://grafana.com/docs/>

Features	Benefits
Cluster API	Automate cluster lifecycle management using Cluster API (see page 84) to simplify the provisioning, upgrading, and operating multiple Kubernetes clusters across a wide range of distributions and virtual and physical environments.
Cluster Autoscaling	Save operational costs by automatically scaling down (see page 1119) capacity when it's not needed and adding capability when there is greater demand, with CAPI enabled autoscaling groups.
Logging	Collect and analyze logs and metrics (see page 899) to ensure optimal performance and troubleshooting with Grafana (see page 985) , Loki , and Fluentbit (see page 923) .
Cloud-Native Scale Testing	Extensive integration and workload testing at massive scale with a wide range of workloads to ensure real-world preparedness.
Networking and Routing	Easily automate and expose application endpoints (see page 954) with Calico , Traefik (see page 946) , and CoreDNS .
Fine-Grained Cluster	Upgrades (see page 1551) reduce operational overhead with non-disruptive patching or parallel worker node upgrades.
Backup and Recovery	Ensure business continuity and disaster recovery (see page 874) with Velero³ .
Declarative Automated Installer With Day 2 Platform Services	Accelerate time-to-production on any infrastructure (see page 147) with consistency and reliability with the required platform applications (see page 639) needed for Day 2 production (see page 582) .
Operate in Air-gapped Environments (see page 310)	Leverage declarative APIs to optimize cluster resources for cost, resilience, and performance.

³ <https://velero.io/docs/main/>

Features	Benefits
End-to-End Support	Enterprise-grade support ⁴ and services for both Kubernetes and its supported platform applications.

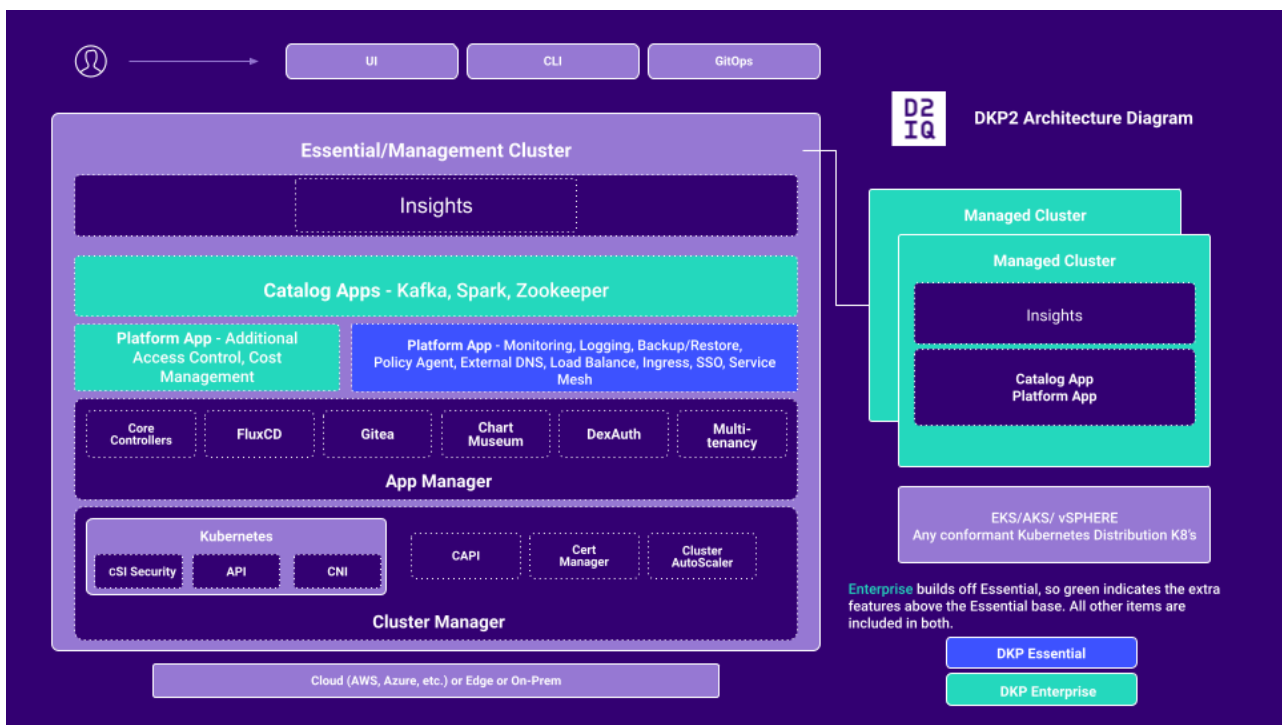
⁴ <https://support.d2iq.com/hc/en-us>

1 Architecture

1.1 Learn the key concepts and architectural components of a DKP cluster

Kubernetes provides the foundation for a DKP cluster. Because of this fundamental relationship, you should be familiar with a few key concepts and terms. The topic in this section provide a brief overview of the native Kubernetes architecture, a simplified view of the DKP architecture - both Essential and Enterprise versions. Also, it shows the operational workflow for a DKP cluster.

The following diagram provides a simplified architectural overview to help you visualize the flow of the key components:



1 DKP Architectural overview

1.2 Components for the Kubernetes Control Plane

The native Kubernetes cluster consists of **components** that provide the cluster’s **control plane** and **worker nodes** that run users’ containers and maintain the runtime environment.

DKP supplements the native Kubernetes cluster by providing a predefined and pre-configured set of applications. Because this predefined set of applications provides critical features for managing a Kubernetes cluster in a production environment, the default set is identified as DKP **platform services**.

See [Platform applications](#) (see page 582) for the full set of DKP platform services.

1.3 Related Information

For information on related topics or procedures, refer to the following [Kubernetes](#)⁵ documentation.

1.4 Next Steps

- [DKP Ports](#) (see page 67)
- [Supported Infrastructure Operating Systems](#) (see page 69)

1.5 DKP Ports

This section describes ports used by the different Kubernetes components in your DKP cluster. For more information regarding ports, refer to the Kubernetes Documentation on this page: [Ports and Protocols](#)⁶

1.5.1 Control-plane nodes

Port	DKP Component	Notes
22		ssh
179	calico-node	BGP
1338	Containerd	metrics
2379	etcd	client
2380	etcd	peer
6443	kube-apiserver	
7472	metallb speaker	monitoring
7946	metallb speaker	memberlist

⁵ <https://kubernetes.io/docs/concepts/overview/components/>

⁶ <https://kubernetes.io/docs/reference/networking/ports-and-protocols/>

Port	DKP Component	Notes
9091	calico-node	felix metrics
9092	calico-node	bird metrics
9099	calico-node	felix liveness
9100	prometheus node-exporter	metrics
10248	kubelet	health
10249	kube-proxy	metrics
10250	kubelet	
10256	kube-proxy	health
10257	kube-controller-manager	secure port
10259	kube-scheduler	secure port
30000-32767	Kubernetes NodePorts	

1.5.2 Worker nodes

Port	DKP Component	Notes
22	Ansible	ssh
179	calico-node	BGP
1338	Containerd	metrics
5473	calico-typha	syncserver

Port	DKP Component	Notes
9091	calico-node	felix metrics
9092	calico-node	bird metrics
9093	calico-typha	metrics
9099	calico-node	felix liveness
9100	prometheus node-exporter	metrics
9400	NVIDIA GPU DCGM	metrics
10248	kubelet	health
10249	kube-proxy	metrics
10250	kubelet	
10256	kube-proxy	health
30000-32767	Kubernetes NodePorts	

1.5.3 Next Topic:

[Supported Infrastructure Operating Systems \(see page 69\)](#)

1.6 Supported Infrastructure Operating Systems



If the full table is not showing below, select the box icon in the upper-right corner to expand the view:

Amazon Web Services (AWS)



1.6.1 Amazon Web Services (AWS)

Operating System	Kernel	Default Config (see page 294)	FIPS (see page 328)	Air Gapped (see page 310)	FIPS with Air Gapped (see page 344)	GPU Support (see page 363)	GPU with Air Gapped (see page 382)	Konvoy Image Builder (see page 1465)
CentOS 7.9 ⁷	3.10.0-1160.80.1.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 7.9 ⁸	3.10.0-1160.80.1.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.4 ⁹	4.18.0-305.88.1.el8_4.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.6 ¹⁰	4.18.0-372.52.1.el8_6.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ubuntu 18.04 (Bionic Beaver) ¹¹	5.4.0-1103-aws	Yes						Yes
Ubuntu 20.04 (Focal Fossa) ¹²	5.15.0-1040-aws	Yes				Yes		Yes

⁷ [https://wiki.centos.org/Manuals\(2f\)ReleaseNotes\(2f\)CentOS7\(2e\)2009.html](https://wiki.centos.org/Manuals(2f)ReleaseNotes(2f)CentOS7(2e)2009.html)

⁸ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/7.9_release_notes/index

⁹ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/8.4_release_notes/index

¹⁰ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/8.6_release_notes/index

¹¹ <https://wiki.ubuntu.com/BionicBeaver/ReleaseNotes>

¹² <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

Operating System	Kernel	Default Config (see page 294)	FIPS (see page 328)	Air Gapped (see page 310)	FIPS with Air Gapped (see page 344)	GPU Support (see page 363)	GPU with Air Gapped (see page 382)	Konvoy Image Builder (see page 1465)
SLES 15 SP3 ¹³	5.3.18-59.37-default	Yes				Yes		Yes
Oracle Linux RHCK 7.9 ¹⁴	3.10.0-1160.92.1.0.el7.x86_64	Yes	Yes					Yes
Rocky Linux 9.1 ¹⁵	5.14.0-162.12.1.el9_1.0.2.x86_64	Yes		Yes				Yes

1.6.2 Microsoft Azure

Operating System	Kernel	Default Config (see page 532)	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1477)
Ubuntu 20.04 (Focal Fossa) ¹⁶	5.15.0-1042-azure	Yes						Yes
Rocky Linux 9.1 ¹⁷	5.14.0-162.12.1.el9_1.0.2.x86_64	Yes						Yes

¹³ <https://documentation.suse.com/en-us/sles/15-SP3/>

¹⁴ <https://docs.oracle.com/en/operating-systems/oracle-linux/7/relnotes7.9/>

¹⁵ https://docs.rockylinux.org/release_notes/9_1/

¹⁶ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

¹⁷ https://docs.rockylinux.org/release_notes/9_0/

1.6.3 Google Cloud Platform (GCP)

Operating System	Kernel	Default Config (see page 563)	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1481)
CentOS 7.9 ¹⁸	3.10.0-1160.62.1.el7.x86_64	Yes						Yes
Ubuntu 18.04 ¹⁹	5.4.0-1072-gcp	Yes						Yes
Ubuntu 20.04 ²⁰	5.13.0-1024-gcp	Yes						Yes

1.6.4 Pre-Provisioned

[Pre-provisioned](#) (see page 149) includes on-premises, vSphere, AWS, Azure, and GCP infrastructures and is described in more detail under [What is Pre-provisioned Infrastructure](#) (see page 86).

Operating System	Kernel	Default Config (see page 151)	FIPS (see page 195)	Air Gapped (see page 170)	FIPS with Air-Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1493)
CentOS 7.9 ²¹	3.10.0-1160.80.1.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes

¹⁸ [https://wiki.centos.org/Manuals\(2f\)ReleaseNotes\(2f\)CentOS7\(2e\)2009.html](https://wiki.centos.org/Manuals(2f)ReleaseNotes(2f)CentOS7(2e)2009.html)

¹⁹ <https://wiki.ubuntu.com/BionicBeaver/ReleaseNotes>

²⁰ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

²¹ [https://wiki.centos.org/Manuals\(2f\)ReleaseNotes\(2f\)CentOS7\(2e\)2009.html](https://wiki.centos.org/Manuals(2f)ReleaseNotes(2f)CentOS7(2e)2009.html)

Operating System	Kernel	Default Config (see page 150)	FIPS (see page 195)	Air Gapped (see page 170)	FIPS with Air-Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1493)
RHEL 7.9 ²²	3.10.0-1160.80.1.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.4 ²³	4.18.0-305.88.1.el8_4.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.6 ²⁴	4.18.0-372.52.1.el8_6.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Flatcar 3033.3 ²⁵ x	5.10.186-flatcar	Yes						Yes
Ubuntu 20.04 ²⁶	5.15.0-1033-aws	Yes				Yes		Yes
SLES 15 ²⁷	5.3.18-59.37-default	Yes				Yes		Yes
Oracle Linux RHCK 7.9 ²⁸	5.4.17-2136.314.6.3.el7uek.x86_64	Yes						Yes

²² https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/7.9_release_notes/index

²³ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/8.4_release_notes/index

²⁴ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/8.6_release_notes/index

²⁵ <https://www.flatcar.org/releases/#lts-release>

²⁶ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

²⁷ <https://documentation.suse.com/en-us/sles/15-SP3/>

²⁸ <https://docs.oracle.com/en/operating-systems/oracle-linux/7/relnotes7.9/>

Operating System	Kernel	Default Config (see page 150)	FIPS (see page 195)	Air Gapped (see page 170)	FIPS with Air-Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1493)
Rocky Linux 9.1 ²⁹	5.14.0-162.12.1.el9_1.0.2.x86_64	Yes		Yes				Yes

1.6.5 Pre-provisioned/Azure

Operating System	Kernel	Default Config (see page 151)	FIPS (see page 195)	Air Gapped (see page 170)	FIPS with Air Gapped (see page 215)	GPU Support (see page 239)	GPU with Air Gapped (see page 262)	Konvoy Image Builder (see page 1493)
RHEL 8.6 ³⁰	4.18.0-372.52.1.el8_6.x86_64	Yes	Yes	Yes	Yes			Yes

1.6.6 vSphere

vCenter Server + ESXi

²⁹ https://docs.rockylinux.org/release_notes/9_1/

³⁰ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/8.6_release_notes/index

Operating System	Kernel	Default Config (see page 438)	FIPS (see page 483)	Air Gapped (see page 459)	FIPS with Air Gapped (see page 505)	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1487)
RHEL 7.9 ³¹	3.10.0-1160.el7.x86_64	Yes	Yes	Yes	Yes			Yes
RHEL 8.4 ³²	4.18.0-305.el8.x86_64	Yes	Yes	Yes	Yes			Yes
RHEL 8.6 ³³	4.18.0-372.9.1.el8.x86_64	Yes	Yes	Yes	Yes			Yes
Ubuntu 20.04 ³⁴	5.4.0-125-generic	Yes		Yes				Yes
Rocky Linux 9.1 ³⁵	5.14.0-162.6.1.el9_1.x86_64	Yes		Yes				Yes

³¹ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/7.9_release_notes/index

³² https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/8.4_release_notes/index

³³ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/8.6_release_notes/index

³⁴ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

³⁵ https://docs.rockylinux.org/release_notes/9_1/

1.6.7 VMware Cloud Director (VCD) - Beta Product

Operating System	Kernel	Default Config (see page 1322)	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1487)
Ubuntu 20.04 ³⁶	5.4.0-125-generic	Yes - Beta						

1.6.8 Amazon Elastic Kubernetes (EKS)

Operating System	Kernel	Default Config (see page 405)	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1476)
Amazon Linux 2 v7.9 ³⁷	5.10.178-162.673.amzn2.x86_64	Yes						

1.6.9 Azure Kubernetes Service (AKS)

Operating System	Kernel	Default Config (see page 547)	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1481)
Ubuntu 22.04.2 LTS ³⁸	5.15.0-1037-azure	Yes						

³⁶ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

³⁷ <https://docs.aws.amazon.com/AL2/latest/relnotes/relnotes-al2.html>

³⁸ <https://discourse.ubuntu.com/t/jammy-jellyfish-release-notes/24668>

1.6.10 Next Step

[Download DKP](#) (see page 78)

2 Download DKP

To download a new version of DKP, you have 2 options:

2.1 Download from the Support Website

Download DKP

From the Download site, select the DKP binary for either Darwin(MacOS) or Linux. After download, extract for your system.

To extract a tar file, either:

- Darwin - Use a file manager program like 7-Zip and right-click the tar file. Select “Open With” and then 7-Zip File Manager. Select “Extract” and choose a location to save the extracted files.
- Linux - Use the command line and type `tar -xvf filename.tar` to extract the files to the current directory. You can add `tar -xvzf` option if the file is compressed with gzip.

2.1.1 AWS Marketplace Only Download

Follow the instructions on the AWS console to download the container image.

After downloading the image, run this command to copy the binaries to your host:

```
docker run -it --rm -u $(id -u):$(id -g) -v $(pwd):/dkp $CONTAINER_IMAGES
```

A successful operation returns this output:

```
dkp binary is placed in the local directory, to run:
./dkp --help
```

And you can view the `dkp` binary in your working directory. Follow the [DKP installation \(see page 108\)](#) instructions using these binaries, and then [add your license \(see page 139\)](#) to DKP. If you have problems downloading or installing DKP, contact your sales representative or sales@d2iq.com³⁹.

2.1.2 Next Step:

[Day 0 - Get Started with DKP \(see page 79\)](#)

³⁹ <mailto:sales@d2iq.com>

3 Day 0 - Get Started with DKP

At D2iQ, we partner with you throughout the entire cloud-native journey, at the beginning on Day 0 all the way through your Day 2 operations. Here is the breakdown:

- **Day 0** is the planning phase where we introduce definitions and concepts.
- **Day 1** is to install the software and get it up and running.
- **Day 2** involves customizations, application, and operations management. In the end, we provide ongoing maintenance, keep the platform resilient, keep it secure, and make sure it can be upgraded.

When installing DKP, first determine the infrastructure on which you want to deploy. Then, ensure you understand the basic concepts and have the prerequisites met, like storage and resource requirements, using the information contained in this section of the documentation.

For example, you can:

- Install on a public cloud infrastructure, such as Amazon Web Services (AWS), GCP, or Azure.
- Install on an internal network, on-premises environment, with a physical or virtual infrastructure.
- Install on an Air-gapped environment.
- Install with or without FIPS and GPU.

After you choose your environment, you [download DKP \(see page 78\)](#), and then select the [Day 1 - Basic Install instructions \(see page 147\)](#) for your Infrastructure provider and environment. The basic installs set up a cluster with the Konvoy component, and then install the Kommander component, so that you can access the dashboards through the DKP UI. You can use these basic clusters to explore DKP and prepare to take your clusters into *production*, where you will deploy and enable the applications that support Day 2 operations. Then you can deploy and test your workloads, and put Kubernetes to work!



After you finish the basic install and are ready to customize, move to [Custom Installation and Additional Infrastructure Tools \(see page 1036\)](#), if desired. Otherwise, proceed to [Day 2 Cluster Operations Management \(see page 582\)](#) to make the software exactly what you want it to be.

Follow this list to ensure you have the resource requirements met for a successful implementation:

- [DKP Concepts and Terms \(see page 80\)](#)
- [Provide Context for Commands with a kubeconfig File \(see page 87\)](#)
- [Storage \(see page 89\)](#)
- [Resource Requirements \(see page 100\)](#)
- [Install Overview \(see page 108\)](#)
- [Prerequisites for Install \(see page 121\)](#)

3.1 DKP Concepts and Terms

Understand the terminology used in association with DKP with the information on this page - DKP concepts and terms.

DKP is composed of three main components, Konvoy, Kommander, and Konvoy Image Builder (KIB) that work together to provide a single and centralized point of control for an organization's application infrastructure. DKP empowers organizations to deploy, manage, and scale Kubernetes workloads in Day 2 production environments more easily.

Each main component specifically manages the following:

- **Konvoy** is the cluster lifecycle manager component of DKP. Konvoy relies on Cluster API, Calico, and other open-source and proprietary software to provide simple cluster lifecycle management for conformant Kubernetes clusters with networking and storage capabilities. Konvoy uses industry-standard tools to provision certified Kubernetes clusters on multiple cloud providers, vSphere, and on-premises hardware in connected and air-gapped environments. Konvoy contains the following components:
 - **Cluster Manager** - Cluster API, CSI, CNI, Cluster AutoScaler, Cert Manager, MetalLB

The Konvoy component is installed according to the **cluster's infrastructure**. See [Day 1 - Basic Installs by Infrastructure \(see page 147\)](#) or [Custom Installation and Additional Infrastructure Tools \(see page 1036\)](#) for more information.

- **Kommander** is the fleet management component of DKP. Kommander delivers centralized observability, control, governance, unified policy, and better operational insights. With DKP Essential, Kommander manages a single Kubernetes cluster. In DKP Enterprise, Kommander supports attaching workload clusters and lifecycle management of clusters using Cluster API. DKP Enterprise also offers lifecycle management of applications through FluxCD. Kommander contains the following components:
 - User interface, Security, Observability, Networking, Application Management
 - *Platform Applications*: Applications that come out of the box with DKP providing functionality like observability, cost management, monitoring, logging, making DKP clusters Day 2 production ready right from install. Platform applications are D2iQ's choice of selected applications from the open source community that are consumed by the platform.
 - **Essential Platform Applications**: Monitoring, Logging, Backup/Restore, Policy Agent, External DNS, Load Balance, Ingress, SSO, Service Mesh.
 - **Enterprise Platform Applications**: All above Essential Platform Applications, plus additional Access Control and Centralized Cost Management.

- *Catalog Applications*: Applications in DKP Enterprise that are deployed to be used for customer workloads, such as Kafka, Spark and ZooKeeper.

The Kommander component is installed according to the **cluster's environment type**. See [Installing Kommander by Environment](#) (see page 1394) for more information.

- **Konvoy Image Builder** (KIB) creates Cluster API-compliant machine images. It configures those images to contain all the necessary software to deploy Kubernetes cluster nodes. See [Konvoy Image Builder](#) (see page 1461) for more information.



DKP Essential and Enterprise also provide a helpful add-on called [DKP Insights](#)⁴⁰.

3.1.1 Next Topic:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20What%20is%20Pre-provisioned%20Infrastructure&linkCreation=true&fromPagelD=273319940>⁴¹


3.1.2 Cluster Types and Concepts

Cluster types such as Management clusters, Managed clusters, and Attached clusters are key concepts in understanding and getting the most out of DKP Essential versus Enterprise environments.

⁴⁰ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274890970/2.6+DKP+Insights+Overview>

⁴¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelD=273319940&linkCreation=true&spaceKey=DENT&title=%282.6%29+What+is+Pre-provisioned+Infrastructure>

3.1.2.1 Multi-cluster Environment



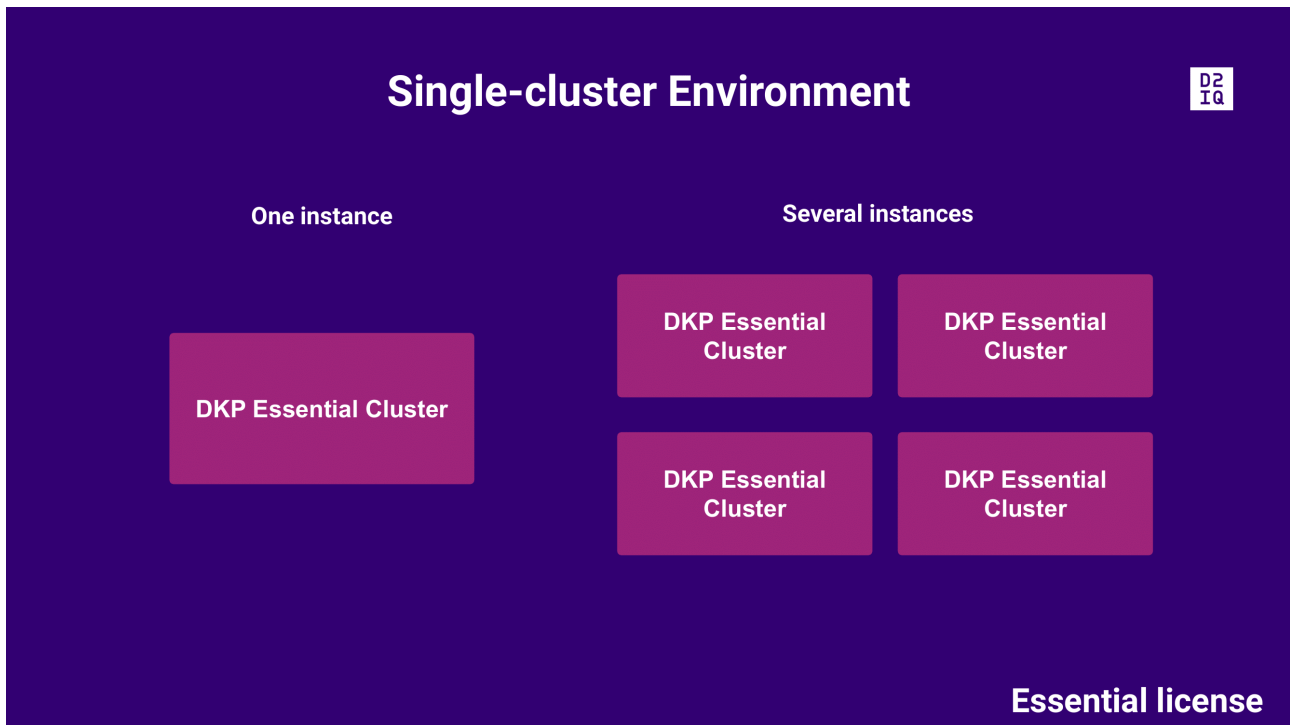
Management Cluster: Is the cluster where you install DKP, and it is self-managed. In a multi-cluster environment, the Management cluster also manages other clusters.

Customers with an [Enterprise license](#) (see page 130) should run workloads on *Managed* and *Attached* clusters, and not on the *Management cluster*.

Managed Cluster: Also called a “DKP cluster,” this is a type of workload cluster that you can create with DKP. The *DKP Management cluster* manages its infrastructure, its lifecycle, and its applications.

Attached Cluster: A type of workload cluster that is created outside of DKP, but is then connected to the *DKP Management Cluster* so that it can be managed by DKP. In these cases, the DKP Management cluster only manages the attached cluster’s applications.

3.1.2.2 Single-cluster Environment



DKP Essential Cluster: Is the cluster where you install DKP. A DKP Essential cluster is a stand-alone cluster. It is self-managed and therefore capable of provisioning itself. In this single-cluster environment, you cannot attach other clusters and all workloads are run on your *DKP Essential cluster*. You can, however, have several separate DKP Essential instances, each with its own license.

Customers with an [Essential license](#) (see page 132) can run workloads on their *DKP Essential cluster*.



If you have not decided which license to get, but plan on adding one or several clusters to your environment, and manage them centrally, D2iQ recommends obtaining an [Enterprise license](#) (see page 130).

3.1.2.3 Other important concepts

Self-managed Cluster: In a DKP landscape, only *DKP Essential* and *DKP Enterprise Management* clusters are self-managed. *Self-managed clusters* are clusters that manage the provisioning, and deployment of its own

nodes through CAPI controllers. The CAPI controllers are a managing entity, which automatically manages the lifecycle of a cluster's nodes based on a customizable definition of the resources.

A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing. As part of the underlying processing using the `--self-managed` flag, the DKP CLI:

- creates a bootstrap cluster
- creates a workload cluster
- moves CAPI controllers from the bootstrap cluster to the workload cluster, making it self-managed
- deletes the bootstrap cluster

Network-Restricted Cluster:

(NOT equivalent to air-gapped)

A network-restricted or firewalled cluster is secured by a firewall, DMZ, NAT gateway, proxy server, or requires additional access information. Network-restricted clusters are usually located in remote locations or at edge, and therefore not in the same network as the Management cluster.

The main difference between *network-restricted* and *air-gapped* clusters is that network-restricted clusters can reach external networks (like the Internet), but its services or ingresses cannot be accessed from outside. Air-gapped clusters, on the other hand, do not allow ingress, nor egress traffic.

In a multi-cluster environment, DKP supports [attaching a network-restricted cluster \(see page 772\)](#) to a DKP Management cluster. You can also [enable a proxied access \(see page 814\)](#) pipeline through the Management cluster, which allows you to access the network-restricted cluster's dashboards without being in the same network.

3.1.2.4 Next Topic:

[Provide Context for Commands with a kubeconfig File \(see page 87\)](#)

3.1.3 CAPI Concepts and Terms

DKP uses CAPI (ClusterAPI) technology for creating and managing the lifecycle of Kubernetes Clusters. A basic understanding of CAPI concepts and terms is helpful in understanding how to install and maintain DKP. You can find a deeper discussion of the architecture in the [ClusterAPI Book](#)⁴².

CAPI makes use of a bootstrap cluster for provisioning and starting clusters. A bootstrap cluster handles the following actions:

1. Generating the cluster certificates, if they are not otherwise specified.
2. Initializing the control plane, and managing the creation of other nodes until it is complete.
3. Joining control plane and worker nodes to the cluster.
4. Installing and configuring networking plugin (Calico CNI), CSI volume provisioners, cluster-autoscaler and other core Kubernetes components.

⁴² <https://cluster-api.sigs.k8s.io/user/concepts.html>

BootstrapData is machine or node role-specific data, such as cloud initialization data, used to bootstrap a “machine” onto a node.

For customers using Kommander for multi-cluster management, a **management cluster** that manages the lifecycle of workload clusters. As the management cluster, DKP Kommander works with bootstrap providers, infrastructure providers, and maintains cluster resources such as bootstrap configurations and templates. If you are working with only one cluster, Kommander will provide you with add-on (platform application) management for that cluster, but not others.

A **workload cluster** is a Kubernetes cluster whose lifecycle is managed by a management cluster, and provides the platform on which you deploy and execute your workloads.

As part of a collection of **Custom Resource Definitions** or **CRDs** that extend the Kubernetes API, these additional concepts are important for understanding the upgrade.

A **ClusterResourceSet** Kubernetes cluster created by CAPI is functionally minimal in nature. Crucial components like CSI and CNI are not a part of the default cluster spec. A ClusterResourceSet is a CRD that can be used to group and deploy core cluster components post-Kubernetes cluster install.

When you create a bootstrap cluster. You can find all the components in the default namespace and we move them to the workload cluster during the process of making the cluster self-managed.

A **machine** is a declarative spec for a platform or infrastructure component that hosts a Kubernetes node such as a bare metal server or a VM. CAPI uses provider-specific controllers to provision and install new hosts which then register as nodes. When you update a machine spec other than for certain values, such as annotations, status, and labels, the controller deletes the host and creates a new one that conforms to the new spec. This is called machine immutability. If you delete a machine, the controller deletes both the infrastructure and the node. Provider-specific information is not portable between providers.

Within CAPI, you use declarative **MachineDeployments** to handle changes to machines by replacing them in much the same way that a core Kubernetes Deployment replaces Pods. MachineDeployments reconcile changes to machine specs by rolling out changes to two **MachineSets** (similar to a ReplicaSet), both the old and the newly-updated.

A **MachineHealthCheck** identifies unhealthy node conditions, and initiates remediation for nodes owned by a MachineSet.

3.1.3.1 Next Topic:

[What is Pre-provisioned Infrastructure \(see page 86\)](#)

3.1.4 Air-gapped or Non-air-gapped Environment?

3.1.4.1 Air-gapped Environments

In an air-gapped environment, your environment is **isolated** from unsecured networks, like the Internet. Running your workloads on an air-gapped environment is common in scenarios where **security** is a determining factor. DKP in air-gapped environments allows you to manage your clusters while shielding it from undesired external influences.

You can create an air-gapped cluster on on-premise environments or any [supported cloud infrastructure \(see page 69\)](#). In this configuration, you are responsible for providing an image registry. You must also retrieve

required artifacts and configure DKP to use those from a local directory when creating and managing DKP clusters.

Secure interactions with other networks

There are a number of ways you can perform actions that require incoming data from other networks in spite of your environment's isolation: in some configurations, air-gapped clusters allow inbound connections, but cannot initiate outbound connections. In other configurations, you can set up a bastion host, which serves as a gateway between the Internet (or other untrusted networks) and your environment, and facilitates the download of install, upgrade, etc. files and images that are required for other machines to run in air-gapped environments.

Common industry synonyms:

Fettered, disconnected, restricted, SIPR, etc.

3.1.4.2 Non-air-gapped Environments

In a non-air-gapped environment, your environment has two-way access to and from the **Internet**. You can create a non-air-gapped cluster on pre-provisioned (on-premise) environments or any cloud infrastructure.

DKP in a non-air-gapped environment allows you to manage your clusters while facilitating **connections** and offering integration with other tools and systems.

Common industry synonyms:

Open, accessible (to the Internet), not restricted, NIPR, etc.

3.1.4.3 Next Step:

[What is Pre-provisioned Infrastructure \(see page 86\)](#)

3.1.5 What is Pre-provisioned Infrastructure

Pre-provisioned as an infrastructure type allows the deployment of Kubernetes using DKP® to pre-existing machines. Other providers such as vSphere®, AWS®, or Azure® create, or "provision," the machines themselves before Kubernetes is deployed. On most infrastructures (including vSphere and cloud providers), DKP provisions the actual nodes automatically as part of deploying a cluster. It creates the VMs using the appropriate image, then handles the networking and installation of Kubernetes.

However, DKP can also work with pre-provisioned infrastructure in which *you* provision the VMs for the nodes themselves. Note that you can pre-provision nodes for DKP on bare metal, on vSphere, or even in the cloud. Pre-provisioned and vSphere are a combination of the physical (on-premises bare metal) and virtual servers (VMware® vSphere).

3.1.5.1 Where is Pre-provisioned used?

Pre-provisioned environments are often used in bare metal deployments, where you deploy [your OS \(see page 69\)](#) (such as RHEL® or Ubuntu™, and so on) on physical machines. Creating a pre-provisioned cluster means that you, as an **Infrastructure Ops Manager**, are responsible for allocating your compute resources, setting up networking, collecting IP and SSH information to be provided to DKP, etc. You can then provide all of th

needed details to the pre-provisioned provider to deploy Kubernetes. These operations are done manually or with the help of other tools.

In pre-provisioned environments, DKP handles your cluster's lifecycle (installation, upgrade, node management, and so on). DKP installs Kubernetes, monitoring and logging apps, as well as its own user interface.

The main use cases for the pre-provisioned provider are:


- On-premises clusters
- Cloud or IAAS environments that do not currently have a D2iQ® supported infrastructure provider
- Cloud environments where you must use predefined infrastructure instead of having one of the supported cloud providers create it for you

In an environment with access to the Internet, you retrieve artifacts from specialized repositories dedicated to them, such as Docker® images contained in DockerHub, and Helm™ Charts that come from a dedicated Helm Chart repository. However, in an air-gapped environment, you need local repositories to store Helm charts, Docker images, and other artifacts. Tools such as JFrog™, Harbor™, and Nexus™ handle multiple types of artifacts in one local repository.

3.1.5.2 Next Topic:

[Cluster Types and Concepts](#) (see page 81)

3.2 Provide Context for Commands with a kubeconfig File

 This page contains some basic recommendations regarding `kubeconfig` as it relates to target clusters and the `--kubeconfig=${CLUSTER_NAME}.conf` flag. Refer to the [Kubernetes documentation](#)⁴³ for more information.

For `kubectl` and `dkp` commands to run, it is often necessary to specify the environment or cluster in which you want to run them. This specially applies to commands that create, delete, or update a cluster's resources.

To specify the context, there are two options:

Export an environment variable	Specify the target cluster in the command
Export an environment variable from a cluster's <code>kubeconfig</code> file, which sets the environment for the commands you run after exporting it.	Specify an environment variable for one command at the time by running it with the <code>--kubeconfig=\${CLUSTER_NAME}.conf</code> flag.

⁴³ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

Export an environment variable	Specify the target cluster in the command
<ul style="list-style-type: none"> • Better suited for single-cluster environments 	<ul style="list-style-type: none"> • Better suited for multi-cluster environments

3.2.1 Single-cluster Environment

In a single-cluster environment, you don't need to switch between clusters to execute commands and perform operations. However, it is still necessary to specify an environment for each terminal session, so the DKP CLI runs the operations on the cluster with the Kommander component.

Use a kubeconfig file to set the environment variable for all your operations

3.2.1.1 Set the environment variable for all your operations

1. When you create a cluster, a `kubeconfig` file is generated automatically. Get this `kubeconfig` file and write it to the `${CLUSTER_NAME}.conf` variable:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

After setting this configuration, whenever you use run a command with the variable `${CLUSTER_NAME}.conf`, it will automatically run on your cluster.

2. Set the context by exporting your `kubeconfig` from the source file:
:note: Execute this command for each terminal session.

```
export KUBECONFIG=${CLUSTER_NAME}.conf
```

For the current terminal session, you don't need to provide the `--kubeconfig` flag.

3.2.2 Multi-cluster Environment


Since having multiple clusters means switching between them to run operations, D2iQ recommends using the flag to set the context.

Use a flag to reference the target cluster

3.2.2.1 Use a flag to reference the target cluster

The `--kubeconfig ${CLUSTER_NAME}.conf` flag defines the configuration file for the cluster that you are configuring and accessing.

When operating and using multiple clusters, this is the easiest way to ensure that you are working on the correct cluster. If you create additional clusters and do not store the name as an environment variable, you can type out the name of the cluster followed by `.conf` to access your cluster.

 Ensure you run `dkp get kubeconfig` for each cluster you create to generate a `kubeconfig` file.


Example flag:

```
--kubeconfig azurecluster1.conf
```

Example command with flag:

```
dkp install kommander --kubeconfig azurecluster1.conf
```

Advanced: Use a kubeconfig file to configure access to multiple clusters using contexts

 For advanced users only.

It is possible to set up a `kubeconfig` file to manage access to several clusters. For more information, refer to <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>.

You can also set your environment variable to multiple `kubeconfig` files by merging them. For more information, refer to <https://kubernetes.io/docs/concepts/configuration/organize-cluster-access-kubeconfig/>.

3.2.3 Next Topic:

[Storage](#) (see page 89)

3.3 Storage

An introduction to persistent storage in Kubernetes

This document describes the model used in Kubernetes for managing persistent, cluster-scoped storage for workloads requiring access to persistent data.

A workload on Kubernetes typically requires two types of storage:

- Ephemeral Storage
- Persistent Volume

3.3.1 Ephemeral Storage

Ephemeral storage, by its name, is ephemeral in the sense that it is cleaned up when the workload is deleted or the container crashes. For example, the following are examples of ephemeral storage provided by Kubernetes:

EmptyDir volume.	Managed by kubelet under <code>/var/lib/kubelet</code> .
Container logs.	Typically under <code>/var/logs/containers</code> .
Container image layers.	Managed by container runtime (e.g., under <code>/var/lib/containerd</code>).
Container writable layers.	Managed by container runtime (e.g., under <code>/var/lib/containerd</code>).

Ephemeral storage is automatically managed by Kubernetes, and typically does not require explicit settings. You may need to express the capacity requests for ephemeral storage so that `kubelet` can use that information to make sure it does not run out of ephemeral storage space on each node.

3.3.2 Persistent Volume

Persistent Volumes are storage resources that can be used by the cluster. Persistent Volumes are volume plug-ins that have lifecycle capabilities that are independent of any Kubernetes Pod or Deployment. A Kubernetes persistent volume (PV) is an object that allows pods to access persistent storage on a storage device and defined via a Kubernetes StorageClass. Unlike regular volumes, which are transient in nature, PVs are persistent, supporting stateful application use cases.

You may have stateful workloads requiring persistent storage whose lifecycle is longer than that of Pods or containers. For instance, a database server needs to recover database files after it crashes. For those cases, the workloads need to use PersistentVolumes (PV).

Persistent Volumes are resources that represent storage in the cluster that has been provisioned by an administrator or dynamically provisioned using Storage Classes. Unlike ephemeral storage, the lifecycle of a PersistentVolume is independent of that of the workload that uses it.

The Persistent Volume API objects capture the details of the implementation of the storage, be that NFS, iSCSI, or a cloud-provider-specific storage system. In order to use a Persistent Volume (PV), your application needs to invoke a Persistent Volume Claim (PVC).

3.3.2.1 Persistent Volume Claim

A PersistentVolumeClaim is a request for storage. For a workload that requires persistent volumes, the workload should use PersistentVolumeClaim (PVC) to express its request on persistent storage. A PersistentVolumeClaim can request specific size and [Access Modes](#)⁴⁴ (for example, they can be mounted after read/write or many times read-only).

Any workload can specify a PersistentVolumeClaim. For example, a Pod may need a volume that is at least 4Gi large or a volume mounted under `/data` in the container's filesystem. If there is a PersistentVolume (PV) that satisfies the specified requirements in the PersistentVolumeClaim (PVC), it will be bound to the PVC before the Pod starts.

3.3.2.2 Related Information

- [Kubernetes Storage](#)⁴⁵
- [Kubernetes persistent storage design document](#)⁴⁶
- [Default Storage Providers in DKP](#) (see page 91)

3.3.3 Next Step:

[Resource Requirements](#) (see page 100)

3.3.4 Default Storage Providers in DKP

Default storage providers in DKP

When deploying DKP using a supported cloud provisioner (AWS, Azure, GCP), DKP automatically configures native storage drivers for the target platform. In addition, DKP deploys a default [StorageClass](#)⁴⁷ for <https://kubernetes.io/docs/concepts/storage/dynamic-provisioning/> creation. The table below lists the driver and default StorageClass for each supported cloud provisioner.

Cloud Provisioner	Version	Driver	Default Storage Class
AWS	1.23	aws-ebs-csi-driver ⁴⁸	<code>ebs-sc</code>
Azure	1.29.0	azuredisk-csi-driver ⁴⁹	<code>azuredisk-sc</code>

⁴⁴ <https://kubernetes.io/docs/concepts/storage/persistent-volumes/#access-modes>

⁴⁵ <https://kubernetes.io/docs/concepts/storage/>

⁴⁶ <https://github.com/kubernetes/design-proposals-archive/tree/main/storage>

⁴⁷ <https://kubernetes.io/docs/concepts/storage/storage-classes/>

⁴⁸ <https://github.com/kubernetes-sigs/aws-ebs-csi-driver>

⁴⁹ <https://github.com/kubernetes-sigs/azuredisk-csi-driver>

Cloud Provisioner	Version	Driver	Default Storage Class
Pre-provisioned	2.5.0	local-static-provisioner ⁵⁰	localvolume-provisioner
vSphere	2.7.2	vsphere-csi-driver ⁵¹	vsphere-raw-block-sc
Cloud Director (VCD)	main	cloud-director-named-csi-driver ⁵²	vcd-disk-sc
GCP	1.10.3	gcp-compute-persistent-disk-csi-driver ⁵³	csi-gce-pd



DKP uses the local static provisioner as the default storage provider for Pre-provisioned clusters. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)⁵⁴ compatible storage that is suitable for production. You can choose from any of the [storage options](#)⁵⁵ available for Kubernetes. To disable the default that Konvoy deploys, set the default StorageClass `localvolume-provisioner` as non-default. Then set your newly created StorageClass to be the default by following the commands in the Kubernetes documentation called [Changing the Default Storage Class](#)⁵⁶.

When a default StorageClass is specified, you can create **PersistentVolumeClaims** (PVCs) without needing to specify the storage class. For instance, to request a volume using the default provisioner, create a PVC with the following:

```
apiVersion: v1
kind: PersistentVolumeClaim
metadata:
  name: my-pv-claim
spec:
  accessModes:
    - ReadWriteOnce
  resources:
```

⁵⁰ <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

⁵¹ <https://github.com/kubernetes-sigs/vsphere-csi-driver>

⁵² <https://github.com/vmware/cloud-director-named-disk-csi-driver>

⁵³ <https://github.com/kubernetes-sigs/gcp-compute-persistent-disk-csi-driver/releases>

⁵⁴ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>


⁵⁵ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁵⁶ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

```
requests:
  storage: 4Gi
```

To start the provisioning of a volume, launch a pod which references the PVC:

```
...
  volumeMounts:
  - mountPath: /data
    name: persistent-storage
...
  volumes:
  - name: persistent-storage
    persistentVolumeClaim:
      claimName: my-pv-claim
```

 To specify a StorageClass that references a storage policy when making a PVC⁵⁷, you can name it in `storageClassName`. If left blank, the default StorageClass is used.

3.3.4.1 Change or Manage Multiple StorageClasses

The default `StorageClass` provisioned with DKP is acceptable for install, but may not be suitable for production. If your workload has different requirements, you can create additional `StorageClass` types with specific configurations. You can change the default `StorageClass` using these steps from the Kubernetes site:

- [Changing the default storage class](#)⁵⁸

Ceph can also be used as CSI storage. Refer to that section in the documentation for information on how to use [Rook Ceph](#)⁵⁹.

3.3.4.2 Driver Information

All default drivers implement the [Container Storage Interface](#)⁶⁰ (CSI). The CSI provides a common abstraction to container orchestrators for interacting with storage subsystems of various types. Each driver has specific configuration parameters which effect PV provisioning. This section details the default configuration for drivers used with DKP. This section also has links to driver documentation, if further customization is required.

⁵⁷ <https://kubernetes.io/docs/concepts/storage/persistent-volumes/#class-1>

⁵⁸ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

⁵⁹ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+Rook+Ceph+in+DKP)

[fromPageId=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+Rook+Ceph+in+DKP](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+Rook+Ceph+in+DKP)

⁶⁰ <https://github.com/container-storage-interface/spec/blob/master/spec.md>

NOTE: `StorageClass` parameters cannot be changed after creation. To use a different volume configuration, you must create a new `StorageClass`.

3.3.4.2.1 Amazon Elastic Block Store (EBS) CSI Driver

DKP EBS default `StorageClass`:

```
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  annotations:
    storageclass.kubernetes.io/is-default-class: "true" # This tells kubernetes to
    make this the default storage class
  name: ebs-sc
provisioner: ebs.csi.aws.com
reclaimPolicy: Delete # volumes are automatically reclaimed when no longer in use
and PVCs are deleted
volumeBindingMode: WaitForFirstConsumer # Physical volumes will not be created until
a pod is created that uses the PVC, required to use CSI's Topology feature
parameters:
  csi.storage.k8s.io/fstype: ext4
  type: gp3 # General Purpose SSD
```

DKP deploys with gp3 (general purpose SSDs) EBS volumes.

- Driver documentation: [aws-ebs-csi-driver](#)⁶¹
- Volume types and pricing: [volume types](#)⁶²

3.3.4.2.2 Azure CSI Driver

DKP deploys with StandardSSD_LRS for Azure Virtual Disks.

- Driver documentation: [azuredisk-csi-driver](#)⁶³
- Volume types and pricing: [volume types](#)⁶⁴

⁶¹ <https://github.com/kubernetes-sigs/aws-ebs-csi-driver>

⁶² <https://aws.amazon.com/ebs/features/>

⁶³ <https://github.com/kubernetes-sigs/azuredisk-csi-driver>

⁶⁴ <https://learn.microsoft.com/en-us/azure/aks/azure-disk-csi>

- Specifics for Azure using Pre-provisioning can be found here: <https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Pre-provisioned%20Azure%20only%20Configurations&linkCreation=true&fromPagelD=273743912>⁶⁵

3.3.4.2.3 vSphere CSI Driver

DKP default storage class for vSphere supports dynamic provisioning and static provisioning of block volumes.

- Driver documentation: <https://docs.vmware.com/en/VMware-vSphere-Container-Storage-Plug-in/index.html>
- Specifics for using vSphere storage driver: <https://docs.vmware.com/en/VMware-vSphere-Container-Storage-Plug-in/2.0/vmware-vsphere-csp-getting-started/GUID-5D144DA0-4806-4DEB-8819-10A1C42E38AB.html>

3.3.4.2.4 Cloud Director

In VMware Cloud Director:

- The [Cloud provider component \(CPI\)](#)⁶⁶ purpose is to manage the lifecycle of the load balancers as well as associate Kubernetes nodes with virtual machines in the infrastructure.
- CAPVCD is a component that runs in a Kubernetes cluster that connects to the VCD API. It uses the CPI to do the work of creating and managing the infrastructure.
- For storage, DKP has a CSI plugin for interfacing with CPI and can create disks dynamically and becomes the StorageClass

3.3.4.2.5 Pre-provisioned CSI Driver

In a [pre-provisioned](#)⁶⁷ environment, DKP will also deploy a CSI compatible driver and configure a default StorageClass - `localvolumeprovisioner`.

- Driver documentation: [local-static-provisioner](#)⁶⁸

⁶⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelD=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+Pre-provisioned+Azure+only+Configurations>

⁶⁶ <https://github.com/vmware/cloud-provider-for-cloud-director/>

⁶⁷ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelD=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+What+is+Pre-provisioned+Infrastructure>

⁶⁸ <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

[-] DKP uses [local-static-provisioner](#)⁶⁹ (`localvolume-provisioner`) as the default storage provider for a pre-provisioned environment. However, [local-static-provisioner](#)⁷⁰ is not suitable for production use. You should use a different [Kubernetes CSI](#)⁷¹ compatible storage that is suitable for production.

To disable the default that Konvoy deploys, set the default StorageClass `localvolume-provisioner` as non-default. Then set your newly created StorageClass by following the steps in the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>. You can choose from any of the storage options available for Kubernetes and make your [storage choice](#)⁷² the default storage.

Ceph can also be used as CSI storage. Refer to that section in the documentation for information on how to use [Rook Ceph](#)⁷³.

3.3.4.2.6 GCP CSI Driver

This driver allows volumes backed by Google Cloud Filestore instances to be dynamically created and mounted by workloads.

- Driver documentation: [gcp-filestore-csi-driver](#)⁷⁴
- Persistent volumes and dynamic provisioning: [volume types](#)⁷⁵

3.3.4.3 Related Information

- [Kubernetes Storage](#)⁷⁶
- [Kubernetes CSI Storage Drivers](#)⁷⁷
- [Kubernetes Local Persistent Volumes](#)⁷⁸
- [Persistent Volumes](#)⁷⁹
- [DKP vSphere Storage](#)⁸⁰

⁶⁹ <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

⁷⁰ <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

⁷¹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁷² <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁷³ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+BYOS+%28Bring+Your+Own+Storage%29+to+DKP+Clusters)

[fromPageId=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+BYOS+%28Bring+Your+Own+Storage%29+to+DKP+Clusters](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+BYOS+%28Bring+Your+Own+Storage%29+to+DKP+Clusters)

⁷⁴ <https://github.com/kubernetes-sigs/gcp-filestore-csi-driver>

⁷⁵ <https://cloud.google.com/kubernetes-engine/docs/concepts/persistent-volumes>

⁷⁶ <https://kubernetes.io/docs/concepts/storage/>

⁷⁷ <https://kubernetes-csi.github.io/docs/drivers.html>

⁷⁸ <https://kubernetes.io/blog/2019/04/04/kubernetes-1.14-local-persistent-volumes-ga/>


⁷⁹ <https://kubernetes.io/docs/concepts/storage/persistent-volumes/#class-1>

⁸⁰ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+vSphere%3A+Storage+Options)

[fromPageId=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+vSphere%3A+Storage+Options](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273743912&linkCreation=true&spaceKey=DENT&title=%282.6%29+vSphere%3A+Storage+Options)

3.3.5 Provision a Static Local Volume

Learn how to provision a static local volume for a DKP cluster

 When creating a pre-provisioned infrastructure cluster, DKP uses `localvolumeprovisioner` as the [default storage provider](#) (see [page 91](#)). However, `localvolumeprovisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)⁸¹ compatible storage that is suitable for production.

You can choose from any of the storage options available for Kubernetes. To disable the default that Konvoy deploys, set the default StorageClass `localvolumeprovisioner` as non-default. Then set your newly-created StorageClass by following the steps in the Kubernetes documentation, [Changing the Default Storage Class](#)⁸².

The `localvolumeprovisioner` component uses the [local volume static provisioner](#)⁸³ to manage persistent volumes for pre-allocated disks. It does this by watching the `/mnt/disks` folder on each host and creating persistent volumes in the `localvolumeprovisioner` storage class for each disk that it discovers in this folder.

- Persistent volumes with a 'Filesystem' volume-mode are discovered if you mount them under `/mnt/disks`.
- Persistent volumes with a 'Block' volume-mode are discovered if you create a symbolic link to the block device in `/mnt/disks`.

For additional DKP documentation regarding StorageClass, see: [Default Storage Providers in DKP](#) (see [page 91](#))

3.3.5.1 Before you Begin

Before starting, verify the following:

- You have access to a Linux, macOS, or Windows computer with a supported operating system version.
- You have a provisioned `dkp` cluster that uses the `localvolumeprovisioner` platform application, but have not added any other Kommander applications to the cluster yet.

This distinction between provisioning and deployment is important because some applications depend on the storage class provided by the `localvolumeprovisioner` component and can fail to start if it is not configured yet.

⁸¹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁸² <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

⁸³ <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

3.3.5.2 Provision the Cluster and a Volume

1. Create a pre-provisioned cluster by following the steps outlined in the [Pre-provisioned Infrastructure](#) (see page 1200) topic.

As volumes are created/mounted on the nodes, the local volume provisioner detects each volume in the `/mnt/disks` directory and adds it as a persistent volume with the `localvolumeprovisioner` storage class. For more information, see the documentation regarding [Kubernetes Local Storage](#)⁸⁴.

2. Create at least one volume in `/mnt/disks` on each host.

For example, mount a `tmpfs` volume:

```
mkdir -p /mnt/disks/example-volume && mount -t tmpfs example-volume /mnt/disks/example-volume
```

3. Verify the persistent volume by running the following command:

```
kubectl get pv
```

The command displays output similar to the following:

NAME	CAPACITY	ACCESS MODES	RECLAIM POLICY	STATUS
CLAIM	STORAGECLASS	REASON	AGE	
local-pv-4c7fc8ba	3986Mi	RWO	Delete	Available
localvolumeprovisioner		2s		

4. Claim the persistent volume using a PersistentVolumeClaim, by running the following command:

```
cat <<EOF | kubectl create -f -
kind: PersistentVolumeClaim
apiVersion: v1
metadata:
  name: example-claim
spec:
  accessModes:
  - ReadWriteOnce
  resources:
    requests:
      storage: 100Mi
```

⁸⁴ <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner/blob/master/docs/operations.md#link-devices-into-directory-to-be-discovered-as-block-pvs>

```
storageClassName: localvolumeprovisioner
EOF
```

5. Reference the persistent volume claim in a pod by running the following command:

```
cat <<EOF | kubectl create -f -
apiVersion: v1
kind: Pod
metadata:
  name: pod-with-persistent-volume
spec:
  containers:
    - name: frontend
      image: nginx
      volumeMounts:
        - name: data
          mountPath: "/var/www/html"
  volumes:
    - name: data
      persistentVolumeClaim:
        claimName: example-claim
EOF
```

6. Verify the persistent volume claim by running the following command:

```
kubectl get pvc
```

The command displays output similar to the following:

NAME	STATUS	VOLUME	CAPACITY	ACCESS MODES
STORAGECLASS		AGE		
example-claim	Bound	local-pv-4c7fc8ba	3986Mi	RWO
localvolumeprovisioner		78s		

And you can also check the persistent volume:

NAME	CAPACITY	ACCESS MODES	RECLAIM POLICY	STATUS
CLAIM		STORAGECLASS	REASON	AGE
local-pv-4c7fc8ba	3986Mi	RWO	Delete	Bound
default /example-claim		localvolumeprovisioner		15m

Upon deletion of the persistent volume claim, the corresponding persistent volume resource uses the *Delete* reclaim policy, which removes all data on the volume.

3.4 Resource Requirements

To ensure a successful DKP install, there are certain requirements to be met in regards to Control Plane nodes as well as Worker nodes. These resource requirements can be slightly different for different Infrastructure Providers.

- [Managed Cluster Requirements \(see page 101\)](#)
- [Management Cluster Application Requirements \(see page 103\)](#)
- [Workspace Platform Application Defaults and Resource Requirements \(see page 105\)](#)



The specific infrastructure provider resources are listed in their install sections which are referenced below.

3.4.1 General Resource Requirements

To [Install DKP \(see page 108\)](#) with the minimum amount of resources, please see the list below before beginning installation. You need at least three control plane nodes and you must have at least four worker nodes. The specific number of worker nodes required for your environment can vary depending on the cluster workload and size of the nodes.

	Control Plane Nodes	Worker Nodes
vCPU Count	4	8
Memory	16 GB	32 GB
Disk Volume	Approximately 80 GiB <ul style="list-style-type: none"> • used for <code>/var/lib/kubelet</code> and <code>/var/lib/containerd</code> 	Approximately 80 GiB <ul style="list-style-type: none"> • used for <code>/var/lib/kubelet</code> and <code>/var/lib/containerd</code>
Root Volume	Disk usage must be below 85%	Disk usage must be below 85%

- If you use the instructions to create a cluster using the DKP default settings without any edits to configuration files or additional flags, your cluster is deployed on an [Ubuntu 20.04 operating system image](#) (see page 69) with 3 control plane nodes, and 4 worker nodes which match the requirements above.

3.4.2 Infrastructure Provider Specific

For Specific Infrastructure Providers additional requirements may apply. For example, DKP on Azure defaults to deploying a `Standard_D4s_v3` virtual machine with an 128 GiB volume for the OS and an 80GiB volume for etcd storage, which meets the above requirements. See Main Page for Provider for infrastructure provider specific additional resource information:

- [Pre-provisioned Install Options](#) (see page 149)
- [AWS Install Options](#) (see page 293)
- [Azure Install Options](#) (see page 531)
- [vSphere Install Options](#) (see page 433)
- [VMware Cloud Director Install Options](#) (see page 529)

3.4.3 Kommander Component Requirements:

- [Management Cluster Application Requirements](#) (see page 103)
- [Workspace Platform Application Defaults and Resource Requirements](#) (see page 105)

3.4.3.1 Next Topic:

[Install Overview](#) (see page 108)

3.4.4 Managed Cluster Requirements

To create additional clusters that will be managed by your [Management cluster](#) (see page 81), please ensure you have at least the minimal recommendation of required resources below.

3.4.4.1 Minimum Recommendation for Managed Clusters:

Worker Nodes - At least four worker nodes with:

- 8 CPU cores each
- 12Gi of memory
- A storage class and disk mounts that can accommodate at least four persistent volumes

- NOTE: Four worker nodes are required in order to support upgrades to the `rook-ceph` platform application. `rook-ceph` is used to support the logging stack, the `velero` backup tool, and DKP Insights. If you have disabled the `rook-ceph` platform application, only three worker nodes are required.

Control Plane Nodes:

- 8 CPU cores each
- 12Gi of memory

- NOTE: One control plane node is acceptable for a non-critical test environment, but any production workload must have at least three control planes.

The cluster needs:

- default Storage Class and four volumes of 32G, 32G, 2G, and 100G or the ability to create those volumes depending on the Storage Class:

```
$ kubectl get pv -A
NAME                                CAPACITY  ACCESS MODES  RECLAIM POLICY
STATUS  CLAIM
STORAGECLASS  REASON  AGE
pvc-08de8c06-bd66-40a3-9dd4-b0aece8ccbe8  32Gi      RWO           Delete
Bound      kommander-default-workspace/kubecost-cost-analyzer
ebs-sc     124m
pvc-64552486-7f4c-476a-a35d-19432b3931af  32Gi      RWO           Delete
Bound      kommander-default-workspace/kubecost-prometheus-server
ebs-sc     124m
pvc-972c3ee3-20bd-449b-84d9-25b7a06a6630  2Gi       RWO           Delete
Bound      kommander-default-workspace/kubecost-prometheus-alertmanager
ebs-sc     124m
pvc-98ab93f1-2c2f-46b6-b7d3-505c55437fbb  100Gi     RWO           Delete
Bound      kommander-default-workspace/db-prometheus-kube-prometheus-stack-prometheus-0
ebs-sc     123m
```

- NOTE: Actual workloads may demand more resources depending on usage.

3.4.5 Management Cluster Application Requirements


3.4.5.1 Management cluster application minimum resources and storage requirements

This section only details requirements for management cluster-specific applications in the Kommander component. For the list of all platform applications, see [Platform Application Configuration Requirements](#) (see page 639).

This table describes the workspace platform applications specific to the management cluster, minimum resource requirements, minimum persistent storage requirements, and default priority class values.

Common Name	App ID (for App versions, see the Release Notes (see page 1760))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required	Default Priority Class
Centralized Grafana*	centralized-grafana	Yes	cpu: 200m memory: 100Mi		DKP Critical (100002000)
Centralized Kubecost*	centralized-kubecost	Yes	cpu: 1200m memory: 4151Mi	# of PVs: 1 PV sizes: 32Gi	DKP High (100001000)
Chartmuseum	chartmuseum	Yes		# of PVs: 1 PV sizes: 2Gi	DKP Critical (100002000)
Dex	dex	Yes	cpu: 100m memory: 50Mi		DKP Critical (100002000)
Dex Authenticator	dex-k8s-authenticator	Yes	cpu: 100m memory: 128Mi		DKP High (100001000)
DKP Insights Management	dcp-insights-management	No	cpu: 100m memory: 128Mi		DKP Critical (100002000)

Common Name	App ID (for App versions, see the Release Notes (see page 1760))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required	Default Priority Class
Gitea	gitea	Yes	cpu: 500m memory: 512Mi	# of PVs: 2 PV sizes: 10Gi	DKP Critical (100002000)
Karma*	karma	Yes			DKP Critical (100002000)
Kommander	kommander	No	cpu: 1100m memory: 896Mi		DKP Critical (100002000)
Kommander AppManagement	kommander-appmanagement	Yes	cpu: 300m memory: 256Mi		DKP Critical (100002000)
Kommander Flux	kommander-flux	Yes	cpu: 5000m memory: 5Gi		DKP Critical (100002000)
Kommander UI	kommander-ui	No	cpu: 100m memory: 256Mi		DKP Critical (100002000)
Kubefed	kubefed	Yes	cpu: 300m memory: 192Mi		DKP Critical (100002000)
Kubetunnel	kubetunnel	Yes	cpu: 200m memory: 148Mi		DKP Critical (100002000)
Thanos*	thanos	Yes			DKP Critical (100002000)
Traefik ForwardAuth	traefik-forward-auth-mgmt	Yes	cpu: 100m memory: 128Mi		DKP Critical (100002000)

 Applications designated with an asterisk (“*”) are DKP Enterprise only apps that are deployed by default for DKP Enterprise Customers only.

3.4.6 Workspace Platform Application Defaults and Resource Requirements

The following table provides a list of platform applications that are available in DKP with the Kommander component. Some of them are deployed by default on attachment, others require [manual installation](#) (see [page 639](#)).

Workspace platform applications require more resources than solely deploying or attaching clusters into a workspace. Your cluster must have sufficient resources when deploying or attaching to ensure that the platform services are installed successfully.

The following table describes all the workspace platform applications that are available to the clusters in a workspace, minimum resource requirements, whether they are enabled by default, and their default priority classes.

Common Name	App ID (for App versions, see the Release Notes (see page 1760))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required	Default Priority Class
Cert Manager	cert-manager	Yes	cpu: 10m memory: 32Mi		system-cluster-critical (2000000000)
External DNS	external-dns	No			DKP High (100001000)
Fluent Bit	fluent-bit	No	cpu: 350m memory: 350Mi		DKP Critical (100002000)
Gatekeeper	gatekeeper	Yes	cpu: 300m memory: 768Mi		system-cluster-critical (2000000000)
Grafana	grafana-logging	No	cpu: 200m memory: 100Mi		DKP Critical (100002000)

Common Name	App ID (for App versions, see the Release Notes (see page 1760))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required	Default Priority Class
Loki	grafana-loki	No		# of PVs: 8 PV sizes: 10Gi x 8 (total: 80Gi)	DKP Critical (100002000)
Istio	istio	No	cpu: 1270m memory: 4500Mi		DKP Critical (100002000)
Jaeger	jaeger	No			DKP High (100001000)
Kiali	kiali	No	cpu: 20m memory: 128Mi		DKP High (100001000)
Knative	knative	No	cpu: 610m memory: 400Mi		DKP High (100001000)
Kube OIDC Proxy	kube-oidc-proxy	Yes			DKP Critical (100002000)
Kube Prometheus Stack	kube-prometheus-stack	Yes	cpu: 1300m memory: 4300Mi	# of PVs: 1 PV sizes: 100Gi	DKP Critical (100002000)
Kubecost	kubecost	Yes	cpu: 700m memory: 1700Mi	# of PVs: 3 PV sizes: 2Gi, 32Gi, 32Gi (total: 66Gi)	DKP High (100001000)
Kubernetes Dashboard	kubernetes-dashboard	Yes	cpu: 250m memory: 300Mi		DKP High (100001000)

Common Name	App ID (for App versions, see the Release Notes (see page 1760))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required	Default Priority Class
Logging Operator	logging-operator	No	cpu: 350m * # of nodes + 600m memory: 228Mi + 350Mi * # of nodes	# of PVs: 1 PV sizes: 10Gi	DKP Critical (100002000)
NFS Server Provisioner	nfs-server-provisioner	No	# of PVs: 1 PV sizes: 100Gi		DKP High (100001000)
NVIDIA GPU Operator	nvidia-gpu-operator	No	cpu: 100m memory: 128Mi		system-cluster-critical (2000000000)
Prometheus Adapter	prometheus-adapter	Yes	cpu: 1000m memory: 1000Mi		DKP Critical (100002000)
Reloader	reloader	Yes	cpu: 100m memory: 128Mi		DKP High (100001000)
Rook Ceph	rook-ceph	Yes	cpu: 100m memory: 128Mi		system-cluster-critical (2000000000)
Rook Ceph Cluster	rook-ceph-cluster	Yes	cpu 2500m mem 8Gi	# of PVs: 4 PV sizes: 40Gi	DKP Critical (100002000) system-cluster-critical (2000000000) system-node-critical
Traefik	traefik	Yes	cpu: 500m		DKP Critical (100002000)

Common Name	App ID (for App versions, see the Release Notes (see page 1760))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required	Default Priority Class
Traefik ForwardAuth	traefik-forward-auth	Yes	cpu: 100m memory: 128Mi		DKP Critical (100002000)
Velero	velero	No	cpu: 1000m memory: 1024Mi		DKP Critical (100002000)



- Currently, DKP only supports a single deployment of `cert-manager` per cluster. Because of this, `cert-manager` cannot be installed on any Konvoy managed clusters or clusters that come with `cert-manager` pre-installed.
- Only a single deployment of `traefik` per cluster is supported.
- DKP automatically manages the deployment of `traefik-forward-auth` and `kube-oidc-proxy` when clusters are attached to the workspace. These applications are not shown in the DKP UI.
- Applications are enabled in DKP and then deployed to attached clusters. To confirm that your enabled application has successfully deployed, you should [verify via the CLI](#) (see page 674).

3.4.6.1 Next Topic:

[Install Overview](#) (see page 108)

3.5 Install Overview

3.5.1 Prepare Your Environment for Install:

Follow the steps below to install the basic package requirements for your environment to achieve a successful install of DKP. Next install DKP, and then you can begin any custom configurations based on your environment.

1. Install required packages. In most cases, you can install the required software using your preferred package manager. For example, on a macOS computer, you can use [Homebrew](#)⁸⁵ to install `kubectl` and the `aws` command-line utility by running the following command:

```
brew install kubernetes-cli awscli
```

2. Check the Kubernetes client version. Many important Kubernetes functions **do not work** if your client is outdated. You can verify that the version of `kubectl` you have installed is supported by running the following command:

```
kubectl version --short=true
```

3. Check [supported Kubernetes versions](#) (see page 117) after finding your version with the command above.
4. For air-gapped, create a [bastion host](#) (see page 1285) for the cluster nodes to use within the air-gapped network. This bastion host needs access to a local registry in lieu of an Internet connection for pulling images. The recommended template naming pattern is `./folder-name/dkp-e2e-bastion-template` or similar. Each infrastructure provider has its own set of bastion host instructions. Refer to your provider's site for details - [Azure](#)⁸⁶, [AWS](#)⁸⁷, [GCP](#)⁸⁸, or [vSphere](#)⁸⁹.
5. For creating DKP machine images, you need to download [Konvoy Image Builder](#) (see page 1461) and extract it.
6. To download DKP, see the [Download](#) (see page 78) topic for information. You will need to download and extract the DKP binary package tarball.
7. Verify you have valid **cloud provider security credentials** to deploy the cluster on that platform.



This step regarding provider security credentials is not required if you are installing DKP on an on-premises environment. For information about installing in an on-premises environment, see [Install Pre-provisioned](#) (see page 1200).

8. Install Konvoy depending on which infrastructure you have. Consult the [Day 1 - Basic Installs by Infrastructure](#) (see page 147) for provider specific installs. To use customized YAML and other advanced features, consult the [Custom Installation and Additional Infrastructure Tools](#) (see page 1036)

⁸⁵ <https://docs.brew.sh/Installation>

⁸⁶ <https://learn.microsoft.com/en-us/azure/bastion/quickstart-host-portal>

⁸⁷ <https://aws.amazon.com/solutions/implementations/linux-bastion/>

⁸⁸ <https://blogs.vmware.com/cloud/2021/06/02/intro-google-cloud-vmware-engine-bastion-host-access-iap/>

⁸⁹ <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-6975426F-56D0-4FE2-8A58-580B40D2F667.html>

section of the documentation for steps on creating a customized cluster with the Konvoy component of DKP using YAML. Find your infrastructure provider listed in that section.

9. Now that you have a basic Kubernetes cluster installed through Konvoy, configure the [Kommander](#) (see page 1394) component.
10. Lastly, continue with initializing the configuration file under the [Kommander Installer Configuration File](#) (see page 1422) component of DKP.

You may also want to test operations by deploying a simple, sample application, customizing the cluster configuration, or checking the status of cluster components.

3.5.2 Next Step:

[Prerequisites for Install](#) (see page 121)

3.5.3 Install OS Version Compatibility

Below is a list of Supported Operating Systems and the versions for the latest release.

3.5.3.1 Amazon Web Services (AWS)

Operating System	Kernel	Default Config (see page 294)	FIPS (see page 328)	Air Gapped (see page 310)	FIPS with Air Gapped (see page 344)	GPU Support (see page 363)	GPU with Air Gapped (see page 382)	Konvoy Image Builder (see page 1465)
CentOS 7.9 ⁹⁰	3.10.0-1160.80.1.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 7.9 ⁹¹	3.10.0-1160.80.1.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.4 ⁹²	4.18.0-305.88.1.el8_4.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes

⁹⁰ [https://wiki.centos.org/Manuals\(2f\)ReleaseNotes\(2f\)CentOS7\(2e\)2009.html](https://wiki.centos.org/Manuals(2f)ReleaseNotes(2f)CentOS7(2e)2009.html)

⁹¹ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/7.9_release_notes/index

⁹² https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/8.4_release_notes/index

Operating System	Kernel	Default Config (see page 294)	FIPS (see page 328)	Air Gapped (see page 310)	FIPS with Air Gapped (see page 344)	GPU Support (see page 363)	GPU with Air Gapped (see page 382)	Konvoy Image Builder (see page 1465)
RHEL 8.6 ⁹³	4.18.0-372.52.1.el8_6.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ubuntu 18.04 (Bionic Beaver) ⁹⁴	5.4.0-1103-aws	Yes						Yes
Ubuntu 20.04 (Focal Fossa) ⁹⁵	5.15.0-1040-aws	Yes				Yes		Yes
SLES 15 SP3 ⁹⁶	5.3.18-59.37-default	Yes				Yes		Yes
Oracle Linux RHCK 7.9 ⁹⁷	3.10.0-1160.92.1.0.1.el7.x86_64	Yes	Yes					Yes
Rocky Linux 9.1 ⁹⁸	5.14.0-162.12.1.el9_1.0.2.x86_64	Yes		Yes				Yes

⁹³ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/8.6_release_notes/index

⁹⁴ <https://wiki.ubuntu.com/BionicBeaver/ReleaseNotes>

⁹⁵ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

⁹⁶ <https://documentation.suse.com/en-us/sles/15-SP3/>

⁹⁷ <https://docs.oracle.com/en/operating-systems/oracle-linux/7/relnotes7.9/>

⁹⁸ https://docs.rockylinux.org/release_notes/9_1/

3.5.3.2 Microsoft Azure

Operating System	Kernel	Default Config (see page 532)	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1477)
Ubuntu 20.04 (Focal Fossa) ⁹⁹	5.15.0-1042-azure	Yes						Yes
Rocky Linux 9.1 ¹⁰⁰	5.14.0-162.12.1.el9_1.0.2.x86_64	Yes						Yes

3.5.3.3 Google Cloud Platform (GCP)

Operating System	Kernel	Default Config (see page 563)	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1481)
CentOS 7.9 ¹⁰¹	3.10.0-1160.62.1.el7.x86_64	Yes						Yes
Ubuntu 18.04 ¹⁰²	5.4.0-1072-gcp	Yes						Yes
Ubuntu 20.04 ¹⁰³	5.13.0-1024-gcp	Yes						Yes

⁹⁹ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

¹⁰⁰ https://docs.rockylinux.org/release_notes/9_0/

¹⁰¹ [https://wiki.centos.org/Manuals\(2f\)ReleaseNotes\(2f\)CentOS7\(2e\)2009.html](https://wiki.centos.org/Manuals(2f)ReleaseNotes(2f)CentOS7(2e)2009.html)

¹⁰² <https://wiki.ubuntu.com/BionicBeaver/ReleaseNotes>

¹⁰³ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

3.5.3.4 Pre-Provisioned

Pre-provisioned (see page 149) includes on-premises, vSphere, AWS, Azure, and GCP infrastructures and is described in more detail under **What is Pre-provisioned Infrastructure** (see page 86).

Operating System	Kernel	Default Config (see page 151)	FIPS (see page 195)	Air Gapped (see page 170)	FIPS with Air-Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1493)
CentOS 7.9 ¹⁰⁴	3.10.0-116.0.80.1.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 7.9 ¹⁰⁵	3.10.0-116.0.80.1.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.4 ¹⁰⁶	4.18.0-305.88.1.el8_4.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.6 ¹⁰⁷	4.18.0-372.52.1.el8_6.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Flatcar 3033.3 ¹⁰⁸ _x	5.10.186-flatcar	Yes						Yes
Ubuntu 20.04 ¹⁰⁹	5.15.0-1033-aws	Yes				Yes		Yes

104 [https://wiki.centos.org/Manuals\(2f\)ReleaseNotes\(2f\)CentOS7\(2e\)2009.html](https://wiki.centos.org/Manuals(2f)ReleaseNotes(2f)CentOS7(2e)2009.html)

105 https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/7.9_release_notes/index

106 https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/8.4_release_notes/index

107 https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/8.6_release_notes/index

108 <https://www.flatcar.org/releases/#lts-release>

109 <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

Operating System	Kernel	Default Config (see page 150)	FIPS (see page 195)	Air Gapped (see page 170)	FIPS with Air-Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1493)
SLES 15 ¹¹⁰	5.3.18-59.37-default	Yes				Yes		Yes
Oracle Linux RHCK 7.9 ¹¹¹	5.4.17-2136.314.6.3.el7uek.x86_64	Yes						Yes
Rocky Linux 9.1 ¹¹²	5.14.0-162.12.1.el9_1.0.2.x86_64	Yes		Yes				Yes

3.5.3.5 Pre-provisioned/Azure

Operating System	Kernel	Default Config (see page 151)	FIPS (see page 195)	Air Gapped (see page 170)	FIPS with Air Gapped (see page 215)	GPU Support (see page 239)	GPU with Air Gapped (see page 262)	Konvoy Image Builder (see page 1493)
RHEL 8.6 ¹¹³	4.18.0-372.52.1.el8_6.x86_64	Yes	Yes	Yes	Yes			Yes

¹¹⁰ <https://documentation.suse.com/en-us/sles/15-SP3/>

¹¹¹ <https://docs.oracle.com/en/operating-systems/oracle-linux/7/relnotes7.9/>

¹¹² https://docs.rockylinux.org/release_notes/9_1/

¹¹³ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/8.6_release_notes/index

3.5.3.6 vSphere

vCenter Server + ESXi

Operating System	Kernel	Default Config (see page 438)	FIPS (see page 483)	Air Gapped (see page 459)	FIPS with Air Gapped (see page 505)	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1487)
RHEL 7.9 ¹¹⁴	3.10.0-1160.el7.x86_64	Yes	Yes	Yes	Yes			Yes
RHEL 8.4 ¹¹⁵	4.18.0-305.el8.x86_64	Yes	Yes	Yes	Yes			Yes
RHEL 8.6 ¹¹⁶	4.18.0-372.9.1.el8.x86_64	Yes	Yes	Yes	Yes			Yes
Ubuntu 20.04 ¹¹⁷	5.4.0-125-generic	Yes		Yes				Yes
Rocky Linux 9.1 ¹¹⁸	5.14.0-162.6.1.el9_1.x86_64	Yes		Yes				Yes

¹¹⁴ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/7.9_release_notes/index

¹¹⁵ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html/8.4_release_notes/index

¹¹⁶ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/8/html-single/8.6_release_notes/index

¹¹⁷ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

¹¹⁸ https://docs.rockylinux.org/release_notes/9_1/

3.5.3.7 VMware Cloud Director (VCD) - [Beta Product](#)

Operating System	Kernel	Default Config (see page 1322)	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1487)
Ubuntu 20.04 ¹¹⁹	5.4.0-125-generic	Yes - Beta						

3.5.3.8 Amazon Elastic Kubernetes (EKS)

Operating System	Kernel	Default Config (see page 405)	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1476)
Amazon Linux 2 v7.9 ¹²⁰	5.10.178-162.673.amzn2.x86_64	Yes						

¹¹⁹ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

¹²⁰ <https://docs.aws.amazon.com/AL2/latest/relnotes/relnotes-al2.html>

3.5.3.9 Azure Kubernetes Service (AKS)

Operating System	Kernel	Default Config (see page 547)	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder (see page 1481)
Ubuntu 22.04.2 LTS ¹²¹	5.15.0-1037-azure	Yes						

3.5.4 Install Related Version Information

Other helpful links for versions required to successfully run with DKP for the latest version can be found below and also in their respective section of Documentation.

3.5.4.1 Supported Kubernetes Version

3.5.4.1.1 Deploying Clusters Versions

The latest supported Kubernetes versions for DKP are listed below. In DKP 2.6.2, the Kubernetes version installed by default is 1.26.14. The newest and oldest Kubernetes versions used with DKP must be within one minor version.

- Latest supported Kubernetes version is at **1.26.14 for deploying clusters**
- Other Kubernetes versions are supported at **1.25.13 for deploying clusters in EKS**

Kubernetes 1.26.14 enables you to benefit from the latest features and security fixes in upstream Kubernetes. This release comes with approximately 40 enhancements that you can benefit from like [Service internal tracking policy](#)¹²², [Kubelet credential provider](#)¹²³, among others. To read more about major features in this release, visit [GitHub](#)¹²⁴ and <https://kubernetes.io/blog/2022/12/09/kubernetes-v1-26-release/>.

3.5.4.1.2 Attaching Clusters Versions

DKP 2.6.2 supports attaching clusters with the following Kubernetes versions:

¹²¹ <https://discourse.ubuntu.com/t/jammy-jellyfish-release-notes/24668>

¹²² <https://github.com/kubernetes/enhancements/issues/2086>

¹²³ <https://github.com/kubernetes/enhancements/issues/2133>

¹²⁴ <https://github.com/kubernetes/sig-release/blob/master/releases/release-1.26/README.md>

Product	Compatible Kubernetes Versions
EKS	1.25.13
AKS	1.26.10
GKE	1.26.x



Attaching Kubernetes clusters with versions greater than n-1 is not recommended.



When attempting to attach a cluster with a lower Kubernetes version than the DKP default, you need to build and use an image with that lower version of Kubernetes. See [Konvoy Image Builder](#) (see page 1461) for documentation on building images. Failure to do this could result in an error.

3.5.4.2 Konvoy Image Builder

Refer to the [main page](#) (see page 1461) for a table of components versions and compatibility.

3.5.4.3 Storage Providers

Default storage providers in DKP

When deploying DKP using a supported cloud provisioner (AWS, Azure, GCP), DKP automatically configures native storage drivers for the target platform. In addition, DKP deploys a default [StorageClass](#)¹²⁵ for <https://kubernetes.io/docs/concepts/storage/dynamic-provisioning/> creation. The table below lists the driver and default StorageClass for each supported cloud provisioner.

Cloud Provisioner	Version	Driver	Default Storage Class
AWS	1.23	aws-ebs-csi-driver ¹²⁶	ebs-sc
Azure	1.29.0	azuredisk-csi-driver ¹²⁷	azuredisk-sc

¹²⁵ <https://kubernetes.io/docs/concepts/storage/storage-classes/>

¹²⁶ <https://github.com/kubernetes-sigs/aws-ebs-csi-driver>

¹²⁷ <https://github.com/kubernetes-sigs/azuredisk-csi-driver>

Cloud Provisioner	Version	Driver	Default Storage Class
Pre-provisioned	2.5.0	local-static-provisioner ¹²⁸	<code>localvolume-provisioner</code>
vSphere	2.7.2	vsphere-csi-driver ¹²⁹	<code>vsphere-raw-block-sc</code>
Cloud Director (VCD)	main	cloud-director-named-csi-driver ¹³⁰	<code>vcd-disk-sc</code>
GCP	1.10.3	gcp-compute-persistent-disk-csi-driver ¹³¹	<code>csi-gce-pd</code>



DKP uses the local static provisioner as the default storage provider for Pre-provisioned clusters. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)¹³² compatible storage that is suitable for production. You can choose from any of the [storage options](#)¹³³ available for Kubernetes. To disable the default that Konvoy deploys, set the default StorageClass `localvolume-provisioner` as non-default. Then set your newly created StorageClass to be the default by following the commands in the Kubernetes documentation called [Changing the Default Storage Class](#)¹³⁴.

3.5.4.4 Deploying a Cluster in FIPS mode

In order to create a cluster in FIPS mode, we must inform the bootstrap controllers of the appropriate image repository and version tags of the official D2iQ FIPS builds of Kubernetes.

¹²⁸ <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

¹²⁹ <https://github.com/kubernetes-sigs/vsphere-csi-driver>

¹³⁰ <https://github.com/vmware/cloud-director-named-disk-csi-driver>

¹³¹ <https://github.com/kubernetes-sigs/gcp-compute-persistent-disk-csi-driver/releases>

¹³² <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

¹³³ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

¹³⁴ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

3.5.4.4.1 Supported FIPS Builds

Component	Repository	Version
Kubernetes	docker.io/mesosphere ¹³⁵	v1.26.14+fips.0
etcd	docker.io/mesosphere ¹³⁶	3.5.10+fips.0

3.6 Prerequisites for Install

Before you create a DKP image and deploy the initial DKP cluster, the operator's machine is required to be either an OSX or Linux-based machine of a supported version. Ensure you have all the prerequisites met for both the Konvoy and Kommander component.

i Additional prerequisites are necessary for **air-gapped**, so verify that you have all non-air-gapped met and then any additional air-gapped prerequisites listed after non-air-gapped.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)

3.6.1 Prerequisites for the Konvoy Component

For DKP and Konvoy Image Builder to run, the operator machine requires:

3.6.1.1 Non-air-gapped (all environments)

In a non-air-gapped environment, your environment has two-way access to and from the **Internet**. Below are the prerequisites required if installing in a non-air-gapped environment:

- An x86_64-based Linux or macOS machine.
- The `dkp` binary on the bastion by [downloading](#) (see page 78). To check which version of DKP you installed for compatibility reasons, run the `dkp version` command ([dkp version](#) (see page 1749)).
- A Container engine/runtime installed is required to install DKP:

¹³⁵ <https://hub.docker.com/layers/mesosphere/kube-apiserver/v1.26.14/images/sha256-3c559c75123fa40d588bc2d9050938a0b7dd1dab4e38a4fd24847cd0c0f24694?context=explore>

¹³⁶ https://hub.docker.com/layers/mesosphere/etcd/v3.5.10_fips.0/images/sha256-61df32b8ec1e6229d8bd94caad07331c7e0c0b8e1947af8152c69eb2f04fb771?context=explore

- Version [Docker®](#)¹³⁷ container engine version 18.09.2 or higher installed for Linux or MacOS - On macOS, Docker runs in a virtual machine which needs configured with at least 8 GB of memory.
- Version 4.0 of [Podman](#)¹³⁸ or higher for Linux. Host requirements found here: [Host Requirements](#)¹³⁹
- [kubectl](#)¹⁴⁰ for interacting with the running cluster on the host where the DKP Konvoy command line interface (CLI) runs.
- [Konvoy Image Builder](#) (see page 1461) (KIB)
- Valid Provider account with credentials configured.
 - AWS [credentials configured](#)¹⁴¹ that can manage CloudFormation Stacks, IAM Policies, IAM Roles, and IAM Instance Profiles
 - Azure [credentials configured](#)¹⁴²
- CLI tooling of the cloud provider being used to deploy DKP commands:
 - [aws-cli](#)¹⁴³
 - [googlecloud-cli](#)¹⁴⁴
 - [azure-cli](#)¹⁴⁵
- EKS and AKS only a self-managed cluster is required:
 - If you follow [Day 1 - Basic Installation](#) (see page 147) instructions for install, you use the `--self-managed` flag and therefore have a self-managed cluster. If you use the instructions under [Custom Installation and Additional Infrastructure Tools](#) (see page 1036), ensure you perform the self-managed process on your new cluster:
 - A [Self-managed AWS Cluster](#) (see page 1098)
 - A [Self-managed Azure Cluster](#) (see page 1177)
- Pre-provisioned on Premises only:
 - Pre-provisioned hosts with SSH access enabled.
 - An unencrypted SSH private key, whose public key is configured on the above hosts.
 - [Pre-provisioned Override Files](#) (see page 1524) if needed.
- vSphere only:
 - A valid VMware vSphere account with credentials configured.

137 <https://docs.docker.com/get-docker/>

138 <https://podman.io/getting-started/installation>

139 <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>

140 <https://kubernetes.io/docs/tasks/tools/#kubectl>

141 <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

142 <https://github.com/kubernetes-sigs/cluster-api-provider-azure/blob/master/docs/book/src/topics/getting-started.md#prerequisites>

143 <https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html>

144 <https://cloud.google.com/sdk/docs/install>

145 <https://learn.microsoft.com/en-us/cli/azure/install-azure-cli>

- DKP requires permissions for the cloud provider that is being used. See the [Supported Infrastructure Operating Systems](#) (see page 69) page for further setup requirements.

3.6.1.2 Air-gapped Only (additional requirements)

In an air-gapped environment, your environment is **isolated** from unsecured networks, like the Internet, and therefore require additional considerations for installation. Below are the **additional** prerequisites required if installing in an air-gapped environment:

- Linux machine (bastion) that has access to the existing VPC. (Instead of An x86_64-based Linux or macOS machine).
- Ability to download artifacts from the internet and then copy those onto your bastion machine.
- [Download](#) (see page 78) the Complete DKP Air-gapped Bundle for this release - `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`.
- An existing [local registry](#) (see page 1533) to seed the air-gapped environment.

3.6.2 Prerequisites for Kommander Component

To install the Kommander component of DKP, ensure you meet the following prerequisites:

3.6.2.1 Non-air-gapped (all environments)

In a non-air-gapped environment, your environment has two-way access to and from the **Internet**. Below are the prerequisites required if installing in a non-air-gapped environment:

- The version of the CLI that matches the DKP version you want to install.
- Review the [Management Cluster Application Requirements](#) (see page 103) and [Workspace Platform Application Defaults and Resource Requirements](#) (see page 105) to ensure that your cluster will have sufficient resources.
- Ensure you have a [default StorageClass](#) (see page 1395) configured (the Konvoy component is responsible for configuring one)
- A load balancer to route external traffic. In cloud environments, this is provided by your cloud provider. In on-premises and vSphere deployments, you can configure MetalLB. It is possible to use [Virtual IP](#) (see page 1236) also. See [Load Balancing](#) (see page 964) topic for more details.

- NOTE:** If you want to customize your cluster's domain or certificate, ensure you review the respective documentation sections:
 - [Configure custom domains and certificates during the Kommander installation](#) (see page 1428)

- [Configure custom domains and certificates after Kommander has been installed](#) (see page 861)

- For pre-provisioned on-premise environments:
 - Ensure you meet the [Storage requirements](#) (see page 89), [default storage class](#) (see page 91) and [Application Storage](#) (see page 1004).
 - Ensure you have added at least 40 GB of [raw storage to each of your worker nodes](#) (see page 1007) in your clusters.

3.6.2.2 Air-gapped Only (additional requirements)

In an air-gapped environment, your environment is **isolated** from unsecured networks, like the Internet, and therefore require additional considerations for installation. Below are the **additional** prerequisites required if installing in an air-gapped environment:

- A [local registry](#) (see page 1533) containing all the necessary installation images, including the Kommander images which were downloaded in the air-gapped bundle above - `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`. See the pages for Load the Registry under each provider in the [Day 1 - Basic Installs by Infrastructure](#) (see page 147) section for how to push the necessary images to this registry.
- Connectivity with the clusters attaching to the management cluster:
 - Both management and attached clusters must be able to connect to the registry.
 - The management cluster must be able to connect to all attached cluster's API servers.
 - The management cluster must be able to connect to any load balancers created for platform services on the management cluster.



NOTE: If you want to customize your cluster's domain or certificate, ensure you review the respective documentation sections:

- [Configure custom domains and certificates during the Kommander installation](#) (see page 1428)
- [Configure custom domains and certificates after Kommander has been installed](#) (see page 861)

- For pre-provisioned on-premise environments:
 - Ensure you meet the [Storage requirements](#)¹⁴⁶.

¹⁴⁶ [https://d2iq.atlassian.net/wiki/spaces/DENT/pages/213027274/2.5+Rook+Ceph+Configuration#Resource-Requirements-\[inlineExtension\]](https://d2iq.atlassian.net/wiki/spaces/DENT/pages/213027274/2.5+Rook+Ceph+Configuration#Resource-Requirements-[inlineExtension])

- Ensure you have added at least 40 GB of [raw storage to each of your worker nodes](#) (see [page 1007](#)) in your cluster.


3.6.2.3 Next Step:

[Day 1 - Basic Installs by Infrastructure](#) (see [page 147](#))

Proceed to the Day 1 - Basic Installation combinations to install and setup DKP based on your infrastructure environment provider.

3.6.3 Container Runtimes

DKP supports both Podman and Docker as container engines for cluster creation. If you choose to use Podman, it works with Linux on DKP.

 Installation of Podman will be different for different OS environments, so please refer to the [Podman Installation](#)¹⁴⁷ site to install.

Version requirements for a container engine:

- Version 4.0 of [Podman](#)¹⁴⁸ or higher for Linux. Host requirements found here: <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>
- Version 20.10.0 of [Docker](#)¹⁴⁹ or higher for Linux or MacOS. On macOS, Docker runs in a virtual machine. Configure this VM with at least 8 GB of memory.

3.6.3.1 TIPS

- `Set`¹⁵⁰ `alias docker=podman` if desired to ensure any Docker specific commands will run without problem, but our testing shows it is not required.
- You will need to restart your machine after following the [Host Requirements](#)¹⁵¹ to add the `/etc/systemd/system/user@.service.d/delegate.conf` and the `/etc/modules-load.d/iptables.conf` files.

¹⁴⁷ <https://podman.io/getting-started/installation>

¹⁴⁸ <https://podman.io/getting-started/installation>

¹⁴⁹ <https://www.docker.com/>

¹⁵⁰ <https://docs.podman.io/en/latest/>

¹⁵¹ <https://kind.sigs.k8s.io/docs/user/rootless/>

4 Licenses

This section contains overviews of the feature sets for all the available DKP licenses and how to acquire and use your license.

- *DKP Essential*¹⁵² supports your single-cluster environment, or a single Management cluster.
- *DKP Enterprise*¹⁵³ supports a multi-cluster environment, with a Management Cluster and one or more Attached or Managed Clusters.
- *DKP Government Essential*¹⁵⁴ consists of DKP Essential with Confirmed State Side Support (CSS)
- *DKP Government Advanced*¹⁵⁵ consists of DKP Enterprise with Confirmed State Side Support (CSS)

All DKP Enterprise and DKP Government Advanced features are marked with the following badges:

Enterprise

Gov Advanced

The type of license you purchase determines how much of DKP you can use. Some features that are compatible with all versions of DKP can be activated by purchasing additional Add-on licenses. The following table shows the differences between the different DKP licenses:

Feature	DKP Essential	DKP Enterprise	DKP Gov Essential	DKP Gov Advanced
Applications				
Workspace Platform Applications	✓	✓	✓	✓
Project Platform Applications (see page 688)		✓		✓
Catalog Applications (Kafka, Zookeeper, Spark)		✓		✓
Custom Applications		✓		✓

¹⁵² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+DKP+Essential>

¹⁵³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+DKP+Enterprise>

¹⁵⁴ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+DKP+Government+Essential>

¹⁵⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+DKP+Government+Advanced>

¹⁵⁴ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+DKP+Government+Essential>

¹⁵⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+DKP+Government+Advanced>

¹⁵⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+DKP+Government+Advanced>

¹⁵⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+DKP+Government+Advanced>

Feature	DKP Essential	DKP Enterprise	DKP Gov Essential	DKP Gov Advanced
DKP Insights ¹⁵⁶				
Cluster Management				
Single Cluster	✓	✓	✓	✓
Multiple Clusters		✓		✓
Attaching Cluster		✓		✓
Provisioning an additional cluster		✓		✓
Upgrades (see page 1551)	✓	✓	✓	✓
GitOps				
Continuous Deployment (FluxCD)		✓		✓
UX				
DKP CLI (see page 1644)	✓	✓	✓	✓
DKP UI	✓	✓	✓	✓
Workspaces		✓		✓

156 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/2.6+DKP+Insights>

Feature	DKP Essential	DKP Enterprise	DKP Gov Essential	DKP Gov Advanced
Projects		✓		✓
Logging				
Workspace Level Logging	✓	✓	✓	✓
Multi-tenant Logging		✓		✓
Monitoring	✓	✓	✓	✓
Backup & Restore (see page 874)	✓	✓	✓	✓
GPU	✓	✓	✓	✓
Security				
Authentication using Dex	✓	✓	✓	✓
Policy control using Gatekeeper	✓	✓	✓	✓
Cluster Provisioning				
DKP on AWS (see page 293)	✓	✓	✓	✓
DKP on Azure (see page 531)	✓	✓	✓	✓
DKP on GCP (see page 563)	✓	✓	✓	✓

Feature	DKP Essential	DKP Enterprise	DKP Gov Essential	DKP Gov Advanced
DKP on vSphere (see page 433)	✓	✓	✓	✓
Pre-provisioned (see page 149)	✓	✓	✓	✓
On-Prem (see page 149)	✓	✓	✓	✓
EKS Provisioning (see page 404)		✓		✓
AKS Provisioning (see page 547)		✓		✓
FIPS Compliance	✓	✓	✓	✓
Konvoy Image Builder (see page 1461)	✓	✓	✓	✓
Air-Gapped Deployments	✓	✓	✓	✓
Confirmed Stateside Support (CSS)			✓	✓

4.1 Buy a License

DKP Licenses are sold in units of cores. To learn more about both, and to obtain a valid license:

- Contact your sales representative at sales@d2iq.com¹⁵⁷ to buy one or more licenses.
- Purchase a license via the [AWS Marketplace](https://aws.amazon.com/marketplace/)¹⁵⁸. After buying your license, the information (ARN) is accessible in the AWS License Manager console.

After purchasing, [download the binary files](#)¹⁵⁹.

¹⁵⁷ <mailto:sales@d2iq.com>

¹⁵⁸ <https://aws.amazon.com/marketplace/>

¹⁵⁹ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+Download+DKP)

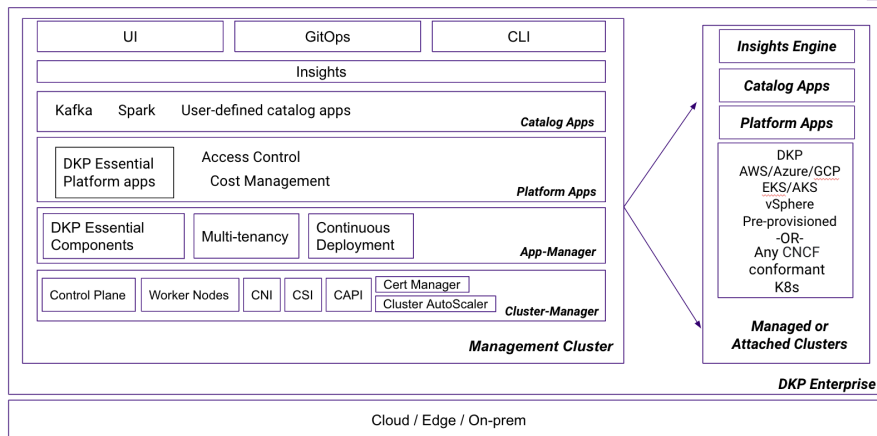
[fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+Download+DKP](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273386200&linkCreation=true&spaceKey=DENT&title=%282.6%29+Download+DKP)

4.2 DKP Enterprise

Enterprise

Gov Advanced

Features and capabilities that are specific to DKP Enterprise are marked with this badge at the top of each page. These items are NOT available to DKP Essential customers.



DKP Enterprise is a multi-cluster lifecycle management Kubernetes solution centered around a management cluster that manages multiple attached or managed Kubernetes clusters via a centralized management dashboard. The management dashboard gives you a single point of observability and control throughout all of your attached or managed clusters. The DKP Enterprise license gives the customer access to the entire Konvoy cluster environment, the DKP UI dashboard that deploys platform and catalog applications, and multi-cluster management, and comprehensive compatibility with our full range of infrastructure deployment options.

4.2.1 Compatible Infrastructure

DKP Enterprise operates across D2iQ's entire range of cloud, on-premises, edge, and air-gapped infrastructures and has support for various OSs, including immutable OSs. See [Supported Operating Systems](#)¹⁶⁰ for a full list of compatible infrastructure.

For the basics on standing up a DKP Enterprise cluster in one of the listed environments of your choice, see [Advanced Configuration](#) (see page 1461).

4.2.2 Platform Applications

Applications can be deployed in any DKP managed cluster, giving you complete flexibility to operate across cloud, on-premises, edge, and air-gapped scenarios. Customers can also use the UI with Kommander to customize which platform applications to deploy to the cluster in a given workspace.

¹⁶⁰ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273748465&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems)

[fromPageId=273748465&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273748465&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems)

4.2.3 Catalog Applications

Quickly and easily deploy applications and complex data services from a centralized service catalog to specific or multiple clusters, with governance. Fast data pipelines can be provisioned automatically from the following catalog of DKP Applications:

- **Kafka:** Primarily used to build real-time streaming data pipelines and applications that adapt to the data streams. It combines messaging, storage, and stream processing to allow storage and analysis of both historical and real-time data.
- **Spark:** An industry standard analytics engine for big data processing and machine learning. Spark enables you to process data for both batch and streaming workloads.
- **Zookeeper:** A centralized service for maintaining configuration information, naming, providing distributed synchronization, and providing group services.

For instructions on how to deploy catalog applications, see [Workspace Catalog Applications \(see page 663\)](#)

4.2.4 Cluster Manager

Konvoy is the Kubernetes installer component of DKP Enterprise that uses industry standard tools to produce a certified Kubernetes cluster. These industry standard tools create a cluster management system that includes:

- **Control Plane:** Manages the worker nodes and pods in the cluster.
- **Worker Nodes:** Used to run containerized applications and handle networking to ensure that traffic between applications across the cluster and from outside of the cluster can be properly facilitated.
- **Container Networking Interface (CNI):** Calico's open source networking and network security solution for containers, virtual machines, and native host-based workloads.
- **Container Storage Interface (CSI):** A common abstraction to container orchestrators for interacting with storage subsystems of various types.
- **Kubernetes Cluster API (CAPI):** Cluster API uses Kubernetes-style APIs and patterns to automate cluster lifecycle management for platform operators. For more on how CAPI is integrated into DKP Enterprise, see [Understanding CAPI Concepts and Terms \(see page 84\)](#)
- **Cert Manager:** A Kubernetes addon to automate the management and issuance of TLS certificates from various issuing sources.
- **Cluster Autoscaler:** A component that automatically adjusts the size of a Kubernetes cluster so that all pods have a place to run and there are no unneeded nodes.

4.2.5 Built-in GitOps

DKP Enterprise comes bundled with GitOps, an operating model for Kubernetes and other cloud native technologies, providing a set of best practices that unify Git deployment, management and monitoring for containerized clusters and applications. GitOps works by using Git as a single source of truth for declarative infrastructure and applications. With GitOps, the use of software agents can alert on any divergence between Git with what is running in a cluster, and if there's a difference, Kubernetes reconcilers automatically update or rollback the cluster depending on the case.

4.2.6 DKP Enterprise Multi-cluster UI

Bundled with DKP Enterprise is a multi-cluster management view in the UI with Kommander that can be used in lieu of the bundled CLI. From the UI you can:

- **Connect to an infrastructure provider:** DKP Enterprise supports on-premises and cloud infrastructure providers such as AWS and Azure for your Konvoy clusters. To automate their provisioning, DKP requires authentication keys to your preferred infrastructure provider entered on the Add Infrastructure Provider form.
- **Setup identity providers:** DKP Enterprise supports GitHub, LDAP, SAML, and standard OIDC identity providers such as Google. These identity management providers support the login and authentication process for your Kubernetes cluster. See [Identity Providers \(see page 597\)](#) for more information.
- **Configure access control:** Role-based authorization control (RBAC) is central to DKP Enterprise and controls access to resources on all connected clusters. The resources are similar to Kubernetes RBAC. You add an identity provider group, and in that group add cluster roles and cluster role bindings for those roles.
- **Deploy applications:** The DKP Enterprise UI allows you to customize your workspace application deployments via the Applications page within the UI.
- **Create a project:** Create projects within a workspace and deploy project-scoped applications. Projects enable teams to deploy configurations and services to their clusters consistently. After configuring roles, ConfigMaps, secrets, and applications for a project, DKP distributes this configuration to each project namespace. For more information concerning projects, see [Projects \(see page 686\)](#).
- **Add a license:** To add a license via the UI with Kommander, see [Add a License \(see page 139\)](#)
- **Kubernetes cost monitoring and management:** The kubecost platform application provides real-time cost visibility and insights for external cloud services such as AWS, helping you continuously reduce your cloud costs.

For more information concerning the global and workspace-level UI with Kommander, see [Workspaces \(see page 650\)](#)

4.3 DKP Essential

DKP Essential is a self-managed single cluster Kubernetes solution that gives you a feature-rich, easy-to-deploy, and easy-to-manage entry-level cloud container platform. The DKP Essential license gives the user access to the entire Konvoy cluster environment, and to the Kommander platform application manager.



Features and capabilities that are specific to DKP Enterprise are marked with this badge at the top of each page. These items are NOT available to DKP Essential users.

4.3.1 Compatible Infrastructure

DKP Essential operates across a range of cloud, on-premise, edge, and air-gapped infrastructures and has support for various OSs, including immutable OSs. See [Supported Operating Systems](#) (see page 69) for a full list of compatible infrastructure.

For the basics on standing up a DKP Essential cluster in one of the listed environments of your choice, see [Advanced Configuration](#) (see page 1461).

4.3.2 Platform Applications

When creating a cluster, the application manager deploys certain platform applications on the newly created cluster. DKP Essential users can use the Kommander UI to customize which platform applications to deploy to the cluster in a given workspace. For a list of available platform applications that are included with DKP, see [Workspace Platform Applications](#) (see page 105).



NOTE: The platform application `kubecost` is not included with DKP Essential, but is included with [DKP Enterprise](#) (see page 129).

4.3.3 Cluster Manager

Konvoy is the Kubernetes installer component of DKP Essential that uses industry standard tools to produce a certified Kubernetes cluster. These industry standard tools create a cluster management system that includes:

- **Control Plane:** Manages the worker nodes and pods in the cluster.
- **Worker Nodes:** Used to run containerized applications and handle networking to ensure that traffic between applications across the cluster and from outside of the cluster can be properly facilitated.
- **Container Networking Interface (CNI):** Calico's open source networking and network security solution for containers, virtual machines, and native host-based workloads.
- **Container Storage Interface (CSI):** A common abstraction to container orchestrators for interacting with storage subsystems of various types.
- **Kubernetes Cluster API (CAPI):** Cluster API uses Kubernetes-style APIs and patterns to automate cluster lifecycle management for platform operators. For more on how CAPI is integrated into DKP Essential, see [Understanding CAPI Concepts and Terms](#) (see page 84)
- **Cert Manager:** A Kubernetes add-on to automate the management and issuance of TLS certificates from various issuing sources.
- **Cluster Autoscaler:** A component that automatically adjusts the size of a Kubernetes cluster so that all pods have a place to run and there are no unneeded nodes.

4.3.4 Built-in GitOps

DKP Essential comes bundled with GitOps, an operating model for Kubernetes and other cloud native technologies, providing a set of best practices that unify Git deployment, management and monitoring for containerized clusters and applications. GitOps works by using Git as a single source of truth for declarative infrastructure and applications. With GitOps, the use of software agents can alert on any divergence between Git with what is running in a cluster, and if there's a difference, Kubernetes reconcilers automatically update or rollback the cluster depending on the case.

4.3.5 DKP Essential Single Cluster UI

Bundled with DKP Essential is a single cluster management view of [DKP UI \(see page 1420\)](#) that can be used in lieu of the bundled CLI. From the UI you can:

- **Setup identity providers:** DKP Essential supports GitHub, LDAP, SAML, and standard OIDC identity providers such as Google. These identity management providers support the login and authentication process for your Kubernetes cluster. See [Identity Providers \(see page 597\)](#) for more information.
- **Deploy applications:** The DKP Essential UI with Kommander allows you to customize your workspace application deployments via the Applications page within the UI.
- **Add a license:** To add a license via the UI with Kommander, see [Add a license to Kommander \(see page 139\)](#)

For more information concerning the global and workspace-level UI, see [Workspaces \(see page 650\)](#)

4.4 DKP Government Advanced

Features and capabilities that are specific to DKP Government Advanced are marked with this badge at the top of each page. These items are NOT available to DKP Government Essential users. In general, every feature available in [DKP Enterprise \(see page 129\)](#) is also available in DKP Government Advanced.

Enterprise

Gov Advanced

DKP Government Advanced (also referred to as DKP Gov Advanced) is a multi-cluster Kubernetes management platform that is optimized for deployment in the government sector. With top-to-bottom declarative programming, DKP Government Advanced provides the most advanced and feature-rich Kubernetes management capability across any infrastructure, whether on-premise, in the cloud, in air-gapped environments, or at the edge.

Ease of use is the hallmark of D2iQ Kubernetes solutions. DKP Government Advanced enables you to be up and running in minutes and hours rather than weeks and months, with complete stability, reliability, military-grade security, and rapid time-to-value. Complexity is reduced through packaging, automation, integration, and elegant design.

Built on pure open-source Kubernetes, DKP Government Advanced enables you to avoid vendor lock-in, easily upgrade services, and take full advantage of open-source community innovation. DKP Government Advanced provides platform services for networking, storage, logging, monitoring, and more, all of which have been carefully selected from the Cloud Native Computing Foundation (CNCF) landscape and rigorously tested to work together.

DKP Government Advanced is a multi-cluster lifecycle management Kubernetes solution centered around a management cluster that manages multiple attached or managed Kubernetes clusters via a centralized management dashboard. The management dashboard gives you a single point of observability and control throughout all of your attached or managed clusters. The DKP Government Advanced license gives the user access to the entire Konvoy cluster environment, the DKP UI dashboard that deploys platform and catalog applications, and multi-cluster management, and comprehensive compatibility with our full range of infrastructure deployment options.

4.4.1 Compatible Infrastructure

DKP Government Advanced operates across D2iQ's entire range of cloud, on-premises, edge, and air-gapped infrastructures and has support for various OSs, including immutable OSs. See [Supported Operating Systems](#) (see page 69) for a full list of compatible infrastructure.

For the basics on standing up a DKP Government Advanced cluster in one of the listed environments of your choice, see [Additional Infrastructure Customized Configurations](#) (see page 1036).

4.4.2 Platform Applications

Applications can be deployed in any DKP managed cluster, giving you complete flexibility to operate across cloud, on-premises, edge, and air-gapped scenarios. Customers can also use the UI with Kommander to customize which platform applications to deploy to the cluster in a given workspace.

4.4.3 Catalog Applications

Quickly and easily deploy applications and complex data services from a centralized service catalog to specific or multiple clusters, with governance. Fast data pipelines can be provisioned automatically from the following catalog of DKP Applications:

- **Kafka:** Primarily used to build real-time streaming data pipelines and applications that adapt to the data streams. It combines messaging, storage, and stream processing to allow storage and analysis of both historical and real-time data.
- **Spark:** An industry standard analytics engine for big data processing and machine learning. Spark enables you to process data for both batch and streaming workloads.
- **Zookeeper:** A centralized service for maintaining configuration information, naming, providing distributed synchronization, and providing group services.

For instructions on how to deploy catalog applications, see [Workspace Catalog Applications](#) (see page 663)

4.4.4 Cluster Manager

Konvoy is the Kubernetes installer component of DKP Government Advanced that uses industry standard tools to produce a certified Kubernetes cluster. These industry standard tools create a cluster management system that includes:

- **Control Plane:** Manages the worker nodes and pods in the cluster.
- **Worker Nodes:** Used to run containerized applications and handle networking to ensure that traffic between applications across the cluster and from outside of the cluster can be properly facilitated.

- **Container Networking Interface (CNI):** Calico’s open source networking and network security solution for containers, virtual machines, and native host-based workloads.
- **Container Storage Interface (CSI):** A common abstraction to container orchestrators for interacting with storage subsystems of various types.
- **Kubernetes Cluster API (CAPI):** Cluster API uses Kubernetes-style APIs and patterns to automate cluster lifecycle management for platform operators. For more on how CAPI is integrated into DKP Government Advanced, see [Understanding CAPI Concepts and Terms \(see page 84\)](#)
- **Cert Manager:** A Kubernetes add-on to automate the management and issuance of TLS certificates from various issuing sources.
- **Cluster Autoscaler:** A component that automatically adjusts the size of a Kubernetes cluster so that all pods have a place to run and there are no unneeded nodes.

4.4.5 Built-in GitOps

DKP Government Advanced comes bundled with GitOps, an operating model for Kubernetes and other cloud native technologies, providing a set of best practices that unify Git deployment, management and monitoring for containerized clusters and applications. GitOps works by using Git as a single source of truth for declarative infrastructure and applications. With GitOps, the use of software agents can alert on any divergence between Git with what is running in a cluster, and if there’s a difference, Kubernetes reconcilers automatically update or rollback the cluster depending on the case.

4.4.6 DKP Government Advanced Multi-cluster UI

Bundled with DKP Government Advanced is a multi-cluster management view in the UI with Kommander that can be used in lieu of the bundled CLI. From the UI you can:

- **Connect to an infrastructure provider:** DKP Government Advanced supports on-premises and cloud infrastructure providers such as AWS and Azure for your Konvoy clusters. To automate their provisioning, DKP requires authentication keys to your preferred infrastructure provider entered on the Add Infrastructure Provider form.
- **Setup identity providers:** DKP Government Advanced supports GitHub, LDAP, SAML, and standard OIDC identity providers such as Google. These identity management providers support the login and authentication process for your Kubernetes cluster. See [Identity Providers \(see page 597\)](#) for more information.
- **Configure access control:** Role-based authorization control (RBAC) is central to DKP Government Advanced and controls access to resources on all connected clusters. The resources are similar to Kubernetes RBAC. You add an identity provider group, and in that group add cluster roles and cluster role bindings for those roles.
- **Deploy applications:** The DKP Government Advanced UI allows you to customize your workspace application deployments via the Applications page within the UI.
- **Create a project:** Create projects within a workspace and deploy project-scoped applications. Projects enable teams to deploy configurations and services to their clusters consistently. After configuring roles, ConfigMaps, secrets, and applications for a project, DKP distributes this configuration to each project namespace. For more information concerning projects, see [Projects \(see page 686\)](#).
- **Add a license:** To add a license via the UI with Kommander, see [Add a License \(see page 139\)](#)

- **Kubernetes cost monitoring and management:** The kubecost platform application provides real-time cost visibility and insights for external cloud services such as AWS, helping you continuously reduce your cloud costs.

For more information concerning the global and workspace-level UI with Kommander, see [Workspaces](#) (see page 650)

4.4.7 Confirmed Stateside Support (CSS)

Because many governmental applications require CSS, DKP Government Advanced includes access to D2iQ's CSS team of U.S. citizens working at U.S. support centers.

4.4.8 Military-Grade Security

DKP Government Advanced can be configured to meet defined security standards. Each component is tested and certified before release. DKP also provides instant platform engineering that reduces the burden of security on DevOps and DevSecOps teams. Being based on pure upstream open-source Kubernetes enables easy and trouble-free upgrades, patches, and bug fixes.

4.5 DKP Government Essential

DKP Government (Gov) Essential is a single-cluster Kubernetes management platform that is optimized for deployment in the government sector. Based on the D2iQ Kubernetes Platform (DKP), DKP Government Essential can easily and seamlessly expand to a multi-cluster Kubernetes management platform as your needs grow.

With top-to-bottom declarative programming, DKP Government Essential provides the most advanced and feature-rich single-cluster Kubernetes management capability across any infrastructure, whether on-premise, in the cloud, in air-gapped environments, or at the edge.

Ease of use is the hallmark of D2iQ Kubernetes solutions. DKP Government Essential enables you to be up and running in minutes and hours rather than weeks and months, with complete stability, reliability, military-grade security, and rapid time-to-value. Complexity is reduced through packaging, automation, integration, and elegant design.

Built on pure open-source Kubernetes, DKP Government Essential enables you to avoid vendor lock-in, easily upgrade services, and take full advantage of open-source community innovation. DKP Government Essential provides platform services for networking, storage, logging, monitoring, and more, all of which have been carefully selected from the Cloud Native Computing Foundation (CNCF) landscape and rigorously tested to work together.

DKP Government Essential is a self-managed single cluster Kubernetes solution that gives you a feature-rich, easy-to-deploy, and easy-to-manage entry-level cloud container platform. The DKP Government Essential license gives the user access to the entire Konvoy cluster environment, and to the Kommander platform application manager. In general, every feature available in [DKP Essential](#) (see page 132) is also available in DKP Government Essential

4.5.1 Compatible Infrastructure

DKP Government Essential operates across a range of cloud, on-premise, edge, and air-gapped infrastructures and has support for various OSs, including immutable OSs. See [Supported Operating Systems](#) (see page 69) for a full list of compatible infrastructure.

For the basics on standing up a DKP Government Essential cluster in one of the listed environments of your choice, see [Additional Infrastructure Configurations](#) (see page 1036).

4.5.2 Platform Applications

When creating a cluster, the application manager deploys certain platform applications on the newly created cluster. DKP Government Essential users can use the Kommander UI to customize which platform applications to deploy to the cluster in a given workspace. For a list of available platform applications that are included with DKP, see [Workspace Platform Applications](#) (see page 105).



NOTE: The platform application `kubecost` is not included with DKP Government Essential, but is included with [DKP Enterprise](#) (see page 132) and [DKP Government Advanced](#) (see page 134).

4.5.3 Cluster Manager

Konvoy is the Kubernetes installer component of DKP Government Essential that uses industry standard tools to produce a certified Kubernetes cluster. These industry standard tools create a cluster management system that includes:

- **Control Plane:** Manages the worker nodes and pods in the cluster.
- **Worker Nodes:** Used to run containerized applications and handle networking to ensure that traffic between applications across the cluster and from outside of the cluster can be properly facilitated.
- **Container Networking Interface (CNI):** Calico's open source networking and network security solution for containers, virtual machines, and native host-based workloads.
- **Container Storage Interface (CSI):** A common abstraction to container orchestrators for interacting with storage subsystems of various types.
- **Kubernetes Cluster API (CAPI):** Cluster API uses Kubernetes-style APIs and patterns to automate cluster lifecycle management for platform operators. For more on how CAPI is integrated into DKP Government Essential, see [Understanding CAPI Concepts and Terms](#) (see page 84)
- **Cert Manager:** A Kubernetes add-on to automate the management and issuance of TLS certificates from various issuing sources.
- **Cluster Autoscaler:** A component that automatically adjusts the size of a Kubernetes cluster so that all pods have a place to run and there are no unneeded nodes.

4.5.4 Built-in GitOps

DKP Government Essential comes bundled with GitOps, an operating model for Kubernetes and other cloud native technologies, providing a set of best practices that unify Git deployment, management and monitoring for containerized clusters and applications. GitOps works by using Git as a single source of truth for declarative infrastructure and applications. With GitOps, the use of software agents can alert on any divergence between Git with what is running in a cluster, and if there's a difference, Kubernetes reconcilers automatically update or rollback the cluster depending on the case.

4.5.5 DKP Government Essential Single Cluster UI

Bundled with DKP Government Essential is a single cluster management view of [DKP UI \(see page 1420\)](#) that can be used in lieu of the bundled CLI. From the UI you can:

- **Setup identity providers:** DKP Government Essential supports GitHub, LDAP, SAML, and standard OIDC identity providers such as Google. These identity management providers support the login and authentication process for your Kubernetes cluster. See [Identity Providers \(see page 597\)](#) for more information.
- **Deploy applications:** The DKP Government Essential UI with Kommander allows you to customize your workspace application deployments via the Applications page within the UI.
- **Add a license:** To add a license via the UI with Kommander, see [Add a license to Kommander \(see page 139\)](#)

For more information concerning the global and workspace-level UI, see [Workspaces \(see page 650\)](#)

4.5.6 Confirmed Stateside Support (CSS)

Because many governmental applications require CSS, DKP Government Essential includes access to D2iQ's CSS team of U.S. citizens working at U.S. support centers.

4.5.7 Military-Grade Security

DKP Government Essential can be configured to meet defined security standards. Each component is tested and certified before release. DKP also provides instant platform engineering that reduces the burden of security on DevOps and DevSecOps teams. Being based on pure upstream open-source Kubernetes enables easy and trouble-free upgrades, patches, and bug fixes.

4.6 Add a DKP license

4.6.1 Prerequisites

For licenses that you purchase from the [AWS Marketplace](#)¹⁶¹, an AWS administrator must attach the [AWS managed policy](#)¹⁶² `AWSLicenseManagerConsumptionPolicy` to the `control-plane.cluster-api-provider-aws.sigs.k8s.io` role created when [configuring AWS IAM policies for Konvoy](#) (see page 1063). If an Administrator does not attach this policy attached to the role, DKP cannot verify the license information provided in the procedure steps that follow.

The machine onto which you install DKP must meet these requirements:

- Docker installed - DKP uses Kubernetes-in-Docker (Kind) to create a bootstrap cluster for creating Management clusters, and thus Docker installation is required. Docker is not used to run your Management or workload clusters.
- Active Internet connectivity
- Existing AWS account
- AWS Administrator logged into the AWS account used to purchase DKP

4.6.2 Download the Container Image and Extract the Binaries

1. When you complete the sales transaction in the AWS Marketplace, select a DKP version from the **Choose a fulfillment option** dropdown. We recommend that you select the latest version.
2. Select **Continue to Launch** to obtain the download instructions, available when you select **Usage Instructions**.
3. Complete the steps in each of the procedures under **Usage instructions**. The links in these steps take you to DKP documentation pages that provide the steps you need to set up AWS to work with DKP, and the download and installation of DKP itself.
4. Expand the **Create a license token and IAM role** link in the **Container images** section of the window. This gives you access to the Create token button to generate the license token and IAM role and exposes an inset code window with the token and role credentials.
5. Use the generated command to log in to AWS using the token and role you just generated. Note that you can omit this step if you have already configured the AWS CLI with valid credentials.
Tip: Using the Copy button at the upper right of this inset code window helps ensure you copy the entire command.
6. Copy the commands at the bottom of the page, and run them sequentially to download the container images. It will look similar to this:

¹⁶¹ <https://aws.amazon.com/marketplace/>

¹⁶² <https://docs.aws.amazon.com/license-manager/latest/userguide/security-iam-awsmanpol.html#security-iam-AWSLicenseManagerConsumptionPolicy>

```
aws ecr get-login-password \
  --region <insert-region> | docker login \
  --username <username> \
  --password-stdin <ecr-address>
```

```
CONTAINER_IMAGES="<ecr-address>/mesosphere/d2iq-dkp-essential-premium-support:v2.5.0"
```

```
for i in $(echo $CONTAINER_IMAGES | sed "s/,/ /g"); do docker pull $i; done
```

7. Access a terminal window on the Linux machine where you plan to run the DKP CLI cluster, and execute the command above, using your environment's specific values. This downloads a container that contains the DKP binary.
8. After downloading the container, run the following command to copy the binaries onto your Linux machine.

```
docker run -it --rm -u $(id -u):$(id -g) -v $(pwd):/dkp $CONTAINER_IMAGES
```

You will then see the following output:

```
dkp binary is placed in the local directory, to run:
./dkp --help
```

You will now see the `dkp` binary in your working directory. Follow the [DKP installation \(see page 108\)](#) instructions using these binaries, and then [add your license \(see page 139\)](#) to DKP.

4.6.3 Obtain the Amazon Resource Number (ARN) from the AWS Marketplace UI

At this point, you need to use the AWS Marketplace user interface to navigate to the license table and configure its settings to display the ARN for your purchased license.

1. From the **Launch this software** page, navigate back to the **Product Detail** page.
2. Select the **View Subscription** button at the upper right (in the blue information box labeled **"You have access to this product"**).
3. Select **Manage subscriptions** from the left-hand navigation pane, and then click the **Manage** button on the specific license. The details page for that license displays.
4. Select the **License** number hyperlink to display the license page, then select **Granted licenses** from the left-hand navigation pane and find your license in the list.
5. Use the **Settings** icon in the upper right corner to display the columns displayed, and then enable the **License arn** value.

6. Select the **Confirm** button to return to the license list with the License ARN displayed in the last column. You can now copy and paste this ARN as needed.

If you have difficulty obtaining your ARN, contact D2iQ Support for assistance.

4.6.4 Enter License Information in the DKP UI

For licenses purchased directly from D2iQ, you can obtain the license token through the [Support Portal](#)¹⁶³. Insert this token in the last step below.

For licenses purchased through the AWS marketplace, you can find the license in the “Granted Licenses” table of AWS License Manager inside the AWS console. If necessary, modify the table view preferences to show the “License ARN.” Copy this value for use in the last step below.



You must be an administrator to add licenses to DKP.

From the the UI with Kommander, complete the following steps:

1. Select **Global** in the workspace header drop-down.
2. In the sidebar menu, select **Administration > Licensing**.
3. Select + Activate License to enter the Activate License form.
4. On the Activate License form page, select **D2iQ** or **AWS Marketplace** depending on where you acquired your license.
5. Paste your license token, or the AWS ARN, in the text box and select **Save**.

If there is an error submitting a license acquired directly from D2iQ, you can activate the license directly through `kubectl`.

4.6.5 Enter a DKP License using kubectl

You can activate a license acquired from D2iQ directly using the `kubectl` utility.

1. Create a secret, replacing the value `MY_LICENSE` in the command that follows with your actual, D2iQ-provided DKP license token:

```
kubectl create secret generic my-license-secret --from-literal=jwt=MY_LICENSE
-n kommander
kubectl label secret my-license-secret kommanderType=license -n kommander
```

2. Create a license object with this command:

¹⁶³ <https://support.d2iq.com/hc/en-us>

```
cat <<EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: License
metadata:
  name: my-license
  namespace: kommander
spec:
  licenseRef:
    name: my-license-secret
EOF
```

3. Return to the license page in the DKP UI to see your valid license display.

4.6.6 Activate a DKP Insights License Key

Instructions for activating an Insights License Key in the DKP UI

4.6.6.1 Prerequisites


Prior to activating an Insights license key, ensure that you have a valid DKP Essential or DKP Enterprise license.

4.6.6.2 Obtain a License Key

DKP Insights license keys to activate the product are provided directly by D2iQ. Contact your account rep to receive a license key.


After you receive a license key, follow the instructions in the next section to activate DKP Insights in your environment.

4.6.6.3 Enter License Information

 An administrator must activate licenses in DKP.

From the DKP UI, complete the following steps:

1. Select **Global** in the workspace header drop-down.
2. In the sidebar menu, select **Administration > Licensing**.
3. Select **+ Activate Add-on License**.
4. In the Add-on License form page, input the Insight license key that was provided to you by DKP.
5. Select **Save**.

 Once you activate an Insights license key, the prompt to activate an Add-on License is removed.

4.6.6.4 Remove an Insights License


If your Insights license information has changed, you may need to remove an existing license from DKP to activate a new one. Only DKP administrators can remove licenses.

In the DKP UI, perform the following:

1. Select **Global** in the workspace header drop-down.
2. In the sidebar menu, select **Administration > Licensing**.
3. Select **Remove License** in the DKP Insights card, then follow the prompts.

4.7 Remove a DKP license

If your license information has changed, you may need to remove an existing license from DKP to add a new one. Only DKP administrators have the ability to remove licenses.

 Original license information can still be obtained from D2iQ or the AWS License Manager console even after removing from DKP.

4.7.1 Remove a License via the UI

In the DKP UI, do the following:

1. Select **Global** in the workspace header drop-down.
2. In the sidebar menu, select **Administration > Licensing**.
3. Your existing licenses will be listed. Click **Remove License** on the license you would like to remove, and follow the prompts.

4.7.2 Manually Remove a License using kubectl

To remove a license from DKP using `kubectl`, you have to delete the `Secret` and `License` objects. In this example, the secret is named “my-license-secret”.

1. Validate that the secret exists in the `kommander` namespace:

```
kubectl describe secret -n kommander my-license-secret
```

Expected output:

```
Name:          my-license-secret
Namespace:     kommander
Labels:        kommanderType=license
Annotations:   <none>

Type:          Opaque

Data
====
jwt:  455 bytes
```

2. Delete the secret from the `kommander` namespace:

```
kubectl delete secret -n kommander my-license-secret
```

Expected output:

```
secret "my-license-secret" deleted
```

3. We do the same with the `License` object. Validate that it exists in the `kommander` namespace:

```
kubectl describe license -n kommander my-license
```

Expected output:

```
Name:          my-license
Namespace:     kommander
Labels:        <none>
Annotations:   kubectl.kubernetes.io/last-applied-configuration:
                {"apiVersion":"kommander.mesosphere.io/v1beta1","kind":"License
                ","metadata":{"annotations":{},"name":"my-license", "namespace":"kommand...
API Version:   kommander.mesosphere.io/v1beta1
Kind:          License
Metadata:
  Creation Timestamp:  2020-03-25T14:57:31Z
  Generate Name:      license-
  Generation:         1
  Resource Version:    17895
  Self Link:          /apis/kommander.mesosphere.io/v1beta1/namespaces/
  kommander/licenses/my-license
```



```

UID:                35ee9254-4094-40eb-a2d8-4687c5d212d9
Spec:
  License Ref:
    Name: my-license-secret
Status:
  Cluster Capacity: 500
  Customer Id:      mesosphere-developer
  End Date:         2020-10-02T14:00:09Z
  License Id:       mesosphere-developer
  Start Date:       2019-10-02T14:00:09Z
  Valid:            true
  Version:          1.0
Events:
  Type      Reason          Age          From          Message
  ----      -
  Normal    LicenseUpdateSuccess  7m7s (x2 over 7m7s)  LicenseSignature  License
updated successfully

```

4. Delete the license from the `kommander` namespace:

```
kubectl delete license -n kommander my-license
```

Expected output:

```
license.kommander.mesosphere.io "my-license" deleted
```

You have now successfully removed a license.

4.8 MSP Program

The D2iQ Managed Service Provider (MSP) Program is a partnership program targeted at companies that provide cloud infrastructure services to their customers. When becoming a member, you can benefit from flexible billing options and enhanced support. For more information, read the full [press release](#)¹⁶⁴.

4.8.1 Benefits

- Transparent and flexible billing options
- Enhanced support experience
- Support inquiries are routed to CKA-certified, senior support engineers

4.8.2 How does it work?

After you sign your MSP agreement, you will be part of the program.

¹⁶⁴ <https://d2iq.com/press-release/managed-service-providers-kubernetes-management-program>

To benefit from the flexible billing options, submit a monthly report in the form of a [telemetry bundle](#) (see [page 1030](#)) for any cluster that ran during the month.

4.8.3 How do I become an MSP partner?

Contact your D2iQ representative or use our [contact form](#)¹⁶⁵ to obtain more information on becoming a partner.

¹⁶⁵ <https://d2iq.com/contact>

5 Day 1 - Basic Installs by Infrastructure

This section provides basic installation instructions for your infrastructure using combinations of providers and other variables. A **basic** cluster contains nodes and a running instance of DKP, but is not yet a *production* cluster. While you will not yet be ready to deploy a workload after these procedures, you will be able to get familiar with DKP and see the cluster structure. You can think of the install procedures in this section as a sort of quick start guide.



For [Custom Installation and Additional Infrastructure Tools](#) (see page 1036) refer to that section of the documentation.

Depending on your selected infrastructure, and the requirements of your specific environment, there are [additional configuration tasks in the Day 2 section](#) (see page 582) to complete. Production cluster configuration lets you deploy and enable the Day 2 applications, and your workload applications, that you need for production operations.

For virtualized environments, DKP can provision the virtual machines necessary to run Kubernetes clusters. If you allow DKP to manage your infrastructure, look for your supported infrastructure provider installation choices below.



In bare metal environments or any time you would like to provision your nodes manually, see the corresponding [Pre-provisioned](#) (see page 1200) section.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.1 Scenario Based Installation Instructions

5.2 Pre-provisioned Install Options

- [Pre-provisioned Non-air-gapped Install](#) (see page 151)
- [Pre-provisioned Air-gapped Install](#) (see page 170)
- [Pre-provisioned FIPS Install](#) (see page 195)
- [Pre-provisioned FIPS Air-gapped Install](#) (see page 215)
- [Pre-provisioned with GPU Install](#) (see page 239)

- [Pre-provisioned Air-gapped with GPU Install \(see page 262\)](#)

5.3 AWS Install Options

- [AWS Install \(see page 294\)](#)
- [AWS Air-gapped Install \(see page 310\)](#)
- [AWS with FIPS Install \(see page 328\)](#)
- [AWS Air-gapped FIPS Install \(see page 344\)](#)
- [AWS with GPU Install \(see page 363\)](#)
- [AWS Air-gapped with GPU Install \(see page 382\)](#)

5.4 EKS Install Options

- [EKS Install \(see page 405\)](#)

5.5 vSphere Install Options

- [vSphere Prerequisites: All Installation Types \(see page 434\)](#)
- [vSphere Install \(see page 438\)](#)
- [vSphere Air-gapped Install \(see page 459\)](#)
- [vSphere FIPS Install \(see page 483\)](#)
- [vSphere FIPS Air-gapped Install \(see page 505\)](#)

5.6 VMware Cloud Director Install Options

- [Cloud Director Service Provider \(see page 530\)](#)

5.7 Azure Install Options

- [Azure Install \(see page 532\)](#)

5.8 AKS Install Options

- [AKS Install \(see page 547\)](#)

5.9 GCP Install Options

- [GCP Install \(see page 563\)](#)

5.10 Pre-provisioned Install Options

For an either non-air-gapped or air-gapped environment that is [Pre-provisioned](#) (see page 86), install options are provided for you in this one location.

Remember, there are always more options for custom YAML in the [Pre-provisioned Infrastructure](#) (see page 1200) section, but this will get you operative with the most common scenarios.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

Below are all the DKP supported environment combinations: (Refer to this page to check if your OS is a [Supported Operating System](#) (see page 69).)

- [Pre-provisioned Non-air-gapped Install](#) (see page 151)
- [Pre-provisioned Air-gapped Install](#) (see page 170)
- [Pre-provisioned FIPS Install](#) (see page 195)
- [Pre-provisioned FIPS Air-gapped Install](#) (see page 215)
- [Pre-provisioned with GPU Install](#) (see page 239)
- [Pre-provisioned Air-gapped with GPU Install](#) (see page 262)

Important Pre-provisioned Topics:

- Pre-provisioned includes on-premises, vSphere, AWS, Azure, and GCP infrastructures and is described in more detail under [What is Pre-provisioned Infrastructure](#) (see page 86).
- For more information regarding CSI Disk storage and changing default StorageClass, see [Default Storage Providers in DKP](#) (see page 91).
- For Azure environment using Pre-provisioned specifics, see [Pre-provisioned Azure only Configurations](#) (see page 1227)

5.10.1 Resources



Additional resource information specific to On Premises and Pre-provisioned is listed below:

- **Control plane machines** -
 - 15% free space on the root file system.
 - Multiple ports open, as described in [DKP Ports](#) (see page 67).
 - `firewalld` systemd service disabled. If it exists and is enabled, use the commands `systemctl stop firewalld` then `systemctl disable firewalld`, so that `firewalld` remains disabled after the machine restarts.

Swap is disabled. The `kubelet` does not have generally-available support for swap. Due to variable commands, refer to your operating system documentation.

- **Worker machines -**

- 15% free space on the root file system
- If you plan to use local volume provisioning to provide persistent volumes for your workloads, you must mount at least four volumes to the `/mnt/disks/` mount point on each machine. Each volume must have at least 100 GiB of capacity.
- Ensure your disk meets the resource requirements for Rook Ceph in `Block` mode for `ObjectStorageDaemons`¹⁶⁶ as specified in the [requirements table](#) (see page 1007).
- Multiple ports open, as described in [DKP Ports](#) (see page 67).
- `firewalld` systemd service disabled. If it exists and is enabled, use the commands `systemctl stop firewalld` then `systemctl disable firewalld`, so that `firewalld` remains disabled after the machine restarts.

Swap is disabled. The `kubelet` does not have generally-available support for swap. Due to variable commands, refer to your operating system documentation.

5.10.2 Pre-provisioned Non-air-gapped Install

This section provides instructions to install DKP in a Pre-provisioned non-air-gapped environment.

If not already done, refer to the [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.10.2.1 Next Step:

[Pre-provisioned: Define Infrastructure](#) (see page 151)

5.10.2.2 Pre-provisioned: Define Infrastructure

Define the cluster hosts and infrastructure

¹⁶⁶ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

Konvoy needs to know how to access your cluster hosts. This is done using inventory resources. For initial cluster creation, you must define a control-plane and at least one worker pool.

5.10.2.2.1 Define your Infrastructure

1. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```
export CONTROL_PLANE_1_ADDRESS="<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<control-plane-address-3>"
export WORKER_1_ADDRESS="<worker-address-1>"
export WORKER_2_ADDRESS="<worker-address-2>"
export WORKER_3_ADDRESS="<worker-address-3>"
export WORKER_4_ADDRESS="<worker-address-4>"
export SSH_USER="<ssh-user>"
export SSH_PRIVATE_KEY_SECRET_NAME="CLUSTER_NAME-ssh-key"
```

2. Use the following template to help you define your infrastructure. The environment variables that you set in the previous step automatically replace the variable names when the inventory YAML file is created.

```
cat <<EOF > preprovisioned_inventory.yaml
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: CLUSTER_NAME-control-plane
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    # Create as many of these as needed to match your infrastructure
    # Note that the command line parameter --control-plane-replicas determines
    # how many control plane nodes will actually be used.
    #
    - address: CONTROL_PLANE_1_ADDRESS
    - address: CONTROL_PLANE_2_ADDRESS
    - address: CONTROL_PLANE_3_ADDRESS
  sshConfig:
    port: 22
    # This is the username used to connect to your infrastructure. This user
    # must be root or
    # have the ability to use sudo without a password
    user: SSH_USER
    privateKeyRef:
```

```

    # This is the name of the secret you created in the previous step. It
    must exist in the same
    # namespace as this inventory object.
    name: $SSH_PRIVATE_KEY_SECRET_NAME
    namespace: default
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: $CLUSTER_NAME-md-0
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: $CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    - address: $WORKER_1_ADDRESS
    - address: $WORKER_2_ADDRESS
    - address: $WORKER_3_ADDRESS
    - address: $WORKER_4_ADDRESS
  sshConfig:
    port: 22
    user: $SSH_USER
    privateKeyRef:
      name: $SSH_PRIVATE_KEY_SECRET_NAME
      namespace: default
EOF

```

5.10.2.2.2 Next Step:

[Pre-provisioned: Define Control Plane Endpoint \(see page 152\)](#)

5.10.2.3 Pre-provisioned: Define Control Plane Endpoint

5.10.2.3.1 Define the Control Plane Endpoint for your cluster as well as connection mechanism

A control plane should have three, five, or seven nodes, so it can remain available if one or more nodes fail. A control plane with one node should not be used in production.

In addition, the control plane should have an endpoint that remains available if some nodes fail.

```

          ----- cp1.example.com:6443
          |
lb.example.com:6443 ----- cp2.example.com:6443
          |
          ----- cp3.example.com:6443

```


In this example, the control plane endpoint host is `lb.example.com`, and the control plane endpoint port is `6443`. The control plane nodes are `cp1.example.com`, `cp2.example.com`, and `cp3.example.com`. The port of each API server is `6443`.

5.10.2.3.2 Select your Connection Mechanism

A virtual IP is the address that the client uses to connect to the service. A load balancer is the device that distributes the client connections to the backend servers. Before you create a new DKP cluster, choose an external load balancer(LB) or virtual IP.

- **External load balancer**

It is recommended that an external load balancer be the control plane endpoint. To distribute request load among the control plane machines, configure the load balancer to send requests to all the control plane machines. Configure the load balancer to send requests only to control plane machines that are responding to API requests.

- **Built-in virtual IP**

If an external load balancer is not available, use the [built-in virtual IP \(see page 1236\)](#). The virtual IP is *not* a load balancer; it does not distribute request load among the control plane machines. However, if the machine receiving requests does not respond to them, the virtual IP automatically moves to another machine.

5.10.2.3.3 Single-Node Control Plane



Do not use a single-node control plane in a production cluster.

A control plane with one node can use its single node as the endpoint, so you will not require an external load balancer, or a built-in virtual IP. At least one control plane node must always be running. Therefore, to upgrade a cluster with one control plane node, a spare machine must be available in the control plane inventory. This machine is used to provision the new node before the old node is deleted.



Modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane \(see page 1538\)](#).

When the API server endpoints are defined, you can create the cluster using the link in Next Step below.

5.10.2.3.4 Known Limitations



Be aware of these limitations in the current release of DKP.

The control plane endpoint port is also used as the API server port on each control plane machine. The default port is 6443. Before you create the cluster, ensure the port is available for use on each control plane machine.

5.10.2.3.5 Next Step:

[Pre-provisioned: Create a New Cluster](#) (see page 155)

5.10.2.4 Pre-provisioned: Create a Management Cluster

Create a new Pre-provisioned Kubernetes cluster in a non-air-gapped environment with the steps below.

5.10.2.4.1 Name Your Cluster



The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)¹⁶⁷ for more naming information.

When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:


1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=preprovisioned-example
```


¹⁶⁷ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

5.10.2.4.2 Create a Kubernetes Cluster

Once you've defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster. This process will create a self-managed cluster to be used as the Management cluster.

 Before you create a new DKP cluster below, choose an external load balancer (LB) or [virtual IP](#) (see page 1236) and use the corresponding `dkp create cluster` command.

In a pre-provisioned environment, use the Kubernetes CSI and third party drivers for local volumes and other storage devices in your data center.

 DKP uses local static provisioner as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)¹⁶⁸ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the [inventory YAML](#) (see page 151) previously created.

The create cluster command below includes the `--self-managed` flag. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.

1. This command uses the default external load balancer (LB) option (see alternative Step 1 for virtual IP):

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host <control plane endpoint host> \
  --control-plane-endpoint-port <control plane endpoint port, if different than
6443> \
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml \
  --ssh-private-key-file <path-to-ssh-private-key> \
  --self-managed
```

¹⁶⁸ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1039\)](#).

1. **ALTERNATIVE Virtual IP** - if you don't have an external LB, and wish to use a VIRTUAL IP provided by kube-vip, specify these flags example below:

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host 196.168.1.10 \
  --virtual-ip-interface eth1
```

The output from this command is shortened here for reading clarity, but should start like this:

```
Generating cluster resources
cluster.cluster.x-k8s.io/preprovisioned-example created
cont.....
```

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:


```
cluster.cluster.x-k8s.io/preprovisioned-example condition met
```


⏳ Depending on the cluster size, it will take a few minutes to create.

When the command completes, you will have a running Kubernetes cluster! For bootstrap and custom YAML cluster creation, refer to the Additional Infrastructure Customization section of the documentation for Pre-provisioned: [Pre-provisioned Infrastructure \(see page 1200\)](#)

Use this command to get the Kubernetes `kubeconfig` for the new cluster and proceed to installing the DKP Kommander UI:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

 If changing the [Calico encapsulation](#) (see page 1231), D2iQ recommends changing it after cluster creation, but before production.

 If you need to increase [Docker Hub's rate limit](#)¹⁶⁹, use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

5.10.2.4.3 Audit logs

To modify Control Plane Audit logs settings using the information contained in the page [Configuring the Control Plane](#)¹⁷⁰.

5.10.2.4.4 Further Steps:

For more customized cluster creation, access the [Custom Installation and Additional Infrastructure Tools](#) (see page 1200) section. That section is for [Pre-Provisioned Override Files](#) (see page 1524), custom flags, and more that specify the secret as part of the create cluster command. If these are not specified, the overrides for your nodes will not be applied.

5.10.2.4.5 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPagelD=273320989>¹⁷¹.

5.10.2.4.6 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Pre-provisioned%3A%20Configure%20MetalLB&linkCreation=true&fromPagelD=273320989>¹⁷²

¹⁶⁹ <https://docs.docker.com/docker-hub/download-rate-limit/>

¹⁷⁰ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelD=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane>

¹⁷¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelD=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

¹⁷² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelD=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Pre-provisioned%3A+Configure+MetalLB>

5.10.2.5 Pre-provisioned: Configure MetalLB

5.10.2.5.1 Create a MetalLB configmap for your pre-provisioned infrastructure.

Choose one of the following two protocols you want to use to announce service IPs. If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can continue the installation process with [Pre-provisioned: Install Kommander \(see page 161\)](#).

To use MetalLB, create a MetalLB configMap for your Pre-provisioned infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)

Select one of the following procedures to create your MetalLB manifest for further editing.

5.10.2.5.2 Layer 2 configuration

Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.



- MetalLB IP address ranges/CIDRs should be within the node's primary network [subnet \(see page 1046\)](#).
- MetalLB IP address ranges/CIDRs and node subnet should not conflict with the Kubernetes cluster pod and service subnets.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:



The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
```

```

config: |
  address-pools:
  - name: default
    protocol: layer2
    addresses:
    - 192.168.1.240-192.168.1.250
EOF

```

Once complete, run the following `kubectl` command.


```
kubectl apply -f metallb-conf.yaml
```

5.10.2.5.3 BGP configuration

For a basic configuration featuring one BGP router and one IP address range, you need 4 pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:

 The following values are generic, enter your specific values into the fields where applicable.

```

cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
    - peer-address: 10.0.0.1
      peer-asn: 64501
      my-asn: 64500
    address-pools:
    - name: default
      protocol: bgp
      addresses:
      - 192.168.10.0/24
EOF

```

Once complete, run the following `kubectl` command.

```
kubectl apply -f metallb-conf.yaml
```

5.10.2.5.3.1 Next Step:

[Pre-provisioned: Install Kommander \(see page 161\)](#)

5.10.2.6 Pre-provisioned: Install Kommander

5.10.2.6.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install \(see page 121\)](#).
- Ensure you have a [default StorageClass \(see page 1395\)](#).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.10.2.6.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file \(see page 1422\)](#) for the deployment:

```
dkp install kommander --init > kommander.yaml
```

4. Edit the installer file to include configuration overrides for the `rook-ceph-cluster`. DKP's default configuration ships Ceph with [PVC based storage](#)¹⁷³ which requires your CSI provider to support PVC

¹⁷³ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/>

with type `volumeMode: Block`. As this is not possible with the default [local static provisioner](#) (see [page 91](#)), you can install Ceph in [host storage](#)¹⁷⁴ mode.

You can choose whether Ceph's object storage daemon (osd) pods should consume all or just some of the devices on your nodes. Include **one** of the following Overrides:

- a. To automatically assign all raw storage devices on all nodes to the Ceph cluster:

```

rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllDevices: true
        useAllNodes: true
        deviceFilter: "<<value>>"

```

- b. To assign specific storage devices on all nodes to the Ceph cluster:

```

rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllNodes: true
        useAllDevices: false
        deviceFilter: "^sdb."

```

Note: If you want to assign specific devices to specific nodes using the `deviceFilter` option, refer to [Specific Nodes and Devices](#)¹⁷⁵. For general information on the `deviceFilter` value, refer to [Storage Selection Settings](#)¹⁷⁶.

5. *If required:* Customize your `kommander.yaml`.


i See [Kommander Customizations](#) (see [page 1422](#)) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, etc.

¹⁷⁴ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/>

¹⁷⁵ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/#specific-nodes-and-devices>


¹⁷⁶ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

5.10.2.6.3 Enable DKP Catalog Applications and Install Kommander

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
        ref:
          tag: v2.6.1
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).

- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.10.2.6.4 Next Step:

[Pre-provisioned: Verify Install and Log in to UI \(see page 164\)](#)

5.10.2.7 Pre-provisioned: Verify Install and Log in to UI

Verify Kommander Install and Log in to the Dashboard UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.



NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
```

```
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.10.2.7.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

5.10.2.7.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{\n}}Password: {{.data.password|base64decode}}
{{\n}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/dashboard{{ "\n"}}
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbJUKVwLJMQL2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.10.2.7.3 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

5.10.2.7.4 Next Step:


[Pre-provisioned: Subsequent Cluster Creation](#) (see [page 166](#))

5.10.2.8 Pre-provisioned: Create Managed Clusters Using the DKP CLI


Enterprise

Gov Advanced

After initial cluster creation, you have the ability to create additional clusters from the CLI. In a previous step, the new cluster was created as Self-managed which allows it to be a Management cluster or a stand alone cluster. Subsequent new clusters are not self-managed as they will likely be Managed or Attached clusters to this Management Cluster.

 When creating [Managed clusters](#) (see [page 81](#)), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see [page 81](#))!

5.10.2.8.1 Make New Cluster Part of a Workspace


1. If you have an existing Workspace name, run this command to find the name:
 **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see [page 650](#)).

```
kubectl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.10.2.8.2 Name Your Cluster

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)¹⁷⁷ for more naming information.

When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:


1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=preprovisioned-additional
```


¹⁷⁷ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

5.10.2.8.3 Create a Kubernetes Cluster

Once you've defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster. This process will create a self-managed cluster to be used as the Management cluster.

 Before you create a new DKP cluster below, choose an external load balancer (LB) or [virtual IP](#) (see page 1236) and use the corresponding `dkp create cluster` command.

In a pre-provisioned environment, use the Kubernetes CSI and third party drivers for local volumes and other storage devices in your data center.

 DKP uses local static provisioner as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)¹⁷⁸ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the [inventory YAML](#) (see page 151) previously created.

1. This command uses the default external load balancer (LB) option:

```
dkp create cluster preprovisioned
  --cluster-name ${CLUSTER_NAME}
  --control-plane-endpoint-host <control plane endpoint host>
  --control-plane-endpoint-port <control plane endpoint port, if different than
  6443>
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml
  --ssh-private-key-file <path-to-ssh-private-key>
```


2. Use the wait command to monitor the cluster control-plane readiness:


¹⁷⁸ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```


Output:

```
cluster.cluster.x-k8s.io/preprovisioned-additional condition met
```

 If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

 NOTE: Depending on the cluster size, it will take a few minutes to create.

5.10.2.8.4 Manually Attach a DKP CLI Cluster to the Management Cluster

 These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > $
{MANAGED_CLUSTER_NAME}.conf
```


- You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

- Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

- Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

- You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

- This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

- Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

- Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
```

```
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubect1 get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.10.2.8.5 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.10.3 Pre-provisioned Air-gapped Install

This section provides instructions to install DKP in a Pre-provisioned air-gapped environment.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)



NOTE: For air-gapped, ensure you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, so you can extract the tarball on next page.

5.10.3.1 Next Step:

[Pre-provisioned Air-gapped: Configure Environment](#) (see page 170)

5.10.3.2 Pre-provisioned Air-gapped: Configure Environment

The instructions below outline how to fulfill the requirements for using pre-provisioned infrastructure when using air-gapped. In order to create a cluster, you must first setup the environment with necessary artifacts. All artifacts for Pre-provisioned Air-gapped need to get onto the bastion host. Artifacts needed by nodes

must be unpacked and distributed on the bastion before other provisioning will work in the absence of an internet connection.

There is a new complete DKP air-gapped bundle available to [download](#) (see page 78) which contains all the DKP components needed for air-gapped installation. (i.e. `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`)

5.10.3.2.1 Setup Process:

1. The bootstrap image must be extracted and loaded onto the [bastion host](#) (see page 1285).
2. Artifacts must be copied onto cluster hosts for nodes to access.
3. If using GPU, those artifacts must be positioned locally.
4. Registry seeded with images locally.

5.10.3.2.2 Load the Bootstrap Image

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz` from the [download site](#) (see page 78) mentioned above, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Load the bootstrap image on your bastion machine:

```
docker load -i konvoy-bootstrap-image-v2.6.2.tar
```

5.10.3.2.3 Copy Air-gapped Artifacts onto Cluster Hosts

Using the [Konvoy Image Builder](#) (see page 1461), you can copy the required artifacts onto your cluster hosts.

Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, continue below:

1. The kubernetes image bundle will be located in `kib/artifacts/images` and you will want to verify image and artifacts.
 - a. Verify the image bundles exist in `artifacts/images`:

```
$ ls artifacts/images/
kubernetes-images-1.26.14-d2iq.1.tar kubernetes-images-1.26.14-d2iq.1-
fips.tar
```

- b. Verify the artifacts for your OS exist in the `artifacts/` directory and export the appropriate variables:

```
$ ls artifacts/
1.26.14_centos_7_x86_64.tar.gz      1.26.14_redhat_8_x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rhel-8.6-x86_64_fips.tar.gz  pip-packages.tar.gz
1.26.14_centos_7_x86_64_fips.tar.gz  1.26.14_rocky_9_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64_fips.tar.gz  containerd-1.6.17-
d2iq.1-rocky-9.0-x86_64.tar.gz
1.26.14_redhat_7_x86_64.tar.gz      1.26.14_ubuntu_20_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rocky-9.1-x86_64.tar.gz
1.26.14_redhat_7_x86_64_fips.tar.gz  containerd-1.6.17-d2iq.1-centos-7.9-
x86_64.tar.gz      containerd-1.6.17-d2iq.1-rhel-8.4-x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-ubuntu-20.04-x86_64.tar.gz
1.26.14_redhat_8_x86_64.tar.gz      containerd-1.6.17-d2iq.1-centos-7.9-
x86_64_fips.tar.gz  containerd-1.6.17-d2iq.1-rhel-8.6-x86_64.tar.gz
images
```

- c. For example, for RHEL 8.4 you would set:

```
export OS_PACKAGES_BUNDLE=1.26.14_redhat_8_x86_64.tar.gz
export CONTAINERD_BUNDLE=containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz
```

2. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```
export CONTROL_PLANE_1_ADDRESS="<<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<<control-plane-address-3>"
export WORKER_1_ADDRESS="<<worker-address-1>"
export WORKER_2_ADDRESS="<<worker-address-2>"
export WORKER_3_ADDRESS="<<worker-address-3>"
export WORKER_4_ADDRESS="<<worker-address-4>"
export SSH_USER="<<ssh-user>"
export SSH_PRIVATE_KEY_FILE="<<private key file>"
```

`SSH_PRIVATE_KEY_FILE` must be either the name of the SSH private key file in your working directory or an absolute path to the file in your user's home directory.

3. Generate an `inventory.yaml` which is automatically picked up by the `konvoy-image upload` in the next step.

```
cat <<EOF > inventory.yaml
all:
```

```

vars:
  ansible_user: $SSH_USER
  ansible_port: 22
  ansible_ssh_private_key_file: $SSH_PRIVATE_KEY_FILE
hosts:
  $CONTROL_PLANE_1_ADDRESS:
    ansible_host: $CONTROL_PLANE_1_ADDRESS
  $CONTROL_PLANE_2_ADDRESS:
    ansible_host: $CONTROL_PLANE_2_ADDRESS
  $CONTROL_PLANE_3_ADDRESS:
    ansible_host: $CONTROL_PLANE_3_ADDRESS
  $WORKER_1_ADDRESS:
    ansible_host: $WORKER_1_ADDRESS
  $WORKER_2_ADDRESS:
    ansible_host: $WORKER_2_ADDRESS
  $WORKER_3_ADDRESS:
    ansible_host: $WORKER_3_ADDRESS
  $WORKER_4_ADDRESS:
    ansible_host: $WORKER_4_ADDRESS
EOF

```

4. Upload the artifacts onto cluster hosts with the following command:

```

konvoy-image upload artifacts \
  --container-images-dir=./artifacts/images/ \
  --os-packages-bundle=./artifacts/$OS_PACKAGES_BUNDLE \
  --containerd-bundle=artifacts/$CONTAINERD_BUNDLE \
  --pip-packages-bundle=./artifacts/pip-packages.tar.gz

```

KIB uses [variable overrides](#) (see [page 1513](#)) to specify base image and container images to use in your new machine image. The variable overrides files for NVIDIA and FIPS can be ignored unless adding an overlay feature.

5.10.3.2.3.1 Next Step:

[Pre-provisioned Air-gapped: Load the Registry](#) (see [page 173](#))

5.10.3.3 Pre-provisioned Air-gapped: Load the Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine](#) (see [page 1285](#)) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.



If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see [page 1533](#)) page for more information.

1. [Download \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz` , and extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

3. Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```



It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.



To increase [Docker Hub's rate limit](#)¹⁷⁹ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

5.10.3.3.1 Kommander Load Images

If you are operating in an air-gapped environment, a [local container registry \(see page 1533\)](#) containing all the necessary installation images, including the Kommander images is required. See below for how to push the necessary images to this registry.

5.10.3.3.2 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

¹⁷⁹ <https://docs.docker.com/docker-hub/download-rate-limit/>

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.10.3.3.3 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.10.3.3.4 Next Step:

[Pre-provisioned Air-gapped: Define Infrastructure \(see page 176\)](#)

5.10.3.4 Pre-provisioned Air-gapped: Define Infrastructure

Define the cluster hosts and infrastructure

Konvoy needs to know how to access your cluster hosts. This is done using inventory resources. For initial cluster creation, you must define a control-plane and at least one worker pool.

5.10.3.4.1 Define your Infrastructure

1. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```
export CONTROL_PLANE_1_ADDRESS="<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<control-plane-address-3>"
export WORKER_1_ADDRESS="<worker-address-1>"
export WORKER_2_ADDRESS="<worker-address-2>"
export WORKER_3_ADDRESS="<worker-address-3>"
export WORKER_4_ADDRESS="<worker-address-4>"
export SSH_USER="<ssh-user>"
export SSH_PRIVATE_KEY_SECRET_NAME="CLUSTER_NAME-ssh-key"
```

2. Use the following template to help you define your infrastructure. The environment variables that you set in the previous step automatically replace the variable names when the inventory YAML file is created.

```

cat <<EOF > preprovisioned_inventory.yaml
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: $CLUSTER_NAME-control-plane
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: $CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    # Create as many of these as needed to match your infrastructure
    # Note that the command line parameter --control-plane-replicas determines
    # how many control plane nodes will actually be used.
    #
    - address: $CONTROL_PLANE_1_ADDRESS
    - address: $CONTROL_PLANE_2_ADDRESS
    - address: $CONTROL_PLANE_3_ADDRESS
  sshConfig:
    port: 22
    # This is the username used to connect to your infrastructure. This user
    # must be root or
    # have the ability to use sudo without a password
    user: $SSH_USER
    privateKeyRef:
      # This is the name of the secret you created in the previous step. It
      # must exist in the same
      # namespace as this inventory object.
      name: $SSH_PRIVATE_KEY_SECRET_NAME
      namespace: default
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: $CLUSTER_NAME-md-0
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: $CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    - address: $WORKER_1_ADDRESS
    - address: $WORKER_2_ADDRESS
    - address: $WORKER_3_ADDRESS
    - address: $WORKER_4_ADDRESS

```



```
sshConfig:
  port: 22
  user: $SSH_USER
  privateKeyRef:
    name: $SSH_PRIVATE_KEY_SECRET_NAME
    namespace: default
EOF
```

5.10.3.4.2 Next Step:

[Pre-provisioned Air-gapped: Define Control Plane Endpoint \(see page 177\)](#)

5.10.3.5 Pre-provisioned Air-gapped: Define Control Plane Endpoint

5.10.3.5.1 Define the Control Plane Endpoint for your cluster as well as connection mechanism

A control plane should have three, five, or seven nodes, so it can remain available if one or more nodes fail. A control plane with one node should not be used in production.

In addition, the control plane should have an endpoint that remains available if some nodes fail.

```

          ----- cp1.example.com:6443
          |
lb.example.com:6443 ----- cp2.example.com:6443
          |
          ----- cp3.example.com:6443

```

In this example, the control plane endpoint host is `lb.example.com`, and the control plane endpoint port is `6443`. The control plane nodes are `cp1.example.com`, `cp2.example.com`, and `cp3.example.com`. The port of each API server is `6443`.

5.10.3.5.2 Select your Connection Mechanism

A virtual IP is the address that the client uses to connect to the service. A load balancer is the device that distributes the client connections to the backend servers. Before you create a new DKP cluster, choose an external load balancer(LB) or virtual IP.

- **External load balancer**

It is recommended that an external load balancer be the control plane endpoint. To distribute request load among the control plane machines, configure the load balancer to send requests to all the control plane machines. Configure the load balancer to send requests only to control plane machines that are responding to API requests.

- **Built-in virtual IP**

If an external load balancer is not available, use the [built-in virtual IP](#) (see page 1236). The virtual IP is *not* a load balancer; it does not distribute request load among the control plane machines. However, if the machine receiving requests does not respond to them, the virtual IP automatically moves to another machine.

5.10.3.5.3 Single-Node Control Plane



Do not use a single-node control plane in a production cluster.

A control plane with one node can use its single node as the endpoint, so you will not require an external load balancer, or a built-in virtual IP. At least one control plane node must always be running. Therefore, to upgrade a cluster with one control plane node, a spare machine must be available in the control plane inventory. This machine is used to provision the new node before the old node is deleted.



Modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane](#) (see page 1538).

When the API server endpoints are defined, you can create the cluster using the link in Next Step below.

5.10.3.5.4 Known Limitations



Be aware of these limitations in the current release of DKP.

The control plane endpoint port is also used as the API server port on each control plane machine. The default port is 6443. Before you create the cluster, ensure the port is available for use on each control plane machine.

5.10.3.5.5 Next Step:

[Pre-provisioned Air-gapped: Create a New Cluster](#) (see page 179)


5.10.3.6 Pre-provisioned Air-gapped: Create a Management Cluster

If your cluster is air-gapped or you have a local registry, you must provide additional arguments when creating the cluster. These tell the cluster where to locate the local registry to use by defining the URL.

```
export REGISTRY_URL=<https/http>://<registry-address>:<registry-port>
export REGISTRY_CA=<path to the CA on the bastion>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

- `REGISTRY_URL` : the address of an existing registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
- `REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. Konvoy will configure the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the AMLs are not already configured to trust this CA.
- `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
- `REGISTRY_PASSWORD` : optional if username is not set.

5.10.3.6.1 Name Your Cluster

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)¹⁸⁰ for more naming information.


When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:


```
export CLUSTER_NAME=<preprovisioned-example>
```

 Before you create a new DKP cluster below, choose an external load balancer(LB) or [virtual IP](#) (see [page 1236](#)) and use the corresponding `dkp create cluster` command.

¹⁸⁰ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

5.10.3.6.2 Create an Air-gapped Kubernetes Cluster

After you have defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster.

 DKP uses local static provisioner as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)¹⁸¹ compatible storage that is suitable for production.


After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the inventory.

The create cluster command below includes the `--self-managed` flag. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.

1. The command uses the default external load balancer (see alternative Step 1 below for virtual IP):

```
dkp create cluster preprovisioned --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host <control plane endpoint host> \
  --control-plane-endpoint-port <control plane endpoint port, if different than
6443> \
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml \
  --ssh-private-key-file <path-to-ssh-private-key> \
  --registry-mirror-url=${REGISTRY_URL} \
  --registry-mirror-cacert=${REGISTRY_CA} \
  --registry-mirror-username=${REGISTRY_USERNAME} \
  --registry-mirror-password=${REGISTRY_PASSWORD} \
  --self-managed
```

 If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

¹⁸¹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

1. **Virtual IP ALTERNATIVE** - if you don't have an external LB, and wish to use a [VIRTUAL IP \(see page 1236\)](#) provided by kube-vip, specify these flags example below:

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host 196.168.1.10 \
  --virtual-ip-interface eth1
```

The output from this command is shortened here for reading clarity, but should start like this:

```
Generating cluster resources
cluster.cluster.x-k8s.io/preprovisioned-example created
cont.....
```

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:

```
cluster.cluster.x-k8s.io/preprovisioned-example condition met
```



NOTE: Depending on the cluster size, it will take a few minutes to create.

When the command completes, you will have a running Kubernetes cluster! For bootstrap and custom YAML cluster creation, refer to the Additional Infrastructure Customization section of the documentation for Pre-provisioned: [Pre-provisioned Infrastructure \(see page 1200\)](#)

Use this command to get the Kubernetes `kubeconfig` for the new cluster and proceed to installing the DKP Kommander UI:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```



If changing the [Calico encapsulation \(see page 1231\)](#), D2iQ recommends changing it after cluster creation, but before production.

5.10.3.6.3 Audit Logs

To modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane](#)¹⁸².

5.10.3.6.4 Further Steps:

For more customized cluster creation, see the [Pre-Provisioned Additional Configurations](#) (see page 1200) section.



NOTE: The section mentioned above is for [overrides for your clusters](#) (see page 1524), custom flags, and more that specify the secret as part of the create cluster command. If these are not specified, the overrides for your nodes will not be applied.

5.10.3.6.5 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>¹⁸³¹⁸⁴.

5.10.3.6.6 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Pre-provisioned%20Air-gapped%3A%20Configure%20MetalLB&linkCreation=true&fromPageId=273681652>¹⁸⁵

¹⁸² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273681652&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane>

¹⁸³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

¹⁸⁴ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

¹⁸⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273681652&linkCreation=true&spaceKey=DENT&title=%282.6%29+Pre-provisioned+Air-gapped%3A+Configure+MetalLB>

5.10.3.7 Pre-provisioned Air-gapped: Configure MetalLB

5.10.3.7.1 Create a MetalLB configmap for your pre-provisioned infrastructure.

Choose one of the following two protocols you want to use to announce service IPs. If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can continue the installation process [Pre-provisioned Air-gapped: Install Kommander](#) (see page 185).

To use MetalLB, create a MetalLB configMap for your Pre-provisioned infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)


Select one of the following procedures to create your MetalLB manifest for further editing.

5.10.3.7.2 Layer 2 configuration

Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:

 The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    address-pools:
    - name: default
      protocol: layer2
      addresses:
      - 192.168.1.240-192.168.1.250
EOF
```

Once complete, run the following `kubectl` command.


```
kubectl apply -f metallb-conf.yaml
```

5.10.3.7.3 BGP configuration

For a basic configuration featuring one BGP router and one IP address range, you need 4 pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:

 The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
      - peer-address: 10.0.0.1
        peer-asn: 64501
        my-asn: 64500
    address-pools:
      - name: default
        protocol: bgp
        addresses:
          - 192.168.10.0/24
EOF
```

Once complete, run the following `kubectl` command.

```
kubectl apply -f metallb-conf.yaml
```

5.10.3.7.3.1 Next Step:

[Pre-provisioned Air-gapped: Install Kommander \(see page 185\)](#)

5.10.3.8 Pre-provisioned Air-gapped: Install Kommander

5.10.3.8.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Ensure you have loaded all necessary images for your configuration. See [Load the Images into Your Registry: Air-gapped Environments](#) (see page 1400).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.10.3.8.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init --airgapped > kommander.yaml
```

4. Edit the installer file to include configuration overrides for the `rook-ceph-cluster`. DKP's default configuration ships Ceph with [PVC based storage](#)¹⁸⁶ which requires your CSI provider to support PVC with type `volumeMode: Block`. As this is not possible with the default [local static provisioner](#) (see page 91), you can install Ceph in [host storage](#)¹⁸⁷ mode.

You can choose whether Ceph's object storage daemon (`osd`) pods should consume all or just some of the devices on your nodes. Include **one** of the following Overrides:

- a. To automatically assign all raw storage devices on all nodes to the Ceph cluster:

```
rook-ceph-cluster:
  enabled: true
```

¹⁸⁶ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/>

¹⁸⁷ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/>

```
values: |
  cephClusterSpec:
    storage:
      storageClassDeviceSets: []
      useAllDevices: true
      useAllNodes: true
      deviceFilter: "<<value>>"
```

- b. To assign specific storage devices on all nodes to the Ceph cluster:

```
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllNodes: true
        useAllDevices: false
        deviceFilter: "^sdb."
```

Note: If you want to assign specific devices to specific nodes using the `deviceFilter` option, refer to [Specific Nodes and Devices](#)¹⁸⁸. For general information on the `deviceFilter` value, refer to [Storage Selection Settings](#)¹⁸⁹.

5. *If required:* Customize your `kommander.yaml`.

i See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, etc.

5.10.3.8.3 Enable DKP Catalog Applications and Install Kommander in an Air-gapped Environment

i If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add the following values to enable DKP Catalog Applications:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
```

¹⁸⁸ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/#specific-nodes-and-devices>

¹⁸⁹ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

```

kind: Installation
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      path: ./dkp-catalog-applications-v2.6.1.tar.gz

```

⚠️ If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```

dkp install kommander --installer-config kommander.yaml --kubeconfig=${
{CLUSTER_NAME}.conf \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz
\
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-
v2.6.1.tar.gz

```



Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.10.3.8.4 Next Step:

[Pre-provisioned Air-gapped: Verify Install and Log in to UI](#) (see page 188)

5.10.3.9 Pre-provisioned Air-gapped: Verify Install and Log in to UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.



NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the install command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
```

```
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.10.3.9.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

5.10.3.9.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{\n}}Password: {{.data.password|base64decode}}
{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/
dashboard{{ "\n"}}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQL2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see page 597):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.10.3.9.3 Dashboard UI Functions


After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see page 582) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

5.10.3.9.4 Next Step:


[Pre-provisioned Air-gapped: Subsequent Cluster Creation](#) (see page 191)

5.10.3.10 Pre-provisioned Air-gapped: Create Managed Clusters Using the DKP CLI

After initial cluster creation, you have the ability to create additional clusters from the CLI. In a previous step, the new cluster was created as Self-managed which allows it to be a Management cluster or a stand alone cluster. Subsequent new clusters are not self-managed as they will likely be Managed or Attached clusters to this Management Cluster.

 When creating [Managed clusters](#) (see page 81), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see page 81)!

5.10.3.10.1 Make New Cluster Part of a Workspace


1. If you have an existing Workspace name, run this command to find the name:
 **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see page 650).

```
kubect! get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.10.3.10.2 Name Your Cluster

-  The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)¹⁹⁰ for more naming information. When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.


Follow these steps:

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=<preprovisioned-additional>
```

5.10.3.10.3 Create a Kubernetes Cluster

Once you've defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster. This process will create a self-managed cluster to be used as the Management cluster.

-  Before you create a new DKP cluster below, choose an external load balancer (LB) or [virtual IP \(see page 1236\)](#) and use the corresponding `dkp create cluster` command.

In a pre-provisioned environment, use the Kubernetes CSI and third party drivers for local volumes and other storage devices in your data center.

¹⁹⁰ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>



DKP uses local static provisioner as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)¹⁹¹ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the [inventory YAML](#) (see page 151) previously created.

1. This command uses the default external load balancer (LB) option:

```
dkp create cluster preprovisioned --cluster-name ${CLUSTER_NAME}
  --control-plane-endpoint-host <control plane endpoint host>
  --control-plane-endpoint-port <control plane endpoint port, if different than
  6443>
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml
  --ssh-private-key-file <path-to-ssh-private-key>
  --registry-mirror-url=${REGISTRY_URL} \
  --registry-mirror-cacert=${REGISTRY_CA} \
  --registry-mirror-username=${REGISTRY_USERNAME} \
  --registry-mirror-password=${REGISTRY_PASSWORD}
```

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
  timeout=30m
```

Output:


```
cluster.cluster.x-k8s.io/preprovisioned-additional condition met
```



NOTE: Depending on the cluster size, it will take a few minutes to create.

¹⁹¹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

5.10.3.10.4 Manually Attach a DKP CLI Cluster to the Management Cluster

 These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

7. This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```

apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>

```

8. Create this secret in the desired workspace:

```

kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}

```

9. Create this `kommandercluster` object to attach the cluster to the workspace:

```

cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF

```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```

kubectl get kommanderclusters -A

```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.10.3.10.5 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.10.4 Pre-provisioned FIPS Install

This section provides instructions to install DKP in a Pre-provisioned non-air-gapped environment using FIPS.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.10.4.1 Next Step:

[Pre-provisioned FIPS: Define Infrastructure](#) (see page 195)

5.10.4.2 Pre-provisioned FIPS: Define Infrastructure

Define the cluster hosts and infrastructure

Konvoy needs to know how to access your cluster hosts. This is done using inventory resources. For initial cluster creation, you must define a control-plane and at least one worker pool.

5.10.4.2.1 Define your Infrastructure

1. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```
export CONTROL_PLANE_1_ADDRESS="<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<control-plane-address-3>"
export WORKER_1_ADDRESS="<worker-address-1>"
export WORKER_2_ADDRESS="<worker-address-2>"
export WORKER_3_ADDRESS="<worker-address-3>"
export WORKER_4_ADDRESS="<worker-address-4>"
export SSH_USER="<ssh-user>"
export SSH_PRIVATE_KEY_SECRET_NAME="CLUSTER_NAME-ssh-key"
```

2. Use the following template to help you define your infrastructure. The environment variables that you set in the previous step automatically replace the variable names when the inventory YAML file is created.

```
cat <<EOF > preprovisioned_inventory.yaml
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
```

```

metadata:
  name: $CLUSTER_NAME-control-plane
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: $CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    # Create as many of these as needed to match your infrastructure
    # Note that the command line parameter --control-plane-replicas determines
    # how many control plane nodes will actually be used.
    #
    - address: $CONTROL_PLANE_1_ADDRESS
    - address: $CONTROL_PLANE_2_ADDRESS
    - address: $CONTROL_PLANE_3_ADDRESS
  sshConfig:
    port: 22
    # This is the username used to connect to your infrastructure. This user
    # must be root or
    # have the ability to use sudo without a password
    user: $SSH_USER
    privateKeyRef:
      # This is the name of the secret you created in the previous step. It
      # must exist in the same
      # namespace as this inventory object.
      name: $SSH_PRIVATE_KEY_SECRET_NAME
      namespace: default
  ---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: $CLUSTER_NAME-md-0
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: $CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    - address: $WORKER_1_ADDRESS
    - address: $WORKER_2_ADDRESS
    - address: $WORKER_3_ADDRESS
    - address: $WORKER_4_ADDRESS
  sshConfig:
    port: 22
    user: $SSH_USER
    privateKeyRef:
      name: $SSH_PRIVATE_KEY_SECRET_NAME
      namespace: default
EOF

```

5.10.4.2 Next Step:


[Pre-provisioned FIPS: Define Control Plane Endpoint \(see page 198\)](#)

5.10.4.3 Pre-provisioned FIPS: Define Control Plane Endpoint

5.10.4.3.1 Define the Control Plane Endpoint for your cluster as well as connection mechanism

A control plane should have three, five, or seven nodes, so it can remain available if one or more nodes fail. A control plane with one node should not be used in production.

In addition, the control plane should have an endpoint that remains available if some nodes fail.



In this example, the control plane endpoint host is `lb.example.com`, and the control plane endpoint port is `6443`. The control plane nodes are `cp1.example.com`, `cp2.example.com`, and `cp3.example.com`. The port of each API server is `6443`.

5.10.4.3.2 Select your Connection Mechanism

A virtual IP is the address that the client uses to connect to the service. A load balancer is the device that distributes the client connections to the backend servers. Before you create a new DKP cluster, choose an external load balancer(LB) or virtual IP.

- **External load balancer**

It is recommended that an external load balancer be the control plane endpoint. To distribute request load among the control plane machines, configure the load balancer to send requests to all the control plane machines. Configure the load balancer to send requests only to control plane machines that are responding to API requests.

- **Built-in virtual IP**

If an external load balancer is not available, use the [built-in virtual IP \(see page 1236\)](#). The virtual IP is *not* a load balancer; it does not distribute request load among the control plane machines. However, if the machine receiving requests does not respond to them, the virtual IP automatically moves to another machine.

5.10.4.3.3 Single-Node Control Plane



Do not use a single-node control plane in a production cluster.

A control plane with one node can use its single node as the endpoint, so you will not require an external load balancer, or a built-in virtual IP. At least one control plane node must always be running. Therefore, to upgrade a cluster with one control plane node, a spare machine must be available in the control plane inventory. This machine is used to provision the new node before the old node is deleted.



Modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane](#) (see page 1538).

When the API server endpoints are defined, you can create the cluster using the link in Next Step below.

5.10.4.3.4 Known Limitations



Be aware of these limitations in the current release of DKP.

The control plane endpoint port is also used as the API server port on each control plane machine. The default port is 6443. Before you create the cluster, ensure the port is available for use on each control plane machine.

5.10.4.3.5 Next Step:

[Pre-provisioned FIPS: Create a New Cluster](#) (see page 199)

5.10.4.4 Pre-provisioned FIPS: Create a Management Cluster

5.10.4.4.1 Deploying a Cluster in FIPS mode

In order to create a cluster in FIPS mode, we must inform the bootstrap controllers of the appropriate image repository and version tags of the official D2iQ FIPS builds of Kubernetes.

5.10.4.4.1.1 Supported FIPS Builds

Component	Repository	Version
Kubernetes	docker.io/mesosphere ¹⁹²	v1.26.14+fips.0
etcd	docker.io/mesosphere ¹⁹³	3.5.10+fips.0

5.10.4.4.1.2 Name Your Cluster

ⓘ The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](https://kubernetes.io/docs/concepts/overview/working-with-objects/names/)¹⁹⁴ for more naming information.

When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=<preprovisioned-example>
```


¹⁹² <https://hub.docker.com/layers/mesosphere/kube-apiserver/v1.26.14/images/sha256-3c559c75123fa40d588bc2d9050938a0b7dd1dab4e38a4fd24847cd0c0f24694?context=explore>

¹⁹³ https://hub.docker.com/layers/mesosphere/etcd/v3.5.10_fips.0/images/sha256-61df32b8ec1e6229d8bd94caad07331c7e0c0b8e1947af8152c69eb2f04fb771?context=explore


¹⁹⁴ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

5.10.4.4.1.3 Create a Kubernetes Cluster

Once you've defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster. This process will create a self-managed cluster to be used as the Management cluster.

 Before you create a new DKP cluster below, choose an external load balancer (LB) or [virtual IP](#) (see [page 1236](#)) and use the corresponding `dkp create cluster` command.

In a pre-provisioned environment, use the Kubernetes CSI and third party drivers for local volumes and other storage devices in your data center.

 DKP uses local static provisioner as the [default storage provider](#) (see [page 91](#)) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)¹⁹⁵ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the [inventory YAML](#) (see [page 151](#)) previously created.

The create cluster command below includes the `--self-managed` flag. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.

1. This command uses the default external load balancer (LB) option (see alternative Step 1 below for virtual IP):

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host <control plane endpoint host> \
  --control-plane-endpoint-port <control plane endpoint port, if different than
6443> \
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml \
  --ssh-private-key-file <path-to-ssh-private-key> \
  --kubernetes-version=v1.26.14+fips.0 \
```

¹⁹⁵ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>


```
--etcd-version=3.5.10+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--self-managed
```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

1. **ALTERNATIVE Virtual IP** - if you don't have an external LB, and wish to use a VIRTUAL IP provided by kube-vip, specify these flags example below:

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host 196.168.1.10 \
  --virtual-ip-interface eth1
```

The output from this command is shortened here for reading clarity, but should start like this:

```
Generating cluster resources
cluster.cluster.x-k8s.io/preprovisioned-example created
cont.....
```

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:


```
cluster.cluster.x-k8s.io/preprovisioned-example condition met
```


= Depending on the cluster size, it will take a few minutes to create.

When the command completes, you will have a running Kubernetes cluster! For bootstrap and custom YAML cluster creation, refer to the Additional Infrastructure Customization section of the documentation for Pre-provisioned: [Pre-provisioned Infrastructure](#) (see page 1200)

Use this command to get the Kubernetes `kubeconfig` for the new cluster and proceed to installing the DKP Kommander UI:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

 If changing the [Calico encapsulation](#) (see page 1231), D2iQ recommends changing it after cluster creation, but before production.

 If you need to increase [Docker Hub's rate limit](#)¹⁹⁶, use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

5.10.4.4.1.4 Audit logs

To modify Control Plane Audit logs settings using the information contained in the page [Configuring the Control Plane](#)^{197,198}.

5.10.4.4.1.5 Further Steps:

For more customized cluster creation, access the [Custom Installation and Additional Infrastructure Tools](#) (see page 1200) section. That section is for [Pre-Provisioned Override Files](#) (see page 1524), custom flags, and more that specify the secret as part of the create cluster command. If these are not specified, the overrides for your nodes will not be applied.

For more information regarding FIPS, refer to the [FIPS 140-2 Compliance](#) (see page 1526) section of the documentation.

5.10.4.4.1.6 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

¹⁹⁶ <https://docs.docker.com/docker-hub/download-rate-limit/>

¹⁹⁷ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane>

¹⁹⁸ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane>

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>¹⁹⁹²⁰⁰

5.10.4.4.1.7 Next Step:

[Pre-provisioned FIPS: Configure MetalLB \(see page 204\)](#)

5.10.4.5 Pre-provisioned FIPS: Configure MetalLB

5.10.4.5.1 Create a MetalLB configmap for your pre-provisioned infrastructure.

Choose one of the following two protocols you want to use to announce service IPs. If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can continue the installation process.

To use MetalLB, create a MetalLB configMap for your Pre-provisioned infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)

Select one of the following procedures to create your MetalLB manifest for further editing.

5.10.4.5.2 Layer 2 configuration

Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:



The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
```

¹⁹⁹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

²⁰⁰ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

```
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    address-pools:
    - name: default
      protocol: layer2
      addresses:
      - 192.168.1.240-192.168.1.250
EOF
```

Once complete, run the following `kubectl` command.


```
kubectl apply -f metallb-conf.yaml
```

5.10.4.5.3 BGP configuration

For a basic configuration featuring one BGP router and one IP address range, you need 4 pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:

 The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
    - peer-address: 10.0.0.1
      peer-asn: 64501
      my-asn: 64500
    address-pools:
```

```

- name: default
  protocol: bgp
  addresses:
    - 192.168.10.0/24
EOF

```

Once complete, run the following `kubectl` command.

```
kubectl apply -f metallb-conf.yaml
```

5.10.4.5.3.1 Next Step:

[Pre-provisioned FIPS: Install Kommander \(see page 205\)](#)

5.10.4.6 Pre-provisioned FIPS: Install Kommander

5.10.4.6.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install \(see page 121\)](#).
- Ensure you have a [default StorageClass \(see page 1395\)](#).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.10.4.6.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file \(see page 1422\)](#) for the deployment:

```
dkp install kommander --init > kommander.yaml
```

4. Edit the installer file to include configuration overrides for the `rook-ceph-cluster`. DKP's default configuration ships Ceph with [PVC based storage](#)²⁰¹ which requires your CSI provider to support PVC with type `volumeMode: Block`. As this is not possible with the default [local static provisioner](#) (see [page 91](#)), you can install Ceph in [host storage](#)²⁰² mode.

You can choose whether Ceph's object storage daemon (osd) pods should consume all or just some of the devices on your nodes. Include **one** of the following Overrides:

- a. To automatically assign all raw storage devices on all nodes to the Ceph cluster:

```
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllDevices: true
        useAllNodes: true
        deviceFilter: "<<value>>"
```

- b. To assign specific storage devices on all nodes to the Ceph cluster:

```
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllNodes: true
        useAllDevices: false
        deviceFilter: "^sdb."
```

Note: If you want to assign specific devices to specific nodes using the `deviceFilter` option, refer to [Specific Nodes and Devices](#)²⁰³. For general information on the `deviceFilter` value, refer to [Storage Selection Settings](#)²⁰⁴.

5. *If required:* Customize your `kommander.yaml`.

See [Kommander Customizations](#) (see [page 1422](#)) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, etc.


²⁰¹ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/>

²⁰² <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/>

²⁰³ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/#specific-nodes-and-devices>


²⁰⁴ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

5.10.4.6.3 Enable DKP Catalog Applications and Install Kommander

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
        ref:
          tag: v2.6.1
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).

- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.10.4.6.4 Next Step:

[Pre-provisioned FIPS: Verify Install and Log in to UI](#) (see page 209)

5.10.4.7 Pre-provisioned FIPS: Verify Install and Log in to UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.



NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
```



```
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.10.4.7.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

5.10.4.7.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{{ "\n"}}Password: {{.data.password|base64decode}}
{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/dashboard{{ "\n"}}
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbJUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.10.4.7.3 Dashboard UI Functions


After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

5.10.4.7.4 Next Step:


[Pre-provisioned FIPS: Subsequent Cluster Creation](#) (see [page 211](#))

5.10.4.8 Pre-provisioned FIPS: Create Managed Clusters Using the DKP CLI

After initial cluster creation, you have the ability to create additional clusters from the CLI. In a previous step, the new cluster was created as Self-managed which allows it to be a Management cluster or a stand alone cluster. Subsequent new clusters are not self-managed as they will likely be Managed or Attached clusters to this Management Cluster.

 When creating [Managed clusters](#) (see [page 81](#)), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see [page 81](#))!

5.10.4.8.1 Make New Cluster Part of a Workspace


1. If you have an existing Workspace name, run this command to find the name:
 **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see [page 650](#)).

```
kubectl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.10.4.8.2 Name Your Cluster

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)²⁰⁵ for more naming information.

When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:


1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=<preprovisioned-additional>
```


²⁰⁵ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

5.10.4.8.3 Create a Kubernetes Cluster

Once you've defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster. This process will create a self-managed cluster to be used as the Management cluster.

 Before you create a new DKP cluster below, choose an external load balancer (LB) or [virtual IP](#) (see page 1236) and use the corresponding `dkp create cluster` command.

In a pre-provisioned environment, use the Kubernetes CSI and third party drivers for local volumes and other storage devices in your data center.

 DKP uses local static provisioner as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)²⁰⁶ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the [inventory YAML](#) (see page 151) previously created.

1. This command uses the default external load balancer (LB) option:

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host <control plane endpoint host> \
  --control-plane-endpoint-port <control plane endpoint port, if different than
6443> \
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml \
  --ssh-private-key-file <path-to-ssh-private-key> \
  --kubernetes-version=v1.26.6+fips.0 \
  --etcd-version=3.5.6+fips.0 \
  --kubernetes-image-repository=docker.io/mesosphere
```

²⁰⁶ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:

```
cluster.cluster.x-k8s.io/preprovisioned-additional condition met
```



NOTE: Depending on the cluster size, it will take a few minutes to create.

5.10.4.8.4 Manually Attach a DKP CLI Cluster to the Management Cluster



These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dcp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > $
{MANAGED_CLUSTER_NAME}.conf
```

- You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

- Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

- Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

- You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

- This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

- Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

- Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
```

```
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubect1 get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.10.4.8.5 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.10.5 Pre-provisioned FIPS Air-gapped Install

This section provides instructions to install DKP in a Pre-provisioned non-air-gapped environment using FIPS.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)



NOTE: For air-gapped, ensure you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.1_linux_amd64.tar.gz`, so you can extract the tarball on next page.

5.10.5.1 Next Step:

[Pre-provisioned Air-gapped FIPS: Configure Environment](#) (see page 215)

5.10.5.2 Pre-provisioned Air-gapped FIPS: Configure Environment

In order to create a cluster in a Pre-provisioned Air-gapped environment with FIPS, you must first prepare the environment.

The instructions below outline how to fulfill the requirements for using pre-provisioned infrastructure when using air-gapped. In order to create a cluster, you must first setup the environment with necessary artifacts. All artifacts for Pre-provisioned Air-gapped need to get onto the bastion host. Artifacts needed by nodes must be unpacked and distributed on the bastion before other provisioning will work in the absence of an internet connection.

There is a new complete DKP air-gapped bundle available to [download](#) (see page 78) which contains all the DKP components needed for air-gapped installation. (i.e. `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`)

5.10.5.2.1 Setup Process:

1. The bootstrap image must be extracted and loaded onto the [bastion host](#) (see page 1285).
2. Artifacts must be copied onto cluster hosts for nodes to access.
3. If using GPU, those artifacts must be positioned locally.
4. Registry seeded with images locally.

5.10.5.2.2 Load the Bootstrap Image

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz` from the [download site](#) (see page 78) mentioned above, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Load the bootstrap image on your bastion machine:

```
docker load -i konvoy-bootstrap-image-v2.6.2.tar
```

5.10.5.2.3 Copy air-gapped artifacts onto cluster hosts

Using the [Konvoy Image Builder](#) (see page 1461), you can copy the required artifacts onto your cluster hosts.

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.1_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.1_linux_amd64.tar.gz && cd dkp-v2.6.1/kib
```

2. The kubernetes image bundle will be located in `kib/artifacts/images` and you will want to verify image and artifacts.
 - a. Verify the image bundles exist in `artifacts/images` :


```
$ ls artifacts/images/
kubernetes-images-1.26.6-d2iq.1.tar kubernetes-images-1.26.6-d2iq.1-
fips.tar
```

- b. Verify the artifacts for your OS exist in the `artifacts/` directory and export the appropriate variables:

```
$ ls artifacts/
1.26.6_centos_7_x86_64.tar.gz          1.26.6_redhat_8_x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rhel-8.6-x86_64_fips.tar.gz  pip-packages.tar.gz
1.26.6_centos_7_x86_64_fips.tar.gz  1.26.6_rocky_9_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64_fips.tar.gz  containerd-1.6.17-
d2iq.1-rocky-9.0-x86_64.tar.gz
1.26.6_redhat_7_x86_64.tar.gz          1.26.6_ubuntu_20_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rocky-9.1-x86_64.tar.gz
1.26.6_redhat_7_x86_64_fips.tar.gz  containerd-1.6.17-d2iq.1-centos-7.9-
x86_64.tar.gz          containerd-1.6.17-d2iq.1-rhel-8.4-x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-ubuntu-20.04-x86_64.tar.gz
1.26.6_redhat_8_x86_64.tar.gz          containerd-1.6.17-d2iq.1-centos-7.9-
x86_64_fips.tar.gz  containerd-1.6.17-d2iq.1-rhel-8.6-x86_64.tar.gz
images
```

- c. For example, for RHEL 8.4 you would set:

```
export OS_PACKAGES_BUNDLE=1.26.6_redhat_8_x86_64_fips.tar.gz
export CONTAINERD_BUNDLE=containerd-1.6.10-d2iq.1-rhel-8.4-x86_64.tar.gz
```

3. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```
export CONTROL_PLANE_1_ADDRESS="<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<control-plane-address-3>"
export WORKER_1_ADDRESS="<worker-address-1>"
export WORKER_2_ADDRESS="<worker-address-2>"
export WORKER_3_ADDRESS="<worker-address-3>"
export WORKER_4_ADDRESS="<worker-address-4>"
export SSH_USER="<ssh-user>"
export SSH_PRIVATE_KEY_FILE="<private key file>"
```

`SSH_PRIVATE_KEY_FILE` must be either the name of the SSH private key file in your working directory or an absolute path to the file in your user's home directory.

4. Generate an `inventory.yaml` which is automatically picked up by the `konvoy-image upload` in the next step.

```
cat <<EOF > inventory.yaml
all:
  vars:
    ansible_user: $SSH_USER
    ansible_port: 22
    ansible_ssh_private_key_file: $SSH_PRIVATE_KEY_FILE
  hosts:
    $CONTROL_PLANE_1_ADDRESS:
      ansible_host: $CONTROL_PLANE_1_ADDRESS
    $CONTROL_PLANE_2_ADDRESS:
      ansible_host: $CONTROL_PLANE_2_ADDRESS
    $CONTROL_PLANE_3_ADDRESS:
      ansible_host: $CONTROL_PLANE_3_ADDRESS
    $WORKER_1_ADDRESS:
      ansible_host: $WORKER_1_ADDRESS
    $WORKER_2_ADDRESS:
      ansible_host: $WORKER_2_ADDRESS
    $WORKER_3_ADDRESS:
      ansible_host: $WORKER_3_ADDRESS
    $WORKER_4_ADDRESS:
      ansible_host: $WORKER_4_ADDRESS
EOF
```

5. Upload the artifacts onto cluster hosts with the following command:

```
konvoy-image upload artifacts \
  --container-images-dir=./artifacts/images/ \
  --os-packages-bundle=./artifacts/$OS_PACKAGES_BUNDLE \
  --containerd-bundle=artifacts/$CONTAINERD_BUNDLE \
  --pip-packages-bundle=./artifacts/pip-packages.tar.gz
```

KIB uses [variable overrides](#) (see [page 1513](#)) to specify base image and container images to use in your new machine image. The variable overrides files for NVIDIA and FIPS can be ignored unless adding an overlay feature.



- Use the `--overrides` flag and reference either `fips.yaml` or `offline-fips.yaml` manifests located in the [overrides directory](#)²⁰⁷ or see these pages in the documentation:
 - [FIPS Overrides](#) (see [page 1514](#))
 - [Create FIPS 140 Images](#) (see [page 1528](#))


²⁰⁷ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

5.10.5.2.3.1 Next Step:

[Pre-provisioned Air-gapped FIPS: Load the Registry \(see page 220\)](#)

5.10.5.3 Pre-provisioned Air-gapped FIPS: Load the Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine \(see page 1285\)](#) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.

 If you do not already have a local registry set up, please refer to [Local Registry Tools \(see page 1533\)](#) page for more information.

1. [Download \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, and extract the tarball to a local directory:


```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

3. Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```

 It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

5.10.5.3.1 Kommander Load Images

If you are operating in an air-gapped environment, a [local container registry](#) (see page 1533) containing all the necessary installation images, including the Kommander images is required. See below for how to push the necessary images to this registry.

5.10.5.3.2 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.10.5.3.3 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```



To increase [Docker Hub's rate limit](#)²⁰⁸ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

5.10.5.3.4 Next Step:

[Pre-provisioned Air-gapped FIPS: Define Infrastructure](#) (see page 221)

²⁰⁸ <https://docs.docker.com/docker-hub/download-rate-limit/>

5.10.5.4 Pre-provisioned Air-gapped FIPS: Define Infrastructure

Define the cluster hosts and infrastructure

Konvoy needs to know how to access your cluster hosts. This is done using inventory resources. For initial cluster creation, you must define a control-plane and at least one worker pool.

5.10.5.4.1 Define your Infrastructure

1. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```
export CONTROL_PLANE_1_ADDRESS="<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<control-plane-address-3>"
export WORKER_1_ADDRESS="<worker-address-1>"
export WORKER_2_ADDRESS="<worker-address-2>"
export WORKER_3_ADDRESS="<worker-address-3>"
export WORKER_4_ADDRESS="<worker-address-4>"
export SSH_USER="<ssh-user>"
export SSH_PRIVATE_KEY_SECRET_NAME="CLUSTER_NAME-ssh-key"
```

2. Use the following template to help you define your infrastructure. The environment variables that you set in the previous step automatically replace the variable names when the inventory YAML file is created.

```
cat <<EOF > preprovisioned_inventory.yaml
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: CLUSTER_NAME-control-plane
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    # Create as many of these as needed to match your infrastructure
    # Note that the command line parameter --control-plane-replicas determines
    # how many control plane nodes will actually be used.
    #
    - address: CONTROL_PLANE_1_ADDRESS
    - address: CONTROL_PLANE_2_ADDRESS
    - address: CONTROL_PLANE_3_ADDRESS
  sshConfig:
    port: 22
```

```

    # This is the username used to connect to your infrastructure. This user
    must be root or
    # have the ability to use sudo without a password
    user: $SSH_USER
    privateKeyRef:
      # This is the name of the secret you created in the previous step. It
      must exist in the same
      # namespace as this inventory object.
      name: $SSH_PRIVATE_KEY_SECRET_NAME
      namespace: default
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: $CLUSTER_NAME-md-0
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: $CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    - address: $WORKER_1_ADDRESS
    - address: $WORKER_2_ADDRESS
    - address: $WORKER_3_ADDRESS
    - address: $WORKER_4_ADDRESS
  sshConfig:
    port: 22
    user: $SSH_USER
    privateKeyRef:
      name: $SSH_PRIVATE_KEY_SECRET_NAME
      namespace: default
EOF

```

5.10.5.4.2 Next Step:

[Pre-provisioned Air-gapped FIPS: Define Control Plane Endpoint \(see page 222\)](#)

5.10.5.5 Pre-provisioned Air-gapped FIPS: Define Control Plane Endpoint

5.10.5.5.1 Define the Control Plane Endpoint for your cluster as well as connection mechanism


A control plane should have three, five, or seven nodes, so it can remain available if one or more nodes fail. A control plane with one node should not be used in production.

In addition, the control plane should have an endpoint that remains available if some nodes fail.

```

----- cp1.example.com:6443

```



In this example, the control plane endpoint host is `lb.example.com`, and the control plane endpoint port is `6443`. The control plane nodes are `cp1.example.com`, `cp2.example.com`, and `cp3.example.com`. The port of each API server is `6443`.

5.10.5.5.2 Select your Connection Mechanism

A virtual IP is the address that the client uses to connect to the service. A load balancer is the device that distributes the client connections to the backend servers. Before you create a new DKP cluster, choose an external load balancer(LB) or virtual IP.

- **External load balancer**

It is recommended that an external load balancer be the control plane endpoint. To distribute request load among the control plane machines, configure the load balancer to send requests to all the control plane machines. Configure the load balancer to send requests only to control plane machines that are responding to API requests.

- **Built-in virtual IP**

If an external load balancer is not available, use the [built-in virtual IP](#) (see page 1236). The virtual IP is *not* a load balancer; it does not distribute request load among the control plane machines. However, if the machine receiving requests does not respond to them, the virtual IP automatically moves to another machine.

5.10.5.5.3 Single-Node Control Plane



Do not use a single-node control plane in a production cluster.

A control plane with one node can use its single node as the endpoint, so you will not require an external load balancer, or a built-in virtual IP. At least one control plane node must always be running. Therefore, to upgrade a cluster with one control plane node, a spare machine must be available in the control plane inventory. This machine is used to provision the new node before the old node is deleted.



Modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane](#) (see page 1538).

When the API server endpoints are defined, you can create the cluster using the link in Next Step below.

5.10.5.5.4 Known Limitations



Be aware of these limitations in the current release of DKP.

The control plane endpoint port is also used as the API server port on each control plane machine. The default port is 6443. Before you create the cluster, ensure the port is available for use on each control plane machine.

5.10.5.5.5 Next Step:

[Pre-provisioned Air-gapped FIPS: Create a New Cluster](#) (see page 225)

5.10.5.6 Pre-provisioned Air-gapped FIPS: Create a Management Cluster

If your cluster is air-gapped or you have a local registry, you must provide additional arguments when creating the cluster. These tell the cluster where to locate the local registry to use by defining the URL.

```
export REGISTRY_URL=<https/http>://<registry-address>:<registry-port>
export REGISTRY_CA=<path to the CA on the bastion>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

- `REGISTRY_URL` : the address of an existing registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
- `REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. Konvoy will configure the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the AMLs are not already configured to trust this CA.
- `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
- `REGISTRY_PASSWORD` : optional if username is not set.

5.10.5.6.1 Name Your Cluster



The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)²⁰⁹ for more naming information.

²⁰⁹ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.


Follow these steps:

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=<preprovisioned-example>
```

5.10.5.6.2 Create an Air-gapped Kubernetes Cluster


After you have defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster.

 DKP uses local static provisioner as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)²¹⁰ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the inventory.

The create cluster command below includes the `--self-managed` flag. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.

 Before you create a new DKP cluster below, choose an external load balancer (LB) or [virtual IP](#) (see page 1236) and use the corresponding `dkp create cluster` command.

²¹⁰ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

1. The command below uses the default external load balancer (See alternative Step 1 for virtual IP):

```
dkp create cluster preprovisioned --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host <control plane endpoint host> \
  --control-plane-endpoint-port <control plane endpoint port, if different than
6443> \
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml \
  --ssh-private-key-file <path-to-ssh-private-key> \
  --registry-mirror-url=${REGISTRY_URL} \
  --registry-mirror-cacert=${REGISTRY_CA} \
  --registry-mirror-username=${REGISTRY_USERNAME} \
  --registry-mirror-password=${REGISTRY_PASSWORD} \
  --kubernetes-version=v1.26.14+fips.0 \
  --etcd-version=3.5.10+fips.0 \
  --kubernetes-image-repository=docker.io/mesosphere \
  --self-managed
```

- i** If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1039\)](#).

1. **Virtual IP ALTERNATIVE** - if you don't have an external LB, and wish to use a [VIRTUAL IP \(see page 1236\)](#) provided by kube-vip, specify these flags example below:

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host 196.168.1.10 \
  --virtual-ip-interface eth1
```

The output from this command is shortened here for reading clarity, but should start like this:

```
Generating cluster resources
cluster.cluster.x-k8s.io/preprovisioned-example created
cont.....
```

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:

```
cluster.cluster.x-k8s.io/preprovisioned-example condition met
```



NOTE: Depending on the cluster size, it will take a few minutes to create.

When the command completes, you will have a running Kubernetes cluster! For bootstrap and custom YAML cluster creation, refer to the Additional Infrastructure Customization section of the documentation for Pre-provisioned: [Pre-provisioned Infrastructure \(see page 1200\)](#)

Use this command to get the Kubernetes `kubeconfig` for the new cluster and proceed to installing the DKP Kommander UI:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```



If changing the [Calico encapsulation \(see page 1231\)](#), D2iQ recommends changing it after cluster creation, but before production.

5.10.5.6.3 Audit Logs

To modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane](#)²¹¹²¹².

5.10.5.6.4 Further Steps:

For more customized cluster creation, see the [Pre-Provisioned Additional Configurations \(see page 1200\)](#) section.



NOTE: The section mentioned above is for [overrides for your clusters \(see page 1524\)](#), custom flags, and more that specify the secret as part of the create cluster command. If these are not specified, the overrides for your nodes will not be applied.

²¹¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273681652&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane>

²¹² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273681652&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane>

5.10.5.6.5 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPagelId=273320989>²¹³²¹⁴.

5.10.5.6.6 Next Step:

[Pre-provisioned Air-gapped FIPS: Configure MetalLB \(see page 229\)](#)

5.10.5.7 Pre-provisioned Air-gapped FIPS: Configure MetalLB

5.10.5.7.1 Create a MetalLB configmap for your pre-provisioned infrastructure.

Choose one of the following two protocols you want to use to announce service IPs. If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can continue the installation process.

To use MetalLB, create a MetalLB configMap for your Pre-provisioned infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)

Select one of the following procedures to create your MetalLB manifest for further editing.

5.10.5.7.2 Layer 2 configuration

Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:



The following values are generic, enter your specific values into the fields where applicable.

²¹³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

²¹⁴ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    address-pools:
    - name: default
      protocol: layer2
      addresses:
      - 192.168.1.240-192.168.1.250
EOF
```

Once complete, run the following `kubectl` command.


```
kubectl apply -f metallb-conf.yaml
```

5.10.5.7.3 BGP configuration

For a basic configuration featuring one BGP router and one IP address range, you need 4 pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:

 The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
    - peer-address: 10.0.0.1
```

```

peer-asn: 64501
my-asn: 64500
address-pools:
- name: default
  protocol: bgp
  addresses:
  - 192.168.10.0/24
EOF

```

Once complete, run the following `kubectl` command.

```
kubectl apply -f metallb-conf.yaml
```

5.10.5.7.3.1 Next Step:

[Pre-provisioned Air-gapped FIPS: Install Kommander](#) (see page 230)

5.10.5.8 Pre-provisioned Air-gapped FIPS: Install Kommander

5.10.5.8.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Ensure you have loaded all necessary images for your configuration. See [Load the Images into Your Registry: Air-gapped Environments](#) (see page 1400).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.10.5.8.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init --airgapped > kommander.yaml
```

4. Edit the installer file to include configuration overrides for the `rook-ceph-cluster`. DKP's default configuration ships Ceph with [PVC based storage](#)²¹⁵ which requires your CSI provider to support PVC with type `volumeMode: Block`. As this is not possible with the default [local static provisioner](#) (see [page 91](#)), you can install Ceph in [host storage](#)²¹⁶ mode.

You can choose whether Ceph's object storage daemon (osd) pods should consume all or just some of the devices on your nodes. Include **one** of the following Overrides:

- a. To automatically assign all raw storage devices on all nodes to the Ceph cluster:

```
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllDevices: true
        useAllNodes: true
        deviceFilter: "<<value>>"
```

- b. To assign specific storage devices on all nodes to the Ceph cluster:

```
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllNodes: true
        useAllDevices: false
        deviceFilter: "^sdb."
```

Note: If you want to assign specific devices to specific nodes using the `deviceFilter` option, refer to [Specific Nodes and Devices](#)²¹⁷. For general information on the `deviceFilter` value, refer to [Storage Selection Settings](#)²¹⁸.

5. *If required:* Customize your `kommander.yaml`.

i See [Kommander Customizations](#) (see [page 1422](#)) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, etc.


²¹⁵ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/>

²¹⁶ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/>

²¹⁷ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/#specific-nodes-and-devices>


²¹⁸ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

5.10.5.8.3 Enable DKP Catalog Applications and Install Kommander in an Air-gapped Environment

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add the following values to enable DKP Catalog Applications:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      path: ./dkp-catalog-applications-v2.6.1.tar.gz
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubecfg=${CLUSTER_NAME}.conf \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-
v2.6.1.tar.gz
```

 **Tips and recommendations**

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File \(see page 87\)](#).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.10.5.8.4 Next Step:

[Pre-provisioned Air-gapped FIPS: Verify Install and Log in to UI \(see page 233\)](#)

5.10.5.9 Pre-provisioned Air-gapped FIPS: Verify Install and Log in to UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.



NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
```

```

helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met

```

5.10.5.9.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```

kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'

```


5.10.5.9.2 Next Step:

[Pre-provisioned Air-gapped FIPS: Subsequent Clusters \(see page 234\)](#)

5.10.5.10 Pre-provisioned Air-gapped FIPS: Create Managed Clusters Using the DKP CLI


After initial cluster creation, you have the ability to create additional clusters from the CLI. In a previous step, the new cluster was created as Self-managed which allows it to be a Management cluster or a stand alone

cluster. Subsequent new clusters are not self-managed as they will likely be Managed or Attached clusters to this Management Cluster.

 When creating [Managed clusters](#) (see page 81), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see page 81)!

5.10.5.10.1 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:


 **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see page 650).

```
kubectl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.10.5.10.2 Name Your Cluster

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)²¹⁹ for more naming information.

When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:


1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

²¹⁹ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>


```
export CLUSTER_NAME=<preprovisioned-additional>
```

5.10.5.10.3 Create a Kubernetes Cluster

Once you've defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster. This process will create a self-managed cluster to be used as the Management cluster.

 Before you create a new DKP cluster below, choose an external load balancer (LB) or [virtual IP](#) (see page 1236) and use the corresponding `dkp create cluster` command.

In a pre-provisioned environment, use the Kubernetes CSI and third party drivers for local volumes and other storage devices in your data center.

 DKP uses local static provisioner as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)²²⁰ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the [inventory YAML](#) (see page 151) previously created.

1. This command uses the default external load balancer (LB) option:

```
dkp create cluster preprovisioned --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host <control plane endpoint host> \
  --control-plane-endpoint-port <control plane endpoint port, if different than
6443> \
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml \
  --ssh-private-key-file <path-to-ssh-private-key> \
  --registry-mirror-url=${REGISTRY_URL} \
  --registry-mirror-cacert=${REGISTRY_CA} \
  --registry-mirror-username=${REGISTRY_USERNAME} \
  --registry-mirror-password=${REGISTRY_PASSWORD}
```

²²⁰ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:

```
cluster.cluster.x-k8s.io/preprovisioned-additional condition met
```



NOTE: Depending on the cluster size, it will take a few minutes to create.

5.10.5.10.4 Manually Attach a DKP CLI Cluster to the Management Cluster



These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > $
{MANAGED_CLUSTER_NAME}.conf
```

- You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

- Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

- Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

- You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

- This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

- Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

- Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
```

```
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubect1 get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.10.5.10.5 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.10.6 Pre-provisioned with GPU Install

This section provides instructions to install DKP in a Pre-provisioned non-air-gapped environment using GPU.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.10.6.1 Next Step:

[Pre-provisioned GPU: Nodepool Secrets and Overrides](#) (see page 240)

5.10.6.2 Pre-provisioned GPU: Nodepool Secrets and Overrides

5.10.6.2.1 For pre-provisioned environments, DKP has introduced the `nvidia-runfile` flag for Air-gapped Pre-provisioned environments. If the [NVIDIA runfile](#)²²¹ installer has not been downloaded, then retrieve and install the download first by running the following command. The first line in the command below downloads and installs the runfile and the second line places it in the artifacts directory (you must create an `artifacts` directory if it doesn't already exist).

```
curl -O https://download.nvidia.com/XFree86/Linux-x86_64/470.82.01/NVIDIA-Linux-x86_64-470.82.01.run
mv NVIDIA-Linux-x86_64-470.82.01.run artifacts
```

 DKP supported [NVIDIA driver](#)²²² version is 470.x.

1. Create the secret that GPU nodepool would use, this secret is populated from the KIB overrides. In this example we have a file called, `overrides/nvidia.yaml`. It should resemble this:

```
gpu:
  types:
    - nvidia
  build_name_extra: "-nvidia"
```

2. Create a secret on the bootstrap cluster that is populated from the above file. We will name it

```
${CLUSTER_NAME}-user-overrides
```

```
kubectl create secret generic ${CLUSTER_NAME}-user-overrides --from-file=overrides.yaml=overrides/nvidia.yaml
```

3. Create an inventory and nodepool with the instructions below and use the `${CLUSTER_NAME}-user-overrides` secret.

Follow these steps:

²²¹ <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

²²² <https://www.nvidia.com/Download/Find.aspx>

1. Create an inventory object that has the same name as the node pool you're creating, and the details of the pre-provisioned machines that you want to add to it. For example, to create a node pool named `gpu-nodepool` an inventory named `gpu-nodepool` must be present in the same namespace:

```

apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: ${MY_NODEPOOL_NAME}
spec:
  hosts:
    - address: ${IP_OF_NODE}
  sshConfig:
    port: 22
    user: ${SSH_USERNAME}
    privateKeyRef:
      name: ${NAME_OF_SSH_SECRET}
      namespace: ${NAMESPACE_OF_SSH_SECRET}

```


2. (Optional) If your pre-provisioned machines have [overrides](#), you must create a secret that includes all of the overrides you want to provide in one file. Create an [override secret](#) using the instructions detailed on this page.
3. Once the `PreprovisionedInventory` object and overrides are created, create a node pool:

```

dkp create nodepool preprovisioned -c ${MY_CLUSTER_NAME} $
${MY_NODEPOOL_NAME} --override-secret-name ${MY_OVERRIDE_SECRET}

```

- Advanced users can use a combination of the `--dry-run` and `--output=yaml` or `--output-directory=<existing-directory>` flags to get a complete set of node pool objects to modify locally or store in version control.

 For more information regarding this flag or others, please refer to the [dkp create nodepool](#) section of the documentation for either cluster or nodepool and select your provider.

For more information, see:

- [Pre-provisioned Create and Delete Node Pools](#)
- [Pre-provisioned Air-gapped Define Environment](#)

Next Step:

[Pre-provisioned GPU: Define Infrastructure](#) (see page 243)

5.10.6.3 Pre-provisioned GPU: Define Infrastructure

Define the cluster hosts and infrastructure

Konvoy needs to know how to access your cluster hosts. This is done using inventory resources. For initial cluster creation, you must define a control-plane and at least one worker pool.

5.10.6.3.1 Define your Infrastructure

1. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```
export CONTROL_PLANE_1_ADDRESS="<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<control-plane-address-3>"
export WORKER_1_ADDRESS="<worker-address-1>"
export WORKER_2_ADDRESS="<worker-address-2>"
export WORKER_3_ADDRESS="<worker-address-3>"
export WORKER_4_ADDRESS="<worker-address-4>"
export SSH_USER="<ssh-user>"
export SSH_PRIVATE_KEY_SECRET_NAME="<CLUSTER_NAME-ssh-key>"
```

2. Use the following template to help you define your infrastructure. The environment variables that you set in the previous step automatically replace the variable names when the inventory YAML file is created.

```
cat <<EOF > preprovisioned_inventory.yaml
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: <CLUSTER_NAME>-control-plane
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: <CLUSTER_NAME>
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    # Create as many of these as needed to match your infrastructure
    # Note that the command line parameter --control-plane-replicas determines
    # how many control plane nodes will actually be used.
    #
    - address: <CONTROL_PLANE_1_ADDRESS>
    - address: <CONTROL_PLANE_2_ADDRESS>
    - address: <CONTROL_PLANE_3_ADDRESS>
```

```

sshConfig:
  port: 22
  # This is the username used to connect to your infrastructure. This user
  # must be root or
  # have the ability to use sudo without a password
  user: $SSH_USER
  privateKeyRef:
    # This is the name of the secret you created in the previous step. It
    # must exist in the same
    # namespace as this inventory object.
    name: $SSH_PRIVATE_KEY_SECRET_NAME
    namespace: default
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: $CLUSTER_NAME-md-0
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: $CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    - address: $WORKER_1_ADDRESS
    - address: $WORKER_2_ADDRESS
    - address: $WORKER_3_ADDRESS
    - address: $WORKER_4_ADDRESS
  sshConfig:
    port: 22
    user: $SSH_USER
    privateKeyRef:
      name: $SSH_PRIVATE_KEY_SECRET_NAME
      namespace: default
EOF

```

5.10.6.3.2 Next Step:


[Pre-provisioned GPU: Define Control Plane Endpoint](#) (see page 243)

5.10.6.4 Pre-provisioned GPU: Define Control Plane Endpoint

5.10.6.4.1 Define the Control Plane Endpoint for your cluster as well as connection mechanism

A control plane should have three, five, or seven nodes, so it can remain available if one or more nodes fail. A control plane with one node should not be used in production.

In addition, the control plane should have an endpoint that remains available if some nodes fail.



In this example, the control plane endpoint host is `lb.example.com`, and the control plane endpoint port is `6443`. The control plane nodes are `cp1.example.com`, `cp2.example.com`, and `cp3.example.com`. The port of each API server is `6443`.

5.10.6.4.2 Select your Connection Mechanism

A virtual IP is the address that the client uses to connect to the service. A load balancer is the device that distributes the client connections to the backend servers. Before you create a new DKP cluster, choose an external load balancer(LB) or virtual IP.

- **External load balancer**

It is recommended that an external load balancer be the control plane endpoint. To distribute request load among the control plane machines, configure the load balancer to send requests to all the control plane machines. Configure the load balancer to send requests only to control plane machines that are responding to API requests.


- **Built-in virtual IP**

If an external load balancer is not available, use the [built-in virtual IP](#) (see page 1236). The virtual IP is *not* a load balancer; it does not distribute request load among the control plane machines. However, if the machine receiving requests does not respond to them, the virtual IP automatically moves to another machine.

5.10.6.4.3 Single-Node Control Plane

 Do not use a single-node control plane in a production cluster.

A control plane with one node can use its single node as the endpoint, so you will not require an external load balancer, or a built-in virtual IP. At least one control plane node must always be running. Therefore, to upgrade a cluster with one control plane node, a spare machine must be available in the control plane inventory. This machine is used to provision the new node before the old node is deleted.

 Modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane](#) (see page 1538).

When the API server endpoints are defined, you can create the cluster using the link in Next Step below.

5.10.6.4.4 Known Limitations



Be aware of these limitations in the current release of DKP.

The control plane endpoint port is also used as the API server port on each control plane machine. The default port is 6443. Before you create the cluster, ensure the port is available for use on each control plane machine.

5.10.6.4.5 Next Step:

[Pre-provisioned GPU: Create a Cluster](#) (see page 246)

5.10.6.5 Pre-provisioned GPU: Create a Management Cluster

5.10.6.5.1 Create a Cluster with GPU AMI

If a custom AMI was created using Konvoy Image Builder, the custom `ami` id is printed and written to `./manifest.json`.

To use the built `ami` with Konvoy, specify it with the `--ami` flag when calling `cluster create`.

For [GPU Steps in Pre-provisioned](#) (see page 0) section of the documentation to use the overrides/`nvidia.yaml`.

Additional helpful information can be found in the [NVIDIA Device Plug-in](#)²²³ for Kubernetes instructions and the [Installation Guide of Supported Platforms](#)²²⁴.

5.10.6.5.2 Name Your Cluster



The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)²²⁵ for more naming information.

When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

²²³ <https://github.com/NVIDIA/k8s-device-plugin/blob/master/README.md>

²²⁴ <https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/install-guide.html>

²²⁵ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.


Follow these steps:

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:


```
export CLUSTER_NAME=<preprovisioned-example>
```

5.10.6.5.3 Create a Kubernetes Cluster

Once you've defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster. This process will create a self-managed cluster to be used as the Management cluster.

 Before you create a new DKP cluster below, choose an external load balancer (LB) or [virtual IP](#) ([see page 1236](#)) and use the corresponding `dkp create cluster` command.

In a pre-provisioned environment, use the Kubernetes CSI and third party drivers for local volumes and other storage devices in your data center.

 DKP uses local static provisioner as the [default storage provider](#) ([see page 91](#)) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)²²⁶ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the [inventory YAML](#) ([see page 151](#)) previously created.

The create cluster command below includes the `--self-managed` flag. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.

²²⁶ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

1. Execute this command to create a cluster with a GPU AMI using the default external load balancer option:

```
dkp create cluster preprovisioned \
  --cluster-name=${CLUSTER_NAME} \
  --control-plane-endpoint-host <control plane endpoint host> \
  --control-plane-endpoint-port <control plane endpoint port, if different than
6443> \
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml \
  --ssh-private-key-file <path-to-ssh-private-key> \
  --ami <ami> \
  --self-managed
```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

1. **Virtual IP ALTERNATIVE** - if you don't have an external LB, and wish to use a [VIRTUAL IP](#) (see page 1236) provided by kube-vip, specify these flags example below:

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host 196.168.1.10 \
  --virtual-ip-interface eth1
```

The output from this command is shortened here for reading clarity, but should start like this:

```
Generating cluster resources
cluster.cluster.x-k8s.io/preprovisioned-example created
cont.....
```

2. Create the node pool after cluster creation:


```
dkp create nodepool aws -c ${CLUSTER_NAME} \
  --instance-type p2.xlarge \
  --ami-id=${AMI_ID_FROM_KIB} \
  --replicas=1 ${NODEPOOL_NAME} \
  --kubeconfig=${CLUSTER_NAME}.conf
```

3. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:

```
cluster.cluster.x-k8s.io/preprovisioned-example condition met
```


 Depending on the cluster size, it will take a few minutes to create.

When the command completes, you will have a running Kubernetes cluster! For bootstrap and custom YAML cluster creation, refer to the Additional Infrastructure Customization section of the documentation for Pre-provisioned: [Pre-provisioned Infrastructure](#) (see page 1200)

Use this command to get the Kubernetes `kubeconfig` for the new cluster and proceed to installing the DKP Kommander UI:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

 If changing the [Calico encapsulation](#) (see page 1231), D2iQ recommends changing it after cluster creation, but before production.

 To increase [Docker Hub's rate limit](#)²²⁷ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

5.10.6.5.4 Audit logs

To modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane](#) (see page 1538).

²²⁷ <https://docs.docker.com/docker-hub/download-rate-limit/>

5.10.6.5.4.1 Further Steps:

For more customized cluster creation, access the [Pre-Provisioned Additional Configurations](#) (see page 1200) section. That section is for [Pre-Provisioned Override Files](#) (see page 246), custom flags, and more that specify the secret as part of the create cluster command. If these are not specified, the overrides for your nodes will not be applied.

5.10.6.5.4.2 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>²²⁸²²⁹.

5.10.6.5.4.3 Next Step:

[Pre-provisioned GPU: Configure MetalLB](#) (see page 249)

5.10.6.6 Pre-provisioned GPU: Configure MetalLB

5.10.6.6.1 Create a MetalLB configmap for your pre-provisioned infrastructure.

Choose one of the following two protocols you want to use to announce service IPs. If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can continue the installation process.

To use MetalLB, create a MetalLB configMap for your Pre-provisioned infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)

Select one of the following procedures to create your MetalLB manifest for further editing.

5.10.6.6.2 Layer 2 configuration


Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

²²⁸ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

²²⁹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:

 The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    address-pools:
    - name: default
      protocol: layer2
      addresses:
      - 192.168.1.240-192.168.1.250
EOF
```

Once complete, run the following `kubectl` command.


```
kubectl apply -f metallb-conf.yaml
```

5.10.6.6.3 BGP configuration

For a basic configuration featuring one BGP router and one IP address range, you need 4 pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:

 The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
      - peer-address: 10.0.0.1
        peer-asn: 64501
        my-asn: 64500
    address-pools:
      - name: default
        protocol: bgp
        addresses:
          - 192.168.10.0/24
EOF
```

Once complete, run the following `kubectl` command.

```
kubectl apply -f metallb-conf.yaml
```

5.10.6.6.3.1 Next Step:

[Pre-provisioned GPU: Install Kommander \(see page 251\)](#)

5.10.6.7 Pre-provisioned GPU: Install Kommander

5.10.6.7.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install \(see page 121\)](#).
- Ensure you have a [default StorageClass \(see page 1395\)](#).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.10.6.7.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

- Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

- Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init > kommander.yaml
```

- Edit the installer file to include configuration overrides for the `rook-ceph-cluster`. DKP's default configuration ships Ceph with [PVC based storage](#)²³⁰ which requires your CSI provider to support PVC with type `volumeMode: Block`. As this is not possible with the default [local static provisioner](#) (see page 91), you can install Ceph in [host storage](#)²³¹ mode.

You can choose whether Ceph's object storage daemon (osd) pods should consume all or just some of the devices on your nodes. Include **one** of the following Overrides:

- To automatically assign all raw storage devices on all nodes to the Ceph cluster:

```
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllDevices: true
        useAllNodes: true
        deviceFilter: "<<value>>"
```

- To assign specific storage devices on all nodes to the Ceph cluster:

```
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllNodes: true
        useAllDevices: false
        deviceFilter: "^sdb."
```

²³⁰ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/>

²³¹ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/>

Note: If you want to assign specific devices to specific nodes using the `deviceFilter` option, refer to [Specific Nodes and Devices](#)²³². For general information on the `deviceFilter` value, refer to [Storage Selection Settings](#)²³³.

5. *If required:* Customize your `kommander.yaml`.

Note: See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, etc.

5.10.6.7.3 Enable GPU Resources

1. In the same `kommander.yaml` file, enable Nvidia platform services.

```
apps:
  nvidia-gpu-operator:
    enabled: true
```

2. Append the correct Toolkit version based on your OS:

The **NVIDIA Container Toolkit** allows users to run GPU accelerated containers. The toolkit includes a container runtime library and utilities to automatically configure containers to leverage NVIDIA GPU and must be configured correctly according to your base operating system.

Centos 7.9/RHEL 7.9:

If you're using Centos 7.9 or RHEL 7.9 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-centos7
```

RHEL 8.4/8.6 and SLES 15 SP3

If you're using RHEL 8.4/8.6 or SLES 15 SP3 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
```

²³² <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/#specific-nodes-and-devices>

²³³ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>


```
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubi8
```

Ubuntu 18.04 and 20.04

If you're using Ubuntu 18.04 or 20.04 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubuntu20.04
```

5.10.6.7.4 Enable DKP Catalog Applications and Install Kommander

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
```

```
ref:
  tag: v2.6.1
```

⚠️ If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

ⓘ Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File \(see page 87\)](#).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.10.6.7.5 Next Step:

[Pre-provisioned GPU: Verify Install and Log in to UI \(see page 255\)](#)

5.10.6.8 Pre-provisioned GPU: Verify Install and Log in to UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.

ⓘ **NOTE:** If the Kommander installation fails or you wish to reconfigure applications, you can rerun the install command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.10.6.8.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:


```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

5.10.6.8.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{\n"}\n"}Password: {{.data.password|base64decode}}
{{ "\n"}\n"}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{\n"}\n"}or .hostname .ip}}{\n"}\n"}dkp/kommander/
dashboard{{ "\n"}\n"}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.10.6.8.3 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as

attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

For more [GPU information \(see page 968\)](#), see the section for GPU under Day 2 - Cluster Operations Management.

5.10.6.8.4 Next Step:

[Pre-provisioned GPU: Subsequent Cluster Creation \(see page 258\)](#)

5.10.6.9 Pre-provisioned GPU: Create Managed Clusters Using the DKP CLI

After initial cluster creation, you have the ability to create additional clusters from the CLI. In a previous step, the new cluster was created as Self-managed which allows it to be a Management cluster or a stand alone cluster. Subsequent new clusters are not self-managed as they will likely be Managed or Attached clusters to this Management Cluster.



When creating [Managed clusters \(see page 81\)](#), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters \(see page 81\)](#)!

5.10.6.9.1 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:

⚠ NOTE: If you need to create a new Workspace, follow the instructions to [Create a New Workspace \(see page 650\)](#).

```
kubect1 get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.10.6.9.2 Name Your Cluster

ⓘ The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes²³⁴](#) for more naming information.

When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=<preprovisioned-additional>
```

If a custom AMI was created using Konvoy Image Builder, the custom `ami` id is printed and written to `./manifest.json`. To use the built `ami` with Konvoy, specify it with the `--ami` flag when calling cluster create command below.


5.10.6.9.3 Create a Kubernetes Cluster

Once you've defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster. This process will create a self-managed cluster to be used as the Management cluster.

ⓘ Before you create a new DKP cluster below, choose an external load balancer (LB) or [virtual IP \(see page 1236\)](#) and use the corresponding `dkp create cluster` command.

In a pre-provisioned environment, use the Kubernetes CSI and third party drivers for local volumes and other storage devices in your data center.

²³⁴ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

 DKP uses local static provisioner as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)²³⁵ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the [inventory YAML](#) (see page 151) previously created.

1. This command uses the default external load balancer (LB) option:

```
dkp create cluster preprovisioned
  --cluster-name ${CLUSTER_NAME}
  --control-plane-endpoint-host <control plane endpoint host>
  --control-plane-endpoint-port <control plane endpoint port, if different than
  6443>
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml
  --ssh-private-key-file <path-to-ssh-private-key>
  --ami <ami>
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
  timeout=30m
```

Output:


```
cluster.cluster.x-k8s.io/preprovisioned-additional condition met
```



NOTE: Depending on the cluster size, it will take a few minutes to create.

²³⁵ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

5.10.6.9.4 Manually Attach a DKP CLI Cluster to the Management Cluster

 These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

7. This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```

apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>

```

8. Create this secret in the desired workspace:

```

kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}

```

9. Create this `kommandercluster` object to attach the cluster to the workspace:

```

cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF

```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```

kubectl get kommanderclusters -A

```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.10.6.9.5 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.10.7 Pre-provisioned Air-gapped with GPU Install

This section provides instructions to install DKP in a Pre-provisioned air-gapped environment using GPU.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

NOTE: For air-gapped, ensure you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, so you can extract the tarball on next page.

5.10.7.1 Next Step:

[Pre-provisioned Air-gapped GPU: Configure Environment](#) (see page 263)

5.10.7.2 Pre-provisioned Air-gapped GPU: Configure Environment

If the [NVIDIA runfile](#)²³⁶ installer has not been downloaded, then retrieve and install the download first by running the following command. The first line in the command below downloads and installs the runfile and the second line places it in the artifacts directory.

```
curl -O https://download.nvidia.com/XFree86/Linux-x86_64/470.82.01/
NVIDIA-Linux-x86_64-470.82.01.run
mv NVIDIA-Linux-x86_64-470.82.01.run artifacts
```

The instructions below outline how to fulfill the requirements for using pre-provisioned infrastructure when using air-gapped. In order to create a cluster, you must first setup the environment with necessary artifacts. All artifacts for Pre-provisioned Air-gapped need to get onto the bastion host. Artifacts needed by nodes must be unpacked and distributed on the bastion before other provisioning will work in the absence of an internet connection.

There is a new complete DKP air-gapped bundle available to [download](#) (see page 78) which contains all the DKP components needed for air-gapped installation. (i.e. `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`)

²³⁶ <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

5.10.7.2.1 Setup Process:

1. The bootstrap image must be extracted and loaded onto the [bastion host](#) (see page 1285).
2. Artifacts must be copied onto cluster hosts for nodes to access.
3. If using GPU, those artifacts must be positioned locally.
4. Registry seeded with images locally.

5.10.7.2.2 Load the Bootstrap Image

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz` from the [download site](#) (see page 78) mentioned above, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Load the bootstrap image on your bastion machine:

```
docker load -i konvoy-bootstrap-image-v2.6.2.tar
```

5.10.7.2.3 Copy air-gapped artifacts onto cluster hosts

Using the [Konvoy Image Builder](#) (see page 1461), you can copy the required artifacts onto your cluster hosts.

Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, continue below:

1. The Kubernetes image bundle will be located in `kib/artifacts/images` and you will want to verify image and artifacts.
 - a. Verify the image bundles exist in `artifacts/images` :

```
$ ls artifacts/images/
kubernetes-images-1.26.14-d2iq.1.tar kubernetes-images-1.26.14-d2iq.1-
fips.tar
```

- b. Verify the artifacts for your OS exist in the `artifacts/` directory and export the appropriate variables:

```
$ ls artifacts/
1.26.14-centos_7_x86_64.tar.gz      1.26.14-redhat_8_x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rhel-8.6-x86_64_fips.tar.gz  pip-packages.tar.gz
```



```

1.26.14_centos_7_x86_64_fips.tar.gz 1.26.14_rocky_9_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64_fips.tar.gz containerd-1.6.17-
d2iq.1-rocky-9.0-x86_64.tar.gz
1.26.14_redhat_7_x86_64.tar.gz 1.26.14_ubuntu_20_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz containerd-1.6.17-
d2iq.1-rocky-9.1-x86_64.tar.gz
1.26.14_redhat_7_x86_64_fips.tar.gz containerd-1.6.17-d2iq.1-centos-7.9-
x86_64.tar.gz containerd-1.6.17-d2iq.1-rhel-8.4-x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-ubuntu-20.04-x86_64.tar.gz NVIDIA-Linux-
x86_64-470.82.01.run
1.26.14_redhat_8_x86_64.tar.gz containerd-1.6.17-d2iq.1-centos-7.9-
x86_64_fips.tar.gz containerd-1.6.17-d2iq.1-rhel-8.6-x86_64.tar.gz
images

```

- c. For example, for RHEL 8.4 you would set:

```

export OS_PACKAGES_BUNDLE=1.26.14_redhat_8_x86_64.tar.gz
export CONTAINERD_BUNDLE=containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz

```

2. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```

export CONTROL_PLANE_1_ADDRESS="<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<control-plane-address-3>"
export WORKER_1_ADDRESS="<worker-address-1>"
export WORKER_2_ADDRESS="<worker-address-2>"
export WORKER_3_ADDRESS="<worker-address-3>"
export WORKER_4_ADDRESS="<worker-address-4>"
export SSH_USER="<ssh-user>"
export SSH_PRIVATE_KEY_FILE="<private key file>"

```

`SSH_PRIVATE_KEY_FILE` must be either the name of the SSH private key file in your working directory or an absolute path to the file in your user's home directory.

3. Generate an `inventory.yaml` for GPU nodes which is automatically picked up by the `konvoy-image upload` in the next step.

```

cat <<EOF > gpu_inventory.yaml
all:
  vars:
    ansible_port: 22
    ansible_ssh_private_key_file: $SSH_PRIVATE_KEY_FILE
    ansible_user: $SSH_USER

  hosts:
    $GPU_WORKER_1_ADDRESS:
      ansible_host: $GPU_WORKER_1_ADDRESS
EOF

```

```

cat <<EOF > inventory.yaml
all:
  vars:
    ansible_user: $SSH_USER
    ansible_port: 22
    ansible_ssh_private_key_file: $SSH_PRIVATE_KEY_FILE
  hosts:
    $CONTROL_PLANE_1_ADDRESS:
      ansible_host: $CONTROL_PLANE_1_ADDRESS
    $CONTROL_PLANE_2_ADDRESS:
      ansible_host: $CONTROL_PLANE_2_ADDRESS
    $CONTROL_PLANE_3_ADDRESS:
      ansible_host: $CONTROL_PLANE_3_ADDRESS
    $WORKER_1_ADDRESS:
      ansible_host: $WORKER_1_ADDRESS
    $WORKER_2_ADDRESS:
      ansible_host: $WORKER_2_ADDRESS
    $WORKER_3_ADDRESS:
      ansible_host: $WORKER_3_ADDRESS
    $WORKER_4_ADDRESS:
      ansible_host: $WORKER_4_ADDRESS
EOF

```

4. Upload the artifacts to the gpu nodepool with the `nvidia-runfile` flag with the following command:

```

konvoy-image upload artifacts --inventory-file=gpu_inventory.yaml \
  --container-images-dir=./artifacts/images/ \
  --os-packages-bundle=./artifacts/$OS_PACKAGES_BUNDLE \
  --containerd-bundle=artifacts/$CONTAINERD_BUNDLE \
  --pip-packages-bundle=./artifacts/pip-packages.tar.gz \
  --nvidia-runfile=./artifacts/NVIDIA-Linux-x86_64-470.82.01.run

```

KIB uses [variable overrides](#) (see page 1513) to specify base image and container images to use in your new machine image. The variable overrides files for NVIDIA and FIPS can be ignored unless adding an overlay feature.



- Use the `--overrides overrides/fips.yaml,overrides/offline-fips.yaml` flag with manifests located in the [overrides directory](#)²³⁷ or see these pages in the documentation:
 - [FIPS Overrides](#) (see page 1514)

²³⁷ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>


- [Create FIPS 140 Images \(see page 1528\)](#)

5.10.7.2.3.1 Next Step:

[Pre-provisioned Air-gapped GPU: Load the Registry \(see page 266\)](#)

5.10.7.3 Pre-provisioned Air-gapped GPU: Load the Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine \(see page 1285\)](#) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.

 If you do not already have a local registry set up, please refer to [Local Registry Tools \(see page 1533\)](#) page for more information.

1. [Download \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, and extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:


```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

3. Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```



It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

 To increase [Docker Hub's rate limit](#)²³⁸ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

5.10.7.3.1 Kommander Load Images

If you are operating in an air-gapped environment, a [local container registry](#) (see page 1533) containing all the necessary installation images, including the Kommander images is required. See below for how to push the necessary images to this registry.

5.10.7.3.2 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.10.7.3.3 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.10.7.3.4 Next Step:

[Pre-provisioned Air-gapped GPU: Nodepool Secrets and Overrides](#) (see page 268)

²³⁸ <https://docs.docker.com/docker-hub/download-rate-limit/>

5.10.7.4 Pre-provisioned Air-gapped GPU: Nodepool Secrets and Overrides

For pre-provisioned environments, DKP has introduced the `nvidia-runfile` flag for Air-gapped Pre-provisioned environments. If the [NVIDIA runfile](#)²³⁹ installer has not been downloaded, then retrieve and install the download first by running the following command. The first line in the command below downloads and installs the runfile and the second line places it in the artifacts directory (you must create an `artifacts` directory if it doesn't already exist).

```
curl -O https://download.nvidia.com/XFree86/Linux-x86_64/470.82.01/NVIDIA-Linux-x86_64-470.82.01.run
mv NVIDIA-Linux-x86_64-470.82.01.run artifacts
```

 DKP supported [NVIDIA driver](#)²⁴⁰ version is 470.x.

1. Create the secret that GPU nodepool would use, this secret is populated from the KIB overrides. In this example we have a file called, `overrides/nvidia.yaml`. It should resemble this:

```
gpu:
  types:
    - nvidia
  build_name_extra: "-nvidia"
```

2. Create a secret on the bootstrap cluster that is populated from the above file. We will name it

`${CLUSTER_NAME}-user-overrides`

```
kubectl create secret generic ${CLUSTER_NAME}-user-overrides --from-file=overrides.yaml=overrides/nvidia.yaml
```

3. Create an inventory and nodepool with the instructions below and use the `${CLUSTER_NAME}-user-overrides` secret.

Follow these steps:

1. Create an inventory object that has the same name as the node pool you're creating, and the details of the pre-provisioned machines that you want to add to it. For example, to create a node pool named `gpu-nodepool` an inventory named `gpu-nodepool` must be present in the same namespace:

²³⁹ <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

²⁴⁰ <https://www.nvidia.com/Download/Find.aspx>

```

apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: ${MY_NODEPOOL_NAME}
spec:
  hosts:
    - address: ${IP_OF_NODE}
  sshConfig:
    port: 22
    user: ${SSH_USERNAME}
    privateKeyRef:
      name: ${NAME_OF_SSH_SECRET}
      namespace: ${NAMESPACE_OF_SSH_SECRET}

```

- (Optional) If your pre-provisioned machines have [overrides](#) (see [page 1513](#)), you must create a secret that includes all of the overrides you want to provide in one file. Create an [override secret](#) (see [page 1213](#)) using the instructions detailed on this page.
- Once the `PreprovisionedInventory` object and overrides are created, create a node pool:

```

dkp create nodepool preprovisioned -c ${MY_CLUSTER_NAME} ${MY_NODEPOOL_NAME} --
override-secret-name ${MY_OVERRIDE_SECRET}

```

- Advanced users can use a combination of the `--dry-run` and `--output=yaml` or `--output-directory=<existing-directory>` flags to get a complete set of node pool objects to modify locally or store in version control.



For more information regarding this flag or others, please refer to the [dkp create nodepool](#) (see [page 1684](#)) section of the documentation for either cluster or nodepool and select your provider.

For more information, see:

- [Pre-provisioned Create and Delete Node Pools](#) (see [page 1245](#))
- [Pre-provisioned Air-gapped Define Environment](#) (see [page 1204](#))

5.10.7.4.1 Next Step:

[Pre-provisioned Air-gapped GPU: Define Infrastructure](#) (see [page 270](#))

5.10.7.5 Pre-provisioned Air-gapped GPU: Define Infrastructure

Define the cluster hosts and infrastructure

Konvoy needs to know how to access your cluster hosts. This is done using inventory resources. For initial cluster creation, you must define a control-plane and at least one worker pool.

5.10.7.5.1 Define your Infrastructure

1. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```
export CONTROL_PLANE_1_ADDRESS="<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<control-plane-address-3>"
export WORKER_1_ADDRESS="<worker-address-1>"
export WORKER_2_ADDRESS="<worker-address-2>"
export WORKER_3_ADDRESS="<worker-address-3>"
export WORKER_4_ADDRESS="<worker-address-4>"
export SSH_USER="<ssh-user>"
export SSH_PRIVATE_KEY_SECRET_NAME="CLUSTER_NAME-ssh-key"
```

2. Use the following template to help you define your infrastructure. The environment variables that you set in the previous step automatically replace the variable names when the inventory YAML file is created.

```
cat <<EOF > preprovisioned_inventory.yaml
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: CLUSTER_NAME-control-plane
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    # Create as many of these as needed to match your infrastructure
    # Note that the command line parameter --control-plane-replicas determines
    # how many control plane nodes will actually be used.
    #
    - address: CONTROL_PLANE_1_ADDRESS
    - address: CONTROL_PLANE_2_ADDRESS
    - address: CONTROL_PLANE_3_ADDRESS
  sshConfig:
    port: 22
    # This is the username used to connect to your infrastructure. This user
    # must be root or
    # have the ability to use sudo without a password
    user: SSH_USER
    privateKeyRef:
```

```

    # This is the name of the secret you created in the previous step. It
    must exist in the same
    # namespace as this inventory object.
    name: $SSH_PRIVATE_KEY_SECRET_NAME
    namespace: default
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: $CLUSTER_NAME-md-0
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: $CLUSTER_NAME
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    - address: $WORKER_1_ADDRESS
    - address: $WORKER_2_ADDRESS
    - address: $WORKER_3_ADDRESS
    - address: $WORKER_4_ADDRESS
  sshConfig:
    port: 22
    user: $SSH_USER
    privateKeyRef:
      name: $SSH_PRIVATE_KEY_SECRET_NAME
      namespace: default
EOF

```

5.10.7.5.2 Next Step:

[Pre-provisioned Air-gapped GPU: Define Control Plane Endpoint \(see page 272\)](#)

5.10.7.6 Pre-provisioned Air-gapped GPU: Define Control Plane Endpoint

5.10.7.6.1 Define the Control Plane Endpoint for your cluster as well as connection mechanism

A control plane should have three, five, or seven nodes, so it can remain available if one or more nodes fail. A control plane with one node should not be used in production.

In addition, the control plane should have an endpoint that remains available if some nodes fail.

```

          ----- cp1.example.com:6443
          |
lb.example.com:6443 ----- cp2.example.com:6443
          |
          ----- cp3.example.com:6443

```


In this example, the control plane endpoint host is `lb.example.com`, and the control plane endpoint port is `6443`. The control plane nodes are `cp1.example.com`, `cp2.example.com`, and `cp3.example.com`. The port of each API server is `6443`.

5.10.7.6.2 Select your Connection Mechanism

A virtual IP is the address that the client uses to connect to the service. A load balancer is the device that distributes the client connections to the backend servers. Before you create a new DKP cluster, choose an external load balancer(LB) or virtual IP.

- **External load balancer**

It is recommended that an external load balancer be the control plane endpoint. To distribute request load among the control plane machines, configure the load balancer to send requests to all the control plane machines. Configure the load balancer to send requests only to control plane machines that are responding to API requests.

- **Built-in virtual IP**

If an external load balancer is not available, use the [built-in virtual IP \(see page 1236\)](#). The virtual IP is *not* a load balancer; it does not distribute request load among the control plane machines. However, if the machine receiving requests does not respond to them, the virtual IP automatically moves to another machine.

5.10.7.6.3 Single-Node Control Plane



Do not use a single-node control plane in a production cluster.

A control plane with one node can use its single node as the endpoint, so you will not require an external load balancer, or a built-in virtual IP. At least one control plane node must always be running. Therefore, to upgrade a cluster with one control plane node, a spare machine must be available in the control plane inventory. This machine is used to provision the new node before the old node is deleted.



Modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane \(see page 1538\)](#).

When the API server endpoints are defined, you can create the cluster using the link in Next Step below.

5.10.7.6.4 Known Limitations



Be aware of these limitations in the current release of DKP.

The control plane endpoint port is also used as the API server port on each control plane machine. The default port is 6443. Before you create the cluster, ensure the port is available for use on each control plane machine.

5.10.7.6.5 Next Step:

[Pre-provisioned Air-gapped GPU: Create a Cluster \(see page 274\)](#)

5.10.7.7 Pre-provisioned Air-gapped GPU: Create a Management Cluster

If your cluster is air-gapped or you have a local registry, you must provide additional arguments when creating the cluster. These tell the cluster where to locate the local registry to use by defining the URL.


```
export REGISTRY_URL=<https/http>://<registry-address>:<registry-port>
export REGISTRY_CA=<path to the CA on the bastion>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

- `REGISTRY_URL` : the address of an existing registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
- `REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. Konvoy will configure the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
- `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
- `REGISTRY_PASSWORD` : optional if username is not set.

5.10.7.7.1 Create a Cluster with GPU AMI


If a custom AMI was created using Konvoy Image Builder, the custom `ami` id is printed and written to `./manifest.json`.

To use the built `ami` with Konvoy, specify it with the `--ami` flag when calling `cluster create` in Step 1.


 For [GPU Steps in Pre-provisioned](#) (see page 0) section of the documentation to use the overrides/`nvidia.yaml`.

For [GPU Steps in Pre-provisioned](#) (see page 0) section of the documentation to use the overrides/`nvidia.yaml`.

Additional helpful information can be found in the [NVIDIA Device Plug-in](#)²⁴¹ for Kubernetes instructions and the [Installation Guide of Supported Platforms](#)²⁴².

 Before you create a new DKP cluster below, choose an external load balancer or [virtual IP](#) (see page 1236) and use the corresponding `dkp create cluster` command.

5.10.7.7.2 Name Your Cluster

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)²⁴³ for more naming information.

When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=<preprovisioned-example>
```


²⁴¹ <https://github.com/NVIDIA/k8s-device-plugin/blob/master/README.md>

²⁴² <https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/install-guide.html>

²⁴³ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

5.10.7.3 Create an Air-gapped Kubernetes Cluster

After you have defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster.

 DKP uses local static provisioner as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)²⁴⁴ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the inventory.

The create cluster command below includes the `--self-managed` flag. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.

1. Execute this command to create a cluster with a GPU AMI using the default external load balancer option:

```
dkp create cluster preprovisioned \
  --cluster-name=${CLUSTER_NAME} \
  --control-plane-endpoint-host <control plane endpoint host> \
  --control-plane-endpoint-port <control plane endpoint port, if different than
6443> \
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml \
  --ssh-private-key-file <path-to-ssh-private-key> \
  --registry-mirror-url=${REGISTRY_URL} \
  --registry-mirror-cacert=${REGISTRY_CA} \
  --registry-mirror-username=${REGISTRY_USERNAME} \
  --registry-mirror-password=${REGISTRY_PASSWORD} \
  --ami <ami> \
  --self-managed
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

²⁴⁴ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

2. **Virtual IP ALTERNATIVE** - if you don't have an external LB, and wish to use a [VIRTUAL IP](#) (see page 1236) provided by kube-vip, specify these flags example below::

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host 196.168.1.10 \
  --virtual-ip-interface eth1
```

The output from this command is shortened here for reading clarity, but should start like this:

```
Generating cluster resources
cluster.cluster.x-k8s.io/preprovisioned-example created
cont.....
```

3. Create the node pool after cluster creation:

```
dkp create nodepool aws -c ${CLUSTER_NAME} \
  --instance-type p2.xlarge \
  --ami-id=${AMI_ID_FROM_KIB} \
  --replicas=1 ${NODEPOOL_NAME} \
  --kubeconfig=${CLUSTER_NAME}.conf
```

4. Use the wait command to monitor the cluster control-plane readiness: Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:

```
cluster.cluster.x-k8s.io/preprovisioned-example condition met
```




NOTE: Depending on the cluster size, it will take a few minutes to create.

When the command completes, you will have a running Kubernetes cluster! For bootstrap and custom YAML cluster creation, refer to the [Custom Installation and Additional Infrastructure Tools](#) (see page 1036) section of the documentation for Pre-provisioned: [Pre-provisioned Infrastructure](#) (see page 1200)

Use this command to get the Kubernetes `kubeconfig` for the new cluster and proceed to installing the DKP Kommander UI:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

-  If changing the [Calico encapsulation](#) (see page 1231), D2iQ recommends changing it after cluster creation, but before production.

5.10.7.7.4 Audit logs

To modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane](#) (see page 1538).

5.10.7.7.4.1 Further Steps:

- [Modify the Calico installation](#)²⁴⁵
- [Use the built-in Virtual IP](#)²⁴⁶
- [Provision on the Flatcar Linux OS](#)²⁴⁷
- [Use an HTTP proxy](#)²⁴⁸
- [Use alternate pod or service subnets](#)²⁴⁹

1. **Virtual IP ALTERNATIVE** - if you don't have an external LB, and wish to use a [VIRTUAL IP](#) (see page 1236) provided by kube-vip, specify these flags example below:

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host 196.168.1.10 \
  --virtual-ip-interface eth1
```

The output from this command is shortened here for reading clarity, but should start like this:

```
Generating cluster resources
cluster.cluster.x-k8s.io/preprovisioned-example created
cont.....
```

²⁴⁵ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/edit-v2/165347706#>

²⁴⁶ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/edit-v2/165347706#>

²⁴⁷ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/edit-v2/165347706#>

²⁴⁸ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/edit-v2/165347706#>

²⁴⁹ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/edit-v2/165347706#>

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:

```
cluster.cluster.x-k8s.io/preprovisioned-example condition met
```



NOTE: Depending on the cluster size, it will take a few minutes to create.

When the command completes, you will have a running Kubernetes cluster! For bootstrap and custom YAML cluster creation, refer to the Additional Infrastructure Customization section of the documentation for Pre-provisioned: [Pre-provisioned Infrastructure \(see page 1200\)](#)

Use this command to get the Kubernetes `kubeconfig` for the new cluster and proceed to installing the DKP Kommander UI:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```



If changing the [Calico encapsulation \(see page 1231\)](#), D2iQ recommends changing it after cluster creation, but before production.

5.10.7.7.4.2 Audit Logs

To modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane](#)²⁵⁰²⁵¹.

5.10.7.7.4.3 Further Steps:

For more customized cluster creation, see the [Pre-Provisioned Additional Configurations \(see page 1200\)](#) section.

²⁵⁰ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273681652&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane>

²⁵¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273681652&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane>

NOTE: The section mentioned above is for [overrides for your clusters \(see page 1524\)](#), custom flags, and more that specify the secret as part of the create cluster command. If these are not specified, the overrides for your nodes will not be applied.

5.10.7.7.4.4 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>²⁵²²⁵³.

5.10.7.7.4.5 Next Step:

[Pre-provisioned Air-gapped GPU: Configure MetalLB \(see page 280\)](#)

5.10.7.8 Pre-provisioned Air-gapped GPU: Configure MetalLB

5.10.7.8.1 Create a MetalLB configmap for your pre-provisioned infrastructure.

Choose one of the following two protocols you want to use to announce service IPs. If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can continue the installation process.

To use MetalLB, create a MetalLB configMap for your Pre-provisioned infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)

Select one of the following procedures to create your MetalLB manifest for further editing.

5.10.7.8.2 Layer 2 configuration


Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.

²⁵² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

²⁵³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:

 The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    address-pools:
    - name: default
      protocol: layer2
      addresses:
      - 192.168.1.240-192.168.1.250
EOF
```

Once complete, run the following `kubectl` command.


```
kubectl apply -f metallb-conf.yaml
```

5.10.7.8.3 BGP configuration

For a basic configuration featuring one BGP router and one IP address range, you need 4 pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:

 The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
```

```
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
      - peer-address: 10.0.0.1
        peer-asn: 64501
        my-asn: 64500
    address-pools:
      - name: default
        protocol: bgp
        addresses:
          - 192.168.10.0/24
EOF
```

Once complete, run the following `kubectl` command.

```
kubectl apply -f metallb-conf.yaml
```

5.10.7.8.3.1 Next Step:

[Pre-provisioned Air-gapped GPU: Install Kommander \(see page 282\)](#)

5.10.7.9 Pre-provisioned Air-gapped GPU: Install Kommander

5.10.7.9.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install \(see page 121\)](#).
- Ensure you have a [default StorageClass \(see page 1395\)](#).
- Ensure you have loaded all necessary images for your configuration. See [Load the Images into Your Registry: Air-gapped Environments \(see page 1400\)](#).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.10.7.9.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

- Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

- Create a [configuration file](#) (see [page 1422](#)) for the deployment:

```
dkp install kommander --init --airgapped > kommander.yaml
```

- Edit the installer file to include configuration overrides for the `rook-ceph-cluster`. DKP's default configuration ships Ceph with [PVC based storage](#)²⁵⁴ which requires your CSI provider to support PVC with type `volumeMode: Block`. As this is not possible with the default [local static provisioner](#) (see [page 91](#)), you can install Ceph in [host storage](#)²⁵⁵ mode.

You can choose whether Ceph's object storage daemon (osd) pods should consume all or just some of the devices on your nodes. Include **one** of the following Overrides:

- To automatically assign all raw storage devices on all nodes to the Ceph cluster:

```
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllDevices: true
        useAllNodes: true
        deviceFilter: "<<value>>"
```

- To assign specific storage devices on all nodes to the Ceph cluster:

```
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllNodes: true
        useAllDevices: false
        deviceFilter: "^sdb."
```

²⁵⁴ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/>

²⁵⁵ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/>

Note: If you want to assign specific devices to specific nodes using the `deviceFilter` option, refer to [Specific Nodes and Devices](#)²⁵⁶. For general information on the `deviceFilter` value, refer to [Storage Selection Settings](#)²⁵⁷.

5. *If required:* Customize your `kommander.yaml`.

Note: See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, etc.

5.10.7.9.3 Enable GPU Resources

1. In the same `kommander.yaml` file, enable Nvidia platform services.

```
apps:
  nvidia-gpu-operator:
    enabled: true
```

2. Append the correct Toolkit version based on your OS:

The **NVIDIA Container Toolkit** allows users to run GPU accelerated containers. The toolkit includes a container runtime library and utilities to automatically configure containers to leverage NVIDIA GPU and must be configured correctly according to your base operating system.

Centos 7.9/RHEL 7.9:

If you're using Centos 7.9 or RHEL 7.9 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-centos7
```

RHEL 8.4/8.6 and SLES 15 SP3

If you're using RHEL 8.4/8.6 or SLES 15 SP3 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
```

²⁵⁶ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/#specific-nodes-and-devices>

²⁵⁷ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

```
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubi8
```

Ubuntu 18.04 and 20.04

If you're using Ubuntu 18.04 or 20.04 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubuntu20.04
```

5.10.7.9.4 Enable DKP Catalog Applications and Install Kommander in an Air-gapped Environment



If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add the following values to enable DKP Catalog Applications:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
```

```
path: ./dkp-catalog-applications-v2.6.1.tar.gz
```

⚠️ If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf \
--kommander-applications-repository ./application-repositories/kommander-applications-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-v2.6.1.tar.gz
```

Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File \(see page 87\)](#).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.10.7.9.5 Next Step:

[Pre-provisioned Air-gapped GPU: Verify Install and Log in to UI \(see page 286\)](#)

5.10.7.10 Pre-provisioned Air-gapped GPU: Verify Install and Log in to UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.

NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the install command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.10.7.10.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelMReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelMRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{\n}Password: {{.data.password|base64decode}}
{\n}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/
dashboard{\n}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.10.7.10.2 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see page 582) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

For more [GPU information](#) (see page 968), see the section for GPU under Day 2 - Cluster Operations Management.

5.10.7.10.3 Next Step:

[Pre-provisioned Air-gapped with GPU: Subsequent Clusters](#) (see page 289)

5.10.7.11 Pre-provisioned Air-gapped with GPU: Create Managed Clusters Using the DKP CLI

After initial cluster creation, you have the ability to create additional clusters from the CLI. In a previous step, the new cluster was created as Self-managed which allows it to be a Management cluster or a stand alone cluster. Subsequent new clusters are not self-managed as they will likely be Managed or Attached clusters to this Management Cluster.



When creating [Managed clusters](#) (see page 81), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see page 81)!

5.10.7.11.1 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:

⚠ NOTE: If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see page 650).

```
kubectl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.10.7.11.2 Name Your Cluster

ⓘ The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)²⁵⁸ for more naming information.

When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=<preprovisioned-additional>
```


5.10.7.11.3 Create a Cluster with GPU AMI

If a custom AMI was created using Konvoy Image Builder, the custom `ami` id is printed and written to `./manifest.json`. To use the built `ami` with Konvoy, specify it with the `--ami` flag when calling `cluster create` in Step 1 in the next section where you create your Kubernetes cluster.


5.10.7.11.4 Create a Kubernetes Cluster

Once you've defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster. This process will create a self-managed cluster to be used as the Management cluster.

²⁵⁸ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

 Before you create a new DKP cluster below, choose an external load balancer (LB) or [virtual IP](#) (see page 1236) and use the corresponding `dkp create cluster` command.

In a pre-provisioned environment, use the Kubernetes CSI and third party drivers for local volumes and other storage devices in your data center.

 DKP uses local static provisioner as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)²⁵⁹ compatible storage that is suitable for production.

After disabling `localvolume-provisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the [inventory YAML](#) (see page 151) previously created.

1. This command uses the default external load balancer (LB) option:

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host <control plane endpoint host> \
  --control-plane-endpoint-port <control plane endpoint port, if different than
6443> \
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml \
  --ssh-private-key-file <path-to-ssh-private-key> \
  --registry-mirror-url=${_REGISTRY_URL} \
  --registry-mirror-cacert=${_REGISTRY_CA} \
  --registry-mirror-username=${_REGISTRY_USERNAME} \
  --registry-mirror-password=${_REGISTRY_PASSWORD}
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

2. Use the wait command to monitor the cluster control-plane readiness:

²⁵⁹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:

```
cluster.cluster.x-k8s.io/preprovisioned-additional condition met
```



NOTE: Depending on the cluster size, it will take a few minutes to create.

5.10.7.11.5 Manually Attach a DKP CLI Cluster to the Management Cluster



These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > $
{MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

7. This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

8. Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

9. Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
```

```
EOF
```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubect1 get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.10.7.11.6 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.11 AWS Install Options

For an environment that is on the AWS Infrastructure, install options are provided for you in this one location.

Remember, there are always more options for custom YAML in the [Custom Installation and Additional Infrastructure Tools](#) (see page 1048) section, but this will get you operative in the most common scenarios.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

Below are all the supported environment combinations: (Refer to this page to check if your OS is a [Supported Operating System](#) (see page 69).)

- [AWS Install](#) (see page 294)
- [AWS Air-gapped Install](#) (see page 310)
- [AWS with FIPS Install](#) (see page 328)
- [AWS Air-gapped FIPS Install](#) (see page 344)
- [AWS with GPU Install](#) (see page 363)
- [AWS Air-gapped with GPU Install](#) (see page 382)



Additional Resource Information specific to AWS is below.

- **Control Plane Nodes** - DKP on AWS defaults to deploying an `m5.xlarge` instance with an 80GiB root volume for control plane nodes, which meets the above [resource requirements](#) (see page 100).
- **Worker Nodes** - DKP on AWS defaults to deploying a `m5.2xlarge` instance with an 80GiB root volume for worker nodes, which meets the above [resource requirements](#) (see page 100).

5.11.1 AWS Install

This section provides instructions to install DKP in an AWS non-air-gapped environment.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.11.1.1 AWS Prerequisites

Before you begin using Konvoy with AWS, you must:

1. Follow the steps to create permissions and roles on the [Minimal Permissions and Role to Create Clusters](#) (see page 1057) page.
2. Create [Cluster IAM Policies and Roles](#) (see page 1063).
3. Export the AWS region where you want to deploy the cluster:

```
export AWS_REGION=us-west-2
```

4. Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```



If using AWS ECR as your local private registry, more information can be found on the [Registry Mirror Tools](#) (see page 1533) page.

To deploy a cluster with a custom image in a region where [CAPI images](#)²⁶⁰ are not provided, you need to use [Konvoy Image Builder](#) (see page 1053) to create your own image for the region.



For multi-tenancy, every tenant should be in a different AWS account to ensure they are truly independent of other tenants in order to enforce security.

²⁶⁰ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

5.11.1.2 Next Step:

[AWS: Create an Image](#) (see page 295)

5.11.1.3 AWS: Create an Image

5.11.1.3.1 Learn how to build a custom AMI for use with DKP

This procedure describes how to use the [Konvoy Image Builder](#) (see page 1461) (KIB) to create a [Cluster API](#)²⁶¹ compliant Amazon Machine Image (AMI). KIB uses [variable overrides](#) (see page 1513) to specify base image and container images to use in your new AMI.



The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to create a custom AMI to take advantage of enhanced cluster operations. Explore the [KIB with AWS](#) (see page 1465) topics for more options.

5.11.1.3.2 Prerequisites

Before you begin, you must:

- Download the [KIB](#) (see page 0) bundle for your version of DKP prefixed with `konvoy-image-bundle` for your OS.
- Check the [Supported Infrastructure Operating Systems](#) (see page 69)
- Check the [Supported Kubernetes Version](#) (see page 1552) for your Provider.
- Create a working registry:
 - Version 4.0 of [Podman](#)²⁶² or higher for Linux. Host requirements found here: <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>
 - Version 20.10.0 of [Docker](#)²⁶³ or higher for Linux or MacOS
- Ensure you have met the minimal set of permissions from the [AWS Image Builder Book](#)²⁶⁴.
- A [Minimal IAM Permissions for KIB](#) (see page 1465) to create an Image for an AWS account using Konvoy Image Builder.

²⁶¹ <https://cluster-api.sigs.k8s.io/>

²⁶² <https://podman.io/getting-started/installation>

²⁶³ <https://www.docker.com/>

²⁶⁴ <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>

5.11.1.3.3 Extract KIB Bundle

Extract the bundle and `cd` into the extracted `konvoy-image-bundle-$VERSION` folder. The bundled version of `konvoy-image` contains an embedded `docker` image that contains all the requirements for building.

ⓘ The `konvoy-image` binary and all supporting folders are also extracted. When run, `konvoy-image` `bind` mounts the current working directory (`${PWD}`) into the container to be used.

- Set environment variables for [AWS access](#)²⁶⁵. The following variables must be set using your credentials including [required IAM](#) (see page 1465):

```
export AWS_ACCESS_KEY_ID
export AWS_SECRET_ACCESS_KEY
export AWS_DEFAULT_REGION
```

- If you have an [override file](#) (see page 1513) to configure specific attributes of your AMI file, add it.

5.11.1.3.4 Build the Image

Depending on which [version of DKP](#) (see page 1461) you are running, steps and flags will be different. To deploy in a region where CAPI images are not provided, you need to use KIB to create your own image for the region. For a list of supported AWS regions, refer to the [Published AMI](#)²⁶⁶ information from AWS.

5.11.1.3.4.1 Execute the following to begin image creation:

Run the `konvoy-image` command to build and validate the image.

```
konvoy-image build aws images/ami/centos-79.yaml
```

By default it builds in the `us-west-2` region. to specify another region set the `--region` flag:

```
konvoy-image build aws --region us-east-1 images/ami/centos-79.yaml
```

²⁶⁵ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-envvars.html>

²⁶⁶ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

Ensure you have named the correct YAML file for your OS in the `konvoy-image build` command.

After KIB provisions the image successfully, the `ami` id is printed and written to the `manifest.json` file. This file has an `artifact_id` field whose value provides the name of the AMI ID as shown in the example below:

```
{
  "name": "rhel-7.9-fips",
  "builder_type": "amazon-ebs",
  "build_time": 1659486130,
  "files": null,
  "artifact_id": "us-west-2:ami-0f2ef742482e1b829",
  "packer_run_uuid": "0ca500d9-a5f0-815c-6f12-aceb4d46645b",
  "custom_data": {
    "containerd_version": "",
    "distribution": "RHEL",
    "distribution_version": "7.9",
    "kubernetes_cni_version": "",
    "kubernetes_version": "1.26.14+fips.0"
  }
}
```

To use a custom AMI when creating your cluster, you must create that AMI using KIB first. Then perform the export and name the custom AMI for use in the command `dkp create cluster`

```
export AWS_AMI_ID=ami-<ami-id-here>
```


5.11.1.3.5 Next Step:

[AWS: Cluster Creation \(see page 298\)](#)

5.11.1.4 AWS: Create the Management Cluster

Use this procedure to create a self-managed AWS Management cluster with DKP. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.

5.11.1.4.1 Name Your Cluster

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)²⁶⁷ for more naming information.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:


1. Give your cluster a unique name suitable for your environment.
In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.
2. Set the environment variable:

```
export CLUSTER_NAME=<aws-example>
```

To increase [Docker Hub's rate limit](#)²⁶⁸ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

5.11.1.4.2 Create a New AWS Kubernetes Cluster

If you use these instructions to create a cluster on AWS using the DKP default settings without any edits to configuration files or additional flags, your cluster is deployed on an [Ubuntu 20.04 operating system image](#)²⁶⁹ with 3 control plane nodes, and 4 worker nodes.

 In previous DKP releases, AMI images provided by the upstream CAPA project would be used if you did not specify an AMI. However, the upstream images are not recommended for production and may not always be available. Therefore, DKP now requires you to specify an AMI when creating a cluster. To create an AMI, use [Konvoy Image Builder](#) (see page 1465).

²⁶⁷ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

²⁶⁸ <https://docs.docker.com/docker-hub/download-rate-limit/>

²⁶⁹ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273747774&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems)

[fromPageId=273747774&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273747774&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems)

There are two approaches to supplying the ID of your AMI. Either provide the ID of the AMI or provide a way for DKP to discover the AMI using location, format and OS information:

1. **Option One** - Provide the ID of your AMI:
 - a. Use the example command below leaving the existing flag that provides the AMI ID: `--ami AMI_ID`
2. **Option Two** - Provide a path for your AMI with the information required for image discover:
 - a. Where the AMI is published using your AWS Account ID: `--ami-owner AWS_ACCOUNT_ID`
 - b. The format or string used to search for matching AMIs and ensure it references the Kubernetes version plus the base OS name: `--ami-base-os ubuntu-20.04`
 - c. The base OS information: `--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*'`

- (Optional) Configure your cluster to use an existing local registry as a mirror when attempting to pull images. Below is an AWS ECR example:

⚠ IMPORTANT: The AMI must be created with [Konvoy Image Builder](#) (see page 1461) in order to use the registry mirror feature.

```
export REGISTRY_MIRROR_URL=<your_registry_url>
```

- `REGISTRY_MIRROR_URL` : the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.

⚠ DKP uses AWS CSI as the [default storage provider](#) (see page 91). You can use a [Kubernetes CSI](#)²⁷⁰ compatible storage solution that is suitable for production. See the Kubernetes documentation called [Changing the Default Storage Class](#)²⁷¹ for more information.

1. Run this command to create your Kubernetes cluster using any relevant flags for **Option One** explained above:

```
dkp create cluster aws \
  --cluster-name=${CLUSTER_NAME} \
  --additional-tags=owner=$(whoami) \
  --with-aws-bootstrap-credentials=true \
  --ami AMI_ID \
  --self-managed
```

OR

2. **Option Two** is to run the command as shown from the explanation above:

²⁷⁰ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

²⁷¹ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

```

dkp create cluster aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--ami-owner AWS_ACCOUNT_ID \
--ami-base-os ubuntu-20.04 \
--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*' \
--self-managed

```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

To understand how this process works step by step, you can find a customizable [Create a New Customized AWS Cluster](#) (see page 1071) under [Custom Installation and Additional Infrastructure Tools](#) (see page 1036).

5.11.1.4.3 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>²⁷²²⁷³.

5.11.1.4.4 Next Step:

[AWS: Install Kommander](#) (see page 301)

5.11.1.5 AWS: Install Kommander

You have installed the Konvoy component and started creating a Management cluster. Now it is time to Install Kommander which completes the Management cluster and allows you to access the UI and attach new or existing managed clusters to monitor.

²⁷² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

²⁷³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

5.11.1.5.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.11.1.5.2 Create and customize your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```


2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init > kommander.yaml
```


4. *If required:* Customize your `kommander.yaml`.

 See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, [Rook Ceph customization \(Pre-provisioned envs\)](#) (see page 1405), etc.

5. *If required:* If your cluster uses a custom **AWS VPC** and requires an internal load-balancer, set the `traefik` annotation to create an internal-facing ELB:


```
apps:
  traefik:
    enabled: true
    values: |
      service:
        annotations:
          service.beta.kubernetes.io/aws-load-balancer-internal: "true"
```

5.11.1.5.3 Enable DKP Catalog Applications and Install Kommander

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
        ref:
          tag: v2.6.1
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.

- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.11.1.5.4 Next Step:

[AWS: Verify Install and Log in the UI](#) (see page 303)

5.11.1.6 AWS: Verify Install and Log in the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.



NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
```



```
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.11.1.6.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

5.11.1.6.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{{ "\n"}}Password: {{.data.password|base64decode}}
{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/dashboard{{ "\n"}}
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQL2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.11.1.6.3 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

Now you can [create additional new clusters](#) (see [page 306](#)), or proceed to [Day 2 Cluster Operations](#) (see [page 582](#)) for attaching existing cluster..

5.11.1.6.4 Next Step:


[AWS: Subsequent New Clusters](#) (see [page 306](#))

5.11.1.7 AWS: Create a Managed Cluster Using the DKP CLI

Enterprise

Gov Advanced


In previous steps, you created a new Management cluster, which is self-managed. When you use these steps to create new Managed clusters, they will become Attached clusters under your Management Cluster.

 When creating [Managed clusters](#) (see [page 81](#)), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see [page 81](#))!

Your new, managed cluster needs to be part of either a new Workspace or an existing Workspace under a management cluster. Create or locate a workspace name to get started with your managed cluster.

5.11.1.7.1 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:


 **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see [page 650](#)).

```
kubectl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.11.1.7.1.1 Name Your Cluster

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)²⁷⁴ for more naming information.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:

1. Give your cluster a unique name suitable for your environment.

In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable:

²⁷⁴ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

```
export MANAGED_CLUSTER_NAME=<aws-additional>
```

5.11.1.7.2 Create a New Cluster with the CLI

ⓘ The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29920384/Attach+a+Kubernetes+Cluster>. (see page 752).

Execute this command to create your additional Kubernetes cluster using any relevant flags. This will create a new non-self-managed cluster that can be managed by [management cluster](#) (see page 298) you created in the previous section:

```
dkp create cluster aws \
--cluster-name=${MANAGED_CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--kubeconfig=<management-cluster-kubeconfig-path> \
--namespace ${WORKSPACE_NAMESPACE}
```

ⓘ If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

5.11.1.7.3 Retrieve the kubeconfig and Explore New AWS Cluster

Follow these steps:


1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:


```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:

 **NOTE:** Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.11.1.7.3.1 Manually Attach a DKP CLI Cluster to the Management Cluster

 These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

- Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

- You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

- This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

- Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

- Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubectrl get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

If you have existing clusters or want to create new clusters to attach, there are many ways to attach a cluster with various requirements and restrictions. To see all the options, visit the section in documentation [Day 2 - Attach an Existing Kubernetes Cluster](#) (see page 752).

5.11.1.7.3.2 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.11.2 AWS Air-gapped Install

This section provides the instructions to install DKP in an AWS air-gapped environment.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.11.2.1 AWS Prerequisites


Before you begin using Konvoy with AWS, you must:

1. Follow the steps to create permissions and roles on the [Minimal Permissions and Role to Create Clusters](#) (see page 1057) page.
2. Create [Cluster IAM Policies and Roles](#) (see page 1063).
3. Export the AWS region where you want to deploy the cluster:


```
export AWS_REGION=us-west-2
```


4. Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```

 If using AWS ECR as your local private registry, more information can be found on the [Registry Mirror Tools](#) (see page 1533) page.

To deploy a cluster with a custom image in a region where [CAPI images](#)²⁷⁵ are not provided, you need to use [Konvoy Image Builder](#) (see page 1053) to create your own image for the region.

 For multi-tenancy, every tenant should be in a different AWS account to ensure they are truly independent of other tenants in order to enforce security.

 For air-gapped, ensure you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, so you can extract the tarball on next page.

5.11.2.2 Next Step:

[AWS Air-gapped: Create an Image](#) (see page 311)

5.11.2.3 AWS Air-gapped: Create an Image

This procedure describes how to use the [Konvoy Image Builder](#) (see page 1461) (KIB) to create a [Cluster API](#)²⁷⁶ compliant Amazon Machine Image (AMI). AMI images contain configuration information and software to create a specific, pre-configured, operating environment. KIB uses [variable overrides](#) (see page 1513) to specify base image and container images to use in your new AMI.

²⁷⁵ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>


²⁷⁶ <https://cluster-api.sigs.k8s.io/>

5.11.2.3.1 Create an Air-gapped AMI

5.11.2.3.1.1 Create an Image using Konvoy Image Builder (KIB) for use in an air-gapped cluster

5.11.2.3.2 Prerequisites

- [Minimal IAM Permissions for KIB \(see page 1465\)](#) to create an Image for an AWS account using Konvoy Image Builder.

 The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to create an AWS Air-gapped AMI to take advantage of enhanced cluster operations. Explore the [KIB with AWS \(see page 1465\)](#) topics for more options.

Using [KIB \(see page 1461\)](#), you can build an AMI without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2/kib
```

2. Follow the instructions below to build an AMI.

Depending on which [version of DKP \(see page 1461\)](#) you are running, steps and flags will be different. To deploy in a region where CAPI images are not provided, you need to use KIB to create your own image for the region. For a list of supported AWS regions, refer to the [Published AMI²⁷⁷](#) information from AWS.

5.11.2.3.2.1 Execute the following to begin image creation:

Run the `konvoy-image` command to build and validate the image.

```
konvoy-image build aws images/ami/centos-79.yaml --overrides overrides/offline.yaml
```

By default it builds in the `us-west-2` region. to specify another region set the `--region` flag:

²⁷⁷ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

```
konvoy-image build aws --region us-east-1 images/ami/centos-79.yaml --overrides
overrides/offline.yaml
```

When the command is complete the `ami` id is printed and written to `./manifest.json`.

5.11.2.3.3 Next Step:


[AWS Air-gapped: Docker Registry](#) (see page 313)

5.11.2.4 AWS Air-gapped: Load the Registry

After you create an image for your air-gapped environment, you will need to load it into your registry.

5.11.2.4.1 Load Images into your Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine](#) (see page 1285) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.

 If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see page 1533) page for more information.

1. [Download](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, and extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

3. Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```

It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

5.11.2.4.2 Kommander Load Images

If you are operating in an air-gapped environment, a [local container registry \(see page 1533\)](#) containing all the necessary installation images, including the Kommander images is required. See below for how to push the necessary images to this registry.

5.11.2.4.3 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.11.2.4.4 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.11.2.4.5 Next Step:

[AWS Air-gapped: Create the Management Cluster \(see page 314\)](#)

5.11.2.5 AWS Air-gapped: Create the Management Cluster

Create a new self-managed AWS Kubernetes cluster in an Air-gapped environment on your AWS infrastructure. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing. When you use existing infrastructure, DKP does *not* create, modify, or delete the following AWS resources:

- Internet Gateways
- NAT Gateways
- Routing tables
- Subnets
- VPC
- VPC Endpoints (for subnets without NAT Gateways)

ⓘ An AWS subnet has Network ACLs that can control traffic in and out of the subnet. DKP does not modify the Network ACLs of an existing subnet. DKP uses Security Groups to control traffic. If a Network ACL denies traffic that is allowed by DKP-managed Security Groups, the cluster may not work correctly.

ⓘ To increase [Docker Hub's rate limit](https://docs.docker.com/docker-hub/download-rate-limit/)²⁷⁸ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io` `--registry-mirror-username=` `--registry-mirror-password=` on the `dkp create cluster` command.

1. Give your cluster a unique name suitable for your environment.

In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable to the name you assigned this cluster:

```
export CLUSTER_NAME=<aws-example>
```

²⁷⁸ <https://docs.docker.com/docker-hub/download-rate-limit/>

NOTE: The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)²⁷⁹ for more naming information.

3. Export variables for the existing infrastructure details:

```
export AWS_VPC_ID=<vpc-...>
export AWS_SUBNET_IDS=<subnet-...,subnet-...,subnet-...>
export AWS_ADDITIONAL_SECURITY_GROUPS=<sg-...>
export AWS_AMI_ID=<ami-...>
```

- `AWS_VPC_ID` : the VPC ID where the cluster will be created. The VPC requires the following AWS VPC Endpoints to be already present:

- `ec2` - `com.amazonaws.{region}.ec2`
- `elasticloadbalancing` - `com.amazonaws.{region}.elasticloadbalancing`
- `secretsmanager` - `com.amazonaws.{region}.secretsmanager`
- `autoscaling` - `com.amazonaws.{region}.autoscaling`
- `ecr` - `com.amazonaws.{region}.ecr.api` - (authentication)
- `ecr` - `com.amazonaws.{region}.ecr.dkr` - (data transfer)

More details about [AWS service using an interface VPC endpoint](#)²⁸⁰ and [AWS VPC endpoints list](#)²⁸¹.

- `AWS_SUBNET_IDS` : a comma-separated list of one or more private Subnet IDs with each one in a different Availability Zone. The cluster control-plane and worker nodes will automatically be spread across these Subnets.
- `AWS_ADDITIONAL_SECURITY_GROUPS` : a comma-separated list of one or more Security Groups IDs to use in addition to the ones automatically created by [CAPA](#)²⁸².
- `AWS_AMI_ID` : the AMI ID to use for control-plane and worker nodes. The AMI must be created by the [konvoy-image-builder](#) (see page 1461).

²⁷⁹ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

²⁸⁰ <https://docs.aws.amazon.com/vpc/latest/privatelink/create-interface-endpoint.html>

²⁸¹ <https://docs.aws.amazon.com/vpc/latest/privatelink/aws-services-privatelink-support.html>

²⁸² <https://github.com/kubernetes-sigs/cluster-api-provider-aws>

! In previous DKP releases, AMI images provided by the upstream CAPA project would be used if you did not specify an AMI. However, the upstream images are not recommended for production and may not always be available. Therefore, DKP now requires you to specify an AMI when creating a cluster. To create an AMI, use [Konvoy Image Builder](#) (see page 1465).

There are two approaches to supplying the ID of your AMI while creating your cluster. **Option One** is to provide the ID of the AMI. **Option Two** is provide a way for DKP to discover the AMI using location, format and OS information using flags.

! **IMPORTANT:** You must tag the subnets as described below to allow for Kubernetes to create External Load Balancers (ELBs) for services of type `LoadBalancer` in those subnets. If the subnets are not tagged, they will not receive an ELB and the following error displays: `Error syncing load balancer, failed to ensure load balancer; could not find any suitable subnets for creating the ELB.`

The tags should be set as follows, where `<CLUSTER_NAME>` corresponds to the name set in `CLUSTER_NAME` environment variable:

```
kubernetes.io/cluster = <CLUSTER_NAME>
kubernetes.io/cluster/<CLUSTER_NAME> = owned
kubernetes.io/role/internal-elb = 1
```

- (Optional) Configure your cluster to use an existing container registry as a mirror when attempting to pull images. The example below is for AWS ECR:


! If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see page 1533) page for more information.

! **IMPORTANT:** The AMI must be created by the [konvoy-image-builder](#) (see page 1461) project in order to use the registry mirror feature.

```
export REGISTRY_MIRROR_URL=<your_registry_url>
```

- `REGISTRY_MIRROR_URL` : the address of an existing registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
- NOTE: Other local registries may use the options below:
 - JFrog - `REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
 - `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
 - `REGISTRY_PASSWORD` : optional if username is not set.

5. Create a Kubernetes cluster. The following example shows a common configuration. See [dkp create cluster aws](#) (see page 1665) reference for the full list of cluster creation options:

 DKP uses AWS CSI as the [default storage provider](#) (see page 91). You can use a [Kubernetes CSI](#)²⁸³ compatible storage solution that is suitable for production. See the Kubernetes documentation called [Changing the Default Storage Class](#)²⁸⁴ for more information.

1. Run this command to create your Kubernetes cluster using any relevant flags for **Option One** explained above providing the AMI ID:

```
dkp create cluster aws --cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--vpc-id=${AWS_VPC_ID} \
--ami=${AWS_AMI_ID} \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=<YOUR_ECR_URL> \
--self-managed
```

OR

2. **Option Two** is to run the command as shown from the explanation above to allow discovery of your AMI:

```
dkp create cluster aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--vpc-id=${AWS_VPC_ID} \
--ami-owner AWS_ACCOUNT_ID \
--ami-base-os ubuntu-20.04 \
--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*' \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=<YOUR_ECR_URL> \
--self-managed
```

²⁸³ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

²⁸⁴ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

You can find a customizable YAML and other custom instructions in [Create a new Air-gapped AWS Cluster](#) (see page 1071) under [Custom Installation and Additional Infrastructure Tools](#) (see page 1036).

5.11.2.5.1 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>²⁸⁵²⁸⁶.

5.11.2.5.2 Next Step:

[AWS Air-gapped: Install Kommander](#) (see page 319)

5.11.2.6 AWS Air-gapped: Install Kommander

You have installed the Konvoy component and created a cluster. Now it is time to Install Kommander which will allow you to access the UI and attach new or existing clusters to monitor.

5.11.2.6.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Ensure you have loaded all necessary images for your configuration. See [Load the Images into Your Registry: Air-gapped Environments](#) (see page 1400).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.11.2.6.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

²⁸⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

²⁸⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>


```
export CLUSTER_NAME=<your-management-cluster-name>
```


2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init --airgapped > kommander.yaml
```


4. *If required:* Customize your `kommander.yaml`.

 See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, [Rook Ceph customization \(Pre-provisioned envs\)](#) (see page 1405), etc.

5. *If required:* If your cluster uses a custom **AWS VPC** and requires an internal load-balancer, set the `traefik` annotation to create an internal-facing ELB:

```
apps:
  traefik:
    enabled: true
    values: |
      service:
        annotations:
          service.beta.kubernetes.io/aws-load-balancer-internal: "true"
```

5.11.2.6.3 Enable DKP Catalog Applications and Install Kommander in an Air-gapped Environment

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add the following values to enable DKP Catalog Applications:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
...
```

```
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      path: ./dkp-catalog-applications-v2.6.1.tar.gz
```

⚠️ If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${
{CLUSTER_NAME}.conf \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz
\
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-
v2.6.1.tar.gz
```



Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File \(see page 87\)](#).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.11.2.6.4 Next Step:

[AWS Air-gapped: Verify Install and Log in the UI \(see page 321\)](#)

5.11.2.7 AWS Air-gapped: Verify Install and Log in the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.

NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the install command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
```

```
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.11.2.7.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

5.11.2.7.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{\n"}Password: {{.data.password|base64decode}}
{\n"}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/
dashboard{\n"}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQL2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see page 597):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.11.2.7.3 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see page 582) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

Now you can [create additional new clusters](#) (see page 323), or proceed to [Day 2 Operations](#) (see page 582).

5.11.2.7.4 Next Step:

[AWS Air-gapped: Subsequent New Clusters](#) (see page 323)

5.11.2.8 AWS Air-gapped: Create a Managed Cluster Using the DKP CLI

Enterprise

Gov Advanced

In previous steps, you created a new Management cluster, which is self-managed. When you use these steps to create new Managed clusters, they will become Attached clusters under your Management Cluster.



When creating [Managed clusters](#) (see page 81), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see page 81)!

5.11.2.8.1 Make New Managed Cluster Part of a Workspace

Your new managed cluster needs to be part of either a new Workspace or an existing Workspace.

To create a new Workspace, follow the instructions to [Create a New Workspace](#)²⁸⁷.


To find the existing Workspaces, use this command:

```
kubectl get workspace -A
```

When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.11.2.8.2 Name Your Cluster

 The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29920384/Attach+a+Kubernetes+Cluster>. (see page 752).

Follow these steps:

1. Give your cluster a unique name suitable for your environment.

In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable:

```
export MANAGED_CLUSTER_NAME=<aws-additional>
```

²⁸⁷ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/213320100>

⚠ The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes²⁸⁸](#) for more naming information.

3. Execute this command to create your Kubernetes cluster using any relevant flags but note that the `--self-managed` is absent so that this cluster can be attached to your previously created Management Cluster.

```
dkp create cluster aws --cluster-name=${MANAGED_CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--namespace ${WORKSPACE_NAMESPACE} \
--kubeconfig=<management-cluster-kubeconfig-path> \
--with-aws-bootstrap-credentials=true \
--vpc-id=${AWS_VPC_ID} \
--ami=${AWS_AMI_ID} \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
--registry-mirror-password=${REGISTRY_PASSWORD} \
--kubeconfig=<management-cluster-kubeconfig-path> \
```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1039\)](#).

5.11.2.8.3 Retrieve the kubeconfig and Explore New AWS Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:


```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:

²⁸⁸ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>


```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:

 **NOTE:** Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.11.2.8.3.1 Manually Attach a DKP CLI Cluster to the Management Cluster

 These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:


```
kubectl get workspaces -A
```

- Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

- You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

- This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

- Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

- Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
```

```
clusterRef:
  capiCluster:
    name: ${MANAGED_CLUSTER_NAME}
EOF
```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubect1 get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.11.2.8.3.2 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.11.3 AWS with FIPS Install

This section provides instructions to install DKP in an AWS non-air-gapped FIPS environment.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.11.3.1 AWS Prerequisites

Before you begin using Konvoy with AWS, you must:

1. Follow the steps to create permissions and roles on the [Minimal Permissions and Role to Create Clusters](#) (see page 1057) page.
2. Create [Cluster IAM Policies and Roles](#) (see page 1063).
3. Export the AWS region where you want to deploy the cluster:

```
export AWS_REGION=us-west-2
```

4. Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```



If using AWS ECR as your local private registry, more information can be found on the [Registry Mirror Tools](#) (see page 1533) page.

To deploy a cluster with a custom image in a region where [CAPI images](#)²⁸⁹ are not provided, you need to use [Konvoy Image Builder](#) (see page 1053) to create your own image for the region.



For multi-tenancy, every tenant should be in a different AWS account to ensure they are truly independent of other tenants in order to enforce security.

5.11.3.2 Next Step:

[AWS FIPS: Create an Image](#) (see page 329)

5.11.3.3 AWS FIPS: Create an Image

5.11.3.3.1 Learn how to build a custom AMI for use with DKP

This procedure describes how to use the [Konvoy Image Builder](#) (see page 1461) (KIB) to create a [Cluster API](#)²⁹⁰ compliant Amazon Machine Image (AMI). KIB uses [variable overrides](#) (see page 1513) to specify base image and container images to use in your new AMI.

²⁸⁹ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

²⁹⁰ <https://cluster-api.sigs.k8s.io/>



The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to create a custom AMI to take advantage of enhanced cluster operations. Explore the [KIB with AWS \(see page 1465\)](#) topics for more options.

5.11.3.3.2 Prerequisites

Before you begin, you must:

- Download the [KIB \(see page 0\)](#) bundle for your version of DKP prefixed with `konvoy-image-bundle` for your OS.
- Check the [Supported Infrastructure Operating Systems \(see page 69\)](#)
- Check the [Supported Kubernetes Version \(see page 1552\)](#) for your Provider.
- Create a working registry:
 - Version 4.0 of [Podman²⁹¹](#) or higher for Linux. Host requirements found here: <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>
 - Version 20.10.0 of [Docker²⁹²](#) or higher for Linux or MacOS
- Ensure you have met the minimal set of permissions from the [AWS Image Builder Book²⁹³](#).
- A [Minimal IAM Permissions for KIB \(see page 1465\)](#) to create an Image for an AWS account using Konvoy Image Builder.

5.11.3.3.3 Extract KIB Bundle

Extract the bundle and `cd` into the extracted `konvoy-image-bundle-$VERSION` folder. The bundled version of `konvoy-image` contains an embedded `docker` image that contains all the requirements for building.



The `konvoy-image` binary and all supporting folders are also extracted. When run, `konvoy-image` `bind` mounts the current working directory (`${PWD}`) into the container to be used.

²⁹¹ <https://podman.io/getting-started/installation>

²⁹² <https://www.docker.com/>

²⁹³ <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>

- Set environment variables for [AWS access](#)²⁹⁴. The following variables must be set using your credentials including [required IAM](#) (see page 1465):

```
export AWS_ACCESS_KEY_ID
export AWS_SECRET_ACCESS_KEY
export AWS_DEFAULT_REGION
```

- If you have an [override file](#) (see page 1513) to configure specific attributes of your AMI file, add it.

5.11.3.3.4 Non-air-gapped Environment Create FIPS-140 images

KIB can produce images containing FIPS-140 compliant binaries. Use the `fips.yaml` [override file](#) (see page 1514) provided with the image bundles.

 You can also find these override files in the [Konvoy Image Builder repo](#)²⁹⁵.

- A non-air-gapped example of override file use is the command below, which produces a FIPS-compliant image on RHEL 8.4 for AWS:

Replace `ami` with your infrastructure provisioner

```
konvoy-image build --overrides overrides/fips.yaml images/ami/rhel-84.yaml
```

5.11.3.3.5 Next Step:

[AWS FIPS Cluster Creation](#) (see page 331)

5.11.3.4 AWS FIPS: Create the Management Cluster

Use this procedure to create a self-managed AWS FIPS Management cluster with DKP. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing. Additional considerations exist when using FIPS. Alter the Create a Kubernetes cluster command by using the FIPS example, or similar, when executing the command to create a cluster.

²⁹⁴ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-envvars.html>

²⁹⁵ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>


5.11.3.4.1 Deploying a Cluster in FIPS mode

In order to create a cluster in FIPS mode, we must inform the bootstrap controllers of the appropriate image repository and version tags of the official D2iQ FIPS builds of Kubernetes.

5.11.3.4.1.1 Supported FIPS Builds

Component	Repository	Version
Kubernetes	docker.io/mesosphere²⁹⁶	v1.26.14+fips.0
etcd	docker.io/mesosphere²⁹⁷	3.5.10+fips.0

5.11.3.4.1.2 Name Your Cluster

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes²⁹⁸](https://kubernetes.io/docs/concepts/overview/working-with-objects/names/) for more naming information.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:

1. Give your cluster a unique name suitable for your environment.

In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable:

```
export CLUSTER_NAME=<aws-example>
```


²⁹⁶ <https://hub.docker.com/layers/mesosphere/kube-apiserver/v1.26.14+fips.0/images/sha256-3c559c75123fa40d588bc2d9050938a0b7dd1dab4e38a4fd24847cd0c0f24694?context=explore>

²⁹⁷ <https://hub.docker.com/layers/mesosphere/etcd/v3.5.10+fips.0/images/sha256-61df32b8ec1e6229d8bd94caad07331c7e0c0b8e1947af8152c69eb2f04fb771?context=explore>


²⁹⁸ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

5.11.3.4.1.3 Create a New AWS Kubernetes Cluster

If you use these instructions to create a cluster on AWS using the DKP default settings without any edits to configuration files or additional flags, your cluster is deployed on an [Ubuntu 20.04 operating system image](#)²⁹⁹³⁰⁰ with 3 control plane nodes, and 4 worker nodes.

 In previous DKP releases, AMI images provided by the upstream CAPA project would be used if you did not specify an AMI. However, the upstream images are not recommended for production and may not always be available. Therefore, DKP now requires you to specify an AMI when creating a cluster. To create an AMI, use [Konvoy Image Builder](#) (see page 1465).

There are two approaches to supplying the ID of your AMI. Either provide the ID of the AMI or provide a way for DKP to discover the AMI using location, format and OS information:


1. **Option One** - Provide the ID of your AMI:
 - a. Use the example command below leaving the existing flag that provides the AMI ID: `--ami AMI_ID`
 2. **Option Two** - Provide a path for your AMI with the information required for image discover:
 - a. Where the AMI is published using your AWS Account ID: `--ami-owner AWS_ACCOUNT_ID`
 - b. The format or string used to search for matching AMIs and ensure it references the Kubernetes version plus the base OS name: `--ami-base-os ubuntu-20.04`
 - c. The base OS information: `--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*'`
- (Optional) Configure your cluster to use an existing local registry as a mirror when attempting to pull images. Below is an AWS ECR example:
 -  **IMPORTANT:** The AMI must be created with [Konvoy Image Builder](#) (see page 1461) in order to use the registry mirror feature.

```
export REGISTRY_MIRROR_URL=<your_registry_url>
```

²⁹⁹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273747774&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems>

³⁰⁰ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273747774&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems>

- `REGISTRY_MIRROR_URL` : the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.

 DKP uses AWS CSI as the [default storage provider](#) (see page 91). You can use a [Kubernetes CSI](#)³⁰¹ compatible storage solution that is suitable for production. See the Kubernetes documentation called [Changing the Default Storage Class](#)³⁰² for more information.


1. Run this command to create your Kubernetes cluster using any relevant flags for **Option One** explained above:

```
dkp create cluster aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--ami AMI_ID \
--kubernetes-version=v1.26.14+fips.0 \
--etcd-version=3.5.10+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--self-managed
```

OR

2. **Option Two** is to run the command as shown from the explanation above:

```
dkp create cluster aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--ami-owner AWS_ACCOUNT_ID \
--ami-base-os ubuntu-20.04 \
--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*' \
--kubernetes-version=v1.26.14+fips.0 \
--etcd-version=3.5.10+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--self-managed
```

 If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

301 <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

302 <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

You can find a customizable [Create a new AWS Cluster](#) (see page 1071) under the [Custom Installation and Additional Infrastructure Tools](#) (see page 1036). For more details about FIPS, refer to the [FIPS 140-2 Compliance](#) (see page 1526) section of the documentation.

5.11.3.4.1.4 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>³⁰³³⁰⁴.

5.11.3.4.1.5 Next Step:

[AWS FIPS: Install Kommander](#) (see page 335)

5.11.3.5 AWS FIPS: Install Kommander

You have installed the Konvoy component and created a cluster. Now it is time to Install Kommander which will allow you to access the UI and attach new or existing clusters to monitor.

5.11.3.5.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.11.3.5.2 Create and customize your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

2. Copy the `kubeconfig` file of your Management cluster to your local directory:

³⁰³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

³⁰⁴ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init > kommander.yaml
```

4. *If required:* Customize your `kommander.yaml`.

i See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, [Rook Ceph customization \(Pre-provisioned envs\)](#) (see page 1405), etc.

5. *If required:* If your cluster uses a custom **AWS VPC** and requires an internal load-balancer, set the `traefik` annotation to create an internal-facing ELB:

```
apps:
  traefik:
    enabled: true
    values: |
      service:
        annotations:
          service.beta.kubernetes.io/aws-load-balancer-internal: "true"
```

5.11.3.5.3 Enable DKP Catalog Applications and Install Kommander



If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
```

```
url: https://github.com/mesosphere/dkp-catalog-applications
ref:
  tag: v2.6.1
```

⚠️ If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

📌 Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.11.3.5.4 Next Step:

[AWS FIPS: Verify Install and Log in the UI](#) (see page 337)

5.11.3.6 AWS FIPS: Verify Install and Log in the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.

📌 **NOTE:** If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.11.3.6.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

5.11.3.6.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{\n"}Password: {{.data.password|base64decode}}
{{ "\n"}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{or .hostname .ip}}{end}}/dkp/kommander/
dashboard{{ "\n"}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbJUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.11.3.6.3 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see page 582) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

Now you can [create additional new clusters](#) (see page 340), or proceed to Day 2 Operations in the link below.

5.11.3.6.4 Next Step:

[AWS FIPS: Subsequent New Clusters](#) (see page 340)

5.11.3.7 AWS FIPS: Create a Managed Cluster Using the DKP CLI

Enterprise

Gov Advanced

In previous steps, you created a new Management cluster, which is self-managed. When you use these steps to create new Managed clusters, they will become Attached clusters under your Management Cluster.



When creating [Managed clusters](#) (see page 81), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see page 81)!

5.11.3.7.1 Create New Managed FIPS Cluster with the CLI



The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29920384/Attach+a+Kubernetes+Cluster>. (see page 752).

Your new, managed cluster needs to be part of either a new Workspace or an existing Workspace under a management cluster. Create or locate a workspace name to get started with your managed cluster.

5.11.3.7.2 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:
⚠ NOTE: If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see page 650).

```
kubectll get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.11.3.7.3 Name Your Cluster

Follow these steps:

1. Give your cluster a unique name suitable for your environment.

In AWS, it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable:

```
export MANAGED_CLUSTER_NAME=<aws-additional>
```

ⓘ The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)³⁰⁵ for more naming information.

3. Execute this command to create your Kubernetes cluster using any relevant flags but note that the `--self-managed` is absent so that this cluster can be attached to your previously created Management Cluster.

Execute this command:

```
dkp create cluster aws \  
--cluster-name=${MANAGED_CLUSTER_NAME} \  
--additional-tags=owner=$(whoami) \  

```

³⁰⁵ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

```
--namespace ${WORKSPACE_NAMESPACE}
--with-aws-bootstrap-credentials=true \
--kubernetes-version=v1.26.6+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere
--kubeconfig=<management-cluster-kubeconfig-path>
```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

5.11.3.7.4 Retrieve the kubeconfig and Explore New AWS Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:

NOTE: Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```


5.11.3.7.4.1 Manually Attach a DKP CLI Cluster to the Management Cluster



These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

- This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new attached-cluster-kubeconfig.yaml file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

- Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

- Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

- You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubectl get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

After you have existing clusters or want to create new clusters to attach, there are many ways to attach a cluster with various requirements and restrictions. To see all the options, visit the section in documentation [Day 2 - Attach an Existing Kubernetes Cluster](#) (see page 752).

5.11.3.7.4.2 Next Step:

[Day 2 - Cluster Operations Management \(see page 582\)](#)

5.11.4 AWS Air-gapped FIPS Install

This section provides instructions to install DKP in an AWS air-gapped FIPS environment.

If not already done, refer to [Get Started \(see page 79\)](#) section of the documentation for:

- [Resource Requirements \(see page 100\)](#)
- [Install Overview \(see page 108\)](#)
- [Prerequisites for Install \(see page 121\)](#)

5.11.4.1 AWS Prerequisites

Before you begin using Konvoy with AWS, you must:

1. Follow the steps to create permissions and roles on the [Minimal Permissions and Role to Create Clusters \(see page 1057\)](#) page.
2. Create [Cluster IAM Policies and Roles \(see page 1063\)](#).
3. Export the AWS region where you want to deploy the cluster:

```
export AWS_REGION=us-west-2
```

4. Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```



If using AWS ECR as your local private registry, more information can be found on the [Registry Mirror Tools \(see page 1533\)](#) page.

To deploy a cluster with a custom image in a region where [CAPI images](#)³⁰⁶ are not provided, you need to use [Konvoy Image Builder \(see page 1053\)](#) to create your own image for the region.

³⁰⁶ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>



For multi-tenancy, every tenant should be in a different AWS account to ensure they are truly independent of other tenants in order to enforce security.



For air-gapped, ensure you have [Download DKP \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, so you can extract the tarball on next page.

5.11.4.2 Next Step:

[AWS Air-gapped FIPS: Create an Image \(see page 346\)](#)

5.11.4.3 AWS Air-gapped FIPS: Create an Image

5.11.4.3.1 Create an Image using Konvoy Image Builder (KIB) for use in an air-gapped cluster

5.11.4.3.2 Prerequisites

- [Minimal IAM Permissions for KIB \(see page 1465\)](#) to create an Image for an AWS account using Konvoy Image Builder.



The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to create an AWS Air-gapped AMI to take advantage of enhanced cluster operations. Explore the [KIB with AWS \(see page 1465\)](#) topics for more options.

Using [KIB \(see page 1461\)](#), you can build an AMI without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2/kib
```

2. Follow the instructions below to build an AMI.

5.11.4.3.2.1 Offline FIPS Override File (Air-gapped)

Add the following FIPS Overrides file to your environment:

```
--overrides overrides/offline-fips.yaml
```

```
# fips os-packages
os_packages_local_bundle_file: "{{ playbook_dir }}/../artifacts/
{{ kubernetes_version }}-{{ ansible_distribution|lower }}
_{{ ansible_distribution_major_version }}_x86_64_fips.tar.gz"
containerd_local_bundle_file: "{{ playbook_dir }}/../artifacts/
{{ containerd_tar_file }}"
pip_packages_local_bundle_file: "{{ playbook_dir }}/../artifacts/pip-packages.tar.gz"
images_local_bundle_dir: "{{ playbook_dir }}/../artifacts/images"
```



You can find all available Overrides files in the [Konvoy Image Builder repo](#)³⁰⁷.

- An air-gapped example of override file use is the command below which produces an AWS FIPS-compliant image on RHEL 8.4:

```
konvoy-image build --overrides overrides/offline-fips.yaml images/ami/rhel-84.yaml
```

5.11.4.3.3 Next Step:

[AWS Air-gapped FIPS: Load the Registry](#) (see page 347)


5.11.4.4 AWS Air-gapped FIPS: Load the Registry

After you create an image for your air-gapped environment, you will need to load it into your registry.

5.11.4.4.1 Load Images into your Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine](#) (see page 1285) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.

³⁰⁷ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

 If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see page 1533) page for more information.

1. [Download](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz` , and extract the tarball to a local directory:


```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

3. Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```

 It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

5.11.4.4.2 Kommander Load Images

If you are operating in an air-gapped environment, a [local container registry](#) (see page 1533) containing all the necessary installation images, including the Kommander images is required. See below for how to push the necessary images to this registry.

5.11.4.4.3 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.11.4.4.4 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:


```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.11.4.4.5 Next Step:

[AWS Air-gapped FIPS: Create the Management Cluster](#) (see page 349)

5.11.4.5 AWS Air-gapped FIPS: Create the Management Cluster

Create a new self-managed AWS Kubernetes cluster in an Air-gapped FIPS environment on your AWS infrastructure. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.

 To increase [Docker Hub's rate limit](#)³⁰⁸ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

1. Give your cluster a unique name suitable for your environment.

In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable to the name you assigned this cluster:

³⁰⁸ <https://docs.docker.com/docker-hub/download-rate-limit/>

```
export CLUSTER_NAME=<aws-example>
```

i NOTE: The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)³⁰⁹ for more naming information.

3. Export variables for the existing infrastructure details:

```
export AWS_VPC_ID=<vpc-...>
export AWS_SUBNET_IDS=<subnet-...,subnet-...,subnet-...>
export AWS_ADDITIONAL_SECURITY_GROUPS=<sg-...>
export AWS_AMI_ID=<ami-...>
```

- `AWS_VPC_ID` : the VPC ID where the cluster will be created. The VPC requires the following AWS VPC Endpoints to be already present:
 - `ec2` - `com.amazonaws.{region}.ec2`
 - `elasticloadbalancing` - `com.amazonaws.{region}.elasticloadbalancing`
 - `secretsmanager` - `com.amazonaws.{region}.secretsmanager`
 - `autoscaling` - `com.amazonaws.{region}.autoscaling`
 - `ecr` - `com.amazonaws.{region}.ecr.api` - (authentication)
 - `ecr` - `com.amazonaws.{region}.ecr.dkr` - (data transfer)

More details about [AWS service using an interface VPC endpoint](#)³¹⁰ and [AWS VPC endpoints list](#)³¹¹.

- `AWS_SUBNET_IDS` : a comma-separated list of one or more private Subnet IDs with each one in a different Availability Zone. The cluster control-plane and worker nodes will automatically be spread across these Subnets.
- `AWS_ADDITIONAL_SECURITY_GROUPS` : a comma-separated list of one or more Security Groups IDs to use in addition to the ones automatically created by [CAPA](#)³¹².
- `AWS_AMI_ID` : the AMI ID to use for control-plane and worker nodes. The AMI must be created by the [konvoy-image-builder](#) (see page 1461).

³⁰⁹ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

³¹⁰ <https://docs.aws.amazon.com/vpc/latest/privatelink/create-interface-endpoint.html>

³¹¹ <https://docs.aws.amazon.com/vpc/latest/privatelink/aws-services-privatelink-support.html>

³¹² <https://github.com/kubernetes-sigs/cluster-api-provider-aws>

⚠ In previous DKP releases, AMI images provided by the upstream CAPA project would be used if you did not specify an AMI. However, the upstream images are not recommended for production and may not always be available. Therefore, DKP now requires you to specify an AMI when creating a cluster. To create an AMI, use [Konvoy Image Builder](#) (see page 1465).

There are two approaches to supplying the ID of your AMI while creating your cluster. **Option One** is to provide the ID of the AMI. **Option Two** is provide a way for DKP to discover the AMI using location, format and OS information using flags.

⚠ IMPORTANT: You must tag the subnets as described below to allow for Kubernetes to create External Load Balancers (ELBs) for services of type `LoadBalancer` in those subnets. If the subnets are not tagged, they will not receive an ELB and the following error displays: `Error syncing load balancer, failed to ensure load balancer; could not find any suitable subnets for creating the ELB.`

The tags should be set as follows, where `<CLUSTER_NAME>` corresponds to the name set in `CLUSTER_NAME` environment variable:

```
kubernetes.io/cluster = <CLUSTER_NAME>
kubernetes.io/cluster/<CLUSTER_NAME> = owned
kubernetes.io/role/internal-elb = 1
```

- (Optional) Configure your cluster to use an existing container registry as a mirror when attempting to pull images. The example below is for AWS ECR:

⚠ If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see page 1533) page for more information.


⚠ IMPORTANT: The AMI must be created by the [konvoy-image-builder](#) (see page 1461) project in order to use the registry mirror feature.

```
export REGISTRY_MIRROR_URL=<your_registry_url>
```

- `REGISTRY_MIRROR_URL` : the address of an existing registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
- NOTE: Other local registries may use the options below:

- `JFrog - REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
- `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
- `REGISTRY_PASSWORD` : optional if username is not set.

5. Create a Kubernetes cluster. The following example shows a common configuration. See [dkp create cluster aws](#) (see page 1665) reference for the full list of cluster creation options:

 DKP uses AWS CSI as the [default storage provider](#) (see page 91). You can use a [Kubernetes CSI](#)³¹³ compatible storage solution that is suitable for production. See the Kubernetes documentation called [Changing the Default Storage Class](#)³¹⁴ for more information.

6. Run this command to create your Kubernetes cluster using any relevant flags for **Option One** explained above providing the AMI ID:

```
dkp create cluster aws --cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--vpc-id=${AWS_VPC_ID} \
--ami=${AWS_AMI_ID} \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
--registry-mirror-password=${REGISTRY_PASSWORD} \
--kubernetes-version=v1.26.14+fips.0 \
--etcd-version=3.5.10+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--self-managed
```

OR

7. **Option Two** is to run the command as shown from the explanation above providing the location, format and base OS of your image:

```
dkp create cluster aws --cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
```

³¹³ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

³¹⁴ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

```

--vpc-id=${AWS_VPC_ID} \
--ami=${AWS_AMI_ID} \
--ami-owner AWS_ACCOUNT_ID \
--ami-base-os ubuntu-20.04 \
--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*' \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
--registry-mirror-password=${REGISTRY_PASSWORD} \
--kubernetes-version=v1.26.14+fips.0 \
--etcd-version=3.5.10+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--self-managed

```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

You can find a customizable [Create a New AWS Cluster](#) (see page 1071) under the [Custom Installation and Additional Infrastructure Tools](#) (see page 1036). For more details about FIPS, refer to the [FIPS 140-2 Compliance](#) (see page 1526) section of the documentation.

5.11.4.5.1 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPagelId=273320989>³¹⁵³¹⁶.

5.11.4.5.2 Next Step:

[AWS Air-gapped FIPS: Install Kommander](#) (see page 353)

³¹⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

³¹⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

5.11.4.6 AWS Air-gapped FIPS: Install Kommander

You have installed the Konvoy component and created a cluster. Now it is time to Install Kommander which will allow you to access the UI and attach new or existing clusters to monitor.

5.11.4.6.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Ensure you have loaded all necessary images for your configuration. See [Load the Images into Your Registry: Air-gapped Environments](#) (see page 1400).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.11.4.6.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```


2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init --airgapped > kommander.yaml
```

4. *If required, customize* your `kommander.yaml`.

 See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include:

- Custom Domains and Certificates
- HTTP proxy
- External Load Balancer
- GPU utilization, etc.
- [Rook Ceph customization for Pre-provisioned environments](#) (see page 1405)

5. *If required:* If your cluster uses a custom **AWS VPC** and requires an internal load-balancer, set the `traefik` annotation to create an internal-facing ELB:


```

...
apps:
  traefik:
    enabled: true
    values: |
      service:
        annotations:
          service.beta.kubernetes.io/aws-load-balancer-internal: "true"
...

```

- Expand **one** of the following sets of instructions, depending on your license and application environments:

5.11.4.6.3 Enable DKP Catalog Applications and Install Kommander in an Air-gapped Environment


 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

- In the same `kommander.yaml` of the previous section, add the following values to enable DKP Catalog Applications:

```

apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      path: ./dkp-catalog-applications-v2.6.1.tar.gz

```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

- Use the customized `kommander.yaml` to install DKP:

```

dkp install kommander --installer-config kommander.yaml --kubeconfig=${
{CLUSTER_NAME}.conf \

```

```
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz
\
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-
v2.6.1.tar.gz
```

Tips and recommendations


- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.11.4.6.4 Next Step:

[AWS Air-gapped FIPS: Verify Install and Log in the UI](#) (see page 355)

5.11.4.7 AWS Air-gapped FIPS: Verify Install and Log in the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.

 **NOTE:** If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.11.4.7.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
```

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": false}]'
```

5.11.4.7.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{"\n"}}Password: {{.data.password|base64decode}}
{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/
dashboard{{ "\n"}}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.11.4.7.3 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2

section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

Now you can [create additional new clusters \(see page 358\)](#), or proceed to Day 2 Operations in the link below.

5.11.4.7.4 Next Step:

[AWS Air-gapped FIPS: Subsequent New Clusters \(see page 358\)](#)

5.11.4.8 AWS Air-gapped FIPS: Create a Managed Cluster Using the DKP CLI

Enterprise

Gov Advanced

In previous steps, you created a new Management cluster, which is self-managed. When you use these steps to create new Managed clusters, they will become Attached clusters under your Management Cluster.



When creating [Managed clusters \(see page 81\)](#), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters \(see page 81\)](#)!

5.11.4.8.1 Create New Managed Air-gapped FIPS Cluster with the CLI



The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29920384/Attach+a+Kubernetes+Cluster>. (see page 752).

Your new, managed cluster needs to be part of either a new Workspace or an existing Workspace under a management cluster. Create or locate a workspace name to get started with your managed cluster.

5.11.4.8.2 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:

! **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace \(see page 650\)](#).

```
kubectl get workspace -A
```

- When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.11.4.8.3 Name Your Cluster

Follow these steps:

- Give your cluster a unique name suitable for your environment.

In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

- Set the environment variable:

```
export MANAGED_CLUSTER_NAME=<aws-additional>
```



The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)³¹⁷ for more naming information.


- Execute this command to create your Kubernetes cluster using any relevant flags but note that the `--self-managed` is absent so that this cluster can be attached to your previously created Management Cluster:

```
dkp create cluster aws
--cluster-name=${MANAGED_CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--namespace ${WORKSPACE_NAMESPACE}
--with-aws-bootstrap-credentials=true \
--vpc-id=${AWS_VPC_ID} \
--ami=${AWS_AMI_ID} \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
```

³¹⁷ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

```
--registry-mirror-password=${REGISTRY_PASSWORD} \  
--kubeconfig=<management-cluster-kubeconfig-path> \  

```

 If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

5.11.4.8.4 Retrieve the kubeconfig and Explore New AWS Cluster

Follow these steps:


1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig  
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $  
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:


```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:

 **NOTE:** Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.11.4.8.4.1 Manually Attach a DKP CLI Cluster to the Management Cluster

 These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

7. This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
```

```
value: <value-you-copied-from-secret-above>
```

8. Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

9. Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubectl get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

When you have existing clusters or want to create other new clusters to attach, there are many ways to attach a cluster with various requirements and restrictions. To see all the options, visit the section in the Day 2- Cluster Operations Management section for documentation on how to [Attach an Existing Kubernetes Cluster](#) (see page 752).

5.11.4.8.4.2 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.11.5 AWS with GPU Install

This section provides instructions to install DKP in an AWS non-air-gapped GPU environment.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.11.5.1 AWS Prerequisites

Before you begin using Konvoy with AWS, you must:

1. Follow the steps to create permissions and roles on the [Minimal Permissions and Role to Create Clusters](#) (see page 1057) page.
2. Create [Cluster IAM Policies and Roles](#) (see page 1063).
3. Export the AWS region where you want to deploy the cluster:

```
export AWS_REGION=us-west-2
```

4. Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```



If using AWS ECR as your local private registry, more information can be found on the [Registry Mirror Tools](#) (see page 1533) page.

To deploy a cluster with a custom image in a region where [CAPI images](#)³¹⁸ are not provided, you need to use [Konvoy Image Builder](#) (see page 1053) to create your own image for the region.



For multi-tenancy, every tenant should be in a different AWS account to ensure they are truly independent of other tenants in order to enforce security.

318 <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

5.11.5.2 GPU Prerequisites

Before you begin, you must:

- Ensure nodes provide an NVIDIA GPU.
- If you are using a public cloud service such as [AWS](#) (see page 1110), create an AMI with KIB using the instructions on the [KIB for GPU](#) (see page 1485) page.
- If you are deploying in a pre-provisioned environment, ensure that you have created the appropriate [secret for your GPU nodepool](#) (see page 240) and have uploaded the appropriate artifacts to each node. See the [GPU only steps section](#) (see page 0) on the Pre-provisioned Prerequisites Air-gapped page for additional information.



Specific instructions must be followed for enabling `nvidia-gpu-operator` depending on if you want to deploy the app on a Management cluster or a Attached or a Managed cluster.

- For instructions on enabling the NVIDIA platform application on a Management cluster, follow the instructions in the [NVIDIA Platform Application Management Cluster](#) (see page 970) section.
- For instructions on enabling the NVIDIA platform application on attached or managed clusters, follow the instructions in the [NVIDIA Platform Application Attached or Managed Cluster](#) (see page 972) section.

After `nvidia-gpu-operator` has been enabled depending on the cluster type, proceed to the **Select the correct Toolkit version for your NVIDIA GPU Operator** on each of those pages.

5.11.5.3 Next Step:

[AWS GPU: Use Node Label Automatic Configuration](#) (see page 364)

5.11.5.4 AWS GPU: Use Node Label Automatic Configuration

When using GPU nodes, it is important they have the proper label identifying them as Nvidia GPU nodes. Node feature discovery (NFD), by default labels PCI hardware as:

```
"feature.node.kubernetes.io/pci-<device label>.present": "true"
```

where `<device label>` is by default as [defined in this topic](#)³¹⁹:

³¹⁹ <https://kubernetes-sigs.github.io/node-feature-discovery/v0.7/get-started/features.html#pci>

```
< class > _ < vendor >
```

However, because there is a wide variety in devices and their assigned PCI classes, you may find that the labels assigned to your GPU nodes do not always properly identify them as containing an Nvidia GPU.

If the default detection does not work, you can manually change the daemonset that the GPU operator creates by running the following command:

```
nodeSelector:
  feature.node.kubernetes.io/pci-< class > _ < vendor>.present: "true"
```

where `class` is any 4 digit number starting with `03xy` and the vendor for Nvidia is `10de`. If this is already deployed, you can always change the `daemonset` and change the `nodeSelector` field so that it deploys to the right nodes.

5.11.5.4.1 Next Step:

[AWS GPU: Create an Image \(see page 365\)](#)

5.11.5.5 AWS GPU: Create an Image

5.11.5.5.1 Learn how to build a custom AMI for use with DKP

This procedure describes how to use the [Konvoy Image Builder \(see page 1461\)](#) (KIB) to create a [Cluster API³²⁰](#) compliant Amazon Machine Image (AMI). KIB uses [variable overrides \(see page 1513\)](#) to specify base image and container images to use in your new AMI.



The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to create a custom AMI to take advantage of enhanced cluster operations. Explore the [KIB with AWS \(see page 1465\)](#) topics for more options.

5.11.5.5.2 Prerequisites


Before you begin, you must:

³²⁰ <https://cluster-api.sigs.k8s.io/>

- Download the [KIB \(see page 0\)](#) bundle for your version of DKP prefixed with `konvoy-image-bundle` for your OS.
- Check the [Supported Infrastructure Operating Systems \(see page 69\)](#)
- Check the [Supported Kubernetes Version \(see page 1552\)](#) for your Provider.
- Create a working registry:
 - Version 4.0 of [Podman](#)³²¹ or higher for Linux. Host requirements found here: <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>
 - Version 20.10.0 of [Docker](#)³²² or higher for Linux or MacOS
- Ensure you have met the minimal set of permissions from the [AWS Image Builder Book](#)³²³.
- A [Minimal IAM Permissions for KIB \(see page 1465\)](#) to create an Image for an AWS account using Konvoy Image Builder.

5.11.5.5.3 Extract KIB Bundle

Extract the bundle and `cd` into the extracted `konvoy-image-bundle-$VERSION` folder. The bundled version of `konvoy-image` contains an embedded `docker` image that contains all the requirements for building.

 The `konvoy-image` binary and all supporting folders are also extracted. When run, `konvoy-image` bind mounts the current working directory (`${PWD}`) into the container to be used.

- Set environment variables for [AWS access](#)³²⁴. The following variables must be set using your credentials including [required IAM \(see page 1465\)](#):

```
export AWS_ACCESS_KEY_ID
export AWS_SECRET_ACCESS_KEY
export AWS_DEFAULT_REGION
```

- If you have an [override file \(see page 1513\)](#) to configure specific attributes of your AMI file, add it.

³²¹ <https://podman.io/getting-started/installation>

³²² <https://www.docker.com/>

³²³ <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>

³²⁴ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-envvars.html>

5.11.5.5.4 Build the GPU Image

Using the [Konvoy Image Builder](#) (see page 1461), you can build an image that has support to use NVIDIA GPU hardware to support GPU workloads.

NOTE: The NVIDIA driver requires a specific Linux kernel version. Make sure that the base image for the OS version has the required kernel version.

See [Supported Infrastructure Operating Systems](#) (see page 69) for a list of OS versions and the corresponding kernel versions known to work with the NVIDIA driver.

If the [NVIDIA runfile](#)³²⁵ installer has not been downloaded, then retrieve and install the download first by running the following command. The first line in the command below downloads and installs the runfile and the second line places it in the artifacts directory (you must create an `artifacts` directory if you don't already have one).

```
curl -O https://download.nvidia.com/XFree86/Linux-x86_64/470.82.01/NVIDIA-Linux-x86_64-470.82.01.run
mv NVIDIA-Linux-x86_64-470.82.01.run artifacts
```

DKP supported [NVIDIA driver](#)³²⁶ version is 470.x.

To build an image for use on GPU enabled hardware, perform the following steps.

1. In your `overrides/nvidia.yaml` file, add the following to enable GPU builds. You can also access and use the overrides [repo](#)³²⁷ or in the documentation under [Nvidia GPU Override File](#) (see page 1518) or [Offline Nvidia Override file](#) (see page 1519).
 - a. Non-air-gapped GPU override:

```
gpu:
  types:
    - nvidia
  build_name_extra: "-nvidia"
```

³²⁵ <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

³²⁶ <https://www.nvidia.com/Download/Find.aspx>

³²⁷ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

b. Air-gapped GPU override:

```
# Use this file when building a machine image, not as a override secret
for preprovisioned environments
nvidia_runfile_local_file: "{{ playbook_dir }}/../artifacts/
{{ nvidia_runfile_installer }}"
gpu:
  types:
    - nvidia

build_name_extra: "-nvidia"
```

NOTE: For [RHEL Pre-provisioned Override Files](#) (see page 1524) used with KIB, see specific note for GPU.

- Build your image using the `build` Konvoy Image Builder command, making sure to include the flag `--instance-type` that specifies an AWS instance that has an available GPU:

AWS Example:

```
konvoy-image build --region us-east-1 --instance-type=p2.xlarge --source-
ami=ami-12345abcdef images/ami/centos-7.yaml --overrides overrides/nvidia.yaml
```

In this example, we chose an instance type with an NVIDIA GPU using the `--instance-type` flag, and we provided the NVIDIA overrides using the `--overrides` flag. See [KIB with AWS](#) (see page 1465) for more information on creating an AML.

Additional helpful information can be found in the [NVIDIA Device Plug-in](#)³²⁸ for Kubernetes instructions and the [Installation Guide of Supported Platforms](#)³²⁹.

See also: [NVIDIA documentation](#)³³⁰

5.11.5.5 Next Steps:

[AWS GPU: Create the Management Cluster](#) (see page 368)

5.11.5.6 AWS GPU: Create the Management Cluster

Use this procedure to create a self-managed AWS GPU Management cluster with DKP. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on


³²⁸ <https://github.com/NVIDIA/k8s-device-plugin/blob/master/README.md>

³²⁹ <https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/install-guide.html>

³³⁰ https://nvidia.custhelp.com/app/answers/detail/a_id/131/kw/driver%20installation%20docs/related/1

the same cluster they are managing. After the GPU compatible image is created with KIB, a cluster can be generated using that custom AMI.

5.11.5.6.1 Name Your Cluster

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)³³¹ for more naming information.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.


Follow these steps:

1. Give your cluster a unique name suitable for your environment.
In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.
2. Set the environment variable:

```
export CLUSTER_NAME=<aws-example>
```

5.11.5.6.2 Create a New AWS Kubernetes Cluster

If you use these instructions to create a cluster on AWS using the DKP default settings without any edits to configuration files or additional flags, your cluster is deployed on an [Ubuntu 20.04 operating system image](#)^{332,333} with 3 control plane nodes, and 4 worker nodes.

 In previous DKP releases, AMI images provided by the upstream CAPA project would be used if you did not specify an AMI. However, the upstream images are not recommended for production and may not always be available. Therefore, DKP now requires you to specify an AMI when creating a cluster. To create an AMI, use [Konvoy Image Builder](#) (see page 1465).

331 <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

332 <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273747774&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems>

333 <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273747774&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems>

There are two approaches to supplying the ID of your AMI. Either provide the ID of the AMI or provide a way for DKP to discover the AMI using location, format and OS information:

1. **Option One** - Provide the ID of your AMI:
 - a. Use the example command below leaving the existing flag that provides the AMI ID: `--ami AMI_ID`
2. **Option Two** - Provide a path for your AMI with the information required for image discover:
 - a. Where the AMI is published using your AWS Account ID: `--ami-owner AWS_ACCOUNT_ID`
 - b. The format or string used to search for matching AMIs and ensure it references the Kubernetes version plus the base OS name: `--ami-base-os ubuntu-20.04`
 - c. The base OS information: `--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*'`

- (Optional) Configure your cluster to use an existing local registry as a mirror when attempting to pull images. Below is an AWS ECR example:

⚠ IMPORTANT: The AMI must be created with [Konvoy Image Builder](#) (see page 1461) in order to use the registry mirror feature.

```
export REGISTRY_MIRROR_URL=<your_registry_url>
```

- `REGISTRY_MIRROR_URL` : the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.

⚠ DKP uses AWS CSI as the [default storage provider](#) (see page 91). You can use a [Kubernetes CSI](#)³³⁴ compatible storage solution that is suitable for production. See the Kubernetes documentation called [Changing the Default Storage Class](#)³³⁵ for more information.

1. Run the **Option One** command as explained above to create a cluster with a GPU AMI and `--self-managed` flag:

```
dkp create cluster aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--ami AMI_ID \
--self-managed
```

³³⁴ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

³³⁵ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

OR

Run the **Option Two** command as explained above to create a cluster with a GPU AMI providing the location, format and base OS:

```
dkp create cluster aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--ami-owner AWS_ACCOUNT_ID \
--ami-base-os ubuntu-20.04 \
--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*' \
--self-managed
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1039\)](#).

2. Create the node pool after cluster creation:

```
dkp create nodepool aws -c ${CLUSTER_NAME} \
--instance-type p2.xlarge \
--ami-id=${AMI_ID_FROM_KIB} \
--replicas=1 ${NODEPOOL_NAME} \
--kubeconfig=${CLUSTER_NAME}.conf
```

You can find a customizable [Create a New AWS Cluster \(see page 1071\)](#) under [Custom Installation and Additional Infrastructure Tools \(see page 1036\)](#).

5.11.5.6.3 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>³³⁶³³⁷.

5.11.5.6.4 Next Step:

[AWS GPU: Install Kommander \(see page 371\)](#)

³³⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

³³⁷ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

5.11.5.7 AWS GPU: Install Kommander

You have installed the Konvoy component and created a cluster. Now it is time to Install Kommander which will allow you to access the UI and attach new or existing clusters to monitor.

5.11.5.7.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.11.5.7.2 Create and customize your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```


2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init > kommander.yaml
```

4. *If required:* Customize your `kommander.yaml`.

 See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, [Rook Ceph customization \(Pre-provisioned envs\)](#) (see page 1405), etc.

5. *If required:* If your cluster uses a custom **AWS VPC** and requires an internal load-balancer, set the `traefik` annotation to create an internal-facing ELB:

```
apps:
  traefik:
    enabled: true
    values: |
      service:
        annotations:
```

```
service.beta.kubernetes.io/aws-load-balancer-internal: "true"
```

5.11.5.7.3 Enable GPU Resources

1. In the same `kommander.yaml` file, enable Nvidia platform services.

```
apps:
  nvidia-gpu-operator:
    enabled: true
```

2. Append the correct Toolkit version based on your OS:

The **NVIDIA Container Toolkit** allows users to run GPU accelerated containers. The toolkit includes a container runtime library and utilities to automatically configure containers to leverage NVIDIA GPU and must be configured correctly according to your base operating system.

Centos 7.9/RHEL 7.9:

If you're using Centos 7.9 or RHEL 7.9 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-centos7
```

RHEL 8.4/8.6 and SLES 15 SP3

If you're using RHEL 8.4/8.6 or SLES 15 SP3 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubi8
```


Ubuntu 18.04 and 20.04

If you're using Ubuntu 18.04 or 20.04 as the base operating system for your GPU enabled nodes, set

the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:


```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubuntu20.04
```

5.11.5.7.4 Enable DKP Catalog Applications and Install Kommander

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
        ref:
          tag: v2.6.1
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

[-] Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.11.5.7.5 Next Step:

[AWS GPU: Verify Install and Log in the UI](#) (see page 378)

5.11.5.8 AWS GPU: Verify Install and Log in the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.

[-] NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
```

```

helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met

```

5.11.5.8.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```

kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'

```

5.11.5.8.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{\n}Password: {{.data.password|base64decode}}
{\n}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{\nor .hostname .ip}}{\end}}/dkp/kommander/
dashboard{\n}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQL2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.11.5.8.3 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

For more [GPU information](#) (see [page 968](#)), see the section for GPU under Day 2 - Cluster Operations Management.

5.11.5.8.4 Next Step:

[AWS GPU: Subsequent New Clusters](#) (see [page 377](#))

5.11.5.9 AWS GPU: Create a Managed Cluster Using the DKP CLI

Enterprise

Gov Advanced

In previous steps, you created a new Management cluster, which is self-managed. When you use these steps to create new Managed clusters, they will become Attached clusters under your Management Cluster.



When creating [Managed clusters](#) (see [page 81](#)), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see [page 81](#))!

5.11.5.9.1 Create New Managed GPU Cluster with the CLI



The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29920384/Attach+a+Kubernetes+Cluster>. (see [page 752](#)).

Your new, managed cluster needs to be part of either a new Workspace or an existing Workspace under a management cluster. Create or locate a workspace name to get started with your managed cluster.

5.11.5.9.2 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:

⚠ NOTE: If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see [page 650](#)).

```
kubectl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.11.5.9.3 Name Your Cluster


Follow these steps:

1. Give your cluster a unique name suitable for your environment.

In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable:

```
export MANAGED_CLUSTER_NAME=<aws-additional>
```

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes³³⁸](#) for more naming information.

3. Execute this command to create your Kubernetes cluster using any relevant flags but note that the `--self-managed` is absent so that this cluster can be attached to your previously created Management Cluster.

- Execute this command to create a cluster with a GPU AMI:

```
dkp create cluster aws \
--cluster-name=${MANAGED_CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--kubeconfig=<management-cluster-kubeconfig-path> \
--namespace ${WORKSPACE_NAMESPACE} \
--with-aws-bootstrap-credentials=true \
--kubeconfig=<management-cluster-kubeconfig-path> \
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful.

More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

- Create the node pool after cluster creation:

```
dkp create nodepool aws -c ${CLUSTER_NAME} \
--instance-type p2.xlarge \
--ami-id=${AMI_ID_FROM_KIB} \
--replicas=1 ${NODEPOOL_NAME} \
--kubeconfig=${CLUSTER_NAME}.conf \
```

³³⁸ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

5.11.5.9.4 Retrieve the kubeconfig and Explore New AWS Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:



NOTE: Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.11.5.9.4.1 Manually Attach a DKP CLI Cluster to the Management Cluster



These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

- Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

- You can now either [\[attach it in the UI\]](#)(link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

- Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

- Set the WORKSPACE_NAMESPACE environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

- You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

- This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new attached-cluster-kubeconfig.yaml file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

- Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace ${WORKSPACE_NAMESPACE}
```

- Create this `kommandercluster` object to attach the cluster to the workspace:


```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubectl get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.11.5.9.4.2 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.11.6 AWS Air-gapped with GPU Install

This section provides instructions to install DKP in an AWS air-gapped GPU environment.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)



For air-gapped, ensure you have [Download DKP](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, so you can extract the tarball on next page.

5.11.6.1 AWS Prerequisites

Before you begin using Konvoy with AWS, you must:

1. Follow the steps to create a [Minimal Permissions and Role to Create Clusters](#) (see page 1057)..
2. Create [Cluster IAM Policies and Roles](#) (see page 1063).
3. Export the AWS region where you want to deploy the cluster:

```
export AWS_REGION=us-west-2
```

4. Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```

5.11.6.2 GPU Prerequisites

Before you begin, you must:

- Ensure nodes provide an NVIDIA GPU.
- If you are using a public cloud service such as [AWS](#) (see page 1110), create an AMI with KIB using the instructions on the [KIB for GPU](#) (see page 1485) page.
- If you are deploying in a pre-provisioned environment, ensure that you have created the appropriate [secret for your GPU nodepool](#) (see page 240) and have uploaded the appropriate artifacts to each node. See the [GPU only steps section](#) (see page 0) on the Pre-provisioned Prerequisites Air-gapped page for additional information.



Specific instructions must be followed for enabling `nvidia-gpu-operator` depending on if you want to deploy the app on a Management cluster or a Attached or a Managed cluster.

- For instructions on enabling the NVIDIA platform application on a Management cluster, follow the instructions in the [NVIDIA Platform Application Management Cluster](#) (see page 970) section.
- For instructions on enabling the NVIDIA platform application on attached or managed clusters, follow the instructions in the [NVIDIA Platform Application Attached or Managed Cluster](#) (see page 972) section.

After `nvidia-gpu-operator` has been enabled depending on the cluster type, proceed to the **Select the correct Toolkit version for your NVIDIA GPU Operator** on each of those pages.

5.11.6.3 Next Step:

[AWS Air-gapped GPU: Use Node Label Automatic Configuration](#) (see page 383)

5.11.6.4 AWS Air-gapped GPU: Use Node Label Automatic Configuration

When using GPU nodes, it is important they have the proper label identifying them as Nvidia GPU nodes. Node feature discovery (NFD), by default labels PCI hardware as:

```
"feature.node.kubernetes.io/pci-<device label>.present": "true"
```

where `<device label>` is by default as [defined in this topic](#)³³⁹:

```
< class > _ < vendor >
```

However, because there is a wide variety in devices and their assigned PCI classes, you may find that the labels assigned to your GPU nodes do not always properly identify them as containing an Nvidia GPU.

If the default detection does not work, you can manually change the daemonset that the GPU operator creates by running the following command:

```
nodeSelector:
  feature.node.kubernetes.io/pci-< class > _ < vendor>.present: "true"
```

where `class` is any 4 digit number starting with `03xy` and the vendor for Nvidia is `10de`. If this is already deployed, you can always change the `daemonset` and change the `nodeSelector` field so that it deploys to the right nodes.

5.11.6.4.1 Next Step:

[AWS Air-gapped GPU: Create an Image \(see page 384\)](#)


5.11.6.5 AWS Air-gapped GPU: Create an Image

5.11.6.5.1 Create an Image using Konvoy Image Builder (KIB) for use in an air-gapped cluster

5.11.6.5.2 Prerequisites

- [Minimal IAM Permissions for KIB \(see page 1465\)](#) to create an Image for an AWS account using Konvoy Image Builder.

³³⁹ <https://kubernetes-sigs.github.io/node-feature-discovery/v0.7/get-started/features.html#pci>

 The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to create an AWS Air-gapped AMI to take advantage of enhanced cluster operations. Explore the [KIB with AWS](#) (see page 1465) topics for more options.

Using [KIB](#) (see page 1461), you can build an AMI without requiring access to the internet by providing an additional `--override` flag.


1. Assuming you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2/kib
```

2. Follow the instructions below to build an AMI.

5.11.6.5.3 Build the GPU Image

Using the [Konvoy Image Builder](#) (see page 1461), you can build an image that has support to use NVIDIA GPU hardware to support GPU workloads.

 **NOTE:** The NVIDIA driver requires a specific Linux kernel version. Make sure that the base image for the OS version has the required kernel version.

See [Supported Infrastructure Operating Systems](#) (see page 69) for a list of OS versions and the corresponding kernel versions known to work with the NVIDIA driver.

If the [NVIDIA runfile](#)³⁴⁰ installer has not been downloaded, then retrieve and install the download first by running the following command. The first line in the command below downloads and installs the runfile and the second line places it in the artifacts directory (you must create an `artifacts` directory if you don't already have one).

```
curl -O https://download.nvidia.com/XFree86/Linux-x86_64/470.82.01/NVIDIA-Linux-x86_64-470.82.01.run
mv NVIDIA-Linux-x86_64-470.82.01.run artifacts
```

³⁴⁰ <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

 DKP supported [NVIDIA driver](#)³⁴¹ version is 470.x.

To build an image for use on GPU enabled hardware, perform the following steps.

1. In your `overrides/nvidia.yaml` file, add the following to enable GPU builds. You can also access and use the overrides [repo](#)³⁴² or in the documentation under [Nvidia GPU Override File](#) (see page 1518) or [Offline Nvidia Override file](#) (see page 1519).

- a. Non-air-gapped GPU override:

```
gpu:
  types:
    - nvidia
  build_name_extra: "-nvidia"
```

- b. Air-gapped GPU override:

```
# Use this file when building a machine image, not as a override secret
for preprovisioned environments
nvidia_runfile_local_file: "{{ playbook_dir }}/../artifacts/
{{ nvidia_runfile_installer }}"
gpu:
  types:
    - nvidia

build_name_extra: "-nvidia"
```

NOTE: For [RHEL Pre-provisioned Override Files](#) (see page 1524) used with KIB, see specific note for GPU.

2. Build your image using the `build` Konvoy Image Builder command, making sure to include the flag `--instance-type` that specifies an AWS instance that has an available GPU:

AWS Example:

```
konvoy-image build --region us-east-1 --instance-type=p2.xlarge --source-ami=ami-12345abcdef images/ami/centos-7.yaml --overrides overrides/nvidia.yaml
```

In this example, we chose an instance type with an NVIDIA GPU using the `--instance-type` flag, and we provided the NVIDIA overrides using the `--overrides` flag. See [KIB with AWS](#) (see page 1465) for more information on creating an AML.

³⁴¹ <https://www.nvidia.com/Download/Find.aspx>

³⁴² <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

Additional helpful information can be found in the [NVIDIA Device Plug-in](#)³⁴³ for Kubernetes instructions and the [Installation Guide of Supported Platforms](#)³⁴⁴.

See also: [NVIDIA documentation](#)³⁴⁵

5.11.6.5.4 Next Steps:

[AWS Air-gapped GPU: Load the Registry](#) (see page 386)

5.11.6.6 AWS Air-gapped GPU: Load the Registry

After you create an image for your air-gapped environment, you will need to load it into your registry.

5.11.6.6.1 Load Images into your Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine](#) (see page 1285) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.



If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see page 1533) page for more information.

1. [Download](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, and extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

3. Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```

³⁴³ <https://github.com/NVIDIA/k8s-device-plugin/blob/master/README.md>

³⁴⁴ <https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/install-guide.html>

³⁴⁵ https://nvidia.custhelp.com/app/answers/detail/a_id/131/kw/driver%20installation%20docs/related/1

ⓘ It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

5.11.6.6.2 Kommander Load Images

If you are operating in an air-gapped environment, a [local container registry](#) (see [page 1533](#)) containing all the necessary installation images, including the Kommander images is required. See below for how to push the necessary images to this registry.

5.11.6.6.3 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.11.6.6.4 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.11.6.6.5 Next Step:

[AWS Air-gapped GPU: Create the Management Cluster](#) (see [page 388](#))

5.11.6.7 AWS Air-gapped GPU: Create the Management Cluster

Use this procedure to create a self-managed air-gapped AWS GPU Management cluster with DKP. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing. To create a cluster in an AWS Air-gapped environment using GPUs, execute the following:

- = To increase [Docker Hub's rate limit](#)³⁴⁶, use your Docker Hub credentials when creating the cluster by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

1. Give your cluster a unique name suitable for your environment.

In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable to the name you assigned this cluster:

```
export CLUSTER_NAME=<aws-example>
```

NOTE: The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)³⁴⁷ for more naming information.

3. Export variables for the existing infrastructure details:

```
export AWS_VPC_ID=<vpc-...>
export AWS_SUBNET_IDS=<subnet-...,subnet-...,subnet-...>
export AWS_ADDITIONAL_SECURITY_GROUPS=<sg-...>
export AWS_AMI_ID=<ami-...>
```

- `AWS_VPC_ID` : the VPC ID where the cluster will be created. The VPC requires the following AWS VPC Endpoints to be already present:
 - `ec2` - `com.amazonaws.{region}.ec2`
 - `elasticloadbalancing` - `com.amazonaws.{region}.elasticloadbalancing`
 - `secretsmanager` - `com.amazonaws.{region}.secretsmanager`
 - `autoscaling` - `com.amazonaws.{region}.autoscaling`
 - `ecr` - `com.amazonaws.{region}.ecr.api` - (authentication)
 - `ecr` - `com.amazonaws.{region}.ecr.dkr` - (data transfer)

³⁴⁶ <https://docs.docker.com/docker-hub/download-rate-limit/>

³⁴⁷ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

More details about [AWS service using an interface VPC endpoint](#)³⁴⁸ and [AWS VPC endpoints list](#)³⁴⁹.

- `AWS_SUBNET_IDS` : a comma-separated list of one or more private Subnet IDs with each one in a different Availability Zone. The cluster control-plane and worker nodes will automatically be spread across these Subnets.
- `AWS_ADDITIONAL_SECURITY_GROUPS` : a comma-separated list of one or more Security Groups IDs to use in addition to the ones automatically created by [CAPA](#)³⁵⁰.
- `AWS_AMI_ID` : the AMI ID to use for control-plane and worker nodes. The AMI must be created by the [konvoy-image-builder](#) (see page 1461).



In previous DKP releases, AMI images provided by the upstream CAPA project would be used if you did not specify an AMI. However, the upstream images are not recommended for production and may not always be available. Therefore, DKP now requires you to specify an AMI when creating a cluster. To create an AMI, use [Konvoy Image Builder](#) (see page 1465).

There are two approaches to supplying the ID of your AMI while creating your cluster. **Option One** is to provide the ID of the AMI. **Option Two** is provide a way for DKP to discover the AMI using location, format and OS information using flags.

! IMPORTANT: You must tag the subnets as described below to allow for Kubernetes to create External Load Balancers (ELBs) for services of type `LoadBalancer` in those subnets. If the subnets are not tagged, they will not receive an ELB and the following error displays: `Error syncing load balancer, failed to ensure load balancer; could not find any suitable subnets for creating the ELB. .`

The tags should be set as follows, where `<CLUSTER_NAME>` corresponds to the name set in `CLUSTER_NAME` environment variable:

```
kubernetes.io/cluster = <CLUSTER_NAME>
kubernetes.io/cluster/<CLUSTER_NAME> = owned
kubernetes.io/role/internal-elb = 1
```

4. (Optional) Configure your cluster to use an existing container registry as a mirror when attempting to pull images. The example below is for AWS ECR:

! If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see page 1533) page for more information.

³⁴⁸ <https://docs.aws.amazon.com/vpc/latest/privatelink/create-interface-endpoint.html>

³⁴⁹ <https://docs.aws.amazon.com/vpc/latest/privatelink/aws-services-privatelink-support.html>

³⁵⁰ <https://github.com/kubernetes-sigs/cluster-api-provider-aws>

⚠ IMPORTANT: The AMI must be created by the [konvoy-image-builder](#) (see page 1461) project in order to use the registry mirror feature.

```
export REGISTRY_MIRROR_URL=<your_registry_url>
```

- `REGISTRY_MIRROR_URL` : the address of an existing registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
 - NOTE: Other local registries may use the options below:
 - JFrog - `REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
 - `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
 - `REGISTRY_PASSWORD` : optional if username is not set.
5. Create a Kubernetes cluster. The following example shows a common configuration. See [dkp create cluster aws](#) (see page 1665) reference for the full list of cluster creation options:

⚠ DKP uses AWS CSI as the [default storage provider](#) (see page 91). You can use a [Kubernetes CSI](#)³⁵¹ compatible storage solution that is suitable for production. See the Kubernetes documentation called [Changing the Default Storage Class](#)³⁵² for more information.

6. Run the **Option One** command as explained above to create a cluster with a GPU AMI and `--self-managed` flag:

```
dkp create cluster aws
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--vpc-id=${AWS_VPC_ID} \
--ami=${AWS_AMI_ID} \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
--registry-mirror-password=${REGISTRY_PASSWORD} \
--self-managed
```


³⁵¹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

³⁵² <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

OR

Run the **Option Two** command as explained above to create a cluster with a GPU AMI providing the location, format and base OS:

```
dkp create cluster aws
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--vpc-id=${AWS_VPC_ID} \
--ami-owner AWS_ACCOUNT_ID \
--ami-base-os ubuntu-20.04 \
--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*' \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
--registry-mirror-password=${REGISTRY_PASSWORD} \
--self-managed
```

 If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

7. After cluster creation, create the node pool after cluster creation:

```
dkp create nodepool aws -c ${CLUSTER_NAME} \
--instance-type p2.xlarge \
--ami-id=${AMI_ID_FROM_KIB} \
--replicas=1 ${NODEPOOL_NAME} \
--kubeconfig=${CLUSTER_NAME}.conf
```

To understand how this process works step by step, you can find a customizable [Create a New AWS Cluster](#) (see page 1071) under [Custom Installation and Additional Infrastructure Tools](#) (see page 1036).

5.11.6.7.1 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>³⁵³³⁵⁴.

5.11.6.7.2 Next Step:

[AWS Air-gapped GPU: Install Kommander](#) (see page 393)

5.11.6.8 AWS Air-gapped GPU: Install Kommander

You have installed the Konvoy component and created a cluster. Now it is time to Install Kommander which will allow you to access the UI and attach new or existing clusters to monitor.

5.11.6.8.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Ensure you have loaded all necessary images for your configuration. See [Load the Images into Your Registry: Air-gapped Environments](#) (see page 1400).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.11.6.8.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init --airgapped > kommander.yaml
```

³⁵³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

³⁵⁴ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

4. If required, **customize** your `kommander.yaml`.

i See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include:

- Custom Domains and Certificates
- HTTP proxy
- External Load Balancer
- GPU utilization, etc.
- [Rook Ceph customization for Pre-provisioned environments](#) (see page 1405)

5. If required: If your cluster uses a custom **AWS VPC** and requires an internal load-balancer, set the `traefik` annotation to create an internal-facing ELB:

```
...
apps:
  traefik:
    enabled: true
    values: |
      service:
        annotations:
          service.beta.kubernetes.io/aws-load-balancer-internal: "true"
...
```

6. Expand **one** of the following sets of instructions, depending on your license and application environments:

5.11.6.8.3 Enable GPU Resources

1. Append the following to the apps section in the `kommander.yaml` file to enable Nvidia platform services.

```
apps:
  nvidia-gpu-operator:
    enabled: true
```

2. Append the correct Nvidia Toolkit version based on your OS:

The **NVIDIA Container Toolkit** allows users to run GPU accelerated containers. The toolkit includes a container runtime library and utilities to automatically configure containers to leverage NVIDIA GPU and must be configured correctly according to your base operating system.

Centos 7.9/RHEL 7.9:

If you're using Centos 7.9 or RHEL 7.9 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your [Kommander Installer Configuration file](#) (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-centos7
```

RHEL 8.4/8.6 and SLES 15 SP3

If you're using RHEL 8.4/8.6 or SLES 15 SP3 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:


```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubi8
```

Ubuntu 18.04 and 20.04

If you're using Ubuntu 18.04 or 20.04 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:


```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubuntu20.04
```

5.11.6.8.4 Enable DKP Catalog Applications and Install Kommander in an Air-gapped Environment

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add the following values to enable DKP Catalog Applications:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      path: ./dkp-catalog-applications-v2.6.1.tar.gz
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubecfg=${CLUSTER_NAME}.conf \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-
v2.6.1.tar.gz
```

 **Tips and recommendations**

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File \(see page 87\)](#).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.11.6.8.5 Next Step:

[AWS Air-gapped GPU: Verify Install and Log in the UI \(see page 396\)](#)

5.11.6.9 AWS Air-gapped GPU: Verify Install and Log in the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.



NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
```



```

helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met

```

5.11.6.9.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```

kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'

```

5.11.6.9.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{{ "\n"}}Password: {{.data.password|base64decode}}
{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/
dashboard{{ "\n"}}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.11.6.9.3 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

For more [GPU information](#) (see [page 968](#)), see the section for GPU under Day 2 - Cluster Operations Management.

5.11.6.9.4 Next Step:


[AWS Air-gapped GPU: Subsequent New Clusters](#) (see [page 399](#))

5.11.6.10 AWS Air-gapped GPU: Create a Managed Cluster Using the DKP CLI


Enterprise

Gov Advanced

In previous steps, you created a new Management cluster, which is self-managed. When you use these steps to create new Managed clusters, they will become Attached clusters under your Management Cluster.

 When creating [Managed clusters](#) (see page 81), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see page 81)!

5.11.6.10.1 Create New Managed Air-gapped GPU Cluster with the CLI


 The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29920384/Attach+a+Kubernetes+Cluster>. (see page 752).

Your new, managed cluster needs to be part of either a new Workspace or an existing Workspace under a management cluster. Create or locate a workspace name to get started with your managed cluster.

5.11.6.10.2 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:

 **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see page 650).

```
kubect1 get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.11.6.10.3 Name Your Cluster


Follow these steps:

1. Give your cluster a unique name suitable for your environment.

In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable:


```
export MANAGED_CLUSTER_NAME=<aws-additional>
```

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)³⁵⁵ for more naming information.

3. Execute this command to create your Kubernetes cluster using any relevant flags but note that the `--self-managed` is absent so that this cluster can be attached to your previously created Management Cluster.

- Execute this command to create a cluster with a GPU AMI:

```
dkp create cluster aws --cluster-name=${MANAGED_CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--namespace ${WORKSPACE_NAMESPACE} \
--with-aws-bootstrap-credentials=true \
--vpc-id=${AWS_VPC_ID} \
--ami=${AWS_AMI_ID} \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
--registry-mirror-password=${REGISTRY_PASSWORD} \
--kubeconfig=<management-cluster-kubeconfig-path> \
```

 If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

- Create the node pool after cluster creation:

³⁵⁵ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

```
dkp create nodepool aws -c ${CLUSTER_NAME} \
--instance-type p2.xlarge \
--ami-id=${AMI_ID_FROM_KIB} \
--replicas=1 ${NODEPOOL_NAME} \
--kubeconfig=${CLUSTER_NAME}.conf
```

5.11.6.10.4 Retrieve the kubeconfig and Explore New AWS Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:



NOTE: Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.11.6.10.4.1 Manually Attach a DKP CLI Cluster to the Management Cluster



These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

7. This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

8. Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace ${WORKSPACE_NAMESPACE}
```

9. Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubectl get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.11.6.10.4.2 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.12 EKS Install Options

Enterprise

Gov Advanced

For an environment that is EKS on the AWS Infrastructure, install options are provided for you in this one location. Remember, there are always more options in the [EKS Infrastructure](#) (see page 1127) section under [Custom Installation and Additional Infrastructure Tools](#) (see page 1036), but this will get you operative in the most common scenario.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)

- [Prerequisites for Install](#) (see page 121)

⚠ An EKS cluster cannot be a [Management](#) (see page 82) or [Essential](#) (see page 83) cluster. To install DKP on your EKS cluster, first ensure you have a Management cluster with DKP and the Kommander component installed, that handles the lifecycle of your EKS cluster.

Below are the supported environment variables. (Refer to this page to check if your OS is a [Supported Operating System](#) (see page 69).)

- [EKS Install](#) (see page 405)

i In order to install Kommander, you need to have CAPI components, cert-manager, etc on a self-managed cluster. The CAPI components mean you can control the lifecycle of the cluster, and other clusters. However, because EKS is semi-managed by AWS, the EKS clusters are under AWS control and don't have those components. Therefore, Kommander will not be installed.

5.12.1 EKS Install

Enterprise

Gov Advanced

This section provides instructions to install DKP with EKS in an AWS non-air-gapped environment.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

⚠ Ensure that the `KUBECONFIG` environment variable is set to the Management cluster by running `export KUBECONFIG=<Management_cluster_kubeconfig>.conf`.

5.12.1.1 AWS prerequisites

Before you begin using DKP with AWS, you must have:

- A Management cluster with the Kommander component installed.

- An EKS cluster cannot be a [Management](#) (see page 82) or [Essential](#) (see page 83) cluster. To install DKP on your EKS cluster, first ensure you have a Management cluster with DKP and the Kommander component installed, that handles the lifecycle of your EKS cluster.

- You have a valid AWS account with [credentials configured](#)³⁵⁶ that can manage CloudFormation Stacks, IAM Policies, and IAM Roles.
- You will need to have the [AWS CLI utility installed](#)³⁵⁷.
- Install [aws-iam-authenticator](#)³⁵⁸. This binary is used to access your cluster using kubectl.

- ⓘ In order to install Kommander, you need to have CAPI components, cert-manager, etc on a self-managed cluster. The CAPI components mean you can control the lifecycle of the cluster, and other clusters. However, because EKS is semi-managed by AWS, the EKS clusters are under AWS control and don't have those components. Therefore, Kommander will not be installed and these clusters will be attached to the management cluster.

5.12.1.2 Next Step:

[EKS: Minimal User Permission for Cluster Creation](#) (see page 406)

5.12.1.3 EKS: Minimal User Permission for Cluster Creation

Enterprise

Gov Advanced

The following is a cloudformation stack which adds a policy named `eks-bootstrapper` to manage EKS cluster to the `dkp-bootstrapper-role` created by the cloudformation stack in the [Minimal Permissions and Role to Create Cluster](#) (see page 1057) section. Consult the [Leveraging the Role](#) (see page 1061) section for an example of how to use this role and how a system administrator wants to expose using the permissions.

```
AWSTemplateFormatVersion: 2010-09-09
Parameters:
  existingBootstrapperRole:
    Type: CommaDelimitedList
```

³⁵⁶ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

³⁵⁷ <https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html>

³⁵⁸ <https://docs.aws.amazon.com/eks/latest/userguide/install-aws-iam-authenticator.html>

Description: 'Name of existing minimal role you want to add to add EKS cluster management permissions to'

Default: dkp-bootstrapper-role

Resources:

EKSMinimumPermissions:

Properties:

Description: Minimal user policy to manage eks clusters

ManagedPolicyName: eks-bootstrapper

PolicyDocument:

Statement:

- Action:

- 'ssm:GetParameter'

Effect: Allow

Resource:

- 'arn:*:ssm:*:*:parameter/aws/service/eks/optimized-ami/*'

- Action:

- 'iam:CreateServiceLinkedRole'

Condition:

StringLike:

'iam:AWSServiceName': eks.amazonaws.com

Effect: Allow

Resource:

- >-

arn:*:iam:*:*:role/aws-service-role/eks.amazonaws.com/

AWSServiceRoleForAmazonEKS

- Action:

- 'iam:CreateServiceLinkedRole'

Condition:

StringLike:

'iam:AWSServiceName': eks-nodegroup.amazonaws.com

Effect: Allow

Resource:

- >-

arn:*:iam:*:*:role/aws-service-role/eks-nodegroup.amazonaws.com/

AWSServiceRoleForAmazonEKSNodegroup

- Action:

- 'iam:CreateServiceLinkedRole'

Condition:

StringLike:

'iam:AWSServiceName': eks-fargate.amazonaws.com

Effect: Allow

Resource:

- >-

arn:aws:iam:*:*:role/aws-service-role/eks-fargate-pods.amazonaws.com/

AWSServiceRoleForAmazonEKSFargate

- Action:

- 'iam:GetRole'

- 'iam:ListAttachedRolePolicies'

Effect: Allow

Resource:

- 'arn:*:iam:*:*:role/*'

- Action:

```

    - 'iam:GetPolicy'
Effect: Allow
Resource:
  - 'arn:aws:iam::aws:policy/AmazonEKSClusterPolicy'
- Action:
  - 'eks:DescribeCluster'
  - 'eks:ListClusters'
  - 'eks:CreateCluster'
  - 'eks:TagResource'
  - 'eks:UpdateClusterVersion'
  - 'eks>DeleteCluster'
  - 'eks:UpdateClusterConfig'
  - 'eks:UntagResource'
  - 'eks:UpdateNodegroupVersion'
  - 'eks:DescribeNodegroup'
  - 'eks>DeleteNodegroup'
  - 'eks:UpdateNodegroupConfig'
  - 'eks:CreateNodegroup'
  - 'eks:AssociateEncryptionConfig'
  - 'eks:ListIdentityProviderConfigs'
  - 'eks:AssociateIdentityProviderConfig'
  - 'eks:DescribeIdentityProviderConfig'
  - 'eks:DisassociateIdentityProviderConfig'
Effect: Allow
Resource:
  - 'arn*:eks*:*:cluster/*'
  - 'arn*:eks*:*:nodegroup/*/*/*'
- Action:
  - 'ec2:AssociateVpcCidrBlock'
  - 'ec2:DisassociateVpcCidrBlock'
  - 'eks:ListAddons'
  - 'eks:CreateAddon'
  - 'eks:DescribeAddonVersions'
  - 'eks:DescribeAddon'
  - 'eks>DeleteAddon'
  - 'eks:UpdateAddon'
  - 'eks:TagResource'
  - 'eks:DescribeFargateProfile'
  - 'eks:CreateFargateProfile'
  - 'eks>DeleteFargateProfile'
Effect: Allow
Resource:
  - '*'
- Action:
  - 'iam:PassRole'
Condition:
  StringEquals:
    'iam:PassedToService': eks.amazonaws.com
Effect: Allow
Resource:
  - '*'
- Action:

```

```

    - 'kms:CreateGrant'
    - 'kms:DescribeKey'
  Condition:
    'ForAnyValue:StringLike':
      'kms:ResourceAliases': alias/cluster-api-provider-aws-*
  Effect: Allow
  Resource:
    - '*'
  Version: 2012-10-17
  Roles: !Ref existingBootstrapperRole
  Type: 'AWS::IAM::ManagedPolicy'

```

ⓘ If your role is not named `dkp-bootstrapper-role` change the parameter on line 6 of the file.

To create the resources in the cloudformation stack, copy the contents above into a file. Before executing the following command, replace `MYFILENAME.yaml` and `MYSTACKNAME` with the intended values for your system:

```
aws cloudformation create-stack --template-body=file://MYFILENAME.yaml --stack-name=MYSTACKNAME --capabilities CAPABILITY_NAMED_IAM
```

5.12.1.3.1 Next Step:

[EKS: Cluster IAM Policies and Roles \(see page 409\)](#)

5.12.1.4 EKS: Cluster IAM Policies and Roles

Enterprise

Gov Advanced

This section guides a DKP user in creating IAM Policies and Instance Profiles that governs who has access to the cluster. The IAM Role is used by the cluster's control plane and worker nodes using the provided AWS CloudFormation Stack specific to EKS. This CloudFormation Stack has additional permissions that are used to delegate access roles for other users.

5.12.1.4.1 Prerequisites from AWS:

- The user you delegate from your role must have a minimum set of permissions, see [User Roles and Instance Profiles \(see page 1057\)](#) page for AWS.

- Create the [Cluster IAM Policies](#) (see page 1063) in your AWS account.

5.12.1.4.2 EKS IAM Artifacts

5.12.1.4.2.1 Policies

- `controllers-eks.cluster-api-provider-aws.sigs.k8s.io` - enumerates the Actions required by the workload cluster to create and modify EKS clusters in the user's AWS Account. It is attached to the existing `control-plane.cluster-api-provider-aws.sigs.k8s.io` role
- `eks-nodes.cluster-api-provider-aws.sigs.k8s.io` - enumerates the Actions required by the EKS workload cluster's worker machines. It is attached to the existing `nodes.cluster-api-provider-aws.sigs.k8s.io`

5.12.1.4.2.2 Roles

- `eks-controlplane.cluster-api-provider-aws.sigs.k8s.io` - is the Role associated with EKS cluster control planes



NOTE: `control-plane.cluster-api-provider-aws.sigs.k8s.io` and `nodes.cluster-api-provider-aws.sigs.k8s.io` roles were created by [Cluster IAM Policies and Roles](#) (see page 1063) in AWS.

Below is a [CloudFormation stack](#)³⁵⁹ that includes IAM policies and roles required to setup EKS Clusters:

```
AWSTemplateFormatVersion: 2010-09-09
Parameters:
  existingControlPlaneRole:
    Type: CommaDelimitedList
    Description: 'Names of existing Control Plane Role you want to add to the newly
created EKS Managed Policy for AWS cluster API controllers'
    Default: control-plane.cluster-api-provider-aws.sigs.k8s.io
  existingNodeRole:
    Type: CommaDelimitedList
    Description: 'ARN of the Nodes Managed Policy to add to the role for nodes'
    Default: nodes.cluster-api-provider-aws.sigs.k8s.io
Resources:
  AWSIAMManagedPolicyControllersEKS:
    Properties:
```

³⁵⁹ <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html>

Description: For the Kubernetes Cluster API Provider AWS Controllers
 ManagedPolicyName: controllers-eks.cluster-api-provider-aws.sigs.k8s.io
 PolicyDocument:

```
Statement:
  - Action:
    - 'ssm:GetParameter'
    Effect: Allow
    Resource:
    - 'arn:*:ssm:*:*:parameter/aws/service/eks/optimized-ami/*'
  - Action:
    - 'iam:CreateServiceLinkedRole'
    Condition:
      StringLike:
        'iam:AWSServiceName': eks.amazonaws.com
    Effect: Allow
    Resource:
    - >-
      arn:*:iam:*:*:role/aws-service-role/eks.amazonaws.com/
```

AWSServiceRoleForAmazonEKS

```
- Action:
  - 'iam:CreateServiceLinkedRole'
  Condition:
    StringLike:
      'iam:AWSServiceName': eks-nodegroup.amazonaws.com
  Effect: Allow
  Resource:
  - >-
    arn:*:iam:*:*:role/aws-service-role/eks-nodegroup.amazonaws.com/
```

AWSServiceRoleForAmazonEKSNodegroup

```
- Action:
  - 'iam:CreateServiceLinkedRole'
  Condition:
    StringLike:
      'iam:AWSServiceName': eks-fargate.amazonaws.com
  Effect: Allow
  Resource:
  - >-
    arn:aws:iam:*:*:role/aws-service-role/eks-fargate-pods.amazonaws.com/
```

AWSServiceRoleForAmazonEKSFargate

```
- Action:
  - 'iam:GetRole'
  - 'iam:ListAttachedRolePolicies'
  Effect: Allow
  Resource:
  - 'arn:*:iam:*:*:role/*'
- Action:
  - 'iam:GetPolicy'
  Effect: Allow
  Resource:
  - 'arn:aws:iam::aws:policy/AmazonEKSClusterPolicy'
- Action:
  - 'eks:DescribeCluster'
```

```

- 'eks:ListClusters'
- 'eks:CreateCluster'
- 'eks:TagResource'
- 'eks:UpdateClusterVersion'
- 'eks>DeleteCluster'
- 'eks:UpdateClusterConfig'
- 'eks:UntagResource'
- 'eks:UpdateNodegroupVersion'
- 'eks:DescribeNodegroup'
- 'eks>DeleteNodegroup'
- 'eks:UpdateNodegroupConfig'
- 'eks:CreateNodegroup'
- 'eks:AssociateEncryptionConfig'
- 'eks:ListIdentityProviderConfigs'
- 'eks:AssociateIdentityProviderConfig'
- 'eks:DescribeIdentityProviderConfig'
- 'eks:DisassociateIdentityProviderConfig'
Effect: Allow
Resource:
- 'arn:*:eks:*:*:cluster/*'
- 'arn:*:eks:*:*:nodegroup/*/*/*'
- Action:
- 'ec2:AssociateVpcCidrBlock'
- 'ec2:DisassociateVpcCidrBlock'
- 'eks:ListAddons'
- 'eks:CreateAddon'
- 'eks:DescribeAddonVersions'
- 'eks:DescribeAddon'
- 'eks>DeleteAddon'
- 'eks:UpdateAddon'
- 'eks:TagResource'
- 'eks:DescribeFargateProfile'
- 'eks:CreateFargateProfile'
- 'eks>DeleteFargateProfile'
Effect: Allow
Resource:
- '*'
- Action:
- 'iam:PassRole'
Condition:
StringEquals:
  'iam:PassedToService': eks.amazonaws.com
Effect: Allow
Resource:
- '*'
- Action:
- 'kms:CreateGrant'
- 'kms:DescribeKey'
Condition:
  'ForAnyValue:StringLike':
    'kms:ResourceAliases': alias/cluster-api-provider-aws-*
Effect: Allow

```

```

    Resource:
      - '*'
    Version: 2012-10-17
    Roles: !Ref existingControlPlaneRole
    Type: 'AWS::IAM::ManagedPolicy'
AWSIAMManagedEKSNodesPolicy:
  Properties:
    Description: Additional Policies to nodes role to work for EKS
    ManagedPolicyName: eks-nodes.cluster-api-provider-aws.sigs.k8s.io
    PolicyDocument:
      Statement:
        - Action:
            - "ec2:AssignPrivateIpAddresses"
            - "ec2:AttachNetworkInterface"
            - "ec2:CreateNetworkInterface"
            - "ec2>DeleteNetworkInterface"
            - "ec2:DescribeInstances"
            - "ec2:DescribeTags"
            - "ec2:DescribeNetworkInterfaces"
            - "ec2:DescribeInstanceTypes"
            - "ec2:DetachNetworkInterface"
            - "ec2:ModifyNetworkInterfaceAttribute"
            - "ec2:UnassignPrivateIpAddresses"
          Effect: Allow
          Resource:
            - '*'
        - Action:
            - ec2:CreateTags
          Effect: Allow
          Resource:
            - arn:aws:ec2:*:*:network-interface/*
        - Action:
            - "ec2:DescribeInstances"
            - "ec2:DescribeInstanceTypes"
            - "ec2:DescribeRouteTables"
            - "ec2:DescribeSecurityGroups"
            - "ec2:DescribeSubnets"
            - "ec2:DescribeVolumes"
            - "ec2:DescribeVolumesModifications"
            - "ec2:DescribeVpcs"
            - "eks:DescribeCluster"
          Effect: Allow
          Resource:
            - '*'
    Version: 2012-10-17
    Roles: !Ref existingNodeRole
    Type: 'AWS::IAM::ManagedPolicy'
AWSIAMRoleEKSCoontrolPlane:
  Properties:
    AssumeRolePolicyDocument:
      Statement:
        - Action:

```



```

    - 'sts:AssumeRole'
  Effect: Allow
  Principal:
    Service:
      - eks.amazonaws.com
  Version: 2012-10-17
  ManagedPolicyArns:
    - 'arn:aws:iam::aws:policy/AmazonEKSClusterPolicy'
  RoleName: eks-controlplane.cluster-api-provider-aws.sigs.k8s.io
  Type: 'AWS::IAM::Role'

```

To create the resources in the CloudFormation stack, copy the contents above into a file and execute the following command after replacing `MYFILENAME.yaml` and `MYSTACKNAME` with the intended values:

```
aws cloudformation create-stack --template-body=file://MYFILENAME.yaml --stack-name=MYSTACKNAME --capabilities CAPABILITY_NAMED_IAM
```

5.12.1.4.2.3 Add EKS CSI Policy

AWS CloudFormation does not support attaching an existing IAM Policy to an existing IAM Role. Add the necessary IAM policy to your worker instance profile using the `aws` CLI:

```
aws iam attach-role-policy --role-name nodes.cluster-api-provider-aws.sigs.k8s.io --policy-arn arn:aws:iam::aws:policy/service-role/AmazonEBSCSIDriverPolicy
```

5.12.1.4.3 Next Step:

[EKS: Create an Image \(see page 414\)](#)

5.12.1.5 EKS: Create an Image

Enterprise

Gov Advanced

AWS EKS best practices discourage building custom images. The [Amazon EKS Optimized AMI](#)³⁶⁰ is the preferred way to deploy containers for EKS. If the image is customized, it breaks some of the autoscaling and security capabilities of EKS.

5.12.1.5.1 Next Step:

[EKS: Cluster Creation \(see page 414\)](#)

³⁶⁰ <https://docs.aws.amazon.com/eks/latest/userguide/eks-optimized-amis.html>

5.12.1.6 EKS: Create an EKS Cluster

Enterprise

Gov Advanced

When creating a [Managed](#) (see [page 82](#)) cluster on your EKS infrastructure, you can choose from multiple configuration types. The steps for creating and accessing your cluster are listed below after setting [minimal permissions](#) (see [page 409](#)) and creating [IAM Policies and Roles](#) (see [page 412](#)). For more information about custom EKS configurations, refer to the [EKS Infrastructure](#) (see [page 1127](#)) under [Custom Installation and Additional Infrastructure Tools](#) (see [page 1036](#)).

5.12.1.6.1 Access Cluster


- Use previously installed [aws-iam-authenticator](#)³⁶¹ to access your cluster using kubectl. Amazon EKS uses IAM to provide authentication to your Kubernetes cluster.
- Export the AWS region where you want to deploy the cluster:

```
export AWS_REGION=us-west-2
```

- Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```

See the AWS site for more information about [AWS credentials](#).³⁶²

 Ensure that the `KUBECONFIG` environment variable is set to the Management cluster by running `export KUBECONFIG=<Management_cluster_kubeconfig>.conf`.

5.12.1.6.2 Name Your Cluster

Follow these steps:

1. Give your cluster a unique name suitable for your environment.


In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable:


³⁶¹ <https://docs.aws.amazon.com/eks/latest/userguide/install-aws-iam-authenticator.html>

³⁶² <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

```
export CLUSTER_NAME=<aws-example>
```

-  The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)³⁶³ for more naming information.

5.12.1.6.3 Create a New EKS Kubernetes Cluster

-  By default, the control-plane Nodes will be created in 3 different Availability Zones. However, the default worker Nodes will reside in a single zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:

1. Set the environment variable to the name you assigned this cluster.

```
export CLUSTER_NAME=eks-example
```

2. Make sure your AWS credentials are up-to-date. Refresh the credentials command is only necessary if you are using [Static Credentials](#) (see page 0) (Access Keys) otherwise, if you are using role-based authentication on a bastion host, proceed to step 3:

```
dkp update bootstrap credentials aws
```

3. Create the cluster:

```
dkp create cluster eks \
  --cluster-name=${CLUSTER_NAME} \
  --additional-tags=owner=$(whoami)
```

⚠ NOTE: If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

³⁶³ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

⚠ Optional flag for ECR: Configure your cluster to use an existing local registry as a mirror when attempting to pull images. Below is an AWS ECR example:

```
export --registry-mirror-url-string=<YOUR_ECR_URL>
```

- `REGISTRY-MIRROR-URL-STRING`: the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images. Users can still pull their own images from ECR directly or use ECR as a mirror. See CLI commands for `--registry-mirror` flag here: [dkp create cluster eks \(see page 1663\)](#)

```
Generating cluster resources
cluster.cluster.x-k8s.io/eks-example created
awsmanagedcontrolplane.controlplane.cluster.x-k8s.io/eks-example-control-plane
created
machinedeployment.cluster.x-k8s.io/eks-example-md-0 created
awsmachinetemplate.infrastructure.cluster.x-k8s.io/eks-example-md-0 created
eksconfigtemplate.bootstrap.cluster.x-k8s.io/eks-example-md-0 created
clusterresourceset.addons.cluster.x-k8s.io/calico-cni-installation-eks-example
created
configmap/calico-cni-installation-eks-example created
configmap/tigera-operator-eks-example created
clusterresourceset.addons.cluster.x-k8s.io/cluster-autoscaler-eks-example
created
configmap/cluster-autoscaler-eks-example created
clusterresourceset.addons.cluster.x-k8s.io/node-feature-discovery-eks-example
created
configmap/node-feature-discovery-eks-example created
clusterresourceset.addons.cluster.x-k8s.io/nvidia-feature-discovery-eks-example
created
configmap/nvidia-feature-discovery-eks-example created
```

4. Inspect or edit the cluster objects:

Use your favorite editor.

NOTE: Editing the cluster objects requires some understanding of Cluster API. Edits can prevent the cluster from deploying successfully.

The objects are [Custom Resources](#)³⁶⁴ defined by Cluster API components, and they belong in three different categories:

a. **Cluster**

A *Cluster* object has references to the infrastructure-specific and control plane objects.

Because this is an AWS cluster, there is an *AWSCluster* object that describes the infrastructure-specific cluster properties. Here, this means the AWS region, the VPC ID, subnet IDs, and security group rules required by the Pod network implementation.

b. **Control Plane**

³⁶⁴ <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

A *AWSManagedControlPlane* object describes the control plane, which is the group of machines that run the Kubernetes control plane components, which include the etcd distributed database, the API server, the core controllers, and the scheduler. The object describes the configuration for these components. The object also has a reference to an infrastructure-specific object that describes the properties of all control plane machines.

c. **Node Pool**

A Node Pool is a collection of machines with identical properties. For example, a cluster might have one Node Pool with large memory capacity, another Node Pool with GPU support. Each Node Pool is described by three objects: The *MachinePool* references an object that describes the configuration of Kubernetes components (for example, kubelet) deployed on each node pool machine, and an infrastructure-specific object that describes the properties of all node pool machines. Here, it references a *KubeadmConfigTemplate*, and an *AWSMachineTemplate* object, which describes the instance type, the type of disk used, the size of the disk, among other properties.

For in-depth documentation about the objects, read [Concepts](#)³⁶⁵ in the Cluster API Book.

5. Wait for the cluster control-plane to be ready:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

```
cluster.cluster.x-k8s.io/eks-example condition met
```

The `READY` status will become `True` after the cluster control-plane becomes ready in one of the following steps.

6. After the objects are created on the API server, the Cluster API controllers reconcile them. They create infrastructure and machines. As they progress, they update the Status of each object. Konvoy provides a command to describe the current status of the cluster:

```
dkp describe cluster -c ${CLUSTER_NAME}
```

```
NAME                                                    READY
SEVERITY REASON SINCE MESSAGE
Cluster/eks-example                                     True
10m
└─ControlPlane - AWSManagedControlPlane/eks-example-control-plane True
10m
```

³⁶⁵ <https://cluster-api.sigs.k8s.io/user/concepts.html>

```

└─Workers
  └─MachineDeployment/eks-example-md-0                               True
26s
  └─Machine/eks-example-md-0-78fcd7c7b7-66ntt                       True
84s
  └─Machine/eks-example-md-0-78fcd7c7b7-b9qmc                       True
84s
  └─Machine/eks-example-md-0-78fcd7c7b7-v5vfq                       True
84s
  └─Machine/eks-example-md-0-78fcd7c7b7-zl6m2                       True
84s

```

7. As they progress, the controllers also create Events. List the Events using this command:

```
kubectl get events | grep ${CLUSTER_NAME}
```

For brevity, the example uses `grep`. It is also possible to use separate commands to get Events for specific objects. For example, `kubectl get events --field-selector involvedObject.kind="AWSCluster"` and `kubectl get events --field-selector involvedObject.kind="AWSMachine"`.

```

46m          Normal    SuccessfulCreateVPC
awsmanagedcontrolplane/eks-example-control-plane   Created new managed VPC
"vpc-05e775702092abf09"
46m          Normal    SuccessfulSetVPCAttributes
awsmanagedcontrolplane/eks-example-control-plane   Set managed VPC attributes
for "vpc-05e775702092abf09"
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane   Created new managed Subnet
"subnet-0419dd3f2dfd95ff8"
46m          Normal    SuccessfulModifySubnetAttributes
awsmanagedcontrolplane/eks-example-control-plane   Modified managed Subnet
"subnet-0419dd3f2dfd95ff8" attributes
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane   Created new managed Subnet
"subnet-0e724b128e3113e47"
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane   Created new managed Subnet
"subnet-06b2b31ea6a8d3962"
46m          Normal    SuccessfulModifySubnetAttributes
awsmanagedcontrolplane/eks-example-control-plane   Modified managed Subnet
"subnet-06b2b31ea6a8d3962" attributes
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane   Created new managed Subnet
"subnet-0626ce238be32bf98"
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane   Created new managed Subnet
"subnet-0f53cf59f83177800"

```

```

46m          Normal    SuccessfulModifySubnetAttributes
awsmanagedcontrolplane/eks-example-control-plane    Modified managed Subnet
"subnet-0f53cf59f83177800" attributes
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane    Created new managed Subnet
"subnet-0878478f6bbf153b2"
46m          Normal    SuccessfulCreateInternetGateway
awsmanagedcontrolplane/eks-example-control-plane    Created new managed Internet
Gateway "igw-09fb52653949d4579"
46m          Normal    SuccessfulAttachInternetGateway
awsmanagedcontrolplane/eks-example-control-plane    Internet Gateway
"igw-09fb52653949d4579" attached to VPC "vpc-05e775702092abf09"
46m          Normal    SuccessfulCreateNATGateway
awsmanagedcontrolplane/eks-example-control-plane    Created new NAT Gateway
"nat-06356aac28079952d"
46m          Normal    SuccessfulCreateNATGateway
awsmanagedcontrolplane/eks-example-control-plane    Created new NAT Gateway
"nat-0429d1cd9d956bf35"
46m          Normal    SuccessfulCreateNATGateway
awsmanagedcontrolplane/eks-example-control-plane    Created new NAT Gateway
"nat-059246bcc9d4e88e7"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Created managed RouteTable
"rtb-01689c719c484fd3c"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane    Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Associated managed
RouteTable "rtb-01689c719c484fd3c" with subnet "subnet-0419dd3fd95ff8"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Created managed RouteTable
"rtb-065af81b9752eeb69"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane    Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Associated managed
RouteTable "rtb-065af81b9752eeb69" with subnet "subnet-0e724b128e3113e47"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Created managed RouteTable
"rtb-03eeff810a89afc98"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane    Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Associated managed
RouteTable "rtb-03eeff810a89afc98" with subnet "subnet-06b2b31ea6a8d3962"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Created managed RouteTable
"rtb-0fab36f8751fdee73"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane    Created route {...

```

```

46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane Associated managed
RouteTable "rtb-0fab36f8751fdee73" with subnet "subnet-0626ce238be32bf98"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane Created managed RouteTable
"rtb-0e5c9c7bbc3740a0f"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane Associated managed
RouteTable "rtb-0e5c9c7bbc3740a0f" with subnet "subnet-0f53cf59f83177800"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane Created managed RouteTable
"rtb-0bf58eb5f73c387af"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane Associated managed
RouteTable "rtb-0bf58eb5f73c387af" with subnet "subnet-0878478f6bbf153b2"
46m          Normal    SuccessfulCreateSecurityGroup
awsmanagedcontrolplane/eks-example-control-plane Created managed
SecurityGroup "sg-0b045c998a120a1b2" for Role "node-eks-additional"
46m          Normal    InitiatedCreateEKSControlPlane
awsmanagedcontrolplane/eks-example-control-plane Initiated creation of a new
EKS control plane default_eks-example-control-plane
37m          Normal    SuccessfulCreateEKSControlPlane
awsmanagedcontrolplane/eks-example-control-plane Created new EKS control
plane default_eks-example-control-plane
37m          Normal    SuccessfulCreateKubeconfig
awsmanagedcontrolplane/eks-example-control-plane Created kubeconfig for
cluster "eks-example"
37m          Normal    SuccessfulCreateUserKubeconfig
awsmanagedcontrolplane/eks-example-control-plane Created user kubeconfig for
cluster "eks-example"
27m          Normal    SuccessfulCreate
awsmachine/eks-example-md-0-4t9nc Created new node instance
with id "i-0aecc1897c93df740"
26m          Normal    SuccessfulDeleteEncryptedBootstrapDataSecrets
awsmachine/eks-example-md-0-4t9nc AWS Secret entries
containing userdata deleted
26m          Normal    SuccessfulSetNodeRef                               machine/
eks-example-md-0-78fcd7c7b7-fn7x9 ip-10-0-88-24.us-west-2.compute.inte
rnal
26m          Normal    SuccessfulSetNodeRef                               machine/
eks-example-md-0-78fcd7c7b7-g64nv ip-10-0-110-219.us-west-2.compute.in
ternal
26m          Normal    SuccessfulSetNodeRef                               machine/
eks-example-md-0-78fcd7c7b7-gwc5j ip-10-0-101-161.us-west-2.compute.in
ternal
26m          Normal    SuccessfulSetNodeRef                               machine/
eks-example-md-0-78fcd7c7b7-j58s4 ip-10-0-127-49.us-west-2.compute.int
ernal

```



```

46m          Normal    SuccessfulCreate
machineset/eks-example-md-0-78fcd7c7b7
example-md-0-78fcd7c7b7-fn7x9"          Created machine "eks-
46m          Normal    SuccessfulCreate
machineset/eks-example-md-0-78fcd7c7b7
example-md-0-78fcd7c7b7-g64nv"          Created machine "eks-
46m          Normal    SuccessfulCreate
machineset/eks-example-md-0-78fcd7c7b7
example-md-0-78fcd7c7b7-j58s4"          Created machine "eks-
46m          Normal    SuccessfulCreate
machineset/eks-example-md-0-78fcd7c7b7
example-md-0-78fcd7c7b7-gwc5j"          Created machine "eks-
27m          Normal    SuccessfulCreate
awsmachine/eks-example-md-0-7whkv
with id "i-06dfc0466b8f26695"          Created new node instance
26m          Normal    SuccessfulDeleteEncryptedBootstrapDataSecrets
awsmachine/eks-example-md-0-7whkv
containing userdata deleted            AWS Secret entries
27m          Normal    SuccessfulCreate
awsmachine/eks-example-md-0-ttgzv
with id "i-0544fce0350fd41fb"          Created new node instance
26m          Normal    SuccessfulDeleteEncryptedBootstrapDataSecrets
awsmachine/eks-example-md-0-ttgzv
containing userdata deleted            AWS Secret entries
27m          Normal    SuccessfulCreate
awsmachine/eks-example-md-0-v2hrf
with id "i-0498906edde162e59"          Created new node instance
26m          Normal    SuccessfulDeleteEncryptedBootstrapDataSecrets
awsmachine/eks-example-md-0-v2hrf
containing userdata deleted            AWS Secret entries
46m          Normal    SuccessfulCreate
machinedeployment/eks-example-md-0
example-md-0-78fcd7c7b7"              Created MachineSet "eks-

```

5.12.1.6.4 Known Limitations



Be aware of these limitations in the current release of Konvoy.

- The Konvoy version used to create a workload cluster must match the Konvoy version used to delete a workload cluster.
- EKS clusters cannot be Self-managed.
- Konvoy supports deploying one workload cluster.
- Konvoy generates a set of objects for one Node Pool.
- Konvoy does not validate edits to cluster objects.

5.12.1.6.4.1 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>³⁶⁶³⁶⁷.

5.12.1.6.4.2 Next Step:

[EKS: Grant Cluster Access](#) (see page 423)

5.12.1.7 EKS: Grant Cluster Access

Enterprise

Gov Advanced

5.12.1.7.1 How to Grant EKS Cluster Access

You can access your cluster using AWS IAM roles in the dashboard. When you create an EKS cluster, the IAM entity is granted `system:masters` permissions in [Kubernetes Role Based Access Control \(RBAC\) configuration](#).³⁶⁸



More information about the configuration of the EKS control plane can be found on the [EKS Cluster IAM Policies and Roles](#) (see page 1132) page.

If the EKS cluster was created as a cluster using a self-managed AWS cluster that uses IAM Instance Profiles, you will need to modify the `IAMAuthenticatorConfig` field in the `AWSManagedControlPlane` API object to allow other IAM entities to access the EKS workload cluster. Follow the steps below:

1. Run the following command with your `KUBECONFIG` configured to select the self-managed cluster previously used to create the workload EKS cluster. Ensure you substitute `${CLUSTER_NAME}` and `${CLUSTER_NAMESPACE}` with their corresponding values for your cluster.

³⁶⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

³⁶⁷ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

³⁶⁸ <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

```
kubectl edit awsmanagedcontrolplane ${CLUSTER_NAME}-control-plane -n ${CLUSTER_NAMESPACE}
```

2. Edit the `IamAuthenticatorConfig` field with the IAM Role to the corresponding Kubernetes Role. In this example, the IAM role `arn:aws:iam::111122223333:role/PowerUser` is granted the cluster role `system:masters`. Note that this example uses example AWS resource [ARNs](#)³⁶⁹, so these values should be substituted for real values in the corresponding AWS account.

```
iamAuthenticatorConfig:
  mapRoles:
    - groups:
      - system:bootstrappers
      - system:nodes
      rolearn: arn:aws:iam::111122223333:role/my-node-role
      username: system:node:{{EC2PrivateDNSName}}
    - groups:
      - system:masters
      rolearn: arn:aws:iam::111122223333:role/PowerUser
      username: admin
```

For further instructions on changing or assigning `roles` or `clusterroles` to which you can map IAM users or roles, see [Amazon Enabling IAM access to your cluster](#)³⁷⁰.

5.12.1.7.2 Next Step:

[EKS: Retrieve kubeconfig for EKS Cluster](#) (see page 424)

5.12.1.8 EKS: Retrieve kubeconfig for EKS Cluster

Enterprise

Gov Advanced

This guide explains how to use the command line to interact with your newly deployed Kubernetes cluster. Before you start, make sure you have created a workload cluster, as described in [EKS: Create an EKS Cluster](#) (see page 418).

5.12.1.8.1 Explore the New Kubernetes Cluster

Follow these steps:

1. Get a `kubeconfig` file for the workload cluster:

³⁶⁹ <https://docs.aws.amazon.com/general/latest/gr/aws-arns-and-namespaces.html>

³⁷⁰ <https://docs.aws.amazon.com/eks/latest/userguide/add-user-role.html>

When the workload cluster is created, the cluster lifecycle services generate a `kubeconfig` file for the workload cluster, and write it to a `Secret`. The `kubeconfig` file is scoped to the cluster administrator.

Get the `kubeconfig` from the `Secret`, and write it to a file, using this command:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

2. List the Nodes using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
ip-10-0-122-211.us-west-2.compute.internal eks-ba74326	Ready	<none>	35m	v1.23.9-
ip-10-0-127-74.us-west-2.compute.internal eks-ba74326	Ready	<none>	35m	v1.23.9-
ip-10-0-71-155.us-west-2.compute.internal eks-ba74326	Ready	<none>	35m	v1.23.9-
ip-10-0-93-47.us-west-2.compute.internal eks-ba74326	Ready	<none>	35m	v1.23.9-

NOTE: It may take a few minutes for the Status to move to `Ready` while the Pod network is deployed. The Nodes' Status should change to `Ready` soon after the `calico-node` DaemonSet Pods are `Ready`.

3. List the Pods using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get --all-namespaces pods
```

NAMESPACE	STATUS	RESTARTS	NAME	AGE	READY
calico-system	Running	0	calico-kube-controllers-7d6749878f-ccsx9	34m	1/1
calico-system	Running	0	calico-node-2r6l8	34m	1/1

calico-system	calico-node-5pdlb	1/1
Running 0	34m	
calico-system	calico-node-n24hh	1/1
Running 0	34m	
calico-system	calico-node-qrh7p	1/1
Running 0	34m	
calico-system	calico-typha-7bbcb87696-7pk45	1/1
Running 0	34m	
calico-system	calico-typha-7bbcb87696-t4c8r	1/1
Running 0	34m	
calico-system	csi-node-driver-bz48k	2/2
Running 0	34m	
calico-system	csi-node-driver-k5mmk	2/2
Running 0	34m	
calico-system	csi-node-driver-nvckc	2/2
Running 0	34m	
calico-system	csi-node-driver-x4xnh	2/2
Running 0	34m	
kube-system	aws-node-2xp86	1/1
Running 0	35m	
kube-system	aws-node-5f2kx	1/1
Running 0	35m	
kube-system	aws-node-6lzm7	1/1
Running 0	35m	
kube-system	aws-node-pz8c6	1/1
Running 0	35m	
kube-system	cluster-autoscaler-789d86b489-sz9x2	0/1
Init:0/1 0	36m	
kube-system	coredns-57ff979f67-pk5cg	1/1
Running 0	75m	
kube-system	coredns-57ff979f67-sf2j9	1/1
Running 0	75m	
kube-system	ebs-csi-controller-5f6bd5d6dc-bplwm	6/6
Running 0	36m	
kube-system	ebs-csi-controller-5f6bd5d6dc-dpjt7	6/6
Running 0	36m	
kube-system	ebs-csi-node-7hmm5	3/3
Running 0	35m	
kube-system	ebs-csi-node-l4vfh	3/3
Running 0	35m	
kube-system	ebs-csi-node-mfr7c	3/3
Running 0	35m	
kube-system	ebs-csi-node-v8krq	3/3
Running 0	35m	
kube-system	kube-proxy-7fc5x	1/1
Running 0	35m	
kube-system	kube-proxy-vvkmk	1/1
Running 0	35m	
kube-system	kube-proxy-x6hcc	1/1
Running 0	35m	
kube-system	kube-proxy-x8frb	1/1
Running 0	35m	

kube-system	snapshot-controller-8ff89f489-4cfxv	1/1
Running 0	36m	
kube-system	snapshot-controller-8ff89f489-78gg8	1/1
Running 0	36m	
node-feature-discovery	node-feature-discovery-master-7d5985467-52fcn	1/1
Running 0	36m	
node-feature-discovery	node-feature-discovery-worker-88hr7	1/1
Running 0	34m	
node-feature-discovery	node-feature-discovery-worker-h95nq	1/1
Running 0	35m	
node-feature-discovery	node-feature-discovery-worker-lfghg	1/1
Running 0	34m	
node-feature-discovery	node-feature-discovery-worker-prc8p	1/1
Running 0	35m	
tigera-operator	tigera-operator-6dcd98c8ff-k97hq	1/1
Running 0	36m	

5.12.1.8.2 Next Step:

[EKS: Attach a Cluster](#) (see page 427)

5.12.1.9 EKS: Attach a Cluster

Enterprise

Gov Advanced

You can attach existing Kubernetes clusters to the Management Cluster. After attaching the cluster, you can use the UI to [examine and manage](#) (see page 745) this cluster. The following procedure shows how to attach an existing Amazon Elastic Kubernetes Service (EKS) cluster.

ⓘ This procedure assumes you have an existing and spun up Amazon EKS cluster(s) with administrative privileges. Refer to the Amazon [EKS](#)³⁷¹ for setup and configuration information.

- Install [aws-iam-authenticator](#)³⁷². This binary is used to access your cluster using kubectl.

5.12.1.9.1 Attach a Pre-existing EKS Cluster

Ensure that the `KUBECONFIG` environment variable is set to the Management cluster before attaching by running:

³⁷¹ <https://aws.amazon.com/eks/>

³⁷² <https://docs.aws.amazon.com/eks/latest/userguide/install-aws-iam-authenticator.html>

```
export KUBECONFIG=<Management_cluster_kubeconfig>.conf
```

5.12.1.9.1.1 Access your EKS clusters

1. Ensure you are connected to your EKS clusters. Enter the following commands for each of your clusters:

```
kubectl config get-contexts
kubectl config use-context <context for first eks cluster>
```

2. Confirm `kubectl` can access the EKS cluster:

```
kubectl get nodes
```

5.12.1.9.1.2 Create a kubeconfig File

To get started, ensure you have `kubecti`³⁷³ set up and configured with `ClusterAdmin`³⁷⁴ for the cluster you want to connect to Kommander.

1. Create the necessary service account:

```
kubectl -n kube-system create serviceaccount kommander-cluster-admin
```

2. Create a token secret for the `serviceaccount`:

```
kubectl -n kube-system create -f - <<EOF
apiVersion: v1
kind: Secret
metadata:
  name: kommander-cluster-admin-sa-token
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
type: kubernetes.io/service-account-token
EOF
```

For more information on Service Account Tokens, refer to [this article](#)³⁷⁵ in our blog.

3. Verify that the `serviceaccount` token is ready by running this command:

³⁷³ <https://kubernetes.io/docs/tasks/tools/#kubectl>

³⁷⁴ <https://kubernetes.io/docs/concepts/cluster-administration/>

³⁷⁵ <https://eng.d2iq.com/blog/service-account-tokens-in-kubernetes-v1.24/#whats-changed-in-kubernetes-v124>

```
kubectl -n kube-system get secret kommander-cluster-admin-sa-token -oyaml
```

Verify that the `data.token` field is populated. The output should be similar to this:

```
apiVersion: v1
data:
  ca.crt: LS0tLS1CRUdJTiBDR...
  namespace: ZGVmYXVsdA==
  token: ZXlKaGJHY2lPaUpTVX...
kind: Secret
metadata:
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
    kubernetes.io/service-account.uid: b62bc32e-b502-4654-921d-94a742e273a8
  creationTimestamp: "2022-08-19T13:36:42Z"
  name: kommander-cluster-admin-sa-token
  namespace: default
  resourceVersion: "8554"
  uid: 72c2a4f0-636d-4a70-9f1c-55a75f15e520
type: kubernetes.io/service-account-token
```

4. Configure the new service account for `cluster-admin` permissions:

```
cat << EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: kommander-cluster-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- kind: ServiceAccount
  name: kommander-cluster-admin
  namespace: kube-system
EOF
```

5. Set up the following environment variables with the access data that is needed for producing a new kubeconfig file:

```
export USER_TOKEN_VALUE=$(kubectl -n kube-system get secret/kommander-cluster-admin-sa-token -o=go-template='{{.data.token}}' | base64 --decode)
export CURRENT_CONTEXT=$(kubectl config current-context)
```



```
export CURRENT_CLUSTER=$(kubectl config view --raw -o=go-
template='{{range .contexts}}{{if eq .name ""${CURRENT_CONTEXT}''''}}
{{ index .context "cluster" }}{{end}}{{end}}')
export CLUSTER_CA=$(kubectl config view --raw -o=go-
template='{{range .clusters}}{{if eq .name ""${CURRENT_CLUSTER}''''}}"{{with
index .cluster "certificate-authority-data" }}{{.}}{{end}}"{{ end }}{{ end }}')
export CLUSTER_SERVER=$(kubectl config view --raw -o=go-
template='{{range .clusters}}{{if eq .name ""${CURRENT_CLUSTER}''''}}
{{ .cluster.server }}{{end}}{{ end }}')
```

6. Confirm these variables have been set correctly:

```
export -p | grep -E 'USER_TOKEN_VALUE|CURRENT_CONTEXT|CURRENT_CLUSTER|
CLUSTER_CA|CLUSTER_SERVER'
```

7. Generate a kubeconfig file that uses the environment variable values from the previous step:

```
cat << EOF > kommander-cluster-admin-config
apiVersion: v1
kind: Config
current-context: ${CURRENT_CONTEXT}
contexts:
- name: ${CURRENT_CONTEXT}
  context:
    cluster: ${CURRENT_CONTEXT}
    user: kommander-cluster-admin
    namespace: kube-system
clusters:
- name: ${CURRENT_CONTEXT}
  cluster:
    certificate-authority-data: ${CLUSTER_CA}
    server: ${CLUSTER_SERVER}
users:
- name: kommander-cluster-admin
  user:
    token: ${USER_TOKEN_VALUE}
EOF
```

8. This process produces a file in your current working directory called `kommander-cluster-admin-config`. The contents of this file are used in Kommander to attach the cluster.

Before importing this configuration, verify the `kubeconfig` file can access the cluster:

```
kubectl --kubeconfig $(pwd)/kommander-cluster-admin-config get all --all-
namespaces
```

5.12.1.9.2 Attach from UI or Attach Manually Through the CLI

There are two options to attach the cluster. Choose from one of the following based on your preference.

- [Through the UI \(see page 431\)](#)
- [Using the DKP CLI \(see page 431\)](#)

5.12.1.9.2.1 Finish Attach EKS Cluster from the UI

Now that you have kubeconfig, go to the DKP UI and follow these steps below:

1. From the top menu bar, select your target workspace.
2. On the Dashboard page, select the **Add Cluster** option in the **Actions** dropdown menu at the top right.
3. Select **Attach Cluster**.
4. Select the **No additional networking restrictions** card. Alternatively, if you must use network restrictions, stop following the steps below, and see the instructions on the page [Attach a cluster WITH network restrictions \(see page 772\)](#).
5. Upload the kubeconfig file you created in the previous section (or copy its contents) into the **Cluster Configuration** section.
6. The **Cluster Name** field automatically populates with the name of the cluster in the kubeconfig. You can edit this field with the name you want for your cluster.
7. Add labels to classify your cluster as needed.
8. Select **Create** to attach your cluster.



If a cluster has limited resources to deploy all the federated platform services, it will fail to stay attached in the DKP UI. If this happens, ensure your system has sufficient resources for all pods.

5.12.1.9.3 Manually Attach a DKP CLI Cluster to the Management Cluster



These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster \(see page 82\)](#) with the DKP CLI, it attaches automatically to the [Management Cluster \(see page 82\)](#) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [\[attach it in the UI\]](#)(link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

7. This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

8. Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

9. Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubectl get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.12.1.9.4 Related Information

For information on related topics or procedures, refer to the following:

- [Create an EKS Cluster from the UI](#) (see page 1144)
- [Configuring and Running Amazon EKS Clusters](#)³⁷⁶
- [Manage Clusters](#) (see page 745)

5.12.1.9.5 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.13 vSphere Install Options

For vSphere, an installation scenario is provided for you in this location.

³⁷⁶ <https://aws.amazon.com/eks/>

Remember, there are always more options for custom YAML in the [Custom Install and Additional Infrastructure Tools](#) (see page 1036) in the [vSphere Infrastructure](#) (see page 1254) section, but this will get you operative in the most common scenarios.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

Below are all the supported environment variable combinations: (Refer to this page to check if your OS is a [Supported Operating System](#) (see page 69).)

- [vSphere Prerequisites: All Installation Types](#) (see page 434)
- [vSphere Install](#) (see page 438)
- [vSphere Air-gapped Install](#) (see page 459)
- [vSphere FIPS Install](#) (see page 483)
- [vSphere FIPS Air-gapped Install](#) (see page 505)

5.13.1 Overview

vSphere is a more complex setup than some of the other providers and infrastructures, so an overview of steps has been provided to help.

The overall process for configuring vSphere and DKP together includes the following steps:

1. Configure vSphere to provide the needed elements, described in the [vSphere Prerequisites: All Installation Types](#) (see page 434).
2. **For air-gapped environments:** Create a [bastion VM host](#) (see page 1285).
3. Create a base OS image (for use in the OVA package containing the disk images packaged with the OVF).
4. Create a CAPI VM image template that uses the base OS image and adds the needed Kubernetes cluster components.
5. Create a new self-managing cluster on vSphere.
6. Install Kommander.
7. Verify and log in to the UI.

5.13.2 Next Step:

[vSphere Prerequisites: All Installation Types](#) (see page 434)

5.13.3 vSphere Prerequisites: All Installation Types

- [vSphere: Minimum User Permissions](#) (see page 435)
- [vSphere: Storage Options](#) (see page 437)
- [vSphere: VMware Prerequisites](#) (see page 437)

5.13.3.1 vSphere: Minimum User Permissions

5.13.3.1.1 Create minimum required roles for provisioning and installing in vSphere

When a user needs permissions less than Admin, a role must be created. The process for configuring a vSphere role with the least permissions for provisioning nodes and installing includes the following steps:

1. Open a vSphere Client connection to the vCenter Server, described in the [Prerequisites](#) (see page 1256).
2. Select Home > Administration > Roles > Add Role.
3. Give the new role a name, then select these Privileges:

Cns		
<input checked="" type="checkbox"/>		Searchable
Datastore		
<input checked="" type="checkbox"/>		Allocate space
<input checked="" type="checkbox"/>		Low level file operations
Host		
		• Configuration
<input checked="" type="checkbox"/>		Storage partition configuration
Profile-driven storage		
<input checked="" type="checkbox"/>		Profile-driven storage view
Network		
<input checked="" type="checkbox"/>		Assign network
Resource		

<input checked="" type="checkbox"/>	Assign virtual machine to resource pool
Virtual machine	
	• Change Configuration - from the list in that section, select these permissions below:
<input checked="" type="checkbox"/>	Add new disk
<input checked="" type="checkbox"/>	Add existing disk
<input checked="" type="checkbox"/>	Add or remove device
<input checked="" type="checkbox"/>	Advanced configuration
<input checked="" type="checkbox"/>	Change CPU count
<input checked="" type="checkbox"/>	Change Memory
<input checked="" type="checkbox"/>	Change Settings
<input checked="" type="checkbox"/>	Reload from path
Edit inventory	
<input checked="" type="checkbox"/>	Create from existing
<input checked="" type="checkbox"/>	Remove
Interaction	
<input checked="" type="checkbox"/>	Power off
<input checked="" type="checkbox"/>	Power on
Provisioning	
<input checked="" type="checkbox"/>	Clone template

<input checked="" type="checkbox"/>	Deploy template
Session	
<input checked="" type="checkbox"/>	ValidateSession

Add the permission at the highest level and set to propagate the permissions.

5.13.3.1.2 Next Step:

[vSphere: Storage Options](#) (see page 437)

5.13.3.2 vSphere: Storage Options

5.13.3.2.1 Explore storage options and considerations for using DKP with VMware vSphere

The [vSphere Container Storage](#)³⁷⁷ plugin supports shared NFS, vNFS, and vSAN. You need to provision your storage options in vCenter prior to [creating a CAPI VM Template](#) (see page 440) for non-air-gapped or in [air-gapped](#) (see page 463) in DKP for use with vSphere.

DKP has integrated the [CSI 2.x driver](#)³⁷⁸ used in vSphere. When creating your DKP cluster, DKP uses whatever configuration you provide for the Datastore name. vSAN is not required. Using NFS can reduce the amount of tagging and permission granting required to configure your cluster.

5.13.3.2.2 Next Step:

[vSphere: VMware Prerequisites](#) (see page 437)

5.13.3.3 vSphere: VMware Prerequisites

Before installing, verify that your [VMware vSphere Client environment](#)³⁷⁹ meets the following basic requirements:

- Access to a [bastion](#) (see page 1285) VM, or other network-connected host, running vSphere Client version v6.7.x with Update 3 or later version.
 - You must be able to reach the vSphere API endpoint from where the Konvoy command line interface (CLI) runs.
- vSphere account with credentials configured - this account must have [Administrator privileges](#) (see page 438).

³⁷⁷ <https://docs.vmware.com/en/VMware-vSphere-Container-Storage-Plug-in/2.0/vmware-vsphere-csp-getting-started/GUID-74AF02D7-1562-48BD-A9FE-C81A53342AC3.html>

³⁷⁸ <https://github.com/kubernetes-sigs/vsphere-csi-driver>

³⁷⁹ https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html

- A RedHat® subscription with username and password for downloading DVD ISOs.
- Valid vSphere values for the following:
 - vCenter API server URL
 - Datacenter name
 - Zone name that contains [ESXi hosts](#)³⁸⁰ for your cluster's nodes.
 - Datastore name for the shared storage resource to be used for the VMs in the cluster.
 - Use of PersistentVolumes in your cluster depends on Cloud Native Storage (CNS), available in vSphere v6.7.x with Update 3 and later versions. CNS depends on this shared Datastore's configuration.
 - Datastore URL from the datastore record for the shared datastore you want your cluster to use.
 - You need this URL value to ensure that the correct Datastore is used when DKP creates VMs for your cluster in vSphere.
 - Folder name
 - Base template name, such as base-rhel-8, or base-rhel-7.
 - Name of a Virtual Network that has DHCP enabled for both air-gapped and non air-gapped environments.
 - Name of your virtual IP interface.
 - Resource Pools - at least one resource pool is needed, with every host in the pool having access to shared storage, such as VSAN.
 - Each host in the resource pool needs access to shared storage, such as NFS or VSAN, to make use of MachineDeployments and high-availability control planes.

5.13.3.3.1 Next Step:

[vSphere Install](#) (see page 438)

5.13.4 vSphere Install

This section provides instructions to install DKP in a vSphere non-air-gapped environment.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

³⁸⁰ <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.esxi.install.doc/GUID-B2F01BF5-078A-4C7E-B505-5DFFED0B8C38.html>

5.13.4.1 Further vSphere Prerequisites

- [vSphere: VMware Prerequisites](#) (see page 440)
- <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/162430985/vSphere+Prerequisites+-+Storage+Options> (see page 440)
- <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/162430977/vSphere+Prerequisites+-+Minimum+User+Permissions> (see page 438)

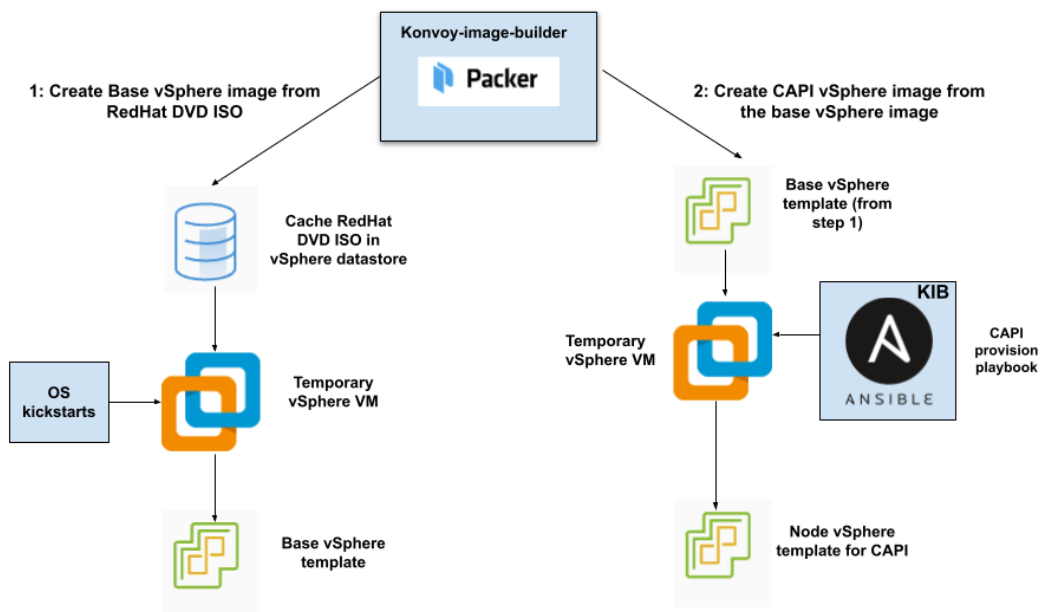
5.13.4.2 Next Step:

[vSphere: Image Creation Overview](#) (see page 439)

5.13.4.3 vSphere: Image Creation Overview

5.13.4.3.1 Image Creation Process

This diagram illustrates the image creation process:



The workflow on the left shows the creation of a base OS image in the vCenter vSphere client using inputs from Packer. The workflow on the right shows how DKP uses that same base OS image to create CAPI-enabled VM images for your cluster.

After creating the base image, the DKP image builder uses it to create a CAPI-enabled vSphere template that includes the Kubernetes objects for the cluster. You can use that resulting template with the DKP `create cluster` command to create the VM nodes in your cluster directly on a vCenter server. From that point, you can use DKP to provision and manage your cluster.

5.13.4.3.2 Next Step:

[vSphere: Create an Image \(see page 440\)](#)

5.13.4.4 vSphere: Create an Image

Creating a base OS image from DVD ISO files is a one-time process. The base OS image file is created in the vSphere Client for use in the vSphere VM template. Therefore, the base OS image is used by [Konvoy Image Builder \(see page 1461\)](#)(KIB) to create a VM template to configure Kubernetes nodes by the DKP vSphere provider.

5.13.4.4.1 Create the Base OS Image

For vSphere, a username is populated by `SSH_USERNAME` and the user can use authorization via `SSH_PASSWORD` or `SSH_PRIVATE_KEY_FILE` environment variables and required by default by packer. This user should have administrator privileges. It is possible to configure a custom user and password when building the OS image, however, that requires the Konvoy Image Builder (KIB) configuration to be [overridden \(see page 1513\)](#).

While creating the base OS image, it is important to take into consideration the following elements:

- **Storage configuration:** D2iQ recommends customizing disk partitions and not configuring a SWAP partition.
- **Network configuration:** as KIB must download and install packages, activating the network is required.
- **Connect to Red Hat:** if using RHEL, registering with Red Hat is required to configure software repositories and install software packages.
- **Software selection:** D2iQ recommends choosing **Minimal Install**.
- DKP recommends to install with the packages provided by the operating system package managers. Use the version that corresponds to the major version of your operating system.

[vSphere: Create a CAPI VM Template Image \(see page 440\)](#)

5.13.4.5 vSphere: Create a CAPI VM Template

5.13.4.5.1 Prerequisites

- Users need to [vSphere: Create an Image \(see page 443\)](#) before starting this procedure.
- Build image template with Konvoy Image Builder (KIB)



- If using GPU, prepare an OS image for your NVIDIA GPU nodepool, using the Konvoy Image Builder instructions here: [KIB for GPU](#) (see page 1485).

5.13.4.5.2 Create a vSphere Template for Your Cluster from a Base OS Image

Using the base OS image created in a previous procedure, DKP creates the new vSphere template directly on the vCenter server.

1. Set the following vSphere environment variables on the bastion VM host:

```
export VSPHERE_SERVER=your_vCenter_APIserver_URL
export VSPHERE_USERNAME=your_vCenter_user_name
export VSPHERE_PASSWORD=your_vCenter_password
```

2. Copy the base OS image file created in the vSphere Client to your desired location on the bastion VM host, and make a note of the path and file name.
3. Create an `image.yaml` file and add the following variables for vSphere. DKP uses this file and these variables as inputs in the next step.

⚠ NOTE: This example is Ubuntu 20.04. You will need to replace OS name below based on your OS. See other [YAML examples](#)³⁸¹ for copy and paste below last step.

```
---
download_images: true
build_name: "ubuntu-2004"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: "<VSPHERE_CLUSTER_NAME>"
  datacenter: "<VSPHERE_DATACENTER_NAME>"
  datastore: "<VSPHERE_DATASTORE_NAME>"
  folder: "<VSPHERE_FOLDER>"
  insecure_connection: "false"
  network: "<VSPHERE_NETWORK>"
  resource_pool: "<VSPHERE_RESOURCE_POOL>"
```

³⁸¹ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

template: "os-qualification-templates/d2iq-base-Ubuntu-20.04" # change
default value with your base template name
vsphere_guest_os_type: "other4xLinux64Guest"
guest_os_type: "ubuntu2004-64"
# goss params
distribution: "ubuntu"
distribution_version: "20.04"
# Use following overrides to select the authentication method that can be used
with base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key
will be ignored

```

4. Create a vSphere VM template with your variation of the following command:

```
konvoy-image build images/ova/<image.yaml>
```

Any additional configurations can be added to this command using `--overrides` flags as shown below:

- a. Any credential overrides: `--overrides overrides.yaml`
 - b. for FIPS, add this flag: `--overrides overrides/fips.yaml`
 - c. for air-gapped, add this flag: `--overrides overrides/offline-fips.yaml`
5. The [Konvoy Image Builder](#) (see page 1487) (KIB) uses the values in `image.yaml` and the input base OS image to create a vSphere template directly on the vCenter server. This template contains the required artifacts needed to create a Kubernetes cluster. When KIB provisions the OS `image`³⁸² successfully, it creates a manifest file. The `artifact_id` field of this file contains the name of the AMI ID (AWS), template name (vSphere), or image name (GCP/Azure), for example:

```

{
  "name": "vsphere-clone",
  "builder_type": "vsphere-clone",
  "build_time": 1644985039,
  "files": null,
  "artifact_id": "konvoy-ova-vsphere-rhel-84-1.21.6-1644983717",
  "packer_run_uuid": "260e8110-77f8-ca94-e29e-ac7a2ae779c8",
  "custom_data": {
    "build_date": "2022-02-16T03:55:17Z",
    "build_name": "vsphere-rhel-84",

```

³⁸² <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

    "build_timestamp": "1644983717",
    [...]
  }
}

```

Recommendation: Now we can now see the template created in our vCenter, it is best to rename it to `dkp-<DKP_VERSION>-k8s-<K8S_VERSION>-<DISTRO>`, like `dkp-2.6.2-k8s-1.26.14-ubuntu` to keep templates organized.

6. Next steps are to deploy a DKP cluster using your vSphere template.

5.13.4.5.2.1 Additional OS YAML files for copy/paste:

RHEL 8.6

```

---
download_images: true
build_name: "rhel-86"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "os-qualification-templates/d2iq-base-RHEL-86" # change default value
with your base template name
vsphere_guest_os_type: "rhel8_64Guest"
guest_os_type: "rhel8-64"
# goss params
distribution: "RHEL"
distribution_version: "8.6"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

Rocky Linux 9.1

```

---
download_images: true
build_name: "rocky-91"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "os-qualification-templates/d2iq-base-RockyLinux-9.1" # change default
value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "rocky9-64"
  # goss params
  distribution: "rocky"
  distribution_version: "9.1"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

5.13.4.5.3 Next Step:

[vSphere: Create the Management Cluster \(see page 444\)](#)

5.13.4.6 vSphere: Create the Management Cluster


Use this procedure to create a self-managed vSphere Management cluster with DKP. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.


5.13.4.6.1 Name your Cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the CLUSTER_NAME environment variable with the command:

```
export CLUSTER_NAME=<my-vsphere-cluster>
```

5.13.4.6.2 Create a New vSphere Kubernetes Cluster

 DKP uses the vsphere CSI driver as the [default storage provider](#) (see page 91). Use a [Kubernetes CSI](#)³⁸³ compatible storage that is suitable for production.


 The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above. If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273744255/2.6%2BAttach%2Ba%2BKubernetes%2BCluster> (see page 752).

Follow these steps to create the cluster:

1. Use the following command to set the environment variables for vSphere:

```
export VSPHERE_SERVER=example.vsphere.url
export VSPHERE_USERNAME=user@example.vsphere.url
export VSPHERE_PASSWORD=example_password
```

2. Generate the Kubernetes cluster objects by copying and editing this command to include the correct values, including the VM template name you assigned in the previous procedure:

-  • To increase [Dockerhub's rate limit](#)³⁸⁴ use your Dockerhub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.
- The following example shows a common configuration. See [dkp create cluster](#) (see page 1662) reference for the full list of cluster creation options:

³⁸³ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

³⁸⁴ <https://docs.docker.com/docker-hub/download-rate-limit/>


```

dkp create cluster vsphere \
  --cluster-name ${CLUSTER_NAME} \
  --network <NETWORK_NAME> \
  --control-plane-endpoint-host <xxx.yyy.zzz.000> \
  --data-center <DATACENTER_NAME> \
  --data-store <DATASTORE_NAME> \
  --folder <FOLDER_NAME> \
  --server <VCENTER_API_SERVER_URL> \
  --ssh-public-key-file <SSH_PUBLIC_KEY_FILE> \
  --resource-pool <RESOURCE_POOL_NAME> \
  --vm-template <TEMPLATE_NAME> \
  --virtual-ip-interface <ip_interface_name> \
  --self-managed

```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1039\)](#).

= For bootstrap and custom YAML cluster creation, refer to the [Custom Installation and Additional Infrastructure Tools \(see page 1036\)](#) section of the documentation for vSphere: [vSphere Infrastructure \(see page 1254\)](#)

5.13.4.6.2.1 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>³⁸⁵³⁸⁶.

³⁸⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

³⁸⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

5.13.4.6.2.2 Known Limitations



Be aware of these limitations in the current release of DKP Konvoy.

- The DKP Konvoy version used to create a bootstrap cluster must match the DKP Konvoy version used to create a workload cluster.
- DKP Konvoy supports deploying one workload cluster.
- DKP Konvoy generates a set of objects for one Node Pool.
- DKP Konvoy does not validate edits to cluster objects.

5.13.4.6.2.3 Next Step:

[vSphere: Configure Metal LB \(see page 447\)](#)

5.13.4.7 vSphere: Configure Metal LB

If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can proceed to [vSphere: Install Kommander \(see page 449\)](#) and continue the installation process.

To use MetalLB, create a MetalLB configMap for your vSphere infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)

Select one of the following procedures to create your MetalLB manifest for further editing.

5.13.4.7.1 Layer 2 configuration


Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.



- MetalLB IP address ranges/CIDRs should be within the node's primary network [subnet \(see page 1046\)](#).
- MetalLB IP address ranges/CIDRs and node subnet should not conflict with the Kubernetes cluster pod and service subnets.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:

 The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    address-pools:
    - name: default
      protocol: layer2
      addresses:
      - 192.168.1.240-192.168.1.250
EOF
```

When complete, run the following `kubectl` command.


```
kubectl apply -f metallb-conf.yaml
```

5.13.4.7.2 BGP configuration

BGP, like any routing protocol, requires a mesh, so you configure BGP between the nodes. If you are in a Data Center, peer with the BGP routers. The instructions for use with Cloud providers is the same. For a standard configuration featuring one BGP router and one IP address range, you need four pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:

 The following values are generic, enter your specific values into the fields where applicable.

Extract the kubeconfig and deploy a config map for MetalLB using the following command:

```
dkp get kubeconfig -c ${DKP_CLUSTER_NAME} > ${DKP_CLUSTER_NAME}.conf
```

Deploy MetalLB Configuration with the command below:

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
      - peer-address: 10.0.0.1
        peer-asn: 64501
        my-asn: 64500
    address-pools:
      - name: default
        protocol: bgp
        addresses:
          - 192.168.10.0/24
EOF
```

When complete, run the following `kubectl` command:

```
kubectl apply -f metallb-conf.yaml
```

5.13.4.7.3 Next Step:

[vSphere: Install Kommander](#) (see page 449)

5.13.4.8 vSphere: Install Kommander

You have installed the Konvoy component and created a cluster. Now it is time to Install Kommander which will allow you to access the UI and attach new or existing clusters to monitor.

5.13.4.8.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.13.4.8.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```


2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```


3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init > kommander.yaml
```

4. *If required:* Customize your `kommander.yaml`.

 See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, [Rook Ceph customization](#) (Pre-provisioned envs (see page 1405)), etc.

5.13.4.8.3 Enable DKP Catalog Applications and Install Kommander

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
```

```
ref:
  tag: v2.6.1
```

⚠ If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

☰ Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.13.4.8.4 Next Step:

[vSphere: Verify Install and Log in to the UI](#) (see page 451)

5.13.4.9 vSphere: Verify Install and Log in to the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.

☰ **NOTE:** If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.13.4.9.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

5.13.4.9.2.2 Next Steps:


[vSphere: Create a Managed Cluster Using the DKP CLI \(see page 454\)](#)

5.13.4.10 vSphere: Create a Managed Cluster Using the DKP CLI


Enterprise

Gov Advanced

Creating a vSphere Managed cluster with the DKP CLI assumes that you already fulfilled all of the prerequisites and successfully created a vSphere Management cluster. Use this procedure to create a Managed vSphere cluster.

 When creating [Managed clusters \(see page 81\)](#), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters \(see page 81\)](#)!

5.13.4.10.1 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:
 **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace \(see page 650\)](#).

```
kubectrl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:


```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```


5.13.4.10.2 Name your Cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the `CLUSTER_NAME` environment variable with the command:

```
export CLUSTER_NAME=<my-managed-vsphere-cluster>
```

5.13.4.10.3 Create a New vSphere Kubernetes Cluster

 DKP uses the vsphere CSI driver as the [default storage provider](#) (see page 91). Use a [Kubernetes CSI](#)³⁸⁷ compatible storage that is suitable for production.


 The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above. If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273744255/2.6%2BAttach%2Ba%2BKubernetes%2BCluster> (see page 752).

Follow these steps to create the cluster:

1. Use the following command to set the environment variables for vSphere:

```
export VSPHERE_SERVER=example.vsphere.url
export VSPHERE_USERNAME=user@example.vsphere.url
export VSPHERE_PASSWORD=example_password
```

2. Generate the Kubernetes cluster objects by copying and editing this command to include the correct values, including the VM template name you assigned in the previous procedure:

-  • To increase [Dockerhub's rate limit](#)³⁸⁸ use your Dockerhub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.
- The following example shows a common configuration. See [dkp create cluster](#) (see page 1662) reference for the full list of cluster creation options:

³⁸⁷ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

³⁸⁸ <https://docs.docker.com/docker-hub/download-rate-limit/>

```

dkp create cluster vsphere \
  --cluster-name ${MANAGED_CLUSTER_NAME} \
  --additional-tags=owner=$(whoami) \
  --namespace ${WORKSPACE_NAMESPACE} \
  --network <NETWORK_NAME> \
  --control-plane-endpoint-host <xxx.yyy.zzz.000> \
  --data-center <DATACENTER_NAME> \
  --data-store <DATASTORE_NAME> \
  --folder <FOLDER_NAME> \
  --server <VCENTER_API_SERVER_URL> \
  --ssh-public-key-file <SSH_PUBLIC_KEY_FILE> \
  --resource-pool <RESOURCE_POOL_NAME> \
  --virtual-ip-interface <ip_interface_name> \
  --vm-template <TEMPLATE_NAME> \
  --kubeconfig=<management-cluster-kubeconfig-path> \

```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1039\)](#).

5.13.4.10.4 Retrieve the kubeconfig and Explore New vSphere Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:

```

dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf

```

2. List the nodes with the following command:

```

kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes


```

3. List the pods with the following command:

NOTE: Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.13.4.10.4.1 Manually Attach a DKP CLI Cluster to the Management Cluster

 These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [\[attach it in the UI\]](#)(link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-
template='{{.data.value}}{{ "\n"}}
```

- This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new attached-cluster-kubeconfig.yaml file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

- Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

- Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

- You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubectl get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

If you have existing clusters or want to create other new clusters to attach, there are many ways to attach a cluster with various requirements and restrictions. To see all the options, visit the section in documentation [Day 2 - Attach an Existing Kubernetes Cluster](#) (see page 752).

At this point, you can create more clusters, perform [other configuration tasks](#) (see page 1036), or proceed to [Day 2 Operations](#) (see page 582).

5.13.5 vSphere Air-gapped Install

This section provides instructions to install DKP in a vSphere **air-gapped** environment.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)



NOTE: For air-gapped, ensure you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, so you can extract the tarball on next page.

5.13.5.1 Further vSphere Prerequisites

- [vSphere: VMware Prerequisites](#) (see page 440)
- <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273354355/2.6%2BvSphere%2BStorage%2BOptions> (see page 440)
- <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273681241/2.6%2BvSphere%2BMinimum%2BUser%2BPermissions> (see page 438)

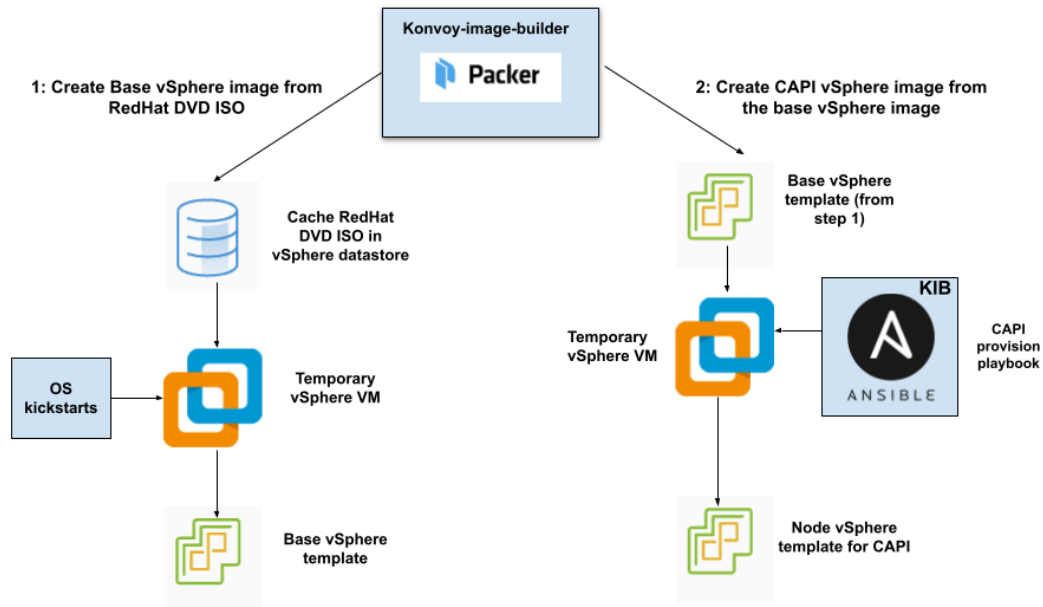
5.13.5.2 Next Step:

[vSphere Air-gapped: Image Creation Overview](#) (see page 459)

5.13.5.3 vSphere Air-gapped: Image Creation Overview

5.13.5.3.1 Image Creation Process

This diagram illustrates the image creation process:



The workflow on the left shows the creation of a base OS image in the vCenter vSphere client using inputs from Packer. The workflow on the right shows how DKP uses that same base OS image to create CAPI-enabled VM images for your cluster.

After creating the base image, the DKP image builder uses it to create a CAPI-enabled vSphere template that includes the Kubernetes objects for the cluster. You can use that resulting template with the DKP `create cluster` command to create the VM nodes in your cluster directly on a vCenter server. From that point, you can use DKP to provision and manage your cluster.

5.13.5.3.2 Next Step:

[vSphere Air-gapped: Create an Image \(see page 460\)](#)

5.13.5.4 vSphere Air-gapped: Create an Image

Creating a base OS image from DVD ISO files is a one-time process. The base OS image file is created in the vSphere Client for use in the vSphere VM template. Therefore, the base OS image is used by [Konvoy Image Builder \(see page 1461\)](#) (KIB) to create a VM template to configure Kubernetes nodes by the DKP vSphere provider.

5.13.5.4.1 Create the Base OS Image

For vSphere, a username is populated by `SSH_USERNAME` and the user can use authorization via `SSH_PASSWORD` or `SSH_PRIVATE_KEY_FILE` environment variables and required by default by packer. This user should have administrator privileges. It is possible to configure a custom user and password when building the OS image, however, that requires the Konvoy Image Builder (KIB) configuration to be [overridden](#) (see page 1513).

While creating the base OS image, it is important to take into consideration the following elements:


- **Storage configuration:** D2iQ recommends customizing disk partitions and not configuring a SWAP partition.
- **Network configuration:** as KIB must download and install packages, activating the network is required.
- **Connect to Red Hat:** if using RHEL, registering with Red Hat is required to configure software repositories and install software packages.
- **Software selection:** D2iQ recommends choosing **Minimal Install**.
- DKP recommends to install with the packages provided by the operating system package managers. Use the version that corresponds to the major version of your operating system.

5.13.5.4.2 Next Step:

vSphere Air-gapped: [Docker Registry](#) (see page 461)

5.13.5.5 vSphere Air-gapped: Load the Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine](#) (see page 1285) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.

 If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see page 1533) page for more information.

1. [Download](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, and extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
```



```
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

- Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```



It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

5.13.5.5.1 Kommander Load Images

If you are operating in an air-gapped environment, a [local container registry](#) (see page 1533) containing all the necessary installation images, including the Kommander images is required. See below for how to push the necessary images to this registry.

5.13.5.5.2 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-
registry $REGISTRY_ADDRESS
```

5.13.5.5.3 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-
v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.13.5.5.4 Next Step:

[vSphere Air-gapped: Create a CAPI VM Template Image](#) (see page 463)

5.13.5.6 vSphere Air-gapped: Create a CAPI VM Template

You must have at least one image before creating a new cluster. As long as you have an image, this step in your configuration is not required each time since that image can be used to spin up a new cluster. However, if you need different images for different environments or providers, you will need to create a new custom image.

Using [KIB](#) (see page 1461), you can create your VM template without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2/kib
```

2. Follow the instructions to build a vSphere template below and set the override `--overrides overrides/offline.yaml` flag.

5.13.5.6.1 Create a vSphere Template for Your Cluster from a Base OS Image

Using the base OS image created in a previous procedure, DKP creates the new vSphere template directly on the vCenter server.

1. Set the following vSphere environment variables on the bastion VM host:

```
export VSPHERE_SERVER=your_vCenter_APIserver_URL
export VSPHERE_USERNAME=your_vCenter_user_name
export VSPHERE_PASSWORD=your_vCenter_password
```

2. Copy the base OS image file created in the vSphere Client to your desired location on the bastion VM host, and make a note of the path and file name.
3. Create an `image.yaml` file and add the following variables for vSphere. DKP uses this file and these variables as inputs in the next step.

⚠ NOTE: This example is Ubuntu 20.04. You will need to replace OS name below based on your OS. See other [YAML examples](#)³⁸⁹ for copy and paste below last step.

³⁸⁹ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

---
download_images: true
build_name: "ubuntu-2004"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-
init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: "<VSPHERE_CLUSTER_NAME>"
  datacenter: "<VSPHERE_DATACENTER_NAME>"
  datastore: "<VSPHERE_DATASTORE_NAME>"
  folder: "<VSPHERE_FOLDER>"
  insecure_connection: "false"
  network: "<VSPHERE_NETWORK>"
  resource_pool: "<VSPHERE_RESOURCE_POOL>"
  template: "os-qualification-templates/d2iq-base-Ubuntu-20.04" # change
default value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "ubuntu2004-64"
  # goss params
  distribution: "ubuntu"
  distribution_version: "20.04"
# Use following overrides to select the authentication method that can be used
with base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key
will be ignored

```

4. Create a vSphere VM template with your variation of the following command:

```
konvoy-image build images/ova/<image.yaml>
```

Any additional configurations can be added to this command using `--overrides` flags as shown below:

- a. Any credential overrides: `--overrides overrides.yaml`
 - b. for FIPS, add this flag: `--overrides overrides/fips.yaml`
 - c. for air-gapped, add this flag: `--overrides overrides/offline-fips.yaml`
5. The [Konvoy Image Builder](#) (see page 1487) (KIB) uses the values in `image.yaml` and the input base OS image to create a vSphere template directly on the vCenter server. This template contains the

required artifacts needed to create a Kubernetes cluster.

When KIB provisions the OS [image](#)³⁹⁰ successfully, it creates a manifest file. The `artifact_id` field of this file contains the name of the AMI ID (AWS), template name (vSphere), or image name (GCP/Azure), for example:

```
{
  "name": "vsphere-clone",
  "builder_type": "vsphere-clone",
  "build_time": 1644985039,
  "files": null,
  "artifact_id": "konvoy-ova-vsphere-rhel-84-1.21.6-1644983717",
  "packer_run_uuid": "260e8110-77f8-ca94-e29e-ac7a2ae779c8",
  "custom_data": {
    "build_date": "2022-02-16T03:55:17Z",
    "build_name": "vsphere-rhel-84",
    "build_timestamp": "1644983717",
    [...]
  }
}
```

Recommendation: Now we can now see the template created in our vCenter, it is best to rename it to `dkp-<DKP_VERSION>-k8s-<K8S_VERSION>-<DISTRO>`, like `dkp-2.6.2-k8s-1.26.14-ubuntu` to keep templates organized.

6. Next steps are to deploy a DKP cluster using your vSphere template.

5.13.5.6.1.1 Additional OS YAML files for copy/paste:

RHEL 8.6

```
---
download_images: true
build_name: "rhel-86"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
```

³⁹⁰ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

network: ""
resource_pool: ""
template: "base-rhel-8.6" # change default value with your base template name
vsphere_guest_os_type: "rhel8_64Guest"
guest_os_type: "rhel8-64"
# goss params
distribution: "RHEL"
distribution_version: "8.6"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

Ubuntu 20.04

```

---
download_images: true
build_name: "ubuntu-2004"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-
vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "base-ubuntu-20.04" # change default value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "ubuntu2004-64"
  # goss params
  distribution: "ubuntu"
  distribution_version: "20.04"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

Rocky Linux 9.1

```

---
download_images: true
build_name: "rocky-91"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "base-rocky-9.1" # change default value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "rocky9-64"
  # goss params
  distribution: "rocky"
  distribution_version: "9.1"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

5.13.5.6.2 Next Step:

[vSphere Air-gapped: Create the Management Cluster](#) (see page 467)

5.13.5.7 vSphere Air-gapped: Create the Management Cluster

This page of instructions will create a self-managed air-gapped management cluster. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.

5.13.5.7.1 Name your Cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the CLUSTER_NAME environment variable with the command:

```
export CLUSTER_NAME=<my-vsphere-cluster>
```



DKP uses the vsphere CSI driver as the [default storage provider](#) (see page 91). Use a [Kubernetes CSI](#)³⁹¹ compatible storage that is suitable for production.

5.13.5.7.1.1 Create a New vSphere Kubernetes Cluster



The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273744255/2.6%2BAttach%2Ba%2BKubernetes%2BCluster> (see page 752).

Use the following steps to create a new, air-gapped vSphere cluster.

1. Configure your cluster to use an existing local registry as a mirror when attempting to pull images: **IMPORTANT:** The image must be created by [Konvoy Image Builder](#) (see page 1487) in order to use the registry mirror feature.

```
export REGISTRY_URL=<https/http>://<registry-address>:<registry-port>
export REGISTRY_CA=<path to the CA on the bastion>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

- `REGISTRY_URL` : the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
- `REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. Konvoy will configure the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
- `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
- `REGISTRY_PASSWORD` : optional if username is not set.

2. Load the image, using either the `docker` or `podman` command:

³⁹¹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

```
docker load -i konvoy-bootstrap-image-v2.6.2.tar
```

OR

```
podman load -i konvoy-bootstrap-image-v2.6.2.tar
```

3. Create a Kubernetes cluster by copying the following command and substituting the valid values for your environment:

```
dkp create cluster vsphere \
  --cluster-name=${MANAGED_CLUSTER_NAME} \
  --additional-tags=owner=$(whoami) \
  --kubeconfig=<management-cluster-kubeconfig-path> \
  --namespace ${WORKSPACE_NAMESPACE} \
  --network <NETWORK_NAME> \
  --control-plane-endpoint-host <CONTROL_PLANE_IP> \
  --data-center <DATACENTER_NAME> \
  --data-store <DATASTORE_NAME> \
  --folder <FOLDER_NAME> \
  --server <VCENTER_API_SERVER_URL> \
  --ssh-public-key-file </path/to/key.pub> \
  --resource-pool <RESOURCE_POOL_NAME> \
  --vm-template konvoy-ova-vsphere-os-release-k8s_release-vsphere-timestamp \
  --virtual-ip-interface <ip_interface_name> \
  --extra-sans "127.0.0.1" \
  --registry-mirror-url=${REGISTRY_URL} \
  --registry-mirror-cacert=${REGISTRY_CA} \
  --registry-mirror-username=${REGISTRY_USERNAME} \
  --registry-mirror-password=${REGISTRY_PASSWORD} \
  --self-managed
```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

= For bootstrap and custom YAML cluster creation, refer to the [Custom Installation and Additional Infrastructure Tools](#) (see page 1036) section of the documentation for vSphere: [vSphere Infrastructure](#) (see page 1254)

5.13.5.7.2 Retrieve the kubeconfig and Explore New vSphere Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:



NOTE: Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.13.5.7.2.1 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>³⁹²³⁹³.

As they progress, the controllers create Events, which you can also monitor using the command:

```
kubectl get events | grep ${CLUSTER_NAME}
```

³⁹² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

³⁹³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

For brevity, this example uses `grep`. You can also use separate commands to get Events for specific objects, such as `kubectl get events --field-selector involvedObject.kind="VSphereCluster"` and `kubectl get events --field-selector involvedObject.kind="VSphereMachine"`.

5.13.5.7.2.2 Next Step:

[vSphere Air-gapped: Configure Metal LB \(see page 471\)](#)

5.13.5.8 vSphere Air-gapped: Configure Metal LB

If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can proceed to [vSphere Air-gapped: Install Kommander \(see page 473\)](#) and continue the installation process.

To use MetalLB, create a MetalLB configMap for your vSphere infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)

Select one of the following procedures to create your MetalLB manifest for further editing.

5.13.5.8.1 Layer 2 configuration

Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.



- MetalLB IP address ranges/CIDRs should be within the node's primary network [subnet \(see page 1046\)](#).
- MetalLB IP address ranges/CIDRs and node subnet should not conflict with the Kubernetes cluster pod and service subnets.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:



The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    address-pools:
    - name: default
      protocol: layer2
      addresses:
      - 192.168.1.240-192.168.1.250
EOF
```

When complete, run the following `kubectl` command.


```
kubectl apply -f metallb-conf.yaml
```

5.13.5.8.2 BGP configuration

BGP, like any routing protocol, requires a mesh, so you configure BGP between the nodes. If you are in a Data Center, peer with the BGP routers. The instructions for use with Cloud providers is the same. For a standard configuration featuring one BGP router and one IP address range, you need four pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:

 The following values are generic, enter your specific values into the fields where applicable.

Extract the kubeconfig and deploy a config map for MetalLB using the following command:

```
dkp get kubeconfig -c ${DKP_CLUSTER_NAME} > ${DKP_CLUSTER_NAME}.conf
```

Deploy MetalLB Configuration with the command below:

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
```

```

kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
      - peer-address: 10.0.0.1
        peer-asn: 64501
        my-asn: 64500
    address-pools:
      - name: default
        protocol: bgp
        addresses:
          - 192.168.10.0/24
EOF

```

When complete, run the following `kubectl` command:

```
kubectl apply -f metallb-conf.yaml
```

5.13.5.8.3 Next Step:

[vSphere Air-gapped: Install Kommander](#) (see page 473)

5.13.5.9 vSphere Air-gapped: Install Kommander

5.13.5.9.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Ensure you have loaded all necessary images for your configuration. See [Load the Images into Your Registry: Air-gapped Environments](#) (see page 1400).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.13.5.9.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```


2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init --airgapped > kommander.yaml
```

4. *If required:* Customize your `kommander.yaml`.

 See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, [Rook Ceph customization](#) (Pre-provisioned envs (see page 1405)), etc.


5.13.5.9.3 Enable DKP Catalog Applications and Install Kommander in an Air-gapped Environment



If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add the following values to enable DKP Catalog Applications:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      path: ./dkp-catalog-applications-v2.6.1.tar.gz
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf \
```

```

--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz
\
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-
v2.6.1.tar.gz

```

5.13.5.9.4 Next Step:

[vSphere Air-gapped: Verify Install and Log in to the UI \(see page 475\)](#)

5.13.5.10 vSphere Air-gapped: Verify Install and Log in to the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.



NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the install command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```

helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met

```

```
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.13.5.10.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

5.13.5.10.2 Log in to the UI with Kommander

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{{ "\n"}}Password: {{.data.password|base64decode}}
{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/dashboard{{ "\n"}}
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.13.5.10.2.1 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

5.13.5.10.2.2 Next Steps:

[vSphere Air-gapped: Create a Managed Cluster Using the DKP CLI](#) (see [page 477](#))

5.13.5.11 vSphere Air-gapped: Create a Managed Cluster Using the DKP CLI

Creating an air-gapped vSphere Managed cluster with the DKP CLI assumes that you already fulfilled all of the prerequisites and successfully created a vSphere Management cluster. Use this procedure to create a Managed vSphere cluster.



When creating [Managed clusters](#) (see [page 81](#)), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see [page 81](#))!

5.13.5.11.1 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:
⚠ NOTE: If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see page 650).

```
kubectl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.13.5.11.2 Name your Cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the `CLUSTER_NAME` environment variable with the command:

```
export CLUSTER_NAME=<my-managed-vsphere-cluster>
```

⚠ DKP uses the vsphere CSI driver as the [default storage provider](#) (see page 91). Use a [Kubernetes CSI](#)³⁹⁴ compatible storage that is suitable for production.

5.13.5.11.2.1 Create a New vSphere Kubernetes Cluster



The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273744255/2.6%2BAAttach%2Ba%2BKubernetes%2BCluster> (see page 752).

³⁹⁴ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

Use the following steps to create a new, air-gapped vSphere cluster.

1. Configure your cluster to use an existing local registry as a mirror when attempting to pull images: **IMPORTANT:** The image must be created by [Konvoy Image Builder](#) (see page 1487) in order to use the registry mirror feature.

```
export REGISTRY_URL=<https/http>://<registry-address>:<registry-port>
export REGISTRY_CA=<path to the CA on the bastion>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

- `REGISTRY_URL` : the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
- `REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. Konvoy will configure the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
- `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
- `REGISTRY_PASSWORD` : optional if username is not set.

2. Load the image, using either the `docker` or `podman` command:

```
docker load -i konvoy-bootstrap-image-v2.6.2.tar
```

OR

```
podman load -i konvoy-bootstrap-image-v2.6.2.tar
```

3. Create a Kubernetes cluster by copying the following command and substituting the valid values for your environment:

```
dkp create cluster vsphere \
  --cluster-name=${MANAGED_CLUSTER_NAME} \
  --additional-tags=owner=$(whoami) \
  --kubeconfig=<management-cluster-kubeconfig-path> \
  --namespace ${WORKSPACE_NAMESPACE} \
  --network <NETWORK_NAME> \
  --control-plane-endpoint-host <CONTROL_PLANE_IP> \
  --data-center <DATACENTER_NAME> \
  --data-store <DATASTORE_NAME> \
  --folder <FOLDER_NAME> \
  --server <VCENTER_API_SERVER_URL> \
  --ssh-public-key-file </path/to/key.pub> \
  --resource-pool <RESOURCE_POOL_NAME> \
```

```
--vm-template konvoy-ova-vsphere-os-release-k8s_release-vsphere-timestamp \  
--virtual-ip=interface <ip_interface_name> \  
--extra-sans "127.0.0.1" \  
--registry-mirror-url=${REGISTRY_URL} \  
--registry-mirror-cacert=${REGISTRY_CA} \  
--registry-mirror-username=${REGISTRY_USERNAME} \  
--registry-mirror-password=${REGISTRY_PASSWORD} \  
--kubeconfig=<management-cluster-kubeconfig-path> \  

```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

5.13.5.11.3 Retrieve the kubeconfig and Explore New vSphere Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig  
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $  
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:

NOTE: Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.13.5.11.3.1 Manually Attach a DKP CLI Cluster to the Management Cluster



These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

7. This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```

apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>

```

8. Create this secret in the desired workspace:

```

kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}

```

9. Create this `kommandercluster` object to attach the cluster to the workspace:

```

cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF

```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```

kubectl get kommanderclusters -A

```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

If you have existing clusters or want to create other new clusters to attach, there are many ways to attach a cluster with various requirements and restrictions. To see all the options, visit the section in documentation [Day 2 - Attach an Existing Kubernetes Cluster](#) (see page 752).

At this point, you can create more clusters, perform [other configuration tasks](#) (see page 1036), or proceed to [Day 2 Operations](#) (see page 582).

5.13.6 vSphere FIPS Install

This section provides instructions to install DKP in a vSphere with FIPS.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.13.6.1 Further vSphere Prerequisites

- [vSphere: VMware Prerequisites](#) (see page 440)
- <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/162430985/vSphere+Prerequisites+-+Storage+Options> (see page 440)
- <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/162430977/vSphere+Prerequisites+-+Minimum+User+Permissions> (see page 438)

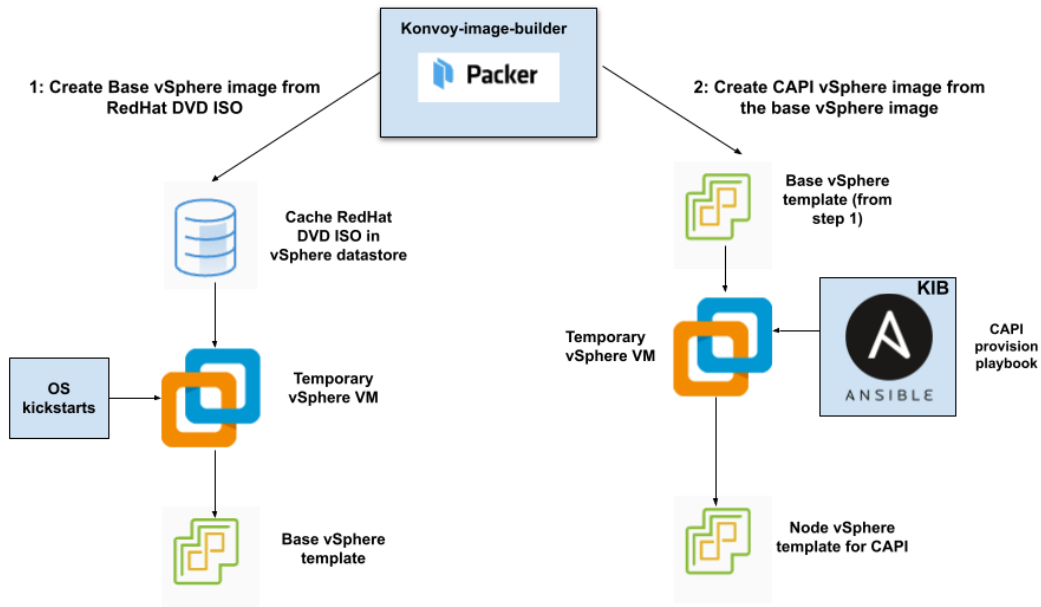
5.13.6.2 Next Step:

[vSphere FIPS: Image Creation Overview](#) (see page 483)

5.13.6.3 vSphere FIPS: Image Creation Overview

5.13.6.3.1 Image Creation Process

This diagram illustrates the image creation process:



The workflow on the left shows the creation of a base OS image in the vCenter vSphere client using inputs from Packer. The workflow on the right shows how DKP uses that same base OS image to create CAPI-enabled VM images for your cluster.

After creating the base image, the DKP image builder uses it to create a CAPI-enabled vSphere template that includes the Kubernetes objects for the cluster. You can use that resulting template with the DKP `create cluster` command to create the VM nodes in your cluster directly on a vCenter server. From that point, you can use DKP to provision and manage your cluster.

5.13.6.3.2 Next Step:

[vSphere FIPS: Create an Image](#) (see page 484)

5.13.6.4 vSphere FIPS: Create an Image

Creating a base OS image from DVD ISO files is a one-time process. The base OS image file is created in the vSphere Client for use in the vSphere VM template. Therefore, the base OS image is used by [Konvoy Image Builder](#) (see page 1461) (KIB) to create a VM template to configure Kubernetes nodes by the DKP vSphere provider.

5.13.6.4.1 Create the Base OS Image

For vSphere, a username is populated by `SSH_USERNAME` and the user can use authorization via `SSH_PASSWORD` or `SSH_PRIVATE_KEY_FILE` environment variables and required by default by packer. This user should have administrator privileges. It is possible to configure a custom user and password when building the OS image, however, that requires the Konvoy Image Builder (KIB) configuration to be [overridden](#) (see page 1513).

While creating the base OS image, it is important to take into consideration the following elements:

- **Storage configuration:** D2iQ recommends customizing disk partitions and not configuring a SWAP partition.
- **Network configuration:** as KIB must download and install packages, activating the network is required.
- **Connect to Red Hat:** if using RHEL, registering with Red Hat is required to configure software repositories and install software packages.
- **Software selection:** D2iQ recommends choosing **Minimal Install**.
- DKP recommends to install with the packages provided by the operating system package managers. Use the version that corresponds to the major version of your operating system.

5.13.6.4.2 Next Step:

The next step is to [vSphere FIPS: Create a CAPI VM Template Image](#) (see page 485). The image contains the CAPI and Kubernetes objects.

5.13.6.5 vSphere FIPS: Create a CAPI VM Template

5.13.6.5.1 Prerequisites

- Users need to create a [vSphere FIPS: Create an Image](#) (see page 487) before starting this procedure.
- Build image with Konvoy Image Builder (KIB)
- In step 4 of the following section, ensure you use `--overrides overrides/fips.yaml`

5.13.6.5.2 Create a vSphere Template for Your Cluster from a Base OS Image

Using the base OS image created in a previous procedure, DKP creates the new vSphere template directly on the vCenter server.

1. Set the following vSphere environment variables on the bastion VM host:

```
export VSPHERE_SERVER=your_vCenter_APIserver_URL
export VSPHERE_USERNAME=your_vCenter_user_name
export VSPHERE_PASSWORD=your_vCenter_password
```


- Copy the base OS image file created in the vSphere Client to your desired location on the bastion VM host, and make a note of the path and file name.
- Create an `image.yaml` file and add the following variables for vSphere. DKP uses this file and these variables as inputs in the next step.

⚠ NOTE: This example is Ubuntu 20.04. You will need to replace OS name below based on your OS. See other [YAML examples](#)³⁹⁵ for copy and paste below last step.

```
---
download_images: true
build_name: "ubuntu-2004"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: "<VSPHERE_CLUSTER_NAME>"
  datacenter: "<VSPHERE_DATACENTER_NAME>"
  datastore: "<VSPHERE_DATASTORE_NAME>"
  folder: "<VSPHERE_FOLDER>"
  insecure_connection: "false"
  network: "<VSPHERE_NETWORK>"
  resource_pool: "<VSPHERE_RESOURCE_POOL>"
  template: "os-qualification-templates/d2iq-base-Ubuntu-20.04" # change
  default value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "ubuntu2004-64"
  # goss params
  distribution: "ubuntu"
  distribution_version: "20.04"
# Use following overrides to select the authentication method that can be used
with base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key
will be ignored
```

- Create a vSphere VM template with your variation of the following command:

```
konvoy-image build images/ova/<image.yaml>
```

³⁹⁵ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

Any additional configurations can be added to this command using `--overrides` flags as shown below:

- a. Any credential overrides: `--overrides overrides.yaml`
 - b. for FIPS, add this flag: `--overrides overrides/fips.yaml`
 - c. for air-gapped, add this flag: `--overrides overrides/offline-fips.yaml`
5. The [Konvoy Image Builder](#) (see page 1487) (KIB) uses the values in `image.yaml` and the input base OS image to create a vSphere template directly on the vCenter server. This template contains the required artifacts needed to create a Kubernetes cluster.

When KIB provisions the OS `image`³⁹⁶ successfully, it creates a manifest file. The `artifact_id` field of this file contains the name of the AMI ID (AWS), template name (vSphere), or image name (GCP/Azure), for example:

```
{
  "name": "vsphere-clone",
  "builder_type": "vsphere-clone",
  "build_time": 1644985039,
  "files": null,
  "artifact_id": "konvoy-ova-vsphere-rhel-84-1.21.6-1644983717",
  "packer_run_uuid": "260e8110-77f8-ca94-e29e-ac7a2ae779c8",
  "custom_data": {
    "build_date": "2022-02-16T03:55:17Z",
    "build_name": "vsphere-rhel-84",
    "build_timestamp": "1644983717",
    [...]
  }
}
```

Recommendation: Now we can now see the template created in our vCenter, it is best to rename it to `dkp-<DKP_VERSION>-k8s-<K8S_VERSION>-<DISTRO>`, like `dkp-2.6.2-k8s-1.26.14-ubuntu` to keep templates organized.

6. Next steps are to deploy a DKP cluster using your vSphere template.

5.13.6.5.2.1 Additional OS YAML files for copy/paste:

RHEL 8.6

```
---
download_images: true
build_name: "rhel-86"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
```

³⁹⁶ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "base-rhel-8.6" # change default value with your base template name
  vsphere_guest_os_type: "rhel8_64Guest"
  guest_os_type: "rhel8-64"
  # goss params
  distribution: "RHEL"
  distribution_version: "8.6"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

Ubuntu 20.04

```

---
download_images: true
build_name: "ubuntu-2004"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-
vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "base-ubuntu-20.04" # change default value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "ubuntu2004-64"
  # goss params
  distribution: "ubuntu"
  distribution_version: "20.04"

```

```
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored
```

Rocky Linux 9.1

```
---
download_images: true
build_name: "rocky-91"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-
vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "base-rocky-9.1" # change default value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "rocky9-64"
  # goss params
  distribution: "rocky"
  distribution_version: "9.1"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored
```

5.13.6.5.3 Next Step:

[vSphere FIPS: Create the Management Cluster \(see page 490\)](#)

5.13.6.6 vSphere FIPS: Create the Management Cluster

This page instructions will create a self-managed cluster for use as the Management cluster. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing. Additional considerations exist when using FIPS. Alter the Create a Kubernetes cluster command by using the FIPS example, or similar, when executing the command to create a cluster.

5.13.6.6.1 Deploying a Cluster in FIPS mode

In order to create a cluster in FIPS mode, we must inform the bootstrap controllers of the appropriate image repository and version tags of the official D2iQ FIPS builds of Kubernetes.

5.13.6.6.1.1 Supported FIPS Builds

Component	Repository	Version
Kubernetes	docker.io/mesosphere³⁹⁷	v1.26.14+fips.0
etcd	docker.io/mesosphere³⁹⁸	3.5.10+fips.0

5.13.6.6.2 Name your Cluster


1. Give your cluster a unique name suitable for your environment.
2. Set the CLUSTER_NAME environment variable with the command:


```
export CLUSTER_NAME=<my-vsphere-cluster>
```

³⁹⁷ <https://hub.docker.com/layers/mesosphere/kube-apiserver/v1.26.14/images/sha256-3c559c75123fa40d588bc2d9050938a0b7dd1dab4e38a4fd24847cd0c0f24694?context=explore>

³⁹⁸ https://hub.docker.com/layers/mesosphere/etcd/v3.5.10_fips.0/images/sha256-61df32b8ec1e6229d8bd94caad07331c7e0c0b8e1947af8152c69eb2f04fb771?context=explore

5.13.6.6.3 Create a New vSphere Kubernetes Cluster

 DKP uses the vsphere CSI driver as the [default storage provider](#) (see page 91). Use a [Kubernetes CSI](#)³⁹⁹ compatible storage that is suitable for production.

 The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.


If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273744255/2.6%2BAttach%2Ba%2BKubernetes%2BCluster> (see page 752).

Follow these steps to create the cluster:

1. Use the following command to set the environment variables for vSphere:

```
export VSPHERE_SERVER=example.vsphere.url
export VSPHERE_USERNAME=user@example.vsphere.url
export VSPHERE_PASSWORD=example_password
```

2. Generate the Kubernetes cluster objects by copying and editing this command to include the correct values, including the VM template name you assigned in the previous procedure:

-  • To increase [Dockerhub's rate limit](#)⁴⁰⁰ use your Dockerhub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.
- The following example shows a common configuration. See [dkp create cluster](#) (see page 1662) reference for the full list of cluster creation options:

```
dkp create cluster vsphere \
```

³⁹⁹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁴⁰⁰ <https://docs.docker.com/docker-hub/download-rate-limit/>

```

--cluster-name ${CLUSTER_NAME} \
--network <NETWORK_NAME> \
--control-plane-endpoint-host <xxx.yyy.zzz.000> \
--data-center <DATACENTER_NAME> \
--data-store <DATASTORE_NAME> \
--folder <FOLDER_NAME> \
--server <VCENTER_API_SERVER_URL> \
--ssh-public-key-file <SSH_PUBLIC_KEY_FILE> \
--resource-pool <RESOURCE_POOL_NAME> \
--vm-template <TEMPLATE_NAME> \
--virtual-ip-interface <ip_interface_name> \
--self-managed \
--kubernetes-version=v1.26.14+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--etcd-image-repository=docker.io/mesosphere \
--etcd-version=3.5.10+fips.0 \
--self-managed

```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

= For bootstrap and custom YAML cluster creation, refer to the [Custom Installation and Additional Infrastructure Tools](#) (see page 1036) section of the documentation for vSphere: [vSphere Infrastructure](#) (see page 1254) or [FIPS 140-2 Compliance](#) (see page 1526) for more FIPS details.

5.13.6.6.3.1 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>⁴⁰¹⁴⁰².

5.13.6.6.3.2 Next Step:

[vSphere FIPS: Configure Metal LB](#) (see page 493)

⁴⁰¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

⁴⁰² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

5.13.6.7 vSphere FIPS: Configure Metal LB

If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can proceed to [vSphere FIPS: Install Kommander \(see page 495\)](#) and continue the installation process.

To use MetalLB, create a MetalLB configMap for your vSphere infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)

Select one of the following procedures to create your MetalLB manifest for further editing.

5.13.6.7.1 Layer 2 configuration

Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.



- MetalLB IP address ranges/CIDRs should be within the node's primary network [subnet \(see page 1046\)](#).
- MetalLB IP address ranges/CIDRs and node subnet should not conflict with the Kubernetes cluster pod and service subnets.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:



The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    address-pools:
    - name: default
      protocol: layer2
```



```

addresses:
- 192.168.1.240-192.168.1.250
EOF

```

When complete, run the following `kubectl` command.


```
kubectl apply -f metallb-conf.yaml
```

5.13.6.7.2 BGP configuration

BGP, like any routing protocol, requires a mesh, so you configure BGP between the nodes. If you are in a Data Center, peer with the BGP routers. The instructions for use with Cloud providers is the same. For a standard configuration featuring one BGP router and one IP address range, you need four pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:

 The following values are generic, enter your specific values into the fields where applicable.

Extract the kubeconfig and deploy a config map for MetalLB using the following command:

```
dkp get kubeconfig -c ${DKP_CLUSTER_NAME} > ${DKP_CLUSTER_NAME}.conf
```

Deploy MetalLB Configuration with the command below:

```

cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
    - peer-address: 10.0.0.1
      peer-asn: 64501
      my-asn: 64500
    address-pools:
    - name: default

```

```

protocol: bgp
addresses:
- 192.168.10.0/24
EOF

```

When complete, run the following `kubectl` command:

```
kubectl apply -f metallb-conf.yaml
```

5.13.6.7.3 Next Step:

[vSphere FIPS: Install Kommander](#) (see page 495)

5.13.6.8 vSphere FIPS: Install Kommander

You have installed the Konvoy component and created a cluster. Now it is time to Install Kommander which will allow you to access the UI and attach new or existing clusters to monitor.

5.13.6.8.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.13.6.8.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```


2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:


```
dkp install kommander --init > kommander.yaml
```

4. *If required:* Customize your `kommander.yaml`.

 See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include:


Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, [Rook Ceph customization \(Pre-provisioned envs \(see page 1405\)\)](#), etc.

5.13.6.8.3 Enable DKP Catalog Applications and Install Kommander

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP \(see page 1458\)](#).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
        ref:
          tag: v2.6.1
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File \(see page 87\)](#).

- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.13.6.8.4 Next Step:

[vSphere FIPS: Verify Install and Log in to the UI](#) (see page 497)

5.13.6.9 vSphere FIPS: Verify Install and Log in to the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.



NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
```

```
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.13.6.9.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

5.13.6.9.2 Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{{ "\n"}}Password: {{.data.password|base64decode}}
{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/dashboard{{ "\n"}}
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbJUKVwLJMQL2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.13.6.9.3 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

5.13.6.9.4 Next Step:


[vSphere FIPS: Create a Managed Cluster Using the DKP CLI](#) (see [page 499](#))

5.13.6.10 vSphere FIPS: Create a Managed Cluster Using the DKP CLI


Enterprise

Gov Advanced

Creating a vSphere FIPS Managed cluster with the DKP CLI assumes that you already fulfilled all of the prerequisites and successfully created a vSphere Management cluster. Use this procedure to create a Managed vSphere cluster.

 When creating [Managed clusters](#) (see page 81), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see page 81)!

5.13.6.10.1 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:
 **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see page 650).

```
kubectl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:


```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.13.6.10.2 Name your Cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the `CLUSTER_NAME` environment variable with the command:

```
export CLUSTER_NAME=<my-managed-vsphere-cluster>
```

5.13.6.10.3 Create a New vSphere Kubernetes Cluster

 DKP uses the vsphere CSI driver as the [default storage provider](#) (see page 91). Use a [Kubernetes CSI](#)⁴⁰³ compatible storage that is suitable for production.

⁴⁰³ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

- The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273744255/2.6%2BAttach%2Ba%2BKubernetes%2BCluster> (see page 752).

Follow these steps to create the cluster:

- Use the following command to set the environment variables for vSphere:

```
export VSPHERE_SERVER=example.vsphere.url
export VSPHERE_USERNAME=user@example.vsphere.url
export VSPHERE_PASSWORD=example_password
```

- Generate the Kubernetes cluster objects by copying and editing this command to include the correct values, including the VM template name you assigned in the previous procedure:

- To increase [Dockerhub's rate limit](https://docs.docker.com/docker-hub/download-rate-limit/)⁴⁰⁴ use your Dockerhub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.
- The following example shows a common configuration. See [dkp create cluster](#) (see page 1662) reference for the full list of cluster creation options:

```
dkp create cluster vsphere \
  --cluster-name=${MANAGED_CLUSTER_NAME} \
  --additional-tags=owner=$(whoami) \
  --kubeconfig=<management-cluster-kubeconfig-path> \
  --namespace ${WORKSPACE_NAMESPACE} \
  --network <NETWORK_NAME> \
  --control-plane-endpoint-host <xxx.yyy.zzz.000> \
  --data-center <DATACENTER_NAME> \
  --data-store <DATASTORE_NAME> \
  --folder <FOLDER_NAME> \
  --server <VCENTER_API_SERVER_URL> \
  --ssh-public-key-file <SSH_PUBLIC_KEY_FILE> \
```

⁴⁰⁴ <https://docs.docker.com/docker-hub/download-rate-limit/>


```
--resource-pool <RESOURCE_POOL_NAME> \
--vm-template <TEMPLATE_NAME> \
--virtual-ip-interface <ip_interface_name> \
--kubernetes-version=v1.26.14+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--etcd-image-repository=docker.io/mesosphere \
--etcd-version=3.5.10+fips.0 \
```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1039\)](#).

5.13.6.10.4 Retrieve the kubeconfig and Explore New vSphere Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:


```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:

NOTE: Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.13.6.10.4.1 Manually Attach a DKP CLI Cluster to the Management Cluster

 These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

7. This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```

apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>

```

8. Create this secret in the desired workspace:

```

kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}

```

9. Create this `kommandercluster` object to attach the cluster to the workspace:

```

cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF

```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```

kubectl get kommanderclusters -A

```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

If you have existing clusters or want to create other new clusters to attach, there are many ways to attach a cluster with various requirements and restrictions. To see all the options, visit the section in documentation [Day 2 - Attach an Existing Kubernetes Cluster](#) (see page 752).

At this point, you can create more clusters, perform [other configuration tasks](#) (see page 1036), or proceed to [Day 2 Operations](#) (see page 582).

5.13.7 vSphere FIPS Air-gapped Install

This section provides instructions to install DKP in a vSphere with FIPS in an **air-gapped** environment.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)



NOTE: For air-gapped, ensure you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, so you can extract the tarball on next page.

5.13.7.1 Further vSphere Prerequisites

- [vSphere: VMware Prerequisites](#) (see page 440)
- <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273354355/2.6%2BvSphere%2BStorage%2BOptions> (see page 440)
- <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273681241/2.6%2BvSphere%2BMinimum%2BUser%2BPermissions> (see page 438)


5.13.7.2 Next Step:

[vSphere FIPS Air-gapped: Image Creation Overview](#) (see page 505)

5.13.7.3 vSphere FIPS Air-gapped: Image Creation Overview

5.13.7.3.1 Image Creation Process

This diagram illustrates the image creation process:



The workflow on the left shows the creation of a base OS image in the vCenter vSphere client using inputs from Packer. The workflow on the right shows how DKP uses that same base OS image to create CAPI-enabled VM images for your cluster.

After creating the base image, the DKP image builder uses it to create a CAPI-enabled vSphere template that includes the Kubernetes objects for the cluster. You can use that resulting template with the DKP `create cluster` command to create the VM nodes in your cluster directly on a vCenter server. From that point, you can use DKP to provision and manage your cluster.

5.13.7.3.2 Next Step:

[vSphere FIPS Air-gapped: Create an Image \(see page 506\)](#)

5.13.7.4 vSphere FIPS Air-gapped: Create an Image

Creating a base OS image from DVD ISO files is a one-time process. The base OS image file is created in the vSphere Client for use in the vSphere VM template. Therefore, the base OS image is used by [Konvoy Image Builder \(see page 1461\)](#) (KIB) to create a VM template to configure Kubernetes nodes by the DKP vSphere provider.

5.13.7.4.1 Create the Base OS Image

For vSphere, a username is populated by `SSH_USERNAME` and the user can use authorization via `SSH_PASSWORD` or `SSH_PRIVATE_KEY_FILE` environment variables and required by default by packer. This user should have administrator privileges. It is possible to configure a custom user and password when building the OS image, however, that requires the Konvoy Image Builder (KIB) configuration to be [overridden](#) (see page 1513).

While creating the base OS image, it is important to take into consideration the following elements:


- **Storage configuration:** D2iQ recommends customizing disk partitions and not configuring a SWAP partition.
- **Network configuration:** as KIB must download and install packages, activating the network is required.
- **Connect to Red Hat:** if using RHEL, registering with Red Hat is required to configure software repositories and install software packages.
- **Software selection:** D2iQ recommends choosing **Minimal Install**.
- DKP recommends to install with the packages provided by the operating system package managers. Use the version that corresponds to the major version of your operating system.

5.13.7.4.2 Next Step:

[vSphere FIPS Air-gapped: Docker Registry](#) (see page 507)

5.13.7.5 vSphere FIPS Air-gapped: Load the Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine](#) (see page 1285) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.

 If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see page 1533) page for more information.

1. [Download](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, and extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
```

```
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

- Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```



It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

5.13.7.5.1 Kommander Load Images

If you are operating in an air-gapped environment, a [local container registry](#) (see [page 1533](#)) containing all the necessary installation images, including the Kommander images is required. See below for how to push the necessary images to this registry.

5.13.7.5.2 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-
registry $REGISTRY_ADDRESS
```

5.13.7.5.3 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-
v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

5.13.7.5.4 Next Step:

[vSphere FIPS Air-gapped: Create a CAPI VM Template Image](#) (see page 509)

5.13.7.6 vSphere FIPS Air-gapped: Create a CAPI VM Template

You must have at least one image before creating a new cluster. As long as you have an image, this step in your configuration is not required each time since that image can be used to spin up a new cluster. However, if you need different images for different environments or providers, you will need to create a new custom image.

Using [KIB](#) (see page 1461), you can create your VM template without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2/kib
```

2. Follow the instructions to build a vSphere template below and set the override `--overrides overrides/offline.yaml` flag.

5.13.7.6.1 Create a vSphere Template for Your Cluster from a Base OS Image

Using the base OS image created in a previous procedure, DKP creates the new vSphere template directly on the vCenter server.

1. Set the following vSphere environment variables on the bastion VM host:

```
export VSPHERE_SERVER=your_vCenter_APIserver_URL
export VSPHERE_USERNAME=your_vCenter_user_name
export VSPHERE_PASSWORD=your_vCenter_password
```

2. Copy the base OS image file created in the vSphere Client to your desired location on the bastion VM host, and make a note of the path and file name.
3. Create an `image.yaml` file and add the following variables for vSphere. DKP uses this file and these variables as inputs in the next step.

⚠ NOTE: This example is Ubuntu 20.04. You will need to replace OS name below based on your OS. See other [YAML examples](#)⁴⁰⁵ for copy and paste below last step.

⁴⁰⁵ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>


```

---
download_images: true
build_name: "ubuntu-2004"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-
init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: "<VSPHERE_CLUSTER_NAME>"
  datacenter: "<VSPHERE_DATACENTER_NAME>"
  datastore: "<VSPHERE_DATASTORE_NAME>"
  folder: "<VSPHERE_FOLDER>"
  insecure_connection: "false"
  network: "<VSPHERE_NETWORK>"
  resource_pool: "<VSPHERE_RESOURCE_POOL>"
  template: "os-qualification-templates/d2iq-base-Ubuntu-20.04" # change
default value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "ubuntu2004-64"
  # goss params
  distribution: "ubuntu"
  distribution_version: "20.04"
# Use following overrides to select the authentication method that can be used
with base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key
will be ignored

```

4. Create a vSphere VM template with your variation of the following command:

```
konvoy-image build images/ova/<image.yaml>
```

Any additional configurations can be added to this command using `--overrides` flags as shown below:

- a. Any credential overrides: `--overrides overrides.yaml`
 - b. for FIPS, add this flag: `--overrides overrides/fips.yaml`
 - c. for air-gapped, add this flag: `--overrides overrides/offline-fips.yaml`
5. The [Konvoy Image Builder](#) (see page 1487) (KIB) uses the values in `image.yaml` and the input base OS image to create a vSphere template directly on the vCenter server. This template contains the

required artifacts needed to create a Kubernetes cluster.

When KIB provisions the OS [image](#)⁴⁰⁶ successfully, it creates a manifest file. The `artifact_id` field of this file contains the name of the AMI ID (AWS), template name (vSphere), or image name (GCP/Azure), for example:

```
{
  "name": "vsphere-clone",
  "builder_type": "vsphere-clone",
  "build_time": 1644985039,
  "files": null,
  "artifact_id": "konvoy-ova-vsphere-rhel-84-1.21.6-1644983717",
  "packer_run_uuid": "260e8110-77f8-ca94-e29e-ac7a2ae779c8",
  "custom_data": {
    "build_date": "2022-02-16T03:55:17Z",
    "build_name": "vsphere-rhel-84",
    "build_timestamp": "1644983717",
    [...]
  }
}
```

Recommendation: Now we can now see the template created in our vCenter, it is best to rename it to `dkp-<DKP_VERSION>-k8s-<K8S_VERSION>-<DISTRO>`, like `dkp-2.6.2-k8s-1.26.14-ubuntu` to keep templates organized.

6. Next steps are to deploy a DKP cluster using your vSphere template.

5.13.7.6.1.1 Additional OS YAML files for copy/paste:

RHEL 8.6

```
---
download_images: true
build_name: "rhel-86"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
```

⁴⁰⁶ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

network: ""
resource_pool: ""
template: "base-rhel-8.6" # change default value with your base template name
vsphere_guest_os_type: "rhel8_64Guest"
guest_os_type: "rhel8-64"
# goss params
distribution: "RHEL"
distribution_version: "8.6"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

Ubuntu 20.04

```

---
download_images: true
build_name: "ubuntu-2004"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-
vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "base-ubuntu-20.04" # change default value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "ubuntu2004-64"
  # goss params
  distribution: "ubuntu"
  distribution_version: "20.04"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

Rocky Linux 9.1

```

---
download_images: true
build_name: "rocky-91"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "base-rocky-9.1" # change default value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "rocky9-64"
  # goss params
  distribution: "rocky"
  distribution_version: "9.1"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

5.13.7.6.2 Next Step:

[vSphere FIPS Air-gapped: Create the Management Cluster \(see page 513\)](#)

5.13.7.7 vSphere FIPS Air-gapped: Create the Management Cluster

The instructions on this page will generate a self-managed air-gapped cluster to be used as the Management cluster. A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing.

5.13.7.7.1 Deploying a Cluster in FIPS mode

In order to create a cluster in FIPS mode, we must inform the bootstrap controllers of the appropriate image repository and version tags of the official D2iQ FIPS builds of Kubernetes.

5.13.7.7.1.1 Supported FIPS Builds

Component	Repository	Version
Kubernetes	docker.io/mesosphere ⁴⁰⁷	v1.26.14+fips.0
etcd	docker.io/mesosphere ⁴⁰⁸	3.5.10+fips.0

5.13.7.7.2 Name your Cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the CLUSTER_NAME environment variable with the command:

```
export CLUSTER_NAME=<my-vsphere-cluster>
```



DKP uses the vsphere CSI driver as the [default storage provider](#) (see page 91). Use a [Kubernetes CSI](#)⁴⁰⁹ compatible storage that is suitable for production.

5.13.7.7.2.1 Create a New vSphere Kubernetes Cluster



The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273744255/2.6%2BAttach%2Ba%2BKubernetes%2BCluster> (see page 752).

⁴⁰⁷ <https://hub.docker.com/layers/mesosphere/kube-apiserver/v1.26.14/images/sha256-3c559c75123fa40d588bc2d9050938a0b7dd1dab4e38a4fd24847cd0c0f24694?context=explore>

⁴⁰⁸ https://hub.docker.com/layers/mesosphere/etcd/v3.5.10_fips.0/images/sha256-61df32b8ec1e6229d8bd94caad07331c7e0c0b8e1947af8152c69eb2f04fb771?context=explore

⁴⁰⁹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

Use the following steps to create a new, air-gapped vSphere cluster.

1. Configure your cluster to use an existing local registry as a mirror when attempting to pull images:
IMPORTANT: The image must be created by [Konvoy Image Builder](#) (see page 1487) in order to use the registry mirror feature.

```
export REGISTRY_URL=<https/http>://<registry-address>:<registry-port>
export REGISTRY_CA=<path to the CA on the bastion>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

- `REGISTRY_URL` : the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
- `REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. Konvoy will configure the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
- `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
- `REGISTRY_PASSWORD` : optional if username is not set.

2. Load the image, using either the `docker` or `podman` command:

```
docker load -i konvoy-bootstrap-image-v2.6.2.tar
```

OR

```
podman load -i konvoy-bootstrap-image-v2.6.2.tar
```

3. Create a Kubernetes cluster by copying the following command and substituting the valid values for your environment:

```
dkp create cluster vsphere
--cluster-name ${CLUSTER_NAME} \
--network <NETWORK_NAME> \
--control-plane-endpoint-host <CONTROL_PLANE_IP> \
--data-center <DATACENTER_NAME> \
--data-store <DATASTORE_NAME> \
--folder <FOLDER_NAME> \
--server <VCENTER_API_SERVER_URL> \
--ssh-public-key-file </path/to/key.pub> \
--resource-pool <RESOURCE_POOL_NAME> \
--vm-template konvoy-ova-vsphere-os-release-k8s_release-vsphere-timestamp \
--virtual-ip-interface <ip_interface_name> \
--extra-sans "127.0.0.1" \
```

```

--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
--registry-mirror-password=${REGISTRY_PASSWORD} \
--kubernetes-version=v1.26.14+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--etcd-image-repository=docker.io/mesosphere --etcd-version=3.5.10+fips.0 \
--self-managed

```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

= For bootstrap and custom YAML cluster creation, refer to the [Custom Installation and Additional Infrastructure Tools](#) (see page 1036) section of the documentation for vSphere: [Install vSphere Air-Gapped](#) (see page 1285) or [FIPS 140-2 Compliance](#) (see page 1526) for more FIPS details.

5.13.7.7.2.2 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>⁴¹⁰⁴¹¹.

As they progress, the controllers create Events, which you can also monitor using the command:

```
kubectl get events | grep ${CLUSTER_NAME}
```

For brevity, this example uses `grep`. You can also use separate commands to get Events for specific objects, such as `kubectl get events --field-selector involvedObject.kind="VSphereCluster"` and `kubectl get events --field-selector involvedObject.kind="VSphereMachine"`.

⁴¹⁰ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

⁴¹¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

5.13.7.7.2.3 Next Step:

[vSphere Air-gapped FIPS: Configure Metal LB](#) (see page 517)

5.13.7.8 vSphere Air-gapped FIPS: Configure Metal LB

If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can proceed to [vSphere FIPS Air-gapped: Install Kommander](#) (see page 519) and continue the installation process.

To use MetalLB, create a MetalLB configMap for your vSphere infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)

Select one of the following procedures to create your MetalLB manifest for further editing.

5.13.7.8.1 Layer 2 configuration

Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.



- MetalLB IP address ranges/CIDRs should be within the node's primary network [subnet](#) (see [page 1046](#)).
- MetalLB IP address ranges/CIDRs and node subnet should not conflict with the Kubernetes cluster pod and service subnets.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:



The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
```



```
data:
  config: |
    address-pools:
      - name: default
        protocol: layer2
        addresses:
          - 192.168.1.240-192.168.1.250
EOF
```

When complete, run the following `kubectl` command.


```
kubectl apply -f metallb-conf.yaml
```

5.13.7.8.2 BGP configuration

BGP, like any routing protocol, requires a mesh, so you configure BGP between the nodes. If you are in a Data Center, peer with the BGP routers. The instructions for use with Cloud providers is the same. For a standard configuration featuring one BGP router and one IP address range, you need four pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:

 The following values are generic, enter your specific values into the fields where applicable.

Extract the kubeconfig and deploy a config map for MetalLB using the following command:

```
dkp get kubeconfig -c ${DKP_CLUSTER_NAME} > ${DKP_CLUSTER_NAME}.conf
```

Deploy MetalLB Configuration with the command below:

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
```

```

- peer-address: 10.0.0.1
  peer-asn: 64501
  my-asn: 64500
address-pools:
- name: default
  protocol: bgp
  addresses:
  - 192.168.10.0/24
EOF

```

When complete, run the following `kubectl` command:

```
kubectl apply -f metallb-conf.yaml
```

5.13.7.8.3 Next Step:

[vSphere FIPS Air-gapped: Install Kommander](#) (see page 519)

5.13.7.9 vSphere FIPS Air-gapped: Install Kommander

5.13.7.9.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Ensure you have loaded all necessary images for your configuration. See [Load the Images into Your Registry: Air-gapped Environments](#) (see page 1400).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.13.7.9.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```


2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```


3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init --airgapped > kommander.yaml
```

4. *If required:* Customize your `kommander.yaml`.


 See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, [Rook Ceph customization \(Pre-provisioned envs\)](#) (see page 1405), etc.

5.13.7.9.3 Enable DKP Catalog Applications and Install Kommander in an Air-gapped Environment

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add the following values to enable DKP Catalog Applications:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      path: ./dkp-catalog-applications-v2.6.1.tar.gz
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubecfg=${CLUSTER_NAME}.conf \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-
v2.6.1.tar.gz
```

Tips and recommendations


- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.13.7.9.4 Next Step:

[vSphere FIPS Air-gapped: Verify Install and Log in to the UI](#) (see page 521)

5.13.7.10 vSphere FIPS Air-gapped: Verify Install and Log in to the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.

 **NOTE:** If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
```

```

helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met

```

5.13.7.10.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```

kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'

```

5.13.7.10.2 Log in to the UI with Kommander

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{{ "\n"}}Password: {{.data.password|base64decode}}
{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/
dashboard{{ "\n"}}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.13.7.10.2.1 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

5.13.7.10.2.2 Next Step:


[vSphere FIPS Air-gapped: Create a Managed Cluster Using the DKP CLI](#) (see [page 523](#))

5.13.7.11 vSphere FIPS Air-gapped: Create a Managed Cluster Using the DKP CLI


Enterprise

Gov Advanced

Creating an air-gapped vSphere FIPS Managed cluster with the DKP CLI assumes that you already fulfilled all of the prerequisites and successfully created a vSphere Management cluster. Use this procedure to create a Managed vSphere cluster.

 When creating [Managed clusters](#) (see page 81), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see page 81)!

5.13.7.11.1 Make New Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:
 **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see page 650).

```
kubectl get workspace -A
```


2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.13.7.11.2 Name your Cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the `CLUSTER_NAME` environment variable with the command:

```
export CLUSTER_NAME=<my-managed-vsphere-cluster>
```

 DKP uses local static provisioner as the [default storage provider](#) (see page 91). However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)⁴¹² compatible storage that is suitable for production.

⁴¹² <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

- You can choose from any of the [storage options](#)⁴¹³ available for Kubernetes. To disable the default that Konvoy deploys, set the default StorageClass `localvolume-provisioner` as non-default. Then set your newly created StorageClass to be the default by following the commands in the Kubernetes documentation called [Changing the Default Storage Class](#)⁴¹⁴.

5.13.7.11.3 Create a New vSphere Kubernetes Cluster

ⓘ The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273744255/2.6%2BAttach%2Ba%2BKubernetes%2BCluster> (see page 752).

Use the following steps to create a new, air-gapped vSphere cluster.

- Configure your cluster to use an existing registry as a mirror when attempting to pull images: **IMPORTANT:** The image must be created by [Konvoy Image Builder](#) (see page 1461) so that it uses the registry mirror feature.

```
export REGISTRY_URL=<https/http>://<registry-address>:<registry-port>
export REGISTRY_CA=<path to the CA on the bastion>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

- `REGISTRY_URL` : the address of an existing Docker registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
 - `REGISTRY_CA` : (optional) the path on the bastion machine to the Docker registry CA. Konvoy will configure the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
 - `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
 - `REGISTRY_PASSWORD` : optional if username is not set.
- Create a Kubernetes cluster by copying the following command and substituting the valid values for your environment:

```
dkp create cluster vsphere \
```

⁴¹³ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁴¹⁴ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>


```

--cluster-name ${MANAGED_CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--namespace ${WORKSPACE_NAMESPACE}
--network <NETWORK_NAME> \
--control-plane-endpoint-host <CONTROL_PLANE_IP> \
--data-center <DATACENTER_NAME> \
--data-store <DATASTORE_NAME> \
--folder <FOLDER_NAME> \
--server <VCENTER_API_SERVER_URL> \
--ssh-public-key-file </path/to/key.pub> \e
--resource-pool <RESOURCE_POOL_NAME> \
--vm-template konvoy-ova-vsphere-os-release-k8s_release-vsphere-timestamp \
--virtual-ip-interface <ip_interface_name> \
--extra-sans "127.0.0.1" \
--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
--registry-mirror-password=${REGISTRY_PASSWORD} \
--kubernetes-version=v1.26.14+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--etcd-image-repository=docker.io/mesosphere \
--etcd-version=3.5.10+fips.0 \
--kubeconfig=<management-cluster-kubeconfig-path> \

```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1039\)](#).

5.13.7.11.4 Retrieve the kubeconfig and Explore New vSphere Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:

```

dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf

```

2. List the nodes with the following command:

```

kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes

```

3. List the pods with the following command:

NOTE: Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.13.7.11.4.1 Manually Attach a DKP CLI Cluster to the Management Cluster

⚠ These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

- You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

- This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new attached-cluster-kubeconfig.yaml file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

- Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

- Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

- You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubectl get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

If you have existing clusters or want to create other new clusters to attach, there are many ways to attach a cluster with various requirements and restrictions. To see all the options, visit the section in documentation [Day 2 - Attach an Existing Kubernetes Cluster](#) (see page 752).

At this point, you can create more clusters, perform [other configuration tasks](#) (see page 1036), or proceed to [Day 2 Operations](#) (see page 582).

5.14 VMware Cloud Director Install Options

For VMware Cloud Director (VCD), an installation scenario is provided for you in this location. Remember, there are always more options in the [Additional Infrastructure Configurations](#) (see page 1254) sections, but this will get you operative in the most common scenarios.

Beta

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.14.1 Special Resource Requirements for VMware Cloud Director:

5.14.1.1 CPU and Memory

1. For CPU and Memory, the VCD **Provider** must create the appropriate VM Sizing Policies.
2. The **Provider** must reference these VM Sizing Policies when they create the cluster, using the `--control-plane-sizing-policy` and `--worker-sizing-policy` flags.
3. See <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-73C48A9C-79AF-402C-9746-4C81CAC2B35A.html> regarding parameters like vCPU Speed or CPU Reservation Guarantee that require consideration in VM Sizing Policy. For example, vCPU speed recommendation is a minimum of 3 GHz.

5.14.1.2 Storage

For storage, the VCD **Provider** must follow [Create a Base OS image in vSphere](#) (see page 1260) when they create the KIB base image. In VCD, the base image determines the minimum storage available for the VM.

For more information about VCD parameters that can affect performance, refer to VMware Documentation: <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-73C48A9C-79AF-402C-9746-4C81CAC2B35A.html>

Below are all the supported environment variable combinations: (Refer to this page to check if your OS is a [Supported Operating System](#) (see page 69).)

- [Cloud Director Service Provider](#) (see page 530)

5.14.1.3 Next Step:

[Cloud Director Service Provider](#) (see page 530)

5.14.2 Cloud Director Service Provider

Beta

There is not a “Basic Install” for VMWare Cloud Director since each tenant will have different needs. For Managed Service Providers (MSPs), please refer to [Cloud Director for Service Providers](#) (see page 1323) section of the documentation regarding installation plus image and OVA export and import for your Tenant Organizations.

Before continuing to install DKP on Cloud Director, verify that your VMware vSphere Client environment is running [vSphere Client version](#)⁴¹⁵v6.7.x with Update 3 or later version with [ESXi](#)⁴¹⁶. You must be able to reach the vSphere API endpoint from where the Konvoy command line interface (CLI) runs and have a vSphere account containing Administrator privileges. A RedHat® subscription is required with username and password for downloading DVD ISOs and valid vSphere values for the following: vCenter API server URL Datacenter name Zone name that contains ESXi hosts for your cluster’s nodes.

5.14.2.1 If Required - Further vSphere Prerequisites

- [vSphere: VMware Prerequisites](#) (see page 440)
- <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/162430985/vSphere+Prerequisites+-+Storage+Options> (see page 440)
- <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/162430977/vSphere+Prerequisites+-+Minimum+User+Permissions> (see page 438)

5.14.2.2 Overview

VMware Cloud Director (VCD) provides a private cloud like experience in data centers. VCD allows creation of virtual data centers for a more self-service of administration tasks. An overview of steps has been provided to help.

The overall process for configuring VCD and DKP together includes the following steps:

⁴¹⁵ https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html

⁴¹⁶ <https://docs.vmware.com/en/VMware-vSphere/index.html>

1. Configure vSphere to provide the needed elements, described in the [vSphere Prerequisites: All Installation Types](#) (see page 437).
2. **For air-gapped environments:** Create a [bastion VM host](#) (see page 1285).
3. Create a base OS image (for use in the OVA package containing the disk images packaged with the OVF).
4. Create a CAPI VM image template that uses the base OS image and adds the needed Kubernetes cluster components.
5. Create a Tenant Organization
6. Configure Virtual Data Centers (VDCs)
7. MSP will upload the appropriate OVA to tenant organization's catalog
 - a. MSP uses `konvoy-image build vsphere` to create an OVA
 - b. MSP creates a vApp from the OVA
 - c. Before the MSP creates a cluster in a VCD tenant, it makes this vApp available to that tenant
8. User must be able to reach the VCD API endpoint from where the DKP CLI runs
9. Create a new self-managing DKP cluster on VCD.
10. Install Kommander
11. Verify and log in to the UI

5.14.2.3 Next Step:

For MSPs, please refer to the [VMware Cloud Director Infrastructure](#) (see page 1322) section of the documentation regarding installation and Tenant Organization creation.

5.15 Azure Install Options

For an environment that is Azure, an install is provided for you in this one location.

Remember, there are always more options for custom YAML in the [Custom Installation and Additional Infrastructure Tools](#) (see page 1036) section under [Azure Infrastructure](#) (see page 1156), but this will get you operative in the most common scenario.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

Below are the supported environment variable combinations: (Refer to this page to check if your OS is a [Supported Operating System](#) (see page 69).)

- [Azure Install](#) (see page 532)



Additional Resource Information specific to Azure is below.

- **Control plane nodes** - DKP on Azure defaults to deploying a `Standard_D4s_v3` virtual machine with an 128 GiB volume for the OS and an 80GiB volume for etcd storage, which meets the above [resource requirements \(see page 100\)](#).
- **Worker nodes** - DKP on Azure defaults to deploying a `Standard_D8s_v3` virtual machine with an 80 GiB volume for the OS, which meets the above [resource requirements \(see page 100\)](#).

5.15.1 Azure Install

For an environment that is on the Azure Infrastructure, install options are provided for you in this one location. Remember, there are always more options in the [Azure Infrastructure \(see page 1156\)](#) section under [Custom Installation and Additional Infrastructure Tools \(see page 1036\)](#), but this will get you up and running in the most common scenario.

If not already done, refer to [Get Started \(see page 79\)](#) section of the documentation for:

- [Resource Requirements \(see page 100\)](#)
- [Install Overview \(see page 108\)](#)
- [Prerequisites for Install \(see page 121\)](#)

5.15.1.1 Azure Prerequisites

1. Sign in to Azure:

```
az login
```

```
[
  {
    "cloudName": "AzureCloud",
    "homeTenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "id": "b1234567-abcd-11a1-a0a0-1234a5678b90",
    "isDefault": true,
    "managedByTenants": [],
    "name": "D2iQ Developer Subscription",
    "state": "Enabled",
    "tenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "user": {
      "name": "user@azuremesosphere.onmicrosoft.com",
      "type": "user"
    }
  }
]
```

```
]

```

2. Create an Azure Service Principal (SP) by running the following commands:
 - a. If you have more than one Azure account, run this command to identify your account:

```
echo $(az account show --query id -o tsv)
```

- b. Run this command to ensure you are pointing to the correct Azure subscription ID:

```
az account set --subscription "D2iQ Developer Subscription"
```

- c. If an SP with the name exists, this command rotates the password.

```
az ad sp create-for-rbac --role contributor --name "$(whoami)-konvoy" --scopes=/subscriptions/$(az account show --query id -o tsv) --query "{ client_id: appId, client_secret: password, tenant_id: tenant }"
```

```
{
  "client_id": "7654321a-1a23-567b-b789-0987b6543a21",
  "client_secret": "Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C",
  "tenant_id": "a1234567-b132-1234-1a11-1234a5678b90"
}
```

3. Set the `AZURE_CLIENT_SECRET` environment variable:

```
export AZURE_CLIENT_SECRET="<azure_client_secret>" #
Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C
export AZURE_CLIENT_ID="<client_id>" # 7654321a-1a23-567b-
b789-0987b6543a21
export AZURE_TENANT_ID="<tenant_id>" # a1234567-b132-1234-1a11-12
34a5678b90
export AZURE_SUBSCRIPTION_ID="<subscription_id>" # b1234567-abcd-11a1-
a0a0-1234a5678b90
```


4. Ensure you have an [override file](#) (see page 1513) to configure specific attributes of your Azure image.

5.15.1.2 Next Step:

[Azure: Create Image](#) (see page 534)

5.15.1.3 Azure: Create Image

This procedure describes how to use the [Konvoy Image Builder \(see page 1461\)](#) (KIB) to create a [Cluster API](#)⁴¹⁷ compliant Azure Virtual Machine (VM) Image. The VM Image contains the base operating system you specify and all the necessary Kubernetes components. The Konvoy Image Builder uses variable `overrides` to specify the base image and container images to use in your new Azure VM image.

 The default Azure image is not recommended for use in production. We suggest using KIB for Azure to build the image in order to take advantage of enhanced cluster operations. To explore more information on this topic refer to the [Azure Infrastructure \(see page 1156\)](#).


5.15.1.3.1 Prerequisites

Before you begin, you must:

- Download the [Konvoy Image Builder \(see page 1461\)](#) bundle for your version of DKP.
- Check the [Supported Infrastructure Operating Systems \(see page 69\)](#)
- Check the [Supported Kubernetes Version \(see page 1552\)](#) for your Provider.
- Create a working `Docker` setup.

5.15.1.3.2 Extract the KIB Bundle

Extract the bundle and `cd` into the extracted `konvoy-image-bundle-$VERSION_$(OS)` folder. The bundled version of `konvoy-image` contains an embedded `docker` image that contains all the requirements for building.

 The `konvoy-image` binary and all supporting folders are also extracted. When extracted, `konvoy-image` bind mounts the current working directory (`$(PWD)`) into the container to be used.

5.15.1.3.3 Build the Image

Run the `konvoy-image` command to build and validate the image.

⁴¹⁷ <https://cluster-api.sigs.k8s.io/>

```
konvoy-image build azure --client-id ${AZURE_CLIENT_ID} --tenant-id $
{AZURE_TENANT_ID} --overrides override-source-image.yaml images/azure/ubuntu-2004.yam
l
```

By default, the image builder builds in the `westus2` location. To specify another location set the `--location` flag (shown in example below is how to change the location to `eastus`):

```
konvoy-image build azure --client-id ${AZURE_CLIENT_ID} --tenant-id $
{AZURE_TENANT_ID} --location eastus --overrides override-source-image.yaml images/
azure/centos-7.yaml
```

When the command is complete, the image id is printed and written to the `./manifest.json` file. This file has an `artifact_id` field whose value provides the name of the image. You should then specify this image id when creating the cluster.

5.15.1.3.4 Image Gallery

By default Konvoy Image Builder will create a Resource Group, Gallery, and Image Name to store the resulting image in. To specify a specific Resource Group, Gallery, or Image Name flags may be specified:

```
--gallery-image-locations string      a list of locations to publish the image
(default same as location)
--gallery-image-name string           the gallery image name to publish the image to
--gallery-image-offer string          the gallery image offer to set (default "dkp")
--gallery-image-publisher string      the gallery image publisher to set (default
"dkp")
--gallery-image-sku string            the gallery image sku to set
--gallery-name string                 the gallery name to publish the image in
(default "dkp")
--resource-group string               the resource group to create the image in
(default "dkp")
```

When creating your cluster, you will then add this flag during the create process for your custom image: `--compute-gallery-id "<Managed Image Shared Image Gallery Id>"`. The SKU and Image Name will default to the values found in the image YAML.

5.15.1.3.4.1 More customization options:

- See [Create a New Azure Cluster \(see page 1163\)](#) for specific consumption of image commands.
- See [KIB for Azure \(see page 1477\)](#) for more flags and Rocky Linux Marketplace.

5.15.1.3.5 Next Step:

[Azure: Create the Management Cluster \(see page 536\)](#)

5.15.1.4 Azure: Create the Management Cluster

5.15.1.4.1 Name your cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=<azure-example>
```

5.15.1.4.2 Encode your Azure Credential Variables:

Base64 encode the Azure environment variables set in the [Azure install prerequisites](#) (see page 535) step:

```
export AZURE_SUBSCRIPTION_ID_B64="$(echo -n "${AZURE_SUBSCRIPTION_ID}" | base64 | tr -d '\n')"
```

```
export AZURE_TENANT_ID_B64="$(echo -n "${AZURE_TENANT_ID}" | base64 | tr -d '\n')"
```

```
export AZURE_CLIENT_ID_B64="$(echo -n "${AZURE_CLIENT_ID}" | base64 | tr -d '\n')"
```

```
export AZURE_CLIENT_SECRET_B64="$(echo -n "${AZURE_CLIENT_SECRET}" | base64 | tr -d '\n')"
```



DKP uses the Azure CSI driver as the [default storage provider](#) (see page 91). Use a [Kubernetes CSI](#)⁴¹⁸ compatible storage that is suitable for production.

5.15.1.4.3 Create an Azure Kubernetes Cluster

If you use these instructions to create an Azure cluster using the DKP default settings, your cluster is deployed with 3 control plane nodes and 4 worker nodes.

Availability zones (AZs) are isolated locations within data center regions from which public cloud services originate and operate. Because all the nodes in a node pool are deployed in a single Availability Zone, you may wish to create additional node pools to ensure your cluster has nodes deployed in multiple Availability Zones.

⁴¹⁸ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

- By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

1. Generate the Kubernetes cluster objects. The following example shows a common configuration. See [dkp create cluster azure](#) (see page 1662) reference for the full list of cluster creation options.

```
dkp create cluster azure \
--cluster-name=${CLUSTER_NAME} \
--self-managed
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#)⁴¹⁹.

Output is similar to below:

```
Generating cluster resources
```

A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing. As part of the underlying processing using the `--self-managed` flag, the DKP CLI:

- creates a bootstrap cluster
- creates a workload cluster
- moves CAPI controllers from the bootstrap cluster to the workload cluster, making it self-managed
- deletes the bootstrap cluster

To understand how this process works step by step, you can find a customizable [Create a Custom Azure Cluster](#)⁴²⁰ under [Custom Installation and Additional Infrastructure Tools](#) (see page 1036).

5.15.1.4.4 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

⁴¹⁹ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273680223&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configuring+an+HTTP%2FHTTPS+Pr)

[oxy](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273680223&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configuring+an+HTTP%2FHTTPS+Pr)

⁴²⁰ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273680223&linkCreation=true&spaceKey=DENT&title=%282.6%29+Create+a+New+Custom+Azure+Clus)

[ter](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273680223&linkCreation=true&spaceKey=DENT&title=%282.6%29+Create+a+New+Custom+Azure+Clus)

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>⁴²¹⁴²².

5.15.1.4.5 Known Limitations

The Konvoy version used to create a bootstrap cluster must match the Konvoy version used to create a workload cluster.

- Konvoy supports deploying one workload cluster.
- Konvoy generates a set of objects for one Node Pool.
- Konvoy does not validate edits to cluster objects.

5.15.1.4.6 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Azure%3A%20Install%20Kommander&linkCreation=true&fromPageId=273680223>⁴²³

5.15.1.5 Azure: Install Kommander

You have installed the Konvoy component and created a cluster. Now it is time to Install Kommander which will allow you to access the UI and attach new or existing clusters to monitor.

5.15.1.5.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.15.1.5.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

⁴²¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

⁴²² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

⁴²³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273680223&linkCreation=true&spaceKey=DENT&title=%282.6%29+Azure%3A+Install+Kommander>


2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```


3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init > kommander.yaml
```

4. *If required:* Customize your `kommander.yaml`.


 See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, [Rook Ceph customization](#) (Pre-provisioned envs (see page 1405)), etc.

5.15.1.5.3 Enable DKP Catalog Applications and Install Kommander

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
        ref:
          tag: v2.6.1
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

Tips and recommendations


- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File \(see page 87\)](#).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.15.1.5.4 Next Step:

[Azure: Verify Install and Log in to the UI \(see page 540\)](#)

5.15.1.6 Azure: Verify Install and Log in to the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.

 **NOTE:** If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

5.15.1.6.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

By default, you can log in to the UI in Kommander with the credentials given using this command:


```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{\n"}Password: {{.data.password|base64decode}}
{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/
dashboard{{ "\n"}}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbJUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.15.1.6.2 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

Now you can [create additional new clusters](#) (see [page 543](#)), or proceed to Day 2 Operations in the link.

5.15.1.6.3 Next Step:

[Azure: Subsequent New Clusters](#) (see [page 543](#))

5.15.1.7 Azure: Create a Managed Cluster Using the DKP CLI

Enterprise

Gov Advanced

In the previous step, the new cluster was created as Self-managed which allows it to be a Management cluster. Subsequent new clusters are not self-managed as they will likely be Managed or Attached clusters to this Management Cluster.



When creating [Managed clusters](#) (see [page 81](#)), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see [page 81](#))!

Your new, managed cluster needs to be part of either a new Workspace or an existing Workspace under a management cluster. Create or locate a workspace name to get started with your managed cluster.

5.15.1.7.1 Make New Managed Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:

⚠ NOTE: If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see [page 650](#)).

```
kubectl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

5.15.1.7.2 Name your Cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export MANAGED_CLUSTER_NAME=<azure-additional>
```

5.15.1.7.3 Use DKP CLI

ⓘ The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above.

If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29920384/Attach+a+Kubernetes+Cluster>. (see page 752).

Execute this command to create an additional cluster without the `self-managed` flag:

```
dkp create cluster azure --cluster-name=${MANAGED_CLUSTER_NAME} \
--namespace=${WORKSPACE_NAMESPACE} \
--additional-tags=owner=$(whoami) \
--kubeconfig=<management-cluster-kubeconfig-path> \
```

ⓘ If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

5.15.1.7.4 Retrieve the kubeconfig and Explore New Azure Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:

NOTE: Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.15.1.7.4.1 Manually Attach a DKP CLI Cluster to the Management Cluster

⚠ These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

- You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

- This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new attached-cluster-kubeconfig.yaml file:

```
apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>
```

- Create this secret in the desired workspace:

```
kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}
```

- Create this `kommandercluster` object to attach the cluster to the workspace:

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF
```

- You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```
kubectl get kommanderclusters -A
```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

You can also [Create a new Azure Cluster Using the DKP UI](#) (see page 746).

5.15.1.7.4.2 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.16 AKS Install Options

Enterprise

Gov Advanced

For an environment that is AKS on the Azure Infrastructure, install options are provided for you in this one location. Remember, there are always more options in the [AKS Infrastructure](#) (see page 1189) section under [Custom Installation and Additional Infrastructure Tools](#) (see page 1036), but this will get you operative in the most common scenario.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

Below are the supported environment variable combinations: (Refer to this page to check if your OS is a [Supported Operating System](#) (see page 69).)

- [AKS Install](#) (see page 547)

i In order to install Kommander, you need to have CAPI components, cert-manager, etc on a self-managed cluster. The CAPI components mean you can control the lifecycle of the cluster, and other clusters. However, because AKS is semi-managed by Azure, the AKS clusters are under Azure control and don't have those components. Therefore, Kommander will not be able to be installed and these clusters will be attached to the management cluster.

5.16.1 AKS Install

Enterprise

Gov Advanced

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:


- [Resource Requirements](#) (see page 100)

- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

5.16.1.1 DKP Prerequisites

Before starting the DKP installation, verify that you have:

- A Management cluster with DKP and the Kommander component installed:

 An AKS cluster cannot be a [Management](#) (see page 82) or [Essential](#) (see page 83) cluster. To install DKP on your AKS cluster, first ensure you have a Management cluster with DKP and the Kommander component installed, that handles the lifecycle of your AKS cluster.

- An x86_64-based Linux or macOS machine with a supported version of the [operating system](#) (see page 69).
- A [Self-managed Azure cluster](#) (see page 1177), if you used the Day 1-Basic Install for Azure instructions, your cluster was created using `--self-managed` flag and therefore is already a self-managed cluster.
- Download the `dkp` [binary](#) (see page 78) for Linux, or macOS. To check which version of DKP you installed for compatibility reasons, run the `dkp version -h` command ([dkp version](#) (see page 1749)).
- [Docker](#)⁴²⁴ version 18.09.2 or later.
- [kubectl](#)⁴²⁵ for interacting with the running cluster.
- The [Azure CLI](#)⁴²⁶.
- A valid Azure account [used to sign in to the Azure CLI](#)⁴²⁷.
- All [Resource requirements](#) (see page 0)

5.16.1.2 AKS Prerequisites

Follow these steps:

1. Log in to Azure:

```
az login
```

⁴²⁴ <https://docs.docker.com/get-docker/>

⁴²⁵ <https://kubernetes.io/docs/tasks/tools/#kubectl>

⁴²⁶ <https://docs.microsoft.com/en-us/cli/azure/install-azure-cli>

⁴²⁷ <https://docs.microsoft.com/en-us/cli/azure/authenticate-azure-cli?view=azure-cli-latest>

```
[
  {
    "cloudName": "AzureCloud",
    "homeTenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "id": "b1234567-abcd-11a1-a0a0-1234a5678b90",
    "isDefault": true,
    "managedByTenants": [],
    "name": "Mesosphere Developer Subscription",
    "state": "Enabled",
    "tenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "user": {
      "name": "user@azuresphere.onmicrosoft.com",
      "type": "user"
    }
  }
]
```

2. Create an Azure Service Principal (SP) by running the following command:

i NOTE: If an SP with the name exists, this command will rotate the password.

```
az ad sp create--for-rbac --role contributor --name "$(whoami)-konvoy" --
scopes=/subscriptions/$(az account show --query id -o tsv)
```

```
{
  "appId": "7654321a-1a23-567b-b789-0987b6543a21",
  "displayName": "azure-cli-2021-03-09-23-17-06",
  "password": "Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C",
  "tenant": "a1234567-b132-1234-1a11-1234a5678b90"
}
```

3. Set the required environment variables:

```
export AZURE_SUBSCRIPTION_ID="<id>"           # b1234567-abcd-11a1-
a0a0-1234a5678b90
export AZURE_TENANT_ID="<tenant>"             # a1234567-b132-1234-1a11-1234a5678b9
0
export AZURE_CLIENT_ID="<appId>"             # 7654321a-1a23-567b-
b789-0987b6543a21
export AZURE_CLIENT_SECRET="<password>"      # Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C
```

4. Base64 encode the same environment variables:


```
export AZURE_SUBSCRIPTION_ID_B64="$(echo -n "${AZURE_SUBSCRIPTION_ID}" | base64
| tr -d '\n')"
export AZURE_TENANT_ID_B64="$(echo -n "${AZURE_TENANT_ID}" | base64 | tr -d
'\n')"
export AZURE_CLIENT_ID_B64="$(echo -n "${AZURE_CLIENT_ID}" | base64 | tr -d
'\n')"
export AZURE_CLIENT_SECRET_B64="$(echo -n "${AZURE_CLIENT_SECRET}" | base64 |
tr -d '\n')"
```

5. Check to see what version of Kubernetes is available in your region. When deploying with AKS, you must pick a version of Kubernetes that is available in AKS and use that version for subsequent steps. To find out the list of available Kubernetes versions in the Azure Region you are using, run the following command, substituting `<your-location>` for the Azure region you're deploying to:

```
az aks get-versions -o table --location <your-location>
```

The output from this command resembles the following:

```
az aks get-versions -o table --location westus
KubernetesVersion    Upgrades
-----
1.27.3(preview)      None available
1.27.1(preview)      1.27.3(preview)
1.26.6               1.27.1(preview), 1.27.3(preview)
1.26.3               1.26.6, 1.27.1(preview), 1.27.3(preview)
1.25.11              1.26.3, 1.26.6
1.25.6               1.25.11, 1.26.3, 1.26.6
1.24.15              1.25.6, 1.25.11
1.24.10              1.24.15, 1.25.6, 1.25.11
```

6. Choose a version of Kubernetes to install from the list of `KubernetesVersion`, choosing a compatible version as documented in the [Supported Kubernetes Versions](#) (see page 1552) for this version of DKP. The version listed in the command is an example:

```
export KUBERNETES_VERSION=1.26.6
```

5.16.1.3 Next Step:

[AKS: Create Image](#) (see page 550)

5.16.1.4 AKS: Create Image

Enterprise

Gov Advanced

AKS best practices discourage building custom images. If the image is customized, it breaks some of the autoscaling and security capabilities of AKS. Since custom virtual machine images are discouraged in AKS, Konvoy Image Builder (KIB) does not include any support for building custom machine images for AKS.

5.16.1.4.1 Next Step:

[AKS: Cluster Creation](#) (see page 551)

5.16.1.5 AKS: Create an AKS Cluster

Enterprise

Gov Advanced

When creating a [Managed](#) (see page 82) cluster on AKS infrastructure, you can choose from multiple configuration types. Create a new Kubernetes cluster in a non-air-gapped environment with the steps below.

5.16.1.5.1 Use DKP to create a new AKS cluster

Ensure that the `KUBECONFIG` environment variable is set to the Management cluster by running `export KUBECONFIG=<Management_cluster_kubeconfig>.conf`.

5.16.1.5.2 Name Your Cluster

Give your cluster a unique name suitable for your environment.

The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)⁴²⁸ for more naming information.

5.16.1.5.3 Create a New AKS Kubernetes Cluster

1. Set the environment variable to a name for this cluster.

⁴²⁸ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

```
export CLUSTER_NAME=<aks-example>
```

2. Check to see what version of Kubernetes is available in your region. When deploying with AKS, you need to declare the version of Kubernetes you wish to use by running the following command, substituting `<your-location>` for the Azure region you're deploying to:

```
az aks get-versions -o table --location <your-location>
```

3. Set the version of Kubernetes you've chosen:

NOTE: Refer to the current release Kubernetes [compatibility table](#) (see page 117) for correct version to use and choose an available 1.26.x version. The version listed in the command is an example.

```
export KUBERNETES_VERSION=1.26.6
```

4. Create the cluster:

```
dkp create cluster aks --cluster-name=${CLUSTER_NAME} --additional-tags=owner=$(whoami) --kubernetes-version=${KUBERNETES_VERSION}
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

```
Generating cluster resources
cluster.cluster.x-k8s.io/aks-example created
azuremanagedcontrolplane.infrastructure.cluster.x-k8s.io/aks-example created
azuremanagedcluster.infrastructure.cluster.x-k8s.io/aks-example created
machinepool.cluster.x-k8s.io/aks-example created
azuremanagedmachinepool.infrastructure.cluster.x-k8s.io/cp6dsz8 created
machinepool.cluster.x-k8s.io/aks-example-md-0 created
azuremanagedmachinepool.infrastructure.cluster.x-k8s.io/mp6gglj created
clusterresourceset.addons.cluster.x-k8s.io/cluster-autoscaler-aks-example
created
configmap/cluster-autoscaler-aks-example created
clusterresourceset.addons.cluster.x-k8s.io/node-feature-discovery-aks-example
created
configmap/node-feature-discovery-aks-example created
clusterresourceset.addons.cluster.x-k8s.io/nvidia-feature-discovery-aks-example
created
configmap/nvidia-feature-discovery-aks-example created
```

5.16.1.5.4 Inspecting or Editing the Cluster Objects

Use your favorite editor.



NOTE: Editing the cluster objects requires some understanding of Cluster API. Edits can prevent the cluster from deploying successfully.

The objects are [Custom Resources](#)⁴²⁹ defined by Cluster API components, and they belong in three different categories:

- Cluster
A *Cluster* object has references to the infrastructure-specific and control plane objects.
- Control Plane
- Node Pool
A Node Pool is a collection of machines with identical properties. For example, a cluster might have one Node Pool with large memory capacity, another Node Pool with GPU support. Each Node Pool is described by three objects: The MachinePool references an object that describes the configuration of Kubernetes components (for example, kubelet) deployed on each node pool machine, and an infrastructure-specific object that describes the properties of all node pool machines. Here, it references a *KubeadmConfigTemplate*.

For in-depth documentation about the objects, read [Concepts](#)⁴³⁰ in the Cluster API Book.

1. Wait for the cluster control-plane to be ready:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

```
cluster.cluster.x-k8s.io/aks-example condition met
```

The `READY` status will become `True` after the cluster control-plane becomes ready.

2. After the objects are created on the API server, the Cluster API controllers reconcile them. They create infrastructure and machines. As they progress, they update the Status of each object. DKP provides a command to describe the current status of the cluster:

```
dkp describe cluster -c ${CLUSTER_NAME}
```

⁴²⁹ <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

⁴³⁰ <https://cluster-api.sigs.k8s.io/user/concepts.html>

```

NAME                                                                 READY  SEVERITY
REASON  SINCE  MESSAGE
Cluster/aks-example                                                                 True
48m
├─ClusterInfrastructure - AzureManagedCluster/aks-example
└─ControlPlane - AzureManagedControlPlane/aks-example

```

3. As they progress, the controllers also create Events. List the Events using this command:

```
kubectl get events | grep ${CLUSTER_NAME}
```

For brevity, the example uses `grep`. It is also possible to use separate commands to get Events for specific objects. For example, `kubectl get events --field-selector involvedObject.kind="AKSCluster"` and `kubectl get events --field-selector involvedObject.kind="AKSMachine"`.

```

48m          Normal    SuccessfulSetNodeRefs          machinepool/aks-
example-md-0          [{Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000000 UID:e3c30389-660d-46f5-b9d7-219f80b5674d APIVersion:
ResourceVersion: FieldPath:} {Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000001 UID:300d71a0-f3a7-4c29-9ff1-1995ffb9cfd3 APIVersion:
ResourceVersion: FieldPath:} {Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000002 UID:8eae2b39-a415-425d-8417-d915a0b2fa52 APIVersion:
ResourceVersion: FieldPath:} {Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000003 UID:3e860b88-f1a4-44d1-b674-a54fad599a9d APIVersion:
ResourceVersion: FieldPath:}]
6m4s          Normal    AzureManagedControlPlane available
azuremanagedcontrolplane/aks-example          successfully reconciled
48m          Normal    SuccessfulSetNodeRefs          machinepool/aks-
example          [{Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000000 UID:e3c30389-660d-46f5-b9d7-219f80b5674d APIVersion:
ResourceVersion: FieldPath:} {Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000001 UID:300d71a0-f3a7-4c29-9ff1-1995ffb9cfd3 APIVersion:
ResourceVersion: FieldPath:} {Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000002 UID:8eae2b39-a415-425d-8417-d915a0b2fa52 APIVersion:
ResourceVersion: FieldPath:}]

```

5.16.1.5.5 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>⁴³¹⁴³².

More information about AKS can be found in the [AKS Infrastructure](#) (see page 1189) under [Custom Installation and Additional Infrastructure Tools](#) (see page 1036).

5.16.1.5.6 Next Steps:

[AKS: Retrieve kubeconfig for AKS Cluster](#) (see page 555)

5.16.1.6 AKS: Retrieve kubeconfig for AKS Cluster

Enterprise

Gov Advanced

5.16.1.6.1 Learn to interact with your AKS Kubernetes cluster

This guide explains how to use the command line to interact with your newly deployed Kubernetes cluster. Before you start, make sure you have created a workload cluster, as described in [AKS: Create an AKS Cluster](#) (see page 554).

5.16.1.6.2 Explore the New AKS Cluster

1. Get a kubeconfig file for the workload cluster:

When the workload cluster is created, the cluster lifecycle services generate a kubeconfig file for the workload cluster, and write it to a *Secret*. The kubeconfig file is scoped to the cluster administrator.

Get the kubeconfig from the *Secret*, and write it to a file, using this command:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

2. List the Nodes using this command:

⁴³¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

⁴³² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
aks-cp6dsz8-41174201-vmss000000	Ready	agent	56m	v1.26.6
aks-cp6dsz8-41174201-vmss000001	Ready	agent	55m	v1.26.6
aks-cp6dsz8-41174201-vmss000002	Ready	agent	56m	v1.26.6
aks-mp6gglj-41174201-vmss000000	Ready	agent	55m	v1.26.6
aks-mp6gglj-41174201-vmss000001	Ready	agent	55m	v1.26.6
aks-mp6gglj-41174201-vmss000002	Ready	agent	55m	v1.26.6
aks-mp6gglj-41174201-vmss000003	Ready	agent	56m	v1.26.6



NOTE: It may take a few minutes for the Status to move to `Ready` while the Pod network is deployed. The Nodes' Status should change to Ready soon after the `calico-node` DaemonSet Pods are Ready.

3. List the Pods using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get --all-namespaces pods
```

NAMESPACE	STATUS	RESTARTS	NAME	AGE	READY
calico-system	Running	0	calico-kube-controllers-5dcd4b47b5-tgslm	3m58s	1/1
calico-system	Running	0	calico-node-46dj9	3m58s	1/1
calico-system	Running	0	calico-node-crdgc	3m58s	1/1
calico-system	Running	0	calico-node-m7s7x	3m58s	1/1
calico-system	Running	0	calico-node-qfkqc	3m57s	1/1
calico-system	Running	0	calico-node-sfqfm	3m57s	1/1
calico-system	Running	0	calico-node-sn67x	3m53s	1/1
calico-system	Running	0	calico-node-w2pvt	3m58s	1/1

calico-system	calico-typha-6f7f59969c-5z4t5	1/1
Running 0	3m51s	
calico-system	calico-typha-6f7f59969c-ddzqb	1/1
Running 0	3m58s	
calico-system	calico-typha-6f7f59969c-rr4lj	1/1
Running 0	3m51s	
kube-system	azure-ip-masq-agent-4f4v6	1/1
Running 0	4m11s	
kube-system	azure-ip-masq-agent-5xfh2	1/1
Running 0	4m11s	
kube-system	azure-ip-masq-agent-9hlk8	1/1
Running 0	4m8s	
kube-system	azure-ip-masq-agent-9vsgg	1/1
Running 0	4m16s	
kube-system	azure-ip-masq-agent-b9wjj	1/1
Running 0	3m57s	
kube-system	azure-ip-masq-agent-kpjtl	1/1
Running 0	3m53s	
kube-system	azure-ip-masq-agent-vr7hd	1/1
Running 0	3m57s	
kube-system	cluster-autoscaler-b4789f4bf-qkfk2	0/1
Init:0/1 0	3m28s	
kube-system	coredns-845757d86-9jf8b	1/1
Running 0	5m29s	
kube-system	coredns-845757d86-h4xfs	1/1
Running 0	4m	
kube-system	coredns-autoscaler-5f85dc856b-xjb5z	1/1
Running 0	5m23s	
kube-system	csi-azuredisk-node-4n4fx	3/3
Running 0	3m53s	
kube-system	csi-azuredisk-node-8pnjj	3/3
Running 0	3m57s	
kube-system	csi-azuredisk-node-sbt6r	3/3
Running 0	3m57s	
kube-system	csi-azuredisk-node-v25wc	3/3
Running 0	4m16s	
kube-system	csi-azuredisk-node-vfbxg	3/3
Running 0	4m11s	
kube-system	csi-azuredisk-node-w5ff5	3/3
Running 0	4m11s	
kube-system	csi-azuredisk-node-zzgqx	3/3
Running 0	4m8s	
kube-system	csi-azurefile-node-2rpcc	3/3
Running 0	3m57s	
kube-system	csi-azurefile-node-4gqkf	3/3
Running 0	4m11s	
kube-system	csi-azurefile-node-f6k8m	3/3
Running 0	4m16s	
kube-system	csi-azurefile-node-k72xq	3/3
Running 0	4m8s	
kube-system	csi-azurefile-node-vx7r4	3/3
Running 0	3m53s	

kube-system	csi-azurefile-node-zc8kr	3/3
Running 0	4m11s	
kube-system	csi-azurefile-node-zkl6b	3/3
Running 0	3m57s	
kube-system	kube-proxy-4fpb6	1/1
Running 0	3m53s	
kube-system	kube-proxy-6qfbf	1/1
Running 0	4m16s	
kube-system	kube-proxy-6wnt2	1/1
Running 0	4m8s	
kube-system	kube-proxy-cspd5	1/1
Running 0	3m57s	
kube-system	kube-proxy-nsgq6	1/1
Running 0	4m11s	
kube-system	kube-proxy-qz2st	1/1
Running 0	4m11s	
kube-system	kube-proxy-zvh9k	1/1
Running 0	3m57s	
kube-system	metrics-server-6bc97b47f7-ltkkj	1/1
Running 0	5m28s	
kube-system	tunnelfront-77d68f78bf-t78ck	1/1
Running 0	5m23s	
node-feature-discovery	node-feature-discovery-master-65dc499cd-fxwb5	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-277xc	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-4dq5k	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-57nb8	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-b4lkl	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-kslst	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-ppjtm	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-x5bgf	1/1
Running 0	3m28s	
tigera-operator	tigera-operator-74c4d9cf84-k7css	1/1
Running 0	5m25s	

5.16.1.6.3 Next Step:

[AKS: Attach Cluster \(see page 558\)](#)

5.16.1.7 AKS: Attach Cluster

ENTERPRISE


5.16.1.7.0.1 Attach an existing AKS cluster

You can attach existing Kubernetes clusters to the Management Cluster. After attaching the cluster, you can use the UI to [examine and manage](#) (see page 745) this cluster. The following procedure shows how to attach an existing Azure Kubernetes Service (AKS) cluster.

5.16.1.7.1 Before you Begin

This procedure requires the following items and configurations:

- A fully configured and running Azure [AKS](#)⁴³³ cluster with administrative privileges.
- The current version DKP Enterprise is [installed](#) (see page 147) on your cluster.
- Ensure you have installed `kubectl` in your Management cluster.

 This procedure assumes you have an existing and spun up Azure AKS cluster(s) with administrative privileges. Refer to the Azure [AKS](#)⁴³⁴ for setup and configuration information.

5.16.1.7.2 Attach AKS Clusters

Ensure that the `KUBECONFIG` environment variable is set to the Management cluster before attaching by running:

```
export KUBECONFIG=<Management_cluster_kubeconfig>.conf
```

5.16.1.7.2.1 Ensure you have access to your AKS clusters

1. Ensure you are connected to your AKS clusters. Enter the following commands for each of your clusters:

```
kubectl config get-contexts
kubectl config use-context <context for first AKS cluster>
```

2. Confirm `kubectl` can access the AKS cluster:

```
kubectl get nodes
```

⁴³³ <https://azure.microsoft.com/en-us/products/kubernetes-service/>

⁴³⁴ <https://azure.microsoft.com/en-us/products/kubernetes-service/>

5.16.1.7.2.2 Create a kubeconfig file for your AKS cluster

To get started, ensure you have [kubectl](#)⁴³⁵ set up and configured with [ClusterAdmin](#)⁴³⁶ for the cluster you want to connect to Kommander.

1. Create the necessary service account:

```
kubectl -n kube-system create serviceaccount kommander-cluster-admin
```

2. Create a token secret for the `serviceaccount`:

```
kubectl -n kube-system create -f - <<EOF
apiVersion: v1
kind: Secret
metadata:
  name: kommander-cluster-admin-sa-token
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
type: kubernetes.io/service-account-token
EOF
```

3. Verify that the `serviceaccount` token is ready by running this command:

```
kubectl -n kube-system get secret kommander-cluster-admin-sa-token -oyaml
```

Verify that the `data.token` field is populated. The output should be similar to this:

```
apiVersion: v1
data:
  ca.crt: LS0tLS1CRUdJTiBDR...
  namespace: ZGVmYXVsdA==
  token: ZXlKaGJHY2lPaUpTVX...
kind: Secret
metadata:
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
    kubernetes.io/service-account.uid: b62bc32e-b502-4654-921d-94a742e273a8
  creationTimestamp: "2022-08-19T13:36:42Z"
  name: kommander-cluster-admin-sa-token
  namespace: default
  resourceVersion: "8554"
  uid: 72c2a4f0-636d-4a70-9f1c-55a75f15e520
```

⁴³⁵ <https://kubernetes.io/docs/tasks/tools/#kubectl>

⁴³⁶ <https://kubernetes.io/docs/concepts/cluster-administration/>

```
type: kubernetes.io/service-account-token
```

4. Configure the new service account for `cluster-admin` permissions:

```
cat << EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: kommander-cluster-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- kind: ServiceAccount
  name: kommander-cluster-admin
  namespace: kube-system
EOF
```

5. Set up the following environment variables with the access data that is needed for producing a new kubeconfig file:

```
export USER_TOKEN_VALUE=$(kubectl -n kube-system get secret/kommander-cluster-admin-sa-token -o=go-template='{{.data.token}}' | base64 --decode)
export CURRENT_CONTEXT=$(kubectl config current-context)
export CURRENT_CLUSTER=$(kubectl config view --raw -o=go-template='{{range .contexts}}{{if eq .name "'${CURRENT_CONTEXT}'"}}{{index .context "cluster" }}{{end}}{{end}}')
export CLUSTER_CA=$(kubectl config view --raw -o=go-template='{{range .clusters}}{{if eq .name "'${CURRENT_CLUSTER}'"}}"{{with index .cluster "certificate-authority-data" }}{{.}}{{end}}"{{ end }}{{ end }}')
export CLUSTER_SERVER=$(kubectl config view --raw -o=go-template='{{range .clusters}}{{if eq .name "'${CURRENT_CLUSTER}'"}}{{.cluster.server }}{{end}}{{ end }}')
```

6. Confirm these variables have been set correctly:

```
export -p | grep -E 'USER_TOKEN_VALUE|CURRENT_CONTEXT|CURRENT_CLUSTER|CLUSTER_CA|CLUSTER_SERVER'
```

7. Generate a kubeconfig file that uses the environment variable values from the previous step:

```
cat << EOF > kommander-cluster-admin-config
apiVersion: v1
kind: Config
current-context: ${CURRENT_CONTEXT}
```

```

contexts:
- name: ${CURRENT_CONTEXT}
  context:
    cluster: ${CURRENT_CONTEXT}
    user: kommander-cluster-admin
    namespace: kube-system
clusters:
- name: ${CURRENT_CONTEXT}
  cluster:
    certificate-authority-data: ${CLUSTER_CA}
    server: ${CLUSTER_SERVER}
users:
- name: kommander-cluster-admin
  user:
    token: ${USER_TOKEN_VALUE}
EOF

```

- This process produces a file in your current working directory called `kommander-cluster-admin-config`. The contents of this file are used in Kommander to attach the cluster.


Before importing this configuration, verify the `kubeconfig` file can access the cluster:

```
kubectl --kubeconfig $(pwd)/kommander-cluster-admin-config get all --all-namespaces
```

5.16.1.7.2.3 Finalize attaching your cluster from the UI

Now that you have kubeconfig, go to the DKP UI and follow these steps below:

- From the top menu bar, select your target workspace.
- On the Dashboard page, select the **Add Cluster** option in the **Actions** dropdown menu at the top right.
- Select **Attach Cluster**.
- Select the **No additional networking restrictions** card. Alternatively, if you must use network restrictions, stop following the steps below, and see the instructions on the page [Attach a cluster WITH network restrictions \(see page 772\)](#).
- Upload the kubeconfig file you created in the previous section (or copy its contents) into the **Cluster Configuration** section.
- The **Cluster Name** field automatically populates with the name of the cluster in the kubeconfig. You can edit this field with the name you want for your cluster.
- Add labels to classify your cluster as needed.
- Select **Create** to attach your cluster.

-  If a cluster has limited resources to deploy all the federated platform services, it will fail to stay attached in the DKP UI. If this happens, ensure your system has sufficient resources for all pods.

Next Steps:

[Day 2 - Cluster Operations Management](#) (see page 582)

5.17 GCP Install Options

For an environment that is Google Cloud Platform, an install is provided for you in this one location.

Remember, there are always more options for custom YAML in the [GCP Infrastructure](#) (see page 1365) section for Google Cloud, but this will get you operative in the most common scenario.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

Below are the supported environment variable combinations: (Refer to this page to check if your OS is a [Supported Operating System](#) (see page 69).)

- [GCP Install](#) (see page 563)

-  Additional Resource Information specific to GCP is below.

- **Control Plane Nodes** - DKP on GCP defaults to deploying an `n2-standard-4` instance with an 80GiB root volume for control plane nodes, which meets the above requirements.
- **Worker Nodes** - DKP on GCP defaults to deploying a `n2-standard-8` instance with an 80GiB root volume for worker nodes, which meets the above requirements.

5.17.1 GCP Install

For an environment that is on the GCP Infrastructure, install options are provided for you in this one location. Remember, there are always more options in the [GCP Infrastructure](#) (see page 1365) section under [Custom Installation and Additional Infrastructure Tools](#) (see page 1036), but this will get you up and running in the most common scenario.


If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)

- [Prerequisites for Install](#) (see page 121)

5.17.1.1 GCP Prerequisites:

- Verify that your Google Cloud project does not have the Enable OS Login feature enabled.

 The Enable OS Login feature is sometimes enabled by default in GCP projects. If the OS login feature is enabled, KIB will not be able to `ssh` to the VM instances it creates and will not be able to successfully create an image. To check if it is enabled, use the commands on this page https://cloud.google.com/compute/docs/metadata/setting-custom-metadata#console_2 to inspect the metadata configured in in your project. If you find the the `enable-oslogin` flag set to TRUE, you must remove (or set it to FALSE) to use KIB.


- The user creating the Service Accounts needs additional privileges in addition to the Editor role.
 - See [GCP Prerequisite Roles](#) (see page 0)

5.17.1.2 Next Step:

[GCP: Create Image](#) (see page 564)

5.17.1.3 GCP: Create Image

This procedure describes how to use the [Konvoy Image Builder](#) (see page 1461) (KIB) to create a [Cluster API](#)⁴³⁷ compliant GCP image. GCP images contain configuration information and software to create a specific, pre-configured, operating environment. For example, you can create a GCP image of your current computer system settings and software. The GCP image can then be replicated and distributed, creating your computer system for other users. KIB uses [variable overrides](#) (see page 1513) to specify base image and container images to use in your new GCP image.

 Google Cloud Platform does not publish images. You must first build the image using Konvoy Image Builder. For more information regarding images and clusters, refer to the [GCP Infrastructure](#) (see page 1365) section of the documentation.

5.17.1.3.1 Prerequisites

Before you begin, you must:

⁴³⁷ <https://cluster-api.sigs.k8s.io/>

- Download the [KIB \(see page 1463\)](#) bundle for your version of DKP.
- Check the [Supported Infrastructure Operating Systems \(see page 69\)](#)
- Check the [Supported Kubernetes Version \(see page 1552\)](#) for your Provider.
- Create a working Docker or other registry setup.
- On Debian-based Linux distributions, install a version of the [cri-tools](#)⁴³⁸ package known to be compatible with both the Kubernetes and container runtime versions.
- Verify that your Google Cloud project does not have the Enable OS Login feature enabled. See below for more information:



The Enable OS Login feature is sometimes enabled by default in GCP projects. If the OS login feature is enabled, KIB will not be able to `ssh` to the VM instances it creates and will not be able to successfully create an image.

To check if it is enabled, use the commands on this page https://cloud.google.com/compute/docs/metadata/setting-custom-metadata#console_2 to inspect the metadata configured in in your project. If you find the the `enable-oslogin` flag set to TRUE, you must remove (or set it to FALSE) to successfully use KIB.

5.17.1.3.2 GCP Prerequisite Roles

If you are creating your image on either a non-GCP instance or one that does not have the [required roles](#)⁴³⁹⁴⁴⁰, you must either:

- Create a [GCP service account](#)⁴⁴¹.
- If you have already created a service account, retrieve the credentials for an existing service account.
- Export the static credentials that will be used to create the cluster:

```
export GCP_B64ENCODED_CREDENTIALS=$(base64 < "$
{GOOGLE_APPLICATION_CREDENTIALS}" | tr -d '\n')
```



Make sure to rotate static credentials for increased security.

- If you are creating your image on either a non-GCP instance or one that does not have the [required roles \(see page 1365\)](#):
 - (option 1) Create a service account using the following `gcloud` commands:

438 <https://github.com/kubernetes-sigs/cri-tools>

439 <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/295272692>

440 <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/295272692>

441 <https://cloud.google.com/iam/docs/service-account-overview>


```

export GCP_PROJECT=<your GCP project ID>
export GCP_SERVICE_ACCOUNT_USER=<some new service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud iam service-accounts create "${GCP_SERVICE_ACCOUNT_USER}" --
project=${GCP_PROJECT}
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
compute.instanceAdmin.v1
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
iam.serviceAccountUser
gcloud iam service-accounts keys create $
{GOOGLE_APPLICATION_CREDENTIALS} --iam-account="$
{GCP_SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com"

```

- (option 2) If you have already created a service account, retrieve the credentials for an existing service account using the following `gcloud` commands:

```

export GCP_PROJECT=<your GCP project ID>
export GCP_SERVICE_ACCOUNT_USER=<existing service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
compute.instanceAdmin.v1
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
iam.serviceAccountUser
gcloud iam service-accounts keys create $
{GOOGLE_APPLICATION_CREDENTIALS} --iam-account="$
{GCP_SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com"

```

- Export the static credentials that will be used to create the cluster:

```

export GCP_B64ENCODED_CREDENTIALS=$(base64 < "$
{GOOGLE_APPLICATION_CREDENTIALS}" | tr -d '\n')

```



Make sure to rotate static credentials for increased security.

5.17.1.3.2.1 Build the GCP image

1. Run the `konvoy-image` command to build and validate the image:

```
./konvoy-image build gcp --project-id ${GCP_PROJECT} --network ${NETWORK_NAME}
images/gcp/ubuntu-2004.yaml
```

2. KIB will run and print out the name of the created image, you will use this name when creating a Kubernetes cluster. See sample output below:

```
...
==> ubuntu-2004-focal-v20220419: Deleting instance...
    ubuntu-2004-focal-v20220419: Instance has been deleted!
==> ubuntu-2004-focal-v20220419: Creating image...
==> ubuntu-2004-focal-v20220419: Deleting disk...
    ubuntu-2004-focal-v20220419: Disk has been deleted!
==> ubuntu-2004-focal-v20220419: Running post-processor: manifest
Build 'ubuntu-2004-focal-v20220419' finished after 7 minutes 46 seconds.

==> Wait completed after 7 minutes 46 seconds

==> Builds finished. The artifacts of successful builds are:
--> ubuntu-2004-focal-v20220419: A disk image was created: konvoy-ubuntu-2004-1-2
3-7-1658523168
--> ubuntu-2004-focal-v20220419: A disk image was created: konvoy-ubuntu-2004-1-2
3-7-1658523168
```

3. To find a list of images you have created in your account, run the following command:


```
gcloud compute images list --no-standard-images
```

5.17.1.3.2.2 Next Step:

[GCP: Cluster Creation](#) (see page 567)

5.17.1.4 GCP: Create the Management Cluster

Create a new Google Cloud Platform Kubernetes cluster in a non-air-gapped environment with the steps below.

 DKP uses the GCP CSI driver as the [default storage provider](#) (see page 91). Use a [Kubernetes CSI](#)⁴⁴² compatible storage that is suitable for production.


5.17.1.4.1 Name your cluster


1. Give your cluster a unique name suitable for your environment.

In GCP it is critical that the name is unique, as no two clusters in the same GCP account can have the same name.

2. Set the environment variable:

```
export CLUSTER_NAME=<gcp-example>
```

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)⁴⁴³ for more naming information.

 To increase [Docker Hub's rate limit](#)⁴⁴⁴ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username=<username> --registry-mirror-password=<password>` on the `dkp create cluster` command.

5.17.1.4.2 Create a new GCP cluster

Availability zones (AZs) are isolated locations within data center regions from which public cloud services originate and operate. Because all the nodes in a node pool are deployed in a single Availability Zone, you may wish to create additional node pools to ensure your cluster has nodes deployed in multiple Availability Zones.

⁴⁴² <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁴⁴³ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

⁴⁴⁴ <https://docs.docker.com/docker-hub/download-rate-limit/>

- By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single zone. You may create additional node pools in other zones with the `dkp create nodepool` command. The default region for the availability zones is `us-west1`.

- Google Cloud Platform does not publish images. You must first build the image using [Konvoy Image Builder](#) (see page 1461).

1. Create an image using [Konvoy Image Builder \(KIB\)](#) (see page 1461) and then export the image name:

```
export IMAGE_NAME=projects/${GCP_PROJECT}/global/images/<image_name_from_kib>
```

2. (Optional) You can modify Control Plane Audit logs settings using the information contained in the page [Configuring the Control Plane](#) (see page 1538).
3. (Optional) Determine what VPC Network to use. All GCP accounts come with a preconfigured VPC Network named `default`, which will be used if you do not specify a different network. To use a different VPC network for your cluster, create one by following these instructions for [Create and Manage VPC Networks](#)⁴⁴⁵. Then specify the `--network <new_vpc_network_name>` option on the create cluster command below. Follow the link for more information on [GCP Cloud Nat](#)⁴⁴⁶ and network flag.
4. Create a Kubernetes cluster. The following example shows a common configuration. See [dkp create cluster gcp](#) (see page 1678) reference for the full list of cluster creation options: Once the cluster creation process has finished, move onto [exploring your new cluster](#). (see page 1376)

- Ensure your [subnets](#) (see page 1046) do not overlap with your host subnet because they cannot be changed after cluster creation. If you need to change the kubernetes subnets, you must do this at cluster creation. The default subnets used in DKP are:

```
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
```

⁴⁴⁵ <https://cloud.google.com/vpc/docs/create-modify-vpc-networks#create-auto-network>

⁴⁴⁶ <https://github.com/kubernetes-sigs/cluster-api-provider-gcp/blob/3406adaae0f8f65a615844e55d856d306eddacc1/docs/book/src/topics/prerequisites.md#cloud-nat>

```

- 192.168.0.0/16
services:
  cidrBlocks:
- 10.96.0.0/12

```

```

dkp create cluster gcp \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--project=${GCP_PROJECT} \
--image=${IMAGE_NAME} \
--self-managed

```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#)⁴⁴⁷.

5. Wait for the cluster control-plane to be ready:

```

kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m

```

6. After the objects are created on the API server, the Cluster API controllers reconcile them. They create infrastructure and machines. As they progress, they update the Status of each object. Konvoy provides a command to describe the current status of the cluster:

```

dkp describe cluster -c ${CLUSTER_NAME}

```

NAME	SEVERITY	REASON	SINCE	MESSAGE	READY
Cluster/gcp-example					True
52s					
ClusterInfrastructure - GCPCluster/gcp-example					
ControlPlane - KubeadmControlPlane/gcp-example-control-plane					True
52s					

⁴⁴⁷ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273352944&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configuring+an+HTTP%2FHTTPS+Pr)

[oxy](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273352944&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configuring+an+HTTP%2FHTTPS+Pr)

```

| | Machine/gcp-example-control-plane-6fbzn True
2m32s
| | MachineInfrastructure - GCPMachine/gcp-example-control-plane-62g6s
| | Machine/gcp-example-control-plane-jf6s2 True
7m36s
| | MachineInfrastructure - GCPMachine/gcp-example-control-plane-bsr2z
| | Machine/gcp-example-control-plane-mnbfs True
54s
| | MachineInfrastructure - GCPMachine/gcp-example-control-plane-s8xsx
MachineDeployment/gcp-example-md-0 True
78s
| | Machine/gcp-example-md-0-68b86fddb8-8glsw True
2m49s
| | MachineInfrastructure - GCPMachine/gcp-example-md-0-zls8d
| | Machine/gcp-example-md-0-68b86fddb8-bvbm7 True
2m48s
| | MachineInfrastructure - GCPMachine/gcp-example-md-0-5zcvc
| | Machine/gcp-example-md-0-68b86fddb8-k9499 True
2m49s
| | MachineInfrastructure - GCPMachine/gcp-example-md-0-k8h5p
| | Machine/gcp-example-md-0-68b86fddb8-l6vfb True
2m49s
| | MachineInfrastructure - GCPMachine/gcp-example-md-0-9h5vn

```

- A self-managed cluster refers to one in which the CAPI resources and controllers that describe and manage it are running on the same cluster they are managing. As part of the underlying processing using the `--self-managed` flag, the DKP CLI:
 - creates a bootstrap cluster
 - creates a workload cluster
 - moves CAPI controllers from the bootstrap cluster to the workload cluster, making it self-managed
 - deletes the bootstrap cluster

To understand how this process works step by step, you can find customizable steps in [GCP Infrastructure](#) (see page 1365) under [Custom Installation and Additional Infrastructure Tools](#) (see page 1036).

5.17.1.4.3 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Verify%20your%20Cluster%20and%20DKP%20Installation&linkCreation=true&fromPageId=273320989>⁴⁴⁸⁴⁴⁹.

⁴⁴⁸ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

⁴⁴⁹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320989&linkCreation=true&spaceKey=DENT&title=%282.6%29+Verify+your+Cluster+and+DKP+Installation>

5.17.1.4.4 Next Steps:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20GCP%3A%20Install%20Kommander&linkCreation=true&fromPageId=273352944>⁴⁵⁰

5.17.1.5 GCP: Install Kommander

You have installed the Konvoy component and created a cluster. Now it is time to Install Kommander which will allow you to access the UI and attach new or existing clusters to monitor.

5.17.1.5.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1395).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

5.17.1.5.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} >> ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:


```
dkp install kommander --init > kommander.yaml
```

4. *If required:* Customize your `kommander.yaml`.

i See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, [Rook Ceph customization \(Pre-provisioned envs\)](#) (see page 1405), etc.


⁴⁵⁰ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273352944&linkCreation=true&spaceKey=DENT&title=%282.6%29+GCP%3A+Install+Kommander>

5.17.1.5.3 Enable DKP Catalog Applications and Install Kommander

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
        ref:
          tag: v2.6.1
```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.

- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

5.17.1.5.4 Next Step:

[GCP: Verify Install and Log in the UI](#) (see page 574)

5.17.1.6 GCP: Verify Install and Log in the UI

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.



NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the `install` command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
```

```

helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met

```

5.17.1.6.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```

kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'

```

Log in to the UI

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```

kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{"\n"}}Password: {{.data.password|base64decode}}
{{ "\n"}}'

```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubect1 -n kommander get svc kommander-traefik -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/
dashboard{{ "\n"}}
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 597](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 597](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

5.17.1.6.2 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 582](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

5.17.1.6.3 Next Step:


[GCP: Subsequent New Clusters](#) (see [page 576](#))

5.17.1.7 GCP: Create a Managed Cluster Using the DKP CLI

Enterprise

Gov Advanced


Use this procedure to create GCP Managed clusters using the DKP command line interface (CLI)

 When creating [Managed clusters](#) (see [page 81](#)), you do not need to create and move CAPI objects, or install the Kommander component. Those tasks are only done on [Management clusters](#) (see [page 81](#))!

Your new, managed cluster needs to be part of either a new Workspace or an existing Workspace under a management cluster. Create or locate a workspace name to get started with your managed cluster.

5.17.1.7.1 Make New Managed Cluster Part of a Workspace

1. If you have an existing Workspace name, run this command to find the name:

 **NOTE:** If you need to create a new Workspace, follow the instructions to [Create a New Workspace](#) (see [page 650](#)).

```
kubectl get workspace -A
```

2. When you have the Workspace name, set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```


5.17.1.7.2 Name your Cluster

1. Give your cluster a unique name suitable for your environment.


In GCP it is critical that the name is unique, as no two clusters in the same GCP account can have the same name.

2. Set the environment variable:

```
export MANAGED_CLUSTER_NAME=<gcp-additional>
```


 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)⁴⁵¹ for more naming information.


⁴⁵¹ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>


-  To increase [Docker Hub's rate limit](#)⁴⁵², use your Docker Hub credentials when creating the cluster by setting the following flag, `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username=<username> --registry-mirror-password=<password>` on the `dkp create cluster` command.

5.17.1.7.3 Create a New GCP Cluster

Availability zones (AZs) are isolated locations within data center regions from which public cloud services originate and operate. Because all the nodes in a node pool are deployed in a single Availability Zone, you may wish to create additional node pools to ensure your cluster has nodes deployed in multiple Availability Zones.

-  By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single zone. You may create additional node pools in other zones with the `dkp create nodepool` command. The default region for the availability zones is `us-west1`.

-  Google Cloud Platform does not publish images. You must first build the image using [Konvoy Image Builder](#) (see page 1461).

-  The below instructions tell you how to create a cluster and have it automatically attach to the workspace you set above. If you do not set a workspace, it will be created in the `default` workspace, and you need to take additional steps to attach to a workspace later. For instructions on how to do this, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273744255/2.6%2BAttach%2Ba%2BKubernetes%2BCluster> (see page 752).

1. Create an image using [Konvoy Image Builder \(KIB\)](#) (see page 1461) and then export the image name:

```
export IMAGE_NAME=projects/${GCP_PROJECT}/global/images/<image_name_from_kib>
```

⁴⁵² <https://docs.docker.com/docker-hub/download-rate-limit/>

2. Create a Kubernetes cluster. The following example shows a common configuration. See [dkp create cluster gcp](#) (see page 1678) reference for the full list of cluster creation options:

```
dkp create cluster gcp \
--cluster-name=${MANAGED_CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--namespace ${WORKSPACE_NAMESPACE} \
--project=${GCP_PROJECT} \
--image=${IMAGE_NAME} \
--kubeconfig=<management-cluster-kubeconfig-path>
```

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1039).

5.17.1.7.4 Retrieve the kubeconfig and Explore New GCP Cluster

Follow these steps:

1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} --kubeconfig
<management-cluster-kubeconfig-path> -n ${WORKSPACE_NAMESPACE} > $
{MANAGED_CLUSTER_NAME}.conf
```

2. List the nodes with the following command:


```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get nodes
```

3. List the pods with the following command:

NOTE: Wait for the Status to move to Ready while `calico-node` pods are being deployed.

```
kubectl --kubeconfig=${MANAGED_CLUSTER_NAME}.conf get pods -A
```

5.17.1.7.4.1 Manually Attach a DKP CLI Cluster to the Management Cluster

 These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

7. This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```

apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>

```

8. Create this secret in the desired workspace:

```

kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}

```

9. Create this `kommandercluster` object to attach the cluster to the workspace:

```

cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF

```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```

kubectl get kommanderclusters -A

```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

5.17.1.7.4.2 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 582)

6 Day 2 - Cluster Operations Management

Use these sections to manage your DKP environment.

- [Operations](#) (see page 582)
- [Applications](#) (see page 621)
- [Workspaces](#) (see page 650)
- [Projects](#) (see page 686)
- [Managing Clusters](#) (see page 745)
- [Backup and Restore](#) (see page 874)
- [Logging](#) (see page 899)
- [Security](#) (see page 934)
- [Networking](#) (see page 954)
- [GPUs](#) (see page 968)
- [Monitoring and Alerts](#) (see page 982)
- [Storage for Applications](#) (see page 1004)
- [DKP Troubleshooting](#) (see page 1019)
- [Telemetry](#) (see page 1028)

6.1 Operations

Manage your cluster and deployed applications using platform applications

After you deploy a DKP cluster and the platform applications you want to use, you are ready to begin managing cluster operations and their application workloads to optimize your organization's productivity.

In most cases, a production cluster requires additional advanced configuration tailored for your environment, ongoing maintenance, authentication and authorization, and other common activities. For example, it is important to monitor cluster activity and collect metrics to ensure application performance and response time, evaluate network traffic patterns, manage user access to services, and verify workload distribution and efficiency.

In addition to advanced configurations, you can also control the appearance of your DKP user interface by adding Banners and Footers. There are different options available depending on the DKP level that you license and install.

- [Access Control](#) (see page 583)
- [Identity Providers](#) (see page 597)
- [Infrastructure Providers](#) (see page 603)
- [Add a Header, Footer, and Logo to Your DKP Implementation](#) (see page 619)

6.1.1 Access Control

6.1.1.1 Centrally Manage Access Across Clusters

You can centrally define role-based authorization within DKP UI to control resource access on the management cluster and a set, or all, of the target clusters. These resources are similar to Kubernetes RBAC but with crucial differences, and they make it possible to define the roles and role bindings once, and have them federated to clusters within a given scope.

DKP UI has two conceptual groups of resources that are used to manage access control:

- **Kommander Roles:** control access to resources on the management cluster.
- **Cluster Roles:** control access to resources on all target clusters in scope.

Use these two groups of resources to manage access control within 3 levels of scope:

Context	Kommander Roles	Cluster Roles
Global	Create ClusterRoles on the management cluster.	Federates ClusterRoles on all target clusters across all workspaces.
Workspace	Create namespaced Roles on the management cluster in the workspace namespace.	Federates ClusterRoles on all target clusters in the workspace.
Project	Create namespaced Roles on the management cluster in the project namespace.	Federates namespaced Roles on all target clusters in the project in the project namespace.

The [role bindings](#) (see page 586) for each level and type create `RoleBindings` or `ClusterRoleBindings` on the clusters that apply to each category.

This approach gives you maximum flexibility over who has access to what resources, conveniently mapped to your existing identity providers' claims.

6.1.1.2 Special Limitation for Kommander Roles

In addition to granting a Kommander Role, you must also grant the appropriate DKP role to allow external users and groups into the UI. See [RBAC - DKP UI Authorization](#) (see page 586) for details about the built-in DKP roles. Here are examples of `ClusterRoleBindings` that grant an IDP group admin access to the Kommander routes:

```
---
```

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: eng-kommander-dashboard
  labels:
    "workspaces.kommander.mesosphere.io/rbac": ""
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: dkp-kommander-admin
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: oidc:engineering
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: eng-dkp-routes
  labels:
    "workspaces.kommander.mesosphere.io/rbac": ""
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: dkp-admin
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: oidc:engineering

```

6.1.1.3 Types of Access Control Objects

Kubernetes role-based access control can be controlled with three different object categories: Groups, Roles and Policies, as explained in more detail below.

6.1.1.3.1 Groups

You can map group and user claims made by your configured identity providers to Kommander groups by selecting Administration / Identity providers in the left sidebar in the global workspace level, and then selecting the **Groups** tab.

6.1.1.3.2 Roles

ClusterRoles are named collections of rules defining which verbs can be applied to which resources.

- Kommander Roles apply specifically to resources on the management cluster.
- Cluster Roles apply to target clusters within their scope at these levels:
 - Global level - this is all target clusters in all workspaces,

- Workspace level - this is all target clusters in the workspace,
- Project level - this is all target clusters that have been added to the project.

6.1.1.3.3 Propagate Workspace Roles to Projects

By default, users granted the Kommander Workspace Admin, Edit, or View roles will also be granted the equivalent Kommander Project Admin, Edit, or View role for any project created in the workspace. Other workspace roles are not automatically propagated to the equivalent role for a project in the workspace.

Each workspace has roles defined using `KommanderWorkspaceRole` resources. Automatic propagation is controlled using the annotation `"workspace.kommander.mesosphere.io/sync-to-project"`: `"true"` on a `KommanderWorkspaceRole` resource. You can manage this only by using the CLI.

```
kubectl get kommanderworkspaceroles -n test-qznrn-6sz52
```

NAME	DISPLAY NAME	AGE
kommander-workspace-admin	Kommander Workspace Admin Role	2m18s
kommander-workspace-edit	Kommander Workspace Edit Role	2m18s
kommander-workspace-view	Kommander Workspace View Role	2m18s

To prevent propagation of the `kommander-workspace-view` role, remove this annotation from the `KommanderWorkspaceRole` resource.

```
kubectl annotate kommanderworkspacerole -n test-qznrn-6sz52 kommander-workspace-view workspace.kommander.mesosphere.io/sync-to-project-
```

To enable propagation of the role, add this annotation to the relevant `KommanderWorkspaceRole` resource.

```
kubectl annotate kommanderworkspacerole -n test-qznrn-6sz52 kommander-workspace-view workspace.kommander.mesosphere.io/sync-to-project=true
```

6.1.1.3.4 Special Limitation for Workspace > Project Role Inheritance

When granting users access to a workspace, you must manually grant access to the projects within that workspace. Each project is created with a set of admin/edit/view roles, and you can choose to add an additional `RoleBinding` to each group or user of the workspace for one of these project roles. Usually these are prefixed `kommander-project-(admin/edit/view)`. Here is an example `RoleBinding` that grants the Kommander Project Admin role access for the project namespace to the engineering group:

```

apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: workspace-admin-project1-admin
  namespace: my-project-namespace-xxxxx
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: kommander-project-admin-xxxxx
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: oidc:engineering

```

6.1.1.3.5 Role Bindings

Kommander role bindings, cluster role bindings, and project role bindings bind a Kommander group to any number of roles. All groups defined in the **Groups** tab will be present at the global, workspace, or project level, and are ready for you to assign roles to them.

6.1.1.4 Related Information

- [Project Role Bindings](#) (see page 728)
- [Workspace Role Bindings](#) (see page 685)
- [Kommander RBAC Tutorial](#) (see page 586)
- [RBAC - DKP UI Authorization](#) (see page 586)
- [Kubernetes RBAC Authorization](#)⁴⁵³

6.1.1.5 Granting Access to Kubernetes and Kommander Resources

6.1.1.5.1 Grant access to Kommander and Kubernetes resources using RBAC

6.1.1.5.2 Granting Access to External Users

Users and groups from an external identity provider initially have no access to kubernetes resources. Privileges must be granted explicitly by interacting with the RBAC API. This section provides some basic examples for general usage. More information about the RBAC API can be found in the [Kubernetes documentation](#)⁴⁵⁴.

⁴⁵³ <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

⁴⁵⁴ <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

6.1.1.5.2.1 The Basics

Kubernetes does not provide an identity database for standard users. Users and group membership must be provided by a trusted identity provider. In Kubernetes, RBAC policies are additive, which means that a subject (user, group, or service account) is denied access to a resource unless explicitly granted access by a cluster administrator. You can grant access by binding a subject to a role, which grants some level of access to one or more resources. Kubernetes ships with some [default roles](#)⁴⁵⁵, which aid in creating broad access control policies.

For example, if you want to make `mary@example.com` a cluster administrator, bind her username to the `cluster-admin` default role as follows:

```
cat << EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: mary-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: mary@example.com
EOF
```

This user now has the highest level of access which can be achieved. Use the `cluster-admin` role and `system:masters` group sparingly.

6.1.1.5.2.2 Restricting a User to Namespace

A more common example would be to grant a user access to a specific namespace, by creating a RoleBinding (RoleBindings are namespace scoped). For example, to make the user `bob@example.com` a reader of the `baz` namespace, bind the user to the `view` role:

```
cat << EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: bob-view
  namespace: baz
roleRef:
  apiGroup: rbac.authorization.k8s.io
```

⁴⁵⁵ <https://kubernetes.io/docs/reference/access-authn-authz/rbac/#default-roles-and-role-bindings>

```

kind: ClusterRole
name: view
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: bob@example.com
EOF

```

The user can now perform non-destructive operations targeting resources in the `baz` namespace only.

6.1.1.5.2.3 Groups

If your external identity provider supports group claims, you can also bind groups to roles. To make the `devops` LDAP group administrators of the `production` namespace bind the group to the `admin` role:

```

cat << EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: devops-admin
  namespace: production
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: admin
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: oidc:devops
EOF

```

One important distinction from adding users is that all external groups are prefixed with `oidc:`, so a group name becomes `oidc:devops`. This prevents collision with locally defined groups.

6.1.1.5.3 DKP UI Authorization

The DKP UI, and other HTTP applications protected by Kommander forward authentication, are also authorized by the Kubernetes RBAC API. In addition to Kubernetes API resources, it is possible to define rules which map to HTTP URIs and HTTP verbs. Kubernetes RBAC calls these `nonResourceURLs`, Kommander forward authentication uses these rules to grant or deny access to HTTP endpoints.

6.1.1.5.3.1 Default Roles

Roles have been created for granting access to the dashboard and select applications which expose an HTTP server through the ingress controller. The `cluster-admin` role is actually a system role that defines grants permission to all actions (verbs) on any resource; including non-resource URLs. The default dashboard user is bound to this role.

Granting user `admin` privileges on `/dkp/*` grants `admin` privileges to all sub-resources, even if bindings exist for sub-resources with less privileges.

Dashboard	Role	Path	access
*	cluster-admin	*	read, write, delete
kommander	dkp-view	/dkp/*	read
kommander	dkp-edit	/dkp/*	read, write
kommander	dkp-admin	/dkp/*	read, write, delete
kommander-dashboard	dkp-kommander-view	/dkp/kommander/ dashboard/*	read
kommander-dashboard	dkp-kommander-edit	/dkp/kommander/ dashboard/*	read, write
kommander-dashboard	dkp-kommander-admin	/dkp/kommander/ dashboard/*	read, write, delete
alertmanager	dkp-kube-prometheus- stack-alertmanager- view	/dkp/alertmanager/*	read
alertmanager	dkp-kube-prometheus- stack-alertmanager-edit	/dkp/alertmanager/*	read, write
alertmanager	dkp-kube-prometheus- stack-alertmanager- admin	/dkp/alertmanager/*	read, write, delete
centralized-grafana	dkp-centralized- grafana-grafana-view	/dkp/kommander/ monitoring/grafana/*	read
centralized-grafana	dkp-centralized- grafana-grafana-edit	/dkp/kommander/ monitoring/grafana/*	read, write

Dashboard	Role	Path	access
centralized-grafana	dkp-centralized-grafana-grafana-admin	/dkp/kommander/monitoring/grafana/*	read, write, delete
centralized-kubecost	dkp-centralized-kubecost-view	/dkp/kommander/kubecost/*	read
centralized-kubecost	dkp-centralized-kubecost-edit	/dkp/kommander/kubecost/*	read, write
centralized-kubecost	dkp-centralized-kubecost-admin	/dkp/kommander/kubecost/*	read, write, delete
grafana	dkp-kube-prometheus-stack-grafana-view	/dkp/grafana/*	read
grafana	dkp-kube-prometheus-stack-grafana-edit	/dkp/grafana/*	read, write
grafana	dkp-kube-prometheus-stack-grafana-admin	/dkp/grafana/*	read, write, delete
grafana-logging	dkp-grafana-logging-view	/dkp/logging/grafana/*	read
grafana-logging	dkp-grafana-logging-edit	/dkp/logging/grafana/*	read, write
grafana-logging	dkp-grafana-logging-admin	/dkp/logging/grafana/*	read, write, delete
karma	dkp-karma-view	/dkp/kommander/monitoring/karma/*	read
karma	dkp-karma-edit	/dkp/kommander/monitoring/karma/*	read, write
karma	dkp-karma-admin	/dkp/kommander/monitoring/karma/*	read, write, delete

Dashboard	Role	Path	access
kubernetes-dashboard	dkp-kubernetes-dashboard-view	/dkp/kubernetes/*	read
kubernetes-dashboard	dkp-kubernetes-dashboard-edit	/dkp/kubernetes/*	read, write
kubernetes-dashboard	dkp-kubernetes-dashboard-admin	/dkp/kubernetes/*	read, write, delete
prometheus	dkp-kube-prometheus-stack-prometheus-view	/dkp/prometheus/*	read
prometheus	dkp-kube-prometheus-stack-prometheus-edit	/dkp/prometheus/*	read, write
prometheus	dkp-kube-prometheus-stack-prometheus-admin	/dkp/prometheus/*	read, write, edit
traefik	dkp-traefik-view	/dkp/traefik/*	read
traefik	dkp-traefik-edit	/dkp/traefik/*	read, edit
traefik	dkp-traefik-admin	/dkp/traefik/*	read, edit, delete
thanos	dkp-thanos-query-view	/dkp/kommander/monitoring/query/*	read
thanos	dkp-thanos-query-edit	/dkp/kommander/monitoring/query/*	read, write
thanos	dkp-thanos-query-admin	/dkp/kommander/monitoring/query/*	read, write, delete

This section provides a few examples of binding subjects to the default roles defined for the DKP UI endpoints.

6.1.1.5.3.2 Examples

User

To grant the user `mary@example.com` administrative access to all Kommander resources, bind the user to the `dkp-admin` role:

```
cat << EOF | kubectl apply -f -
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: dkp-admin-mary
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: dkp-admin
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: mary@example.com
EOF
```

If you inspect the role, you see what access is now granted:

```
kubectl describe clusterroles dkp-admin
```

```
Name:          dkp-admin
Labels:        app.kubernetes.io/instance=kommander
               app.kubernetes.io/managed-by=Helm
               app.kubernetes.io/version=v2.0.0
               helm.toolkit.fluxcd.io/name=kommander
               helm.toolkit.fluxcd.io/namespace=kommander
               rbac.authorization.k8s.io/aggregate-to-admin=true
Annotations:   meta.helm.sh/release-name: kommander
               meta.helm.sh/release-namespace: kommander
PolicyRule:
  Resources  Non-Resource URLs  Resource Names  Verbs
  -----  -
  [/dkp/*]   []                 []             [delete]
  [/dkp]     []                 []             [delete]
  [/dkp/*]   []                 []             [get]
  [/dkp]     []                 []             [get]
  [/dkp/*]   []                 []             [head]
  [/dkp]     []                 []             [head]
```

[/dkp/*]	[]	[post]
[/dkp]	[]	[post]
[/dkp/*]	[]	[put]
[/dkp]	[]	[put]

The user can now use the HTTP verbs HEAD, GET, DELETE, POST, and PUT when accessing any URL at or under `/dkp`. Provided the downstream application follows REST conventions, this effectively allows read, edit, and delete privileges.

NOTE: To allow users to access the DKP UI, ensure they have the appropriate `dkp-kommander-` role in addition to the Kommander roles granted in the DKP UI. For more information, see the [Access Control section of the Kommander documentation](#) (see page 586).

Group

In order to grant view access to the `/dkp/*` endpoints and edit access to the grafana logging endpoint to group `logging-ops`, create the following ClusterRoleBindings:

```
cat << EOF | kubectl apply -f -
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: dkp-view-logging-ops
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: dkp-view
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: oidc:logging-ops
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: dkp-logging-edit-logging-ops
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: dkp-logging-edit
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: oidc:logging-ops
EOF
```

Note: external groups must be prefixed by `oidc:`

Members of `logging-ops` are now able to `view` all resources under `/dkp` and edit all resources under `/dkp/logging/grafana`.

6.1.1.5.4 Accessing the Kubernetes Dashboard

The Kubernetes dashboard offloads authorization directly to the Kubernetes API server. Once authenticated, all users may access the dashboard at `/dkp/kubernetes/` without needing a `dkp` role. However, access to the underlying kubernetes resources exposed by the dashboard are protected by the cluster RBAC policy.

6.1.1.5.5 Further Reading

This page has provides some basic examples of operations which provide the building blocks of creating an access control policy. For more information about creating your own roles and advanced policies, we highly recommend reading the Kubernetes [RBAC documentation](#)⁴⁵⁶.

For information on how you can add a user to a cluster as an admin, refer to [Onboarding Users onto a DKP Cluster](#) (see page 594).


6.1.1.6 Onboarding Users onto a DKP Cluster

After you install DKP and create a cluster, use this procedure to add new users to your environment.

6.1.1.6.1 Prerequisites

- A Valid DKP License (Essential or Enterprise)
- A running Cluster

6.1.1.6.2 Onboarding a New User to a DKP Cluster

 This procedure assumes you are using a LDAP Connector and that you are a Cluster Admin. For information on how you add users using other types of connectors, refer to the following pages:

- [OIDC](#)⁴⁵⁷
- [SAML](#)⁴⁵⁸
- [GitHub](#)⁴⁵⁹

456 <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

457 <https://dexidp.io/docs/connectors/oidc/>

458 <https://dexidp.io/docs/connectors/saml/>

459 <https://dexidp.io/docs/connectors/github/>

1. Create an LDAP Connector definition and name it `ldap.yaml` . An example is shown below:

```

apiVersion: v1
kind: Secret
metadata:
  name: ldap-password
  namespace: kommander
type: Opaque
stringData:
  password: superSecret
---
apiVersion: dex.mesosphere.io/v1alpha1
kind: Connector
metadata:
  name: ldap
  namespace: kommander
spec:
  enabled: true
  type: ldap
  displayName: LDAP Test Connector
  ldap:
    host: ldapdce.testdomain
    insecureNoSSL: true
    bindDN: cn=ldapconnector,cn=testgroup,ou=testorg,dc=testdomain
    bindSecretRef:
      name: ldap-password
    userSearch:
      baseDN: dc=testdomain
      filter: "(objectClass=inetOrgPerson)"
      username: uid
      idAttr: uid
      emailAttr: uid
    groupSearch:
      baseDN: ou=testorg,dc=testdomain
      filter: "(objectClass=posixGroup)"
      userMatchers:
        - userAttr: uid
          groupAttr: memberUid
          nameAttr: cn

```

2. Add the connector using the following command:

```
kubectl apply -f ldap.yaml
```

- a. The output should look similar to the following:

```

secret/ldap-password created
connector.dex.mesosphere.io/ldap created

```

3. Add the appropriate role bindings and name the file `new_user.yaml`. Examples for both Single User and Group Bindings are shown below:

- For Single Users:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: cluster-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: newUser
```

- Group Binding

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: cluster-admin
  namespace: ml
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: Group
  name: oidc:kommanderAdmins
```

4. Add the role binding(s) using the following command:

```
kubectl apply -f new_user.yaml
```



ClusterRoleBindings permissions are applicable on a global level.
RoleBindings permissions are applicable on a namespace level.



This procedure assumes that the user being added is an admin.

For additional information about the other roles in DKP and their permissions, refer to [Granting Access to Kubernetes and Kommander Resources](#) (see page 589)

6.1.2 Identity Providers

6.1.2.1 Grant access to users in your organization

DKP supports [GitHub](#) (see page 598), [LDAP](#) (see page 599), SAML and standard OIDC identity providers such as Google. These identity management providers support the login and authentication process for DKP and your Kubernetes clusters. You can configure as many identity providers as you want, and users can select from any method when logging in.

6.1.2.2 Benefits of Using an External Identity Provider

Using an external identity provider is beneficial for many reasons:

- Centralized management of multiple users, to multiple clusters
- Centralized management of password rotation, expiration, etc.
- Support of 2-factor-authentication methods for increased security
- Separate storage of user credentials

6.1.2.3 Prerequisites

To get started with DKP, you must:

- [Log in to the UI](#) (see page 1420)

6.1.2.3.1 Limit Access

- The [GitHub](#) (see page 598) provider allows you to specify any of the organizations and teams are eligible for access.
- The [LDAP](#) (see page 599) provider allows you to configure search filters for either users or groups.
- The OIDC provider cannot limit users based on identity.
- The SAML provider allows users to log in using a single sign-on (SSO) profile.

6.1.2.3.2 Configure an Identity Provider via the UI

1. From the drop down, select the **Global** workspace.
2. Select **Administration > Identity Providers**.
3. Select the **Identity Providers** tab.
4. Select **+ Add Identity Provider**.

5. Select an identity provider and complete the form field with the relevant details.
6. Select **Save** to create your Identity Provider.

6.1.2.3.3 Temporarily Disabling a Provider

Select the three dot button on the Identity Providers table and select **Disable** from the drop-down menu. The provider option no longer appears on the login screen.

6.1.2.4 Groups

Access control groups are configured in the Groups tab of the Identity Providers page. See [Access Control](#) (see page 586) for an overview of groups in DKP.

6.1.2.5 Authorize a Group in Github

Enterprise

Gov Advanced

Install GitHub as an identity provider and grant access to all developers.

6.1.2.5.1 Authorize Access to Your Clusters

Ensure every developer in your GitHub organization has access to your Kubernetes clusters. To authorize all developers to have read access to your clusters:

1. Set up GitHub as an identity provider. Start by creating a new OAuth Application in our GitHub Organization by filling out [this form](#)⁴⁶⁰.

Important: Use your cluster URL followed by `/dex/callback` as the Authorization callback URL.

2. After you create the application, you will be taken to a settings page. You will need the **Client ID** and **Client Secret** from this page for the DKP UI. Select the **Generate a new client secret** button if you do not already have a Client Secret for the application.
3. From the top menu bar in the DKP UI, select the **Global** workspace.
4. Select **Identity Providers** in the **Administration** section of the sidebar menu.
5. Select the **Identity Providers** tab, and then select the **+ Add Identity Provider** button.
6. Ensure GitHub is selected as the identity provider type, and copy the **Client ID** and **Client Secret** values into the form.
7. Select **Save** to create your Identity Provider.

D2iQ configured the identity provider to load all groups, so you must map these groups to the Kubernetes groups.

⁴⁶⁰ <https://github.com/settings/applications/new>

6.1.2.5.1.1 Map the Identity Provider Groups to the Kubernetes Groups

1. Select the **Groups** tab, and then select the **Create Group** button.
2. Give your group a descriptive name and add the groups from your GitHub provider under **Identity Provider Groups**.
3. Click **Save** to create the group, which creates it on the management cluster and federated to all target clusters, and also describes the developers for your organization.

To enable this group, you need to first create a role which allows you to view every resource.

6.1.2.5.1.2 Create a “Read Everything” Role

1. Select **Access Control** in the **Administration** section of the sidebar menu.
2. Select the **Cluster Roles** tab, and then select the **+ Create Role** button.
3. Give the role a descriptive name, and ensure that **Cluster Role** is selected as the type.
4. For a read-only role, select **+ Add Rule**, then select **All Resource Types** in the **Resources** input, and select the **get**, **list**, and **watch** verbs.

Now you can assign the “Read Everything” role to the developers group.

6.1.2.5.1.3 Assign the Role to the Developers Group

1. Select the **Cluster Role Bindings** tab, and then select the **Add roles** button for your group.
2. Select “Read Everything” role from the **Roles** drop-down.

Lastly, follow the example in the [Access Control documentation](#) (see page 586) to grant users access to Kommander routes on your cluster.

When you check your attached clusters and login as a user from your matched groups, you can see every resource, but neither delete or edit them, as intended.


6.1.2.6 External LDAP directory

6.1.2.6.1 How to connect your cluster to an external LDAP directory

This guide shows you how to configure your DKP cluster so that users can log in with the credentials stored in an external LDAP directory service.

6.1.2.6.2 Step 1: Add LDAP connector

Each LDAP directory is set up in its own specific manner, so these steps are important. The LDAP authentication mechanism can be added using the CLI or the UI in Kommander.

 The following example does not cover all possible configurations. Refer to the [Dex LDAP connector reference documentation](#)⁴⁶¹ for more details.

The example below configures a DKP cluster to connect to the [Online LDAP Test Server](#)⁴⁶² and for demonstration purposes, the configuration shown uses `insecureNoSSL: true`. In production, you should protect LDAP communication with a properly-configured transport layer security (TLS). When using TLS, the admin can add `insecureSkipVerify: true` to `spec.ldap` to skip server certificate verification, if needed.

Below are steps for the CLI followed by UI in the event you prefer to use it.

1. Create a YAML file (`ldap.yaml`) similar to the following:

```
apiVersion: v1
kind: Secret
metadata:
  name: ldap-password
  namespace: kommander
type: Opaque
stringData:
  password: password
---
apiVersion: dex.mesosphere.io/v1alpha1
kind: Connector
metadata:
  name: ldap
  namespace: kommander
spec:
  enabled: true
  type: ldap
  displayName: LDAP Test
  ldap:
    host: ldap.forumsys.com:389
    insecureNoSSL: true
    bindDN: cn=read-only-admin,dc=example,dc=com
    bindSecretRef:
      name: ldap-password
    userSearch:
      baseDN: dc=example,dc=com
      filter: "(objectClass=inetOrgPerson)"
      username: uid
      idAttr: uid
      emailAttr: mail
    groupSearch:
```

⁴⁶¹ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/ldap.md>

⁴⁶² <https://www.forumsys.com/tutorials/integration-how-to/ldap/online-ldap-test-server/>

```
baseDN: dc=example,dc=com
filter: "(objectClass=groupOfUniqueNames)"
userMatchers:
- userAttr: DN
  groupAttr: uniqueMember
  nameAttr: ou
```

NOTE: The value for the LDAP connector `name` parameter (here: `LDAP Test`) appears on one of the login buttons in the DKP UI. You should choose an expressive name.

2. Run the following command to deploy the LDAP connector.

```
kubectl apply -f ldap.yaml
```

NOTE: UI - Deploy LDAP through **Add Identity Provider** screen within the UI.

3. Retrieve a list of connectors:

```
kubectl get connector.dex.mesosphere.io -A
```

4. Run the following command to verify that the LDAP connector was created successfully:

```
kubectl get Connector.dex.mesosphere.io -n kommander <LDAP-CONNECTOR-NAME> -o
yaml
```

6.1.2.6.3 Step 2: Log in

1. Visit `https://<YOUR-CLUSTER-HOST>/token` and initiate a login flow.
2. On the login page choose the `Log in with <ldap-name>` button.
3. Enter the LDAP credentials, and log in.



UI - While LDAP authentication has been enabled, additional access rights will need to be configured through the Add Identity Provider screen in the UI using the documentation for [Granting Access to Kubernetes and Kommander Resources](#)⁴⁶³.

⁴⁶³ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273288422&linkCreation=true&spaceKey=DENT&title=%282.6%29+Granting+Access+to+Kubernetes+and+Kommander+Resources)

[fromPageId=273288422&linkCreation=true&spaceKey=DENT&title=%282.6%29+Granting+Access+to+Kubernetes+and+Kommander+Resources](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273288422&linkCreation=true&spaceKey=DENT&title=%282.6%29+Granting+Access+to+Kubernetes+and+Kommander+Resources)

6.1.2.6.4 Troubleshooting

It is likely that the Dex LDAP connector configuration is not quite right from the start. In that case, you need to be able to debug the problem and iterate on it. The Dex log output contains helpful error messages as indicated by the following examples.

6.1.2.6.4.1 Errors during Dex startup

If the Dex configuration fragment provided results in an invalid Dex config, Dex does not properly start up. In that case, reviewing the Dex logs will provide error details. Use the following command to retrieve the Dex logs:

```
kubectl logs -f dex-66675fcb7c-snxb8 -n kommander
```

You may see an error similar to the following:

```
error parse config file /etc/dex/cfg/config.yaml: error unmarshaling JSON: parse
connector config: illegal base64 data at input byte 0
```

Another reason for Dex not starting up correctly is that `https://<YOUR-CLUSTER-HOST>/token` throws a 5xx HTTP error response after timing out.

6.1.2.6.4.2 Errors upon login

Most problems with the Dex LDAP connector configuration become apparent only after a login attempt. A login failing from misconfiguration will result in an error page showing only `Internal Server Error` and `Login error`. You can then usually find the root cause by reading the Dex log, as shown in the following example:

```
kubectl logs -f dex-5d55b6b94b-9pm2d -n kommander
```

You can look for output similar to this example:

```
[...]
time="2019-07-29T13:03:57Z" level=error msg="Failed to login user: failed to connect:
LDAP Result Code 200 \"Network Error\": dial tcp: lookup freeipa.example.com on
10.255.0.10:53: no such host"
```

Here, the directory's DNS name was misconfigured, which should be easy to address.

A more difficult problem occurs when a login through Dex through LDAP fails because Dex was not able to find the specified user unambiguously in the directory. That could be the result of an invalid LDAP user search configuration. Here's an example error message from the Dex log:

```
time="2019-07-29T14:21:27Z" level=info msg="performing ldap search
cn=users,cn=compat,dc=demo1,dc=freeipa,dc=org sub (&(objectClass=posixAccount)
(uid=employee))"
time="2019-07-29T14:21:27Z" level=error msg="Failed to login user: ldap: filter
returned multiple (2) results: \"(&(objectClass=posixAccount)(uid=employee))\""
```

Solving problems like this requires you to review the directory structure carefully. (Directory structures can be very different between different LDAP setups.) Then you must carefully assemble a user search configuration matching the directory structure.

Notably, with some directories, it can be hard to distinguish between the cases “properly configured, and user not found” (login fails in an expected way) and “not properly configured, and therefore user not found” (login fails in an unexpected way).

6.1.2.6.4.3 Example for successful login

For comparison, here are some sample log lines issued by Dex after successful login:

```
time="2019-07-29T15:35:51Z" level=info msg="performing ldap search
cn=accounts,dc=demo1,dc=freeipa,dc=org sub (&(objectClass=posixAccount)
(uid=employee))"
time="2019-07-29T15:35:52Z" level=info msg="username \"employee\" mapped to entry
uid=employee,cn=users,cn=accounts,dc=demo1,dc=freeipa,dc=org"
time="2019-07-29T15:35:52Z" level=info msg="login successful: connector \"ldap\",
username=\"\", email=\"employee@demo1.freeipa.org\", groups=[]"
```

6.1.3 Infrastructure Providers

Enterprise

Gov Advanced

Infrastructure providers, like AWS, Azure, and vSphere, provide the infrastructure for your Management clusters. You may have many accounts for a single infrastructure provider. To automate cluster provisioning, DKP needs authentication keys for your preferred infrastructure provider.

To provision new clusters and manage them, DKP also needs infrastructure provider credentials. Currently, you can create infrastructure providers records for [AWS](#) (see page 0), [Azure](#) (see page 746), and [vSphere](#) (see page 0) in the DKP user interface.



Infrastructure provider credentials are configured in each workspace. The name you assign must be unique across all namespaces in your cluster.

6.1.3.1 View and Modify Infrastructure Providers

You can use the DKP user interface to view, create, and delete infrastructure provider records.

1. From the top menu bar, select your target workspace.
2. Select **Infrastructure Providers** in the **Administration** section of the sidebar menu.

Refer to the information in the linked topics that follow for more information about creating infrastructure providers and clusters.

6.1.3.1.1 AWS

- [Configure AWS Provider with Role Credentials \(see page 604\)](#) (Recommended if using AWS)
- [Configure AWS Provider with Static Credentials \(see page 611\)](#)

6.1.3.1.2 Azure

- [Create an Azure Infrastructure Provider in the DKP UI \(see page 618\)](#)
- [Create a new Azure Cluster from the DKP UI \(see page 746\)](#)

6.1.3.1.3 vSphere

- [Create a vSphere Infrastructure Provider in DKP \(see page 619\)](#)
- [Create a Managed Cluster on vSphere from the DKP UI \(see page 747\)](#)

6.1.3.2 Delete an infrastructure provider

Before deleting an infrastructure provider, DKP verifies whether any existing managed clusters were created using this provider. You cannot delete an infrastructure provider until you delete all clusters created with that infrastructure provider. This ensures DKP has access to your infrastructure provider to remove all resources created for a managed cluster.

6.1.3.3 Configure an AWS Infrastructure Provider with a User Role

Enterprise

Gov Advanced

Configure your provider to add resources to your AWS account



We highly recommend using the role-based method as this is more secure.

 The role authentication method can only be used if your management cluster is running in AWS.

For more flexible credential configuration, we offer a [role-based authentication](#)⁴⁶⁴ method with an optional External ID for third party access.

6.1.3.3.1 Create a Role Manually

The role should grant permissions to create the following resources in the AWS account:

- EC2 Instances
- VPC
- Subnets
- Elastic Load Balancer (ELB)
- Internet Gateway
- NAT Gateway
- Elastic Block Storage (EBS) Volumes
- Security Groups
- Route Tables
- IAM Roles

The user you delegate from your role must have a minimum set of permissions. Below is the minimal IAM policy required:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:AllocateAddress",
        "ec2:AssociateRouteTable",
        "ec2:AttachInternetGateway",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateInternetGateway",
        "ec2:CreateNatGateway",
        "ec2:CreateRoute",
        "ec2:CreateRouteTable",
        "ec2:CreateSecurityGroup",
        "ec2:CreateSubnet",
        "ec2:CreateTags",
        "ec2:CreateVpc",
        "ec2:ModifyVpcAttribute",
        "ec2>DeleteInternetGateway",
        "ec2>DeleteNatGateway",

```

⁴⁶⁴ <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>


```

"ec2:DeleteRouteTable",
"ec2:DeleteSecurityGroup",
"ec2:DeleteSubnet",
"ec2:DeleteTags",
"ec2:DeleteVpc",
"ec2:DescribeAccountAttributes",
"ec2:DescribeAddresses",
"ec2:DescribeAvailabilityZones",
"ec2:DescribeInstances",
"ec2:DescribeInternetGateways",
"ec2:DescribeImages",
"ec2:DescribeNatGateways",
"ec2:DescribeNetworkInterfaces",
"ec2:DescribeNetworkInterfaceAttribute",
"ec2:DescribeRouteTables",
"ec2:DescribeSecurityGroups",
"ec2:DescribeSubnets",
"ec2:DescribeVpcs",
"ec2:DescribeVpcAttribute",
"ec2:DescribeVolumes",
"ec2:DetachInternetGateway",
"ec2:DisassociateRouteTable",
"ec2:DisassociateAddress",
"ec2:ModifyInstanceAttribute",
"ec2:ModifyNetworkInterfaceAttribute",
"ec2:ModifySubnetAttribute",
"ec2:ReleaseAddress",
"ec2:RevokeSecurityGroupIngress",
"ec2:RunInstances",
"ec2:TerminateInstances",
"tag:GetResources",
"elasticloadbalancing:AddTags",
"elasticloadbalancing:CreateLoadBalancer",
"elasticloadbalancing:ConfigureHealthCheck",
"elasticloadbalancing>DeleteLoadBalancer",
"elasticloadbalancing:DescribeLoadBalancers",
"elasticloadbalancing:DescribeLoadBalancerAttributes",
"elasticloadbalancing:ApplySecurityGroupsToLoadBalancer",
"elasticloadbalancing:DescribeTags",
"elasticloadbalancing:ModifyLoadBalancerAttributes",
"elasticloadbalancing:RegisterInstancesWithLoadBalancer",
"elasticloadbalancing:DeregisterInstancesFromLoadBalancer",
"elasticloadbalancing:RemoveTags",
"autoscaling:DescribeAutoScalingGroups",
"autoscaling:DescribeInstanceRefreshes",
"ec2:CreateLaunchTemplate",
"ec2:CreateLaunchTemplateVersion",
"ec2:DescribeLaunchTemplates",
"ec2:DescribeLaunchTemplateVersions",
"ec2>DeleteLaunchTemplate",
"ec2>DeleteLaunchTemplateVersions",
"ec2:DescribeKeyPairs"

```

```

    ],
    "Resource": ["*"]
  },
  {
    "Effect": "Allow",
    "Action": [
      "autoscaling:CreateAutoScalingGroup",
      "autoscaling:UpdateAutoScalingGroup",
      "autoscaling:CreateOrUpdateTags",
      "autoscaling:StartInstanceRefresh",
      "autoscaling>DeleteAutoScalingGroup",
      "autoscaling>DeleteTags"
    ],
    "Resource": [
      "arn:*:autoscaling:*:*:autoScalingGroup:*:autoScalingGroupName/*"
    ]
  },
  {
    "Effect": "Allow",
    "Action": ["iam:CreateServiceLinkedRole"],
    "Resource": [
      "arn:*:iam:*:role/aws-service-role/autoscaling.amazonaws.com/
AWSServiceRoleForAutoScaling"
    ],
    "Condition": {
      "StringLike": { "iam:AWSServiceName": "autoscaling.amazonaws.com" }
    }
  },
  {
    "Effect": "Allow",
    "Action": ["iam:CreateServiceLinkedRole"],
    "Resource": [
      "arn:*:iam:*:role/aws-service-role/elasticloadbalancing.amazonaws.com/
AWSServiceRoleForElasticLoadBalancing"
    ],
    "Condition": {
      "StringLike": {
        "iam:AWSServiceName": "elasticloadbalancing.amazonaws.com"
      }
    }
  },
  {
    "Effect": "Allow",
    "Action": ["iam:CreateServiceLinkedRole"],
    "Resource": [
      "arn:*:iam:*:role/aws-service-role/spot.amazonaws.com/
AWSServiceRoleForEC2Spot"
    ],
    "Condition": {
      "StringLike": { "iam:AWSServiceName": "spot.amazonaws.com" }
    }
  },
  },

```

```

{
  "Effect": "Allow",
  "Action": ["iam:PassRole"],
  "Resource": ["arn::*:iam::*:role/*cluster-api-provider-aws.sigs.k8s.io"]
},
{
  "Effect": "Allow",
  "Action": [
    "secretsmanager:CreateSecret",
    "secretsmanager>DeleteSecret",
    "secretsmanager:TagResource"
  ],
  "Resource": ["arn::*:secretsmanager::*:secret:aws.cluster.x-k8s.io/*"]
},
{
  "Effect": "Allow",
  "Action": ["ssm:GetParameter"],
  "Resource": ["arn::*:ssm::*:parameter/aws/service/eks/optimized-ami/*"]
},
{
  "Effect": "Allow",
  "Action": ["iam:CreateServiceLinkedRole"],
  "Resource": [
    "arn::*:iam::*:role/aws-service-role/eks.amazonaws.com/
AWSServiceRoleForAmazonEKS"
  ],
  "Condition": {
    "StringLike": { "iam:AWSServiceName": "eks.amazonaws.com" }
  }
},
{
  "Effect": "Allow",
  "Action": ["iam:CreateServiceLinkedRole"],
  "Resource": [
    "arn::*:iam::*:role/aws-service-role/eks-nodegroup.amazonaws.com/
AWSServiceRoleForAmazonEKSNodegroup"
  ],
  "Condition": {
    "StringLike": { "iam:AWSServiceName": "eks-nodegroup.amazonaws.com" }
  }
},
{
  "Effect": "Allow",
  "Action": ["iam:CreateServiceLinkedRole"],
  "Resource": [
    "arn:aws:iam::*:role/aws-service-role/eks-fargate-pods.amazonaws.com/
AWSServiceRoleForAmazonEKSFargate"
  ],
  "Condition": {
    "StringLike": { "iam:AWSServiceName": "eks-fargate.amazonaws.com" }
  }
},

```

```

{
  "Effect": "Allow",
  "Action": ["iam:GetRole", "iam:ListAttachedRolePolicies"],
  "Resource": ["arn:*:iam::*:role/*"]
},
{
  "Effect": "Allow",
  "Action": ["iam:GetPolicy"],
  "Resource": ["arn:aws:iam::aws:policy/AmazonEKSClusterPolicy"]
},
{
  "Effect": "Allow",
  "Action": [
    "eks:DescribeCluster",
    "eks:ListClusters",
    "eks:CreateCluster",
    "eks:TagResource",
    "eks:UpdateClusterVersion",
    "eks>DeleteCluster",
    "eks:UpdateClusterConfig",
    "eks:UntagResource",
    "eks:UpdateNodegroupVersion",
    "eks:DescribeNodegroup",
    "eks>DeleteNodegroup",
    "eks:UpdateNodegroupConfig",
    "eks:CreateNodegroup",
    "eks:AssociateEncryptionConfig"
  ],
  "Resource": ["arn:*:eks:*:*:cluster/*", "arn:*:eks:*:*:nodegroup/*/*/*"]
},
{
  "Effect": "Allow",
  "Action": [
    "eks:ListAddons",
    "eks:CreateAddon",
    "eks:DescribeAddonVersions",
    "eks:DescribeAddon",
    "eks>DeleteAddon",
    "eks:UpdateAddon",
    "eks:TagResource",
    "eks:DescribeFargateProfile",
    "eks:CreateFargateProfile",
    "eks>DeleteFargateProfile"
  ],
  "Resource": ["*"]
},
{
  "Effect": "Allow",
  "Action": ["iam:PassRole"],
  "Resource": ["*"],
  "Condition": {
    "StringEquals": { "iam:PassedToService": "eks.amazonaws.com" }
  }
}

```

```

    }
  },
  {
    "Effect": "Allow",
    "Action": ["kms:CreateGrant", "kms:DescribeKey"],
    "Resource": ["*"],
    "Condition": {
      "ForAnyValue:StringLike": {
        "kms:ResourceAliases": "alias/cluster-api-provider-aws-*"
      }
    }
  }
]
}

```

Make sure to also add a correct [trust relationship](#)⁴⁶⁵ to the created role. This example allows everyone within the same account to `AssumeRole` with the created role. `YOURACCOUNTRESTRICTION` must be replaced with the AWS Account ID you would like to `AssumeRole` from.



IMPORTANT: Never add a `*/` wildcard. This opens your account to the whole world.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "ec2.amazonaws.com",
        "AWS": "arn:aws:iam::YOURACCOUNTRESTRICTION:root"
      },
      "Action": "sts:AssumeRole"
    }
  ]
}

```

To use the role created, attach the following policy to the role which is already attached to your Managed or Attached cluster. Replace `YOURACCOUNTRESTRICTION` with the AWS Account ID where the role you like to `AssumeRole` is saved. Also, replace `THEROLEYOUCREATED` with the AWS Role name.

```

{
  "Version": "2012-10-17",
  "Statement": [
    {

```

⁴⁶⁵ <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

```

    "Sid": "AssumeRoleKommander",
    "Effect": "Allow",
    "Action": "sts:AssumeRole",
    "Resource": "arn:aws:iam::YOURACCOUNTRESTRICTION:role/THEROLEYOUCREATED"
  }
]
}

```

6.1.3.3.2 Create an AWS Infrastructure Provider in DKP

1. From the top menu bar, select your target workspace.
2. Select **Infrastructure Providers** in the **Administration** section of the sidebar menu.
3. Select the **Add Infrastructure Provider** button.
4. Select the **Amazon Web Services (AWS)** option.
5. Ensure **Role** is selected as the **Authentication Method**.
6. Enter a name for your infrastructure provider. Select a name that matches the AWS user.
7. Enter the **Role ARN**.
8. You can add an **External ID** if you share the Role with a 3rd party. External IDs secure your environment from accidentally used roles. [Read more about External IDs](#)⁴⁶⁶.
9. Select **Save** to save your provider.

6.1.3.4 Configure an AWS Infrastructure Provider with Static Credentials

Enterprise

Gov Advanced

When configuring an infrastructure provider with static credentials, you need an access ID and secret key for a user with a set of minimum capabilities.

6.1.3.4.1 Create a New User Using CLI Commands

You will need to have the [AWS CLI utility installed](#)⁴⁶⁷. Create a new user with the AWS CLI commands below:

```
aws iam create-user --user-name Kommander
```

```
aws iam create-policy --policy-name kommander-policy --policy-document
'{"Version":"2012-10-17","Statement":[{"Effect":"Allow","Action":
["ec2:AllocateAddress","ec2:AssociateRouteTable","ec2:AttachInternetGateway","ec2:AuthorizeSecurityG
roupIngress","ec2:CreateInternetGateway","ec2:CreateNatGateway","ec2:CreateRoute","ec2:CreateRouteTa
ble","ec2:CreateSecurityGroup","ec2:CreateSubnet","ec2:CreateTags","ec2:CreateVpc","ec2:ModifyVpcAtt
ribute","ec2:DeleteInternetGateway","ec2:DeleteNatGateway","ec2:DeleteRouteTable","ec2:DeleteSecurit
```

⁴⁶⁶ https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user_externalid.html

⁴⁶⁷ <https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html>

```

yGroup", "ec2:DeleteSubnet", "ec2:DeleteTags", "ec2:DeleteVpc", "ec2:DescribeAccountAttributes", "ec2:DescribeAddresses", "ec2:DescribeAvailabilityZones", "ec2:DescribeInstances", "ec2:DescribeInternetGateways", "ec2:DescribeImages", "ec2:DescribeNatGateways", "ec2:DescribeNetworkInterfaces", "ec2:DescribeNetworkInterfaceAttribute", "ec2:DescribeRouteTables", "ec2:DescribeSecurityGroups", "ec2:DescribeSubnets", "ec2:DescribeVpcs", "ec2:DescribeVpcAttribute", "ec2:DescribeVolumes", "ec2:DetachInternetGateway", "ec2:DisassociateRouteTable", "ec2:DisassociateAddress", "ec2:ModifyInstanceAttribute", "ec2:ModifyNetworkInterfaceAttribute", "ec2:ModifySubnetAttribute", "ec2:ReleaseAddress", "ec2:RevokeSecurityGroupIngress", "ec2:RunInstances", "ec2:TerminateInstances", "tag:GetResources", "elasticloadbalancing:AddTags", "elasticloadbalancing:CreateLoadBalancer", "elasticloadbalancing:ConfigureHealthCheck", "elasticloadbalancing:DeleteLoadBalancer", "elasticloadbalancing:DescribeLoadBalancers", "elasticloadbalancing:DescribeLoadBalancerAttributes", "elasticloadbalancing:ApplySecurityGroupsToLoadBalancer", "elasticloadbalancing:DescribeTags", "elasticloadbalancing:ModifyLoadBalancerAttributes", "elasticloadbalancing:RegisterInstancesWithLoadBalancer", "elasticloadbalancing:DeregisterInstancesFromLoadBalancer", "elasticloadbalancing:RemoveTags", "autoscaling:DescribeAutoScalingGroups", "autoscaling:DescribeInstanceRefreshes", "ec2:CreateLaunchTemplate", "ec2:CreateLaunchTemplateVersion", "ec2:DescribeLaunchTemplates", "ec2:DescribeLaunchTemplateVersions", "ec2:DeleteLaunchTemplate", "ec2:DeleteLaunchTemplateVersions", "ec2:DescribeKeyPairs"], "Resource": ["*"]}, {"Effect": "Allow", "Action": ["autoscaling:CreateAutoScalingGroup", "autoscaling:UpdateAutoScalingGroup", "autoscaling:CreateOrUpdateTags", "autoscaling:StartInstanceRefresh", "autoscaling:DeleteAutoScalingGroup", "autoscaling:DeleteTags"], "Resource": ["arn:*:autoscaling:*:*:autoScalingGroup:*:autoScalingGroupName/*"]}, {"Effect": "Allow", "Action": ["iam:CreateServiceLinkedRole"], "Resource": ["arn:*:iam:*:*:role/aws-service-role/autoscaling.amazonaws.com/AWSServiceRoleForAutoScaling"], "Condition": {"StringLike": {"iam:AWSServiceName": "autoscaling.amazonaws.com"}}}, {"Effect": "Allow", "Action": ["iam:CreateServiceLinkedRole"], "Resource": ["arn:*:iam:*:*:role/aws-service-role/elasticloadbalancing.amazonaws.com/AWSServiceRoleForElasticLoadBalancing"], "Condition": {"StringLike": {"iam:AWSServiceName": "elasticloadbalancing.amazonaws.com"}}}, {"Effect": "Allow", "Action": ["iam:CreateServiceLinkedRole"], "Resource": ["arn:*:iam:*:*:role/aws-service-role/spot.amazonaws.com/AWSServiceRoleForEC2Spot"], "Condition": {"StringLike": {"iam:AWSServiceName": "spot.amazonaws.com"}}}, {"Effect": "Allow", "Action": ["iam:PassRole"], "Resource": ["arn:*:iam:*:*:role/*:cluster-api-provider-aws.sigs.k8s.io"]}, {"Effect": "Allow", "Action": ["secretsmanager:CreateSecret", "secretsmanager:DeleteSecret", "secretsmanager:TagResource"], "Resource": ["arn:*:secretsmanager:*:*:secret:aws.cluster.x-k8s.io/*"]}, {"Effect": "Allow", "Action": ["ssm:GetParameter"], "Resource": ["arn:*:ssm:*:*:parameter:aws/service/eks/optimized-ami/*"]}, {"Effect": "Allow", "Action": ["iam:CreateServiceLinkedRole"], "Resource": ["arn:*:iam:*:*:role/aws-service-role/eks.amazonaws.com/AWSServiceRoleForAmazonEKS"], "Condition": {"StringLike": {"iam:AWSServiceName": "eks.amazonaws.com"}}}, {"Effect": "Allow", "Action": ["iam:CreateServiceLinkedRole"], "Resource": ["arn:*:iam:*:*:role/aws-service-role/eks-nodegroup.amazonaws.com/AWSServiceRoleForAmazonEKSNodegroup"], "Condition": {"StringLike": {"iam:AWSServiceName": "eks-nodegroup.amazonaws.com"}}}, {"Effect": "Allow", "Action": ["iam:CreateServiceLinkedRole"], "Resource": ["arn:aws:iam:*:*:role/aws-service-role/eks-fargate-pods.amazonaws.com/AWSServiceRoleForAmazonEKSFargate"], "Condition": {"StringLike": {"iam:AWSServiceName": "eks-fargate.amazonaws.com"}}}, {"Effect": "Allow", "Action": ["iam:GetRole", "iam:ListAttachedRolePolicies"], "Resource": ["arn:*:iam:*:*:role/*"]}, {"Effect": "Allow", "Action": ["iam:GetPolicy"], "Resource": ["arn:aws:iam:*:aws:policy/AmazonEKSClusterPolicy"]}, {"Effect": "Allow", "Action": ["eks:DescribeCluster", "eks:ListClusters", "eks:CreateCluster", "eks:TagResource", "eks:UpdateClusterVersion", "eks:DeleteCluster", "eks:UpdateClusterConfig", "eks:UntagResource", "eks:UpdateNodegroupVersion", "eks:DescribeNodegroup", "eks:DeleteNodegroup", "eks:UpdateNodegroupConfig", "eks:CreateNodegroup", "eks:AssociateEncryptionConfig"], "Resource": ["arn:*:eks:*:*:cluster/*", "arn:*:eks:*:*:nodegroup/*/*/*"]}, {"Effect": "Allow", "Action": ["eks:ListAddons", "eks:CreateAddon", "eks:DescribeAddonVersions", "eks:DescribeAddon", "eks:DeleteAddon", "eks:UpdateAddon", "eks:TagResource", "eks:DescribeFargateProfile", "eks:CreateFargateProfile", "eks:DeleteFargateProfile"], "Resource": ["*"]}, {"Effect": "Allow", "Action": ["iam:PassRole"], "Resource": ["*"], "Condition": {"StringEquals": {"iam:PassedToService": "eks.amazonaws.com"}}, {"Effect": "Allow", "Action": ["kms:CreateGrant", "kms:DescribeKey"], "Resource": ["*"], "Condition": {"ForAnyValue:StringLike": {"kms:ResourceAliases": "alias/cluster-api-provider-aws-*"}}}]'

```

```

aws iam attach-user-policy --user-name Kommander --policy-arn $(aws iam list-policies --query 'Policies[?PolicyName=`kommander-policy`].Arn' | grep -o '".*"' | tr -d '|')

```

```
aws iam create-access-key --user-name Kommander
```

6.1.3.4.1.1 Using an Existing User

You can use an existing AWS user with [credentials configured](#)⁴⁶⁸. The user must be authorized to create the following resources in the AWS account:

- EC2 Instances
- VPC
- Subnets
- Elastic Load Balancer (ELB)
- Internet Gateway
- NAT Gateway
- Elastic Block Storage (EBS) Volumes
- Security Groups
- Route Tables
- IAM Roles

Minimum IAM Policy Example

The following is the minimal IAM policy required:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": [
        "ec2:AllocateAddress",
        "ec2:AssociateRouteTable",
        "ec2:AttachInternetGateway",
        "ec2:AuthorizeSecurityGroupIngress",
        "ec2:CreateInternetGateway",
        "ec2:CreateNatGateway",
        "ec2:CreateRoute",
        "ec2:CreateRouteTable",
        "ec2:CreateSecurityGroup",
        "ec2:CreateSubnet",
        "ec2:CreateTags",
        "ec2:CreateVpc",
        "ec2:ModifyVpcAttribute",
        "ec2>DeleteInternetGateway",
        "ec2>DeleteNatGateway",
        "ec2>DeleteRouteTable",
        "ec2>DeleteSecurityGroup",

```

⁴⁶⁸ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>


```

"ec2:DeleteSubnet",
"ec2:DeleteTags",
"ec2:DeleteVpc",
"ec2:DescribeAccountAttributes",
"ec2:DescribeAddresses",
"ec2:DescribeAvailabilityZones",
"ec2:DescribeInstances",
"ec2:DescribeInternetGateways",
"ec2:DescribeImages",
"ec2:DescribeNatGateways",
"ec2:DescribeNetworkInterfaces",
"ec2:DescribeNetworkInterfaceAttribute",
"ec2:DescribeRouteTables",
"ec2:DescribeSecurityGroups",
"ec2:DescribeSubnets",
"ec2:DescribeVpcs",
"ec2:DescribeVpcAttribute",
"ec2:DescribeVolumes",
"ec2:DetachInternetGateway",
"ec2:DisassociateRouteTable",
"ec2:DisassociateAddress",
"ec2:ModifyInstanceAttribute",
"ec2:ModifyNetworkInterfaceAttribute",
"ec2:ModifySubnetAttribute",
"ec2:ReleaseAddress",
"ec2:RevokeSecurityGroupIngress",
"ec2:RunInstances",
"ec2:TerminateInstances",
"tag:GetResources",
"elasticloadbalancing:AddTags",
"elasticloadbalancing:CreateLoadBalancer",
"elasticloadbalancing:ConfigureHealthCheck",
"elasticloadbalancing>DeleteLoadBalancer",
"elasticloadbalancing:DescribeLoadBalancers",
"elasticloadbalancing:DescribeLoadBalancerAttributes",
"elasticloadbalancing:ApplySecurityGroupsToLoadBalancer",
"elasticloadbalancing:DescribeTags",
"elasticloadbalancing:ModifyLoadBalancerAttributes",
"elasticloadbalancing:RegisterInstancesWithLoadBalancer",
"elasticloadbalancing:DeregisterInstancesFromLoadBalancer",
"elasticloadbalancing:RemoveTags",
"autoscaling:DescribeAutoScalingGroups",
"autoscaling:DescribeInstanceRefreshes",
"ec2:CreateLaunchTemplate",
"ec2:CreateLaunchTemplateVersion",
"ec2:DescribeLaunchTemplates",
"ec2:DescribeLaunchTemplateVersions",
"ec2:DeleteLaunchTemplate",
"ec2:DeleteLaunchTemplateVersions",
"ec2:DescribeKeyPairs"
],
"Resource": ["*"]

```

```

    },
    {
      "Effect": "Allow",
      "Action": [
        "autoscaling:CreateAutoScalingGroup",
        "autoscaling:UpdateAutoScalingGroup",
        "autoscaling:CreateOrUpdateTags",
        "autoscaling:StartInstanceRefresh",
        "autoscaling>DeleteAutoScalingGroup",
        "autoscaling>DeleteTags"
      ],
      "Resource": [
        "arn:*:autoscaling:*:*:autoScalingGroup:*:autoScalingGroupName/*"
      ]
    },
    {
      "Effect": "Allow",
      "Action": ["iam:CreateServiceLinkedRole"],
      "Resource": [
        "arn:*:iam:*:role/aws-service-role/autoscaling.amazonaws.com/
AWSServiceRoleForAutoScaling"
      ],
      "Condition": {
        "StringLike": { "iam:AWSServiceName": "autoscaling.amazonaws.com" }
      }
    },
    {
      "Effect": "Allow",
      "Action": ["iam:CreateServiceLinkedRole"],
      "Resource": [
        "arn:*:iam:*:role/aws-service-role/elasticloadbalancing.amazonaws.com/
AWSServiceRoleForElasticLoadBalancing"
      ],
      "Condition": {
        "StringLike": {
          "iam:AWSServiceName": "elasticloadbalancing.amazonaws.com"
        }
      }
    },
    {
      "Effect": "Allow",
      "Action": ["iam:CreateServiceLinkedRole"],
      "Resource": [
        "arn:*:iam:*:role/aws-service-role/spot.amazonaws.com/
AWSServiceRoleForEC2Spot"
      ],
      "Condition": {
        "StringLike": { "iam:AWSServiceName": "spot.amazonaws.com" }
      }
    },
    {
      "Effect": "Allow",

```

```

    "Action": ["iam:PassRole"],
    "Resource": ["arn::*:iam::*:role/*cluster-api-provider-aws.sigs.k8s.io"]
  },
  {
    "Effect": "Allow",
    "Action": [
      "secretsmanager:CreateSecret",
      "secretsmanager>DeleteSecret",
      "secretsmanager:TagResource"
    ],
    "Resource": ["arn::*:secretsmanager::*:secret:aws.cluster.x-k8s.io/*"]
  },
  {
    "Effect": "Allow",
    "Action": ["ssm:GetParameter"],
    "Resource": ["arn::*:ssm::*:parameter/aws/service/eks/optimized-ami/*"]
  },
  {
    "Effect": "Allow",
    "Action": ["iam:CreateServiceLinkedRole"],
    "Resource": [
      "arn::*:iam::*:role/aws-service-role/eks.amazonaws.com/
AWSServiceRoleForAmazonEKS"
    ],
    "Condition": {
      "StringLike": { "iam:AWSServiceName": "eks.amazonaws.com" }
    }
  },
  {
    "Effect": "Allow",
    "Action": ["iam:CreateServiceLinkedRole"],
    "Resource": [
      "arn::*:iam::*:role/aws-service-role/eks-nodegroup.amazonaws.com/
AWSServiceRoleForAmazonEKSNodegroup"
    ],
    "Condition": {
      "StringLike": { "iam:AWSServiceName": "eks-nodegroup.amazonaws.com" }
    }
  },
  {
    "Effect": "Allow",
    "Action": ["iam:CreateServiceLinkedRole"],
    "Resource": [
      "arn:aws:iam::*:role/aws-service-role/eks-fargate-pods.amazonaws.com/
AWSServiceRoleForAmazonEKSFargate"
    ],
    "Condition": {
      "StringLike": { "iam:AWSServiceName": "eks-fargate.amazonaws.com" }
    }
  },
  {
    "Effect": "Allow",

```

```

    "Action": ["iam:GetRole", "iam:ListAttachedRolePolicies"],
    "Resource": ["arn:*:iam:*:role/*"]
  },
  {
    "Effect": "Allow",
    "Action": ["iam:GetPolicy"],
    "Resource": ["arn:aws:iam::aws:policy/AmazonEKSClusterPolicy"]
  },
  {
    "Effect": "Allow",
    "Action": [
      "eks:DescribeCluster",
      "eks:ListClusters",
      "eks:CreateCluster",
      "eks:TagResource",
      "eks:UpdateClusterVersion",
      "eks>DeleteCluster",
      "eks:UpdateClusterConfig",
      "eks:UntagResource",
      "eks:UpdateNodegroupVersion",
      "eks:DescribeNodegroup",
      "eks>DeleteNodegroup",
      "eks:UpdateNodegroupConfig",
      "eks:CreateNodegroup",
      "eks:AssociateEncryptionConfig"
    ],
    "Resource": ["arn:*:eks:*:cluster/*", "arn:*:eks:*:nodegroup/*/*/*"]
  },
  {
    "Effect": "Allow",
    "Action": [
      "eks:ListAddons",
      "eks:CreateAddon",
      "eks:DescribeAddonVersions",
      "eks:DescribeAddon",
      "eks>DeleteAddon",
      "eks:UpdateAddon",
      "eks:TagResource",
      "eks:DescribeFargateProfile",
      "eks:CreateFargateProfile",
      "eks>DeleteFargateProfile"
    ],
    "Resource": ["*"]
  },
  {
    "Effect": "Allow",
    "Action": ["iam:PassRole"],
    "Resource": ["*"],
    "Condition": {
      "StringEquals": { "iam:PassedToService": "eks.amazonaws.com" }
    }
  },

```

```

{
  "Effect": "Allow",
  "Action": ["kms:CreateGrant", "kms:DescribeKey"],
  "Resource": ["*"],
  "Condition": {
    "ForAnyValue:StringLike": {
      "kms:ResourceAliases": "alias/cluster-api-provider-aws-*"
    }
  }
}
]
}

```

6.1.3.4.1.2 Create an AWS Infrastructure Provider in DKP

Use the following steps to create an AWS infrastructure provider with static credentials as configured above:

1. In DKP, select the Workspace associated with the credentials you are adding.
2. Navigate to **Administration > Infrastructure Providers** and click the **Add Infrastructure Provider** button.
3. Select the Amazon Web Services (AWS) option.
4. Ensure **Static** is selected as the Authentication Method.
5. Enter a name for your infrastructure provider for later reference. Consider choosing a name that matches the AWS user.
6. Fill out the access ID and secret keys using the keys generated above.
7. Select **Save** to save your provider.

6.1.3.5 Create an Azure Infrastructure Provider in the DKP UI

Before you provision Azure clusters using the DKP UI, you must first create an Azure infrastructure provider to contain your Azure credentials:

1. Log in to the Azure command line:

```
az login
```

2. Create an Azure Service Principal (SP) by running the following command:

```
az ad sp create-for-rbac --role contributor --name "$(whoami)-konvoy" --scopes=/subscriptions/$(az account show --query id -o tsv)
```

3. Select **Infrastructure Providers** from the Dashboard menu.
4. Select **Add Infrastructure Provider**.

5. [Workspaces \(see page 650\)](#) If you are already in a workspace, the provider is automatically created in that workspace.
6. Select **Microsoft Azure**.
7. Add a **Name** for your Infrastructure Provider.
8. Copy and paste the following values into the indicated fields:
 - Copy the `id` output from the login command above and paste it into the **Subscription ID** field.
 - Copy the `tenant` used in Step 2 and paste it into the **Tenant ID** field.
 - Copy the `appId` used in Step 2 and paste it into the **Client ID** field.
 - Copy the `password` used in Step 2 and paste it into the **Client Secret** field.
9. Select **Save**.


6.1.3.6 Create a vSphere Infrastructure Provider in the DKP UI

Next, you must create a vSphere infrastructure provider:

1. Log in to your DKP Enterprise user interface and access the DKP home page.
2. Select **Infrastructure Providers** from the navigation menu on the left.
3. Select **Add Infrastructure Provider**.
4. [Choose a workspace. \(see page 650\)](#) If you are already in a workspace, the new infrastructure provider is automatically created in that workspace.
5. Select **vSphere** and add the following information:
 - Add a **Name** for your Infrastructure Provider.
 - Enter a valid vSphere vCenter user's name in the **Username** field.
 - Enter a valid vSphere vCenter user's password in the **Password** field.
 - Enter the vCenter Server URL in the **Host URL** field.
 - Enter a valid **TLS Certificate Thumbprint** value in the field.
The TLS Certificate Thumbprint helps to create a secure connection to VMware vCenter. If you do not have a thumbprint, your connection will be marked as insecure. This field is optional because you might not have a self-signed vCenter instance, and you only need the thumbprint if you do. The command to obtain this SHA-1 thumbprint for the vSphere's server's TLS certificate is listed under the field in the interface.
6. Select the **Save** button in the upper right corner of the page.

6.1.4 Add a Header, Footer, and Logo to Your DKP Implementation

Use the customizable Banners page to add header banners, footer banners, and select the colors for them. You can define header and/or footer banners for your DKP pages and turn them on and off, as needed.

 DKP displays your header and footer banner in a default typeface and size, which cannot be changed.

6.1.4.1 Creating a Header Banner

The text you type in the **Text** field appears centered at the top of the screen. The text length is limited to 100 characters, including spaces. The text color is determined by the background color and automatically calculates an appropriate light or dark color for you.

The **Color** selection control uses the style of your browser for its color picker tool. This control allows you to select a color for your header banner in one of three ways:

- Enter the color's Hex code.
- Select a general color range, and then select a specific shade or tint. The color input uses the style of your browser for its color selection tool.
- Select the eyedropper, move it to a sample of the color you want and select once to select that color's location.

6.1.4.2 Creating a Footer Banner

The text you type in the **Text** field appears centered at the bottom of the screen.

The **Color** selection control uses the style of your browser for its color picker tool. This control allows you to select a color for your footer banner in one of three ways:

- Enter the color's Hex code.
- Select a general color range from the slider bar, and then select a specific shade or tint with your mouse cursor.
- Select the eyedropper, move it to a sample of the color you want and select once to select that color's location.

6.1.4.3 Adding Your Organization's Logo to the Header

Enterprise


Gov Advanced

When you license and install DKP Enterprise or Gov Advanced, you also have the option to add your organization's logo to the header. The width of the header banner automatically adjusts to contain your logo. DKP automatically places your logo on the left side of the header, and centers it vertically.

Your logo graphic must meet the following criteria:


- Use a suggested file format - PNG, SVG, or JPEG.
- The file size cannot exceed 200 KB.

Error messages affecting the file to upload appear below the image in **red**, inside the shaded logo area.

-  To provide security against certain kinds of malicious activity, your browser has a “same-origin policy” for accessing resources. When you upload a file, the browser creates a unique identifier for the file. This will prevent you from selecting a file more than once.

6.1.4.3.1 Upload a logo image using drag-and-drop

1. Locate the desired file in the MacOS Finder or Windows File Explorer.
2. Drag and drop an image of the appropriate file type into the shaded area to see a preview of the image and display the filename.
You can select the **X** button (upper-right) or the **Remove** link (lower-right) to clear the image, if needed.
3. Select the **Save** button in the upper-right corner to retain your changes.

-  You will not be able to select a file for drag-and-drop if it does not have a valid image format.

6.1.4.3.2 Upload a logo image by browsing your files

1. Select the Browse Files button.
2. Navigate to, and select, the desired file to see a preview of the image and display the filename.
You can select the **X** button (upper-right) or the **Remove** link (lower-right) to clear the image, if needed.
3. Select the **Save** button in the upper-right corner to retain your changes.

6.2 Applications

This section includes information on the applications you can deploy in DKP.

- [AppDeployment resources](#) (see page 622)
- [Logging Stack Application Sizing Recommendations](#) (see page 625)
- [Rook Ceph Cluster Sizing Recommendations](#) (see page 629)
- [Manage an Application using the UI](#) (see page 633)
- [Platform Applications](#) (see page 639)
- [Setting Priority Classes in DKP Applications](#) (see page 647)

6.2.1 AppDeployment resources

Use AppDeployments to deploy, and customize platform, DKP catalog, and custom applications in your environment

An `AppDeployment` is a [Custom Resource](#)⁴⁶⁹ created by DKP with the purpose of deploying applications (platform, DKP catalog and custom applications) in the management cluster, managed clusters, or both. Customers of both Essential and Enterprise products use `AppDeployments`, regardless of their setup (air-gapped, non-air-gapped, etc.), and their infrastructure provider.

When installing DKP, an `AppDeployment` resource is created for each enabled **Platform Application**. This `AppDeployment` resource references a `ClusterApp`, which then references the repository that contains a concrete declarative and preconfigured setup of an application, usually in the form of a `HelmRelease`. `ClusterApps` are cluster-scoped so that these platform applications are deployable to all workspaces or projects.

In the case of **DKP catalog** and **custom applications**, the `AppDeployment` references an `App` instead of a `ClusterApp`, which also references the repository containing the installation and deployment information. `Apps` are namespace-scoped and are meant to only be deployable to the workspace or project in which they have been created.

For example, this is the default `AppDeployment` for the Kube Prometheus Stack platform application:

```
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: kube-prometheus-stack
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: kube-prometheus-stack-46.8.1
    kind: ClusterApp
```

6.2.1.1 Customization

6.2.1.2 Prerequisites

Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace where the cluster is attached:

⁴⁶⁹ <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

```
export WORKSPACE_NAMESPACE=<your_workspace_namespace>
```

You are now able to copy the following commands without having to replace the placeholder with your workspace namespace every time you run a command.

6.2.1.3 Customize Your Application

If you want to customize an application, or change how a specific app is deployed, you can create a `ConfigMap` to change or add values to the information that is stored in the `HelmRelease` . Override the default configuration of an application by setting the `configOverrides` field on the `AppDeployment` to that `ConfigMap` . This overrides the configuration of the app for all clusters within the workspace.

For workspace applications, you can also enable and customize them on a per-cluster basis. Refer to the [cluster-scoped configuration \(see page 654\)](#) page for instructions on how to enable and customize an application per cluster in a given workspace.

This is an example, of how to customize the `AppDeployment` of Kube Prometheus Stack:

1. Provide the name of a `ConfigMap` with the custom configuration in the `AppDeployment` :

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: kube-prometheus-stack
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: kube-prometheus-stack-46.8.1
    kind: ClusterApp
  configOverrides:
    name: kube-prometheus-stack-overrides-attached
EOF
```

2. Create the `ConfigMap` with the name provided in the previous step, which provides the custom configuration on top of the default configuration:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: kube-prometheus-stack-overrides-attached
data:
  values.yaml: |
    prometheus:
```

```

prometheusSpec:
  storageSpec:
    volumeClaimTemplate:
      spec:
        resources:
          requests:
            storage: 150Gi
EOF

```

6.2.1.4 Print and Review the Current State of an `AppDeployment` Resource

If you want to know how the `AppDeployment` resource is currently configured, use the commands below to print a table of the declared information. If the `AppDeployment` is configured for several clusters in a workspace, a column will display a list of the clusters.

6.2.1.4.1 Review all `AppDeployments` in a workspace

To review the state of the `AppDeployment` resource for a specific workspace, run the `get` command with the name of your workspace, as in this example:

```
dkp get appdeployments -w kommander-workspace
```

The output should contain a list of all your applications, here is an example:

NAME	APP	CLUSTERS
[...]		
kube-oidc-proxy		kube-oidc-proxy-0.3.2 host-cluster
kube-prometheus-stack		kube-prometheus-stack-46.8.1 host-cluster
kubecost		kubecost-0.35.1 host-cluster
[...]		

6.2.1.4.2 Review a specific `AppDeployment` of an application in a workspace

To review the state of a specific `AppDeployment` of an application, run the `get` command with the name of the application and your workspace, as in this example:

```
dkp get appdeployment kube-prometheus-stack -w kommander-workspace
```

The output should look similar to this:

```
NAME APP CLUSTERS
kube-prometheus-stack kube-prometheus-stack-46.8.1 host-cluster
```

6.2.1.5 Deployment Scope

In a single-cluster environment with an **Essential license**, `AppDeployments` enable customizing any platform application.

In a multi-cluster environment with an **Enterprise license**, `AppDeployments` enable [workspace-level](#) (see [page 105](#)), [project-level](#) (see [page 690](#)), and [per-cluster deployment and customization of workspace applications](#) (see [page 654](#)).

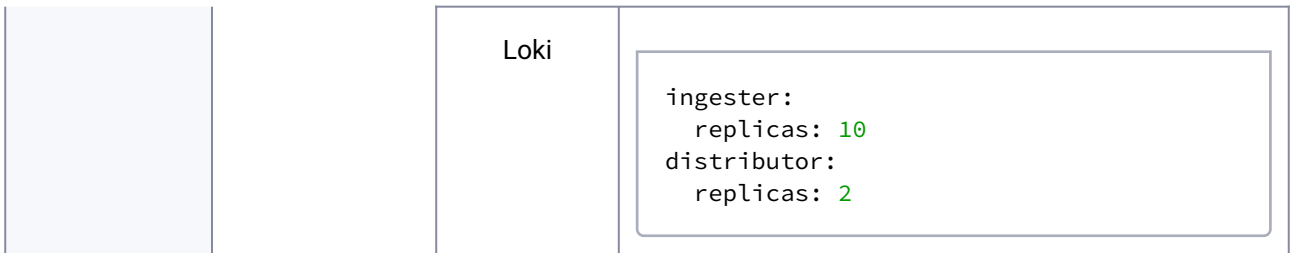
6.2.1.6 More Information

Refer to the [CLI documentation](#) (see [page 1644](#)) for more information on how to `create`, or `get` an `AppDeployment`.

6.2.2 Logging Stack Application Sizing Recommendations

# of worker nodes	Log Generating Load	Application	Suggested Configuration

50	1.4 MB/s	Logging Operator Logging Override Config	<pre> values.yaml: - clusterOutputs: - name: loki spec: loki: # change \$ {WORKSPACE_NAMESPACE} to the actual value of your workspace namespace url: http://grafana-loki-loki- distributed-gateway.\$ {WORKSPACE_NAMESPACE}.svc.cluster.local: 80 extract_kubernetes_labels: true configure_kubernetes_labels: true buffer: disabled: true retry_forever: false retry_max_times: 5 flush_mode: interval flush_interval: 10s flush_thread_count: 8 extra_labels: log_source: kubernetes_container fluentbit: inputTail: Mem_Buf_Limit: 512MB fluentd: bufferStorageVolume: emptyDir: medium: Memory disablePvc: true scaling: replicas: 10 resources: requests: memory: 1000Mi cpu: 1000m limits: memory: 2000Mi cpu: 1000m </pre>
----	----------	--	---



100	8.5 MB/s	Logging Operator Logging Override Config	<pre> values.yaml: - clusterOutputs: - name: loki spec: loki: # change \$ {WORKSPACE_NAMESPACE} to the actual value of your workspace namespace url: http://grafana-loki-loki- distributed-gateway.\$ {WORKSPACE_NAMESPACE}.svc.cluster.local: 80 extract_kubernetes_labels: true configure_kubernetes_labels: true buffer: disabled: true retry_forever: false retry_max_times: 5 flush_mode: interval flush_interval: 10s flush_thread_count: 8 extra_labels: log_source: kubernetes_container fluentbit: inputTail: Mem_Buf_Limit: 512MB fluentd: bufferStorageVolume: emptyDir: medium: Memory disablePvc: true scaling: replicas: 15 resources: requests: memory: 1000Mi cpu: 1000m limits: memory: 2000Mi cpu: 1000m </pre>
-----	----------	--	---

	<p>Loki Override Config</p>	<pre> ingester: replicas: 12 distributor: replicas: 3 loki: structuredConfig: limits_config: ingestion_rate_mb: 1024 ingestion_burst_size_mb: 1024 per_stream_rate_limit: 30MB per_stream_rate_limit_burst: 60MB max_streams_per_user: 0 max_global_streams_per_user: 0 max_label_names_per_series: 40 gateway: nginxConfig: httpSnippet: - client_max_body_size 50M; serverSnippet: - client_max_body_size 50M; </pre>
--	-------------------------------------	---


i Refer to [AppDeployment resources](#) (see page 625) for information on how you customize your AppDeployments.

= When configuring storage for `logging-operator-logging-overrides`, ensure that you create a ConfigMap in your workspace namespace for every cluster in that workspace. Keep in mind that `logging-operator-logging-overrides` can only be configured via the CLI.

6.2.3 Rook Ceph Cluster Sizing Recommendations

# of worker nodes	Application	Suggested Configuration
-------------------	-------------	-------------------------

50	Rook Ceph Cluster	<pre> cephClusterSpec: labels: monitoring: prometheus.kommander.d2iq.io/ select: "true" storage: storageClassDeviceSets: - name: rook-ceph-osd- set1 count: 4 portable: true encrypted: false placement: topologySpreadConstraints: - maxSkew: 1 topologyKey: topology.kubernetes.io/zone # The nodes in the same rack have the same topology.kubernetes.io/zone label. whenUnsatisfiable: ScheduleAnyway labelSelector: matchExpressions: - key: app operator: In values: - rook- ceph-osd - rook- ceph-osd-prepare - maxSkew: 1 topologyKey: kubernetes.io/hostname whenUnsatisfiable: ScheduleAnyway labelSelector: matchExpressions: - key: app operator: In values: - rook- ceph-osd - rook- ceph-osd-prepare volumeClaimTemplates: </pre>
----	-------------------	---



```
# If there are some
faster devices and some slower
devices, it is more efficient
to use
# separate metadata,
wal, and data devices.
# Refer https://
rook.io/docs/rook/v1.10/CRDs/
Cluster/pvc-cluster/#dedicated-
metadata-and-wal-device-for-
osd-on-pvc
- metadata:
  name: data
  spec:
    resources:
      requests:
        storage:
120Gi
    volumeMode: Block
    accessModes:
      - ReadWriteOnce
```

100

```

dkp:
  grafana-loki:
    additionalConfig:
      maxSize: "1000G"

cephClusterSpec:
  labels:
    monitoring:

prometheus.kommander.d2iq.io/
select: "true"
storage:
  storageClassDeviceSets:
    - name: rook-ceph-osd-
      set1
        count: 8
        portable: true
        encrypted: false
        placement:

topologySpreadConstraints:
  - maxSkew: 1
    topologyKey:
topology.kubernetes.io/zone #
The nodes in the same rack have
the same
topology.kubernetes.io/zone
label.
    whenUnsatisfiable:
ScheduleAnyway
    labelSelector:
      matchExpressions:
        - key: app
          operator: In
          values:
            - rook-
ceph-osd
            - rook-
ceph-osd-prepare
  - maxSkew: 1
    topologyKey:
kubernetes.io/hostname
    whenUnsatisfiable:
ScheduleAnyway
    labelSelector:
      matchExpressions:
        - key: app
          operator: In
          values:

```

```

- rook-
ceph-osd
- rook-
ceph-osd-prepare
  volumeClaimTemplates:
    - metadata:
      name: data
      spec:
        resources:
          requests:
            storage:
200Gi
          volumeMode:
Block
        accessModes:
-
ReadWriteOnce

```

i Refer to [AppDeployment resources](#) (see page 625) for information on how you customize your AppDeployments.

f To add more storage to `rook-ceph-cluster`, copy and paste the `storageClassDeviceSets` list from the `rook-ceph-cluster-1.10.3-d2iq-defaults` ConfigMap into your workspace where `rook-ceph-cluster` is present and then modify `count` and `volumeClaimTemplates.spec.resource.requests.storage`.

6.2.4 Manage an Application using the UI

Choose your license type for instructions on how to enable, and customize an application and then verify it has been deployed correctly.

- [Enterprise - Manage an Application using the UI](#) (see page 634)
 - [Enable an Application using the UI](#) (see page 634)
 - [Customize an Application using the UI](#) (see page 635)
 - [Verify an Application using the UI](#) (see page 636)
- [Essential - Manage an Application using the UI](#) (see page 637)
 - [Enable an Application using the UI - Essential](#) (see page 637)
 - [Customize an Application using the UI - Essential](#) (see page 638)
 - [Verify an Application via UI - Essential](#) (see page 638)

6.2.4.1 Enterprise - Manage an Application using the UI

Enterprise

Gov Advanced



To use the CLI to deploy or uninstall applications, see [Application Deployment \(see page 639\)](#).

- [Enable an Application using the UI \(see page 634\)](#)
- [Customize an Application using the UI \(see page 635\)](#)
- [Verify an Application using the UI \(see page 636\)](#)

6.2.4.1.1 Enable an Application using the UI

Enterprise

Gov Advanced

Follow these steps to enable your platform applications from the UI:

1. From the top menu bar, select your target workspace.
2. Select **Applications** from the sidebar to browse through the available applications from your configured repositories.
3. Select the three-dot button of the desired application card > **Enable**.
4. If available, select a version from the drop-down menu. This drop-down menu is only visible if there is more than one version to choose from.
5. Select the clusters where you want to deploy the application.
6. **For customizations only:** to override the default configuration values, select **Configuration** in the sidebar.

Note: If there are customization *Overrides* at the workspace and cluster level, they are combined for implementation. *Cluster-level Overrides* take precedence over *Workspace Overrides*.


- a. To customize an application for all clusters in a workspace, copy your customized values into the text editor under **Workspace Application Configuration** or upload your YAML file that contains the values:

```
someField: someValue
```

- b. To add a customization per cluster, copy the customized values into the text editor of each cluster under **Cluster Application Configuration Override** or upload your YAML file that contains the values:

```
someField: someValue
```


7. Confirm the details are correct, and then select the **Enable** button.

 There may be dependencies between the applications, which are listed in the [Workspace Platform Application Dependencies](#) (see page 641) documentation. Review them carefully prior to customizing to ensure that the applications are deployed successfully.

6.2.4.1.1.1 Next topic:

[Customize an App using the UI](#) (see page 635)

6.2.4.1.2 Customize an Application using the UI

 If you want to enable an application for the first time and customize it, refer to [Enable an Application using the UI](#) (see page 637).

You can customize the applications that are deployed to a workspace's cluster using the UI:

1. From the top menu bar, select your target workspace.
2. Select **Applications** from the sidebar menu to browse all applications.
3. In the **Application** card you want to customize, select the three dot menu and **Edit**.
4. To override the default configuration values, select **Configuration** in the sidebar.

Note: If there are customization *Overrides* at the workspace and cluster level, they are combined for implementation. *Cluster-level Overrides* take precedence over *Workspace Overrides*.

 - a. To customize an application for all clusters in a workspace, copy your customized values into the text editor under **Workspace Application Configuration** or upload your YAML file that contains the values:

```
someField: someValue
```

- b. To add a customization per cluster, copy the customized values into the text editor of each cluster under **Cluster Application Configuration Override** or upload your YAML file that contains the values:

```
someField: someValue
```

5. Confirm the details are correct, and select **Save**.

6.2.4.1.2.1 Customize an App for a Specific Cluster from the **Clusters** page

You can also customize an application *for a specific cluster* from the **Clusters** view:

1. Select **Clusters** from the sidebar menu.
2. Selecting the target cluster.
3. Select the **Applications** tab.
4. Navigate to the target **Application** card.
5. Select the three-dot menu > **Edit**.

6.2.4.1.2.2 Next topic:

[Verify an App using the UI](#) (see page 636)

6.2.4.1.3 Verify an Application using the UI

Enterprise

Gov Advanced

The application has now been enabled. Follow these steps to ensure the application was deployed correctly:

1. Select your target workspace from the top menu bar.
2. Select the cluster you want to verify from the left sidebar menu:
 - a. Select **Management Cluster** if your target cluster is the Management Cluster Workspace.
 - b. Otherwise, select **Clusters**, and choose your target cluster.
3. Select the **Applications** tab and navigate to the application you want to verify.
4. If the application was deployed successfully, the status **Deployed** appears in the application card. Otherwise, hover over the failed status to obtain more information on why the application failed to deploy.



It can take several minutes for the application to deploy completely. If the **Deployed** or **Failed** status is not displayed, the deployment process is not finished.

6.2.4.2 Essential - Manage an Application using the UI



To use the CLI to deploy or uninstall applications, see [Application Deployment \(see page 639\)](#).

- [Enable an Application using the UI - Essential \(see page 637\)](#)
- [Customize an Application using the UI - Essential \(see page 638\)](#)
- [Verify an Application via UI - Essential \(see page 638\)](#)

6.2.4.2.1 Enable an Application using the UI - Essential

Follow these steps to enable your platform applications from the UI in Kommander:

1. Select **Applications** from the sidebar to browse through the available applications from your configured repositories.
2. Select the three-dot button of the desired application card > **Enable**.
3. If available, select a version from the drop-down menu. This drop-down menu is only visible if there is more than one version to choose from.
4. **For customizations only:**
 - To override the default configuration values, select **Configuration** in the sidebar.
 - Copy your customized values into the text editor under **Workspace Application Configuration** or upload your YAML file that contains the values:

```
someField: someValue
```

5. Confirm the details are correct, and then select the **Enable** button.



There may be dependencies between the applications, which are listed in the [Workspace Platform Application Dependencies \(see page 641\)](#) documentation. Review them carefully prior to customizing to ensure that the applications are deployed successfully.

6.2.4.2.1.1 Next topic:

[Customize an App using the UI - Essential \(see page 638\)](#)

6.2.4.2.2 Customize an Application using the UI - Essential



If you want to enable an application for the first time and customize it, refer to [Enable an Application using the UI - Essential](#) (see page 640).

You can customize the applications that are deployed to your Management Cluster using the UI:

1. Select **Applications** from the sidebar to browse all applications.
2. In the **Application** card you want to customize, select the three dot menu and **Edit**.
3. To override the default configuration values, select **Configuration** in the sidebar.
 - a. To customize an application, copy your customized values into the text editor under **Workspace Application Configuration** or upload your YAML file that contains the values:

```
someField: someValue
```

4. Confirm the details are correct, and select **Save**.

6.2.4.2.2.1 Next topic:

[Verify an App via UI - Essential](#) (see page 638)

6.2.4.2.3 Verify an Application via UI - Essential

The application has now been enabled. Follow these steps to ensure the application was deployed correctly:

1. Select **Management Cluster** from the sidebar.
2. Select the **Applications** tab and navigate to the application you want to verify.
3. If the application was deployed successfully, the status **Deployed** appears in the application card. Otherwise, hover over the failed status to obtain more information on why the application failed to deploy.



It can take several minutes for the application to deploy completely. If the **Deployed** or **Failed** status is not displayed, the deployment process is not finished.

6.2.5 Platform Applications

How platform applications work

When attaching a cluster, DKP deploys certain platform applications on the newly attached cluster. Operators can use the DKP UI to customize which platform applications to deploy to the attached clusters in a given workspace.

Currently, the monitoring stack is deployed by default. The logging stack is not.

Review the [Workspace Platform Application Defaults and Resource Requirements \(see page 105\)](#) to ensure that the attached clusters have sufficient resources.

When deploying and upgrading applications, platform applications come as a bundle; they are tested as a single unit, and you must deploy or upgrade them in a single process, for each workspace. This means all clusters in a workspace have the same set and versions of platform applications deployed.

6.2.5.1 Related Topics

- [Deploy Platform Applications via CLI \(see page 639\)](#)
- [Cluster-scoped Configuration for Existing AppDeployments \(see page 654\)](#)
- [Manage an Application using the UI \(see page 636\)](#)
- [Cluster-scoped Application Configuration from the DKP UI \(see page 652\)](#)
- [Platform Application Dependencies \(see page 641\)](#)
- [Workspace Platform Application Defaults and Resource Requirements \(see page 105\)](#)

6.2.5.2 Deploy Platform Applications via CLI

This topic describes how to use the CLI to enable an application to deploy to managed and attached clusters in a workspace.

6.2.5.2.1 Related Topics

- [Available Workspace Platform Applications \(see page 105\)](#)
- [Customize an existing AppDeployment \(see page 625\)](#)
- [Cluster-scoped configuration via CLI \(see page 654\)](#)
- [Cluster-scoped configuration via UI \(see page 652\)](#)
- [Enterprise Upgrade Platform Applications \(see page 1569\) OR Essential: Upgrade the Essential Cluster and Platform Applications \(see page 1590\)](#) depending on your license type.

6.2.5.2.2 Prerequisites

Before you begin, you must have:

- A running cluster with [Kommander installed \(see page 78\)](#).
- An [existing Kubernetes cluster attached to Kommander \(see page 752\)](#).

- Determine the name of the workspace where you wish to perform the deployments. You can use the `dkp get workspaces` command to see the list of workspace names and their corresponding namespaces.
- Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace where the cluster is attached:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

- Set the `WORKSPACE_NAME` environment variable to the name of the workspace where the cluster is attached:

```
export WORKSPACE_NAME=<workspace_name>
```



From the CLI, you can enable applications to deploy in the workspace. Verify that an application has successfully deployed [via the CLI \(see page 641\)](#).

6.2.5.2.3 Create the `AppDeployment` to Enable the Application

Enable a supported application to deploy to your existing attached or managed cluster with an `AppDeployment` resource (see page 625).

1. Obtain the **APP ID** and **Version** of the application from the [applications table in the latest Release Notes \(see page 1760\)](#). You will need them in the `<APP-ID>-<Version>` format, for example, `istio-1.17.2`.
2. Run the following command and define the `--app` flag to specify which platform application and version will be enabled:

```
dkp create appdeployment istio --app istio-1.17.2 --workspace ${WORKSPACE_NAME}
```



- The `--app` flag must match the `APP NAME` from the list of available platform applications.
- Observe that the `dkp create` command must be run with the `WORKSPACE_NAME` instead of the `WORKSPACE_NAMESPACE` flag.

This instructs Kommander to create and deploy the `AppDeployment` to the `KommanderClusters` in the specified `WORKSPACE_NAME`.

6.2.5.2.4 Verify Applications

The applications are now enabled. Connect to the attached cluster and watch the `HelmReleases` to verify the deployment. In this example, we are checking if `istio` got deployed correctly:

```
kubectl get helmreleases istio -n ${WORKSPACE_NAMESPACE} -w
```

You should eventually see the `HelmRelease` marked as `Ready`:

NAMESPACE	NAME	READY	STATUS	AGE
workspace-test-vjsfq	istio	True	Release reconciliation succeeded	7m3s



Some supported applications have dependencies on other applications. See [Workspace Platform Application Dependencies \(see page 641\)](#) for that table.

6.2.5.3 Platform Application Dependencies

6.2.5.3.1 Dependencies between workspace applications

Platform applications that are deployed to a workspace's attached clusters can depend on each other. It is important to note these dependencies when customizing the workspace platform applications to ensure that your applications are properly deployed to the clusters. For more information on how to customize workspace platform applications, see [Workspace Platform Applications \(see page 642\)](#).

When deploying or troubleshooting platform applications, it helps to understand how platform applications interact and may require other platform applications as dependencies.

If a platform application's dependency does not successfully deploy, the platform application requiring that dependency does not successfully deploy.

The following sections detail information about the workspace platform application.

6.2.5.3.2 Foundational Applications

Provides the foundation for all platform application capabilities and deployments on managed clusters. These applications must be enabled for any platform applications to work properly.

The foundational applications are comprised of the following platform application:

- [cert-manager](#)⁴⁷⁰: Automates TLS certificate management and issuance.
- [reloader](#)⁴⁷¹: A controller that watches changes on ConfigMaps and Secrets, and automatically triggers updates on the dependent applications.
- [traefik](#)⁴⁷²: Provides an HTTP reverse proxy and load balancer. Requires cert-manager and reloader.

Platform Application	Required Dependencies
cert-manager	-
reloader	-
traefik	cert-manager, reloader

6.2.5.3.3 Logging

Collects logs over time from Kubernetes and applications deployed on managed clusters. Also provides the ability to visualize and query the aggregated logs.

- [fluent-bit](#)⁴⁷³: Open source and multi-platform log processor tool which aims to be a generic Swiss knife for logs processing and distribution.
- [grafana-logging](#)⁴⁷⁴: Logging dashboard used to view logs aggregated to Grafana Loki.
- [grafana-loki](#)⁴⁷⁵: A horizontally-scalable, highly-available, multi-tenant log aggregation system inspired by Prometheus.
- [logging-operator](#)⁴⁷⁶: Automates the deployment and configuration of a Kubernetes logging pipeline.
- [rook-ceph](#)⁴⁷⁷ and [rook-ceph-cluster](#)⁴⁷⁸: A Kubernetes-native high performance object store with an S3-compatible API that supports deploying into private and public cloud infrastructures.

Platform Application	Required Dependencies	Optional Dependencies
fluent-bit	-	-
grafana-logging	grafana-loki	-

470 <https://cert-manager.io/docs>

471 <https://github.com/stakater/Reloader>

472 <https://traefik.io/>

473 <https://docs.fluentbit.io/manual/>


474 <https://grafana.com/oss/grafana/>

475 <https://grafana.com/oss/loki/>

476 <https://banzaicloud.com/docs/one-eye/logging-operator/>

477 <https://rook.io/docs/rook/v1.10/Helm-Charts/operator-chart/>

478 <https://rook.io/docs/rook/v1.10/Helm-Charts/ceph-cluster-chart/>

Platform Application	Required Dependencies	Optional Dependencies
grafana-loki	rook-ceph-cluster	-
logging-operator	-	-
rook-ceph	-	-
rook-ceph-cluster	rook-ceph	kube-prometheus-stack <div style="border: 1px solid purple; padding: 10px; margin-top: 10px;">  Users can override the configuration to remove the dependency, as needed. </div>

6.2.5.3.4 Monitoring

Provides monitoring capabilities by collecting metrics, including cost metrics, for Kubernetes and applications deployed on managed clusters. Also provides visualization of metrics and evaluates rule expressions to trigger alerts when specific conditions are observed.

- [kubecost](https://kubecost.com/)⁴⁷⁹: provides real-time cost visibility and insights for teams using Kubernetes, helping you continuously reduce your cloud costs.
- [kubernetes-dashboard](https://kubernetes.io/docs/tasks/access-application-cluster/web-ui-dashboard/)⁴⁸⁰: A general purpose, web-based UI for Kubernetes clusters. It allows users to manage applications running in the cluster, troubleshoot them and manage the cluster itself.
- [kube-prometheus-stack](https://github.com/prometheus-community/helm-charts/tree/main/charts/kube-prometheus-stack)⁴⁸¹: A stack of applications that collect metrics and provide visualization and alerting capabilities.
NOTE: [Prometheus](https://prometheus.io/)⁴⁸², [Prometheus Alertmanager](https://prometheus.io/docs/alerting/latest/alertmanager)⁴⁸³ and [Grafana](https://grafana.com/)⁴⁸⁴ are included in the bundled installation.
- [nvidia-gpu-operator](https://catalog.ngc.nvidia.com/orgs/nvidia/containers/gpu-operator)⁴⁸⁵: The NVIDIA GPU Operator manages NVIDIA GPU resources in a Kubernetes cluster and automates tasks related to bootstrapping GPU nodes.
- [prometheus-adapter](https://github.com/DirectXMan12/k8s-prometheus-adapter)⁴⁸⁶: Provides cluster metrics from Prometheus.

479 <https://kubecost.com/>

480 <https://kubernetes.io/docs/tasks/access-application-cluster/web-ui-dashboard/>

481 <https://github.com/prometheus-community/helm-charts/tree/main/charts/kube-prometheus-stack>

482 <https://prometheus.io/>

483 <https://prometheus.io/docs/alerting/latest/alertmanager>

484 <https://grafana.com/>

485 <https://catalog.ngc.nvidia.com/orgs/nvidia/containers/gpu-operator>

486 <https://github.com/DirectXMan12/k8s-prometheus-adapter>

Platform Application	Required Dependencies
kubecost	-
kubernetes-dashboard	-
kube-prometheus-stack	-
prometheus-adapter	kube-prometheus-stack
nvidia-gpu-operator	-

6.2.5.3.5 Security

Allows management of security constraints and capabilities for the clusters and users.

- [gatekeeper](#)⁴⁸⁷: A policy Controller for Kubernetes.

Platform Application	Required Dependencies
gatekeeper	-

6.2.5.3.6 Single Sign On (SSO)

Group of platform applications that allow enabling SSO on attached clusters. SSO is a centralized system for connecting attached clusters to the centralized authority on the management cluster.

- [kube-oidc-proxy](#)⁴⁸⁸: A reverse proxy server that authenticates users using OIDC to Kubernetes API servers where OIDC authentication is not available.
- [traefik-forward-auth](#)⁴⁸⁹: Installs a forward authentication application providing Google OAuth based authentication for Traefik.

Platform Application	Required Dependencies
kube-oidc-proxy	cert-manager, traefik

487 <https://github.com/open-policy-agent/gatekeeper>

488 <https://github.com/jetstack/kube-oidc-proxy>


489 <https://github.com/thomseddon/traefik-forward-auth>

Platform Application	Required Dependencies
traefik-forward-auth	traefik

6.2.5.3.7 Backup

This platform application assists you with backing up and restoring your environment.

- [velero](#)⁴⁹⁰: An open source tool for safely backing up and restoring resources in a Kubernetes cluster, performing disaster recovery, and migrating resources and persistent volumes to another Kubernetes cluster.

Platform Application	Dependencies
velero	rook-ceph-cluster <div style="border: 1px solid purple; padding: 10px; margin-top: 10px;">  This is an optional dependency. Users can override the config to remove the dependency, as needed. </div>

6.2.5.3.8 Service Mesh

Allows deploying service mesh on clusters, enabling the management of microservices in cloud-native applications. Service mesh can provide a number of benefits, such as providing observability into communications, providing secure connections, or automating retries and backoff for failed requests.

- [istio](#)⁴⁹¹: Addresses the challenges developers and operators face with a distributed or microservices architecture.
- [jaeger](#)⁴⁹²: A distributed tracing system used for monitoring and troubleshooting microservices-based distributed systems.
- [kiali](#)⁴⁹³: A management console for an Istio-based service mesh. It provides dashboards, observability, and lets you operate your mesh with robust configuration and validation capabilities.

⁴⁹⁰ <https://velero.io/>

⁴⁹¹ <https://istio.io/latest/about/service-mesh/>

⁴⁹² <https://www.jaegertracing.io/>

⁴⁹³ <https://kiali.io/>

Platform Application	Required Dependencies	Optional Dependencies
istio	kube-prometheus-stack	-
jaeger	istio	-
kiali	istio	jaeger (optional for monitoring purposes)
knative	istio	-

6.2.5.3.9 DKP Insights

Analyses your cluster's health, detects current and future anomalies, and creates alerts that offer root cause analysis and recommended solutions.

- dkp-insights: Also called **DKP Insights Engine**, collects events and metrics on your cluster and uses rule-based heuristics to detect potential anomalies of varying criticality.
- dkp-insights-management: Also called **DKP Insights Management**, collects anomaly alerts from one or several clusters and displays root cause analysis and recommended solutions in the UI.
-

Platform Application	Required Dependencies	Optional Dependencies
dkp-insights	Management/Essential Cluster <ul style="list-style-type: none"> • dkp-insights-management Workspace <ul style="list-style-type: none"> • rook-ceph • rook-ceph-cluster • kube-prometheus-stack 	-
dkp-insights-management	kubefed	-

6.2.5.3.10 Related Information

- [Kommander security architecture \(see page 934\)](#)

- [Traefik Ingress controller](#)⁴⁹⁴
- [Logging and audits](#) (see page 899)

6.2.5.4 Workspace Platform Application Resource Requirements

Refer to [Workspace Platform Application Defaults and Resource Requirements](#) (see page 105) for a list of all platform applications, their default deployment configuration, required resources, and storage minimums.

6.2.6 Setting Priority Classes in DKP Applications

In DKP, your workloads can be prioritized to ensure that your critical components stay running in any situation.

Priority Classes can be set for any application in DKP, including Platform Applications, Catalog applications, and even your own Custom Applications.



By default, the priority classes of Platform Applications are set by DKP.

See the following pages for more information about the default priority classes for DKP applications:

- [Workspace Platform Application Defaults and Resource Requirements](#) (see page 105)
- [Management Cluster Application Requirements](#) (see page 103)
- [Project Platform Application Configuration Requirements](#) (see page 694)

This page gives instructions on how you can override the default priority class of any application in DKP to a different one.

6.2.6.1 DKP Priority Classes

Here are the priority classes that are available in DKP:

Class	Value Name	Value	Description
DKP High	<code>dkp-high-priority</code>	100001000	This is the priority class that is used for high priority DKP workloads.

⁴⁹⁴ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273648782&linkCreation=true&spaceKey=DENT&title=%282.6%29+Authentication+and+authorization+architecture)

[fromPageId=273648782&linkCreation=true&spaceKey=DENT&title=%282.6%29+Authentication+and+authorization+architecture](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273648782&linkCreation=true&spaceKey=DENT&title=%282.6%29+Authentication+and+authorization+architecture)

DKP Critical	<code>dkp-critical-priority</code>	100002000	This is the highest priority class that is used for critical priority DKP workloads.
--------------	------------------------------------	-----------	--

6.2.6.2 Prerequisites

Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace where the cluster is attached:

```
export WORKSPACE_NAMESPACE=<your_workspace_namespace>
```

You are now able to copy the following commands without having to replace the placeholder with your workspace namespace every time you run a command.

6.2.6.3 Set the Priority Class

Follow these steps.

- Keep in mind that the overrides for each application appears differently and is dependent on how the application's helm chart values are configured.

For more information about the helm chart values used in the DKP Platform Applications, see [DKP 2.6.0 Components and Applications](#) (see page 1760).

Generally speaking, performing a search for the `priorityClassName` field allows you to find out how you can set the priority class for a component.

In the example below which uses the helm chart values in Grafana Loki, the referenced `priorityClassName` field is nested under the `ingester` component. That said, if you reference the [Grafana Loki helm chart values](#)⁴⁹⁵, you'll see that the priority class can be set for several other components, including `distributor`, `ruler`, and on a `global` level.

1. Create a `ConfigMap` with custom priority class configuration values for Grafana Loki. The following example sets the priority class of ingester component to the DKP critical priority class:

```
cat <<EOF | kubectl apply -f -
```

⁴⁹⁵ <https://github.com/grafana/helm-charts/blob/main/charts/loki-distributed/values.yaml>

```

apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: grafana-loki-overrides
data:
  values.yaml: |
    ingester:
      priorityClassName: dkp-critical-priority
EOF

```

2. Edit the `grafana-loki` `AppDeployment` to set the value of `spec.configOverrides.name` to `grafana-loki-overrides` (Refer to [Deploy a service with a custom configuration](#) (see page 0) for more information).

```

dkp edit appdeployment -n ${WORKSPACE_NAMESPACE} grafana-loki

```

After your editing is complete, the `AppDeployment` resembles this example:

```

apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: grafana-loki
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: grafana-loki-0.69.16
    kind: ClusterApp
  configOverrides:
    name: grafana-loki-overrides

```

3. It will take a few minutes to reconcile, but you can check the ingester pod's priority class after reconciling:

```

kubectl get pods -n ${WORKSPACE_NAMESPACE} -o custom-
columns=NAME:.metadata.name,PRIORITY:.spec.priorityClassName,PRIORITY:.spec.pri
ority |grep ingester

```

It will show something like this:

NAME	PRIORITY
grafana-loki-loki-distributed-ingester-0	dkp-critical-
priority	100002000

For more information about priority classes, see <https://kubernetes.io/docs/concepts/scheduling-eviction/pod-priority-preemption/>

6.3 Workspaces

Enterprise

Gov Advanced

Allow teams to manage their own clusters using workspaces

Workspaces are a logical grouping of clusters that maintain a similar configuration, with certain configurations automatically federated to those clusters. Workspaces give you the flexibility to represent your organization in a way that makes sense for your team and configuration. For example, you can create separate workspaces for each department, product, or business function. Each workspace manages its own clusters and role-based permissions, while retaining an overall organization-level view of all clusters in operation.

6.3.1 Global / Workspace UI

The UI is designed to be accessible for different roles at different levels:

- **Global:** At the top level, IT administrators manage all clusters across all workspaces.
- **Workspace:** DevOps administrators manage multiple clusters within a workspace.
- **Projects:** DevOps administrators or developers manage configuration and services across multiple clusters.

6.3.2 Default Workspace

To get started immediately, you can use the default workspace deployed in DKP. Your workspace delegation can be done at a later time.

6.3.3 Create a Workspace


In DKP you can create your own Workspaces. The following steps describe this procedure.

1. From the workspace selection dropdown in the top menu bar, select **Create Workspace**.
2. Type a name and description and select **Save**. The workspace is now accessible from the workspace selection drop down.

6.3.4 Add, Edit, and Delete Workspace Annotations and Labels


When creating or editing a workspace, you can use the **Advanced Options** to add, edit, or delete annotations and labels to your workspace. Both the annotations and labels are applied to the workspace namespace.

1. From the top menu bar, select your target workspace.
2. Select the **Actions** dropdown button in the top-right, and select **Edit Workspace**.
3. Type in new **Key** and **Value** labels for your workspace or edit existing **Key** and **Value** labels.

 Labels that are added to a workspace, are also applied to the `kommanderclusters` object as well as to all of the clusters in the workspace.

6.3.5 Delete a Workspace

In DKP you can delete existing Workspaces. The following steps describe this procedure.

 Workspaces can only be deleted if all the clusters in the workspace have been deleted or detached.

1. From the top menu bar, select **Global**.
2. From the sidebar menu, select **Workspaces**.
3. Select the three dot button to the right of the workspace you want to delete and then select **Delete**.
4. Confirm deleting the Workspace in the **Delete Workspace** dialog box.

The following procedures are supported for workspaces:

- [Application Deployment using the CLI \(see page 642\)](#)
- [Platform Applications \(see page 642\)](#)
- [Platform Application Dependencies \(see page 644\)](#)
- [Workspace Platform Application Defaults and Resource Requirements \(see page 105\)](#)

6.3.6 Workspace Applications

Enterprise

Gov Advanced

This section documents the applications and application types that you can use with DKP.

Application types are:

- [Catalog Applications \(see page 663\)](#) are either pre-packaged applications from the D2iQ Application Catalog, or custom applications that you maintain for your teams or organization.
- [Platform Applications \(see page 642\)](#) are applications integrated into DKP.

- [Cluster-scoped Application Configuration from the DKP UI](#) (see page 652)
- [Cluster-scoped Configuration for Existing AppDeployments](#) (see page 654)

6.3.6.1 Cluster-scoped Application Configuration from the DKP UI

Enterprise

Gov Advanced

When you enable an application for a workspace, you deploy this application to all clusters within that workspace. You can also choose to enable or customize an application on certain clusters within a workspace.

This functionality allows you to use DKP in a multi-cluster scenario without restricting the management of multiple clusters from a single workspace.



NOTE: DKP Essential users are only be able to configure and deploy applications to a single cluster within a workspace. Selecting an application to deploy to a cluster skips cluster selection and takes you directly to the workspace configuration overrides page.

6.3.6.1.1 Prerequisites

Ensure that you've provisioned or attached clusters in one of the following environments:

- [Amazon Web Services \(AWS\)](#) (see page 1101)
- [Amazon Elastic Kubernetes Service \(EKS\)](#) (see page 1144)
- [Microsoft Azure](#) (see page 746)

Refer to the current list of Catalog and Platform Applications:

- [Workspace Catalog Applications](#) (see page 663)
- [Workspace Platform Applications](#) (see page 105)

6.3.6.1.2 Enable a Cluster-scoped Application

Navigate to the workspace containing the clusters you want to deploy to by selecting the appropriate workspace name from the drop-down menu at the top of the DKP dashboard.

1. Select **Applications** from the left navigation bar and find the application you want to deploy to the cluster.
2. Select the the three-dot menu in the desired application's tile and select **Enable** to reach the application enablement page.

NOTE: The application enablement page can also be reached by selecting **View Details** from the three-dot menu, and then selecting **Enable** from the application's details page.

3. Select the cluster(s) that you're deploying the application to. The available clusters can be sorted by **Name, Type, Provider** and any **Labels** you've added.
4. Deploy the application to the clusters by selecting **Enable** in the top right corner of the application enablement page.

You are automatically redirected to either the **Applications** or **View Details** page. To view the application enabled in your chosen cluster, navigate to the **Clusters** page on the left navigation bar. The application appears in the **Applications** pane of the appropriate cluster.



NOTE: Once you enable an application at the workspace level, DKP automatically enables that app on any other cluster you create or attach.

6.3.6.1.3 Configure a Cluster-scoped Application

For scenarios where applications require different configurations on a per-cluster basis, navigate to the **Applications** page and select **Edit** from the three-dot menu of the appropriate application to return to the application enablement page.

1. Select the cluster(s) that you're deploying the application to. The available clusters can be sorted by **Name, Type, Provider** and any **Labels** you've added.
2. Select the **Configuration** tab.
3. The **Configuration** tab contains two separate types of code editors, where you can enter your specified overrides and configurations:
 - Workspace Application Configuration:** A workspace-level code editor that applies all configurations and overrides to the entirety of the workspace and its clusters for this application.
 - Cluster Application Configuration Override:** A cluster-scoped code editor that applies configurations and overrides to the cluster specified. These customizations will merge with the workspace application configuration. If there is no cluster-scoped configuration, the workspace configuration applies.
4. If you already have a configuration to apply in a text or .yaml file, you can upload the file by selecting **Upload File**. If you want to download the displayed set of configurations, select **Download File**.
5. Finish configuring the cluster-scoped applications by selecting **Save** in the top right corner of the application enablement page.

You are automatically redirected to either the **Applications** or **View Details** page. To view the custom configurations of the application in the cluster, select the **Configurations** tab on the details page of the application.

NOTE: Editing is disabled in the code boxes displayed in the application’s details page. To edit the configuration, select the **Edit** button in the top right of the page and repeat the steps in this section.

6.3.6.1.4 Remove a Cluster-scoped Application

Navigate to the cluster you’ve deployed your applications to by selecting **Clusters** from the left navigation bar.

1. Click on the Applications tab.
2. Select the three-dot menu in the desired application’s tile and select **Uninstall**.
3. A prompt appears to confirm your decision to uninstall the application. Follow the instructions in the prompt and select **Uninstall**.
4. Refresh the page to confirm that the application has been removed from the cluster.

This process only removes the application from the specific cluster you’ve navigated to. To remove this application from other clusters, navigate to the **Clusters** page and repeat the process.

6.3.6.2 Cluster-scoped Configuration for Existing AppDeployments

Enterprise

Gov Advanced

This topic describes how to enable cluster-scoped configuration of applications for **existing** AppDeployments . For more information on how to **create** AppDeployments , refer to the [Application Deployment \(see page 1659\)](#) section.

6.3.6.2.1 Enable or Disable an App per Cluster and Define a Custom Configuration

Edit the `AppDeployment` resource to modify the configuration of an application per cluster.

- [Prerequisites \(per-cluster application configuration\) \(see page 655\)](#)
- [Enable an Application per Cluster \(see page 655\)](#)
- [Customizations per Cluster \(see page 657\)](#)
- [Verify the Current Application Configuration \(see page 663\)](#)

6.3.6.2.2 For Better Understanding

When you enable an application for a workspace, you deploy this application to all clusters within that workspace. You can also choose to enable or customize an application on certain clusters within a workspace. This functionality allows you to use DKP in a multi-cluster scenario without restricting the management of multiple clusters from a single workspace.

Your DKP cluster comes bundled with a set of default application configurations. If you want to override the default configuration of your applications, you can define workspace `configOverrides` on top of the default workspace configuration. And if you want to further customize your workplace by enabling applications on a per-cluster basis or by defining per-cluster customizations, you can create and apply `clusterConfigOverrides`.

The cluster-scoped enablement and customization of applications is an **Enterprise-only** feature, which allows the configuration of all workspace [Platform](#) (see page 642), [DKP catalog](#) (see page 663) and [Custom Applications](#) (see page 676) via the **CLI** in your managed and attached clusters regardless of your environment setup (air-gapped or non-air-gapped). This capability is not provided for project applications.

6.3.6.2.3 Prerequisites (per-cluster application configuration)

Enterprise

Gov Advanced

- Any application you wish to enable or customize at a cluster level, first needs to be enabled at the workspace-level via an `AppDeployment` ([platform application deployment](#) (see page 642), [workspace catalog application deployment](#) (see page 671)).
- For custom configurations, you have [created a ConfigMap](#) (see page 0) with all the required `spec` fields for each customization you would like to add to an application in a cluster. You can apply a `ConfigMap` to several clusters, or create a `ConfigMap` for each cluster, but the `ConfigMap` object must exist in the Management cluster.
- Determine the name of the workspace where you wish to perform the deployments. You can use the `dkp get workspaces` command to see the list of workspace names and their corresponding namespaces.
- Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace where the cluster is attached:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

- Set the `WORKSPACE_NAME` environment variable to the name of the workspace where the cluster is attached:

```
export WORKSPACE_NAME=<workspace_name>
```

6.3.6.2.4 Enable an Application per Cluster

Enterprise

Gov Advanced

6.3.6.2.4.1 Prerequisites

Ensure you have reviewed the [prerequisites \(see page 658\)](#) section.

6.3.6.2.4.2 Enable an Application per Cluster *for the First Time*

When you enable an application on a workspace, it is deployed to all clusters in the workspace by default. If you would like to only deploy it to a subset of clusters when enabling it on a workspace for the first time, you can follow the steps below:

1. Create an `AppDeployment` for your application, selecting a subset of clusters within the workspace to enable it on. You can use the `dkp get clusters --workspace ${WORKSPACE_NAME}` command to see the list of clusters in the workspace.

i The following snippet is an example. Replace the application name, version, workspace name and cluster names according to your requirements. See [Components and Applications \(see page 1762\)](#) for compatible versions.

```
dkp create appdeployment kube-prometheus-stack --app kube-prometheus-
stack-46.8.1 --workspace ${WORKSPACE_NAME} --clusters attached-
cluster1,attached-cluster2
```

2. **Optional:** Check the [current status \(see page 663\)](#) of the `AppDeployment` to see the names of the clusters where the application is currently enabled.

6.3.6.2.4.3 Enable or Disable an Application per Cluster *after it has been enabled at the workspace level*

You can enable or disable applications at any time. After you have [enabled the application at the workspace level \(see page 0\)](#), the `spec.clusterSelector` [field \(see page 1605\)](#) populates.

i For clusters that are newly attached into the workspace, all applications enabled for the workspace are automatically enabled on and deployed to the new clusters.

If you want to see on what clusters your application is currently deployed, refer to the [Print and Review the Current State of your AppDeployment \(see page 627\)](#) documentation.

1. Edit the `AppDeployment` YAML by adding or removing the names of the clusters where you want to enable your application in the `clusterSelector` section:

i The following snippet is an example. Replace the application name, version, workspace name and cluster names according to your requirements. See [Components and Applications \(see page 1762\)](#) for the compatible versions.

```

cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: kube-prometheus-stack
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: kube-prometheus-stack-46.8.1
    kind: ClusterApp
  clusterSelector:
    matchExpressions:
      - key: kommander.d2iq.io/cluster-name
        operator: In
        values:
          - attached-cluster1
          - attached-cluster3-new
EOF

```

6.3.6.2.4.4 Verify the Current Configuration of your Application

Refer to the [Verify applications help](#) (see page 0) to connect to the managed or attached cluster and check the status of the deployments.

If you want to know how the `AppDeployment` resource is currently configured, refer to the [Print and review the state of your AppDeployment](#) (see page 627) section.

6.3.6.2.5 Customizations per Cluster

Enable or disable a customization per cluster.

- [Customize an Application per Cluster](#) (see page 657)
- [Disable a Custom Configuration of an Application for a Cluster](#) (see page 660)

6.3.6.2.5.1 Customize an Application per Cluster

Enterprise

Gov Advanced

Prerequisites

Ensure you have reviewed the [prerequisites](#) (see page 658) section.

Enable a Custom Configuration of an Application for a Cluster

You can customize the application for each cluster occurrence of said application. If you want to customize the application for a cluster that is not yet attached, refer to the [instructions below](#) (see page 660), so the application is deployed with the custom configuration on attachment.

To enable per-cluster customizations:

1. Reference the name of the `ConfigMap` to be applied per cluster in the `spec.clusterConfigOverrides` fields. In this example, you have three different customizations specified in three different `ConfigMaps` for three different clusters in one workspace:

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: kube-prometheus-stack
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: kube-prometheus-stack-46.8.0
    kind: ClusterApp
  clusterSelector:
    matchExpressions:
      - key: kommander.d2iq.io/cluster-name
        operator: In
        values:
          - attached-cluster1
          - attached-cluster2
          - attached-cluster3-new
  clusterConfigOverrides:
    - configMapName: kps-cluster1-overrides
      clusterSelector:
        matchExpressions:
          - key: kommander.d2iq.io/cluster-name
            operator: In
            values:
              - attached-cluster1
    - configMapName: kps-cluster2-overrides
      clusterSelector:
        matchExpressions:
          - key: kommander.d2iq.io/cluster-name
            operator: In
            values:
              - attached-cluster2
    - configMapName: kps-cluster3-overrides
      clusterSelector:
        matchExpressions:
          - key: kommander.d2iq.io/cluster-name
            operator: In
            values:
              - attached-cluster3-new
EOF
```

2. If you have not done so yet, [create the ConfigMaps](#) (see page 0) referenced in each `clusterConfigOverrides` entry.



- The changes are applied only if the YAML file has a valid syntax.
- Set up only one cluster override `ConfigMap` per cluster. If there are several `ConfigMaps` configured for a cluster, only one will be applied.
- Cluster override `ConfigMaps` must be created on the Management cluster.

Enable a Custom Configuration of an Application for a Cluster at Attachment

You can customize the application configuration for a cluster prior to its attachment, so that the application is deployed with this custom configuration on attachment. This is preferable, if you do not want to redeploy the application with an updated configuration after it has been initially installed, which may cause downtime.

To enable per-cluster customizations, follow these steps **before attaching the cluster**:

1. Set the `CLUSTER_NAME` environment variable to the cluster name that you will give your to-be-attached cluster:

```
export CLUSTER_NAME=<your_attached_cluster_name>
```

2. Reference the name of the `ConfigMap` you want to apply to this cluster in the `spec.clusterConfigOverrides` fields. **You do not need to update the `spec.clusterSelector` field.** In this example, you have the `kps-cluster1-overrides` customization specified for `attached-cluster-1` and a different customization (in `kps-your-attached-cluster-overrides` `ConfigMap`) for your to-be-attached cluster:

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: kube-prometheus-stack
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: kube-prometheus-stack-46.8.0
    kind: ClusterApp
  clusterSelector:
    matchExpressions:
      - key: kommander.d2iq.io/cluster-name
```

```

operator: In
values:
- attached-cluster1
clusterConfigOverrides:
- configMapName: kps-cluster1-overrides
  clusterSelector:
    matchExpressions:
    - key: kommander.d2iq.io/cluster-name
      operator: In
      values:
      - attached-cluster1
- configMapName: kps-your-attached-cluster-overrides
  clusterSelector:
    matchExpressions:
    - key: kommander.d2iq.io/cluster-name
      operator: In
      values:
      - ${CLUSTER_NAME}
EOF

```

3. If you have not done so yet, [create the ConfigMap](#) (see page 625) referenced for your to-be-attached cluster.



- The changes are applied only if the YAML file has a valid syntax.
- Cluster override `ConfigMaps` must be created on the Management cluster.

Verify the Current Configuration of your Application

Refer to the [Verify applications help](#) (see page 0) to connect to the managed or attached cluster and check the status of the deployments.

If you want to know how the `AppDeployment` resource is currently configured, refer to the [Print and review the state of your AppDeployment](#) (see page 627) section.

6.3.6.2.5.2 Disable a Custom Configuration of an Application for a Cluster

Enabled customizations are defined in a `ConfigMap`, which, in turn, is referenced in the `spec.clusterConfigOverrides` object of your `AppDeployment`.

1. Review your current configuration to establish what you want to remove:

```
kubectl get appdeployment -n ${WORKSPACE_NAMESPACE} kube-prometheus-stack -o yamll
```

The output looks similar to this:

```

apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: kube-prometheus-stack
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: kube-prometheus-stack-46.8.1
    kind: ClusterApp
  configOverrides:
    name: kube-prometheus-stack-overrides-attached
  clusterSelector:
    matchExpressions:
      - key: kommander.d2iq.io/cluster-name
        operator: In
        values:
          - attached-cluster1
          - attached-cluster2
  clusterConfigOverrides:
    - configMapName: kps-cluster1-overrides
      clusterSelector:
        matchExpressions:
          - key: kommander.d2iq.io/cluster-name
            operator: In
            values:
              - attached-cluster1
    - configMapName: kps-cluster2-overrides
      clusterSelector:
        matchExpressions:
          - key: kommander.d2iq.io/cluster-name
            operator: In
            values:
              - attached-cluster2

```

Here you can see that `kube-prometheus-stack` has been enabled for the `attached-cluster1` and `attached-cluster2`. There is also a custom configuration for each of the clusters: `kps-cluster1-overrides` and `kps-cluster2-overrides`.

2. Edit the file by deleting the `configMapName` entry of the cluster for which you would like to delete the customization. This is located under the `clusterConfigOverrides`:

```

cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: kube-prometheus-stack

```



```

namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    kind: ClusterApp
    name: kube-prometheus-stack-46.8.1
  configOverrides:
    name: kube-prometheus-stack-ws-overrides
  clusterSelector:
    matchExpressions:
      - key: kommander.d2iq.io/cluster-name
        operator: In
        values:
          - attached-cluster1
  clusterConfigOverrides:
    - configMapName: kps-cluster1-overrides
      clusterSelector:
        matchExpressions:
          - key: kommander.d2iq.io/cluster-name
            operator: In
            values:
              - attached-cluster1
EOF

```

Compare steps one and two for a reference of how an entry should be deleted.

3. Before deleting a `ConfigMap` that contains your customization, **ensure you will NOT require it at a later time**. It is not possible to restore a deleted `ConfigMap`.

If you choose to delete it, run:

```
kubectl delete configmap <name_configmap> -n ${WORKSPACE_NAMESPACE}
```

- It is **NOT** possible to delete a `ConfigMap` that is being actively used and referenced in the `configOverride` of any `AppDeployment`.

Verify the Current Configuration of your Application

Refer to the [Verify applications help \(see page 0\)](#) to connect to the managed or attached cluster and check the status of the deployments.

If you want to know how the `AppDeployment` resource is currently configured, refer to the [Print and review the state of your AppDeployment \(see page 627\)](#) section.

6.3.6.2.6 Verify the Current Application Configuration

Enterprise

Gov Advanced

6.3.6.2.6.1 Verify the Current Configuration of your Application

Refer to the [Verify applications help](#) (see page 0) to connect to the managed or attached cluster and check the status of the deployments.

If you want to know how the `AppDeployment` resource is currently configured, refer to the [Print and review the state of your AppDeployment](#) (see page 627) section.

6.3.7 Workspace Catalog Applications

Enterprise

Gov Advanced

Catalog applications are any third-party or open source applications that appear in the Catalog. These applications are deployed to be used for customer workloads.

D2iQ provides [DKP Catalog Applications](#) (see page 663) for use in your environment.

- [Workspace DKP Catalog Applications](#) (see page 663)
- [Deployment of Catalog Applications in Workspaces](#) (see page 671)
- [Workspace Catalog Application Upgrades](#) (see page 674)
- [Custom Applications](#) (see page 676)

6.3.7.1 Workspace DKP Catalog Applications

Enterprise

Gov Advanced

Catalog applications are applications provided by D2iQ for use in your environment.

6.3.7.1.1 Prerequisites

- Ensure your clusters run on a supported Kubernetes version for [this release of DKP](#) (see page 1758), and that said Kubernetes version is also compatible with your [catalog application version](#) (see page 664).
- For customers with an [Enterprise license](#) (see page 129) and a multi-cluster environment, we recommend keeping all clusters on the same Kubernetes version. This ensures your DKP catalog application can run on all clusters in a given workspace.

6.3.7.1.1.1 Install the DKP Catalog via the CLI

Follow these steps to install the DKP catalog from the CLI.

1. Refer to [Install Kommander in an Air-gapped Environment \(see page 1397\)](#), if you are running in air-gapped environment.
2. Set the `WORKSPACE_NAMESPACE` environment variable to the name of your workspace's namespace:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```

3. Create the `GitRepository`:

```
kubectl apply -f - <<EOF
apiVersion: source.toolkit.fluxcd.io/v1beta1
kind: GitRepository
metadata:
  name: dkp-catalog-applications
  namespace: ${WORKSPACE_NAMESPACE}
  labels:
    kommander.d2iq.io/gitapps-gitrepository-type: catalog
    kommander.d2iq.io/gitrepository-type: catalog
spec:
  interval: 1m0s
  ref:
    tag: v2.6.2
  timeout: 20s
  url: https://github.com/mesosphere/dkp-catalog-applications
EOF
```

4. Verify that you can see the DKP workspace catalog `Apps` available in the UI (in the Applications section in said workspace), and in the CLI, using `kubectl`:

```
kubectl get apps -n ${WORKSPACE_NAMESPACE}
```

6.3.7.1.1.2 Compatibility

Ensure that your DKP [Catalog application \(see page 0\)](#) is compatible with:

- The Kubernetes version in all the Managed and Attached clusters of the workspace where you want to install the catalog application.
- The range of Kubernetes versions [supported in this release of DKP \(see page 1758\)](#).

If your current Catalog application version is not compatible, upgrade the application to a compatible version.



From DKP 2.6, the following versions of Catalog applications are deprecated:

- kafka-operator-0.20.0
- kafka-operator-0.20.2
- spark-operator-1.1.6
- zookeeper-operator-0.2.13

If you plan on upgrading to DKP 2.6 or above, ensure that you upgrade these applications to the latest compatible version.

- To find what versions of applications are available for upgrade, use `kubectl` :

```
kubectl get apps -n ${WORKSPACE_NAMESPACE}
```

Name	App ID	Compatible Kubernetes versions
kafka-operator-0.20.0	kafka-operator	1.21-1.24
kafka-operator-0.20.2	kafka-operator	1.21 - 1.24
kafka-operator-0.23.0-dev.0	kafka-operator	1.21 - 1.26
kafka-operator-0.25.1	kafka-operator	1.21 - 1.26
spark-operator-1.1.6	spark-operator	1.21
spark operator-1.1.17	spark-operator	1.21 - 1.26
zookeeper-operator-0.2.13	zookeeper-operator	1.21 - 1.24
zookeeper-operator-0.2.14	zookeeper-operator	1.21 - 1.26
zookeeper-operator-0.2.15	zookeeper-operator	1.21 - 1.26

6.3.7.1.2 Kafka in a Workspace

Enterprise


Gov Advanced

Information about the Kafka Operator

[Apache Kafka](#)⁴⁹⁶ is an open-source distributed event streaming platform used for high-performance data pipelines, streaming analytics, data integration, and mission-critical applications. The [Kafka Operator](#)⁴⁹⁷ is a Kubernetes operator to automate provisioning, management, autoscaling, and operations of Apache Kafka clusters deployed to Kubernetes. It works by watching custom resources, such as `KafkaClusters`, `KafkaUsers`, and `KafkaTopics`, to provision underlying Kubernetes resources (i.e. `StatefulSets`) required for a production-ready Kafka Cluster.

6.3.7.1.2.1 Install

Follow these steps to install the Kafka operator in a workspace. This procedure results in a Kafka operator running in a workspace namespace, ready to manage and create Kafka clusters in any project namespaces. See the [Kafka in a Project](#) (see page 698) custom resource documentation for more information on creating Kafka clusters.

 Only install the Kafka operator once per workspace.

1. Follow the generic installation instructions for workspace catalog applications on the [Application Deployment](#) (see page 671) page.
2. Within the `AppDeployment`, update the `appRef` to specify the correct `kafka-operator` App. You can find the `appRef.name` by listing the available Apps in the workspace namespace:

```
kubectl get apps -n ${WORKSPACE_NAMESPACE}
```

Refer to the [Kafka operator Helm Chart documentation](#)⁴⁹⁸ for details on custom configuration for the operator.

6.3.7.1.2.2 Uninstall via the CLI

Uninstalling the Kafka operator does not affect existing `KafkaCluster` deployments. After uninstalling the operator, you must manually remove any remaining Custom Resource Definitions (CRDs) from the operator.

1. Delete all of the deployed Kafka custom resources (see [Deleting Kafka Custom Resources](#) (see page 0)).
2. Uninstall a Kafka operator `AppDeployment` :

⁴⁹⁶ <https://kafka.apache.org/>

⁴⁹⁷ <https://banzaicloud.com/docs/supertubes/kafka-operator/>

⁴⁹⁸ <https://github.com/banzaicloud/koperator/tree/master/charts/kafka-operator#configuration>

```
kubectl -n <workspace namespace> delete AppDeployment <name of AppDeployment>
```

3. Remove Kafka CRDs:

NOTE: The CRDs are not finalized for deletion until you delete the associated custom resources.

```
kubectl delete crds kafkaclusters.kafka.banzaicloud.io
kafkausers.kafka.banzaicloud.io kafkatopics.kafka.banzaicloud.io
```

6.3.7.1.2.3 Resources

- [Kafka Operator Documentation](#)⁴⁹⁹
- [Kafka Operator GitHub Repository](#)⁵⁰⁰
- [Sample Kafka Operator Custom Resources](#)⁵⁰¹
- [Apache Kafka Documentation](#)⁵⁰²

6.3.7.1.3 Spark Operator in a Workspace

Enterprise

Gov Advanced

How to spin up your Spark Operator

The Kubernetes Operator for Apache Spark aims to make specifying and running Spark applications as easy and idiomatic as running other workloads on Kubernetes. It uses Kubernetes custom resources for specifying, running, and surfacing status of Spark applications. For a complete reference of the custom resource definitions, refer to the API Definition. For details on its design, refer to the design documentation. It requires Spark 2.3 and above that supports Kubernetes as a native scheduler backend.



The default installation is basic, provide your override configmap to enable desired Spark Operator features.

6.3.7.1.3.1 Install Spark Operator


You can find generic installation instructions for workspace catalog applications on the [Application Deployment](#) (see page 671) topic.

⁴⁹⁹ <https://banzaicloud.com/docs/supertubes/kafka-operator/>

⁵⁰⁰ <https://github.com/banzaicloud/koperator>

⁵⁰¹ <https://github.com/banzaicloud/koperator/tree/master/config/samples>

⁵⁰² <https://kafka.apache.org/documentation/>

 Only install the Spark operator once per workspace.

For details on custom configuration for the operator, refer to the [Spark Operator Helm Chart documentation](#)⁵⁰³.

After you finish the installation, see [Spark Operator in a Project](#) (see page 701) custom resource documentation for more information about how to submit your Spark jobs.

6.3.7.1.3.2 Sample Override Configuration File

 Ensure you configure the AppDeployment with the appropriate override configmap.

- Using UI

```
podLabels:
  owner: john
  team: operations
```

- Using CLI

See [Application Deployment](#) (see page 671) for details.

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: spark-operator-overrides
data:
  values.yaml: |
    configInline:
      podLabels:
        owner: john
        team: operations
EOF
```

⁵⁰³ <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/charts/spark-operator-chart/README.md>

6.3.7.1.3.3 Uninstall via the CLI



Uninstalling the Spark Operator does not affect existing **SparkApplication** and **ScheduledSparkApplication** custom resources. You need to manually remove any leftover custom resources and CRDs from the operator. refer to [deleting Spark Operator custom resources](#) (see page 0).

Follow these steps:

1. Uninstall the Spark Operator `AppDeployment` :

```
kubectl -n <your workspace namespace> delete AppDeployment <name of AppDeployment>
```

2. Remove the Spark Operator Service Account:

```
# <name of service account> is spark-operator-service-account if you didn't
override the RBAC resources.
kubectl -n <your workspace namespace> delete serviceaccounts <name of service
account>
```

3. Remove the Spark Operator CRDs:

NOTE: The CRDs are not finalized for deletion until you delete the associated custom resources.

```
kubectl delete crds scheduledsparkapplications.sparkoperator.k8s.io
sparkapplications.sparkoperator.k8s.io
```

6.3.7.1.3.4 Resources

Here are some resources to learn more about Spark Operator:

- [Spark Operator Documentation](#)⁵⁰⁴
- [Spark Operator Quick Start Guide](#)⁵⁰⁵
- [Spark Operator User Guide](#)⁵⁰⁶
- [Spark Documentation](#)⁵⁰⁷

504 <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/README.md>

505 <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/docs/quick-start-guide.md>

506 <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/docs/user-guide.md>

507 <https://spark.apache.org/docs/latest/>

6.3.7.1.4 Zookeeper Operator in a Workspace


Enterprise

Gov Advanced

[ZooKeeper](#)⁵⁰⁸ is a centralized service for maintaining configuration information, naming, providing distributed synchronization, and providing group services. The [ZooKeeper operator](#)⁵⁰⁹ is a Kubernetes operator that handles the provisioning and management of ZooKeeper clusters. It works by watching custom resources, such as `ZookeeperClusters`, to provision the underlying Kubernetes resources (`StatefulSets`) required for a production-ready ZooKeeper Cluster.

6.3.7.1.4.1 Install Zookeeper

Follow these steps to install the ZooKeeper operator in a workspace. This procedure results in a ZooKeeper operator running in a workspace namespace, ready to manage and create ZooKeeper clusters in any project namespaces. See the [ZooKeeper in a Project](#) (see page 696) custom resource documentation for more information on creating ZooKeeper clusters.

 Only install the ZooKeeper operator once per workspace.

1. Follow the generic installation instructions for workspace catalog applications on the [Application Deployment](#) (see page 671) page.
2. Within the `AppDeployment`, update the `appRef` to specify the correct `zookeeper-operator` App. You can find the `appRef.name` by listing the available `Apps` in the workspace namespace:

```
kubectl get apps -n ${WORKSPACE_NAMESPACE}
```

For details on custom configuration for the operator, refer to the [ZooKeeper operator Helm Chart documentation](#)⁵¹⁰.

6.3.7.1.4.2 Uninstall via the CLI

Uninstalling the ZooKeeper operator will not directly affect any running `ZookeeperClusters`. By default, the operator waits for any `ZookeeperClusters` to be deleted before it will fully uninstall (you can set `hooks.delete: true` in the application configuration to disable this behavior). After uninstalling the operator, you need to manually clean up any leftover Custom Resource Definitions (CRDs).

⁵⁰⁸ <https://zookeeper.apache.org/index.html>

⁵⁰⁹ <https://github.com/pravega/zookeeper-operator>

⁵¹⁰ <https://github.com/pravega/zookeeper-operator/tree/master/charts/zookeeper-operator#configuration>

1. Delete all `ZookeeperClusters`, see [Deleting ZooKeeper Clusters](#) (see page 698).
2. Uninstall a ZooKeeper operator `AppDeployment` :

```
kubectl -n <workspace namespace> delete AppDeployment <name of AppDeployment>
```

3. Remove ZooKeeper CRDs:

WARNING: After you remove the CRDs, all deployed ZookeeperClusters will be deleted!

```
kubectl delete crds zookeeperclusters.zookeeper.pravega.io
```

6.3.7.1.4.3 Resources

- [ZooKeeper Operator Documentation](#)⁵¹¹
- [ZooKeeper Cluster Helm Chart](#)⁵¹²
- [ZooKeeper Documentation](#)⁵¹³

6.3.7.2 Deployment of Catalog Applications in Workspaces

Enterprise

Gov Advanced

Deploy applications to attached clusters using the CLI

This topic describes how to use the CLI to deploy a workspace catalog application to attached clusters within a workspace. To deploy an application to selected clusters within a workspace, refer to the [Cluster-scoped Configuration](#) (see page 657) section of the documentation.

6.3.7.2.1 Prerequisites

Before you begin, you must have:

- A running cluster with [Kommander installed](#) (see page 1422). The cluster should be on a supported Kubernetes version for [this release of DKP](#) (see page 1752), and also compatible with the [catalog application version](#) (see page 667) you wish to install.
- An [Attach an Existing Kubernetes Cluster](#) (see page 752) section of the documentation completed.

Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace the attached cluster exists in:

⁵¹¹ <https://github.com/pravega/zookeeper-operator>

⁵¹² <https://github.com/pravega/zookeeper-operator/tree/master/charts/zookeeper>

⁵¹³ <https://zookeeper.apache.org/documentation>

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

After creating a GitRepository, use either the DKP UI or the CLI to enable your catalog applications.



From within a workspace, you can enable applications to deploy. Verify that an application has successfully deployed [via the CLI \(see page 674\)](#).

6.3.7.2.2 Enable (and Customize) the Application using the DKP UI

Follow these steps to enable your catalog applications from the DKP UI:

1. **Enterprise only:** from the top menu bar, select your target workspace.
2. Select **Applications** from the sidebar menu to browse the available applications from your configured repositories.
3. Select the three dot button from the bottom-right corner of the desired application card, and then select **Enable**.
4. If available, select a version from the drop-down menu. This drop-down menu will only be visible if there is more than one version.
5. (Optional) If you want to override the default configuration values, copy your customized values into the text editor under **Configure Service** or upload your YAML file that contains the values:

```
someField: someValue
```

6. Confirm the details are correct, and then select the **Enable** button.

For all applications, you must provide a display name and an ID which is automatically generated based on what you enter for the display name, unless or until you edit the ID directly. The ID must be compliant with [Kubernetes DNS subdomain name validation rules](#)⁵¹⁴.

Alternately, you can [use the CLI to enable your catalog applications \(see page 0\)](#).

6.3.7.2.3 Enable the Application using the CLI

See [Workspace Catalog Applications \(see page 666\)](#) for the list of available applications that you can deploy on the attached cluster.

1. Enable a supported application to deploy to your [Attached Cluster \(see page 752\)](#) with an `AppDeployment` resource.
2. Within the `AppDeployment`, define the `appRef` to specify which `App` to enable:

⁵¹⁴ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#dns-subdomain-names>

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: spark-operator
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: spark-operator-1.1.17
    kind: App
EOF
```



- The `appRef.name` must match the app `name` from the list of available catalog applications.
- Create the resource in the workspace you just created, which instructs Kommander to deploy the `AppDeployment` to the `KommanderCluster`s in the same workspace.

6.3.7.2.4 Enable an Application with a Custom Configuration using the CLI

Follow these steps:

Provide the name of a `ConfigMap` in the `AppDeployment`, which provides custom configuration on top of the default configuration:

1.

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: spark-operator
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: spark-operator-1.1.17
    kind: App
  configOverrides:
    name: spark-operator-overrides
EOF
```

2. Create the `ConfigMap` with the name provided in the step above, with the custom configuration:

```
cat <<EOF | kubectl apply -f -
```

```

apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: spark-operator-overrides
data:
  values.yaml: |
    configInline:
      uiService:
        enable: false
EOF

```

Kommander waits for the `ConfigMap` to be present before deploying the `AppDeployment` to the managed or attached clusters.

6.3.7.2.5 Verify Applications

The applications are now enabled. Connect to the attached cluster and check the `HelmReleases` to verify the deployment:

```
kubectl get helmreleases -n ${WORKSPACE_NAMESPACE}
```

The output looks similar to this:

NAMESPACE	NAME	READY	STATUS
AGE			
workspace-test-vjsfq 7m3s	spark-operator	True	Release reconciliation succeeded

6.3.7.3 Workspace Catalog Application Upgrades

Enterprise

Gov Advanced

Upgrade catalog applications using the CLI and UI

6.3.7.3.1 Prerequisites

See [Upgrade Prerequisites](#) (see page 1555).

6.3.7.3.2 Upgrade Catalog Applications

Before upgrading, keep in mind the distinction between Platform applications and Catalog applications. Platform applications are deployed and upgraded as a set for each cluster or workspace. Catalog

applications are deployed separately, so that you can deploy and upgrade them individually for each workspace or project.

6.3.7.3.2.1 Upgrade with UI

Follow these steps to upgrade an application from the DKP UI:

1. From the top menu bar, select your target workspace.
2. Select **Applications** from the sidebar menu.
3. Select the three dot button from the bottom-right corner of the desired application tile, and then select **Edit**.
4. Select the **Version** drop-down, and select a new version. This drop-down will only be available if there is a newer version to upgrade to.
5. Select **Save**.

6.3.7.3.2.2 Upgrade with CLI



Please note that the below commands are using the workspace **name** and not namespace. You can retrieve the workspace name by running the command:

```
dkp get workspaces
```

To view a list of the deployed apps to your workspace, you can run the command:

```
dkp get appdeployments --workspace=<workspace-name>
```

1. To see what app(s) and app versions are available to upgrade, run the following command:
NOTE: The **APP ID** column displays the available apps and the versions available to upgrade.

```
kubectl get apps -n <workspace-namespace>
```

2. Run the following command to upgrade an application from the DKP CLI:

```
dkp upgrade catalogapp <appdeployment-name> --workspace=<my-workspace-name> --to-version=<version.number>
```

As an example, the following command upgrades the Kafka Operator application, named `kafka-operator-abc`, in a workspace to version `0.25.1`:

```
dkp upgrade catalogapp kafka-operator-abc --workspace=my-workspace --to-version=0.25.1
```

- Platform applications cannot be upgraded on a one-off basis, and must be upgraded in a single process for each workspace. If you attempt to upgrade a platform application with these commands, you receive an error and the application is not upgraded.

6.3.7.4 Custom Applications

Enterprise

Gov Advanced

Custom applications are third-party applications you have added to the DKP Catalog.

Custom applications are any third-party applications that are not provided in the DKP Application Catalog. Custom applications can leverage applications from the DKP Catalog or be fully-customized. There is no expectation of support by D2iQ for a Custom application. Custom applications can be deployed on Konvoy clusters or on any D2iQ supported 3rd party Kubernetes distribution.

- [Create a Git Repository \(see page 676\)](#)
- [Git Repository Structure \(see page 678\)](#)
- [Workspace Application Metadata \(see page 680\)](#)
- [Enable a Custom Application from the Workspace Catalog \(see page 682\)](#)

6.3.7.4.1 Create a Git Repository

Enterprise

Gov Advanced

Create a Git Repository in the Workspace namespace

Use the CLI to create the GitRepository resource and add a new repository to your Workspace.

1. Refer to [Air-gapped Environment \(see page 1397\)](#) setup instructions section, if you are running in air-gapped environment.
2. Set the `WORKSPACE_NAMESPACE` environment variable to the name of your workspace's namespace:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

3. Adapt the URL of your Git repository:

```
kubectl apply -f - <<EOF
apiVersion: source.toolkit.fluxcd.io/v1beta1
kind: GitRepository
```

```

metadata:
  name: example-repo
  namespace: ${WORKSPACE_NAMESPACE}
  labels:
    kommander.d2iq.io/gitapps-gitrepository-type: dkp
    kommander.d2iq.io/gitrepository-type: catalog
spec:
  interval: 1m0s
  ref:
    branch: <your-target-branch-name> # e.g., main
  timeout: 20s
  url: https://github.com/<example-org>/<example-repo>
EOF

```

4. Ensure the status of the `GitRepository` signals a ready state:

```
kubectl get gitrepository example-repo -n ${WORKSPACE_NAMESPACE}
```

The repository commit also displays the ready state:

NAME	URL	READY
example-repo	https://github.com/example-org/example-repo	True
	Fetches revision: master/6c54bd1722604bd03d25dcac7a31c44ff4e03c6a	11m

For more information on the `GitRepository` resource fields and how to make Flux aware of credentials required to access a private Git repository, see the [Flux documentation](#)⁵¹⁵.

6.3.7.4.1.1 Troubleshoot

To troubleshoot issues with adding the `GitRepository`, review the following logs:

```

kubectl -n kommander-flux logs -l app=source-controller
[...]
kubectl -n kommander-flux logs -l app=kustomize-controller
[...]
kubectl -n kommander-flux logs -l app=helm-controller
[...]

```

6.3.7.4.1.2 Related Information

- [Flux](#)⁵¹⁶

⁵¹⁵ <https://fluxcd.io/flux/components/source/gitrepositories/#secret-reference>

⁵¹⁶ <https://fluxcd.io/>

- [Flux docs](#)⁵¹⁷

6.3.7.4.2 Git Repository Structure

Enterprise

Gov Advanced

Git repositories must be structured in a specific manner for defined applications to be processed by Kommander.

You must structure your git repository based on the following guidelines, for your applications to be processed properly by Kommander so that they can be deployed.

6.3.7.4.2.1 Git Repository Directory Structure

Use the following basic directory structure for your git repository:

```

├── helm-repositories
│   ├── <helm repository 1>
│   │   ├── kustomization.yaml
│   │   └── <helm repository name>.yaml
│   └── <helm repository 2>
│       ├── kustomization.yaml
│       └── <helm repository name>.yaml
├── services
│   ├── <app name>
│   │   ├── <app version1> # semantic version of the app helm chart. e.g., 1.2.3
│   │   │   ├── defaults
│   │   │   │   ├── cm.yaml
│   │   │   │   └── kustomization.yaml
│   │   │   ├── <app name>.yaml
│   │   │   └── kustomization.yaml
│   │   └── <app version2> # another semantic version of the app helm chart. e.g.,
│   │       2.3.4
│   │       ├── defaults
│   │       │   ├── cm.yaml
│   │       │   └── kustomization.yaml
│   │       ├── <app name>.yaml
│   │       └── kustomization.yaml
│   │       └── metadata.yaml
│   └── <another app name>
│       ...

```

- Define applications in the `services/` directory.
- You can define multiple versions of an application, under different directories nested under the `services/<app name>/` directory.

⁵¹⁷ <https://fluxcd.io/docs>

- Define application manifests, such as a `HelmRelease`⁵¹⁸, under each versioned directory `services/<app name>/<version>/` in `<app name>.yaml` which is listed in the `kustomization.yaml` Kubernetes Kustomization file. For more information, see the [Kubernetes Kustomization docs](#)⁵¹⁹.
- Define the default values ConfigMap for `HelmReleases` in the `services/<app name>/<version>/defaults` directory, accompanied by a `kustomization.yaml` Kubernetes Kustomization file pointing to the `ConfigMap` file.
- Define the `metadata.yaml` of each application under the `services/<app name>/` directory. For more information, see the [Workspace Application Metadata docs](#) (see page 680).

See the [DKP Catalog repository](#)⁵²⁰ for an example of how to structure custom catalog Git repositories.

6.3.7.4.2 Helm Repositories

You must include the `HelmRepository` that is referenced in each `HelmRelease`'s `Chart` spec.

Each `services/<app name>/<version>/kustomization.yaml` must include the path of the YAML file that defines the `HelmRepository`. For example:

```
# services/<app name>/<version>/kustomization.yaml
apiVersion: kustomize.config.k8s.io/v1beta1
kind: Kustomization
resources:
  - <app name>.yaml
  - ../../../../helm-repositories/<helm repository 1>
```

For more information, see the flux documentation about [HelmRepositories](#)⁵²¹.

6.3.7.4.2.3 Substitution Variables

Some [substitution variables](#)⁵²² are provided.

- `${releaseName}` : For each App deployment, this variable is set to the `AppDeployment` name. Use this variable to prefix the names of any resources that are defined in the application directory in the Git repository so that multiple instances of the same application can be deployed. If you create resources without using the `releaseName` prefix (or suffix) in the name field, there can be conflicts if the same named resource is created in that same namespace.
- `${releaseNamespace}` : The namespace of the Workspace.
- `${workspaceNamespace}` : The namespace of the Workspace that the Workspace belongs to.

518 <https://fluxcd.io/docs/components/helm/helmreleases/>

519 <https://kubectl.docs.kubernetes.io/references/kustomize/kustomization/>

520 <https://github.com/mesosphere/dkp-catalog-applications>

521 <https://fluxcd.io/docs/components/source/helmrepositories/>

522 <https://fluxcd.io/docs/components/kustomize/kustomization/#variable-substitution>

6.3.7.4.2.4 Related Information

- [Flux](#)⁵²³
- [Flux docs](#)⁵²⁴
- [Getting started with Flux guide](#)⁵²⁵

6.3.7.4.3 Workspace Application Metadata

Enterprise

Gov Advanced



To display more information about custom applications in the UI, define a `metadata.yaml` file for each application in the Git repository.

You can define how custom applications display in the DKP UI by defining a `metadata.yaml` file for each application in the git repository. You must define this file at `services/<application>/metadata.yaml` for it to process correctly.

You can define the following fields:

Field	Default	Description
<code>displayName</code>	falls back to App ID	Display name of the application for the UI.
<code>description</code>	""	Short description, should be a sentence or two, displayed in the UI on the application card.
<code>category</code>	general	1 or more categories for this application. Categories are used to group applications in the UI.
<code>overview</code>		Markdown overview used on the application detail page in the UI.

⁵²³ <https://fluxcd.io/>

⁵²⁴ <https://fluxcd.io/docs>

⁵²⁵ <https://fluxcd.io/docs/get-started/>

Field	Default	Description
icon		Base64 encoded icon SVG file used for application logos in the UI.
scope	project	List of scopes, can be set only to project or workspace currently.

None of these fields are required for the application to display in the UI.

Here is an example `metadata.yaml` file:

```

displayName: Prometheus Monitoring Stack
description: Stack of applications that collect metrics and provides visualization
and alerting capabilities. Includes Prometheus, Prometheus Alertmanager and Grafana.
category:
  - monitoring
overview: >
  # Overview
  A stack of applications that collects metrics and provides visualization and
  alerting capabilities. Includes Prometheus, Prometheus Alertmanager and Grafana.

  ## Dashboards
  By deploying the Prometheus Monitoring Stack, the following platform applications
  and their respective dashboards are deployed. After deployment to clusters in a
  workspace, the dashboards are available to access from a respective cluster's detail
  page.

  ### Prometheus

  A software application for event monitoring and alerting. It records real-time
  metrics in a time series database built using a HTTP pull model, with flexible and
  real-time alerting.

  - [Prometheus Documentation - Overview](https://prometheus.io/docs/introduction/
  overview/)

  ### Prometheus Alertmanager

  A Prometheus component that enables you to configure and manage alerts sent by the
  Prometheus server and to route them to notification, paging, and automation systems.

  - [Prometheus Alertmanager Documentation - Overview](https://prometheus.io/docs/
  alerting/latest/alertmanager/)

  ### Grafana

  A monitoring dashboard from Grafana that can be used to visualize metrics collected
  by Prometheus.

```

- [Grafana Documentation] (<https://grafana.com/docs/>)
 icon:
 PHN2ZyB4bWxucz0iaHR0cDovL3d3dy53My5vcmcvMjAwMC9zdmc iIHZpZXhCb3g9IjAgMCAzMdAgMzAwIiBzd
 HlsZT0iZW5hYmxlLWJhY2tncm91bmQ6bmV3IDAgMCAzMdAgMzAwIiB4bWw6c3BhY2U9InByZXNlc nZlIj48cG
 F0aCBkPSJNMTUwIDUwQzk0LjggNTAgNTAgOTQuOCA1MCAxNTBzNDQuOCAxMDAgMTAwIDEwMCAxMDAtNDQuOCA
 xMDAtMTAwUzIwNS4yIDUwIDE1MCA1MHptMCAxODcuMmMtMTUuNyAwLTI4LjUtMTAuNS0yOCA1LTIzLjRoNTYu
 OWMuMSAxMi45LTEyLjcgMjMuNC0yOCA0IDIZLjR6bTQ3LTMxLjJoLTk0di0xN2g5NHYxN3ptLS4zLTI1LjloL
 TkzLjRjLS4zLS40LS42LS43LS45LTEuMS05LjYtMTEuNy0xMS45LTE3LjgtMTQuMS0yNCAwLS4yIDExLjcgMi
 40IDIwIDQuMyAwIDA gNC4zIDEgMTAuNSAyLjEtNi03LTkuNi0xNi05LjYtMjUuMSAwLTIwIDE1LjQtMzcuNiA
 5LjgtNTEuNyA1LjQuNCAxMS4yIDExLjQgMTEuNiAyOCA1IDUuNy03LjkgOC4xLTIyLjQgOC4xLTMxLjMgMC05
 LjI gNi4xLTE5LjkgMTIuMS0yMCA4LTUuNCA4LjkgMS40IDE2LjUgNy40IDM1LjUgMi4zIDcuMSAyIDE5LjEgM
 y43IDI2LjcuNi0xNS44IDMuMy0zOCA4IDEzLjMtNDYuNy00LjQgMTAgLjcgMjUuNSA0LjEgMjguNSA1LjYgOS
 43IDkgMTEuMSA5IDMxIDA gOS4zLTMuNCAxOC4xLTKuMyAyNSA2LjYtMS4yIDExLjItMi40IDExLjItMi40bDI
 xLjQtNC4yYy4xIDA tMyAxMi44LTE0LjkgMjUuMXoiIHh0eWw lPSJmaWxsOiNmODQzMTEiLz48L3N2Zz4=

6.3.7.4.4 Enable a Custom Application from the Workspace Catalog

Enterprise

Gov Advanced

Enable a Custom Application from the Workspace Catalog

After creating a GitRepository, you can either use the DKP UI or the CLI to enable your custom applications. To deploy an application to selected clusters within a workspace, refer to the [cluster-scoped configuration](#) (see page 657) section.



From within a workspace, you can enable applications to deploy. Verify that an application has successfully deployed [via the CLI](#) (see page 0).

6.3.7.4.4.1 Enable the Application Using the UI

Follow these steps:

1. From the top menu bar, select your target workspace.
2. Select **Applications** from the sidebar menu to browse the available applications from your configured repositories.
3. Select the three dot button from the bottom-right corner of the desired application tile, and then select **Enable**.
4. If available, select a version from the drop-down menu. This drop-down menu will only be visible if there is more than one version.
5. (Optional) If you want to override the default configuration values, copy your customized values into the text editor under **Configure Service** or upload your YAML file that contains the values:

```
someField: someValue
```

6. Confirm the details are correct, and then select the **Enable** button.

For all applications, you must provide a display name and an ID which is automatically generated based on what you enter for the display name, unless or until you edit the ID directly. The ID must be compliant with [Kubernetes DNS subdomain name validation rules](#)⁵²⁶.

Alternately, you can [use the CLI to enable your custom applications](#) (see page 683).

6.3.7.4.4.2 Prerequisites

- Determine the name of the workspace where you wish to perform the deployments. You can use the `dkp get workspaces` command to see the list of workspace names and their corresponding namespaces.
- Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace where the cluster is attached:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

6.3.7.4.4.3 Enable the Application using the CLI

Follow these steps:

1. Get the list of available applications to enable using the following command:

```
kubectl get apps -n ${WORKSPACE_NAMESPACE}
```

2. Deploy one of the supported applications from the list with an `AppDeployment` resource.
3. Within the `AppDeployment`, define the `appRef` to specify which `App` will be enabled:

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: my-custom-app
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: custom-app-0.0.1
    kind: App
EOF
```

⁵²⁶ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#dns-subdomain-names>



- The `appRef.name` must match the `app name` from the list of available catalog applications.
- Create the resource in the workspace you just created, which instructs Kommander to deploy the `AppDeployment` to the `KommanderCluster`s in the same workspace.

6.3.7.4.4.4 Enable an Application with a Custom Configuration using the CLI

Follow these steps:

1. Provide the name of a `ConfigMap` in the `AppDeployment`, which provides custom configuration on top of the default configuration:

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: my-custom-app
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: custom-app-0.0.1
    kind: App
  configOverrides:
    name: my-custom-app-overrides
EOF
```

2. Create the `ConfigMap` with the name provided in the step above, with the custom configuration:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: my-custom-app-overrides
data:
  values.yaml: |
    someField: someValue
EOF
```

Kommander waits for the `ConfigMap` to be present before deploying the `AppDeployment` to the managed or attached clusters in the Workspace.

6.3.7.4.5 Verify Applications

After completing the previous steps, your applications are enabled. Connect to the attached cluster and check the `HelmsReleases` to verify the deployments:

```
kubect1 get helmreleases -n ${WORKSPACE_NAMESPACE}
```

The output looks similar to this:

NAMESPACE	NAME	READY	STATUS
AGE			
workspace-test-vjsfq 7m3s	my-custom-app	True	Release reconciliation succeeded

6.3.8 Workspace Role Bindings

Enterprise

Gov Advanced

Workspace Role Bindings grant access to specified Workspace Roles for a specified group of people.

6.3.8.1 Configure Workspace Role Bindings

Before you can create a Workspace Role Binding, ensure you have created a Group. A Kommander Group can contain one or several Identity Provider users, groups or both.

You can assign a role to this Kommander Group:

1. From the top menu bar, select your target workspace.
2. Select **Access Control** in the **Administration** section of the sidebar menu.
3. Select the **Cluster Role Bindings** tab, and then select the **Add Roles** button next to the group you want.
4. Select the Role, or Roles, you want from the drop-down menu and select **Save**.



It will take a few minutes for the resource to be created.

6.4 Projects

Enterprise

Gov Advanced

Multi-cluster Configuration Management

Projects support the management of configMaps, continuous deployments, secrets, services, quotas, and role-based access control and multi-tenant logging by leveraging federated resources. When a Project is created, DKP creates a federated namespace that is propagated to the Kubernetes clusters associated with this Project.

Federation in this context means that a common configuration is pushed out from a central location (DKP) to all Kubernetes clusters, or a pre-defined subset group, under DKP management. This pre-defined subset group of Kubernetes clusters is called a Project.

Projects enable teams to deploy their configurations and services to clusters in a consistent way. Projects enable central IT or a business unit to share their Kubernetes clusters among several teams. Using Projects, DKP leverages Kubernetes Cluster Federation (KubeFed) to coordinate the configuration of multiple Kubernetes clusters.

Kommander allows a user to use labels to select, manually or dynamically, the Kubernetes clusters associated with a Project.

6.4.1 Project Namespace

Project Namespaces isolate configurations across clusters. Individual standard Kubernetes namespaces are automatically created on all clusters belonging to the project. When creating a new project, you can customize the Kubernetes namespace name that is created. It is the grouping of all of these individual standard Kubernetes namespaces that make up the concept of a Project Namespace. A Project Namespace is a Kommander specific concept.

6.4.2 Create a Project

When you create a Project, you must specify a Project Name, a Namespace Name (optional) and a way to allow Kommander to determine which Kubernetes clusters will be part of this project.

As mentioned previously, a Project Namespace corresponds to a Kubernetes Federated Namespace. By default, the name of the namespace is auto-generated based on the project name (first 57 characters) plus 5 unique alphanumeric characters. You can specify a namespace name, but you must ensure it does not conflict with any existing namespace on the target Kubernetes clusters, that will be a part of the Project.

To determine which Kubernetes clusters will be part of this project, you can either select manually existing clusters or define labels that Kommander will use to dynamically add clusters. The latter is recommended because it will allow you to deploy additional Kubernetes clusters later and to have them automatically associated with Projects based on their labels.

To create a Project, you can either use the DKP UI or create a Project object on the Kubernetes cluster where Kommander is running (using kubectl or the Kubernetes API). The latter allows you to configure Kommander resources in a declarative way. It is available for all kinds of Kommander resources.

6.4.3 Create a Project - UI Method

Here is an example of what it looks like to create a project using the DKP UI:

Cancel
Create Project
Create

General

Project Name *

ID / Namespace

By default, a unique ID / Namespace will be auto-generated based on the project name (first 57 characters) plus 5 unique alphanumeric characters. You can also specify a custom ID / Namespace.

Description

Clusters

Add one or more clusters to this project. Project configurations, including deployed platform services, will be applied to all clusters in this project.

Manually Select Clusters

Explicitly select cluster(s) to include in this project.

Dynamically Select Clusters

Cluster(s) will be dynamically added & removed for this project based on the cluster labels included that match respective clusters.

Key **Value** x

+ Add Label

1 Cluster

Name	Type	Provider	Labels
cluster1	Managed	aws	dev: true

6.4.4 Create a Project - CLI Method

The following sample is a YAML Kubernetes object for creating a Kommander Project. This example does not work verbatim because it depends on a workspace name that has been previously created and does not exist by default in your cluster. Use this as an example format and fill in the workspace name and namespace name appropriately along with the proper labels.

```

apiVersion: workspaces.kommander.mesosphere.io/v1alpha1
kind: Project
metadata:
  name: My-Project-Name
  namespace: my-project-k8s-namespace-name
spec:
  workspaceRef:
    name: myworkspacename
    namespaceName: myworkspacename-di3tx
  placement:
    clusterSelector:
      matchLabels:
        cluster: prod

```

The following procedures are supported for projects:

- [Project Applications](#) (see page 688)
- [Project Deployments](#) (see page 713)
- [Project Role Bindings](#) (see page 728)
- [Project Roles](#) (see page 732)
- [Project ConfigMaps](#) (see page 735)
- [Project Secrets](#) (see page 736)
- [Project Quotas & Limit Ranges](#) (see page 738)
- [Project Network Policies](#) (see page 740)

6.4.5 Project Applications

Enterprise

Gov Advanced

This section documents the applications and application types that you can utilize with DKP.

Application types are:

- [Catalog Applications](#) (see page 694) are either pre-packaged applications from the D2iQ Application Catalog, or custom applications that you maintain for your teams or organization.
 - [DKP Applications](#) (see page 696) are applications that are provided by D2iQ and added to the Catalog.
 - [Custom Applications](#) (see page 704) are applications you create and add to the Catalog.
- [Platform Applications](#) (see page 688) are applications integrated into Kommander.



When deploying and upgrading applications, platform applications come as a bundle; they are tested as a single unit and you must deploy or upgrade them in a single process, for each workspace. This means all clusters in a workspace have the same set and versions of platform applications deployed. Whereas catalog applications are individual, so you can deploy and upgrade them individually, for each project.

6.4.5.1 Project Platform Applications

Enterprise

Gov Advanced

How project Platform applications work

The following table describes the list of applications that can be deployed to attached clusters within a project.

Review the [Project Platform Application Requirements](#) (see page 694) to ensure that the attached clusters in the project have sufficient resources.



From within a project, you can enable applications to deploy. Verify that an application has successfully deployed [via the CLI \(see page 642\)](#).

6.4.5.1.1 Enable Applications in a Project Using the UI

1. From the top menu bar, select your target workspace.
2. Select **Projects** from the sidebar menu.
3. Select your project from the list.
4. Select the **Applications** tab to browse the available applications.
5. Select the three dot button from the bottom-right corner of the desired application tile, and then select **Enable**.
6. To override the default configuration values, copy your values content into the text editor under **Configure Service** or just upload your yaml file that contains the values:

```
someField: someValue
```

7. Select the **Enable** button.

To use the CLI to enable or disable applications, see [Application Deployment \(see page 642\)](#)



There may be dependencies between the applications, which are listed in [Project Platform Application Dependencies \(see page 692\)](#). Review them carefully prior to customizing to ensure that the applications are deployed successfully.

6.4.5.1.2 Platform Applications

NAME	APP ID	Deployed by default
project-grafana-logging-6.57.4	project-grafana-logging	False
project-grafana-loki-0.69.16	project-grafana-loki	False
project-logging-1.0.3	project-logging	False

6.4.5.1.3 Upgrade Platform Applications from the CLI

Platform Applications within a Project are automatically upgraded when the Workspace that a Project belongs to is upgraded. See [Enterprise: Upgrade the Management Cluster and Platform Applications](#) (see page 1565) for more information on how to upgrade these applications.

6.4.5.1.4 Deploy Project Platform Applications

Enterprise

Gov Advanced

Deploy applications to attached clusters in a project using the CLI

This topic describes how to use the CLI to deploy an application to attached clusters within a project. To use the DKP UI to deploy applications, see [Deploy Applications in a Project](#) (see page 691).

See [Project Platform Applications](#) (see page 691) for a list of all applications and those that are enabled by default.

6.4.5.1.4.1 Prerequisites

Before you begin, you must have:

- A running cluster with [Kommander installed](#) (see page 108).
- An [existing Kubernetes cluster attached to Kommander](#) (see page 752).

Set the `WORKSPACE_NAME` environment variable to the name of the workspace where the cluster is attached:

```
export WORKSPACE_NAME=<workspace_name>
```

Set the `WORKSPACE_NAMESPACE` environment variable to the namespace of the above workspace:

```
export WORKSPACE_NAMESPACE=$(kubectl get namespace --
selector='workspaces.kommander.mesosphere.io/workspace-name=${WORKSPACE_NAME}' -o
jsonpath='{.items[0].metadata.name}')
```

Set the `PROJECT_NAME` environment variable to the name of the project in which the cluster is included:

```
export PROJECT_NAME=<project_name>
```

Set the `PROJECT_NAMESPACE` environment variable to the name of the above project's namespace:

```
export PROJECT_NAMESPACE=$(kubectl get project ${PROJECT_NAMESPACE} -n ${
WORKSPACE_NAMESPACE} -o jsonpath='{.status.namespaceRef.name}')
```

6.4.5.1.4.2 Deploy the Application

The list of available applications that can be deployed on the attached clusters in a project can be found in [the project platform applications topic](#) (see page 691).

1. Deploy one of the supported applications to [your existing attached cluster](#) (see page 752) with an `AppDeployment` resource. Provide the `appRef` and application version to specify which `App` is deployed:

```
dkp create appdeployment project-grafana-logging --app project-grafana-
logging-6.38.1 --workspace ${WORKSPACE_NAME} --project ${PROJECT_NAME}
```

2. Create the resource in the project you just created, which instructs Kommander to deploy the `AppDeployment` to the `KommanderClusters` in the same project.



- The `appRef.name` must match the `app name` from the list of available applications.
- Observe that the `dkp create` command must be run with both the `--workspace` and `--project` flags for project platform applications.

6.4.5.1.4.3 Deploy an Application with a Custom Configuration

Follow these steps:

1. Create the `AppDeployment` and provide the name of a `ConfigMap`, which provides custom configuration on top of the default configuration:

```
dkp create appdeployment project-grafana-logging --app project-grafana-
logging-6.38.1 --config-overrides project-grafana-logging-overrides --workspace
${WORKSPACE_NAME} --project ${PROJECT_NAMESPACE}
```

2. Create the `ConfigMap` with the name provided in the step above, with the custom configuration:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
```

```

metadata:
  namespace: ${PROJECT_NAMESPACE}
  name: project-grafana-logging-overrides
data:
  values.yaml: |
    datasources:
      datasources.yaml:
        apiVersion: 1
        datasources:
        - name: Loki
          type: loki
          url: "http://project-grafana-loki-loki-distributed-gateway"
          access: proxy
          isDefault: false
EOF

```

Kommander waits for the `ConfigMap` to be present before deploying the `AppDeployment` to the attached clusters.

6.4.5.1.4.4 Verify Applications

The applications are now enabled.

Export the `project_namespace` with this command:

```
export PROJECT_NAMESPACE=<project_namespace>
```

Connect to the attached cluster and check the `HelmsReleases` to verify the deployment:

```

kubectl get helmreleases -n ${PROJECT_NAMESPACE}

```

NAMESPACE	NAME	READY	STATUS
project-test-vjsfq	project-grafana-logging	True	Release reconciliation succeeded 7m3s



Some of the supported applications have dependencies on other applications. See [Project Application Dependencies \(see page 692\)](#) for that table.

6.4.5.1.5 Project Platform Application Dependencies

Enterprise

Gov Advanced

Dependencies between project platform applications

There are many dependencies between the applications that are deployed to a project's attached clusters. It is important to note these dependencies when customizing the platform applications to ensure that your services are properly deployed to the clusters. For more information on how to customize platform applications, see [Platform Application Deployment](#) (see page 693).

6.4.5.1.5.1 Application Dependencies

When deploying or troubleshooting applications, it helps to understand how applications interact and may require other applications as dependencies.

If an application's dependency does not successfully deploy, the application requiring that dependency does not successfully deploy.

The following sections detail information about the platform applications.

6.4.5.1.5.2 Logging

Collects logs over time from Kubernetes pods deployed in the project namespace. Also provides the ability to visualize and query the aggregated logs.

- [project-logging](#)⁵²⁷: Defines resources for the Logging Operator which uses them to direct the project's logs to its respective Grafana Loki application.
- [project-grafana-loki](#)⁵²⁸: A horizontally-scalable, highly-available, multi-tenant log aggregation system inspired by Prometheus.
- [project-grafana-logging](#)⁵²⁹: Logging dashboard used to view logs aggregated to Grafana Loki.



The project logging applications depend on the [workspace logging applications](#) (see page 902) being deployed.

Application	Required Dependencies
project-logging	logging-operator (workspace)
project-grafana-loki	project-logging, grafana-loki (workspace), logging-operator (workspace)
project-grafana-logging	project-grafana-loki

⁵²⁷ <https://grafana.com/oss/grafana/>

⁵²⁸ <https://grafana.com/oss/loki/>

⁵²⁹ <https://grafana.com/oss/grafana/>

6.4.5.1.6 Project Platform Application Configuration Requirements

Enterprise

Gov Advanced

6.4.5.1.6.1 Project Platform Application Descriptions and Resource Requirements

Platform applications require more resources than solely deploying or attaching clusters into a project. Your cluster must have sufficient resources when deploying or attaching to ensure that the applications are installed successfully.

The following table describes all the platform applications that are available to the clusters in a project, minimum resource and persistent storage requirements, and whether they are enabled by default.

Name	Minimum Resources Suggested	Minimum Persistent Storage Required	Deployed by Default	Default Priority Class
project-grafana-logging	cpu: 200m memory: 100Mi		No	DKP Critical (100002000)
project-grafana-loki		# of PVs: 3 PV sizes: 10Gi x 3 (total: 30Gi)	No	DKP Critical (100002000)
project-logging			No	DKP Critical (100002000)

6.4.5.2 Project Catalog Applications

Enterprise

Gov Advanced

Catalog applications are any third-party or open source applications that appear in the Catalog. These can be DKP applications provided by D2iQ for use in your environment, or [Custom Applications \(see page 704\)](#) that can be used but are not supported by D2iQ.

- [Project Catalog Applications Upgrades \(see page 695\)](#)
- [Project-level DKP Applications \(see page 696\)](#)
- [Use Custom Resources with Workspace Catalog Applications \(see page 704\)](#)
- [Custom Project Applications \(see page 704\)](#)

6.4.5.2.1 Project Catalog Applications Upgrades

Enterprise

Gov Advanced

Before upgrading your catalog applications, verify the current and supported versions of the application. Also, keep in mind the distinction between Platform applications and Catalog applications. Platform applications are deployed and upgraded as a set for each cluster or workspace. Catalog applications are deployed separately, so that you can deploy and upgrade them individually for each project.



Catalog applications must be upgraded to the latest version BEFORE upgrading the Konvoy component for Managed clusters or Kubernetes version for attached clusters.

6.4.5.2.1.1 Upgrade with the UI

Follow these steps to upgrade an application from the DKP UI:

1. From the top menu bar, select your target workspace.
2. From the side menu bar, select **Projects**.
3. Select your target project.
4. Select **Applications** from the project menu bar.
5. Select the three dot button from the bottom-right corner of the desired application tile, and then select **Edit**.
6. Select the **Version** drop-down, and select a new version. This drop-down will only be available if there is a newer version to upgrade to.
7. Select **Save**.

6.4.5.2.1.2 Upgrade with the CLI

1. To see what app(s) and app versions are available to upgrade, run the following command:

NOTE: The `APP ID` column displays the available apps and the versions available to upgrade.

```
kubectl get apps -n ${PROJECT_NAMESPACE}
```

2. Run the following command to upgrade an application from the DKP CLI:

```
dkp upgrade catalogapp <appdeployment-name> --workspace=my-workspace --
project=my-project --to-version=<version.number>
```

As an example, the following command upgrades the Kafka Operator application, named `kafka-operator-abc`, in a workspace to version `0.25.1`:

```
dkp upgrade catalogapp kafka-operator-abc --workspace=my-workspace --to-version=0.25.1
```

- Platform applications cannot be upgraded on a one-off basis, and must be upgraded in a single process for each workspace. If you attempt to upgrade a platform application with these commands, you receive an error and the application is not upgraded.

6.4.5.2.2 Project-level DKP Applications

Enterprise

Gov Advanced

DKP applications are catalog applications provided by D2iQ for use in your environment.

Some DKP workspace catalog applications will provision `CustomResourceDefinitions`, which allow you to deploy Custom Resources to a Project. See your DKP workspace catalog application's documentation for instructions.

6.4.5.2.2.1 Zookeeper in a Project

Enterprise

Gov Advanced

Deploying ZooKeeper in a project

Get started

To get started with creating ZooKeeper clusters in your project namespace, you first need to deploy the [ZooKeeper operator](#) (see page 673) in the workspace where the project exists.

After you deploy the ZooKeeper operator, you can create ZooKeeper Clusters by applying a `ZookeeperCluster` custom resource **on each attached cluster** in a project's namespace.

A [Helm chart](#)⁵³⁰ exists in the ZooKeeper operator repository that can assist with deploying ZooKeeper clusters.

Example Deployment

- If you need to manage these custom resources across all clusters in a project, it is recommended you use [Project Deployments](#) (see page 713) which enables you to leverage GitOps to deploy the resources. Otherwise, you will need to create the resources manually in each cluster.

Follow these steps to deploy a ZooKeeper cluster in a project namespace. This procedure results in a running ZooKeeper cluster, ready for use in your project's namespace.

1. Set the `PROJECT_NAMESPACE` environment variable to the name of your project's namespace:

```
export PROJECT_NAMESPACE=<project namespace>
```

2. Create a ZooKeeper Cluster custom resource in your project namespace

```
kubectl apply -f - <<EOF
apiVersion: zookeeper.pravega.io/v1beta1
kind: ZookeeperCluster
metadata:
  name: zookeeper
  namespace: ${PROJECT_NAMESPACE}
spec:
  replicas: 1
EOF
```

3. Check the status of your ZooKeeper cluster using `kubectl` :

```
kubectl get zookeeperclusters -n ${PROJECT_NAMESPACE}
```

NAME	REPLICAS	READY	REPLICAS	VERSION	DESIRED	VERSION	INTERNAL
ENDPOINT	EXTERNAL	ENDPOINT	AGE				
zookeeper	1	1		0.2.15	0.2.15		10.100.200
.18:2181	N/A		94s				

⁵³⁰ <https://github.com/pravega/zookeeper-operator/tree/master/charts/zookeeper>

Delete ZooKeeper clusters

Follow these steps to delete the Zookeeper clusters.

1. View `ZookeeperClusters` in all namespaces:

```
kubectl get zookeeperclusters -A
```

2. Delete a specific `ZookeeperCluster` :

```
kubectl -n ${PROJECT_NAMESPACE} delete zookeepercluster <name of zookeepercluster>
```

Resources

- [ZooKeeper Operator Documentation](#)⁵³¹
- [ZooKeeper Cluster Helm Chart](#)⁵³²
- [ZooKeeper Documentation](#)⁵³³

6.4.5.2.2.2 Kafka in a Project

Enterprise

Gov Advanced

Deploying Kafka in a project

Prerequisites

To get started with creating and managing a Kafka Cluster in a project, you must:

- Deploy the [Kafka operator](#) (see page 668) in the workspace where the project exists.
- Deploy the [ZooKeeper operator](#) (see page 673) in the workspace where the project exists.
- Deploy [Zookeeper in a Project](#) (see page 699) in the same project where you want to enable Kafka.

Get Started

After deploying the Kafka operator, create Kafka Clusters by applying a `KafkaCluster` custom resource **on each attached cluster** in a project's namespace. Refer to the [Kafka operator repository](#)⁵³⁴ for examples of the custom resources and their configurations.


⁵³¹ <https://github.com/pravega/zookeeper-operator>

⁵³² <https://github.com/pravega/zookeeper-operator/tree/master/charts/zookeeper>

⁵³³ <https://zookeeper.apache.org/documentation>

⁵³⁴ <https://github.com/banzaicloud/koperator/tree/master/config/samples>

Example Deployment


-  If you need to manage these custom resources across all clusters in a project, it is recommended you use [Project Deployments](#) (see page 713) which enables you to leverage GitOps to deploy the resources. Otherwise, you will need to create the custom resources manually in each cluster.

This example deployment walks you through deploying a Kafka cluster in a project namespace. The result of this procedure is a running ZooKeeper cluster and Kafka cluster ready for use in your project's namespace.

1. Ensure you deployed [Zookeeper in a Project](#) (see page 699).
2. Set the `PROJECT_NAMESPACE` environment variable to the name of your project's namespace:

```
export PROJECT_NAMESPACE=<project namespace>
```


3. Obtain the Kafka Operator version you deployed in the workspace:

-  Replace `<target_namespace>` with the namespace where you deployed Kafka:

```
kubectl get appdeployments.apps.kommander.d2iq.io -n <target_namespace> -o
template="{{{ .spec.appRef.name }}" kafka-operator
```



The output prints the Kafka Operator version.

4. Use the Kafka Operator version to download the `simplekafkacluster.yaml` file you require:

-  In the following URL, replace `/v0.25.1/` with the Kafka version you obtained in the previous step and download the file.

<https://raw.githubusercontent.com/banzaicloud/koperator/v0.25.1/config/samples/simplekafkacluster.yaml>

5. Open and edit the downloaded file to use the correct Zookeeper Cluster address:

-  Replace `<project_namespace>` with the target project namespace.
-  Replace `<zookeeper-server-client.zookeeper:2181 >` with `zookeeper-client.<project_namespace>:2181`

```
zkAddresses:
```

```
- "zookeeper-client.<project_namespace>:2181"
```

6. Apply the `KafkaCluster` configuration to your project's namespace:

```
kubectl apply -n ${PROJECT_NAMESPACE} -f simplekafkacluster.yaml
```

7. Check the status of your Kafka cluster using `kubectl`:

```
kubectl -n ${PROJECT_NAMESPACE} get kafkaclusters
```

The output should look similar to this:

NAME	CLUSTER STATE	CLUSTER ALERT COUNT	LAST SUCCESSFUL UPGRADE
kafka	ClusterRunning	0	0
	UPGRADE ERROR COUNT	AGE	
	79m		

With both the ZooKeeper cluster and Kafka cluster running in your project's namespace, refer to the [Kafka Operator documentation](#)⁵³⁵ for information on how to test and verify they are working as expected in. When performing those steps, ensure you substitute: `zookeeper-client.<project namespace>:2181` anywhere that the zookeeper client address is mentioned.

Delete Kafka Custom Resources

Follow these steps to delete the Kafka custom resources.

1. View all Kafka resources in the cluster:

```
kubectl get kafkaclusters -A
kubectl get kafkausers -A
Kubectl get kafkatopics -A
```

2. Delete a `KafkaCluster` example:

```
kubectl -n ${PROJECT_NAMESPACE} delete kafkacluster <name of KafkaCluster>
```

Resources

- [Kafka Operator Documentation](#)⁵³⁶

⁵³⁵ <https://banzaicloud.com/docs/supertubes/kafka-operator/test/>

⁵³⁶ <https://banzaicloud.com/docs/supertubes/kafka-operator/>

- [Kafka Operator GitHub Repository](#)⁵³⁷
- [Sample Kafka Operator Custom Resources](#)⁵³⁸
- [Apache Kafka Documentation](#)⁵³⁹

6.4.5.2.2.3 Spark in a Project

Deploying Spark in a project

To run your Spark workloads with Spark Operator, apply the Spark Operator specific custom resources. The Spark Operator works with the following kinds of custom resources:

- `SparkApplication`
- `ScheduledSparkApplication`

See [Spark Operator API documentation](#)⁵⁴⁰ for more details.



If you need to manage these custom resources and RBAC resources across all clusters in a project, it is recommended you use [Project Deployments](#) (see page 713) which enables you to leverage GitOps to deploy the resources. Otherwise, you will need to create the resources manually in each cluster.

Prerequisites

Follow these steps:

1. Deploy your Spark Operator. See the [Spark Operator](#) (see page 670) documentation for more information.
2. Ensure the necessary RBAC resources referenced in your custom resources exist, otherwise the custom resources can fail. See the [Spark Operator documentation](#)⁵⁴¹ for details.
 - This is an example of commands for you to create the RBAC resources needed in your project namespace:

```
export PROJECT_NAMESPACE=<project namespace>

kubectl apply -f - <<EOF
apiVersion: v1
```

⁵³⁷ <https://github.com/banzaicloud/koperator>

⁵³⁸ <https://github.com/banzaicloud/koperator/tree/master/config/samples>

⁵³⁹ <https://kafka.apache.org/documentation/>

⁵⁴⁰ <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/docs/api-docs.md>

⁵⁴¹ <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/docs/quick-start-guide.md#about-the-spark-job-namespace>


```

kind: ServiceAccount
metadata:
  name: spark-service-account
  namespace: ${PROJECT_NAMESPACE}
---
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  namespace: ${PROJECT_NAMESPACE}
  name: spark-role
rules:
- apiGroups: [""]
  resources: ["pods"]
  verbs: ["*"]
- apiGroups: [""]
  resources: ["services"]
  verbs: ["*"]
- apiGroups: [""]
  resources: ["configmaps"]
  verbs: ["*"]
- apiGroups: [""]
  resources: ["persistentvolumeclaims"]
  verbs: ["*"]
---
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  name: spark-role-binding
  namespace: ${PROJECT_NAMESPACE}
subjects:
- kind: ServiceAccount
  name: spark-service-account
  namespace: ${PROJECT_NAMESPACE}
roleRef:
  kind: Role
  name: spark-role
  apiGroup: rbac.authorization.k8s.io
EOF

```

Deploy a Simple SparkApplication

Follow these steps:

1. Create your [Project](#) (see page 689) if you don't already have one.
2. Set the `PROJECT_NAMESPACE` environment variable to the name of your project's namespace:

```
export PROJECT_NAMESPACE=<project namespace>
```

3. Set the SPARK_SERVICE_ACCOUNT environment variable to one of the following:
 - a. `${PROJECT_NAMESPACE}` , if you skipped the step in [Prerequisites](#) (see page 704) to create RBAC resources.

```
# This service account is automatically created when you create a
project and has access to everything in the project namespace.
export SPARK_SERVICE_ACCOUNT=${PROJECT_NAMESPACE}
```

- b. Or set to `spark-service-account`

```
export SPARK_SERVICE_ACCOUNT=spark-service-account
```

4. Apply the SparkApplication custom resource in your project namespace

```
kubectl apply -f - <<EOF
apiVersion: "sparkoperator.k8s.io/v1beta2"
kind: SparkApplication
metadata:
  name: pyspark-pi
  namespace: ${PROJECT_NAMESPACE}
spec:
  type: Python
  pythonVersion: "3"
  mode: cluster
  image: "gcr.io/spark-operator/spark-py:v3.1.1"
  imagePullPolicy: Always
  mainApplicationFile: local:///opt/spark/examples/src/main/python/pi.py
  sparkVersion: "3.1.1"
  restartPolicy:
    type: OnFailure
    onFailureRetries: 3
    onFailureRetryInterval: 10
    onSubmissionFailureRetries: 5
    onSubmissionFailureRetryInterval: 20
  driver:
    cores: 1
    coreLimit: "1200m"
    memory: "512m"
    labels:
      version: 3.1.1
      serviceAccount: ${SPARK_SERVICE_ACCOUNT}
  executor:
    cores: 1
    instances: 1
    memory: "512m"
    labels:
      version: 3.1.1
EOF
```

Clean up

Follow these steps:

1. View `SparkApplications` in all namespaces:

```
kubectl get sparkapp -A
```

2. Deleting a specific `SparkApplication`:

```
kubectl -n ${PROJECT_NAMESPACE} delete sparkapp <name of sparkapplication>
```

6.4.5.2.3 Use Custom Resources with Workspace Catalog Applications

Enterprise

Gov Advanced

Some workspace catalog applications provision some `CustomResourceDefinition`, which allow you to deploy Custom Resources. Refer to your workspace catalog application's documentation for instructions.

6.4.5.2.4 Custom Project Applications

Enterprise

Gov Advanced

Custom applications are third-party applications you have added to the Kommander Catalog.

Custom applications are any third-party applications that are not provided in the DKP Application Catalog. Custom applications can leverage applications from the DKP Catalog or be fully-customized. There is no expectation of support by D2iQ for a Custom application. Custom applications can be deployed on Konvoy clusters or on any D2iQ supported 3rd party Kubernetes distribution.

- [Projects - Create a Git Repository \(see page 704\)](#)
- [Projects - Git repository structure \(see page 706\)](#)
- [Projects - Application Metadata \(see page 708\)](#)
- [Enable a Custom Application from the Project Catalog \(see page 710\)](#)

6.4.5.2.4.1 Projects - Create a Git Repository

Enterprise

Gov Advanced

Create a Git Repository in the Project namespace

Use the CLI to create the GitRepository resource and add a new repository to your Project.

1. Refer to [air-gapped setup instructions](#) (see page 1397), if you are running in air-gapped environment.
2. Set the `PROJECT_NAMESPACE` environment variable to the name of your project's namespace:

```
export PROJECT_NAMESPACE=<project_namespace>
```

3. Adapt the URL of your Git repository.

```
kubectl apply -f - <<EOF
apiVersion: source.toolkit.fluxcd.io/v1beta1
kind: GitRepository
metadata:
  name: example-repo
  namespace: ${PROJECT_NAMESPACE}
spec:
  interval: 1m0s
  ref:
    branch: <your-target-branch-name> # e.g., main
    timeout: 20s
    url: https://github.com/<example-org>/<example-repo>
EOF
```

4. Ensure the status of the `GitRepository` signals a ready state:

```
kubectl get gitrepository example-repo -n ${PROJECT_NAMESPACE}
```

The repository commit also displays the ready state:

NAME	URL	READY
STATUS		AGE
example-repo	https://github.com/example-org/example-repo	True
	Fetches revision: master/6c54bd1722604bd03d25dcac7a31c44ff4e03c6a	11m

For more information on the GitRepository resource fields and how to make Flux aware of credentials required to access a private Git repository, see the [Flux documentation](#)⁵⁴².

Troubleshoot

To troubleshoot issues with adding the GitRepository, review the following logs:

⁵⁴² <https://fluxcd.io/docs/components/source/gitrepositories/>

```
kubectl -n kommander-flux logs -l app=source-controller
[...]
kubectl -n kommander-flux logs -l app=kustomize-controller
[...]
kubectl -n kommander-flux logs -l app=helm-controller
[...]
```

Related Information

- [Flux](#)⁵⁴³
- [Flux docs](#)⁵⁴⁴

6.4.5.2.4.2 Projects - Git repository structure

Enterprise

Gov Advanced

Git repositories must be structured in a specific manner for defined applications to be processed by Kommander.

You must structure your git repository based on the following guidelines, for your applications to be processed properly by Kommander so that they can be deployed.

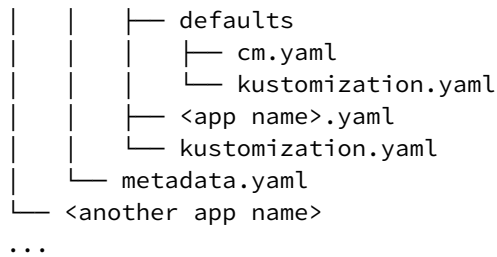
Git Repository Directory Structure

Use the following basic directory structure for your git repository:

```
├─ helm-repositories
│   ├── <helm repository 1>
│   │   ├── kustomization.yaml
│   │   └── <helm repository name>.yaml
│   └── <helm repository 2>
│       ├── kustomization.yaml
│       └── <helm repository name>.yaml
└─ services
    ├── <app name>
    │   ├── <app version1> # semantic version of the app helm chart. e.g., 1.2.3
    │   │   ├── defaults
    │   │   │   ├── cm.yaml
    │   │   │   └── kustomization.yaml
    │   │   └── <app name>.yaml
    │   └── kustomization.yaml
    └── <app version2> # another semantic version of the app helm chart. e.g.,
        2.3.4
```

⁵⁴³ <https://fluxcd.io/>

⁵⁴⁴ <https://fluxcd.io/docs>



- Define applications in the `services/` directory.
- You can define multiple versions of an application, under different directories nested under the `services/<app name>/` directory.
- Define application manifests, such as a [HelmRelease](#)⁵⁴⁵, under each versioned directory `services/<app name>/<version>/` in `<app name>.yaml` which is listed in the `kustomization.yaml` Kubernetes Kustomization file. For more information, see the [Kubernetes Kustomization docs](#)⁵⁴⁶.
- Define the default values ConfigMap for HelmReleases in the `services/<app name>/<version>/defaults` directory, accompanied by a `kustomization.yaml` Kubernetes Kustomization file pointing to the ConfigMap file.
- Define the `metadata.yaml` of each application under the `services/<app name>/` directory. For more information, see the [Application Metadata docs](#) (see page 683).

See [the DKP Catalog repository](#)⁵⁴⁷ for an example of how to structure custom catalog Git repositories.

Helm Repositories

You must include the `HelmRepository` that is referenced in each `HelmRelease`'s `Chart` spec.

Each `services/<app name>/<version>/kustomization.yaml` must include the path of the YAML file that defines the `HelmRepository`. For example:

```

# services/<app name>/<version>/kustomization.yaml
apiVersion: kustomize.config.k8s.io/v1beta1
kind: Kustomization
resources:
  - <app name>.yaml
  - ../../../../helm-repositories/<helm repository 1>
  
```

For more information, see the flux documentation about [HelmRepositories](#)⁵⁴⁸.

Substitution Variables

⁵⁴⁵ <https://fluxcd.io/docs/components/helm/helmreleases/>

⁵⁴⁶ <https://kubectl.docs.kubernetes.io/references/kustomize/kustomization/>

⁵⁴⁷ <https://github.com/mesosphere/dkp-catalog-applications>

⁵⁴⁸ <https://fluxcd.io/docs/components/source/helmrepositories/>

Some [substitution variables](#)⁵⁴⁹ are provided.

- `${releaseName}` : For each App deployment, this variable is set to the `AppDeployment` name. Use this variable to prefix the names of any resources that are defined in the application directory in the Git repository so that multiple instances of the same application can be deployed. If you create resources without using the `releaseName` prefix (or suffix) in the name field, there can be conflicts if the same named resource is created in that same namespace.
- `${releaseNamespace}` : The namespace of the Project.
- `${workspaceNamespace}` : The namespace of the Workspace that the Project belongs to.

Related Information

- [Flux](#)⁵⁵⁰
- [Flux docs](#)⁵⁵¹
- [Getting started with Flux guide](#)⁵⁵²

6.4.5.2.4.3 Projects - Application Metadata

Enterprise

Gov Advanced

To display more information about custom applications in the UI, define a `metadata.yaml` file for each application in the git repository.

You can define how custom applications display in the DKP UI by defining a `metadata.yaml` file for each application in your git repository. You must define this file at `services/<application>/metadata.yaml` for it to process correctly.

You can define the following fields:

Field	Default	Description
<code>displayName</code>	falls back to App ID	Display name of the application for the UI.
<code>description</code>	""	Short description, should be a sentence or two, displayed in the UI on the application card.

⁵⁴⁹ <https://fluxcd.io/docs/components/kustomize/kustomization/#variable-substitution>

⁵⁵⁰ <https://fluxcd.io/>

⁵⁵¹ <https://fluxcd.io/docs>

⁵⁵² <https://fluxcd.io/docs/get-started/>

Field	Default	Description
category	general	1 or more categories for this application. Categories are used to group applications in the UI.
overview		Markdown overview used on the application detail page in the UI.
icon		Base64 encoded icon SVG file used for application logos in the UI.
scope	project	List of scopes, can be set only to project or workspace currently.

None of these fields are required for the application to display in the UI.

Here is an example `metadata.yaml` file:

```

displayName: Prometheus Monitoring Stack
description: Stack of applications that collect metrics and provides visualization
and alerting capabilities. Includes Prometheus, Prometheus Alertmanager and Grafana.
category:
  - monitoring
overview: >
  # Overview
  A stack of applications that collects metrics and provides visualization and
  alerting capabilities. Includes Prometheus, Prometheus Alertmanager and Grafana.

  ## Dashboards
  By deploying the Prometheus Monitoring Stack, the following platform applications
  and their respective dashboards are deployed. After deployment to clusters in a
  workspace, the dashboards are available to access from a respective cluster's detail
  page.

  ### Prometheus

  A software application for event monitoring and alerting. It records real-time
  metrics in a time series database built using a HTTP pull model, with flexible and
  real-time alerting.

  - [Prometheus Documentation - Overview](https://prometheus.io/docs/introduction/
  overview/)

  ### Prometheus Alertmanager

```


A Prometheus component that enables you to configure and manage alerts sent by the Prometheus server and to route them to notification, paging, and automation systems.

- [Prometheus Alertmanager Documentation - Overview](<https://prometheus.io/docs/alerting/latest/alertmanager/>)

Grafana

A monitoring dashboard from Grafana that can be used to visualize metrics collected by Prometheus.

- [Grafana Documentation](<https://grafana.com/docs/>)

icon:

```
PHN2ZyB4bWxucz0iaHR0cDovL3d3dy53My5vcmcvMjAwMC9zdmc iIHZpZXdCb3g9IjAgMCAzMDAgMzAwIiBzd
HlsZT0iZW5hYm xLLWJhY2tncm91bmQ6bmV3IDAgMCAzMDAgMzAwIiB4bWw6c3BhY2U9InByZXNlcnZlIj48cG
F0aCBkPSJNMTUwIDUwQzk0LjggNTAgNTAgOTQuOCA1MCAxNTBzNDQuOCAxMDAgMTAwIDEwMCAxMDAtNDQuOCA
xMDAtMTAwUzIwNS4yIDUwIDE1MCA1MHptMCAxODcuMmMtMTUuNyAwLTI4LjUtMTAuNS0yOC41LTIzLjRoNTYu
OWMuMSAxMi45L TEyLjcgMjMuNC0yOC40IDIZLjR6bTQ3LTMxLjJoLTk0di0xN2g5NHYxN3ptLS4zLTI1LjloL
TkzLjRjLS4zLS40LS42LS43LS45L TEuMS05LjYtMTEuNy0xMS45L TE3LjgtMTQuMS0yNCAwLS4yIDExLjcgMi
40IDIWIDQuMyAwIDA gNC4zIDEGMTAuNSAyLjEtNi03LTkuNi0xNi05LjYtMjUuMSAwLTlWIDE1LjQtMzcuNiA
5LjgtNTEuNyA1LjQuNCAxMS4yIDExLjQgMTEuNiAyOC41IDUuNy03LjkgOC4xLTIyLjQgOC4xLTMxLjMgMC05
LjI gNi4xL TE5LjkgMTIuMS0yMC4yL TUuNCA4LjkgMS40IDE2LjUgNy40IDM1LjUgMi4zIDcuMSAyIDE5LjEgM
y43IDI2LjcuNi0xNS44IDMuMy0zOC44IDExLjMtNDYuNy00LjQgMTAgLjcgMjUuNSA0LjEgMjguNSA1LjYgOS
43IDkgMTcuMSA5IDMxIDA gOS4zLTMuNCAxOC4xLTKuMyAyNSA2LjYtMjUuMS4yIDExLjItMi40IDExLjItMi40bDI
xLjQtNC4yYy4xIDA tMyAxMi44L TE0LjkgMjUuMXoiIHNoeWwPSJmaWxsOiNmODQzMTEiLz48L3N2Zz4=
```

6.4.5.2.4.4 Enable a Custom Application from the Project Catalog

Enterprise

Gov Advanced

Enable a Custom Application from the Project Catalog

After creating a GitRepository, you can either use the DKP UI or the CLI to enable your custom applications.



From within a project, you can enable applications to deploy. Verify that an application has successfully deployed via the CLI.

Enable the Application using the UI

1. From the top menu bar, select your target workspace.
2. Select **Projects** from the sidebar menu.
3. Select your project from the list.
4. Select the **Applications** tab to browse the available applications from your configured repositories.

5. Select the three dot button from the bottom-right corner of the desired application tile, and then select **Enable**.
6. If available, select a version from the drop-down menu. This drop-down menu will only be visible if there is more than one version.
7. (Optional) If you want to override the default configuration values, copy your values content into the text editor under **Configure Service** or just upload your YAML file that contains the values:

```
someField: someValue
```

8. Confirm the details are correct, and then select the **Enable** button.

For all applications, you must provide a display name and an ID which is automatically generated based on what you enter for the display name, unless or until you edit the ID directly. The ID must be compliant with [Kubernetes DNS subdomain name validation rules](#)⁵⁵³.

Alternately, you can use the [CLI](#) (see page 0) to enable your custom applications in the section below.

Enable the Application using the CLI

Follow these steps:

1. Set the `PROJECT_NAMESPACE` environment variable to the name of the project's namespace:

```
export PROJECT_NAMESPACE=<project_namespace>
```

2. Get the list of available applications to enable using the following command:

```
kubectl get apps -n ${PROJECT_NAMESPACE}
```

3. Enable one of the supported applications from the list with an `AppDeployment` resource.
4. Within the `AppDeployment` resource. Provide the `appRef` and application version to specify which `App` is deployed:

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: my-custom-app
  namespace: ${PROJECT_NAMESPACE}
spec:
  appRef:
    name: custom-app-0.0.1
```

⁵⁵³ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#dns-subdomain-names>

```
kind: App
EOF
```

 The `appRef.name` must match the app `name` from the list of available applications.

Enable an Application with a Custom Configuration using the CLI

Follow these steps:

1. Provide the name of a `ConfigMap` in the `AppDeployment`, which provides custom configuration on top of the default configuration:

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: my-custom-app
  namespace: ${PROJECT_NAMESPACE}
spec:
  appRef:
    name: custom-app-0.0.1
    kind: App
  configOverrides:
    name: my-custom-app-overrides
EOF
```

2. Create the `ConfigMap` with the name provided in the step above, with the custom configuration:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${PROJECT_NAMESPACE}
  name: my-custom-app-overrides
data:
  values.yaml: |
    someField: someValue
EOF
```

Kommander waits for the `ConfigMap` to be present before enabling the `AppDeployment` to the attached clusters in the Project.

Verify Applications

After completing the previous steps, your applications are enabled. Connect to the attached cluster and check the `HelMReLeases` to verify the deployments:

```
kubectL get helmreleases -n ${PROJECT_NAMESPACE}
```

The output looks similar to this:

NAMESPACE	NAME	READY	STATUS
AGE			
project-test-vjsfq 7m3s	my-custom-app	True	Release reconciliation succeeded

Upgrade Custom Applications

You must maintain your custom applications manually. When upgrading DKP, ensure you validate for compatibility issues any custom applications you run against the current version of Kubernetes. We recommend upgrading to the latest compatible application versions as soon as possible.

6.4.5.3 Project AppDeployments

Enterprise

Gov Advanced

An `AppDeployment` is a Custom Resource created by DKP with the purpose of deploying applications (platform, DKP catalog and custom applications). For more information about these Custom Resources and how to customize them, refer to the [AppDeployments](#) (see page 625) documentation.

6.4.6 Project Deployments

Enterprise

Gov Advanced

Use Project Deployments to manage GitOps based Continuous Deployments

You can configure Kommander Projects with GitOps-based Continuous Deployments for federation of your Applications to associated clusters of the project. This is backed by Flux, which enables software and applications to be continuously deployed (CD) using GitOps processes. GitOps enables the application to be deployed as per a manifest that is stored in a Git repository. This ensures that the application deployment can be automated, audited, and declaratively deployed to the infrastructure.

- [What is GitOps?](#) (see page 714)
- [Continuous Delivery with GitOps](#) (see page 714)

- [Continuous Deployment \(see page 724\)](#)
- [Project Deployments Troubleshooting \(see page 727\)](#)
- [View Helm Releases \(see page 728\)](#)

6.4.6.1 What is GitOps?

GitOps is a modern software deployment strategy. The configuration that describes how your application is deployed to a cluster are stored in a Git repository. The configuration is continuously synchronized from the Git repository to the cluster, ensuring that the specified state of the cluster always matches what is defined in the “GitOps” Git repository.

The benefits of using a GitOps deployment strategy are:

- Familiar, collaborative change and review process. Engineers are intimately familiar with Git-based workflows: branches, pull requests, code reviews, etc. GitOps leverages this experience to control the deployment of software and updates to catch issues early.
- Clear change log and audit trail. The Git commit log serves as an audit trail to answer the question: “who changed what, and when?” Having such information available, you can contact the right people when fixing or prioritizing a production incident to determine the why and correctly resolve the issue as quickly as possible. Additionally, Kommander’s CD component (Flux CD) maintains a separate audit trail in the form of Kubernetes Events, as changes to a Git repository don’t include exactly when those changes were deployed.
- Avoid configuration drift. The scope of manual changes made by operators expands over time. It soon becomes difficult to know which cluster configuration is critical and which is left over from temporary workarounds or live debugging. Over time, changing a project configuration or replicating a deployment to a new environment becomes a daunting task. GitOps supports simple, reproducible deployment to multiple different clusters by having a single source of truth for cluster and application configuration.

That said, there are some cases when live debugging is necessary in order to resolve an incident in the minimum amount of time. In such cases, pull-request-based workflow adds precious time to resolution for critical production outages. Kommander’s CD strategy supports this scenario by letting you disable the auto sync feature. After auto sync is disabled, Flux will stop synchronizing the cluster state from the GitOps git repository. This lets you use `kubectl`, `helm`, or whichever tool you need to resolve the issue.

6.4.6.2 Continuous Delivery with GitOps

DKP enables software and applications to be continuously delivered (CD) using GitOps processes. GitOps enables you to deploy an application according to a manifest that is stored in a Git repository. This ensures that the application deployment can be automated, audited, and declaratively deployed to the infrastructure.

This section contains step-by-step tutorials for performing some common deployment-related tasks using DKP. All tutorials begin with a Prerequisites section that contains links to any steps that need to be taken first. This means you can visit any tutorial to get started.

- [Store secrets in GitOps repository using SealedSecrets \(see page 715\)](#)
- [Deploy a Sample App from DKP GitOps \(see page 718\)](#)

6.4.6.2.1 Store secrets in GitOps repository using SealedSecrets

6.4.6.2.1.1 Securely managing secrets in a GitOps workflow using SealedSecrets

For security reasons, Kubernetes secrets are usually the only resource that cannot be managed with a GitOps workflow. Instead of managing secrets outside of GitOps and having to use a third-party tool like Vault, SealedSecrets provides a way to keep all the advantages of using a GitOps workflow while avoiding exposing secrets. SealedSecrets is composed of two main components:

- A CLI (Kubeseal) to encrypt secrets.
- A cluster-side controller to decrypt the sealed secrets into regular Kubernetes secrets. Only this controller can decrypt sealed secrets, not even the original author.

This tutorial describes how to install these two components, configure the controller, and add or remove sealed secrets.

6.4.6.2.1.2 Set up

These instructions are used as an example. For instructions on the latest release, see the [release page](#)⁵⁵⁴. For full documentation on SealedSecrets, see the [GitHub repo](#)⁵⁵⁵.

6.4.6.2.1.3 Install Kubeseal CLI to Encrypt Your Secrets

- On MacOS

```
brew install kubeseal
```

- On Linux

```
wget https://github.com/bitnami-labs/sealed-secrets/releases/download/v0.18.1/kubeseal-0.18.1-linux-amd64.tar.gz -O kubeseal
sudo install -m 755 kubeseal /usr/local/bin/kubeseal
```

6.4.6.2.1.4 Install the SealedSecrets Controller on Your Cluster

This controller will be able to decrypt SealedSecrets and create Kubernetes secrets.

1. Create the controller:

```
kubectl apply -f https://github.com/bitnami-labs/sealed-secrets/releases/download/v0.18.1/controller.yaml
```

⁵⁵⁴ <https://github.com/bitnami-labs/sealed-secrets/releases>

⁵⁵⁵ <https://github.com/bitnami-labs/sealed-secrets>

2. Fetch the certificate that you will use to encrypt your secrets into sealed secrets:

```
kubeseal --fetch-cert > mycert.pem
```

3. Commit `mycert.pem` to your git repo.

6.4.6.2.1.5 Add a Secret

Secrets can be securely added to Git using sealed secrets:

1. Export the project namespace that you are using for your GitOps repository

```
export PROJECT_NAMESPACE=<your-project-with-gitops>
```

2. Create a Kubernetes secret and pipe it into `kubeseal`⁵⁵⁶ using the certificate `mycert.pem` that you fetched from the controller in the setup:

```
echo '---' >> secrets.yaml
kubectl create secret -n ${PROJECT_NAMESPACE} generic mysecret --dry-run=client
-o yaml --from-literal=my-secret=value | \
  kubeseal --format yaml --cert mycert.pem >> secrets.yaml
```

3. Go to the end of `secrets.yaml` where you just added your new sealed secret. Remove any “creationTimestamp” fields from the YAML.
4. Apply the `secrets.yaml` file to your namespace. If you do not have permission, commit your changes to the repo and let FluxCD apply the changes for you.

```
kubectl apply -f secrets.yaml
```

5. The sealed secret controller will then decrypt the sealed secret and generate a Kubernetes secret from it. Your secret got successfully created by running:

```
kubectl get secret mysecret -n ${PROJECT_NAMESPACE} -o yaml
```

6. If your sealed secret got created successfully but did not generate the matching secret, look at the logs of the controller:

```
kubectl logs -l=name=sealed-secrets-controller -n kube-system
```

⁵⁵⁶ <https://github.com/bitnami-labs/sealed-secrets#usage>

7. Commit `secrets.yaml` to your repo if you have not already done so in step 3.

6.4.6.2.1.6 Remove a Secret

1. Following the same example from above in “Adding a secret”, now remove the manifest for `mysecret` in `secrets.yaml` and commit those changes to the repo.
2. Delete the SealedSecret in the cluster:

```
kubectl delete SealedSecret -n ${PROJECT_NAMESPACE} mysecret
```

3. Delete the secret itself:

```
kubectl delete secret -n ${PROJECT_NAMESPACE} mysecret
```

6.4.6.2.1.7 Rotate the Controller’s Sealing Key

For added security, it is a good practice to rotate the key the controller uses to decrypt sealed secrets. By default, the controller generates a new key every 30 days. When this happens, you need to update the certificate you use to create sealed secrets by fetching the latest one:

```
kubeseal --fetch-cert > mycert.pem
```

 Do not forget to commit it back to the repo!

In a disaster case, let’s say your cluster gets destroyed, you would lose all your sealing keys, so you would not be able to recreate all the secrets from the sealed secrets in your GitOps repo. For this reason, you might want to back up the sealing keys. To do this every time a new sealing key is generated, run:

```
kubectl get secret -n kube-system -l sealedsecrets.bitnami.com/sealed-secrets-key -o yml > sealing-key
```

Then store `sealing-key` with the others in a safe location such as OneLogin Notes or Vault. To restore from a backup after a disaster, recreate all of the sealing keys with `kubectl apply -f sealing-key1 sealing-key2 ...` before starting the controller. If the controller was already started, restart it: `kubectl delete pod -n kube-system -l name=sealed-secrets-controller`

To disable sealing key rotation For example, configure the controller’s command in the pod template with `--key-renew-period=0`. See the following YAML file.


```
Pod Template:
Labels:          name=sealed-secrets-controller
Service Account: sealed-secrets-controller
Containers:
  sealed-secrets-controller:
    Image:        docker.io/bitnami/sealed-secrets-controller:v0.18.1
    Port:         8080/TCP
    Host Port:    0/TCP
    Command:     controller
                 --key-renew-period=0
```

If required, edit the controller's manifest with:

```
kubectl edit deployment.apps/sealed-secrets-controller -n kube-system
```

6.4.6.2.2 Deploy a Sample App from DKP GitOps

Enterprise


Gov Advanced

Use this procedure to deploy a sample `podinfo` application from DKP Enterprise GitOps.

6.4.6.2.2.1 Prerequisites

- [Enterprise DKP installed](#) (see page 108)
- [Kommander installed](#) (see page 1714)
- Github account and [personal access token](#)⁵⁵⁷
- [Add cluster to Kommander](#) (see page 752)
- [Setup Workspace and Projects](#) (see page 653)

6.4.6.2.2.2 Deployment Steps

 This procedure was run on an AWS cluster with DKP installed.

Follow these steps:

1. Ensure you are on the **Default Workspace** (or other workspace you have access to) so that you can create a project.

⁵⁵⁷ <https://docs.github.com/en/authentication/keeping-your-account-and-data-secure/creating-a-personal-access-token>

2. [Create a project \(see page 689\)](#).

In the working example we name the project **pod-info**. When you create a namespace, Kommander appends five alphanumeric characters. You can opt to select a target cluster for this project from one of the available attached clusters, and then this (**pod-info-xxxxx**) is the namespace used for deployments under the project, for example:

Name	Namespace	Description	Clusters	Cluster L
pod info	pod-info-xt2sz	No description provided		

3. [Optional] Create a secret in order to pull from the repository, for private repositories.
 - a. Select the Secrets tab and set up your secret according to the [Continuous Deployment documentation \(see page 724\)](#).
 - b. Add a **key** and **value** pair for the GitHub personal access token and then select **Create**.

CancelCreate SecretCreate

ID (name)

podinfo-secret

Must be unique within this project, and cannot be modified once saved.

Type

Opaque

Description

Data (1)

Key *✕

password

Value *

.....

[+ Add Key Value Pair](#)

- Verify that the secret `podinfo-secret` is created on the project namespace in the **managed or attached** cluster:

```
kubectl get secrets -n pod-info-xt2sz --kubeconfig=${CLUSTER_NAME}.conf
```

NAME	TYPE	DATA	AGE
default -token-k685t	kubernetes.io/service-account-token	3	94m
pod-info-xt2sz-token-p9k5z	kubernetes.io/service-account-token	3	94m
podinfo-secret	Opaque	1	1s
tls-root-ca	Opaque	1	93m

5. Select your project and then select the **CD** tab.
6. Add a GitOps Source, complete the required fields, and then **Save**.
There are several configurable options such as selecting the **Git Ref Type** but in this example we use the master branch. The **Path** value should contain where the manifests are located. Additionally, the **Primary Git Secret** is the secret (**podinfo-secret**) that you created in the previous step, if you need to access private repositories. This can be disregarded for public repositories.

Cancel
Create GitOps Source
Save

General

ID (name) *

Must be unique within this project, and cannot be modified once saved.

Repository URL *

Git Ref Type

Branch Name

If no value is entered, the Git reference will use the default branch for the repo.

Path

Primary Git Secret

Credentials used to fetch and administer the Git repository.

7.
 - a. Verify the status of `gitrepository` creation with this command (on the attached or managed cluster), and if **READY** is marked as **True**:

```
kubectl get gitrepository -A --kubeconfig=${CLUSTER_NAME}.conf
```

```

NAMESPACE      NAME      URL
AGE    READY  STATUS
kommander-flux  management  https://gitea-http.kommander.svc/kommander/
kommander.git  134m  True   stored artifact for revision 'main/
4fbee486076778c85e14f3196e49b8766e50e6ce'
pod-info-xt2sz  podinfo-source https://github.com/stefanprodan/podinfo
116m  True   stored artifact for revision 'master/
b3b00fe35424a45d373bf4c7214178bc36fd7872'

```

8. Verify the **Kustomization** with this command below (on the attached or managed cluster), and if **READY** is marked as **True**:

```
kubectl get kustomizations -n pod-info-xt2sz --kubeconfig=${CLUSTER_NAME}.conf
```

```

NAME              AGE    READY  STATUS
originalpodinfo  10m   True   Applied revision: master/
b3b00fe35424a45d373bf4c7214178bc36fd7872
podinfo-source   113m  True   Applied revision: master/
b3b00fe35424a45d373bf4c7214178bc36fd7872
project          116m  True   Applied revision: main/
4fbee486076778c85e14f3196e49b8766e50e6ce
project-tls-root-ca 117m  True   Applied revision: main/
4fbee486076778c85e14f3196e49b8766e50e6ce

```

Note the **port** so that you can use to verify if the app is deployed correctly (on the attached or managed cluster):

```
kubectl get deployments,services -n pod-info-xt2sz --kubeconfig=${CLUSTER_NAME}.conf
```

```

NAME              READY  UP-TO-DATE  AVAILABLE  AGE
deployment.apps/podinfo  2/2    2            2          118m

NAME      TYPE        CLUSTER-IP    EXTERNAL-IP  PORT(S)
AGE
service/podinfo  ClusterIP  10.99.239.120 <none>       9898/TCP,9999/TCP
118m

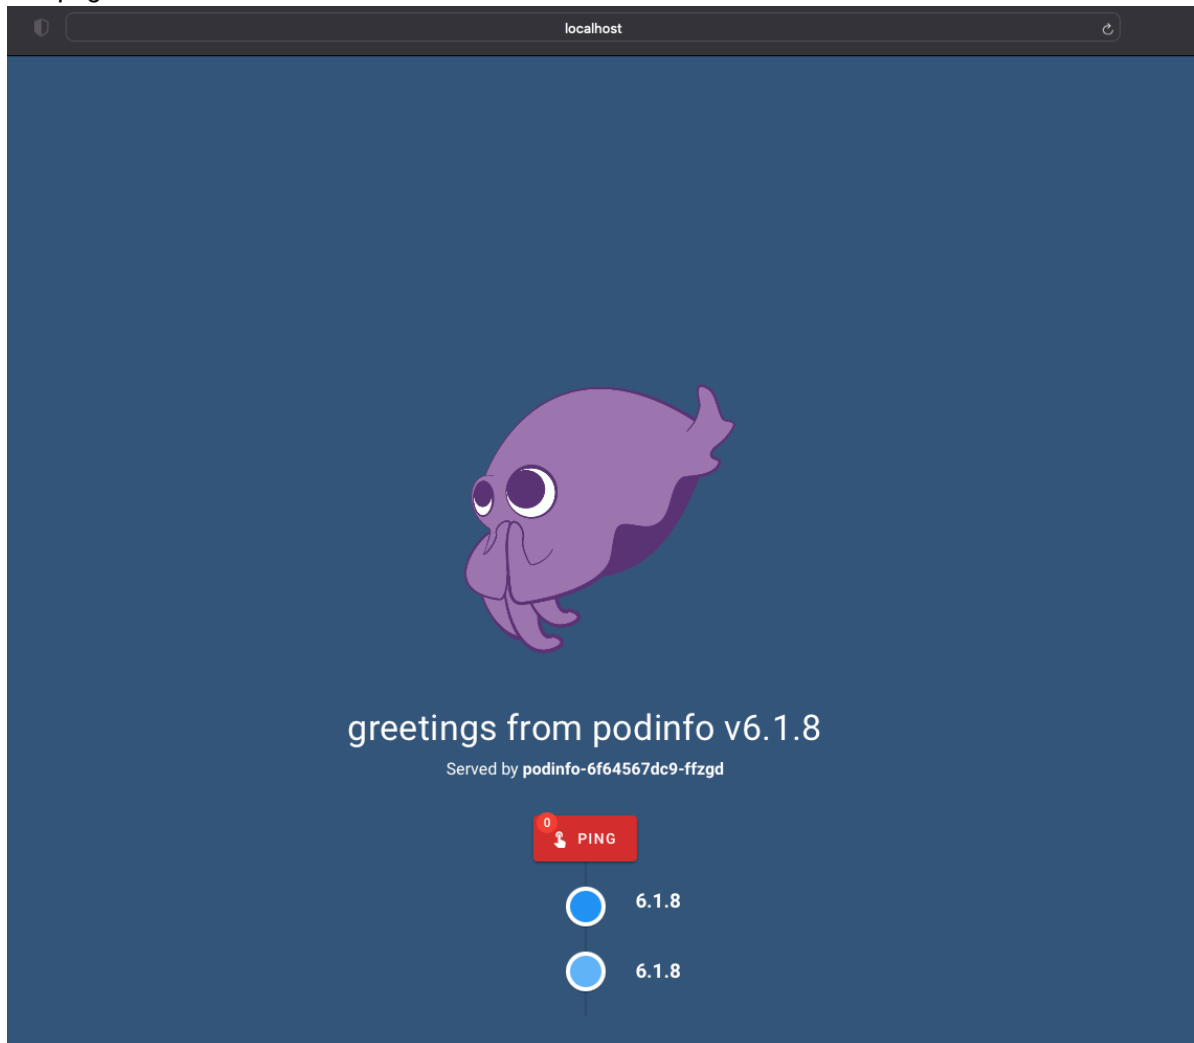
```

9. Port forward the **podinfo** service (port **9898**) to verify (on the attached or managed cluster):

```
kubectl port-forward svc/podinfo -n pod-info-xt2sz 9898:9898 --kubeconfig=${CLUSTER_NAME}.conf
```

```
Forwarding from 127.0.0.1:9898 -> 9898  
Forwarding from [::1]:9898 -> 9898  
Handling connection for 9898  
Handling connection for 9898  
Handling connection for 9898
```

10. Open a browser and type in **localhost:9898**. A successful deployment of the podinfo app gives you this page:



6.4.6.3 Continuous Deployment

Enterprise

Gov Advanced

After installing Kommander and [configuring your project and its clusters](#) (see page 745), navigate to the **Continuous Deployment (CD)** tab under your Project. Here you create a GitOps source which is a source code management (SCM) repository hosting the application definition. D2iQ recommends that you create a secret first then create a GitOps source accessed by the secret.

6.4.6.3.1 Prerequisites

You must have [Kommander installed](#) (see page 1422) to run this procedure.

6.4.6.3.2 Set up a Secret for Accessing GitOps

Create a secret that Kommander uses to deploy the contents of your GitOps repository:

ⓘ This dialog box creates a `types.kubefed.io/v1beta1, Kind=FederatedSecret` and this is not yet supported by DKP CLI. Use the GUI, as described above, to create a federated secret or create a `FederatedSecret` manifest and apply it to the project namespace. Learn more about [FederatedSecrets](#) (see page 736).

Kommander secrets (for CD) can be configured to support any of the following three authentication methods:

- HTTPS Authentication (described above)
- HTTPS self-signed certificates
- SSH Authentication

The following table describes the fields required for each authentication method:

HTTP Auth	HTTPS Auth (Self-signed)	SSH Auth
username	username	identity
password	password	identity.pub
	caFile	known_hosts

If you are using GitOps by using a GitHub repo as your source, you can create your secret with a [personal access token](#)⁵⁵⁸. Then, in the DKP UI, in your project, create a Secret, with a key:value pair of `password : <your-token-created-on-github>`. If you are using a GitHub personal access token, you do not need to have a key:value pair of `username : <your-github-username>`.

If you are using a secret with your GitHub username and your password, you will need one secret created in the DKP UI, with key:value pairs of `username : <your-github-username>` and `password : <your-github-password>`. **Note:** if you have multi-factor authentication turned on in your GitHub account, this will not work.



Using a token without a username is valid for GitHub, but other providers (such as GitLab) require both username and tokens.



If you are using a public GitHub repository, you do not need to use a secret.

6.4.6.3.3 Create GitOps Source

After the secret is created, you can view it in the `Secrets` tab. Configure the GitOps source accessed by the secret.



If using an SSH secret, the SCM repo URL needs to be an SSH address. It does not support SCP syntax. The URL format is `ssh://user@host:port/org/repository`.

It takes a few moments for the GitOps Source to be reconciled and the manifests from the SCM repository at the given path to be federated to attached clusters. After the sync is complete, manifests from GitOps source are created in attached clusters.

After a GitOps Source is created, there are various commands that can be executed from the CLI to check the various stages of syncing the manifests.

On the management cluster, check your `GitopsRepository` to ensure that the `CD` manifests have been created successfully.

```
kubectl describe gitopsrepositories.dispatch.d2iq.io -n<PROJECT_NAMESPACE> gitopsdemo
```

⁵⁵⁸ <https://docs.github.com/en/authentication/keeping-your-account-and-data-secure/creating-a-personal-access-token>


```
Name:          gitopsdemo
Namespace:    <PROJECT_NAMESPACE>
...
Events:
  Type      Reason                Age   From                                     Message
  ----      -
  Normal    ManifestSyncSuccess   1m7s  GitopsRepositoryController             manifests synced to
bootstrap repo
...
```

On the attached cluster, check for your `Kustomization` and `GitRepository` resources. The `status` field reflects the syncing of manifests:

```
kubectl get kustomizations.kustomize.toolkit.fluxcd.io -n<PROJECT_NAMESPACE>
<GITOPS_SOURCE_NAME> -oyaml
```

```
...
status:
  conditions:
  - reason: ReconciliationSucceeded
    status: "True"
    type: Ready
    ...
  ...
```

Similarly, with `GitRepository` resource:

```
kubectl get gitrepository.source.toolkit.fluxcd.io -n<PROJECT_NAMESPACE>
<GITOPS_SOURCE_NAME> -oyaml
```

```
...
status:
  conditions:
  - reason: GitOperationSucceed
    status: "True"
    type: Ready
    ...
  ...
```

If there are errors creating the manifests, those events are populated in the `status` field of the `GitopsRepository` resource on the management cluster, the `GitRepository` and `Kustomization` resources on the attached cluster(s), or both.

6.4.6.3.4 Suspend GitOps Source

There may be times when you need to suspend the auto-sync between the GitOps repository and the associated clusters. This *live debugging* may be necessary to resolve an incident in the minimum amount of time without the overhead of pull request based workflows.

To Suspend the GitOps Source from the DKP UI:

1. From the top menu bar, select your target workspace.
2. Select **Projects** from the sidebar menu.
3. Select your project from the list.
4. Select the **Continuous Deployment (CD)** tab.
5. Select the three dot button to the right of the desired GitOps Source.
6. Select **Suspend** to manually suspend the GitOps reconciliation.

This lets you use `kubectl`, `helm`, or another tool to resolve the issue. After the issue is resolved select **Resume** to sync the updated contents of the GitOps source to the associated clusters.

Similar to **Suspend/Resume**, you can use the **Delete** action to remove the GitOps source. Removing the GitOps source results in removal of all the manifests applied from the GitOps source.

You can have more than one GitOps Source in your Project to deploy manifests from various sources.

Kommander deployments are backed by FluxCD. Please refer to Flux [Source Controller](#)⁵⁵⁹ and [Kustomize controller](#)⁵⁶⁰ docs for advanced configuration and more examples.

6.4.6.4 Project Deployments Troubleshooting

Enterprise

Gov Advanced

- Events related to federation are stored in respective `FederatedGitRepository`, `FederatedKustomization`, or both resources.
- View the events and logs for `deployments/kommander-repository-controller` in Kommander namespace, if there are any unexpected errors.
- Enabling the Kommander repository controller for your project namespace causes a number of [related Flux controller components](#)⁵⁶¹ to deploy into the namespace. These are necessary for the proper operation of the repository controller and should not be removed.
- Ensure your GitOps repository does not contain any manifests that are cluster-scoped - for example, `Namespace`, `ClusterRole`, `ClusterRoleBinding`, etc. All of the manifests must be namespace-scoped.

⁵⁵⁹ <https://fluxcd.io/docs/components/source/>

⁵⁶⁰ <https://fluxcd.io/docs/components/kustomize/>

⁵⁶¹ <https://toolkit.fluxcd.io/components/>

- Ensure your GitOps repository does not contain any `HelmRelease` and `Kustomization` resources that are targeting a different namespace than the project namespace.

6.4.6.5 View Helm Releases

Enterprise

Gov Advanced

View Helm Releases for Continuous Deployments

In addition to viewing the current GitOps Sources, you can also view the current Helm Releases that have been deployed.

1. From the top menu bar, select your target workspace.
2. Select **Projects** from the sidebar menu.
3. Select your project from the list.
4. Select the **Continuous Deployment (CD)** tab.
5. Select the **Helm Releases** button.

All of the current Helm Release charts are displayed with their Chart Version and the names of the clusters. In this example *daily* is the name of the current cluster being displayed.

6.4.7 Project Role Bindings

Enterprise

Gov Advanced

Project Role Bindings grant access to a specified Project Role for a specified group of people

6.4.7.1 Configure Project Role Bindings - UI Method

Before you can create a Project Role Binding, ensure you have created a Group. A Kommander Group can contain one or several Identity Provider users or groups.

You can assign a role to this Kommander Group:

1. From the **Projects** page, select your project.
2. Select the **Role Bindings** tab, then select **Add Roles** next to the group you want.
3. Select the **Role**, or Roles, you want from the drop-down menu, and then select **Save**.

6.4.7.2 Configure Project Role Bindings - CLI Method

A Project role binding can also be created using kubectl:

```
cat << EOF | kubectl create -f -
apiVersion: workspaces.kommander.mesosphere.io/v1alpha1
kind: VirtualGroupProjectRoleBinding
metadata:
  generateName: projectpolicy-
  namespace: ${projectns}
spec:
  projectRoleRef:
    name: ${projectrole}
  virtualGroupRef:
    name: ${virtualgroup}
EOF
```

6.4.7.3 Configure Project Role Bindings to Bind to WorkspaceRoles - CLI Method

You can also create a Project role binding to bind to a WorkspaceRole in certain instances. To list the WorkspaceRoles that you can bind to a Project, run the following command:

```
kubectl get workspaceRoles -n ${workspacens} -o=jsonpath="{.items[?
(@.metadata.annotations.workspace\.kommander\.d2iq\.io\/project-default-workspace-
role-for==\"${projectns}\")].metadata.name}"
```

You can bind to any of the above WorkspaceRoles by setting `spec.workspaceRoleRef` in the project role binding:

```
cat << EOF | kubectl create -f -
apiVersion: workspaces.kommander.mesosphere.io/v1alpha1
kind: VirtualGroupProjectRoleBinding
metadata:
  generateName: projectpolicy-
  namespace: ${projectns}
spec:
  workspaceRoleRef:
    name: ${workspaceRole}
  virtualGroupRef:
    name: ${virtualgroup}
EOF
```

Note that you must specify either `workspaceRoleRef` or `projectRoleRef` to be validated by the admission webhook. Specifying both values is not valid and will cause an error.

Ensure the `projectns`, `workspacens`, `projectrole` (or `workspaceRole`) and the `virtualgroup` variables are set before executing the command.

When a Project Role Binding is created, Kommander creates a Kubernetes `FederatedRoleBinding` on the Kubernetes cluster where Kommander is running. You can view this by first finding the name of the

project role binding that you created: `kubectl -n ${projectns} get federatedrolebindings.types.kubefed.io`.

Then, view the details like in this example:

```
kubectl -n ${projectns} get federatedrolebindings.types.kubefed.io projectpolicy-gtct4-rdkwq -o yaml
```

Output:

```
apiVersion: types.kubefed.io/v1beta1
kind: FederatedRoleBinding
metadata:
  creationTimestamp: "2020-06-04T16:19:27Z"
  finalizers:
  - kubefed.io/sync-controller
  generation: 1
  name: projectpolicy-gtct4-rdkwq
  namespace: project1-5ljs9-lhvjl
  ownerReferences:
  - apiVersion: workspaces.kommander.mesosphere.io/v1alpha1
    blockOwnerDeletion: true
    controller: true
    kind: VirtualGroupProjectRoleBinding
    name: projectpolicy-gtct4
    uid: 19614de2-4593-433e-82fa-96dc9470e07a
  resourceVersion: "196270"
  selfLink: /apis/types.kubefed.io/v1beta1/namespaces/project1-5ljs9-lhvjl/federatedrolebindings/projectpolicy-gtct4-rdkwq
  uid: beaffc29-edec-4258-9813-3a17ba27a2a6
spec:
  placement:
    clusterSelector: {}
  template:
    roleRef:
      apiGroup: rbac.authorization.k8s.io
      kind: Role
      name: admin-dbfpj-l6s9g
    subjects:
    - apiGroup: rbac.authorization.k8s.io
      kind: User
      name: user1@d2iq.lab
status:
  clusters:
  - name: konvoy-5nr5h
  conditions:
  - lastTransitionTime: "2020-06-04T16:19:27Z"
    lastUpdateTime: "2020-06-04T16:19:27Z"
    status: "True"
```

```
type: Propagation
observedGeneration: 1
```

Then, if you run the following command on a Kubernetes cluster associated with the Project, you'll see a Kubernetes RoleBinding Object, in the corresponding namespace:

```
kubectl -n ${projectns} get rolebinding projectpolicy-gtct4-rdkwq -o yaml
```

Output:

```
apiVersion: rbac.authorization.k8s.io/v1
kind: RoleBinding
metadata:
  creationTimestamp: "2020-06-04T16:19:27Z"
  labels:
    kubefed.io/managed: "true"
  name: projectpolicy-gtct4-rdkwq
  namespace: project1-5ljs9-lhvjl
  resourceVersion: "125392"
  selfLink: /apis/rbac.authorization.k8s.io/v1/namespaces/project1-5ljs9-lhvjl/
  rolebindings/projectpolicy-gtct4-rdkwq
  uid: 2938398d-437b-4f3a-9cb9-c92e50139196
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: Role
  name: admin-dbfpj-l6s9g
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: user1@d2iq.lab
```

6.4.7.4 Role Binding with VirtualGroup

In DKP, `VirtualGroup` is a list of subjects that can be assigned to several different kinds of roles, including:

- `ClusterRole` for cluster-scoped objects
- `WorkspaceRole` for workspace-scoped objects
- `ProjectRole` for project-scoped objects

In order to define which `VirtualGroup` (s) is assigned to one of these roles, administrators can create corresponding role bindings such as `VirtualGroupClusterRoleBinding`, `VirtualGroupWorkspaceRoleBinding`, and `VirtualGroupProjectRoleBinding`.

For more information about configuring `VirtualGroup`, please refer to the [DKP API Documentation](#) (see page 1625).

Note that for `WorkspaceRole` and `ProjectRole`, the referenced `VirtualGroup` and corresponding role and role binding objects need to be in the same namespace. If they are not in the same namespace, the role will not bind to the `VirtualGroup` since it is assumed that the rules set in the role apply to objects that live in that namespace. Whereas for `ClusterRole` which is cluster-scoped, the `VirtualGroupClusterRoleBinding` is also cluster-scoped, even though it references a namespace-scoped `VirtualGroup`.

6.4.8 Project Roles

Enterprise

Gov Advanced

Project Roles are used to define permissions at the namespace level.

6.4.8.1 Configure Project Role - UI Method

In the example below, a Project Role is created with a single Rule. This Project Role corresponds to an admin role.

You can create a Project Role with several Rules.

6.4.8.2 Configure Project Role - CLI Method

The same Project Role can also be created using `kubectl`:

```

cat << EOF | kubectl create -f -
apiVersion: workspaces.kommander.mesosphere.io/v1alpha1
kind: ProjectRole
metadata:
  annotations:
    kommander.mesosphere.io/display-name: Admin
  generateName: admin-
  namespace: ${projectns}
spec:
  rules:
  - apiGroups:
    - '*'
    resources:
    - '*'
    verbs:
    - '*'
EOF

```

Ensure the `projectns` variable is set before executing the command.

You can set it using the following command (for a Kommander Project called `project1`, and after setting the `workspacens` as explained in the previous section):

```

projectns=$(kubectl -n ${workspacens} get projects.workspaces.kommander.mesosphere.io
-o jsonpath='{.items[?
(@.metadata.generateName=="project1-")].status.namespaceRef.name}')

```

When a Project Role is created, Kommander creates a Kubernetes `FederatedRole` on the Kubernetes cluster where Kommander is running:

```

kubectl -n ${projectns} get federatedroles.types.kubefed.io admin-dbfpj-l6s9g -o yaml
apiVersion: types.kubefed.io/v1beta1
kind: FederatedRole
metadata:
  creationTimestamp: "2020-06-04T11:54:26Z"
  finalizers:
  - kubefed.io/sync-controller
  generation: 1
  name: admin-dbfpj-l6s9g
  namespace: project1-5ljs9-lhvjl
  ownerReferences:
  - apiVersion: workspaces.kommander.mesosphere.io/v1alpha1
    blockOwnerDeletion: true
    controller: true
    kind: ProjectRole
    name: admin-dbfpj
    uid: e5f3b2ca-16bf-474d-8305-7be04c034793
  resourceVersion: "75680"

```



```

selfLink: /apis/types.kubefed.io/v1beta1/namespaces/project1-5ljs9-lhvjl/
federatedroles/admin-dbfpj-l6s9g
uid: 1e5a3d98-b223-4605-bba1-16276a3eb47c
spec:
  placement:
    clusterSelector: {}
  template:
    rules:
      - apiGroups:
          - '*'
        resourceNames:
          - '*'
        resources:
          - '*'
        verbs:
          - '*'
status:
  clusters:
    - name: konvoy-5nr5h
  conditions:
    - lastTransitionTime: "2020-06-04T11:54:26Z"
      lastUpdateTime: "2020-06-04T11:54:26Z"
      status: "True"
      type: Propagation
  observedGeneration: 1

```

Then, if you run the following command on a Kubernetes cluster associated with the Project, you see a Kubernetes Role object in the corresponding namespace:

```

kubectl -n ${projectns} get role admin-dbfpj-l6s9g -o yaml
apiVersion: rbac.authorization.k8s.io/v1
kind: Role
metadata:
  creationTimestamp: "2020-06-04T11:54:26Z"
  labels:
    kubefed.io/managed: "true"
  name: admin-dbfpj-l6s9g
  namespace: project1-5ljs9-lhvjl
  resourceVersion: "29218"
  selfLink: /apis/rbac.authorization.k8s.io/v1/namespaces/project1-5ljs9-lhvjl/roles/
  admin-dbfpj-l6s9g
  uid: f05b998c-4649-4e73-bbfe-c12bc4c86a3c
rules:
- apiGroups:
  - '*'
  resourceNames:
  - '*'
  resources:
  - '*'
  verbs:
  - '*'

```

6.4.9 Project ConfigMaps

Enterprise

Gov Advanced

Use ConfigMaps to automate ConfigMaps creation on your clusters

Project ConfigMaps can be created to make sure Kubernetes ConfigMaps are automatically created on all Kubernetes clusters associated with the Project, in the corresponding namespace.

As reference, a ConfigMap is a key-value pair to store some type of non-confidential data like “name=bob” or “state=CA”. For a full reference to the concept, consult the Kubernetes documentation on the topic of [ConfigMaps](#)⁵⁶².

6.4.9.1 Configuring Project ConfigMaps - UI Method

The below Project ConfigMap form can be navigated to by:

1. From the top menu bar, select your target workspace.
2. Select **Projects** from the sidebar menu.
3. Select your project from the list.
4. Select the **ConfigMaps** tab to browse the deployed ConfigMaps.
5. Select **+ Create ConfigMap** button.
6. Enter an ID, Description and Data for the ConfigMap, and select the **Create** button.

6.4.9.2 Configuring Project ConfigMaps - CLI Method

A Project ConfigMap is simply a Kubernetes FederatedConfigMap and can be created using kubectl with YAML:

```
cat << EOF | kubectl create -f -
apiVersion: types.kubefed.io/v1beta1
kind: FederatedConfigMap
metadata:
  generateName: cm1-
  namespace: ${projectns}
spec:
  placement:
    clusterSelector: {}
  template:
    data:
      key: value
EOF
```

⁵⁶² <https://kubernetes.io/docs/concepts/configuration/configmap/>

Ensure the `projectns` variable is set before executing the command. This variable is the project namespace (the Kubernetes Namespace associated with the project) that was defined/created when the project itself was initially created.

```
projectns=$(kubectl -n ${workspacens} get projects.workspaces.kommander.mesosphere.io
-o jsonpath='{.items[?
(@.metadata.generateName=="project1-")].status.namespaceRef.name}')
```

Then, if you run the following command on a Kubernetes cluster associated with the Project, you'll see a Kubernetes ConfigMap Object, in the corresponding namespace:

```
kubectl -n ${projectns} get configmap cm1-8469c -o yaml
apiVersion: v1
data:
  key: value
kind: ConfigMap
metadata:
  creationTimestamp: "2020-06-04T16:37:10Z"
  labels:
    kubefed.io/managed: "true"
  name: cm1-8469c
  namespace: project1-5ljs9-lhvjl
  resourceVersion: "131844"
  selfLink: /api/v1/namespaces/project1-5ljs9-lhvjl/configmaps/cm1-8469c
  uid: d32acb98-3d57-421f-a677-016da5dab980
```

6.4.10 Project Secrets

Enterprise

Gov Advanced

Project Secrets can be created to make sure a Kubernetes Secrets are automatically created on all the Kubernetes clusters associated with the Project, in the corresponding namespace.

6.4.10.1 Configure Project Secrets - UI Method

1. Select the workspace your project was created in from the workspace selection dropdown in the header.
2. In the sidebar menu, select **Projects**.
3. Select the project you want to configure from the table.
4. Select the **Secrets** tab, and then select the **Create Secret** button.
5. Complete the form and select the **Create** button.


6.4.10.2 Configure Project Secrets - CLI Method

A Project Secret is simply a Kubernetes FederatedConfigSecret and can also be created using kubectl:

```
cat << EOF | kubectl create -f -
apiVersion: types.kubefed.io/v1beta1
kind: FederatedSecret
metadata:
  generateName: secret1-
  namespace: ${projectns}
spec:
  placement:
    clusterSelector: {}
  template:
    data:
      key: dmFsdWU=
EOF
```

Ensure the `projectns` variable is set before executing the command.

```
projectns=$(kubectl -n ${workspacens} get projects.workspaces.kommander.mesosphere.io
-o jsonpath='{.items[?
(@.metadata.generateName=="project1-")].status.namespaceRef.name}')
```

 The value of the key is base64 encoded.

If you run the following command on a Kubernetes cluster associated with the Project, you see a Kubernetes Secret Object, in the corresponding namespace:

```
kubectl -n ${projectns} get secret secret1-r9vk2 -o yaml
apiVersion: v1
data:
  key: dmFsdWU=
kind: Secret
metadata:
  creationTimestamp: "2020-06-04T16:51:59Z"
  labels:
    kubefed.io/managed: "true"
  name: secret1-r9vk2
  namespace: project1-5ljs9-lhvjl
  resourceVersion: "137215"
  selfLink: /api/v1/namespaces/project1-5ljs9-lhvjl/secrets/secret1-r9vk2
```

```
uid: e5c6fc1d-93e7-47fe-ae1e-f418f8e35d72
type: Opaque
```

6.4.11 Project Quotas & Limit Ranges

Enterprise

Gov Advanced

Project Quotas and Limit Ranges can be set up to limit the number of resources the Project team uses.

6.4.11.1 Creating Project Quotas & Limit Ranges- UI Method

Project Quotas and Limit Ranges can be set up to limit the number of resources the Project team uses. Quotas and Limit Ranges are applied to all project clusters.

1. Select the workspace your project was created in from the workspace selection dropdown in the header.
2. In the sidebar menu, select **Projects**.
3. Select the project you want to configure from the table.
4. Select the **Quotas & Limit Ranges** tab, and then select the **Edit** button.

Kommander provides a set of default resources for which you can set Quotas. You can also define Quotas for custom resources. We recommend that you set Quotas for CPU and Memory. By using Limit Ranges, you can restrict the resource consumption of individual Pods, Containers, and Persistent Volume Claims in the project namespace. You can also constrain memory and CPU resources consumed by Pods and Containers, and storage resources consumed by Persistent Volume Claims.

5. To add a custom quota, scroll to the bottom of the form and select **Add Quota**.
6. When you are finished, select the **Save** button.

6.4.11.2 Create Project Quotas & Limit Ranges - CLI Method

All the Project Quotas are defined using a Kubernetes FederatedResourceQuota called `kommander` which you can also create/update using `kubectl`:

```
cat << EOF | kubectl apply -f -
apiVersion: types.kubefed.io/v1beta1
kind: FederatedResourceQuota
metadata:
  name: kommander
  namespace: ${projectns}
spec:
  placement:
```

```

clusterSelector: {}
template:
  spec:
    hard:
      limits.cpu: "10"
      limits.memory: 1024.000Mi
EOF

```

Ensure the `projectns` variable is set before executing the command.

```

projectns=$(kubectl -n ${workspacens} get projects.workspaces.kommander.mesosphere.io
-o jsonpath='{.items[?
(@.metadata.generateName=="project1-")].status.namespaceRef.name}')

```

Then, if you run the following command on a Kubernetes cluster associated with the Project, you'll see a Kubernetes Secret Object in the corresponding namespace:

```

kubectl -n ${projectns} get resourcequota kommander -o yaml
apiVersion: v1
kind: ResourceQuota
metadata:
  creationTimestamp: "2020-06-05T08:04:37Z"
  labels:
    kubefed.io/managed: "true"
  name: kommander
  namespace: project1-5ljs9-lhvjl
  resourceVersion: "470822"
  selfLink: /api/v1/namespaces/project1-5ljs9-lhvjl/resourcequotas/kommander
  uid: 925b61b4-134b-4c45-915c-96a05b63d3c3
spec:
  hard:
    limits.cpu: "10"
    limits.memory: 1Gi
status:
  hard:
    limits.cpu: "10"
    limits.memory: 1Gi
  used:
    limits.cpu: "0"
    limits.memory: "0"

```

Similarly, Project Limit Ranges are defined using a FederatedLimitRange object with name `kommander` in the project namespace:

```

cat << EOF | kubectl apply -f -
apiVersion: types.kubefed.io/v1beta1
kind: FederatedLimitRange
metadata:

```

```

name: kommander
namespace: ${projectns}
spec:
  placement:
    clusterSelector: {}
  template:
    spec:
      limits:
        - type: "Pod"
          max:
            cpu: 500m
            memory: 50Gi
          min:
            cpu: 100m
            memory: 10Gi
        - type: "Container"
          max:
            cpu: 2
            memory: 100Mi
          min:
            cpu: 1
            memory: 10Mi
        - type: "PersistentVolumeClaim"
          max:
            storage: 3Gi
          min:
            storage: 1Gi
EOF

```

6.4.12 Project Network Policies

Enterprise

Gov Advanced

Projects are created with a secure-by-default network policy and users needing more flexibility can edit or add more policies to tailor to their unique security needs.

Cluster networking is a critical and central part of Kubernetes that can also be quite challenging. All network communication within and between clusters depends on the presence of a Container Network Interface (CNI) plugin.

6.4.12.1 About Network Plugins

Network plugins ensure that Kubernetes networking requirements are met and surface features needed by network administrators, such as enforcing network policies. Common Network Plugins include Flannel, Calico, and Weave among many others. As an example, D2iQ® Konvoy® uses the Calico CNI plugin by default, but can support others.

Since pods are short-lived, the cluster needs a way to configure the network dynamically as pods are created and destroyed. Plugins provision and manage IP addresses to interfaces and let administrators manage IPs and their assignments to containers, in addition to connections to more than one host, when needed.

6.4.12.2 About Network Policies

By default, and to enable fluid communications within and between clusters, all traffic is authorized between nodes and pods. Most production environments require some kind of traffic flow control at the IP address or port level. An application-centric approach to this uses Network Policies. Pods are isolated by having a Network Policy that selects them, and the configuration of the policy limits the kind of traffic they can receive or send. Network policies do not conflict because they are additive. Pods selected by more than one policy are subject to the union of the policies' ingress and egress rules.

A Network Policy's rules define ingress and egress for network communications between pods and across namespaces. Successful traffic control using network policies is bi-directional. You have to configure both the egress policy on the source pod and the ingress policy on the destination pod to enable the traffic. If either end denies the traffic, it will not flow between the pods.

Since the Kubernetes default is to allow all traffic, it is a common practice to create a default "deny all traffic" rule, and then specifically open up some combination of the pods, ports, and applications as needed.

6.4.12.3 Creating Network Policies

You create network policies in three main parts:

- General information
- Ingress rules
- Egress rules

6.4.12.3.1 General information section

The fields in this part of the form allow you to create a name and description for this policy. Creating a detailed **Description** helps to keep policy functions understandable for additional use and maintenance.

This section also contains the **Pod Selector** fields for selecting pods using either Labels or Expressions. **Labels** added to pod declarations are a common means of identifying individual pods, or creating groups of pods, in a Kubernetes cluster. **Expressions** are similar to Labels, but allow you to define parameters that identify a range of pods.

The **Policy Types** selections help to define the type of Network Policy you are creating:

- **Default** - automatically includes ingress, and egress is set only if the network policy defines egress rules.
- **Ingress** - this policy applies to ingress traffic for the selected pods, to namespaces using the options you define below, or both.
- **Egress** - this policy applies to egress traffic for the selected pods, to namespaces using the options you define below, or both.

If the Default policy type is too rigid or does not offer what you need, you can select the Ingress or Egress type, or both, and explicitly define the policy with the options that follow. For example, if you do not want this policy to apply to ingress traffic, you would select only Egress, and then define the policy.

To deny all ingress traffic, select the **Ingress** option here and then leave the ingress rules empty.

To deny all egress traffic, select the **Egress** option here and then leave the egress rules empty.

6.4.12.3.2 Ingress rules section

Ingress rules use a combination of **Port / Protocol** and **Source** to define the incoming traffic allowed to some or all of the pods in this namespace.

The options under **Sources: From** enable you to define a source either by using the pod selector or by defining an IP block. When using the pod selector method, you can define the namespace, the pods within that namespace, or both.

Namespaces - Selecting a namespace in an ingress rule source permits the pods selected by the pod selector, in your selected namespaces, to receive incoming traffic that meets the other defined criteria. If you have not selected any pods, the rule permits traffic from all pods in the selected namespaces.

Pods - This option selects specific Pods which should be allowed as ingress sources or egress destinations. If you have not selected any namespaces in the namespace selector, this option selects all matching pods in the project namespace. Otherwise, this option selects all matching pods in the selected namespaces.

There also are options to select all namespaces, all pods, or both.

When defining ingress rules using the IP Block method, you define a CIDR and exception conditions. CIDR stands for Classless Inter-Domain Routing and is an IP standard for creating unique network and device identifiers. When grouped together so that they share an initial sequence of bits in their binary representation, the range of addresses creates a CIDR block. The block identity is in an IPv4-like notation including a dotted-decimal address, followed by a slash, then a colon and a number from 0 through 32, for example, 127.0.26.33:31.

6.4.12.3.3 Egress rules section

Egress rules use a similar combination of options to define the outgoing traffic from pods, ranges of pods, or namespaces in a Kommander Project. Port, Protocol, and Destination options for egress rules define the outgoing traffic. You can define your egress rules under **Destination: To**. Ensure the egress policy on the source pods, and the ingress policy on the destination pods, permit traffic in order for the pods to be able to communicate over the network.

6.4.12.4 Network Policy Examples

IMPORTANT: Before you begin each example, ensure you're on the Network Policy page for your project

To navigate to your project's Network Policy page:

1. From the top menu bar, select your target workspace.
2. Select **Projects** from the sidebar menu.
3. Select your project from the list.
4. Select the **Network Policy** tab.

6.4.12.4.1 Ingress: Permit access to API service pods from all namespaces

Suppose you need to create a network policy to permit incoming traffic to API service pods in a specific Kommander Project's namespace from any other pod in any namespace that has the label, `service.corp/users-api-role: client`. For this example, API service pods are those pods created with the Label, `service.corp/users-api-role: api`.

You can limit the policy to just incoming traffic from select namespaces by adding an ingress rule with these characteristics:

- Use Port 8080 to receive incoming TCP traffic
- Refuse traffic from pods unless they are client pods that have a specific Label, such as `service.corp/users-api-role: client`. This example follows a common microservice architecture pattern, `microservice-tier-role: access_mode`

6.4.12.4.1.1 Configure the general information to access API service pods

1. Select the **+ Create Network Policy** button.
2. Type "microsvc-users-api-allow" in the **ID Name** field.
3. Type "Allow Users microservice clients to reach the APIs provided in this namespace" in the **Description** field.
4. Select **Add** under **Pod Selector** and then select **Match Label**.
5. Set the **Key** to "service.corp/users-api-role" and the **Value** to "API".

6.4.12.4.1.2 Create an Ingress rule to access API service pods

1. Leave **Policy Types** set to Default.
2. Scroll down to **Ingress Rules** and select **+ Add an Ingress Rule**.
3. Select **+ Add Port**, and set the **Port** to 8080 and the **Protocol** to TCP.

6.4.12.4.1.3 Add Sources to access API service pods

1. Select **+ Add Source** and mark the **Select All Namespaces** check box.
2. Select **+ Add Pod Selector**.
3. Select **Match Label**.
4. Set the **Key** value to "service.corp/users-api-role" and set the **Value** to "client".
5. Scroll up and select **Save**.

6.4.12.4.2 Ingress: Limit pods that access a database to this namespace

Suppose that while deploying an application in a project, you want to protect its database pods by permitting ingress only from API service pods in the current namespace, and prevent ingress from pods in any other namespace.

You can limit the database pods to just the incoming traffic from the current namespaces by adding an ingress rule with these characteristics:

- Use Port 3306 to receive incoming TCP traffic for pods that have the label, `tier: database`
- Refuse traffic from pods unless they have the label, `tier: api`

6.4.12.4.2.1 Configure the general information to access a database

1. Select the **+ Create Network Policy** button.
2. Type “database-access-api-only” in the **ID name** field.
3. Type “Allow MySQL access only from API pods in this namespace” in the **Description** field.
4. Select **Pod Selector** then select **Match Label**.
5. Set the **Key** to “tier” and the **Value** to “database”.

6.4.12.4.2.2 Create an Ingress rule to access a database

1. Select Default for the **Policy Types**.
2. Scroll down to the **Ingress** section.
3. Select **+ Add Ingress Rule**, then select **+ Add Port**.
4. Set the **Port** to “3306” and leave the **Protocol** set to “TCP”.

6.4.12.4.2.3 Add Sources to access a database

1. Select **+ Add Source**.
2. Select **+ Add Pod Selector**.
3. Select **Match Label** and set the **Key** to “tier” and the **Value** to “API”.
4. Scroll up and select **Save**.

6.4.12.4.3 Ingress: Disable but don't delete ingress rules

Suppose that you want to disable ingress rules temporarily for testing or triaging purposes.

First, you need to create a network policy with one or more ingress rules. You could follow one of the preceding procedures, for example. Then, you need to edit the policy to match the following example:

6.4.12.4.3.1 Edit your Network Policy

1. In the table row belonging to your network policy, click the context menu at the right of the row and select **Edit**.

6.4.12.4.3.2 Disable Ingress rules

1. Update the **Policy Types** so that only **Egress** is selected. If you don't want to deny *all* egress traffic, ensure that you add an egress rule that suits your preferred level of access. You can add an empty rule to allow all egress traffic.
2. Scroll up and select **Save**.

6.4.12.4.4 Egress: Deny all egress traffic from restricted pods

Suppose that you need to deny all egress traffic from a group of restricted pods. This is a simple egress rule and you can create it following this example and steps:

6.4.12.4.4.1 Configure the General Information to deny Egress

1. Select **+ Add Network Policy**.
2. Type "deny-restricted-egress" in the **ID name** field.
3. Type "Deny egress traffic from restricted pods" in the **Description** field.
4. Select **Pod selector** then select **Match Label**.
5. Set the **Key** to "access" and the **Value** to "restricted".

6.4.12.4.4.2 Deny Egress traffic

1. Update the **Policy Types** so that only **Egress** is selected. Do not add any egress rules.
2. Scroll up and select **Save**.

6.5 Managing Clusters

View clusters created with Kommander or any connected Kubernetes cluster

Kommander allows you to monitor and manage very large numbers of clusters. Use the features described in this area to connect existing clusters, or to create new clusters whose lifecycle is managed by Konvoy. You can view clusters from the Clusters tab in the navigation pane on the left. You can see the details for a cluster by selecting the **View Details** link at the bottom of the cluster card or the cluster name in either the card or the table view.

- [Create a Managed Azure Cluster from the DKP UI \(see page 746\)](#)

- [Create a Managed Cluster on vSphere from the DKP UI](#) (see page 747)
- [Attach a Kubernetes Cluster](#) (see page 752)
- [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827)
- [Advanced Creation of CLI Clusters](#) (see page 858)
- [Custom Domains and Certificates Configuration for All Cluster Types](#) (see page 858)
- [Disconnect or Delete Clusters](#) (see page 867)
- [Management Cluster](#) (see page 868)
- [Cluster Statuses](#) (see page 869)
- [Cluster Resources](#) (see page 870)
- [DKP Platform Applications](#) (see page 871)
- [Cluster Applications and Statuses](#) (see page 871)
- [Kubernetes Cluster Federation \(KubeFed\)](#) (see page 874)

6.5.1 Create a Managed Azure Cluster from the DKP UI

Enterprise

Gov Advanced

DKP UI allows you to provision an Azure cluster from your browser.

6.5.1.1 Prerequisites




Before you provision an Azure cluster using the DKP UI, you must first [create an Azure infrastructure provider](#) (see page 621) to hold your Azure credentials.

6.5.1.2 Provision an Azure Cluster

Follow these steps to provision an Azure cluster:

1. From the top menu bar, select your target workspace.
2. Select **Clusters > Add Cluster**.
3. Choose **Create Cluster**.
4. Enter the **Cluster Name**.
5. From **Select Infrastructure Provider**, choose the provider created in the prerequisites section.
6. If available, choose a **Kubernetes Version**. Otherwise, the [DKP 2.6.0 Supported Kubernetes Versions](#) (see page 1758) installs.
7. Select a data center **location** or specify a custom **location**.

8. Edit your worker **Node Pools** as necessary. You can choose the **Number of Nodes**, the **Machine Type**, and for the worker nodes you can choose a **Worker Availability Zone**.
9. Add any additional **Labels** or **Infrastructure Provider Tags** as necessary.
10. Review your inputs to ensure they meet the predefined criteria, and select **Create**.

 It can take up to 15 minutes for your cluster to appear in the **Provisioned** status.

You are then redirected to the **Clusters** page, where you'll see your new cluster in the **Provisioning** status. Hover over the status to view the details.

6.5.2 Create a Managed Cluster on vSphere from the DKP UI


Enterprise

Gov Advanced

You can use the DKP user interface to provision a vSphere cluster quickly and easily.

6.5.2.1 Prerequisites

Ensure that you have fulfilled the vSphere vCenter configuration prerequisites described in [vSphere: VMware Prerequisites \(see page 440\)](#) before you begin these procedures.

 You must also [create a vSphere infrastructure provider \(see page 622\)](#) before you can create additional vSphere clusters.

6.5.2.2 Provision a vSphere Cluster

Provisioning a production-ready cluster in vSphere requires you to specify a fairly large number of parameters. Breaking up the sections of the form, as done below, makes it a little easier to complete.

Complete these procedures to provision a vSphere cluster:

6.5.2.2.1 Provide Basic Cluster Information

In the section of the provisioning form, you give the cluster a name and provide some basic information:

1. In the selected workspace Dashboard, select the **Add Cluster** button at the top right to display the Add Cluster page.

2. Select the **Create Cluster** card.
3. Provide these cluster details in the form:
 - **Cluster Name:** A valid Kubernetes name for the cluster.
 - **Add Labels:** Add any required Labels the cluster needs for your environment by selecting the `+ Add Label` link.
By default, your cluster has labels that reflect the infrastructure provider provisioning. For example, your vSphere cluster may have a label for the data center and `provider: vsphere`. Cluster labels are matched to the selectors created for [Projects](#) (see page 689). Changing a cluster label may add or remove the cluster from projects.
 - **Infrastructure Provider:** This field's value corresponds to the vSphere infrastructure provider you created while fulfilling the prerequisites.
 - **Kubernetes Version:** Select a supported version of Kubernetes for this version of DKP.
 - **SSH Public Key:** Paste into this field the public key value for a user who is authorized to create vSphere clusters.
 - **Workspace:** The workspace where this cluster belongs (if within the Global workspace).

6.5.2.2.2 Specify the Cluster Resources and Network Information

This section of the form identifies already existing resources in your VMware vCenter configuration. Refer to your vCenter configuration to find the necessary values.

1. Provide the following values for the **Resources** that are specific to vSphere:
 - **Datacenter:** Select an existing data center name.
The data-center is the top level organizational unit in vSphere.
 - **Datastore:** Enter a valid vSphere datastore name.
Datastores in vSphere are storage resources that provide storage infrastructure for virtual machines within a data center. They are a subset of data center resources, with each datastore being associated with a specific data center.
 - **Folder:** Enter a valid, existing folder name, or leave it blank to use the vSphere root folder.
When provisioning a Kubernetes cluster on vSphere using Cluster API and `clusterctl`, vSphere uses the folder parameter to specify the vSphere folder where it creates and manages the virtual machines for the Kubernetes cluster. Specifying the folder helps maintain an organized inventory of your virtual machines and other resources in your vSphere environment.
2. Enter the values for the network information in the lower half of this section:
 - **Network:** Enter an existing network name you want the new cluster to use.
You need to create required network resources, such as port groups or distributed port groups, in the vSphere Client or using the vSphere API before you use DKP to create a new cluster.
 - **Resource Pool:** Enter the name of a logical resource pool for the cluster's resources.
In vSphere, resource pools are a logical abstraction that allow you to allocate and manage

computing resources, such as CPU and memory, for a group of virtual machines. Use resource pools only when needed, as they can add complexity to your environment.

- **Virtual Machine Template:** Enter the name of the virtual machine template to use for the managed cluster's virtual machines.
In vSphere, a virtual machine (VM) template is a pre-configured virtual machine that you can use to create new virtual machines with identical configurations quickly. The template contains the basic configuration settings for the VM, such as the operating system, installed software, and hardware configurations.
- **Storage Policy:** Enter the name of a valid vSphere storage policy. This field is optional.
A storage policy in vSphere specifies the storage requirements for virtual machine disks and files. It consists of a rule-set that defines the storage capabilities required, tags to identify it, profiles that collect settings and requirements, and storage requirements that include storage performance, capacity, redundancy, and other attributes necessary for the virtual machine to function properly. By creating and applying a storage policy to a specific datastore or group of datastores, you can ensure that virtual machines using that datastore meet the specified storage requirements.

6.5.2.2.3 Configure Node Pool Information

You need to configure node pool information for both your control plane nodes and your worker nodes. The form splits these information sets into two groups.

1. Provide the control plane node pool name and resource sizing information:
 - **Node Pool Name:** DKP sets this field's value, `control-plane`, and you cannot change it.
 - **Disk:** Enter the amount of disk space allocated for each control plane node. The default value is 80 GB.
 - **Memory:** The amount of memory for each control plane node, in GB. The default value is 16 GB.
 - **Number of CPUs:** Enter the number of virtual processors in each control plane node. The default value is 4 CPUs per control plane node.
 - **Replicas:** Enter the number of control plane nodes to create for your new cluster. Valid values for production clusters are 3 or 5. You can enter 1 if you are creating a test cluster, but a single control plane is not a valid production configuration. You must enter an odd number to allow for internal leader selection processes to provide proper failover for high available. The default value is 3 control plane nodes.
2. Provide the worker node pool name and resource sizing information:
 - **Node Pool Name:** Enter a node pool name for the worker nodes. DKP sets this field's default value to `worker-0`.
 - **Disk:** Enter the amount of disk space allotted for each worker node. The default value is 80GB.
 - **Memory:** The amount of memory for each worker node, in GB. The default value is 32 GB.
 - **Number of CPUs:** Enter the number of virtual processors in each worker node. The default value is 8 CPUs per node.

- **Replicas:** Enter the number of worker nodes to create for your new cluster. The default value is 4 worker nodes.

6.5.2.2.4 Set Virtual IP Parameters

In this section of the form, you configure the built-in virtual IP.

1. Provide the Virtual IP information needed for managing this cluster with DKP:
 - **Interface:** Enter the name of the network used for the virtual IP control plane endpoint. This value is specific to your environment and cannot be inferred by DKP. An example value is `eth0` or `ens5`.
 - **Host:** Enter the control plane endpoint address. To use an external load balancer, set this value to the load balancer's IP address or hostname. To use the built-in virtual IP, set to a static IPv4 address in the Layer 2 network of the control plane machines.
 - **Port:** Enter the control plane's endpoint port. The default port value is 6443. To use an external load balancer, set this value to the load balancer's listening port.

6.5.2.2.5 Supply MetalLB Information

The MetalLB load balancer is needed for cluster installation, and requires these values:

1. Provide a **Starting IP** address range value for the load balancing allocation.
2. Provide an **Ending IP** address range value for the load balancing allocation.

6.5.2.2.6 Configure the StorageClass Options

In this section of the form, you configure the storage options for your vSphere cluster. The StorageClass defines the provisioning properties and requirements for the storage used to store the persistent data of the Kubernetes application.



You can provide either the Datastore URL or the Storage Policy Name in this section.

- Select the **Datastore URL** button if it is not already highlighted, and then in the **Datastore URL** field, enter a unique identifier in URL format used by vSphere to access specific storage locations. A typical example of the field's format is `ds:///vmfs/volumes/<datastore_uuid>/`.
- Select the **Storage Policy Name** button if it is not already highlighted, and then in the **Storage Policy Name** field, enter the name of the storage policy to use with the cluster's StorageClass.

6.5.2.3 Advanced Configuration Parameters

You can open the Advanced configuration parameters sections by selecting the **Show Advanced** link.

6.5.2.3.1 Configure CIDR Values for the Pod Network and Kubernetes Services

In this section of the form, you configure Classless Inter-Domain Routing (CIDR) Values that your vSphere cluster uses.

1. Enter a CIDR value for the Pod network in the **Pod Network CIDR** field. The default value is 192.168.0.0/16.
2. Enter a CIDR value for Kubernetes Services in the **Service CIDR** field. The default value is 10.96.0.0/12.

6.5.2.3.2 Configure the Docker Registry Mirror

In this section, you configure a registry mirror for container images. The first time you request an image from your local registry mirror, it pulls the image from a public registry and stores it locally before handing it back to you. On subsequent requests, the local registry mirror serves the image from its own storage.

1. Configure the image registry mirror:
 - **Registry Mirror URL:** Enter the URL of a container registry to use as a registry mirror.
 - **Registry Mirror Username:** Enter the name of a user who can authenticate to the registry mirror.
 - **Registry Mirror Password:** Enter the password for the username in the previous entry.
 - **Registry Mirror CA Cert:** Upload a certificate file, or copy the CA certificate chain value into the provided field to use while communicating with the registry mirror using Transport Layer Security (TLS).
This value is a trusted root certificate (or chain of certificates) that validates the SSL/TLS connection between clients and the registry mirror, ensuring secure and trustworthy communications.

6.5.2.3.3 Create the Managed Cluster on vSphere

Select the **Create** button (at the page's top right corner) to begin provisioning the cluster.

This step may take a few minutes, taking time for the cluster to be ready and fully deploy its components. The cluster automatically tries to join the management cluster for federation and fleet operations, and should resolve after it is fully provisioned.

While DKP provisions the new cluster, you can access the **Clusters** page to view the new cluster. A new cluster card with the name of your cluster appears, and shows a "Pending" cluster status while the cluster comes up and joins the management cluster.

6.5.2.4 Next Actions

Select the **View Details** link (on the cluster card's bottom left corner) to see additional information about this cluster.

6.5.3 Attach a Kubernetes Cluster

Enterprise

Gov Advanced

Attach an existing Kubernetes cluster using kubeconfig

6.5.3.1 Attach Kubernetes Cluster

You can attach an existing cluster (whether it is a cluster created with the DKP CLI or by means of another platform) to DKP. At the time of attachment, certain namespaces are created on the cluster, and workspace platform applications are deployed automatically into the newly-created namespaces.

Review the [Workspace Platform Application Configuration Requirements](#)⁵⁶³ to ensure the attached cluster has sufficient resources. For more information on platform applications and customizing them, see [Workspace Applications](#)⁵⁶⁴.

If the cluster you want to attach was created using Amazon EKS, Azure AKS or Google GKE, create a service account as described in <https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Attach%20a%20Cluster%20with%20no%20Networking%20Restrictions%20%28UI%29&linkCreation=true&fromPageId=273744255>⁵⁶⁵.



Starting with DKP 2.6.0, DKP supports the attachment of all Kubernetes Conformant clusters, but only x86-64 architecture is supported, not ARM.

Platform applications extend the functionality of Kubernetes and provide ready-to-use logging and monitoring stacks by deploying platform applications when attaching a cluster to Kommander. For more information, refer to [Workspace Applications](#)⁵⁶⁶.

- [Requirements for Attaching an Existing Cluster](#) (see page 753)

⁵⁶³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273744255&linkCreation=true&spaceKey=DENT&title=%282.6%29+Workspace+Platform+Application+Defaults+and+Resource+Requirements>

⁵⁶⁴ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273744255&linkCreation=true&spaceKey=DENT&title=%282.6%29+Workspace+Applications>

⁵⁶⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273744255&linkCreation=true&spaceKey=DENT&title=%282.6%29+Attach+a+Cluster+with+no+Networking+Restrictions+%28UI%29>

⁵⁶⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273744255&linkCreation=true&spaceKey=DENT&title=%282.6%29+Workspace+Applications>

- [Create a kubeconfig File for Attachment](#) (see page 755)
- [Attach a Cluster with no Networking Restrictions \(UI\)](#) (see page 758)
- [Attach a Cluster with Networking Restrictions](#) (see page 772)
- [Attach a DKP-created Kubernetes Cluster](#) (see page 824)
- [Access a Managed or Attached Cluster](#) (see page 826)
- [Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827)

6.5.3.2 Requirements for Attaching an Existing Cluster

Enterprise

Gov Advanced

6.5.3.2.1 Basic Requirements

To attach an existing cluster in the UI, the Application Management cluster must be able to reach the services and the `api-server` of the target cluster.

The cluster you want to attach can be a **DKP-CLI-created** cluster (which will become a [Managed cluster](#) (see page 82) upon attachment), or **another Kubernetes cluster** like AKS, EKS, or GKE (which will become an [Attached cluster](#) (see page 82) upon attachment).

- Starting with DKP 2.6.0, DKP supports the attachment of all Kubernetes Conformant clusters, but only x86-64 architecture is supported, not ARM.

For attaching existing clusters without networking restrictions, the requirements depend on which DKP version you are using. Each version of DKP supports a specific range of [Kubernetes versions](#) (see page 1758). You must ensure that the target cluster is running a compatible version.

- As of this version of DKP, there is no official method of uninstalling Kommander.

6.5.3.2.2 Create a Default StorageClass

To deploy many of the services on the attached cluster, there must be a default `StorageClass` configured. Run the following command on the cluster you want to attach:

```
kubectl get sc
```

The output should look similar to this. Note the `(default)` after the name:

NAME	PROVISIONER	RECLAIMPOLICY	VOLUMEBINDINGMODE
ALLOWVOLUMEEXPANSION	AGE		
ebs-sc (default)	ebs.csi.aws.com	Delete	WaitForFirstConsumer false
41s			

If the `StorageClass` is not set as default, add the following annotation to the `StorageClass` manifest:

```

annotations:
  storageclass.kubernetes.io/is-default-class: "true"

```

6.5.3.2.3 Creating Projects and Workspaces

Before you attach clusters, you need to create one or more Workspaces, and we recommend that you also create Projects within your Workspaces. [Workspaces](#) (see page 653) give you a logical way to represent your teams and specific configurations. [Projects](#) (see page 689) let you define one or more clusters as a group to which Kommander pushes a common configuration. Grouping your existing clusters in Kommander projects and workspaces makes managing their platform services and resources easier and supports monitoring and logging.



Do not attach a cluster in the "Management Cluster Workspace" workspace. This workspace is reserved for your Application Management cluster only.

6.5.3.2.4 Platform Application Requirements

In addition to the basic cluster requirements, the platform services you want DKP to manage on those clusters will have an impact on the total cluster requirements. The specific combinations of platform applications will make a difference in the requirements for the cluster nodes and their resources (CPU, memory, and storage).

See [this table of platform applications](#) (see page 642) that DKP provides by default.

6.5.3.2.5 Attach Existing AKS, EKS and GKE clusters

Attaching an existing AWS cluster requires that the cluster be fully [configured and running](#) (see page 1076). You must create a separate service account when attaching existing AKS, EKS or Google GKE Kubernetes clusters. This is necessary because the kubeconfig files generated from those clusters are not usable out-of-the-box by Kommander. The kubeconfig files call CLI commands, such as `azure`, `aws` or `gcloud`, and use locally-obtained authentication tokens. Having a separate service account also allows you to keep access to the cluster specific and isolated to Kommander.

The suggested default cluster configuration includes a control plane pool containing three (3) m5.xlarge nodes and a worker pool containing four (4) m5.2xlarge nodes.

Consider the additional resource requirements for running the platform services you want DKP to manage, and ensure that your existing clusters comply.

To attach an existing EKS cluster, refer to the specific information in [Attach Amazon EKS Cluster](#) (see page 1144).

To attach an existing GKE cluster, refer to the specific information in [Attach GKE Cluster](#) (see page 768).

6.5.3.2.6 Attach Clusters with an Existing cert-manager Installation

If you are attaching clusters that already have cert-manager installed, the cert-manager `HelmsRelease` provided by DKP will fail to deploy, due to the existing cert-manager installation. As long as the pre-existing cert-manager functions as expected, you can ignore this failure. It will have no impact on the operation of the cluster.

6.5.3.3 Create a kubeconfig File for Attachment

Create a separate service account when attaching existing clusters (for example Amazon EKS, or Azure AKS clusters)

If you already have a `kubeconfig` file to attach your cluster, go directly to [Attach a Cluster with no Networking Restrictions \(UI\)](#) (see page 758) or [Attach a Cluster with Networking Restrictions](#) (see page 772).

The `kubeconfig` files generated from existing clusters are not usable out-of-the-box, because they call provisioner-specific CLI commands (like `aws` commands), and use locally-obtained authentication tokens that are not compatible with DKP. Having a separate service account also allows you to have a dedicated identity for all DKP operations.

To get started, ensure you have `kubectl`⁵⁶⁷ set up and configured with `ClusterAdmin`⁵⁶⁸ for the cluster you want to connect to Kommander.

1. Create the necessary service account:

```
kubectl -n kube-system create serviceaccount kommander-cluster-admin
```

2. Create a token secret for the `serviceaccount`:

```
kubectl -n kube-system create -f - <<EOF
apiVersion: v1
```

⁵⁶⁷ <https://kubernetes.io/docs/tasks/tools/#kubectl>

⁵⁶⁸ <https://kubernetes.io/docs/concepts/cluster-administration/>

```
kind: Secret
metadata:
  name: kommander-cluster-admin-sa-token
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
type: kubernetes.io/service-account-token
EOF
```

3. Verify that the `serviceaccount` token is ready by running this command:

```
kubectl -n kube-system get secret kommander-cluster-admin-sa-token -oyaml
```

Verify that the `data.token` field is populated. The output should be similar to this:

```
apiVersion: v1
data:
  ca.crt: LS0tLS1CRUdJTiBDR...
  namespace: ZGVmYXVsdA==
  token: ZXlKaGJHY2lPaUpTVX...
kind: Secret
metadata:
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
    kubernetes.io/service-account.uid: b62bc32e-b502-4654-921d-94a742e273a8
  creationTimestamp: "2022-08-19T13:36:42Z"
  name: kommander-cluster-admin-sa-token
  namespace: default
  resourceVersion: "8554"
  uid: 72c2a4f0-636d-4a70-9f1c-55a75f15e520
type: kubernetes.io/service-account-token
```

4. Configure the new service account for `cluster-admin` permissions:

```
cat << EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: kommander-cluster-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- kind: ServiceAccount
  name: kommander-cluster-admin
  namespace: kube-system
EOF
```

- Set up the following environment variables with the access data that is needed for producing a new kubeconfig file:

```
export USER_TOKEN_VALUE=$(kubectl -n kube-system get secret/kommander-cluster-admin-sa-token -o=go-template='{{.data.token}}' | base64 --decode)
export CURRENT_CONTEXT=$(kubectl config current-context)
export CURRENT_CLUSTER=$(kubectl config view --raw -o=go-template='{{range .contexts}}{{if eq .name "'${CURRENT_CONTEXT}'"}}{{ index .context "cluster" }}{{end}}{{end}}')
export CLUSTER_CA=$(kubectl config view --raw -o=go-template='{{range .clusters}}{{if eq .name "'${CURRENT_CLUSTER}'"}}"{{with index .cluster "certificate-authority-data" }}{{.}}{{end}}"{{ end }}{{ end }}')
export CLUSTER_SERVER=$(kubectl config view --raw -o=go-template='{{range .clusters}}{{if eq .name "'${CURRENT_CLUSTER}'"}}{{ .cluster.server }}{{end}}{{ end }}')
```

- Confirm these variables have been set correctly:

```
export -p USER_TOKEN_VALUE CURRENT_CONTEXT CURRENT_CLUSTER CLUSTER_CA CLUSTER_SERVER
```

- Generate a `kubeconfig` file that uses the environment variable values from the previous step:

```
cat << EOF > kommander-cluster-admin-config
apiVersion: v1
kind: Config
current-context: ${CURRENT_CONTEXT}
contexts:
- name: ${CURRENT_CONTEXT}
  context:
    cluster: ${CURRENT_CONTEXT}
    user: kommander-cluster-admin
    namespace: kube-system
clusters:
- name: ${CURRENT_CONTEXT}
  cluster:
    certificate-authority-data: ${CLUSTER_CA}
    server: ${CLUSTER_SERVER}
users:
- name: kommander-cluster-admin
  user:
    token: ${USER_TOKEN_VALUE}
EOF
```

- This process produces a file in your current working directory called `kommander-cluster-admin-config`. The contents of this file are used in Kommander to attach the cluster.

Before importing this configuration, verify the `kubeconfig` file can access the cluster:


```
kubectl --kubeconfig $(pwd)/kommander-cluster-admin-config get all --all-namespaces
```

6.5.3.3.1 Next Step:

Use this `kubeconfig` to:

- Attach a cluster with [no additional networking restrictions](#) (see page 758)
- Attach a cluster that [has networking restrictions](#) (see page 772)

ⓘ If a cluster has limited resources to deploy all the federated platform services, it will fail to stay attached in the DKP UI. If this happens, check if there are any pods that are not getting the resources required.

6.5.3.4 Attach a Cluster with no Networking Restrictions (UI)

Enterprise

Gov Advanced

For provider-specific instructions, refer to:

- [Attach Amazon EKS Cluster](#) (see page 759)
- [Attach AKS Cluster](#) (see page 763)
- [Attach GKE Cluster](#) (see page 768)

Using the **Add Cluster** option, you can attach an existing Kubernetes or DKP cluster directly to DKP. This enables you to access the multi-cluster management and monitoring benefits that DKP provides, while keeping your existing cluster on its current provider and infrastructure.

How to attach an existing cluster that has no additional networking restrictions

Use this option when you want to attach a cluster that does not require additional access information.

1. From the top menu bar, select your target workspace.
2. On the Dashboard page, select the **Add Cluster** option in the **Actions** dropdown menu at the top right.
3. Select **Attach Cluster**.
4. Select the **No additional networking restrictions** card. Alternatively, if you MUST use network restrictions, stop following the steps below, and see the instructions on the page [Attach a Cluster with Network Restrictions](#) (see page 772).

5. In the **Cluster Configuration** section, paste your `kubeconfig` file into the field, or select the **Upload kubeconfig File** button to specify the file.
6. The **Cluster Name** field will automatically populate with the name of the cluster is in the `kubeconfig`. You can edit this field with the name you want for your cluster.
7. The **Context** select list is populated from the `kubeconfig`. Select the desired context with admin privileges from the **Context** select list.
8. Add labels to classify your cluster as needed.
9. Select **Create** to attach your cluster.

6.5.3.4.1 Attach Amazon EKS Cluster

Enterprise

Gov Advanced

6.5.3.4.1.1 Attach an existing EKS cluster

You can attach existing Kubernetes clusters to the Management Cluster. After attaching the cluster, you can use the UI to [examine and manage](#) (see page 748) this cluster. The following procedure shows how to attach an existing Amazon Elastic Kubernetes Service (EKS) cluster.

6.5.3.4.1.2 Before you Begin

This procedure requires the following items and configurations:

- A fully configured and running Amazon [EKS](#)⁵⁶⁹ cluster with administrative privileges.
- The current version DKP Enterprise is [installed](#) (see page 1394) on your cluster.
- Ensure you have installed `kubectl` in your Management cluster.



This procedure assumes you have an existing and spun up Amazon EKS cluster(s) with administrative privileges. Refer to the Amazon [EKS](#)⁵⁷⁰ for setup and configuration information.

6.5.3.4.1.3 Attach Amazon EKS Clusters

Access your EKS clusters

⁵⁶⁹ <https://aws.amazon.com/eks/>

⁵⁷⁰ <https://aws.amazon.com/eks/>

1. Ensure you are connected to your EKS clusters. Enter the following commands for each of your clusters:

```
kubectl config get-contexts
kubectl config use-context <context for first eks cluster>
```

2. Confirm `kubectl` can access the EKS cluster:

```
kubectl get nodes
```

Create a `kubeconfig` File for your EKS cluster

To get started, ensure you have `kubectl`⁵⁷¹ set up and configured with `ClusterAdmin`⁵⁷² for the cluster you want to connect to Kommander.

1. Create the necessary service account:

```
kubectl -n kube-system create serviceaccount kommander-cluster-admin
```

2. Create a token secret for the `serviceaccount`:

```
kubectl -n kube-system create -f - <<EOF
apiVersion: v1
kind: Secret
metadata:
  name: kommander-cluster-admin-sa-token
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
type: kubernetes.io/service-account-token
EOF
```

For more information on Service Account Tokens, refer to [this article](#)⁵⁷³ in our blog.

3. Verify that the `serviceaccount` token is ready by running this command:

```
kubectl -n kube-system get secret kommander-cluster-admin-sa-token -oyaml
```

Verify that the `data.token` field is populated. The output should be similar to this:

⁵⁷¹ <https://kubernetes.io/docs/tasks/tools/#kubectl>

⁵⁷² <https://kubernetes.io/docs/concepts/cluster-administration/>

⁵⁷³ <https://eng.d2iq.com/blog/service-account-tokens-in-kubernetes-v1.24/#whats-changed-in-kubernetes-v124>

```

apiVersion: v1
data:
  ca.crt: LS0tLS1CRUdJTiBDR...
  namespace: ZGVmYXVsdA==
  token: ZXlKaGJHY2lPaUpTVX...
kind: Secret
metadata:
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
    kubernetes.io/service-account.uid: b62bc32e-b502-4654-921d-94a742e273a8
  creationTimestamp: "2022-08-19T13:36:42Z"
  name: kommander-cluster-admin-sa-token
  namespace: default
  resourceVersion: "8554"
  uid: 72c2a4f0-636d-4a70-9f1c-55a75f15e520
type: kubernetes.io/service-account-token

```

4. Configure the new service account for `cluster-admin` permissions:

```

cat << EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: kommander-cluster-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- kind: ServiceAccount
  name: kommander-cluster-admin
  namespace: kube-system
EOF

```

5. Set up the following environment variables with the access data that is needed for producing a new kubeconfig file:

```

export USER_TOKEN_VALUE=$(kubectl -n kube-system get secret/kommander-cluster-admin-sa-token -o=go-template='{{.data.token}}' | base64 --decode)
export CURRENT_CONTEXT=$(kubectl config current-context)
export CURRENT_CLUSTER=$(kubectl config view --raw -o=go-template='{{range .contexts}}{{if eq .name "'${CURRENT_CONTEXT}'"}}{{index .context "cluster"}}{{end}}{{end}}')
export CLUSTER_CA=$(kubectl config view --raw -o=go-template='{{range .clusters}}{{if eq .name "'${CURRENT_CLUSTER}'"}}"{{with index .cluster "certificate-authority-data"}}{{.}}{{end}}"{{end}}')

```

```
export CLUSTER_SERVER=$(kubectl config view --raw -o=go-template='{{range .clusters}}{{if eq .name ""}}{{CURRENT_CLUSTER}}'')}}
{{.cluster.server}}'')}}{{end}}'')}}')
```

6. Confirm these variables have been set correctly:

```
export -p | grep -E 'USER_TOKEN_VALUE|CURRENT_CONTEXT|CURRENT_CLUSTER|
CLUSTER_CA|CLUSTER_SERVER'
```

7. Generate a kubeconfig file that uses the environment variable values from the previous step:

```
cat << EOF > kommander-cluster-admin-config
apiVersion: v1
kind: Config
current-context: ${CURRENT_CONTEXT}
contexts:
- name: ${CURRENT_CONTEXT}
  context:
    cluster: ${CURRENT_CONTEXT}
    user: kommander-cluster-admin
    namespace: kube-system
clusters:
- name: ${CURRENT_CONTEXT}
  cluster:
    certificate-authority-data: ${CLUSTER_CA}
    server: ${CLUSTER_SERVER}
users:
- name: kommander-cluster-admin
  user:
    token: ${USER_TOKEN_VALUE}
EOF
```

8. This process produces a file in your current working directory called `kommander-cluster-admin-config`. The contents of this file are used in Kommander to attach the cluster.

Before importing this configuration, verify the `kubeconfig` file can access the cluster:

```
kubectl --kubeconfig $(pwd)/kommander-cluster-admin-config get all --all-
namespaces
```

Finalize Attachment of your Cluster from the UI

Now that you have kubeconfig, go to the DKP UI and follow these steps below:

1. From the top menu bar, select your target workspace.
2. On the Dashboard page, select the **Add Cluster** option in the **Actions** dropdown menu at the top right.

3. Select **Attach Cluster**.
4. Select the **No additional networking restrictions** card. Alternatively, if you must use network restrictions, stop following the steps below, and see the instructions on the page [Attach a cluster WITH network restrictions](#) (see page 772).
5. Upload the kubeconfig file you created in the previous section (or copy its contents) into the **Cluster Configuration** section.
6. The **Cluster Name** field automatically populates with the name of the cluster in the kubeconfig. You can edit this field with the name you want for your cluster.
7. Add labels to classify your cluster as needed.
8. Select **Create** to attach your cluster.



If a cluster has limited resources to deploy all the federated platform services, it will fail to stay attached in the DKP UI. If this happens, ensure your system has sufficient resources for all pods.

6.5.3.4.1.4 Related Information

For information on related topics or procedures, refer to the following:

- [Configuring and Running Amazon EKS Clusters](#)⁵⁷⁴
- [Installing and Configuring Konvoy v2.0 or above](#) (see page 1127)
- [Installing and Configuring Kommander v2.0 or above](#) (see page 1402)
- [Manage Clusters](#) (see page 748)

6.5.3.4.2 Attach AKS Cluster

ENTERPRISE

6.5.3.4.2.1 Attach an existing AKS cluster

You can attach existing Kubernetes clusters to the Management Cluster. After attaching the cluster, you can use the UI to [examine and manage](#) (see page 748) this cluster. The following procedure shows how to attach an existing Azure Kubernetes Service (AKS) cluster.

6.5.3.4.2.2 Before you Begin

This procedure requires the following items and configurations:

- A fully configured and running Azure [AKS](#)⁵⁷⁵ cluster with administrative privileges.

⁵⁷⁴ <https://aws.amazon.com/eks/>

⁵⁷⁵ <https://azure.microsoft.com/en-us/products/kubernetes-service/>

- The current version DKP Enterprise is [installed](#) (see page 147) on your cluster.
- Ensure you have installed `kubectl` in your Management cluster.

ⓘ This procedure assumes you have an existing and spun up Azure AKS cluster(s) with administrative privileges. Refer to the Azure [AKS](#)⁵⁷⁶ for setup and configuration information.

6.5.3.4.2.3 Attach AKS Clusters

Ensure that the `KUBECONFIG` environment variable is set to the Management cluster before attaching by running:

```
export KUBECONFIG=<Management_cluster_kubeconfig>.conf
```

Ensure you have access to your AKS clusters

1. Ensure you are connected to your AKS clusters. Enter the following commands for each of your clusters:

```
kubectl config get-contexts
kubectl config use-context <context for first AKS cluster>
```

2. Confirm `kubectl` can access the AKS cluster:

```
kubectl get nodes
```

Create a kubeconfig file for your AKS cluster

To get started, ensure you have `kubect`⁵⁷⁷ set up and configured with `ClusterAdmin`⁵⁷⁸ for the cluster you want to connect to Kommander.

1. Create the necessary service account:

```
kubectl -n kube-system create serviceaccount kommander-cluster-admin
```

⁵⁷⁶ <https://azure.microsoft.com/en-us/products/kubernetes-service/>

⁵⁷⁷ <https://kubernetes.io/docs/tasks/tools/#kubectl>

⁵⁷⁸ <https://kubernetes.io/docs/concepts/cluster-administration/>

2. Create a token secret for the `serviceaccount` :

```
kubectl -n kube-system create -f - <<EOF
apiVersion: v1
kind: Secret
metadata:
  name: kommander-cluster-admin-sa-token
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
type: kubernetes.io/service-account-token
EOF
```

3. Verify that the `serviceaccount` token is ready by running this command:

```
kubectl -n kube-system get secret kommander-cluster-admin-sa-token -oyaml
```

Verify that the `data.token` field is populated. The output should be similar to this:

```
apiVersion: v1
data:
  ca.crt: LS0tLS1CRUdJTiBDR...
  namespace: ZGVmYXVsdA==
  token: ZXlKaGJHY2lPaUpTVX...
kind: Secret
metadata:
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
    kubernetes.io/service-account.uid: b62bc32e-b502-4654-921d-94a742e273a8
  creationTimestamp: "2022-08-19T13:36:42Z"
  name: kommander-cluster-admin-sa-token
  namespace: default
  resourceVersion: "8554"
  uid: 72c2a4f0-636d-4a70-9f1c-55a75f15e520
type: kubernetes.io/service-account-token
```

4. Configure the new service account for `cluster-admin` permissions:

```
cat << EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: kommander-cluster-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
```



```

subjects:
- kind: ServiceAccount
  name: kommander-cluster-admin
  namespace: kube-system
EOF

```

5. Set up the following environment variables with the access data that is needed for producing a new kubeconfig file:

```

export USER_TOKEN_VALUE=$(kubectl -n kube-system get secret/kommander-cluster-admin-sa-token -o=go-template='{{.data.token}}' | base64 --decode)
export CURRENT_CONTEXT=$(kubectl config current-context)
export CURRENT_CLUSTER=$(kubectl config view --raw -o=go-template='{{range .contexts}}{{if eq .name "'${CURRENT_CONTEXT}'"}}{{index .context "cluster"}}{{end}}{{end}}')
export CLUSTER_CA=$(kubectl config view --raw -o=go-template='{{range .clusters}}{{if eq .name "'${CURRENT_CLUSTER}'"}}"{{with index .cluster "certificate-authority-data"}}{{.}}"{{end}}"{{end}}')
export CLUSTER_SERVER=$(kubectl config view --raw -o=go-template='{{range .clusters}}{{if eq .name "'${CURRENT_CLUSTER}'"}}{{.cluster.server}}"{{end}}')

```

6. Confirm these variables have been set correctly:

```

export -p | grep -E 'USER_TOKEN_VALUE|CURRENT_CONTEXT|CURRENT_CLUSTER|CLUSTER_CA|CLUSTER_SERVER'

```

7. Generate a kubeconfig file that uses the environment variable values from the previous step:

```

cat << EOF > kommander-cluster-admin-config
apiVersion: v1
kind: Config
current-context: ${CURRENT_CONTEXT}
contexts:
- name: ${CURRENT_CONTEXT}
  context:
    cluster: ${CURRENT_CONTEXT}
    user: kommander-cluster-admin
    namespace: kube-system
clusters:
- name: ${CURRENT_CONTEXT}
  cluster:
    certificate-authority-data: ${CLUSTER_CA}
    server: ${CLUSTER_SERVER}
users:
- name: kommander-cluster-admin
  user:
    token: ${USER_TOKEN_VALUE}

```

```
EOF
```

- This process produces a file in your current working directory called `kommander-cluster-admin-config`. The contents of this file are used in Kommander to attach the cluster.

Before importing this configuration, verify the `kubeconfig` file can access the cluster:

```
kubectl --kubeconfig $(pwd)/kommander-cluster-admin-config get all --all-namespaces
```

Finalize attaching your cluster from the UI

Now that you have kubeconfig, go to the DKP UI and follow these steps below:

- From the top menu bar, select your target workspace.
- On the Dashboard page, select the **Add Cluster** option in the **Actions** dropdown menu at the top right.
- Select **Attach Cluster**.
- Select the **No additional networking restrictions** card. Alternatively, if you must use network restrictions, stop following the steps below, and see the instructions on the page [Attach a cluster WITH network restrictions](#) (see page 772).
- Upload the kubeconfig file you created in the previous section (or copy its contents) into the **Cluster Configuration** section.
- The **Cluster Name** field automatically populates with the name of the cluster in the kubeconfig. You can edit this field with the name you want for your cluster.
- Add labels to classify your cluster as needed.
- Select **Create** to attach your cluster.



If a cluster has limited resources to deploy all the federated platform services, it will fail to stay attached in the DKP UI. If this happens, ensure your system has sufficient resources for all pods.

6.5.3.4.2.4 Related Information

For information on related topics or procedures, refer to the following:

- [Installing and Configuring Konvoy v2.0 or above](#) (see page 1189)
- [Installing and Configuring Kommander v2.0 or above](#) (see page 1396)
- [Manage Clusters](#) (see page 748)

6.5.3.4.3 Attach GKE Cluster

Enterprise

Gov Advanced

Attach an existing GKE cluster in DKP

You can attach existing Kubernetes clusters to the Management Cluster. After attaching the cluster, you can use the UI to [examine and manage](#) (see page 748) this cluster. The following procedure shows how to attach an existing **standard** GKE cluster.

6.5.3.4.3.1 Before you Begin

This procedure requires the following items and configurations:

- A fully configured and running [GKE cluster](#)⁵⁷⁹ with a [supported Kubernetes version](#)⁵⁸⁰, and administrative privileges.
- The current version of DKP Enterprise [installed](#) (see page 147) on your cluster.
- Ensure you have installed `kubectl` in your Management cluster.



This procedure assumes you have an existing and spun up GKE cluster with administrator privileges.

6.5.3.4.3.2 Attach GKE Clusters

Ensure you have access to your GKE clusters

1. Ensure you are connected to your GKE clusters. Enter the following commands for each of your clusters:

```
kubectl config get-contexts
kubectl config use-context <context for first gcloud cluster>
```

2. Confirm `kubectl` can access the GKE cluster.

```
kubectl get nodes
```

⁵⁷⁹ <https://cloud.google.com/kubernetes-engine/docs/concepts/types-of-clusters>

⁵⁸⁰ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/150765792/DKP+2.5.0+Supported+Kubernetes+Versions>

Configure a kubeconfig file

To get started, ensure you have [kubectl](#)⁵⁸¹ set up and configured with [ClusterAdmin](#)⁵⁸² for the cluster you want to connect to Kommander.

1. Create the necessary service account:

```
kubectl -n kube-system create serviceaccount kommander-cluster-admin
```

2. Create a token secret for the `serviceaccount`:

```
kubectl -n kube-system create -f - <<EOF
apiVersion: v1
kind: Secret
metadata:
  name: kommander-cluster-admin-sa-token
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
type: kubernetes.io/service-account-token
EOF
```

For more information on Service Account Tokens, refer to [this article](#)⁵⁸³ in our blog.

3. Verify that the `serviceaccount` token is ready by running this command:

```
kubectl -n kube-system get secret kommander-cluster-admin-sa-token -oyaml
```

Verify that the `data.token` field is populated. The output should be similar to this:

```
apiVersion: v1
data:
  ca.crt: LS0tLS1CRUdJTiBDR...
  namespace: ZGVmYXVsdA==
  token: ZXlKaGJHY2lPaUpTVX...
kind: Secret
metadata:
  annotations:
    kubernetes.io/service-account.name: kommander-cluster-admin
    kubernetes.io/service-account.uid: b62bc32e-b502-4654-921d-94a742e273a8
  creationTimestamp: "2022-08-19T13:36:42Z"
```

581 <https://kubernetes.io/docs/tasks/tools/#kubectl>

582 <https://kubernetes.io/docs/concepts/cluster-administration/>

583 <https://eng.d2iq.com/blog/service-account-tokens-in-kubernetes-v1.24/#whats-changed-in-kubernetes-v124>

```

name: kommander-cluster-admin-sa-token
namespace: default
resourceVersion: "8554"
uid: 72c2a4f0-636d-4a70-9f1c-55a75f15e520
type: kubernetes.io/service-account-token

```

4. Configure the new service account for `cluster-admin` permissions:

```

cat << EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: kommander-cluster-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- kind: ServiceAccount
  name: kommander-cluster-admin
  namespace: kube-system
EOF

```

5. Set up the following environment variables with the access data that is needed for producing a new kubeconfig file:

```

export USER_TOKEN_VALUE=$(kubectl -n kube-system get secret/kommander-cluster-admin-sa-token -o=go-template='{{.data.token}}' | base64 --decode)
export CURRENT_CONTEXT=$(kubectl config current-context)
export CURRENT_CLUSTER=$(kubectl config view --raw -o=go-template='{{range .contexts}}{{if eq .name "'${CURRENT_CONTEXT}'"}}{{index .context "cluster"}}{{end}}{{end}}')
export CLUSTER_CA=$(kubectl config view --raw -o=go-template='{{range .clusters}}{{if eq .name "'${CURRENT_CLUSTER}'"}}{{with index .cluster "certificate-authority-data"}}{{.}}{{end}}"{{end}}')
export CLUSTER_SERVER=$(kubectl config view --raw -o=go-template='{{range .clusters}}{{if eq .name "'${CURRENT_CLUSTER}'"}}{{.cluster.server}}{{end}}{{end}}')

```

6. Confirm these variables have been set correctly:

```

export -p | grep -E 'USER_TOKEN_VALUE|CURRENT_CONTEXT|CURRENT_CLUSTER|CLUSTER_CA|CLUSTER_SERVER'

```

7. Generate a kubeconfig file that uses the environment variable values from the previous step:

```

cat << EOF > kommander-cluster-admin-config
apiVersion: v1
kind: Config
current-context: ${CURRENT_CONTEXT}
contexts:
- name: ${CURRENT_CONTEXT}
  context:
    cluster: ${CURRENT_CONTEXT}
    user: kommander-cluster-admin
    namespace: kube-system
clusters:
- name: ${CURRENT_CONTEXT}
  cluster:
    certificate-authority-data: ${CLUSTER_CA}
    server: ${CLUSTER_SERVER}
users:
- name: kommander-cluster-admin
  user:
    token: ${USER_TOKEN_VALUE}
EOF

```

8. This process produces a file in your current working directory called `kommander-cluster-admin-config`. The contents of this file are used in Kommander to attach the cluster. Before importing this configuration, verify the `kubeconfig` file can access the cluster:

```

kubectl --kubeconfig $(pwd)/kommander-cluster-admin-config get all --all-namespaces

```

Attach the Cluster

Now that you have a kubeconfig file, go to the DKP UI and follow these steps:

1. Select your target workspace from the top menu bar.
2. Select **Add Cluster** in the **Actions** dropdown menu from the Dashboard page, located at the top-right.
3. Select **Attach Cluster**.
4. Select the **No additional networking restrictions** card. Alternatively, if you must use network restrictions, stop following the steps below, and see the instructions on the page [Attach a cluster WITH network restrictions \(see page 772\)](#).
5. Upload the kubeconfig file you created in the previous section (or copy its contents) into the **Cluster Configuration** section.
6. The **Cluster Name** field automatically populates with the name of the cluster in the kubeconfig. You can edit this field with the name you want for your cluster.
7. Add labels to classify your cluster as needed.

8. Select **Create** to attach your cluster.



If a cluster has limited resources to deploy all the federated platform services, it will fail to stay attached in the DKP UI. If this happens, ensure your system has sufficient resources for all pods.

6.5.3.4.3 Related Information

For information on related topics or procedures, refer to the following:

- [Installing and Configuring Kommander](#) (see page 1396)
- [Manage Clusters](#) (see page 748)

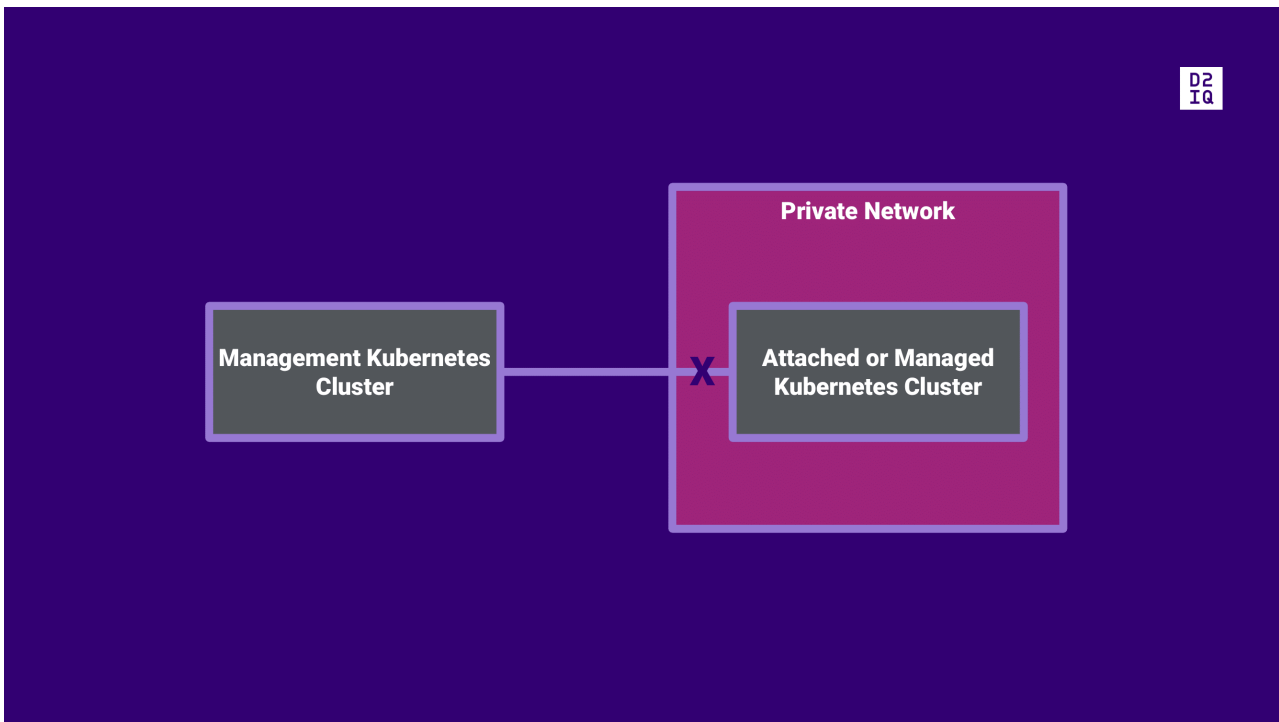
6.5.3.5 Attach a Cluster with Networking Restrictions

Enterprise

Gov Advanced

6.5.3.5.1 When do I Need to use a Secure Tunnel?

When attaching a cluster to DKP, the Management cluster initiates an outbound connection to the cluster you wish to attach. This is not possible if the cluster you want to attach ([Managed or Attached](#) (see page 81)) has networking restrictions and is not exposed, for example, because it is in a private network, or its API is not accessible from the same network as the Management cluster. This is what we call a [network-restricted cluster](#) (see page 84).



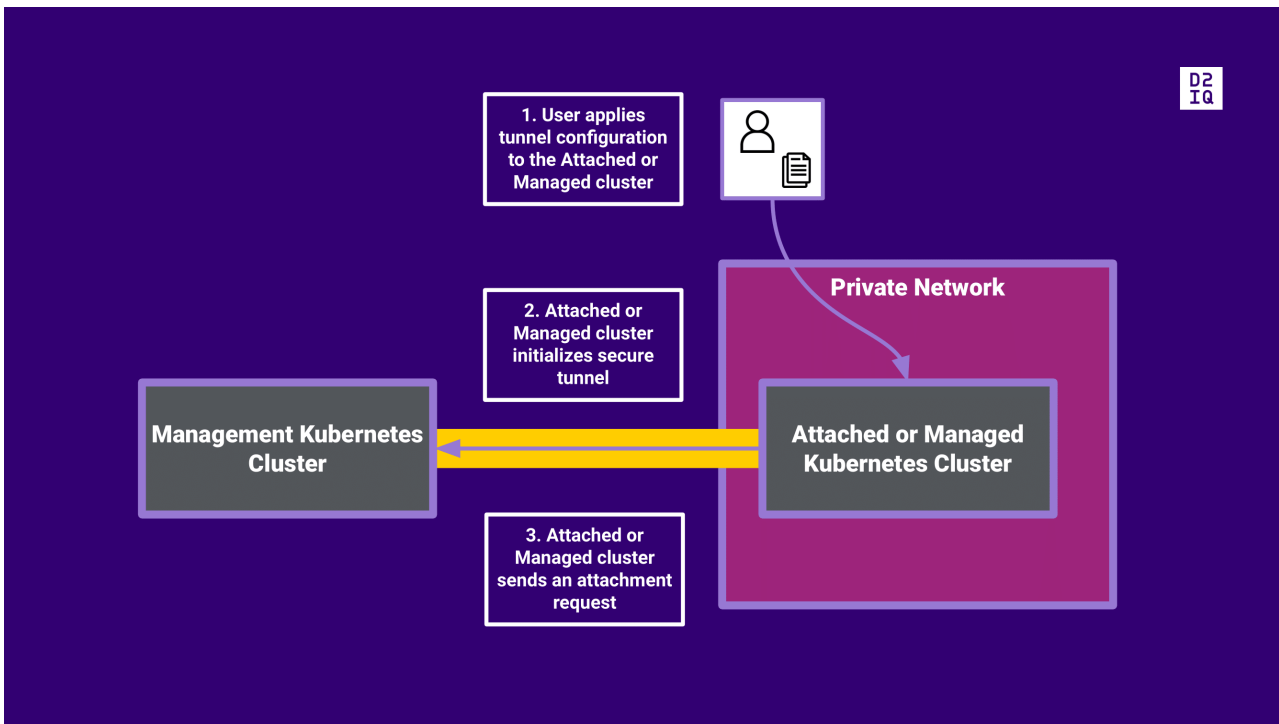
2 Diagram: Attachment not possible when cluster is in a private network.

Use a tunnel when you want to attach a cluster that is in a DMZ, behind a NAT gateway, behind a proxy server, firewall, or requires additional access information.

6.5.3.5.2 How Does the Tunneled Attachment work?

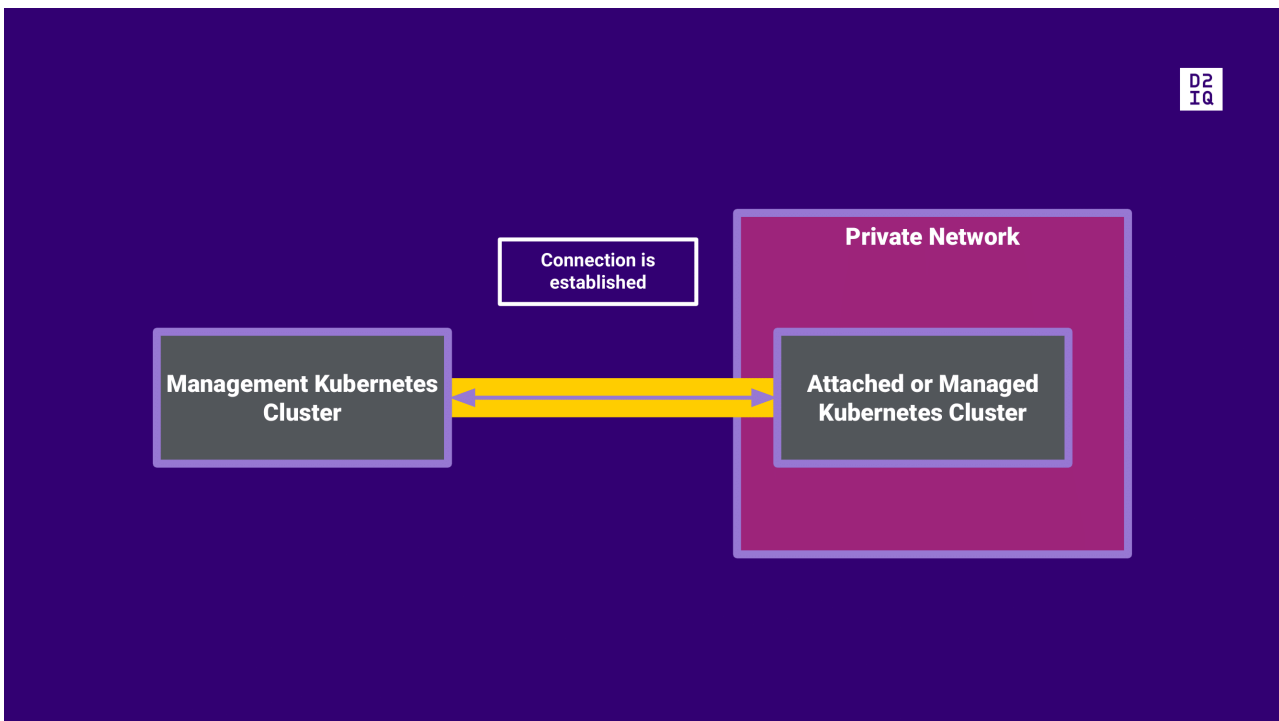
DKP can create a secure tunnel to enable the attachment of clusters that are not directly reachable.

To create a secure tunnel, you must provide a configuration for the tunnel in the cluster you want to attach. After you apply that configuration, the cluster you want to attach will establish a secure tunnel with the Management cluster, and make an attachment request.



3 Diagram: Attachment through a secure tunnel

After the attachment request is accepted and the connection between clusters is established, both clusters will allow bilateral communication.



4 Diagram: Connection between clusters is established

6.5.3.5.3 What are the Steps to Attach a Cluster via a Tunnel?

Ensure you meet the prerequisites for a tunneled attachment, and then attach the cluster with the UI or CLI:

- [Prerequisites for a Tunneled Attachment](#) (see page 775)
- [Attach a Cluster using a Tunnel via the UI](#) (see page 776)
- [Attach a Cluster Using a Tunnel via the CLI](#) (see page 787)
- [API documentation \(v1alpha1\)](#) (see page 803)
- [Proxied Access to Network-Restricted Clusters](#) (see page 814)

6.5.3.5.4 Prerequisites for a Tunneled Attachment

Enterprise

Gov Advanced

6.5.3.5.4.1 Before you Begin

- Gain more understanding of this approach by reviewing [Attach a Cluster with Networking Restrictions](#) (see page 775).
- Ensure you have reviewed the general [Requirements for Attaching an Existing Cluster](#) (see page 756).

6.5.3.5.4.2 Prerequisites

To enable a tunneled attachment, you have the following **additional** prerequisites:

- Ensure that `kubetunnel` is deployed on the Management Cluster.
Use the following command to check if `kubetunnel` is deployed:

```
kubectl get appdeployments.apps.kommander.d2iq.io -n kommander kubetunnel
```

The output should look similar to this:

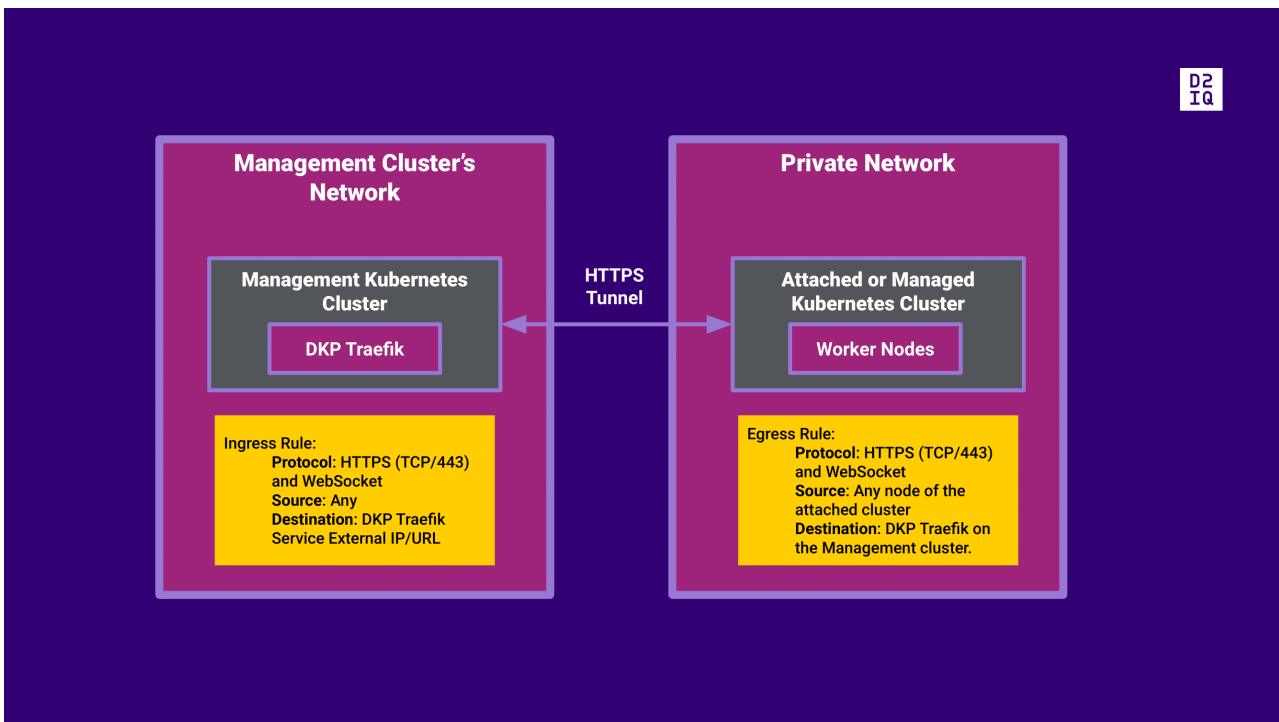
NAME	APP	AGE
kubetunnel	kubetunnel-<version>	5h14m



Kubetunnel is deployed by default. If you need to install it manually, see [Deploy Platform Applications via CLI](#) (see page 642) and ensure you install it on the Management cluster.

• **Firewall rules:**

	The ingress rule on the Management cluster network must allow:	The egress rule on the Attached or Managed cluster private network must allow:
Protocol	HTTPS (TCP/443) and WebSocket	HTTPS (TCP/443) and WebSocket
Source	Any	Any node of the Attached or Managed cluster
Destination	DKP Traefik Service External IP/URL	DKP Traefik Service on the Management cluster



Choose your Next Step:

[Attach a Cluster using a Tunnel via the UI \(see page 776\)](#)

[Attach a Cluster Using a Tunnel via the CLI \(see page 787\)](#)

6.5.3.5.5 Attach a Cluster using a Tunnel via the UI

Use the UI to attach a [network-restricted cluster](#) (see page 84)



Ensure you have reviewed and followed the steps in [Prerequisites for a Tunneled Attachment](#) (see page 778) before proceeding.

6.5.3.5.5.1 Attach a Cluster Using a Tunnel

Here is a high-level overview of the steps required to attach a cluster via the UI:

- [UI: Attach a Network-Restricted Cluster](#) (see page 777)
- [UI: Finish Attaching the Existing Cluster](#) (see page 779)
- [Use the Network-restricted Cluster](#) (see page 780)

6.5.3.5.5.2 UI: Attach a Network-Restricted Cluster

Attach a Cluster

To attach a [network-restricted cluster](#) (see page 84) to your DKP landscape:

1. From the top menu bar, select your target workspace.
2. On the Dashboard page, select the **Add Cluster** option in the **Actions** dropdown menu at the top right.
3. Select **Attach Cluster**.
4. Select the **Cluster has networking restrictions** card to display the configuration page.

Configure the Attachment

Establish the configuration parameters for the attachment:

5. Enter the **Cluster Name** of the cluster you're attaching.
6. Create additional new **Labels** as needed.
7. Select the hostname that is the Ingress for the cluster from the **Load Balancer Hostname** dropdown menu. The hostname must match the Kommander Host cluster to which you are attaching your existing cluster with network restrictions.
8. Specify the **URL Path Prefix** for your Load Balancer Hostname. This URL path will serve as the prefix for the specific tunnel services you want to expose on the Kommander management cluster. If no value is specified, the value defaults to `/dkp/tunnel`.

Kommander uses Traefik 2 ingress, which requires explicit definition of strip prefix middleware as a Kubernetes API object, opposed to a simple annotation. Kommander provides default middleware that supports creating tunnels only on the `/dkp/tunnel` URL prefix. This is indicated by using the extra annotation, `traefik.ingress.kubernetes.io/router.middlewares: kommander-striprefixes-kubetunnel@kubernetescrd` as shown in the code sample that follows. If you

want to expose a tunnel on a different URL prefix, you must manage your own middleware configuration.

9. **Optional:** Enter a value for the **Hostname** field.
10. Provide a secret for your certificate in the **Root CA Certificate** drop-down menu.
 - a. For environments where the Management cluster uses a publicly-signed CA (like ZeroSSL or Let's Encrypt), select **Use Publicly Trusted CA**.
 - b. If you manually created a secret in advance, select it from the drop-down menu.
 - c. For all other cases, select **Create a new secret**. Then, execute the following command on the Management cluster to obtain the `caBundle` key:

```
kubectl get kommandercluster -n kommander host-cluster -o go-template='{{ .status.ingress.caBundle }}'
```

Copy and paste the output into the **Root CA Certificate** field.

11. Add any **Extra Annotations** as needed.

Optional: Enable a Proxied Access

Activate a [proxied access](#) (see page 814) to enable kubectl access and dashboard observability for the network-restricted cluster from the [Management cluster](#) (see page 82). For more information, see [Proxied Access to Network-Restricted Clusters](#) (see page 814).

12. Select **Show Advanced**.
13. Add a **Cluster Proxy Domain**.

- If you previously configured a [domain wildcard](#) (see page 823) for your cluster, a Cluster Proxy Domain is suggested automatically based on your cluster name. Replace the suggestion if you want to assign a different domain for the proxied cluster.
- If you want to use the `external-dns` service, specify a **Cluster Proxy Domain** that is within the zones specified in the `--domain-filter` argument of the [external-dns deployment manifest](#) (see page 1443) stored on the **Management cluster**. For example, if the filter is set to `example.com`, a possible domain for the `TUNNEL_PROXY_EXTERNAL_DOMAIN` would be `myclusterproxy.example.com`.

14. Establish a DNS record and certificate configuration for the **Cluster Proxy Domain**. You can choose between the default and a custom option:

	DNS record creation	Certificate Management
--	----------------------------	-------------------------------

Default settings box checked <input checked="" type="checkbox"/>	Automatic, handled by <code>external-dns</code>	Automatic, handled by <code>kommander-ca</code>
Custom settings box unchecked <input type="checkbox"/>	Manually create a DNS record. The record's A/CNAME value must point to the Management cluster's Traefik IP address, URL or domain. OR Enable <code>external-dns</code> with an annotation that points to the Cluster Proxy Domain .	Select an existing TLS certificate. OR Select an existing Issuer or ClusterIssuer .

15. Select the **Save & Generate kubeconfig** button to generate a file required to finish attaching the cluster.

A new window appears with instructions on how to finalize attaching the cluster. See [UI: Finish Attaching the Existing Cluster](#) (see page 779) for further instructions.

Next Step:

[UI: Finish Attaching the Existing Cluster](#) (see page 779)

Related Topic:

For information on the TunnelGateway review the [API documentation \(v1alpha1\)](#) (see page 803).

6.5.3.5.5.3 UI: Finish Attaching the Existing Cluster

Enterprise

Gov Advanced

How to apply the kubeconfig file to create the network tunnel to attach a [network-restricted cluster](#) (see page 84)

After you have configured your cluster's attachment in [UI: Attach a Network-Restricted Cluster](#) (see page 780), finalize attaching the cluster. Now you must apply the generated manifest to create the network tunnel and complete the attachment process:

1. Select the **Download Manifest** link to download the file you [generated previously](#) (see page 780).
2. Copy the `kubectl apply` command from the UI and paste into your terminal session, DO NOT run it yet.
3. Ensure you substitute the actual name of the file for the variable. Also ensure you use the `--` `kubeconfig=<managed_cluster_kubeconfig.conf>` flag to run the command on the

Attached or Managed cluster. Run the command.

Note: Running this command starts the attachment process, which may take several minutes to complete. The Cluster details page will be displayed automatically when the cluster attachment process completes.

4. **Optional:** Select the **Verify Connection to Cluster** button to send a request to Kommander to refresh the connection information. You can use this option to check to see if the connection is complete, though the Cluster Details page displays automatically when the connection is complete.

After the initial connection is made, and your cluster becomes viewable as attached in the DKP UI, the attachment, federated add-ons, and platform services will still need to be completed. This may take several additional minutes. If a cluster has limited resources to deploy all the federated platform services, the installation of the federated resources will fail, and the cluster may become unreachable in the DKP UI. If this happens, check whether there are any pods that are not getting the resources required.

Next Step:

Now you can learn how to [Use the Network-restricted Cluster](#) (see page 780) (this approach uses the CLI).

6.5.3.5.5.4 Use the Network-restricted Cluster

Enterprise

Gov Advanced

To access services running on the remote, edge or [network-restricted cluster](#) (see page 84) from the Management cluster, connect to the tunnel proxy.

You can use these three methods:

1. If the client program supports use of a kubeconfig file, use the network-restricted cluster's kubeconfig.
2. If the client program supports SOCKS5 proxies, use the proxy directly.
3. Otherwise, deploy a proxy server on the Management cluster.

Network-restricted Cluster Service

These sections require a service to run on the Attached or Managed network-restricted cluster.

As an example, start the following service:

```
service_namespace=test
```

```

service_name=webserver
service_port=8888
service_endpoint=${service_name}.${service_namespace}.svc.cluster.local:${service_port}

cat > nginx.yaml <<EOF
apiVersion: v1
kind: Namespace
metadata:
  name: ${service_namespace}
---
apiVersion: apps/v1
kind: Deployment
metadata:
  namespace: ${service_namespace}
  name: nginx-deployment
  labels:
    app: nginx-deployment
spec:
  replicas: 3
  selector:
    matchLabels:
      app: nginx-app
  template:
    metadata:
      labels:
        app: nginx-app
    spec:
      containers:
        - name: nginx
          image: nginx:1.14.2
          ports:
            - containerPort: 80
---
apiVersion: v1
kind: Service
metadata:
  namespace: ${service_namespace}
  name: ${service_name}
spec:
  selector:
    app: nginx-app
  type: ClusterIP
  ports:
    - targetPort: 80
      port: ${service_port}
EOF

kubectl apply -f nginx.yaml

kubectl rollout status deploy -n ${service_namespace} nginx-deployment

```

On the Attached or Managed cluster, a client Job can access this service using:


```

cat > curl.yaml <<EOF
apiVersion: batch/v1
kind: Job
metadata:
  name: curl
spec:
  template:
    spec:
      containers:
      - name: curl
        image: curlimages/curl:7.76.0
        command: ["curl", "--silent", "--show-error", "http://${service_endpoint}"]
        restartPolicy: Never
      backoffLimit: 4
EOF

kubectl apply -f curl.yaml

kubectl wait --for=condition=complete job curl

podname=$(kubectl get pods --selector=job-name=curl --field-
selector=status.phase=Succeeded -o jsonpath='{.items[0].metadata.name}')

kubectl logs ${podname}

```

The final command returns the default Nginx web page:

```

<!DOCTYPE html>
<html>
<head>
<title>Welcome to nginx!</title>
<style>
  body {
    width: 35em;
    margin: 0 auto;
    font-family: Tahoma, Verdana, Arial, sans-serif;
  }
</style>
</head>
<body>
<h1>Welcome to nginx!</h1>
<p>If you see this page, the nginx web server is successfully installed and
working. Further configuration is required.</p>

<p>For online documentation and support please refer to
<a href="http://nginx.org/">nginx.org</a>.<br/>
Commercial support is available at
<a href="http://nginx.com/">nginx.com</a>.</p>

<p><em>Thank you for using nginx.</em></p>

```

```
</body>
</html>
```

Use of kubeconfig file

This is primarily useful for running `kubectl` commands on the Management cluster to monitor the network-restricted, Managed or Attached cluster.

On the Management cluster, a `kubeconfig` file for the Attached or Managed cluster configured to use the tunnel proxy is available as a Secret. The Secret's name can be identified using:

```
kubeconfig_secret=$(kubectl get tunnelconnector -n ${namespace} ${connector} -o
jsonpath='{.status.kubeconfigRef.name}')
```

After setting `service_namespace` and `service_name` to the service resource, run this command on the Management cluster:

```
cat > get-service.yaml <<EOF
apiVersion: batch/v1
kind: Job
metadata:
  name: get-service
spec:
  template:
    spec:
      containers:
      - name: kubectl
        image: bitnami/kubectl:1.19
        command: ["kubectl", "get", "service", "-n", "${service_namespace}", "${
service_name}"]
        env:
        - name: KUBECONFIG
          value: /tmp/kubeconfig/kubeconfig
      volumeMounts:
      - name: kubeconfig
        mountPath: /tmp/kubeconfig
      volumes:
      - name: kubeconfig
        secret:
          secretName: "${kubeconfig_secret}"
        restartPolicy: Never
      backoffLimit: 4
EOF

kubectl apply -n ${namespace} -f get-service.yaml

kubectl wait --for=condition=complete --timeout=5m job -n ${namespace} get-service
```

```
podname=$(kubectl get pods -n ${namespace} --selector=job-name=get-service --field-selector=status.phase=Succeeded -o jsonpath='{.items[0].metadata.name}')
```

```
kubectl logs -n ${namespace} ${podname}
```

Direct use of SOCKS5 proxy

To use the SOCKS5 proxy directly, obtain the SOCKS5 proxy endpoint using:

```
proxy_service=$(kubectl get tunnelconnector -n ${namespace} ${connector} -o jsonpath='{.status.tunnelServer.serviceRef.name}')
```

```
socks_proxy=$(kubectl get service -n ${namespace} "${proxy_service}" -o jsonpath='{.spec.clusterIP}{":"}{.spec.ports[?(@.name=="proxy")].port}')
```

Provide the value of `${socks_proxy}` as the SOCKS5 proxy to your client.

For example, since `curl` supports SOCKS5 proxies, the Attached or Managed service started above can be accessed from the Management cluster by adding the SOCKS5 proxy to the `curl` command. After setting `service_endpoint` to the service endpoint, on the Management cluster run:

```
cat > curl.yaml <<EOF
apiVersion: batch/v1
kind: Job
metadata:
  name: curl
spec:
  template:
    spec:
      containers:
      - name: curl
        image: curlimages/curl:7.76.0
        command: ["curl", "--silent", "--show-error", "--socks5-hostname", "${socks_proxy}", "http://${service_endpoint}"]
        restartPolicy: Never
      backoffLimit: 4
EOF
```

```
kubectl apply -f curl.yaml
```

```
kubectl wait --for=condition=complete --timeout=5m job curl
```

```
podname=$(kubectl get pods --selector=job-name=curl --field-selector=status.phase=Succeeded -o jsonpath='{.items[0].metadata.name}')
```

```
kubectl logs ${podname}
```

The final command returns the same output as for the job on the Attached or Managed cluster, demonstrating that the job on the Management cluster accessed the service running on the Attached or Managed cluster.

Use of deployed proxy on Management cluster

To deploy a proxy on the Management cluster, obtain the SOCKS5 proxy endpoint using:

```
proxy_service=$(kubectl get tunnelconnector -n ${namespace} ${connector} -o
jsonpath='{.status.tunnelServer.serviceRef.name}')

socks_proxy=$(kubectl get service -n ${namespace} "${proxy_service}" -o jsonpath='{.spec.clusterIP}{" ":"}{.spec.ports[?(@.name=="proxy")].port}')
```

Provide the value of `${socks_proxy}` as the SOCKS5 proxy to a proxy deployed on the Management cluster. After setting `service_endpoint` to the service endpoint, on the Management cluster run:

```
cat > nginx-proxy.yaml <<EOF
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
  name: nginx-proxy-crt
spec:
  secretName: nginx-proxy-crt-secret
  dnsNames:
  - nginx-proxy-service.${namespace}.svc.cluster.local
  issuerRef:
    group: cert-manager.io
    kind: ClusterIssuer
    name: kubernetes-ca
---
apiVersion: apps/v1
kind: Deployment
metadata:
  name: nginx-proxy
  labels:
    app: nginx-proxy-deployment
spec:
  replicas: 1
  selector:
    matchLabels:
      app: nginx-proxy-app
  template:
    metadata:
      labels:
        app: nginx-proxy-app
    spec:
      containers:
```

```

- name: nginx-proxy
  image: mesosphere/ghostunnel:v1.5.3-server-backend-proxy
  args:
  - "server"
  - "--listen=:443"
  - "--target=${service_endpoint}"
  - "--cert=/etc/certs/tls.crt"
  - "--key=/etc/certs/tls.key"
  - "--cacert=/etc/certs/ca.crt"
  - "--unsafe-target"
  - "--disable-authentication"
  env:
  - name: ALL_PROXY
    value: socks5://${socks_proxy}
  ports:
  - containerPort: 443
  volumeMounts:
  - name: certs
    mountPath: /etc/certs
  volumes:
  - name: certs
    secret:
      secretName: nginx-proxy-crt-secret
---
apiVersion: v1
kind: Service
metadata:
  name: nginx-proxy-service
spec:
  selector:
    app: nginx-proxy-app
  type: ClusterIP
  ports:
  - targetPort: 443
    port: 8765
EOF

kubectl apply -n ${namespace} -f nginx-proxy.yaml

kubectl rollout status deploy -n ${namespace} nginx-proxy

proxy_port=$(kubectl get service -n ${namespace} nginx-proxy-service -o
jsonpath='{.spec.ports[0].port}')

```

Any client running on the Management cluster can now access the service running on the Attached or Managed cluster using the proxy service endpoint. Note in the following that the `curl` job runs in the same namespace as the proxy, to provide access to the CA certificate secret.

```

cat > curl.yaml <<EOF
apiVersion: batch/v1
kind: Job

```

```

metadata:
  name: curl
spec:
  template:
    spec:
      containers:
      - name: curl
        image: curlimages/curl:7.76.0
        command:
        - curl
        - --silent
        - --show-error
        - --cacert
        - /etc/certs/ca.crt
        - https://nginx-proxy-service.${namespace}.svc.cluster.local:${proxy_port}
      volumeMounts:
      - name: certs
        mountPath: /etc/certs
      volumes:
      - name: certs
        secret:
          secretName: nginx-proxy-crt-secret
        restartPolicy: Never
      backoffLimit: 4
EOF

kubectl apply -n ${namespace} -f curl.yaml

kubectl wait --for=condition=complete --timeout=5m job -n ${namespace} curl

podname=$(kubectl get pods -n ${namespace} --selector=job-name=curl --field-selector=status.phase=Succeeded -o jsonpath='{.items[0].metadata.name}')

kubectl logs -n ${namespace} ${podname}

```

The final command returns the same output as the job on the Attached or Managed cluster, demonstrating that the job on the Management cluster accessed the service running on the network-restricted cluster.

Related Topic:


For information on the TunnelGateway review the [API documentation \(v1alpha1\)](#) (see page 803).

6.5.3.5.6 Attach a Cluster Using a Tunnel via the CLI

Use the CLI to attach a [network-restricted cluster](#) (see page 84) with a Tunnel

Enterprise

Gov Advanced

 Ensure you have reviewed and followed the steps in [Prerequisites for a Tunneled Attachment](#) (see page 778) before proceeding.

6.5.3.5.6.1 Attach a Cluster Using a Tunnel

Here is a high-level overview of the steps required to attach a cluster via the CLI:

- [CLI: Prepare the Management Cluster](#) (see page 788)
- [CLI: Create and Configure the Tunnel](#) (see page 790)
- [CLI: Use the Network-Restricted Cluster](#) (see page 795)

6.5.3.5.6.2 CLI: Prepare the Management Cluster

Enterprise

Gov Advanced

Identify the Management Cluster Endpoint

Execute the following command on the Management cluster to obtain the hostname and CA certificate:

```
hostname=$(kubectl get service -n kommander kommander-traefik -o go-template='{{with index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}')
b64ca_cert=$(kubectl get secret -n cert-manager kommander-ca -o=go-template='{{index .data "tls.crt"}}')
```

Specify a Workspace Namespace

Obtain the desired workspace namespace on the Management cluster for the tunnel gateway:

```
namespace=$(kubectl get workspace default-workspace -o jsonpath="{.status.namespaceRef.name}")
```

Alternatively, you can create a new workspace instead of using an existing workspace:

Run the following command, and replace the `<workspace_name>` with the new workspace name:

```
workspace=<workspace_name>
```

Finish creating the workspace:

```
namespace=${workspace}

cat > workspace.yaml <<EOF
apiVersion: workspaces.kommander.mesosphere.io/v1alpha1
kind: Workspace
metadata:
  annotations:
    kommander.mesosphere.io/display-name: ${workspace}
  name: ${workspace}
spec:
  namespaceName: ${namespace}
EOF

kubectl apply -f workspace.yaml
```

You can verify the workspace exists using:

```
kubectl get workspace ${workspace}
```

Create a Tunnel Gateway

Create a [tunnel gateway](#) (see page 0) on the Management cluster to listen for tunnel agents on remote clusters:

- Kommander uses Traefik 2 ingress, which requires explicit definition of strip prefix middleware as a Kubernetes API object, opposed to a simple annotation. Kommander provides default middleware that supports creating tunnels only on the `/dkp/tunnel` URL prefix. This is indicated by using the extra annotation, `traefik.ingress.kubernetes.io/router.middlewares: kommander-striprefixes-kubetunnel@kubernetescrd` as shown in the code sample that follows. If you want to expose a tunnel on a different URL prefix, you must manage your own middleware configuration.

Establish variables for the certificate secret and gateway. Replace the `<gateway_name>` placeholder with the name of the gateway:

```
cacert_secret=kubetunnel-ca
gateway=<gateway_name>
```

Create the `Secret` and `TunnelGateway` objects:

```
cat > gateway.yaml <<EOF
```



```

apiVersion: v1
kind: Secret
metadata:
  namespace: ${namespace}
  name: ${cacert_secret}
data:
  ca.crt:
    ${b64ca_cert}
---
apiVersion: kubernetes.d2iq.io/v1alpha1
kind: TunnelGateway
metadata:
  namespace: ${namespace}
  name: ${gateway}
spec:
  ingress:
    caSecretRef:
      namespace: ${namespace}
      name: ${cacert_secret}
    loadBalancer:
      hostname: ${hostname}
      urlPathPrefix: /dkp/tunnel
    extraAnnotations:
      kubernetes.io/ingress.class: kommander-traefik
      traefik.ingress.kubernetes.io/router.tls: "true"
      traefik.ingress.kubernetes.io/router.middlewares: kommander-striprefixes-
kubetunnel@kubernetescrd
EOF

kubectl apply -f gateway.yaml

```

You can verify the gateway exists using the command:

```
kubectl get tunnelgateway -n ${namespace} ${gateway}
```

Next Step:

[CLI: Create and Configure the Tunnel \(see page 790\)](#)

6.5.3.5.6.3 CLI: Create and Configure the Tunnel

Enterprise

Gov Advanced

Connect a remote, edge or network-restricted cluster

Create a Tunnel Connector

Create a [tunnel connector](#) (see page 810) on the **Management cluster** for the remote cluster.

1. Establish a variable for the connector. Provide the name of the connector, by replacing the `<connector_name>` placeholder:

```
connector=<connector_name>
```

2. Create the `TunnelConnector` object:

```
cat > connector.yaml <<EOF
apiVersion: kubetunnel.d2iq.io/v1alpha1
kind: TunnelConnector
metadata:
  namespace: ${namespace}
  name: ${connector}
spec:
  gatewayRef:
    name: ${gateway}
EOF

kubectl apply -f connector.yaml
```

After you create the `TunnelConnector` object, DKP creates a `manifest.yaml`. This `manifest.yaml` contains the configuration information for the components required by the tunnel for a specific cluster.

3. Verify the connector exists:

```
kubectl get tunnelconnector -n ${namespace} ${connector}
```

4. Wait for the tunnel connector to reach the `Listening` state and export the agent manifest:

```
while [ "$(kubectl get tunnelconnector -n ${namespace} ${connector} -o
jsonpath="{.status.state}")" != "Listening" ]
do
  sleep 5
done

manifest=$(kubectl get tunnelconnector -n ${namespace} ${connector} -o
jsonpath="{.status.tunnelAgent.manifestsRef.name}")
while [ -z ${manifest} ]
do
  sleep 5
  manifest=$(kubectl get tunnelconnector -n ${namespace} ${connector} -o
jsonpath="{.status.tunnelAgent.manifestsRef.name}")
done
```

The `manifest.yaml` is applied successfully after the command completes.

5. Fetch the `manifest.yaml` to use it in the following section:

```
kubectl get secret -n ${namespace} ${manifest} -o
jsonpath='{.data.manifests\.yaml}' | base64 -d > manifest.yaml
```

- When attaching several clusters, ensure that you fetch the `manifest.yaml` of the cluster you are attempting to attach. Using the wrong combination of `manifest.yaml` and cluster will cause the attachment to fail.

Set up the Network-restricted Cluster

- In the following commands, the `--kubeconfig` flag ensures that you set the context to the Attached or Managed cluster. For alternatives and recommendations around setting your context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

1. Apply the `manifest.yaml` file to the **Attached or Managed** cluster and deploy the tunnel agent:

```
kubectl apply --kubeconfig=<managed_cluster_kubeconfig.conf> -f manifest.yaml
```

2. Check the status of the created pods using:

```
kubectl get pods --kubeconfig=<managed_cluster_kubeconfig.conf> -n kubetunnel
```

After a short time, expect to see a `post-kubeconfig` pod that reaches `Completed` state and a `tunnel-agent` pod that stays in `Running` state.

NAME	READY	STATUS	RESTARTS	AGE
post-kubeconfig-j2ghk	0/1	Completed	0	14m
tunnel-agent-f8d9f4cb4-thx8h	1/1	Running	0	14m

Add the Network-restricted Cluster into Kommander

When you create a cluster using the DKP CLI, it does not attach automatically.

1. On the Management cluster, wait for the tunnel to be connected by the tunnel agent:

```
while [ "$(kubectl get tunnelconnector -n ${namespace} ${connector} -o
jsonpath="{.status.state}")" != "Connected" ]
do
  sleep 5
done
```

2. Establish variables for the managed cluster. Replace the `<private_cluster>` placeholder with the name of the managed cluster:

```
managed=<private-cluster>
display_name=${managed}
```

3. Update the `KommanderCluster` object:

```
cat > kommander.yaml <<EOF
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  namespace: ${namespace}
  name: ${managed}
  annotations:
    kommander.mesosphere.io/display-name: ${display_name}
spec:
  clusterTunnelConnectorRef:
    name: ${connector}
EOF

kubectl apply -f kommander.yaml
```

4. Wait for the Attached or Managed cluster to join:

```
while [ "$(kubectl get kommandercluster -n ${namespace} ${managed} -o
jsonpath='{.status.phase}')" != "Joined" ]
do
  sleep 5
done

kubefed=$(kubectl get kommandercluster -n ${namespace} ${managed} -o
jsonpath="{.status.kubefedclusterRef.name}")
while [ -z "${kubefed}" ]
do
  sleep 5
```


```
kubefed=$(kubectl get kommandercluster -n ${namespace} ${managed} -o
jsonpath="{.status.kubefedclusterRef.name}")
done

kubectl wait --for=condition=ready --timeout=60s kubefedcluster -n kube-
federation-system ${kubefed}
```

After the command completes, your cluster becomes visible in the DKP UI and you can start using it. Its metrics will be accessible through different dashboards such as Grafana, Karma, etc.

Create a Network Policy for the Tunnel Server

This step is optional but improves **security** by restricting which remote hosts can connect to the tunnel.

1. Apply a network policy that restricts tunnel access to specific namespaces and IP blocks.
 -  The following *example* permits connections from
 - Pods running in the `kommander` and `kube-federation-system` namespace.
 - Remote clusters with IP addresses in the ranges 192.0.2.0 to 192.0.2.255 and 203.0.113.0 to 203.0.113.255.
 - Pods running in namespaces with a label `kubetunnel.d2iq.io/networkpolicy` that match the tunnel name and namespace.

```
cat > net.yaml <<EOF
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  namespace: ${namespace}
  name: ${connector}-deny
  labels:
    kubetunnel.d2iq.io/tunnel-connector: ${connector}
    kubetunnel.d2iq.io/networkpolicy-type: "tunnel-server"
spec:
  podSelector:
    matchLabels:
      kubetunnel.d2iq.io/tunnel-connector: ${connector}
  policyTypes:
  - Ingress
---
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  namespace: ${namespace}
  name: ${connector}-allow
  labels:
    kubetunnel.d2iq.io/tunnel-connector: ${connector}
    kubetunnel.d2iq.io/networkpolicy-type: "tunnel-server"
spec:
  podSelector:
```

```

    matchLabels:
      kubetunnel.d2iq.io/tunnel-connector: ${connector}
policyTypes:
- Ingress
ingress:
- from:
  - namespaceSelector:
      matchLabels:
        kubernetes.io/metadata.name: "kube-federation-system"
  - namespaceSelector:
      matchLabels:
        kubernetes.io/metadata.name: "kommander"
  - namespaceSelector:
      matchLabels:
        kubetunnel.d2iq.io/networkpolicy: ${connector}-${namespace}
- ipBlock:
    cidr: 192.0.2.0/24
- ipBlock:
    cidr: 203.0.113.0/24
EOF
kubectl apply -f net.yaml

```

- To enable applications running in another namespace to access the attached cluster, add the label `kubetunnel.d2iq.io/networkpolicy=${connector}-${namespace}` to the target namespace:

```
kubectl label ns ${namespace} kubetunnel.d2iq.io/networkpolicy=${connector}-${namespace}
```

All pods in the target namespace can now reach the attached cluster services.

Next Steps:

- Optional:** If you want to access the network-restricted attached cluster from the Management cluster, [enable proxied access](#) (see page 816).
- Alternatively,** start [using your remote cluster](#) (see page 783).

6.5.3.5.6.4 CLI: Use the Network-Restricted Cluster

Enterprise

Gov Advanced

To access services running on the remote, edge or [network-restricted cluster](#) (see page 84) from the Management cluster, connect to the tunnel proxy.

You can use these three methods:

1. If the client program supports use of a kubeconfig file, use the network-restricted cluster's kubeconfig.
2. If the client program supports SOCKS5 proxies, use the proxy directly.
3. Otherwise, deploy a proxy server on the Management cluster.

Network-restricted Cluster Service

These sections require a service to run on the Attached or Managed network-restricted cluster.

As an example, start the following service:

```

service_namespace=test
service_name=webserver
service_port=8888
service_endpoint=${service_name}.${service_namespace}.svc.cluster.local:${
service_port}

cat > nginx.yaml <<EOF
apiVersion: v1
kind: Namespace
metadata:
  name: ${service_namespace}
---
apiVersion: apps/v1
kind: Deployment
metadata:
  namespace: ${service_namespace}
  name: nginx-deployment
  labels:
    app: nginx-deployment
spec:
  replicas: 3
  selector:
    matchLabels:
      app: nginx-app
  template:
    metadata:
      labels:
        app: nginx-app
    spec:
      containers:
        - name: nginx
          image: nginx:1.14.2
          ports:
            - containerPort: 80
---
apiVersion: v1
kind: Service
metadata:

```

```

namespace: ${service_namespace}
name: ${service_name}
spec:
  selector:
    app: nginx-app
  type: ClusterIP
  ports:
  - targetPort: 80
    port: ${service_port}
EOF

kubectl apply -f nginx.yaml

kubectl rollout status deploy -n ${service_namespace} nginx-deployment

```

On the Attached or Managed cluster, a client Job can access this service using:

```

cat > curl.yaml <<EOF
apiVersion: batch/v1
kind: Job
metadata:
  name: curl
spec:
  template:
    spec:
      containers:
      - name: curl
        image: curlimages/curl:7.76.0
        command: ["curl", "--silent", "--show-error", "http://${service_endpoint}"]
        restartPolicy: Never
      backoffLimit: 4
EOF

kubectl apply -f curl.yaml

kubectl wait --for=condition=complete job curl

podname=$(kubectl get pods --selector=job-name=curl --field-
selector=status.phase=Succeeded -o jsonpath='{.items[0].metadata.name}')

kubectl logs ${podname}

```

The final command returns the default Nginx web page:

```

<!DOCTYPE html>
<html>
<head>
<title>Welcome to nginx!</title>
<style>
  body {

```



```

        width: 35em;
        margin: 0 auto;
        font-family: Tahoma, Verdana, Arial, sans-serif;
    }
</style>
</head>
<body>
<h1>Welcome to nginx!</h1>
<p>If you see this page, the nginx web server is successfully installed and
working. Further configuration is required.</p>

<p>For online documentation and support please refer to
<a href="http://nginx.org/">nginx.org</a>.<br/>
Commercial support is available at
<a href="http://nginx.com/">nginx.com</a>.</p>

<p><em>Thank you for using nginx.</em></p>
</body>
</html>

```

Use of kubeconfig file

This is primarily useful for running `kubectl` commands on the Management cluster to monitor the network-restricted, Managed or Attached cluster.

On the Management cluster, a `kubeconfig` file for the Attached or Managed cluster configured to use the tunnel proxy is available as a Secret. The Secret's name can be identified using:

```
kubeconfig_secret=$(kubectl get tunnelconnector -n ${namespace} ${connector} -o
jsonpath='{.status.kubeconfigRef.name}')
```

After setting `service_namespace` and `service_name` to the service resource, run this command on the Management cluster:

```

cat > get-service.yaml <<EOF
apiVersion: batch/v1
kind: Job
metadata:
  name: get-service
spec:
  template:
    spec:
      containers:
      - name: kubectl
        image: bitnami/kubectl:1.19
        command: ["kubectl", "get", "service", "-n", "${service_namespace}", "${
service_name}"]
        env:

```

```

- name: KUBECONFIG
  value: /tmp/kubeconfig/kubeconfig
volumeMounts:
- name: kubeconfig
  mountPath: /tmp/kubeconfig
volumes:
- name: kubeconfig
  secret:
    secretName: "${kubeconfig_secret}"
  restartPolicy: Never
  backoffLimit: 4
EOF

kubectl apply -n ${namespace} -f get-service.yaml

kubectl wait --for=condition=complete --timeout=5m job -n ${namespace} get-service
podname=$(kubectl get pods -n ${namespace} --selector=job-name=get-service --field-selector=status.phase=Succeeded -o jsonpath='{.items[0].metadata.name}')

kubectl logs -n ${namespace} ${podname}

```

Direct use of SOCKS5 proxy

To use the SOCKS5 proxy directly, obtain the SOCKS5 proxy endpoint using:

```

proxy_service=$(kubectl get tunnelconnector -n ${namespace} ${connector} -o
jsonpath='{.status.tunnelServer.serviceRef.name}')

socks_proxy=$(kubectl get service -n ${namespace} "${proxy_service}" -o jsonpath='{.spec.clusterIP}{" ":"}{.spec.ports[?(@.name=="proxy")].port}')

```

Provide the value of `${socks_proxy}` as the SOCKS5 proxy to your client.

For example, since `curl` supports SOCKS5 proxies, the Attached or Managed service started above can be accessed from the Management cluster by adding the SOCKS5 proxy to the `curl` command. After setting `service_endpoint` to the service endpoint, on the Management cluster run:

```

cat > curl.yaml <<EOF
apiVersion: batch/v1
kind: Job
metadata:
  name: curl
spec:
  template:
    spec:
      containers:
      - name: curl

```

```

    image: curlimages/curl:7.76.0
    command: ["curl", "--silent", "--show-error", "--socks5-hostname", "${socks_proxy}", "http://${service_endpoint}"]
    restartPolicy: Never
    backoffLimit: 4
EOF
kubectl apply -f curl.yaml

kubectl wait --for=condition=complete --timeout=5m job curl

podname=$(kubectl get pods --selector=job-name=curl --field-selector=status.phase=Succeeded -o jsonpath='{.items[0].metadata.name}')

kubectl logs ${podname}

```

The final command returns the same output as for the job on the Attached or Managed cluster, demonstrating that the job on the Management cluster accessed the service running on the Attached or Managed cluster.

Use of deployed proxy on Management cluster

To deploy a proxy on the Management cluster, obtain the SOCKS5 proxy endpoint using:

```

proxy_service=$(kubectl get tunnelconnector -n ${namespace} ${connector} -o jsonpath='{.status.tunnelServer.serviceRef.name}')

socks_proxy=$(kubectl get service -n ${namespace} "${proxy_service}" -o jsonpath='{.spec.clusterIP}{" ":"}{.spec.ports[?(@.name=="proxy")].port}')

```

Provide the value of `${socks_proxy}` as the SOCKS5 proxy to a proxy deployed on the Management cluster. After setting `service_endpoint` to the service endpoint, on the Management cluster run:

```

cat > nginx-proxy.yaml <<EOF
apiVersion: cert-manager.io/v1
kind: Certificate
metadata:
  name: nginx-proxy-crt
spec:
  secretName: nginx-proxy-crt-secret
  dnsNames:
  - nginx-proxy-service.${namespace}.svc.cluster.local
  issuerRef:
    group: cert-manager.io
    kind: ClusterIssuer
    name: kubernetes-ca
---
apiVersion: apps/v1

```

```

kind: Deployment
metadata:
  name: nginx-proxy
  labels:
    app: nginx-proxy-deployment
spec:
  replicas: 1
  selector:
    matchLabels:
      app: nginx-proxy-app
  template:
    metadata:
      labels:
        app: nginx-proxy-app
    spec:
      containers:
      - name: nginx-proxy
        image: mesosphere/ghostunnel:v1.5.3-server-backend-proxy
        args:
          - "server"
          - "--listen=:443"
          - "--target=${service_endpoint}"
          - "--cert=/etc/certs/tls.crt"
          - "--key=/etc/certs/tls.key"
          - "--cacert=/etc/certs/ca.crt"
          - "--unsafe-target"
          - "--disable-authentication"
        env:
          - name: ALL_PROXY
            value: socks5://${socks_proxy}
        ports:
          - containerPort: 443
        volumeMounts:
          - name: certs
            mountPath: /etc/certs
      volumes:
      - name: certs
        secret:
          secretName: nginx-proxy-crt-secret
---
apiVersion: v1
kind: Service
metadata:
  name: nginx-proxy-service
spec:
  selector:
    app: nginx-proxy-app
  type: ClusterIP
  ports:
    - targetPort: 443
      port: 8765
EOF

```

```
kubectl apply -n ${namespace} -f nginx-proxy.yaml

kubectl rollout status deploy -n ${namespace} nginx-proxy

proxy_port=$(kubectl get service -n ${namespace} nginx-proxy-service -o
jsonpath='{.spec.ports[0].port}')
```

Any client running on the Management cluster can now access the service running on the Attached or Managed cluster using the proxy service endpoint. Note in the following that the `curl` job runs in the same namespace as the proxy, to provide access to the CA certificate secret.

```
cat > curl.yaml <<EOF
apiVersion: batch/v1
kind: Job
metadata:
  name: curl
spec:
  template:
    spec:
      containers:
      - name: curl
        image: curlimages/curl:7.76.0
        command:
        - curl
        - --silent
        - --show-error
        - --cacert
        - /etc/certs/ca.crt
        - https://nginx-proxy-service.${namespace}.svc.cluster.local:${proxy_port}
        volumeMounts:
        - name: certs
          mountPath: /etc/certs
      volumes:
      - name: certs
        secret:
          secretName: nginx-proxy-crt-secret
        restartPolicy: Never
      backoffLimit: 4
EOF

kubectl apply -n ${namespace} -f curl.yaml

kubectl wait --for=condition=complete --timeout=5m job -n ${namespace} curl

podname=$(kubectl get pods -n ${namespace} --selector=job-name=curl --field-
selector=status.phase=Succeeded -o jsonpath='{.items[0].metadata.name}')

kubectl logs -n ${namespace} ${podname}
```

The final command returns the same output as the job on the Attached or Managed cluster, demonstrating that the job on the Management cluster accessed the service running on the network-restricted cluster.

Next Step:

Optional: If you want to manage the attached cluster from the Management cluster, [enable proxied access](#) (see page 816).

Related Topic:

For information on the TunnelGateway review the [API documentation \(v1alpha1\)](#) (see page 803).

6.5.3.5.7 API documentation (v1alpha1)

API Documentation (v1alpha1)

This document is automatically generated from the API definition in the code.

Page Contents

6.5.3.5.7.1 CertificateSpec

CertificateSpec holds settings for the TunnelProxy's ingress exposed Certificate.

Field	Description	Scheme	Required
issuerRef	IssuerRef is a reference to an <code>Issuer</code> or <code>ClusterIssuer</code> on the target cluster to be used to create a <code>Certificate</code> for the cluster's Ingress. If the type is an <code>Issuer</code> , it must be located in the namespace of the <code>TunnelProxy</code> .	IssuerReference (see page 804)	false
certificateSecretRef	CertificateSecretRef is a reference to a secret of type <code>TLS</code> that holds a TLS certificate, private key and CA certificate. The certificate must be valid for the Ingress hostname. The secret must be located in the same namespace of the <code>TunnelProxy</code> .	<code>v1.LocalObjectReference</code>	false

6.5.3.5.7.2 IngressSpec

IngressSpec holds settings for the TunnelProxy managed Ingress.

Field	Description	Scheme	Required
annotations	Annotations allows to set-up important metadata to configure DNS records and certificate generation	map[string]string	false
certificate	Certificate holds settings for the cluster's TunnelProxy exposed Certificate.	CertificateSpec (see page 806)	false

6.5.3.5.7.3 IssuerReference

IssuerReference is a reference to an issuer with a given name, kind and group.

Field	Description	Scheme	Required
name	Name of the issuer being referred to.	string	true
kind	Kind of the issuer being referred to.	string	false
group	Group of the issuer being referred to.	string	false

6.5.3.5.7.4 TunnelProxy

Describes the local endpoint for the tunnel. A remote cluster will connect to this endpoint to create a tunnel.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ⁵⁸⁴	false
spec		TunnelProxySpec (see page 805)	false

⁵⁸⁴ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.26/#objectmeta-v1-meta>

Field	Description	Scheme	Required
status		TunnelProxyStatus (see page 806)	false

6.5.3.5.7.5 TunnelProxyList

Contains a list of `TunnelProxy`.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ⁵⁸⁵	false
items		[]TunnelProxy (see page 807)	true

6.5.3.5.7.6 TunnelProxySpec

`TunnelProxy` describes the exposition of an remote `KommanderCluster` through the Management cluster.

Field	Description	Scheme	Required
clusterProxyDomain	ClusterProxyDomain sets the desired domain to expose the remote cluster. It expects a domain name without scheme and without port. Ex. "traefik.kommander.cluster", "traefik.kommander.cluster" "10.0.0.1"	string	true
caBundle	CABundle is a PEM encoded CA bundle which should be used to verify the TLS connection for the Ingress-served end-entity certificate.	[]byte	false

⁵⁸⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.26/#listmeta-v1-meta>

Field	Description	Scheme	Required
remoteClusterEndpoint	RemoteClusterEndpoint points to the remote cluster Endpoint (i.e. Traefik instance). It expects a URL without scheme and (optionally) a port. Ex. "traefik.kommander.cluster", "traefik.kommander.cluster:8080" "10.0.0.1"	string	false
ingress	Ingress holds settings for the TunnelProxy managed Ingress for the proxied cluster.	IngressSpec (see page 807)	false
tunnelConnectorRef	TunnelConnectorRef points to the TunnelConnector object to be used by the ReverseProxy	v1.LocalObjectReference	true

6.5.3.5.7.7 TunnelProxyStatus

Field	Description	Scheme	Required
conditions	Conditions contains the status conditions of the object.	[]metav1.Condition	false
clusterProxyDomain	ClusterProxyDomain returns the actual domain to reach the remote cluster.	string	true
clientAuthSecretName	ClientAuthSecretName returns the name of the secret containing the mTLS client-auth Certificate Authority.	string	true
reverseProxyReleaseName	ReverseProxyReleaseName returns the name of the Helm release for the ReverseProxy.	string	true

6.5.3.5.7.8 TunnelGateway

Provides an endpoint for remote clusters to connect to the management cluster.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ⁵⁸⁶	false
spec		TunnelGatewaySpec (see page 809)	false

6.5.3.5.7.9 TunnelGatewayIngressSpec

Field	Description	Scheme	Required
loadBalancer	Ingress point for the load-balancer. Traffic intended for the service should be sent to these ingress points. If not specified, the controller will derive from the Ingress record status field.	corev1.LoadBalancerIngress	false
host	Restrict access to requests addressed to a specific host or domain using the <code>IngressRule</code> format. Defaults to allow all hosts.	string	false

⁵⁸⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.26/#objectmeta-v1-meta>

Field	Description	Scheme	Required
urlPathPrefix	URL path prefix to prepend to all endpoints. For example, if this field is set to <code>/ops/portal/kt</code> , the ingresses created will have URL paths like <code>/ops/portal/kt/default/cluster1/tunnel-server</code> and <code>/ops/portal/kt/default/cluster1/kubeconfig</code> . Defaults to root path (<code>/</code>).	string	false
caSecretRef	A secret reference to the root CA required to verify the ingress endpoints. The secret should have type <code>Opaque</code> and contain the key <code>ca.crt</code> . If not specified, remote hosts will use their system root CA's to verify the endpoints.	corev1.ObjectReference	false
extraAnnotations	Extra annotations to set on the Ingress object.	map[string]string	false

6.5.3.5.7.10 TunnelGatewayList

Contains a list of `TunnelGateway`.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ⁵⁸⁷	false
items		[]TunnelGateway	true

6.5.3.5.7.11 TunnelGatewaySpec

If no ingress is set, the services will only be accessible on `localhost`.

Field	Description	Scheme	Required
ingress	Expose services using an Ingress as specified in the <code>TunnelGatewayIngressSpec</code> .	TunnelGatewayIngressSpec (see page 810)	false

6.5.3.5.7.12 KubeconfigWebhookStatus

Status of the kubeconfig webhook.

Field	Description	Scheme	Required
deploymentRef	A reference to the deployment for the kubeconfig webhook.	corev1.LocalObjectReference	false
serviceRef	A reference to the service for the kubeconfig webhook.	corev1.LocalObjectReference	false
ingressRef	A reference to the ingress for the kubeconfig webhook.	corev1.LocalObjectReference	false

⁵⁸⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.26/#listmeta-v1-meta>

6.5.3.5.7.13 TunnelAgentStatus

Status of the tunnel agent.

Field	Description	Scheme	Required
manifestsRef	A reference to a secret holding YAML manifests for launching the tunnel agent on the target cluster. The secret is a generic typed secret with filenames as the keys. There might be multiple files in the secret.	corev1.LocalObjectReference	false

6.5.3.5.7.14 TunnelConnector

Describes the local endpoint for the tunnel. A remote cluster will connect to this endpoint to create a tunnel.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ⁵⁸⁸	false
spec		TunnelConnectorSpec (see page 811)	false
status		TunnelConnectorStatus (see page 812)	false

6.5.3.5.7.15 TunnelConnectorList

Contains a list of `TunnelConnector`.

⁵⁸⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.26/#listmeta-v1-meta>

Field	Description	Scheme	Required
metadata		metav1.ListMeta ⁵⁸⁹	false
items		[]TunnelConnector (see page 813)	true

6.5.3.5.7.16 TunnelConnectorSpec

Field	Description	Scheme	Required
gatewayRef	A reference to the <code>TunnelGateway</code> object which describes how tunnel services will be exposed outside the current cluster.	<code>corev1.LocalObjectReference</code>	false
proxyPort	The port for the tunnel proxy.	<code>int32</code>	false

⁵⁸⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.26/#listmeta-v1-meta>

6.5.3.5.7.17 TunnelConnectorStatus

Field	Description	Scheme	Required
state	State of the tunnel connector: Starting - the initial state; Listening - the local tunnel server is waiting for the remote agent to connect; Pending - the remote agent has connected but the local proxy is not ready; Connected - the tunnel is configured and contact to the remote API server succeeded; Disconnected - the tunnel is configured but contact to the remote API server failed; Failed - an unexpected error occurred, such as not being able to parse the kubeconfig.	TunnelConnectorState	false
tunnelServer	Status of the tunnel server.	TunnelServerStatus (see page 813)	false
kubeconfigWebhook	Status of the kubeconfig webhook.	KubeconfigWebhookStatus (see page 812)	false
tunnelAgent	Status of the tunnel agent.	TunnelAgentStatus (see page 813)	false
serviceAccountRef	A reference to the service account that will be used for registration (of the tunnel agent) and authentication purpose.	corev1.LocalObjectReference	false

Field	Description	Scheme	Required
roleRef	A reference to the role that will be bound to the service account for authorization purpose.	corev1.LocalObjectReference	false
roleBindingRef	A reference to the rolebinding that will be created to bind the service account and the role.	corev1.LocalObjectReference	false
kubeconfigRef	A reference to the secret holding the KUBECONFIG that the clients can use to talk to the API server of the target cluster when it becomes available.	corev1.LocalObjectReference	false
gatewayObservedGeneration	The generation of the linked TunnelGateway object associated with this object. When the linked TunnelGateway object is updated, a controller will update this status field which will in turn trigger a reconciliation of this object.	int64	false

6.5.3.5.7.18 TunnelServerStatus

Status of the tunnel server.

Field	Description	Scheme	Required
deploymentRef	A reference to the deployment for the tunnel server.	corev1.LocalObjectReference	false

Field	Description	Scheme	Required
serviceRef	A reference to the service for the tunnel server.	corev1.LocalObjectReference	false
ingressRef	A reference to the ingress for the tunnel server.	corev1.LocalObjectReference	false

6.5.3.5.8 Proxied Access to Network-Restricted Clusters

Enterprise

Gov Advanced

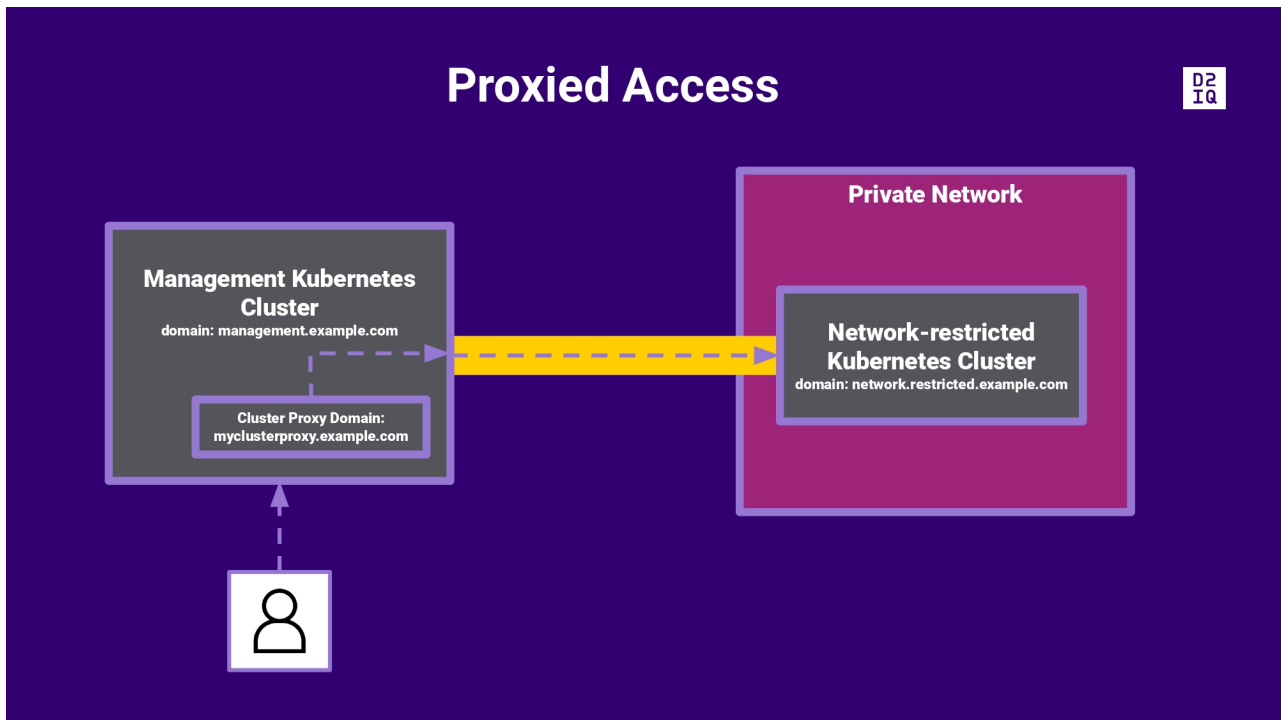
Enabling a proxied access allows you to access [Attached](#) (see page 82) and [Managed](#) (see page 82) clusters that are [network-restricted](#) (see page 84), in a private network, firewalled, or at edge.



This section only applies to clusters with networking restrictions that were [attached through a secure tunnel](#) (see page 775).

You can [attach clusters that are in a private network](#) (see page 775) (clusters that have networking restrictions or are at the edge). D2iQ provides the option of using a secure tunnel or a tunneled attachment to attach a Kubernetes cluster to the [Management cluster](#) (see page 82). To **access** these attached clusters through `kubectl` or monitor its resources through the Management cluster, you have to be in the same network, or **enable a proxied access**.

6.5.3.5.8.1 Proxied Access



5 Access to a network-restricted cluster through a proxy in the Management cluster

Enabling the proxied access for a [network-restricted cluster](#) (see page 84) makes it possible for DKP to authenticate user requests (regardless of the [identity provider](#) (see page 600)) through the Management cluster's authentication proxy. This is helpful, when the cluster you are trying to reach is in a different network. The proxied access allows you to:

- Access and observe the cluster's monitoring and logging services from the Management cluster, for example:
 - Access the cluster's Grafana, Kubernetes and KubeCost dashboards from the Management cluster.
 - Use the CLI to print a cluster's service URL, so you can access the cluster's dashboards.
- Access and perform operations on the network-restricted cluster from the Management cluster, for example:
 - Generate an API token (**Generate Token** option, from the upper right corner of the UI) that allows you to authenticate to the network-restricted cluster.
 - Upon authentication, use `kubect1` to manage the network-restricted cluster.

You can perform the previous actions, without being in the same network as the network-restricted cluster.

6.5.3.5.8.2 Next Step:

[Enable Proxied Access Using the CLI](#) (see page 816)

6.5.3.5.8.3 Enable Proxied Access Using the CLI

Enterprise

Gov Advanced

Enabling a proxied access allows you to access [Attached](#) (see page 82) and [Managed](#) (see page 82) clusters that are [network-restricted](#) (see page 84), in a private network, firewalled, or at edge.



This section only applies to clusters with networking restrictions that were [attached through a secure tunnel](#) (see page 775).

Continue with the following procedures to enable and configure a proxied access for your cluster:

- [1. Prerequisites: CLI Proxied Access](#) (see page 816)
- [2. Prepare your Environment: CLI Proxied Access](#) (see page 816)
- [3. Create a TunnelProxy Object: CLI Proxied Access](#) (see page 818)
- [4. Enable the TunnelProxy Object in the KommanderCluster: CLI Proxied Access](#) (see page 821)
- [5. Verify the Proxy: CLI Proxied Access](#) (see page 822)

1. Prerequisites: CLI Proxied Access

- You have [attached a network-restricted cluster](#) (see page 775).
- The [Management](#) (see page 82) and network-restricted cluster are on the same DKP version.
- You have a domain that you can use to put on top of the network-restricted cluster's domain to redirect user requests (Cluster Proxy Domain).
- A DNS record to map your domain to your cluster. There are two supported options for this:
 - **Manual** DNS record creation:
Create a DNS record manually. The record's A/CNAME value must point to the Management cluster's Traefik IP address, URL or domain. Use one record per proxied cluster.
 - **Automatic** DNS record creation:
A service that creates and maintains your DNS record automatically. For this method, enable the [external-dns service on the Management cluster](#) (see page 1443) before configuring the proxy.

The following pages walk you through enabling the proxied access on the network-restricted cluster.

Next Step:

[2. Prepare your Environment: CLI Proxied Access](#) (see page 816)

2. Prepare your Environment: CLI Proxied Access

Establish the following environment variables on the **Management cluster**.

i See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for more information around switching cluster contexts.

The following commands allow you to run most commands without replacing the information manually.

1. Set the `WORKSPACE_NAMESPACE` environment variable to the name of your network-restricted cluster's workspace namespace:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```

2. Set the variable to the proxy domain through which your cluster should be available:

```
TUNNEL_PROXY_EXTERNAL_DOMAIN=<myclusterproxy.example.com>
```

i If you want to use the `external-dns` service, specify a `TUNNEL_PROXY_EXTERNAL_DOMAIN` that is within the zones specified in the `--domain-filter` argument of the [external-dns deployment manifest \(see page 1443\)](#) stored on the **Management cluster**.

For example, if the filter is set to `example.com`, a possible domain for the `TUNNEL_PROXY_EXTERNAL_DOMAIN` would be `myclusterproxy.example.com`.

3. Establish a variable that points to the name of the network-restricted cluster:

i The name of the network-restricted cluster is established in the [KommanderCluster \(see page 1430\)](#) object.

```
NETWORK_RESTRICTED_CLUSTER=<name_of_restricted_cluster>
```

4. Given that each cluster can only have one proxy domain, reuse the name of the network-restricted cluster for the proxy object:

```
TUNNEL_PROXY_NAME=${NETWORK_RESTRICTED_CLUSTER}
```

5. Obtain the name of the connector and set it to a variable:

```
TUNNEL_CONNECTOR_NAME=$(kubectl get kommandercluster -n ${WORKSPACE_NAMESPACE}
${NETWORK_RESTRICTED_CLUSTER} -o
template='{{ .spec.clusterTunnelConnectorRef.name }}')
```

Next Step:[3. Create a TunnelProxy Object: CLI Proxied Access \(see page 818\)](#)**3. Create a TunnelProxy Object: CLI Proxied Access**

On the **Management cluster**, create a `TunnelProxy` object for your proxied cluster and assign it a unique domain. This domain forwards all user authentication requests through the Management cluster, and is used to generate a URL that exposes the cluster's dashboards (`clusterProxyDomain`).

You require both a **certificate** and a **DNS record** to back the domain. If you choose the default configuration, DKP will handle the certificate creation (self-signed certificate), but you must create a DNS record manually.

Alternatively, you can set up a different Certificate Authority to handle the certificate creation and rotation for your domain. You can also set up the `external-dns` service to automatically create a DNS record.

Here are some examples of possible configuration combinations:



- Ensure you have set the required variables in [2. Prepare your Environment: CLI Proxied Access \(see page 819\)](#) before you run the following commands.
- Ensure you run the following commands on the Management cluster. See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for more information around switching cluster contexts.

Domain with default certificate and automatic DNS record creation (requires External DNS)

In this example, the following configuration applies:

- **Certificate** - The domain uses a self-signed certificate created by DKP.
- **DNS record** - The `external-dns` manages the creation of a DNS record automatically. For it to work, ensure you have enabled [External DNS in your Management cluster \(see page 1443\)](#).

```
cat > tunnelproxy.yaml <<EOF | kubectl apply -f -
apiVersion: kubetunnel.d2iq.io/v1alpha1
kind: TunnelProxy
metadata:
  name: ${TUNNEL_PROXY_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  clusterProxyDomain: ${TUNNEL_PROXY_EXTERNAL_DOMAIN}
  tunnelConnectorRef:
    name: ${TUNNEL_CONNECTOR_NAME}
  ingress:
    annotations:
      external-dns.alpha.kubernetes.io/hostname: ${TUNNEL_PROXY_EXTERNAL_DOMAIN}
```

EOF

- The `spec.ingress.annotations` field contains the annotation required for DNS record management. For more information, see [DNS Record Creation with External DNS](#) (see page 1442).

Domain with default certificate and default DNS setup (requires manually-created DNS record)

In this example, the following configuration applies:

- **Certificate** - The domain uses a self-signed certificate created by DKP.
- **DNS record** - For the domain to be recognized by the cluster, ensure you manually create a DNS record. The record's A/CNAME value must point to the Management cluster's Traefik IP address, URL or domain. Create a record per proxied cluster.

```
cat > tunnelproxy.yaml <<EOF | kubectl apply -f -
apiVersion: kubetunnel.d2iq.io/v1alpha1
kind: TunnelProxy
metadata:
  name: ${TUNNEL_PROXY_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  clusterProxyDomain: ${TUNNEL_PROXY_EXTERNAL_DOMAIN}
  tunnelConnectorRef:
    name: ${TUNNEL_CONNECTOR_NAME}
EOF
```

Domain with automatically-generated ACME certificate and automatic DNS record creation (requires External DNS)

In this example, the following configuration applies:

- **Certificate** - The domain uses `cert-manager` to enable an ACME-based Certificate Authority. This CA automatically issues and rotates your certificates. By default, DKP uses Let's Encrypt.
- **DNS record** - The `external-dns` manages the creation of a DNS record automatically. For it to work, ensure you have enabled [External DNS in your Management cluster](#) (see page 1443).

1. Set the environment variable for your issuing object:

i This can be a `ClusterIssuer` or `Issuer`. See [Advanced Configuration: ClusterIssuer](#) (see page 1437) for more information.

```
ISSUER_KIND=ClusterIssuer
```

2. Set the environment variable for your CA:

i Replace `letsEncrypt` if you are using another ACME-based certificate authority.

```
ISSUER_NAME=letsEncrypt
```

3. Create the TunnelProxy :

```
cat > tunnelproxy.yaml <<EOF | kubectl apply -f -
apiVersion: kubetunnel.d2iq.io/v1alpha1
kind: TunnelProxy
metadata:
  name: ${TUNNEL_PROXY_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  clusterProxyDomain: ${TUNNEL_PROXY_EXTERNAL_DOMAIN}
  tunnelConnectorRef:
    name: ${TUNNEL_CONNECTOR_NAME}
  ingress:
    annotations:
      external-dns.alpha.kubernetes.io/hostname: ${
TUNNEL_PROXY_EXTERNAL_DOMAIN}
    certificate:
      issuerRef:
        kind: ${ISSUER_KIND}
        name: ${ISSUER_NAME}
EOF
```



For more information, see [Configure the Kommander Installation with a Custom Domain and Certificate](#) (see page 1432) and [DNS Record Creation with External DNS](#) (see page 1442).

Domain with a custom certificate (requires certificate secret) and automatic DNS record creation (requires External DNS)

In this example, the following configuration applies:

- **Certificate** - The domain uses a custom certificate created manually. Ensure you reference the `<certificate_secret_name>`.
- **DNS record** - The `external-dns` manages the creation of a DNS record automatically. For it to work, ensure you have enabled [External DNS in your Management cluster](#) (see page 1443).

1. Set an environment variable for the name of your custom certificate:

See [Manually-generated certificate](#) (see page 1436) for more information.

```
CERTIFICATE_SECRET_NAME=<custom_certificate_secret_name>
```

2. Optional: If you do not have a secret yet and wish to create one pointing at the certificate, execute the following command:

```
kubectl create secret tls ${CERTIFICATE_SECRET_NAME} -n ${WORKSPACE_NAMESPACE}
--key="tls.key" --cert="tls.crt"
```

3. Create the TunnelProxy :

```
cat > tunnelproxy.yaml <<EOF | kubectl apply -f -
apiVersion: kubetunnel.d2iq.io/v1alpha1
kind: TunnelProxy
metadata:
  name: ${TUNNEL_PROXY_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  clusterProxyDomain: ${TUNNEL_PROXY_EXTERNAL_DOMAIN}
  tunnelConnectorRef:
    name: ${TUNNEL_CONNECTOR_NAME}
  ingress:
    annotations:
      external-dns.alpha.kubernetes.io/hostname: ${
{TUNNEL_PROXY_EXTERNAL_DOMAIN}
    certificate:
      certificateSecretRef:
        name: ${CERTIFICATE_SECRET_NAME}
EOF
```



For more information, see [manually-created certificates](#) (see page 863) and [DNS Record Creation with External DNS](#) (see page 1442).

Next Step:

4. Enable the TunnelProxy Object in the KommanderCluster: CLI Proxied Access (see page 821)

4. Enable the TunnelProxy Object in the KommanderCluster: CLI Proxied Access



- Ensure you have set the required variables in [2. Prepare your Environment: CLI Proxied Access](#) (see page 819) and created a TunnelProxy object in [3. Create a TunnelProxy Object: CLI Proxied Access](#) (see page 821) before you run the following commands.
- Ensure you run the following command on the Management cluster. See [Provide Context for Commands with a kubeconfig File](#) (see page 87) for more information around switching cluster contexts.

To enable the TunnelProxy, reference the object in the KommanderCluster object.

On the **Management cluster**, patch the `KommanderCluster` object with the name of the `TunnelProxy` you created in the previous page:

```
kubectl patch --type merge kommanderclusters -n ${WORKSPACE_NAMESPACE} $
{NETWORK_RESTRICTED_CLUSTER} --patch "{\"spec\": {\"clusterTunnelProxyConnectorRef\":
{ \"name\": \"${TUNNEL_PROXY_NAME}\"}}}"
```

Next Step:

[5. Verify the Proxy: CLI Proxied Access \(see page 822\)](#)

5. Verify the Proxy: CLI Proxied Access

On the **Management cluster**:



Ensure you run the following command on the Management cluster. See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for more information around switching cluster contexts.

1. Verify that the following conditions for the `TunnelProxy` configuration are met:

```
kubectl wait --for=condition=ClientAuthReady=true --timeout=300s -n $
{WORKSPACE_NAMESPACE} tunnelproxy/${TUNNEL_PROXY_NAME}
kubectl wait --for=condition=ReverseProxyReady=true --timeout=300s -n $
{WORKSPACE_NAMESPACE} tunnelproxy/${TUNNEL_PROXY_NAME}
kubectl wait --for=condition=available -n ${WORKSPACE_NAMESPACE} deploy -l
control-plane=${TUNNEL_PROXY_NAME}-kubetunnel-reverse-proxy-rp
```

The output should look like this:

```
tunnelproxy.kubetunnel.d2iq.io/test condition met
tunnelproxy.kubetunnel.d2iq.io/test condition met
deployment.apps/${TUNNEL_PROXY_NAME}-kubetunnel-reverse-proxy-rp condition met
```

2. Verify that the `TunnelProxy` is correctly assigned and connected to your cluster:

```
curl -Lk -s -o /dev/null -w "%{http_code}" https://$
{TUNNEL_PROXY_EXTERNAL_DOMAIN}/dkp/grafana
```

The output should return a successful HTTP response status:

```
200
```

Now you can access the network-restricted cluster's dashboards and [use kubectl to manage its resources from the Management cluster](#) (see page 0).

6.5.3.5.8.4 Enable Proxied Access Using the UI

See [UI: Attach a Network-Restricted Cluster](#) (see page 780) for instructions on how to enable a proxied access while you attach a cluster with the UI.

If you attached your Kubernetes cluster to your DKP environment already, it is not possible to enable the proxied access using the UI. Use the CLI as explained in [Enable Proxied Access Using the CLI](#) (see page 819).

Configure a Wildcard Domain for your Proxied Cluster

Wildcard domains are helpful in multi-cluster environments. When you set up a wildcard domain, every time you attach an additional [network-restricted cluster](#) (see page 84) through a proxy, the DKP UI pre-fills a domain for the cluster automatically.

After you set up a wildcard domain in the `kommander.yaml` as explained in the following section, DKP will suggest domains for attached clusters automatically based on the wildcard domain + the name of the cluster, for example:

Wildcard Domain	Cluster Name	Cluster Domain
*.example.com	development.cluster	development.cluster.example.com
*.example.com	janedoe	janedoe.example.com

You can override the domain by overwriting the suggested domain in the UI.

Set up a Wildcard Domain

1. Open the `kommander.yaml` file:
 - a. If you have not installed the Kommander component yet, [initialize the configuration file](#) (see page 1422), so you can edit it in the following steps.

⚠ WARNING: Initialize this file only ONCE, otherwise you will overwrite previous customizations.
 - b. If you have installed the Kommander component already, open the existing `kommander.yaml` with the editor of your choice.
2. Adjust the `apps` section of your `kommander.yaml` file to include these values:

```

apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
[...]
  kubetunnel:
    enabled: true
    values: |
      proxiedAccess:
        clusterDomainWildcard: "*.example.com"

```

3. Use the configuration file to install or update the Kommander component:

```

dkp install kommander --installer-config kommander.yaml --kubecfg=${CLUSTER_NAME}.conf

```

Whenever you attach a network-restricted cluster, the UI will suggest a new domain based on the wildcard domain and cluster name.

Configure a DNS Record or DNS Service

The clusters will not be available through the established domain until you manually create a DNS record or enable `external-dns` to automatically manage the creation of records.

D2iQ recommends enabling the [External DNS \(see page 1442\)](#) service to manage your records automatically. However, you can choose to create your records manually.

6.5.3.6 Attach a DKP-created Kubernetes Cluster

This section contains the following topics:

- [Make a DKP-CLI-created Cluster Managed \(see page 824\)](#)

Starting with DKP 2.6, when you create a [Managed Cluster \(see page 82\)](#) with the DKP CLI, it attaches automatically to the [Management Cluster \(see page 82\)](#) after a few moments.


If you have several [Essential Clusters \(see page 0\)](#) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster \(see page 827\)](#).

6.5.3.6.1 Make a DKP-CLI-created Cluster Managed

Enterprise

Gov Advanced

6.5.3.6.1.1 Manually Attach a DKP CLI Cluster to the Management Cluster

 These steps are only applicable if you do not set a `WORKSPACE_NAMESPACE` when creating a cluster. If you already set a `WORKSPACE_NAMESPACE`, then you do not need to perform these steps since the cluster is already attached to the workspace.

Starting with DKP 2.6, when you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) after a few moments.

However, if you do not set a workspace, the attached cluster will be created in the `default` workspace. To ensure that the attached cluster is created in your desired workspace namespace, follow these instructions:

1. Confirm you have your `MANAGED_CLUSTER_NAME` variable set with the following command:

```
echo ${MANAGED_CLUSTER_NAME}
```

2. Retrieve your kubeconfig from the cluster you have created without setting a workspace:

```
dkp get kubeconfig --cluster-name ${MANAGED_CLUSTER_NAME} > ${MANAGED_CLUSTER_NAME}.conf
```

3. You can now either [attach it in the UI](link to attaching it to workspace via UI that was earlier), or attach your cluster to the workspace you want in the CLI.

NOTE: This is only necessary if you never set the workspace of your cluster upon creation.

4. Retrieve the workspace where you want to attach the cluster:

```
kubectl get workspaces -A
```

5. Set the `WORKSPACE_NAMESPACE` environment variable:

```
export WORKSPACE_NAMESPACE=<workspace-namespace>
```

6. You need to create a secret in the desired workspace before attaching the cluster to that workspace. Retrieve the kubeconfig secret value of your cluster:

```
kubectl -n default get secret ${MANAGED_CLUSTER_NAME}-kubeconfig -o go-template='{{.data.value}}{{ "\n"}}
```

7. This will return a lengthy value. Copy this entire string for a secret using the template below as a reference. Create a new `attached-cluster-kubeconfig.yaml` file:

```

apiVersion: v1
kind: Secret
metadata:
  name: <your-managed-cluster-name>-kubeconfig
  labels:
    cluster.x-k8s.io/cluster-name: <your-managed-cluster-name>
type: cluster.x-k8s.io/secret
data:
  value: <value-you-copied-from-secret-above>

```

8. Create this secret in the desired workspace:

```

kubectl apply -f attached-cluster-kubeconfig.yaml --namespace $
{WORKSPACE_NAMESPACE}

```

9. Create this `kommandercluster` object to attach the cluster to the workspace:

```

cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
  name: ${MANAGED_CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  kubeconfigRef:
    name: ${MANAGED_CLUSTER_NAME}-kubeconfig
  clusterRef:
    capiCluster:
      name: ${MANAGED_CLUSTER_NAME}
EOF

```

10. You can now view this cluster in your Workspace in the UI and you can confirm its status by running the below command. It may take a few minutes to reach "Joined" status:

```

kubectl get kommanderclusters -A

```

If you have several [Essential Clusters](#) (see page 0) and want to turn one of them to a Managed Cluster to be centrally administrated by a Management Cluster, refer to [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 827).

6.5.3.7 Access a Managed or Attached Cluster

Enterprise

Gov Advanced

How to access a Managed or Attached cluster with credentials

6.5.3.7.1 Access your Clusters Using your UI Administrator Credentials

After the cluster is attached, retrieve a custom `kubeconfig` file from the UI.

1. Select the username in the top right corner, and then select **Generate Token**.
2. Select the cluster name, and follow the instructions to assemble a `kubeconfig` for accessing its Kubernetes API.



If the UI prompts you to log in, use the credentials you normally use to access the UI.

You can also retrieve a custom `kubeconfig` file by visiting the `/token` endpoint on the Kommander cluster domain (example URL: `https://your-server-name.your-region-2.elb.service.com/token/`). Selecting the cluster's name displays the instructions to assemble a `kubeconfig` for accessing its Kubernetes API.

6.5.3.8 Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster

Enterprise

Gov Advanced

If you are a DKP Essential customer, you can easily convert your independent clusters into a multi-cluster environment if you upgrade to an Enterprise license.

The pages in this section provides information on how you can turn your Essential Clusters into Managed clusters.

For all the instructions regarding converting an existing Essential cluster into an Enterprise cluster, refer to that entire [section of documentation](#) (see page 827).

6.5.4 Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster

If you are a DKP Essential customer, you can easily convert your independent clusters into a multi-cluster environment if you upgrade to an Enterprise license.

The pages in this section provides information on how you can turn your Essential Clusters into Managed clusters.

- [Prerequisites: General Prerequisites for your Cluster Conversion](#) (see page 828)
- [Prerequisites: Prepare your Cluster's Existing Configurations](#) (see page 830)

- [Prerequisites: Clone Git Repository from Gitea](#) (see page 831)
- [Back up your Cluster's Applications and Persistent Volumes](#) (see page 832)
- [UI: Convert an Essential Cluster into a Managed Cluster](#) (see page 842)
- [CLI: Convert an Essential Cluster into a Managed Cluster](#) (see page 844)
- [Post Conversion Cleanup: Cluster Autoscaler Configuration](#) (see page 845)
- [Post Conversion Cleanup: Clusters run on Different Cloud Platforms](#) (see page 848)
- [Troubleshooting](#) (see page 851)

6.5.4.1 Prerequisites: General Prerequisites for your Cluster Conversion

Information that you need to know prior to converting your DKP Essential Clusters into DKP Enterprise Managed Clusters.

Starting with DKP 2.5, you have the option to convert your DKP Essential Clusters into DKP Managed Enterprise Clusters.

6.5.4.1.1 Requirements

To convert your DKP Essential clusters into DKP Managed Enterprise clusters, ensure the following is valid:

- A DKP Management cluster with a valid DKP Enterprise License installed
- At least one installed and running standalone DKP Essential cluster
- All DKP Essential Clusters are [upgraded](#) (see page 1551) to the same DKP version as the DKP Managed Enterprise cluster
- The DKP Essential Cluster you want to convert is [self-managed](#) (see page 83)
- The DKP Essential Cluster you want to convert only contains its own Cluster API resources and does not contain Cluster API resources from other clusters



See the [Licenses](#) (see page 125) page for information on how you can purchase a DKP Enterprise license.



Attaching DKP Enterprise clusters is not supported.

6.5.4.1.2 Downtime Considerations

Your DKP Essential cluster will not be accessible externally for several minutes during the expansion process. Any configuration of the cluster's `Ingress` that requires `traefik-forward-auth` authentication will be affected.

- Access from within the cluster via Kubernetes service hostname (for example, `http://SERVICE_NAME.NAMESPACE:PORT`) is not affected.

6.5.4.1.2.1 Affected DKP Services

- `dkp-ceph-cluster-dashboard`
- `grafana-logging`
- `kube-prometheus-stack-alertmanager`
- `kube-prometheus-stack-grafana`
- `kube-prometheus-stack-prometheus`
- `kubernetes-dashboard`
- `traefik-dashboard`

6.5.4.1.2.2 Other Services

To verify if your services are affected by `traefik-forward-auth`'s downtime, run the following command:

```
kubectl get ingress -n NS <your_customized_ingress_name>
```

Look for the `traefik.ingress.kubernetes.io/router.middlewares` field in the output. If this field contains the value `kommander-forwardauth@kubernetescrd`, your service will be affected by the downtime.

6.5.4.1.2.3 Duration

The `traefik-forward-auth` service is affected starting with the `PreAttachmentCleanup` conversion stage, and will run normally again after `ResumeFluxOperations` is completed. [Observe the conversion progress \(see page 852\)](#) to monitor your cluster's current status.

6.5.4.1.3 Special Considerations for DKP Insights

If you are a participant in the Technical Preview of DKP Insights, you must uninstall Insights prior to converting your DKP Essential clusters into Managed Clusters.

Following conversion, you can re-install Insights.

See the following page for information about installing and uninstalling Insights in your environment: (2.6) [DKP Insights Setup and Configuration](#)⁵⁹⁰

6.5.4.1.4 Next Step

<https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273286363/2.6%2BPrerequisites%2BPrepare%2Byour%2BCluster%2Bs%2BExisting%2BConfigurations> (see page 830)

6.5.4.1.5 See Also

[Troubleshooting](#) (see page 851)

6.5.4.2 Prerequisites: Prepare your Cluster's Existing Configurations

Enterprise

Gov Advanced

Before you convert an *Essential* cluster into a *Managed* cluster, review the following information.

6.5.4.2.1 SSO Configuration


After attachment, the SSO configuration of the Management cluster applies to the Managed (formerly Essential) cluster. Any SSO configuration of the former Essential cluster will be deleted.

- If the **Essential cluster has SSO configured**, but the Management cluster has NO configuration, you can copy your Essential cluster's SSO configuration (`dex-controller` resources) to the Management cluster before conversion.
- If your **Management cluster has SSO configured** and the **Essential cluster has another SSO configuration**, you can choose to keep one, or both. To keep the configuration of your Essential cluster, manually copy the `dex-controller` resources to the Management cluster before conversion. DKP maintains the SSO configuration of your Management cluster automatically, unless you manually delete it.

6.5.4.2.2 Domains and Certificates

Any custom domain or certificate configuration you have set up for your Essential cluster remains functional after you turn it into a Managed cluster.

⁵⁹⁰ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274759846/%282.6%29+DKP+Insights+Setup+and+Configuration>

 After conversion, any domain or certificate customizations you would like to apply to your Managed cluster must be done via the `KommanderCluster`. This object is now stored in the Management cluster.


6.5.4.2.3 Next Step

Prerequisites: [Delete Gitea \(see page 831\)](#)

6.5.4.3 Prerequisites: Clone Git Repository from Gitea

After your DKP Essential Cluster is converted into a DKP Enterprise Managed cluster, the old instance of Gitea that is used to host all Git repositories in the DKP Essential Cluster, will not be preserved.

Perform the steps in this page to ensure that you have a local copy of the Management Git Repository in the state it was in prior to undergoing the expansion process.

 All DKP Platform Applications will be migrated from the DKP Essential Cluster to the DKP Enterprise Managed Cluster.

1. Prior to turning an DKP Essential cluster to a DKP Enterprise Managed Cluster, clone Gitea using the following command:

```
dkp experimental gitea clone
```

2. Verify that the Git Repository has been successfully cloned to your local environment.

```
cd kommander
git remote -v
# output from git remote -v look like
# origin https://<YOUR_CLUSTER_INGRESS_HOSTNAME>/dkp/kommander/git/kommander/
kommander (fetch)
# origin https://<YOUR_CLUSTER_INGRESS_HOSTNAME>/dkp/kommander/git/kommander/
kommander (push)
```

6.5.4.3.1 Next Steps

[Back up your Cluster's Applications and Persistent Volumes \(see page 832\)](#)

6.5.4.4 Back up your Cluster's Applications and Persistent Volumes

Enterprise

Gov Advanced

Back up and restore your cluster's applications before attempting converting your DKP Essential cluster into a DKP Enterprise Managed cluster in the scope of [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 830).

The instructions differ depending on the infrastructure provider of your **DKP Essential cluster**.

For AWS, refer to: [Back up an AWS Cluster](#) (see page 832)

For Azure, vSphere, GCP and pre-provisioned environments, refer to: [Back up an Azure, vSphere, GCP or Pre-provisioned Cluster](#) (see page 837)

6.5.4.4.1 Back up an AWS Cluster

This section contains the instructions necessary to back up and restore a DKP Essential cluster on the AWS environment in the scope of [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 830).

Here is a summary of the steps:

- [Prepare your Cluster for Backup - AWS](#) (see page 832)
- [Back up a Cluster - AWS Environment](#) (see page 835)

6.5.4.4.1.1 Prepare your Cluster for Backup - AWS

This section describes how to prepare your cluster on an AWS environment, so it can be backed up before you begin with [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster](#) (see page 830).

Prerequisites

- Ensure Velero is installed on your Essential cluster
- [Install the Velero CLI](#) (see page 894) (Use at least Velero CLI version 1.10.1)
- Ensure [kubectl](#)⁵⁹¹ is installed
- Ensure you have admin rights to the DKP Essential cluster

Prepare your Cluster

⁵⁹¹ <https://kubernetes.io/docs/tasks/tools/>



Run the following commands in the **DKP Essential cluster**. For general guidelines on how to set the context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

Prepare Velero

Enable the CSI snapshotting plug-in by providing a custom configuration of Velero.

1. Create an Override with the custom configuration:

```
cat << EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: velero-overrides
  namespace: kommander
data:
  values.yaml: |
    ---
    configuration:
      features: EnableCSI
    initContainers:
      - name: velero-plugin-for-aws
        image: velero/velero-plugin-for-aws:v1.5.2
        imagePullPolicy: IfNotPresent
        volumeMounts:
          - mountPath: /target
            name: plugins
      - name: velero-plugin-for-csi
        image: velero/velero-plugin-for-csi:v0.4.2
        imagePullPolicy: IfNotPresent
        volumeMounts:
          - mountPath: /target
            name: plugins
EOF
```

2. Update the `AppDeployment` to apply the new configuration:

```
cat << EOF | kubectl -n kommander patch appdeployment velero --type='merge' --
patch-file=/dev/stdin
spec:
  configOverrides:
    name: velero-overrides
EOF
```

3. Verify the configuration has been updated before proceeding with the next section:

```
kubectl -n kommander wait --for=condition=Ready kustomization velero
```

The output should look similar to this:

```
kustomization.kustomize.toolkit.fluxcd.io/velero condition met
```

Prepare the AWS IAM Permission

When creating a cluster on AWS, you provided an additional permission as specified in [Cluster IAM Policies, Roles, and Artifacts](#) (see page 1063).

For the CSI plugin to function correctly, you must update the existing IAM role to include an additional policy.

Add the [AmazonEBSCSIDriverPolicy](#)⁵⁹² policy to the control plane role `control-plane.cluster-api-provider-aws.sigs.k8s.io`:

```
aws iam attach-role-policy \
  --role-name control-plane.cluster-api-provider-aws.sigs.k8s.io \
  --policy-arn arn:aws:iam::aws:policy/service-role/AmazonEBSCSIDriverPolicy
```

This will allow the EBS CSI driver, a volume manager, to have enough permissions to create volume snapshots.



The default control plane role name is `control-plane.cluster-api-provider-aws.sigs.k8s.io`. If you customized this name when creating the AWS cluster, replace the default control plane role with the name you assigned to it.

Prepare the CSI Configuration

Configure a `VolumeSnapshotClass` object on the cluster, so Velero can create a volume snapshot:

```
cat << EOF | kubectl apply -f -
apiVersion: snapshot.storage.k8s.io/v1
kind: VolumeSnapshotClass
metadata:
  name: aws
  labels:
    velero.io/csi-volumesnapshot-class: "true"
driver: ebs.csi.aws.com
```

⁵⁹² <https://docs.aws.amazon.com/aws-managed-policy/latest/reference/AmazonEBSCSIDriverPolicy.html>

```
deletionPolicy: Delete
parameters:
EOF
```

Next Step:

[Back up a Cluster - AWS Environment](#) (see page 835)

6.5.4.4.1.2 Back up a Cluster - AWS Environment

With this workflow, you can back up and restore your cluster's applications. The backup contains all Kubernetes objects and the Persistent Volumes of your DKP applications.

When backing up a cluster that runs on a cloud provider, Velero captures the state of your cluster in a snapshot. Review Velero's list of cloud providers for [CSI⁵⁹³](#) compatibility.



Run the following commands in the **DKP Essential cluster**. For general guidelines on how to set the context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

Back up Instructions

1. Configure Velero to use CSI snapshotting:

```
velero client config set features=EnableCSI
```

2. Create a backup with Velero. Use the following flags to reduce the scope of the backup and only include the applications that are affected during the expansion:

```
velero backup create pre-expansion \
  --include-namespaces="kommander,kommander-default-workspace,kommander-
  flux,kubecost" \
  --include-cluster-resources \
  --wait
```

After completion, the output should look similar to this:

```
Backup request "pre-expansion" submitted successfully.
Waiting for backup to complete. You may safely press ctrl-c to stop waiting -
your backup will continue in the background.
```

⁵⁹³ <https://velero.io/docs/v1.10/supported-providers/>

```

.....
.....
.....
.....
Backup completed with status: Completed. You may check for more information
using the commands `velero backup describe pre-expansion` and `velero backup
logs pre-expansion`.
    
```

Verify the Backup

Review the backup has completed successfully:

```
velero backup describe pre-expansion
```

The following example output will vary depending on your cloud provider. Verify that it shows no errors and the Phase is Completed :

```

Name:          pre-expansion
Namespace:     kommander
Labels:        velero.io/storage-location=default
Annotations:   velero.io/source-cluster-k8s-gitversion=v1.25.5
               velero.io/source-cluster-k8s-major-version=1
               velero.io/source-cluster-k8s-minor-version=25

Phase:  Completed

Errors:  0
Warnings: 0

Namespaces:
  Included:  kommander, kommander-default-workspace, kommander-flux, kubecost
  Excluded:  <none>

Resources:
  Included:  *
  Excluded:  <none>
  Cluster-scoped: included

Label selector: <none>

Storage Location: default

Velero-Native Snapshot PVs: auto

TTL: 720h0m0s

CSISnapshotTimeout: 10m0s
    
```

```

Hooks: <none>

Backup Format Version: 1.1.0

Started:    2023-03-15 10:40:25 -0400 EDT
Completed:  2023-03-15 10:44:39 -0400 EDT

Expiration: 2023-04-14 10:40:24 -0400 EDT

Total items to be backed up: 5188
Items backed up:             5188

Velero-Native Snapshots: <none included>

```

Next Step

Expand your platform. Use the UI or CLI:

[UI: Convert an Essential Cluster into a Managed Cluster \(see page 842\)](#)

[CLI: Convert an Essential Cluster into a Managed Cluster \(see page 844\)](#)

Related Topics

[Restore your Backup and Retry the Cluster Expansion \(see page 855\)](#)

6.5.4.4.2 Back up an Azure, vSphere, GCP or Pre-provisioned Cluster

This section contains the instructions necessary to back up and restore a DKP Essential cluster on Azure, vSphere, GCP or Pre-provisioned environments in the scope of [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster \(see page 830\)](#).

Here is a summary of the steps:

- [Prepare your Cluster for Backup - Azure, vSphere, GCP, Pre-provisioned Environments \(see page 837\)](#)
- [Back up a Cluster - Azure, vSphere, GCP, and Pre-provisioned Environments \(see page 839\)](#)

6.5.4.4.2.1 Prepare your Cluster for Backup - Azure, vSphere, GCP, Pre-provisioned Environments

This section describes how to prepare your cluster on AWS, Azure, vSphere, Google Cloud or pre-provisioned environment, so it can be backed up before you begin with [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster \(see page 830\)](#).

Prerequisites

- Ensure Velero is installed on your Essential cluster
- [Install the Velero CLI \(see page 894\)](#) (Use at least Velero CLI version 1.10.1)
- Ensure [kubectl](#)⁵⁹⁴ is installed

⁵⁹⁴ <https://kubernetes.io/docs/tasks/tools/>

- Ensure you have admin rights to the DKP Essential cluster

Prepare your Cluster



Run the following commands in the **DKP Essential cluster**. For general guidelines on how to set the context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

Prepare Velero

Enable `restic` backup capabilities by providing a custom configuration of Velero.

1. Create an Override with a custom configuration for Velero. This custom configuration deploys the `node-agent` service, which enables `restic`:

```
cat << EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: velero-overrides
  namespace: kommander
data:
  values.yaml: |
    ---
    deployNodeAgent: true
EOF
```

2. Reference the created Override in Velero's `AppDeployment` to apply the new configuration:

```
cat << EOF | kubectl -n kommander patch appdeployment velero --type='merge' --
patch-file=/dev/stdin
spec:
  configOverrides:
    name: velero-overrides
EOF
```

3. Wait until the `node-agent` has been deployed:

```
until kubectl get daemonset -A | grep -m 1 "node-agent"; do 0.1 ; done
```

The `node-agent` is ready after a similar output appears:

```
kommander          node-agent
3                 0         3         0         <none>
```

4. Verify the configuration has been updated before proceeding with the next section:

```
kubectl -n kommander wait --for=condition=Ready kustomization velero
```

The output should look similar to this:

```
kustomization.kustomize.toolkit.fluxcd.io/velero condition met
```

Prepare your Pods for Backup



Run the following commands in the **DKP Essential cluster**. For general guidelines on how to set the context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

Annotate the `gitea` pods to ensure `restic` backs up the Persistent Volumes (PVs) of the pods that will be affected during the expansion process. These volumes contain the Git repository information of your DKP Essential cluster.

```
kubectl -n kommander annotate pod/gitea-0 backup.velero.io/backup-volumes=data
```

```
kubectl -n kommander annotate pod/gitea-postgresql-0 backup.velero.io/backup-volumes=data
```

Next Step:

[Back up a Cluster - Azure, vSphere, GCP, and Pre-provisioned Environments](#) (see page 839)

6.5.4.4.2 Back up a Cluster - Azure, vSphere, GCP, and Pre-provisioned Environments

With this workflow, you can back up and restore your cluster's applications. This backup contains Kubernetes objects and the Persistent Volumes (PVs) of Gitea pods. Given that Gitea's PVs store information on your cluster's state, you will be able to [Restore your Cluster](#) (see page 855), if required.



Run the following commands in the **DKP Essential cluster**. For general guidelines on how to set the context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

Back up Instructions

1. Create a backup with Velero. Use the following flags to reduce the scope of the backup and only include the applications that are affected during the expansion:

```
velero backup create pre-expansion \
  --include-namespaces="kommander,kommander-default-workspace,kommander-
flux,kubecost" \
  --include-cluster-resources \
  --snapshot-volumes=false --wait \
  --namespace kommander
```

After completion, the output should look similar to this:

```
Backup request "pre-expansion" submitted successfully.
Waiting for backup to complete. You may safely press ctrl-c to stop waiting -
your backup will continue in the background.
.....
.....
.....
.....
Backup completed with status: Completed. You may check for more information
using the commands `velero backup describe pre-expansion` and `velero backup
logs pre-expansion`.
```

Verify the Backup

1. Ensure the backup has completed successfully:

```
velero backup describe pre-expansion --namespace kommander
```

The following example output will vary depending on your cloud provider. Verify that it shows no errors and the Phase is Completed :

```
Name:          pre-expansion
Namespace:     kommander
Labels:        velero.io/storage-location=default
Annotations:   velero.io/source-cluster-k8s-gitversion=v1.25.5
```

```

velero.io/source-cluster-k8s-major-version=1
velero.io/source-cluster-k8s-minor-version=25

```

Phase: Completed

Errors: 0

Warnings: 0

Namespaces:

Included: kommander, kommander-**default**-workspace, kommander-flux, kubecost

Excluded: <none>

Resources:

Included: *

Excluded: <none>

Cluster-scoped: included

Label selector: <none>

Storage Location: **default**

Velero-Native Snapshot PVs: auto

TTL: 720h0m0s

CSISnapshotTimeout: 10m0s

Hooks: <none>

Backup Format Version: 1.1.0

Started: 2023-03-15 10:40:25 -0400 EDT

Completed: 2023-03-15 10:44:39 -0400 EDT

Expiration: 2023-04-14 10:40:24 -0400 EDT

Total items to be backed up: 5188

Items backed up: 5188

Velero-Native Snapshots: <none included>

2. Ensure that the `PodVolumeBackup` objects have been created:

```
kubectl get podvolumebackups -A
```

The output should look similar to this:

NAMESPACE	NAME	STATUS	CREATED	NAMESPACE	POD
VOLUME	REPOSITORY ID				
UPLOADER TYPE	STORAGE LOCATION	AGE			
kommander	ash-59ntg	Completed	39s	kommander	gitea-postgresql-0
data	s3:https://a54904d80411e4d64b572b96cb3ddb62-477717230.us-west-2.elb.amazonaws.com:8085/dkp-velero/restic/kommander	restic			
default		39s			
kommander	ash-5vsbf	Completed	42s	kommander	gitea-0
data	s3:https://a54904d80411e4d64b572b96cb3ddb62-477717230.us-west-2.elb.amazonaws.com:8085/dkp-velero/restic/kommander	restic			
default		42s			

Next Step

Expand your platform. Use the UI or CLI:

UI: [Convert an Essential Cluster into a Managed Cluster \(see page 842\)](#)

CLI: [Convert an Essential Cluster into a Managed Cluster \(see page 844\)](#)

Related Topics

[Restore your Backup and Retry the Cluster Expansion \(see page 855\)](#)

6.5.4.5 UI: Convert an Essential Cluster into a Managed Cluster

Enterprise

Gov Advanced

Instructions on how you can convert an Essential Cluster to a Managed cluster with no Networking Restrictions with the DKP UI.

Ingress that contains [Traefik-Forward-Authentication \(see page 946\)](#) configuration will be not available during the expansion process, therefore, your DKP Essential cluster will not be accessible externally for several minutes. Access from within the cluster via Kubernetes service hostname (for example, `http://SERVICE_NAME.NAMESPACE:PORT`) is not affected. See [Downtime Considerations \(see page 831\)](#) for more information.

How to attach an existing cluster that has no additional networking restrictions

Use this option when you want to attach a cluster that does not require additional access information.

1. From the top menu bar, select your target workspace.

2. On the Dashboard page, select the **Add Cluster** option in the **Actions** dropdown menu at the top right.
3. Select **Attach Cluster**.
4. Select the **No additional networking restrictions** card. Alternatively, if you **MUST** use network restrictions, stop following the steps below, and see the instructions on the page [Attach a Cluster with Network Restrictions \(see page 775\)](#).
5. In the **Cluster Configuration** section, paste your `kubeconfig` file into the field, or select the **Upload kubeconfig File** button to specify the file.
6. The **Cluster Name** field will automatically populate with the name of the cluster is in the `kubeconfig`. You can edit this field with the name you want for your cluster.
7. The **Context** select list is populated from the `kubeconfig`. Select the desired context with admin privileges from the **Context** select list.
8. Add labels to classify your cluster as needed.
9. Select **Create** to attach your cluster.

6.5.4.5.1 Verify the Conversion

 Run the following commands in the **Management cluster**. For general guidelines on how to set the context, refer to [Provide Context for Commands with a kubeconfig File \(see page 87\)](#).

1. Export the environment variable for the workspace namespace:


```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

2. To verify that the conversion is successful, check the `KommanderCluster` object:

```
kubectl wait --for=condition=AttachmentCompleted kommandercluster <cluster name> -n ${WORKSPACE_NAMESPACE} --timeout 30m
```

The following output appears if the conversion is successful:

```
kommandercluster.kommander.mesosphere.io/<cluster name> condition met
```

 After conversion, all Platform Applications will be in the Kommander Namespace in the Managed Cluster.

6.5.4.5.2 Next Steps

[Post Conversion Cleanup: Cluster Autoscaler Configuration](#) (see page 845)

6.5.4.5.3 See Also

[Troubleshooting](#) (see page 851)

6.5.4.6 CLI: Convert an Essential Cluster into a Managed Cluster

Enterprise

Gov Advanced

Instructions on how you can convert an DKP Essential Cluster to a DKP Enterprise Managed Cluster with no Networking Restrictions via the CLI.



Ingress that contains [Traefik-Forward-Authentication](#) (see page 946) configuration will be not available during the expansion process, therefore, your DKP Essential cluster will not be accessible externally for several minutes. Access from within the cluster via Kubernetes service hostname (for example, `http://SERVICE_NAME.NAMESPACE:PORT`) is not affected. See [Downtime Considerations](#) (see page 831) for more information.

1. Run the following command in the DKP Enterprise cluster you will have managing your clusters. The cluster name and kubeconfig is from the cluster you are attaching and want to convert to initiate the process of turning an DKP Essential cluster to a DKP Enterprise Managed cluster. The workspace is the workspace name you want the attached cluster to go into.

```
dkp attach cluster --name <essential-cluster-name> --attached-kubeconfig
<kubeconfig-file-of-essential-cluster> --workspace <workspace-name>
```

6.5.4.6.1 Verify the Conversion



Run the following commands in the **Management cluster**. For general guidelines on how to set the context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

1. Export the environment variable for the workspace namespace:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

- To verify that the conversion is successful, check the `KommanderCluster` object:

```
kubect! wait --for=condition=AttachmentCompleted kommandercluster <cluster name> -n ${WORKSPACE_NAMESPACE} --timeout 30m
```

The following output appears if the conversion is successful:

```
kommandercluster.kommander.mesosphere.io/<cluster name> condition met
```



After conversion, all Platform Applications will be in the `kommander` namespace in the Managed Cluster.

6.5.4.6.2 Next Steps

[Post Conversion Cleanup: Cluster Autoscaler Configuration](#) (see page 845)

6.5.4.6.3 See Also

[Troubleshooting](#) (see page 851)

6.5.4.7 Post Conversion Cleanup: Cluster Autoscaler Configuration

Follow the steps in this page to ensure that the Cluster Autoscaler is able to function properly after turning your DKP Essential cluster to a DKP Enterprise Managed cluster.

After converting your cluster from Essential to Enterprise, the Cluster Lifecycle Management responsibilities is moved to a single Management cluster.

The Cluster Autoscaler feature also depends on the same Cluster Lifecycle Management components. If you are using the Cluster Autoscaler feature in DKP, you must perform the following steps for this feature to continue to work correctly:



Run the following commands in the **Management cluster**. For general guidelines on how to set the context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

- Set the following environment variables with your cluster's details:


```
export CLUSTER_NAME=<>
export WORKSPACE_NAMESPACE=<>
```

2. Apply the Cluster Autoscaler Deployment and supporting resources:

```
cat <<EOF | kubectl apply -f -
---
apiVersion: apps/v1
kind: Deployment
metadata:
  labels:
    app: cluster-autoscaler-`${CLUSTER_NAME}`
    name: cluster-autoscaler-`${CLUSTER_NAME}`
    namespace: `${WORKSPACE_NAMESPACE}`
spec:
  replicas: 1
  selector:
    matchLabels:
      app: cluster-autoscaler-`${CLUSTER_NAME}`
  template:
    metadata:
      labels:
        app: cluster-autoscaler-`${CLUSTER_NAME}`
    spec:
      containers:
        - args:
            - --cloud-provider=clusterapi
            - --node-group-auto-discovery=clusterapi:clusterName=${CLUSTER_NAME}
            - --kubeconfig=/workload-cluster/kubeconfig
            - --clusterapi-cloud-config-authoritative
            - -v5
          command:
            - /cluster-autoscaler
          image: us.gcr.io/k8s-artifacts-prod/autoscaling/cluster-
autoscaler:v1.25.0
          name: cluster-autoscaler
          volumeMounts:
            - mountPath: /workload-cluster
              name: kubeconfig
              readOnly: true
        serviceAccountName: cluster-autoscaler-`${CLUSTER_NAME}`
        terminationGracePeriodSeconds: 10
      tolerations:
        - effect: NoSchedule
          key: node-role.kubernetes.io/master
        - effect: NoSchedule
          key: node-role.kubernetes.io/control-plane
      volumes:
        - name: kubeconfig
```

```

    secret:
      items:
        - key: value
          path: kubeconfig
          secretName: ${CLUSTER_NAME}-kubeconfig
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: cluster-autoscaler-management-${CLUSTER_NAME}
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-autoscaler-management-${CLUSTER_NAME}
subjects:
- kind: ServiceAccount
  name: cluster-autoscaler-${CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
---
apiVersion: v1
kind: ServiceAccount
metadata:
  name: cluster-autoscaler-${CLUSTER_NAME}
  namespace: ${WORKSPACE_NAMESPACE}
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: cluster-autoscaler-management-${CLUSTER_NAME}
rules:
- apiGroups:
  - cluster.x-k8s.io
  resources:
  - machinedeployments
  - machinedeployments/scale
  - machines
  - machinesets
  verbs:
  - get
  - list
  - update
  - watch
EOF

```

3. Verify the output is similar to the following:

```

deployment.apps/cluster-autoscaler-<cluster-name> created
clusterrolebinding.rbac.authorization.k8s.io/cluster-autoscaler-management-
<cluster-name> created
serviceaccount/cluster-autoscaler-<cluster-name> created

```

```
clusterrole.rbac.authorization.k8s.io/cluster-autoscaler-management-<cluster-name> created
```

- To check that the status of the deployment has the expected `AVAILABLE` count of `1`, run the following command and verify that the output is similar:

```
$ kubectl get deployment -n $WORKSPACE_NAMESPACE cluster-autoscaler-
$CLUSTER_NAME
NAME                                READY    UP-TO-DATE    AVAILABLE    AGE
cluster-autoscaler-<cluster-name>  1/1      1              1            1m
```

6.5.4.7.1 Next Step

[Post Conversion Cleanup: Clusters run on Different Cloud Platforms \(see page 848\)](#)

6.5.4.7.2 See Also

[Troubleshooting \(see page 851\)](#)

6.5.4.8 Post Conversion Cleanup: Clusters run on Different Cloud Platforms

For a DKP Enterprise Management cluster to manage a cluster hosted in another cloud provider, you must ensure the Management cluster has all the necessary permissions.

6.5.4.8.1 Prerequisites

Prior to running these commands, you must ensure that the DKP Management Enterprise cluster is configured with the necessary platform specific permissions to manage the incoming CAPI objects that backs the infrastructure resources in the target cloud platform.

For example, for the DKP Enterprise Managed cluster to manage CAPI clusters in AWS, refer to <https://cluster-api-aws.sigs.k8s.io/topics/iam-permissions.html>.

DKP supports expanding your platform in the following scenarios:

<i>DKP Enterprise Management cluster host provider</i>	<i>DKP Enterprise Management cluster IAM permissions</i>	<i>DKP Essential cluster host provider</i>
AWS	https://cluster-api-aws.sigs.k8s.io/topics/iam-permissions.html	AWS, GCP, vSphere, Pre-provisioned

DKP Enterprise Management cluster host provider	DKP Enterprise Management cluster IAM permissions	DKP Essential cluster host provider
GCP	https://cloud.google.com/iam/docs/overview	AWS, GCP, vSphere, Pre-provisioned
vSphere	https://docs.vmware.com/en/vRealize-Operations/Cloud/com.vmware.vcom.config.doc/GUID-F85638E3-937E-4E31-90D0-9D4A5E479292.html	AWS, GCP, vSphere, Pre-provisioned
Azure	https://learn.microsoft.com/en-us/azure/active-directory/fundamentals/active-directory-ops-guide-iam	Azure
Pre-provisioned	NA	AWS, GCP, vSphere, Pre-provisioned

6.5.4.8.2 Moving the CAPI Resources

1. Following the conversion into a DKP Enterprise managed cluster, run the following command to move the CAPI Objects:

```
dkp move capi-resources --from-kubeconfig <essential_cluster_kubeconfig> --to-kubeconfig <enterprise_cluster_kubeconfig> --to-namespace $
{WORKSPACE_NAMESPACE}
```

2. Verify that the output looks similar to the following:

```
✓ Moving cluster resources
You can now view resources in the moved cluster by using the --kubeconfig flag
with kubectl. For example: kubectl --kubeconfig=<enterprise_cluster_kubeconfig>
get nodes
```

3. Since cert-manager is not managed by the Kommander component of DKP in an Essential cluster, you must change the ownership of the cert-manager application and give ownership to Kommander.

On the expanded, now considered a *Managed Cluster*, **annotate** the elements so they are managed by the Kommander installation of cert-manager. First run the command to annotate `release-namespace` :

```
kubectl annotate
deployment,sa,clusterrole,clusterrolebinding,role,service,MutatingWebhookConfiguration,ValidatingWebhookConfiguration,configmap,RoleBinding -l
app.kubernetes.io/instance=cert-manager meta.helm.sh/release-namespace=cert-manager --all-namespaces
```

Next run the command to annotate `release-name` :

```
kubectl annotate
deployment,sa,clusterrole,clusterrolebinding,role,service,MutatingWebhookConfiguration,ValidatingWebhookConfiguration,configmap,RoleBinding -l
app.kubernetes.io/instance=cert-manager meta.helm.sh/release-name=cert-manager --all-namespaces
```

4. On the expanded (*Managed Cluster*), **label** the cert-manager to be managed by the Kommander installation of cert-manager by running:

```
kubectl label
deployment,sa,clusterrole,clusterrolebinding,role,service,MutatingWebhookConfiguration,ValidatingWebhookConfiguration,configmap,RoleBinding -l
app.kubernetes.io/instance=cert-manager app.kubernetes.io/managed-by=Helm --all-namespaces
```

5. While still using the Managed Cluster, check to see if the other cert-manager objects are still ready:

```
kubectl get
Issuers,ClusterIssuers,Certificates,CertificateRequests,Orders,Challenges --all-namespaces
```

Output will be similar to this:

NAMESPACE	NAME	READY	AGE
	clusterissuer.cert-manager.io/kommander-ca	True	111m
	clusterissuer.cert-manager.io/selfsigned-issuer	True	111m
NAMESPACE	NAME	READY	AGE
SECRET			
cert-manager	certificate.cert-manager.io/kommander-ca	True	
kommander-ca			111m
kommander	certificate.cert-manager.io/dex	True	
tls			101m
			dex-

```

kommander      certificate.cert-manager.io/git-tls      True    git-
tls            111m
kommander      certificate.cert-manager.io/kommander-traefik-tls  True
kommander-traefik-tls      102m
kommander      certificate.cert-manager.io/kube-oidc-proxy      True
kube-oidc-proxy-server-tls  101m

NAMESPACE      NAME
APPROVED  DENIED  READY  ISSUER  REQUESTOR
AGE
cert-manager  certificaterequest.cert-manager.io/kommander-ca-kk9zw
True          True    selfsigned-issuer  system:serviceaccount:cert-
manager:cert-manager  111m
kommander    certificaterequest.cert-manager.io/git-tls-dgcxv
True          True    kommander-ca      system:serviceaccount:cert-
manager:cert-manager  111m
kommander    certificaterequest.cert-manager.io/kommander-traefik-tls-lqn8z
True          True    kommander-ca      system:serviceaccount:cert-
manager:cert-manager  102m
kommander    certificaterequest.cert-manager.io/kube-oidc-proxy-j5st6
True          True    kommander-ca      system:serviceaccount:cert-
manager:cert-manager  101m

```

6. The cert-manager application will still be running, but you want to enable the **Workspace** application for the cluster. This ensures that any updates you make to your **Workspace** will be enforced by Kommander and will upgrade as expected in the future.
 - a. In the DKP UI, go to the **Workspace** you attached your Essential cluster.
 - b. Select **Applications** from the left navigation menu.
 - c. Locate the cert-manager in the **Security** section and select the three button menu at the bottom of the card and select **Edit**.
 - d. From the edit cert-manager Workspace Application page, select the checkbox for the newly managed cluster and select the **Enable/Save** button.
7. After moving the resources, run the following command to remove the CAPI controller manager deployments:

```
dkp delete capi-components --kubeconfig <essential_cluster_kubeconfig>
```

6.5.4.8.3 Next Step

[Troubleshooting](#) (see page 851)

6.5.4.9 Troubleshooting

Enterprise

Gov Advanced

When an error or failure occurs when converting a DKP Essential cluster to a DKP Enterprise Managed cluster, DKP automatically keeps retrying the cluster's conversion and attachment process. You do not need to trigger it manually.

If the state does not improve after a while, here are some ways in which you can check or troubleshoot the failed conversion:

- [Verify the Conversion Status of your Cluster](#) (see page 852)
- [Errors Related to CAPI Resources](#) (see page 853)
- [Restore your Backup and Retry the Cluster Expansion](#) (see page 855)

6.5.4.9.1 Verify the Conversion Status of your Cluster

Enterprise

Gov Advanced

6.5.4.9.1.1 Verify the Conversion



Run the following commands in the **Management cluster**. For general guidelines on how to set the context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

1. Export the environment variable for the workspace namespace:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

2. To verify that the conversion is successful, check the `KommanderCluster` object:

```
kubectrl wait --for=condition=AttachmentCompleted kommandercluster <cluster name> -n ${WORKSPACE_NAMESPACE} --timeout 30m
```

The following output appears if the conversion is successful:

```
kommandercluster.kommander.mesosphere.io/<cluster name> condition met
```

If the condition is not met yet, you can observe the conversion process:

6.5.4.9.1.2 Observe the Conversion Process

1. Export the environment variable for the workspace namespace:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

- Print the state of your cluster's `KommanderCluster` object through the CLI, and observe the cluster conversion process:

```
kubectl get kommandercluster -n ${WORKSPACE_NAMESPACE} <essential_cluster_name>
-o go-template='{{range .status.conditions }}type: {{.type}} {"\n"}}status:
{{.status}} {"\n"}}reason: {{.reason}} {"\n"}}lastTxTime:
{{.lastTransitionTime}} {"\n"}}message: {{.message}} {"\n\n"}}{{end}}'
```

The output looks similar to this:

```
type: IngressAddressReady
status: False
reason: IngressServiceNotFound
lastTxTime: 2023-02-24T14:58:18Z
message: Ingress service object was not found in the cluster

type: IngressCertificateReady
status: True
reason: Ready
lastTxTime: 2023-02-24T14:49:09Z
message: Certificate is up to date and has not expired

type: CAPIResourceMoved
status: True
reason: Succeeded
lastTxTime: 2023-02-24T14:50:56Z
message: Moved CAPI resources from attached cluster to management cluster

type: PreAttachmentCleanup
status: True
reason: Succeeded
lastTxTime: 2023-02-24T14:54:47Z
message: pre-attach cleanup succeeded

# [...]
```

6.5.4.9.1.3 Next Step:

[Errors Related to CAPI Resources \(see page 853\)](#)

6.5.4.9.2 Errors Related to CAPI Resources

Enterprise

Gov Advanced

6.5.4.9.2.1 Failed Condition Reason: FailedToIdentifyCAPIResources

kubeconfigRef points to the wrong secret

Verify the `KommanderCluster` object of both your Essential and Enterprise clusters. The `spec.kubeconfigRef.name` of each object should point to a valid `kubeconfig` secret.

1. Download the referenced `kubeconfig` to your local machine:

```
kubectl get secret -n <WORKSPACE_NAMESPACE> <cluster_name>-kubeconfig -o
jsonpath='{.data.value}' | base64 --decode > <cluster_name>-kubeconfig
```

2. Verify if the `kubeconfig` is valid:

```
kubectl get namespaces -A --kubeconfig <cluster_name>-kubeconfig
```

3. Verify the output of the previous command:

No errors in the output

If your output shows no errors, the error message is not related to a `kubeconfig`.

Error in the output

If the output shows an error, delete the `KommanderCluster` object via CLI:

```
kubectl delete kommandercluster -n <WORKSPACE_NAMESPACE> <WRONG_KOMMANDER_CLUSTER>
```

i At this particular stage and in the context of converting your cluster, deleting your `KommanderCluster` will not affect your environment. However, DO NOT delete your `KommanderCluster` in other scenarios, as it detaches the referenced cluster from the Management cluster.

Finally, restart the cluster conversion process with the [UI \(see page 845\)](#) or [CLI \(see page 847\)](#).

Essential cluster has more than one instance of `v1beta1/clusters.cluster.x-k8s.io`



D2iQ does not support converting Essential clusters that contain the Cluster API resources of more than one cluster.

Ensure your Essential cluster only contains its own CAPI resources and does not contain the CAPI resources of other clusters.

6.5.4.9.2.2 Next Topic:

[Restore your Backup and Retry the Cluster Expansion \(see page 855\)](#)

6.5.4.9.3 Restore your Backup and Retry the Cluster Expansion

When an error or failure occurs when converting a DKP Essential cluster to a DKP Enterprise Managed cluster, DKP automatically keeps retrying the cluster's conversion and attachment process. You do not need to trigger it manually.

If you must interrupt the expansion process to restore your cluster and retry the expansion procedure, follow these instructions:

- [Restore your Cluster \(see page 855\)](#)
- [Move your CAPI Resources Back to the DKP Essential Cluster \(see page 856\)](#)
- [Verify the Restore Process and Retry the Expansion \(see page 857\)](#)

6.5.4.9.3.1 Restore your Cluster

Prerequisites

- You have [backed up your cluster \(see page 835\)](#).
- The [cluster expansion \(see page 830\)](#) you attempted was not successful.

Restore the Cluster



Switch between **DKP Enterprise Management** and **DKP Essential** clusters for the following commands. For general guidelines on how to set the context, refer to [Provide Context for Commands with a kubeconfig File \(see page 87\)](#).

1. Delete the `KommanderCluster` object on the **DKP Enterprise Management** cluster:

```
kubectl -n <WORKSPACE_NAMESPACE> delete kommandercluster
<KOMMANDER_CLUSTER_NAME> --wait=false
```

2. Disable the Flux controllers on the **DKP Essential** cluster to interrupt the expansion process:

```
kubectl -n kommander-flux delete deployment -l app.kubernetes.io/instance=kommander-flux
```

3. Delete the `kube-federation-system` namespace on the **DKP Essential** cluster:

```
kubectl get ns kube-federation-system -o json | jq '.spec.finalizers = []' | kubectl replace --raw "/api/v1/namespaces/kube-federation-system/finalize" -f -
```

4. Restore your cluster's configuration on the **DKP Essential** cluster:

```
velero restore create pre-expansion --from-backup pre-expansion --existing-resource-policy update --wait --namespace kommander
```

Next Step:

[Move your CAPI Resources Back to the DKP Essential Cluster \(see page 856\)](#)


6.5.4.9.3.2 Move your CAPI Resources Back to the DKP Essential Cluster

Previous Step:

[Restore your Cluster \(see page 858\)](#)

Move your cluster's CAPI Resources

Because the created backup does not include CAPI resources, you will also have to move the resources back to the DKP Essential cluster.

 Ensure you replace `<CAPI_CLUSTER_NAME>` with the name of the DKP Essential cluster you were converting in the current workflow. If you accidentally provide the CAPI cluster name of another Managed cluster, the command will move the CAPI resources of the incorrect cluster to the DKP Essential cluster.

1. Retrieve your Managed cluster's kubeconfig and write it to the `<essential.conf>` file:

```
KUBECONFIG=management.conf ./dkp get kubeconfig -n <WORKSPACE_NAMESPACE> -c <CAPI_CLUSTER_NAME> > essential.conf
```

2. Move your CAPI resources back to the DKP Essential cluster:

```
dkp move capi-resources --from-kubeconfig management.conf --to-kubeconfig
<essential.conf> -n <WORKSPACE_NAMESPACE> --to-namespace default
```

Next Step:

[Verify the Restore Process and Retry the Expansion \(see page 857\)](#)

6.5.4.9.3.3 Verify the Restore Process and Retry the Expansion

Previous Step:

[Move your CAPI Resources Back to the DKP Essential Cluster \(see page 859\)](#)

Verify the Restore Process

Verify that you successfully restored your cluster:

```
velero restore describe pre-expansion --namespace kommander
```

The output looks similar to this:

```
Name:          pre-expansion
Namespace:     kommander
Labels:        <none>
Annotations:   <none>

Phase:                Completed
Total items to be restored: 2411
Items restored:       2411

...
```

Retry the Expansion Process

Execute the cluster expansion again, as shown in [Platform Expansion: Convert a DKP Essential Cluster to a DKP Enterprise Managed Cluster \(see page 830\)](#).

Related Topic:

[Troubleshooting \(see page 854\)](#)

6.5.5 Advanced Creation of CLI Clusters

Create CLI clusters



This feature is for advanced users and users in unique environments only. We highly recommend using other documented methods to create clusters whenever possible.

6.5.5.1 Generate Cluster Objects

Depending on your infrastructure, DKP CLI can generate a set of cluster objects that can be customized for unusual use cases. Visit the [Advanced Configuration documentation for AWS](#) (see page 1076) for an example on how to use the `--dry-run` and `--output` flags to create a set of cluster objects.

6.5.5.2 Use the Upload YAML Form

1. In the selected workspace Dashboard, select the **Add Cluster** option in the **Actions** dropdown menu at the top right.
2. On the Add Cluster page, select **Upload YAML to Create a Cluster** and provide advanced cluster details:
 - **Workspace:** The workspace where this cluster belongs (if within the Global workspace).
 - **Cluster YAML:** Paste or upload your customized set of cluster objects into this field. Only valid YAML is accepted.
 - **Add Labels:** By default, your cluster has labels that reflect the infrastructure provider provisioning. For example, your AWS cluster may have a label for the data center region and `provider: aws`. Cluster labels are matched to the selectors created for [Projects](#) (see page 689). Changing a cluster label may add or remove the cluster from projects.
3. Click **Create** to begin provisioning the DKP CLI cluster. This step may take a few minutes, taking time for the cluster to be ready and fully deploy its components. The cluster automatically tries to join, and should resolve after it is fully-provisioned.

6.5.6 Custom Domains and Certificates Configuration for All Cluster Types

Configure a custom domain and certificate after installing the Kommander component.

For any cluster type, including [Managed](#) (see page 82) and [Attached](#) (see page 82) clusters.

Configuration Method	WHILE installing the Kommander component	AFTER installing the Kommander component
----------------------	--	--

Supported cluster types	Only Essential or Management clusters	All cluster types
Documentation	Go to Custom Domains and Certificates for Management and Essential Clusters (see page 1428)	Remain in this section

DKP supports configuring a custom **domain** name for accessing the UI and other platform services, as well as setting up manual or automatic **certificate renewal** or rotation. This section provides instructions and examples on how to configure a customized domain and certificate on **Essential** (see page 83), **Management** (see page 82), **Managed** (see page 82) or **Attached** (see page 82) clusters.

Section Contents:

- [Why Should You Set Up a Custom Domain or Certificate?](#) (see page 859)
- [KommanderCluster and Certificate Issuer Concepts](#) (see page 860)
- [Configure Custom Domains or Custom Certificates post Kommander Installation](#) (see page 861)

6.5.6.1 Why Should You Set Up a Custom Domain or Certificate?

6.5.6.1.1 Reasons for Using a Custom DNS Domain

DKP supports the customization of domains to allow you to use your own domain or hostname for your services. For example, you can set up your DKP UI or any of your clusters to be accessible with your custom domain name instead of the domain provided by default.

- To set up a custom domain (without a custom certificate), refer to [Configure a Custom Domain without a Custom Certificate](#) (see page 1440).

6.5.6.1.2 Reasons for Using a Custom Certificate

DKP's default CA identity supports the encryption of data exchange and traffic (between your client and your environment's server). To configure an additional security layer that validates your environment's server authenticity, DKP supports configuring a custom certificate issued by a trusted Certificate Authority either directly in a Secret or managed automatically using the ACME protocol (for example, Let's Encrypt).

Changing the default certificate for any of your clusters can be helpful. For example, you can adapt it to classify your DKP UI or any other type of service as trusted (when accessing a service via a browser).

To set up a custom domain and certificate, refer to the following pages respectively:

- Configure a custom domain and certificate as part of the [cluster's installation process](#) (see page 1432). This is only possible for your [Management/Essential cluster](#) (see page 81).
- Update your cluster's current domain and certificate configuration as part of your [cluster's Day 2 operations](#) (see page 861). You can do this for any [cluster type](#) (see page 81) in your environment.

Using Let's Encrypt or other public ACME certificate authorities **does not work in air-gapped scenarios**, as these services require connection to the Internet for their setup. For air-gapped environments, you can either use self-signed certificates issued by the cluster (the default configuration), or a certificate created manually using a trusted Certificate Authority.

6.5.6.1.3 Next Topic:

[KommanderCluster and Certificate Issuer Concepts](#) (see page 860)

6.5.6.2 KommanderCluster and Certificate Issuer Concepts

6.5.6.2.1 KommanderCluster Object

The `KommanderCluster` resource is an object that contains key information for [all types of clusters](#) (see page 81) that are part of your environment, such as:

- Cluster access and endpoint information
- Cluster attachment information
- Cluster status and configuration information

6.5.6.2.2 Issuer Objects:

`Issuer`, `ClusterIssuer` or `certificateSecret` ?

If you use a certificate issued and managed automatically by `cert-manager`, you need an `Issuer` or `ClusterIssuer` that you reference in your `KommanderCluster` resource. The referenced object must contain the information of your certificate provider.

If you want to use a manually-created certificate, you need a `certificateSecret` that you reference in your `KommanderCluster` resource.

6.5.6.2.3 Location of the KommanderCluster and Issuer Objects: Management, Managed or Attached cluster?

In the [Management](#) (see page 82) or [Essential](#) (see page 83) cluster, both the `KommanderCluster` and `issuer` objects are stored on the same cluster. The `issuer` can be referenced as an `Issuer`, `ClusterIssuer` or `certificateSecret`.

In **Managed** (see page 82) and **Attached** (see page 82) clusters, the `KommanderCluster` object is stored on the Management cluster. The `Issuer`, `ClusterIssuer` or `certificateSecret` is stored on the Managed or Attached cluster.

6.5.6.2.4 HTTP or DNS solver?

When configuring a certificate for your DKP cluster, you can set up an *HTTP solver* or a *DNS solver*. The HTTP protocol exposes your cluster to the public Internet, whereas DNS keeps your traffic hidden. If you use HTTP, your cluster must be publically accessible (via the ingress or load balancer). If you use DNS, this is not a requirement. See [Advanced Configuration: ClusterIssuer](#) (see page 1437) for HTTP and DNS configuration options.



If you are enabling a [proxied access](#) (see page 817) for a [network-restricted cluster](#) (see page 84), this configuration is restricted to DNS.

6.5.6.2.5 Next Step:

[Configure Custom Domains or Custom Certificates post Kommander Installation](#) (see page 861)

6.5.6.3 Configure Custom Domains or Custom Certificates post Kommander Installation

This page contains instructions on how to set up custom certificates for [any cluster type](#) (see page 81) after installing DKP.

There are two configuration methods:

Configuration Method	<i>WHILE</i> installing the Kommander component	<i>AFTER</i> installing the Kommander component
Supported cluster types	Only Essential or Management clusters	All cluster types
Documentation	Go to Configure the Kommander Installation with a Custom Domain and Certificate (see page 1432)	Remain in this page

6.5.6.3.1 Configuration Options

After you have installed the Kommander component of DKP, you can configure a custom domain and certificate by modifying the `KommanderCluster` object of your cluster. You have several options to establish a custom domain and certificate.

- ☐ Choose an ACME-supported Certificate Authority, if you want the `cert-manager` can automatically handle **certificate renewal** and **rotation**.

I want to use an automatically-generated certificate with ACME

6.5.6.3.1.1 Automatically-generated Certificate

Use a certificate that is managed automatically and supported by `cert-manager` :

1. Create an `Issuer` or `ClusterIssuer` with your certificate provider information. Store this object in the cluster where you want to customize the certificate and domain.
 - a. See <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/137593048/Configuration+Example+with+Let+s+Encrypt> (see page 864) if you want to use DKP's default certificate authority.
 - b. For an advanced configuration example, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/175636481/Configure+the+Kommander+Installation+with+a+Custom+Domain+and+Certificate#basic> (see page 1433), and select *I want to use an automatically-generated certificate with ACME and require advanced configuration*.
2. Update the `KommanderCluster` by referencing the **name of the created** `Issuer` or `ClusterIssuer` in the `spec.ingress.issuerRef` field.

Enter the custom domain name in the `spec.ingress.hostname` field:

```
cat <<EOF | kubectl -n <workspace_namespace> --kubeconfig
<management_cluster_kubeconfig> patch \
kommandercluster <cluster_name> --type='merge' --patch-file=/dev/stdin
spec:
  ingress:
    hostname: <cluster_hostname>
    issuerRef:
      name: <issuer_name>
      kind: Issuer # or ClusterIssuer depending on the issuer config
EOF
```

⚠ Certificates issued by another Issuer

You can also configure a certificate issued by another Certificate Authority. In this case, the CA will determine which information to include in the configuration.

- Refer to <https://cert-manager.io/docs/configuration/> for configuration examples.
- The `ClusterIssuer`'s name **MUST BE** `kommander-acme-issuer`.

I have a manually-generated certificate

6.5.6.3.1.2 Manually-generated Certificate

Use a manually-created certificate that is **customized for your hostname**.

1. Obtain or create a certificate that is customized for your hostname. Store this object in the workspace namespace of the target cluster.
2. Create a secret with the certificate in the cluster's namespace. Give it a name by replacing `<certificate_secret_name>`:

```
kubectl create secret generic -n "${WORKSPACE_NAMESPACE}"
<certificate_secret_name> \
  --from-file=ca.crt=$CERT_CA_PATH \
  --from-file=tls.crt=$CERT_PATH \
  --from-file=tls.key=$CERT_KEY_PATH \
  --type=kubernetes.io/tls
```

3. Update the `KommanderCluster` by referencing this **secret** in the `spec.ingress.certificateSecretRef` field and provide the custom domain name in the `spec.ingress.hostname`:

```
cat <<EOF | kubectl -n <workspace_namespace> --kubeconfig
<management_cluster_kubeconfig> patch \
kommandercluster <cluster_name> --type='merge' --patch-file=/dev/stdin
spec:
  ingress:
    hostname: <cluster_hostname>
    certificateSecretRef:
      name: <certificate_secret_name>
EOF
```

- For Kommander to access the secret containing the certificate, it must be located in the workspace namespace of the target cluster.

6.5.6.3.1.3 Next Step:

[Verify and Troubleshoot Domain and Certificate Customization](#) (see page 865)

6.5.6.3.1.4 Related topics

[Why Should You Set Up a Custom Domain or Certificate?](#) (see page 862)

<https://d2iq.atlassian.net/wiki/spaces/DENT/pages/137593048/Configuration+Example+with+Let+s+Encrypt> (see page 864)

[Advanced Configuration: ClusterIssuer](#) (see page 1437)

6.5.6.3.2 Configuration Example with Let's Encrypt

6.5.6.3.2.1 Configure a Custom Certificate with Let's Encrypt

Let's Encrypt is one of the Certificate Authorities (CA) supported by `cert-manager`. To set up a Let's Encrypt certificate, create an `Issuer` or `ClusterIssuer` in the target cluster and then reference it in the `issuerRef` field of the `KommanderCluster` resource.

1. Create the Let's Encrypt ACME cert-manager issuer:

```
cat <<EOF | kubectl apply -f -
apiVersion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
  name: kommander-acme-issuer
spec:
  acme:
    email: <your_email>
    server: https://acme-v02.api.letsencrypt.org/directory
    privateKeySecretRef:
      name: kommander-acme-issuer-account
    solvers:
      - dns01:
          route53:
            region: us-east-1
            role: arn:aws:iam::YYYYYYYYYYYY:role/dns-manager
```

```
EOF
```

2. Configure the Management cluster to use your `custom-domain.example.com` with a certificate issued by Let's Encrypt by referencing the created `ClusterIssuer` :

```
cat <<EOF | kubectl -n kommander --kubeconfig <management_cluster_kubeconfig>
patch \ kommandercluster host-cluster --type='merge' --patch-file=/dev/stdin
spec:
  ingress:
    hostname: custom-domain.example.com
    issuerRef:
      name: custom-acme-issuer
      kind: ClusterIssuer
EOF
```

6.5.6.3.2.2 Next Step:

[Verify and Troubleshoot Domain and Certificate Customization \(see page 865\)](#)

6.5.6.3.2.3 Related Topics

[Why Should You Set Up a Custom Domain or Certificate? \(see page 862\)](#)

[Configure Custom Domains or Custom Certificates post Kommander Installation \(see page 864\)](#)

6.5.6.3.3 Verify and Troubleshoot Domain and Certificate Customization

If you want to ensure the customization for a domain and certificate is completed, or if you want to obtain more information on the status of the customization, display the status information for the `KommanderCluster` . On the [Management cluster \(see page 82\)](#):

1. Inspect the modified `KommanderCluster` object:

```
kubectl describe kommandercluster -n <workspace_name> <cluster_name>
```

2. If the ingress is still being provisioned, the output looks similar to this:

```
[...]
Conditions:
  Last Transition Time: 2022-06-24T07:48:31Z
  Message:             Ingress service object was not found in the cluster
  Reason:              IngressServiceNotFound
  Status:              False
```

```
Type:          IngressAddressReady
[...]
```

If the provisioning has been completed, the output looks similar to this:

```
[...]
Conditions:
  Last Transition Time: 2022-06-28T13:43:33Z
  Message:             Ingress service address has been provisioned
  Reason:              IngressServiceAddressFound
  Status:              True
  Type:                IngressAddressReady
  Last Transition Time: 2022-06-28T13:42:24Z
  Message:             Certificate is up to date and has not expired
  Reason:              Ready
  Status:              True
  Type:                IngressCertificateReady
[...]
```

The same command also prints the actual customized values for the `KommanderCluster.Status.Ingress`. Here is an example:

```
[...]
  ingress:
    address: 172.20.255.180
    caBundle: LS0tLS1CRUdJTiBD...<output has been shortened>...DQVRFLS0tLS0K
[...]
```

6.5.6.3.3.1 Related Topics

[Why Should You Set Up a Custom Domain or Certificate?](#) (see page 862)

<https://d2iq.atlassian.net/wiki/spaces/DENT/pages/137593048/Configuration+Example+with+Let+s+Encrypt> (see page 867)


[Configure Custom Domains or Custom Certificates post Kommander Installation](#) (see page 864)

6.5.7 Disconnect or Delete Clusters

6.5.7.1 Disconnect or delete a cluster


6.5.7.2 Disconnect vs. Delete

When you attach a cluster to Kommander that was **not created with Kommander**, you can later detach it. This does not alter the running state of the cluster, but simply removes it from the DKP UI. User workloads, platform services, and other Kubernetes resources are not cleaned up at detach.

 After successfully detaching the cluster, manually disconnect the attached cluster's Flux installation from the management Git repository. Otherwise, changes to apps in the managed cluster's workspace will still be reflected on the cluster you just detached. Ensure your `dkp` configuration references the target cluster. You can do this by setting the `KUBECONFIG` environment variable [to the appropriate kubeconfig file location](#)⁵⁹⁵. An alternative to initializing the `KUBECONFIG` environment variable is to use the `-kubeconfig=cluster_name.conf` flag. Then, run `kubectl -n kommander-flux patch gitrepo management -p '{"spec":{"suspend":true}}' --type merge` to make the cluster's workloads not managed by Kommander, anymore.

If you **created a managed cluster with Kommander**, you cannot disconnect the cluster, but you can **delete** the cluster. This completely removes the cluster and all of its cloud assets.

We recommend deleting a managed cluster via the DKP UI.

 If you delete the Management (Konvoy) cluster, you can not use Kommander to delete any Managed clusters created by Kommander. If you want to delete all clusters, ensure you delete any Managed clusters before finally deleting the Management cluster.

6.5.7.2.1 Statuses

See [Statuses](#) (see page 869) for a list of possible states a cluster can have when it is getting disconnected or deleted.

⁵⁹⁵ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

6.5.7.3 Troubleshooting

I cannot detach an attached cluster that is “Pending”

OR

The cluster I deleted via CLI still appears in the UI with “Error” state

Sometimes detaching or deleting a Kubernetes cluster causes that cluster to get stuck in a “Pending” or “Error” state. This can happen because the wrong `kubeconfig` file is used, or the cluster is just not reachable. In order to detach the cluster, so it does not show in the UI, follow these steps:

1. Determine the `KommanderCluster` resource backing the cluster you tried to attach/detach:

```
kubectl -n WORKSPACE_NAMESPACE get kommandercluster
```

Replace `WORKSPACE_NAMESPACE` with the actual current workspace name. You can find this name by going to `https://YOUR_CLUSTER_DOMAIN_OR_IP_ADDRESS/dkp/kommander/dashboard/workspaces` in your browser.

2. Delete the cluster:

```
kubectl -n WORKSPACE_NAMESPACE delete kommandercluster CLUSTER_NAME
```

3. If the resource does not go after a short time, remove its finalizers:

```
kubectl -n WORKSPACE_NAMESPACE patch kommandercluster CLUSTER_NAME --type json -p '[{"op": "remove", "path": "/metadata/finalizers"}]'
```

This removes the cluster from the DKP UI.

6.5.8 Management Cluster

A guide for the Management Cluster and the Management Cluster Workspace

When you install Kommander, the host cluster is attached in the **Management Cluster Workspace** as the **Management** (see page 82) **Cluster**, called **Management Cluster** in the Global workspace dashboard and **Kommander Host** inside the Management Cluster Workspace. This allows the Management Cluster to be included in **Projects** (see page 689) and enables the management of its **Platform Applications** (see page 691) from the Management Cluster Workspace.



Do not attach a cluster in the "Management Cluster Workspace" workspace. This workspace is reserved for your Kommander Management cluster only.

6.5.8.1 Editing

As an attached cluster, you can edit the Management Cluster to add or remove Labels, then you can use these labels to include the Management Cluster in Projects inside of the Management Cluster Workspace.

6.5.8.2 Disconnecting

The Management Cluster cannot be disconnected from the GUI like other attached clusters. Because of this, the Management Cluster Workspace cannot be deleted from the GUI as it will always have the Management Cluster inside itself.

6.5.9 Cluster Statuses

A cluster card's status line displays both the current status and the version of Kubernetes running in the cluster.

The status list includes these values:

Status	Description
Pending	This is the initial state when a cluster is created or connected.
Pending Setup	The cluster has networking restrictions that require additional setup, and is not yet connected or attached.
Loading Data	The cluster has been added to Kommander and we are fetching details about the cluster. This is the status before <code>Active</code> .
Active	The cluster is connected to API server.
Provisioning*	The cluster is being created on your cloud provider. This process may take some time.
Provisioned*	The cluster's infrastructure has been created and configured.
Joining	The cluster is being joined to the management cluster for federation.
Joined	The join process is done, and waiting for the first data from the cluster to arrive.
Deleting*	The cluster and its resources are being removed from your cloud provider. This process may take some time.

Status	Description
Error	There has been an error connecting to the cluster or retrieving data from the cluster.
Join Failed	This status can appear when kubefed does not have permission to create entities in the target cluster.
Unjoining	Kubefed is cleaning up after itself, removing all installed resources on the target cluster.
Unjoined	The cluster has been disconnected from the management cluster.
Unjoin Failed	The Unjoin from kubefed failed or there is some other error with deleting or disconnecting.
Unattached*	The cluster was created manually and the infrastructure has been created and configured. However, the cluster is not attached. Review the Manually attach a CLI-created cluster (see page 827) page to resolve this status.

*These statuses only appear on Managed clusters.

6.5.10 Cluster Resources

The Resources graphs on a cluster card show you a cluster's resource requests, limits, and usage. This allows a quick, visual scan of cluster health. Hover over each resource to get specific details for that specific cluster resource.

Resource	Description
CPU Requests	The requested portion of the total allocatable CPU resource for the cluster, measured in number of cores, such 0.5 cores.
CPU Limits	The portion of the total allocatable CPU resource to which the cluster is limited, measured in number of cores, such as 0.5 cores.
CPU Usage	The amount of the allocatable CPU resource being consumed. Cannot be higher than the configured CPU limit. Measured in number of cores, such as 0.5 cores)
Memory Requests	The requested portion of the total allocatable memory resource for the cluster, measured in bytes, such as 64 GiB.

Resource	Description
Memory Limits	The portion of the allocatable memory resource to which the cluster is limited, measured in bytes, such as 64 GiB.
Memory Usage	The amount of the allocatable memory resource being consumed. Cannot be higher than the configured memory limit. Measured in bytes, such as 64 GiB.
Disk Requests	The requested portion of the allocatable ephemeral storage resource for the cluster, measured in bytes, such as 64 GiB.
Disk Limits	The portion of the allocatable ephemeral storage resource to which the cluster is limited, measured in bytes, such as 64 GiB.

For more detailed information, see the [Kubernetes documentation](#)⁵⁹⁶ about resources.

6.5.11 DKP Platform Applications

Platform Applications are applications that DKP provides out-of-the-box with functionality such as observability, cost management, monitoring, logging, making DKP clusters day 2 production-ready right from install.

Platform applications are the applications selected by D2iQ from the open source community for use by the DKP platform. You can visit a cluster's detail page to see which platform applications are enabled under the "Platform Applications" section.

Review the [Workspace Platform Configuration Requirements](#) (see page 105) to ensure that the attached clusters have sufficient resources. For more information on platform applications and how to customize them, refer to the pages in the [Platform Applications](#) (see page 642) section.

6.5.12 Cluster Applications and Statuses

Applications are installed by the management cluster. You can visit a cluster's detail page to see the application dashboards enabled from the deployed applications under the **Application Dashboards** section.

Under the **Applications** section of the cluster's detail page, you can view the workspace applications enabled for the cluster, grouped by category.

In this section, you can also view the current status of the enabled applications on the cluster, on each application card. Hovering on the status displays details about the status of the application.

Cluster applications can have one of the following statuses:

⁵⁹⁶ <https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/>

Status	Description
Enabled	The application is enabled, but the status on the cluster is not available.
Pending	The application is waiting to be deployed.
Deploying	The application is currently being deployed to the cluster.
Deployed	The application has successfully been deployed to the cluster.
Deploy Failed	The application failed to deploy to the cluster.

Review the [Workspace Platform Application Configuration Requirements](#) (see page 105) to ensure that the attached clusters have sufficient resources. For more information on applications and how to customize them, see [Workspace Applications](#) (see page 666).

6.5.12.1 Custom Cluster Application Dashboard Cards

You can add custom application dashboard cards to the cluster detail page's Applications section by creating a `ConfigMap` on the cluster. The `ConfigMap` must have a `kommander.d2iq.io/application` label applied through the CLI, and must contain both `name` and `dashboardLink` data keys to be displayed. Upon creation of the `ConfigMap`, the DKP UI displays a card corresponding to the data provided in the `ConfigMap`. Custom application cards have a Kubernetes icon, and can link to a service running in the cluster, or use an absolute URL to link to any accessible URL.

6.5.12.1.1 ConfigMap Example

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: "my-app"
  namespace: "app-namespace"
  labels:
    "kommander.d2iq.io/application": "my-app"
data:
  name: "My Application"
  dashboardLink: "/path/to/app"
```

Key	Description	Required
<code>metadata.labels."kommander.d2iq.io/application"</code>	The application name (ID).	X
<code>data.name</code>	The display name that describes the application and displays on the custom application card in the user interface.	X
<code>data.dashboardLink</code>	The link to the application. This can be an absolute link, <code>https://www.d2iq.com</code> or a relative link, <code>/dkp/kommander/dashboard</code> . If you use a relative link, the link is built using the cluster's path as the base of the URL to the application.	X
<code>data.docsLink</code>	Link to documentation about the application. This is displayed on the application card, but omitted if not present.	
<code>data.category</code>	Category with which to group the custom application. If not provided, the application is grouped under the category, "None."	
<code>data.version</code>	A version string for the application. If not provided, "N/A" is displayed on the application card in the user interface.	

Use a command similar to this to create a new custom application ConfigMap :

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: "my-app"
```

```

namespace: "default"
labels:
  "kommander.d2iq.io/application": "my-app"
data:
  name: "My Application"
  dashboardLink: "/path/to/app"
EOF

```

6.5.13 Kubernetes Cluster Federation (KubeFed)

[Kubernetes Cluster Federation \(KubeFed\)](#)⁵⁹⁷ allows you to coordinate the configuration of multiple Kubernetes clusters from a single set of APIs in a hosting cluster. KubeFed aims to provide mechanisms for expressing which clusters should have their configuration managed and what that configuration should be. The mechanisms that KubeFed provides are intentionally low-level, and intended to be foundational for more complex multicluster use cases, such as deploying multi-geo applications and disaster recovery.

DKP uses KubeFed to manage multiple clusters from the management cluster and also to federate various resources. A `KubefedCluster` object is automatically created for each attached cluster and joined to the management cluster. After they are joined, namespaces can be federated to the clusters - this is how you get workspace and project namespaces created on the attached clusters. From here other resources can be federated into those namespaces, such as ConfigMaps, RBAC, and so on.

See these pages for more information:

- <https://github.com/kubernetes-sigs/kubefed/blob/master/docs/concepts.md>
- <https://github.com/kubernetes-sigs/kubefed/blob/master/docs/userguide.md>

6.6 Backup and Restore

For production clusters, regular maintenance should include routine backup operations to ensure data integrity and reduce the risk of data loss due to unexpected events. Backup operations should include the cluster state, application state, and the running configuration of both stateless and stateful applications in the cluster.

DKP stores all data as CRDs in the Kubernetes API and you can back up and restore it. Choose a procedure depending on your infrastructure provider:

- [Configure Velero](#) (see page 874)
- [Back up with Velero](#) (see page 893)

6.6.1 Configure Velero

For default installations, DKP deploys [Velero](#)⁵⁹⁸ integrated with [Rook Ceph](#)⁵⁹⁹, operating inside the same cluster.

⁵⁹⁷ <https://github.com/kubernetes-sigs/kubefed>

⁵⁹⁸ <https://velero.io/>

⁵⁹⁹ <https://rook.io/>

For production use-cases, D2iQ advises providing an *external* storage class to use with [Rook Ceph](#)⁶⁰⁰. See the [Rook Ceph in DKP](#) (see page 1004) for more information.

6.6.1.1 Cloud Provider Storage

If you do not want to use Rook Ceph for storing Velero backups, you can [configure Velero to use cloud provider storage](#) (see page 875).

- [Use Velero with AWS](#) (see page 875)
- [Use Velero with Azure](#) (see page 881)
- [Use Velero with GCP](#) (see page 887)

6.6.1.2 Use Velero with AWS

Configure Velero to use S3 buckets as storage for its backup operations.

6.6.1.2.1 Overview

Follow these procedures to set up Velero with AWS S3:

- [Velero with AWS S3 - Prepare your Environment](#) (see page 875)
- [Velero with AWS S3 - Configure Velero](#) (see page 877)
- [Velero with AWS S3 - Establish a Backup Location](#) (see page 880)

6.6.1.2.2 Velero with AWS S3 - Prepare your Environment

6.6.1.2.2.1 Prerequisites

- Ensure you have installed Velero (included in the default DKP installation).
- Ensure you have [installed the Velero CLI](#) (see page 894).
- You have [created an S3 bucket with AWS](#)⁶⁰¹.

Environment variables:

- Set the `BUCKET` environment variable to the name of the S3 bucket you want to use as backup storage:

```
export BUCKET=<aws-bucket-name>
```

- Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace. Replace `<workspace_namespace>` with the name of the target workspace:

⁶⁰⁰ <https://rook.io/>

⁶⁰¹ <https://docs.aws.amazon.com/AmazonS3/latest/userguide/creating-bucket.html>

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

i This can be the `kommander` namespace for the [Management](#) (see page 82) cluster or any other additional workspace namespace for [Attached](#) (see page 82) or [Managed](#) (see page 82) clusters. To list all available workspaces and clusters, use the `dkp get clusters -A` (see page 1712) command.

- Set the `CLUSTER_NAME` environment variable. Replace `<target_cluster>` with the name of the cluster where you want to set up Velero:

```
export CLUSTER_NAME=<target_cluster>
```

6.6.1.2.2.2 Prepare your AWS Credentials

= See the [official docs](#)⁶⁰² for details on how to use IAM roles instead of static credentials.

1. Create a file containing your static AWS credentials:

In this example, the file's name is `credentials-velero`.

```
cat << EOF > aws-credentials
[default]
aws_access_key_id=<REDACTED>
aws_secret_access_key=<REDACTED>
EOF
```

2. Create a secret on the cluster where you are installing and configuring Velero by referencing the file created in the previous step. This can be the [Management](#) (see page 82), a [Managed](#) (see page 82) or an [Attached](#) (see page 82) cluster:

In this example, the secret's name is `velero-aws-credentials`.

```
kubectl create secret generic -n ${WORKSPACE_NAMESPACE} velero-aws-credentials
--from-file=aws=aws-credentials --kubeconfig=${CLUSTER_NAME}.conf
```

6.6.1.2.2.3 Next Step:

[Velero with AWS S3 - Configure Velero](#) (see page 877)

⁶⁰² <https://github.com/vmware-tanzu/velero-plugin-for-aws>

6.6.1.2.3 Velero with AWS S3 - Configure Velero

Customize Velero to allow the configuration of a non-default backup location.

1. Create a `ConfigMap` for the Velero configuration:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: velero-overrides
data:
  values.yaml: |
    configuration:
      backupStorageLocation:
        - bucket: ${BUCKET}
          provider: "aws"
          config:
            region: <AWS_REGION> # such as us-west-2
            s3ForcePathStyle: "false"
            insecureSkipTLSVerify: "false"
            s3Url: ""
            # profile should be set to the AWS profile name mentioned in the
secret
            profile: default
          credentials:
            # With the proper IAM permissions with access to the S3 bucket,
            # you can attach the EC2 instances using the IAM Role, OR fill in
            "existingSecret" OR "secretContents" below.
            #
            # Name of a pre-existing secret (if any) in the Velero namespace
            # that should be used to get IAM account credentials.
            existingSecret: velero-aws-credentials
            # The key must be named "cloud", and the value corresponds to the entire
            content of your IAM credentials file.
            # For more information, consult the documentation for the velero plugin
for AWS at:
            # [AWS] https://github.com/vmware-tanzu/velero-plugin-for-aws/blob/main/README.md
            secretContents:
              # cloud: |
              #   [default]
              #   aws_access_key_id=<REDACTED>
              #   aws_secret_access_key=<REDACTED>
EOF
```

2. Patch the Velero `AppDeployment` to reference the created `ConfigMap` with the Velero overrides:
 - a. To update Velero in **all clusters in a workspace**:


```
cat << EOF | kubectl -n ${WORKSPACE_NAMESPACE} patch appdeployment
velero --type="merge" --patch-file=/dev/stdin
spec:
  configOverrides:
    name: velero-overrides
EOF
```

- b. To update Velero for a **specific cluster** in a workspace, see [Customize an Application per Cluster](#) (see page 660).

3. Check the `ConfigMap` on the `HelmRelease` object:

```
kubectl wait --for=jsonpath='{.spec.valuesFrom[1].name}'=velero-overrides
HelmRelease/velero -n ${WORKSPACE_NAMESPACE}
```

4. Verify that the Velero pod is running:

```
kubectl get pods -A --kubeconfig=${CLUSTER_NAME}.conf |grep velero
```

You can also configure Velero by editing the `kommander.yaml` and rerunning the installation. To follow this *alternative* configuration path, expand the following section:

Alternative Configuration path for Management/Essential clusters

6.6.1.2.3.1 Configure Velero on the Management Cluster

1. Output the Kommander configuration to `kommander.yaml`

```
dkp install kommander -o yaml --init > kommander.yaml
```

- a. Add the Velero configuration to the `apps.velero.values` section:

NOTE: This process has been tested to work with plugins for AWS v1.1.0. Newer versions of these plugins can be used, but have not been tested by D2iQ.

```
...
velero:
  values: |
    configuration:
      backupStorageLocation:
        bucket: ${BUCKET}
      config:
        region: <AWS_REGION> # such as us-west-2
        s3ForcePathStyle: "false"
```

```

    insecureSkipTLSVerify: "false"
    s3Url: ""
    # profile should be set to the AWS profile name mentioned in
the secret
    profile: default
    credentials:
    # With the proper IAM permissions with access to the S3 bucket,
    # you can attach the EC2 instances using the IAM Role, OR fill
in "existingSecret" OR "secretContents" below.
    #
    # Name of a pre-existing secret (if any) in the Velero namespace
    # that should be used to get IAM account credentials.
    existingSecret: velero-aws-credentials
    # The key must be named "cloud", and the value corresponds to
the entire content of your IAM credentials file.
    # For more information, consult the documentation for the velero
plugin for AWS at:
    # [AWS] https://github.com/vmware-tanzu/velero-plugin-for-aws/
    # blob/main/README.md
    secretContents:
    # cloud: |
    #   [default]
    #   aws_access_key_id=<REDACTED>
    #   aws_secret_access_key=<REDACTED>
    ...

```

2. Run `dkp install kommander` using the `kommander.yaml` configuration file:

```

dkp install kommander --installer-config kommander.yaml --kubeconfig=${
CLUSTER_NAME}.conf

```

3. After running `dkp install kommander` see the ConfigMap on the `HelmRelease` object:

```

kubectl wait --for=jsonpath='{.spec.valuesFrom[1].name}'=velero-overrides
HelmRelease/velero -n kommander

```

4. Check the Helm releases that the new Velero configuration has been applied:

```

kubectl get helmrelease -n ${WORKSPACE_NAMESPACE} --kubeconfig=${
CLUSTER_NAME}.conf

```

5. Verify that the Velero pod is running:

```

kubectl get pods -A --kubeconfig=${CLUSTER_NAME}.conf |grep velero

```

6. Create the backup storage location:

```
velero backup-location create -n ${WORKSPACE_NAMESPACE} <aws-backup-location-name> \
  --provider aws \
  --bucket ${BUCKET} \
  --config region=<AWS_REGION> \
  --credential=velero-aws-credentials=aws
```

7. Check that the backup storage location is `Available` and referencing the correct S3 bucket:

```
kubectl get backupstoragelocations -n kommander -oyaml
```

6.6.1.2.3.2 Next Step:

[Velero with AWS S3 - Establish a Backup Location \(see page 880\)](#)

6.6.1.2.4 Velero with AWS S3 - Establish a Backup Location

6.6.1.2.4.1 Create a Backup Storage Location

1. Create a location for the backup by pointing to an existing S3 bucket:

```
velero backup-location create -n ${WORKSPACE_NAMESPACE} <aws-backup-location-name> \
  --provider aws \
  --bucket ${BUCKET} \
  --config region=<AWS_REGION> \
  --credential=velero-aws-credentials=aws
```

2. Check that the backup storage location is `Available` and that it references the correct S3 bucket:

```
kubectl get backupstoragelocations -n ${WORKSPACE_NAMESPACE} -oyaml
```



If the BackupStorageLocation is not `Available`, view any error events by using: `kubectl describe backupstoragelocations -n ${WORKSPACE_NAMESPACE}`

6.6.1.2.4.2 Create a Test Backup

1. Create a test backup for AWS:

```
velero backup create aws-velero-testbackup -n ${WORKSPACE_NAMESPACE} --
kubeconfig=${CLUSTER_NAME}.conf --storage-location <aws-backup-location-name>
--snapshot-volumes=false
```

2. View your backup:

```
velero backup describe aws-velero-testbackup
```

Next Step:

[Back up with Velero \(see page 893\)](#)

6.6.1.3 Use Velero with Azure

Configure Velero to use Azure Blob Storage as storage for its backup operations.

6.6.1.3.1 Overview

Follow these procedures to set up Velero with Azure Blob Storage:

- [Velero with Azure Blob Containers - Prepare your Environment \(see page 881\)](#)
- [Velero with Azure Blob Containers - Configure Velero \(see page 884\)](#)
- [Velero with Azure Blob Containers - Establish a Backup Location \(see page 886\)](#)

6.6.1.3.2 Velero with Azure Blob Containers - Prepare your Environment

6.6.1.3.2.1 Prerequisites

- Ensure you have installed Velero (included in the default DKP installation).
- Ensure you have [installed the Velero CLI \(see page 894\)](#).
- Ensure you have [installed the Azure CLI](#)⁶⁰³.
- Ensure you have sufficient access rights to the Azure storage environment and blob container you want to use for backup. For more information about data authorization, see the [official Azure blob storage documentation](#)⁶⁰⁴.

⁶⁰³ <https://learn.microsoft.com/en-us/cli/azure/install-azure-cli?view=azure-cli-latest>

⁶⁰⁴ [https://learn.microsoft.com/en-us/azure/storage/common/authorize-data-access?](https://learn.microsoft.com/en-us/azure/storage/common/authorize-data-access?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&bc=%2Fazure%2Fstorage%2Fblobs%2Fbreadcrumb%2Ftoc.json)

[toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&bc=%2Fazure%2Fstorage%2Fblobs%2Fbreadcrumb%2Ftoc.json](https://learn.microsoft.com/en-us/azure/storage/common/authorize-data-access?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&bc=%2Fazure%2Fstorage%2Fblobs%2Fbreadcrumb%2Ftoc.json)

6.6.1.3.2.2 Prepare your Environment

1. [Create a container in Azure blob storage](#)⁶⁰⁵.
2. Set the `BLOB_CONTAINER` environment variable to the name of the blob container you created to use as backup storage:

```
export BLOB_CONTAINER=<Azure-blob-container-name>
```

3. Set up a [storage account and resource group](#)⁶⁰⁶.
4. Set the `AZURE_BACKUP_RESOURCE_GROUP` variable to the name of the resource group you created:

```
AZURE_BACKUP_RESOURCE_GROUP=<azure-resource-group-name>
```

5. Set the `AZURE_STORAGE_ACCOUNT_ID` variable to the unique identifier of the storage account you want to use for the backup:
 - i** See <https://learn.microsoft.com/en-us/azure/storage/common/storage-account-get-info?toc=%2Fazure%2Fstorage%2Fblobs%2Ftoc.json&bc=%2Fazure%2Fstorage%2Fblobs%2Fbreadcrumb%2Ftoc.json&tabs=azure-cli#get-the-resource-id-for-a-storage-account> to obtain the ID. The output shows the entire location path of the storage account. You only need the last part, or storage account name, to set the variable.

```
AZURE_STORAGE_ACCOUNT_ID=<storage-account-name>
```

6. Set the `AZURE_BACKUP_SUBSCRIPTION_ID` variable to the unique identifier of the subscription you want to use for the backup:
 - i** See <https://learn.microsoft.com/en-us/cli/azure/account?view=azure-cli-latest#az-account-list> to obtain the ID.

```
AZURE_BACKUP_SUBSCRIPTION_ID=<azure-subscription-id>
```

7. Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace.
 - i** Replace `<workspace_namespace>` with the name of the target workspace:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

⁶⁰⁵ <https://learn.microsoft.com/en-us/azure/storage/blobs/storage-quickstart-blobs-portal#create-a-container>

⁶⁰⁶ <https://learn.microsoft.com/en-us/azure/storage/common/storage-account-create?tabs=azure-cli#create-a-storage-account-1>

i This can be the `kommander` namespace for the [Management](#) (see page 82) cluster or any other additional workspace namespace for [Attached](#) (see page 82) or [Managed](#) (see page 82) clusters. To list all available workspace namespaces, use the `kubectl get kommandercluster -A` command.

8. Set the `CLUSTER_NAME` environment variable. Replace `<target_cluster>` with the name of the cluster where you want to set up Velero:

```
export CLUSTER_NAME=<target_cluster>
```

6.6.1.3.2.3 Prepare your Azure Credentials

i See <https://learn.microsoft.com/en-us/azure/storage/blobs/authorize-data-operations-portal> for more details on authorization.

1. Create a `credentials-velero` file with the information required to create a secret. Use the same credentials that you employed [when creating the cluster](#) (see page 1156).

i These credentials **should not** be Base64 encoded, because Velero will not read them properly.

i Replace the variables in `<...>` with your environment's information. See your Microsoft Azure account to look up the values.

```
cat << EOF > ./credentials-velero
AZURE_SUBSCRIPTION_ID=${AZURE_BACKUP_SUBSCRIPTION_ID}
AZURE_TENANT_ID=<AZURE_TENANT_ID>
AZURE_CLIENT_ID=<AZURE_CLIENT_ID>
AZURE_CLIENT_SECRET=<AZURE_CLIENT_SECRET>
AZURE_BACKUP_RESOURCE_GROUP=${AZURE_BACKUP_RESOURCE_GROUP}
AZURE_CLOUD_NAME=AzurePublicCloud
EOF
```

2. Use the `credentials-velero` file to create the secret:

```
kubectl create secret generic -n ${WORKSPACE_NAMESPACE} velero-azure-
credentials --from-file=azure=credentials-velero --kubeconfig=${
CLUSTER_NAME}.conf
```

6.6.1.3.2.4 Next Step:

[Velero with Azure Blob Containers - Configure Velero](#) (see page 884)

6.6.1.3.3 Velero with Azure Blob Containers - Configure Velero

Customize Velero to allow the configuration of a non-default backup location.

1. Create a `ConfigMap` to allow Velero to use Azure blob containers as backup storage location:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: velero-overrides
data:
  values.yaml: |
    initContainers:
      - name: velero-plugin-for-microsoft-azure
        image: velero/velero-plugin-for-microsoft-azure:v1.5.1
        imagePullPolicy: IfNotPresent
        volumeMounts:
          - mountPath: /target
            name: plugins
    credentials:
      extraSecretRef: velero-azure-credentials
EOF
```

2. Patch the Velero `AppDeployment` to reference the created `ConfigMap` with the Velero overrides:
 - a. To update Velero in **all clusters in a workspace**:

```
cat << EOF | kubectl -n ${WORKSPACE_NAMESPACE} patch appdeployment
velero --type="merge" --patch-file=/dev/stdin
spec:
  configOverrides:
    name: velero-overrides
EOF
```

- b. To update Velero for a **specific cluster** in a workspace, see [Customize an Application per Cluster \(see page 660\)](#).
3. Check the `ConfigMap` on the `HelmRelease` object:

```
kubectl get hr -n kommander velero -o jsonpath='{.spec.valuesFrom[?
(@.name=="velero-overrides")]}'
```

The output looks like this if the deployment is successful:

```
{"kind":"ConfigMap","name":"velero-overrides"}
```

4. Verify that the Velero pod is running:

```
kubectl get pods -A --kubeconfig=${CLUSTER_NAME}.conf |grep velero
```

You can also configure Velero by editing the `kommander.yaml` and rerunning the installation. To follow this *alternative* configuration path, expand the following section:

Alternative Configuration Path for Management/Essential Clusters

1. Refresh the `kommander.yaml` to add the customization of Velero:

⚠ Before running this command, ensure the `kommander.yaml` is the configuration file you are currently using for your environment. Otherwise, your previous configuration will be lost. ⚠

```
dkp install kommander -o yaml --init > kommander.yaml
```

2. Configure DKP to load the plugins and to include the secret in the `apps.velero` section:

ℹ This process has been tested to work with plugin Azure v1.5.1. More recent versions of these plugins can be used, but have not been tested by D2iQ.

```
...
  velero:
    values: |
      initContainers:
        - name: velero-plugin-for-microsoft-azure
          image: velero/velero-plugin-for-microsoft-azure:v1.5.1
          imagePullPolicy: IfNotPresent
          volumeMounts:
            - mountPath: /target
              name: plugins
      credentials:
        extraSecretRef: velero-azure-credentials
...

```

3. Use the modified `kommander.yaml` configuration to install this Velero configuration:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

4. Check the `ConfigMap` on the `HelMRelease` object:


```
kubectl get hr -n kommander velero -o jsonpath='{.spec.valuesFrom[?(@.name=="velero-overrides")]}'
```

The output looks like this if the deployment is successful:

```
{"kind":"ConfigMap","name":"velero-overrides"}
```

5. Verify that the Velero pod is running:



```
kubectl get pods -A --kubeconfig=${CLUSTER_NAME}.conf |grep velero
```

6.6.1.3.3.1 Next Step:

[Velero with Azure Blob Containers - Establish a Backup Location \(see page 886\)](#)

6.6.1.3.4 Velero with Azure Blob Containers - Establish a Backup Location

6.6.1.3.4.1 Create a Backup Storage Location

1. Create a location for the backup by pointing to an existing Azure container:
 -  Ensure you set the required environment variables as specified in [Velero with Azure Blob Containers - Prepare your Environment \(see page 884\)](#).
 -  Replace `<azure-backup-location-name>` with a name for the backup location.

```
velero backup-location create <azure-backup-location-name> -n $
{WORKSPACE_NAMESPACE} \
--provider azure \
--bucket ${BLOB_CONTAINER} \
--config resourceGroup=${AZURE_BACKUP_RESOURCE_GROUP},storageAccount=${
{AZURE_STORAGE_ACCOUNT_ID},subscriptionId=${AZURE_BACKUP_SUBSCRIPTION_ID} \
--credential=velero-azure-credentials=azure --kubeconfig=${CLUSTER_NAME}.conf
```

2. Check that the backup storage location is `Available` and that it references the correct Azure container:

```
kubectl get backupstoragelocations -n ${WORKSPACE_NAMESPACE} -oyaml
```

6.6.1.3.4.2 Create a Test Backup

1. Create a test backup that is stored in the location you created in the previous section:

```
velero backup create azure-velero-testbackup -n ${WORKSPACE_NAMESPACE} \
--kubeconfig=${CLUSTER_NAME}.conf \
--storage-location <azure-backup-location-name> \
--snapshot-volumes=false
```

2. View your backup:

```
velero backup describe azure-velero-testbackup
```

Troubleshooting

1. If your backup wasn't created, Velero may have had an issue installing the plugin. If the plugin was not installed, run this command:

```
velero plugin add velero/velero-plugin-for-microsoft-azure:v1.5.1 -n $
{WORKSPACE_NAMESPACE}
```

2. Confirm your `backupstoragelocation` was configured correctly

```
kubectl get backupstoragelocations -n ${WORKSPACE_NAMESPACE}
```

If your backup storage location is "Available", proceed creating a test backup.

NAME	PHASE	LAST VALIDATED	AGE	DEFAULT
<azure-backup-location-name>		Available	38s	60m

6.6.1.3.4.3 Next Step:

[Back up with Velero \(see page 893\)](#)

6.6.1.4 Use Velero with GCP

Configure Velero to use Google Cloud Storage Buckets as storage for its backup operations.

6.6.1.4.1 Overview

Follow these procedures to set up Velero with Azure Blob Storage:

- [Velero with Google Cloud Storage Buckets - Prepare your Environment](#) (see page 888)
- [Velero with Google Cloud Storage Buckets - Configure Velero](#) (see page 889)
- [Velero with Google Cloud Storage Buckets - Establish a Backup Location](#) (see page 892)

6.6.1.4.2 Velero with Google Cloud Storage Buckets - Prepare your Environment

6.6.1.4.2.1 Prerequisites

- Ensure you have installed Velero (included in the default DKP installation).
- Ensure you have [installed the Velero CLI](#) (see page 894).
- Ensure you have [installed the gcloud CLI](#)⁶⁰⁷.
- **Optional:** You can [install the gsutil CLI](#)⁶⁰⁸ (or opt to create buckets through the GCS Console)
- Ensure you have created a [GCS bucket](#)⁶⁰⁹.
- Ensure you have sufficient access rights to the bucket you want to use for backup. For more information about GCP-related access control, refer to the official documentation of the [Google Cloud Storage platform](#)⁶¹⁰.

6.6.1.4.2.2 Set Environment Variables

- Set the `BUCKET` environment variable to the name of the GCS container you want to use as backup storage:

```
export BUCKET=<GCS-bucket-name>
```

- Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace. Replace `<workspace_namespace>` with the name of the target workspace:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

- **i** This can be the `kommander` namespace for the [Management](#) (see page 82) cluster or any other additional workspace namespace for [Attached](#) (see page 82) or [Managed](#) (see page 82) clusters. To list all available workspace namespaces, use the `kubectl get kommandercluster -A` command.

607 <https://cloud.google.com/sdk/docs/install>

608 https://cloud.google.com/storage/docs/gsutil_install

609 <https://cloud.google.com/storage/docs/creating-buckets>

610 <https://cloud.google.com/storage/docs/access-control>

- Set the `CLUSTER_NAME` environment variable. Replace `<target_cluster>` with the name of the cluster where you want to set up Velero:

```
export CLUSTER_NAME=<target_cluster>
```

6.6.1.4.2.3 Prepare your Google Credentials

You can store your backups in **Google Cloud Platform/GCS buckets**.

- See <https://cloud.google.com/storage/docs/creating-buckets#required-roles> for more information on setting up access to your bucket.

1. Create a `credentials-velero` file with the information required to create a secret.
 - Replace `<service-account-email>` with the email address you used to grant permissions to your bucket. The address usually follows the format `<service-account-user>@<gcp-project>.iam.gserviceaccount.com`.

```
gcloud iam service-accounts keys create credentials-velero \
  --iam-account <service-account-email>
```

2. Use the `credentials-velero` file to create the secret:

```
kubectl create secret generic -n ${WORKSPACE_NAMESPACE} velero-gcp-credentials
--from-file=gcp=credentials-velero --kubeconfig=${CLUSTER_NAME}.conf
```

6.6.1.4.2.4 Next Step:

[Velero with Google Cloud Storage Buckets - Configure Velero \(see page 889\)](#)

6.6.1.4.3 Velero with Google Cloud Storage Buckets - Configure Velero

Customize Velero to allow the configuration of a non-default backup location.

1. Create a `ConfigMap` to allow Velero to use GCS buckets as backup storage location:

```
cat <<EOF | kubectl apply -f -
```

```

apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: velero-overrides
data:
  values.yaml: |
    initContainers:
      - name: velero-plugin-for-gcp
        image: velero/velero-plugin-for-gcp:v1.5.0
        imagePullPolicy: IfNotPresent
        volumeMounts:
          - mountPath: /target
            name: plugins
    credentials:
      extraSecretRef: velero-gcp-credentials
EOF

```

2. Patch the Velero `AppDeployment` to reference the created `ConfigMap` with the Velero overrides:
 - a. To update Velero in **all clusters in a workspace**:

```

cat << EOF | kubectl -n ${WORKSPACE_NAMESPACE} patch appdeployment
velero --type="merge" --patch-file=/dev/stdin
spec:
  configOverrides:
    name: velero-overrides
EOF

```

- b. To update Velero for a **specific cluster** in a workspace, see [Customize an Application per Cluster](#) (see page 660).
3. Check the `ConfigMap` on the `HelmRelease` object:

```

kubectl get hr -n kommander velero -o jsonpath='{.spec.valuesFrom[?
(@.name=="velero-overrides")]}'

```

The output looks like this if the deployment is successful:

```

{"kind":"ConfigMap","name":"velero-overrides"}

```

4. Ensure that the Velero pod is running:

```

kubectl get pods -A --kubeconfig=${CLUSTER_NAME}.conf |grep velero

```

You can also configure Velero by editing the `kommander.yaml` and rerunning the installation. To follow this *alternative* configuration path, expand the following section:

Alternative Configuration Path for Management/Essential Clusters

6.6.1.4.3.1 Configure Velero on the Management Cluster

1. Refresh the `kommander.yaml` to add the customization of Velero:

⚠ Before running this command, ensure the `kommander.yaml` is the configuration file you are currently using for your environment. Otherwise, your previous configuration will be lost. ⚠

```
dkp install kommander -o yaml --init > kommander.yaml
```

2. Configure DKP to load the plugins and to include the secret in the `apps.velero` section:

ℹ This process has been tested to work with plugin GCP v1.5.0. More recent versions of these plugins can be used, but have not been tested by D2iQ.

```
...
  velero:
    values: |
      initContainers:
        - name: velero-plugin-for-gcp
          image: velero/velero-plugin-for-gcp:v1.5.0
          imagePullPolicy: IfNotPresent
          volumeMounts:
            - mountPath: /target
              name: plugins
      credentials:
        extraSecretRef: velero-gcp-credentials
...

```

3. Use the modified `kommander.yaml` configuration to install this Velero configuration:

```
dkp install kommander --installer-config kommander.yaml --kubecfg=${CLUSTER_NAME}.conf
```

4. Check the `ConfigMap` on the `HelmRelease` object:

```
kubectl get hr -n kommander velero -o jsonpath='{.spec.valuesFrom[?(@.name=="velero-overrides")]}'
```

The output looks like this if the deployment is successful:

```
{"kind":"ConfigMap","name":"velero-overrides"}
```

5. Verify that the Velero pod is running:

```
kubectl get pods -A --kubeconfig=${CLUSTER_NAME}.conf |grep velero
```

6.6.1.4.3.2 Next Step:

[Velero with Google Cloud Storage Buckets - Establish a Backup Location](#) (see page 892)

6.6.1.4.4 Velero with Google Cloud Storage Buckets - Establish a Backup Location

6.6.1.4.4.1 Create a Backup Storage Location

1. Create a location for the backup by pointing to an existing GCS bucket:
 - i** Ensure you set the required environment variables as specified in [Velero with Google Cloud Storage Buckets - Prepare your Environment](#) (see page 891).
 - i** Replace `<gcp-backup-location-name>` with a name for the backup location.

```
velero backup-location create <gcp-backup-location-name> -n $
{WORKSPACE_NAMESPACE} \
--provider gcp \
--bucket $BUCKET \
--credential=velero-gcp-credentials=gcp
```

2. Check that the backup storage location is `Available` and that it references the correct GCS bucket:

```
kubectl get backupstoragelocations -n ${WORKSPACE_NAMESPACE} -oyaml
```

6.6.1.4.4.2 Create a Test Backup

1. Create a test backup that is stored in the location you created in the previous section:

```
velero backup create gcp-velero-testbackup -n ${WORKSPACE_NAMESPACE} \
--kubeconfig=${CLUSTER_NAME}.conf \
--storage-location <gcp-backup-location-name> \
```

```
--snapshot-volumes=false
```

2. View your backup:

```
velero backup describe gcp-velero-testbackup
```

Troubleshooting

1. If your backup wasn't created, Velero may have had an issue installing the plugin. If the plugin was not installed, run this command:

```
velero plugin add velero/velero-plugin-for-gcp:v1.5.0 -n ${WORKSPACE_NAMESPACE}
```

2. Confirm your `backupstoragelocation` was configured correctly

```
kubectl get backupstoragelocations -n ${WORKSPACE_NAMESPACE}
```

If your backup storage location is "Available", proceed creating a test backup.

NAME	PHASE	LAST VALIDATED	AGE	DEFAULT
<gcp-backup-location-name>		Available	38s	60m

6.6.1.4.4.3 Next Step:

[Back up with Velero \(see page 893\)](#)

6.6.2 Back up with Velero

DKP provides [Velero](#)⁶¹¹ by default, to support backup and restore operations for your Kubernetes clusters and persistent volumes.

- [Install the Velero CLI \(see page 894\)](#)
- [Regular Backup Operations \(see page 894\)](#)
- [Restore a cluster from backup \(see page 897\)](#)
- [Backup Service Diagnostics \(see page 899\)](#)

⁶¹¹ <https://velero.io/>

6.6.2.1 Install the Velero CLI

Although installing the Velero command-line interface is optional and independent of deploying a DKP cluster, having access to the command-line interface provides several benefits. For example, you can use the Velero command-line interface to back up or restore a cluster on demand, or to modify certain settings without changing the Velero configuration.

- By default, DKP sets up Velero to use Rook Ceph over TLS using a self-signed certificate.
- As a result, when using certain commands, you may be asked to use the `--insecure-skip-tls-verify` flag.

Again, the default setup is not suitable for production use-cases.

See the instructions to [install the Velero command-line interface](#)⁶¹² for more information.

In DKP, the Velero platform application is installed in the `kommander` namespace, instead of `velero`.

Thus, after installing the CLI, we recommend that you set the Velero CLI `namespace` config option so that subsequent Velero CLI invocations will use the correct namespace:

```
velero client config set namespace=kommander
```

6.6.2.2 Regular Backup Operations

For backing up production clusters, you should be familiar with the following basic administrative functions Velero provides:

- [Set a Backup Schedule](#) (see page 0)
- [Run On-demand Backups](#) (see page 0)
- [Restore from a Backup archive](#) (see page 897)



- If you want to back up your cluster in the scope of [converting a DKP Essential cluster to a DKP Enterprise Managed cluster](#) (see page 830), see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273351033/2.6+Back+up+your+Cluster+s+Applications+and+Persistent+Volumes> (see page 835).
- If you require a custom backup location, see how to create one for [AWS](#) (see page 883), [Azure](#) (see page 889), or [GCP](#) (see page 895).

⁶¹² <https://velero.io/docs/v1.5/basic-install/#install-the-cli>

6.6.2.2.1 Prepare your Environment

Before you modify a schedule or create an on-demand backup, set the following environment variables:

1. Specify the workspace namespace of the cluster for which you want to configure the backup:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

2. Specify the cluster for which you want to create the backup:

```
export CLUSTER_NAME=<target_cluster_name>
```

6.6.2.2.2 Set a Backup Schedule

By default, DKP configures a regular, automatic backup of the cluster's state in Velero. The default settings do the following:

- Create daily backups
- Save the data from all namespaces



DKP default backups do not support the creation of Volume Snapshots.

These default settings take effect after the cluster is created. If you install DKP with the default platform services deployed, the initial backup starts after the cluster is successfully provisioned and ready for use.

6.6.2.2.2.1 Create Backup Schedules

The Velero CLI provides an easy way to create alternate backup schedules. For example:

```
velero create schedule <backup-schedule-name> -n ${WORKSPACE_NAMESPACE} \
--kubeconfig=${CLUSTER_NAME}.conf \
--snapshot-volumes=false \
--schedule="@every 8h"
```

6.6.2.2.2.2 Change Default Backup Service Settings

1. Check the backup schedules currently configured for the cluster:

```
velero get schedules
```

2. Delete the `velero-default` schedule:

```
velero delete schedule velero-default
```

3. Replace the default schedule with your custom settings:

```
velero create schedule velero-default -n ${WORKSPACE_NAMESPACE} \
--kubeconfig=${CLUSTER_NAME}.conf \
--snapshot-volumes=false \
--schedule="@every 24h"
```

6.6.2.2.3 Create Backup Schedule for a Specific Namespace

You can also create backup schedules for specific namespaces. Creating a backup for a specific namespace can be useful for clusters running multiple apps operated by multiple teams. For example:

```
velero create schedule <backup-schedule-name> \
--include-namespaces=kube-system,kube-public,kommander \
--snapshot-volumes=false \
--schedule="@every 24h"
```

The Velero command line interface provides many more options worth exploring. You can also find tutorials for [disaster recovery](#)⁶¹³ and [cluster migration](#)⁶¹⁴ on the Velero community site.

6.6.2.2.3 Back up on Demand

In some cases, you might find it necessary to create a backup outside the regularly-scheduled interval. For example, if you are preparing to upgrade any components or modify your cluster configuration, perform a backup before taking that action.

Create a backup by running the following command:

```
velero backup create <backup-name> -n ${WORKSPACE_NAMESPACE} \
--kubeconfig=${CLUSTER_NAME}.conf \
--snapshot-volumes=false
```

⁶¹³ <https://velero.io/docs/v0.11.0/disaster-case>

⁶¹⁴ <https://velero.io/docs/v0.11.0/migration-case>

6.6.2.3 Restore a cluster from backup

6.6.2.3.1 Prerequisites

Before attempting to restore the cluster state using the Velero command-line interface, verify the following requirements:

- The backend storage, Rook Ceph Cluster, is still operational.
- The Velero platform service in the cluster is still operational.
- The Velero platform service is set to a `restore-only-mode` to avoid having backups run while restoring.


6.6.2.3.2 Restoring from Backup

To list the available backup archives for your cluster, run the following command:

```
velero backup get
```

To set Velero to a `restore-only-mode`, run the following command:

```
velero server --restore-only=true
```

 This mode is being deprecated and will be removed in Velero in Velero v2.0. Use read-only backup storage locations instead.

Finally, check your deployment to verify that the configuration change was applied correctly:

```
helm get values -n kommander velero
```

To restore cluster data on demand from a selected backup snapshot available in the cluster, run a command similar to the following:

```
velero restore create --from-backup <BACKUP-NAME>
```

- Important:** If you are restoring using Velero from the default setup (and not using an [external bucket or blob to store your backups](#) (see page 878)), you may see an error when describing or viewing the logs of your backup restore. This is a known issue when restoring from an object store that is not accessible from outside your cluster. However, you can review the success of the backup restore by confirming the Phase is `Completed` and not in error, and viewing the logs by running these `kubectl` commands:

```
kubectl logs -l name=velero -n kommander --tail -1
```

6.6.2.3.3 Restore your DKP Management Cluster

- When restoring a backup to the Management Cluster, you must adjust configuration to avoid restore errors.

6.6.2.3.3.1 Prior to restoring a backup

1. Ensure the following `ResourceQuota` setup is not configured on your cluster (this `ResourceQuota` will be automatically restored)

```
kubectl -n kommander delete resourcequota one-kommandercluster-per-kommander-workspace
```

2. Turn off the `Workspace` validation webhooks. Otherwise, you will not restore Workspaces with pre-configured namespaces. If the validation webhook named `kommander-validating` is present, it must be modified with this command:

```
kubectl patch validatingwebhookconfigurations kommander-validating \
  --type json \
  --patch '[
    {
      "op": "remove",
      "path": "/webhooks/0/rules/3/operations/0"
    }
  ]'
```

6.6.2.3.3.2 After restoring a backup

1. Verify that the `ResourceQuota` named `one-kommandercluster-per-kommander-workspace` is restored
2. Add the removed `CREATE` webhook rule operation with:

```
kubectl patch validatingwebhookconfigurations kommander-validating \
  --type json \
  --patch '[
    {
      "op": "add",
      "path": "/webhooks/0/rules/3/operations/0",
      "value": "CREATE"
    }
  ]'
```

6.6.2.4 Backup Service Diagnostics

You can check whether the Velero service is currently running on your cluster through the Kubernetes dashboard (accessible through the DKP UI on the Management Cluster), or by running the following `kubectl` command:

```
kubectl get all -A | grep velero
```

If the Velero platform service application is currently running, you can generate diagnostic information about Velero backup and restore operations. For example, you can run the following commands to retrieve, back up, and restore information that you can use to assess the overall health of Velero in your cluster:

```
velero get schedules
velero get backups
velero get restores
velero get backup-locations
velero get snapshot-locations
```

6.7 Logging

D2iQ Kubernetes Platform (DKP) comes with a pre-configured logging stack that allows you to collect and visualize pod and admin log data at the Workspace level. The logging stack is also multi-tenant capable. Multi-tenancy is enabled at the Project level through role-based access control (RBAC).

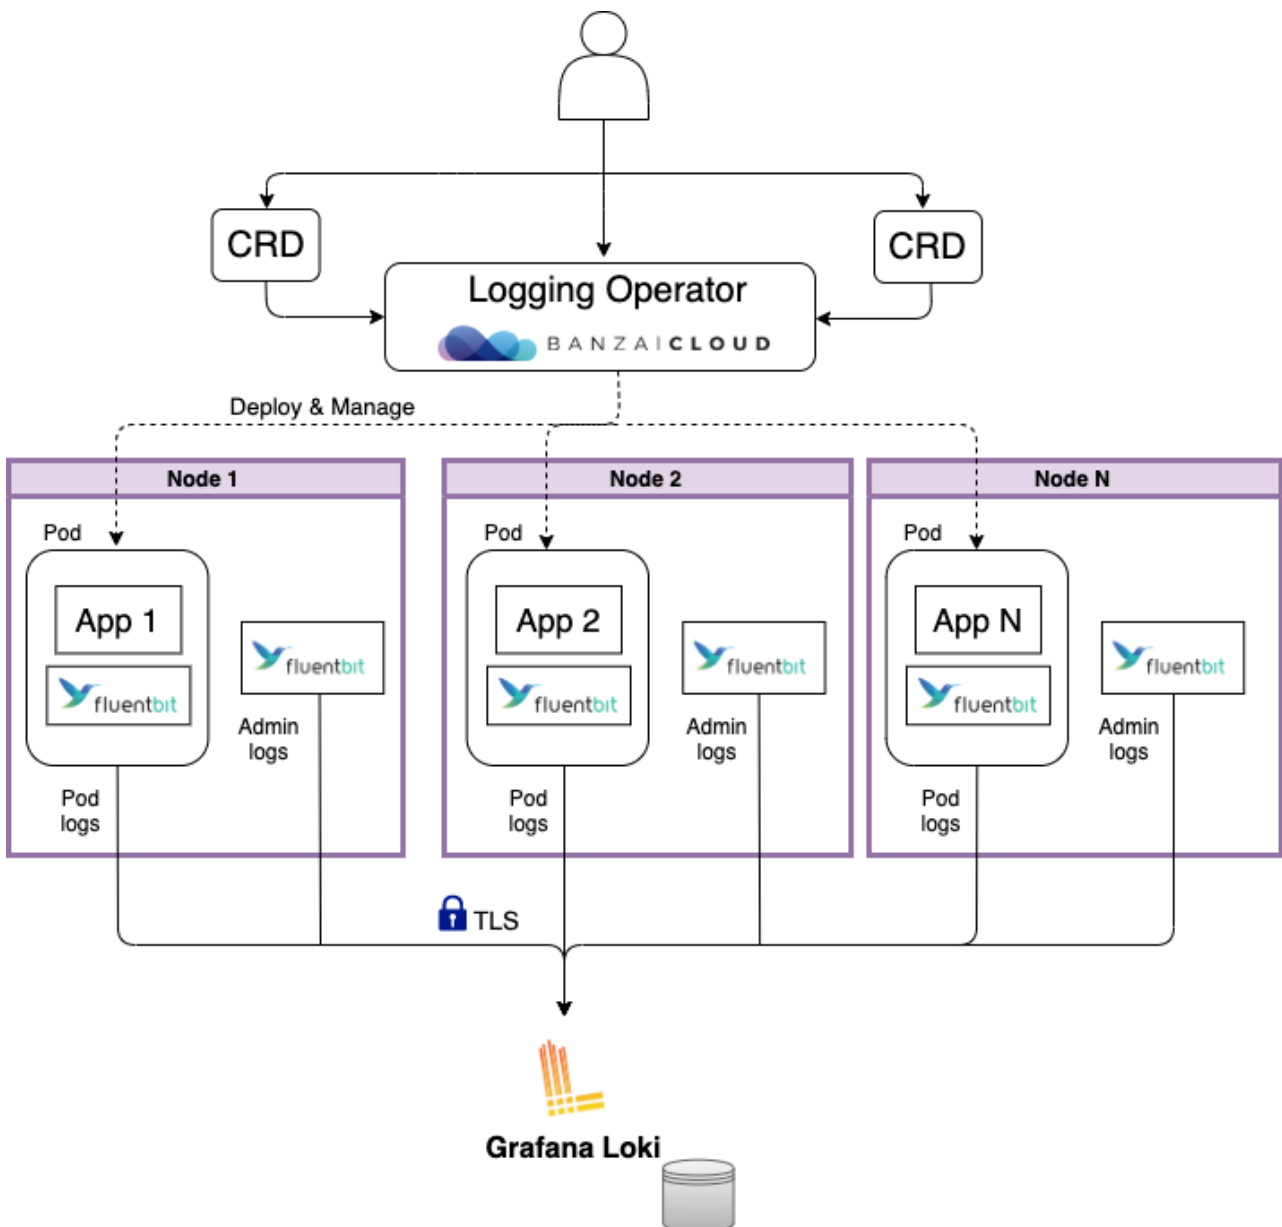
By default, logging is disabled on managed and attached clusters. You need to enable the logging stack applications explicitly on the workspace to make use of these capabilities.

The primary components of the logging stack include these platform services:

- BanzaiCloud Logging-operator
- Grafana and Grafana Loki
- Fluent Bit and Fluentd

In addition to these platform services, logging relies on other software and system facilities, including the container runtime, the journald facility, and systemd configuration are used to collect logs and messages from all the machines in the cluster.

The following diagram illustrates how different components of the logging stack collect log data and provide information about clusters:



The DKP logging stack aggregates logs from applications and nodes running inside your cluster.

DKP uses the BanzaiCloud Logging-operator to manage the Fluent Bit and Fluentd deployments that collect pod logs, using Kubernetes API extensions called custom resources. The custom resources allow users to declare logging configurations using `kubectl` commands. The Fluent Bit instance deployed by Logging-operator gathers pod logs data and sends it to Fluentd, which forwards it to the appropriate Grafana Loki servers based on the configuration defined in custom resources.

Loki then indexes the log data by label and stores it for querying. Loki maintains log order integrity but does not index the log messages themselves, which improves its efficiency and lowers its footprint.

6.7.1 Logging Operator

This section contains information about setting up the logging operator to manage the Fluent Bit and Fluentd resources.

6.7.1.1 Fluent Bit Buffer Information

Fluent Bit collects container logs from the host filesystem and performs the following:

- Maintains a small buffer of logs (5 MB in memory)
- Maintains a checkpoint on each host for each file it's consumed, so if it's restarted, it can resumed where it left off.
- Flushes the buffer every one second.
- Has a five second grace period to flush the buffer on exit (along with a 30 second termination grace period on the pod).

Every pod restart should result in the buffer being flushed to `fluentd` (allows up to five seconds for that flush to happen).

If the buffer is not fully flushed during that five seconds, a small amount of log data may be dropped, but if `fluentd` is functional, it is unlikely that fluent-bit would be unable to flush its logs on pod termination.

Fluent Bit can be configured to use a `hostPath` volume to store the buffer information, so it can be picked up again when Fluent Bit restarts.

For more information about fluent-bit, see <https://kube-logging.dev/docs/logging-infrastructure/fluentbit/#hostpath-volumes-for-buffers-and-positions>.

For more information about Logging in relation to how it is used in DKP, refer to these pages in our Help Center:

- [Admin-level logs](#) (see page 902)
- [Enable Workspace-level Logging](#) (see page 902)
- [Multi-Tenant Logging Overview](#) (see page 913)
- [Fluent Bit](#) (see page 923)
- [Scaling the Logging Stack](#) (see page 929)
- [Configuring Loki to use AWS S3 Storage in DKP](#) (see page 932)

6.7.2 Admin-level logs

DKP also includes a Fluentbit instance to collect admin-level log information which is sent to the workspace Grafana Loki that's running on the cluster. The admin log information includes:

- Logs for host processes managed by systemd
- Kernel logs
- Kubernetes audit logs

This approach helps to isolate the more sensitive logs from Logging-operator, eliminating the possibility that users might gain inadvertent access to that data.

See [Fluent Bit \(see page 923\)](#) for more information about these logs.



On the Management cluster, the Fluentbit application is disabled by default. The amount of admin logs ingested to Loki requires additional disk space to be configured on the `rook-ceph-cluster`. Enabling admin logs may use around 2GB/day per node. See [Rook Ceph in DKP \(see page 1004\)](#) for more details on how to configure the Ceph Cluster.

6.7.3 Enable Workspace-level Logging

6.7.3.1 How to enable Workspace-level Logging for use with DKP.

Logging is disabled by default on managed and attached clusters. You will need to enable logging features explicitly at the Workspace level, if you want to capture and view log data.



You must perform these procedures to enable multi-tenant logging at the Project level as well.

- [Logging Prerequisites and Architecture \(see page 903\)](#)
- [Enable Logging Applications through the UI \(see page 904\)](#)
- [Create AppDeployments to Enable Workspace Logging \(see page 905\)](#)
- [Override ConfigMap to Restrict Logging to Specific Namespaces \(see page 907\)](#)
- [Override ConfigMap to Modify the Storage Retention Period in Workspace Grafana Loki \(see page 909\)](#)
- [Verify Cluster Logging Stack Installation \(see page 910\)](#)
- [View Cluster Log Data \(see page 911\)](#)

6.7.3.2 Logging Prerequisites and Architecture

Before you begin, you must:

- Be a cluster administrator with permissions to configure cluster-level platform services.
- Set a [default storage class](#) (see [page 1395](#)) on each attached cluster for successful Loki deployment.

For more information, see [Default Storage Providers in DKP](#) (see [page 91](#)) .

6.7.3.2.1 DKP Logging Stack Architecture

The DKP logging stack architecture provides a comprehensive logging solution for the DKP platform. It combines Fluent Bit, Fluentd, Loki, and Grafana components to collect, process, store, and visualize log data. The architecture establishes a robust logging solution by assigning specific roles to each of those components.

6.7.3.2.1.1 Components

1. **Fluent Bit** - Fluent Bit is a lightweight log processor and forwarder that collects log data from various sources, such as application logs or Kubernetes components. It forwards the collected logs to Fluentd for further processing.
2. **Fluentd** - Fluentd is a powerful and flexible log aggregator that receives log data from Fluent Bit, processes it, and forwards it to the Loki Distributor. Fluentd can handle various log formats and enrich the log data with additional metadata before forwarding it.
3. **Loki** - Loki is a horizontally-scalable, highly-available, multi-tenant log aggregation system. Loki components include:-
 - Compactor: Responsible for compacting index files and chunks to improve query performance and reduce storage usage.
 - Distributor: Receives log streams, partitions them into chunks, and forwards these chunks to the Loki Ingestor component.
 - Gateway: Acts as a single access point to various Loki components, routing requests between Distributor, Query Frontend, and other components as needed.
 - Ingestor: Compresses, indexes, and persists received log chunks.
 - Querier: Fetches log chunks from the Ingestor, decompresses and filters them based on the query, and returns the results to the Query Frontend.
 - Query Frontend: Splits incoming queries into smaller parts, forwards these to the Loki Querier component, and combines the results from all Queriers before returning the final result.
4. **Grafana** - Grafana is a visualization and analytics platform that supports Loki as one of its data sources. Grafana provides a user-friendly interface for querying and visualizing log data based on user-defined dashboards and panels.

6.7.3.2.1.2 Workflow

- **Write Path:**

- Fluent Bit instances running on each node collect log data from various sources, like application logs or Kubernetes components.
- Fluent Bit forwards the collected log data to the Fluentd instance.
- Fluentd processes the received log data and forwards it to the Loki Distributor through the Loki Gateway.
- The Loki Distributor receives the log streams, partitions them into chunks, and forwards these chunks to the Loki Ingestor component.
- Loki Ingestors are responsible for compressing, indexing, and persisting the received log chunks.
- **Read Path:**
 - When a user queries logs through Grafana, the request goes to the Loki Gateway, which routes it to the Loki Query Frontend.
 - The Query Frontend splits the query into smaller parts and forwards these to the Loki Querier component.
 - Loki Queriers fetch the log chunks from the Loki Ingestor, decompress and filter them based on the query, and return the results to the Query Frontend.
 - The Query Frontend combines the results from all Queriers and returns the final result to Grafana through the Loki Gateway.
 - Grafana visualizes the log data based on the user's dashboard configuration

6.7.3.3 Enable Logging Applications through the UI


6.7.3.3.1 How to enable the logging stack through the UI for Workspace-level logging

You can enable the Workspace logging stack to all attached clusters within the Workspace through the UI. If you prefer to enable the logging stack with `kubectl`, review how you [create AppDeployments to Enable Workspace Logging \(see page 905\)](#).

To enable workspace-level logging in DKP using the UI, follow these steps:

1. From the top menu bar, select your target workspace.
2. Select **Applications** from the sidebar menu.
3. Ensure `traefik` and `cert-manager` are enabled on your cluster. These are deployed by default unless you modified your configuration.
4. Scroll to the **Logging** applications section.
5. Select the three dot button from the bottom-right corner of the cards for Rook Ceph and Rook Ceph Cluster, then click **Enable**. On the **Enable Workspace Platform Application** page, you can add a customized configuration for settings that best fit your organization. You can leave the configuration settings unchanged to enable with default settings.
6. Select **Enable** at the top right of the page.

7. Repeat the process for the Grafana Loki, Logging Operator, and Grafana Logging applications.
8. You can verify the cluster logging stack installation by waiting until the cards have a Deployed checkmark on the [Cluster Application](#) (see page 874) page, or you can [verify the Cluster Logging Stack installation via the CLI](#) (see page 910).
9. Then, you can [view cluster log data](#) (see page 911).

 We do not recommend installing Fluent Bit, which is responsible for collecting admin logs, unless you have configured the Grafana Loki Ceph Cluster Bucket with sufficient storage space. The amount of admin logs ingested to Loki requires additional disk space to be configured on the `rook-ceph-cluster`. Enabling admin logs may use around 2GB/day per node. See [Rook Ceph in DKP](#) (see page 1004) for more details on how to configure the Ceph Cluster.

6.7.3.4 Create AppDeployments to Enable Workspace Logging

How to create AppDeployments to enable Workspace-level logging

Workspace logging AppDeployments enable and deploy the logging stack to all attached clusters within the workspace. Use the DKP UI to enable the logging applications, or, alternately, use the CLI to create the AppDeployments.

To enable logging in DKP using the CLI, follow these steps on the **management** cluster:

1. Execute the following command to get the name and namespace of your workspace:

```
dkp get workspaces
```

And copy the values under the `NAME` and `NAMESPACE` columns for your workspace.

2. Export the `WORKSPACE_NAME` variable:

```
export WORKSPACE_NAME=<WORKSPACE_NAME>
```

3. Export the `WORKSPACE_NAMESPACE` variable:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

4. Ensure that Cert-Manager and Traefik are enabled in the workspace. If you want to find if the applications are enabled on the management cluster workspace, you can run:

```
dkp get appdeployments --workspace ${WORKSPACE_NAME}
```

5. You can confirm that the applications are deployed on the **managed** or **attached** cluster by running this `kubectl` command in that cluster. (Ensure you switch to the correct [context or kubeconfig](#)⁶¹⁵ of the attached cluster for the following `kubectl` command.)

```
kubectl get helmreleases -n ${WORKSPACE_NAMESPACE}
```

6. Copy these commands and execute them on the **management** cluster from a command line to create the Logging-operator, Grafana-loki, and Grafana-logging AppDeployments:

```
dkp create appdeployment logging-operator --app logging-operator-4.2.3 --
workspace ${WORKSPACE_NAME}
dkp create appdeployment rook-ceph --app rook-ceph-1.11.9 --workspace $
${WORKSPACE_NAME}
dkp create appdeployment rook-ceph-cluster --app rook-ceph-cluster-1.11.9 --
workspace ${WORKSPACE_NAME}
dkp create appdeployment grafana-loki --app grafana-loki-0.69.16 --workspace $
${WORKSPACE_NAME}
dkp create appdeployment grafana-logging --app grafana-logging-6.57.4 --
workspace ${WORKSPACE_NAME}
```

Then, you can [verify the cluster logging stack installation](#) (see page 910).



To deploy the applications to selected clusters within the workspace, refer to the [cluster-scoped configuration](#) (see page 657) section.



We do not recommend installing Fluent Bit, which is responsible for collecting admin logs, unless you have configured the Rook Ceph Cluster with sufficient storage space. Enabling admin logs via Fluent Bit may use around 2GB/day per node. See [Rook Ceph in DKP](#) (see page 1004) for more details on how to configure the Rook Ceph Cluster.

To install Fluent Bit, create the AppDeployment:

```
dkp create appdeployment fluent-bit --app fluent-bit-0.34.1 --workspace $
${WORKSPACE_NAME}
```

⁶¹⁵ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

6.7.3.5 Override ConfigMap to Restrict Logging to Specific Namespaces

 How to override the logging configMap to restrict logging to specific namespaces.

As a cluster administrator, you may have a need to limit, or restrict, logging activities only to certain namespaces. Kommander allows you to do this by creating an override configMap that modifies the logging configuration created in the [Create AppDeployment for Workspace Logging \(see page 908\)](#) procedure.

6.7.3.5.1 Prerequisites

- Implement each of the steps listed in [Enable Workspace-level Logging \(see page 905\)](#).
- Ensure that log data is available before you execute this procedure.

6.7.3.5.2 Create and use the Override Entries

To create and use the override configMap entries, follow these steps:

1. Execute the following command to get the namespace of your workspace:

```
dkp get workspaces
```

And copy the value under the `NAMESPACE` column for your workspace.

2. Set the `WORKSPACE_NAMESPACE` variable to the namespace copied in the previous step:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

3. Identify one or more namespaces to which you want to restrict logging.
4. Create a file named `logging-operator-logging-overrides.yaml` and paste the following YAML code into it to create the overrides configMap:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: logging-operator-logging-overrides
  namespace: ${WORKSPACE_NAMESPACE}
data:
  values.yaml: |
    ---
    clusterFlows:
```

```

- name: cluster-containers
  spec:
    globalOutputRefs:
      - loki
    match:
      - exclude:
          namespaces:
            - <your-namespace>
            - <your-other-namespace>

```

5. Add the relevant namespace values for `metadata.namespace` and the `clusterFlows[0].spec.match[0].exclude.namespaces` values at the end of the file, and save the file.
6. Use the following command to apply the YAML file:

```
kubectl apply -f logging-operator-logging-overrides.yaml
```

7. Edit the `logging-operator` `AppDeployment` to set the value of `spec.configOverrides.name` to `logging-operator-logging-overrides`.
(Refer to [Deploy an application with a custom configuration](#) (see page 0) for more information)

```
dkp edit appdeployment -n ${WORKSPACE_NAMESPACE} logging-operator
```

After your editing is complete, the `AppDeployment` resembles this example:

```


apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: logging-operator
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: logging-operator-3.17.7
    kind: ClusterApp
  configOverrides:
    name: logging-operator-logging-overrides

```

8. Perform actions that generate log data, both in the specified namespaces *and* the namespaces you mean to exclude.
9. Verify that the log data contains only the data you expected to receive.

6.7.3.6 Override ConfigMap to Modify the Storage Retention Period in Workspace Grafana Loki

DKP configures Grafana Loki to retain log metadata and logs for 1 week, using the [Compactor](#)⁶¹⁶. This retention policy can be customized.

 The minimum retention period is 24h.

To customize the retention policy using `configOverrides`, run these commands on the **management** cluster:

1. Execute the following command to get the namespace of your workspace:

```
dkp get workspaces
```

Copy the value under the `NAMESPACE` column for your workspace.

2. Set the `WORKSPACE_NAMESPACE` variable to the namespace copied in the previous step:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

3. Create a `ConfigMap` with custom configuration values for Grafana Loki. Since the retention configuration is nested in a `config` string, you must copy the entire block. The following example sets the retention period to 360 hours (15 days). Refer to [Grafana Loki's Retention Configuration documentation](#)⁶¹⁷ for more information about this field.

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: grafana-loki-custom-overrides
  namespace: ${WORKSPACE_NAMESPACE}
data:
  values.yaml: |
    loki:
      structuredConfig:
        limits_config:
          retention_period: 360h
```

⁶¹⁶ <https://grafana.com/docs/loki/latest/operations/storage/boltdb-shipper/#compactor>

⁶¹⁷ <https://grafana.com/docs/loki/latest/operations/storage/retention/#retention-configuration>


```
EOF
```

4. Edit the `grafana-loki` AppDeployment to set the value of `spec.configOverrides.name` to `grafana-loki-custom-overrides` (Refer to [Deploy a service with a custom configuration](#) (see [page 0](#)) for more information).

```
dkp edit appdeployment -n ${WORKSPACE_NAMESPACE} grafana-loki
```

After your editing is complete, the AppDeployment resembles this example:

```
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: grafana-loki
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: grafana-loki-0.69.10
    kind: ClusterApp
  configOverrides:
    name: grafana-loki-custom-overrides
```

6.7.3.7 Verify Cluster Logging Stack Installation

 How to verify the cluster's logging stack installed successfully.

You must wait for the cluster's logging stack `HelmsReleases` to deploy before attempting to configure or use the logging features.

Run the following commands on the **management** cluster:

1. Execute the following command to get the namespace of your workspace:

```
dkp get workspaces
```

And copy the value under the `NAMESPACE` column for your workspace.

2. Set the `WORKSPACE_NAMESPACE` variable to the namespace copied in the previous step:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

Ensure you switched to the correct [context or kubeconfig](#)⁶¹⁸ of the attached cluster for the following kubectl commands.

3. Check the deployment status using this command on the **attached** cluster:

```
kubectl get helmreleases -n ${WORKSPACE_NAMESPACE}
```

NOTE: It may take some time for these changes to take effect, based on the duration configured for the Flux GitRepository reconciliation.


When the logging stack is successfully deployed, you will see output that includes the following

`HelmReleases` :

NAME	READY	STATUS	AGE
grafana-logging	True	Release reconciliation succeeded	15m
logging-operator	True	Release reconciliation succeeded	15m
logging-operator-logging	True	Release reconciliation succeeded	15m
grafana-loki	True	Release reconciliation succeeded	15m
rook-ceph	True	Release reconciliation succeeded	15m
rook-ceph-cluster	True	Release reconciliation succeeded	15m
object-bucket-claims	True	Release reconciliation succeeded	15m

Then, you can [View Cluster Log Data](#) (see page 911).

6.7.3.8 View Cluster Log Data

 How to view the cluster's log data after enabling logging.

Though you enable logging at the Workspace level, viewing the log data is done at the cluster level, using the cluster's Grafana logging URL.

Run the following commands on the **management** cluster:

1. Execute the following command to get the namespace of your workspace:

```
dkp get workspaces
```

And copy the value under the `NAMESPACE` column for your workspace.

2. Set the `WORKSPACE_NAMESPACE` variable to the namespace copied in the previous step:

⁶¹⁸ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

Run the following commands on the **attached** cluster to access the Grafana UI:

Ensure you switched to the correct [context or kubeconfig](#)⁶¹⁹ of the attached cluster for the following kubectl commands:

3. Get the Grafana URL:

```
kubectl get ingress -n ${WORKSPACE_NAMESPACE} grafana-logging -o go-
template='https://{{with index .status.loadBalancer.ingress 0}}
{{or .hostname .ip}}{{end}}{{with index .spec.rules 0}}{{with index .http.paths
0}}{{.path }}{{end}}{{end}}{\n\''
```

To view logs in Grafana:

1. Go to the Explore tab:


```
kubectl get ingress -n ${WORKSPACE_NAMESPACE} grafana-logging -o go-
template='https://{{with index .status.loadBalancer.ingress 0}}
{{or .hostname .ip}}{{end}}{{with index .spec.rules 0}}{{with index .http.paths
0}}{{.path }}{{end}}{{end}}/explore{\n\''
```

2. You may be prompted to log in using the SSO flow. See [Authentication and Authorization](#) (see page 934) for more information.
3. At the top of the page, change the datasource to `Loki` .

See the [Grafana Loki documentation](#)⁶²⁰ for more on how to use the interface to view and query logs.

⁶¹⁹ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

⁶²⁰ <https://grafana.com/docs/grafana/v7.5/datasources/loki/>



Cert-Manager and Traefik must be deployed in the attached cluster to be able to access the Grafana UI. These are deployed by default on the workspace.

6.7.4 Multi-Tenant Logging Overview

Enterprise

Gov Advanced

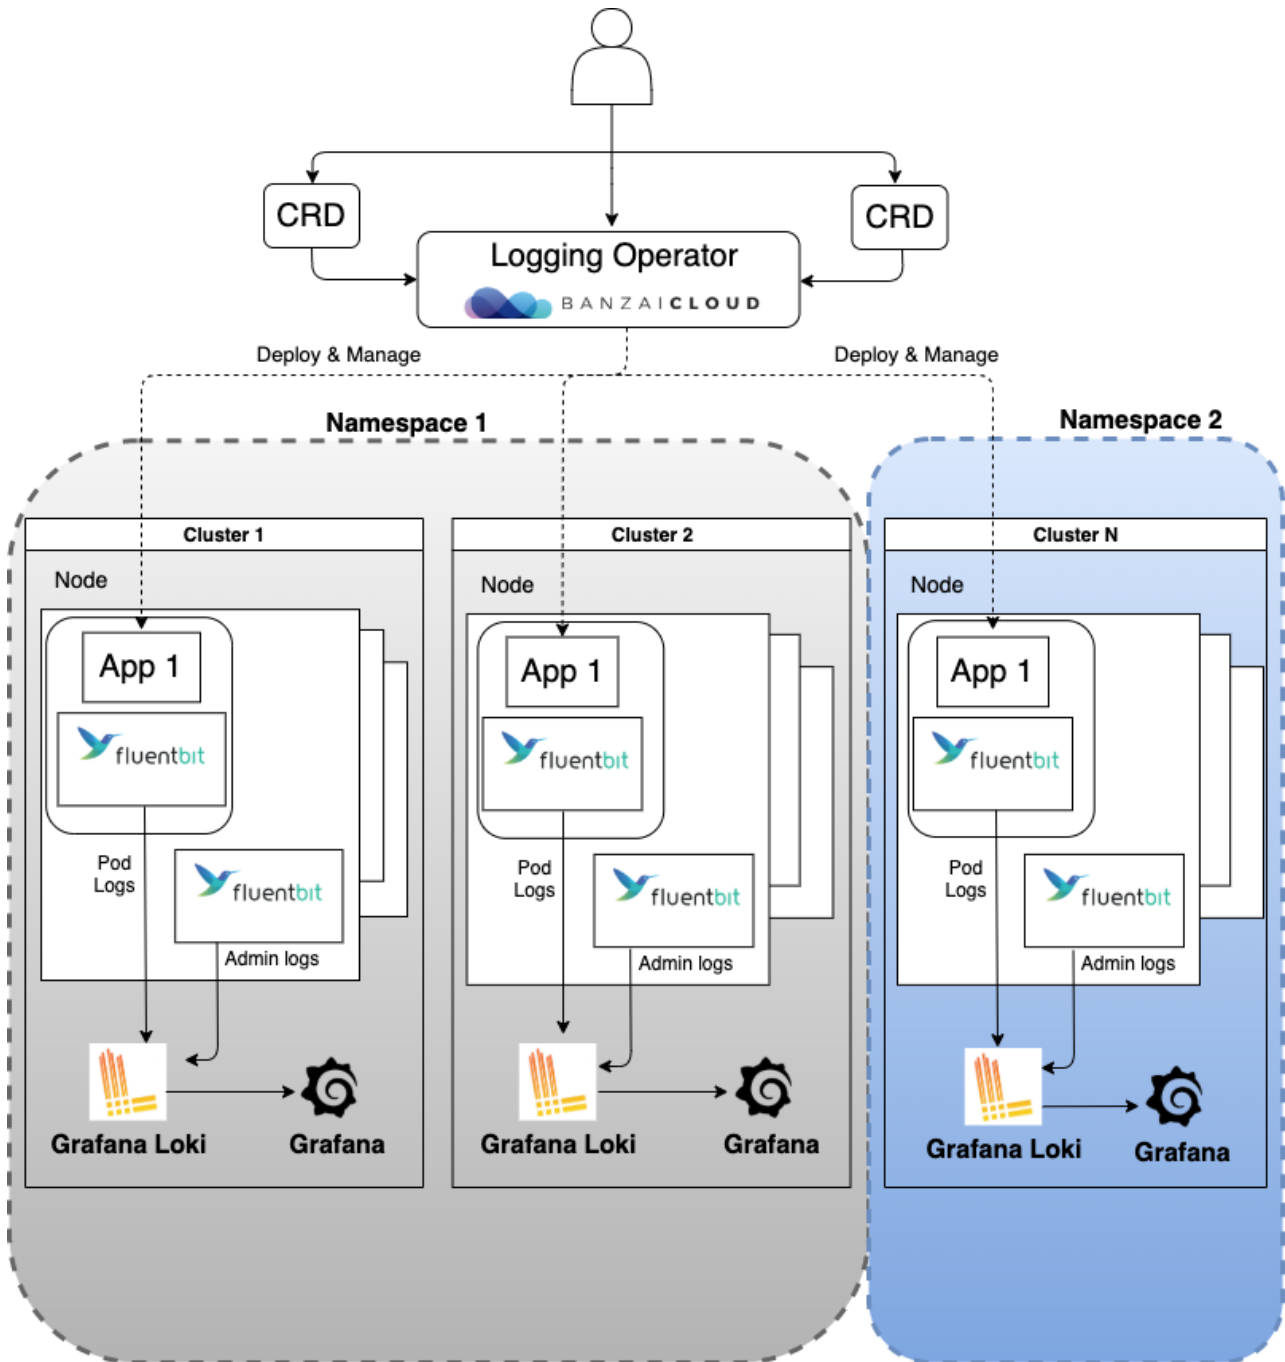
Multi-tenant logging in Kommander is built on the DKP Logging architecture. Multi-tenant logging uses the same components as the base logging architecture:

- BanzaiCloud logging-operator
- Grafana Loki
- Grafana



You must perform the [Workspace-level Logging](#) (see page 905) procedures as a prerequisite to enable multi-tenant logging at the Project level as well.

Access to log data is done at the namespace level, through the use of Projects within Kommander as shown in the diagram:



Each Project namespace has a logging-operator “Flow” that sends its pod logs to its own Loki server. A custom controller deploys corresponding Loki and Grafana servers in each namespace, and defines a logging-operator Flow in each namespace that forwards its pod logs to its respective Loki server. There is a corresponding Grafana server for visualizations for each namespace.

For the convenience of cluster Administrators, a cluster-scoped Loki/Grafana instance pair is deployed with a corresponding Logging-operator `ClusterFlow` that directs pod logs from all namespaces to the pair. A cluster Administrator can grant access either to none of the logs, or to all logs collected from all pods in a given namespace. Assigning teams to specific namespaces enables the team members to see only the logs for the namespaces they own.

As with any endpoint, if an Ingress controller is in use in the environment, take care that the ingress rules do not supersede the RBAC permissions and thus prevent access to the logs.



Cluster Administrators will need to monitor and adjust resource usage to prevent operational difficulties or excessive use on a *per namespace* basis.

- [Enable Multi-tenant Logging](#) (see page 915)
- [Create a Project for Logging](#) (see page 916)
- [Create Project-level Logging AppDeployments](#) (see page 916)
- [Override ConfigMap to Modify the Storage Retention Period in Project Grafana Loki](#) (see page 917)
- [Verify Project Logging Stack Installation](#) (see page 919)
- [View Project Log Data](#) (see page 920)

6.7.4.1 Enable Multi-tenant Logging

Enterprise

Gov Advanced

6.7.4.1.1 Prerequisites

Before you begin, you must:

- [Enable Workspace-level Logging](#) (see page 905) before you can configure multi-tenant logging.
- Be a cluster administrator with permissions to configure cluster level platform services.

6.7.4.1.2 Multi-tenant Logging Enablement Process

The steps required to enable multi-tenant logging include:

1. [Create a Project.](#) (see page 916)
2. [Create the required Project-level AppDeployments.](#) (see page 916)
3. [Verify that the Project logging stack is installed.](#) (see page 919)
4. [View a Project's log data.](#) (see page 920)

Get started with multi-tenant logging by [Creating a Project](#) (see page 916).

6.7.4.2 Create a Project for Logging

Enterprise

Gov Advanced

To enable multi-tenant logging, you must first [Create a Project](#) (see page 689) and its namespace. Users assigned to this namespace will be able to access log data for only that namespace and not others.

Then, you can [create project-level AppDeployments for use in multi-tenant logging](#) (see page 916).

6.7.4.3 Create Project-level Logging AppDeployments

Enterprise

Gov Advanced

How to create Project-level AppDeployments for use in multi-tenant logging.

You must create AppDeployments in the Project namespace to enable and deploy the logging stack to all clusters within a Project. You can use the CLI to do this, or use the DKP UI to enable the logging applications.

To create the AppDeployments needed for Project-level logging, follow these steps **on the management cluster**:

1. Determine the name and namespace of the workspace that your project is in. You can use the `dkp get workspaces` command to see the list of workspace names and their corresponding namespaces.

```
dkp get workspaces
```

Copy the values under the `NAME` and `NAMESPACE` columns for your workspace.

2. Export the `WORKSPACE_NAME` variable:

```
export WORKSPACE_NAME=<WORKSPACE_NAME>
```

3. Export the `WORKSPACE_NAMESPACE` variable:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

4. Execute the following command to get the namespace of your project:

```
kubectl get projects -n ${WORKSPACE_NAMESPACE}
```

Copy the value under the `NAME` column for your project. This may NOT be identical to the Display Name of the `Project`.

5. Export the `PROJECT_NAME` variable:

```
export PROJECT_NAME=<PROJECT_NAME>
```

6. Copy these commands and execute them from a command line:

```
dkp create appdeployment project-grafana-loki --app project-grafana-loki-0.69.1
6 --workspace ${WORKSPACE_NAME} --project ${PROJECT_NAME}
dkp create appdeployment project-grafana-logging --app project-grafana-
logging-6.57.4 --workspace ${WORKSPACE_NAME} --project ${PROJECT_NAME}
dkp create appdeployment project-logging --app project-logging-1.0.3 --
workspace ${WORKSPACE_NAME} --project ${PROJECT_NAME}
```

After completing these steps, you can [Verify the Project Logging Stack Installation](#) (see page 919) for multi-tenant logging.

6.7.4.4 Override ConfigMap to Modify the Storage Retention Period in Project Grafana Loki

Enterprise

Gov Advanced

DKP configures Project Grafana Loki to retain log metadata and logs for 1 week, using the [Compactor](#)⁶²¹. This retention policy can be customized.

 The minimum retention period is 24h.

To customize the retention policy using `configOverrides`, run these commands on the **management** cluster:

1. Determine the namespace of the workspace that your project is in. You can use the `dkp get workspaces` command to see the list of workspace names and their corresponding namespaces.

```
dkp get workspaces
```

⁶²¹ <https://grafana.com/docs/loki/latest/operations/storage/boltdb-shipper/#compactor>

Copy the value under the `NAMESPACE` column for your workspace. This may NOT be identical to the Display Name of the `Workspace`.

2. Set the `WORKSPACE_NAMESPACE` variable to the namespace copied in the previous step:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

3. Get the namespace of your project:

```
kubectl get projects --namespace ${WORKSPACE_NAMESPACE}
```

Copy the value under the `PROJECT_NAMESPACE` column for your workspace. This may NOT be identical to the Display Name of the `Project`.

4. Set the `PROJECT_NAMESPACE` variable to the namespace copied in the previous step:

```
export PROJECT_NAMESPACE=<PROJECT_NAMESPACE>
```

5. Create a `ConfigMap` with custom configuration values for Grafana Loki. Since the retention configuration is nested in a `config` string, you must copy the entire block. The following example sets the retention period under to 360 hours (15 days). Refer to [Grafana Loki's Retention Configuration documentation](#)⁶²² for more information about this field.

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: project-grafana-loki-custom-overrides
  namespace: ${PROJECT_NAMESPACE}
data:
  values.yaml: |
    loki:
      structuredConfig:
        limits_config:
          retention_period: 360h
EOF
```

6. Run the following command on the **management** cluster to reference the `configOverrides` in the `project-grafana-loki` AppDeployment:

```
cat <<EOF | kubectl apply -f -
```

⁶²² <https://grafana.com/docs/loki/latest/operations/storage/retention/#retention-configuration>

```

apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: project-grafana-loki
  namespace: ${PROJECT_NAMESPACE}
spec:
  appRef:
    name: project-grafana-loki-0.48.6
    kind: ClusterApp
  configOverrides:
    name: project-grafana-loki-custom-overrides
EOF

```

6.7.4.5 Verify Project Logging Stack Installation

Enterprise

Gov Advanced

How to verify the project logging stack installation for multi-tenant logging

You must wait for the Project's logging stack `HelmsReleases` to deploy before configuring or using the Project-level logging features, including multi-tenancy:

Run the following commands on the **management** cluster:

1. Determine the namespace of the workspace that your project is in. You can use the `dkp get workspaces` command to see the list of workspace names and their corresponding namespaces.

```
dkp get workspaces
```

Copy the value under the `NAMESPACE` column for your workspace.

2. Export the `WORKSPACE_NAMESPACE` variable:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

3. Execute the following command to get the namespace of your project

```
kubectl get projects -n ${WORKSPACE_NAMESPACE}
```

Copy the value under `PROJECT_NAMESPACE` column for your project. This may NOT be identical to the Display Name of the `Project`.

4. Export the `PROJECT_NAMESPACE` variable:

```
export PROJECT_NAMESPACE=<PROJECT_NAMESPACE>
```

Run the following commands on the **managed** or **attached** cluster. For this, ensure you switch to the correct [context or kubeconfig](#)⁶²³ of the cluster for the following `kubectl` commands:

5. Check the deployment status using this command on the attached cluster:

```
kubectl get helmreleases -n ${PROJECT_NAMESPACE}
```

NOTE: It may take some time for these changes to take effect, based on the duration configured for the Flux GitRepository reconciliation.

When successfully deployed, you will see output that includes the following `HelmReleases` :

NAMESPACE	NAME	READY	STATUS
AGE			
\${PROJECT_NAMESPACE}	project-grafana-logging	True	Release
reconciliation succeeded	15m		
\${PROJECT_NAMESPACE}	project-grafana-loki	True	Release
reconciliation succeeded	11m		
\${PROJECT_NAMESPACE}	project-loki-object-bucket-claims	True	Release
reconciliation succeeded	11m		

Then, you can [View Project Log Data](#) (see page 920) within multi-tenant logging.

6.7.4.6 View Project Log Data

Enterprise

Gov Advanced

How to view project log data within multi-tenant logging

You can only view the log data for a Project to which you have been granted access.

6.7.4.6.1 Access Project Grafana's UI

Run the following commands on the **management** cluster:

1. Determine the namespace of the workspace that your project is in. You can use the `dkp get workspaces` command to see the list of workspace names and their corresponding namespaces.

⁶²³ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

```
dkp get workspaces
```

Copy the value under the `NAMESPACE` column for your workspace.

2. Export the `WORKSPACE_NAMESPACE` variable:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

3. Execute the following command to get the namespace of your project

```
kubectl get projects -n ${WORKSPACE_NAMESPACE}
```

Copy the value under `PROJECT_NAMESPACE` column for your project. This may NOT be identical to the Display Name of the `Project`.

4. Export the `PROJECT_NAMESPACE` variable:

```
export PROJECT_NAMESPACE=<PROJECT_NAMESPACE>
```

Run the following commands **on the attached cluster** to access the Project Grafana's UI:

Ensure you switched to the correct [context or kubeconfig](#)⁶²⁴ of the attached cluster for the following `kubectl` commands:

5. Get the Grafana URL:

```
kubectl get ingress -n ${PROJECT_NAMESPACE} ${PROJECT_NAMESPACE}-project-
grafana-logging -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}{{with
index .spec.rules 0}}{{with index .http.paths 0}}{{.path }}{{end}}{{end}}
{{"\n"}}'
```

6.7.4.6.2 View Logs in Grafana

1. Go to the `Explore` tab:

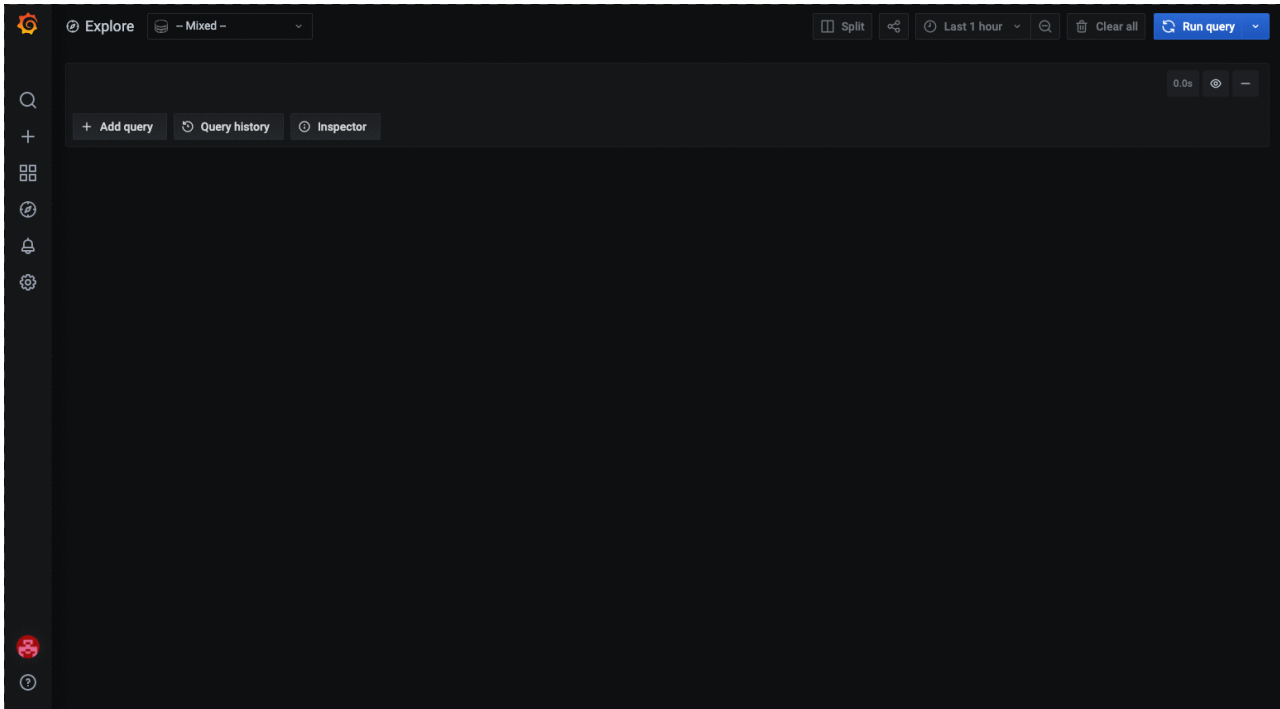
```
kubectl get ingress -n ${PROJECT_NAMESPACE} ${PROJECT_NAMESPACE}-project-
grafana-logging -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}{{with
```

⁶²⁴ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

```
index .spec.rules 0}}>{{with index .http.paths 0}}{.path }}{{end}}}}/
explore{{"\n"}}'
```

2. You may be prompted to log in using the SSO flow. See [Authentication and Authorization](#) (see page 934) for more information.
3. At the top of the page, change the data source to `Loki`.


See the [Grafana Loki documentation](#)⁶²⁵ for more on how to use the interface to view and query logs.



⚠ Cert-Manager and Traefik must be deployed in the attached cluster to be able to access the Grafana UI. These are deployed by default on the workspace.

You can [configure workspace policy](#) (see page 902) to restrict access to the Project logging Grafana UI. Each Grafana instance in a Project has a unique URL at the cluster level. Consider creating a `WorkspaceRoleBinding` that maps to a `ClusterRoleBinding`, on attached cluster(s), for each Project level Grafana instance. For example, If you have a group named `sample-group` and two projects named `first-project` and `second-project` in `sample-workspace` workspace, then the Role Bindings look similar to the following:

⁶²⁵ <https://grafana.com/docs/grafana/v7.5/datasources/loki/>



Select the correct role bindings for each group for a project at the workspace level.

6.7.5 Fluent Bit

Fluent Bit⁶²⁶ is the DKP choice of open-source log collection and forwarding tool.



On the Management cluster, the Fluentbit application is disabled by default. The amount of admin logs ingested to Loki requires additional disk space to be configured on the `rook-ceph-cluster`. Enabling admin logs may use around 2GB/day per node. See [Rook Ceph in DKP \(see page 1004\)](#) for more details on how to configure the Rook Ceph Cluster.

6.7.5.1 Audit Log Collection

Auditing⁶²⁷ in Kubernetes provides a way to chronologically document the actions taken on a cluster. On Kommander, by default, audit logs are collected and stored for quick indexing. Viewing and accessing can be done via the Grafana logging UI.

To adjust the default Audit Policy [log backend](#)⁶²⁸ configuration, you must modify the log retention settings by [Configuring the Control Plane](#)⁶²⁹ before creating the cluster. This needs to be done prior to creating the cluster since it cannot be edited after creation.

- [Collecting systemd Logs from a Non-default Path \(see page 924\)](#)

626 <https://fluentbit.io/>

627 <https://kubernetes.io/docs/tasks/debug-application-cluster/audit/>

628 <https://kubernetes.io/docs/tasks/debug/debug-cluster/audit/#log-backend>

629 [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320304&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane)

[fromPageId=273320304&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273320304&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configure+the+Control+Plane)

6.7.5.2 Related Information

For information on related topics or procedures, refer to the following:

- [Logging \(see page 902\)](#)

6.7.5.3 Collecting systemd Logs from a Non-default Path

By default, Fluent Bit pods are configured to collect `systemd` logs from the `/var/log/journal/` path on cluster nodes.

If `systemd-journald` running as a part of the OS on the nodes uses a different path for writing logs, you will need to override configuration of the `fluent-bit` AppDeployment to make Fluent Bit collect `systemd` logs.

To configure the Fluent Bit AppDeployment to collect `systemd` logs from a non-default path, follow these steps (all `kubectl` and `dkp` invocations refer to the **management** cluster):

1. Execute the following command to get the namespace of the workspace in which you would like to configure Fluent Bit:

```
dkp get workspaces
```

And copy the value under the `NAMESPACE` column for your workspace.

2. Set the `WORKSPACE_NAMESPACE` variable to the namespace copied in the previous step:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

3. Identify the `systemd-journald` log data storage path on the nodes of the clusters in the workspace by using the OS documentation and examining the `systemd` configuration.

Usually it will be either `/var/log/journal` (typically used when `systemd-journald` is configured to store logs permanently; in this case the default Fluent Bit configuration should work) or `/run/log/journal` (typically used when `systemd-journald` is configured to use a volatile storage).

4. Extract the default Helm values used by the Fluent Bit App:

```
kubectl get -n ${WORKSPACE_NAMESPACE} configmaps fluent-bit-0.20.9-d2iq-defaults -o=jsonpath='{.data.values\.yaml}' > fluent-bit-values.yaml
```

5. Edit the resulting file `fluent-bit-values.yaml` by removing all sections except for `extraVolumes`, `extraVolumeMounts` and `config.inputs`. The result should look similarly to this:

```
extraVolumes:
# we create this to have a persistent tail-db directory an all nodes
# otherwise a restarted fluent-bit would rescrape all tails
- name: tail-db
  hostPath:
    path: /var/log/tail-db
    type: DirectoryOrCreate
# we create this to get rid of error messages that would appear on non control-
plane nodes
- name: kubernetes-audit
  hostPath:
    path: /var/log/kubernetes/audit
    type: DirectoryOrCreate
# needed for kmsg input plugin
- name: uptime
  hostPath:
    path: /proc/uptime
    type: File
- name: kmsg
  hostPath:
    path: /dev/kmsg
    type: CharDevice

extraVolumeMounts:
- name: tail-db
  mountPath: /tail-db
- name: kubernetes-audit
  mountPath: /var/log/kubernetes/audit
- name: uptime
  mountPath: /proc/uptime
- name: kmsg
  mountPath: /dev/kmsg

config:
  inputs: |
    # Collect audit logs, systemd logs, and kernel logs.
    # Pod logs are collected by the fluent-bit deployment managed by logging-
operator.
    [INPUT]
      Name tail
      Alias kubernetes_audit
      Path /var/log/kubernetes/audit/*.log
      Parser kubernetes-audit
      DB /tail-db/audit.db
      Tag audit.*
      Refresh_Interval 10
```



```

Rotate_Wait 5
Mem_Buf_Limit 135MB
Buffer_Chunk_Size 5MB
Buffer_Max_Size 20MB
Skip_Long_Lines Off
[INPUT]
Name systemd
Alias kubernetes_host
DB /tail-db/journal.db
Tag host.*
Max_Entries 1000
Read_From_Tail On
Strip_Underscores On
[INPUT]
Name kmsg
Alias kubernetes_host_kernel
Tag kernel

```

6. Add the following item to the list under the `extraVolumes` key:

```

- name: kubernetes-host
  hostPath:
    path: <path to systemd logs on the node>
    type: Directory

```

7. Add the following item to the list under the `extraVolumeMounts` key:

```

- name: kubernetes-host
  mountPath: <path to systemd logs on the node>

```

These items will make Kubernetes mount systemd logs into Fluent Bit pods.

8. Add the following line into the `[INPUT]` entry identified by `Name systemd` and `Alias kubernetes_host`.

```

Path <path to systemd logs on the node>

```

This is needed to make Fluent Bit actually collect the mounted logs

9. Assuming that the path to systemd logs on the node is `/run/log/journal`, the result will look similarly to this:

```

extraVolumes:
# we create this to have a persistent tail-db directory an all nodes
# otherwise a restarted fluent-bit would rescrape all tails
- name: tail-db

```

```

    hostPath:
      path: /var/log/tail-db
      type: DirectoryOrCreate
# we create this to get rid of error messages that would appear on non control-
plane nodes
- name: kubernetes-audit
  hostPath:
    path: /var/log/kubernetes/audit
    type: DirectoryOrCreate
# needed for kmsg input plugin
- name: uptime
  hostPath:
    path: /proc/uptime
    type: File
- name: kmsg
  hostPath:
    path: /dev/kmsg
    type: CharDevice
- name: kubernetes-host
  hostPath:
    path: /run/log/journal
    type: Directory

extraVolumeMounts:
- name: tail-db
  mountPath: /tail-db
- name: kubernetes-audit
  mountPath: /var/log/kubernetes/audit
- name: uptime
  mountPath: /proc/uptime
- name: kmsg
  mountPath: /dev/kmsg
- name: kubernetes-host
  mountPath: /run/log/journal

config:
  inputs: |
    # Collect audit logs, systemd logs, and kernel logs.
    # Pod logs are collected by the fluent-bit deployment managed by logging-
operator.
    [INPUT]
      Name tail
      Alias kubernetes_audit
      Path /var/log/kubernetes/audit/*.log
      Parser kubernetes-audit
      DB /tail-db/audit.db
      Tag audit.*
      Refresh_Interval 10
      Rotate_Wait 5
      Mem_Buf_Limit 135MB
      Buffer_Chunk_Size 5MB
      Buffer_Max_Size 20MB

```

```

    Skip_Long_Lines Off
  [INPUT]
    Name systemd
    Alias kubernetes_host
    Path /run/log/journal
    DB /tail-db/journal.db
    Tag host.*
    Max_Entries 1000
    Read_From_Tail On
    Strip_Underscores On
  [INPUT]
    Name kmsg
    Alias kubernetes_host_kernel
    Tag kernel

```

10. Create a `ConfigMap` manifest with override values from `fluent-bit-values.yaml`:

```

cat <<EOF >fluent-bit-overrides.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: fluent-bit-overrides
data:
  values.yaml: |
$(cat fluent-bit-values.yaml | sed 's/^/  /g')
EOF

```

11. Create a `ConfigMap` from the manifest above:

```
kubectl apply -f fluent-bit-overrides.yaml
```

12. Edit the `fluent-bit` `AppDeployment` to set the value of `spec.configOverrides.name` to the name of the created `ConfigMap`. (You can use the steps in the procedure, [Deploy an Application with a Custom Configuration](#) (see page 927) as a guide.)

```
dkp edit appdeployment -n ${WORKSPACE_NAMESPACE} fluent-bit
```

After your editing is complete, the `AppDeployment` resembles this example:

```

apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: fluent-bit
  namespace: ${WORKSPACE_NAMESPACE}

```

```
spec:
  appRef:
    name: fluent-bit-0.20.9
    kind: ClusterApp
  configOverrides:
    name: fluent-bit-overrides
```

- Log in into the Grafana logging UI of your workspace and verify that logs with a label `log_source=kubernetes_host` are now present in Loki.

6.7.6 Scaling the Logging Stack

6.7.6.1 Introduction

Depending on the application workloads you run on your clusters, you may find that the default settings for the DKP logging stack do not meet your needs. In particular, if your workloads produce lots of log traffic, you may find you need to adjust the logging stack components to properly capture all the log traffic. Follow the suggestions below to tune the logging stack components as needed.

6.7.6.2 Logging Operator

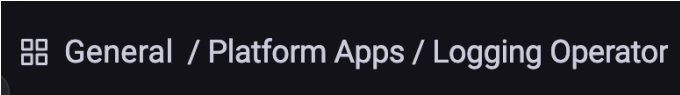
In a high log traffic environment, `fluentd` usually becomes the bottleneck of the logging stack. According to <https://banzaicloud.com/docs/one-eye/logging-operator/operation/scaling/>:

The typical sign of this is when `fluentd` cannot handle its `buffer`⁶³⁰ directory size growth for more than the configured or calculated (`timekey + timekey_wait`) flush interval.

For metrics to monitor, refer to <https://docs.fluentd.org/monitoring-fluentd/monitoring-prometheus#metrics-to-monitor>.

6.7.6.2.1 Grafana dashboard

In DKP, if the `Prometheus Monitoring` (`kube-prometheus-stack`) platform application is enabled, you can view the Logging Operator dashboard in the Grafana UI.



☰ General / Platform Apps / Logging Operator

You can also improve `fluentd` throughput by disabling the buffering for `Loki` `clusterOutput`.

6.7.6.2.2 Example Configuration

You can see an example configuration of the logging operator in [Logging Stack Application Sizing Recommendations](#) (see page 628).

⁶³⁰ <https://banzaicloud.com/docs/one-eye/logging-operator/configuration/plugins/outputs/buffer/>

For more information, refer to:

- <https://docs.fluentd.org/deployment/performance-tuning-single-process>

6.7.6.3 Grafana Loki

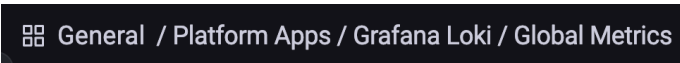
DKP deploys Loki in [Microservice mode](#)⁶³¹ – this provides you with the highest flexibility in terms of scaling.

In a high log traffic environment, we recommend:

- Ingestor should be the first component to be considered for scaling up.
- Distributor should be scaled up only when the existing Distributor is experiencing stress due to high computing resource usage.
 - Usually, the number of Distributor pods should be much lower than the number of Ingestor pods

6.7.6.3.1 Grafana dashboard

In DKP, if [Prometheus Monitoring](#) (kube-prometheus-stack) platform app is enabled, you can view the Loki dashboards in Grafana UI and here is one of the Loki dashboard:



☰ General / Platform Apps / Grafana Loki / Global Metrics

6.7.6.3.2 Example Configuration

You can see an example config of Loki at [Logging Stack Application Sizing Recommendations](#) (see page 628).

For more information, refer to:

- <https://grafana.com/docs/loki/latest/fundamentals/architecture/components/>
- <https://grafana.com/docs/loki/latest/operations/scalability/>
- <https://grafana.com/docs/loki/latest/best-practices/>

6.7.6.4 Rook Ceph

Ceph is the default S3 storage provider. In DKP, a Rook Ceph Operator and a Rook Ceph Cluster are deployed together to have a Ceph Cluster.

6.7.6.4.1 Storage

The default configuration of Rook Ceph Cluster in DKP has a 33% overhead in data storage for redundancy. Meaning, if the data disks allocated for your Rook Ceph Cluster is 1000Gb, 750Gb will be used to store your data. Thus, it is important to account for that in planning the capacity of your data disks to prevent issues.

⁶³¹ <https://grafana.com/docs/loki/latest/fundamentals/architecture/deployment-modes/#microservices-mode>

6.7.6.4.1.1 ObjectBucketClaim storage limit

ObjectBucketClaim has a storage limit option to prevent your S3 bucket from growing over a limit. In DKP this is enabled by default.

Thus, after you size up your Rook Ceph Cluster for more storage, it is important to also increase the storage limit of your ObjectBucketClaims of your `grafana-loki` and/or `project-grafana-loki`.

To change it for `grafana-loki`, provide an override configmap in `rook-ceph-cluster` platform app to override `dkp.grafana-loki.maxSize`

To change it for `project-grafana-loki`, provide an override configmap in `project-grafana-loki` platform app to override `dkp.project-grafana-loki.maxSize`

6.7.6.4.2 Example Configuration

You can see an example config at [Rook Ceph Cluster Sizing Recommendations](#) (see page 632).

6.7.6.4.3 Ceph OSD CPU considerations

`ceph-osd` is the object storage daemon for the Ceph distributed file system. It is responsible for storing objects on a local file system and providing access to them over the network.

If you determine that the Ceph OSD component is the bottleneck, then you may wish to consider increasing the CPU allocated to it.

See this page for more info: <https://ceph.io/en/news/blog/2022/ceph-osd-cpu-scaling/>

6.7.6.4.4 Grafana dashboard

In DKP, if the `Prometheus Monitoring` (`kube-prometheus-stack`) platform app is enabled, you can view the Ceph dashboards in the Grafana UI. Below is one of the Ceph dashboard:



General / Platform Apps / Rook Ceph Cluster

6.7.6.5 Audit Log

6.7.6.5.1 Overhead

Enabling audit logging requires additional computing and storage resources.

When you enable audit logging by enabling the `kommander-fluent-bit` AppDeployment, inbound traffic to the logging stack increases the log traffic by approximately 3-4 more times.

Thus, when enabling the audit log, consider scaling up all components in the logging stack mentioned above.

6.7.6.5.2 Fine-tuning audit log Fluent Bit

If you are certain that you only need to collect a subset of the logs that the default config makes the `kommander-fluent-bit` pods collect, you can add your own override configmap to `kommander-fluent-bit` with proper Fluent Bit `INPUT`, `FILTER`, `OUTPUT` settings. This helps reduce the audit log traffic.

To see the default configuration of Fluent Bit, see the [Release Notes \(see page 1752\)](#) > **Components and Applications**.

For more information, refer to:

- [Admin-level logs \(see page 905\)](#)
- <https://docs.fluentbit.io/manual/administration/configuring-fluent-bit/classic-mode/configuration-file>

6.7.7 Configuring Loki to use AWS S3 Storage in DKP

Follow the instructions on this page to configure Loki to use AWS S3 Storage in DKP

6.7.7.1 Configuring Loki

1. Execute the following command to get the namespace of your workspace:

```
dkp get workspaces
```

2. Set the `WORKSPACE_NAMESPACE` variable to the namespace copied in the previous step:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

3. Create a secret containing the static AWS S3 credentials. The secret is then mounted into each of the Grafana Loki pods as environment variables.

```
kubectl create secret generic dkp-aws-s3-creds -n${WORKSPACE_NAMESPACE} \
--from-literal=AWS_ACCESS_KEY_ID=<key id> \
--from-literal=AWS_SECRET_ACCESS_KEY=<secret key>
```

4. Create a config overrides ConfigMap to update the storage configuration:

NOTE: This could also be added to the installer configuration if you are configuring Grafana Loki on the Management Cluster

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
```

```

kind: ConfigMap
metadata:
  name: grafana-loki-overrides
  namespace: ${WORKSPACE_NAMESPACE}
data:
  values.yaml: |
    loki:
      annotations:
        secret.reloader.stakater.com/reload: dkp-aws-s3-creds
      structuredConfig:
        storage_config:
          aws:
            s3: s3://<region>/<bucket name>
    ingester:
      extraEnvFrom:
        - secretRef:
            name: dkp-aws-s3-creds
    querier:
      extraEnvFrom:
        - secretRef:
            name: dkp-aws-s3-creds
    queryFrontend:
      extraEnvFrom:
        - secretRef:
            name: dkp-aws-s3-creds
    compactor:
      extraEnvFrom:
        - secretRef:
            name: dkp-aws-s3-creds
    ruler:
      extraEnvFrom:
        - secretRef:
            name: dkp-aws-s3-creds
    distributor:
      extraEnvFrom:
        - secretRef:
            name: dkp-aws-s3-creds
EOF

```

- Update the `grafana-loki` AppDeployment to apply the configuration override.

NOTE: If you use the Kommander CLI installation configuration file, you don't need this step

```

cat << EOF | kubectl -n ${WORKSPACE_NAMESPACE} patch appdeployment grafana-loki
--type="merge" --patch-file=/dev/stdin
spec:
  configOverrides:
    name: grafana-loki-overrides
EOF

```


6.8 Security

Security architecture.

- [Authentication and authorization architecture](#) (see page 934)
- [OpenID Connect \(OIDC\) Introduction](#) (see page 935)
- [SAML Connector](#) (see page 939)
- [Policy Controls](#) (see page 942)
- [Traefik-Forward-Authentication in DKP \(TFA\)](#) (see page 946)
- [Create Local Access](#) (see page 948)

6.8.1 Authentication and authorization architecture

6.8.1.1 Details on distributed authentication and authorization between clusters

6.8.1.2 Authentication

DKP UI comes with a pre-configured authentication [Dex](#)⁶³² identity broker and provider.



Kubernetes, Kommander, and Dex do not store any user identities. The Kommander installation comes with default admin static credentials. These credentials should only be used to access the **DKP UI** for configuring an external identity provider. Currently, there is no way to update these credentials, so they should be treated as backup credentials and not used for normal access.


The DKP UI admin credentials are stored as a secret. They never leave the boundary of the UI cluster and are never shared to any other cluster.

The Dex service issues an [OIDC ID token](#)⁶³³ on successful user authentication. Other platform services use the ID token as an authentication proof. User identity to the Kubernetes API server is provided by the `kube-oidc-proxy` platform service that reads the identity from an ID token. Web requests to DKP UI access are authenticated by the [traefik forward auth](#)⁶³⁴ platform service.

⁶³² <https://github.com/dexidp/dex>

⁶³³ https://openid.net/specs/openid-connect-core-1_0.html#IDToken

⁶³⁴ <https://github.com/mesosphere/traefik-forward-auth>

 The `kube-oidc-proxy`⁶³⁵ service authenticates kubectl CLI requests using the Kubernetes API Server Go library. This library requires that if an **email_verified** claim is present, it must be set to **true**, even if the `insecureSkipEmailVerified: true` flag is configured in the Dex connector. Thus, ensure that the OIDC provider is configured to set the `email_verified` field to 'true'.

A user identity is shared across a UI cluster and all other attached clusters.

6.8.1.2.1 Attached clusters

A newly attached cluster has federated `kube-oidc-proxy`, `dex-k8s-authenticator`, and `traefik-forward-auth` platform applications. These platform applications are configured to accept the [Management/Essential](#) (see page 81) cluster Dex issued ID tokens.

When the `traefik-forward-auth` is used as a [Traefik Ingress authenticator](#)⁶³⁶, it checks if the user identity was issued by the Kommander cluster Dex service. An anonymous user is redirected to the Management/Essential cluster Dex service to authenticate and confirm their identity.

Never enter your own credentials on any of the attached clusters. On the the Management/Essential cluster use the static admin credentials or an external identity provider (IDP).

6.8.1.3 Authorization

There is no centralized authorization component in Kommander. Each component and service makes its own authorization decisions based on user identity.

6.8.2 OpenID Connect (OIDC) Introduction

6.8.2.1 An introduction to OpenID Connect (OIDC) Authentication in Kubernetes

All Kubernetes clusters have two categories of users: service accounts and normal users. Kubernetes manages authentication for service accounts, but the cluster administrator, or a separate service, manages authentication for normal users.

Kommander configures the cluster to use OpenID Connect (OIDC), a popular and extensible user authentication method, and installs Dex, a popular, open-source software product that integrates your existing Identity Providers with Kubernetes.

To begin, set up an Identity Provider with Dex, then use OIDC as the Authentication method.

⁶³⁵ <https://github.com/jetstack/kube-oidc-proxy>

⁶³⁶ <https://docs.traefik.io/v2.4/providers/kubernetes-ingress/>

6.8.2.2 Identity Provider

An Identity Provider (IdP) is a service that lets you manage identity information for users, including groups. A cluster created in Kommander uses Dex as its IdP. Dex, in turn, delegates to one or more external IdPs.

If you already use one or more of the following IdPs, you can configure Dex to use them:

Name	Supports Refresh Tokens	Supports Groups Claim	Supports preferred_username Claim	Status	Notes
LDAP ⁶³⁷	yes	yes	yes	stable	
GitHub ⁶³⁸	yes	yes	yes	stable	
SAML 2.0 ⁶³⁹	no	yes	no	stable	
GitLab ⁶⁴⁰	yes	yes	yes	beta	
OpenID Connect ⁶⁴¹	yes	yes	yes	beta	Includes Salesforce, Azure, etc.
Google ⁶⁴²	yes	yes	yes	alpha	
LinkedIn ⁶⁴³	yes	no	no	beta	
Microsoft ⁶⁴⁴	yes	yes	no	beta	

⁶³⁷ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/ldap.md>

⁶³⁸ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/github.md>

⁶³⁹ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/saml.md>

⁶⁴⁰ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/gitlab.md>

⁶⁴¹ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/oidc.md>

⁶⁴² <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/google.md>

⁶⁴³ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/linkedin.md>

⁶⁴⁴ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/microsoft.md>

Name	Supports Refresh Tokens	Supports Groups Claim	Supports preferred_username Claim	Status	Notes
AuthProxy ⁶⁴⁵	no	no	no	alpha	Authentication proxies such as Apache2 mod_auth, etc.
Bitbucket Cloud ⁶⁴⁶	yes	yes	no	alpha	
OpenShift ⁶⁴⁷	no	yes	no	stable	

NOTE: These are the Identity Providers supported by Dex [2.22.0](#)⁶⁴⁸, the version used by DKP.

6.8.2.3 Add Login Connectors

Kommander uses Dex to provide OpenID Connect single sign-on (SSO) to the cluster. Dex can be configured to use multiple connectors, including GitHub, LDAP, and SAML 2.0. The [Dex Connector documentation](#)⁶⁴⁹ describes how to configure different connectors. You can add the configuration as the values field in the Dex application. An example Dex configuration provided to the Kommander CLI's `install` command would look similar to this:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
  dex:
    values: |
      config:
        connectors:
          - type: oidc
            id: google
            name: Google
            config:
              issuer: https://accounts.google.com/o/oauth2/v2/auth
              clientID: YOUR_CLIENT_ID
              clientSecret: YOUR_CLIENT_SECRET
```

⁶⁴⁵ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/authproxy.md>

⁶⁴⁶ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/bitbucketcloud.md>

⁶⁴⁷ <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/openshift.md>

⁶⁴⁸ <https://github.com/dexidp/dex/blob/v2.22.0/README.md>

⁶⁴⁹ <https://dexidp.io/docs/connectors/>

```

redirectURI: https://DKP_CLUSTER_DOMAIN/dex/callback
scopes:
- openid
- profile
- email
insecureSkipEmailVerified: true
insecureEnableGroups: true
userIDKey: email
userNameKey: email

```

[...]

6.8.2.4 Change the Access Token Lifetime

By default, the client access token lifetime is 24 hours. After this time, the token expires and cannot be used to authenticate. See the [Dex documentation](#)⁶⁵⁰ for more information on access token expiration and rotation settings.

Here is an example configuration for extending the token lifetime to 48 hours:

```

apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
  dex:
    values: |
      config:
        expiry:
          idTokens: "48h"

```

[...]

6.8.2.5 Authentication

OpenID Connect is an extension of the [OAuth2 authentication protocol](#)⁶⁵¹. As required by OAuth2, the client must be registered with Dex. Do this by passing the name of the application and a callback/redirect URI. These handle the processing of the OpenID token after the user authenticates successfully. After registration, Dex returns a `client_id` and a `secret`. Authentication requests use these between the client and Dex to identify the client.

Users access Kommander in two ways:

- To interact with Kubernetes API, usually through `kubectl`.
- To interact with the DKP UI, which has GUI dashboards for Prometheus, Grafana, etc.

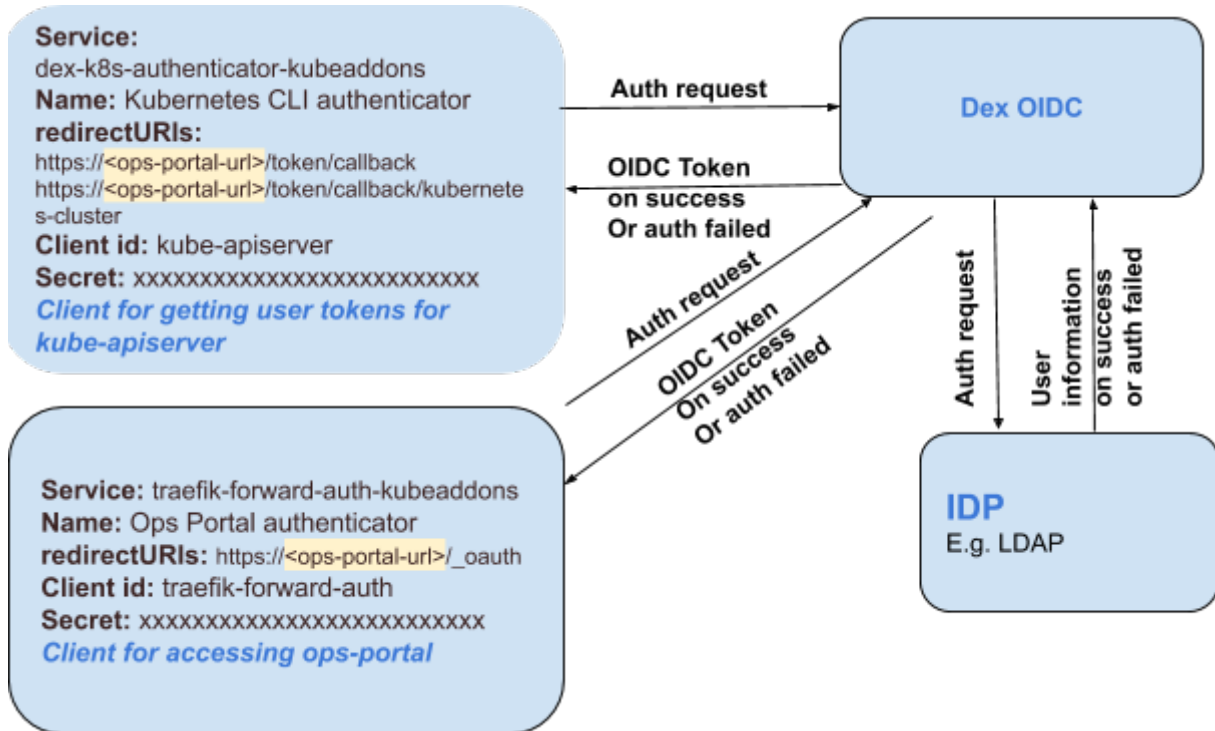
In Kommander, Dex comes pre-configured with a client for these access use cases. The clients talk to Dex for authentication. Dex talks to the configured Identity Provider, or IdP, (for example LDAP, SAML, etc) to perform the actual task of authenticating the user.

⁶⁵⁰ <https://dexidp.io/docs/id-tokens/#expiration-and-rotation-settings>

⁶⁵¹ <https://oauth.net/2/>

If the user authenticates successfully, Dex pulls the user’s information from the IdP and forms an OpenID token. The token contains this information and returns it to the respective client’s callback URL. The client or end user uses this token for communicating with the DKP UI or Kubernetes API respectively.

This figure illustrates these components and their interaction at a high level:



6.8.3 SAML Connector

6.8.3.1 Connect your Kommander cluster to an IdP using SAML

6.8.3.2 Connect Kommander to an IdP Using SAML

This procedure configures your Kommander cluster to use SAML, to connect to an identity provider (IdP).

1. [Install DKP \(see page 108\)](#).
2. Configure the IdP

Provide the issuer URL and the Assertion Consumer Service (ACS) or callback URL to your IdP. The issuer URL points to the authentication endpoint at the service provider (Dex), which issues a request towards the IdP via the user agent.

The issuer URL follows this schema:

```
https://<your-cluster-host>/dex
```

The ACS URL points to the service provider (Dex) endpoint that receives SAML assertions issued by the IdP.

The ACS or callback URL should look like this:

```
https://<your-cluster-host>/dex/callback
```

Depending on the IdP, you might be asked to provide the configuration in some form of an XML snippet. See the following example, making sure to replace `<your-cluster-host>` with your URL:

```
<?xml version="1.0" encoding="UTF-8"?>
<EntityDescriptor xmlns="urn:oasis:names:tc:SAML:2.0:metadata" entityID="https://<your-cluster-host>/dex">
  <SPSSODescriptor AuthnRequestsSigned="false" WantAssertionsSigned="true"
    protocolSupportEnumeration="urn:oasis:names:tc:SAML:2.0:protocol">
    <NameIDFormat>urn:oasis:names:tc:SAML:2.0:nameid-format:persistent</
    NameIDFormat>
    <AssertionConsumerService index="0" isDefault="true" Binding="urn:oasis:
    names:tc:SAML:2.0:bindings:HTTP-POST" Location="https://<your-cluster-host>/
    dex/callback" />
  </SPSSODescriptor>
</EntityDescriptor>
```

3. Modify the `dex` configuration:

For this step, get the following from your IdP:

- single sign-on URL or SAML URL -> `ssoURL`
- base64 encoded, PEM encoded CA certificate -> `caData`
- username attribute name in SAML response -> `usernameAttr`
- email attribute name in SAML response -> `emailAttr`

From above you need:

- issuer URL -> `entityIssuer`
- callback URL -> `redirectURI`

Ensure you base64 encode the contents of the PEM file. As an example, the prefix of the contents will result into this exact base64 prefix:

```
LS0tLS1CRUdJTiBDRVJUSUZJQ0FURS0tLS0tC[...]
```

You can add the configuration as the values field in the `dex` application. An example `dex` configuration provided to the [Kommander CLI's install command](#) (see page 1422) should look similar to:

```

apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
  dex:
    values: |
      config:
        connectors:
          - type: saml
            id: saml
            name: SAML
            config:
              ssoURL: < url for POST request >
              caData: < base64 PEM encoded CA for the IdP server >
              redirectURI: https://<your-cluster-host>/dex/callback
              entityIssuer: https://<your-cluster-host>/dex
              usernameAttr: < user attribute in saml response >
              emailAttr: < email attribute in saml response >
[...]
```

4. Modify the `traefik-forward-auth-mgmt` configuration and add a whitelist:

This step is required to give access to a user to the DKP UI. For each user, you must give [Access to Kubernetes resources](#) (see page 586) and add an entry in the `whitelist` below.

```

apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
  ...
  traefik-forward-auth-mgmt:
    values: |
      traefikForwardAuth:
        allowedUser:
          valueFrom:
            secretKeyRef: null
        whitelist:
          - < allowed email addresses >
```

5. Run `kommander install --installer-config kommander.yaml` to deploy modified `dex`.
6. Visit `https://<your-cluster-host>/dkp/kommander/dashboard` to login to the DKP UI.
7. Select `Launch Console` and follow the authentication steps to complete the procedure.

6.8.4 Policy Controls

6.8.4.1 Workload Policy Controls

6.8.4.2 Enforce Policies Using Gatekeeper

6.8.4.2.1 Learn how to enforce policies using Gatekeeper

[Gatekeeper](#)⁶⁵² is the policy controller for Kubernetes, allowing organizations to enforce configurable policies using the [Open Policy Agent](#)⁶⁵³, a policy engine for Cloud Native environments hosted by CNCF as a graduated-level project.

This tutorial describes how to use Gatekeeper to enforce policies by rejecting non-compliant resources. Specifically, this tutorial describes two constraints as a way to use Gatekeeper as an alternative to [Pod Security Policies](#)⁶⁵⁴:

- [Prevent the running of privileged pods](#) (see page 0)
- [Prevent mounting host path volumes](#) (see page 0)

6.8.4.2.2 Before you Begin

Before starting this tutorial, verify the following:

- You must have access to a Linux, macOS, or Windows computer with a supported operating system version.
- You must have a properly deployed and running cluster. For information about deploying Kubernetes with default settings on different types of infrastructures, see the [Choose Infrastructure](#) (see page 1036) topic.
- If you install Kommander with a custom configuration, make sure you enable Gatekeeper.



If you intend to disable Gatekeeper, keep in mind that the app is deployed pre-configured with constraint templates that enforce multi-tenancy in projects.

652 <https://github.com/open-policy-agent/gatekeeper>

653 <https://github.com/open-policy-agent/opa>

654 <https://kubernetes.io/docs/concepts/policy/pod-security-policy/>

6.8.4.2.3 Use Gatekeeper

Gatekeeper uses the [OPA Constraint Framework](#)⁶⁵⁵ to describe and enforce policy. Before you can define a constraint, you must first define a `ConstraintTemplate`, which describes both the `Rego` (a powerful query language) that enforces the constraint and the schema of the constraint. The schema of the constraint allows an admin to fine-tune the behavior of a constraint, much like arguments to a function.

The Gatekeeper repository includes a [library of policies](#)⁶⁵⁶ to replace Pod Security Policies which you will use in the following tutorials.

6.8.4.2.3.1 Prevent privileged pods

Define the ConstraintTemplate

Create the privileged pod policy constraint template `k8spspprivilegedcontainer` by running the following command:

```
kubectl apply -f https://raw.githubusercontent.com/open-policy-agent/gatekeeper-library/master/library/pod-security-policy/privileged-containers/template.yaml
```

Define the Constraint

Constraints are then used to inform Gatekeeper that the admin wants to enforce a `ConstraintTemplate`, and how.

Create the privileged pod policy constraint `psp-privileged-container` by running the following command:

```
kubectl apply -f https://raw.githubusercontent.com/open-policy-agent/gatekeeper-library/master/library/pod-security-policy/privileged-containers/samples/psp-privileged-container/constraint.yaml
```

Test that the constraint is enforced

Create a privileged pod by running the following command:

```
kubectl apply -f https://raw.githubusercontent.com/open-policy-agent/gatekeeper-library/master/library/pod-security-policy/privileged-containers/samples/psp-privileged-container/example_disallowed.yaml
```

You should see the following output:

⁶⁵⁵ <https://github.com/open-policy-agent/frameworks/tree/master/constraint>

⁶⁵⁶ <https://github.com/open-policy-agent/gatekeeper-library/tree/master/library/pod-security-policy>

```
Error from server ([denied by psp-privileged-container] Privileged container is not
allowed: nginx, securityContext: {"privileged": true}): error when creating "https://
raw.githubusercontent.com/open-policy-agent/gatekeeper-library/master/library/pod-
security-policy/privileged-containers/samples/psp-privileged-container/
example_disallowed.yaml": admission webhook "validation.gatekeeper.sh" denied the
request: [denied by psp-privileged-container] Privileged container is not allowed:
nginx, securityContext: {"privileged": true}
```

6.8.4.2.3.2 Prevent host path volumes

Define the ConstraintTemplate

Create the host path volume policy constraint template `k8spsphostfilesystem` by running the following command:

```
kubectl apply -f https://raw.githubusercontent.com/open-policy-agent/gatekeeper-
library/master/library/pod-security-policy/host-filesystem/template.yaml
```

Define the Constraint

Constraints are then used to inform Gatekeeper that the admin wants to enforce a ConstraintTemplate, and how.

Create the host path volume policy constraint `psp-host-filesystem` by running the following command to only allow `/foo` to be mounted as a host path volume:

```
cat <<EOF | kubectl apply -f -
apiVersion: constraints.gatekeeper.sh/v1beta1
kind: K8sPSPHostFilesystem
metadata:
  name: psp-host-filesystem
spec:
  match:
    kinds:
      - apiGroups: [""]
        kinds: ["Pod"]
  parameters:
    allowedHostPaths:
      - readOnly: true
        pathPrefix: "/foo"
EOF
```

Test that the constraint is enforced

Create a privileged pod by running the following command:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: Pod
metadata:
  name: nginx-host-filesystem
  labels:
    app: nginx-host-filesystem
spec:
  containers:
  - name: nginx
    image: nginx
    volumeMounts:
    - mountPath: /cache
      name: cache-volume
      readOnly: true
  volumes:
  - name: cache-volume
    hostPath:
      path: /tmp # directory location on host
EOF
```

You should see the following output:

```
Error from server ([denied by psp-host-filesystem] HostPath volume {"hostPath": {"path": "/tmp", "type": ""}, "name": "cache-volume"} is not allowed, pod: nginx-host-filesystem. Allowed path: [{"readOnly": true, "pathPrefix": "/foo"}]): error when creating "STDIN": admission webhook "validation.gatekeeper.sh" denied the request: [denied by psp-host-filesystem] HostPath volume {"hostPath": {"path": "/tmp", "type": ""}, "name": "cache-volume"} is not allowed, pod: nginx-host-filesystem. Allowed path: [{"readOnly": true, "pathPrefix": "/foo"}]
```

Test that the constraint allows the allowed host paths

Create a pod that mounts an allowed host path by running the following command:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: Pod
metadata:
  name: nginx-host-filesystem
  labels:
```

```
  app: nginx-host-filesystem
spec:
  containers:
  - name: nginx
    image: nginx
    volumeMounts:
    - mountPath: /cache
      name: cache-volume
      readOnly: true
  volumes:
  - name: cache-volume
    hostPath:
      path: /foo # directory location on host
EOF
```

You should see the following output:

```
pod/nginx-host-filesystem created
```


6.8.5 Traefik-Forward-Authentication in DKP (TFA)

6.8.5.1 TFA Overview

Traefik-Forward-Authentication (TFA) is a vital component in DKP. It handles authentication and authorization of web (HTTP) requests and how users can access the DKP UI. Regardless of the authentication method you choose (user and password, SSO, 2FA, etc.), TFA is a component that authorizes HTTP clients to enter your environment.

TFA is one of the standard applications in the Kommander component of DKP. It is deployed by a controller on all attached, managed, and management clusters.

6.8.5.2 TFA Authentication Workflow



6.8.5.3 Default TFA Configuration in DKP

In the default configuration in DKP, TFA stores all authentication information about users via the browser cookies. When TFA authenticates users, it stores the user’s metadata in encrypted browser cookies.

Subsequent requests following initial authentication will use these cookies to recognize users without the need to re-authenticate them.



NOTE:

- The cookies are securely encrypted, so they cannot be modified by users.
- The cookies contain the RBAC username.
- The cookies contain a list of groups that users are associated with.

6.8.5.4 Cluster Storage Option

The browser cookie storage is limited to a maximum of 4Kb per cookie. However, if a user is assigned to a large number of groups, this limitation can be exceeded and this will return a response 500-internal server error, meaning that the user will be unable to access any web services.

To work around the cookie storage size limit, TFA can be configured to store the metadata claims in the cluster as a Kubernetes secret instead of in the browser. To do so, the `clusterStorage` option can be configured in the Traefik Forward Auth application when installing Kommander.

In order to enable the `clusterStorage` option, add the following to the `cluster.yaml` when installing Kommander:

```
traefik-forward-auth-mgmt:
  values: |
    clusterStorage:
      enabled: true
      namespace: kommander
```

If the `clusterStorage` feature is enabled, automatic garbage collection will delete the secrets after 12 hours. Keep in mind that enabling this feature will have performance implications for web requests, because TFA needs to load the secret to retrieve the user groups for each HTTP API request. Because of this, we recommend first trying to reduce the number of groups associated users and only enabling this option if that cannot be accomplished.

For additional information about traefik-forward-authentication, see the [TFA page on Github](#)⁶⁵⁷.

6.8.6 Create Local Access

D2iQ recommends configuring an **external identity provider** to manage access to your cluster.

⁶⁵⁷ <https://github.com/mesosphere/traefik-forward-auth>

See <https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Identity%20Providers&linkCreation=true&fromPagelD=273385909>⁶⁵⁸ for an overview of the benefits, supported providers, and instructions on how to configure an external identity provider.

However, you can create **local users** as an alternative, for example, when you want to quickly test RBAC policies. DKP automatically creates a unique local user, the `admin` user, but you can create additional local users and assign them other RBAC roles.

This section shows how to create and manage local users to access your Essential or Management cluster.

6.8.6.1 Next Step:

There are two ways in which you can create local users:

- [Create Local Users during the Kommander Installation](#)⁶⁵⁹
- [Create Local Users after Installing Kommander](#)⁶⁶⁰

6.8.6.2 Create Local Users during the Kommander Installation

Create local users before you install the Kommander component of DKP. If you have already installed Kommander, see [Create Local Users after Installing Kommander](#) (see page 950).




D2iQ does not recommend creating local users for production clusters. See [Identity Providers](#) (see page 600) for instructions on how to configure an external identity provider to manage your users.

1. Open the **Kommander Installer Configuration File** or `kommander.yaml` file. If you do not have the `kommander.yaml` file, [initialize the configuration file](#) (see page 1422), so you can edit it in the following steps.
 - ⚠ Initialize this file only ONE time, otherwise you will overwrite previous customizations.
2. In that file, add the following customization for `dex` and create local users by establishing their credentials:
 - 📘 Replace `<example_email>` with the user's email address or username.

⁶⁵⁸ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelD=273385909&linkCreation=true&spaceKey=DENT&title=%282.6%29+Identity+Providers>

⁶⁵⁹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelD=273385909&linkCreation=true&spaceKey=DENT&title=%282.6%29+Create+Local+Users+during+the+Kommander+Installation>


⁶⁶⁰ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelD=273385909&linkCreation=true&spaceKey=DENT&title=%282.6%29+Create+Local+Users+after+Installing+Kommander>

 Replace `<password_bcrypt_hash>` with the bcrypt hash of the password you want to assign. You can use the `htpasswd` CLI to create the hash of a specific password. For example, by running `htpasswd -bnBC 10 "" password | tr -d ':\n' && echo`, you can generate the hash for the password “password”.

```
app:
  dex:
    enabled: true
    values: |
      config:
        staticPasswords:
          - email: <example_email>
            hash: <password_bcrypt_hash>
          - email: <example_email2>
            hash: <password_bcrypt_hash2>
```

3. Save the `kommander.yaml` file.
4. Install Kommander using the customized installation file. See <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29892161/Kommander+Customizations#install-config-file> (see page 1424) for more information.

After Kommander finishes installing, you have created local users, but you cannot use them until you have [assigned them permissions](#) (see page 952).


 You have created a user that does not have any permissions to see or manage your DKP cluster yet. [Add RBAC Roles to Local Users](#) (see page 952) to complete the configuration.

6.8.6.2.1 Next Step:

[Add RBAC Roles to Local Users](#) (see page 952)

6.8.6.3 Create Local Users after Installing Kommander

Create local users after you install the Kommander component of DKP. If you have not installed Kommander yet, and want to create additional users during the installation, see [Create Local Users during the Kommander Installation](#) (see page 952).

 D2iQ does not recommend creating local users for production clusters. See [Identity Providers](#) (see page 600) for instructions on how to configure an external identity provider to manage your users.

Customize the Dex AppDeployment, and add a `configOverrides` section.

1. Create a `configMap` resource with the credentials of the new local user:

i Replace `<example_email>` with the user's email address or a username.

i Replace `<password_bcrypt_hash>` with the bcrypt hash of the password you want to assign. You can use the `htpasswd` CLI to create the hash of a specific password. For example, by running `htpasswd -bnBC 10 "" password | tr -d ':\n' && echo` you can generate the hash for the password "password".

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: dex-overrides
  namespace: kommander
data:
  values.yaml: |
    config:
      staticPasswords:
        - email: <example_email>
          hash: <password_bcrypt_hash>
EOF
```

2. Open the Dex AppDeployment to edit it:

```
kubectl edit -n kommander appdeployment dex
```

The editor displays the AppDeployment.

3. Copy the following values and paste them in a location in the file where they are nested in the `spec` field:

```
configOverrides:
  name: dex-overrides
```

Example:

```
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  ...
spec:
  appRef:
    kind: ClusterApp
    name: dex-2.11.1
```

```

clusterConfigOverrides:
- clusterSelector:
  matchExpressions:
  - key: kommander.d2iq.io/cluster-name
    operator: In
    values:
    - host-cluster
  configMapName: dex-kommander-overrides
  configOverrides: # Copy and paste this section.
  name: dex-overrides
status:
...

```

Editing the AppDeployment restarts the HelmRelease for Dex. The new users will be created after the reconciliation. However, the user creation is not completed until you [assign it permissions](#) (see page 952).



You have created a user that does not have any permissions to see or manage your DKP cluster yet. [Add RBAC Roles to Local Users](#) (see page 952) to complete the configuration.

6.8.6.3.1 Next Step:

[Add RBAC Roles to Local Users](#) (see page 952)

6.8.6.4 Add RBAC Roles to Local Users

Manage access to your cluster and its resources by assigning Kubernetes RBAC roles to local users. If you have not created local users yet, see [Create Local Users after Installing Kommander](#) (see page 953) or [Create Local Users during the Kommander Installation](#) (see page 952).

6.8.6.4.1 Assign a Role - Cluster Admin Example

Create the following `ClusterRoleBinding` resource:



- Replace `<example_email>` with the user's email address or a username.
- Replace `cluster-admin` with the RBAC role you want to assign to a user.

```
cat <<EOF | kubectl apply -f -
```

```

apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding
metadata:
  name: cluster-admin
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: cluster-admin
subjects:
- apiGroup: rbac.authorization.k8s.io
  kind: User
  name: <example_email>
EOF

```

After assigning the previous role to `<example_email>`, the user is able to log in to the cluster using the credentials you assigned in [Create Local Users after Installing Kommander \(see page 953\)](#) or [Create Local Users during the Kommander Installation \(see page 952\)](#).



The **Login page** and cluster **URL** is the same for the default admin user and the local users you create with this method.

6.8.6.4.2 More Information

For more information on RBAC resources in DKP, see [Granting Access to Kubernetes and Kommander Resources \(see page 589\)](#).

For general information on RBAC as a Kubernetes resource, see the official Kubernetes [RBAC⁶⁶¹](#) documentation.

6.8.6.4.3 Next Topic:

[Modify or Delete Local Users \(see page 953\)](#)

6.8.6.5 Modify or Delete Local Users

6.8.6.5.1 Modify Local Users

To change the password or username of a user, edit the `dex-overrides` ConfigMap.

If you change the email address or username of a user, ensure you update any `RoleBindings` or RBAC resources associated with this user.

⁶⁶¹ <https://kubernetes.io/docs/reference/access-authn-authz/rbac>

6.8.6.5.2 Delete Local Users

To delete local users, edit the `dex-overrides` ConfigMap and remove the email and hash fields for the user. Also, ensure you delete any `RoleBindings` or RBAC resources associated with this user.

6.9 Networking

6.9.1 Configure networking for Konvoy cluster

This section describes different networking components that come together to form a Konvoy networking stack. It assumes familiarity with Kubernetes networking.

- [Networking Service](#) (see page 954)
- [Required Domains](#) (see page 960)
- [Configure Ingress for load balancing](#) (see page 961)
- [Ingress](#) (see page 963)
- [Load Balancing](#) (see page 964)
- [Use Istio as a Microservice](#) (see page 965)

6.9.2 Networking Service

A [Service](#)⁶⁶² is an API resource that defines a logical set of pods and a policy by which to access them, and is an abstracted manner to expose applications as network services.

Kubernetes gives pods their own IP addresses and a single DNS name for a set of pods. Services are used as endpoints to load-balance the traffic across the pods. A selector determines the set of Pods targeted by a Service.

For example, if you have a set of pods that each listen on TCP port `9191` and carry a label `app=MyKonvoyApp`, as configured in the following:

```
apiVersion: v1
kind: Service
metadata:
  name: my-konvoy-service
  namespace: default
spec:
  selector:
    app: MyKonvoyApp
  ports:
    - protocol: TCP
      port: 80
```

⁶⁶² <https://kubernetes.io/docs/concepts/services-networking/service/#service-resource>

```
targetPort: 9191
```

This specification creates a new `Service` object named `"my-konvoy-service"`, that targets TCP port `9191` on any pod with the `app=MyKonvoyApp` label.

Kubernetes assigns this `Service` an IP address. In particular, the `kube-proxy` implements a form of virtual IP for `Services` of type other than `ExternalName`.



- The name of a `Service` object must be a valid DNS label name.
- A `Service` is not a Platform `Service`.

6.9.2.1 Service Topology

[Service Topology](#)⁶⁶³ is a mechanism in Kubernetes to route traffic based upon the Node topology of the cluster. For example, you can configure a `Service` to route the traffic to endpoints on specific nodes, or even based on the region or availability zone of the node's location.

To enable this new feature in your Kubernetes cluster, use the feature gates `--feature-gates="ServiceTopology=true,EndpointSlice=true"` flag. After enabling, you can control `Service` traffic routing by defining the `topologyKeys` field in the `Service` API object.

In the following example, a `Service` defines `topologyKeys` to be routed to endpoints only in the same zone:

```
apiVersion: v1
kind: Service
metadata:
  name: my-konvoy-service
  namespace: default
spec:
  selector:
    app: MyKonvoyApp
  ports:
  - protocol: TCP
    port: 80
    targetPort: 9191
  topologyKeys:
  - "topology.kubernetes.io/zone"
```

⁶⁶³ <https://kubernetes.io/docs/concepts/services-networking/service-topology/#using-service-topology>

ⓘ If the value of the `topologyKeys` field does not match any pattern, the traffic is rejected.

6.9.2.2 EndpointSlices

EndpointSlices⁶⁶⁴ are an API resource that appears as a scalable and more manageable solution to network endpoints within a Kubernetes cluster. They allow for distributing network endpoints across multiple resources with a limit of 100 endpoints per EndpointSlice.

An EndpointSlice contains references to a set of endpoints, and the control plane takes care of creating EndpointSlices for any Service that has a selector specified. These EndpointSlices include references to all the pods that match the Service selector.

Like Services, the name of a EndpointSlice object must be a valid DNS subdomain name.

In this example, here's a sample EndpointSlice resource for the example Kubernetes Service:

```
apiVersion: discovery.k8s.io/v1beta1
kind: EndpointSlice
metadata:
  name: konvoy-endpoint-slice
  namespace: default
  labels:
    kubernetes.io/service-name: my-konvoy-service
addressType: IPv4
ports:
- name: http
  protocol: TCP
  port: 80
endpoints:
- addresses:
  - "192.168.126.168"
  conditions:
    ready: true
  hostname: ip-10-0-135-39.us-west-2.compute.internal
  topology:
    kubernetes.io/hostname: ip-10-0-135-39.us-west-2.compute.internal
    topology.kubernetes.io/zone: us-west2-b
```

6.9.2.3 DNS for Services and Pods

Every new Service object in Kubernetes gets assigned a DNS name. The Kubernetes DNS component schedules a DNS name for the pods and services created on the cluster, and then the Kubelets are configured so containers can resolve these DNS names.

⁶⁶⁴ <https://kubernetes.io/docs/concepts/services-networking/endpoint-slices/#endpointslice-resource>

Considering previous examples, assume there is a Service named `my-konvoy-service` in the Kubernetes namespace `default`. A Pod running in namespace `default` can look up this service by performing a DNS query for `my-konvoy-service`. A Pod running in namespace `kommander` can look up this service by performing a DNS query for `my-konvoy-service.default`.

In general, a pod has the following DNS resolution:

```
pod-ip-address.namespace-name.pod-name.cluster-domain.example.
```

Similarly, a service has the following DNS resolution:

```
service-name.namespace-name.svc.cluster-domain.example.
```

You can find additional information about all the possible record types and layout [here](#)⁶⁶⁵.

6.9.2.4 Ingress and Networking

[Ingress](#)⁶⁶⁶ is an API resource that manages external access to the services in a cluster through HTTP or HTTPS. It offers name-based virtual hosting, SSL termination and load balancing when exposing HTTP/HTTPS routes from outside to services in the cluster.

The traffic policies are controlled by rules as part of the Ingress definition. Each rule defines the following details:

- An optional host to which apply the rules.
- A list of paths or routes which has an associated backend defined with a Service `name`, a port `name` and `number`.
- A backend is a combo of a Service and port names, or a custom resource backend defined as a CRD. Consequently HTTP/HTTPS requests to the Ingress that matches the host and path of the rule are sent to the listed backend.

An example of an Ingress specification is:

```
apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
  name: konvoy-ingress
  namespace: default
  annotations:
    nginx.ingress.kubernetes.io/rewrite-target: /
spec:
  rules:
  - http:
```

⁶⁶⁵ <https://kubernetes.io/docs/concepts/services-networking/dns-pod-service/>

⁶⁶⁶ <https://kubernetes.io/docs/concepts/services-networking/ingress/#what-is-ingress>


```

paths:
- path: /path
  pathType: Prefix
  backend:
    service:
      name: my-konvoy-service
      port:
        number: 80

```

In Kommander, you can expose services to the outside world using Ingress objects.

6.9.2.4.1 Ingress Controllers

In contrast with the controllers in the Kubernetes control plane, Ingress controllers are not started with a cluster so you need to choose the desired Ingress controller.

An Ingress controller has to be deployed in a cluster for the Ingress definitions to work.

Kubernetes as a project currently supports and maintains GCE and nginx controllers.

These are four of the most known [Ingress controllers](#)⁶⁶⁷:

- [HAProxy Ingress](#)⁶⁶⁸ is a highly customizable community-driven ingress controller for HAProxy.
- NGINX offers support and maintenance for the [NGINX Ingress Controller](#)⁶⁶⁹ for Kubernetes.
- [Traefik](#)⁶⁷⁰ is a fully featured Ingress controller (Let's Encrypt, secrets, http2, websocket), and has commercial support.
- [Ambassador API Gateway](#)⁶⁷¹ EXPERIMENTAL is an Envoy based Ingress controller with community and commercial support.

In Kommander, `Traefik` deploys by default as a well-suited Ingress controller.

6.9.2.5 Network Policies

NetworkPolicy is an API resource that controls the traffic flow at port level 3 or 4, or at the IP address level. It enables defining constraints on how a pod communicates with various network services such as `endpoints` and `services`.

A Pod can be restricted to talk to other network services through a selection of the following identifiers:

- Namespaces that have to access. There can be pods that are not allowed to talk to other namespaces.
- Other allowed IP blocks regardless of the node or IP address assigned to the targeted Pod.
- Other allowed Pods.

An example of a NetworkPolicy specification is:

⁶⁶⁷ <https://kubernetes.io/docs/concepts/services-networking/ingress-controllers/#additional-controllers>

⁶⁶⁸ <https://haproxy-ingress.github.io/>

⁶⁶⁹ <https://www.nginx.com/products/nginx-ingress-controller>

⁶⁷⁰ <https://github.com/containous/traefik>

⁶⁷¹ <https://www.getambassador.io/>

```

apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
  name: network-konvoy-policy
  namespace: default
spec:
  podSelector:
    matchLabels:
      role: db
  policyTypes:
  - Ingress
  - Egress
  ingress:
  - from:
    - ipBlock:
        cidr: 172.17.0.0/16
        except:
        - 172.17.1.0/24
    - namespaceSelector:
        matchLabels:
          app: MyKonvoyApp
    - podSelector:
        matchLabels:
          app: MyKonvoyApp
    ports:
    - protocol: TCP
      port: 6379
  egress:
  - to:
    - ipBlock:
        cidr: 10.0.0.0/24
    ports:
    - protocol: TCP
      port: 5978

```

As shown in the example, when defining a pod or namespace based NetworkPolicy, you use a selector to specify what traffic is allowed to and from the Pod(s).

6.9.2.5.1 Adding Entries to Pod /etc/hosts with HostAliases

The Pod API resource definition has a `HostAliases` field that allows adding entries to the Pod's container `/etc/hosts` file. This field overrides the hostname resolution when DNS and other options are not applicable.

For example, to resolve `foo.node.local`, `bar.node.local` to `127.0.0.1` and `foo.node.remote`, `bar.node.remote` to `10.1.2.3`, configure the `HostAliases` values as follows:

```

apiVersion: v1

```

```
kind: Pod
metadata:
  name: hostaliases-konvoy-pod
spec:
  restartPolicy: Never
  hostAliases:
  - ip: "127.0.0.1"
    hostnames:
    - "foo.node.local"
    - "bar.node.local"
  - ip: "10.1.2.3"
    hostnames:
    - "foo.node.remote"
    - "bar.node.remote"
  containers:
  - name: cat-hosts
    image: busybox
    command:
    - cat
    args:
    - "/etc/hosts"
```

6.9.3 Required Domains

6.9.3.1 This section describes the required domains for DKP

You must have access to the following domains through the customer networking rules so that Kommander can download all required images:

- <http://docker.io>
- <http://gcr.io>
- <http://k8s.gcr.io>
- mcr.microsoft.com⁶⁷²
- nvcr.io
- <http://quay.io>
- <http://us.gcr.io>
- registry.k8s.io⁶⁷³



In an air-gapped installation, these domains do not need to be accessible.

⁶⁷² <http://mcr.microsoft.com>

⁶⁷³ <http://registry.k8s.io>

6.9.4 Configure Ingress for load balancing

6.9.4.1 Learn how to configure Ingress settings for load balancing (layer-7)

Ingress is the name used to describe an API object that manages external access to the services in a cluster. Typically, an Ingress exposes HTTP and HTTPS routes from outside the cluster to services running within the cluster.

The object is called an Ingress because it acts as a gateway for inbound traffic. The Ingress receives inbound requests and routes them according to the rules you defined for the **Ingress resource** as part of your cluster configuration.

Expose an application running on your cluster by configuring an Ingress for load balancing (layer-7).

6.9.4.2 Prerequisites

Before you begin, you must:

- Have access to a Linux, macOS, or Windows computer with a supported operating system version.
- Have a properly deployed and running cluster.

6.9.4.3 Expose a pod using an Ingress (L7)

1. Deploy two web application Pods on your Kubernetes cluster by running the following command:

```
kubectl run --restart=Never --image hashicorp/http-echo --labels app=http-echo-1 --port 80 http-echo-1 -- -listen=:80 --text="Hello from http-echo-1"
kubectl run --restart=Never --image hashicorp/http-echo --labels app=http-echo-2 --port 80 http-echo-2 -- -listen=:80 --text="Hello from http-echo-2"
```

2. Expose the Pods with a service type of ClusterIP by running the following commands:

```
kubectl expose pod http-echo-1 --port 80 --target-port 80 --name "http-echo-1"
kubectl expose pod http-echo-2 --port 80 --target-port 80 --name "http-echo-2"
```

3. Create the Ingress to expose the application to the outside world by running the following command:

```
cat <<EOF | kubectl create -f -
apiVersion: networking.k8s.io/v1
kind: Ingress
metadata:
  annotations:
    kubernetes.io/ingress.class: kommander-traefik
    traefik.ingress.kubernetes.io/router.tls: "true"
```

```

generation: 7
name: echo
namespace: default
spec:
  rules:
  - http:
      paths:
      - backend:
          service:
            name: http-echo-1
            port:
              number: 80
          path: /echo1
          pathType: Prefix
  - http:
      paths:
      - backend:
          service:
            name: http-echo-2
            port:
              number: 80
          path: /echo2
          pathType: Prefix
EOF

```

The configuration settings in this example illustrate:

- setting the `kind` to `Ingress`.
- setting the `service.name` to be exposed as each `backend`.

4. Run the following command to get the URL of the load balancer created on AWS for the Traefik service:

```
kubectl get svc kommander-traefik -n kommander
```

This command displays the internal and external IP addresses for the exposed service. (Note that IP addresses and host names are for illustrative purposes. Always use the information from your own cluster)

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP
kommander-traefik	LoadBalancer	10.0.24.215	
abf2e5bda6ca811e982140acb7ee21b7-37522315			9/TCP,443:32297/TCP,8080:31923/TCP
		4h22m	80:3116

5. Validate that you can access the web application Pods by running the following commands: (Note that IP addresses and host names are for illustrative purposes. Always use the information from your own cluster)

```
curl -k https://abf2e5bda6ca811e982140acb7ee21b7-37522315.us-west-2.elb.amazonaws.com/echo1
curl -k https://abf2e5bda6ca811e982140acb7ee21b7-37522315.us-west-2.elb.amazonaws.com/echo2
```

6.9.5 Ingress

6.9.5.1 Traefik Ingress Controller

Kubernetes Ingress resources expose HTTP and HTTPS routes from outside the cluster to services within the cluster. In Kommander, the [Traefik](#)⁶⁷⁴ Ingress controller is installed by default and provides access to the DKP UI.

An Ingress performs the following:

- Gives Services externally-reachable URLs
- Load balances traffic
- Terminates SSL/TLS sessions
- Offers name-based virtual hosting

An Ingress controller fulfills the Ingress with a load balancer.

A cluster can have multiple Ingress controllers. D2iQ recommends adding your own Ingress controllers for your applications. The Traefik Ingress controller that Kommander installs for access to the DKP UI can be replaced later if a different solution is a better fit. Using your own Ingress controller in parallel for your own business requirements ensures that you are not limited by any future changes in Kommander.

6.9.5.2 Traefik v2.4

Traefik is a modern HTTP reverse proxy and load balancer that deploys microservices with ease. Kommander currently installs Traefik v2.4 by default on every cluster. Traefik creates a service of type `LoadBalancer`. In the cloud, the cloud provider creates the appropriate load balancer. In an on-premises deployment, by default, it uses MetalLB.

Traefik listens to the Kubernetes API and automatically generates and updates the routes without any further configuration or intervention so that the Services selected by the Ingress resources are connected to the outside world. Further, Traefik supports a rich set of functionality such as Name-based routing, Path-based routing, Traffic splitting, etc.

Major features highlighted in the Traefik documentation:

- Continuously updates its configuration (No restarts!)
- Supports multiple load balancing algorithms
- Provides HTTPS to your microservices
- Circuit breakers, retry
- A clean web UI
- Websocket, HTTP/2, GRPC ready

⁶⁷⁴ <https://landscape.cncf.io/card-mode?category=service-proxy&grouping=category&selected=traefik>

- Provides metrics (Rest, Prometheus, Datadog, StatsD, InfluxDB)
- Keeps access logs (JSON, CLF)
- Exposes a Rest API
- Packaged as a single binary file (made with go) and available as a docker image

6.9.5.3 Related Information

For information on related topics or procedures, refer to the following:

- [List of Ingress controllers](#)⁶⁷⁵
- [Traefik v2.4 docs](#)⁶⁷⁶
- [Configure Ingress for load balancing \(see page 964\)](#)
- [Load Balancing \(see page 964\)](#)

6.9.6 Load Balancing

In a Kubernetes cluster, depending on the flow of traffic direction, there are two kinds of load balancing:

- Internal load balancing for the traffic within a Kubernetes cluster
- External load balancing for the traffic coming from outside the cluster

6.9.6.1 Load Balancing for Internal Traffic

Load balancing within a Kubernetes cluster is accessed through a `ClusterIP` service type. `ClusterIP` presents a single IP address to the client and load balances the traffic going to this IP to the backend servers. The actual load balancing happens using `iptables` or IPVS configuration. The `kube-proxy` Kubernetes component programs these. The `iptables` mode of operation uses `DNAT`⁶⁷⁷ rules to distribute direct traffic to real servers, whereas `IPVS`⁶⁷⁸ leverages in-kernel transport-layer load-balancing. Read a [comparison between these two methods](#)⁶⁷⁹. By default, `kube-proxy` runs in `iptables` mode. The `kube-proxy` configuration can be altered by updating the `kube-proxy` configmap in the `kube-system` namespace.

6.9.6.2 Load Balancing for External Traffic

External traffic destined for the Kubernetes service requires a service of type `LoadBalancer`, through which external clients connect to your internal service. Under the hood, it uses a load balancer provided by the underlying infrastructure to direct the traffic.

⁶⁷⁵ <https://kubernetes.io/docs/concepts/services-networking/ingress-controllers/>

⁶⁷⁶ <https://doc.traefik.io/traefik/v2.4/>

⁶⁷⁷ https://www.linuxtopia.org/Linux_Firewall_iptables/x4013.html

⁶⁷⁸ https://en.wikipedia.org/wiki/IP_Virtual_Server

⁶⁷⁹ <https://www.projectcalico.org/comparing-kube-proxy-modes-iptables-or-ipvs/>



In DKP environments, the external load balancer must be configured without TLS termination.

6.9.6.2.1 In the Cloud

In cloud deployments, the load balancer is provided by the cloud provider. For example, in AWS, the service controller communicates with the AWS API to provision an ELB service which targets the cluster nodes.

6.9.6.2.2 On-premises

For an on-premises Pre-provisioned deployment, DKP ships with [MetalLB](#)⁶⁸⁰, which provides load-balancing services. The environments that use MetalLB are pre-provisioned and vSphere infrastructures.

For more information on how to configure MetalLB for these environments, refer to the following pages:

- <https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Pre-provisioned%20Configure%20MetalLB&linkCreation=true&fromPageId=273353624>⁶⁸¹
- [Configure MetalLB for a vSphere infrastructure](#) (see page 1282)

6.9.6.2.3 Custom Load Balancer for External Traffic

If you want to use a non-DKP load balancer for external traffic, see [External Load Balancer](#) (see page 1450).

6.9.7 Use Istio as a Microservice

Istio helps you manage cloud-based deployments by providing an open-source service mesh to connect, secure, control, and observe microservices.

This topic describes how to expose an application running on the DKP cluster using the LoadBalancer (layer-4) service type.

6.9.7.1 Prerequisites

Before you begin, verify the following:

- You must have access to a Linux, macOS, or Windows computer with a supported operating system version.
- You must have a properly deployed and running cluster.
- Determine the name of the workspace where you wish to perform the deployment. You can use the `dkp get workspaces` command to see the list of workspace names and their corresponding namespaces.

⁶⁸⁰ <https://metallb.universe.tf/concepts/>

⁶⁸¹ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273353624&linkCreation=true&spaceKey=DENT&title=%282.6%29+Pre-provisioned+Configure+MetalLB)

[fromPageId=273353624&linkCreation=true&spaceKey=DENT&title=%282.6%29+Pre-provisioned+Configure+MetalLB](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273353624&linkCreation=true&spaceKey=DENT&title=%282.6%29+Pre-provisioned+Configure+MetalLB)

- Set the `WORKSPACE_NAME` environment variable to the name of the workspace where the cluster is attached:

```
export WORKSPACE_NAME=<workspace_name>
```

6.9.7.2 Deploy Istio Using DKP

Follow these steps to prepare DKP to run Istio:

1. Review the [list of available applications \(see page 1760\)](#) to obtain the current *APP ID* and *version* for Istio and its [dependencies \(see page 648\)](#) – you need this information to execute the following commands.
2. Install [Istio's dependencies \(see page 648\)](#) with an `AppDeployment` resource. Replace the `<APPID>` and `<APPID-version>` variables with the information of the application:

i The required value for the `--app` flag consists of the *APP ID* and *version* in the `APPID-version` format.

```
dkp create appdeployment <APPID> --app <APPID-version> --workspace $
{WORKSPACE_NAME}
```

3. Install Istio. Replace `<APPID-version>` with the version of Istio you want to deploy:

```
dkp create appdeployment istio --app <APPID-version> --workspace $
{WORKSPACE_NAME}
```



- Create the resource in the workspace you just created, which instructs Kommander to deploy the `AppDeployment` to the `KommanderCluster`s in the same workspace.
- Observe that the `dkp create` command must be run with the `WORKSPACE_NAME` instead of the `WORKSPACE_NAMESPACE` flag.

6.9.7.2.1 Download the Istio Command Line Utility

Follow these steps to download and run Istio:

1. Pull a copy of the corresponding Istio command line to your system. Replace `<your_istio_version_here>` with the Istio version you want to deploy:

```
curl -L https://istio.io/downloadIstio |
ISTIO_VERSION=<your_istio_version_here> sh -
```

2. Change to the Istio directory and set your PATH environment variable by running the following commands:

```
cd istio*
export PATH=$PWD/bin:$PATH
```

3. Run the following `istioctl` command and view the subsequent output:

```
istioctl version
```

```
client version: <your_istio_version_here>
control plane version: <your_istio_version_here>
data plane version: <your_istio_version_here> (1 proxies)
```

6.9.7.2.2 Deploy a Sample Application on Istio

The Istio `bookinfo` sample application is composed of four separate microservices that demonstrate various Istio features.

1. Deploy the sample `bookinfo` application on the Kubernetes cluster by running the following commands:

IMPORTANT: Ensure your `dkp` configuration references the cluster where you deployed Istio by setting the `KUBECONFIG` environment variable, or using the `--kubeconfig` flag, [in accordance with Kubernetes conventions](#)⁶⁸².

```
kubectl apply -f <(istioctl kube-inject -f samples/bookinfo/platform/kube/
bookinfo.yaml)
kubectl apply -f samples/bookinfo/networking/bookinfo-gateway.yaml
```

2. Get the URL of the load balancer created on AWS for this service by running the following command:

```
kubectl get svc istio-ingressgateway -n istio-system
```

The command displays output similar to the following:

⁶⁸² <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

NAME	TYPE	CLUSTER-IP	EXTERNAL-IP
istio-ingressgateway	LoadBalancer	10.0.29.241	
a682d13086ccf11e982140acb7ee21b7-2083182676.us-west-2.elb.amazonaws.com			15020:30380/TCP,80:31380/TCP,443:31390/TCP,31400:31400/TCP,15029:30756/TCP,15030:31420/TCP,15031:31948/TCP,15032:32061/TCP,15443:31232/TCP
			110s

- Open a browser and navigate to the external IP address for the load balancer to access the application.

For example, the external IP address in the sample output is

`a682d13086ccf11e982140acb7ee21b7-2083182676.us-west-2.elb.amazonaws.com`, enabling you to access the application using the following URL: `http://a682d13086ccf11e982140acb7ee21b7-2083182676.us-west-2.elb.amazonaws.com/productpage`

- Follow the steps in the Istio [BookInfo Application](#)⁶⁸³ documentation to understand the different Istio features.

For more information, see [Istio's official documentation](#)⁶⁸⁴.

6.10 GPUs

This section describes NVIDIA GPU support on Kommander. DKP supported [NVIDIA driver](#)⁶⁸⁵ version is 470.x. GPUs, such as those made by AMD, are not currently supported. This document assumes familiarity with Kubernetes GPU support. More information about GPUs in AWS environment can be found in the [Advanced AWS \(see page 1110\)](#) section.

6.10.1 Kommander GPU Overview

GPU support on Kommander uses the [NVIDIA](#)⁶⁸⁶ GPU operator. Through the NVIDIA GPU operator application, Kommander configures the container runtime to run GPU containers, and installs all the necessary items to power up the NVIDIA GPU devices.

The following components provide NVIDIA GPU support on Kommander:

- `libnvidia-container` and `nvidia-container-runtime`: GPU Support in Kommander depends on the containerd runtime. `libnvidia-container` and `nvidia-container-runtime` fit between `containerd` and `runc`, simplifying the container runtime integration with the GPU.

⁶⁸³ <https://istio.io/docs/examples/bookinfo/>

⁶⁸⁴ <https://istio.io/latest/docs/>

⁶⁸⁵ <https://www.nvidia.com/Download/Find.aspx>

⁶⁸⁶ <https://github.com/NVIDIA/gpu-operator>

- [NVIDIA Device Plugin](#)⁶⁸⁷: Kommander makes use of NVIDIA GPUs using this Kubernetes device plugin. It allows GPU enabled containers to run on Kubernetes, tracking the number of available GPUs on each node and their health.
- [NVIDIA Data Center GPU Manager](#)⁶⁸⁸: Contains a Prometheus exporter that provides NVIDIA GPU metrics.

Kommander runs these components as daemonsets, making them easier to manage and upgrade across all GPU nodes.

For more information from NVIDIA, see the [Getting Started](#)⁶⁸⁹ page for NVIDIA.

6.10.2 GPU Prerequisites

6.10.2.1 Configure GPU for Kommander clusters

6.10.2.2 Prerequisites

Before you begin, you must:

- Ensure nodes provide an NVIDIA GPU.
- If you are using a public cloud service such as [AWS](#) (see page 1110), create an AMI with KIB using the instructions on the [KIB for GPU](#) (see page 1485) page.
- If you are deploying in a pre-provisioned environment, ensure that you have created the appropriate [secret for your GPU nodepool](#) (see page 240) and have uploaded the appropriate artifacts to each node. See the [GPU only steps section](#) (see page 0) on the Pre-provisioned Prerequisites Air-gapped page for additional information.



Specific instructions must be followed for enabling `nvidia-gpu-operator` depending on if you want to deploy the app on a Management cluster or a Attached or a Managed cluster.

- For instructions on enabling the NVIDIA platform application on a Management cluster, follow the instructions in the [NVIDIA Platform Application Management Cluster](#) (see page 970) section.
- For instructions on enabling the NVIDIA platform application on attached or managed clusters, follow the instructions in the [NVIDIA Platform Application Attached or Managed Cluster](#) (see page 972) section.

After `nvidia-gpu-operator` has been enabled depending on the cluster type, proceed to the **Select the correct Toolkit version for your NVIDIA GPU Operator** on each of those pages.

687 <https://github.com/NVIDIA/k8s-device-plugin>

688 <https://developer.nvidia.com/dcgm>

689 <https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/getting-started.html#chart-customization-options>

6.10.3 NVIDIA Platform Application Management Cluster

Instructions on enabling the NVIDIA platform application on a Management cluster

6.10.3.1 Enable NVIDIA Platform Application on Kommander for Management Cluster

If you intend to run applications that make use of GPU's on your cluster, you should install the NVIDIA GPU operator. To enable NVIDIA GPU support when installing Kommander on a management cluster, perform the following steps:

1. Create an installation configuration file:

```
dkp install kommander --init > install.yaml
```

2. Append the following to the apps section in the `install.yaml` file to enable Nvidia platform services.

```
apps:  
  nvidia-gpu-operator:  
    enabled: true
```

3. Install Kommander using the configuration file you created:

```
dkp install kommander --installer-config ./install.yaml --kubeconfig=${  
CLUSTER_NAME}.conf
```

In the previous command, the `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you set the context to **install Kommander on the right cluster**. For alternatives and recommendations around setting your context, refer to [Provide a context for your commands \(see page 87\)](#).

4. Proceed to the [Select the correct Toolkit version for your NVIDIA GPU Operator \(see page 973\)](#) section.



TIP: Sometimes, applications require a longer period of time to deploy, which causes the installation to time out. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.

6.10.3.2 Select the Correct Toolkit Version for your NVIDIA GPU Operator

The NVIDIA Container Toolkit allows users to run GPU accelerated containers. The toolkit includes a container runtime library and utilities to automatically configure containers to leverage NVIDIA GPU and must be configured correctly according to your base operating system.

6.10.3.2.1 Kommander (Management Cluster) Customization

1. Select the correct Toolkit version based on your OS:

The **NVIDIA Container Toolkit** allows users to run GPU accelerated containers. The toolkit includes a container runtime library and utilities to automatically configure containers to leverage NVIDIA GPU and must be configured correctly according to your base operating system.

Centos 7.9/RHEL 7.9:

If you're using Centos 7.9 or RHEL 7.9 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-centos7
```

RHEL 8.4/8.6 and SLES 15 SP3

If you're using RHEL 8.4/8.6 or SLES 15 SP3 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubi8
```

Ubuntu 18.04 and 20.04

If you're using Ubuntu 18.04 or 20.04 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your *Kommander Installer Configuration file* (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubuntu20.04
```

2. Install Kommander, using the configuration file you created:

```
dkp install kommander --installer-config ./install.yaml
```

In the previous command, the `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you set the context to install Kommander on the right cluster. For alternatives and recommendations around setting your context, refer to [Provide a context for your commands \(see page 87\)](#).

TIP: Sometimes, applications require a longer period of time to deploy, which causes the installation to time out. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.

6.10.4 NVIDIA Platform Application Attached or Managed Cluster

Instructions on enabling the NVIDIA platform application on attached or managed clusters

6.10.4.1 Enable NVIDIA Platform Application on Attached or Managed Clusters

If you intend to run applications that utilize GPU's on Attached or Managed clusters, you must enable the `nvidia-gpu-operator` platform application in the workspace.

To use the UI to enable the application, refer to the <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29919665/Platform+Applications#Customize-a-workspace%E2%80%99s-applications> (see page 0) page.

To use the CLI, refer to the [Deploy Platform Applications via CLI \(see page 642\)](#) page.

If only a subset of attached or managed clusters in the workspace are utilizing GPU's, refer to [Enable an Application per Cluster](#) (see page 658) on how to only enable the `nvidia-gpu-operator` on specific clusters.

After you have enabled the `nvidia-gpu-operator` app in the workspace on the necessary clusters, proceed to the next section.

6.10.4.2 Select the Correct Toolkit Version for your NVIDIA GPU Operator

The NVIDIA Container Toolkit allows users to run GPU accelerated containers. The toolkit includes a container runtime library and utilities to automatically configure containers to leverage NVIDIA GPU and must be configured correctly according to your base operating system.

6.10.4.2.1 Workspace (Attached and Managed clusters) Customization

Refer to [AppDeployment resources](#) (see page 625) for how to use the CLI to customize the platform application on a workspace.

If specific attached/managed clusters in the workspace require different configurations, refer to [Customize an Application per Cluster](#) (see page 660) for how to do this.

1. Select the correct Toolkit version based on your OS and create a `ConfigMap` with these configuration override values:

Centos 7.9/RHEL 7.9:

If you're using Centos 7.9 or RHEL 7.9 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your `install.yaml` to the following:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: nvidia-gpu-operator-overrides-attached
data:
  values.yaml: |
    toolkit:
      version: v1.10.0-centos7
EOF
```


RHEL 8.4/8.6 and SLES 15 SP3

If you're using RHEL 8.4/8.6 or SLES 15 SP3 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your `install.yaml` to the following:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: nvidia-gpu-operator-overrides-attached
data:
  values.yaml: |
    toolkit:
      version: v1.10.0-ubi8
EOF
```

Ubuntu 18.04 and 20.04

If you're using Ubuntu 18.04 or 20.04 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your `install.yaml` to the following:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: ${WORKSPACE_NAMESPACE}
  name: nvidia-gpu-operator-overrides-attached
data:
  values.yaml: |
    toolkit:
      version: v1.11.0-ubuntu20.04
EOF
```

- Note the name of this ConfigMap (`nvidia-gpu-operator-overrides-attached`) and use it to set the necessary `nvidia-gpu-operator` AppDeployment spec fields depending on the scope of the override. Alternatively, you can also use the UI to pass in the configuration overrides for the app per workspace or per cluster.

6.10.5 Kommander GPU Settings and Troubleshoot

6.10.5.1 Validate and troubleshoot GPU

6.10.5.2 Validate that the Application has Started Correctly

Run the following command to validate that your application has started correctly:

```
kubectl get pods -A | grep nvidia
```

The output should be similar to the following:

```
nvidia-container-toolkit-daemonset-7h2l5    1/1    Running    0    150m
nvidia-container-toolkit-daemonset-mm65g    1/1    Running    0    150m
nvidia-container-toolkit-daemonset-mv7xj    1/1    Running    0    150m
nvidia-cuda-validator-pdlz8                0/1    Completed  0    150m
nvidia-cuda-validator-r7qc4                0/1    Completed  0    150m
nvidia-cuda-validator-xvtqm                0/1    Completed  0    150m
nvidia-dcgm-exporter-9r6rl                 1/1    Running    1 (149m ago) 150m
nvidia-dcgm-exporter-hn6hn                 1/1    Running    1 (149m ago) 150m
nvidia-dcgm-exporter-j7g7g                 1/1    Running    0    150m
nvidia-dcgm-jpr57                           1/1    Running    0    150m
nvidia-dcgm-jwldh                           1/1    Running    0    150m
nvidia-dcgm-qg2vc                           1/1    Running    0    150m
nvidia-device-plugin-daemonset-2gv8h       1/1    Running    0    150m
nvidia-device-plugin-daemonset-tcmgk       1/1    Running    0    150m
nvidia-device-plugin-daemonset-vqj88       1/1    Running    0    150m
nvidia-device-plugin-validator-9xdqr        0/1    Completed  0    149m
nvidia-device-plugin-validator-jjhdr        0/1    Completed  0    149m
nvidia-device-plugin-validator-llxjk        0/1    Completed  0    149m
nvidia-operator-validator-9kzv4             1/1    Running    0    150m
nvidia-operator-validator-fvsr7             1/1    Running    0    150m
nvidia-operator-validator-qr9cj             1/1    Running    0    150m
```

If you are seeing errors, ensure that you set the container toolkit version appropriately based on your OS, as described in the previous section.

6.10.5.3 NVIDIA GPU Monitoring

Kommander uses the [NVIDIA Data Center GPU Manager](https://developer.nvidia.com/dcgm)⁶⁹⁰ to export GPU metrics towards Prometheus. By default, Kommander has a Grafana dashboard called `NVIDIA DCGM Exporter Dashboard` to monitor GPU metrics. This GPU dashboard is shown in Kommander's Grafana UI.

⁶⁹⁰ <https://developer.nvidia.com/dcgm>

6.10.5.4 NVIDIA MIG Settings

MIG stands for Multi-Instance-GPU. It is a mode of operation for future NVIDIA GPUs that allows the user to partition a GPU into a set of MIG devices. Each set appears to the software that is consuming them as a mini-GPU with a fixed partition of memory and a fixed partition of compute resources.

NOTE: MIG is only available for the following NVIDIA devices: H100, A100, and A30.

6.10.5.4.1 To Configure MIG

1. Set the MIG strategy according to your GPU topology.
 - `mig.strategy` should be set to `mixed` when MIG mode is not enabled on all GPUs on a node.
 - `mig.strategy` should be set to `single` when MIG mode is enabled on all GPUs on a node and they have the same MIG device types across all of them.

For the Management Cluster, this can be set at install time by modifying the Kommander configuration file to add configuration for the `nvidia-gpu-operator` application:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
  nvidia-gpu-operator:
    values: |
      mig:
        strategy: single
  ...
```

Or by modifying the `clusterPolicy` object for the GPU operator once it has already been installed.

2. Set the MIG profile for the GPU you are using. In our example, we are using the A30 GPU that supports the following MIG profiles:

```
4 GPU instances @ 6GB each
2 GPU instances @ 12GB each
1 GPU instance @ 24GB
```

Set the mig profile by labeling the node `$(NODE)` with the profile as in the example below:

```
kubectl label nodes ${NODE} nvidia.com/mig.config=all-1g.6gb --overwrite
```

3. Check the node labels to see if the changes were applied to your MIG enabled GPU node

```
kubectl get no -o json | jq .items[0].metadata.labels
```

```
"nvidia.com/mig.config": "all-1g.6gb",
  "nvidia.com/mig.config.state": "success",
  "nvidia.com/mig.strategy": "single"
```

4. Deploy a sample workload:

```
apiVersion: v1
kind: Pod
metadata:
  name: cuda-vector-add
spec:
  restartPolicy: OnFailure
  containers:
    - name: cuda-vectoradd
      image: "nvidia/samples:vectoradd-cuda11.2.1"
      resources:
        limits:
          nvidia.com/gpu: 1

  nodeSelector:
    "nvidia.com/gpu.product": NVIDIA-A30-MIG-1g.6gb
```

If the workload successfully finishes, then your GPU has been properly MIG partitioned.

6.10.5.5 Troubleshooting NVIDIA GPU Operator on Kommander

In case you run into any errors with NVIDIA GPU Operator, here are a couple commands you can run to troubleshoot:

1. Connect (using SSH or similar) to your GPU enabled nodes and run the `nvidia-smi` command. Your output should be similar to the following example:

```
[ec2-user@ip-10-0-0-241 ~]$ nvidia-smi
```

```

Thu Nov  3 22:52:59 2022
-----+-----
| NVIDIA-SMI 470.82.01    Driver Version: 470.82.01    CUDA Version: 11.4
|
|-----+-----
-----+-----
| GPU  Name          Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC
| Fan  Temp  Perf    Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M.
|                                         |              | GPU-Util  Compute M.
|                                         |              |              |   MIG M.
|-----+-----+-----+-----+-----+-----+-----+-----+-----|
|   0  Tesla T4             On   | 00000000:00:1E.0 Off  |
| N/A   54C    P8     11W / 70W |      0MiB / 15109MiB |      0%      Default
|                                         |              |              |
|                                         |              |              |   N/A
-----+-----
-----+-----
| Processes:
|
| GPU  GI  CI          PID  Type  Process name          GPU Memory
|      ID  ID
|                                         |
|-----+-----+-----+-----+-----+-----+-----+-----+-----|
| No running processes found
|
-----+-----

```

- Another common issue is having a misconfigured toolkit version, resulting in NVIDIA pods in a bad state

For example:

```

nvidia-container-toolkit-daemonset-jrqt2          1/1
Running                                           0          29s
nvidia-dcgm-exporter-b4mww                       0/1      Error
1 (9s ago)   16s
nvidia-dcgm-pqs8                                 0/1
CrashLoopBackOff    1 (13s ago)  27s
nvidia-device-plugin-daemonset-7fkzr            0/1      Init:0
/1                                           0          14s

```

```
nvidia-operator-validator-zxn4w 0/1
Init:CrashLoopBackOff 1 (7s ago) 11s
```

To modify the toolkit version, run the following commands to modify the `AppDeployment` for the `nvidia gpu operator` application:

- Provide the name of a `ConfigMap` with the custom configuration in the `AppDeployment` :

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: nvidia-gpu-operator
  namespace: kommander
spec:
  appRef:
    kind: ClusterApp
    name: nvidia-gpu-operator-1.11.1
  configOverrides:
    name: nvidia-gpu-operator-overrides
EOF
```

- Create the `ConfigMap` with the name provided in the previous step, which provides the custom configuration on top of the default configuration in the config map, set the version appropriately:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: kommander
  name: nvidia-gpu-operator-overrides
data:
  values.yaml: |
    toolkit:
      version: v1.10.0-centos7
EOF
```

3. If a node has an NVIDIA GPU installed and the `nvidia-gpu-operator` application is enabled on the cluster, but the node is still not accepting GPU workloads, it's possible that the nodes do not have a label that indicates there is an NVIDIA GPU present. By default the GPU operator will attempt to configure nodes with the following labels present, which are usually applied by the node feature discovery component:

```
"feature.node.kubernetes.io/pci-10de.present": "true",
```

```
"feature.node.kubernetes.io/pci-0302_10de.present": "true",
"feature.node.kubernetes.io/pci-0300_10de.present": "true",
```

If these labels are not present on a node that you know contains an NVIDIA GPU, you can manually label the node using the following command:

```
kubectl label node ${NODE} feature.node.kubernetes.io/
pci-0302_10de.present=true
```

6.10.5.6

Disable NVIDIA GPU Operator Platform Application on Kommander

1. Delete all GPU workloads on the GPU nodes where the NVIDIA GPU Operator platform application is present.
2. Delete the existing NVIDIA GPU Operator AppDeployment using the following command:

```
kubectl delete appdeployment -n kommander nvidia-gpu-operator
```

3. Wait for all NVIDIA related resources in the `Terminating` state to be cleaned up. You can check pod status with the following command:

```
kubectl get pods -A | grep nvidia
```



For information on how to delete nodepools, refer to [Pre-provisioned Create and Delete Node Pools](#) (see page 1245)

6.10.6 GPU Toolkit Versions

The **NVIDIA Container Toolkit** allows users to run GPU accelerated containers. The toolkit includes a container runtime library and utilities to automatically configure containers to leverage NVIDIA GPU and must be configured correctly according to your base operating system.

Centos 7.9/RHEL 7.9:

If you're using Centos 7.9 or RHEL 7.9 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your [Kommander Installer Configuration file](#) (see page 1422) or `<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-centos7
```

RHEL 8.4/8.6 and SLES 15 SP3

If you're using RHEL 8.4/8.6 or SLES 15 SP3 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your [Kommander Installer Configuration file](#) (see page 1422) or

`<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubi8
```

Ubuntu 18.04 and 20.04

If you're using Ubuntu 18.04 or 20.04 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter in your [Kommander Installer Configuration file](#) (see page 1422) or

`<kommander.yaml>` to the following:

```
kind: Installation
apps:
  nvidia-gpu-operator:
    enabled: true
    values: |
      toolkit:
        version: v1.13.1-ubuntu20.04
```

6.10.7 Enable GPU Usage After Installing DKP

If you want to enable your cluster to run GPU nodes after installing DKP, enable the correct Toolkit version for your operating system in the `nvidia-gpu-operator` AppDeployment.



If you have not installed the Kommander component of DKP yet, [set the Toolkit version in the Kommander Installer Configuration file](#) (see page 983) and skip this section.

1. Create a `ConfigMap` with the necessary configuration overrides to [set the correct Toolkit version](#) (see page 983). For example, if you're using Centos 7.9 or RHEL 7.9 as the base operating system for your GPU enabled nodes, set the `toolkit.version` parameter:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: kommander
  name: nvidia-gpu-operator-overrides
data:
  values.yaml: |
    toolkit:
      version: v1.10.0-centos7
EOF
```

Take the correct GPU Toolkit version from this page: [GPU Toolkit Versions](#) (see page 983)

2. Update the `nvidia-gpu-operator` `AppDeployment` in the `kommander` namespace to reference the `ConfigMap` you created:

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: nvidia-gpu-operator
  namespace: kommander
spec:
  appRef:
    name: nvidia-gpu-operator-1.11.1
    kind: ClusterApp
  configOverrides:
    name: nvidia-gpu-operator-overrides
EOF
```

6.11 Monitoring and Alerts

Using DKP you can monitor the state of the cluster and the health and availability of the processes running on the cluster. By default, Kommander provides monitoring services using a pre-configured monitoring stack based on the Prometheus open-source project and its broader ecosystem.

The default DKP monitoring stack:

- Provides in-depth monitoring of Kubernetes components and platform services.
- Includes a default set of Grafana dashboards to visualize the status of the cluster and its platform services.

- Supports predefined critical error and warning alerts. These alerts notify immediately if there is a problem with cluster operations or availability.

By incorporating Prometheus, Kommander visualizes all the exposed metrics from your different nodes, Kubernetes objects, and platform service applications running in your cluster. The default monitoring stack also enables you to add metrics from any of your deployed applications, making those applications part of the overall Prometheus metrics stream.

- [Recommendations](#) (see page 983)
- [Grafana Dashboards](#) (see page 985)
- [Cluster Metrics](#) (see page 988)
- [Configure Alerts Using AlertManager](#) (see page 988)
- [Centralized Monitoring](#) (see page 994)
- [Centralized Cost Monitoring](#) (see page 998)
- [Monitor Applications using Prometheus](#) (see page 1001)
- [Set Storage Capacity for Prometheus](#) (see page 1003)

6.11.1 Recommendations

Enterprise

Gov Advanced

Recommended settings for monitoring and collecting metrics for Kubernetes, platform services, and applications deployed on the cluster

D2iQ conducts routine performance testing of Kommander. The following table provides recommended settings, based on cluster size and increasing workloads, that maintain a healthy Prometheus monitoring deployment.



The resource settings reflect some settings but do not represent the exact structure to be used in the platform service configuration.

6.11.1.1 Prometheus

Cluster Size	Number of Pods	Number of Services	Resource settings
--------------	----------------	--------------------	-------------------

10	1k	250	<pre>resources: limits: cpu: 500m memory: 2192Mi requests: cpu: 100m memory: 500Mi storage: 35Gi</pre>
25	1k	250	<pre>resources: limits: cpu: 2 memory: 6Gi requests: cpu: 1 memory: 3Gi storage: 60Gi</pre>
50	1.5k	500	<pre>resources: limits: cpu: 7 memory: 28Gi requests: cpu: 2 memory: 8Gi storage: 100Gi</pre>
100	3k	1k	<pre>resources: limits: cpu: 12 memory: 50Gi requests: cpu: 10 memory: 48Gi storage: 100Gi</pre>

200	10k	3k	<pre>resources: limits: cpu: 20 memory: 80Gi requests: cpu: 15 memory: 50Gi storage: 100Gi</pre>
300	15k	6k	<pre>resources: limits: cpu: 35 memory: 150Gi requests: cpu: 25 memory: 120Gi storage: 100Gi</pre>

6.11.2 Grafana Dashboards

With Grafana, you can query and view collected metrics in easy-to-read graphs. Kommander ships with a set of default dashboards including:

- Kubernetes Components: API Server, Nodes, Pods, Kubelet, Scheduler, StatefulSets and Persistent Volumes
- Kubernetes USE method: Cluster and Nodes
- Calico
- etcd
- Prometheus

Find the complete list of default enabled dashboards [on GitHub](#)⁶⁹¹.

To disable all of the default dashboards, follow these steps to define an overrides ConfigMap:

1. Create a file named `kube-prometheus-stack-overrides.yaml` and paste the following YAML code into it to create the overrides ConfigMap:

⁶⁹¹ <https://github.com/prometheus-community/helm-charts/tree/main/charts/kube-prometheus-stack/templates/grafana/dashboards-1.14>

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: kube-prometheus-stack-overrides
  namespace: <your-workspace-namespace>
data:
  values.yaml: |
    ---
    grafana:
      defaultDashboardsEnabled: false

```

2. Use the following command to apply the YAML file:

```
kubectl apply -f kube-prometheus-stack-overrides.yaml
```

3. Edit the `kube-prometheus-stack` `AppDeployment` to replace the `spec.configOverrides.name` value with `kube-prometheus-stack-overrides`. (You can use the steps in the procedure, [Deploy an application with a custom configuration](#) (see page 660) as a guide.) When your editing is complete, the `AppDeployment` will resemble this code sample:

```

apiVersion: apps.kommander.d2iq.io/v1alpha2
kind: AppDeployment
metadata:
  name: kube-prometheus-stack
  namespace: <your-workspace-namespace>
spec:
  appRef:
    name: kube-prometheus-stack-34.9.3
    kind: ClusterApp
  configOverrides:
    name: kube-prometheus-stack-overrides

```

To access the Grafana UI, browse to the landing page and then search for the Grafana dashboard, for example, `https://<CLUSTER_URL>/dkp/grafana`.

6.11.2.1 Add Custom Dashboards

In Kommander you can define your own custom dashboards. You can use a few methods to [import dashboards](#)⁶⁹² to Grafana.

One method is to [use ConfigMaps to import dashboards](#)⁶⁹³. Below are steps on how to create a `ConfigMap` with your dashboard definition.

For simplicity, this section assumes the desired dashboard definition is in `json` format:

⁶⁹² <https://github.com/grafana/helm-charts/tree/main/charts/grafana#import-dashboards>

⁶⁹³ <https://github.com/grafana/helm-charts/tree/main/charts/grafana#sidecar-for-dashboards>

```
{
  "annotations": {
    "list": []
  },
  "description": "etcd sample Grafana dashboard with Prometheus",
  "editable": true,
  "gnetId": null,
  "hideControls": false,
  "id": 6,
  "links": [],
  "refresh": false,
  ...
}
```

After creating your custom dashboard json, insert it into a ConfigMap and save it as `etcd-custom-dashboard.yaml` :

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: etcd-custom-dashboard
  labels:
    grafana_dashboard: "1"
data:
  etcd.json: |
    {
      "annotations": {
        "list": []
      },
      "description": "etcd sample Grafana dashboard with Prometheus",
      "editable": true,
      "gnetId": null,
      "hideControls": false,
      "id": 6,
      "links": [],
      "refresh": false,
      ...
    }
}
```

Apply the ConfigMap, which automatically gets imported to Grafana using the [Grafana dashboard sidecar](#)⁶⁹⁴.

```
kubectl apply -f etcd-custom-dashboard.yaml
```

⁶⁹⁴ <https://github.com/grafana/helm-charts/tree/main/charts/grafana#sidecar-for-dashboards>

6.11.3 Cluster Metrics

The `kube-prometheus-stack` is deployed by default on the management cluster and attached clusters. This stack deploys the following Prometheus components to expose metrics from nodes, Kubernetes units, and running apps:

- `prometheus-operator`: orchestrates various components in the monitoring pipeline.
- `prometheus`: collects metrics, saves them in a time series database, and serves queries.
- `alertmanager`: handles alerts sent by client applications such as the Prometheus server.
- `node-exporter`: deployed on each node to collect the machine hardware and OS metrics.
- `kube-state-metrics`: simple service that listens to the Kubernetes API server and generates metrics about the state of the objects.
- `grafana`: monitors and visualizes metrics.
- `service monitors`: collects internal Kubernetes components.



DKP has a listener on the `metrics.k8s.io/v1beta1/nodes` resource, which updates your backend store when that value changes. We then poll that backend store every 5 seconds, so the metrics are updated in realtime every 5-seconds without the need to refresh your view.

A detailed description of the exposed metrics can be found [in the kube-state-metrics documentation on GitHub](#)⁶⁹⁵. The `service-monitors` collect internal Kubernetes components but can also be extended to monitor customer apps as explained in the section Monitor Applications, on this page below.

6.11.4 Configure Alerts Using AlertManager

To keep your clusters and applications healthy and drive productivity forward, you need to stay informed of all events occurring in your cluster. DKP helps you to stay informed of these events by using the `alertmanager` of the `kube-prometheus-stack`.

Kommander is configured with pre-defined alerts to monitor four specific events. You receive alerts related to:

- State of your nodes
- System services managing the Kubernetes cluster
- Resource events from specific system services
- Prometheus expressions exceeding some pre-defined thresholds

Some examples of the alerts currently available are:

- `CPUThrottlingHigh`
- `TargetDown`
- `KubeletNotReady`

⁶⁹⁵ <https://github.com/kubernetes/kube-state-metrics/tree/master/docs#exposed-metrics>

- KubeAPIDown
- CoreDNSDown
- KubeVersionMismatch

A complete list with all the pre-defined alerts can be found [on GitHub](#)⁶⁹⁶.

6.11.4.1 Prerequisites

- Determine the name of the workspace where you wish to perform the actions. You can use the `dkp get workspaces` command to see the list of workspace names and their corresponding namespaces.
- Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace where the cluster is attached:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

6.11.4.1.1 Configure Alert Rules

Use override `ConfigMaps` to configure alert rules.

You can enable or disable the default alert rules by providing the desired configuration in an overrides `ConfigMap`. For example, if you want to disable the default `node` alert rules, follow these steps to define an overrides `ConfigMap`:

1. Create a file named `kube-prometheus-stack-overrides.yaml` and paste the following YAML code into it to create the overrides `ConfigMap`:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: kube-prometheus-stack-overrides
  namespace: ${WORKSPACE_NAMESPACE}
data:
  values.yaml: |
    ---
    defaultRules:
      rules:
        node: false
```

2. Use the following command to apply the YAML file:

⁶⁹⁶ <https://github.com/prometheus-community/helm-charts/tree/main/charts/kube-prometheus-stack/templates/prometheus/rules-1.14>


```
kubectl apply -f kube-prometheus-stack-overrides.yaml
```

3. Edit the `kube-prometheus-stack` `AppDeployment` to replace the `spec.configOverrides.name` value with `kube-prometheus-stack-overrides`. (You can use the steps in the procedure, [Deploy an application with a custom configuration](#) (see page 991) as a guide.)

```
dkp edit appdeployment -n ${WORKSPACE_NAMESPACE} kube-prometheus-stack
```

After your editing is complete, the `AppDeployment` resembles this example:

```
apiVersion: apps.kommander.d2iq.io/v1alpha2
kind: AppDeployment
metadata:
  name: kube-prometheus-stack
  namespace: ${WORKSPACE_NAMESPACE}
spec:
  appRef:
    name: kube-prometheus-stack-34.9.3
    kind: ClusterApp
  configOverrides:
    name: kube-prometheus-stack-overrides
```

4. To disable all rules, create an overrides `ConfigMap` with this YAML code:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: kube-prometheus-stack-overrides
  namespace: ${WORKSPACE_NAMESPACE}
data:
  values.yaml: |
    ---
    defaultRules:
      create: false
```

5. Alert rules for the Velero platform service are turned off by default. You can enable them with the following overrides `ConfigMap`. They should be enabled only if the `velero` platform service is enabled. If platform services are disabled disable the alert rules to avoid alert misfires.

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: kube-prometheus-stack-overrides
```

```

namespace: ${WORKSPACE_NAMESPACE}
data:
  values.yaml: |
    ---
    mesosphereResources:
      rules:
        velero: true

```

6. To create a custom alert rule named `my-rule-name`, create the overrides ConfigMap with this YAML code:

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: kube-prometheus-stack-overrides
  namespace: ${WORKSPACE_NAMESPACE}
data:
  values.yaml: |
    ---
    additionalPrometheusRulesMap:
      my-rule-name:
        groups:
          - name: my_group
            rules:
              - record: my_record
                expr: 100 * my_record

```

After you set up your alerts, you can manage each alert using the Prometheus web console to mute or unmute firing alerts, and perform other operations. For more information about configuring `alertmanager`, see the [Prometheus website](#)⁶⁹⁷.

To access the Prometheus Alertmanager UI, browse to the landing page and then search for the Prometheus Alertmanager dashboard, for example `https://<CLUSTER_URL>/dkp/alertmanager`.

6.11.4.1.2 Notify Prometheus Alerts in Slack

To hook up the Prometheus `alertmanager` notification system, you need to overwrite the existing configuration.

1. The following file, named `alertmanager.yaml`, configures `alertmanager` to use the Incoming Webhooks feature of Slack (`slack_api_url: https://hooks.slack.com/services/<HOOK_ID>`) to fire all the alerts to a specific channel `#MY-SLACK-CHANNEL-NAME`.

```

global:

```

⁶⁹⁷ <https://prometheus.io/docs/alerting/configuration/>

```

resolve_timeout: 5m
slack_api_url: https://hooks.slack.com/services/<HOOK_ID>

route:
  group_by: ['alertname']
  group_wait: 2m
  group_interval: 5m
  repeat_interval: 1h

# If an alert isn't caught by a route, send it to slack.
receiver: slack_general
routes:
  - match:
      alertname: Watchdog
      receiver: "null"

receivers:
  - name: "null"
  - name: slack_general
    slack_configs:
      - channel: '#MY-SLACK-CHANNEL-NAME'
        icon_url: https://avatars3.githubusercontent.com/u/3380462
        send_resolved: true
        color: '{{ if eq .Status "firing" }}danger{{ else }}good{{ end }}'
        title: '{{ template "slack.default.title" . }}'
        title_link: '{{ template "slack.default.titlelink" . }}'
        pretext: '{{ template "slack.default.pretext" . }}'
        text: '{{ template "slack.default.text" . }}'
        fallback: '{{ template "slack.default.fallback" . }}'
        icon_emoji: '{{ template "slack.default.iconemoji" . }}'

templates:
  - '*.tpl'

```

- The following file, named `notification.tpl`, is a template that defines a pretty format for the fired notifications:

```

{{ define "__titlelink" }}
{{ .ExternalURL }}/#/alerts?receiver={{ .Receiver }}
{{ end }}

{{ define "__title" }}
[{{ .Status | toUpper }}]{{ if eq .Status "firing" }}:{{ .Alerts.Firing | len }}
{{ end }}] {{ .GroupLabels.SortedPairs.Values | join " " }}
{{ end }}

{{ define "__text" }}
{{ range .Alerts }}
{{ range .Labels.SortedPairs }}*{{ .Name }}*: `{{ .Value }}`
{{ end }} {{ range .Annotations.SortedPairs }}*{{ .Name }}*: {{ .Value }}

```

```

{{ end }} *source*: {{ .GeneratorURL }}
{{ end }}
{{ end }}

{{ define "slack.default.title" }}{{ template "__title" . }}{{ end }}
{{ define "slack.default.username" }}{{ template "__alertmanager" . }}{{ end }}
{{ define "slack.default.fallback" }}{{ template "slack.default.title" . }} |
{{ template "slack.default.titlelink" . }}{{ end }}
{{ define "slack.default.pretext" }}{{ end }}
{{ define "slack.default.titlelink" }}{{ template "__titlelink" . }}{{ end }}
{{ define "slack.default.iconemoji" }}{{ end }}
{{ define "slack.default.iconurl" }}{{ end }}
{{ define "slack.default.text" }}{{ template "__text" . }}{{ end }}

```

3. Finally, apply these changes to `alertmanager` as follows. Set `${WORKSPACE_NAMESPACE}` to the workspace namespace that `kube-prometheus-stack` is deployed in:

```

kubectl create secret generic -n ${WORKSPACE_NAMESPACE} \
  alertmanager-kube-prometheus-stack-alertmanager \
  --from-file=alertmanager.yaml \
  --from-file=notification.tpl \
  --dry-run=client --save-config -o yaml | kubectl apply -f -

```

6.11.4.1.3 Notify Prometheus Alerts in E-Mail

To configure the Prometheus `alertmanager` notification system to send an email for alerts, you need to overwrite the existing configuration. The steps below configure Alertmanager to send all configured alerts to a gmail account named `test@gmail.com`⁶⁹⁸.

1. Create a file named `alertmanager.yaml` with the following contents:

```

global:
  resolve_timeout: 5m
inhibit_rules: []
receivers:
- name: "null"
- name: test_gmail
  email_configs:
  - to: test@gmail.com
    from: test@gmail.com
    auth_username: test@gmail.com
    auth_password: password
    send_resolved: true
    require_tls: true
    smarthost: smtp.gmail.com:587

```

⁶⁹⁸ <mailto:test@gmail.com>

```

route:
  receiver: test_gmail
  group_by:
  - namespace
  group_interval: 5m
  group_wait: 30s
  repeat_interval: 12h
  routes:
  - matchers:
    - alertname =~ "InfoInhibitor|Watchdog"
    receiver: "null"
templates:
- /etc/alertmanager/config/*.tmpl

```

2. Apply these changes to `alertmanager` as follows. Set `${WORKSPACE_NAMESPACE}` to the workspace namespace that `kube-prometheus-stack` is deployed in (typically the `kommander` namespace):

```

kubectl create secret generic -n ${WORKSPACE_NAMESPACE} \
  alertmanager-kube-prometheus-stack-alertmanager \
  --from-file=alertmanager.yaml \
  --dry-run=client --save-config -o yaml | kubectl apply -f -

```

3. Allow some time for the configuration to take effect. You can then use the following command to verify that the configuration took effect:

```

kubectl exec -it alertmanager-kube-prometheus-stack-alertmanager-0 -n kommander \
  -- cat /etc/alertmanager/config_out/alertmanager.env.yaml

```

For more information on configuring email alerting, refer to the [Alertmanager documentation](#)⁶⁹⁹.

6.11.5 Centralized Monitoring

Enterprise

Gov Advanced

Monitor clusters, created with Kommander, on any attached cluster

Kommander provides centralized monitoring, in a multi-cluster environment, using the monitoring stack running on any attached clusters. Centralized monitoring is provided by default in every managed or attached cluster.

Managed or attached clusters are distinguished by a monitoring ID. The monitoring ID corresponds to the kube-system namespace UID of the cluster. To find a cluster's monitoring ID, you can go to the Clusters tab on the DKP UI (in the relevant workspace), or go to the **Clusters** page in the **Global** workspace:

⁶⁹⁹ <https://prometheus.io/docs/alerting/latest/configuration/>

```
https://<CLUSTER_URL>/dkp/kommander/dashboard/clusters
```

Select the `View Details` link on the attached cluster card, and then select the **Configuration** tab, and find the monitoring ID under **Monitoring ID (clusterId)**.

You may also search or filter by monitoring IDs on the Clusters page, linked above.

You can also run this `kubectl` command, **using the correct cluster's context or kubeconfig**, to look up the cluster's kube-system namespace UID to determine which cluster the metrics and alerts correspond to:

```
kubectl get namespace kube-system -o jsonpath='{.metadata.uid}'
```

6.11.5.1 Centralized Metrics

Managed and attached clusters collect and present metrics from all attached clusters remotely using Thanos. You can visualize these metrics in Grafana using a set of provided dashboards.

The [Thanos Query](#)⁷⁰⁰ component is installed on attached and managed clusters. Thanos Query queries the Prometheus instances on the attached clusters, using a Thanos sidecar running alongside each Prometheus container. Grafana is configured with Thanos Query as its datasource, and comes with pre-installed dashboards for a global view of all attached clusters. The `Thanos Query` dashboard is also installed, by default, to monitor the Thanos Query component.

NOTE: Metrics from clusters are read remotely from Kommander; they are not backed up. If a attached cluster goes down, Kommander no longer collects or presents its metrics, including past data.

You can access the centralized Grafana UI at:

```
https://<CLUSTER_URL>/dkp/kommander/monitoring/grafana
```

NOTE: This is a separate Grafana instance than the one installed on all attached clusters. It is dedicated specifically to components related to centralized monitoring.

Optionally, if you want to access the Thanos Query UI (essentially the Prometheus UI), the UI is accessible at:

```
https://<CLUSTER_URL>/dkp/kommander/monitoring/query/
```

You can also check that the attached cluster's Thanos sidecars are successfully added to Thanos Query by going to:

```
https://<CLUSTER_URL>/dkp/kommander/monitoring/query/stores
```

The preferred method to view the metrics for a specific cluster is to go directly to that cluster's Grafana UI.

⁷⁰⁰ <https://thanos.io/v0.5/components/query/#query>

6.11.5.1.1 Adding Custom Dashboards

You can also define custom dashboards for centralized monitoring on Kommander. There are a few methods to [import dashboards](https://github.com/mesosphere/charts/tree/master/stable/grafana#import-dashboards)⁷⁰¹ to Grafana. For simplicity, assume the desired dashboard definition is in `json` format:

```
{
  "annotations":
  ...
  # Complete json file here
  ...
  "title": "Some Dashboard",
  "uid": "abcd1234",
  "version": 1
}
```

After creating your custom dashboard json, insert it into a ConfigMap and save it as `some-dashboard.yaml`:

```
apiVersion: v1
kind: ConfigMap
metadata:
  name: some-dashboard
  labels:
    grafana_dashboard_kommander: "1"
data:
  some_dashboard.json: |
    {
      "annotations":
      ...
      # Complete json file here
      ...
      "title": "Some Dashboard",
      "uid": "abcd1234",
      "version": 1
    }
```

Apply the ConfigMap, which will automatically get imported to Grafana via the Grafana dashboard sidecar:

```
kubectl apply -f some-dashboard.yaml
```

⁷⁰¹ <https://github.com/mesosphere/charts/tree/master/stable/grafana#import-dashboards>


6.11.5.2 Centralized Alerts

A centralized view of alerts, from attached clusters, is provided using an alert dashboard called [Karma](#)⁷⁰². Karma aggregates all alerts from the Alertmanagers running in the attached clusters, allowing you to visualize these alerts on one page. Using the Karma dashboard, you can get an overview of each alert and filter by alert type, cluster, and more.

 Silencing alerts using the Karma UI is currently not supported.

You can access the Karma dashboard UI at:

```
https://<CLUSTER_URL>/dkp/kommander/monitoring/karma
```

 When there are no attached clusters, the Karma UI displays an error message `Get https://placeholder.invalid/api/v2/status: dial tcp: lookup placeholder.invalid on 10.0.0.10:53: no such host`. This is expected, and the error disappears when clusters are connected.

6.11.5.2.1 Federating Prometheus Alerting Rules

You can define additional [Prometheus alerting rules](#)⁷⁰³ on attached and managed clusters and federate them to all of the attached clusters by following these instructions. To use these instructions you must [install the kubefedctl CLI](#)⁷⁰⁴.

1. Enable the PrometheusRule type for federation.

```
kubefedctl enable PrometheusRules --kubefed-namespace kommander
```

2. Modify the existing alertmanager configuration.

```
kubectl edit PrometheusRules/kube-prometheus-stack-alertmanager.rules -n kommander
```

⁷⁰² <https://github.com/primitive/karma>

⁷⁰³ https://prometheus.io/docs/prometheus/latest/configuration/alerting_rules/

⁷⁰⁴ <https://github.com/kubernetes-sigs/kubefed/blob/master/docs/installation.md#kubefedctl-cli>

3. Append a sample rule.

```
- alert: MyFederatedAlert
  annotations:
    message: A custom alert that will always fire.
  expr: vector(1)
  labels:
    severity: warning
```

4. Federate the rules you just modified.

```
kubefedctl federate PrometheusRules kube-prometheus-stack-alertmanager.rules --
kubefed-namespace kommander -n kommander
```

5. Ensure that the clusters selection (`status.clusters`) is appropriately set for your desired federation strategy and check the [propagation status](#)⁷⁰⁵.

```
kubectl get federatedprometheusrules kube-prometheus-stack-alertmanager.rules
-n kommander -oyaml
```

6.11.6 Centralized Cost Monitoring

Enterprise

Gov Advanced

Monitoring costs of all attached clusters with Kubecost

[Kubecost](#)⁷⁰⁶, running on Kommander, provides centralized cost monitoring for all attached clusters. This feature, installed by default in the management cluster, provides a centralized view of Kubernetes resources used on all attached clusters.



By default, up to 15 days of cost metrics are retained, with no backup to an external store.


⁷⁰⁵ <https://github.com/kubernetes-sigs/kubefed/blob/master/docs/userguide.md#propagation-status>

⁷⁰⁶ <https://kubecost.com/>

6.11.6.1 Adding a Kubecost License Key to your Clusters

By default, DKP is deployed with the free version of Kubecost, which provides cost monitoring for individual clusters and metric retention for up to 15 days. Licensed plans from Kubecost with additional features are also available. For more information, refer to the [Kubecost Pricing](#)⁷⁰⁷ page.

The following instructions provide information on how you can add a Kubecost license key to your clusters if you have purchased one.

 The license key must be applied to the Centralized Kubecost application running on the Management cluster.


1. Obtain license key from Kubecost. For more information, refer to the [Kubecost Pricing](#)⁷⁰⁸ page.
2. Access the `centralized-kubecost` dashboard with the following link:

```
https://<CLUSTER_URL>/dkp/kommander/kubecost/frontend/
```

3. From the dashboard, select the **Settings** icon, then select “Add License Key”.

NOTE: Alternatively, the dashboard can be accessed via the following link: `https://<CLUSTER_URL>/dkp/kommander/kubecost/frontend/settings`

4. Input your license key, then select “Update.”
5. After the license key has been added, the licensed features become available in the Kubecost UI.

 There is a free trial of the Kubecost Enterprise plan that allows you to preview enterprise features for 30 days. To access this, select “Start Free Trial” in the **Settings** pane.

6.11.6.1.1 Considerations when Adding a License Key

Until you add a license key, Kubecost caches the context of the cluster you first navigate from. This means that if you navigate to the Kubecost UI via the dashboard button in the Application Dashboards tab of any cluster, you will not be able to access the centralized Kubecost UI until you clear your browser cookies/cache.

⁷⁰⁷ <https://www.kubecost.com/pricing/>

⁷⁰⁸ <https://www.kubecost.com/pricing/>

6.11.6.2 Centralized Costs

Using Thanos, the management cluster collects cost metrics remotely from each attached cluster. Costs from the last day and the last 7 days are displayed for each cluster, workspace, and project in the respective DKP UI pages. Further cost analysis and details can be found in the centralized KubeCost UI running on Kommander, at:

```
https://<CLUSTER_URL>/dkp/kommander/kubecost/frontend/detail.html#&agg=cluster
```

For more information on cost allocation metrics and how to navigate this view in the KubeCost UI, see the [KubeCost docs on Kubernetes Cost Allocation](#)⁷⁰⁹.

To identify the clusters in KubeCost, use the cluster's monitoring ID. The monitoring ID corresponds to the kube-system namespace UID of the cluster. To find the cluster's monitoring ID, select the **Clusters** tab on the DKP UI in the relevant workspace, or go to the **Clusters** page in the **Global** workspace:

```
https://<CLUSTER_URL>/dkp/kommander/dashboard/clusters
```

Select **View Details** on the attached cluster card. Select the **Configuration** tab, and find the monitoring ID under **Monitoring ID (clusterId)**.

You can also search or filter by monitoring ID on the **Clusters** page.

To look up a cluster's kube-system namespace UID directly using the CLI, run the following kubectl command, **using the cluster's context or kubeconfig**.

```
kubectl get namespace kube-system -o jsonpath='{.metadata.uid}'
```

6.11.6.2.1 KubeCost

KubeCost integrates directly with the Kubernetes API and cloud billing APIs to give you real-time visibility into your Kubernetes spend and cost allocation. By monitoring your Kubernetes spend across clusters, you can avoid overspend caused by uncaught bugs or oversights. With a cost monitoring solution in place you can realize the full potential and cost of these resources and avoid over-provisioning resources.

To customize pricing and out of cluster costs for AWS, you must apply these settings using the KubeCost UI running on each attached cluster. You can access the attached cluster's KubeCost Settings page at:

```
https://<MANAGED_CLUSTER_URL>/dkp/kubecost/frontend/settings.html
```



Make sure you access the cluster's KubeCost UI linked above, not the centralized KubeCost UI running on the Kommander management cluster.

⁷⁰⁹ <https://docs.kubecost.com/using-kubecost/getting-started/cost-allocation>

6.11.6.2.1.1 AWS

For more accurate AWS Spot pricing, follow [these steps](#)⁷¹⁰ to configure a data feed for the AWS Spot instances.

To allocate out of cluster costs for AWS, visit [this guide](#)⁷¹¹.

6.11.6.2.2 Grafana dashboards

A set of Grafana dashboards providing visualization of cost metrics is provided in the centralized Grafana UI:

```
https://<CLUSTER_URL>/dkp/kommander/monitoring/grafana
```

These dashboards give a global view of accumulated costs from all attached clusters. From the navigation in Grafana, you can find these dashboards by selecting those tagged with `cost`, `metrics`, and `utilization`.

6.11.6.3 Related Information

For information on related topics or procedures, refer to the following:

- [Kubecost documentation](#)⁷¹²

6.11.7 Monitor Applications using Prometheus

Before attempting to monitor your own applications, you should be familiar with the Prometheus conventions for exposing metrics. In general, there are two key recommendations:

- You should expose metrics using an HTTP endpoint named `/metrics`.
- The metrics you expose must be in a format that Prometheus can consume.

By following these conventions, you ensure that your application metrics can be consumed by Prometheus itself or by any Prometheus-compatible tool that can retrieve metrics, using the Prometheus client endpoint.

The `kube-prometheus-stack` for Kubernetes provides easy monitoring definitions for Kubernetes services and deployment and management of Prometheus instances. It provides a Kubernetes resource called `ServiceMonitor`.

By default, the `kube-prometheus-stack` provides the following service monitors to collect internal Kubernetes components:

- `kube-apiserver`
- `kube-scheduler`

⁷¹⁰ <https://docs.kubecost.com/using-kubecost/getting-started/spot-checklist#implementing-spot-nodes-in-your-cluster>

⁷¹¹ <https://docs.kubecost.com/install-and-configure/advanced-configuration/cloud-integration/aws-cloud-integrations/aws-out-of-cluster>

⁷¹² <https://docs.kubecost.com/>

- kube-controller-manager
- etcd
- kube-dns/coredns
- kube-proxy

The operator is in charge of iterating over all of these `ServiceMonitor` objects and collecting the metrics from these defined components.

The following example illustrates how to retrieve application metrics. In this example, there are:

- Three instances of a simple app named `my-app`
- The sample app listens and exposes metrics on port 8080
- The app is assumed to already be running

To prepare for monitoring of the sample app, create a service that selects the pods that have `my-app` as the value defined for their app label setting.

The service object also specifies the port on which the metrics are exposed. The `ServiceMonitor` has a label selector to select services and their underlying endpoint objects. For example:

```
kind: Service
apiVersion: v1
metadata:
  name: my-app
  namespace: my-namespace
  labels:
    app: my-app
spec:
  selector:
    app: my-app
  ports:
  - name: metrics
    port: 8080
```

This service object is discovered by a `ServiceMonitor`, which defines the selector to match the labels with those defined in the service. The app label must have the value `my-app`.

In this example, in order for `kube-prometheus-stack` to discover this `ServiceMonitor`, add a specific label `prometheus.kommander.d2iq.io/select: "true"` in the `yaml`:

```
apiVersion: monitoring.coreos.com/v1
kind: ServiceMonitor
metadata:
  name: my-app-service-monitor
  namespace: my-namespace
  labels:
    prometheus.kommander.d2iq.io/select: "true"
spec:
```

```

selector:
  matchLabels:
    app: my-app
endpoints:
- port: metrics

```

In this example, you would modify the Prometheus settings to have the operator collect metrics from the service monitor by appending the following configuration to the overrides ConfigMap:

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: kube-prometheus-stack-overrides
  namespace: ${WORKSPACE_NAMESPACE}
data:
  values.yaml: |
    ---
    prometheus:
      additionalServiceMonitors:
        - name: my-app-service-monitor
          selector:
            matchLabels:
              app: my-app
            namespaceSelector:
              matchNames:
                - my-namespace
          endpoints:
            - port: metrics
              interval: 30s

```

Official documentation about using a `ServiceMonitor` to monitor an app with the Prometheus-operator on Kubernetes can be found [on this GitHub repository](#)⁷¹³.

6.11.8 Set Storage Capacity for Prometheus

Follow the steps in this page to set a specific storage capacity for Prometheus.

When defining the requirements of a cluster, you can specify the capacity and resource requirements of Prometheus by modifying the settings in the overrides ConfigMap definition as shown below:

```

apiVersion: v1
kind: ConfigMap
metadata:
  name: kube-prometheus-stack-overrides
  namespace: ${WORKSPACE_NAMESPACE}
data:

```

⁷¹³ <https://github.com/coreos/prometheus-operator/blob/master/Documentation/user-guides/getting-started.md#related-resources>

```

values.yaml: |
---
prometheus:
  prometheusSpec:
    resources:
      limits:
        cpu: "4"
        memory: "8Gi"
      requests:
        cpu: "2"
        memory: "6Gi"
    storageSpec:
      volumeClaimTemplate:
        spec:
          resources:
            requests:
              storage: "100Gi"

```

6.12 Storage for Applications

DKP ships with a Rook Ceph cluster, that is used as the primary blob storage for various DKP components in the logging stack and backups.

The pages in this section provide an overview of the Rook Ceph application in DKP, including information about its components, resource requirements, storage configuration information, and dashboard.

- [Rook Ceph in DKP \(see page 1004\)](#)
- [BYOS \(Bring Your Own Storage\) to DKP Clusters \(see page 1011\)](#)

6.12.1 Rook Ceph in DKP

DKP ships with a Rook Ceph cluster, that is used as the primary blob storage for various DKP components in the logging stack, backups, and DKP Insights.

The pages in this section provide an overview of the Rook Ceph application in DKP, including information about its components, resource requirements, storage configuration information, and dashboard.



The Ceph instance installed by DKP is intended only for use by [DKP Insights](#)⁷¹⁴ and the logging stack and `velero` platform applications.

If you have an instance of Ceph that is managed outside of the DKP lifecycle, see [Bring Your Own Storage to DKP Clusters \(see page 1011\)](#).


- [Rook Ceph in DKP - Prerequisites \(see page 1005\)](#)

⁷¹⁴ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274890970/2.6+DKP+Insights+Overview>

- [Rook Ceph Configuration](#) (see page 1006)
- [Rook Ceph Dashboard](#) (see page 1010)

6.12.1.1 Rook Ceph in DKP - Prerequisites

Important information you should know prior to using Rook Ceph in DKP

 The Ceph instance installed by DKP is intended only for use by the logging stack and `velero` platform applications.


If you have an instance of Ceph that is managed outside of the DKP lifecycle, see [Bring Your Own Storage to DKP Clusters](#) (see page 1011).

If you do not plan on using any of the logging stack components such as `grafana-loki` or `velero` for backups, then you do not need Rook Ceph for your installation and you can disable it by adding the following to your installer config file:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
  ...
  ...
  grafana-loki:
    enabled: false
  ...
  rook-ceph:
    enabled: false
  rook-ceph-cluster:
    enabled: false
  ...
  velero:
    enabled: false
  ...
```

You must enable `rook-ceph` and `rook-ceph-cluster` if any of the following is true:

- If `grafana-loki` is enabled.
- If `velero` is enabled. If you applied config overrides for `velero` to use a storage that is [external to your cluster](#), (see page 878) then you do not need Ceph to be installed.

 For more information about Ceph, refer to the following:

- [Intro to Ceph – Ceph Documentation](#)⁷¹⁵
- [Rook – Rook Ceph Documentation](#)⁷¹⁶

i See [Rook Ceph Configuration](#) (see page 1006) for information on how you can configure Ceph for you Ceph Environment.

6.12.1.2 Rook Ceph Configuration

This page contains information about configuring Rook Ceph in your DKP Environment.

ⓘ The Ceph instance installed by DKP is intended only for use by the logging stack, DKP Insights and `veLero` platform applications.

If you have an instance of Ceph that is managed outside of the DKP lifecycle, see [Bring Your Own Storage to DKP Clusters](#) (see page 1011).

If you intend to use Ceph in conjunction with DKP Insights, see [DKP Insights BYOS \(Bring Your Own Storage\) to Insights](#)⁷¹⁷

6.12.1.2.1 Components of a Rook Ceph Cluster

Ceph supports creating clusters in different modes as listed in [CephCluster CRD - Rook Ceph Documentation](#)⁷¹⁸. DKP, specifically is shipped with a PVC Cluster, as documented in [PVC Storage Cluster - Rook Ceph Documentation](#)⁷¹⁹. It is recommended to use the PVC mode to keep the deployment and upgrades simple and agnostic to technicalities with node draining.

Ceph cannot be your CSI Provisioner when installing in PVC mode as Ceph relies on an existing CSI provisioner to bind the PVCs created by it. It is possible to use Ceph as your CSI provisioner, but that is outside the scope of this document. If you have an instance of Ceph that acts as the CSI Provisioner, then it is possible to reuse it for your DKP Storage needs. See [BYOS \(Bring Your Own Storage\) to DKP Clusters](#) (see page 1011) for information on reusing existing Ceph.

When you create `AppDeployments` for `rook-ceph` and `rook-ceph-cluster` platform applications results in the deployment of various components as listed in the following diagram:


⁷¹⁵ <https://docs.ceph.com/en/quincy/start/intro/>

⁷¹⁶ <https://rook.io/docs/rook/v1.10/Getting-Started/intro/>

⁷¹⁷ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/225214512/DKP+Insights+BYOS+%28Bring+Your+Own+Storage%29+to+Insights>

⁷¹⁸ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/>

⁷¹⁹ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/#pvc-storage-only-for-monitors>



6 Rook Ceph Cluster Components

Items highlighted in green are user-facing and configurable.

■ Please refer to [Rook Ceph Storage Architecture](#)⁷²⁰ and [Ceph Architecture](#)⁷²¹ for an in-depth explanation of the inner workings of the components outlined in the above diagram.

For additional details about the data model, refer to the [Rook Ceph Data Model](#)⁷²² page.

6.12.1.2.2 Resource Requirements

The following is a non-exhaustive list of the resource requirements for long running components of Ceph:

720 <https://rook.io/docs/rook/v1.10/Getting-Started/storage-architecture/>
 721 <https://docs.ceph.com/en/quincy/architecture/>
 722 <https://github.com/rook/rook/blob/release-1.10/design/ceph/data-model.md>

Type	Resources	Total
CPUs	100m x # of mgr instances (default 2) 250m x # of mon instances (default 3) 250m x # of osd instances (default 4) 100m x # of crashcollector instances (Daemonset i.e., # of nodes) 250m x # of rados gateway replicas (default 2)	~2000m CPU
Memory	512Mi x # of mgr instances (default 2) 512Gi x # of mon instances (default 3) 1Gi x # of osd instances (default 4) 500Mi x # of rados gateway replicas (default 2)	~8Gi Memory
Disk	4 x 40Gi PVCs with <code>Block</code> mode for <code>ObjectStorageDaemons</code> ⁷²³ 3 x 10Gi PVCs with <code>Block</code> or <code>FileSystem</code> mode for <code>Mons</code> ⁷²⁴	190Gi

Your default `StorageClass` should support creation of `PersistentVolume`s that satisfy the `PersistentVolumeClaim`s created by Ceph with `volumeMode: Block`.

6.12.1.2.3 Ceph Storage Configuration

Ceph is highly configurable and can support Replication or Erasure Coding to ensure [data durability](#)⁷²⁵. DKP is configured to use Erasure Coding for maximum efficiency.

Primer on Replication Strategies

Replication and Erasure Coding are the two primary methods for storing data in a durable fashion in any distributed system.

6.12.1.2.3.1 Replication⁷²⁶

- For a replication factor of N, data has N copies (including the original copy)
- Smallest possible replication factor is 2 (usually this means 2 storage nodes).
 - With replication factor of 2, data has 2 copies and this tolerates loss of one copy of data.
- Storage efficiency: $(1/N) * 100$ percentage. For example,
 - If `N=2`, then efficiency is `50%`.
 - If `N=3`, then efficiency is `33%` so on.

⁷²³ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

⁷²⁴ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#mon-settings>

⁷²⁵ [https://en.wikipedia.org/wiki/Durability_\(database_systems\)](https://en.wikipedia.org/wiki/Durability_(database_systems))

⁷²⁶ [https://en.wikipedia.org/wiki/Replication_\(computing\)](https://en.wikipedia.org/wiki/Replication_(computing))

- Fault Tolerance : $N-1$ nodes can be lost without loss of data. For example,
 - If $N=2$, then atmost 1 node can be lost without data loss.
 - If $N=3$, then atmost 2 nodes can be lost without data loss and so on.

6.12.1.2.3.2 Erasure Coding⁷²⁷

- Slices an object into k data fragments and computes m parity fragments. The erasure coding scheme gaurentees that data can be recreated using any k fragments out of $k+m$ fragments.
- The $k + m = n$ fragments are spread across ($>=n$) Storage Nodes to offer durability.
- Since k out of n fragments (could be parity or could be data fragments) are needed for recreation of data, at most m fragments can be lost without loss of data.
- The smallest possible count is $k = 2$, $m = 1$ i.e., $n = k + m = 3$. This works only if there are at least $n = 3$ storage nodes.
- Storage efficiency: $k / (k+m) * 100$ percentage. For example,
 - If $k=2$, $m=1$, then efficiency is 67%
 - If $k=3$, $m=1$, then efficiency is 75% and so on.
- Fault Tolerance: m nodes can be lost without loss of data. For example:
 - If $k=3$, $m=1$ then atmost 1 out of 4 nodes can be lost without data loss.
 - If $k=4$, $m=2$ then atmost 2 out of 6 nodes can be lost without data loss and so on.

The default configuration creates a `CephCluster` that creates 4 x `PersistentVolumeClaims` of 40G each, resulting in 160G of raw storage. Erasure coding ensures durability with $k=3$ data bits and $m=1$ parity bits. This gives a storage efficiency of 75% (refer to the primer above for calculation), which means 120G of disk space is available for consumption by services like `grafana-loki` , `project-grafana-loki` , and `velero` .

It is possible to override replication strategy for logging stack (`grafana-loki`) and `velero` backups.

Refer to the default configmap for the `CephObjectStore` at [services/rook-ceph-cluster/1.10.3/defaults/cm.yaml#L126-L175](https://github.com/mesosphere/kommander-applications/blob/v2.4.0/services/rook-ceph-cluster/1.10.3/defaults/cm.yaml#L126-L175)⁷²⁸ and override the replication strategy according to your needs by referring to [CephObjectStore CRD](https://www.rook.io/docs/rook/v1.11/CRDs/Object-Storage/ceph-object-store-crd/)⁷²⁹ documentation.

For more information about configuring storage in Rook Ceph, refer to the following pages:

- <https://www.rook.io/docs/rook/v1.11/Storage-Configuration/Advanced/ceph-osd-mgmt/> - for general information on how to configure Object Storage Daemons (OSDs).

⁷²⁷ https://en.wikipedia.org/wiki/Erasure_code

⁷²⁸ <https://github.com/mesosphere/kommander-applications/blob/v2.4.0/services/rook-ceph-cluster/1.10.3/defaults/cm.yaml#L126-L175>

⁷²⁹ <https://www.rook.io/docs/rook/v1.11/CRDs/Object-Storage/ceph-object-store-crd/>

- <https://www.rook.io/docs/rook/v1.11/Storage-Configuration/Advanced/ceph-configuration/?h=expa#auto-expansion-of-osds> - for information on how to set up auto-expansion of OSDs.

See Also:

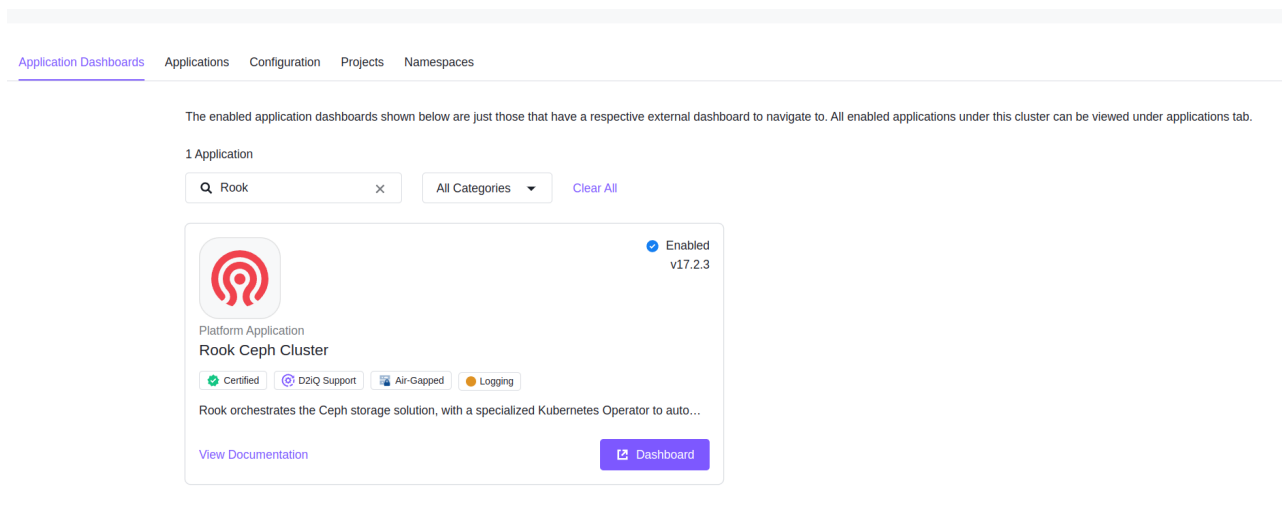
[Rook Ceph in DKP - Prerequisites](#) (see page 1008)

[Rook Ceph Dashboard](#) (see page 1010)

6.12.1.3 Rook Ceph Dashboard

6.12.1.3.1 Rook Ceph Dashboard Login

Rook Ceph Cluster provides a Dashboard which can be accessed from the UI in the Application Dashboards tab inside Kommander.



7 Rook Ceph Dashboard Link in the UI


The dashboard can be used to view current cluster health and logs.

6.12.1.3.1.1 How to Access the Rook Ceph Dashboard

1. Go to the applications dashboard
2. Select the Dashboard button
3. Username is `admin`
4. To retrieve your password, in the command line **and using the kubeconfig of the Kubernetes cluster you have Rook Ceph deployed to**, run the following command:
NOTE: Set the NAMESPACE variable according to your environment (`kommander` on management cluster or workspace namespace on attached clusters and managed clusters).

```
kubectl get secret -n ${NAMESPACE} rook-ceph-dashboard-password -o go-template="{{.data.password|base64decode}}"
```

- Copy the password and paste it into the UI to access the dashboard. After successful login, a Healthy Rook Ceph Cluster dashboard will look similar to:



8 Rook Ceph Cluster Dashboard

6.12.2 BYOS (Bring Your Own Storage) to DKP Clusters

You can use Ceph as the CSI Provisioner in some environments. For environments where Ceph was installed before installing DKP, you can reuse your existing Ceph installation to satisfy the storage requirements of DKP Applications.

i This guide assumes you have a Ceph cluster that is not managed by DKP. Refer to [Rook Ceph Configuration \(see page 1009\)](#) for information on how to configure the Ceph instance installed by DKP for use by DKP platform applications.

6.12.2.1 Disable DKP Managed Ceph

Disable `rook-ceph` in your installer config since the default config of DKP has already installed a Ceph Cluster.

Disable `rook-ceph` in the installer config to prevent DKP from installing a Ceph cluster:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
```


```
apps:
  rook-ceph:
    enabled: false
  rook-ceph-cluster:
    enabled: false
  ...
  ...
```

The DKP instances of `velero` and `grafana-loki` rely on the storage provided by Ceph. Before installing the Kommander component of DKP, be sure to configure appropriate Ceph resources for their usage as detailed in the next section.

6.12.2.2 Creating DKP Compatible Ceph Resources

This section walks you through the creation of `CephObjectStore` and then a set of `ObjectBucketClaims`, which can be consumed by either `velero` and `grafana-loki`.

Typically, Ceph is installed in the `rook-ceph` namespace, which is the default namespace if you have followed the [Quickstart - Rook Ceph Documentation](#)⁷³⁰ guide.

 This guide assumes your Ceph instance is installed in the `rook-ceph` namespace. Configure the variable `CEPH_NAMESPACE` in subsequent steps as it is applicable to your environment.

6.12.2.2.1 Creating `CephObjectStore`

There are two ways to install Ceph:

- Using [Helm Charts](#) (see page 1014)
- [Directly applying Kubernetes manifests](#) (see page 1014)

6.12.2.2.1.1 Helm Chart Values

This section is relevant if you have installed Ceph using `helm install` or some other managed Helm resource mechanism.

If you have applied any configuration overrides to your Rook Ceph operator, ensure it was deployed with `currentNamespaceOnly` set to `false`⁷³¹ (It is the default value, so unless you have applied any overrides, it will be `false` by default). This ensures that the Ceph Operator in the `rook-ceph` namespace is able to monitor and manage resources in other namespaces such as `kommander`.

⁷³⁰ <https://www.rook.io/docs/rook/v1.10/Getting-Started/quickstart/#create-a-ceph-cluster>

⁷³¹ <https://www.rook.io/docs/rook/v1.10/Helm-Charts/operator-chart/>

Ensure the following [configuration](#)⁷³² for `rook-ceph` [helm chart](#)⁷³³:

```
# This is the default value, so need to overwrite if you are just using the defaults
as-is
currentNamespaceOnly: false
```

You must enable the following [configuration overrides](#)⁷³⁴ for the `rook-ceph-cluster` [helm chart](#)⁷³⁵:

```
cephObjectStores:
  - name: dkp-object-store
    # see https://github.com/rook/rook/blob/master/Documentation/CRDs/Object-Storage/
    # ceph-object-store-crd.md#object-store-settings for available configuration
    spec:
      metadataPool:
        # The failure domain: osd/host/(region or zone if available) - technically
        # also any type in the crush map
        failureDomain: osd
        # Must use replicated pool ONLY. Erasure coding is not supported.
        replicated:
          size: 3
      dataPool:
        # The failure domain: osd/host/(region or zone if available) - technically
        # also any type in the crush map
        failureDomain: osd
        # Data pool can use either replication OR erasure coding. Consider the
        # following example scenarios:
        # Erasure Coding is used here with 3 data chunks and 1 parity chunks which
        # assumes 4 OSDs exist.
        # Configure this according to your CephCluster specification.
        erasureCoded:
          dataChunks: 3
          codingChunks: 1
      preservePoolsOnDelete: false
      gateway:
        port: 80
        instances: 2
        priorityClassName: system-cluster-critical
      resources:
        limits:
          cpu: "750m"
          memory: "1Gi"
        requests:
          cpu: "250m"
          memory: "500Mi"
```

⁷³² <https://www.rook.io/docs/rook/v1.10/Helm-Charts/operator-chart/#configuration>

⁷³³ <https://www.rook.io/docs/rook/v1.10/Helm-Charts/operator-chart/#release>

⁷³⁴ <https://www.rook.io/docs/rook/v1.10/Helm-Charts/ceph-cluster-chart/#ceph-object-stores>

⁷³⁵ <https://www.rook.io/docs/rook/v1.10/Helm-Charts/ceph-cluster-chart/#release>


```

healthCheck:
  bucket:
    interval: 60s
storageClass:
  enabled: true
  name: dkp-object-store
  reclaimPolicy: Delete

```

6.12.2.2.1.2 Managing resources directly

1. Set a variable to refer to the namespace the `AppDeployment`s are created in.

NOTE: This is the `kommander` namespace on the management cluster or `Workspace` namespace on all other clusters.

```

export CEPH_NAMESPACE=rook-ceph
export NAMESPACE=kommander

```

2. Create `CephObjectStore` in the same namespace as the `CephCluster` :

```

cat <<EOF | kubectl apply -f -
apiVersion: ceph.rook.io/v1
kind: CephObjectStore
metadata:
  name: dkp-object-store
  namespace: ${CEPH_NAMESPACE}
spec:
  metadataPool:
    # The failure domain: osd/host/(region or zone if available) - technically,
    any type in the crush map
    failureDomain: osd
    # Must use replicated pool ONLY. Erasure coding is not supported.
    replicated:
      size: 3
  dataPool:
    # The failure domain: osd/host/(region or zone if available) - technically,
    any type in the crush map
    failureDomain: osd
    # Data pool can use either replication OR erasure coding. Consider the
    following example scenarios:
    # Erasure Coding is used here with 3 data chunks and 1 parity chunks which
    assumes 4 OSDs exist.
    # Configure this according to your CephCluster specification.
    erasureCoded:
      dataChunks: 3
      codingChunks: 1
  preservePoolsOnDelete: false
gateway:

```

```

port: 80
instances: 2
priorityClassName: system-cluster-critical
resources:
  limits:
    cpu: "750m"
    memory: "1Gi"
  requests:
    cpu: "250m"
    memory: "500Mi"
healthCheck:
  bucket:
    interval: 60s
EOF

```

3. Wait for the `CephObjectStore` to be `Connected` :

```

$ kubectl get cephobjectstore -A
NAMESPACE   NAME                PHASE
rook-ceph   dkp-object-store    Progressing
...
...
rook-ceph   dkp-object-store    Connected

```

4. Create a `StorageClass` to consume the object storage:

```

cat <<EOF | kubectl apply -f -
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  name: dkp-object-store
parameters:
  objectStoreName: dkp-object-store
  objectStoreNamespace: ${CEPH_NAMESPACE}
provisioner: ${CEPH_NAMESPACE}.ceph.rook.io/bucket
reclaimPolicy: Delete
volumeBindingMode: Immediate
EOF

```

6.12.2.2.2 Creating ObjectBucketClaims

1. After connecting the Object Store, create the `ObjectBucketClaim` in the same namespace as `velero` and `grafana-loki` .

This results in the creation of `ObjectBucket` , that then creates `Secret` s that are consumed by `velero` and `grafana-loki` .

- a. For grafana-loki :

```
cat <<EOF | kubectl apply -f -
apiVersion: objectbucket.io/v1alpha1
kind: ObjectBucketClaim
metadata:
  name: dkp-loki
  namespace: ${NAMESPACE}
spec:
  additionalConfig:
    maxSize: 80G
  bucketName: dkp-loki
  storageClassName: dkp-object-store
EOF
```

- b. For velero :

```
cat <<EOF | kubectl apply -f -
apiVersion: objectbucket.io/v1alpha1
kind: ObjectBucketClaim
metadata:
  name: dkp-velero
  namespace: ${NAMESPACE}
spec:
  additionalConfig:
    maxSize: 10G
  bucketName: dkp-velero
  storageClassName: dkp-object-store
EOF
```

2. Wait for the ObjectBucket s to be Bound by executing the following command:

```
kubectl get objectbucketclaim -n${NAMESPACE} --ocustom-
columns='NAME:.metadata.name,PHASE:.status.phase'
```

which should display something similar to:

NAME	PHASE
dkp-loki	Bound
dkp-velero	Bound

6.12.2.3 Configure Loki to use S3 Compatible Storage

If you wish to use your own storage in DKP that is S3 compatible, create a secret that contains your AWS secret credentials.

```

apiVersion: v1
data:
  AWS_ACCESS_KEY_ID: base64EncodedValue
  AWS_SECRET_ACCESS_KEY: base64EncodedValue
kind: Secret
metadata:
  name: dkp-loki #If you want to configure a custom name here also use it in the step
  below
  namespace: kommander

```

6.12.2.4 Overriding `velero` and `grafana-loki` Configuration

After all the buckets are in the `Bound` state, DKP applications are now ready to be installed with the following configuration overrides populated in the installer config:

```

cat <<EOF | kubectl apply -f -
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
  grafana-loki:
    enabled: true
    values: |
      loki:
        structuredConfig:
          storage_config:
            aws:
              s3: "http://rook-ceph-rgw-dkp-object-store.${CEPH_NAMESPACE}.svc:80/
dkp-loki"
    ingester:
      extraEnvFrom:
        - secretRef:
            name: dkp-loki # Optional: This is the default value
    querier:
      extraEnvFrom:
        - secretRef:
            name: dkp-loki # Optional: This is the default value
    queryFrontend:
      extraEnvFrom:
        - secretRef:
            name: dkp-loki # Optional: This is the default value
    compactor:
      extraEnvFrom:
        - secretRef:
            name: dkp-loki # Optional: This is the default value
    ruler:
      extraEnvFrom:
        - secretRef:
            name: dkp-loki # Optional: This is the default value

```

```

distributor:
  extraEnvFrom:
    - secretRef:
        name: dkp-loki # Optional: This is the default value
velero:
  enabled: true
  values: |
    configuration:
      backupStorageLocation:
        - bucket: dkp-velero
          provider: "aws"
          config:
            region: dkp-object-store
            s3Url: http://rook-ceph-rgw-dkp-object-store.${CEPH_NAMESPACE}.svc:80/
      credentials:
        # This secret is owned by the ObjectBucketClaim. A ConfigMap and a Secret
        with same name as bucket are created.
        extraSecretRef: dkp-velero
EOF

```

This installer config can be merged with your installer config with any other relevant configuration before installing DKP.

6.12.2.5 Overriding `project-grafana-loki` Configuration

When installing project level grafana loki, its configuration needs to be overridden in a similar manner to workspace level grafana loki, so that the project logs can be persisted in Ceph storage.

The following overrides need to be applied to `project-grafana-loki` :

```

loki:
  structuredConfig:
    storage_config:
      aws:
        s3: "http://rook-ceph-rgw-dkp-object-store.${CEPH_NAMESPACE}.svc:80/dkp-loki"

```

These overrides can be applied from the UI directly while substituting the ``${CEPH_NAMESPACE}`` appropriately.

If you are using using CLI, follow these steps:

1. Set `NAMESPACE` to project namespace and `CEPH_NAMESPACE` to Ceph install namespace:

```

export CEPH_NAMESPACE=rook-ceph
export NAMESPACE=my-project

```

2. Create a `ConfigMap` to apply the configuration overrides:

```

cat <<EOF | kubectl apply -f -
apiVersion: v1
data:
  values.yaml: |
    loki:
      structuredConfig:
        storage_config:
          aws:
            s3: "http://rook-ceph-rgw-dkp-object-store.$
{CEPH_NAMESPACE}.svc:80/proj-loki-${NAMESPACE}"
kind: ConfigMap
metadata:
  name: project-grafana-loki-ceph
  namespace: ${NAMESPACE}
EOF

```

3. Create the `AppDeployment` with a reference to the above `ConfigMap` :

NOTE: The `clusterSelector` can be adjusted according to your needs.

```

cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
  name: project-grafana-loki
  namespace: ${NAMESPACE}
spec:
  appRef:
    kind: ClusterApp
    name: project-grafana-loki-0.48.6
  clusterSelector: {}
  configOverrides:
    name: project-grafana-loki-ceph
EOF

```

The project level Grafana Loki creates an `ObjectBucketClaim` and assumes that the Ceph operator is monitoring the project namespace, so there is no need to create `ObjectBucketClaim` manually.

6.13 DKP Troubleshooting

The following pages provide the steps to troubleshoot a Kubernetes cluster installation.

- [Generate a Support Bundle](#) (see page 1020)
- [Custom Collectors](#) (see page 1024)

6.13.1 Generate a Support Bundle

Follow these instructions to generate a support bundle with data collected for the last 48 hours of the life of the cluster.

6.13.1.1 Prerequisites

Before generating a support bundle, verify that you have:

- An AMD64-based Linux or macOS machine with a supported version of the operating system.
- A running Kubernetes cluster.
- Access to the [DKP CLI \(see page 1644\)](#).

6.13.1.2 Diagnostic Bundle

`dkp diagnose` was developed by D2iQ and builds on the open source `troubleshoot.sh` project.

■ The command `dkp diagnose` is based on version `0.13.16` of `troubleshoot.sh` with custom modifications. The D2iQ fork is open source and available from [this public GitHub repository](#)⁷³⁶.

`dkp diagnose` supports [multiple support bundle collectors](#)⁷³⁷ and can be configured as a `SupportBundle` Kubernetes resource in a YAML file.

The following list is the minimum set of resources that is required to debug a cluster, but can be further customized.

The bundle uses the following collectors:

- [clusterInfo](#)⁷³⁸ collects basic information about the cluster
- [clusterResources](#)⁷³⁹ collects a subset of available resources in the cluster
- [configMap](#)⁷⁴⁰ collects the values of Kubernetes ConfigMaps
- [secrets](#)⁷⁴¹ collects the names and metadata of secrets, but NOT the secret values or keys
- [execCopyFromHost](#) (see page 0) runs a container on each node on the cluster and copies the created data
- [allLogs](#) (see page 1024) is capable of collecting logs from all containers on the cluster

⁷³⁶ <https://github.com/mesosphere/troubleshoot>

⁷³⁷ <https://troubleshoot.sh/docs/collect/all/>

⁷³⁸ <https://troubleshoot.sh/docs/collect/cluster-info/>

⁷³⁹ <https://troubleshoot.sh/docs/collect/cluster-resources/>

⁷⁴⁰ <https://troubleshoot.sh/docs/collect/configmap/>

⁷⁴¹ <https://troubleshoot.sh/docs/collect/secret/>

- `cpuMemoryUtilisation` collects the average CPU and memory utilization of the last 30 days.

6.13.1.3 Generate a Support Bundle

F The command `dkp diagnose` uses the same Kubernetes configuration as `kubectl`. `dkp diagnose` can also be pointed at a specific configuration by using the `--kubeconfig` parameter.

To generate the support bundle, perform the following steps:

1. Run the `dkp diagnose` command by running the default collectors configuration.

```
dkp diagnose
```

The output looks similar to this:

```
Collecting support bundle ...
support-bundle-2021-08-13T14_44_23.tar.gz
```

2. To view the bundle contents, extract the bundle (replacing `support-bundle-2021-08-13T14_44_23.tar.gz` with the location from the previous step):

```
tar -xzvf support-bundle-2021-08-13T14_44_23.tar.gz
```

3. A new directory named `support-bundle-<date-created>` is created. This directory contains the files specified:

```
ls support-bundle-2021-08-13T14_44_23
```

The output looks similar to this:

```
cluster-info cluster-resources configmaps node-diagnostics pod-logs
secrets version.yaml
```


6.13.1.3.1 Collect Information from a Bootstrap Cluster

In the case where your bootstrap cluster has not yet pivoted towards your Konvoy cluster, you can collect log information from that bootstrap cluster as well, and there are a preconfigured set of relevant collectors. Specify an additional bootstrap cluster kubeconfig using the `--bootstrap-kubeconfig` parameter to activate bootstrap cluster diagnostics. You will receive an additional support bundle named `bootstrap-support-bundle-<date created>`.

Note that the bootstrap cluster diagnostics are independent of the configuration of the “main” or Konvoy cluster diagnostics. We run a static collector set that collects the following bootstrap cluster information:

- ClusterInfo
- ClusterResources
- AllLogs
- ConfigMaps
- Secrets

1. Run the `dkp diagnose` command with bootstrap bundle configuration.


```
dkp diagnose bundle.yaml
```

6.13.1.3.2 Customizations

To print the default collectors configuration, run the following command:

```
dkp diagnose default-config > bundle.yaml
```

Edit the file to make appropriate modifications.

-  By default, `dkp diagnose` does not require that you supply a configuration. You can print the default bundle by running `dkp diagnose default-config`.

6.13.1.4 SSH Fallback

In some cases the Kubernetes API is not available for the cluster. In those cases you can collect node level information using SSH access to the diagnosed nodes. Be aware that not all clusters have SSH access configured. If they do not then access using SSH fallback is not possible.

To get node level information from your cluster using SSH access, perform the following steps:

1. Enter the following command:

```
dkp diagnose ssh <path/to/ansible-inventory.yaml>
```

The `ansible-inventory.yaml` file specifies the nodes to access for data collection.

ⓘ This collector does not use the full Ansible `inventory.yaml` format only a limited subset to describe the infrastructure.

Only the following attributes of the `ansible-inventory.yaml` are supported. All other group definitions are ignored.

- Support for `all` shared variables.
- Support for `hosts` key in `all` groups.
- Supported behavioral inventory is limited to:
 - `ansible_host`
 - `ansible_port`
 - `ansible_user`
 - `ansible_ssh_private_key_file`

The following is an example `inventory.yaml` file:

```
all:
  vars:
    ansible_user: centos
  hosts:
    host-1:
      ansible_host: 192.168.10.1
    host-2:
      ansible_host: 192.168.10.22
      ansible_port: 2222
```

More information on these Ansible parameters can be found in the [Ansible user guide](https://docs.ansible.com/ansible/latest/user_guide/intro_inventory.html#connecting-to-hosts-behavioral-inventory-parameters)⁷⁴².

ⓘ All other group definitions in the `inventory.yaml` file are ignored.

Refer to the following example file:

⁷⁴² https://docs.ansible.com/ansible/latest/user_guide/intro_inventory.html#connecting-to-hosts-behavioral-inventory-parameters

```

all:
  vars:
    ansible_user: centos
  hosts:
    host-1:
      ansible_host: 192.168.10.1
    host-2:
      ansible_host: 192.168.10.22
      ansible_port: 2222

```

The fallback collector runs a bash script over SSH and copies the collected data. The format of the created bundle matches that of `dkp diagnose collector` generated bundles.

```

node-diagnostics/<HOSTNAME_PORT>/data/
  - dmesg
  - ....

```

Redactors are supported and are in the same format as the main `dkp diagnose` command. Per node collection timeouts are supported using the `--timeout` parameter.

See also: [dkp-cli](#) (see page 1644)

6.13.2 Custom Collectors

For creating diagnostic bundles, D2iQ is using a customized version of `troubleshoot.sh` integrated into `dkp-diagnose`.

6.13.2.1 Customizations

To meet the specific needs of diagnosing DKP 2 clusters we have developed custom collectors and modified the behavior of upstream collectors. Go to our repository for more information on the [details](#)⁷⁴³ of the changes.

6.13.2.1.1 ExecCopyFromHost Collector

This is a new collector created specifically for gathering host level information from cluster nodes. The collector allows you to run a provided container image in a privileged mode, as a root user, with additional Linux capabilities and with the host filesystem mounted in the container.

You can collect host level information other than copying host level files. (This is already possible with the `CopyFromHost` collector.) Like the `CopyFromHost` collector, this collector runs as a Kubernetes `DaemonSet` executed on all nodes in the system. The data produced by the container are copied from a

⁷⁴³ <https://github.com/mesosphere/troubleshoot/blob/v0.13-d2iq/README.d2iq.md>

pre-defined directory into the diagnostics bundle under each node name. The name of the parent directory, in the diagnostics bundle, is determined by the name of the collector specified in its configuration.

The data written into the diagnostics bundle follows this format:

```
<collector-name> / <node-name> / data / (file1|file2|...)
```

The following is a sample configuration file:

```
spec:
  collectors:
    - execCopyFromHost:
      name: node-diagnostics
      image: mesosphere/dkp-diagnostics-node-collector:latest
      timeout: 30s
      command:
        - "/bin/bash"
        - "-c"
        - "/diagnostics/container.sh --hostroot /host --hostpath ${PATH} --
outputroot /output"
      workingDir: "/diagnostics"
      includeControlPlane: true
      privileged: true
      capabilities:
        - AUDIT_CONTROL
        - AUDIT_READ
        - BLOCK_SUSPEND
        - BPF
        - CHECKPOINT_RESTORE
        - DAC_READ_SEARCH
        - IPC_LOCK
        - IPC_OWNER
        - LEASE
        - LINUX_IMMUTABLE
        - MAC_ADMIN
        - MAC_OVERRIDE
        - NET_ADMIN
        - NET_BROADCAST
        - PERFMON
        - SYS_ADMIN
        - SYS_BOOT
        - SYS_MODULE
        - SYS_NICE
        - SYS_PACCT
        - SYS_PTRACE
        - SYS_RAWIO
        - SYS_RESOURCE
        - SYS_TIME
        - SYS_TTY_CONFIG
        - SYSLOG
```

```
- WAKE_ALARM
extractArchive: true
```

The following is an example of the data produced by running this collector:

```
├─ node-diagnostics
│  └─ troubleshoot-control-plane
│     └─ data
│        ├── certs_expiration_kubeadm
│        ├── containerd_config.toml
│        └─ ...
│  └─ troubleshoot-worker
│     └─ data
│        ├── containerd_config.toml
│        ├── containers_crictl
│        └─ ...
│     └─ whoami_validate
```

In the event that an error occurs while collecting node diagnostics, the `node-diagnostics/<node>/pod-collector.json` file contains the serialized JSON representations of the running pod. This helps debug the reasons for the collection failure. The `node-diagnostics/<node>/pod-collector.log` file contains stdout from the collector container that runs the diagnostics script. In addition, the command may also produce certain `-error.txt` files. `file-copy-error.txt` and `pod-collector-files-copy-error.txt` are two file examples. These files contain error messages generated while trying to fetch log files from the collector.

When using this collector for node level information you must run additional docker containers and must have the following docker images:

- `mesosphere/pause-alpine:3.2`
- `mesosphere/dkp-diagnostics-node-collector:$(dkp-diagnose version)`

For more information on the configuration options see the `ExecCopyFromHost` in the `pkg/apis/troubleshoot/v1beta2/exec_copy_from_host.go` file.

6.13.2.1.2 AllLogs Collector

This collector gathers pod logs from specified namespaces or from all namespaces if none are specified. You can collect logs of all the pods from all the namespaces. The pod logs are collected under the `allPodLogs` directory.

The data written into the diagnostics bundle follows this format:

```
<collector-name> / <namespace-name> / <pod-name> - (container1|container2|...)
```

The following is a sample configuration file to collect logs from all the pods from all the namespaces:

```
spec:
  collectors:
    - allLogs:
      namespaces:
        - "*"

```

The following is a sample configuration file to collect logs from all the pods from specific namespaces:

```
spec:
  collectors:
    - allLogs:
      namespaces:
        - default
        - dev
        - prod

```

The following is an example of the data produced by running this collector:

```
├─ node-diagnostics
│  └─ troubleshoot-control-plane
│     └─ data
│        ├── certs_expiration_kubeadm
│        ├── containerd_config.toml
│        └─ ...
│  └─ troubleshoot-worker
│     └─ data
│        ├── containerd_config.toml
│        ├── containers_crictl
│        └─ ...
│     └─ whoami_validate

```

In the event that an error occurs while collecting node diagnostics, the `node-diagnostics/<node>/pod-collector.json` file contains the serialized JSON representations of the running pod. This helps debug the reasons for the collection failure. The `node-diagnostics/<node>/pod-collector.log` file contains stdout from the collector container that runs the diagnostics script.

When using this collector for node level information you must run additional docker containers and must have the following docker images:

- `mesosphere/pause-alpine:3.2`
- `mesosphere/dkp-diagnostics-node-collector:$(dkp-diagnose version)`

For more information on the configuration options see the `ExecCopyFromHost` in the `pkg/apis/troubleshoot/v1beta2/exec_copy_from_host.go` file.

6.13.2.1.3 Collect from all Namespaces for ConfigMap and Secret Collector

Support for collecting from all namespaces for ConfigMap and Secret collector.

In the original collectors `namespace` there is a required parameter. This adds support for collecting from all namespaces by not setting the `namespace` (or setting it to `""`). Note: To collect all config maps / secrets an empty selector must be used (`selector: [""]`).

6.13.2.1.4 Support for Optional support-bundle Name Prefix

When generating a support bundle, you need naming defaults to provide deterministic bundle identifiers. This feature is especially useful for our convenience extension of providing diagnostics for both, a bootstrap, Konvoy, or other K8s cluster. Using an empty prefix keeps the original naming convention.

6.13.2.1.5 ClusterResources Collector

Another customization is added to collect custom resource definitions and all custom resources in the cluster.

6.14 Telemetry

With the telemetry function, you can provide environment data to D2iQ.

You can benefit from providing your data, as it will help us **improve your installation, upgrade and support experiences**. Given that these experiences are part of your cluster's lifecycle, it is important for us to constantly improve them, make them user-friendlier and seamless. This is only possible if we can take a closer look at your environment data, learn more about how you use our product, and recognize the need for new features and prioritize them.

You can choose to send data bundles to D2iQ for:

- [Troubleshooting and support assistance \(see page 1023\)](#)
- [Product improvement \(see page 1030\)](#) (improvement of installation and upgrade experiences)

6.14.1 Information collected by D2iQ

With the help of [Collectors \(see page 1029\)](#), D2iQ gathers **environment information** that includes, but is not limited to:

- Infrastructure (cloud provider, pre-provisioned environment, etc.)
- Kubernetes version
- DKP version and license type
- Operating system
- Number of nodes
 - Number of cores per node / vCPU

- CPU & memory utilization
- FIPS or GPU enablement

6.14.2 Data Storage Location

The telemetry command generates a bundle and automatically sends it to a storage service managed by D2iQ. The CLI determines the storage location for the bundle based on the geographical location of the user running the command. Bundles from users running the command in Europe are stored in European data centers. Bundles generated outside Europe are stored in US data centers.

6.14.3 Size Limitations

The bundle upload size limit is 10 GB.

If the upload fails because the bundle is too large, [generate a support bundle locally](#) (see page 1023), and provide it to D2iQ through a [support ticket](#)⁷⁴⁴.

6.14.4 Used Collectors

Extended Bundle Collectors

6.14.4.1 Extended Bundle Collectors

The bundle uses the following collectors:

- [clusterInfo](#)⁷⁴⁵ collects basic information about the cluster
- [clusterResources](#)⁷⁴⁶ collects a subset of available resources in the cluster
- [configMap](#)⁷⁴⁷ collects the values of Kubernetes ConfigMaps
- [secrets](#)⁷⁴⁸ collects the names and metadata of secrets, but NOT the secret values or keys
- [execCopyFromHost](#) (see page 0) runs a container on each node on the cluster and copies the created data
- [allLogs](#) (see page 1027) is capable of collecting logs from all containers on the cluster
- `cpuMemoryUtilisation` collects the average CPU and memory utilization of the last 30 days.

See <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273387782/2.6%2BSending%2BTelemetry%2BData%2Bfor%2BAnalysis#full> (see page 1031) for the exact command.

Telemetry Bundle Collectors

⁷⁴⁴ <https://support.d2iq.com/hc/en-us/articles/4408231804564-How-to-Submit-a-Case>

⁷⁴⁵ <https://troubleshoot.sh/docs/collect/cluster-info/>

⁷⁴⁶ <https://troubleshoot.sh/docs/collect/cluster-resources/>

⁷⁴⁷ <https://troubleshoot.sh/docs/collect/configmap/>

⁷⁴⁸ <https://troubleshoot.sh/docs/collect/secret/>

6.14.4.2 Telemetry Bundle Collectors

This bundle contains a subset of the data contained in the diagnostics bundle:

- [clusterInfo](#)⁷⁴⁹ collects basic information about the cluster
- [clusterResources](#)⁷⁵⁰ collects a subset of available resources in the cluster (only in the `kommander`, `kubsystem`, and `default` namespaces)
- `cpuMemoryUtilisation` collects the average CPU and memory utilization of the last 30 days.

See <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273387782/2.6%2BSending%2BTelemetry%2BData%2Bfor%2BAnalysis#partial> (see page 1031) for the exact command.

6.14.5 Sending Telemetry Data for Analysis

Generate and send Telemetry data to D2iQ for analysis.

For more information around the gathered data, collection purposes, and data storage location, see [Telemetry](#) (see page 1031).



This feature is currently not available in air-gapped environments.


6.14.5.1 Prerequisites

- Your environment must have a valid license. See [Add a DKP license](#) (see page 139) for more information.
- Your host machine (where the commands are executed) must have access to the Internet.
- Ensure your firewall rules allow HTTPS calls and connections to the <https://api.secure-upload.d2iq.cloud/secure-upload> endpoint.

⁷⁴⁹ <https://troubleshoot.sh/docs/collect/cluster-info/>

⁷⁵⁰ <https://troubleshoot.sh/docs/collect/cluster-resources/>

6.14.5.2 Generate and Send a Bundle

-  A telemetry bundle is tied to a specific cluster. If there is more than one cluster in your environment, run the following instructions for each cluster by providing a context for your commands. You can either use the `--kubeconfig` flag as shown below, or set an environment variable as shown in [Provide Context for Commands with a kubeconfig File](#) (see page 87).


For more information around the gathered data, collection purposes, and data storage location, see [Telemetry](#) (see page 1031).

6.14.5.2.1 Telemetry Bundle

Execute the following command to upload a telemetry bundle for analysis:


```
dkp upload telemetry --bundle-profile telemetry --kubeconfig=${CLUSTER_NAME}.conf
```

This command automatically generates a bundle and sends it to a storage service managed by D2iQ.

-  This is the default profile, if no flag is set when running `dkp upload telemetry`.

If the upload fails due to size limitations, [generate a support bundle locally](#) (see page 1023), and provide it to D2iQ through a [support ticket](#)⁷⁵¹.

6.14.5.2.2 Alternative: Extended Bundle

-  See [Used collectors](#) (see page 1032) if you want to know more about how these two bundles differ.

Execute the following command to upload an extended bundle for analysis:

```
dkp upload telemetry --bundle-profile diagnostics --kubeconfig=${CLUSTER_NAME}.conf
```

⁷⁵¹ <https://support.d2iq.com/hc/en-us/articles/4408231804564-How-to-Submit-a-Case>

This command automatically generates a bundle and sends it to a storage service managed by D2iQ.

6.14.5.3 Recommended Telemetry Upload Frequency

The recommended frequency for data upload is every three months.

If you want to upload telemetry data on a schedule, see [Scheduling Telemetry Upload](#) (see page 1032) for a configuration aid.

6.14.6 Scheduling Telemetry Upload

If you want to upload telemetry data on a schedule, you can use the following example as an aid to set up a Cron Job that runs the telemetry command for you on a regular basis. The recommended frequency for data upload is every three months.

For more information around the gathered data, collection purposes, and data storage location, see [Telemetry](#) (see page 1031).



A bundle is tied to a specific cluster. If you want to set up a schedule for more than one cluster in your environment, set up a CronJob for each cluster individually.



This feature is currently not available in air-gapped environments.

6.14.6.1 Prerequisites

- Your environment must have a valid license. See [Add a DKP license](#) (see page 139) for more information.
- Your host machine (where the commands are executed) must have access to the Internet.
- Ensure your firewall rules allow HTTPS calls and connections to the <https://api.secure-upload.d2iq.cloud/secure-upload> endpoint.
- Access to the cluster's `kubeconfig` file.
- Your cluster nodes must have access to the Internet.

6.14.6.2 Create and Enable the CronJob

The following example shows a [Kubernetes CronJob](#)⁷⁵² that manages running a command on an established schedule inside the cluster.

⁷⁵² <https://kubernetes.io/docs/concepts/workloads/controllers/cron-jobs/>

1. Create a file with the schedule using the [Cron Syntax](https://en.wikipedia.org/wiki/Cron)⁷⁵³. In this example, the job is called `telemetry` and runs four times a year, starting on the first day of the third month from the configuration date.

```

apiVersion: batch/v1
kind: CronJob
metadata:
  name: telemetry
spec:
  schedule: "0 0 1 */3 *"
  jobTemplate:
    spec:
      template:
        spec:
          serviceAccountName: telemetry-collector
          restartPolicy: OnFailure
          containers:
            - name: telemetry
              image: alpine:3.6
              imagePullPolicy: IfNotPresent
              env:
                - name: DKP_VERSION
                  value: v2.6.1
              command:
                - /bin/sh
                - -c
                - |
                  apk update && apk add curl
                  curl -L https://downloads.d2iq.com/dkp/${DKP_VERSION}/dkp_${DKP_VERSION}_linux_amd64.tar.gz | tar -xz
                  ./dkp upload telemetry --bundle-profile diagnostics
---
apiVersion: v1
kind: ServiceAccount
metadata:
  name: telemetry-collector
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRole
metadata:
  name: telemetry-collector
rules:
  - apiGroups: ["*"]
    resources: ["*"]
    verbs: ["get", "watch", "list"]
---
apiVersion: rbac.authorization.k8s.io/v1
kind: ClusterRoleBinding

```

753 <https://en.wikipedia.org/wiki/Cron>

```

metadata:
  name: telemetry-collector
  namespace: default
roleRef:
  apiGroup: rbac.authorization.k8s.io
  kind: ClusterRole
  name: telemetry-collector
subjects:
  - kind: ServiceAccount
    namespace: default
    name: telemetry-collector

```

i Specify a different `--bundle-profile` flag if you want to limit the amount of shared data. See [Sending Telemetry Data for Analysis \(see page 1033\)](#) for more information on the *telemetry* profile.

2. Apply the created file to the cluster where you would like to automate the command. Replace the placeholder `<target_cluster_kubeconfig>` with the path to the target cluster's `kubeconfig` file:

```
kubectl apply -f CronJob.yaml --kubeconfig=<target_cluster_kubeconfig>
```

Alternative: Set up GitHub Actions to execute a workflow on a schedule for your cluster

6.14.6.3 GitHub Actions

In this section, we provide an alternative to set up a workflow that manages running the command on a schedule. This is an alternative for organizations that already use GitHub Actions and wish to use it to manage this workflow as well.

6.14.6.3.1 Prerequisites

- Access to the cluster's `kubeconfig` file.
- A GitHub Actions account.

6.14.6.3.2 Create and Enable GitHub Actions

Here is an example of how you could configure a GitHub Actions workflow for the telemetry command. In this example, the command runs four times a year, starting on the first day of the third month from the configuration date.

Replace the placeholder `<target_cluster_kubeconfig>` with the path to the target cluster's `kubeconfig` file:

```
name: Run DKP telemetry upload command on the cluster
on:
  schedule:
    - cron: "0 0 1 */3 *"

jobs:
  telemetry:
    runs-on:
      - ubuntu-latest

    steps:
      ...
      - name: Upload telemetry data
        env:
          DKP_VERSION: v2.6.1
        run: |
          # DKP_TAR_FILE_URL is a download link to DKP binary
          curl -L https://downloads.d2iq.com/dkp/${DKP_VERSION}/dkp_${DKP_VERSION}
          _linux_amd64.tar.gz | tar -xz
          ./dkp upload telemetry --bundle-profile diagnostics --kubeconfig
          <target_cluster_kubeconfig>
```

- Specify a different `--bundle-profile` flag if you want to limit the amount of shared data. See [Sending Telemetry Data for Analysis](#) (see page 1033) for more information on the restricted telemetry profile.

7 Custom Installation and Additional Infrastructure Tools

The Konvoy component of DKP can be customized for different environments and infrastructures depending on your network technology choices. If you have already installed using [Day 1 - Basic Installs by Infrastructure](#) (see page 147) instructions, you may find helpful tools in this section. However, if you have not already installed DKP with the Basic Install instructions, find your infrastructure and begin the process of custom DKP installation in this area. See the following sections for more information on custom settings:

7.1 Universal Configurations for all Infrastructure Providers

- [Container Runtime Engine \(CRE\)](#) (see page 1039)
- [Configuring an HTTP/HTTPS Proxy](#) (see page 1039)
- [Working with Load Balancers](#) (see page 1045)
- [Registry and Registry Mirrors](#) (see page 1046)
- [Subnets and Pods](#) (see page 1046)

7.2 AWS Infrastructure

- [AWS Prerequisites](#) (see page 1051)
- [AWS Konvoy Image Builder](#) (see page 1053)
- [Minimal Permissions and Role to Create Clusters](#) (see page 1057)
- [Cluster IAM Policies, Roles, and Artifacts](#) (see page 1063)
- [AWS Cluster Creation Choices](#) (see page 1071)
- [Delete an AWS Cluster](#) (see page 1102)
- [Multiple AWS Accounts](#) (see page 1107)
- [GPUs in an AWS environment](#) (see page 1110)
- [Manage AWS Node Pools](#) (see page 1113)
- [AWS Certificate Renewal](#) (see page 1121)
- [Configure Infrastructure in UI](#) (see page 1123)
- [Replace an AWS Node](#) (see page 1124)

7.3 EKS Infrastructure

- [EKS Introduction](#) (see page 1128)
- [Minimal User Permission for EKS Cluster Creation](#) (see page 1129)
- [EKS Cluster IAM Permissions and Roles](#) (see page 1132)
- [Create an EKS Cluster from the CLI](#) (see page 1137)
- [Create an EKS Cluster from the DKP UI](#) (see page 1144)
- [Grant Cluster Access](#) (see page 1147)
- [Explore EKS Cluster](#) (see page 1148)

- [Manage EKS Nodepools](#) (see page 1151)
- [Delete EKS Cluster from CLI](#) (see page 1153)
- [Delete EKS Cluster from the DKP UI](#) (see page 1154)

7.4 Azure Infrastructure

- [Azure Prerequisites](#) (see page 1156)
- [Azure using Konvoy Image Builder](#) (see page 1159)
- [Azure Bootstrap](#) (see page 1161)
- [Create a New Custom Azure Cluster](#) (see page 1163)
- [Explore new Azure Cluster](#) (see page 1173)
- [Azure Make new Cluster Self-Managed](#) (see page 1177)
- [Azure Certificate Renewal](#) (see page 1179)
- [Azure Replace a Node](#) (see page 1182)
- [Azure Delete Cluster](#) (see page 1185)

7.5 AKS Infrastructure

- [Create a New AKS Cluster](#) (see page 1189)
- [Explore New AKS Cluster](#) (see page 1193)
- [Delete AKS Cluster](#) (see page 1197)
- [Create a new AKS Cluster from the DKP UI](#) (see page 1198)

7.6 Pre-provisioned Infrastructure

- [Advanced Workflow](#) (see page 1201)
- [Pre-provisioned Prerequisite Configuration](#) (see page 1201)
- [Pre-provisioned Air-gapped Define Environment](#) (see page 1204)
- [Pre-provisioned Set Infrastructure](#) (see page 1209)
- [Pre-provisioned Define Control Plane Endpoint](#) (see page 1211)
- [Pre-provisioned Create Secrets and Overrides](#) (see page 1213)
- [Pre-provisioned Bootstrap Cluster](#) (see page 1219)
- [Pre-provisioned Create a New Cluster](#) (see page 1220)
- [Pre-provisioned Azure only Configurations](#) (see page 1227)
- [Pre-provisioned Modify the Calico Installation](#) (see page 1231)
- [Pre-provisioned Built-in Virtual IP](#) (see page 1236)
- [Provision on the Flatcar Linux OS](#) (see page 1236)
- [Pre-provisioned Use HTTP Proxy](#) (see page 1237)
- [Pre-provisioned Use Alternate Pod or Service Subnets](#) (see page 1238)
- [Pre-provisioned: Make Cluster Self-managed](#) (see page 1240)
- [Pre-provisioned Configure MetalLB](#) (see page 1243)
- [Pre-provisioned Create and Delete Node Pools](#) (see page 1245)

- [Pre-provisioned Add Nodes to Existing Node Pool](#) (see page 1247)
- [GPU Nodepools in a Pre-provisioned Environment](#) (see page 1249)
- [Pre-provisioned Delete Cluster](#) (see page 1251)

7.7 vSphere Infrastructure

- [vSphere Prerequisites](#) (see page 1256)
- [Create a Base OS image in vSphere](#) (see page 1260)
- [Create a VM Template](#) (see page 1262)
- [vSphere Bootstrap](#) (see page 1266)
- [Create new vSphere Cluster](#) (see page 1268)
- [Explore a vSphere Cluster](#) (see page 1275)
- [Make vSphere Cluster Self-Managed](#) (see page 1279)
- [Configure MetalLB for a vSphere infrastructure](#) (see page 1282)
- [Install vSphere Air-Gapped](#) (see page 1285)
- [vSphere Certificate Renewal](#) (see page 1305)
- [Delete vSphere Cluster](#) (see page 1307)
- [Manage vSphere Node Pools](#) (see page 1311)

7.8 VMware Cloud Director Infrastructure

- [Cloud Director for Service Providers](#) (see page 1323)

7.9 GCP Infrastructure

- [GCP Prerequisites](#) (see page 1365)
- [GCP Konvoy Image Builder](#) (see page 1368)
- [Bootstrap GCP](#) (see page 1371)
- [Create a New GCP Cluster](#) (see page 1373)
- [Explore the GCP Cluster](#) (see page 1376)
- [Make the New GCP Cluster Self-Managed](#) (see page 1380)
- [Manage GCP Node Pools](#) (see page 1382)
- [Delete a GCP Cluster](#) (see page 1390)

7.10 Universal Configurations for all Infrastructure Providers

Several areas of DKP configuration are common amongst all amongst all infrastructure providers. Some of the universal configurations are described in this section, and some pages include links to expanded information for these topics.

- [Container Runtime Engine \(CRE\)](#) (see page 1039)

- [Configuring an HTTP/HTTPS Proxy](#) (see page 1039)
- [Working with Load Balancers](#) (see page 1045)
- [Registry and Registry Mirrors](#) (see page 1046)
- [Subnets and Pods](#) (see page 1046)

7.10.1 Additional Configurations

You can find more information regarding global configurations or customization of specific components in the [Additional Configurations](#) (see page 1461) section of the documentation.

- [Konvoy Image Builder](#) (see page 1461)
- [FIPS 140-2 Compliance](#) (see page 1526)
- [Local Registry Tools](#) (see page 1533)
- [Air-gapped Seed the Registry](#) (see page 1537)
- [Configure the Control Plane](#) (see page 1538)
- [Update Cluster Nodepools](#) (see page 1545)

7.10.2 Container Runtime Engine (CRE)

A Container Runtime Engine (CRE) installed is required to install DKP. You can choose one of these supported CREs:

- [Docker](#)⁷⁵⁴ container engine version 18.09.2 or higher installed for Linux or MacOS
 - On macOS, Docker runs in a virtual machine which needs to be configured with at least 8 GB of memory.
- [Podman](#)⁷⁵⁵ version 4.0 or higher for Linux. You can find the host requirements here: [Host Requirements](#)⁷⁵⁶.

7.10.3 Configuring an HTTP/HTTPS Proxy

When creating a DKP cluster in environments that use an HTTP/HTTPS proxy, you must provide proxy details. The proxy values are strings that list a set of proxy servers, URLs, or wildcard addresses that is specific to your environment.

When creating a DKP cluster in a proxied environment, you need to specify proxy settings for the:

- Bootstrap cluster
- CAPI components
- DKP Kommander component

When you create a DKP cluster using the `--self-managed` flag, the bootstrap cluster and CAPI components are created for you automatically, and use the HTTP and HTTPS proxy settings you specify in the `dkp create cluster <provider>...` command.

⁷⁵⁴ <https://docs.docker.com/get-docker/>

⁷⁵⁵ <https://podman.io/getting-started/installation>

⁷⁵⁶ <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>

You can also create the bootstrap cluster and CAPI components manually, using the appropriate commands, `dkp create bootstrap` and `dkp create capi-components` respectively, combined with the command line flags to include your HTTP/S proxy information.

You can also specify HTTP/S proxy information in an override file when using [Konvoy Image Builder \(KIB\)](#) (see page 1513).

Without these values provided as part of the relevant `dkp create` command, DKP cannot create the requisite parts of your new cluster correctly. This is true of both [management and managed clusters](#) (see page 81) alike.



For DKP installation, you should create the bootstrap cluster from within the same network in which the new cluster will run. Using a bootstrap cluster on a laptop with different proxy settings, for example, or residing in a different network, could cause problems.

You can define HTTP/HTTPS proxy information using the steps in these pages:

- [Configure the Bootstrap Cluster HTTP Proxy Settings](#) (see page 1040)
- [Create CAPI Components with HTTP/HTTPS Proxy Settings](#) (see page 1041)
- [Create a Cluster with HTTP/HTTPS Proxy](#) (see page 1042)
- [Configure an HTTP/HTTPS Proxy for the DKP Kommander Component](#) (see page 1044)

7.10.3.1 Configure the Bootstrap Cluster HTTP Proxy Settings

When creating a bootstrap cluster, you must locate the device used to create the bootstrap in the same proxied environment in which the workload cluster will run. D2iQ does not recommend creating a bootstrap cluster from outside a proxied environment.

When you first install DKP, the API server doesn't exist yet in the bootstrap environment, because the API server is created *during* cluster creation. To create a bootstrap server in a proxied environment, you need to include the following flags:

- `--http-proxy <<http proxy list>>`
- `--https-proxy <<https proxy list>>`
- `--no-proxy <<no proxy list>>`

The following is an example `dkp create bootstrap` command's syntax with the HTTP proxy settings included:

```
dkp create bootstrap --http-proxy <<http proxy list>> --https-proxy <<https proxy list>> --no-proxy <<no proxy list>>
```

7.10.3.1.1 Create a Bootstrap Cluster with HTTP Proxy Settings

Note that the delimiter between each proxy value within a flag is a comma (,) with no space character following it. The flags can include a mix of IP addresses and domain names.

1. If an HTTP proxy is required, locate the values to use for the `http_proxy`, `https_proxy`, and `no_proxy` flags. They will be built into the bootstrap cluster during cluster creation.
2. Create a bootstrap cluster using this command syntax, in addition to any other flags you may need:

```
dkp create bootstrap --kubeconfig $HOME/.kube/config \
  --http-proxy <string> \
  --https-proxy <string> \
  --no-proxy <string>
```

This code sample shows the command with example values for the proxy settings:

```
dkp create bootstrap \
  --http-proxy 10.0.0.15:3128 \
  --https-proxy 10.0.0.15:3128 \
  --no-proxy 127.0.0.1,192.168.0.0/16,10.0.0.0/16,10.96.0.0/12,169.254.169.254,169.254.0.0/24,localhost,kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluster.local,.svc.cluster.local.,kubecost-prometheus-server.kommander,logging-operator-logging-fluentd.kommander.svc.cluster.local,elb.amazonaws.com
```

7.10.3.2 Create CAPI Components with HTTP/HTTPS Proxy Settings

Creating CAPI components for a DKP cluster from the command line requires HTTP/HTTPS proxy information, if your environment is proxied.

If you created a cluster without using the `--self-managed` flag, the cluster will not have any of the CAPI controllers or the cert-manager component. This means that the cluster will be managed from the context of the cluster from which it was created such as the bootstrap cluster. However, you can transform the cluster to a self-managed cluster by performing the commands:

```
dkp create capi-components --kubeconfig=<newcluster>
```

and...

```
dkp move --to-kubeconfig=<newcluster>
```

This combination of actions is sometimes called a **pivot**.

When creating the CAPI components for a proxied environment using the DKP command line interface, you must include the following flags :

- `--http-proxy <<http proxy list>>`

- `--https-proxy <<https proxy list>>`
- `--no-proxy <<no proxy list>>`

The following is an example `dkp create capi-components` command's syntax with the HTTP proxy settings included:

```
dkp create capi-components --http-proxy <<http proxy list>> --https-proxy <<https proxy list>> --no-proxy <<no proxy list>>
```

7.10.3.2.1 Create CAPI Components with HTTP Proxy Settings

Note that the delimiter between each proxy value within a flag is a comma (,) with no space character following it. The flags can include a mix of IP addresses and domain names.

1. If an HTTP proxy is required, locate the values to use for the `http_proxy`, `https_proxy`, and `no_proxy` flags. They will be built into the CAPI components during their creation.
2. Create CAPI components using this command syntax, in addition to any other flags you may need:

```
dkp create capi-components --kubeconfig $HOME/.kube/config \
  --http-proxy <string> \
  --https-proxy <string> \
  --no-proxy <string>
```

This code sample shows the command with example values for the proxy settings:

```
dkp create capi-components \
  --http-proxy 10.0.0.15:3128 \
  --https-proxy 10.0.0.15:3128 \
  --no-proxy 127.0.0.1,192.168.0.0/16,10.0.0.0/16,10.96.0.0/12,169.254.169.254,169.254.0.0/24,localhost,kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluster.local,.svc.cluster.local.,kubecost-prometheus-server.kommander,logging-operator-logging-fluentd.kommander.svc.cluster.local,elb.amazonaws.com
```

7.10.3.3 Create a Cluster with HTTP/HTTPS Proxy

During cluster creation, you may need to configure the control plane and worker nodes to use an HTTP proxy. This can occur during installation of the Konvoy component of DKP, or when creating a [managed](#) (see page 81) cluster.

If you require HTTP proxy configurations, you can apply them during the `create` operation by adding the appropriate flags to the `create cluster` command example below:

Proxy configuration	Flag
HTTP proxy for control plane machines	<code>--control-plane-http-proxy string</code>
HTTPS proxy for control plane machines	<code>--control-plane-https-proxy string</code>
No Proxy list for control plane machines	<code>--control-plane-no-proxy strings</code>
HTTP proxy for worker machines	<code>--worker-http-proxy string</code>
HTTPS proxy for worker machines	<code>--worker-https-proxy string</code>
No Proxy list for worker machines	<code>--worker-no-proxy strings</code>

- You must apply the same configuration to any custom machine images built with the Konvoy Image Builder (KIB) by using an HTTP [override](#) (see [page 1524](#)) file. For more information, refer to [Use Override Files with Konvoy Image Builder](#) (see [page 1520](#)) section of the documentation.

7.10.3.3.1 Configure the Control plane and Worker Nodes to Use HTTP/S proxy

This method for configuring the HTTP proxy values uses environment variables. (You are not required to use this method.)

Review this sample code for configuring environment variables for the control plane and worker nodes, taking into account the list of considerations that follows the sample.

```
export CONTROL_PLANE_HTTP_PROXY=http://example.org:8080
export CONTROL_PLANE_HTTPS_PROXY=http://example.org:8080
export
CONTROL_PLANE_NO_PROXY="example.org,example.com,example.net,localhost,127.0.0.1,10.96
.0.0/12,192.168.0.0/16,kubernetes,kubernetes.default,kubernetes.default.svc,kubernet
s.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.clu
ster.local,169.254.169.254,.elb.amazonaws.com"
```

```
export WORKER_HTTP_PROXY=http://example.org:8080
export WORKER_HTTPS_PROXY=http://example.org:8080
export
WORKER_NO_PROXY="example.org,example.com,example.net,localhost,127.0.0.1,10.96.0.0/12
,192.168.0.0/16,kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.defau
lt.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluster.lo
cal,169.254.169.254,.elb.amazonaws.com"
```

HTTP proxy configuration considerations to ensure the core components work correctly

- Replace `example.org,example.com,example.net` with your internal addresses
- `localhost` and `127.0.0.1` addresses should not use the proxy
- `10.96.0.0/12` is the default Kubernetes service subnet
- `192.168.0.0/16` is the default Kubernetes pod subnet
- `kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local` is the internal Kubernetes kube-apiserver service
- The entries `.svc,.svc.cluster,.svc.cluster.local` are the internal Kubernetes services
- Auto-IP addresses `169.254.169.254` for any cloud provider

7.10.3.3.1 Create a Cluster Using the Configured HTTP Proxy Variables

The following is an example of a `dkp create cluster...` command that uses the values set in the environment variables from the code sample above. Use the appropriate infrastructure provider name in line 1 from the choices listed:

```
dkp create cluster [aws, azure, gcp, preprovisioned, vsphere] \
--cluster-name ${CLUSTER_NAME} \
--control-plane-http-proxy="${CONTROL_PLANE_HTTP_PROXY}" \
--control-plane-https-proxy="${CONTROL_PLANE_HTTPS_PROXY}" \
--control-plane-no-proxy="${CONTROL_PLANE_NO_PROXY}" \
--worker-http-proxy="${WORKER_HTTP_PROXY}" \
--worker-https-proxy="${WORKER_HTTPS_PROXY}" \
--worker-no-proxy="${WORKER_NO_PROXY}"
```

7.10.3.4 Configure an HTTP/HTTPS Proxy for the DKP Kommander Component

7.10.3.4.1 HTTP Proxy for the Kommander Component of DKP

After the cluster is running in the Konvoy component from above, you need to configure the `NO_PROXY` variable for each provider.

For example, in addition to the values above for AWS, you need the following settings:

- The default VPC CIDR range of `10.0.0.0/16`
- `kube-apiserver` internal/external ELB address



- The `NO_PROXY` variable contains the Kubernetes Services CIDR. This example uses the default CIDR, `10.96.0.0/12`. If your cluster's CIDR is different, update the value in the `NO_PROXY` field.

Set the `httpProxy` and `httpsProxy` environment variables to the address of the HTTP and HTTPS proxy servers, respectively. (Frequently, environments use the same values for both.) Set the `noProxy` environment variable to the addresses that should be accessed directly, and not through the proxy.

For the Kommander component of DKP, refer to more HTTP Proxy information in [HTTP Proxy](#) (see page 1452).

7.10.4 Working with Load Balancers

In a Kubernetes cluster, depending on the flow of traffic direction, there are two kinds of load balancing:

- Internal load balancing for the traffic within a Kubernetes cluster
- External load balancing for the traffic coming from outside the cluster

7.10.4.1 External Load Balancers

DKP includes a load balancing solution for the [supported cloud infrastructure providers](#) (see page 69) and for pre-provisioned environments. For more information, see [Load Balancing for external traffic](#) (see page 967) in DKP.

If you want to use a **non-DKP load balancer** (for example, as an alternative to MetalLB in pre-provisioned environments), DKP supports setting up an **external load balancer**.

When enabled, the external load balancer routes incoming traffic requests to a single point of entry in your cluster. Users and services can then access the **DKP UI** through an established IP or DNS address.

7.10.4.2 Select your Connection Mechanism

A virtual IP is the address that the client uses to connect to the service. A load balancer is the device that distributes the client connections to the backend servers. Before you create a new DKP cluster, choose an external load balancer(LB) or virtual IP.

7.10.4.2.1 External load balancer

It is recommended that an external load balancer be the control plane endpoint. To distribute request load among the control plane machines, configure the load balancer to send requests to all the control plane machines. Configure the load balancer to send requests only to control plane machines that are responding to API requests.

7.10.4.2 Built-in virtual IP (option for Pre-provisioned or vSphere)

If an external load balancer is not available, use the [built-in virtual IP](#) (see page 1041). The virtual IP is *not* a load balancer; it does not distribute request load among the control plane machines. However, if the machine receiving requests does not respond to them, the virtual IP automatically moves to another machine.

7.10.4.3 External Load Balancer for the DKP Kommander Component

If you want to use a **non-DKP load balancer** (for example, as an alternative to MetalLB in pre-provisioned environments), DKP supports setting up an **external load balancer**.

When enabled, the external load balancer routes incoming traffic requests to a single point of entry in your cluster. Users and services can then access the **DKP UI** through an established IP or DNS address. Refer to [External Load Balancer](#) (see page 1450) page of documentation for further details.

For MetalLB, refer to the page: [Pre-provisioned Configure MetalLB](#) (see page 1243)

7.10.5 Registry and Registry Mirrors

If you need to set up a private registry with a registry mirror, see [this page](#) (see page 1536) for details on using that flag.

Container registries are collections of container repositories, and can also offer API paths and access rules.

Container repositories are a collection of related container images. The container image has everything a piece of software may need to run, including code, resources, and tools. Container repositories store container images for setup and deployment, and you use the repositories to manage, pull, and push images during cluster operations.

DKP supports operation with several [local registry tools](#) (see page 1533).

7.10.6 Subnets and Pods

Some subnets are reserved by Kubernetes and can prevent proper cluster deployment if you unknowingly configure DKP so that the Node subnet collides with either the Pod or Service subnet.



Ensure your subnets do not overlap with your host subnet because they cannot be changed after cluster creation. If you need to change the Kubernetes subnets, you must do this at cluster creation.

The default subnets used in DKP are:

```
spec:
  clusterNetwork:
    pods:
```

```
cidrBlocks:
- 192.168.0.0/16
services:
  cidrBlocks:
  - 10.96.0.0/12
```

In Konvoy, the default pod subnet is 192.168.0.0/16, and the default service subnet is 10.96.0.0/12. If you wish to change the subnets you can do so with the following steps:

1. Generate the YAML manifests for the cluster using the `--dry-run` and `-o yaml` flags, along with the desired `dkp cluster create` command:

```
dkp create cluster preprovisioned --cluster-name ${CLUSTER_NAME} --control-
plane-endpoint-host <control plane endpoint host> --control-plane-endpoint-port
<control plane endpoint port, if different than 6443> --dry-run -o yaml >
cluster.yaml
```

2. To modify the service subnet, add or edit the `spec.clusterNetwork.services.cidrBlocks` field of the `Cluster` object:

```
kind: Cluster
spec:
  clusterNetwork:
    services:
      cidrBlocks:
      - 10.0.0.0/12
```

3. To modify the pod subnet, edit the `Cluster` and `calico-cni ConfigMap` resources:

`Cluster`: Add or edit the `spec.clusterNetwork.pods.cidrBlocks` field:

```
kind: Cluster
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
      - 172.16.0.0/16
```

`ConfigMap`: Edit the `data."custom-`

`resources.yaml".spec.calicoNetwork.ipPools.cidr` field with your desired pod subnet:

```
apiVersion: v1
data:
  custom-resources.yaml: |
    apiVersion: operator.tigera.io/v1
```

```

kind: Installation
metadata:
  name: default
spec:
  # Configures Calico networking.
  calicoNetwork:
    # Note: The ipPools section cannot be modified post-install.
    ipPools:
      - blockSize: 26
        cidr: 172.16.0.0/16
kind: ConfigMap
metadata:
  name: calico-cni-<cluster-name>

```

When you provision the cluster, the configured pod and service subnets will be applied.

7.10.6.1 Next Step:

If using the [Custom Installation and Additional Infrastructure Tools](#) (see page 1039) instructions, proceed to the **infrastructure provider** you are using.

If using the [Day 1 - Basic Installs by Infrastructure](#) (see page 147) instructions, proceed (or return) to that section to install and setup DKP based on your **infrastructure environment provider**.

7.11 AWS Infrastructure

7.11.1 Configuration Types

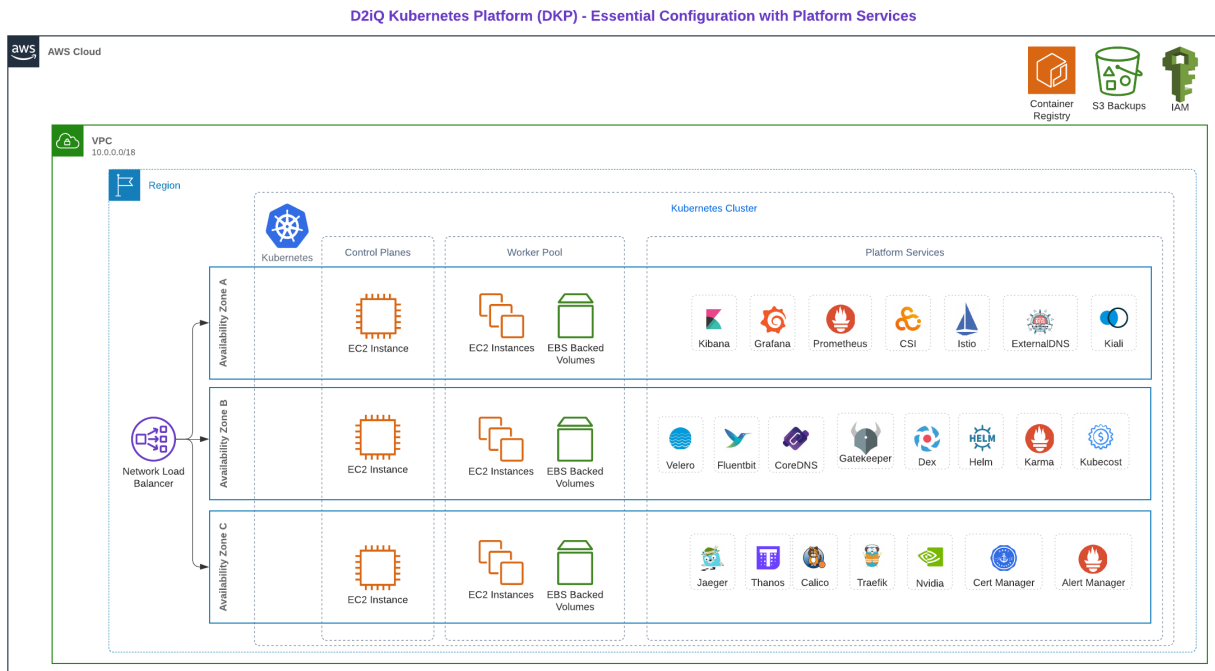
When installing DKP on AWS infrastructure, you can choose from multiple configuration types. The different types of AWS configuration types supported in DKP are listed below.

- [AWS Prerequisites](#) (see page 1051)
- [AWS Konvoy Image Builder](#) (see page 1053)
- [Minimal Permissions and Role to Create Clusters](#) (see page 1057)
- [Cluster IAM Policies, Roles, and Artifacts](#) (see page 1063)
- [AWS Cluster Creation Choices](#) (see page 1071)
- [Delete an AWS Cluster](#) (see page 1102)
- [Multiple AWS Accounts](#) (see page 1107)
- [GPUs in an AWS environment](#) (see page 1110)
- [Manage AWS Node Pools](#) (see page 1113)
- [AWS Certificate Renewal](#) (see page 1121)
- [Configure Infrastructure in UI](#) (see page 1123)
- [Replace an AWS Node](#) (see page 1124)

7.11.2 AWS Diagrams


The following diagrams show the two different ways you can implement DKP Essential and DKP Enterprise on AWS.

This diagram shows the granular detail of a single Kubernetes cluster running in AWS Cloud:



* Platform Services are not restricted to specific availability zones

This diagram shows a higher-level view of DKP Enterprise, and assumes a multi-cluster environment, where each cluster might look like the single cluster example above:



7.11.3 AWS Pricing Considerations

Deploying AWS services can incur costs to your organization, depending on how and what you deploy. For more information, see the [AWS Pricing Calculator](https://calculator.aws/#/)⁷⁵⁷.

7.11.4 AWS Service Limits

When using DKP on AWS, you need to be aware of the possibility of errors due to AWS service limits. For more information, see the [AWS Service Limits](https://aws.amazon.com/premiumsupport/knowledge-center/manage-service-limits/)⁷⁵⁸.

⁷⁵⁷ <https://calculator.aws/#/>

⁷⁵⁸ <https://aws.amazon.com/premiumsupport/knowledge-center/manage-service-limits/>

7.11.5 AWS Prerequisites

7.11.5.1 Konvoy Prerequisites

Before you begin using Konvoy, you must have:

- An x86_64-based Linux or macOS machine.
- The `dkp` binary for Linux, or macOS.
- A Container engine/runtime installed is required to install DKP:
 - Version [Docker®](#)⁷⁵⁹ container engine version 18.09.2 or higher installed for Linux or MacOS - On macOS, Docker runs in a virtual machine which needs configured with at least 8 GB of memory.
 - Version 4.0 of [Podman](#)⁷⁶⁰ or higher for Linux. Host requirements found here: [Host Requirements](#)⁷⁶¹
- [kubect](#)⁷⁶² for interacting with the running cluster.
- A valid AWS account with [credentials configured](#)⁷⁶³.
- For air-gapped environment:
 - Linux machine (bastion) that has access to the existing VPC.
 - The `dkp` binary on the bastion.
 - [kubect](#)⁷⁶⁴ for interacting with the running cluster on the bastion.
 - An existing [local registry](#) (see page 1533).
 - Ability to download artifacts from the internet and then copy those onto your bootstrap machine.
 - An [AWS Air-Gapped Machine Image](#) (see page 1474)



On macOS, Docker runs in a virtual machine. Configure this virtual machine with at least 8GB of memory.

7.11.5.2 Control Plane Nodes

You should have at least three control plane nodes. Each control plane node should have at least:

⁷⁵⁹ <https://docs.docker.com/get-docker/>

⁷⁶⁰ <https://podman.io/getting-started/installation>

⁷⁶¹ <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>

⁷⁶² <https://kubernetes.io/docs/tasks/tools/#kubectl>

⁷⁶³ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

⁷⁶⁴ <https://kubernetes.io/docs/tasks/tools/#kubectl>

- 4 cores
- 16 GiB memory
- Approximately 80 GiB of free space for the volume used for `/var/lib/kubelet` and `/var/lib/containerd`.
- Disk usage must be below 85% on the root volume.

DKP on AWS defaults to deploying an `m5.xlarge` instance with an 80GiB root volume for control plane nodes, which meets the above requirements.

7.11.5.3 Worker Nodes

You should have at least four worker nodes. The specific number of worker nodes required for your environment can vary depending on the cluster workload and size of the nodes. Each worker node should have at least:

- 8 cores
- 32 GiB memory
- Around 80 GiB of free space for the volume used for `/var/lib/kubelet` and `/var/lib/containerd`.
- Disk usage must be below 85% on the root volume.

DKP on AWS defaults to deploying a `m5.2xlarge` instance with an 80GiB root volume for worker nodes, which meets the above requirements.

If you use these instructions to create a cluster on AWS using the DKP default settings without any edits to configuration files or additional flags, your cluster is deployed on an [Ubuntu 20.04 operating system image](#) (see page 69) with 3 control plane nodes, and 4 worker nodes which match the requirements above.



Using these default images work, but due to missing optimizations, the created cluster will have certain limits. We suggest using [Konvoy Image Builder to create a custom AMI](#)⁷⁶⁵ to take advantage of enhanced cluster operations.

7.11.5.4 AWS Prerequisites

Before you begin using Konvoy with AWS, you must:

- Create [Minimal Permissions and Role to Create Clusters](#) (see page 1057).
- Create [AWS Cluster IAM Policies, Roles, and Artifacts](#)⁷⁶⁶.
- For multi-tenancy, every tenant should be in a different AWS account to ensure they are truly independent of other tenants in order to enforce security.
- Export the AWS region where you want to deploy the cluster:

⁷⁶⁵ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/143891417>

⁷⁶⁶ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29559226/AWS+Cluster+IAM+Policies%2C+Roles%2C+and+Artifacts>

```
export AWS_REGION=us-west-2
```

- Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```

7.11.6 AWS Konvoy Image Builder

You must have at least one image before creating a new cluster. As long as you have an image, this step in your configuration is not required each time since that image can be used to spin up a new cluster. However, if you need different images for different environments or providers, you will need to create a new custom image.

The following section describes how to use Konvoy Image Builder (KIB) with Amazon Web Services (AWS). There are two options:

- [1. Create a Custom AMI \(see page 0\)](#)
- [2. Air-gapped AMI \(see page 1056\)](#)



[Konvoy Image Builder \(see page 1461\)](#) has a more detailed section in the documentation if you need to refer there for compatible versions with DKP and other specific information.

7.11.6.1 Learn how to build a custom AMI for use with DKP

This procedure describes how to use the [Konvoy Image Builder \(see page 1461\)](#) (KIB) to create a [Cluster API](#)⁷⁶⁷ compliant Amazon Machine Image (AMI). KIB uses [variable overrides \(see page 1513\)](#) to specify base image and container images to use in your new AMI.



The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to create a custom AMI to take advantage of enhanced cluster operations. Explore the [KIB with AWS \(see page 1465\)](#) topics for more options.

7.11.6.2 Prerequisites


Before you begin, you must:

⁷⁶⁷ <https://cluster-api.sigs.k8s.io/>

- Download the [KIB \(see page 0\)](#) bundle for your version of DKP prefixed with `konvoy-image-bundle` for your OS.
- Check the [Supported Infrastructure Operating Systems \(see page 69\)](#)
- Check the [Supported Kubernetes Version \(see page 1552\)](#) for your Provider.
- Create a working registry:
 - Version 4.0 of [Podman⁷⁶⁸](#) or higher for Linux. Host requirements found here: <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>
 - Version 20.10.0 of [Docker⁷⁶⁹](#) or higher for Linux or MacOS
- Ensure you have met the minimal set of permissions from the [AWS Image Builder Book⁷⁷⁰](#).
- A [Minimal IAM Permissions for KIB \(see page 1465\)](#) to create an Image for an AWS account using Konvoy Image Builder.

7.11.6.3 Extract KIB Bundle

Extract the bundle and `cd` into the extracted `konvoy-image-bundle-$VERSION` folder. The bundled version of `konvoy-image` contains an embedded `docker` image that contains all the requirements for building.

 The `konvoy-image` binary and all supporting folders are also extracted. When run, `konvoy-image` binds mounts the current working directory (`$(PWD)`) into the container to be used.

- Set environment variables for [AWS access⁷⁷¹](#). The following variables must be set using your credentials including [required IAM \(see page 1465\)](#):

```
export AWS_ACCESS_KEY_ID
export AWS_SECRET_ACCESS_KEY
export AWS_DEFAULT_REGION
```

- If you have an [override file \(see page 1513\)](#) to configure specific attributes of your AMI file, add it.

7.11.6.4 1. Create a Custom AMI

Depending on which [version of DKP \(see page 1461\)](#) you are running, steps and flags will be different. To deploy in a region where CAPI images are not provided, you need to use KIB to create your own image for the region. For a list of supported AWS regions, refer to the [Published AMI⁷⁷²](#) information from AWS.

⁷⁶⁸ <https://podman.io/getting-started/installation>

⁷⁶⁹ <https://www.docker.com/>

⁷⁷⁰ <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>

⁷⁷¹ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-envvars.html>

⁷⁷² <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

7.11.6.4.1 Execute the following to begin image creation:

Run the `konvoy-image` command to build and validate the image.

```
konvoy-image build aws images/ami/centos-79.yaml
```

By default it builds in the `us-west-2` region. to specify another region set the `--region` flag:

```
konvoy-image build aws --region us-east-1 images/ami/centos-79.yaml
```

- Ensure you have named the correct YAML file for your OS in the `konvoy-image build` command.

After KIB provisions the image successfully, the `ami` id is printed and written to the `manifest.json` file. This file has an `artifact_id` field whose value provides the name of the AMI ID as shown in the example below:

```
{
  "name": "rhel-7.9-fips",
  "builder_type": "amazon-eks",
  "build_time": 1659486130,
  "files": null,
  "artifact_id": "us-west-2:ami-0f2ef742482e1b829",
  "packer_run_uuid": "0ca500d9-a5f0-815c-6f12-aceb4d46645b",
  "custom_data": {
    "containerd_version": "",
    "distribution": "RHEL",
    "distribution_version": "7.9",
    "kubernetes_cni_version": "",
    "kubernetes_version": "1.26.14+fips.0"
  }
}
```

To use a custom AMI when creating your cluster, you must create that AMI using KIB first. Then perform the export and name the custom AMI for use in the command `dkp create cluster`


```
export AWS_AMI_ID=ami-<ami-id-here>
```

7.11.6.5 2. Air-gapped AMI

7.11.6.5.1 Create an Image using Konvoy Image Builder (KIB) for use in an air-gapped cluster

7.11.6.6 Prerequisites

- [Minimal IAM Permissions for KIB \(see page 1465\)](#) to create an Image for an AWS account using Konvoy Image Builder.

 The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to create an AWS Air-gapped AMI to take advantage of enhanced cluster operations. Explore the [KIB with AWS \(see page 1465\)](#) topics for more options.

Using [KIB \(see page 1461\)](#), you can build an AMI without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2/kib
```

2. Follow the instructions below to build an AMI.

Depending on which [version of DKP \(see page 1461\)](#) you are running, steps and flags will be different. To deploy in a region where CAPI images are not provided, you need to use KIB to create your own image for the region. For a list of supported AWS regions, refer to the [Published AMI⁷⁷³](#) information from AWS.

7.11.6.6.1 Execute the following to begin image creation:

Run the `konvoy-image` command to build and validate the image.

```
konvoy-image build aws images/ami/centos-79.yaml --overrides overrides/offline.yaml
```

By default it builds in the `us-west-2` region. to specify another region set the `--region` flag:

⁷⁷³ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

```
konvoy-image build aws --region us-east-1 images/ami/centos-79.yaml --overrides
overrides/offline.yaml
```

When the command is complete the `ami` id is printed and written to `./manifest.json`.

7.11.6.7 Next Step:

[Minimal Permissions and Role to Create Clusters](#) (see page 1057)

7.11.7 Minimal Permissions and Role to Create Clusters

Configure IAM Prerequisites before starting a cluster

This section guides you in creating and using a minimally-scoped policy to create DKP clusters on an AWS account. For multi-tenancy, every tenant should be in a different AWS account to ensure they are truly independent of other tenants in order to enforce security.

7.11.7.1 Prerequisites

Before applying the IAM Policies, verify the following:

- You have a valid AWS account with [credentials configured](#)⁷⁷⁴ that can manage CloudFormation Stacks, IAM Policies, IAM Roles, and IAM Instance Profiles.
- The [AWS CLI utility](#)⁷⁷⁵ is installed.

7.11.7.2 Minimal Permissions

The following is an AWS CloudFormation stack that creates:

- A policy named `dkp-bootstrapper-policy` that enumerates the minimal permissions for a user that can create dkp aws clusters.
- A role named `dkp-bootstrapper-role` that uses the `dkp-bootstrapper-policy` with a trust policy to allow IAM users and ec2 instances from `MYAWSACCOUNTID` to use the role via STS.
- An instance profile `DKPBootstrapInstanceProfile` that wraps the `dkp-bootstrapper-role` to be used by ec2 instances.

7.11.7.3 Create Resources in CloudFormation Stack

To create the resources in the CloudFormation stack:

1. Copy the following contents into a file:

⁷⁷⁴ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

⁷⁷⁵ <https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html>

```

AWSTemplateFormatVersion: 2010-09-09
Resources:
  AWSIAMInstanceProfileDKPBootstrapper:
    Properties:
      InstanceProfileName: DKPBootstrapInstanceProfile
      Roles:
        - Ref: DKPBootstrapRole
      Type: AWS::IAM::InstanceProfile
  AWSIAMManagedPolicyDKPBootstrapper:
    Properties:
      Description: Minimal policy to create dkp clusters in AWS
      ManagedPolicyName: dkp-bootstrapper-policy
      PolicyDocument:
        Statement:
          - Action:
              - ec2:AllocateAddress
              - ec2:AssociateRouteTable
              - ec2:AttachInternetGateway
              - ec2:AuthorizeSecurityGroupIngress
              - ec2:CreateInternetGateway
              - ec2:CreateNatGateway
              - ec2:CreateRoute
              - ec2:CreateRouteTable
              - ec2:CreateSecurityGroup
              - ec2:CreateSubnet
              - ec2:CreateTags
              - ec2:CreateVpc
              - ec2:ModifyVpcAttribute
              - ec2>DeleteInternetGateway
              - ec2>DeleteNatGateway
              - ec2>DeleteRouteTable
              - ec2>DeleteSecurityGroup
              - ec2>DeleteSubnet
              - ec2>DeleteTags
              - ec2>DeleteVpc
              - ec2:DescribeAccountAttributes
              - ec2:DescribeAddresses
              - ec2:DescribeAvailabilityZones
              - ec2:DescribeInstanceTypes
              - ec2:DescribeInternetGateways
              - ec2:DescribeImages
              - ec2:DescribeNatGateways
              - ec2:DescribeNetworkInterfaces
              - ec2:DescribeNetworkInterfaceAttribute
              - ec2:DescribeRouteTables
              - ec2:DescribeSecurityGroups
              - ec2:DescribeSubnets
              - ec2:DescribeVpcs
              - ec2:DescribeVpcAttribute
              - ec2:DescribeVolumes
              - ec2:DetachInternetGateway

```

- ec2:DisassociateRouteTable
 - ec2:DisassociateAddress
 - ec2:ModifyInstanceAttribute
 - ec2:ModifyInstanceMetadataOptions
 - ec2:ModifyNetworkInterfaceAttribute
 - ec2:ModifySubnetAttribute
 - ec2:ReleaseAddress
 - ec2:RevokeSecurityGroupIngress
 - ec2:RunInstances
 - ec2:TerminateInstances
 - tag:GetResources
 - elasticloadbalancing:AddTags
 - elasticloadbalancing:CreateLoadBalancer
 - elasticloadbalancing:ConfigureHealthCheck
 - elasticloadbalancing>DeleteLoadBalancer
 - elasticloadbalancing:DescribeLoadBalancers
 - elasticloadbalancing:DescribeLoadBalancerAttributes
 - elasticloadbalancing:DescribeTargetGroups
 - elasticloadbalancing:ApplySecurityGroupsToLoadBalancer
 - elasticloadbalancing:DescribeTags
 - elasticloadbalancing:ModifyLoadBalancerAttributes
 - elasticloadbalancing:RegisterInstancesWithLoadBalancer
 - elasticloadbalancing:DeregisterInstancesFromLoadBalancer
 - elasticloadbalancing:RemoveTags
 - autoscaling:DescribeAutoScalingGroups
 - autoscaling:DescribeInstanceRefreshes
 - ec2:CreateLaunchTemplate
 - ec2:CreateLaunchTemplateVersion
 - ec2:DescribeLaunchTemplates
 - ec2:DescribeLaunchTemplateVersions
 - ec2>DeleteLaunchTemplate
 - ec2>DeleteLaunchTemplateVersions
 - ec2:DescribeKeyPairs
- Effect: Allow
- Resource:
- '*'
- Action:
 - autoscaling:CreateAutoScalingGroup
 - autoscaling:UpdateAutoScalingGroup
 - autoscaling:CreateOrUpdateTags
 - autoscaling:StartInstanceRefresh
 - autoscaling>DeleteAutoScalingGroup
 - autoscaling>DeleteTags
- Effect: Allow
- Resource:
- arn:*:autoscaling:*:*:autoScalingGroup:*:autoScalingGroupName/*
- Action:
 - ecr:DescribeRepositories
 - ecr:CreateRepository
 - ecr:PutLifecyclePolicy
 - ecr:CompleteLayerUpload
 - ecr:GetAuthorizationToken

```

- ecr:UploadLayerPart
- ecr:InitiateLayerUpload
- ecr:BatchCheckLayerAvailability
- ecr:BatchGetImage
- ecr:GetDownloadUrlForLayer
- ecr:PutImage
Effect: Allow
Resource:
- arn:aws:ecr:*:MYAWSACCOUNT:repository/*
- Action:
- iam:CreateServiceLinkedRole
Condition:
  StringLike:
    iam:AWSServiceName: autoscaling.amazonaws.com
Effect: Allow
Resource:
- arn:*:iam::*:role/aws-service-role/autoscaling.amazonaws.com/
AWSServiceRoleForAutoScaling
- Action:
- iam:CreateServiceLinkedRole
Condition:
  StringLike:
    iam:AWSServiceName: elasticloadbalancing.amazonaws.com
Effect: Allow
Resource:
- arn:*:iam::*:role/aws-service-role/
elasticloadbalancing.amazonaws.com/AWSServiceRoleForElasticLoadBalancing
- Action:
- iam:CreateServiceLinkedRole
Condition:
  StringLike:
    iam:AWSServiceName: spot.amazonaws.com
Effect: Allow
Resource:
- arn:*:iam::*:role/aws-service-role/spot.amazonaws.com/
AWSServiceRoleForEC2Spot
- Action:
- iam:PassRole
Effect: Allow
Resource:
- arn:*:iam::*:role/*.cluster-api-provider-aws.sigs.k8s.io
- Action:
- secretsmanager:CreateSecret
- secretsmanager>DeleteSecret
- secretsmanager:TagResource
Effect: Allow
Resource:
- arn:*:secretsmanager::*:secret:aws.cluster.x-k8s.io/*
Version: 2012-10-17
Roles:
- Ref: DKPBootstrapRole
Type: AWS::IAM::ManagedPolicy

```

```

DKPBootstrapRole:
  Properties:
    AssumeRolePolicyDocument:
      Statement:
        - Action:
            - sts:AssumeRole
          Effect: Allow
          Principal:
            Service:
              - ec2.amazonaws.com
        - Action:
            - sts:AssumeRole
          Effect: Allow
          Principal:
            AWS: arn:aws:iam::MYAWSACCOUNT:root
      Version: 2012-10-17
    RoleName: dkp-bootstrapper-role
    Type: AWS::IAM::Role

```

2. Replace the following with the correct values:
 - a. `MYFILENAME.yaml` - give your file a meaningful name.
 - b. `MYSTACKNAME` - give your cloudformation stack a meaningful name.
 - c. `MYAWSACCOUNT` - replace with an AWS Account ID number such as: `111122223333`
3. Run the following command to create the stack :

```
aws cloudformation create-stack --template-body=file://MYFILENAME.yaml --stack-name=MYSTACKNAME --capabilities CAPABILITY_NAMED_IAM
```

7.11.7.4 Leverage the Role

Use temporary User Access Keys via STS.

The created `dkp-bootstrapper-role` can be assumed by IAM users for temporary credentials via STS by running the command below:

```
aws sts assume-role --role-arn arn:aws:iam::MYAWSACCOUNT:role/dkp-bootstrapper-role --role-session-name EXAMPLE
```

Which returns something similar to this:

```
{
  "Credentials": {
    "AccessKeyId": "ASIA6RTF53ZH5B52EVM5",
```



```

    "SecretAccessKey": "BSsylvSsdfJY74jubsadfdsafdsaH7x1L+8Vk/",
    "SessionToken": "IQoJb3JpZ2Z5cyChb9PtJvP0S6KAi",
    "Expiration": "2022-07-14T20:19:13+00:00"
  },
  "AssumedRoleUser": {
    "AssumedRoleId": "ASIA6RTF53ZH5B52EVM5:test",
    "Arn": "arn:aws:sts::MYAWSACCOUNTID:assumed-role/dkp-bootstrapper-role/test"
  }
}

```

And then `export` the following environment variables with the results:

```

export AWS_ACCESS_KEY_ID=(.Credentials.AccessKeyId)

export AWS_SECRET_ACCESS_KEY=(.Credentials.SecretAccessKey)

export AWS_SESSION_TOKEN=(.Credentials.SessionToken)

```

 **These credentials are short lived and would need to be updated in the bootstrap cluster**

7.11.7.5 Use EC2 Instance Profiles

The created `dkp-bootstrapper-role` can be assumed by an ec2 instance a user would run `dkp create cluster` commands from. To do this, specify the IAM Instance Profile `DKPBootstrapInstanceProfile` on creation.

7.11.7.6 Use Access Keys

AWS administrators can attach the `dkp-bootstrapper-policy` to an [existing IAM user](#)⁷⁷⁶ and authenticate with [Access Keys](#)⁷⁷⁷ on the work station they would run `dkp create cluster` commands from by exporting the following environment variables with the appropriate values for the IAM user.

```

export AWS_ACCESS_KEY_ID=AKIAIOSFODNN7EXAMPLE
export AWS_SECRET_ACCESS_KEY=wJalrXUtnFEMI/K7MDENG/bPxrFiCYEXAMPLEKEY
export AWS_DEFAULT_REGION=us-west-2

```

In regards to Access Keys usage, a system administrator should always consider AWS's [Best practices](#)⁷⁷⁸.

⁷⁷⁶ https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_manage-attach-detach.html

⁷⁷⁷ https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html

⁷⁷⁸ <https://docs.aws.amazon.com/general/latest/gr/aws-access-keys-best-practices.html>

- = EKS cluster minimal permissions are required if attaching an EKS cluster. Refer to [Minimal User Permission for EKS Cluster Creation](#) (see page 1129). The CloudFormation stack on that page adds a policy named `eks-bootstrapter` to manage EKS cluster to the `dkp-bootstrapter-role` created by the CloudFormation stack on this page.

If your organization uses encrypted AMI's (<https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIEncryption.html>), then you will need to add additional permissions to the control plane policy to allow access to the Amazon Key Management Services. See the following documentation for information on the necessary policies you may need: [AWS Key Policies](#)⁷⁷⁹.

7.11.7.7 Next Step:

[Cluster IAM Policies, Roles, and Artifacts](#) (see page 1063)

7.11.8 Cluster IAM Policies, Roles, and Artifacts

This guides a DKP user in creating IAM Policies and Instance Profiles used by the cluster's control plane and worker nodes using the provided AWS CloudFormation Stack. For multi-tenancy, every tenant should be in a different AWS account to ensure they are truly independent of other tenants in order to enforce security.

7.11.8.1 Prerequisites

Before applying the IAM Policies, verify the following:

- You have a valid AWS account with [credentials configured](#)⁷⁸⁰ that can manage CloudFormation Stacks, IAM Policies, IAM Roles, and IAM Instance Profiles.
- You have the [AWS CLI utility installed](#)⁷⁸¹.

Below is information regarding the set up for policies and roles. After reading the information for each of these areas, you will find the CloudFormation Stack that creates these policies and roles in the IAM Artifacts drop-down section below:

Policies

- `AWSIAMManagedPolicyCloudProviderControlPlane` enumerates the Actions required by the workload cluster control plane machines. It is attached to the `AWSIAMRoleControlPlane` Role.
- `AWSIAMManagedPolicyCloudProviderNodes` enumerates the Actions required by the workload cluster worker machines. It is attached to the `AWSIAMRoleNodes` Role.

⁷⁷⁹ <https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-default.html#key-policy-service-integration>

⁷⁸⁰ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-files.html>

⁷⁸¹ <https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html>

3. `AWSIAMManagedPolicyControllers` enumerates the Actions required by the workload cluster worker machines. It is attached to the `AWSIAMRoleControlPlane` Role.

Roles

1. `AWSIAMRoleControlPlane` is the Role associated with the `AWSIAMInstanceProfileControlPlane` Instance Profile.
2. `AWSIAMRoleNodes` is the Role associated with the `AWSIAMInstanceProfileNodes` Instance Profile.

For more information on learning how to grant cluster access to IAM users and roles, see the official [AWS Documentation](#)⁷⁸².

IAM Artifacts

7.11.8.2 Instance Profiles

1. `AWSIAMInstanceProfileControlPlane`, assigned to workload cluster control plane machines.

ⓘ If the name is changed from the default, used below, it must be passed to `dkp create cluster` with the `--control-plane-iam-instance-profile` flag.

2. `AWSIAMInstanceProfileNodes`, assigned to workload cluster worker machines.

ⓘ If the name is changed from the default, used below, it must be passed to `dkp create cluster` with the `--worker-iam-instance-profile` flag.

```
AWSTemplateFormatVersion: 2010-09-09
Resources:
  AWSIAMInstanceProfileControlPlane:
    Properties:
      InstanceProfileName: control-plane.cluster-api-provider-aws.sigs.k8s.io
      Roles:
        - Ref: AWSIAMRoleControlPlane
    Type: AWS::IAM::InstanceProfile
  AWSIAMInstanceProfileNodes:
    Properties:
```

⁷⁸² <https://docs.aws.amazon.com/eks/latest/userguide/add-user-role.html>

```

InstanceProfileName: nodes.cluster-api-provider-aws.sigs.k8s.io
Roles:
- Ref: AWSIAMRoleNodes
Type: AWS::IAM::InstanceProfile
AWSIAMManagedPolicyCloudProviderControlPlane:
Properties:
  Description: For the Kubernetes Cloud Provider AWS Control Plane
  ManagedPolicyName: control-plane.cluster-api-provider-aws.sigs.k8s.io
  PolicyDocument:
    Statement:
      - Action:
        - autoscaling:DescribeAutoScalingGroups
        - autoscaling:DescribeLaunchConfigurations
        - autoscaling:DescribeTags
        - ec2:DescribeInstances
        - ec2:DescribeImages
        - ec2:DescribeRegions
        - ec2:DescribeRouteTables
        - ec2:DescribeSecurityGroups
        - ec2:DescribeSubnets
        - ec2:DescribeVolumes
        - ec2:CreateSecurityGroup
        - ec2:CreateTags
        - ec2:CreateVolume
        - ec2:ModifyInstanceAttribute
        - ec2:ModifyVolume
        - ec2:AttachVolume
        - ec2:AuthorizeSecurityGroupIngress
        - ec2:CreateRoute
        - ec2>DeleteRoute
        - ec2>DeleteSecurityGroup
        - ec2>DeleteVolume
        - ec2:DetachVolume
        - ec2:RevokeSecurityGroupIngress
        - ec2:DescribeVpcs
        - elasticloadbalancing:AddTags
        - elasticloadbalancing:AttachLoadBalancerToSubnets
        - elasticloadbalancing:ApplySecurityGroupsToLoadBalancer
        - elasticloadbalancing:CreateLoadBalancer
        - elasticloadbalancing:CreateLoadBalancerPolicy
        - elasticloadbalancing:CreateLoadBalancerListeners
        - elasticloadbalancing:ConfigureHealthCheck
        - elasticloadbalancing>DeleteLoadBalancer
        - elasticloadbalancing>DeleteLoadBalancerListeners
        - elasticloadbalancing:DescribeLoadBalancers
        - elasticloadbalancing:DescribeLoadBalancerAttributes
        - elasticloadbalancing:DetachLoadBalancerFromSubnets
        - elasticloadbalancing:DeregisterInstancesFromLoadBalancer
        - elasticloadbalancing:ModifyLoadBalancerAttributes
        - elasticloadbalancing:RegisterInstancesWithLoadBalancer
        - elasticloadbalancing:SetLoadBalancerPoliciesForBackendServer
        - elasticloadbalancing:AddTags

```

```

- elasticloadbalancing:CreateListener
- elasticloadbalancing:CreateTargetGroup
- elasticloadbalancing>DeleteListener
- elasticloadbalancing>DeleteTargetGroup
- elasticloadbalancing:DescribeListeners
- elasticloadbalancing:DescribeLoadBalancerPolicies
- elasticloadbalancing:DescribeTargetGroups
- elasticloadbalancing:DescribeTargetHealth
- elasticloadbalancing:ModifyListener
- elasticloadbalancing:ModifyTargetGroup
- elasticloadbalancing:RegisterTargets
- elasticloadbalancing:SetLoadBalancerPoliciesOfListener
- iam:CreateServiceLinkedRole
- kms:DescribeKey
- kms:CreateGrant
Effect: Allow
Resource:
- '*'

Version: 2012-10-17
Roles:
- Ref: AWSIAMRoleControlPlane
Type: AWS::IAM::ManagedPolicy
AWSIAMManagedPolicyCloudProviderNodes:
Properties:
Description: For the Kubernetes Cloud Provider AWS nodes
ManagedPolicyName: nodes.cluster-api-provider-aws.sigs.k8s.io
PolicyDocument:
Statement:
- Action:
  - ec2:DescribeInstances
  - ec2:DescribeRegions
  - ecr:GetAuthorizationToken
  - ecr:BatchCheckLayerAvailability
  - ecr:GetDownloadUrlForLayer
  - ecr:GetRepositoryPolicy
  - ecr:DescribeRepositories
  - ecr:ListImages
  - ecr:BatchGetImage
Effect: Allow
Resource:
- '*'
- Action:
  - secretsmanager>DeleteSecret
  - secretsmanager:GetSecretValue
Effect: Allow
Resource:
- arn::*:secretsmanager::*:secret:aws.cluster.x-k8s.io/*
- Action:
  - ssm:UpdateInstanceInformation
  - ssmmessages:CreateControlChannel
  - ssmmessages:CreateDataChannel
  - ssmmessages:OpenControlChannel

```

```

- ssm:messages:OpenDataChannel
- s3:GetEncryptionConfiguration
Effect: Allow
Resource:
- '*'
Version: 2012-10-17
Roles:
- Ref: AWSIAMRoleControlPlane
- Ref: AWSIAMRoleNodes
Type: AWS::IAM::ManagedPolicy
AWSIAMManagedPolicyControllers:
Properties:
Description: For the Kubernetes Cluster API Provider AWS Controllers
ManagedPolicyName: controllers.cluster-api-provider-aws.sigs.k8s.io
PolicyDocument:
Statement:
- Action:
- ec2:AllocateAddress
- ec2:AssociateRouteTable
- ec2:AttachInternetGateway
- ec2:AuthorizeSecurityGroupIngress
- ec2:CreateInternetGateway
- ec2:CreateNatGateway
- ec2:CreateRoute
- ec2:CreateRouteTable
- ec2:CreateSecurityGroup
- ec2:CreateSubnet
- ec2:CreateTags
- ec2:CreateVpc
- ec2:ModifyVpcAttribute
- ec2>DeleteInternetGateway
- ec2>DeleteNatGateway
- ec2>DeleteRouteTable
- ec2>DeleteSecurityGroup
- ec2>DeleteSubnet
- ec2>DeleteTags
- ec2>DeleteVpc
- ec2:DescribeAccountAttributes
- ec2:DescribeAddresses
- ec2:DescribeAvailabilityZones
- ec2:DescribeInstanceTypes
- ec2:DescribeInternetGateways
- ec2:DescribeImages
- ec2:DescribeNatGateways
- ec2:DescribeNetworkInterfaces
- ec2:DescribeNetworkInterfaceAttribute
- ec2:DescribeRouteTables
- ec2:DescribeSecurityGroups
- ec2:DescribeSubnets
- ec2:DescribeVpcs
- ec2:DescribeVpcAttribute
- ec2:DescribeVolumes

```

- ec2:DetachInternetGateway
 - ec2:DisassociateRouteTable
 - ec2:DisassociateAddress
 - ec2:ModifyInstanceAttribute
 - ec2:ModifyInstanceMetadataOptions
 - ec2:ModifyNetworkInterfaceAttribute
 - ec2:ModifySubnetAttribute
 - ec2:ReleaseAddress
 - ec2:RevokeSecurityGroupIngress
 - ec2:RunInstances
 - ec2:TerminateInstances
 - tag:GetResources
 - elasticloadbalancing:AddTags
 - elasticloadbalancing:CreateLoadBalancer
 - elasticloadbalancing:ConfigureHealthCheck
 - elasticloadbalancing>DeleteLoadBalancer
 - elasticloadbalancing:DescribeLoadBalancers
 - elasticloadbalancing:DescribeLoadBalancerAttributes
 - elasticloadbalancing:DescribeTargetGroups
 - elasticloadbalancing:ApplySecurityGroupsToLoadBalancer
 - elasticloadbalancing:DescribeTags
 - elasticloadbalancing:ModifyLoadBalancerAttributes
 - elasticloadbalancing:RegisterInstancesWithLoadBalancer
 - elasticloadbalancing:DeregisterInstancesFromLoadBalancer
 - elasticloadbalancing:RemoveTags
 - autoscaling:DescribeAutoScalingGroups
 - autoscaling:DescribeInstanceRefreshes
 - ec2:CreateLaunchTemplate
 - ec2:CreateLaunchTemplateVersion
 - ec2:DescribeLaunchTemplates
 - ec2:DescribeLaunchTemplateVersions
 - ec2>DeleteLaunchTemplate
 - ec2>DeleteLaunchTemplateVersions
- Effect: Allow
- Resource:
- '*'
- Action:
 - autoscaling:CreateAutoScalingGroup
 - autoscaling:UpdateAutoScalingGroup
 - autoscaling:CreateOrUpdateTags
 - autoscaling:StartInstanceRefresh
 - autoscaling>DeleteAutoScalingGroup
 - autoscaling>DeleteTags
- Effect: Allow
- Resource:
- arn:*:autoscaling:*:*:autoScalingGroup:*:autoScalingGroupName/*
- Action:
 - iam:CreateServiceLinkedRole
- Condition:
- StringLike:
- iam:AWSserviceName: autoscaling.amazonaws.com
- Effect: Allow

```

    Resource:
      - arn::iam::*:role/aws-service-role/autoscaling.amazonaws.com/
AWSServiceRoleForAutoScaling
  - Action:
    - iam:CreateServiceLinkedRole
  Condition:
    StringLike:
      iam:AWSServiceName: elasticloadbalancing.amazonaws.com
  Effect: Allow
  Resource:
    - arn::iam::*:role/aws-service-role/elasticloadbalancing.amazonaws.com/
AWSServiceRoleForElasticLoadBalancing
  - Action:
    - iam:CreateServiceLinkedRole
  Condition:
    StringLike:
      iam:AWSServiceName: spot.amazonaws.com
  Effect: Allow
  Resource:
    - arn::iam::*:role/aws-service-role/spot.amazonaws.com/
AWSServiceRoleForEC2Spot
  - Action:
    - iam:PassRole
  Effect: Allow
  Resource:
    - arn::iam::*:role/*-cluster-api-provider-aws.sigs.k8s.io
  - Action:
    - secretsmanager:CreateSecret
    - secretsmanager>DeleteSecret
    - secretsmanager:TagResource
  Effect: Allow
  Resource:
    - arn::secretsmanager::*:secret:aws.cluster.x-k8s.io/*
  Version: 2012-10-17
  Roles:
    - Ref: AWSIAMRoleControlPlane
  Type: AWS::IAM::ManagedPolicy
AWSIAMRoleControlPlane:
  Properties:
    AssumeRolePolicyDocument:
      Statement:
        - Action:
          - sts:AssumeRole
        Effect: Allow
        Principal:
          Service:
            - ec2.amazonaws.com
        Version: 2012-10-17
      RoleName: control-plane.cluster-api-provider-aws.sigs.k8s.io
    Type: AWS::IAM::Role
AWSIAMRoleNodes:
  Properties:


```




```
AssumeRolePolicyDocument:
  Statement:
  - Action:
    - sts:AssumeRole
    Effect: Allow
    Principal:
      Service:
      - ec2.amazonaws.com
    Version: 2012-10-17
  RoleName: nodes.cluster-api-provider-aws.sigs.k8s.io
  Type: AWS::IAM::Role
```

To create the resources in the cloudformation stack copy the contents above into a file replacing `MYFILENAME.yaml` and `MYSTACKNAME` with the intended values. Then execute the following command:

```
aws cloudformation create-stack --template-body=file://MYFILENAME.yaml --stack-
name=MYSTACKNAME --capabilities CAPABILITY_NAMED_IAM
```

 If your organization uses [encrypted AMIs](#)⁷⁸³, then you will need to add additional permissions to the control plane policy `control-plane.cluster-api-provider-aws.sigs.k8s.io` to allow access to the Amazon Key Management Services. The code snippet shows how to add a particular key ARN that is used to encrypt and decrypt AMIs.

```
---
Action:
- kms:CreateGrant
- kms:DescribeKey
- kms:Encrypt
- kms:Decrypt
- kms:ReEncrypt*
- kms:GenerateDataKey*
Resource:
- arn:aws:kms:us-west-2:111122223333:key/key-arn
Effect: Allow
```

 EKS IAM Policies and Instance Profiles that govern who has access to the cluster are found here: [EKS Cluster IAM Permissions and Roles](#) (see page 1132). If attaching an EKS cluster, that CloudStack Formation will also need to be run to create the ARN of the bootstrapper role.

⁷⁸³ <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AMIEncryption.html>

7.11.8.3 Next Steps:

[AWS Cluster Creation Choices](#) (see page 1071)

7.11.9 AWS Cluster Creation Choices

Create a new AWS Kubernetes cluster in your AWS infrastructure with the following choices:

- [Custom AMI in Cluster Creation](#) (see page 1071)
- [AWS Non-air-gapped Create a New Cluster](#) (see page 1073)
- [AWS Air-gapped Create a New Cluster](#) (see page 1089)
- [Explore New AWS Cluster](#) (see page 1097)
- [Make the AWS Cluster Self-Managed](#) (see page 1098)
- [Creating DKP Clusters on AWS from the DKP UI](#) (see page 1101)

7.11.9.1 Custom AMI in Cluster Creation

7.11.9.1.1 Launch a DKP Cluster with a Custom AMI

To use the built `ami` with DKP, specify it with the `--ami` flag when calling cluster create.

```
dkp create cluster aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--ami AMI_ID
```

7.11.9.1.2 Launch a DKP Cluster with Custom AMI Lookup

By default, the `konvoy-image` command in KIB will name the AMI in such a way that `dkp` can discover the latest AMI for a base OS and Kubernetes version. To create a cluster that will use the latest AMI, specify the `--ami-format`, `--ami-base-os` and `--ami-owner` flags:

```
dkp create cluster aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--ami-owner AWS_ACCOUNT_ID \
--ami-base-os ubuntu-20.04 \
--ami-format 'konvoy-ami-{{.BaseOS}}-?{{.K8sVersion}}-*'
```

7.11.9.1.3 Using Custom Source AMIs

When using KIB for building machine images to Amazon, the default source AMIs that we provide are modeled by looking up an AMI based on the owner. Then we apply a filter for that operating system and version.

You can view an example of that with the provided `centos-79.yaml` snippet below:

```
download_images: true

packer:
  ami_filter_name: "CentOS Linux 7"
  ami_filter_owners: "125523088429"
  distribution: "CentOS"
  distribution_version: "7.9"
  source_ami: ""
  ssh_username: "centos"
  root_device_name: "/dev/sda1"
  ...
```

At times, a particular upstream AMI may not be available in your region or something could be renamed. Other times you want to provide a custom AMI. If this is the case, you will want to edit or create your own YAML file that looks up based on the `source_ami` field. For example, you can select images that are otherwise deprecated.

Once you select the source AMI that you want, you can declare that when running your build command:

```
konvoy-image build aws path/to/ami/centos-79.yaml --source-ami ami-0123456789
```

Alternatively, if you want to add it to your YAML file, or make your own file, you can do that as well. You add that AMI ID into the `source_ami` in the YAML file:

```
download_images: true

packer:
  ami_filter_name: ""
  ami_filter_owners: ""
  distribution: "CentOS"
  distribution_version: "7.9"
  source_ami: "ami-123456789"
  ssh_username: "centos"
  root_device_name: "/dev/sda1"
  ...
```

When you're done selecting your `source_ami`, you can build your KIB image as you would normally:

```
konvoy-image build aws path/to/ami/centos-79.yaml
```

7.11.9.2 AWS Non-air-gapped Create a New Cluster

This section provides instructions to install DKP in an AWS non-air-gapped environment with custom settings. First you create an [Bootstrap Cluster](#) (see page 1074) and then move the CAPI resources to the workload cluster and delete the bootstrap cluster.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

7.11.9.2.1 AWS Prerequisites

Before you begin using Konvoy with AWS, you must:

1. Follow the steps to create permissions and roles on the [Minimal Permissions and Role to Create Clusters](#) (see page 1060) page.
2. Create [Cluster IAM Policies and Roles](#) (see page 1066).
3. Export the AWS region where you want to deploy the cluster:

```
export AWS_REGION=us-west-2
```

4. Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```



If using AWS ECR as your local private registry, more information can be found on the [Registry Mirror Tools](#) (see page 1533) page.

To deploy a cluster with a custom image in a region where [CAPI images](#)⁷⁸⁴ are not provided, you need to use [Konvoy Image Builder](#) (see page 1056) to create your own image for the region.

⁷⁸⁴ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>



For multi-tenancy, every tenant should be in a different AWS account to ensure they are truly independent of other tenants in order to enforce security.

7.11.9.2.2 Next Step:

[AWS Bootstrap Cluster](#) (see page 1074)

7.11.9.2.3 AWS Bootstrap Cluster

To create Kubernetes clusters, Konvoy uses [Cluster API](#)⁷⁸⁵ (CAPI) controllers. These controllers run on a Kubernetes cluster. To get started, you need a *bootstrap* cluster. By default, Konvoy creates a bootstrap cluster for you in a Docker container using the Kubernetes-in-Docker ([KIND](#)⁷⁸⁶) tool.

7.11.9.2.3.1 Prerequisites

Before you begin, you must:

- Complete the steps in [Prerequisites](#) (see page 121).
- Ensure the `dkp` binary can be found in your `$PATH`.

7.11.9.2.3.2 Bootstrap Cluster Lifecycle Services

1. If an HTTP proxy is required for the bootstrap cluster, set the local `http_proxy`, `https_proxy`, and `no_proxy` environment variables. They are copied into the bootstrap cluster.
2. Create a bootstrap cluster:

```
dkp create bootstrap --kubeconfig $HOME/.kube/config
```

- ✓ Creating a bootstrap cluster
- ✓ Initializing **new** CAPI components

Konvoy creates a bootstrap cluster using [KIND](#)⁷⁸⁷ as a library. Konvoy then deploys the following [Cluster API](#)⁷⁸⁸ providers on the cluster:

⁷⁸⁵ <https://cluster-api.sigs.k8s.io/>

⁷⁸⁶ <https://github.com/kubernetes-sigs/kind>

⁷⁸⁷ <https://github.com/kubernetes-sigs/kind>

⁷⁸⁸ <https://cluster-api.sigs.k8s.io/>

- [Core Provider](#)⁷⁸⁹
- [AWS Infrastructure Provider](#)⁷⁹⁰
- [Kubeadm Bootstrap Provider](#)⁷⁹¹
- [Kubeadm ControlPlane Provider](#)⁷⁹²

Konvoy waits until the controller-manager and webhook deployments of these providers are ready. List these deployments using this command:

```
kubectl get --all-namespaces deployments -l=clusterctl.cluster.x-k8s.io
```

NAMESPACE	READY	UP-TO-DATE	AVAILABLE	AGE	NAME
capa-system	1/1	1	1	2m8s	capa-controller-manager
capi-kubeadm-bootstrap-system	1/1	1	1	2m10s	capi-kubeadm-bootstrap-controller-manager
capi-kubeadm-control-plane-system	1/1	1	1	2m10s	capi-kubeadm-control-plane-controller-manager
capi-system	1/1	1	1	2m11s	capi-controller-manager
capp-system	1/1	1	1	2m6s	capp-controller-manager
capv-system	1/1	1	1	2m5s	capv-controller-manager
capz-system	1/1	1	1	2m7s	capz-controller-manager
cert-manager	1/1	1	1	2m21s	cert-manager
cert-manager	1/1	1	1	2m21s	cert-manager-cainjector
cert-manager	1/1	1	1	2m21s	cert-manager-webhook

7.11.9.2.3.3 Using a Custom AWS CA

You need to add the custom CAs into two places:

- The `capa-controller-manager` pod, because CAPA controllers interact with AWS API when creating and deleting infrastructure.

⁷⁸⁹ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

⁷⁹⁰ <https://github.com/kubernetes-sigs/cluster-api-provider-aws>

⁷⁹¹ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

⁷⁹² <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>

- The trusted root CAs in the AWS AMI used as Kubernetes nodes. The first step of the node bootstrap process is to fetch the sensitive information from the AWS Secrets Manager service, so the `aws` client on the instances needs to trust this custom CA. This process is unique to your environment but a general flow can be similar to what is documented in [adding trusted root certificates to the server](https://manuals.gfi.com/en/kerio/connect/content/server-configuration/ssl-certificates/adding-trusted-root-certificates-to-the-server-1605.html)⁷⁹³.

1. Place the AWS CA file as `ca.pem` in your working directory
2. Create a ConfigMap with the contents of the file:

```
kubectl create configmap -n capa-system aws-ca --from-file=ca.pem
```

3. Update the `capa-controller-manager` to set an environment variable `AWS_CA_BUNDLE` in `capa-controller-manager`:

```
kubectl patch deployment -n capa-system capa-controller-manager --patch
'{"spec":{"template":{"spec":{"$setElementOrder/containers":
[{"name":"manager"}, {"name":"kube-rbac-proxy"}], "$setElementOrder/volumes":
[{"name":"cert"}, {"name":"credentials"}, {"name":"aws-ca"}], "containers":
[{"$setElementOrder/env": [{"name":"AWS_SHARED_CREDENTIALS_FILE"},
{"name":"AWS_CA_BUNDLE"}], "$setElementOrder/volumeMounts": [{"mountPath":"/tmp/
k8s-webhook-server/serving-certs"}, {"mountPath":"/home/.aws"}, {"mountPath":"/
home/.konvoy/aws-ca.pem"}], "env": [{"name":"AWS_CA_BUNDLE", "value":"/
home/.konvoy/aws-ca.pem"}], "name":"manager", "volumeMounts": [{"mountPath":"/
home/.konvoy/aws-ca.pem", "name":"aws-ca", "subPath":"ca.pem"}]}], "volumes":
[{"configMap":{"name":"aws-ca"}, {"name":"aws-ca"}]}}}]'
```

7.11.9.2.3.4 Next Step:

[Create a New Customized AWS Cluster](#) (see page 1076)

7.11.9.2.4 Create a New Customized AWS Cluster

7.11.9.2.4.1 Prerequisites

Before you begin, make sure you have created a [Bootstrap](#) (see page 1077) cluster.

Name Your Cluster

⁷⁹³ <https://manuals.gfi.com/en/kerio/connect/content/server-configuration/ssl-certificates/adding-trusted-root-certificates-to-the-server-1605.html>

The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)⁷⁹⁴ for more naming information. Follow these steps:

1. Give your cluster a unique name suitable for your environment.

In AWS it is critical that the name is unique, as no two clusters in the same AWS account can have the same name.

2. Set the environment variable:

```
export CLUSTER_NAME=<aws-example>
```

Tips and Tricks

Below are a few ways to customize your setup. If you prefer to do a basic setup, skip Tips and Tricks and proceed to [Create a New AWS Cluster](#) (see page 1079) section.

1. To get a list of names used in your AWS account, use the `aws CLI`⁷⁹⁵. After downloading, use the following command:

```
aws ec2 describe-vpcs --filter "Name=tag-key,Values=kubernetes.io/cluster" --query "Vpcs[*].Tags[?Key=='kubernetes.io/cluster'].Value | sort(@[*][0])"
```

```
"alex-aws-cluster-afe98",
"sam-aws-cluster-8if9q"
```

2. (Optional) To create a cluster name that is unique, use the following command:

```
export CLUSTER_NAME=aws-example-$(LC_CTYPE=C tr -dc 'a-z0-9' </dev/urandom | fold -w 5 | head -n1)
echo $CLUSTER_NAME
```

```
aws-example-pf4a3
```

This will create a unique name every time you run it, so use it with forethought.

3. (Optional) Configure your cluster to use an existing local registry as a mirror when attempting to pull images by adding the flag to the `dkp create cluster` command.

Below is an AWS ECR example:


⁷⁹⁴ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

⁷⁹⁵ <https://aws.amazon.com/cli/>


```
export REGISTRY_MIRROR_URL=<your_registry_url>
```

- `REGISTRY_MIRROR_URL` : the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
 - NOTE: Other local registries may use the options below:
 - JFrog - `CONTAINER_REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
 - `CONTAINER_REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
 - `CONTAINER_REGISTRY_PASSWORD` : optional if username is not set.
4. To use a custom AMI when creating your cluster, you must create that AMI using [KIB \(see page 1469\)](#) first. Then perform the export and name the custom AMI for use in the command `dkp create cluster` after this step:

```
export AWS_AMI_ID=ami-<ami-id-here>
```

 **IMPORTANT:** The AMI must be created with [Konvoy Image Builder \(see page 1461\)](#) in order to use the registry mirror feature.

5. Set the environment variable for the AMI you choose during the `dkp create cluster` command. This will output the generated manifest into a new file:

```
dkp create cluster aws
--cluster-name=${CLUSTER_NAME} \
--ami=${AWS_AMI_ID} \
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml
```

- (Optional) Allow lookup of the AMI by providing location, format and OS with the following flags instead of AMI ID `--ami=${AWS_AMI_ID}` :
 - `--ami-owner AWS_ACCOUNT_ID \`
 - `--ami-base-os ubuntu-20.04 \`
 - `--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*' \`

- (Optional) Add the registry mirror flag to the command above to pull from your local registry:
`--registry-mirror-url=<YOUR_ECR_URL>`

= If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

6. Ensure your [subnets](#) (see page 1049) do not overlap with your host subnet because they cannot be changed after cluster creation. If you need to change the kubernetes subnets, you must do this at cluster creation. The default subnets used in DKP are:

```
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
        - 192.168.0.0/16
    services:
      cidrBlocks:
        - 10.96.0.0/12
```

7. (Optional) Alternatively, you can create individual files with different smaller manifests for ease in editing using the `--output-directory` flag. This will create multiple files in the specified directory which must already exist:

```
dkp create cluster aws
--cluster-name=${CLUSTER_NAME} \
--ami=${AWS_AMI_ID} \
--dry-run \
--output=yaml \
--output-directory=<existing-directory>
```

= For more information regarding this flag or others, please refer to the CLI section of the documentation for the [dkp create cluster](#) (see page 1662) command and select your provider.

Create a New AWS Kubernetes Cluster

- By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

1. Ensure your AWS credentials are up to date. If you are using Static Credentials, use the following command to refresh the credentials. Otherwise, proceed to step 2:

```
dkp update bootstrap credentials aws
```

2. Generate the Kubernetes cluster objects. The following example shows a common configuration. See [dkp create cluster aws](#) (see page 1665) reference for the full list of cluster creation options:

- To increase [Docker Hub's rate limit](#)⁷⁹⁶ use your Docker Hub credentials when creating the cluster, by setting the following flags `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username=<your-username> --registry-mirror-password=<your-password>` on the `dkp create cluster` command.

```
dkp create cluster aws --cluster-name=${CLUSTER_NAME} \
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml
```

Generating cluster resources

- Refer to the Cluster Creation Customization Choices section for more information on how to use optional flags such as the `--output-directory` flag.

3. (Optional) To configure the Control Plane and Worker nodes to use an HTTP proxy:

⁷⁹⁶ <https://docs.docker.com/docker-hub/download-rate-limit/>

```

export CONTROL_PLANE_HTTP_PROXY=http://example.org:8080
export CONTROL_PLANE_HTTPS_PROXY=http://example.org:8080
export
CONTROL_PLANE_NO_PROXY="example.org,example.com,example.net,localhost,127.0.0.1
,10.96.0.0/12,192.168.0.0/16,kubernetes,kubernetes.default,kubernetes.default.s
vc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.sv
c.cluster,.svc.cluster.local,169.254.169.254,.elb.amazonaws.com"

export WORKER_HTTP_PROXY=http://example.org:8080
export WORKER_HTTPS_PROXY=http://example.org:8080
export
WORKER_NO_PROXY="example.org,example.com,example.net,localhost,127.0.0.1,10.96.
0.0/12,192.168.0.0/16,kubernetes,kubernetes.default,kubernetes.default.svc,kube
rnetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.clust
er,.svc.cluster.local,169.254.169.254,.elb.amazonaws.com"

```

- Replace `example.org,example.com,example.net` with you internal addresses
 - `localhost` and `127.0.0.1` addresses should not use the proxy
 - `10.96.0.0/12` is the default Kubernetes service subnet
 - `192.168.0.0/16` is the default Kubernetes pod subnet
 - `kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.defa
ult.svc.cluster,kubernetes.default.svc.cluster.local` is the internal
Kubernetes kube-apiserver service
 - `.svc,.svc.cluster,.svc.cluster.local` is the internal Kubernetes services
 - `169.254.169.254` is the AWS metadata server
 - `.elb.amazonaws.com` is for the worker nodes to allow them to communicate directly to
the kube-apiserver ELB
4. (Optional) Create a Kubernetes cluster with HTTP proxy configured. This step assumes you did not already create a cluster in the previous steps:

```

dkp create cluster aws --cluster-name=${CLUSTER_NAME} \
--control-plane-http-proxy="${CONTROL_PLANE_HTTP_PROXY}" \
--control-plane-https-proxy="${CONTROL_PLANE_HTTPS_PROXY}" \
--control-plane-no-proxy="${CONTROL_PLANE_NO_PROXY}" \
--worker-http-proxy="${WORKER_HTTP_PROXY}" \
--worker-https-proxy="${WORKER_HTTPS_PROXY}" \
--worker-no-proxy="${WORKER_NO_PROXY}" \
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml

```

[-] To increase [Docker Hub's rate limit](#)⁷⁹⁷ use your Docker Hub credentials when creating the cluster, by setting flags `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username=<your-username> --registry-mirror-password=<your-password>` when running `dkp create cluster`.

5. Inspect or edit the cluster objects:

NOTE: Familiarize yourself with Cluster API before editing the cluster objects as edits can prevent the cluster from deploying successfully.

The objects are [Custom Resources](#)⁷⁹⁸ defined by Cluster API components, and they belong in three different categories:

a. Cluster

A *Cluster* object has references to the infrastructure-specific and control plane objects.

Because this is an AWS cluster, there is an *AWSCluster* object that describes the infrastructure-specific cluster properties. Here, this means the AWS region, the VPC ID, subnet IDs, and security group rules required by the Pod network implementation.

b. Control Plane

A *KubeadmControlPlane* object describes the control plane, which is the group of machines that run the Kubernetes control plane components, which include the etcd distributed database, the API server, the core controllers, and the scheduler. The object describes the configuration for these components. The object also has a reference to an infrastructure-specific object that describes the properties of all control plane machines. Here, it references an *AWSMachineTemplate* object, which describes the instance type, the type of disk used, and the size of the disk, among other properties.

c. Node Pool

A Node Pool is a collection of machines with identical properties. For example, a cluster might have one Node Pool with large memory capacity, another Node Pool with GPU support. Each Node Pool is described by three objects: The *MachinePool* references an object that describes the configuration of Kubernetes components (for example, kubelet) deployed on each node pool machine, and an infrastructure-specific object that describes the properties of all node pool machines. Here, it references a *KubeadmConfigTemplate*, and an *AWSMachineTemplate* object, which describes the instance type, the type of disk used, the size of the disk, among other properties.

For in-depth documentation about the objects, read [Concepts](#)⁷⁹⁹ in the Cluster API Book.

6. Modify Control Plane Audit logs settings using the information contained in the page [Configuring the Control Plane](#) (see page 1538).

7. Create the cluster from the objects. A warning will appear in the console if the resource already exists and will require you to remove the resource or update your YAML.

⁷⁹⁷ <https://docs.docker.com/docker-hub/download-rate-limit/>

⁷⁹⁸ <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

⁷⁹⁹ <https://cluster-api.sigs.k8s.io/user/concepts.html>

```
kubectl create -f ${CLUSTER_NAME}.yaml
```

NOTE: If you used the `--output-directory` flag in your `dkp create .. --dry-run` step above, create the cluster from the objects you created by specifying the directory:

```
kubectl create -f <existing-directory>/
```

Output will be similar to output shown below:

```
cluster.cluster.x-k8s.io/aws-example created
awscluster.infrastructure.cluster.x-k8s.io/aws-example created
kubeadmcontrolplane.controlplane.cluster.x-k8s.io/aws-example-control-plane
created
awsmachinetemplate.infrastructure.cluster.x-k8s.io/aws-example-control-plane
created
secret/aws-example-etcd-encryption-config created
machinedeployment.cluster.x-k8s.io/aws-example-md-0 created
awsmachinetemplate.infrastructure.cluster.x-k8s.io/aws-example-md-0 created
kubeadmconfigtemplate.bootstrap.cluster.x-k8s.io/aws-example-md-0 created
clusterresourceset.addons.cluster.x-k8s.io/calico-cni-installation-aws-example
created
configmap/calico-cni-installation-aws-example created
configmap/tigera-operator-aws-example created
clusterresourceset.addons.cluster.x-k8s.io/aws-ebs-csi-aws-example created
configmap/aws-ebs-csi-aws-example created
clusterresourceset.addons.cluster.x-k8s.io/cluster-autoscaler-aws-example
created
configmap/cluster-autoscaler-aws-example created
clusterresourceset.addons.cluster.x-k8s.io/node-feature-discovery-aws-example
created
configmap/node-feature-discovery-aws-example created
clusterresourceset.addons.cluster.x-k8s.io/nvidia-feature-discovery-aws-example
created
configmap/nvidia-feature-discovery-aws-example created
```

8. Wait for the cluster control-plane to be ready:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

```
cluster.cluster.x-k8s.io/aws-example condition met
```

The `READY` status becomes `True` after the cluster control-plane becomes ready in one of the following steps.

- After the objects are created on the API server, the Cluster API controllers reconcile them. They create infrastructure and machines. As they progress, they update the Status of each object. Konvoy provides a command to describe the current status of the cluster:

```
dkp describe cluster -c ${CLUSTER_NAME}
```

```

NAME                                                                 READY
SEVERITY REASON SINCE MESSAGE
Cluster/aws-example                                               True
60s
├─ClusterInfrastructure - AWSCluster/aws-example                 True
5m23s
├─ControlPlane - KubeadmControlPlane/aws-example-control-plane  True
60s
│ └─Machine/aws-example-control-plane-55jh4                     True
4m59s
│ └─Machine/aws-example-control-plane-6sn97                     True
2m49s
│ └─Machine/aws-example-control-plane-nx9v5                     True
66s
└─Workers
  └─MachineDeployment/aws-example-md-0                            True
117s
    └─Machine/aws-example-md-0-cb9c9bbf7-hcl8z                   True
3m1s
    └─Machine/aws-example-md-0-cb9c9bbf7-rtdqw                   True
3m2s
    └─Machine/aws-example-md-0-cb9c9bbf7-t894m                   True
3m1s
    └─Machine/aws-example-md-0-cb9c9bbf7-td29r                   True
3m1s

```

- As they progress, the controllers also create Events. List the Events using this command:

```
kubectl get events | grep ${CLUSTER_NAME}
```

For brevity, the example uses `grep`. It is also possible to use separate commands to get Events for specific objects. For example, `kubectl get events --field-selector`

```
involvedObject.kind="AWSCluster" and kubectl get events --field-selector
involvedObject.kind="AWSMachine".
```

```
7m26s      Normal    SuccessfulSetNodeRef                                     machine/
aws-example-control-plane-2wb9q      ip-10-0-182-218.us-west-2.compute.internal
11m        Normal    SuccessfulCreate                                       awsmachine/aws-example-control-plane-vcjkr    Created new control-plane instance
with id "i-0dde024e80ae3de7a"
11m        Normal    SuccessfulAttachControlPlaneELB                       awsmachine/aws-example-control-plane-vcjkr    Control plane instance
"i-0dde024e80ae3de7a" is registered with load balancer
7m25s      Normal    SuccessfulDeleteEncryptedBootstrapDataSecrets         awsmachine/aws-example-control-plane-vcjkr    AWS Secret entries containing
userdata deleted
7m6s       Normal    FailedDescribeInstances                                awsmachinepool/aws-example-mp-0               No Auto Scaling Groups with aws-
example-mp-0 found
7m3s       Warning   FailedLaunchTemplateReconcile                          awsmachinepool/aws-example-mp-0               Failed to reconcile launch
template: ValidationError: AutoScalingGroup name not found - AutoScalingGroup
aws-example-mp-0 not found
7m1s       Warning   FailedLaunchTemplateReconcile                          awsmachinepool/aws-example-mp-0               Failed to reconcile launch
template: ValidationError: AutoScalingGroup name not found - AutoScalingGroup
aws-example-mp-0 not found
6m59s      Warning   FailedLaunchTemplateReconcile                          awsmachinepool/aws-example-mp-0               Failed to reconcile launch
template: ValidationError: AutoScalingGroup name not found - AutoScalingGroup
aws-example-mp-0 not found
6m57s      Warning   FailedLaunchTemplateReconcile                          awsmachinepool/aws-example-mp-0               Failed to reconcile launch
template: ValidationError: AutoScalingGroup name not found - AutoScalingGroup
aws-example-mp-0 not found
6m55s      Warning   FailedLaunchTemplateReconcile                          awsmachinepool/aws-example-mp-0               Failed to reconcile launch
template: ValidationError: AutoScalingGroup name not found - AutoScalingGroup
aws-example-mp-0 not found
6m53s      Warning   FailedLaunchTemplateReconcile                          awsmachinepool/aws-example-mp-0               Failed to reconcile launch
template: ValidationError: AutoScalingGroup name not found - AutoScalingGroup
aws-example-mp-0 not found
6m51s      Warning   FailedLaunchTemplateReconcile                          awsmachinepool/aws-example-mp-0               Failed to reconcile launch
template: ValidationError: AutoScalingGroup name not found - AutoScalingGroup
aws-example-mp-0 not found
6m49s      Warning   FailedLaunchTemplateReconcile                          awsmachinepool/aws-example-mp-0               Failed to reconcile launch
template: ValidationError: AutoScalingGroup name not found - AutoScalingGroup
aws-example-mp-0 not found
```



```

6m47s      Warning   FailedLaunchTemplateReconcile
awsmachinepool/aws-example-mp-0      Failed to reconcile launch
template: ValidationError: AutoScalingGroup name not found - AutoScalingGroup
aws-example-mp-0 not found
74s       Warning   FailedLaunchTemplateReconcile
awsmachinepool/aws-example-mp-0      (combined from similar events):
Failed to reconcile launch template: ValidationError: AutoScalingGroup name not
found - AutoScalingGroup aws-example-mp-0 not found
16m       Normal    SuccessfulCreateVPC
awscluster/aws-example                 Created new managed VPC
"vpc-032fff0fe06a85035"
16m       Normal    SuccessfulSetVPCAttributes
awscluster/aws-example                 Set managed VPC attributes for
"vpc-032fff0fe06a85035"
16m       Normal    SuccessfulCreateSubnet
awscluster/aws-example                 Created new managed Subnet
"subnet-0677a4fbd7d170dfe"
16m       Normal    SuccessfulModifySubnetAttributes
awscluster/aws-example                 Modified managed Subnet
"subnet-0677a4fbd7d170dfe" attributes
16m       Normal    SuccessfulCreateSubnet
awscluster/aws-example                 Created new managed Subnet
"subnet-04fc9deb4fa9f8333"
16m       Normal    SuccessfulCreateSubnet
awscluster/aws-example                 Created new managed Subnet
"subnet-0a266c15dd211ce6c"
16m       Normal    SuccessfulModifySubnetAttributes
awscluster/aws-example                 Modified managed Subnet
"subnet-0a266c15dd211ce6c" attributes
16m       Normal    SuccessfulCreateSubnet
awscluster/aws-example                 Created new managed Subnet
"subnet-06269d5b52d50840f"
16m       Normal    SuccessfulCreateSubnet
awscluster/aws-example                 Created new managed Subnet
"subnet-0fc41ffef7dceface"
16m       Normal    SuccessfulModifySubnetAttributes
awscluster/aws-example                 Modified managed Subnet
"subnet-0fc41ffef7dceface" attributes
16m       Normal    SuccessfulCreateSubnet
awscluster/aws-example                 Created new managed Subnet
"subnet-0725068cca16ad9f9"
16m       Normal    SuccessfulCreateInternetGateway
awscluster/aws-example                 Created new managed Internet
Gateway "igw-07cd7ad3e6c7c1ca7"
16m       Normal    SuccessfulAttachInternetGateway
awscluster/aws-example                 Internet Gateway
"igw-07cd7ad3e6c7c1ca7" attached to VPC "vpc-032fff0fe06a85035"
16m       Normal    SuccessfulCreateNATGateway
awscluster/aws-example                 Created new NAT Gateway
"nat-0a0cf17d29150cf9a"

```

```

16m          Normal    SuccessfulCreateNATGateway
awscluster/aws-example          Created new NAT Gateway
"nat-065e5e383e6f23320"
16m          Normal    SuccessfulCreateNATGateway
awscluster/aws-example          Created new NAT Gateway
"nat-01c4a6fed2a31ed4c"
13m          Normal    SuccessfulCreateRouteTable
awscluster/aws-example          Created managed RouteTable
"rtb-09f4e2eecb7462d22"
13m          Normal    SuccessfulCreateRoute
awscluster/aws-example          Created route {
13m          Normal    SuccessfulAssociateRouteTable
awscluster/aws-example          Associated managed RouteTable
"rtb-09f4e2eecb7462d22" with subnet "subnet-0677a4fbd7d170dfe"
13m          Normal    SuccessfulCreateRouteTable
awscluster/aws-example          Created managed RouteTable
"rtb-0007b98b36f37d1e4"
13m          Normal    SuccessfulCreateRoute
awscluster/aws-example          Created route {
13m          Normal    SuccessfulAssociateRouteTable
awscluster/aws-example          Associated managed RouteTable
"rtb-0007b98b36f37d1e4" with subnet "subnet-04fc9deb4fa9f8333"
13m          Normal    SuccessfulCreateRouteTable
awscluster/aws-example          Created managed RouteTable
"rtb-079a1d7d3667c2525"
13m          Normal    SuccessfulCreateRoute
awscluster/aws-example          Created route {
13m          Normal    SuccessfulAssociateRouteTable
awscluster/aws-example          Associated managed RouteTable
"rtb-079a1d7d3667c2525" with subnet "subnet-0a266c15dd211ce6c"
13m          Normal    SuccessfulCreateRouteTable
awscluster/aws-example          Created managed RouteTable
"rtb-0e5ebc8ec29848a17"
13m          Normal    SuccessfulCreateRoute
awscluster/aws-example          Created route {
13m          Normal    SuccessfulAssociateRouteTable
awscluster/aws-example          Associated managed RouteTable
"rtb-0e5ebc8ec29848a17" with subnet "subnet-06269d5b52d50840f"
13m          Normal    SuccessfulCreateRouteTable
awscluster/aws-example          Created managed RouteTable
"rtb-087f0c400675c4bce"
13m          Normal    SuccessfulCreateRoute
awscluster/aws-example          Created route {
13m          Normal    SuccessfulAssociateRouteTable
awscluster/aws-example          Associated managed RouteTable
"rtb-087f0c400675c4bce" with subnet "subnet-0fc41ffef7dceface"
13m          Normal    SuccessfulCreateRouteTable
awscluster/aws-example          Created managed RouteTable
"rtb-05a05080bbb3cead9"
13m          Normal    SuccessfulCreateRoute
awscluster/aws-example          Created route {

```

```

13m          Normal    SuccessfulAssociateRouteTable
awscluster/aws-example          Associated managed RouteTable
"rtb-05a05080bbb3cead9" with subnet "subnet-0725068cca16ad9f9"
13m          Normal    SuccessfulCreateSecurityGroup
awscluster/aws-example          Created managed SecurityGroup
"sg-0379bf77211472854" for Role "bastion"
13m          Normal    SuccessfulCreateSecurityGroup
awscluster/aws-example          Created managed SecurityGroup
"sg-0a4e0635f68a2f57d" for Role "apiserver-lb"
13m          Normal    SuccessfulCreateSecurityGroup
awscluster/aws-example          Created managed SecurityGroup
"sg-022da9dfc21ef3d5e" for Role "lb"
13m          Normal    SuccessfulCreateSecurityGroup
awscluster/aws-example          Created managed SecurityGroup
"sg-00db2e847c0b49d6e" for Role "controlplane"
13m          Normal    SuccessfulCreateSecurityGroup
awscluster/aws-example          Created managed SecurityGroup
"sg-01fe3426404f94708" for Role "node"
13m          Normal    SuccessfulAuthorizeSecurityGroupIngressRules
awscluster/aws-example          Authorized security group ingress
rules [protocol=tcp/range=[22-22]/description=SSH] for SecurityGroup
"sg-0379bf77211472854"
13m          Normal    SuccessfulAuthorizeSecurityGroupIngressRules
awscluster/aws-example          Authorized security group ingress
rules [protocol=tcp/range=[6443-6443]/description=Kubernetes API] for
SecurityGroup "sg-0a4e0635f68a2f57d"
13m          Normal    SuccessfulAuthorizeSecurityGroupIngressRules
awscluster/aws-example          Authorized security group ingress
rules [protocol=tcp/range=[5473-5473]/description=typha (calico) protocol=tcp/
range=[179-179]/description=bgp (calico) protocol=4/range=[-1-65535]/
description=IP-in-IP (calico) protocol=tcp/range=[22-22]/description=SSH
protocol=tcp/range=[6443-6443]/description=Kubernetes API protocol=tcp/range=[2
379-2379]/description=etcd protocol=tcp/range=[2380-2380]/description=etcd
peer] for SecurityGroup "sg-00db2e847c0b49d6e"
13m          Normal    SuccessfulAuthorizeSecurityGroupIngressRules
awscluster/aws-example          Authorized security group ingress
rules [protocol=tcp/range=[5473-5473]/description=typha (calico) protocol=tcp/
range=[179-179]/description=bgp (calico) protocol=4/range=[-1-65535]/
description=IP-in-IP (calico) protocol=tcp/range=[22-22]/description=SSH
protocol=tcp/range=[30000-32767]/description=Node Port Services protocol=tcp/
range=[10250-10250]/description=Kubelet API] for SecurityGroup
"sg-01fe3426404f94708"

```

Known Limitations



Be aware of these limitations in the current release of Konvoy.

- The Konvoy version used to create a bootstrap cluster must match the Konvoy version used to create a workload cluster.
- Konvoy supports deploying one workload cluster.
- Konvoy generates a set of objects for one Node Pool.
- Konvoy does not validate edits to cluster objects.

7.11.9.3 AWS Air-gapped Create a New Cluster

This section provides the instructions to install DKP in an AWS air-gapped environment with custom settings. First you create an [Air-gapped Bootstrap Cluster](#) (see page 1090) and then move the CAPI resources to the workload cluster and delete the bootstrap cluster.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

7.11.9.3.1 AWS Prerequisites

Before you begin using Konvoy with AWS, you must:

1. Follow the steps to create permissions and roles on the [Minimal Permissions and Role to Create Clusters](#) (see page 1060) page.
2. Create [Cluster IAM Policies and Roles](#) (see page 1066).
3. Export the AWS region where you want to deploy the cluster:

```
export AWS_REGION=us-west-2
```

4. Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```



If using AWS ECR as your local private registry, more information can be found on the [Registry Mirror Tools](#) (see page 1533) page.

To deploy a cluster with a custom image in a region where [CAPI images](#)⁸⁰⁰ are not provided, you need to use [Konvoy Image Builder](#) (see page 1056) to create your own image for the region.

⁸⁰⁰ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>



For multi-tenancy, every tenant should be in a different AWS account to ensure they are truly independent of other tenants in order to enforce security.



For air-gapped, ensure you have [downloaded \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, so you can extract the tarball on next page.

7.11.9.3.2 Next Step:

[AWS Air-gapped Bootstrap \(see page 1090\)](#)

7.11.9.3.3 AWS Air-gapped Bootstrap

7.11.9.3.3.1 Bootstrap a kind cluster and CAPI controllers

Konvoy deploys all cluster lifecycle services to a bootstrap cluster, which deploys a workload cluster. When the workload cluster is ready, move the cluster lifecycle services to the workload cluster, after which the workload cluster manages its own lifecycle.

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address with this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
export REGISTRY_CA=<path to the cacert file on the bastion>
```

3. Seed the registry by running the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```

4. Load the bootstrap container image on your bastion machine using Docker or Podman command:

```
docker load -i konvoy-bootstrap-image-v2.6.2.tar
```

```
podman load -i konvoy-bootstrap-image-v2.6.2.tar
```

5. Create a bootstrap cluster:

```
dkp create bootstrap --kubeconfig $HOME/.kube/config
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#)⁸⁰¹.

6. (Optional) Refresh the credentials used by the AWS provider at any time, using the command:

```
dkp update bootstrap credentials aws
```

7.11.9.3.3.2 Using a Custom AWS CA

You need to add the custom CAs into two places:

- The `capa-controller-manager` pod, because CAPA controllers interact with AWS API when creating and deleting infrastructure.
- The trusted root CAs in the AWS AMI used as Kubernetes nodes. The first step of the node bootstrap process is to fetch the sensitive information from the AWS Secrets Manager service, so the `aws` client on the instances needs to trust this custom CA. This process is unique to your environment but a general flow can be similar to what is documented in [adding trusted root certificates to the server](#)⁸⁰².

1. Place the AWS CA file as `ca.pem` in your working directory
2. Create a ConfigMap with the contents of the file:

```
kubectl create configmap -n capa-system aws-ca --from-file=ca.pem
```

3. Update the `capa-controller-manager` to set an environment variable `AWS_CA_BUNDLE` in `capa-controller-manager`:

⁸⁰¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273352012&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configuring+an+HTTP%2FHTTPS+Pr>

[oxy](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273352012&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configuring+an+HTTP%2FHTTPS+Pr)

⁸⁰² <https://manuals.gfi.com/en/kerio/connect/content/server-configuration/ssl-certificates/adding-trusted-root-certificates-to-the-server-1605.html>

```
kubectl patch deployment -n capa-system capa-controller-manager --patch
'{"spec":{"template":{"spec":{"$setElementOrder/containers":
[{"name":"manager"}, {"name":"kube-rbac-proxy"}], "$setElementOrder/volumes":
[{"name":"cert"}, {"name":"credentials"}, {"name":"aws-ca"}], "containers":
[{"$setElementOrder/env": [{"name":"AWS_SHARED_CREDENTIALS_FILE"},
{"name":"AWS_CA_BUNDLE"}], "$setElementOrder/volumeMounts": [{"mountPath":"/tmp/
k8s-webhook-server/serving-certs"}, {"mountPath":"/home/.aws"}, {"mountPath":"/
home/.konvoy/aws-ca.pem"}], "env": [{"name":"AWS_CA_BUNDLE", "value":"/
home/.konvoy/aws-ca.pem"}], "name":"manager", "volumeMounts": [{"mountPath":"/
home/.konvoy/aws-ca.pem", "name":"aws-ca", "subPath":"ca.pem"}]}], "volumes":
[{"configMap":{"name":"aws-ca"}, {"name":"aws-ca"}]}}}]}'
```

Konvoy creates a bootstrap cluster using [KIND](#)⁸⁰³ as a library. Konvoy then deploys the following [Cluster API](#)⁸⁰⁴ providers on the cluster:

- [Core Provider](#)⁸⁰⁵
- [vSphere Infrastructure Provider](#)⁸⁰⁶
- [Kubeadm Bootstrap Provider](#)⁸⁰⁷
- [Kubeadm ControlPlane Provider](#)⁸⁰⁸

7.11.9.3.3 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Create%20a%20New%20Air-gapped%20AWS%20Cluster&linkCreation=true&fromPageId=273352012>⁸⁰⁹

7.11.9.3.4 Create a New Air-gapped AWS Cluster

7.11.9.3.4.1 Prerequisites

Before you begin, make sure you have created a [Bootstrap](#) (see page 1093) cluster.

⁸⁰³ <https://github.com/kubernetes-sigs/kind>

⁸⁰⁴ <https://cluster-api.sigs.k8s.io/>


⁸⁰⁵ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

⁸⁰⁶ <https://github.com/kubernetes-sigs/cluster-api-provider-vsphere>

⁸⁰⁷ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

⁸⁰⁸ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>

⁸⁰⁹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273352012&linkCreation=true&spaceKey=DENT&title=%282.6%29+Create+a+New+Air-gapped+AWS+Cluster>


 DKP uses `localvolumeplugin` as the [default storage provider](#) (see page 91). However, `localvolumeplugin` is not suitable for production use. You should use a [Kubernetes CSI](#)⁸¹⁰ compatible storage that is suitable for production.

You can choose from any of the [storage options](#)⁸¹¹ available for Kubernetes. To disable the default that Konvoy deploys, set the default StorageClass `localvolumeplugin` as non-default. Then set your newly created StorageClass to be the default by following the commands in the Kubernetes documentation called [Changing the Default Storage Class](#)⁸¹².

7.11.9.3.4.2 Create a New Cluster


Create a new AWS Kubernetes cluster in an Air-gapped environment in your AWS infrastructure. When you use existing infrastructure, DKP does *not* create, modify, or delete the following AWS resources:

- Internet Gateways
- NAT Gateways
- Routing tables
- Subnets
- VPC
- VPC Endpoints (for subnets without NAT Gateways)

 An AWS subnet has Network ACLs that can control traffic in and out of the subnet. DKP does not modify the Network ACLs of an existing subnet. DKP uses Security Groups to control traffic. If a Network ACL denies traffic that is allowed by DKP-managed Security Groups, the cluster may not work correctly.

1. Set the environment variable to the name you assigned this cluster:

```
export CLUSTER_NAME=<aws-example>
```

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)⁸¹³ for more naming information.

⁸¹⁰ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁸¹¹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁸¹² <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

⁸¹³ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

2. Export variables for the existing infrastructure details:

```
export AWS_VPC_ID=<vpc-...>
export AWS_SUBNET_IDS=<subnet-...,subnet-...,subnet-...>
export AWS_ADDITIONAL_SECURITY_GROUPS=<sg-...>
export AWS_AMI_ID=<ami-...>
```

- `AWS_VPC_ID` : the VPC ID where the cluster will be created. The VPC requires the following AWS VPC Endpoints to be already present:
 - `ec2` - `com.amazonaws.{region}.ec2`
 - `elasticloadbalancing` - `com.amazonaws.{region}.elasticloadbalancing`
 - `secretsmanager` - `com.amazonaws.{region}.secretsmanager`
 - `autoscaling` - `com.amazonaws.{region}.autoscaling`
 - `ecr` - `com.amazonaws.{region}.ecr.api` - (authentication)
 - `ecr` - `com.amazonaws.{region}.ecr.dkr` - (data transfer)

More details about [AWS service using an interface VPC endpoint](#)⁸¹⁴ and [AWS VPC endpoints list](#)⁸¹⁵.

- `AWS_SUBNET_IDS` : a comma-separated list of one or more private Subnet IDs with each one in a different Availability Zone. The cluster control-plane and worker nodes will automatically be spread across these Subnets.
- `AWS_ADDITIONAL_SECURITY_GROUPS` : a comma-separated list of one or more Security Groups IDs to use in addition to the ones automatically created by [CAPA](#)⁸¹⁶.
- `AWS_AMI_ID` : the AMI ID to use for control-plane and worker nodes. The AMI must be created by the [konvoy-image-builder](#) (see page 1461).

! IMPORTANT: Ensure your [subnets](#) (see page 1049) do not overlap with your host subnet because they cannot be changed after cluster creation. If you need to change the kubernetes subnets, you must do this at cluster creation. The default subnets used in DKP are:

```
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
        - 192.168.0.0/16
    services:
```

⁸¹⁴ <https://docs.aws.amazon.com/vpc/latest/privatelink/create-interface-endpoint.html>

⁸¹⁵ <https://docs.aws.amazon.com/vpc/latest/privatelink/aws-services-privatelink-support.html>

⁸¹⁶ <https://github.com/kubernetes-sigs/cluster-api-provider-aws>

```
cidrBlocks:
- 10.96.0.0/12
```

⚠ IMPORTANT: You must tag the subnets as described below to allow for Kubernetes to create ELBs for services of type `LoadBalancer` in those subnets. If the subnets are not tagged, they will not receive an ELB and the following error displays: `Error syncing load balancer, failed to ensure load balancer; could not find any suitable subnets for creating the ELB.`

The tags should be set as follows, where `<CLUSTER_NAME>` corresponds to the name set in `CLUSTER_NAME` environment variable:

```
kubernetes.io/cluster = <CLUSTER_NAME>
kubernetes.io/cluster/<CLUSTER_NAME> = owned
kubernetes.io/role/internal-elb = 1
```

- a. **⚠ IMPORTANT:** The AMI must be created with [Konvoy Image Builder](#) (see page 1461) in order to use the registry mirror feature.
 - b. **⚠** If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see page 1533) page for more information.
3. (Optional) Configure your cluster to use an existing [local registry](#) (see page 1533) as a mirror when attempting to pull images previously [pushed to your registry](#) (see page 314). Below is an example command for **AWS ECR**:

```
export REGISTRY_MIRROR_URL=<your_registry_url>
```

- `REGISTRY_MIRROR_URL` : the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
 - NOTE: Other local registries may use the options below:
 - `JFrog - CONTAINER_REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
 - `CONTAINER_REGISTRY_USERNAME` : optional-set to a user that has pull access to this registry.
 - `CONTAINER_REGISTRY_PASSWORD` : optional if username is not set.
4. Create a Kubernetes cluster. The following example shows a common configuration. See [dkp create cluster aws](#) (see page 1665) reference for the full list of cluster creation options:

```
dkp create cluster aws --cluster-name=${CLUSTER_NAME} \
```

```

--vpc-id=${AWS_VPC_ID} \
--ami=${AWS_AMI_ID} \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=<YOUR_ECR_URL>

```

(Optional) Allow lookup of the AMI by providing *location*, *format* and *OS* with the following flags instead of AMI ID `--ami=${AWS_AMI_ID}` :

- `--ami-owner AWS_ACCOUNT_ID \`
- `--ami-base-os ubuntu-20.04 \`
- `--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*' \`

5. (Optional) Add the registry mirror flag to the command above to pull from your local registry:

```
--registry-mirror-url=<YOUR_ECR_URL>
```

6. (Optional) The Control Plane and Worker nodes can be configured to use an HTTP proxy:

```

export CONTROL_PLANE_HTTP_PROXY=http://example.org:8080
export CONTROL_PLANE_HTTPS_PROXY=http://example.org:8080
export
CONTROL_PLANE_NO_PROXY="example.org,example.com,example.net,localhost,127.0.0.1
,10.96.0.0/12,192.168.0.0/16,kubernetes,kubernetes.default,kubernetes.default.s
vc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.sv
c.cluster,.svc.cluster.local,169.254.169.254,.elb.amazonaws.com"

export WORKER_HTTP_PROXY=http://example.org:8080
export WORKER_HTTPS_PROXY=http://example.org:8080
export
WORKER_NO_PROXY="example.org,example.com,example.net,localhost,127.0.0.1,10.96.
0.0/12,192.168.0.0/16,kubernetes,kubernetes.default,kubernetes.default.svc,kube
rnetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.clust
er,.svc.cluster.local,169.254.169.254,.elb.amazonaws.com"

```

- Replace `example.org,example.com,example.net` with you internal addresses
- `localhost` and `127.0.0.1` addresses should not use the proxy
- `10.96.0.0/12` is the default Kubernetes service subnet
- `192.168.0.0/16` is the default Kubernetes pod subnet
- `kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local` is the internal Kubernetes kube-apiserver service
- `.svc,.svc.cluster,.svc.cluster.local` is the internal Kubernetes services
- `169.254.169.254` is the AWS metadata server

- `.elb.amazonaws.com` is for the worker nodes to allow them to communicate directly to the kube-apiserver ELB

7. (Optional) Create a Kubernetes cluster with HTTP proxy configured. This step assumes you did not already create a cluster in the previous steps:

```
dkp create cluster aws --cluster-name=${CLUSTER_NAME} \
--vpc-id=${AWS_VPC_ID} \
--ami=${AWS_AMI_ID} \
--subnet-ids=${AWS_SUBNET_IDS} \
--internal-load-balancer=true \
--additional-security-group-ids=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=<YOUR_ECR_URL> \
--control-plane-http-proxy="${CONTROL_PLANE_HTTP_PROXY}" \
--control-plane-https-proxy="${CONTROL_PLANE_HTTPS_PROXY}" \
--control-plane-no-proxy="${CONTROL_PLANE_NO_PROXY}" \
--worker-http-proxy="${WORKER_HTTP_PROXY}" \
--worker-https-proxy="${WORKER_HTTPS_PROXY}" \
--worker-no-proxy="${WORKER_NO_PROXY}"
```

8. Inspect the created cluster resources:

```
kubectl get clusters,kubeadmcontrolplanes,machinedeployments
```

9. Wait for the cluster control-plane to be ready:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=60m
```

Then, [Explore New AWS Cluster \(see page 1097\)](#)

7.11.9.4 Explore New AWS Cluster

This guide explains how to use the command line to interact with your newly deployed Kubernetes cluster.

Follow these steps:

1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

2. List the Nodes with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

Note: wait for the Status to move to `Ready` while `calico-node` pods are being deployed.

3. List the Pods with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get pods -A
```

7.11.9.5 Make the AWS Cluster Self-Managed

Konvoy deploys all cluster lifecycle services to a bootstrap cluster, which then deploys a workload cluster. When the workload cluster is ready, move the cluster lifecycle services to the workload cluster, which makes the workload cluster self-managed. This section describes how to make a workload cluster self-managed.

This page contains instructions on how to make your cluster self-managed. This is necessary if there is only one cluster in your environment, or if this cluster should become the Management cluster in a multi-cluster environment.



If you already have a self-managed or Management cluster in your environment, skip this page.

7.11.9.5.1 Make the New Kubernetes Cluster Manage Itself

1. Deploy cluster lifecycle services on the workload cluster:

```
dkp create capi-components --kubeconfig ${CLUSTER_NAME}.conf
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1042\)](#).

```
✓ Initializing new CAPI components
```

2. Move the Cluster API objects from the bootstrap to the workload cluster:

The cluster lifecycle services on the workload cluster are ready, but the workload cluster configuration is on the bootstrap cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the bootstrap to the workload cluster. This process is also called a [Pivot](#)⁸¹⁷.

⁸¹⁷ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```
dkp move capi-resources --to-kubeconfig ${CLUSTER_NAME}.conf
```

✓ Moving cluster resources

You can now view resources in the moved cluster by using the `--kubeconfig` flag with `kubectl`. For example: `kubectl --kubeconfig=aws-example.conf get nodes`

NOTE: To ensure only one set of cluster lifecycle services manages the workload cluster, Konvoy first pauses reconciliation of the objects on the bootstrap cluster, then creates the objects on the workload cluster. As Konvoy copies the objects, the cluster lifecycle services on the workload cluster reconcile the objects. The workload cluster becomes self-managed after Konvoy creates all the objects. If it fails, the `move` command can be safely retried.

3. Wait for the cluster control-plane to be ready:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --timeout=20m
```

```
cluster.cluster.x-k8s.io/aws-example condition met
```

4. Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

NOTE: After moving the cluster lifecycle services to the workload cluster, remember to use Konvoy with the workload cluster kubeconfig.

```
dkp describe cluster --kubeconfig ${CLUSTER_NAME}.conf -c ${CLUSTER_NAME}
```

NAME	SEVERITY	REASON	SINCE	MESSAGE	READY
Cluster/aws-example			109s		True
└ClusterInfrastructure - AWSCluster/aws-example			112s		True
└ControlPlane - KubeadmControlPlane/aws-example-control-plane			109s		True

```

| | Machine/aws-example-control-plane-55jh4      True
111s
| | Machine/aws-example-control-plane-6sn97      True
111s
| | Machine/aws-example-control-plane-nx9v5      True
110s
| | Workers
| | MachineDeployment/aws-example-md-0           True
114s
| | | Machine/aws-example-md-0-cb9c9bbf7-hcl8z   True
111s
| | | Machine/aws-example-md-0-cb9c9bbf7-rtdqw   True
111s
| | | Machine/aws-example-md-0-cb9c9bbf7-t894m   True
111s
| | | Machine/aws-example-md-0-cb9c9bbf7-td29r   True
111s

```

5. Remove the bootstrap cluster, as the workload cluster is now self-managed:

```
dkp delete bootstrap
```

```
✓ Deleting bootstrap cluster
```

7.11.9.5.2 Install and Log in to the Kommander UI

You can now proceed to installing the [UI with Kommander](#) (see page 1422) and applications. After installation, you will be able to [log in](#) (see page 1420) to the UI to explore it.

7.11.9.5.3 DKP Install Configuration

You can [configure](#) (see page 1422) the Kommander component of DKP during the initial installation, and also post-installation using the Kommander CLI.

The following explains how to run DKP on top of an [air-gapped DKP cluster](#) (see page 1095) installation with catalog applications.



Depending on your configuration, there are different ways you can install DKP to an air-gapped environment.

Ensure you follow the correct procedure for your configuration and license type, and ignore the other sections that do not pertain to your environment:

[Install Kommander in an Air-gapped Environment \(see page 1397\)](#)

7.11.9.5.4 Known Limitations

- Before making a workload cluster self-managed, be sure that its control plane nodes have sufficient permissions for running Cluster API controllers. See [IAM Policy Configuration \(see page 1066\)](#).
- Konvoy supports moving only one set of cluster objects from the bootstrap cluster to the workload cluster, or vice-versa.
- Konvoy only supports moving all namespaces in the cluster; Konvoy does not support migration of individual namespaces.

7.11.9.6 Creating DKP Clusters on AWS from the DKP UI

A guide for creating DKP clusters on AWS console

7.11.9.6.1 Before you begin

Configure an [AWS Infrastructure \(see page 1051\)](#), or add credentials using [AWS role credentials \(see page 607\)](#).

7.11.9.6.2 Simplified Cluster Creation on AWS in the UI

1. In the selected workspace Dashboard, select the **Add Cluster** option in the **Actions** dropdown menu at the top right.
2. On the Add Cluster page, select the **Create DKP Cluster**.
3. Provide some basic cluster details within the form:
 - **Workspace:** The workspace where this cluster belongs (if within the Global workspace).
 - **Kubernetes Version:** The initial version of Kubernetes to install on the cluster.
 - **Name:** A valid Kubernetes name for the cluster.
 - **Add Labels:** By default, your cluster has labels that reflect the infrastructure provider provisioning. For example, your AWS cluster may have a label for the data center region and `provider: aws`. Cluster labels are matched to the selectors created for [Projects \(see page 689\)](#). Changing a cluster label may add or remove the cluster from projects.

4. Select the pre-configured [AWS Infrastructure](#) (see page 1051), or [AWS role credentials](#) (see page 607) to display the remaining options specific to AWS.
 - **Region:** Select a data center region or specify a custom region.
 - **Configure Node Pools:** Specify pools of nodes, their machine types, quantity, and the IAM instance profile.
 - **Machine Type:** Machine instance type.
 - **Quantity:** Number of nodes. The control plane must be an odd number.
 - **IAM instance profile:** Name of the IAM instance profile to assign to the machines.
 - **AMI Image:** Specify the AMI to use for each node pool using either method (a) OR method (b) below:
 - (a) AMI ID specifying by name.
 - **AMI ID:** AMI ID to use for all nodes.
 - (b) AMI Lookup - Owner ID, Base OS, Lookup Format (all fields are required, if any is used)
 - **Base OS:** Base OS for Lookup search.
 - **Lookup Format:** Lookup Format string to generate AMI search name from.
 - **Owner ID:** Owner ID for AMI Lookup search.
 - **Worker Availability Zone:** Available under “Show Advanced” - Worker Zones for a region (worker nodes only).
 - **Add Infrastructure Provider Tags:** Specify tags applied on all resources created in your infrastructure for this cluster. Different infrastructure providers have varying restrictions on the usable tags. See the [AWS Tags User Guide](#)⁸¹⁸ for more information on using tags in AWS.
5. Select **Create** to begin provisioning the cluster. This step may take a few minutes, taking time for the cluster to be ready and fully deploy its components. The cluster automatically tries to join, and should resolve after it is fully-provisioned.

7.11.10 Delete an AWS Cluster



A self-managed workload cluster cannot delete itself. If your workload cluster is self-managed, you must first create a bootstrap cluster and move the cluster lifecycle services to it, before deleting the workload cluster.

If you did not make your workload cluster self-managed, as described in [Make New Cluster Self-Managed](#) (see page 1101), see the instructions in [Delete the workload cluster](#) (see page 1105).

⁸¹⁸ https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html

7.11.10.1 Create a Bootstrap Cluster and Move CAPI Resources

Follow these steps to create a bootstrap cluster and move CAPI resources:

1. Make sure your AWS credentials are up to date. Refresh the credentials using this command:

```
dkp update bootstrap credentials aws --kubeconfig $HOME/.kube/config
```

2. The bootstrap cluster will host the Cluster API controllers that reconcile the cluster objects marked for deletion. Create a bootstrap cluster:

i NOTE: To avoid using the wrong kubeconfig, the following steps use explicit kubeconfig paths and contexts.

```
dkp create bootstrap --kubeconfig $HOME/.kube/config --with-aws-bootstrap-credentials=true
```

- ✓ Creating a bootstrap cluster
- ✓ Initializing **new** CAPI components

3. Move the Cluster API objects from the workload to the bootstrap cluster: The cluster lifecycle services on the bootstrap cluster are ready, but the workload cluster configuration is on the workload cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the workload to the bootstrap cluster. This process is also called a **Pivot**⁸¹⁹.

```
dkp move capi-resources \
  --from-kubeconfig ${CLUSTER_NAME}.conf \
  --from-context ${CLUSTER_NAME}-admin@${CLUSTER_NAME} \
  --to-kubeconfig $HOME/.kube/config \
  --to-context kind-konvoy-capi-bootstrapper
```

- ✓ Moving cluster resources
- You can now view resources in the moved cluster by using the `--kubeconfig` flag with `kubectl`. For example: `kubectl --kubeconfig $HOME/.kube/config get nodes`

4. Use the cluster lifecycle services on the workload cluster to check the workload cluster's status:

⁸¹⁹ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```
dkp describe cluster --kubeconfig $HOME/.kube/config -c ${CLUSTER_NAME}
```

NAME	REASON	SINCE	MESSAGE	READY	SEVERITY
Cluster/aws-example		91s		True	
└─ClusterInfrastructure - AWSCluster/aws-example		103s		True	
└─ControlPlane - KubeadmControlPlane/aws-example-control-plane		91s		True	
└─Machine/aws-example-control-plane-55jh4		102s		True	
└─Machine/aws-example-control-plane-6sn97		102s		True	
└─Machine/aws-example-control-plane-nx9v5		102s		True	
└─Workers					
└─MachineDeployment/aws-example-md-0		108s		True	
└─Machine/aws-example-md-0-cb9c9bbf7-hcl8z		102s		True	
└─Machine/aws-example-md-0-cb9c9bbf7-rtdqw		102s		True	
└─Machine/aws-example-md-0-cb9c9bbf7-td29r		102s		True	
└─Machine/aws-example-md-0-cb9c9bbf7-w64kg		102s		True	

i NOTE: After moving the cluster lifecycle services to the workload cluster, remember to use DKP with the workload cluster kubeconfig and use DKP with the bootstrap cluster to delete the workload cluster.

5. Wait for the cluster control-plane to be ready:

```
kubectl --kubeconfig $HOME/.kube/config wait --for=condition=controlplaneready "clusters/${CLUSTER_NAME}" --timeout=20m
```

```
cluster.cluster.x-k8s.io/aws-example condition met
```

7.11.10.2 Delete the Workload Cluster

1. Make sure your AWS credentials are up to date. Refresh the credentials using this command:

```
dkp update bootstrap credentials aws --kubeconfig $HOME/.kube/config
```

i NOTE: Persistent Volumes (PVs) are not deleted automatically by design to preserve your data. However, the PVs take up storage space if not deleted. You must delete PVs manually. Information for backup of a cluster and PVs is on the page, [Back up your Cluster's Applications and Persistent Volumes](#) (see page 835).

2. To delete a cluster, you would use `dkp delete cluster` and pass in the name of the cluster you are trying to delete with `--cluster-name` flag. You would use `kubectl get clusters` to get those details (`--cluster-name` and `--namespace`) of the Kubernetes cluster to delete it.

NOTE: Do not use `dkp get clusters` since that gets you Kommander cluster details rather than Konvoy kubernetes cluster details.

```
kubectl get clusters
```

3. Delete the Kubernetes cluster and wait a few minutes:

NOTE: Before deleting the cluster, dkp deletes all Services of type LoadBalancer on the cluster. Each Service is backed by an AWS Classic ELB. Deleting the Service deletes the ELB that backs it. To skip this step, use the flag `--delete-kubernetes-resources=false`. Do not skip this step if the VPC is managed by DKP. When DKP deletes the cluster, it deletes the VPC. If the VPC has any AWS Classic ELBs, AWS does not allow the VPC to be deleted, and DKP cannot delete the cluster.

```
dkp delete cluster --cluster-name=${CLUSTER_NAME} --kubeconfig $HOME/.kube/config
```

```
✓ Deleting Services with type LoadBalancer for Cluster default/azure-example
✓ Deleting ClusterResourceSets for Cluster default/azure-example
✓ Deleting cluster resources
✓ Waiting for cluster to be fully deleted
Deleted default/azure-example cluster
```

After the workload cluster is deleted, you can delete the bootstrap cluster.

7.11.10.3 Delete the Bootstrap Cluster

After you have moved the workload resources back to a bootstrap cluster and deleted the workload cluster, you no longer need the bootstrap cluster. You can safely delete the bootstrap cluster with these steps:

1. Make sure your AWS credentials are up to date. Refresh the credentials using this command:

```
dkp update bootstrap credentials aws --kubeconfig $HOME/.kube/config
```

2. Delete the bootstrap cluster:

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

```
✓ Deleting bootstrap cluster
```

7.11.10.3.1 Air-gapped environments:

1. Make sure your AWS credentials are up to date. Refresh the credentials using this command:

```
dkp update bootstrap credentials aws --kubeconfig $HOME/.kube/config
```

2. Delete the provisioned Kubernetes cluster and wait a few minutes:

```
dkp delete cluster --cluster-name=${CLUSTER_NAME}
```

3. Delete the `kind` Kubernetes cluster:

```
dkp delete bootstrap
```

7.11.10.4 Known Limitations

- The Konvoy version used to create the workload cluster must match the Konvoy version used to delete the workload cluster.

7.11.10.5 Next Topic:

[Multiple AWS Accounts](#) (see page 1107)

7.11.11 Multiple AWS Accounts

Leverage Multiple AWS Accounts for Kubernetes Cluster Deployments

7.11.11.1 Objective

You can leverage multiple AWS accounts in your organization to meet specific business purposes, reflect your organizational structure, or implement a multi-tenancy strategy. Specific scenarios include:

- Implementing isolation between environment tiers such as development, testing, acceptance, and production.
- Implementing separation of concerns between management clusters, and workload clusters.
- Reducing the impact of security events and incidents.

For additional benefits of using multiple AWS accounts, refer to the following [white paper](#)⁸²⁰.

This document describes how to leverage the D2iQ Kubernetes Platform (DKP) to deploy a management cluster, and multiple workload clusters, leveraging multiple AWS accounts.

7.11.11.2 Assumptions

This guide assumes you have some understanding of Cluster API concepts and basic DKP provisioning workflows on AWS.

Cluster API Concepts - [cluster API concepts](#)⁸²¹

[AWS Install Options](#) (see page 294)

7.11.11.3 Glossary

- **Management cluster** - The cluster that runs in AWS and is used to create target clusters in different AWS accounts.
- **Target account** - The account where the target cluster is created.
- **Source account** - The AWS account where the CAPA controllers for the management cluster runs.

7.11.11.4 Prerequisites

Before you begin deploying DKP on AWS, you configure the prerequisites for the environment you use either non-air-gapped or air-gapped.

⁸²⁰ <https://docs.aws.amazon.com/whitepapers/latest/organizing-your-aws-environment/benefits-of-using-multiple-aws-accounts.html>

⁸²¹ <https://cluster-api.sigs.k8s.io/user/concepts.html>

[Prerequisites for Install](#) (see page 121)

7.11.11.5 Deploy DKP on AWS

1. Deploy a management cluster in your AWS source account.
AWS: [create Kubernetes AWS cluster](#) (see page 1079)
2. Configure a trusted relationship between source and target accounts and create a management cluster:

7.11.11.5.1 Step 1:

DKP leverages the Cluster API provider for AWS (CAPA) to provision Kubernetes clusters in a declarative way. Customers declare the desired state of the cluster through a cluster configuration YAML file which is generated using:

(AWS)

```
dkp create cluster aws --cluster-name=${CLUSTER_NAME} \
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml
```

7.11.11.5.2 Step 2:

Configure a trust relationship between the source and target accounts.

Follow all the prerequisite steps in both the source and target accounts

1. Create all policies and roles in management and workload accounts a. The prerequisite IAM policies for DKP are documented here: [Configure AWS IAM policies](#) (see page 1066).
2. Establish a trust relationship in workload account for the management account.
 - a. Go to your target (workload) account
 - b. Search for the role control-plane.cluster-api-provider-aws.sigs.k8s.io
 - c. Navigate to the Trust Relationship tab and select Edit Trust Relationship
 - d. Add the following relationship:

```
{
  "Effect": "Allow",
  "Principal": {
    "AWS": "arn:aws:iam::${mgmt-aws-account}:role/control-plane.cluster-api-provider-aws.sigs.k8s.io"
  },
  "Action": "sts:AssumeRole"
}
```

3. Give permission to role in the source (management cluster) account to call the `sts:AssumeRole` API
 - a. Log in to the source AWS account and attach the following inline policy to `control-plane.cluster-api-provider-aws.sigs.k8s.io` role:

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Action": "sts:AssumeRole",
      "Resource": [
        "arn:aws:iam::${workload-aws-account}:role/control-plane.cluster-api-provider-aws.sigs.k8s.io"
      ]
    }
  ]
}
```

4. Modify the management cluster configuration file and update the `AWSCluster` object with following details:

```
apiVersion: infrastructure.cluster.x-k8s.io/v1alpha3
kind: AWSCluster
metadata:
spec:
  identityRef:
    kind: AWSClusterRoleIdentity
    name: cross-account-role
...
---
apiVersion: infrastructure.cluster.x-k8s.io/v1alpha3
kind: AWSClusterRoleIdentity
metadata:
  name: cross-account-role
spec:
  allowedNamespaces: {}
  roleARN: "arn:aws:iam::${workload-aws-account}:role/control-plane.cluster-api-provider-aws.sigs.k8s.io"
  sourceIdentityRef:
    kind: AWSClusterControllerIdentity
    name: default
```

After performing the above steps, your Management cluster will be configured to create new managed clusters in the target AWS workload account.

7.11.11.6 Next Steps:

[AWS Install Options](#) (see page 294)

7.11.12 GPUs in an AWS environment

7.11.12.1 Understanding GPUs

Before working with GPUs, ensure you are familiar with the following:

- Install GPU support on supported distributions on AWS
- GPU node labeling specifications described in [Configure Konvoy Automatic GPU Node Labels](#) (see page 1111)

Kommander also accesses [Kommander GPU configuration](#) (see page 978) resources.



For using GPUs in an air-gapped on-premises environment, D2iQ recommends setting up [Pod Disruption Budget](#)⁸²² before [Update Cluster Nodepools](#) (see page 1545).

7.11.12.2 Install GPU Support for Supported Distributions on AWS

Using the [Konvoy Image Builder](#) (see page 1461), you can build an image that has support to use NVIDIA GPU hardware to support GPU workloads. DKP supported [NVIDIA driver](#)⁸²³ version is 470.x.

1. In your `overrides/nvidia.yaml` file add the following to enable GPU builds. You can also access and use the overrides [repo](#).⁸²⁴

```
gpu:
  type:
  - nvidia
```

Build your image using the following Konvoy image builder commands:

⁸²² <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>


⁸²³ <https://www.nvidia.com/Download/Find.aspx>

⁸²⁴ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

```
konvoy-image build --region us-west-2 --source-ami=ami-12345abcdef images/ami/centos-7.yaml --overrides overrides/nvidia.yaml
```

By default, your image builds in the `us-west-2` region. To specify another region, set the `--region` flag:


```
konvoy-image build --region us-east-1 --overrides override-source-ami.yaml images/ami/<Your OS>.yaml
```

 Ensure that an AMI file is available for your OS selection.

1. When the command is complete the `ami` id is printed and written to `./manifest.json`.

To use the built `ami` with Konvoy Image Builder, specify it with the `--ami` flag when calling `cluster create`.

```
dkp create cluster aws --cluster-name=$(whoami)-aws-cluster --region us-west-2 --ami <ami>
```

 If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

Additional helpful information can be found in the [NVIDIA Device Plug-in](#)⁸²⁵ for Kubernetes instructions and the [Installation Guide of Supported Platforms](#)⁸²⁶.

See also: [NVIDIA documentation](#)⁸²⁷

7.11.12.3 Configure Konvoy Automatic GPU Node Labels

When using GPU nodes, it is important they have the proper label identifying them as Nvidia GPU nodes. Node feature discovery (NFD), by default labels PCI hardware as:

⁸²⁵ <https://github.com/NVIDIA/k8s-device-plugin/blob/master/README.md>

⁸²⁶ <https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/install-guide.html>

⁸²⁷ https://nvidia.custhelp.com/app/answers/detail/a_id/131/kw/driver%20installation%20docs/related/1

```
"feature.node.kubernetes.io/pci-<device label>.present": "true"
```

where `<device label>` is by default as [defined in this topic](#)⁸²⁸.

```
< class > _ < vendor >
```

However, because there is a wide variety in devices and their assigned PCI classes, you may find that the labels assigned to your GPU nodes do not always properly identify them as containing an Nvidia GPU.

If the default detection does not work, you can manually change the daemonset that the GPU operator creates by running the following command:

```
nodeSelector:
  feature.node.kubernetes.io/pci-< class > _ < vendor>.present: "true"
```

where `class` is any 4 digit number starting with `03xy` and the vendor for Nvidia is `10de`. If this is already deployed, you can always change the `daemonset` and change the `nodeSelector` field so that it deploys to the right nodes.

7.11.12.4 Update NVIDIA GPU Clusters

Upgrading a node pool involves draining the existing nodes in the node pool and replacing them with new nodes. In order to ensure minimum downtime and maintain high availability of the critical application workloads during the upgrade process, we recommend deploying Pod Disruption Budget ([Disruptions](#)⁸²⁹) for your critical applications.

The Pod Disruption Budget will prevent any impact on critical applications as a result of misconfiguration or failures during the upgrade process.

7.11.12.4.1 Prerequisites:

- Deploy [Pod Disruption Budget](#)⁸³⁰ (PDB)
- [Konvoy Image Builder](#) (see page 1461) (KIB)


7.11.12.4.2 Steps:

1. Deploy Pod Disruption Budget for your critical applications. If your application can tolerate only one replica to be unavailable at a time, then you can set Pod disruption budget as shown in the following example. The example below is for NVIDIA GPU node pools, but the process is the same for all node pools.

⁸²⁸ <https://kubernetes-sigs.github.io/node-feature-discovery/v0.7/get-started/features.html#pci>

⁸²⁹ <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

⁸³⁰ <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

 Repeat this step for each additional node pool.

Create the file: `pod-disruption-budget-nvidia.yaml`

```
apiVersion: policy/v1
kind: PodDisruptionBudget
metadata:
  name: nvidia-critical-app
spec:
  maxUnavailable: 1
  selector:
    matchLabels:
      app: nvidia-critical-app
```

Apply the YAML file above using the following command:

```
kubectl create -f pod-disruption-budget-nvidia.yaml
```

2. Prepare OS image for your node pool using [Konvoy Image Builder](#) (see page 1461).

7.11.13 Manage AWS Node Pools

Node pools are part of a cluster and managed as a group, and you can use a node pool to manage a group of machines using the same common properties. When Konvoy creates a new default cluster, there is one node pool for the worker nodes and all nodes in that new node pool have the same configuration. You can create additional node pools for more specialized hardware or configuration. For example, if you want to tune your memory usage on a cluster where you need maximum memory for some machines and minimal memory on other machines, you would create a new node pool with those specific resource needs.

DKP implements node pools using Cluster API [MachineDeployments](#)⁸³¹. For more information on node pools, see these sections:

- [Create AWS Node Pools](#) (see page 1114)
- [List AWS Node Pools](#) (see page 1115)
- [Scale AWS Node Pools](#) (see page 1116)
- [Delete AWS Node Pools](#) (see page 1118)
- [AWS Cluster Autoscaler](#) (see page 1119)

⁸³¹ <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

7.11.13.1 Create AWS Node Pools

Creating a node pool is useful when you need to run workloads that require machines with specific resources, such as a GPU, additional memory, or specialized network or storage hardware.

7.11.13.1.1 Prepare the Environment

1. Set the environment variable to the name you assigned this cluster.

```
export CLUSTER_NAME=aws-example
```

See [AWS Install Options \(see page 294\)](#) for information on naming your cluster.

2. If your workload cluster is self-managed, as described in [Make the New Cluster Self-Managed \(see page 1101\)](#), configure `kubectl` to use the kubeconfig for the cluster.

```
export KUBECONFIG=${CLUSTER_NAME}.conf
```

3. Define your node pool name.

```
export NODEPOOL_NAME=example
```

7.11.13.1.2 Create an AWS Node Pool

Availability zones (AZs) are isolated locations within data center regions from which public cloud services originate and operate. Because all the nodes in a node pool are deployed in a single Availability Zone, you may wish to create additional node pools to ensure your cluster has nodes deployed in multiple Availability Zones.



By default, the first Availability Zone in the region will be used for the Nodes in the node pool, set `--availability-zone` to create the Nodes in a different Availability Zone.

Create a new AWS node pool with 3 replicas using this command:

```
dkp create nodepool aws ${NODEPOOL_NAME} \
  --cluster-name=${CLUSTER_NAME} \
  --replicas=3
```

This example uses default values for brevity. Use flags to define custom instance types, AMIs, and other properties.

Advanced users can use a combination of the `--dry-run` and `--output=yaml` or `--output-directory=<existing-directory>` flags to get a complete set of node pool objects to modify locally or store in version control.



For more information regarding this flag or others, please refer to the [DKP CLI Create](#) (see page 1684) section of the documentation for either cluster or nodepool and select your provider.

7.11.13.2 List AWS Node Pools

List node pools for a cluster

Use this command to list the node pools of a given cluster. This returns specific properties of each node pool so that you can see the name of the MachineDeployments.

7.11.13.2.1 Managed Clusters

To list all node pools for a managed cluster, run:

```
dkp get nodepools --cluster-name=${CLUSTER_NAME} --kubeconfig=${CLUSTER_NAME}.conf
```

The expected output is similar to the following example, indicating the desired size of the node pool, the number of replicas ready in the node pool, and the Kubernetes version those nodes are running:

NODEPOOL	DESIRED	READY	KUBERNETES VERSION
example	3	3	v1.26.6
aws-example-md-0	4	4	v1.26.6

7.11.13.2.2 Attached Clusters

Before running the command to list the attached clusters, ensure that you know the following:

- [KUBECONFIG](#) (see page 87) for the [management cluster](#) (see page 81) - In order to find the KUBECONFIG for a cluster from the UI, refer to this section in the documentation: [Access a Managed or Attached Cluster](#) (see page 829)
- The `CLUSTER_NAME` of the attached cluster
- The `NAMESPACE` of the attached cluster

To list all node pools for an attached cluster, run:

```
dkp get nodepools --cluster-name=${ATTACHED_CLUSTER_NAME} --kubeconfig=${CLUSTER_NAME}.conf -n ${ATTACHED_CLUSTER_NAMESPACE}
```

The expected output is similar to below:

NODEPOOL VERSION	DESIRED	READY	KUBERNETES
aws-attached-md-0	4	4	v1.26.6

7.11.13.3 Scale AWS Node Pools

Scale node pools in a cluster

While you can run [Cluster Autoscaler](#) (see page 1119), you can also manually scale your node pools up or down when you need more finite control over your environment. For example, if you require 10 machines to run a process, you can manually set the scaling to run those 10 machines only. However, if also using the Cluster Autoscaler, you must stay within your minimum and maximum bounds.

7.11.13.3.1 Prerequisite

Environment variables, such as define the node pool name, are set in the [Prepare the Environment](#) (see page 1117) section on the previous page. If need, refer to that page to set those variables.

7.11.13.3.2 Scaling Up Node Pools

To scale up a node pool in a cluster, run:

```
dkp scale nodepools ${NODEPOOL_NAME} --replicas=5 --cluster-name=${CLUSTER_NAME}
```

OR for attached clusters:

```
dkp scale nodepools ${ATTACHED_NODEPOOL_NAME} --replicas=5 --cluster-name=${ATTACHED_CLUSTER_NAME} --kubeconfig=${CLUSTER_NAME}.conf -n ${ATTACHED_CLUSTER_WORKSPACE}
```

Your output should be similar to this example, indicating the scaling is in progress:

```
✓ Scaling node pool example to 5 replicas
```

After a few minutes, you can list the node pools to:

```
dkp get nodepools --cluster-name=${CLUSTER_NAME} --kubeconfig=${CLUSTER_NAME}.conf
```

Your output should be similar to this example, with the number of DESIRED and READY replicas increased to 5:

NODEPOOL	DESIRED	READY	
KUBERNETES VERSION			
aws-example-md-0	5	5	v1.26.6
aws-attached-md-0	5	5	v1.26.6

7.11.13.3 Scaling Down Node Pools

To scale down a node pool, run:

```
dkp scale nodepools ${NODEPOOL_NAME} --replicas=4 --cluster-name=${CLUSTER_NAME}
```

OR for attached clusters:

```
dkp scale nodepools ${ATTACHED_NODEPOOL_NAME} --replicas=4 --cluster-name=${ATTACHED_CLUSTER_NAME} --kubeconfig=${CLUSTER_NAME}.conf -n ${ATTACHED_CLUSTER_WORKSPACE}
```

✓ Scaling node pool example to 4 replicas

After a few minutes, you can list the node pools using this command:

```
dkp get nodepools --cluster-name=${CLUSTER_NAME} --kubeconfig=${CLUSTER_NAME}.conf
```

Your output should be similar to this example, with the number of DESIRED and READY replicas decreased to 4:

NODEPOOL	DESIRED	READY	
KUBERNETES VERSION			
example	4	4	v1.26.6
aws-example-md-0	4	4	v1.26.6

In a default cluster, the nodes to delete are selected at random. This behavior is controlled by [CAPI's delete policy](#)⁸³². However, when using the DKP CLI to scale down a node pool, it is also possible to specify the Kubernetes Nodes you want to delete.

⁸³² https://github.com/kubernetes-sigs/cluster-api/blob/v0.4.0/api/v1alpha4/machineset_types.go#L85-L105

To do this, set the flag `--nodes-to-delete` with a list of nodes as below. This adds an annotation `cluster.x-k8s.io/delete-machine=yes` to the matching Machine object that contains `status.NodeRef` with the node names from `--nodes-to-delete`.

```
dkp scale nodepools ${NODEPOOL_NAME} --replicas=3 --nodes-to-delete=<> --cluster-name=${CLUSTER_NAME}
```

✓ Scaling node pool example to 3 replicas

7.11.13.3.4 Scaling Node Pools When Using Cluster Autoscaler

If you [configured the cluster autoscaler \(see page 1119\)](#) for the `demo-cluster-md-0` node pool, the value of `--replicas` must be within the minimum and maximum bounds.

For example, assuming you have these annotations:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment ${NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-autoscaler-node-group-min-size=2
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment ${NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size=6
```

Try to scale the node pool to 7 replicas with the command:

```
dkp scale nodepools ${NODEPOOL_NAME} --replicas=7 -c demo-cluster
```

Which results in an error similar to:

```
✗ Scaling node pool example to 7 replicas
failed to scale nodepool: scaling MachineDeployment is forbidden: desired replicas 7 is greater than the configured max size annotation cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size: 6
```

Similarly, scaling down to a number of replicas less than the configured `min-size` also returns an error.

7.11.13.4 Delete AWS Node Pools

Deleting a node pool deletes the Kubernetes nodes and the underlying infrastructure. All nodes will be drained prior to deletion and the pods running on those nodes will be rescheduled.

To delete a node pool from a managed cluster, run:

```
dkp delete nodepool ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME}
```

Here, `example` is the node pool to be deleted.

The expected output will be similar to the following example, indicating the node pool is being deleted:

```
✓ Deleting default/example nodepool resources
```

Deleting an invalid node pool results in output similar to this example:

```
dkp delete nodepool ${CLUSTER_NAME}-md-invalid --cluster-name=${CLUSTER_NAME}
```

```
MachineDeployments or MachinePools.infrastructure.cluster.x-k8s.io "no
MachineDeployments or MachinePools found for cluster aws-example" not found
```

7.11.13.5 AWS Cluster Autoscaler

This page explains how to configure autoscaler for node pools. [Cluster Autoscaler](#)⁸³³ provides the ability to automatically scale-up or scale-down the number of worker nodes in a cluster, based on the number of pending pods to be scheduled. Running the Cluster Autoscaler is optional.

Unlike [Horizontal-Pod Autoscaler](#)⁸³⁴, Cluster Autoscaler does not depend on any Metrics server and does not need Prometheus or any other metrics source.

The Cluster Autoscaler looks at the following annotations on a MachineDeployment to determine its scale-up and scale-down ranges:

```
cluster.x-k8s.io/cluster-api-autoscaler-node-group-min-size
cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size
```

The full list of command line arguments to the Cluster Autoscaler controller is [on the Kubernetes public GitHub repository](#)⁸³⁵.

For more information about how Cluster Autoscaler works, see these documents:

- [What is Cluster Autoscaler](#)⁸³⁶
- [How does scale-up work](#)⁸³⁷
- [How does scale-down work](#)⁸³⁸

833 <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

834 <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-fast-is-hpa-when-combined-with-ca>

835 <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-are-the-parameters-to-ca>

836 <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-is-cluster-autoscaler>

837 <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-up-work>

838 <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-down-work>

- [CAPI Provider for Cluster Autoscaler](#)⁸³⁹

7.11.13.5.1 Cluster Autoscaler Prerequisites

Before you begin, you must have:

- A [Bootstrap Cluster Lifecycle](#) (see page 1077).
- [Created a new Kubernetes Cluster](#) (see page 1079).
- A [Self-Managed Cluster](#) (see page 1101).

7.11.13.5.2 Run Cluster Autoscaler

The Cluster Autoscaler controller runs on the workload cluster. Upon creation of the workload cluster, this controller does not have all the objects required to function correctly until after a `dkp move` is issued from the bootstrap cluster.

Run the following steps to enable Cluster Autoscaler:

1. Ensure the Cluster Autoscaler controller is up and running (no restarts and no errors in the logs)

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf logs deployments/cluster-autoscaler
cluster-autoscaler -n kube-system -f
```

2. Enable Cluster Autoscaler by setting the min & max ranges

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment $
{NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-autoscaler-node-group-min-size=2
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment $
{NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size=6
```

3. The Cluster Autoscaler logs will show that the worker nodes are associated with node-groups and that pending pods are being watched.
4. To demonstrate that it is working properly, create a large deployment which will trigger pending pods (For this example we used AWS m5.2xlarge worker nodes. If you have larger worker-nodes, you should scale up the number of replicas accordingly).

```
cat <<EOF | kubectl --kubeconfig=${CLUSTER_NAME}.conf apply -f -
apiVersion: apps/v1
kind: Deployment
metadata:
  name: busybox-deployment
  labels:
    app: busybox
```

⁸³⁹ <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

```
spec:
  replicas: 600
  selector:
    matchLabels:
      app: busybox
  template:
    metadata:
      labels:
        app: busybox
    spec:
      containers:
      - name: busybox
        image: busybox:latest
        command:
          - sleep
          - "3600"
        imagePullPolicy: IfNotPresent
        restartPolicy: Always
EOF
```

5. Cluster Autoscaler will scale up the number of Worker Nodes until there are no pending pods.
6. Scale down the number of replicas for `busybox-deployment`.

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf scale --replicas=30 deployment/
busybox-deployment
```

7. Cluster Autoscaler starts to scale down the number of Worker Nodes after the default timeout of 10 minutes.



Enterprise Customers: In previous DKP versions, the Cluster-Autoscaler did not work on **managed** clusters. Starting with this release, an instance of Cluster-Autoscaler is deployed on the **management cluster** for each **managed** cluster. During an upgrade, the upgrade procedure on the **management cluster** will create the necessary deployments.

7.11.14 AWS Certificate Renewal

During cluster creation, Kubernetes establishes a Public Key Infrastructure (PKI) for generating the TLS certificates needed for securing cluster communication for various components such as `etcd`, `kubernetes-apiserver` and `kube-proxy`. The certificates created by these components have a default expiration of one year and are renewed when an administrator updates the cluster.

Kubernetes provides a facility to renew all certificates automatically during control plane updates. For administrators who need long-running clusters or clusters that are not upgraded often, `dkp` provides automated certificate renewal, without a cluster upgrade.

7.11.14.1 Prerequisites

- This feature requires Python 3.5 or greater to be installed on all control plane hosts.
- Complete the [Bootstrap Cluster](#) (see page 1077) topic.

7.11.14.2 Create a Cluster with Automated Certificate Renewal

To enable the automated certificate renewal, create a Konvoy cluster using the `certificate-renew-interval` flag:

```
dkp create cluster aws --certificate-renew-interval=60 --cluster-name=long-running
```

The `certificate-renew-interval` is the number of days after which Kubernetes-managed PKI certificates will be renewed. For example, an `certificate-renew-interval` value of 60 means the certificates will be renewed every 60 days.

7.11.14.3 Technical Details

The following manifests are modified on the control plane hosts, and are located at `/etc/kubernetes/manifests`. Modifications to these files requires SUDO access.

```
kube-controller-manager.yaml
kube-apiserver.yaml
kube-scheduler.yaml
kube-proxy.yaml
```

The following annotation indicates the time each component was reset:

```
metadata:
  annotations:
    konvoy.d2iq.io/restartedAt: $(date +%s)
```

This only occurs when the PKI certificates are older than the interval given at cluster creation time. This is activated by a `systemd timer` called `renew-certs.timer` that triggers an associated `systemd service` called `renew-certs.service` that runs on all of the control plane hosts.

7.11.14.4 Debug

To debug the automatic certificate renewal feature, a cluster administrator can look at several different components to see if the certificates were renewed. For example, an administrator might start with a look at the control plane pod definition to check the last reset time. To determine if a scheduler pod was properly reset, run the command:

```
kubectl get pod -n kube-system kube-scheduler-ip-10-0-xx-xx.us-west-2.compute.interna  
l -o yaml
```

The output of the command will be similar to the following:

```
apiVersion: v1  
kind: Pod  
metadata:  
  annotations:  
    konvoy.d2iq.io/restartedAt: "1626124940.735733"
```

Administrators who want more details on the execution of the `systemd` service can use `ssh` to connect to the control plane hosts, and then use the `systemctl` and `journalctl` commands that follow to help diagnose potential issues.

To check the status of the timers, when they last ran, and when they are scheduled to run next, use the command:

```
systemctl list-timers
```

To check the status of the `renew-certs` service, use the command:

```
systemctl status renew-certs
```

To get the logs of the last run of the service, use the command:

```
journalctl logs -u renew-certs
```

7.11.15 Configure Infrastructure in UI

This page refers to working inside the AWS Console in order to add Infrastructure Provider.

7.11.15.1 Create Infrastructure Provider

To configure an AWS Provider with a User Role in the AWS Console (UI), perform the following steps:

1. From the top menu bar, select your target workspace.
2. Select **Infrastructure Providers** in the **Administration** section of the sidebar menu.
3. Select the **Add Infrastructure Provider** button.
4. Select the **Amazon Web Services (AWS)** option.
5. Ensure **Role** is selected as the **Authentication Method**.
6. Enter a name for your infrastructure provider. Select a name that matches the AWS user.
7. Enter the **Role ARN**.
8. You can add an **External ID** if you share the Role with a 3rd party. External IDs secure your environment from accidentally used roles. [Read more about External IDs](#)⁸⁴⁰.
9. Select **Save** to save your provider.

7.11.15.2 Fill out the Add Infrastructure Provider Form

To fill out the Add Infrastructure Provider form in the AWS Console using Static Credentials, perform the following steps:

1. In Kommander, select the Workspace associated with the credentials you are adding.
2. Navigate to **Administration > Infrastructure Providers** and click the **Add Infrastructure Provider** button.
3. Select the Amazon Web Services (AWS) option.
4. Ensure **Static** is selected as the Authentication Method.
5. Enter a name for your infrastructure provider for later reference. Consider choosing a name that matches the AWS user.
6. Fill out the access and secret keys using the keys generated above.
7. Select **Save** to save your provider.

7.11.16 Replace an AWS Node

7.11.16.1 Prerequisites

Before you begin, you must:

⁸⁴⁰ https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user_externalid.html

- [Create a workload cluster](#) (see page 1079).
- [Make the new cluster self-managed](#) (see page 1101).

7.11.16.2 Replace a Worker Node

In certain situations, you may want to delete a worker node and have [Cluster API](#)⁸⁴¹ replace it with a newly provisioned machine.

1. Identify the name of the node to delete.

List the nodes:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf get nodes
```

The output from this command resembles the following:

NAME	STATUS	ROLES
AGE VERSION		
ip-10-0-100-85.us-west-2.compute.internal 16m v1.26.6	Ready	<none>
ip-10-0-106-183.us-west-2.compute.internal 15m v1.26.6	Ready	control-plane,master
ip-10-0-158-104.us-west-2.compute.internal 17m v1.26.6	Ready	control-plane,master
ip-10-0-203-138.us-west-2.compute.internal 16m v1.26.6	Ready	control-plane,master
ip-10-0-70-169.us-west-2.compute.internal 16m v1.26.6	Ready	<none>
ip-10-0-77-176.us-west-2.compute.internal 16m v1.26.6	Ready	<none>
ip-10-0-96-61.us-west-2.compute.internal 16m v1.26.6	Ready	<none>

2. Export a variable with the node name to use in the next steps:

This example uses the name `ip-10-0-100-85.us-west-2.compute.internal`.

```
export NAME_NODE_TO_DELETE="ip-10-0-100-85.us-west-2.compute.internal"
```

3. Delete the Machine resource

⁸⁴¹ <https://cluster-api.sigs.k8s.io/>


```
export NAME_MACHINE_TO_DELETE=$(kubectl --kubeconfig ${CLUSTER_NAME}.conf get
machine -ojsonpath="{.items[?
(@.status.nodeRef.name==\"$NAME_NODE_TO_DELETE\")].metadata.name}")
kubectl --kubeconfig ${CLUSTER_NAME}.conf delete machine
"$NAME_MACHINE_TO_DELETE"
```

```
machine.cluster.x-k8s.io "aws-example-1-md-0-cb9c9bbf7-t894m" deleted
```

The command will not return immediately. It will return after the Machine resource has been deleted. A few minutes after the Machine resource is deleted, the corresponding Node resource is also deleted.

- Observe that the Machine resource is being replaced using this command:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf get machinedeployment
```

NAME	PHASE	AGE	CLUSTER	VERSION	REPLICAS	READY	UPDATED	UNAVAILABLE
aws-example-md-0	ScalingUp	7m53s	aws-example	v1.26.3	4	3	4	1

In this example, there are two replicas, but only 1 is ready. One replica is unavailable, and the `ScalingUp` phase means a new Machine is being created.

- Identify the replacement Machine using this command:

```
export NAME_NEW_MACHINE=$(kubectl --kubeconfig ${CLUSTER_NAME}.conf get
machines \
  -l=cluster.x-k8s.io/deployment-name=${CLUSTER_NAME}-md-0 \
  -ojsonpath='{.items[?(.status.phase=="Running")].metadata.name}'
echo "$NAME_NEW_MACHINE"
```

```
aws-example-md-0-cb9c9bbf7-hcl8z aws-example-md-0-cb9c9bbf7-rtdqw aws-example-
md-0-cb9c9bbf7-td29r aws-example-md-0-cb9c9bbf7-w64kg
```

If the output is empty, the new Machine has probably exited the `Provisioning` phase and entered the `Running` phase.

6. Identify the replacement Node using this command:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf get nodes
```

NAME	STATUS	ROLES
AGE VERSION		
ip-10-0-106-183.us-west-2.compute.internal 20m v1.26.6	Ready	control-plane,master
ip-10-0-158-104.us-west-2.compute.internal 23m v1.26.6	Ready	control-plane,master
ip-10-0-203-138.us-west-2.compute.internal 22m v1.26.6	Ready	control-plane,master
ip-10-0-70-169.us-west-2.compute.internal 22m v1.26.6	Ready	<none>
ip-10-0-77-176.us-west-2.compute.internal 22m v1.26.6	Ready	<none>
ip-10-0-86-58.us-west-2.compute.internal 57s v1.26.6	Ready	<none>
ip-10-0-96-61.us-west-2.compute.internal 22m v1.26.6	Ready	<none>

If the output is empty, the Node resource is not yet available, or does not yet have the expected annotation. Wait a few minutes, then repeat the command.

7.12 EKS Infrastructure

Enterprise

Gov Advanced

When installing DKP on your EKS infrastructure, you can choose from multiple configuration types.

If not already done, refer to [Get Started \(see page 79\)](#) section of the documentation for:

- [Resource Requirements \(see page 100\)](#)
- [Install Overview \(see page 108\)](#)
- [Prerequisites for Install \(see page 121\)](#)

The steps for creating and accessing your cluster are listed below:

- [EKS Introduction](#) (see page 1128)
- [Minimal User Permission for EKS Cluster Creation](#) (see page 1129)
- [EKS Cluster IAM Permissions and Roles](#) (see page 1132)
- [Create an EKS Cluster from the CLI](#) (see page 1137)
- [Create an EKS Cluster from the DKP UI](#) (see page 1144)
- [Grant Cluster Access](#) (see page 1147)
- [Explore EKS Cluster](#) (see page 1148)
- [Manage EKS Nodepools](#) (see page 1151)
- [Delete EKS Cluster from CLI](#) (see page 1153)
- [Delete EKS Cluster from the DKP UI](#) (see page 1154)

7.12.1 EKS Introduction

Enterprise

Gov Advanced

DKP brings value to EKS customers by providing all components needed for a production-ready Kubernetes environment. DKP provides the capability to provision EKS clusters using the DKP UI. In addition to above, DKP also provides the ability to upgrade your EKS clusters using the DKP platform, making it possible to manage the complete lifecycle of EKS clusters from a centralized platform.

DKP adds value to Amazon EKS through:

- **Time to Value** in hours/days to get to production, instead of weeks/months, or even failure. Particularly in complex environments like air-gapped, customers tried various options and even after spending millions did not see success, or saw Day 2 later than they could have. We delivered results in hours/days.
- **Less Risk**
 - **Cloud-Native Expertise** eliminates the lack of skills issue. Our industry-leading expertise closes skill gaps on the customer side, avoids costly mistakes, transfers skills, and improves project success rates while shortening timelines.
 - **Simplicity** mitigates operational complexity. We focus on a great user experience and automate the hard parts of cloud-native operations to get customers to Day 2 faster and meet all Day 2 operational challenges. This frees up time on the customer side to build what differentiates them, instead of reinventing the wheel for Kubernetes operations.
 - **Military-Grade Security** alleviates security concerns. The D2iQ Kubernetes Platform can be configured to meet NSA Kubernetes security hardening guidelines. D2iQ Kubernetes Platform and supported add-on components are security scanned and secure out of the box. Encryption of data-at-rest, FIPS compliance, and fully supported air-gapped deployments round out D2iQ offerings.
- **Lower TCO** with operational insights and a simpler platform that curates needed capabilities from Amazon EKS and the open source community that reduces the time and cost of consulting engagements as well as ongoing support costs.
- **Enterprise-grade Kubernetes** - comes with a curated list of Day 2 applications necessary for running Kubernetes in production.

- **One platform for all** - Single platform to manage multiple clusters on any infrastructure cloud, on-premise, and edge.
- **D2iQ GitOps and EKS** - Delivering business value through applications is the primary goal of any Kubernetes cluster. While EKS provides the hosted framework that leads the market, delivering applications to your environment requires a mature and integrated approach. D2iQ DKP provides workspace and project level constructs to a Kubernetes cluster so that application teams have division of resources, security and cost optimization at the the project and namespace level.
 - Projects deliver applications via the built in GitOps via FluxCD - Just provide a Git repository and DKP does the rest
 - Through integration with Kubecost, DKP monitors utilization of project resources and proves real time reporting for performance and cost optimization
 - Security of projects is defined through DKP integration of customer authentication methods and is reinforced through several layers of application security
- **Cluster Lifecycle Management through CAPI** - Through the use of cluster API, DKP gives customers full lifecycle management of their EKS clusters with the ability to instantiate new EKS clusters through a unified API. This allows administrators to deploy new EKS clusters through code and deliver consistent cluster configurations.
- Time to application value is greatly reduced by minimizing the steps necessary to provision a cluster segment clusters through integrated permissions
- Secure and reliable cluster deployments
- Automatic day 2 operations of EKS clusters (Monitoring, Logging, Central Management, Security, Cost Optimization)
- Day 2 GitOps integration with every EKS cluster

7.12.2 Minimal User Permission for EKS Cluster Creation

The following is a cloudformation stack which adds a policy named `eks-bootstrapper` to manage EKS cluster to the `dkp-bootstrapper-role` created by the cloudformation stack in the [Minimal Permissions and Role to Create Cluster](#) (see page 1060) section. Consult the [Leveraging the Role](#) (see page 1064) section for an example of how to use this role and how a system administrator wants to expose using the permissions.

```

AWSTemplateFormatVersion: 2010-09-09
Parameters:
  existingBootstrapperRole:
    Type: CommaDelimitedList
    Description: 'Name of existing minimal role you want to add to add EKS cluster
management permissions to'
    Default: dkp-bootstrapper-role
Resources:
  EKSMinimumPermissions:
    Properties:
      Description: Minimal user policy to manage eks clusters
      ManagedPolicyName: eks-bootstrapper
      PolicyDocument:

```

Statement:

- Action:
 - 'ssm:GetParameter'
- Effect: Allow
- Resource:
 - 'arn::*:ssm::*:parameter/aws/service/eks/optimized-ami/*'
- Action:
 - 'iam:CreateServiceLinkedRole'
- Condition:
 - StringLike:
 - 'iam:AWSServiceName': eks.amazonaws.com
- Effect: Allow
- Resource:
 - >-
 - arn::*:iam::*:role/aws-service-role/eks.amazonaws.com/

AWSServiceRoleForAmazonEKS

- Action:
 - 'iam:CreateServiceLinkedRole'
- Condition:
 - StringLike:
 - 'iam:AWSServiceName': eks-nodegroup.amazonaws.com
- Effect: Allow
- Resource:
 - >-
 - arn::*:iam::*:role/aws-service-role/eks-nodegroup.amazonaws.com/

AWSServiceRoleForAmazonEKSNodegroup

- Action:
 - 'iam:CreateServiceLinkedRole'
- Condition:
 - StringLike:
 - 'iam:AWSServiceName': eks-fargate.amazonaws.com
- Effect: Allow
- Resource:
 - >-
 - arn:aws:iam::*:role/aws-service-role/eks-fargate-pods.amazonaws.com/

AWSServiceRoleForAmazonEKSFargate

- Action:
 - 'iam:GetRole'
 - 'iam:ListAttachedRolePolicies'
- Effect: Allow
- Resource:
 - 'arn::*:iam::*:role/*'
- Action:
 - 'iam:GetPolicy'
- Effect: Allow
- Resource:
 - 'arn:aws:iam::aws:policy/AmazonEKSClusterPolicy'
- Action:
 - 'eks:DescribeCluster'
 - 'eks:ListClusters'
 - 'eks:CreateCluster'
 - 'eks:TagResource'

```

- 'eks:UpdateClusterVersion'
- 'eks>DeleteCluster'
- 'eks:UpdateClusterConfig'
- 'eks:UntagResource'
- 'eks:UpdateNodegroupVersion'
- 'eks:DescribeNodegroup'
- 'eks>DeleteNodegroup'
- 'eks:UpdateNodegroupConfig'
- 'eks:CreateNodegroup'
- 'eks:AssociateEncryptionConfig'
- 'eks:ListIdentityProviderConfigs'
- 'eks:AssociateIdentityProviderConfig'
- 'eks:DescribeIdentityProviderConfig'
- 'eks:DisassociateIdentityProviderConfig'
Effect: Allow
Resource:
- 'arn*:eks:*:*:cluster/*'
- 'arn*:eks:*:*:nodegroup/*/*/*'
- Action:
- 'ec2:AssociateVpcCidrBlock'
- 'ec2:DisassociateVpcCidrBlock'
- 'eks:ListAddons'
- 'eks:CreateAddon'
- 'eks:DescribeAddonVersions'
- 'eks:DescribeAddon'
- 'eks>DeleteAddon'
- 'eks:UpdateAddon'
- 'eks:TagResource'
- 'eks:DescribeFargateProfile'
- 'eks:CreateFargateProfile'
- 'eks>DeleteFargateProfile'
Effect: Allow
Resource:
- '*'
- Action:
- 'iam:PassRole'
Condition:
StringEquals:
  'iam:PassedToService': eks.amazonaws.com
Effect: Allow
Resource:
- '*'
- Action:
- 'kms:CreateGrant'
- 'kms:DescribeKey'
Condition:
  'ForAnyValue:StringLike':
    'kms:ResourceAliases': alias/cluster-api-provider-aws-*
Effect: Allow
Resource:
- '*'
Version: 2012-10-17

```

```
Roles: !Ref existingBootstrapperRole
Type: 'AWS::IAM::ManagedPolicy'
```

If your role is not named `dkp-bootstrapper-role` change the parameter on line 6 of the file.

To create the resources in the cloudformation stack, copy the contents above into a file. Before executing the following command, replace `MYFILENAME.yaml` and `MYSTACKNAME` with the intended values for your system:

```
aws cloudformation create-stack --template-body=file://MYFILENAME.yaml --stack-name=MYSTACKNAME --capabilities CAPABILITY_NAMED_IAM
```

7.12.2.1 Next Step:

[EKS: Cluster IAM Policies and Roles \(see page 412\)](#)

7.12.3 EKS Cluster IAM Permissions and Roles

This section guides a DKP user in creating IAM Policies and Instance Profiles that governs who has access to the cluster. The IAM Role is used by the cluster's control plane and worker nodes using the provided AWS CloudFormation Stack specific to EKS. This CloudFormation Stack has additional permissions that are used to delegate access roles for other users.

7.12.3.1 Prerequisites from AWS:

- The user you delegate from your role must have a minimum set of permissions, see [User Roles and Instance Profiles \(see page 1060\)](#) page for AWS.
- Create the [Cluster IAM Policies \(see page 1066\)](#) in your AWS account.

7.12.3.2 EKS IAM Artifacts

7.12.3.2.1 Policies

- `controllers-eks.cluster-api-provider-aws.sigs.k8s.io` - enumerates the Actions required by the workload cluster to create and modify EKS clusters in the user's AWS Account. It is attached to the existing `control-plane.cluster-api-provider-aws.sigs.k8s.io` role

- `eks-nodes.cluster-api-provider-aws.sigs.k8s.io` - enumerates the Actions required by the EKS workload cluster's worker machines. It is attached to the existing `nodes.cluster-api-provider-aws.sigs.k8s.io`

7.12.3.2.2 Roles

- `eks-controlplane.cluster-api-provider-aws.sigs.k8s.io` - is the Role associated with EKS cluster control planes

NOTE: `control-plane.cluster-api-provider-aws.sigs.k8s.io` and `nodes.cluster-api-provider-aws.sigs.k8s.io` roles were created by [Cluster IAM Policies and Roles](#) (see page 1066) in AWS.

Below is a [CloudFormation stack](#)⁸⁴² that includes IAM policies and roles required to setup EKS Clusters:

```
AWSTemplateFormatVersion: 2010-09-09
Parameters:
  existingControlPlaneRole:
    Type: CommaDelimitedList
    Description: 'Names of existing Control Plane Role you want to add to the newly created EKS Managed Policy for AWS cluster API controllers'
    Default: control-plane.cluster-api-provider-aws.sigs.k8s.io
  existingNodeRole:
    Type: CommaDelimitedList
    Description: 'ARN of the Nodes Managed Policy to add to the role for nodes'
    Default: nodes.cluster-api-provider-aws.sigs.k8s.io
Resources:
  AWSIAMManagedPolicyControllersEKS:
    Properties:
      Description: For the Kubernetes Cluster API Provider AWS Controllers
      ManagedPolicyName: controllers-eks.cluster-api-provider-aws.sigs.k8s.io
      PolicyDocument:
        Statement:
          - Action:
              - 'ssm:GetParameter'
            Effect: Allow
            Resource:
              - 'arn:*:ssm:*:*:parameter/aws/service/eks/optimized-ami/*'
          - Action:
              - 'iam:CreateServiceLinkedRole'
            Condition:
              StringLike:
```

⁸⁴² <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html>


```

    'iam:AWSServiceName': eks.amazonaws.com
  Effect: Allow
  Resource:
    - >-
      arn:*:iam::*:role/aws-service-role/eks.amazonaws.com/
AWSServiceRoleForAmazonEKS
- Action:
  - 'iam:CreateServiceLinkedRole'
  Condition:
    StringLike:
      'iam:AWSServiceName': eks-nodegroup.amazonaws.com
  Effect: Allow
  Resource:
    - >-
      arn:*:iam::*:role/aws-service-role/eks-nodegroup.amazonaws.com/
AWSServiceRoleForAmazonEKSNodegroup
- Action:
  - 'iam:CreateServiceLinkedRole'
  Condition:
    StringLike:
      'iam:AWSServiceName': eks-fargate.amazonaws.com
  Effect: Allow
  Resource:
    - >-
      arn:aws:iam::*:role/aws-service-role/eks-fargate-pods.amazonaws.com/
AWSServiceRoleForAmazonEKSFargate
- Action:
  - 'iam:GetRole'
  - 'iam:ListAttachedRolePolicies'
  Effect: Allow
  Resource:
    - 'arn:*:iam::*:role/*'
- Action:
  - 'iam:GetPolicy'
  Effect: Allow
  Resource:
    - 'arn:aws:iam::aws:policy/AmazonEKSClusterPolicy'
- Action:
  - 'eks:DescribeCluster'
  - 'eks:ListClusters'
  - 'eks:CreateCluster'
  - 'eks:TagResource'
  - 'eks:UpdateClusterVersion'
  - 'eks>DeleteCluster'
  - 'eks:UpdateClusterConfig'
  - 'eks:UntagResource'
  - 'eks:UpdateNodegroupVersion'
  - 'eks:DescribeNodegroup'
  - 'eks>DeleteNodegroup'
  - 'eks:UpdateNodegroupConfig'
  - 'eks:CreateNodegroup'
  - 'eks:AssociateEncryptionConfig'

```

```

    - 'eks:ListIdentityProviderConfigs'
    - 'eks:AssociateIdentityProviderConfig'
    - 'eks:DescribeIdentityProviderConfig'
    - 'eks:DisassociateIdentityProviderConfig'
  Effect: Allow
  Resource:
    - 'arn:*:eks:*:*:cluster/*'
    - 'arn:*:eks:*:*:nodegroup/*/*/*'
- Action:
  - 'ec2:AssociateVpcCidrBlock'
  - 'ec2:DisassociateVpcCidrBlock'
  - 'eks:ListAddons'
  - 'eks:CreateAddon'
  - 'eks:DescribeAddonVersions'
  - 'eks:DescribeAddon'
  - 'eks>DeleteAddon'
  - 'eks:UpdateAddon'
  - 'eks:TagResource'
  - 'eks:DescribeFargateProfile'
  - 'eks:CreateFargateProfile'
  - 'eks>DeleteFargateProfile'
  Effect: Allow
  Resource:
    - '*'
- Action:
  - 'iam:PassRole'
  Condition:
    StringEquals:
      'iam:PassedToService': eks.amazonaws.com
  Effect: Allow
  Resource:
    - '*'
- Action:
  - 'kms:CreateGrant'
  - 'kms:DescribeKey'
  Condition:
    'ForAnyValue:StringLike':
      'kms:ResourceAliases': alias/cluster-api-provider-aws-*
  Effect: Allow
  Resource:
    - '*'
  Version: 2012-10-17
  Roles: !Ref existingControlPlaneRole
  Type: 'AWS::IAM::ManagedPolicy'
AWSIAMManagedEKSNodesPolicy:
  Properties:
    Description: Additional Policies to nodes role to work for EKS
    ManagedPolicyName: eks-nodes.cluster-api-provider-aws.sigs.k8s.io
    PolicyDocument:
      Statement:
        - Action:
            - "ec2:AssignPrivateIpAddresses"

```

```

    - "ec2:AttachNetworkInterface"
    - "ec2:CreateNetworkInterface"
    - "ec2>DeleteNetworkInterface"
    - "ec2:DescribeInstances"
    - "ec2:DescribeTags"
    - "ec2:DescribeNetworkInterfaces"
    - "ec2:DescribeInstanceTypes"
    - "ec2:DetachNetworkInterface"
    - "ec2:ModifyNetworkInterfaceAttribute"
    - "ec2:UnassignPrivateIpAddresses"
  Effect: Allow
  Resource:
    - '*'
- Action:
  - ec2:CreateTags
  Effect: Allow
  Resource:
    - arn:aws:ec2:*:*:network-interface/*
- Action:
  - "ec2:DescribeInstances"
  - "ec2:DescribeInstanceTypes"
  - "ec2:DescribeRouteTables"
  - "ec2:DescribeSecurityGroups"
  - "ec2:DescribeSubnets"
  - "ec2:DescribeVolumes"
  - "ec2:DescribeVolumesModifications"
  - "ec2:DescribeVpcs"
  - "eks:DescribeCluster"
  Effect: Allow
  Resource:
    - '*'
  Version: 2012-10-17
  Roles: !Ref existingNodeRole
  Type: 'AWS::IAM::ManagedPolicy'
AWSIAMRoleEKSCoontrolPlane:
  Properties:
    AssumeRolePolicyDocument:
      Statement:
        - Action:
            - 'sts:AssumeRole'
          Effect: Allow
          Principal:
            Service:
              - eks.amazonaws.com
    Version: 2012-10-17
  ManagedPolicyArns:
    - 'arn:aws:iam::aws:policy/AmazonEKSClusterPolicy'
  RoleName: eks-controlplane.cluster-api-provider-aws.sigs.k8s.io
  Type: 'AWS::IAM::Role'

```

To create the resources in the CloudFormation stack, copy the contents above into a file and execute the following command after replacing `MYFILENAME.yaml` and `MYSTACKNAME` with the intended values:

```
aws cloudformation create-stack --template-body=file://MYFILENAME.yaml --stack-name=MYSTACKNAME --capabilities CAPABILITY_NAMED_IAM
```

7.12.3.2.3 Add EKS CSI Policy

AWS CloudFormation does not support attaching an existing IAM Policy to an existing IAM Role. Add the necessary IAM policy to your worker instance profile using the `aws` CLI:

```
aws iam attach-role-policy --role-name nodes.cluster-api-provider-aws.sigs.k8s.io --policy-arn arn:aws:iam::aws:policy/service-role/AmazonEBSCSIDriverPolicy
```

7.12.4 Create an EKS Cluster from the CLI

Enterprise

Gov Advanced

- Ensure that the `KUBECONFIG` environment variable is set to the Management cluster by running `export KUBECONFIG=<Management_cluster_kubeconfig>.conf`.

7.12.4.1 Create a New EKS Kubernetes Cluster

- By default, the control-plane Nodes will be created in 3 different Availability Zones. However, the default worker Nodes will reside in a single zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:

1. Set the environment variable to the name you assigned this cluster.

```
export CLUSTER_NAME=eks-example
```

2. Make sure your AWS credentials are up-to-date. Refresh the credentials command is only necessary if you are using [Static Credentials \(see page 0\)](#) (Access Keys) otherwise, if you are using role-based authentication on a bastion host, proceed to step 3:

```
dkp update bootstrap credentials aws
```

3. Create the cluster:

```
dkp create cluster eks \
  --cluster-name=${CLUSTER_NAME} \
  --additional-tags=owner=$(whoami)
```

⚠ NOTE: If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1042\)](#).

⚠ Optional flag for ECR: Configure your cluster to use an existing local registry as a mirror when attempting to pull images. Below is an AWS ECR example:

```
export --registry-mirror-url-string=<YOUR_ECR_URL>
```

- `REGISTRY-MIRROR-URL-STRING`: the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images. Users can still pull their own images from ECR directly or use ECR as a mirror. See CLI commands for `--registry-mirror` flag here: [dkp create cluster eks \(see page 1663\)](#)

```
Generating cluster resources
cluster.cluster.x-k8s.io/eks-example created
awsmanagedcontrolplane.controlplane.cluster.x-k8s.io/eks-example-control-plane
  created
machinedeployment.cluster.x-k8s.io/eks-example-md-0 created
awsmachine.template.infrastructure.cluster.x-k8s.io/eks-example-md-0 created
eksconfig.template.bootstrap.cluster.x-k8s.io/eks-example-md-0 created
clusterresourceset.addons.cluster.x-k8s.io/calico-cni-installation-eks-example
  created
configmap/calico-cni-installation-eks-example created
configmap/tigera-operator-eks-example created
clusterresourceset.addons.cluster.x-k8s.io/cluster-autoscaler-eks-example
  created
configmap/cluster-autoscaler-eks-example created
clusterresourceset.addons.cluster.x-k8s.io/node-feature-discovery-eks-example
  created
configmap/node-feature-discovery-eks-example created
clusterresourceset.addons.cluster.x-k8s.io/nvidia-feature-discovery-eks-example
  created
```

```
configmap/nvidia-feature-discovery-eks-example created
```

4. Inspect or edit the cluster objects:

Use your favorite editor.

NOTE: Editing the cluster objects requires some understanding of Cluster API. Edits can prevent the cluster from deploying successfully.

The objects are [Custom Resources](#)⁸⁴³ defined by Cluster API components, and they belong in three different categories:

a. **Cluster**

A *Cluster* object has references to the infrastructure-specific and control plane objects.

Because this is an AWS cluster, there is an *AWSCluster* object that describes the infrastructure-specific cluster properties. Here, this means the AWS region, the VPC ID, subnet IDs, and security group rules required by the Pod network implementation.

b. **Control Plane**

A *AWSMangedControlPlane* object describes the control plane, which is the group of machines that run the Kubernetes control plane components, which include the etcd distributed database, the API server, the core controllers, and the scheduler. The object describes the configuration for these components. The object also has a reference to an infrastructure-specific object that describes the properties of all control plane machines.

c. **Node Pool**

A Node Pool is a collection of machines with identical properties. For example, a cluster might have one Node Pool with large memory capacity, another Node Pool with GPU support. Each Node Pool is described by three objects: The *MachinePool* references an object that describes the configuration of Kubernetes components (for example, kubelet) deployed on each node pool machine, and an infrastructure-specific object that describes the properties of all node pool machines. Here, it references a *KubeadmConfigTemplate*, and an *AWSMachineTemplate* object, which describes the instance type, the type of disk used, the size of the disk, among other properties.

For in-depth documentation about the objects, read [Concepts](#)⁸⁴⁴ in the Cluster API Book.

5. Wait for the cluster control-plane to be ready:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

```
cluster.cluster.x-k8s.io/eks-example condition met
```

⁸⁴³ <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

⁸⁴⁴ <https://cluster-api.sigs.k8s.io/user/concepts.html>

The `READY` status will become `True` after the cluster control-plane becomes ready in one of the following steps.

- After the objects are created on the API server, the Cluster API controllers reconcile them. They create infrastructure and machines. As they progress, they update the Status of each object. Konvoy provides a command to describe the current status of the cluster:

```
dkp describe cluster -c ${CLUSTER_NAME}
```

```
NAME                                                    READY
SEVERITY REASON SINCE MESSAGE
Cluster/eks-example                                    True
10m
├─ControlPlane - AWSManagedControlPlane/eks-example-control-plane True
10m
├─Workers
  └─MachineDeployment/eks-example-md-0                  True
26s
    └─Machine/eks-example-md-0-78fcd7c7b7-66ntt       True
84s
    └─Machine/eks-example-md-0-78fcd7c7b7-b9qmc       True
84s
    └─Machine/eks-example-md-0-78fcd7c7b7-v5vfq       True
84s
    └─Machine/eks-example-md-0-78fcd7c7b7-zl6m2       True
84s
```

- As they progress, the controllers also create Events. List the Events using this command:

```
kubectl get events | grep ${CLUSTER_NAME}
```

For brevity, the example uses `grep`. It is also possible to use separate commands to get Events for specific objects. For example, `kubectl get events --field-selector involvedObject.kind="AWSCluster"` and `kubectl get events --field-selector involvedObject.kind="AWSMachine"`.

```
46m          Normal    SuccessfulCreateVPC
awsmanagedcontrolplane/eks-example-control-plane   Created new managed VPC
"vpc-05e775702092abf09"
46m          Normal    SuccessfulSetVPCAttributes
awsmanagedcontrolplane/eks-example-control-plane   Set managed VPC attributes
for "vpc-05e775702092abf09"
```

```

46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane    Created new managed Subnet
"subnet-0419dd3f2dfd95ff8"
46m          Normal    SuccessfulModifySubnetAttributes
awsmanagedcontrolplane/eks-example-control-plane    Modified managed Subnet
"subnet-0419dd3f2dfd95ff8" attributes
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane    Created new managed Subnet
"subnet-0e724b128e3113e47"
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane    Created new managed Subnet
"subnet-06b2b31ea6a8d3962"
46m          Normal    SuccessfulModifySubnetAttributes
awsmanagedcontrolplane/eks-example-control-plane    Modified managed Subnet
"subnet-06b2b31ea6a8d3962" attributes
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane    Created new managed Subnet
"subnet-0626ce238be32bf98"
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane    Created new managed Subnet
"subnet-0f53cf59f83177800"
46m          Normal    SuccessfulModifySubnetAttributes
awsmanagedcontrolplane/eks-example-control-plane    Modified managed Subnet
"subnet-0f53cf59f83177800" attributes
46m          Normal    SuccessfulCreateSubnet
awsmanagedcontrolplane/eks-example-control-plane    Created new managed Subnet
"subnet-0878478f6bbf153b2"
46m          Normal    SuccessfulCreateInternetGateway
awsmanagedcontrolplane/eks-example-control-plane    Created new managed Internet
Gateway "igw-09fb52653949d4579"
46m          Normal    SuccessfulAttachInternetGateway
awsmanagedcontrolplane/eks-example-control-plane    Internet Gateway
"igw-09fb52653949d4579" attached to VPC "vpc-05e775702092abf09"
46m          Normal    SuccessfulCreateNATGateway
awsmanagedcontrolplane/eks-example-control-plane    Created new NAT Gateway
"nat-06356aac28079952d"
46m          Normal    SuccessfulCreateNATGateway
awsmanagedcontrolplane/eks-example-control-plane    Created new NAT Gateway
"nat-0429d1cd9d956bf35"
46m          Normal    SuccessfulCreateNATGateway
awsmanagedcontrolplane/eks-example-control-plane    Created new NAT Gateway
"nat-059246bcc9d4e88e7"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Created managed RouteTable
"rtb-01689c719c484fd3c"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane    Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Associated managed
RouteTable "rtb-01689c719c484fd3c" with subnet "subnet-0419dd3f2dfd95ff8"

```



```

46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Created managed RouteTable
"rtb-065af81b9752eeb69"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane    Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Associated managed
RouteTable "rtb-065af81b9752eeb69" with subnet "subnet-0e724b128e3113e47"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Created managed RouteTable
"rtb-03eeff810a89afc98"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane    Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Associated managed
RouteTable "rtb-03eeff810a89afc98" with subnet "subnet-06b2b31ea6a8d3962"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Created managed RouteTable
"rtb-0fab36f8751fdee73"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane    Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Associated managed
RouteTable "rtb-0fab36f8751fdee73" with subnet "subnet-0626ce238be32bf98"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Created managed RouteTable
"rtb-0e5c9c7bbc3740a0f"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane    Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Associated managed
RouteTable "rtb-0e5c9c7bbc3740a0f" with subnet "subnet-0f53cf59f83177800"
46m          Normal    SuccessfulCreateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Created managed RouteTable
"rtb-0bf58eb5f73c387af"
46m          Normal    SuccessfulCreateRoute
awsmanagedcontrolplane/eks-example-control-plane    Created route {...
46m          Normal    SuccessfulAssociateRouteTable
awsmanagedcontrolplane/eks-example-control-plane    Associated managed
RouteTable "rtb-0bf58eb5f73c387af" with subnet "subnet-0878478f6bbf153b2"
46m          Normal    SuccessfulCreateSecurityGroup
awsmanagedcontrolplane/eks-example-control-plane    Created managed
SecurityGroup "sg-0b045c998a120a1b2" for Role "node-eks-additional"
46m          Normal    InitiatedCreateEKSControlPlane
awsmanagedcontrolplane/eks-example-control-plane    Initiated creation of a new
EKS control plane default_eks-example-control-plane
37m          Normal    SuccessfulCreateEKSControlPlane
awsmanagedcontrolplane/eks-example-control-plane    Created new EKS control
plane default_eks-example-control-plane
37m          Normal    SuccessfulCreateKubeconfig
awsmanagedcontrolplane/eks-example-control-plane    Created kubeconfig for
cluster "eks-example"

```


```

37m          Normal    SuccessfulCreateUserKubeconfig
awsmanagedcontrolplane/eks-example-control-plane    Created user kubeconfig for
  cluster "eks-example"
27m          Normal    SuccessfulCreate
awsmachine/eks-example-md-0-4t9nc                    Created new node instance
with id "i-0aecc1897c93df740"
26m          Normal    SuccessfulDeleteEncryptedBootstrapDataSecrets
awsmachine/eks-example-md-0-4t9nc                    AWS Secret entries
containing userdata deleted
26m          Normal    SuccessfulSetNodeRef                                machine/
eks-example-md-0-78fcd7c7b7-fn7x9                    ip-10-0-88-24.us-west-2.compute.inte
rnal
26m          Normal    SuccessfulSetNodeRef                                machine/
eks-example-md-0-78fcd7c7b7-g64nv                    ip-10-0-110-219.us-west-2.compute.in
ternal
26m          Normal    SuccessfulSetNodeRef                                machine/
eks-example-md-0-78fcd7c7b7-gwc5j                    ip-10-0-101-161.us-west-2.compute.in
ternal
26m          Normal    SuccessfulSetNodeRef                                machine/
eks-example-md-0-78fcd7c7b7-j58s4                    ip-10-0-127-49.us-west-2.compute.int
ernal
46m          Normal    SuccessfulCreate
machineset/eks-example-md-0-78fcd7c7b7              Created machine "eks-
example-md-0-78fcd7c7b7-fn7x9"
46m          Normal    SuccessfulCreate
machineset/eks-example-md-0-78fcd7c7b7              Created machine "eks-
example-md-0-78fcd7c7b7-g64nv"
46m          Normal    SuccessfulCreate
machineset/eks-example-md-0-78fcd7c7b7              Created machine "eks-
example-md-0-78fcd7c7b7-j58s4"
46m          Normal    SuccessfulCreate
machineset/eks-example-md-0-78fcd7c7b7              Created machine "eks-
example-md-0-78fcd7c7b7-gwc5j"
27m          Normal    SuccessfulCreate
awsmachine/eks-example-md-0-7whkv                    Created new node instance
with id "i-06dfc0466b8f26695"
26m          Normal    SuccessfulDeleteEncryptedBootstrapDataSecrets
awsmachine/eks-example-md-0-7whkv                    AWS Secret entries
containing userdata deleted
27m          Normal    SuccessfulCreate
awsmachine/eks-example-md-0-ttgzv                    Created new node instance
with id "i-0544fce0350fd41fb"
26m          Normal    SuccessfulDeleteEncryptedBootstrapDataSecrets
awsmachine/eks-example-md-0-ttgzv                    AWS Secret entries
containing userdata deleted
27m          Normal    SuccessfulCreate
awsmachine/eks-example-md-0-v2hrf                    Created new node instance
with id "i-0498906edde162e59"
26m          Normal    SuccessfulDeleteEncryptedBootstrapDataSecrets
awsmachine/eks-example-md-0-v2hrf                    AWS Secret entries
containing userdata deleted

```

```
46m          Normal    SuccessfulCreate
machinedeployment/eks-example-md-0          Created MachineSet "eks-
example-md-0-78fcd7c7b7"
```

7.12.4.2 Known Limitations

 Be aware of these limitations in the current release of Konvoy.

- The Konvoy version used to create a workload cluster must match the Konvoy version used to delete a workload cluster.
- EKS clusters cannot be Self-managed.
- Konvoy supports deploying one workload cluster.
- Konvoy generates a set of objects for one Node Pool.
- Konvoy does not validate edits to cluster objects.

7.12.5 Create an EKS Cluster from the DKP UI

Enterprise

Gov Advanced

The DKP user interface allows you to provision a Cluster from your browser quickly and easily.

7.12.5.1 Create an AWS Infrastructure Provider

Before you create a Cluster, you first need to create an AWS infrastructure provider to hold your AWS/EKS Credentials:

1. [Get the AWS RoleARN \(see page 1060\)](#).

```
aws iam get-role --role-name <role-name> --query 'Role.[RoleName, Arn]' --
output text
```

2. Select **Infrastructure Providers** from the Dashboard menu.
3. Select **Add Infrastructure Provider**.
4. Choose a workspace. If you are already in a workspace, the provider is automatically created in that workspace.
5. Ensure you select **Amazon Web Services**.
6. Add a **Name** for your Infrastructure Provider and include the **Role ARN** from Step 1 above.

7. Select **Save**.



If you choose to, you can use static credentials. However, this method is not as secure so it is not recommended.

7.12.5.2 Provision an EKS Cluster

Follow these steps to provision the EKS cluster:

1. From the top menu bar, select your target workspace.
2. Select **Clusters > Add Cluster**.
This begins the provisioning workflow.
3. Choose **Create Cluster**.
4. Enter the **Cluster Name**.
5. Select **EKS** from the **Choose Infrastructure** choices.
6. If available, choose a **Kubernetes Version**. Otherwise, the [default Kubernetes version](#)⁸⁴⁵ installs.
7. Select a data center **region** or specify a custom **region**.
8. Edit your worker **Node Pools** as necessary. You can choose the **Number of Nodes**, the **Machine Type**, and our **IAM Instance Profile**. For the worker pool, you can also choose a **Worker Availability Zone**.
9. Add any additional **Labels** or **Infrastructure Provider Tags** as necessary.
10. Validate your inputs, and then select **Create**.

You are redirected to the **Clusters** page, where you see your **Cluster** in the **Provisioning** status. Hover over the status to view the details.

After about 15 minutes, your **Cluster** should be in the **Provisioned** status.

See [AWS RoleARN](#)⁸⁴⁶ for more information from the AWS site.

7.12.5.3 Access EKS Cluster

After the cluster is successfully attached(managed), you can retrieve a custom `kubeconfig` file from the UI using your Kommander administrator credentials.

⁸⁴⁵ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/150765792/DKP+2.5.0+Supported+Kubernetes+Versions>

⁸⁴⁶ https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_identifiers.html#identifiers-arns

7.12.5.4 IAM User and Role Access for EKS Clusters

When creating an EKS cluster through the UI, the `kubeconfig` that is returned using the download `kubeconfig` button allows access for 15 minutes. To follow best practices for AWS security, you should configure accessing the EKS cluster using IAM role or user based authentication. This allows account administrators to monitor all actions made.

To enable IAM based cluster access, follow the steps below:

1. Download the `kubeconfig` by selecting the Download `kubeconfig` button on the top section of the UI.
2. Using that `kubeconfig`, edit the config map with a command similar to the one below:

```
kubectl --kubeconfig=MYCLUSTER.conf edit cm -n kube-system aws-auth
```

3. Modify the `mapRoles` and `mapUsers` objects according to the permissions as needed. The example below is mapping the `arn:aws:iam::MYAWSACCOUNTID:role/PowerUser` role to `systems:masters` on the Kubernetes cluster:

```
apiVersion: v1
data:
  mapRoles: |
    - groups:
      - system:bootstrappers
      - system:nodes
      rolearn: arn:aws:iam::MYAWSACCOUNTID:role/nodes.cluster-api-provider-
aws.sigs.k8s.io
      username: system:node:{{EC2PrivateDNSName}}
    - groups:
      - system:masters
      rolearn: arn:aws:iam::MYAWSACCOUNTID:role/PowerUser
      username: admin
kind: ConfigMap
```

For more information, consult the [Enabling IAM user and role access](#)⁸⁴⁷ to your cluster guide and the [Kubernetes RBAC guide](#)⁸⁴⁸.

4. From your management cluster run the following command to fetch a `kubeconfig` that uses IAM based permissions by running:

⁸⁴⁷ <https://docs.aws.amazon.com/eks/latest/userguide/add-user-role.html>

⁸⁴⁸ <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

```
dkp get kubeconfig -c ${EKS_CLUSTER_NAME} -n ${KOMMANDER_WORKSPACE_NAMESPACE}
>> ${EKS_CLUSTER_NAME}.conf
```

7.12.6 Grant Cluster Access

Enterprise

Gov Advanced

7.12.6.1 How to Grant EKS Cluster Access

You can access your cluster using AWS IAM roles in the dashboard. When you create an EKS cluster, the IAM entity is granted `system:masters` permissions in [Kubernetes Role Based Access Control \(RBAC\) configuration](#).⁸⁴⁹



More information about the configuration of the EKS control plane can be found on the [EKS Cluster IAM Policies and Roles](#) (see page 1135) page.

If the EKS cluster was created as a cluster using a self-managed AWS cluster that uses IAM Instance Profiles, you will need to modify the `IAMAuthenticatorConfig` field in the `AWSManagedControlPlane` API object to allow other IAM entities to access the EKS workload cluster. Follow the steps below:

1. Run the following command with your `KUBECONFIG` configured to select the self-managed cluster previously used to create the workload EKS cluster. Ensure you substitute `${CLUSTER_NAME}` and `${CLUSTER_NAMESPACE}` with their corresponding values for your cluster.

```
kubectl edit awsmanagedcontrolplane ${CLUSTER_NAME}-control-plane -n ${CLUSTER_NAMESPACE}
```

2. Edit the `IamAuthenticatorConfig` field with the IAM Role to the corresponding Kubernetes Role. In this example, the IAM role `arn:aws:iam::111122223333:role/PowerUser` is granted the cluster role `system:masters`. Note that this example uses example AWS resource [ARNs](#)⁸⁵⁰, so these values should be substituted for real values in the corresponding AWS account.

⁸⁴⁹ <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

⁸⁵⁰ <https://docs.aws.amazon.com/general/latest/gr/aws-arns-and-namespaces.html>

```
iamAuthenticatorConfig:
  mapRoles:
    - groups:
      - system:bootstrappers
      - system:nodes
      rolearn: arn:aws:iam::111122223333:role/my-node-role
      username: system:node:{{EC2PrivateDNSName}}
    - groups:
      - system:masters
      rolearn: arn:aws:iam::111122223333:role/PowerUser
      username: admin
```

For further instructions on changing or assigning `roles` or `clusterroles` to which you can map IAM users or roles, see [Amazon Enabling IAM access to your cluster](#)⁸⁵¹.

7.12.7 Explore EKS Cluster

Enterprise

Gov Advanced

This guide explains how to use the command line to interact with your newly deployed Kubernetes cluster. Before you start, make sure you have created a workload cluster, as described in [Create a New Cluster](#) (see [page 1140](#)).

7.12.7.1 Explore the new Kubernetes cluster

Follow these steps:

1. Get a kubeconfig file for the workload cluster:

When the workload cluster is created, the cluster lifecycle services generate a kubeconfig file for the workload cluster, and write it to a `Secret`. The kubeconfig file is scoped to the cluster administrator.

Get the kubeconfig from the `Secret`, and write it to a file, using this command:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

2. List the Nodes using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

⁸⁵¹ <https://docs.aws.amazon.com/eks/latest/userguide/add-user-role.html>

NAME	STATUS	ROLES	AGE	VERSION
ip-10-0-122-211.us-west-2.compute.internal eks-ba74326	Ready	<none>	35m	v1.25.13-
ip-10-0-127-74.us-west-2.compute.internal eks-ba74326	Ready	<none>	35m	v1.25.13-
ip-10-0-71-155.us-west-2.compute.internal eks-ba74326	Ready	<none>	35m	v1.25.13-
ip-10-0-93-47.us-west-2.compute.internal eks-ba74326	Ready	<none>	35m	v1.25.13-

NOTE: It may take a few minutes for the Status to move to `Ready` while the Pod network is deployed. The Nodes' Status should change to Ready soon after the `calico-node` DaemonSet Pods are Ready.

- List the Pods using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get --all-namespaces pods
```

NAMESPACE	STATUS	RESTARTS	NAME	AGE	READY
calico-system	Running	0	calico-kube-controllers-7d6749878f-ccsx9	34m	1/1
calico-system	Running	0	calico-node-2r6l8	34m	1/1
calico-system	Running	0	calico-node-5pdlb	34m	1/1
calico-system	Running	0	calico-node-n24hh	34m	1/1
calico-system	Running	0	calico-node-qrh7p	34m	1/1
calico-system	Running	0	calico-typha-7bbcb87696-7pk45	34m	1/1
calico-system	Running	0	calico-typha-7bbcb87696-t4c8r	34m	1/1
calico-system	Running	0	csi-node-driver-bz48k	34m	2/2
calico-system	Running	0	csi-node-driver-k5mmk	34m	2/2
calico-system	Running	0	csi-node-driver-nvcck	34m	2/2
calico-system	Running	0	csi-node-driver-x4xnh	34m	2/2
kube-system	Running	0	aws-node-2xp86	35m	1/1

kube-system	aws-node-5f2kx	1/1
Running 0	35m	
kube-system	aws-node-6lzm7	1/1
Running 0	35m	
kube-system	aws-node-pz8c6	1/1
Running 0	35m	
kube-system	cluster-autoscaler-789d86b489-sz9x2	0/1
Init:0/1 0	36m	
kube-system	coredns-57ff979f67-pk5cg	1/1
Running 0	75m	
kube-system	coredns-57ff979f67-sf2j9	1/1
Running 0	75m	
kube-system	ebs-csi-controller-5f6bd5d6dc-bplwm	6/6
Running 0	36m	
kube-system	ebs-csi-controller-5f6bd5d6dc-dpjt7	6/6
Running 0	36m	
kube-system	ebs-csi-node-7hmm5	3/3
Running 0	35m	
kube-system	ebs-csi-node-l4vfh	3/3
Running 0	35m	
kube-system	ebs-csi-node-mfr7c	3/3
Running 0	35m	
kube-system	ebs-csi-node-v8krq	3/3
Running 0	35m	
kube-system	kube-proxy-7fc5x	1/1
Running 0	35m	
kube-system	kube-proxy-vvkmk	1/1
Running 0	35m	
kube-system	kube-proxy-x6hcc	1/1
Running 0	35m	
kube-system	kube-proxy-x8frb	1/1
Running 0	35m	
kube-system	snapshot-controller-8ff89f489-4cfv	1/1
Running 0	36m	
kube-system	snapshot-controller-8ff89f489-78gg8	1/1
Running 0	36m	
node-feature-discovery	node-feature-discovery-master-7d5985467-52fcn	1/1
Running 0	36m	
node-feature-discovery	node-feature-discovery-worker-88hr7	1/1
Running 0	34m	
node-feature-discovery	node-feature-discovery-worker-h95nq	1/1
Running 0	35m	
node-feature-discovery	node-feature-discovery-worker-lfghg	1/1
Running 0	34m	
node-feature-discovery	node-feature-discovery-worker-prc8p	1/1
Running 0	35m	
tigera-operator	tigera-operator-6dcd98c8ff-k97hq	1/1
Running 0	36m	

7.12.8 Manage EKS Nodepools

Enterprise

Gov Advanced

- Ensure that the `KUBECONFIG` environment variable is set to the self-managed cluster by running `export KUBECONFIG={SELF_MANAGED_AWS_CLUSTER}.conf`

Node pools are part of a cluster and managed as a group, and can be used to manage a group of machines using common properties. New default clusters created by Konvoy contain one node pool of worker nodes that have the same configuration.

You can create additional node pools for specialized hardware or other configurations. For example, if you want to tune your memory usage on a cluster where you need maximum memory for some machines and minimal memory on others, you can create a new node pool with those specific resource needs.

- NOTE:** Konvoy implements node pools using Cluster API [MachineDeployments](#)⁸⁵².

7.12.8.1 Create a node pool

Availability zones (AZs) are isolated locations within data center regions from which public cloud services originate and operate. Because all the nodes in a node pool are deployed in a single Availability Zone, you may wish to create additional node pools to ensure your cluster has nodes deployed in multiple Availability Zones.

- By default, the first [Availability Zone](#)⁸⁵³ in the region is used for the nodes in the node pool. To create the nodes in a different Availability Zone set the appropriate `--availability-zone`.

To create a new EKS node pool with 3 replicas, run:

```
dkp create nodepool eks ${NODEPOOL_NAME} \
  --cluster-name=${CLUSTER_NAME} \
```

⁸⁵² <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

⁸⁵³ https://aws.amazon.com/about-aws/global-infrastructure/regions_az/

```
--replicas=3
```

```
machinedeployment.cluster.x-k8s.io/example created
awscloudformation.infrastructure.cluster.x-k8s.io/example created
eksconfigtemplate.bootstrap.cluster.x-k8s.io/example created
✓ Creating default/example nodepool resources
```

Advanced users can use a combination of the `--dry-run` and `--output=yaml` flags to get a complete set of node pool objects to modify locally or store in version control.

7.12.8.2 Scaling Up Node Pools

To scale up a node pool in a cluster, run:

```
dkp scale nodepools ${NODEPOOL_NAME} --replicas=5 --cluster-name=${CLUSTER_NAME}
```

Your output should be similar to this example, indicating the scaling is in progress:

```
✓ Scaling node pool example to 5 replicas
```

After a few minutes, you can list the node pools to:

```
dkp get nodepools --cluster-name=${CLUSTER_NAME}
```

Your output should be similar to this example, with the number of DESIRED and READY replicas increased to 5:

NODEPOOL	DESIRED	READY	KUBERNETES VERSION
example	5	5	v1.25.13
eks-example-md-0	4	4	v1.25.13

7.12.8.3 Delete EKS Node Pools

Delete node pools in a cluster

Deleting a node pool deletes the Kubernetes nodes and the underlying infrastructure. All nodes are drained prior to deletion and the pods running on those nodes are rescheduled.

To delete a node pool from a managed cluster, run:

```
dkp delete nodepool ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME}
```

Here, `example` is the node pool to be deleted.

The expected output will be similar to the following example, indicating the node pool is being deleted:

```
✓ Deleting default/example nodepool resources
```

Deleting an invalid node pool results in output similar to this example:

```
dkp delete nodepool ${CLUSTER_NAME}-md-invalid --cluster-name=${CLUSTER_NAME}
```

```
MachineDeployments or MachinePools.infrastructure.cluster.x-k8s.io "no
MachineDeployments or MachinePools found for cluster eks-example" not found
```

7.12.9 Delete EKS Cluster from CLI

Enterprise

Gov Advanced

ⓘ Ensure that the `KUBECONFIG` environment variable is set to the self-managed cluster by running `export KUBECONFIG={SELF_MANAGED_AWS_CLUSTER}.conf`

7.12.9.1 Delete the EKS cluster

Follow these steps:

1. Ensure your AWS credentials are up to date. If you are using user profiles, then refresh the credentials using the command below. Otherwise, proceed to step 2.

```
dkp update bootstrap credentials aws
```

2. Delete the Kubernetes cluster and wait a few minutes:

Before deleting the cluster, `dkp` deletes all Services of type LoadBalancer on the cluster. Each Service is backed by an AWS Classic ELB. Deleting the Service deletes the ELB that backs it. To skip this step, use the flag `--delete-kubernetes-resources=false`.

NOTE: Do not skip this step if the VPC is managed by DKP. When DKP deletes the cluster, it deletes the VPC. If the VPC has any EKS Classic ELBs, EKS does not allow the VPC to be deleted, and DKP cannot delete the cluster.

```
dkp delete cluster --cluster-name=${CLUSTER_NAME}
```

```
✓ Deleting Services with type LoadBalancer for Cluster default/eks-example
✓ Deleting ClusterResourceSets for Cluster default/eks-example
✓ Deleting cluster resources
✓ Waiting for cluster to be fully deleted
Deleted default/eks-example cluster
```

7.12.9.2 Next Step:

After your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will install the [Kommander](#) (see page 1422) component of DKP to see your dashboard and continue customization.

7.12.9.3 Known Limitations

NOTE: Be aware of these limitations in the current release of DKP.

- The DKP version used to create the workload cluster must match the DKP version used to delete the workload cluster.


7.12.10 Delete EKS Cluster from the DKP UI

Enterprise

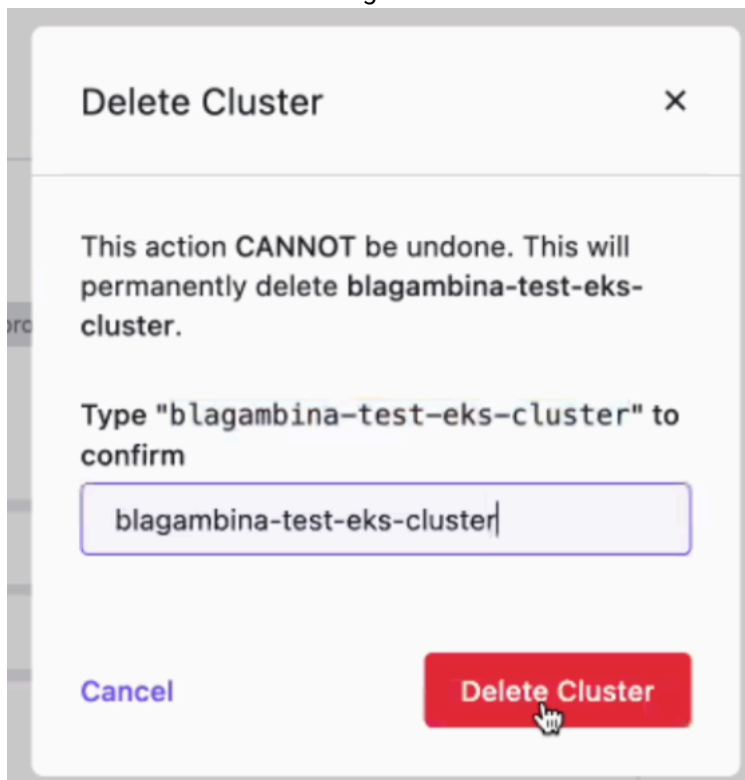
Gov Advanced

In order to delete a cluster in the UI, you must first have [created a cluster](#) (see page 1147) and have permissions to delete. Otherwise, an administrator privilege could delete a cluster as well.

1. Open the dashboard and select Clusters in the left menu.
2. Select the cluster you wish to delete and click the dotted icon in the bottom right corner.
3. Then select Delete in red.



4. When the next screen appears, copy the name of your cluster and paste it into the empty box.
5. Now execute the deletion using Delete Cluster button.



6. You will see the status as "Deleting" in the top left corner of the cluster you selected for deletion.

This removes the cluster from the DKP UI. For a generic overview of deleting clusters within the UI as well as troubleshooting, see the [Disconnect or Delete Clusters](#) (see page 870) instructions.

7.13 Azure Infrastructure

This section describes how you create a DKP cluster in Azure.

- [Azure Prerequisites](#) (see page 1156)
- [Azure using Konvoy Image Builder](#) (see page 1159)
- [Azure Bootstrap](#) (see page 1161)
- [Create a New Custom Azure Cluster](#) (see page 1163)
- [Explore new Azure Cluster](#) (see page 1173)
- [Azure Make new Cluster Self-Managed](#) (see page 1177)
- [Azure Certificate Renewal](#) (see page 1179)
- [Azure Replace a Node](#) (see page 1182)
- [Azure Delete Cluster](#) (see page 1185)

7.13.1 Azure Prerequisites

Prepare your machine and environment to run DKP

7.13.1.1 DKP Prerequisites

Before you begin using DKP you must have:

- An x86_64-based Linux or macOS machine.
- [Download](#) (see page 78) the `dkp` binary for Linux, or macOS. To check which version of DKP you installed for compatibility reasons, run the `dkp version -h` command ([dkp version](#) (see page 1749)).
- A Container engine/runtime installed is required to install DKP:
 - Version [Docker®](#)⁸⁵⁴ container engine version 18.09.2 or higher installed for Linux or MacOS - On macOS, Docker runs in a virtual machine which needs configured with at least 8 GB of memory.
 - Version 4.0 of [Podman](#)⁸⁵⁵ or higher for Linux. Host requirements found here: [Host Requirements](#)⁸⁵⁶
- [kubectl](#)⁸⁵⁷ for interacting with the running cluster.
- [Azure CLI](#)⁸⁵⁸.
- A valid Azure account with [credentials configured](#)⁸⁵⁹.
- Create a custom Azure image using [KIB](#) (see page 1461).

854 <https://docs.docker.com/get-docker/>

855 <https://podman.io/getting-started/installation>

856 <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>

857 <https://kubernetes.io/docs/tasks/tools/#kubectl>

858 <https://docs.microsoft.com/en-us/cli/azure/install-azure-cli>

859 <https://github.com/kubernetes-sigs/cluster-api-provider-azure/blob/master/docs/book/src/topics/getting-started.md#prerequisites>

- On macOS, Docker runs in a virtual machine. Configure this virtual machine with at least 8GB of memory.

7.13.1.1.1 Control plane nodes

You should have at least three control plane nodes. Each control plane node should have at least:

- 4 cores
- 16 GiB memory
- Approximately 80 GiB of free space for the volume used for `/var/lib/kubelet` and `/var/lib/containerd`.
- Disk usage must be below 85% on the root volume.

DKP on Azure defaults to deploying a `Standard_D4s_v3` virtual machine with an 128 GiB volume for the OS and an 80GiB volume for etcd storage, which meets the above requirements.

7.13.1.1.2 Worker nodes

You should have at least four worker nodes. The specific number of worker nodes required for your environment can vary depending on the cluster workload and size of the nodes. Each worker node should have at least:

- 8 cores
- 32 GiB memory
- Around 80 GiB of free space for the volume used for `/var/lib/kubelet` and `/var/lib/containerd`.
- Disk usage must be below 85% on the root volume.

DKP on Azure defaults to deploying a `Standard_D8s_v3` virtual machine with an 80 GiB volume for the OS, which meets the above requirements.

If you use these instructions to create a cluster on Azure using the DKP default settings without any edits to configuration files or additional flags, your cluster is deployed on an [Ubuntu 20.04 operating system image \(see page 69\)](#) with 3 control plane nodes, and 4 worker nodes which match the requirements above.

7.13.1.2 Azure Prerequisites

Before you begin using Konvoy with Azure, you must:

1. Log in to Azure:

```
az login
```



```
[
  {
    "cloudName": "AzureCloud",
    "homeTenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "id": "b1234567-abcd-11a1-a0a0-1234a5678b90",
    "isDefault": true,
    "managedByTenants": [],
    "name": "Mesosphere Developer Subscription",
    "state": "Enabled",
    "tenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "user": {
      "name": "user@azuremesosphere.onmicrosoft.com",
      "type": "user"
    }
  }
]
```

2. Create an Azure Service Principal (SP) by running the following command:



If an SP with the name exists, this command will rotate the password.

```
az ad sp create-for-rbac --role contributor --name "$(whoami)-konvoy" --scopes=/
subscriptions/$(az account show --query id -o tsv)
```

```
{
  "appId": "7654321a-1a23-567b-b789-0987b6543a21",
  "displayName": "azure-cli-2021-03-09-23-17-06",
  "password": "Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C",
  "tenant": "a1234567-b132-1234-1a11-1234a5678b90"
}
```

3. Set the required environment variables:

```
export AZURE_SUBSCRIPTION_ID="<id>"           # b1234567-abcd-11a1-a0a0-1234a5678b90
export AZURE_TENANT_ID="<tenant>"             # a1234567-b132-1234-1a11-1234a5678b90
export AZURE_CLIENT_ID="<appId>"             # 7654321a-1a23-567b-b789-0987b6543a21
export AZURE_CLIENT_SECRET="<password>"      # Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C
```

4. Base64 encode the same environment variables:

```
export AZURE_SUBSCRIPTION_ID_B64="$(echo -n "${AZURE_SUBSCRIPTION_ID}" | base64 | tr
-d '\n')"
```

```
export AZURE_TENANT_ID_B64="$(echo -n "${AZURE_TENANT_ID}" | base64 | tr -d '\n')"
```

```
export AZURE_CLIENT_ID_B64="$(echo -n "${AZURE_CLIENT_ID}" | base64 | tr -d '\n')"
```

```
export AZURE_CLIENT_SECRET_B64="$(echo -n "${AZURE_CLIENT_SECRET}" | base64 | tr -d '\n')"
```

When you completed, move on to the [Bootstrap](#) (see page 1161) section.

7.13.2 Azure using Konvoy Image Builder

This procedure describes how to use the [Konvoy Image Builder](#) (see page 1461) (KIB) to create a [Cluster API](#)⁸⁶⁰ compliant Azure Virtual Machine (VM) Image. The VM Image contains the base operating system you specify and all the necessary Kubernetes components. The Konvoy Image Builder uses variable `overrides` to specify the base image and container images to use in your new Azure VM image.



The default Azure image is not recommended for use in production. We suggest using KIB for Azure to build the image in order to take advantage of enhanced cluster operations. To explore more information on this topic refer to the [Azure Infrastructure](#) (see page 1159).

7.13.2.1 Prerequisites

Before you begin, you must:

- Download the [Konvoy Image Builder](#) (see page 1461) bundle for your version of DKP.
- Check the [Supported Infrastructure Operating Systems](#) (see page 69)
- Check the [Supported Kubernetes Version](#) (see page 1552) for your Provider.
- Create a working `Docker` setup.

7.13.2.2 Extract the KIB Bundle

Extract the bundle and `cd` into the extracted `konvoy-image-bundle-$VERSION_$OS` folder. The bundled version of `konvoy-image` contains an embedded `docker` image that contains all the requirements for building.



The `konvoy-image` binary and all supporting folders are also extracted. When extracted, `konvoy-image` bind mounts the current working directory (`${PWD}`) into the container to be used.

⁸⁶⁰ <https://cluster-api.sigs.k8s.io/>

7.13.2.3 Configure Azure Prerequisites

- If you have already followed the [Azure Prerequisites \(see page 1159\)](#) topic steps, then the environment variables needed by KIB ([AZURE_CLIENT_SECRET , AZURE_CLIENT_ID , AZURE_TENANT_ID , AZURE_SUBSCRIPTION_ID]) are set and do not need repeated if you are still working in the same window.

If you have not executed the [Azure Prerequisite \(see page 1159\)](#) steps, they are listed below.

1. Sign in to Azure:

```
az login
```

```
[
  {
    "cloudName": "AzureCloud",
    "homeTenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "id": "b1234567-abcd-11a1-a0a0-1234a5678b90",
    "isDefault": true,
    "managedByTenants": [],
    "name": "Mesosphere Developer Subscription",
    "state": "Enabled",
    "tenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "user": {
      "name": "user@azuresmesosphere.onmicrosoft.com",
      "type": "user"
    }
  }
]
```

2. Create an Azure Service Principal (SP) by running the following command:

If an SP with the name exists, this command will rotate the password.

```
az ad sp create-for-rbac --role contributor --name "$(whoami)-konvoy" --
scopes=/subscriptions/$(az account show --query id -o tsv) --query
"{ client_id: appId, client_secret: password, tenant_id: tenant }"
```

```
{
  "client_id": "7654321a-1a23-567b-b789-0987b6543a21",
  "client_secret": "Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C",
  "tenant_id": "a1234567-b132-1234-1a11-1234a5678b90"
}
```

3. Set the `AZURE_CLIENT_SECRET` environment variable:

```
export AZURE_CLIENT_SECRET="<azure_client_secret>" #
Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C
export AZURE_CLIENT_ID="<client_id>" # 7654321a-1a23-567b-
b789-0987b6543a21
export AZURE_TENANT_ID="<tenant_id>" # a1234567-b132-1234-1a11-12
34a5678b90
export AZURE_SUBSCRIPTION_ID="<subscription_id>" # b1234567-abcd-11a1-
a0a0-1234a5678b90
```

4. Ensure you have an [override file](#) (see page 1513) to [configure specific attributes](#)⁸⁶¹ of your Azure image.

7.13.2.4 Next Step:

[Azure Bootstrap](#) (see page 1161)

7.13.3 Azure Bootstrap

7.13.3.1 Prepare to deploy Kubernetes clusters

To create Kubernetes clusters, Konvoy uses [Cluster API](#)⁸⁶² (CAPI) controllers. These controllers run on a Kubernetes cluster. To get started, you need a *bootstrap* cluster. By default, Konvoy creates a bootstrap cluster for you in a Docker container using the Kubernetes-in-Docker ([KIND](#)⁸⁶³) tool.

7.13.3.2 Prerequisites

Before you begin, you must:

- Complete the steps in [Prerequisites](#) (see page 1159).
- Ensure the `dkp` binary can be found in your `$PATH`.

⁸⁶¹ <https://github.com/mesosphere/konvoy-image-builder/tree/7447074a6d910e71ad2e61fc4a12820d073ae5ae/images/azure>

⁸⁶² <https://cluster-api.sigs.k8s.io/>

⁸⁶³ <https://github.com/kubernetes-sigs/kind>

7.13.3.3 Bootstrap Cluster Lifecycle Services

1. If an HTTP proxy is required for the bootstrap cluster, set the local `http_proxy`, `https_proxy`, and `no_proxy` environment variables. They are copied into the bootstrap cluster.
2. Create a bootstrap cluster:

```
dkp create bootstrap --kubeconfig $HOME/.kube/config
```

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

Konvoy creates a bootstrap cluster using [KIND](#)⁸⁶⁴ as a library. Konvoy then deploys the following [Cluster API](#)⁸⁶⁵ providers on the cluster:

- [Core Provider](#)⁸⁶⁶
- [Azure Infrastructure Provider](#)⁸⁶⁷
- [kubeadm Bootstrap Provider](#)⁸⁶⁸
- [kubeadm ControlPlane Provider](#)⁸⁶⁹

Konvoy waits until the controller-manager and webhook deployments of these providers are ready. List these deployments using this command:

```
kubectl get --all-namespaces deployments -l=clusterctl.cluster.x-k8s.io
```

NAMESPACE	READY	UP-TO-DATE	AVAILABLE	NAME	AGE
capa-system	1/1	1	1	capa-controller-manager	69s
capi-kubeadm-bootstrap-system	1/1	1	1	capi-kubeadm-bootstrap-controller-manager	71s

⁸⁶⁴ <https://github.com/kubernetes-sigs/kind>

⁸⁶⁵ <https://cluster-api.sigs.k8s.io/>

⁸⁶⁶ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

⁸⁶⁷ <https://github.com/kubernetes-sigs/cluster-api-provider-azure>

⁸⁶⁸ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

⁸⁶⁹ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>

```

capi-kubeadm-control-plane-system      capi-kubeadm-control-plane-controller-
manager 1/1      1          1          70s
capi-system                             capi-controller-manager
1/1      1          1          73s
capp-system                             capp-controller-manager
1/1      1          1          66s
capv-system                             capv-controller-manager
1/1      1          1          65s
capz-system                             capz-controller-manager
1/1      1          1          67s
cert-manager                           cert-manager
1/1      1          1          16m
cert-manager                           cert-manager-cainjector
1/1      1          1          16m
cert-manager                           cert-manager-webhook
1/1      1          1          16m

```

7.13.3.4 (Optional) Create identity secret for Azure

If your bootstrap cluster resides on a Virtual machine inside Azure, create an identity secret that uses the cappz-controller:

```

export AZURE_CLUSTER_IDENTITY_SECRET_NAME="cluster-identity-secret"
export CLUSTER_IDENTITY_NAME="cluster-identity"
export AZURE_CLUSTER_IDENTITY_SECRET_NAMESPACE="default"

kubectl create secret generic "${AZURE_CLUSTER_IDENTITY_SECRET_NAME}" --from-
literal=clientSecret="${AZURE_CLIENT_SECRET}"

```

7.13.3.5 Next Step:

[Create a New Custom Azure Cluster \(see page 1163\)](#)

7.13.4 Create a New Custom Azure Cluster

7.13.4.1 Prerequisites

- Before you begin, make sure you have created a [Bootstrap \(see page 1164\)](#) cluster.

7.13.4.2 Name your cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=<azure-example>
```

7.13.4.3 Tips and Tricks

Below are a few ways to customize your setup which are optional. If you prefer to do a basic setup, skip Tips and Tricks and proceed to [Create a New Azure Cluster](#) (see page 1165) section.

Important to remember related to the options below: the `--compute-gallery-id` image will be in the format `--compute-gallery-id /subscriptions/<subscription id>/resourceGroups/<resource group name>/providers/Microsoft.Compute/galleries/<gallery name>/images/<image definition name>/versions/<version id>`.

- Option 1: To create a cluster name that is unique, use the following command. This creates a unique name every time you run it, so use it with forethought:

```
export CLUSTER_NAME=azure-example-$(LC_CTYPE=C tr -dc 'a-z0-9' </dev/urandom |
fold -w 5 | head -n1)
echo $CLUSTER_NAME
```

```
azure-example-pf4a3
```

- Option 2: To use a custom Azure Image when creating your cluster, you must create that Azure Image using [KIB](#) (see page 1477) first.

```
dkp create cluster azure --cluster-name=${CLUSTER_NAME} \
--compute-gallery-id "<Managed Image Shared Image Gallery Id>"
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

- Option 3: To create individual files with different smaller manifests for ease in editing, you can add the `--output-directory` flag. This will create multiple files in the specified directory which must already exist.

```
dkp create cluster azure --cluster-name=${CLUSTER_NAME} \
--compute-gallery-id "<Managed Image Shared Image Gallery Id>"
--dry-run \
--output=yaml \
--output-directory=<existing-directory>
```

- Option 4: To use a custom DNS on Azure, you need a DNS name in your control. Then create a DKP cluster using the standard method described below with the `--self-managed` flag. Once the resource group has been created, you can create your hosted zone with the command below:

```
az network dns zone create --resource-group "d2iq-professional-services" --name
```

You no longer need to create a cluster issuer. There are several documents that explain [custom DNS](#) (see page 1432) in the Kommander component.

- Option 5: To allow DKP to create a cluster with Marketplace based images such as for Rocky Linux, the following flags are available. If these fields were specified in the [override file](#) (see page 1513) during [image creation](#) (see page 1477), the flags must be used in cluster creation:
 - `--plan-offer`, `--plan-publisher` and `--plan-sku`

```
--plan-offer rockylinux-9
--plan-publisher erockyenterprisesoftwarefoundationinc1653071250513
--plan-sku rockylinux-9
```



If you see a similar error to "Creating a virtual machine from Marketplace image or a custom image sourced from a Marketplace image requires Plan information in the request." when creating a cluster, you must also set the following flags `--plan-offer`, `--plan-publisher`, `--plan-sku`. For example when creating a cluster with Rocky Linux VMs, add the following flags to your `dkp create cluster azure` command:

- `--plan-offer`, `--plan-publisher` and `--plan-sku`



For more information regarding this flag or others, please refer to the CLI for the [dkp create cluster](#) (see page 1662) section of the documentation and select your provider.

7.13.4.4 Create a new Azure Kubernetes cluster

Availability zones (AZs) are isolated locations within data center regions from which public cloud services originate and operate. Because all the nodes in a node pool are deployed in a single Availability Zone, you

may wish to create additional node pools to ensure your cluster has nodes deployed in multiple Availability Zones.

- By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command. See Microsoft's documentation for more information on [Availability Options for Azure VM](#)⁸⁷⁰.



If you are using Azure as a pre-provisioned environment: DKP uses

`localvolumeprovisioner` as the [default storage provider](#) (see page 91) if creating a [pre-provisioned Azure cluster](#) (see page 1227). However, `localvolumeprovisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)⁸⁷¹ compatible storage that is suitable for production.

You can choose from any of the [storage options](#)⁸⁷² available for Kubernetes. To disable the default that Konvoy deploys, set the default StorageClass `localvolumeprovisioner` as non-default. Then set your newly created StorageClass to be the default by following the commands in the Kubernetes documentation called [Changing the Default Storage Class](#)⁸⁷³.

- Ensure your [subnets](#) (see page 1049) do not overlap with your host subnet because they cannot be changed after cluster creation. If you need to change the kubernetes subnets, you must do this at cluster creation. The default subnets used in DKP are:

```
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
        - 192.168.0.0/16
    services:
      cidrBlocks:
        - 10.96.0.0/12
```

The [below cluster create directions](#) instead describes how to create a cluster using Azure as the infrastructure provider provisioning clusters, which uses Azure Disks Container Storage Interface as the default StorageClass.

⁸⁷⁰ <https://docs.microsoft.com/en-us/azure/virtual-machines/manage-availability#configure-multiple-virtual-machines-in-an-availability-set-for-redundancy>

⁸⁷¹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>


⁸⁷² <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁸⁷³ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

1. Generate the Kubernetes cluster objects. The following example shows a common configuration. See [dkp create cluster azure](#) (see page 1667) reference for the full list of cluster creation options.

```
dkp create cluster azure --cluster-name=${CLUSTER_NAME} \
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml
```

Generating cluster resources

 Refer to the Cluster Creation Customization Choices section for more information on how to use optional flags such as the `--output-directory` flag.

2. (Optional) To configure the Control Plane and Worker nodes to use an HTTP proxy:

```
export CONTROL_PLANE_HTTP_PROXY=http://example.org:8080
export CONTROL_PLANE_HTTPS_PROXY=http://example.org:8080
export
CONTROL_PLANE_NO_PROXY="example.org,example.com,example.net,localhost,127.0.0.1
,10.96.0.0/12,192.168.0.0/16,kubernetes,kubernetes.default,kubernetes.default.s
vc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.sv
c.cluster,.svc.cluster.local,169.254.169.254,.cloudapp.azure.com"

export WORKER_HTTP_PROXY=http://example.org:8080
export WORKER_HTTPS_PROXY=http://example.org:8080
export
WORKER_NO_PROXY="example.org,example.com,example.net,localhost,127.0.0.1,10.96.
0.0/12,192.168.0.0/16,kubernetes,kubernetes.default,kubernetes.default.svc,kube
rnetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.clust
er,.svc.cluster.local,169.254.169.254,.cloudapp.azure.com"
```

- Replace `example.org,example.com,example.net` with your internal addresses
- `localhost` and `127.0.0.1` addresses should not use the proxy
- `10.96.0.0/12` is the default Kubernetes service subnet
- `192.168.0.0/16` is the default Kubernetes pod subnet
- `kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local` is the internal Kubernetes kube-apiserver service
- `.svc,.svc.cluster,.svc.cluster.local` is the internal Kubernetes services

- `169.254.169.254` is the Azure metadata server
 - `.cloudapp.azure.com` is for the worker nodes to allow them to communicate directly to the kube-apiserver load balancer
3. (Optional) Create a Kubernetes cluster with HTTP proxy configured. This step assumes you did not already create a cluster in the previous steps:

```
dkp create cluster azure --cluster-name=${CLUSTER_NAME} \
--control-plane-http-proxy="${CONTROL_PLANE_HTTP_PROXY}" \
--control-plane-https-proxy="${CONTROL_PLANE_HTTPS_PROXY}" \
--control-plane-no-proxy="${CONTROL_PLANE_NO_PROXY}" \
--worker-http-proxy="${WORKER_HTTP_PROXY}" \
--worker-https-proxy="${WORKER_HTTPS_PROXY}" \
--worker-no-proxy="${WORKER_NO_PROXY}" \
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml
```

4. Inspect or edit the cluster objects:

NOTE: Familiarize yourself with Cluster API before editing the cluster objects as edits can prevent the cluster from deploying successfully.

The objects are [Custom Resources](#)⁸⁷⁴ defined by Cluster API components, and they belong in three different categories:

- Cluster**
A *Cluster* object has references to the infrastructure-specific and control plane objects. Because this is an Azure cluster, there is an *AzureCluster* object that describes the infrastructure-specific cluster properties. Here, this means the Azure region, the VPC ID, subnet IDs, and security group rules required by the Pod network implementation.
- Control Plane**
A *KubeadmControlPlane* object describes the control plane, which is the group of machines that run the Kubernetes control plane components, which include the etcd distributed database, the API server, the core controllers, and the scheduler. The object describes the configuration for these components. The object also has a reference to an infrastructure-specific object that describes the properties of all control plane machines. Here, it references an *AzureMachineTemplate* object, which describes the instance type, the type of disk used, and the size of the disk, among other properties.
- Node Pool**
A Node Pool is a collection of machines with identical properties. For example, a cluster might have one Node Pool with large memory capacity, another Node Pool with GPU support. Each Node Pool is described by three objects: The *MachinePool* references an object that describes the configuration of Kubernetes components (for example, kubelet) deployed on each node pool machine, and an infrastructure-specific object that describes the properties of

⁸⁷⁴ <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

all node pool machines. Here, it references a *KubeadmConfigTemplate*, and an *AzureMachineTemplate* object, which describes the instance type, the type of disk used, the size of the disk, among other properties.

For in-depth documentation about the objects, read [Concepts](#)⁸⁷⁵ in the Cluster API Book.

5. Modify Control Plane Audit logs settings using the information contained in the page [Configuring the Control Plane](#) (see page 1538).
6. Create the cluster from the objects. A warning will appear in the console if the resource already exists and will require you to remove the resource or update your YAML.

```
kubectl create -f ${CLUSTER_NAME}.yaml
```

NOTE: If you used the `--output-directory` flag in your `dkp create .. --dry-run` step above, create the cluster from the objects you created by specifying the directory:

```
kubectl create -f <existing-directory>/
```

Output will be similar to output below:

```
cluster.cluster.x-k8s.io/azure-example created
azurecluster.infrastructure.cluster.x-k8s.io/azure-example created
kubeadmcontrolplane.controlplane.cluster.x-k8s.io/azure-example-control-plane
created
azuremachinetemplate.infrastructure.cluster.x-k8s.io/azure-example-control-
plane created
secret/azure-example-etcd-encryption-config created
machinedeployment.cluster.x-k8s.io/azure-example-md-0 created
azuremachinetemplate.infrastructure.cluster.x-k8s.io/azure-example-md-0 created
kubeadmconfigtemplate.bootstrap.cluster.x-k8s.io/azure-example-md-0 created
clusterresourceset.addons.cluster.x-k8s.io/calico-cni-installation-azure-
example created
configmap/calico-cni-installation-azure-example created
configmap/tigera-operator-azure-example created
clusterresourceset.addons.cluster.x-k8s.io/azure-disk-csi-azure-example created
configmap/azure-disk-csi-azure-example created
clusterresourceset.addons.cluster.x-k8s.io/cluster-autoscaler-azure-example
created
configmap/cluster-autoscaler-azure-example created
clusterresourceset.addons.cluster.x-k8s.io/node-feature-discovery-azure-example
created
configmap/node-feature-discovery-azure-example created
clusterresourceset.addons.cluster.x-k8s.io/nvidia-feature-discovery-azure-
example created
configmap/nvidia-feature-discovery-azure-example created
```

⁸⁷⁵ <https://cluster-api.sigs.k8s.io/user/concepts.html>

7. Wait for the cluster control-plane to be ready:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

```
cluster.cluster.x-k8s.io/azure-example condition met
```

8. After the objects are created on the API server, the Cluster API controllers reconcile them. They create infrastructure and machines. As they progress, they update the Status of each object. Konvoy provides a command to describe the current status of the cluster:

```
dkp describe cluster -c ${CLUSTER_NAME}
```

```
NAME                                                                 READY
SEVERITY REASON SINCE MESSAGE
Cluster/azure-example                                             True
3m4s
├─ClusterInfrastructure - AzureCluster/azure-example             True
8m26s
├─ControlPlane - KubeadmControlPlane/azure-example-control-plane True
3m4s
| └─Machine/azure-example-control-plane-l8j9r                    True
3m9s
| └─Machine/azure-example-control-plane-slprd                    True
7m17s
| └─Machine/azure-example-control-plane-xhxxg                    True
5m9s
└─Workers
  └─MachineDeployment/azure-example-md-0                          True
4m31s
    └─Machine/azure-example-md-0-d67567c8b-2674r                 True
5m19s
    └─Machine/azure-example-md-0-d67567c8b-mbmhk                 True
5m17s
    └─Machine/azure-example-md-0-d67567c8b-pzg8k                 True
5m17s
    └─Machine/azure-example-md-0-d67567c8b-z8km9                 True
5m17s
```

9. As they progress, the controllers also create Events. List the Events using this command:

```
kubectl get events | grep ${CLUSTER_NAME}
```

For brevity, the example uses `grep`. It is also possible to use separate commands to get Events for specific objects. For example, `kubectl get events --field-selector involvedObject.kind="AzureCluster"` and `kubectl get events --field-selector involvedObject.kind="AzureMachine"`.

```

15m          Normal    AzureClusterObjectNotFound          azurecluster
AzureCluster object default/azure-example not found
15m          Normal    AzureManagedControlPlaneObjectNotFound
azuremanagedcontrolplane          AzureManagedControlPlane
object default/azure-example not found
15m          Normal    AzureClusterObjectNotFound          azurecluster
AzureCluster.infrastructure.cluster.x-k8s.io "azure-example" not found
8m22s        Normal    SuccessfulSetNodeRef                machine/
azure-example-control-plane-bmc9b    azure-example-control-plane-fdvm
10m          Normal    Machine controller dependency not yet met azuremachine/
azure-example-control-plane-fdvm      Machine Controller has not yet set
OwnerRef
12m          Normal    SuccessfulSetNodeRef                machine/
azure-example-control-plane-msftd      azure-example-control-plane-z9q45
10m          Normal    SuccessfulSetNodeRef                machine/
azure-example-control-plane-nrvff      azure-example-control-plane-vmqwx
12m          Normal    Machine controller dependency not yet met azuremachine/
azure-example-control-plane-vmqwx      Machine Controller has not yet set
OwnerRef
14m          Normal    Machine controller dependency not yet met azuremachine/
azure-example-control-plane-z9q45      Machine Controller has not yet set
OwnerRef
14m          Warning   VMIdentityNone
azuremachinetemplate/azure-example-control-plane You are using Service
Principal authentication for Cloud Provider Azure which is less secure than
Managed Identity. Your Service Principal credentials will be written to a file
on the disk of each VM in order to be accessible by Cloud Provider. To learn
more, see https://capz.sigs.k8s.io/topics/identities-use-cases.html#azure-host-identity
12m          Warning   ControlPlaneUnhealthy
kubeadmcontrolplane/azure-example-control-plane Waiting for control plane to
pass preflight checks to continue reconciliation: [machine azure-example-
control-plane-msftd does not have APIServerPodHealthy condition, machine azure-
example-control-plane-msftd does not have ControllerManagerPodHealthy
condition, machine azure-example-control-plane-msftd does not have
SchedulerPodHealthy condition, machine azure-example-control-plane-msftd does
not have EtcdPodHealthy condition, machine azure-example-control-plane-msftd
does not have EtcdMemberHealthy condition]
11m          Warning   ControlPlaneUnhealthy
kubeadmcontrolplane/azure-example-control-plane Waiting for control plane to
pass preflight checks to continue reconciliation: [machine azure-example-
control-plane-nrvff does not have APIServerPodHealthy condition, machine azure-
example-control-plane-nrvff does not have ControllerManagerPodHealthy
condition, machine azure-example-control-plane-nrvff does not have
SchedulerPodHealthy condition, machine azure-example-control-plane-nrvff does

```

```

not have EtcdPodHealthy condition, machine azure-example-control-plane-nrvff
does not have EtcdMemberHealthy condition]
9m52s      Normal    SuccessfulSetNodeRef      machine/
azure-example-md-0-84bd8b5f5b-b8cnq      azure-example-md-0-bsc82
9m53s      Normal    SuccessfulSetNodeRef      machine/
azure-example-md-0-84bd8b5f5b-j8ldg      azure-example-md-0-mjcbn
9m52s      Normal    SuccessfulSetNodeRef      machine/
azure-example-md-0-84bd8b5f5b-lx89f      azure-example-md-0-pmq8f
10m        Normal    SuccessfulSetNodeRef      machine/
azure-example-md-0-84bd8b5f5b-pcv7q      azure-example-md-0-vzprf
15m        Normal    SuccessfulCreate          machineset/
azure-example-md-0-84bd8b5f5b      Created machine "azure-example-
md-0-84bd8b5f5b-j8ldg"
15m        Normal    SuccessfulCreate          machineset/
azure-example-md-0-84bd8b5f5b      Created machine "azure-example-
md-0-84bd8b5f5b-lx89f"
15m        Normal    SuccessfulCreate          machineset/
azure-example-md-0-84bd8b5f5b      Created machine "azure-example-
md-0-84bd8b5f5b-pcv7q"
15m        Normal    SuccessfulCreate          machineset/
azure-example-md-0-84bd8b5f5b      Created machine "azure-example-
md-0-84bd8b5f5b-b8cnq"
15m        Normal    Machine controller dependency not yet met  azuremachine/
azure-example-md-0-bsc82      Machine Controller has not yet set
OwnerRef
15m        Normal    Machine controller dependency not yet met  azuremachine/
azure-example-md-0-mjcbn      Machine Controller has not yet set
OwnerRef
15m        Normal    Machine controller dependency not yet met  azuremachine/
azure-example-md-0-pmq8f      Machine Controller has not yet set
OwnerRef

```



If changing the [Calico encapsulation](#) (see page 1231), D2iQ recommends changing it after cluster creation, but before production.

7.13.4.5 Known Limitations



Be aware of these limitations in the current release of Konvoy.

The Konvoy version used to create a bootstrap cluster must match the Konvoy version used to create a workload cluster.

- Konvoy supports deploying one workload cluster.
- Konvoy generates a set of objects for one Node Pool.
- Konvoy does not validate edits to cluster objects.

Next, you can [Explore the New Cluster](#) (see page 1173) or [Make it Self-managed](#) (see page 1177).

7.13.5 Explore new Azure Cluster

7.13.5.1 Learn to interact with your Kubernetes cluster

This guide explains how to use the command line to interact with your newly deployed Kubernetes cluster. Before you start, make sure you have created a workload cluster, as described in [Create a New Cluster](#) (see page 1166).

7.13.5.2 Explore the new Kubernetes cluster

1. Get a kubeconfig file for the workload cluster:

When the workload cluster is created, the cluster lifecycle services generate a kubeconfig file for the workload cluster, and write it to a *Secret*. The kubeconfig file is scoped to the cluster administrator.

Get the kubeconfig from the *Secret*, and write it to a file, using this command:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

2. List the Nodes using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

NAME	STATUS	ROLES	AGE	
VERSION				
azure-example-control-plane-7ffnl v1.26.14	Ready	control-plane,master	6m18s	
azure-example-control-plane-l4bv8 v1.26.14	Ready	control-plane,master	14m	
azure-example-control-plane-n4g4l v1.26.14	Ready	control-plane,master	18m	
azure-example-md-0-mpctb 6.14	Ready	<none>	15m	v1.2
azure-example-md-0-qglp9 6.14	Ready	<none>	15m	v1.2


```

azure-example-md-0-sgrd6      Ready    <none>      16m    v1.2
6.14
azure-example-md-0-wzbnk1    Ready    <none>      16m    v1.2
6.14

```

NOTE: It may take a few minutes for the Status to move to `Ready` while the Pod network is deployed. The Nodes' Status should change to Ready soon after the `calico-node` DaemonSet Pods are Ready.

3. List the Pods using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get --all-namespaces pods
```

```

NAMESPACE          NAME
READY  STATUS  RESTARTS  AGE
calico-system      calico-kube-controllers-57fbd7bd59-v4tss
1/1    Running  0         19m
calico-system      calico-node-59llv
1/1    Running  0         17m
calico-system      calico-node-7t7wj
1/1    Running  0         16m
calico-system      calico-node-pf8q8
1/1    Running  0         17m
calico-system      calico-node-sh2b7
1/1    Running  0         8m17s
calico-system      calico-node-tmxl5
1/1    Running  0         19m
calico-system      calico-node-vt5fh
1/1    Running  0         18m
calico-system      calico-node-whfs8
1/1    Running  0         18m
calico-system      calico-typha-797c9666d5-5w99r
1/1    Running  0         19m
calico-system      calico-typha-797c9666d5-hj6mj
1/1    Running  0         18m
calico-system      calico-typha-797c9666d5-s7rc6
1/1    Running  0         17m
capa-system        capa-controller-manager-74fffb5676-ch6xd
1/1    Running  0         11m

```

capi-kubeadm-bootstrap-system manager-867759cc67-vg4lh	1/1	Running	0	15m	capi-kubeadm-bootstrap-controller-
capi-kubeadm-control-plane-system manager-5df55579c4-pc8x9	1/1	Running	1 (11m ago)	15m	capi-kubeadm-control-plane-controller-
capi-system					capi-controller-manager-79cc58bf5f-xsp9t
1/1	Running	0		15m	
cappp-system					cappp-controller-manager-85b5c77497-8ss8r
1/1	Running	0		14m	
capv-system					capv-controller-manager-7bf4d8b66-6x2mx
1/1	Running	0		14m	
capz-system					capz-controller-manager-5d4c6468bf-wfhcc
1/1	Running	0		14m	
capz-system					capz-nmi-2cbrg
1/1	Running	0		14m	
capz-system					capz-nmi-8dllm
1/1	Running	0		14m	
capz-system					capz-nmi-95dfk
1/1	Running	0		14m	
capz-system					capz-nmi-rtnd4
1/1	Running	0		14m	
cert-manager					cert-manager-848f547974-gjc5p
1/1	Running	1 (10m ago)		15m	
cert-manager					cert-manager-cainjector-54f4cc6b5-rnh4f
1/1	Running	0		15m	
cert-manager					cert-manager-webhook-7c9588c76-rn2sd
1/1	Running	0		15m	
kube-system					cluster-autoscaler-68c759fbf6-6vg5r
1/1	Running	1 (11m ago)		20m	
kube-system					coredns-78fcd69978-6gx44
1/1	Running	0		20m	
kube-system					coredns-78fcd69978-gr5q7
1/1	Running	0		20m	
kube-system					csi-azuredisk-controller-c8fb44c8b-jhmfz
6/6	Running	5 (11m ago)		20m	
kube-system					csi-azuredisk-controller-c8fb44c8b-lpbbs
6/6	Running	0		20m	
kube-system					csi-azuredisk-node-2g7vw
3/3	Running	0		8m17s	
kube-system					csi-azuredisk-node-6rdqc
3/3	Running	0		18m	
kube-system					csi-azuredisk-node-99c6q
3/3	Running	0		17m	
kube-system					csi-azuredisk-node-9b4ms
3/3	Running	0		17m	
kube-system					csi-azuredisk-node-mz5pr
3/3	Running	0		18m	
kube-system					csi-azuredisk-node-r2t99
3/3	Running	0		16m	
kube-system					csi-azuredisk-node-t7gfs
3/3	Running	0		20m	
kube-system					etcd-azure-example-control-plane-7ffnl
1/1	Running	0		8m15s	

```

kube-system          etcd-azure-example-control-plane-l4bv8
1/1      Running    0          16m
kube-system          etcd-azure-example-control-plane-n4g4l
1/1      Running    0          19m
kube-system          kube-apiserver-azure-example-control-plane-7ffn1
1/1      Running    0          8m16s
kube-system          kube-apiserver-azure-example-control-plane-l4bv8
1/1      Running    0          16m
kube-system          kube-apiserver-azure-example-control-plane-n4g4l
1/1      Running    0          19m
kube-system          kube-controller-manager-azure-example-control-
plane-7ffn1          1/1      Running    0          8m17s
kube-system          kube-controller-manager-azure-example-control-
plane-l4bv8          1/1      Running    0          16m
kube-system          kube-controller-manager-azure-example-control-
plane-n4g4l          1/1      Running    1 (17m ago) 19m
kube-system          kube-proxy-82zd1
1/1      Running    0          8m17s
kube-system          kube-proxy-fd9f9
1/1      Running    0          18m
kube-system          kube-proxy-l6lgc
1/1      Running    0          17m
kube-system          kube-proxy-lzswH
1/1      Running    0          16m
kube-system          kube-proxy-ndfmt
1/1      Running    0          20m
kube-system          kube-proxy-nxlp9
1/1      Running    0          18m
kube-system          kube-proxy-v9sxp
1/1      Running    0          17m
kube-system          kube-scheduler-azure-example-control-plane-7ffn1
1/1      Running    0          8m16s
kube-system          kube-scheduler-azure-example-control-plane-l4bv8
1/1      Running    0          16m
kube-system          kube-scheduler-azure-example-control-plane-n4g4l
1/1      Running    1 (17m ago) 19m
node-feature-discovery node-feature-discovery-master-84c67dcbb6-d2gm7
1/1      Running    0          20m
node-feature-discovery node-feature-discovery-worker-drgf6
1/1      Running    0          17m
node-feature-discovery node-feature-discovery-worker-hcz6k
1/1      Running    0          17m
node-feature-discovery node-feature-discovery-worker-pgbcD
1/1      Running    0          16m
node-feature-discovery node-feature-discovery-worker-vhj96
1/1      Running    0          16m
tigera-operator      tigera-operator-d499f5c8f-jnj8b
1/1      Running    1 (18m ago) 19m

```

7.13.6 Azure Make new Cluster Self-Managed

Konvoy deploys all cluster lifecycle services to a bootstrap cluster, which then deploys a workload cluster. When the workload cluster is ready, move the cluster lifecycle services to the workload cluster, which makes the workload cluster self-managed. This section describes how to make a workload cluster self-managed.

This page contains instructions on how to make your cluster self-managed. This is necessary if there is only one cluster in your environment, or if this cluster should become the Management cluster in a multi-cluster environment.



If you already have a self-managed or Management cluster in your environment, skip this page.

7.13.6.1 Make the new Kubernetes cluster manage itself

1. Deploy cluster lifecycle services on the workload cluster:

```
dkp create capi-components --kubeconfig ${CLUSTER_NAME}.conf
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

The output is similar to:

```
✓ Initializing new CAPI components
```

2. Move the Cluster API objects from the bootstrap to the workload cluster:

The cluster lifecycle services on the workload cluster are ready, but the workload cluster configuration is on the bootstrap cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the bootstrap to the workload cluster. This process is also called a [Pivot](#)⁸⁷⁶.

```
dkp move capi-resources --to-kubeconfig ${CLUSTER_NAME}.conf
```

⁸⁷⁶ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

✓ Moving cluster resources

You can now view resources in the moved cluster by using the `--kubeconfig` flag with `kubectl`. For example: `kubectl --kubeconfig=azure-example.conf get nodes`

- To ensure only one set of cluster lifecycle services manages the workload cluster, Konvoy first pauses reconciliation of the objects on the bootstrap cluster, then creates the objects on the workload cluster. As Konvoy copies the objects, the cluster lifecycle services on the workload cluster reconcile the objects. The workload cluster becomes self-managed after Konvoy creates all the objects. If it fails, the `move` command can be safely retried.

3. Wait for the cluster control-plane to be ready:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf wait --for=condition=ControlPlaneReady
"clusters/${CLUSTER_NAME}" --timeout=20m
```

```
cluster.cluster.x-k8s.io/azure-example condition met
```

4. Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

- After moving the cluster lifecycle services to the workload cluster, remember to use Konvoy with the workload cluster kubeconfig.

```
dkp describe cluster --kubeconfig ${CLUSTER_NAME}.conf -c ${CLUSTER_NAME}
```

NAME	REASON	SINCE	MESSAGE	READY	SEVERITY
Cluster/azure-example		55s		True	
ClusterInfrastructure - AzureCluster/azure-example		67s		True	
ControlPlane - KubeadmControlPlane/azure-example-control-plane		55s		True	
Machine/azure-example-control-plane-67f47		58s		True	

```

| |─Machine/azure-example-control-plane-7pllh           True
65s
| |─Machine/azure-example-control-plane-jtfgv           True
65s
└─Workers
  └─MachineDeployment/azure-example-md-0                 True
67s
    └─Machine/azure-example-md-0-f9cb9c79b-6nsb9       True
59s
    └─Machine/azure-example-md-0-f9cb9c79b-jxwl6       True
58s
    └─Machine/azure-example-md-0-f9cb9c79b-ktg7z       True
59s
    └─Machine/azure-example-md-0-f9cb9c79b-nxcm2       True
66s

```

5. Remove the bootstrap cluster, as the workload cluster is now self-managed:

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

```
✓ Deleting bootstrap cluster
```

7.13.6.2 Known Limitations

- Konvoy supports moving only one set of cluster objects from the bootstrap cluster to the workload cluster, or vice-versa.
- Konvoy only supports moving all namespaces in the cluster; Konvoy does not support migration of individual namespaces.

7.13.7 Azure Certificate Renewal

7.13.7.1 Configure Automated Renewal for Managed Kubernetes PKI Certificates

During cluster creation, Kubernetes establishes a Public Key Infrastructure (PKI) for generating the TLS certificates needed for securing cluster communication for various components such as `etcd`, `kubernetes-apiserver` and `kube-proxy`. The certificates created by these components have a default expiration of one year and are renewed when an administrator updates the cluster.

Kubernetes provides a facility to renew all certificates automatically during control plane updates. For administrators who need long-running clusters or clusters that are not upgraded often, `dkp` provides automated certificate renewal, without a cluster upgrade.

7.13.7.2 Requirements

This feature requires Python 3.5 or greater to be installed on all control plane hosts.

7.13.7.3 Prerequisites

Prerequisite:

- Complete the [bootstrap Cluster Lifecycle](#) (see page 1164) topic.

7.13.7.3.1 Create a cluster with automated certificate renewal

To enable the automated certificate renewal, create a Konvoy cluster using the `certificate-renew-interval` flag:

```
dkp create cluster azure --certificate-renew-interval=60 --cluster-name=long-running
```

The `certificate-renew-interval` is the number of days after which Kubernetes-managed PKI certificates will be renewed. For example, a `certificate-renew-interval` value of 60 means the certificates will be renewed every 60 days.

7.13.7.3.2 Technical details

The following manifests are modified on the control plane hosts, and are located at `/etc/kubernetes/manifests`. Modifications to these files requires SUDO access.

```
kube-controller-manager.yaml
kube-apiserver.yaml
kube-scheduler.yaml
kube-proxy.yaml
```

The following annotation indicates the time each component was reset:

```
metadata:
  annotations:
    konvoy.d2iq.io/restartedAt: $(date +%s)
```

This only occurs when the PKI certificates are older than the interval given at cluster creation time. This is activated by a `systemd` timer called `renew-certs.timer` that triggers an associated `systemd` service called `renew-certs.service` that runs on all of the control plane hosts.

7.13.7.3.3 Debugging

To debug the automatic certificate renewal feature, a cluster administrator can look at several different components to see if the certificates were renewed. For example, an administrator might start with a look at the control plane pod definition to check the last reset time. To determine if a scheduler pod was properly reset, run the command:

```
kubectl get pod -n kube-system kube-scheduler-ip-10-0-xx-xx.us-west-2.compute.interna  
l -o yaml
```

The output of the command will be similar to the following:

```
apiVersion: v1  
kind: Pod  
metadata:  
  annotations:  
    konvoy.d2iq.io/restartedAt: "1626124940.735733"
```

Administrators who want more details on the execution of the `systemd` service can use `ssh` to connect to the control plane hosts, and then use the `systemctl` and `journalctl` commands that follow to help diagnose potential issues.

To check the status of the timers, when they last ran, and when they are scheduled to run next, use the command:

```
systemctl list-timers
```

To check the status of the `renew-certs` service, use the command:

```
systemctl status renew-certs
```

To get the logs of the last run of the service, use the command:

```
journalctl logs -u renew-certs
```


7.13.8 Azure Replace a Node

7.13.8.1 Replace a worker node

7.13.8.2 Prerequisites

Before you begin, you must:

- [Create a workload cluster \(see page 1166\)](#).
- [Make the new cluster self-managed \(see page 1180\)](#).

7.13.8.3 Replace a worker node

In certain situations, you may want to delete a worker node and have [Cluster API](#)⁸⁷⁷ replace it with a newly provisioned machine.

1. Identify the name of the node to delete.

List the nodes:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf get nodes
```

The output from this command resembles the following:

NAME	STATUS	ROLES	AGE	VERSION
azure-example-control-plane-ckwm4	Ready	control-plane,master	35m	v1.26.14
azure-example-control-plane-d4fdf	Ready	control-plane,master	31m	v1.26.14
azure-example-control-plane-qrvm9	Ready	control-plane,master	33m	v1.26.14
azure-example-md-0-4w7gq	Ready	<none>	33m	v1.26.14
azure-example-md-0-6gb9k	Ready	<none>	33m	v1.26.14
azure-example-md-0-p2n8c	Ready	<none>	11m	v1.26.14
azure-example-md-0-s5zbh	Ready	<none>	33m	v1.26.14

⁸⁷⁷ <https://cluster-api.sigs.k8s.io/>

- Export a variable with the node name to use in the next steps:

This example uses the name `azure-example-control-plane-ckwm4`.

```
export NAME_NODE_TO_DELETE=<azure-example-control-plane-ckwm4>
```

- Delete the Machine resource

```
export NAME_MACHINE_TO_DELETE=$(kubectl --kubeconfig ${CLUSTER_NAME}.conf get
machine -ojsonpath="{.items[?
(@.status.nodeRef.name==\"$NAME_NODE_TO_DELETE\")].metadata.name}")
kubectl --kubeconfig ${CLUSTER_NAME}.conf delete machine
"$NAME_MACHINE_TO_DELETE"
```

```
machine.cluster.x-k8s.io "azure-example-control-plane-slprd" deleted
```

The command will not return immediately. It will return once the Machine resource has been deleted.

A few minutes after the Machine resource is deleted, the corresponding Node resource is also deleted.

- Observe that the Machine resource is being replaced using this command:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf get machinedeployment
```

NAME	AGE	CLUSTER	REPLICAS	READY	UPDATED	UNAVAILABLE
azure-example-md-0		azure-example	4	3	4	1
ScalingUp	7m30s	v1.26.14				
long-running-md-0		long-running	4	4	4	0
Running	7m28s	v1.26.14				

In this example, there are two replicas, but only 1 is ready. One replica is unavailable, and the `ScalingUp` phase means a new Machine is being created.

- Identify the replacement Machine using this command:

```
export NAME_NEW_MACHINE=$(kubectl --kubeconfig ${CLUSTER_NAME}.conf get
machines \
  -l=cluster.x-k8s.io/deployment-name=${CLUSTER_NAME}-md-0 \
  -ojsonpath='{.items[?(@.status.phase=="Running")].metadata.name}{"\n"}')
echo $NAME_NEW_MACHINE
```

```
azure-example-md-0-d67567c8b-2674r azure-example-md-0-d67567c8b-n276j azure-
example-md-0-d67567c8b-pzg8k azure-example-md-0-d67567c8b-z8km9
```

If the output is empty, the new Machine has probably exited the **Provisioning** phase and entered the **Running** phase.

- Identify the replacement Node using this command:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf get nodes
```

NAME	STATUS	ROLES	AGE	
VERSION				
azure-example-control-plane-d4fdf	Ready	control-plane,master	43m	
v1.26.14				
azure-example-control-plane-qrvm9	Ready	control-plane,master	45m	
v1.26.14				
azure-example-control-plane-tz56m	Ready	control-plane,master	8m22s	
v1.26.14				
azure-example-md-0-4w7gq	Ready	<none>	45m	v1.2
6.14				
azure-example-md-0-6gb9k	Ready	<none>	45m	v1.2
6.14				
azure-example-md-0-p2n8c	Ready	<none>	22m	v1.2
6.14				
azure-example-md-0-s5zbh	Ready	<none>	45m	v1.2
6.14				

If the output is empty, the Node resource is not yet available, or does not yet have the expected annotation. Wait a few minutes, then repeat the command.

7.13.9 Azure Delete Cluster

7.13.9.1 Delete the Kubernetes cluster and clean up your environment

7.13.9.2 Prepare to Delete a Workload Cluster

NOTE: A self-managed workload cluster cannot delete itself. If your workload cluster is self-managed, you must create a bootstrap cluster and move the cluster lifecycle services to the bootstrap cluster before deleting the workload cluster.

If you did not make your workload cluster self-managed, as described in [Make New Cluster Self-Managed](#) (see [page 1180](#)), proceed to [Delete the workload cluster](#) (see [page 1187](#)) section below.

1. Create a bootstrap cluster:

The bootstrap cluster will host the Cluster API controllers that reconcile the cluster objects marked for deletion:

NOTE: To avoid using the wrong kubeconfig, the following steps use explicit kubeconfig paths and contexts.

```
dkp create bootstrap --kubeconfig $HOME/.kube/config
```

- ✓ Creating a bootstrap cluster
- ✓ Initializing **new** CAPI components

2. Move the Cluster API objects from the workload to the bootstrap cluster: The cluster lifecycle services on the bootstrap cluster are ready, but the workload cluster configuration is on the workload cluster. The `move`

command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the workload to the bootstrap cluster. This process is also called a [Pivot](#)⁸⁷⁸.

```
dkp move capi-resources \
  --from-kubeconfig ${CLUSTER_NAME}.conf \
  --from-context ${CLUSTER_NAME}-admin@${CLUSTER_NAME} \
  --to-kubeconfig $HOME/.kube/config \
  --to-context kind-konvoy-capi-bootstrapper
```

✓ Moving cluster resources

3. Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

```
dkp describe cluster --kubeconfig $HOME/.kube/config -c ${CLUSTER_NAME}
```

NAME	READY	SEVERITY
REASON SINCE MESSAGE		
Cluster/azure-example	True	
15s		
└─ClusterInfrastructure - AzureCluster/azure-example	True	
29s		
└─ControlPlane - KubeadmControlPlane/azure-example-control-plane	True	
15s		
└─Machine/azure-example-control-plane-gvj5d	True	
22s		
└─Machine/azure-example-control-plane-l8j9r	True	
23s		
└─Machine/azure-example-control-plane-xhxxg	True	
23s		
└─Workers		
└─MachineDeployment/azure-example-md-0	True	
35s		
└─Machine/azure-example-md-0-d67567c8b-2674r	True	
24s		
└─Machine/azure-example-md-0-d67567c8b-n276j	True	
25s		
└─Machine/azure-example-md-0-d67567c8b-pzg8k	True	
23s		
└─Machine/azure-example-md-0-d67567c8b-z8km9	True	
24s		

⁸⁷⁸ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

- **NOTE:** After moving the cluster lifecycle services to the workload cluster, remember to use `dkp` with the workload cluster `kubeconfig`.

1. Wait for the cluster control-plane to be ready:

```
kubectl --kubeconfig $HOME/.kube/config wait --for=condition=controlplaneready
"clusters/${CLUSTER_NAME}" --timeout=60m
```

```
cluster.cluster.x-k8s.io/azure-example condition met
```

- Persistent Volumes (PVs) are not deleted automatically by design in order to preserve your data. However, they take up storage space if not deleted. You must delete PVs manually. Information for backup of a cluster and PVs is on the page in documentation called [Back up your Cluster's Applications and Persistent Volumes \(see page 835\)](#).

7.13.9.3 Delete the Workload Cluster

1. Make sure your Azure credentials are up to date. Refresh the credentials using this command:

```
dkp update bootstrap credentials azure --kubeconfig $HOME/.kube/config
```

2. To delete a cluster, you would use `dkp delete cluster` and pass in the name of the cluster you are trying to delete with `--cluster-name` flag. You would use `kubectl get clusters` to get those details (`--cluster-name` and `--namespace`) of the Kubernetes cluster to delete it.

NOTE: Do not use `dkp get clusters` since that gets you Kommander cluster details rather than Konvoy kubernetes cluster details.

```
kubectl get clusters
```

3. Delete the Kubernetes cluster and wait a few minutes:

NOTE: Before deleting the cluster, dkp deletes all Services of type LoadBalancer on the cluster. To skip this step, use the flag `--delete-kubernetes-resources=false`.

```
dkp delete cluster --cluster-name=${CLUSTER_NAME} --kubeconfig $HOME/.kube/config
```

```
✓ Deleting Services with type LoadBalancer for Cluster default/azure-example
✓ Deleting ClusterResourceSets for Cluster default/azure-example
✓ Deleting cluster resources
✓ Waiting for cluster to be fully deleted
Deleted default/azure-example cluster
```

After the workload cluster is deleted, delete the bootstrap cluster.

7.13.9.4 Delete the Bootstrap Cluster

1. Use `dkp` with the bootstrap cluster to delete the workload cluster.

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

```
✓ Deleting bootstrap cluster
```

7.13.9.5 Next Step:

Once your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will install the [Kommander](#) (see page 1422) component of DKP to see your dashboard and continue customization.

7.13.9.6 Known Limitations



NOTE: Be aware of these limitations in the current release of Konvoy.

- The Konvoy version used to create the workload cluster must match the Konvoy version used to delete the workload cluster.

7.14 AKS Infrastructure

Enterprise

Gov Advanced

When installing DKP on Azure Kubernetes Service (**AKS**) infrastructure, you can choose from multiple configuration types.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

The different types of AKS configuration types supported in DKP are covered in this section.

ⓘ An AKS cluster cannot be a [Management](#) (see page 82) or [Essential](#) (see page 83) cluster. To install DKP on your AKS cluster, first ensure you have a Management cluster with DKP and the Kommander component installed, that handles the lifecycle of your AKS cluster.

- [Create a New AKS Cluster](#) (see page 1189)
- [Explore New AKS Cluster](#) (see page 1193)
- [Delete AKS Cluster](#) (see page 1197)
- [Create a new AKS Cluster from the DKP UI](#) (see page 1198)

7.14.1 Create a New AKS Cluster

Enterprise


Gov Advanced

7.14.1.1 Use DKP to create a new AKS cluster

ⓘ Ensure that the `KUBECONFIG` environment variable is set to the Management cluster by running `export KUBECONFIG=<Management_cluster_kubeconfig>.conf`.

7.14.1.2 Name Your Cluster

Give your cluster a unique name suitable for your environment.

-  The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)⁸⁷⁹ for more naming information.

7.14.1.3 Create a New AKS Kubernetes Cluster


1. Set the environment variable to a name for this cluster.

```
export CLUSTER_NAME=<aks-example>
```

2. Check to see what version of Kubernetes is available in your region. When deploying with AKS, you need to declare the version of Kubernetes you wish to use by running the following command, substituting `<your-location>` for the Azure region you're deploying to:

```
az aks get-versions -o table --location <your-location>
```

3. Set the version of Kubernetes you've chosen:

 **NOTE:** Refer to the current release Kubernetes [compatibility table](#) (see page 117) for correct version to use and choose an available 1.26.x version. The version listed in the command is an example.

```
export KUBERNETES_VERSION=1.26.6
```

4. Create the cluster:

```
dkp create cluster aks --cluster-name=${CLUSTER_NAME} --additional-tags=owner=$(whoami) --kubernetes-version=${KUBERNETES_VERSION}
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

⁸⁷⁹ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

```

Generating cluster resources
cluster.cluster.x-k8s.io/aks-example created
azuremanagedcontrolplane.infrastructure.cluster.x-k8s.io/aks-example created
azuremanagedcluster.infrastructure.cluster.x-k8s.io/aks-example created
machinepool.cluster.x-k8s.io/aks-example created
azuremanagedmachinepool.infrastructure.cluster.x-k8s.io/cp6dsz8 created
machinepool.cluster.x-k8s.io/aks-example-md-0 created
azuremanagedmachinepool.infrastructure.cluster.x-k8s.io/mp6gglj created
clusterresourceset.addons.cluster.x-k8s.io/cluster-autoscaler-aks-example
created
configmap/cluster-autoscaler-aks-example created
clusterresourceset.addons.cluster.x-k8s.io/node-feature-discovery-aks-example
created
configmap/node-feature-discovery-aks-example created
clusterresourceset.addons.cluster.x-k8s.io/nvidia-feature-discovery-aks-example
created
configmap/nvidia-feature-discovery-aks-example created

```

7.14.1.4 Inspecting or Editing the Cluster Objects

Use your favorite editor.



NOTE: Editing the cluster objects requires some understanding of Cluster API. Edits can prevent the cluster from deploying successfully.

The objects are [Custom Resources](#)⁸⁸⁰ defined by Cluster API components, and they belong in three different categories:

- Cluster

A *Cluster* object has references to the infrastructure-specific and control plane objects.
- Control Plane
- Node Pool

A Node Pool is a collection of machines with identical properties. For example, a cluster might have one Node Pool with large memory capacity, another Node Pool with GPU support. Each Node Pool is described by three objects: The MachinePool references an object that describes the configuration of Kubernetes components (for example, kubelet) deployed on each node pool machine, and an infrastructure-specific object that describes the properties of all node pool machines. Here, it references a *KubeadmConfigTemplate*.

For in-depth documentation about the objects, read [Concepts](#)⁸⁸¹ in the Cluster API Book.

1. Wait for the cluster control-plane to be ready:

⁸⁸⁰ <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

⁸⁸¹ <https://cluster-api.sigs.k8s.io/user/concepts.html>

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

```
cluster.cluster.x-k8s.io/aks-example condition met
```

The `READY` status will become `True` after the cluster control-plane becomes ready.

- After the objects are created on the API server, the Cluster API controllers reconcile them. They create infrastructure and machines. As they progress, they update the Status of each object. DKP provides a command to describe the current status of the cluster:

```
dkp describe cluster -c ${CLUSTER_NAME}
```

```
NAME                                                    READY  SEVERITY
REASON  SINCE  MESSAGE
Cluster/aks-example                                     True
48m
├─ClusterInfrastructure - AzureManagedCluster/aks-example
└─ControlPlane - AzureManagedControlPlane/aks-example
```

- As they progress, the controllers also create Events. List the Events using this command:

```
kubectl get events | grep ${CLUSTER_NAME}
```

For brevity, the example uses `grep`. It is also possible to use separate commands to get Events for specific objects. For example, `kubectl get events --field-selector involvedObject.kind="AKSCluster"` and `kubectl get events --field-selector involvedObject.kind="AKSMachine"`.

```
48m          Normal    SuccessfulSetNodeRefs          machinepool/aks-
example-md-0          [{Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000000 UID:e3c30389-660d-46f5-b9d7-219f80b5674d APIVersion:
ResourceVersion: FieldPath:} {Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000001 UID:300d71a0-f3a7-4c29-9ff1-1995ffb9cfd3 APIVersion:
ResourceVersion: FieldPath:} {Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000002 UID:8eae2b39-a415-425d-8417-d915a0b2fa52 APIVersion:
ResourceVersion: FieldPath:} {Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000003 UID:3e860b88-f1a4-44d1-b674-a54fad599a9d APIVersion:
ResourceVersion: FieldPath:}]
```

```

6m4s      Normal      AzureManagedControlPlane  available
azuremanagedcontrolplane/aks-example      successfully reconciled
48m       Normal      SuccessfulSetNodeRefs      machinepool/aks-
example                                     [{Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000000 UID:e3c30389-660d-46f5-b9d7-219f80b5674d APIVersion:
ResourceVersion: FieldPath:} {Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000001 UID:300d71a0-f3a7-4c29-9ff1-1995ffb9cfd3 APIVersion:
ResourceVersion: FieldPath:} {Kind: Namespace: Name:aks-mp6gglj-41174201-
vmss000002 UID:8eae2b39-a415-425d-8417-d915a0b2fa52 APIVersion:
ResourceVersion: FieldPath:}]

```

7.14.1.5 Known Limitations

- The DKP version used to create a workload cluster must match the DKP version used to create a workload cluster.
- DKP supports deploying one workload cluster.
- DKP generates a single nodepool is deployed by default, but you can add additional nodepools.
- DKP does not validate edits to cluster objects.

When complete, you can [explore the new cluster](#) (see page 1193).

7.14.2 Explore New AKS Cluster

Enterprise

Gov Advanced

7.14.2.1 Learn to interact with your AKS Kubernetes cluster

This guide explains how to use the command line to interact with your newly deployed Kubernetes cluster.

Before you start, make sure you have created a workload cluster, as described in [Create a New Cluster](#) (see page 1192).

7.14.2.2 Explore the New AKS Cluster

1. Get a kubeconfig file for the workload cluster:

When the workload cluster is created, the cluster lifecycle services generate a kubeconfig file for the workload cluster, and write it to a *Secret*. The kubeconfig file is scoped to the cluster administrator.

Get the kubeconfig from the *Secret*, and write it to a file, using this command:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

- List the Nodes using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
aks-cp6dsz8-41174201-vmss000000	Ready	agent	56m	v1.26.6
aks-cp6dsz8-41174201-vmss000001	Ready	agent	55m	v1.26.6
aks-cp6dsz8-41174201-vmss000002	Ready	agent	56m	v1.26.6
aks-mp6gglj-41174201-vmss000000	Ready	agent	55m	v1.26.6
aks-mp6gglj-41174201-vmss000001	Ready	agent	55m	v1.26.6
aks-mp6gglj-41174201-vmss000002	Ready	agent	55m	v1.26.6
aks-mp6gglj-41174201-vmss000003	Ready	agent	56m	v1.26.6



NOTE: It may take a few minutes for the Status to move to `Ready` while the Pod network is deployed. The Nodes' Status should change to Ready soon after the `calico-node` DaemonSet Pods are Ready.

- List the Pods using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get --all-namespaces pods
```

NAMESPACE	STATUS	RESTARTS	NAME	AGE	READY
calico-system	Running	0	calico-kube-controllers-5dcd4b47b5-tgslm	3m58s	1/1
calico-system	Running	0	calico-node-46dj9	3m58s	1/1
calico-system	Running	0	calico-node-crdgc	3m58s	1/1
calico-system	Running	0	calico-node-m7s7x	3m58s	1/1

calico-system	calico-node-qfkqc	1/1
Running 0	3m57s	
calico-system	calico-node-sfqfm	1/1
Running 0	3m57s	
calico-system	calico-node-sn67x	1/1
Running 0	3m53s	
calico-system	calico-node-w2pvt	1/1
Running 0	3m58s	
calico-system	calico-typha-6f7f59969c-5z4t5	1/1
Running 0	3m51s	
calico-system	calico-typha-6f7f59969c-ddzqb	1/1
Running 0	3m58s	
calico-system	calico-typha-6f7f59969c-rr4lj	1/1
Running 0	3m51s	
kube-system	azure-ip-masq-agent-4f4v6	1/1
Running 0	4m11s	
kube-system	azure-ip-masq-agent-5xfh2	1/1
Running 0	4m11s	
kube-system	azure-ip-masq-agent-9hlk8	1/1
Running 0	4m8s	
kube-system	azure-ip-masq-agent-9vsgg	1/1
Running 0	4m16s	
kube-system	azure-ip-masq-agent-b9wjj	1/1
Running 0	3m57s	
kube-system	azure-ip-masq-agent-kpjtl	1/1
Running 0	3m53s	
kube-system	azure-ip-masq-agent-vr7hd	1/1
Running 0	3m57s	
kube-system	cluster-autoscaler-b4789f4bf-qkfk2	0/1
Init:0/1 0	3m28s	
kube-system	coredns-845757d86-9jf8b	1/1
Running 0	5m29s	
kube-system	coredns-845757d86-h4xfs	1/1
Running 0	4m	
kube-system	coredns-autoscaler-5f85dc856b-xjb5z	1/1
Running 0	5m23s	
kube-system	csi-azuredisk-node-4n4fx	3/3
Running 0	3m53s	
kube-system	csi-azuredisk-node-8pnjj	3/3
Running 0	3m57s	
kube-system	csi-azuredisk-node-sbt6r	3/3
Running 0	3m57s	
kube-system	csi-azuredisk-node-v25wc	3/3
Running 0	4m16s	
kube-system	csi-azuredisk-node-vfbxg	3/3
Running 0	4m11s	
kube-system	csi-azuredisk-node-w5ff5	3/3
Running 0	4m11s	
kube-system	csi-azuredisk-node-zzgqx	3/3
Running 0	4m8s	
kube-system	csi-azurefile-node-2rpcc	3/3
Running 0	3m57s	

kube-system	csi-azurefile-node-4gqkf	3/3
Running 0	4m11s	
kube-system	csi-azurefile-node-f6k8m	3/3
Running 0	4m16s	
kube-system	csi-azurefile-node-k72xq	3/3
Running 0	4m8s	
kube-system	csi-azurefile-node-vx7r4	3/3
Running 0	3m53s	
kube-system	csi-azurefile-node-zc8kr	3/3
Running 0	4m11s	
kube-system	csi-azurefile-node-zkl6b	3/3
Running 0	3m57s	
kube-system	kube-proxy-4fpb6	1/1
Running 0	3m53s	
kube-system	kube-proxy-6qfbf	1/1
Running 0	4m16s	
kube-system	kube-proxy-6wnt2	1/1
Running 0	4m8s	
kube-system	kube-proxy-cspd5	1/1
Running 0	3m57s	
kube-system	kube-proxy-nsgq6	1/1
Running 0	4m11s	
kube-system	kube-proxy-qz2st	1/1
Running 0	4m11s	
kube-system	kube-proxy-zvh9k	1/1
Running 0	3m57s	
kube-system	metrics-server-6bc97b47f7-ltkkj	1/1
Running 0	5m28s	
kube-system	tunnelfront-77d68f78bf-t78ck	1/1
Running 0	5m23s	
node-feature-discovery	node-feature-discovery-master-65dc499cd-fxwb5	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-277xc	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-4dq5k	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-57nb8	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-b4lkl	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-kslst	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-ppjtm	1/1
Running 0	3m28s	
node-feature-discovery	node-feature-discovery-worker-x5bgf	1/1
Running 0	3m28s	
tigera-operator	tigera-operator-74c4d9cf84-k7css	1/1
Running 0	5m25s	

When ready, you can [delete the cluster](#) (see page 1197).

7.14.3 Delete AKS Cluster

Enterprise

Gov Advanced

Ensure that the `KUBECONFIG` environment variable is set to the self-managed cluster by running `export KUBECONFIG={SELF_MANAGED_AZURE_CLUSTER}.conf`

7.14.3.1 Delete the AKS cluster and clean up your environment

7.14.3.2 Delete the Workload Cluster

1. Delete the Kubernetes cluster and wait a few minutes:

Before deleting the cluster, DKP deletes all Services of type LoadBalancer on the cluster. Deleting the Service deletes the Azure LoadBalancer that backs it. To skip this step, use the flag `--delete-kubernetes-resources=false`.

NOTE: Do not skip this step if the Azure Network is managed by DKP. When DKP deletes cluster, it deletes the Network.

```
dkp delete cluster --cluster-name=${CLUSTER_NAME}
```

```
✓ Deleting Services with type LoadBalancer for Cluster default/aks-example
✓ Deleting ClusterResourceSets for Cluster default/aks-example
✓ Deleting cluster resources
✓ Waiting for cluster to be fully deleted
Deleted default/aks-example cluster
```


7.14.3.3 Next Step:

After your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will [Install the Kommander \(see page 1422\)](#) component of DKP to see your dashboard and continue customization.

7.14.3.4 Known Limitations

NOTE: Be aware of these limitations in the current release of DKP.

- The DKP version used to create the workload cluster must match the DKP version used to delete the workload cluster.

7.14.4 Create a new AKS Cluster from the DKP UI

Enterprise

Gov Advanced

The DKP user interface allows you to provision an AKS cluster from your browser quickly and easily.

7.14.4.1 Prerequisites

7.14.4.1.1 Create an AKS Infrastructure Provider

Before provisioning a cluster via the UI, first create an AKS infrastructure provider to hold your AKS credentials:

1. Log in to the Azure command line:

```
az login
```

2. Create an Azure Service Principal (SP) by running the following command:

```
az ad sp create-for-rbac --role contributor --name "$(whoami)-konvoy" --scopes=/subscriptions/$(az account show --query id -o tsv)
```

3. Select **Infrastructure Providers** from the Dashboard menu.
4. Select **Add Infrastructure Provider**.

5. [Choose a workspace.](#) (see page 653) If you are already in a workspace, the provider is automatically created in that workspace.
6. Select **Microsoft Azure**.
7. Add a **Name** for your Infrastructure Provider.
8. Take the `id` output from the log in command above and put it into the **Subscription ID** field.
9. Take the `tenant` used in Step 2 and put it into the **Tenant ID** field.
10. Take the `appId` used in Step 2 and put it into the **Client ID** field.
11. Take the `password` used in Step 2 and put it into the **Client Secret** field.
12. Select **Save**.

7.14.4.2 Provision an AKS Cluster

Follow these steps to provision an AKS cluster:

1. From the top menu bar, select your target workspace.
2. Select **Clusters > Add Cluster**.
3. Choose **Create Cluster**.
4. Enter the **Cluster Name**.
5. From **Select Infrastructure Provider**, choose the provider created in the prerequisites section.
6. In order to create a **Kubernetes Version**, run the command below in the `az` cli, and then select the version of AKS you want to use.

```
az aks get-versions -o table --location <location>
```

7. Select a data center **location** or specify a custom **location**.
8. Edit your worker **Node Pools**, as necessary. You can choose the **Number of Nodes**, the **Machine Type**, and for the worker nodes, you can choose a **Worker Availability Zone**.
9. Add any additional **Labels** or **Infrastructure Provider Tags**, as necessary.
10. Validate your inputs, then select **Create**.



It can take up to 15 minutes for your cluster to appear in the **Provisioned** status.

You are then redirected to the **Clusters** page, where you'll see your new cluster in the **Provisioning** status. Hover over the status to view the details.

7.14.4.3 Access an AKS Cluster

After the cluster is successfully attached (managed), you can retrieve a custom `kubeconfig` file from the UI using your DKP UI administrator credentials.

7.15 Pre-provisioned Infrastructure

7.15.1 Create a Kubernetes cluster on pre-provisioned nodes in a bare metal infrastructure

The following procedure describes creating a DKP cluster on a pre-provisioned infrastructure using SSH.

Completing this procedure results in a Kubernetes cluster that includes a [Container Networking Interface \(CNI\)](#)⁸⁸² and a [Local Persistence Volume Static Provisioner](#)⁸⁸³, and that is ready for workload deployment.

Before moving to a production environment, you may want to add applications for logging and monitoring, storage, security, and other functions. You can use DKP to select and [deploy applications](#) (see page 642), or deploy your own.

To get started, see these topics.

- [Advanced Workflow](#) (see page 1201)
- [Pre-provisioned Prerequisite Configuration](#) (see page 1201)
- [Pre-provisioned Air-gapped Define Environment](#) (see page 1204)
- [Pre-provisioned Set Infrastructure](#) (see page 1209)
- [Pre-provisioned Define Control Plane Endpoint](#) (see page 1211)
- [Pre-provisioned Create Secrets and Overrides](#) (see page 1213)
- [Pre-provisioned Bootstrap Cluster](#) (see page 1219)
- [Pre-provisioned Create a New Cluster](#) (see page 1220)
- [Pre-provisioned Azure only Configurations](#) (see page 1227)
- [Pre-provisioned Modify the Calico Installation](#) (see page 1231)
- [Pre-provisioned Built-in Virtual IP](#) (see page 1236)
- [Provision on the Flatcar Linux OS](#) (see page 1236)
- [Pre-provisioned Use HTTP Proxy](#) (see page 1237)
- [Pre-provisioned Use Alternate Pod or Service Subnets](#) (see page 1238)
- [Pre-provisioned: Make Cluster Self-managed](#) (see page 1240)
- [Pre-provisioned Configure MetalLB](#) (see page 1243)
- [Pre-provisioned Create and Delete Node Pools](#) (see page 1245)
- [Pre-provisioned Add Nodes to Existing Node Pool](#) (see page 1247)
- [GPU Nodepools in a Pre-provisioned Environment](#) (see page 1249)
- [Pre-provisioned Delete Cluster](#) (see page 1251)

⁸⁸² <https://docs.projectcalico.org/>

⁸⁸³ <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

7.15.2 Advanced Workflow

In the Day 1 - Basic Installations, the installation process for Pre-provisioned on any provider uses a `self-managed` flag. This process eliminates errors of creating a bootstrap and not managing the KUBECONFIG properly to run commands against the correct cluster when you have multiples.

The more advanced workflow for Pre-provisioned, allows for custom YAML files to pass various overrides and other custom configurations during cluster creation. The pages in this section describe these more complex scenarios.

7.15.3 Pre-provisioned Prerequisite Configuration

In order to fulfill all the prerequisites for a successful implementation on a Pre-provisioned environment, there will be infrastructure requirements as well as machine requirements. Please read all the sections on this page to ensure you have met all prerequisites.

7.15.3.1 Prerequisites for a Pre-provisioned infrastructure

Before you begin using DKP, you must have:

- An x86_64-based Linux or macOS machine.
- The `dkp` binary for Linux, or macOS.
- [kubecti](#)⁸⁸⁴ for interacting with the running cluster.
- Pre-provisioned hosts with SSH access enabled.
- An unencrypted SSH private key, whose public key is configured on the above hosts.
- A Container engine/runtime installed is required to install DKP:
 - Version [Docker](#)⁸⁸⁵ container engine version 18.09.2 or higher installed for Linux or MacOS - On macOS, Docker runs in a virtual machine which needs configured with at least 8 GB of memory.
 - Version 4.0 of [Podman](#)⁸⁸⁶ or higher for Linux. Host requirements found here: [Host Requirements](#)⁸⁸⁷
- For air-gapped, create a working registry:
 - [local registry](#) (see page 1533) on bastion or other machine
- Resource [requirements](#) (see page 100)
- [Pre-provisioned Override Files](#) (see page 1524)


When in an air-gapped environment, you must follow the steps described in the [Air-gapped Define Environment](#) (see page 1204) and Docker Registry also as a prerequisite.

884 <https://kubernetes.io/docs/tasks/tools/#kubecti>

885 <https://docs.docker.com/get-docker/>

886 <https://podman.io/getting-started/installation>

887 <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>

 DKP uses `localvolumeplugin` as the [default storage provider](#) (see page 91). However, `localvolumeplugin` is not suitable for production use. You should use a [Kubernetes CSI](#)⁸⁸⁸ compatible storage that is suitable for production.

You can choose from any of the [storage options](#)⁸⁸⁹ available for Kubernetes. To disable the default that Konvoy deploys, set the default StorageClass `localvolumeplugin` as non-default. Then set your newly created StorageClass to be the default by following the commands in the Kubernetes documentation called [Changing the Default Storage Class](#)⁸⁹⁰.


7.15.3.2 Machine Specifications

7.15.3.2.1 Control plane machines

You should have at least three control plane machines.

Each control plane machine must have:

- 4 cores
- 16 GiB memory
- Approximately 80 GiB of free space for the volume used for `/var/lib/kubelet` and `/var/lib/containerd`.
- 15% free space on the root file system.
- Multiple ports open, as described in [DKP Ports](#) (see page 67).
- `firewalld` `systemd` service disabled. If it exists and is enabled, use the commands `systemctl stop firewalld` then `systemctl disable firewalld`, so that `firewalld` remains disabled after the machine restarts.
- For a Pre-provisioned environment using Ubuntu 20.04, ensure the machine has the `/run` directory mounted with `exec` permissions.

 Swap is disabled. The `kubelet` does not have generally-available support for swap. Due to variable commands, refer to your operating system documentation.

⁸⁸⁸ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁸⁸⁹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>


⁸⁹⁰ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

7.15.3.2 Worker machines

You should have at least four worker machines. The specific number of worker machines required for your environment can vary depending on the cluster workload and size of the machines.

Each worker machine must have:

- 8 cores
- 32 GiB memory
- Around 80 GiB of free space for the volume used for `/var/lib/kubelet` and `/var/lib/containerd`.
- 15% free space on the root file system
- If you plan to use local volume provisioning to provide persistent volumes for your workloads, you must mount at least four volumes to the `/mnt/disks/` mount point on each machine. Each volume must have at least 55 GiB of capacity.
- Ensure your disk meets the resource requirements for Rook Ceph in `Block` mode for [ObjectStorageDaemons](#)⁸⁹¹ as specified in the [requirements table](#) (see page 1009).
- Multiple ports open, as described in [DKP Ports](#) (see page 67).
- `firewalld` `systemd` service disabled. If it exists and is enabled, use the commands `systemctl stop firewalld` then `systemctl disable firewalld`, so that `firewalld` remains disabled after the machine restarts.
- For a Pre-provisioned environment using Ubuntu 20.04, ensure the machine has the `/run` directory mounted with `exec` permissions.

 Swap is disabled. The `kubelet` does not have generally-available support for swap. Due to variable commands, refer to your operating system documentation.

7.15.3.3 Next Step:

[Pre-provisioned Set Infrastructure](#) (see page 1209)

⁸⁹¹ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

7.15.4 Pre-provisioned Air-gapped Define Environment

7.15.4.1 Fulfill the prerequisites for using a pre-provisioned infrastructure when Air-Gapped

The instructions below outline how to fulfill the prerequisites for using pre-provisioned infrastructure when using an air-gapped environment. In the current release of DKP, there is a complete DKP air-gapped bundle available to [download](#) (see page 78) which contains all the DKP components needed for an air-gapped [Install Overview](#) (see page 108) environment installation. (i.e. `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`)

7.15.4.2 Air-Gapped Registry Prerequisites

DKP in an air-gapped environment requires a [local container registry](#) (see page 1533) of trusted images to enable production level Kubernetes cluster management. In an environment with access to the internet, you retrieve artifacts from specialized repositories dedicated to them such as Docker images contained in DockerHub and Helm Charts that come from a dedicated Helm Chart repository. However, in an air-gapped environment, you need [local repositories](#) (see page 1533) to store Helm charts, Docker images and other artifacts. Tools such as ECR, jFrog, Harbor and Nexus handle multiple types of artifacts in one local repository.

- AWS ECR - AWS ECR (Elastic Container Registry) is supported as your air-gapped image registry. DKP added support for using AWS ECR as a default registry when uploading image bundles.
- JFrog Artifactory - If you use Jfrog Artifactory or Jfrog Container Registry, you must update to a new version of the software. Any build newer than version 7.11 will work, as we have confirmed that older versions are not compatible.
- Nexus Registry - If you use Nexus Registry, there was an issue that prevented usage with DKP 2.X and OCI Images, but support for OCI Images was added here in this publicly available Jira ticket: [\[NEXUS-21087\] Support OCI registry format - Sonatype JIRA](#)⁸⁹²
- Harbor Registry - Any newer version than **Harbor Registry v2.1.1-5f52168e** will support OCI images.

7.15.4.2.1 Bastion Host

If you have not set up a [Bastion Host](#) (see page 1285) yet, refer to that section of the Documentation.

7.15.4.3 Load the bootstrap image

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz` from the [download site](#) (see page 78) mentioned above, extract the tarball to a local directory:

⁸⁹² <https://issues.sonatype.org/browse/NEXUS-21087>

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz
```

- The directory structure after extraction can be accessed in subsequent steps using commands to access files from different directories. For the bootstrap, change your directory to the `dkp-<version>` directory similar to example below depending on your current location:

```
cd dkp-v2.6.2
```

- Load the bootstrap Docker image on your bastion machine:

```
docker load -i konvoy-bootstrap-image-v2.6.2.tar
```

7.15.4.4 Copy air-gapped artifacts onto cluster hosts

Using the [Konvoy Image Builder](#) (see page 1461), you can copy the required artifacts (such as charts, java or OS packages like RPM or Deb) onto your cluster hosts.

- Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz` and extracted the tarball from the section above, the Kubernetes image bundle will be located in `kib/artifacts/images`. Verify image and artifacts with the commands below:
 - Verify the image bundles exist in `kib/artifacts/images`:

```
$ ls kib/artifacts/images/
kubernetes-images-1.26.14-d2iq.1.tar kubernetes-images-1.26.14-d2iq.1-
fips.tar
```

- Verify the artifacts for your OS exist in the `artifacts/` directory and export the appropriate variables:

```
$ ls kib/artifacts/
1.26.14-centos_7_x86_64.tar.gz      1.26.14-redhat_8_x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rhel-8.6-x86_64_fips.tar.gz  pip-packages.tar.gz
1.26.14-centos_7_x86_64_fips.tar.gz  1.26.14-rocky_9_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64_fips.tar.gz  containerd-1.6.17-
d2iq.1-rocky-9.0-x86_64.tar.gz
1.26.14-redhat_7_x86_64.tar.gz      1.26.14-ubuntu_20_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rocky-9.1-x86_64.tar.gz
1.26.14-redhat_7_x86_64_fips.tar.gz  containerd-1.6.17-d2iq.1-centos-7.9-
x86_64.tar.gz      containerd-1.6.17-d2iq.1-rhel-8.4-x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-ubuntu-20.04-x86_64.tar.gz
```



```
1.26.14_redhat_8_x86_64.tar.gz      containerd-1.6.17-d2iq.1-centos-7.9-
x86_64_fips.tar.gz  containerd-1.6.17-d2iq.1-rhel-8.6-x86_64.tar.gz
images
```

- c. For example, for RHEL 8.4 you would set:

```
export OS_PACKAGES_BUNDLE=1.26.14_redhat_8_x86_64.tar.gz
export CONTAINERD_BUNDLE=containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz
```

2. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```
export CONTROL_PLANE_1_ADDRESS="<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<control-plane-address-3>"
export WORKER_1_ADDRESS="<worker-address-1>"
export WORKER_2_ADDRESS="<worker-address-2>"
export WORKER_3_ADDRESS="<worker-address-3>"
export WORKER_4_ADDRESS="<worker-address-4>"
export SSH_USER="<ssh-user>"
export SSH_PRIVATE_KEY_FILE="<private key file>"
```

SSH_PRIVATE_KEY_FILE must be either the name of the SSH private key file in your working directory or an absolute path to the file in your user's home directory.

3. Generate an `inventory.yaml` which is automatically picked up by the `konvoy-image upload` in the next step. This `inventory.yaml` **should exclude any GPU workers**, which will be handled in steps #6-7.

```
cat <<EOF > inventory.yaml
all:
  vars:
    ansible_user: $SSH_USER
    ansible_port: 22
    ansible_ssh_private_key_file: $SSH_PRIVATE_KEY_FILE
  hosts:
    $CONTROL_PLANE_1_ADDRESS:
      ansible_host: $CONTROL_PLANE_1_ADDRESS
    $CONTROL_PLANE_2_ADDRESS:
      ansible_host: $CONTROL_PLANE_2_ADDRESS
    $CONTROL_PLANE_3_ADDRESS:
      ansible_host: $CONTROL_PLANE_3_ADDRESS
    $WORKER_1_ADDRESS:
      ansible_host: $WORKER_1_ADDRESS
    $WORKER_2_ADDRESS:
      ansible_host: $WORKER_2_ADDRESS
    $WORKER_3_ADDRESS:
```

```

    ansible_host: $WORKER_3_ADDRESS
$WORKER_4_ADDRESS:
    ansible_host: $WORKER_4_ADDRESS
EOF

```

4. Upload the artifacts onto cluster hosts with the following command:

```

konvoy-image upload artifacts \
    --container-images-dir=./kib/artifacts/images/ \
    --os-packages-bundle=./kib/artifacts/$OS_PACKAGES_BUNDLE \
    --containerd-bundle=./kib/artifacts/$CONTAINERD_BUNDLE \
    --pip-packages-bundle=./kib/artifacts/pip-packages.tar.gz

```

The `konvoy-image upload artifacts` command copies all OS packages and other artifacts onto each of the machines in your inventory. When you create the cluster, the provisioning process connects to each node and runs commands to install those artifacts and consequently Kubernetes running. KIB uses [variable overrides](#) (see [page 1513](#)) to specify base image and container images to use in your new machine image. The variable overrides files for NVIDIA and FIPS can be ignored unless adding an overlay feature.

- = • Use the overrides flag (EX: `--overrides overrides/fips.yaml`) and reference either `fips.yaml` or `offline-fips.yaml` manifests located in the [overrides directory](#)⁸⁹³ or see these pages in the documentation:
 - [FIPS Overrides](#) (see [page 1514](#))
 - [Create FIPS 140 Images](#) (see [page 1529](#))

7.15.4.4.1 GPU Only Steps

- = If the [NVIDIA runfile](#)⁸⁹⁴ installer has not been downloaded, then retrieve and install the download first by running the following command. The first line in the command below downloads and installs the runfile and the second line places it in the artifacts directory.

```

• curl -O https://download.nvidia.com/XFree86/Linux-x86_64/470.82.01/
  NVIDIA-Linux-x86_64-470.82.01.run

```

⁸⁹³ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

⁸⁹⁴ <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

```
mv NVIDIA-Linux-x86_64-470.82.01.run artifacts
```

6. Create an inventory for GPU Nodes.

```
cat <<EOF > gpu_inventory.yaml
all:
  vars:
    ansible_port: 22
    ansible_ssh_private_key_file: $SSH_PRIVATE_KEY_FILE
    ansible_user: $SSH_USER

  hosts:
    $GPU_WORKER_1_ADDRESS:
      ansible_host: $GPU_WORKER_1_ADDRESS
EOF
```

7. Upload the artifacts to the gpu nodepool with the `nvidia-runfile` flag

```
konvoy-image upload artifacts \
  --inventory-file=gpu_inventory.yaml \
  --container-images-dir=./kib/artifacts/images/ \
  --os-packages-bundle=./kib/artifacts/$OS_PACKAGES_BUNDLE \
  --containerd-bundle=./kib/artifacts/$CONTAINERD_BUNDLE \
  --pip-packages-bundle=./kib/artifacts/pip-packages.tar.gz \
  --nvidia-runfile=./kib/artifacts/NVIDIA-Linux-x86_64-470.82.01.run
```

KIB uses [variable overrides](#) (see page 1513) to specify base image and container images to use in your new machine image. The variable overrides files for NVIDIA and FIPS can be ignored unless adding an overlay feature.

7.15.4.4.2 Seed your Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine](#) (see page 1285) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.



If you do not already have a local registry set up, please refer to [Local Registry Tools](#) (see page 1533) page for more information.

1. [Download](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, and extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

3. Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```



It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

7.15.4.4.3 Next Step:

[Pre-provisioned Set Infrastructure](#) (see page 1209)

7.15.5 Pre-provisioned Set Infrastructure

Define the cluster hosts and infrastructure

Konvoy needs to know how to access your cluster hosts. This is done using inventory resources. For initial cluster creation, you must define a control-plane and at least one worker pool.

7.15.5.1 Define your infrastructure

1. Export the following environment variables, ensuring that all control plane and worker nodes are included:

```
export CONTROL_PLANE_1_ADDRESS="<control-plane-address-1>"
export CONTROL_PLANE_2_ADDRESS="<control-plane-address-2>"
export CONTROL_PLANE_3_ADDRESS="<control-plane-address-3>"
export WORKER_1_ADDRESS="<worker-address-1>"
export WORKER_2_ADDRESS="<worker-address-2>"
export WORKER_3_ADDRESS="<worker-address-3>"
```

```
export WORKER_4_ADDRESS="<worker-address-4>"
export SSH_USER="<ssh-user>"
export SSH_PRIVATE_KEY_SECRET_NAME="$(CLUSTER_NAME)-ssh-key"
```

2. Use the following template to help you define your infrastructure. The environment variables that you set in the previous step automatically replace the variable names when the file is created.

```
cat <<EOF > preprovisioned_inventory.yaml
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: $(CLUSTER_NAME)-control-plane
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: $(CLUSTER_NAME)
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    # Create as many of these as needed to match your infrastructure
    # Note that the command line parameter --control-plane-replicas determines
    # how many control plane nodes will actually be used.
    #
    - address: $(CONTROL_PLANE_1_ADDRESS)
    - address: $(CONTROL_PLANE_2_ADDRESS)
    - address: $(CONTROL_PLANE_3_ADDRESS)
  sshConfig:
    port: 22
    # This is the username used to connect to your infrastructure. This user
    # must be root or
    # have the ability to use sudo without a password
    user: $SSH_USER
    privateKeyRef:
      # This is the name of the secret you created in the previous step. It
      # must exist in the same
      # namespace as this inventory object.
      name: $SSH_PRIVATE_KEY_SECRET_NAME
      namespace: default
---
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: $(CLUSTER_NAME)-md-0
  namespace: default
  labels:
    cluster.x-k8s.io/cluster-name: $(CLUSTER_NAME)
    clusterctl.cluster.x-k8s.io/move: ""
spec:
  hosts:
    - address: $WORKER_1_ADDRESS
```

```

- address: $WORKER_2_ADDRESS
- address: $WORKER_3_ADDRESS
- address: $WORKER_4_ADDRESS
sshConfig:
  port: 22
  user: $SSH_USER
  privateKeyRef:
    name: $SSH_PRIVATE_KEY_SECRET_NAME
    namespace: default
EOF

```

- To tell the bootstrap cluster which nodes you want to be control plane nodes and which nodes are worker nodes, you use the `kubectl apply` command to apply the file to the bootstrap cluster:

```
kubectl apply -f preprovisioned_inventory.yaml
```

Output:

```

preprovisionedinventory.infrastructure.cluster.konvoy.d2iq.io/preprovisioned-
example-control-plane created
preprovisionedinventory.infrastructure.cluster.konvoy.d2iq.io/preprovisioned-
example-md-0 created

```

After defining the infrastructure, [define the control plane endpoint](#) (see page 1211).

7.15.6 Pre-provisioned Define Control Plane Endpoint

7.15.6.1 Define the Control Plane Endpoint for your cluster

A control plane should have three, five, or seven nodes, so it can remain available if one, two, or three nodes fail. A control plane with one node should not be used in production.

In addition, the control plane should have an endpoint that remains available if some nodes fail.

```

          ----- cp1.example.com:6443
          |
lb.example.com:6443 ----- cp2.example.com:6443
          |
          ----- cp3.example.com:6443

```

In this example, the control plane endpoint host is `lb.example.com`, and the control plane endpoint port is `6443`. The control plane nodes are `cp1.example.com`, `cp2.example.com`, and `cp3.example.com`. The port of each API server is `6443`.

7.15.6.2 External Load Balancer

It is recommended that an external load balancer(LB) be the control plane endpoint. To distribute request load among the control plane machines, configure the load balancer to send requests to all the control plane machines. Configure the load balancer to send requests only to control plane machines that are responding to API requests.

7.15.6.3 Built-in virtual IP

If an external load balancer is not available, use the [built-in virtual IP \(see page 1220\)](#). The virtual IP is *not* a load balancer; it does not distribute request load among the control plane machines. However, if the machine receiving requests does not respond to them, the virtual IP automatically moves to another machine.

7.15.6.4 Single-Node control plane



Do not use a single-node control plane in a production cluster.

A control plane with one node can use its single node as the endpoint, so you will not require an external load balancer, or a built-in virtual IP. At least one control plane node must always be running. Therefore, to upgrade a cluster with one control plane node, a spare machine must be available in the control plane inventory. This machine is used to provision the new node before the old node is deleted.



Modify Control Plane Audit logs settings using the information contained in the page [Configuring the Control Plane \(see page 1538\)](#).

When the API server endpoints are defined, you can [create the cluster \(see page 1220\)](#).

7.15.6.5 Known limitations



Be aware of these limitations in the current release of DKP.

The control plane endpoint port is also used as the API server port on each control plane machine. The default port is 6443. Before you create the cluster, ensure the port is available for use on each control plane machine.

7.15.6.5.1 Next Step:

[Pre-provisioned Bootstrap Cluster](#) (see page 1219) unless you need to [Create Secrets and Overrides](#) (see page 1213) first.

7.15.7 Pre-provisioned Create Secrets and Overrides

7.15.7.1 Create necessary secrets and overrides for pre-provisioned clusters

Most applications deployed through [Kubernetes](#)⁸⁹⁵ require access to databases, services, and other resources located externally. The easiest way to manage the login information necessary to access those resources is using secrets in order to help organize and distribute sensitive information across a cluster while minimizing the risk of sensitive information exposure.

DKP requires SSH access to your infrastructure with superuser privileges. You must provide an unencrypted SSH private key to DKP so secrets are a good way to achieve this. Populate the key and create the required secret, on your bootstrap cluster using the following procedure.

7.15.7.2 Create a unique cluster name

Give your cluster a unique name suitable for your environment.

Set the environment variable to be used throughout this procedure:

```
export CLUSTER_NAME=preprovisioned-example
```

(Optional) If you want to create a unique cluster name, use this command. This creates a unique name every time you run it, so use it carefully.

```
export CLUSTER_NAME=preprovisioned-example-$(LC_CTYPE=C tr -dc 'a-z0-9' </dev/urandom
| fold -w 5 | head -n1)
echo $CLUSTER_NAME
```

```
preprovisioned-example-pf4a3
```

7.15.7.3 Create a secret

Create a secret that contains the SSH key with these commands:

⁸⁹⁵ <https://kubernetes.io/docs/concepts/configuration/secret/>


```
export SSH_PRIVATE_KEY_FILE="<path-to-ssh-private-key>"
```

```
export SSH_PRIVATE_KEY_SECRET_NAME=$CLUSTER_NAME-ssh-key
```

```
kubectl create secret generic ${SSH_PRIVATE_KEY_SECRET_NAME} --from-file=ssh-
privatekey=${SSH_PRIVATE_KEY_FILE}
kubectl label secret ${SSH_PRIVATE_KEY_SECRET_NAME} clusterctl.cluster.x-k8s.io/move=
```

```
secret/preprovisioned-example-ssh-key created
secret/preprovisioned-example-ssh-key labeled
```

7.15.7.4 Create Overrides

In these steps, you will point your machines at the desired Registry to obtain the container images. If your pre-provisioned machines need to have [Custom Override Files](#) (see page 1520), create a secret that includes all the overrides you want to provide in one file.

1. Example CentOS7 and Docker:

If you want to provide an override with Docker credentials and a different source for EPEL on a CentOS7 machine, you should create a file like this:

```
cat > overrides.yaml << EOF
image_registries_with_auth:
- host: "registry-1.docker.io"
  username: "my-user"
  password: "my-password"
  auth: ""
  identityToken: ""

epel_centos_7_rpm: https://my-rpm-repository.org/epel/epel-release-
latest-7.noarch.rpm
EOF
```

You can then create the related secret by running the following command:

```
kubectl create secret generic $CLUSTER_NAME-user-overrides --from-
file=overrides.yaml=overrides.yaml
kubectl label secret $CLUSTER_NAME-user-overrides clusterctl.cluster.x-k8s.io/
move=
```

2. Example:

When using Oracle 7 OS, you may wish to deploy the RHCK kernel instead of the default UEK kernel.

To do so, add the following text to your `overrides.yaml`:

```
cat > overrides.yaml << EOF
---
oracle_kernel: RHCK
EOF
```

You can then create the related secret by running the following command:

```
kubectl create secret generic $CLUSTER_NAME-user-overrides --from-
file=overrides.yaml=overrides.yaml
kubectl label secret $CLUSTER_NAME-user-overrides clusterctl.cluster.x-k8s.io/
move=
```

7.15.7.5 Next Step:

[Pre-provisioned Set Infrastructure](#) (see page 1212)

7.15.7.6 Pre-provisioned FIPS Create Secrets and Overrides

7.15.7.6.1 Create necessary secrets and overrides for pre-provisioned clusters

Most applications deployed through [Kubernetes](#)⁸⁹⁶ require access to databases, services, and other resources located externally. The easiest way to manage the login information necessary to access those resources is using secrets in order to help organize and distribute sensitive information across a cluster while minimizing the risk of sensitive information exposure.

DKP requires SSH access to your infrastructure with superuser privileges. You must provide an unencrypted SSH private key to DKP so secrets are a good way to achieve this. Populate the key and create the required secret, on your bootstrap cluster using the following procedure.

7.15.7.6.2 Create a unique cluster name

Give your cluster a unique name suitable for your environment.

Set the environment variable to be used throughout this procedure:

```
export CLUSTER_NAME=preprovisioned-example
```

⁸⁹⁶ <https://kubernetes.io/docs/concepts/configuration/secret/>

(Optional) If you want to create a unique cluster name, use this command. This creates a unique name every time you run it, so use it carefully.

```
export CLUSTER_NAME=preprovisioned-example-$(LC_CTYPE=C tr -dc 'a-z0-9' </dev/urandom
| fold -w 5 | head -n1)
echo $CLUSTER_NAME
```

```
preprovisioned-example-pf4a3
```

7.15.7.6.3 Create a secret

Create a secret that contains the SSH key with these commands:

```
export SSH_PRIVATE_KEY_FILE="<path-to-ssh-private-key>"
```

```
export SSH_PRIVATE_KEY_SECRET_NAME=$CLUSTER_NAME-ssh-key
```

```
kubectl create secret generic ${SSH_PRIVATE_KEY_SECRET_NAME} --from-file=ssh-
privatekey=${SSH_PRIVATE_KEY_FILE}
kubectl label secret ${SSH_PRIVATE_KEY_SECRET_NAME} clusterctl.cluster.x-k8s.io/move=
```

```
secret/preprovisioned-example-ssh-key created
secret/preprovisioned-example-ssh-key labeled
```

7.15.7.6.4 Non-air-gapped Environment Create FIPS-140 images

KIB can produce images containing FIPS-140 compliant binaries. Use the `fips.yaml` [override file](#) (see [page 1514](#)) provided with the image bundles.



You can also find these override files in the [Konvoy Image Builder repo](#)⁸⁹⁷.

⁸⁹⁷ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

7.15.7.6.5 Create overrides

1. Create a secret that includes the customization Overrides for FIPS compliance:

Note: Get the latest values for FIPS from the [Konvoy Image Builder repo](#)⁸⁹⁸.

```
cat > overrides.yaml << EOF
---
k8s_image_registry: docker.io/mesosphere

fips:
  enabled: true

build_name_extra: -fips
kubernetes_build_metadata: fips.0
default_image_repo: hub.docker.io/mesosphere
kubernetes_rpm_repository_url: "https://packages.d2iq.com/konvoy/stable/linux/
repos/el/kubernetes-v{{ kubernetes_version }}-fips/x86_64"
docker_rpm_repository_url: "\
https://containerd-fips.s3.us-east-2.amazonaws.com\
/{{ ansible_distribution_major_version|int }}\
/x86_64"
EOF
```

2. If your pre-provisioned machines need to have a customization with alternate package libraries, Docker image or other container registry image repos, or other [Custom Override Files](#) (see page 1520), **add** more lines to the same Overrides file.

- a. Example:

If you want to provide an override with Docker credentials and a different source for EPEL on a CentOS7 machine, you should create a file like this:

```
cat > overrides.yaml << EOF
---
# fips configuration
k8s_image_registry: docker.io/mesosphere

fips:
  enabled: true

build_name_extra: -fips
kubernetes_build_metadata: fips.0
default_image_repo: hub.docker.io/mesosphere
kubernetes_rpm_repository_url: "https://packages.d2iq.com/konvoy/stable/
linux/repos/el/kubernetes-v{{ kubernetes_version }}-fips/x86_64"
```

⁸⁹⁸ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

```

docker_rpm_repository_url: "\
https://containerd-fips.s3.us-east-2.amazonaws.com\
/{{ ansible_distribution_major_version|int }}\
/x86_64"

# custom configuration
image_registries_with_auth:
- host: "registry-1.docker.io"
  username: "my-user"
  password: "my-password"
  auth: ""
  identityToken: ""

epel_centos_7_rpm: https://my-rpm-repostory.org/epel/epel-release-
latest-7.noarch.rpm
EOF

```

b. Example:

When using Oracle 7 OS, you may wish to deploy the RHCK kernel instead of the default UEK kernel. To do so, add the following text to your `overrides.yaml` :

```

cat > overrides.yaml << EOF
---
# fips configuration
k8s_image_registry: docker.io/mesosphere

fips:
  enabled: true

build_name_extra: -fips
kubernetes_build_metadata: fips.0
default_image_repo: hub.docker.io/mesosphere
kubernetes_rpm_repository_url: "https://packages.d2iq.com/konvoy/stable/
linux/repos/el/kubernetes-v{{ kubernetes_version }}-fips/x86_64"
docker_rpm_repository_url: "\
https://containerd-fips.s3.us-east-2.amazonaws.com\
/{{ ansible_distribution_major_version|int }}\
/x86_64"

# custom configuration
oracle_kernel: RHCK
EOF

```

3. Create the related secret by running the following command:

```

kubectl create secret generic $CLUSTER_NAME-user-overrides --from-
file=overrides.yaml=overrides.yaml

```

```
kubectl label secret $CLUSTER_NAME-user-overrides clusterctl.cluster.x-k8s.io/
move=
```

7.15.8 Pre-provisioned Bootstrap Cluster

A bootstrap cluster refers to a special type of local Kubernetes cluster used to bootstrap other clusters. The bootstrap cluster is required because the controllers that create other Kubernetes clusters require a Kubernetes cluster to run.

Konvoy deploys all cluster lifecycle services to a bootstrap cluster, which deploys a workload cluster. When the workload cluster is ready, move the cluster lifecycle services to the workload cluster. The workload cluster then manages its own lifecycle.

7.15.8.1 Bootstrap a kind cluster and CAPI controllers

Use the following command to create a bootstrap cluster:

```
dkp create bootstrap
```

- ✓ Creating a bootstrap cluster
- ✓ Initializing **new** CAPI components

Konvoy creates a bootstrap cluster using [KIND](#)⁸⁹⁹ as a library. Konvoy then deploys the following [Cluster API](#)⁹⁰⁰ providers on the cluster:

- [Core Provider](#)⁹⁰¹
- [vSphere Infrastructure Provider](#)⁹⁰²
- [Kubeadm Bootstrap Provider](#)⁹⁰³
- [Kubeadm ControlPlane Provider](#)⁹⁰⁴

7.15.8.2 Next Step:

[Pre-provisioned Create a New Cluster](#) (see page 1220)

⁸⁹⁹ <https://github.com/kubernetes-sigs/kind>

⁹⁰⁰ <https://cluster-api.sigs.k8s.io/>

⁹⁰¹ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

⁹⁰² <https://github.com/kubernetes-sigs/cluster-api-provider-vsphere>

⁹⁰³ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

⁹⁰⁴ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>

7.15.9 Pre-provisioned Create a New Cluster

7.15.9.1 Create a Kubernetes cluster using the infrastructure definition

After you have defined the [infrastructure](#) (see page 1212) and [control plane endpoints](#) (see page 1214), you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster.

i Before you create a new DKP cluster below, you may choose an external load balancer or virtual IP and use the corresponding `dkp create cluster` command example from that page in the docs from the links below. Other customizations are available, but require different flags during `dkp create cluster` command also. Refer to corresponding section for custom cluster creation:

- [Virtual IP](#) (see page 1236)
- [Provision on the Flatcar Linux OS](#) (see page 1236)
- [Use an HTTP proxy](#) (see page 1237)
- [Use alternate pod or service subnets](#) (see page 1238)

! DKP uses `localvolumeprovisioner` as the [default storage provider](#) (see page 91). However, `localvolumeprovisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)⁹⁰⁵ compatible storage that is suitable for production.

You can choose from any of the [storage options](#)⁹⁰⁶ available for Kubernetes. To disable the default that Konvoy deploys, set the default StorageClass `localvolumeprovisioner` as non-default. Then set your newly created StorageClass to be the default by following the commands in the Kubernetes documentation called [Changing the Default Storage Class](#)⁹⁰⁷.

1. The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the inventory.

NOTE: When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

NOTE: (Optional) If you have [overrides for your clusters](#)⁹⁰⁸, you must specify the secret as part of the create cluster command. If these are not specified, the overrides for your nodes will not be applied.

```
--override-secret-name=$CLUSTER_NAME-user-overrides
```

⁹⁰⁵ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁹⁰⁶ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

⁹⁰⁷ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

⁹⁰⁸ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29918736>

NOTE: To increase [Docker Hub's rate limit](https://docs.docker.com/docker-hub/download-rate-limit/)⁹⁰⁹ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io` `--registry-mirror-username=` `--registry-mirror-password=` on the `dkp create cluster` command.

NOTE: Ensure your [subnets](#) (see page 1238) do not overlap with your host subnet because they cannot be changed after cluster creation. If you need to change the kubernetes subnets, you must do this at cluster creation. The default subnets used in DKP are:

```
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
        - 192.168.0.0/16
    services:
      cidrBlocks:
        - 10.96.0.0/12
```

Create cluster command may take a while depending on the size of the cluster:

```
dkp create cluster preprovisioned \
--cluster-name ${CLUSTER_NAME} \
--control-plane-endpoint-host <control plane endpoint host> \
--control-plane-endpoint-port <control plane endpoint port, if different than 6443> \
--override-secret-name=${CLUSTER_NAME}-user-overrides
```

The output from this command is shortened here for reading clarity, but should start like this:

```
Generating cluster resources
cluster.cluster.x-k8s.io/preprovisioned-example created
cont.....
```

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=30m
```

Output:

```
cluster.cluster.x-k8s.io/preprovisioned-example condition met
```

⁹⁰⁹ <https://docs.docker.com/docker-hub/download-rate-limit/>

3. If your cluster is in an air-gapped environment or you have a local registry, you must provide additional arguments when creating the cluster. These tell the cluster where to locate the local registry to use by defining the URL.

```
export REGISTRY_URL="<https/http>://<registry-address>:<registry-port>"
export REGISTRY_CA="<path to the CA on the bastion>"
export REGISTRY_USERNAME="<username>"
export REGISTRY_PASSWORD="<password>"
```

- `REGISTRY_URL` : the address of an existing container registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
- `REGISTRY_CA` : (optional) the path on the bastion machine to the container registry CA. Konvoy will configure the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
- `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
- `REGISTRY_PASSWORD` : optional if username is not set.

```
dkp create cluster preprovisioned --cluster-name ${CLUSTER_NAME} \
--control-plane-endpoint-host <control plane endpoint host> \
--control-plane-endpoint-port <control plane endpoint port, if different than 6443> \
--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
--registry-mirror-password=${REGISTRY_PASSWORD}
```

NOTE: Depending on the cluster size, it will take a few minutes to create.

4. After the creation, use this command to get the Kubernetes kubeconfig for the new cluster and begin deploying workloads:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```



Azure requires changing the CNI encapsulation type of Calico from the default of IPtoIP to VXlan. If changing the [Calico encapsulation](#) (see [page 1231](#)), D2iQ recommends changing it after cluster creation, but before production.

7.15.9.2 Audit logs

To modify Control Plane Audit logs settings using the information contained in the page [Configuring the Control Plane](#) (see [page 1538](#)).

7.15.9.2.1 Further Optional Steps:

- [Azure Only Configurations](#) (see page 1227)
- [Modify the Calico installation](#) (see page 1231)

7.15.9.2.2 Next Step:

i When you complete this procedure, move on to [Pre-provisioned: Make Cluster Self-managed](#) (see page 1240) to continue the process.

7.15.9.3 Pre-provisioned Air-gapped New Cluster

If your cluster is in an air-gapped environment or you have a local registry, you must provide additional arguments when creating the cluster. These tell the cluster where to locate the local registry to use by defining the URL.

```
export REGISTRY_URL="https/http://<registry-address>:<registry-port>"
export REGISTRY_CA="path to the CA on the bastion"
export REGISTRY_USERNAME="username"
export REGISTRY_PASSWORD="password"
```


- `REGISTRY_URL` : the address of an existing local registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
- `REGISTRY_CA` : (optional) the path on the bastion machine to the local registry CA. Konvoy will configure the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
- `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
- `REGISTRY_PASSWORD` : optional if username is not set.

i Before you create a new DKP cluster below, you may choose an external load balancer or virtual IP and use the corresponding `dkp create cluster` command example from that page in the docs from the links below. Other customizations are available, but require different flags during `dkp create cluster` command also. Refer to corresponding section for custom cluster creation:

- [Virtual IP](#) (see page 1236)
- [Provision on the Flatcar Linux OS](#) (see page 1236)

- [Use an HTTP proxy \(see page 1237\)](#)
- [Use alternate pod or service subnets \(see page 1238\)](#)
- [Alternative Mirror \(see page 1536\)](#)

7.15.9.3.1 Name Your Cluster

 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)⁹¹⁰ for more naming information.


When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.

By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single Availability Zone. You may create additional node pools in other Availability Zones with the `dkp create nodepool` command.

Follow these steps:

1. Give your cluster a unique name suitable for your environment.
2. Set the environment variable:

```
export CLUSTER_NAME=preprovisioned-example
```

 Before you create a new DKP cluster below, choose an external load balancer or [virtual IP \(see page 1236\)](#) and use the corresponding `dkp create cluster` command.

7.15.9.3.2 Create an Air-gapped Kubernetes Cluster

Once you've defined the infrastructure and control plane endpoints, you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster.

⁹¹⁰ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

NOTE: (Optional) If you have [overrides for your clusters](#) (see page 1226), you must specify the secret as part of the create cluster command. If these are not specified, the overrides for your nodes will not be applied.

⚠️ DKP uses `localvolumeprovisioner` as the [default storage provider](#) (see page 91) for a pre-provisioned environment. However, `localvolumeprovisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)⁹¹¹ compatible storage that is suitable for production.

After disabling `localvolumeprovisioner`, you can choose from any of the storage options available for Kubernetes. To make that storage the default storage, use the commands shown in this section of the Kubernetes documentation: <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

The following command relies on the pre-provisioned cluster API infrastructure provider to initialize the Kubernetes control plane and worker nodes on the hosts defined in the inventory.

- **NOTE:** When specifying the `cluster-name`, you must use the same `cluster-name` as used when defining your inventory objects.
- **NOTE:** (Optional) If you have [overrides for your clusters](#)⁹¹², you must specify the secret as part of the create cluster command. If these are not specified, the overrides for your nodes will not be applied.
`--override-secret-name=$CLUSTER_NAME-user-overrides`
- **NOTE:** To increase [Docker Hub's rate limit](#)⁹¹³ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.
- **NOTE:** Ensure your [subnets](#) (see page 1238) do not overlap with your host subnet because they cannot be changed after cluster creation. If you need to change the kubernetes subnets, you must do this at cluster creation. The default subnets used in DKP are:

```
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
        - 192.168.0.0/16
    services:
      cidrBlocks:
        - 10.96.0.0/12
```

911 <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

912 <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29918736>

913 <https://docs.docker.com/docker-hub/download-rate-limit/>

Create cluster command may take a while depending on the size of the cluster:

1. The command below uses the default external load balancer:

```
dkp create cluster preprovisioned --cluster-name ${CLUSTER_NAME}
  --control-plane-endpoint-host <control plane endpoint host>
  --control-plane-endpoint-port <control plane endpoint port, if different than
  6443>
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml
  --ssh-private-key-file <path-to-ssh-private-key>
  --registry-mirror-url=${REGISTRY_URL} \
  --registry-mirror-cacert=${REGISTRY_CA} \
  --registry-mirror-username=${REGISTRY_USERNAME} \
  --registry-mirror-password=${REGISTRY_PASSWORD}
```

- a. **Virtual IP ALTERNATIVE** - if you don't have an external LB, and wish to use a VIRTUAL IP provided by kube-vip, specify these flags example below:

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host 196.168.1.10 \
  --virtual-ip-interface eth1
```

The output from this command is shortened here for reading clarity, but should start like this:

```
Generating cluster resources
cluster.cluster.x-k8s.io/preprovisioned-example created
cont.....
```

2. Use the wait command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
  timeout=30m
```

Output:


```
cluster.cluster.x-k8s.io/preprovisioned-example condition met
```



NOTE: Depending on the cluster size, it will take a few minutes to create.

When the command completes, you will have a running Kubernetes cluster! Use this command to get the Kubernetes `kubeconfig` for the new cluster and proceed to installing the DKP Kommander UI:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

-  Azure requires changing the CNI encapsulation type of Calico from the default of IPtoIP to VXlan. If changing the [Calico encapsulation](#) (see page 1231), D2iQ recommends changing it after cluster creation, but before production.

7.15.9.3.3 Audit logs

To modify Control Plane Audit logs settings using the information contained in the page [Configure the Control Plane](#) (see page 1538).

7.15.9.3.3.1 Further Optional Steps:

- [Azure Only Configurations](#) (see page 1227)
- [Modify the Calico installation](#) (see page 1231)

7.15.9.3.3.2 Cluster Verification

If you want to **monitor** or verify the installation of your clusters, refer to [Verify your Cluster and DKP Installation](#) (see page 1547).

7.15.9.3.3.3 Next Step:

[Pre-provisioned: Make Cluster Self-managed](#) (see page 1240)

7.15.10 Pre-provisioned Azure only Configurations

After your bootstrap is running and your cluster is created, you will need to install the [Azure Disk CSI Driver](#)⁹¹⁴ on your pre-provisioned Azure Kubernetes cluster. The DKP pre-provisioned provider installs by default the [storage-local-static-provisioner](#)⁹¹⁵ CSI driver, which is not suitable for production environments. For this reason, it needs to be replaced by the [Azure Disk CSI Driver](#)⁹¹⁶.

⁹¹⁴ <https://github.com/kubernetes-sigs/azuredisk-csi-driver/tree/master>

⁹¹⁵ <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

⁹¹⁶ <https://github.com/kubernetes-sigs/azuredisk-csi-driver/blob/master/docs/install-azuredisk-csi-driver.md>

7.15.10.1 Prerequisites:

Before you begin using DKP you must have:

- An x86_64-based Linux or macOS machine.
- [Download \(see page 78\)](#) the `dkp` binary for Linux, or macOS. To check which version of DKP you installed for compatibility reasons, run the `dkp version` command ([dkp version \(see page 1749\)](#)).
- A container engine:
 - Version 4.0 of [Podman](#)⁹¹⁷ or higher for Linux
 - Version 18.09.2 of [Docker](#)⁹¹⁸ or higher for Linux or MacOS
- [kubecti](#)⁹¹⁹ for interacting with the running cluster.
- [Azure CLI](#)⁹²⁰.
- A valid Azure account with [credentials configured](#)⁹²¹.
- Create a custom Azure image using [KIB \(see page 1477\)](#).
- **For air-gapped environments only** -
 - Ability to download artifacts from the internet and then copy those onto your bastion machine.
 - [Download \(see page 78\)](#) the Complete DKP Air-gapped Bundle for this release - `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`.
 - An existing [local registry \(see page 1533\)](#) to seed the air-gapped environment.



On macOS, Docker runs in a virtual machine. Configure this virtual machine with at least 8GB of memory.

7.15.10.2 Set Environment Variables with Credentials:

An Azure Service Principal is needed for deploying resources. To [configure your Azure environment](#)⁹²², follow below:

1. Log in to Azure:

⁹¹⁷ <https://podman.io/getting-started/installation>

⁹¹⁸ <https://www.docker.com/>

⁹¹⁹ <https://kubernetes.io/docs/tasks/tools/#kubecti>

⁹²⁰ <https://docs.microsoft.com/en-us/cli/azure/install-azure-cli>

⁹²¹ <https://github.com/kubernetes-sigs/cluster-api-provider-azure/blob/master/docs/book/src/topics/getting-started.md#prerequisites>

⁹²² <https://github.com/kubernetes-sigs/cluster-api-provider-azure/blob/main/docs/book/src/topics/getting-started.md#setting-up-your-azure-environment>

```
az login
```

```
[
  {
    "cloudName": "AzureCloud",
    "homeTenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "id": "b1234567-abcd-11a1-a0a0-1234a5678b90",
    "isDefault": true,
    "managedByTenants": [],
    "name": "Mesosphere Developer Subscription",
    "state": "Enabled",
    "tenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "user": {
      "name": "user@azuresphere.onmicrosoft.com",
      "type": "user"
    }
  }
]
```

2. Create an Azure Service Principal (SP) by running the following command:

Note: If an SP with the name exists, this command will rotate the password.

```
az ad sp create-for-rbac --role contributor --name "$(whoami)-konvoy" --scopes=/subscriptions/$(az account show --query id -o tsv)
```

```
{
  "appId": "7654321a-1a23-567b-b789-0987b6543a21",
  "displayName": "azure-cli-2021-03-09-23-17-06",
  "password": "Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C",
  "tenant": "a1234567-b132-1234-1a11-1234a5678b90"
}
```

For air-gapped environments, you need to create a [resource management private link](#)⁹²³ with a private endpoint to ensure the Azure CSI driver will run correctly in further steps. Private links enable you to access Azure services over a private endpoint in your virtual network.

To set up a private link resource, use the following process.

- a. Create the [resource management private link](#)⁹²⁴ using Azure CLI.

⁹²³ <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/create-private-link-access-portal>

⁹²⁴ <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/create-private-link-access-commands?tabs=azure-cli#create-resource-management-private-link>

- b. Create a [private link association](#)⁹²⁵ for the root management group which also references the resource ID for the resource management private link.
 - c. Add a private endpoint that references the resource management private link using the [Azure Documentation](#)⁹²⁶.
3. Set the required environment variables using that output:

```
export AZURE_SUBSCRIPTION_ID="<id>"           # b1234567-abcd-11a1-
a0a0-1234a5678b90
export AZURE_TENANT_ID="<tenant>"           # a1234567-b132-1234-1a11-1234a5678b9
0
export AZURE_CLIENT_ID="<appId>"           # 7654321a-1a23-567b-
b789-0987b6543a21
export AZURE_CLIENT_SECRET="<password>"     # Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C
export AZURE_RESOURCE_GROUP="<resource group name>" # set to the name of the
resource group
export AZURE_LOCATION="westus"           # set to the location you are using
```

4. Set your KUBECONFIG environment variable:

```
export kubeconfig=${CLUSTER_NAME}.conf
```

5. Create the Secret with the Azure credentials, this will be used by the Azure CSI driver:
 - a. Create an `azure.json` file:

```
cat <<EOF > azure.json
{
  "cloud": "AzurePublicCloud",
  "tenantId": "$AZURE_TENANT_ID",
  "subscriptionId": "$AZURE_SUBSCRIPTION_ID",
  "aadClientId": "$AZURE_CLIENT_ID",
  "aadClientSecret": "$AZURE_CLIENT_SECRET",
  "resourceGroup": "$AZURE_RESOURCE_GROUP",
  "location": "$AZURE_LOCATION"
}
EOF
```

- b. Create the Secret:

```
kubectl create secret generic azure-cloud-provider --namespace=kube-
system --type=Opaque --from-file=cloud-config=azure.json
```

6. Install the Azure Disk CSI driver:

⁹²⁵ <https://learn.microsoft.com/en-us/azure/azure-resource-manager/management/create-private-link-access-commands?tabs=azure-cli#create-private-link-association>

⁹²⁶ <https://learn.microsoft.com/en-us/azure/private-link/create-private-endpoint-cli?tabs=dynamic-ip>

```
$ curl -skSL https://raw.githubusercontent.com/kubernetes-sigs/azuredisk-csi-driver/v1.26.2/deploy/install-driver.sh | bash -s v1.26.2 snapshot -
```

7. Check the status to see if the driver is ready for use:

```
kubectl -n kube-system get pod -o wide --watch -l app=csi-azuredisk-controller
kubectl -n kube-system get pod -o wide --watch -l app=csi-azuredisk-node
```

8. Now Kubernetes knows that this is Azure disk, and will create clusters on Azure. You are ready to create the StorageClass for the Azure Disk CSI Driver:

```
kubectl create -f https://raw.githubusercontent.com/kubernetes-sigs/azuredisk-csi-driver/master/deploy/example/storageclass-azuredisk-csi.yaml
```

9. Change the default storage class to this new StorageClass so that every new disk will be created in the Azure environment:

```
kubectl patch sc/localvolumeprovisioner -p '{"metadata": {"annotations": {"storageclass.kubernetes.io/is-default-class":"false"}}}'
kubectl patch sc/managed-csi -p '{"metadata": {"annotations": {"storageclass.kubernetes.io/is-default-class":"true"}}}'
```

10. Verify that the StorageClass chosen is currently the default:

```
kubectl get storageclass
```

For more information about Azure Disk CSI for persistent storage and changing the default StorageClass, refer to that page in the documentation: [Default Storage Providers in DKP \(see page 91\)](#)

7.15.10.3 Next Step:

[Pre-provisioned Modify the Calico Installation \(see page 1231\)](#)

7.15.11 Pre-provisioned Modify the Calico Installation

7.15.11.1 Set the Interface

By default, Calico automatically detects the IP address to use for each node using the `first-found method`⁹²⁷. This is not always appropriate for your particular nodes. In that case, you must modify Calico's

⁹²⁷ <https://projectcalico.docs.tigera.io/reference/node/configuration#ip-autodetection-methods>

configuration to use a different method. An alternative is to use the `interface` method by providing the interface ID. Follow the steps outlined in this section to modify Calico's configuration.

ⓘ Azure does **not** set the interface. Proceed to [Change the Encapsulation Type](#) (see page 1234) section below.

In this example, all cluster nodes use `ens192` as the interface name.

1. Get the pods running on your cluster with this command:

```
kubectl get pods -A --kubeconfig ${CLUSTER_NAME}.conf
```

```


NAMESPACE          NAME
READY  STATUS      RESTARTS      AGE
calico-system      calico-kube-controllers-57fbd7bd59-vpn8b
1/1    Running      0              16m
calico-system      calico-node-5tbvl
1/1    Running      0              16m
calico-system      calico-node-nbdwd
1/1    Running      0              4m40s
calico-system      calico-node-tw6b
0/1    PodInitializing  0              9s
calico-system      calico-node-wtkh
1/1    Running      0              5m35s
calico-system      calico-typha-54f46b998d-52pt2
1/1    Running      0              16m
calico-system      calico-typha-54f46b998d-9tzb8
1/1    Running      0              4m31s
default            cuda-vectoradd
0/1    Pending      0              0s
kube-system        coredns-78fcd69978-frwx4
1/1    Running      0              16m
kube-system        coredns-78fcd69978-kkf44
1/1    Running      0              16m
kube-system        etcd-ip-10-0-121-16.us-west-2.compute.internal
0/1    Running      0              8s
kube-system        etcd-ip-10-0-46-17.us-west-2.compute.internal
1/1    Running      1              16m
kube-system        etcd-ip-10-0-88-238.us-west-2.compute.internal
1/1    Running      1              5m35s
kube-system        kube-apiserver-ip-10-0-121-16.us-west-2.compute.internal
0/1    Running      6              7s
kube-system        kube-apiserver-ip-10-0-46-17.us-west-2.compute.internal
1/1    Running      1              16m

```

```

kube-system          kube-apiserver-ip-10-0-88-238.us-west-2.compute.internal
1/1      Running          1              5m34s
kube-system          kube-controller-manager-ip-10-0-121-16.us-west-2.compute.int
ernal  0/1      Running          0              7s
kube-system          kube-controller-manager-ip-10-0-46-17.us-west-2.compute.inte
rnal  1/1      Running          1 (5m25s ago)  15m
kube-system          kube-controller-manager-ip-10-0-88-238.us-west-2.compute.int
ernal  1/1      Running          0              5m34s
kube-system          kube-proxy-gclmt
1/1      Running          0              16m
kube-system          kube-proxy-gptd4
1/1      Running          0              9s
kube-system          kube-proxy-mwkg1
1/1      Running          0              4m40s
kube-system          kube-proxy-zcqx4
1/1      Running          0              5m35s
kube-system          kube-scheduler-ip-10-0-121-16.us-west-2.compute.internal
0/1      Running          1              7s
kube-system          kube-scheduler-ip-10-0-46-17.us-west-2.compute.internal
1/1      Running          3 (5m25s ago)  16m
kube-system          kube-scheduler-ip-10-0-88-238.us-west-2.compute.internal
1/1      Running          1              5m34s
kube-system          local-volume-provisioner-2mv7z
1/1      Running          0              4m10s
kube-system          local-volume-provisioner-vcrg
1/1      Running          0              4m53s
kube-system          local-volume-provisioner-wsjrt
1/1      Running          0              16m
node-feature-discovery node-feature-discovery-master-84c67dccb6-m78vr
1/1      Running          0              16m
node-feature-discovery node-feature-discovery-worker-ppvpl
1/1      Running          0              4m10s
tigera-operator      tigera-operator-d499f5c8f-79dc4
1/1      Running          1 (5m24s ago)  16m

```

 If a `calico-node` pod is not ready on your cluster, you must edit the `default` Installation resource. To edit the Installation resource, run the command:

```
kubectl edit installation default --kubeconfig ${CLUSTER_NAME}.conf
```

2. Change the value for `spec.calicoNetwork.nodeAddressAutodetectionV4` to `interface: ens192`, and save the resource:

```
spec:
  calicoNetwork:
```

```
...
nodeAddressAutodetectionV4:
  interface: ens192
```

3. Save this resource. You may need to delete the node feature discovery worker pod in the `node-feature-discovery` namespace if that pod has failed. After you delete it, Kubernetes replaces the pod as part of its normal reconciliation.

7.15.11.2 Change the Encapsulation Type

Calico can leverage different network encapsulation methods to route traffic for your workloads. Encapsulation is useful when running on top of an underlying network that is not aware of workload IPs. Common examples of this include:

- Public cloud environments where you don't own the hardware.
- AWS across VPC subnet boundaries.
- Environments where you cannot peer Calico over BGP to the underlay or easily configure static routes.



WARNING: Switching encapsulation modes can cause disruption to in-progress connections. You can do this safely when the cluster is first deployed. However, if user workloads are already running on the cluster, plan accordingly for interruption.

7.15.11.3 Provider Specific Settings

The encapsulation type to be used for networking depends on the cloud provider. IP-in-IP is the default encapsulation method for Calico which most providers use, but not Azure.



Azure only supports **VXLAN** encapsulation type. Therefore, if you install on Azure pre-provisioned VMs, you must set the encapsulation mode to VXLAN.

D2iQ recommends changing the encapsulation type after cluster creation but before production using the method below. To change the encapsulation type, follow these steps:

1. First, remove the existing `default-ipv4-ippool` IPPool resource from `kubeconfig`. The resource must be deleted, so that it can re-create after you edit the Installation resource. Execute the command below to delete:

```
kubectl delete ippool default-ipv4-ippool
```

2. Run the following command to edit:

```
kubectl edit installation default --kubeconfig ${CLUSTER_NAME}.conf
```

3. Change the value for encapsulation – `encapsulation:` as shown below:

```
spec:
  calicoNetwork:
    ipPools:
      - encapsulation: VXLAN
```

7.15.11.3.1 VXLAN

VXLAN is a tunneling protocol that encapsulates layer 2 Ethernet frames in UDP packets, enabling you to create virtualized layer 2 subnets that span Layer 3 networks. It has a slightly larger header than IP-in-IP which creates a slight reduction in performance over IP-in-IP.

7.15.11.3.2 IPIP

IP-in-IP is an IP tunneling protocol that encapsulates one IP packet in another IP packet. An outer packet header is added with the tunnel entry point and the tunnel exit point. The calico implementation of this protocol uses BGP to determine the exit point making this protocol unusable on networks that don't pass BGP.

For more information, see:

- [Calico Overlay Networking](#)⁹²⁸
- [Calico Routing for VXLAN](#)⁹²⁹
- [IP-in-IP RFC 2003](#)⁹³⁰
- [VXLAN RFC 7348](#)⁹³¹



If using Windows, see this documentation on Calico site regarding limitations: [Calico for Windows VXLAN](#)⁹³²

928 <https://docs.projectcalico.org/networking/vxlan-ipip>

929 <https://joshrosso.com/docs/2020/2020-10-01-calico-routing-modes/>

930 <https://datatracker.ietf.org/doc/html/rfc2003>

931 <https://datatracker.ietf.org/doc/html/rfc7348>

932 <https://projectcalico.docs.tigera.io/getting-started/windows-calico/limitations#calico-vxlan-networking-limitations>

7.15.12 Pre-provisioned Built-in Virtual IP

As explained in [Define the Control Plane Endpoint](#) (see page 152), we recommend using an external load balancer for the control plane endpoint, but provide a built-in virtual IP when an external load balancer is not available. The built-in virtual IP uses the [kube-vip](#)⁹³³ project. To use the virtual IP, add these flags to the `create cluster` command:

Virtual IP Configuration	Flag
Network interface to use for Virtual IP. <i>Must exist on all control plane machines.</i>	<code>--virtual-ip-interface string</code>
IPv4 address. <i>Reserved for use by the cluster.</i>	<code>--control-plane-endpoint string</code>

7.15.12.1 Virtual IP example

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-endpoint-host 196.168.1.10 \
  --virtual-ip-interface eth1
  --pre-provisioned-inventory-file preprovisioned_inventory.yaml
  --ssh-private-key-file <path-to-ssh-private-key>
  --self-managed
```

7.15.13 Provision on the Flatcar Linux OS

When provisioning onto the Flatcar Container Linux distribution, you must instruct the bootstrap cluster to make some changes related to the installation paths. To accomplish this, add the `--os-hint flatcar` flag to the above `create cluster` command.

7.15.13.1 Flatcar Linux example

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --os-hint flatcar
```

⁹³³ <https://kube-vip.io/>

- For provisioning DKP on Flatcar, DKP configures cluster nodes to use [Control Groups \(cgroups\) version 1](#)⁹³⁴. In versions prior to Flatcar 3033.3.x, a restart is required in order to apply the changes to the kernel. For more information, refer to the [Flatcar documentation](#)⁹³⁵.

7.15.14 Pre-provisioned Use HTTP Proxy

If you require HTTP proxy configurations, you can apply them during the `create` operation by adding the appropriate flags to the `create cluster` command:

Proxy configuration	Flag
HTTP proxy for control plane machines	<code>--control-plane-http-proxy string</code>
HTTPS proxy for control plane machines	<code>--control-plane-https-proxy string</code>
No Proxy list for control plane machines	<code>--control-plane-no-proxy strings</code>
HTTP proxy for worker machines	<code>--worker-http-proxy string</code>
HTTPS proxy for worker machines	<code>--worker-https-proxy string</code>
No Proxy list for worker machines	<code>--worker-no-proxy strings</code>

- You must also add the same configuration as an [override](#) (see page 1524). For more information, refer to [this documentation](#) (see page 1520).

⁹³⁴ <https://www.kernel.org/doc/html/latest/admin-guide/cgroup-v1/cgroups.html#what-are-cgroups>

⁹³⁵ <https://www.flatcar.org/docs/latest/container-runtimes/switching-to-unified-cgroups/#starting-new-nodes-with-legacy-cgroups>

7.15.14.1 HTTP proxy example

- = To increase [Docker Hub's rate limit](#)⁹³⁶ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --control-plane-http-proxy http://proxy.example.com:8080 \
  --control-plane-https-proxy https://proxy.example.com:8080 \
  --control-plane-no-proxy
"127.0.0.1,10.96.0.0/12,192.168.0.0/16,kubernetes,kubernetes.default.svc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluster.local" \
  --worker-http-proxy http://proxy.example.com:8080 \
  --worker-https-proxy https://proxy.example.com:8080 \
  --worker-no-proxy
"127.0.0.1,10.96.0.0/12,192.168.0.0/16,kubernetes,kubernetes.default.svc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluster.local"
```

- i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

7.15.15 Pre-provisioned Use Alternate Pod or Service Subnets

In Konvoy, the default pod subnet is 192.168.0.0/16, and the default service subnet is 10.96.0.0/12. If you wish to change the subnets you can do so with the following steps:

1. Generate the YAML manifests for the cluster using the `--dry-run` and `-o yaml` flags, along with the desired `dkp cluster create` command:

⁹³⁶ <https://docs.docker.com/docker-hub/download-rate-limit/>

```
dkp create cluster preprovisioned --cluster-name ${CLUSTER_NAME} --control-
plane-endpoint-host <control plane endpoint host> --control-plane-endpoint-port
<control plane endpoint port, if different than 6443> --dry-run -o yaml >
cluster.yaml
```

- To modify the service subnet, add or edit the `spec.clusterNetwork.services.cidrBlocks` field of the `Cluster` object:

```
kind: Cluster
spec:
  clusterNetwork:
    services:
      cidrBlocks:
        - 10.0.0.0/18
```

- To modify the pod subnet, edit the `Cluster` and `calico-cni ConfigMap` resources:

`Cluster`: Add or edit the `spec.clusterNetwork.pods.cidrBlocks` field:

```
kind: Cluster
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
        - 172.16.0.0/16
```

`ConfigMap`: Edit the `data."custom-resources.yaml".spec.calicoNetwork.ipPools.cidr` field with your desired pod subnet:

```
apiVersion: v1
data:
  custom-resources.yaml: |
    apiVersion: operator.tigera.io/v1
    kind: Installation
    metadata:
      name: default
    spec:
      # Configures Calico networking.
      calicoNetwork:
        # Note: The ipPools section cannot be modified post-install.
        ipPools:
          - blockSize: 26
            cidr: 172.16.0.0/16
kind: ConfigMap
metadata:
  name: calico-cni-<cluster-name>
```


When you provision the cluster, the configured pod and service subnets will be applied.

7.15.16 Pre-provisioned: Make Cluster Self-managed


Make the new Kubernetes cluster manage itself

Konvoy deploys all cluster lifecycle services to a bootstrap cluster, which then deploys a workload cluster. When the workload cluster is ready, move the cluster lifecycle services to the workload cluster, which is now self-managed. This guide describes how to make a workload cluster self-managed.

This page contains instructions on how to make your cluster self-managed. This is necessary if there is only one cluster in your environment, or if this cluster should become the Management cluster in a multi-cluster environment.

 If you already have a self-managed or Management cluster in your environment, skip this page.

7.15.16.1 Make the new Kubernetes cluster manage itself

 If you have not already retrieved the kubeconfig after creating the cluster, use this command before proceeding: `dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf`

1. Deploy cluster lifecycle services on the workload cluster:

```
dkp create capi-components --kubeconfig ${CLUSTER_NAME}.conf
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1042\)](#).

```
✓ Initializing new CAPI components
```

2. Move the Cluster API objects from the bootstrap to the workload cluster:

The cluster lifecycle services on the workload cluster are ready, but the workload cluster configuration is on the bootstrap cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the bootstrap to the workload cluster. This process is also called a [Pivot](#)⁹³⁷. First unset the kubeconfig and then move the CAPI:

```
unset KUBECONFIG
```

```
dkp move capi-resources --to-kubeconfig ${CLUSTER_NAME}.conf
```

Output:

```
✓ Moving cluster resources
You can now view resources in the moved cluster by using the --kubeconfig flag
with kubectl. For example: kubectl --kubeconfig=preprovisioned-example.conf get
nodes
```

NOTE: To ensure only one set of cluster lifecycle services manages the workload cluster, Konvoy first pauses reconciliation of the objects on the bootstrap cluster, then creates the objects on the workload cluster. As Konvoy copies the objects, the cluster lifecycle services on the workload cluster reconcile the objects. The workload cluster becomes self-managed after Konvoy creates all the objects. If it fails, the `move` command can be safely retried.

3. Wait for the cluster control-plane to be ready:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --timeout=20m
```

```
cluster.cluster.x-k8s.io preprovisioned-example condition met
```

4. Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

NOTE: After moving the cluster lifecycle services to the workload cluster, remember to use Konvoy with the workload cluster kubeconfig.

```
dkp describe cluster --kubeconfig ${CLUSTER_NAME}.conf -c ${CLUSTER_NAME}
```

⁹³⁷ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```

NAME
READY SEVERITY REASON SINCE MESSAGE
Cluster/preprovisioned-example True
2m31s
├─ClusterInfrastructure - PreprovisionedCluster/preprovisioned-example
├─ControlPlane - KubeadmControlPlane/preprovisioned-example-control-plane True
2m31s
| ├─Machine/preprovisioned-example-control-plane-6g6nr True
2m33s
| ├─Machine/preprovisioned-example-control-plane-8lhcv True
2m33s
| └─Machine/preprovisioned-example-control-plane-kk2kg True
2m33s
└─Workers
    └─MachineDeployment/preprovisioned-example-md-0 True
2m34s
        └─Machine/preprovisioned-example-md-0-77f667cd9-tnctd True
2m33s


```

- Remove the bootstrap cluster, as the workload cluster is now self-managed:

```
dkp delete bootstrap
```

```
✓ Deleting bootstrap cluster
```

7.15.16.1.1 Known limitations

 Be aware of these limitations in the current release of Konvoy.

- DKP supports moving only one set of cluster objects from the bootstrap cluster to the workload cluster, or vice-versa.
- DKP only supports moving all namespaces in the cluster; DKP does not support migration of individual namespaces.

7.15.17 Pre-provisioned Configure MetalLB

7.15.17.1 Create a MetalLB configmap for your pre-provisioned infrastructure.

Choose one of the following two protocols you want to use to announce service IPs. If your environment is not currently equipped with a load balancer, you can use MetalLB. Otherwise, your own load balancer will work and you can continue the installation process.

To use MetalLB, create a MetalLB configMap for your Pre-provisioned infrastructure. MetalLB uses one of two protocols for exposing Kubernetes services:

- Layer 2, with Address Resolution Protocol (ARP)
- Border Gateway Protocol (BGP)

Select one of the following procedures to create your MetalLB manifest for further editing.

7.15.17.2 Layer 2 Configuration

Layer 2 mode is the easiest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.



- MetalLB IP address ranges/CIDRs should be within the node's primary network [subnet](#) (see [page 1049](#)).
- MetalLB IP address ranges/CIDRs and node subnet should not conflict with the Kubernetes cluster pod and service subnets.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:



The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
```

```

config: |
  address-pools:
  - name: default
    protocol: layer2
    addresses:
    - 192.168.1.240-192.168.1.250
EOF

```

After completion, run the following `kubectl` command.

```
kubectl apply -f metallb-conf.yaml
```

7.15.17.3 BGP Configuration

For a basic configuration featuring one BGP router and one IP address range, you need 4 pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,
- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:



The following values are generic, enter your specific values into the fields where applicable.

```

cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
    - peer-address: 10.0.0.1
      peer-asn: 64501
      my-asn: 64500
    address-pools:
    - name: default
      protocol: bgp
      addresses:
      - 192.168.10.0/24
EOF

```

After completion, run the following `kubectl` command.

```
kubectl apply -f metallb-conf.yaml
```

7.15.18 Pre-provisioned Create and Delete Node Pools

Node pools are part of a cluster and managed as a group, and can be used to manage a group of machines using common properties. New default clusters created by Konvoy contain one node pool of worker nodes that have the same configuration.

You can create additional node pools for specialized hardware or other configurations. For example, if you want to tune your memory usage on a cluster where you need maximum memory for some machines and minimal memory on others, you could create a new node pool with those specific resource needs.

NOTE: Konvoy implements node pools using Cluster API [MachineDeployments](#)⁹³⁸.

7.15.18.1 Create a Pre-provisioned node pool

Follow these steps:

1. Create an inventory object that has the same name as the node pool you're creating, and the details of the pre-provisioned machines that you want to add to it. For example, to create a node pool named `gpu-nodepool` an inventory named `gpu-nodepool` must be present in the same namespace:

```
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: ${MY_NODEPOOL_NAME}
spec:
  hosts:
    - address: ${IP_OF_NODE}
  sshConfig:
    port: 22
    user: ${SSH_USERNAME}
  privateKeyRef:
    name: ${NAME_OF_SSH_SECRET}
    namespace: ${NAMESPACE_OF_SSH_SECRET}
```

⁹³⁸ <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

- (Optional) If your pre-provisioned machines have [overrides](#) (see page 1513), you must create a secret that includes all of the overrides you want to provide in one file. Create an [override secret](#) (see page 1216) using the instructions detailed on this page.
- Once the `PreprovisionedInventory` object and overrides are created, create a node pool:

```
dkp create nodepool preprovisioned -c ${MY_CLUSTER_NAME} ${MY_NODEPOOL_NAME} --
  override-secret-name ${MY_OVERRIDE_SECRET}
```

- Advanced users can use a combination of the `--dry-run` and `--output=yaml` or `--output-directory=<existing-directory>` flags to get a complete set of node pool objects to modify locally or store in version control.



For more information regarding this flag or others, please refer to the [dkp create nodepool](#) (see page 1684) section of the documentation for either cluster or nodepool and select your provider.

7.15.18.2 Delete a node pool

Deleting a node pool deletes both the Kubernetes nodes and their underlying infrastructure. DKP drains all nodes prior to deletion and reschedules the pods running on those nodes.

To delete a node pool from a managed cluster, run the following command:

```
dkp delete nodepool ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME}
```

The expected output is similar to the following example, indicating the `example` node pool is being deleted:

```
INFO[2021-07-28T17:14:26-07:00] Running nodepool delete command
Nodepool=example clusterName=d2iq-e2e-cluster-1 managementClusterKubeconfig=
namespace=default src="nodepool/delete.go:80"
```

Deleting an invalid node pool results in an output similar to this example command output:

```
dkp delete nodepool ${CLUSTER_NAME}-md-invalid --cluster-name=${CLUSTER_NAME}

INFO[2021-07-28T17:11:44-07:00] Running nodepool delete command
Nodepool=demo-cluster-md-invalid clusterName=d2iq-e2e-cluster-1
managementClusterKubeconfig= namespace=default src="nodepool/delete.go:80"
Error: failed to get nodepool with name demo-cluster-md-invalid in namespace default
: failed to get nodepool with name demo-cluster-md-invalid in namespace default :
machedeployments.cluster.x-k8s.io "demo-cluster-md-invalid" not found
```

7.15.19 Pre-provisioned Add Nodes to Existing Node Pool

This section covers the prerequisites and procedure you need to scale-up or scale-down nodes in an existing DKP cluster.

7.15.19.1 Prerequisites

- You must have the bootstrap node running with the SSH key/secrets created.
- The export values in the environment variables section should contain the addresses of the nodes that you need to add [Pre-provisioned: Define Infrastructure](#) (see page 151).
- Update the `preprovisioned_inventory.yaml` with the new host addresses.
- Run the `kubectl apply` command.

7.15.19.2 Scale up a Cluster Node

Follow these steps:

1. Fetch the existing `preprovisioned_inventory`:

```
$ kubectl get preprovisionedinventory
```

2. Edit the `preprovisioned_inventory` to add additional IPs needed for additional worker nodes in the `spec.hosts` section:

```
$ kubectl edit preprovisionedinventory <preprovisioned_inventory> -n default
```

3. Add any additional IPs that you require:

```
spec:
  hosts:
    - address: <worker.ip.add.1>
    - address: <worker.ip.add.2>
```

After you edit `preprovisioned_inventory`, fetch the machine deployment. The naming convention with `md` means that it is for worker machines.

For example:

```
$ kubectl --kubeconfig ${CLUSTER_NAME}.conf get machinedeployment
```

NAME	CLUSTER	AGE	PHASE	REPLICAS	READY
UPDATED UNAVAILABLE machinedeployment-md-0	cluster-name	9m10s	Running	4	4

- Scale the worker node to the required number. In this example we are scaling from 4 to 6 worker nodes:

```
$ kubectl --kubeconfig ${CLUSTER_NAME}.conf scale --replicas=6 machinedeployment
machinedeployment-md-0
machinedeployment-md-0 scaled
```

- Monitor the scaling with this command, by adding `-w` option to watch:

```
$ kubectl --kubeconfig ${CLUSTER_NAME}.conf get machinedeployment -w
```

NAME	CLUSTER	AGE	PHASE	REPLICAS	READY
UPDATED UNAVAILABLE machinedeployment-md-0	cluster-name	20m	ScalingUp	6	4

- Also you can check the machine deployment if it is already scaled. The output should resemble this example:

```
$ kubectl --kubeconfig ${CLUSTER_NAME}.conf get machinedeployment
```

NAME	CLUSTER	AGE	PHASE	REPLICAS	READY
UPDATED UNAVAILABLE machinedeployment-md-0	cluster-name	3h33m	Running	6	6

- Alternately, you can use this command and verify the `NODENAME` column and you should see the additional worker nodes added and in `Running` state:

```
$ kubectl --kubeconfig ${CLUSTER_NAME}.conf get machines -o wide
```

NAME	CLUSTER	AGE	PROVIDERID	PHASE	VERSION	NODENAME
------	---------	-----	------------	-------	---------	----------

7.15.19.3 Scale Down a Cluster Node

Follow these steps

- Run this command on your worker nodes:

```
kubectl scale machinedeployment <machinedeployment-name> --replicas <new
number>
```

2. For control plane nodes, execute the following command:

```
kubectl scale kubeadmcontrolplane ${CLUSTER_NAME}-control-plane --replicas <new
number>
```

7.15.19.3.1 Additional Notes for Scaling Down

It is possible for machines to get stuck in the provisioning stage when you scaling down. You can utilize a delete operation to clear the stale machine deployment:

```
kubectl delete machine ${CLUSTER_NAME}-control-plane-<hash>
```

```
kubectl delete machine <machinedeployment-name>-<hash>
```

7.15.20 GPU Nodepools in a Pre-provisioned Environment

For pre-provisioned environments, DKP has introduced the `nvidia-runfile` flag for Air-gapped Pre-provisioned environments. If the [NVIDIA runfile](#)⁹³⁹ installer has not been downloaded, then retrieve and install the download first by running the following command. The first line in the command below downloads and installs the runfile and the second line places it in the artifacts directory (you must create an `artifacts` directory if it doesn't already exist).

```
curl -O https://download.nvidia.com/XFree86/Linux-x86_64/470.82.01/NVIDIA-Linux-
x86_64-470.82.01.run
mv NVIDIA-Linux-x86_64-470.82.01.run artifacts
```

 DKP supported [NVIDIA driver](#)⁹⁴⁰ version is 470.x.

1. Create the secret that GPU nodepool would use, this secret is populated from the KIB overrides. In this example we have a file called, `overrides/nvidia.yaml`. It should resemble this:

⁹³⁹ <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

⁹⁴⁰ <https://www.nvidia.com/Download/Find.aspx>

```
gpu:
  types:
    - nvidia
  build_name_extra: "-nvidia"
```

2. Create a secret on the bootstrap cluster that is populated from the above file. We will name it

```
`${CLUSTER_NAME}-user-overrides`
```

```
kubectl create secret generic `${CLUSTER_NAME}-user-overrides` --from-
file=overrides.yaml=overrides/nvidia.yaml
```

3. Create an inventory and nodepool with the instructions below and use the ``${CLUSTER_NAME}-user-overrides`` secret.

Follow these steps:


1. Create an inventory object that has the same name as the node pool you're creating, and the details of the pre-provisioned machines that you want to add to it. For example, to create a node pool named `gpu-nodepool` an inventory named `gpu-nodepool` must be present in the same namespace:

```
apiVersion: infrastructure.cluster.konvoy.d2iq.io/v1alpha1
kind: PreprovisionedInventory
metadata:
  name: `${MY_NODEPOOL_NAME}`
spec:
  hosts:
    - address: `${IP_OF_NODE}`
  sshConfig:
    port: 22
    user: `${SSH_USERNAME}`
    privateKeyRef:
      name: `${NAME_OF_SSH_SECRET}`
      namespace: `${NAMESPACE_OF_SSH_SECRET}`
```

2. (Optional) If your pre-provisioned machines have [overrides](#) (see page 1513), you must create a secret that includes all of the overrides you want to provide in one file. Create an [override secret](#) (see page 1216) using the instructions detailed on this page.
3. Once the `PreprovisionedInventory` object and overrides are created, create a node pool:

```
dkp create nodepool preprovisioned -c `${MY_CLUSTER_NAME}` `${MY_NODEPOOL_NAME}` --
override-secret-name `${MY_OVERRIDE_SECRET}`
```


- Advanced users can use a combination of the `--dry-run` and `--output=yaml` or `--output-directory=<existing-directory>` flags to get a complete set of node pool objects to modify locally or store in version control.

 For more information regarding this flag or others, please refer to the [dkp create nodepool](#) (see [page 1684](#)) section of the documentation for either cluster or nodepool and select your provider.

For more information, see:

- [Pre-provisioned Create and Delete Node Pools](#) (see [page 1248](#))
- [Pre-provisioned Air-gapped Define Environment](#) (see [page 1207](#))

7.15.21 Pre-provisioned Delete Cluster

 **NOTE:** A self-managed workload cluster cannot delete itself. If your workload cluster is self-managed, you must create a bootstrap cluster and move the cluster lifecycle services to the bootstrap cluster before deleting the workload cluster.


If you did not make your workload cluster self-managed, as described in [Make New Cluster Self-Managed](#) (see [page 1243](#)), see [Delete the workload cluster](#) (see [page 1253](#)).

7.15.21.1 Prepare to Delete the Pre-provisioned Cluster

Follow these steps:

1. Create a bootstrap cluster:

The bootstrap cluster will host the Cluster API controllers that reconcile the cluster objects marked for deletion:

 **NOTE:** To avoid using the wrong kubeconfig, the following steps use explicit kubeconfig paths and contexts.

```
dkp create bootstrap --kubeconfig $HOME/.kube/config
```

- ✓ Creating a bootstrap cluster
- ✓ Initializing **new** CAPI components

2. Move the Cluster API objects from the workload to the bootstrap cluster: The cluster lifecycle services on the bootstrap cluster are ready, but the workload cluster configuration is on the workload cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the workload to the bootstrap cluster. This process is also called a [Pivot](#)⁹⁴¹.

```
dkp move capi-resources \
  --from-kubeconfig ${CLUSTER_NAME}.conf \
  --from-context ${CLUSTER_NAME}-admin@${CLUSTER_NAME} \
  --to-kubeconfig $HOME/.kube/config \
  --to-context kind-konvoy-capi-bootstrapper
```

✓ Moving cluster resources

You can now view resources in the moved cluster by using the `--kubeconfig` flag with `kubectl`. For example: `kubectl --kubeconfig $HOME/.kube/config get nodes`

3. Use the cluster lifecycle services on the workload cluster to check the workload cluster status by running the following command:

```
dkp describe cluster --kubeconfig $HOME/.kube/config -c ${CLUSTER_NAME}
```

```
NAME
READY SEVERITY REASON SINCE MESSAGE
Cluster/preprovisioned-example True
2m31s
├─ClusterInfrastructure - PreprovisionedCluster/preprovisioned-example
├─ControlPlane - KubeadmControlPlane/preprovisioned-example-control-plane True
2m31s
| ├─Machine/preprovisioned-example-control-plane-6g6nr True
2m33s
| ├─Machine/preprovisioned-example-control-plane-8lhcv True
2m33s
| └─Machine/preprovisioned-example-control-plane-kk2kg True
2m33s
└─Workers
  └─MachineDeployment/preprovisioned-example-md-0 True
2m34s
    └─Machine/preprovisioned-example-md-0-77f667cd9-tnctd True
2m33s
```

⁹⁴¹ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

i NOTE: After moving the cluster lifecycle services to the workload cluster, remember to use `dkp` with the workload cluster kubeconfig.

4. Wait for the cluster control-plane to be ready. Run the command below and wait for the condition to be met:

```
kubectl --kubeconfig $HOME/.kube/config wait --for=condition=controlplaneready
"clusters/${CLUSTER_NAME}" --timeout=20m
```

```
cluster.cluster.x-k8s.io/preprovisioned-example condition met
```

= Persistent Volumes (PVs) are not deleted automatically by design in order to preserve your data. However, they take up storage space if not deleted. You must delete PVs manually. Information for backup of a cluster and PVs is on the page in documentation called [Back up your Cluster's Applications and Persistent Volumes \(see page 835\)](#) .

7.15.21.2 Delete the Workload Cluster

If you have a need to remove the Kubernetes cluster, such as for environment cleanup, use this command to delete the provisioned Kubernetes cluster.

1. To delete a cluster, you would use `dkp delete cluster` and pass in the name of the cluster you are trying to delete with `--cluster-name` flag. You would use `kubectl get clusters` to get those details (`--cluster-name` and `--namespace`) of the Kubernetes cluster to delete it.
NOTE: Do not use `dkp get clusters` since that gets you Kommander cluster details rather than Konvoy kubernetes cluster details.

```
kubectl get clusters
```

2. Delete the Kubernetes cluster and wait a few minutes:

NOTE: Before deleting the cluster, `dkp` deletes all Services of type LoadBalancer on the cluster. Each Service is backed by an AWS Classic ELB. Deleting the Service deletes the ELB that backs it. To skip this step, use the flag `--delete-kubernetes-resources=false` . Do not skip this step if the VPC is managed by DKP. When DKP deletes the cluster, it deletes the VPC. If the VPC has any AWS Classic ELBs, AWS does not allow the VPC to be deleted, and DKP cannot delete the cluster.


```
dkp delete cluster --cluster-name=${CLUSTER_NAME} --kubeconfig $HOME/.kube/
config
```

```
✓ Deleting Services with type LoadBalancer for Cluster default/preprovisioned-
example
✓ Deleting ClusterResourceSets for Cluster default/preprovisioned-example
✓ Deleting cluster resources
✓ Waiting for cluster to be fully deleted
Deleted default/preprovisioned-example cluster
```

7.15.21.2.1 Delete the bootstrap cluster

After you have moved the workload resources back to a bootstrap cluster and deleted the workload cluster, you no longer need the bootstrap cluster. You can safely delete the bootstrap cluster with this command:

Use `dkp` with the bootstrap cluster to delete the workload cluster. Delete the `kind` Kubernetes cluster:

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

```
✓ Deleting bootstrap cluster
```

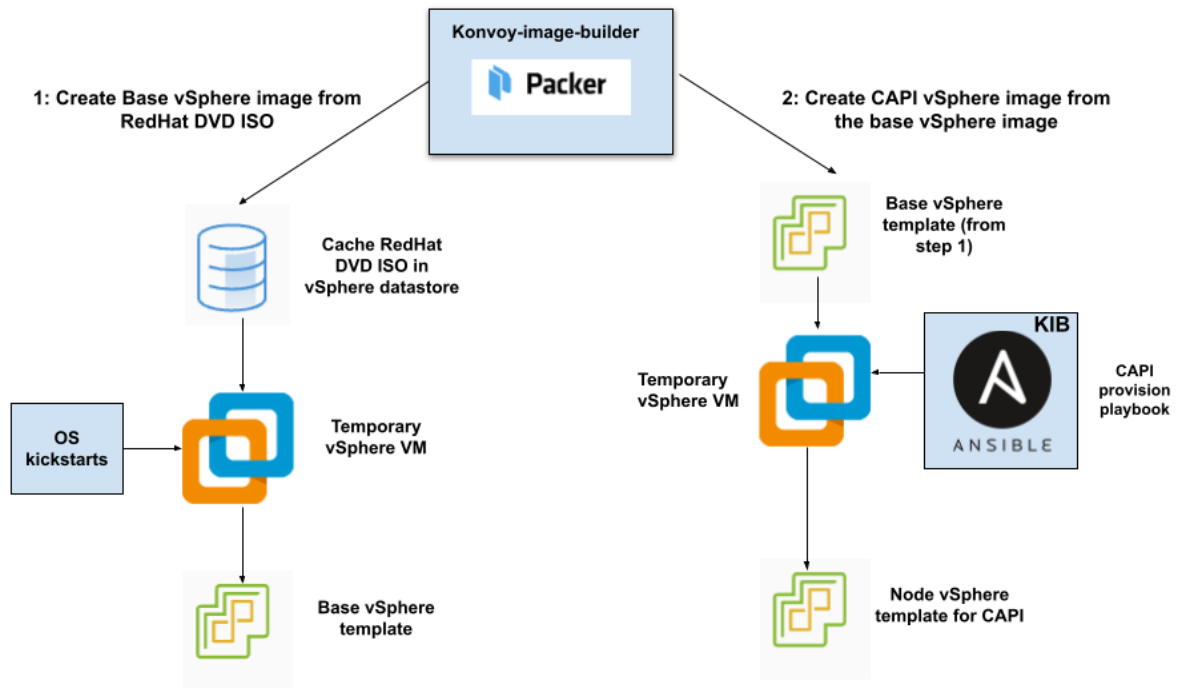
7.16 vSphere Infrastructure

7.16.1 Creating DKP clusters in a VMware vSphere environment

The overall process for configuring vSphere and DKP together includes the following steps:

1. Configure vSphere to provide the needed elements, described in the [Prerequisites \(see page 1256\)](#).
2. Create a [bastion VM host \(see page 1285\)](#) if you are using an air-gapped environment.
3. Create a base OS image (for use in the OVA package containing the disk images packaged with the OVF).
4. Create a CAPI VM image template that uses the base OS image and adds the needed Kubernetes cluster components.
5. Create a bootstrap cluster.
6. Create a new DKP cluster on vSphere.
7. Make the cluster self-managing.
8. Explore the cluster and perform other functions as needed.

This diagram illustrates the image creation process:



The workflow on the left shows the creation of a base OS image in the vCenter vSphere client using inputs from Packer. The workflow on the right shows how DKP uses that same base OS image to create CAPI-enabled VM images for your cluster.

After creating the base image, the [Konvoy Image Builder](#) (see page 1461) uses it to create a CAPI-enabled vSphere template that includes the Kubernetes objects for the cluster. You can use that resulting template with the DKP `create cluster` command to create the VM nodes in your cluster directly on a vCenter server. From that point, you can use DKP to provision and manage your cluster.

DKP communicates with the code in [vCenter Server](#)⁹⁴² as the management layer for creating and managing virtual machines after [ESXi 6.7 Update 3 or later](#)⁹⁴³ is installed and configured.

7.16.2 Next Step:

To get started, fulfill the [Prerequisites](#) (see page 1256) on the following page.

⁹⁴² <https://docs.vmware.com/en/VMware-vSphere/index.html>

⁹⁴³ <https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.esxi.install.doc/GUID-B2F01BF5-078A-4C7E-B505-5DFFED0B8C38.html>

7.16.3 vSphere Prerequisites

7.16.3.1 Prepare your environment to run DKP with VMware vSphere

Fulfilling the prerequisites involves completing these two areas:

- DKP prerequisites
- vSphere prerequisites - vCenter Server + ESXi

7.16.3.2 DKP Prerequisites

Before using DKP to create a vSphere cluster, verify that you have:

- An x86_64-based Linux® or macOS® machine.
- [Download DKP binaries and Konvoy Image Builder \(KIB\)](#) (see page 78) image bundle for Linux or macOS.
- A Container engine/runtime installed is required to install DKP and bootstrap:
 - Version [Docker®](#)⁹⁴⁴ container engine version 18.09.2 or higher installed for Linux or MacOS.
 - Version 4.0 of [Podman](#)⁹⁴⁵ or higher for Linux. Host requirements found here: [Host Requirements](#)⁹⁴⁶
- A registry needs installed on the host where the DKP Konvoy CLI runs. For example, if you are installing Konvoy on your laptop, ensure the laptop has a supported version of Docker or other registry.



On macOS, Docker runs in a virtual machine. Configure this virtual machine with at least 8GB of memory.

- [kubecti](#)⁹⁴⁷ 1.21.6 for interacting with the running cluster, installed on the host where the DKP Konvoy command line interface (CLI) runs.
- A valid VMware vSphere account with credentials configured.



DKP uses the vsphere CSI driver as the [default storage provider](#) (see page 91). Use a [Kubernetes CSI](#)⁹⁴⁸ compatible storage that is suitable for production.

944 <https://docs.docker.com/get-docker/>

945 <https://podman.io/getting-started/installation>

946 <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>

947 <https://kubernetes.io/docs/tasks/tools/#kubectl>

948 <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

You can choose from any of the [storage options](#)⁹⁴⁹ available for Kubernetes. To disable the default that Konvoy deploys, set the default StorageClass as non-default. Then set your newly created StorageClass to be the default by following the commands in the Kubernetes documentation called [Changing the Default Storage Class](#)⁹⁵⁰.

7.16.3.3 VMware vSphere Prerequisites

Before installing, verify that your [VMware vSphere Client environment](#)⁹⁵¹ meets the following basic requirements:

- Access to a bastion VM, or other network connected host, running vSphere Client version v6.7.x with Update 3 or later version.
 - You must be able to reach the vSphere API endpoint from where the Konvoy command line interface (CLI) runs.
- vSphere account with credentials configured - this account must have Administrator privileges.
- A RedHat® subscription with user name and password for downloading DVD ISOs.
- For air-gapped environments, a [bastion VM host template](#) (see page 108) with access to a configured [local registry](#) (see page 1533). VMware site has more information on [vSphere Bastion Hosts](#)⁹⁵². The recommended template naming pattern is `./folder-name/dkp-e2e-bastion-template` or similar. Each infrastructure provider has its own set of bastion host instructions. Refer to your provider's site for details - [Azure](#)⁹⁵³, [AWS](#)⁹⁵⁴, [GCP](#)⁹⁵⁵, or [vSphere](#)⁹⁵⁶.
- Valid vSphere values for the following:
 - vCenter API server URL.
 - Datacenter name.
 - Zone name that contains [ESXi hosts](#)⁹⁵⁷ for your cluster's nodes.
 - Datastore name for the shared storage resource to be used for the VMs in the cluster.
 - Use of PersistentVolumes in your cluster depends on Cloud Native Storage (CNS), available in vSphere v6.7.x with Update 3 and later versions. CNS depends on this shared Datastore's configuration.
 - Datastore URL from the datastore record for the shared datastore you want your cluster to use.

949 <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

950 <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

951 https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html

952 <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-6975426F-56D0-4FE2-8A58-580B40D2F667.html>

953 <https://learn.microsoft.com/en-us/azure/bastion/quickstart-host-portal>

954 <https://aws.amazon.com/solutions/implementations/linux-bastion/>

955 <https://blogs.vmware.com/cloud/2021/06/02/intro-google-cloud-vmware-engine-bastion-host-access-iap/>

956 <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-6975426F-56D0-4FE2-8A58-580B40D2F667.html>

957 <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.esxi.install.doc/GUID-B2F01BF5-078A-4C7E-B505-5DFFED0B8C38.html>

- You need this URL value to ensure that the correct Datastore is used when DKP creates VMs for your cluster in vSphere.
- Folder name.
- Base template name, such as base-rhel-8, or base-rhel-7.
- Name of a Virtual Network that has DHCP enabled for both air-gapped and non-air-gapped environments.
- Resource Pools - at least one resource pool needed, with every host in the pool having access to shared storage, such as VSAN.
 - Each host in the resource pool needs access to shared storage, such as NFS or VSAN, to make use of MachineDeployments and high-availability control planes.

7.16.3.4 Next Steps:

- Non-air-gapped [Create a Base OS image in vSphere \(see page 1260\)](#)
- Air-gapped [Create a Base Air-gapped OS VM Image \(see page 1288\)](#)

7.16.3.5 Minimum User Permissions

7.16.3.5.1 Create minimum required roles for provisioning and installing in vSphere

When a user needs permissions less than Admin, a role must be created. The process for configuring a vSphere role with the least permissions for provisioning nodes and installing includes the following steps:

1. Open a vSphere Client connection to the vCenter Server, described in the [Prerequisites \(see page 1259\)](#).
2. Select Home > Administration > Roles > Add Role.
3. Give the new role a name, then select these Privileges:

Cns		
<input checked="" type="checkbox"/>		Searchable
Datastore		
<input checked="" type="checkbox"/>		Allocate space
<input checked="" type="checkbox"/>		Low level file operations
Host		

		• Configuration
	<input checked="" type="checkbox"/>	Storage partition configuration
Profile-driven storage		
	<input checked="" type="checkbox"/>	Profile-driven storage view
Network		
	<input checked="" type="checkbox"/>	Assign network
Resource		
	<input checked="" type="checkbox"/>	Assign virtual machine to resource pool
Virtual machine		
		• Change Configuration - from the list in that section, select these permissions below:
	<input checked="" type="checkbox"/>	Add new disk
	<input checked="" type="checkbox"/>	Add existing disk
	<input checked="" type="checkbox"/>	Add or remove device
	<input checked="" type="checkbox"/>	Advanced configuration
	<input checked="" type="checkbox"/>	Change CPU count
	<input checked="" type="checkbox"/>	Change Memory
	<input checked="" type="checkbox"/>	Change Settings
	<input checked="" type="checkbox"/>	Reload from path
Edit inventory		

<input checked="" type="checkbox"/>	Create from existing
<input checked="" type="checkbox"/>	Remove
Interaction	
<input checked="" type="checkbox"/>	Power off
<input checked="" type="checkbox"/>	Power on
Provisioning	
<input checked="" type="checkbox"/>	Clone template
<input checked="" type="checkbox"/>	Deploy template
Session	
<input checked="" type="checkbox"/>	ValidateSession

Add the permission at the highest level and set to propagate the permissions.

7.16.3.6 vSphere Storage Options

7.16.3.6.1 Explore storage options and considerations for using DKP with VMware vSphere

The [vSphere Container Storage](#)⁹⁵⁸ plugin supports shared NFS, vNFS, and vSAN. You need to provision your storage options in vCenter prior to [creating a CAPI image \(see page 1289\)](#) in DKP for use with vSphere.

DKP has integrated the CSI 2.x driver used in vSphere. When creating your DKP cluster, DKP uses whatever configuration you provide for the Datastore name. vSAN is not required. Using NFS can reduce the amount of tagging and permission granting required to configure your cluster.

7.16.4 Create a Base OS image in vSphere

Creating a base OS image from DVD ISO files is a one-time process. Building a base OS image creates a base vSphere template in your vSphere environment. The base OS image is used by [Konvoy Image Builder](#)

⁹⁵⁸ <https://docs.vmware.com/en/VMware-vSphere-Container-Storage-Plug-in/2.0/vmware-vsphere-csp-getting-started/GUID-74AF02D7-1562-48BD-A9FE-C81A53342AC3.html>

(see page 1461)(KIB) to create a VM template to configure Kubernetes nodes by the DKP vSphere provider. For more information about images in vSphere vCenter, refer to [VMware documentation](#)⁹⁵⁹.

7.16.4.1 Create the Base OS Image

For vSphere, a username and password is populated by `SSH_USERNAME` and the user can use authorization via `SSH_PASSWORD` or `SSH_PRIVATE_KEY_FILE` environment variables and [required by default by packer](#)⁹⁶⁰. This user should have administrator privileges. It is possible to configure a custom user and password when building the OS image, however, that requires the Konvoy Image Builder (KIB) configuration to be [overridden](#) (see page 1513).

While creating the base OS image, it is important to take into consideration the following elements:

- **Storage configuration:** D2iQ recommends customizing disk partitions and not configuring a SWAP partition.
- **Network configuration:** as KIB must download and install packages, activating the network is required.
- **Connect to Red Hat:** if using RHEL, registering with Red Hat is required to configure software repositories and install software packages.
- **Software selection:** D2iQ recommends choosing **Minimal Install**.
- DKP recommends to install with the packages provided by the operating system package managers. Use the version that corresponds to the major version of your operating system.

7.16.4.1.1 Disk Size

For each cluster you create using this base OS image, ensure you establish the disk size of the **root file system** based on:

- The minimum [DKP storage requirements](#) (see page 100).
- The minimum storage requirements for your organization.

7.16.4.1.1.1 Defaults

Clusters are created with a default disk size of 80 GB.



For clusters created with the default disk size, the base OS image root file system must be exactly 80 GB. The root file system cannot be reduced automatically when a machine first boots.

⁹⁵⁹ <https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/vcf-deploy/GUID-2674DA5A-8DF7-4212-A4A9-88CD798DC303.html>

⁹⁶⁰ <https://github.com/mesosphere/konvoy-image-builder/blob/0523fd2c5e6e1ad1d4962f60a47039aa145a6e42/pkg/packer/manifests/vsphere/packer.pkr.hcl>

7.16.4.1.1.2 Customization

You can also specify a custom disk size when you [create a cluster](#) (see page 0) (see the flags of the [Create new vSphere Cluster](#) (see page 1670) command). This allows you to use one base OS image to create multiple clusters that have different storage requirements.



Before specifying a disk size when you create a cluster, take into account:

1. **For some base OS images, the custom disk size option has no effect on the size of the root file system.** This is because some root file systems, for example, those contained in an LVM Logical Volume, cannot be resized automatically when a machine first boots.
2. **The specified custom disk size must be equal to, or larger than the size of the base OS image root file system.** This is because a root file system cannot be reduced automatically when a machine first boots.
3. In [VMware Cloud Director](#) (see page 1322), the base image determines the minimum storage available for the VM.

7.16.4.2 Next Step:

[Create a VM Template](#) (see page 1262)

7.16.5 Create a VM Template

7.16.5.1 Create a vSphere Template for Your Cluster from a Base OS Image

Using the base OS image created in a previous procedure, DKP creates the new vSphere template directly on the vCenter server.

1. Set the following vSphere environment variables on the bastion VM host:

```
export VSPHERE_SERVER=your_vCenter_APIserver_URL
export VSPHERE_USERNAME=your_vCenter_user_name
export VSPHERE_PASSWORD=your_vCenter_password
```

2. Copy the base OS image file created in the vSphere Client to your desired location on the bastion VM host, and make a note of the path and file name.

3. Create an `image.yaml` file and add the following variables for vSphere. DKP uses this file and these variables as inputs in the next step.

⚠ **NOTE:** This example is Ubuntu 20.04. You will need to replace OS name below based on your OS. See other [YAML examples](#)⁹⁶¹ for copy and paste below last step.

```
---
download_images: true
build_name: "ubuntu-2004"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: "<VSPHERE_CLUSTER_NAME>"
  datacenter: "<VSPHERE_DATACENTER_NAME>"
  datastore: "<VSPHERE_DATASTORE_NAME>"
  folder: "<VSPHERE_FOLDER>"
  insecure_connection: "false"
  network: "<VSPHERE_NETWORK>"
  resource_pool: "<VSPHERE_RESOURCE_POOL>"
  template: "os-qualification-templates/d2iq-base-Ubuntu-20.04" # change
  default value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "ubuntu2004-64"
  # goss params
  distribution: "ubuntu"
  distribution_version: "20.04"
  # Use following overrides to select the authentication method that can be used
  # with base template
  # ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
  # ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
  # ssh_private_key_file = "" # can be exported as environment variable
  # 'SSH_PRIVATE_KEY_FILE'
  # ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key
  # will be ignored
```

4. Create a vSphere VM template with your variation of the following command:

```
konvoy-image build images/ova/<image.yaml>
```

Any additional configurations can be added to this command using `--overrides` flags as shown below:

- a. Any credential overrides: `--overrides overrides.yaml`

⁹⁶¹ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

- b. for FIPS, add this flag: `--overrides overrides/fips.yaml`
 - c. for air-gapped, add this flag: `--overrides overrides/offline-fips.yaml`
5. The [Konvoy Image Builder](#) (see page 1487) (KIB) uses the values in `image.yaml` and the input base OS image to create a vSphere template directly on the vCenter server. This template contains the required artifacts needed to create a Kubernetes cluster.
- When KIB provisions the OS [image](#)⁹⁶² successfully, it creates a manifest file. The `artifact_id` field of this file contains the name of the AMI ID (AWS), template name (vSphere), or image name (GCP/Azure), for example:

```
{
  "name": "vsphere-clone",
  "builder_type": "vsphere-clone",
  "build_time": 1644985039,
  "files": null,
  "artifact_id": "konvoy-ova-vsphere-rhel-84-1.21.6-1644983717",
  "packer_run_uuid": "260e8110-77f8-ca94-e29e-ac7a2ae779c8",
  "custom_data": {
    "build_date": "2022-02-16T03:55:17Z",
    "build_name": "vsphere-rhel-84",
    "build_timestamp": "1644983717",
    [...]
  }
}
```

Recommendation: Now we can now see the template created in our vCenter, it is best to rename it to `dkp-<DKP_VERSION>-k8s-<K8S_VERSION>-<DISTRO>`, like `dkp-2.6.2-k8s-1.26.14-ubuntu` to keep templates organized.

- 6. Next steps are to deploy a DKP cluster using your vSphere template.

Create a Kubernetes Bootstrap Cluster to enable creating your vSphere cluster and moving CAPI objects to it.

7.16.5.2 Specific to VMware Cloud Director

For storage, the VCD **Provider** must follow [Create a Base OS image in vSphere](#) (see page 1263) when they create the KIB base image. In VCD, the base image determines the minimum storage available for the VM.

7.16.5.2.1 Additional OS YAML files for copy/paste:

RHEL 8.6

```
---
download_images: true
```

⁹⁶² <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

build_name: "rhel-86"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false" # can be "true"
  network: "<VSPHERE_NETWORK>"
  resource_pool: "<VSPHERE_RESOURCE_POOL>"
  template: "os-qualification-templates/d2iq-base-RHEL-86" # change default value
with your base template name
vsphere_guest_os_type: "rhel8_64Guest"
guest_os_type: "rhel8-64"
# goss params
distribution: "RHEL"
distribution_version: "8.6"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

Rocky Linux 9.1

```

---
download_images: true
build_name: "rocky-91"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""

```

```

template: "os-qualification-templates/d2iq-base-RockyLinux-9.1" # change default
value with your base template name
vsphere_guest_os_type: "other4xLinux64Guest"
guest_os_type: "rocky9-64"
# goss params
distribution: "rocky"
distribution_version: "9.1"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored

```

[YAML Files \(see page 1492\)](#) found in Konvoy Image Builder section.

7.16.5.2.2 Next Step:

[vSphere Bootstrap \(see page 1266\)](#)

7.16.6 vSphere Bootstrap

7.16.6.1 Prepare to deploy Kubernetes clusters

To create Kubernetes clusters, Konvoy uses [Cluster API](#)⁹⁶³ (CAPI) controllers, which run on a Kubernetes cluster. To get started creating your vSphere cluster, you need a *bootstrap* cluster. By default, Konvoy creates a bootstrap cluster for you in a Docker container using the Kubernetes-in-Docker ([KIND](#)⁹⁶⁴) tool.

7.16.6.2 Prerequisites

Before you begin, you must:

- Ensure the `dkp` binary can be found in your `$PATH`.
- Complete the steps in [Create a CAPI VM template \(see page 1265\)](#)

7.16.6.3 Bootstrap cluster lifecycle services

1. If an HTTP proxy is required for the bootstrap cluster, set the local `http_proxy`, `https_proxy`, and `no_proxy` environment variables. They are copied into the bootstrap cluster.
2. Create a bootstrap cluster:

⁹⁶³ <https://cluster-api.sigs.k8s.io/>

⁹⁶⁴ <https://github.com/kubernetes-sigs/kind>

```
dkp create bootstrap --kubeconfig $HOME/.kube/config
```

The output resembles this example:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

Konvoy creates a bootstrap cluster using [KIND](#)⁹⁶⁵ as a library. Konvoy then deploys the following [Cluster API](#)⁹⁶⁶ providers on the cluster:

- [Core Provider](#)⁹⁶⁷
- [vSphere Infrastructure Provider](#)⁹⁶⁸
- [Kubeadm Bootstrap Provider](#)⁹⁶⁹
- [Kubeadm ControlPlane Provider](#)⁹⁷⁰

3. Ensure that the CAPV controllers are present with the command:

```
kubectl get pods -n capv-system
```

The output resembles the following:

NAME	READY	STATUS	RESTARTS	AGE
capv-controller-manager-785c5978f-nnfn5	1/1	Running	0	13h

4. Konvoy waits until the controller-manager and webhook deployments of these providers are ready. List these deployments using this command:

```
kubectl get --all-namespaces deployments -l=clusterctl.cluster.x-k8s.i
```

The output resembles the following:

NAMESPACE	NAME
capa-system	capa-controller-manager
1/1	1
1	1
22m	

965 <https://github.com/kubernetes-sigs/kind>

966 <https://cluster-api.sigs.k8s.io/>

967 <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

968 <https://github.com/kubernetes-sigs/cluster-api-provider-vsphere>

969 <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

970 <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>

capi-kubeadm-bootstrap-system				capi-kubeadm-bootstrap-controller-manager
1/1	1	1	22m	
capi-kubeadm-control-plane-system				capi-kubeadm-control-plane-controller-
manager	1/1	1	1	22m
capi-system				capi-controller-manager
1/1	1	1	22m	
capp-system				capp-controller-manager
1/1	1	1	22m	
capv-system				capv-controller-manager
1/1	1	1	22m	
capz-system				capz-controller-manager
1/1	1	1	22m	
cert-manager				cert-manager
1/1	1	1	22m	
cert-manager				cert-manager-cainjector
1/1	1	1	22m	
cert-manager				cert-manager-webhook
1/1	1	1	22m	

7.16.6.4 Next Step:

[Create new vSphere Cluster \(see page 1268\)](#)

7.16.7 Create new vSphere Cluster

7.16.7.1 Prerequisites

Before you begin, make sure you have created a [vSphere Bootstrap \(see page 1269\)](#) cluster.

7.16.7.2 Name your cluster

1. Give your cluster a unique name suitable for your environment.
2. Set the CLUSTER_NAME environment variable with the command:

```
export CLUSTER_NAME=<my-vsphere-cluster>
```

7.16.7.3 Create a New vSphere Kubernetes Cluster

Follow these steps:

1. Use the following command to set the environment variables for vSphere:

```
export VSPHERE_SERVER=example.vsphere.url
export VSPHERE_USERNAME=user@example.vsphere.url
export VSPHERE_PASSWORD=example_password
```

2. Ensure your vSphere credentials are up-to-date by refreshing the credentials with the command:

```
dkp update bootstrap credentials vsphere
```

3. Generate the Kubernetes cluster objects by copying and editing this command to include the correct values, including the VM template name you assigned in the previous procedure:



- To increase [Dockerhub's rate limit](#)⁹⁷¹ use your Dockerhub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.
- Ensure your [subnets](#) (see [page 1049](#)) do not overlap with your host subnet because they cannot be changed after cluster creation. If you need to change the kubernetes subnets, you must do this at cluster creation. The default subnets used in DKP are:

```
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
        - 192.168.0.0/16
    services:
      cidrBlocks:
        - 10.96.0.0/12
```

- The following `dkp create cluster` example shows a common configuration. See [dkp create cluster vsphere](#) (see [page 1670](#)) reference for the full list of cluster creation options:

```
dkp create cluster vsphere \
  --cluster-name ${CLUSTER_NAME} \
  --network <NETWORK_NAME> \
  --control-plane-endpoint-host <xxx.yyy.zzz.000> \
  --data-center <DATACENTER_NAME> \
  --data-store <DATASTORE_NAME> \
  --folder <FOLDER_NAME> \
```

⁹⁷¹ <https://docs.docker.com/docker-hub/download-rate-limit/>


```

--server <VCENTER_API_SERVER_URL> \
--ssh-public-key-file <SSH_PUBLIC_KEY_FILE> \
--resource-pool <RESOURCE_POOL_NAME> \
--virtual-ip-interface <ip_interface_name> \
--vm-template <TEMPLATE_NAME>
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml

```

- (Optional) Alternatively, you can create individual files with different smaller manifests for ease in editing using the `--output-directory` flag. This will create multiple files in the specified directory which must already exist:

```
--output-directory=<existing-directory>
```



For more information regarding this flag or others, please refer to the CLI section of the documentation for [dkp create cluster](#) (see [page 1662](#)) and select your provider.

4. (Optional) To configure the Control Plane and Worker nodes to use an HTTP proxy:

```

export CONTROL_PLANE_HTTP_PROXY=http://example.org:8080
export CONTROL_PLANE_HTTPS_PROXY=http://example.org:8080
export
CONTROL_PLANE_NO_PROXY="example.org,example.com,example.net,localhost,127.0.0.1,10.96
.0.0/12,192.168.0.0/16,kubernetes,kubernetes.default,kubernetes.default.svc,kubernet
s.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.clu
ster.local,169.254.169.254,.elb.amazonaws.com"


export WORKER_HTTP_PROXY=http://example.org:8080
export WORKER_HTTPS_PROXY=http://example.org:8080
export
WORKER_NO_PROXY="example.org,example.com,example.net,localhost,127.0.0.1,10.96.0.0/12
,192.168.0.0/16,kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.defau
lt.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluster.lo
cal,169.254.169.254,.elb.amazonaws.com"

```

- Replace `example.org,example.com,example.net` with your internal addresses
- `localhost` and `127.0.0.1` addresses should not use the proxy
- `10.96.0.0/12` is the default Kubernetes service subnet
- `192.168.0.0/16` is the default Kubernetes pod subnet


- `kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local` is the internal Kubernetes kube-apiserver service
- `.svc,.svc.cluster,.svc.cluster.local` is the internal Kubernetes services
- `169.254.169.254` is the AWS metadata server
- `.elb.amazonaws.com` is for the worker nodes to allow them to communicate directly to the kube-apiserver ELB

5. (Optional) Create a Kubernetes cluster with HTTP proxy configured. This step assumes you did not already create a cluster in the previous steps:

 To increase [Dockerhub's rate limit](#)⁹⁷² use your Dockerhub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` on the `dkp create cluster` command.

```
dkp create cluster vsphere --cluster-name=${CLUSTER_NAME} \
--control-plane-http-proxy="${CONTROL_PLANE_HTTP_PROXY}" \
--control-plane-https-proxy="${CONTROL_PLANE_HTTPS_PROXY}" \
--control-plane-no-proxy="${CONTROL_PLANE_NO_PROXY}" \
--worker-http-proxy="${WORKER_HTTP_PROXY}" \
--worker-https-proxy="${WORKER_HTTPS_PROXY}" \
--worker-no-proxy="${WORKER_NO_PROXY}" \
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml
```

6. Inspect or edit the cluster objects:

 Familiarize yourself with Cluster API before editing the cluster objects as edits can prevent the cluster from deploying successfully.

The objects are [Custom Resources](#)⁹⁷³ defined by Cluster API components, and they belong in three different categories:

1. Cluster

972 <https://docs.docker.com/docker-hub/download-rate-limit/>

973 <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

A *Cluster* object has references to the infrastructure-specific and control plane objects. Because this is a vSphere cluster, there is an object that describes the infrastructure-specific cluster properties.

2. Control plane

A *KubeadmControlPlane* object describes the control plane, which is the group of machines that run the Kubernetes control plane components, which include the etcd distributed database, the API server, the core controllers, and the scheduler. The object describes the configuration for these components. The object also has a reference to an infrastructure-specific object that describes the properties of all control plane machines. Here, it references an *vSphereMachineTemplate* object.

3. Node pool

A node pool is a collection of machines with identical properties. For example, a cluster might have one node pool with large memory capacity, another node pool with GPU support. Each node pool is described by three objects: The *MachinePool* references an object that describes the configuration of Kubernetes components (for example, kubelet) deployed on each node pool machine, and an infrastructure-specific object that describes the properties of all node pool machines. Here, it references a *KubeadmConfigTemplate*, and a *vSphereMachineTemplate* object.

For in-depth documentation about the objects, read [Concepts](#)⁹⁷⁴ in the Cluster API Book.

7. Modify control plane audit logs settings using the information contained in the page [Configuring the Control Plane](#) (see page 1538).

8. Create the cluster from the objects. A warning will appear in the console if the resource already exists and will require you to remove the resource or update your YAML.

```
kubectl create -f ${CLUSTER_NAME}.yaml
```

:note: NOTE: If you used the `--output-directory` flag in your `dkp create .. --dry-run` step above, create the cluster from the objects you created by specifying the directory:

```
kubectl create -f <existing-directory>/
```

Output will be similar to below:

```
cluster.cluster.x-k8s.io/vsphere-example created
cluster.infrastructure.cluster.x-k8s.io/vsphere-example created
kubeadmcontrolplane.controlplane.cluster.x-k8s.io/vsphere-example-control-plane
created
machinedeployment.cluster.x-k8s.io/vsphere-example-mp-0 created
kubeadmconfigtemplate.bootstrap.cluster.x-k8s.io/vsphere-example-mp-0 created
```

9. Use the `wait` command to monitor the cluster control-plane readiness:

⁹⁷⁴ <https://cluster-api.sigs.k8s.io/user/concepts.html>

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

```
cluster.cluster.x-k8s.io/${CLUSTER_NAME} condition met
```

The `READY` status becomes `True` after the cluster control-plane becomes Ready in one of the following steps.

After DKP creates the objects on the API server, the Cluster API controllers reconcile them, creating infrastructure and machines. As the controllers progress, they update the Status of each object.

10. Run the DKP describe command to monitor the current status of the cluster:

```
dkp describe cluster -c ${CLUSTER_NAME}
```

NAME	READY	SEVERITY
REASON SINCE MESSAGE		
Cluster/d2iq-e2e-cluster_name-1	True	
13h		
└─ClusterInfrastructure - VSphereCluster/d2iq-e2e-cluster_name-1	True	
13h		
└─ControlPlane - KubeadmControlPlane/d2iq-control-plane	True	
13h		
├─Machine/d2iq--control-plane-7llgd	True	
13h		
├─Machine/d2iq--control-plane-vncbl	True	
13h		
├─Machine/d2iq--control-plane-wbgrm	True	
13h		
└─Workers		
├─MachineDeployment/d2iq--md-0	True	
13h		
├─Machine/d2iq--md-0-74c849dc8c-67rv4	True	
13h		
├─Machine/d2iq--md-0-74c849dc8c-n2skc	True	
13h		
├─Machine/d2iq--md-0-74c849dc8c-nkftv	True	
13h		
├─Machine/d2iq--md-0-74c849dc8c-sqklv	True	
13h		

11. Check all machines has `NODE_NAME` assigned:

```
kubectl get machines
```

The output appears similar to the following:

NAME	CLUSTER	NODENAME
d2iq-e2e-cluster-1-control-plane-7llgd	d2iq-e2e-cluster-1	d2iq-e2e-cluster-1-control-plane-7llgd
13h v1.26.6	vsphere://421638e2-e776-9af6-f683-5e105de5da5a	Running
d2iq-e2e-cluster-1-control-plane-vncbl	d2iq-e2e-cluster-1	d2iq-e2e-cluster-1-control-plane-vncbl
13h v1.26.6	vsphere://42168835-7fef-95c4-3652-ebcad3e10d36	Running
d2iq-e2e-cluster-1-control-plane-wbgrm	d2iq-e2e-cluster-1	d2iq-e2e-cluster-1-control-plane-wbgrm
13h v1.26.6	vsphere://421642df-afc4-b6c2-9e61-5b86e7c37eac	Running
d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4	d2iq-e2e-cluster-1	d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4
14h v1.26.6	vsphere://4216f467-8483-73cb-a8b6-8d6a4a71e4b4	Running
d2iq-e2e-cluster-1-md-0-74c849dc8c-n2skc	d2iq-e2e-cluster-1	d2iq-e2e-cluster-1-md-0-74c849dc8c-n2skc
14h v1.26.6	vsphere://42161cde-9904-4dd2-7a3e-cdfc7655f090	Running
d2iq-e2e-cluster-1-md-0-74c849dc8c-nkftv	d2iq-e2e-cluster-1	d2iq-e2e-cluster-1-md-0-74c849dc8c-nkftv
14h v1.26.6	vsphere://42163a0d-eb8d-b5a6-82d5-188e24817c00	Running
d2iq-e2e-cluster-1-md-0-74c849dc8c-sqklv	d2iq-e2e-cluster-1	d2iq-e2e-cluster-1-md-0-74c849dc8c-sqklv
14h v1.26.6	vsphere://42161dff-92a5-6da9-7ac1-e987e2c8fed2	Running

12. Verify that the kubeadm control plane is ready with the command

```
kubectl get kubeadmcontrolplane
```

The output appears similar to the following:

NAME	CLUSTER	INITIALIZED	API SERVER
d2iq-e2e-cluster-1-control-plane	d2iq-e2e-cluster-1	true	true
3	14h v1.26.6		

13. Describe the kubeadm control plane and check its status and events with the command:

```
kubectl describe kubeadmcontrolplane
```

14. As they progress, the controllers also create Events, which you can list using the command:

```
kubectl get events | grep ${CLUSTER_NAME}
```

For brevity, this example uses `grep`. You can also use separate commands to get Events for specific objects, such as `kubectl get events --field-selector`

```
involvedObject.kind="VSphereCluster" and kubectl get events --field-selector
involvedObject.kind="VSphereMachine" .
```

7.16.7.4 Known Limitations



Be aware of these limitations in the current release of DKP Konvoy.

- The DKP Konvoy version used to create a bootstrap cluster must match the DKP Konvoy version used to create a workload cluster.
- DKP Konvoy supports deploying one workload cluster.
- DKP Konvoy generates a set of objects for one Node Pool.
- DKP Konvoy does not validate edits to cluster objects.

The optional next step is to [Make the vSphere Cluster Self-managed](#) (see page 1279). The step is optional because, as an example, if you are using an existing, self-managed cluster to create a managed cluster, you would not want the managed cluster to be self-managed.

7.16.7.5 Next Steps:

[Explore a vSphere Cluster](#) (see page 1275)

[Make vSphere Cluster Self-Managed](#) (see page 1279)

7.16.8 Explore a vSphere Cluster

This guide explains how to use the command line interface to interact with your newly-deployed Kubernetes cluster.

Before you start, make sure you have [created a workload cluster](#) (see page 1271) and, if needed, that you have [made the cluster self-managing](#) (see page 1279).

7.16.8.1 Get the kubeconfig file for the new Kubernetes cluster

1. Get a kubeconfig file for the workload cluster:

When the workload cluster is created, the cluster lifecycle services generate a kubeconfig file for the workload cluster, and write it to a *Secret*. The kubeconfig file is scoped to the cluster Administrator.

Get the kubeconfig from the *Secret*, and write it to a file using this command:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

7.16.8.2 Create a StorageClass with a vSphere datastore

Follow these steps:

1. Access the Datastore tab in the vSphere client and select a datastore by name.
2. Copy the URL of that datastore from the information dialog that displays.
3. Return to the DKP CLI, and delete the existing `StorageClass` with the command:

```
kubectl delete storageclass vsphere-raw-block-sc
```

4. Run the following command to create a new `StorageClass`, supplying the correct values for your environment:

```
cat <<EOF > vsphere-raw-block-sc.yaml
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  annotations:
    storageclass.kubernetes.io/is-default-class: "true"
  name: vsphere-raw-block-sc
provisioner: csi.vsphere.vmware.com
parameters:
  datastoreurl: "<url>"
volumeBindingMode: WaitForFirstConsumer
EOF
```

7.16.8.3 Explore nodes and pods in the new cluster

1. List the nodes using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

NOTE: It may take a few minutes for the Status to move to `Ready` while the Pod network is deployed. The Node's Status should change to `Ready` soon after the `calico-node` DaemonSet Pods are `Ready`.

The output resembles the following example:

NAME VERSION	STATUS	ROLES	AGE
d2iq-e2e-cluster-1-control-plane-7llgd v1.26.6	Ready	control-plane,master	20h
d2iq-e2e-cluster-1-control-plane-vncbl v1.26.6	Ready	control-plane,master	19h
d2iq-e2e-cluster-1-control-plane-wbgrm v1.26.6	Ready	control-plane,master	19h
d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4 v1.26.6	Ready	<none>	19h
d2iq-e2e-cluster-1-md-0-74c849dc8c-n2skc v1.26.6	Ready	<none>	19h
d2iq-e2e-cluster-1-md-0-74c849dc8c-nkftv v1.26.6	Ready	<none>	19h
d2iq-e2e-cluster-1-md-0-74c849dc8c-sqklv v1.26.6	Ready	<none>	19h

- List the pods using this command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get --all-namespaces pods
```

The output resembles the following example:

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
calico-system	calico-kube-controllers-57fbd7bd59-qqd96	1/1	Running	0	20h
calico-system	calico-node-2m524	1/1	Running	3 (19h ago)	19h
calico-system	calico-node-bbhg5	1/1	Running	0	20h
calico-system	calico-node-cc5lf	1/1	Running	2 (19h ago)	19h
calico-system	calico-node-cwg7x	1/1	Running	1 (19h ago)	19h
calico-system	calico-node-d59hn	1/1	Running	1 (19h ago)	19h
calico-system	calico-node-qmmc	1/1	Running	0	19h
calico-system	calico-node-wdqhx	1/1	Running	0	19h
calico-system	calico-typha-655489d8cc-b5jnt	1/1	Running	0	20h
calico-system	calico-typha-655489d8cc-q92x9	1/1	Running	0	19h

calico-system			calico-typha-655489d8cc-vjlkx
1/1	Running	0	19h
kube-system			cluster-autoscaler-68c759fbf6-7d2ck
0/1	Init:0/1	0	20h
kube-system			coredns-78fcd69978-qn4qt
1/1	Running	0	20h
kube-system			coredns-78fcd69978-wqpmg
1/1	Running	0	20h
kube-system			etcd-d2iq-e2e-cluster-1-control-plane-7llgd
1/1	Running	0	20h
kube-system			etcd-d2iq-e2e-cluster-1-control-plane-vncbl
1/1	Running	0	19h
kube-system			etcd-d2iq-e2e-cluster-1-control-plane-wbgrm
1/1	Running	0	19h
kube-system			kube-apiserver-d2iq-e2e-cluster-1-control-plane-7llgd
1/1	Running	0	20h
kube-system			kube-apiserver-d2iq-e2e-cluster-1-control-plane-vncbl
1/1	Running	0	19h
kube-system			kube-apiserver-d2iq-e2e-cluster-1-control-plane-wbgrm
1/1	Running	0	19h
kube-system			kube-controller-manager-d2iq-e2e-cluster-1-control-
plane-7llgd	1/1	Running	1 (19h ago) 20h
kube-system			kube-controller-manager-d2iq-e2e-cluster-1-control-
plane-vncbl	1/1	Running	0 19h
kube-system			kube-controller-manager-d2iq-e2e-cluster-1-control-
plane-wbgrm	1/1	Running	0 19h
kube-system			kube-proxy-cpscs
1/1	Running	0	19h
kube-system			kube-proxy-hhmxq
1/1	Running	0	19h
kube-system			kube-proxy-hxhmk
1/1	Running	0	19h
kube-system			kube-proxy-nsrbp
1/1	Running	0	19h
kube-system			kube-proxy-scxfg
1/1	Running	0	20h
kube-system			kube-proxy-tth4k
1/1	Running	0	19h
kube-system			kube-proxy-x2xfx
1/1	Running	0	19h
kube-system			kube-scheduler-d2iq-e2e-cluster-1-control-plane-7llgd
1/1	Running	1 (19h ago)	20h
kube-system			kube-scheduler-d2iq-e2e-cluster-1-control-plane-vncbl
1/1	Running	0	19h
kube-system			kube-scheduler-d2iq-e2e-cluster-1-control-plane-wbgrm
1/1	Running	0	19h
kube-system			kube-vip-d2iq-e2e-cluster-1-control-plane-7llgd
1/1	Running	1 (19h ago)	20h
kube-system			kube-vip-d2iq-e2e-cluster-1-control-plane-vncbl
1/1	Running	0	19h
kube-system			kube-vip-d2iq-e2e-cluster-1-control-plane-wbgrm
1/1	Running	0	19h

```

kube-system          vsphere-cloud-controller-manager-4zj7q
1/1    Running    0                19h
kube-system          vsphere-cloud-controller-manager-87tgm
1/1    Running    0                19h
kube-system          vsphere-cloud-controller-manager-xqmn4
1/1    Running    1 (19h ago)    20h
node-feature-discovery node-feature-discovery-master-84c67dcbb6-txfw9
1/1    Running    0                20h
node-feature-discovery node-feature-discovery-worker-8tg2l
1/1    Running    3 (19h ago)    19h
node-feature-discovery node-feature-discovery-worker-c5f6q
1/1    Running    0                19h
node-feature-discovery node-feature-discovery-worker-fjfk
1/1    Running    0                19h
node-feature-discovery node-feature-discovery-worker-x6tz8
1/1    Running    0                19h
tigera-operator      tigera-operator-d499f5c8f-r2srj
1/1    Running    1 (19h ago)    20h
vmware-system-csi    vsphere-csi-controller-7ffd6884cc-d7rql
7/7    Running    5 (19h ago)    20h
vmware-system-csi    vsphere-csi-controller-7ffd6884cc-k82cm
7/7    Running    2 (19h ago)    20h
vmware-system-csi    vsphere-csi-controller-7ffd6884cc-qttkp
7/7    Running    1 (19h ago)    20h
vmware-system-csi    vsphere-csi-node-678hw
3/3    Running    0                19h
vmware-system-csi    vsphere-csi-node-6tbsh
3/3    Running    0                19h
vmware-system-csi    vsphere-csi-node-9htwr
3/3    Running    5 (20h ago)    20h
vmware-system-csi    vsphere-csi-node-g8r6l
3/3    Running    0                19h
vmware-system-csi    vsphere-csi-node-ghmr6
3/3    Running    0                19h
vmware-system-csi    vsphere-csi-node-jhvgn
3/3    Running    0                19h
vmware-system-csi    vsphere-csi-node-rp77r
3/

```

7.16.9 Make vSphere Cluster Self-Managed

Konvoy deploys all cluster lifecycle services to a bootstrap cluster, which then deploys a workload cluster. When the workload cluster is ready, move the cluster lifecycle services to the workload cluster, which makes the workload cluster self-managed. This section describes how to make a workload cluster self-managed.



Before starting, ensure you create a workload cluster as described in [Create a New vSphere Cluster](#) or [Create a New Air-gapped vSphere Cluster](#).

This page contains instructions on how to make your cluster self-managed. This is necessary if there is only one cluster in your environment, or if this cluster should become the Management cluster in a multi-cluster environment.

 If you already have a self-managed or Management cluster in your environment, skip this page.

7.16.9.1 Make the new Kubernetes cluster manage itself

Follow these steps:

1. Deploy cluster lifecycle services on the workload cluster:

```
dkp create capi-components --kubeconfig ${CLUSTER_NAME}.conf
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1042\)](#).

✓ Initializing **new** CAPI components

2. Move the Cluster API objects from the bootstrap to the workload cluster:

The cluster lifecycle services on the workload cluster are ready, but the workload cluster configuration is on the bootstrap cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the bootstrap to the workload cluster. This process is also called a [Pivot](#)⁹⁷⁵.

```
dkp move capi-resources --to-kubeconfig ${CLUSTER_NAME}.conf
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1042\)](#).

✓ Moving cluster resources

⁹⁷⁵ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```
You can now view resources in the moved cluster by using the --kubeconfig flag
with kubectl. For example: kubectl --kubeconfig=vsphere-example.conf get nodes
```

NOTE: To ensure only one set of cluster lifecycle services manages the workload cluster, Konvoy first pauses reconciliation of the objects on the bootstrap cluster, then creates the objects on the workload cluster. As Konvoy copies the objects, the cluster lifecycle services on the workload cluster reconcile the objects. The workload cluster becomes self-managed after Konvoy creates all the objects. If it fails, the move command can be safely retried.

3. Wait for the cluster control-plane to be ready:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf wait --for=condition=ControlPlaneReady
"clusters/${CLUSTER_NAME}" --timeout=20m
```

```
cluster.cluster.x-k8s.io/vsphere-example condition met
```

4. Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

NOTE: After moving the cluster lifecycle services to the workload cluster, remember to use Konvoy with the workload cluster kubeconfig.

```
dkp describe cluster --kubeconfig ${CLUSTER_NAME}.conf -c ${CLUSTER_NAME}
```

```
NAME                                                    READY
SEVERITY REASON SINCE MESSAGE
Cluster/vsphere-example-1                               True
13h
├ClusterInfrastructure - VSphereCluster/vsphere-example-1 True
13h
├ControlPlane - KubeadmControlPlane/vsphere-example-control-plane True
13h
| ├Machine/vsphere-example-control-plane-7llgd          True
13h
| ├Machine/vsphere-example-control-plane-vncbl          True
13h
| └Machine/vsphere-example-control-plane-wbgrm          True
13h
└Workers
  └MachineDeployment/vsphere-example-md-0                True
13h
```

13h	—Machine/vsphere-example-md-0-74c849dc8c-67rv4	True
13h	—Machine/vsphere-example-md-0-74c849dc8c-n2skc	True
13h	—Machine/vsphere-example-md-0-74c849dc8c-nkftv	True
13h	—Machine/vsphere-example-md-0-74c849dc8c-sqklv	True

5. Remove the bootstrap cluster, if desired, as the workload cluster is now self-managed:

```
dkp delete bootstrap
```

```
✓ Deleting bootstrap cluster
```

7.16.9.2 Known limitations

Be aware of these limitations in the current release of DKP Konvoy.

- Before making a workload cluster self-managed, be sure that its control plane nodes have sufficient permissions for running Cluster API controllers.
- DKP Konvoy supports moving only one set of cluster objects from the bootstrap cluster to the workload cluster, or vice-versa.
- DKP Konvoy only supports moving all namespaces in the cluster; DKP does not support migration of individual namespaces.

7.16.10 Configure MetalLB for a vSphere infrastructure

7.16.10.1 Create a MetalLB configmap for your vSphere infrastructure.

Choose one of the following two protocols you want to use to announce service IPs, either Layer 2 or BGP configurations.

7.16.10.2 Layer 2 configuration

Layer 2 mode is the simplest to configure: in many cases, you don't need any protocol-specific configuration, only IP addresses.

Layer 2 mode does not require the IPs to be bound to the network interfaces of your worker nodes. It works by responding to ARP requests on your local network directly, to give the machine's MAC address to clients.



- MetalLB IP address ranges/CIDRs should be within the node's primary network [subnet](#) (see [page 1049](#)).
- MetalLB IP address ranges/CIDRs and node subnet should not conflict with the Kubernetes cluster pod and service subnets.

For example, the following configuration gives MetalLB control over IPs from 192.168.1.240 to 192.168.1.250, and configures Layer 2 mode:



The following values are generic, enter your specific values into the fields where applicable.

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    address-pools:
    - name: default
      protocol: layer2
      addresses:
      - 192.168.1.240-192.168.1.250
EOF
```

When complete, run the following `kubectl` command.

```
kubectl apply -f metallb-conf.yaml
```

7.16.10.3 BGP configuration

BGP, like any routing protocol, requires a mesh, so you configure BGP between the nodes. If you are in a Data Center, peer with the BGP routers. The instructions for use with Cloud providers is the same. For a standard configuration featuring one BGP router and one IP address range, you need four pieces of information:

- The router IP address that MetalLB should connect to,
- The router's AS number,
- The AS number MetalLB should use,

- An IP address range expressed as a CIDR prefix.

As an example, if you want to give MetalLB the range 192.168.10.0/24 and AS number 64500, and connect it to a router at 10.0.0.1 with AS number 64501, your configuration will look like:



The following values are generic, enter your specific values into the fields where applicable.

Extract the kubeconfig and deploy a config map for MetalLB using the following command:

```
dkp get kubeconfig -c ${DKP_CLUSTER_NAME} > ${DKP_CLUSTER_NAME}.conf
```

Deploy MetalLB Configuration with the command below:

```
cat << EOF > metallb-conf.yaml
apiVersion: v1
kind: ConfigMap
metadata:
  namespace: metallb-system
  name: config
data:
  config: |
    peers:
      - peer-address: 10.0.0.1
        peer-asn: 64501
        my-asn: 64500
    address-pools:
      - name: default
        protocol: bgp
        addresses:
          - 192.168.10.0/24
EOF
```

When complete, run the following `kubectl` command:

```
kubectl apply -f metallb-conf.yaml
```

7.16.11 Install vSphere Air-Gapped

7.16.11.1 Create a Kubernetes vSphere cluster in a private network with no access to the Internet (air-gapped)

This section provides the instructions to create a Kubernetes cluster in a private, DHCP-enabled network with no access to the Internet. Complete the list of [Prerequisites](#) (see page 1259) before you begin the procedures in this section.

7.16.11.2 Bastion Host

When creating an air-gapped vSphere cluster, the bastion VM hosts the installation of the DKP Konvoy bundles and images, as well as the Docker or other [local registry](#) (see page 1533), needed to create and operate your vSphere cluster. The bastion VM must have access to the vSphere API Server (vCenter Server). Ensure the items below are installed and the environment matches the requirements below:

- Create a bastion VM host template for the cluster nodes to use within the air-gapped network. This bastion VM host also needs access to a local registry in lieu of an Internet connection for pulling images. The recommended template naming pattern is `../folder-name/dkp-e2e-bastion-template` or similar.
- Find and record the bastion VM's IP or host name.
- [Download](#) (see page 78) the following required DKP Konvoy binaries and installation bundles discussed in step 5 below.
- A [local registry](#) (see page 1533) or [Docker](#)⁹⁷⁶ version 18.09.2 or later installed. You must have Docker installed on the host where the DKP Konvoy CLI runs. For example, if you are installing Konvoy on your laptop, ensure the laptop has a supported version of Docker. On macOS, Docker runs in a virtual machine which you configure with at least 8GB of memory.
- [kubect](#)⁹⁷⁷ for interacting with the running cluster, installed on the host where the DKP Konvoy command line interface (CLI) runs.

Depending on your OS, there are various commands for setting up your own bastion host for use with air-gapped vSphere.

This would be a **generic** example for RHEL Bastion nodes using Docker:

1. Once `base-rhel-os` boots, open an `ssh` terminal to host and install the tools and packages:

```
sudo yum install -y yum-utils bzip2 wget
```

2. Install `kubect` as mentioned above, below is a RHEL example:

⁹⁷⁶ <https://docs.docker.com/get-docker/>

⁹⁷⁷ <https://kubernetes.io/docs/tasks/tools/#kubect>


```
cat <<EOF | sudo tee /etc/yum.repos.d/kubernetes.repo
[kubernetes]
name=Kubernetes
baseurl=https://packages.cloud.google.com/yum/repos/kubernetes-el7-\\$basearch
enabled=1
gpgcheck=1
gpgkey=https://packages.cloud.google.com/yum/doc/rpm-package-key.gpg
EOF
sudo yum install -y kubectl
```

3. Install for Docker as an example registry (only on Bastion Host) and add the repo for upstream Docker:

```
sudo yum-config-manager --add-repo https://download.docker.com/linux/rhel/
docker-ce.repo
```

NOTE: Other Docker repo downloads are available on [docker.com](https://download.docker.com/linux/)⁹⁷⁸: <https://download.docker.com/linux/>

4. Install example for Docker:

```
sudo yum install -y docker-ce docker-ce-cli containerd.io
```

5. Create directory for Konvoy Image Builder and DKP CLI:

```
mkdir kib && mkdir dkp
```

6. Get the needed D2iQ Software by downloading the air-gapped bundle:

[Download \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2/kib
```

7. Install DKP CLI:

```
cd ..
cd dkp
wget https://downloads.d2iq.com/dkp/v2.6.2/dkp_v2.6.2_linux_amd64.tar.gz
```

8. Set the following vSphere environment variables on the bastion VM host:

978 <http://docker.com>

```
export VSPHERE_SERVER=your_vCenter_APIserver_URL
export VSPHERE_USERNAME=your_vCenter_user_name
export VSPHERE_PASSWORD=your_vCenter_password
```

9. Set the following environment variables to enable connection to an existing Docker or other registry. You must create the VM template with the [Konvoy Image Builder](#) (see page 1461) to be able to use the registry mirror feature:

```
export REGISTRY_ADDRESS=<https/http>://<registry-address>:<registry-port>
export REGISTRY_CA=<path to the CA on the bastion host>
```

- **REGISTRY_ADDRESS** : the address of an existing registry accessible in the vSphere Zone where the new cluster nodes will be configured, to use a mirror registry when pulling images.
- **REGISTRY_CA** : (optional) the path on the bastion host to the registry CA. Konvoy configures the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the VMs are not already configured to trust this CA.

7.16.11.2.1 More information:

Each infrastructure provider has its own set of bastion host instructions. Refer to your own OS instructions to setup a bastion host like [AWS Bastion](#)⁹⁷⁹, [Azure](#)⁹⁸⁰, [GCP](#)⁹⁸¹, or [vSphere](#)⁹⁸².

⁹⁷⁹ <https://aws.amazon.com/solutions/implementations/linux-bastion/>

⁹⁸⁰ <https://learn.microsoft.com/en-us/azure/bastion/quickstart-host-portal>

⁹⁸¹ <https://blogs.vmware.com/cloud/2021/06/02/intro-google-cloud-vmware-engine-bastion-host-access-iap/>

⁹⁸² <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-6975426F-56D0-4FE2-8A58-580B40D2F667.html>

7.16.11.3 Create a Base Air-gapped OS VM Image

7.16.11.3.1 Create a Base OS VM Image in the VMware vSphere Client (air-gapped)

7.16.11.3.2 Creating a base OS image from DVD ISO files is a one-time process. Building a base OS image creates a base vSphere template in your vSphere environment. The base OS image is used by [Konvoy Image Builder](#) (KIB) to create a VM template to configure Kubernetes nodes by the DKP vSphere provider. For more information about images in vSphere vCenter, refer to [VMware documentation](#)⁹⁸³.

7.16.11.3.3 Create the Base OS Image

For vSphere, a username and password is populated by `SSH_USERNAME` and the user can use authorization via `SSH_PASSWORD` or `SSH_PRIVATE_KEY_FILE` environment variables and [required by default by packer](#)⁹⁸⁴. This user should have administrator privileges. It is possible to configure a custom user and password when building the OS image, however, that requires the Konvoy Image Builder (KIB) configuration to be [overridden](#) (see page 1513).

While creating the base OS image, it is important to take into consideration the following elements:

- [Storage configuration](#): D2iQ recommends customizing disk partitions and not configuring a SWAP partition.
- [Network configuration](#): as KIB must download and install packages, activating the network is required.
- [Connect to Red Hat](#): if using RHEL, registering with Red Hat is required to configure software repositories and install software packages.
- [Software selection](#): D2iQ recommends choosing **Minimal Install**.
- DKP recommends to install with the packages provided by the operating system package managers. Use the version that corresponds to the major version of your operating system.

7.16.11.3.3.1 Disk Size

For each cluster you create using this base OS image, ensure you establish the disk size of the **root file system** based on:

- The minimum [DKP storage requirements](#) (see page 100).
- The minimum storage requirements for your organization.

Defaults

⁹⁸³ <https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/vcf-deploy/GUID-2674DA5A-8DF7-4212-A4A9-88CD798DC303.html>

⁹⁸⁴ <https://github.com/mesosphere/konvoy-image-builder/blob/0523fd2c5e6e1ad1d4962f60a47039aa145a6e42/pkg/packer/manifests/vsphere/packer.pkr.hcl>

Clusters are created with a default disk size of 80 GB.



For clusters created with the default disk size, the base OS image root file system must be exactly 80 GB. The root file system cannot be reduced automatically when a machine first boots.

Customization

You can also specify a custom disk size when you [create a cluster \(see page 0\)](#) (see the flags of the [Create new vSphere Cluster \(see page 1670\)](#) command). This allows you to use one base OS image to create multiple clusters that have different storage requirements.



Before specifying a disk size when you create a cluster, take into account:

1. **For some base OS images, the custom disk size option has no effect on the size of the root file system.** This is because some root file systems, for example, those contained in an LVM Logical Volume, cannot be resized automatically when a machine first boots.
2. **The specified custom disk size must be equal to, or larger than the size of the base OS image root file system.** This is because a root file system cannot be reduced automatically when a machine first boots.
3. In [VMware Cloud Director \(see page 1322\)](#), the base image determines the minimum storage available for the VM.

7.16.11.3.3.2 Next Step:

[Create a CAPI VM Template \(see page 1289\)](#)

7.16.11.4 Create an Air-gapped CAPI VM Template

You must have at least one image before creating a new cluster. As long as you have an image, this step in your configuration is not required each time since that image can be used to spin up a new cluster. However, if you need different images for different environments or providers, you will need to create a new custom image.

Create a vSphere template for your cluster from a base OS image

Using [KIB \(see page 1461\)](#), you can build an image without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2/kib
```

2. Follow the instructions to build a vSphere template below and set the override `--overrides overrides/offline.yaml` flag.

7.16.11.4.1 Create a vSphere Template for Your Cluster from a Base OS Image

Using the base OS image created in a previous procedure, DKP creates the new vSphere template directly on the vCenter server.

1. Set the following vSphere environment variables on the bastion VM host:

```
export VSPHERE_SERVER=your_vCenter_APIserver_URL
export VSPHERE_USERNAME=your_vCenter_user_name
export VSPHERE_PASSWORD=your_vCenter_password
```

2. Copy the base OS image file created in the vSphere Client to your desired location on the bastion VM host, and make a note of the path and file name.
3. Create an `image.yaml` file and add the following variables for vSphere. DKP uses this file and these variables as inputs in the next step.

⚠ NOTE: This example is Ubuntu 20.04. You will need to replace OS name below based on your OS. See other [YAML examples](#)⁹⁸⁵ for copy and paste below last step.

```
---
download_images: true
build_name: "ubuntu-2004"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: "<VSPHERE_CLUSTER_NAME>"
  datacenter: "<VSPHERE_DATACENTER_NAME>"
  datastore: "<VSPHERE_DATASTORE_NAME>"
  folder: "<VSPHERE_FOLDER>"
  insecure_connection: "false"
  network: "<VSPHERE_NETWORK>"
  resource_pool: "<VSPHERE_RESOURCE_POOL>"
```

⁹⁸⁵ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

template: "os-qualification-templates/d2iq-base-Ubuntu-20.04" # change
default value with your base template name
vsphere_guest_os_type: "other4xLinux64Guest"
guest_os_type: "ubuntu2004-64"
# goss params
distribution: "ubuntu"
distribution_version: "20.04"
# Use following overrides to select the authentication method that can be used
with base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key
will be ignored

```

4. Create a vSphere VM template with your variation of the following command:

```
konvoy-image build images/ova/<image.yaml>
```

Any additional configurations can be added to this command using `--overrides` flags as shown below:

- a. Any credential overrides: `--overrides overrides.yaml`
 - b. for FIPS, add this flag: `--overrides overrides/fips.yaml`
 - c. for air-gapped, add this flag: `--overrides overrides/offline-fips.yaml`
5. The [Konvoy Image Builder](#) (see page 1487) (KIB) uses the values in `image.yaml` and the input base OS image to create a vSphere template directly on the vCenter server. This template contains the required artifacts needed to create a Kubernetes cluster.

When KIB provisions the OS `image`⁹⁸⁶ successfully, it creates a manifest file. The `artifact_id` field of this file contains the name of the AMI ID (AWS), template name (vSphere), or image name (GCP/Azure), for example:

```

{
  "name": "vsphere-clone",
  "builder_type": "vsphere-clone",
  "build_time": 1644985039,
  "files": null,
  "artifact_id": "konvoy-ova-vsphere-rhel-84-1.21.6-1644983717",
  "packer_run_uuid": "260e8110-77f8-ca94-e29e-ac7a2ae779c8",
  "custom_data": {
    "build_date": "2022-02-16T03:55:17Z",
    "build_name": "vsphere-rhel-84",

```

⁹⁸⁶ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

    "build_timestamp": "1644983717",
    [...]
  }
}

```

Recommendation: Now we can now see the template created in our vCenter, it is best to rename it to `dkp-<DKP_VERSION>-k8s-<K8S_VERSION>-<DISTRO>`, like `dkp-2.6.2-k8s-1.26.14-ubuntu` to keep templates organized.

6. Next steps are to deploy a DKP cluster using your vSphere template.

7.16.11.4.2 Next Step:

[vSphere Air-gapped Seed the Registry \(see page 1292\)](#)

7.16.11.5 vSphere Air-gapped Seed the Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine \(see page 1288\)](#) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.



If you do not already have a local registry set up, please refer to [Local Registry Tools \(see page 1533\)](#) page for more information.

1. [Download \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, and extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

3. Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-registry-password $REGISTRY_PASSWORD
```

ⓘ It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

7.16.11.5.1 Next Step:

[vSphere Air-gapped Bootstrap](#) (see page 1293)

7.16.11.6 vSphere Air-gapped Bootstrap

7.16.11.6.1 Prerequisites

Before you perform this procedure, ensure that you have [created a CAPI VM template](#) (see page 1292).

7.16.11.6.2 Bootstrap a kind cluster and CAPI controllers

Konvoy deploys all cluster lifecycle services to a bootstrap cluster, which deploys a workload cluster. When the workload cluster is ready, move the cluster lifecycle services to the workload cluster, after which the workload cluster manages its own lifecycle.

1. Assuming you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.1_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.1_linux_amd64.tar.gz && cd dkp-v2.6.1
```

2. Load the bootstrap image on your bastion machine using either Docker or Podman command:

```
docker load -i konvoy-bootstrap-image-v2.6.1.tar
```

```
podman load -i konvoy-bootstrap-image-v2.6.1.tar
```

3. Create a bootstrap cluster:

```
dkp create bootstrap --kubeconfig $HOME/.kube/config
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful.

More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

The output resembles this example:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

4. Ensure that the CAPV controllers are present with the command:

```
kubectl get pods -n capv-system
```

The output resembles the following:

NAME	READY	STATUS	RESTARTS	AGE
capv-controller-manager-785c5978f-nnfns	1/1	Running	0	13h

7.16.11.6.3 Next Step:

[Create a new Air-gapped vSphere Cluster](#) (see page 1294)

7.16.11.7 Create a new Air-gapped vSphere Cluster

7.16.11.7.1 Prerequisites

Before you begin, be sure that you have created a [Bootstrap Cluster](#) (see page 1296)

7.16.11.7.2 Create a new vSphere Kubernetes cluster

Use the following steps to create a new, air-gapped vSphere cluster.

1. Configure your cluster to use an existing registry as a mirror when attempting to pull images:

IMPORTANT: The image must be created by the [konvoy-image-builder](#)⁹⁸⁷ project in order to use the registry mirror feature.

```
export REGISTRY_URL=<https/http>://<registry-address>:<registry-port>
export REGISTRY_CA=<path to the CA on the bastion>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

⁹⁸⁷ <https://github.com/mesosphere/konvoy-image-builder>

- `REGISTRY_URL` : the address of an existing registry accessible in the VPC that the new cluster nodes will be configured to use a mirror registry when pulling images.
 - `REGISTRY_CA` : (optional) the path on the bastion machine to the registry CA. Konvoy will configure the cluster nodes to trust this CA. This value is only needed if the registry is using a self-signed certificate and the AMIs are not already configured to trust this CA.
 - `REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
 - `REGISTRY_PASSWORD` : optional if username is not set.
2. Create a Kubernetes cluster by copying the following command and substituting the valid values for your environment:

NOTE: Ensure your [subnets](#) (see page 1049) do not overlap with your host subnet because they cannot be changed after cluster creation. If you need to change the kubernetes subnets, you must do this at cluster creation. The default subnets used in DKP are:

```
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
        - 192.168.0.0/16
    services:
      cidrBlocks:
        - 10.96.0.0/12
```

Create cluster with common configurations shown below:

```
dkp create cluster vsphere
--cluster-name ${CLUSTER_NAME} \
--network <NETWORK_NAME> \
--control-plane-endpoint-host <CONTROL_PLANE_IP> \
--data-center <DATACENTER_NAME> \
--data-store <DATASTORE_NAME> \
--folder <FOLDER_NAME> \
--server <VCENTER_API_SERVER_URL> \
--ssh-public-key-file </path/to/key.pub> \
--resource-pool <RESOURCE_POOL_NAME> \
--vm-template konvoy-ova-vsphere-os-release-k8s_release-vsphere-timestamp \
--virtual-ip-interface eth0 \
--extra-sans "127.0.0.1" \
--registry-mirror-url=${REGISTRY_URL} \
--registry-mirror-cacert=${REGISTRY_CA} \
--registry-mirror-username=${REGISTRY_USERNAME} \
--registry-mirror-password=${REGISTRY_PASSWORD}
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

3. Inspect the created cluster resources with the command:

```
kubectl get clusters,kubeadmcontrolplanes,machinedeployments
```

4. Use the `wait` command to monitor the cluster control-plane readiness:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

Output:

```
cluster.cluster.x-k8s.io/${CLUSTER_NAME} condition met
```

The `READY` status becomes `True` after the cluster control-plane becomes Ready in one of the following steps.

After DKP creates the objects on the API server, the Cluster API controllers reconcile them, creating infrastructure and machines. As the controllers progress, they update the Status of each object.

5. Run the DKP `describe` command to monitor the current status of the cluster:

```
dkp describe cluster -c ${CLUSTER_NAME}
```

```
NAME                                                    READY
SEVERITY REASON SINCE MESSAGE
Cluster/e2e-airgapped-1                                True
13h
└─ClusterInfrastructure - VSphereCluster/e2e-airgapped-1 True
13h
└─ControlPlane - KubeadmControlPlane/e2e-airgapped-1-control-plane True
13h
| └─Machine/e2e-airgapped-1-control-plane-7llgd         True
13h
| └─Machine/e2e-airgapped-1-control-plane-vncbl         True
13h
| └─Machine/e2e-airgapped-1-control-plane-wbgrm         True
13h
└─Workers
  └─MachineDeployment/e2e-airgapped-1-md-0               True
13h
```

```

    |Machine/e2e-airgapped-1-md-0-74c849dc8c-67rv4           True
13h
    |Machine/e2e-airgapped-1-md-0-74c849dc8c-n2skc          True
13h
    |Machine/e2e-airgapped-1-md-0-74c849dc8c-nkftv         True
13h
    |Machine/e2e-airgapped-1-md-0-74c849dc8c-sqklv         True
13h

```

6. As they progress, the controllers also create Events, which you can list using the command:

```
kubectl get events | grep ${CLUSTER_NAME}
```

For brevity, this example uses `grep`. You can also use separate commands to get Events for specific objects, such as `kubectl get events --field-selector involvedObject.kind="VSphereCluster"` and `kubectl get events --field-selector involvedObject.kind="VSphereMachine"`.

You can [make the cluster self-managed](#) (see page 1282) with the information in the linked page.

7.16.11.7.2.1 Next Step:

You can [explore your new cluster](#) (see page 1301).

7.16.11.7.3 Known limitations

NOTE: Be aware of these limitations in the current release of DKP Konvoy.

- The DKP Konvoy version used to create a bootstrap cluster must match the DKP Konvoy version used to create a workload cluster.
- DKP Konvoy supports deploying one workload cluster.
- DKP Konvoy generates a set of objects for one Node Pool.
- DKP Konvoy does not validate edits to cluster objects.

7.16.11.8 Make vSphere Air-gapped Cluster Self-Managed

Konvoy deploys all cluster lifecycle services to a bootstrap cluster, which then deploys a workload cluster. When the workload cluster is ready, move the cluster lifecycle services to the workload cluster, which makes the workload cluster self-managed. This section describes how to make a workload cluster self-managed.

- Before starting, ensure you create a workload cluster as described in [Create a New vSphere Cluster](#) or [Create a New Air-gapped vSphere Cluster](#).

This page contains instructions on how to make your cluster self-managed. This is necessary if there is only one cluster in your environment, or if this cluster should become the Management cluster in a multi-cluster environment.

- If you already have a self-managed or Management cluster in your environment, skip this page.

7.16.11.8.1 Make the new Kubernetes cluster manage itself

Follow these steps:

1. Deploy cluster lifecycle services on the workload cluster:

```
dkp create capi-components --kubeconfig ${CLUSTER_NAME}.conf
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

```
✓ Initializing new CAPI components
```

2. Move the Cluster API objects from the bootstrap to the workload cluster:

The cluster lifecycle services on the workload cluster are ready, but the workload cluster configuration is on the bootstrap cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the bootstrap to the workload cluster. This process is also called a [Pivot](#)⁹⁸⁸.

```
dkp move capi-resources --to-kubeconfig ${CLUSTER_NAME}.conf
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--`

⁹⁸⁸ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

`https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy \(see page 1042\)](#).

✓ Moving cluster resources

You can now view resources in the moved cluster by using the `--kubeconfig` flag with `kubectl`. For example: `kubectl --kubeconfig=vsphere-example.conf get nodes`

NOTE: To ensure only one set of cluster lifecycle services manages the workload cluster, Konvoy first pauses reconciliation of the objects on the bootstrap cluster, then creates the objects on the workload cluster. As Konvoy copies the objects, the cluster lifecycle services on the workload cluster reconcile the objects. The workload cluster becomes self-managed after Konvoy creates all the objects. If it fails, the move command can be safely retried.

3. Wait for the cluster control-plane to be ready:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --timeout=20m
```

```
cluster.cluster.x-k8s.io/vsphere-example condition met
```

4. Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

NOTE: After moving the cluster lifecycle services to the workload cluster, remember to use Konvoy with the workload cluster `kubeconfig`.

```
dkp describe cluster --kubeconfig ${CLUSTER_NAME}.conf -c ${CLUSTER_NAME}
```

```
NAME                                                                 READY
SEVERITY REASON SINCE MESSAGE
Cluster/vsphere-example-1                                          True
13h
|─ClusterInfrastructure - VSphereCluster/vsphere-example-1       True
13h
|─ControlPlane - KubeadmControlPlane/vsphere-example-control-plane True
13h
| |─Machine/vsphere-example-control-plane-7llgd                   True
13h
```

	Machine/vsphere-example-control-plane-vncbl	True
13h		
	Machine/vsphere-example-control-plane-wbgrm	True
13h		
	└─Workers	
	└─MachineDeployment/vsphere-example-md-0	True
13h		
	Machine/vsphere-example-md-0-74c849dc8c-67rv4	True
13h		
	Machine/vsphere-example-md-0-74c849dc8c-n2skc	True
13h		
	Machine/vsphere-example-md-0-74c849dc8c-nkftv	True
13h		
	└─Machine/vsphere-example-md-0-74c849dc8c-sqklv	True
13h		

- Remove the bootstrap cluster, if desired, as the workload cluster is now self-managed:

```
dkp delete bootstrap
```

```
✓ Deleting bootstrap cluster
```

7.16.11.8.2 Known limitations

Be aware of these limitations in the current release of DKP Konvoy.

- Before making a workload cluster self-managed, be sure that its control plane nodes have sufficient permissions for running Cluster API controllers.
- DKP Konvoy supports moving only one set of cluster objects from the bootstrap cluster to the workload cluster, or vice-versa.
- DKP Konvoy only supports moving all namespaces in the cluster; DKP does not support migration of individual namespaces.

7.16.11.8.3 Next Step:

[Explore vSphere Air-gapped Cluster \(see page 1301\)](#)

7.16.11.9 Explore vSphere Air-gapped Cluster

7.16.11.9.1 Get the `kubeconfig` File for the New Kubernetes Cluster

1. Fetch the kubeconfig file with the command:

```
dkp get kubeconfig -c ${air-gapped_NAME} > ${air-gapped_NAME}.conf
```

7.16.11.9.2 Create a StorageClass with a vSphere Datastore

Follow these steps:

1. Access the Datastore tab in the vSphere client and select a datastore by name.
2. Copy the URL of that datastore from the information dialog that displays.
3. Return to the DKP CLI, and delete the existing `StorageClass` with the command:

```
kubectl delete storageclass vsphere-raw-block-sc
```

4. Run the following command to create a new StorageClass, supplying the correct values for your environment:

```
cat <<EOF > vsphere-raw-block-sc.yaml
kind: StorageClass
apiVersion: storage.k8s.io/v1
metadata:
  annotations:
    storageclass.kubernetes.io/is-default-class: "true"
  name: vsphere-raw-block-sc
provisioner: csi.vsphere.vmware.com
parameters:
  datastoreurl: "<url>"
volumeBindingMode: WaitForFirstConsumer
EOF
```

7.16.11.9.3 Explore Nodes and Pods in the New Cluster

1. List the Nodes with this command:


```
kubectl --kubeconfig=${air-gapped_NAME}.conf get nodes
```



NOTE: It may take a few minutes for the Status to move to Ready while the Pod network is deployed. The Node's Status should change to Ready soon after the calico-node DaemonSet Pods are Ready.

The output resembles this example:

NAME	STATUS	ROLES	AGE
VERSION			
d2iq-e2e-air-gapped-1-control-plane-7llgd v1.26.6	Ready	control-plane,master	20h
d2iq-e2e-air-gapped-1-control-plane-vncbl v1.26.6	Ready	control-plane,master	19h
d2iq-e2e-air-gapped-1-control-plane-wbgrm v1.26.6	Ready	control-plane,master	19h
d2iq-e2e-air-gapped-1-md-0-74c849dc8c-67rv4 v1.26.6	Ready	<none>	19h
d2iq-e2e-air-gapped-1-md-0-74c849dc8c-n2skc v1.26.6	Ready	<none>	19h
d2iq-e2e-air-gapped-1-md-0-74c849dc8c-nkftv v1.26.6	Ready	<none>	19h
d2iq-e2e-air-gapped-1-md-0-74c849dc8c-sqklv v1.26.6	Ready	<none>	19h

2. List the pods with the command:

```
kubectl --kubeconfig=${air-gapped_NAME}.conf get pods -A
```

The output resembles the following example:

NAMESPACE	NAME	READY	STATUS	RESTARTS	AGE
calico-system	calico-kube-controllers-57fbd7bd59-qqd96	1/1	Running	0	20h
calico-system	calico-node-2m524	1/1	Running	3 (19h ago)	19h

```

calico-system          calico-node-bbhg5
1/1      Running      0              20h
calico-system          calico-node-cc5lf
1/1      Running      2 (19h ago)   19h
calico-system          calico-node-cwg7x
1/1      Running      1 (19h ago)   19h
calico-system          calico-node-d59hn
1/1      Running      1 (19h ago)   19h
calico-system          calico-node-qmmcz
1/1      Running      0              19h
calico-system          calico-node-wdqhx
1/1      Running      0              19h
calico-system          calico-typha-655489d8cc-b5jnt
1/1      Running      0              20h
calico-system          calico-typha-655489d8cc-q92x9
1/1      Running      0              19h
calico-system          calico-typha-655489d8cc-vjlkx
1/1      Running      0              19h
kube-system            cluster-autoscaler-68c759fbf6-7d2ck
0/1      Init:0/1      0              20h
kube-system            coredns-78fcd69978-qn4qt
1/1      Running      0              20h
kube-system            coredns-78fcd69978-wqpmg
1/1      Running      0              20h
kube-system            etcd-d2iq-e2e-air-gapped-1-control-plane-7llgd
1/1      Running      0              20h
kube-system            etcd-d2iq-e2e-air-gapped-1-control-plane-vncbl
1/1      Running      0              19h
kube-system            etcd-d2iq-e2e-air-gapped-1-control-plane-wbgrm
1/1      Running      0              19h
kube-system            kube-apiserver-d2iq-e2e-air-gapped-1-control-plane-7llgd
1/1      Running      0              20h
kube-system            kube-apiserver-d2iq-e2e-air-gapped-1-control-plane-vncbl
1/1      Running      0              19h
kube-system            kube-apiserver-d2iq-e2e-air-gapped-1-control-plane-wbgrm
1/1      Running      0              19h
kube-system            kube-controller-manager-d2iq-e2e-air-gapped-1-control-
plane-7llgd  1/1      Running      1 (19h ago)   20h
kube-system            kube-controller-manager-d2iq-e2e-air-gapped-1-control-plane-
vncbl  1/1      Running      0              19h
kube-system            kube-controller-manager-d2iq-e2e-air-gapped-1-control-plane-
wbgrm  1/1      Running      0              19h
kube-system            kube-proxy-cpscs
1/1      Running      0              19h
kube-system            kube-proxy-hhmxq
1/1      Running      0              19h
kube-system            kube-proxy-hxhnk
1/1      Running      0              19h
kube-system            kube-proxy-nsrbp
1/1      Running      0              19h
kube-system            kube-proxy-scxfg
1/1      Running      0              20h

```

kube-system			kube-proxy-tth4k	
1/1	Running	0	19h	
kube-system			kube-proxy-x2xfx	
1/1	Running	0	19h	
kube-system			kube-scheduler-d2iq-e2e-air-gapped-1-control-plane-7llgd	
1/1	Running	1 (19h ago)	20h	
kube-system			kube-scheduler-d2iq-e2e-air-gapped-1-control-plane-vncbl	
1/1	Running	0	19h	
kube-system			kube-scheduler-d2iq-e2e-air-gapped-1-control-plane-wbgrm	
1/1	Running	0	19h	
kube-system			kube-vip-d2iq-e2e-air-gapped-1-control-plane-7llgd	
1/1	Running	1 (19h ago)	20h	
kube-system			kube-vip-d2iq-e2e-air-gapped-1-control-plane-vncbl	
1/1	Running	0	19h	
kube-system			kube-vip-d2iq-e2e-air-gapped-1-control-plane-wbgrm	
1/1	Running	0	19h	
kube-system			vsphere-cloud-controller-manager-4zj7q	
1/1	Running	0	19h	
kube-system			vsphere-cloud-controller-manager-87tgm	
1/1	Running	0	19h	
kube-system			vsphere-cloud-controller-manager-xqmn4	
1/1	Running	1 (19h ago)	20h	
node-feature-discovery			node-feature-discovery-master-84c67dccb6-txfw9	
1/1	Running	0	20h	
node-feature-discovery			node-feature-discovery-worker-8tg2l	
1/1	Running	3 (19h ago)	19h	
node-feature-discovery			node-feature-discovery-worker-c5f6q	
1/1	Running	0	19h	
node-feature-discovery			node-feature-discovery-worker-fjfkf	
1/1	Running	0	19h	
node-feature-discovery			node-feature-discovery-worker-x6tz8	
1/1	Running	0	19h	
tigera-operator			tigera-operator-d499f5c8f-r2srj	
1/1	Running	1 (19h ago)	20h	
vmware-system-csi			vsphere-csi-controller-7ffd6884cc-d7rql	
7/7	Running	5 (19h ago)	20h	
vmware-system-csi			vsphere-csi-controller-7ffd6884cc-k82cm	
7/7	Running	2 (19h ago)	20h	
vmware-system-csi			vsphere-csi-controller-7ffd6884cc-qttkp	
7/7	Running	1 (19h ago)	20h	
vmware-system-csi			vsphere-csi-node-678hw	
3/3	Running	0	19h	
vmware-system-csi			vsphere-csi-node-6tbsh	
3/3	Running	0	19h	
vmware-system-csi			vsphere-csi-node-9htwr	
3/3	Running	5 (20h ago)	20h	
vmware-system-csi			vsphere-csi-node-g8r6l	
3/3	Running	0	19h	
vmware-system-csi			vsphere-csi-node-ghmr6	
3/3	Running	0	19h	
vmware-system-csi			vsphere-csi-node-jhvgr	
3/3	Running	0	19h	

```
vmware-system-csi      vsphere-csi-node-rp77r
3/3      Running      0      19h
```

When you are ready, [delete your cluster and clean up your environment](#) (see page 1307).

7.16.12 vSphere Certificate Renewal

7.16.12.1 Configure Automated Renewal for Managed Kubernetes PKI Certificates

7.16.12.2 Certificate Renewal

During cluster creation, Kubernetes establishes a Public Key Infrastructure (PKI) for generating the TLS certificates needed for securing cluster communication for various components such as `etcd`, `kubernetes-apiserver` and `kube-proxy`. The certificates created by these components have a default expiration of one year and are renewed when an administrator updates the cluster.

Kubernetes provides a facility to renew all certificates automatically during control plane updates. For administrators who need long-running clusters or clusters that are not upgraded often, `dkp` provides automated certificate renewal, without a cluster upgrade.

7.16.12.2.1 Requirements

This feature requires that you install Python 3.5 or higher version on all control plane hosts.

7.16.12.3 Prerequisites

- Complete the [bootstrap Cluster Lifecycle](#) (see page 1269) topic.

7.16.12.3.1 Create a cluster with automated certificate renewal

To enable the automated certificate renewal, create a Konvoy cluster using the `certificate-renew-interval` flag:

```
dkp create cluster vsphere --certificate-renew-interval=60 --cluster-name=long-running
```

The `certificate-renew-interval` is the number of days after which Kubernetes-managed PKI certificates will be renewed. For example, an `certificate-renew-interval` value of 30 means the certificates are renewed every 30 days.

7.16.12.3.2 Technical details

The following manifests are modified on the control plane hosts, and are located at `/etc/kubernetes/manifests`. Modifications to these files requires SUDO access.

```
kube-controller-manager.yaml
kube-apiserver.yaml
kube-scheduler.yaml
kube-proxy.yaml
```

The following annotation indicates the time each component was reset:

```
metadata:
  annotations:
    konvoy.d2iq.io/restartedAt: $(date +%s)
```

This only occurs when the PKI certificates are older than the interval given at cluster creation time. This is activated by a `systemd` timer called `renew-certs.timer` that triggers an associated `systemd` service called `renew-certs.service` that runs on all of the control plane hosts.

7.16.12.3.3 Debugging

To debug the automatic certificate renewal feature, a cluster administrator can look at several different components to see if the certificates were renewed. For example, an administrator might start with a look at the control plane pod definition to check the last reset time. To determine if a scheduler pod was properly reset, run the command:

```
kubectl get pod -n kube-system kube-scheduler-nodename -o yaml
```

The output of the command is similar to the following:

```
apiVersion: v1
kind: Pod
metadata:
  annotations:
    konvoy.d2iq.io/restartedAt: "1626124940.735733"
```

Administrators who want more details on the execution of the `systemd` service can use `ssh` to connect to the control plane hosts, and then use the `systemctl` and `journalctl` commands that follow to help diagnose potential issues.

7.16.12.3.3.1 Check timer status

To check the status of the timers, when they last ran, and when they are scheduled to run next, use the command:

```
systemctl list-timers
```

7.16.12.3.3.2 Check renew certs status

To check the status of the `renew-certs` service, use the command:

```
systemctl status renew-certs
```

7.16.12.3.3.3 Review logs

To get the logs of the last run of the service, use the command:

```
journalctl logs -u renew-certs
```

7.16.13 Delete vSphere Cluster

7.16.13.1 Prepare to delete a self-managed workload cluster



A self-managed workload cluster cannot delete itself. If your workload cluster is self-managed, you must create a bootstrap cluster and move the cluster lifecycle services to the bootstrap cluster before deleting the workload cluster.

If you did not make your workload cluster self-managed, as described in [Make New Cluster Self-Managed](#) (see [page 1282](#)), see [Delete the workload cluster](#) (see [page 1310](#)).

1. Create a bootstrap cluster:

The bootstrap cluster will host the Cluster API controllers that reconcile the cluster objects marked for deletion. To avoid using the wrong kubeconfig, the following steps use **explicit** kubeconfig paths and contexts.

```
dkp create bootstrap --kubeconfig $HOME/.kube/config
```

The output resembles this example:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

2. Move the Cluster API objects from the workload to the bootstrap cluster: The cluster lifecycle services on the bootstrap cluster are ready, but the workload cluster configuration is on the workload cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the workload to the bootstrap cluster. This process is also called a [Pivot](#)⁹⁸⁹.

```
dkp move \
  --from-kubeconfig ${CLUSTER_NAME}.conf \
  --from-context konvoy-${CLUSTER_NAME}-admin@konvoy-${CLUSTER_NAME} \
  --to-kubeconfig $HOME/.kube/config \
  --to-context kind-konvoy-capi-bootstrapper
```

```
INFO[2021-06-09T11:47:11-07:00] Running pivot command
fromClusterKubeconfig=aws-example.conf fromClusterContext= src="move/
move.go:83" toClusterKubeconfig=/home/clusteradmin/.kube/config
toClusterContext=
INFO[2021-06-09T11:47:36-07:00] Pivot operation complete.
src="move/move.go:108"
INFO[2021-06-09T11:47:36-07:00] You can now view resources in the moved cluster
by using the --kubeconfig flag with kubectl. For example: kubectl --
kubeconfig=/home/clusteradmin/.kube/config get nodes src="move/move.go:155"
```

3. Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

⁹⁸⁹ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```
dkp describe cluster --kubeconfig $HOME/.kube/config -c ${CLUSTER_NAME}
```

```

NAME                                                                 READY
SEVERITY REASON SINCE MESSAGE
Cluster/d2iq-e2e-cluster_name-1                                     True
13h
└─ClusterInfrastructure - VSphereCluster/d2iq-e2e-cluster_name-1   True
13h
└─ControlPlane - KubeadmControlPlane/d2iq-control-plane           True
13h
| └─Machine/d2iq--control-plane-7llgd                               True
13h
| └─Machine/d2iq--control-plane-vncbl                               True
13h
| └─Machine/d2iq--control-plane-wbgrm                               True
13h
└─Workers
  └─MachineDeployment/d2iq--md-0                                     True
13h
    └─Machine/d2iq--md-0-74c849dc8c-67rv4                           True
13h
    └─Machine/d2iq--md-0-74c849dc8c-n2skc                             True
13h
    └─Machine/d2iq--md-0-74c849dc8c-nkftv                             True
13h
    └─Machine/d2iq--md-0-74c849dc8c-sqklv                             True
13h

```

After moving the cluster lifecycle services to the workload cluster, remember to use `dkp` with the workload cluster kubeconfig. Use `DKP` with the bootstrap cluster to delete the workload cluster.

4. Wait for the cluster control-plane to be ready:

```
kubectl --kubeconfig $HOME/.kube/config wait --for=condition=controlplaneready
"clusters/${CLUSTER_NAME}" --timeout=60m
```

The output should be similar to this example:

```
d2iq-e2e-cluster-1-control-plane/vsphere-example condition met
```


- Persistent Volumes (PVs) are not deleted automatically by design in order to preserve your data. However, they take up storage space if not deleted. You must delete PVs manually. Information for backup of a cluster and PVs is on the page in documentation called [Back up your Cluster's Applications and Persistent Volumes](#) (see page 835) .

With Vsphere clusters, `dkp delete` doesn't delete the virtual disks backing the PVs for DKP add ons. Therefore internal VMware cluster runs out of storage eventually. These PVs are only visible if VSAN is installed which gives users a Container Native Storage tab.

7.16.13.2 Delete the Workload Cluster

1. Make sure your vSphere credentials are up-to-date. Refresh the credentials using this command:

```
dkp update bootstrap credentials vsphere --kubeconfig $HOME/.kube/config
```

2. To delete a cluster, you would use `dkp delete cluster` and pass in the name of the cluster you are trying to delete with `--cluster-name` flag. You would use `kubectl get clusters` to get those details (`--cluster-name` and `--namespace`) of the Kubernetes cluster to delete it.

NOTE: Do not use `dkp get clusters` since that gets you Kommander cluster details rather than Konvoy kubernetes cluster details.

```
kubectl get clusters
```

3. Delete the Kubernetes cluster and wait a few minutes:

Before deleting the cluster, DKP deletes all Services of type LoadBalancer on the cluster.

To skip this step, use the flag `--delete-kubernetes-resources=false` .

```
dkp delete cluster --cluster-name=${CLUSTER_NAME} --kubeconfig $HOME/.kube/config
```

```
INFO[2022-03-30T11:53:42-07:00] Running cluster delete command
clusterName=d2iq-e2e-cluster-1 managementClusterKubeconfig= namespace=default
src="cluster/delete.go:95"
```

```
INFO[2022-03-30T11:53:42-07:00] Waiting for cluster to be fully deleted
src="cluster/delete.go:123"
INFO[2022-03-30T12:14:03-07:00] Deleted default/d2iq-e2e-cluster-1 cluster
src="cluster/delete.go:129"
```

7.16.13.3 Delete the Bootstrap Cluster

1. After the workload cluster is deleted, delete the bootstrap cluster with the following command.


```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

```
INFO[2021-06-09T12:15:20-07:00] Deleting bootstrap cluster
src="bootstrap/bootstrap.go:182"
```

7.16.13.4 Next Step:

Once your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will install the [Kommander](#) (see page 1422) component of DKP to see your dashboard and continue customization.

7.16.13.5 Known limitations

 Be aware of these limitations in the current release of DKP Konvoy.

- The DKP Konvoy version used to create the workload cluster must match the DKP Konvoy version used to delete the workload cluster.

7.16.14 Manage vSphere Node Pools

Node pools are part of a cluster and managed as a group, and you can use a node pool to manage a group of machines using the same common properties. When Konvoy creates a new default cluster, there is one node pool for the worker nodes, and all nodes in that new node pool have the same configuration.

You can create additional node pools for more specialized hardware or configuration. For example, if you want to tune your memory usage on a cluster where you need maximum memory for some machines and minimal memory on other machines, you would create a new node pool with those specific resource needs.

Konvoy implements node pools using Cluster API [MachineDeployments](#)⁹⁹⁰.

7.16.14.1 Create vSphere Node Pools

Creating a node pool is useful when you need to run workloads that require machines with specific resources, such as a GPU, additional memory, or specialized network or storage hardware.

7.16.14.1.1 Prepare the environment

Follow these steps:

1. Set the environment variable to the name you assigned this cluster with the command:

```
export CLUSTER_NAME=my-vsphere-cluster
```

2. If your workload cluster is self-managed, as described in [Make the New Cluster Self-Managed](#) (see [page 1282](#)), configure `kubectl` to use the kubeconfig for the cluster:

```
export KUBECONFIG=${CLUSTER_NAME}.conf
```

3. Define your node pool name:

```
export NODEPOOL_NAME=example
```

7.16.14.1.2 Create a vSphere node pool

Create a new vSphere node pool with 3 replicas using this command:

```
dkp create nodepool vsphere ${NODEPOOL_NAME} \
  --cluster-name=${CLUSTER_NAME} \
  --network=example_network \
  --data-center=example_datacenter \
  --data-store=example_datastore \
  --folder=example_folder \
  --server=example_vsphere_api_server_url \
  --resource-pool=example_resource_pool \
  --vm-template=example_vm_template \
```

⁹⁹⁰ <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

```
--replicas=3
```

The output resembles this example:

```
machinedeployment.cluster.x-k8s.io/example created
vspheremachinetemplate.infrastructure.cluster.x-k8s.io/example created
kubeadmconfigtemplate.bootstrap.cluster.x-k8s.io/example created
✓ Creating default/example nodepool resources
```

This example uses default values for brevity. Advanced users can use a combination of the `--dry-run` and `--output=yaml` or `--output-directory=<your-target-directory>/` flags to get a complete set of node pool objects to modify locally or store in version control.

7.16.14.2 List vSphere Node Pools

7.16.14.2.1 Listing node pools

Use this command to list the node pools of a given cluster. This returns specific properties of each node pool so that you can see the name of the MachineDeployments.

To list all node pools for a managed cluster, run the command:

```
dkp get nodepools --cluster-name=${CLUSTER_NAME} --kubeconfig=${CLUSTER_NAME}.conf
```

The expected output is similar to the following example, indicating the desired size of the node pool, the number of replicas ready in the node pool, and the Kubernetes version those nodes are running:

NODEPOOL VERSION	DESIRED	READY	KUBERNETES
demo-cluster-md-0	4	4	v1.26.6
example	3	0	v1.26.6

7.16.14.3 Scale vSphere Node Pools

7.16.14.3.1 Scaling node pools

While you can run [Cluster Autoscaler](#) (see page 1316), you can also manually scale your node pools up or down when you need more finite control over your environment. For example, if you require 10 machines to run a process, you can manually set the scaling to run those 10 machines only. However, if also using the Cluster Autoscaler, you must stay within your minimum and maximum bounds.

7.16.14.3.1.1 Scaling up node pools

To scale up a node pool in a cluster, run the command that follows, replacing the value 5 with the actual number of replicas you need:

```
dkp scale nodepools ${NODEPOOL_NAME} --replicas=5 --cluster-name=${CLUSTER_NAME}
```

Your output should be similar to this example, indicating the scaling is in progress:

```
INFO[2021-07-26T08:54:35-07:00] Running scale nodepool command
clusterName=demo-cluster managementClusterKubeconfig= namespace=default src="nodepool
/scale.go:82"
INFO[2021-07-26T08:54:35-07:00] Nodepool example scaled to 5 replicas
clusterName=demo-cluster managementClusterKubeconfig= namespace=default src="nodepool
/scale.go:94"
```

After a few minutes, you can list the node pools with the command:

```
dkp get nodepools --cluster-name=${CLUSTER_NAME} --kubeconfig=${CLUSTER_NAME}.conf
```

Your output should be similar to this example, with the number of DESIRED and READY replicas increased to 5:

NODEPOOL	DESIRED	READY	KUBERNETES
VERSION			
example	5	5	v1.26.6
demo-cluster-md-0	4	4	v1.26.6

7.16.14.3.1.2 Scaling down node pools

To scale down a node pool, run the command:

```
dkp scale nodepools ${NODEPOOL_NAME} --replicas=4 --cluster-name=${CLUSTER_NAME}
```

```
INFO[2021-07-26T08:54:35-07:00] Running scale nodepool command
clusterName=demo-cluster managementClusterKubeconfig= namespace=default src="nodepool
/scale.go:82"
INFO[2021-07-26T08:54:35-07:00] Nodepool example scaled to 4 replicas
clusterName=demo-cluster managementClusterKubeconfig= namespace=default src="nodepool
/scale.go:94"
```

In a default cluster, the nodes to delete are selected at random. This behavior is controller by [CAPI's delete policy](#)⁹⁹¹. However, when using the Konvoy CLI to scale down a node pool, you can specify the Kubernetes Nodes you want to delete.

To do this, set the flag `--nodes-to-delete` with a list of nodes as shown in the next command. This adds an annotation `cluster.x-k8s.io/delete-machine=yes` to the matching Machine object that contains `status.NodeRef` with the node names from `--nodes-to-delete`.

```
dkp scale nodepools ${NODEPOOL_NAME} --replicas=3 --nodes-to-delete=<> --cluster-
name=${CLUSTER_NAME}
```

```
INFO[2021-07-26T08:54:35-07:00] Running scale nodepool command
clusterName=demo-cluster managementClusterKubeconfig= namespace=default src="nodepool
/scale.go:82"
INFO[2021-07-26T08:54:35-07:00] Nodepool example scaled to 3 replicas
clusterName=demo-cluster managementClusterKubeconfig= namespace=default src="nodepool
/scale.go:94"
```

7.16.14.3.1.3 Scaling node pools when using cluster autoscaler

If you [configured the cluster autoscaler](#)⁹⁹² for the `demo-cluster-md-0` node pool, the value of `--replicas` must be within the minimum and maximum bounds.

For example, assuming you have the these annotations:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment ${NODEPOOL_NAME}
cluster.x-k8s.io/cluster-api-autoscaler-node-group-min-size=2
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment ${NODEPOOL_NAME}
cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size=6
```

⁹⁹¹ https://github.com/kubernetes-sigs/cluster-api/blob/v0.4.0/api/v1alpha4/machineset_types.go#L85-L105

⁹⁹² <https://docs.d2iq.com/dkp/konvoy/2.2/choose-infrastructure/vsphere/nodepools/cluster-autoscaler>

Try to scale the node pool to 7 replicas with the command:

```
dkp scale nodepools ${NODEPOOL_NAME} --replicas=7 -c demo-cluster
```

This action results in an error similar to:

```
INFO[2021-07-26T09:46:37-07:00] Running scale nodepool command
clusterName=demo-cluster managementClusterKubeconfig= namespace=default src="nodepool
/scale.go:82"
Error: failed to scale nodepool: scaling MachineDeployment is forbidden: desired
replicas 7 is greater than the configured max size annotation cluster.x-k8s.io/
cluster-api-autoscaler-node-group-max-size: 6
```

Similarly, scaling down to a number of replicas less than the configured `min-size` also returns an error.

7.16.14.4 vSphere Cluster Autoscaler

7.16.14.4.1 Cluster Autoscaler

[Cluster Autoscaler](#)⁹⁹³ provides the ability to automatically scale-up or scale-down the number of worker nodes in a cluster, based on the number of pending pods to be scheduled. Running the Cluster Autoscaler is optional.

Unlike [Horizontal-Pod Autoscaler](#)⁹⁹⁴, Cluster Autoscaler does not depend on any Metrics server and does not need Prometheus or any other metrics source.

The Cluster Autoscaler looks at the following annotations on a MachineDeployment to determine its scale-up and scale-down ranges:

```
cluster.x-k8s.io/cluster-api-autoscaler-node-group-min-size
cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size
```

The full list of command line arguments to the Cluster Autoscaler controller is found [this topic](#)⁹⁹⁵.

For more information about how Cluster Autoscaler works, see these documents:

- [What is Cluster Autoscaler](#)⁹⁹⁶

⁹⁹³ <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

⁹⁹⁴ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-fast-is-hpa-when-combined-with-ca>

⁹⁹⁵ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-are-the-parameters-to-ca>

⁹⁹⁶ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-is-cluster-autoscaler>

- [How does scale-up work](#)⁹⁹⁷
- [How does scale-down work](#)⁹⁹⁸
- [CAPI Provider for Cluster Autoscaler](#)⁹⁹⁹

7.16.14.4.1.1 Cluster Autoscaler Prerequisites

Before you begin, you must have:

- A [Bootstrap Cluster Lifecycle](#) (see page 1269).
- [Created a new Kubernetes Cluster](#) (see page 1271).
- A [Self-Managed Cluster](#) (see page 1282).

7.16.14.4.1.2 Run Cluster Autoscaler

The Cluster Autoscaler controller runs on the workload cluster. Upon creation of the workload cluster, this controller does not have all the objects required to function correctly until after a `dkp move` is issued from the bootstrap cluster.

Run the following steps to enable Cluster Autoscaler:

1. Ensure the Cluster Autoscaler controller is up and running (no restarts and no errors in the logs)

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf logs deployments/cluster-autoscaler
cluster-autoscaler -n kube-system -f
```

2. Enable Cluster Autoscaler by setting the min & max ranges

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment $
{NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-autoscaler-node-group-min-size=2
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment $
{NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size=6
```

3. The Cluster Autoscaler logs will show that the worker nodes are associated with node-groups and that pending pods are being watched.
4. To demonstrate that it is working properly, create a large deployment which will trigger pending pods.

```
cat <<EOF | kubectl --kubeconfig=${CLUSTER_NAME}.conf apply -f -
apiVersion: apps/v1
```

⁹⁹⁷ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-up-work>

⁹⁹⁸ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-down-work>

⁹⁹⁹ <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>


```

kind: Deployment
metadata:
  name: busybox-deployment
  labels:
    app: busybox
spec:
  replicas: 600
  selector:
    matchLabels:
      app: busybox
  template:
    metadata:
      labels:
        app: busybox
    spec:
      containers:
      - name: busybox
        image: busybox:latest
        command:
          - sleep
          - "3600"
        imagePullPolicy: IfNotPresent
        restartPolicy: Always
EOF

```

Cluster Autoscaler scales up the number of Worker Nodes until there are no pending pods.

5. Scale down the number of replicas for `busybox-deployment` with the command:

```

kubectl --kubeconfig ${CLUSTER_NAME}.conf scale --replicas=30 deployment/
busybox-deployment

```

6. Cluster Autoscaler starts to scale down the number of Worker Nodes after the default timeout of 10 minutes.



Enterprise Customers: In previous DKP versions, the Cluster-Autoscaler did not work on **managed** clusters. Starting with this release, an instance of Cluster-Autoscaler is deployed on the **management cluster** for each **managed** cluster. During an upgrade, the upgrade procedure on the **management cluster** will create the necessary deployments.

7.16.14.5 Replace a vSphere Node

7.16.14.5.1 Prerequisites

Before you begin, you must:

- [Create a workload cluster \(see page 1271\)](#) or [create an air-gapped workload cluster \(see page 1297\)](#).
- [Make the new cluster self-managed \(see page 1282\)](#).

7.16.14.5.2 Replace a worker node

In certain situations, you may want to delete a worker node and have [Cluster API](#)¹⁰⁰⁰ replace it with a newly-provisioned machine.

1. Identify the name of the node to delete.

List the nodes:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf get nodes
```

The output from this command resembles the following:

NAME	STATUS	ROLES	AGE
VERSION			
d2iq-e2e-cluster-1-control-plane-7llgd v1.26.6	Ready	control-plane,master	20h
d2iq-e2e-cluster-1-control-plane-vncbl v1.26.6	Ready	control-plane,master	20h
d2iq-e2e-cluster-1-control-plane-wbgrm v1.26.6	Ready	control-plane,master	19h
d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4 v1.26.6	Ready	<none>	20h
d2iq-e2e-cluster-1-md-0-74c849dc8c-n2skc v1.26.6	Ready	<none>	20h
d2iq-e2e-cluster-1-md-0-74c849dc8c-nkftv v1.26.6	Ready	<none>	20h
d2iq-e2e-cluster-1-md-0-74c849dc8c-sqklv v1.26.6	Ready	<none>	20h

2. Export a variable with the node name to use in the next steps:

This example uses the name `d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4`.

¹⁰⁰⁰ <https://cluster-api.sigs.k8s.io/>

```
export NAME_NODE_TO_DELETE="d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4"
```

3. Delete the Machine resource with the command:

```
NAME_MACHINE_TO_DELETE=$(kubectl --kubeconfig ${CLUSTER_NAME}.conf get machine
-ojsonpath="{.items[?
(@.status.nodeRef.name==\"$NAME_NODE_TO_DELETE\")].metadata.name}")
kubectl --kubeconfig ${CLUSTER_NAME}.conf delete machine
"$NAME_MACHINE_TO_DELETE"
```

```
machine.cluster.x-k8s.io "d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4" deleted
```

The command does not return immediately, but it does return after the Machine resource is deleted.

A few minutes after the Machine resource is deleted, the corresponding Node resource is also deleted.

4. Observe the Machine resource replacement using this command:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf get machinedeployment
```

NAME	UNAVAILABLE	PHASE	CLUSTER	AGE	VERSION	REPLICAS	READY	UPDATED
d2iq-e2e-cluster-1-md-0	1	ScalingUp	d2iq-e2e-cluster-1	20h	1.26.6	4	3	4

In this example, there exist 4 replicas, but only 3 are ready. One replica is unavailable, and the `ScalingUp` phase means a new Machine is being created.

5. Identify the replacement Machine using this command:

```
export NAME_NEW_MACHINE=$(kubectl --kubeconfig ${CLUSTER_NAME}.conf get
machines \
  -l=cluster.x-k8s.io/deployment-name=${CLUSTER_NAME}-md-0 \
  -ojsonpath='{.items[?(@.status.phase=="Provisioning")].metadata.name}
{"\n"}')
```

```
echo "$NAME_NEW_MACHINE"
```

If the output is empty, the new Machine has probably exited the `Provisioning` phase and entered the `Running` phase.

- Identify the replacement Node using this command:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf get nodes \
  -o=jsonpath="{.items[?(@.metadata.annotations.cluster\.x-k8s\.io/
  machine==\"$NAME_NEW_MACHINE\").metadata.name]}"
```

The output should be similar to this example:

```
d2iq-e2e-cluster-1-md-0-74c849dc8c-rc528
```

If the output is empty, the Node resource is not yet available, or does not yet have the expected annotation. Wait a few minutes, then repeat the command.

7.16.14.6 Delete vSphere Node Pools

Deleting a node pool deletes the Kubernetes nodes and the underlying infrastructure. DKP drains all nodes prior to deletion and reschedules the pods running on those nodes.

To delete a node pool from a managed cluster, run the command:

```
dkp delete nodepool ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME}
```

Here, `example` is the node pool to be deleted.

The expected output is similar to the following example, indicating the node pool is being deleted:

```
INFO[2021-07-28T17:14:26-07:00] Running nodepool delete command
Nodepool=example clusterName=d2iq-e2e-cluster-1 managementClusterKubeconfig=
namespace=default src="nodepool/delete.go:80"
```

Deleting an invalid node pool results in output similar to this example command output:

```
dkp delete nodepool ${CLUSTER_NAME}-md-invalid --cluster-name=${CLUSTER_NAME}
```

```
INFO[2021-07-28T17:11:44-07:00] Running nodepool delete command
Nodepool=demo-cluster-md-invalid clusterName=d2iq-e2e-cluster-1
managementClusterKubeconfig= namespace=default src="nodepool/delete.go:80"
Error: failed to get nodepool with name demo-cluster-md-invalid in namespace default
: failed to get nodepool with name demo-cluster-md-invalid in namespace default :
machinedeployments.cluster.x-k8s.io "demo-cluster-md-invalid" not found
```

7.17 VMware Cloud Director Infrastructure

Beta

VMware Cloud Director (VCD) is a cloud services platform that delivers virtual data center computation, network, storage, and security in a self-service model for tenants. This feature has been released as a Beta currently. For a definition of what to expect with Beta features, refer to our [Version Support Policy](#) (see page 1853) page in the documentation.

If not already done, refer to [Get Started](#) (see page 79) section of the documentation for:

- [Resource Requirements](#) (see page 100)
- [Install Overview](#) (see page 108)
- [Prerequisites for Install](#) (see page 121)

7.17.1 Special Resource Requirements for VCD

7.17.1.1 VM Sizing Policy - CPU and Memory

For CPU and Memory, the VCD **Provider** creates the appropriate VM Sizing Policies. The VM Sizing Policy defines CPU and memory resources available to the VM.

1. When creating the cluster, reference these VM Sizing Policies by using the `--control-plane-sizing-policy` and `--worker-sizing-policy` flags.
2. See <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-73C48A9C-79AF-402C-9746-4C81CAC2B35A.html> regarding parameters like *vCPU Speed* or *CPU Reservation Guarantee* that require consideration in VM Sizing Policy. For example, vCPU speed recommendation is a minimum of 3 GHz.

7.17.1.2 VM Placement Policy

The VCD Provider(SP) determines the [VM Placement Policy](#)¹⁰⁰¹. This policy defines the infrastructure where the VM should run meaning its placement of a VM on a host. When you assign a **VM Placement Policy** to a virtual machine, the placement engine adds this virtual machine to the corresponding VM group of the cluster on which it resides.

¹⁰⁰¹ <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-236A070E-83E6-4648-8F2F-557248C9735D.html>

7.17.1.3 Storage

For storage, follow [Create a Base OS image in vSphere](#) (see page 1263) when creating the KIB base image. In VCD, the base image determines the minimum storage available for the VM.



For more information about VCD parameters that can affect performance, refer to VMware Documentation: <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-73C48A9C-79AF-402C-9746-4C81CAC2B35A.html>

The following procedure advises on Organization setup as it relates to installing DKP and creating clusters on VCD.

- [Cloud Director for Service Providers](#) (see page 1323)
 - [Cloud Director Prerequisites](#) (see page 1325)
 - [Cloud Director Configure the Organization](#) (see page 1329)
 - [Cloud Director Roles, Rights and Rights Bundles](#) (see page 1332)
 - [Cloud Director CAPVCD User Rights](#) (see page 1335)
 - [Cloud Director Create Image and Template](#) (see page 1339)
 - [Cloud Director Bootstrap the Cluster](#) (see page 1340)
 - [Cloud Director Create a New Cluster](#) (see page 1342)
 - [Cloud Director Explore the Cluster](#) (see page 1346)
 - [Cloud Director Make the New Cluster Self-Managed](#) (see page 1349)
 - [Cloud Director Install Kommander](#) (see page 1352)
 - [Cloud Director Manage Node Pools](#) (see page 1354)
 - [Cloud Director Delete a Cluster](#) (see page 1361)

More information can be found directly in the [VMWare Cloud Director documentation](#)¹⁰⁰².

7.17.2 Next Step:

[Cloud Director for Service Providers](#) (see page 1323)

7.17.3 Cloud Director for Service Providers

Beta

¹⁰⁰² <https://docs.vmware.com/en/VMware-Cloud-Director/index.html>.

[VMware Cloud Director](#)¹⁰⁰³ is a platform that turns physical data center into virtual data centers and runs on one or more vCenters. After you have a working [vSphere vCenter](#)¹⁰⁰⁴, a brief overview of the general workflow for VMware Cloud Director is below:

- Use vCenter and VCD Service Provider Admin portal to create and configure your virtual infrastructure, including:
 - Creating the **Organization**¹⁰⁰⁵(tenant) and its users
 - Creating Roles, Rights, and other related permissions necessary for the tenant users and other software components
 - Creating Base Images and VM Templates that will be uploaded to the [vApp catalog](#)¹⁰⁰⁶ of the VCD tenant **Organization**
- Use DKP to create and manage clusters and node pools
- [Cloud Director Prerequisites](#) (see page 1325)
- [Cloud Director Configure the Organization](#) (see page 1329)
- [Cloud Director Roles, Rights and Rights Bundles](#) (see page 1332)
- [Cloud Director CAPVCD User Rights](#) (see page 1335)
- [Cloud Director Create Image and Template](#) (see page 1339)
- [Cloud Director Bootstrap the Cluster](#) (see page 1340)
- [Cloud Director Create a New Cluster](#) (see page 1342)
- [Cloud Director Explore the Cluster](#) (see page 1346)
- [Cloud Director Make the New Cluster Self-Managed](#) (see page 1349)
- [Cloud Director Install Kommander](#) (see page 1352)
- [Cloud Director Manage Node Pools](#) (see page 1354)
- [Cloud Director Delete a Cluster](#) (see page 1361)



To see a visual architecture structure as it relates to Tenants, Organization(or Tenant) VDCs and Provider VDCs for VMWare Cloud Director, refer to the page linked here for their documentation under the title: [New to VMware Cloud Director](#)¹⁰⁰⁷.

1003 <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/rn/vmware-cloud-director-104-release-notes/index.html>

1004 <https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.esxi.install.doc/GUID-B2F01BF5-078A-4C7E-B505-5DFFED0B8C38.html>

1005 <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-2F217F99-48C1-42F3-BF06-5ABBACADE2BA.html>

1006 <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-D5737821-C3A4-4C73-8959-CA293C12A7DE.html>

1007 <https://docs.vmware.com/en/VMware-Cloud-Director/index.html>

7.17.3.1 Related Topics:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20MSP%20Program&linkCreation=true&fromPagelId=273383945>¹⁰⁰⁸

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Version%20Support%20Policy&linkCreation=true&fromPagelId=273383945>¹⁰⁰⁹ - defines a Beta feature

7.17.3.2 Helpful VMWare Resources:

- [Setting Up Cloud Director Environment](#)¹⁰¹⁰ in GitHub
- <https://www.youtube.com/@VMwareCloudProvider>

7.17.3.3 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Cloud%20Director%20Prerequisites&linkCreation=true&fromPagelId=273383945>¹⁰¹¹

7.17.3.4 Cloud Director Prerequisites

Beta

Before continuing to install DKP on VMware Cloud Director(VCD), verify that your VMware vSphere Client environment is running [vSphere Client version](#)¹⁰¹² v6.7.x with Update 3 or later version with [ESXi](#)¹⁰¹³. You must be able to reach the vSphere API endpoint from where the Konvoy command line interface (CLI) runs. See the listed Prerequisites for all the specific details below.

7.17.3.4.1 Prerequisites Related to VMware Cloud Director

- Ensure you have meet VMware Cloud Director requirements: <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/rn/vmware-cloud-director-104-release-notes/index.html>

¹⁰⁰⁸ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273383945&linkCreation=true&spaceKey=DENT&title=%282.6%29+MSP+Program>

¹⁰⁰⁹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273383945&linkCreation=true&spaceKey=DENT&title=%282.6%29+Version+Support+Policy>

¹⁰¹⁰ https://github.com/mesosphere/cluster-api-provider-cloud-director/blob/fc216d559ff875da8ee0c61cf82454b759bb16aa/docs/VCD_SETUP.md#L44

¹⁰¹¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273383945&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Prerequisites>

¹⁰¹² https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html

¹⁰¹³ <https://docs.vmware.com/en/VMware-vSphere/index.html>

- Download and install VMware Cloud Director: [Download VMware Cloud Director](#)¹⁰¹⁴
- Considerations regarding CPU and Memory during configuration:
 - For CPU and Memory, the VCD **Provider** creates the appropriate *VM Sizing Policies*.
 - The **Provider** or tenant Organization user with cluster creation rights references these *VM Sizing Policies* when creating the cluster, using the `--control-plane-sizing-policy` and `--worker-sizing-policy` flags.
 - See <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-73C48A9C-79AF-402C-9746-4C81CAC2B35A.html> regarding parameters like *vCPU Speed* or *CPU Reservation Guarantee* that require consideration in VM Sizing Policy. For example, vCPU minimum speed is recommended to be at least 3 GHz.
- Considerations for Storage during configuration:
 - For storage, the VCD **Provider** must follow [Create a Base OS image in vSphere](#) (see page 1263) when creating the KIB base image. In VCD, the base image determines the minimum storage available for the VM.



Values required to create *VM Sizing Policies* and other fields are left to the Provider's discretion. For more information about VCD parameters that can affect performance, refer to VMware Documentation: <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-73C48A9C-79AF-402C-9746-4C81CAC2B35A.html>

7.17.3.4.2 Prerequisites for DKP

Before beginning a DKP installation, verify that you have:

- An x86_64-based Linux or macOS machine with a supported version of the operating system to execute the commands needed to run DKP and to create a cluster.
- [Download](#) (see page 78) the `dkp` binary for Linux or macOS with KIB image bundle. To check which version of DKP you installed for compatibility reasons, run the `dkp version -h` command ([dkp version](#) (see page 1749)).
- A Container engine/runtime installed is required to install DKP:
 - Version [Docker®](#)¹⁰¹⁵ container engine version 18.09.2 or higher installed for Linux or MacOS - On macOS, Docker runs in a virtual machine which needs configured with at least 8 GB of memory.

¹⁰¹⁴ https://customerconnect.vmware.com/downloads/info/slug/datacenter_cloud_infrastructure/vmware_cloud_director/10_4

¹⁰¹⁵ <https://docs.docker.com/get-docker/>

- Version 4.0 of [Podman](#)¹⁰¹⁶ or higher for Linux. Host requirements found here: [Host Requirements](#)¹⁰¹⁷
- [kubectl](#)¹⁰¹⁸ for interacting with the running cluster
- Valid VMware [Cloud Director version 10.4.x](#)¹⁰¹⁹ or later account as mentioned in Cloud Director prerequisites
- You as the Service Provider (SP) have created the appropriate images in vSphere vCenter available for upload to the tenant organization’s catalog:
 - a. Use `konvoy-image build vsphere` to create one or more OVA’s
 - b. Create a [vApp Template](#)¹⁰²⁰ from the OVA
 - c. Before the SP creates a cluster in a VCD tenant, make this vApp Template available to that tenant
- You must be able to reach the VCD API endpoint from where the DKP CLI runs

7.17.3.4.3 Machine Specifications

7.17.3.4.3.1 Control plane nodes

You must have at least three control plane nodes. Each control plane node should have at least:

- 4 cores
- 16 GiB memory
- Approximately 80 GiB of free space for the volume used for `/var/lib/kubelet` and `/var/lib/containerd`.
- Disk usage must be below 85% on the root volume.

DKP on VCD defaults to deploying an `n2-standard-4` instance with an 80GiB root volume for control plane nodes, which meets the above requirements.

7.17.3.4.3.2 Worker nodes

You must have at least four worker nodes. The specific number of worker nodes required for your environment can vary depending on the cluster workload and size of the nodes. Each worker node should have at least:

- 8 cores
- 32 GiB memory
- Around 80 GiB of free space for the volume used for `/var/lib/kubelet` and `/var/lib/containerd`.

¹⁰¹⁶ <https://podman.io/getting-started/installation>

¹⁰¹⁷ <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>

¹⁰¹⁸ <https://kubernetes.io/docs/tasks/tools/#kubectl>

¹⁰¹⁹ <https://docs.vmware.com/en/VMware-Cloud-Director/10.4.1/rn/vmware-cloud-director-1041-release-notes/index.html>

¹⁰²⁰ <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-D5737821-C3A4-4C73-8959-CA293C12A7DE.html?hWord=N4lghgNiBclG4EEAOSAEAXApgWyrMWAziAL5A>

- Disk usage must be below 85% on the root volume.

DKP on VCD defaults to deploying a `n2-standard-8` instance with an 80GiB root volume for worker nodes, which meets the above requirements.

7.17.3.4.4 Next Step:

[Cloud Director Configure the Organization](#) (see page 1329)

7.17.3.4.5 Cloud Director Concepts and Terms

Beta

Before attempting to create DKP cluster in VMware Cloud Director(VCD), we recommend taking time to familiarize yourself with the following concepts and related terminology.

7.17.3.4.5.1 VMware Cloud Director Concepts

The VMware Cloud Director is based on the following concepts:

- **Provider:** Service Provider(SP) that administrates the data centers and provisions virtual infrastructure for Organizations(tenants).
- **Organizations (Tenants):** A unit of administration for users, groups, and computing resources. Tenant users are managed at the **Organization** level.
- **System Administrators:** This role exists only in the Provider organization (SP) and can create and provision tenant **Organizations** in the [Service Provider](#)¹⁰²¹ as well as the [Organization](#)¹⁰²² portal. By default, the **System Administrator** role has all VMware Cloud Director rights.
- **Organization Administrators:** This role creates users, groups, and service catalogs. Tenant [Organization Administrator](#)¹⁰²³ is a predefined role that can use the [VCD tenant portal](#)¹⁰²⁴ to manage users in their **Organization** and assign them roles.
- **Rights:** Each right provides view or manage access to a particular object type in VCD. Also see: <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-816FBBBC-2CDA-4B1D-9B1A-C22BC31B46F2.html>
- **Rights Bundle:** A collection of rights for the **Organization** as a whole.
- **Roles:** A role is a set of rights that is assignable to one or more users and groups. When you create or import a user or group, you must assign it a role. Also see: <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/>

¹⁰²¹ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-F8F4B534-49B2-43B2-AEEE-7BAEE8CE1844.html>

¹⁰²² <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-120992A9-4FCB-4900-B19C-9AACFCB3F40B.html>

¹⁰²³ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-1CACBB2E-FE35-4662-A08D-D2BCB174A43C.html>

¹⁰²⁴ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-120992A9-4FCB-4900-B19C-9AACFCB3F40B.html>

[GUID-816FBBBC-2CDA-4B1D-9B1A-C22BC31B46F2.html](#) mentioned above or the section in DKP documentation: [Cloud Director Roles, Rights and Rights Bundles](#) (see page 1332)

- **Users and Groups:** Administrators can create users manually, programmatically, or integrate with a directory service like LDAP to import user accounts and user groups at scale.
- **Virtual Data Centers (VDC):** An isolated environment provided to a cloud user, in which they can provision resources, deploy, store and operate applications.
- **Organization VMware Cloud Director Networks:** Similar to the Amazon concept of Virtual Private Cloud, a VMware Cloud Director network is available only to a specific VMware Cloud Director and available to all vApps in the **Organization**. It can be connected to external networks as needed.
- **vApp Networks:** Similar to the concept of a subnet, a vApp network is an isolated network within a VMware Cloud Director network that allows specific vApps to communicate with each other.
- **vApp:** One or more virtual machines (VMs) that come preconfigured to provide a specific type of cloud service. vApp is a virtual app that defines compute, storage and networking metadata.
- **Media Files and Catalogs:** VMware Cloud Director organizes deployable resources via media files. These are virtual machine and vApp templates (machine images) that can be used as a sort of initial boot program for a VM. The **Organization Administrator** organizes these files into catalogs, allowing users within the **Organization** to provision the resources they need.
- **Storage Profiles:** VCD concept to organize [storage](#)¹⁰²⁵ (ex. Gold, Platinum)
- **NSXT gateway:** A Logical Router configured in a traditional hardware switch. Gateways are used to provide connectivity to external networks and between different logical networks.
- **Tier-0 Gateway:** The Tier-0 gateway provides connectivity to external networks, using static routes or BGP. The Tier-0 gateway is primarily in charge of handling the North-South traffic between the virtualized environment and the external physical network.
- **Tier-1 Gateway:** The Tier-1 Gateway acts as a tenant router. The Tier-1 gateway is optimized for East-West traffic.
- **Edge Gateway:** A gateway that provides a VDC with connectivity and other features. It can provide NAT, firewall and other network features.
- **Provider Gateway:** A logical gateway representing Tier-0 gateway and managed by the SP.
- **Organization Edge Gateway:** An organization specific gateway in the Provider Gateway. They are created using network segments in the provider Gateway.

7.17.3.4.5.2 Next Step:

[Cloud Director Configure the Organization](#) (see page 1329)

7.17.3.5 Cloud Director Configure the Organization

Beta

¹⁰²⁵<https://docs.vmware.com/en/VMware-Cloud-Director/10.5/VMware-Cloud-Director-Tenant-Guide/GUID-17DDD3AC-0ABB-49EC-B3AE-28DFB2D2B80B.html?hWord=N4IghgNiBcIMoBcD2AnMBzApgAgMITAGdCQBfIA>

The following steps document the minimum configuration needed when creating and configuring an **Organization** for use with DKP. Refer to the [VCD documentation](#)¹⁰²⁶ for additional tenant **Organization** configuration information.

The two main sections of the VMware Cloud Director documentation are for system administration and the tenant configuration and access:

- **Service Provider** Admin Portal - <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-F8F4B534-49B2-43B2-AEEE-7BAEE8CE1844.html>
- **Organization** (tenant) Portal - <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-120992A9-4FCB-4900-B19C-9AACFCB3F40B.html>

7.17.3.5.1 Prerequisites to Configure an Organization

1. Add a [vCenter Server](#)¹⁰²⁷ to VCD - Log in to the Cloud Director portal, select the **Resources** tab across the top menu - **Infrastructure Resources** - **vCenter Server Instances** and select <ADD> to begin the process of connecting your [vCenter server](#)¹⁰²⁸ by following the on screen instructions.
2. Create **Organization**¹⁰²⁹ (Tenant) - You [create a new organization](#)¹⁰³⁰ from the VMware Cloud Director Admin Portal.

7.17.3.5.2 Configuring the Organization

1. Create the **Organization's VDCs**¹⁰³¹
2. Configure **Edge Gateway**¹⁰³²
3. Create the **Tenant Network**¹⁰³³
4. Configure **Policies**¹⁰³⁴

¹⁰²⁶ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-1AE706A9-79F2-4A0E-8B0D-FFE9E87109CD.html>

¹⁰²⁷ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-F107933D-0E66-497E-BCB1-34A670E79D7C.html>

¹⁰²⁸ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-88470D75-4899-45DF-B01D-49C847CA4945.html>

¹⁰²⁹ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-03E72753-1BEB-45C0-BAF8-AC71A089F14E.html>

¹⁰³⁰ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-03E72753-1BEB-45C0-BAF8-AC71A089F14E.html>

¹⁰³¹ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-09851831-142E-46B9-A278-6488784D5B6A.html>

¹⁰³² <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-45C0FEDF-84F2-4487-8DB8-3BC281EB25CD.html>

¹⁰³³ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-8F806B38-2489-4D36-82FF-B23BAFC3B294.html>

¹⁰³⁴ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-DBA11D01-E102-47A4-926C-BFDB681B75F2.html>

After a tenant **Organization** is created, use the menus in the **System Administrator** portal to configure the following settings:

1. Under the Data Center tab - **Virtual Data Center**: this is the location in which to define CPU size, memory and storage.
2. Under the Data Center tab - **Networking - Edges**: select **Configuration-External Networks** and supply the publicly accessible IP address range under Subnet column.
3. Under the Libraries tab - **Content Libraries** in left menu: specify the *vApp Templates* to import the VM Templates from vCenter that you want to make available to the tenant. EX: KIB templates from vCenter
NOTE: The Service Provider(SP) can create a shared *Catalog* where items are placed to be automatically imported.
4. Under the Networking tab - **Networks**: the tenant **Organization Administrator** will configure a network. Select the name of the network to be taken to its *General* properties such as Gateway CIDR address where all VMs will receive an IP address from the private IP space.
NOTE: The LoadBalancer (LB) will use the routable external network and is automatically created using the CAPVCD controller.
5. Under the Resources tab - **Edge Gateway - Services - NAT - External IP**: the Specify the IP address that will be used to allow VMs to access external networks. You should either create a SNAT rule or provide Egresses to the VM's.
6. **Edge Gateway Firewall** - Allow port access from outside the cluster for TCP 6443 to the control plane endpoint load balancers and TCP 443 to Kubernetes load balancers (e.g. to reach the DKP dashboard). Other ports must be allowed access among the machines in the cluster.
 For more DKP Ports, refer to that section under Architecture [DKP Ports](#) (see page 67).



The tenant is required to get one public IP to create an Edge Gateway. The Service Provider(SP) will allocate the pool of IPs from which the tenant pulls. After you have associated an external network through a gateway, the tenant can ask for IPs. Otherwise, if you have chosen an IP address, you can specify it.

7.17.3.5.3 Production Policies

After the tenant **Organization** is in production, various policies will need to be defined for storage and resources. The [Configure Organization Policy Section](#)¹⁰³⁵ of VMware documentation will provide more detail.

¹⁰³⁵ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-DBA11D01-E102-47A4-926C-BFDB681B75F2.html>

7.17.3.5.4 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Cloud%20Director%20Roles%2C%20Rights%20and%20Rights%20Bundles&linkCreation=true&fromPageId=273646407>¹⁰³⁶

7.17.3.6 Cloud Director Roles, Rights and Rights Bundles

Beta

VMware Cloud Director has a variety of ways to assign permissions. As a Service Provider(Cloud Provider), you have numerous options when it comes to giving tenants access to VMware Cloud Director (VCD) features. There are:

- **Rights**¹⁰³⁷ - provide view or manage access to a particular object type in VMware Cloud Director and belong to different categories depending on the objects to which they relate such as vApp, Catalog, or an Organization
- **Roles**¹⁰³⁸ - a collection of rights for a User and defines what an individual user has access to
- **Rights Bundles**¹⁰³⁹ - a collection of rights for the tenant **Organization** as a whole and defines what a tenant **Organization** has access to

Various **Rights** are common to multiple predefined Global Roles. These **Rights** are granted by default to all new organizations, and are available for use in other **Roles** created by the tenant **Organization Administrator**. There are some predefined **Roles** for both Provider and Tenant explained in the VMware documentation: <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-BC504F6B-3D38-4F25-AACF-ED584063754F.html>.

7.17.3.6.1 Service Provider(SP) System Administrator

The **System Administrator**¹⁰⁴⁰ role exists only in the provider organization.

As a Service Provider(SP) you will have roles for vSphere vCenter as well as VMware Cloud Director(VCD).

- **vCenter/NXT/AVI Infrastructure System Administrator**: Manages physical infra for vCenter, NXT network fabric, AVI load balancers (D2IQ SRE / Service provider(OVH cloud))
- **System Administrator - Provider(SP)** : Manages Virtual infra in VCD that uses vCenter(s), NXT(s), AVI(s) etc. (D2IQ SRE)

¹⁰³⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273646407&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Roles%2C+Rights+and+Rights+Bundles>

¹⁰³⁷ <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-9B462637-E350-43B6-989A-621F226A56D4.html>

¹⁰³⁸ <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-9B462637-E350-43B6-989A-621F226A56D4.html>

¹⁰³⁹ <https://blogs.vmware.com/cloudprovider/2019/12/effective-rights-bundles.html>

¹⁰⁴⁰ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-438B2F8C-65B0-4895-AF40-6506E379A89D.html>

- **Organization(Tenant) Administrator:** Manages Virtual infra (org, orgvdc, network, catalogs, templates, users etc) for a tenant. Users that can create k8s cluster

Through the VMware Cloud Director(VCD) [Service Provider Admin Portal](#)¹⁰⁴¹, the SP can add **System Administrators** for Cloud Director and see the predefined list of rights included in any role. The **System Administrator** manages the virtual infrastructure in VCD that uses vCenter(s), NXT(s), AVI(s) and other components of the VCD environment.

Also, refer to [Managing System Administrator and Roles](#)¹⁰⁴² documentation from the VMware documentation site.

7.17.3.6.2 Organization(Tenant) Administrator

As an **Organization Administrator**, from the [tenant portal](#)¹⁰⁴³ you can create, edit, import, and delete users. The tenant **Organization Administrator** manages the virtual infrastructure that includes the organization itself which includes the related network, catalogs, templates and such for the Tenant Organization. Important predefined role information is below:

- Tenant **Organization Administrator**¹⁰⁴⁴ is predefined and can use VCD tenant portal to [create](#)¹⁰⁴⁵ and [manage users](#)¹⁰⁴⁶ in their organization and assign them roles.
- **vApp Author**¹⁰⁴⁷ role can use catalogs and create vApps

A tenant **Organization Administrator** can access roles if allowed. They can only view the global tenant roles that a **System Administrator** has published to the organization, but cannot modify them. The **Organization Administrator** can create **custom** tenant roles with similar rights and assign them to the users within their own tenant **Organization**.

There are some predefined Global Tenant Roles as well which are explained in the VMware documentation: <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-BC504F6B-3D38-4F25-AACF-ED584063754F.html> as well as the <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-AE42A8F6-868C-4FC0-B224-87CA0F3D6350.html>.

¹⁰⁴¹ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-F8F4B534-49B2-43B2-AEEE-7BAEE8CE1844.html>

¹⁰⁴² <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-9DFB2238-23FB-4D07-B563-144AC4E9EDAF.html>

¹⁰⁴³ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-74C9E10D-9197-43B0-B469-126FFBCB5121.html>

¹⁰⁴⁴ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-BC504F6B-3D38-4F25-AACF-ED584063754F.html#GUID-BC504F6B-3D38-4F25-AACF-ED584063754F>

¹⁰⁴⁵ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-1CACBB2E-FE35-4662-A08D-D2BCB174A43C.html>

¹⁰⁴⁶ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-A358E190-BFC0-4187-9406-66E82C92564A.html>

¹⁰⁴⁷ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-BC504F6B-3D38-4F25-AACF-ED584063754F.html>

7.17.3.6.3 Tenant Roles and Rights

Several predefined global tenant roles are described in the [VMware documentation](#)¹⁰⁴⁸ regarding which components they can access and change. User rights may be available to a role, but a rights bundle needs to be published to a tenant Organization in order for those user role permissions to work.

1. Create a [Rights Bundle](#)¹⁰⁴⁹ (a collection of rights) that includes all required Rights. See the page for the <https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Cloud%20Director%20CAPVCD%20User%20Rights&linkCreation=true&fromPagelId=273745264>^{1050, 1051}.
2. Create a [Tenant Role](#)¹⁰⁵² that uses all these Rights.

The SP should publish these to tenant **Organizations** that need to create clusters. Assign the Tenant Role to a VCD user and then create an API Token. The API token will be used in the DKP CLI commands to authenticate the CLI to vCenter. They will assign the Role to a VCD User, create an API Token, and pass the Token to the DKP CLI.

7.17.3.6.4 Related Topics:

The CAPVCD provider uses a related component called [CSE](#)¹⁰⁵³. Some of the permissions necessary to create a VCD cluster are defined using this component. Note that the term [Role Based Access Control \(RBAC\)](#)¹⁰⁵⁴ used in the CSE documentation refers ONLY to the VCD rights and permissions necessary to perform lifecycle management of Kubernetes clusters using VCD. It has no impact on the RBAC configuration of any clusters created using VCD

[Role Based Access Control \(RBAC\)](#)¹⁰⁵⁵ from GitHub - The RBAC in that page refers to the roles and rights required for the tenants to perform the life cycle management of Kubernetes clusters. It does not have anything to do with the RBAC inside the Kubernetes cluster itself.

1048 <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-AE42A8F6-868C-4FC0-B224-87CA0F3D6350.html>

1049 <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-CFB0EFEE-0D4C-498D-A937-390811F11B8E.html>

1050 <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273745264&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+CAPVCD+User+Rights>

1051 <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273745264&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+CAPVCD+User+Rights>

1052 <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-0D991FCF-3800-461D-B123-FAE7CFF34216.html>

1053 https://vmware.github.io/container-service-extension/cse3_1/INTRO.html

1054 https://vmware.github.io/container-service-extension/cse3_1/RBAC.html#additional-required-rights

1055 https://vmware.github.io/container-service-extension/cse3_1/RBAC.html#additional-required-rights

7.17.3.6.5 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Cloud%20Director%20CAPVCD%20User%20Rights&linkCreation=true&fromPageId=273745264>¹⁰⁵⁶

7.17.3.7 Cloud Director CAPVCD User Rights

Beta

CAPVCD requires specific user rights set in order for the tenant Organization to have a **Role** that can successfully execute DKP clusters, those **Rights** must be specified. When a DKP workload cluster is created, each machine needs to register with **VCD Cloud Provider Interface(CPI)**¹⁰⁵⁷ and get node references.

Refer to <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-816FBBBC-2CDA-4B1D-9B1A-C22BC31B46F2.html> for terminology related to Rights and Roles. The remainder of this page describes what the CAPVCD User requires in relation to **Rights** and **Rights Bundles**.

7.17.3.7.1 CAPVCD User

CAPVCD uses the credential (username/password or API token) of a VCD User to manage the cluster. The **VCD Cloud Provider Interface(CPI)**¹⁰⁵⁸ and CSI controllers also use the same credential. This same User needs specific API permissions, known as **Rights**, for the CAPVCD, CPI, and CSI controllers to work correctly. For a User to be granted a **Right**, the User must be associated with a **Role** that consumes this **Right**, meaning the **Role** grants **Rights** to the User.

If the User belongs to an **Organization**, then the **Right** must *also* be published by the **Provider** to the tenant **Organization** in a Rights Bundle. The Rights Bundle grants Rights to the tenant **Organization**.

7.17.3.7.1.1 Create the VCD User Required by CAPVCD

CAPVCD requires all the Rights in the default `vApp Author` **Global Role**¹⁰⁵⁹, plus the **Rights** required by the VCD CPI, the VCD CSI, and some **Rights** required by CAPVCD itself.

1. A Provider administrator creates a **Rights Bundle**¹⁰⁶⁰ that enumerates all the below rights. We recommend the name `DKP Cluster Admin` for the **Rights Bundle**.

¹⁰⁵⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273745264&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+CAPVCD+User+Rights>

¹⁰⁵⁷ <https://github.com/vmware/cloud-provider-for-cloud-director/tree/main#terminology>

¹⁰⁵⁸ <https://github.com/vmware/cloud-provider-for-cloud-director/tree/main#terminology>

¹⁰⁵⁹ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-BC504F6B-3D38-4F25-AACF-ED584063754F.html>

¹⁰⁶⁰ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-CFB0EFEE-0D4C-498D-A937-390811F11B8E.html>

2. A Provider administrator creates a **Global Role** that enumerates all the below rights. We recommend the name `DKP Cluster Admin` for the **Global Role**.
3. A Provider administrator publishes the both the **Rights Bundle** and **Global Role** to every Organization that will deploy DKP clusters.
4. An Organization administrator creates a User, and associates it with the **Global Role**.

The procedure for Creating a Rights Bundle is found in VMware Documentation: <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-CFB0EFEE-0D4C-498D-A937-390811F11B8E.html> and the steps for creating the Role: <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Service-Provider-Admin-Portal-Guide/GUID-0D991FCF-3800-461D-B123-FAE7CFF34216.html>

7.17.3.7.2 List of Rights

CAPVCD requires the following **Rights**:

- The **Rights** cataloged by the predefined `vApp Author` **Global Role**.¹⁰⁶¹
- Additional **Rights** listed in the [CAPVCD documentation](#)¹⁰⁶².
- Additional **Rights** listed in the [VCD CPI documentation](#)¹⁰⁶³.
- Additional **Rights** listed in the [VCD CSI documentation](#)¹⁰⁶⁴.

The majority of the **Rights** are from the `vApp Author` Global Role. When independent components, like CAPVCD and CPI, need the same **Right**, that **Right** appears in multiple sources.

Below are the lists of the **Rights** from the above sources as well as **Rights** required, but not documented by CAPVCD. The last list includes the **Rights** from all sources, with duplicates removed:

7.17.3.7.2.1 Required by the `vApp Author` Role

Some Rights appear in more than one source. This list includes the Rights from all sources, with duplicates removed in the final list below:

```
Catalog: Add vApp from My Cloud
Catalog: View Private and Shared Catalogs
Organization vDC Compute Policy: View
Organization vDC Disk: View IOPS
Organization vDC Named Disk: Create
Organization vDC Named Disk: Delete
Organization vDC Named Disk: Edit Properties
```

¹⁰⁶¹ <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-AE42A8F6-868C-4FC0-B224-87CA0F3D6350.html#GUID-AE42A8F6-868C-4FC0-B224-87CA0F3D6350>

¹⁰⁶² <https://github.com/vmware/cloud-director-named-disk-csi-driver/blob/1.3.2/README.md#additional-rights-for-csi>

¹⁰⁶³ <https://github.com/vmware/cloud-provider-for-cloud-director/blob/1.3.0/README.md#additional-rights-for-cpi>

¹⁰⁶⁴ https://github.com/vmware/cluster-api-provider-cloud-director/blob/b315c9c5e4b1b05600ccec20eb5116cc6b173f845/docs/VCD_SETUP.md#publish-the-rights-to-the-tenant-organizations

```

Organization vDC Named Disk: View Encryption Status
Organization vDC Named Disk: View Properties
Organization vDC Network: View Properties
Organization vDC: VM-VM Affinity Edit
Organization: View
UI Plugins: View
vApp Template / Media: Copy
vApp Template / Media: Edit
vApp Template / Media: View
vApp Template: Checkout
VAPP_VM_METADATA_TO_VCENTER
vApp: Copy
vApp: Create / Reconfigure
vApp: Delete
vApp: Download
vApp: Edit Properties
vApp: Edit VM Compute Policy
vApp: Edit VM CPU
vApp: Edit VM Hard Disk
vApp: Edit VM Memory
vApp: Edit VM Network
vApp: Edit VM Properties
vApp: Manage VM Password Settings
vApp: Power Operations
vApp: Sharing
vApp: Snapshot Operations
vApp: Upload
vApp: Use Console
vApp: View ACL
vApp: View VM and VM's Disks Encryption Status
vApp: View VM metrics
vApp: VM Boot Options

```

7.17.3.7.2.2 Additional Rights Required by CAPVCD

```

API Tokens: Manage
Organization vDC Gateway: Configure Load Balancer
Organization vDC Gateway: Configure NAT
Organization vDC Gateway: View
vApp: Allow All Extra Config

```

7.17.3.7.2.3 Rights Required to List Catalogs (Not Documented by CAPVCD)

These Rights are not documented by CAPVCD, but are required. They allow CAPVCD to list the Catalogs in the Organization.

```

General: Administrator View

```

Access All Organization VDCs

7.17.3.7.2.4 Implied Rights (Not Documented by CAPVCD)

These Rights are not documented by CAPVCD, but are implied by the documented Rights.

Certificate Library: View
Organization vDC Gateway: View NAT
Organization vDC Gateway: View Load Balancer

7.17.3.7.2.5 Additional Rights Required by CPI

API Tokens: Manage
Organization vDC Gateway: Configure Load Balancer
Organization vDC Gateway: Configure NAT
Organization vDC Gateway: View

7.17.3.7.2.6 Additional Rights Required by CSI

API Tokens: Manage
Organization vDC Shared Named Disk: Create
Organization vDC Named Disk: Delete
Organization vDC Named Disk: Edit Properties
Organization vDC Named Disk: View Encryption Status
Organization vDC Named Disk: View Properties

7.17.3.7.2.7 All Sources Merged with Duplicates Removed

Access All Organization VDCs
API Tokens: Manage
Catalog: Add vApp from My Cloud
Catalog: View Private and Shared Catalogs
Certificate Library: View
General: Administrator View
Organization vDC Compute Policy: View
Organization vDC Disk: View IOPS
Organization vDC Gateway: Configure Load Balancer
Organization vDC Gateway: Configure NAT
Organization vDC Gateway: View
Organization vDC Gateway: View Load Balancer
Organization vDC Gateway: View NAT

```

Organization vDC Named Disk: Create
Organization vDC Named Disk: Delete
Organization vDC Named Disk: Edit Properties
Organization vDC Named Disk: View Encryption Status
Organization vDC Named Disk: View Properties
Organization vDC Network: View Properties
Organization vDC Shared Named Disk: Create
Organization vDC: VM-VM Affinity Edit
Organization: View
UI Plugins: View
vApp Template / Media: Copy
vApp Template / Media: Edit
vApp Template / Media: View
vApp Template: Checkout
VAPP_VM_METADATA_TO_VCENTER
vApp: Allow All Extra Config
vApp: Copy
vApp: Create / Reconfigure
vApp: Delete
vApp: Download
vApp: Edit Properties
vApp: Edit VM Compute Policy
vApp: Edit VM CPU
vApp: Edit VM Hard Disk
vApp: Edit VM Memory
vApp: Edit VM Network
vApp: Edit VM Properties
vApp: Manage VM Password Settings
vApp: Power Operations
vApp: Sharing
vApp: Snapshot Operations
vApp: Upload
vApp: Use Console
vApp: View ACL
vApp: View VM and VM's Disks Encryption Status
vApp: View VM metrics
vApp: VM Boot Options

```

7.17.3.7.3 Next Step:

[Cloud Director Create Image and Template \(see page 1339\)](#)

7.17.3.8 Cloud Director Create Image and Template

Beta

7.17.3.8.1 Create Image for Tenants

As a vSphere vCenter administrator, you will create the images for your VMWare Cloud Director (VCD) tenants.

1. The [vSphere Base OS Image](#) (see page 1263) and [VM Template](#) (see page 1265) are made in vCenter for Ubuntu 20.04. For storage, the base image determines the minimum storage available for the VM.
2. An OVA would be created from image for each VCD tenant accordingly.
3. That image is exported from vSphere and imported into the catalog of the VCD tenant **Organization** catalog of [vApp Templates](#)¹⁰⁶⁵.

Refer to the VMWare documentation for specifics on creating and deploying OVA for your tenants: [Deploying OVF and OVA Templates](#)¹⁰⁶⁶ and [Working with vApp Templates](#)¹⁰⁶⁷.

If you have not already done so, create your base image and template in vSphere and import into the VCD tenant catalog to make available for tenant use.

7.17.3.8.2 Next Step:

[Cloud Director Bootstrap the Cluster](#) (see page 1340)

7.17.3.9 Cloud Director Bootstrap the Cluster

Beta

To create Kubernetes clusters, DKP uses [Cluster API](#)¹⁰⁶⁸ (CAPI) controllers. These controllers run on a Kubernetes cluster. To get started, you need a *bootstrap* cluster. Bootstrapping is the process of creating a new Kubernetes cluster from scratch and involves setting up the control plane and worker nodes. It also determines which node has the correct information for synchronization with all other nodes.

7.17.3.9.1 Prerequisites

Before you begin, you must:

- Complete the steps in [Prerequisites](#) (see page 1328).
- Ensure the `dkp` binary can be found in your `$PATH`.
- For tenants, you need to set their Organization name and Network name in order to create a cluster.

¹⁰⁶⁵ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-D5737821-C3A4-4C73-8959-CA293C12A7DE.html>

¹⁰⁶⁶ https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-AFEDC48B-C96F-4088-9C1F-4F0A30E965DE.html

¹⁰⁶⁷ <https://docs.vmware.com/en/VMware-Cloud-Director/10.3/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-D5737821-C3A4-4C73-8959-CA293C12A7DE.html>

¹⁰⁶⁸ <https://cluster-api.sigs.k8s.io/>

7.17.3.9.2 Bootstrap Cluster Lifecycle Services

1. If an HTTP proxy is required for the bootstrap cluster, set the local `http_proxy`, `https_proxy`, and `no_proxy` environment variables. They are copied into the bootstrap cluster.
2. Create a bootstrap cluster:

```
dkp create bootstrap
```

3. The output looks similar to:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

4. DKP creates a bootstrap cluster using [KIND](#)¹⁰⁶⁹ as a library. DKP then deploys the following [Cluster API](#)¹⁰⁷⁰ providers on the cluster:
 - [Core Provider](#)¹⁰⁷¹
 - [Cloud Director Infrastructure Provider](#)¹⁰⁷²
 - [Kubeadm Bootstrap Provider](#)¹⁰⁷³
 - [Kubeadm ControlPlane Provider](#)¹⁰⁷⁴
5. DKP waits until the controller-manager and webhook deployments of these providers are ready. List these deployments using this command:

```
kubectl get --all-namespaces deployments -l=clusterctl.cluster.x-k8s.io
```

You will then receive the following output:

NAMESPACE	READY	UP-TO-DATE	AVAILABLE	AGE	NAME
capa-system	1/1	1	1	1h	capa-controller-manager
capg-system	1/1	1	1	1h	capg-controller-manager
capi-kubeadm-bootstrap-system	1/1	1	1	1h	capi-kubeadm-bootstrap-controller-manager

¹⁰⁶⁹ <https://github.com/kubernetes-sigs/kind>

¹⁰⁷⁰ <https://cluster-api.sigs.k8s.io/>

¹⁰⁷¹ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

¹⁰⁷² <https://github.com/vmware/cluster-api-provider-cloud-director>

¹⁰⁷³ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

¹⁰⁷⁴ <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>

capi-kubeadm-control-plane-system manager	1/1	1	1	capi-kubeadm-control-plane-controller- 1h
capi-system				capi-controller-manager
1/1	1	1	1h	
capp-system				capp-controller-manager
1/1	1	1	1h	
capv-system				capv-controller-manager
1/1	1	1	1h	
capz-system				capz-controller-manager
1/1	1	1	1h	
capvcd-system				capvcd-controller-manager
1/1	1	1	1h	
cert-manager				cert-manager
1/1	1	1	1h	
cert-manager				cert-manager-cainjector
1/1	1	1	1h	
cert-manager				cert-manager-webhook
1/1	1	1	1h	

7.17.3.9.3 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Cloud%20Director%20Create%20a%20New%20%20Cluster&linkCreation=true&fromPageId=273648304>¹⁰⁷⁵

7.17.3.10 Cloud Director Create a New Cluster

Beta


Before beginning to create your new cluster, confirm you can reach any of the IPs in the range allocated to the load balancer pool so the CAPVCD drivers can reach it during cluster creation.

7.17.3.10.1 Flags Specific to VMware Cloud Director Cluster Creation

When creating a VMware Cloud Director(VCD) cluster, CPU and Memory flags are needed to reference the VM Sizing Policies:

- For CPU and Memory, the VCD **Provider** creates the appropriate VM Sizing Policies. Then either the Provider or Tenant User with cluster creation rights references these VM Sizing Policies when creating the cluster by using those flags:
 - `--control-plane-sizing-policy`
 - `--worker-sizing-policy`

¹⁰⁷⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273648304&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Create+a+New+Cluster>

-  If the Service Provider(SP) has given a tenant user the [permissions \(see page 1335\)](#) to create clusters inside their own **Organization**, then that tenant user will need to reference those flags are well.


7.17.3.10.2 Name your Cluster

1. Give your cluster a unique name suitable for your environment. The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes¹⁰⁷⁶](#) for more naming information.

In Cloud Director it is critical that the name is unique, as no two clusters in the same Organization can have the same name.


2. Set the environment variable:

```
export CLUSTER_NAME=vcd-example
```

-  To increase [Docker Hub's rate limit¹⁰⁷⁷](#) use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io` `--registry-mirror-username=<username>` `--registry-mirror-password=<password>` on the `dkp create cluster` command.

7.17.3.10.3 Create a New Cloud Director Cluster

After you have met all the prerequisites, begin creating a cluster:

-  Cloud Director needs the [Base OS image \(see page 1487\)](#) and [VM Template \(see page 1489\)](#) from vSphere to be converted to a vApp Template and listed in the tenant **Organization's** catalog.

¹⁰⁷⁶ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

¹⁰⁷⁷ <https://docs.docker.com/docker-hub/download-rate-limit/>

1. Create a [vApp template](#)¹⁰⁷⁸ using the instructions from the previous section of documentation: [Cloud Director Create Image and Template](#) (see page 1342)
2. Set Refresh API token using the instructions from VMware:

```
export VCD_REFRESH_TOKEN=<REFRESH TOKEN>
```

- ⓘ** You must provide the refresh tokens (<REFRESH TOKEN>) when provisioning a cluster. See the VMware Documentation for details: <https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-A1B3B2FA-7B2C-4EE1-9D1B-188BE703EEDE.html#procedure-2>

3. Create a Kubernetes cluster:
Reference the VM Sizing Policies during cluster creation using the flags:

- `--control-plane-sizing-policy`
- `--worker-sizing-policy`

```
dkp create cluster vcd --cluster-name ${CLUSTER_NAME} \
--site "" \
--organization "" \
--data-center "" \
--catalog "" \
--vapp-template "" \
--network "" \
--ssh-public-key-file "" \
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml
```

- (Optional) Alternatively, you can create individual files with different smaller manifests for ease in editing using the `--output-directory` flag. This will create multiple files in the specified directory which must already exist:

```
--output-directory=<existing-directory>
```

Description of parameters:

- `--site` : the cloud director's endpoint with the format https://VCD_HOST.
- `--organization` : the tenant Organization name

¹⁰⁷⁸<https://docs.vmware.com/en/VMware-Cloud-Director/10.4/VMware-Cloud-Director-Tenant-Portal-Guide/GUID-D5737821-C3A4-4C73-8959-CA293C12A7DE.html?hWord=N4lghgNiBcIG4EEAOSAEAXApgWyrMWAziAL5A>

- `--data-center` : the Virtual Data Center(VDC) name
- `--catalog` : the folder name where vApp templates are located
- `--network` : your private network

See [dkp create cluster vcd](#) (see page 1680) for Cloud Director reference to see the full list of cluster creation options and their descriptions.

4. Create the cluster from the objects. A warning will appear in the console if the resource already exists and will require you to remove the resource or update your YAML.

```
kubectl create -f ${CLUSTER_NAME}.yaml
```

- ⓘ** If you used the `--output-directory` flag in your `dkp create .. --dry-run` step above, create the cluster from the objects by specifying the directory in the command:

```
kubectl create -f <existing-directory>
```

5. Wait for the cluster control-plane to be ready:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

6. Run the DKP describe command to monitor the current status of the cluster:

```
dkp describe cluster -c ${CLUSTER_NAME}
```

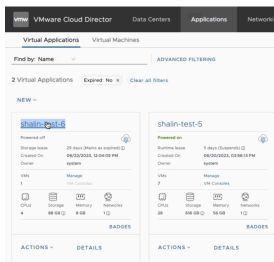
```
NAME                                                                 READY
SEVERITY REASON SINCE MESSAGE
Cluster/vcd-example                                                                 True
52s
├─ClusterInfrastructure - vcdCluster/vcd-example
├─ControlPlane - KubeadmControlPlane/vcd-example-control-plane                                                                 True
52s
│ └─Machine/vcd-example-control-plane-6fbzn                                                                 True
2m32s
│ │ └─MachineInfrastructure - vcdMachine/vcd-example-control-plane-62g6s
│ └─Machine/vcd-example-control-plane-jf6s2                                                                 True
7m36s
```

```

| | └─ MachineInfrastructure - vcdMachine/vcd-example-control-plane-bsr2z
| └─ Machine/vcd-example-control-plane-mnbfs True
54s
| └─ MachineInfrastructure - vcdMachine/vcd-example-control-plane-s8xsx
└─ Workers
  └─ MachineDeployment/vcd-example-md-0 True
78s
  └─ Machine/vcd-example-md-0-68b86fddb8-8glsw True
2m49s
  | └─ MachineInfrastructure - vcdMachine/vcd-example-md-0-zls8d
  └─ Machine/vcd-example-md-0-68b86fddb8-bvbm7 True
2m48s
  | └─ MachineInfrastructure - vcdMachine/vcd-example-md-0-5zcvc
  └─ Machine/vcd-example-md-0-68b86fddb8-k9499 True
2m49s
  | └─ MachineInfrastructure - vcdMachine/vcd-example-md-0-k8h5p
  └─ Machine/vcd-example-md-0-68b86fddb8-l6vfb True
2m49s
  └─ MachineInfrastructure - vcdMachine/vcd-example-md-0-9h5vn
    
```

NOTE: In the Cloud Director interface, you can find this new cluster under the **Applications** tab - **Virtual Applications** (vApp). One vApp per cluster will be created into CAPVCD with virtual machine information as well as worker node and control plane node information.

EXAMPLE:



7.17.3.10.4 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20CloudDirector%20Explore%20the%20Cluster&linkCreation=true&fromPagelId=273681351>¹⁰⁷⁹

7.17.3.11 Cloud Director Explore the Cluster

Beta

¹⁰⁷⁹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273681351&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Explore+the+Cluster>

This guide explains how to use the command line to interact with your newly deployed Kubernetes cluster. Before you start, make sure you have created a workload cluster, as described in [Create a New VCD Cluster](#) (see page 1345).

7.17.3.11.1 Explore the new Kubernetes cluster

1. Get a kubeconfig file for the workload cluster:

When the workload cluster is created, the cluster lifecycle services generate a kubeconfig file for the workload cluster, and write it to a *Secret*. The kubeconfig file is scoped to the cluster administrator.

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

2. Verify the API server is up :

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

NOTE: It may take a few minutes for the Status to move to Ready while the Pod network is deployed. The Nodes' Status should change to Ready soon after the calico-node DaemonSet Pods are Ready. You should receive the following output:

NAME	STATUS	ROLES	AGE	
VERSION				
vcd-example-control-plane-9z77w	Ready	control-plane,master	4m44s	
v1.26.6				
vcd-example-control-plane-rtj9h	Ready	control-plane,master	104s	
v1.26.6				
vcd-example-control-plane-zbf9w	Ready	control-plane,master	3m23s	
v1.26.6				
vcd-example-md-0-88c46	Ready	<none>	3m28s	v1.26
.6				
vcd-example-md-0-fp8s7	Ready	<none>	3m28s	v1.26
.6				
vcd-example-md-0-qvnx7	Ready	<none>	3m28s	v1.26
.6				
vcd-example-md-0-wjdrq	Ready	<none>	3m27s	v1.26
.6				

3. List the Pods with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get pods -A
```

4. The output has been shortened, but will resemble the following:

NAMESPACE	READY	STATUS	RESTARTS	NAME	AGE
calico-system	1/1	Running	0	calico-kube-controllers-577c696df9-v2nzv	5m23s
calico-system	1/1	Running	0	calico-node-4x5rk	4m22s
calico-system	1/1	Running	0	calico-node-cxsgc	4m23s
calico-system	1/1	Running	0	calico-node-dvlnm	4m23s
calico-system	1/1	Running	0	calico-node-h6nlt	4m23s
calico-system	1/1	Running	0	calico-node-jmkwq	5m23s
calico-system	1/1	Running	0	calico-node-tnf54	4m18s
calico-system	1/1	Running	0	calico-node-v6bwq	2m39s
calico-system	1/1	Running	0	calico-typha-6d8c94bfd-dkfvq	5m23s
calico-system	1/1	Running	0	calico-typha-6d8c94bfd-fdfn2	3m43s
calico-system	1/1	Running	0	calico-typha-6d8c94bfd-kjgzj	3m43s
capa-system	1/1	Running	0	capa-controller-manager-6468bc488-w7nj9	67s
capg-system	1/1	Running	0	capg-controller-manager-5fb47f869b-6jgms	53s
capi-kubeadm-bootstrap-system	1/1	Running	0	capi-kubeadm-bootstrap-controller-manager-65ffc94457-7cjd	74s
capi-kubeadm-control-plane-system	1/1	Running	0	capi-kubeadm-control-plane-controller-manager-bc7b688d4-vv8wg	72s
capi-system	1/1	Running	0	capi-controller-manager-dbfc7b49-dzvw8	77s
capp-system	1/1	Running	0	capp-controller-manager-8444d67568-rmms2	59s
capv-system	1/1	Running	0	capv-controller-manager-58b8ccf868-rbscn	56s
capz-system	1/1	Running	0	capz-controller-manager-6467f986d8-dnvj4	62s
capvcd-system	1/1	Running	0	capvcd-controller-manager-dbfc7b49-dzvw8	77s
cert-manager	1/1	Running	0	cert-manager-6888d6b69b-7b7m9	91s
kube-system	1/1	Running	0	cluster-autoscaler-7f695dc48f-v5kvh	5m40s
kube-system	1/1	Running	0	coredns-64897985d-hbkqd	5m38s
kube-system	1/1	Running	0	coredns-64897985d-m8g5j	5m38s

```

kube-system          etcd-vcd-example-control-plane-9z77w
1/1 Running 0 5m32s
kube-system          etcd-vcd-example-control-plane-rtj9h
1/1 Running 0 2m37s
kube-system          etcd-vcd-example-control-plane-zbf9w
1/1 Running 0 4m17s
kube-system          kube-apiserver-vcd-example-control-
plane-9z77w          1/1 Running 0 5m32s
kube-system          kube-controller-manager-vcd-example-
control-plane-9z77w 1/1 Running 0 5m33s
kube-system          kube-proxy-gdkn5
1/1 Running 0 4m23s
kube-system          kube-scheduler-vcd-example-control-plane-
zbf9w              1/1 Running 0 4m17s
node-feature-discovery node-feature-discovery-master-7d5985467-
lh7dc              1/1 Running 0 5m40s
tigera-operator      tigera-operator-5f9bdc5c59-j9tnr
1/1 Running 0 5m38s

```

7.17.3.11.2 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Cloud%20Director%20Make%20the%20New%20%20Cluster%20Self-Managed&linkCreation=true&fromPageId=273715108>¹⁰⁸⁰

7.17.3.12 Cloud Director Make the New Cluster Self-Managed

Beta

DKP deploys all cluster lifecycle services to a bootstrap cluster, which then deploys a workload cluster. When the workload cluster is ready, move the cluster lifecycle services to the workload cluster, which makes the workload cluster self-managed. This section describes how to make a workload cluster self-managed.

Before starting, ensure you create a workload cluster as described in [Create a New VCD Cluster](#) (see page 570) and that you as the Managed Service Provider(MSP) have a [Bastion Host](#) (see page 1288) configured so that the CAPVCD drivers can reach it.

This page contains instructions on how to make your cluster self-managed. This is necessary if there is only one cluster in your environment, or if this cluster should become the Management cluster in a multi-cluster environment.



If you already have a self-managed or Management cluster in your environment, skip this page.

¹⁰⁸⁰ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273715108&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Make+the+New+Cluster+Self-Managed>

7.17.3.12.1 Make the New Kubernetes Cluster Manage Itself

Follow these steps to create a management cluster in Enterprise or multiple stand alone Essential clusters:

1. Deploy cluster lifecycle services on the workload cluster:

```
dkp create capi-components --kubeconfig ${CLUSTER_NAME}.conf
```

```
✓ Initializing new CAPI components
```

2. Move the Cluster API objects from the bootstrap to the workload cluster:

The cluster lifecycle services on the workload cluster are ready, but the workload cluster configuration is on the bootstrap cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the bootstrap to the workload cluster. This process is also called a [Pivot](#)¹⁰⁸¹.

```
dkp move capi-resources --to-kubeconfig ${CLUSTER_NAME}.conf
```

```
✓ Moving cluster resources
```

You can now view resources in the moved cluster by using the `--kubeconfig` flag with `kubectl`. For example: `kubectl --kubeconfig=vcd-example.conf get nodes`

i NOTE: To ensure only one set of cluster lifecycle services manages the workload cluster, DKP first pauses reconciliation of the objects on the bootstrap cluster, then creates the objects on the workload cluster. As DKP copies the objects, the cluster lifecycle services on the workload cluster reconcile the objects. The workload cluster becomes self-managed after DKP creates all the objects. If it fails, the `move` command can be safely retried.

3. Wait for the cluster control-plane to be ready:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --timeout=20m
```

¹⁰⁸¹ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```
cluster.cluster.x-k8s.io/vcd-example condition met
```

- Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

i NOTE: After moving the cluster lifecycle services to the workload cluster, remember to use DKP with the workload cluster kubeconfig.

```
dkp describe cluster --kubeconfig ${CLUSTER_NAME}.conf -c ${CLUSTER_NAME}
```

```
NAME                                                                 READY
SEVERITY REASON SINCE MESSAGE
Cluster/vcd-example                                                                 True
14s
├─ClusterInfrastructure - vcdCluster/vcd-example
├─ControlPlane - KubeadmControlPlane/vcd-example-control-plane
14s
│ └─Machine/vcd-example-control-plane-6fbzn
17s
│ │ └─MachineInfrastructure - vcdMachine/vcd-example-control-plane-62g6s
│ └─Machine/vcd-example-control-plane-jf6s2
17s
│ │ └─MachineInfrastructure - vcdMachine/vcd-example-control-plane-bsr2z
│ └─Machine/vcd-example-control-plane-mnbfs
17s
│ └─MachineInfrastructure - vcdMachine/vcd-example-control-plane-s8xsx
└─Workers
  └─MachineDeployment/vcd-example-md-0
17s
    └─Machine/vcd-example-md-0-68b86fddb8-8glsw
17s
      └─MachineInfrastructure - vcdMachine/vcd-example-md-0-zls8d
        └─Machine/vcd-example-md-0-68b86fddb8-bvbm7
17s
          └─MachineInfrastructure - vcdMachine/vcd-example-md-0-5zcvc
            └─Machine/vcd-example-md-0-68b86fddb8-k9499
17s
              └─MachineInfrastructure - vcdMachine/vcd-example-md-0-k8h5p
                └─Machine/vcd-example-md-0-68b86fddb8-l6vfb
17s
                  └─MachineInfrastructure - vcdMachine/vcd-example-md-0-9h5vn
```

- Remove the bootstrap cluster, as the workload cluster is now self-managed:

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
✓ Deleting bootstrap cluster
```

7.17.3.12.2 Known Limitations

DKP only supports moving all namespaces in the cluster; DKP does not support migration of individual namespaces.

7.17.3.12.3 Next Step:

[Cloud Director Install Kommander](#) (see page 1352)

7.17.3.13 Cloud Director Install Kommander

Beta

The Kommander component of DKP can be configured differently depending on your environment type and other desired customizations. This installation allows access to the DKP Dashboard for cluster management and insights.

7.17.3.13.1 Installing Kommander by Environment

This section contains instructions for installing the Kommander component on *air-gapped*, *non-air-gapped*, or *pre-provisioned* environments. It also contains instructions on how to enable *DKP catalog apps* or *DKP Insights*, if desired.

7.17.3.13.2 Kommander Customizations

This section contains instructions to enable a custom configuration of a Kommander component such as custom domains or certificates, an HTTP Proxy, an external load balancer, etc.

7.17.3.13.3 Log in to the UI with Kommander:

After the Kommander component is installed according to the environment you selected, log in to the UI.

7.17.3.13.4 By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{{ "\n"}}Password: {{.data.password|
base64decode}}{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://
{{with index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/
dkp/kommander/dashboard{{ "\n"}}'
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password:
kqZ31lMBSClCbJUKVwLJMQL2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

7.17.3.13.5 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see page 585) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

7.17.3.13.6 Next Step:

[Day 2 - Cluster Operations Management](#) (see page 585)

7.17.3.14 Cloud Director Manage Node Pools

Beta

Node pools are part of a cluster and managed as a group, and you can use a node pool to manage a group of machines using the same common properties. When Konvoy creates a new default cluster, there is one node pool for the worker nodes and all nodes in that new node pool have the same configuration. You can create additional node pools for more specialized hardware or configuration. For example, if you want to tune your memory usage on a cluster where you need maximum memory for some machines and minimal memory on other machines, you would create a new node pool with those specific resource needs.

Konvoy implements node pools using Cluster API [MachineDeployments](#)¹⁰⁸².

7.17.3.14.1 Cloud Director Create Node Pools

Beta

Creating a node pool is useful when you need to run workloads that require machines with specific resources, such as additional memory, or specialized network or storage hardware.

7.17.3.14.1.1 Prerequisites

Before you begin, make sure you have created a [New VCD Cluster](#). (see page 1345)

Prepare the Environment

Follow these steps:

1. Set the environment variable to the name you assigned this cluster.

```
export CLUSTER_NAME=vcd-example
```

2. If your workload cluster is self-managed, as described in [Make the New Cluster Self-Managed](#) (see page 1352), configure `kubectl` to use the kubeconfig for the cluster.

```
export KUBECONFIG=${CLUSTER_NAME}.conf
```

¹⁰⁸² <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

3. Define your node pool name.

```
export NODEPOOL_NAME=example
```

Create a VCD Node Pool

Availability zones (AZs) are isolated locations within data center regions from which public cloud services originate and operate. Because all the nodes in a node pool are deployed in a single Availability Zone, you may wish to create additional node pools to ensure your cluster has nodes deployed in multiple Availability Zones.

Create a new AWS node pool with 3 replicas using this command:

Set the `--zone` flag to a [zone](#)¹⁰⁸³ in the same region as your cluster.

```
dkp create nodepool vcd ${NODEPOOL_NAME} \
  --cluster-name=${CLUSTER_NAME} \
  --image $IMAGE_NAME \
  --zone us-west1-b \
  --replicas=3
```

```
machinedeployment.cluster.x-k8s.io/example created
.. Creating default/example nodepool resources
vcdmachinetemplate.infrastructure.cluster.x-k8s.io/example created
kubeadmconfigtemplate.bootstrap.cluster.x-k8s.io/example created
✓ Creating default/example nodepool resources
```

This example uses default values for brevity. Use flags to define custom instance types, and other properties.

Advanced users can use a combination of the `--dry-run` and `--output=yaml` or `--output-directory=<your-target-directory>/` flags to get a complete set of node pool objects to modify locally or store in version control.

7.17.3.14.2 Cloud Director List Node Pools

Beta

Use this command to list the node pools of a given cluster. This returns specific properties of each node pool so that you can see the name of the MachineDeployments.

To list all node pools for a managed cluster, run:

¹⁰⁸³ <https://cloud.google.com/compute/docs/regions-zones>

```
dkp get nodepool --cluster-name=${CLUSTER_NAME}
```

The expected output is similar to the following example, indicating the desired size of the node pool, the number of replicas ready in the node pool, and the Kubernetes version those nodes are running:

NODEPOOL VERSION	DESIRED	READY	KUBERNETES
example	3	3	v1.26.6
vcd-example-2-md-0	4	4	v1.26.6

7.17.3.14.3 Cloud Director Scale Node Pools

Beta

While you can run [Cluster Autoscaler](#)¹⁰⁸⁴, you can also manually scale your node pools up or down when you need more finite control over your environment. For example, if you require 10 machines to run a process, you can manually set the scaling to run those 10 machines only. However, if also using the Cluster Autoscaler, you must stay within your minimum and maximum bounds.

7.17.3.14.3.1 Scaling Up Node Pools

To scale up a node pool in a cluster, run:

```
dkp scale nodepool ${NODEPOOL_NAME} --replicas=5 --cluster-name=${CLUSTER_NAME}
```

Your output should be similar to this example, indicating the scaling is in progress:

```
✓ Scaling node pool example to 5 replicas
```

After a few minutes, you can list the node pools to:

```
dkp get nodepool --cluster-name=${CLUSTER_NAME}
```

Your output should be similar to this example, with the number of DESIRED and READY replicas increased to 5:

NODEPOOL KUBERNETES VERSION	DESIRED	READY	
example	5	5	v1.26.6
vcd-example-md-0	4	4	v1.26.6

¹⁰⁸⁴ <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

7.17.3.14.3.2 Scaling Down Node Pools

To scale down a node pool, run:

```
dkp scale nodepool ${NODEPOOL_NAME} --replicas=4 --cluster-name=${CLUSTER_NAME}
```

✓ Scaling node pool example to 4 replicas

After a few minutes, you can list the node pools using this command:

```
dkp get nodepool --cluster-name=${CLUSTER_NAME}
```

Your output should be similar to this example, with the number of DESIRED and READY replicas decreased to 4:

NODEPOOL	DESIRED	READY	
KUBERNETES VERSION			
example	4	4	v1.26.6
vcd-example-md-0	4	4	v1.26.6

In a default cluster, the nodes to delete are selected at random. This behavior is controlled by [CAPI's delete policy](#)¹⁰⁸⁵. However, when using the DKP CLI to scale down a node pool, it is also possible to specify the Kubernetes Nodes you want to delete.

To do this, set the flag `--nodes-to-delete` with a list of nodes as below. This adds an annotation `cluster.x-k8s.io/delete-machine=yes` to the matching Machine object that contains `status.NodeRef` with the node names from `--nodes-to-delete`.

```
dkp scale nodepool ${NODEPOOL_NAME} --replicas=3 --cluster-name=${CLUSTER_NAME}
```

✓ Scaling node pool example to 3 replicas

7.17.3.14.3.3 Scaling Node Pools When Using Cluster Autoscaler

If you configured [the cluster autoscaler](#)¹⁰⁸⁶ for the `demo-cluster-md-0` node pool, the value of `--replicas` must be within the minimum and maximum bounds.

For example, assuming you have these annotations:

¹⁰⁸⁵ https://github.com/kubernetes-sigs/cluster-api/blob/v0.4.0/api/v1alpha4/machineset_types.go#L85-L105

¹⁰⁸⁶ <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>


```
kubectl annotate machinedeployment ${NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-
autoscaler-node-group-min-size=2
kubectl annotate machinedeployment ${NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-
autoscaler-node-group-max-size=6
```

Try to scale the node pool to 7 replicas with the command:

```
dkp scale nodepool ${NODEPOOL_NAME} --replicas=7 --cluster-name=${CLUSTER_NAME}
```

Which results in an error similar to:

```
✗ Scaling node pool example to 7 replicas
failed to scale nodepool: scaling MachineDeployment is forbidden: desired replicas 7
is greater than the configured max size annotation cluster.x-k8s.io/cluster-api-
autoscaler-node-group-max-size: 6
```

Similarly, scaling down to a number of replicas less than the configured `min-size` also returns an error.

7.17.3.14.4 Cloud Director Delete Node Pools

Beta

Deleting a node pool deletes the Kubernetes nodes and the underlying infrastructure. All nodes will be drained prior to deletion and the pods running on those nodes will be rescheduled.

To delete a node pool from a managed cluster, run:

```
dkp delete nodepool ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME}
```

Here, `example` is the node pool to be deleted.

The expected output will be similar to the following example, indicating the node pool is being deleted:

```
✓ Deleting default/example nodepool resources
```

Deleting an invalid node pool results in output similar to this example:

```
dkp delete nodepool ${CLUSTER_NAME}-md-invalid --cluster-name=${CLUSTER_NAME}
```

```
MachineDeployments or MachinePools.infrastructure.cluster.x-k8s.io "no
MachineDeployments or MachinePools found for cluster vcd-example" not found
```

7.17.3.14.5 Cloud Director Cluster Autoscaler

Beta

[Cluster Autoscaler](#)¹⁰⁸⁷ provides the ability to automatically scale-up or scale-down the number of worker nodes in a cluster, based on the number of pending pods to be scheduled. Running the Cluster Autoscaler is optional.

Unlike [Horizontal-Pod Autoscaler](#)¹⁰⁸⁸, Cluster Autoscaler does not depend on any Metrics server and does not need Prometheus or any other metrics source.

The Cluster Autoscaler looks at the following annotations on a MachineDeployment to determine its scale-up and scale-down ranges:

```
cluster.x-k8s.io/cluster-api-autoscaler-node-group-min-size
cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size
```

The full list of command line arguments to the Cluster Autoscaler controller is [on the Kubernetes public GitHub repository](#)¹⁰⁸⁹.

For more information about how Cluster Autoscaler works, see these documents:

- [What is Cluster Autoscaler](#)¹⁰⁹⁰
- [How does scale-up work](#)¹⁰⁹¹
- [How does scale-down work](#)¹⁰⁹²
- [CAPI Provider for Cluster Autoscaler](#)¹⁰⁹³

7.17.3.14.5.1 Cluster Autoscaler Prerequisites

Before you begin, you must have:

- A [Bootstrap Cluster Lifecycle](#)¹⁰⁹⁴.
- A [New Kubernetes Cluster](#)¹⁰⁹⁵.

¹⁰⁸⁷ <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

¹⁰⁸⁸ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-fast-is-hpa-when-combined-with-ca>

¹⁰⁸⁹ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-are-the-parameters-to-ca>

¹⁰⁹⁰ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-is-cluster-autoscaler>

¹⁰⁹¹ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-up-work>

¹⁰⁹² <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-down-work>

¹⁰⁹³ <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

¹⁰⁹⁴ <https://d2iq.atlassian.net/wiki/pages/createtime.action?fromPageId=273714206&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Bootstrap+the+Cluster>

¹⁰⁹⁵ <https://d2iq.atlassian.net/wiki/pages/createtime.action?fromPageId=273714206&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Create+a+New+Cluster>

¹⁰⁹⁵ <https://d2iq.atlassian.net/wiki/pages/createtime.action?fromPageId=273714206&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Create+a+New+Cluster>

¹⁰⁹⁵ <https://d2iq.atlassian.net/wiki/pages/createtime.action?fromPageId=273714206&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Create+a+New+Cluster>

- A [Self-Managed Cluster](#)¹⁰⁹⁶.

7.17.3.14.5.2 Run Cluster Autoscaler

The Cluster Autoscaler controller runs on the workload cluster. Upon creation of the workload cluster, this controller does not have all the objects required to function correctly until after a `dkp move` is issued from the bootstrap cluster.

Run the following steps to enable Cluster Autoscaler:

1. Ensure the Cluster Autoscaler controller is up and running (no restarts and no errors in the logs)

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf logs deployments/cluster-autoscaler
cluster-autoscaler -n kube-system -f
```

2. Enable Cluster Autoscaler by setting the min & max ranges

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment $
{NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-autoscaler-node-group-min-size=2
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment $
{NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size=6
```

3. The Cluster Autoscaler logs will show that the worker nodes are associated with node-groups and that pending pods are being watched.
4. To demonstrate that it is working properly, create a large deployment which will trigger pending pods (For this example we used Cloud Director n2-standard-8 worker nodes. If you have larger worker-nodes, you should scale up the number of replicas accordingly).

```
cat <<EOF | kubectl --kubeconfig=${CLUSTER_NAME}.conf apply -f -
apiVersion: apps/v1
kind: Deployment
metadata:
  name: busybox-deployment
  labels:
    app: busybox
spec:
  replicas: 600
  selector:
    matchLabels:
      app: busybox
  template:
    metadata:
```

¹⁰⁹⁶[https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273714206&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Make+the+New+Cluster+Self-Managed)

[fromPageId=273714206&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Make+the+New+Cluster+Self-Managed](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273714206&linkCreation=true&spaceKey=DENT&title=%282.6%29+Cloud+Director+Make+the+New+Cluster+Self-Managed)

```

labels:
  app: busybox
spec:
  containers:
  - name: busybox
    image: busybox:latest
    command:
      - sleep
      - "3600"
    imagePullPolicy: IfNotPresent
  restartPolicy: Always
EOF

```

5. Cluster Autoscaler will scale up the number of Worker Nodes until there are no pending pods.
6. Scale down the number of replicas for `busybox-deployment`.

```

kubectl --kubeconfig ${CLUSTER_NAME}.conf scale --replicas=30 deployment/
busybox-deployment

```

7. Cluster Autoscaler starts to scale down the number of Worker Nodes after the default timeout of 10 minutes.

7.17.3.15 Cloud Director Delete a Cluster

Beta

7.17.3.15.1 Prepare to Delete a Self-managed Workload Cluster

ⓘ A self-managed workload cluster cannot delete itself. If your workload cluster is self-managed, you must create a bootstrap cluster and move the cluster lifecycle services to the bootstrap cluster before deleting the workload cluster.

If you did not make your workload cluster self-managed, as described in [Make the New Cloud Director Cluster Self-Managed](#) (see page 1352), proceed to the **Delete the Workload Cluster** section below.

1. Create a bootstrap cluster:

The bootstrap cluster will host the Cluster API controllers that reconcile the cluster objects marked for deletion:

ⓘ NOTE: To avoid using the wrong `kubeconfig`, the following steps use explicit `kubeconfig` paths and contexts.

```
dkp create bootstrap --vcd-bootstrap-credentials=true --kubeconfig $HOME/.kube/
config
```

- ✓ Creating a bootstrap cluster
- ✓ Initializing **new** CAPI components

2. Move the Cluster API objects from the workload to the bootstrap cluster: The cluster lifecycle services on the bootstrap cluster are ready, but the workload cluster configuration is on the workload cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the workload to the bootstrap cluster. This process is also called a [Pivot](#)¹⁰⁹⁷.

```
dkp move capi-resources \
  --from-kubeconfig ${CLUSTER_NAME}.conf \
  --to-kubeconfig $HOME/.kube/config
```

3. Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

```
dkp describe cluster --kubeconfig $HOME/.kube/config -c ${CLUSTER_NAME}
```

```
NAME                                                                 READY
SEVERITY REASON SINCE MESSAGE
Cluster/vcd-example                                                                 True
34s
├─ClusterInfrastructure - vcdCluster/vcd-example
├─ControlPlane - KubeadmControlPlane/vcd-example-control-plane           True
34s
| └─Machine/vcd-example-control-plane-6fbzn                               True
37s
| | └─MachineInfrastructure - vcdMachine/vcd-example-control-plane-62g6s
| └─Machine/vcd-example-control-plane-jf6s2                               True
37s
| | └─MachineInfrastructure - vcdMachine/vcd-example-control-plane-bsr2z
| └─Machine/vcd-example-control-plane-mnbfs                               True
37s
| └─MachineInfrastructure - vcdMachine/vcd-example-control-plane-s8xsx
└─Workers
  └─MachineDeployment/vcd-example-md-0                                     True
37s
```

¹⁰⁹⁷ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```

    |─Machine/vcd-example-md-0-68b86fddb8-8glsw                                True
37s |   └─MachineInfrastructure - vcdMachine/vcd-example-md-0-zls8d
    |─Machine/vcd-example-md-0-68b86fddb8-bvbm7                                True
37s |   └─MachineInfrastructure - vcdMachine/vcd-example-md-0-5zcvc
    |─Machine/vcd-example-md-0-68b86fddb8-k9499                                True
37s |   └─MachineInfrastructure - vcdMachine/vcd-example-md-0-k8h5p
    └─Machine/vcd-example-md-0-68b86fddb8-l6vfb                                True
37s   └─MachineInfrastructure - vcdMachine/vcd-example-md-0-9h5vn

```



- After moving the cluster lifecycle services to the workload cluster, remember to use `dkp` with the workload cluster `kubeconfig` to point at the correct cluster.
- Persistent Volumes (PVs) are not deleted automatically by design in order to preserve your data. However, they take up storage space if not deleted. You must delete PVs manually. Information for backup of a cluster and PVs is on the page in documentation called [Back up your Cluster's Applications and Persistent Volumes](#) (see page 835).

7.17.3.15.2 Delete the Workload Cluster

1. Use `dkp` with the bootstrap cluster to delete the workload cluster. Delete the Kubernetes cluster and wait a few minutes:

Before deleting the cluster, `dkp` deletes all Services of type LoadBalancer on the cluster. To skip this step, use the flag `--delete-kubernetes-resources=false`.

```

dkp delete cluster --kubeconfig $HOME/.kube/config --cluster-name $
{CLUSTER_NAME}

```

```

✓ Deleting Services with type LoadBalancer for Cluster default/vcd-example
✓ Deleting ClusterResourceSets for Cluster default/vcd-example
✓ Deleting cluster resources
✓ Waiting for cluster to be fully deleted
Deleted default/vcd-example cluster

```

After the workload cluster is deleted, delete the bootstrap cluster.

7.17.3.15.3 Delete the Bootstrap Cluster

1. Delete the bootstrap cluster using `dkp delete`:

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

✓ Deleting bootstrap cluster

7.17.3.15.4 Delete a Managed Cluster

In the CLI, using the `kubeconfig` for the **Managed** cluster that you want to delete. In order to properly delete a **Managed** cluster, you must first delete the PVCs using these steps:

1. Export your workspace namespace:

```
export WORKSPACE_NAMESPACE=<your-workspace-namespace>
```

2. Get all of the PVCs that you need to delete from the Applications deployed by Kommander:

```
kubectl get pvc --namespace ${WORKSPACE_NAMESPACE}
```

3. Delete those PVCs and the Managed cluster deletion process will continue:

```
kubectl delete pvc --namespace ${WORKSPACE_NAMESPACE} <pvc-name-1> <pvc-name-2>
```



User can delete the worker node pool(s) before deleting the cluster itself and Cluster API will drain the Nodes. See also: [Cloud Director Delete Node Pools \(see page 1361\)](#)

7.17.3.15.5 Known Limitations

The DKP version used to create the workload cluster must match the DKP version used to delete the workload cluster.

7.18 GCP Infrastructure

The following procedure describes creating a DKP cluster on GCP.

- [GCP Prerequisites](#) (see page 1365)
- [GCP Konvoy Image Builder](#) (see page 1368)
- [Bootstrap GCP](#) (see page 1371)
- [Create a New GCP Cluster](#) (see page 1373)
- [Explore the GCP Cluster](#) (see page 1376)
- [Make the New GCP Cluster Self-Managed](#) (see page 1380)
- [Manage GCP Node Pools](#) (see page 1382)
- [Delete a GCP Cluster](#) (see page 1390)

7.18.1 GCP Prerequisites

7.18.1.1 Prerequisites

Before beginning a DKP installation, verify that you have:

- An x86_64-based Linux or macOS machine with a supported version of the operating system.
- [Download](#) (see page 78) the `dkp` binary for Linux, or macOS. To check which version of DKP you installed for compatibility reasons, run the `dkp version -h` command ([dkp version](#) (see page 1749)).
- A Container engine/runtime installed is required to bootstrap:
 - Version [Docker®](#)¹⁰⁹⁸ container engine version 18.09.2 or higher installed for Linux or MacOS - On macOS, Docker runs in a virtual machine which needs configured with at least 8 GB of memory.
 - Version 4.0 of [Podman](#)¹⁰⁹⁹ or higher for Linux. Host requirements found here: [Host Requirements](#)¹¹⁰⁰
- [kubecti](#)¹¹⁰¹ for interacting with the running cluster.
- Install the GCP `gcloud` CLI by following the <https://cloud.google.com/sdk/docs/install>

7.18.1.2 Control plane nodes

You must have at least three control plane nodes. Each control plane node should have at least:

- 4 cores
- 16 GiB memory

¹⁰⁹⁸ <https://docs.docker.com/get-docker/>

¹⁰⁹⁹ <https://podman.io/getting-started/installation>

¹¹⁰⁰ <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>

¹¹⁰¹ <https://kubernetes.io/docs/tasks/tools/#kubectl>

- Approximately 80 GiB of free space for the volume used for `/var/lib/kubelet` and `/var/lib/containerd`.
- Disk usage must be below 85% on the root volume.

DKP on GCP defaults to deploying an `n2-standard-4` instance with an 80GiB root volume for control plane nodes, which meets the above requirements.

7.18.1.3 Worker nodes

You must have at least four worker nodes. The specific number of worker nodes required for your environment can vary depending on the cluster workload and size of the nodes. Each worker node should have at least:

- 8 cores
- 32 GiB memory
- Around 80 GiB of free space for the volume used for `/var/lib/kubelet` and `/var/lib/containerd`.
- Disk usage must be below 85% on the root volume.


DKP on GCP defaults to deploying a `n2-standard-8` instance with an 80GiB root volume for worker nodes, which meets the above requirements.

7.18.1.4 GCP Prerequisite Roles

If you are creating your image on either a non-GCP instance or one that does not have the required roles, you must either:

- Create a [GCP service account](#)¹¹⁰².
- If you have already created a service account, retrieve the credentials for an existing service account.
- Export the static credentials that will be used to create the cluster:

```
export GCP_B64ENCODED_CREDENTIALS=$(base64 < "$
{GOOGLE_APPLICATION_CREDENTIALS}" | tr -d '\n')
```

 Make sure to rotate static credentials for increased security.

- If you are creating the cluster on a non-GCP instance or one that does not have the required `Editor` role:
 - (option 1) Create a GCP Service Account using the following `gcloud` commands:

¹¹⁰² <https://cloud.google.com/iam/docs/service-account-overview>

```
export GCP_PROJECT=<your GCP project ID>
export GCP_SERVICE_ACCOUNT_USER=<some new service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud iam service-accounts create "$GCP_SERVICE_ACCOUNT_USER" --
project=$GCP_PROJECT
gcloud projects add-iam-policy-binding $GCP_PROJECT --
member="serviceAccount:
$GCP_SERVICE_ACCOUNT_USER@$GCP_PROJECT.iam.gserviceaccount.com" --
role=roles/editor
gcloud iam service-accounts keys create $GOOGLE_APPLICATION_CREDENTIALS
--iam-
account="$GCP_SERVICE_ACCOUNT_USER@$GCP_PROJECT.iam.gserviceaccount.com"
```

- (option 2) Retrieve the credentials for an existing service account using the following `gcloud` commands:

```
export GCP_PROJECT=<your GCP project ID>
export GCP_SERVICE_ACCOUNT_USER=<existing service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud iam service-accounts keys create $GOOGLE_APPLICATION_CREDENTIALS
--iam-
account="$GCP_SERVICE_ACCOUNT_USER@$GCP_PROJECT.iam.gserviceaccount.com"
```

- Export the static credentials that will be used to create the cluster:

```
export GCP_B64ENCODED_CREDENTIALS=$(base64 < "$
{GOOGLE_APPLICATION_CREDENTIALS}" | tr -d '\n')
```

To create a **GCP Service Account** with the `Editor` role, the user creating the GCP Service Account needs the `Editor`, `RoleAdministrator`, and `SecurityAdmin` roles. However, those pre-defined roles grant more permissions than the minimum set needed to create a DKP cluster.

ⓘ For **DKP** cluster creation, a **minimal** set of roles and permissions needed for the user creating the GCP Service Account is the `Editor` role **plus** the following additional permissions:

- `compute.disks.setIamPolicy`
- `compute.instances.setIamPolicy`
- `iam.roles.create`

- `iam.roles.delete`
- `iam.roles.update`
- `iam.serviceAccounts.setIamPolicy`
- `resourcemanager.projects.setIamPolicy`


For more information on GCP service accounts, see GCP's documentation:

- <https://cloud.google.com/iam/docs/creating-managing-service-accounts>
- <https://cloud.google.com/iam/docs/best-practices-service-accounts>

7.18.2 GCP Konvoy Image Builder

Konvoy Image Builder (KIB) is a complete solution for building Cluster API compliant images.

This procedure describes how to use the [Konvoy Image Builder](#) (see page 1461) (KIB) to create a [Cluster API](#)¹¹⁰³ compliant GCP image. GCP images contain configuration information and software to create a specific, pre-configured, operating environment. For example, you can create a GCP image of your current computer system settings and software. The GCP image can then be replicated and distributed, creating your computer system for other users. KIB uses [variable overrides](#) (see page 1513) to specify base image and container images to use in your new GCP image.

 Google Cloud Platform does not publish images. You must first build the image using Konvoy Image Builder. For more information regarding images and clusters, refer to the [GCP Infrastructure](#) (see page 1368) section of the documentation.


7.18.2.1 Prerequisites

Before you begin, you must:

- Download the [KIB](#) (see page 1463) bundle for your version of DKP.
- Check the [Supported Infrastructure Operating Systems](#) (see page 69)
- Check the [Supported Kubernetes Version](#) (see page 1552) for your Provider.
- Create a working Docker or other registry setup.
- On Debian-based Linux distributions, install a version of the [cri-tools](#)¹¹⁰⁴ package known to be compatible with both the Kubernetes and container runtime versions.
- Verify that your Google Cloud project does not have the Enable OS Login feature enabled. See below for more information:

¹¹⁰³ <https://cluster-api.sigs.k8s.io/>

¹¹⁰⁴ <https://github.com/kubernetes-sigs/cni-plugins>


 The Enable OS Login feature is sometimes enabled by default in GCP projects. If the OS login feature is enabled, KIB will not be able to `ssh` to the VM instances it creates and will not be able to successfully create an image. To check if it is enabled, use the commands on this page https://cloud.google.com/compute/docs/metadata/setting-custom-metadata#console_2 to inspect the metadata configured in in your project. If you find the the `enable-oslogin` flag set to TRUE, you must remove (or set it to FALSE) to successfully use KIB.

7.18.2.2 GCP Prerequisite Roles

If you are creating your image on either a non-GCP instance or one that does not have the [required roles](#)¹¹⁰⁵¹¹⁰⁶, you must either:

- Create a [GCP service account](#)¹¹⁰⁷.
- If you have already created a service account, retrieve the credentials for an existing service account.
- Export the static credentials that will be used to create the cluster:

```
export GCP_B64ENCODED_CREDENTIALS=$(base64 < "{GOOGLE_APPLICATION_CREDENTIALS}" | tr -d '\n')
```

 Make sure to rotate static credentials for increased security.

- If you are creating your image on either a non-GCP instance or one that does not have the [required roles](#) (see page 1368):
 - (option 1) Create a service account using the following `gcloud` commands:

```
export GCP_PROJECT=<your GCP project ID>
export GCP_SERVICE_ACCOUNT_USER=<some new service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud iam service-accounts create "${GCP_SERVICE_ACCOUNT_USER}" --
project=${GCP_PROJECT}
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
```

1105 <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/295272692>

1106 <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/295272692>

1107 <https://cloud.google.com/iam/docs/service-account-overview>

```
{GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
compute.instanceAdmin.v1
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
iam.serviceAccountUser
gcloud iam service-accounts keys create $
{GOOGLE_APPLICATION_CREDENTIALS} --iam-account="$
{GCP_SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com"
```

- (option 2) If you have already created a service account, retrieve the credentials for an existing service account using the following `gcloud` commands:

```
export GCP_PROJECT=<your GCP project ID>
export GCP_SERVICE_ACCOUNT_USER=<existing service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
compute.instanceAdmin.v1
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
iam.serviceAccountUser
gcloud iam service-accounts keys create $
{GOOGLE_APPLICATION_CREDENTIALS} --iam-account="$
{GCP_SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com"
```

- Export the static credentials that will be used to create the cluster:

```
export GCP_B64ENCODED_CREDENTIALS=$(base64 < "$
{GOOGLE_APPLICATION_CREDENTIALS}" | tr -d '\n')
```



Make sure to rotate static credentials for increased security.

7.18.2.2.1 Build the GCP image

1. Run the `konvoy-image` command to build and validate the image:

```
./konvoy-image build gcp --project-id ${GCP_PROJECT} --network ${NETWORK_NAME}
images/gcp/ubuntu-2004.yaml
```

2. KIB will run and print out the name of the created image, you will use this name when creating a Kubernetes cluster. See sample output below:

```
...
==> ubuntu-2004-focal-v20220419: Deleting instance...
ubuntu-2004-focal-v20220419: Instance has been deleted!
==> ubuntu-2004-focal-v20220419: Creating image...
==> ubuntu-2004-focal-v20220419: Deleting disk...
ubuntu-2004-focal-v20220419: Disk has been deleted!
==> ubuntu-2004-focal-v20220419: Running post-processor: manifest
Build 'ubuntu-2004-focal-v20220419' finished after 7 minutes 46 seconds.

==> Wait completed after 7 minutes 46 seconds

==> Builds finished. The artifacts of successful builds are:
--> ubuntu-2004-focal-v20220419: A disk image was created: konvoy-ubuntu-2004-1-23-7-1658523168
--> ubuntu-2004-focal-v20220419: A disk image was created: konvoy-ubuntu-2004-1-23-7-1658523168
```

3. To find a list of images you have created in your account, run the following command:

```
gcloud compute images list --no-standard-images
```



[Konvoy Image Builder](#) (see page 1461) has a more detailed section in the documentation if you need to refer there for compatible versions with DKP and other specific information.

7.18.3 Bootstrap GCP

To create Kubernetes clusters, DKP uses [Cluster API](#)¹¹⁰⁸ (CAPI) controllers. These controllers run on a Kubernetes cluster. To get started, you need a *bootstrap* cluster. By default, DKP creates a bootstrap cluster for you in a Docker container using the Kubernetes-in-Docker ([KIND](#)¹¹⁰⁹) tool.

7.18.3.1 Prerequisites

Before you begin, you must:

- Complete the steps in [Prerequisites](#) (see page 1368).
- Ensure the `dkp` binary can be found in your `$PATH`.

¹¹⁰⁸ <https://cluster-api.sigs.k8s.io/>

¹¹⁰⁹ <https://github.com/kubernetes-sigs/kind>

7.18.3.2 Bootstrap Cluster Lifecycle Services

1. If an HTTP proxy is required for the bootstrap cluster, set the local `http_proxy`, `https_proxy`, and `no_proxy` environment variables. They are copied into the bootstrap cluster.
2. Create a bootstrap cluster:

```
dkp create bootstrap --with-gcp-bootstrap-credentials=true
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

3. The output looks similar to:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

4. DKP creates a bootstrap cluster using [KIND](#)¹¹¹⁰ as a library. DKP then deploys the following [Cluster API](#)¹¹¹¹ providers on the cluster:
 - [Core Provider](#)¹¹¹²
 - [GCP Infrastructure Provider](#)¹¹¹³
 - [Kubeadm Bootstrap Provider](#)¹¹¹⁴
 - [Kubeadm ControlPlane Provider](#)¹¹¹⁵
5. DKP waits until the controller-manager and webhook deployments of these providers are ready. List these deployments using this command:

```
kubectl get --all-namespaces deployments -l=clusterctl.cluster.x-k8s.io
```

You will then receive the following output:

NAMESPACE		NAME	
READY	UP-TO-DATE	AVAILABLE	AGE

1110 <https://github.com/kubernetes-sigs/kind>

1111 <https://cluster-api.sigs.k8s.io/>

1112 <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

1113 <https://github.com/kubernetes-sigs/cluster-api-provider-gcp>

1114 <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

1115 <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>

```

capa-system          capa-controller-manager
1/1      1          1          1h
capg-system          capg-controller-manager
1/1      1          1          1h
capi-kubeadm-bootstrap-system  capi-kubeadm-bootstrap-controller-manager
1/1      1          1          1h
capi-kubeadm-control-plane-system  capi-kubeadm-control-plane-controller-
manager  1/1      1          1          1h
capi-system          capi-controller-manager
1/1      1          1          1h
capp-system          capp-controller-manager
1/1      1          1          1h
capv-system          capv-controller-manager
1/1      1          1          1h
capz-system          capz-controller-manager
1/1      1          1          1h
cert-manager         cert-manager
1/1      1          1          1h
cert-manager         cert-manager-cainjector
1/1      1          1          1h
cert-manager         cert-manager-webhook
1/1      1          1          1h

```

After completion, go to [creating a new GCP cluster](#) (see page 1373).

7.18.4 Create a New GCP Cluster


7.18.4.1 Name your cluster

1. Give your cluster a unique name suitable for your environment.


In GCP it is critical that the name is unique, as no two clusters in the same GCP account can have the same name.

2. Set the environment variable:

```
export CLUSTER_NAME=<gcp-example>
```


 The cluster name may only contain the following characters: `a-z`, `0-9`, `.`, and `-`. Cluster creation will fail if the name has capital letters. See [Kubernetes](#)¹¹¹⁶ for more naming information.


¹¹¹⁶ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

-  To increase [Docker Hub's rate limit](#)¹¹¹⁷ use your Docker Hub credentials when creating the cluster, by setting the following flag `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username=<username> --registry-mirror-password=<password>` on the `dkp create cluster` command.

7.18.4.2 Create a new GCP cluster

Availability zones (AZs) are isolated locations within data center regions from which public cloud services originate and operate. Because all the nodes in a node pool are deployed in a single Availability Zone, you may wish to create additional node pools to ensure your cluster has nodes deployed in multiple Availability Zones.

-  By default, the control-plane Nodes will be created in 3 different zones. However, the default worker Nodes will reside in a single zone. You may create additional node pools in other zones with the `dkp create nodepool` command. The default region for the availability zones is `us-west1`.

-  Google Cloud Platform does not publish images. You must first build the image using [Konvoy Image Builder](#) (see page 1461).

1. Create an image using [Konvoy Image Builder \(KIB\)](#) (see page 1461) and then export the image name:

```
export IMAGE_NAME=projects/${GCP_PROJECT}/global/images/<image_name_from_kib>
```

2. (Optional) You can modify Control Plane Audit logs settings using the information contained in the page [Configuring the Control Plane](#) (see page 1538).
3. (Optional) Determine what VPC Network to use. All GCP accounts come with a preconfigured VPC Network named `default`, which will be used if you do not specify a different network. To use a different VPC network for your cluster, create one by following these instructions for [Create and Manage VPC Networks](#)¹¹¹⁸. Then specify the `--network <new_vpc_network_name>` option

¹¹¹⁷ <https://docs.docker.com/docker-hub/download-rate-limit/>

¹¹¹⁸ <https://cloud.google.com/vpc/docs/create-modify-vpc-networks#create-auto-network>

on the create cluster command below. Follow the link for more information on [GCP Cloud Nat](#)¹¹¹⁹ and network flag.

4. Create a Kubernetes cluster. The following example shows a common configuration. See [dkp create cluster gcp](#) (see page 1678) reference for the full list of cluster creation options: Once the cluster creation process has finished, move onto [exploring your new cluster](#). (see page 1376)



- Ensure your [subnets](#) (see page 1049) do not overlap with your host subnet because they cannot be changed after cluster creation. If you need to change the kubernetes subnets, you must do this at cluster creation. The default subnets used in DKP are:

```
spec:
  clusterNetwork:
    pods:
      cidrBlocks:
        - 192.168.0.0/16
    services:
      cidrBlocks:
        - 10.96.0.0/12
```

```
dkp create cluster gcp \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-gcp-bootstrap-credentials=true \
--project=${GCP_PROJECT} \
--image=${IMAGE_NAME}
```



If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

5. Wait for the cluster control-plane to be ready:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

¹¹¹⁹<https://github.com/kubernetes-sigs/cluster-api-provider-gcp/blob/3406adaae0f8f65a615844e55d856d306eddacc1/docs/book/src/topics/prerequisites.md#cloud-nat>

6. After the objects are created on the API server, the Cluster API controllers reconcile them. They create infrastructure and machines. As they progress, they update the Status of each object. Konvoy provides a command to describe the current status of the cluster:

```
dkp describe cluster -c ${CLUSTER_NAME}
```

```

NAME                                                                 READY
SEVERITY REASON SINCE MESSAGE
Cluster/gcp-example                                                                 True
52s
├─ClusterInfrastructure - GCPCluster/gcp-example
├─ControlPlane - KubeadmControlPlane/gcp-example-control-plane           True
52s
| └─Machine/gcp-example-control-plane-6fbzn                               True
2m32s
| | └─MachineInfrastructure - GCPMachine/gcp-example-control-plane-62g6s
| └─Machine/gcp-example-control-plane-jf6s2                               True
7m36s
| | └─MachineInfrastructure - GCPMachine/gcp-example-control-plane-bsr2z
| └─Machine/gcp-example-control-plane-mnbfs                               True
54s
| └─MachineInfrastructure - GCPMachine/gcp-example-control-plane-s8xsx
└─Workers
  └─MachineDeployment/gcp-example-md-0                                     True
78s
  └─Machine/gcp-example-md-0-68b86fddb8-8glsw                             True
2m49s
    └─MachineInfrastructure - GCPMachine/gcp-example-md-0-zls8d
    └─Machine/gcp-example-md-0-68b86fddb8-bvbm7                             True
2m48s
    └─MachineInfrastructure - GCPMachine/gcp-example-md-0-5zcvc
    └─Machine/gcp-example-md-0-68b86fddb8-k9499                             True
2m49s
    └─MachineInfrastructure - GCPMachine/gcp-example-md-0-k8h5p
    └─Machine/gcp-example-md-0-68b86fddb8-l6vfb                             True
2m49s
    └─MachineInfrastructure - GCPMachine/gcp-example-md-0-9h5vn

```

7.18.5 Explore the GCP Cluster

This guide explains how to use the command line to interact with your newly deployed Kubernetes cluster.

Before you start, make sure you have created a workload cluster, as described in [Create a New GCP Cluster \(see page 1376\)](#).

7.18.5.1 Explore the new Kubernetes cluster

1. Get a kubeconfig file for the workload cluster:

When the workload cluster is created, the cluster lifecycle services generate a kubeconfig file for the workload cluster, and write it to a *Secret*. The kubeconfig file is scoped to the cluster administrator.

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

2. Verify the API server is up :

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

NOTE: It may take a few minutes for the Status to move to Ready while the Pod network is deployed. The Nodes' Status should change to Ready soon after the calico-node DaemonSet Pods are Ready. You should receive the following output:

NAME	STATUS	ROLES	AGE	
VERSION				
gcp-example-control-plane-9z77w	Ready	control-plane,master	4m44s	
v1.26.6				
gcp-example-control-plane-rtj9h	Ready	control-plane,master	104s	
v1.26.6				
gcp-example-control-plane-zbf9w	Ready	control-plane,master	3m23s	
v1.26.6				
gcp-example-md-0-88c46	Ready	<none>	3m28s	v1.26
.6				
gcp-example-md-0-fp8s7	Ready	<none>	3m28s	v1.26
.6				
gcp-example-md-0-qvnx7	Ready	<none>	3m28s	v1.26
.6				
gcp-example-md-0-wjdrq	Ready	<none>	3m27s	v1.26
.6				

3. List the Pods with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get pods -A
```

4. View the following output:

NAMESPACE	READY	STATUS	RESTARTS	NAME	AGE
calico-system				calico-kube-controllers-577c696df9-v2nzv	
1/1	Running	0		5m23s	

```

calico-system                                calico-node-4x5rk
1/1      Running  0      4m22s
calico-system                                calico-node-cxsgc
1/1      Running  0      4m23s
calico-system                                calico-node-dvlnm
1/1      Running  0      4m23s
calico-system                                calico-node-h6nlt
1/1      Running  0      4m23s
calico-system                                calico-node-jmkwq
1/1      Running  0      5m23s
calico-system                                calico-node-tnf54
1/1      Running  0      4m18s
calico-system                                calico-node-v6bwq
1/1      Running  0      2m39s
calico-system                                calico-typha-6d8c94bfd-dkfvq
1/1      Running  0      5m23s
calico-system                                calico-typha-6d8c94bfd-fdfn2
1/1      Running  0      3m43s
calico-system                                calico-typha-6d8c94bfd-kjgzj
1/1      Running  0      3m43s
capa-system                                capa-controller-manager-6468bc488-w7nj9
1/1      Running  0      67s
capg-system                                capg-controller-manager-5fb47f869b-6jgms
1/1      Running  0      53s
capi-kubeadm-bootstrap-system               capi-kubeadm-bootstrap-controller-
manager-65ffc94457-7cjd                      1/1      Running  0      74s
capi-kubeadm-control-plane-system           capi-kubeadm-control-plane-controller-
manager-bc7b688d4-vv8wg                      1/1      Running  0      72s
capi-system                                capi-controller-manager-dbfc7b49-dzvws
1/1      Running  0      77s
cappp-system                                cappp-controller-manager-8444d67568-rmms2
1/1      Running  0      59s
capv-system                                capv-controller-manager-58b8ccf868-rbscn
1/1      Running  0      56s
capz-system                                capz-controller-manager-6467f986d8-dnvj4
1/1      Running  0      62s
cert-manager                                cert-manager-6888d6b69b-7b7m9
1/1      Running  0      91s
cert-manager                                cert-manager-cainjector-76f7798c9-gnp8f
1/1      Running  0      91s
cert-manager                                cert-manager-webhook-7d4b5d8484-gn5dr
1/1      Running  0      91s
gce-pd-csi-driver                           csi-gce-pd-controller-5bd587fbfb-lrx29
5/5      Running  0      5m40s
gce-pd-csi-driver                           csi-gce-pd-node-4cgd8
2/2      Running  0      4m22s
gce-pd-csi-driver                           csi-gce-pd-node-5qsfk
2/2      Running  0      4m23s
gce-pd-csi-driver                           csi-gce-pd-node-5w4bq
2/2      Running  0      4m18s
gce-pd-csi-driver                           csi-gce-pd-node-fbdbw
2/2      Running  0      4m23s

```

```

gce-pd-csi-driver          csi-gce-pd-node-h82lx
2/2      Running  0      4m23s
gce-pd-csi-driver          csi-gce-pd-node-jzq58
2/2      Running  0      5m39s
gce-pd-csi-driver          csi-gce-pd-node-k6bz9
2/2      Running  0      2m39s
kube-system                cluster-autoscaler-7f695dc48f-v5kvh
1/1      Running  0      5m40s
kube-system                coredns-64897985d-hbkqd
1/1      Running  0      5m38s
kube-system                coredns-64897985d-m8g5j
1/1      Running  0      5m38s
kube-system                etcd-gcp-example-control-plane-9z77w
1/1      Running  0      5m32s
kube-system                etcd-gcp-example-control-plane-rtj9h
1/1      Running  0      2m37s
kube-system                etcd-gcp-example-control-plane-zbf9w
1/1      Running  0      4m17s
kube-system                kube-apiserver-gcp-example-control-
plane-9z77w          1/1      Running  0      5m32s
kube-system                kube-apiserver-gcp-example-control-plane-
rtj9h          1/1      Running  0      2m38s
kube-system                kube-apiserver-gcp-example-control-plane-
zbf9w          1/1      Running  0      4m17s
kube-system                kube-controller-manager-gcp-example-
control-plane-9z77w  1/1      Running  0      5m33s
kube-system                kube-controller-manager-gcp-example-
control-plane-rtj9h  1/1      Running  0      2m37s
kube-system                kube-controller-manager-gcp-example-
control-plane-zbf9w  1/1      Running  0      4m17s
kube-system                kube-proxy-bskz2
1/1      Running  0      4m18s
kube-system                kube-proxy-gdkn5
1/1      Running  0      4m23s
kube-system                kube-proxy-knvh9
1/1      Running  0      4m22s
kube-system                kube-proxy-tcj7r
1/1      Running  0      4m23s
kube-system                kube-proxy-thdpl
1/1      Running  0      5m38s
kube-system                kube-proxy-txxmb
1/1      Running  0      4m23s
kube-system                kube-proxy-vq6kv
1/1      Running  0      2m39s
kube-system                kube-scheduler-gcp-example-control-
plane-9z77w          1/1      Running  0      5m33s
kube-system                kube-scheduler-gcp-example-control-plane-
rtj9h          1/1      Running  0      2m37s
kube-system                kube-scheduler-gcp-example-control-plane-
zbf9w          1/1      Running  0      4m17s
node-feature-discovery     node-feature-discovery-master-7d5985467-
lh7dc          1/1      Running  0      5m40s

```

```

node-feature-discovery           node-feature-discovery-worker-5qtvq
1/1      Running      0           3m40s
node-feature-discovery           node-feature-discovery-worker-66rwx
1/1      Running      0           3m40s
node-feature-discovery           node-feature-discovery-worker-7h92d
1/1      Running      0           3m35s
node-feature-discovery           node-feature-discovery-worker-b4666
1/1      Running      0           3m40s
tigera-operator                  tigera-operator-5f9bdc5c59-j9tnr
1/1      Running      0           5m38s

```

When you're done exploring the cluster, move onto [making your cluster self-managed](#). (see page 1380)

7.18.6 Make the New GCP Cluster Self-Managed

DKP deploys all cluster lifecycle services to a bootstrap cluster, which then deploys a workload cluster. When the workload cluster is ready, move the cluster lifecycle services to the workload cluster, which makes the workload cluster self-managed. This section describes how to make a workload cluster self-managed.

Before starting, ensure you create a workload cluster as described in [Create a New GCP Cluster](#) (see page 570).

This page contains instructions on how to make your cluster self-managed. This is necessary if there is only one cluster in your environment, or if this cluster should become the Management cluster in a multi-cluster environment.



If you already have a self-managed or Management cluster in your environment, skip this page.

7.18.6.1 Make the new Kubernetes cluster manage itself

Follow these steps:

1. Deploy cluster lifecycle services on the workload cluster:

```
dkp create capi-components --kubeconfig ${CLUSTER_NAME}.conf
```

If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

```
✓ Initializing new CAPI components
```

2. Move the Cluster API objects from the bootstrap to the workload cluster:

The cluster lifecycle services on the workload cluster are ready, but the workload cluster configuration is on the bootstrap cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the bootstrap to the workload cluster. This process is also called a [Pivot](#)¹¹²⁰.

```
dkp move capi-resources --to-kubeconfig ${CLUSTER_NAME}.conf
```

✓ Moving cluster resources

You can now view resources in the moved cluster by using the `--kubeconfig` flag with `kubectl`. For example: `kubectl --kubeconfig=gcp-example.conf get nodes`

i NOTE: To ensure only one set of cluster lifecycle services manages the workload cluster, DKP first pauses reconciliation of the objects on the bootstrap cluster, then creates the objects on the workload cluster. As DKP copies the objects, the cluster lifecycle services on the workload cluster reconcile the objects. The workload cluster becomes self-managed after DKP creates all the objects. If it fails, the `move` command can be safely retried.

3. Wait for the cluster control-plane to be ready:

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --timeout=20m
```

```
cluster.cluster.x-k8s.io/gcp-example condition met
```

4. Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

i NOTE: After moving the cluster lifecycle services to the workload cluster, remember to use DKP with the workload cluster kubeconfig.

```
dkp describe cluster --kubeconfig ${CLUSTER_NAME}.conf -c ${CLUSTER_NAME}
```

NAME	SEVERITY	REASON	SINCE	MESSAGE	READY

¹¹²⁰ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>


```

Cluster/gcp-example                                     True
14s
├─ClusterInfrastructure - GCPCluster/gcp-example
├─ControlPlane - KubeadmControlPlane/gcp-example-control-plane   True
14s
| └─Machine/gcp-example-control-plane-6fbzn                       True
17s
| | └─MachineInfrastructure - GCPMachine/gcp-example-control-plane-62g6s
| └─Machine/gcp-example-control-plane-jf6s2                       True
17s
| | └─MachineInfrastructure - GCPMachine/gcp-example-control-plane-bsr2z
| └─Machine/gcp-example-control-plane-mnbfs                       True
17s
| └─MachineInfrastructure - GCPMachine/gcp-example-control-plane-s8xsx
└─Workers
  └─MachineDeployment/gcp-example-md-0                             True
17s
    └─Machine/gcp-example-md-0-68b86fddb8-8glsw                   True
17s
      | └─MachineInfrastructure - GCPMachine/gcp-example-md-0-zls8d
      └─Machine/gcp-example-md-0-68b86fddb8-bvbm7                 True
17s
        | └─MachineInfrastructure - GCPMachine/gcp-example-md-0-5zcvc
        └─Machine/gcp-example-md-0-68b86fddb8-k9499               True
17s
          | └─MachineInfrastructure - GCPMachine/gcp-example-md-0-k8h5p
          └─Machine/gcp-example-md-0-68b86fddb8-l6vfb             True
17s
            └─MachineInfrastructure - GCPMachine/gcp-example-md-0-9h5vn

```

5. Remove the bootstrap cluster, as the workload cluster is now self-managed:

```

dkp delete bootstrap --kubeconfig $HOME/.kube/config
✓ Deleting bootstrap cluster

```

7.18.6.2 Known Limitations

DKP only supports moving all namespaces in the cluster; DKP does not support migration of individual namespaces.

7.18.7 Manage GCP Node Pools

Node pools are part of a cluster and managed as a group, and you can use a node pool to manage a group of machines using the same common properties. When Konvoy creates a new default cluster, there is one node pool for the worker nodes and all nodes in that new node pool have the same configuration. You can create additional node pools for more specialized hardware or configuration. For example, if you want to tune your memory usage on a cluster where you need maximum memory for some machines and minimal memory on other machines, you would create a new node pool with those specific resource needs.

Konvoy implements node pools using Cluster API [MachineDeployments](#)¹¹²¹.

7.18.7.1 Create GCP Node Pools

Creating a node pool is useful when you need to run workloads that require machines with specific resources, such as additional memory, or specialized network or storage hardware.

7.18.7.1.1 Prerequisites

Before you begin, make sure you have created a [GCP cluster](#). (see page 1376)

7.18.7.1.1.1 Prepare the environment

Follow these steps:

1. Set the environment variable to the name you assigned this cluster.

```
export CLUSTER_NAME=gcp-example
```

2. If your workload cluster is self-managed, as described in [Make the New Cluster Self-Managed](#) (see page 1383), configure `kubectl` to use the kubeconfig for the cluster.

```
export KUBECONFIG=${CLUSTER_NAME}.conf
```

3. Define your node pool name.

```
export NODEPOOL_NAME=example
```

7.18.7.1.1.2 Create a GCP node pool

Availability zones (AZs) are isolated locations within data center regions from which public cloud services originate and operate. Because all the nodes in a node pool are deployed in a single Availability Zone, you may wish to create additional node pools to ensure your cluster has nodes deployed in multiple Availability Zones.

Create a new AWS node pool with 3 replicas using this command:

Set the `--zone` flag to a [zone](#)¹¹²² in the same same region as your cluster.

```
dkp create nodepool gcp ${NODEPOOL_NAME} \
  --cluster-name=${CLUSTER_NAME} \
```

¹¹²¹ <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

¹¹²² <https://cloud.google.com/compute/docs/regions-zones>

```
--image $IMAGE_NAME \
--zone us-west1-b \
--replicas=3
```

```
machinedeployment.cluster.x-k8s.io/example created
.. Creating default/example nodepool resources
gcpmachinetemplate.infrastructure.cluster.x-k8s.io/example created
kubeadmconfigtemplate.bootstrap.cluster.x-k8s.io/example created
✓ Creating default/example nodepool resources
```

This example uses default values for brevity. Use flags to define custom instance types, and other properties.

Advanced users can use a combination of the `--dry-run` and `--output=yaml` or `--output-directory=<your-target-directory>/` flags to get a complete set of node pool objects to modify locally or store in version control.

7.18.7.2 List GCP Node Pools

List node pools for a cluster

Use this command to list the node pools of a given cluster. This returns specific properties of each node pool so that you can see the name of the MachineDeployments.

To list all node pools for a managed cluster, run:

```
dkp get nodepool --cluster-name=${CLUSTER_NAME}
```

The expected output is similar to the following example, indicating the desired size of the node pool, the number of replicas ready in the node pool, and the Kubernetes version those nodes are running:

NODEPOOL VERSION	DESIRED	READY	KUBERNETES
example	3	3	v1.26.6
gcp-example-2-md-0	4	4	v1.26.6

7.18.7.3 Scale GCP Node Pools

While you can run [Cluster Autoscaler](#)¹¹²³, you can also manually scale your node pools up or down when you need more finite control over your environment. For example, if you require 10 machines to run a process, you can manually set the scaling to run those 10 machines only. However, if also using the Cluster Autoscaler, you must stay within your minimum and maximum bounds.

¹¹²³ <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

7.18.7.3.1 Scaling Up Node Pools

To scale up a node pool in a cluster, run:

```
dkp scale nodepool ${NODEPOOL_NAME} --replicas=5 --cluster-name=${CLUSTER_NAME}
```

Your output should be similar to this example, indicating the scaling is in progress:

```
✓ Scaling node pool example to 5 replicas
```

After a few minutes, you can list the node pools to:

```
dkp get nodepool --cluster-name=${CLUSTER_NAME}
```

Your output should be similar to this example, with the number of DESIRED and READY replicas increased to 5:

NODEPOOL	DESIRED	READY	
KUBERNETES VERSION			
example	5	5	v1.26.6
gcp-example-md-0	4	4	v1.26.6

7.18.7.3.2 Scaling Down Node Pools

To scale down a node pool, run:

```
dkp scale nodepool ${NODEPOOL_NAME} --replicas=4 --cluster-name=${CLUSTER_NAME}
```

```
✓ Scaling node pool example to 4 replicas
```

After a few minutes, you can list the node pools using this command:

```
dkp get nodepool --cluster-name=${CLUSTER_NAME}
```

Your output should be similar to this example, with the number of DESIRED and READY replicas decreased to 4:

NODEPOOL	DESIRED	READY	
KUBERNETES VERSION			
example	4	4	v1.26.6

```
gcp-example-md-0          4          4          v1.26.6
```

In a default cluster, the nodes to delete are selected at random. This behavior is controlled by [CAPI's delete policy](#)¹¹²⁴. However, when using the DKP CLI to scale down a node pool, it is also possible to specify the Kubernetes Nodes you want to delete.

To do this, set the flag `--nodes-to-delete` with a list of nodes as below. This adds an annotation `cluster.x-k8s.io/delete-machine=yes` to the matching Machine object that contains `status.NodeRef` with the node names from `--nodes-to-delete`.

```
dkp scale nodepool ${NODEPOOL_NAME} --replicas=3 --cluster-name=${CLUSTER_NAME}
```

```
✓ Scaling node pool example to 3 replicas
```

7.18.7.3.3 Scaling Node Pools When Using Cluster Autoscaler

If you configured [the cluster autoscaler](#)¹¹²⁵ for the `demo-cluster-md-0` node pool, the value of `--replicas` must be within the minimum and maximum bounds.

For example, assuming you have these annotations:

```
kubectl annotate machinedeployment ${NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-
autoscaler-node-group-min-size=2
kubectl annotate machinedeployment ${NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-
autoscaler-node-group-max-size=6
```

Try to scale the node pool to 7 replicas with the command:

```
dkp scale nodepool ${NODEPOOL_NAME} --replicas=7 --cluster-name=${CLUSTER_NAME}
```

Which results in an error similar to:

```
✗ Scaling node pool example to 7 replicas
failed to scale nodepool: scaling MachineDeployment is forbidden: desired replicas 7
is greater than the configured max size annotation cluster.x-k8s.io/cluster-api-
autoscaler-node-group-max-size: 6
```

Similarly, scaling down to a number of replicas less than the configured `min-size` also returns an error.

¹¹²⁴ https://github.com/kubernetes-sigs/cluster-api/blob/v0.4.0/api/v1alpha4/machineset_types.go#L85-L105

¹¹²⁵ <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

7.18.7.4 Delete GCP Node Pools

Delete node pools in a cluster

Deleting a node pool deletes the Kubernetes nodes and the underlying infrastructure. All nodes will be drained prior to deletion and the pods running on those nodes will be rescheduled.

To delete a node pool from a managed cluster, run:

```
dkp delete nodepool ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME}
```

Here, `example` is the node pool to be deleted.

The expected output will be similar to the following example, indicating the node pool is being deleted:

```
✓ Deleting default/example nodepool resources
```

Deleting an invalid node pool results in output similar to this example:

```
dkp delete nodepool ${CLUSTER_NAME}-md-invalid --cluster-name=${CLUSTER_NAME}
```

```
MachineDeployments or MachinePools.infrastructure.cluster.x-k8s.io "no
MachineDeployments or MachinePools found for cluster gcp-example" not found
```

7.18.7.5 GCP Cluster Autoscaler

Configure autoscaler for node pools

[Cluster Autoscaler](https://github.com/kubernetes/autoscaler)¹¹²⁶ provides the ability to automatically scale-up or scale-down the number of worker nodes in a cluster, based on the number of pending pods to be scheduled. Running the Cluster Autoscaler is optional.

Unlike [Horizontal-Pod Autoscaler](https://github.com/kubernetes/autoscaler)¹¹²⁷, Cluster Autoscaler does not depend on any Metrics server and does not need Prometheus or any other metrics source.

The Cluster Autoscaler looks at the following annotations on a MachineDeployment to determine its scale-up and scale-down ranges:

```
cluster.x-k8s.io/cluster-api-autoscaler-node-group-min-size
cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size
```

¹¹²⁶ <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

¹¹²⁷ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-fast-is-hpa-when-combined-with-ca>

The full list of command line arguments to the Cluster Autoscaler controller is [on the Kubernetes public GitHub repository](#)¹¹²⁸.

For more information about how Cluster Autoscaler works, see these documents:

- [What is Cluster Autoscaler](#)¹¹²⁹
- [How does scale-up work](#)¹¹³⁰
- [How does scale-down work](#)¹¹³¹
- [CAPI Provider for Cluster Autoscaler](#)¹¹³²

7.18.7.5.1 Cluster Autoscaler Prerequisites

Before you begin, you must have:

- A [Bootstrap Cluster Lifecycle](#) (see page 1374).
- A [New Kubernetes Cluster](#) (see page 1376).
- A [Self-Managed Cluster](#) (see page 1383).

7.18.7.5.2 Run Cluster Autoscaler

The Cluster Autoscaler controller runs on the workload cluster. Upon creation of the workload cluster, this controller does not have all the objects required to function correctly until after a `dkp move` is issued from the bootstrap cluster.

Run the following steps to enable Cluster Autoscaler:

1. Ensure the Cluster Autoscaler controller is up and running (no restarts and no errors in the logs)

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf logs deployments/cluster-autoscaler
cluster-autoscaler -n kube-system -f
```

2. Enable Cluster Autoscaler by setting the min & max ranges

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment $
{NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-autoscaler-node-group-min-size=2
kubectl --kubeconfig=${CLUSTER_NAME}.conf annotate machinedeployment $
{NODEPOOL_NAME} cluster.x-k8s.io/cluster-api-autoscaler-node-group-max-size=6
```

3. The Cluster Autoscaler logs will show that the worker nodes are associated with node-groups and that pending pods are being watched.
4. To demonstrate that it is working properly, create a large deployment which will trigger pending pods (For this example we used GCP n2-standard-8 worker nodes. If you have larger worker-nodes, you

¹¹²⁸ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-are-the-parameters-to-ca>

¹¹²⁹ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-is-cluster-autoscaler>

¹¹³⁰ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-up-work>

¹¹³¹ <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-down-work>

¹¹³² <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

should scale up the number of replicas accordingly).

```
cat <<EOF | kubectl --kubeconfig=${CLUSTER_NAME}.conf apply -f -
apiVersion: apps/v1
kind: Deployment
metadata:
  name: busybox-deployment
  labels:
    app: busybox
spec:
  replicas: 600
  selector:
    matchLabels:
      app: busybox
  template:
    metadata:
      labels:
        app: busybox
    spec:
      containers:
      - name: busybox
        image: busybox:latest
        command:
          - sleep
          - "3600"
        imagePullPolicy: IfNotPresent
        restartPolicy: Always
EOF
```

5. Cluster Autoscaler will scale up the number of Worker Nodes until there are no pending pods.
6. Scale down the number of replicas for `busybox-deployment` .

```
kubectl --kubeconfig ${CLUSTER_NAME}.conf scale --replicas=30 deployment/
busybox-deployment
```

7. Cluster Autoscaler starts to scale down the number of Worker Nodes after the default timeout of 10 minutes.



Enterprise Customers: In previous DKP versions, the Cluster-Autoscaler did not work on **managed** clusters. Starting with this release, an instance of Cluster-Autoscaler is deployed on the **management cluster** for each **managed** cluster. During an upgrade, the upgrade procedure on the **management cluster** will create the necessary deployments.

7.18.8 Delete a GCP Cluster

7.18.8.1 Prepare to Delete a Workload Cluster

ⓘ A self-managed workload cluster cannot delete itself. If your workload cluster is self-managed, you must create a bootstrap cluster and move the cluster lifecycle services to the bootstrap cluster before deleting the workload cluster.

If you did not make your workload cluster self-managed, as described in [Make the New GCP Cluster Self-Managed](#) (see page 1383), proceed to the [Delete the workload cluster](#) (see page 0) section below.

1. Create a bootstrap cluster:

The bootstrap cluster will host the Cluster API controllers that reconcile the cluster objects marked for deletion:

ⓘ NOTE: To avoid using the wrong `kubeconfig`, the following steps use explicit `kubeconfig` paths and contexts.

```
dkp create bootstrap --with-gcp-bootstrap-credentials=true --kubeconfig
$HOME/.kube/config
```

- ✓ Creating a bootstrap cluster
- ✓ Initializing **new** CAPI components

2. Move the Cluster API objects from the workload to the bootstrap cluster: The cluster lifecycle services on the bootstrap cluster are ready, but the workload cluster configuration is on the workload cluster. The `move` command moves the configuration, which takes the form of Cluster API Custom Resource objects, from the workload to the bootstrap cluster. This process is also called a [Pivot](#)¹¹³³.

```
dkp move capi-resources \
  --from-kubeconfig ${CLUSTER_NAME}.conf \
  --to-kubeconfig $HOME/.kube/config
```

3. Use the cluster lifecycle services on the workload cluster to check the workload cluster status:

¹¹³³ <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```
dkp describe cluster --kubeconfig $HOME/.kube/config -c ${CLUSTER_NAME}
```

```

NAME                                                                 READY
SEVERITY REASON SINCE MESSAGE
Cluster/gcp-example                                                                 True
34s
├─ClusterInfrastructure - GCPCluster/gcp-example
├─ControlPlane - KubeadmControlPlane/gcp-example-control-plane           True
34s
| └─Machine/gcp-example-control-plane-6fbzn                               True
37s
| | └─MachineInfrastructure - GCPMachine/gcp-example-control-plane-62g6s
| | └─Machine/gcp-example-control-plane-jf6s2                               True
37s
| | └─MachineInfrastructure - GCPMachine/gcp-example-control-plane-bsr2z
| | └─Machine/gcp-example-control-plane-mnbfs                               True
37s
| └─MachineInfrastructure - GCPMachine/gcp-example-control-plane-s8xsx
└─Workers
  └─MachineDeployment/gcp-example-md-0                                     True
37s
    └─Machine/gcp-example-md-0-68b86fddb8-8glsw                             True
37s
      └─MachineInfrastructure - GCPMachine/gcp-example-md-0-zls8d
      └─Machine/gcp-example-md-0-68b86fddb8-bvbm7                             True
37s
        └─MachineInfrastructure - GCPMachine/gcp-example-md-0-5zcvc
        └─Machine/gcp-example-md-0-68b86fddb8-k9499                             True
37s
          └─MachineInfrastructure - GCPMachine/gcp-example-md-0-k8h5p
          └─Machine/gcp-example-md-0-68b86fddb8-l6vfb                             True
37s
            └─MachineInfrastructure - GCPMachine/gcp-example-md-0-9h5vn

```

i NOTE: After moving the cluster lifecycle services to the workload cluster, remember to use `dkp` with the workload cluster kubeconfig.



Persistent Volumes (PVs) are not deleted automatically by design in order to preserve your data. However, they take up storage space if not deleted. You must delete PVs manually. Information for backup of a cluster and PVs is on the page in documentation called [Back up your Cluster's Applications and Persistent Volumes](#) (see page 835) .

7.18.8.2 Delete the Workload Cluster

1. To delete a cluster, you would use `dkp delete cluster` and pass in the name of the cluster you are trying to delete with `--cluster-name` flag. You would use `kubectl get clusters` to get those details (`--cluster-name` and `--namespace`) of the Kubernetes cluster to delete it.

NOTE: Do not use `dkp get clusters` since that gets you Kommander cluster details rather than Konvoy cluster details.

```
kubectl get clusters
```

2. Use `dkp` with the bootstrap cluster to delete the workload cluster. Delete the Kubernetes cluster and wait a few minutes:

Before deleting the cluster, `dkp` deletes all Services of type LoadBalancer on the cluster. To skip this step, use the flag `--delete-kubernetes-resources=false`.

```
dkp delete cluster --kubeconfig $HOME/.kube/config --cluster-name $
{CLUSTER_NAME}
```

```
✓ Deleting Services with type LoadBalancer for Cluster default/gcp-example
✓ Deleting ClusterResourceSets for Cluster default/gcp-example
✓ Deleting cluster resources
✓ Waiting for cluster to be fully deleted
Deleted default/gcp-example cluster
```

After the workload cluster is deleted, delete the bootstrap cluster.

7.18.8.3 Delete the Bootstrap Cluster

1. Delete the bootstrap cluster using `dkp delete`:

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

```
✓ Deleting bootstrap cluster
```

7.18.8.4 Known Limitations

The DKP version used to create the workload cluster must match the DKP version used to delete the workload cluster.

8 Additional Kommander Configuration

The Kommander component of DKP can be configured differently depending on your environment type and other desired customizations.

8.1 Installing Kommander by Environment

This section contains instructions for installing the Kommander component in *air-gapped*, *non-air-gapped*, or *pre-provisioned* environments. It also contains instructions on how to enable *DKP catalog apps* or *DKP Insights*, if desired.

8.2 Kommander Customizations

This section contains instructions for enabling a custom configuration of a Kommander component such as custom domains or certificates, an HTTP Proxy, an external load balancer, etc.

8.3 Installing Kommander by Environment

This section provides installation instructions for the Kommander component of DKP according to your environment type.

BEFORE you install the Kommander component, answer the following three questions:

8.3.1 1. Is your Environment Air-gapped or Non-Air-gapped?

In an **air-gapped environment**, your environment is isolated from unsecured networks, like the Internet.

In a **non-air-gapped environment**, your environment has two-way access to and from the Internet.

See <https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Air-gapped%20or%20Non-air-gapped%20Environment%3F&linkCreation=true&fromPageId=273351190>¹¹³⁴ for more information.

8.3.2 2. What License do you Have?

DKP Essential and **DKP Government Essential** are self-managed single-cluster Kubernetes solutions that give you a feature-rich, easy-to-deploy, and easy-to-manage entry-level cloud container platform. The DKP Essential and DKP Gov licenses give the user access to the entire Konvoy cluster environment, and to the Kommander platform application manager.

¹¹³⁴<https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273351190&linkCreation=true&spaceKey=DENT&title=%282.6%29+Air-gapped+or+Non-air-gapped+Environment%3F>

DKP Enterprise and **DKP Government Advanced** are multi-cluster solutions centered around a management cluster that manage multiple attached or managed Kubernetes clusters via a centralized management dashboard. For this license type, you will determine whether or not to use the [DKP Catalog applications](#)¹¹³⁵.

See <https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Licenses&linkCreation=true&fromPagelId=273351190>¹¹³⁶ for more information.

8.3.3 3. Do you Want to Enable DKP Insights?

DKP Insights is a predictive analytics capability that **detects anomalies** that occur either in the present or future and generates an alert in the DKP UI.

See [\(2.6\) DKP Insights](#)¹¹³⁷ for more information.

8.3.4 Next Step:

<https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Identify%20or%20Change%20your%20Default%20StorageClass&linkCreation=true&fromPagelId=273351190>¹¹³⁸

8.3.5 Identify or Change your Default StorageClass

Kommander requires a default `StorageClass`.

For the supported **cloud providers**, the Konvoy component handles the creation of a default `StorageClass`.

For **pre-provisioned** environments, the Konvoy component handles the creation of a `StorageClass` in the form of a `localvolume-provisioner`, which is not suitable for production use. Before installing the Kommander component, you should identify and install a [Kubernetes CSI](#)¹¹³⁹ compatible storage provider that is suitable for production, and then ensure it is set as the default as shown below. See [Provision a Static Local Volume](#) (see [page 97](#)) for more information.

For infrastructure driver specifics, see [Default Storage Providers in DKP](#) (see [page 91](#)).

¹¹³⁵ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273351190&linkCreation=true&spaceKey=DENT&title=%282.6%29+Workspace+Catalog+Applications>

¹¹³⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273351190&linkCreation=true&spaceKey=DENT&title=%282.6%29+Licenses>

¹¹³⁷ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/%282.6%29+DKP+Insights>

¹¹³⁸ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPagelId=273351190&linkCreation=true&spaceKey=DENT&title=%282.6%29+Identify+or+Change+your+Default+StorageClass>

¹¹³⁹ <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

8.3.5.1 Identify your StorageClass

This StorageClass is required to install Kommander.

1. Execute the following command to verify one is configured:

```
kubectl get sc --kubeconfig ${CLUSTER_NAME}.conf
```

The output should look similar to this. Note the (default) after the name:

NAME	PROVISIONER	RECLAIMPOLICY	VOLUMEBINDINGMODE
ALLOWVOLUMEEXPANSION	AGE		
ebs-sc (default)	ebs.csi.aws.com	Delete	WaitForFirstConsumer
false	41s		

2. If the desired StorageClass is not set as default, add the following annotation to the StorageClass manifest:

```
annotations:
  storageclass.kubernetes.io/is-default-class: "true"
```

More information on setting a StorageClass as default can be found at [Changing the default storage class](#)¹¹⁴⁰ in the Kubernetes documentation.

8.3.5.2 Next Step:

[Install Kommander According to your Environment](#) (see page 1396)

8.3.6 Install Kommander According to your Environment

Select the installation type according to your environment:

- [Install Kommander in an Air-gapped Environment](#) (see page 1397)
- [Install Kommander in a Non-air-gapped Environment](#) (see page 1402)
- [Install Kommander in a Pre-provisioned, Air-gapped Environment](#) (see page 1405)
- [Install Kommander in a Pre-provisioned, Non-air-gapped Environment](#) (see page 1412)
- [Install Kommander in a Small Environment](#) (see page 1416)

¹¹⁴⁰ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class>

8.3.6.1 Install Kommander in an Air-gapped Environment

Follow these steps to install the Kommander component of DKP in an [air-gapped environment](#)¹¹⁴¹.

8.3.6.1.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install](#) (see page 121).
- Ensure you have a [default StorageClass](#) (see page 1398).
- Ensure you have loaded all necessary images for your configuration. See [Load the Images into Your Registry: Air-gapped Environments](#) (see page 1400).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

8.3.6.1.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```


2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init --airgapped > kommander.yaml
```

4. *If required*, **customize** your `kommander.yaml`.

 See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include:


- Custom Domains and Certificates
- HTTP proxy
- External Load Balancer
- GPU utilization, etc.
- [Rook Ceph customization for Pre-provisioned environments](#) (see page 1405)

5. *If required*: If your cluster uses a custom **AWS VPC** and requires an internal load-balancer, set the `traefik` annotation to create an internal-facing ELB:

¹¹⁴¹ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/225902613#TermAirgapped>


```
...
apps:
  traefik:
    enabled: true
    values: |
      service:
        annotations:
          service.beta.kubernetes.io/aws-load-balancer-internal: "true"
...
```

- Expand **one** of the following sets of instructions, depending on your license and application environments:

 If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

Essential License: Install Kommander in an Air-Gapped Environment


8.3.6.1.3 Essential License: Install Kommander in an Air-gapped Environment

Use the customized `kommander.yaml` to install DKP in an air-gapped environment:

```
dkp install kommander \
--installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz
```

Enterprise License: Install Kommander in an Air-gapped Environment with DKP Catalog Applications

8.3.6.1.4 Enterprise License: Install Kommander in an Air-gapped Environment with DKP Catalog Applications

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add the following values to enable DKP Catalog Applications:

```
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      path: ./dkp-catalog-applications-v2.6.1.tar.gz
...

```

⚠️ If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander \
--installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz \
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-
v2.6.1.tar.gz

```



Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

8.3.6.1.5 Next Step:

[Verify Kommander Installation](#) (see page 1419)



See [\(2.6\) DKP Insights](#)¹¹⁴²¹¹⁴³ if you want to enable a solution that detects current and future anomalies in workload configurations or Kubernetes clusters.

8.3.6.1.6 Load the Images into Your Registry: Air-gapped Environments

Because air-gapped environments do not have direct access to the Internet, you must download, extract and load several required images to your [local container registry](#) (see page 1533), before installing DKP.

8.3.6.1.6.1 Download all Images for Air-gapped Deployments

If you are operating in an air-gapped environment, a [local container registry](#) (see page 1533) containing all the necessary installation images, including the Kommander images is required. See below for prerequisites to download and then how to push the necessary images to this registry.

1. [Download](#) (see page 78) the Complete DKP Air-gapped Bundle for this release (i.e. `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`) to load registry images as explained below.
2. Connectivity with clusters attaching to the management cluster is required:
 - Both management and attached clusters must be able to connect to the local registry.
 - The management cluster must be able to connect to all attached cluster's API servers.
 - The management cluster must be able to connect to any load balancers created for platform services on the management cluster.

8.3.6.1.6.2 Extract Air-gapped Images and Set Variables

Follow these steps to **extract** the air-gapped image bundles into your private registry:

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address:


¹¹⁴² <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/%282.6%29+DKP+Insights>

¹¹⁴³ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/%282.6%29+DKP+Insights>

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```


8.3.6.1.6.3 Load Images to your Private Registry - Konvoy

Before creating or upgrading a Kubernetes cluster, you need to load the required images in a local registry if operating in an air-gapped environment. This registry must be accessible from both the [bastion machine \(see page 1288\)](#) and either the AWS EC2 instances or other machines that will be created for the Kubernetes cluster.

 If you do not already have a local registry set up, refer to [Local Registry Tools \(see page 1533\)](#) page for more information.

Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-registry-password $REGISTRY_PASSWORD
```

 It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

8.3.6.1.6.4 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** kommander image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

8.3.6.1.6.5 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-
v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

It can take a while to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

8.3.6.1.6.6 Next Step:

[Install Kommander in an Air-gapped Environment \(see page 1400\)](#)

8.3.6.2 Install Kommander in a Non-air-gapped Environment

Install the Kommander component of DKP in [non-air-gapped environments](#)¹¹⁴⁴.

8.3.6.2.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install \(see page 121\)](#).
- Ensure you have a [default StorageClass \(see page 1398\)](#).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

8.3.6.2.2 Create your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

3. Create a [configuration file \(see page 1422\)](#) for the deployment:

```
dkp install kommander --init > kommander.yaml
```

¹¹⁴⁴ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/225902613#TermNonairgapped>

4. If required, customize your `kommander.yaml`.

i See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include:

- Custom Domains and Certificates
- HTTP proxy
- External Load Balancer
- GPU utilization, etc.
- [Rook Ceph customization for Pre-provisioned environments](#) (see page 1405)

5. *If required:* If your cluster uses a custom **AWS VPC** and requires an internal load-balancer, set the `traefik` annotation to create an internal-facing ELB:

```
...
apps:
  traefik:
    enabled: true
    values: |
      service:
        annotations:
          service.beta.kubernetes.io/aws-load-balancer-internal: "true"
...

```

6. Expand **one** of the following sets of instructions, depending on your license and application environments:

i If your environment uses HTTP/HTTPS proxies, you must include the flags `--http-proxy`, `--https-proxy`, and `--no-proxy` and their related values in this command for it to be successful. More information is available in [Configuring an HTTP/HTTPS Proxy](#) (see page 1042).

Essential License: Install Kommander in a Non-Air-Gapped Environment


8.3.6.2.3 Essential License: Install Kommander

Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

Enterprise License: Install Kommander in a Non-air-gapped Environment with DKP Catalog Applications


8.3.6.2.4 Enterprise License: Install Kommander with DKP Catalog Applications

 If you want to enable DKP Catalog applications *after* installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
        ref:
          tag: v2.6.0
...

```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubecfg=${CLUSTER_NAME}.conf
```

8.3.6.2.5



Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

8.3.6.2.6 Next Step:

[Verify Kommander Installation \(see page 1419\)](#)



See [\(2.6\) DKP Insights](#)¹¹⁴⁵¹¹⁴⁶ if you want to enable a solution that detects current and future anomalies in workload configurations or Kubernetes clusters.

8.3.6.3 Install Kommander in a Pre-provisioned, Air-gapped Environment

8.3.6.3.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install \(see page 121\)](#).
- Ensure you have a [default StorageClass \(see page 1398\)](#).
- Ensure you have loaded all necessary images for your configuration. See [Load the Images into Your Registry: Air-gapped Environments \(see page 1403\)](#).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.

¹¹⁴⁵ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/%282.6%29+DKP+Insights>

¹¹⁴⁶ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/%282.6%29+DKP+Insights>

8.3.6.3.2 Create and Customize your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

```
export CLUSTER_NAME=<your-management-cluster-name>
```

2. Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

3. Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init --airgapped > kommander.yaml
```

4. Edit the `kommander.yaml` file to include configuration overrides for the `rook-ceph-cluster`. DKP's default configuration ships Ceph with [PVC based storage](#)¹¹⁴⁷ which requires your CSI provider to support PVC with type `volumeMode: Block`. As this is not possible with the default local static provisioner, you can install Ceph in [host storage](#)¹¹⁴⁸ mode.

You can choose whether Ceph's object storage daemon (osd) pods should consume all or just some of the devices on your nodes. Include **one** of the following Overrides:

- a. To automatically assign all raw storage devices on all nodes to the Ceph cluster:

```
...
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllDevices: true
        useAllNodes: true
        deviceFilter: "<<value>>"
...
```

- b. To assign specific storage devices on all nodes to the Ceph cluster:

```
...
rook-ceph-cluster:
  enabled: true
```

¹¹⁴⁷ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/>

¹¹⁴⁸ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/>

```

values: |
  cephClusterSpec:
    storage:
      storageClassDeviceSets: []
      useAllNodes: true
      useAllDevices: false
      deviceFilter: "^sdb."
  ...

```

Note: If you want to assign specific devices to specific nodes using the `deviceFilter` option, refer to [Specific Nodes and Devices](#)¹¹⁴⁹. For general information on the `deviceFilter` value, refer to [Storage Selection Settings](#)¹¹⁵⁰.

5. *If required:* **Customize** your `kommander.yaml`.

i See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, etc.

6. *If required:* If your cluster uses a custom **AWS VPC** and requires an internal load-balancer, set the `traefik` annotation to create an internal-facing ELB:

```

...
apps:
  traefik:
    enabled: true
    values: |
      service:
        annotations:
          service.beta.kubernetes.io/aws-load-balancer-internal: "true"
  ...

```

7. Expand **one** of the following sets of instructions, depending on your license and application environments:

Essential License: Install Kommander in a Pre-provisioned, Air-Gapped Environment

8.3.6.3.3 Essential License: Install Kommander in an Air-gapped Environment

Use the customized `kommander.yaml` to install DKP in a pre-provisioned, air-gapped environment:

```

dkp install kommander --installer-config kommander.yaml --kubeconfig=${
{CLUSTER_NAME}.conf \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.2.tar.gz \

```


¹¹⁴⁹ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/#specific-nodes-and-devices>

¹¹⁵⁰ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

```
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.2.tar.gz
```

Enterprise License: Install Kommander in a Pre-provisioned, Air-gapped Environment with DKP Catalog Applications


8.3.6.3.4 Enterprise License: Install Kommander in an Air-gapped Environment with DKP Catalog Applications

 If you want to enable DKP Catalog applications **AFTER** installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add the following values to enable DKP Catalog Applications:

```
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      path: ./dkp-catalog-applications-v2.6.2.tar.gz
...

```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubecfg=${
{CLUSTER_NAME}.conf \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.2.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.2.tar.gz
\
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-
v2.6.2.tar.gz
```

8.3.6.3.5 Next Step:

[Verify Kommander Installation](#) (see page 1419)



See [\(2.6\) DKP Insights](#)¹¹⁵¹¹¹⁵² if you want to enable a solution that detects current and future anomalies in workload configurations or Kubernetes clusters.

8.3.6.3.6 Load the Images into Your Registry: Air-gapped, Pre-provisioned Environments

Because air-gapped environments do not have direct access to the Internet, you must download, extract and load several required images to your [local container registry](#) (see page 1533), before installing DKP.

8.3.6.3.6.1 Download all Images for Air-gapped Deployments

If you are operating in an air-gapped environment, a [local container registry](#) (see page 1533) containing all the necessary installation images, including the Kommander images is required. See below for prerequisites to download and then how to push the necessary images to this registry.

1. [Download](#) (see page 78) the Complete DKP Air-gapped Bundle for this release (i.e. `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`) to load registry images as explained below.
2. Connectivity with clusters attaching to the management cluster is required:
 - Both management and attached clusters must be able to connect to the local registry.
 - The management cluster must be able to connect to all attached cluster's API servers.
 - The management cluster must be able to connect to any load balancers created for platform services on the management cluster.

8.3.6.3.6.2 Extract Air-gapped Images and Set Variables

Follow these steps to **extract** the air-gapped image bundles into your private registry:

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address:

¹¹⁵¹ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/%282.6%29+DKP+Insights>

¹¹⁵² <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/%282.6%29+DKP+Insights>

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

8.3.6.3.6.3 Only Pre-provisioned: Load Images for Deployments - Konvoy

For **Pre-provisioned air-gapped environments only**, you must run `konvoy-image upload artifacts` to [copy the artifacts](#) (see page 1207) onto the cluster hosts **before** you begin the [Upgrade the CAPI Components](#) (see page 1576) process later in the upgrade steps.

1. The Kubernetes image bundle will be located in `kib/artifacts/images` and you will want to verify image and artifacts.
 - a. Verify the image bundles exist in `kib/artifacts/images` :

```
$ ls kib/artifacts/images/
kubernetes-images-1.26.14-d2iq.1.tar kubernetes-images-1.26.14-d2iq.1-
fips.tar
```

- b. Verify the artifacts for your OS exist in the `kib/artifacts/` directory and export the appropriate variables:

```
$ ls kib/artifacts/
1.26.14_centos_7_x86_64.tar.gz      1.26.14_redhat_8_x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rhel-8.6-x86_64_fips.tar.gz  pip-packages.tar.gz
1.26.14_centos_7_x86_64_fips.tar.gz  1.26.14_rocky_9_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64_fips.tar.gz  containerd-1.6.17-
d2iq.1-rocky-9.0-x86_64.tar.gz
1.26.14_redhat_7_x86_64.tar.gz      1.26.14_ubuntu_20_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rocky-9.1-x86_64.tar.gz
1.26.14_redhat_7_x86_64_fips.tar.gz  containerd-1.6.17-d2iq.1-centos-7.9-
x86_64.tar.gz      containerd-1.6.17-d2iq.1-rhel-8.4-x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-ubuntu-20.04-x86_64.tar.gz
1.26.14_redhat_8_x86_64.tar.gz      containerd-1.6.17-d2iq.1-centos-7.9-
x86_64_fips.tar.gz  containerd-1.6.17-d2iq.1-rhel-8.6-x86_64.tar.gz
images
```

- c. Set the bundle values with the name from the private registry location:

```
export OS_PACKAGES_BUNDLE=<name_of_the_OS_package>
export CONTAINERD_BUNDLE=<name_of_the_containerd_bundle>
```

For example, for RHEL 8.4 you would set:


```
export OS_PACKAGES_BUNDLE=1.26.14_redhat_8_x86_64.tar.gz
export CONTAINERD_BUNDLE=containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz
```

2. Copy the artifacts:

```
konvoy-image upload artifacts \
  --container-images-dir=./kib/artifacts/images/ \
  --os-packages-bundle=./kib/artifacts/$OS_PACKAGES_BUNDLE \
  --containerd-bundle=./kib/artifacts/$CONTAINERD_BUNDLE \
  --pip-packages-bundle=./kib/artifacts/pip-packages.tar.gz
```


8.3.6.3.6.4 Load Images to your Private Registry - Konvoy

Before creating or upgrading a Kubernetes cluster, you need to load the required images in a local registry if operating in an air-gapped environment. This registry must be accessible from both the [bastion machine \(see page 1288\)](#) and either the AWS EC2 instances or other machines that will be created for the Kubernetes cluster.

 If you do not already have a local registry set up, refer to [Local Registry Tools \(see page 1533\)](#) page for more information.

Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-registry-password $REGISTRY_PASSWORD
```

 It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

8.3.6.3.6.5 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

8.3.6.3.6.6 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

8.3.6.3.6.7 Next Step:

[Install Kommander in a Pre-provisioned, Air-gapped Environment \(see page 1408\)](#)

8.3.6.4 Install Kommander in a Pre-provisioned, Non-air-gapped Environment

Install the Kommander component of DKP in [pre-provisioned \(see page 86\)](#), [non-air-gapped](#)¹¹⁵³ environments.

8.3.6.4.1 Prerequisites

- Ensure you have reviewed all [Prerequisites for Install \(see page 121\)](#).
- Ensure you have a [default StorageClass \(see page 1398\)](#).
- Note down the name of the cluster, where you want to install Kommander. If you do not know it, use `kubectl get clusters -A` to display it.



You **MUST** modify the Kommander installer configuration file (`kommander.yaml`) before installing the Kommander component of DKP in a pre-provisioned environment.

8.3.6.4.2 Create and Customize your Kommander Installer Configuration File

1. Set the environment variable for your cluster:

¹¹⁵³ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/225902613#TermNonairgapped>

```
export CLUSTER_NAME=<your-management-cluster-name>
```

- Copy the `kubeconfig` file of your Management cluster to your local directory:

```
dkp get kubeconfig -c ${CLUSTER_NAME} > ${CLUSTER_NAME}.conf
```

- Create a [configuration file](#) (see page 1422) for the deployment:

```
dkp install kommander --init > kommander.yaml
```

- Edit the installer file to include configuration overrides for the `rook-ceph-cluster`. DKP's default configuration ships Ceph with [PVC based storage](#)¹¹⁵⁴ which requires your CSI provider to support PVC with type `volumeMode: Block`. As this is not possible with the default [local static provisioner](#) (see page 91), you can install Ceph in [host storage](#)¹¹⁵⁵ mode.

You can choose whether Ceph's object storage daemon (osd) pods should consume all or just some of the devices on your nodes. Include **one** of the following Overrides:

- To automatically assign all raw storage devices on all nodes to the Ceph cluster:

```
...
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllDevices: true
        useAllNodes: true
        deviceFilter: "<<value>>"
...
```

- To assign specific storage devices on all nodes to the Ceph cluster:

```
...
rook-ceph-cluster:
  enabled: true
  values: |
    cephClusterSpec:
      storage:
        storageClassDeviceSets: []
        useAllNodes: true
        useAllDevices: false
```

¹¹⁵⁴ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/>

¹¹⁵⁵ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/>


```
...
    deviceFilter: "^sdb."
...

```

Note: If you want to assign specific devices to specific nodes using the `deviceFilter` option, refer to [Specific Nodes and Devices](#)¹¹⁵⁶. For general information on the `deviceFilter` value, refer to [Storage Selection Settings](#)¹¹⁵⁷.

5. *If required:* Customize your `kommander.yaml`.

i See [Kommander Customizations](#) (see page 1422) for customization options. Some of them include: Custom Domains and Certificates, HTTP proxy, External Load Balancer, GPU utilization, etc.

6. *If required:* If your cluster uses a custom **AWS VPC** and requires an internal load-balancer, set the `traefik` annotation to create an internal-facing ELB:

```
...
apps:
...
  traefik:
    enabled: true
    values: |
      service:
        annotations:
          service.beta.kubernetes.io/aws-load-balancer-internal: "true"
...

```

7. Expand **one** of the following sets of instructions, depending on your license and application environments:

Essential License: Install Kommander in a Pre-provisioned, Non-Air-Gapped Environment

8.3.6.4.3 Essential License: Install Kommander

Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf


```

Enterprise License: Install Kommander in a Pre-provisioned, Non-air-gapped Environment with DKP Catalog Applications

¹¹⁵⁶ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/host-cluster/#specific-nodes-and-devices>

¹¹⁵⁷ <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>


8.3.6.4.4 Enterprise License: Install Kommander with DKP Catalog Applications

 If you want to enable DKP Catalog applications after installing DKP, see [Enable DKP Catalog Applications after Installing DKP](#) (see page 1458).

1. In the same `kommander.yaml` of the previous section, add these values for `dkp-catalog-applications`:

```
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
        ref:
          tag: v2.6.0
...

```

 If you only want to enable catalog applications to an **existing** configuration, add these values to an **existing** installer configuration file to maintain your Management cluster's settings.

2. Use the customized `kommander.yaml` to install DKP:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).

- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

8.3.6.4.5 Next Step:

[Verify Kommander Installation \(see page 1419\)](#)



See [\(2.6\) DKP Insights](#)¹¹⁵⁸¹¹⁵⁹ if you want to enable a solution that detects current and future anomalies in workload configurations or Kubernetes clusters.

8.3.6.5 Install Kommander in a Small Environment

You can install the Kommander component of DKP on a small, single-cluster environment with smaller memory, storage, and CPU requirements for testing and demo purposes. This topic describes methods for installing Kommander in these environments.



Enterprise considerations: D2iQ recommends performing testing and demo tasks in a single-cluster environment. The Enterprise license is designed for multi-cluster environments and fleet management, which require a [minimum amount of resources \(see page 103\)](#). Applying an Enterprise license key to the previous installation adds modifications to your environment that can exhaust a small environment's resources.

8.3.6.5.1 Prerequisites

Ensure you have done the following:

- You have acquired a DKP license.
- You have installed the [Konvoy component \(see page 147\)](#).
- You have reviewed the [prerequisite section \(see page 121\)](#) pertaining to your air-gapped, or networked environment.

¹¹⁵⁸ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/%282.6%29+DKP+Insights>

¹¹⁵⁹ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/%282.6%29+DKP+Insights>

8.3.6.5.2 Minimal Kommander installation

The YAML file that is used to install a minimal configuration of Kommander contains the bare minimum setup that allows you to deploy applications, and access the DKP UI. It does **NOT** include applications for cost monitoring, logging, alerting, object storage, etc.

In this YAML file you can find the lines that correspond to all platform applications which would be included in a normal Kommander setup. Applications that have `enabled` set to `false` are not taken into account during installation. If you want to test an additional application, you can enable it individually to be installed by setting `enabled` to `true` on the corresponding line in the YAML file.

For example, if you want to enable the logging stack, set `enabled` to `true` for `grafana-logging`, `grafana-loki`, `logging-operator`, `rook-ceph` and `rook-ceph-cluster`. Note that depending on the size of your cluster, enabling several platform applications could exhaust your cluster's resources.



Some applications depend on other applications to work properly. Refer to the [dependencies documentation](#) (see page 644) to find out which other applications you need to enable to test the target application.

1. Initialize your Kommander installation and name it `kommander_minimal.yaml`:

```
dkp install kommander --init --kubeconfig=${CLUSTER_NAME}.conf -oyaml >
kommander_minimal.yaml
```

2. Edit your `kommander_minimal.yaml` file to match the following example:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
  dex:
    enabled: true
  dex-k8s-authenticator:
    enabled: true
  dkp-insights-management:
    enabled: false
  gatekeeper:
    enabled: true
  gitea:
    enabled: true
  grafana-logging:
    enabled: false
  grafana-loki:
    enabled: false
```

```

kommander:
  enabled: true
kommander-ui:
  enabled: true
kube-prometheus-stack:
  enabled: false
kubernetes-dashboard:
  enabled: false
kubefed:
  enabled: true
kubetunnel:
  enabled: false
logging-operator:
  enabled: false
prometheus-adapter:
  enabled: false
reloader:
  enabled: true
rook-ceph:
  enabled: false
rook-ceph-cluster:
  enabled: false
traefik:
  enabled: true
traefik-forward-auth-mgmt:
  enabled: true
velero:
  enabled: false
ageEncryptionSecretName: sops-age
clusterHostname: ""

```

3. Install Kommander on your cluster with the following command:

```

dkp install kommander --installer-config ./kommander_minimal.yaml --
kubeconfig=${CLUSTER_NAME}.conf

```

In the previous command, the `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you set the context to **install Kommander on the right cluster**. For alternatives and recommendations around setting your context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

TIP: Sometimes, applications require a longer period of time to deploy, which causes the installation to time out. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.

If the Kommander installation fails, or you wish to reconfigure applications, you can rerun the install command, and you can view the progress by increasing the log verbosity by adding the flag `-v 2`.

8.3.6.5.3 Next Step:

[Verify Kommander Installation \(see page 1419\)](#)

8.3.7 Verify Kommander Installation

After the Konvoy cluster is built and Kommander has been installed, you will want to verify your installation of Kommander. After the CLI successfully installs the components, it will wait for all applications to be ready by default.



NOTE: If the Kommander installation fails or you wish to reconfigure applications, you can rerun the install command to retry the installation.

If you prefer the CLI to not wait for all applications to become ready, you can set the `--wait=false` flag.

If you choose not to wait via the DKP CLI, you can check the status of the installation using the following command:

```
kubectl -n kommander wait --for condition=Released helmreleases --all --timeout 15m
```

This will wait for each of the helm charts to reach their `Released` condition, eventually resulting in an output resembling below:

```
helmrelease.helm.toolkit.fluxcd.io/centralized-grafana condition met
helmrelease.helm.toolkit.fluxcd.io/dex condition met
helmrelease.helm.toolkit.fluxcd.io/dex-k8s-authenticator condition met
helmrelease.helm.toolkit.fluxcd.io/fluent-bit condition met
helmrelease.helm.toolkit.fluxcd.io/gitea condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-logging condition met
helmrelease.helm.toolkit.fluxcd.io/grafana-loki condition met
helmrelease.helm.toolkit.fluxcd.io/karma condition met
helmrelease.helm.toolkit.fluxcd.io/kommander condition met
helmrelease.helm.toolkit.fluxcd.io/kommander-appmanagement condition met
helmrelease.helm.toolkit.fluxcd.io/kube-prometheus-stack condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost condition met
helmrelease.helm.toolkit.fluxcd.io/kubecost-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/kubefed condition met
helmrelease.helm.toolkit.fluxcd.io/kubernetes-dashboard condition met
helmrelease.helm.toolkit.fluxcd.io/kubetunnel condition met
```

```
helmrelease.helm.toolkit.fluxcd.io/logging-operator condition met
helmrelease.helm.toolkit.fluxcd.io/logging-operator-logging condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-adapter condition met
helmrelease.helm.toolkit.fluxcd.io/prometheus-thanos-traefik condition met
helmrelease.helm.toolkit.fluxcd.io/reloader condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph condition met
helmrelease.helm.toolkit.fluxcd.io/rook-ceph-cluster condition met
helmrelease.helm.toolkit.fluxcd.io/thanos condition met
helmrelease.helm.toolkit.fluxcd.io/traefik condition met
helmrelease.helm.toolkit.fluxcd.io/traefik-forward-auth-mgmt condition met
helmrelease.helm.toolkit.fluxcd.io/velero condition met
```

8.3.7.1 Failed HelmReleases

If any application fails to successfully deploy, you can check the status of a `HelmRelease` with:

```
kubectl -n kommander get helmrelease <HELMRELEASE_NAME>
```

If you find any `HelmReleases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": false}]'
```

8.3.7.2 Next Step:

[Log in to the UI with Kommander \(see page 1420\)](#)

8.3.8 Log in to the UI with Kommander

By default, you can log in to the UI in Kommander with the credentials given using this command:

```
dkp open dashboard --kubeconfig=${CLUSTER_NAME}.conf
```

You can also retrieve your credentials at any time using the following command:

```
kubectl -n kommander get secret dkp-credentials -o go-template='Username:
{{.data.username|base64decode}}{"\n"}}Password: {{.data.password|base64decode}}
{{ "\n"}}'
```

You can also retrieve the URL used for accessing the UI using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='https://{{with index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}/dkp/kommander/dashboard{{ "\n"}}
```

You should only use these static credentials to access the UI for configuring an [external identity provider](#) (see [page 600](#)). Treat them as back up credentials rather than use them for normal access.

You rotate the password using the following command:

```
dkp experimental rotate dashboard-password
```

The output displays the new password:

```
Password: kqZ31lMBSClCbjUKVwLJMQl2PxaIipIzZw5Pjyw09wDqjWV3dz2wPSSBYi09JGJp
```

You can perform the following operations on [Identity Providers](#) (see [page 600](#)):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

8.3.8.1 Dashboard UI Functions

After installing Konvoy component and building a cluster as well as successfully installing Kommander and logging into the UI, you are now ready to customize configurations using the [Day 2 Cluster Operations Management](#) (see [page 585](#)) section of the documentation. The majority of this customization such as attaching clusters and deploying applications will take place in the dashboard or UI of DKP. The Day 2 section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.



See [\(2.6\) DKP Insights](#)¹¹⁶⁰ if you want to enable a solution that detects current and future anomalies in workload configurations or Kubernetes clusters.

8.3.8.2 Next Step

[Day 2 - Cluster Operations Management](#) (see [page 585](#))

¹¹⁶⁰ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274956432/%282.6%29+DKP+Insights>

8.4 Kommander Customizations

You can configure the Kommander component of DKP during the initial installation, and also post-installation using the DKP CLI.

8.4.1 Prerequisites

- Review the [Management cluster application requirements \(see page 103\)](#) to ensure your cluster has enough resources.
- Ensure you have a default `StorageClass`, as shown in [Identify or Change your Default StorageClass \(see page 1398\)](#).

8.4.2 Initialize a Kommander Installer Configuration File

To begin configuring Kommander, run the following command to initialize a default configuration file:

- For a non-air-gapped environment, run the following command:

```
dkp install kommander --init > kommander.yaml
```

- For an air-gapped environment, run the following command:

```
dkp install kommander --init --airgapped > kommander.yaml
```

8.4.3 Configure Applications

After you have a default configuration file, you can then configure each `app` either inline **or** by referencing another YAML file. The configuration values for each `app` correspond to the Helm Chart values for the application.

After the initial deployment of Kommander, you can find the application Helm Charts by checking the `spec.chart.spec.sourceRef` field of the associated `HelmRelease` :

```
kubectl get helmreleases <application> -o yaml -n kommander
```

8.4.3.1 Inline configuration (using values)

In this example, you configure the `centralized-grafana` application with resource limits by defining the Helm Chart values in the Kommander configuration file.

```

apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
  centralized-grafana:
    values: |
      grafana:
        resources:
          limits:
            cpu: 150m
            memory: 100Mi
          requests:
            cpu: 100m
            memory: 50Mi
  ...

```

8.4.3.2 Reference another YAML file (using valuesFrom)

Alternatively, you could create another YAML file containing the configuration for `centralized-grafana` and reference that using `valuesFrom`. Point to this file by using either a relative path (from the configuration file location) or by using an absolute path.

```

cat > centralized-grafana.yaml <<EOF
grafana:
  resources:
    limits:
      cpu: 150m
      memory: 100Mi
    requests:
      cpu: 100m
      memory: 50Mi
EOF

```

```

apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
  centralized-grafana:
    valuesFrom: centralized-grafana.yaml
  ...

```

8.4.4 Minimal Kommander Installation

You can install Kommander with a bare minimum of applications on a small environment with smaller memory, storage, and CPU requirements for testing and demo purposes. Refer to the [Install DKP on a Small Environment](#) (see page 1419) documentation for more information.

8.4.5 Install with Configuration File

In the following command, the `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you set the context to **install Kommander on the right cluster**. For alternatives and recommendations around setting your context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).

Add the `--installer-config` flag to the `kommander install` command to use a custom configuration file. To reconfigure applications, you can also run this command after the initial installation.

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

TIP: Sometimes, applications require a longer period of time to deploy, which causes the installation to time out. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.

8.4.6 Verify Installation

After building the Konvoy cluster and installing Kommander, you can [verify your Kommander installation](#) (see page 1422), then you can [Log in to the Kommander UI](#). (see page 1423)


8.4.7 DKP Kommander Configuration Reference

8.4.7.1 Configuration Parameters



This page contains the configuration parameters for the Kommander component of DKP.

For additional information about configuring the Kommander component of DKP during initial installation, see [Kommander Install Configuration](#) (see page 1425).

Parameter	Description	Default Value
<p>apps (AppConfig list (see page 1427))</p>	<p>List of platform applications that will be installed on the management cluster.</p>	<pre>apps: dex: enabled: true dex-k8s-authenticator: enabled: true dkp-insights-management: enabled: true gatekeeper: enabled: true gitea: enabled: true grafana-logging: enabled: true grafana-loki: enabled: true kommander: enabled: true kube-prometheus-stack: enabled: true values: <shortened for brevity> kubefed: enabled: true kubernetes-dashboard: enabled: true kubetunnel: enabled: true logging-operator: enabled: true prometheus-adapter: enabled: true reloader: enabled: true rook-ceph: enabled: true rook-ceph-cluster: enabled: true traefik: enabled: true traefik-forward-auth- mgmt: enabled: true velero: enabled: true</pre>

Parameter	Description	Default Value
ageEncryptionSecretName	Defines the name of the secret to store the Age encryption in.	sops-age
clusterHostName	Allows users to provide a hostname that is used for accessing the cluster's ingresses.	
ingressCertificate (see page 1428)	Allows users to provide a custom certificate that's used for TLS in the cluster's ingresses.	
acme	Enable automatic ingress certificate management via ACME.	
appManagementImageTag	Specifies image tag of the AppManagement container.	
appManagementImageRepository	Specifies the image repository of AppManagement container	
appManagementKubetoolsImageRepository	Specifies the image repository of AppManagement Kubetools container	
kommanderChartsVersion	Specifies DKP Kommander Helm chart version.	
airgapped (see page 1428)	Specifies parameters for an airgapped environment.	
catalog Enterprise Gov Advanced	Specifies parameters for installing default catalog repositories. <div style="border: 1px solid purple; padding: 10px; margin-top: 10px;">  For additional information, see Configure an Enterprise catalog (see page 1458). </div>	

8.4.7.2 AppConfig

Parameter	Description	Default Value
enabled	Denotes whether the specific app should be deployed or not.	false
valuesFrom	<p>File path containing the values that are passed onto the application's <code>helmRelease</code>.</p> <p>This a Helm values files for all applications at the moment. The path in this field must either be a relative file location which is then interpreted to be relative to the location of the configuration file or an absolute path.</p> <div style="border: 1px solid purple; padding: 5px; margin-top: 10px;"> <p> Only one of <code>valuesFrom</code> or <code>values</code> may be set; both cannot be set.</p> </div>	
values	<p>Contains the values that are passed to the the application's <code>helmRelease</code>.</p> <div style="border: 1px solid purple; padding: 5px; margin-top: 10px;"> <p> Only one of <code>valuesFrom</code> or <code>values</code> may be set; both cannot be set.</p> </div>	

8.4.7.3 IngressCertificate

Parameter	Description	Default Value
certificate	The path to a certificate PEM file.	
private_key	The path to the certificate's private key (PEM).	
ca	The path to the certificate's CA bundle; a PEM file containing root and intermediate certificates.	

8.4.7.4 Airgapped

Parameter	Description	Default Value
enabled	Specifies if installation happens in an air-gapped environment.	
helmMirrorImageTag	Specifies an image tag of Helm-mirror container.	
helmMirrorImageRepository	Specifies image repository of Helm-mirror container.	

8.4.7.4.1 Next Step:

[Configure HTTP Proxy \(see page 1452\)](#)

8.4.8 Custom Domains and Certificates for Management and Essential Clusters

Configure a custom domain and certificate during the installation of the Kommander component.

For [Management](#) (see page 82) or [Essential](#) (see page 83) clusters only.

There are two configuration methods:

Configuration Method	<i>WHILE</i> installing the Kommander component	<i>AFTER</i> installing the Kommander component
Supported cluster types	Only Essential or Management clusters	All cluster types
Documentation	Remain in this section	Go to Custom Domains and Certificates Configuration for All Cluster Types (see page 861)

DKP supports configuring a custom **domain** name for accessing the UI and other platform services, as well as setting up manual or automatic **certificate renewal** or rotation. This section provides instructions and examples on how to configure the DKP installation to add a customized domain and certificate on your **Essential cluster** (see page 83) or your **Management cluster** (see page 82).

Section contents:

- [Why to set up a Custom Domain or Certificate?](#) (see page 1429)
- [Certificate Issuer and KommanderCluster Concepts](#) (see page 1430)
- [Certificate Authority \(CA\) Specifics](#) (see page 1431)
- [Configure the Kommander Installation with a Custom Domain and Certificate](#) (see page 1432)
- [Verify and Troubleshoot the Domain and Certificate Customization](#) (see page 1441)

8.4.8.1 Why to set up a Custom Domain or Certificate?

8.4.8.1.1 Reasons for Using a Custom DNS Domain

DKP supports the customization of domains to allow you to use your own domain or hostname for your services. For example, you can set up your DKP UI or any of your clusters to be accessible with your custom domain name instead of the domain provided by default.

- To set up a custom domain (without a custom certificate), refer to [Configure a Custom Domain without a Custom Certificate](#) (see page 1440).

8.4.8.1.2 Reasons for Using a Custom Certificate

DKP's default CA identity supports the encryption of data exchange and traffic (between your client and your environment's server). To configure an additional security layer that validates your environment's server authenticity, DKP supports configuring a custom certificate issued by a trusted Certificate Authority either directly in a Secret or managed automatically using the ACME protocol (for example, Let's Encrypt).

Changing the default certificate for any of your clusters can be helpful. For example, you can adapt it to classify your DKP UI or any other type of service as trusted (when accessing a service via a browser).

To set up a custom domain and certificate, refer to the following pages respectively:

- Configure a custom domain and certificate as part of the [cluster's installation process](#) (see page 1432). This is only possible for your [Management/Essential cluster](#) (see page 81).

- Update your cluster's current domain and certificate configuration as part of your [cluster's Day 2 operations](#) (see page 864). You can do this for any [cluster type](#) (see page 81) in your environment.



Using Let's Encrypt or other public ACME certificate authorities **does not work in air-gapped scenarios**, as these services require connection to the Internet for their setup. For air-gapped environments, you can either use self-signed certificates issued by the cluster (the default configuration), or a certificate created manually using a trusted Certificate Authority.

8.4.8.1.3 Next Topic:

[Certificate Issuer and KommanderCluster Concepts](#) (see page 1430)

8.4.8.2 Certificate Issuer and KommanderCluster Concepts

If you set up an [advanced configuration of your custom domain](#) (see page 1434), ensure you understand the following concepts.

8.4.8.2.1 KommanderCluster Object

The `KommanderCluster` resource is an object that contains key information for [all types of clusters](#) (see page 81) that are part of your environment, such as:

- Cluster access and endpoint information
- Cluster attachment information
- Cluster status and configuration information

8.4.8.2.2 Issuer Objects:

`Issuer` , `ClusterIssuer` or `certificateSecret` ?

If you use a certificate issued and managed automatically by `cert-manager` , you need an `Issuer` or `ClusterIssuer` that you reference in your `KommanderCluster` resource. The referenced object must contain the information of your certificate provider.

If you want to use a manually-created certificate, you need a `certificateSecret` that you reference in your `KommanderCluster` resource.

8.4.8.2.3 Location of the `KommanderCluster` and `Issuer` Objects: Management, Managed or Attached cluster?

In the **Management** (see page 82) or **Essential** (see page 83) cluster, both the `KommanderCluster` and `issuer` objects are stored on the same cluster. The `issuer` can be referenced as an `Issuer`, `ClusterIssuer` or `certificateSecret`.

In **Managed** (see page 82) and **Attached** (see page 82) clusters, the `KommanderCluster` object is stored on the Management cluster. The `Issuer`, `ClusterIssuer` or `certificateSecret` is stored on the Managed or Attached cluster.

8.4.8.2.4 HTTP or DNS solver?

When configuring a certificate for your DKP cluster, you can set up an *HTTP solver* or a *DNS solver*. The HTTP protocol exposes your cluster to the public Internet, whereas DNS keeps your traffic hidden. If you use HTTP, your cluster must be publically accessible (via the ingress or load balancer). If you use DNS, this is not a requirement. See [Advanced Configuration: ClusterIssuer](#) (see page 1437) for HTTP and DNS configuration options.



If you are enabling a [proxied access](#) (see page 817) for a [network-restricted cluster](#) (see page 84), this configuration is restricted to DNS.

8.4.8.2.5 Related topics:

[Why to set up a Custom Domain or Certificate?](#) (see page 1432)

[Configure the Kommander Installation with a Custom Domain and Certificate](#) (see page 1432)

[Advanced Configuration: ClusterIssuer](#) (see page 1437)

[Certificate Authority \(CA\) Specifics](#) (see page 1431)

8.4.8.3 Certificate Authority (CA) Specifics

The following table defines some values you require to set up a custom certificate according to your Certificate Authority (CA).

Certificate Authority	Prerequisites	<i>Kommander Installer values</i>
Let's Encrypt	None	Generated automatically by Kommander when <code>acme</code> is enabled
ZeroSSL	An access and a secret key provided by ZeroSSL	<code>acme.server: https://acme.zerossll.com/v2/DV90</code>
SSL.com	An access and a secret key provided by SSL	<code>acme.server: https://acme.ssl.com/sslcom-dv-rsa</code>

Use these values to [Configure your Custom Domain and Certificate](#) (see page 1432).

8.4.8.3.1 Next Topic:

[Configure your Custom Domain and Certificate](#) (see page 1432)

8.4.8.4 Configure the Kommander Installation with a Custom Domain and Certificate

This page contains instructions on how to set up custom certificates for your [Management](#) (see page 82) or [Essential](#) (see page 83) cluster during the installation of DKP.

There are two configuration methods:

Configuration Method	<i>WHILE</i> installing the Kommander component	<i>AFTER</i> installing the Kommander component
Supported cluster types	Only Essential or Management clusters	All cluster types
Documentation	Remain in this page	Go to Configure Custom Domains or Custom Certificates post Kommander Installation (see page 864)

8.4.8.4.1 Configuration Options

- Choose an ACME-supported Certificate Authority, if you want the `cert-manager` to automatically handle **certificate renewal** and **rotation**.
- Refer to [Certificate Authority \(CA\) Specifics](#) (see page 1434) for more information on values that are specific to your Certificate Authority or CA.

I want to use an automatically-generated certificate with ACME and require basic configuration*

8.4.8.4.1.1 I want to use an **automatically-generated** certificate with **ACME** and require **basic** configuration*

When you enable ACME, by default DKP generates an ACME-supported certificate with an HTTP01 solver. The `cert-manager` automatically issues a trusted certificate for the configured custom domain, and takes care of renewing the certificate before expiration.

1. Open the *Kommander Installer Configuration File* or `<kommander.yaml>` file:
 - a. If you do not have the `<kommander.yaml>` file, [initialize the configuration file](#) (see page 1425), so you can edit it in the following steps. **WARNING:** Initialize this file only ONE time, otherwise you will overwrite previous customizations.
 - b. If you have initialized the configuration file already, open the `<kommander.yaml>` with the editor of your choice.
2. In that file, configure the custom domain for your cluster:

```
[...]
clusterHostname: <mycluster.example.com>
[...]
```

3. Enable ACME by adding `acme` value, the issuer's server and your e-mail. If you don't provide a server, DKP sets up **Let's Encrypt** as your certificate provider:

```
acme:
  email: <your_email>
  server: <your_server>
[...]
```


4. [Use the configuration file to install Kommander](#) (see page 1427).

*basic configuration: ACME server without EAB (External Account Bindings) and HTTP solver

I want to use an automatically-generated certificate with ACME and require advanced configuration (e.g. EAB, DNS solver, etc.)

8.4.8.4.1.2 I want to use an **automatically-generated** certificate with **ACME** and require **advanced** configuration

If you require additional configuration options like DNS solver, EAB, among others, create a `ClusterIssuer` with the required configurations before you run the installation of Kommander. The `cert-manager` automatically issues a trusted certificate for the configured custom domain, and takes care of renewing the certificate before expiration.

 To read more about the `ClusterIssuer`, other objects, and where to store them, refer to [Advanced Configuration: ClusterIssuer \(see page 1437\)](#) and [Certificate Issuer and KommanderCluster Concepts \(see page 1433\)](#).

1. Create a `ClusterIssuer` and store it in the target cluster. It **must** be called `kommander-acme-issuer`:
 - a. If you require an **HTTP** solver, adapt the following example with the properties required for your certificate and execute the command:

```
cat <<EOF | kubectl apply -f -
apiVersion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
  name: kommander-acme-issuer # This part is important
spec:
  acme:
    email: <your_email>
    server: <https://acme.server.example>
    skipTLSVerify: true
    privateKeySecretRef:
      name: kommander-acme-issuer-account # Set this to <name>-account
  solvers:
  - http01:
      ingress:
        ingressTemplate:
          metadata:
            annotations:
              kubernetes.io/ingress.class: kommander-traefik
              "traefik.ingress.kubernetes.io/router.priority":
"2147483647"
EOF
```

⚠ The values `kommander-acme-issuer`, `kommander-acme-issuer-account` and `"traefik.ingress.kubernetes.io/router.priority": "2147483647"` are not placeholders and MUST be filled out exactly as in the example.

📘 In **on-premise** environments, replace the annotation in the previous example with `traefik.ingress.kubernetes.io/router.tls: "true"`.

- b. If you require a **DNS** solver, adapt the following example with the properties required for your certificate and execute the command:

```
cat <<EOF | kubectl apply -f -
apiVersion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
  name: kommander-acme-issuer # This part is important
spec:
  acme:
    email: <your_email>
    server: <https://acme.server.example>
    privateKeySecretRef:
      name: kommander-acme-issuer-account # Set this to <name>-account
    solvers:
      - dns01:
          route53:
            region: us-east-1
            role: arn:aws:iam::YYYYYYYYYYYYY:role/dns-manager
EOF
```

⚠ The values `kommander-acme-issuer`, `kommander-acme-issuer-account` are not placeholders and MUST be filled out exactly as in the example.

2. Optional: If you require External Account Bindings to link your ACME account to an external database, refer to <https://cert-manager.io/docs/configuration/acme/#external-account-bindings>.
3. Optional: Create a DNS record, by setting up the `external-dns` service (see page 1442). This way, the `external-dns` will take care of pointing the DNS record to the ingress of the cluster automatically.
 - 📘 You can also create a DNS record manually, that maps your domain name or IP address to the cluster ingress. If you choose to create a DNS record manually, finish installing the Kommander component, and then manually create a DNS record that points to the load balancer address.
4. Open the **Kommander Installer Configuration File** or `kommander.yaml` file:
 - a. If you do not have the `kommander.yaml` file, [initialize the configuration file](#) (see page 1425), so you can edit it in the following steps. **WARNING:** Initialize this file only ONCE, otherwise you will overwrite previous customizations.
 - b. If you have initialized the configuration file already, open the `kommander.yaml` with the editor of your choice.
5. In that file, configure the cluster to use your custom domain:

```
[...]
clusterHostname: <mycluster.example.com>
[...]
```

6. Enable ACME by configuring the issuer's server and your e-mail:

```
[...]
acme:
  email: <your_email>
  server: <your_server>
[...]
```

7. [Use the configuration file to install Kommander \(see page 1427\)](#).

I have a manually-generated certificate

8.4.8.4.1.3 I have a **manually-generated** certificate

D2iQ supports the use of a manually-created certificate. In this case, there is no certificate controller that handles the renewal and update of your certificate automatically, so you will have to take care of these tasks manually.

Prerequisites:

- Obtain the PEM files of your certificate and store them in the target cluster's namespace:
 - Certificate
 - certificate's private key
 - CA bundle (containing the root and intermediate certificates)

Configure the manually-generated certificate

1. Open the **Kommander Installer Configuration File** or `<kommander.yaml>` file:
 - a. If you do not have the `<kommander.yaml>` file, [initialize the configuration file \(see page 1425\)](#), so you can edit it in the following steps. **WARNING:** Initialize this file only ONCE, otherwise you will overwrite previous customizations.
 - b. If you have initialized the configuration file already, open the `<kommander.yaml>` with the editor of your choice.
2. In the *Kommander Installer Configuration file*, provide your custom domain and the paths to the PEM files of your certificate:

```
[...]
clusterHostname: <mycluster.example.com>
ingressCertificate:
```

```
certificate: <certs/cert.pem>
private_key: <certs/key.pem>
ca: <certs/ca.pem>
[...]
```

3. Use the configuration file to install Kommander (see page 1427).



Certificates issued by another Issuer

You can also configure a certificate issued by another Certificate Authority. In this case, the CA will determine which information to include in the configuration.

- Refer to <https://cert-manager.io/docs/configuration/> for configuration examples.
- The `ClusterIssuer` 's name **MUST BE** `kommander-acme-issuer` .

8.4.8.4.1.4 Next Step:

[Verify and Troubleshoot the Domain and Certificate Customization](#) (see page 1441)

8.4.8.4.1.5 Related Topics:

- [Advanced Configuration: ClusterIssuer](#) (see page 1437)
- [Configure a Custom Domain without a Custom Certificate](#) (see page 1440)

8.4.8.4.2 Advanced Configuration: ClusterIssuer



Ensure you review [Certificate Issuer and KommanderCluster Concepts](#) (see page 1433) to understand the basic concepts for certificate configuration.

When you enable ACME (see page 1436), by default DKP generates an ACME-supported certificate with an HTTP01 solver that is provided by Let's Encrypt.


You can also set up an [advanced configuration for a Custom Domain and Certificate](#) (see page 1437). In these cases, the custom configuration cannot be done completely via the `installer config` file, but must be specified further in a `ClusterIssuer` .

Whether it is sufficient to establish the configuration of your custom certificate in the `installer config` file only, or you require a `ClusterIssuer` to define further configuration options depends on the degree of customization.

⚠ If you require a `ClusterIssuer`, you MUST create it before you run the Kommander installation.

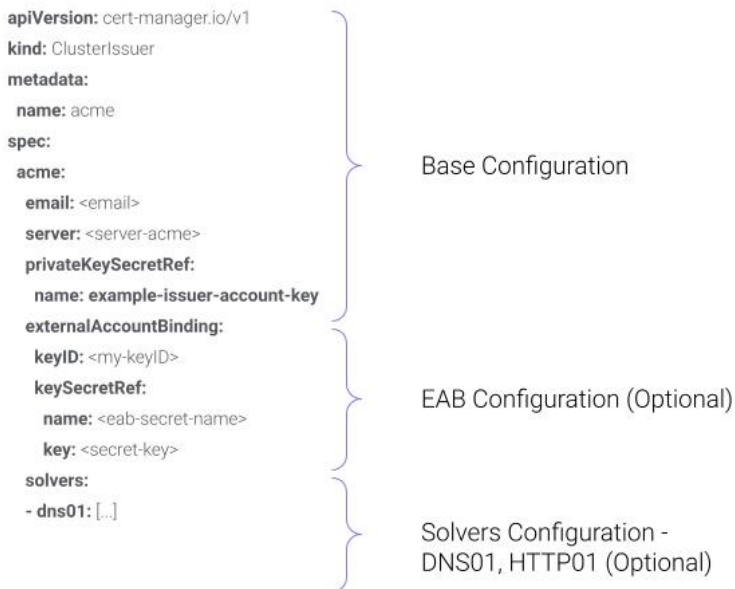
8.4.8.4.2.1 When do You Need a ClusterIssuer?

The configuration of the `ClusterIssuer` resource depends on your DKP landscape:



8.4.8.4.2.2 How do You Configure a ClusterIssuer?


The following image describes the configurable fields of a `ClusterIssuer` :



For more information on the available options, refer to the [ACME section in the cert-manager documentation](#)¹¹⁶¹.

8.4.8.4.2.3 Examples:

Refer to [Configure the Kommander Installation with a Custom Domain and Certificate](#) (see page 1435) for configuration steps and examples.

 If you need to make changes in the configuration of your domain or certificate *after* you have installed DKP, or if you want to set up a custom domain and certificate for [Attached or Managed clusters](#) (see page 81), modify the `ingress` in the `KommanderCluster` object as shown in the [Custom domains and certificates configuration](#) (see page 864) section.

8.4.8.4.2.4 Related topics:

[Why to set up a Custom Domain or Certificate?](#) (see page 1432)

[Configure the Kommander Installation with a Custom Domain and Certificate](#) (see page 1435)

[Certificate Issuer and KommanderCluster Concepts](#) (see page 1433)

¹¹⁶¹ <https://cert-manager.io/docs/configuration/acme/>

8.4.8.4.3 Configure a Custom Domain without a Custom Certificate

To configure Kommander to use a custom domain, the domain name must be provided in an installation config file. If you want to set up a custom domain and certificate, refer to [Configure your Custom Domain and Certificate](#) (see page 1435).

1. Open the *Kommander Installer Configuration File* or `<kommander.yaml>` file:
 - a. If you do not have the `<kommander.yaml>` file, [initialize the configuration file](#) (see page 1425), so you can edit it in the following steps. **WARNING:** Initialize this file only ONCE, otherwise you will overwrite previous customizations.
 - b. If you have initialized the configuration file already, open the `<kommander.yaml>` with the editor of your choice.
2. In that file, configure the custom domain for your cluster by adding this line:

```
[...]
clusterHostname: <mycluster.example.com>
[...]
```

3. This configuration can be used when installing or reconfiguring Kommander by passing it to the `dkp install kommander` command:


```
dkp install kommander --installer-config <kommander.yaml> --kubeconfig=${
{CLUSTER_NAME}.conf
```

Note: To ensure Kommander is installed on the right cluster, use the `--kubeconfig=cluster_name.conf` flag as an alternative to KUBECONFIG.

4. After the command completes, obtain the cluster ingress IP address or hostname using the following command:

```
kubectl -n kommander get svc kommander-traefik -o go-template='{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}{{ "\n"}}
```

If required, create a DNS record (for example, by using [external-dns](#) (see page 1443)) for your custom hostname that resolves to the cluster ingress load balancer hostname or IP address. If the previous command returns a hostname, you should create a CNAME DNS entry that resolves to that hostname. If the cluster ingress is an IP address, create a DNS A record.

 The domain must be resolvable from the client (your browser) and from the cluster. If you set up an `external-dns` service, it will take care of pointing the DNS record to the ingress of the cluster automatically. If you are manually creating a DNS record, you have to install Kommander first to obtain the load balancer address required for the DNS record. The [Configure a Custom Certificate](#) (see page 1435) page contains more details and examples on how and when to set up the DNS record.

8.4.8.4.3.1 Related topics:

[Why to set up a Custom Domain or Certificate?](#) (see page 1432)

[Configure your Custom Domain and Certificate](#) (see page 1435)

[Advanced Configuration: Important Concepts](#) (see page 1433)

[Advanced Configuration: ClusterIssuer](#) (see page 1440)

8.4.8.5 Verify and Troubleshoot the Domain and Certificate Customization

If you want to ensure the customization for a domain and certificate is completed, or if you want to obtain more information on the status of the customization, display the status information for the `KommanderCluster`. On the [Management cluster](#) (see page 82):

1. Inspect the modified `KommanderCluster` object:

```
kubectl describe kommandercluster -n <workspace_name> <cluster_name>
```

2. If the ingress is still being provisioned, the output looks similar to this:

```
[...]
Conditions:
  Last Transition Time: 2022-06-24T07:48:31Z
  Message:             Ingress service object was not found in the cluster
  Reason:              IngressServiceNotFound
  Status:              False
  Type:                IngressAddressReady
[...]
```

If the provisioning has been completed, the output looks similar to this:

```
[...]
Conditions:
  Last Transition Time: 2022-06-28T13:43:33Z
```

```

Message:      Ingress service address has been provisioned
Reason:      IngressServiceAddressFound
Status:      True
Type:      IngressAddressReady
Last Transition Time: 2022-06-28T13:42:24Z
Message:      Certificate is up to date and has not expired
Reason:      Ready
Status:      True
Type:      IngressCertificateReady
[...]

```

The same command also prints the actual customized values for the `KommanderCluster.Status.Ingress`. Here is an example:

```

[...]
  ingress:
    address: 172.20.255.180
    caBundle: LS0tLS1CRUdJTiBD...<output has been shortened>...DQVRFLS0tLS0K
[...]

```

8.4.8.5.1 Related topics:

[Why to set up a Custom Domain or Certificate? \(see page 1432\)](#)

[Configure the Kommander Installation with a Custom Domain and Certificate \(see page 1435\)](#)

[Certificate Issuer and KommanderCluster Concepts \(see page 1433\)](#)

[Advanced Configuration: ClusterIssuer \(see page 1440\)](#)

8.4.9 DNS Record Creation with External DNS

This section describes how to use `external-dns` to maintain your DNS records

When you set up a [custom domain \(see page 861\)](#) for your cluster, you require a DNS record that maps the configured domain or IP address to the cluster's ingress. You can either create one manually, or set up the `external-dns` service to manage your DNS record automatically.

If you choose to use `external-dns` to maintain your DNS records, the `external-dns` will take care of pointing the DNS record to the ingress of the cluster automatically.

Select one of the following options to configure the `external-dns` service for [Management \(see page 82\)/Essential \(see page 83\)](#) clusters, or for [Managed \(see page 82\)/Attached \(see page 82\)](#) clusters:

- [Configure External DNS with the CLI: Management or Essential Cluster \(see page 1443\)](#)
- [Configure External DNS with the UI: All Clusters \(see page 1445\)](#)
- [Verify your External DNS Configuration \(see page 1447\)](#)

If you choose to create a DNS record manually, finish installing the Kommander component, and then manually create a DNS record that points to the load balancer address.

8.4.9.1 Configure External DNS with the CLI: Management or Essential Cluster

This page contains information on how to configure an `external-dns` service to manage DNS records automatically in your [Management](#) (see page 82) or [Essential](#) (see page 83) Cluster.

8.4.9.1.1 Prerequisite

- Ensure you have configured a DNS zone with your cloud provider.

8.4.9.1.2 Configure External DNS and Customize Traefik

The configuration varies depending on your cloud provider.

1. Open the `kommander.yaml` file:
 - a. If you have not installed the Kommander component yet, [initialize the configuration file](#) (see page 1425), so you can edit it in the following steps.

⚠ WARNING: Initialize this file only ONCE, otherwise you will overwrite previous customizations.
 - b. If you have installed the Kommander component already, open the existing `kommander.yaml` with the editor of your choice.
2. Adjust the `app` section of your `kommander.yaml` file to include these values:

AWS Example

- i** Replace the placeholders `<...>` with your environment's information.
- i** The following example shows how to configure `external-dns` to manage DNS records in AWS Route 53 automatically.

```
apps:
  external-dns:
    enabled: true
    values: |
      aws:
        credentials:
          secretKey: <secret-key>
          accessKey: <access-key>
        region: <provider-region>
        preferCNAME: true
    policy: upsert-only
    txtPrefix: local-
    domainFilters:
```

```
- <example.com>
```

Azure Example

i Replace the placeholders `<...>` with your environment's information.

```
apps:
  external-dns:
    enabled: true
    values: |
      azure:
        cloud: AzurePublicCloud
        resourceGroup: <resource-group>
        tenantId: <tenant-id>
        subscriptionId: <your-subscription-id>
        aadClientId: <client-id>
        aadClientSecret: <client-secret>
      domainFilters:
        - <example.com>
    txtPrefix: txt-
    policy: sync
    provider: azure
```

3. In the same `app` section, adjust the `traefik` section to include the following:

```
traefik:
  enabled: true
  values: |
    service:
      annotations:
        external-dns.alpha.kubernetes.io/hostname: <mycluster.example.com>
```

4. Use the configuration file to install or update the Kommander component:

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

Refer to the [external-dns documentation](#)¹¹⁶² for more information, as well as further instructions on how to configure `external-dns` to use other DNS providers like Google Cloud DNS, CloudFlare, or on-site providers.

8.4.9.1.2.1 Next Step:

[Verify your External DNS Configuration](#) (see page 1447)

¹¹⁶² <https://artifacthub.io/packages/helm/bitnami/external-dns>

8.4.9.2 Configure External DNS with the UI: All Clusters

Enterprise

Gov Advanced

This page contains information on how to configure an `external-dns` service to manage DNS records automatically and applies to [all cluster types](#) (see page 81).

8.4.9.2.1 Prerequisite

- Ensure you have configured a DNS zone with your cloud provider.

8.4.9.2.2 Configure External DNS Using the UI

The configuration varies depending on your cloud provider.

1. Select the target workspace from the top navigation bar.
 - i** It must be the workspace that contains the cluster, for which you want to configure External DNS. In the case of the Management cluster, it would be the Management cluster workspace.
2. Select **Applications** from the sidebar menu.
3. Search for the **External DNS** application.
4. On the application card, select the *three-dot menu* > **Enable**.
5. On the **Enable Workspace Platform Application** page, select **Configuration** from the sidebar menu.
6. Copy and paste the following contents into the code editor and replace the placeholders `<...>` with your environment's information.

Here is an example configuration:

AWS Example

i Replace the placeholders `<...>` with your environment's information.

i The following example shows how to configure `external-dns` to manage DNS records in AWS Route 53 automatically.

```
aws:
  credentials:
    secretKey: <secret-key>
    accessKey: <access-key>
  region: <provider-region>
  preferCNAME: true
policy: upsert-only
txtPrefix: local-
domainFilters:
```



```
- <example.com>
```

Azure Example

i Replace the placeholders `<...>` with your environment's information.

```
azure:
  cloud: AzurePublicCloud
  resourceGroup: <resource-group>
  tenantId: <tenant-id>
  subscriptionId: <your-subscription-id>
  aadClientId: <client-id>
  aadClientSecret: <client-secret>
domainFilters:
  - <example.com>
txtPrefix: txt-
policy: sync
provider: azure
```

= Refer to [external-dns documentation](#)¹¹⁶³ for more configuration options.


8.4.9.2.3 Customize the Traefik Deployment Using the UI

= DKP deploys Traefik to all clusters by default.

1. Select the target workspace from the top navigation bar.
 - i** It must be the workspace that contains the cluster, for which you want to configure External DNS. In the case of the Management cluster, it would be the Management cluster workspace.
2. Select **Applications** from the sidebar menu.
3. Search for the *Traefik* application.
4. On the application card, select the *three-dot menu* > **Edit**.
5. Select the **Configuration** lateral tab to add a customization.
6. Copy and paste the following configuration into the code editor:
 - i** Use the **Cluster Application Configuration Override** code editor to apply a configuration per cluster.

¹¹⁶³ <https://artifacthub.io/packages/helm/bitnami/external-dns>

```
service:
  annotations:
    external-dns.alpha.kubernetes.io/hostname: <mycluster.example.com>
```

 Ensure you set up a domain per cluster, for example: `<mycluster1.example.com>` , `<mycluster2.example.com>` and `<mycluster3.example.com>` .

8.4.9.2.3.1 Next Step:

[Verify your External DNS Configuration \(see page 1447\)](#)


8.4.9.3 Verify your External DNS Configuration

This page contains commands to verify that the `external-dns` service is functioning correctly.

If the `external-dns` service is not working properly, these commands also provide aids to find the cause or identify the issue.

8.4.9.3.1 Verify the Deployment

Verify that the deployment was triggered.

1. Set the environment variable to the **Management/Essential cluster** by exporting the `kubeconfig` file in your terminal window or using the `--kubeconfig=${CLUSTER_NAME}.conf` as explained in [Provide Context for Commands with a kubeconfig File \(see page 87\)](#).
2. Verify that the `external-dns` deployment is present:
 -  Replace `<target_WORKSPACE_NAMESPACE>` in the namespace `-n` flag with the target cluster's workspace namespace.


```
kubectl get appdeployments.apps.kommander.d2iq.io -n
<target_WORKSPACE_NAMESPACE> external-dns
```

The output should look like this:

NAME	APP	AGE
external-dns	external-dns-<app_version>	36s

The CLI has triggered the deployment of the application. However, this does not mean that the application has been installed completely and successfully.

Verify that the deployment was successful.

1. Set the environment variable to the **target cluster** (where you enabled `external-dns`) by exporting the `kubeconfig` file in your terminal window or using the `--kubeconfig=$ {CLUSTER_NAME}.conf` as explained in [Provide Context for Commands with a kubeconfig File](#) (see [page 87](#)).
2. Verify that the `external-dns` deployment is ready:
 -  Replace `<target_WORKSPACE_NAMESPACE>` in the namespace `-n` flag with the target cluster's workspace namespace.


```
kubectl get deployments.apps -n <target_WORKSPACE_NAMESPACE> external-dns
```

The deployment should display a ready state:

NAME	READY	UP-TO-DATE	AVAILABLE	AGE
external-dns	1/1	1	1	42s


The CLI has deployed the application completely and successfully.

8.4.9.3.2 Examine the Cluster's Ingress

1. Set the environment variable to the **target cluster** (where you enabled `external-dns`) by exporting the `kubeconfig` file in your terminal window or using the `--kubeconfig=$ {CLUSTER_NAME}.conf` as explained in [Provide Context for Commands with a kubeconfig File](#) (see [page 87](#)).
2. Verify that the cluster's ingress contains the correct hostname annotation:
 -  Replace `<target_WORKSPACE_NAMESPACE>` in the namespace `-n` flag with the target cluster's workspace namespace.

```
kubectl get services -n <target_WORKSPACE_NAMESPACE> kommander-traefik -o yaml
```

The output looks like this:

-  Ensure that the service object contains the `external-dns.alpha.kubernetes.io/hostname: <mycluster.example.com>` annotation.


```
apiVersion: v1
kind: Service
```

```

metadata:
  annotations:
    meta.helm.sh/release-name: kommander-traefik
    meta.helm.sh/release-namespace: kommander
    external-dns.alpha.kubernetes.io/hostname: <mycluster.example.com>
  creationTimestamp: "2023-06-21T04:52:49Z"
  finalizers:
[...]
```

The `external-dns` service has been linked to the cluster correctly.


8.4.9.3.3 Verify the DNS Record

 It can take a few minutes for the `external-dns` service to create a DNS record. The delay depends on your cloud provider.

Verify that the `external-dns` service has created a DNS record.

1. Set the environment variable to the **target cluster** (where you enabled `external-dns`) by exporting the `kubeconfig` file in your terminal window or using the `--kubeconfig=${CLUSTER_NAME}.conf` as explained in [Provide Context for Commands with a kubeconfig File](#) (see [page 87](#)).
2. Access and execute the required image:

```
kubectl run -it --image=nicolaka/netshoot --rm test-dns -- /bin/bash
```

3. Use the image to check your domain and see the record:
 -  Replace `<mycluster.example.com>` with the domain you assigned to your target cluster.

```
nslookup <mycluster.example.com>
```

The output should look like this:

```

Server:      192.168.178.1
Address:     192.168.178.1#53

Non-authoritative answer:
Name:   <mycluster.example.com>
Address: 134.568.789.12
```

The `external-dns` service is working and the DNS provider recognizes the record created by the service. If the command displays an error, the configuration is failing on the end of the DNS provider.

Troubleshooting

If your deployment has not succeeded and the previous steps have not helped you identify the issue, you can also check the logs for the `external-dns` deployment:

1. Set the environment variable to the **target cluster** (where you enabled `external-dns`) by exporting the `kubeconfig` file in your terminal window or using the `--kubeconfig=${CLUSTER_NAME}.conf` as explained in [Provide Context for Commands with a kubeconfig File](#) (see [page 87](#)).
2. Verify the `external-dns` logs:
 - 📘 Replace `<target_WORKSPACE_NAMESPACE>` in the namespace `-n` flag with the target cluster's workspace namespace.

```
kubectl logs -n kommander deployment/external-dns
```

The output displays the pod's logs for the `external-dns` deployment. Here is an example:

```
...
time="2023-07-04T06:56:35Z" level=info msg="Instantiating new Kubernetes
client"
time="2023-07-04T06:56:35Z" level=info msg="Using inCluster-config based on
serviceaccount-token"
time="2023-07-04T06:56:35Z" level=info msg="Created Kubernetes client https://
10.96.0.1:443"
time="2023-07-04T06:56:35Z" level=error msg="records retrieval failed: failed
to list hosted zones:
..."
```

8.4.10 External Load Balancer

8.4.10.1 Load Balancing for External Traffic in DKP

DKP includes a load balancing solution for the [supported cloud infrastructure providers](#) (see [page 69](#)) and for pre-provisioned environments. For more information, see [Load Balancing for external traffic](#) (see [page 967](#)) in DKP.

If you want to use a **non-DKP load balancer** (for example, as an alternative to MetalLB in pre-provisioned environments), DKP supports setting up an **external load balancer**.


When enabled, the external load balancer routes incoming traffic requests to a single point of entry in your cluster. Users and services can then access the **DKP UI** through an established IP or DNS address.

 In DKP environments, the external load balancer must be configured without TLS termination.

8.4.10.2 Configure Kommander to use an External Load Balancer


To configure an external load balancer, configure a custom hostname (static IP or dynamic DNS address) and specify the target `nodePorts` for your cluster.

1. Open the **Kommander Installer Configuration File** or `kommander.yaml` file:
 - a. If you do not have the `kommander.yaml` file, [initialize the configuration file \(see page 1425\)](#), so you can edit it in the following steps. **WARNING:** Initialize this file only ONCE, otherwise you will overwrite previous customizations.
 - b. If you have initialized the configuration file already, open the `kommander.yaml` with the editor of your choice.
2. In that file, add the following line for the IP address or DNS name:

 ACME does not support the automatic creation of a certificate if you select an IP address for your `clusterHostname`.

```
[...]
clusterHostname: <mycluster.example.com OR IP_address>
[...]
```

3. Optional: If you require a custom certificate for your `clusterHostname`, see [Configure the Kommander Installation with a Custom Domain and Certificate \(see page 1435\)](#).
4. In the same **Kommander Installer Configuration File**, configure Kommander to use the `NodePort` service by adding a custom configuration under `traefik`:

 You can specify the `nodePort` entry points for the load balancer. Ensure the port is within the Kubernetes default (30 000 - 32 768). If not specified, Kommander assigns a port dynamically.

```
traefik:
  enabled: true
  values: |-
    ports:
      web:
        nodePort: 32080 #if not specified, will be assigned dynamically
      websecure:
        nodePort: 32443 #if not specified, will be assigned dynamically
```

```
service:
  type: NodePort
```

5. Use the configuration file to install Kommander (see page 1427).

8.4.10.3 Configure the External Load Balancer to Target the Specified Ports

The `traefik` service of the Kommander component now actively listens on the pod IPs, and is accessible through the specified ports on every node.

Configure the load balancer targets to include every worker node address (DNS name or IP address) and node port combination by following this format:

```
<node1>:<nodePort_web> # for example, my.node1.internal:32080
<node2>:<nodePort_web>
<node3>:<nodePort_web>
[...]
<node1>:<nodePort_websecure> # for example, my.node1.internal:32443
<node2>:<nodePort_websecure>
<node3>:<nodePort_websecure>
[...]
```



The exact configuration depends on your load balancer provider.

8.4.11 Installing Kommander with an HTTP Proxy

Configure HTTP proxy for the Kommander clusters

Kommander supports environments that connect through an HTTP/HTTPS proxy, when access to the Internet is restricted. Use the information in this section to configure the Kommander component of DKP correctly.

In these environments, you must configure Kommander to use the HTTP/HTTPS proxy. In turn, Kommander configures all platform services to use the HTTP/HTTPS proxy.



Kommander follows a common convention for using an HTTP proxy server. The convention is based on three environment variables, and is supported by many, though not all, applications.

- `HTTP_PROXY` : the HTTP proxy server address
- `HTTPS_PROXY` : the HTTPS proxy server address

- `NO_PROXY` : a list of IPs and domain names that are not subject to proxy settings

8.4.11.1 Prerequisites

In the examples below:

1. The `curl` command-line tool is available on the host.
2. The proxy server address is `http://proxy.company.com:3128`.
3. The HTTP and HTTPS proxy server addresses use the `http` scheme.
4. The proxy server can reach `www.google.com` using HTTP or HTTPS.

Verify the cluster nodes can access the Internet through the proxy server. On each cluster node, run:

```
curl --proxy http://proxy.company.com:3128 --head http://www.google.com
curl --proxy http://proxy.company.com:3128 --head https://www.google.com
```

If the proxy is working for HTTP and HTTPS, respectively, the `curl` command returns a `200 OK HTTP` response.

8.4.11.2 Enable Gatekeeper

Gatekeeper acts as a [Kubernetes mutating webhook](#)¹¹⁶⁴. You can use this to mutate the Pod resources with `HTTP_PROXY`, `HTTPS_PROXY` and `NO_PROXY` environment variables.

1. Create (if necessary) or update the Kommander installation configuration file. If one does not already exist, then create it using the following commands:

```
dkp install kommander --init > kommander.yaml
```

2. Append this `apps` section to the `kommander.yaml` file with the following values to enable Gatekeeper and configure it to add HTTP proxy settings to the pods.

NOTE: Only pods created after applying this setting will be mutated. Also, this will only affect pods in the namespace with the `"gatekeeper.d2iq.com/mutate=pod-proxy"` label.

```
apps:
  gatekeeper:
    values: |
      disableMutation: false
    mutations:
```

¹¹⁶⁴ <https://kubernetes.io/docs/reference/access-authn-authz/admission-controllers/#mutatingadmissionwebhook>


```

enablePodProxy: true
podProxySettings:
  noProxy:
"127.0.0.1,192.168.0.0/16,10.0.0.0/16,10.96.0.0/12,169.254.169.254,169.254.0.0/
24,localhost,kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.de
fault.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.c
luster.local,.svc.cluster.local.,kubecost-prometheus-server.kommander,logging-
operator-logging-fluentd.kommander.svc.cluster.local,elb.amazonaws.com"
  httpProxy: "http://proxy.company.com:3128"
  httpsProxy: "http://proxy.company.com:3128"
  excludeNamespacesFromProxy: []
  namespaceSelectorForProxy:
    "gatekeeper.d2iq.com/mutate": "pod-proxy"

```

3. Create the `kommander` and `kommander-flux` namespaces, or the namespace where Kommander will be installed. Label the namespaces to activate the Gatekeeper mutation on them:

```

kubectl create namespace kommander
kubectl label namespace kommander gatekeeper.d2iq.com/mutate=pod-proxy

kubectl create namespace kommander-flux
kubectl label namespace kommander-flux gatekeeper.d2iq.com/mutate=pod-proxy

```

8.4.11.3 Create Gatekeeper ConfigMap in the Kommander Namespace

To configure Gatekeeper so that these environment variables are mutated in the pods, create the following `gatekeeper-overrides` ConfigMap in the `kommander` Workspace you created in a previous step:

```
export NAMESPACE=kommander
```

```

cat << EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: gatekeeper-overrides
  namespace: ${NAMESPACE}
data:
  values.yaml: |
    ---
    # enable mutations
    disableMutation: false
    mutations:
      enablePodProxy: true
      podProxySettings:

```

```

    noProxy:
"127.0.0.1,192.168.0.0/16,10.0.0.0/16,10.96.0.0/12,169.254.169.254,169.254.0.0/24,localhost,kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluster.local,.svc.cluster.local.,kubecost-prometheus-server.kommander,logging-operator-logging-fluentd.kommander.svc.cluster.local,elb.amazonaws.com"
    httpProxy: "http://proxy.company.com:3128"
    httpsProxy: "http://proxy.company.com:3128"
    excludeNamespacesFromProxy: []
    namespaceSelectorForProxy:
"gatekeeper.d2iq.com/mutate": "pod-proxy"
EOF

```

Set the `httpProxy` and `httpsProxy` environment variables to the address of the HTTP and HTTPS proxy servers, respectively. Set the `noProxy` environment variable to the addresses that should be accessed directly, not through the proxy.

Performing this step before installing Kommander allows the Flux components to respect the proxy configuration in this ConfigMap.

8.4.11.4 HTTP Proxy Configuration Considerations

To ensure that core components work correctly, always add these addresses to the `noProxy`:

- Loopback addresses (`127.0.0.1` and `localhost`)
- Kubernetes API Server addresses
- Kubernetes Pod IPs (for example, `192.168.0.0/16`). This comes from two places:
 - Calico pod CIDR - Defaults to `192.168.0.0/16`
 - The `podSubnet` is configured in CAPI objects and needs to match above Calico's - Defaults to `192.168.0.0/16` (same as above)
- Kubernetes Service addresses (for example, `10.96.0.0/12` , `kubernetes` , `kubernetes.default` , `kubernetes.default.svc` , `kubernetes.default.svc.cluster` , `kubernetes.default.svc.cluster.local` , `.svc` , `.svc.cluster` , `.svc.cluster.local` , `.svc.cluster.local` .)
- Auto-IP addresses `169.254.169.254,169.254.0.0/24`

In addition to the values above, the following settings are needed when installing on AWS:

- The default VPC CIDR range of `10.0.0.0/16`
- `kube-apiserver` internal/external ELB address



- The `NO_PROXY` variable contains the Kubernetes Services CIDR. This example uses the default CIDR, `10.96.0.0/12`. If your cluster's CIDR is different, update the value in the `NO_PROXY` field.
- Based on the order in which the Gatekeeper Deployment is Ready (in relation to other Deployments), not all the core services are guaranteed to be mutated with the proxy environment variables. Only the user deployed workloads are guaranteed to be mutated with the proxy environment variables. If you need a core service to be mutated with your proxy environment variables, you can restart the AppDeployment for that core service.

8.4.11.5 Install Kommander

Kommander installs with the DKP CLI. Install Kommander using the configuration files and ConfigMap from previous steps:

NOTE: To ensure Kommander is installed on the workload cluster, use the `--kubecfg=cluster_name.conf` flag:

```
dkp install kommander --installer-config kommander.yaml
```

8.4.11.6 Configure Workspace or Project

Configure the Workspace or Project in which you want to use the proxy. To have Gatekeeper mutate the manifests, create the `Workspace` (or `Project`) with the following label:

```
labels:
  gatekeeper.d2iq.com/mutate: "pod-proxy"
```

This can be done when creating the Workspace (or Project) from the UI OR by running the following command from the CLI after creating the namespace:

```
kubectl label namespace <NAMESPACE> "gatekeeper.d2iq.com/mutate=pod-proxy"
```

8.4.11.7 Configure HTTP Proxy in Attached Clusters

To ensure that Gatekeeper is deployed before everything else in the attached clusters that you want to configure with proxy configuration, you must manually create the exact Namespace of the Workspace in which the cluster is going to be attached, *before* attaching the cluster:

Execute the following command in the attached cluster before attaching it to the host cluster:

```
kubectl create namespace <NAMESPACE>
```

Then, to configure the pods in this namespace to use proxy configuration, you must label the Workspace with `gatekeeper.d2iq.com/mutate=pod-proxy` when creating it so that Gatekeeper deploys a `validatingwebhook` to mutate the pods with proxy configuration.

```
kubectl label namespace <NAMESPACE> "gatekeeper.d2iq.com/mutate=pod-proxy"
```

8.4.11.8 Create Gatekeeper ConfigMap in the Workspace Namespace

To configure Gatekeeper so that these environment variables are mutated in the pods, create the following `gatekeeper-overrides` ConfigMap in the Workspace Namespace:

```
export NAMESPACE=<NAMESPACE>
```

```
cat << EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: gatekeeper-overrides
  namespace: ${NAMESPACE}
data:
  values.yaml: |
    ---
    # enable mutations
    disableMutation: false
    mutations:
      enablePodProxy: true
      podProxySettings:
        noProxy:
"127.0.0.1,192.168.0.0/16,10.0.0.0/16,10.96.0.0/12,169.254.169.254,169.254.0.0/24,localhost,kubernetes,kubernetes.default,kubernetes.default.svc,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluster.local,.svc.cluster.local.,kubecost-prometheus-server.kommander,logging-operator-logging-fluentd.kommander.svc.cluster.local,elb.amazonaws.com"
        httpProxy: "http://proxy.company.com:3128"
        httpsProxy: "http://proxy.company.com:3128"
      excludeNamespacesFromProxy: []
      namespaceSelectorForProxy:
        "gatekeeper.d2iq.com/mutate": "pod-proxy"
EOF
```

Set the `httpProxy` and `httpsProxy` environment variables to the address of the HTTP and HTTPS proxy servers, respectively. Set the `noProxy` environment variable to the addresses that should be

accessed directly, not through the proxy. The list of the recommended settings is in the section *HTTP Proxy Configuration Considerations* above.


8.4.11.9 Configure Your Applications

In a default installation with `gatekeeper` enabled, you can have proxy environment variables applied to all your pods automatically by adding the following label to your namespace:

```
"gatekeeper.d2iq.com/mutate": "pod-proxy"
```

No further manual changes are required.

8.4.11.10 Manually Configure Your Application

 If Gatekeeper is not installed, and you need to use an HTTP proxy, you must manually configure your applications.

Some applications follow the convention of `HTTP_PROXY`, `HTTPS_PROXY`, and `NO_PROXY` environment variables.

In this example, the environment variables are set for a container in a Pod:

See [Define Environment Variables for a Container](#)¹¹⁶⁵ for more details.

8.4.11.10.1 Next Steps:

Now select your environment, and finish your Kommander Installation using one of the following:

[Install Kommander in an Air-gapped Environment](#) (see page 1400)

[Install Kommander in a Non-air-gapped Environment](#) (see page 1405)

[Install Kommander in a Small Environment](#) (see page 1419)

8.4.12 Enable DKP Catalog Applications after Installing DKP

Enterprise


Gov Advanced

DKP supports configuring DKP Catalog applications for environments with an Enterprise license. There are two possible ways of enabling the Enterprise catalog:

¹¹⁶⁵ <https://kubernetes.io/docs/tasks/inject-data-application/define-environment-variable-container/#define-an-environment-variable-for-a-container>



Enable an Enterprise catalog during the installation of DKP.	See Install Kommander in an Air-gapped Environment (see page 1400), or Install Kommander in a Non-air-gapped Environment (see page 1405) depending on your environment.
Enable an Enterprise catalog after installing DKP.	Follow the instructions on this page.

8.4.12.1 Configure a Default Enterprise Catalog after Installing DKP

1. Locate the **Kommander Installer Configuration file** you are using for your current deployment. This file is stored locally on your computer.
 -  The default file name is `kommander.yaml`, but your file's name could be different.
2. Open that file and go into edit mode.
3. Enable the default catalog repository by adding the following values to your existing and already configured **Kommander Installer Configuration file**:

```
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      gitRepositorySpec:
        url: https://github.com/mesosphere/dkp-catalog-applications
        ref:
          tag: v2.6.1
...

```

4. Reconfigure DKP by reinstalling the Kommander component. In the following example, replace `kommander.yaml` with the name of your **Kommander Installer Configuration file**.
 -  Ensure you are using the correct name for the **Kommander Installer Configuration file** to maintain your clusters settings. Installing a different `kommander.yaml` than the one your environment is using overwrites all of your previous configurations and customizations. 

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=${CLUSTER_NAME}.conf
```

Tips and recommendations

- The `--kubeconfig=${CLUSTER_NAME}.conf` flag ensures that you **install Kommander on the correct cluster**. For alternatives, see [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- Applications can take longer to deploy, and time out the installation. Add the `--wait-timeout <time to wait>` flag and specify a period of time (for example, `1h`) to allocate more time to the deployment of applications.
- If the Kommander installation fails, or you wish to reconfigure applications, rerun the `install` command to retry.

8.4.12.2 DKP Catalog Application labels

The following section describes each label:

Label	Description
<code>kommander.d2iq.io/project-default-catalog-repository</code>	Indicates this acts as a Catalog Repository in all projects
<code>kommander.d2iq.io/workspace-default-catalog-repository</code>	Indicates this acts as a Catalog Repository in all workspaces
<code>kommander.d2iq.io/gitapps-gitrepository-type</code>	Indicates this Catalog Repository (and all its Applications) are certified to run on DKP

8.4.12.3 Next Step:

[Custom Domains and Certificates for Management and Essential Clusters](#) (see page 1431)

9 Additional Konvoy Configurations

When installing DKP for a project, line-of-business, or enterprise, the first step is to determine the infrastructure on which you want to deploy.

The infrastructure you select then determines the specific requirements for a successful installation. [Day 1 - Basic Install](#) (see page 147) section gave you a recommended install. For further customization or advanced questions, this Day 1 section of the documentation will provide answers.

If you have decided to uninstall DKP, refer to the same infrastructure documentation you have selected in this Day 1 section for specific steps to take down your cluster.

The different Infrastructures and components for further configuration are listed below:

9.1 Section Contents

- [Konvoy Image Builder](#) (see page 1461)
- [FIPS 140-2 Compliance](#) (see page 1526)
- [Registry Mirror Tools](#) (see page 1533)
- [Air-gapped Seed the Registry](#) (see page 1537)
- [Configure the Control Plane](#) (see page 1538)
- [Update Cluster Nodepools](#) (see page 1545)
- [Verify your Cluster and DKP Installation](#) (see page 1547)
- [GPU for Konvoy](#) (see page 1549)
- [Delete a DKP Cluster with One Command](#) (see page 1549)

9.2 Konvoy Image Builder

Konvoy Image Builder (KIB) is a complete solution for building Cluster API compliant images. The goal of Konvoy Image Builder is to produce a common operating surface to run Konvoy across heterogeneous infrastructure. KIB relies on ansible to install software, configure, and sanitize systems for running Konvoy. Packer is used to build images for cloud environments. Goss is used to validate systems are capable of running Konvoy.

This section describes how to use KIB to create a [Cluster API](#)¹¹⁶⁶ compliant machine images. Machine images contain [configuration information](#)¹¹⁶⁷ and software to create a specific, pre-configured, operating environment. For example, you can create an image of your current computer system settings and software. The machine image can then be replicated and distributed, creating your computer system for other users. KIB uses [variable overrides](#)¹¹⁶⁸ to specify base image and container images to use in your new machine image. The variable overrides files for Nvidia and FIPS can be ignored unless adding an overlay feature.

¹¹⁶⁶ <https://cluster-api.sigs.k8s.io/>

¹¹⁶⁷ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images>

¹¹⁶⁸ [https://d2iq.atlassian.net/wiki/pages/createpage.action?](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Use+Override+Files+with+Konvoy+Image+Builder)

[fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Use+Override+Files+with+Konvoy+Image+Builder](https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Use+Override+Files+with+Konvoy+Image+Builder)

9.2.1 How KIB Works

You can use KIB to build machine images, but first you need to be aware of the default behaviors of KIB. Stated very simply, KIB installs `kubeadm` and the other basic components you need, so that when the machine boots for the first time, it becomes a Kubernetes control plane or worker node, and then can form or join a cluster.

KIB does this by booting a computer using a stripped-down base image, like an AMI, and then runs a series of steps to install all of the components that DKP needs. When the installation completes, KIB takes a snapshot or backup of that machine image and saves it. This becomes the image or AMI, and so on, that you use when building the cluster.

9.2.1.1 Using Override files

You can [use override files](#)¹¹⁶⁹ to customize some of the components installed on this machine image. (The KIB base override files are located in [this Github repository](#)¹¹⁷⁰.) For example, you could tell KIB to install the FIPS versions of the Kubernetes components. You could also add `offline` overrides, used for air-gapped environments, that will upload the necessary OS packages and other files to the image, so that KIB doesn't try to look for those using the Internet. Or you could use an override file to provide similar configurations for using a GPU-enabled cluster.

You can only change certain, predefined parameters using this mechanism. The set of parameters you can change depends on which of the supported infrastructure providers you are using. For example, for vSphere, there are many settings that you use to tell the KIB how to talk to your vCenter implementation, where to store the images, which credentials to use, and so on.

9.2.1.2 Using HTTP/S Proxies

In some networked environments, the machines used for building images can reach the Internet, but only through an HTTP/S proxy. For DKP to operate in these networks, you need a way to [specify what proxies to use](#)¹¹⁷¹. You can use an HTTP proxy override file to specify that proxy. When KIB is trying to install a particular OS package, it uses that proxy to reach the Internet to download that package.



The proxy setting specified here is NOT “baked into” the image - it is only used while the image is being built. The settings are removed before the image is finalized.

While it might seem logical to include the proxy information in the image, the reality is that many companies have multiple proxies - one perhaps for each geographical region, or maybe even a proxy per Datacenter or office. All network traffic to the Internet goes through the proxy. If you were in Germany, you would probably

¹¹⁶⁹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Use+Override+Files+with+Konvoy+Image+Builder>

¹¹⁷⁰ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

¹¹⁷¹ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Configuring+an+HTTP%2FHTTPS+Proxy>

not want to send all your traffic to a U.S.-based proxy. Doing that would slow traffic down and consume too many network resources. If you baked the proxy settings into the image, you would have to create a separate image for each specific region. It makes more sense to create an image with no proxy included, but remember that you still need a proxy to access the Internet. Thus, when creating the cluster, (and also when installing the Kommander component of DKP), you must specify the correct proxy settings for the network environment into which you are installing the cluster. You will use the same base image for *that* cluster as one installed in an environment with different proxy settings.

9.2.2 Prerequisites

Before you begin, you must ensure your versions of KIB and DKP are compatible:

- Download the Konvoy Image Builder bundle from the KIB Version column of the chart below for your version of DKP prefixed with `konvoy-image-bundle` for your Operating System.
- Check the <https://d2iq.atlassian.net/wiki/pages/createpage.action?spaceKey=DENT&title=%282.6%29%20Supported%20Infrastructure%20Operating%20Systems&linkCreation=true&fromPageId=273318360>¹¹⁷²
- Check the [Supported Kubernetes Version](#)¹¹⁷³ for your Provider.
- An x86_64-based Linux or MacOS machine
- A Container engine installed:
 - Version [Docker®](#)¹¹⁷⁴ container engine version 18.09.2 or higher installed for Linux or MacOS - On macOS, Docker runs in a virtual machine which needs configured with at least 8 GB of memory.
 - Version 4.0 of [Podman](#)¹¹⁷⁵ or higher for Linux.
- For air-gapped only - a [local registry](#)¹¹⁷⁶

9.2.2.1 Compatible DKP to KIB Versions

Along with the KIB Bundle, we publish a file containing checksums for the bundle files. The recommendation for using these checksums is to verify the integrity of the downloads.

- On the corresponding link page, download the package prefixed with `konvoy-image-bundle` for your OS.

¹¹⁷² <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Supported+Infrastructure+Operating+Systems>

¹¹⁷³ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Upgrade+Compatibility+Tables>

¹¹⁷⁴ <https://docs.docker.com/get-docker/>

¹¹⁷⁵ <https://podman.io/getting-started/installation>

¹¹⁷⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Local+Registry+Tools>

¹¹⁷⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Local+Registry+Tools>

¹¹⁷⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Local+Registry+Tools>

¹¹⁷⁶ <https://d2iq.atlassian.net/wiki/pages/createpage.action?fromPageId=273318360&linkCreation=true&spaceKey=DENT&title=%282.6%29+Local+Registry+Tools>

DKP Version	KIB Version (Click for bundle download)
v2.6.2	2.5.5 ¹¹⁷⁷
v2.6.1	2.5.4 ¹¹⁷⁸
v2.6.0	2.5.3 ¹¹⁷⁹
v2.5.2	2.2.13 ¹¹⁸⁰
v2.5.1	2.2.8 ¹¹⁸¹
v2.5.0	2.2.6 ¹¹⁸²
v2.4.0	1.24.5 ¹¹⁸³



KIB will run and print out the name of the created image for your infrastructure provider as shown on specific provider KIB pages below. Use this name when creating a Kubernetes cluster.

9.2.2.2 Next Steps:

Either return to [Basic Install \(see page 147\)](#) or [Custom Install \(see page 1039\)](#) instructions, or for more KIB specific provider information you can continue to the provider link below for additional information:

- [KIB with AWS \(see page 1465\)](#)
- [KIB for Azure \(see page 1477\)](#)
- [KIB with GCP \(see page 1481\)](#)
- [KIB for GPU \(see page 1485\)](#)
- [KIB with vSphere \(see page 1487\)](#)
- [KIB with Pre-provisioned Environments \(see page 1493\)](#)

1177 <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.5.5>

1178 <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.5.4>

1179 <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.5.3>

1180 <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.2.13>

1181 <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.2.8>

1182 <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.2.6>

1183 <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.24.5>

- [Konvoy Image Builder CLI](#) (see page 1494)
- [Use Override Files with Konvoy Image Builder](#) (see page 1513)

9.2.3 KIB with AWS

The following section describes how to use Konvoy Image Builder (KIB) with Amazon Web Services (AWS). A minimal set of permissions from the [AWS Image Builder Book](#)¹¹⁸⁴ should be met before beginning.

Section Contents

- [Minimal IAM Permissions for KIB](#) (see page 1465)
- [AWS Image Integrated with DKP CLI](#) (see page 1468)
- [Create a Custom AMI](#) (see page 1469)
- [AWS Air-gapped AMI](#) (see page 1474)
- [Add Custom Tags to Image](#) (see page 1475)
- [KIB for EKS](#) (see page 1476)

9.2.3.1 Minimal IAM Permissions for KIB

Configure IAM Prerequisites before building an AWS Image

This section guides you in creating and using a minimally-scoped policy to create an Image for an AWS account using Konvoy Image Builder.

9.2.3.1.1 Prerequisites

Before applying the IAM Policies, verify the following:

- You have a valid AWS account with [credentials configured](#)¹¹⁸⁵ that can manage CloudFormation Stacks, IAM Policies, IAM Roles, and IAM Instance Profiles.
- The [AWS CLI utility](#)¹¹⁸⁶ is installed.

9.2.3.1.2 Minimal Permissions

The following is an AWS CloudFormation stack that creates the minimal policy to run KIB in AWS.

1. Copy the following contents into a file:

```
AWSTemplateFormatVersion: 2010-09-09
Resources:
  AWSIAMInstanceKIBUser:
    Properties:
      InstanceProfileName: KIBUserInstnaceProfile
```

¹¹⁸⁴ <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>

¹¹⁸⁵ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

¹¹⁸⁶ <https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html>

```
Roles:
  - Ref: KIBUserRole
Type: AWS::IAM::InstanceProfile
AWSIAMManagedPolicyKIBPolicy:
Properties:
  Description: Minimal policy to run KIB in AWS
  ManagedPolicyName: kib-policy
  PolicyDocument:
    Statement:
      - Action:
          - ec2:AssociateRouteTable
          - ec2:AssociateRouteTable
          - ec2:AttachInternetGateway
          - ec2:AttachVolume
          - ec2:AuthorizeSecurityGroupIngress
          - ec2:CreateImage
          - ec2:CreateInternetGateway
          - ec2:CreateKeyPair
          - ec2:CreateRoute
          - ec2:CreateRouteTable
          - ec2:CreateSecurityGroup
          - ec2:CreateSubnet
          - ec2:CreateTags
          - ec2:CreateVolume
          - ec2:CreateVpc
          - ec2>DeleteInternetGateway
          - ec2>DeleteKeyPair
          - ec2>DeleteRouteTable
          - ec2>DeleteSecurityGroup
          - ec2>DeleteSnapshot
          - ec2>DeleteSubnet
          - ec2>DeleteVolume
          - ec2>DeleteVpc
          - ec2:DeregisterImage
          - ec2:DescribeAccountAttributes
          - ec2:DescribeImages
          - ec2:DescribeInstances
          - ec2:DescribeInternetGateways
          - ec2:DescribeKeyPairs
          - ec2:DescribeNetworkAcls
          - ec2:DescribeNetworkInterfaces
          - ec2:DescribeRegions
          - ec2:DescribeRouteTables
          - ec2:DescribeSecurityGroups
          - ec2:DescribeSubnets
          - ec2:DescribeVolume
          - ec2:DescribeVpcAttribute
          - ec2:DescribeVpcClassicLink
          - ec2:DescribeVpcClassicLinkDnsSupport
          - ec2:DescribeVpcs
          - ec2:DetachInternetGateway
          - ec2:DetachVolume
```

```

    - ec2:DisassociateRouteTable
    - ec2:ModifyImageAttribute
    - ec2:ModifySnapshotAttribute
    - ec2:ModifySubnetAttribute
    - ec2:ModifyVpcAttribute
    - ec2:RegisterImage
    - ec2:RevokeSecurityGroupEgress
    - ec2:RunInstances
    - ec2:StopInstances
    - ec2:TerminateInstances
  Effect: Allow
  Resource:
    - '*'
  Version: 2012-10-17
  Roles:
    - Ref: KIBUserRole
  Type: AWS::IAM::ManagedPolicy
  Version: 2012-10-17
KIBUserRole:
  Properties:
    AssumeRolePolicyDocument:
      Statement:
        - Action:
            - sts:AssumeRole
          Effect: Allow
          Principal:
            Service:
              - ec2.amazonaws.com
        - Action:
            - sts:AssumeRole
          Effect: Allow
          Principal:
            AWS: arn:aws:iam::MYAWSACCOUNTID:root
      Version: 2012-10-17
    RoleName: kib-user-role
  Type: AWS::IAM::Role

```

2. Replace the following with the correct values:

- MYFILENAME.yaml - give your file a meaningful name.
- MYSTACKNAME - give your cloudformation stack a meaningful name.


3. Run the following command to create the stack:

```
aws cloudformation create-stack --template-body=file://MYFILENAME.yaml --stack-name=MYSTACKNAME --capabilities CAPABILITY_NAMED_IAM
```

9.2.3.2 AWS Image Integrated with DKP CLI

This procedure describes how to create a [Cluster API](#)¹¹⁸⁷ compliant Amazon Machine Image (AMI).

A customized image requires the [Konvoy Image Builder](#) (see page 1464) tool to be [downloaded](#) (see page 78) and use [variable overrides](#) (see page 1513) to specify the base image and container images to use in your new AMI. To create a custom AMI and take advantage of enhanced cluster operations, explore the [KIB with AWS](#) (see page 1468) topics for more options.

 In previous DKP releases, AMI images provided by the upstream CAPA project would be used if you did not specify an AMI. However, the upstream images are not recommended for production and may not always be available. Therefore, DKP now requires you to specify an AMI when creating a cluster. To create an AMI, use [Konvoy Image Builder](#) (see page 1469). A customized image requires the [Konvoy Image Builder](#) (see page 1464) tool to be [downloaded](#) (see page 78) where you can use [variable overrides](#) (see page 1513) to specify the base image and container images for use in your new [custom AMI](#) (see page 1469).

9.2.3.2.1 Prerequisites

Before you begin, you must:

- Check the [Supported Infrastructure Operating Systems](#) (see page 69)
- Check the [Supported Kubernetes Version](#) (see page 1552) for your Provider.
- Create a working `Docker` or other Registry setup.
- Ensure you have met the minimal set of permissions from the [AWS Image Builder Book](#)¹¹⁸⁸.

9.2.3.2.2 Build the Image

Depending on which [version of DKP](#) (see page 1464) you are running, steps and flags will be different.

9.2.3.2.2.1 Execute the following to begin image creation:

Run the `dkp-image-builder create` command to build and validate the image.

```
dkp-image-builder create image aws/centos-79.yaml
```

By default it builds in the `us-west-2` region. to specify another region set the `--region` flag:

¹¹⁸⁷ <https://cluster-api.sigs.k8s.io/>

¹¹⁸⁸ <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>

```
dkp-image-builder create image aws --region us-east-1 images/ami/centos-79.yaml
```

Once DKP provisions the image successfully, the `ami_id` is printed and written to the `manifest.json` file. This file has an `artifact_id` field whose value provides the name of the AMI ID as shown in the example below:

```
{
  "name": "rhel-7.9-fips",
  "builder_type": "amazon-eks",
  "build_time": 1659486130,
  "files": null,
  "artifact_id": "us-west-2:ami-0f2ef742482e1b829",
  "packer_run_uuid": "0ca500d9-a5f0-815c-6f12-aceb4d46645b",
  "custom_data": {
    "containerd_version": "",
    "distribution": "RHEL",
    "distribution_version": "7.9",
    "kubernetes_cni_version": "",
    "kubernetes_version": "1.24.5+fips.0"
  }
}
```

9.2.3.3 Create a Custom AMI

9.2.3.3.1 Learn how to build a custom AMI for use with DKP

This procedure describes how to use the [Konvoy Image Builder \(see page 1464\)](#) (KIB) to create a [Cluster API](#)¹¹⁸⁹ compliant Amazon Machine Image (AMI). KIB uses [variable overrides \(see page 1513\)](#) to specify base image and container images to use in your new AMI.



The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to create a custom AMI to take advantage of enhanced cluster operations. Explore the [KIB with AWS \(see page 1468\)](#) topics for more options.

9.2.3.3.2 Prerequisites

Before you begin, you must:


- Download the [KIB \(see page 0\)](#) bundle for your version of DKP prefixed with `konvoy-image-bundle` for your OS.

¹¹⁸⁹ <https://cluster-api.sigs.k8s.io/>

- Check the [Supported Infrastructure Operating Systems](#) (see page 69)
- Check the [Supported Kubernetes Version](#) (see page 1552) for your Provider.
- Create a working registry:
 - Version 4.0 of [Podman](#)¹¹⁹⁰ or higher for Linux. Host requirements found here: <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>
 - Version 20.10.0 of [Docker](#)¹¹⁹¹ or higher for Linux or MacOS
- Ensure you have met the minimal set of permissions from the [AWS Image Builder Book](#)¹¹⁹².
- A [Minimal IAM Permissions for KIB](#) (see page 1468) to create an Image for an AWS account using Konvoy Image Builder.

9.2.3.3.3 Extract KIB Bundle

Extract the bundle and `cd` into the extracted `konvoy-image-bundle- $\$$ VERSION` folder. The bundled version of `konvoy-image` contains an embedded `docker` image that contains all the requirements for building.

 The `konvoy-image` binary and all supporting folders are also extracted. When run, `konvoy-image` bind mounts the current working directory (`$\${$ PWD}`) into the container to be used.

- Set environment variables for [AWS access](#)¹¹⁹³. The following variables must be set using your credentials including [required IAM](#) (see page 1468):

```
export AWS_ACCESS_KEY_ID
export AWS_SECRET_ACCESS_KEY
export AWS_DEFAULT_REGION
```

- If you have an [override file](#) (see page 1513) to configure specific attributes of your AMI file, add it.

AMI images contain configuration information and software to create a specific, pre-configured, operating environment. For example, you can create an AMI image of your current computer system settings and software. The AMI image can then be replicated and distributed, creating your computer system for other users.

¹¹⁹⁰ <https://podman.io/getting-started/installation>

¹¹⁹¹ <https://www.docker.com/>

¹¹⁹² <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>

¹¹⁹³ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-envvars.html>

9.2.3.3.4 Build the Image

Depending on which [version of DKP](#) (see [page 1464](#)) you are running, steps and flags will be different. To deploy in a region where CAPI images are not provided, you need to use KIB to create your own image for the region. For a list of supported AWS regions, refer to the [Published AMI](#)¹¹⁹⁴ information from AWS.

9.2.3.3.4.1 Execute the following to begin image creation:

Run the `konvoy-image` command to build and validate the image.

```
konvoy-image build aws images/ami/centos-79.yaml
```

By default it builds in the `us-west-2` region. to specify another region set the `--region` flag:

```
konvoy-image build aws --region us-east-1 images/ami/centos-79.yaml
```

ⓘ Ensure you have named the correct YAML file for your OS in the `konvoy-image build` command.

After KIB provisions the image successfully, the `ami` id is printed and written to the `manifest.json` file. This file has an `artifact_id` field whose value provides the name of the AMI ID as shown in the example below:

```
{
  "name": "rhel-7.9-fips",
  "builder_type": "amazon-ebs",
  "build_time": 1659486130,
  "files": null,
  "artifact_id": "us-west-2:ami-0f2ef742482e1b829",
  "packer_run_uuid": "0ca500d9-a5f0-815c-6f12-aceb4d46645b",
  "custom_data": {
    "containerd_version": "",
    "distribution": "RHEL",
    "distribution_version": "7.9",
    "kubernetes_cni_version": "",
    "kubernetes_version": "1.26.14+fips.0"
  }
}
```

¹¹⁹⁴ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

To use a custom AMI when creating your cluster, you must create that AMI using KIB first. Then perform the export and name the custom AMI for use in the command `dkp create cluster`

```
export AWS_AMI_ID=ami-<ami-id-here>
```

9.2.3.3.5 Related Information

For information on related topics or procedures, refer to the following:

- [Creating GPU enabled on-premises configurations \(see page 1113\)](#)

9.2.3.3.6 Use Custom AMI to Build Cluster

9.2.3.3.6.1 Launch a DKP Cluster with a Custom AMI

To use the built `ami` with DKP, specify it with the `--ami` flag when calling cluster create.

```
dkp create cluster aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--ami AMI_ID
```

9.2.3.3.6.2 Launch a DKP Cluster with Custom AMI Lookup

By default, the `konvoy-image` command in KIB will name the AMI in such a way that `dkp` can discover the latest AMI for a base OS and Kubernetes version. To create a cluster that will use the latest AMI, specify the `--ami-format`, `--ami-base-os` and `--ami-owner` flags:

```
dkp create cluster aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--ami-owner AWS_ACCOUNT_ID \
--ami-base-os ubuntu-20.04 \
--ami-format 'konvoy-ami-{{.BaseOS}}-?{{.K8sVersion}}-*
```

9.2.3.3.6.3 Using Custom Source AMIs

When using KIB for building machine images to Amazon, the default source AMIs that we provide are modeled by looking up an AMI based on the owner. Then we apply a filter for that operating system and version.

You can view an example of that with the provided `centos-79.yaml` snippet below:

```
download_images: true

packer:
  ami_filter_name: "CentOS Linux 7"
  ami_filter_owners: "125523088429"
  distribution: "CentOS"
  distribution_version: "7.9"
  source_ami: ""
  ssh_username: "centos"
  root_device_name: "/dev/sda1"
  ...
```

At times, a particular upstream AMI may not be available in your region or something could be renamed. Other times you want to provide a custom AMI. If this is the case, you will want to edit or create your own YAML file that looks up based on the `source_ami` field. For example, you can select images that are otherwise deprecated.

Once you select the source AMI that you want, you can declare that when running your build command:

```
konvoy-image build aws path/to/ami/centos-79.yaml --source-ami ami-0123456789
```

Alternatively, if you want to add it to your YAML file, or make your own file, you can do that as well. You add that AMI ID into the `source_ami` in the YAML file:

```
download_images: true

packer:
  ami_filter_name: ""
  ami_filter_owners: ""
  distribution: "CentOS"
  distribution_version: "7.9"
  source_ami: "ami-123456789"
  ssh_username: "centos"
  root_device_name: "/dev/sda1"
  ...
```

When you're done selecting your `source_ami`, you can build your KIB image as you would normally:


```
konvoy-image build aws path/to/ami/centos-79.yaml
```

9.2.3.4 AWS Air-gapped AMI

9.2.3.4.1 Create an Image using Konvoy Image Builder (KIB) for use in an air-gapped cluster

9.2.3.4.2 Prerequisites

- [Minimal IAM Permissions for KIB](#) (see page 1468) to create an Image for an AWS account using Konvoy Image Builder.

 The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to create an AWS Air-gapped AMI to take advantage of enhanced cluster operations. Explore the [KIB with AWS](#) (see page 1468) topics for more options.

Using [KIB](#) (see page 1464), you can build an AMI without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded](#) (see page 78) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2/kib
```

2. Follow the instructions below to build an AMI.

Depending on which [version of DKP](#) (see page 1464) you are running, steps and flags will be different. To deploy in a region where CAPI images are not provided, you need to use KIB to create your own image for the region. For a list of supported AWS regions, refer to the [Published AMI](#)¹¹⁹⁵ information from AWS.

9.2.3.4.2.1 Execute the following to begin image creation:

Run the `konvoy-image` command to build and validate the image.

```
konvoy-image build aws images/ami/centos-79.yaml --overrides overrides/offline.yaml
```

By default it builds in the `us-west-2` region. to specify another region set the `--region` flag:

¹¹⁹⁵ <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

```
konvoy-image build aws --region us-east-1 images/ami/centos-79.yaml --overrides
overrides/offline.yaml
```

When the command is complete the `ami` id is printed and written to `./manifest.json`.

After you complete these steps, you can [Air-gapped Seed the Registry](#) (see page 1537).



If using AWS ECR as your local private registry, more information can be found on the [Registry Mirror Tools](#) (see page 1533) page.

9.2.3.5 Add Custom Tags to Image

As certain cloud environments enforce the tagging of resources, DKP platform users have inquired about whether is possible to add custom tags on an AMI that was built with Konvoy Image Builder (KIB). Currently, adding a custom tag as an attribute to the AMI is only possible by modifying the `packer.json` file.

In order to add a tag, perform the commands below:

1. Generate manifest files for KIB by executing the following command. In this example, we're building a CentOS 7.9 image:

```
./konvoy-image build images/ami/centos-79.yaml
```

2. Send `SIGINT` to the process to kill it after seeing the output `...writing new packer configuration` to `work/centos-7-xxxxxxx-yyyyy`. During the attempt to build the AMI a `packer.json` is generated under the path `work/centos-7-xxxxxxx-yyyyy/packer.json`.
3. Edit the `packer.json` file by adding the parameter `"run_tags"` to the `packer.json` file as seen in the example below:

```
"builders": [
{
"name": "{{(user `distribution`) | lower}}-{{user `distribution_version`}}
{{user `build_name_extra`}}",
"type": "amazon-eks",
"run_tags": {"my_custom_tag": "tag"},
"instance_type": "{{user `aws_instance_type`}}",
"source_ami_filter": {
"filters": {
"virtualization-type": "hvm",
"name": "{{user `ami_filter_name`}}",
```

```
"root-device-type": "ebs",
"architecture": "x86_64"
},
```

4. Use the packer manifest option to build the AMI:

```
./konvoy-image build aws images/ami/centos-79.yaml --packer-manifest=work/
centos-7-1658174984-TycMM/packer.json
2022/09/07 18:23:33 writing new packer configuration to work/centos-7-166257501
3-zJUHP
2022/09/07 18:23:33 starting packer build
centos-7.9: output will be in this color.

==> centos-7.9: Prevalidating any provided VPC information
==> centos-7.9: Prevalidating AMI Name: konvoy-ami-centos-7-1.26.3-1662575013
centos-7.9: Found Image ID: ami-0686851c4e7b1a8e1
==> centos-7.9: Creating temporary keypair: packer_6318e1a6-9f45-c01b-c7ca-
f5404735f709
==> centos-7.9: Creating temporary security group for this instance:
packer_6318e1aa-7448-d74b-9b90-166f26d21619
==> centos-7.9: Authorizing access to port 22 from [0.0.0.0/0] in the temporary
security groups...
==> centos-7.9: Launching a source AWS instance...
centos-7.9: Adding tag: "my_custom_tag": "tag"
centos-7.9: Instance ID: i-04d457b20713dcf7d
==> centos-7.9: Waiting for instance (i-04d457b20713dcf7d) to become ready...
==> centos-7.9: Using SSH communicator to connect: 34.222.153.107
==> centos-7.9: Waiting for SSH to become available...
```

Now that your tags have been added to the AMI, you can continue to [Advanced Configuration](#) (see page 1051) for AWS and finish building your clusters.

9.2.3.6 KIB for EKS

AWS EKS best practices discourage building custom images. The [Amazon EKS Optimized AMI](#)¹¹⁹⁶ is the preferred way to deploy containers for EKS. If the image is customized, it breaks some of the autoscaling and security capabilities of EKS.

¹¹⁹⁶ <https://docs.aws.amazon.com/eks/latest/userguide/eks-optimized-amis.html>

9.2.4 KIB for Azure

9.2.4.1 Learn how to build a custom Azure Image for use with DKP

This procedure describes how to use the [Konvoy Image Builder \(see page 1464\)](#) (KIB) to create a [Cluster API](#)¹¹⁹⁷ compliant Azure Virtual Machine (VM) Image. The VM Image contains the base operating system you specify and all the necessary Kubernetes components. The Konvoy Image Builder uses variable `overrides` to specify the base image and container images to use in your new Azure VM image.



The default Azure image is not recommended for use in production. We suggest using KIB for Azure to build the image in order to take advantage of enhanced cluster operations. To explore more information on this topic refer to the [Azure Infrastructure \(see page 1159\)](#).

9.2.4.2 Prerequisites

Before you begin, you must:

- Download the [Konvoy Image Builder \(see page 1464\)](#) bundle for your version of DKP.
- Check the [Supported Infrastructure Operating Systems \(see page 69\)](#)
- Check the [Supported Kubernetes Version \(see page 1552\)](#) for your Provider.
- Create a working `Docker` setup.

9.2.4.3 Extract the KIB Bundle

Extract the bundle and `cd` into the extracted `konvoy-image-bundle-$VERSION_$(OS)` folder. The bundled version of `konvoy-image` contains an embedded `docker` image that contains all the requirements for building.



The `konvoy-image` binary and all supporting folders are also extracted. When extracted, `konvoy-image` bind mounts the current working directory (`$(PWD)`) into the container to be used.

¹¹⁹⁷ <https://cluster-api.sigs.k8s.io/>

9.2.4.4 Configure Azure Prerequisites

- If you have already followed the [Azure Prerequisites \(see page 1159\)](#) topic steps, then the environment variables needed by KIB ([AZURE_CLIENT_SECRET , AZURE_CLIENT_ID , AZURE_TENANT_ID , AZURE_SUBSCRIPTION_ID]) are set and do not need repeated if you are still working in the same window.

If you have not executed the [Azure Prerequisite \(see page 1159\)](#) steps, they are listed below.

1. Sign in to Azure:

```
az login
```

```
[
  {
    "cloudName": "AzureCloud",
    "homeTenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "id": "b1234567-abcd-11a1-a0a0-1234a5678b90",
    "isDefault": true,
    "managedByTenants": [],
    "name": "Mesosphere Developer Subscription",
    "state": "Enabled",
    "tenantId": "a1234567-b132-1234-1a11-1234a5678b90",
    "user": {
      "name": "user@azuresmesosphere.onmicrosoft.com",
      "type": "user"
    }
  }
]
```

2. Create an Azure Service Principal (SP) by running the following command:

If an SP with the name exists, this command will rotate the password.

```
az ad sp create-for-rbac --role contributor --name "$(whoami)-konvoy" --
scopes=/subscriptions/$(az account show --query id -o tsv) --query
"{ client_id: appId, client_secret: password, tenant_id: tenant }"
```

```
{
  "client_id": "7654321a-1a23-567b-b789-0987b6543a21",
  "client_secret": "Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C",
  "tenant_id": "a1234567-b132-1234-1a11-1234a5678b90"
}
```

3. Set the `AZURE_CLIENT_SECRET` environment variable:

```
export AZURE_CLIENT_SECRET="<azure_client_secret>" #
Z79yVstq_E.R0R7RUUck718vEHSuyhAB0C
export AZURE_CLIENT_ID="<client_id>" # 7654321a-1a23-567b-
b789-0987b6543a21
export AZURE_TENANT_ID="<tenant_id>" # a1234567-b132-1234-1a11-12
34a5678b90
export AZURE_SUBSCRIPTION_ID="<subscription_id>" # b1234567-abcd-11a1-
a0a0-1234a5678b90
```

4. Ensure you have an [override file](#) (see page 1513) to [configure specific attributes](#)¹¹⁹⁸ of your Azure image.

9.2.4.5 Build the Image

Run the `konvoy-image` command to build and validate the image.

```
konvoy-image build azure --client-id ${AZURE_CLIENT_ID} --tenant-id $
{AZURE_TENANT_ID} --overrides override-source-image.yaml images/azure/ubuntu-2004.yam
l
```

By default, the image builder builds in the `westus2` location. To specify another location set the `--location` flag (shown in example below is how to change the location to `eastus`):

```
konvoy-image build azure --client-id ${AZURE_CLIENT_ID} --tenant-id $
{AZURE_TENANT_ID} --location eastus --overrides override-source-image.yaml images/
azure/centos-7.yaml
```

When the command is complete, the image id is printed and written to the `./manifest.json` file. This file has an `artifact_id` field whose value provides the name of the image. You should then specify this image id when creating the cluster.

¹¹⁹⁸ <https://github.com/mesosphere/konvoy-image-builder/tree/7447074a6d910e71ad2e61fc4a12820d073ae5ae/images/azure>

9.2.4.6 Image Gallery

By default Konvoy Image Builder will create a Resource Group, Gallery, and Image Name to store the resulting image in. To specify a specific Resource Group, Gallery, or Image Name flags may be specified:

```

--gallery-image-locations string    a list of locations to publish the image
(default same as location)
--gallery-image-name string        the gallery image name to publish the image to
--gallery-image-offer string       the gallery image offer to set (default "dkp")
--gallery-image-publisher string   the gallery image publisher to set (default
"dkp")
--gallery-image-sku string         the gallery image sku to set
--gallery-name string              the gallery name to publish the image in
(default "dkp")
--resource-group string           the resource group to create the image in
(default "dkp")

```

When creating your cluster, you will then add this flag during the create process for your custom image: `--compute-gallery-id "<Managed Image Shared Image Gallery Id>"`. See [Create a New Azure Cluster \(see page 1166\)](#) for specific consumption of image commands.

The SKU and Image Name will default to the values found in the image YAML.

9.2.4.7 Marketplace Images for Rocky Linux

Similar to Image Gallery, additional flags allow DKP to create a cluster with Marketplace based images for Rocky Linux 9.0: `--plan-offer`, `--plan-publisher` and `--plan-sku`.

- If you use these fields in the [override file \(see page 1513\)](#) when you create a machine image with KIB, you must also set the corresponding flags when you create your cluster with DKP.
- Conversely, if you do not use these fields when you create a machine image with KIB, you do not need to set these flags when you create your cluster with DKP.

Rocky Linux YAML

```

---
download_images: true

packer:
  distribution: "rockylinux-9"          # Offer
  distribution_version: "rockylinux-9" # SKU
  # Azure Rocky linux official image: https://portal.azure.com/#view/
  # Microsoft_Azure_Marketplace/GalleryItemDetailsBladeNopdl/id/
  # erockyenterprisesoftwarefoundationinc1653071250513.rockylinux-9
  image_publisher: "erockyenterprisesoftwarefoundationinc1653071250513"
  image_version: "latest"
  ssh_username: "azureuser"

```

```

plan_image_sku: "rockylinux-9" # SKU
plan_image_offer: "rockylinux-9" # offer
plan_image_publisher: "erockyenterprisesoftwarefoundationinc1653071250513" #
publisher

build_name: "rocky-90-az"
packer_builder_type: "azure"
python_path: ""

```

9.2.4.8 KIB for AKS

Because custom virtual machine images are discouraged in AKS, Konvoy Image Builder (KIB) does not include any support for building custom machine images for AKS. If the image is customized, it breaks some of the autoscaling and security capabilities of AKS. The AKS managed image ensures everything is patched and tuned for container workload.

9.2.5 KIB with GCP

This procedure describes how to use the [Konvoy Image Builder](#) (see page 1464) (KIB) to create a [Cluster API](#)¹¹⁹⁹ compliant GCP image. GCP images contain configuration information and software to create a specific, pre-configured, operating environment. For example, you can create a GCP image of your current computer system settings and software. The GCP image can then be replicated and distributed, creating your computer system for other users. KIB uses [variable overrides](#) (see page 1513) to specify base image and container images to use in your new GCP image.



Google Cloud Platform does not publish images. You must first build the image using Konvoy Image Builder. For more information regarding images and clusters, refer to the [GCP Infrastructure](#) (see page 1368) section of the documentation.


9.2.5.1 Prerequisites

Before you begin, you must:

- Download the [KIB](#) (see page 1466) bundle for your version of DKP.
- Check the [Supported Infrastructure Operating Systems](#) (see page 69)
- Check the [Supported Kubernetes Version](#) (see page 1552) for your Provider.
- Create a working Docker or other registry setup.
- On Debian-based Linux distributions, install a version of the [cri-tools](#)¹²⁰⁰ package known to be compatible with both the Kubernetes and container runtime versions.
- Verify that your Google Cloud project does not have the Enable OS Login feature enabled. See below for more information:

¹¹⁹⁹ <https://cluster-api.sigs.k8s.io/>

¹²⁰⁰ <https://github.com/kubernetes-sigs/cni-plugins>


 The Enable OS Login feature is sometimes enabled by default in GCP projects. If the OS login feature is enabled, KIB will not be able to `ssh` to the VM instances it creates and will not be able to successfully create an image. To check if it is enabled, use the commands on this page https://cloud.google.com/compute/docs/metadata/setting-custom-metadata#console_2 to inspect the metadata configured in in your project. If you find the the `enable-oslogin` flag set to TRUE, you must remove (or set it to FALSE) to successfully use KIB.

9.2.5.2 GCP Prerequisite Roles

If you are creating your image on either a non-GCP instance or one that does not have the [required roles](#)¹²⁰¹, you must either:

- Create a [GCP service account](#)¹²⁰².
- If you have already created a service account, retrieve the credentials for an existing service account.
- Export the static credentials that will be used to create the cluster:

```
export GCP_B64ENCODED_CREDENTIALS=$(base64 < "$  
{GOOGLE_APPLICATION_CREDENTIALS}" | tr -d '\n')
```

 Make sure to rotate static credentials for increased security.

- If you are creating your image on either a non-GCP instance or one that does not have the [required roles](#) (see page 1368):
 - (option 1) Create a service account using the following `gcloud` commands:

```
export GCP_PROJECT=<your GCP project ID>  
export GCP_SERVICE_ACCOUNT_USER=<some new service account user>  
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"  
  
gcloud iam service-accounts create "${GCP_SERVICE_ACCOUNT_USER}" --  
project=${GCP_PROJECT}  
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --  
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${  
{GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/  
compute.instanceAdmin.v1
```

¹²⁰¹ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/295272692>

¹²⁰² <https://cloud.google.com/iam/docs/service-account-overview>

```
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
iam.serviceAccountUser
gcloud iam service-accounts keys create $
{GOOGLE_APPLICATION_CREDENTIALS} --iam-account="$
{GCP_SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com"
```


- (option 2) If you have already created a service account, retrieve the credentials for an existing service account using the following `gcloud` commands:

```
export GCP_PROJECT=<your GCP project ID>
export GCP_SERVICE_ACCOUNT_USER=<existing service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
compute.instanceAdmin.v1
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${GCP_SERVICE_ACCOUNT_USER}@${
GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/
iam.serviceAccountUser
gcloud iam service-accounts keys create $
{GOOGLE_APPLICATION_CREDENTIALS} --iam-account="$
{GCP_SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com"
```

- Export the static credentials that will be used to create the cluster:

```
export GCP_B64ENCODED_CREDENTIALS=$(base64 < "$
{GOOGLE_APPLICATION_CREDENTIALS}" | tr -d '\n')
```

 Make sure to rotate static credentials for increased security.

9.2.5.2.1 Build the GCP image

1. Run the `konvoy-image` command to build and validate the image:

```
./konvoy-image build gcp --project-id ${GCP_PROJECT} --network ${NETWORK_NAME}
images/gcp/ubuntu-2004.yaml
```

2. KIB will run and print out the name of the created image, you will use this name when creating a Kubernetes cluster. See sample output below:

```
...
==> ubuntu-2004-focal-v20220419: Deleting instance...
    ubuntu-2004-focal-v20220419: Instance has been deleted!
==> ubuntu-2004-focal-v20220419: Creating image...
==> ubuntu-2004-focal-v20220419: Deleting disk...
    ubuntu-2004-focal-v20220419: Disk has been deleted!
==> ubuntu-2004-focal-v20220419: Running post-processor: manifest
Build 'ubuntu-2004-focal-v20220419' finished after 7 minutes 46 seconds.

==> Wait completed after 7 minutes 46 seconds

==> Builds finished. The artifacts of successful builds are:
--> ubuntu-2004-focal-v20220419: A disk image was created: konvoy-ubuntu-2004-1-23-7-1658523168
--> ubuntu-2004-focal-v20220419: A disk image was created: konvoy-ubuntu-2004-1-23-7-1658523168
```

3. To find a list of images you have created in your account, run the following command:

```
gcloud compute images list --no-standard-images
```

9.2.5.2.2 Create a Network (optional)

Building an image requires a [Network](#)¹²⁰³ with firewall rules that allow SSH access to the VM instance.

1. Set your GCP Project ID for your `gcp` account unless already set previously:

```
export GCP_PROJECT=<your GCP project ID>
```

2. Run the following to create a new network:

```
export NETWORK_NAME=kib-ssh-network
gcloud compute networks create "${NETWORK_NAME}" --project="${GCP_PROJECT}" --
subnet-mode=auto --mtu=1460 --bgp-routing-mode=regional
```

3. Create the firewall rule to allow Ingress access on port 22:

```
gcloud compute firewall-rules create "${NETWORK_NAME}-allow-ssh" --project="${GCP_PROJECT}" --network="projects/${GCP_PROJECT}/global/networks/${NETWORK_NAME}" --description="Allows TCP connections from any source to any
```

¹²⁰³ <https://cloud.google.com/vpc/docs/vpc>

```
instance on the network using port 22." --direction=INGRESS --priority=65534 --
source-ranges=0.0.0.0/0 --action=ALLOW --rules=tcp:22
```

With your KIB image now created, you can now move on to create your cluster and set up your [Cluster API](#)¹²⁰⁴ (CAPI) controllers..

9.2.6 KIB for GPU

Create a GPU supported OS image using Konvoy Image Builder

Using the [Konvoy Image Builder](#) (see [page 1464](#)), you can build an image that has support to use NVIDIA GPU hardware to support GPU workloads.

NOTE: The NVIDIA driver requires a specific Linux kernel version. Make sure that the base image for the OS version has the required kernel version.

See [Supported Infrastructure Operating Systems](#) (see [page 69](#)) for a list of OS versions and the corresponding kernel versions known to work with the NVIDIA driver.

If the [NVIDIA runfile](#)¹²⁰⁵ installer has not been downloaded, then retrieve and install the download first by running the following command. The first line in the command below downloads and installs the runfile and the second line places it in the artifacts directory (you must create an `artifacts` directory if you don't already have one).

```
curl -O https://download.nvidia.com/XFree86/Linux-x86_64/470.82.01/NVIDIA-Linux-
x86_64-470.82.01.run
mv NVIDIA-Linux-x86_64-470.82.01.run artifacts
```

DKP supported NVIDIA driver¹²⁰⁶ version is 470.x.

To build an image for use on GPU enabled hardware, perform the following steps.

1. In your `overrides/nvidia.yaml` file, add the following to enable GPU builds. You can also access and use the overrides [repo](#)¹²⁰⁷ or in the documentation under [Nvidia GPU Override File](#) (see [page 1518](#)) or [Offline Nvidia Override file](#) (see [page 1519](#)).

¹²⁰⁴ <https://cluster-api.sigs.k8s.io/>

¹²⁰⁵ <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

¹²⁰⁶ <https://www.nvidia.com/Download/Find.aspx>

¹²⁰⁷ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

- a. Non-air-gapped GPU override:

```
gpu:
  types:
    - nvidia
  build_name_extra: "-nvidia"
```

- b. Air-gapped GPU override:

```
# Use this file when building a machine image, not as a override secret
for preprovisioned environments
nvidia_runfile_local_file: "{{ playbook_dir }}/../artifacts/
{{ nvidia_runfile_installer }}"
gpu:
  types:
    - nvidia

build_name_extra: "-nvidia"
```

NOTE: For [RHEL Pre-provisioned Override Files](#) (see page 1524) used with KIB, see specific note for GPU.

2. Build your image using the `build` Konvoy Image Builder command, making sure to include the flag `--instance-type` that specifies an AWS instance that has an available GPU:

AWS Example:

```
konvoy-image build --region us-east-1 --instance-type=p2.xlarge --source-ami=ami-12345abcdef images/ami/centos-7.yaml --overrides overrides/nvidia.yaml
```

In this example, we chose an instance type with an NVIDIA GPU using the `--instance-type` flag, and we provided the NVIDIA overrides using the `--overrides` flag. See [KIB with AWS](#) (see page 1468) for more information on creating an AML.

Additional helpful information can be found in the [NVIDIA Device Plug-in](#)¹²⁰⁸ for Kubernetes instructions and the [Installation Guide of Supported Platforms](#)¹²⁰⁹.

See also: [NVIDIA documentation](#)¹²¹⁰

9.2.6.1 Verification

To verify that the NVIDIA driver is working, connect to the node and execute this command:

```
nvidia-smi
```

¹²⁰⁸ <https://github.com/NVIDIA/k8s-device-plugin/blob/master/README.md>

¹²⁰⁹ <https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/install-guide.html>

¹²¹⁰ https://nvidia.custhelp.com/app/answers/detail/a_id/131/kw/driver%20installation%20docs/related/1

When drivers are successfully installed, the display looks like the following:

```

Fri Jun 11 09:05:31 2021
+-----+
| NVIDIA-SMI 460.73.01    Driver Version: 460.73.01    CUDA Version: 11.2    |
+-----+-----+-----+-----+-----+-----+
| GPU  Name          Persistence-M| Bus-Id        Disp.A | Volatile Uncorr. ECC |
| Fan  Temp   Perf   Pwr:Usage/Cap|      Memory-Usage | GPU-Util  Compute M. |
|                                           MIG M.         |
+-----+-----+-----+-----+-----+-----+
|   0   Tesla K80          Off          | 00000000:00:1E:0  Off |                    0 |
| N/A   35C    P0        73W / 149W |      0MiB / 11441MiB |      99%      Default |
|                                           N/A           |
+-----+-----+-----+-----+-----+

+-----+
| Processes:
| GPU  GI  CI           PID  Type  Process name          GPU Memory
|     ID  ID                                   Usage
+-----+-----+-----+-----+-----+
| No running processes found
+-----+

```

9.2.7 KIB with vSphere

vSphere for DKP using Konvoy Image Builder

This section assumes you have access to a vSphere environment and have credentials with proper privileges and access. Refer to the [vCenter](#)¹²¹¹ and [vSphere Client](#)¹²¹² documentation for details.

- [Create a vSphere Base OS Image](#) (see page 1487)
- [Create a vSphere Virtual Machine Template](#) (see page 1489)

9.2.7.1 Create a vSphere Base OS Image

Creating a base OS image from DVD ISO files is a one-time process. Building a base OS image creates a base vSphere template in your vSphere environment. The base OS image is used by [Konvoy Image Builder](#) (see page 1464)(KIB) to create a VM template to configure Kubernetes nodes by the DKP vSphere provider. For more information about images in vSphere vCenter, refer to [VMware documentation](#)¹²¹³.

1211 <https://www.vmware.com/products/vcenter-server.html>

1212 <https://docs.vmware.com/en/VMware-vSphere/index.html>

1213 <https://docs.vmware.com/en/VMware-Cloud-Foundation/4.5/vcf-deploy/GUID-2674DA5A-8DF7-4212-A4A9-88CD798DC303.html>

9.2.7.1.1 Create the Base OS Image

For vSphere, a username and password is populated by `SSH_USERNAME` and the user can use authorization via `SSH_PASSWORD` or `SSH_PRIVATE_KEY_FILE` environment variables and [required by default by packer](#)¹²¹⁴. This user should have administrator privileges. It is possible to configure a custom user and password when building the OS image, however, that requires the Konvoy Image Builder (KIB) configuration to be [overridden](#) (see page 1513).

While creating the base OS image, it is important to take into consideration the following elements:

- **Storage configuration:** D2iQ recommends customizing disk partitions and not configuring a SWAP partition.
- **Network configuration:** as KIB must download and install packages, activating the network is required.
- **Connect to Red Hat:** if using RHEL, registering with Red Hat is required to configure software repositories and install software packages.
- **Software selection:** D2iQ recommends choosing **Minimal Install**.
- DKP recommends to install with the packages provided by the operating system package managers. Use the version that corresponds to the major version of your operating system.

9.2.7.1.1.1 Disk Size

For each cluster you create using this base OS image, ensure you establish the disk size of the **root file system** based on:

- The minimum [DKP storage requirements](#) (see page 100).
- The minimum storage requirements for your organization.

Defaults

Clusters are created with a default disk size of 80 GB.



For clusters created with the default disk size, the base OS image root file system must be exactly 80 GB. The root file system cannot be reduced automatically when a machine first boots.

Customization

You can also specify a custom disk size when you [create a cluster](#) (see page 0) (see the flags of the [Create new vSphere Cluster](#) (see page 1670) command). This allows you to use one base OS image to create multiple clusters that have different storage requirements.

¹²¹⁴ <https://github.com/mesosphere/konvoy-image-builder/blob/0523fd2c5e6e1ad1d4962f60a47039aa145a6e42/pkg/packer/manifests/vsphere/packer.pkr.hcl>



Before specifying a disk size when you create a cluster, take into account:

1. **For some base OS images, the custom disk size option has no effect on the size of the root file system.** This is because some root file systems, for example, those contained in an LVM Logical Volume, cannot be resized automatically when a machine first boots.
2. **The specified custom disk size must be equal to, or larger than the size of the base OS image root file system.** This is because a root file system cannot be reduced automatically when a machine first boots.
3. In [VMware Cloud Director](#) (see page 1325), the base image determines the minimum storage available for the VM.

For additional information and examples of KIB vSphere overrides, see this page: [Base Image Override Files](#) (see page 1520)

The next step is to [create a vSphere VM template](#) (see page 1489) that contains the CAPI and Kubernetes objects.

Refer to the [vCenter](#)¹²¹⁵ and [vSphere Client](#)¹²¹⁶ documentation for details.

9.2.7.1.2 Next Step:

[Create a vSphere Virtual Machine Template](#) (see page 1489)

9.2.7.2 Create a vSphere Virtual Machine Template

Create a vSphere template for your cluster from a base OS image

You must have at least one image before creating a new cluster. As long as you have an image, this step in your configuration is not required each time since that image can be used to spin up a new cluster. However, if you need different images for different environments or providers, you will need to create a new custom image.

The [Konvoy Image Builder](#) (see page 1490) (KIB) uses the values in `image.yaml` and the input base OS image to create a vSphere template directly on the vCenter server. Using KIB, you can build an image without requiring access to the internet by providing an additional `--override` flag.

9.2.7.2.1 Prerequisites

- Users need to create a [base OS image](#) (see page 1490) in their vSphere client before starting this procedure.
- Konvoy Image Builder (KIB) [downloaded](#) (see page 78) and extracted
- **Air-gapped Only**

¹²¹⁵ <https://www.vmware.com/products/vcenter-server.html>

¹²¹⁶ <https://docs.vmware.com/en/VMware-vSphere/index.html>

- Assuming you have [downloaded](#) (see [page 78](#)) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2/kib
```

- Follow the instructions to build a vSphere template below and set the override `--overrides overrides/offline.yaml` flag.

9.2.7.2.2 Create a vSphere Template for Your Cluster from a Base OS Image

Using the base OS image created in a previous procedure, DKP creates the new vSphere template directly on the vCenter server.

1. Set the following vSphere environment variables on the bastion VM host:

```
export VSPHERE_SERVER=your_vCenter_APIserver_URL
export VSPHERE_USERNAME=your_vCenter_user_name
export VSPHERE_PASSWORD=your_vCenter_password
```

2. Copy the base OS image file created in the vSphere Client to your desired location on the bastion VM host, and make a note of the path and file name.
3. Create an `image.yaml` file and add the following variables for vSphere. DKP uses this file and these variables as inputs in the next step.

⚠ NOTE: This example is Ubuntu 20.04. You will need to replace OS name below based on your OS. See other [YAML examples](#)¹²¹⁷ for copy and paste below last step.

```
---
download_images: true
build_name: "ubuntu-2004"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: "<VSPHERE_CLUSTER_NAME>"
  datacenter: "<VSPHERE_DATACENTER_NAME>"
```

¹²¹⁷ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

datastore: "<VSPHERE_DATASTORE_NAME>"
folder: "<VSPHERE_FOLDER>"
insecure_connection: "false"
network: "<VSPHERE_NETWORK>"
resource_pool: "<VSPHERE_RESOURCE_POOL>"
template: "os-qualification-templates/d2iq-base-Ubuntu-20.04" # change
default value with your base template name
vsphere_guest_os_type: "other4xLinux64Guest"
guest_os_type: "ubuntu2004-64"
# goss params
distribution: "ubuntu"
distribution_version: "20.04"
# Use following overrides to select the authentication method that can be used
with base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key
will be ignored

```

4. Create a vSphere VM template with your variation of the following command:

```
konvoy-image build images/ova/<image.yaml>
```

Any additional configurations can be added to this command using `--overrides` flags as shown below:

- a. Any credential overrides: `--overrides overrides.yaml`
 - b. for FIPS, add this flag: `--overrides overrides/fips.yaml`
 - c. for air-gapped, add this flag: `--overrides overrides/offline-fips.yaml`
5. The [Konvoy Image Builder](#) (see page 1490) (KIB) uses the values in `image.yaml` and the input base OS image to create a vSphere template directly on the vCenter server. This template contains the required artifacts needed to create a Kubernetes cluster.

When KIB provisions the OS [image](#)¹²¹⁸ successfully, it creates a manifest file. The `artifact_id` field of this file contains the name of the AMI ID (AWS), template name (vSphere), or image name (GCP/Azure), for example:

```

{
  "name": "vsphere-clone",
  "builder_type": "vsphere-clone",
  "build_time": 1644985039,
  "files": null,

```

¹²¹⁸ <https://github.com/mesosphere/konvoy-image-builder/tree/main/images/ova>

```

    "artifact_id": "konvoy-ova-vsphere-rhel-84-1.21.6-1644983717",
    "packer_run_uuid": "260e8110-77f8-ca94-e29e-ac7a2ae779c8",
    "custom_data": {
      "build_date": "2022-02-16T03:55:17Z",
      "build_name": "vsphere-rhel-84",
      "build_timestamp": "1644983717",
      [...]
    }
  }
}

```

Recommendation: Now we can now see the template created in our vCenter, it is best to rename it to `dkp-<DKP_VERSION>-k8s-<K8S_VERSION>-<DISTRO>`, like `dkp-2.6.2-k8s-1.26.14-ubuntu` to keep templates organized.

6. Next steps are to deploy a DKP cluster using your vSphere template.

9.2.7.2.2.1 Additional OS YAML files for copy/paste:

RHEL 8.6

```

---
download_images: true
build_name: "rhel-86"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "os-qualification-templates/d2iq-base-RHEL-86" # change default value
with your base template name
vsphere_guest_os_type: "rhel8_64Guest"
guest_os_type: "rhel8-64"
# goss params
distribution: "RHEL"
distribution_version: "8.6"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'

```

```
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored
```

Rocky Linux 9.1

```
---
download_images: true
build_name: "rocky-91"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  cluster: ""
  datacenter: ""
  datastore: ""
  folder: ""
  insecure_connection: "false"
  network: ""
  resource_pool: ""
  template: "os-qualification-templates/d2iq-base-RockyLinux-9.1" # change default
value with your base template name
  vsphere_guest_os_type: "other4xLinux64Guest"
  guest_os_type: "rocky9-64"
  # goss params
  distribution: "rocky"
  distribution_version: "9.1"
# Use following overrides to select the authentication method that can be used with
base template
# ssh_username: "" # can be exported as environment variable 'SSH_USERNAME'
# ssh_password: "" # can be exported as environment variable 'SSH_PASSWORD'
# ssh_private_key_file = "" # can be exported as environment variable
'SSH_PRIVATE_KEY_FILE'
# ssh_agent_auth: false # is set to true, ssh_password and ssh_private_key will be
ignored
```

9.2.8 KIB with Pre-provisioned Environments

In Pre-provisioned environments, KIB runs inside the `bootstrap cluster`, against the list of nodes that you define. Whereas with other providers, you run KIB outside the cluster manually to build your images.

For Pre-provisioned, you define a set of nodes that already exist. During the cluster creation process, KIB is built into DKP and automatically runs the machine configuration process (which KIB uses to build images with other providers) against the set of nodes that you defined. This results in your pre-existing/'pre-provisioned' nodes being configured properly. The remainder of the cluster provisioning happens automatically after that.

- In a **Pre-provisioned environment** (see page 149), you have existing machines and DKP consumes them to form a cluster.
- When you have another provisioner (for example, cloud providers such [AWS](#) (see page 294), [vSphere](#) (see page 436) and others), you build images with KIB and DKP consumes the images to provision machines and form a cluster.

9.2.9 Konvoy Image Builder CLI

- [konvoy-image build](#) (see page 1494)
- [konvoy-image completion](#) (see page 1501)
- [konvoy-image generate-docs](#) (see page 1506)
- [konvoy-image generate](#) (see page 1506)
- [konvoy-image provision](#) (see page 1510)
- [konvoy-image upload](#) (see page 1511)
- [konvoy-image validate](#) (see page 1512)
- [konvoy-image version](#) (see page 1513)

9.2.9.1 Other resources:

- Main [Github KIB CLI](#)¹²¹⁹
- [DKP CLI](#) (see page 1644)

9.2.9.2 konvoy-image build

build and provision images

9.2.9.2.1 Synopsis

Build and Provision images. Specifying AWS arguments is deprecated and will be removed in a future version. Use the `aws` subcommand instead.

```
konvoy-image build <image.yaml> [flags]
```

9.2.9.2.2 Options

```
--ami-groups stringArray      a list of AWS groups which are allowed use
the image, using 'all' result in a public image
--ami-regions stringArray     a list of regions to publish amis
```

¹²¹⁹ <https://github.com/mesosphere/konvoy-image-builder/tree/main/docs/cli>

```

--ami-users stringArray      a list AWS user accounts which are allowed
use the image
--containerd-version string  the version of containerd to install
--dry-run                    do not create artifacts, or delete them
after creating. Recommended for tests.
--extra-vars strings        flag passed Ansible's extra-vars
-h, --help                  help for build
--instance-type string      instance type used to build the AMI; the
type must be present in the region in which the AMI is built
--kubernetes-version string  The version of kubernetes to install.
Example: 1.21.6
--overrides strings         a comma separated list of override YAML
files
--packer-manifest string     provide the path to a custom packer manifest
--packer-on-error string     [advanced] set error strategy for packer.
strategies [cleanup, abort, run-cleanup-provisioner]
--packer-path string        the location of the packer binary (default
"packer")
--region string             the region in which to build the AMI
--source-ami string         the ID of the AMI to use as the source; must
be present in the region in which the AMI is built
--source-ami-filter-name string restricts the set of source AMIs to ones
whose Name matches filter
--source-ami-filter-owner string restricts the source AMI to ones with this
owner ID
--work-dir string          path to custom work directory generated by
the generate command

```

9.2.9.2.3 Options inherited from parent commands

```

--color    enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose   enable debug level logging (same as --v 5)

```

9.2.9.2.4 SEE ALSO

- [konvoy-image build aws](#) (see page 1495) - build and provision aws images
- [konvoy-image build azure](#) (see page 1497) - build and provision azure images
- [konvoy-image build gcp](#) (see page 1498) - build and provision gcp images
- [konvoy-image build vsphere](#) (see page 1499) - build and provision vsphere images

9.2.9.2.5 konvoy-image build aws

build and provision aws images

```
konvoy-image build aws <image.yaml> [flags]
```

9.2.9.2.5.1 Examples

```
aws --region us-west-2 --source-ami=ami-12345abcdef images/ami/centos-79.yaml
```

9.2.9.2.5.2 Options

```

    --ami-groups stringArray      a list of AWS groups which are allowed use
the image, using 'all' result in a public image
    --ami-regions stringArray     a list of regions to publish amis
    --ami-users stringArray       a list AWS user accounts which are allowed
use the image
    --containerd-version string   the version of containerd to install
    --dry-run                     do not create artifacts, or delete them
after creating. Recommended for tests.
    --extra-vars strings          flag passed Ansible's extra-vars
    -h, --help                   help for aws
    --instance-type string        instance type used to build the AMI; the
type must be present in the region in which the AMI is built
    --kubernetes-version string   The version of kubernetes to install.
Example: 1.21.6
    --overrides strings          a comma separated list of override YAML
files
    --packer-manifest string      provide the path to a custom packer manifest
    --packer-on-error string      [advanced] set error strategy for packer.
strategies [cleanup, abort, run-cleanup-provisioner]
    --packer-path string         the location of the packer binary (default
"packer")
    --region string              the region in which to build the AMI
    --source-ami string          the ID of the AMI to use as the source; must
be present in the region in which the AMI is built
    --source-ami-filter-name string restricts the set of source AMIs to ones
whose Name matches filter
    --source-ami-filter-owner string restricts the source AMI to ones with this
owner ID
    --work-dir string            path to custom work directory generated by
the generate command

```

9.2.9.2.5.3 Options inherited from parent commands

```
--color      enable color output (default true)
```

```
-v, --v int    select verbosity level, should be between 0 and 6
--verbose      enable debug level logging (same as --v 5)
```

9.2.9.2.5.4 SEE ALSO

- [konvoy-image build](#) (see page 1497) - build and provision images

9.2.9.2.6 konvoy-image build azure

build and provision azure images

```
konvoy-image build azure <image.yaml> [flags]
```

9.2.9.2.6.1 Examples

```
azure --location westus2 --subscription-id <sub_id> images/azure/centos-79.yaml
```

9.2.9.2.6.2 Options

```
--client-id string          the client id to use for the build
--cloud-endpoint string    Azure cloud endpoint. Which can be one
of [Public USGovernment China] (default "Public")
--containerd-version string the version of containerd to install
--dry-run                  do not create artifacts, or delete them
after creating. Recommended for tests.
--extra-vars strings       flag passed Ansible's extra-vars
--gallery-image-locations stringArray a list of locations to publish the
image (default [westus])
--gallery-image-name string the gallery image name to publish the
image to
--gallery-image-offer string the gallery image offer to set (default
"dkp")
--gallery-image-publisher string the gallery image publisher to set
(default "dkp")
--gallery-image-sku string  the gallery image sku to set
--gallery-name string       the gallery name to publish the image
in (default "dkp")
-h, --help                  help for azure
--instance-type string      the Instance Type to use for the build
(default "Standard_D2s_v3")
--kubernetes-version string The version of kubernetes to install.
Example: 1.21.6
```

<code>--location</code> string	the location in which to build the
image (default "westus2")	
<code>--overrides</code> strings	a comma separated list of override YAML
files	
<code>--packer-manifest</code> string	provide the path to a custom packer
manifest	
<code>--packer-on-error</code> string	[advanced] set error strategy for
packer. strategies [cleanup, abort, run-cleanup-provisioner]	
<code>--packer-path</code> string	the location of the packer binary
(default "packer")	
<code>--resource-group</code> string	the resource group to create the image
in (default "dkp")	
<code>--subscription-id</code> string	the subscription id to use for the
build	
<code>--tenant-id</code> string	the tenant id to use for the build
<code>--work-dir</code> string	path to custom work directory generated
by the generate command	

9.2.9.2.6.3 Options inherited from parent commands

<code>--color</code>	enable color output (default true)
<code>-v, --v int</code>	select verbosity level, should be between 0 and 6
<code>--verbose</code>	enable debug level logging (same as <code>--v 5</code>)

9.2.9.2.6.4 SEE ALSO

- [konvoy-image build](#) (see page 1497) - build and provision images

9.2.9.2.7 konvoy-image build gcp

build and provision gcp images

```
konvoy-image build gcp <image.yaml> [flags]
```

9.2.9.2.7.1 Examples

```
gcp ... images/gcp/centos-79.yaml
```

9.2.9.2.7.2 Options

```

--containerd-version string  the version of containerd to install
--dry-run                    do not create artifacts, or delete them after
creating. Recommended for tests.
--extra-vars strings        flag passed Ansible's extra-vars
-h, --help                  help for gcp
--kubernetes-version string The version of kubernetes to install. Example:
1.21.6
--network string            the network to use when creating an image
--overrides strings         a comma separated list of override YAML files
--packer-manifest string    provide the path to a custom packer manifest
--packer-on-error string    [advanced] set error strategy for packer.
strategies [cleanup, abort, run-cleanup-provisioner]
--packer-path string        the location of the packer binary (default
"packer")
--project-id string         the project id to use when storing created image
--region string             the region in which to launch the instance
(default "us-west1")
--work-dir string           path to custom work directory generated by the
generate command

```

9.2.9.2.7.3 Options inherited from parent commands

```

--color    enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose   enable debug level logging (same as --v 5)

```

9.2.9.2.7.4 SEE ALSO

- [konvoy-image build](#) (see page 1497) - build and provision images

9.2.9.2.8 konvoy-image build vsphere

build and provision vsphere images

```
konvoy-image build vsphere <image.yaml> [flags]
```

9.2.9.2.8.1 Examples

```
vsphere --datacenter dc1 --cluster zone1 --datastore nfs-store1 --network public --
template=d2iq-base-templates/d2iq-base-CentOS-7.9 images/ami/centos-79.yaml
```

9.2.9.2.8.2 Options

```

--cluster string          vSphere cluster to be used. Alternatively set
host. Required value: you can pass the cluster name through an override file or image
definition file.
--containerd-version string  the version of containerd to install
--datacenter string        The vSphere datacenter. Required value: you can
pass the datacenter name through an override file or image definition file.
--datastore string         vSphere datastore used to build and store the
image template. Required value: you can pass the datastore name through an override
file or image definition file.
--dry-run                  do not create artifacts, or delete them after
creating. Recommended for tests.
--extra-vars strings       flag passed Ansible's extra-vars
--folder string            vSphere folder to store the image template
-h, --help                 help for vsphere
--host string              vSphere host to be used. Alternatively set
cluster. Required value: you can pass the host name through an override file or image
definition file.
--kubernetes-version string The version of kubernetes to install. Example:
1.21.6
--network string           vSphere network used to build image template.
Ensure the host running the command has access to this network. Required value: you
can pass the network name through an override file or image definition file.
--overrides strings        a comma separated list of override YAML files
--packer-manifest string   provide the path to a custom packer manifest
--packer-on-error string   [advanced] set error strategy for packer.
strategies [cleanup, abort, run-cleanup-provisioner]
--packer-path string       the location of the packer binary (default
"packer")
--resource-pool string     vSphere resource pool to be used to build image
template
--ssh-privatekey-file string Path to ssh private key which will be used to
log into the base image template
--ssh-publickey string     Path to SSH public key which will be copied to
the image template. Ensure to set ssh-privatekey-file or load the private key into
ssh-agent
--ssh-username string      username to be used with the vSphere image
template

```

```

--template string          Base template to be used. Can include folder.
                           <templatename> or <folder>/<templatename>. Required value: you can pass the template
                           name through an override file or image definition file.
--work-dir string          path to custom work directory generated by the
generate command

```

9.2.9.2.8.3 Options inherited from parent commands

```

--color          enable color output (default true)
-v, --v int    select verbosity level, should be between 0 and 6
--verbose        enable debug level logging (same as --v 5)

```

9.2.9.2.8.4 SEE ALSO

- [konvoy-image build](#) (see page 1497) - build and provision images

9.2.9.3 konvoy-image completion

Generate the autocompletion script for the specified shell

9.2.9.3.1 Synopsis

Generate the autocompletion script for konvoy-image for the specified shell. See each sub-command's help for details on how to use the generated script.

9.2.9.3.2 Options

```

-h, --help  help for completion

```

9.2.9.3.3 Options inherited from parent commands

```

--color          enable color output (default true)
-v, --v int    select verbosity level, should be between 0 and 6
--verbose        enable debug level logging (same as --v 5)

```

9.2.9.3.4 SEE ALSO

- [konvoy-image completion bash](#) (see page 1502) - Generate the autocompletion script for bash
- [konvoy-image completion fish](#) (see page 1503) - Generate the autocompletion script for fish

- [konvoy-image completion powershell](#) (see page 1504) - Generate the autocompletion script for powershell
- [konvoy-image completion zsh](#) (see page 1505) - Generate the autocompletion script for zsh

9.2.9.3.5 konvoy-image completion bash

Generate the autocompletion script for bash

9.2.9.3.5.1 Synopsis

Generate the autocompletion script for the bash shell.

This script depends on the 'bash-completion' package. If it is not installed already, you can install it via your OS's package manager.

To load completions in your current shell session:

```
source <(konvoy-image completion bash)
```

To load completions for every new session, execute once:

Linux:

```
konvoy-image completion bash > /etc/bash_completion.d/konvoy-image
```

macOS:

```
konvoy-image completion bash > $(brew --prefix)/etc/bash_completion.d/konvoy-image
```

You will need to start a new shell for this setup to take effect.

```
konvoy-image completion bash
```

9.2.9.3.5.2 Options

```
-h, --help          help for bash
--no-descriptions  disable completion descriptions
```

9.2.9.3.5.3 Options inherited from parent commands

```

--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)

```

9.2.9.3.5.4 SEE ALSO

- [konvoy-image completion](#) (see page 1504) - Generate the autocompletion script for the specified shell

9.2.9.3.6 konvoy-image completion fish

Generate the autocompletion script for fish

9.2.9.3.6.1 Synopsis

Generate the autocompletion script for the fish shell.

To load completions in your current shell session:

```
konvoy-image completion fish | source
```

To load completions for every new session, execute once:

```
konvoy-image completion fish > ~/.config/fish/completions/konvoy-image.fish
```

You will need to start a new shell for this setup to take effect.

```
konvoy-image completion fish [flags]
```

9.2.9.3.6.2 Options

```

-h, --help          help for fish
--no-descriptions  disable completion descriptions

```

9.2.9.3.6.3 Options inherited from parent commands

```

--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)

```

9.2.9.3.6.4 SEE ALSO

- [konvoy-image completion](#) (see page 1504) - Generate the autocompletion script for the specified shell

9.2.9.3.7 konvoy-image completion powershell

Generate the autocompletion script for powershell

9.2.9.3.7.1 Synopsis

Generate the autocompletion script for powershell.

To load completions in your current shell session:

```
konvoy-image completion powershell | Out-String | Invoke-Expression
```

To load completions for every new session, add the output of the above command to your powershell profile.

```
konvoy-image completion powershell [flags]
```

9.2.9.3.7.2 Options

```

-h, --help          help for powershell
--no-descriptions  disable completion descriptions

```

9.2.9.3.7.3 Options inherited from parent commands

```

--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)

```

9.2.9.3.7.4 SEE ALSO

- [konvoy-image completion](#) (see page 1504) - Generate the autocompletion script for the specified shell

9.2.9.3.8 konvoy-image completion zsh

Generate the autocompletion script for zsh

9.2.9.3.8.1 Synopsis

Generate the autocompletion script for the zsh shell.

If shell completion is not already enabled in your environment you will need to enable it. You can execute the following once:

```
echo "autoload -U compinit; compinit" >> ~/.zshrc
```

To load completions in your current shell session:

```
source <(konvoy-image completion zsh)
```

To load completions for every new session, execute once:

Linux:

```
konvoy-image completion zsh > "${fpath[1]}/_konvoy-image"
```

macOS:

```
konvoy-image completion zsh > $(brew --prefix)/share/zsh/site-functions/_konvoy-image
```

You will need to start a new shell for this setup to take effect.

```
konvoy-image completion zsh [flags]
```

9.2.9.3.8.2 Options

```
-h, --help          help for zsh
```

```
--no-descriptions  disable completion descriptions
```

9.2.9.3.8.3 Options inherited from parent commands

```
--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)
```

9.2.9.3.8.4 SEE ALSO

- [konvoy-image completion \(see page 1504\)](#) - Generate the autocompletion script for the specified shell

9.2.9.4 konvoy-image generate-docs

generate docs in path

```
konvoy-image generate-docs <PATH> [flags]
```

9.2.9.4.1 Examples

```
generate-docs /tmp/docs
```

9.2.9.4.2 Options

```
-h, --help  help for generate-docs
```

9.2.9.4.3 Options inherited from parent commands

```
--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)
```

9.2.9.5 konvoy-image generate

generate files relating to building images

9.2.9.5.1 Synopsis

Generate files relating to building images. Specifying AWS arguments is deprecated and will be removed in a future version. Use the `aws` subcommand instead.

```
konvoy-image generate <image.yaml> [flags]
```

9.2.9.5.2 Options

```

--ami-groups stringArray      a list of AWS groups which are allowed use
the image, using 'all' result in a public image
--ami-regions stringArray     a list of regions to publish amis
--ami-users stringArray       a list AWS user accounts which are allowed
use the image
--containerd-version string   the version of containerd to install
--extra-vars strings          flag passed Ansible's extra-vars
-h, --help                    help for generate
--instance-type string        instance type used to build the AMI; the
type must be present in the region in which the AMI is built
--kubernetes-version string   The version of kubernetes to install.
Example: 1.21.6
--overrides strings           a comma separated list of override YAML
files
--region string               the region in which to build the AMI
--source-ami string           the ID of the AMI to use as the source; must
be present in the region in which the AMI is built
--source-ami-filter-name string restricts the set of source AMIs to ones
whose Name matches filter
--source-ami-filter-owner string restricts the source AMI to ones with this
owner ID

```

9.2.9.5.3 Options inherited from parent commands

```

--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)

```

9.2.9.5.4 SEE ALSO

- [konvoy-image generate aws](#) (see page 1508) - generate files relating to building aws images
- [konvoy-image generate azure](#) (see page 1509) - generate files relating to building azure images

9.2.9.5.5 konvoy-image generate aws

generate files relating to building aws images

```
konvoy-image generate aws <image.yaml> [flags]
```

9.2.9.5.5.1 Examples

```
aws --region us-west-2 --source-ami=ami-12345abcdef images/ami/centos-79.yaml
```

9.2.9.5.5.2 Options

```

    --ami-groups stringArray      a list of AWS groups which are allowed use
the image, using 'all' result in a public image
    --ami-regions stringArray     a list of regions to publish amis
    --ami-users stringArray       a list AWS user accounts which are allowed
use the image
    --containerd-version string   the version of containerd to install
    --extra-vars strings          flag passed Ansible's extra-vars
    -h, --help                    help for aws
    --instance-type string        instance type used to build the AMI; the
type must be present in the region in which the AMI is built
    --kubernetes-version string   The version of kubernetes to install.
Example: 1.21.6
    --overrides strings           a comma separated list of override YAML
files
    --region string               the region in which to build the AMI
    --source-ami string           the ID of the AMI to use as the source; must
be present in the region in which the AMI is built
    --source-ami-filter-name string restricts the set of source AMIs to ones
whose Name matches filter
    --source-ami-filter-owner string restricts the source AMI to ones with this
owner ID

```

9.2.9.5.5.3 Options inherited from parent commands

```

    --color      enable color output (default true)
    -v, --v int select verbosity level, should be between 0 and 6
    --verbose    enable debug level logging (same as --v 5)

```

9.2.9.5.5.4 SEE ALSO

- [konvoy-image generate](#) (see page 1509) - generate files relating to building images

9.2.9.5.6 konvoy-image generate azure

generate files relating to building azure images

```
konvoy-image generate azure <image.yaml> [flags]
```

9.2.9.5.6.1 Examples

```
azure --location westus2 --subscription-id <sub_id> images/azure/centos-79.yaml
```

9.2.9.5.6.2 Options

<code>--client-id string</code>	the client id to use for the build
<code>--cloud-endpoint string</code>	Azure cloud endpoint. Which can be one of [Public USGovernment China] (default "Public")
<code>--containerd-version string</code>	the version of containerd to install
<code>--extra-vars strings</code>	flag passed Ansible's extra-vars
<code>--gallery-image-locations stringArray</code>	a list of locations to publish the image (default [westus])
<code>--gallery-image-name string</code>	the gallery image name to publish the image to
<code>--gallery-image-offer string</code>	the gallery image offer to set (default "dkp")
<code>--gallery-image-publisher string</code>	the gallery image publisher to set (default "dkp")
<code>--gallery-image-sku string</code>	the gallery image sku to set
<code>--gallery-name string</code>	the gallery name to publish the image in (default "dkp")
<code>-h, --help</code>	help for azure
<code>--instance-type string</code>	the Instance Type to use for the build (default "Standard_D2s_v3")
<code>--kubernetes-version string</code>	The version of kubernetes to install.
Example: <code>1.21.6</code>	
<code>--location string</code>	the location in which to build the image (default "westus2")
<code>--overrides strings</code>	a comma separated list of override YAML files
<code>--resource-group string</code>	the resource group to create the image in (default "dkp")


```

--subscription-id string      the subscription id to use for the
build
--tenant-id string           the tenant id to use for the build

```

9.2.9.5.6.3 Options inherited from parent commands

```

--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)

```

9.2.9.5.6.4 SEE ALSO

- [konvoy-image generate \(see page 1509\)](#) - generate files relating to building images

9.2.9.6 konvoy-image provision

provision to an inventory.yaml or hostname, note the comma at the end of the hostname

```
konvoy-image provision <inventory.yaml|hostname,> [flags]
```

9.2.9.6.1 Examples

```
provision --inventory-file inventory.yaml
```

9.2.9.6.2 Options

```

--containerd-version string  the version of containerd to install
--extra-vars strings        flag passed Ansible's extra-vars
-h, --help                  help for provision
--inventory-file string     an ansible inventory defining your infrastructure
--kubernetes-version string The version of kubernetes to install. Example:
1.21.6
--overrides strings         a comma separated list of override YAML files
--provider string           specify a provider if you wish to install
provider specific utilities
--work-dir string           path to custom work directory generated by the
generate command

```

9.2.9.6.3 Options inherited from parent commands

```

--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)

```

9.2.9.7 konvoy-image upload

Upload one of [artifacts]

9.2.9.7.1 Options

```

-h, --help  help for upload

```

9.2.9.7.2 Options inherited from parent commands

```

--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)

```

9.2.9.7.3 SEE ALSO

- [konvoy-image upload artifacts \(see page 1511\)](#) - upload offline artifacts to hosts defined in inventory-file

9.2.9.7.4 konvoy-image upload artifacts

upload offline artifacts to hosts defined in inventory-file

```

konvoy-image upload artifacts [flags]

```

9.2.9.7.4.1 Options

```

--container-images-dir string  path to container images for install on remote
hosts.
--containerd-bundle string     path to Containerd tar file for install on
remote hosts.

```

```

--extra-vars strings      flag passed Ansible's extra-vars
-h, --help               help for artifacts
--inventory-file string  an ansible inventory defining your
infrastructure (default "inventory.yaml")
--nvidia-runfile string  path to nvidia runfile to place on remote
hosts.
--os-packages-bundle string path to os-packages tar file for install on
remote hosts.
--overrides strings     a comma separated list of override YAML files
--pip-packages-bundle string path to pip-packages tar file for install on
remote hosts.
--work-dir string       path to custom work directory generated by the
command

```

9.2.9.7.4.2 Options inherited from parent commands

```

--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)

```

9.2.9.7.4.3 SEE ALSO

- [konvoy-image upload](#) (see page 1514) - Upload one of [artifacts]

9.2.9.8 konvoy-image validate

validate existing infrastructure

```
konvoy-image validate [flags]
```

9.2.9.8.1 Options

```

--apiserver-port int  apiserver port (default 6443)
-h, --help           help for validate
--inventory-file string an ansible inventory defining your infrastructure
(default "inventory.yaml")
--pod-subnet string   ip addresses used for the pod subnet (default
"192.168.0.0/16")
--service-subnet string ip addresses used for the service subnet (default
"10.96.0.0/12")

```

9.2.9.8.2 Options inherited from parent commands

```

--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)

```

9.2.9.9 konvoy-image version

Version for konvoy-image

```
konvoy-image version [flags]
```

9.2.9.9.1 Options

```
-h, --help  help for version
```

9.2.9.9.2 Options inherited from parent commands

```

--color      enable color output (default true)
-v, --v int  select verbosity level, should be between 0 and 6
--verbose    enable debug level logging (same as --v 5)

```

9.2.10 Use Override Files with Konvoy Image Builder

9.2.10.1 Learn how to use override files with Konvoy Image builder

[Konvoy Image Builder](#) (see page 1464) uses [YAML](#)¹²²⁰ `override` files to configure specific attributes. These files provide information to override default values for certain parts of your image file. These `override` files can modify the version and parameters of the image description and the Ansible playbook.

There are 2 types of override files:

- [Default override](#) (see page 1514) files provided by Konvoy Image Builder for you to use for specific purposes.
- [Custom override files](#) (see page 1520) you create to use with Konvoy Image Builder.

¹²²⁰ <https://github.com/mesosphere/konvoy-image-builder/tree/7447074a6d910e71ad2e61fc4a12820d073ae5ae/images>

9.2.10.2 Default Override Files

9.2.10.2.1 Learn how to use override files with DKP

DKP comes with default override files:

- [FIPS Override Files](#) (see page 1514)
- [Air-gapped Environment FIPS Override Files](#) (see page 1516)
- [Offline Override File](#) (see page 1517)
- [Nvidia GPU Override File](#) (see page 1518)
- [Offline Nvidia GPU Override file](#) (see page 1519)
- [Oracle Redhat Linux Kernel Override File](#) (see page 1520)

You can find these override files in the [Konvoy Image Builder repo](#)¹²²¹.

9.2.10.2.2 FIPS Override Files

9.2.10.2.2.1 Cloud Provisioners Override File:

Online FIPS Override File (Non-air-gapped)

Add the following FIPS Overrides file to your environment:

```
--overrides overrides/fips.yaml
```

```
---
k8s_image_registry: docker.io/mesosphere

fips:
  enabled: true

build_name_extra: -fips
kubernetes_build_metadata: fips.0
default_image_repo: hub.docker.io/mesosphere
kubernetes_rpm_repository_url: "https://packages.d2iq.com/konvoy/stable/linux/repos/el/kubernetes-v{{ kubernetes_version }}-fips/x86_64"
docker_rpm_repository_url: "\
  https://containerd-fips.s3.us-east-2.amazonaws.com\
  /{{ ansible_distribution_major_version|int }}\
  /x86_64"
```

¹²²¹ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>



You can find all available Overrides files in the [Konvoy Image Builder repo](#)¹²²².

9.2.10.2.2.2 Pre-provisioned Environments Override File:

Online FIPS Override File (Pre-provisioned)

Add the following FIPS Overrides file to your environment:

1. If your pre-provisioned machines need to have a default Override file like FIPS, create a secret that includes the overrides in a file:

```
cat > fips.yaml << EOF
---
k8s_image_registry: docker.io/mesosphere

fips:
  enabled: true

build_name_extra: -fips
kubernetes_build_metadata: fips.0
default_image_repo: hub.docker.io/mesosphere
kubernetes_rpm_repository_url: "https://packages.d2iq.com/konvoy/stable/linux/
repos/el/kubernetes-v{{ kubernetes_version }}-fips/x86_64"
docker_rpm_repository_url: "\
https://containerd-fips.s3.us-east-2.amazonaws.com\
/{{ ansible_distribution_major_version|int }}\
/x86_64"
EOF
```

2. Create the related secret by running the following command:

```
kubectl create secret generic $CLUSTER_NAME-user-overrides --from-
file=fips.yaml=fips.yaml
kubectl label secret $CLUSTER_NAME-user-overrides clusterctl.cluster.x-k8s.io/
move=
```

¹²²² <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>



You can find all available Overrides files in the [Konvoy Image Builder repo](#)¹²²³.

9.2.10.2.3 Air-gapped Environment FIPS Override Files

9.2.10.2.3.1 Cloud Provisioners Override File:

Offline FIPS Override File (Air-gapped)

Add the following FIPS Overrides file to your environment:

```
--overrides overrides/offline-fips.yaml
```

```
# fips os-packages
os_packages_local_bundle_file: "{{ playbook_dir }}/../artifacts/
{{ kubernetes_version }}_{{ ansible_distribution|lower }}
_{{ ansible_distribution_major_version }}_x86_64_fips.tar.gz"
containerd_local_bundle_file: "{{ playbook_dir }}/../artifacts/
{{ containerd_tar_file }}"
pip_packages_local_bundle_file: "{{ playbook_dir }}/../artifacts/pip-packages.tar.gz"
images_local_bundle_dir: "{{ playbook_dir }}/../artifacts/images"
```



You can find all available Overrides files in the [Konvoy Image Builder repo](#)¹²²⁴.

9.2.10.2.3.2 Pre-provisioned Environments Override File:

Offline FIPS Override File (Air-gapped)

Add the following FIPS Overrides file to your environment:

1. If your pre-provisioned machines need to have a default Override file like FIPS, create a secret that includes the overrides in a file:

¹²²³ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

¹²²⁴ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

```
cat > fips.yaml << EOF
# fips os-packages
os_packages_local_bundle_file: "{{ playbook_dir }}/../artifacts/
{{ kubernetes_version }}_{{ ansible_distribution|lower }}
_{{ ansible_distribution_major_version }}_x86_64_fips.tar.gz"
containerd_local_bundle_file: "{{ playbook_dir }}/../artifacts/
{{ containerd_tar_file }}"
pip_packages_local_bundle_file: "{{ playbook_dir }}/../artifacts/pip-
packages.tar.gz"
images_local_bundle_dir: "{{ playbook_dir }}/../artifacts/images"
EOF
```

2. Create the related secret by running the following command:

```
kubectl create secret generic $CLUSTER_NAME-user-overrides --from-
file=fips.yaml=fips.yaml
kubectl label secret $CLUSTER_NAME-user-overrides clusterctl.cluster.x-k8s.io/
move=
```



You can find all available Overrides files in the [Konvoy Image Builder repo](#)¹²²⁵.

9.2.10.2.4 Offline Override File

You can find these override files in the [Konvoy Image Builder repo](#)¹²²⁶.

```
--overrides overrides/offline.yaml
```

```
os_packages_local_bundle_file: "{{ playbook_dir }}/../artifacts/
{{ kubernetes_version }}_{{ ansible_distribution|lower }}
_{{ ansible_distribution_major_version }}_x86_64.tar.gz"
containerd_local_bundle_file: "{{ playbook_dir }}/../artifacts/
{{ containerd_tar_file }}"
pip_packages_local_bundle_file: "{{ playbook_dir }}/../artifacts/pip-packages.tar.gz"
images_local_bundle_dir: "{{ playbook_dir }}/../artifacts/images"
```

¹²²⁵ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

¹²²⁶ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

- For Ubuntu 20.04, when Konvoy Image Builder runs, it will temporarily disable all defined Debian repositories by appending a `.disabled` suffix. At the end of installation, each repository will revert to its original name. In case of failures, the files will not be renamed back.

For [GPU Offline Override File](#) (see page 1519), you can use `nvidia-runfile` flag for GPU support if you have downloaded the [runfile installer](#)¹²²⁷.

9.2.10.2.5 Nvidia GPU Override File

9.2.10.2.5.1 GPU Override File

This override file should be used to create images for use with DKP that can use GPU hardware. These override files can also be located in the [Konvoy Image Builder repo](#)¹²²⁸.

```
--overrides overrides/nvidia.yaml
```

```
---
gpu:
  types:
    - nvidia
build_name_extra: "-nvidia"
```

9.2.10.2.5.2 GPU Override File for Pre-provisioned environments

If you require your environment to consume GPU resources, add the following FIPS Overrides file to your environment:

1. If your pre-provisioned machines need to have a default Override file like FIPS, create a secret that includes the overrides in a file:

```
cat > nvidia.yaml << EOF
---
gpu:
  types:
    - nvidia
build_name_extra: "-nvidia"
EOF
```

¹²²⁷ <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

¹²²⁸ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

2. Create the related secret by running the following command:

```
kubectl create secret generic $CLUSTER_NAME-user-overrides --from-
file=nvidia.yaml=nvidia.yaml
kubectl label secret $CLUSTER_NAME-user-overrides clusterctl.cluster.x-k8s.io/
move=
```

9.2.10.2.6 Offline Nvidia GPU Override file

This override file should be used to create images for use with DKP that can use GPU hardware. These override files can also be located in the [Konvoy Image Builder repo](#)¹²²⁹.

```
--overrides overrides/offline-nvidia.yaml
```

```
# Use this file when building a machine image, not as a override secret for
preprovisioned environments
nvidia_runfile_local_file: "{{ playbook_dir }}/../artifacts/
{{ nvidia_runfile_installer }}"
gpu:
  types:
    - nvidia

build_name_extra: "-nvidia"
```

9.2.10.2.6.1 GPU Override File for Air-gapped Pre-provisioned Environments

If you require your environment to consume GPU resources, add the following GPU Overrides file to your environment:

1. If your pre-provisioned machines need to have a default Override file like GPU, create a secret that includes the overrides in a file:

```
cat > nvidia.yaml << EOF
---
nvidia_runfile_local_file: "{{ playbook_dir }}/../artifacts/
{{ nvidia_runfile_installer }}"
gpu:
  types:
    - nvidia
build_name_extra: "-nvidia"
EOF
```

2. Create the related secret by running the following command:

¹²²⁹ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

```
kubectl create secret generic $CLUSTER_NAME-user-overrides --from-
file=nvidia.yaml=nvidia.yaml
kubectl label secret $CLUSTER_NAME-user-overrides clusterctl.cluster.x-k8s.io/
move=
```

9.2.10.2.7 Oracle Redhat Linux Kernel Override File

You can find these override files in the [Konvoy Image Builder repo](#)¹²³⁰.

```
--overrides overrides/rhck.yaml
```

```
---
oracle_kernel: RHCK
```

9.2.10.3 Custom Override Files

You can customize your KIB images using override files to specify alternate package libraries, Docker image repos, HTTP/S proxy information, and other customizations.

Some example custom override files are described in the following topics:

- [Base Image Override Files](#) (see page 1520)
- [HTTP Proxy Override Files](#) (see page 1524)
- [Pre-Provisioned Override Files](#) (see page 1524)
- [Private Registry in Air-gapped Override](#) (see page 1525)

9.2.10.3.1 Base Image Override Files

When using KIB to create an OS image that is compliant with DKP, the instructions on how to build and configure the image are included in the file located in `images/<builder-type>/<os-version>.yaml`, as seen below:

```
./konvoy-image build images/<builder-type>/{os}-{version}.yaml
```



In previous versions of KIB, only Azure and AWS providers were available, but now vSphere is an option as well.

Although there are several parameters specified by default in the Packer templates for each provider, it is possible to override the default values.

¹²³⁰ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

1. One option is to execute KIB with specific flags to override the values of the source AMI (`--source-ami`), AMI region (`--ami-regions`), AWS EC2 instance type (`--aws-instance-type`), and so on. For a comprehensive list of these flags, please run:

```
./konvoy-image build --help
```

2. Another option is by creating a file with the parameters to be overridden and specify the `--overrides` flag as shown below:

```
./konvoy-image build images/<builder-type>/<os>-<version>.yaml --overrides
overrides.yaml
```



While CLI flags can be used in combination with override files, CLI flags take priority over any override files.

9.2.10.3.1.1 Using Overrides flag:

AWS Example:

For example, when using the AWS Packer builder to override the above base image with another base image, create an override file and set the `source_ami` under the `packer` key. This overrides the image search and forces the use of the specified `source_ami`.

```
---
packer:
  source_ami: "ami-0123456789"
```

After creating the override file for our `source_ami`, we can pass our override file by using the `--overrides` flag when building our image:

```
./konvoy-image build aws images/ami/centos-7.yaml --overrides override-source-
ami.yaml
```

vSphere Example:

Create your vsphere overrides file `overrides/vcenter.yaml` and fill in the relevant details for your vSphere environment:

```

---
packer:
  vcenter_server: <FQDN of your vcenter>
  vsphere_username: <username for vcenter e.g. administrator@vsphere.local>
  vsphere_password: <password for vcenter>
  ssh_username: builder
  ssh_password: <password for your VMs builder user>
  linked_clone: false
  cluster: <vsphere cluster to use>
  datacenter: <vsphere datacenter to use>
  datastore: <vsphere datastore to use for templates>
  folder: <vsphere folder to store the template in>
  insecure_connection: "true"
  network: <a DHCP-enabled network in vcenter>
  resource_pool: <vsphere resource pool to use>

```

After creating the override file for our `source_ova`, we can pass our override file by using the `--overrides` flag when building our image:

```

konvoy-image build images/ova/ubuntu-2004.yaml \
  --overrides overrides/kube-version.yaml \
  --overrides overrides/vcenter.yaml

```

9.2.10.3.1.2 Override Credentials in Packer:

To abide to security practices, a user could set their own username and password while creating the base OS image and override the default credentials in KIB. To do this, create a file with the following content:

```

---
packer:
  ssh_username: "<USERNAME>"
  ssh_password: "<PASSWORD>"

```

For a complete list of the variables that can be modified for each Packer builder, users can refer to:

- [AWS Packer template](#)¹²³¹
- [Azure Packer template](#)¹²³²
- [GCP Packer template](#)¹²³³
- [vSphere Packer template](#)¹²³⁴

AWS Example:

¹²³¹ <https://github.com/mesosphere/konvoy-image-builder/blob/main/pkg/packer/manifests/aws/packer.pkr.hcl>

¹²³² <https://github.com/mesosphere/konvoy-image-builder/blob/main/pkg/packer/manifests/azure/packer.pkr.hcl>

¹²³³ <https://github.com/mesosphere/konvoy-image-builder/blob/main/pkg/packer/manifests/gcp/packer.pkr.hcl>

¹²³⁴ <https://github.com/mesosphere/konvoy-image-builder/blob/main/pkg/packer/manifests/vsphere/packer.pkr.hcl>

An AWS example would be the current base image description at `images/ami/centos-7.yaml` which is similar to the following:

```
---
# Example images/ami/centos-7.yaml
download_images: true
packer:
  ami_filter_name: "CentOS 7.9.2009 x86_64"
  ami_filter_owners: "125523088429"
  distribution: "CentOS"
  distribution_version: "7"
  source_ami: ""
  ssh_username: "centos"
  root_device_name: "/dev/sda1"
build_name: "centos-7"
packer_builder_type: "amazon"
python_path: ""
```

or

Azure Example:

An Azure example would be the current base image description at `images/azure/centos-79.yaml` which is similar to the following:

```
---
# Example images/azure/centos-79.yaml
download_images: true
packer:
  distribution: "centos" #offer
  distribution_version: "7_9-gen2" #SKU
  image_publisher: "openlogic"
  image_version: "latest"
  ssh_username: "centos"

build_name: "centos-7"
packer_builder_type: "azure"
python_path: ""
```

or

vSphere Example:

A vSphere example would be the current base image description at `images/vsphere/rhel-79.yaml` which is similar to the following:

```
---
download_images: true
build_name: "rhel-79"
```

```

packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
cluster: "zone1"
  datacenter: "dc1"
  datastore: "esxi-06-disk1"
  folder: "cluster-api"
  insecure_connection: "false"
  network: "Airgapped"
  resource_pool: "Users"
  template: "base-rhel-7"
  vsphere_guest_os_type: "rhel7_64Guest"
  guest_os_type: "rhel7-64"
  #goss params
  distribution: "RHEL"
  distribution_version: "7.9"

```

See [Supported Operating Systems](#) (see page 69) for details.

9.2.10.3.2 HTTP Proxy Override Files

You can use an HTTP proxy configuration when creating your image. The Ansible playbooks create `systemd` drop-in files for `containerd` and `kubelet` to configure the `http_proxy`, `https_proxy`, and `no_proxy` environment variables for the service from the file `/etc/konvoy_http_proxy.conf`.

To configure a proxy for use during image creation, create a new override file and specify the following:

```

# Example override-proxy.yaml
---
http_proxy: http://example.org:8080
https_proxy: http://example.org:8081
no_proxy: example.org,example.com,example.net

```

These values are only used for the image creation. After the image is created, the Ansible playbooks remove the `/etc/konvoy_http_proxy.conf` file.

You can use the `dkp` command to configure the `KubeadmConfigTemplate` object to create this file on start-up of the image with values supplied during the `dkp` invocation. This enables using different proxy settings for image creation and runtime.

9.2.10.3.3 Pre-Provisioned Override Files

Pre-provisioned environments require certain override files to work properly. The override files can contain HTTP Proxy info and other factors, including whether you want the image to be FIPS-enabled, NVIDIA optimized, and so on.

Those override files must also have a Secret that includes all of the overrides you wish to provide in one file.

Prior to DKP 2.6, you had to specify the proxy in the KIB override setup and then again in the `create cluster` command, even though in this case they would always both use the exact same proxy setting. As of DKP 2.6, an HTTP proxy gets created from the Konvoy flags for the control plane proxy and workers proxy values. The flags in the DKP command for Preprovisioned clusters populate a Secret automatically in the bootstrap cluster. That Secret has a known name that the Preprovisioned controller finds and applies when it runs the KIB provisioning job.

The nodes that get Kubernetes on them through CAPPP automatically get the HTTP proxy secrets that you set using the flags. You no longer have to put the proxy info in both the overrides AND the create cluster command as an argument.

For example, if you wish to provide an override with Docker credentials and a different source for EPEL on a CentOS7 machine, you can create a file like this:

```
image_registries_with_auth:
- host: "registry-1.docker.io"
  username: "my-user"
  password: "my-password"
  auth: ""
  identityToken: ""
```

- For RHEL Pre-provisioned using GPU, provide the following additional lines of information in your override file:

```
rhsm_user: ""
rhsm_password: ""
```

See Example below:

```
image_registries_with_auth:
- host: "registry-1.docker.io"
  username: "my-user"
  password: "my-password"
  auth: ""
  identityToken: ""

rhsm_user: ""
rhsm_password: ""
```

9.2.10.3.4 Private Registry in Air-gapped Override

To create an Air-gapped registry override for Docker hub, run the following command:


```
[plugins."io.containerd.grpc.v1.cri".registry.mirrors."docker.io"]
  endpoint = ["https://my-local-registry.local/v2/harbor-registry","https://
registry-1.docker.io"]
```

Tell the override how to create a wildcard mirror with this command:

```
[plugins."io.containerd.grpc.v1.cri".registry.mirrors."*"]
  endpoint = ["https://my-local-registry.local/v2/harbor-registry"]
```

9.3 FIPS 140-2 Compliance

Understand FIPS-140 Operating Mode and Requirements

Developed by a working group of government, industry operators, and vendors, the Federal Information Processing Standard (FIPS), FIPS-140 defines security requirements for cryptographic modules. FIPS defines what cryptographic cyphers can be used. Kubernetes uses encryption by default between various components and FIPS support ensures that the ciphers used for those communications meet those standards. The standard provides for a wide spectrum of data sensitivity, transaction values, and a diversity of application environment security situations. The standard specifies four security levels for each of eleven requirement areas. Each successive level offers increased security.

NIST introduced FIPS 140-2 validation, by accredited third party laboratories, as a formal, rigorous process to protect sensitive digitally-stored information not under Federal security classifications.

9.3.1 FIPS Support in DKP

DKP supports provisioning a FIPS-enabled Kubernetes control plane. Core Kubernetes components are compiled using a version of Go, called goboring, which uses a FIPS-certified [cryptographic module](#)¹²³⁵ for all cryptographic functions.

Before provisioning DKP, you will need to follow your OS vendor's instructions to ensure that your OS, or OS images, are [prepared for operating in FIPS mode](#)¹²³⁶.



You cannot apply FIPS-mode to an existing cluster, you must create a new cluster with FIPS enabled. Similarly, a FIPS-mode cluster must remain a FIPS-mode cluster; you cannot change the cluster's FIPS status after you create it.

¹²³⁵ <https://csrc.nist.gov/CSRC/media/projects/cryptographic-module-validation-program/documents/security-policies/140sp3702.pdf>

¹²³⁶ https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/security_guide/chap-federal_standards_and_regulations

9.3.2 Infrastructure Requirements for FIPS-140-2 Mode

To ensure proper operations in FIPS mode, be sure that your environment meets these requirements:

[FIPS 140 Mode Performance Impact](#) (see page 1533)

9.3.2.1 Supported Operating Systems

Supported Operating Systems for FIPS mode are Red Hat Enterprise Linux and CentOS. See the [Supported Operating Systems](#) (see page 69) for details on the tested and supported versions.

9.3.3 Deploying a Cluster in FIPS mode

In order to create a cluster in FIPS mode, we must inform the bootstrap controllers of the appropriate image repository and version tags of the official D2iQ FIPS builds of Kubernetes.

9.3.3.1 Supported FIPS Builds

Component	Repository	Version
Kubernetes	docker.io/mesosphere ¹²³⁷	v1.26.14+fips.0
etcd	docker.io/mesosphere ¹²³⁸	3.5.10+fips.0

AWS Example:

When creating a cluster, use the following command line options:

- `--ami <fips enabled AMI created in the previous step>` (AWS only)
- `--kubernetes-version <version>+fips.<build>`
- `--etcd-version <version>+fips.<build>`
- `--kubernetes-image-repository docker.io/mesosphere`
- `--etcd-image-repository docker.io/mesosphere`

For example:

```
dkp create cluster aws --cluster-name myFipsCluster \
--ami=ami-03dcaa75d45aca36f \
```

¹²³⁷ <https://hub.docker.com/layers/mesosphere/kube-apiserver/v1.26.14/images/sha256-3c559c75123fa40d588bc2d9050938a0b7dd1dab4e38a4fd24847cd0c0f24694?context=explore>

¹²³⁸ https://hub.docker.com/layers/mesosphere/etcd/v3.5.10_fips.0/images/sha256-61df32b8ec1e6229d8bd94caad07331c7e0c0b8e1947af8152c69eb2f04fb771?context=explore

```
--kubernetes-version=v1.26.6+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--etcd-image-repository=docker.io/mesosphere \
--etcd-version=3.5.6+fips.0
```

vSphere Example:


```
dkp create cluster vsphere \
  --cluster-name ${CLUSTER_NAME} \
  --network <NETWORK_NAME> \
  --control-plane-endpoint-host <xxx.yyy.zzz.000> \
  --data-center <DATACENTER_NAME> \
  --data-store <DATASTORE_NAME> \
  --folder <FOLDER_NAME> \
  --server <VCENTER_API_SERVER_URL> \
  --ssh-public-key-file <SSH_PUBLIC_KEY_FILE> \
  --resource-pool <RESOURE_POOL_NAME> \
  --vm-template <TEMPLATE_NAME> \
  --self-managed \
  --kubernetes-version=v1.26.6+fips.0 \
  --kubernetes-image-repository=docker.io/mesosphere \
  --etcd-image-repository=docker.io/mesosphere --etcd-version=3.5.6+fips.0
```

9.3.4 Create FIPS 140 Images: Non-air-gapped Environments

9.3.4.1 Use [Konvoy Image Builder](#) to create images with FIPS-compliant binaries


9.3.4.2 Non-air-gapped Environment Create FIPS-140 images

KIB can produce images containing FIPS-140 compliant binaries. Use the `fips.yaml` [override file](#) (see [page 1517](#)) provided with the image bundles.

 You can also find these override files in the [Konvoy Image Builder repo](#)¹²³⁹.

¹²³⁹ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

9.3.4.2.1 Examples:

 The below snippets will create images with FIPS-compliant Kubernetes components. If you need the underlying OS to be FIPS-compliant, then you will need to provide the specific FIPS-compliant OS image, using the `--source-ami` flag for AWS.

- A non-air-gapped environment example of override file use is the command below, which produces a FIPS-compliant image on RHEL 8.4 for AWS:

Replace `ami` with your infrastructure provisioner

```
konvoy-image build --overrides overrides/fips.yaml images/ami/rhel-84.yaml
```

- vSphere FIPS-complaint using `image.yaml` created during [VM Template \(see page 1492\)](#) configuration:

```
konvoy-image build --overrides overrides/fips.yaml images/ova/<image.yaml>
```

Here is a list of [FIPS Override Files](#).

9.3.5 Create FIPS 140 Images: Air-gapped Environment

KIB can produce images containing FIPS-140 compliant binaries. Use the `fips.yaml` [override file \(see page 1517\)](#) provided with the image bundles.

 You can also find these override files in the [Konvoy Image Builder repo](#)¹²⁴⁰.

¹²⁴⁰ <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

9.3.5.1 Examples:

■ The below snippets will create images with FIPS-compliant Kubernetes components. If you need the underlying OS to be FIPS-compliant, then you will need to provide the specific FIPS-compliant OS image, using the `--source-ami` flag for AWS.

- An air-gapped environment example of override file use is the command below which produces an AWS FIPS-compliant image on RHEL 8.4:

```
konvoy-image build --overrides offline-fips.yaml --overrides overrides/fips.yaml
images/ami/rhel-84.yaml
```

- vSphere FIPS-compliant air-gapped environment example:

```
konvoy-image build --overrides offline-fips.yaml --overrides overrides/fips.yaml
images/ova/<image.yaml>
```

9.3.5.2 Pre-provisioned FIPS Infrastructure

If you are targeting a [Pre-provisioned Installs](#) (see page 149), you can create a FIPS-compliant cluster by doing the following:

1. Create a [Pre-provisioned: Bootstrap Cluster](#) (see page 1222)
2. Create a secret on the bootstrap cluster with the contents from `fips.yaml` [override file](#)¹²⁴¹ and any other user overrides you wish to provide

```
kubectl create secret generic $CLUSTER_NAME-fips-overrides --from-
file=overrides.yaml=overrides.yaml
kubectl label secret $CLUSTER_NAME-fips-overrides clusterctl.cluster.x-k8s.io/move=
```

Here is a list of [FIPS Override Files](#) (see page 1517).

¹²⁴¹ <https://github.com/mesosphere/konvoy-image-builder/blob/main/overrides/fips.yaml>

9.3.6 Validate FIPS in Cluster

You can use the FIPS validation tool to verify that specific components and services are FIPS-compliant. The tool checks the components by comparing their file signatures against ones stored in a signed signature file, and by checking that services are using the certified algorithms.

9.3.6.1 Run FIPS validation

To verify the cluster is FIPS compliant, run `dkp check cluster fips`. This command reads from the signature files embedded in the `dkp` executable in order to validate that specific components and services are FIPS-compliant. Run the command:

```
dkp check cluster fips
```

Upon successful completion, the command's output displays details about the deployment in JSON format. If validation fails, the output will say which components fail and a list of the nodes that failed validation will return.

The full command usage and flags include:

```
dkp check cluster fips [flags]
```

Flags:

```
-h, --help                Help for fips
--kubeconfig string       Path to the kubeconfig file for the fips
cluster. If unspecified, default discovery rules apply.
-n, --namespace string     If present, the namespace scope for this CLI
request. (default "default")
--output-configmap string  ConfigMap to store result of the fips check.
(default "check-cluster-fips-output") (DEPRECATED: This flag will be removed in a
future release.)
--signature-configmap string ConfigMap with fips signature data to verify.
--signature-file string    File containing fips signature data.
--timeout duration         The length of time to wait before giving up.
Zero means wait forever (e.g. 1s, 2m, 3h). (default 10m0s)
```

9.3.6.1.1 Run FIPS validation with custom signature file

To validate FIPS-mode operation with the a custom signature file, you can use the `signature-file` flag, as in the following command. You also need to use the `signature-configmap` flag to set the name of the ConfigMap used to store your custom signature file.

```
dkp check cluster fips \
  --signature-file custom.json.asc \
  --signature-configmap custom-signature-file
```

9.3.6.1.2 Run FIPS validation with existing ConfigMap

If you already have a signature file stored in a ConfigMap, you can omit the `signature-file` flag, as in the following command:

```
dkp check cluster fips \
  --signature-configmap prod-rhel8-fips-signatures
```

9.3.6.2 Signature Files

The following signature files are already embedded in the `dkp` executable. They are provided for reference. You do not need to download them to run the FIPS check.

9.3.6.2.1 DKP Version 2.6.0

Operating System version	Kubernetes version	containerd version	Signature File URL
CentOS 7.9	v1.26.14	1.6.28	CentOS 7.9 ¹²⁴²
Oracle 7.9	v1.26.14	1.6.28	Oracle 7.9 ¹²⁴³
RHEL 7.9	v1.26.14	1.6.28	RHEL 7.9 ¹²⁴⁴
RHEL 8.4	v1.26.14	1.6.28	RHEL 8.4 ¹²⁴⁵
RHEL 8.6	v1.26.14	1.6.28	RHEL 8.6 ¹²⁴⁶

¹²⁴² <https://downloads.d2iq.com/dkp/fips/v2.6.2/manifest-centos-79.json.asc>

¹²⁴³ <https://downloads.d2iq.com/dkp/fips/v2.6.2/manifest-oracle-79.json.asc>

¹²⁴⁴ <https://downloads.d2iq.com/dkp/fips/v2.6.2/manifest-rhel-79.json.asc>

¹²⁴⁵ <https://downloads.d2iq.com/dkp/fips/v2.6.2/manifest-rhel-84.json.asc>

¹²⁴⁶ <https://downloads.d2iq.com/dkp/fips/v2.6.2/manifest-rhel-86.json.asc>

9.3.7 FIPS 140 Mode Performance Impact

9.3.7.1 Understand the performance impact from operating your cluster in FIPS 140 mode

The Go language cryptographic module, Goboring, relies on CGO's foreign function interface to call C-language functions exposed by the cryptographic module. Each call into the C library starts with a base overhead of 200ns.

One [benchmark](#)¹²⁴⁷ finds that the time to encrypt a single AES-128 block increased from 13ns to 209ns over the internal golang implementation. The preferred mode of D2iQ's FIPS module is

```
TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 .
```

The aggregate impact on a stable control plane seems to be an increase of around 10% CPU utilization over default operation. Workloads that do not directly interact with the control plane are not affected.

9.4 Registry Mirror Tools

Kubernetes does not natively provide a registry for hosting the container images you will use to run the applications you want to deploy on Kubernetes. Instead, Kubernetes requires you to use an external solution for storing and sharing container images. There are a variety of Kubernetes-compatible registry options that are compatible with DKP.

In an environment with access to the Internet, you can retrieve artifacts from specialized repositories dedicated to them, such as Docker images contained in DockerHub and Helm Charts that come from a dedicated Helm Chart repository. You can also create your own local repository to hold the container images you need, including any custom images you've created with the [Konvoy Image Builder](#) (see page 1464) tool.

In an air-gapped environment, you need a local repository to store Helm charts, Docker images, and other artifacts. Tools such as JFrog™ Artifactory, Amazon® AWS ECR, Harbor™, and Nexus™ handle multiple types of artifacts in one local repository.

9.4.1 Air-Gapped Registry Prerequisites

Private registries provide security and privacy into enterprise container image storage, whether hosted remotely or on-premises locally in an air-gapped environment. DKP in an air-gapped environment requires a local container registry of trusted images to enable production-level Kubernetes cluster management.

If you want to use images from this local registry to deploy applications inside your Kubernetes cluster, you'll need to set up a secret for a private registry. The secret contains your login data, which Kubernetes needs to connect to your private repository. Local registry tools compatible with DKP are listed below.

¹²⁴⁷ <https://github.com/golang/go/issues/21525>

9.4.2 AWS ECR

AWS ECR (Elastic Container Registry) is supported as your air-gapped image registry or a non-air-gapped registry mirror. DKP added support for using AWS ECR as a default registry when uploading image bundles in AWS.

9.4.2.1 Prerequisites

- Ensure you have followed the steps to create proper permissions in [Minimal Permissions and Role to Create Clusters](#) (see page 1060)
- Ensure you have created [Cluster IAM Policies, Roles, and Artifacts](#) (see page 1066)

9.4.2.2 Upload the Air-gapped Image Bundle to the Local ECR Registry:

A cluster administrator uses DKP CLI commands to upload the image bundle to ECR with parameters:

```
dkp push bundle --bundle <bundle> --to-registry=<ecr-registry-address>/<ecr-registry-name>
```

Parameter definitions:

- `--bundle <bundle>` the group of images. The example below is for the DKP air-gapped environment bundle
- `--to-registry=<ecr-registry-address>/<ecr-registry-name>` to provide registry location for push

An **example** command would be:

```
dkp push bundle --bundle container-images/konvoy-image-bundle-v2.6.2.tar --to-registry=333000009999.dkr.ecr.us-west-2.amazonaws.com/can-test
```

9.4.2.3 Air-gapped Environment Information regarding your AWS ECR Account

- The cluster administrator uses existing DKP CLI commands to create the cluster and refer to their internal ECR for image repository. The administrator does not need to provide static ECR registry credentials. See [Use a Registry Mirror](#) (see page 1536) and [Create an EKS Cluster from the CLI](#) (see page 1140) for more details.

9.4.3 JFrog Artifactory

JFrog Artifactory can function as a container registry, as well as an automated management tool for binaries and artifacts of all types. If you use [JFrog Artifactory](#)¹²⁴⁸ or **JFrog Container Registry**, you must update to a new version of the software. Use a build newer than version 7.11; older versions are not compatible.

9.4.4 Nexus Registry

[Nexus Repository](#)¹²⁴⁹ is a package registry for your Docker images and Helm Chart repositories and supports Proxy, Hosted, and Group repositories. It can be used a single registry for all your Kubernetes deployments.

9.4.5 Harbor Registry

[Install Harbor](#)¹²⁵⁰ and configure any HTTP access required, as well as the system level parameters in the `harbor.yml` file. Then run the installer script. If you are upgrading from a previous version of Harbor, you update the configuration file and migrate your data to fit the database schema of the later version. For information about upgrading, see [Upgrading Harbor](#)¹²⁵¹. A version than **Harbor Registry v2.1.1-5f52168e** will support OCI images.

While seeding you may see error messages such as the following:

```
2023/09/12 20:01:18 retrying without mount: POST https://harbor-registry.daclusta/v2/harbor-registry/mesosphere/kube-proxy/blobs/uploads/?from=mesosphere%2Fkube-proxy&mount=sha256%3A9fd5070b83085808ed850ff84acc98a116e839cd5dcfefa12f2906b7d9c6e50d&origin=REDACTED: UNAUTHORIZED: project not found, name: mesosphere: project not found, name: mesosphere
```

This appears to indicate that the image was not successfully pushed to your Harbor docker registry, but it is a false positive error message. This will only affect version of the DKP binary newer than DKP 2.4.0. This does not affect any other Local Registry solution such as Nexus or Artifactory. You can safely ignore these error messages.

9.4.6 Bastion Host

If you have not set up a [Bastion Host](#) (see page 1288) yet, refer to that section of the documentation.

9.4.7 Registry Mirror

If you need to set up a private registry with a registry mirror, see [this page](#) (see page 1536) for details on using that flag.

¹²⁴⁸ <https://jfrog.com/artifactory/>

¹²⁴⁹ <https://www.nexusregistry.com/info/>

¹²⁵⁰ <https://goharbor.io/docs/2.0.0/install-config/download-installer/>

¹²⁵¹ <https://goharbor.io/docs/2.0.0/administration/upgrade/>

9.4.8 Next Topic:

[Air-gapped Seed the Registry](#) (see page 1537)

9.4.9 Use a Registry Mirror

Registry mirrors are local copies of images from a public registry that follows (or mirrors) the file structure of a public registry. You can push container images to a local registry from downloaded images or images that you create with the Konvoy Image Builder. If your environment allows Internet access, when an image is not found, the mirror registry pulls it from it would consult its upstream registries. This kind of registry does not contain any images other than the ones requested.

To apply private registry configurations during the create operation, add the appropriate flags to the `create cluster` command. Only use one image registry per cluster:

Registry configuration	Flag
CA certificate chain to use while communicating with the registry mirror using TLS	<code>--registry-mirror-cacert file</code>
URL of a container registry to use as a mirror in the cluster	<code>--registry-mirror-url string</code>
Set to a user that has pull access to this registry	<code>--registry-mirror-username</code>
Password	<code>--registry-mirror-password</code>

This is useful when using an internal registry and when Internet access is not available such as in an air-gapped environment. However, registry mirrors can be used in non-air-gapped environments as well for security and speed.



AWS ECR - Adding the mirror flags to EKS would enable new clusters to also use ECR as image mirror. If you set the `--registry-mirror` flag, the Kubelet will now send requests to the `dynamic-credential-provider` with a different config. You can still pull your own images from ECR directly or use ECR as a mirror.

When the cluster is up and running, you can deploy and test workloads.


9.4.9.1 Alternative mirror example

```
dkp create cluster preprovisioned \
  --cluster-name ${CLUSTER_NAME} \
  --registry-mirror-cacert /tmp/registry.pem \
  --registry-mirror-url https://registry.example.com
```

More information is found in the [Custom Installation and Additional Infrastructure Tools \(see page 1039\)](#) sections under the Create a New Cluster section of each Infrastructure Provider. Mirrors can be used in both air-gapped and non-air-gapped environments by adding the flag to the `dkp create cluster` command.

9.5 Air-gapped Seed the Registry

Before creating an air-gapped Kubernetes cluster, you need to load the required images in a local registry. This registry must be accessible from both the [bastion machine \(see page 1288\)](#) and either the AWS EC2 instances (if deploying to AWS) or other machines that will be created for the Kubernetes cluster.

 If you do not already have a local registry set up, please refer to [Local Registry Tools \(see page 1536\)](#) page for more information.

1. [Download \(see page 78\)](#) `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, and extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address using this command:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

3. Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --
to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-
registry-password $REGISTRY_PASSWORD
```

- It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

You are now ready to create an air-gapped bootstrap cluster for a custom cluster for your [infrastructure provider](#) (see page 1039), or create an air-gapped cluster from the [Day 1 - Basic Installs](#) (see page 147) section for your provider.

9.6 Configure the Control Plane

9.6.1 Prerequisites

Before you begin, make sure you have created your cluster using a bootstrap cluster from the respective [Infrastructure Providers](#) (see page 1039) section.

Users can make modifications to the `KubeadmControlPlane` cluster-api object to configure different kubelet options. Please see the following guide if you wish to configure your control plane beyond the existing options that are available from flags.

9.6.1.1 Modifying Audit Logs

In order to modify control plane option, get the appropriate `cluster-api` objects that describe the cluster by running the following command:

- The following example uses AWS, but can be used for `gcp`, `azure`, `preprovisioned`, and `vsphere` clusters.

```
dkp create cluster aws -c {MY_CLUSTER_NAME} -o yaml --dry-run >>
{MY_CLUSTER_NAME}.yaml
```

When you open `{MY_CLUSTER_NAME}.yaml` with your favorite text editor, look for the `KubeadmControlPlane` object for your cluster. Below is an example object:

```
apiVersion: controlplane.cluster.x-k8s.io/v1beta1
kind: KubeadmControlPlane
metadata:
  name: my-cluster-control-plane
```

```

namespace: default
spec:
  kubeadmConfigSpec:
    clusterConfiguration:
      apiServer:
        extraArgs:
          audit-log-maxage: "30"
          audit-log-maxbackup: "10"
          audit-log-maxsize: "100"
          audit-log-path: /var/log/audit/kube-apiserver-audit.log
          audit-policy-file: /etc/kubernetes/audit-policy/apiserver-audit-policy.yaml
          cloud-provider: aws
          encryption-provider-config: /etc/kubernetes/pki/encryption-config.yaml
        extraVolumes:
          - hostPath: /etc/kubernetes/audit-policy/
            mountPath: /etc/kubernetes/audit-policy/
            name: audit-policy
          - hostPath: /var/log/kubernetes/audit
            mountPath: /var/log/audit/
            name: audit-logs
      controllerManager:
        extraArgs:
          cloud-provider: aws
          configure-cloud-routes: "false"
      dns: {}
      etcd:
        local:
          imageTag: 3.5.6
      networking: {}
      scheduler: {}
    files:
      - content: |
          # Taken from https://github.com/kubernetes/kubernetes/blob/master/cluster/gce/gci/configure-helper.sh
          # Recommended in Kubernetes docs
          apiVersion: audit.k8s.io/v1
          kind: Policy
          rules:
            # The following requests were manually identified as high-volume and low-
            risk,
            # so drop them.
            - level: None
              users: ["system:kube-proxy"]
              verbs: ["watch"]
              resources:
                - group: "" # core
                  resources: ["endpoints", "services", "services/status"]
            - level: None
              # Ingress controller reads 'configmaps/ingress-uid' through the unsecured
              port.
              # TODO(#46983): Change this to the ingress controller service account.
              users: ["system:unsecured"]

```

```

namespaces: ["kube-system"]
verbs: ["get"]
resources:
  - group: "" # core
    resources: ["configmaps"]
- level: None
users: ["kubelet"] # legacy kubelet identity
verbs: ["get"]
resources:
  - group: "" # core
    resources: ["nodes", "nodes/status"]
- level: None
userGroups: ["system:nodes"]
verbs: ["get"]
resources:
  - group: "" # core
    resources: ["nodes", "nodes/status"]
- level: None
users:
  - system:kube-controller-manager
  - system:kube-scheduler
  - system:serviceaccount:kube-system:endpoint-controller
verbs: ["get", "update"]
namespaces: ["kube-system"]
resources:
  - group: "" # core
    resources: ["endpoints"]
- level: None
users: ["system:apiserver"]
verbs: ["get"]
resources:
  - group: "" # core
    resources: ["namespaces", "namespaces/status", "namespaces/finalize"]
- level: None
users: ["cluster-autoscaler"]
verbs: ["get", "update"]
namespaces: ["kube-system"]
resources:
  - group: "" # core
    resources: ["configmaps", "endpoints"]
# Don't log HPA fetching metrics.
- level: None
users:
  - system:kube-controller-manager
verbs: ["get", "list"]
resources:
  - group: "metrics.k8s.io"
# Don't log these read-only URLs.
- level: None
nonResourceURLs:
  - /healthz*
  - /version

```

```

    - /swagger*
# Don't log events requests.
- level: None
  resources:
    - group: "" # core
      resources: ["events"]
# node and pod status calls from nodes are high-volume and can be large,
don't log responses for expected updates from nodes
- level: Request
  users: ["kubelet", "system:node-problem-detector",
"system:serviceaccount:kube-system:node-problem-detector"]
  verbs: ["update","patch"]
  resources:
    - group: "" # core
      resources: ["nodes/status", "pods/status"]
omitStages:
  - "RequestReceived"
- level: Request
  userGroups: ["system:nodes"]
  verbs: ["update","patch"]
  resources:
    - group: "" # core
      resources: ["nodes/status", "pods/status"]
omitStages:
  - "RequestReceived"
# deletecollection calls can be large, don't log responses for expected
namespace deletions
- level: Request
  users: ["system:serviceaccount:kube-system:namespace-controller"]
  verbs: ["deletecollection"]
omitStages:
  - "RequestReceived"
# Secrets, ConfigMaps, and TokenReviews can contain sensitive & binary
data,
# so only log at the Metadata level.
- level: Metadata
  resources:
    - group: "" # core
      resources: ["secrets", "configmaps"]
    - group: authentication.k8s.io
      resources: ["tokenreviews"]
omitStages:
  - "RequestReceived"
# Get responses can be large; skip them.
- level: Request
  verbs: ["get", "list", "watch"]
  resources:
    - group: "" # core
    - group: "admissionregistration.k8s.io"
    - group: "apiextensions.k8s.io"
    - group: "apiregistration.k8s.io"
    - group: "apps"

```



```

- group: "authentication.k8s.io"
- group: "authorization.k8s.io"
- group: "autoscaling"
- group: "batch"
- group: "certificates.k8s.io"
- group: "extensions"
- group: "metrics.k8s.io"
- group: "networking.k8s.io"
- group: "node.k8s.io"
- group: "policy"
- group: "rbac.authorization.k8s.io"
- group: "scheduling.k8s.io"
- group: "settings.k8s.io"
- group: "storage.k8s.io"
omitStages:
  - "RequestReceived"
# Default level for known APIs
- level: RequestResponse
resources:
  - group: "" # core
  - group: "admissionregistration.k8s.io"
  - group: "apiextensions.k8s.io"
  - group: "apiregistration.k8s.io"
  - group: "apps"
  - group: "authentication.k8s.io"
  - group: "authorization.k8s.io"
  - group: "autoscaling"
  - group: "batch"
  - group: "certificates.k8s.io"
  - group: "extensions"
  - group: "metrics.k8s.io"
  - group: "networking.k8s.io"
  - group: "node.k8s.io"
  - group: "policy"
  - group: "rbac.authorization.k8s.io"
  - group: "scheduling.k8s.io"
  - group: "settings.k8s.io"
  - group: "storage.k8s.io"
omitStages:
  - "RequestReceived"
# Default level for all other requests.
- level: Metadata
omitStages:
  - "RequestReceived"
path: /etc/kubernetes/audit-policy/apiserver-audit-policy.yaml
permissions: "0600"
- content: |
  #!/bin/bash
  # CAPI does not expose an API to modify KubeProxyConfiguration
  # this is a workaround to use a script with preKubeadmCommand to modify the
  kubeadm config files
  # https://github.com/kubernetes-sigs/cluster-api/issues/4512

```

```

do
  for i in $(ls /run/kubeadm/ | grep 'kubeadm.yaml\|kubeadm-join-config.yaml');
    cat <<EOF>> "/run/kubeadm//${i}"
    ---
    kind: KubeProxyConfiguration
    apiVersion: kubeproxy.config.k8s.io/v1alpha1
    metricsBindAddress: "0.0.0.0:10249"
    EOF
  done
  path: /run/kubeadm/konvoy-set-kube-proxy-configuration.sh
  permissions: "0700"
- content: |
  [metrics]
    address = "0.0.0.0:1338"
    grpc_histogram = false
  path: /etc/containerd/conf.d/konvoy-metrics.toml
  permissions: "0644"
- content: |
  #!/bin/bash
  systemctl restart containerd

  SECONDS=0
  until crictl info
  do
    if (( SECONDS > 60 ))
    then
      echo "Containerd is not running. Giving up..."
      exit 1
    fi
    echo "Containerd is not running yet. Waiting..."
    sleep 5
  done
  path: /run/konvoy/restart-containerd-and-wait.sh
  permissions: "0700"
- contentFrom:
  secret:
    key: value
    name: my-cluster-etcd-encryption-config
  owner: root:root
  path: /etc/kubernetes/pki/encryption-config.yaml
  permissions: "0640"
format: cloud-config
initConfiguration:
  localAPIEndpoint: {}
  nodeRegistration:
    kubeletExtraArgs:
      cloud-provider: aws
    name: '{{ ds.meta_data.local_hostname }}'
joinConfiguration:
  discovery: {}
  nodeRegistration:
    kubeletExtraArgs:

```

```

    cloud-provider: aws
    name: '{{ ds.meta_data.local_hostname }}'
preKubeadmCommands:
- systemctl daemon-reload
- /run/konvoy/restart-containerd-and-wait.sh
- /run/kubeadm/konvoy-set-kube-proxy-configuration.sh
machineTemplate:
  infrastructureRef:
    apiVersion: infrastructure.cluster.x-k8s.io/v1beta1
    kind: AWSMachineTemplate
    name: my-cluster-control-plane
    namespace: default
  metadata: {}
  replicas: 3
  rolloutStrategy:
    rollingUpdate:
      maxSurge: 1
    type: RollingUpdate
  version: v1.26.6

```



If you use the above example as-is, ensure you update your Kubereters version number on the final line by replacing the x.

Now a user can configure the fields below for the log backend. The log backend will write audit events to a file in [JSON](#)¹²⁵² format. You can configure the log audit backend using the `kube-apiserver` flags shown below:

```

audit-log-maxage
audit-log-maxbackup
audit-log-maxsize
audit-log-path

```



See [upstream documentation](#)¹²⁵³ for more information.

After modifying the values appropriately, you can create the cluster by running the command below:

```
kubectl create -f {MY_CLUSTER_NAME}.yaml
```

¹²⁵² <https://jsonlines.org/>

¹²⁵³ <https://kubernetes.io/docs/tasks/debug/debug-cluster/audit/#log-backend>

Once the cluster is created, users can get the corresponding kubeconfig for the cluster by running the following command:

```
dkp get kubeconfig -c {MY_CLUSTER_NAME} >> {MY_CLUSTER_NAME}.conf
```

9.6.1.2 Viewing the Audit Logs

Fluent Bit is disabled on the management cluster by default, to view the audit logs run the command below:

```
dkp diagnose --kubeconfig={MY_CLUSTER_NAME}.conf
```

A file similar to `support-bundle-2022-08-15T02_28_48.tar.gz` will be created. Untar the file using a command similar to the example below:

```
tar -xzf support-bundle-2022-08-15T02_28_48.tar.gz
```

Navigate to the node-diagnostics sub directory from the extracted file with a command like the one shown below:

```
cd support-bundle-2022-08-15T02_28_48/node-diagnostics
```

Finally, to find the audit logs run the following command:

```
$ find . -type f | grep audit.log
./ip-10-0-142-117.us-west-2.compute.internal/data/kube_apiserver_audit.log
./ip-10-0-148-139.us-west-2.compute.internal/data/kube_apiserver_audit.log
./ip-10-0-128-181.us-west-2.compute.internal/data/kube_apiserver_audit.log
```

See other related pages below:

[Fluent bit](#) (see page 926)

9.7 Update Cluster Nodepools

Upgrading a node pool involves draining the existing nodes in the node pool and replacing them with new nodes. In order to ensure minimum downtime and maintain high availability of the critical application workloads during the upgrade process, we recommend deploying Pod Disruption Budget ([Disruptions](#)¹²⁵⁴) for your critical applications.

¹²⁵⁴ <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>


The Pod Disruption Budget will prevent any impact on critical applications as a result of misconfiguration or failures during the upgrade process.

9.7.1 Prerequisites:

- Deploy [Pod Disruption Budget](#)¹²⁵⁵ (PDB)
- [Konvoy Image Builder](#) (see page 1464) (KIB)

9.7.2 Steps:

1. Deploy Pod Disruption Budget for your critical applications. If your application can tolerate only one replica to be unavailable at a time, then you can set Pod disruption budget as shown in the following example. The example below is for NVIDIA GPU node pools, but the process is the same for all node pools.

 Repeat this step for each additional node pool.

Create the file: `pod-disruption-budget-nvidia.yaml`

```
apiVersion: policy/v1
kind: PodDisruptionBudget
metadata:
  name: nvidia-critical-app
spec:
  maxUnavailable: 1
  selector:
    matchLabels:
      app: nvidia-critical-app
```

Apply the YAML file above using the following command:

```
kubectl create -f pod-disruption-budget-nvidia.yaml
```

2. Prepare OS image for your node pool using [Konvoy Image Builder](#) (see page 1464).

9.7.3 Related Topics:

[Upgrade Konvoy](#) (see page 1588)

¹²⁵⁵ <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

9.8 Verify your Cluster and DKP Installation

Check DKP components to verify the status of your cluster

This section contains information on how to verify a DKP installation.

9.8.1 Check the Cluster Infrastructure and Nodes

DKP ships with some default diagnosis tools to check your cluster, such as the `describe` command. You can use those tools to validate your installation.

If you have not done so already, set the environment variable for your cluster name, substituting `dkp-example` with the name of your cluster:

```
export CLUSTER_NAME=dkp-example
```

Then, run this command to check the health of the cluster infrastructure:

```
dkp describe cluster --cluster-name=${CLUSTER_NAME}
```



For more details on the `dkp describe cluster` command, refer to [dkp describe cluster \(see page 1701\)](#).

A healthy cluster returns an output similar to:

```
NAME                                     READY  SEVERITY
REASON  SINCE  MESSAGE
Cluster/dkp-example                       True
121m
├─ClusterInfrastructure - AWSCluster/dkp-example   True
121m
├─ControlPlane - KubeadmControlPlane/dkp-example-control-plane True
121m
| └─Machine/dkp-example-control-plane-h52t6       True
121m
| └─Machine/dkp-example-control-plane-knrh        True
121m
| └─Machine/dkp-example-control-plane-zmjyx       True
121m
└─Workers
```

```

└─MachineDeployment/dkp-example-md-0           True
121m
  └─Machine/dkp-example-md-0-88488cb74-2vxjq   True
121m
  └─Machine/dkp-example-md-0-88488cb74-84xsd   True
121m
  └─Machine/dkp-example-md-0-88488cb74-9xmc6   True
121m
    └─Machine/dkp-example-md-0-88488cb74-mjf6s True
121m

```

Use this `kubectl` command to check to see if all cluster nodes are ready:

```
kubectl get nodes
```

The output should look similar to this, with all statuses set to `Ready`

NAME	STATUS	ROLES	AGE
VERSION			
ip-10-0-112-116.us-west-2.compute.internal v1.21.6	Ready	<none>	135m
ip-10-0-122-142.us-west-2.compute.internal v1.21.6	Ready	<none>	135m
ip-10-0-186-214.us-west-2.compute.internal v1.21.6	Ready	control-plane,master	133m
ip-10-0-231-82.us-west-2.compute.internal v1.21.6	Ready	control-plane,master	135m
ip-10-0-71-114.us-west-2.compute.internal v1.21.6	Ready	<none>	135m
ip-10-0-71-207.us-west-2.compute.internal v1.21.6	Ready	<none>	135m
ip-10-0-85-253.us-west-2.compute.internal v1.21.6	Ready	control-plane,master	137m

To verify a successful installation, all of the previous commands should return as `Ready` or `True`.

9.8.2 Monitor the CAPI resources

Obtain a list of all resources that comprise a cluster and their statuses:

```
kubectl get cluster-api
```

9.8.3 Verify all Pods

All pods installed by DKP should be in `Running` or `Completed` status if the install is successful.

```
kubectll get pods --all-namespaces
```

9.8.4 Troubleshooting

If any pod is not in `Running` or `Completed` status, you need to investigate further. If it appears that something has not deployed properly or fully, run a [diagnostic bundle](#) (see page 1023):

```
dkp diagnose
```

This collects information from pods and infrastructure. For more information, see [more about generating a support bundle and the different configurations that it supports](#) (see page 1023).

9.9 GPU for Konvoy

In this version of DKP, the nodes with NVIDIA GPUs are configured with `nvidia-gpu-operator` ([Overview – NVIDIA Cloud Native Technologies documentation](#)¹²⁵⁶) and NVIDIA drivers to support the container runtime.

This page will link you to all the relevant GPU pages necessary such as [Updating Cluster Nodepools](#) (see page 1548) or [KIB for GPU](#) (see page 1488). The remainder of [GPU information](#) (see page 971) is found in Day 2 Operations under Kommander component section of documentation.

9.10 Delete a DKP Cluster with One Command

You can use a single command line entry to delete a Kubernetes cluster on any of the platforms supported by DKP. Deleting a cluster means removing the cluster, all of its nodes, and all of the platform applications that were deployed on it as part of its creation. Use this command with extreme care, as it is not reversible!



You need to delete the attachment for any clusters attached in Kommander before running the `delete` command.



Persistent Volumes (PVs) are not deleted automatically by design in order to preserve your data. However, they take up storage space if not deleted. You must delete PVs manually. Information for backup of a cluster and PVs is on the page in documentation called [Back up your Cluster's Applications and Persistent Volumes](#) (see page 835) .

¹²⁵⁶ <https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/overview.html>

With Vsphere clusters, `dkp delete` doesn't delete the virtual disks backing the PVs for DKP add ons. Therefore internal VMware cluster runs out of storage eventually. These PVs are only visible if VSAN is installed which gives users a Container Native Storage tab.

Set the environment variable to be used throughout this documentation:

```
export CLUSTER_NAME=cluster-example
```

The basic DKP `delete` command structure is:

```
dkp delete cluster --cluster-name=${CLUSTER_NAME} --self-managed --kubeconfig=${CLUSTER_NAME}.conf
```

When you use the `--self-managed` flag, the prerequisite components and resources are moved from the self-managed cluster before deleting. When you omit this flag (the default value is false) the resources are assumed to be installed in a management cluster. The default value is false, or no flag.

This command performs the following actions:

- Creates a local bootstrap cluster
- Moves controllers to it
- Deletes the management cluster
- Deletes the local bootstrap cluster

10 Upgrade DKP

DKP Upgrade Path

The DKP upgrade represents an important step of your environment’s lifecycle, as it ensures that you are up-to-date with the latest features and can benefit from the most recent improvements, enhanced cluster management, and better performance. This section describes how to upgrade your air-gapped and non-air-gapped environment to the latest version of DKP compatible with the latest Kubernetes version.

10.1 Prerequisite

Check what version of DKP you have downloaded currently using cli command `dkp version` (see page 1749).

```
dkp version
```

10.1.1 Supported Upgrade Paths:

Use this table to determine your correct upgrade path:

		Upgrading from Release...						
...to Release		2.4.0	2.4.1	2.5.0	2.5.1	2.5.2	2.6.0	2.6.1
	2.4.0	NA	NA	NA	NA	NA	NA	NA
	2.4.1	Yes	NA	NA	NA	NA	NA	NA
	2.5.0	Yes	Yes	NA	NA	NA	NA	NA
	2.5.1	Yes	Yes	Yes	NA	NA	NA	NA
	2.5.2	Yes	Yes	Yes	Yes	NA	NA	NA
	2.6.0	No	No	Yes	Yes	Yes	NA	NA
	2.6.1	No	No	Yes	Yes	Yes	Yes	NA
	2.6.2	No	No	Yes	Yes	Yes	Yes	Yes

10.1.2 Upgrade Order

Perform the upgrade sequentially, beginning with the Kommander component and then moving to upgrading clusters and CAPI components in the Konvoy component.

When upgrading your entire DKP product, the process is different depending on your environment.

- **Essential** (see page 1588): A [stand-alone Management Cluster](#) (see page 83)
- **Enterprise** (see page 1562): A [multi-cluster environment](#) (see page 81) that includes a combination of a Management cluster(s) and Managed or Attached workspace clusters.



- Before upgrading, we **strongly recommend** reading the [release notes](#) (see page 1752) and verifying your current setup against any possible breaking changes.
- Review the [list of major Kubernetes changes](#) (see page 1759) that may affect your system.

10.1.3 Next Step

[Upgrade Compatibility Tables](#) (see page 1552)

10.2 Upgrade Compatibility Tables

The latest DKP versions for Kubernetes, applications, and components are available on this page. Individual sections will have links to pages with further information if required.

10.2.1 Supported Operating Systems

The supported Operating System (OS) is important for upgrades. During an upgrade, you would need to create new AMI's or images with the Kubernetes version to which you want to upgrade. That requires selecting an OS and ensuring you have the right version. Check the [Supported Infrastructure Operating Systems](#) (see page 69) page if you have not already done so.

10.2.2 Konvoy Image Builder

Refer to the [main page](#) (see page 1464) for a table of components versions and compatibility.

10.2.3 Kubernetes Version Supported

These are the supported versions:

10.2.3.1 Deploying Clusters Versions

The latest supported Kubernetes versions for DKP are listed below. In DKP 2.6.2, the Kubernetes version installed by default is 1.26.14. The newest and oldest Kubernetes versions used with DKP must be within one minor version.

- Latest supported Kubernetes version is at **1.26.14 for deploying clusters**
- Other Kubernetes versions are supported at **1.25.13 for deploying clusters in EKS**

Kubernetes 1.26.14 enables you to benefit from the latest features and security fixes in upstream Kubernetes. This release comes with approximately 40 enhancements that you can benefit from like [Service internal tracking policy](#)¹²⁵⁷, [Kubelet credential provider](#)¹²⁵⁸, among others. To read more about major features in this release, visit [GitHub](#)¹²⁵⁹ and <https://kubernetes.io/blog/2022/12/09/kubernetes-v1-26-release/>.

10.2.3.2 Attaching Clusters Versions

DKP 2.6.2 supports attaching clusters with the following Kubernetes versions:

Product	Compatible Kubernetes Versions
EKS	1.25.13
AKS	1.26.10
GKE	1.26.x



Attaching Kubernetes clusters with versions greater than n-1 is not recommended.



When attempting to attach a cluster with a lower Kubernetes version than the DKP default, you need to build and use an image with that lower version of Kubernetes. See [Konvoy Image Builder](#) (see page 1464) for documentation on building images. Failure to do this could result in an error.

¹²⁵⁷ <https://github.com/kubernetes/enhancements/issues/2086>

¹²⁵⁸ <https://github.com/kubernetes/enhancements/issues/2133>

¹²⁵⁹ <https://github.com/kubernetes/sig-release/blob/master/releases/release-1.26/README.md>

10.2.4 Next Step:

[Upgrade Prerequisites](#) (see page 1555)

10.2.5 Upgrade Cluster Node Pools

Upgrading a node pool involves draining the existing nodes in the node pool and replacing them with new nodes. In order to ensure minimum downtime and maintain high availability of the critical application workloads during the upgrade process, we recommend deploying Pod Disruption Budget ([Disruptions](#)¹²⁶⁰) for your critical applications.


The Pod Disruption Budget will prevent any impact on critical applications as a result of misconfiguration or failures during the upgrade process.

10.2.5.1 Prerequisites:

- Deploy [Pod Disruption Budget](#)¹²⁶¹ (PDB)
- [Konvoy Image Builder](#) (see page 1464) (KIB)

10.2.5.2 Steps:

1. Deploy Pod Disruption Budget for your critical applications. If your application can tolerate only one replica to be unavailable at a time, then you can set Pod disruption budget as shown in the following example. The example below is for NVIDIA GPU node pools, but the process is the same for all node pools.

 Repeat this step for each additional node pool.

Create the file: `pod-disruption-budget-nvidia.yaml`

```
apiVersion: policy/v1
kind: PodDisruptionBudget
metadata:
  name: nvidia-critical-app
spec:
  maxUnavailable: 1
  selector:
```

¹²⁶⁰ <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

¹²⁶¹ <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

```
matchLabels:
  app: nvidia-critical-app
```

Apply the YAML file above using the following command:

```
kubectl create -f pod-disruption-budget-nvidia.yaml
```

2. Prepare OS image for your node pool using [Konvoy Image Builder](#) (see page 1464).

10.2.5.3 Related Topics:

[Enterprise: Upgrade the Management Cluster Kubernetes Version](#) (see page 1580)

10.3 Upgrade Prerequisites

The prerequisites vary depending on your environment. Ensure you follow the correct lists for your specific configurations:

- [Prerequisites for the Kommander Component](#) (see page 1555)
 - [All Kommander Environments](#) (see page 1555)
 - [Air-gapped Kommander Environments Only](#) (see page 1556)
 - [DKP Insights Only](#) (see page 1556)
 - [DKP Workspace Catalog Applications Only](#) (see page 1557)
- [Prerequisites for the Konvoy Component](#) (see page 1557)
 - [All Konvoy Environments](#) (see page 1557)
 - [Air-gapped Konvoy Environments Only](#) (see page 1558)

10.3.1 Prerequisites for the Kommander Component

The prerequisites for Kommander consists of two parts, the first section for all environments, and the subsections depend on air-gapped or non-air-gapped, and optional applications. Ensure you follow the correct section for your specific environment type.

10.3.1.1 All Kommander Environments

For all Kommander environments, regardless of network type and applications, complete these prerequisites:

- Before upgrading, we **strongly recommend** reading the [release notes](#) (see page 1752) and verifying your current setup against any possible breaking changes.
- Review the [DKP 2.6.0 Components and Applications](#) (see page 1760) versions that are part of this upgrade.
- Ensure your applications you are running will be compatible with Kubernetes v1.26.x.

- If you have attached clusters, ensure your applications you are running will be compatible with Kubernetes v1.26.x.
- Review the [list of major Kubernetes changes \(see page 1759\)](#) that may affect your system.
- Ensure you are attempting to run a [supported DKP upgrade \(see page 0\)](#).
- **REQUIRED** Before upgrading, create an [on-demand backup \(see page 877\)](#) of your current configuration with Velero.
- [Download \(see page 78\)](#) and install the supported DKP CLI binary of [this release \(see page 1752\)](#) on your computer.

The **remaining prerequisites** are necessary if you have one of the following environments or additions to basic Kommander. Otherwise, proceed to the prerequisites for the Konvoy component.

10.3.1.1.1 Air-gapped Kommander Environments Only Additional Prerequisites

If you operate Kommander in an air-gapped environment, complete these prerequisites in addition to the ones above:

- There are several [sets of images \(see page 1558\)](#) you will need to push to your local registry:
 - The Kommander component - kommander-image-bundle-v2.6.1.tar
 - The Konvoy component for Kubernetes cluster - konvoy-image-bundle-v2.6.1.tar
 - **Optional:** Catalog Applications (Enterprise only) - dkp-catalog-applications-image-bundle-v2.6.1.tar
- **For air-gapped environments with catalog applications:** Ensure you have updated your catalog repository before upgrading. The catalog repository should contain the Docker registry with the [DKP catalog application \(see page 677\)](#) images and a charts bundle file containing [DKP catalog applications \(see page 0\)](#) charts.

10.3.1.1.2 DKP Insights Only

10.3.1.1.2.1 For 2.5.x > 2.6.0 Upgrades

Ensure you uninstall DKP Insights before upgrading DKP. See [\(2.6\) DKP Insights Upgrade from DKP 2.5.x to 2.6.0¹²⁶²](#) for more information.



The **Insights Management** and **Engine** are not able to communicate with each other until they are upgraded to the same DKP Insights version. D2iQ recommends not resolving Insight alerts until the full DKP Insights upgrade has been completed.

¹²⁶² <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274989271/%282.6%29+DKP+Insights+Upgrade+from+DKP+2.5.x+to+2.6.0>

10.3.1.1.3 DKP Workspace Catalog Applications Only

If you use any of the DKP Workspace Catalog Applications, complete these prerequisites:

- Ensure your clusters run on a supported Kubernetes version for [this release of DKP \(see page 1752\)](#), and that said Kubernetes version is also compatible with your [catalog application version \(see page 0\)](#).
- For customers with an [Enterprise license \(see page 129\)](#) and a multi-cluster environment, we recommend keeping all clusters on the same Kubernetes version. This ensures your DKP catalog application can run on all clusters in a given workspace.
- Konvoy installs a new Kubernetes version with each upgrade, therefore it is important that your [Catalog application \(see page 667\)](#) is compatible with the Kubernetes version that comes with this release of DKP.



If your current Catalog application version is not compatible with this release's Kubernetes version, upgrade the application to a compatible version BEFORE upgrading the Konvoy component on [Managed clusters \(see page 82\)](#) or BEFORE upgrading the Kubernetes version on [Attached clusters \(see page 82\)](#).

- **For air-gapped environments with catalog applications:** Ensure you have updated your catalog repository before upgrading. The catalog repository should contain the Docker registry with the [DKP catalog application \(see page 677\)](#) images and a charts bundle file containing [DKP catalog applications \(see page 0\)](#) charts.

10.3.2 Prerequisites for the Konvoy Component

The prerequisites for Konvoy also consists of two parts, the first section for all environments, and the subsections depend on your network type, air-gapped or non-air-gapped, and optional applications. Ensure you follow the correct section for your specific environment type.

10.3.2.1 All Konvoy Environments

For all Konvoy environments, regardless of network type and applications, complete these prerequisites:

- Check what version of DKP you have downloaded using the CLI command [dkp version \(see page 1749\)](#).
- Set the appropriate environment variables:
 - For Air-gapped, download the required bundles at our [support site](#)¹²⁶³.
 - For Azure, set the required [environment variables \(see page 1159\)](#).
 - For [AWS and EKS \(see page 1060\)](#), set the required [environment variables \(see page 295\)](#).
 - For vSphere, set the required [environment variables \(see page 1259\)](#).

¹²⁶³ <https://support.d2iq.com/hc/en-us>

- For GCP, set the required [environment variables](#) (see page 1368).
- **vSphere only:** If you want to resize your disk, ensure you have reviewed [Create a vSphere Base OS Image](#) (see page 1490).

10.3.2.1.1 Air-gapped Konvoy Environments Only Additional Prerequisites

If you operate Konvoy in an air-gapped environment, complete these prerequisites in addition to the ones listed above:

- Check what version of DKP you have downloaded currently using cli command `dkp version`¹²⁶⁴.
- Ensure that all platform applications in the management cluster have been upgraded to avoid compatibility issues with the [Kubernetes version](#) (see page 1758) included in this release. This is done automatically when [upgrading Kommander](#) (see page 1558), so ensure that you upgrade Kommander prior to upgrading Konvoy.
- For air-gapped environments, seed the registry as explained here: [Upgrade: For Air-gapped Environments Only](#) (see page 1558)

10.3.3 Third-Party and Open Source Attributions

See <https://d2iq.com/legal/3rd> for a list of third-party and open source attributions.

10.3.3.1 Next Steps:

Depending on your license type, you will follow the relevant link.

- [Upgrade DKP Enterprise](#) (see page 1562)
- [Upgrade DKP Essential](#) (see page 1588)

10.4 Upgrade: For Air-gapped Environments Only

Because air-gapped environments do not have direct access to the Internet, you must download, extract and load several required images to your [local container registry](#) (see page 1536), before installing or upgrading DKP. The information below will be covered as a step during either **Enterprise** or **Essential** Upgrade steps, but feel free to familiarize yourself with the concept below if desired. Otherwise, depending on your license type, follow the relevant link to begin upgrading:

- [Upgrade DKP Enterprise](#) (see page 1562)
- [Upgrade DKP Essential](#) (see page 1588)

10.4.1 Overview of Seeding the Registry for Air-gapped Environment

¹²⁶⁴ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/213157688>

10.4.2 Download all Images for Air-gapped Deployments

If you are operating in an air-gapped environment, a [local container registry](#) (see page 1536) containing all the necessary installation images, including the Kommander images is required. See below for prerequisites to download and then how to push the necessary images to this registry.

1. [Download](#) (see page 78) the Complete DKP Air-gapped Bundle for this release (i.e. `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`) to load registry images as explained below.
2. Connectivity with clusters attaching to the management cluster is required:
 - Both management and attached clusters must be able to connect to the local registry.
 - The management cluster must be able to connect to all attached cluster's API servers.
 - The management cluster must be able to connect to any load balancers created for platform services on the management cluster.

10.4.3 Extract Air-gapped Images and Set Variables

Follow these steps to **extract** the air-gapped image bundles into your private registry:

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

10.4.4 Only Pre-provisioned: Load Images for Deployments - Konvoy

For **Pre-provisioned air-gapped environments only**, you must run `konvoy-image upload artifacts` to [copy the artifacts](#) (see page 1207) onto the cluster hosts **before** you begin the [Upgrade the CAPI Components](#) (see page 1576) process later in the upgrade steps.

1. The Kubernetes image bundle will be located in `kib/artifacts/images` and you will want to verify image and artifacts.
 - a. Verify the image bundles exist in `kib/artifacts/images` :

```
$ ls kib/artifacts/images/
kubernetes-images-1.26.14-d2iq.1.tar kubernetes-images-1.26.14-d2iq.1-
fips.tar
```

- b. Verify the artifacts for your OS exist in the `kib/artifacts/` directory and export the appropriate variables:

```
$ ls kib/artifacts/
1.26.14_centos_7_x86_64.tar.gz      1.26.14_redhat_8_x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rhel-8.6-x86_64_fips.tar.gz  pip-packages.tar.gz
1.26.14_centos_7_x86_64_fips.tar.gz  1.26.14_rocky_9_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-7.9-x86_64_fips.tar.gz  containerd-1.6.17-
d2iq.1-rocky-9.0-x86_64.tar.gz
1.26.14_redhat_7_x86_64.tar.gz      1.26.14_ubuntu_20_x86_64.tar.gz
containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz  containerd-1.6.17-
d2iq.1-rocky-9.1-x86_64.tar.gz
1.26.14_redhat_7_x86_64_fips.tar.gz  containerd-1.6.17-d2iq.1-centos-7.9-
x86_64.tar.gz      containerd-1.6.17-d2iq.1-rhel-8.4-x86_64_fips.tar.gz
containerd-1.6.17-d2iq.1-ubuntu-20.04-x86_64.tar.gz
1.26.14_redhat_8_x86_64.tar.gz      containerd-1.6.17-d2iq.1-centos-7.9-
x86_64_fips.tar.gz  containerd-1.6.17-d2iq.1-rhel-8.6-x86_64.tar.gz
images
```

- c. Set the bundle values with the name from the private registry location:

```
export OS_PACKAGES_BUNDLE=<name_of_the_OS_package>
export CONTAINERD_BUNDLE=<name_of_the_containerd_bundle>
```

For example, for RHEL 8.4 you would set:

```
export OS_PACKAGES_BUNDLE=1.26.14_redhat_8_x86_64.tar.gz
export CONTAINERD_BUNDLE=containerd-1.6.17-d2iq.1-rhel-8.4-x86_64.tar.gz
```

2. Copy the artifacts:

```
konvoy-image upload artifacts \
  --container-images-dir=./kib/artifacts/images/ \
  --os-packages-bundle=./kib/artifacts/$OS_PACKAGES_BUNDLE \
  --containerd-bundle=./kib/artifacts/$CONTAINERD_BUNDLE \
  --pip-packages-bundle=./kib/artifacts/pip-packages.tar.gz
```

10.4.5 Load Images to your Private Registry - Konvoy

Before creating or upgrading a Kubernetes cluster, you need to load the required images in a local registry if operating in an air-gapped environment. This registry must be accessible from both the [bastion machine \(see page 1288\)](#) and either the AWS EC2 instances or other machines that will be created for the Kubernetes cluster.



If you do not already have a local registry set up, refer to [Local Registry Tools \(see page 1536\)](#) page for more information.

Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-registry-password $REGISTRY_PASSWORD
```



It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

10.4.6 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

10.4.7 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-  
v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

10.4.8 Next Step:

Depending on your license type, you will follow the relevant link:


- [Upgrade DKP Enterprise \(see page 1562\)](#)
- [Upgrade DKP Essential \(see page 1588\)](#)

10.5 Upgrade DKP Enterprise

Enterprise

Gov Advanced

10.5.1 Steps to upgrade Konvoy via CLI

 Before you begin the upgrade, ensure you review the [Upgrade Prerequisites \(see page 1558\)](#).

This section describes how to upgrade your DKP clusters to the latest DKP version in an Enterprise, multi-cluster, environment. The DKP upgrade represents an important step of your environment's lifecycle, as it ensures that you are up-to-date with the latest features and can benefit from the most recent improvements, enhanced cluster management, and better performance. If you are on this scenario of the upgrade section of documentation, you have an [Enterprise License \(see page 125\)](#). For an [Essential \(see page 125\)](#) license, refer to the [DKP Upgrade Essential \(see page 1588\)](#) section instead.



Ensure that all platform applications in the management cluster have been upgraded to avoid compatibility issues with the [Kubernetes version \(see page 1555\)](#) included in this release. This is done automatically when upgrading Kommander, so ensure that you follow these sections in order to upgrade the Kommander component prior to upgrading the Konvoy component.

10.5.2 Overview of Steps in Order

- [Enterprise: For Air-gapped Environments Only](#) (see page 1563)
- [Enterprise: Upgrade the Management Cluster and Platform Applications](#) (see page 1565)
- [Enterprise: Upgrade Platform Applications on Managed and Attached Clusters](#) (see page 1569)
- [Enterprise: Upgrade Workspace Catalog Applications](#) (see page 1571)
- [Enterprise: Upgrade Project Catalog Applications](#) (see page 1574)
- [Enterprise: Upgrade Custom Applications](#) (see page 1576)
- [Enterprise: Upgrade the Management Cluster CAPI Components](#) (see page 1576)
- [Enterprise: Upgrade the Management Cluster Core Addons](#) (see page 1578)
- [Enterprise: Upgrade the Management Cluster Kubernetes Version](#) (see page 1580)
- [Enterprise: Upgrade Managed Clusters](#) (see page 1585)

10.5.3 Next Step:

If you are in an air-gapped environment, be sure to select that choice to find out how to load your local registry first. Otherwise, your next step will be [Upgrade the Management Cluster and Platform Applications](#) (see page 1565).

- [Enterprise: For Air-gapped Environments Only](#) (see page 1563)

10.5.4 Enterprise: For Air-gapped Environments Only

Enterprise

Gov Advanced

If you are operating in an air-gapped environment, a [local container registry](#) (see page 1536) containing all the necessary installation images, including the Kommander images is required. See below for how to push the necessary images to this registry.

- [Download](#) (see page 78) the Complete DKP Air-gapped Bundle for this release (i.e. `dkp-air-gapped-bundle_v2.6.1_linux_amd64.tar.gz`)
- Connectivity with clusters attaching to the management cluster:
 - Both management and attached clusters must be able to connect to the local registry.
 - The management cluster must be able to connect to all attached cluster's API servers.
 - The management cluster must be able to connect to any load balancers created for platform services on the attached cluster.



Note: You can also attach clusters when there are networking restrictions between the management cluster and attached cluster. See [Attach a Cluster with Networking Restrictions](#) (see page 775) for more information.

10.5.4.1 Extract Air-gapped Images and Set Variables

Follow these steps to **extract** the air-gapped image bundles into your private registry:

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

10.5.4.2 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

10.5.4.3 Load Images to your Private Registry - DKP Catalog Applications

Optional: This step is required only if you have an Enterprise license.

For **DKP Catalog Applications**, also perform this image load:

Run the following command to load the `dkp-catalog-applications` image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/dkp-catalog-applications-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

10.5.4.4 Load Images to your Private Registry - Konvoy

Before creating or upgrading a Kubernetes cluster, you need to load the required images in a local registry if operating in an air-gapped environment. This registry must be accessible from both the [bastion machine](#) (see

[page 1288](#)) and either the AWS EC2 instances or other machines that will be created for the Kubernetes cluster.



If you do not already have a local registry set up, refer to [Local Registry Tools \(see page 1536\)](#) page for more information.

Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-registry-password $REGISTRY_PASSWORD
```



It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

10.5.4.5 Next Step:

[Enterprise: Upgrade the Management Cluster and Platform Applications \(see page 1565\)](#)

10.5.5 Enterprise: Upgrade the Management Cluster and Platform Applications

Enterprise

Gov Advanced

This section describes how to upgrade your Kommander Management cluster and all Platform Applications to their supported versions in air-gapped, non-air-gapped and on-prem environments. To prevent compatibility issues, you must first upgrade Kommander on your Management Cluster before upgrading to DKP.



It is important you upgrade Kommander BEFORE upgrading the Kubernetes version (or Konvoy version for Managed Konvoy clusters) in attached clusters. This ensures that any changes required for new or changed Kubernetes API's are already present.

10.5.5.1 Upgrade Kommander

10.5.5.1.1 Prerequisites

- Use the `--kubeconfig=${CLUSTER_NAME}.conf` flag or set the `KUBECONFIG` environment variable to ensure that you **upgrade Kommander on the right cluster**. For alternatives and recommendations around setting your context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- If you have DKP Insights installed, ensure you **uninstall** it before upgrading DKP. See [\(2.6\) DKP Insights Upgrade from DKP 2.5.x to 2.6.0](#)¹²⁶⁵ for more information.



Limited availability of DKP resources

- If you have configured a **custom domain**, running the `upgrade` command could result in an inaccessibility of your services via your custom domain for a few minutes.
- The **DKP UI** and other APIs may be inconsistent or unavailable until the upgrade is complete.

Use the DKP CLI to upgrade Kommander and all the Platform Applications in the Management Cluster:

For non-air-gapped environments

```
dkp upgrade kommander
```



If you want to disable the AI Navigator, add the flag `--disable-appdeployments ai-navigator-app` to this command.

The output looks similar to this:

```
✓ Ensuring upgrading conditions are met
✓ Ensuring application definitions are updated
✓ Ensuring helm-mirror implementation is migrated to ChartMuseum
...
```

¹²⁶⁵ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274989271/%282.6%29+DKP+Insights+Upgrade+from+DKP+2.5.x+to+2.6.0>

After the upgrade, if you have DKP Catalog Applications deployed, update the GitRepository for `dkp-catalog-applications`.

Update the GitRepository with the tag of your updated DKP version on the `kommander` workspace:

```
kubectl patch gitrepository -n kommander dkp-catalog-applications --type merge --
patch '{"spec":{"ref":{"tag":"v2.6.2"}}}'
```

 This command updates the catalog application repositories for all workspaces.

For any additional workspaces created outside the `kommander.yaml` configuration:

Set the `WORKSPACE_NAMESPACE` environment variable to the namespace of the workspace:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```

Update the GitRepository for the additional workspace:

```
kubectl patch gitrepository -n ${WORKSPACE_NAMESPACE} dkp-catalog-applications --type
merge --patch '{"spec":{"ref":{"tag":"v2.6.2"}}}'
```

For air-gapped environments

```
dkp upgrade kommander \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.2.tar.gz \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.2.tar.gz
```

The output looks similar to this:

```
✓ Ensuring upgrading conditions are met
✓ Ensuring application definitions are updated
✓ Ensuring helm-mirror implementation is migrated to ChartMuseum
...
```

For air-gapped environments with DKP Catalog Applications


```
dkp upgrade kommander \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.2.tar.gz \
```

```
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-v2.6.2.tar.gz \
--kommander-applications-repository ./application-repositories/kommander-applications-v2.6.2.tar.gz
```

The output looks similar to this:

```
✓ Ensuring upgrading conditions are met
✓ Ensuring application definitions are updated
✓ Ensuring helm-mirror implementation is migrated to ChartMuseum
...
```


After the upgrade, if you have DKP Catalog Applications deployed, update the GitRepository for `dkp-catalog-applications`.

 For the following section, ensure you modify the **most recent** `kommander.yaml` configuration file. It must be the file that reflects the current state of your environment. Reinstalling Kommander with an outdated `kommander.yaml` overwrites the list of platform applications that are currently running in your cluster.

In the `kommander.yaml` you are currently using for your environment, update the DKP Catalog Applications by setting the correct DKP version:

```
...
# The list of enabled/disabled apps here should reflect the current state of the
environment, including configuration overrides!
...
catalog:
  repositories:
    - name: dkp-catalog-applications
      labels:
        kommander.d2iq.io/project-default-catalog-repository: "true"
        kommander.d2iq.io/workspace-default-catalog-repository: "true"
        kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
      path: ./dkp-catalog-applications-v2.6.2.tar.gz # modify this version to match
the DKP upgrade version
...
```

Refresh the `kommander.yaml` to apply the updated tarball:

 Ensure the `kommander.yaml` is the Kommander Installer Configuration file you are currently using for your environment. Otherwise, your configuration will be overwritten and previous configuration lost.

```
dkp install kommander --installer-config kommander.yaml
```

10.5.5.1.1 Troubleshooting

1. If the upgrade fails, run the following command to get more information on the upgrade process:

```
dkp upgrade kommander -v 6
```

2. If you find any `HelMReLeases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelMReLease` using the following commands:

```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": false}]'
```

10.5.5.1.2 Next Step:

<https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273748688/2.6%2BEnterprise%2BUpgrade%2BPlatform%2BApplications%2Bon%2BManaged%2Band%2BAttached%2BClusters> (see page 1569)

10.5.6 Enterprise: Upgrade Platform Applications on Managed and Attached Clusters

Enterprise

Gov Advanced

For the [Management Cluster](#) (see page 82) (and the `kommander` workspace namespace), `dkp upgrade` handles all Platform applications; no other steps are necessary. However, for [managed](#) (see page 82) or [attached clusters](#) (see page 82), you MUST manually upgrade the Platform applications by upgrading the workspace.



If you are upgrading your Platform applications as part of the [DKP upgrade](#) (see page 1554), upgrade your Platform applications on any additional Workspaces before proceeding with the Konvoy upgrade. Some applications in the previous release are not compatible with the [Kubernetes version](#) (see page 1758) of this release, and upgrading Kubernetes is part of the DKP Konvoy upgrade process.

10.5.6.1 Prerequisites

Before you begin, you must:

- Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace where the cluster is attached:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

- Set the `WORKSPACE_NAME` environment variable to the name of the workspace where the cluster is attached:

```
export WORKSPACE_NAME=<workspace_name>
```

10.5.6.2 Upgrade Platform Applications from the CLI

10.5.6.2.1 Execute the DKP Upgrade Command

Upgrade all Platform applications in a workspace and its projects to the same version as the platform applications running on the management cluster:

```
dkp upgrade workspace ${WORKSPACE_NAME}
```

An output similar to this appears:

```
✓ Ensuring HelmReleases are upgraded on clusters in namespace "${WORKSPACE_NAME}"
...
```

10.5.6.2.2 Troubleshooting

If the upgrade fails or times out, retry the command with higher verbosity to get more information on the upgrade process:

```
dkp upgrade workspace ${WORKSPACE_NAME} -v 4
```

- = If you find any `HelmReleases` in a “broken” release state, such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HelmRelease` using the following commands:

```
kubectl -n ${WORKSPACE_NAMESPACE} patch helmrelease <HELMRELEASE_NAME> --type='json'
-p='[{"op": "replace", "path": "/spec/suspend", "value": true}]'
kubectl -n 4${WORKSPACE_NAMESPACE} patch helmrelease <HELMRELEASE_NAME> --type='json'
-p='[{"op": "replace", "path": "/spec/suspend", "value": false}]'
```

10.5.6.3 Next Step:

[Enterprise: Upgrade Workspace Catalog Applications \(see page 1571\)](#)

10.5.7 Enterprise: Upgrade Workspace Catalog Applications

Upgrade catalog applications using the CLI and UI

Enterprise

Gov Advanced

Upgrade the workspace Catalog Applications in `kommander` on each additional workspace on which catalog apps are deployed. Ensure that you have updated the actual DKP Catalog GitRepository from the [previous page \(see page 1568\)](#) in order for the new apps to be available for upgrade as shown below:

Update the GitRepository with the tag of your updated DKP version on the `kommander` workspace:

```
kubectl patch gitrepository -n kommander dkp-catalog-applications --type merge --
patch '{"spec":{"ref":{"tag":"v2.6.2"}}}'
```

- = This command updates the catalog application repositories for all workspaces.

For any additional Catalog `GitRepository` resources created outside the `kommander.yaml` configuration (e.g. such as in a project or workspace namespace), follow these steps:

Set the `WORKSPACE_NAMESPACE` environment variable to the namespace of the workspace:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```

Update the `GitRepository` for the additional workspace:

```
kubectl patch gitrepository -n ${WORKSPACE_NAMESPACE} dkp-catalog-applications --type merge --patch '{"spec":{"ref":{"tag":"v2.6.2"}}}'
```

10.5.7.1 Upgrade Catalog Applications

Before upgrading also, keep in mind the distinction between Platform applications and Catalog applications. Platform applications are deployed and upgraded as a set for each cluster or workspace. Catalog applications are deployed separately, so that you can deploy and upgrade them individually for each project.

10.5.7.1.1 Upgrade with UI

Follow these steps to upgrade an application from the DKP UI:

1. From the top menu bar, select your target workspace.
2. Select **Applications** from the sidebar menu.
3. Select the three dot button from the bottom-right corner of the desired application tile, and then select **Edit**.
4. Select the **Version** drop-down, and select a new version. This drop-down will only be available if there is a newer version to upgrade to.
5. Select **Save**.

10.5.7.1.2 Upgrade with CLI



Please note that the below commands are using the workspace **name** and not namespace. You can retrieve the workspace name by running the command:

```
dkp get workspaces
```

To view a list of the deployed apps to your workspace, you can run the command:

```
dkp get appdeployments --workspace=<workspace-name>
```

1. To see what app(s) and app versions are available to upgrade, run the following command:

NOTE: You can reference the app version by going into the app name (e.g. `<APP ID>-<APP VERSION>`)

```
kubectl get apps -n ${WORKSPACE_NAMESPACE}
```

You can also use this command to display the apps and app versions, for example:

```
kubectl get apps -n ${WORKSPACE_NAMESPACE} -o jsonpath='{range .items[*]}{@.spec.appId}{"----"}{@.spec.version}{"\n"}{end}'
```

Below is an example of an output that shows the different apps and apps versions.

```
kafka-operator----0.20.0
kafka-operator----0.20.2
kafka-operator----0.23.0-dev.0
spark-operator----1.1.17
spark-operator----1.1.6
zookeeper-operator----0.2.13
zookeeper-operator----0.2.14
```

2. Run the following command to upgrade an application from the DKP CLI:

```
dkp upgrade catalogapp <appdeployment-name> --workspace=${WORKSPACE_NAME} --to-version=<version.number>
```

As an example, the following command upgrades the Kafka Operator application, named `kafka-operator-abc` , in a workspace to version `0.25.1` :

```
dkp upgrade catalogapp kafka-operator-abc --workspace=${WORKSPACE_NAME} --to-version=0.25.1
```



Platform applications cannot be upgraded on a one-off basis, and must be upgraded in a single process for each workspace. If you attempt to upgrade a platform application with these commands, you receive an error and the application is not upgraded.



From DKP 2.6, the following versions of Catalog applications are deprecated:

- kafka-operator-0.20.0
- kafka-operator-0.20.2
- spark-operator-1.1.6
- zookeeper-operator-0.2.13

If you plan on upgrading to DKP 2.6 or above, ensure that you upgrade these applications to the latest compatible version.

- To find what versions of applications are available for upgrade, use `kubectl` :

```
kubectl get apps -n ${WORKSPACE_NAMESPACE}
```

For more information, see [Workspace DKP Catalog Applications \(see page 667\)](#).

10.5.7.2 Next Step:

[Enterprise: Upgrade Project Catalog Applications \(see page 1574\)](#)

10.5.8 Enterprise: Upgrade Project Catalog Applications

Enterprise

Gov Advanced

Before upgrading your catalog applications, verify the current and supported versions of the application. Also, keep in mind the distinction between Platform applications and Catalog applications. Platform applications are deployed and upgraded as a set for each cluster or workspace. Catalog applications are deployed separately, so that you can deploy and upgrade them individually for each project.



Catalog applications must be upgraded to the latest version BEFORE upgrading the Konvoy components for Managed clusters or Kubernetes version for attached clusters.

10.5.8.1 Upgrade with the UI

Follow these steps to upgrade an application from the DKP UI:

1. From the top menu bar, select your target workspace.
2. From the side menu bar, select **Projects**.
3. Select your target project.
4. Select **Applications** from the project menu bar.

5. Select the three dot button from the bottom-right corner of the desired application tile, and then select **Edit**.
6. Select the **Version** drop-down, and select a new version. This drop-down will only be available if there is a newer version to upgrade to.
7. Select **Save**.

10.5.8.2 Upgrade with the CLI

1. To see what app(s) and app versions are available to upgrade, run the following command:

NOTE: You can reference the app version by going into the app name (e.g. `<APP ID>-<APP VERSION>`)

```
kubectl get apps -n ${PROJECT_NAMESPACE}
```

You can also use this command to display the apps and app versions, for example:

```
kubectl get apps -n ${PROJECT_NAMESPACE} -o jsonpath='{range .items[*]}{@.spec.appId}{"----"}{@.spec.version}{"\n"}{end}'
```

Below is an example of an output that shows the different apps and apps versions.


```
zookeeper-operator----0.2.13
zookeeper-operator----0.2.14
zookeeper-operator----0.2.15
```

2. Run the following command to upgrade an application from the DKP CLI:

```
dkp upgrade catalogapp <appdeployment-name> --workspace=my-workspace --project=my-project --to-version=<version.number>
```

As an example, the following command upgrades the Zookeeper Operator application, named `zookeeper-operator-abc` , in a workspace to version `0.2.15` :

```
dkp upgrade catalogapp zookeeper-operator-abc --workspace=my-workspace --to-version=0.2.15
```

 Platform applications cannot be upgraded on a one-off basis, and must be upgraded in a single process for each workspace. If you attempt to upgrade a platform application with these commands, you receive an error and the application is not upgraded.

10.5.8.3 Next Step:

[Enterprise: Upgrade Custom Applications \(see page 1576\)](#)


10.5.9 Enterprise: Upgrade Custom Applications

Enterprise

Gov Advanced

10.5.9.1 Verify the compatibility of Custom Applications with the current Kubernetes version

We recommend upgrading your Custom Applications to the latest compatible version as soon as possible. Since Custom Applications are not created, maintained or supported by D2iQ, you must upgrade them manually.

 Ensure you validate any Custom Applications you run for compatibility issues against the [Kubernetes version \(see page 1752\)](#) in the new release. If the Custom Application's version is not compatible with the Kubernetes version, do not continue with the Konvoy upgrade. Otherwise, your custom Applications may stop running.

After you have ensured your Custom Applications are compatible with the current Kubernetes version, return to the [DKP Upgrade overview \(see page 0\)](#) documentation, to review the next steps depending on your environment and license type.

10.5.9.2 Next Step:

[Enterprise: Upgrade the Management Cluster CAPI Components \(see page 1576\)](#)

10.5.10 Enterprise: Upgrade the Management Cluster CAPI Components

Upgrade the CAPI Components on the Management cluster.

Enterprise

Gov Advanced

- For a Pre-provisioned air-gapped environment only, ensure you have [uploaded the artifacts \(see page 1561\)](#).

10.5.10.1 Upgrade the CAPI Components

New versions of DKP come pre-bundled with newer versions of CAPI, newer versions of infrastructure providers, or new infrastructure providers. When using a new version of the DKP CLI, upgrade all of these components first.

If you are running on more than one management cluster, you must upgrade the CAPI components on each of these clusters.

- Ensure your `dkp` configuration references the management cluster where you want to run the upgrade by setting the `KUBECONFIG` environment variable, or using the `--kubeconfig` flag, [in accordance with Kubernetes conventions](#)¹²⁶⁶.

Execute the upgrade command for the CAPI components.

- If you created CAPI components using flags to specify values, use those same flags during **Upgrade** to preserve existing values while setting additional values.
 - Refer to [dkp create cluster aws \(see page 1665\)](#) for flag descriptions for `--with-aws-bootstrap-credentials` and `--aws-service-endpoints`
 - Refer to the HTTP section for details: [Create CAPI Components with HTTP/HTTPS Proxy Settings \(see page 1044\)](#)

```
dkp upgrade capi-components
```

The output resembles the following:

```
✓ Upgrading CAPI components
✓ Waiting for CAPI components to be upgraded
```

¹²⁶⁶ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

✓ Initializing **new** CAPI components

If the upgrade fails, review the prerequisites section and ensure that you've followed the steps in the [DKP Upgrade](#) (see page 1554) overview. Furthermore, ensure you have adhered to the Prerequisites at the top of this page.

10.5.10.2 Next Step:

[Enterprise: Upgrade the Management Cluster Core Addons](#) (see page 1578)

10.5.11 Enterprise: Upgrade the Management Cluster Core Addons

Upgrade the Core Addons on the **Management Cluster**.

Enterprise

Gov Advanced

For Pre-provisioned air-gapped environments only, you must run `konvoy-image upload artifacts` to [copy the artifacts](#) (see page 1207) onto the cluster hosts **before** you begin the Upgrade the CAPI Components section below.

```
konvoy-image upload artifacts \
  --container-images-dir=./artifacts/images/ \
  --os-packages-bundle=./artifacts/$OS_PACKAGES_BUNDLE \
  --containerd-bundle=artifacts/$CONTAINERD_BUNDLE \
  --pip-packages-bundle=./artifacts/pip-packages.tar.gz
```

To install the core addons, DKP relies on the `ClusterResourceSet` [Cluster API feature](#)¹²⁶⁷. In the CAPI component upgrade, we deleted the previous set of outdated global `ClusterResourceSets` because in past releases, some addons were installed using a global configuration. In order to support individual cluster upgrades, DKP now installs all addons with a unique set of `ClusterResourceSets` and corresponding referenced resources, all named using the cluster's name as a suffix. For example: `calico-cni-installation-my-aws-cluster`.



If you have modified any of the `ClusterResourceSet` definitions, these changes will **not** be preserved when running the command `dkp upgrade addons <provider>`. You must define the cloud provider before you use the `--dry-run -o yaml` options to save the new configuration to a file and remake the same changes upon each upgrade.

¹²⁶⁷ <https://cluster-api.sigs.k8s.io/>

Your cluster comes preconfigured with a few different core addons that provide functionality to your cluster upon creation. These include: CSI, CNI, Cluster Autoscaler, and Node Feature Discovery. New versions of DKP may come pre-bundled with newer versions of these addons.

Perform the following steps to update these addons:

1. If you have any additional managed clusters, you will need to upgrade the core addons and Kubernetes version for each one.
2. Ensure your `dkp` configuration references the management cluster where you want to run the upgrade by setting the `KUBECONFIG` environment variable, or using the `--kubeconfig` flag, [in accordance with Kubernetes conventions](#)¹²⁶⁸.
3. Upgrade the core addons in a cluster using the `dkp upgrade addons` command specifying the cluster infrastructure (choose `aws`, `azure`, `vsphere`, `eks`, `gcp`, `preprovisioned`) and the name of the cluster.



If you need to verify or discover your cluster name to use with this example, first run the `kubectl get clusters` command.

Examples for upgrade core addons commands:

```
export CLUSTER_NAME=my-azure-cluster
dkp upgrade addons azure --cluster-name=${CLUSTER_NAME}
```

OR

```
export CLUSTER_NAME=my-aws-cluster
dkp upgrade addons aws --cluster-name=${CLUSTER_NAME}
```

The output for the AWS example should be similar to:

```
Generating addon resources
clusterresourceset.addons.cluster.x-k8s.io/calico-cni-installation-my-aws-cluster
upgraded
configmap/calico-cni-installation-my-aws-cluster upgraded
clusterresourceset.addons.cluster.x-k8s.io/tigera-operator-my-aws-cluster upgraded
configmap/tigera-operator-my-aws-cluster upgraded
clusterresourceset.addons.cluster.x-k8s.io/aws-ebs-csi-my-aws-cluster upgraded
configmap/aws-ebs-csi-my-aws-cluster upgraded
clusterresourceset.addons.cluster.x-k8s.io/cluster-autoscaler-my-aws-cluster upgraded
configmap/cluster-autoscaler-my-aws-cluster upgraded
```

¹²⁶⁸ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

```

clusterresourceset.addons.cluster.x-k8s.io/node-feature-discovery-my-aws-cluster
upgraded
configmap/node-feature-discovery-my-aws-cluster upgraded
clusterresourceset.addons.cluster.x-k8s.io/nvidia-feature-discovery-my-aws-cluster
upgraded
configmap/nvidia-feature-discovery-my-aws-cluster upgraded

```

10.5.11.1 See Also:

[DKP upgrade addons](#) (see page 1742) for more CLI command help.

10.5.11.2 Next Step:

[Enterprise: Upgrade the Management Cluster Kubernetes Version](#) (see page 1580)

10.5.12 Enterprise: Upgrade the Management Cluster Kubernetes Version

Upgrade the Kubernetes version on the **Management cluster**.

Enterprise

Gov Advanced

10.5.12.1 Upgrade the Kubernetes Version

When upgrading the Kubernetes version of a cluster:

1. Upgrade the control plane first using the infrastructure specific command.
 - a. NOTE the additional considerations for FIPS if using FIPS configuration.
2. Upgrade the node pools second using the infrastructure specific command.
 - a. NOTE the additional considerations for FIPS if using FIPS configuration.
3. If you have any additional managed, you need to upgrade the core addons and Kubernetes version for each one. Attached clusters would need the Kubernetes version upgraded to a supported DKP version using the tool that created that cluster.
4. During an upgrade, you would need to create new AMI's or images with the K8S version to which you want to upgrade which requires selecting your OS and ensuring you have the current [supported version](#) (see page 69). Build a new image if applicable.
 - If an AMI was specified when initially creating a cluster for AWS, you must build a new one with [Konvoy Image Builder](#) (see page 1472) and set the flag(s) in the update commands. Either AMI ID `--ami AMI_ID`, or the lookup image flags: `--ami-owner AWS_ACCOUNT_ID`, `--ami-base-os ubuntu-20.04`, and `--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*'`.

⚠ The AMI lookup method will return an error if the lookup uses the upstream CAPA account ID.

- If an Azure Machine Image was specified for Azure, you must build a new one with [Konvoy Image Builder](#) (see page 1480).
 - If a vSphere template Image was specified for vSphere, you must build a new one with [Konvoy Image Builder](#) (see page 1490).
 - You must build a new GCP image with [Konvoy Image Builder](#) (see page 1371).
5. Upgrade the Kubernetes version of the control plane. Each cloud provider has distinctive commands. Below is the AWS command example. Select the drop-down menu next to your provider for compliant CLI.

NOTE: The first example below is for AWS. If you created your initial cluster with a custom AMI using the `--ami` flag, it is required to set the `--ami` flag during the Kubernetes upgrade.

```
dkp update controlplane aws \
--cluster-name=${CLUSTER_NAME} \
--ami AMI_ID \
--kubernetes-version=v1.26.14
```

OR for AWS AMI lookup ID

```
dkp update controlplane aws \
--cluster-name=${CLUSTER_NAME} \
--ami-owner AWS_ACCOUNT_ID \
--ami-base-os ubuntu-20.04 \
--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*' \
--kubernetes-version=v1.26.14
```

Azure

```
dkp update controlplane azure --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --compute-gallery-id <Azure Compute Gallery built by KIB for
Kubernetes v1.26.14>
```

- If these fields were specified in the [override file](#) (see page 1516) during [image creation](#) (see page 1480), the flags must be used in upgrade:

- `--plan-offer`, `--plan-publisher` and `--plan-sku`

- `--plan-offer rockylinux-9`
`--plan-publisher erockyenterprisesoftwarefoundationinc1653071250513`
`--plan-sku rockylinux-9`

vSphere


```
dkp update controlplane vsphere --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --vm-template <vSphere template built by KIB for Kubernetes v1.26.14>
```

GCP

```
dkp update controlplane gcp --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --image=projects/${GCP_PROJECT}/global/images/<GCP image built by
KIB for Kubernetes v1.26.14>
```

Pre-provisioned

```
dkp update controlplane preprovisioned --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14
```

EKS

```
dkp update controlplane eks --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.25.11
```

Additional Considerations for upgrading a FIPS cluster:

If upgrading a FIPS cluster, to correctly upgrade the Kubernetes version, instead run the command shown below which includes the etcd FIPS information as well:

```
dkp update controlplane aws --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14+fips.0 --ami=<ami-with-fips-id>
```

The output should be similar to the below example, with the provider name corresponding to the CLI you executed from the choices above:

```
Updating control plane resource controlplane.cluster.x-k8s.io/v1beta1,
Kind=KubeadmControlPlane default/my-aws-cluster-control-plane
Waiting for control plane update to finish.
✓ Updating the control plane
```

- Some advanced options are available for various providers. To see all the options for your particular provider, run this command `dkp update controlplane aws|vsphere|preprovisioned|azure|gcp|eks --help` for more advance options like the example below:

This example for AWS AMI instance type: `aws: --ami, --instance-type` would be some of the options mentioned in the note above.

NOTE: The command `dkp update controlplane {provider}` has a 30 minute default timeout for the update process to finish. If you see the error " `timed out waiting for the condition` ", you can check the control plane nodes version using the command `kubectl get machines -o wide $KUBECONFIG` before trying again.

6. Upgrade the Kubernetes version of your node pools. Upgrading a nodepool involves draining the existing nodes in the nodepool and replacing them with new nodes. In order to ensure minimum downtime and maintain high availability of the critical application workloads during the upgrade process, we recommend deploying Pod Disruption Budget ([Disruptions](#)¹²⁶⁹) for your critical applications. For more information, refer to [Update Cluster Nodepools](#) (see [page 1548](#)) documentation.

a. First, get a list of all node pools available in your cluster by running the following command:

```
dkp get nodepool --cluster-name ${CLUSTER_NAME}
```

b. Select the nodepool you want to upgrade with the command below:

```
export NODEPOOL_NAME=my-nodepool
```

c. Then update the selected nodepool using the command below. The first example command shows AWS language, so select the drop-down menu for your provider for the correct command. Execute the `update` command for each of the node pools listed in the previous command:

NOTE: The first example below is for AWS. If you created your initial cluster with a custom AMI using the `--ami` flag, it is required to set the `--ami` flag during the Kubernetes upgrade.

```
dkp update nodepool aws ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-version=v1.26.14
```

Azure

```
dkp update nodepool azure ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-version=v1.26.14 --compute-gallery-id <Azure Compute Gallery built by KIB for Kubernetes v1.26.14>
```

- If these fields were specified in the [override file](#) (see [page 1516](#)) during [image creation](#) (see [page 1480](#)), the flags must be used in upgrade:

- `--plan-offer` , `--plan-publisher` and `--plan-sku`

- `--plan-offer rockylinux-9`

¹²⁶⁹ <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

```
--plan-publisher erockyenterprisesoftwarefoundationinc1653071250513
--plan-sku rockylinux-9
```

vSphere

```
dkp update nodepool vsphere ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --
kubernetes-version=v1.26.14 --vm-template <vSphere template built by KIB for
Kubernetes v1.26.14>
```

GCP

```
dkp update nodepool gcp ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --image=projects/${GCP_PROJECT}/global/images/<GCP image built by
KIB for Kubernetes v1.26.14>
```

Pre-provisioned

```
dkp update nodepool preprovisioned ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --
kubernetes-version=v1.26.14
```

EKS

```
dkp update nodepool eks ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.25.11
```

The output should be similar to the following, with the name of the infrastructure provider shown accordingly:

```
Updating node pool resource cluster.x-k8s.io/v1beta1, Kind=MachineDeployment default/
my-aws-cluster-my-nodepool
Waiting for node pool update to finish.
✓ Updating the my-aws-cluster-my-nodepool node pool
```

d. Repeat this step for each additional node pool.

Additional Considerations for upgrading a FIPS cluster:

If upgrading a FIPS cluster, to correctly upgrade the Kubernetes version, instead run the command shown below:

```
dkp update nodepool aws ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14+fips.0 --ami=<ami-with-fips-id>
```

When all nodepools have been updated, your upgrade is complete. For the overall process for upgrading to the latest version of DKP, refer back to [Upgrade DKP \(see page 1554\)](#) for more details.

10.5.12.2 Next Step:

[Enterprise: Upgrade Managed Clusters](#) (see page 1585)

10.5.13 Enterprise: Upgrade Managed Clusters

10.5.13.1 Upgrade Managed Clusters

Enterprise

Gov Advanced

If you have managed clusters, follow these steps to upgrade each cluster similar to the management cluster process order:

1. Using the kubeconfig of your management cluster, find your cluster name and be sure to copy the information for all of your clusters:

```
kubectl get clusters -A
```

2. Set your cluster variable:

```
export CLUSTER_NAME=<your-managed-cluster-name>
```

3. Set your cluster's workspace variable:


```
export CLUSTER_WORKSPACE=<your-workspace-namespace>
```

4. Then, upgrade the core addons (replacing `aws` with whatever infrastructure provider you would be using):

```
dkp upgrade addons aws --cluster-name=${CLUSTER_NAME} -n ${CLUSTER_WORKSPACE}
```

10.5.13.2 Upgrade Kubernetes Version on a Managed Cluster

After you complete the previous steps for all managed clusters and you update your core addons, begin upgrading the Kubernetes version.

-  You should first complete the upgrade of your Kommander Management Cluster before upgrading any managed clusters.

1. Use this command to begin upgrading the Kubernetes version:

```
dkp update controlplane aws --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 -n ${CLUSTER_WORKSPACE}
```

For EKS

```
dkp update controlplane eks --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.25.11 -n ${CLUSTER_WORKSPACE}
```

For Azure

```
dkp update controlplane azure --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --compute-gallery-id <Azure Compute Gallery built by KIB for
Kubernetes v1.26.14> -n ${CLUSTER_WORKSPACE}
```

For vSphere

```
dkp update controlplane vsphere --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --vm-template <vSphere template built by KIB for Kubernetes v1.26.14>
-n ${CLUSTER_WORKSPACE}
```

For GCP

```
dkp update controlplane gcp --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --image=projects/${GCP_PROJECT}/global/images/<GCP image built by
KIB for Kubernetes v1.26.14> -n ${CLUSTER_WORKSPACE}
```

For Pre-provisioned

```
dkp update controlplane preprovisioned --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 -n ${CLUSTER_WORKSPACE}
```

2. Get a list of all node pools available in your cluster by running the following command:

```
dkp get nodepools -c ${CLUSTER_NAME} -n ${CLUSTER_WORKSPACE}
```

```
export NODEPOOL_NAME=<my-nodepool>
```

3. Use this command to upgrade the node pools:

```
dkp update nodepool aws ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 -n ${CLUSTER_WORKSPACE}
```

For EKS

```
dkp update nodepool eks ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.25.11 -n ${CLUSTER_WORKSPACE}
```

For Azure

```
dkp update nodepool azure ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --
kubernetes-version=v1.26.14 --compute-gallery-id <Azure Compute Gallery built by KIB
for Kubernetes v1.26.14> -n ${CLUSTER_WORKSPACE}
```

For vSphere

```
dkp update nodepool vsphere ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --
kubernetes-version=v1.26.14 --vm-template <vSphere template built by KIB for
Kubernetes v1.26.14> -n ${CLUSTER_WORKSPACE}
```

For GCP

```
dkp update nodepool gcp ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --image=projects/${GCP_PROJECT}/global/images/<GCP image built by
KIB for Kubernetes v1.26.14> -n ${CLUSTER_WORKSPACE}
```

For Pre-provisioned

```
dkp update nodepool preprovisioned ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --
kubernetes-version=v1.26.14 -n ${CLUSTER_WORKSPACE}
```

10.5.13.3 Upgrade Attached Cluster

Since any attached clusters you have are managed by their corresponding cloud provider, none of the components would be upgraded using the DKP process. The tool used to create the cluster is used to upgrade it. Ensure that it has the Kubernetes version that is compatible with the versions we support.

10.6 Upgrade DKP Essential

10.6.1 Steps to upgrade DKP via CLI

The DKP upgrade represents an important step of your environment's lifecycle, as it ensures that you are up-to-date with the latest features and can benefit from the most recent improvements, enhanced cluster management, and better performance. If you have are on this scenario of the upgrade section of documentation, you have an Essential License. If you haven an [Enterprise license \(see page 125\)](#), please jump to that section of the [Upgrade DKP Enterprise \(see page 1565\)](#) portion of documentation.

Ensure that all platform applications in the management cluster have been upgraded to avoid compatibility issues with the [Kubernetes version \(see page 1555\)](#) included in this release. This is done automatically when upgrading Kommander, so ensure that you follow these sections in order to upgrade the Kommander component prior to upgrading the Konvoy component.

If you have more than one Essential cluster, repeat all of these steps for each Essential cluster. For a full list of DKP Essential features, see [DKP Essential \(see page 132\)](#).



Before you begin the upgrade, ensure you review the [Upgrade Prerequisites \(see page 1558\)](#).

10.6.2 Overview of Steps in Order

This section describes how to upgrade your DKP clusters to the latest DKP version.

- [Essential: For Air-gapped Environments Only \(see page 1589\)](#)
- [Essential: Upgrade the Essential Cluster and Platform Applications \(see page 1590\)](#)
- [Essential: Upgrade the Cluster CAPI Components \(see page 1592\)](#)
- [Essential: Upgrade the Cluster Core Addons \(see page 1594\)](#)
- [Essential Upgrade Kubernetes Version \(see page 1596\)](#)

10.6.3 Next Step:

If you are in an air-gapped environment, be sure to select that choice to find out how to load your local registry first. Otherwise, your next step will be [Essential: Upgrade the Essential Cluster and Platform Applications \(see page 1590\)](#).

- [Essential: For Air-gapped Environments Only \(see page 1589\)](#)

10.6.4 Essential: For Air-gapped Environments Only

If you are operating in an air-gapped environment, a [local container registry](#) (see page 1536) containing all the necessary installation images, including the Kommander images is required. See below for how to push the necessary images to this registry.

- [Download](#) (see page 78) the Complete DKP Air-gapped Bundle for this release (i.e. `dkp-air-gapped-bundle_v2.6.1_linux_amd64.tar.gz`)
- Connectivity:
 - Your Essential cluster must be able to connect to the local registry.

10.6.4.1 Extract Air-gapped Images and Set Variables

Follow these steps to **extract** the air-gapped image bundles into your private registry:

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzf dkp-air-gapped-bundle_v2.6.2_linux_amd64.tar.gz && cd dkp-v2.6.2
```

2. Set an environment variable with your registry address:

```
export REGISTRY_ADDRESS=<registry-address>:<registry-port>
export REGISTRY_USERNAME=<username>
export REGISTRY_PASSWORD=<password>
```

10.6.4.2 Load Images to your Private Registry - Kommander

Load Kommander images to your Private Registry

For the **air-gapped** `kommander` image bundle, run the command below:

Run the following command to load the image bundle:

```
dkp push bundle --bundle ./container-images/kommander-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS
```

10.6.4.3 Load Images to your Private Registry - Konvoy

Before creating or upgrading a Kubernetes cluster, you need to load the required images in a local registry if operating in an air-gapped environment. This registry must be accessible from both the [bastion machine](#) (see

[page 1288](#)) and either the AWS EC2 instances or other machines that will be created for the Kubernetes cluster.



If you do not already have a local registry set up, refer to [Local Registry Tools \(see page 1536\)](#) page for more information.

Execute the following command to load the air-gapped image bundle into your private registry:

```
dkp push bundle --bundle ./container-images/konvoy-image-bundle-v2.6.2.tar --to-registry $REGISTRY_ADDRESS --to-registry-username $REGISTRY_USERNAME --to-registry-password $REGISTRY_PASSWORD
```



It may take some time to push all the images to your image registry, depending on the performance of the network between the machine you are running the script on and the registry.

10.6.4.4 Next Step:

[Essential: Upgrade the Essential Cluster and Platform Applications \(see page 1590\)](#)

10.6.5 Essential: Upgrade the Essential Cluster and Platform Applications

This section describes how to upgrade your Kommander Management cluster and all Platform Applications to their supported versions in air-gapped, non-air-gapped and on-prem environments. To prevent compatibility issues, you must first upgrade the Kommander component on your [Essential Cluster \(see page 83\)](#) before upgrading to DKP.



It is important you upgrade Kommander BEFORE upgrading the Kubernetes version. This ensures that any changes required for new or changed Kubernetes API's are already present.

10.6.5.1 Upgrade Kommander

10.6.5.1.1 Prerequisite

- Use the `--kubeconfig=${CLUSTER_NAME}.conf` flag or set the `KUBECONFIG` environment variable to ensure that you **upgrade Kommander on the right cluster**. For alternatives and recommendations around setting your context, refer to [Provide Context for Commands with a kubeconfig File](#) (see page 87).
- If you have DKP Insights installed, ensure you **uninstall** it before upgrading DKP. See [\(2.6\) DKP Insights Upgrade from DKP 2.5.x to 2.6.0](#)¹²⁷⁰ for more information.



Limited availability of DKP resources

- If you have configured a **custom domain**, running the `upgrade` command could result in an inaccessibility of your services via your custom domain for a few minutes.
- The **DKP UI** and other APIs may be inconsistent or unavailable until the upgrade is complete.

10.6.5.1.2 Upgrade

Use the DKP CLI to upgrade Kommander and all the Platform Applications in the Management Cluster:

For air-gapped environments:

```
dkp upgrade kommander \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.6.1.tar.gz \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.6.1.tar.gz
```


The output looks similar to this:

```
✓ Ensuring upgrading conditions are met
✓ Ensuring application definitions are updated
✓ Ensuring helm-mirror implementation is migrated to ChartMuseum
...
```

For non-air-gapped environments:

¹²⁷⁰ <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/274989271/%282.6%29+DKP+Insights+Upgrade+from+DKP+2.5.x+to+2.6.0>

```
dkp upgrade kommander
```

 If you want to disable the AI Navigator, add the flag `--disable-appdeployments ai-navigator-app` to this command.

The output looks similar to this:

```
✓ Ensuring upgrading conditions are met
✓ Ensuring application definitions are updated
✓ Ensuring helm-mirror implementation is migrated to ChartMuseum
...
```

10.6.5.1.3 Troubleshooting

1. If the upgrade fails, run the following command to get more information on the upgrade process:

```
dkp upgrade kommander -v 6
```

2. If you find any `HeLmReLeases` in a “broken” release state such as “exhausted” or “another rollback/release in progress”, you can trigger a reconciliation of the `HeLmReLease` using the following commands:


```
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": true}]'
kubectl -n kommander patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": false}]'
```

10.6.5.1.4 Next Step:

[Essential: Upgrade the Cluster CAPI Components](#) (see page 1592)

10.6.6 Essential: Upgrade the Cluster CAPI Components


Upgrade the CAPI Components on the Management cluster.

 For a Pre-provisioned air-gapped environment only, ensure you have [uploaded the artifacts](#) (see [page 1561](#)).


10.6.6.1 Upgrade the CAPI Components

New versions of DKP come pre-bundled with newer versions of CAPI, newer versions of infrastructure providers, or new infrastructure providers. When using a new version of the DKP CLI, upgrade all of these components first.

If you are running on more than one management cluster, you must upgrade the CAPI components on each of these clusters.

 Ensure your `dkp` configuration references the management cluster where you want to run the upgrade by setting the `KUBECONFIG` environment variable, or using the `--kubeconfig` flag, [in accordance with Kubernetes conventions](#)¹²⁷¹.

Execute the upgrade command for the CAPI components.

 If you created CAPI components using flags to specify values, use those same flags during **Upgrade** to preserve existing values while setting additional values.

- Refer to [dkp create cluster aws](#) (see [page 1665](#)) for flag descriptions for `--with-aws-bootstrap-credentials` and `--aws-service-endpoints`
- Refer to the HTTP section for details: [Create CAPI Components with HTTP/HTTPS Proxy Settings](#) (see [page 1044](#))

```
dkp upgrade capi-components
```

The output resembles the following:

```
✓ Upgrading CAPI components
✓ Waiting for CAPI components to be upgraded
✓ Initializing new CAPI components
```

¹²⁷¹ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

If the upgrade fails, review the prerequisites section and ensure that you've followed the steps in the [DKP Upgrade \(see page 1554\)](#) overview. Furthermore, ensure you have adhered to the Prerequisites at the top of this page.

10.6.6.2 Next Step:

[Essential: Upgrade the Cluster Core Addons \(see page 1594\)](#)

10.6.7 Essential: Upgrade the Cluster Core Addons

For Pre-provisioned air-gapped environments only, you must run `konvoy-image upload artifacts` to [copy the artifacts \(see page 1207\)](#) onto the cluster hosts **before** you begin the Upgrade the CAPI Components section below.

```
konvoy-image upload artifacts \
  --container-images-dir=./artifacts/images/ \
  --os-packages-bundle=./artifacts/$OS_PACKAGES_BUNDLE \
  --containerd-bundle=artifacts/$CONTAINERD_BUNDLE \
  --pip-packages-bundle=./artifacts/pip-packages.tar.gz
```


10.6.7.1 Upgrade the Core Addons

To install the core addons, DKP relies on the `ClusterResourceSet` [Cluster API feature](#)¹²⁷². In the CAPI component upgrade, we deleted the previous set of outdated global `ClusterResourceSets` because in past releases, some addons were installed using a global configuration. In order to support individual cluster upgrades, DKP now installs all addons with a unique set of `ClusterResourceSets` and corresponding referenced resources, all named using the cluster's name as a suffix. For example: `calico-cni-installation-my-aws-cluster`.

If you modify any of the `ClusterResourceSet` definitions, these changes are **not** be preserved when running the command `dkp upgrade addons`. You must use the `--dry-run -o yaml` options to save the new configuration to a file and continue the same changes upon each upgrade.

Your cluster comes preconfigured with a few different core addons that provide functionality to your cluster upon creation. These include: CSI, CNI, Cluster Autoscaler, and Node Feature Discovery. New versions of DKP may come pre-bundled with newer versions of these addons.

¹²⁷² <https://cluster-api.sigs.k8s.io/>

 If you have more than one essential cluster, ensure your `dkp` configuration references the cluster where you want to run the upgrade by setting the `KUBECONFIG` environment variable, or using the `--kubeconfig` flag, [in accordance with Kubernetes conventions](#)¹²⁷³.

Perform the following steps to update your addons.

1. Ensure your `dkp` configuration references the cluster where you want to run the upgrade by setting the `KUBECONFIG` environment variable, or using the `--kubeconfig` flag, [in accordance with Kubernetes conventions](#)¹²⁷⁴.
2. Upgrade the core addons in a cluster using the `dkp upgrade addons` command specifying the cluster infrastructure (choose `aws`, `azure`, `vsphere`, `eks`, `gcp`, `preprovisioned`) and the name of the cluster.

Examples for upgrade core addons commands:

```
export CLUSTER_NAME=my-azure-cluster
dkp upgrade addons azure --cluster-name=${CLUSTER_NAME}
```

OR

```
export CLUSTER_NAME=my-aws-cluster
dkp upgrade addons aws --cluster-name=${CLUSTER_NAME}
```

The output for the AWS example should be similar to:

```
Generating addon resources
clusterresourceset.addons.cluster.x-k8s.io/calico-cni-installation-my-aws-cluster
upgraded
configmap/calico-cni-installation-my-aws-cluster upgraded
clusterresourceset.addons.cluster.x-k8s.io/tigera-operator-my-aws-cluster upgraded
configmap/tigera-operator-my-aws-cluster upgraded
clusterresourceset.addons.cluster.x-k8s.io/aws-ebs-csi-my-aws-cluster upgraded
configmap/aws-ebs-csi-my-aws-cluster upgraded
clusterresourceset.addons.cluster.x-k8s.io/cluster-autoscaler-my-aws-cluster upgraded
configmap/cluster-autoscaler-my-aws-cluster upgraded
clusterresourceset.addons.cluster.x-k8s.io/node-feature-discovery-my-aws-cluster
upgraded
configmap/node-feature-discovery-my-aws-cluster upgraded
clusterresourceset.addons.cluster.x-k8s.io/nvidia-feature-discovery-my-aws-cluster
upgraded
```

¹²⁷³ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

¹²⁷⁴ <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

```
configmap/nvidia-feature-discovery-my-aws-cluster upgraded
```

10.6.7.1.1 See Also

[DKP upgrade addons](#) (see page 1742) for more CLI command help.

10.6.7.2 Next Step:

[Essential Upgrade Kubernetes Version](#) (see page 1596)

10.6.8 Essential Upgrade Kubernetes Version

10.6.8.1 Upgrade the Kubernetes Version

When upgrading the Kubernetes version of a cluster:

1. Upgrade the control plane first using the infrastructure specific command.
 - a. NOTE the additional considerations for FIPS if using FIPS configuration.
2. Upgrade the node pools second using the infrastructure specific command.
 - a. NOTE the additional considerations for FIPS if using FIPS configuration.
3. Build a new image if applicable.
 - If an AMI was specified when initially creating a cluster for AWS, you must build a new one with [Konvoy Image Builder](#) (see page 1472) and set the flag(s) in the update commands. Either AMI ID `--ami AMI_ID`, or the lookup image flags: `--ami-owner AWS_ACCOUNT_ID`, `--ami-base-os ubuntu-20.04`, and `--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*'`.

⚠ The AMI lookup method will return an error if the lookup uses the upstream CAPA account ID.
 - If an Azure Machine Image was specified for Azure, you must build a new one with [Konvoy Image Builder](#) (see page 1480).
 - If a vSphere template Image was specified for vSphere, you must build a new one with [Konvoy Image Builder](#) (see page 1490).
 - You must build a new GCP image with [Konvoy Image Builder](#) (see page 1371).
4. Upgrade the Kubernetes version of the control plane. Each cloud provider has distinctive commands. Below is the AWS command example. Select the drop-down menu next to your provider for compliant CLI.

NOTE: The first example below is for AWS. If you created your initial cluster with a custom AMI using the `--ami` flag, it is required to set the `--ami` flag during the Kubernetes upgrade.

```
dkp update controlplane aws \
--cluster-name=${CLUSTER_NAME} \
--ami AMI_ID \
--kubernetes-version=v1.26.14
```

OR for AWS AMI lookup ID

```
dkp update controlplane aws \
--cluster-name=${CLUSTER_NAME} \
--ami-owner AWS_ACCOUNT_ID \
--ami-base-os ubuntu-20.04 \
--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*' \
--kubernetes-version=v1.26.14
```

Azure

```
dkp update controlplane azure --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --compute-gallery-id <Azure Compute Gallery built by KIB for
Kubernetes v1.26.14>
```

- If these fields were specified in the [override file](#) (see page 1516) during [image creation](#) (see page 1480), the flags must be used in upgrade:
 - `--plan-offer`, `--plan-publisher` and `--plan-sku`

- `--plan-offer rockylinux-9`
`--plan-publisher erockyenterprisesoftwarefoundationinc1653071250513`
`--plan-sku rockylinux-9`

vSphere

```
dkp update controlplane vsphere --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --vm-template <vSphere template built by KIB for Kubernetes v1.26.14>
```

GCP

```
dkp update controlplane gcp --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --image=projects/${GCP_PROJECT}/global/images/<GCP image built by
KIB for Kubernetes v1.26.14>
```

Pre-provisioned

```
dkp update controlplane preprovisioned --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14
```


EKS

```
dkp update controlplane eks --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.25.11
```

Additional Considerations for upgrading a FIPS cluster:

If upgrading a FIPS cluster, to correctly upgrade the Kubernetes version, instead run the command shown below which contains the etcd version as well:

```
dkp update controlplane aws --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14+fips.0 --ami=<ami-with-fips-id>
```

The output should be similar to the below example, with the provider name corresponding to the CLI you executed from the choices above:

```
Updating control plane resource controlplane.cluster.x-k8s.io/v1beta1,
Kind=KubeadmControlPlane default/my-aws-cluster-control-plane
Waiting for control plane update to finish.
✓ Updating the control plane
```

Some advanced options are available for various providers. To see all the options for your particular provider, run this command `dkp update controlplane aws|vsphere|preprovisioned|azure|gcp|eks --help` for more advance options like the example below:

This example for AWS AMI instance type: `aws: --ami, --instance-type` would be some of the options mentioned in the note above.

NOTE: The command `dkp update controlplane {provider}` has a 30 minute default timeout for the update process to finish. If you see the error "timed out waiting for the condition", you can check the control plane nodes version using the command `kubectl get machines -o wide --kubeconfig $KUBECONFIG` before trying again.

5. Upgrade the Kubernetes version of your node pools. Upgrading a nodepool involves draining the existing nodes in the nodepool and replacing them with new nodes. In order to ensure minimum downtime and maintain high availability of the critical application workloads during the upgrade process, we recommend deploying Pod Disruption Budget ([Disruptions](https://kubernetes.io/docs/concepts/workloads/pods/disruptions/)¹²⁷⁵) for your critical applications. For more information, refer to [Update Cluster Nodepools](https://d2iq.atlassian.net/wiki/spaces/DENT/pages/213191226)¹²⁷⁶ documentation.

a. First, get a list of all node pools available in your cluster by running the following command:

¹²⁷⁵ <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

¹²⁷⁶ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/213191226>

```
dkp get nodepool --cluster-name ${CLUSTER_NAME}
```

b. Select the nodepool you want to upgrade with the command below:

```
export NODEPOOL_NAME=my-nodepool
```

c. Then update the selected nodepool using the command below. Upgrading a node pool involves draining the existing nodes in the node pool and replacing them with new nodes. we recommend deploying Pod Disruption Budget ([Disruptions](#)¹²⁷⁷) for your critical applications. Refer to [Update Cluster Nodepools](#) (see page 1548) for more information.

The first example command shows AWS language, so select the drop-down menu for your provider for the correct command. Execute the `update` command for each of the node pools listed in the previous command:

NOTE: The first example below is for AWS. If you created your initial cluster with a custom AMI using the `--ami` flag, it is required to set the `--ami` flag during the Kubernetes upgrade.

```
dkp update nodepool aws ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-version=v1.26.14
```

Azure

```
dkp update nodepool azure ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-version=v1.26.14 --compute-gallery-id <Azure Compute Gallery built by KIB for Kubernetes v1.26.14>
```

- If these fields were specified in the [override file](#) (see page 1516) during [image creation](#) (see page 1480), the flags must be used in upgrade:
 - `--plan-offer`, `--plan-publisher` and `--plan-sku`

```
--plan-offer rockylinux-9
--plan-publisher erockyenterprisesoftwarefoundationinc1653071250513
--plan-sku rockylinux-9
```

vSphere

```
dkp update nodepool vsphere ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-version=v1.26.14 --vm-template <vSphere template built by KIB for Kubernetes v1.26.14>
```

GCP

¹²⁷⁷ <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

```
dkp update nodepool gcp ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14 --image=projects/${GCP_PROJECT}/global/images/<GCP image built by
KIB for Kubernetes v1.26.14>
```

Pre-provisioned

```
dkp update nodepool preprovisioned ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --
kubernetes-version=v1.26.14
```

EKS

```
dkp update nodepool eks ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.25.11
```

Additional Considerations for upgrading a FIPS cluster:

If upgrading a FIPS cluster, to correctly upgrade the Kubernetes version, instead run the command shown below:

```
dkp update nodepool aws ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.26.14+fips.0 --ami=<ami-with-fips-id>
```

The output should be similar to the following, with the name of the infrastructure provider shown accordingly:

```
Updating node pool resource cluster.x-k8s.io/v1beta1, Kind=MachineDeployment default/
my-aws-cluster-my-nodepool
Waiting for node pool update to finish.
✓ Updating the my-aws-cluster-my-nodepool node pool
```

d. Repeat this step for each additional node pool.

When all nodepools have been updated, your upgrade is complete. For the overall process for upgrading to the latest version of DKP, refer back to [Upgrade DKP \(see page 1554\)](#) for more details.

11 CLI and API Tools

CLI commands and API tools.

- [API Documentation](#) (see page 1601)
- [CLI Commands](#) (see page 1644)

11.1 API Documentation

This document is automatically generated from the API definition in the code.

- [apps.kommander.mesosphere.io/v1alpha2](#) (see page 1601)
- [apps.kommander.mesosphere.io/v1alpha3](#) (see page 1605)
- [dispatch.d2iq.io/v1alpha2](#) (see page 1612)
- [kommander.mesosphere.io/v1beta1](#) (see page 1614)
- [workspaces.kommander.mesosphere.io/v1alpha1](#) (see page 1627)

11.1.1 apps.kommander.mesosphere.io/v1alpha2

11.1.1.1 App

App is the Schema for a Service or Application in Kommander.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹²⁷⁸	false
spec		GenericAppSpec (see page 1604)	false
status		object	false

11.1.1.2 AppDeployment

AppDeployment is the Schema for a concrete installation of a Service or Application in Kommander.

¹²⁷⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹²⁷⁹	false
spec		AppDeploymentSpec (see page 1602)	false
status		object	false

11.1.1.3 AppDeploymentList

AppDeploymentList contains a list of AppDeployments.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹²⁸⁰	false
items		[...]AppDeployment (see page 1604)	true

11.1.1.4 AppDeploymentSpec

AppDeploymentSpec defines an instance of an Application.

Field	Description	Scheme	Required
appRef	AppRef provides reference to a ClusterApp or App object.	TypedLocalObjectReference (see page 1605)	true
configOverrides	ConfigOverrides allows a user to define a ConfigMap that contains configuration overrides	corev1.LocalObjectReference ¹²⁸¹	false

¹²⁷⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹²⁸⁰ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹²⁸¹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

11.1.1.5 AppList

AppList contains a list of Apps.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹²⁸²	false
items		[...]App (see page 1604)	true

11.1.1.6 ClusterApp

ClusterApp is the Schema for a Service or Application in Kommander.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹²⁸³	false
spec		GenericAppSpec (see page 1604)	false
status		object	false

11.1.1.7 ClusterAppList

ClusterAppList contains a list of ClusterApps.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹²⁸⁴	false
items		[...]ClusterApp (see page 1606)	true

¹²⁸² <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹²⁸³ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹²⁸⁴ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

11.1.1.8 CrossNamespaceGitRepositoryReference

CrossNamespaceGitRepositoryReference contains enough information to let you locate the typed referenced object at cluster level.

Field	Description	Scheme	Required
apiVersion	API version of the referent	string	false
kind	Kind of the referent	string	true
name	Name of the referent	string	true
namespace	Namespace of the referent, defaults to the Kustomization namespace	string	false

11.1.1.9 GenericAppSpec

GenericAppSpec defines the actual Application spec.

Field	Description	Scheme	Required
appId	AppID specifies the name of the application/workload	string	true
gitRepositoryRef	GitRepository is reference to the Flux's GitRepository, which in turn describes Git repository where the service resides	CrossNamespaceGitRepositoryReference (see page 1607)	true
version	Version depicts the version of the service in the semantic versioning format.	string	true

11.1.1.10 TypedLocalObjectReference

TypedLocalObjectReference contains enough information to let you locate the typed referenced object at cluster or local level.

Field	Description	Scheme	Required
apiVersion	API version of the referent	string	false
kind	Kind of the referent	string	true
name	Name of the referent	string	true

11.1.2 apps.kommander.mesosphere.io/v1alpha3

11.1.2.1 App

App is the Schema for a Service or Application in Kommander.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹²⁸⁵	false
spec		GenericAppSpec (see page 1611)	false
status		object	false

11.1.2.2 AppDeployment

AppDeployment is the Schema for a concrete installation of a Service or Application in Kommander.

¹²⁸⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹²⁸⁶	false
spec		AppDeploymentSpec (see page 1607)	false
status		object	false

11.1.2.3 AppDeploymentClusterCondition

Field	Description	Scheme	Required
lastTransitionTime	LastTransitionTime is the last time the condition transitioned from one status to another.	metav1.Time ¹²⁸⁷	false
status	Status of the condition, one of True, False, Unknown	string	true
type	Type indicates type of condition in CamelCase or in foo.example.com/CamelCase.	string	true

11.1.2.4 AppDeploymentList

AppDeploymentList contains a list of AppDeployments.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹²⁸⁸	false

¹²⁸⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹²⁸⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

¹²⁸⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

Field	Description	Scheme	Required
items		[...]AppDeployment (see page 1608)	true

11.1.2.5 AppDeploymentSpec

AppDeploymentSpec defines an instance of an Application.

Field	Description	Scheme	Required
appRef	AppRef provides reference to a ClusterApp or App object.	TypedLocalObjectReference (see page 1611)	true
clusterConfigOverrides	ClusterConfigOverrides defines a list of config overrides configmaps and label selectors to select clusters on which to apply the config overrides.	[...]ClusterConfigOverrides (see page 1609)	false
clusterSelector	ClusterSelector defines the label selectors to select clusters on which to enable the AppDeployment.	metav1.LabelSelector ¹²⁸⁹	false
configOverrides	ConfigOverrides allows a user to define a ConfigMap that contains configuration overrides at the workspace level.	corev1.LocalObjectReference ¹²⁹⁰	false

11.1.2.6 AppDeploymentStatus

AppDeploymentStatus defines the current state of the AppDeployment.

¹²⁸⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#labelselector-v1-meta>

¹²⁹⁰ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

Field	Description	Scheme	Required
clusters	Clusters tracks clusters that have the AppDeployment enabled along with a ref to the cluster config overrides configmap in the management cluster. Existence of cluster means that the cluster's default AppDeployment state has been processed. Nonexistence of cluster means that the cluster was newly attached and needs to be selected by default to enable the AppDeployment on it.	[...]ClusterDeploymentStatus (see page 1610)	false
observedGeneration	ObservedGeneration is the most recent generation observed for this AppDeployment. It corresponds to the AppDeployment's generation, which is updated on mutation by the API Server.	int64	false

11.1.2.7 AppList

AppList contains a list of Apps.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹²⁹¹	false
items		[...]App (see page 1608)	true

¹²⁹¹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

11.1.2.8 ClusterApp

ClusterApp is the Schema for a Service or Application in Kommander.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹²⁹²	false
spec		GenericAppSpec (see page 1611)	false
status		object	false

11.1.2.9 ClusterAppList

ClusterAppList contains a list of ClusterApps.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹²⁹³	false
items		[...]ClusterApp (see page 1612)	true

11.1.2.10 ClusterConfigOverrides

ClusterConfigOverrides defines the cluster config overrides configmap and label selector to select clusters on which to apply the overrides configmap.

Field	Description	Scheme	Required
clusterSelector	ClusterSelector defines the label selectors to select clusters on which to deploy the configmap.	metav1.LabelSelector ¹²⁹⁴	false

¹²⁹² <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹²⁹³ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹²⁹⁴ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#labelselector-v1-meta>

Field	Description	Scheme	Required
configMapName	ConfigMapName is the name of the cluster config overrides configmap.	string	true

11.1.2.11 ClusterDeploymentStatus

Field	Description	Scheme	Required
clusterConfigOverridesRef	ClusterConfigOverridesRef is set to reference the overrides ConfigMap in the management cluster	corev1.LocalObjectReference ¹²⁹⁵	false
conditions		[...]AppDeploymentClusterCondition (see page 1609)	false
name		string	true

11.1.2.12 CrossNamespaceGitRepositoryReference

CrossNamespaceGitRepositoryReference contains enough information to let you locate the typed referenced object at cluster level.

Field	Description	Scheme	Required
apiVersion	API version of the referent	string	false
kind	Kind of the referent	string	true
name	Name of the referent	string	true

¹²⁹⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

Field	Description	Scheme	Required
namespace	Namespace of the referent, defaults to the Kustomization namespace	string	false

11.1.2.13 GenericAppSpec

GenericAppSpec defines the actual Application spec.

Field	Description	Scheme	Required
appId	AppID specifies the name of the application/workload	string	true
gitRepositoryRef	GitRepository is reference to the Flux's GitRepository, which in turn describes Git repository where the service resides	CrossNamespaceGitRepositoryReference (see page 1613)	true
version	Version depicts the version of the service in the semantic versioning format.	string	true

11.1.2.14 TypedLocalObjectReference

TypedLocalObjectReference contains enough information to let you locate the typed referenced object at cluster or local level.

Field	Description	Scheme	Required
apiVersion	API version of the referent	string	false
kind	Kind of the referent	string	true

Field	Description	Scheme	Required
name	Name of the referent	string	true

11.1.3 dispatch.d2iq.io/v1alpha2

11.1.3.1 GitopsRepository

GitopsRepository represents an SCM repository that backs a gitops resource. A gitops resource can be a FluxCD HelmRelease or Kustomization resource.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹²⁹⁶	false
spec		GitopsRepositorySpec (see page 1612)	true

11.1.3.2 GitopsRepositoryList

GitopsRepositoryList contains a list of Repository.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹²⁹⁷	false
items		[...] GitopsRepository (see page 1615)	true

11.1.3.3 GitopsRepositorySpec

GitopsRepositorySpec defines the desired state of GitopsRepository.

¹²⁹⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹²⁹⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

Field	Description	Scheme	Required
cloneUrl	URL to clone the repository from.	string	false
secret	Reference to the secret to use when interacting with the Git provider API. The secret should contain the following fields: username: the username (if password is not a token). password: the password or token (required).	string	false
template	Template of the gitops resource backing the repository.	GitopsRolloutTemplate (see page 1613)	false

11.1.3.4 GitopsRolloutTemplate

Field	Description	Scheme	Required
path	Root path to fetch manifests from in the Git repository.	string	false
ref	The Git reference to checkout and monitor for changes, defaults to master branch. This field supersedes the <code>revision</code> field in v1alpha1 API and is a breaking change.	sourcectrl/v1beta1.GitRepositoryRef ¹²⁹⁸	false
suspend	Whether to suspend periodic or webhook-notified sync.	bool	false

¹²⁹⁸ <https://fluxcd.io/flux/components/source/api/#source.toolkit.fluxcd.io/v1beta1.GitRepositoryRef>

11.1.4 kommander.mesosphere.io/v1beta1

11.1.4.1 CAPIClusterReference

Field	Description	Scheme	Required
name		string	true
namespace		string	false

11.1.4.2 ClusterReference

ClusterReference holds a single reference to clusters provisioned via Kommander. Only one field is allowed to be set. Currently, only CAPI clusters are creatable, but this is left extensible for other provider types in the future.

Field	Description	Scheme	Required
capiCluster		CAPIClusterReference (see page 1617)	false

11.1.4.3 GenericClusterReference

GenericClusterReference sets the name of the cluster.

Field	Description	Scheme	Required
name		string	true

11.1.4.4 IngressSpec

IngressSpec holds settings for the cluster's Kommander managed Ingress.

Field	Description	Scheme	Required
certificateSecretRef	CertificateSecretRef is a reference to a secret of type TLS that holds a TLS certificate, private key and CA certificate. The certificate must be valid for the Ingress hostname. The secret must be located in the workspace's namespace on the target cluster.	corev1.LocalObjectReference ¹²⁹⁹	false
hostname	Hostname can be set to configure the cluster's Ingress with a custom domain name.	string	false
issuerRef	IssuerRef is a reference to an <code>Issuer</code> or <code>ClusterIssuer</code> on the target cluster to be used to create a <code>Certificate</code> for the cluster's Ingress. If the type is an <code>Issuer</code> , it must be located in the workspace's namespace on the cluster.	IssuerReference (see page 1616)	false

11.1.4.5 IngressStatus

IngressSpec holds information about the cluster's Kommander managed Ingress.

¹²⁹⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

Field	Description	Scheme	Required
address	Address is the main DNS name or IP address of the cluster's Ingress that should be used to access services on the cluster. If the field is empty, it means that the load balancer hasn't been populated yet.	string	false
caBundle	CABundle is a PEM encoded CA bundle which should be used to verify the TLS connection for the Ingress-served end-entity certificate.	[...]byte	false
useSystemCA	UseSystemCA is set to true if the Ingress end-entity certificate was issued by a public CA and it is expected that the root CA cert is included in a OS CA bundle (which should be used for TLS verification in that case). If this field is true the CABundle field is empty. default: false	bool	false

11.1.4.6 IssuerReference

IssuerReference is a reference to an issuer with a given name, kind and group.

Field	Description	Scheme	Required
group	Group of the issuer being referred to.	string	false

Field	Description	Scheme	Required
kind	Kind of the issuer being referred to.	string	false
name	Name of the issuer being referred to.	string	true

11.1.4.7 KommanderCluster

KommanderCluster is the Schema for the kommander clusters API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³⁰⁰	false
spec		KommanderClusterSpec (see page 1617)	false
status		KommanderClusterStatus (see page 1619)	false

11.1.4.8 KommanderClusterList

KommanderClusterList contains a list of Cluster.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³⁰¹	false
items		[...] KommanderCluster (see page 1620)	true

11.1.4.9 KommanderClusterSpec

KommanderClusterSpec defines the desired state of Cluster.

¹³⁰⁰ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹³⁰¹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

Field	Description	Scheme	Required
clusterRef		ClusterReference (see page 1617)	false
clusterTunnelConnectorRef	ClusterTunnelConnectorRef is a reference to TunnelConnector that should be used for connecting to cluster.	corev1.LocalObjectReference ¹³⁰²	false
clusterTunnelProxyConnectorRef	ClusterTunnelProxyConnectorRef is a reference to TunnelProxy that should be used for connecting to the cluster.	corev1.LocalObjectReference ¹³⁰³	false
ingress	Ingress configures the Kommander managed Ingress. If no value is provided, the default ingress will be provisioned.	IngressSpec (see page 1617)	false
kubeconfigRef		corev1.LocalObjectReference ¹³⁰⁴	false

1302 <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

1303 <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

1304 <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

Field	Description	Scheme	Required
namespaceRef	NamespaceRef will override the namespace on target cluster for the given Workspace. This is used in special circumstances such as when an essential cluster is attached and thus implicitly demoted from being a management cluster to a managed cluster thus forcing the kommander namespace to be recycled as the underlying workspace namespace.	corev1.LocalObjectReference ¹³⁰⁵	false

11.1.4.10 KommanderClusterStatus

KommanderClusterStatus defines the observed state of Cluster.

Field	Description	Scheme	Required
type	ClusterType represents the type of the cluster. E.g. EKS, GKE, etc.	string	false
conditions		[...] metav1.Condition ¹³⁰⁶	false
dextfaclientRef	DexTFAClientRef holds a reference to a DexClient provisioned for Traefik Forward Auth running on managed cluster.	corev1.ObjectReference ¹³⁰⁷	false

¹³⁰⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁰⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#condition-v1-meta>

¹³⁰⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectreference-v1-core>

Field	Description	Scheme	Required
ingress	Ingress holds information about the cluster's Ingress	IngressStatus (see page 1618)	false
kubefedclusterRef	KubefedClusterRef holds a reference to a kubefedcluster in the kubefed system namespace.	corev1.LocalObjectReference ¹³⁰⁸	false
clusterId	KubernetesClusterID is the stable cluster ID of the Kubernetes cluster that this Kommander cluster represents.	string	false
kubernetesVersion	KubernetesVersion is the Kubernetes version of the cluster.	string	false
phase	Phase represents the current phase of cluster actuation. E.g. Pending, Provisioning, Provisioned, Deleting, Failed, etc.	string	false
serviceEndpoints	ServiceEndpoints will be the addresses assigned to the Kubernetes exposed services	object	false

11.1.4.11 License

License is the Schema for the licenses API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³⁰⁹	false

¹³⁰⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁰⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

Field	Description	Scheme	Required
spec		LicenseSpec (see page 1622)	false
status		LicenseStatus (see page 1623)	false

11.1.4.12 LicenseCondition

Field	Description	Scheme	Required
lastTransitionTime		metav1.Time ¹³¹⁰	false
message		string	false
reason		string	false
status		string	true
type		string	true

11.1.4.13 LicenseExternalAWS

Field	Description	Scheme	Required
licenseArn	LicenseArn is a fully qualified AWS arn to a license for Kommander.	string	false

¹³¹⁰ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

11.1.4.14 LicenseExternalReference

Field	Description	Scheme	Required
awsLicense	AWSLicense holds a single reference to a license in AWS's License Manager	LicenseExternalAWS (see page 1624)	false
type	Type is the source of the external license referenced, i.e. AWS, GCP, Azure	string	true

11.1.4.15 LicenseList

LicenseList contains a list of License.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³¹¹	false
items		[...] License (see page 1623)	true

11.1.4.16 LicenseSpec

LicenseSpec defines the desired state of License.

Field	Description	Scheme	Required
externalLicenseRef	ExternalLicenseRef holds a reference to an external license	LicenseExternalReference (see page 1625)	false

¹³¹¹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

Field	Description	Scheme	Required
licenseRef	LicenseReference holds a single reference to the secret holding the license JWT	corev1.LocalObjectReference ¹³¹²	false

11.1.4.17 LicenseStatus

LicenseStatus defines the observed state of License.

Field	Description	Scheme	Required
clusterCapacity	Maximum number of clusters that the license allows.	int32	true
conditions	Conditions relevant to the license (currently used to track term breaches)	[...]LicenseCondition (see page 1624)	false
coreCapacity	Maximum number of cores that the license allows.	int32	true
customerId	The customer's ID. This is the customer name provided from Salesforce.	string	true
dkpLevel	The DKP license level.	string	true
endDate	End date of the licensing period.	metav1.Time ¹³¹³	false
licenseId	The license's ID as provided from Salesforce.	string	true

¹³¹² <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³¹³ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

Field	Description	Scheme	Required
productName	The product name the license is for.	string	true
productSKU	The product SKU the license is for.	string	true
startDate	Start date of the licensing period.	metav1.Time ¹³¹⁴	false
valid	Indicates whether the license is valid, i.e. the secret containing the JWT exists and the JWT carries a valid D2iQ signature. This does NOT indicate whether the license has expired or terms have been breached.	bool	true
version	The license's version ID as provided when the license was created.	string	true

11.1.4.18 PlacementSelector

PlacementSelector defines the fields to select clusters where to apply the project.

Field	Description	Scheme	Required
clusterSelector		metav1.LabelSelector ¹³¹⁵	false
clusters		[...] GenericClusterReference (see page 1617)	false

¹³¹⁴ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

¹³¹⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#labelselector-v1-meta>

11.1.4.19 VirtualGroup

VirtualGroup is the Schema for the virtualgroups API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³¹⁶	false
spec		VirtualGroupSpec (see page 1626)	false

11.1.4.20 VirtualGroupClusterRoleBinding

VirtualGroupClusterRoleBinding is the Schema for the virtualgroupclusterrolebindings API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³¹⁷	false
spec		VirtualGroupClusterRoleBindingSpec (see page 1626)	true

11.1.4.21 VirtualGroupClusterRoleBindingList

VirtualGroupClusterRoleBindingList contains a list of VirtualGroupClusterRoleBinding.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³¹⁸	false
items		[...] VirtualGroupClusterRoleBinding (see page 1628)	true

¹³¹⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹³¹⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹³¹⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

11.1.4.22 VirtualGroupClusterRoleBindingSpec

VirtualGroupClusterRoleBindingSpec defines the desired state of VirtualGroupClusterRoleBinding.

Field	Description	Scheme	Required
clusterRoleRef		corev1.LocalObjectReference ¹³¹⁹	true
placement		PlacementSelector (see page 1627)	false
virtualGroupRef		corev1.LocalObjectReference ¹³²⁰	true

11.1.4.23 VirtualGroupList

VirtualGroupList contains a list of VirtualGroup.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³²¹	false
items		[...]VirtualGroup (see page 1628)	true

11.1.4.24 VirtualGroupSpec

VirtualGroupSpec defines the desired state of VirtualGroup.

Field	Description	Scheme	Required
subjects		[...]rbacv1.Subject ¹³²²	false

¹³¹⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³²⁰ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³²¹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹³²² <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#subject-v1-rbac-authorization-k8s-io>

11.1.5 workspaces.kommander.mesosphere.io/v1alpha1

11.1.5.1 KommanderProjectRole

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³²³	false
spec		KommanderProjectRole Spec (see page 1627)	false
status		KommanderProjectRole Status (see page 1628)	false

11.1.5.2 KommanderProjectRoleList

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³²⁴	false
items		[...] KommanderProjectRole (see page 1630)	true

11.1.5.3 KommanderProjectRoleSpec

Field	Description	Scheme	Required
projectObjectVerbs		[...]string	false
rules		[...] rbacv1.PolicyRule ¹³²⁵	false

¹³²³ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹³²⁴ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹³²⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#policyrule-v1-rbac-authorization-k8s-io>

11.1.5.4 KommanderProjectRoleStatus

Field	Description	Scheme	Required
roleInProjectRef		corev1.LocalObjectReference ¹³²⁶	false
roleInWorkspaceRef		corev1.LocalObjectReference ¹³²⁷	false

11.1.5.5 KommanderWorkspaceRole

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³²⁸	false
spec		KommanderWorkspaceRoleSpec (see page 1629)	false
status		KommanderWorkspaceRoleStatus (see page 1629)	false

11.1.5.6 KommanderWorkspaceRoleList

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³²⁹	false
items		[...] KommanderWorkspaceRole (see page 1631)	true

¹³²⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³²⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³²⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹³²⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

11.1.5.7 KommanderWorkspaceRoleSpec

Field	Description	Scheme	Required
rules		[...] rbacv1.PolicyRule ¹³³⁰	false
workspaceObjectVerbs		[...]string	false

11.1.5.8 KommanderWorkspaceRoleStatus

Field	Description	Scheme	Required
clusterRoleRef		corev1.LocalObjectReference ¹³³¹	false
roleInWorkspaceRef		corev1.LocalObjectReference ¹³³²	false

11.1.5.9 Project

Project is a logical top-level container for a set of Kommander resources.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³³³	false
spec		ProjectSpec (see page 1632)	false
status		ProjectStatus (see page 1632)	false

¹³³⁰ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#policyrule-v1-rbac-authorization-k8s-io>

¹³³¹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³³² <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³³³ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

11.1.5.10 ProjectCondition

Field	Description	Scheme	Required
lastTransitionTime	Last time the condition transitioned from one status to another.	metav1.Time ¹³³⁴	false
message	A human readable message indicating details about the transition.	string	false
reason	The reason for the condition's last transition.	string	false
status	Status of the condition, one of True, False, Unknown.	string	true
type	Type of project condition.	ProjectConditionType	true

11.1.5.11 ProjectList

ProjectList is a list of Project objects.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³³⁵	false
items		[...] Project (see page 1632)	true

¹³³⁴ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

¹³³⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

11.1.5.12 ProjectRole

ProjectRole is the Schema for the workspaces ProjectRole API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³³⁶	false
spec		ProjectRoleSpec (see page 1631)	false
status		ProjectRoleStatus (see page 1632)	false

11.1.5.13 ProjectRoleList

ProjectRoleList contains a list of ProjectRole.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³³⁷	false
items		[...] ProjectRole (see page 1634)	true

11.1.5.14 ProjectRoleSpec

Field	Description	Scheme	Required
rules		[...] rbacv1.PolicyRule ¹³³⁸	false

¹³³⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹³³⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹³³⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#policyrule-v1-rbac-authorization-k8s-io>

11.1.5.15 ProjectRoleStatus

Field	Description	Scheme	Required
federatedRoleRef		corev1.LocalObjectReference ¹³³⁹	false

11.1.5.16 ProjectSpec

ProjectSpec describes the attributes on a Project.

Field	Description	Scheme	Required
namespaceName	NamespaceName specifies the optional namespace name to use for the project. This field is immutable, only settable on create.	string	false
placement		v1beta1.PlacementSelector ¹³⁴⁰	false
workspaceRef		corev1.LocalObjectReference ¹³⁴¹	false

11.1.5.17 ProjectStatus

ProjectStatus is information about the current status of a Project.

¹³³⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁴⁰ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/110854693/kommander.mesosphere.io+v1beta1#PlacementSelector>

¹³⁴¹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

Field	Description	Scheme	Required
conditions	Represents the latest available observations of a project's current state.	[...]ProjectCondition (see page 1633)	false
namespaceRef		corev1.LocalObjectReference ¹³⁴²	false

11.1.5.18 VirtualGroupKommanderClusterRoleBinding

VirtualGroupKommanderClusterRoleBinding is the Schema for the VirtualGroupWorkspaceRoleBinding API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³⁴³	false
spec		VirtualGroupKommanderClusterRoleBindingSpec (see page 1634)	true
status		VirtualGroupKommanderClusterRoleBindingStatus (see page 1634)	false

11.1.5.19 VirtualGroupKommanderClusterRoleBindingList

VirtualGroupKommanderClusterRoleBindingList contains a list of VirtualGroupKommanderClusterRoleBinding.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³⁴⁴	false

¹³⁴² <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁴³ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹³⁴⁴ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

Field	Description	Scheme	Required
items		[...]VirtualGroupKommanderClusterRoleBinding (see page 1636)	true

11.1.5.20 VirtualGroupKommanderClusterRoleBindingSpec

Field	Description	Scheme	Required
clusterRoleRef		corev1.LocalObjectReference ¹³⁴⁵	true
virtualGroupRef		corev1.LocalObjectReference ¹³⁴⁶	true

11.1.5.21 VirtualGroupKommanderClusterRoleBindingStatus

Field	Description	Scheme	Required
clusterRoleBindingRef		corev1.LocalObjectReference ¹³⁴⁷	false

11.1.5.22 VirtualGroupKommanderProjectRoleBinding

VirtualGroupKommanderProjectRoleBinding is the Schema to be used in the API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³⁴⁸	false

¹³⁴⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁴⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁴⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁴⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

Field	Description	Scheme	Required
spec		VirtualGroupKommanderProjectRoleBindingSpec (see page 1635)	true
status		VirtualGroupKommanderProjectRoleBindingStatus (see page 1636)	false

11.1.5.23 VirtualGroupKommanderProjectRoleBindingList

VirtualGroupKommanderProjectRoleBindingList contains a list of VirtualGroupKommanderProjectRoleBinding.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³⁴⁹	false
items		[...] VirtualGroupKommanderProjectRoleBinding (see page 1637)	true

11.1.5.24 VirtualGroupKommanderProjectRoleBindingSpec

Field	Description	Scheme	Required
kommanderProjectRoleRef		corev1.LocalObjectReference ¹³⁵⁰	true
virtualGroupRef		corev1.LocalObjectReference ¹³⁵¹	true

¹³⁴⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹³⁵⁰ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁵¹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

11.1.5.25 VirtualGroupKommanderProjectRoleBindingStatus

Field	Description	Scheme	Required
roleBindingInProjectRef		corev1.LocalObjectReference ¹³⁵²	false
roleBindingInWorkspaceRef		corev1.LocalObjectReference ¹³⁵³	false

11.1.5.26 VirtualGroupKommanderWorkspaceRoleBinding

VirtualGroupKommanderWorkspaceRoleBinding is the Schema to be used in the API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³⁵⁴	false
spec		VirtualGroupKommanderWorkspaceRoleBindingSpec (see page 1637)	true
status		VirtualGroupKommanderWorkspaceRoleBindingStatus (see page 1637)	false

11.1.5.27 VirtualGroupKommanderWorkspaceRoleBindingList

VirtualGroupKommanderWorkspaceRoleBindingList contains a list of VirtualGroupKommanderWorkspaceRoleBinding.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³⁵⁵	false

¹³⁵² <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁵³ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁵⁴ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹³⁵⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

Field	Description	Scheme	Required
items		[...]VirtualGroupKommanderWorkspaceRoleBinding (see page 1639)	true

11.1.5.28 VirtualGroupKommanderWorkspaceRoleBindingSpec

Field	Description	Scheme	Required
kommanderWorkspaceRoleRef		corev1.LocalObjectReference ¹³⁵⁶	true
virtualGroupRef		corev1.LocalObjectReference ¹³⁵⁷	true

11.1.5.29 VirtualGroupKommanderWorkspaceRoleBindingStatus

Field	Description	Scheme	Required
clusterRoleBindingRef		corev1.LocalObjectReference ¹³⁵⁸	false
roleBindingInWorkspaceRef		corev1.LocalObjectReference ¹³⁵⁹	false

11.1.5.30 VirtualGroupProjectRoleBinding

VirtualGroupProjectRoleBinding is the Schema for the VirtualGroupProjectRoleBinding API.

¹³⁵⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁵⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁵⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁵⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³⁶⁰	false
spec		VirtualGroupProjectRoleBindingSpec (see page 1638)	true
status		VirtualGroupProjectRoleBindingStatus (see page 1639)	false

11.1.5.31 VirtualGroupProjectRoleBindingList

VirtualGroupProjectRoleBindingList contains a list of VirtualGroupProjectRoleBinding.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³⁶¹	false
items		[...]VirtualGroupProjectRoleBinding (see page 1640)	true

11.1.5.32 VirtualGroupProjectRoleBindingSpec

Field	Description	Scheme	Required
projectRoleRef		corev1.LocalObjectReference ¹³⁶²	false
virtualGroupRef		corev1.LocalObjectReference ¹³⁶³	true

¹³⁶⁰ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹³⁶¹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹³⁶² <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁶³ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

Field	Description	Scheme	Required
workspaceRoleRef	WorkspaceRoleRef maybe a LocalObjectReference but the WorkspaceRole is not created in project namespace but in Workspace namespace. Local in LocalObjectReference means Local to project's workspace since there can only be one workspace the project is in.	corev1.LocalObjectReference ¹³⁶⁴	false

11.1.5.33 VirtualGroupProjectRoleBindingStatus

Field	Description	Scheme	Required
federatedRoleBindingRef		corev1.LocalObjectReference ¹³⁶⁵	false

11.1.5.34 VirtualGroupWorkspaceRoleBinding

VirtualGroupWorkspaceRoleBinding is the Schema for the VirtualGroupWorkspaceRoleBinding API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³⁶⁶	false
spec		VirtualGroupWorkspaceRoleBindingSpec (see page 1640)	true

¹³⁶⁴ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁶⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁶⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

Field	Description	Scheme	Required
status		VirtualGroupWorkspaceRoleBindingStatus (see page 1641)	false

11.1.5.35 VirtualGroupWorkspaceRoleBindingList

VirtualGroupWorkspaceRoleBindingList contains a list of VirtualGroupWorkspaceRoleBinding.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³⁶⁷	false
items		[...] VirtualGroupWorkspaceRoleBinding (see page 1642)	true

11.1.5.36 VirtualGroupWorkspaceRoleBindingSpec

Field	Description	Scheme	Required
placement		v1beta1.PlacementSelector ¹³⁶⁸	false
virtualGroupRef		corev1.LocalObjectReference ¹³⁶⁹	true
workspaceRoleRef		corev1.LocalObjectReference ¹³⁷⁰	true

¹³⁶⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹³⁶⁸ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/110854693/kommander.mesosphere.io+v1beta1#PlacementSelector>

¹³⁶⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁷⁰ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

11.1.5.37 VirtualGroupWorkspaceRoleBindingStatus

Field	Description	Scheme	Required
federatedClusterRoleBindingRef		corev1.LocalObjectReference ¹³⁷¹	false

11.1.5.38 Workspace

Workspace is the Schema for the workspaces API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³⁷²	false
spec		WorkspaceSpec (see page 1644)	false
status		WorkspaceStatus (see page 1644)	false

11.1.5.39 WorkspaceCondition

Field	Description	Scheme	Required
lastTransitionTime	Last time the condition transitioned from one status to another.	metav1.Time ¹³⁷³	false
message	A human readable message indicating details about the transition.	string	false

¹³⁷¹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

¹³⁷² <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

¹³⁷³ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

Field	Description	Scheme	Required
reason	The reason for the condition's last transition.	string	false
status	Status of the condition, one of True, False, Unknown.	string	true
type	Type of workspace condition.	string	true

11.1.5.40 WorkspaceList

WorkspaceList contains a list of Workspace.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³⁷⁴	false
items		[...]Workspace (see page 1644)	true

11.1.5.41 WorkspaceRole

WorkspaceRole is the Schema for the workspaces API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta ¹³⁷⁵	false
spec		WorkspaceRoleSpec (see page 1643)	false
status		WorkspaceRoleStatus (see page 1643)	false

¹³⁷⁴ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹³⁷⁵ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

11.1.5.42 WorkspaceRoleList

WorkspaceRoleList contains a list of WorkspaceRole.

Field	Description	Scheme	Required
metadata		metav1.ListMeta ¹³⁷⁶	false
items		[...] WorkspaceRole (see page 1645)	true

11.1.5.43 WorkspaceRoleSpec

Field	Description	Scheme	Required
aggregationRule		rbacv1.AggregationRule ¹³⁷⁷	false
rules		[...] rbacv1.PolicyRule ¹³⁷⁸	false

11.1.5.44 WorkspaceRoleStatus

Field	Description	Scheme	Required
federatedClusterRoleRef		corev1.LocalObjectReference ¹³⁷⁹	false

¹³⁷⁶ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

¹³⁷⁷ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#aggregationrule-v1-rbac-authorization-k8s-io>

¹³⁷⁸ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#policyrule-v1-rbac-authorization-k8s-io>

¹³⁷⁹ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

11.1.5.45 WorkspaceSpec

Field	Description	Scheme	Required
clusterLabels		object	false
namespaceName	NamespaceName specifies the optional namespace name to use for the workspace. This field is immutable, only settable on create.	string	false

11.1.5.46 WorkspaceStatus

Field	Description	Scheme	Required
conditions	Represents the latest available observations of a workspace's current state.	[...]WorkspaceCondition (see page 1644)	false
namespaceRef		corev1.LocalObjectReference ¹³⁸⁰	false

11.2 CLI Commands

11.2.1 CLI Commands for DKP

The table of contents on the left lists the command groupings. Inside each section, you will find the individual commands for most actions. For KIB CLI, refer to the documentation section for [Konvoy Image Builder CLI](#) (see page 1497) commands.

- [dkp attach](#) (see page 1645)
- [dkp check](#) (see page 1646)
- [dkp cluster](#) (see page 1648)
- [dkp completion](#) (see page 1649)

¹³⁸⁰ <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

- [dkp config](#) (see page 1654)
- [dkp create](#) (see page 1657)
- [dkp delete](#) (see page 1697)
- [dkp describe](#) (see page 1701)
- [dkp detach](#) (see page 1702)
- [dkp diagnose](#) (see page 1703)
- [dkp edit](#) (see page 1706)
- [dkp get](#) (see page 1707)
- [dkp import](#) (see page 1712)
- [dkp install](#) (see page 1713)
- [dkp move](#) (see page 1715)
- [dkp open](#) (see page 1716)
- [dkp push](#) (see page 1717)
- [dkp scale](#) (see page 1720)
- [dkp serve](#) (see page 1721)
- [dkp update](#) (see page 1722)
- [dkp upgrade](#) (see page 1738)
- [dkp version](#) (see page 1749)
- [dkp upload](#) (see page 1749)

11.2.2 dkp attach

Attach one of [cluster]

11.2.2.1 Options

<code>--config string</code>	Config file to use (default <code>"/home/runner/.kommander/config"</code>)
<code>--context string</code>	The name of the kubeconfig context to use
<code>-h, --help</code>	help for attach
<code>--kubeconfig string</code>	Path to the kubeconfig file to use for CLI requests.
<code>--request-timeout string</code>	The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
<code>-v, --verbose int</code>	Output verbosity

11.2.2.2 SEE ALSO

- [dkp attach cluster](#) (see page 1645) - Attach a cluster

11.2.2.3 dkp attach cluster

Attach a cluster


```
dkp attach cluster -n NAME --attached-kubeconfig FILENAME [flags]
```

11.2.2.3.1 Options

```

    --attached-kubeconfig string  Path of the kubeconfig file of the cluster to be
attached
    -h, --help                    help for cluster
    -n, --name string             Desired name of the attached cluster
    -w, --workspace string        Name of the workspace of the attached cluster

```

11.2.2.3.2 Options inherited from parent commands

```

    --config string              Config file to use (default "/home/
runner/.kommander/config")
    --context string            The name of the kubeconfig context to use
    --kubeconfig string         Path to the kubeconfig file to use for CLI requests.
    --request-timeout string    The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g.
1s, 2m, 3h). A value of zero means don't timeout requests.
    -v, --verbose int         Output verbosity

```

11.2.2.3.3 SEE ALSO

- [dkp attach](#) (see page 1648) - Attach one of [cluster]

11.2.3 dkp check

Check, one of [cluster]

11.2.3.1 Options

```

    -h, --help                    help for check
    -v, --verbose int           Output verbosity

```

11.2.3.2 SEE ALSO

- [dkp check cluster](#) (see page 1647) - Check a cluster, one of [fips]

11.2.3.3 dkp check cluster

Check a cluster, one of [fips]

11.2.3.3.1 Options

```
-h, --help help for cluster
```

11.2.3.3.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.3.3.3 SEE ALSO

- [dkp check \(see page 1649\)](#) - Check, one of [cluster]
- [dkp check cluster fips \(see page 1647\)](#) - Validate the components in your cluster are FIPS compliant

11.2.3.3.4 dkp check cluster fips

Validate the components in your cluster are FIPS compliant

11.2.3.3.4.1 Synopsis

The check cluster fips command is used to validate that specific components and services are FIPS compliant by checking the signatures of the files against a signed signature file, and checking that services are using the certified algorithms.

Examples:

To use the built-in signature files **for** supported operating systems:

```
dkp check cluster fips
```

To use a custom signature file, named "manifest-rhel-84.json.asc":

```
dkp check cluster fips \  
  --signature-file manifest-rhel-84.json.asc \  
  --signature-configmap myconfigmap
```

The file will be copied to the ConfigMap. To use an existing ConfigMap:

```
dkp check cluster fips \
  --signature-configmap myconfigmap
```

The validation will be re-checked against the existing signature data.

```
dkp check cluster fips [flags]
```

11.2.3.3.4.2 Options

<code>-h, --help</code>	help for fips
<code>--kubeconfig string</code>	Path to the kubeconfig file for the fips cluster. If unspecified, default discovery rules apply.
<code>-n, --namespace string</code>	If present, the namespace scope for this CLI request. (default "default")
<code>--output-configmap string</code>	ConfigMap to store result of the fips check. (default "check-cluster-fips-output") (DEPRECATED: This flag will be removed in a future release.)
<code>--signature-configmap string</code>	ConfigMap with fips signature data to verify.
<code>--signature-file string</code>	File containing fips signature data.
<code>--timeout duration</code>	The length of time to wait before giving up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 10m0s)

11.2.3.3.4.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.3.3.4.4 SEE ALSO

- [dkp check cluster](#) (see page 1650) - Check a cluster, one of [fips]

11.2.4 dkp cluster

Get cluster information

11.2.4.1 Options

<code>--config string</code>	Config file to use (default "/home/runner/.kommander/config")
<code>--context string</code>	The name of the kubeconfig context to use

```

-h, --help                help for cluster
--kubeconfig string       Path to the kubeconfig file to use for CLI requests.
--request-timeout string  The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g.
1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int       Output verbosity

```

11.2.4.2 SEE ALSO

- [dkp cluster type](#) (see page 1649) - Retrieve cluster type

11.2.4.3 dkp cluster type

Retrieve cluster type

```
dkp cluster type [flags]
```

11.2.4.3.1 Options

```
-h, --help  help for type
```

11.2.4.3.2 Options inherited from parent commands

```

--config string          Config file to use (default "/home/
runner/.kommander/config")
--context string         The name of the kubeconfig context to use
--kubeconfig string      Path to the kubeconfig file to use for CLI requests.
--request-timeout string The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g.
1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int      Output verbosity

```

11.2.4.3.3 SEE ALSO

- [dkp cluster](#) (see page 1651) - Get cluster information

11.2.5 dkp completion

Generate the autocompletion script for the specified shell

11.2.5.1 Synopsis

Generate the autocompletion script for dkp for the specified shell. See each sub-command's help for details on how to use the generated script.

11.2.5.2 Options

```
-h, --help help for completion
```

11.2.5.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.5.4 SEE ALSO

- [dkp completion bash \(see page 1651\)](#) - Generate the autocompletion script for bash
- [dkp completion fish \(see page 1650\)](#) - Generate the autocompletion script for fish
- [dkp completion powershell \(see page 1653\)](#) - Generate the autocompletion script for powershell
- [dkp completion zsh \(see page 1652\)](#) - Generate the autocompletion script for zsh

11.2.5.5 dkp completion fish

Generate the autocompletion script for fish

11.2.5.5.1 Synopsis

Generate the autocompletion script for the fish shell.

To load completions in your current shell session:

```
dkp completion fish | source
```

To load completions for every new session, execute once:

```
dkp completion fish > ~/.config/fish/completions/dkp.fish
```

You will need to start a new shell for this setup to take effect.

```
dkp completion fish [flags]
```

11.2.5.5.2 Options

```
-h, --help           help for fish
--no-descriptions   disable completion descriptions
```

11.2.5.5.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.5.5.4 SEE ALSO

- [dkp completion](#) (see page 1652) - Generate the autocompletion script for the specified shell

11.2.5.6 dkp completion bash

Generate the autocompletion script for bash

11.2.5.6.1 Synopsis

Generate the autocompletion script for the bash shell.

This script depends on the 'bash-completion' package. If it is not installed already, you can install it via your OS's package manager.

To load completions in your current shell session:

```
source <(dkp completion bash)
```

To load completions for every new session, execute once:

11.2.5.6.1.1 Linux:

```
dkp completion bash > /etc/bash_completion.d/dkp
```

11.2.5.6.1.2 macOS:

```
dkp completion bash > $(brew --prefix)/etc/bash_completion.d/dkp
```

You will need to start a new shell for this setup to take effect.

```
dkp completion bash
```

11.2.5.6.2 Options

```
-h, --help          help for bash
--no-descriptions  disable completion descriptions
```

11.2.5.6.3 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.5.6.4 SEE ALSO

- [dkp completion](#) (see page 1652) - Generate the autocompletion script for the specified shell

11.2.5.7 dkp completion zsh

Generate the autocompletion script for zsh

11.2.5.7.1 Synopsis

Generate the autocompletion script for the zsh shell.

If shell completion is not already enabled in your environment you will need to enable it. You can execute the following once:

```
echo "autoload -U compinit; compinit" >> ~/.zshrc
```

To load completions in your current shell session:

```
source <(dkp completion zsh); compdef _dkp dkp
```

To load completions for every new session, execute once:

11.2.5.7.1.1 Linux:

```
dkp completion zsh > "${fpath[1]}/_dkp"
```

11.2.5.7.1.2 macOS:

```
dkp completion zsh > $(brew --prefix)/share/zsh/site-functions/_dkp
```

You will need to start a new shell for this setup to take effect.

```
dkp completion zsh [flags]
```

11.2.5.7.2 Options

```
-h, --help           help for zsh
--no-descriptions   disable completion descriptions
```

11.2.5.7.3 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.5.7.4 SEE ALSO

- [dkp completion](#) (see page 1652) - Generate the autocompletion script for the specified shell

11.2.5.8 dkp completion powershell

Generate the autocompletion script for powershell

11.2.5.8.1 Synopsis

Generate the autocompletion script for powershell.

To load completions in your current shell session:


```
dkp completion powershell | Out-String | Invoke-Expression
```

To load completions for every new session, add the output of the above command to your powershell profile.

```
dkp completion powershell [flags]
```

11.2.5.8.2 Options

```
-h, --help           help for powershell
--no-descriptions   disable completion descriptions
```

11.2.5.8.3 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.5.8.4 SEE ALSO

- [dkp completion](#) (see page 1652) - Generate the autocompletion script for the specified shell

11.2.6 dkp config

Manage DKP Kommander's configuration

11.2.6.1 Options

```
--config string      Config file to use (default "/home/runner/.kommander/config")
--context string     The name of the kubeconfig context to use
-h, --help           help for config
--kubeconfig string  Path to the kubeconfig file to use for CLI requests.
--request-timeout string The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int  Output verbosity
```

11.2.6.2 SEE ALSO

- [dkp config get](#) (see page 1655) - Retrieve DKP Kommander's configuration

- [dkp config set](#) (see page 1656) - Modify DKP Kommander's configuration

11.2.6.3 dkp config get

Retrieve DKP Kommander's configuration

11.2.6.3.1 Options

```
-h, --help  help for get
```

11.2.6.3.2 Options inherited from parent commands

```

--config string          Config file to use (default "/home/
runner/.kommander/config")
--context string         The name of the kubeconfig context to use
--kubeconfig string      Path to the kubeconfig file to use for CLI requests.
--request-timeout string The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g.
1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int       Output verbosity

```

11.2.6.3.3 SEE ALSO

- [dkp config](#) (see page 1657) - Manage DKP Kommander's configuration
- [dkp config get default-workspace](#) (see page 1655) - Displays the name of the configured default Workspace

11.2.6.3.4 dkp config get default-workspace

Displays the name of the configured default Workspace

```
dkp config get default-workspace [flags]
```

11.2.6.3.4.1 Options

```
-h, --help  help for default-workspace
```

11.2.6.3.4.2 Options inherited from parent commands

<code>--config string</code>	Config file to use (default <code>"/home/runner/.kommander/config"</code>)
<code>--context string</code>	The name of the kubeconfig context to use
<code>--kubeconfig string</code>	Path to the kubeconfig file to use for CLI requests.
<code>--request-timeout string</code>	The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
<code>-v, --verbose int</code>	Output verbosity

11.2.6.3.4.3 SEE ALSO

- [dkp config get](#) (see page 1658) - Retrieve DKP Kommander's configuration

11.2.6.4 dkp config set

Modify DKP Kommander's configuration

11.2.6.4.1 Options

```
-h, --help help for set
```

11.2.6.4.2 Options inherited from parent commands

<code>--config string</code>	Config file to use (default <code>"/home/runner/.kommander/config"</code>)
<code>--context string</code>	The name of the kubeconfig context to use
<code>--kubeconfig string</code>	Path to the kubeconfig file to use for CLI requests.
<code>--request-timeout string</code>	The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
<code>-v, --verbose int</code>	Output verbosity

11.2.6.4.3 SEE ALSO

- [dkp config](#) (see page 1657) - Manage DKP Kommander's configuration
- [dkp config set default-workspace](#) (see page 1657) - Set the name of the default Workspace to use in all commands when none is provided

11.2.6.4.4 dkp config set default-workspace

Set the name of the default Workspace to use in all commands when none is provided

```
dkp config set default-workspace [WORKSPACE] [flags]
```

11.2.6.4.4.1 Options

```
-h, --help help for default-workspace
```

11.2.6.4.4.2 Options inherited from parent commands

```

--config string          Config file to use (default "/home/
runner/.kommander/config")
--context string         The name of the kubeconfig context to use
--kubeconfig string      Path to the kubeconfig file to use for CLI requests.
--request-timeout string The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g.
1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int      Output verbosity

```

11.2.6.4.4.3 SEE ALSO

- [dkp config set](#) (see page 1659) - Modify DKP Kommander's configuration

11.2.7 dkp create

Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

11.2.7.1 Options

```
-h, --help help for create
```

11.2.7.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.7.3 SEE ALSO

- [dkp create appdeployment](#) (see page 1659) - Create an AppDeployment
- [dkp create bootstrap](#) (see page 1661) - Create bootstrap cluster
- [dkp create capi-components](#) (see page 1660) - Create the CAPI components in the cluster
- [dkp create chart-bundle](#) (see page 1683) - Create charts bundle based on a catalog applications git repository
- [dkp create cluster](#) (see page 1662) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]
- [dkp create image-bundle](#) (see page 1661) - Create an image bundle
- [dkp create nodepool](#) (see page 1684) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]
- [dkp create workspace](#) (see page 1658) - Create a Workspace

11.2.7.4 dkp create workspace

Create a Workspace

```
dkp create workspace WORKSPACE_NAME [flags]
```

11.2.7.4.1 Options

```

--config string          Config file to use (default "/home/runner/.kommander/config")
--context string        The name of the kubeconfig context to use
--dry-run              Export in YAML format to stdout
-h, --help             help for workspace
--kubeconfig string    Path to the kubeconfig file to use for CLI requests.
-n, --namespace string Name of the Namespace to create for the workspace
-o, --output string     Output format. One of: yaml|json (default "yaml")
--request-timeout string The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int    Output verbosity

```

11.2.7.4.2 SEE ALSO

- [dkp create](#) (see page 1660) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

11.2.7.5 dkp create appdeployment

Create an AppDeployment

11.2.7.5.1 Synopsis

Creates an AppDeployment in a workspace or project.

When [-clusters C1,C2..] is not specified, it targets all existing clusters in the specified workspace or project.

```
dkp create appdeployment APPDEPLOYMENT_NAME --app NAME [flags]
```

11.2.7.5.2 Options

```

    --add-cluster-config-overrides strings  Comma-separated list of mappings of
kommanderCluster name to cluster configuration override ConfigMap name to apply to
the targeting clusters. (e.g., cluster-1:override-1-cm,cluster-2:override-2-cm)(only
valid for dkp clusters version 2.3.x and up) (default [])
    -a, --app string                        Name of the App to deploy
    --clusters strings                      List of names of kommanderClusters to
select for the command. (e.g., cluster-1,cluster-2)(only valid for dkp clusters
version 2.3.x and up) (default [])
    --config string                        Config file to use (default "/home/
runner/.kommander/config")
    -c, --config-overrides string          Name of the ConfigMap used to override
default configuration of the App
    --context string                       The name of the kubeconfig context to
use
    --dry-run                              Export in YAML format to stdout
    -h, --help                             help for appdeployment
    --kubeconfig string                    Path to the kubeconfig file to use for
CLI requests.
    -o, --output string                    Output format. One of: yaml|json
(default "yaml")
    -p, --project string                   Name of the project to create the
AppDeployment in. Requires workspace flag (workspace that the project belongs to).
    --request-timeout string               The length of time to wait before
giving up on a single server request. Non-zero values should contain a corresponding
time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
    -v, --verbose int                    Output verbosity

```

```
-w, --workspace string          Name of the workspace to create the
AppDeployment in
```

11.2.7.5.3 SEE ALSO

- [dkp create](#) (see page 1660) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

11.2.7.6 dkp create capi-components

Create the CAPI components in the cluster

```
dkp create capi-components [flags]
```

11.2.7.6.1 Options

```
--aws-service-endpoints string    Custom AWS service endpoints in a semi-colon
separated format: ${SigningRegion1}:${ServiceID1}=${URL},${ServiceID2}=${URL};$
{SigningRegion2}...
-h, --help                        help for capi-components
--http-proxy string              HTTP proxy for CAPI controllers
--https-proxy string            HTTPS proxy for CAPI controllers
--kubeconfig string             Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--no-proxy strings              No Proxy list for CAPI controllers (default
[])
--timeout duration              The length of time to wait before giving up.
Zero means wait forever (e.g. 300s, 30m, 3h). (default 10m0s)
-v, --verbose int              Output verbosity
--wait                          If true, wait for operations to complete
before returning. (default true)
--with-aws-bootstrap-credentials Set true to use AWS bootstrap credentials
from your environment. When false, the instance profile of the EC2 instance where the
CAPA controller is scheduled on will be used instead.
--with-gcp-bootstrap-credentials Set true to use GCP bootstrap credentials
from your environment. When false, the service account of the VM instance where the
CAPG controller is scheduled on will be used instead.
```

11.2.7.6.2 SEE ALSO

- [dkp create](#) (see page 1660) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

11.2.7.7 dkp create image-bundle

Create an image bundle

```
dkp create image-bundle [flags]
```

11.2.7.7.1 Options

```
-h, --help                help for image-bundle
--image-pull-concurrency int Image pull concurrency (default 1)
--images-file string      File containing list of images to create bundle
from, either as YAML configuration or a simple list of images
--output-file string      Output file to write image bundle to (default
"images.tar")
--overwrite              Overwrite image bundle file if it already exists
--platform platformSlice platforms to download images (required format:
<os>/<arch>[/<variant>]) (default [linux/amd64])
-v, --verbose int        Output verbosity
```

11.2.7.7.2 SEE ALSO

- [dkp create](#) (see page 1660) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

11.2.7.8 dkp create bootstrap

Create bootstrap cluster

```
dkp create bootstrap [flags]
```

11.2.7.8.1 Options

```
--aws-service-endpoints string Custom AWS service endpoints in a semi-colon
separated format: ${SigningRegion1}:${ServiceID1}=${URL},${ServiceID2}=${URL};$
{SigningRegion2}...
-h, --help                help for bootstrap
--http-proxy string       HTTP proxy for CAPI controllers
--https-proxy string      HTTPS proxy for CAPI controllers
```



```

--kind-cluster-image string      Kind node image for the bootstrap cluster (d
efault "mesosphere/konvoy-bootstrap:v2.6.2-rc.4")
--kubeconfig string             Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--no-proxy strings              No Proxy list for CAPI controllers (default
[])
--timeout duration              The length of time to wait before giving up.
Zero means wait forever (e.g. 300s, 30m, 3h). (default 10m0s)
-v, --verbose int              Output verbosity
--wait                           If true, wait for operations to complete
before returning. (default true)
--with-aws-bootstrap-credentials Set true to use AWS bootstrap credentials
from your environment. When false, the instance profile of the EC2 instance where the
CAPA controller is scheduled on will be used instead.
--with-gcp-bootstrap-credentials Set true to use GCP bootstrap credentials
from your environment. When false, the service account of the VM instance where the
CAPG controller is scheduled on will be used instead.

```

11.2.7.8.2 SEE ALSO

- [dkp create](#) (see page 1660) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

11.2.7.9 dkp create cluster

Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.9.1 Options

```

-h, --help                help for cluster
-v, --verbose int        Output verbosity

```

11.2.7.9.2 SEE ALSO

- [dkp create](#) (see page 1660) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]
- [dkp create cluster aks](#) (see page 1673) - Create a Konvoy cluster in AKS
- [dkp create cluster aws](#) (see page 1665) - Create a Konvoy cluster in AWS
- [dkp create cluster azure](#) (see page 1667) - Create a Konvoy cluster in Azure
- [dkp create cluster eks](#) (see page 1663) - Create a Konvoy cluster in EKS
- [dkp create cluster gcp](#) (see page 1678) - Create a Konvoy cluster in GCP
- [dkp create cluster preprovisioned](#) (see page 1675) - Create a Konvoy cluster on pre-provisioned infrastructure
- [dkp create cluster vcd](#) (see page 1680) - Create a Konvoy cluster in VMware Cloud Director

- [dkp create cluster vsphere](#) (see page 1670) - Create a Konvoy cluster in vSphere

11.2.7.9.3 dkp create cluster eks

Create a Konvoy cluster in EKS

```
dkp create cluster eks [flags]
```

11.2.7.9.3.1 Options

```

--additional-security-group-ids strings  A comma separated list of existing
security group IDs to use for machines in addition to those created automatically (de
fault [])
--additional-tags stringToString         Tags to apply to the provisioned
infrastructure (default [])
--allow-missing-template-keys           If true, ignore any errors in
templates when a field or map key is missing in the template. Only applies to goLang
and jsonpath output formats. (default true)
--aws-service-endpoints string          Custom AWS service endpoints in a
semi-colon separated format: ${SigningRegion1}:${ServiceID1}=${URL},${ServiceID2}=${
URL};${SigningRegion2}...
-c, --cluster-name name                 Name used to prefix the cluster and
all the created resources.
--dry-run                               Only print the objects that would be
created, without creating them.
-h, --help                             help for eks
--http-proxy string                    HTTP proxy for CAPI controllers
--https-proxy string                   HTTPS proxy for CAPI controllers
--kind-cluster-image string            Kind node image for the bootstrap
cluster (default "mesosphere/konvoy-bootstrap:v2.6.2-rc.4")
--kubeconfig string                    Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply.
--kubernetes-pod-network-cidr cidr      The Kubernetes Pod network CIDR to
use in the cluster (default 192.168.0.0/16)
--kubernetes-service-cidr cidr         The Kubernetes Service CIDR to use in
the cluster (default 10.96.0.0/12)
--kubernetes-version string            Kubernetes version (default "1.25.11")
-n, --namespace string                 If present, the namespace scope for
this CLI request. (default "default")
--no-proxy strings                      No Proxy list for CAPI controllers (d
efault [])
-o, --output string                    Output format. One of: (json, yaml,
name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-
json, jsonpath-file).
--output-directory string              Used with --output=json|yaml. The
directory where to output resources to files. The directory must already exist.

```

```

--region string                AWS region to deploy cluster to
(default "us-west-2")
--registry-mirror-cacert file  CA certificate chain to use while
communicating with the registry mirror using TLS
--registry-mirror-password string Password to authenticate to the
registry mirror with
--registry-mirror-url url      URL of a container registry to use as
a mirror in the cluster
--registry-mirror-username string Username to authenticate to the
registry mirror with
--show-managed-fields          If true, keep the managedFields when
printing objects in JSON or YAML format.
--subnet-ids strings           A comma separated list of existing
subnet IDs to use for the kube-apiserver ELB and all control-plane and worker nodes (d
efault [])
--template string              Template string or path to template
file to use when -o=go-template, -o=go-template-file. The template format is goLang
templates [http://golang.org/pkg/text/template/#pkg-overview].
--timeout duration             The length of time to wait before
giving up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
--vpc-id string                Existing VPC ID to use for the
cluster
--wait                          If true, wait for operations to
complete before returning. (default true)
--with-aws-bootstrap-credentials Set true to use AWS bootstrap
credentials from your environment. When false, the instance profile of the EC2
instance where the CAPA controller is scheduled on will be used instead.
--with-gcp-bootstrap-credentials Set true to use GCP bootstrap
credentials from your environment. When false, the service account of the VM instance
where the CAPG controller is scheduled on will be used instead.
--worker-availability-zone string The AvailabilityZone in the region to
deploy the worker nodes to, if not set a random one will be selected (ex. us-west-2a)
--worker-iam-instance-profile string Name of the IAM instance profile to
assign to worker machines. (default "nodes.cluster-api-provider-aws.sigs.k8s.io")
--worker-instance-type string   Worker machine instance type (default
"m5.2xlarge")
--worker-replicas int         Number of workers (default 4)

```

11.2.7.9.3.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.7.9.3.3 SEE ALSO

- [dkp create cluster](#) (see page 1665) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.9.4 dkp create cluster aws

Create a Konvoy cluster in AWS

```
dkp create cluster aws [flags]
```

11.2.7.9.4.1 Options

<code>--additional-security-group-ids</code> strings	A comma separated list of existing security group IDs to use for machines in addition to those created automatically (default [])
<code>--additional-tags</code> stringToString	Infrastructure tags to apply to the provisioned infrastructure (default [])
<code>--allow-missing-template-keys</code>	If true , ignore any errors in templates when a field or map key is missing in the template. Only applies to goyaml and jsonpath output formats. (default true)
<code>--ami</code> string	AMI ID to use for machines
<code>--ami-base-os</code> string	Base OS used in search of AMIs.
Examples: 'centos-7', 'ubuntu-18.04', 'ubuntu-20.04'	
<code>--ami-format</code> string	Query string used in search of AMIs. Example: When <code>--ami-base-os='rhel8.4'</code> , then the string <code>'prefix-{{.BaseOS}}-?{{.K8sVersion}}-*'</code> matches any AMIs with the Name <code>'prefix-rhel8.4-1.26.14'</code>
<code>--ami-owner</code> string	ID of AWS account used in search of AMIs
<code>--aws-service-endpoints</code> string	Custom AWS service endpoints in a semi-colon separated format: <code> \${SigningRegion1}:\${ServiceID1}=\${URL},\${ServiceID2}=\${URL};\${SigningRegion2}...</code>
<code>--certificate-renew-interval</code> int	The interval number of days Kubernetes managed PKI certificates are renewed. For example, an Interval value of <code>30</code> means the certificates will be refreshed every <code>30</code> days. A value of <code>0</code> disables the feature. (default 0)
<code>-c, --cluster-name</code> name	Name used to prefix the cluster and all the created resources.
<code>--control-plane-http-proxy</code> string	HTTP proxy for control plane machines
<code>--control-plane-https-proxy</code> string	HTTPS proxy for control plane machines
<code>--control-plane-iam-instance-profile</code> string	Name of the IAM instance profile to assign to control plane machines. (default "control-plane.cluster-api-provider-aws.sigs.k8s.io")
<code>--control-plane-instance-type</code> string	Control Plane machine instance type (default "m5.xlarge")
<code>--control-plane-no-proxy</code> strings	No Proxy list for control plane machines (default [])
<code>--control-plane-replicas</code> int	Number of control plane nodes (default 3)

```

--dry-run                Only print the objects that would
be created, without creating them.
--etcd-image-repository string  The image repository to use for
pulling the etcd image
--etcd-version string        The version of etcd to use.
--extra-sans strings        A comma separated list of
additional Subject Alternative Names for the API Server signing cert (default [])
-h, --help                help for aws
--http-proxy string        HTTP proxy for CAPI controllers
--https-proxy string       HTTPS proxy for CAPI controllers
--internal-load-balancer    Make the control plane load
balancer internal, i.e., reachable only within the VPC.
--kind-cluster-image string  Kind node image for the bootstrap
cluster (default "mesosphere/konvoy-bootstrap:v2.6.2-rc.4")
--kubeconfig string        Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply. This flag is
ignored if used with the --self-managed flag.
--kubernetes-image-repository string  The image repository to use for
pulling kubernetes images
--kubernetes-pod-network-cidr cidr  The Kubernetes Pod network CIDR
to use in the cluster (default 192.168.0.0/16)
--kubernetes-service-cidr cidr     The Kubernetes Service CIDR to
use in the cluster (default 10.96.0.0/12)
--kubernetes-version string        Kubernetes version (default
"1.26.14")
-n, --namespace string            If present, the namespace scope
for this CLI request. (default "default")
--no-proxy strings                No Proxy list for CAPI
controllers (default [])
-o, --output string                Output format. One of: (json,
yaml, name, go-template, go-template-file, template, templatefile, jsonpath,
jsonpath-as-json, jsonpath-file).
--output-directory string          Used with --output=json|yaml. The
directory where to output resources to files. The directory must already exist.
--region string                    AWS region to deploy cluster to
(default "us-west-2")
--registry-mirror-cacert file      CA certificate chain to use while
communicating with the registry mirror using TLS
--registry-mirror-password string  Password to authenticate to the
registry mirror with
--registry-mirror-url url          URL of a container registry to
use as a mirror in the cluster
--registry-mirror-username string  Username to authenticate to the
registry mirror with
--self-managed                    When set to true, the required
prerequisites are created before creating the cluster and the resulting cluster has
all necessary components deployed onto itself, so it can manage its own cluster
lifecycle. When set to false, a management cluster is used. (default false)
--show-managed-fields              If true, keep the managedFields
when printing objects in JSON or YAML format.
--ssh-public-key-file string       Path to the authorized SSH key
for the user

```

```

--ssh-username string           Name of the user to create on the
instance (default "konvoy")
--subnet-ids strings           A comma separated list of
existing subnet IDs to use for the kube-apiserver ELB and all control-plane and
worker nodes (default [])
--template string              Template string or path to
template file to use when -o=go-template, -o=go-template-file. The template format is
golang templates [http://golang.org/pkg/text/template/#pkg-overview].
--timeout duration             The length of time to wait before
giving up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
--vpc-id string                Existing VPC ID to use for the
cluster
--wait                          If true, wait for operations to
complete before returning. This flag is ignored and will always be 'true' if used
with the --self-managed flag. (default true)
--with-aws-bootstrap-credentials Set true to use AWS bootstrap
credentials from your environment. When false, the instance profile of the EC2
instance where the CAPA controller is scheduled on will be used instead.
--with-gcp-bootstrap-credentials Set true to use GCP bootstrap
credentials from your environment. When false, the service account of the VM instance
where the CAPG controller is scheduled on will be used instead.
--worker-availability-zone string The AvailabilityZone in the
region to deploy the worker nodes to, if not set a random one will be selected (ex.
us-west-2a)
--worker-http-proxy string      HTTP proxy for nodes
--worker-https-proxy string     HTTPS proxy for nodes
--worker-iam-instance-profile string Name of the IAM instance profile
to assign to worker machines. (default "nodes.cluster-api-provider-aws.sigs.k8s.io")
--worker-instance-type string   Worker machine instance type
(default "m5.2xlarge")
--worker-no-proxy strings       No Proxy list for nodes (default
[])
--worker-replicas int           Number of workers (default 4)

```

11.2.7.9.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.7.9.4.3 SEE ALSO

- [dkp create cluster](#) (see page 1665) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.9.5 dkp create cluster azure

Create a Konvoy cluster in Azure

```
dkp create cluster azure [flags]
```

11.2.7.9.5.1 Options

```

    --additional-tags stringToString      Tags to apply to the provisioned
    infrastructure (default [])
    --allow-missing-template-keys         If true, ignore any errors in templates
    when a field or map key is missing in the template. Only applies to goolang and
    jsonpath output formats. (default true)
    --aws-service-endpoints string       Custom AWS service endpoints in a semi-
    colon separated format: ${SigningRegion1}:${ServiceID1}=${URL},${ServiceID2}=${URL};${
    {SigningRegion2}}...
    --certificate-renew-interval int    The interval number of days Kubernetes
    managed PKI certificates are renewed. For example, an Interval value of 30 means the
    certificates will be refreshed every 30 days. A value of 0 disables the feature. (def
    ault 0)
    -c, --cluster-name name              Name used to prefix the cluster and all
    the created resources.
    --compute-gallery-id string          Compute Gallery ID of a custom image,
    e.g., '/subscriptions/<subscription id>/resourceGroups/<resource group name>/
    providers/Microsoft.Compute/galleries/<gallery name>/images/<image definition name>/
    versions/<version id>' (replacing placeholders with the values used when creating the
    image)
    --control-plane-http-proxy string    HTTP proxy for control plane machines
    --control-plane-https-proxy string   HTTPS proxy for control plane machines
    --control-plane-machine-size string  Control Plane machine size (default
    "Standard_D4s_v3")
    --control-plane-no-proxy strings     No Proxy list for control plane machines
    (default [])
    --control-plane-replicas int       Number of control plane nodes (default 3)
    --dry-run                             Only print the objects that would be
    created, without creating them.
    --etcd-image-repository string       The image repository to use for pulling
    the etcd image
    --etcd-version string                The version of etcd to use.
    --extra-sans strings                 A comma separated list of additional
    Subject Alternative Names for the API Server signing cert (default [])
    -h, --help                           help for azure
    --http-proxy string                  HTTP proxy for CAPI controllers
    --https-proxy string                 HTTPS proxy for CAPI controllers
    --kind-cluster-image string          Kind node image for the bootstrap
    cluster (default "mesosphere/konvoy-bootstrap:v2.6.2-rc.4")
    --kubeconfig string                  Path to the kubeconfig for the
    management cluster. If unspecified, default discovery rules apply. This flag is
    ignored if used with the --self-managed flag.
    --kubernetes-image-repository string The image repository to use for pulling
    kubernetes images

```

```

--kubernetes-pod-network-cidr cidr      The Kubernetes Pod network CIDR to use
in the cluster (default 192.168.0.0/16)
--kubernetes-service-cidr cidr         The Kubernetes Service CIDR to use in
the cluster (default 10.96.0.0/12)
--kubernetes-version string            Kubernetes version (default "1.26.14")
--location string                       Azure location to deploy cluster to
(default "westus")
-n, --namespace string                 If present, the namespace scope for this
CLI request. (default "default")
--no-proxy strings                      No Proxy list for CAPI controllers (default [])
-o, --output string                    Output format. One of: (json, yaml,
name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-
json, jsonpath-file).
--output-directory string               Used with --output=json|yaml. The
directory where to output resources to files. The directory must already exist.
--plan-offer string                     The offer for a Marketplace image or a
custom image sourced from a Marketplace image requiring Plan information.
--plan-publisher string                 The publisher for a Marketplace image or
a custom image sourced from a Marketplace image requiring Plan information.
--plan-sku string                       The SKU for a Marketplace image or a
custom image sourced from a Marketplace image requiring Plan information.
--registry-mirror-cacert file           CA certificate chain to use while
communicating with the registry mirror using TLS
--registry-mirror-password string       Password to authenticate to the registry
mirror with
--registry-mirror-url url               URL of a container registry to use as a
mirror in the cluster
--registry-mirror-username string       Username to authenticate to the registry
mirror with
--self-managed                          When set to true, the required
prerequisites are created before creating the cluster and the resulting cluster has
all necessary components deployed onto itself, so it can manage its own cluster
lifecycle. When set to false, a management cluster is used. (default false)
--show-managed-fields                   If true, keep the managedFields when
printing objects in JSON or YAML format.
--ssh-public-key-file string            Path to the authorized SSH key for the
user
--ssh-username string                   Name of the user to create on the
instance (default "konvoy")
--template string                       Template string or path to template file
to use when -o=go-template, -o=go-template-file. The template format is go lang
templates [http://golang.org/pkg/text/template/#pkg-overview].
--timeout duration                      The length of time to wait before giving
up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
--wait                                  If true, wait for operations to complete
before returning. This flag is ignored and will always be 'true' if used with the --
self-managed flag. (default true)
--with-aws-bootstrap-credentials        Set true to use AWS bootstrap
credentials from your environment. When false, the instance profile of the EC2
instance where the CAPA controller is scheduled on will be used instead.

```



```

--with-gcp-bootstrap-credentials      Set true to use GCP bootstrap
credentials from your environment. When false, the service account of the VM instance
where the CAPG controller is scheduled on will be used instead.
--worker-availability-zone string      The availability zone in the region to
deploy the worker nodes to, if not set a random one will be selected (ex. 1). Not all
locations, including the default 'westus', support setting this flag, see https://
docs.microsoft.com/en-us/azure/availability-zones/az-overview.
--worker-http-proxy string             HTTP proxy for nodes
--worker-https-proxy string           HTTPS proxy for nodes
--worker-machine-size string          Worker machine size (default
"Standard_D8s_v3")
--worker-no-proxy strings             No Proxy list for nodes (default [])
--worker-replicas int               Number of workers (default 4)

```

11.2.7.9.5.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.7.9.5.3 SEE ALSO

- [dkp create cluster](#) (see page 1665) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.9.6 dkp create cluster vsphere

Create a Konvoy cluster in vSphere

```
dkp create cluster vsphere [flags]
```

11.2.7.9.6.1 Options

```

--allow-missing-template-keys          If true, ignore any errors in templates
when a field or map key is missing in the template. Only applies to golang and
jsonpath output formats. (default true)
--aws-service-endpoints string        Custom AWS service endpoints in a semi-
colon separated format: ${SigningRegion1}:${ServiceID1}=${URL},${ServiceID2}=${URL};$
{SigningRegion2}...
--certificate-renew-interval int    The interval number of days Kubernetes
managed PKI certificates are renewed. For example, an Interval value of 30 means the
certificates will be refreshed every 30 days. A value of 0 disables the feature. (def
ault 0)
-c, --cluster-name name              Name used to prefix the cluster and all
the created resources.

```

```

--control-plane-cpus int           The number of virtual processors in a
control plane machine (default 4)
--control-plane-disk-size int       The size of a control plane machine's
disk, in GB (default 80)
--control-plane-endpoint-host string  The control plane endpoint address. To
use an external load balancer, set to its IP or hostname. To use the built-in virtual
IP, set to a static IPv4 address in the Layer 2 network of the control plane
machines. [Not for production use: To use a single-machine control plane, set to the
IP or hostname of the machine.]
--control-plane-endpoint-port int   The control plane endpoint port. To use
an external load balancer, set to its listening port. (default 6443)
--control-plane-http-proxy string     HTTP proxy for control plane machines
--control-plane-https-proxy string   HTTPS proxy for control plane machines
--control-plane-memory int         The size of a control plane machine's
memory, in GB (default 16)
--control-plane-no-proxy strings     No Proxy list for control plane machines
(default [])
--control-plane-replicas int       Number of control plane nodes (default 3)
--data-center string                 The vSphere datacenter to deploy the
workload cluster on.
--data-store string                  The vSphere datastore to deploy the
workload cluster on.
--dry-run                             Only print the objects that would be
created, without creating them.
--etcd-image-repository string        The image repository to use for pulling
the etcd image
--etcd-version string                 The version of etcd to use.
--extra-sans strings                  A comma separated list of additional
Subject Alternative Names for the API Server signing cert (default [])
--folder string                       The vSphere folder for your VMs. Set to
"" to use the root vSphere folder.
-h, --help                             help for vsphere
--http-proxy string                   HTTP proxy for CAPI controllers
--https-proxy string                  HTTPS proxy for CAPI controllers
--kind-cluster-image string           Kind node image for the bootstrap
cluster (default "mesosphere/konvoy-bootstrap:v2.6.2-rc.4")
--kubeconfig string                   Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply. This flag is
ignored if used with the --self-managed flag.
--kubernetes-image-repository string  The image repository to use for pulling
kubernetes images
--kubernetes-pod-network-cidr cidr    The Kubernetes Pod network CIDR to use
in the cluster (default 192.168.0.0/16)
--kubernetes-service-cidr cidr       The Kubernetes Service CIDR to use in
the cluster (default 10.96.0.0/12)
--kubernetes-version string           Kubernetes version (default "1.26.14")
-n, --namespace string                If present, the namespace scope for this
CLI request. (default "default")
--network string                       The vSphere network to deploy the
workload cluster on.
--no-proxy strings                     No Proxy list for CAPI controllers (defa
ult [])

```

```

-o, --output string                Output format. One of: (json, yaml,
name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-
json, jsonpath-file).
--output-directory string          Used with --output=json|yaml. The
directory where to output resources to files. The directory must already exist.
--registry-mirror-cacert file      CA certificate chain to use while
communicating with the registry mirror using TLS
--registry-mirror-password string  Password to authenticate to the registry
mirror with
--registry-mirror-url url         URL of a container registry to use as a
mirror in the cluster
--registry-mirror-username string  Username to authenticate to the registry
mirror with
--resource-pool string            The vSphere resource pool for the
workload cluster's virtual machines.
--self-managed                    When set to true, the required
prerequisites are created before creating the cluster and the resulting cluster has
all necessary components deployed onto itself, so it can manage its own cluster
lifecycle. When set to false, a management cluster is used. (default false)
--server string                  The vCenter server address. Accepted
formats: host, host:port, http[s]://host[:port]. Accepted host formats: IPv4, IPv6,
or DNS name.
--show-managed-fields            If true, keep the managedFields when
printing objects in JSON or YAML format.
--ssh-public-key-file string     Path to the authorized SSH key for the
user
--ssh-username string            Name of the user to create on the
instance (default "konvoy")
--storage-policy string          This is the vSphere storage policy. Set
it to "" if you don't want to use a storage policy.
--template string                Template string or path to template file
to use when -o=go-template, -o=go-template-file. The template format is golang
templates [http://golang.org/pkg/text/template/#pkg-overview].
--timeout duration               The length of time to wait before giving
up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
--tls-thumb-print string         sha1 thumbprint of the vcenter
certificate: openssl x509 -sha1 -fingerprint -in ca.crt -noout
--virtual-ip-interface string   The network interface, e.g. 'eth0' or
'ens5', to use for the built-in virtual IP control plane endpoint. This interface
must be available on every control plane machine. If the value is empty, the flag
does nothing. If the value is not empty, the built-in virtual IP control plane
endpoint is created, using values from --control-plane-endpoint-host and --control-
plane-endpoint-port.
--vm-template string             The virtual machine template to use for
the workload cluster's virtual machines.
--wait                           If true, wait for operations to complete
before returning. This flag is ignored and will always be 'true' if used with the --
self-managed flag. (default true)
--with-aws-bootstrap-credentials Set true to use AWS bootstrap
credentials from your environment. When false, the instance profile of the EC2
instance where the CAPA controller is scheduled on will be used instead.

```

<code>--with-gcp-bootstrap-credentials</code>	Set true to use GCP bootstrap credentials from your environment. When false , the service account of the VM instance where the CAPG controller is scheduled on will be used instead.
<code>--worker-cpus</code> int	The number of virtual processors in a worker machine (default 8)
<code>--worker-disk-size</code> int	The size of a worker machine's disk, in GB (default 80)
<code>--worker-http-proxy</code> string	HTTP proxy for nodes
<code>--worker-https-proxy</code> string	HTTPS proxy for nodes
<code>--worker-memory</code> int	The size of a worker machine's memory, in GB (default 32)
<code>--worker-no-proxy</code> strings	No Proxy list for nodes (default [])
<code>--worker-replicas</code> int	Number of workers (default 4)

11.2.7.9.6.2 Options inherited from parent commands

`-v, --verbose` **int** Output verbosity

11.2.7.9.6.3 SEE ALSO

- [dkp create cluster](#) (see page 1665) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.9.7 dkp create cluster aks

Create a Konvoy cluster in AKS

```
dkp create cluster aks [flags]
```

11.2.7.9.7.1 Options

<code>--additional-tags</code> stringToString	Tags to apply to the provisioned infrastructure (default [])
<code>--allow-missing-template-keys</code>	If true , ignore any errors in templates when a field or map key is missing in the template. Only applies to goLang and jsonPath output formats. (default true)
<code>--aws-service-endpoints</code> string	Custom AWS service endpoints in a semi-colon separated format: <code>\${SigningRegion1}:\${ServiceID1}=\${URL},\${ServiceID2}=\${URL};\${SigningRegion2}...</code>
<code>--certificate-renew-interval</code> int	The interval number of days Kubernetes managed PKI certificates are renewed. For example, an Interval value of 30 means the certificates will be refreshed every 30 days. A value of 0 disables the feature. (default 0)

```

-c, --cluster-name name           Name used to prefix the cluster and all
the created resources.
  --dry-run                       Only print the objects that would be
created, without creating them.
-h, --help                       help for aks
  --http-proxy string             HTTP proxy for CAPI controllers
  --https-proxy string            HTTPS proxy for CAPI controllers
  --kind-cluster-image string     Kind node image for the bootstrap cluster
(default "mesosphere/konvoy-bootstrap:v2.6.2-rc.4")
  --kubeconfig string             Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
  --kubernetes-pod-network-cidr cidr The Kubernetes Pod network CIDR to use in
the cluster (default 192.168.0.0/16)
  --kubernetes-service-cidr cidr  The Kubernetes Service CIDR to use in the
cluster (default 10.96.0.0/12)
  --kubernetes-version string     Kubernetes version. Run 'az aks get-
versions -o table --location <location>' to see available versions. See https://
docs.microsoft.com/en-us/azure/aks/supported-kubernetes-versions for more details.
Must be a patch version for v1.26.x.
  --location string               Azure location to deploy cluster to
(default "westus")
-n, --namespace string           If present, the namespace scope for this
CLI request. (default "default")
  --no-proxy strings              No Proxy list for CAPI controllers (default
t [])
-o, --output string              Output format. One of: (json, yaml, name,
go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
  --output-directory string        Used with --output=json|yaml. The
directory where to output resources to files. The directory must already exist.
  --show-managed-fields            If true, keep the managedFields when
printing objects in JSON or YAML format.
  --ssh-public-key-file string     Path to the authorized SSH key for the
user
  --ssh-username string           Name of the user to create on the instance
(default "konvoy")
  --system-machine-size string     System node pool machine size (default
"Standard_D4s_v3")
  --system-replicas int         Number of system nodes (default 3)
  --template string               Template string or path to template file
to use when -o=go-template, -o=go-template-file. The template format is golang
templates [http://golang.org/pkg/text/template/#pkg-overview].
  --timeout duration              The length of time to wait before giving
up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
  --wait                          If true, wait for operations to complete
before returning. (default true)
  --with-aws-bootstrap-credentials Set true to use AWS bootstrap credentials
from your environment. When false, the instance profile of the EC2 instance where the
CAPA controller is scheduled on will be used instead.
  --with-gcp-bootstrap-credentials Set true to use GCP bootstrap credentials
from your environment. When false, the service account of the VM instance where the
CAPG controller is scheduled on will be used instead.

```

<code>--worker-http-proxy</code> string	HTTP proxy for nodes
<code>--worker-https-proxy</code> string	HTTPS proxy for nodes
<code>--worker-machine-size</code> string	Worker machine size (default
<code>"Standard_D8s_v3"</code>)	
<code>--worker-no-proxy</code> strings	No Proxy list for nodes (default [])
<code>--worker-replicas</code> int	Number of workers (default 4)

11.2.7.9.7.2 Options inherited from parent commands

`-v, --verbose` **int** Output verbosity

11.2.7.9.7.3 SEE ALSO

- [dkp create cluster](#) (see page 1665) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.9.8 dkp create cluster preprovisioned

Create a Konvoy cluster on pre-provisioned infrastructure

```
dkp create cluster preprovisioned [flags]
```

11.2.7.9.8.1 Options

<code>--allow-missing-template-keys</code>	If true , ignore any errors in templates when a field or map key is missing in the template. Only applies to goyaml and jsonpath output formats. (default true)
<code>--aws-service-endpoints</code> string	Custom AWS service endpoints in a semi-colon separated format: <code>\${SigningRegion1}:\${ServiceID1}=\${URL},\${ServiceID2}=\${URL};\${SigningRegion2}...</code>
<code>--certificate-renew-interval</code> int	The interval number of days Kubernetes managed PKI certificates are renewed. For example, an Interval value of 30 means the certificates will be refreshed every 30 days. A value of 0 disables the feature. (default 0)
<code>-c, --cluster-name</code> name	Name used to prefix the cluster and all the created resources.
<code>--control-plane-endpoint-host</code> string	The control plane endpoint address. To use an external load balancer, set to its IP or hostname. To use the built-in virtual IP, set to a static IPv4 address in the Layer 2 network of the control plane machines. [Not for production use: To use a single-machine control plane, set to the IP or hostname of the machine.]
<code>--control-plane-endpoint-port</code> int	The control plane endpoint port. To use an external load balancer, set to its listening port. (default 6443)

```

--control-plane-http-proxy string      HTTP proxy for control plane machines
--control-plane-https-proxy string     HTTPS proxy for control plane
machines
--control-plane-no-proxy strings       No Proxy list for control plane
machines (default [])
--control-plane-replicas int         Number of control plane nodes (default
t 3)
--dry-run                               Only print the objects that would be
created, without creating them.
--etcd-image-repository string         The image repository to use for
pulling the etcd image
--etcd-version string                 The version of etcd to use.
--extra-sans strings                   A comma separated list of additional
Subject Alternative Names for the API Server signing cert (default [])
-h, --help                             help for preprovisioned
--http-proxy string                   HTTP proxy for CAPI controllers
--https-proxy string                  HTTPS proxy for CAPI controllers
--kind-cluster-image string            Kind node image for the bootstrap
cluster (default "mesosphere/konvoy-bootstrap:v2.6.2-rc.4")
--kubeconfig string                   Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply. This flag is
ignored if used with the --self-managed flag.
--kubernetes-image-repository string   The image repository to use for
pulling kubernetes images
--kubernetes-pod-network-cidr cidr     The Kubernetes Pod network CIDR to
use in the cluster (default 192.168.0.0/16)
--kubernetes-service-cidr cidr        The Kubernetes Service CIDR to use in
the cluster (default 10.96.0.0/12)
--kubernetes-version string            Kubernetes version (default "1.26.14")
-n, --namespace string                 If present, the namespace scope for
this CLI request. (default "default")
--no-proxy strings                     No Proxy list for CAPI controllers (d
efault [])
--os-hint flatcar                       A hint which will allow the installer
to generate appropriate configurations for a target OS. Presently, only the hint for
flatcar is supported.
-o, --output string                     Output format. One of: (json, yaml,
name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-
json, jsonpath-file).
--output-directory string               Used with --output=json|yaml. The
directory where to output resources to files. The directory must already exist.
--override-secret-name string           Name of the secret for any provided
overrides on a preprovisioned cluster. All overrides defined at provisioning should
be present in this secret.
--pre-provisioned-inventory-file string Path to PreprovisionedInventory
inventory file
--registry-mirror-cacert file           CA certificate chain to use while
communicating with the registry mirror using TLS
--registry-mirror-password string       Password to authenticate to the
registry mirror with
--registry-mirror-url url               URL of a container registry to use as
a mirror in the cluster

```

<code>--registry-mirror-username</code> string	Username to authenticate to the registry mirror with
<code>--self-managed</code>	When set to true , the required prerequisites are created before creating the cluster and the resulting cluster has all necessary components deployed onto itself, so it can manage its own cluster lifecycle. When set to false , a management cluster is used. (default false)
<code>--show-managed-fields</code>	If true , keep the managedFields when printing objects in JSON or YAML format.
<code>--ssh-private-key-file</code> string	Path to the private SSH key file used by Konvoy to access the pre-provisioned hosts
<code>--ssh-public-key-file</code> string	Path to the authorized SSH key for the user
<code>--ssh-username</code> string	Name of the user to create on the instance (default "konvoy")
<code>--template</code> string	Template string or path to template file to use when <code>-o=go-template</code> , <code>-o=go-template-file</code> . The template format is go lang templates [http://golang.org/pkg/text/template/#pkg-overview].
<code>--timeout</code> duration	The length of time to wait before giving up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
<code>--virtual-ip-interface</code> string	The network interface , e.g, 'eth0' or 'ens5', to use for the built-in virtual IP control plane endpoint. This interface must be available on every control plane machine. If the value is empty, the flag does nothing. If the value is not empty, the built-in virtual IP control plane endpoint is created, using values from <code>--control-plane-endpoint-host</code> and <code>--control-plane-endpoint-port</code> .
<code>--wait</code>	If true , wait for operations to complete before returning. This flag is ignored and will always be 'true' if used with the <code>--self-managed</code> flag. (default true)
<code>--with-aws-bootstrap-credentials</code>	Set true to use AWS bootstrap credentials from your environment. When false , the instance profile of the EC2 instance where the CAPA controller is scheduled on will be used instead.
<code>--with-gcp-bootstrap-credentials</code>	Set true to use GCP bootstrap credentials from your environment. When false , the service account of the VM instance where the CAPG controller is scheduled on will be used instead.
<code>--worker-http-proxy</code> string	HTTP proxy for nodes
<code>--worker-https-proxy</code> string	HTTPS proxy for nodes
<code>--worker-no-proxy</code> strings	No Proxy list for nodes (default [])
<code>--worker-replicas</code> int	Number of workers (default 4)

11.2.7.9.8.2 Options inherited from parent commands

`-v, --verbose` **int** Output verbosity

11.2.7.9.8.3 SEE ALSO

- [dkp create cluster](#) (see page 1665) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.9.9 dkp create cluster gcp

Create a Konvoy cluster in GCP

```
dkp create cluster gcp [flags]
```

11.2.7.9.9.1 Options

<code>--additional-tags</code> stringToString infrastructure (default [])	Tags to apply to the provisioned infrastructure
<code>--allow-missing-template-keys</code> templates when a field or map key is missing in the template. Only applies to goyaml and jsonpath output formats. (default true)	If true , ignore any errors in templates
<code>--associate-public-ip-address</code> machines. When set to false the specified network must have Cloud NAT configured to provide internet access. (default true)	Associate a public IP for all machines
<code>--aws-service-endpoints</code> string a semi-colon separated format: <code>\${SigningRegion1}:\${ServiceID1}=\${URL},\${ServiceID2}=\${URL};\${SigningRegion2}...</code>	Custom AWS service endpoints in a semi-colon separated format
<code>--certificate-renew-interval</code> int Kubernetes managed PKI certificates are renewed. For example, an Interval value of 30 means the certificates will be refreshed every 30 days. A value of 0 disables the feature. (default 0)	The interval number of days Kubernetes managed PKI certificates are renewed
<code>-c, --cluster-name</code> name and all the created resources.	Name used to prefix the cluster and all the created resources
<code>--control-plane-http-proxy</code> string machines	HTTP proxy for control plane machines
<code>--control-plane-https-proxy</code> string machines	HTTPS proxy for control plane machines
<code>--control-plane-instance-type</code> string type (default "n2-standard-4")	Control Plane machine instance type
<code>--control-plane-no-proxy</code> strings machines (default [])	No Proxy list for control plane machines
<code>--control-plane-replicas</code> int default 3)	Number of control plane nodes (default 3)
<code>--control-plane-service-account-email</code> string email address (default "default")	Control Plane Service Account email address
<code>--dry-run</code> would be created, without creating them.	Only print the objects that would be created, without creating them
<code>--etcd-image-repository</code> string pulling the etcd image	The image repository to use for pulling the etcd image
<code>--etcd-version</code> string	The version of etcd to use
<code>--extra-sans</code> strings additional Subject Alternative Names for the API Server signing cert (default [])	A comma separated list of additional Subject Alternative Names for the API Server signing cert
<code>-h, --help</code>	help for gcp
<code>--http-proxy</code> string	HTTP proxy for CAPI controllers
<code>--https-proxy</code> string	HTTPS proxy for CAPI controllers

<p>--image string</p> <p>use for all nodes (set either this or --image-family) (ex. 'projects/my-project/global/images/konvoy-ubuntu-2004-1-99-99-1234567890')</p>	<p>Full reference to an image to</p>
<p>--image-family string</p> <p>family to use for all nodes (set either this or --image) (ex. 'projects/my-project/global/images/family/konvoy-ubuntu-2004-{{.K8sVersion}}')</p>	<p>Full reference to an image</p>
<p>--kind-cluster-image string</p> <p>bootstrap cluster (default "mesosphere/konvoy-bootstrap:v2.6.2-rc.4")</p>	<p>Kind node image for the</p>
<p>--kubeconfig string</p> <p>management cluster. If unspecified, default discovery rules apply. This flag is ignored if used with the --self-managed flag.</p>	<p>Path to the kubeconfig for the</p>
<p>--kubernetes-image-repository string</p> <p>pulling kubernetes images</p>	<p>The image repository to use for</p>
<p>--kubernetes-pod-network-cidr cidr</p> <p>to use in the cluster (default 192.168.0.0/16)</p>	<p>The Kubernetes Pod network CIDR</p>
<p>--kubernetes-service-cidr cidr</p> <p>use in the cluster (default 10.96.0.0/12)</p>	<p>The Kubernetes Service CIDR to</p>
<p>--kubernetes-version string</p> <p>"1.26.14")</p>	<p>Kubernetes version (default</p>
<p>-n, --namespace string</p> <p>for this CLI request. (default "default")</p>	<p>If present, the namespace scope</p>
<p>--network string</p> <p>the cluster to (default "default")</p>	<p>The GCP network name to deploy</p>
<p>--no-proxy strings</p> <p>controllers (default [])</p>	<p>No Proxy list for CAPI</p>
<p>-o, --output string</p> <p>yaml, name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json, jsonpath-file).</p>	<p>Output format. One of: (json,</p>
<p>--output-directory string</p> <p>The directory where to output resources to files. The directory must already exist.</p>	<p>Used with --output=json yaml.</p>
<p>--project string</p> <p>the cluster to</p>	<p>The GCP project name to deploy</p>
<p>--region string</p> <p>(default "us-west1")</p>	<p>GCP region to deploy cluster to</p>
<p>--registry-mirror-cacert file</p> <p>while communicating with the registry mirror using TLS</p>	<p>CA certificate chain to use</p>
<p>--registry-mirror-password string</p> <p>registry mirror with</p>	<p>Password to authenticate to the</p>
<p>--registry-mirror-url url</p> <p>use as a mirror in the cluster</p>	<p>URL of a container registry to</p>
<p>--registry-mirror-username string</p> <p>registry mirror with</p>	<p>Username to authenticate to the</p>
<p>--self-managed</p> <p>prerequisites are created before creating the cluster and the resulting cluster has all necessary components deployed onto itself, so it can manage its own cluster lifecycle. When set to false, a management cluster is used. (default false)</p>	<p>When set to true, the required</p>
<p>--show-managed-fields</p> <p>when printing objects in JSON or YAML format.</p>	<p>If true, keep the managedFields</p>
<p>--ssh-public-key-file string</p> <p>for the user</p>	<p>Path to the authorized SSH key</p>
<p>--ssh-username string</p> <p>the instance (default "konvoy")</p>	<p>Name of the user to create on</p>

```

--template string                Template string or path to
template file to use when -o=go-template, -o=go-template-file. The template format is
golang templates [http://golang.org/pkg/text/template/#pkg-overview].
--timeout duration              The length of time to wait
before giving up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
--wait                          If true, wait for operations to
complete before returning. This flag is ignored and will always be 'true' if used
with the --self-managed flag. (default true)
--with-aws-bootstrap-credentials Set true to use AWS bootstrap
credentials from your environment. When false, the instance profile of the EC2
instance where the CAPA controller is scheduled on will be used instead.
--with-gcp-bootstrap-credentials Set true to use GCP bootstrap
credentials from your environment. When false, the service account of the VM instance
where the CAPG controller is scheduled on will be used instead.
--worker-http-proxy string      HTTP proxy for nodes
--worker-https-proxy string     HTTPS proxy for nodes
--worker-instance-type string   Worker machine instance type
(default "n2-standard-8")
--worker-no-proxy strings       No Proxy list for nodes (default
[])
--worker-replicas int           Number of workers (default 4)
--worker-service-account-email string Worker machine Service Account
email address (default "default")
--worker-zone string            Zone in the region to deploy the
worker nodes to, if not set a random one will be selected (ex. us-west1-a)

```

11.2.7.9.9.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.7.9.9.3 SEE ALSO

- [dkp create cluster](#) (see page 1665) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.9.10 dkp create cluster vcd

Create a Konvoy cluster in VMware Cloud Director

```
dkp create cluster vcd [flags]
```

11.2.7.9.10.1 Options

<code>--allow-missing-template-keys</code>	If true , ignore any errors in templates when a field or map key is missing in the template. Only applies to <code>golang</code> and <code>jsonpath</code> output formats. (default true)
<code>--aws-service-endpoints string</code>	Custom AWS service endpoints in a semi-colon separated format: <code>\${SigningRegion1}:\${ServiceID1}=\${URL},\${ServiceID2}=\${URL};\${SigningRegion2}...</code>
<code>--catalog string</code>	Catalog hosting virtual machine templates.
<code>--certificate-renew-interval int</code>	The interval number of days Kubernetes managed PKI certificates are renewed. For example, an Interval value of <code>30</code> means the certificates will be refreshed every <code>30</code> days. A value of <code>0</code> disables the feature. (default 0)
<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>--control-plane-endpoint-host string</code>	The control plane endpoint IP or hostname.
<code>--control-plane-endpoint-port int</code>	The control plane endpoint port. To use an external load balancer, set to its listening port. (default 6443)
<code>--control-plane-http-proxy string</code>	HTTP proxy for control plane machines
<code>--control-plane-https-proxy string</code>	HTTPS proxy for control plane machines
<code>--control-plane-no-proxy strings</code>	No Proxy list for control plane machines (default [])
<code>--control-plane-placement-policy string</code>	The placement policy for control plane to be used on this machine
<code>--control-plane-replicas int</code>	Number of control plane nodes (default 3)
<code>--control-plane-sizing-policy string</code>	The sizing policy for control plane to be used on this machine
<code>--data-center string</code>	The organization's virtual datacenter name where the cluster should be deployed.
<code>--dry-run</code>	Only print the objects that would be created, without creating them.
<code>--etcd-image-repository string</code>	The image repository to use for pulling the etcd image
<code>--etcd-version string</code>	The version of etcd to use.
<code>--extra-sans strings</code>	A comma separated list of additional Subject Alternative Names for the API Server signing cert (default [])
<code>-h, --help</code>	help for vcd
<code>--http-proxy string</code>	HTTP proxy for CAPI controllers
<code>--https-proxy string</code>	HTTPS proxy for CAPI controllers
<code>--kind-cluster-image string</code>	Kind node image for the bootstrap cluster (default "mesosphere/konvoy-bootstrap:v2.6.2-rc.4")
<code>--kubeconfig string</code>	Path to the kubeconfig for the management cluster. If unspecified, default discovery rules apply. This flag is ignored if used with the <code>--self-managed</code> flag.

<pre> --kubernetes-image-repository string pulling kubernetes images --kubernetes-pod-network-cidr cidr use in the cluster (default 192.168.0.0/16) --kubernetes-service-cidr cidr the cluster (default 10.96.0.0/12) --kubernetes-version string --loadbalancer-network-cidr cidr both the control plane nodes and kubernetes external services -n, --namespace string this CLI request. (default "default") --network string network to be used by the cluster. --no-proxy strings default []) --organization string cluster should be deployed. -o, --output string name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as- json, jsonpath-file). --output-directory string directory where to output resources to files. The directory must already exist. --registry-mirror-cacert file communicating with the registry mirror using TLS --registry-mirror-password string registry mirror with --registry-mirror-url url a mirror in the cluster --registry-mirror-username string registry mirror with --self-managed prerequisites are created before creating the cluster and the resulting cluster has all necessary components deployed onto itself, so it can manage its own cluster lifecycle. When set to false, a management cluster is used. (default false) --show-managed-fields printing objects in JSON or YAML format. --site url the format https://VCD_HOST. --ssh-public-key-file string the user --ssh-username string instance (default "konvoy") --storage-profile string virtual machine (default "*") --template string file to use when -o=go-template, -o=go-template-file. The template format is golang templates [http://golang.org/pkg/text/template/#pkg-overview]. --timeout duration giving up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s) --vapp-template string virtual machine. </pre>	<pre> The image repository to use for The Kubernetes Pod network CIDR to The Kubernetes Service CIDR to use in Kubernetes version (default "1.26.14") The load balancer CIDR to use for If present, the namespace scope for The organization's virtual datacenter No Proxy list for CAPI controllers (d The organization name where the Output format. One of: (json, yaml, Used with --output=json yaml. The CA certificate chain to use while Password to authenticate to the URL of a container registry to use as Username to authenticate to the When set to true, the required If true, keep the managedFields when The cloud director's endpoint with Path to the authorized SSH key for Name of the user to create on the The storage profile to be used for a Template string or path to template The length of time to wait before The vApp template used to create the </pre>
--	--

<code>--wait</code>	If true , wait for operations to complete before returning. This flag is ignored and will always be 'true' if used with the <code>--self-managed</code> flag. (default true)
<code>--with-aws-bootstrap-credentials</code>	Set true to use AWS bootstrap credentials from your environment. When false , the instance profile of the EC2 instance where the CAPA controller is scheduled on will be used instead.
<code>--with-gcp-bootstrap-credentials</code>	Set true to use GCP bootstrap credentials from your environment. When false , the service account of the VM instance where the CAPG controller is scheduled on will be used instead.
<code>--worker-http-proxy string</code>	HTTP proxy for nodes
<code>--worker-https-proxy string</code>	HTTPS proxy for nodes
<code>--worker-no-proxy strings</code>	No Proxy list for nodes (default [])
<code>--worker-placement-policy string</code>	The placement policy for workers to be used on this machine
<code>--worker-replicas int</code>	Number of workers (default 4)
<code>--worker-sizing-policy string</code>	The sizing policy for workers to be used on this machine

11.2.7.9.10.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

11.2.7.9.10.3 SEE ALSO

- [dkp create cluster](#) (see page 1665) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.10 dkp create chart-bundle

Create charts bundle based on a catalog applications git repository

```
dkp create chart-bundle [flags]
```

11.2.7.10.1 Options

<code>--catalog-repository string</code>	Git repository containing catalog application definitions
<code>--config string</code>	Config file to use (default "/home/runner/.kommander/config")
<code>--context string</code>	The name of the kubeconfig context to use
<code>--dry-run</code>	Export in YAML format to stdout

```

    --extra-charts strings           Extra charts to include in the bundle, in
the format <repo-url1>|<chart-name1>,<repo-url2>|<chart-name2>:<chart-version2>,...
(default [])
  -h, --help                       help for chart-bundle
    --kommander-charts-version string Kommander helm charts version to download.
Default: download all available versions
    --kubeconfig string            Path to the kubeconfig file to use for CLI
requests.
  -o, --output string              File path to write charts bundle to
(default "charts-bundle.tar.gz")
    --request-timeout string       The length of time to wait before giving up
on a single server request. Non-zero values should contain a corresponding time unit
(e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
    --skip-charts strings          Charts to not to include in the bundle, in
the format <chart-name1>,<chart-name2>:<chart-version2>,... (default [])
  -v, --verbose int              Output verbosity

```

11.2.7.10.2 SEE ALSO

- [dkp create](#) (see page 1660) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

11.2.7.11 dkp create nodepool

Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.11.1 Options

```

-h, --help           help for nodepool
-v, --verbose int  Output verbosity

```

11.2.7.11.2 SEE ALSO

- [dkp create](#) (see page 1660) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]
- [dkp create nodepool aks](#) (see page 1688) - Create a nodepool in AKS
- [dkp create nodepool aws](#) (see page 1689) - Create a nodepool in AWS
- [dkp create nodepool azure](#) (see page 1685) - Create a nodepool in Azure
- [dkp create nodepool eks](#) (see page 1691) - Create a nodepool for EKS
- [dkp create nodepool gcp](#) (see page 1693) - Create a nodepool in GCP
- [dkp create nodepool preprovisioned](#) (see page 1694) - Create a nodepool for a Pre Provisioned Cluster.
- [dkp create nodepool vcd](#) (see page 1696) - Create a nodepool in VCD
- [dkp create nodepool vsphere](#) (see page 1686) - Create a nodepool in vSphere

11.2.7.11.3 dkp create nodepool azure

Create a nodepool in Azure

```
dkp create nodepool azure name [flags]
```

11.2.7.11.3.1 Options

```

    --additional-tags stringToString    Tags to apply to the provisioned
    infrastructure (default [])
    --allow-missing-template-keys      If true, ignore any errors in templates
    when a field or map key is missing in the template. Only applies to golang and
    jsonpath output formats. (default true)
    --availability-zone string         The availability zone in the region to
    deploy the worker nodes to, if not set a random one will be selected (ex. 1). Not all
    locations, including the default 'westus', support setting this flag, see https://docs.microsoft.com/en-us/azure/availability-zones/az-overview.
    -c, --cluster-name name           Name used to prefix the cluster and all the
    created resources.
    --compute-gallery-id string       Compute Gallery ID of a custom image, e.g.,
    '/subscriptions/<subscription id>/resourceGroups/<resource group name>/providers/
    Microsoft.Compute/galleries/<gallery name>/images/<image definition name>/versions/
    <version id>' (replacing placeholders with the values used when creating the image)
    --dry-run                          Only print the objects that would be
    created, without creating them.
    -h, --help                        help for azure
    --http-proxy string               HTTP proxy for nodes
    --https-proxy string              HTTPS proxy for nodes
    --kubeconfig string               Path to the kubeconfig for the management
    cluster. If unspecified, default discovery rules apply.
    --kubernetes-version string       Kubernetes version (default "1.26.14")
    --machine-size string              Worker machine size (default
    "Standard_D8s_v3")
    -n, --namespace string            If present, the namespace scope for this
    CLI request. (default "default")
    --no-proxy strings                No Proxy list for nodes (default [])
    -o, --output string                Output format. One of: (json, yaml, name,
    go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
    jsonpath-file).
    --output-directory string          Used with --output=json|yaml. The directory
    where to output resources to files. The directory must already exist.
    --plan-offer string                The offer for a Marketplace image or a
    custom image sourced from a Marketplace image requiring Plan information.
    --plan-publisher string            The publisher for a Marketplace image or a
    custom image sourced from a Marketplace image requiring Plan information.
    --plan-sku string                  The SKU for a Marketplace image or a custom
    image sourced from a Marketplace image requiring Plan information.

```



```

--registry-mirror-cacert file      CA certificate chain to use while
communicating with the registry mirror using TLS
--registry-mirror-password string Password to authenticate to the registry
mirror with
--registry-mirror-url url         URL of a container registry to use as a
mirror in the cluster
--registry-mirror-username string Username to authenticate to the registry
mirror with
--replicas int                  Number of replicas (default 1)
--show-managed-fields             If true, keep the managedFields when
printing objects in JSON or YAML format.
--ssh-public-key-file string     Path to the authorized SSH key for the user
--ssh-username string            Name of the user to create on the instance
(default "konvoy")
--template string                Template string or path to template file to
use when -o=go-template, -o=go-template-file. The template format is go lang templates
[http://golang.org/pkg/text/template/#pkg-overview].
--timeout duration              The length of time to wait before giving
up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
--wait                           If true, wait for operations to complete
before returning.

```

11.2.7.11.3.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.7.11.3.3 SEE ALSO

- [dkp create nodepool](#) (see page 1687) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.11.4 dkp create nodepool vsphere

Create a nodepool in vSphere

```
dkp create nodepool vsphere name [flags]
```

11.2.7.11.4.1 Options

```
--additional-tags stringToString  Tags to apply to the provisioned
infrastructure (default [])
```

```

--allow-missing-template-keys      If true, ignore any errors in templates
when a field or map key is missing in the template. Only applies to goLang and
jsonpath output formats. (default true)
-c, --cluster-name name           Name used to prefix the cluster and all the
created resources.
--cpus int                       The number of virtual processors in a
worker machine (default 8)
--data-center string              The vSphere datacenter to deploy the
workload cluster on.
--data-store string               The vSphere datastore to deploy the
workload cluster on.
--disk-size int                  The size of a worker machine's disk, in GB
(default 80)
--dry-run                          Only print the objects that would be
created, without creating them.
--folder string                   The vSphere folder for your VMs. Set to ""
to use the root vSphere folder.
-h, --help                         help for vsphere
--http-proxy string               HTTP proxy for nodes
--https-proxy string              HTTPS proxy for nodes
--kubeconfig string              Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--kubernetes-version string       Kubernetes version (default "1.26.14")
--memory int                     The size of a worker machine's memory, in
GB (default 32)
-n, --namespace string            If present, the namespace scope for this
CLI request. (default "default")
--network string                  The vSphere network to deploy the workload
cluster on.
--no-proxy strings                No Proxy list for nodes (default [])
-o, --output string               Output format. One of: (json, yaml, name,
go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-directory string         Used with --output=json|yaml. The directory
where to output resources to files. The directory must already exist.
--registry-mirror-cacert file      CA certificate chain to use while
communicating with the registry mirror using TLS
--registry-mirror-password string Password to authenticate to the registry
mirror with
--registry-mirror-url url          URL of a container registry to use as a
mirror in the cluster
--registry-mirror-username string Username to authenticate to the registry
mirror with
--replicas int                   Number of replicas (default 1)
--resource-pool string            The vSphere resource pool for the workload
cluster's virtual machines.
--server string                   The vCenter server address. Accepted
formats: host, host:port, http[s]://host[:port]. Accepted host formats: IPv4, IPv6,
or DNS name.
--show-managed-fields             If true, keep the managedFields when
printing objects in JSON or YAML format.
--ssh-public-key-file string      Path to the authorized SSH key for the user

```

```

--ssh-username string          Name of the user to create on the instance
(default "konvoy")
--storage-policy string        This is the vSphere storage policy. Set it
to "" if you don't want to use a storage policy.
--template string              Template string or path to template file to
use when -o=go-template, -o=go-template-file. The template format is go lang templates
[http://golang.org/pkg/text/template/#pkg-overview].
--timeout duration             The length of time to wait before giving
up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
--tls-thumb-print string       sha1 thumbprint of the vcenter certificate:
openssl x509 -sha1 -fingerprint -in ca.crt -noout
--vm-template string           The virtual machine template to use for the
workload cluster's virtual machines.
--wait                          If true, wait for operations to complete
before returning.

```

11.2.7.11.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.7.11.4.3 SEE ALSO

- [dkp create nodepool](#) (see page 1687) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.11.5 dkp create nodepool aks

Create a nodepool in AKS

```
dkp create nodepool aks name [flags]
```

11.2.7.11.5.1 Options

```

--additional-tags stringToString  Tags to apply to the provisioned
infrastructure (default [])
--allow-missing-template-keys      If true, ignore any errors in templates when
a field or map key is missing in the template. Only applies to go lang and jsonpath
output formats. (default true)
-c, --cluster-name name            Name used to prefix the cluster and all the
created resources.
--dry-run                          Only print the objects that would be
created, without creating them.
-h, --help                          help for aks

```

```

--http-proxy string          HTTP proxy for nodes
--https-proxy string        HTTPS proxy for nodes
--kubeconfig string         Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--kubernetes-version string  Kubernetes version. Run 'az aks get-versions
-o table --location <location>' to see available versions. See https://
docs.microsoft.com/en-us/azure/aks/supported-kubernetes-versions for more details.
Must be a patch version for v1.26.x.
--machine-size string       Worker machine size (default
"Standard_D8s_v3")
-n, --namespace string      If present, the namespace scope for this CLI
request. (default "default")
--no-proxy strings          No Proxy list for nodes (default [])
-o, --output string         Output format. One of: (json, yaml, name,
go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-directory string   Used with --output=json|yaml. The directory
where to output resources to files. The directory must already exist.
--replicas int            Number of replicas (default 1)
--show-managed-fields      If true, keep the managedFields when
printing objects in JSON or YAML format.
--ssh-public-key-file string Path to the authorized SSH key for the user
--ssh-username string      Name of the user to create on the instance
(default "konvoy")
--template string          Template string or path to template file to
use when -o=go-template, -o=go-template-file. The template format is go lang templates
[http://golang.org/pkg/text/template/#pkg-overview].
--timeout duration         The length of time to wait before giving up.
Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
--wait                     If true, wait for operations to complete
before returning.

```

11.2.7.11.5.2 Options inherited from parent commands

```
-v, --verbose int    Output verbosity
```

11.2.7.11.5.3 SEE ALSO

- [dkp create nodepool](#) (see page 1687) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.11.6 dkp create nodepool aws

Create a nodepool in AWS

```
dkp create nodepool aws [flags]
```

11.2.7.11.6.1 Options

```

--additional-security-group-ids strings  A comma separated list of existing
security group IDs to use for machines in addition to those created automatically (de
fault [])
--additional-tags stringToString          Tags to apply to the provisioned
infrastructure (default [])
--allow-missing-template-keys            If true, ignore any errors in
templates when a field or map key is missing in the template. Only applies to goLang
and jsonpath output formats. (default true)
--ami string                              AMI ID to use for machines
--ami-base-os string                      Base OS used in search of AMIs.
Examples: 'centos-7', 'ubuntu-18.04', 'ubuntu-20.04'
--ami-format string                      Query string used in search of AMIs.
Example: When --ami-base-os='rhel8.4', then the string 'prefix-{{.BaseOS}}-?
{{.K8sVersion}}-*' matches any AMIs with the Name 'prefix-rhel8.4-1.26.14
--ami-owner string                       ID of AWS account used in search of
AMIs
--availability-zone string               The AvailabilityZone in the region to
deploy the worker nodes to, if not set a random one will be selected (ex. us-west-2a)
-c, --cluster-name name                 Name used to prefix the cluster and
all the created resources.
--dry-run                                Only print the objects that would be
created, without creating them.
-h, --help                               help for aws
--http-proxy string                     HTTP proxy for nodes
--https-proxy string                    HTTPS proxy for nodes
--iam-instance-profile string            Name of the IAM instance profile to
assign to worker machines. (default "nodes.cluster-api-provider-aws.sigs.k8s.io")
--instance-type string                  Worker machine instance type (default
"m5.2xlarge")
--kubeconfig string                     Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply.
--kubernetes-version string             Kubernetes version (default "1.26.14")
-n, --namespace string                  If present, the namespace scope for
this CLI request. (default "default")
--no-proxy strings                      No Proxy list for nodes (default [])
-o, --output string                     Output format. One of: (json, yaml,
name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-
json, jsonpath-file).
--output-directory string                Used with --output=json|yaml. The
directory where to output resources to files. The directory must already exist.
--registry-mirror-cacert file           CA certificate chain to use while
communicating with the registry mirror using TLS
--registry-mirror-password string       Password to authenticate to the
registry mirror with
--registry-mirror-url url               URL of a container registry to use as
a mirror in the cluster

```

<code>--registry-mirror-username</code> string	Username to authenticate to the registry mirror with
<code>--replicas</code> int	Number of replicas (default 1)
<code>--show-managed-fields</code>	If true , keep the managedFields when printing objects in JSON or YAML format.
<code>--ssh-public-key-file</code> string	Path to the authorized SSH key for the user
<code>--ssh-username</code> string	Name of the user to create on the instance (default "konvoy")
<code>--template</code> string	Template string or path to template file to use when <code>-o=go-template</code> , <code>-o=go-template-file</code> . The template format is go lang templates [http://golang.org/pkg/text/template/#pkg-overview].
<code>--timeout</code> duration	The length of time to wait before giving up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
<code>--wait</code>	If true , wait for operations to complete before returning.

11.2.7.11.6.2 Options inherited from parent commands

`-v`, `--verbose` **int** Output verbosity

11.2.7.11.6.3 SEE ALSO

- [dkp create nodepool](#) (see page 1687) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.11.7 dkp create nodepool eks

Create a nodepool for EKS

```
dkp create nodepool eks [flags]
```

11.2.7.11.7.1 Options

<code>--additional-security-group-ids</code> strings	A comma separated list of existing security group IDs to use for machines in addition to those created automatically (default [])
<code>--additional-tags</code> stringToString	Tags to apply to the provisioned infrastructure (default [])
<code>--allow-missing-template-keys</code>	If true , ignore any errors in templates when a field or map key is missing in the template. Only applies to go lang and jsonpath output formats. (default true)

```

    --availability-zone string           The AvailabilityZone in the region to
deploy the worker nodes to, if not set a random one will be selected (ex. us-west-2a)
    -c, --cluster-name name             Name used to prefix the cluster and
all the created resources.
    --dry-run                            Only print the objects that would be
created, without creating them.
    -h, --help                           help for eks
    --http-proxy string                  HTTP proxy for nodes
    --https-proxy string                 HTTPS proxy for nodes
    --iam-instance-profile string        Name of the IAM instance profile to
assign to worker machines. (default "nodes.cluster-api-provider-aws.sigs.k8s.io")
    --instance-type string               Worker machine instance type (default
"m5.2xlarge")
    --kubeconfig string                  Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply.
    --kubernetes-version string          Kubernetes version (default "1.25.11")
    -n, --namespace string               If present, the namespace scope for
this CLI request. (default "default")
    --no-proxy strings                   No Proxy list for nodes (default [])
    -o, --output string                   Output format. One of: (json, yaml,
name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-
json, jsonpath-file).
    --output-directory string            Used with --output=json|yaml. The
directory where to output resources to files. The directory must already exist.
    --replicas int                       Number of replicas (default 1)
    --show-managed-fields                 If true, keep the managedFields when
printing objects in JSON or YAML format.
    --ssh-public-key-file string          Path to the authorized SSH key for
the user
    --ssh-username string                 Name of the user to create on the
instance (default "konvoy")
    --template string                     Template string or path to template
file to use when -o=go-template, -o=go-template-file. The template format is golang
templates [http://golang.org/pkg/text/template/#pkg-overview].
    --timeout duration                    The length of time to wait before
giving up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
    --wait                                If true, wait for operations to
complete before returning.

```

11.2.7.11.7.2 Options inherited from parent commands

```
-v, --verbose int    Output verbosity
```

11.2.7.11.7.3 SEE ALSO

- [dkp create nodepool](#) (see page 1687) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.11.8 dkp create nodepool gcp

Create a nodepool in GCP

```
dkp create nodepool gcp name [flags]
```

11.2.7.11.8.1 Options

```

--additional-tags stringToString  Tags to apply to the provisioned
infrastructure (default [])
--allow-missing-template-keys     If true, ignore any errors in templates
when a field or map key is missing in the template. Only applies to golang and
jsonpath output formats. (default true)
--associate-public-ip-address     Associate a public IP for all machines.
When set to false the specified network must have Cloud NAT configured to provide
internet access. (default true)
-c, --cluster-name name          Name used to prefix the cluster and all the
created resources.
--dry-run                        Only print the objects that would be
created, without creating them.
-h, --help                       help for gcp
--http-proxy string              HTTP proxy for nodes
--https-proxy string             HTTPS proxy for nodes
--image string                   Full reference to an image to use for all
nodes (set either this or --image-family) (ex. 'projects/my-project/global/images/
konvoy-ubuntu-2004-1-99-99-1234567890')
--image-family string            Full reference to an image family to use
for all nodes (set either this or --image) (ex. 'projects/my-project/global/images/
family/konvoy-ubuntu-2004-{{.K8sVersion}}')
--instance-type string           Worker machine instance type (default "n2-
standard-8")
--kubeconfig string              Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--kubernetes-version string      Kubernetes version (default "1.26.14")
-n, --namespace string           If present, the namespace scope for this
CLI request. (default "default")
--no-proxy strings               No Proxy list for nodes (default [])
-o, --output string              Output format. One of: (json, yaml, name,
go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-directory string        Used with --output=json|yaml. The directory
where to output resources to files. The directory must already exist.
--registry-mirror-cacert file     CA certificate chain to use while
communicating with the registry mirror using TLS
--registry-mirror-password string Password to authenticate to the registry
mirror with

```



```

--registry-mirror-url url           URL of a container registry to use as a
mirror in the cluster
--registry-mirror-username string   Username to authenticate to the registry
mirror with
--replicas int                     Number of replicas (default 1)
--service-account-email string     Worker machine Service Account email
address (default "default")
--show-managed-fields              If true, keep the managedFields when
printing objects in JSON or YAML format.
--ssh-public-key-file string       Path to the authorized SSH key for the user
--ssh-username string              Name of the user to create on the instance
(default "konvoy")
--template string                  Template string or path to template file to
use when -o=go-template, -o=go-template-file. The template format is go lang templates
[http://golang.org/pkg/text/template/#pkg-overview].
--timeout duration                 The length of time to wait before giving
up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
--wait                              If true, wait for operations to complete
before returning.
--zone string                       Zone in the region to deploy the worker
nodes to, if not set a random one will be selected (ex.us-west1-a)

```

11.2.7.11.8.2 Options inherited from parent commands

```
-v, --verbose int    Output verbosity
```

11.2.7.11.8.3 SEE ALSO

- [dkp create nodepool](#) (see page 1687) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.11.9 dkp create nodepool preprovisioned

Create a nodepool for a Pre Provisioned Cluster.

```
dkp create nodepool preprovisioned name [flags]
```

11.2.7.11.9.1 Options

```

--additional-tags stringToString   Tags to apply to the provisioned
infrastructure (default [])

```

```

--allow-missing-template-keys      If true, ignore any errors in templates
when a field or map key is missing in the template. Only applies to goLang and
jsonpath output formats. (default true)
-c, --cluster-name name           Name used to prefix the cluster and all the
created resources.
--dry-run                          Only print the objects that would be
created, without creating them.
-h, --help                        help for preprovisioned
--http-proxy string              HTTP proxy for nodes
--https-proxy string             HTTPS proxy for nodes
--kubeconfig string             Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--kubernetes-version string      Kubernetes version (default "1.26.14")
-n, --namespace string          If present, the namespace scope for this
CLI request. (default "default")
--no-proxy strings              No Proxy list for nodes (default [])
--os-hint flatcar               A hint which will allow the installer to
generate appropriate configurations for a target OS. Presently, only the hint for
flatcar is supported.
-o, --output string             Output format. One of: (json, yaml, name,
go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-directory string       Used with --output=json|yaml. The directory
where to output resources to files. The directory must already exist.
--override-secret-name string   Name of the secret for any provided
overrides on a preprovisioned cluster. All overrides defined at provisioning should
be present in this secret.
--registry-mirror-cacert file   CA certificate chain to use while
communicating with the registry mirror using TLS
--registry-mirror-password string Password to authenticate to the registry
mirror with
--registry-mirror-url url      URL of a container registry to use as a
mirror in the cluster
--registry-mirror-username string Username to authenticate to the registry
mirror with
--replicas int                Number of replicas (default 1)
--show-managed-fields          If true, keep the managedFields when
printing objects in JSON or YAML format.
--ssh-public-key-file string   Path to the authorized SSH key for the user
--ssh-username string          Name of the user to create on the instance
(default "konvoy")
--template string             Template string or path to template file to
use when -o=go-template, -o=go-template-file. The template format is goLang templates
[http://goLang.org/pkg/text/template/#pkg-overview].
--timeout duration            The length of time to wait before giving
up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
--wait                        If true, wait for operations to complete
before returning.

```

11.2.7.11.9.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.7.11.9.3 SEE ALSO

- [dkp create nodepool](#) (see page 1687) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.7.11.10 dkp create nodepool vcd

Create a nodepool in VCD

```
dkp create nodepool vcd name [flags]
```

11.2.7.11.10.1 Options

<code>--allow-missing-template-keys</code>	If true , ignore any errors in templates when a field or map key is missing in the jsonpath output formats. (default true)
<code>--catalog string</code>	Catalog hosting virtual machine templates.
<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>--dry-run</code>	Only print the objects that would be created, without creating them.
<code>-h, --help</code>	help for vcd
<code>--http-proxy string</code>	HTTP proxy for nodes
<code>--https-proxy string</code>	HTTPS proxy for nodes
<code>--kubeconfig string</code>	Path to the kubeconfig for the management cluster. If unspecified, default discovery rules apply.
<code>--kubernetes-version string</code>	Kubernetes version (default "1.26.14")
<code>-n, --namespace string</code>	If present, the namespace scope for this CLI request. (default "default")
<code>--no-proxy strings</code>	No Proxy list for nodes (default [])
<code>-o, --output string</code>	Output format. One of: (json, yaml, name, go-template, go-template-file, template, jsonpath, jsonpath-as-json, jsonpath-file).
<code>--output-directory string</code>	Used with <code>--output=json yaml</code> . The directory where to output resources to files. The directory must already exist.
<code>--registry-mirror-cacert file</code>	CA certificate chain to use while communicating with the registry mirror using TLS
<code>--registry-mirror-password string</code>	Password to authenticate to the registry mirror with

<code>--registry-mirror-url url</code>	URL of a container registry to use as a mirror in the cluster
<code>--registry-mirror-username string</code>	Username to authenticate to the registry mirror with
<code>--replicas int</code>	Number of replicas (default 1)
<code>--show-managed-fields</code>	If true , keep the managedFields when printing objects in JSON or YAML format.
<code>--ssh-public-key-file string</code>	Path to the authorized SSH key for the user
<code>--ssh-username string</code>	Name of the user to create on the instance (default "konvoy")
<code>--storage-profile string</code>	The storage profile to be used for a virtual machine (default "*")
<code>--template string</code>	Template string or path to template file to use when <code>-o=go-template</code> , <code>-o=go-template-file</code> . The template format is go lang templates [http://golang.org/pkg/text/template/#pkg-overview].
<code>--timeout duration</code>	The length of time to wait before giving up. Zero means wait forever (e.g. 300s, 30m, 3h). (default 30m0s)
<code>--vapp-template string</code>	The vApp template used to create the virtual machine.
<code>--wait</code>	If true , wait for operations to complete before returning.
<code>--worker-placement-policy string</code>	The placement policy for workers to be used on this machine
<code>--worker-sizing-policy string</code>	The sizing policy for workers to be used on this machine

11.2.7.11.10.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

11.2.7.11.10.3 SEE ALSO

- [dkp create nodepool](#) (see page 1687) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vcd, vsphere]

11.2.8 dkp delete

Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

11.2.8.1 Options

`-h, --help` help **for** delete

11.2.8.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.8.3 SEE ALSO

- [dkp delete bootstrap](#) (see page 1698) - Delete bootstrap cluster
- [dkp delete capi-components](#) (see page 1698) - Delete the CAPI components from the cluster
- [dkp delete chart](#) (see page 1700) - Delete a chart from the repository
- [dkp delete cluster](#) (see page 1699) - Delete a Kubernetes cluster
- [dkp delete nodepool](#) (see page 1701) - Delete a nodepool for a given cluster

11.2.8.4 dkp delete bootstrap

Delete bootstrap cluster

```
dkp delete bootstrap [flags]
```

11.2.8.4.1 Options

```
-h, --help                help for bootstrap
--kind-cluster-name string Kind cluster name for the bootstrap cluster (default "konvoy-capi-bootstrapper")
--kubeconfig string       Path to the kubeconfig for the management cluster.
If unspecified, default discovery rules apply.
--timeout duration        The length of time to wait before giving up. Zero
means wait forever (e.g. 300s, 30m, 3h). (default 5m0s)
-v, --verbose int       Output verbosity
--wait                    If true, wait for operations to complete before
returning. (default true)
```

11.2.8.4.2 SEE ALSO

- [dkp delete](#) (see page 1700) - Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

11.2.8.5 dkp delete capi-components

Delete the CAPI components from the cluster

```
dkp delete capi-components [flags]
```

11.2.8.5.1 Options

```
-h, --help                help for capi-components
--kubeconfig string      Path to the kubeconfig for the management cluster. If
unspecified, default discovery rules apply.
--timeout duration       The length of time to wait before giving up. Zero means
wait forever (e.g. 300s, 30m, 3h). (default 5m0s)
-v, --verbose int       Output verbosity
--wait                   If true, wait for operations to complete before
returning. (default true)
```

11.2.8.5.2 SEE ALSO

- [dkp delete](#) (see page 1700) - Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

11.2.8.6 dkp delete cluster

Delete a Kubernetes cluster

```
dkp delete cluster [flags]
```

11.2.8.6.1 Options

```
--aws-service-endpoints string    Custom AWS service endpoints in a semi-colon
separated format: ${SigningRegion1}:${ServiceID1}=${URL},${ServiceID2}=${URL};$
${SigningRegion2}...
-c, --cluster-name name           Name used to prefix the cluster and all the
created resources.
--delete-kubernetes-resources      Delete Kubernetes resources on the cluster
before deleting that cluster (Services with type LoadBalancer) (default true)
-h, --help                       help for cluster
--http-proxy string              HTTP proxy for CAPI controllers
--https-proxy string             HTTPS proxy for CAPI controllers
--kind-cluster-image string       Kind node image for the bootstrap cluster (d
efault "mesosphere/konvoy-bootstrap:v2.6.2-rc.4")
--kubeconfig string              Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
-n, --namespace string           If present, the namespace scope for this CLI
request. (default "default")
```

```

--no-proxy strings          No Proxy list for CAPI controllers (default
[])
--self-managed              When set to true, the required prerequisites
and resources are moved from the self managed cluster before deleting. When set to
false, the resources are assumed installed in a management cluster. (default false)
--timeout duration         The length of time to wait before giving up.
Zero means wait forever (e.g. 300s, 30m, 3h). (default 15m0s)
-v, --verbose int        Output verbosity
--wait                      If true, wait for operations to complete
before returning. This flag is ignored and will always be 'true' if used with the --
self-managed flag. (default true)
--with-aws-bootstrap-credentials Set true to use AWS bootstrap credentials
from your environment. When false, the instance profile of the EC2 instance where the
CAPA controller is scheduled on will be used instead.
--with-gcp-bootstrap-credentials Set true to use GCP bootstrap credentials
from your environment. When false, the service account of the VM instance where the
CAPG controller is scheduled on will be used instead.

```

11.2.8.6.2 SEE ALSO

- [dkp delete](#) (see page 1700) - Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

11.2.8.7 dkp delete chart

Delete a chart from the repository

```
dkp delete chart [chartName] [chartVersion] [flags]
```

11.2.8.7.1 Options

```

--config string            Config file to use (default "/home/
runner/.kommander/config")
--context string          The name of the kubeconfig context to use
-h, --help                help for chart
--kubeconfig string       Path to the kubeconfig file to use for CLI requests.
--request-timeout string  The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g.
1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int        Output verbosity

```

11.2.8.7.2 SEE ALSO

- [dkp delete](#) (see page 1700) - Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

11.2.8.8 dkp delete nodepool

Delete a nodepool for a given cluster

```
dkp delete nodepool name [flags]
```

11.2.8.8.1 Options

```
-c, --cluster-name name  Name used to prefix the cluster and all the created
resources.
--dry-run                Only print the objects that would be created, without
creating them.
-h, --help               help for nodepool
--kubeconfig string     Path to the kubeconfig for the management cluster. If
unspecified, default discovery rules apply.
-n, --namespace string  If present, the namespace scope for this CLI request. (de
fault "default")
-v, --verbose int      Output verbosity
```

11.2.8.8.2 SEE ALSO

- [dkp delete](#) (see page 1700) - Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

11.2.9 dkp describe

Describe one of [cluster]

11.2.9.1 Options

```
-h, --help               help for describe
-v, --verbose int      Output verbosity
```

11.2.9.2 SEE ALSO

- [dkp describe cluster](#) (see page 1701) - Describe a Kubernetes cluster status

11.2.9.3 dkp describe cluster

Describe a Kubernetes cluster status


```
dkp describe cluster [flags]
```

11.2.9.3.1 Options

```
-c, --cluster-name name    Name used to prefix the cluster and all the created
resources.
-h, --help                help for cluster
--kubeconfig string       Path to the kubeconfig for the management cluster. If
unspecified, default discovery rules apply.
-n, --namespace string     If present, the namespace scope for this CLI request. (default "default")
```

11.2.9.3.2 Options inherited from parent commands

```
-v, --verbose int    Output verbosity
```

11.2.9.3.3 SEE ALSO

- [dkp describe](#) (see page 1704) - Describe one of [cluster]

11.2.10 dkp detach

Detach one of [cluster]

11.2.10.1 Options

```
--config string          Config file to use (default "/home/runner/.kommander/config")
--context string         The name of the kubeconfig context to use
-h, --help              help for detach
--kubeconfig string      Path to the kubeconfig file to use for CLI requests.
--request-timeout string The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int       Output verbosity
```

11.2.10.2 SEE ALSO

- [dkp detach cluster](#) (see page 1703) - Detach a cluster

11.2.10.3 dkp detach cluster

Detach a cluster

```
dkp detach cluster CLUSTER_NAME [flags]
```

11.2.10.3.1 Options

```
-h, --help                help for cluster
--timeout duration       The length of time to wait before giving up on a detach,
defaults to wait forever (default 0s)
--wait                   If true, wait for resources to be gone before returning.
This waits for finalizers. (default true)
-w, --workspace string   Name of the workspace of the attached cluster
```

11.2.10.3.2 Options inherited from parent commands

```
--config string          Config file to use (default "/home/
runner/.kommander/config")
--context string         The name of the kubeconfig context to use
--kubeconfig string      Path to the kubeconfig file to use for CLI requests.
--request-timeout string The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g.
1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int       Output verbosity
```

11.2.10.3.3 SEE ALSO

- [dkp detach](#) (see page 1705) - Detach one of [cluster]

11.2.11 dkp diagnose

Generate a support bundle

```
dkp diagnose [flags]
```

11.2.11.1 Options

```

--allow-insecure-connections    When set, do not verify TLS certs when
retrieving spec and reporting results
--as string                     Username to impersonate for the operation.
User could be a regular user or a service account in a namespace.
--as-group stringArray         Group to impersonate for the operation, this
flag can be repeated to specify multiple groups. (default [])
--as-uid string                UID to impersonate for the operation.
--bootstrap-kubeconfig string  Path to the kubeconfig file to use for
requests towards an additional bootstrap cluster
--bundle-name-prefix string    Prefix added to support bundle name created
--cache-dir string            Default cache directory (default "/home/
runner/.kube/cache")
--certificate-authority string  Path to a cert file for the certificate
authority
--client-certificate string    Path to a client certificate file for TLS
--client-key string           Path to a client key file for TLS
--cluster string              The name of the kubeconfig cluster to use
--collect-without-permissions Always generate a support bundle, even if it
some require additional permissions (default true)
--context string              The name of the kubeconfig context to use
--disable-compression         If true, opt-out of response compression for
all requests to the server
-h, --help                   help for diagnose
--insecure-skip-tls-verify    If true, the server's certificate will not be
checked for validity. This will make your HTTPS connections insecure
--kubeconfig string           Path to the kubeconfig file to use for CLI
requests.
-n, --namespace string        If present, the namespace scope for this CLI
request
--redactors strings           Names of the additional redactors to use
(default [])
--request-timeout string      The length of time to wait before giving up on
a single server request. Non-zero values should contain a corresponding time unit
(e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
-s, --server string           The address and port of the Kubernetes API
server
--since string                Force pod logs collectors to return logs newer
than a relative duration like 5s, 2m, or 3h.
--since-time string           Force pod logs collectors to return logs after
a specific date (RFC3339)
--tls-server-name string      Server name to use for server certificate
validation. If it is not provided, the hostname used to contact the server is used
--token string                Bearer token for authentication to the API
server
--user string                 The name of the kubeconfig user to use
-v, --verbose int           Output verbosity

```

11.2.11.2 SEE ALSO

- [dkp diagnose default-config](#) (see page 1705) - Prints the default configuration of the diagnostics bundle collectors
- [dkp diagnose ssh](#) (see page 1705) - Collect node-level diagnostics data over SSH

11.2.11.3 dkp diagnose default-config

Prints the default configuration of the diagnostics bundle collectors

```
dkp diagnose default-config [flags]
```

11.2.11.3.1 Options

```
-h, --help help for default-config
```

11.2.11.3.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.11.3.3 SEE ALSO

- [dkp diagnose](#) (see page 1706) - Generate a support bundle

11.2.11.4 dkp diagnose ssh

Collect node-level diagnostics data over SSH

```
dkp diagnose ssh path/to/inventory-file.yaml [flags]
```

11.2.11.4.1 Options

```
--bundle-name-prefix string Prefix added to support bundle name created  
-h, --help help for ssh
```

```

--redactors strings      Names of the additional redactors to use (default
[])
--timeout duration      Timeout for collecting bundle per node (default
5m0s)

```

11.2.11.4.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.11.4.3 SEE ALSO

- [dkp diagnose](#) (see page 1706) - Generate a support bundle

11.2.12 dkp edit

Edit a resource on the server

11.2.12.1 Synopsis

Edit a resource from the default editor.

The edit command allows you to directly edit any API resource you can retrieve via the command-line tools. It will open the editor defined by your KUBE_EDITOR, or EDITOR environment variables, or fall back to 'vi' for Linux or 'notepad' for Windows. You can edit multiple objects, although changes are applied one at a time. The command accepts file names as well as command-line arguments, although the files you point to must be previously saved versions of resources.

```
dkp edit (RESOURCE/NAME | -f FILENAME)
```

11.2.12.2 Options

```

--allow-missing-template-keys  If true, ignore any errors in templates when a
field or map key is missing in the template. Only applies to golang and jsonpath
output formats. (default true)
--config string                Config file to use (default "/home/
runner/.kommander/config")
--context string              The name of the kubeconfig context to use
--field-manager string        Name of the manager used to track field
ownership. (default "kommander-cli")
-f, --filename strings        Filename, directory, or URL to files to use to
edit the resource (default [])
-h, --help                    help for edit

```

```

--kubeconfig string      Path to the kubeconfig file to use for CLI
requests.
-k, --kustomize string   Process the kustomization directory. This flag
can't be used together with -f or -R.
-n, --namespace string   namespace of the resource (default "default")
-o, --output string      Output format. One of: (json, yaml, name, go-
template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-patch           Output the patch if the resource is edited.
-R, --recursive          Process the directory used in -f, --filename
recursively. Useful when you want to manage related manifests organized within the
same directory.
--request-timeout string The length of time to wait before giving up on
a single server request. Non-zero values should contain a corresponding time unit
(e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
--save-config            If true, the configuration of current object
will be saved in its annotation. Otherwise, the annotation will be unchanged. This
flag is useful when you want to perform kubectl apply on this object in the future.
--show-managed-fields    If true, keep the managedFields when printing
objects in JSON or YAML format.
--subresource string     If specified, edit will operate on the
subresource of the requested object. Must be one of [status]. This flag is alpha and
may change in the future.
--template string        Template string or path to template file to use
when -o=go-template, -o=go-template-file. The template format is go lang templates
[http://golang.org/pkg/text/template/#pkg-overview].
--validate string[="strict"] Must be one of: strict (or true), warn, ignore
(or false).
                        "true" or "strict" will use a schema to
validate the input and fail the request if invalid. It will perform server side
validation if ServerSideFieldValidation is enabled on the api-server, but will fall
back to less reliable client-side validation if not.
                        "warn" will warn about unknown or
duplicate fields without blocking the request if server-side field validation is
enabled on the API server, and behave as "ignore" otherwise.
                        "false" or "ignore" will not perform any
schema validation, silently dropping any unknown or duplicate fields. (default
"strict")
-v, --verbose int      Output verbosity
--windows-line-endings   Defaults to the line ending native to your
platform.

```

11.2.13 dkp get

Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

11.2.13.1 Options

```
-h, --help    help for get
```

11.2.13.2 Options inherited from parent commands

```
-v, --verbose int    Output verbosity
```

11.2.13.3 SEE ALSO

- [dkp get appdeployments](#) (see page 1709) - Get AppDeployments from a Workspace, Project, or all Workspaces and Projects
- [dkp get chart](#) (see page 1711) - Obtain information about charts stored in the repository
- [dkp get clusters](#) (see page 1712) - Get clusters from specified Workspace
- [dkp get kubeconfig](#) (see page 1708) - Retrieve cluster kubeconfig and modify local kubeconfig file
- [dkp get nodepools](#) (see page 1710) - Get nodepools for a given cluster
- [dkp get workspaces](#) (see page 1710) - Get Workspaces

11.2.13.4 dkp get kubeconfig

Retrieve cluster kubeconfig and modify local kubeconfig file

```
dkp get kubeconfig [flags]
```

11.2.13.4.1 Options

```

    --cluster string      Kommander Cluster to get kubeconfig for
-c, --cluster-name name  Name used to prefix the cluster and all the created
resources.
-h, --help              help for kubeconfig
    --kubeconfig string  Path to the kubeconfig for the management cluster. If
unspecified, default discovery rules apply.
-n, --namespace string  If present, the namespace scope for this CLI request. (de
fault "default")
-w, --workspace string  Name of the workspace to show clusters from

```

11.2.13.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.13.4.3 SEE ALSO

- [dkp get \(see page 1710\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

11.2.13.5 dkp get appdeployments

Get AppDeployments from a Workspace, Project, or all Workspaces and Projects

11.2.13.5.1 Synopsis

Prints a table of the most important information about the specified resources. In case the AppDeployments are configured with cluster scoped specification, an optional CLUSTERS column (when printing in table format) will display enabled clusters for each AppDeployment.

```
dkp get appdeployments [APPDEPLOYMENT_NAME] [flags]
```

11.2.13.5.2 Options

-A, --all-namespaces	If present, list the requested object(s) across all namespaces.
--config string	Config file to use (default "/home/runner/.kommander/config")
--context string	The name of the kubeconfig context to use
-h, --help	help for appdeployments
--kubeconfig string	Path to the kubeconfig file to use for CLI requests.
-o, --output string	Output format. One of: table yaml
-p, --project string	Name of the project to show AppDeployments from. Requires workspace flag (workspace that the project belongs to).
--request-timeout string	The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int	Output verbosity
-w, --workspace string	Name of the workspace to show AppDeployments from

11.2.13.5.3 SEE ALSO

- [dkp get \(see page 1710\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

11.2.13.6 dkp get workspaces

Get Workspaces

```
dkp get workspaces [flags]
```

11.2.13.6.1 Options

-A, --all-namespaces	If present, list the requested object(s) across all namespaces.
--config string	Config file to use (default "/home/runner/.kommander/config")
--context string	The name of the kubeconfig context to use
-h, --help	help for workspaces
--kubeconfig string	Path to the kubeconfig file to use for CLI requests.
-o, --output string	Output format. One of: table yaml
--request-timeout string	The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int	Output verbosity

11.2.13.6.2 SEE ALSO

- [dkp get \(see page 1710\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

11.2.13.7 dkp get nodepools

Get nodepools for a given cluster

```
dkp get nodepools [flags]
```

11.2.13.7.1 Options

```

-c, --cluster-name name    Name used to prefix the cluster and all the created
resources.
-h, --help                help for nodepools
  --kubeconfig string      Path to the kubeconfig for the management cluster. If
unspecified, default discovery rules apply.
-n, --namespace string     If present, the namespace scope for this CLI request. (default "default")
-v, --verbose int         Output verbosity

```

11.2.13.7.2 SEE ALSO

- [dkp get \(see page 1710\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

11.2.13.8 dkp get chart

Obtain information about charts stored in the repository

```
dkp get chart [chartName] [chartVersion] [flags]
```

11.2.13.8.1 Options

```

-A, --all-namespaces      If present, list the requested object(s) across all
namespaces.
  --config string         Config file to use (default "/home/runner/.kommander/config")
  --context string        The name of the kubeconfig context to use
-h, --help                help for chart
  --kubeconfig string     Path to the kubeconfig file to use for CLI requests.
-o, --output string        Output format. One of: table|yaml
  --request-timeout string The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int         Output verbosity

```

11.2.13.8.2 SEE ALSO

- [dkp get \(see page 1710\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

11.2.13.9 dkp get clusters

Get clusters from specified Workspace

```
dkp get clusters [CLUSTER_NAME] [flags]
```

11.2.13.9.1 Options

-A, --all-namespaces	If present, list the requested object(s) across all namespaces.
--config string	Config file to use (default "/home/runner/.kommander/config")
--context string	The name of the kubeconfig context to use
-h, --help	help for clusters
--kubeconfig string	Path to the kubeconfig file to use for CLI requests.
-o, --output string	Output format. One of: table yaml
--request-timeout string	The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int	Output verbosity
-w, --workspace string	Name of the workspace to show clusters from

11.2.13.9.2 SEE ALSO

- [dkp get \(see page 1710\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

11.2.14 dkp import

Import images from an image bundle into Containerd

11.2.14.1 Options

-h, --help	help for import
-v, --verbose int	Output verbosity

11.2.14.2 SEE ALSO

- [dkp import image-bundle \(see page 1713\)](#) - Import images from image bundles into Containerd

11.2.14.3 dkp import image-bundle

Import images from image bundles into Containerd

```
dkp import image-bundle [flags]
```

11.2.14.3.1 Options

```

    --containerd-namespace string  Containerd namespace to import images into (default "k8s.io")
    -h, --help                    help for image-bundle
    --image-bundle strings        Tarball containing list of images to import.
    Can also be a glob pattern. (default [])

```

11.2.14.3.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.14.3.3 SEE ALSO

- [dkp import](#) (see page 1715) - Import images from an image bundle into Containerd

11.2.15 dkp install

Install one of [kommander]

11.2.15.1 Options

```

    --config string          Config file to use (default "/home/runner/.kommander/config")
    --context string        The name of the kubeconfig context to use
    -h, --help             help for install
    --kubeconfig string     Path to the kubeconfig file to use for CLI requests.
    --request-timeout string The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
    -v, --verbose int    Output verbosity

```

11.2.15.2 SEE ALSO

- [dkp install kommander \(see page 1714\)](#) - Install kommander

11.2.15.3 dkp install kommander

Install kommander

```
dkp install kommander [flags]
```

11.2.15.3.1 Options

```

--airgapped                Enable airgapped mode.
--bootstrap-repository string git repository with bootstrap
definitions
--charts-bundle stringArray Path to charts-bundle to upload to
chartmuseum, apart from parsing the kommander applications repository (default [])
--disallow-charts-download make CLI rely solely on provided
chart bundles and do not try to download charts from the Internet
--gitea-kommander-repository-name string gitea kommander repository name
(default "kommander")
-h, --help                help for kommander
--init                    Initialize default configuration
file, print it and exit without installing Kommander.
--installer-config file   Path to installation configuration
file
--kommander-applications-repository string git repository with application
definitions (default "v2.6.2-rc.4")
-o, --output string       Output format for the
configuration file generated by --init. One of: yaml|json (default "yaml")
--wait                    Wait for all enabled applications
to be ready (default true)
--wait-timeout duration   Time to wait for all enabled
applications to be ready (30m, 1h, 2h) (default 1h0m0s)

```

11.2.15.3.2 Options inherited from parent commands

```

--config string           Config file to use (default "/home/
runner/.kommander/config")
--context string          The name of the kubeconfig context to use
--kubeconfig string       Path to the kubeconfig file to use for CLI requests.

```

```

--request-timeout string  The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g.
1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int      Output verbosity

```

11.2.15.3.3 SEE ALSO

- [dkp install](#) (see page 1716) - Install one of [kommander]

11.2.16 dkp move

Move one of [capi-resources]

11.2.16.1 Synopsis

Command "move" is deprecated, use "dkp move capi-resources" instead Move one of [capi-resources]

```
dkp move [flags]
```

11.2.16.2 Options

```

--from-context string  Context to be used within the from-cluster's
kubecfg file. If empty, current context will be used. (DEPRECATED: use "dkp move
capi-resources" instead)
--from-kubecfg file    Path to the kubecfg for pivot's source cluster. If
unspecified, default discovery rules apply. (DEPRECATED: use "dkp move capi-
resources" instead)
-h, --help            help for move
-n, --namespace string If present, the namespace scope for this CLI request.
(default "default")
--to-context string    Context to be used within the to-cluster's kubecfg
file. If empty, current context will be used. (DEPRECATED: use "dkp move capi-
resources" instead)
--to-kubecfg file      Path to the kubecfg for pivot's destination cluster
(DEPRECATED: use "dkp move capi-resources" instead)
--to-namespace string  Resources are moved to this namespace in the to-
cluster. By default, the same as the from-cluster namespace.
-v, --verbose int    Output verbosity

```

11.2.16.3 SEE ALSO

- [dkp move capi-resources](#) (see page 1716) - Move controllers and objects from one cluster to the other

11.2.16.4 dkp move capi-resources

Move controllers and objects from one cluster to the other

```
dkp move capi-resources [flags]
```

11.2.16.4.1 Options

```

    --from-context string      Context to be used within the from-cluster's
    kubeconfig file. If empty, current context will be used.
    --from-kubeconfig file     Path to the kubeconfig for pivot's source cluster. If
    unspecified, default discovery rules apply.
    -h, --help                 help for capi-resources
    -n, --namespace string     If present, the namespace scope for this CLI request.
    (default "default")
    --to-context string        Context to be used within the to-cluster's kubeconfig
    file. If empty, current context will be used.
    --to-kubeconfig file       Path to the kubeconfig for pivot's destination cluster
    --to-namespace string      Resources are moved to this namespace in the to-
    cluster. By default, the same as the from-cluster namespace.

```

11.2.16.4.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.16.4.3 SEE ALSO

- [dkp move \(see page 1718\)](#) - Move one of [capi-resources]

11.2.17 dkp open

Open one of [dashboard]

11.2.17.1 Options

```

    --config string           Config file to use (default "/home/
    runner/.kommander/config")
    --context string          The name of the kubeconfig context to use

```

```

-h, --help                help for open
--kubeconfig string      Path to the kubeconfig file to use for CLI requests.
--request-timeout string  The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g.
1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int       Output verbosity

```

11.2.17.2 SEE ALSO

- [dkp open dashboard](#) (see page 1717) - Open DKP UI in your browser

11.2.17.3 dkp open dashboard

Open DKP UI in your browser

```
dkp open dashboard [flags]
```

11.2.17.3.1 Options

```
-h, --help    help for dashboard
```

11.2.17.3.2 Options inherited from parent commands

```

--config string          Config file to use (default "/home/
runner/.kommander/config")
--context string         The name of the kubeconfig context to use
--kubeconfig string      Path to the kubeconfig file to use for CLI requests.
--request-timeout string The length of time to wait before giving up on a
single server request. Non-zero values should contain a corresponding time unit (e.g.
1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int       Output verbosity

```

11.2.17.3.3 SEE ALSO

- [dkp open](#) (see page 1719) - Open one of [dashboard]

11.2.18 dkp push

Push one of [bundle, chart, chart-bundle, image-bundle]

11.2.18.1 Options

```
-h, --help    help for push
```

11.2.18.2 Options inherited from parent commands

```
-v, --verbose int    Output verbosity
```

11.2.18.3 SEE ALSO

- [dkp push bundle \(see page 1719\)](#) - Push from bundles into an existing OCI registry
- [dkp push chart \(see page 1719\)](#) - Upload charts to the repository
- [dkp push chart-bundle \(see page 1718\)](#) - Upload chart bundles to the repository

11.2.18.4 dkp push chart-bundle

Upload chart bundles to the repository

```
dkp push chart-bundle [chartsTarball]... [flags]
```

11.2.18.4.1 Options

```

--config string          Config file to use (default "/home/runner/.kommander/config")
--context string         The name of the kubeconfig context to use
-h, --help              help for chart-bundle
--kubeconfig string     Path to the kubeconfig file to use for CLI requests.
--request-timeout string The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int     Output verbosity

```

11.2.18.4.2 SEE ALSO

- [dkp push \(see page 1720\)](#) - Push one of [bundle, chart, chart-bundle, image-bundle]

11.2.18.5 dkp push chart

Upload charts to the repository

```
dkp push chart [chartTarball]... [flags]
```

11.2.18.5.1 Options

```

    --config string           Config file to use (default "/home/runner/.kommander/config")
    --context string         The name of the kubeconfig context to use
    -h, --help               help for chart
    --kubeconfig string      Path to the kubeconfig file to use for CLI requests.
    --request-timeout string The length of time to wait before giving up on a
    single server request. Non-zero values should contain a corresponding time unit (e.g.
    1s, 2m, 3h). A value of zero means don't timeout requests.
    -v, --verbose int      Output verbosity

```

11.2.18.5.2 SEE ALSO

- [dkp push](#) (see page 1720) - Push one of [bundle, chart, chart-bundle, image-bundle]

11.2.18.6 dkp push bundle

Push from bundles into an existing OCI registry

```
dkp push bundle [flags]
```

11.2.18.6.1 Options

```

    --bundle strings          Tarball containing list of images to
    push. Can also be a glob pattern. (default [])
    --ecr-lifecycle-policy-file string File containing ECR lifecycle policy
for newly created repositories (only applies if target registry is hosted on ECR,
    ignored otherwise)
    -h, --help               help for bundle
    --image-push-concurrency int Image push concurrency (default 1)
    --on-existing-tag string how to handle existing tags: one of
    "overwrite", "error", or "skip" (default "overwrite")

```

```

--to-registry string           Registry to push images to. TLS
verification will be skipped when using an http:// registry.
--to-registry-ca-cert-file string CA certificate file used to verify TLS
verification of registry to push images to
--to-registry-insecure-skip-tls-verify Skip TLS verification of registry to
push images to (also use for non-TLS http registries)
--to-registry-password string Password to use to log in to
destination registry
--to-registry-username string Username to use to log in to
destination registry
-v, --verbose int           Output verbosity

```

11.2.18.6.2 SEE ALSO

- [dkp push \(see page 1720\)](#) - Push one of [bundle, chart, chart-bundle, image-bundle]

11.2.19 dkp scale

Scale one of [nodepool]

11.2.19.1 Options

```

-h, --help           help for scale
-v, --verbose int Output verbosity

```

11.2.19.2 SEE ALSO

- [dkp scale nodepool \(see page 1720\)](#) - Scale a nodepool of a given cluster to the number of replicas

11.2.19.3 dkp scale nodepool

Scale a nodepool of a given cluster to the number of replicas

```
dkp scale nodepool name [flags]
```

11.2.19.3.1 Options

```

-c, --cluster-name name Name used to prefix the cluster and all the created
resources.
-h, --help           help for nodepool

```

```

--kubeconfig string      Path to the kubeconfig for the management cluster.
If unspecified, default discovery rules apply.
-n, --namespace string  If present, the namespace scope for this CLI
request. (default "default")
--nodes-to-delete strings A list of node names to mark for deletion when
scaling down a node pool. If left empty, the nodes to delete will be selected at
random. (default [])
--replicas int        The new desired number of replicas.

```

11.2.19.3.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.19.3.3 SEE ALSO

- [dkp scale](#) (see page 1723) - Scale one of [nodepool]

11.2.20 dkp serve

Serve image or Helm chart bundles from an OCI registry

11.2.20.1 Options

```

-h, --help            help for serve
-v, --verbose int  Output verbosity

```

11.2.20.2 SEE ALSO

- [dkp serve bundle](#) (see page 1721) - Serve an OCI registry from previously created bundles

11.2.20.3 dkp serve bundle

Serve an OCI registry from previously created bundles

```
dkp serve bundle [flags]
```

11.2.20.3.1 Options

<code>--bundle</code> strings	Bundle to serve. Can also be a glob pattern.
(default [])	
<code>-h, --help</code>	help for bundle
<code>--listen-address</code> string	Address to listen on (default "127.0.0.1")
<code>--listen-port</code> uint16	Port to listen on (0 means use any free port)
<code>--tls-cert-file</code> string	TLS certificate file
<code>--tls-private-key-file</code> string	TLS private key file

11.2.20.3.2 Options inherited from parent commands

`-v, --verbose` **int** Output verbosity

11.2.20.3.3 SEE ALSO

- [dkp serve](#) (see page 1724) - Serve image or Helm chart bundles from an OCI registry

11.2.21 dkp update

Update one of [bootstrap (cluster), controlplane, nodepool]

11.2.21.1 Options

<code>-h, --help</code>	help for update
<code>-v, --verbose</code> int	Output verbosity

11.2.21.2 SEE ALSO

- [dkp update bootstrap](#) (see page 1734) - Update bootstrap cluster
- [dkp update controlplane](#) (see page 1728) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]
- [dkp update nodepool](#) (see page 1722) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.3 dkp update nodepool

Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.3.1 Options

```
-h, --help help for nodepool
```

11.2.21.3.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.21.3.3 SEE ALSO

- [dkp update](#) (see page 1725) - Update one of [bootstrap (cluster), controlplane, nodepool]
- [dkp update nodepool aws](#) (see page 1724) - Update a Konvoy cluster node pool in AWS
- [dkp update nodepool azure](#) (see page 1723) - Update a Konvoy cluster node pool in Azure
- [dkp update nodepool eks](#) (see page 1726) - Update a Konvoy cluster node pool in EKS
- [dkp update nodepool gcp](#) (see page 1727) - Update a Konvoy cluster node pool in GCP
- [dkp update nodepool preprovisioned](#) (see page 1725) - Update a Konvoy cluster node pool in Preprovisioned
- [dkp update nodepool vsphere](#) (see page 1726) - Update a Konvoy cluster node pool in vSphere

11.2.21.3.4 dkp update nodepool azure

Update a Konvoy cluster node pool in Azure

```
dkp update nodepool azure [flags]
```

11.2.21.3.4.1 Options

```
-c, --cluster-name name          Name used to prefix the cluster and all the
created resources.
--compute-gallery-id string      Compute Gallery ID of a custom image, e.g., '/
subscriptions/<subscription id>/resourceGroups/<resource group name>/providers/
Microsoft.Compute/galleries/<gallery name>/images/<image definition name>/versions/
<version id>' (replacing placeholders with the values used when creating the image)
-h, --help                       help for azure
--kubeconfig string              Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--kubernetes-version string      Kubernetes version
--machine-size string            Worker machine size (ex. 'Standard_D2s_v3')
```

```

-n, --namespace string          If present, the namespace scope for this CLI
request. (default "default")
  --plan-offer string           The offer for a Marketplace image or a custom
image sourced from a Marketplace image requiring Plan information.
  --plan-publisher string      The publisher for a Marketplace image or a custom
image sourced from a Marketplace image requiring Plan information.
  --plan-sku string            The SKU for a Marketplace image or a custom image
sourced from a Marketplace image requiring Plan information.
  --use-context string         Use a specific context in a kubeconfig file.
  --wait                       If true, wait for operations to complete before
returning. (default true)

```

11.2.21.3.4.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.21.3.4.3 SEE ALSO

- [dkp update nodepool](#) (see page 1725) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.3.5 dkp update nodepool aws

Update a Konvoy cluster node pool in AWS

```
dkp update nodepool aws [flags]
```

11.2.21.3.5.1 Options

```

--ami string                AMI ID to use for machines
--ami-base-os string        Base OS used in search of AMIs. Examples:
'centos-7', 'ubuntu-18.04', 'ubuntu-20.04'
--ami-format string         Query string used in search of AMIs. Example:
When --ami-base-os='rhel8.4', then the string 'prefix-{{.BaseOS}}-?{{.K8sVersion}}-*'
matches any AMIs with the Name 'prefix-rhel8.4-1.26.14'
--ami-owner string         ID of AWS account used in search of AMIs
-c, --cluster-name name    Name used to prefix the cluster and all the
created resources.
-h, --help                 help for aws
--instance-type string     Instance type to use for node pool machines
--kubeconfig string        Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--kubernetes-version string Kubernetes version

```

<p><code>-n, --namespace string</code> request. (default "default")</p> <p><code>--use-context string</code></p> <p><code>--wait</code> returning. (default true)</p>	<p>If present, the namespace scope for this CLI</p> <p>Use a specific context in a kubeconfig file.</p> <p>If true, wait for operations to complete before</p>
---	---

11.2.21.3.5.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

11.2.21.3.5.3 SEE ALSO

- [dkp update nodepool](#) (see page 1725) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.3.6 dkp update nodepool preprovisioned

Update a Konvoy cluster node pool in Preprovisioned

`dkp update nodepool preprovisioned [flags]`

11.2.21.3.6.1 Options

<p><code>-c, --cluster-name name</code> created resources.</p> <p><code>-h, --help</code></p> <p><code>--kubeconfig string</code> cluster. If unspecified, default discovery rules apply.</p> <p><code>--kubernetes-version string</code> Kubernetes version</p> <p><code>-n, --namespace string</code> request. (default "default")</p> <p><code>--use-context string</code></p> <p><code>--wait</code> returning. (default true)</p>	<p>Name used to prefix the cluster and all the</p> <p>help for preprovisioned</p> <p>Path to the kubeconfig for the management</p> <p>If present, the namespace scope for this CLI</p> <p>Use a specific context in a kubeconfig file.</p> <p>If true, wait for operations to complete before</p>
---	--

11.2.21.3.6.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

11.2.21.3.6.3 SEE ALSO

- [dkp update nodepool](#) (see page 1725) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.3.7 dkp update nodepool eks

Update a Konvoy cluster node pool in EKS

```
dkp update nodepool eks [flags]
```

11.2.21.3.7.1 Options

-c, --cluster-name name	Name used to prefix the cluster and all the created resources.
-h, --help	help for eks
--instance-type string	Instance type to use for node pool machines
--kubeconfig string	Path to the kubeconfig for the management cluster. If unspecified, default discovery rules apply.
--kubernetes-version string	Kubernetes version
-n, --namespace string	If present, the namespace scope for this CLI request. (default "default")
--use-context string	Use a specific context in a kubeconfig file.
--wait	If true , wait for operations to complete before returning. (default true)

11.2.21.3.7.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.21.3.7.3 SEE ALSO

- [dkp update nodepool](#) (see page 1725) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.3.8 dkp update nodepool vsphere

Update a Konvoy cluster node pool in vSphere

```
dkp update nodepool vsphere [flags]
```

11.2.21.3.8.1 Options

<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>--cpus int</code>	The number of virtual processors in a virtual machine.
<code>--disk-size int</code>	The size of a virtual machine's disk, in GB.
<code>-h, --help</code>	help for vsphere
<code>--kubeconfig string</code>	Path to the kubeconfig for the management cluster. If unspecified, default discovery rules apply.
<code>--kubernetes-version string</code>	Kubernetes version
<code>--memory int</code>	The size of a virtual machine's memory, in GB.
<code>-n, --namespace string</code>	If present, the namespace scope for this CLI request. (default "default")
<code>--use-context string</code>	Use a specific context in a kubeconfig file.
<code>--vm-template string</code>	The virtual machine template to use for a virtual machine.
<code>--wait</code>	If true , wait for operations to complete before returning. (default true)

11.2.21.3.8.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.21.3.8.3 SEE ALSO

- [dkp update nodepool](#) (see page 1725) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.3.9 dkp update nodepool gcp

Update a Konvoy cluster node pool in GCP

```
dkp update nodepool gcp [flags]
```

11.2.21.3.9.1 Options

<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>-h, --help</code>	help for gcp
<code>--image string</code>	Full reference to an image to use for all nodes (set either this or <code>--image-family</code>) (ex. <code>'projects/my-project/global/images/konvoy-ubuntu-2004-1-99-99-1234567890'</code>)
<code>--image-family string</code>	Full reference to an image family to use for all nodes (set either this or <code>--image</code>) (ex. <code>'projects/my-project/global/images/family/konvoy-ubuntu-2004-{{.K8sVersion}}'</code>)
<code>--instance-type string</code>	Worker machine instance type (ex. <code>"n2-standard-8"</code>)
<code>--kubeconfig string</code>	Path to the kubeconfig for the management cluster. If unspecified, default discovery rules apply.
<code>--kubernetes-version string</code>	Kubernetes version
<code>-n, --namespace string</code>	If present, the namespace scope for this CLI request. (default <code>"default"</code>)
<code>--use-context string</code>	Use a specific context in a kubeconfig file.
<code>--wait</code>	If true , wait for operations to complete before returning. (default true)

11.2.21.3.9.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

11.2.21.3.9.3 SEE ALSO

- [dkp update nodepool](#) (see page 1725) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.4 dkp update controlplane

Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.4.1 Options

`-h, --help` help **for** controlplane

11.2.21.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.21.4.3 SEE ALSO

- [dkp update](#) (see page 1725) - Update one of [bootstrap (cluster), controlplane, nodepool]
- [dkp update controlplane aws](#) (see page 1730) - Update a Konvoy cluster control plane in AWS
- [dkp update controlplane azure](#) (see page 1731) - Update a Konvoy cluster control plane in Azure
- [dkp update controlplane eks](#) (see page 1732) - Update a Konvoy cluster control plane in EKS
- [dkp update controlplane gcp](#) (see page 1733) - Update a Konvoy cluster control plane in GCP
- [dkp update controlplane preprovisioned](#) (see page 1729) - Update a Konvoy cluster control plane in Preprovisioned
- [dkp update controlplane vsphere](#) (see page 1732) - Update a Konvoy cluster control plane in vSphere

11.2.21.4.4 dkp update controlplane preprovisioned

Update a Konvoy cluster control plane in Preprovisioned

```
dkp update controlplane preprovisioned [flags]
```

11.2.21.4.4.1 Options

```
-c, --cluster-name name          Name used to prefix the cluster and all the
created resources.
-h, --help                      help for preprovisioned
--kubeconfig string            Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--kubernetes-version string    Kubernetes version
-n, --namespace string         If present, the namespace scope for this CLI
request. (default "default")
--use-context string           Use a specific context in a kubeconfig file.
--wait                         If true, wait for operations to complete before
returning. (default true)
```

11.2.21.4.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.21.4.4.3 SEE ALSO

- [dkp update controlplane](#) (see page 1731) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.4.5 dkp update controlplane aws

Update a Konvoy cluster control plane in AWS

```
dkp update controlplane aws [flags]
```

11.2.21.4.5.1 Options

```

--ami string                AMI ID to use for machines
--ami-base-os string        Base OS used in search of AMIs. Examples:
'centos-7', 'ubuntu-18.04', 'ubuntu-20.04'
--ami-format string         Query string used in search of AMIs. Example:
When --ami-base-os='rhel8.4', then the string 'prefix-{{.BaseOS}}-?{{.K8sVersion}}-*'
matches any AMIs with the Name 'prefix-rhel8.4-1.26.14'
--ami-owner string          ID of AWS account used in search of AMIs
-c, --cluster-name name     Name used to prefix the cluster and all the
created resources.
-h, --help                  help for aws
--instance-type string      Instance type to use for control plane machines
--kubeconfig string         Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--kubernetes-version string Kubernetes version
-n, --namespace string      If present, the namespace scope for this CLI
request. (default "default")
--use-context string        Use a specific context in a kubeconfig file.
--wait                      If true, wait for operations to complete before
returning. (default true)

```

11.2.21.4.5.2 Options inherited from parent commands

```
-v, --verbose int      Output verbosity
```

11.2.21.4.5.3 SEE ALSO

- [dkp update controlplane](#) (see page 1731) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.4.6 dkp update controlplane azure

Update a Konvoy cluster control plane in Azure

```
dkp update controlplane azure [flags]
```

11.2.21.4.6.1 Options

```
-c, --cluster-name name          Name used to prefix the cluster and all the
created resources.
--compute-gallery-id string      Compute Gallery ID of a custom image, e.g., '/
subscriptions/<subscription id>/resourceGroups/<resource group name>/providers/
Microsoft.Compute/galleries/<gallery name>/images/<image definition name>/versions/
<version id>' (replacing placeholders with the values used when creating the image)
-h, --help                       help for azure
--kubeconfig string              Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
--kubernetes-version string      Kubernetes version
--machine-size string            Worker machine size (ex. 'Standard_D2s_v3')
-n, --namespace string           If present, the namespace scope for this CLI
request. (default "default")
--plan-offer string              The offer for a Marketplace image or a custom
image sourced from a Marketplace image requiring Plan information.
--plan-publisher string          The publisher for a Marketplace image or a custom
image sourced from a Marketplace image requiring Plan information.
--plan-sku string                The SKU for a Marketplace image or a custom image
sourced from a Marketplace image requiring Plan information.
--use-context string             Use a specific context in a kubeconfig file.
--wait                           If true, wait for operations to complete before
returning. (default true)
```

11.2.21.4.6.2 Options inherited from parent commands

```
-v, --verbose int    Output verbosity
```

11.2.21.4.6.3 SEE ALSO

- [dkp update controlplane](#) (see page 1731) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.4.7 dkp update controlplane eks

Update a Konvoy cluster control plane in EKS

```
dkp update controlplane eks [flags]
```

11.2.21.4.7.1 Options

<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>-h, --help</code>	help for eks
<code>--kubeconfig string</code>	Path to the kubeconfig for the management cluster. If unspecified, default discovery rules apply.
<code>--kubernetes-version string</code>	Kubernetes version
<code>-n, --namespace string</code>	If present, the namespace scope for this CLI request. (default "default")
<code>--use-context string</code>	Use a specific context in a kubeconfig file.
<code>--wait</code>	If true , wait for operations to complete before returning. (default true)

11.2.21.4.7.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.21.4.7.3 SEE ALSO

- [dkp update controlplane](#) (see page 1731) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.4.8 dkp update controlplane vsphere

Update a Konvoy cluster control plane in vSphere

```
dkp update controlplane vsphere [flags]
```

11.2.21.4.8.1 Options

<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>--cpus int</code>	The number of virtual processors in a virtual machine.
<code>--disk-size int</code>	The size of a virtual machine's disk, in GB.
<code>-h, --help</code>	help for vsphere
<code>--kubeconfig string</code>	Path to the kubeconfig for the management cluster. If unspecified, default discovery rules apply.
<code>--kubernetes-version string</code>	Kubernetes version
<code>--memory int</code>	The size of a virtual machine's memory, in GB.
<code>-n, --namespace string</code>	If present, the namespace scope for this CLI request. (default "default")
<code>--use-context string</code>	Use a specific context in a kubeconfig file.
<code>--vm-template string</code>	The virtual machine template to use for a virtual machine.
<code>--wait</code>	If true , wait for operations to complete before returning. (default true)

11.2.21.4.8.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

11.2.21.4.8.3 SEE ALSO

- [dkp update controlplane](#) (see page 1731) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.4.9 dkp update controlplane gcp

Update a Konvoy cluster control plane in GCP

```
dkp update controlplane gcp [flags]
```

11.2.21.4.9.1 Options

<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>-h, --help</code>	help for gcp

<code>--image</code> string	Full reference to an image to use for all nodes (set either this or <code>--image-family</code>) (ex. 'projects/my-project/global/images/konvoy-ubuntu-2004-1-99-99-1234567890')
<code>--image-family</code> string	Full reference to an image family to use for all nodes (set either this or <code>--image</code>) (ex. 'projects/my-project/global/images/family/konvoy-ubuntu-2004-{{.K8sVersion}}')
<code>--instance-type</code> string	Control Plane machine instance type (ex. "n2-standard-4")
<code>--kubeconfig</code> string	Path to the kubeconfig for the management cluster. If unspecified, default discovery rules apply.
<code>--kubernetes-version</code> string	Kubernetes version
<code>-n</code> , <code>--namespace</code> string	If present, the namespace scope for this CLI request. (default "default")
<code>--use-context</code> string	Use a specific context in a kubeconfig file.
<code>--wait</code>	If true , wait for operations to complete before returning. (default true)

11.2.21.4.9.2 Options inherited from parent commands

`-v`, `--verbose` **int** Output verbosity

11.2.21.4.9.3 SEE ALSO

- [dkp update controlplane](#) (see page 1731) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.21.5 dkp update bootstrap

Update bootstrap cluster

11.2.21.5.1 Options

`-h`, `--help` help **for** bootstrap

11.2.21.5.2 Options inherited from parent commands

`-v`, `--verbose` **int** Output verbosity

11.2.21.5.3 SEE ALSO

- [dkp update](#) (see page 1725) - Update one of [bootstrap (cluster), controlplane, nodepool]
- [dkp update bootstrap credentials](#) (see page 1735) - Update credentials in the cluster

11.2.21.5.4 dkp update bootstrap credentials

Update credentials in the cluster

11.2.21.5.4.1 Options

```
-h, --help    help for credentials
```

11.2.21.5.4.2 Options inherited from parent commands

```
-v, --verbose int    Output verbosity
```

11.2.21.5.4.3 SEE ALSO

- [dkp update bootstrap](#) (see page 1737) - Update bootstrap cluster
- [dkp update bootstrap credentials aws](#) (see page 1735) - Update AWS credentials in the cluster and restart CAPA controllers
- [dkp update bootstrap credentials azure](#) (see page 1736) - Update Azure credentials in the cluster and restart CAPZ controllers
- [dkp update bootstrap credentials gcp](#) (see page 1737) - Update GCP credentials in the cluster and restart CAPG controllers
- [dkp update bootstrap credentials vsphere](#) (see page 1737) - Update VSphere credentials in the cluster and restart CAPV controllers

11.2.21.5.4.4 dkp update bootstrap credentials aws

Update AWS credentials in the cluster and restart CAPA controllers

```
dkp update bootstrap credentials aws [flags]
```

Options

```

    --context string      The name of the kubeconfig context to use
  -h, --help            help for aws
    --kubeconfig string  Path to the kubeconfig for the management cluster. If
unspecified, default discovery rules apply.
    --print-only         Print the credentials and exit the function. Without
modifying cluster

```

Options inherited from parent commands

```

-v, --verbose int    Output verbosity

```

SEE ALSO

- [dkp update bootstrap credentials](#) (see page 1738) - Update credentials in the cluster

11.2.21.5.4.5 dkp update bootstrap credentials azure

Update Azure credentials in the cluster and restart CAPZ controllers

```

dkp update bootstrap credentials azure [flags]

```

Options

```

    --context string      The name of the kubeconfig context to use
  -h, --help            help for azure
    --kubeconfig string  Path to the kubeconfig for the management cluster. If
unspecified, default discovery rules apply.
    --print-only         Print the credentials and exit the function. Without
modifying cluster

```

Options inherited from parent commands

```

-v, --verbose int    Output verbosity

```

SEE ALSO

- [dkp update bootstrap credentials](#) (see page 1738) - Update credentials in the cluster

11.2.21.5.4.6 dkp update bootstrap credentials gcp

Update GCP credentials in the cluster and restart CAPG controllers

```
dkp update bootstrap credentials gcp [flags]
```

Options

```

--context string      The name of the kubeconfig context to use
-h, --help           help for gcp
--kubeconfig string  Path to the kubeconfig for the management cluster. If
unspecified, default discovery rules apply.
--print-only        Print the credentials and exit the function. Without
modifying cluster

```

Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

SEE ALSO

- [dkp update bootstrap credentials](#) (see page 1738) - Update credentials in the cluster

11.2.21.5.4.7 dkp update bootstrap credentials vsphere

Update VSphere credentials in the cluster and restart CAPV controllers

```
dkp update bootstrap credentials vsphere [flags]
```

Options

```

--context string      The name of the kubeconfig context to use
-h, --help           help for vsphere
--kubeconfig string  Path to the kubeconfig for the management cluster. If
unspecified, default discovery rules apply.
--print-only        Print the credentials and exit the function. Without
modifying cluster

```

Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

SEE ALSO

- [dkp update bootstrap credentials](#) (see page 1738) - Update credentials in the cluster

11.2.22 dkp upgrade

Upgrade one of [addons, capi-components, catalogapp, cluster, kommander, workspace]

11.2.22.1 Options

```
-h, --help help for upgrade
```

11.2.22.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.22.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1742) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]
- [dkp upgrade capi-components](#) (see page 1740) - Upgrade the CAPI components in the cluster
- [dkp upgrade catalogapp](#) (see page 1740) - Upgrade a Catalog Application to a newer version
- [dkp upgrade kommander](#) (see page 1738) - Upgrade the Kommander version of the targeted cluster
- [dkp upgrade workspace](#) (see page 1741) - Upgrade all platform applications in the given workspace and its projects to the same version as platform applications running on the management cluster

11.2.22.4 dkp upgrade kommander

Upgrade the Kommander version of the targeted cluster

11.2.22.4.1 Synopsis

Upgrades all Kommander components and platform applications running on the targeted cluster. No attached clusters and applications running on them are affected by this action.

```
dkp upgrade kommander [flags]
```

11.2.22.4.2 Options

<code>--charts-bundle</code> stringArray	Path to charts-bundle to upload to chartmuseum, apart from parsing the kommander applications repository (default [])
<code>--config</code> string	Config file to use (default "/home/runner/.kommander/config")
<code>--context</code> string	The name of the kubeconfig context to use
<code>--core-app-timeout</code> duration	Timeout to wait for upgrade of each kommander core application (default 20m0s)
<code>--disable-appdeployments</code> strings	List of AppDeployments to be disabled during upgrade (default [])
<code>--disallow-charts-download</code>	make CLI rely solely on provided chart bundles and do not try to download charts from the Internet
<code>--gitea-kommander-repository-name</code> string	gitea kommander repository name (default "kommander")
<code>-h, --help</code>	help for kommander
<code>--kommander-applications-repository</code> string	git repository with application definitions (default "v2.6.2-rc.4")
<code>--kommander-charts-version</code> string	Kommander helm charts version to download. Default: download all available versions
<code>--kubeconfig</code> string	Path to the kubeconfig file to use for CLI requests.
<code>--platform-apps-timeout</code> duration	Timeout to wait for upgrade of the set of platform applications (default 30m0s)
<code>--request-timeout</code> string	The length of time to wait before giving up on a single server request. Non-zero values should contain a corresponding time unit (e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
<code>-v, --verbose</code> int	Output verbosity

11.2.22.4.3 SEE ALSO

- [dkp upgrade](#) (see page 1741) - Upgrade one of [addons, capi-components, catalogapp, cluster, kommander, workspace]

11.2.22.5 dkp upgrade capi-components

Upgrade the CAPI components in the cluster

```
dkp upgrade capi-components [flags]
```

11.2.22.5.1 Options

```

    --aws-service-endpoints string      Custom AWS service endpoints in a semi-colon
separated format: ${SigningRegion1}:${ServiceID1}=${URL},${ServiceID2}=${URL};${
{SigningRegion2}...
    -h, --help                          help for capi-components
    --http-proxy string                  HTTP proxy for CAPI controllers
    --https-proxy string                 HTTPS proxy for CAPI controllers
    --kubeconfig string                  Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
    --no-proxy strings                   No Proxy list for CAPI controllers (default
[])
    --timeout duration                  The length of time to wait before giving up.
Zero means wait forever (e.g. 300s, 30m, 3h). (default 10m0s)
    -v, --verbose int                  Output verbosity
    --wait                               If true, wait for operations to complete
before returning. (default true)
    --with-aws-bootstrap-credentials    Set true to use AWS bootstrap credentials
from your environment. When false, the instance profile of the EC2 instance where the
CAPA controller is scheduled on will be used instead.
    --with-gcp-bootstrap-credentials    Set true to use GCP bootstrap credentials
from your environment. When false, the service account of the VM instance where the
CAPG controller is scheduled on will be used instead.

```

11.2.22.5.2 SEE ALSO

- [dkp upgrade](#) (see page 1741) - Upgrade one of [addons, capi-components, catalogapp, cluster, kommander, workspace]

11.2.22.6 dkp upgrade catalogapp

Upgrade a Catalog Application to a newer version

```
dkp upgrade catalogapp CATALOGAPP_NAME --to-version VERSION [--workspace WORKSPACE |
--project PROJECT] [flags]
```

11.2.22.6.1 Options

```

--config string                Config file to use (default "/home/
runner/.kommander/config")
--context string              The name of the kubeconfig context to use
--core-app-timeout duration   Timeout to wait for upgrade of each
kommander core application (default 20m0s)
--disable-appdeployments strings List of AppDeployments to be disabled during
upgrade (default [])
-h, --help                    help for catalogapp
--kubeconfig string          Path to the kubeconfig file to use for CLI
requests.
--platform-apps-timeout duration Timeout to wait for upgrade of the set of
platform applications (default 30m0s)
--project string             Name of the Project to upgrade the Catalog
App in
--request-timeout string     The length of time to wait before giving up
on a single server request. Non-zero values should contain a corresponding time unit
(e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
--to-version string          Version the Catalog App should be upgraded
to
-v, --verbose int           Output verbosity
-w, --workspace string       Name of the Workspace to upgrade the Catalog
App in

```

11.2.22.6.2 SEE ALSO

- [dkp upgrade](#) (see page 1741) - Upgrade one of [addons, capi-components, catalogapp, cluster, kommander, workspace]

11.2.22.7 dkp upgrade workspace

Upgrade all platform applications in the given workspace and its projects to the same version as platform applications running on the management cluster

```
dkp upgrade workspace WORKSPACE_NAME [--dry-run] [flags]
```

11.2.22.7.1 Options

```

--config string                Config file to use (default "/home/
runner/.kommander/config")
--context string              The name of the kubeconfig context to use

```



```

--core-app-timeout duration      Timeout to wait for upgrade of each
kommander core application (default 20m0s)
--disable-appdeployments strings List of AppDeployments to be disabled during
upgrade (default [])
--dry-run                        Do not upgrade, just list the AppDeployments
that would be upgraded
-h, --help                      help for workspace
--kubeconfig string            Path to the kubeconfig file to use for CLI
requests.
--platform-apps-timeout duration Timeout to wait for upgrade of the set of
platform applications (default 30m0s)
--request-timeout string       The length of time to wait before giving up
on a single server request. Non-zero values should contain a corresponding time unit
(e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
-v, --verbose int           Output verbosity

```

11.2.22.7.2 SEE ALSO

- [dkp upgrade](#) (see page 1741) - Upgrade one of [addons, capi-components, catalogapp, cluster, kommander, workspace]

11.2.22.8 dkp upgrade addons

Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.22.8.1 Options

```

-h, --help          help for addons
-v, --verbose int Output verbosity

```

11.2.22.8.2 SEE ALSO

- [dkp upgrade](#) (see page 1741) - Upgrade one of [addons, capi-components, catalogapp, cluster, kommander, workspace]
- [dkp upgrade addons aws](#) (see page 1743) - Upgrade the core Addons in a AWS cluster
- [dkp upgrade addons azure](#) (see page 1744) - Upgrade the core Addons in a Azure cluster
- [dkp upgrade addons eks](#) (see page 1746) - Upgrade the core Addons in a EKS cluster
- [dkp upgrade addons gcp](#) (see page 1748) - Upgrade the core Addons in a GCP cluster
- [dkp upgrade addons preprovisioned](#) (see page 1745) - Upgrade the core Addons in a Preprovisioned cluster
- [dkp upgrade addons vsphere](#) (see page 1747) - Upgrade the core Addons in a vSphere cluster

11.2.22.8.3 dkp upgrade addons aws

Upgrade the core Addons in a AWS cluster

```
dkp upgrade addons aws [flags]
```

11.2.22.8.3.1 Options

```

--allow-missing-template-keys  If true, ignore any errors in templates when a
field or map key is missing in the template. Only applies to goLang and jsonpath
output formats. (default true)
-c, --cluster-name name      Name used to prefix the cluster and all the
created resources.
--dry-run                    Only print the objects that would be created,
without creating them.
-h, --help                  help for aws
--kubeconfig string         Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
-n, --namespace string      If present, the namespace scope for this CLI
request. (default "default")
-o, --output string         Output format. One of: (json, yaml, name, go-
template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-directory string   Used with --output=json|yaml. The directory
where to output resources to files. The directory must already exist.
--show-managed-fields       If true, keep the managedFields when printing
objects in JSON or YAML format.
--template string           Template string or path to template file to use
when -o=go-template, -o=go-template-file. The template format is goLang templates
[http://goLang.org/pkg/text/template/#pkg-overview].

```

11.2.22.8.3.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.22.8.3.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1745) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.22.8.4 dkp upgrade addons azure

Upgrade the core Addons in a Azure cluster

```
dkp upgrade addons azure [flags]
```

11.2.22.8.4.1 Options

```

--allow-missing-template-keys  If true, ignore any errors in templates when a
field or map key is missing in the template. Only applies to goLang and jsonpath
output formats. (default true)
-c, --cluster-name name      Name used to prefix the cluster and all the
created resources.
--dry-run                    Only print the objects that would be created,
without creating them.
-h, --help                  help for azure
--kubeconfig string         Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
-n, --namespace string      If present, the namespace scope for this CLI
request. (default "default")
-o, --output string         Output format. One of: (json, yaml, name, go-
template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-directory string   Used with --output=json|yaml. The directory
where to output resources to files. The directory must already exist.
--show-managed-fields       If true, keep the managedFields when printing
objects in JSON or YAML format.
--template string           Template string or path to template file to use
when -o=go-template, -o=go-template-file. The template format is goLang templates
[http://goLang.org/pkg/text/template/#pkg-overview].

```

11.2.22.8.4.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.22.8.4.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1745) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.22.8.5 dkp upgrade addons preprovisioned

Upgrade the core Addons in a Preprovisioned cluster

```
dkp upgrade addons preprovisioned [flags]
```

11.2.22.8.5.1 Options

```

--allow-missing-template-keys  If true, ignore any errors in templates when a
field or map key is missing in the template. Only applies to goLang and jsonpath
output formats. (default true)
-c, --cluster-name name      Name used to prefix the cluster and all the
created resources.
--dry-run                    Only print the objects that would be created,
without creating them.
-h, --help                  help for preprovisioned
--kubeconfig string         Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
-n, --namespace string      If present, the namespace scope for this CLI
request. (default "default")
-o, --output string         Output format. One of: (json, yaml, name, go-
template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-directory string   Used with --output=json|yaml. The directory
where to output resources to files. The directory must already exist.
--show-managed-fields       If true, keep the managedFields when printing
objects in JSON or YAML format.
--template string           Template string or path to template file to use
when -o=go-template, -o=go-template-file. The template format is goLang templates
[http://goLang.org/pkg/text/template/#pkg-overview].

```

11.2.22.8.5.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.22.8.5.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1745) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.22.8.6 dkp upgrade addons eks

Upgrade the core Addons in a EKS cluster

```
dkp upgrade addons eks [flags]
```

11.2.22.8.6.1 Options

```

--allow-missing-template-keys  If true, ignore any errors in templates when a
field or map key is missing in the template. Only applies to goLang and jsonpath
output formats. (default true)
-c, --cluster-name name       Name used to prefix the cluster and all the
created resources.
--dry-run                     Only print the objects that would be created,
without creating them.
-h, --help                   help for eks
--kubeconfig string          Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
-n, --namespace string       If present, the namespace scope for this CLI
request. (default "default")
-o, --output string          Output format. One of: (json, yaml, name, go-
template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-directory string    Used with --output=json|yaml. The directory
where to output resources to files. The directory must already exist.
--show-managed-fields        If true, keep the managedFields when printing
objects in JSON or YAML format.
--template string            Template string or path to template file to use
when -o=go-template, -o=go-template-file. The template format is goLang templates
[http://goLang.org/pkg/text/template/#pkg-overview].

```

11.2.22.8.6.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.22.8.6.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1745) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.22.8.7 dkp upgrade addons vsphere

Upgrade the core Addons in a vSphere cluster

```
dkp upgrade addons vsphere [flags]
```

11.2.22.8.7.1 Options

```

--allow-missing-template-keys  If true, ignore any errors in templates when a
field or map key is missing in the template. Only applies to goLang and jsonpath
output formats. (default true)
-c, --cluster-name name       Name used to prefix the cluster and all the
created resources.
--dry-run                     Only print the objects that would be created,
without creating them.
-h, --help                   help for vsphere
--kubeconfig string          Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
-n, --namespace string       If present, the namespace scope for this CLI
request. (default "default")
-o, --output string          Output format. One of: (json, yaml, name, go-
template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-directory string    Used with --output=json|yaml. The directory
where to output resources to files. The directory must already exist.
--show-managed-fields        If true, keep the managedFields when printing
objects in JSON or YAML format.
--template string            Template string or path to template file to use
when -o=go-template, -o=go-template-file. The template format is goLang templates
[http://goLang.org/pkg/text/template/#pkg-overview].

```

11.2.22.8.7.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.22.8.7.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1745) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.22.8.8 dkp upgrade addons gcp

Upgrade the core Addons in a GCP cluster

```
dkp upgrade addons gcp [flags]
```

11.2.22.8.8.1 Options

```

--allow-missing-template-keys  If true, ignore any errors in templates when a
field or map key is missing in the template. Only applies to goLang and jsonpath
output formats. (default true)
-c, --cluster-name name       Name used to prefix the cluster and all the
created resources.
--dry-run                     Only print the objects that would be created,
without creating them.
-h, --help                   help for gcp
--kubeconfig string          Path to the kubeconfig for the management
cluster. If unspecified, default discovery rules apply.
-n, --namespace string       If present, the namespace scope for this CLI
request. (default "default")
-o, --output string          Output format. One of: (json, yaml, name, go-
template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json,
jsonpath-file).
--output-directory string     Used with --output=json|yaml. The directory
where to output resources to files. The directory must already exist.
--show-managed-fields        If true, keep the managedFields when printing
objects in JSON or YAML format.
--template string            Template string or path to template file to use
when -o=go-template, -o=go-template-file. The template format is goLang templates
[http://goLang.org/pkg/text/template/#pkg-overview].

```

11.2.22.8.8.2 Options inherited from parent commands

```
-v, --verbose int  Output verbosity
```

11.2.22.8.8.3 SEE ALSO

- [dkp upgrade addons \(see page 1745\)](#) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

11.2.23 dkp version

Print version information

11.2.23.1 Synopsis

Print version information

```
dkp version [flags]
```

11.2.23.2 Options

```
-h, --help           help for version
--long              If true, print additional version information.
-o, --output string One of 'yaml' or 'json'.
```

11.2.23.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.24 dkp upload

Collect and upload telemetry information

11.2.24.1 Options

```
-h, --help           help for upload
-v, --verbose int Output verbosity
```

11.2.24.2 SEE ALSO

- [dkp upload telemetry](#) (see page 1750) - Upload telemetry information to a destination

11.2.24.3 dkp upload telemetry

Upload telemetry information to a destination

```
dkp upload telemetry [flags]
```

11.2.24.3.1 Options

```

  --as string                Username to impersonate for the operation.
User could be a regular user or a service account in a namespace.
  --as-group stringArray    Group to impersonate for the operation, this
flag can be repeated to specify multiple groups. (default [])
  --as-uid string           UID to impersonate for the operation.
  --bundle-profile string   available profiles are: [telemetry,
diagnostics] (default "telemetry")
  --cache-dir string        Default cache directory (default "/home/
runner/.kube/cache")
  --certificate-authority string Path to a cert file for the certificate
authority
  --client-certificate string Path to a client certificate file for TLS
  --client-key string       Path to a client key file for TLS
  --cluster string          The name of the kubeconfig cluster to use
  --context string         The name of the kubeconfig context to use
  --disable-compression    If true, opt-out of response compression for
all requests to the server
  -h, --help               help for telemetry
  --insecure-skip-tls-verify If true, the server's certificate will not be
checked for validity. This will make your HTTPS connections insecure
  --kubeconfig string       Path to the kubeconfig file to use for CLI
requests.
  -n, --namespace string   If present, the namespace scope for this CLI
request
  --request-timeout string  The length of time to wait before giving up on
a single server request. Non-zero values should contain a corresponding time unit
(e.g. 1s, 2m, 3h). A value of zero means don't timeout requests.
  -s, --server string       The address and port of the Kubernetes API
server
  --tls-server-name string  Server name to use for server certificate
validation. If it is not provided, the hostname used to contact the server is used
  --token string            Bearer token for authentication to the API
server
  --user string             The name of the kubeconfig user to use

```

11.2.24.3.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

11.2.24.3.3 SEE ALSO

- [dkp upload \(see page 1752\)](#) - Collect and upload telemetry information

12 Release Notes

View release-specific information for DKP

Download DKP

- ☐ You must be a registered user and logged on to the support portal to download this product. New customers must contact their sales representative or sales@d2iq.com¹³⁸¹ before attempting to download or install DKP.

Release Notes for specific releases:

- [DKP 2.6.0 Release Notes](#) (see page 1752)
- [DKP 2.6.1 Release Notes](#) (see page 1785)
- [DKP 2.6.2 Release Notes](#) (see page 1810)

12.1 DKP 2.6.0 Release Notes

DKP® version 2.6 was released on August 14, 2023.

Download DKP

- ☐ You must be a registered user and logged on to the support portal to download this product. New customers must contact their sales representative or sales@d2iq.com¹³⁸² before attempting to download or install DKP.

12.1.1 Release Summary

Welcome to D2iQ Kubernetes Platform (DKP) 2.6.0! This release focuses on... Below, you will find information about the following:

- [DKP 2.6.0 Features and Enhancements](#) (see page 1753)
- [DKP 2.6.0 Supported Kubernetes Versions](#) (see page 1758)
- [DKP 2.6.0 Kubernetes Major Updates and Deprecations](#) (see page 1759)
- [DKP 2.6.0 Components and Applications](#) (see page 1760)
- [DKP 2.6.0 Known Issues and Limitations](#) (see page 1769)

¹³⁸¹ <mailto:sales@d2iq.com>

¹³⁸² <mailto:sales@d2iq.com>

- [DKP 2.6.0 Customer Incidents](#) (see page 1780)
- [DKP 2.6.0 Deprecations](#) (see page 1784)

12.1.2 Additional resources

- For more information about working with native Kubernetes, see the [Kubernetes documentation](#)¹³⁸³.
- For a full list of attributed 3rd party software, see <http://d2iq.com/legal/3rd>.

12.1.3 DKP 2.6.0 Features and Enhancements

The following improvements are included in this release.

12.1.3.1 Podman Support

Podman is now supported by DKP! If you choose to use Podman, it works with Linux.

- Version 4.0 of [Podman](#)¹³⁸⁴ or higher for Linux.
- Host requirements found here: <https://kind.sigs.k8s.io/docs/user/rootless/#host-requirements>
- `Set`¹³⁸⁵ `alias docker=podman` - In DKP documentation we use `Docker` command line syntax in some procedures. If you are using Podman, we suggest you set `alias docker=podman`.

12.1.3.2 Kubernetes 1.26.6 Support

Kubernetes 1.26.6 enables you to benefit from the latest features and security fixes in upstream Kubernetes. This release of Kubernetes comes with approximately 40 enhancements that you can benefit from like [Service internal tracking policy](#)¹³⁸⁶, [Kubelet credential provider](#)¹³⁸⁷, among others.

To read more about major features in this release, visit [this page](#)¹³⁸⁸ and <https://kubernetes.io/blog/2022/12/09/kubernetes-v1-26-release/>.

12.1.3.3 Telemetry

With the telemetry function, you can send environment data to D2iQ. By gaining deeper insights into how the product is used, D2iQ can continuously improve the product, as well as the upgrade, support, and customer experiences. You can send data bundles for support assistance, cluster health monitoring, product improvement, or billing purposes.

For instructions on how to send data and more information about the collected data, available bundles, and use cases, see [Telemetry](#) (see page 1031).

¹³⁸³ <https://kubernetes.io/docs/home/>

¹³⁸⁴ <https://podman.io/getting-started/installation>

¹³⁸⁵ <https://docs.podman.io/en/latest/index.html>

¹³⁸⁶ <https://github.com/kubernetes/enhancements/issues/2086>

¹³⁸⁷ <https://github.com/kubernetes/enhancements/issues/2133>


¹³⁸⁸ <https://github.com/kubernetes/sig-release/blob/master/releases/release-1.26/README.md>

12.1.3.4 DKP AI Navigator

Experimental

We're thrilled to unveil our new DKP AI Navigator feature. Engineered with the power of AI and trained meticulously on our product documentation and Support knowledge base, this chatbot ensures real-time, interactive communication to answer a wide range of user queries, spanning basic instructions to complex functionalities. The chatbot is offered for Internet-connected environments regardless of DKP license type.

Use of the DKP AI Navigator requires that you are logged in to a licensed copy of DKP 2.6. You can access the DKP AI Navigator application from the icon in the lower right corner of the DKP user interface. Use of the chatbot is pursuant to the following agreement:

 This DKP AI Navigator is a tool through which I can obtain information or assistance for the products or services offered by D2iQ, Inc. Use of this chatbot is governed by the terms of the [privacy policy](#)¹³⁸⁹ and I agree not to share any personal, sensitive, or confidential information through this chatbot.

12.1.3.5 High Availability of DKP Workloads

DKP allows you to set priority classes for DKP workloads and applications. This ensures that core DKP components stay running in any situation.

For more information, see [Setting Priority Classes in DKP Applications](#) (see page 650).

12.1.3.6 Auto-attach CAPI Clusters

When you create a [Managed Cluster](#) (see page 82) with the DKP CLI, it attaches automatically to the [Management Cluster](#) (see page 82) - saving you time and allowing you to set up your DKP environments more efficiently.

12.1.3.7 AWS ECR Supported

We added support for AWS ECR registry. This feature enables more secure deployments with DKP on AWS by allowing our customers to host the container images necessary for deployment in Amazon's Elastic Container Registry (ECR). The ability to [modify instance metadata](#)¹³⁹⁰ is available as well through DKP using the [IAM Policies in the CloudFormation stack](#) (see page 1060). You can utilize AWS ECR as a registry mirror in non-air-gapped environments for faster retrieval of artifacts. For specific information on using ECR, refer to [Registry Mirror Tools](#) (see page 1536), which contains helpful information regarding mirrors.

¹³⁸⁹ <https://d2iq.com/privacy-policy>

¹³⁹⁰ https://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_ModifyInstanceMetadataOptions.html

12.1.3.8 Customizable Banners in the Dashboard

DKP allows you to create custom banners to display in the header, footer, or both. This feature is available for [Essential](#) (see page 622) and [Enterprise](#) (see page 622) installations at the cluster level. Choose a background color for your banners, or craft custom messaging to display throughout the DKP UI. Government organizations can leverage this feature to create classification banners, for example.

For Enterprise customers, an advanced option is available to upload a logo for added brand recognition. This feature is particularly useful for Managed Service Providers (MSPs) hosting multiple customer clusters. By using custom banners, you can effectively communicate your brand across DKP UI. Find more information on this feature through the links provided in this description.

12.1.3.9 Support for Provisioning in vCloud Director

Beta

VMWare vCloud Director is supported for our Managed Service Providers (MSP) to use with DKP. This allows our MSPs to setup virtual data centers from their vSphere environment and grant access to their tenant users. By providing a private cloud-like experience in data centers, VCD enables MSPs to deliver the benefits of a multi-tenant cloud environment while leveraging their existing on-prem infrastructure and vSphere virtualization layer. Integrating DKP with VCD streamlines the process of creating Kubernetes clusters on VCD infrastructure, expanding the adoption of DKP for customers offering Kubernetes-as-a-Service. This integration enables organizations to achieve greater agility and scalability in managing their cloud resources, leading to increased innovation and competitiveness in the marketplace.



This feature is currently in [Beta release](#) (see page 1853) for DKP 2.6.0.

12.1.3.10 Observability and Monitoring of Network-restricted Clusters at Edge

Previous versions of DKP include the capability of attaching network-restricted clusters to your DKP environment [using a tunnel](#) (see page 775). However, not all cluster monitoring options were available from the Management cluster, out of the box. DKP includes an option to enable a proxied access through the Management cluster, which allows you to access monitoring and logging dashboards and perform `kubectl` operations on [network-restricted clusters](#) (see page 84) without being in the same network. This feature is particularly valuable for edge use cases where clusters at the edge are often behind firewalls.

To enable the proxied access to a Managed or Attached cluster, see [Proxied Access to Network-Restricted Clusters](#) (see page 817).

12.1.3.11 Provision VMware vSphere Clusters from the DKP User Interface

This release offers a new way to provision [managed](#) (see page 81) DKP clusters for VMware vSphere - from the DKP user interface. After fulfilling the prerequisites listed in [vSphere: VMware Prerequisites](#) (see page 440), follow the instructions to [Create a vSphere Infrastructure Provider](#) (see page 622), then go on to the steps

[Create a Managed Cluster on vSphere from the DKP UI](#) (see page 750) to create a DKP-managed vSphere cluster.

12.1.3.12 Cluster Autoscaler can Manage Multiple Clusters

Enterprise

Gov Advanced

In previous DKP versions, the Cluster-Autoscaler did not work on managed clusters. Starting with this release, an instance of Cluster-Autoscaler is deployed on the management cluster for each managed cluster. During an upgrade, the upgrade procedure on the management cluster will create the necessary deployments.

12.1.3.13 Advanced Scenario-Based Kommander Installations

In earlier versions of the documentation, the installation options for the Kommander component of DKP were located in different sections. The DKP documentation now includes a comprehensive [guideline on the decisions you must make](#) (see page 1397) before installing the Kommander component, along with [clear installation workflows by environment type](#) (see page 1399). Each installation workflow also includes license, customization, and application information relevant for the installation. This makes it easier to find all available environment and customization options in one place.

You can also find other customization options, for example, how to set up custom domains, custom certificates, load balancers, HTTP, and so on in [Kommander Customizations](#) (see page 1425).

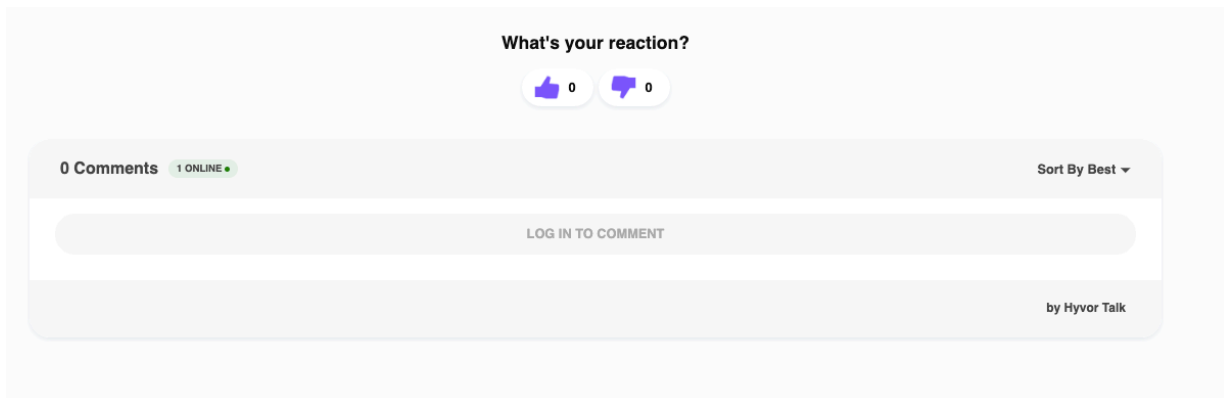
12.1.3.14 Scenario Based DKP Upgrades

Previously, we separated upgrades by the Kommander and Konvoy components as well as various air-gapped environment steps inserted along the way, making the order of steps difficult to determine. In the current release, we have divided them into Enterprise and Essential Upgrade for all components. The prerequisites for both components are centrally-located at the beginning of the process, along with some helpful compatibility tables. This new scenario install process should make upgrading easier to follow in the proper order.


12.1.3.15 Updated Links to Documentation in the DKP UI

The documentation landing page is now called the D2iQ Help Center, with access to multiple product's information. In conjunction with this change, the *View Documentation* links in the DKP UI are renamed to *Visit Help Center*.

12.1.3.16 Documentation Page Rating and Feedback



You can now rate and provide D2iQ feedback on every page in the documentation. This feature runs on a 3rd-party plug-in that requires registration to leave comments, which helps reduce spam comments on our website and also track the comments so that we can communicate directly back with you, our customer!

 Your data is NOT shared and is only used to verify your identity.

To create an account, select **Log In to Comment** from the bottom of any page, and then select **Sign Up** to get started.

We welcome your feedback; positive, negative, anything - let us know what you think about our documentation pages.

12.1.3.17 Konvoy Image Builder (KIB) [Release v2.5.0](https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.5.0)¹³⁹¹

The new release of Konvoy Image Builder features:

- Upgrade Kubernetes version to v1.26.6
- Upgrade CRI tools to 1.26.0

12.1.3.18 Display Application Dependencies in the UI

With this enhancement, you no longer need to manually ensure all dependencies are enabled before enabling an application. The UI automatically validates this, flagging all dependencies that are not enabled for an application.

¹³⁹¹ <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.5.0>

12.1.3.19 DKP Insights Release Notes

This release provides extended etcd Insight alert details, and enables users to use DKP Insights as a platform application. It also maintains compatibility and support for other packages used in Insights.

12.1.3.19.1 Extended Root Cause Analysis and Solutions for etcd Insights

etcd serves as the central data store for a Kubernetes cluster. Monitoring etcd allows you to ensure the health, availability and stability of your cluster. By monitoring key metrics, DKP Insights can detect and flag any issues or anomalies that can impact the cluster's operation.

This version of DKP Insights comes with extended alerts for [etcd](#)¹³⁹²-grouped issues. The alert details now include information on the root cause (RCA) and suggestions for solutions. The user can leverage this information to identify potential etcd issues and react to them proactively.

12.1.3.19.2 DKP Insights as a Platform Application

In previous versions, the GitRepository for Insights had to be added to your DKP environment to make the application available for deployment. This version of DKP includes DKP Insights as a platform application, making it ready to deploy using the DKP UI or CLI.

As a result, the application versions are now tied to specific DKP versions and upgrades *from* DKP 2.6.0 instances will include the upgrade of Insights.

12.1.3.19.3 Supported Kubernetes Versions

Insights supports the same [Kubernetes version as the DKP platform](#)¹³⁹³¹³⁹⁴.

12.1.4 DKP 2.6.0 Supported Kubernetes Versions

12.1.4.1 Kubernetes Version Supported

These are the supported versions:

12.1.4.1.1 Deploying Clusters Versions

The latest supported Kubernetes versions for DKP are listed below. In DKP 2.6.0, the Kubernetes version installed by default is 1.26.6. The newest and oldest Kubernetes versions used with DKP must be within one minor version.

- Latest supported Kubernetes version is at **1.26.6 for deploying clusters**
- Other Kubernetes versions are supported at **1.25.13 for deploying clusters in EKS**

¹³⁹² <https://etcd.io/>

¹³⁹³ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273350882/2.6+DKP+2.6.0+Features+and+Enhancements>

¹³⁹⁴ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273350882/2.6+DKP+2.6.0+Features+and+Enhancements>

Kubernetes 1.26.6 enables you to benefit from the latest features and security fixes in upstream Kubernetes. This release comes with approximately 40 enhancements that you can benefit from like [Service internal tracking policy](#)¹³⁹⁵, [Kubelet credential provider](#)¹³⁹⁶, among others. To read more about major features in this release, visit [GitHub](#)¹³⁹⁷ and <https://kubernetes.io/blog/2022/12/09/kubernetes-v1-26-release/>.

12.1.4.1.2 Attaching Clusters Versions

DKP 2.6.0 supports attaching clusters with the following Kubernetes versions:

Product	Compatible Kubernetes Versions
EKS	1.25.13
AKS	1.26.x
GKE	1.26.x



Attaching Kubernetes clusters with versions greater than n-1 is not recommended.



When attempting to attach a cluster with a lower Kubernetes version than the DKP default, you need to build and use an image with that lower version of Kubernetes. See [Konvoy Image Builder](#) (see page 1464) for documentation on building images. Failure to do this could result in an error.

12.1.5 DKP 2.6.0 Kubernetes Major Updates and Deprecations



Before upgrading, we **strongly recommend** verifying your current setup against the information on this page and reading more about Kubernetes' new features in [Kubernetes upgrade notes](#)¹³⁹⁸.

¹³⁹⁵ <https://github.com/kubernetes/enhancements/issues/2086>

¹³⁹⁶ <https://github.com/kubernetes/enhancements/issues/2133>

¹³⁹⁷ <https://github.com/kubernetes/sig-release/blob/master/releases/release-1.26/README.md>

¹³⁹⁸ <https://github.com/kubernetes/kubernetes/blob/master/CHANGELOG/CHANGELOG-1.26.md>

12.1.5.1 Deprecated API Services

With this version, Kubernetes stopped serving several API versions for Flow control resources and HorizontalPodAutoscaler services. For more information on the deprecated API versions for each of these services, refer to <https://kubernetes.io/docs/reference/using-api/deprecation-guide/#v1-26>.

12.1.5.2 Removed OCI Image Registry

DKP uses the upstream `registry.k8s.io` registry instead of the previous `k8s.gcr.io` registry for all new and upgraded clusters. The Kubernetes image is only available through the new image registry starting with 1.26.0.

Ensure you add `registry.k8s.io` to any firewall or proxy allowlists.

12.1.5.3 Migration to External Cloud-Controller-Managers

The deprecation of in-tree `cloud-controller-managers` that are part of `kube-controller-manager` has started. Starting with DKP 2.6 and Kubernetes 1.26.0, the in-tree `cloud-controller-manager` is disabled and migrated to the external `cloud-controller-manager`. If you had custom configuration for the in-tree `cloud-controller-manager`, you may need to add an equivalent configuration in the new deployments.

12.1.6 DKP 2.6.0 Components and Applications

The following are component and application versions for DKP.

12.1.6.1 Components

Component Name	Version
Cluster API Core (CAPI)	1.4.3-d2iq.0
Cluster API Amazon AWS Infrastructure Provider (CAPA)	2.2.0
Cluster API Google Cloud Infrastructure Provider (CAPG)	1.3.1
Cluster API Pre-provisioned Infrastructure Provider (CAPPP)	0.16.0

Component Name	Version
Cluster API vSphere Infrastructure Provider (CAPV)	1.6.1
Cluster API Azure Infrastructure Provider (CAPZ)	1.8.5
Cluster API for VMware Cloud Director (CAPVCD)	1.1.0-d2iq.0 ¹³⁹⁹
Konvoy Image Builder (KIB)	2.5.3 ¹⁴⁰⁰
containerd	1.6.17
etcd	3.5.6
Calico	3.26.1 ¹⁴⁰¹
Cluster Autoscaler	1.26.3
CSI_VERSION	Default Storage Providers in DKP (see page 91) See this page for versions and driver table.
Metal LB	0.12.1
Node Feature Discovery	0.13.2

1399 <https://github.com/vmware/cluster-api-provider-cloud-director>

1400 <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.5.3>

1401 <https://docs.tigera.io/archive>

12.1.6.2 Applications

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Centralized Grafana	centralized-grafana	46.8.0	<ul style="list-style-type: none"> • chart: 46.8.0¹⁴⁰² • prometheus-operator: 0.65.2 • grafana: 9.5.5 	Link ¹⁴⁰³	Link ¹⁴⁰⁴
Centralized Kubecost	centralized-kubecost	0.35.1	<ul style="list-style-type: none"> • chart: 0.35.1¹⁴⁰⁵ • kubecost: 1.104.0 	Link ¹⁴⁰⁶	Link ¹⁴⁰⁷
Cert Manager	cert-manager	1.12.2	<ul style="list-style-type: none"> • chart: 1.12.2¹⁴⁰⁸ • cert-manager: 1.12.2 	Link ¹⁴⁰⁹	Link ¹⁴¹⁰
Chartmuseum	chartmuseum	3.10.0	<ul style="list-style-type: none"> • chart: 3.10.0¹⁴¹¹ • chartmuseum: 3.10.0 	Link ¹⁴¹²	Link ¹⁴¹³
Dex	dex	2.12.3	<ul style="list-style-type: none"> • chart: 2.12.3¹⁴¹⁴ • dex: 2.35.2 	Link ¹⁴¹⁵	Link ¹⁴¹⁶

¹⁴⁰² <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.0>

¹⁴⁰³ <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.0?modal=values>

¹⁴⁰⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/centralized-grafana/46.8.0/defaults/cm.yaml>

¹⁴⁰⁵ <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1>

¹⁴⁰⁶ <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1?modal=values>

¹⁴⁰⁷ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/centralized-kubecost/0.35.1/defaults/cm.yaml>

¹⁴⁰⁸ <https://artifacthub.io/packages/helm/cert-manager/cert-manager/1.12.2>

¹⁴⁰⁹ <https://artifacthub.io/packages/helm/cert-manager/cert-manager/1.12.2?modal=values>

¹⁴¹⁰ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/cert-manager/1.12.2/defaults/cm.yaml>

¹⁴¹¹ <https://artifacthub.io/packages/helm/chartmuseum/chartmuseum/3.10.0>

¹⁴¹² <https://artifacthub.io/packages/helm/chartmuseum/chartmuseum/3.10.0?modal=values>

¹⁴¹³ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/chartmuseum/3.10.0/defaults/cm.yaml>

¹⁴¹⁴ <https://artifacthub.io/packages/helm/mesosphere-stable/dex/2.12.3>

¹⁴¹⁵ <https://artifacthub.io/packages/helm/mesosphere-stable/dex/2.12.3?modal=values>

¹⁴¹⁶ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/dex/2.12.3/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Dex K8s Authenticator	dex-k8s-authenticator	1.2.15	<ul style="list-style-type: none"> chart: 1.2.15¹⁴¹⁷ dex-k8s-authenticator: 1.2.4 	Link ¹⁴¹⁸	Link ¹⁴¹⁹
Dkp Insights	dkp-insights	0.5.0	<ul style="list-style-type: none"> chart: 0.5.0 dkp-insights: 0.5.0 	Link (see page 1763)	Link ¹⁴²⁰
DKP Insights Management	dkp-insights-management	0.5.0	<ul style="list-style-type: none"> chart: 0.5.0 dkp-insights-management: 0.5.0 	N/A	Link ¹⁴²¹
External DNS	external-dns	6.2.0.4	<ul style="list-style-type: none"> chart: 6.20.4¹⁴²² external-dns: 0.13.4 	Link ¹⁴²³	Link ¹⁴²⁴
Fluent Bit	fluent-bit	0.34.1	<ul style="list-style-type: none"> chart: 0.34.1¹⁴²⁵ fluent-bit: 2.1.6 	Link ¹⁴²⁶	Link ¹⁴²⁷
Gatekeeper	gatekeeper	3.12.0	<ul style="list-style-type: none"> chart: 3.12.0¹⁴²⁸ gatekeeper: 3.12.0 	Link ¹⁴²⁹	Link ¹⁴³⁰

¹⁴¹⁷ <https://artifacthub.io/packages/helm/mesosphere/dex-k8s-authenticator/1.2.15>

¹⁴¹⁸ <https://artifacthub.io/packages/helm/mesosphere/dex-k8s-authenticator/1.2.15?modal=values>

¹⁴¹⁹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/dex-k8s-authenticator/1.2.15/defaults/cm.yaml>

¹⁴²⁰ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/dkp-insights/0.5.0/defaults/cm.yaml>

¹⁴²¹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/dkp-insights-management/0.5.0/defaults/cm.yaml>

¹⁴²² <https://artifacthub.io/packages/helm/bitnami/external-dns/6.20.4>

¹⁴²³ <https://artifacthub.io/packages/helm/bitnami/external-dns/6.20.4?modal=values>

¹⁴²⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/external-dns/6.20.4/defaults/cm.yaml>

¹⁴²⁵ <https://artifacthub.io/packages/helm/fluent/fluent-bit/0.34.1>

¹⁴²⁶ <https://artifacthub.io/packages/helm/fluent/fluent-bit/0.34.1?modal=values>

¹⁴²⁷ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/fluent-bit/0.34.1/defaults/cm.yaml>

¹⁴²⁸ <https://artifacthub.io/packages/helm/gatekeeper/gatekeeper/3.12.0>

¹⁴²⁹ <https://artifacthub.io/packages/helm/gatekeeper/gatekeeper/3.12.0?modal=values>

¹⁴³⁰ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/gatekeeper/3.12.0/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Gitea	gitea	8.2.0	<ul style="list-style-type: none"> • chart: 8.2.0¹⁴³¹ • gitea: 1.19.2 	Link ¹⁴³²	Link ¹⁴³³
Grafana Logging	grafana-logging	6.57.4	<ul style="list-style-type: none"> • chart: 6.57.4¹⁴³⁴ • grafana: 9.5.5 	Link ¹⁴³⁵	Link ¹⁴³⁶
Grafana Loki	grafana-loki	0.69.16	<ul style="list-style-type: none"> • chart: 0.69.16¹⁴³⁷ • loki: 2.8.2 	Link ¹⁴³⁸	Link ¹⁴³⁹
Istio	istio	1.17.2	<ul style="list-style-type: none"> • chart: 1.17.2¹⁴⁴⁰ • istio: 1.17.2 	Link ¹⁴⁴¹	Link ¹⁴⁴²
Jaeger	jaeger	2.46.2	<ul style="list-style-type: none"> • chart: 2.46.2¹⁴⁴³ • jaeger: 1.46.0 	Link ¹⁴⁴⁴	Link ¹⁴⁴⁵
Karma	karma	2.0.2	<ul style="list-style-type: none"> • chart: 2.0.2¹⁴⁴⁶ • karma: 0.70 	Link ¹⁴⁴⁷	Link ¹⁴⁴⁸

¹⁴³¹ <https://artifacthub.io/packages/helm/gitea/gitea/8.2.0>

¹⁴³² <https://artifacthub.io/packages/helm/gitea/gitea/8.2.0?modal=values>

¹⁴³³ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/gitea/8.2.0/defaults/cm.yaml>

¹⁴³⁴ <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4>

¹⁴³⁵ <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4?modal=values>

¹⁴³⁶ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/grafana-logging/6.57.4/defaults/cm.yaml>

¹⁴³⁷ <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16>

¹⁴³⁸ <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16?modal=values>

¹⁴³⁹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/grafana-loki/0.69.16/defaults/cm.yaml>

¹⁴⁴⁰ <https://artifacthub.io/packages/helm/mesosphere/istio/1.17.2>

¹⁴⁴¹ <https://artifacthub.io/packages/helm/mesosphere/istio/1.17.2?modal=values>

¹⁴⁴² <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/istio/1.17.2/defaults/cm.yaml>

¹⁴⁴³ <https://artifacthub.io/packages/helm/jaegertracing/jaeger-operator/2.46.2>

¹⁴⁴⁴ <https://artifacthub.io/packages/helm/jaegertracing/jaeger-operator/2.46.2?modal=values>

¹⁴⁴⁵ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/jaeger/2.46.2/defaults/cm.yaml>

¹⁴⁴⁶ <https://artifacthub.io/packages/helm/mesosphere-stable/karma/2.0.2>

¹⁴⁴⁷ <https://artifacthub.io/packages/helm/mesosphere-stable/karma/2.0.2?modal=values>

¹⁴⁴⁸ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/karma/2.0.2/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Kiali	kiali	1.70.0	<ul style="list-style-type: none"> • chart: 1.70.0¹⁴⁴⁹ • kiali: 1.70.0 	Link ¹⁴⁵⁰	Link ¹⁴⁵¹
Knative	knative	1.10.4	<ul style="list-style-type: none"> • chart: 1.10.4¹⁴⁵² • knative: 1.10.0 	Link ¹⁴⁵³	Link ¹⁴⁵⁴
Flux	kommander-flux	0.41.2	<ul style="list-style-type: none"> • chart: N/A • flux: 0.41.2 	N/A	N/A
Kommander Ui	kommander-ui	9.61.2	<ul style="list-style-type: none"> • chart: 9.61.2 • kommander-ui: 9.61.2 	Link (see page 1763)	Link ¹⁴⁵⁵
Kube OIDC Proxy	kube-oidc-proxy	0.3.2	<ul style="list-style-type: none"> • chart: 0.3.2¹⁴⁵⁶ • kube-oidc-proxy: 0.3.0 	Link ¹⁴⁵⁷	Link ¹⁴⁵⁸
Kube Prometheus Stack	kube-prometheus-stack	46.8.0	<ul style="list-style-type: none"> • chart: 46.8.0¹⁴⁵⁹ • prometheus-operator: 0.65.2 • grafana: 9.5.5 • prometheus: 2.44.0 • prometheus-alertmanager: 0.25.0 	Link ¹⁴⁶⁰	Link ¹⁴⁶¹

1449 <https://artifacthub.io/packages/helm/kiali/kiali-operator/1.70.0>

1450 <https://artifacthub.io/packages/helm/kiali/kiali-operator/1.70.0?modal=values>

1451 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/kiali/1.70.0/defaults/cm.yaml>

1452 <https://artifacthub.io/packages/helm/mesosphere/knative/1.10.4>

1453 <https://artifacthub.io/packages/helm/mesosphere/knative/1.10.4?modal=values>

1454 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/knative/1.10.4/defaults/cm.yaml>

1455 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/kommander-ui/9.61.2/defaults/cm.yaml>

1456 <https://artifacthub.io/packages/helm/mesosphere/kube-oidc-proxy/0.3.2>

1457 <https://artifacthub.io/packages/helm/mesosphere/kube-oidc-proxy/0.3.2?modal=values>

1458 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/kube-oidc-proxy/0.3.2/defaults/cm.yaml>

1459 <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.0>

1460 <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.0?modal=values>

1461 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/kube-prometheus-stack/46.8.0/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Kubecost	kubecost	0.35.1	<ul style="list-style-type: none"> chart: 0.35.1¹⁴⁶² kubecost: 1.104.0 	Link ¹⁴⁶³	Link ¹⁴⁶⁴
Kubefed	kubefed	0.10.3	<ul style="list-style-type: none"> chart: 0.10.3¹⁴⁶⁵ kubefed: 0.10.3 	Link ¹⁴⁶⁶	Link ¹⁴⁶⁷
Kubernetes Dashboard	kubernetes-dashboard	6.0.8	<ul style="list-style-type: none"> chart: 6.0.8¹⁴⁶⁸ kubernetes-dashboard: 2.7.0 	Link ¹⁴⁶⁹	Link ¹⁴⁷⁰
Kubetunnel	kubetunnel	0.0.27	<ul style="list-style-type: none"> chart: 0.0.27 kubetunnel: 0.0.27 	N/A	Link ¹⁴⁷¹
Logging Operator	logging-operator	4.2.3	<ul style="list-style-type: none"> chart: 4.2.3¹⁴⁷² logging-operator: 4.2.2 logging-operator-logging: 4.2.3 	Link ¹⁴⁷³	Link ¹⁴⁷⁴

¹⁴⁶² <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1>

¹⁴⁶³ <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1?modal=values>

¹⁴⁶⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/kubecost/0.35.1/defaults/cm.yaml>

¹⁴⁶⁵ <https://artifacthub.io/packages/helm/kubefed/kubefed/0.10.3>

¹⁴⁶⁶ <https://artifacthub.io/packages/helm/kubefed/kubefed/0.10.3?modal=values>

¹⁴⁶⁷ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/kubefed/0.10.3/defaults/cm.yaml>

¹⁴⁶⁸ <https://artifacthub.io/packages/helm/k8s-dashboard/kubernetes-dashboard/6.0.8>

¹⁴⁶⁹ <https://artifacthub.io/packages/helm/k8s-dashboard/kubernetes-dashboard/6.0.8?modal=values>

¹⁴⁷⁰ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/kubernetes-dashboard/6.0.8/defaults/cm.yaml>

¹⁴⁷¹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/kubetunnel/0.0.27/defaults/cm.yaml>

¹⁴⁷² <https://artifacthub.io/packages/helm/banzaicloud-stable/logging-operator/4.2.3>

¹⁴⁷³ <https://artifacthub.io/packages/helm/banzaicloud-stable/logging-operator/4.2.3?modal=values>

¹⁴⁷⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/logging-operator/4.2.3/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
NFS Server Provisioner	nfs-server-provisioner	0.6.1	<ul style="list-style-type: none"> chart: 0.6.1¹⁴⁷⁵ nfs-server-provisioner: 2.3.0 	Link ¹⁴⁷⁶	Link ¹⁴⁷⁷
NVIDIA GPU Operator	nvidia-gpu-operator	23.3.2	<ul style="list-style-type: none"> chart: 23.3.2¹⁴⁷⁸ nvidia-gpu-operator: 23.3.2 	Link ¹⁴⁷⁹	Link ¹⁴⁸⁰
Grafana (project)	project-grafana-logging	6.5 7.4	<ul style="list-style-type: none"> chart: 6.57.4¹⁴⁸¹ grafana: 9.5.5 	Link ¹⁴⁸²	Link ¹⁴⁸³
Grafana Loki (project)	project-grafana-loki	0.6 9.1 6	<ul style="list-style-type: none"> chart: 0.69.16¹⁴⁸⁴ loki: 2.8.2 	Link ¹⁴⁸⁵	Link ¹⁴⁸⁶
Prometheus Adapter	prometheus-adapter	4.2.1	<ul style="list-style-type: none"> chart: 4.2.0¹⁴⁸⁷ prometheus-adapter: 0.10.0 	Link ¹⁴⁸⁸	Link ¹⁴⁸⁹

¹⁴⁷⁵ <https://artifacthub.io/packages/helm/mesosphere/nfs-server-provisioner/0.6.1>

¹⁴⁷⁶ <https://artifacthub.io/packages/helm/mesosphere/nfs-server-provisioner/0.6.1?modal=values>

¹⁴⁷⁷ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/nfs-server-provisioner/0.6.1/defaults/cm.yaml>

¹⁴⁷⁸ <https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/getting-started.html#chart-customization-options>

¹⁴⁷⁹ <https://raw.githubusercontent.com/NVIDIA/gpu-operator/v23.3.2/v%25s/deployments/gpu-operator/values.yaml>

¹⁴⁸⁰ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/nvidia-gpu-operator/23.3.2/defaults/cm.yaml>

¹⁴⁸¹ <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4>

¹⁴⁸² <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4?modal=values>

¹⁴⁸³ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/project-grafana-logging/6.57.4/defaults/cm.yaml>

¹⁴⁸⁴ <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16>

¹⁴⁸⁵ <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16?modal=values>

¹⁴⁸⁶ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/project-grafana-loki/0.69.16/defaults/cm.yaml>

¹⁴⁸⁷ <https://artifacthub.io/packages/helm/prometheus-community/prometheus-adapter/4.2.0>

¹⁴⁸⁸ <https://artifacthub.io/packages/helm/prometheus-community/prometheus-adapter/4.2.0?modal=values>

¹⁴⁸⁹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/prometheus-adapter/4.2.1/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Reloader	reloader	1.0.29	<ul style="list-style-type: none"> chart: 1.0.29¹⁴⁹⁰ reloader: 1.0.29 	Link ¹⁴⁹¹	Link ¹⁴⁹²
Rook Ceph	rook-ceph	1.11.9	<ul style="list-style-type: none"> chart: 1.11.9¹⁴⁹³ rook-ceph: 1.11.9 	Link ¹⁴⁹⁴	Link ¹⁴⁹⁵
Rook Ceph Cluster	rook-ceph-cluster	1.11.9	<ul style="list-style-type: none"> chart: 1.11.9¹⁴⁹⁶ rook-ceph: 1.11.9 rook-ceph-cluster: 17.2.6 	Link ¹⁴⁹⁷	Link ¹⁴⁹⁸
Thanos	thanos	12.8.3	<ul style="list-style-type: none"> chart: 12.8.3¹⁴⁹⁹ thanos: 0.31.0 	Link ¹⁵⁰⁰	Link ¹⁵⁰¹
Traefik	traefik	23.1.0	<ul style="list-style-type: none"> chart: 23.1.0¹⁵⁰² traefik: 2.10.1 	Link ¹⁵⁰³	Link ¹⁵⁰⁴

1490 <https://github.com/stakater/Reloader/tree/v1.0.29/README.md#helm-charts>

1491 <https://artifacthub.io/packages/helm/stakater/reloader/1.0.29?modal=values>

1492 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/reloader/1.0.29/defaults/cm.yaml>

1493 <https://github.com/rook/rook/tree/vv1.11.9/Documentation/Helm-Charts/operator-chart.md>

1494 <https://raw.githubusercontent.com/rook/rook/v1.11.9/v%25s/deploy/charts/rook-ceph/values.yaml>

1495 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/rook-ceph/1.11.9/defaults/cm.yaml>

1496 <https://github.com/rook/rook/tree/vv1.11.9/Documentation/Helm-Charts/ceph-cluster-chart.md>

1497 <https://raw.githubusercontent.com/rook/rook/v1.11.9/v%25s/deploy/charts/rook-ceph-cluster/values.yaml>

1498 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/rook-ceph-cluster/1.11.9/defaults/cm.yaml>

1499 <https://artifacthub.io/packages/helm/bitnami/thanos/12.8.3>

1500 <https://artifacthub.io/packages/helm/bitnami/thanos/12.8.3?modal=values>

1501 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/thanos/12.8.3/defaults/cm.yaml>

1502 <https://artifacthub.io/packages/helm/traefik/traefik/23.1.0>

1503 <https://artifacthub.io/packages/helm/traefik/traefik/23.1.0?modal=values>

1504 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/traefik/23.1.0/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Traefik ForwardAuth	traefik-forward-auth	0.3.9	<ul style="list-style-type: none"> • chart: 0.3.9¹⁵⁰⁵ • traefik-forward-auth: 3.1.0 	Link ¹⁵⁰⁶	Link ¹⁵⁰⁷
Velero	velero	4.1.1	<ul style="list-style-type: none"> • chart: 4.1.1¹⁵⁰⁸ • velero: 1.11.0 	Link ¹⁵⁰⁹	Link ¹⁵¹⁰

12.1.7 DKP 2.6.0 Known Issues and Limitations

The following items are known issues with this release.

12.1.7.1 AWS `additionalTags` cannot contain spaces

Due to an upstream bug in the `cluster-api-provider-aws` component, it is not possible to specify tags with spaces in their name in the `additionalTags` section of an `AWSCluster`. If you have any tags like this during an upgrade of the `capi-components`, you may receive a validation error, and will need to remove any such tags. This issue will be corrected in a future DKP release.

12.1.7.2 Increased Time to Attach Several Clusters

When [attaching clusters](#) (see page 755) in a batch, the wait time for applications to deploy and become ready might take longer compared to previous releases. This issue is actively being addressed and is scheduled to be resolved in an upcoming DKP patch release. Please do not interrupt the attachment process.

12.1.7.3 Use Static Credentials to Provision an Azure Cluster

Only static credentials can be used when [provisioning an Azure cluster](#) (see page 1166).

¹⁵⁰⁵ <https://artifacthub.io/packages/helm/mesosphere/traefik-forward-auth/0.3.9>

¹⁵⁰⁶ <https://artifacthub.io/packages/helm/mesosphere/traefik-forward-auth/0.3.9?modal=values>

¹⁵⁰⁷ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/traefik-forward-auth/0.3.9/defaults/cm.yaml>

¹⁵⁰⁸ <https://artifacthub.io/packages/helm/vmware-tanzu/velero/4.1.1>

¹⁵⁰⁹ <https://artifacthub.io/packages/helm/vmware-tanzu/velero/4.1.1?modal=values>

¹⁵¹⁰ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.0/services/velero/4.1.1/defaults/cm.yaml>

12.1.7.4 Containerd File Limit Issue

In this version of DKP, we introduced containerd 1.6.17. The systemd unit for containerd 1.6.17 provided upstream removes all file number limits (`LimitNOFILE=infinity`). In our testing, we found that removing these limits broke some IO-sensitive applications such as Rook Ceph and HAProxy. Because of this, the KIB version included in this release sets the `LimitNOFILE` value in the containerd systemd unit to the value (`1048576`), which was used in previous containerd 1.4.13 version releases.

12.1.7.5 Intermittent Error Status when Creating EKS Clusters in the UI

When provisioning an EKS cluster through the UI, you might receive a brief error state because the EKS cluster can sporadically lose connectivity with the management cluster, which results in the following symptoms:

- The UI shows the cluster is in an error state.
- The kubeconfig generated and retrieved from Kommander ceases to work.
- Applications created on the management cluster may not be immediately federated to managed EKS clusters.

After a few moments, the error self-resolves, without any action on your part. A new kubeconfig generated and retrieved from Kommander then works properly, and the UI shows that it is working again. In the meantime, you can continue to use the UI to work on the cluster such as deploy applications, create projects, and add roles.

12.1.7.6 Limitations to Disk Resizing in vSphere

The DKP CLI flags `--control-plane-disk-size` and `--worker-disk-size` are unable to resize the root file system of VMs, which are created using OS images. The flags work by resizing the primary disk of the VM. When the VM boots, the root file system expands to fill the disk, but that expansion does not work for some file systems. For example, for file systems contained in an LVM Logical Volume. To ensure your root file system has the size you expect, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273384734/2.6%2BCreate%2Ba%2BvSphere%2BBase%2BOS%2BImage#Disk-Size> (see page 0).

12.1.7.7 Error Status in Grafana Logging Dashboard with EKS Clusters

Currently, it is not possible to use FluentBit to collect [Admin-level logs](#) (see page 905) on a managed EKS cluster.

If you have these logs enabled, the following message appears when you access the Kubernetes Audit Dashboard in the Grafana Logging Dashboard:

```
Cannot read properties of undefined (reading '0')
```

12.1.7.8 Rook Ceph Install Error

An issue might emerge when installing `rook-ceph` on vSphere clusters using RHEL operating systems.

This issue occurs during initial installation of `rook-ceph`, causing the object store used by Velero and Grafana Loki, to be unavailable. If the installation of Kommander component of DKP is unsuccessful due to `rook-ceph` failing, you might need to apply the following workaround:

1. Run this command to check if the cluster is affected by this issue.

```
kubectl describe CephObjectStores dkp-object-store -n kommander
```

2. If this output appears, the workaround needs to be applied so continue with the next step. If you do not see this output, you can stop at this step.

```
Name:          dkp-object-store
Namespace:    kommander
...
Warning ReconcileFailed      7m55s (x19 over 52m)
rook-ceph-object-controller failed to reconcile CephObjectStore
"kommander/dkp-object-store". failed to create object store deployments:
failed
to configure multisite for object store: failed create ceph multisite for
object-store ["dkp-object-store"]: failed to commit config changes after
creating multisite config for CephObjectStore "kommander/dkp-object-store":
failed to commit RGW configuration period changes%(EXTRA []string=[]):
signal: interrupt
```

3. Kubectl exec into the `rook-ceph-tools` pod.

```
export WORKSPACE_NAMESPACE=<workspace namespace>
CEPH_TOOLS_POD=$(kubectl get pods -l app=rook-ceph-tools -n $
{WORKSPACE_NAMESPACE} -o name)
kubectl exec -it -n ${WORKSPACE_NAMESPACE} $CEPH_TOOLS_POD bash
```

4. Run these commands to set `dkp-object-store` as the default zonegroup.

i **NOTE:** The `period update` command may take a few minutes to complete

```
radosgw-admin zonegroup default --rgw-zonegroup=dkp-object-store
radosgw-admin period update --commit
```

5. Next, restart the `rook-ceph-operator` deployment for the `CephobjectStore` to be reconciled.

```
kubectl rollout restart deploy -n${WORKSPACE_NAMESPACE} rook-ceph-operator
```

6. After running the commands above, the `CephObjectStore` should be `Connected` once the `rook-ceph` operator reconciles the object (this may take some time).

```
kubectl wait CephObjectStore --for=jsonpath='{.status.phase}'=Connected dkp-object-store -n ${WORKSPACE_NAMESPACE} --timeout 10m
```

12.1.7.9 Kommander Installation Configuration File Changes

In this release, the `kommander-ui`, which provides the DKP Dashboard, is now deployed in the same manner as the other Platform Applications. Therefore, the contents of the default Kommander installation configuration file, which is the file produced by the `dkp install kommander --init` command, have changed. If you installed previous versions of DKP using a customized Kommander installation configuration file, we recommend:

1. Generate a new template for this release.
2. Reapply your customizations rather than reusing a file created by older DKP versions.

12.1.7.10 FluentD Logging Operator Error

An error may occur with the FluentD Operator not appearing when installing Kommander on a GCP Environment.

This is due to a limitation of the GCP API. As a result, you need to rename the Fluentd buffer name because the default name exceeds 63 characters and disable buffer volume metrics for FluentD in order for it to operate properly.

Follow these procedures to resolve this issue for the Management Cluster, and Managed or Attached Clusters.

12.1.7.10.1 Management Cluster on GCP with the Logging Stack enabled

1. Create the `kommander` namespace:

```
kubectl create namespace kommander
```

2. Create the logging-operator-logging-overrides ConfigMap

```
cat <<EOF | kubectl apply -n kommander -f -
apiVersion: v1
kind: ConfigMap
```

```

metadata:
  name: logging-operator-logging-overrides
data:
  values.yaml: |
    fluentd:
      # disable buffer metrics
      bufferVolumeMetrics: null
      bufferStorageVolume:
        pvc:
          source:
            # update this name to make the PV <= 63 chars long
            # logging-operator-logging-buf-logging-operator-logging-fluentd-0
            claimName: buf
          spec:
            accessModes:
              - ReadWriteOnce
            resources:
              requests:
                storage: 10Gi
            volumeMode: Filesystem
EOF

```

3. Continue installing Kommander as normal.

12.1.7.10.2 Managed/Attached Cluster in GCP with Logging Stack enabled

1. Set the `WORKSPACE_NAMESPACE` environment variable to the name of your workspace's namespace:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```

2. Create the logging-operator-logging-overrides ConfigMap on the managed or attached cluster prior to enabling the logging-operator application:

```

cat <<EOF | kubectl apply -n ${WORKSPACE_NAMESPACE} -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: logging-operator-logging-overrides
data:
  values.yaml: |
    fluentd:
      # disable buffer metrics
      bufferVolumeMetrics: null
      bufferStorageVolume:
        pvc:
          source:

```



```

# update this name to make the PV <= 63 chars long
# logging-operator-logging-buf-logging-operator-logging-fluentd-0
claimName: buf
spec:
  accessModes:
    - ReadWriteOnce
  resources:
    requests:
      storage: 10Gi
    volumeMode: Filesystem
EOF

```

3. Proceed with installing the logging-operator as normal.

12.1.7.11 Generate a Token or Enable the Konvoy Credentials Plugin for Proxied Clusters

If you have attached a [network-restricted cluster](#) (see page 81) and enabled a [proxied access](#) (see page 817) to make its resources available through the Management cluster, the “Generate Token” and “Konvoy credentials plugin instructions” options required for user authentication in the DKP UI require further steps.

Expand the set of instructions depending on the authentication configuration you want to configure:

12.1.7.11.1 A. Generate a Token

Generate Token

- 12.1.7.11.1.1 I. Execute the commands established in the “Generate Token” page.



You can find the “Generate Token” page in the DKP UI by selecting your username in the top right corner of any page.

- 12.1.7.11.1.2 II. Prepare your Environment

1. Write down the name of your credentials. You can obtain this name by accessing the **DKP UI** > username in the top right corner > **Generate token** page. The credentials are displayed in the `kubectl config set-credentials` command, next to `set-credentials`.

In the following example, `jane_doe-`

```
az189abcd7ab34a2abc0ab845a67a7ea-0123456789.us-west-2.elb.amazonaws.com
```

would be the name of your credentials:

```
kubectl config set-credentials jane_doe-
az189abcd7ab34a2abc0ab845a67a7ea-0123456789.us-west-2.elb.amazonaws.com \
  --token=[...]
```

2. Set the `CONFIG_CREDENTIALS_NAME` environment variable to the credentials you obtained in the previous step:

```
CONFIG_CREDENTIALS_NAME=<CREDENTIALS_NAME>
```

3. Set `WORKSPACE_NAMESPACE` to the namespace of your network-restricted cluster:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```


4. Set `TUNNEL_PROXY_NAME` to the name of the TunnelProxy:

```
TUNNEL_PROXY_NAME=<...>
```

5. Set `DEFAULT_CONFIG_PATH` to the path file of the `kubeconfig` where you executed the “Generate Token” steps:

```
DEFAULT_CONFIG_PATH=<path_of_cluster's_kubeconfig>
```

6. With your `kubeconfig` set to the **Management cluster**, set the `TUNNEL_PROXY_DOMAIN` environment variable to the proxied access URL:

 See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.

```
TUNNEL_PROXY_DOMAIN=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n
$WORKSPACE_NAMESPACE -o template='{{ .status.clusterProxyDomain }}')
```

12.1.7.11.1.3 III. Reconfigure the Token

Select one of the following options according to the certificate you configured for the proxied access:


12.1.7.11.1.4 A. Clusters with ACME-supported Certificates for the TunnelProxy

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
${TUNNEL_PROXY_DOMAIN}/dkp/api-server --kubeconfig ${DEFAULT_CONFIG_PATH}
```

```
kubectl config set-context ${TUNNEL_PROXY_NAME} --cluster=${TUNNEL_PROXY_NAME} --
user=${CONFIG_CREDENTIALS_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
```

12.1.7.11.1.5 B. Clusters with Self-Signed Certificates for the TunnelProxy

1. With your `kubeconfig` set to the **Management cluster**, retrieve the `ca.crt` (not needed in case of ACME certificates):

 See [Provide Context for Commands with a kubeconfig File](#) (see page 87) for information on setting the kubeconfig to a cluster.

```
LABEL_KEY=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n $WORKSPACE_NAMESPACE
-o template='{{ .status.reverseProxyReleaseName }}')
SECRET_NAME=$(kubectl get ingress -n $WORKSPACE_NAMESPACE -l app.kubernetes.io/
instance=$LABEL_KEY -o jsonpath='{.items[0].spec.tls[0].secretName}')
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
kubectl get secrets $SECRET_NAME -n $WORKSPACE_NAMESPACE -o
template='{{ index .data "ca.crt" | base64decode }}' > ${HOME}/.kube/certs/$
{TUNNEL_PROXY_NAME}/ca.crt
```

2. Configure the `kubeconfig` access via the proxied access:

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
{TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/
certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --cluster=${TUNNEL_PROXY_NAME}
--user=${CONFIG_CREDENTIALS_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig $
{DEFAULT_CONFIG_PATH}
```

12.1.7.11.1.6 C. Clusters with Custom Certificates for the TunnelProxy

1. Copy the `ca.crt` used to sign your certificates to the `kubeconfig` directory:

```
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
cp ca.crt ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}/ca.crt
```

2. Configure the `kubeconfig` access via the proxied access:

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
{TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/
certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --cluster=${TUNNEL_PROXY_NAME}
--user=${CONFIG_CREDENTIALS_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig $
{DEFAULT_CONFIG_PATH}
```

After the commands complete, you can perform `kubectl` actions on the network-restricted cluster.

12.1.7.11.2 B. Konvoy credentials plugin instructions

Konvoy credentials plugin instructions

12.1.7.11.2.1 I. Execute the commands provided in the “Konvoy credentials plugin instructions” page > “Method 1” section.



You can find a link to the “Konvoy credentials plugin instructions” page in the DKP UI when logging in to your environment, or by adding `/token/plugin` to the DKP UI URL of your environment.

12.1.7.11.2.2 II. Prepare your Environment

1. Set `WORKSPACE_NAMESPACE` to the namespace of your network-restricted cluster:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```

2. Set `TUNNEL_PROXY_NAME` to the name of the TunnelProxy:

```
TUNNEL_PROXY_NAME=<...>
```

3. Set `DEFAULT_CONFIG_PATH` to the path of the kubeconfig file you downloaded from the dashboard:

```
DEFAULT_CONFIG_PATH=<path_of_cluster's_kubeconfig>
```

4. With your kubeconfig set to the **Management cluster**, set the `TUNNEL_PROXY_DOMAIN` environment variable to the proxied access URL:

i See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.

```
TUNNEL_PROXY_DOMAIN=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n
$WORKSPACE_NAMESPACE -o template='{{ .status.clusterProxyDomain }}')
```

12.1.7.11.2.3 III. Reconfigure the Token

Select one of the following options according to the certificate you configured for the proxied access:

12.1.7.11.2.4 A. Clusters with ACME-supported Certificates for the TunnelProxy

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
${TUNNEL_PROXY_DOMAIN}/dkp/api-server --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --user default-profile --kubeconfig $
${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
```

12.1.7.11.2.5 B. Clusters with Self-Signed Certificates for the TunnelProxy

1. With your kubeconfig set to the **Management cluster**, retrieve the `ca.crt` (not needed in case of ACME certificates):

i See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.

```
LABEL_KEY=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n $WORKSPACE_NAMESPACE
-o template='{{ .status.reverseProxyReleaseName }}')
SECRET_NAME=$(kubectl get ingress -n $WORKSPACE_NAMESPACE -l app.kubernetes.io/
instance=$LABEL_KEY -o jsonpath='{.items[0].spec.tls[0].secretName}')
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
kubectl get secrets $SECRET_NAME -n $WORKSPACE_NAMESPACE -o
template='{{ index .data "ca.crt" | base64decode }}' > ${HOME}/.kube/certs/$
{TUNNEL_PROXY_NAME}/ca.crt
```

2. Configure the `kubeconfig` access via the proxied access:

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
{TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/
certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --user default-profile --
kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig $
{DEFAULT_CONFIG_PATH}
```

12.1.7.11.2.6 C. Clusters with Custom Certificates for the TunnelProxy

1. Copy the `ca.crt` used to sign your certificates to the `kubeconfig` directory:

```
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
cp ca.crt ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}/ca.crt
```

2. Configure the `kubeconfig` access via the proxied access:

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
{TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/
certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --user default-profile --
kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig $
{DEFAULT_CONFIG_PATH}
```



“Method 2” is not available for network-restricted clusters.

After the commands complete, you can perform `kubectl` actions on the network-restricted cluster.

12.1.7.11.3 Known Common Vulnerabilities and Exposures (CVE)

Starting with DKP 2.6, Catalog apps are scanned for CVEs. However, only the CVEs for app versions that are compatible with the default Kubernetes version, currently 1.26.6, are mitigated. For more information about the known CVEs for compatible catalog apps, see [D2iQ Security Updates \(see page 1851\)](#).

12.1.7.12 Potential Logging Interruption during Upgrade

The configuration setting `enableRecreateWorkloadOnImmutableFieldChange` is enabled by default on the `logging-operator`. This means the operator automatically triggers a recreation of the resource if

any immutable fields are changed on the underlying fluent-bit/fluentd resources. Previously, any changes to immutable fields would fail silently and the failures would only be observable in the `logging-operator` logs. This was done to handle a breaking change introduced in the `logging-operator`, which added an additional selector label on the `fluent-bit` DaemonSet.

Because of this breaking change, during an upgrade to 2.6, `fluent-bit` is restarted, and some log data in the `fluent-bit` buffer might be lost if `fluentd` is unavailable while `fluent-bit` attempts to flush its buffer prior to the pods being terminated. To prevent issues related to data loss, we recommend you run `>1 fluentd pod`. Alternatively, Fluent Bit can also be configured to use a `hostPath` volume to store the buffer information, so it can be picked up again when Fluent Bit restarts. For more information about how to change the host path, see [Fluent Bit log collector](#)¹⁵¹¹.

12.1.7.13 Gatekeeper Uninstallation Error

If you choose to disable Gatekeeper, you can run into an error where the app is still present on your clusters.

For pre-existing attached or managed clusters with Gatekeeper installed when they don't want it to be, you need to manually cleanup Gatekeeper on those clusters, post upgrade.

Follow these steps to manually remove Gatekeeper:

1. For every **Attached/Managed cluster** that has Gatekeeper installed that you have disabled it on otherwise (via appdeployment), run the following commands to remove it:

NOTE: This needs to be run on the attached/managed cluster using the correct kubeconfig. Set your namespace of your **attached** cluster's workspace on your attached cluster:

```
export WORKSPACE_NAMESPACE=<workspace-name>
```

2. Delete the kustomizations:

```
kubectl --kubeconfig <attached-cluster-kubeconfig> delete kustomizations -n $
{WORKSPACE_NAMESPACE} gatekeeper-constraint-templates gatekeeper-constraints
gatekeeper-release
```

12.1.8 DKP 2.6.0 Customer Incidents

The following resolved incidents are included in this release.

12.1.8.1 kubernetes/crictl Cannot Pull Images from Harbor Registry

D2iQ-93950

¹⁵¹¹ <https://kube-logging.dev/docs/logging-infrastructure/fluentbit/#hostpath-volumes-for-buffers-and-positions>

Attempting to pull images from a separate registry using `crictl` would fail due to an issue pertaining to the configuration of that registry. An updated configuration has been included with this release, resolving this issue.

12.1.8.2 Mis-configuration of OIDC Identity Provider from CLI or Kommander UI Crashes Dex and Locks User out of Kommander UI

D2IQ-95131

<Daniel L to add content>

12.1.8.3 Kommander Failed to Install on Microsoft Azure

D2IQ-96225

Kommander failed to install on cluster in Microsoft Azure the first time, but installed successfully the second time.

12.1.8.4 Failing Kommander Upgrade from DKP 2.2.3 to 2.3.2

D2IQ-96059

Kommander upgrade fails from 2.2.3 to 2.3.2 in a Flatcar environment, running on-premises.

12.1.8.5 Add Note about Deploying with a Lower Kubernetes Version

D2IQ-95873

Customer could not deploy a cluster using a lower Kubernetes version than the 2.4 default.

12.1.8.6 DKP Missing Required Images for kubetunnel in Air-gapped Environments

D2IQ-95830

When attempting to attach an air-gapped cluster with the “network restricted” configuration, the managed cluster remains stuck in a “Pending” status.

12.1.8.7 Missing registry_2.tar File in 2.5 Beta Air-gapped Bundle

D2IQ-95598

The air-gapped bundle is missing `registry_2.tar` in the 2.5 Beta air-gapped bundle.

12.1.8.8 Scaling Nodes for Pre-provisioned Clusters

D2IQ-95555

Documentation doesn't have a section on growing an existing pre-provisioned cluster.

12.1.8.9 Velero 'backup logs' Command Doesn't Work

D2IQ-95509

The `velero backup logs <backup-name>` command does not work in DKP 2.4 as a result of the change from Minio to Rook-Ceph.

12.1.8.10 Unable to Build Ubuntu KIB Images in GCP

D2IQ-95418

KIB does not build a GCP Ubuntu image and fails with an SSH error.

12.1.8.11 Unable to Deploy Kommander with HTTP_PROXY

D2IQ-95415

Kommander deployed with the mutations block added, but the `kommander-flux` components cannot pull any Helm charts including those for Gatekeeper/OPA, because the proxy environment variables are not set.

12.1.8.12 Traefik entry point TCP 9090 is tagged Velero-minio

D2IQ-95330

The entry point for Rook-ceph object storage is still tagged from DKP 2.3, which used MinIO.

12.1.8.13 Kommander Install on Small Environment Needs Reference to Platform Application Dependencies

D2IQ-95311

To enable an application, you need to refer to the Platform Application Dependencies to make sure that all of the needed dependencies are enabled.

12.1.8.14 Microsoft Azure Upgrade from DKP 2.3.0 to DKP 2.4.0 Breaks Volume Attachments for Pods

D2IQ-95191

The upgrade process is broken for Microsoft Azure clusters when moving from DKP 2.3.0 to DKP 2.4.0.

12.1.8.15 License Counts the Control Plane Cores after the 2.4 Upgrade

D2IQ-95188

The licensing core count includes cores associated with the Control Plane, whereas it did not in DKP 2.3.

12.1.8.16 EKS Cluster Creation Fails when AWS User Credentials do not have Full IAM Access

D2IQ-95137

An attempt to create an AWS cluster failed because the otherwise-properly credentialed account did not have full IAM access.

12.1.8.17 Increase Disk size for vSphere Cluster does not Increase ContainerD Space

D2IQ-95132

Use of the flag, `using flag --worker-disk-size 120` when creating vSphere cluster only increases disk size, but does not increase mount points. This results in a lack of additional storage available for containerD.

12.1.8.18 Erroneous Configuration of OIDC Identity Provider Prevents User from Accessing the UI

D2IQ-95131

When OIDC is configured incorrectly, the `dex` service crashes, locking the user out of the UI.

12.1.8.19 Calico Not Updated During DKP Upgrade from 2.1.x to 2.2.2 on Flatcar

D2IQ-94862

When upgrading a DKP cluster on Flatcar, after what seems to be a successful upgrade, the Calico service is not updated to a new version as expected and uses a previous image.

12.1.8.20 Attaching a DKP Essential Cluster via Kubetunnel not Supported

D2IQ-94859

In DKP versions 2.4.x and earlier, it is not possible to attach a DKP cluster with Kommander to another DKP cluster with Kommander. This was the case regardless of whether this connection was done using a `kubetunnel` for clusters with or without networking restrictions.

12.1.8.21 Instructions on Configuring a Custom Domain and Certificate After Installing Kommander are Missing

D2IQ-93560

Documentation on how to configure a manually-generated custom domain and certificate post Kommander installation is not available. DKP 2.5 now includes this use case: [Configure Custom Domains or Custom Certificates post Kommander Installation](#) (see page 864).

12.1.8.22 `dkp create cluster` Fails when the Kubeconfig file from a Previous Cluster is set

D2IQ-91934

Running `dkp create cluster --self-managed` for an additional DKP cluster results in an error message if you still have an active kubeconfig file from a previous DKP cluster.

12.1.8.23 Sizing Guidelines for Logging Stack and Rook Ceph Applications

D2IQ-90805

Guidelines with specifications for the size of `logging-stack` and `rook-ceph` applications are missing.

12.1.9 DKP 2.6.0 Deprecations

The following items are deprecated or removed from this version of DKP.

12.1.9.1 Deprecated `dkp push` Subcommand and Flag

To improve consistency, simplify, and streamline DKP CLI commands, this version of DKP deprecates the `dkp push image-bundle` command in favor of `dkp push bundle`.

Deprecated subcommand and flag	New subcommand and flag
<code>dkp push image-bundle --image-bundle</code>	<code>dkp push bundle --bundle</code>

You can still run the `dkp push image-bundle` with this version of DKP, but it will be removed in future DKP versions.

If you have scripts, GitHub Actions, or other cluster creation and upgrade automations, D2iQ recommends updating to the new subcommand and flag as soon as possible.

12.1.9.2 Catalog Applications Compatibility

From DKP 2.6, the following versions of Catalog applications are deprecated:

- kafka-operator-0.20.0
- kafka-operator-0.20.2

- [spark-operator-1.1.6](#)
- [zookeeper-operator-0.2.13](#)

If you plan on upgrading to DKP 2.6, ensure that you upgrade these applications to the latest compatible version.

For more information, see [Workspace DKP Catalog Applications - Compatibility](#) (see page 667).

12.1.9.3 Spark Operator to be Removed in DKP 2.7

Starting from DKP 2.7, D2iQ will discontinue support for all versions of Spark Operator in the DKP Catalog, because the upstream operator is no longer maintained. As a result, the Spark Operator will no longer be available for use in future DKP releases.

Spark Operator remains available in DKP 2.6 and previous versions.

12.2 DKP 2.6.1 Release Notes

DKP® version 2.6.1 was released on October 19, 2023.

[Download DKP](#)

- You must be a registered user and logged on to the support portal to download this product. New customers must contact their sales representative or sales@d2iq.com¹⁵¹² before attempting to download or install DKP.

12.2.1 Release Summary

Welcome to D2iQ Kubernetes Platform (DKP) 2.6.1! This release focuses on... Below, you will find information about the following:

- [DKP 2.6.1 Features and Enhancements](#) (see page 1786)
- [DKP 2.6.1 Supported Kubernetes Versions](#) (see page 1786)
- [DKP 2.6.1 Kubernetes Major Updates and Deprecations](#) (see page 1787)
- [DKP 2.6.1 Components and Applications](#) (see page 1788)
- [DKP 2.6.1 Known Issues and Limitations](#) (see page 1796)
- [DKP 2.6.1 Customer Incidents and Resolved Issues](#) (see page 1808)
- [DKP 2.6.1 Deprecations](#) (see page 1809)

¹⁵¹² <mailto:sales@d2iq.com>

12.2.2 Additional resources

- For more information about working with native Kubernetes, see the [Kubernetes documentation](#)¹⁵¹³.
- For a full list of attributed 3rd party software, see <http://d2iq.com/legal/3rd>.

12.2.3 DKP 2.6.1 Features and Enhancements

The following improvements are included in this release.

12.2.3.1 AI Navigator Optional Install

With this release, the AI Navigator is optional and can be disabled at install time or uninstalled. To disable AI Navigator, [modify the Kommander configuration file after generating it](#)¹⁵¹⁴, to disable the `ai-navigator-` app entry in the `apps:` section of the file.

Currently, the AI Navigator is not available for air-gapped environments. For more information on the AI Navigator, refer to the documentation section for that feature: [AI Navigator \(see page 1839\)](#)

12.2.4 DKP 2.6.1 Supported Kubernetes Versions

12.2.4.1 Kubernetes Version Supported

These are the supported versions:

12.2.4.1.1 Deploying Clusters Versions

The latest supported Kubernetes versions for DKP are listed below. In DKP 2.6.1, the Kubernetes version installed by default is 1.26.6. The newest and oldest Kubernetes versions used with DKP must be within one minor version.

- Latest supported Kubernetes version is at **1.26.6 for deploying clusters**
- Other Kubernetes versions are supported at **1.25.13 for deploying clusters in EKS**

Kubernetes 1.26.6 enables you to benefit from the latest features and security fixes in upstream Kubernetes. This release comes with approximately 40 enhancements that you can benefit from like [Service internal tracking policy](#)¹⁵¹⁵, [Kubelet credential provider](#)¹⁵¹⁶, among others. To read more about major features in this release, visit [GitHub](#)¹⁵¹⁷ and <https://kubernetes.io/blog/2022/12/09/kubernetes-v1-26-release/>.

¹⁵¹³ <https://kubernetes.io/docs/home/>

¹⁵¹⁴ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/233177089>

¹⁵¹⁵ <https://github.com/kubernetes/enhancements/issues/2086>

¹⁵¹⁶ <https://github.com/kubernetes/enhancements/issues/2133>

¹⁵¹⁷ <https://github.com/kubernetes/sig-release/blob/master/releases/release-1.26/README.md>

12.2.4.1.2 Attaching Clusters Versions

DKP 2.6.1 supports attaching clusters with the following Kubernetes versions:

Product	Compatible Kubernetes Versions
EKS	1.25.13
AKS	1.26.x
GKE	1.26.x



Attaching Kubernetes clusters with versions greater than n-1 is not recommended.



When attempting to attach a cluster with a lower Kubernetes version than the DKP default, you need to build and use an image with that lower version of Kubernetes. See [Konvoy Image Builder](#) (see page 1464) for documentation on building images. Failure to do this could result in an error.

12.2.5 DKP 2.6.1 Kubernetes Major Updates and Deprecations



Before upgrading, we **strongly recommend** verifying your current setup against the information on this page and reading more about Kubernetes' new features in [Kubernetes upgrade notes](#)¹⁵¹⁸.

12.2.5.1 Deprecated API Services

With this version, Kubernetes stopped serving several API versions for Flow control resources and HorizontalPodAutoscaler services. For more information on the deprecated API versions for each of these services, refer to <https://kubernetes.io/docs/reference/using-api/deprecation-guide/#v1-26>.

¹⁵¹⁸ <https://github.com/kubernetes/kubernetes/blob/master/CHANGELOG/CHANGELOG-1.26.md>

12.2.5.2 Removed OCI Image Registry

DKP uses the upstream `registry.k8s.io` registry instead of the previous `k8s.gcr.io` registry for all new and upgraded clusters. The Kubernetes image is only available through the new image registry starting with 1.26.0.

Ensure you add `registry.k8s.io` to any firewall or proxy allowlists.

12.2.5.3 Migration to External Cloud-Controller-Managers

The deprecation of in-tree `cloud-controller-managers` that are part of `kube-controller-manager` has started. Starting with DKP 2.6 and Kubernetes 1.26.0, the in-tree `cloud-controller-manager` is disabled and migrated to the external `cloud-controller-manager`. If you had custom configuration for the in-tree `cloud-controller-manager`, you may need to add an equivalent configuration in the new deployments.

12.2.6 DKP 2.6.1 Components and Applications

The following are component and application versions for DKP.

12.2.6.1 Components

Component Name	Version
Cluster API Core (CAPI)	1.4.3-d2iq.0
Cluster API Amazon AWS Infrastructure Provider (CAPA)	2.2.0
Cluster API Google Cloud Infrastructure Provider (CAPG)	1.3.1
Cluster API Pre-provisioned Infrastructure Provider (CAPPP)	0.16.2
Cluster API vSphere Infrastructure Provider (CAPV)	1.6.1
Cluster API Azure Infrastructure Provider (CAPZ)	1.8.5

Component Name	Version
Cluster API for VMware Cloud Director (CAPVCD)	1.1.0-d2iq.0 ¹⁵¹⁹
Konvoy Image Builder (KIB)	2.5.4 ¹⁵²⁰
containerd	1.6.17
etcd	3.5.6
Calico	3.26.1 ¹⁵²¹
Cluster Autoscaler	1.27.3
CSI_VERSION	Default Storage Providers in DKP (see page 91) See this page for versions and driver table.
Metal LB	0.12.1
Node Feature Discovery	0.13.2

12.2.6.2 Applications

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Centralized Grafana	centralized-grafana	46.8.0	<ul style="list-style-type: none"> chart: 46.8.0¹⁵²² prometheus-operator: 0.65.2 grafana: 9.5.5 	Link ¹⁵²³	Link ¹⁵²⁴

¹⁵¹⁹ <https://github.com/vmware/cluster-api-provider-cloud-director>

¹⁵²⁰ <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.5.4>

¹⁵²¹ https://github.com/projectcalico/calico/blob/v3.26.1/calico/_includes/release-notes/v3.26.1-release-notes.md

¹⁵²² <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.0>

¹⁵²³ <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.0?modal=values>

¹⁵²⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/centralized-grafana/46.8.0/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Centralized Kubecost	centralized-kubecost	0.3 5.1	<ul style="list-style-type: none"> chart: 0.35.1¹⁵²⁵ kubecost: 1.104.0 	Link ¹⁵²⁶	Link ¹⁵²⁷
Cert Manager	cert-manager	1.1 2.2	<ul style="list-style-type: none"> chart: 1.12.2¹⁵²⁸ cert-manager: 1.12.2 	Link ¹⁵²⁹	Link ¹⁵³⁰
Chartmuseum	chartmuseum	3.1 0.0	<ul style="list-style-type: none"> chart: 3.10.0¹⁵³¹ chartmuseum: 3.10.0 	Link ¹⁵³²	Link ¹⁵³³
Dex	dex	2.1 2.3	<ul style="list-style-type: none"> chart: 2.12.3¹⁵³⁴ dex: 2.35.2 	Link ¹⁵³⁵	Link ¹⁵³⁶
Dex K8s Authenticator	dex-k8s-authenticator	1.2. 15	<ul style="list-style-type: none"> chart: 1.2.15¹⁵³⁷ dex-k8s-authenticator: 1.2.4 	Link ¹⁵³⁸	Link ¹⁵³⁹
Dkp Insights	dkp-insights	0.5. 0	<ul style="list-style-type: none"> chart: 0.5.0 dkp-insights: 0.5.0 	Link (see page 1791)	Link ¹⁵⁴⁰

1525 <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1>

1526 <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1?modal=values>

1527 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/centralized-kubecost/0.35.1/defaults/cm.yaml>

1528 <https://artifacthub.io/packages/helm/cert-manager/cert-manager/1.12.2>

1529 <https://artifacthub.io/packages/helm/cert-manager/cert-manager/1.12.2?modal=values>

1530 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/cert-manager/1.12.2/defaults/cm.yaml>

1531 <https://artifacthub.io/packages/helm/chartmuseum/chartmuseum/3.10.0>

1532 <https://artifacthub.io/packages/helm/chartmuseum/chartmuseum/3.10.0?modal=values>

1533 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/chartmuseum/3.10.0/defaults/cm.yaml>

1534 <https://artifacthub.io/packages/helm/mesosphere-stable/dex/2.12.3>

1535 <https://artifacthub.io/packages/helm/mesosphere-stable/dex/2.12.3?modal=values>

1536 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/dex/2.12.3/defaults/cm.yaml>

1537 <https://artifacthub.io/packages/helm/mesosphere/dex-k8s-authenticator/1.2.15>

1538 <https://artifacthub.io/packages/helm/mesosphere/dex-k8s-authenticator/1.2.15?modal=values>

1539 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/dex-k8s-authenticator/1.2.15/defaults/cm.yaml>

1540 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/dkp-insights/0.5.0/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
DKP Insights Management	dkp-insights-management	0.5.0	<ul style="list-style-type: none"> chart: 0.5.0 dkp-insights-management: 0.5.0 	N/A	Link ¹⁵⁴¹
External DNS	external-dns	6.2.0.4	<ul style="list-style-type: none"> chart: 6.20.4¹⁵⁴² external-dns: 0.13.4 	Link ¹⁵⁴³	Link ¹⁵⁴⁴
Fluent Bit	fluent-bit	0.34.1	<ul style="list-style-type: none"> chart: 0.34.1¹⁵⁴⁵ fluent-bit: 2.1.6 	Link ¹⁵⁴⁶	Link ¹⁵⁴⁷
Gatekeeper	gatekeeper	3.12.0	<ul style="list-style-type: none"> chart: 3.12.0¹⁵⁴⁸ gatekeeper: 3.12.0 	Link ¹⁵⁴⁹	Link ¹⁵⁵⁰
Gitea	gitea	8.2.0	<ul style="list-style-type: none"> chart: 8.2.0¹⁵⁵¹ gitea: 1.19.2 	Link ¹⁵⁵²	Link ¹⁵⁵³
Grafana Logging	grafana-logging	6.57.4	<ul style="list-style-type: none"> chart: 6.57.4¹⁵⁵⁴ grafana: 9.5.5 	Link ¹⁵⁵⁵	Link ¹⁵⁵⁶

¹⁵⁴¹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/dkp-insights-management/0.5.0/defaults/cm.yaml>

¹⁵⁴² <https://artifacthub.io/packages/helm/bitnami/external-dns/6.20.4>

¹⁵⁴³ <https://artifacthub.io/packages/helm/bitnami/external-dns/6.20.4?modal=values>

¹⁵⁴⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/external-dns/6.20.4/defaults/cm.yaml>

¹⁵⁴⁵ <https://artifacthub.io/packages/helm/fluent/fluent-bit/0.34.1>

¹⁵⁴⁶ <https://artifacthub.io/packages/helm/fluent/fluent-bit/0.34.1?modal=values>

¹⁵⁴⁷ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/fluent-bit/0.34.1/defaults/cm.yaml>

¹⁵⁴⁸ <https://artifacthub.io/packages/helm/gatekeeper/gatekeeper/3.12.0>

¹⁵⁴⁹ <https://artifacthub.io/packages/helm/gatekeeper/gatekeeper/3.12.0?modal=values>

¹⁵⁵⁰ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/gatekeeper/3.12.0/defaults/cm.yaml>

¹⁵⁵¹ <https://artifacthub.io/packages/helm/gitea/gitea/8.2.0>

¹⁵⁵² <https://artifacthub.io/packages/helm/gitea/gitea/8.2.0?modal=values>

¹⁵⁵³ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/gitea/8.2.0/defaults/cm.yaml>

¹⁵⁵⁴ <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4>

¹⁵⁵⁵ <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4?modal=values>

¹⁵⁵⁶ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/grafana-logging/6.57.4/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Grafana Loki	grafana-loki	0.6 9.1 7	<ul style="list-style-type: none"> • chart: 0.69.16¹⁵⁵⁷ • loki: 2.8.2 	Link ¹⁵⁵⁸	Link ¹⁵⁵⁹
Istio	istio	1.1 7.2	<ul style="list-style-type: none"> • chart: 1.17.2¹⁵⁶⁰ • istio: 1.17.2 	Link ¹⁵⁶¹	Link ¹⁵⁶²
Jaeger	jaeger	2.4 6.2	<ul style="list-style-type: none"> • chart: 2.46.2¹⁵⁶³ • jaeger: 1.46.0 	Link ¹⁵⁶⁴	Link ¹⁵⁶⁵
Karma	karma	2.0. 2	<ul style="list-style-type: none"> • chart: 2.0.2¹⁵⁶⁶ • karma: 0.70 	Link ¹⁵⁶⁷	Link ¹⁵⁶⁸
Kiali	kiali	1.7 0.0	<ul style="list-style-type: none"> • chart: 1.70.0¹⁵⁶⁹ • kiali: 1.70.0 	Link ¹⁵⁷⁰	Link ¹⁵⁷¹
Knative	knative	1.1 0.4	<ul style="list-style-type: none"> • chart: 1.10.4¹⁵⁷² • knative: 1.10.0 	Link ¹⁵⁷³	Link ¹⁵⁷⁴

1557 <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16>

1558 <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16?modal=values>

1559 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/grafana-loki/0.69.17/defaults/cm.yaml>

1560 <https://artifacthub.io/packages/helm/mesosphere/istio/1.17.2>

1561 <https://artifacthub.io/packages/helm/mesosphere/istio/1.17.2?modal=values>

1562 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/istio/1.17.2/defaults/cm.yaml>

1563 <https://artifacthub.io/packages/helm/jaegertracing/jaeger-operator/2.46.2>

1564 <https://artifacthub.io/packages/helm/jaegertracing/jaeger-operator/2.46.2?modal=values>

1565 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/jaeger/2.46.2/defaults/cm.yaml>

1566 <https://artifacthub.io/packages/helm/mesosphere-stable/karma/2.0.2>

1567 <https://artifacthub.io/packages/helm/mesosphere-stable/karma/2.0.2?modal=values>

1568 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/karma/2.0.2/defaults/cm.yaml>

1569 <https://artifacthub.io/packages/helm/kiali/kiali-operator/1.70.0>

1570 <https://artifacthub.io/packages/helm/kiali/kiali-operator/1.70.0?modal=values>

1571 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/kiali/1.70.0/defaults/cm.yaml>

1572 <https://artifacthub.io/packages/helm/mesosphere/knative/1.10.4>

1573 <https://artifacthub.io/packages/helm/mesosphere/knative/1.10.4?modal=values>

1574 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/knative/1.10.4/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Flux	kommander-flux	0.4 1.2	<ul style="list-style-type: none"> chart: N/A flux: 0.41.2 	N/A	N/A
Kommander Ui	kommander-ui	9.6 4.4	<ul style="list-style-type: none"> chart: 9.64.4 kommander-ui: 9.64.4 	Link (see page 1791)	Link ¹⁵⁷⁵
Kube OIDC Proxy	kube-oidc-proxy	0.3. 2	<ul style="list-style-type: none"> chart: 0.3.2¹⁵⁷⁶ kube-oidc-proxy: 0.3.0 	Link ¹⁵⁷⁷	Link ¹⁵⁷⁸
Kube Prometheus Stack	kube-prometheus-stack	46. 8.1	<ul style="list-style-type: none"> chart: 46.8.1¹⁵⁷⁹ prometheus-operator: 0.65.2 grafana: 9.5.5 prometheus: 2.44.0 prometheus-alertmanager: 0.25.0 	Link ¹⁵⁸⁰	Link ¹⁵⁸¹
Kubecost	kubecost	0.3 5.1	<ul style="list-style-type: none"> chart: 0.35.1¹⁵⁸² kubecost: 1.104.0 	Link ¹⁵⁸³	Link ¹⁵⁸⁴
Kubefed	kubefed	0.1 0.3	<ul style="list-style-type: none"> chart: 0.10.3¹⁵⁸⁵ kubefed: 0.10.3 	Link ¹⁵⁸⁶	Link ¹⁵⁸⁷

¹⁵⁷⁵ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/kommander-ui/9.64.4/defaults/cm.yaml>

¹⁵⁷⁶ <https://artifacthub.io/packages/helm/mesosphere/kube-oidc-proxy/0.3.2>

¹⁵⁷⁷ <https://artifacthub.io/packages/helm/mesosphere/kube-oidc-proxy/0.3.2?modal=values>

¹⁵⁷⁸ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/kube-oidc-proxy/0.3.2/defaults/cm.yaml>

¹⁵⁷⁹ <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.1>

¹⁵⁸⁰ <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.1?modal=values>

¹⁵⁸¹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/kube-prometheus-stack/46.8.1/defaults/cm.yaml>

¹⁵⁸² <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1>

¹⁵⁸³ <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1?modal=values>

¹⁵⁸⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/kubecost/0.35.1/defaults/cm.yaml>

¹⁵⁸⁵ <https://artifacthub.io/packages/helm/kubefed/kubefed/0.10.3>

¹⁵⁸⁶ <https://artifacthub.io/packages/helm/kubefed/kubefed/0.10.3?modal=values>

¹⁵⁸⁷ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/kubefed/0.10.3/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Kubernetes Dashboard	kubernetes-dashboard	6.0.8	<ul style="list-style-type: none"> • chart: 6.0.8¹⁵⁸⁸ • kubernetes-dashboard: 2.7.0 	Link ¹⁵⁸⁹	Link ¹⁵⁹⁰
Kubetunnel	kubetunnel	0.0.29	<ul style="list-style-type: none"> • chart: 0.0.29 • kubetunnel: 0.0.29 	N/A	Link ¹⁵⁹¹
Logging Operator	logging-operator	4.2.3	<ul style="list-style-type: none"> • chart: 4.2.3¹⁵⁹² • logging-operator: 4.2.2 • logging-operator-logging: 4.2.3 	Link ¹⁵⁹³	Link ¹⁵⁹⁴
NFS Server Provisioner	nfs-server-provisioner	0.6.1	<ul style="list-style-type: none"> • chart: 0.6.1¹⁵⁹⁵ • nfs-server-provisioner: 2.3.0 	Link ¹⁵⁹⁶	Link ¹⁵⁹⁷
NVIDIA GPU Operator	nvidia-gpu-operator	23.3.2	<ul style="list-style-type: none"> • chart: 23.3.2¹⁵⁹⁸ • nvidia-gpu-operator: 23.3.2 	Link ¹⁵⁹⁹	Link ¹⁶⁰⁰

1588 <https://artifacthub.io/packages/helm/k8s-dashboard/kubernetes-dashboard/6.0.8>

1589 <https://artifacthub.io/packages/helm/k8s-dashboard/kubernetes-dashboard/6.0.8?modal=values>

1590 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/kubernetes-dashboard/6.0.8/defaults/cm.yaml>

1591 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/kubetunnel/0.0.29/defaults/cm.yaml>

1592 <https://artifacthub.io/packages/helm/banzaicloud-stable/logging-operator/4.2.3>

1593 <https://artifacthub.io/packages/helm/banzaicloud-stable/logging-operator/4.2.3?modal=values>

1594 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/logging-operator/4.2.3/defaults/cm.yaml>

1595 <https://artifacthub.io/packages/helm/mesosphere/nfs-server-provisioner/0.6.1>

1596 <https://artifacthub.io/packages/helm/mesosphere/nfs-server-provisioner/0.6.1?modal=values>

1597 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/nfs-server-provisioner/0.6.1/defaults/cm.yaml>

1598 <https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/getting-started.html#chart-customization-options>

1599 <https://raw.githubusercontent.com/NVIDIA/gpu-operator/v23.3.2/v%25s/deployments/gpu-operator/values.yaml>

1600 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/nvidia-gpu-operator/23.3.2/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Grafana (project)	project-grafana-logging	6.5 7.4	<ul style="list-style-type: none"> • chart: 6.57.4¹⁶⁰¹ • grafana: 9.5.5 	Link ¹⁶⁰²	Link ¹⁶⁰³
Grafana Loki (project)	project-grafana-loki	0.6 9.1 7	<ul style="list-style-type: none"> • chart: 0.69.16¹⁶⁰⁴ • loki: 2.8.2 	Link ¹⁶⁰⁵	Link ¹⁶⁰⁶
Prometheus Adapter	prometheus-adapter	4.2. 1	<ul style="list-style-type: none"> • chart: 4.2.0¹⁶⁰⁷ • prometheus-adapter: 0.10.0 	Link ¹⁶⁰⁸	Link ¹⁶⁰⁹
Reloader	reloader	1.0. 29	<ul style="list-style-type: none"> • chart: 1.0.29¹⁶¹⁰ • reloader: 1.0.29 	Link ¹⁶¹¹	Link ¹⁶¹²
Rook Ceph	rook-ceph	1.1 1.9	<ul style="list-style-type: none"> • chart: 1.11.9¹⁶¹³ • rook-ceph: 1.11.9 	Link ¹⁶¹⁴	Link ¹⁶¹⁵

¹⁶⁰¹ <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4>

¹⁶⁰² <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4?modal=values>

¹⁶⁰³ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/project-grafana-logging/6.57.4/defaults/cm.yaml>

¹⁶⁰⁴ <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16>

¹⁶⁰⁵ <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16?modal=values>

¹⁶⁰⁶ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/project-grafana-loki/0.69.17/defaults/cm.yaml>

¹⁶⁰⁷ <https://artifacthub.io/packages/helm/prometheus-community/prometheus-adapter/4.2.0>

¹⁶⁰⁸ <https://artifacthub.io/packages/helm/prometheus-community/prometheus-adapter/4.2.0?modal=values>

¹⁶⁰⁹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/prometheus-adapter/4.2.1/defaults/cm.yaml>

¹⁶¹⁰ <https://github.com/stakater/Reloader/tree/v1.0.29/README.md#helm-charts>

¹⁶¹¹ <https://artifacthub.io/packages/helm/stakater/reloader/1.0.29?modal=values>

¹⁶¹² <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/reloader/1.0.29/defaults/cm.yaml>

¹⁶¹³ <https://github.com/rook/rook/tree/vv1.11.9/Documentation/Helm-Charts/operator-chart.md>

¹⁶¹⁴ <https://raw.githubusercontent.com/rook/rook/v1.11.9/v%25s/deploy/charts/rook-ceph/values.yaml>

¹⁶¹⁵ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/rook-ceph/1.11.9/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Rook Ceph Cluster	rook-ceph-cluster	1.1 1.9	<ul style="list-style-type: none"> • chart: 1.11.9¹⁶¹⁶ • rook-ceph: 1.11.9 • rook-ceph-cluster: 17.2.6 	Link ¹⁶¹⁷	Link ¹⁶¹⁸
Thanos	thanos	12.8.3	<ul style="list-style-type: none"> • chart: 12.8.3¹⁶¹⁹ • thanos: 0.31.0 	Link ¹⁶²⁰	Link ¹⁶²¹
Traefik	traefik	23.1.0	<ul style="list-style-type: none"> • chart: 23.1.0¹⁶²² • traefik: 2.10.1 	Link ¹⁶²³	Link ¹⁶²⁴
Traefik ForwardAuth	traefik-forward-auth	0.3.10	<ul style="list-style-type: none"> • chart: 0.3.10¹⁶²⁵ • traefik-forward-auth: 3.1.0 	Link ¹⁶²⁶	Link ¹⁶²⁷
Velero	velero	4.1.1	<ul style="list-style-type: none"> • chart: 4.1.1¹⁶²⁸ • velero: 1.11.0 	Link ¹⁶²⁹	Link ¹⁶³⁰

12.2.7 DKP 2.6.1 Known Issues and Limitations

The following items are known issues with this release.

¹⁶¹⁶ <https://github.com/rook/rook/tree/vv1.11.9/Documentation/Helm-Charts/ceph-cluster-chart.md>

¹⁶¹⁷ <https://raw.githubusercontent.com/rook/rook/v1.11.9/v%25s/deploy/charts/rook-ceph-cluster/values.yaml>

¹⁶¹⁸ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/rook-ceph-cluster/1.11.9/defaults/cm.yaml>

¹⁶¹⁹ <https://artifacthub.io/packages/helm/bitnami/thanos/12.8.3>

¹⁶²⁰ <https://artifacthub.io/packages/helm/bitnami/thanos/12.8.3?modal=values>

¹⁶²¹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/thanos/12.8.3/defaults/cm.yaml>

¹⁶²² <https://artifacthub.io/packages/helm/traefik/traefik/23.1.0>

¹⁶²³ <https://artifacthub.io/packages/helm/traefik/traefik/23.1.0?modal=values>

¹⁶²⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/traefik/23.1.0/defaults/cm.yaml>

¹⁶²⁵ <https://artifacthub.io/packages/helm/mesosphere/traefik-forward-auth/0.3.10>

¹⁶²⁶ <https://artifacthub.io/packages/helm/mesosphere/traefik-forward-auth/0.3.10?modal=values>

¹⁶²⁷ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/traefik-forward-auth/0.3.10/defaults/cm.yaml>

¹⁶²⁸ <https://artifacthub.io/packages/helm/vmware-tanzu/velero/4.1.1>

¹⁶²⁹ <https://artifacthub.io/packages/helm/vmware-tanzu/velero/4.1.1?modal=values>

¹⁶³⁰ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.1/services/velero/4.1.1/defaults/cm.yaml>

12.2.7.1 Known Issues

12.2.7.1.1 AWS `additionalTags` cannot contain spaces

Due to an upstream bug in the `cluster-api-provider-aws` component, it is not possible to specify tags with spaces in their name in the `additionalTags` section of an `AWSCluster`. If you have any tags like this during an upgrade of the `capi-components`, you may receive a validation error, and will need to remove any such tags. This issue will be corrected in a future DKP release.

12.2.7.1.2 Use Static Credentials to Provision an Azure Cluster

Only static credentials can be used when [provisioning an Azure cluster](#) (see page 1166).

12.2.7.1.3 Containerd File Limit Issue

This version of DKP uses containerd 1.6.17. The systemd unit for containerd 1.6.17 provided upstream removes all file number limits (`LimitNOFILE=infinity`). In our testing, we found that removing these limits broke some IO-sensitive applications such as Rook Ceph and HAProxy. Because of this, the KIB version included in this release sets the `LimitNOFILE` value in the containerd systemd unit to the value (`1048576`), which was used in previous containerd 1.4.13 version releases.

12.2.7.1.4 Intermittent Error Status when Creating EKS Clusters in the UI

When provisioning an EKS cluster through the UI, you might receive a brief error state because the EKS cluster can sporadically lose connectivity with the management cluster, which results in the following symptoms:

- The UI shows the cluster is in an error state.
- The kubeconfig generated and retrieved from Kommander ceases to work.
- Applications created on the management cluster may not be immediately federated to managed EKS clusters.

After a few moments, the error self-resolves, without any action on your part. A new kubeconfig generated and retrieved from Kommander then works properly, and the UI shows that it is working again. In the meantime, you can continue to use the UI to work on the cluster such as deploy applications, create projects, and add roles.

12.2.7.1.5 Limitations to Disk Resizing in vSphere

The DKP CLI flags `--control-plane-disk-size` and `--worker-disk-size` are unable to resize the root file system of VMs, which are created using OS images. The flags work by resizing the primary disk of the VM. When the VM boots, the root file system expands to fill the disk, but that expansion does not work for some file systems. For example, for file systems contained in an LVM Logical Volume. To ensure your

root file system has the size you expect, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273384734/2.6%2BCreate%2Ba%2BvSphere%2BBase%2BOS%2BImage#Disk-Size> (see page 0).

12.2.7.1.6 Error Status in Grafana Logging Dashboard with EKS Clusters

Currently, it is not possible to use FluentBit to collect [Admin-level logs](#) (see page 905) on a managed EKS cluster.

If you have these logs enabled, the following message appears when you access the Kubernetes Audit Dashboard in the Grafana Logging Dashboard:

```
Cannot read properties of undefined (reading '0')
```

12.2.7.1.7 Spark Image Upstream Removal

Google has removed the images from their [Spark registry](#)¹⁶³¹ upstream `gcr.io/spark-operator/spark-py:v3.1.1`. Therefore, you are not able to pull these images for non-air-gapped environments. Refer also to our [Deprecations](#) (see page 1809) notice for more information.

12.2.7.1.8 Rook Ceph Install Error

An issue might emerge when installing `rook-ceph` on vSphere clusters using RHEL operating systems.

This issue occurs during initial installation of `rook-ceph`, causing the object store used by Velero and Grafana Loki, to be unavailable. If the installation of Kommander component of DKP is unsuccessful due to `rook-ceph` failing, you might need to apply the following workaround:

1. Run this command to check if the cluster is affected by this issue.

```
kubectl describe CephObjectStores dkp-object-store -n kommander
```

2. If this output appears, the workaround needs to be applied so continue with the next step. If you do not see this output, you can stop at this step.

```
Name:          dkp-object-store
Namespace:     kommander
...
Warning ReconcileFailed      7m55s (x19 over 52m)
rook-ceph-object-controller failed to reconcile CephObjectStore
"kommander/dkp-object-store". failed to create object store deployments:
failed
to configure multisite for object store: failed create ceph multisite for
```

¹⁶³¹ <http://gcr.io/spark-operator/spark-py:v3.1.1>

```
object-store ["dkp-object-store"]: failed to commit config changes after
creating multisite config for CephObjectStore "kommander/dkp-object-store":
failed to commit RGW configuration period changes%(EXTRA []string=[]):
signal: interrupt
```

3. Kubectl exec into the `rook-ceph-tools` pod.

```
export WORKSPACE_NAMESPACE=<workspace namespace>
CEPH_TOOLS_POD=$(kubectl get pods -l app=rook-ceph-tools -n $
{WORKSPACE_NAMESPACE} -o name)
kubectl exec -it -n ${WORKSPACE_NAMESPACE} ${CEPH_TOOLS_POD} bash
```

4. Run these commands to set `dkp-object-store` as the default zonegroup.

i NOTE: The `period update` command may take a few minutes to complete

```
radosgw-admin zonegroup default --rgw-zonegroup=dkp-object-store
radosgw-admin period update --commit
```

5. Next, restart the `rook-ceph-operator` deployment for the `CephObjectStore` to be reconciled.

```
kubectl rollout restart deploy -n${WORKSPACE_NAMESPACE} rook-ceph-operator
```

6. After running the commands above, the `CephObjectStore` should be `Connected` once the `rook-ceph` operator reconciles the object (this may take some time).

```
kubectl wait CephObjectStore --for=jsonpath='{.status.phase}'=Connected dkp-
object-store -n ${WORKSPACE_NAMESPACE} --timeout 10m
```

12.2.7.1.9 Kommander Installation Configuration File Changes

In this release, the `kommander-ui`, which provides the DKP Dashboard, is now deployed in the same manner as the other Platform Applications. Also, the `ai-navigator-app`, which provides the AI Navigator, is deployed in the same manner. Therefore, the contents of the default Kommander installation configuration file, which is the file produced by the `dkp install kommander --init` command, have changed. If you installed previous versions of DKP using a customized Kommander installation configuration file, we recommend:

1. Generate a new template for this release.
2. Reapply your customizations rather than reusing a file created by older DKP versions.
3. The `ai-navigator-app` will be enabled by default in the new configuration file

12.2.7.1.10 FluentD Logging Operator Error

An error may occur with the FluentD Operator not appearing when installing Kommander on a GCP Environment.

This is due to a limitation of the GCP API. As a result, you need to rename the Fluentd buffer name because the default name exceeds 63 characters and disable buffer volume metrics for FluentD in order for it to operate properly.

Follow these procedures to resolve this issue for the Management Cluster, and Managed or Attached Clusters.

12.2.7.1.10.1 Management Cluster on GCP with the Logging Stack enabled

1. Create the `kommander` namespace:

```
kubectl create namespace kommander
```

2. Create the logging-operator-logging-overrides ConfigMap

```
cat <<EOF | kubectl apply -n kommander -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: logging-operator-logging-overrides
data:
  values.yaml: |
    fluentd:
      # disable buffer metrics
      bufferVolumeMetrics: null
      bufferStorageVolume:
        pvc:
          source:
            # update this name to make the PV <= 63 chars long
            # logging-operator-logging-buf-logging-operator-logging-fluentd-0
            claimName: buf
          spec:
            accessModes:
              - ReadWriteOnce
            resources:
              requests:
                storage: 10Gi
            volumeMode: Filesystem
EOF
```

3. Continue installing Kommander as normal.

12.2.7.1.10.2 Managed/Attached Cluster in GCP with Logging Stack enabled

1. Set the `WORKSPACE_NAMESPACE` environment variable to the name of your workspace's namespace:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```

2. Create the logging-operator-logging-overrides ConfigMap on the managed or attached cluster prior to enabling the logging-operator application:

```
cat <<EOF | kubectl apply -n ${WORKSPACE_NAMESPACE} -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: logging-operator-logging-overrides
data:
  values.yaml: |
    fluentd:
      # disable buffer metrics
      bufferVolumeMetrics: null
      bufferStorageVolume:
        pvc:
          source:
            # update this name to make the PV <= 63 chars long
            # logging-operator-logging-buf-logging-operator-logging-fluentd-0
            claimName: buf
          spec:
            accessModes:
              - ReadWriteOnce
            resources:
              requests:
                storage: 10Gi
            volumeMode: Filesystem
EOF
```

3. Proceed with installing the logging-operator as normal.


12.2.7.1.11 Generate a Token or Enable the Konvoy Credentials Plugin for Proxied Clusters

If you have attached a [network-restricted cluster](#) (see page 81) and enabled a [proxied access](#) (see page 817) to make its resources available through the Management cluster, the “Generate Token” and “Konvoy credentials plugin instructions” options required for user authentication in the DKP UI require further steps.

Expand the set of instructions depending on the authentication configuration you want to configure:

12.2.7.1.11.1 A. Generate a Token

Generate Token**I. Execute the commands established in the “Generate Token” page.**

 You can find the “Generate Token” page in the DKP UI by selecting your username in the top right corner of any page.

II. Prepare your Environment

1. Write down the name of your credentials. You can obtain this name by accessing the **DKP UI** > username in the top right corner > **Generate token** page. The credentials are displayed in the `kubectl config set-credentials` command, next to `set-credentials`.

In the following example, `jane_doe-`

`az189abcd7ab34a2abc0ab845a67a7ea-0123456789.us-west-2.elb.amazonaws.com` would be the name of your credentials:

```
kubectl config set-credentials jane_doe-
az189abcd7ab34a2abc0ab845a67a7ea-0123456789.us-west-2.elb.amazonaws.com \
  --token=[...]
```

2. Set the `CONFIG_CREDENTIALS_NAME` environment variable to the credentials you obtained in the previous step:

```
CONFIG_CREDENTIALS_NAME=<CREDENTIALS_NAME>
```

3. Set `WORKSPACE_NAMESPACE` to the namespace of your network-restricted cluster:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```


4. Set `TUNNEL_PROXY_NAME` to the name of the TunnelProxy:

```
TUNNEL_PROXY_NAME=<...>
```

5. Set `DEFAULT_CONFIG_PATH` to the path file of the `kubeconfig` where you executed the “Generate Token” steps:

```
DEFAULT_CONFIG_PATH=<path_of_cluster's_kubeconfig>
```

- With your `kubeconfig` set to the **Management cluster**, set the `TUNNEL_PROXY_DOMAIN` environment variable to the proxied access URL:

 See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.

```
TUNNEL_PROXY_DOMAIN=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n
$WORKSPACE_NAMESPACE -o template='{{ .status.clusterProxyDomain }}')
```

III. Reconfigure the Token


Select one of the following options according to the certificate you configured for the proxied access:

A. Clusters with ACME-supported Certificates for the `TunnelProxy`

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
${TUNNEL_PROXY_DOMAIN}/dkp/api-server --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --cluster=${TUNNEL_PROXY_NAME} --
user=${CONFIG_CREDENTIALS_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
```

B. Clusters with Self-Signed Certificates for the `TunnelProxy`

- With your `kubeconfig` set to the **Management cluster**, retrieve the `ca.crt` (not needed in case of ACME certificates):

 See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.

```
LABEL_KEY=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n $WORKSPACE_NAMESPACE
-o template='{{ .status.reverseProxyReleaseName }}')
SECRET_NAME=$(kubectl get ingress -n $WORKSPACE_NAMESPACE -l app.kubernetes.io/
instance=$LABEL_KEY -o jsonpath='{.items[0].spec.tls[0].secretName}')
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
kubectl get secrets $SECRET_NAME -n $WORKSPACE_NAMESPACE -o
template='{{ index .data "ca.crt" | base64decode }}' > ${HOME}/.kube/certs/$
{TUNNEL_PROXY_NAME}/ca.crt
```

- Configure the `kubeconfig` access via the proxied access:

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
{TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/
certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --cluster=${TUNNEL_PROXY_NAME}
--user=${CONFIG_CREDENTIALS_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig $
{DEFAULT_CONFIG_PATH}
```

C. Clusters with Custom Certificates for the TunnelProxy

1. Copy the `ca.crt` used to sign your certificates to the `kubeconfig` directory:

```
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
cp ca.crt ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}/ca.crt
```

2. Configure the `kubeconfig` access via the proxied access:

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
{TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/
certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --cluster=${TUNNEL_PROXY_NAME}
--user=${CONFIG_CREDENTIALS_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig $
{DEFAULT_CONFIG_PATH}
```

After the commands complete, you can perform `kubectl` actions on the network-restricted cluster.

12.2.7.1.11.2 B. Konvoy credentials plugin instructions

Konvoy credentials plugin instructions

- I. Execute the commands provided in the “Konvoy credentials plugin instructions” page > “Method 1” section.



You can find a link to the “Konvoy credentials plugin instructions” page in the DKP UI when logging in to your environment, or by adding `/token/plugin` to the DKP UI URL of your environment.

II. Prepare your Environment

1. Set `WORKSPACE_NAMESPACE` to the namespace of your network-restricted cluster:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```


2. Set `TUNNEL_PROXY_NAME` to the name of the TunnelProxy:

```
TUNNEL_PROXY_NAME=<...>
```

3. Set `DEFAULT_CONFIG_PATH` to the path of the kubeconfig file you downloaded from the dashboard:

```
DEFAULT_CONFIG_PATH=<path_of_cluster's_kubeconfig>
```

4. With your kubeconfig set to the **Management cluster**, set the `TUNNEL_PROXY_DOMAIN` environment variable to the proxied access URL:

 See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.

```
TUNNEL_PROXY_DOMAIN=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n
$WORKSPACE_NAMESPACE -o template='{{ .status.clusterProxyDomain }}')
```

III. Reconfigure the Token


Select one of the following options according to the certificate you configured for the proxied access:

A. Clusters with ACME-supported Certificates for the TunnelProxy

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
${TUNNEL_PROXY_DOMAIN}/dkp/api-server --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --user default-profile --kubeconfig $
${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
```

B. Clusters with Self-Signed Certificates for the TunnelProxy

1. With your kubeconfig set to the **Management cluster**, retrieve the `ca.crt` (not needed in case of ACME certificates):

 See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.


```

LABEL_KEY=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n $WORKSPACE_NAMESPACE
-o template='{.status.reverseProxyReleaseName }')
SECRET_NAME=$(kubectl get ingress -n $WORKSPACE_NAMESPACE -l app.kubernetes.io/
instance=$LABEL_KEY -o jsonpath='{.items[0].spec.tls[0].secretName}')
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
kubectl get secrets $SECRET_NAME -n $WORKSPACE_NAMESPACE -o
template='{ index .data "ca.crt" | base64decode }' > ${HOME}/.kube/certs/$
{TUNNEL_PROXY_NAME}/ca.crt

```

2. Configure the `kubeconfig` access via the proxied access:

```

kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
{TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/
certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --user default-profile --
kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig $
{DEFAULT_CONFIG_PATH}

```

C. Clusters with Custom Certificates for the `TunnelProxy`

1. Copy the `ca.crt` used to sign your certificates to the `kubeconfig` directory:

```

mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
cp ca.crt ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}/ca.crt


```

2. Configure the `kubeconfig` access via the proxied access:

```

kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
{TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/
certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --user default-profile --
kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig $
{DEFAULT_CONFIG_PATH}

```

 “Method 2” is not available for network-restricted clusters.

After the commands complete, you can perform `kubectl` actions on the network-restricted cluster.

12.2.7.1.12 Known Common Vulnerabilities and Exposures (CVE)

Starting with DKP 2.6, Catalog apps are scanned for CVEs. However, only the CVEs for app versions that are compatible with the default Kubernetes version, currently 1.26.6, are mitigated. For more information about the known CVEs for compatible catalog apps, see [D2iQ Security Updates \(see page 1851\)](#).

12.2.7.1.13 Potential Logging Interruption during Upgrade

The configuration setting `enableRecreateWorkloadOnImmutableFieldChange` is enabled by default on the `logging-operator`. This means the operator automatically triggers a recreation of the resource if any immutable fields are changed on the underlying fluent-bit/fluentd resources. Previously, any changes to immutable fields would fail silently and the failures would only be observable in the `logging-operator` logs. This was done to handle a breaking change introduced in the `logging-operator`, which added an additional selector label on the `fluent-bit` DaemonSet.

Because of this breaking change, during an upgrade to 2.6, `fluent-bit` is restarted, and some log data in the `fluent-bit` buffer might be lost if `fluentd` is unavailable while `fluent-bit` attempts to flush its buffer prior to the pods being terminated. To prevent issues related to data loss, we recommend you run `>1 fluentd pod`. Alternatively, Fluent Bit can also be configured to use a `hostPath` volume to store the buffer information, so it can be picked up again when Fluent Bit restarts. For more information about how to change the host path, see [Fluent Bit log collector](#)¹⁶³².

12.2.7.1.14 Gatekeeper Uninstallation Error

If you choose to disable Gatekeeper, you can run into an error where the app is still present on your clusters.

For pre-existing attached or managed clusters with Gatekeeper installed when they don't want it to be, you need to manually cleanup Gatekeeper on those clusters, post upgrade.

Follow these steps to manually remove Gatekeeper:

1. For every **Attached/Managed cluster** that has Gatekeeper installed that you have disabled it on otherwise (via appdeployment), run the following commands to remove it:

NOTE: This needs to be run on the attached/managed cluster using the correct kubeconfig. Set your namespace of your **attached** cluster's workspace on your attached cluster:

```
export WORKSPACE_NAMESPACE=<workspace-name>
```

2. Delete the kustomizations:

¹⁶³² <https://kube-logging.dev/docs/logging-infrastructure/fluentbit/#hostpath-volumes-for-buffers-and-positions>

```
kubectl --kubeconfig <attached-cluster-kubeconfig> delete kustomizations -n $
{WORKSPACE_NAMESPACE} gatekeeper-constraint-templates gatekeeper-constraints
gatekeeper-release
```

12.2.7.1.15 Pre-provisioned Ubuntu Non-air-gapped Konvoy Image Builder Provisioning Error

Due to an upstream repository change, Konvoy Image Builder (KIB) cannot install the OS packages required for DKP when provisioning each machine. This can result in failures when creating Pre-provisioned Ubuntu non-air-gapped clusters, or when modifying their configuration in a way that machines must be provisioned.

New DKP patch releases 2.6.2 and 2.7.1 contain corresponding Konvoy Image Builder releases that include the changes necessary to provision machines successfully.

12.2.8 DKP 2.6.1 Customer Incidents and Resolved Issues

The following resolved incidents are included in this release.

12.2.8.1 Prometheus has duplicate targets for the apiserver

[D2IQ-98647](#)

The default values in the kube-prometheus-stack chart included a duplicate scrape target for the Kubernetes API server. The duplicate scrape target is removed.

12.2.8.2 Running the DKP upgrade command twice causes a regression in the associated containerd secret

[D2IQ-98947](#)

In DKP 2.6, if a cluster was configured with a registry-mirror and credentials, running the `update controlplane` command twice corrupted the value of the containerd secret. This issue has been corrected.

12.2.8.3 Pre-provisioned Upgrade failed with Flatcar

An issue was resolved that caused Kubernetes upgrades to fail when using Flatcar OS with the pre-provisioned provider.

[D2IQ-98979](#)

12.2.8.4 Editing Credentials of a vSphere cluster in the DKP UI were not updated

An issue was resolved that could corrupt the configuration of vSphere clusters by editing their settings in the UI.

[D2IQ-98762](#)

12.2.8.5 Limits on Capi-controller needed raised

A performance problem was resolved after being observed to occur when creating large numbers of workload clusters.

[D2IQ-99403](#)

12.2.8.6 Insight alerts from an Essentials cluster converted into Managed land into the management cluster workspace

DKP resolved a problem where Insights were not displayed for **Managed Clusters** that had been expanded from Essential clusters.

[D2IQ-98597](#)

12.2.8.7 Grafana-Loki telemetry sending data back to Grafana

DKP disabled telemetry collection settings in the configuration of Grafana-Loki and Grafana.

[D2IQ-99007](#)

12.2.8.8 Increased Time to Attach Several Clusters in 2.6.0

In the 2.6.0 DKP release, when [attaching clusters](#) (see page 755) in a batch, there was an issue with increased time to attach several clusters. [This issue](#) (see page 0) has been resolved.

[D2IQ-98664](#)

12.2.9 DKP 2.6.1 Deprecations

The following items are deprecated or removed from this version of DKP.

12.2.9.1 Deprecated `dkp push` Subcommand and Flag

To improve consistency, simplify, and streamline DKP CLI commands, this version of DKP deprecates the `dkp push image-bundle` command in favor of `dkp push bundle`.

Deprecated subcommand and flag	New subcommand and flag
<code>dkp push image-bundle --image-bundle</code>	<code>dkp push bundle --bundle</code>

You can still run the `dkp push image-bundle` with this version of DKP, but it will be removed in future DKP versions.

If you have scripts, GitHub Actions, or other cluster creation and upgrade automations, D2iQ recommends updating to the new subcommand and flag as soon as possible.

12.2.9.2 Catalog Applications Compatibility

From DKP 2.6, the following versions of Catalog applications are deprecated:

- kafka-operator-0.20.0
- kafka-operator-0.20.2
- spark-operator-1.1.6
- zookeeper-operator-0.2.13

If you plan on upgrading to DKP 2.6, ensure that you upgrade these applications to the latest compatible version.

For more information, see [Workspace DKP Catalog Applications - Compatibility \(see page 667\)](#).

12.2.9.3 Spark Operator to be Removed in DKP 2.7

Starting from DKP 2.7, D2iQ will discontinue support for all versions of Spark Operator in the DKP Catalog, because the upstream operator is no longer maintained. As a result, the Spark Operator will no longer be available for use in future DKP releases.

Spark Operator remains available in DKP 2.6.1 and previous versions.

12.3 DKP 2.6.2 Release Notes

DKP® version 2.6.2 was released on March 28, 2024.

[Download DKP](#)



You must be a registered user and logged on to the support portal to download this product. New customers must contact their sales representative or sales@d2iq.com¹⁶³³ before attempting to download or install DKP.

12.3.1 Release Summary

Welcome to D2iQ Kubernetes Platform (DKP) 2.6.2! This release focuses on... Below, you will find information about the following:

¹⁶³³ <mailto:sales@d2iq.com>

- [DKP 2.6.2 Features and Enhancements](#) (see page 1811)
- [DKP 2.6.2 Supported Kubernetes Versions](#) (see page 1812)
- [DKP 2.6.2 Kubernetes Major Updates and Deprecations](#) (see page 1813)
- [DKP 2.6.2 Components and Applications](#) (see page 1814)
- [DKP 2.6.2 Known Issues and Limitations](#) (see page 1823)
- [DKP 2.6.2 Customer Incidents and Resolved Issues](#) (see page 1835)
- [DKP 2.6.2 Deprecations](#) (see page 1835)

12.3.2 Additional resources

- For more information about working with native Kubernetes, see the [Kubernetes documentation](#)¹⁶³⁴.
- For a full list of attributed 3rd party software, see <http://d2iq.com/legal/3rd>.

12.3.3 DKP 2.6.2 Features and Enhancements

The following improvements are included in this release.

12.3.3.1 Kubernetes 1.26.14 Support

Kubernetes 1.26.14 enables you to benefit from the latest features and security fixes in upstream Kubernetes. This release of Kubernetes comes with approximately [40 enhancements](#)¹⁶³⁵ from which you can benefit.

To read more about major features in this release, visit [this page](#)¹⁶³⁶ and [Kubernetes v1.26 Electrifying](#)¹⁶³⁷.

12.3.3.2 AWS Custom AMI Required

Previous versions of DKP would default to using upstream AMIs published by the CAPA (Cluster API AWS) project when building AWS clusters if you did not specify your own AMI. However, those images are not currently available for the Kubernetes version used in the 2.6.2 patch release.

As a result, starting with this release of DKP, the behavior of the DKP `create cluster aws` command has been changed. It no longer defaults to using the upstream AMIs and instead requires that you specify an AMI built using [Konvoy Image Builder \(KIB\)](#) (see page 1468), or by explicitly requesting that it use the upstream images.

For more information on using a custom AMI in cluster creation or during the upgrade process, refer to these topics:

- [Create a Custom AMI](#) (see page 1472)
- [AWS Air-gapped AMI](#) (see page 1477)

¹⁶³⁴ <https://kubernetes.io/docs/home/>

¹⁶³⁵ <https://github.com/orgs/kubernetes/projects/98>

¹⁶³⁶ <https://github.com/kubernetes/sig-release/blob/master/releases/release-1.26/README.md>

¹⁶³⁷ <https://kubernetes.io/blog/2022/12/09/kubernetes-v1-26-release/>

When applying the AMI during cluster creation, there are two approaches to supplying the ID of your AMI. Either provide the ID of the AMI or provide a way for DKP to discover the AMI using location, format and OS information:

1. **Option One** - Provide the ID of your AMI:
 - a. Use the example command below leaving the existing flag that provides the AMI ID: `--ami AMI_ID`
2. **Option Two** - Provide a path for your AMI with the information required for image discover:
 - a. Where the AMI is published using your AWS Account ID: `--ami-owner AWS_ACCOUNT_ID`
 - b. The format or string used to search for matching AMIs and ensure it references the Kubernetes version plus the base OS name: `--ami-base-os ubuntu-20.04`
 - c. The base OS information: `--ami-format 'example-{{.BaseOS}}-?{{.K8sVersion}}-*'`

12.3.3.3 Konvoy Image Builder (KIB) [Release v2.5.5](#)¹⁶³⁸

The new release of Konvoy Image Builder features:

- Upgrade Kubernetes version to v1.26.14

12.3.3.4 AI Navigator Optional Install

With this release, the AI Navigator is optional and can be disabled at install time or uninstalled. To disable AI Navigator, [modify the Kommander configuration file after generating it](#)¹⁶³⁹, to disable the `ai-navigator-app` entry in the `apps:` section of the file.

Currently, the AI Navigator is not available for air-gapped environments. For more information on the AI Navigator, refer to the documentation section for that feature: [AI Navigator \(see page 1839\)](#)

12.3.4 DKP 2.6.2 Supported Kubernetes Versions

12.3.4.1 Deploying Clusters Versions

The latest supported Kubernetes versions for DKP are listed below. In DKP 2.6.2, the Kubernetes version installed by default is 1.26.14. The newest and oldest Kubernetes versions used with DKP must be within one minor version.

- Latest supported Kubernetes version is at **1.26.14 for deploying clusters**
- Other Kubernetes versions are supported at **1.25.13 for deploying clusters in EKS**

¹⁶³⁸ <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.5.5>

¹⁶³⁹ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/233177089>

Kubernetes 1.26.14 enables you to benefit from the latest features and security fixes in upstream Kubernetes. This release comes with approximately 40 enhancements that you can benefit from like [Service internal tracking policy](#)¹⁶⁴⁰, [Kubelet credential provider](#)¹⁶⁴¹, among others. To read more about major features in this release, visit [GitHub](#)¹⁶⁴² and <https://kubernetes.io/blog/2022/12/09/kubernetes-v1-26-release/>.

12.3.4.2 Attaching Clusters Versions

DKP 2.6.2 supports attaching clusters with the following Kubernetes versions:

Product	Compatible Kubernetes Versions
EKS	1.25.13
AKS	1.26.10
GKE	1.26.x



Attaching Kubernetes clusters with versions greater than n-1 is not recommended.



When attempting to attach a cluster with a lower Kubernetes version than the DKP default, you need to build and use an image with that lower version of Kubernetes. See [Konvoy Image Builder](#) (see page 1464) for documentation on building images. Failure to do this could result in an error.

12.3.5 DKP 2.6.2 Kubernetes Major Updates and Deprecations



Before upgrading, we **strongly recommend** verifying your current setup against the information on this page and reading more about Kubernetes' new features in [Kubernetes upgrade notes](#)¹⁶⁴³.

¹⁶⁴⁰ <https://github.com/kubernetes/enhancements/issues/2086>

¹⁶⁴¹ <https://github.com/kubernetes/enhancements/issues/2133>

¹⁶⁴² <https://github.com/kubernetes/sig-release/blob/master/releases/release-1.26/README.md>

¹⁶⁴³ <https://github.com/kubernetes/kubernetes/blob/master/CHANGELOG/CHANGELOG-1.26.md>

12.3.5.1 Deprecated API Services

With this version, Kubernetes stopped serving several API versions for Flow control resources and HorizontalPodAutoscaler services. For more information on the deprecated API versions for each of these services, refer to <https://kubernetes.io/docs/reference/using-api/deprecation-guide/#v1-26>.

12.3.5.2 Removed OCI Image Registry

DKP uses the upstream `registry.k8s.io` registry instead of the previous `k8s.gcr.io` registry for all new and upgraded clusters. The Kubernetes image is only available through the new image registry starting with 1.26.0.

Ensure you add `registry.k8s.io` to any firewall or proxy allowlists.

12.3.5.3 Migration to External Cloud-Controller-Managers

The deprecation of in-tree `cloud-controller-managers` that are part of `kube-controller-manager` has started. Starting with DKP 2.6 and Kubernetes 1.26.0, the in-tree `cloud-controller-manager` is disabled and migrated to the external `cloud-controller-manager`. If you had custom configuration for the in-tree `cloud-controller-manager`, you may need to add an equivalent configuration in the new deployments.

12.3.6 DKP 2.6.2 Components and Applications

The following are component and application versions for DKP.

12.3.6.1 Components

Component Name	Version
Cluster API Core (CAPI)	1.4.3-d2iq.0
Cluster API Amazon AWS Infrastructure Provider (CAPA)	2.2.0
Cluster API Google Cloud Infrastructure Provider (CAPG)	1.3.1
Cluster API Pre-provisioned Infrastructure Provider (CAPPP)	0.16.3

Component Name	Version
Cluster API vSphere Infrastructure Provider (CAPV)	1.6.1
Cluster API Azure Infrastructure Provider (CAPZ)	1.8.5
Cluster API for VMware Cloud Director (CAPVCD)	1.1.0-d2iq.0 ¹⁶⁴⁴
Konvoy Image Builder (KIB)	2.5.5 ¹⁶⁴⁵
containerd	1.6.28
etcd	3.5.6
Calico	3.26.1 ¹⁶⁴⁶
Cluster Autoscaler	1.27.3
CSI_VERSION	Default Storage Providers in DKP (see page 91) See this page for versions and driver table.
Metal LB	0.12.1
Node Feature Discovery	0.13.2

¹⁶⁴⁴ <https://github.com/vmware/cluster-api-provider-cloud-director>

¹⁶⁴⁵ <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v2.5.5>

¹⁶⁴⁶ https://github.com/projectcalico/calico/blob/v3.26.1/calico/_includes/release-notes/v3.26.1-release-notes.md

12.3.6.2 Applications

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Centralized Grafana	centralized-grafana	46.8.0	<ul style="list-style-type: none"> • chart: 46.8.0¹⁶⁴⁷ • prometheus-operator: 0.65.2 • grafana: 9.5.5 	Link ¹⁶⁴⁸	Link ¹⁶⁴⁹
Centralized Kubecost	centralized-kubecost	0.35.1	<ul style="list-style-type: none"> • chart: 0.35.1¹⁶⁵⁰ • kubecost: 1.104.0 	Link ¹⁶⁵¹	Link ¹⁶⁵²
Cert Manager	cert-manager	1.12.2	<ul style="list-style-type: none"> • chart: 1.12.2¹⁶⁵³ • cert-manager: 1.12.2 	Link ¹⁶⁵⁴	Link ¹⁶⁵⁵
Chartmuseum	chartmuseum	3.10.0	<ul style="list-style-type: none"> • chart: 3.10.0¹⁶⁵⁶ • chartmuseum: 3.10.0 	Link ¹⁶⁵⁷	Link ¹⁶⁵⁸
Dex	dex	2.12.3	<ul style="list-style-type: none"> • chart: 2.12.3¹⁶⁵⁹ • dex: 2.35.2 	Link ¹⁶⁶⁰	Link ¹⁶⁶¹

¹⁶⁴⁷ <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.0>

¹⁶⁴⁸ <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.0?modal=values>

¹⁶⁴⁹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/centralized-grafana/46.8.0/defaults/cm.yaml>

¹⁶⁵⁰ <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1>

¹⁶⁵¹ <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1?modal=values>

¹⁶⁵² <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/centralized-kubecost/0.35.1/defaults/cm.yaml>

¹⁶⁵³ <https://artifacthub.io/packages/helm/cert-manager/cert-manager/1.12.2>

¹⁶⁵⁴ <https://artifacthub.io/packages/helm/cert-manager/cert-manager/1.12.2?modal=values>

¹⁶⁵⁵ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/cert-manager/1.12.2/defaults/cm.yaml>

¹⁶⁵⁶ <https://artifacthub.io/packages/helm/chartmuseum/chartmuseum/3.10.0>

¹⁶⁵⁷ <https://artifacthub.io/packages/helm/chartmuseum/chartmuseum/3.10.0?modal=values>

¹⁶⁵⁸ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/chartmuseum/3.10.0/defaults/cm.yaml>

¹⁶⁵⁹ <https://artifacthub.io/packages/helm/mesosphere-stable/dex/2.12.3>

¹⁶⁶⁰ <https://artifacthub.io/packages/helm/mesosphere-stable/dex/2.12.3?modal=values>

¹⁶⁶¹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/dex/2.12.3/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Dex K8s Authenticator	dex-k8s-authenticator	1.2.15	<ul style="list-style-type: none"> • chart: 1.2.15¹⁶⁶² • dex-k8s-authenticator: 1.2.4 	Link ¹⁶⁶³	Link ¹⁶⁶⁴
Dkp Insights	dkp-insights	0.5.1	<ul style="list-style-type: none"> • chart: 0.5.1 • dkp-insights: 0.5.1 	Link (see page 1817)	Link ¹⁶⁶⁵
DKP Insights Management	dkp-insights-management	0.5.1	<ul style="list-style-type: none"> • chart: 0.5.1 • dkp-insights-management: 0.5.1 	N/A	Link ¹⁶⁶⁶
External DNS	external-dns	6.2.0.4	<ul style="list-style-type: none"> • chart: 6.20.4¹⁶⁶⁷ • external-dns: 0.13.4 	Link ¹⁶⁶⁸	Link ¹⁶⁶⁹
Fluent Bit	fluent-bit	0.34.1	<ul style="list-style-type: none"> • chart: 0.34.1¹⁶⁷⁰ • fluent-bit: 2.1.6 	Link ¹⁶⁷¹	Link ¹⁶⁷²
Gatekeeper	gatekeeper	3.12.0	<ul style="list-style-type: none"> • chart: 3.12.0¹⁶⁷³ • gatekeeper: 3.12.0 	Link ¹⁶⁷⁴	Link ¹⁶⁷⁵

¹⁶⁶² <https://artifacthub.io/packages/helm/mesosphere/dex-k8s-authenticator/1.2.15>

¹⁶⁶³ <https://artifacthub.io/packages/helm/mesosphere/dex-k8s-authenticator/1.2.15?modal=values>

¹⁶⁶⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/dex-k8s-authenticator/1.2.15/defaults/cm.yaml>

¹⁶⁶⁵ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/dkp-insights/0.5.1/defaults/cm.yaml>

¹⁶⁶⁶ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/dkp-insights-management/0.5.1/defaults/cm.yaml>

¹⁶⁶⁷ <https://artifacthub.io/packages/helm/bitnami/external-dns/6.20.4>

¹⁶⁶⁸ <https://artifacthub.io/packages/helm/bitnami/external-dns/6.20.4?modal=values>

¹⁶⁶⁹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/external-dns/6.20.4/defaults/cm.yaml>

¹⁶⁷⁰ <https://artifacthub.io/packages/helm/fluent/fluent-bit/0.34.1>

¹⁶⁷¹ <https://artifacthub.io/packages/helm/fluent/fluent-bit/0.34.1?modal=values>

¹⁶⁷² <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/fluent-bit/0.34.1/defaults/cm.yaml>

¹⁶⁷³ <https://artifacthub.io/packages/helm/gatekeeper/gatekeeper/3.12.0>

¹⁶⁷⁴ <https://artifacthub.io/packages/helm/gatekeeper/gatekeeper/3.12.0?modal=values>

¹⁶⁷⁵ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/gatekeeper/3.12.0/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Gitea	gitea	8.2.0	<ul style="list-style-type: none"> • chart: 8.2.0¹⁶⁷⁶ • gitea: 1.19.2 	Link ¹⁶⁷⁷	Link ¹⁶⁷⁸
Grafana Logging	grafana-logging	6.57.4	<ul style="list-style-type: none"> • chart: 6.57.4¹⁶⁷⁹ • grafana: 9.5.5 	Link ¹⁶⁸⁰	Link ¹⁶⁸¹
Grafana Loki	grafana-loki	0.69.17	<ul style="list-style-type: none"> • chart: 0.69.16¹⁶⁸² • loki: 2.8.2 	Link ¹⁶⁸³	Link ¹⁶⁸⁴
Istio	istio	1.17.2	<ul style="list-style-type: none"> • chart: 1.17.2¹⁶⁸⁵ • istio: 1.17.2 	Link ¹⁶⁸⁶	Link ¹⁶⁸⁷
Jaeger	jaeger	2.46.2	<ul style="list-style-type: none"> • chart: 2.46.2¹⁶⁸⁸ • jaeger: 1.46.0 	Link ¹⁶⁸⁹	Link ¹⁶⁹⁰
Karma	karma	2.0.2	<ul style="list-style-type: none"> • chart: 2.0.2¹⁶⁹¹ • karma: 0.70 	Link ¹⁶⁹²	Link ¹⁶⁹³

¹⁶⁷⁶ <https://artifacthub.io/packages/helm/gitea/gitea/8.2.0>

¹⁶⁷⁷ <https://artifacthub.io/packages/helm/gitea/gitea/8.2.0?modal=values>

¹⁶⁷⁸ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/gitea/8.2.0/defaults/cm.yaml>

¹⁶⁷⁹ <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4>

¹⁶⁸⁰ <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4?modal=values>

¹⁶⁸¹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/grafana-logging/6.57.4/defaults/cm.yaml>

¹⁶⁸² <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16>

¹⁶⁸³ <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16?modal=values>

¹⁶⁸⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/grafana-loki/0.69.17/defaults/cm.yaml>

¹⁶⁸⁵ <https://artifacthub.io/packages/helm/mesosphere/istio/1.17.2>

¹⁶⁸⁶ <https://artifacthub.io/packages/helm/mesosphere/istio/1.17.2?modal=values>

¹⁶⁸⁷ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/istio/1.17.2/defaults/cm.yaml>

¹⁶⁸⁸ <https://artifacthub.io/packages/helm/jaegertracing/jaeger-operator/2.46.2>

¹⁶⁸⁹ <https://artifacthub.io/packages/helm/jaegertracing/jaeger-operator/2.46.2?modal=values>

¹⁶⁹⁰ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/jaeger/2.46.2/defaults/cm.yaml>

¹⁶⁹¹ <https://artifacthub.io/packages/helm/mesosphere-stable/karma/2.0.2>

¹⁶⁹² <https://artifacthub.io/packages/helm/mesosphere-stable/karma/2.0.2?modal=values>

¹⁶⁹³ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/karma/2.0.2/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Kiali	kiali	1.70.0	<ul style="list-style-type: none"> • chart: 1.70.0¹⁶⁹⁴ • kiali: 1.70.0 	Link ¹⁶⁹⁵	Link ¹⁶⁹⁶
Knative	knative	1.10.4	<ul style="list-style-type: none"> • chart: 1.10.4¹⁶⁹⁷ • knative: 1.10.0 	Link ¹⁶⁹⁸	Link ¹⁶⁹⁹
Flux	kommander-flux	0.41.2	<ul style="list-style-type: none"> • chart: N/A • flux: 0.41.2 	N/A	N/A
Kommander Ui	kommander-ui	9.65.2	<ul style="list-style-type: none"> • chart: 9.65.2 • kommander-ui: 9.65.2 	Link (see page 1817)	Link ¹⁷⁰⁰
Kube OIDC Proxy	kube-oidc-proxy	0.3.2	<ul style="list-style-type: none"> • chart: 0.3.2¹⁷⁰¹ • kube-oidc-proxy: 0.3.0 	Link ¹⁷⁰²	Link ¹⁷⁰³
Kube Prometheus Stack	kube-prometheus-stack	46.8.1	<ul style="list-style-type: none"> • chart: 46.8.1¹⁷⁰⁴ • prometheus-operator: 0.65.2 • grafana: 9.5.5 • prometheus: 2.44.0 • prometheus-alertmanager: 0.25.0 	Link ¹⁷⁰⁵	Link ¹⁷⁰⁶

1694 <https://artifacthub.io/packages/helm/kiali/kiali-operator/1.70.0>

1695 <https://artifacthub.io/packages/helm/kiali/kiali-operator/1.70.0?modal=values>

1696 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/kiali/1.70.0/defaults/cm.yaml>

1697 <https://artifacthub.io/packages/helm/mesosphere/knative/1.10.4>

1698 <https://artifacthub.io/packages/helm/mesosphere/knative/1.10.4?modal=values>

1699 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/knative/1.10.4/defaults/cm.yaml>

1700 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/kommander-ui/9.65.2/defaults/cm.yaml>

1701 <https://artifacthub.io/packages/helm/mesosphere/kube-oidc-proxy/0.3.2>

1702 <https://artifacthub.io/packages/helm/mesosphere/kube-oidc-proxy/0.3.2?modal=values>

1703 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/kube-oidc-proxy/0.3.2/defaults/cm.yaml>

1704 <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.1>

1705 <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/46.8.1?modal=values>

1706 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/kube-prometheus-stack/46.8.1/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Kubecost	kubecost	0.35.1	<ul style="list-style-type: none"> chart: 0.35.1¹⁷⁰⁷ kubecost: 1.104.0 	Link ¹⁷⁰⁸	Link ¹⁷⁰⁹
Kubefed	kubefed	0.10.3	<ul style="list-style-type: none"> chart: 0.10.3¹⁷¹⁰ kubefed: 0.10.3 	Link ¹⁷¹¹	Link ¹⁷¹²
Kubernetes Dashboard	kubernetes-dashboard	6.0.8	<ul style="list-style-type: none"> chart: 6.0.8¹⁷¹³ kubernetes-dashboard: 2.7.0 	Link ¹⁷¹⁴	Link ¹⁷¹⁵
Kubetunnel	kubetunnel	0.0.29	<ul style="list-style-type: none"> chart: 0.0.29 kubetunnel: 0.0.29 	N/A	Link ¹⁷¹⁶
Logging Operator	logging-operator	4.2.3	<ul style="list-style-type: none"> chart: 4.2.3¹⁷¹⁷ logging-operator: 4.2.2 logging-operator-logging: 4.2.3 	Link ¹⁷¹⁸	Link ¹⁷¹⁹

¹⁷⁰⁷ <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1>

¹⁷⁰⁸ <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.35.1?modal=values>

¹⁷⁰⁹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/kubecost/0.35.1/defaults/cm.yaml>

¹⁷¹⁰ <https://artifacthub.io/packages/helm/kubefed/kubefed/0.10.3>

¹⁷¹¹ <https://artifacthub.io/packages/helm/kubefed/kubefed/0.10.3?modal=values>

¹⁷¹² <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/kubefed/0.10.3/defaults/cm.yaml>

¹⁷¹³ <https://artifacthub.io/packages/helm/k8s-dashboard/kubernetes-dashboard/6.0.8>

¹⁷¹⁴ <https://artifacthub.io/packages/helm/k8s-dashboard/kubernetes-dashboard/6.0.8?modal=values>

¹⁷¹⁵ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/kubernetes-dashboard/6.0.8/defaults/cm.yaml>

¹⁷¹⁶ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/kubetunnel/0.0.29/defaults/cm.yaml>

¹⁷¹⁷ <https://artifacthub.io/packages/helm/banzaicloud-stable/logging-operator/4.2.3>

¹⁷¹⁸ <https://artifacthub.io/packages/helm/banzaicloud-stable/logging-operator/4.2.3?modal=values>

¹⁷¹⁹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/logging-operator/4.2.3/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
NFS Server Provisioner	nfs-server-provisioner	0.6.1	<ul style="list-style-type: none"> chart: 0.6.1¹⁷²⁰ nfs-server-provisioner: 2.3.0 	Link ¹⁷²¹	Link ¹⁷²²
NVIDIA GPU Operator	nvidia-gpu-operator	23.3.2	<ul style="list-style-type: none"> chart: 23.3.2¹⁷²³ nvidia-gpu-operator: 23.3.2 	Link ¹⁷²⁴	Link ¹⁷²⁵
Grafana (project)	project-grafana-logging	6.5.7.4	<ul style="list-style-type: none"> chart: 6.57.4¹⁷²⁶ grafana: 9.5.5 	Link ¹⁷²⁷	Link ¹⁷²⁸
Grafana Loki (project)	project-grafana-loki	0.6.9.17	<ul style="list-style-type: none"> chart: 0.69.16¹⁷²⁹ loki: 2.8.2 	Link ¹⁷³⁰	Link ¹⁷³¹
Prometheus Adapter	prometheus-adapter	4.2.1	<ul style="list-style-type: none"> chart: 4.2.0¹⁷³² prometheus-adapter: 0.10.0 	Link ¹⁷³³	Link ¹⁷³⁴

¹⁷²⁰ <https://artifacthub.io/packages/helm/mesosphere/nfs-server-provisioner/0.6.1>

¹⁷²¹ <https://artifacthub.io/packages/helm/mesosphere/nfs-server-provisioner/0.6.1?modal=values>

¹⁷²² <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/nfs-server-provisioner/0.6.1/defaults/cm.yaml>

¹⁷²³ <https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/getting-started.html#chart-customization-options>

¹⁷²⁴ <https://raw.githubusercontent.com/NVIDIA/gpu-operator/v23.3.2/v%25s/deployments/gpu-operator/values.yaml>

¹⁷²⁵ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/nvidia-gpu-operator/23.3.2/defaults/cm.yaml>

¹⁷²⁶ <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4>

¹⁷²⁷ <https://artifacthub.io/packages/helm/grafana/grafana/6.57.4?modal=values>

¹⁷²⁸ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/project-grafana-logging/6.57.4/defaults/cm.yaml>

¹⁷²⁹ <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16>

¹⁷³⁰ <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.69.16?modal=values>

¹⁷³¹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/project-grafana-loki/0.69.17/defaults/cm.yaml>

¹⁷³² <https://artifacthub.io/packages/helm/prometheus-community/prometheus-adapter/4.2.0>

¹⁷³³ <https://artifacthub.io/packages/helm/prometheus-community/prometheus-adapter/4.2.0?modal=values>

¹⁷³⁴ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/prometheus-adapter/4.2.1/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Reloader	reloader	1.0.29	<ul style="list-style-type: none"> • chart: 1.0.29¹⁷³⁵ • reloader: 1.0.29 	Link ¹⁷³⁶	Link ¹⁷³⁷
Rook Ceph	rook-ceph	1.11.9	<ul style="list-style-type: none"> • chart: 1.11.9¹⁷³⁸ • rook-ceph: 1.11.9 	Link ¹⁷³⁹	Link ¹⁷⁴⁰
Rook Ceph Cluster	rook-ceph-cluster	1.11.9	<ul style="list-style-type: none"> • chart: 1.11.9¹⁷⁴¹ • rook-ceph: 1.11.9 • rook-ceph-cluster: 17.2.6 	Link ¹⁷⁴²	Link ¹⁷⁴³
Thanos	thanos	12.8.3	<ul style="list-style-type: none"> • chart: 12.8.3¹⁷⁴⁴ • thanos: 0.31.0 	Link ¹⁷⁴⁵	Link ¹⁷⁴⁶
Traefik	traefik	23.1.0	<ul style="list-style-type: none"> • chart: 23.1.0¹⁷⁴⁷ • traefik: 2.10.1 	Link ¹⁷⁴⁸	Link ¹⁷⁴⁹

¹⁷³⁵ <https://github.com/stakater/Reloader/tree/v1.0.29/README.md#helm-charts>

¹⁷³⁶ <https://artifacthub.io/packages/helm/stakater/reloader/1.0.29?modal=values>

¹⁷³⁷ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/reloader/1.0.29/defaults/cm.yaml>

¹⁷³⁸ <https://github.com/rook/rook/tree/vv1.11.9/Documentation/Helm-Charts/operator-chart.md>

¹⁷³⁹ <https://raw.githubusercontent.com/rook/rook/v1.11.9/v%25s/deploy/charts/rook-ceph/values.yaml>

¹⁷⁴⁰ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/rook-ceph/1.11.9/defaults/cm.yaml>

¹⁷⁴¹ <https://github.com/rook/rook/tree/vv1.11.9/Documentation/Helm-Charts/ceph-cluster-chart.md>

¹⁷⁴² <https://raw.githubusercontent.com/rook/rook/v1.11.9/v%25s/deploy/charts/rook-ceph-cluster/values.yaml>

¹⁷⁴³ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/rook-ceph-cluster/1.11.9/defaults/cm.yaml>

¹⁷⁴⁴ <https://artifacthub.io/packages/helm/bitnami/thanos/12.8.3>

¹⁷⁴⁵ <https://artifacthub.io/packages/helm/bitnami/thanos/12.8.3?modal=values>

¹⁷⁴⁶ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/thanos/12.8.3/defaults/cm.yaml>

¹⁷⁴⁷ <https://artifacthub.io/packages/helm/traefik/traefik/23.1.0>

¹⁷⁴⁸ <https://artifacthub.io/packages/helm/traefik/traefik/23.1.0?modal=values>

¹⁷⁴⁹ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/traefik/23.1.0/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Traefik ForwardAuth	traefik-forward-auth	0.3.10	<ul style="list-style-type: none"> chart: 0.3.10¹⁷⁵⁰ traefik-forward-auth: 3.1.0 	Link ¹⁷⁵¹	Link ¹⁷⁵²
Velero	velero	4.1.1	<ul style="list-style-type: none"> chart: 4.1.1¹⁷⁵³ velero: 1.11.0 	Link ¹⁷⁵⁴	Link ¹⁷⁵⁵

12.3.7 DKP 2.6.2 Known Issues and Limitations

The following items are known issues with this release.

12.3.7.1 Known Issues

12.3.7.1.1 AWS `additionalTags` cannot contain spaces

Due to an upstream bug in the `cluster-api-provider-aws` component, it is not possible to specify tags with spaces in their name in the `additionalTags` section of an `AWSCluster`. If you have any tags like this during an upgrade of the `capi-components`, you may receive a validation error, and will need to remove any such tags. This issue will be corrected in a future DKP release.

12.3.7.1.2 AWS Custom AMI Required for Kubernetes Version

Previous versions of DKP would default to using upstream AMIs published by the CAPA (Cluster API AWS) project when building AWS clusters if you did not specify your own AMI. However, those images are not currently available for the Kubernetes version used in the 2.6.2 patch release.

As a result, starting with this release of DKP, the behavior of the DKP `create cluster aws` command has been changed. It no longer defaults to using the upstream AMIs and instead requires that you specify an AMI built using [Konvoy Image Builder \(KIB\)](#) (see page 1468), or by explicitly requesting that it use the upstream images.

¹⁷⁵⁰ <https://artifacthub.io/packages/helm/mesosphere/traefik-forward-auth/0.3.10>

¹⁷⁵¹ <https://artifacthub.io/packages/helm/mesosphere/traefik-forward-auth/0.3.10?modal=values>

¹⁷⁵² <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/traefik-forward-auth/0.3.10/defaults/cm.yaml>

¹⁷⁵³ <https://artifacthub.io/packages/helm/vmware-tanzu/velero/4.1.1>

¹⁷⁵⁴ <https://artifacthub.io/packages/helm/vmware-tanzu/velero/4.1.1?modal=values>

¹⁷⁵⁵ <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.6.2/services/velero/4.1.1/defaults/cm.yaml>

For more information on using a custom AMI in cluster creation or during the upgrade process, refer to these topics:

- [AWS: Create the Management Cluster](#) (see page 298)
- [Enterprise: Upgrade the Management Cluster Kubernetes Version](#) (see page 1583)

12.3.7.1.3 Use Static Credentials to Provision an Azure Cluster

Only static credentials can be used when [provisioning an Azure cluster](#) (see page 1166).

12.3.7.1.4 Containerd File Limit Issue

This version of DKP uses containerd 1.6.17. The systemd unit for containerd 1.6.17 provided upstream removes all file number limits (`LimitNOFILE=infinity`). In our testing, we found that removing these limits broke some IO-sensitive applications such as Rook Ceph and HAProxy. Because of this, the KIB version included in this release sets the `LimitNOFILE` value in the containerd systemd unit to the value (`1048576`), which was used in previous containerd 1.4.13 version releases.

12.3.7.1.5 Intermittent Error Status when Creating EKS Clusters in the UI

When provisioning an EKS cluster through the UI, you might receive a brief error state because the EKS cluster can sporadically lose connectivity with the management cluster, which results in the following symptoms:

- The UI shows the cluster is in an error state.
- The kubeconfig generated and retrieved from Kommander ceases to work.
- Applications created on the management cluster may not be immediately federated to managed EKS clusters.

After a few moments, the error self-resolves, without any action on your part. A new kubeconfig generated and retrieved from Kommander then works properly, and the UI shows that it is working again. In the meantime, you can continue to use the UI to work on the cluster such as deploy applications, create projects, and add roles.

12.3.7.1.6 Limitations to Disk Resizing in vSphere

The DKP CLI flags `--control-plane-disk-size` and `--worker-disk-size` are unable to resize the root file system of VMs, which are created using OS images. The flags work by resizing the primary disk of the VM. When the VM boots, the root file system expands to fill the disk, but that expansion does not work for some file systems. For example, for file systems contained in an LVM Logical Volume. To ensure your root file system has the size you expect, see <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/273384734/2.6%2BCreate%2Ba%2BvSphere%2BBase%2BOS%2BImage#Disk-Size> (see page 0).

12.3.7.1.7 Error Status in Grafana Logging Dashboard with EKS Clusters

Currently, it is not possible to use FluentBit to collect [Admin-level logs](#) (see page 905) on a managed EKS cluster.

If you have these logs enabled, the following message appears when you access the Kubernetes Audit Dashboard in the Grafana Logging Dashboard:

```
Cannot read properties of undefined (reading '0')
```

12.3.7.1.8 Spark Image Upstream Removal

Google has removed the images from their [Spark registry](https://gcr.io/spark-operator/spark-py:v3.1.1)¹⁷⁵⁶ upstream `gcr.io/spark-operator/spark-py:v3.1.1`. Therefore, you are not able to pull these images for non-air-gapped environments. Refer also to our [Deprecations \(see page 1835\)](#) notice for more information.

12.3.7.1.9 Rook Ceph Install Error

An issue might emerge when installing `rook-ceph` on vSphere clusters using RHEL operating systems. This issue occurs during initial installation of rook-ceph, causing the object store used by Velero and Grafana Loki, to be unavailable. If the installation of Kommander component of DKP is unsuccessful due to `rook-ceph` failing, you might need to apply the following workaround:

1. Run this command to check if the cluster is affected by this issue.

```
kubectl describe CephObjectStores dkp-object-store -n kommander
```

2. If this output appears, the workaround needs to be applied so continue with the next step. If you do not see this output, you can stop at this step.

```
Name:          dkp-object-store
Namespace:     kommander
...
Warning ReconcileFailed      7m55s (x19 over 52m)
rook-ceph-object-controller failed to reconcile CephObjectStore
"kommander/dkp-object-store". failed to create object store deployments:
failed
to configure multisite for object store: failed create ceph multisite for
object-store ["dkp-object-store"]: failed to commit config changes after
creating multisite config for CephObjectStore "kommander/dkp-object-store":
failed to commit RGW configuration period changes%(EXTRA []string=[]):
signal: interrupt
```

3. Kubectl exec into the `rook-ceph-tools` pod.

¹⁷⁵⁶ <http://gcr.io/spark-operator/spark-py:v3.1.1>

```
export WORKSPACE_NAMESPACE=<workspace namespace>
CEPH_TOOLS_POD=$(kubectl get pods -l app=rook-ceph-tools -n $
{WORKSPACE_NAMESPACE} -o name)
kubectl exec -it -n ${WORKSPACE_NAMESPACE} $CEPH_TOOLS_POD bash
```

4. Run these commands to set `dkp-object-store` as the default zonegroup.

i NOTE: The `period update` command may take a few minutes to complete

```
radosgw-admin zonegroup default --rgw-zonegroup=dkp-object-store
radosgw-admin period update --commit
```

5. Next, restart the `rook-ceph-operator` deployment for the `CephObjectStore` to be reconciled.

```
kubectl rollout restart deploy -n${WORKSPACE_NAMESPACE} rook-ceph-operator
```

6. After running the commands above, the `CephObjectStore` should be `Connected` once the `rook-ceph` operator reconciles the object (this may take some time).

```
kubectl wait CephObjectStore --for=jsonpath='{.status.phase}'=Connected dkp-
object-store -n ${WORKSPACE_NAMESPACE} --timeout 10m
```

12.3.7.1.10 Kommander Installation Configuration File Changes

In this release, the `kommander-ui`, which provides the DKP Dashboard, is now deployed in the same manner as the other Platform Applications. Also, the `ai-navigator-app`, which provides the AI Navigator, is deployed in the same manner. Therefore, the contents of the default Kommander installation configuration file, which is the file produced by the `dkp install kommander --init` command, have changed. If you installed previous versions of DKP using a customized Kommander installation configuration file, we recommend:

1. Generate a new template for this release.
2. Reapply your customizations rather than reusing a file created by older DKP versions.
3. The `ai-navigator-app` will be enabled by default in the new configuration file

12.3.7.1.11 FluentD Logging Operator Error

An error may occur with the FluentD Operator not appearing when installing Kommander on a GCP Environment.

This is due to a limitation of the GCP API. As a result, you need to rename the Fluentd buffer name because the default name exceeds 63 characters and disable buffer volume metrics for FluentD in order for it to operate properly.

Follow these procedures to resolve this issue for the Management Cluster, and Managed or Attached Clusters.

12.3.7.1.11.1 Management Cluster on GCP with the Logging Stack enabled

1. Create the `kommander` namespace:

```
kubectl create namespace kommander
```

2. Create the logging-operator-logging-overrides ConfigMap

```
cat <<EOF | kubectl apply -n kommander -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: logging-operator-logging-overrides
data:
  values.yaml: |
    fluentd:
      # disable buffer metrics
      bufferVolumeMetrics: null
      bufferStorageVolume:
        pvc:
          source:
            # update this name to make the PV <= 63 chars long
            # logging-operator-logging-buf-logging-operator-logging-fluentd-0
            claimName: buf
          spec:
            accessModes:
              - ReadWriteOnce
            resources:
              requests:
                storage: 10Gi
            volumeMode: Filesystem
EOF
```

3. Continue installing Kommander as normal.

12.3.7.1.11.2 Managed/Attached Cluster in GCP with Logging Stack enabled

1. Set the `WORKSPACE_NAMESPACE` environment variable to the name of your workspace's namespace:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```

2. Create the logging-operator-logging-overrides ConfigMap on the managed or attached cluster prior to enabling the logging-operator application:

```
cat <<EOF | kubectl apply -n ${WORKSPACE_NAMESPACE} -f -
apiVersion: v1
kind: ConfigMap
metadata:
  name: logging-operator-logging-overrides
data:
  values.yaml: |
    fluentd:
      # disable buffer metrics
      bufferVolumeMetrics: null
      bufferStorageVolume:
        pvc:
          source:
            # update this name to make the PV <= 63 chars long
            # logging-operator-logging-buf-logging-operator-logging-fluentd-0
            claimName: buf
          spec:
            accessModes:
              - ReadWriteOnce
            resources:
              requests:
                storage: 10Gi
              volumeMode: Filesystem
EOF
```

3. Proceed with installing the logging-operator as normal.

12.3.7.1.12 Generate a Token or Enable the Konvoy Credentials Plugin for Proxied Clusters

If you have attached a [network-restricted cluster](#) (see page 81) and enabled a [proxied access](#) (see page 817) to make its resources available through the Management cluster, the “Generate Token” and “Konvoy credentials plugin instructions” options required for user authentication in the DKP UI require further steps.

Expand the set of instructions depending on the authentication configuration you want to configure:

12.3.7.1.12.1 A. Generate a Token

Generate Token

- I. Execute the commands established in the “Generate Token” page.

- You can find the “Generate Token” page in the DKP UI by selecting your username in the top right corner of any page.

II. Prepare your Environment

- Write down the name of your credentials. You can obtain this name by accessing the **DKP UI** > username in the top right corner > **Generate token** page. The credentials are displayed in the `kubectl config set-credentials` command, next to `set-credentials`.

In the following example, `jane_doe-`

`az189abcd7ab34a2abc0ab845a67a7ea-0123456789.us-west-2.elb.amazonaws.com` would be the name of your credentials:

```
kubectl config set-credentials jane_doe-
az189abcd7ab34a2abc0ab845a67a7ea-0123456789.us-west-2.elb.amazonaws.com \
--token=[...]
```

- Set the `CONFIG_CREDENTIALS_NAME` environment variable to the credentials you obtained in the previous step:

```
CONFIG_CREDENTIALS_NAME=<CREDENTIALS_NAME>
```

- Set `WORKSPACE_NAMESPACE` to the namespace of your network-restricted cluster:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```

- Set `TUNNEL_PROXY_NAME` to the name of the TunnelProxy:

```
TUNNEL_PROXY_NAME=<...>
```

- Set `DEFAULT_CONFIG_PATH` to the path file of the `kubeconfig` where you executed the “Generate Token” steps:

```
DEFAULT_CONFIG_PATH=<path_of_cluster's_kubeconfig>
```

- With your `kubeconfig` set to the **Management cluster**, set the `TUNNEL_PROXY_DOMAIN` environment variable to the proxied access URL:

i See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.

```
TUNNEL_PROXY_DOMAIN=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n
$WORKSPACE_NAMESPACE -o template='{{ .status.clusterProxyDomain }}')
```

III. Reconfigure the Token

Select one of the following options according to the certificate you configured for the proxied access:

A. Clusters with ACME-supported Certificates for the TunnelProxy

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
${TUNNEL_PROXY_DOMAIN}/dkp/api-server --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --cluster=${TUNNEL_PROXY_NAME} --
user=${CONFIG_CREDENTIALS_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
```

B. Clusters with Self-Signed Certificates for the TunnelProxy

1. With your kubeconfig set to the **Management cluster**, retrieve the `ca.crt` (not needed in case of ACME certificates):

i See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.

```
LABEL_KEY=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n $WORKSPACE_NAMESPACE
-o template='{{ .status.reverseProxyReleaseName }}')
SECRET_NAME=$(kubectl get ingress -n $WORKSPACE_NAMESPACE -l app.kubernetes.io/
instance=$LABEL_KEY -o jsonpath='{.items[0].spec.tls[0].secretName}')
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
kubectl get secrets $SECRET_NAME -n $WORKSPACE_NAMESPACE -o
template='{{ index .data "ca.crt" | base64decode }}' > ${HOME}/.kube/certs/$
{TUNNEL_PROXY_NAME}/ca.crt
```

2. Configure the kubeconfig access via the proxied access:

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
${TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/
certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --cluster=${TUNNEL_PROXY_NAME}
--user=${CONFIG_CREDENTIALS_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
```

```
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig ${
DEFAULT_CONFIG_PATH}
```

C. Clusters with Custom Certificates for the TunnelProxy

1. Copy the `ca.crt` used to sign your certificates to the `kubeconfig` directory:

```
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
cp ca.crt ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}/ca.crt
```

2. Configure the `kubeconfig` access via the proxied access:

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
{TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/
certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --cluster=${TUNNEL_PROXY_NAME}
--user=${CONFIG_CREDENTIALS_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig $
{DEFAULT_CONFIG_PATH}
```

After the commands complete, you can perform `kubectl` actions on the network-restricted cluster.

12.3.7.1.12.2 B. Konvoy credentials plugin instructions

Konvoy credentials plugin instructions

- I. Execute the commands provided in the “Konvoy credentials plugin instructions” page > “Method 1” section.



You can find a link to the “Konvoy credentials plugin instructions” page in the DKP UI when logging in to your environment, or by adding `/token/plugin` to the DKP UI URL of your environment.

II. Prepare your Environment

1. Set `WORKSPACE_NAMESPACE` to the namespace of your network-restricted cluster:

```
export WORKSPACE_NAMESPACE=<workspace namespace>
```


2. Set `TUNNEL_PROXY_NAME` to the name of the TunnelProxy:

```
TUNNEL_PROXY_NAME=<...>
```

3. Set `DEFAULT_CONFIG_PATH` to the path of the kubeconfig file you downloaded from the dashboard:

```
DEFAULT_CONFIG_PATH=<path_of_cluster's_kubeconfig>
```

4. With your kubeconfig set to the **Management cluster**, set the `TUNNEL_PROXY_DOMAIN` environment variable to the proxied access URL:

 See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.

```
TUNNEL_PROXY_DOMAIN=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n $WORKSPACE_NAMESPACE -o template='{{ .status.clusterProxyDomain }}')
```

III. Reconfigure the Token


Select one of the following options according to the certificate you configured for the proxied access:

A. Clusters with ACME-supported Certificates for the TunnelProxy

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://$
${TUNNEL_PROXY_DOMAIN}/dkp/api-server --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --user default-profile --kubeconfig $
${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
```

B. Clusters with Self-Signed Certificates for the TunnelProxy

1. With your kubeconfig set to the **Management cluster**, retrieve the `ca.crt` (not needed in case of ACME certificates):

 See [Provide Context for Commands with a kubeconfig File \(see page 87\)](#) for information on setting the kubeconfig to a cluster.

```
LABEL_KEY=$(kubectl get tunnelproxy $TUNNEL_PROXY_NAME -n $WORKSPACE_NAMESPACE
-o template='{{ .status.reverseProxyReleaseName }}')
```

```
SECRET_NAME=$(kubectl get ingress -n $WORKSPACE_NAMESPACE -l app.kubernetes.io/instance=$LABEL_KEY -o jsonpath='{.items[0].spec.tls[0].secretName}')
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
kubectl get secrets $SECRET_NAME -n $WORKSPACE_NAMESPACE -o
template='{{ index .data "ca.crt" | base64decode }}' > ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}/ca.crt
```

2. Configure the `kubeconfig` access via the proxied access:

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://${TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --user default-profile --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
```

C. Clusters with Custom Certificates for the `TunnelProxy`

1. Copy the `ca.crt` used to sign your certificates to the `kubeconfig` directory:

```
mkdir -p ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}
cp ca.crt ${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}/ca.crt
```

2. Configure the `kubeconfig` access via the proxied access:

```
kubectl config set-cluster ${TUNNEL_PROXY_NAME} --server=https://${TUNNEL_PROXY_DOMAIN}/dkp/api-server --certificate-authority=${HOME}/.kube/certs/${TUNNEL_PROXY_NAME}/ca.crt --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config set-context ${TUNNEL_PROXY_NAME} --user default-profile --kubeconfig ${DEFAULT_CONFIG_PATH}
kubectl config use-context ${TUNNEL_PROXY_NAME} --kubeconfig ${DEFAULT_CONFIG_PATH}
```



“Method 2” is not available for network-restricted clusters.

After the commands complete, you can perform `kubectl` actions on the network-restricted cluster.

12.3.7.1.13 Known Common Vulnerabilities and Exposures (CVE)

Starting with DKP 2.6, Catalog apps are scanned for CVEs. However, only the CVEs for app versions that are compatible with the default Kubernetes version, currently 1.26.6, are mitigated. For more information about the known CVEs for compatible catalog apps, see [D2iQ Security Updates \(see page 1851\)](#).

12.3.7.1.14 Potential Logging Interruption during Upgrade

The configuration setting `enableRecreateWorkloadOnImmutableFieldChange` is enabled by default on the `logging-operator`. This means the operator automatically triggers a recreation of the resource if any immutable fields are changed on the underlying fluent-bit/fluentd resources. Previously, any changes to immutable fields would fail silently and the failures would only be observable in the `logging-operator` logs. This was done to handle a breaking change introduced in the `logging-operator`, which added an additional selector label on the `fluent-bit` DaemonSet.

Because of this breaking change, during an upgrade to 2.6, `fluent-bit` is restarted, and some log data in the `fluent-bit` buffer might be lost if `fluentd` is unavailable while `fluent-bit` attempts to flush its buffer prior to the pods being terminated. To prevent issues related to data loss, we recommend you run `>1 fluentd pod`. Alternatively, Fluent Bit can also be configured to use a `hostPath` volume to store the buffer information, so it can be picked up again when Fluent Bit restarts. For more information about how to change the host path, see [Fluent Bit log collector¹⁷⁵⁷](#).

12.3.7.1.15 Gatekeeper Uninstallation Error

If you choose to disable Gatekeeper, you can run into an error where the app is still present on your clusters.

For pre-existing attached or managed clusters with Gatekeeper installed when they don't want it to be, you need to manually cleanup Gatekeeper on those clusters, post upgrade.

Follow these steps to manually remove Gatekeeper:

1. For every **Attached/Managed cluster** that has Gatekeeper installed that you have disabled it on otherwise (via appdeployment), run the following commands to remove it:

NOTE: This needs to be run on the attached/managed cluster using the correct kubeconfig. Set your namespace of your **attached** cluster's workspace on your attached cluster:

```
export WORKSPACE_NAMESPACE=<workspace-name>
```

2. Delete the kustomizations:

¹⁷⁵⁷ <https://kube-logging.dev/docs/logging-infrastructure/fluentbit/#hostpath-volumes-for-buffers-and-positions>

```
kubectl --kubeconfig <attached-cluster-kubeconfig> delete kustomizations -n $
{WORKSPACE_NAMESPACE} gatekeeper-constraint-templates gatekeeper-constraints
gatekeeper-release
```

12.3.8 DKP 2.6.2 Customer Incidents and Resolved Issues

The following resolved incidents are included in this release.

12.3.8.1 CAPI Move Causes Worker Kubelet Issue (CAPPP)

D2IQ-97212

In a Pre-provisioned DKP cluster, you may have experienced a worker node in a `not ready` condition after performing a `dkp move capi-resources` to make the cluster self-managed. This was due to a race condition in the CAPPP (pre-provisioned) provider, which has been fixed.

12.3.8.2 Upgrade Could Fail for EKS or AKS Clusters

D2IQ-100149

When upgrading from prior versions of DKP, the `dkp update nodepool` command could fail when updating EKS or AKS clusters. The bug has been fixed.

12.3.8.3 CAPA AMIs are Not Available for the Kubernetes Version used in DKP 2.6.2

D2IQ-100175

AWS AMIs for some Kubernetes versions may not be available. We recommend using [Konvoy Image Builder](#) (see [page 1464](#)) (KIB) to create a custom AMI.

- Do not use the upstream values as defaults for the `--ami-*` flags.
- If you specifies any of these flags for AMI ID lookup, specify all of them: `--ami-base-os`, `--ami-format`, and `--ami-owner`.

12.3.9 DKP 2.6.2 Deprecations

The following items are deprecated or removed from this version of DKP.

12.3.9.1 Deprecated `dkp push` Subcommand and Flag

To improve consistency, simplify, and streamline DKP CLI commands, this version of DKP deprecates the `dkp push image-bundle` command in favor of `dkp push bundle`.

Deprecated subcommand and flag	New subcommand and flag
<code>dkp push image-bundle --image-bundle</code>	<code>dkp push bundle --bundle</code>

You can still run the `dkp push image-bundle` with this version of DKP, but it will be removed in future DKP versions.

If you have scripts, GitHub Actions, or other cluster creation and upgrade automations, D2iQ recommends updating to the new subcommand and flag as soon as possible.

12.3.9.2 Catalog Applications Compatibility

From DKP 2.6.2, the following versions of Catalog applications are deprecated:

- kafka-operator-0.20.0
- kafka-operator-0.20.2
- spark-operator-1.1.6
- zookeeper-operator-0.2.13

If you plan on upgrading to DKP 2.6.2, ensure that you upgrade these applications to the latest compatible version.

For more information, see [Workspace DKP Catalog Applications - Compatibility](#) (see page 667).

12.3.9.3 Spark Operator to be Removed in DKP 2.7

Starting from DKP 2.6.2, D2iQ will discontinue support for all versions of Spark Operator in the DKP Catalog, because the upstream operator is no longer maintained. As a result, the Spark Operator will no longer be available for use in future DKP releases.

Spark Operator remains available in DKP 2.6.1 and previous versions.

For instructions on how to keep Spark on your cluster after upgrading, expand this section:

How to keep Spark for existing instances of Spark, but it is unmanaged by DKP or Flux

1. Set the `WORKSPACE_NAMESPACE` variable to the namespace where the spark-operator is deployed:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

2. Set `suspend: true` on the `cluster` Kustomization

NOTE: This needs to be reverted in the last step.

```
kubectl -n ${WORKSPACE_NAMESPACE} patch kustomization cluster --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": true}]'
```

3. Set `suspend: true` and `prune: false` on the `spark-operator` Kustomizations

```
export SPARK_OPERATOR_APPDEPLOYMENT_NAME=<spark operator AppDeployment name>
export SPARK_OPERATOR_APPDEPLOYMENT_VERSION=<spark operator AppDeployment
version>
kubectl -n ${WORKSPACE_NAMESPACE} patch kustomization $
{SPARK_OPERATOR_APPDEPLOYMENT_NAME} --type='json' -p='[{"op": "replace",
"path": "/spec/suspend", "value": true}, {"op": "replace", "path": "/spec/
prune", "value": false}]'
kubectl -n ${WORKSPACE_NAMESPACE} patch kustomization spark-operator-$
{SPARK_OPERATOR_APPDEPLOYMENT_VERSION}-defaults --type='json' -p='[{"op":
"replace", "path": "/spec/suspend", "value": true}, {"op": "replace", "path": "/
spec/prune", "value": false}]'
```

4. Update the `dkp-catalog-applications` [GitRepository with the new version](#)¹⁷⁵⁸
5. Unsuspend the `cluster` Kustomization
NOTE: If the Kustomization is left suspended, Kommander will be unable to function properly.

```
kubectl -n ${WORKSPACE_NAMESPACE} patch kustomization cluster --type='json' -p='
[{"op": "replace", "path": "/spec/suspend", "value": false}]'
```

12.3.9.3.1 Remove Spark from the UI

Follow these steps to remove the `spark-operator` AppDeployment from the DKP UI:

1. From the top menu bar, select your target workspace.
2. Select **Applications** from the sidebar menu.
3. Select the three dot button from the bottom-right corner of the Spark application tile, and then select **Uninstall**.
4. Select **Save**.

OR

12.3.9.3.2 Remove Spark with CLI

1. Execute the following command to get the `WORKSPACE_NAMESPACE` of your workspace:

```
dkp get workspaces
```

And copy the values under the `NAMESPACE` column for your workspace.

¹⁷⁵⁸ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/369397443/2.7+Enterprise+Upgrade+Workspace+DKP+Catalog+Applications#update-gitrepo>

2. Export the `WORKSPACE_NAMESPACE` variable:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAME>
```

3. Run the following command to delete the `spark-operator` AppDeployment:

```
kubectl delete AppDeployment <spark operator appdeployment name> -n $  
{WORKSPACE_NAMESPACE}
```

This results in the `spark-operator` Kustomizations being deleted, but the `HelmRelease` and default `ConfigMap` will remain cluster.

From here, you can continue to manage the `spark-operator` via the `HelmRelease`.

More information on Spark can be found in this section of the documentation [Enterprise: Upgrade Workspace Catalog Applications](#) (see page 1574).

13 AI Navigator

First released in DKP 2.6.0, D2iQ engineered the AI Navigator with the power of AI. D2iQ then trained the chatbot on both our product documentation and Support knowledge base. This AI chatbot ensures real-time, interactive communication to answer a wide range of user queries, spanning basic instructions to complex functionalities.

Use of AI Navigator requires that you are logged in to a licensed copy of DKP version 2.6.0 or later, and that you [accept the User Agreement](#) (see page 1839). The AI chatbot is offered for Internet-connected environments regardless of DKP license type.



To help support heightened security in the public sector, defense, and network-restricted customer environments, AI Navigator is **not** offered for air-gapped implementations.

- [AI Navigator User Agreement and Guidelines](#) (see page 1839)
- [AI Navigator in Installations and Upgrades](#) (see page 1841)
- [Access the AI Navigator](#) (see page 1842)
- [How to Create Good Queries](#) (see page 1843)

13.1 AI Navigator User Agreement and Guidelines

Introducing the AI Navigator - engineered with the power of AI and models that include our product documentation and Support knowledge base, this chatbot ensures real-time, interactive communication to answer a wide range of user queries. The chatbot is offered for Internet-connected environments regardless of DKP license type.

Use of the AI Navigator requires that you are logged in to a licensed copy of DKP 2.6.0 or later.

You can [access the AI Navigator](#) (see page 1842) application from the icon in the lower right corner of the DKP UI.

13.1.1 Accept the User Agreement

When you first access the DKP AI Navigator, it displays the user agreement for acceptance:

Acknowledgement

This chatbot is a tool through which I can obtain information or assistance for the products or services offered by D2iQ, Inc.

Use of this chatbot is governed by the terms of the [D2iQ privacy policy](#) and I agree not to share any personal, sensitive, or confidential information through this chatbot.

This service is powered by Microsoft's Azure OpenAI and is governed by [their privacy policy](#).

Yes, I agree

No, not now

Any use of the chatbot is pursuant to your acceptance of this agreement. If you select **No, not now** to decline, you cannot use AI Navigator.

13.1.2 Guidelines for Using AI Navigator

As with all AI engines, you should **NOT** enter sensitive information, including:

- Full names, addresses, or other personally identifying information (PII)
- Technical secrets, such as:
 - Actual server names
 - Real IP addresses
 - Private keys
 - Access tokens
 - ...any other technical details that might be useful to malevolent actors
- Government ID numbers
- Corporate, financial, or investment information
- Company Confidential information
- Company credentials, including actual usernames or passwords
- Secured or classified information

...and so on.



Your company or organization may have specific rules about the kinds of prompts you should enter. Refer to your company or organization's specific guidelines for additional information.

You can [read more here \(see page 1843\)](#) on creating good prompts for use with AI Navigator.

13.1.2.1 Next Topic

[AI Navigator in Installations and Upgrades](#) (see page 1841)

13.2 AI Navigator in Installations and Upgrades

Use of AI Navigator requires that you are logged in to a licensed copy of DKP version 2.6.0 or later. The AI chatbot is offered for Internet-connected environments regardless of DKP license type.



To help support heightened security in public sector, defense, and network-restricted customer environments, AI Navigator is **not** offered for air-gapped implementations.

13.2.1 Install DKP 2.6.1 and Later Releases

When installing DKP 2.6.1 or later releases for the first time, you have a choice as to whether you install AI Navigator. To install the chatbot application, perform the installation according to the procedures for your chosen infrastructure provider, and DKP installs AI Navigator by default.

To disable AI Navigator, [modify the Kommander configuration file after generating it](#)¹⁷⁵⁹, to disable the `ai-navigator-app` entry in the `apps:` section of the file.

You can find more information about the Kommander configuration file in the topic, [DKP Kommander Configuration Reference](#)¹⁷⁶⁰.

13.2.2 Upgrade DKP 2.6.0 to 2.6.1

By default, DKP installs the AI Navigator application as part of installing the Kommander component. DKP 2.6.0 installed AI Navigator by default as an experimental technical preview. Beginning in DKP 2.6.1, you have a choice. To disable AI Navigator, [modify the Kommander configuration file after generating it](#)¹⁷⁶¹, to disable the `ai-navigator-app` entry in the `apps:` section of the file.

13.2.2.1 For Upgrade in DKP Enterprise Environments

Under the DKP Enterprise [upgrade](#) (see page 1568) section, expand the section for “For non-air-gapped environments.” Follow the instructions to add a CLI flag to the listed `dkp upgrade kommander` command to disable the chatbot.

¹⁷⁵⁹ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/233177089/Install+Kommander+in+a+Non-air-gapped+Environment>

¹⁷⁶⁰ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/94765060/DKP+Kommander+Configuration+Reference>

¹⁷⁶¹ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/233177089/Install+Kommander+in+a+Non-air-gapped+Environment>

13.2.2.2 For Upgrade in DKP Essential Environments

Under the DKP Enterprise [upgrade](#) (see page 1593) section, expand the section for “For non-air-gapped environments.” Follow the instructions to add a CLI flag to the listed `dkp upgrade kommander` command to disable the chatbot.

13.2.2.3 Next Topic

[Access the AI Navigator](#) (see page 1842)

13.3 Access the AI Navigator

If you chose to install the AI Navigator, you can easily access the application in environments running DKP 2.6.0 or later releases. The chatbot is installed by default in Internet-connected environments.



DKP does not offer the AI Navigator in air-gapped environments.

13.3.1 Open the AI Navigator

You can access the AI Navigator application from the icon in the lower right corner of the DKP user interface.

Welcome to DKP Enterprise, D2iQ's solution for managing multiple Kubernetes clusters

DKP Enterprise provides oversight, governance and visibility across your organization's on-premise and cloud footprint, in addition to offering a multi-cluster and multi-tenant experience. From initial setup to final configuration, DKP Enterprise streamlines your Kubernetes operations at every level.

[Get Started](#) [Learn More About DKP](#) Do not show again

Management Cluster	
Active	v1.26.6
CPU	Requests 82%
Mem	Limits 201%
Disk	Usage 13%

Clusters (1)	
Resources	Cost
Clusters with Errors	0
Clusters with Warnings	1
Total Clusters	1

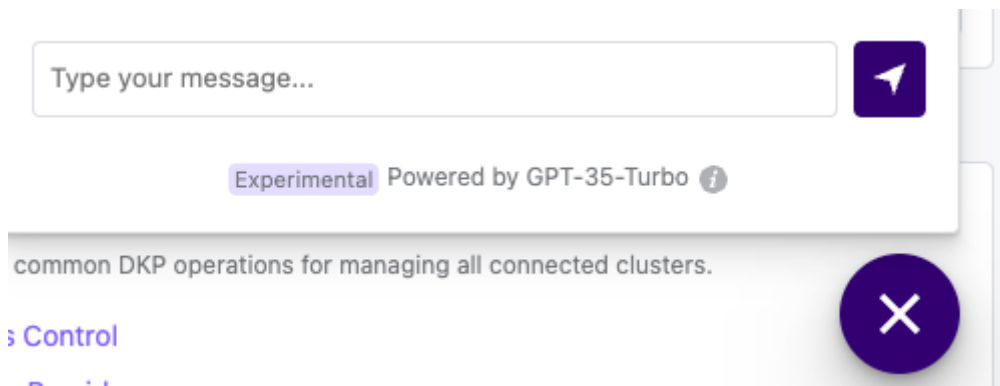
Workspaces (2)	
Default Workspace	N/A
Management Cluster Workspace	\$25.25

When you open the chatbot for the first time, it prompts you with a [user agreement](#) (see page 1842). You must agree to the terms to use AI Navigator. If you decline, the chatbot window closes and you cannot use it.

AI Navigator maintains your query history for the duration of your browser session, whether or not you close the AI Navigator application.

13.3.2 Close the AI Navigator

To close the AI Navigator, select the **Close** button just below the prompt entry field:



13.3.2.1 Next Topic

[How to Create Good Queries](#) (see page 1843)

13.4 How to Create Good Queries

Prompt engineering is the process of creating a chatbot prompt that returns the data you really want. For AI chatbots trained on very large language models, there are some tips you can use to get better information with fewer prompt attempts. Some of those techniques are useful with the AI Navigator.

When you create a prompt, an AI chatbot breaks down your entry into discrete parts to help it search its model. The more precise you can make your prompt, the better the chatbot is at returning the information you wanted.

Another technique for getting to the desired information is called fine-tuning - the limiting of the model - the database - that the chatbot has to search. DKP fine-tunes the AI Navigator model by adding both the DKP and Insights documentation, and the DKP Support Knowledge Base. This helps to keep answers focused and fast.

13.4.1 What Goes in a Prompt

You can include straight text, inline code snippets such as kubectl commands, and code block snippets, within a limit of 4096 characters.

AI Navigator maintains your query history for the duration of your browser session, whether or not you close the AI Navigator application.

13.4.1.1 Inline Commands or Code Snippets

Use single back-tick marks to enclose a word, phrase, or command, for example:

What are the parameters I can use with the `dkp create cluster` command?

...creates this output in the chatbot window:



Hi, what can I help you with?

What are the parameters I can use with the `dkp create cluster` command?

The `dkp create cluster` command has several parameters that you can use. Here are the parameters available:

- `--cluster-name` : Specifies the name of the cluster.
- `--additional-tags` : Specifies additional tags for the cluster.
- `--with-aws-bootstrap-credentials` : Specifies whether to use AWS bootstrap credentials.
- `--kubernetes-version` : Specifies the version of Kubernetes to use.
- `--etcd-version` : Specifies the version of etcd to use.
- `--kubernetes-image-repository` : Specifies the image repository for Kubernetes.
- `--self-managed` : Specifies whether the cluster is self-managed.

13.4.1.2 Code Blocks

Use a triple back-tick to open a code block for entry, in a manner similar to many popular chat and collaboration applications, like this:

Where **do** I need to use **this** code:

```

```
dkp create cluster aws \
 --http-proxy=127.0.0.1
 --https-proxy=127.0.0.2
```

...creates this entry and answer pair:

Where do I need to use this code:

```
dkp create cluster aws \
 --http-proxy=12.7.0.0.1\
 --https-proxy=127.0.0.1
```

The code "dkp create cluster aws --http-proxy=12.7.0.0.1 --https-proxy=127.0.0.1" should be used when creating a Kubernetes cluster on AWS. It sets the HTTP and HTTPS proxies for the cluster.

- [Create a Cluster with HTTP/HTTPS Proxy](#)
- [Create a New Customized AWS Cluster](#)
- [AWS: Create the Management Cluster](#)
- [Create a New Air-gapped AWS Cluster](#)



### 13.4.2 Selected Prompt Examples

You can use the AI Navigator to help you with a variety of tasks related to creating and managing clusters, and prompts can be fairly complex, and not necessarily limited to DKP. Consider these examples:

#### 13.4.2.1 Evaluate a YAML Manifest

For this query:

Explain **this** yaml:



```

apiVersion: policy/v1
kind: PodDisruptionBudget
metadata:
 name: nvidia-critical-app
spec:
 maxUnavailable: 1
 selector:
 matchLabels:
 app: nvidia-critical-app

```

...the AI Navigator returns output like this:

This YAML file is defining a Pod Disruption Budget (PDB) **for** an application named "nvidia-critical-app".

A Pod Disruption Budget is a Kubernetes resource that specifies the minimum number of pods that must be available during a disruption, such as a node upgrade or maintenance. It helps ensure high availability of critical applications by preventing too many pods from being unavailable at the same time.

In **this** specific YAML, the PDB is configured with the following properties:

**maxUnavailable: 1:** This specifies that only one pod of the "nvidia-critical-app" can be unavailable at a time during a disruption. This ensures that at least one pod is always available to handle requests.

**selector:** This specifies the labels used to select the pods that are part of the "nvidia-critical-app". In **this case**, it selects the pods with the label `app: nvidia-critical-app`.

Overall, **this** YAML is creating a Pod Disruption Budget that allows only one pod of the "nvidia-critical-app" to be unavailable at a time during disruptions.

#### Links

-----

[Update Cluster Nodepools](<https://beta-docs.d2iq.com/dkp/2.6/update-nvidia-gpu-cluster-nodepools>)

[Setting Priority Classes in DKP Applications](<https://beta-docs.d2iq.com/dkp/2.6/setting-priority-classes-in-dkp-applications>)

[Project Quotas & Limit Ranges](<https://beta-docs.d2iq.com/dkp/2.6/project-quotas-limit-ranges>)

[Configure the Control Plane](<https://beta-docs.d2iq.com/dkp/2.6/configuring-the-control-plane>)

### 13.4.2.2 Updating a TLS Thumbprint

For this query:

How **do** I update a TLS Thumbprint?

...the AI Navigator returns output like this:

To update a TLS thumbprint in DKP, you need to update the following objects:

1. Update the vspherecluster object:

- Use the command: `kubectl patch vspherecluster <CLUSTER_NAME> --type=merge -p '{"spec": {"thumbprint": "<TLS_THUMBPRINT>"}}'`

2. Update the vsphere-config-secret secret in the vmware-system-csi namespace:

- Use the command: `kubectl patch secret vsphere-config-secret -n vmware-system-csi --type=json -p='[{"op" : "replace" ,"path" : "/data/csi-vsphere.conf" ,"value" : "<BASE64 Encoded>"}]'`

- Make sure to encode the value of csi-vsphere.conf in base64.

3. Update the vsphere-cloud-config configmap in the kube-system namespace:

- Use the command: `kubectl --kubeconfig <CLUSTER_NAME>-workload.conf patch cm vsphere-cloud-config -n kube-system --type=merge -p '{"data": {"vsphere.conf": "global:\n secretName: cloud-provider-vsphere-credentials\n secretNamespace: kube-system\n thumbprint: <TLS_THUMBPRINT>\nvcenter:\n <vCenter_Address>:\n datacenters:\n - 'dc1'\n secretName: cloud-provider-vsphere-credentials\n secretNamespace: kube-system\n server: '<vCenter_Address>'\n thumbprint: <TLS_THUMBPRINT>\n" }}'`

Replace <CLUSTER\_NAME>, <TLS\_THUMBPRINT>, and <BASE64 Encoded> with the actual values specific to your environment.

Links

-----

[Updating vCenter Server TLS Certificate Thumbprint in DKP - D2iQ](#)

[Updating vCenter Server TLS Certificate Thumbprint in DKP - D2iQ](#)

### 13.4.2.3 Describe Clock Skew in Ceph Cluster Status

For this query:

What is clock skew in Ceph cluster status?

...the AI Navigator returns output like this:

Clock skew refers to the difference in time between the clocks of different nodes in a Ceph cluster. In the context of a Ceph cluster, clock skew can affect the monitors, which are responsible **for** maintaining the cluster state. The monitors rely on a tight time alignment to coordinate with other Ceph daemons. If there is a significant clock skew between the monitor nodes, it can lead to issues in the cluster, such as inconsistent cluster maps and coordination problems. It is important to keep the time in-sync across nodes to ensure the proper functioning of the Ceph cluster.

Links

-----

How to detect **if** a clock skew is affecting Rook-Ceph monitors in DKP? - D2iQ  
Create **new** cluster fails to complete due to time drift - D2iQ

## 14 Access Documentation

The following sections describe how to access other versions or forms of our documentation.

### 14.1 Supported Documentation

You can access all n-2 supported documentation at [D2iQ Help Center](#)<sup>1762</sup>.

### 14.2 Archived Documentation

In accordance with our [version support policy](#) (see page 1853), we regularly archive older, unsupported versions of our documentation. At this time, this includes documentation for:

- DKP and DKP Insights 2.3 documentation at [D2iQ Documentation Archive](#)<sup>1763</sup>.
- Konvoy, Kommander, Kaptain, DC/OS and Dispatch at our [Public GitHub Repository](#)<sup>1764</sup>.

### 14.3 Download Documentation PDF

In an air-gapped environment, it is helpful to have a local PDF of the documentation. To download it, refer to the [Download Documentation](#) (see page 1849) page.

### 14.4 Download Documentation

This PDF contains the entire DKP documentation space. The file name contains the version and date created.

---

<sup>1762</sup> <http://docs.d2iq.com/>

<sup>1763</sup> <https://archive-docs.d2iq.com/>

<sup>1764</sup> <https://github.com/mesosphere/dcos-docs-site/tags>

DKP\_2.6.2\_3-18-24.pdf

[\(see page 1852\)](#)

## 15 D2iQ Security Updates

View CVE security scans and update information for D2iQ products. The following information describes mitigated or fixed CVEs found in DKP. The "Vulnerability IDs" column contains the list of known mitigated CVE ids.

| Product                  | Version | Vulnerability Reports                        | Vulnerability IDs                             |
|--------------------------|---------|----------------------------------------------|-----------------------------------------------|
| DKP                      | v2.7.0  | <a href="#">Download CSV</a> <sup>1765</sup> | <a href="#">Download File</a> <sup>1766</sup> |
| DKP Catalog Applications | v2.7.0  | <a href="#">Download CSV</a> <sup>1767</sup> | <a href="#">Download File</a> <sup>1768</sup> |
| DKP                      | v2.6.0  | <a href="#">Download CSV</a> <sup>1769</sup> | <a href="#">Download File</a> <sup>1770</sup> |
| DKP Catalog Applications | v2.6.0  | <a href="#">Download CSV</a> <sup>1771</sup> | <a href="#">Download File</a> <sup>1772</sup> |
| DKP                      | v2.5.0  | <a href="#">Download CSV</a> <sup>1773</sup> | <a href="#">Download File</a> <sup>1774</sup> |
| DKP                      | v2.4.0  | <a href="#">Download CSV</a> <sup>1775</sup> | <a href="#">Download File</a> <sup>1776</sup> |

This report is generated using the scanner Trivy v0.44.0 and the CVE database is updated to 21 Mar 24 12:13 UTC.

### 15.1 CVE Policy

At D2iQ, our commitment to providing secure software solutions is paramount. We understand the critical importance of promptly addressing and mitigating security vulnerabilities to ensure the safety and trust of our customers and partners. This document outlines our policies and procedures regarding CVEs (Common

<sup>1765</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160705\\_DKP\\_v2.7.0.csv](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160705_DKP_v2.7.0.csv)

<sup>1766</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160706\\_DKP\\_v2.7.0\\_ids.txt](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160706_DKP_v2.7.0_ids.txt)

<sup>1767</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160706\\_DKP\\_Catalog\\_Applications\\_v2.7.0.csv](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160706_DKP_Catalog_Applications_v2.7.0.csv)

<sup>1768</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160706\\_DKP\\_Catalog\\_Applications\\_v2.7.0\\_ids.txt](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160706_DKP_Catalog_Applications_v2.7.0_ids.txt)

<sup>1769</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160703\\_DKP\\_v2.6.0.csv](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160703_DKP_v2.6.0.csv)

<sup>1770</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160704\\_DKP\\_v2.6.0\\_ids.txt](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160704_DKP_v2.6.0_ids.txt)

<sup>1771</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160704\\_DKP\\_Catalog\\_Applications\\_v2.6.0.csv](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160704_DKP_Catalog_Applications_v2.6.0.csv)

<sup>1772</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160705\\_DKP\\_Catalog\\_Applications\\_v2.6.0\\_ids.txt](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160705_DKP_Catalog_Applications_v2.6.0_ids.txt)

<sup>1773</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160702\\_DKP\\_v2.5.0.csv](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160702_DKP_v2.5.0.csv)

<sup>1774</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160703\\_DKP\\_v2.5.0\\_ids.txt](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160703_DKP_v2.5.0_ids.txt)

<sup>1775</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160700\\_DKP\\_v2.4.0.csv](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160700_DKP_v2.4.0.csv)

<sup>1776</sup> [https://downloads.d2iq.com/dkp/cves/vulnerability\\_report\\_20240321\\_160702\\_DKP\\_v2.4.0\\_ids.txt](https://downloads.d2iq.com/dkp/cves/vulnerability_report_20240321_160702_DKP_v2.4.0_ids.txt)

Vulnerabilities and Exposures) to demonstrate our dedication to delivering software that is as secure as possible.

### 15.1.1 CVE Management:

Our procedure for managing CVE's is explained in the sections below.

#### 15.1.1.1 Scanning Policy:

- Our primary objective is to deliver software that is free from critical security vulnerabilities (CVEs).
- To achieve this, we conduct regular scans of our software components, including:
  - Kubernetes
  - D2iQ Platform applications (Traefik, Istio, ...)
  - D2iQ Catalog applications (*only* versions that are compatible with the default Kubernetes version supported with that DKP release, shown in our docs: [Workspace DKP Catalog Applications](#) (see page 666))
  - DKP Insights Add-on
- Scans are performed every 24 hours using the latest CVE database to identify and address potential vulnerabilities promptly. The scanner version and the CVE database version are published among the security scan results.

#### 15.1.1.2 Shipping Policy:

- Our goal is to ship software releases with no unmitigated Critical CVEs.
- For major and minor releases, we aim to ensure that there are no known, unmitigated critical CVEs.
- In the case of patch releases, if a critical CVE exists, it may be unmitigated if the only resolution involves upgrading a component to a new minor version.
- Critical CVEs in patch releases will be listed in the release notes' "Known Issues" section.
- We prioritize resolving these issues in the next minor release to maintain our commitment to security.
- In the event that we discover a critical CVE for a release that is already Generally Available (GA), and if the mitigation requires a patch release, we commit to releasing a patch with the CVE mitigation within a maximum of 60 days from the date of discovery. This timeframe ensures that our customers receive prompt attention and protection against critical vulnerabilities, even for GA releases.

#### 15.1.1.3 Known Issues:

- If we become aware of a Critical CVE in any of our shipped applications, even if it is not (yet) mitigated, we will promptly add it to the D2iQ Security Updates page.
- For each CVE, we will provide mitigation steps that customers should take to protect their systems.
- We will also outline our plan for resolution in the "Mitigation" column of the report.

## 16 Version Support Policy

D2iQ® supports N-2 of the latest MAJOR.MINOR version of DKP. For example, if the current GA version of DKP® is 2.6.0, then D2iQ supports all patch versions of DKP 2.5.0, and 2.4.0

When the next version releases, support continues for 2.6.x, and 2.5.x Support for DKP version 2.4.x expires. You should upgrade DKP with every new release to stay up-to-date with the latest features and bug fixes.

### 16.1 Supported Kubernetes Versions

Each DKP release supports a range of Kubernetes versions. Details for supported Kubernetes versions on DKP can be found in the [Release Notes \(see page 1813\)](#) on the Supported Kubernetes Version as well as the Supported operating systems. Details for supported operating systems on DKP can be found in [Supported Operating Systems \(see page 69\)](#).

### 16.2 Features in Patches

Occasionally, to make new features available at a faster rate, D2iQ releases features as part of a patch release. If the Release Notes indicate a feature you need and do not yet have, consider upgrading to the latest version to take full advantage of new features and functions.

### 16.3 Experimental Status

“Experimental” means software, features, functionality, sample configurations, or other speculative content that is still under exploration, development, or testing by D2iQ. Experimental components carry no guarantee of eventual release as GA and therefore must not be used in Production Environments. Experimental components qualify for limited, Severity 4 support only and may be discontinued at any time, with or without notice.

Since Experimental components are not intended for Production Environment use, D2iQ cannot assume any responsibility for errors occurring during their use in Production. We can offer only these services in a commercially-reasonable manner, based on the availability of relevant subject matter experts (SMEs):

- General operational guidance for the Experimental component.
- Identifying and diagnosing of errors in configuration or implementation, if possible.
- Advice on preventing and recovering from failures and troubleshooting, as available.

Support for Experimental components is provided on a Standard level, Severity 4 basis only.

This software is provided “as is” and without any express or implied warranties including, without limitation, the implied warranties of merchantability and fitness for a particular purpose.



## 16.4 Beta Feature

“Beta” feature means functionality that it is released to D2iQ’s customer base for testing and feedback before its official release(GA). Beta features are not considered fully complete or stable, and may contain known issues, gaps or limitations. The purpose of releasing a feature in Beta is to gather user feedback, identify bugs, gather insights on usability, and make necessary improvements before the feature is released General Availability (GA) as a stable and **production-ready** version. Beta features will receive a best effort support from D2iQ’s support & engineering teams with no SLA on response time. While it is our full intention to give a fantastic support experience, the early nature of the software may lead to longer support and fix times.



A Beta feature is:

- Not supported for use in Production Environments.
- Intended to eventually graduate to GA, though details may change.
- May not be supported for Upgrade to production GA.

## 16.5 Technical Preview

We provide Technical Preview, or Tech Preview, features to showcase capabilities that might be added to future versions of the product. As they are not yet production-ready, the support terms are the same as defined for experimental features. Technical Preview features are not guaranteed to move forward, so could be removed from future versions of the product.

## 16.6 End-of-Life (EoL)

In D2iQ, End-of-Life refers to versions of our software that are no longer fully supported. In relation to third party applications, EoL refers to applications that are not compatible with the latest version of Kubernetes that is supported in a release.

## 16.7 Support Definitions

### 16.7.1 Secondary Support

The following section describes D2iQ’s support for secondary applications, such as platform applications. All platform applications that D2iQ ships with DKP products are covered under secondary support:

| Type               | Scope Example                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | Support Offered                             |
|--------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------|
| Configuration      | <ul style="list-style-type: none"> <li>• Guidance for base technology and DKP interoperability configuration questions and troubleshooting for different components of the DKP platform.</li> <li>• No support for base technology’s configuration that is unrelated to its integration with DKP.</li> <li>• No support for performance issues with the base technology that is unrelated to its integration with DKP.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                              | Supported with severity 3 & 4 support terms |
| Failure Assistance | <ul style="list-style-type: none"> <li>• Assistance with installation, and upgrade failures of the service as captured in the supported DKP product upgrade pathway.</li> <li>• Assistance with service failures due to platform issues. For example: DKP Enterprise or DKP Essential</li> <li>• Support is limited to troubleshooting for root cause up to DKP product limit. Root causes that are identified to be beyond this limit will need to be pursued by the company who creates the platform application base technology. Note, if the RCA for the failure is due to a non-standard configuration or non-DKP use of the platform application, we will be unable to provide assistance beyond basic identification.</li> <li>• No assistance for base technology’s failures that is unrelated to its integration with DKP.</li> </ul> | Supported with all severities               |

| Type                 | Scope Example                                                                                                                                                                                                                                                                                                                                                 | Support Offered                                                                    |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------|
| Bug Fixes            | <ul style="list-style-type: none"> <li>• Bug fixes for service integration with DKP.</li> <li>• Upstream bug fixes to identified issues in the base technology of the platform application on a best effort basis.</li> <li>• No guarantees that our changes to upstream will be accepted.</li> <li>• No commitment to maintaining forks upstream.</li> </ul> | RCA supported with all severities, Fix supported with severity 3 & 4 support terms |
| Documentation errors | <ul style="list-style-type: none"> <li>• Documentation fixes for life cycle management of services and integration with DKP.</li> <li>• Upstream documentation fixes to reported and identified issues in base technology of the platform application via a best effort basis.</li> <li>• No guarantees that our changes will be accepted.</li> </ul>         | Supported with severity 4 support terms                                            |

## 16.8 Standard Level & Severity Definitions

To view our severity level and support terms refer to [D2iQ Support and Maintenance Terms](#)<sup>1777</sup>.

## 16.9 Operating System (OS) Version Support Policy

This page contains details about D2iQ's support policy in relation to different operating systems.

### 16.9.1 Red Hat Enterprise Linux (RHEL) Version Support Policy

As part of our commitment to delivering the best possible experience to our customers, D2iQ aligns its support for RHEL versions based on their [Extended Update Support](#).<sup>1778</sup> We prioritize supporting RHEL minor

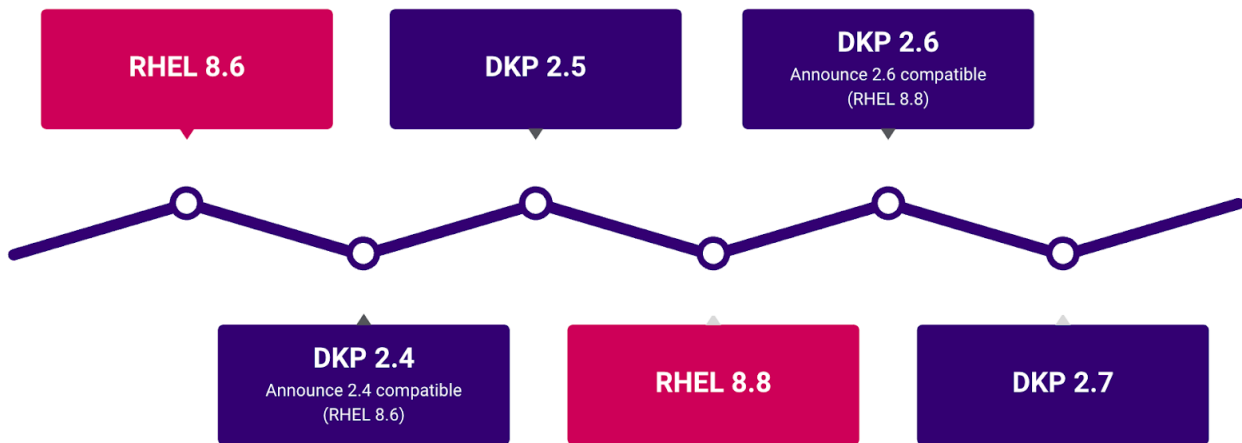
<sup>1777</sup> <https://d2iq.com/legal/support-terms>

<sup>1778</sup> <https://access.redhat.com/articles/rhel-eus>

versions with Extended Update Support, ensuring customers can benefit from long-lasting stability and security.

For each DKP version, we will ensure compatibility and support with the latest released RHEL 8 minor version that has Extended Update Support at the start of the DKP version's development phase. This approach allows us to maintain a seamless integration with the most secure and well-supported RHEL releases, enabling our customers to confidently operate their DKP environments with optimal reliability and peace of mind.

### 16.9.1.1 Example:



## 17 Legal Notices

D2iQ® and its licensors are the owners of all right, title, and interest in and to D2iQ software products, the documentation, all updates, upgrades, and derivative works thereto, and all intellectual property rights therein. D2iQ software products may additionally include third-party open source software.

See the [D2iQ Legal page](#)<sup>1779</sup> for additional details.

---

<sup>1779</sup> <https://d2iq.com/legal/3rd>