

Introducing the D2iQ Kubernetes Platform

D2iQ Kubernetes Platform (DKP)

Exported on 09/21/2023

Table of Contents

1	Quick Start Guides	47
1.1	AWS Quick Start.....	47
1.1.1	Prerequisites	47
1.1.2	Configure AWS Prerequisites.....	47
1.1.3	Create a New AWS Kubernetes Cluster.....	48
1.1.4	Explore the New Kubernetes Cluster	50
1.1.5	Kommander Deployment.....	51
1.1.6	Explore the New Kubernetes Cluster	51
1.1.7	Log in to the UI through Kommander	52
1.1.8	Delete the Kubernetes Cluster and Cleanup your Environment.....	52
1.2	Azure Quick Start	52
1.2.1	Prerequisites	52
1.2.2	Configure Azure Prerequisites	53
1.2.3	Create a New Azure Kubernetes Cluster	54
1.2.4	Explore the New Kubernetes Cluster	56
1.2.5	Kommander Deployment.....	57
1.2.6	Log in to the UI through Kommander	57
1.2.7	Delete the Kubernetes Cluster and Cleanup your Environment.....	57
1.3	GCP Quick Start.....	58
1.3.1	Prerequisites	58
1.3.2	GCP Prerequisites.....	58
1.3.3	Create a New GCP Kubernetes Cluster	59
1.3.4	Explore the New Kubernetes Cluster	61
1.3.5	Install and Log in to the UI.....	64
1.3.6	Delete the Kubernetes Cluster and Clean up Your Environment.....	64
1.4	On Premises Quick Start	65
1.5	vSphere Quick Start	65

1.5.1	DKP Prerequisites	65
1.5.2	vSphere Prerequisites.....	66
1.5.2.1	Create directory for KIB and DKP CLI	67
1.5.2.2	Get the needed D2iQ Software.....	67
1.5.2.3	Create a folder and resource pool in vCenter for DKP cluster	68
1.5.2.4	Build template using KIB	68
1.5.2.5	Adjust the Packer file for your vSphere cluster.....	68
1.5.2.6	Create overrides for docker credentials	69
1.5.2.7	Build VM template using KIB.....	69
1.5.3	Create DKP Cluster on vSphere.....	70
1.5.3.1	Export your vSphere Environment Variables	70
1.5.3.2	Build the Bootstrap Cluster	70
1.5.3.3	Create the DKP cluster deployment YAML.....	70
1.5.4	Create a New vSphere Kubernetes cluster.....	71
1.5.4.1	Pivot the Cluster Controllers and Create CAPI Controllers on Cluster	71
1.5.4.2	After created, move the configuration to the new cluster.....	71
1.5.5	Kommander Deployment.....	72
1.5.6	Explore the New Kubernetes Cluster	72
1.5.7	Log in to the UI through Kommander	73
1.5.8	Delete the Kubernetes Cluster and Cleanup your Environment	73
1.6	DKP Enterprise Quick Start.....	73
1.6.1	Prerequisites	73
1.6.2	Create a new Kubernetes cluster	74
1.6.2.1	Pivot the Cluster Controllers and Create CAPI Controllers on Cluster with this command:	74
1.6.2.2	Once created, move the configuration to the new cluster using the command below:	74
1.6.3	AKS Quick Start.....	74
1.6.3.1	DKP Prerequisites	75
1.6.3.2	AKS Prerequisites	75

1.6.3.3	Name Your Cluster	77
1.6.3.4	Create a New AKS Kubernetes Cluster	78
1.6.3.5	Explore the New Kubernetes Cluster	79
1.6.3.6	Install and Log in to the UI	82
1.6.3.7	Delete the Kubernetes Cluster and Cleanup your Environment	83
1.6.4	EKS Quick Start	83
1.6.4.1	DKP Prerequisites	83
1.6.4.2	AWS Prerequisites	84
1.6.4.3	Configure EKS Prerequisites	84
1.6.4.4	Name Your Cluster	85
1.6.4.5	Create a New EKS Kubernetes Cluster	85
1.6.4.6	Explore the New Kubernetes Cluster	87
1.6.4.7	Install Kommander and Log in to the UI	89
1.6.4.8	Delete the Kubernetes Cluster and Cleanup Your Environment	89
2	Architecture	90
2.1	Learn the key concepts and architectural components of a DKP cluster	90
2.2	Components for the Kubernetes Control Plane	90
2.3	Related Information	91
2.4	DKP Ports	91
2.4.1	Understand the Configured Ports for DKP Deployment	91
2.4.1.1	Control-plane nodes	91
2.4.1.2	Worker nodes	92
2.5	Supported Infrastructure Operating Systems	93
2.5.1	Amazon Web Services (AWS) Amazon Web Services (AWS)	94
2.5.2	Azure	95
2.5.3	GCP	96
2.5.4	Pre-Provisioned/On Premises	96
2.5.5	vSphere	98
2.5.6	EKS	98

2.5.7	AKS	99
3	Download DKP	100
3.1	Download from the Support Website	100
3.2	Download from the AWS Marketplace	100
4	Licenses.....	101
4.1	Purchase a License.....	104
4.2	DKP Enterprise	104
4.2.1	Compatible Infrastructure	105
4.2.2	Platform Applications.....	105
4.2.3	Catalog Applications	105
4.2.4	Cluster Manager.....	105
4.2.5	Built-in GitOps.....	106
4.2.6	DKP Enterprise Multi-cluster UI.....	106
4.3	DKP Essential	107
4.3.1	Compatible Infrastructure	107
4.3.2	Platform Applications.....	107
4.3.3	Cluster Manager.....	107
4.3.4	Built-in GitOps.....	108
4.3.5	DKP Essential Single Cluster UI	108
4.4	Add a DKP license.....	109
4.4.1	Prerequisites	109
4.4.2	Obtain a License Token or AWS License ARN	109
4.4.3	Enter License Information	109
4.4.4	Enter a DKP License via kubectl.....	110
4.5	Remove a DKP license.....	110
4.5.1	Remove a License.....	110
4.5.2	Manually Remove a License using kubectl	111
5	Get Started with DKP	113
5.1	DKP Concepts and Terms	113

5.1.1	Next Topic:	114
5.2	CAPI Concepts and Terms	114
5.2.1	Next Topic:	115
5.3	Resource Requirements	116
5.3.1	Infrastructure Provider Specific	116
5.3.2	General Resource Requirements	116
5.3.3	Control Plane Nodes	116
5.3.4	Worker Nodes.....	116
5.3.5	More Requirement Information for the Kommander Component:	117
5.3.5.1	Next Topic:	117
5.3.6	Management Cluster Application Requirements.....	117
5.3.6.1	Management cluster application minimum resources and storage requirements	117
5.3.7	Workspace Platform Application Defaults and Resource Requirements	118
5.4	Install Overview	121
5.4.1	Prerequisites	121
5.4.2	Prepare Your Environment for Install:.....	122
5.4.3	Next Step	123
5.5	Deploy a Cluster with Konvoy.....	123
5.5.1	Deploy a Cluster with Konvoy.....	124
5.5.2	Infrastructure-specific Flags	124
5.5.3	Next Step	125
5.6	Install Kommander.....	125
5.6.1	Prerequisites	125
5.6.2	Configure a Default StorageClass	125
5.6.3	Install Kommander.....	126
5.6.4	Verify Installation	126
5.6.5	Next Step	127
6	Advanced Konvoy Configuration	128

6.1	Universal Configurations for all Infrastructure Providers.....	128
6.1.1	Use HTTP Proxy	128
6.1.2	Bootstrap Proxy Settings.....	129
6.1.3	Create CAPI Components with HTTP Proxy	129
6.1.4	Configure and Create a Cluster with HTTP Proxy.....	130
6.1.4.1	Example of how to configure the control plane and worker nodes to use HTTP proxy:.....	131
6.1.4.2	Example of create cluster using the configured HTTP Proxy above:.....	132
6.1.5	HTTP Proxy for the Kommander Component of DKP	132
6.1.6	Load Balancer	133
6.1.6.1	External Load Balancer	133
6.1.7	Select your Connection Mechanism	133
6.1.8	Additional Configurations.....	133
6.2	Local Registry Tools for Air-gapped Environments	134
6.2.1	Air-Gapped Registry Prerequisites	134
6.2.1.1	JFrog Artifactory	134
6.2.1.2	Nexus Registry	134
6.2.1.3	Harbor Registry	134
6.2.1.4	Bastion Host.....	135
6.3	AKS Infrastructure	135
6.3.1	AKS Prerequisites	135
6.3.2	Create a New AKS Cluster	136
6.3.2.1	DKP to create a new AKS cluster	136
6.3.2.2	Create a New AKS Kubernetes Cluster.....	136
6.3.2.3	Inspecting or Editing the Cluster Objects.....	137
6.3.2.4	Known Limitations	139
6.3.3	Explore New AKS Cluster	139
6.3.3.1	Learn to interact with your AKS Kubernetes cluster	139
6.3.3.2	Explore the New AKS Cluster	139

6.3.4	Delete AKS Cluster	143
6.3.4.1	Delete the AKS cluster and clean up your environment	143
6.3.4.2	Delete the Workload Cluster.....	143
6.3.4.3	Next Step:	144
6.3.4.4	Known Limitations	144
6.4	AWS Infrastructure	144
6.4.1	Configuration Types	144
6.4.2	AWS Diagrams	144
6.4.3	AWS Pricing Considerations	146
6.4.4	AWS Service Limits.....	146
6.4.5	AWS Konvoy Image Builder	146
6.4.5.1	1. Create a Custom AMI	147
6.4.5.2	Prerequisites	147
6.4.5.3	Extract AMI Bundle	147
6.4.5.4	Build the Image	148
6.4.5.5	Launch a DKP Cluster with a Custom AMI	148
6.4.5.6	Launch a DKP Cluster with Custom AMI Lookup.....	149
6.4.5.7	Using Custom Source AMIs	149
6.4.5.8	Related Information	150
6.4.5.9	2. Air-gapped AMI	151
6.4.6	Minimal Permissions and Role to Create Clusters	151
6.4.6.1	Prerequisites	151
6.4.6.2	Minimal Permissions	152
6.4.6.3	Create Resources in Cloudformation Stack.....	152
6.4.6.4	Leverage the Role	155
6.4.6.5	Use EC2 Instance Profiles	156
6.4.6.6	Use Access Keys.....	157
6.4.7	Cluster IAM Policies and Roles	157
6.4.7.1	Prerequisites	157

6.4.7.2	Next Step:	157
6.4.7.3	IAM Artifacts	158
6.4.8	Multiple AWS Accounts	164
6.4.8.1	Objective	164
6.4.8.2	Assumptions	165
6.4.8.3	Glossary	165
6.4.8.4	Prerequisites	165
6.4.8.5	Deploy DKP on AWS.....	165
6.4.8.6	Next Steps:	167
6.4.9	Advanced AWS Install.....	167
6.4.9.1	AWS Install Prerequisites	168
6.4.9.2	Custom AMI	169
6.4.9.3	Bootstrap AWS.....	173
6.4.9.4	Create a New AWS Cluster	175
6.4.9.5	Explore New AWS Cluster	186
6.4.9.6	Make the New AWS Cluster Self-Managed	189
6.4.9.7	AWS Certificate Renewal.....	192
6.4.9.8	Configure Infrastructure in UI.....	194
6.4.9.9	Delete an AWS Cluster	195
6.4.9.10	Replace an AWS Node	199
6.4.10	Install AWS Air-Gapped	202
6.4.10.1	Create a Kubernetes cluster in a private subnet with no access to the Internet (air-gapped)	202
6.4.10.2	AWS Air-gapped Prerequisites	203
6.4.10.3	Air-gapped AMI	203
6.4.10.4	AWS Air-gapped Seed Docker Registry	204
6.4.10.5	AWS Air-gapped Bootstrap.....	205
6.4.10.6	Create a New AWS Air-gapped Cluster.....	206
6.4.10.7	Explore New AWS Air-gapped Cluster	209

6.4.10.8	Make the Air-gapped AWS Cluster Self-Managed	210
6.4.10.9	Delete AWS Air-gapped Cluster.....	213
6.4.11	GPUs in an AWS environment	214
6.4.11.1	Understanding GPUs.....	214
6.4.11.2	Install GPU Support for Supported Distributions on AWS.....	215
6.4.11.3	Configure Konvoy Automatic GPU Node Labels.....	216
6.4.11.4	Update Nvidia GPU Clusters.....	216
6.4.12	Manage AWS Node Pools	217
6.4.12.1	Create AWS Node Pools	218
6.4.12.2	List AWS Node Pools.....	219
6.4.12.3	Scale AWS Node Pools.....	220
6.4.12.4	Delete AWS Node Pools	222
6.4.12.5	AWS Cluster Autoscaler	223
6.4.13	Creating DKP Clusters on AWS.....	225
6.4.13.1	Before you begin	225
6.4.13.2	Simplified Cluster Creation on AWS	225
6.5	Azure Infrastructure.....	226
6.5.1	Azure Prerequisites.....	226
6.5.1.1	DKP Prerequisites	226
6.5.1.2	Azure Prerequisites.....	228
6.5.2	Azure using Konvoy Image Builder	229
6.5.2.1	Prerequisites	229
6.5.2.2	Extract the Bundle.....	230
6.5.2.3	Configure Azure Prerequisites	230
6.5.2.4	Build the Image	231
6.5.2.5	Image Gallery	232
6.5.3	Azure Bootstrap	232
6.5.3.1	Prepare to deploy Kubernetes clusters	232
6.5.3.2	Prerequisites	233

6.5.3.3	Bootstrap Cluster Lifecycle Services.....	233
6.5.3.4	(Optional) Create identity secret for Azure.....	234
6.5.4	Create a New Azure Cluster	235
6.5.4.1	Prerequisites	235
6.5.4.2	Name your cluster.....	235
6.5.4.3	Tips and Tricks.....	235
6.5.4.4	Create a new Azure Kubernetes cluster	236
6.5.4.5	Known Limitations	242
6.5.5	Explore new Azure Cluster	242
6.5.5.1	Learn to interact with your Kubernetes cluster	242
6.5.5.2	Explore the new Kubernetes cluster	242
6.5.6	Azure Make new Cluster Self-Managed	246
6.5.6.1	Make the new Kubernetes cluster manage itself	246
6.5.6.2	Known Limitations	248
6.5.7	Azure Certificate Renewal	248
6.5.7.1	Configure Automated Renewal for Managed Kubernetes PKI Certificates.....	248
6.5.7.2	Requirements	249
6.5.7.3	Prerequisites	249
6.5.8	Azure Replace a Node	251
6.5.8.1	Replace a worker node	251
6.5.8.2	Prerequisites	251
6.5.8.3	Replace a worker node	251
6.5.9	Azure Delete Cluster	254
6.5.9.1	Delete the Kubernetes cluster and clean up your environment	254
6.5.9.2	Prepare to Delete a Workload Cluster	254
6.5.9.3	Delete the Workload Cluster.....	256
6.5.9.4	Delete the Bootstrap Cluster	257
6.5.9.5	Next Step:	257

6.5.9.6	Known Limitations	257
6.5.10	Create a new Azure Cluster via UI.....	258
6.5.10.1	Prerequisites	258
6.5.10.2	Provision an Azure Cluster	259
6.6	EKS Infrastructure.....	259
6.6.1	EKS Introduction	260
6.6.2	EKS Prerequisites.....	261
6.6.2.1	DKP Prerequisites	261
6.6.2.2	AWS prerequisites	261
6.6.2.3	Access Cluster	265
6.6.3	EKS Cluster IAM Policies and Roles	265
6.6.3.1	Prerequisites:	265
6.6.3.2	EKS IAM Artifacts.....	265
6.6.4	Create an EKS Cluster from the CLI.....	270
6.6.4.1	Create a New EKS Kubernetes Cluster	270
6.6.4.2	Known Limitations	276
6.6.5	Grant Cluster Access	277
6.6.5.1	How to Grant Cluster Access	277
6.6.6	Explore EKS Cluster	278
6.6.6.1	Explore the new Kubernetes cluster	278
6.6.7	Manage EKS Nodepools.....	280
6.6.7.1	Create a node pool.....	281
6.6.7.2	Scaling Up Node Pools	282
6.6.7.3	Delete EKS Node Pools.....	282
6.6.8	Delete EKS Cluster from CLI.....	283
6.6.8.1	Delete the EKS cluster	283
6.6.8.2	Next Step:	284
6.6.8.3	Known Limitations	284
6.6.9	Create an EKS Cluster from the UI.....	284

6.6.9.1	Create an AWS Infrastructure Provider	284
6.6.9.2	Provision an EKS Cluster	285
6.6.9.3	Access EKS Cluster	285
6.6.9.4	IAM User and Role Access for EKS Clusters.....	285
6.6.10	Delete EKS Cluster from UI.....	286
6.7	Pre-provisioned Infrastructure	288
6.7.1	Create a Kubernetes cluster on pre-provisioned infrastructure.....	288
6.7.2	Pre-provisioned Prerequisites	289
6.7.2.1	Fulfill the prerequisites for using a pre-provisioned infrastructure	289
6.7.2.2	Control plane machines.....	289
6.7.2.3	Worker machines	290
6.7.3	Pre-provisioned Prerequisites Air-gapped.....	291
6.7.3.1	Fulfill the prerequisites for using a pre-provisioned infrastructure when Air-Gapped.....	291
6.7.3.2	Air-Gapped Registry Prerequisites	291
6.7.3.3	Load the bootstrap image	291
6.7.3.4	Copy air-gapped artifacts onto cluster hosts.....	292
6.7.4	Pre-provisioned Bootstrap.....	296
6.7.4.1	Bootstrap a Kind Cluster and CAPI controllers	296
6.7.5	Pre-provisioned Create Necessary Secrets and Overrides	296
6.7.5.1	Create necessary secrets and overrides for pre-provisioned clusters ...	296
6.7.5.2	Name your cluster.....	296
6.7.5.3	Create a unique cluster name	297
6.7.5.4	Create a secret.....	297
6.7.5.5	Create overrides.....	297
6.7.6	Pre-provisioned Define Infrastructure	298
6.7.6.1	Define your infrastructure.....	299
6.7.7	Pre-provisioned Define Control Plane Endpoint.....	300
6.7.7.1	Define the Control Plane Endpoint for your cluster	300

6.7.7.2	External load balancer	301
6.7.7.3	Built-in virtual IP	301
6.7.7.4	Single-Node control plane	301
6.7.7.5	Known limitations	302
6.7.8	Pre-provisioned Create new Cluster	302
6.7.8.1	Create a Kubernetes cluster using the infrastructure definition	302
6.7.8.2	Audit logs.....	304
6.7.8.3	Modify the Calico installation.....	305
6.7.8.4	Use the built-in Virtual IP	308
6.7.8.5	Provision on the Flatcar Linux OS.....	309
6.7.8.6	Use an HTTP proxy	309
6.7.8.7	Use alternate pod or service subnets	311
6.7.9	Pre-provisioned Make Cluster Self-managed.....	313
6.7.9.1	Make the New Kubernetes Cluster Manage Itself	313
6.7.10	Pre-provisioned Configure MetalLB.....	316
6.7.10.1	Create a MetalLB configmap for your pre-provisioned infrastructure....	316
6.7.10.2	Layer 2 configuration.....	316
6.7.10.3	BGP configuration.....	317
6.7.11	Pre-provisioned Create and Delete Node Pools.....	317
6.7.11.1	Create a Pre-provisioned node pool	318
6.7.11.2	Delete a node pool	319
6.7.12	Pre-provisioned Add Nodes to Existing Node Pools	319
6.7.12.1	Prerequisites	319
6.7.12.2	Scale up a Cluster Node	319
6.7.12.3	Scale Down a Cluster Node.....	321
6.7.13	GPU Nodepools in a Pre-provisioned Environment	321
6.7.14	Pre-provisioned Delete Cluster.....	323
6.7.14.1	Prepare to Delete the Pre-provisioned Cluster	323
6.7.14.2	Delete the Workload Cluster.....	325

6.8	vSphere Infrastructure	326
6.8.1	Creating DKP clusters in a VMware vSphere environment	326
6.8.2	vSphere Prerequisites.....	327
6.8.2.1	Prepare your environment to run DKP with VMware vSphere.....	327
6.8.2.2	DKP Prerequisites	328
6.8.2.3	VMware vSphere Prerequisites.....	328
6.8.2.4	Next Steps:	329
6.8.2.5	Minimum User Permissions.....	329
6.8.2.6	vSphere Storage Options.....	332
6.8.3	Create a Base OS image in vSphere	332
6.8.3.1	Create the Base OS Image	332
6.8.4	Create a VM Template.....	333
6.8.4.1	Prerequisites	334
6.8.4.2	Create a vSphere Template for Your Cluster from a Base OS Image.....	334
6.8.5	vSphere Bootstrap	335
6.8.5.1	Prepare to deploy Kubernetes clusters	335
6.8.5.2	Prerequisites	336
6.8.5.3	Bootstrap cluster lifecycle services.....	336
6.8.6	Create new vSphere Cluster	338
6.8.6.1	Prerequisites	338
6.8.6.2	Name your cluster.....	338
6.8.6.3	Create a New vSphere Kubernetes Cluster	338
6.8.6.4	Known Limitations	344
6.8.7	Explore a vSphere Cluster	344
6.8.7.1	Get the kubeconfig file for the new Kubernetes cluster	345
6.8.7.2	Create a StorageClass with a vSphere datastore	345
6.8.7.3	Explore nodes and pods in the new cluster	346
6.8.8	Make vSphere Cluster Self-Managed	349
6.8.8.1	Make the new Kubernetes cluster manage itself	349

6.8.8.2	Known limitations	351
6.8.9	Configure MetalLB for a vSphere infrastructure	352
6.8.9.1	Create a MetalLB configmap for your vSphere infrastructure	352
6.8.9.2	Layer 2 configuration	352
6.8.9.3	BGP configuration	353
6.8.10	Install vSphere Air-Gapped	354
6.8.10.1	Create a Kubernetes vSphere cluster in a private network with no access to the Internet (air-gapped)	354
6.8.10.2	Create a Base Air-gapped OS VM Image	354
6.8.10.3	Create a CAPI VM Template	356
6.8.10.4	vSphere Air-gapped Seed Docker Registry	356
6.8.10.5	vSphere Air-gapped Bootstrap	357
6.8.10.6	Create a new Air-gapped vSphere Cluster	358
6.8.10.7	Explore vSphere Air-gapped Cluster	361
6.8.11	vSphere Certificate Renewal	365
6.8.11.1	Configure Automated Renewal for Managed Kubernetes PKI Certificates	365
6.8.11.2	Certificate Renewal	365
6.8.11.3	Prerequisites	366
6.8.12	Delete vSphere Cluster	368
6.8.12.1	Prepare to delete a self-managed workload cluster	368
6.8.12.2	Delete the Workload Cluster	370
6.8.12.3	Delete the bootstrap cluster	371
6.8.12.4	Next Step:	372
6.8.12.5	Known limitations	372
6.8.13	Manage vSphere Node Pools	372
6.8.13.1	Create vSphere Node Pools	372
6.8.13.2	List vSphere Node Pools	374
6.8.13.3	Scale vSphere Node Pools	374
6.8.13.4	Delete vSphere Node Pools	377

6.8.13.5	vSphere Cluster Autoscaler	377
6.8.14	Replace a vSphere Node	380
6.8.14.1	Prerequisites	380
6.8.14.2	Replace a worker node	380
6.9	GCP Infrastructure	382
6.9.1	GCP Prerequisites	382
6.9.1.1	Prerequisites	382
6.9.1.2	Control plane nodes.....	383
6.9.1.3	Worker nodes	383
6.9.1.4	GCP Prerequisite Roles	383
6.9.2	GCP Konvoy Image Builder	385
6.9.2.1	Prerequisites	385
6.9.2.2	GCP Prerequisites	386
6.9.2.3	Create a Network (optional)	387
6.9.2.4	Build the GCP image	387
6.9.3	Bootstrap GCP	388
6.9.3.1	Prerequisites	388
6.9.3.2	Bootstrap Cluster Lifecycle Services.....	389
6.9.4	Create a New GCP Cluster.....	390
6.9.4.1	Prerequisites	390
6.9.4.2	Name your cluster.....	390
6.9.4.3	Create a new GCP cluster.....	391
6.9.5	Explore the GCP Cluster	393
6.9.5.1	Explore the new Kubernetes cluster	393
6.9.6	Make the New GCP Cluster Self-Managed.....	397
6.9.6.1	Make the new Kubernetes cluster manage itself	397
6.9.6.2	Known Limitations	399
6.9.7	Manage GCP Node Pools	399
6.9.7.1	Create GCP Node Pools.....	399

6.9.7.2	List GCP Node Pools	401
6.9.7.3	Scale GCP Node Pools	401
6.9.7.4	Delete GCP Node Pools.....	403
6.9.7.5	GCP Cluster Autoscaler.....	404
6.9.8	Delete a GCP Cluster	406
6.9.8.1	Prepare to Delete a Workload Cluster	406
6.9.8.2	Delete the Workload Cluster.....	408
6.9.8.3	Delete the Bootstrap Cluster	409
6.9.8.4	Next Step:	409
6.9.8.5	Known Limitations	409
6.10	Verify DKP installation	409
6.10.1	Check the cluster infrastructure and nodes.....	410
6.10.2	Verify all pods are running.....	411
6.10.3	Troubleshooting.....	411
6.11	Konvoy Image Builder.....	411
6.11.1	Prerequisites	412
6.11.2	Compatible Versions	412
6.11.3	KIB with AWS.....	414
6.11.3.1	Create a Custom AMI	414
6.11.3.2	AWS Air-gapped AMI	418
6.11.3.3	Add Custom Tags to Image	419
6.11.4	KIB for Azure	420
6.11.4.1	Learn how to build a custom Azure Image for use with DKP.....	420
6.11.4.2	Prerequisites	420
6.11.4.3	Extract the Bundle.....	421
6.11.4.4	Configure Azure Prerequisites	421
6.11.4.5	Build the Image	423
6.11.4.6	Image Gallery	423
6.11.5	KIB with GCP	424

6.11.5.1	Prerequisites	424
6.11.5.2	GCP Prerequisites.....	425
6.11.5.3	Create a Network (optional).....	426
6.11.5.4	Build the GCP image.....	426
6.11.6	KIB for GPU.....	427
6.11.6.1	Verification	429
6.11.7	KIB with vSphere	430
6.11.7.1	Create a Base OS Image.....	430
6.11.7.2	Create a vSphere Virtual Machine Template.....	431
6.11.8	Konvoy Image Builder CLI	433
6.11.8.1	Other resources:.....	433
6.11.8.2	konvoy-image build.....	434
6.11.8.3	konvoy-image completion	439
6.11.8.4	konvoy-image generate-docs	443
6.11.8.5	konvoy-image generate	444
6.11.8.6	konvoy-image provision.....	448
6.11.8.7	konvoy-image upload.....	448
6.11.8.8	konvoy-image validate.....	450
6.11.8.9	konvoy-image version.....	450
6.11.9	Use Override files with Konvoy Image Builder	451
6.11.9.1	Learn how to use override files with Konvoy Image builder	451
6.11.9.2	Default Override Files	451
6.11.9.3	Custom Override Files	454
6.12	GPU for Konvoy.....	460
6.13	FIPS 140-2 Compliance	460
6.13.1	FIPS Support in DKP	460
6.13.2	Infrastructure Requirements for FIPS-140-2 Mode	461
6.13.2.1	Supported Operating Systems	461
6.13.3	Deploying a Cluster in FIPS mode.....	461

6.13.3.1	Supported FIPS Builds	461
6.13.4	Create FIPS 140 Images	462
6.13.4.1	Use Konvoy Image Builder to create images with FIPS-compliant binaries	462
6.13.4.2	Create FIPS-140 images	462
6.13.5	Validate FIPS in Cluster	462
6.13.5.1	Run FIPS validation	463
6.13.5.2	Download Signature Files	464
6.13.6	FIPS 140 Mode Performance Impact	465
6.13.6.1	Understand the performance impact from operating your cluster in FIPS 140 mode	465
6.14	Delete a DKP Cluster with One Command	465
6.15	Configure the Control Plane	466
6.15.1	Prerequisites	466
6.15.1.1	Modifying Audit Logs	466
6.15.1.2	Viewing the Audit Logs	472
6.16	Update Cluster Nodepools	473
6.16.1	Prerequisites:	473
6.16.2	Steps:	474
7	Advanced Kommander Configuration	475
7.1	Kommander Install Configuration	475
7.1.1	Initialize a Configuration File	475
7.1.2	Configure Applications	475
7.1.2.1	Inline configuration (using values)	476
7.1.2.2	Reference another YAML file (using valuesFrom)	476
7.1.3	Minimal Kommander Installation	477
7.1.4	Install with Configuration File	477
7.1.5	Verify Installation	477
7.1.6	Custom Domains and Certificates	478
7.1.6.1	Why to set up a Custom Domain or Certificate?	478

7.1.6.2	Certificate Authority (CA) Specifics.....	479
7.1.6.3	Configure your Custom Domain and Certificate.....	480
7.1.6.4	Verification and Troubleshooting for Custom Certificates.....	488
7.1.7	DKP Kommander Configuration Reference.....	490
7.1.7.1	Configuration Parameters.....	490
7.1.7.2	AppConfig.....	493
7.1.7.3	IngressCertificate.....	494
7.1.7.4	Airgapped	494
7.1.8	Configure HTTP Proxy.....	494
7.1.8.1	Prerequisites	495
7.1.8.2	Enable Gatekeeper	495
7.1.8.3	Create Gatekeeper ConfigMap in the kommander Namespace.....	496
7.1.8.4	HTTP Proxy Configuration Considerations	497
7.1.8.5	Install Kommander.....	498
7.1.8.6	Configure Workspace or Project.....	498
7.1.8.7	Configure HTTP Proxy in Attached Clusters.....	498
7.1.8.8	Create Gatekeeper ConfigMap in the Workspace Namespace	499
7.1.8.9	Configure Your Applications	500
7.1.8.10	Manually Configure Your Application	500
7.1.9	External Load Balancer.....	501
7.1.9.1	Load Balancing for External Traffic in DKP.....	501
7.1.9.2	Configure Kommander to use an External Load Balancer	501
7.1.9.3	Configure the External Load Balancer to Target the Specified Ports.....	502
7.1.10	Configure an Enterprise catalog	502
7.1.10.1	Configure a Default Enterprise Catalog.....	503
7.1.10.2	Configure an Enterprise Catalog after installing or upgrading DKP	503
7.1.10.3	Labels	503
7.1.10.4	Next Step:	504
7.2	Install Kommander in an Air-gapped Environment.....	504

7.2.1	Prerequisites	504
7.2.2	Load the Docker Images into Your Docker Registry	505
7.2.2.1	Next Step	505
7.2.3	Air-gapped Essential	506
7.2.3.1	Kommander in an Air-gapped Environment	506
7.2.3.2	Kommander in Air-gapped with DKP Insights	507
7.2.4	Air-gapped Enterprise	509
7.2.4.1	Install Air-gapped Kommander with DKP Catalog Applications	509
7.2.4.2	Install Air-gapped Kommander with DKP Insights and DKP Catalog Applications	511
7.3	Install Kommander in a Non-air-gapped Environment.....	513
7.3.1	Prerequisites	513
7.3.2	Configure a Default StorageClass	514
7.3.3	Install Kommander.....	514
7.3.4	Verify Installation	515
7.3.4.1	Next Step	515
7.4	Install Kommander in a Pre-provisioned Environment.....	515
7.4.1	Pre-provisioned Prerequisites for Kommander	516
7.4.1.1	Prerequisites	516
7.4.1.2	Next Step:	516
7.4.2	Prepare the Kommander Installer Configuration File.....	516
7.5	Install Kommander on a Small Environment.....	518
7.5.1	Prerequisites	518
7.5.2	Minimal Kommander installation.....	518
7.5.3	Verify Installation	520
7.6	Verify Kommander Installation	521
7.6.1	Failed HelmReleases	522
7.6.2	Next Step	522
7.7	Log in to the UI with Kommander	522

7.7.1	Next Step	524
7.8	Helm and Chart Bundle CLI Commands.....	524
7.8.1	Kommander Charts Bundle	524
7.8.2	DKP Internal Helm Repository.....	524
8	Day 2 Operations.....	526
8.1	Deploy a Sample Application	526
8.1.1	Learn how to deploy a sample application on a DKP cluster.....	526
8.1.2	Before You Begin	526
8.1.3	Deploy a Sample Application	527
8.1.4	Related Information	528
8.2	Operations.....	529
8.2.1	Access Control.....	529
8.2.1.1	Centrally Manage Access Across Clusters	529
8.2.1.2	Special Limitation for Kommander Roles.....	530
8.2.1.3	Types of Access Control Objects.....	530
8.2.1.4	Related Information	532
8.2.1.5	Granting Access to Kubernetes and Kommander Resources.....	533
8.2.2	Identity Providers	540
8.2.2.1	Grant access to users in your organization.....	540
8.2.2.2	Prerequisites	540
8.2.2.3	Groups	541
8.2.2.4	Authorize a Group in Github	541
8.2.2.5	External LDAP directory.....	543
8.2.3	Infrastructure Providers.....	546
8.2.3.1	View and Modify Infrastructure Providers.....	547
8.2.3.2	AWS	547
8.2.3.3	Delete an infrastructure provider	547
8.2.3.4	Configure an AWS Provider with a User Role.....	547
8.2.3.5	AWS Static Credentials.....	554

8.3	Applications	561
8.3.1	AppDeployment resources	561
8.3.1.1	Customization	562
8.3.1.2	Prerequisites	562
8.3.1.3	Customize Your Application.....	562
8.3.1.4	Print and Review the Current State of an AppDeployment Resource.....	563
8.3.1.5	Deployment Scope	564
8.3.1.6	More Information	564
8.3.2	Manage an Application using the UI.....	564
8.3.2.1	Enterprise - Manage an Application using the UI.....	565
8.3.2.2	Essential - Manage an Application using the UI.....	568
8.3.3	Platform Applications.....	570
8.3.3.1	Related Topics	570
8.3.3.2	Deploy Platform Applications via CLI	570
8.3.3.3	Upgrade Platform Applications.....	572
8.3.3.4	Platform Application Dependencies	575
8.3.3.5	Workspace Platform Application Resource Requirements.....	580
8.4	Workspaces.....	580
8.4.1	Global / Workspace UI	580
8.4.2	Default Workspace.....	581
8.4.3	Create a Workspace.....	581
8.4.4	Add, Edit, and Delete Workspace Annotations and Labels	581
8.4.5	Delete a Workspace	581
8.4.6	Workspace Applications.....	582
8.4.6.1	Cluster-scoped Application Configuration via the UI.....	582
8.4.6.2	Cluster-scoped Configuration for Existing AppDeployments	584
8.4.6.3	Workspace Catalog Applications.....	593
8.4.7	Workspace Role Bindings.....	614
8.4.7.1	Configure Workspace Role Bindings	614

8.5	Projects.....	615
8.5.1	Project Namespace	615
8.5.2	Create a Project	615
8.5.3	Create a Project - UI Method	616
8.5.4	Create a Project - CLI Method	616
8.5.5	Project Applications.....	617
8.5.5.1	Project Platform Applications.....	617
8.5.5.2	Project Catalog Applications.....	623
8.5.5.3	Project AppDeployments.....	641
8.5.6	Project Deployments	641
8.5.6.1	What is GitOps?.....	642
8.5.6.2	Continuous Delivery with GitOps.....	642
8.5.6.3	Continuous Deployment	652
8.5.6.4	Project Deployments Troubleshooting.....	656
8.5.6.5	View Helm Releases	656
8.5.7	Project Role Bindings.....	657
8.5.7.1	Configure Project Role Bindings - UI Method.....	657
8.5.7.2	Configure Project Role Bindings - CLI Method.....	657
8.5.7.3	Configure Project Role Bindings to Bind to WorkspaceRoles - CLI Method.....	657
8.5.7.4	Role Binding with VirtualGroup	660
8.5.8	Project Roles	660
8.5.8.1	Configure Project Role - UI Method	660
8.5.8.2	Configure Project Role - CLI Method	661
8.5.9	Project ConfigMaps	663
8.5.9.1	Configuring Project ConfigMaps - UI Method	663
8.5.9.2	Configuring Project ConfigMaps - CLI Method	664
8.5.10	Project Secrets.....	665
8.5.10.1	Configure Project Secrets - UI Method	665

8.5.10.2	Configure Project Secrets - CLI Method	665
8.5.11	Project Quotas & Limit Ranges	666
8.5.11.1	Creating Project Quotas & Limit Ranges- UI Method.....	666
8.5.11.2	Create Project Quotas & Limit Ranges - CLI Method.....	667
8.5.12	Project Network Policies	669
8.5.12.1	About Network Plugins.....	669
8.5.12.2	About Network Policies.....	669
8.5.12.3	Creating Network Policies.....	669
8.5.12.4	Network Policy Examples.....	671
8.6	Manage Clusters	674
8.6.1	Cluster Types	674
8.6.2	Cluster Statuses.....	674
8.6.3	Cluster Resources.....	676
8.6.4	Platform Application.....	677
8.6.5	Kubernetes Cluster Federation (KubeFed).....	677
8.6.6	Attach an Existing Kubernetes Cluster	678
8.6.6.1	Attach Kubernetes Cluster	678
8.6.6.2	Requirements for Attaching an Existing Cluster	678
8.6.6.3	Create a kubeconfig File for your Cluster	680
8.6.6.4	Generate a kubeconfig File.....	683
8.6.6.5	Attach a Cluster with no Networking Restrictions.....	685
8.6.6.6	Attach a Cluster with Networking Restrictions	686
8.6.6.7	Finish Attaching the Existing Cluster.....	708
8.6.6.8	Attach Amazon EKS Cluster.....	708
8.6.6.9	Attach GKE Cluster	713
8.6.6.10	Manually Attach a CLI-created Cluster	717
8.6.6.11	Access a Managed Cluster	718
8.6.6.12	Finalize Attaching your Cluster From the UI.....	719
8.6.7	Advanced Creation of CLI Clusters.....	719

8.6.7.1	Generate Cluster Objects	719
8.6.7.2	Use the Upload YAML Form	720
8.6.8	Management Cluster	720
8.6.8.1	Editing.....	720
8.6.8.2	Disconnecting	721
8.6.9	Cluster Applications.....	721
8.6.9.1	Custom Cluster Application Dashboard Cards	721
8.6.10	Custom Domains and Certificates Configuration.....	723
8.6.10.1	Why Should you set up a Custom Domain or Certificate?.....	724
8.6.10.2	Configure Custom Domains or Custom Certificates	725
8.6.10.3	Configuration Example with Let's Encrypt.....	726
8.6.10.4	Verify and Troubleshoot Configuration Status	727
8.6.11	Disconnect or Delete Clusters.....	729
8.6.11.1	Disconnect or delete a cluster	729
8.6.11.2	Disconnect vs. Delete	729
8.6.11.3	Troubleshooting.....	730
8.7	Backup and Restore.....	730
8.7.1	Configure Velero with Rook Ceph Storage	731
8.7.1.1	Cloud Provider Storage.....	731
8.7.2	Configure Velero with Azure, or GCP storage	731
8.7.2.1	Configure Velero to use Azure Blob Storage.....	731
8.7.2.2	Configure Velero to use Google Cloud Buckets.....	735
8.7.3	Configure Velero with AWS	740
8.7.3.1	Overview	740
8.7.3.2	Velero with AWS S3 - Prepare your Environment.....	741
8.7.3.3	Velero with AWS S3 - Configure Velero	742
8.7.3.4	Velero with AWS S3 - Establish a Backup Location.....	745
8.7.4	Back up with Velero	746
8.7.4.1	Install the Velero CLI.....	747

8.7.4.2	Regular Backup Operations.....	747
8.7.4.3	Restore a cluster from backup.....	749
8.7.4.4	Backup Service Diagnostics.....	751
8.8	Logging.....	751
8.8.1	Admin-level logs.....	753
8.8.2	Enable Workspace-level Logging.....	753
8.8.2.1	How to enable Workspace-level Logging for use with DKP.....	753
8.8.2.2	Logging Prerequisites.....	754
8.8.2.3	Enable Logging Applications through the UI.....	754
8.8.2.4	Create AppDeployments to Enable Workspace Logging.....	755
8.8.2.5	Override ConfigMap to Restrict Logging to Specific Namespaces.....	756
8.8.2.6	Override ConfigMap to Modify the Storage Retention Period in Workspace Grafana Loki.....	758
8.8.2.7	Verify Cluster Logging Stack Installation.....	760
8.8.2.8	View Cluster Log Data.....	761
8.8.3	Multi-Tenant Logging Overview.....	763
8.8.3.1	Enable Multi-tenant Logging.....	765
8.8.3.2	Create a Project for Logging.....	765
8.8.3.3	Create Project-level Logging AppDeployments.....	766
8.8.3.4	Override ConfigMap to Modify the Storage Retention Period in Project Grafana Loki.....	767
8.8.3.5	Verify Project Logging Stack Installation.....	769
8.8.3.6	View Project Log Data.....	770
8.8.4	Fluent Bit.....	773
8.8.4.1	Audit Log Collection.....	773
8.8.4.2	Related Information.....	773
8.8.4.3	Collecting systemd Logs from a Non-default Path.....	773
8.8.5	Configuring Loki to use AWS S3 Storage in DKP.....	778
8.8.5.1	Configuring Loki.....	778
8.9	Security.....	780

8.9.1	Authentication and authorization architecture	781
8.9.1.1	Details on distributed authentication and authorization between clusters	781
8.9.1.2	Authentication	781
8.9.1.3	Authorization	782
8.9.2	OpenID Connect (OIDC) Introduction.....	782
8.9.2.1	An introduction to OpenID Connect (OIDC) Authentication in Kubernetes	782
8.9.2.2	Identity Provider	782
8.9.2.3	Add Login Connectors	784
8.9.2.4	Change the Access Token Lifetime	785
8.9.2.5	Authentication	785
8.9.3	SAML Connector	786
8.9.3.1	Connect your Kommander cluster to an IdP using SAML.....	786
8.9.3.2	Connect Kommander to an IdP Using SAML	786
8.9.4	Policy Controls	789
8.9.4.1	Workload Policy Controls	789
8.9.4.2	Enforce Policies Using Gatekeeper.....	789
8.9.5	Traefik-Forward-Authentication in DKP (TFA).....	793
8.9.5.1	TFA Overview	793
8.9.5.2	TFA Authentication Workflow	794
8.9.5.3	Default TFA Configuration in DKP.....	794
8.9.5.4	Cluster Storage Option	795
8.10	Networking	796
8.10.1	Configure networking for Konvoy cluster.....	796
8.10.2	Networking Service	796
8.10.2.1	Service Topology	797
8.10.2.2	EndpointSlices	798
8.10.2.3	DNS for Services and Pods	798
8.10.2.4	Ingress and Networking.....	799

8.10.2.5	Network Policies	800
8.10.3	Required Domains.....	802
8.10.3.1	This section describes the required domains for DKP	802
8.10.4	Configure Ingress for load balancing	802
8.10.4.1	Learn how to configure Ingress settings for load balancing (layer-7)....	802
8.10.4.2	Prerequisites	803
8.10.4.3	Expose a pod using an Ingress (L7)	803
8.10.5	Ingress	805
8.10.5.1	Traefik Ingress Controller.....	805
8.10.5.2	Traefik v2.4.....	805
8.10.5.3	Related Information	806
8.10.6	Load Balancing	806
8.10.6.1	Load Balancing for Internal Traffic	806
8.10.6.2	Load Balancing for External Traffic	806
8.10.7	External DNS	807
8.10.8	Use Istio as a Microservice	808
8.10.8.1	Prerequisites	808
8.10.8.2	Deploy Istio Using DKP	809
8.11	GPUs	811
8.11.1	Kommander GPU Overview	812
8.11.2	Kommander GPU configuration.....	812
8.11.2.1	Configure GPU for Kommander clusters.....	812
8.11.2.2	Prerequisites	812
8.11.2.3	Enable NVIDIA Platform Application on Kommander (Management Cluster)	813
8.11.2.4	Enable NVIDIA Platform Application on Attached or Managed Clusters.....	813
8.11.2.5	Select the Correct Toolkit Version for your NVIDIA GPU Operator.....	814
8.11.2.6	Validate that the Application has Started Correctly.....	817
8.11.2.7	NVIDIA GPU Monitoring.....	817

8.11.2.8	NVIDIA MIG Settings.....	817
8.11.2.9	Troubleshooting NVIDIA GPU Operator on Kommander.....	820
8.11.2.10	Disable NVIDIA GPU Operator Platform Application on Kommander	822
8.12	Monitoring and Alerts	822
8.12.1	Recommendations.....	823
8.12.1.1	Prometheus.....	823
8.12.2	Grafana Dashboards.....	825
8.12.2.1	Add Custom Dashboards	826
8.12.3	Cluster Metrics.....	827
8.12.4	Configure Alerts Using AlertManager.....	828
8.12.4.1	Prerequisites	828
8.12.5	Centralized Monitoring	833
8.12.5.1	Centralized Metrics.....	833
8.12.5.2	Centralized Alerts.....	835
8.12.6	Centralized Cost Monitoring	837
8.12.6.1	Centralized Costs.....	837
8.12.6.2	Related Information	838
8.12.7	Monitor Applications using Prometheus.....	839
8.12.8	Set Storage Capacity for Prometheus	841
8.13	DKP Troubleshooting.....	841
8.13.1	Generate a Support Bundle	841
8.13.1.1	Prerequisites	842
8.13.1.2	Create a Diagnostic Bundle	842
8.13.1.3	Generate a Support Bundle	843
8.13.1.4	SSH Fallback	844
8.13.2	Custom Collectors	846
8.13.2.1	Customizations	846
8.14	Storage	850
8.14.1	Ephemeral Storage	850

8.14.2	Persistent Volume.....	851
8.14.2.1	Persistent Volume Claim.....	851
8.14.2.2	Related Information.....	851
8.14.3	Provision a Static Local Volume.....	851
8.14.3.1	Before you Begin.....	852
8.14.3.2	Provision the Cluster and a Volume.....	852
8.14.4	Default Storage Providers in DKP.....	854
8.14.4.1	Multiple Storage Classes.....	855
8.14.4.2	Driver Information.....	856
8.14.4.3	On-Premises and other Storage Options.....	858
8.14.4.4	Related Information.....	858
8.14.5	Rook Ceph in DKP.....	858
8.14.5.1	Rook Ceph in DKP - Prerequisites.....	859
8.14.5.2	Rook Ceph Configuration.....	860
8.14.5.3	Rook Ceph Dashboard.....	865
8.14.5.4	BYOS (Bring Your Own Storage) to DKP Clusters.....	866
9	CLI and API Tools.....	876
9.1	API Documentation.....	876
9.1.1	apps.kommander.mesosphere.io/v1alpha2	876
9.1.1.1	App.....	876
9.1.1.2	AppDeployment.....	876
9.1.1.3	AppDeploymentList.....	876
9.1.1.4	AppDeploymentSpec.....	876
9.1.1.5	AppList.....	877
9.1.1.6	ClusterApp.....	877
9.1.1.7	ClusterAppList.....	877
9.1.1.8	CrossNamespaceGitRepositoryReference.....	878
9.1.1.9	GenericAppSpec.....	878
9.1.1.10	TypedLocalObjectReference.....	879

9.1.2	apps.kommander.mesosphere.io/v1alpha3	880
9.1.2.1	App.....	880
9.1.2.2	AppDeployment.....	880
9.1.2.3	AppDeploymentClusterCondition	880
9.1.2.4	AppDeploymentList	880
9.1.2.5	AppDeploymentSpec	882
9.1.2.6	AppDeploymentStatus.....	882
9.1.2.7	AppList.....	882
9.1.2.8	ClusterApp.....	883
9.1.2.9	ClusterAppList.....	883
9.1.2.10	ClusterConfigOverrides	883
9.1.2.11	ClusterDeploymentStatus	885
9.1.2.12	CrossNamespaceGitRepositoryReference.....	885
9.1.2.13	GenericAppSpec	886
9.1.2.14	TypedLocalObjectReference	886
9.1.3	dispatch.d2iq.io/v1alpha2.....	887
9.1.3.1	GitopsRepository	887
9.1.3.2	GitopsRepositoryList	887
9.1.3.3	GitopsRepositorySpec	887
9.1.3.4	GitopsRolloutTemplate.....	887
9.1.4	kommander.mesosphere.io/v1beta1	888
9.1.4.1	CAPIClusterReference	888
9.1.4.2	ClusterReference	888
9.1.4.3	GenericClusterReference.....	888
9.1.4.4	IngressSpec.....	888
9.1.4.5	IngressStatus	889
9.1.4.6	IssuerReference	890
9.1.4.7	KommanderCluster.....	892
9.1.4.8	KommanderClusterCondition	892

9.1.4.9	KommanderClusterList.....	892
9.1.4.10	KommanderClusterSpec	892
9.1.4.11	KommanderClusterStatus.....	893
9.1.4.12	License	895
9.1.4.13	LicenseCondition	895
9.1.4.14	LicenseExternalAWS.....	896
9.1.4.15	LicenseExternalReference.....	896
9.1.4.16	LicenseList	896
9.1.4.17	LicenseSpec	896
9.1.4.18	LicenseStatus.....	896
9.1.4.19	PlacementSelector.....	899
9.1.4.20	VirtualGroup	899
9.1.4.21	VirtualGroupClusterRoleBinding	899
9.1.4.22	VirtualGroupClusterRoleBindingList	900
9.1.4.23	VirtualGroupClusterRoleBindingSpec.....	900
9.1.4.24	VirtualGroupList	900
9.1.4.25	VirtualGroupSpec.....	900
9.1.5	workspaces.kommander.mesosphere.io/v1alpha1	901
9.1.5.1	KommanderProjectRole	901
9.1.5.2	KommanderProjectRoleList	901
9.1.5.3	KommanderProjectRoleSpec	901
9.1.5.4	KommanderProjectRoleStatus	901
9.1.5.5	KommanderWorkspaceRole	902
9.1.5.6	KommanderWorkspaceRoleList	902
9.1.5.7	KommanderWorkspaceRoleSpec	902
9.1.5.8	KommanderWorkspaceRoleStatus.....	903
9.1.5.9	Project.....	903
9.1.5.10	ProjectCondition	903
9.1.5.11	ProjectList	905

9.1.5.12	ProjectRole	905
9.1.5.13	ProjectRoleList	906
9.1.5.14	ProjectRoleSpec	906
9.1.5.15	ProjectRoleStatus	906
9.1.5.16	ProjectSpec	906
9.1.5.17	ProjectStatus	906
9.1.5.18	VirtualGroupKommanderClusterRoleBinding	906
9.1.5.19	VirtualGroupKommanderClusterRoleBindingList	907
9.1.5.20	VirtualGroupKommanderClusterRoleBindingSpec	907
9.1.5.21	VirtualGroupKommanderClusterRoleBindingStatus	908
9.1.5.22	VirtualGroupKommanderProjectRoleBinding	908
9.1.5.23	VirtualGroupKommanderProjectRoleBindingList	908
9.1.5.24	VirtualGroupKommanderProjectRoleBindingSpec	909
9.1.5.25	VirtualGroupKommanderProjectRoleBindingStatus	909
9.1.5.26	VirtualGroupKommanderWorkspaceRoleBinding	909
9.1.5.27	VirtualGroupKommanderWorkspaceRoleBindingList	911
9.1.5.28	VirtualGroupKommanderWorkspaceRoleBindingSpec	911
9.1.5.29	VirtualGroupKommanderWorkspaceRoleBindingStatus	911
9.1.5.30	VirtualGroupProjectRoleBinding	911
9.1.5.31	VirtualGroupProjectRoleBindingList	911
9.1.5.32	VirtualGroupProjectRoleBindingSpec	913
9.1.5.33	VirtualGroupProjectRoleBindingStatus	913
9.1.5.34	VirtualGroupWorkspaceRoleBinding	913
9.1.5.35	VirtualGroupWorkspaceRoleBindingList	913
9.1.5.36	VirtualGroupWorkspaceRoleBindingSpec	914
9.1.5.37	VirtualGroupWorkspaceRoleBindingStatus	914
9.1.5.38	Workspace	914
9.1.5.39	WorkspaceCondition	916
9.1.5.40	WorkspaceList	916

9.1.5.41	WorkspaceRole	917
9.1.5.42	WorkspaceRoleList	917
9.1.5.43	WorkspaceRoleSpec	917
9.1.5.44	WorkspaceRoleStatus	917
9.1.5.45	WorkspaceSpec	917
9.1.5.46	WorkspaceStatus	918
9.2	CLI Commands	919
9.2.1	CLI Commands for DKP	919
9.2.2	dkp attach	920
9.2.2.1	Options	920
9.2.2.2	Options inherited from parent commands	920
9.2.2.3	SEE ALSO	920
9.2.2.4	dkp attach cluster	920
9.2.3	dkp check	921
9.2.3.1	Options	921
9.2.3.2	Options inherited from parent commands	921
9.2.3.3	SEE ALSO	921
9.2.3.4	dkp check cluster	921
9.2.4	dkp completion	923
9.2.4.1	Synopsis	923
9.2.4.2	Options	923
9.2.4.3	Options inherited from parent commands	923
9.2.4.4	SEE ALSO	923
9.2.4.5	dkp completion fish	924
9.2.4.6	dkp completion bash	925
9.2.4.7	dkp completion zsh	926
9.2.4.8	dkp completion powershell	927
9.2.5	dkp config	928
9.2.5.1	Options	928

9.2.5.2	Options inherited from parent commands	928
9.2.5.3	SEE ALSO.....	928
9.2.5.4	dkp config get	929
9.2.5.5	dkp config set.....	930
9.2.6	dkp create.....	931
9.2.6.1	Options	931
9.2.6.2	Options inherited from parent commands	931
9.2.6.3	SEE ALSO.....	931
9.2.6.4	dkp create workspace	932
9.2.6.5	dkp create appdeployment.....	933
9.2.6.6	dkp create capi-components	934
9.2.6.7	dkp create image-bundle	935
9.2.6.8	dkp create bootstrap	936
9.2.6.9	dkp create chart-bundle.....	937
9.2.6.10	dkp create cluster	938
9.2.6.11	dkp create nodepool	956
9.2.7	dkp delete	968
9.2.7.1	Options	968
9.2.7.2	Options inherited from parent commands	968
9.2.7.3	SEE ALSO.....	968
9.2.7.4	dkp delete bootstrap.....	969
9.2.7.5	dkp delete capi-components.....	969
9.2.7.6	dkp delete cluster.....	970
9.2.7.7	dkp delete chart	971
9.2.7.8	dkp delete nodepool	972
9.2.8	dkp describe	973
9.2.8.1	Options	973
9.2.8.2	Options inherited from parent commands	973
9.2.8.3	SEE ALSO.....	973

9.2.8.4	dkp describe cluster	973
9.2.9	dkp detach.....	974
9.2.9.1	Options	974
9.2.9.2	Options inherited from parent commands.....	974
9.2.9.3	SEE ALSO.....	974
9.2.9.4	dkp detach cluster	974
9.2.10	dkp diagnose.....	975
9.2.10.1	Synopsis	975
9.2.10.2	Options	975
9.2.10.3	Options inherited from parent commands.....	976
9.2.10.4	SEE ALSO.....	976
9.2.10.5	dkp diagnose ssh.....	976
9.2.10.6	dkp diagnose default-config	976
9.2.11	dkp get.....	978
9.2.11.1	Options	978
9.2.11.2	Options inherited from parent commands.....	978
9.2.11.3	SEE ALSO.....	978
9.2.11.4	dkp get kubeconfig	978
9.2.11.5	dkp get appdeployments	979
9.2.11.6	dkp get workspaces.....	980
9.2.11.7	dkp get nodepools	981
9.2.11.8	dkp get chart	981
9.2.11.9	dkp get clusters.....	982
9.2.12	dkp import	983
9.2.12.1	Options	983
9.2.12.2	Options inherited from parent commands.....	983
9.2.12.3	SEE ALSO.....	983
9.2.12.4	dkp import image-bundle	983
9.2.13	dkp install	984

9.2.13.1	Options	984
9.2.13.2	Options inherited from parent commands	984
9.2.13.3	SEE ALSO	985
9.2.13.4	dkp install kommander	985
9.2.14	dkp move	986
9.2.14.1	Synopsis	986
9.2.14.2	Options	986
9.2.14.3	Options inherited from parent commands	986
9.2.14.4	SEE ALSO	986
9.2.14.5	dkp move capi-resources	987
9.2.15	dkp open	987
9.2.15.1	Options	988
9.2.15.2	Options inherited from parent commands	988
9.2.15.3	SEE ALSO	988
9.2.15.4	dkp open dashboard	988
9.2.16	dkp push	989
9.2.16.1	Options	989
9.2.16.2	Options inherited from parent commands	989
9.2.16.3	SEE ALSO	989
9.2.16.4	dkp push chart-bundle	989
9.2.16.5	dkp push image-bundle	990
9.2.16.6	dkp push chart	991
9.2.17	dkp scale	992
9.2.17.1	Options	992
9.2.17.2	Options inherited from parent commands	992
9.2.17.3	SEE ALSO	992
9.2.17.4	dkp scale nodepool	992
9.2.18	dkp serve	993
9.2.18.1	Options	993

9.2.18.2	Options inherited from parent commands	993
9.2.18.3	SEE ALSO.....	993
9.2.18.4	dkp serve image-bundle	993
9.2.19	dkp update.....	994
9.2.19.1	Options	994
9.2.19.2	Options inherited from parent commands	994
9.2.19.3	SEE ALSO.....	994
9.2.19.4	dkp update nodepool	994
9.2.19.5	dkp update controlplane.....	1000
9.2.19.6	dkp update bootstrap	1005
9.2.20	dkp upgrade.....	1009
9.2.20.1	Options	1009
9.2.20.2	Options inherited from parent commands	1009
9.2.20.3	SEE ALSO.....	1009
9.2.20.4	dkp upgrade kommander	1009
9.2.20.5	dkp upgrade capi-components	1011
9.2.20.6	dkp upgrade catalogapp.....	1012
9.2.20.7	dkp upgrade workspace	1013
9.2.20.8	dkp upgrade addons	1013
9.2.21	dkp edit.....	1020
9.2.21.1	Synopsis	1020
9.2.21.2	Options	1020
9.2.21.3	Options inherited from parent commands	1021
9.2.22	dkp version	1021
9.2.22.1	Synopsis	1022
9.2.22.2	Options	1022
9.2.22.3	Options inherited from parent commands	1022
9.2.23	dkp cluster.....	1022
9.2.23.1	Options	1022

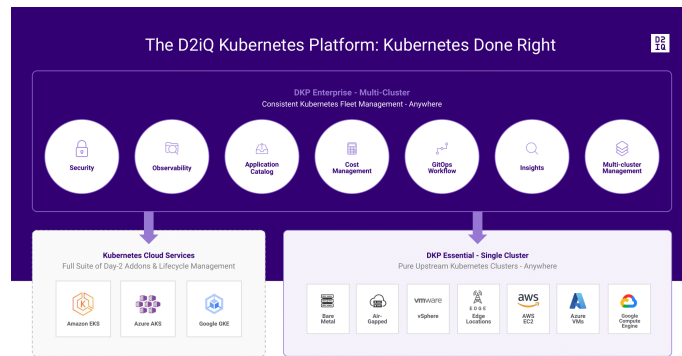
9.2.23.2	Options inherited from parent commands	1022
9.2.23.3	SEE ALSO	1023
9.2.23.4	dkp cluster type	1023
10	Upgrade DKP	1024
10.1	Prerequisite	1024
10.2	Supported Upgrade Paths	1024
10.2.1	Understand the Upgrade Process	1025
10.3	Upgrade Kommander	1026
10.3.1	Prerequisites	1026
10.3.1.1	Prerequisites for Configurations with NVIDIA GPU Nodes	1027
10.3.2	Upgrade Kommander	1029
10.3.3	Troubleshooting	1031
10.3.4	Upgrade Kommander in a Pre-provisioned Environment	1031
10.3.4.1	Prerequisites	1031
10.4	Upgrade Konvoy	1033
10.4.1	Steps to upgrade Konvoy via CLI	1033
10.4.2	DKP Enterprise Upgrade	1033
10.4.2.1	Prerequisites	1033
10.4.2.2	Overview	1034
10.4.2.3	Upgrade the CAPI Components	1034
10.4.2.4	Upgrade the Core Addons	1034
10.4.2.5	Upgrade the Kubernetes Version	1037
10.4.2.6	Upgrade Managed Clusters	1042
10.4.3	DKP Essential Upgrade	1043
10.4.3.1	Prerequisites	1043
10.4.3.2	Overview	1043
10.4.3.3	Upgrade the Core Addons	1045
10.4.3.4	Upgrade the Kubernetes version	1047
10.4.4	Upgrade Cluster Node Pools	1052

10.4.4.1	Prerequisites:	1052
10.4.4.2	Steps:	1052
10.5	Upgrade Custom Applications	1053
10.5.1	Verify the compatibility of Custom Applications with the current Kubernetes version	1053
11	Release Notes	1055
11.1	DKP 2.4.0 Release Notes.....	1055
11.1.1	Release Summary	1055
11.1.2	Additional resources.....	1056
11.1.3	DKP 2.4.0 Supported Kubernetes Versions.....	1056
11.1.3.1	Deploying Clusters Versions	1056
11.1.3.2	Attaching Clusters Versions.....	1056
11.1.4	DKP 2.4.0 Features and Enhancements.....	1057
11.1.4.1	DKP Upgrades	1057
11.1.4.2	NVIDIA GPUs	1057
11.1.4.3	RHEL 8.6 Support.....	1057
11.1.4.4	Installation Enhancements.....	1057
11.1.4.5	Cluster-scoped Application Configuration via the UI.....	1057
11.1.4.6	Enhanced KIB Documentation	1057
11.1.4.7	Complete DKP Air-gapped Bundle for Download.....	1057
11.1.4.8	Get Started and DKP Terms Section of Documentation	1059
11.1.4.9	Ability to Upgrade GCP	1059
11.1.4.10	Provision DKP Clusters on Azure Using DKP UI.....	1059
11.1.4.11	Rook Ceph Default Storage	1059
11.1.4.12	DKP Insights Enhancements.....	1059
11.1.5	DKP 2.4.0 Security Enhancements	1062
11.1.5.1	Security Enhancements for DKP and Kommander roles.....	1062
11.1.6	DKP 2.4.0 Components and Applications	1063
11.1.6.1	Components.....	1063

11.1.6.2	Applications	1064
11.1.7	DKP 2.4.0 Customer Incidents	1071
11.1.7.1	kubernetes/crictl cannot pull images from harbor registry	1071
11.1.7.2	DKP users can download cluster admin kubeconfigs	1071
11.1.7.3	nvidia-feature-discovery-gpu - error opening libnvidia-ml.so.1	1072
11.1.7.4	Cannot specify custom Packer directory in KIB 1.5.0	1072
11.1.8	DKP 2.4.0 Known Issues and Limitations	1072
11.1.8.1	Known issues and limitations	1072
11.1.9	DKP 2.4.0 Deprecations.....	1076
11.1.9.1	MinIO.....	1076
11.1.10	Kubernetes major updates and deprecations.....	1077
11.1.10.1	LegacyServiceAccountTokenNoAutoGeneration Feature Gate.....	1077
11.1.10.2	Control Plane Node Label and Taint.....	1077
11.2	DKP 2.4.1 Release Notes.....	1078
11.2.1	Release Summary	1078
11.2.2	Additional resources.....	1078
11.2.3	DKP 2.4.1 Enhancements.....	1079
11.2.3.1	Improved Installation Times for Kommander	1079
11.2.3.2	External Load Balancer Support.....	1079
11.2.4	DKP 2.4.1 Fixes and Updates	1079
11.2.4.1	Kube-oidc-proxy issue was resolved	1079
11.2.5	DKP 2.4.1 Customer Incidents	1079
11.2.5.1	Improved workflow for Deploying a Lower Version of Kubernetes	1079
11.2.5.2	KIB 1.24.x Fails to Create Ubuntu Images.....	1079
11.2.5.3	Traefik entry point TCP 9090 is tagged Velero-minio.....	1079
11.2.5.4	Upgrading Azure Breaks Volume Attachments for some Pods.....	1079
11.2.5.5	Kommander Install and Upgrade fail on Pre-provisioned	1079
11.2.5.6	Logging-operator fluentd Error.....	1080
11.2.5.7	Calico Not Updated during DKP Upgrade on Flatcar	1080

11.2.5.8	Cannot Connect Cluster with Kommander installed to an existing Management Cluster	1080
11.2.5.9	Sizing Guidelines for Loki and Logging Operator	1080
11.2.5.10	License counts Control Plane Cores post 2.4 upgrade	1080
11.2.5.11	Kommander Upgrade Fails with grafana-loki.....	1080
11.2.6	DKP 2.4.1 Components and Applications	1081
11.2.6.1	Components.....	1081
11.2.6.2	Applications	1082
11.2.7	DKP 2.4.1 Supported Kubernetes Versions.....	1089
11.2.7.1	Deploying Clusters Versions	1089
11.2.7.2	Attaching Clusters Versions.....	1089
11.2.8	DKP 2.4.1 Known Issues and Limitations	1090
11.2.8.1	Known issues and limitations	1090
12	Access Documentation	1096
12.1	Supported Documentation	1096
12.2	Archived Documentation	1096
13	Download Documentation.....	1097
14	Version support policy.....	1098
14.1	Supported Kubernetes Versions	1098
14.2	Supported Operating Systems	1098
14.3	Features in Patches	1098
14.4	Experimental Status.....	1098
14.5	Technical Preview	1099
14.6	Support Definitions	1099
14.6.1	Secondary Support	1099
14.7	Standard Level & Severity Definitions.....	1101
15	Legal Notices	1102

As the leading independent platform for Kubernetes 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 .
Cluster API	Automate cluster lifecycle management using Cluster API ³ to simplify the provisioning, upgrading, and operating multiple Kubernetes clusters across a wide range of distributions and virtual and physical environments.

1 <https://helm.sh/>

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

3 <https://cluster-api.sigs.k8s.io/>

Features	Benefits
Cluster Autoscaling	Save operational costs by automatically scaling down (see page 223) 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 751) to ensure optimal performance and troubleshooting with Grafana, Loki, and Fluentbit.
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 796) with Calico, Traefik, and CoreDNS.
Fine-Grained Cluster	Upgrades (see page 1024) Reduce operational overhead with non-disruptive patching or parallel worker node upgrades.
Backup and Recovery	Ensure business continuity and disaster recovery (see page 730) with Velero ⁴ .
Declarative Automated Installer With Day 2 Platform Services	Accelerate time-to-production on any infrastructure with consistency and reliability with the required platform services (see page 570) needed for Day 2 production.
Operate in Air-gapped Environments (see page 202)	Leverage declarative APIs to optimize cluster resources for cost, resilience, and performance.
End-to-End Support	Enterprise-grade support ⁵ and services for both Kubernetes and its supporting platform services. the entire platform including all components.

4 <https://velero.io/docs/main/>

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

1 Quick Start Guides

Use the following guides to get started in your infrastructure.

- [AWS Quick Start](#) (see page 47)
- [Azure Quick Start](#) (see page 52)
- [GCP Quick Start](#) (see page 58)
- [On Premises Quick Start](#) (see page 65)
- [vSphere Quick Start](#) (see page 65)
- [DKP Enterprise Quick Start](#) (see page 73)

1.1 AWS Quick Start

This page provides instructions for getting started with DKP, to get your Kubernetes cluster up and running with basic configuration requirements on an Amazon Web Services (AWS) public cloud instances. If you want to customize your AWS environment, see [Install AWS Advanced](#) (see page 167).

1.1.1 Prerequisites

Before starting the DKP installation, verify that you have:

- An x86_64-based Linux or macOS machine with a supported version of the operating system.
- The `dkp` binary for Linux, or macOS available on [Download DKP](#) (see page 100) page.
- [Docker](#)⁶ version 18.09.2 or later.
- [kubectl](#)⁷ for interacting with the running cluster.
- A valid AWS account with [credentials configured](#)⁸.
- [Cluster IAM Policies and Roles](#) (see page 157) set up
- [Resource requirements](#) (see page 0)

1.1.2 Configure AWS Prerequisites

Follow these steps:

1. Follow the steps in [IAM Policy Minimal Permissions to Create Clusters](#) (see page 151). This will provide the permissions to launch a DKP Cluster.
2. Follow the steps in [Cluster IAM Policy Policies and Roles](#) (see page 157) to provide the IAM configurations the DKP Cluster will use.
3. Export the AWS region where you want to deploy the cluster:

⁶ <https://docs.docker.com/get-docker/>

⁷ <https://kubernetes.io/docs/tasks/tools/#kubectl>

⁸ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.htm>

```
export AWS_REGION=us-west-2
```

4. Export the AWS Profile with the credentials that you want to use to create the Kubernetes cluster:


```
export AWS_PROFILE=<profile>
```

5. Name your cluster.


Give your cluster a unique name suitable for your environment. In AWS, it is critical that the name be unique as no two clusters in the same AWS account can have the same name.

Set the environment variable to be used throughout this documentation:


```
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 [Kubernetes⁹](#) for more naming information.

1.1.3 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.

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 \(see page 93\)](#) with 3 control plane nodes, and 4 worker nodes.

-  The default AWS image is not recommended for use in production. We suggest using [DKP Image Builder to create a custom AMI \(see page 414\)](#) to take advantage of enhanced cluster operations, and to explore the [advanced AWS installation \(see page 167\)](#) topics for more options.

⁹ <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

1. Create a Kubernetes cluster:

NOTE: To increase [Docker Hub's rate limit](https://docs.docker.com/docker-hub/download-rate-limit/)¹⁰ 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 aws \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-aws-bootstrap-credentials=true \
--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 [Configure HTTP Proxy](#) (see page 494).

2. Verify your output is similar to the following:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
Generating cluster resources
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
```

¹⁰ <https://docs.docker.com/docker-hub/download-rate-limit/>

```
configmap/nvidia-feature-discovery-aws-example created
```

- ✓ Waiting **for** cluster infrastructure to be ready
- ✓ Waiting **for** cluster control-planes to be ready
- ✓ Waiting **for** machines to be ready
- ✓ Initializing **new** CAPI components
- ✓ 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`

- ✓ Deleting bootstrap cluster

Cluster **default**/aws-example kubeconfig was written to to the filesystem.

You can now view resources in the **new** cluster by using the `--kubeconfig` flag with `kubectl`.

For example: `kubectl --kubeconfig=aws-example.conf get nodes`

As part of the underlying processing, 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 follow the workflow in [Install AWS Advanced](#) (see page 167).

1.1.4 Explore the New Kubernetes Cluster

The kubeconfig file is written to your local directory and you can now explore the cluster.

1. List the Nodes with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

2. Verify the output is similar to:

NAME	STATUS	ROLES
AGE VERSION		
ip-10-0-108-63.us-west-2.compute.internal	Ready	<none>
59m v1.22.8		
ip-10-0-115-181.us-west-2.compute.internal	Ready	<none>
59m v1.22.8		
ip-10-0-118-159.us-west-2.compute.internal	Ready	<none>
59m v1.22.8		
ip-10-0-122-136.us-west-2.compute.internal	Ready	control-plane,master
60m v1.22.8		

```
ip-10-0-122-6.us-west-2.compute.internal    Ready    <none>
59m    v1.22.8
ip-10-0-154-239.us-west-2.compute.internal Ready    control-plane,master
59m    v1.22.8
ip-10-0-199-233.us-west-2.compute.internal Ready    control-plane,master
57m    v1.22.8
```

- List the Pods with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get pods -A
```

- Verify the output is similar to:

```

NAMESPACE                                NAME
READY   STATUS    RESTARTS   AGE
calico-system                             calico-typha-665d976df-rf7jg
1/1     Running  0          60m
capa-system                               capa-controller-manager-697b7df888-vhcbj
2/2     Running  0          57m
capi-kubeadm-bootstrap-system            capi-kubeadm-bootstrap-controller-
manager-67d8fc9688-5p65s                 1/1     Running  0          57m
capi-kubeadm-control-plane-system        capi-kubeadm-control-plane-controller-
manager-846ff8b565-jqmhd                 1/1     Running  0          57m
capi-system                               capi-controller-manager-865fddc84c-9g7bb
1/1     Running  0          57m
cappp-system                             cappp-controller-manager-7859fbbb7f-xjh6k
1/1     Running  0          56m
...

```

1.1.5 Kommander Deployment

Deploy Kommander to the DKP Cluster:

```
./dkp install kommander --kubeconfig ${DKP_CLUSTER_NAME}.conf
```

If you would like to watch the Helm Releases Deploy, run the following command:

```
watch kubectl get hr -A --kubeconfig ${DKP_CLUSTER_NAME}.conf
```

1.1.6 Explore the New Kubernetes Cluster

The kubeconfig file is written to your local directory and you can now explore the cluster.

List the Nodes with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

1.1.7 Log in to the UI through Kommander

You can now [log in to the UI](#) (see page 522) to explore.

1.1.8 Delete the Kubernetes Cluster and Cleanup your Environment

Follow these steps:

Delete the provisioned Kubernetes cluster and wait a few minutes:

```
dkp delete cluster \
--cluster-name=${CLUSTER_NAME} \
--kubeconfig=${CLUSTER_NAME}.conf \
--with-aws-bootstrap-credentials=true \
--self-managed
```

Similar to `create cluster`, use the flag `--self-managed` with the `delete cluster` command:

- Creates a bootstrap cluster.
- Moves the CAPI controllers from the workload cluster back to the bootstrap cluster.
- Deletes the workload cluster.
- Deletes the bootstrap cluster.

To understand how this process works step by step, you can follow the workflow in [Delete an AWS Cluster](#) (see page 195).

If you want to customize your AWS environment, see [Install AWS Advanced](#) (see page 167) or [Install AWS Air-gapped](#) (see page 202).

1.2 Azure Quick Start

Get started by installing a cluster with default configuration settings on Azure

This Quick Start guide provides simplified instructions for using DKP to get your Kubernetes cluster up and running with minimal configuration requirements on an Azure public cloud instance. To customize your Azure installation, refer to [Azure Advanced](#) (see page 226) installation.

1.2.1 Prerequisites

Before starting the DKP installation, verify that you have:

- An x86_64-based Linux or macOS machine with a supported version of the operating system.
- The `dkp` binary on this machine available on [Download DKP](#) (see page 100) page.

- [Docker](#)¹¹ version 18.09.2 or later.
- [kubect](#)¹² for interacting with the running cluster.
- [Azure CLI](#)¹³.
- A valid Azure account with [credentials configured](#)¹⁴.
- [Resource requirements](#) (see page 0)

1.2.2 Configure Azure Prerequisites

Follow these steps:

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 commands:
 - a. If you have more than one Azure account, run this command to identify your account:

```
$(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 "Mesosphere Developer Subscription"
```

¹¹ <https://docs.docker.com/get-docker/>

¹² <https://kubernetes.io/docs/tasks/tools/#kubectl>

¹³ <https://docs.microsoft.com/en-us/cli/azure/install-azure-cli>

¹⁴ <https://github.com/kubernetes-sigs/cluster-api-provider-azure/blob/master/docs/book/src/topics/getting-started.md#prerequisites>

- 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 required environment variables:

```
export AZURE_SUBSCRIPTION_ID="<<id>"           # b1234567-abcd-11a1-
a0a0-1234a5678b90
export AZURE_TENANT_ID="<<tenant>"           # a1234567-b132-1234-1a11-1234a5678
b90
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')"
```

1.2.3 Create a New Azure Kubernetes Cluster

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 will be deployed on an [Ubuntu 20.04 operating system image](#)¹⁵ with 3 control plane nodes, and 4 worker nodes.

¹⁵ <https://docs.d2iq.com/dkp/konvoy/2.2/supported-operating-systems>

- The default Azure image is not recommended for use in production. We suggest using [Konvoy Image Builder to create a custom image](#) (see page 420) to take advantage of enhanced cluster operations, and to explore the [advanced Azure installation](#) (see page 226) topics for more options.

1. Give your cluster a name suitable for your environment:

```
export CLUSTER_NAME=azure-example
```

2. Create a Kubernetes cluster:

- To increase [Docker Hub's rate limit](#)¹⁶ 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 azure \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--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 [Configure HTTP Proxy](#) (see page 494).

You will see output similar to the following:

```
Generating cluster resources
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
```

¹⁶ <https://docs.docker.com/docker-hub/download-rate-limit/>

```

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

```

As part of the underlying processing, 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

1.2.4 Explore the New Kubernetes Cluster

The kubeconfig file is written to your local directory and you can now explore the cluster.

1. List the Nodes with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

You will see output similar to:

NAME	STATUS	ROLES	AGE
VERSION			
azure-example-control-plane-84htt v1.22.7	Ready	control-plane,master	8m11s
azure-example-control-plane-r8srg v1.22.7	Ready	control-plane,master	4m17s
azure-example-control-plane-wrdql v1.22.7	Ready	control-plane,master	6m15s
azure-example-md-0-9crp9 v1.22.7	Ready	<none>	6m47s
azure-example-md-0-dvx5d v1.22.7	Ready	<none>	6m42s
azure-example-md-0-gc9mx v1.22.7	Ready	<none>	5m27s
azure-example-md-0-tkqf7 v1.22.7	Ready	<none>	4m48s

2. List the Pods with the command:


```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get pods -A
```

You will see output similar to:

NAMESPACE	READY	STATUS	RESTARTS	AGE	NAME
calico-system	1/1	Running	0	60m	calico-typha-665d976df-rf7jg
capa-system	2/2	Running	0	57m	capa-controller-manager-697b7df888-vhcbj
capi-kubeadm-bootstrap-system	1/1	Running	0	57m	manager-67d8fc9688-5p65s
capi-kubeadm-control-plane-system	1/1	Running	0	57m	manager-846ff8b565-jqmhd
capi-system	1/1	Running	0	57m	capi-controller-manager-865fddc84c-9g7bb
capp-system	1/1	Running	0	56m	capp-controller-manager-7859fbbb7f-xjh6k
...					

1.2.5 Kommander Deployment

Deploy Kommander to the DKP Cluster:

```
./dkp install kommander --kubeconfig ${CLUSTER_NAME}.conf
```

If you would like to watch the Helm Releases Deploy, run the following command:

```
watch kubectl get hr -A --kubeconfig ${CLUSTER_NAME}.conf
```

1.2.6 Log in to the UI through Kommander

You can now [log in to the UI \(see page 522\)](#) to explore.

1.2.7 Delete the Kubernetes Cluster and Cleanup your Environment

Follow these steps:

Delete the provisioned Kubernetes cluster and wait a few minutes:

```
dkp delete cluster \
--cluster-name=${CLUSTER_NAME} \
```

```
--kubeconfig=${CLUSTER_NAME}.conf \
--self-managed
```

Similar to `create cluster`, use the flag `--self-managed` with the `delete cluster` command:

- Creates a bootstrap cluster.
- Moves the CAPI controllers from the workload cluster back to the bootstrap cluster.
- Deletes the workload cluster.
- Deletes the bootstrap cluster.

To understand how this process works step by step, you can follow the workflow in [Delete Azure Cluster \(see page 254\)](#).

To customize your Azure installation, refer to [Azure Infrastructure \(see page 226\)](#) installation.

1.3 GCP Quick Start

This Quick Start guide provides simplified instructions for using DKP to get your Kubernetes cluster up and running with minimal configuration requirements on Google Cloud Platform (GCP). For customization during setup, see [Advanced GCP Install \(see page 382\)](#).

1.3.1 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.
- The `dkp` binary on this machine available on [Download DKP \(see page 100\)](#) page.
- [Docker](#)¹⁷ version 18.09.2 or later.
- [kubect1](#)¹⁸ for interacting with the running cluster.
- Install the GCP `gcloud` CLI by following the [GCP install documentation](#)^{19,20}.

1.3.2 GCP Prerequisites

- If you are creating the bootstrap cluster on a non-GCP instance or one that does not have the required `editor` role:
 - (option 1) Create a service account using the following `gcloud` commands:

```
export GCP_PROJECT=<your GCP project ID>
export SERVICE_ACCOUNT_USER=<some new service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"
```

¹⁷ <https://docs.docker.com/get-docker/>

¹⁸ <https://kubernetes.io/docs/tasks/tools/#kubect1>

¹⁹ <https://cloud.google.com/sdk/docs/install>

²⁰ <https://cloud.google.com/sdk/docs/install>.

```
gcloud iam service-accounts create "$SERVICE_ACCOUNT_USER" --
project=$GCP_PROJECT
gcloud projects add-iam-policy-binding $GCP_PROJECT --
member="serviceAccount:
$SERVICE_ACCOUNT_USER@$GCP_PROJECT.iam.gserviceaccount.com" --
role=roles/editor
gcloud iam service-accounts keys create $GOOGLE_APPLICATION_CREDENTIALS
--iam-
account="$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 SERVICE_ACCOUNT_USER=<existing service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud iam service-accounts keys create $GOOGLE_APPLICATION_CREDENTIALS
--iam-
account="$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')
```

- Export the GCP location where you want to deploy the cluster, the default is set to `us-west1` :

```
export GCP_REGION=us-west1
```

1.3.3 Create a New GCP Kubernetes Cluster

If you use these instructions to create a cluster on GCP using the DKP default settings without any edits to configuration files or additional flags, your cluster will be deployed with 3 control plane nodes, and 4 worker nodes in the default `us-west1` region.

Follow these steps:

1. Create an image using [Konvoy Image Builder \(KIB\)](#) ([see page 424](#)) and export the image name:

```
export IMAGE_NAME=projects/$GCP_PROJECT/global/images/<image-name>
```

2. Give your cluster a name suitable for your environment:

```
export CLUSTER_NAME=gcp-example
```

3. Create a Kubernetes cluster:

NOTE: To increase [Docker Hub's rate limit](https://docs.docker.com/docker-hub/download-rate-limit/)²¹, 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 gcp \
--cluster-name=${CLUSTER_NAME} \
--additional-tags=owner=$(whoami) \
--with-gcp-bootstrap-credentials=true \
--project=$GCP_PROJECT \
--image=$IMAGE_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 [Configure HTTP Proxy \(see page 494\)](#).

You should see output similar to this example:

```
Generating cluster resources
cluster.cluster.x-k8s.io/gcp-example created
gcpcluster.infrastructure.cluster.x-k8s.io/gcp-example created
kubeadmcontrolplane.controlplane.cluster.x-k8s.io/gcp-example-control-plane
created
gcpmachinetemplate.infrastructure.cluster.x-k8s.io/gcp-example-control-plane
created
secret/gcp-example-etcd-encryption-config created
machinedeployment.cluster.x-k8s.io/gcp-example-md-0 created
gcpmachinetemplate.infrastructure.cluster.x-k8s.io/gcp-example-md-0 created
kubeadmconfigtemplate.bootstrap.cluster.x-k8s.io/gcp-example-md-0 created
clusterresourceset.addons.cluster.x-k8s.io/calico-cni-installation-gcp-example
created
configmap/calico-cni-installation-gcp-example created
configmap/tigera-operator-gcp-example created
clusterresourceset.addons.cluster.x-k8s.io/gcp-persistent-disk-gcp-example
created
configmap/gcp-persistent-disk-csi-gcp-example created
clusterresourceset.addons.cluster.x-k8s.io/cluster-autoscaler-gcp-example
created
configmap/cluster-autoscaler-gcp-example created
```

²¹ <https://docs.docker.com/docker-hub/download-rate-limit/>

```
clusterresourceset.addons.cluster.x-k8s.io/node-feature-discovery-gcp-example
created
configmap/node-feature-discovery-gcp-example created
clusterresourceset.addons.cluster.x-k8s.io/nvidia-feature-discovery-gcp-example
created
configmap/nvidia-feature-discovery-gcp-example created
```

As part of the underlying processing, 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

1.3.4 Explore the New Kubernetes Cluster

The `kubeconfig` file is written to your local directory and you can now explore the cluster.

Follow these steps:

1. List the Nodes:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

You should see output similar to this example:

NAME	STATUS	ROLES	AGE	
VERSION				
gcp-example-control-plane-9z77w	Ready	control-plane,master	4m44s	
v1.24.6				
gcp-example-control-plane-rtj9h	Ready	control-plane,master	104s	
v1.24.6				
gcp-example-control-plane-zbf9w	Ready	control-plane,master	3m23s	
v1.24.6				
gcp-example-md-0-88c46	Ready	<none>	3m28s	v1.24
.6				
gcp-example-md-0-fp8s7	Ready	<none>	3m28s	v1.24
.6				
gcp-example-md-0-qvnx7	Ready	<none>	3m28s	v1.24
.6				
gcp-example-md-0-wjdrq	Ready	<none>	3m27s	v1.24
.6				

2. List the Pods with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get pods -A
```

You should see output similar to this example:

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-dvlm	4m23s
calico-system	1/1	Running	0	calico-node-h6nlt	4m23s
calico-system	1/1	Running	0	calico-node-jmkwq	4m23s
calico-system	1/1	Running	0	calico-node-tnf54	5m23s
calico-system	1/1	Running	0	calico-node-v6bwq	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
cert-manager	1/1	Running	0	cert-manager-6888d6b69b-7b7m9	91s
cert-manager	1/1	Running	0	cert-manager-cainjector-76f7798c9-gnp8f	91s
cert-manager	1/1	Running	0	cert-manager-webhook-7d4b5d8484-gn5dr	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
lh7dc                1/1      Running    0          5m40s
node-feature-discovery
1/1      Running    0          3m40s
node-feature-discovery
1/1      Running    0          3m40s
node-feature-discovery
1/1      Running    0          3m35s
node-feature-discovery
1/1      Running    0          3m40s
tigera-operator      tigera-operator-5f9bdc5c59-j9tnr
1/1      Running    0          5m38s

```

1.3.5 Install and Log in to the UI

You can now proceed to installing the [UI in Kommander](#) (see page 475) and applications. After installation, you will be able to [log in](#) (see page 522) to the UI to explore it.

1.3.6 Delete the Kubernetes Cluster and Clean up Your Environment

Delete the provisioned Kubernetes cluster and wait a few minutes:

```

dkp delete cluster \
--cluster-name=${CLUSTER_NAME} \
--with-gcp-bootstrap-credentials=true \
--kubeconfig=${CLUSTER_NAME}.conf \
--self-managed

```

You should see output similar to this example:

```

✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
✓ 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-bootstrap.conf get nodes
✓ Waiting for cluster infrastructure to be ready
✓ Waiting for cluster control-planes to be ready
✓ Waiting for machines to be ready
✓ 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
- ✓ Deleting bootstrap cluster

1.4 On Premises Quick Start

Configure a Cluster for On Premises Operation

To create a cluster for your on premises infrastructure, follow the procedure that describes creating a cluster on a [pre-provisioned infrastructure](#) (see page 288) using SSH access.

Cluster API (CAPI) has a concept of an [Infrastructure Provider](#)²², a set of controllers tasked with managing provider-specific resources for cluster infrastructure and machines via an API. In an on premises environment, there may not be a programmatic API to create and delete infrastructure. Other teams or departments may be responsible for pre-provisioning this infrastructure. In these cases, uses its own pre-provisioned infrastructure provider that relies on having SSH access to the machines and executes the bootstrap steps through SSH.

To use the pre-provisioned infrastructure provider, use [Konvoy Image Builder](#) (see page 411) to ensure that the correct versions of kubeadm, kubelet, and containerd are installed with their relevant configurations to enable to run on your infrastructure. Then use the `dkp` CLI to create a DKP cluster on your infrastructure.

Completing this procedure produces a Kubernetes cluster, including a [Container Networking Interface \(CNI\)](#)²³ and a [Local Persistence Volume Static Provisioner](#)²⁴, 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 the [Application Manager](#) (see page 526) to select and deploy applications, or deploy your own.

To get started, follow the procedure that describes creating a DKP cluster on a [pre-provisioned infrastructure](#) (see page 288).

1.5 vSphere Quick Start

Enterprise

This guide provides instructions for getting started with DKP to get your Kubernetes cluster up and running with basic configuration requirements in a vSphere environment. If you want to customize your vSphere environment, see [vSphere Advanced Install](#) (see page 0).

1.5.1 DKP Prerequisites


Before using DKP to create a vSphere cluster, verify that you have:

²² <https://cluster-api.sigs.k8s.io/reference/glossary.html#infrastructure-provider>

²³ <https://docs.projectcalico.org/>

²⁴ <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

- An x86_64-based Linux® or macOS® machine.
- The `dkp` binaries available on [Download DKP \(see page 100\)](#) page.
- [Konvoy Image Builder \(KIB\) \(see page 411\)](#) image bundle for Linux or macOS.
- [Docker®](#)²⁵ 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. Configure this virtual machine with at least 8GB of memory.

- [kubectl](#)²⁶ 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 \(see page 329\)](#).
- [Resource requirements \(see page 116\)](#)

1.5.2 vSphere Prerequisites

Before installing, verify that your [VMware vSphere Client environment](#)²⁷ meets the following basic requirements:

- Access to a bastion VM or other network connected host.
 - 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 \(see page 329\)](#).
- A RedHat® subscription with user name and password for downloading DVD ISOs.
- For air-gapped environments, prepare your environment using a [bastion VM host template \(see page 122\)](#) with access to a configured Docker registry.
- Valid vSphere values for the following:
 - vCenter API server URL.
 - Datacenter name.
 - vCenter Cluster name that contains [ESXi hosts](#)²⁸ for your cluster's nodes.
 - Datastore name for the shared storage resource to be used for the VMs in the cluster.

²⁵ <https://docs.docker.com/get-docker/>

²⁶ <https://kubernetes.io/docs/tasks/tools/#kubectl>

²⁷ https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html

²⁸ <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.esxi.install.doc/GUID-B2F01BF5-078A-4C7E-B505-5DFFED0B8C38.html>

- 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.
- 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.

The next step is:

- For non air-gapped environments, [Create a Base OS image in vSphere](#) (see page 332) .
- For air-gapped environments, after you create and prepare a bastion VM, [Create a Base Air-gapped OS VM Image](#) (see page 354) . (Refer to your provider's site for details on bastion host set up of vSphere²⁹.)

1.5.2.1 Create directory for KIB and DKP CLI

This command creates directories for working with images created in KIB as well as a directory for running the commands for DKP:

```
mkdir kib && mkdir dkp
```

1.5.2.2 Get the needed D2iQ Software

Download and decompress KIB by running the following command:

```
cd kib
wget https://github.com/mesosphere/konvoy-image-builder/releases/download/v1.12.0/konvoy-image-bundle-v1.12.0_linux_amd64.tar.gz
tar -xvf konvoy-image-bundle-v1.12.0_linux_amd64.tar.gz
```

²⁹ <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-6975426F-56D0-4FE2-8A58-580B40D2F667.html>

Use this link to [Download and decompress DKP](#) (see page 100) and copy or move the dkp binary into the `dkp` subdirectory you created above..

1.5.2.3 Create a folder and resource pool in vCenter for DKP cluster

To create a folder in vCenter, follow the creation steps below:

1. Right click on the datacenter
2. Select New Folder
3. Select host and cluster folder
4. Name folder "D2iQ"

In order to create a Resource Pool, follow the steps below:

1. Right click on the vCenter cluster you plan to use for DKP
2. Select New Resource Pool
3. Adjust values if you need to restrict resources for DKP

1.5.2.4 Build template using KIB

Run the following command replacing `rhel-84.yaml` with `rhel-79.yaml` if necessary to create the compatible image:

```
cd ..
cd kib/images/ova
vi rhel-84.yaml
```

For troubleshooting building an image, see the following [solutions guide](#)³⁰.

1.5.2.5 Adjust the Packer file for your vSphere cluster

Execute this command to adjust the Packer file:

```
---
download_images: true
build_name: "vsphere-rhel-84"
packer_builder_type: "vsphere"
guestinfo_datasource_slug: "https://raw.githubusercontent.com/vmware/cloud-init-vmware-guestinfo"
guestinfo_datasource_ref: "v1.4.0"
```

³⁰ <https://support.d2iq.com/hc/en-us/articles/7315058530196-Common-issues-encountered-when-building-a-compliant-RHEL-template-to-be-used-with-the-vSphere-provider-in-DKP-2-2-X>

```

guestinfo_datasource_script: "{{guestinfo_datasource_slug}}/
{{guestinfo_datasource_ref}}/install.sh"
packer:
  vcenter_server: "10.0.1.52"
  vsphere_username: "administrator@vsphere.local"
  vsphere_password: "Password"
  cluster: "cluster1"
  datacenter: "dc1"
  datastore: "vsanDatastore"
  folder: "d2iq"
  insecure_connection: "true"
  network: "VM Network"
  resource_pool: "D2IQ"
  template: "rhel-boot-8.4"
  vsphere_guest_os_type: "rhel8_64Guest"
  guest_os_type: "rhel8-64"
  # goss params
  distribution: "RHEL"
  distribution_version: "8.4"

```

1.5.2.6 Create overrides for docker credentials

Navigate to the directory that contains your image and then create the override files:

```

cd ..
cd ..

```

To override the docker credentials, run the overrides file command below:

```

vi overrides.yaml
image_registries_with_auth:
- host: "registry-1.docker.io"
  username: "<dockerhub-user>"
  password: "<dockerhub-password>"
  auth: ""
  identityToken: ""

```

1.5.2.7 Build VM template using KIB

Run the following command, to build your template:

```

./konvoy-image build images/ova/rhel-84.yaml --overrides overrides.yaml

```

For more specific information regarding creating a vSphere Virtual Machine Template, you can refer to the [KIB for vSphere \(see page 430\)](#) or [Advanced Configuration for vSphere \(see page 333\)](#) sections of the documentation.

1.5.3 Create DKP Cluster on vSphere

1.5.3.1 Export your vSphere Environment Variables

Copy the set of “exports” below to a text document in a notepad so that you can modify them and copy/paste into the CLI terminal. It is recommended to save this information for later reference. Export with this command:

```
export VSPHERE_SERVER="<vcenter-server-ip-address>"
export VSPHERE_PASSWORD='<password>'
export VSPHERE_USERNAME="<administrator@vsphere.local>"
export export CLUSTER_NAME=dkp
openssl s_client -connect <vcenter_ip.:443 < /dev/null 2>/dev/null | openssl x509
-fingerprint -noout -in /dev/stdin
```

1.5.3.2 Build the Bootstrap Cluster

Execute this command to create your bootstrap cluster:

```
cd ..
cd dkp
./dkp create bootstrap --with-aws-bootstrap-credentials=false
```

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 [Configure HTTP Proxy](#) (see page 494).

1.5.3.3 Create the DKP cluster deployment YAML

If you are not using self-signed certificates, remove the last line of the command `--tls-thumb-print=` before executing the command below:

```
dkp create cluster vsphere --cluster-name="dkp" --network="VM Network" --control-
plane-endpoint-host="<vip_for_api>" --virtual-ip-interface="eth0" --data-center="<dc1>"
" --data-store="vsanDatastore" --folder="${VSPHERE_FOLDER}" --server="<vsphere_server
_ip>" --ssh-public-key-file=/root/.ssh/id_rsa.pub --resource-pool="DKP" --vm-
template=konvoy-ova-vsphere-rhel-84-1.24.6-1649344885 --tls-thumb-print="tls-
thumbprint" --dry-run -o yaml > ${DKP_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 [Configure HTTP Proxy \(see page 494\)](#).

1.5.4 Create a New vSphere Kubernetes cluster

Create/deploy a cluster using the command below:

```
kubectl create -f ${DKP_CLUSTER_NAME}.yaml
```

If you wish to watch the cluster build, run the command below:

```
dkp describe cluster -c ${DKP_CLUSTER_NAME}
```



DKP deploys MetalLB on vSphere. This is an advanced step as it will require the IP addresses managed that needed to be managed by the MetalLB and therefore not necessary for Quick Start. However, know that it is an option and see the [Configure MetalLB \(see page 352\)](#) for vSphere documentation if necessary for your configuration.

1.5.4.1 Pivot the Cluster Controllers and Create CAPI Controllers on Cluster

Perform the pivot with this command:

```
./dkp create capi-components --kubeconfig ${DKP_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 [Configure HTTP Proxy \(see page 494\)](#).

1.5.4.2 After created, move the configuration to the new cluster

Move the configuration using the command below:

```
./dkp move --to-kubeconfig ${DKP_CLUSTER_NAME}.conf
```

You now have a Self-Managing Kubernetes Cluster deployed on vSphere.

Adjust the Storage class to only allow PVs on a specific VMware datastore:

```
kubectl delete sc vsphere-raw-block-sc --kubeconfig ${DKP_CLUSTER_NAME}.conf
```

Create a storage class YAML with the URL of the VMware datastore you want to use:

```
vi sc.yaml
apiVersion: storage.k8s.io/v1
kind: StorageClass
metadata:
  annotations:
    storageclass.kubernetes.io/is-default-class: "true"
  name: vsphere-raw-block-sc
provisioner: csi.vsphere.vmware.com
allowVolumeExpansion: true
parameters:
  datastoreurl: "ds:///vmfs/volumes/vsan:5238a205736fdb1f-c71f7ec7a0353662/"
reclaimPolicy: Delete
volumeBindingMode: WaitForFirstConsumer
```

Apply the Storage class YAML to create a new default StorageClass:

```
kubectl apply -f sc.yaml --kubeconfig ${DKP_CLUSTER_NAME}.conf
```

1.5.5 Kommander Deployment

Deploy the Kommander component to the DKP Cluster:

```
./dkp install kommander --kubeconfig ${DKP_CLUSTER_NAME}.conf
```

If you would like to watch the Helm Releases Deploy, run the following command:

```
watch kubectl get hr -A --kubeconfig ${DKP_CLUSTER_NAME}.conf
```

1.5.6 Explore the New Kubernetes Cluster

The kubeconfig file is written to your local directory and you can now explore the cluster.

List the Nodes with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```


1.5.7 Log in to the UI through Kommander

You can now [log in to the UI \(see page 522\)](#) to explore.

1.5.8 Delete the Kubernetes Cluster and Cleanup your Environment

Follow these steps:

Delete the provisioned Kubernetes cluster and wait a few minutes:

```
dkp delete cluster \  
--cluster-name=${CLUSTER_NAME} \  
--kubeconfig=${CLUSTER_NAME}.conf \  
--self-managed
```

1.6 DKP Enterprise Quick Start

This page provides generic instructions for getting started with DKP, to get your Kubernetes cluster up and running with basic configuration requirements. Each Infrastructure Provider will have different steps. If your provider is listed in the [Table of Contents \(see page 47\)](#) with a Quick Start, please begin there instead. For AKS and EKS, please follow the steps below first before beginning a Quick Start for either AKS or EKS.



If your provider or environment is different than AKS or EKS, please see that provider's [Quick Start \(see page 47\)](#) page instead.

1.6.1 Prerequisites

Before starting the DKP installation, verify that you have:

- An x86_64-based Linux or macOS machine with a supported version of the operating system.
- The `dkp` binary for Linux, or macOS.
- [Docker](#)³¹ version 18.09.2 or later.
- [kubectl](#)³² for interacting with the running cluster.
- A valid infrastructure provider account with [credentials configured \(see page 128\)](#).
- [Resource requirements \(see page 116\)](#)

³¹ <https://docs.docker.com/get-docker/>

³² <https://kubernetes.io/docs/tasks/tools/#kubectl>

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 commands for creating clusters or CAPI objects for the command to be successful. More information is available in [Configure HTTP Proxy \(see page 494\)](#).

1.6.2 Create a new Kubernetes cluster

Create/deploy a cluster using the command below:

```
kubectl create -f ${DKP_CLUSTER_NAME}.yaml
```

If you wish to watch the cluster build, run the command below:

```
dkp describe cluster -c ${DKP_CLUSTER_NAME}
```

1.6.2.1 Pivot the Cluster Controllers and Create CAPI Controllers on Cluster with this command:

```
./dkp create capi-components --kubeconfig ${DKP_CLUSTER_NAME}.conf
```

1.6.2.2 Once created, move the configuration to the new cluster using the command below:

```
./dkp move --to-kubeconfig ${DKP_CLUSTER_NAME}.conf
```

You now have a Self-Managing Kubernetes Cluster deployed. Proceed to the Quick Start Guide of the Infrastructure Provider of your choice below:

- [AKS Quick Start \(see page 74\)](#)
- [EKS Quick Start \(see page 83\)](#)

1.6.3 AKS Quick Start

Enterprise

Get started by installing a cluster with the default configuration settings on AKS.

This guide provides instructions for getting started with DKP to get your Kubernetes cluster up and running with basic configuration requirements on an Azure Kubernetes Service (AKS) public cloud instance. If you want to customize your AKS environment, see [Install AKS Advanced](#) (see page 135).

1.6.3.1 DKP Prerequisites

Before starting the DKP installation, verify that you have:

- An x86_64-based Linux or macOS machine with a supported version of the [operating system](#) (see page 93).
- A [Self-managed Azure cluster](#) (see page 246).
- The `dkp` [binary](#) (see page 100) for Linux, or macOS available on [Download DKP](#) (see page 100) page.
- [Docker](#)³³ version 18.09.2 or later.
- [kubecti](#)³⁴ for interacting with the running cluster.
- The [Azure CLI](#)³⁵.
- A valid Azure account [used to sign in to the Azure CLI](#)³⁶.
- All [Resource requirements](#) (see page 0)

1.6.3.2 AKS Prerequisites

Follow these steps:

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"
    }
  }
]
```

³³ <https://docs.docker.com/get-docker/>

³⁴ <https://kubernetes.io/docs/tasks/tools/#kubecti>

³⁵ <https://docs.microsoft.com/en-us/cli/azure/install-azure-cli>

³⁶ <https://docs.microsoft.com/en-us/cli/azure/authenticate-azure-cli?view=azure-cli-latest>

```

    }
  }
]

```

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"
}

```

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 declare the version of Kubernetes you want 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>
```

6. Set the version of Kubernetes you chose. The version listed in the command is an example:


```
export KUBERNETES_VERSION=1.24.3
```

 Using Kubernetes v1.23.x is recommended for AKS.

Kubernetes Version	Upgrades
1.24.6	None available
1.24.3	1.24.6
1.23.12	1.24.3, 1.24.6
1.23.8	1.23.12, 1.24.3, 1.24.6
1.22.15	1.23.8, 1.23.12
1.22.11	1.22.15, 1.23.8, 1.23.12

1.6.3.3 Name Your Cluster

Give your cluster a unique name suitable for your environment. In AKS it is critical that the name is unique as no two clusters in the same AKS account can have the same name.

 To increase Docker Hub's rate limit 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 dkp create cluster .
```

1. Set the environment variable to be used throughout this documentation:

```
export CLUSTER_NAME=aks-example
```

2. Modify `dkp create cluster aks` to include the Kubernetes version and execute against the self-managed cluster:

```
dkp create cluster aks --cluster-name=${CLUSTER_NAME} --additional-tags=owner=$(whoami) --kubernetes-version=${KUBERNETES_VERSION} KUBECONFIG=${SELF_MANAGED_AZ_CLUSTER}.yaml
```

1.6.3.4 Create a New AKS Kubernetes Cluster

1. Create a Kubernetes cluster and execute this command:

```
dkp create cluster aks --cluster-name=${CLUSTER_NAME} --additional-tags=owner=$(whoami)
```

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 [Configure HTTP Proxy \(see page 494\)](#).

The output resembles:

```
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/cp4n2bm created
machinepool.cluster.x-k8s.io/aks-example-md-0 created
azuremanagedmachinepool.infrastructure.cluster.x-k8s.io/mpn6l25 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
```

2. (Optional) Specify an authorized key file to have SSH access to the machines.

The file must contain exactly one entry, as described in this [manual](#)³⁷.

You can use the `.pub` file that complements your private ssh key. For example, use the public key that complements your RSA private key:

```
--ssh-public-key-file=${HOME}/.ssh/id_rsa.pub
```


The default username for SSH access is `konvoy`. For example, use your own username:

```
--ssh-username=$(whoami)
```

3. 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
```

 Pivoting is not supported in AKS.

1.6.3.5 Explore the New Kubernetes Cluster

Follow these steps:

1. Fetch the kubeconfig file:

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

2. List the Nodes with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

³⁷ https://man7.org/linux/man-pages/man8/sshd.8.html#AUTHORIZED_KEYS_FILE_FORMAT

NAME	STATUS	ROLES	AGE	VERSION
aks-cp4n2bm-57672902-vmss000000	Ready	agent	28m	v1.22.6
aks-cp4n2bm-57672902-vmss000001	Ready	agent	29m	v1.22.6
aks-cp4n2bm-57672902-vmss000002	Ready	agent	29m	v1.22.6
aks-mpn6l25-57672902-vmss000000	Ready	agent	29m	v1.22.6
aks-mpn6l25-57672902-vmss000001	Ready	agent	29m	v1.22.6
aks-mpn6l25-57672902-vmss000002	Ready	agent	29m	v1.22.6
aks-mpn6l25-57672902-vmss000003	Ready	agent	29m	v1.22.6

NOTE: It may take a couple of minutes for the Status to move to `Ready` while `calico-node` pods are being deployed.

- List the Pods with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get pods -A
```

NAMESPACE	STATUS	RESTARTS	NAME	AGE	READY
calico-system	Running	0	calico-kube-controllers-78f65cd5dd-5m6t2	28m	1/1
calico-system	Running	0	calico-node-27h8r	28m	1/1
calico-system	Running	0	calico-node-cn4zw	28m	1/1
calico-system	Running	0	calico-node-fgsqx	28m	1/1
calico-system	Running	0	calico-node-htr4f	28m	1/1
calico-system	Running	0	calico-node-l7skw	28m	1/1
calico-system	Running	0	calico-node-mn67v	28m	1/1
calico-system	Running	0	calico-node-z626n	28m	1/1
calico-system	Running	0	calico-typha-b6c9c78f4-hcnmd	28m	1/1
calico-system	Running	0	calico-typha-b6c9c78f4-pz52w	28m	1/1
calico-system	Running	0	calico-typha-b6c9c78f4-xknwt	28m	1/1
kube-system	Running	0	azure-ip-masq-agent-9hxsr	30m	1/1
kube-system	Running	0	azure-ip-masq-agent-bh5m6	31m	1/1
kube-system	Running	0	azure-ip-masq-agent-c6s4v	31m	1/1

kube-system	azure-ip-masq-agent-gg77k	1/1
Running 0	30m	
kube-system	azure-ip-masq-agent-k5sl8	1/1
Running 0	31m	
kube-system	azure-ip-masq-agent-mmpsp	1/1
Running 0	31m	
kube-system	azure-ip-masq-agent-z4n24	1/1
Running 0	31m	
kube-system	cloud-node-manager-42shm	1/1
Running 0	31m	
kube-system	cloud-node-manager-b9scr	1/1
Running 0	30m	
kube-system	cloud-node-manager-ccmw1	1/1
Running 0	31m	
kube-system	cloud-node-manager-csrm1	1/1
Running 0	31m	
kube-system	cloud-node-manager-gkv6x	1/1
Running 0	31m	
kube-system	cloud-node-manager-ttxz7	1/1
Running 0	30m	
kube-system	cloud-node-manager-twlh8	1/1
Running 0	31m	
kube-system	cluster-autoscaler-68c759fbf6-cnkkp	0/1
Init:0/1 0	29m	
kube-system	coredns-845757d86-brpzs	1/1
Running 0	33m	
kube-system	coredns-845757d86-nqmlc	1/1
Running 0	31m	
kube-system	coredns-autoscaler-7d56cd888-8bc28	1/1
Running 0	33m	
kube-system	csi-azuredisk-node-4bl85	3/3
Running 0	30m	
kube-system	csi-azuredisk-node-8dw5n	3/3
Running 0	31m	
kube-system	csi-azuredisk-node-bg2kb	3/3
Running 0	31m	
kube-system	csi-azuredisk-node-fr9bm	3/3
Running 0	31m	
kube-system	csi-azuredisk-node-nm4k9	3/3
Running 0	31m	
kube-system	csi-azuredisk-node-twvcv	3/3
Running 0	31m	
kube-system	csi-azuredisk-node-wgds6	3/3
Running 0	30m	
kube-system	csi-azurefile-node-5xv28	3/3
Running 0	31m	
kube-system	csi-azurefile-node-9nl7n	3/3
Running 0	31m	
kube-system	csi-azurefile-node-c6mn9	3/3
Running 0	31m	
kube-system	csi-azurefile-node-q69zr	3/3
Running 0	31m	

kube-system	csi-azurefile-node-q894n	3/3
Running 0	31m	
kube-system	csi-azurefile-node-v2rmj	3/3
Running 0	30m	
kube-system	csi-azurefile-node-wkgck	3/3
Running 0	30m	
kube-system	kube-proxy-5kd77	1/1
Running 0	31m	
kube-system	kube-proxy-96jfn	1/1
Running 0	30m	
kube-system	kube-proxy-96pj6	1/1
Running 0	30m	
kube-system	kube-proxy-b8vzs	1/1
Running 0	31m	
kube-system	kube-proxy-fqnw4	1/1
Running 0	31m	
kube-system	kube-proxy-rvpp8	1/1
Running 0	31m	
kube-system	kube-proxy-sfqnm	1/1
Running 0	31m	
kube-system	metrics-server-6576d9ccf8-kfm5q	1/1
Running 0	33m	
kube-system	tunnelfront-78777b4fd6-g84wp	1/1
Running 0	27m	
node-feature-discovery	node-feature-discovery-master-84c67dcbb6-vgxfm	1/1
Running 0	29m	
node-feature-discovery	node-feature-discovery-worker-2htgg	1/1
Running 0	29m	
node-feature-discovery	node-feature-discovery-worker-5cpnt	1/1
Running 0	29m	
node-feature-discovery	node-feature-discovery-worker-6cjbv	1/1
Running 0	29m	
node-feature-discovery	node-feature-discovery-worker-jdmkj	1/1
Running 0	29m	
node-feature-discovery	node-feature-discovery-worker-ms749	1/1
Running 0	29m	
node-feature-discovery	node-feature-discovery-worker-p2z55	1/1
Running 0	29m	
node-feature-discovery	node-feature-discovery-worker-wnwfx	1/1
Running 0	29m	
tigera-operator	tigera-operator-74d785cb58-vbr4d	1/1
Running 0	33m	

1.6.3.6 Install and Log in to the UI

You can now install the [UI with Kommander](#) (see page 475) and applications. After installation, you will be able to [log in](#) (see page 522) to the UI to explore it.

1.6.3.7 Delete the Kubernetes Cluster and Cleanup your Environment

Follow these steps:

1. Delete the provisioned Kubernetes cluster and wait a few minutes:

```
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
```

1.6.4 EKS Quick Start

Enterprise

This guide provides instructions for getting started with DKP to get your Kubernetes cluster up and running with basic configuration requirements on an Elastic Kubernetes Service (EKS) public cloud instance. If you want to customize your EKS environment, see [EKS Advanced Install \(see page 259\)](#).

1.6.4.1 DKP Prerequisites

Before you begin using DKP, you must have:

- An x86_64-based Linux or macOS machine with a supported version of the [operating system \(see page 93\)](#).
- The `dkp` [binary \(see page 100\)](#) for Linux, or macOS available on [Download DKP \(see page 100\)](#) page.
- [Docker](#)³⁸ version 18.09.2 or later installed.
- [kubecti](#)³⁹ for interacting with the running cluster.
- [All resource requirements \(see page 0\)](#).



On macOS, Docker runs in a virtual machine. Configure this virtual machine with at least 8GB of memory.

³⁸ <https://docs.docker.com/get-docker/>

³⁹ <https://kubernetes.io/docs/tasks/tools/install-kubecti/>

1.6.4.2 AWS Prerequisites

Before you begin using DKP with AWS, you must have:

- A valid AWS account with [credentials configured](#) (see page 151).
- Installation of [aws-iam-authenticator](#)⁴⁰. This binary is used to access your cluster using kubectl. Amazon EKS uses IAM to provide authentication to your Kubernetes cluster.
- Create an [IAM policy configuration](#) (see page 265).
- 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](#).⁴¹



EKS for DKP 2.4.x is compatible with Kubernetes 1.23.x because EKS uses its own Kubernetes release cycle.

1.6.4.3 Configure EKS Prerequisites

Follow these steps:

1. Follow the steps in [EKS Cluster IAM Policy and Roles Configuration](#) (see page 0).
2. Export the AWS region where you want to deploy the EKS cluster:

```
export AWS_REGION=us-west-2
```

3. Export the AWS Profile with the credentials that you want to use to create the EKS Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```

⁴⁰ <https://docs.aws.amazon.com/eks/latest/userguide/install-aws-iam-authenticator.html>

⁴¹ <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

1.6.4.4 Name Your Cluster

Give your cluster a unique name suitable for your environment. In EKS it is critical that the name is unique as no two clusters in the same EKS account can have the same name.

Set the environment variable to be used throughout this documentation:

```
export CLUSTER_NAME=eks-example
```

Follow these steps:

1. (Optional) To get a list of names in use in your EKS account, use the `aws` CLI tool. For example:

```
aws ec2 describe-vpcs --filter "Name=tag-key,Values=kubernetes.io/cluster" --
query "Vpcs[*].Tags[?Key=='kubernetes.io/cluster'].Value | sort(@[*][0])"
```

```
"alex-eks-cluster-afe98",
"sam-aws-cluster-8if9q"
```

2. (Optional) If you want to create a cluster name that matches the example above, use this command. This creates a unique name every time you run it, so use the command with forethought.

```
export CLUSTER_NAME=eks-example-$(LC_CTYPE=C tr -dc 'a-z0-9' </dev/urandom |
fold -w 5 | head -n1)
echo $CLUSTER_NAME
```

```
eks-example-i0516
```

1.6.4.5 Create a New EKS Kubernetes Cluster

Follow these steps:

1. Make sure your AWS credentials are up to date. If you are using User Profiles, you must refresh the credentials using the command in Step 1. Otherwise, proceed to Step 2.

```
dkp update bootstrap credentials aws
```

2. Create a Kubernetes cluster:

```
dkp create cluster eks --cluster-name=${CLUSTER_NAME} --additional-tags=owner=$(whoami)
```

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 [Configure HTTP Proxy \(see page 494\)](#).

The output resembles:

```
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
```

3. (Optional) Specify an authorized key file to have SSH access to the machines.

The file must contain exactly one entry, as described in this [manual](#)⁴².

You can use the `.pub` file that complements your private ssh key. For example, use the public key that complements your RSA private key:

```
--ssh-public-key-file=${HOME}/.ssh/id_rsa.pub
```

The default username for SSH access is `konvoy`. For example, use your own username:

```
--ssh-username=$(whoami)
```

4. Wait for the cluster control-plane to be ready:

⁴² https://man7.org/linux/man-pages/man8/sshd.8.html#AUTHORIZED_KEYS_FILE_FORMAT

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

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

1.6.4.6 Explore the New Kubernetes Cluster

Follow these steps:

1. Fetch the kubeconfig file:

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

2. List the Nodes with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
ip-10-0-122-211.us-west-2.compute.internal eks-9017834	Ready	<none>	32s	v1.21.5-
ip-10-0-127-74.us-west-2.compute.internal eks-9017834	Ready	<none>	42s	v1.21.5-
ip-10-0-71-155.us-west-2.compute.internal eks-9017834	Ready	<none>	46s	v1.21.5-
ip-10-0-93-47.us-west-2.compute.internal eks-9017834	Ready	<none>	51s	v1.21.5-

NOTE: It may take a couple of minutes 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
```

NAMESPACE	NAME	READY
STATUS	RESTARTS AGE	

calico-system	calico-kube-controllers-69845d4df5-sc9vq	1/1
Running 0	44s	
calico-system	calico-node-5lppw	1/1
Running 0	44s	
calico-system	calico-node-dwbfj	1/1
Running 0	44s	
calico-system	calico-node-q6tg6	1/1
Running 0	44s	
calico-system	calico-node-rbm7c	1/1
Running 0	44s	
calico-system	calico-typha-68c68c96d-tcrxn	1/1
Running 0	35s	
calico-system	calico-typha-68c68c96d-xhrjv	1/1
Running 0	44s	
kube-system	aws-node-25bnt	1/1
Running 0	80s	
kube-system	aws-node-dr4b7	1/1
Running 0	89s	
kube-system	aws-node-mmn87	1/1
Running 0	70s	
kube-system	aws-node-z6cdb	1/1
Running 0	84s	
kube-system	cluster-autoscaler-68c759fbf6-zszxr	0/1
Init:0/1 0	9m50s	
kube-system	coredns-85d5b4454c-n54rq	1/1
Running 0	12m	
kube-system	coredns-85d5b4454c-xzd9w	1/1
Running 0	12m	
kube-system	kube-proxy-4bhzp	1/1
Running 0	84s	
kube-system	kube-proxy-5hkv9	1/1
Running 0	80s	
kube-system	kube-proxy-g82d7	1/1
Running 0	70s	
kube-system	kube-proxy-h2jv5	1/1
Running 0	89s	
node-feature-discovery	node-feature-discovery-master-84c67dcbb6-s6874	1/1
Running 0	9m50s	
node-feature-discovery	node-feature-discovery-worker-677hh	1/1
Running 0	69s	
node-feature-discovery	node-feature-discovery-worker-fvjwz	1/1
Running 0	49s	
node-feature-discovery	node-feature-discovery-worker-xcgvz	1/1
Running 0	64s	
node-feature-discovery	node-feature-discovery-worker-zctnz	1/1
Running 0	60s	
tigera-operator	tigera-operator-d499f5c8f-b56xn	1/1
Running 1	9m47s	

1.6.4.7 Install Kommander and Log in to the UI

You can now proceed to installing the [UI with Kommander](#) (see page 475) and applications. After installation, you will be able to [log in](#) (see page 522) to the UI to explore it.

1.6.4.8 Delete the Kubernetes Cluster and Cleanup Your Environment

1. Delete the provisioned Kubernetes cluster and wait a few minutes:

```
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
```

2. Delete the `kind` Kubernetes cluster:

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

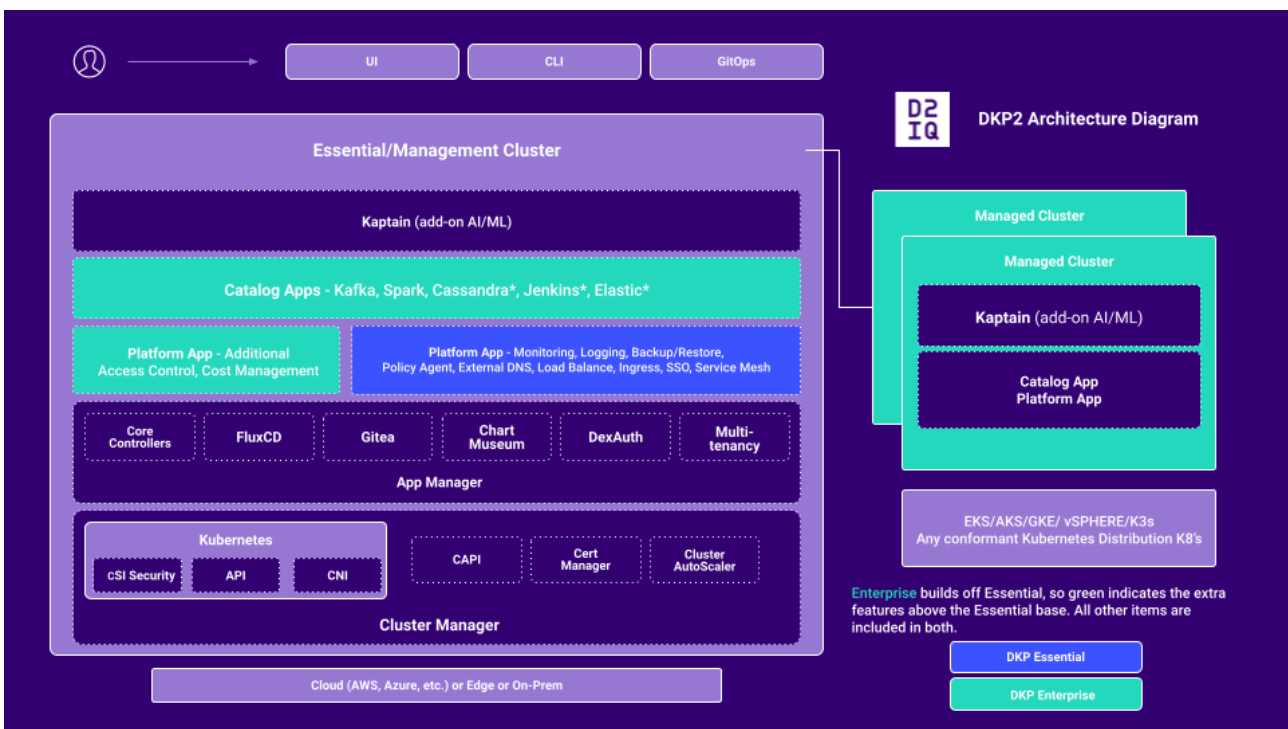
```
✓ Deleting bootstrap cluster
```

2 Architecture

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

Kubernetes provides the foundation of a DKP cluster. Because of this fundamental relationship, you should be familiar with a few key concepts and terms. The topics 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:



DKP Architectural overview

2.2 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 \(see page 526\)](#) for the full set of DKP platform services.

2.3 Related Information

For information on related topics or procedures, refer to the following [Kubernetes](#)⁴³ documentation.

2.4 DKP Ports

2.4.1 Understand the Configured Ports for DKP Deployment

This section describes ports used by the different Kubernetes components in your DKP cluster.

2.4.1.1 Control-plane nodes

Port	DKP Component	Notes
22	Ansible	ssh
179	calico-node	BGP
1338	Containerd	metrics
2379	etcd	client
2380	etcd	peer
6443	kube-apiserver	
9091	calico-node	felix metrics
9092	calico-node	bird metrics
9099	calico-node	felix liveness
9100	prometheus node-exporter	metrics

⁴³ <https://kubernetes.io/docs/concepts/overview/components/>

Port	DKP Component	Notes
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	

2.4.1.2 Worker nodes

Port	DKP Component	Notes
22	Ansible	ssh
179	calico-node	BGP
1338	Containerd	metrics
5473	calico-typha	syncserver
9091	calico-node	felix metrics
9092	calico-node	bird metrics
9093	calico-typha	metrics
9099	calico-node	felix liveness

Port	DKP Component	Notes
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	

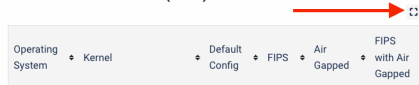
2.5 Supported Infrastructure Operating Systems

DKP supports the following base 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)



2.5.1 Amazon Web Services (AWS) Amazon Web Services (AWS)

Operating System	Kernel	Default Config	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder
CentOS 7.9 ⁴⁴	3.10.0-1160.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 7.9 ⁴⁵	3.10.0-1160.80.1.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.4 ⁴⁶	4.18.0-305.12.1.el8_4.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.6 ⁴⁷	4.18.0-372.9.1.el8.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Ubuntu 18.04 (Bionic Beaver) ⁴⁸	5.4.0-1088-aws	Yes						Yes
Ubuntu 20.04 (Focal Fossa) ⁴⁹	5.15.0-1022-aws	Yes				Yes		Yes

44 <https://wiki.centos.org/action/show/Manuals/ReleaseNotes/CentOS7.2003>

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

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

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

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

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

Operating System	Kernel	Default Config	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder
SLES 15 SP3 ⁵⁰	5.3.18-150300 .59.63-default	Yes				Yes		Yes
Oracle Linux RHCK 7.9 ⁵¹	3.10.0-1160.8 0.1.0.1.el7.x86_64	Yes	Yes					Yes

2.5.2 Azure

Operating System	Kernel	Default Config	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder
Ubuntu 20.04 (Focal Fossa) ⁵²	5.15.0-1020-azure	Yes						Yes

⁵⁰ <https://documentation.suse.com/en-us/sles/15-SP3/>

⁵¹ <https://docs.oracle.com/en/operating-systems/oracle-linux/7/relnotes7.9/>

⁵² <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

2.5.3 GCP

Operating System	Kernel	Default Config	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder
CentOS 7.9 ⁵³	3.10.0-1160.6.2.1.el7.x86_64	Yes						Yes
Ubuntu 18.04 ⁵⁴	5.4.0-1072-gcp	Yes						Yes
Ubuntu 20.04 ⁵⁵	5.13.0-1024-gcp	Yes						Yes

2.5.4 Pre-Provisioned/On Premises

Operating System	Kernel	Default Config	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder
CentOS 7.9 ⁵⁶	3.10.0-1160.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes

⁵³ <https://wiki.centos.org/action/show/Manuals/ReleaseNotes/CentOS7.2003>

⁵⁴ <https://wiki.ubuntu.com/BionicBeaver/ReleaseNotes>

⁵⁵ <https://wiki.ubuntu.com/FocalFossa/ReleaseNotes>

⁵⁶ <https://wiki.centos.org/action/show/Manuals/ReleaseNotes/CentOS7.2003>

Operating System	Kernel	Default Config	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder
RHEL 7.9 ⁵⁷	3.10.0-1160.6.6.1.el7.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.4 ⁵⁸	4.18.0-305.12.1.el8_4.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
RHEL 8.6 ⁵⁹	4.18.0-372.9.1.el8.x86_64	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Flatcar 3033.2.4 ⁶⁰	5.10.84-flatcar	Yes						Yes
Ubuntu 20.04	5.15.0-1020-aws	Yes				Yes		Yes
SLES 15 SP3 ⁶¹	5.3.18-150300.59.63-default	Yes				Yes		Yes
Oracle Linux RHCK 7.9 ⁶²	5.4.17-2011.6.2.el7uek.x86_64	Yes						Yes

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

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

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

60 <https://www.flatcar.org/releases/#stable-release>

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

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

2.5.5 vSphere

Operating System	Kernel	Default Config	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder
RHEL 7.9 ⁶³	3.10.0-1160.el7.x86_64	Yes	Yes	Yes	Yes			Yes
RHEL 8.4 ⁶⁴	4.18.0-305.el8.x86_64	Yes	Yes	Yes	Yes			Yes
RHEL 8.6 ⁶⁵	4.18.0-372.9.1.el8.x86_64	Yes	Yes	Yes	Yes			Yes

2.5.6 EKS

Operating System	Kernel	Default Config	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder
Amazon Linux 2 v7.9 ⁶⁶	5.4.204-113.362.amzn2.x86_64	Yes						

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

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

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

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

2.5.7 AKS

Operating System	Kernel	Default Config	FIPS	Air Gapped	FIPS with Air Gapped	GPU Support	GPU with Air Gapped	Konvoy Image Builder
Ubuntu 18.04 (Bionic Beaver) ⁶⁷	5.4.0-1085-azure	Yes						

⁶⁷ <https://wiki.ubuntu.com/BionicBeaver/ReleaseNotes>

3 Download DKP

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

3.1 Download from the Support Website

[Download DKP](#)

3.2 Download from the AWS Marketplace

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

After downloading the image, run the following command to copy the binaries onto your host.

```
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 \(see page 121\)](#) instructions using these binaries, and then [add your license \(see page 109\)](#) to DKP. If you have problems downloading or installing DKP, contact your sales representative or sales@d2iq.com⁶⁸.

⁶⁸ <mailto:sales@d2iq.com>

4 Licenses

This section contains overviews of the feature sets for [DKP Essential](#) (see page 107) and [DKP Enterprise](#) (see page 104) licenses, and how to acquire and use your license.

All DKP Enterprise features are marked with the Enterprise logo:

Enterprise

The type of license you purchase determines how much of DKP you can utilize. In general, DKP Essential helps you maintain your single-cluster environment, a single Management cluster, in other words. DKP Enterprise helps you maintain a multi-cluster environment, with a Management Cluster and one or more Attached or Managed Clusters. The following table shows the differences between the two types of licenses:

Feature	DKP Essential	DKP Enterprise
Applications		
Workspace Platform Applications	✓	✓
Project Platform Applications (see page 617)		✓
Catalog Applications (Kafka, Zookeeper, Spark)		✓
Custom Applications		✓
Kaptain ⁶⁹	✓	✓
DKP Insights ⁷⁰	✓	✓
Cluster Management		
Single Cluster	✓	✓
Multiple Clusters		✓

⁶⁹ <https://d2iq.atlassian.net/wiki/spaces/DKAP>

⁷⁰ <https://d2iq.atlassian.net/wiki/spaces/DINS>

Feature	DKP Essential	DKP Enterprise
Attaching Cluster		✓
Provisioning an additional cluster		✓
Upgrades (see page 1024)	✓	✓
GitOps		
Continuous Deployment (FluxCD)		✓
UX		
DKP CLI (see page 919)	✓	✓
DKP UI	✓	✓
Workspaces		✓
Projects		✓
Logging		
Workspace Level Logging	✓	✓
Multi-tenant Logging		✓
Monitoring	✓	✓

Feature	DKP Essential	DKP Enterprise
Backup & Restore (see page 730)	✓	✓
GPU	✓	✓
Security		
Authentication using Dex	✓	✓
Policy control using Gatekeeper	✓	✓
Cluster Provisioning		
DKP on AWS (see page 47)	✓	✓
DKP on Azure (see page 52)	✓	✓
DKP on GCP (see page 58)	✓	✓
DKP on vSphere (see page 65)	✓	✓
Pre-provisioned (see page 288)	✓	✓
On-Prem (see page 65)	✓	✓
EKS Provisioning (see page 83)		✓
AKS Provisioning (see page 74)		✓
FIPS Compliance	✓	✓

Feature	DKP Essential	DKP Enterprise
Konvoy Image Builder (see page 411)	✓	✓
Air-Gapped Deployments	✓	✓

4.1 Purchase a License

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

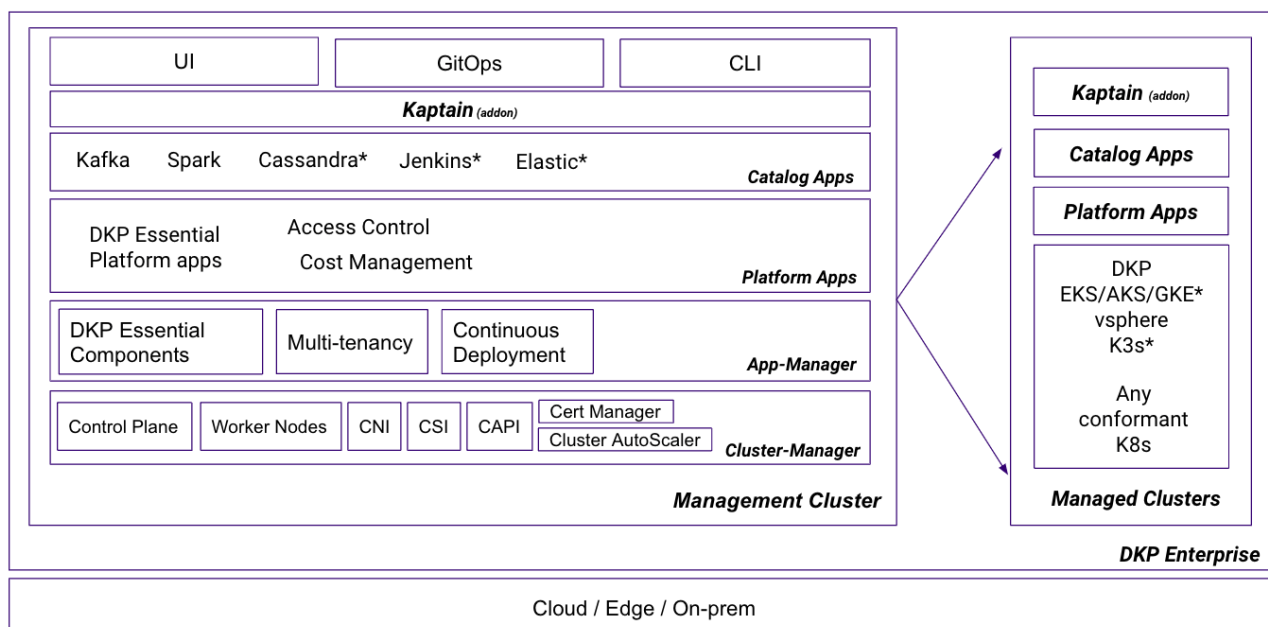
- Contact your sales representative at sales@d2iq.com⁷¹ to purchase one or more licenses.
- Purchase a license via the [AWS Marketplace](https://aws.amazon.com/marketplace/)⁷². After purchase, your license information (ARN) is accessible in the AWS License Manager console.

After purchase, [download the binary files](#) (see page 100).

4.2 DKP Enterprise



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.



⁷¹ <mailto:sales@d2iq.com>

⁷² <https://aws.amazon.com/marketplace/>

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 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.

4.2.1 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](#) (see page 93) 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 128).

4.2.2 Platform Applications

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

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 593)

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 \(see page 114\)](#)
- **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.

4.2.5 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 of DKP UI 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 \(see page 540\)](#) 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 \(see page 615\)](#).
- **Add a license:** To add a license via the UI, see [Add a License \(see page 109\)](#)

- **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, see [Workspaces](#) (see page 580)

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 93) 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 128).

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 118).

NOTE: The platform application `kubecost` is not included with DKP Essential, but is included with [DKP Enterprise](#) (see page 104).

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 \(see page 114\)](#)
- **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 \(see page 522\)](#) 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 \(see page 540\)](#) for more information.
- **Deploy applications:** The DKP Essential UI 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, see [Add a license to Kommander \(see page 109\)](#)

For more information concerning the global and workspace-level UI, see [Workspaces \(see page 580\)](#)

4.4 Add a DKP license

4.4.1 Prerequisites

For licenses that you purchase from the [AWS Marketplace](#)⁷³, an AWS administrator must attach the [AWS managed policy](#)⁷⁴ `AWSLicenseManagerConsumptionPolicy` to the `control-plane.cluster-api-provider-aws.sigs.k8s.io` role created when [configuring AWS IAM policies for Konvoy](#)⁷⁵.

Without this policy attached to the role, DKP cannot verify the license information provided in the steps below.

4.4.2 Obtain a License Token or AWS License ARN

For licenses purchased directly from D2iQ, obtain the license token via the customer support portal. Input this token in the last step below.

For licenses purchased via the AWS marketplace, 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 “License ARN.” Copy this value for use in the last step below.

4.4.3 Enter License Information



An administrator must add licenses to 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 **+ Add License** to enter the Add License form.
4. On the Add License form page, select D2iQ or AWS Marketplace depending on where you acquired your license.
5. Paste your license token or AWS ARN in the text box and select **Save**.

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

⁷³ <https://aws.amazon.com/marketplace/>

⁷⁴ <https://docs.aws.amazon.com/license-manager/latest/userguide/security-iam-awsmanpol.html#security-iam-AWSLicenseManagerConsumptionPolicy>

⁷⁵ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/9897837/Configure+AWS+IAM+Policies#Configure-IAM-Prerequisites-before-starting-a-cluster>

4.4.4 Enter a DKP License via kubectl

You can add a license acquired from D2iQ directly using `kubectl`.

1. Create a secret, replacing `MY_LICENSE` in the below command with your D2iQ-provided DKP license:

```
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:

```
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
```

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

4.5 Remove a DKP license

4.5.1 Remove a 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.

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.

4.5.2 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:  kubectll.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.

5 Get Started with DKP

When installing DKP, the first step is to determine the infrastructure on which you want to deploy.

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 edge environment.

DKP is a tool for provisioning Kubernetes clusters with a suite of pre-selected Cloud Native Computing Foundation [CNCF](#)⁷⁶ and community-contributed tools. By combining a native Kubernetes cluster as its foundation with a default set of cluster extensions, DKP provides a complete out-of-the-box solution for organizations that want to deploy production-ready Kubernetes.

After your environment is chosen, you [download DKP \(see page 100\)](#), setup a cluster with the [Konvoy \(see page 128\)](#) component, and then install the [Kommander \(see page 475\)](#) component in order to access the dashboard through the UI.



For information about installing on a specific provider, see [Advanced Konvoy Configuration \(see page 128\)](#).

Follow this list to help you get started:

- [DKP Concepts and Terms \(see page 113\)](#)
- [CAPI Concepts and Terms \(see page 114\)](#)
- [Resource Requirements \(see page 116\)](#)
- [Install Overview \(see page 121\)](#)
- [Deploy a Cluster with Konvoy \(see page 123\)](#)
- [Install Kommander \(see page 125\)](#)

5.1 DKP Concepts and Terms

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

Each main component specifically manages the following:

⁷⁶ <https://www.cncf.io/>

1. **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

2. **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
- 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 Centralised Cost Management. DKP Enterprise additionally also comes with Catalog Applications - Kafka, Spark and ZooKeeper

3. **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.



DKP Essential and Enterprise also provide some helpful add-ons called [Kaptain](#)⁷⁷ (AI/ML) and [DKP Insights](#)⁷⁸.

5.1.1 Next Topic:

[CAPI Concepts and Terms](#) (see page 114)

5.2 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](#)⁷⁹.

⁷⁷ <https://d2iq.atlassian.net/wiki/spaces/DKAP>

⁷⁸ <https://d2iq.atlassian.net/wiki/spaces/DINS>

⁷⁹ <https://cluster-api.sigs.k8s.io/user/concepts.html>

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

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.

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.

5.2.1 Next Topic:

[Resource Requirements](#) (see page 116)

5.3 Resource Requirements

To ensure a successful DKP install, there are certain requirements to be met. These resource requirements can be slightly different for different Infrastructure Providers.



The general resource requirements are listed after specific providers information.

5.3.1 Infrastructure Provider Specific

For Specific Infrastructure Providers additional requirements apply. An example would be that 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. For additional information regarding your Infrastructure Provider, see the related links below:

- Amazon ([AWS \(see page 168\)](#), [AWS Air-gapped \(see page 203\)](#) or [EKS \(see page 261\)](#))
- Microsoft Azure ([Azure \(see page 226\)](#) or [AKS \(see page 135\)](#))
- Google ([GCP \(see page 382\)](#))
- [Pre-provisioned \(see page 289\)](#)
- [vSphere \(see page 327\)](#)

5.3.2 General Resource Requirements

To [Install DKP \(see page 121\)](#) with the correct amount of resources, please see the list below before beginning installation.

5.3.3 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.

5.3.4 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.

5.3.5 More Requirement Information for the Kommander Component:

- [Management Cluster Application Requirements \(see page 117\)](#)
- [Workspace Platform Application Defaults and Resource Requirements \(see page 118\)](#)

5.3.5.1 Next Topic:

[Install Overview \(see page 121\)](#)

5.3.6 Management Cluster Application Requirements

5.3.6.1 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 \(see page 570\)](#).

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

Common Name	App ID (for App versions, see the Release Notes (see page 1064))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required
Centralized Grafana	centralized-grafana	Yes	cpu: 200m memory: 100Mi	
Centralized Kubecost	centralized-kubecost	Yes	cpu: 1200m memory: 4151Mi	# of PVs: 1 PV sizes: 32Gi
Chartmuseum	chartmuseum	Yes		# of PVs: 1 PV sizes: 2Gi
Dex	dex	Yes	cpu: 100m memory: 50Mi	

Common Name	App ID (for App versions, see the Release Notes (see page 1063))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required
Dex Authenticator	dex-k8s-authenticator	Yes	cpu: 100m memory: 128Mi	
Gitea	gitea	Yes	cpu: 500m memory: 512Mi	# of PVs: 2 PV sizes: 10Gi
Karma	karma	Yes		
Flux	kommander-flux	Yes	cpu: 5000m memory: 5Gi	
Kubefed	kubefed	Yes	cpu: 300m memory: 192Mi	
Thanos	thanos	Yes		
Traefik ForwardAuth	traefik-forward-auth-mgmt	Yes	cpu: 100m memory: 128Mi	

5.3.7 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 570).

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, and whether they are enabled by default.

Common Name	App ID (for App versions, see the Release Notes (see page 1064))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required
Cert Manager	cert-manager	Yes	cpu: 10m memory: 32Mi	
External DNS	external-dns	No		
Fluent Bit	fluent-bit	No	cpu: 350m memory: 350Mi	
Gatekeeper	gatekeeper	Yes	cpu: 300m memory: 768Mi	
Grafana	grafana-logging	No	cpu: 200m memory: 100Mi	
Loki	grafana-loki	No		# of PVs: 8 PV sizes: 10Gi x 8 (total: 80Gi)
Istio	istio	No	cpu: 1270m memory: 4500Mi	
Jaeger	jaeger	No		
Kiali	kiali	No	cpu: 20m memory: 128Mi	
Knative	knative	No	cpu: 610m memory: 400Mi	
Kube OIDC Proxy	kube-oidc-proxy	Yes		
Kube Prometheus Stack	kube-prometheus-stack	Yes	cpu: 1210m memory: 4150Mi	# of PVs: 1 PV sizes: 100Gi

Common Name	App ID (for App versions, see the Release Notes (see page 1063))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required
Kubecost	kubecost	Yes	cpu: 700m memory: 1700Mi	# of PVs: 3 PV sizes: 2Gi, 32Gi, 32Gi (total: 66Gi)
Kubernetes Dashboard	kubernetes-dashboard	Yes	cpu: 250m memory: 300Mi	
Logging Operator	logging-operator	No	cpu: 350m * # of nodes + 600m memory: 228Mi + 350Mi * # of nodes	# of PVs: 1 PV sizes: 10Gi
NFS Server Provisioner	nfs-server-provisioner	No	# of PVs: 1 PV sizes: 100Gi	
NVIDIA GPU Operator	nvidia-gpu-operator	No	cpu: 100m memory: 128Mi	
Prometheus Adapter	prometheus-adapter	Yes	cpu: 1000m memory: 1000Mi	
Reloader	reloader	Yes	cpu: 100m memory: 128Mi	
Rook Ceph	rook-ceph	Yes	cpu: 100m memory: 128Mi	
Rook Ceph Cluster	rook-ceph-cluster	Yes	cpu 2500m mem 8Gi	# of PVs: 4 PV sizes: 40Gi
Traefik	traefik	Yes	cpu: 500m	

Common Name	App ID (for App versions, see the Release Notes (see page 1063))	Deployed by default	Minimum Resources Suggested	Minimum Persistent Storage Required
Traefik ForwardAuth	traefik-forward-auth	Yes	cpu: 100m memory: 128Mi	
Velero	velero	No	cpu: 1000m memory: 1024Mi	

- 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 603).

5.4 Install Overview

5.4.1 Prerequisites

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.

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

- Version 20.10.0 of [Docker](#)⁸⁰ or higher
- CLI tooling of the cloud provider being used to deploy DKP commands:

⁸⁰ <https://www.docker.com/>

- [aws-cli](#)⁸¹
- [googlecloud-cli](#)⁸²
- [azure](#)⁸³
- Version of [kubectl](#)⁸⁴ that is within one minor version difference of the Kubernetes installed by this release.
- [Konvoy Image Builder](#) (see page 411)

DKP requires permissions for the cloud provider that is being used. See [Advanced Konvoy Configuration](#) (see page 128) for more information on your provider permissions needed. Also see the [Supported Operating Systems](#) (see page 93) page for further setup requirements.

5.4.2 Prepare Your Environment for Install:

Follow the steps below to install the basic package requirements. Then install DKP and finally you can begin any custom configuration 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. For air-gapped, create a bastion host for the cluster nodes to use within the air-gapped network. This bastion host needs access to a Docker registry in lieu of an Internet connection for pulling Docker 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.

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

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

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

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

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>

GUID-6975426F-56D0-4FE2-8A58-580B40D2F667.html

4. For creating DKP machine images, you need to download Konvoy Image Builder as mentioned in the prerequisites above and extract it.
5. To download DKP, see the [Download \(see page 100\)](#) topic for information. You will need to download and extract the DKP binary package tarball.
6. Verify you have valid **cloud provider security credentials** to deploy the cluster on that platform.

NOTE: 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 on-premises \(see page 288\)](#).

7. Install Konvoy depending on which infrastructure you have. Consult the [Advanced Konvoy Configuration \(see page 128\)](#) section of the documentation for further steps on creating a cluster with the Konvoy component of DKP on your infrastructure provider listed in that section.
8. Now that you have a basic Kubernetes cluster installed through Konvoy, configure the [Kommander \(see page 475\)](#) component.
9. Lastly, continue with install configuration through sections of documentation under the [Kommander \(see page 475\)](#) component of DKP.

5.4.3 Next Step

[Deploy a Cluster with Konvoy \(see page 123\)](#)

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

For more details, see the following topics:

- [Deploy a sample application \(see page 526\)](#)
- [Platform application deployment \(see page 619\)](#)
- [Troubleshooting \(see page 841\)](#)

5.5 Deploy a Cluster with Konvoy

You can use a single command line entry to create a Kubernetes cluster on any of the infrastructures supported by DKP. Within your environment, each cluster that you create with the `dkp create cluster` command requires a globally-unique cluster name that you specify as a flag. Each Infrastructure Provider has more specifics, but basic DKP deploy steps are listed after the infrastructure provider's Advanced Configuration section of the documentation.

Refer to your environment or provider sections for details using the links directly below. Otherwise proceed to the Create a Cluster section in chosen infrastructure provider below:

- [AKS \(see page 136\)](#)
- [AWS \(see page 167\)](#)
- [EKS \(see page 270\)](#)
- [Azure \(see page 235\)](#)

- [GCP](#) (see page 390)
- [vSphere](#) (see page 338)
- [Pre-provisioned](#) (see page 302)

5.5.1 Deploy a Cluster with Konvoy

1. The basic cluster creation command is:

```
dkp create cluster <provider> --cluster-name=clustername --self-managed --flag1=value --flag2=value ... --flagn=value
```

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](#)⁹⁰.

2. For a complete list of supported providers, enter the command:

```
dkp create cluster --help
```

The default value for the `--self-managed` flag is `false`, so you must specify the flag as `true` it to enable cluster creation with a single command.

When you execute it, this command performs the following actions:

- Creates a bootstrap cluster, if one is not present
- Deploys CAPI controllers on the bootstrap cluster
- Waits for the cluster to be created, moves the CAPI controllers, and deletes the bootstrap cluster

5.5.2 Infrastructure-specific Flags

Additional flags are available to enable needed features on supported cluster providers, and for on-premises and pre-provisioned clusters. You can view additional provider-specific flags and their descriptions with one of the following commands:

```
dkp create cluster <provider> --help
```

For more information on Infrastructure Providers and specific steps for each, see the [Advanced Configuration](#) (see page 128) section's list of providers.

[Verify Installation](#) (see page 409) before beginning the installation of Kommander.

⁹⁰ <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/246644769>

5.5.3 Next Step

Once you have done any [Advanced Configuration \(see page 128\)](#) for your provider, proceed to the [Install Kommander \(see page 125\)](#) section using the steps for your environment.

5.6 Install Kommander

5.6.1 Prerequisites

Before installing Kommander:

- Ensure you have the version of the CLI that matches the DKP version you want to install.
- You have configured a Konvoy cluster using the [Advanced Konvoy Configuration \(see page 128\)](#) for infrastructure provider specific instructions on building a cluster-based on environment.
- Review the [Management Cluster Application Requirements \(see page 117\)](#) and [Workspace Platform Application Defaults and Resource Requirements \(see page 118\)](#) to ensure that your cluster has sufficient resources.
- Ensure you have a default `StorageClass` configured as shown below.
- 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 \(see page 478\)](#)
 - [Configure custom domains and certificates after Kommander has been installed \(see page 723\)](#)

5.6.2 Configure a Default StorageClass

The cluster where Kommander is installed must have a default `StorageClass` configured. Use the following command to verify one is configured:

```
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 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](#)⁹¹ in the Kubernetes documentation.

5.6.3 Install Kommander

To customize your Kommander installation, see the [Kommander Install Configuration](#) (see page 475) for more details.

Before running the commands below, ensure that your `kubectl` configuration **references the cluster on which you want to install Kommander**, otherwise it will install on the bootstrap cluster. You can do this by setting the `KUBECONFIG` environment variable [to the appropriate kubeconfig file's location](#)⁹².

NOTE: An alternative to initializing the `KUBECONFIG` environment variable as stated earlier is to use the `--kubeconfig=cluster_name.conf` flag. This ensures that Kommander is installed on the correct cluster.

Install Kommander:

```
dkp install kommander
```

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 for the deployment of applications.

5.6.4 Verify Installation

Once the Konvoy cluster is built and Kommander has been installed, you will want to [verify your installation of Kommander](#) (see page 521) in that section of documentation.

For other environments, see the section of documentation called [Advanced Kommander Configuration](#) (see page 475).

⁹¹ <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class>

⁹² <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

5.6.5 Next Step

[Log in to the UI with Kommander \(see page 522\)](#)

6 Advanced Konvoy Configuration

When configuring DKP, the first step is to determine your infrastructure

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. If you have decided to uninstall DKP, refer to the same infrastructure documentation you have selected for specific steps to take down your cluster.

The different types of infrastructures supported in DKP are listed in this section.

Section Contents

- [Universal Configurations for all Infrastructure Providers \(see page 128\)](#)
- [Local Registry Tools for Air-gapped Environments \(see page 134\)](#)
- [AKS Infrastructure \(see page 135\)](#)
- [AWS Infrastructure \(see page 144\)](#)
- [Azure Infrastructure \(see page 226\)](#)
- [EKS Infrastructure \(see page 259\)](#)
- [Pre-provisioned Infrastructure \(see page 288\)](#)
- [vSphere Infrastructure \(see page 326\)](#)
- [GCP Infrastructure \(see page 382\)](#)
- [Verify DKP installation \(see page 409\)](#)
- [Konvoy Image Builder \(see page 411\)](#)
- [GPU for Konvoy \(see page 460\)](#)
- [FIPS 140-2 Compliance \(see page 460\)](#)
- [Delete a DKP Cluster with One Command \(see page 465\)](#)
- [Configure the Control Plane \(see page 466\)](#)
- [Update Cluster Nodepools \(see page 473\)](#)

6.1 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.

6.1.1 Use HTTP Proxy

Production environments can deny direct access to the Internet and instead have an HTTP proxy available. To configure DKP to use these proxies, set the standard environment variables in the configuration YAML file. Steps include giving valid data from the environment where the bootstrap is created. Next give the location for the production cluster. Finally the address of the API load balancer from Kommander after the cluster is created using the Konvoy component.

6.1.2 Bootstrap Proxy Settings

When creating the bootstrap, it will not be part of the workload cluster. It is possible that the portion of the bootstrap is located on a different network so you may need different HTTP proxy variables. The API server doesn't exist yet in the bootstrap environment because it is created during cluster creation.

To create a bootstrap server in a proxy environment, you need to include the following flags:

- `--http-proxy <<http proxy list>>`
- `--https-proxy <<https proxy list>>`
- `--no-proxy <<no proxy list>>`

Example:

- ```
dkp create bootstrap --http-proxy <<http proxy list>> --https-proxy <<https proxy list>> --no-proxy <<no proxy list>>
```

1. If an HTTP proxy is required, 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 \
 --http-proxy <string> \
 --https-proxy <string> \
 --no-proxy <string>
```

Example:

```
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
```

## 6.1.3 Create CAPI Components with HTTP Proxy

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>>
```

## 6.1.4 Configure and Create a Cluster with HTTP Proxy

During installation of the Konvoy component of DKP, the Control Plane and Worker nodes can be configured to use an HTTP proxy when creating a 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>        |

| Proxy configuration               | Flag                                   |
|-----------------------------------|----------------------------------------|
| No Proxy list for worker machines | <code>--worker-no-proxy strings</code> |

- The same configuration needs to be applied to the custom machine images built with Konvoy Image Builder (KIB) by using a the `http override` (see page 459) file. For more information, refer to [Use Override Files with Konvoy Image Builder](#) (see page 451) section of the documentation.

### 6.1.4.1 Example of how to configure the control plane and worker nodes to use HTTP proxy:

1. Configure control plane and worker nodes:

```
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
e
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 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
- Auto-IP addresses `169.254.169.254` for any cloud provider

### 6.1.4.2 Example of create cluster using the configured HTTP Proxy above:

1. Use your infrastructure provider name in line one 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}"
```

**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 [Configure HTTP Proxy \(see page 494\)](#).

**≡** If you have an AMI with `/etc/environment` and then add `- export $(xargs < /etc/environment)` using `preKubeadmCommands`.

### 6.1.5 HTTP Proxy for the Kommander Component of DKP

After the cluster is running in the Konvoy component from above, you will need to configure the `NO_PROXY` variable for each provider.

For example, in addition to the values above for AWS, the following settings are needed:

- 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. Set the `noProxy` environment variable to the addresses that should be accessed directly, not through the proxy. For the Kommander component of DKP, refer to more HTTP Proxy information in this document: [Configure HTTP Proxy \(see page 494\)](#)

## 6.1.6 Load Balancer

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

### 6.1.6.1 External Load Balancer

DKP includes a load balancing solution for the [supported cloud infrastructure providers \(see page 93\)](#) and for pre-provisioned environments. For more information, see [Load Balancing for external traffic \(see page 806\)](#) 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.

## 6.1.7 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 (option for Pre-provisioned or vSphere)**

If an external load balancer is not available, use the [built-in virtual IP \(see page 302\)](#). 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.

## 6.1.8 Additional Configurations

More information regarding global configurations or customization of specific components can be found in the [Additional Konvoy Configurations \(see page 128\)](#) section of the documentation.

- [Konvoy Image Builder \(see page 411\)](#)
- [FIPS 140-2 Compliance \(see page 460\)](#)

- [Configure the Control Plane](#) (see page 466)
- [Update Cluster Nodepools](#) (see page 473)

## 6.2 Local Registry Tools for Air-gapped Environments

Kubernetes does not natively provide a registry for hosting container images which contain the applications you want to deploy on Kubernetes. Instead, Kubernetes expects you to use an external solution for storing and sharing container images.

There are a variety of Kubernetes registry options out there that are compatible with DKP. The list below refers to the local registry tools required if running an air-gapped environment.

### 6.2.1 Air-Gapped Registry Prerequisites

DKP in an air-gapped environment requires a local container registry 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 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.

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.

#### 6.2.1.1 JFrog Artifactory

JFrog 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](#)<sup>93</sup> 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.

#### 6.2.1.2 Nexus Registry

[Nexus Repository](#)<sup>94</sup> is a package registry for all of 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.

#### 6.2.1.3 Harbor Registry

[Install Harbor](#)<sup>95</sup> and configure any https 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

---

<sup>93</sup> <https://jfrog.com/artifactory/>

<sup>94</sup> <https://www.nexusregistry.com/info/>

<sup>95</sup> <https://goharbor.io/docs/2.0.0/install-config/download-installer/>

information about upgrading, see [Upgrading Harbor](#)<sup>96</sup>. Any newer version than **Harbor Registry v2.1.1-5f52168e** will support OCI images.

#### 6.2.1.4 Bastion Host

If you have not set up a Bastion Host yet, refer to that section of the Documentation.

## 6.3 AKS Infrastructure

### Enterprise

When installing DKP on Azure Kubernetes Service (**AKS**) infrastructure, you can choose from multiple configuration types. The different types of AKS configuration types supported in DKP are covered in this section.

- [AKS Prerequisites](#) (see page 135)
- [Create a New AKS Cluster](#) (see page 136)
- [Explore New AKS Cluster](#) (see page 139)
- [Delete AKS Cluster](#) (see page 143)

### 6.3.1 AKS Prerequisites

#### Enterprise

Before you begin using DKP with AKS, you must have:

- An x86\_64-based Linux or macOS machine.
- A [Self-managed Azure Cluster](#) (see page 246)
- Installed [kubectl](#)<sup>97</sup> for interacting with the running cluster.
- Installed [Azure CLI](#)<sup>98</sup>.
- A valid Azure account with [credentials configured](#)<sup>99</sup>.

Ensure that the `KUBECONFIG` environment variable is set to the self-managed cluster by running `export KUBECONFIG={SELF_MANAGED_AZURE_CLUSTER}.conf`

<sup>96</sup> <https://goharbor.io/docs/2.0.0/administration/upgrade/>

<sup>97</sup> <https://kubernetes.io/docs/tasks/tools/#kubectl>

<sup>98</sup> <https://docs.microsoft.com/en-us/cli/azure/install-azure-cli>

<sup>99</sup> <https://github.com/kubernetes-sigs/cluster-api-provider-azure/blob/main/docs/book/src/topics/getting-started.md#prerequisites>

## 6.3.2 Create a New AKS Cluster

Enterprise

### 6.3.2.1 DKP to create a new AKS cluster

Ensure that the `KUBECONFIG` environment variable is set to the self-managed cluster by running `export KUBECONFIG=${SELF_MANAGED_AZURE_CLUSTER}.conf`

### 6.3.2.2 Create a New AKS Kubernetes Cluster

1. Set the environment variable to a name for this cluster.

```
export CLUSTER_NAME=aks-example
```

See [Get Started with AKS \(see page 74\)](#) for information on naming your cluster.

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:** Using Kubernetes v1.24.x is recommended. The version listed in the command is an example.

```
export KUBERNETES_VERSION=1.24.6
```

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 [Configure HTTP Proxy \(see page 494\)](#).



```

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

```

### 6.3.2.3 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](#)<sup>100</sup> 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](#)<sup>101</sup> in the Cluster API Book.

1. Wait for the cluster control-plane to be ready:

<sup>100</sup> <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

<sup>101</sup> <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.

- Once 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:}]

```

### 6.3.2.4 Known Limitations



**NOTE:** Be aware of these limitations in the current release of DKP.

- 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 139).

## 6.3.3 Explore New AKS Cluster

Enterprise

### 6.3.3.1 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 136).

### 6.3.3.2 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:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get nodes
```

| NAME                            | STATUS | ROLES | AGE | VERSION |
|---------------------------------|--------|-------|-----|---------|
| aks-cp6dsz8-41174201-vmss000000 | Ready  | agent | 56m | v1.24.6 |
| aks-cp6dsz8-41174201-vmss000001 | Ready  | agent | 55m | v1.24.6 |
| aks-cp6dsz8-41174201-vmss000002 | Ready  | agent | 56m | v1.24.6 |
| aks-mp6gglj-41174201-vmss000000 | Ready  | agent | 55m | v1.24.6 |
| aks-mp6gglj-41174201-vmss000001 | Ready  | agent | 55m | v1.24.6 |
| aks-mp6gglj-41174201-vmss000002 | Ready  | agent | 55m | v1.24.6 |
| aks-mp6gglj-41174201-vmss000003 | Ready  | agent | 56m | v1.24.6 |



**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-tgs1m | 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 | calico-node-m7s7x                   | 1/1 |
| Running 0     | 3m58s                               |     |
| 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 143).

## 6.3.4 Delete AKS Cluster

### Enterprise

- Ensure that the `KUBECONFIG` environment variable is set to the self-managed cluster by running `export KUBECONFIG={SELF_MANAGED_AZURE_CLUSTER}.conf`

### 6.3.4.1 Delete the AKS cluster and clean up your environment

### 6.3.4.2 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. 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
```

### 6.3.4.3 Next Step:

Once your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will [Install the Kommander \(see page 475\)](#) component of DKP to see your dashboard and continue customization.

### 6.3.4.4 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.

## 6.4 AWS Infrastructure

### 6.4.1 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.

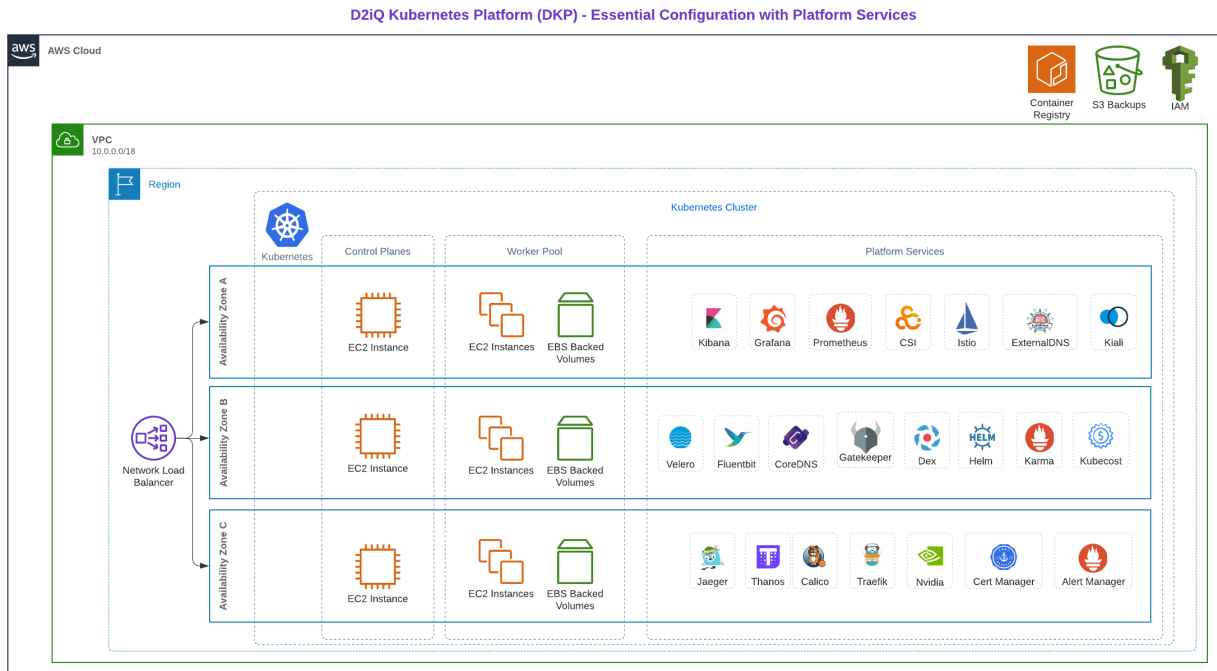
- [AWS Konvoy Image Builder \(see page 146\)](#)
- [Minimal Permissions and Role to Create Clusters \(see page 151\)](#)
- [Cluster IAM Policies and Roles \(see page 157\)](#)
- [Multiple AWS Accounts \(see page 164\)](#)
- [Advanced AWS Install \(see page 167\)](#)
- [Install AWS Air-Gapped \(see page 202\)](#)
- [GPUs in an AWS environment \(see page 214\)](#)
- [Manage AWS Node Pools \(see page 217\)](#)
- [Creating DKP Clusters on AWS \(see page 225\)](#)

### 6.4.2 AWS Diagrams

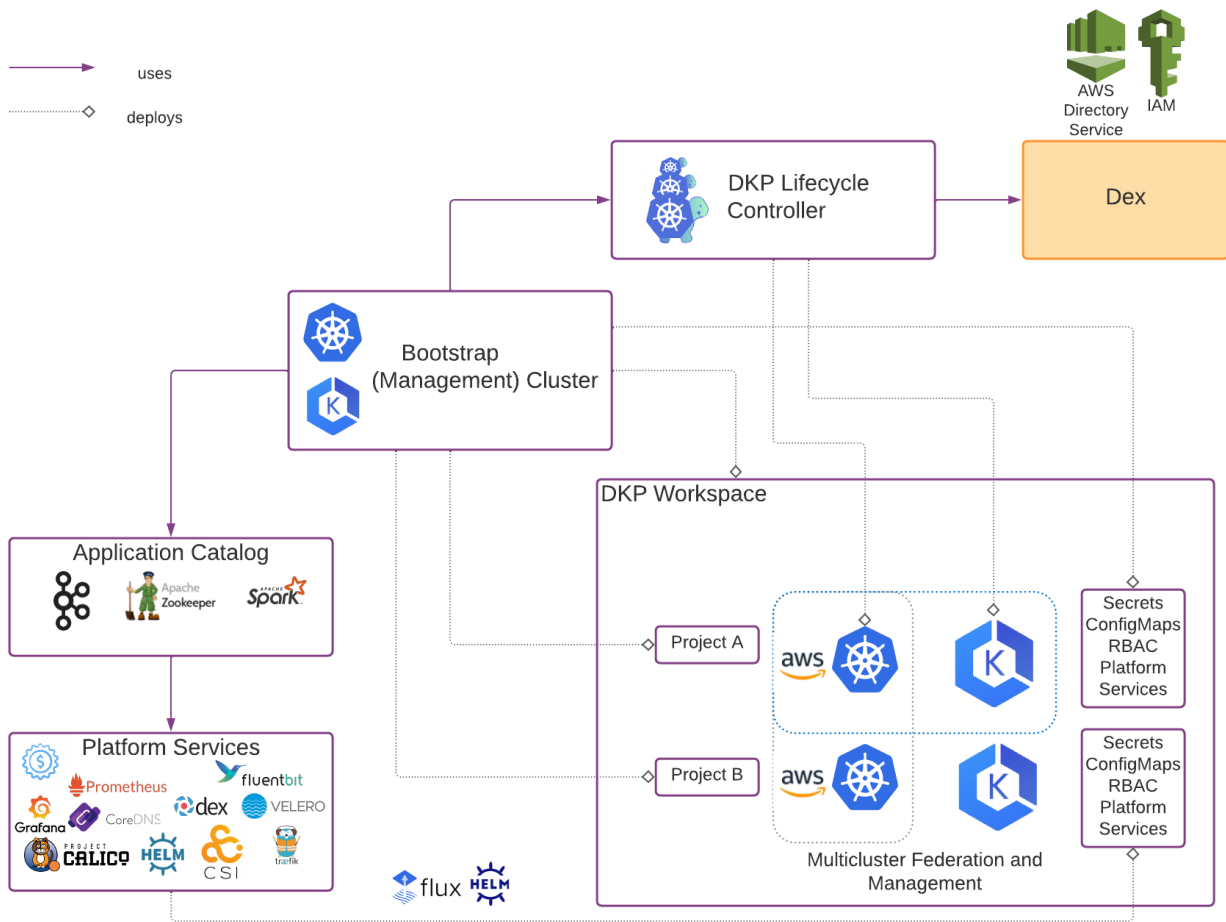
The following diagrams show the two different ways you can implement DKP Essential and DKP Enterprise on AWS.

This diagram shows the granular detail of a single Kubernetes cluster running in AWS Cloud:





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:



### 6.4.3 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](#)<sup>102</sup>.

### 6.4.4 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](#)<sup>103</sup>.

### 6.4.5 AWS Konvoy Image Builder

The following section describes how to use Konvoy Image Builder (KIB) with Amazon Web Services (AWS). There are two options:

102 <https://calculator.aws/#/>


103 <https://aws.amazon.com/premiumsupport/knowledge-center/manage-service-limits/>

 [Konvoy Image Builder](#) (see page 411) has a more detailed section in the documentation if you need to refer there for compatible versions with DKP and other specific information.

## 6.4.5.1 1. Create a Custom AMI

### 6.4.5.1.1 Learn how to build a custom AMI for use with DKP

This procedure describes how to use the [Konvoy Image Builder](#) (see page 411) (KIB) to create a [Cluster API](#)<sup>104</sup> compliant Amazon Machine Image (AMI). 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. The KIB uses [variable overrides](#) (see page 451) 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](#)<sup>105</sup><sup>106</sup> to take advantage of enhanced cluster operations, and to explore the [Advanced AWS Install](#) (see page 167) topics for more options.

## 6.4.5.2 Prerequisites

Before you begin, you must:

- Check the [DKP Supported Kubernetes Versions](#) (see page 1056) and download the [KIB](#) (see page 0) bundle (prefixed with `konvoy-image-bundle`) for your OS. Do not use the release prefixed with `konvoy-image-builder`.
- Create a working `Docker` setup.
- Ensure you have met the minimal set of permissions from the [AWS Image Builder Book](#)<sup>107</sup>.

## 6.4.5.3 Extract AMI 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.

<sup>104</sup> <https://cluster-api.sigs.k8s.io/>

<sup>105</sup> <https://d2iq.atlassian.net/wiki/pages/resumedraft.action?draftId=143891417>

<sup>106</sup> <https://d2iq.atlassian.net/wiki/pages/resumedraft.action?draftId=143891417>

<sup>107</sup> <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>

**[-]** 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](#)<sup>108</sup>. The following variables must be set using your credentials including [required IAM](#) (see page 157):

```
export AWS_ACCESS_KEY_ID
export AWS_SECRET_ACCESS_KEY
export AWS_DEFAULT_REGION
```

- Ensure you have an [override file](#) (see page 451) to configure specific attributes of your AMI file.

## 6.4.5.4 Build the Image

Depending on which [version of DKP](#) (see page 411) 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](#)<sup>109</sup> information from AWS.

### 6.4.5.4.1 Execute the following to begin image creation:

Run the `konvoy-image` command to build and validate the image.

```
konvoy-image build 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 --region us-east-1 images/ami/centos-79.yaml
```

When the command is complete the `ami` id is printed and written to `./manifest.json`.

## 6.4.5.5 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=$(whoami)-aws-cluster --ami ami-0123456789
```

<sup>108</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-envvars.html>

<sup>109</sup> <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

### 6.4.5.6 Launch a DKP Cluster with Custom AMI Lookup

By default `konvoy-image` 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=$(whoami)-aws-cluster --ami-format "konvoy-ami-{{.BaseOS}}-?{{.K8sVersion}}-*" --ami-base-os centos-7 --ami-owner 123456789012
```

Using custom source AMIs When using KIB **for** building machine images to Amazon, the **default** source AMIs that we provide are based on looking up an AMI based on the owner, and a filter **for** that operating system and version. You can view an example of that with the provided `centos-79.yaml` snippet below:

```
yaml
download_images:
 truepacker:
 ami_filter_name: "CentOS 7.9.2009 x86_64"
 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, or perhaps you want to provide a custom AMI **for** whatever reason you need. 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, [CentOS also provides an image](<https://wiki.centos.org/Cloud/AWS>) on the [AWS marketplace](<https://aws.amazon.com/marketplace/pp/prodview-foff247vr2zfw>) which you can subscribe to for free. Once you select the source AMI that you want, you can declare that when running your build command:

```
bash
konvoy-image build path/to/ami/centos-79.yaml --source-ami ami-0123456789
```

Alternatively, if you want to add it to your yaml file, or are making your own file, you can do that as well. You just need to add that AMI ID into the `source_ami` in the yaml file:

```
yaml
download_images:
 truepacker:
 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:

```
bash
konvoy-image build path/to/ami/centos-79.yaml
```

### 6.4.5.7 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:

```
yaml
```

```
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, [CentOS](#)<sup>110</sup> also provides an image on the [AWS marketplace](#)<sup>111</sup> or you can select other 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 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:

```
yaml
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 path/to/ami/centos-79.yaml
```

### 6.4.5.8 Related Information

For information on related topics or procedures, refer to the following:


<sup>110</sup> <https://wiki.centos.org/Cloud/AWS>

<sup>111</sup> <https://aws.amazon.com/marketplace/pp/prodview-foff247vr2zfw>

- [Creating GPU enabled on-premises configurations](#) (see page 214)

## 6.4.5.9 2. Air-gapped AMI

### 6.4.5.9.1 Create an AMI using Konvoy Image Builder (KIB) for use in an air-gapped cluster

 The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to [AWS Air-gapped AMI](#)<sup>112113</sup> to take advantage of enhanced cluster operations, and to explore the [Install AWS Air-Gapped](#) (see page 202) topics for more options.

Using [KIB](#) (see page 411), you can build an AMI without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded](#) (see page 100) `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2/kib
```

2. Follow the instructions to [build an AMI](#) (see page 414) and set the override `--overrides overrides/offline.yaml` flag.

After you complete these steps, you can [seed your docker registry](#) (see page 204).

## 6.4.6 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.

#### 6.4.6.1 Prerequisites

Before applying the IAM Policies, verify the following:

- You have a valid AWS account with [credentials configured](#)<sup>114</sup> that can manage CloudFormation Stacks, IAM Policies, IAM Roles, and IAM Instance Profiles.

<sup>112</sup> <https://d2iq.atlassian.net/wiki/pages/resumedraft.action?draftId=144117789>

<sup>113</sup> <https://d2iq.atlassian.net/wiki/pages/resumedraft.action?draftId=144117789>

<sup>114</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

- The [AWS CLI utility](https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html)<sup>115</sup> is installed.

### 6.4.6.2 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.

### 6.4.6.3 Create Resources in CloudFormation Stack

To create the resources in the cloudformation stack:

1. Copy the following contents into a file:

```
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
```

<sup>115</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html>



- 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: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
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: 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
```

#### 6.4.6.4 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": "BSssyvSsdfJY74jubSadfsafdsaH7x1L+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**

### 6.4.6.5 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.

### 6.4.6.6 Use Access Keys

AWS administrators can attach the `dkp-bootstrapper-policy` to an [existing IAM user](#)<sup>116</sup> and authenticate with [Access Keys](#)<sup>117</sup> 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](#)<sup>118</sup>.

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](#)<sup>119</sup>.

## 6.4.7 Cluster IAM Policies and Roles

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.

### 6.4.7.1 Prerequisites

Before applying the IAM Policies, verify the following:

- You have a valid AWS account with [credentials configured](#)<sup>120</sup> that can manage CloudFormation Stacks, IAM Policies, IAM Roles, and IAM Instance Profiles.
- You have the [AWS CLI utility installed](#)<sup>121</sup>.

### 6.4.7.2 Next Step:

After you meet the prerequisites, proceed to the next page for the Role, Instance Profile and Policy information:

- [IAM Artifacts](#) (see page 158)

<sup>116</sup> [https://docs.aws.amazon.com/IAM/latest/UserGuide/access\\_policies\\_manage-attach-detach.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/access_policies_manage-attach-detach.html)

<sup>117</sup> [https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_credentials\\_access-keys.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_credentials_access-keys.html)

<sup>118</sup> <https://docs.aws.amazon.com/general/latest/gr/aws-access-keys-best-practices.html>

<sup>119</sup> <https://docs.aws.amazon.com/kms/latest/developerguide/key-policy-default.html#key-policy-service-integration>

<sup>120</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

<sup>121</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html>

### 6.4.7.3 IAM Artifacts

Below is information about the following areas of setup. After reading the information for each of these areas, you will find the CloudFormation Stack that creates:

#### 6.4.7.3.1 Policies

1. `AWSIAMManagedPolicyCloudProviderControlPlane` enumerates the Actions required by the workload cluster control plane machines. It is attached to the `AWSIAMRoleControlPlane` Role.
2. `AWSIAMManagedPolicyCloudProviderNodes` enumerates the Actions required by the workload cluster worker machines. It is attached to the `AWSIAMRoleNodes` Role.
3. `AWSIAMManagedPolicyControllers` enumerates the Actions required by the workload cluster worker machines. It is attached to the `AWSIAMRoleControlPlane` Role.

#### 6.4.7.3.2 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](#)<sup>122</sup>.

#### 6.4.7.3.3 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.

1. `AWSIAMInstanceProfileNodes`, assigned to workload cluster worker machines.

<sup>122</sup> <https://docs.aws.amazon.com/eks/latest/userguide/add-user-role.html>

- 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:

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
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
- ssmessages:CreateControlChannel
- ssmessages:CreateDataChannel
- ssmessages:OpenControlChannel
- ssmessages: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: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
- 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 and run the following command:

```

aws cloudformation create-stack --template-body=file://MYFILENAME.yaml --stack-
name=MYSTACKNAME --capabilities CAPABILITY_NAMED_IAM

```

replacing `MYFILENAME.yaml` and `MYSTACKNAME` with the intended values.

#### 6.4.7.3.4 Next Steps:

- For non-air-gapped AWS, refer to the [Advanced AWS Install](#) (see page 167) documentation.
- For air-gapped AWS, refer to the [Install AWS Air-Gapped](#) (see page 202) documentation.

## 6.4.8 Multiple AWS Accounts

### Leverage Multiple AWS Accounts for Kubernetes Cluster Deployments

#### 6.4.8.1 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](#)<sup>123</sup>.

This document describes how to leverage the D2iQ Kubernetes Platform (DKP) to deploy a management cluster, and multiple workload clusters, leveraging multiple AWS accounts.

## 6.4.8.2 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](#)<sup>124</sup>

Getting Started with DKP on AWS - [AWS Quick Start](#) (see page 47)

## 6.4.8.3 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.

## 6.4.8.4 Prerequisites

Before you begin deploying DKP on AWS, you configure the prerequisites for the environment you use either non-air-gapped or air-gapped.

- [AWS Install Prerequisites](#) (see page 168)
- [AWS Air-gapped Prerequisites](#) (see page 203)

## 6.4.8.5 Deploy DKP on AWS

1. Deploy a management cluster in your AWS source account.  
**AWS:** [create Kubernetes AWS cluster](#) (see page 175)
2. Configure a trusted relationship between source and target accounts and create a management cluster:

---

<sup>123</sup> <https://docs.aws.amazon.com/whitepapers/latest/organizing-your-aws-environment/benefits-of-using-multiple-aws-accounts.html>

<sup>124</sup> <https://cluster-api.sigs.k8s.io/user/concepts.html>

### 6.4.8.5.1 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
```

### 6.4.8.5.2 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 \(see page 157\)](#).
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.

#### 6.4.8.6 Next Steps:

- For non-air-gapped AWS, refer to the [Advanced AWS Install](#) (see page 167) documentation.
- For air-gapped AWS, refer to the [Install AWS Air-Gapped](#) (see page 202) documentation.

### 6.4.9 Advanced AWS Install

This section provides AWS advanced configuration information to use with DKP.

- [AWS Install Prerequisites](#) (see page 168)
- [Custom AMI](#) (see page 169)
- [Bootstrap AWS](#) (see page 173)

- [Create a New AWS Cluster](#) (see page 175)
- [Explore New AWS Cluster](#) (see page 186)
- [Make the New AWS Cluster Self-Managed](#) (see page 189)
- [AWS Certificate Renewal](#) (see page 192)
- [Configure Infrastructure in UI](#) (see page 194)
- [Delete an AWS Cluster](#) (see page 195)
- [Replace an AWS Node](#) (see page 199)

## 6.4.9.1 AWS Install Prerequisites

### 6.4.9.1.1 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.
- [Docker](#)<sup>125</sup> version 18.09.2 or later installed.
- [kubect](#)<sup>126</sup> for interacting with the running cluster.
- A valid AWS account with [credentials configured](#)<sup>127</sup>.



On macOS, Docker runs in a virtual machine. Configure this virtual machine with at least 8GB of memory.

### 6.4.9.1.2 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 AWS defaults to deploying an `m5.xlarge` instance with an 80GiB root volume for control plane nodes, which meets the above requirements.

---

<sup>125</sup> <https://docs.docker.com/get-docker/>

<sup>126</sup> <https://kubernetes.io/docs/tasks/tools/#kubect>

<sup>127</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>



### 6.4.9.1.3 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 93) 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](#) (see page 414) to take advantage of enhanced cluster operations.

### 6.4.9.1.4 AWS Prerequisites

Before you begin using Konvoy with AWS, you must:

- Create an [IAM policy configuration](#) (see page 151).
- Create [IAM Policies and Roles](#) (see page 157).
- 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>
```


## 6.4.9.2 Custom AMI

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.

#### 6.4.9.2.1 Learn how to build a custom AMI for use with DKP

This procedure describes how to use the [Konvoy Image Builder](#) (see page 411) (KIB) to create a [Cluster API](#)<sup>128</sup> compliant Amazon Machine Image (AMI). 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. The KIB uses [variable overrides](#) (see page 451) 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](#)<sup>129</sup><sup>130</sup> to take advantage of enhanced cluster operations, and to explore the [Advanced AWS Install](#) (see page 167) topics for more options.


#### 6.4.9.2.2 Prerequisites

Before you begin, you must:

- Check the [DKP Supported Kubernetes Versions](#) (see page 1056) and download the [KIB](#) (see page 0) bundle (prefixed with `konvoy-image-bundle`) for your OS. Do not use the release prefixed with `konvoy-image-builder`.
- Create a working `Docker` setup.
- Ensure you have met the minimal set of permissions from the [AWS Image Builder Book](#)<sup>131</sup>.

#### 6.4.9.2.3 Extract AMI 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.

<sup>128</sup> <https://cluster-api.sigs.k8s.io/>

<sup>129</sup> <https://d2iq.atlassian.net/wiki/pages/resumedraft.action?draftId=143891417>

<sup>130</sup> <https://d2iq.atlassian.net/wiki/pages/resumedraft.action?draftId=143891417>

<sup>131</sup> <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>

- Set environment variables for [AWS access](#)<sup>132</sup>. The following variables must be set using your credentials including [required IAM](#) (see page 157):

```
export AWS_ACCESS_KEY_ID
export AWS_SECRET_ACCESS_KEY
export AWS_DEFAULT_REGION
```

- Ensure you have an [override file](#) (see page 451) to configure specific attributes of your AMI file.

#### 6.4.9.2.4 Build the Image

Depending on which [version of DKP](#) (see page 411) 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](#)<sup>133</sup> information from AWS.

##### 6.4.9.2.4.1 Execute the following to begin image creation:

Run the `konvoy-image` command to build and validate the image.

```
konvoy-image build 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 --region us-east-1 images/ami/centos-79.yaml
```

When the command is complete the `ami` id is printed and written to `./manifest.json`.

#### 6.4.9.2.5 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=$(whoami)-aws-cluster --ami ami-0123456789
```

#### 6.4.9.2.6 Launch a DKP Cluster with Custom AMI Lookup

By default `konvoy-image` 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:

<sup>132</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-envvars.html>

<sup>133</sup> <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

```
dkp create cluster aws --cluster-name=$(whoami)-aws-cluster --ami-format "konvoy-ami-{{.BaseOS}}-?{{.K8sVersion}}-*" --ami-base-os centos-7 --ami-owner 123456789012
```

Using custom source AMIs When using KIB **for** building machine images to Amazon, the **default** source AMIs that we provide are based on looking up an AMI based on the owner, and a filter **for** that operating system and version. You can view an example of that with the provided `centos-79.yaml` snippet below:

```
yaml
download_images:
 truepacker:
 ami_filter_name: "CentOS 7.9.2009 x86_64"
 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, or perhaps you want to provide a custom AMI **for** whatever reason you need. 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, [CentOS also provides an image](<https://wiki.centos.org/Cloud/AWS>) on the [AWS marketplace](<https://aws.amazon.com/marketplace/pp/prodview-foff247vr2zfw>) which you can subscribe to for free. Once you select the source AMI that you want, you can declare that when running your build command:

```
bash ./konvoy-image build path/to/ami/centos-79.yaml --source-ami ami-0123456789
```

Alternatively, if you want to add it to your yaml file, or are making your own file, you can do that as well. You just need to add that AMI ID into the `source_ami` in the yaml file:

```
yaml
download_images:
 truepacker:
 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:

```
bash konvoy-image build path/to/ami/centos-79.yaml
```

#### 6.4.9.2.7 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:

```
yaml
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, [CentOS](#)<sup>134</sup> also provides an image on the [AWS marketplace](#)<sup>135</sup> or you can select other 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 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:

```
yaml
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 path/to/ami/centos-79.yaml
```

#### 6.4.9.2.8 Related Information

For information on related topics or procedures, refer to the following:

- [Creating GPU enabled on-premises configurations](#) (see page 214)

#### 6.4.9.3 Bootstrap AWS

To create Kubernetes clusters, Konvoy uses [Cluster API](#)<sup>136</sup> (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](#)<sup>137</sup>) tool.

<sup>134</sup> <https://wiki.centos.org/Cloud/AWS>

<sup>135</sup> <https://aws.amazon.com/marketplace/pp/prodview-foff247vr2zfw>

<sup>136</sup> <https://cluster-api.sigs.k8s.io/>

<sup>137</sup> <https://github.com/kubernetes-sigs/kind>

### 6.4.9.3.1 Prerequisites

Before you begin, you must:

- Complete the steps in [Prerequisites](#) (see page 168).
- Ensure the `dkp` binary can be found in your `$PATH`.

### 6.4.9.3.2 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
```

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 [Configure HTTP Proxy](#) (see page 494).

The output resembles:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

Konvoy creates a bootstrap cluster using [KIND](#)<sup>138</sup> as a library. Konvoy then deploys the following [Cluster API](#)<sup>139</sup> providers on the cluster:

- [Core Provider](#)<sup>140</sup>
- [AWS Infrastructure Provider](#)<sup>141</sup>
- [Kubeadm Bootstrap Provider](#)<sup>142</sup>
- [Kubeadm ControlPlane Provider](#)<sup>143</sup>

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
```

138 <https://github.com/kubernetes-sigs/kind>

139 <https://cluster-api.sigs.k8s.io/>

140 <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

141 <https://github.com/kubernetes-sigs/cluster-api-provider-aws>

142 <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

143 <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>

| 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                       |
| cappp-system                      | 1/1   | 1          | 1         | 2m6s  | cappp-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                          |

## 6.4.9.4 Create a New AWS Cluster

### 6.4.9.4.1 Prerequisites

Before you begin, make sure you have created a [Bootstrap](#) (see page 173) cluster.

### 6.4.9.4.2 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:

```
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](#)<sup>144</sup> for more naming information.

### 6.4.9.4.3 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 177) section.

- To get a list of names used in your AWS account, use the `aws CLI`<sup>145</sup>. 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"
```

- (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.

- To use a custom AMI when creating your cluster, you must create that AMI using [KIB](#) (see page 414) first. Then perform the export and name the custom AMI. Then set the environment variable for the AMI you choose for use in the second command `dkp create cluster` :

```
export AWS_AMI_ID=ami-<ami-id-here>
```

<sup>144</sup> <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

<sup>145</sup> <https://aws.amazon.com/cli/>



After export, run the following command:

```
dkp create cluster aws --cluster-name=${CLUSTER_NAME} \
--ami=${AWS_AMI_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 [Configure HTTP Proxy](#) (see page 494).

#### 6.4.9.4.4 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 940) reference for the full list of cluster creation options:

**NOTE:** To increase [Dockerhub's rate limit](#)<sup>146</sup> 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 aws --cluster-name=${CLUSTER_NAME} \
--dry-run \
--output=yaml \
> ${CLUSTER_NAME}.yaml
```

<sup>146</sup> <https://docs.docker.com/docker-hub/download-rate-limit/>

The output resembles:

```
Generating cluster resources
```

3. (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,.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:

**NOTE:** To increase [Dockerhub's rate limit](#)<sup>147</sup> use your Dockerhub credentials when creating the cluster, by setting flags `--registry-mirror-url=https://registry-1.docker.io --registry-mirror-username= --registry-mirror-password=` then running `dkp create cluster`

```
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
```

#### 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](#)<sup>148</sup> 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

<sup>147</sup> <https://docs.docker.com/docker-hub/download-rate-limit/>

<sup>148</sup> <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

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](#)<sup>149</sup> 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 466).
7. Create the cluster from the objects.

```
kubectl create -f ${CLUSTER_NAME}.yaml
```

```
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
```

<sup>149</sup> <https://cluster-api.sigs.k8s.io/user/concepts.html>

```
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"

```

#### 6.4.9.4.5 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.

### 6.4.9.5 Explore New AWS Cluster

This guide explains how to use the command line to interact with your newly deployed Kubernetes cluster.

Before you begin, ensure you created a workload cluster as described in [Create a New AWS Cluster](#) (see page 175).

#### 6.4.9.5.1 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.

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                |
|--------------------------------------------|--------|----------------------|
| ip-10-0-100-85.us-west-2.compute.internal  | Ready  | <none>               |
| 8m13s v1.24.6                              |        |                      |
| ip-10-0-106-183.us-west-2.compute.internal | Ready  | control-plane,master |
| 6m22s v1.24.6                              |        |                      |
| ip-10-0-158-104.us-west-2.compute.internal | Ready  | control-plane,master |
| 8m58s v1.24.6                              |        |                      |
| ip-10-0-203-138.us-west-2.compute.internal | Ready  | control-plane,master |
| 7m52s v1.24.6                              |        |                      |
| ip-10-0-70-169.us-west-2.compute.internal  | Ready  | <none>               |
| 8m14s v1.24.6                              |        |                      |
| ip-10-0-77-176.us-west-2.compute.internal  | Ready  | <none>               |
| 8m11s v1.24.6                              |        |                      |

```
ip-10-0-96-61.us-west-2.compute.internal Ready <none>
8m12s v1.24.6
```

**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     | NAME                                     | READY | STATUS   | RESTARTS | AGE   |
|---------------|------------------------------------------|-------|----------|----------|-------|
| calico-system | calico-kube-controllers-57fbd7bd59-xlnfq | 1/1   | Running  | 0        | 9m41s |
| calico-system | calico-node-7lfmj                        | 1/1   | Running  | 0        | 9m2s  |
| calico-system | calico-node-8jm52                        | 1/1   | Running  | 0        | 9m41s |
| calico-system | calico-node-cgrwx                        | 1/1   | Running  | 0        | 9m23s |
| calico-system | calico-node-g96f6                        | 1/1   | Running  | 0        | 9m22s |
| calico-system | calico-node-kvwgq                        | 1/1   | Running  | 0        | 9m21s |
| calico-system | calico-node-nbjr6                        | 1/1   | Running  | 0        | 7m32s |
| calico-system | calico-node-w5sb5                        | 1/1   | Running  | 0        | 9m24s |
| calico-system | calico-typha-8699c7db75-645pm            | 1/1   | Running  | 0        | 9m22s |
| calico-system | calico-typha-8699c7db75-7czwp            | 1/1   | Running  | 0        | 9m12s |
| calico-system | calico-typha-8699c7db75-xk6wn            | 1/1   | Running  | 0        | 9m41s |
| kube-system   | cluster-autoscaler-68c759fbf6-2bc1x      | 0/1   | Init:0/1 | 0        | 10m   |
| kube-system   | coredns-78fcd69978-882jg                 | 1/1   | Running  | 0        | 10m   |
| kube-system   | coredns-78fcd69978-ntp5f                 | 1/1   | Running  | 0        | 10m   |
| kube-system   | ebs-csi-controller-77574b5cf5-84z7l      | 6/6   | Running  | 0        | 10m   |
| kube-system   | ebs-csi-controller-77574b5cf5-tsr5q      | 6/6   | Running  | 0        | 10m   |

```

kube-system ebs-csi-node-2j9pd
3/3 Running 0 9m23s
kube-system ebs-csi-node-6t4sr
3/3 Running 0 9m2s
kube-system ebs-csi-node-8qkpr
3/3 Running 0 10m
kube-system ebs-csi-node-kbq57
3/3 Running 0 9m24s
kube-system ebs-csi-node-qp9mh
3/3 Running 0 9m22s
kube-system ebs-csi-node-v84dd
3/3 Running 0 9m21s
kube-system ebs-csi-node-zqrtp
3/3 Running 0 7m32s
kube-system etcd-ip-10-0-106-183.us-west-2.compute.internal
1/1 Running 0 7m31s
kube-system etcd-ip-10-0-158-104.us-west-2.compute.internal
1/1 Running 0 10m
kube-system etcd-ip-10-0-203-138.us-west-2.compute.internal
1/1 Running 0 9m1s
kube-system kube-apiserver-ip-10-0-106-183.us-west-2.compute.inter
nal 1/1 Running 0 7m31s
kube-system kube-apiserver-ip-10-0-158-104.us-west-2.compute.inter
nal 1/1 Running 0 10m
kube-system kube-apiserver-ip-10-0-203-138.us-west-2.compute.inter
nal 1/1 Running 0 9m1s
kube-system kube-controller-manager-ip-10-0-106-183.us-west-2.comp
ute.internal 1/1 Running 0 7m31s
kube-system kube-controller-manager-ip-10-0-158-104.us-west-2.comp
ute.internal 1/1 Running 1 (8m51s ago) 10m
kube-system kube-controller-manager-ip-10-0-203-138.us-west-2.comp
ute.internal 1/1 Running 0 9m2s
kube-system kube-proxy-4j9s5
1/1 Running 0 9m23s
kube-system kube-proxy-64svf
1/1 Running 0 9m2s
kube-system kube-proxy-9xghm
1/1 Running 0 7m32s
kube-system kube-proxy-cwbqm
1/1 Running 0 9m24s
kube-system kube-proxy-p6thh
1/1 Running 0 9m22s
kube-system kube-proxy-pgs47
1/1 Running 0 9m21s
kube-system kube-proxy-zrplb
1/1 Running 0 10m
kube-system kube-scheduler-ip-10-0-106-183.us-west-2.compute.inter
nal 1/1 Running 0 7m31s
kube-system kube-scheduler-ip-10-0-158-104.us-west-2.compute.inter
nal 1/1 Running 1 (8m51s ago) 10m
kube-system kube-scheduler-ip-10-0-203-138.us-west-2.compute.inter
nal 1/1 Running 0 9m2s

```

```

kube-system snapshot-controller-545b6bf98-k2fbv
1/1 Running 0 10m
kube-system snapshot-controller-545b6bf98-nf2m8
1/1 Running 0 10m
node-feature-discovery node-feature-discovery-master-84c67dcbb6-gqfq8
1/1 Running 0 10m
node-feature-discovery node-feature-discovery-worker-46b9r
1/1 Running 0 8m12s
node-feature-discovery node-feature-discovery-worker-g6dw7
1/1 Running 0 8m12s
node-feature-discovery node-feature-discovery-worker-jzkmt
1/1 Running 0 8m12s
node-feature-discovery node-feature-discovery-worker-vmsbm
1/1 Running 0 8m12s
tigera-operator tigera-operator-d499f5c8f-gl25
1/1 Running 1 (8m51s ago) 10m

```

### 6.4.9.6 Make the New 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.

Before starting, ensure you create a workload cluster as described in [Create a New Cluster](#) (see page 175).

#### 6.4.9.6.1 Make the New Kubernetes Cluster Manage Itself

1. Deploy cluster lifecycle services on the workload cluster:

By default, `create bootstrap controllers` configures the Cluster API controllers to use the AWS credentials from your environment. We recommend you use the `--with-aws-bootstrap-credentials=false` flag to configure the Cluster API controllers of your self-managed AWS cluster to use AWS IAM Instance Profiles, instead of the AWS credentials from your environment.

```

dkp create capi-components --with-aws-bootstrap-credentials=false --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 [Configure HTTP Proxy](#) (see page 494).

The output resembles:

```

✓ 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](#)<sup>150</sup>.

```
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}
```

<sup>150</sup> <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```

NAME READY
SEVERITY REASON SINCE MESSAGE
Cluster/aws-example True
109s
|─ClusterInfrastructure - AWSCluster/aws-example True
112s
|─ControlPlane - KubeadmControlPlane/aws-example-control-plane True
109s
| |─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

```

- Remove the bootstrap cluster, as the workload cluster is now self-managed:

```
dkp delete bootstrap
```

```
✓ Deleting bootstrap cluster
```

#### 6.4.9.6.2 Install and Log in to the DKP UI

You can now proceed to installing the [DKP UI \(see page 475\)](#) and applications. After installation, you will be able to [log in \(see page 522\)](#) to the DKP UI to explore it.

#### 6.4.9.6.3 DKP Install Configuration

You can [configure \(see page 475\)](#) 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 \(see page 206\)](#) installation with catalog applications.

- Depending on your configuration, there are different ways you can install DKP to an environment.  
**Ensure you follow the correct procedure for your configuration type, and ignore the section that does not pertain to your environment:**
  - [Install Kommander in a Non-air-gapped Environment \(see page 513\)](#)
  - [Install Kommander in an Air-gapped Environment \(see page 504\)](#)

#### 6.4.9.6.4 Known Limitations

- Be aware of these limitations in the current release of Konvoy.

- 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 \(see page 157\)](#).
- 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.

#### 6.4.9.7 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.

##### 6.4.9.7.1 Prerequisites

- This feature requires Python 3.5 or greater to be installed on all control plane hosts.
- Complete the [Bootstrap Cluster \(see page 173\)](#) topic.



### 6.4.9.7.2 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.

### 6.4.9.7.3 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.

### 6.4.9.7.4 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.internal -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
```

## 6.4.9.8 Configure Infrastructure in UI

This page refers to working inside the AWS Console in order to add Infrastructure Provider.

### 6.4.9.8.1 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](#)<sup>151</sup>.
9. Select **Save** to save your provider.

#### 6.4.9.8.2 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.

#### 6.4.9.9 Delete an AWS 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 189](#)), see [Delete the workload cluster](#) (see [page 198](#)).

Follow these steps to prepare to delete an AWS 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
```

2. Create a bootstrap cluster:

<sup>151</sup> [https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_create\\_for-user\\_externalid.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user_externalid.html)

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 --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](#)<sup>152</sup>.

```
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 status:

```
dkp describe cluster --kubeconfig $HOME/.kube/config -c ${CLUSTER_NAME}
```

| NAME | REASON | SINCE | MESSAGE | READY | SEVERITY |
|------|--------|-------|---------|-------|----------|
|------|--------|-------|---------|-------|----------|

<sup>152</sup> <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```

Cluster/aws-example True
91s
├─ClusterInfrastructure - AWSCluster/aws-example True
103s
├─ControlPlane - KubeadmControlPlane/aws-example-control-plane True
91s
│ ├─Machine/aws-example-control-plane-55jh4 True
102s
│ ├─Machine/aws-example-control-plane-6sn97 True
102s
│ └─Machine/aws-example-control-plane-nx9v5 True
102s
└─Workers
 └─MachineDeployment/aws-example-md-0 True
108s
 ├─Machine/aws-example-md-0-cb9c9bbf7-hcl8z True
102s
 ├─Machine/aws-example-md-0-cb9c9bbf7-rtdqw True
102s
 ├─Machine/aws-example-md-0-cb9c9bbf7-td29r True
102s
 └─Machine/aws-example-md-0-cb9c9bbf7-w64kg True
102s

```

**NOTE:** 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.

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
```



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.

### 6.4.9.1 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
```

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, delete the bootstrap cluster.

### 6.4.9.9.2 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
```

### 6.4.9.9.3 Next Step:

After your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will install the [Kommander](#) (see page 475) component of DKP to see your dashboard and continue customization.

### 6.4.9.9.4 Known Limitations

 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.

## 6.4.9.10 Replace an AWS Node

### 6.4.9.10.1 Prerequisites

Before you begin, you must:

- [Create a workload cluster](#) (see page 175).
- [Make the new cluster self-managed](#) (see page 189).

### 6.4.9.10.2 Replace a Worker Node

In certain situations, you may want to delete a worker node and have [Cluster API](#)<sup>153</sup> 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<br>16m v1.24.6  | Ready  | <none>               |
| ip-10-0-106-183.us-west-2.compute.internal<br>15m v1.24.6 | Ready  | control-plane,master |
| ip-10-0-158-104.us-west-2.compute.internal<br>17m v1.24.6 | Ready  | control-plane,master |
| ip-10-0-203-138.us-west-2.compute.internal<br>16m v1.24.6 | Ready  | control-plane,master |
| ip-10-0-70-169.us-west-2.compute.internal<br>16m v1.24.6  | Ready  | <none>               |
| ip-10-0-77-176.us-west-2.compute.internal<br>16m v1.24.6  | Ready  | <none>               |
| ip-10-0-96-61.us-west-2.compute.internal<br>16m v1.24.6   | 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

<sup>153</sup> <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 once the Machine resource has been deleted.

A few minutes after the Machine resource is deleted, the corresponding Node resource is also deleted.

4. 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 | v1.24.6 | 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.

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=="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<br>20m v1.24.6 | Ready  | control-plane,master |
| ip-10-0-158-104.us-west-2.compute.internal<br>23m v1.24.6 | Ready  | control-plane,master |
| ip-10-0-203-138.us-west-2.compute.internal<br>22m v1.24.6 | Ready  | control-plane,master |
| ip-10-0-70-169.us-west-2.compute.internal<br>22m v1.24.6  | Ready  | <none>               |
| ip-10-0-77-176.us-west-2.compute.internal<br>22m v1.24.6  | Ready  | <none>               |
| ip-10-0-86-58.us-west-2.compute.internal<br>57s v1.24.6   | Ready  | <none>               |
| ip-10-0-96-61.us-west-2.compute.internal<br>22m v1.24.6   | 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.

## 6.4.10 Install AWS Air-Gapped

### 6.4.10.1 Create a Kubernetes cluster in a private subnet with no access to the Internet (air-gapped)

This section provides the instructions to create a Kubernetes cluster in an existing VPC with the nodes and the kube-apiserver ELB in a private subnet with no access to the Internet.

- [AWS Air-gapped Prerequisites](#) (see page 203)
- [Air-gapped AMI](#) (see page 203)
- [AWS Air-gapped Seed Docker Registry](#) (see page 204)
- [AWS Air-gapped Bootstrap](#) (see page 205)
- [Create a New AWS Air-gapped Cluster](#) (see page 206)
- [Explore New AWS Air-gapped Cluster](#) (see page 209)

- [Make the Air-gapped AWS Cluster Self-Managed](#) (see page 210)
- [Delete AWS Air-gapped Cluster](#) (see page 213)

## 6.4.10.2 AWS Air-gapped Prerequisites

### 6.4.10.2.1 Air-gapped installation prerequisites

#### 6.4.10.2.2 Prerequisites

Before you begin, you must have:

- Linux machine (bastion) that has access to the existing VPC.
- The `dkp` binary on the bastion.
- [Docker](#)<sup>154</sup> version 18.09.2 or later installed on the bastion.
- [kubectl](#)<sup>155</sup> for interacting with the running cluster on the bastion.
- Valid AWS account with [credentials configured](#)<sup>156</sup>.
- An existing docker registry.
- Ability to download artifacts from the internet and then copy those onto your bootstrap machine.
- An [AWS Air-Gapped Machine Image](#) (see page 418)

#### 6.4.10.2.2.1 Configure AWS prerequisites

1. Follow the steps in [IAM Policy Configuration](#) (see page 151).
2. Export the AWS region where you want to deploy the cluster:

```
export AWS_REGION=us-west-2
```

3. Export the AWS profile with the credentials you want to use to create the Kubernetes cluster:

```
export AWS_PROFILE=<profile>
```

## 6.4.10.3 Air-gapped AMI


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.

<sup>154</sup> <https://docs.docker.com/get-docker/>

<sup>155</sup> <https://kubernetes.io/docs/tasks/tools/#kubectl>

<sup>156</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

### 6.4.10.3.1 Create an AMI using Konvoy Image Builder (KIB) for use in an air-gapped cluster

 The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to [AWS Air-gapped AMI](#)<sup>157158</sup> to take advantage of enhanced cluster operations, and to explore the [Install AWS Air-Gapped](#) (see page 202) topics for more options.

Using [KIB](#) (see page 411), you can build an AMI without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded](#) (see page 100) `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2/kib
```

2. Follow the instructions to [build an AMI](#) (see page 414) and set the override `--overrides overrides/offline.yaml` flag.

After you complete these steps, you can [seed your docker registry](#) (see page 204).

## 6.4.10.4 AWS Air-gapped Seed Docker Registry

### 6.4.10.4.1 Seed your Docker Registry

Before creating a Kubernetes cluster you must have the required images in a local docker registry. This registry must be accessible from both the bastion machine and the AWS EC2 instances that will be created for the Kubernetes cluster.

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2
```

2. Set an environment variable with your registry address.

157 <https://d2iq.atlassian.net/wiki/pages/resumedraft.action?draftId=144117789>

158 <https://d2iq.atlassian.net/wiki/pages/resumedraft.action?draftId=144117789>

```
export DOCKER_REGISTRY_ADDRESS=<registry-address>:<registry-port>
export DOCKER_REGISTRY_USERNAME=<username>
export DOCKER_REGISTRY_PASSWORD=<password>
```

3. Run the following command to load the air-gapped image bundle into your private Docker registry.

```
dkp push image-bundle --image-bundle ./container-images/konvoy-image-bundle-
v2.4.2.tar --to-registry $DOCKER_REGISTRY_ADDRESS --to-registry-username
$DOCKER_REGISTRY_USERNAME --to-registry-password $DOCKER_REGISTRY_PASSWORD
```

It may 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 Docker registry.

Then, begin by [Creating the Bootstrap cluster](#) (see page 205).

## 6.4.10.5 AWS Air-gapped Bootstrap

### 6.4.10.5.1 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_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2
```

2. Load the bootstrap docker image on your bastion machine.

```
docker load -i konvoy-bootstrap-image-v2.4.2.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 [Configure HTTP Proxy](#) (see page 494).

4. (Optional) Refresh the credentials used by the AWS provider at any time, using the command:

```
dkp update bootstrap credentials aws
```

Then, you can [create a cluster](#) (see page 206).

### 6.4.10.6 Create a New AWS Air-gapped Cluster

Create a new AWS Kubernetes cluster in an Air-gapped environment in an existing 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
```

See [Get Started with AWS](#) (see page 47) for information on naming your cluster.

**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](#)<sup>159</sup> for more naming information.

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 `ec2`, `elasticloadbalancing`, `secretsmanager` and `autoscaling` VPC endpoints to be already present.

<sup>159</sup> <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>


- `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<sup>160</sup>.
- `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 411).

**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
```

3. Configure your cluster to use an existing Docker registry as a mirror when attempting to pull images:

 **IMPORTANT:** The AMI must be created by the [konvoy-image-builder](#)<sup>161</sup> project in order to use the registry mirror feature.

```
export DOCKER_REGISTRY_URL=<https/http>://<registry-address>:<registry-port>
export DOCKER_REGISTRY_CA=<path to the CA on the bastion>
export DOCKER_REGISTRY_USERNAME=<username>
export DOCKER_REGISTRY_PASSWORD=<password>
```

- `DOCKER_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.
- `DOCKER_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.
- `DOCKER_REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.

<sup>160</sup> <https://github.com/kubernetes-sigs/cluster-api-provider-aws>

<sup>161</sup> <https://github.com/mesosphere/konvoy-image-builder>

- `DOCKER_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 940) 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=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=${DOCKER_REGISTRY_URL} \
--registry-mirror-cacert=${DOCKER_REGISTRY_CA} \
--registry-mirror-username=${DOCKER_REGISTRY_USERNAME} \
--registry-mirror-password=${DOCKER_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 [Configure HTTP Proxy](#) (see page 494).

5. (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.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

6. (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=${AWS_ADDITIONAL_SECURITY_GROUPS} \
--registry-mirror-url=${DOCKER_REGISTRY_URL} \
--registry-mirror-cacert=${DOCKER_REGISTRY_CA} \
--registry-mirror-username=${DOCKER_REGISTRY_USERNAME} \
--registry-mirror-password=${DOCKER_REGISTRY_PASSWORD} \
--control-plane-http-proxy="${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. Inspect the created cluster resources:

```
kubectl get clusters,kubeadmcontrolplanes,machinedeployments
```

8. Wait for the cluster control-plane to be ready:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=60m
```

Then, [Explore the New Cluster \(see page 209\)](#).

### 6.4.10.7 Explore New AWS Air-gapped Cluster

Explore the new Kubernetes cluster using the commands below.

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
```

When you're ready, [delete your cluster and clean up your environment \(see page 213\)](#).

## 6.4.10.8 Make the Air-gapped 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.

Before starting, ensure you create a workload cluster as described in [Create a New Cluster \(see page 175\)](#).

### 6.4.10.8.1 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 [Configure HTTP Proxy \(see page 494\)](#).

```
✓ 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](#)<sup>162</sup>.

```
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 |
|------|----------|--------|-------|---------|-------|
|      |          |        |       |         |       |

<sup>162</sup> <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```

Cluster/aws-example True
109s
├─ClusterInfrastructure - AWSCluster/aws-example True
112s
├─ControlPlane - KubeadmControlPlane/aws-example-control-plane True
109s
│ └─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

```

- Remove the bootstrap cluster, as the workload cluster is now self-managed:

```
dkp delete bootstrap
```

```
✓ Deleting bootstrap cluster
```

#### 6.4.10.8.2 Install and Log in to the DKP UI

You can now proceed to installing the [DKP UI \(see page 475\)](#) and applications. After installation, you will be able to [log in \(see page 522\)](#) to the DKP UI to explore it.

#### 6.4.10.8.3 DKP Install Configuration

You can [configure \(see page 475\)](#) 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 \(see page 206\)](#) installation with catalog applications.

- = Depending on your configuration, there are three 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:**

Essential:

- [Kommander in an Air-gapped Environment \(see page 506\)](#)
- [Kommander in Air-gapped with DKP Insights \(see page 507\)](#)

Enterprise:

- [Install Air-gapped Kommander with DKP Catalog Applications \(see page 509\)](#)
- [Install Air-gapped Kommander with DKP Insights and DKP Catalog Applications \(see page 511\)](#)

#### 6.4.10.8.4 Known Limitations

- = **NOTE:** Be aware of these limitations in the current release of Konvoy.

- 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 \(see page 157\)](#).
- 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.

#### 6.4.10.9 Delete AWS Air-gapped Cluster

##### 6.4.10.9.1 Delete the Kubernetes cluster and cleanup your environment

Follow these steps:

1. To delete a cluster, you would use `dcp 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 `dcp 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}
```

3. Delete the `kind` Kubernetes cluster:

```
dkp delete bootstrap
```

### 6.4.10.9.2 Next Step:

Once your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will install the [Kommander](#) (see page 475) component of DKP to see your dashboard and continue customization.

## 6.4.11 GPUs in an AWS environment

### 6.4.11.1 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](#) (see page 216)

Kommander also accesses [GPU](#) (see page 812) resources.



GPUs in an air-gapped on-premises environment recommend setting up [Pod Disruption Budget](#)<sup>163</sup> before [updating nodepools](#) (see page 473).

<sup>163</sup> <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

### 6.4.11.2 Install GPU Support for Supported Distributions on AWS

Using the [Konvoy Image Builder](#) (see page 427), you can build an image that has support to use Nvidia GPU hardware to support GPU workloads. DKP supported [Nvidia driver](#)<sup>164</sup> 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](#).<sup>165</sup>

```
gpu:
 type:
 - nvidia
```

Build your image using the following Konvoy image builder commands:

```
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, 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>
```

<sup>164</sup> <https://www.nvidia.com/Download/Find.aspx>

<sup>165</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

Additional helpful information can be found in the [Nvidia Device Plug-in](#)<sup>166</sup> for Kubernetes instructions and the [Installation Guide of Supported Platforms](#)<sup>167</sup>.

See also: [Nvidia documentation](#)<sup>168</sup>

### 6.4.11.3 Configure Konvoy Automatic GPU Node Labels

When using GPU node, 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](#)<sup>169</sup>:

```
< class > _ < vendor >
```

However, due to the 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 will deploy to the right nodes.

### 6.4.11.4 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](#)<sup>170</sup>) 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.

<sup>166</sup> <https://github.com/NVIDIA/k8s-device-plugin/blob/master/README.md>

<sup>167</sup> <https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/install-guide.html>

<sup>168</sup> [https://nvidia.custhelp.com/app/answers/detail/a\\_id/131/kw/driver%20installation%20docs/related/1](https://nvidia.custhelp.com/app/answers/detail/a_id/131/kw/driver%20installation%20docs/related/1)

<sup>169</sup> <https://kubernetes-sigs.github.io/node-feature-discovery/v0.7/get-started/features.html#pci>

<sup>170</sup> <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>




#### 6.4.11.4.1 Prerequisites:

- Deploy [Pod Disruption Budget](#)<sup>171</sup> (PDB)
- [Konvoy Image Builder](#) (see page 411) (KIB)

#### 6.4.11.4.2 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.

 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 411).

## 6.4.12 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

<sup>171</sup> <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

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](#)<sup>172</sup>. For more information on node pools, see these sections:

- [Create AWS Node Pools](#) (see page 218)
- [List AWS Node Pools](#) (see page 219)
- [Scale AWS Node Pools](#) (see page 220)
- [Delete AWS Node Pools](#) (see page 222)
- [AWS Cluster Autoscaler](#) (see page 223)

### 6.4.12.1 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.

#### 6.4.12.1.1 Prepare the environment

1. Set the environment variable to the name you assigned this cluster.

```
export CLUSTER_NAME=aws-example
```

See [Get Started with AWS](#) (see page 47) 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 189), configure `kubectl` to use the kubeconfig for the cluster.

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

3. Define your node pool name.

```
export NODEPOOL_NAME=example
```

#### 6.4.12.1.2 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.

---

<sup>172</sup> <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

- 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, AMLs, and other properties.

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.

## 6.4.12.2 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.

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     | v1.24.6 |
| aws-example-md-0   | 4       | 4     | v1.24.6 |

### 6.4.12.2.1 Attached Clusters

Before running the command to list the attached clusters, ensure that you know the following:

- KUBECONFIG for the management cluster
- 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<br>VERSION | DESIRED | READY | KUBERNETES |
|---------------------|---------|-------|------------|
| aws-attached-md-0   | 4       | 4     | v1.24.6    |

### 6.4.12.3 Scale AWS Node Pools

#### Scale node pools in a cluster

While you can run [Cluster Autoscaler](#) (see page 223), 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.

#### 6.4.12.3.1 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=${CLUSTER_NAME} --kubeconfig=${CLUSTER_NAME}.conf
```

Your output should be similar to this example, with the number of DESIRED and READY replicas increased to 5:

| NODEPOOL<br>KUBERNETES VERSION | DESIRED | READY |         |
|--------------------------------|---------|-------|---------|
| example                        | 5       | 5     | v1.24.6 |
| aws-example-md-0               | 4       | 4     | v1.24.6 |

#### 6.4.12.3.2 Scaling Down Node Pools

To scale down a node pool, run:

```
dkp scale nodepools ${NODEPOOL_NAME} --replicas=4 --cluster-name=${CLUSTER_NAME}
```

✓ Scaling node pool example to 4 replicas

After a few minutes, you can list the node pools using this command:

```
dkp get 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     | v1.24.6 |
| aws-example-md-0   | 4       | 4     | v1.24.6 |

In a default cluster, the nodes to delete are selected at random. This behavior is controlled by [CAPI's delete policy](#)<sup>173</sup>. 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 nodepools ${NODEPOOL_NAME} --replicas=3 --nodes-to-delete=<> --cluster-name=${CLUSTER_NAME}
```

✓ Scaling node pool example to 3 replicas

### 6.4.12.3.3 Scaling Node Pools When Using Cluster Autoscaler

If you [configured the cluster autoscaler](#) (see page 223) 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:

<sup>173</sup> [https://github.com/kubernetes-sigs/cluster-api/blob/v0.4.0/api/v1alpha4/machineset\\_types.go#L85-L105](https://github.com/kubernetes-sigs/cluster-api/blob/v0.4.0/api/v1alpha4/machineset_types.go#L85-L105)

```
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.

## 6.4.12.4 Delete AWS 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 aws-example" not found
```

## 6.4.12.5 AWS Cluster Autoscaler

### Configure autoscaler for node pools

[Cluster Autoscaler](#)<sup>174</sup> 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](#)<sup>175</sup>, 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](#)<sup>176</sup>.

For more information about how Cluster Autoscaler works, see these documents:

- [What is Cluster Autoscaler](#)<sup>177</sup>
- [How does scale-up work](#)<sup>178</sup>
- [How does scale-down work](#)<sup>179</sup>
- [CAPI Provider for Cluster Autoscaler](#)<sup>180</sup>

### 6.4.12.5.1 Cluster Autoscaler Prerequisites

Before you begin, you must have:

- A [Bootstrap Cluster Lifecycle](#) (see page 173).
- [Created a new Kubernetes Cluster](#) (see page 175).
- A [Self-Managed Cluster](#) (see page 189).

### 6.4.12.5.2 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.

<sup>174</sup> <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

<sup>175</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-fast-is-hpa-when-combined-with-ca>

<sup>176</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-are-the-parameters-to-ca>

<sup>177</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-is-cluster-autoscaler>

<sup>178</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-up-work>

<sup>179</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-down-work>

<sup>180</sup> <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

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
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.

## 6.4.13 Creating DKP Clusters on AWS

### A guide for creating DKP clusters on AWS

#### 6.4.13.1 Before you begin

- Configure an [AWS Infrastructure \(see page 144\)](#), or add credentials using [AWS role credentials \(see page 547\)](#).

#### 6.4.13.2 Simplified Cluster Creation on AWS

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 \(see page 615\)](#). Changing a cluster label may add or remove the cluster from projects.
4. Select the pre-configured [AWS Infrastructure \(see page 144\)](#), or [AWS role credentials \(see page 547\)](#) 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 Lookup:** You can also specify the base OS, lookup format, and owner ID of the AMI Image, or the AMI ID as part of each pool. Selecting 'Show Advanced' within the Configure Node Pools section will show these options.
    - **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.
    - **AMI ID:** AMI ID to use for all nodes.
  - **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](#)<sup>181</sup> 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 once it is fully-provisioned.

## 6.5 Azure Infrastructure

This section describes how you create a DKP cluster in Azure.

- [Azure Prerequisites](#) (see page 226)
- [Azure using Konvoy Image Builder](#) (see page 229)
- [Azure Bootstrap](#) (see page 232)
- [Create a New Azure Cluster](#) (see page 235)
- [Explore new Azure Cluster](#) (see page 242)
- [Azure Make new Cluster Self-Managed](#) (see page 246)
- [Azure Certificate Renewal](#) (see page 248)
- [Azure Replace a Node](#) (see page 251)
- [Azure Delete Cluster](#) (see page 254)
- [Create a new Azure Cluster via UI](#) (see page 258)

### 6.5.1 Azure Prerequisites

**Prepare your machine and environment to run DKP**

#### 6.5.1.1 DKP Prerequisites

Before you begin using DKP you must have:

- An x86\_64-based Linux or macOS machine.
- The `dkp` binary for Linux, or macOS.

---

<sup>181</sup> [https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using\\_Tags.html](https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/Using_Tags.html)

- [Docker](#)<sup>182</sup> version 18.09.2 or later installed.
- [kubect](#)<sup>183</sup> for interacting with the running cluster.
- [Azure CLI](#)<sup>184</sup>.
- A valid Azure account with [credentials configured](#)<sup>185</sup>.
- Create a custom Azure image using [KIB](#) (see page 411).



On macOS, Docker runs in a virtual machine. Configure this virtual machine with at least 8GB of memory.

#### 6.5.1.1.1 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.

#### 6.5.1.1.2 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 93) with 3 control plane nodes, and 4 worker nodes which match the requirements above.

<sup>182</sup> <https://docs.docker.com/get-docker/>

<sup>183</sup> <https://kubernetes.io/docs/tasks/tools/#kubectl>

<sup>184</sup> <https://docs.microsoft.com/en-us/cli/azure/install-azure-cli>

<sup>185</sup> <https://github.com/kubernetes-sigs/cluster-api-provider-azure/blob/master/docs/book/src/topics/getting-started.md#prerequisites>

### 6.5.1.2 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 232) section.

## 6.5.2 Azure using Konvoy Image Builder

This procedure describes how to use the [KIB](#) (see page 420) to create a [Cluster API](#)<sup>186</sup> 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 [Konvoy Image Builder](#) (see page 411) 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 226).

### 6.5.2.1 Prerequisites


Before you begin, you must:

- Check the [DKP Supported Kubernetes Version](#) (see page 1056) and download the [KIB](#) (see page 411) bundle (prefixed with `konvoy-image-bundle`) for your OS. Do not use the release prefixed with `konvoy-image-builder`.
- Create a working `Docker` setup.


<sup>186</sup> <https://cluster-api.sigs.k8s.io/>

### 6.5.2.2 Extract the 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 run, `konvoy-image` binds mounts the current working directory ( `$(PWD)` ) into the container to be used.

### 6.5.2.3 Configure Azure Prerequisites

 If you have already followed the [Azure Prerequisites](#) (see page 226) 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 226) 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@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) --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 451) to configure specific attributes of your Azure image.

#### 6.5.2.4 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 `./manifest.json`. You should then specify this image id when creating the cluster.

### 6.5.2.5 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 235) for specific consumption of image commands.

The SKU and Image Name will default to the values found in the image YAML.

## 6.5.3 Azure Bootstrap

### 6.5.3.1 Prepare to deploy Kubernetes clusters

To create Kubernetes clusters, Konvoy uses [Cluster API](#)<sup>187</sup> (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](#)<sup>188</sup>) tool.

<sup>187</sup> <https://cluster-api.sigs.k8s.io/>

<sup>188</sup> <https://github.com/kubernetes-sigs/kind>



### 6.5.3.2 Prerequisites

Before you begin, you must:

- Complete the steps in [Prerequisites](#) (see page 226).
- Ensure the `dkp` binary can be found in your `$PATH`.

### 6.5.3.3 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 or CLI flags. They are copied into the bootstrap cluster.
2. 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 [Configure HTTP Proxy](#) (see page 494).

The output resembles:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

Konvoy creates a bootstrap cluster using [KIND](#)<sup>189</sup> as a library. Konvoy then deploys the following [Cluster API](#)<sup>190</sup> providers on the cluster:

- [Core Provider](#)<sup>191</sup>
- [Azure Infrastructure Provider](#)<sup>192</sup>
- [kubeadm Bootstrap Provider](#)<sup>193</sup>
- [kubeadm ControlPlane Provider](#)<sup>194</sup>

Konvoy waits until the controller-manager and webhook deployments of these providers are ready. List these deployments using this command:

<sup>189</sup> <https://github.com/kubernetes-sigs/kind>

<sup>190</sup> <https://cluster-api.sigs.k8s.io/>

<sup>191</sup> <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

<sup>192</sup> <https://github.com/kubernetes-sigs/cluster-api-provider-azure>

<sup>193</sup> <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

<sup>194</sup> <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>

```
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         | 69s | capa-controller-manager                       |
| capi-kubeadm-bootstrap-system     | 1/1   | 1          | 1         | 71s | capi-kubeadm-bootstrap-controller-manager     |
| capi-kubeadm-control-plane-system | 1/1   | 1          | 1         | 70s | capi-kubeadm-control-plane-controller-manager |
| capi-system                       | 1/1   | 1          | 1         | 73s | capi-controller-manager                       |
| capp-system                       | 1/1   | 1          | 1         | 66s | capp-controller-manager                       |
| capv-system                       | 1/1   | 1          | 1         | 65s | capv-controller-manager                       |
| capz-system                       | 1/1   | 1          | 1         | 67s | capz-controller-manager                       |
| cert-manager                      | 1/1   | 1          | 1         | 16m | cert-manager                                  |
| cert-manager                      | 1/1   | 1          | 1         | 16m | cert-manager-cainjector                       |
| cert-manager                      | 1/1   | 1          | 1         | 16m | cert-manager-webhook                          |

### 6.5.3.4 (Optional) Create identity secret for Azure

If your bootstrap cluster resides on a Virtual machine inside Azure, create an identity secret that uses the capz-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}"
```

When complete, move on to the [Create a New Cluster](#) (see page 235) section.

## 6.5.4 Create a New Azure Cluster

### 6.5.4.1 Prerequisites

- Before you begin, make sure you have created a [Bootstrap](#) (see page 232) cluster.

### 6.5.4.2 Name your cluster

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

```
export CLUSTER_NAME=azure-example
```

### 6.5.4.3 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 Azure Cluster](#) (see page 0) section.

**Important to remember** related to the options below: `--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 420) 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
```

- Option 3: 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 478) in the Kommander component.

#### 6.5.4.4 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.

1. Generate the Kubernetes cluster objects. The following example shows a common configuration. See [dkp create cluster azure](#) (see page 943) 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
```

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,kubernetes.default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.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
```

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 [Configure HTTP Proxy \(see page 494\)](#).

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](#)<sup>195</sup> 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 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.

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 *AzureMachineTemplate* 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 *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](#)<sup>196</sup> 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 466).
6. Create the cluster from the objects.

```
kubectl create -f ${CLUSTER_NAME}.yaml
```

```
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
```

<sup>195</sup> <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

<sup>196</sup> <https://cluster-api.sigs.k8s.io/user/concepts.html>

```
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
```

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

```

### 6.5.4.5 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 242).

## 6.5.5 Explore new Azure Cluster

### 6.5.5.1 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 235).

### 6.5.5.2 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 | Ready  | control-plane,master | 6m18s |      |
| v1.24.6                           |        |                      |       |      |
| azure-example-control-plane-l4bv8 | Ready  | control-plane,master | 14m   |      |
| v1.24.6                           |        |                      |       |      |
| azure-example-control-plane-n4g4l | Ready  | control-plane,master | 18m   |      |
| v1.24.6                           |        |                      |       |      |
| azure-example-md-0-mpctb          | Ready  | <none>               | 15m   | v1.2 |
| 4.6                               |        |                      |       |      |
| azure-example-md-0-qglp9          | Ready  | <none>               | 15m   | v1.2 |
| 4.6                               |        |                      |       |      |
| azure-example-md-0-sgrd6          | Ready  | <none>               | 16m   | v1.2 |
| 4.6                               |        |                      |       |      |
| azure-example-md-0-wzbkl          | Ready  | <none>               | 16m   | v1.2 |
| 4.6                               |        |                      |       |      |



**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     | READY | STATUS  | RESTARTS | AGE   | NAME                                     |
|---------------|-------|---------|----------|-------|------------------------------------------|
| calico-system | 1/1   | Running | 0        | 19m   | calico-kube-controllers-57fbd7bd59-v4tss |
| calico-system | 1/1   | Running | 0        | 17m   | calico-node-59llv                        |
| calico-system | 1/1   | Running | 0        | 16m   | calico-node-7t7wj                        |
| calico-system | 1/1   | Running | 0        | 17m   | calico-node-pf8q8                        |
| calico-system | 1/1   | Running | 0        | 8m17s | calico-node-sh2b7                        |
| calico-system | 1/1   | Running | 0        | 19m   | calico-node-tmxl5                        |

```

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-control-plane-system
manager-5df55579c4-pc8x9 1/1 Running 1 (11m ago) 15m
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-7ffn1
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-82zdl
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-lzswl
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

```

## 6.5.6 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.

Before starting, ensure you create a workload cluster as described in [Create a New Cluster](#) (see page 235).

### 6.5.6.1 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 [Configure HTTP Proxy](#) (see page 494).

The output resembles:

```
✓ 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](#)<sup>197</sup>.

```
dkp move capi-resources --to-kubeconfig ${CLUSTER_NAME}.conf
```

<sup>197</sup> <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                      |        | 65s   |         | True  |          |

```

| └─ 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-jxw16 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

## 6.5.6.2 Known Limitations



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

- 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.

## 6.5.7 Azure Certificate Renewal

### 6.5.7.1 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.

### 6.5.7.2 Requirements

This feature requires Python 3.5 or greater to be installed on all control plane hosts.

### 6.5.7.3 Prerequisites

Prerequisite:

- Complete the [bootstrap Cluster Lifecycle](#) (see page 232) topic.

#### 6.5.7.3.1 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.

#### 6.5.7.3.2 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.

### 6.5.7.3.3 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
```

## 6.5.8 Azure Replace a Node

### 6.5.8.1 Replace a worker node

### 6.5.8.2 Prerequisites

Before you begin, you must:

- [Create a workload cluster](#) (see page 235).
- [Make the new cluster self-managed](#) (see page 246).

### 6.5.8.3 Replace a worker node

In certain situations, you may want to delete a worker node and have [Cluster API](#)<sup>198</sup> 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 | v1.24.6 |
| azure-example-control-plane-d4fdf | Ready  | control-plane,master | 31m | v1.24.6 |
| azure-example-control-plane-qrvm9 | Ready  | control-plane,master | 33m | v1.24.6 |
| azure-example-md-0-4w7gq          | Ready  | <none>               | 33m | v1.24.6 |
| azure-example-md-0-6gb9k          | Ready  | <none>               | 33m | v1.24.6 |
| azure-example-md-0-p2n8c          | Ready  | <none>               | 11m | v1.24.6 |
| azure-example-md-0-s5zbh          | Ready  | <none>               | 33m | v1.24.6 |

<sup>198</sup> <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               | PHASE     | AGE   | CLUSTER       | VERSION | REPLICAS | READY | UPDATED | UNAVAILABLE |
|--------------------|-----------|-------|---------------|---------|----------|-------|---------|-------------|
| azure-example-md-0 | ScalingUp | 7m30s | azure-example | v1.24.6 | 4        | 3     | 4       | 1           |
| long-running-md-0  | Running   | 7m28s | long-running  | v1.24.6 | 4        | 4     | 4       | 0           |

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   |      |
| v1.24.6                           |        |                      |       |      |
| azure-example-control-plane-qrvm9 | Ready  | control-plane,master | 45m   |      |
| v1.24.6                           |        |                      |       |      |
| azure-example-control-plane-tz56m | Ready  | control-plane,master | 8m22s |      |
| v1.24.6                           |        |                      |       |      |
| azure-example-md-0-4w7gq          | Ready  | <none>               | 45m   | v1.2 |
| 4.6                               |        |                      |       |      |
| azure-example-md-0-6gb9k          | Ready  | <none>               | 45m   | v1.2 |
| 4.6                               |        |                      |       |      |
| azure-example-md-0-p2n8c          | Ready  | <none>               | 22m   | v1.2 |
| 4.6                               |        |                      |       |      |
| azure-example-md-0-s5zbh          | Ready  | <none>               | 45m   | v1.2 |
| 4.6                               |        |                      |       |      |

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.

## 6.5.9 Azure Delete Cluster

### 6.5.9.1 Delete the Kubernetes cluster and clean up your environment

### 6.5.9.2 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 246](#)), proceed to [Delete the workload cluster](#) (see [page 256](#)) 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](#)<sup>199</sup>.

```
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                                                            | REASON | SINCE | MESSAGE | READY | SEVERITY |
|-----------------------------------------------------------------|--------|-------|---------|-------|----------|
| 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                                                             |        |       |         |       |          |

<sup>199</sup> <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.

Use `dkp` with the bootstrap cluster to delete the workload cluster.

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
```

**NOTE:** 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.

### 6.5.9.3 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:

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.

#### 6.5.9.4 Delete the Bootstrap Cluster

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

```
✓ Deleting bootstrap cluster
```

#### 6.5.9.5 Next Step:

Once your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will install the [Kommander](#) (see page 475) component of DKP to see your dashboard and continue customization.

#### 6.5.9.6 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.

## 6.5.10 Create a new Azure Cluster via UI

### Enterprise

DKP UI allows you to provision an Azure cluster from your browser.

### 6.5.10.1 Prerequisites

#### 6.5.10.1.1 Create an Azure Infrastructure Provider

Before you provision a cluster via the UI, first create an Azure infrastructure provider to hold your Azure credentials:

1. Login 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 580) 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` used in Step 1 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**.

### 6.5.10.2 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 [default Kubernetes version \(see page 1056\)](#) 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. Validate your inputs, and then select **Create**.



It can take up to 15 minutes for your cluster to appear in the **Provisioned** status.

You're 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.

## 6.6 EKS Infrastructure

### Enterprise

When installing DKP on your EKS infrastructure, you can choose from multiple configuration types. The steps for creating and accessing your cluster are listed below:

- [EKS Introduction \(see page 260\)](#)
- [EKS Prerequisites \(see page 261\)](#)
- [EKS Cluster IAM Policies and Roles \(see page 265\)](#)
- [Create an EKS Cluster from the CLI \(see page 270\)](#)
- [Grant Cluster Access \(see page 277\)](#)
- [Explore EKS Cluster \(see page 278\)](#)
- [Manage EKS Nodepools \(see page 280\)](#)
- [Delete EKS Cluster from CLI \(see page 283\)](#)
- [Create an EKS Cluster from the UI \(see page 284\)](#)

- [Delete EKS Cluster from UI \(see page 286\)](#)

## 6.6.1 EKS Introduction

### Enterprise

DKP brings value to EKS customers by providing all components needed for a production-ready Kubernetes environment. DKP 2.4 provides the capability to provision EKS clusters using the DKP UI, In addition to above, DKP 2.4 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 provides 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

## 6.6.2 EKS Prerequisites

Enterprise

### 6.6.2.1 DKP Prerequisites

Before you begin using DKP, you must have:RIP

- An x86\_64-based Linux or macOS machine.
- The `dkp` binary for Linux, or macOS.
- A [Self-managed AWS Cluster](#) (see page 189)
- [kubecti](#)<sup>200</sup> for interacting with the running cluster.

Ensure that the `KUBECONFIG` environment variable is set to the self-managed cluster by running `export KUBECONFIG={SELF_MANAGED_AWS_CLUSTER}.conf`.

### 6.6.2.2 AWS prerequisites

Before you begin using DKP with AWS, you must have:

- You have a valid AWS account with [credentials configured](#)<sup>201</sup> that can manage CloudFormation Stacks, IAM Policies, and IAM Roles.

<sup>200</sup> <https://kubernetes.io/docs/tasks/tools/install-kubecti/>

<sup>201</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

- You will need to have the [AWS CLI utility installed](https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html)<sup>202</sup>.
- Install [aws-iam-authenticator](https://docs.aws.amazon.com/eks/latest/userguide/install-aws-iam-authenticator.html)<sup>203</sup>. This binary is used to access your cluster using kubectl.

### 6.6.2.2.1 Minimal User Permissions for Creating EKS Clusters:

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 151) section. Consult the [Leveraging the Role](#) (see page 0) 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:
 - >-

```

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

203 <https://docs.aws.amazon.com/eks/latest/userguide/install-aws-iam-authenticator.html>

```


 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 and run the following command after replacing `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
```



### 6.6.2.3 Access Cluster

- Use previously installed [aws-iam-authenticator](#)<sup>204</sup> 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](#).<sup>205</sup>

## 6.6.3 EKS Cluster IAM Policies and Roles

### Enterprise

This section guides a DKP user in creating IAM Policies and Instance Profiles that determines type of 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 code for an additional role for EKS not needed in other AWS environments.

### 6.6.3.1 Prerequisites:

- The user you delegate from your role must have a minimum set of permissions, see [User Roles and Instance Profiles](#) (see page 151) page under AWS for details.
- Create the [Cluster IAM Policies](#) (see page 157) in your AWS account
- You will need to have the [AWS CLI utility installed](#)<sup>206</sup>.

### 6.6.3.2 EKS IAM Artifacts

#### 6.6.3.2.1 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

<sup>204</sup> <https://docs.aws.amazon.com/eks/latest/userguide/install-aws-iam-authenticator.html>

<sup>205</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

<sup>206</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html>

- `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`

### 6.6.3.2.2 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 157) in AWS.

Below is a [CloudFormation stack](#)<sup>207</sup> 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:
 'iam:AWSServiceName': eks.amazonaws.com
```

<sup>207</sup> <https://docs.aws.amazon.com/AWSCloudFormation/latest/UserGuide/Welcome.html>

```

 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:CreateTags"

```

```

 - "ec2:DeleteNetworkInterface"
 - "ec2:DescribeInstances"
 - "ec2:DescribeTags"
 - "ec2:DescribeNetworkInterfaces"
 - "ec2:DescribeInstanceTypes"
 - "ec2:DetachNetworkInterface"
 - "ec2:ModifyNetworkInterfaceAttribute"
 - "ec2:UnassignPrivateIpAddresses"
 Effect: Allow
 Resource:
 - '*'
- 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 run 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
```

### 6.6.3.2.3 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
```

## 6.6.4 Create an EKS Cluster from the CLI

### Enterprise

- Ensure that the `KUBECONFIG` environment variable is set to the self-managed cluster by running `export KUBECONFIG={SELF_MANAGED_AWS_CLUSTER}.conf`

### 6.6.4.1 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.

Ensure you have the proper [role permissions](#) (see page 265), then follow these steps:

1. Set the environment variable to the name you assigned this cluster.

```
export CLUSTER_NAME=eks-example
```

See [EKS Quick Start](#) (see page 0) for information on naming your cluster.

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}
```

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 [Configure HTTP Proxy \(see page 494\)](#).

The output resembles:

```
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](#)<sup>208</sup> 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**

<sup>208</sup> <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](#)<sup>209</sup> 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. Once 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
```

<sup>209</sup> <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-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
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-

```

## 6.6.4.2 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.


The next step would be to [attach this cluster](#) (see page 717) to the management cluster.

## 6.6.5 Grant Cluster Access

Enterprise

### 6.6.5.1 How to Grant 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](#)<sup>210</sup> (RBAC).

 Configuration in the control plane as discussed in the [EKS Cluster IAM Policies and Roles](#) (see [page 265](#)) page.

If the EKS cluster was created as a cluster against 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 EKS workload cluster. Follow the steps below:

1. Execute the following command against the self-managed AWS cluster, you are creating the workload EKS cluster against. 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](#)<sup>211</sup>, 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/nodes.cluster-api-provider-aws.sigs.k8s.io
```

<sup>210</sup> <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

<sup>211</sup> <https://docs.aws.amazon.com/general/latest/gr/aws-arns-and-namespaces.html>

```

username: system:node:{{EC2PrivateDNSName}}
- groups:
- system:masters
rolearn: arn:aws:iam::111122223333:role/Mesosphere-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](#)<sup>212</sup>.

## 6.6.6 Explore EKS Cluster

### Enterprise

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 270](#)).

### 6.6.6.1 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
```

| NAME                                                      | STATUS | ROLES  | AGE | VERSION  |
|-----------------------------------------------------------|--------|--------|-----|----------|
| ip-10-0-122-211.us-west-2.compute.internal<br>eks-ba74326 | Ready  | <none> | 35m | v1.23.9- |
| ip-10-0-127-74.us-west-2.compute.internal<br>eks-ba74326  | Ready  | <none> | 35m | v1.23.9- |

<sup>212</sup> <https://docs.aws.amazon.com/eks/latest/userguide/add-user-role.html>

```
ip-10-0-71-155.us-west-2.compute.internal Ready <none> 35m v1.23.9-eks-ba74326
ip-10-0-93-47.us-west-2.compute.internal Ready <none> 35m v1.23.9-eks-ba74326
```

**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-nvccck                   | 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   | Running  | 0        | aws-node-5f2kx                           | 35m | 1/1   |
| kube-system   | Running  | 0        | aws-node-6lzm7                           | 35m | 1/1   |
| kube-system   | Running  | 0        | aws-node-pz8c6                           | 35m | 1/1   |
| kube-system   | Init:0/1 | 0        | cluster-autoscaler-789d86b489-sz9x2      | 36m | 0/1   |

|                        |                                               |     |
|------------------------|-----------------------------------------------|-----|
| 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                                           |     |

## 6.6.7 Manage EKS Nodepools

Enterprise



- 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](#)<sup>213</sup>.

### 6.6.7.1 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](#)<sup>214</sup> 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} \
 --replicas=3
```

```
machinedeployment.cluster.x-k8s.io/example created
awscloudformation.cluster.x-k8s.io/example created
eksconfigtemplate.bootstrap.cluster.x-k8s.io/example created
✓ Creating default/example nodepool resources
```

<sup>213</sup> <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

<sup>214</sup> [https://aws.amazon.com/about-aws/global-infrastructure/regions\\_az/](https://aws.amazon.com/about-aws/global-infrastructure/regions_az/)

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.

### 6.6.7.2 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     | v1.23.6 |
| eks-example-md-0   | 4       | 4     | v1.23.6 |

### 6.6.7.3 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
```

## 6.6.8 Delete EKS Cluster from CLI

### Enterprise

- Ensure that the `KUBECONFIG` environment variable is set to the self-managed cluster by running `export KUBECONFIG={SELF_MANAGED_AWS_CLUSTER}.conf`

### 6.6.8.1 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
```

### 6.6.8.2 Next Step:

Once your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will install the [Kommander](#) (see page 475) component of DKP to see your dashboard and continue customization.

### 6.6.8.3 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.

## 6.6.9 Create an EKS Cluster from the UI

Enterprise

DKP UI allows you to quickly and easily provision a Cluster from your browser.

### 6.6.9.1 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 151).

```
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.

### 6.6.9.2 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](#) (see page 1056) 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](#)<sup>215</sup> for more information from the AWS site.

### 6.6.9.3 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.

### 6.6.9.4 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

---

<sup>215</sup> [https://docs.aws.amazon.com/IAM/latest/UserGuide/reference\\_identifiers.html#identifiers-arns](https://docs.aws.amazon.com/IAM/latest/UserGuide/reference_identifiers.html#identifiers-arns)

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
 roleName: arn:aws:iam::MYAWSACCOUNTID:role/nodes.cluster-api-provider-
aws.sigs.k8s.io
 username: system:node:{{EC2PrivateDNSName}}
 - groups:
 - system:masters
 roleName: arn:aws:iam::MYAWSACCOUNTID:role/PowerUser
 username: admin
kind: ConfigMap
```

For more information, consult the [Enabling IAM user and role access](#)<sup>216</sup> to you cluster guide and the [Kubernetes RBAC guide](#)<sup>217</sup>.

4. From your management cluster run the following command to fetch a kubeconfig that uses IAM based permissions by running:

```
dkp get kubeconfig -c ${EKS_CLUSTER_NAME} -n ${KOMMANDER_WORKSPACE_NAMESPACE}
>> ${EKS_CLUSTER_NAME}.conf
```

## 6.6.10 Delete EKS Cluster from UI

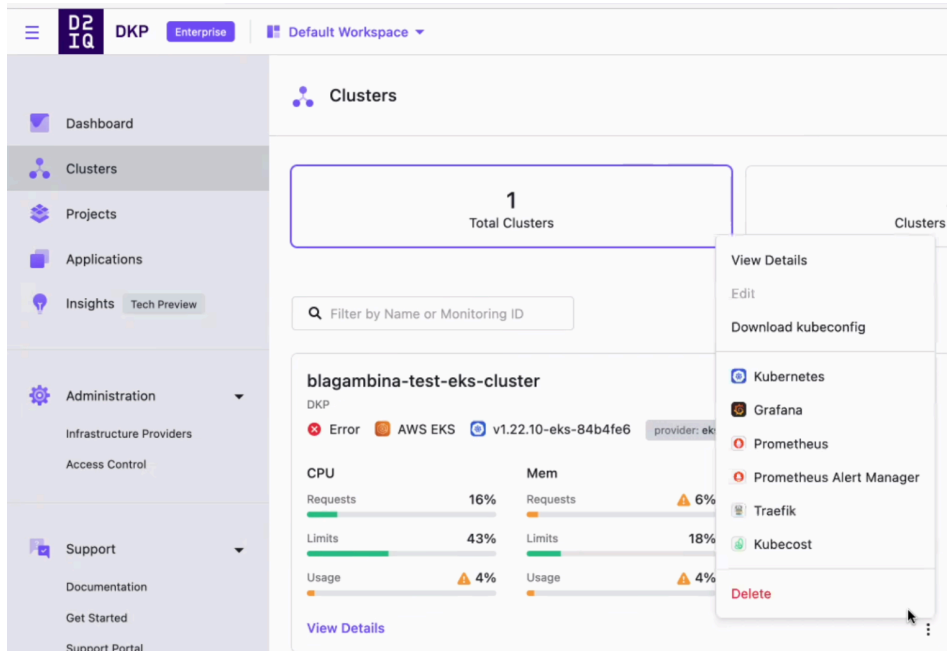
Enterprise

<sup>216</sup> <https://docs.aws.amazon.com/eks/latest/userguide/add-user-role.html>

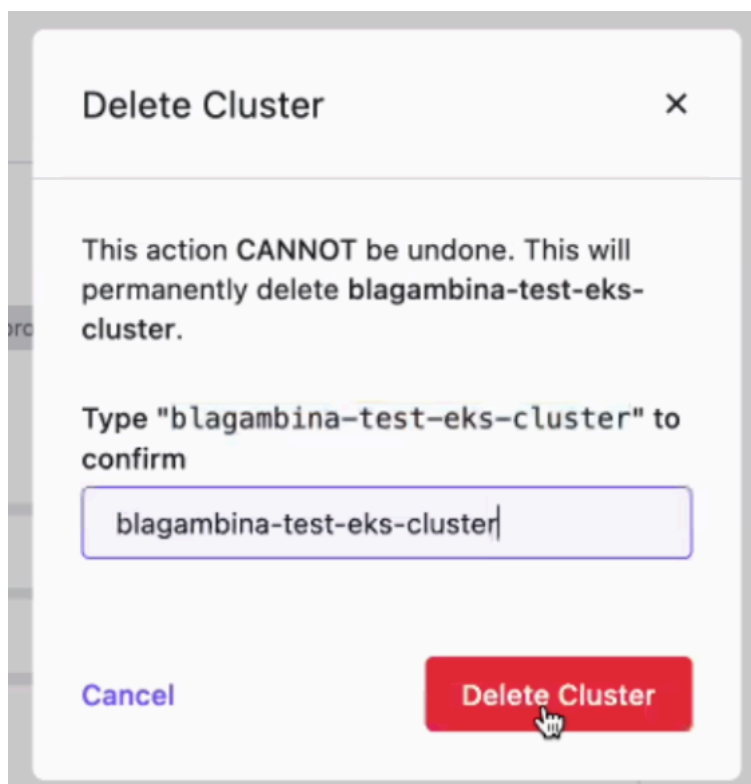
<sup>217</sup> <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

In order to delete a cluster in the UI, you must first have [created a cluster](#) (see page 284) 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 729) instructions.

## 6.7 Pre-provisioned Infrastructure

### 6.7.1 Create a Kubernetes cluster on pre-provisioned 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 \(CNI\)](#)<sup>218</sup> and a [Local Persistence Volume Static Provisioner](#)<sup>219</sup>, 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 570), or deploy your own.

To get started, see these topics.

- [Pre-provisioned Prerequisites](#) (see page 289)
- [Pre-provisioned Prerequisites Air-gapped](#) (see page 291)
- [Pre-provisioned Bootstrap](#) (see page 296)
- [Pre-provisioned Create Necessary Secrets and Overrides](#) (see page 296)

<sup>218</sup> <https://docs.projectcalico.org/>

<sup>219</sup> <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>



- [Pre-provisioned Define Infrastructure](#) (see page 298)
- [Pre-provisioned Define Control Plane Endpoint](#) (see page 300)
- [Pre-provisioned Create new Cluster](#) (see page 302)
- [Pre-provisioned Make Cluster Self-managed](#) (see page 313)
- [Pre-provisioned Configure MetalLB](#) (see page 316)
- [Pre-provisioned Create and Delete Node Pools](#) (see page 317)
- [Pre-provisioned Add Nodes to Existing Node Pools](#) (see page 319)
- [GPU Nodepools in a Pre-provisioned Environment](#) (see page 321)
- [Pre-provisioned Delete Cluster](#) (see page 323)

## 6.7.2 Pre-provisioned Prerequisites

### 6.7.2.1 Fulfill the prerequisites for using 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.
- `kubectl`<sup>220</sup> 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.
- Resource [requirements](#) (see page 116)
- [Pre-provisioned Override Files](#) (see page 459)

When air-gapped, you must follow the steps described in the [Air-gapped Prerequisites](#) (see page 291) page, otherwise you may [begin creating the bootstrap cluster](#) (see page 296).

### 6.7.2.2 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 91).
- `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.

---

<sup>220</sup> <https://kubernetes.io/docs/tasks/tools/#kubectl>

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

### 6.7.2.3 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](https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings)<sup>221</sup> as specified in the [requirements table](https://docs.d2iq.com/dkp/2.4/rook-ceph-configuration#id-(2.4)RookCephConfiguration-ResourceRequirementstrue)<sup>222</sup>.
- Multiple ports open, as described in [DKP Ports \(see page 91\)](#).
- `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.

<sup>221</sup> <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

<sup>222</sup> [https://docs.d2iq.com/dkp/2.4/rook-ceph-configuration#id-\(2.4\)RookCephConfiguration-ResourceRequirementstrue](https://docs.d2iq.com/dkp/2.4/rook-ceph-configuration#id-(2.4)RookCephConfiguration-ResourceRequirementstrue)

## 6.7.3 Pre-provisioned Prerequisites Air-gapped

### 6.7.3.1 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 air-gapped. In DKP 2.4.0, there is a new complete DKP air-gapped bundle available to [download](#) (see [page 100](#)) which contains all the DKP components needed for air-gapped installation. (i.e. `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`)

### 6.7.3.2 Air-Gapped Registry Prerequisites

#### 6.7.3.2.1 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.

#### 6.7.3.2.2 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](#)<sup>223</sup>

#### 6.7.3.2.3 Harbor Registry

Any newer version than **Harbor Registry v2.1.1-5f52168e** will support OCI images.

### 6.7.3.3 Load the bootstrap image

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz` from the [download site](#) (see [page 100](#)) mentioned above, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2
```

2. Load the bootstrap Docker image on your bastion machine:

```
docker load -i konvoy-bootstrap-image-v2.4.2.tar
```

---

<sup>223</sup> <https://issues.sonatype.org/browse/NEXUS-21087>

### 6.7.3.4 Copy air-gapped artifacts onto cluster hosts

Using the [Konvoy Image Builder](#) (see page 411), you can copy the required artifacts onto your cluster hosts.

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2/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-1.24.6-d2iq.1.tar kubernetes-images-1.24.6-d2iq.1-
fips.tar
```

- b. Verify the artifacts for your OS exist in the `artifacts/` directory and export the appropriate variables:

```
$ ls artifacts/
1.24.6-centos_7_x86_64_fips.tar.gz 1.24.6-redhat_8_x86_64_fips.tar.gz
containerd-1.4.13-d2iq.1-rhel-7.9-x86_64_fips.tar.gz containerd-1.4.13-
d2iq.1-rhel-8.4-x86_64_fips.tar.gz images
1.24.6-centos_7_x86_64.tar.gz 1.24.6-redhat_8_x86_64.tar.gz
containerd-1.4.13-d2iq.1-rhel-7.9-x86_64.tar.gz containerd-1.4.13-
d2iq.1-rhel-8.4-x86_64.tar.gz NVIDIA-Linux-x86_64-470.82.01.run
1.24.6-redhat_7_x86_64_fips.tar.gz containerd-1.4.13-d2iq.1-centos-7.9-
x86_64_fips.tar.gz containerd-1.4.13-d2iq.1-rhel-8.2-x86_64_fips.tar.gz
containerd-1.4.13-d2iq.1-rhel-8.6-x86_64_fips.tar.gz pip-
packages.tar.gz
1.24.6-redhat_7_x86_64.tar.gz containerd-1.4.13-d2iq.1-centos-7.9-
x86_64.tar.gz containerd-1.4.13-d2iq.1-rhel-8.2-x86_64.tar.gz
containerd-1.4.13-d2iq.1-rhel-8.6-x86_64.tar.gz
```

- c. For example, for RHEL 8.4 you would set:

```
export OS_PACKAGES_BUNDLE=1.24.6-redhat_8_x86_64.tar.gz
export CONTAINERD_BUNDLE=containerd-1.4.13-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. 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
```

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 451) 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](#)<sup>224</sup> or see these pages in the documentation:
  - [FIPS Overrides](#) (see page 291)
  - [Create FIPS 140 Images](#) (see page 291)

#### 6.7.3.4.1 GPU Only Steps

- If the [NVIDIA runfile](#)<sup>225</sup> 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
```

#### 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

<sup>224</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

<sup>225</sup> <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

```
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 451) 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](#)<sup>226</sup> or see these pages in the documentation:
  - [FIPS Overrides](#) (see page 452)
  - [Create FIPS 140 Images](#) (see page 462)

#### 6.7.3.4.2 Seed your docker registry

Before creating a Kubernetes cluster you must have the required images in a local docker registry. This registry must be accessible from both the bastion machine and the machines that will be created for the Kubernetes cluster.

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2
```

2. Set an environment variable with your registry address:

```
export DOCKER_REGISTRY_ADDRESS=<registry-address>:<registry-port>
```

3. Run the following command to load the air-gapped image bundle into your private Docker registry:

```
./dkp push image-bundle --image-bundle ./container-images/konvoy-image-bundle-v2.4.2.tar --to-registry $DOCKER_REGISTRY_ADDRESS
```

<sup>226</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

It may 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 Docker registry.

Then [begin creating the bootstrap cluster \(see page 296\)](#).

## 6.7.4 Pre-provisioned Bootstrap

### Bootstrap a kind 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.

#### 6.7.4.1 Bootstrap a Kind Cluster and CAPI controllers

Use the following command to create a bootstrap cluster:

```
dkp create bootstrap
```

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 [Configure HTTP Proxy \(see page 494\)](#).

The output resembles:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

When the bootstrap cluster is up, [create the correct secrets and overrides \(see page 296\)](#).

## 6.7.5 Pre-provisioned Create Necessary Secrets and Overrides

### 6.7.5.1 Create necessary secrets and overrides for pre-provisioned clusters

DKP requires SSH access to your infrastructure with superuser privileges. You must provide an unencrypted SSH private key to DKP.

Populate this key and create the required secret, on your bootstrap cluster using the following procedure.

#### 6.7.5.2 Name your cluster

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
```

### 6.7.5.3 Create a unique cluster name

(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
```

### 6.7.5.4 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
```

### 6.7.5.5 Create overrides

If your pre-provisioned machines need to have [override](#) (see page 451) files, you must create a secret that includes all of the overrides you want to provide in one file. For 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
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=

```

```

secret/preprovisioned-example-user-overrides-md-1 created
secret/preprovisioned-example-user-overrides-md-1 labeled

```

When using Oracle 7 OS, you may wish to 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=

```

When this step is complete, [define the infrastructure nodes and partitions \(see page 298\)](#).

## 6.7.6 Pre-provisioned 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.

### 6.7.6.1 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

```

3. Apply the new infrastructure file with the command:

```
envsubst < preprovisioned_inventory.yaml | kubectl apply -f -
```

```

preprovisionedinventory.infrastructure.cluster.konvoy.d2iq.io/preprovisioned-
example-control-plane created
preprovisionedinventory.infrastructure.cluster.konvoy.d2iq.io/preprovisioned-
example-md-0 created

```

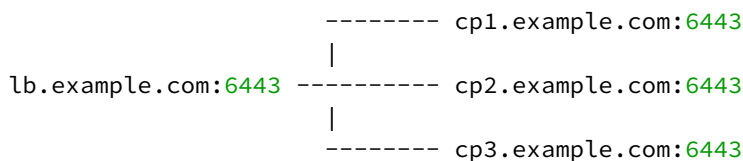
After defining the infrastructure, [define the control plane endpoint \(see page 300\)](#).

## 6.7.7 Pre-provisioned Define Control Plane Endpoint

### 6.7.7.1 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.



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`.

### 6.7.7.2 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.

### 6.7.7.3 Built-in virtual IP

If an external load balancer is not available, use the [built-in virtual IP](#) (see page 302). 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.

### 6.7.7.4 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 466).

When the API server endpoints are defined, you can [create the cluster](#) (see page 302).

### 6.7.7.5 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.

## 6.7.8 Pre-provisioned Create new Cluster

### 6.7.8.1 Create a Kubernetes cluster using the infrastructure definition

Once you've defined the [infrastructure](#) (see page 298) and [control plane endpoints](#) (see page 300), you can proceed to creating the cluster by following these steps to create a new pre-provisioned cluster:

1. With the inventory, and the control plane endpoint defined, use the `dkp` binary to create a Konvoy cluster. 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:** To increase [Docker Hub's rate limit](#)<sup>227</sup> 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-endpoint-host <control plane endpoint host> --control-plane-endpoint-port
<control plane endpoint port, if different than 6443>
```

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 [Configure HTTP Proxy](#) (see page 494).

The output resembles:

<sup>227</sup> <https://docs.docker.com/docker-hub/download-rate-limit/>

```

Generating cluster resources
cluster.cluster.x-k8s.io/preprovisioned-example created
kubeadmcontrolplane.controlplane.cluster.x-k8s.io/preprovisioned-example-
control-plane created
preprovisionedcluster.infrastructure.cluster.konvoy.d2iq.io/preprovisioned-
example created
preprovisionedmachinetemplate.infrastructure.cluster.konvoy.d2iq.io/
preprovisioned-example-control-plane created
secret/preprovisioned-example-etcd-encryption-config created
machinedeployment.cluster.x-k8s.io/preprovisioned-example-md-0 created
preprovisionedmachinetemplate.infrastructure.cluster.konvoy.d2iq.io/
preprovisioned-example-md-0 created
kubeadmconfigtemplate.bootstrap.cluster.x-k8s.io/preprovisioned-example-md-0
created
clusterresourceset.addons.cluster.x-k8s.io/calico-cni-installation-
preprovisioned-example created
configmap/calico-cni-installation-preprovisioned-example created
configmap/tigera-operator-preprovisioned-example created
clusterresourceset.addons.cluster.x-k8s.io/local-volume-provisioner-
preprovisioned-example created
configmap/local-volume-provisioner-preprovisioned-example created
clusterresourceset.addons.cluster.x-k8s.io/node-feature-discovery-
preprovisioned-example created
configmap/node-feature-discovery-preprovisioned-example created
clusterresourceset.addons.cluster.x-k8s.io/nvidia-feature-discovery-
preprovisioned-example created
configmap/nvidia-feature-discovery-preprovisioned-example created
clusterresourceset.addons.cluster.x-k8s.io/metallb-preprovisioned-example
created
configmap/metallb-installation-preprovisioned-example created

```

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

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

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

3. Run the `dkp create cluster` command.

**NOTE:** (Optional) If you have [overrides for your clusters](#) (see page 296), 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 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
```

**NOTE:** If your cluster is air-gapped or you have a local docker registry you must provide additional arguments when creating the cluster.

```
export DOCKER_REGISTRY_URL="<https/http>://<registry-address>:<registry-port>"
export DOCKER_REGISTRY_CA="<path to the CA on the bastion>"
export DOCKER_REGISTRY_USERNAME="<username>"
export DOCKER_REGISTRY_PASSWORD="<password>"
```

- `DOCKER_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.
- `DOCKER_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.
- `DOCKER_REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
- `DOCKER_REGISTRY_PASSWORD` : optional if username is not set.

```
dkp create cluster preprovisioned --cluster-name ${CLUSTER_NAME} \
--control-plane-endpoint-host <control plane endpoint host> \
--control-plane-endpoint-port <control plane endpoint port, if different than 6443> \
--registry-mirror-url=${DOCKER_REGISTRY_URL} \
--registry-mirror-cacert=${DOCKER_REGISTRY_CA} \
--registry-mirror-username=${DOCKER_REGISTRY_USERNAME} \
--registry-mirror-password=${DOCKER_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
```

## 6.7.8.2 Audit logs

To modify Control Plane Audit logs settings using the information contained in the page [Configuring the Control Plane](#) (see page 466).



## 6.7.8.3 Modify the Calico installation

### 6.7.8.3.1 Set the interface

Before exploring the new cluster, confirm your `calico` installation is correct. By default, Calico automatically detects the IP to use for each node using the `first-found` [method](#)<sup>228</sup>. This is not always appropriate for your particular nodes. In that case, you must modify Calico's configuration to use a different method. An alternative is to use the `interface` method by providing the interface ID to use. Follow the steps outlined in this section to modify Calico's configuration. In this example, all cluster nodes use `ens192` as the interface name.

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-twl6b                              | 0/1   | PodInitializing | 0        | 9s    |
| calico-system  | calico-node-wktkh                              | 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 |
| <b>default</b> | 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 |

<sup>228</sup> <https://projectcalico.docs.tigera.io/reference/node/configuration#ip-autodetection-methods>

```

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
ernal 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-zcqx
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-vpvpl
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 `installation` file.

To edit the installation file, run the command:

```
kubectl edit installation default --kubeconfig ${CLUSTER_NAME}.conf
```

Change the value for `spec.calicoNetwork.nodeAddressAutodetectionV4` to `interface: ens192`, and save the file:

```
spec:
 calicoNetwork:
 ...
 nodeAddressAutodetectionV4:
 interface: ens192
```

Save this file. 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.

### 6.7.8.3.2 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.

IPIP is the default encapsulation method.

To change the encapsulation, run the following command:

```
kubectl edit installation default --kubeconfig ${CLUSTER_NAME}.conf
```

Change the value for `spec.calicoNetwork.ipPools[0].encapsulation`

```
spec:
 calicoNetwork:
 ipPools:
 - encapsulation: VXLAN
```

The supported values are "IIPCrossSubnet", "IPIP", "VXLAN", "VXLANCrossSubnet", and "None".

#### 6.7.8.3.2.1 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.

### 6.7.8.3.2.2 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 endpoint 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.



Be aware that switching encapsulation modes can cause disruption to in-progress connections. Plan accordingly.

For more information, see:

- [Calico Overlay Networking](#)<sup>229</sup>
- [IP-in-IP RFC 2003](#)<sup>230</sup>
- [VXLAN RFC 7348](#)<sup>231</sup>

### 6.7.8.4 Use the built-in Virtual IP

As explained in [Define the Control Plane Endpoint](#) (see page 300), 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](#)<sup>232</sup> 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> |

#### 6.7.8.4.1 Virtual IP example

```
dkp create cluster preprovisioned \
 --cluster-name ${CLUSTER_NAME} \
```

<sup>229</sup> <https://docs.projectcalico.org/networking/vxlan-ipip>

<sup>230</sup> <https://datatracker.ietf.org/doc/html/rfc2003>

<sup>231</sup> <https://datatracker.ietf.org/doc/html/rfc7348>

<sup>232</sup> <https://kube-vip.io/>

```
--control-plane-endpoint-host 196.168.1.10 \
--virtual-ip-interface eth1
```

Confirm that your [Calico installation is correct](#) (see page 305).

## 6.7.8.5 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.

### 6.7.8.5.1 Flatcar Linux example

```
dkp create cluster preprovisioned \
 --cluster-name ${CLUSTER_NAME} \
 --os-hint flatcar
```

Confirm that your [Calico installation is correct](#) (see page 305).

- For provisioning DKP on Flatcar, DKP configures cluster nodes to use [Control Groups \(cgroups\) version 1](#)<sup>233</sup>. In versions prior to Flatcar 3033.2.4, a restart is required in order to apply the changes to the kernel. For more information, refer to the [Flatcar documentation](#)<sup>234</sup>.

## 6.7.8.6 Use an 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> |


<sup>233</sup> <https://www.kernel.org/doc/html/latest/admin-guide/cgroup-v1/cgroups.html#what-are-cgroups>

<sup>234</sup> <https://www.flatcar.org/docs/latest/container-runtimes/switching-to-unified-cgroups/#starting-new-nodes-with-legacy-cgroups>

| Proxy configuration                      | Flag                                          |
|------------------------------------------|-----------------------------------------------|
| 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 296). For more information, refer to [this documentation](#) (see page 459).

#### 6.7.8.6.1 HTTP proxy example

 To increase [Docker Hub's rate limit](#)<sup>235</sup> 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.d
efault.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluste
r.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.d
```

<sup>235</sup> <https://docs.docker.com/docker-hub/download-rate-limit/>

```
default.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluster.local"
```

Confirm that your [Calico installation is correct](#) (see page 305).

### 6.7.8.6.2 Use an alternative mirror

To apply Docker registry configurations during the create operation, add the appropriate flags to the `create cluster` command:

| Docker 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>  |

This is useful when using an internal registry and when Internet access is not available (air-gapped installations).

When the cluster is up and running, you can deploy and test workloads.

### 6.7.8.6.3 Alternative mirror example

```
dkp create cluster preprovisioned \
 --cluster-name ${CLUSTER_NAME} \
 --registry-mirror-cacert /tmp/registry.pem \
 --registry-mirror-url https://registry.example.com
```

Confirm that your [Calico installation is correct](#) (see page 305).

### 6.7.8.7 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:

```
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.



Confirm that your [Calico installation is correct](#) (see page 305).

**i** When you complete this procedure, move on to [\(2.4\) Pre-provisioned make Cluster Self-managed](#) (see page 313) to continue the process.

## 6.7.9 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.

Before you start, make sure you have created a workload cluster, as described in [Create the Cluster](#) (see page 302).

#### 6.7.9.1 Make the New Kubernetes Cluster Manage Itself

**i** 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 [Configure HTTP Proxy](#) (see page 494).

The output resembles:

```
✓ 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](#)<sup>236</sup>. First unset the kubeconfig and then move the CAPI:

```
unset KUBECONFIG
```

```
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 [Configure HTTP Proxy](#) (see page 494).

The output resembles:

```
✓ 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.

<sup>236</sup> <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```
dkp describe cluster --kubeconfig ${CLUSTER_NAME}.conf -c ${CLUSTER_NAME}
```

```
NAME
READY SEVERITY REASON SINCE MESSAGE True
Cluster/preprovisioned-example 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
```

### 6.7.9.1.1 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.

## 6.7.10 Pre-provisioned Configure MetalLB

### 6.7.10.1 Create a MetalLB configmap for your pre-provisioned infrastructure.


Choose one of the following two protocols you want to use to announce service IPs, either Layer 2 or BGP configurations.

#### 6.7.10.2 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.

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
```

### 6.7.10.3 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
```

## 6.7.11 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](#)<sup>237</sup>.

### 6.7.11.1 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}
```

2. (Optional) If your pre-provisioned machines have [overrides](#) (see page 451), 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](#) (see page 296).
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}
```

<sup>237</sup> <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

### 6.7.11.2 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 :
machinedeployments.cluster.x-k8s.io "demo-cluster-md-invalid" not found
```

## 6.7.12 Pre-provisioned Add Nodes to Existing Node Pools

This section covers the prerequisites and procedure you need to scale-up or scale-down nodes in an existing DKP cluster.

### 6.7.12.1 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 298).
- Update the `preprovisioned_inventory.yaml` with the new host addresses.
- Run the `kubectl apply` command.

### 6.7.12.2 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 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.cluster.x-k8s.io/<name-of-machine-deployment> scaled
```

5. 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 6
2
```

6. 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 coytestds 3h33m Running 6 6 6
```

7. 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 PHASE VERSION NODENAME
PROVIDERID
```

### 6.7.12.3 Scale Down a Cluster Node

Follow these steps

1. Run this command on your worker nodes:

```
kubectl scale machinedeployment ${CLUSTER_NAME}-md-0 --replicas <new number>
```

2. For control plane nodes, execute the following command:

```
kubectl scale kubeadmcontrolplane ${CLUSTER_NAME}-control-plane --replicas <new number>
```

#### 6.7.12.3.1 Additional Notes for Scaling Down

It is possible for machines to get stuck in the preprovisioning 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 ${CLUSTER_NAME}-md-0-<hash>
```

### 6.7.13 GPU Nodepools in a Pre-provisioned Environment

For pre-provisioned environments, DKP has introduced `nvidia-runfile` flag for Air-gapped Pre-provisioned.

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
```

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 451), 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](#) (see page 296).
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}
```

For more information, see:

- [Pre-provisioned Create and Delete Node Pools](#) (see page 317)
- [Pre-provisioned Prerequisites Air-gapped](#) (see page 291)

## 6.7.14 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 313), see [Delete the workload cluster](#) (see page 325).

### 6.7.14.1 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](#)<sup>238</sup>.

<sup>238</sup> <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```

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

```

**NOTE:** 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. 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.

### 6.7.14.2 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

```

#### 6.7.14.2.1 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:

Delete the `kind` Kubernetes cluster:

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

```
✓ Deleting bootstrap cluster
```

#### 6.7.14.2.2 Next Step:

Once your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will install the [Kommander](#) (see page 475) component of DKP to see your dashboard and continue customization.

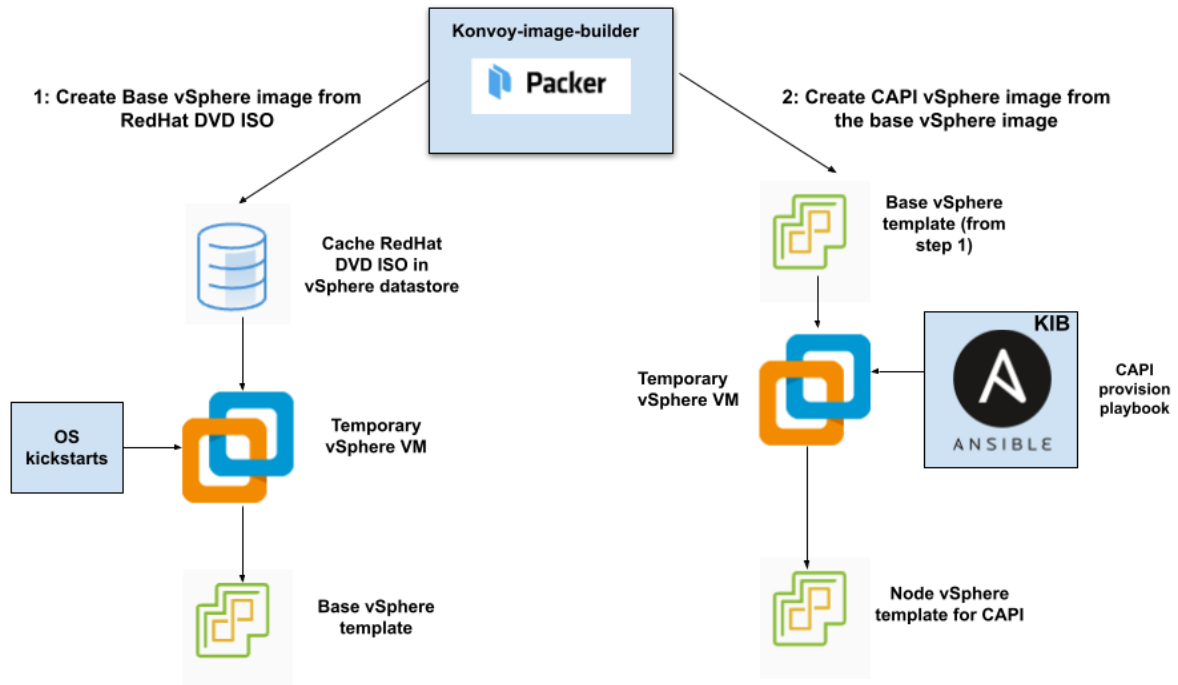
## 6.8 vSphere Infrastructure

### 6.8.1 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 327).
2. Create a bastion VM host 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 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.

To get started, fulfill the [prerequisites](#) (see page 327).

## 6.8.2 vSphere Prerequisites

### 6.8.2.1 Prepare your environment to run DKP with VMware vSphere


Fulfilling the prerequisites involves completing these two areas:

- DKP prerequisites
- vSphere prerequisites

## 6.8.2.2 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 100) image bundle for Linux or macOS.
- [Docker](#)<sup>239</sup> 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. Configure this virtual machine with at least 8GB of memory.

- [kubecti](#)<sup>240</sup> 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.

## 6.8.2.3 VMware vSphere Prerequisites

Before installing, verify that your [VMware vSphere Client environment](#)<sup>241</sup> 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 121) with access to a configured Docker registry. VMWare site has more information on [vSphere Bastion Hosts](#)<sup>242</sup>.
- Valid vSphere values for the following:
  - vCenter API server URL.
  - Datacenter name.

<sup>239</sup> <https://docs.docker.com/get-docker/>

<sup>240</sup> <https://kubernetes.io/docs/tasks/tools/#kubecti>

<sup>241</sup> [https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm\\_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html](https://docs.vmware.com/en/VMware-vSphere/6.7/com.vmware.vsphere.vm_admin.doc/GUID-55238059-912E-411F-A0E9-A7A536972A91.html)

<sup>242</sup> <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.security.doc/GUID-6975426F-56D0-4FE2-8A58-580B40D2F667.html>



- Zone name that contains [ESXi hosts](#)<sup>243</sup> 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.
- 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.

#### 6.8.2.4 Next Steps:

- Non-air-gapped [Create a Base OS image in vSphere](#) (see page 332)
- Air-gapped [Create a Base Air-gapped OS VM Image](#) (see page 354)

#### 6.8.2.5 Minimum User Permissions

##### 6.8.2.5.1 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 327).
2. Select Home > Administration > Roles > Add Role.
3. Give the new role a name, then select these Privileges:

|            |
|------------|
| <b>Cns</b> |
|------------|

<sup>243</sup> <https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.esxi.install.doc/GUID-B2F01BF5-078A-4C7E-B505-5DFFED0B8C38.html>

|                                     |                                                                                         |
|-------------------------------------|-----------------------------------------------------------------------------------------|
| <input checked="" type="checkbox"/> | Searchable                                                                              |
| <b>Datastore</b>                    |                                                                                         |
| <input checked="" type="checkbox"/> | Allocate space                                                                          |
| <input checked="" type="checkbox"/> | Low level file operations                                                               |
| <b>Host</b>                         |                                                                                         |
|                                     | • Configuration                                                                         |
| <input checked="" type="checkbox"/> | Storage partition configuration                                                         |
| <b>Profile-driven storage</b>       |                                                                                         |
| <input checked="" type="checkbox"/> | Profile-driven storage view                                                             |
| <b>Network</b>                      |                                                                                         |
| <input checked="" type="checkbox"/> | Assign network                                                                          |
| <b>Resource</b>                     |                                                                                         |
| <input checked="" type="checkbox"/> | Assign virtual machine to resource pool                                                 |
| <b>Virtual machine</b>              |                                                                                         |
|                                     | • 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     |
| <b>Edit inventory</b>               |                      |
| <input checked="" type="checkbox"/> | Create from existing |
| <input checked="" type="checkbox"/> | Remove               |
| <b>Interaction</b>                  |                      |
| <input checked="" type="checkbox"/> | Power off            |
| <input checked="" type="checkbox"/> | Power on             |
| <b>Provisioning</b>                 |                      |
| <input checked="" type="checkbox"/> | Clone template       |
| <input checked="" type="checkbox"/> | Deploy template      |
| <b>Session</b>                      |                      |
| <input checked="" type="checkbox"/> | ValidateSession      |

Add the permission at the highest level and set to propagate the permissions.

## 6.8.2.6 vSphere Storage Options

### 6.8.2.6.1 Explore storage options and considerations for using DKP with VMware vSphere

The [vSphere Container Storage](#)<sup>244</sup> plugin supports shared NFS, vNFS, and vSAN. You need to provision your storage options in vCenter prior to [creating a CAPI image](#) (see page 356) 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.

## 6.8.3 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](#) (see page 411)(KIB) to create a VM template to configure Kubernetes nodes by the DKP vSphere provider.

### 6.8.3.1 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 for packer. This user should have administrator privileges. It is possible to configure a custom user and password when building the OS image, however, it requires the Konvoy Image Builder (KIB) configuration to be [overridden](#) (see page 451).

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.
- DKP advises not to use usernames and passwords for security reasons, but instead to use private and public keys. If that is not possible, passwords should be generated and a minimum of 20 characters long.

---

<sup>244</sup><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>

### 6.8.3.1.1 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 116).
- The minimum storage requirements for your organization.

#### 6.8.3.1.1.1 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.

#### 6.8.3.1.1.2 Customization

You can specify a disk size when you [create a cluster](#) (see page 338) (see the flags of the [create cluster vsphere](#) (see page 946) 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 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 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.

#### 6.8.3.1.2 Next Step:

[Create a VM Template](#) (see page 333)

## 6.8.4 Create a VM Template

### 6.8.4.1 Prerequisites

- Users need to create a [base OS image \(see page 430\)](#) in vSphere before starting this procedure.
- Build image 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 \(see page 427\)](#)

### 6.8.4.2 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.

```

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: ""
 datacenter: ""
 datastore: ""
 folder: ""
 insecure_connection: "false"
 network: ""
```

```

resource_pool: ""
template: "base-rhel-7.9"
vsphere_guest_os_type: "rhel7_64Guest"
guest_os_type: "rhel7-64"
goss params
distribution: "RHEL"
distribution_version: "7.9"
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 # if 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>
```

The Konvoy Image Builder uses the values in `image.yaml` and the input base OS image to create a vSphere template that contains the required artifacts needed to create a Kubernetes cluster. Give the file a suitable name using this suggested naming convention: `creator-ova-vsphere-OS-ver-k8sver-unique_identifier`. As an example, the filename you create might resemble `dkp-ova-vsphere-rhel-79-1.24.6-1646938922`.

NOTE: To build an image in a GPU environment, use the `overrides/nvidia.yaml` and see the referenced link [KIB for GPU \(see page 427\)](#).

5. DKP creates the new vSphere template directly on the vCenter server.
6. Next steps are to deploy a DKP cluster using your vSphere template.

Next, create a Kubernetes [Bootstrap Cluster \(see page 335\)](#) to enable creating your vSphere cluster and moving CAPI objects to it.

## 6.8.5 vSphere Bootstrap

### 6.8.5.1 Prepare to deploy Kubernetes clusters

To create Kubernetes clusters, Konvoy uses [Cluster API](#)<sup>245</sup> (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](#)<sup>246</sup>) tool.

<sup>245</sup> <https://cluster-api.sigs.k8s.io/>

<sup>246</sup> <https://github.com/kubernetes-sigs/kind>

### 6.8.5.2 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 \(see page 333\)](#)

### 6.8.5.3 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
```

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 [Configure HTTP Proxy \(see page 494\)](#).

The output resembles this example:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

Konvoy creates a bootstrap cluster using [KIND](#)<sup>247</sup> as a library. Konvoy then deploys the following [Cluster API](#)<sup>248</sup> providers on the cluster:

- [Core Provider](#)<sup>249</sup>
  - [vSphere Infrastructure Provider](#)<sup>250</sup>
  - [Kubeadm Bootstrap Provider](#)<sup>251</sup>
  - [Kubeadm ControlPlane Provider](#)<sup>252</sup>
3. Ensure that the CAPV controllers are present with the command:

<sup>247</sup> <https://github.com/kubernetes-sigs/kind>

<sup>248</sup> <https://cluster-api.sigs.k8s.io/>

<sup>249</sup> <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

<sup>250</sup> <https://github.com/kubernetes-sigs/cluster-api-provider-vsphere>

<sup>251</sup> <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

<sup>252</sup> <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>



```
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 |

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
```

| NAMESPACE                         | READY | UP-TO-DATE | AVAILABLE | NAME                                          | AGE |
|-----------------------------------|-------|------------|-----------|-----------------------------------------------|-----|
| capa-system                       | 1/1   | 1          | 1         | capa-controller-manager                       | 22m |
| capi-kubeadm-bootstrap-system     | 1/1   | 1          | 1         | capi-kubeadm-bootstrap-controller-manager     | 22m |
| capi-kubeadm-control-plane-system | 1/1   | 1          | 1         | capi-kubeadm-control-plane-controller-manager | 22m |
| capi-system                       | 1/1   | 1          | 1         | capi-controller-manager                       | 22m |
| cappp-system                      | 1/1   | 1          | 1         | cappp-controller-manager                      | 22m |
| capv-system                       | 1/1   | 1          | 1         | capv-controller-manager                       | 22m |
| capz-system                       | 1/1   | 1          | 1         | capz-controller-manager                       | 22m |
| cert-manager                      | 1/1   | 1          | 1         | cert-manager                                  | 22m |
| cert-manager                      | 1/1   | 1          | 1         | cert-manager-cainjector                       | 22m |
| cert-manager                      | 1/1   | 1          | 1         | cert-manager-webhook                          | 22m |

### 6.8.5.3.1 Next Step:

When the bootstrap cluster is running, you are ready to [Create a new vSphere cluster \(see page 338\)](#).

## 6.8.6 Create new vSphere Cluster

### 6.8.6.1 Prerequisites

Before you begin, make sure you have created a [vSphere Bootstrap](#) (see page 335) cluster.

### 6.8.6.2 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
```

### 6.8.6.3 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](https://docs.docker.com/docker-hub/download-rate-limit/)<sup>253</sup> 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 vsphere](#) (see [page 946](#)) 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> \
 --server <VCENTER_API_SERVER_URL> \
 --ssh-public-key-file <SSH_PUBLIC_KEY_FILE> \
 --resource-pool <RESOURE_POOL_NAME> \
 --virtual-ip-interface <ip_interface_name> \
 --vm-template <TEMPLATE_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 [Configure HTTP Proxy](#) (see [page 494](#)).

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
```

<sup>253</sup> <https://docs.docker.com/docker-hub/download-rate-limit/>

```
lt.svc.cluster,kubernetes.default.svc.cluster.local,.svc,.svc.cluster,.svc.cluster.local,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](https://docs.docker.com/docker-hub/download-rate-limit/)<sup>254</sup> 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:

<sup>254</sup> <https://docs.docker.com/docker-hub/download-rate-limit/>

- =
 Familiarize yourself with Cluster API before editing the cluster objects as edits can prevent the cluster from deploying successfully.

The objects are [Custom Resources](#)<sup>255</sup> defined by Cluster API components, and they belong in three different categories:

1. Cluster

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](#)<sup>256</sup> 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 466).

8. Create the cluster from the objects.

```
kubectl apply -f ${CLUSTER_NAME}.yaml
```

```
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
```

<sup>255</sup> <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

<sup>256</sup> <https://cluster-api.sigs.k8s.io/user/concepts.html>

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

```
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.

9. 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
```

10. 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   |
| vsphere://421638e2-e776-9af6-f683-5e105de5da5a | Running            | 13h v1.22.8                              |
| d2iq-e2e-cluster-1-control-plane-vncbl         | d2iq-e2e-cluster-1 | d2iq-e2e-cluster-1-control-plane-vncbl   |
| vsphere://42168835-7fef-95c4-3652-ebcad3e10d36 | Running            | 13h v1.22.8                              |
| d2iq-e2e-cluster-1-control-plane-wbgrm         | d2iq-e2e-cluster-1 | d2iq-e2e-cluster-1-control-plane-wbgrm   |
| vsphere://421642df-afc4-b6c2-9e61-5b86e7c37eac | Running            | 13h v1.22.8                              |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4       | d2iq-e2e-cluster-1 | d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4 |
| vsphere://4216f467-8483-73cb-a8b6-8d6a4a71e4b4 | Running            | 14h v1.22.8                              |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-n2skc       | d2iq-e2e-cluster-1 | d2iq-e2e-cluster-1-md-0-74c849dc8c-n2skc |
| vsphere://42161cde-9904-4dd2-7a3e-cdfc7655f090 | Running            | 14h v1.22.8                              |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-nkftv       | d2iq-e2e-cluster-1 | d2iq-e2e-cluster-1-md-0-74c849dc8c-nkftv |
| vsphere://42163a0d-eb8d-b5a6-82d5-188e24817c00 | Running            | 14h v1.22.8                              |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-sqklv       | d2iq-e2e-cluster-1 | d2iq-e2e-cluster-1-md-0-74c849dc8c-sqklv |
| vsphere://42161dff-92a5-6da9-7ac1-e987e2c8fed2 | Running            | 14h v1.22.8                              |

## 11. 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       |
| AVAILABLE                        | REPLICAS           | READY       | UPDATED    |
| 3                                | 3                  | 3           | 0          |
|                                  | UNAVAILABLE        | AGE         | VERSION    |
|                                  | 14h                | v1.22.8     |            |

## 12. Describe the kubeadm control plane and check its status and events with the command:

```
kubectl describe kubeadmcontrolplane
```

13. 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"`.

#### 6.8.6.4 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 \(see page 349\)](#). 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.

##### 6.8.6.4.1 Next Step:

[Explore a vSphere Cluster \(see page 344\)](#)

## 6.8.7 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 \(see page 338\)](#) and, if needed, that you have [made the cluster self-managing \(see page 349\)](#).



### 6.8.7.1 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
```

### 6.8.7.2 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
```

### 6.8.7.3 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                                     | STATUS | ROLES                | AGE |
|------------------------------------------|--------|----------------------|-----|
| VERSION                                  |        |                      |     |
| d2iq-e2e-cluster-1-control-plane-7llgd   | Ready  | control-plane,master | 20h |
| v1.24.6                                  |        |                      |     |
| d2iq-e2e-cluster-1-control-plane-vncbl   | Ready  | control-plane,master | 19h |
| v1.24.6                                  |        |                      |     |
| d2iq-e2e-cluster-1-control-plane-wbgrm   | Ready  | control-plane,master | 19h |
| v1.24.6                                  |        |                      |     |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4 | Ready  | <none>               | 19h |
| v1.24.6                                  |        |                      |     |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-n2skc | Ready  | <none>               | 19h |
| v1.24.6                                  |        |                      |     |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-nkftv | Ready  | <none>               | 19h |
| v1.24.6                                  |        |                      |     |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-sqklv | Ready  | <none>               | 19h |
| v1.24.6                                  |        |                      |     |

2. 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-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-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-fjfkj                   |  |
| 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-qttkq               |  |
| 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-jhvgm
3/3 Running 0 19h
vmware-system-csi vsphere-csi-node-rp77r
3/3 Running 0 19h
```

## 6.8.8 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 Cluster](#) (see page 338).

### 6.8.8.1 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 [Configure HTTP Proxy](#) (see page 494).

The output resembles:

```
✓ 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](#)<sup>257</sup>.

<sup>257</sup> <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```
dkp move --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 [Configure HTTP Proxy \(see page 494\)](#).

The output resembles:

```
INFO[2021-08-11T12:09:36-07:00] Pivot operation complete.
src="move/move.go:154"
INFO[2021-08-11T12:09:36-07:00] 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"
```

**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
```

```
d2iq-e2e-cluster-1/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 | SEVERITY | REASON | SINCE | MESSAGE | READY |
|------|----------|--------|-------|---------|-------|
|------|----------|--------|-------|---------|-------|

```

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-sqkly True
13h

```

5. Remove the bootstrap cluster, if desired, as the workload cluster is now self-managed:

```
dkp delete bootstrap
```

```
INFO[2022-03-30T17:53:36-07:00] Deleting bootstrap cluster
src="bootstrap/bootstrap.go:182"
```

### 6.8.8.2 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.

Next, you can [explore the new cluster \(see page 344\)](#) or [explore the new air-gapped cluster \(see page 361\)](#).

## 6.8.9 Configure MetalLB for a vSphere infrastructure

### 6.8.9.1 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.

### 6.8.9.2 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
```

When complete, run the following `kubectl` command.

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




### 6.8.9.3 BGP configuration

For a basic 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
```

## 6.8.10 Install vSphere Air-Gapped

### 6.8.10.1 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 327) before you begin the procedures in this section.

### 6.8.10.2 Create a Base Air-gapped OS VM Image

#### 6.8.10.2.1 Create a Base OS VM Image in the VMware vSphere Client (air-gapped)

**6.8.10.2.2** 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.

#### 6.8.10.2.3 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 for packer. This user should have administrator privileges. It is possible to configure a custom user and password when building the OS image, however, it requires the Konvoy Image Builder (KIB) configuration to be [overridden](#) (see page 451).

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.

- DKP advises not to use usernames and passwords for security reasons, but instead to use private and public keys. If that is not possible, passwords should be generated and a minimum of 20 characters long.

#### 6.8.10.2.3.1 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 116).
- 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 specify a disk size when you [create a cluster](#) (see page 338) (see the flags of the [create cluster vsphere](#) (see page 946) 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 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 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.

#### 6.8.10.2.3.2 Next Step:

[Create a CAPI VM Template](#) (see page 356)

### 6.8.10.3 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.

#### Create a vSphere template for your cluster from a base OS image

Using [KIB \(see page 411\)](#), you can build an image without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded \(see page 100\)](#) `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2/kib
```

2. Follow the instructions to [build a vSphere template \(see page 431\)](#) and set the override `--overrides overrides/offline.yaml` flag.

#### 6.8.10.3.1 Next Step:

Now you can [Seed your Air-gapped Docker Registry \(see page 356\)](#).

### 6.8.10.4 vSphere Air-gapped Seed Docker Registry

#### 6.8.10.4.1 Seed your docker registry

Before creating a Kubernetes cluster, you need the required images in a local docker registry. This registry must be accessible from both the bastion machine and the other machines that will be created for the Kubernetes cluster.

1. Assuming you have [downloaded \(see page 100\)](#) `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2
```

2. Set an environment variable with your registry address with this command:

```
export DOCKER_REGISTRY_ADDRESS=<registry-address>:<registry-port>
export DOCKER_REGISTRY_USERNAME=<username>
export DOCKER_REGISTRY_PASSWORD=<password>
```

- Execute the following command to load the air-gapped image bundle into your private Docker registry:

```
dkp push image-bundle --image-bundle ./container-images/konvoy-image-bundle-
v2.4.2.tar --to-registry $DOCKER_REGISTRY_ADDRESS --to-registry-username
$DOCKER_REGISTRY_USERNAME --to-registry-password $DOCKER_REGISTRY_PASSWORD
```

- It may 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 Docker registry.

#### 6.8.10.4.1.1 Next Step:

Once you have seeded your registry, you begin the cluster building process by [Creating the Bootstrap cluster](#) (see page 357).

## 6.8.10.5 vSphere Air-gapped Bootstrap

### 6.8.10.5.1 Prerequisites

Before you perform this procedure, ensure that you have [created a CAPI VM template](#) (see page 356).

### 6.8.10.5.2 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.

- Assuming you have [downloaded](#) (see page 100) `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzvf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2
```

- Load the bootstrap docker image on your bastion machine.

```
docker load -i konvoy-bootstrap-image-v2.4.2.tar
```

- 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 [Configure HTTP Proxy \(see page 494\)](#).

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-nfnfs | 1/1   | Running | 0        | 13h |

#### 6.8.10.5.2.1 Next Step:

Create a new [vSphere Air-gapped Kubernetes cluster \(see page 358\)](#).

### 6.8.10.6 Create a new Air-gapped vSphere Cluster

#### 6.8.10.6.1 Prerequisites

Before you begin, be sure that you have created a [Bootstrap Cluster \(see page 357\)](#)

#### 6.8.10.6.2 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 Docker registry as a mirror when attempting to pull images:

**IMPORTANT:** The image must be created by the [konvoy-image-builder<sup>258</sup>](#) project in order to use the registry mirror feature.

<sup>258</sup> <https://github.com/mesosphere/konvoy-image-builder>

```
export DOCKER_REGISTRY_URL=<https/http>://<registry-address>:<registry-port>
export DOCKER_REGISTRY_CA=<path to the CA on the bastion>
export DOCKER_REGISTRY_USERNAME=<username>
export DOCKER_REGISTRY_PASSWORD=<password>
```

- `DOCKER_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.
  - `DOCKER_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.
  - `DOCKER_REGISTRY_USERNAME` : optional, set to a user that has pull access to this registry.
  - `DOCKER_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:

```
dkp create cluster vsphere
 --cluster-name ${CLUSTER_NAME} \
 --network <NETWORK_NAME> \
 --control-plane-endpoint-host <CONTROL_PLANE_IP> \
 --data-center <DATACENTER_NAME> \
 --data-store <DATASTORE_NAME> \
 --folder <FOLDER_NAME> \
 --server <VCENTER_API_SERVER_URL> \
 --ssh-public-key-file </path/to/key.pub> \
 --resource-pool <RESOURCE_POOL_NAME> \
 --vm-template konvoy-ova-vsphere-os-release-k8s_release-vsphere-timestamp \
 --virtual-ip-interface <ip_interface_name> \
 --extra-sans "127.0.0.1" \
 --registry-mirror-url=${DOCKER_REGISTRY_URL} \
 --registry-mirror-cacert=${DOCKER_REGISTRY_CA} \
 --registry-mirror-username=${DOCKER_REGISTRY_USERNAME} \
 --registry-mirror-password=${DOCKER_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 [Configure HTTP Proxy](#) (see page 494).

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
```

```
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 349) with the information in the linked page.

#### 6.8.10.6.2.1 Next Step:

You can [explore your new cluster](#) (see page 361).

#### 6.8.10.6.3 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.

### 6.8.10.7 Explore vSphere Air-gapped Cluster

#### 6.8.10.7.1 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
```

#### 6.8.10.7.2 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
```

### 6.8.10.7.3 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<br>v1.24.6   | Ready  | control-plane,master | 20h |
| d2iq-e2e-air-gapped-1-control-plane-vncbl<br>v1.24.6   | Ready  | control-plane,master | 19h |
| d2iq-e2e-air-gapped-1-control-plane-wbgrm<br>v1.24.6   | Ready  | control-plane,master | 19h |
| d2iq-e2e-air-gapped-1-md-0-74c849dc8c-67rv4<br>v1.24.6 | Ready  | <none>               | 19h |
| d2iq-e2e-air-gapped-1-md-0-74c849dc8c-n2skc<br>v1.24.6 | Ready  | <none>               | 19h |
| d2iq-e2e-air-gapped-1-md-0-74c849dc8c-nkftv<br>v1.24.6 | Ready  | <none>               | 19h |
| d2iq-e2e-air-gapped-1-md-0-74c849dc8c-sqklv<br>v1.24.6 | 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                                     | STATUS  | RESTARTS    | AGE |
|---------------|------------------------------------------|---------|-------------|-----|
| calico-system | calico-kube-controllers-57fbd7bd59-qqd96 | Running | 0           | 20h |
| calico-system | calico-node-2m524                        | Running | 3 (19h ago) | 19h |
| calico-system | calico-node-bbhg5                        | Running | 0           | 20h |
| calico-system | calico-node-cc5lf                        | Running | 2 (19h ago) | 19h |
| calico-system | calico-node-cwg7x                        | Running | 1 (19h ago) | 19h |
| calico-system | calico-node-d59hn                        | Running | 1 (19h ago) | 19h |
| calico-system | calico-node-qmmcz                        | Running | 0           | 19h |
| calico-system | calico-node-wdqhx                        | Running | 0           | 19h |
| calico-system | calico-typha-655489d8cc-b5jnt            | Running | 0           | 20h |
| calico-system | calico-typha-655489d8cc-q92x9            | Running | 0           | 19h |
| calico-system | calico-typha-655489d8cc-vjlkx            | 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-jhvqm
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 368).

## 6.8.11 vSphere Certificate Renewal

### 6.8.11.1 Configure Automated Renewal for Managed Kubernetes PKI Certificates

### 6.8.11.2 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.

### 6.8.11.2.1 Requirements

This feature requires that you install Python 3.5 or higher version on all control plane hosts.

### 6.8.11.3 Prerequisites

- Complete the [bootstrap Cluster Lifecycle \(see page 335\)](#) topic.

#### 6.8.11.3.1 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.

#### 6.8.11.3.2 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.

### 6.8.11.3.3 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.

#### 6.8.11.3.3.1 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
```

#### 6.8.11.3.3.2 Check renew certs status

To check the status of the `renew-certs` service, use the command:

```
systemctl status renew-certs
```

#### 6.8.11.3.3 Review logs

To get the logs of the last run of the service, use the command:

```
journalctl logs -u renew-certs
```

## 6.8.12 Delete vSphere Cluster

### 6.8.12.1 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 349), see [Delete the workload cluster](#) (see page 370).

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](#)<sup>259</sup>.

```
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:

```
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                                                             |       |

<sup>259</sup> <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```

└─ 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. With Vsphere clusters, `dkp delete` doesn't delete the virtual disks backing the PVs for DKP addons. 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.

### 6.8.12.2 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"
```

### 6.8.12.3 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"
```

### 6.8.12.4 Next Step:

Once your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will install the [Kommander](#) (see page 475) component of DKP to see your dashboard and continue customization.

### 6.8.12.5 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.

## 6.8.13 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](#)<sup>260</sup>.

### 6.8.13.1 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.

#### 6.8.13.1.1 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
```

<sup>260</sup> <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

2. If your workload cluster is self-managed, as described in [Make the New Cluster Self-Managed](#) (see [page 349](#)), configure `kubectl` to use the kubeconfig for the cluster:

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

3. Define your node pool name:

```
export NODEPOOL_NAME=example
```

### 6.8.13.1.2 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 \
 --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` flags to get a complete set of node pool objects to modify locally or store in version control.

## 6.8.13.2 List vSphere Node Pools

### 6.8.13.2.1 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<br>VERSION | DESIRED | READY | KUBERNETES |
|---------------------|---------|-------|------------|
| demo-cluster-md-0   | 4       | 4     | v1.24.6    |
| example             | 3       | 0     | v1.24.6    |

## 6.8.13.3 Scale vSphere Node Pools

### 6.8.13.3.1 Scaling node pools

While you can run [Cluster Autoscaler](#) (see page 377), 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.

#### 6.8.13.3.1.1 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<br>VERSION | DESIRED | READY | KUBERNETES |
|---------------------|---------|-------|------------|
| example             | 5       | 5     | v1.24.6    |
| demo-cluster-md-0   | 4       | 4     | v1.24.6    |

#### 6.8.13.3.1.2 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 controlled by [CAPI's delete policy](#)<sup>261</sup>. 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`.

<sup>261</sup> [https://github.com/kubernetes-sigs/cluster-api/blob/v0.4.0/api/v1alpha4/machineset\\_types.go#L85-L105](https://github.com/kubernetes-sigs/cluster-api/blob/v0.4.0/api/v1alpha4/machineset_types.go#L85-L105)

```
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"
```

### 6.8.13.3.1.3 Scaling node pools when using cluster autoscaler

If you [configured the cluster autoscaler](#)<sup>262</sup> 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
```

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
```

<sup>262</sup> <https://docs.d2iq.com/dkp/konvoy/2.2/choose-infrastructure/vsphere/nodepools/cluster-autoscaler>



Similarly, scaling down to a number of replicas less than the configured `min-size` also returns an error.

### 6.8.13.4 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
```

### 6.8.13.5 vSphere Cluster Autoscaler

#### 6.8.13.5.1 Cluster Autoscaler

[Cluster Autoscaler](#)<sup>263</sup> 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](#)<sup>264</sup>, 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:

<sup>263</sup> <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

<sup>264</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-fast-is-hpa-when-combined-with-ca>

```
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](#)<sup>265</sup>.

For more information about how Cluster Autoscaler works, see these documents:

- [What is Cluster Autoscaler](#)<sup>266</sup>
- [How does scale-up work](#)<sup>267</sup>
- [How does scale-down work](#)<sup>268</sup>
- [CAPI Provider for Cluster Autoscaler](#)<sup>269</sup>

#### 6.8.13.5.1.1 Cluster Autoscaler Prerequisites

Before you begin, you must have:

- A [Bootstrap Cluster Lifecycle](#) (see page 335).
- [Created a new Kubernetes Cluster](#) (see page 338).
- A [Self-Managed Cluster](#) (see page 349).

#### 6.8.13.5.1.2 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
```

<sup>265</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-are-the-parameters-to-ca>

<sup>266</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-is-cluster-autoscaler>

<sup>267</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-up-work>

<sup>268</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-down-work>

<sup>269</sup> <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

```
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
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.

## 6.8.14 Replace a vSphere Node

### 6.8.14.1 Prerequisites

Before you begin, you must:

- [Create a workload cluster \(see page 338\)](#) or [create an air-gapped workload cluster \(see page 358\)](#).
- [Make the new cluster self-managed \(see page 349\)](#).

### 6.8.14.2 Replace a worker node

In certain situations, you may want to delete a worker node and have [Cluster API](#)<sup>270</sup> 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<br>v1.24.6   | Ready  | control-plane,master | 20h |
| d2iq-e2e-cluster-1-control-plane-vncbl<br>v1.24.6   | Ready  | control-plane,master | 20h |
| d2iq-e2e-cluster-1-control-plane-wbgrm<br>v1.24.6   | Ready  | control-plane,master | 19h |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-67rv4<br>v1.24.6 | Ready  | <none>               | 20h |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-n2skc<br>v1.24.6 | Ready  | <none>               | 20h |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-nkftv<br>v1.24.6 | Ready  | <none>               | 20h |
| d2iq-e2e-cluster-1-md-0-74c849dc8c-sqklv<br>v1.24.6 | 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`.

<sup>270</sup> <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 | v1.24.6 | 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.

6. 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.

## 6.9 GCP Infrastructure

The following procedure describes creating a DKP cluster on GCP.

- [GCP Prerequisites](#) (see page 382)
- [GCP Konvoy Image Builder](#) (see page 385)
- [Bootstrap GCP](#) (see page 388)
- [Create a New GCP Cluster](#) (see page 390)
- [Explore the GCP Cluster](#) (see page 393)
- [Make the New GCP Cluster Self-Managed](#) (see page 397)
- [Manage GCP Node Pools](#) (see page 399)
- [Delete a GCP Cluster](#) (see page 406)

### 6.9.1 GCP Prerequisites

#### 6.9.1.1 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.
- The `dkp` binary on this machine.

- [Docker](#)<sup>271</sup> version 18.09.2 or later.
- [kubect](#)<sup>272</sup> for interacting with the running cluster.
- Install the GCP `gcloud` CLI by following <https://cloud.google.com/sdk/docs/install>.

### 6.9.1.2 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
- 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.

### 6.9.1.3 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.

### 6.9.1.4 GCP Prerequisite Roles

- If you are creating the bootstrap cluster on a non-GCP instance or one that does not have the required `editor` role:
  - (option 1) Create a service account using the following `gcloud` commands:

```
export GCP_PROJECT=<your GCP project ID>
export SERVICE_ACCOUNT_USER=<some new service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud iam service-accounts create "$SERVICE_ACCOUNT_USER" --
project=$GCP_PROJECT
```

<sup>271</sup> <https://docs.docker.com/get-docker/>

<sup>272</sup> <https://kubernetes.io/docs/tasks/tools/#kubectl>

```
gcloud projects add-iam-policy-binding $GCP_PROJECT --
member="serviceAccount:
$SERVICE_ACCOUNT_USER@$GCP_PROJECT.iam.gserviceaccount.com" --
role=roles/editor
gcloud iam service-accounts keys create $GOOGLE_APPLICATION_CREDENTIALS
--iam-
account="$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 SERVICE_ACCOUNT_USER=<existing service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud iam service-accounts keys create $GOOGLE_APPLICATION_CREDENTIALS
--iam-
account="$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')
```

- In order 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.



**NOTE:** 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>

## 6.9.2 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 411) (KIB) to create a [Cluster API](#)<sup>273</sup> 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 451) 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 382) section of the documentation.

### 6.9.2.1 Prerequisites

Before you begin, you must:

- Check the [supported DKP version](#) (see page 1056) and download the [KIB](#) (see page 412) bundle (prefixed with `konvoy-image-bundle`) for your OS. Do not use the release prefixed with `konvoy-image-builder`.
- Create a working `Docker` setup.
- On Debian-based Linux distributions, install a version of the [cri-tools](#)<sup>274</sup> 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.

<sup>273</sup> <https://cluster-api.sigs.k8s.io/>

<sup>274</sup> <https://github.com/kubernetes-sigs/cri-tools>

To check if it is enabled, use the commands on this page [https://cloud.google.com/compute/docs/metadata/setting-custom-metadata#console\\_2](https://cloud.google.com/compute/docs/metadata/setting-custom-metadata#console_2) to inspect the metadata configured in your project. If you find the `enable-oslogin` flag set to `TRUE`, you must remove (or set it to `FALSE`) to successfully use KIB.

## 6.9.2.2 GCP Prerequisites

- If you are creating your image on either a non-GCP instance or one that does not have the [required roles](#) (see page 382):
  - (option 1) Create a service account using the following `gcloud` commands:

```
export GCP_PROJECT=<your GCP project ID>
export SERVICE_ACCOUNT_USER=<some new service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud iam service-accounts create "${SERVICE_ACCOUNT_USER}" --project=${GCP_PROJECT}
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/compute.instanceAdmin.v1
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/iam.serviceAccountUser
gcloud iam service-accounts keys create $GOOGLE_APPLICATION_CREDENTIALS --iam-account="${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 SERVICE_ACCOUNT_USER=<existing service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/compute.instanceAdmin.v1
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/iam.serviceAccountUser
```

```
gcloud iam service-accounts keys create $
{GOOGLE_APPLICATION_CREDENTIALS} --iam-account="${SERVICE_ACCOUNT_USER}
@${GCP_PROJECT}.iam.gserviceaccount.com"
```

### 6.9.2.3 Create a Network (optional)

Building an image requires a [Network](#)<sup>275</sup> 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
instance on the network using port 22." --direction=INGRESS --priority=65534 --
source-ranges=0.0.0.0/0 --action=ALLOW --rules=tcp:22
```

### 6.9.2.4 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...
```

<sup>275</sup> <https://cloud.google.com/vpc/docs/vpc>

```

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
```

With your KIB image now created, you can now move onto [Bootstrap GCP](#) (see page 388) and set up your [Cluster API](#)<sup>276</sup> (CAPI) controllers, or run [GCP Quick Start](#) (see page 58) to create a cluster with little customization.



[Konvoy Image Builder](#) (see page 411) has a more detailed section in the documentation if you need to refer there for compatible versions with DKP and other specific information.

## 6.9.3 Bootstrap GCP

To create Kubernetes clusters, DKP uses [Cluster API](#)<sup>277</sup> (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](#)<sup>278</sup>) tool.

### 6.9.3.1 Prerequisites

Before you begin, you must:

- Complete the steps in [Prerequisites](#) (see page 382).
- Ensure the `dkp` binary can be found in your `$PATH`.

<sup>276</sup> <https://cluster-api.sigs.k8s.io/>

<sup>277</sup> <https://cluster-api.sigs.k8s.io/>

<sup>278</sup> <https://github.com/kubernetes-sigs/kind>

### 6.9.3.2 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 [Configure HTTP Proxy \(see page 494\)](#).

The output looks similar to:

```
✓ Creating a bootstrap cluster
✓ Initializing new CAPI components
```

3. DKP creates a bootstrap cluster using [KIND](#)<sup>279</sup> as a library. DKP then deploys the following [Cluster API](#)<sup>280</sup> providers on the cluster:
  - [Core Provider](#)<sup>281</sup>
  - [GCP Infrastructure Provider](#)<sup>282</sup>
  - [Kubeadm Bootstrap Provider](#)<sup>283</sup>
  - [Kubeadm ControlPlane Provider](#)<sup>284</sup>
4. 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 | NAME                    | AGE |
|-------------|-------|------------|-----------|-------------------------|-----|
| capa-system | 1/1   | 1          | 1         | capa-controller-manager | 1h  |

<sup>279</sup> <https://github.com/kubernetes-sigs/kind>

<sup>280</sup> <https://cluster-api.sigs.k8s.io/>

<sup>281</sup> <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/>

<sup>282</sup> <https://github.com/kubernetes-sigs/cluster-api-provider-gcp>

<sup>283</sup> <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/bootstrap/kubeadm>

<sup>284</sup> <https://github.com/kubernetes-sigs/cluster-api/tree/v0.3.20/controlplane/kubeadm>

```

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

```

Once completed, move onto [creating a new GCP cluster](#) (see page 390).

## 6.9.4 Create a New GCP Cluster

### 6.9.4.1 Prerequisites

- Before you begin, ensure you have created a [Bootstrap](#) (see page 388) cluster.


### 6.9.4.2 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 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](#)<sup>285</sup> for more naming information.

 To increase [Docker Hub's rate limit](#)<sup>286</sup> 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.

### 6.9.4.3 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 411).

1. Create an image using [Konvoy Image Builder \(KIB\)](#) (see page 424) 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 466).

<sup>285</sup> <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/>

<sup>286</sup> <https://docs.docker.com/docker-hub/download-rate-limit/>

- (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](#)<sup>287</sup>. 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](#)<sup>288</sup> and network flag.
- Create a Kubernetes cluster. The following example shows a common configuration. See [dkp create cluster gcp](#) (see page 954) reference for the full list of cluster creation options:

```
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 [Configure HTTP Proxy](#) (see page 494).

- Wait for the cluster control-plane to be ready:

```
kubectl wait --for=condition=ControlPlaneReady "clusters/${CLUSTER_NAME}" --
timeout=20m
```

- 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
```

<sup>287</sup> <https://cloud.google.com/vpc/docs/create-modify-vpc-networks#create-auto-network>

<sup>288</sup> <https://github.com/kubernetes-sigs/cluster-api-provider-gcp/blob/3406adaae0f8f65a615844e55d856d306eddacc1/docs/book/src/topics/prerequisites.md#cloud-nat>



```

| | └─ 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

```

After the cluster creation process has finished, move onto [exploring your new cluster](#). (see page 393)

## 6.9.5 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 390).

### 6.9.5.1 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 |       |
| v1.24.6                         |        |                      |       |       |
| gcp-example-control-plane-rtj9h | Ready  | control-plane,master | 104s  |       |
| v1.24.6                         |        |                      |       |       |
| gcp-example-control-plane-zbf9w | Ready  | control-plane,master | 3m23s |       |
| v1.24.6                         |        |                      |       |       |
| gcp-example-md-0-88c46          | Ready  | <none>               | 3m28s | v1.24 |
| .6                              |        |                      |       |       |
| gcp-example-md-0-fp8s7          | Ready  | <none>               | 3m28s | v1.24 |
| .6                              |        |                      |       |       |
| gcp-example-md-0-qvnx7          | Ready  | <none>               | 3m28s | v1.24 |
| .6                              |        |                      |       |       |
| gcp-example-md-0-wjdrq          | Ready  | <none>               | 3m27s | v1.24 |
| .6                              |        |                      |       |       |

- List the Pods with the command:

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf get pods -A
```

- View the following output:

| 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                     |         |         | 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-dzv8    |
| 1/1                               | Running | 0       | 77s                                      |
| capp-system                       |         |         | capp-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-knzb9
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
lh7dc 1/1 Running 0 5m40s
node-feature-discovery
1/1 Running 0 3m40s
node-feature-discovery
1/1 Running 0 3m40s
node-feature-discovery
1/1 Running 0 3m35s
node-feature-discovery
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 397)

## 6.9.6 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 390).

### 6.9.6.1 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 [Configure HTTP Proxy](#) (see page 494).

The output resembles:

```
✓ 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](#)<sup>289</sup>.

```
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
```

**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

<sup>289</sup> <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

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:

**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/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

```

## 6.9.6.2 Known Limitations



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

DKP only supports moving all namespaces in the cluster; DKP does not support migration of individual namespaces.

## 6.9.7 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](https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html)<sup>290</sup>.

### 6.9.7.1 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.

#### 6.9.7.1.1 Prerequisites

Before you begin, make sure you have created a [GCP cluster](#). (see page 390)

<sup>290</sup> <https://cluster-api.sigs.k8s.io/developer/architecture/controllers/machine-deployment.html>

### 6.9.7.1.1.1 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 397](#)), configure `kubectl` to use the kubeconfig for the cluster.

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

3. Define your node pool name.

```
export NODEPOOL_NAME=example
```

### 6.9.7.1.1.2 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](#)<sup>291</sup> in the same same region as your cluster.

```
dkp create nodepool gcp ${NODEPOOL_NAME} \
 --cluster-name=${CLUSTER_NAME} \
 --image $IMAGE_NAME \
 --zone us-west1-b \
 --replicas=3
```

```
machedeployment.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.

<sup>291</sup> <https://cloud.google.com/compute/docs/regions-zones>



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

## 6.9.7.2 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<br>VERSION | DESIRED | READY | KUBERNETES |
|---------------------|---------|-------|------------|
| example             | 3       | 3     | v1.24.6    |
| gcp-example-2-md-0  | 4       | 4     | v1.24.6    |

## 6.9.7.3 Scale GCP Node Pools

While you can run [Cluster Autoscaler](#)<sup>292</sup>, 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.

### 6.9.7.3.1 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:

<sup>292</sup> <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

```
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     | v1.24.6 |
| gcp-example-md-0   | 4       | 4     | v1.24.6 |

### 6.9.7.3.2 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     | v1.22.8 |
| gcp-example-md-0   | 4       | 4     | v1.22.8 |

In a default cluster, the nodes to delete are selected at random. This behavior is controlled by [CAPI's delete policy](#)<sup>293</sup>. 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}
```

<sup>293</sup> [https://github.com/kubernetes-sigs/cluster-api/blob/v0.4.0/api/v1alpha4/machineset\\_types.go#L85-L105](https://github.com/kubernetes-sigs/cluster-api/blob/v0.4.0/api/v1alpha4/machineset_types.go#L85-L105)

✓ Scaling node pool example to 3 replicas

### 6.9.7.3.3 Scaling Node Pools When Using Cluster Autoscaler

If you configured [the cluster autoscaler](#)<sup>294</sup> 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.

## 6.9.7.4 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:

<sup>294</sup> <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

```
✓ 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
```

### 6.9.7.5 GCP Cluster Autoscaler

#### Configure autoscaler for node pools

[Cluster Autoscaler](#)<sup>295</sup> 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](#)<sup>296</sup>, 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](#)<sup>297</sup>.

For more information about how Cluster Autoscaler works, see these documents:

- [What is Cluster Autoscaler](#)<sup>298</sup>
- [How does scale-up work](#)<sup>299</sup>
- [How does scale-down work](#)<sup>300</sup>
- [CAPI Provider for Cluster Autoscaler](#)<sup>301</sup>

<sup>295</sup> <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

<sup>296</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-fast-is-hpa-when-combined-with-ca>

<sup>297</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-are-the-parameters-to-ca>

<sup>298</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#what-is-cluster-autoscaler>

<sup>299</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-up-work>

<sup>300</sup> <https://github.com/kubernetes/autoscaler/blob/master/cluster-autoscaler/FAQ.md#how-does-scale-down-work>

<sup>301</sup> <https://github.com/kubernetes/autoscaler/tree/master/cluster-autoscaler/cloudprovider/clusterapi>

### 6.9.7.5.1 Cluster Autoscaler Prerequisites

Before you begin, you must have:

- A [Bootstrap Cluster Lifecycle](#) (see page 388).
- A [New Kubernetes Cluster](#) (see page 390).
- A [Self-Managed Cluster](#) (see page 397).

### 6.9.7.5.2 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 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.

## 6.9.8 Delete a GCP Cluster

### 6.9.8.1 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 397), 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](#)<sup>302</sup>.

```
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/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
```

<sup>302</sup> <https://cluster-api.sigs.k8s.io/reference/glossary.html?highlight=pivot#pivot>

```

├─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

```

**NOTE:** 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.

**ⓘ** 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.

### 6.9.8.2 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 kubernetes 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.

### 6.9.8.3 Delete the Bootstrap Cluster

1. Delete the bootstrap cluster using `dkp delete`:

```
dkp delete bootstrap --kubeconfig $HOME/.kube/config
```

```
✓ Deleting bootstrap cluster
```

### 6.9.8.4 Next Step:

Once your cluster is built in the Konvoy component of DKP for your infrastructure/environment, you will install the [Kommander](#) (see page 475) component of DKP to see your dashboard and continue customization.

### 6.9.8.5 Known Limitations



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.

## 6.10 Verify DKP installation

### Check DKP components to verify the status of your cluster

In this section we show you how to verify a DKP installation.

## 6.10.1 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}
```

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<br>v1.21.6 | Ready  | <none>               | 135m |
| ip-10-0-122-142.us-west-2.compute.internal<br>v1.21.6 | Ready  | <none>               | 135m |
| ip-10-0-186-214.us-west-2.compute.internal<br>v1.21.6 | Ready  | control-plane,master | 133m |
| ip-10-0-231-82.us-west-2.compute.internal<br>v1.21.6  | Ready  | control-plane,master | 135m |
| ip-10-0-71-114.us-west-2.compute.internal<br>v1.21.6  | Ready  | <none>               | 135m |
| ip-10-0-71-207.us-west-2.compute.internal<br>v1.21.6  | Ready  | <none>               | 135m |
| ip-10-0-85-253.us-west-2.compute.internal<br>v1.21.6  | Ready  | control-plane,master | 137m |

To verify a successful installation, all of the previous commands should return as `Ready` or `True`.

## 6.10.2 Verify all pods are running

All pods installed by DKP should be in `Running` or `Completed` status if the install is successful.

```
kubectl get pods --all-namespaces
```

## 6.10.3 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 841):

```
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 841).

## 6.11 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 system's are capable of running Konvoy.

This section describes how to use the KIB to create a [Cluster API](https://cluster-api.sigs.k8s.io/)<sup>303</sup> compliant machine images. Machine images contain configuration information 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. The KIB uses [variable overrides](#) (see page 451) 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.

### 6.11.1 Prerequisites

Before you begin, you must ensure your versions of KIB and DKP are compatible:

- Check the [supported DKP version](#) (see page 1056) and download the Konvoy Image Builder bundle from the KIB Version section of the chart below (prefixed with `konvoy-image-bundle`) for your Operating System.
- Create a working `Docker` setup.

### 6.11.2 Compatible 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.

| DKP Version | KIB Version (Click for bundle download) |
|-------------|-----------------------------------------|
| v2.4.1      | <a href="#">v1.24.8</a> <sup>304</sup>  |
| v2.4.2      | <a href="#">v1.24.10</a> <sup>305</sup> |
| v2.4.0      | <a href="#">v1.24.5</a> <sup>306</sup>  |
| v2.3.3      | <a href="#">v1.19.14</a> <sup>307</sup> |
| v2.3.2      | <a href="#">v1.19.14</a> <sup>308</sup> |
| v2.3.1      | <a href="#">v1.19.12</a> <sup>309</sup> |

<sup>303</sup> <https://cluster-api.sigs.k8s.io/>

<sup>304</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.24.8>

<sup>305</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.24.10>

<sup>306</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.24.5>

<sup>307</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.19.14>

<sup>308</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.19.14>

<sup>309</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.19.12>

| DKP Version | KIB Version (Click for bundle download) |
|-------------|-----------------------------------------|
| v2.3.0      | <a href="#">v1.19.12</a> <sup>310</sup> |
| v2.2.3      | <a href="#">v1.17.4</a> <sup>311</sup>  |
| v2.2.2      | <a href="#">v1.17.4</a> <sup>312</sup>  |
| v2.2.1      | <a href="#">v1.13.2</a> <sup>313</sup>  |
| v2.2.0      | <a href="#">v1.11.0</a> <sup>314</sup>  |
| v2.1.5      | <a href="#">v1.5.2</a> <sup>315</sup>   |
| v2.1.4      | <a href="#">v1.5.0</a> <sup>316</sup>   |
| v2.1.3      | <a href="#">v1.5.0</a> <sup>317</sup>   |
| v2.1.2      | <a href="#">v1.5.0</a> <sup>318</sup>   |
| v2.1.1      | <a href="#">v1.5.0</a> <sup>319</sup>   |
| v2.1.0      | <a href="#">v1.5.0</a> <sup>320</sup>   |



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.

<sup>310</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.19.12>

<sup>311</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.17.4>

<sup>312</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.17.4>

<sup>313</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.13.2>

<sup>314</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.11.0>

<sup>315</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.5.1>

<sup>316</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.5.1>

<sup>317</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.5.0>

<sup>318</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.5.0>

<sup>319</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.5.0>

<sup>320</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.5.0>

- [KIB with AWS \(see page 414\)](#)
- [KIB for Azure \(see page 420\)](#)
- [KIB with GCP \(see page 424\)](#)
- [KIB for GPU \(see page 427\)](#)
- [KIB with vSphere \(see page 430\)](#)
- [Konvoy Image Builder CLI \(see page 433\)](#)
- [Use Override files with Konvoy Image Builder \(see page 451\)](#)

### 6.11.3 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](#)<sup>321</sup> should be met before beginning.

#### Section Contents

- [Create a Custom AMI \(see page 414\)](#)
- [AWS Air-gapped AMI \(see page 418\)](#)
- [Add Custom Tags to Image \(see page 419\)](#)

#### 6.11.3.1 Create a Custom AMI

##### 6.11.3.1.1 Learn how to build a custom AMI for use with DKP

This procedure describes how to use the [Konvoy Image Builder \(see page 411\)](#) (KIB) to create a [Cluster API](#)<sup>322</sup> compliant Amazon Machine Image (AMI). 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. The KIB uses [variable overrides \(see page 451\)](#) 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](#)<sup>323</sup> to take advantage of enhanced cluster operations, and to explore the [Advanced AWS Install \(see page 167\)](#) topics for more options.

##### 6.11.3.1.2 Prerequisites

Before you begin, you must:

---

<sup>321</sup> <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>


<sup>322</sup> <https://cluster-api.sigs.k8s.io/>

<sup>323</sup> <https://d2iq.atlassian.net/wiki/pages/resumedraft.action?draftId=143891417>

- Check the [DKP Supported Kubernetes Versions](#) (see page 1056) and download the [KIB](#) (see page 0) bundle (prefixed with `konvoy-image-bundle`) for your OS. Do not use the release prefixed with `konvoy-image-builder`.
- Create a working `Docker` setup.
- Ensure you have met the minimal set of permissions from the [AWS Image Builder Book](#)<sup>324</sup>.

### 6.11.3.1.3 Extract AMI 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](#)<sup>325</sup>. The following variables must be set using your credentials including [required IAM](#) (see page 157):

```
export AWS_ACCESS_KEY_ID
export AWS_SECRET_ACCESS_KEY
export AWS_DEFAULT_REGION
```

- Ensure you have an [override file](#) (see page 451) to configure specific attributes of your AMI file.

### 6.11.3.1.4 Build the Image

Depending on which [version of DKP](#) (see page 411) 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](#)<sup>326</sup> information from AWS.

#### 6.11.3.1.4.1 Execute the following to begin image creation:

Run the `konvoy-image` command to build and validate the image.

```
konvoy-image build images/ami/centos-79.yaml
```

<sup>324</sup> <https://image-builder.sigs.k8s.io/capi/providers/aws.html#required-permissions-to-build-the-aws-amis>

<sup>325</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-envvars.html>

<sup>326</sup> <https://cluster-api-aws.sigs.k8s.io/topics/images/built-amis.html>

By default it builds in the `us-west-2` region. to specify another region set the `--region` flag:

```
./konvoy-image build --region us-east-1 images/ami/centos-79.yaml
```

When the command is complete the `ami` id is printed and written to `./manifest.json`.

### 6.11.3.1.5 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=$(whoami)-aws-cluster --ami ami-0123456789
```

### 6.11.3.1.6 Launch a DKP Cluster with Custom AMI Lookup

By default `konvoy-image` 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=$(whoami)-aws-cluster --ami-format "konvoy-ami-{{.BaseOS}}-?{{.K8sVersion}}-*" --ami-base-os centos-7 --ami-owner 123456789012
```

Using custom source AMIsWhen using KIB **for** building machine images to Amazon, the **default** source AMIs that we provide are based on looking up an AMI based on the owner, and a filter **for** that operating system and version.You can view an example of that with the provided `centos-79.yaml` snippet below:`yaml`download\_images: truepacker: ami\_filter\_name: "CentOS 7.9.2009 x86\_64" 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, or perhaps you want to provide a custom AMI **for** whatever reason you need.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, [CentOS also provides an image](https://wiki.centos.org/Cloud/AWS) on the [AWS marketplace](https://aws.amazon.com/marketplace/pp/prodview-foff247vr2zfw) which you can subscribe to for free.Once you select the source AMI that you want, you can declare that when running your build command:`bash`./konvoy-image build path/to/ami/centos-79.yaml --source-ami ami-0123456789``Alternatively, if you want to add it to your yaml file, or are making your own file, you can do that as well.You just need to add that AMI ID into the `source_ami` in the yaml file:`yaml`download\_images: truepacker: 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:``bashkonvoy-image build path/to/ami/centos-79.yaml``
```

### 6.11.3.1.7 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:

```
yaml
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, [CentOS](https://wiki.centos.org/Cloud/AWS)<sup>327</sup> also provides an image on the [AWS marketplace](https://aws.amazon.com/marketplace)<sup>328</sup> or you can select other 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 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:

```
yaml
download_images: true

packer:
 ami_filter_name: ""
 ami_filter_owners: ""
 distribution: "CentOS"
 distribution_version: "7.9"
 source_ami: "ami-123456789"
```

<sup>327</sup> <https://wiki.centos.org/Cloud/AWS>

<sup>328</sup> <https://aws.amazon.com/marketplace/pp/prodview-foff247vr2zfw>

```
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 path/to/ami/centos-79.yaml
```


### 6.11.3.1.8 Related Information

For information on related topics or procedures, refer to the following:

- [Creating GPU enabled on-premises configurations \(see page 214\)](#)

## 6.11.3.2 AWS Air-gapped AMI

### 6.11.3.2.1 Create an AMI using Konvoy Image Builder (KIB) for use in an air-gapped cluster

 The default AWS image is not recommended for use in production. We suggest using Konvoy Image Builder to [AWS Air-gapped AMI](#)<sup>329</sup> to take advantage of enhanced cluster operations, and to explore the [Install AWS Air-Gapped \(see page 202\)](#) topics for more options.

Using [KIB \(see page 411\)](#), you can build an AMI without requiring access to the internet by providing an additional `--override` flag.

1. Assuming you have [downloaded \(see page 100\)](#) `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2/kib
```

2. Follow the instructions to [build an AMI \(see page 414\)](#) and set the override `--overrides overrides/offline.yaml` flag.

After you complete these steps, you can [seed your docker registry \(see page 204\)](#).

<sup>329</sup> <https://d2iq.atlassian.net/wiki/pages/resumedraft.action?draftId=144117789>

### 6.11.3.3 Add Custom Tags to Image

As certain cloud environments enforce the tagging of resources, DKP platform users have inquired about whether it 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-ebs",
 "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-7-1.24.6-1662575013
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 144) for AWS and finish building your clusters.

## 6.11.4 KIB for Azure

### 6.11.4.1 Learn how to build a custom Azure Image for use with DKP

This procedure describes how to use the [KIB](#) (see page 420) to create a [Cluster API](#)<sup>330</sup> 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 [Konvoy Image Builder](#) (see page 411) 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 226).

### 6.11.4.2 Prerequisites

Before you begin, you must:


---

<sup>330</sup> <https://cluster-api.sigs.k8s.io/>


- Check the [DKP Supported Kubernetes Version](#) (see page 1056) and download the [KIB](#) (see page 411) bundle (prefixed with `konvoy-image-bundle`) for your OS. Do not use the release prefixed with `konvoy-image-builder`.
- Create a working `Docker` setup.

### 6.11.4.3 Extract the 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 run, `konvoy-image` `bind` mounts the current working directory ( `$(PWD)` ) into the container to be used.

### 6.11.4.4 Configure Azure Prerequisites

 If you have already followed the [Azure Prerequisites](#) (see page 226) 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 226) 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@azuresphere.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 451) to configure specific attributes of your Azure image.

### 6.11.4.5 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 `./manifest.json`. You should then specify this image id when creating the cluster.

### 6.11.4.6 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 \(see page 235\)](#) for specific consumption of image commands.

The SKU and Image Name will default to the values found in the image YAML.

## 6.11.5 KIB with GCP

This procedure describes how to use the [Konvoy Image Builder](#) (see page 411) (KIB) to create a [Cluster API](#)<sup>331</sup> 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 451) 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 382) section of the documentation.

### 6.11.5.1 Prerequisites

Before you begin, you must:

- Check the [supported DKP version](#) (see page 1056) and download the [KIB](#) (see page 412) bundle (prefixed with `konvoy-image-bundle`) for your OS. Do not use the release prefixed with `konvoy-image-builder`.
- Create a working `Docker` setup.
- On Debian-based Linux distributions, install a version of the [cri-tools](#)<sup>332</sup> 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](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.

<sup>331</sup> <https://cluster-api.sigs.k8s.io/>

<sup>332</sup> <https://github.com/kubernetes-sigs/cri-tools>



## 6.11.5.2 GCP Prerequisites

- If you are creating your image on either a non-GCP instance or one that does not have the [required roles](#) (see page 382):
  - (option 1) Create a service account using the following `gcloud` commands:

```
export GCP_PROJECT=<your GCP project ID>
export SERVICE_ACCOUNT_USER=<some new service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud iam service-accounts create "${SERVICE_ACCOUNT_USER}" --project=${GCP_PROJECT}
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/compute.instanceAdmin.v1
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/iam.serviceAccountUser
gcloud iam service-accounts keys create $GOOGLE_APPLICATION_CREDENTIALS --iam-account="${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 SERVICE_ACCOUNT_USER=<existing service account user>
export GOOGLE_APPLICATION_CREDENTIALS="$HOME/.gcloud/credentials.json"

gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/compute.instanceAdmin.v1
gcloud projects add-iam-policy-binding ${GCP_PROJECT} --
member="serviceAccount:${SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com" --role=roles/iam.serviceAccountUser
gcloud iam service-accounts keys create $GOOGLE_APPLICATION_CREDENTIALS --iam-account="${SERVICE_ACCOUNT_USER}@${GCP_PROJECT}.iam.gserviceaccount.com"
```

### 6.11.5.3 Create a Network (optional)

Building an image requires a [Network](#)<sup>333</sup> 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
instance on the network using port 22." --direction=INGRESS --priority=65534 --
source-ranges=0.0.0.0/0 --action=ALLOW --rules=tcp:22
```

### 6.11.5.4 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.
```

<sup>333</sup> <https://cloud.google.com/vpc/docs/vpc>

```

==> 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
```

With your KIB image now created, you can now move onto [Bootstrap GCP](#) (see page 388) and set up your [Cluster API](#)<sup>334</sup> (CAPI) controllers, or run [GCP Quick Start](#) (see page 58) to create a cluster with little customization.

## 6.11.6 KIB for GPU

### Create a GPU supported OS image using Konvoy Image Builder

Using the [Konvoy Image Builder](#) (see page 411), you can build an image that has support to use NVIDIA GPU hardware to support GPU workloads.

If the [NVIDIA runfile](#)<sup>335</sup> 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

```

 DKP supported [NVIDIA driver](#)<sup>336</sup> version is 470.x.

To build an image for use on GPU enabled hardware, perform the following steps.

334 <https://cluster-api.sigs.k8s.io/>

335 <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

336 <https://www.nvidia.com/Download/Find.aspx>

1. In your `overrides/nvidia.yaml` file, add the following to enable GPU builds. You can also access and use the overrides [repo](#)<sup>337</sup> or in the documentation under [NVIDIA GPU Overrides](#) (see page 453) or [Offline NVIDIA Override](#) (see page 453).

```
gpu:
 type:
 - nvidia
```

2. Build your image using the following Konvoy Image Builder commands, 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-west-2 --instance-type=p2.xlarge --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 --instance-type=p2.xlarge --overrides override-source-ami.yaml images/ami/<Your OS>.yaml
```



**NOTE:** Ensure that an AMI file is available for your OS selection.

When the command is complete the `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.

```
dkp create cluster aws --cluster-name=$(whoami)-aws-cluster --region us-west-2 --ami <ami>
```

<sup>337</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

For [GPU Steps in Pre-provisioned](#) (see page 0) section of the documentation to use the `overrides/nvidia.yaml`.

### 6.11.6.1 Verification

To verify that the NVIDIA driver is working, connect to the node and execute this command:

```
nvidia-smi
```

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. |
|====+=====+====+=====+=====+=====+=====+=====+
| 0 Tesla K80 Off | 00000000:00:1E:0 Off | 0 |
| N/A 35C P0 73W / 149W | 0MiB / 11441MiB | 99% Default |
+-----+-----+-----+-----+-----+-----+
|
| Processes:
| GPU GI CI PID Type Process name GPU Memory
| ID ID | Usage |
+-----+-----+-----+-----+-----+-----+
| No running processes found
+-----+-----+-----+-----+-----+-----+

```

Additional helpful information can be found in the [NVIDIA Device Plug-in](#)<sup>338</sup> for Kubernetes instructions and the [Installation Guide of Supported Platforms](#)<sup>339</sup>.

See also: [NVIDIA documentation](#)<sup>340</sup>

338 <https://github.com/NVIDIA/k8s-device-plugin/blob/master/README.md>

339 <https://docs.nvidia.com/datacenter/cloud-native/container-toolkit/install-guide.html>

340 [https://nvidia.custhelp.com/app/answers/detail/a\\_id/131/kw/driver%20installation%20docs/related/1](https://nvidia.custhelp.com/app/answers/detail/a_id/131/kw/driver%20installation%20docs/related/1)

## 6.11.7 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](#)<sup>341</sup> and [vSphere Client](#)<sup>342</sup> documentation for details.

- [Create a Base OS Image](#) (see page 430)
- [Create a vSphere Virtual Machine Template](#) (see page 431)

### 6.11.7.1 Create a 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 411)(KIB) to create a VM template to configure Kubernetes nodes by the DKP vSphere provider.

#### 6.11.7.1.1 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 for packer. This user should have administrator privileges. It is possible to configure a custom user and password when building the OS image, however, it requires the Konvoy Image Builder (KIB) configuration to be [overridden](#) (see page 451).

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.
- DKP advises not to use usernames and passwords for security reasons, but instead to use private and public keys. If that is not possible, passwords should be generated and a minimum of 20 characters long.

##### 6.11.7.1.1.1 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 116).

---

341 <https://www.vmware.com/products/vcenter-server.html>

342 <https://docs.vmware.com/en/VMware-vSphere/index.html>

- 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 specify a disk size when you [create a cluster](#) (see page 338) (see the flags of the [create cluster vsphere](#) (see page 946) 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 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 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.

The next step is to [create a vSphere VM template](#) (see page 431) that contains the CAPI and Kubernetes objects.

Refer to the [vCenter](#)<sup>343</sup> and [vSphere Client](#)<sup>344</sup> documentation for details.


## 6.11.7.2 Create a vSphere Virtual Machine Template

### 6.11.7.2.1 Prerequisites

- Users need to create a [base OS image](#) (see page 430) in vSphere before starting this procedure.
- Build image with Konvoy Image Builder (KIB)

<sup>343</sup> <https://www.vmware.com/products/vcenter-server.html>

<sup>344</sup> <https://docs.vmware.com/en/VMware-vSphere/index.html>

-  If using GPU, prepare an OS image for your NVIDIA GPU nodepool, using the Konvoy Image Builder instructions here: [KIB for GPU \(see page 427\)](#)

### 6.11.7.2.2 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.

```

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: ""
 datacenter: ""
 datastore: ""
 folder: ""
 insecure_connection: "false"
 network: ""
 resource_pool: ""
 template: "base-rhel-7.9"
 vsphere_guest_os_type: "rhel7_64Guest"
 guest_os_type: "rhel7-64"
 # goss params
 distribution: "RHEL"
 distribution_version: "7.9"
```



```
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 # if 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>
```

The Konvoy Image Builder uses the values in `image.yaml` and the input base OS image to create a vSphere template that contains the required artifacts needed to create a Kubernetes cluster. Give the file a suitable name using this suggested naming convention: `creator-ova-vsphere-OS-ver-k8sver-unique_identifier`. As an example, the filename you create might resemble `dkp-ova-vsphere-rhel-79-1.24.6-1646938922`.

NOTE: To build an image in a GPU environment, use the `overrides/nvidia.yaml` and see the referenced link [KIB for GPU \(see page 427\)](#).

5. DKP creates the new vSphere template directly on the vCenter server.
6. Next steps are to deploy a DKP cluster using your vSphere template.

## 6.11.8 Konvoy Image Builder CLI

- [konvoy-image build \(see page 434\)](#)
- [konvoy-image completion \(see page 439\)](#)
- [konvoy-image generate-docs \(see page 443\)](#)
- [konvoy-image generate \(see page 444\)](#)
- [konvoy-image provision \(see page 448\)](#)
- [konvoy-image upload \(see page 448\)](#)
- [konvoy-image validate \(see page 450\)](#)
- [konvoy-image version \(see page 450\)](#)

### 6.11.8.1 Other resources:

- Main [Github KIB CLI](#)<sup>345</sup>
- [DKP CLI \(see page 919\)](#)

<sup>345</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/docs/cli>

## 6.11.8.2 konvoy-image build

build and provision images

### 6.11.8.2.1 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]
```

### 6.11.8.2.2 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
 --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

```

### 6.11.8.2.3 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)

```

### 6.11.8.2.4 SEE ALSO

- [konvoy-image build aws](#) (see page 435) - build and provision aws images
- [konvoy-image build azure](#) (see page 436) - build and provision azure images
- [konvoy-image build gcp](#) (see page 438) - build and provision gcp images

### 6.11.8.2.5 konvoy-image build aws

build and provision aws images

```
konvoy-image build aws <image.yaml> [flags]
```

#### 6.11.8.2.5.1 Examples

```
aws --region us-west-2 --source-ami=ami-12345abcdef images/ami/centos-79.yaml
```

#### 6.11.8.2.5.2 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

```

### 6.11.8.2.5.3 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)

```

### 6.11.8.2.5.4 SEE ALSO

- [konvoy-image build](#) (see page 434) - build and provision images

### 6.11.8.2.6 konvoy-image build azure

build and provision azure images

```
konvoy-image build azure <image.yaml> [flags]
```

#### 6.11.8.2.6.1 Examples

```
azure --location westus2 --subscription-id <sub_id> images/azure/centos-79.yaml
```

#### 6.11.8.2.6.2 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
--location string the location in which to build the
image (default "westus2")
--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-group string 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
--work-dir string path to custom work directory generated
by the generate command

```

### 6.11.8.2.6.3 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)

```

#### 6.11.8.2.6.4 SEE ALSO

- [konvoy-image build](#) (see page 434) - build and provision images

#### 6.11.8.2.7 konvoy-image build gcp

build and provision gcp images

```
konvoy-image build gcp <image.yaml> [flags]
```

##### 6.11.8.2.7.1 Examples

```
gcp ... images/gcp/centos-79.yaml
```

##### 6.11.8.2.7.2 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

```

##### 6.11.8.2.7.3 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)
```

#### 6.11.8.2.7.4 SEE ALSO

- [konvoy-image build](#) (see page 434) - build and provision images

### 6.11.8.3 konvoy-image completion

Generate the autocompletion script for the specified shell

#### 6.11.8.3.1 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.

#### 6.11.8.3.2 Options

```
-h, --help help for completion
```

#### 6.11.8.3.3 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)
```

#### 6.11.8.3.4 SEE ALSO

- [konvoy-image completion bash](#) (see page 439) - Generate the autocompletion script for bash
- [konvoy-image completion fish](#) (see page 441) - Generate the autocompletion script for fish
- [konvoy-image completion powershell](#) (see page 441) - Generate the autocompletion script for powershell
- [konvoy-image completion zsh](#) (see page 442) - Generate the autocompletion script for zsh

#### 6.11.8.3.5 konvoy-image completion bash

Generate the autocompletion script for bash

##### 6.11.8.3.5.1 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
```

#### 6.11.8.3.5.2 Options

```
-h, --help help for bash
--no-descriptions disable completion descriptions
```

#### 6.11.8.3.5.3 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)
```

#### 6.11.8.3.5.4 SEE ALSO

- [konvoy-image completion](#) (see page 439) - Generate the autocompletion script for the specified shell



### 6.11.8.3.6 konvoy-image completion fish

Generate the autocompletion script for fish

#### 6.11.8.3.6.1 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]
```

#### 6.11.8.3.6.2 Options

```
-h, --help help for fish
--no-descriptions disable completion descriptions
```

#### 6.11.8.3.6.3 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)
```

#### 6.11.8.3.6.4 SEE ALSO

- [konvoy-image completion](#) (see page 439) - Generate the autocompletion script for the specified shell

### 6.11.8.3.7 konvoy-image completion powershell

Generate the autocompletion script for powershell

### 6.11.8.3.7.1 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]
```

### 6.11.8.3.7.2 Options

```
-h, --help help for powershell
--no-descriptions disable completion descriptions
```

### 6.11.8.3.7.3 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)
```

### 6.11.8.3.7.4 SEE ALSO

- [konvoy-image completion](#) (see page 439) - Generate the autocompletion script for the specified shell

## 6.11.8.3.8 konvoy-image completion zsh

Generate the autocompletion script for zsh

### 6.11.8.3.8.1 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); compdef _konvoy-image konvoy-image
```

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]
```

#### 6.11.8.3.8.2 Options

```
-h, --help help for zsh
--no-descriptions disable completion descriptions
```

#### 6.11.8.3.8.3 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)
```

#### 6.11.8.3.8.4 SEE ALSO

- [konvoy-image completion](#) (see page 439) - Generate the autocompletion script for the specified shell

### 6.11.8.4 konvoy-image generate-docs

generate docs in path

```
konvoy-image generate-docs <PATH> [flags]
```

#### 6.11.8.4.1 Examples

```
generate-docs /tmp/docs
```

#### 6.11.8.4.2 Options

```
-h, --help help for generate-docs
```

#### 6.11.8.4.3 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)
```

### 6.11.8.5 konvoy-image generate

generate files relating to building images

#### 6.11.8.5.1 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]
```

#### 6.11.8.5.2 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

```

### 6.11.8.5.3 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)

```

### 6.11.8.5.4 SEE ALSO

- [konvoy-image generate aws](#) (see page 445) - generate files relating to building aws images
- [konvoy-image generate azure](#) (see page 446) - generate files relating to building azure images

### 6.11.8.5.5 konvoy-image generate aws

generate files relating to building aws images

```
konvoy-image generate aws <image.yaml> [flags]
```

#### 6.11.8.5.5.1 Examples

```
aws --region us-west-2 --source-ami=ami-12345abcdef images/ami/centos-79.yaml
```

### 6.11.8.5.2 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

```

### 6.11.8.5.3 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)

```

### 6.11.8.5.4 SEE ALSO

- [konvoy-image generate](#) (see page 444) - generate files relating to building images

### 6.11.8.5.6 konvoy-image generate azure

generate files relating to building azure images

```
konvoy-image generate azure <image.yaml> [flags]
```

### 6.11.8.5.6.1 Examples

```
azure --location westus2 --subscription-id <sub_id> images/azure/centos-79.yaml
```

### 6.11.8.5.6.2 Options

|                                                           |                                                 |
|-----------------------------------------------------------|-------------------------------------------------|
| <code>--client-id</code> string                           | the client id to use <b>for</b> the build       |
| <code>--cloud-endpoint</code> string                      | Azure cloud endpoint. Which can be one          |
| of [Public USGovernment China] ( <b>default</b> "Public") |                                                 |
| <code>--containerd-version</code> string                  | the version of containerd to install            |
| <code>--extra-vars</code> strings                         | flag passed Ansible's extra-vars                |
| <code>--gallery-image-locations</code> stringArray        | a list of locations to publish the              |
| image ( <b>default</b> [westus])                          |                                                 |
| <code>--gallery-image-name</code> string                  | the gallery image name to publish the           |
| image to                                                  |                                                 |
| <code>--gallery-image-offer</code> string                 | the gallery image offer to set ( <b>default</b> |
| "dkp")                                                    |                                                 |
| <code>--gallery-image-publisher</code> string             | the gallery image publisher to set              |
| ( <b>default</b> "dkp")                                   |                                                 |
| <code>--gallery-image-sku</code> string                   | the gallery image sku to set                    |
| <code>--gallery-name</code> string                        | the gallery name to publish the image           |
| in ( <b>default</b> "dkp")                                |                                                 |
| <code>-h, --help</code>                                   | help <b>for</b> azure                           |
| <code>--instance-type</code> string                       | the Instance Type to use <b>for</b> the build   |
| ( <b>default</b> "Standard_D2s_v3")                       |                                                 |
| <code>--kubernetes-version</code> string                  | The version of kubernetes to install.           |
| Example: 1.21.6                                           |                                                 |
| <code>--location</code> string                            | the location in which to build the              |
| image ( <b>default</b> "westus2")                         |                                                 |
| <code>--overrides</code> strings                          | a comma separated list of override YAML         |
| files                                                     |                                                 |
| <code>--resource-group</code> string                      | the resource group to create the image          |
| in ( <b>default</b> "dkp")                                |                                                 |
| <code>--subscription-id</code> string                     | the subscription id to use <b>for</b> the       |
| build                                                     |                                                 |
| <code>--tenant-id</code> string                           | the tenant id to use <b>for</b> the build       |

### 6.11.8.5.6.3 Options inherited from parent commands

|                          |                                                   |
|--------------------------|---------------------------------------------------|
| <code>--color</code>     | enable color output ( <b>default</b> true)        |
| <code>-v, --v int</code> | select verbosity level, should be between 0 and 6 |
| <code>--verbose</code>   | enable debug level logging (same as --v 5)        |

#### 6.11.8.5.6.4 SEE ALSO

- [konvoy-image generate](#) (see page 444) - generate files relating to building images

#### 6.11.8.6 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]
```

##### 6.11.8.6.1 Examples

```
provision --inventory-file inventory.yaml
```

##### 6.11.8.6.2 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

```

##### 6.11.8.6.3 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)

```

#### 6.11.8.7 konvoy-image upload

Upload one of [artifacts]



### 6.11.8.7.1 Options

```
-h, --help help for upload
```

### 6.11.8.7.2 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)
```

### 6.11.8.7.3 SEE ALSO

- [konvoy-image upload artifacts \(see page 449\)](#) - upload offline artifacts to hosts defined in inventory-file

### 6.11.8.7.4 konvoy-image upload artifacts

upload offline artifacts to hosts defined in inventory-file

```
konvoy-image upload artifacts [flags]
```

#### 6.11.8.7.4.1 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
```

#### 6.11.8.7.4.2 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)

```

#### 6.11.8.7.4.3 SEE ALSO

- [konvoy-image upload](#) (see page 448) - Upload one of [artifacts]

### 6.11.8.8 konvoy-image validate

validate existing infrastructure

```
konvoy-image validate [flags]
```

#### 6.11.8.8.1 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")

```

#### 6.11.8.8.2 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)

```

### 6.11.8.9 konvoy-image version

Version for konvoy-image

```
konvoy-image version [flags]
```

### 6.11.8.9.1 Options

```
-h, --help help for version
```

### 6.11.8.9.2 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)
```

## 6.11.9 Use Override files with Konvoy Image Builder

### 6.11.9.1 Learn how to use override files with Konvoy Image builder

[Konvoy Image Builder](#) (see page 411) uses YAML `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 451) files provided by Konvoy Image Builder for you to use for specific purposes.
- [Custom override files](#) (see page 454) you create to use with Konvoy Image Builder.

### 6.11.9.2 Default Override Files

#### 6.11.9.2.1 Learn how to use override files with DKP

DKP comes with default override files:

- [FIPS Override Files](#) (see page 452)
- [Offline Override File](#) (see page 452)
- [Nvidia GPU Override File](#) (see page 453)
- [Offline Nvidia Override file](#) (see page 453)
- [Oracle Redhat Linux Kernel Override File](#) (see page 453)

You can find these override files in the [Konvoy Image Builder repo](#)<sup>346</sup>.

### 6.11.9.2.2 FIPS Override Files

You can find these override files in the [Konvoy Image Builder repo](#)<sup>347</sup>.

```

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"
```

#### 6.11.9.2.2.1 Offline FIPS Override File

```
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"
```

#### 6.11.9.2.3 Offline Override File

You can find these override files in the [Konvoy Image Builder repo](#)<sup>348</sup>.

```
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 }}"
```

<sup>346</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

<sup>347</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

<sup>348</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

```

pip_packages_local_bundle_file: "{{ playbook_dir }}/../artifacts/pip-packages.tar.gz"
images_local_bundle_dir: "{{ playbook_dir }}/../artifacts/images"

```

For [GPU Offline Override File](#) (see page 453), you can use `nvidia-runfile` flag for GPU support if you have downloaded the [runfile installer](#)<sup>349</sup>.

#### 6.11.9.2.4 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](#)<sup>350</sup>.

```

gpu:
 types:
 - nvidia
build_name_extra: "-nvidia"

```

#### 6.11.9.2.5 Offline Nvidia 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](#)<sup>351</sup>.

```

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"

```

#### 6.11.9.2.6 Oracle Redhat Linux Kernel Override File

You can find these override files in the [Konvoy Image Builder repo](#)<sup>352</sup>.

```

oracle_kernel: RHCK

```

<sup>349</sup> <https://docs.nvidia.com/datacenter/tesla/tesla-installation-notes/index.html#runfile>

<sup>350</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

<sup>351</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

<sup>352</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/overrides>

### 6.11.9.3 Custom Override Files

You can specify customization of your KIB images through the use of override files, which are used to specify alternate package libraries, Docker image repos, and other customizations. Some example custom override files are described in the following topics:

- [Base Image Override Files](#) (see page 454)
- [Container Image Override Files](#) (see page 457)
- [HTTP Proxy Override Files](#) (see page 459)
- [Pre-Provisioned Override Files](#) (see page 459)
- [Private Registry in Air-gapped Override](#) (see page 459)

#### 6.11.9.3.1 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. Therefore, specifying a vSphere was not necessary.

Although there are several parameters specified by default in the Packer templates for each provider, it is possible to override the default values.

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
```

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.

#### 6.11.9.3.1.1 Example 1:

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. Replace infrastructure provider in the command below before running it (`aws`, `azure`, `vsphere`):

```
./konvoy-image build aws images/ami/centos-7.yaml --overrides override-source-ami.yaml
```

#### 6.11.9.3.1.2 Example 2:

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](#)<sup>353</sup>
- [Azure Packer template](#)<sup>354</sup>

<sup>353</sup> <https://github.com/mesosphere/konvoy-image-builder/blob/v1.24.2/pkg/packer/manifests/aws/packer.json.tmpl#L2>

<sup>354</sup> <https://github.com/mesosphere/konvoy-image-builder/blob/v1.24.2/pkg/packer/manifests/azure/packer.json.tmpl#L2>

- [GCP Packer template](#)<sup>355</sup>
- [vSphere Packer template](#)<sup>356</sup>

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

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

A vSphere example would be the current base image description at `images/vsphere/rhel-79.yaml` which is similar to the following:

---

<sup>355</sup> <https://github.com/mesosphere/konvoy-image-builder/blob/v1.24.2/pkg/packer/manifests/gcp/packer.json.tpl#L2>

<sup>356</sup> <https://github.com/mesosphere/konvoy-image-builder/blob/v1.24.2/pkg/packer/manifests/vsphere/packer.json.tpl#L2>



```

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 93) for details.

### 6.11.9.3.2 Container Image Override Files

Your machine image uses a container image. The Ansible playbooks pull a minimal set of [container images](#)<sup>357</sup> for use. You can add or delete additional images by specifying an `override` file for the `extra_images` variable. Konvoy requires several additional images be present. Create a new override file and specify the following `extra_images` :

```

Example override-images.yaml

extra_images:
 - docker.io/bitnami/kubectl:1.24.6
 - docker.io/calico/apiserver:v3.24.1
 - docker.io/calico/cni:v3.24.1
 - docker.io/calico/csi:v3.24.1
 - docker.io/calico/ctl:v3.24.1
 - docker.io/calico/kube-controllers:v3.24.1
 - docker.io/calico/node-driver-registrar:v3.24.1
 - docker.io/calico/node:v3.24.1
 - docker.io/calico/pod2daemon-flexvol:v3.24.1
 - docker.io/calico/typha:v3.24.1
 - docker.io/library/busybox:1

```

<sup>357</sup> <https://github.com/mesosphere/konvoy-image-builder/tree/main/ansible/roles/images>

- docker.io/mesosphere/capd-manager:v1.2.4-d2iq.0
- docker.io/mesosphere/cluster-api-controller:v1.2.4-d2iq.0
- docker.io/mesosphere/cluster-api-preprovisioned-controller:v0.12.1
- docker.io/mesosphere/dkp-diagnostics-node-collector:v0.5.1
- docker.io/mesosphere/kfips:v0.2.0
- docker.io/mesosphere/konvoy-image-builder:v1.24.1
- docker.io/mesosphere/kube-proxy:v1.24.6\_fips.0
- docker.io/mesosphere/kubeadm-bootstrap-controller:v1.2.4-d2iq.0
- docker.io/mesosphere/kubeadm-control-plane-controller:v1.2.4-d2iq.0
- docker.io/mesosphere/pause-busybox:3.2
- docker.io/plndr/kube-vip:v0.3.7
- gcr.io/cloud-provider-vsphere/cpi/release/manager:v1.23.1
- gcr.io/cloud-provider-vsphere/csi/release/driver:v2.5.2
- gcr.io/cloud-provider-vsphere/csi/release/syncer:v2.5.2
- gcr.io/cluster-api-provider-vsphere/release/manager:v1.3.1
- gcr.io/k8s-staging-sig-storage/snapshot-controller:v5.0.1
- ghcr.io/kube-vip/kube-vip:v0.3.9
- k8s.gcr.io/cloud-provider-gcp/gcp-compute-persistent-disk-csi-driver:v1.7.2
- k8s.gcr.io/coredns/coredns:v1.8.4
- k8s.gcr.io/kube-proxy:v1.24.6
- k8s.gcr.io/nfd/node-feature-discovery:v0.11.1-minimal
- k8s.gcr.io/sig-storage/csi-attacher:v3.4.0
- k8s.gcr.io/sig-storage/csi-node-driver-registrar:v2.5.0
- k8s.gcr.io/sig-storage/csi-provisioner:v3.1.0
- k8s.gcr.io/sig-storage/csi-resizer:v1.4.0
- k8s.gcr.io/sig-storage/csi-snapshotter:v4.0.1
- k8s.gcr.io/sig-storage/csi-snapshotter:v5.0.1
- k8s.gcr.io/sig-storage/livenessprobe:v2.6.0
- mcr.microsoft.com/oss/azure/aad-pod-identity/nmi:v1.8.8
- mcr.microsoft.com/oss/kubernetes-csi/azuredisk-csi:v1.19.0
- mcr.microsoft.com/oss/kubernetes-csi/csi-attacher:v3.4.0
- mcr.microsoft.com/oss/kubernetes-csi/csi-node-driver-registrar:v2.5.0
- mcr.microsoft.com/oss/kubernetes-csi/csi-provisioner:v3.1.0
- mcr.microsoft.com/oss/kubernetes-csi/csi-resizer:v1.4.0
- mcr.microsoft.com/oss/kubernetes-csi/csi-snapshotter:v5.0.1
- mcr.microsoft.com/oss/kubernetes-csi/livenessprobe:v2.6.0
- public.ecr.aws/ebs-csi-driver/aws-ebs-csi-driver:v1.8.0
- public.ecr.aws/eks-distro/kubernetes-csi/external-attacher:v3.1.0-eks-1-18-13
- public.ecr.aws/eks-distro/kubernetes-csi/external-provisioner:v2.1.1-eks-1-18-13
- public.ecr.aws/eks-distro/kubernetes-csi/external-resizer:v1.1.0-eks-1-18-13
- public.ecr.aws/eks-distro/kubernetes-csi/external-snapshotter/csi-snapshotter:v5.0.1-eks-1-22-7
- public.ecr.aws/eks-distro/kubernetes-csi/livenessprobe:v2.2.0-eks-1-18-13
- public.ecr.aws/eks-distro/kubernetes-csi/node-driver-registrar:v2.1.0-eks-1-18-13
- quay.io/external\_storage/local-volume-provisioner:v2.4.0
- quay.io/jetstack/cert-manager-cainjector:v1.9.1
- quay.io/jetstack/cert-manager-controller:v1.9.1
- quay.io/jetstack/cert-manager-webhook:v1.9.1
- quay.io/metallb/controller:v0.12.1
- quay.io/metallb/speaker:v0.12.1
- quay.io/tigera/operator:v1.28.1
- registry.k8s.io/cluster-api-aws/cluster-api-aws-controller:v1.5.1
- registry.k8s.io/cluster-api-azure/cluster-api-azure-controller:v1.5.2

- `us.gcr.io/k8s-artifacts-prod/autoscaling/cluster-autoscaler:v1.23.1`
- `us.gcr.io/k8s-artifacts-prod/cluster-api-gcp/cluster-api-gcp-controller:v1.1.0`

### 6.11.9.3.3 HTTP Proxy Override Files

An HTTP proxy configuration can be used when creating your AMI image. The Ansible playbooks will 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.

The `dkp` command can be used to configure the `KubeadmConfigTemplate` object to create this file on bootup of the image with values supplied during the `dkp` invocation. This enables using different proxy settings for image creation and runtime.

### 6.11.9.3.4 Pre-Provisioned Override Files

Pre-provisioned environments require certain override files to work properly. Those override files must also have a secret `secret` that includes all of the overrides you wish to provide in one file. 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: ""

epel_centos_7_rpm: https://my-rpm-repository.org/epel/epel-release-latest-7.noarch.rpm
```

### 6.11.9.3.5 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"]
```

## 6.12 GPU for Konvoy

With DKP 2.4, the nodes with NVIDIA GPUs are configured with `nvidia-gpu-operator` ([Overview – NVIDIA Cloud Native Technologies documentation](#)<sup>358</sup>) 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 473) or [KIB for GPU](#) (see page 427). The remainder of [GPU information](#) (see page 811) is found in Day 2 Operations under Kommander component section of documentation.

## 6.13 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. 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.

### 6.13.1 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](#)<sup>359</sup> 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](#)<sup>360</sup>.

<sup>358</sup> <https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/overview.html>

<sup>359</sup> <https://csrc.nist.gov/CSRC/media/projects/cryptographic-module-validation-program/documents/security-policies/140sp3702.pdf>

<sup>360</sup> [https://access.redhat.com/documentation/en-us/red\\_hat\\_enterprise\\_linux/7/html/security\\_guide/chap-federal\\_standards\\_and\\_regulations](https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/7/html/security_guide/chap-federal_standards_and_regulations)

- 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.

## 6.13.2 Infrastructure Requirements for FIPS-140-2 Mode

To ensure proper operations in FIPS mode, be sure that your environment meets these requirements:

### 6.13.2.1 Supported Operating Systems

Supported Operating Systems for FIPS mode are Red Hat Enterprise Linux and CentOS. See the [Supported Operating Systems](#) (see page 93) for details on the tested and supported versions.

## 6.13.3 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.

### 6.13.3.1 Supported FIPS Builds

| Component  | Repository                                                                                                                                                                                                             | Version        |
|------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------|
| Kubernetes | <a href="https://hub.docker.com/layers/mesosphere/kube-apiserver/v1.24.6_fips.0/images/sha256-5196d46d810bdcfb02737bb4dcfa29632db041c705575c41439d672032bcfeb3?context=explore">docker.io/mesosphere<sup>361</sup></a> | v1.24.6+fips.0 |
| etcd       | <a href="https://hub.docker.com/layers/mesosphere/etcd/v3.4.13_fips.0/images/sha256-01dbd850ff17782e41b2fedd67e60d8f277b8d45dca2d38831f971223cd893d0?context=explore">docker.io/mesosphere<sup>362</sup></a>           | 3.4.13+fips.0  |

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:

<sup>361</sup> [https://hub.docker.com/layers/mesosphere/kube-apiserver/v1.24.6\\_fips.0/images/sha256-5196d46d810bdcfb02737bb4dcfa29632db041c705575c41439d672032bcfeb3?context=explore](https://hub.docker.com/layers/mesosphere/kube-apiserver/v1.24.6_fips.0/images/sha256-5196d46d810bdcfb02737bb4dcfa29632db041c705575c41439d672032bcfeb3?context=explore)

<sup>362</sup> [https://hub.docker.com/layers/mesosphere/etcd/v3.4.13\\_fips.0/images/sha256-01dbd850ff17782e41b2fedd67e60d8f277b8d45dca2d38831f971223cd893d0?context=explore](https://hub.docker.com/layers/mesosphere/etcd/v3.4.13_fips.0/images/sha256-01dbd850ff17782e41b2fedd67e60d8f277b8d45dca2d38831f971223cd893d0?context=explore)

```
dkp create cluster aws --cluster-name myFipsCluster \
--ami=ami-03dcaa75d45aca36f \
--kubernetes-version=v1.24.6+fips.0 \
--kubernetes-image-repository=docker.io/mesosphere \
--etcd-image-repository=docker.io/mesosphere \
--etcd-version=3.4.13+fips.0
```

## 6.13.4 Create FIPS 140 Images

### 6.13.4.1 Use [Konvoy Image Builder](#) to create images with FIPS-compliant binaries

### 6.13.4.2 Create FIPS-140 images

KIB can produce images containing FIPS-140 compliant binaries. Use the `fips.yaml` [override file](#) (see page 451) provided with the image bundles.

For example, this command produces a FIPS-compliant image on RHEL 8.4:

```
konvoy-image build --overrides overrides/fips.yaml images/ami/rhel-84.yaml
```

#### 6.13.4.2.1 Pre-provisioned infrastructure

If you are targeting a [pre-provisioned infrastructure](#) (see page 288), you can create a FIPS-compliant cluster by doing the following:

1. Create a [bootstrap cluster](#) (see page 296)
2. Create a secret on the bootstrap cluster with the contents from `fips.yaml` [override file](#)<sup>363</sup> 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=
```

## 6.13.5 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.

<sup>363</sup> <https://github.com/mesosphere/konvoy-image-builder/blob/main/overrides/fips.yaml>

### 6.13.5.1 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)
```

#### 6.13.5.1.1 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
```

### 6.13.5.1.2 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
```

### 6.13.5.2 Download 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.

#### 6.13.5.2.1 DKP Version 2.4.0

| Operating System version | Kubernetes version | containerd version | Signature File URL                        |
|--------------------------|--------------------|--------------------|-------------------------------------------|
| CentOS 7.9               | v1.24.6            | 1.4.13             | <a href="#">CentOS 7.9</a> <sup>364</sup> |
| Oracle 7.9               | v1.24.6            | 1.4.13             | <a href="#">Oracle 7.9</a> <sup>365</sup> |
| RHEL 7.9                 | v1.24.6            | 1.4.13             | <a href="#">RHEL 7.9</a> <sup>366</sup>   |
| RHEL 8.4                 | v1.24.6            | 1.4.13             | <a href="#">RHEL 8.4</a> <sup>367</sup>   |
| RHEL 8.6                 | v1.24.6            | 1.4.13             | <a href="#">RHEL 8.6</a> <sup>368</sup>   |

<sup>364</sup> <https://downloads.d2iq.com/dkp/fips/v2.4.1/manifest-centos-79.json.asc>

<sup>365</sup> <https://downloads.d2iq.com/dkp/fips/v2.4.1/manifest-oracle-79.json.asc>

<sup>366</sup> <https://downloads.d2iq.com/dkp/fips/v2.4.1/manifest-rhel-79.json.asc>

<sup>367</sup> <https://downloads.d2iq.com/dkp/fips/v2.4.1/manifest-rhel-84.json.asc>

<sup>368</sup> <https://downloads.d2iq.com/dkp/fips/v2.4.1/manifest-rhel-86.json.asc>



## 6.13.6 FIPS 140 Mode Performance Impact

### 6.13.6.1 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](#)<sup>369</sup> 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.

## 6.14 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.

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.

<sup>369</sup> <https://github.com/golang/go/issues/21525>

This command performs the following actions:

- Creates a local bootstrap cluster
- Moves controllers to it
- Deletes the management cluster
- Deletes the local bootstrap cluster

## 6.15 Configure the Control Plane


### 6.15.1 Prerequisites

Before you begin, make sure you have created a bootstrap cluster from the respective [Infrastructure Providers](#) (see page 128) 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.

#### 6.15.1.1 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: 3.4.13-0
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
 for i in $(ls /run/kubeadm/ | grep 'kubeadm.yaml\|kubeadm-join-config.yaml');
do
 cat <<EOF>> "/run/kubeadm//${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: v1.24.6

```

Now a user can configure the fields below for the log backend. The log backend will write audit events to a file in [JSON<sup>370</sup>](#) 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<sup>371</sup>](#) 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
```

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
```

### 6.15.1.2 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
```

<sup>370</sup> <https://jsonlines.org/>

<sup>371</sup> <https://kubernetes.io/docs/tasks/debug/debug-cluster/audit/#log-backend>



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 773)

## 6.16 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](#)<sup>372</sup>) 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.

### 6.16.1 Prerequisites:

- Deploy [Pod Disruption Budget](#)<sup>373</sup> (PDB)
- [Konvoy Image Builder](#) (see page 411) (KIB)

---


<sup>372</sup> <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

<sup>373</sup> <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

## 6.16.2 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 411).

## 7 Advanced Kommander Configuration

The Kommander component of DKP can be customized for different environments and installation types. See the following sections for more information on custom settings:

- [Kommander Install Configuration](#) (see page 475)
- [Install Kommander in an Air-gapped Environment](#) (see page 504)
- [Install Kommander in a Non-air-gapped Environment](#) (see page 513)
- [Install Kommander in a Pre-provisioned Environment](#) (see page 515)
- [Install Kommander on a Small Environment](#) (see page 518)
- [Verify Kommander Installation](#) (see page 521)
- [Log in to the UI with Kommander](#) (see page 522)
- [Helm and Chart Bundle CLI Commands](#) (see page 524)

### 7.1 Kommander Install Configuration

You can configure the Kommander component of DKP during the initial installation, and also post-installation using the DKP CLI.



Review the [Management cluster application requirements](#) (see page 117) and [Workspace platform application requirements](#) (see page 118) to ensure that your cluster has sufficient resources.

#### 7.1.1 Initialize a Configuration File

To begin configuring Kommander, run the following command to initialize a default configuration file:

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

#### 7.1.2 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
```

### 7.1.2.1 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
 ...
```

### 7.1.2.2 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
...
```

### 7.1.3 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 518) documentation for more information.

### 7.1.4 Install with Configuration File

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.

**Info** An alternative to using the `--kubeconfig=<cluster-config>` flag is to initialize the KUBECONFIG environment variable. You can do this by running `export KUBECONFIG=<cluster-config>`. Setting your KUBECONFIG (either by flag or by environment variable) ensures that Kommander is installed on the workload cluster.

```
dkp install kommander --installer-config kommander.yaml --kubeconfig=<cluster-
kubeconfig>
```

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 [Configure HTTP Proxy](#) (see page 494).

**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 for the deployment of applications.

### 7.1.5 Verify Installation

After the Konvoy cluster is built and Kommander has been installed, you will want to [verify your installation of Kommander](#) (see page 521) in that section of documentation.

Then you will be able to [log in in to Kommander UI](#). (see page 522)

## 7.1.6 Custom Domains and Certificates

### Configure a custom domain during installation

DKP supports configuring a custom domain name and certificate for accessing the UI and other platform services.

This section provides instructions and examples on how to configure the **DKP installation** to add a customized domain and certificate on your **Essential cluster** or your **Management cluster**.



If you want to customize the domain and certificate after installing DKP or on an **Attached** or **Managed** cluster, this sections does not apply to your use case. For this, refer to [Configure Custom Domains or Custom Certificates](#) (see page 725) .

#### Section contents:

- [Why to set up a Custom Domain or Certificate?](#) (see page 478)
- [Certificate Authority \(CA\) Specifics](#) (see page 479)
- [Configure your Custom Domain and Certificate](#) (see page 480)
- [Verification and Troubleshooting for Custom Certificates](#) (see page 488)

### 7.1.6.1 Why to set up a Custom Domain or Certificate?

#### 7.1.6.1.1 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 [\(2.4\)Configure a Custom Domain without a Custom Certificate](#) (see page 487).

#### 7.1.6.1.2 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 480). This is only possible for your Management/Essential cluster.

- Update your cluster's current domain and certificate configuration as part of your [cluster's Day 2 operations](#) (see page 723). You can do this for any cluster type 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.

### 7.1.6.1.3 Next Topic:

[\(2.4\)Certificate Authority \(CA\) Specifics](#) (see page 479)

### 7.1.6.2 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</i> values                                      |
|-----------------------|------------------------------------------------|------------------------------------------------------------------------|
| 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 [\(2.4\)Configure your Custom Domain and Certificate](#) (see page 480).

### 7.1.6.2.1 Next Topic:

[\(2.4\)Configure your Custom Domain and Certificate](#) (see page 480)

### 7.1.6.3 Configure your Custom Domain and Certificate

This page contains instructions on how to set up custom certificates for your cluster during the installation of DKP. This allows most browsers to validate the certificate for the cluster when users try to log into the operations portal.



Refer to [Certificate Authority \(CA\) Specifics](#) (see page 479) for more information on values that are specific to your Certificate Authority or CA.

There are three main options:

**I want to use an automatically-generated certificate with ACME and require basic configuration\***

#### 7.1.6.3.1 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 475), 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:

```
[...]
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 477).




\*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.)**

### 7.1.6.3.2 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 \(see page 484\)](#) and [Advanced Configuration: Important Concepts \(see page 486\)](#).

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
```

**Note:** 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.

- 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
```

**Note:** 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.

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 an [external-dns service](#) (see page 807). This way, the `external-dns` will take care of pointing the DNS record to the ingress of the cluster automatically.
 

**Note:** You can also create a DNS record manually, that maps your domain name or IP address to the cluster ingress. In this case, finish installing Kommander and then manually create the DNS record pointing 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 475), 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 \(see page 477\)](#).

## I have a manually-generated certificate

### 7.1.6.3.3 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.

#### 7.1.6.3.3.1 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)

#### 7.1.6.3.3.2 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 \(see page 475\)](#), 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 477).

#### 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` .

#### 7.1.6.3.4 Next Step:

[Verification and Troubleshooting for Custom Certificates](#) (see page 488)

#### 7.1.6.3.5 Related Topics:


- [Advanced Configuration: ClusterIssuer](#) (see page 484)
- [Advanced Configuration: Important Concepts](#) (see page 486)
- [Configure a Custom Domain without a Custom Certificate](#) (see page 487)

#### 7.1.6.3.6 Advanced Configuration: ClusterIssuer

When you enable ACME (see page 480), by default DKP generates an ACME-supported certificate with an HTTP01 solver that is provided by Let's Encrypt.

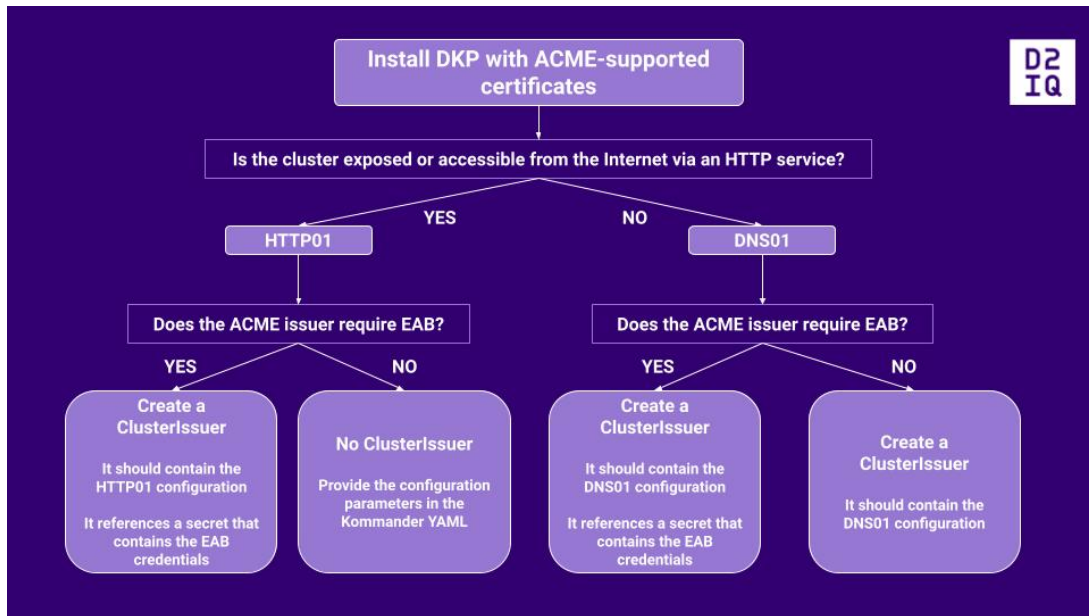
You can also set up an [advanced configuration for a Custom Domain and Certificate](#) (see page 481). In these cases, the custom configuration cannot be done completely via the `installer config` file, but must be specified further in a `ClusterIssuer` .

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.

### 7.1.6.3.6.1 When do You Need a ClusterIssuer?

The configuration of the `ClusterIssuer` resource depends on your DKP landscape:



### 7.1.6.3.6.2 How do You Configure a ClusterIssuer?

The following image describes the configurable fields of a `ClusterIssuer` :

```

apiVersion: cert-manager.io/v1
kind: ClusterIssuer
metadata:
 name: acme
spec:
 acme:
 email: <email>
 server: <server-acme>
 privateKeySecretRef:
 name: example-issuer-account-key
 externalAccountBinding:
 keyID: <my-keyID>
 keySecretRef:
 name: <eab-secret-name>
 key: <secret-key>
 solvers:
 - dns01: [...]

```

Base Configuration

EAB Configuration (Optional)

Solvers Configuration - DNS01, HTTP01 (Optional)

For more information on the available options, refer to the [ACME section in the cert-manager documentation](https://cert-manager.io/docs/configuration/acme/)<sup>374</sup>.

<sup>374</sup> <https://cert-manager.io/docs/configuration/acme/>

### 7.1.6.3.6.3 Examples:

Refer to [Configure your Custom Domain and Certificate](#) (see page 480) 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, modify the `ingress` in the `KommanderCluster` object as shown in the [Custom domains and certificates configuration](#) (see page 723) section.

### 7.1.6.3.6.4 Related topics:

[Why to set up a Custom Domain or Certificate?](#) (see page 478)

[Configure your Custom Domain and Certificate](#) (see page 480)

[Advanced Configuration: Important Concepts](#) (see page 486)

### 7.1.6.3.7 Advanced Configuration: Important Concepts

If you set up an [advanced configuration of your custom domain](#) (see page 481), ensure you understand the following concepts.

#### 7.1.6.3.7.1 IssuerRef, ClusterIssuerRef or certificateSecretRef?

If you use a certificate issued and managed automatically by `cert-manager`, you need an *Issuer* or *Cluster Issuer* 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 *secret* that you reference in your `KommanderCluster` resource.

#### 7.1.6.3.7.2 Management, Managed or Attached cluster? Location of the KommanderCluster and Issuer objects

In the **Management** or **Essential** cluster, both the `KommanderCluster` and *issuer* objects are stored on the same cluster. The *issuer* can be referenced as an `Issuer`, `ClusterIssuer` or `certificateSecretRef`.

In **Managed** and **Attached** clusters, the `KommanderCluster` object is stored on the Management cluster. The `Issuer`, `ClusterIssuer` or `certificateSecretRef` is stored on the Managed or Attached cluster.

For more information on `ClusterIssuer` objects, refer to [Advanced Configuration: ClusterIssuer](#) (see page 484).

#### 7.1.6.3.7.3 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.

#### 7.1.6.3.7.4 Related topics:

[Why to set up a Custom Domain or Certificate?](#) (see page 478)

[Configure your Custom Domain and Certificate](#) (see page 480)

[Advanced Configuration: ClusterIssuer](#) (see page 484)

#### 7.1.6.3.8 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 480).

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 475), 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 807)) 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 480) page contains more details and examples on how and when to set up the DNS record.

#### 7.1.6.3.8.1 Related topics:

[Why to set up a Custom Domain or Certificate?](#) (see page 478)

[Configure your Custom Domain and Certificate](#) (see page 480)

[Advanced Configuration: Important Concepts](#) (see page 486)

[Advanced Configuration: ClusterIssuer](#) (see page 484)

### 7.1.6.4 Verification and Troubleshooting for Custom Certificates

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`.

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
[...]
```

#### 7.1.6.4.1 Related topics:

[Why to set up a Custom Domain or Certificate? \(see page 478\)](#)

[Configure your Custom Domain and Certificate \(see page 480\)](#)

[Advanced Configuration: Important Concepts \(see page 486\)](#)

[Advanced Configuration: ClusterIssuer \(see page 484\)](#)


## 7.1.7 DKP Kommander Configuration Reference

### 7.1.7.1 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 475).

| Parameter                                                       | Description                                                                            | Default Value                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
|-----------------------------------------------------------------|----------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p>apps<br/>(<a href="#">AppConfig list</a> (see page 493))</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: &lt;shortened for brevity&gt;   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.                                                                                                                                                                                                                                                               |               |
| <a href="#">ingressCertificate</a> (see page 494) | 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.                                                                                                                                                                                                                                                                                                          |               |
| <a href="#">airgapped</a> (see page 494)          | Specifies parameters for an airgapped environment.                                                                                                                                                                                                                                                                                                   |               |
| catalog<br><b>Enterprise</b>                      | Specifies parameters for installing default catalog repositories.<br><br><div style="border: 1px solid #800080; padding: 10px; margin-top: 10px;"> <p> For additional information, see <a href="#">Configure an Enterprise catalog</a> (see page 502).</p> </div> |               |

## 7.1.7.2 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>HelRelease</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>HelRelease</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>                                                                                                                                                                                                                                                      |               |

### 7.1.7.3 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. |               |

### 7.1.7.4 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.            |               |

#### 7.1.7.4.1 Next Step:

[Configure HTTP Proxy \(see page 494\)](#)

## 7.1.8 Configure HTTP Proxy

### Configure HTTP proxy for the Kommander cluster(s)

Kommander supports environments where access to the Internet is restricted, and must be made through an HTTP/HTTPS proxy.

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

### 7.1.8.1 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.

### 7.1.8.2 Enable Gatekeeper

Gatekeeper acts as a [Kubernetes mutating webhook](#)<sup>375</sup>. 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 > install.yaml
```

2. Append this `apps` section to the `install.yaml` file with the following values to enable Gatekeeper and configure it to add HTTP proxy settings to the pods.

<sup>375</sup> <https://kubernetes.io/docs/reference/access-authn-authz/admission-controllers/#mutatingadmissionwebhook>

**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:
 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"
```

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
```

### 7.1.8.3 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.

#### 7.1.8.4 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.

### 7.1.8.5 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 ./install.yaml
```

### 7.1.8.6 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 once the namespace is created:

```
kubectl label namespace <NAMESPACE> "gatekeeper.d2iq.com/mutate=pod-proxy"
```

### 7.1.8.7 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"
```

### 7.1.8.8 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.

### 7.1.8.9 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.

### 7.1.8.10 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](#)<sup>376</sup> for more details.

#### 7.1.8.10.1 Next Steps:

Now select your environment, and finish your Kommander Installation in one of the following:

[Install Kommander in an Air-gapped Environment](#) (see page 504)

[Install Kommander in a Non-air-gapped Environment](#) (see page 513)

[Install Kommander on a Small Environment](#) (see page 518)

---

<sup>376</sup> <https://kubernetes.io/docs/tasks/inject-data-application/define-environment-variable-container/#define-an-environment-variable-for-a-container>

## 7.1.9 External Load Balancer

### 7.1.9.1 Load Balancing for External Traffic in DKP

DKP includes a load balancing solution for the [supported cloud infrastructure providers \(see page 93\)](#) and for pre-provisioned environments. For more information, see [Load Balancing for external traffic \(see page 45\)](#) 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.

### 7.1.9.2 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 \(see page 45\)](#), 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 your Custom Domain and Certificate \(see page 480\)](#).
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 \(see page 477\).](#)

### 7.1.9.3 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.

### 7.1.10 Configure an Enterprise catalog

Enterprise

#### Configure an Enterprise catalog for DKP

DKP supports configuring default catalogs for clusters with Enterprise license.

### 7.1.10.1 Configure a Default Enterprise Catalog

To configure DKP to use a default catalog repository, create the following YAML file:

```
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: v2.4.1
```

Use this configuration when installing or reconfiguring DKP by passing it to the `dkp install kommander` command:

To ensure DKP is installed on the workload cluster, use the `--kubeconfig=cluster_name.conf` flag.

```
dkp install kommander --installer-config <config_file.yaml>
```

**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 for the deployment of applications.

### 7.1.10.2 Configure an Enterprise Catalog after installing or upgrading DKP

When configuring the catalog repository post-upgrade, run `dkp install kommander --init > install.yaml` and update it accordingly with any custom configuration. This ensures you are using the proper default configuration values for the new DKP version.

### 7.1.10.3 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 |

#### 7.1.10.4 Next Step:

[Custom Domains and Certificates](#) (see page 478)

## 7.2 Install Kommander in an Air-gapped Environment

### How to install Kommander using an Air-gapped installation

Depending on your configuration, there are three different ways you can install DKP to an air-gapped environment.

**Ensure you follow the correct procedure for your configuration type below:**

DKP Essential:

- [Install Kommander in an air-gapped environment](#) (see page 506)
- [Install air-gapped Kommander with DKP Insights](#) (see page 507)

DKP Enterprise:

- [Install air-gapped Kommander with DKP Catalog Applications](#) (see page 509)
- [Install air-gapped Kommander with DKP Insights and DKP Catalog Applications](#) (see page 511)

### 7.2.1 Prerequisites

Before installing, ensure you have:

- A Docker registry containing all the necessary Docker installation images, including the Kommander images. See below for how to push the necessary images to this registry.
- [Download](#) (see page 100) the Complete DKP Air-gapped Bundle for this release (i.e. `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`)
- Connectivity with clusters attaching to the management cluster:



- Both management and attached clusters must be able to connect to the Docker 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.
- All the prerequisites covered in [air-gapped DKP installation](#) (see page 0).
- Sufficient resources on your cluster to run Kommander. Review the [Management cluster application requirements](#) (see page 117) and [Workspace platform application requirements](#) (see page 118) for application requirements.
- A load balancer to route external traffic which is provided by your cloud provider. For on-premises deployments, you must configure MetalLB. See [Load Balancing](#) (see page 806) topic for more details.

## 7.2.2 Load the Docker Images into Your Docker Registry

Follow these steps:

1. Assuming you have downloaded `dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz`, extract the tarball to a local directory:

```
tar -xzf dkp-air-gapped-bundle_v2.4.2_linux_amd64.tar.gz && cd dkp-v2.4.2
```

2. See the `NOTICES.txt` file for 3rd party software attributions in the `container-images/` directory.
3. Set an environment variable with your registry address:

```
export DOCKER_REGISTRY_ADDRESS=<registry-address>:<registry-port>
```

4. Run the following command to load the air-gapped image bundle into your private Docker registry:

```
./dkp push image-bundle --image-bundle ./container-images/kommander-image-bundle-v2.4.2.tar --to-registry $DOCKER_REGISTRY_ADDRESS
```

It may 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 Docker registry.

### 7.2.2.1 Next Step

Choose the correct air-gapped environment from the pages in this section.

- [Air-gapped Essential](#) (see page 506)
- [Air-gapped Enterprise](#) (see page 509)

## 7.2.3 Air-gapped 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.

Depending on your configuration, there are two different ways you can install DKP to an air-gapped environment in DKP Essential:

- [Kommander in an Air-gapped Environment](#) (see page 506)
- [Kommander in Air-gapped with DKP Insights](#) (see page 507)

For more information regarding Essential vs. Enterprise, refer to the [Licenses](#) (see page 101) section of the documentation.

### 7.2.3.1 Kommander in an Air-gapped Environment

Follow these steps to install the Kommander component of DKP Essential in an air-gapped environment with basic setup.

#### 7.2.3.1.1 Installation

1. Create the [configuration file](#) (see page 475) by running `dkp install kommander --init --airgapped > install.yaml` for the air-gapped deployment.
2. In the same file, if you are installing Kommander in an AWS VPC, set the Traefik annotation to create an internal facing ELB by setting the following:

```
apps:
 traefik:
 enabled: true
 values: |
 service:
 annotations:
 service.beta.kubernetes.io/aws-load-balancer-internal: "true"
```

3. To install Kommander in your air-gapped environment using the above configuration file, enter the following command:

```
dkp install kommander --installer-config ./install.yaml \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.4.2.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.4.2.tar.gz
```

**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 for the deployment of applications.

#### 7.2.3.1.1.1 Next Steps:

- Verify Installation
- Log in to the Kommander UI

#### Verify Installation

Once the Konvoy cluster is built and Kommander has been installed, you will want to [verify your installation of Kommander](#) (see page 521) in that section of documentation.

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

#### Log in to the UI

Then you will be able to [log in to Kommander UI](#). (see page 522)

### 7.2.3.2 Kommander in Air-gapped with DKP Insights

Follow these steps to install the Kommander component of DKP Essential in an air-gapped environment with DKP Insights.

#### 7.2.3.2.1 Install Kommander

Follow these steps:

1. Create the [configuration file](#) (see page 475) by running `dkp install kommander --init --airgapped > install.yaml` for the air-gapped deployment.
2. In `install.yaml`, if you are installing Kommander in an AWS VPC, set the Traefik annotation to create an internal facing ELB by setting the following:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
```

```

kind: Installation
apps:
 ...
 traefik:
 values: |
 service:
 annotations:
 service.beta.kubernetes.io/aws-load-balancer-internal: "true"
 ...

```

3. In `install.yaml`, enable DKP Insights by setting the following:

```

apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
 ...
 dkp-insights-management:
 enabled: true
 ...
catalog:
 repositories:
 - name: insights-catalog-applications
 labels:
 kommander.d2iq.io/workspace-default-catalog-repository: "true"
 kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
 path: ./application-repositories/dkp-insights-v2.4.2.tar.gz
 ...

```

Install DKP with Insights enabled by running:

```

dkp install kommander --installer-config ./install.yaml \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.4.2.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.4.2.tar.gz \
--charts-bundle ./application-charts/dkp-insights-charts-bundle-v2.4.2.tar.gz\

```



**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 for the deployment of applications.

### 7.2.3.2.2 Next Steps:

- Verify Installation
- Log in to the Kommander UI

#### 7.2.3.2.2.1 Verify Installation

After the Konvoy cluster is built and Kommander has been installed, you will want to [verify your installation of Kommander](#) (see page 521) in that section of documentation.



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

#### 7.2.3.2.2.2 Log in to the UI

Then you will be able to [log in to Kommander UI](#). (see page 522)

See [Deployment of Catalog Applications in Workspaces](#) (see page 600) to deploy Insights Engine.

## 7.2.4 Air-gapped Enterprise

### Enterprise

DKP Enterprise is a multi-cluster solution centered around a management cluster that manages multiple attached or managed Kubernetes clusters via a centralized management dashboard.

Depending on your configuration, there are two different ways you can install DKP to an air-gapped environment for DKP Enterprise:

- [Install Air-gapped Kommander with DKP Catalog Applications](#) (see page 509)
- [Install Air-gapped Kommander with DKP Insights and DKP Catalog Applications](#) (see page 511)

For more information regarding Essential vs. Enterprise, refer to the [Licenses](#) (see page 101) section of the documentation.

### 7.2.4.1 Install Air-gapped Kommander with DKP Catalog Applications

#### Enterprise

Follow these steps to install the Kommander component of DKP Enterprise in an air-gapped environment with Catalog Applications.

### 7.2.4.1.1 Install Kommander

Follow these steps:

1. Create the [configuration file](#) (see page 475) by running `dkp install kommander --init --airgapped > install.yaml` for the air-gapped deployment.
2. In `install.yaml`, if you are installing Kommander in an AWS VPC, set the Traefik annotation to create an internal facing ELB by setting the following:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
 ...
 traefik:
 values: |
 service:
 annotations:
 service.beta.kubernetes.io/aws-load-balancer-internal: "true"
 ...
```

3. In `install.yaml`, enable DKP Catalog Applications by setting the following:

```
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-v2.4.2.tar.gz
```

4. To install the Kommander component of DKP in your air-gapped environment using the above configuration file, run the following command:

```
dkp install kommander --installer-config ./install.yaml \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.4.2.tar.gz \
--charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.4.2.tar.gz
\
--charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-
v2.4.2.tar.gz
```

**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 for the deployment of applications.

### 7.2.4.1.2 Next Steps:

- Verify Installation
- Log in to the Kommander UI

#### 7.2.4.1.2.1 Verify Installation

After the Konvoy cluster is built and Kommander has been installed, you will want to [verify your installation of Kommander](#) (see page 521) in that section of documentation.

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

#### 7.2.4.1.2.2 Log in to the UI

Then you will be able to [log in to Kommander UI](#). (see page 522)

See [Deployment of Catalog Applications in Workspaces](#) (see page 600) to deploy the Catalog Applications.

## 7.2.4.2 Install Air-gapped Kommander with DKP Insights and DKP Catalog Applications

### Enterprise

Follow these steps to install the Kommander component of DKP Enterprise in an air-gapped environment with DKP Insights and Catalog Applications.

#### 7.2.4.2.1 Install Kommander

Follow these steps:

1. Create the [configuration file](#) (see page 475) by running `dkp install kommander --init --airgapped > install.yaml` for the air-gapped deployment.

- In `install.yaml`, if you are installing Kommander in an AWS VPC, set the Traefik annotation to create an internal facing ELB by setting the following:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
 ...
 traefik:
 values: |
 service:
 annotations:
 service.beta.kubernetes.io/aws-load-balancer-internal: "true"
 ...
```

- In `install.yaml`, enable DKP Insights and DKP Catalog Applications by setting the following:


```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
apps:
 ...
 dkp-insights-management:
 enabled: true
 ...
catalog:
 repositories:
 - name: insights-catalog-applications
 labels:
 kommander.d2iq.io/workspace-default-catalog-repository: "true"
 kommander.d2iq.io/gitapps-gitrepository-type: "dkp"
 path: ./application-repositories/dkp-insights-v2.4.2.tar.gz
 - 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: ./application-repositories/dkp-catalog-applications-v2.4.2.tar.gz
EOF
```

- Follow the steps on the [Configure an Enterprise catalog](#) (see page 502) page to configure the DKP catalog applications.
- To install the Kommander component of DKP in your air-gapped environment using the above configuration file, run the following command:

```
dkp install kommander --installer-config ./install.yaml \
--kommander-applications-repository ./application-repositories/kommander-
applications-v2.4.2.tar.gz \
--charts-bundle dkp-kommander-charts-bundle-v2.4.2.tar.gz \
```



```
--charts-bundle dkp-catalog-applications-charts-bundle-v2.4.2.tar.gz \
--charts-bundle dkp-insights-charts-bundle-v2.4.2.tar.gz
```


-  **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 for the deployment of applications.

### 7.2.4.2.2 Next Steps:

- Verify Installation
- Log in to the Kommander UI

#### 7.2.4.2.2.1 Verify Installation

After the Konvoy cluster is built and Kommander has been installed, you will want to [verify your installation of Kommander](#) (see page 521) in that section of documentation.

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

#### 7.2.4.2.2.2 Log in to the UI

Then you will be able to [log in to Kommander UI](#). (see page 522)

See [Deployment of Catalog Applications in Workspaces](#) (see page 600) to deploy Insights Engine.

## 7.3 Install Kommander in a Non-air-gapped Environment

To install the Kommander component of DKP in a Non-air-gapped Environment, you will need to ensure you have the prerequisites below met.

### 7.3.1 Prerequisites

Before installing Kommander:

- Ensure you have the version of the CLI that matches the DKP version you want to install.

- You have configured a Konvoy cluster using the [Advanced Konvoy Configuration](#) (see page 128) for infrastructure provider specific instructions on building a cluster-based on environment.
- Review the [Management Cluster Application Requirements](#) (see page 117) and [Workspace Platform Application Defaults and Resource Requirements](#) (see page 118) to ensure that your cluster has sufficient resources.
- Ensure you have a default `StorageClass` configured as shown below.
- 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 478)
  - [Configure custom domains and certificates after Kommander has been installed](#) (see page 723)

### 7.3.2 Configure a Default StorageClass

The cluster where Kommander is installed must have a default `StorageClass` configured. Use the following command to verify one is configured:

```
kubectl get sc
```

The output should look similar to this. Note the `(default)` after the name:

| NAME                      | PROVISIONER     | RECLAIMPOLICY | VOLUMEBINDINGMODE    | ALLOWVOLUMEEXPANSION | AGE |
|---------------------------|-----------------|---------------|----------------------|----------------------|-----|
| ebs-sc ( <b>default</b> ) | ebs.csi.aws.com | Delete        | WaitForFirstConsumer | <b>false</b>         | 41s |

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](#)<sup>377</sup> in the Kubernetes documentation.

### 7.3.3 Install Kommander

To customize your Kommander installation, see the [Kommander Install Configuration](#) (see page 475) for more details.

<sup>377</sup> <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class>

Before running the commands below, ensure that your `kubectl` configuration **references the cluster on which you want to install Kommander**, otherwise it will install on the bootstrap cluster. You can do this by setting the `KUBECONFIG` environment variable [to the appropriate kubeconfig file's location](#)<sup>378</sup>.

**NOTE:** An alternative to initializing the `KUBECONFIG` environment variable as stated earlier is to use the `--kubeconfig=cluster_name.conf` flag. This ensures that Kommander is installed on the correct cluster.

Install Kommander:

```
dkp install kommander
```

**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 for the deployment of applications.

## 7.3.4 Verify Installation

Once the Konvoy cluster is built and Kommander has been installed, you will want to [verify your installation of Kommander](#) (see page 521) in that section of documentation.

### 7.3.4.1 Next Step

[Access the Kommander UI](#) (see page 522)

## 7.4 Install Kommander in a Pre-provisioned Environment

The Kommander installation for pre-provisioned environment consists on the following high-level steps:

- [Pre-provisioned Prerequisites for Kommander](#) (see page 516)
- [Prepare the Kommander Installer Configuration File](#) (see page 516)

---

<sup>378</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

## 7.4.1 Pre-provisioned Prerequisites for Kommander

To install the Kommander component of DKP in a pre-provisioned environment, ensure you meet the following prerequisites:

### 7.4.1.1 Prerequisites

As *part* of the installation of DKP:

- Ensure you meet the general pre-provisioned prerequisites specified in [Pre-provisioned Prerequisites](#) (see page 289), and have configured a Konvoy cluster using [Pre-provisioned Infrastructure](#) (see page 288) instructions.


After installing **Konvoy** and before installing **Kommander**, ensure you have:

- The version of the CLI that matches the DKP version you want to install.
- Sufficient resources on your cluster to run Kommander. Review the [Management Cluster Application Requirements](#) (see page 117) and [Workspace Platform Application Defaults and Resource Requirements](#) (see page 118) for application requirements.
- Ensure you meet the [storage requirements](#) (see page 860).
- A load balancer to route external traffic, which is provided by your cloud provider. For on-premises deployments, you must configure MetalLB. See [Load Balancing](#) (see page 806) topic for more details.
- Ensure you have configured a [default StorageClass](#) (see page 514).
- Ensure you have added at least 40 GB of [raw storage to each of your worker nodes](#) (see page 862) in your cluster.

### 7.4.1.2 Next Step:

[Prepare the Kommander Installer Configuration File](#) (see page 516)

## 7.4.2 Prepare the Kommander Installer Configuration File

 You **MUST** modify the Kommander installer configuration file ( `<kommander.yaml>` ) before installing the Kommander component of DKP.

1. [Initialize a Kommander configuration file](#) (see page 475).
2. Edit the **Kommander installer configuration file** to include configuration overrides for the `rook-ceph-cluster`. DKP default configuration ships ceph with [PVC based storage](#)<sup>379</sup> which requires your CSI provider to support PVC with type `volumeMode: Block`. As this is not possible with the

<sup>379</sup> <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/>

default local static provisioner, you can install ceph in [host storage](#)<sup>380</sup> 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](#)<sup>381</sup>. For general information on the `deviceFilter` value, refer to [Storage Selection Settings](#)<sup>382</sup>.

3. **Optional:** You can add other configuration overrides to your **Kommander installer configuration file**, for example:
  - If you want to **customize your cluster's domain or certificate**, review [Configure custom domains and certificates during the Kommander installation](#) (see page 478).
  - If you require an Enterprise catalog, review [Configure an Enterprise catalog](#) (see page 502).
  - If you require an HTTPS proxy, review [Configure HTTP Proxy](#) (see page 494).
4. Run the following command by replacing the placeholder `<kommander.yaml>` with the name of your **Kommander installer configuration file**:

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

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

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

```
dkp install kommander --installer-config <kommander.yaml> --wait-timeout 1h
```

**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 for the deployment of applications.

## 7.5 Install Kommander on 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 \(see page 117\)](#). Applying an Enterprise license key to the previous installation adds modifications to your environment that can exhaust a small environment's resources.

### 7.5.1 Prerequisites

Ensure you have done the following:

- You have acquired a DKP license.
- You have installed the [Konvoy component \(see page 123\)](#).
- You have reviewed the prerequisite section pertaining to your [air-gapped \(see page 504\)](#), or [networked \(see page 513\)](#) environment.

### 7.5.2 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 575) 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
 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

```

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 [Configure HTTP Proxy \(see page 494\)](#).

**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 for the deployment of applications.

### 7.5.3 Verify Installation

After the Konvoy cluster is built and Kommander has been installed, you will want to [verify your installation of Kommander \(see page 521\)](#) in that section of documentation. 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`.



**NOTE:** An alternative to using the `--kubeconfig=<cluster-kubeconfig>.conf` flag is to initialize the KUBECONFIG environment variable. You can do this by running `export KUBECONFIG=<cluster-kubeconfig>`. Setting your KUBECONFIG (either by flag or by environment variable) ensures that Kommander is installed on the workload cluster.

Then you will be able to [log in in to Kommander UI](#). (see page 522)

## 7.6 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 something resembling this:

```
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

```

## 7.6.1 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}]'

```

## 7.6.2 Next Step

Log in to the [UI with Kommander](#) (see page 522).

## 7.7 Log in to the UI with Kommander

When DKP is installed, a static user account and credentials are created that can be used to access the DKP Dashboard. 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"}}' --kubeconfig=${CLUSTER_NAME}.conf
```

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"}}' --kubeconfig=${CLUSTER_NAME}.conf
```

D2iQ recommends that you only use these static credentials as a backup to credentials configured using an [external identity provider](#) (see page 540).

You can perform the following operations on [Identity Providers](#) (see page 522):

- Create an Identity Provider
- Temporarily Disable an Identity Provider
- Create Groups

**NOTE:** An alternative to using the `--kubeconfig=${CLUSTER_NAME}.conf` flag is to initialize the KUBECONFIG environment variable. You can do this by running `export KUBECONFIG=${CLUSTER_NAME}.conf`. Setting your KUBECONFIG (either by flag or by environment variable) ensures that Kommander is installed on the workload cluster. The below command assumes you've sent your KUBECONFIG.

Always log in with your own identity from external identity providers that provide additional security features like Multi-Factor Authentication. If the need arises to change the password for these credentials, the following shell script can be used to update the secret that contains the static credentials:

```
#!/usr/bin/env bash
set -o errexit pipefail nounset

case "$OSTYPE" in darwin)
 new_password=$(head -c45 /dev/urandom | base64 --break=0)
 ;;
*)
 new_password=$(head -c45 /dev/urandom | base64 --wrap=0)
 ;;
esac

kubectl patch \
 --namespace kommander \
 secret/dkp-credentials \
 --type='json' \
 --patch="[{ \"op\": \"replace\", \"path\": \"/data/password\", \"value\": \"$new_password\" }]"

printf 1>&2 "Password: %s\n" "$new_password"
```

## 7.7.1 Next Step

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 Operations](#) (see page 526) section of the documentation. This section allows you to manage cluster operations and their application workloads to optimize your organization's productivity.

## 7.8 Helm and Chart Bundle CLI Commands

### 7.8.1 Kommander Charts Bundle

The charts bundle is a gzipped tar archive containing Helm charts, which are required during Kommander installation. Create the charts bundle with the Kommander CLI or downloaded along with the DKP CLI. Execute this command to create the charts bundle:

```
dkp create chart-bundle
```

DKP creates `charts-bundle.tar.gz`. Optionally, specify the output using the `-o` parameter:

```
dkp create chart-bundle -o [name of the output file]
```

### 7.8.2 DKP Internal Helm Repository

The DKP charts bundle is pushed to DKP's internal Helm repository. To inspect the contents:

```
dkp get charts
```

Individual charts can be removed using:

```
dkp delete chart [chartName] [chartVersion]
```

It is possible to push new charts as well:

```
dkp push chart [chartTarball]
```

Or push a new bundle:

```
dkp push chart-bundle [chartsTarball]
```

Check the built-in help text for each command for more information.

## 8 Day 2 Operations

Use these sections to manage your DKP environment.

- [Deploy a Sample Application](#) (see page 526)
- [Operations](#) (see page 529)
- [Applications](#) (see page 561)
- [Workspaces](#) (see page 580)
- [Projects](#) (see page 615)
- [Manage Clusters](#) (see page 674)
- [Backup and Restore](#) (see page 730)
- [Logging](#) (see page 751)
- [Security](#) (see page 780)
- [Networking](#) (see page 796)
- [GPUs](#) (see page 811)
- [Monitoring and Alerts](#) (see page 822)
- [DKP Troubleshooting](#) (see page 841)
- [Storage](#) (see page 850)

### 8.1 Deploy a Sample Application

#### 8.1.1 Learn how to deploy a sample application on a DKP cluster

After you have a basic DKP cluster installed and ready to use, you might want to test operations by deploying a simple, sample application. This task is **optional** and is only intended to demonstrate the basic steps for deploying applications in a production environment. If you are configuring the DKP cluster for a production deployment, you can use this section to learn the deployment process. However, deploying applications on a production cluster typically involves more planning and custom configuration than covered in this example.

This tutorial shows how to deploy a simple application that connects to the `redis` service. The sample application used in this tutorial is a condensed form of the Kubernetes sample [guestbook](#)<sup>383</sup> application.

#### 8.1.2 Before You Begin

You must have a [DKP cluster running](#) (see page 128).

Before running the commands below, ensure that your `kubectl` configuration references the DKP cluster on which you want to install the application. You can do this by setting the `KUBECONFIG` environment variable to the appropriate kubeconfig file's location.

---

<sup>383</sup> <https://kubernetes.io/docs/tutorials/stateless-application/guestbook/>

### 8.1.3 Deploy a Sample Application

1. Deploy the Redis pods and service by running the following commands:

```
kubectl apply -f https://k8s.io/examples/application/guestbook/redis-leader-
deployment.yaml
kubectl apply -f https://k8s.io/examples/application/guestbook/redis-leader-
service.yaml
```

2. Deploy Redis followers. The leader deployment created above is a single pod. Adding followers (or replicas) makes it highly available to meet greater traffic demands. You must then setup the guestbook application to communicate with the Redis followers to read the data. To do this, set up another service (the `redis-follower-service.yaml` below). Do this by running the following commands:

```
kubectl apply -f https://k8s.io/examples/application/guestbook/redis-follower-
deployment.yaml
kubectl apply -f https://k8s.io/examples/application/guestbook/redis-follower-
service.yaml
```

3. Deploy the web app frontend by running the following command:

```
kubectl apply -f https://k8s.io/examples/application/guestbook/frontend-
deployment.yaml
```

4. Confirm that there are three frontend replicas running:

```
kubectl get pods -l app=guestbook -l tier=frontend
```

5. Apply the frontend Service:

```
kubectl apply -f https://k8s.io/examples/application/guestbook/frontend-
service.yaml
```

6. Configure the frontend Service to use a cloud load balancer:

```
cat << EOF | kubectl apply -f -
apiVersion: v1
kind: Service
metadata:
 name: frontend
 labels:
```

```

 app: guestbook
 tier: frontend
spec:
 type: LoadBalancer
 ports:
 - port: 80
 selector:
 app: guestbook
 tier: frontend
EOF

```

7. View the frontend service via the LoadBalancer by running the following command to get the IP address for the frontend Service:

```
kubectl get service frontend
```

8. the external IP address, and load the page in your browser to view your guestbook.


The service properties provide the name of the load balancer. You can connect to the application by accessing that load balancer address in your web browser. Because this sample deployment creates a **cloud load balancer**, you should keep in mind that creating the load balancer can take up to a few minutes. You also might experience a slight delay before it is running properly due to DNS propagation and synchronization.

9. Remove the sample application by running the following commands:

```

kubectl delete deployment -l app=redis
kubectl delete service -l app=redis
kubectl delete deployment frontend
kubectl delete service frontend

```

 **WARNING:** This step is **required** because the sample deployment attaches a **cloud provider load balancer** to the DKP cluster. Therefore, you **must delete** the sample application before tearing down the cluster.

## 8.1.4 Related Information

- [Example: Deploying PHP Guestbook application with Redis](#)<sup>384</sup>

<sup>384</sup> <https://kubernetes.io/docs/tutorials/stateless-application/guestbook/>



## 8.2 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.

- [Access Control](#) (see page 529)
- [Identity Providers](#) (see page 540)
- [Infrastructure Providers](#) (see page 546)

### 8.2.1 Access Control

#### 8.2.1.1 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 532) 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.

### 8.2.1.2 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 533) 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

```

### 8.2.1.3 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.

### 8.2.1.3.1 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.

### 8.2.1.3.2 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.

### 8.2.1.3.3 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
```

#### 8.2.1.3.4 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
```

#### 8.2.1.3.5 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.

#### 8.2.1.4 Related Information

- [Project Role Bindings](#) (see page 657)
- [Workspace Role Bindings](#) (see page 614)
- [Kommander RBAC Tutorial](#) (see page 533)
- [RBAC - DKP UI Authorization](#) (see page 533)
- [Kubernetes RBAC Authorization](#)<sup>385</sup>

<sup>385</sup> <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

## 8.2.1.5 Granting Access to Kubernetes and Kommander Resources

### 8.2.1.5.1 Grant access to Kommander and Kubernetes resources using RBAC

#### 8.2.1.5.2 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](#)<sup>386</sup>.

##### 8.2.1.5.2.1 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](#)<sup>387</sup>, 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.

<sup>386</sup> <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

<sup>387</sup> <https://kubernetes.io/docs/reference/access-authn-authz/rbac/#default-roles-and-role-bindings>

### 8.2.1.5.2.2 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
 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.

### 8.2.1.5.2.3 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.

### 8.2.1.5.3 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.

#### 8.2.1.5.3.1 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/<br>dashboard/* | read                |
| kommander-dashboard | dkp-kommander-edit  | /dkp/kommander/<br>dashboard/* | read, write         |
| kommander-dashboard | dkp-kommander-admin | /dkp/kommander/<br>dashboard/* | read, write, delete |

| Dashboard            | Role                                         | Path                                | access              |
|----------------------|----------------------------------------------|-------------------------------------|---------------------|
| 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         |
| 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                |



| Dashboard            | Role                                       | Path                              | access              |
|----------------------|--------------------------------------------|-----------------------------------|---------------------|
| 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 |
| 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  |

| Dashboard | Role                   | Path                                  | access              |
|-----------|------------------------|---------------------------------------|---------------------|
| thanos    | dkp-thanos-query-view  | /dkp/kommander/<br>monitoring/query/* | read                |
| thanos    | dkp-thanos-query-edit  | /dkp/kommander/<br>monitoring/query/* | read, write         |
| thanos    | dkp-thanos-query-admin | /dkp/kommander/<br>monitoring/query/* | read, write, delete |

This section provides a few examples of binding subjects to the default roles defined for the DKP UI endpoints.

### 8.2.1.5.3.2 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 529).

### 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`.

#### 8.2.1.5.4 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.

#### 8.2.1.5.5 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](#)<sup>388</sup>.

## 8.2.2 Identity Providers

### 8.2.2.1 Grant access to users in your organization

DKP supports GitHub, [LDAP](#) (see page 543), 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.

### 8.2.2.2 Prerequisites

To get started with DKP, you must:

- [Log in to the UI](#) (see page 522)

---

<sup>388</sup> <https://kubernetes.io/docs/reference/access-authn-authz/rbac/>

### 8.2.2.2.1 Limit Access

- The [GitHub](#) (see page 541) provider allows you to specify any of the organizations and teams are eligible for access.
- The [LDAP](#) (see page 543) 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.

### 8.2.2.2.2 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.

### 8.2.2.2.3 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.

### 8.2.2.3 Groups

Access control groups are configured in the Groups tab of the Identity Providers page. See [Access Control](#) (see page 529) for an overview of groups in DKP.

### 8.2.2.4 Authorize a Group in Github

Enterprise

**Install GitHub as an identity provider and grant access to all developers.**

#### 8.2.2.4.1 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](#)<sup>389</sup>.
- 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.

#### 8.2.2.4.1.1 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.

#### 8.2.2.4.1.2 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.

#### 8.2.2.4.1.3 Assign the Role to the Developers Group

1. Select the **Cluster Role Bindings** tab, and then select the **Add roles** button for your group.

---

<sup>389</sup> <https://github.com/settings/applications/new>

2. Select “Read Everything” role from the **Roles** drop-down.

Lastly, follow the example in the [Access Control documentation](#) (see page 529) 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.


## 8.2.2.5 External LDAP directory

### 8.2.2.5.1 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.

#### 8.2.2.5.2 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](#)<sup>390</sup> for more details.

The example below configures a DKP cluster to connect to the [Online LDAP Test Server](#)<sup>391</sup> 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
```

<sup>390</sup> <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/ldap.md>

<sup>391</sup> <https://www.forumsys.com/tutorials/integration-how-to/ldap/online-ldap-test-server/>

```

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:
 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
```



### 8.2.2.5.3 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](#) (see page 533).

### 8.2.2.5.4 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.

#### 8.2.2.5.4.1 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.

#### 8.2.2.5.4.2 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).

#### 8.2.2.5.4.3 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=[]"
```


## 8.2.3 Infrastructure Providers

Enterprise

Managing infrastructure providers used by Kommander

Infrastructure providers, like AWS, provide the infrastructure for your Management clusters. To automate their provisioning, Kommander needs authentication keys for your preferred infrastructure provider. You may have many accounts for a single infrastructure provider. Kommander needs authentication keys for your preferred infrastructure provider. You may have many accounts for a single infrastructure provider.

To provision new clusters and manage them, Kommander needs infrastructure provider credentials. Currently, AWS is supported.

 Infrastructure provider credentials are configured in each workspace. The name you assign must be unique across all namespaces in your cluster.

### 8.2.3.1 View and Modify Infrastructure Providers

1. From the top menu bar, select your target workspace.
2. Select **Infrastructure Providers** in the **Administration** section of the sidebar menu.

### 8.2.3.2 AWS


- [Configure AWS Provider with Role Credentials \(see page 547\)](#) (Recommended if using AWS)
- [Configure AWS Provider with Static Credentials \(see page 554\)](#)


### 8.2.3.3 Delete an infrastructure provider

Before deleting an infrastructure provider, Kommander verifies if any existing managed clusters were created using this provider. The infrastructure provider cannot be deleted until all clusters created with the infrastructure provider have been deleted. This ensures Kommander has access to your infrastructure provider to remove all resources created for a managed cluster.

### 8.2.3.4 Configure an AWS Provider with a User Role

**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](#)<sup>392</sup> method with an optional External ID for third party access.

#### 8.2.3.4.1 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",
 "ec2>DeleteRouteTable",
 "ec2>DeleteSecurityGroup",
 "ec2>DeleteSubnet",
 "ec2>DeleteTags",
 "ec2>DeleteVpc",
 "ec2:DescribeAccountAttributes",

```

<sup>392</sup> <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

```

 "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](#)<sup>393</sup> 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": [
 {
 "Sid": "AssumeRoleKommander",
 "Effect": "Allow",
 "Action": "sts:AssumeRole",
 "Resource": "arn:aws:iam::YOURACCOUNTRESTRICTION:role/THEROLEYOUCREATED"
 }
]
}

```

#### 8.2.3.4.2 Create Infrastructure Provider in the UI

1. From the top menu bar, select your target workspace.

<sup>393</sup> <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/iam-roles-for-amazon-ec2.html>

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](#)<sup>394</sup>.
9. Select **Save** to save your provider.

### 8.2.3.5 AWS Static Credentials

#### Enterprise

Configuring an AWS Infrastructure Provider with static credentials

#### 8.2.3.5.1 Configure an AWS Infrastructure Provider with Static Credentials

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.

#### 8.2.3.5.2 Create a New User via CLI Commands

You will need to have the [AWS CLI utility installed](#)<sup>395</sup>. Create a new user via 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
yGroup","ec2>DeleteSubnet","ec2>DeleteTags","ec2>DeleteVpc","ec2:DescribeAccountAttributes","ec2:Des
cribeAddresses","ec2:DescribeAvailabilityZones","ec2:DescribeInstances","ec2:DescribeInternetGateway
s","ec2:DescribeImages","ec2:DescribeNatGateways","ec2:DescribeNetworkInterfaces","ec2:DescribeNetwo
rkInterfaceAttribute","ec2:DescribeRouteTables","ec2:DescribeSecurityGroups","ec2:DescribeSubnets","
ec2:DescribeVpcs","ec2:DescribeVpcAttribute","ec2:DescribeVolumes","ec2:DetachInternetGateway","ec2:
DisassociateRouteTable","ec2:DisassociateAddress","ec2:ModifyInstanceAttribute","ec2:ModifyNetworkIn
terfaceAttribute","ec2:ModifySubnetAttribute","ec2:ReleaseAddress","ec2:RevokeSecurityGroupIngress",
"ec2:RunInstances","ec2:TerminateInstances","tag:GetResources","elasticloadbalancing:AddTags","elast
icloadbalancing:CreateLoadBalancer","elasticloadbalancing:ConfigureHealthCheck","elasticloadbalancin
g>DeleteLoadBalancer","elasticloadbalancing:DescribeLoadBalancers","elasticloadbalancing:DescribeLoa
dBalancerAttributes","elasticloadbalancing:ApplySecurityGroupsToLoadBalancer","elasticloadbalancing:
DescribeTags","elasticloadbalancing:ModifyLoadBalancerAttributes","elasticloadbalancing:RegisterInst
```

<sup>394</sup> [https://docs.aws.amazon.com/IAM/latest/UserGuide/id\\_roles\\_create\\_for-user\\_externalid.html](https://docs.aws.amazon.com/IAM/latest/UserGuide/id_roles_create_for-user_externalid.html)

<sup>395</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-chap-install.html>

```

ancesWithLoadBalancer", "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
```

### 8.2.3.5.3 Using an Existing User

You can use an existing AWS user with [credentials configured](#)<sup>396</sup>. The user must be authorized to create the following resources in the AWS account:

<sup>396</sup> <https://docs.aws.amazon.com/cli/latest/userguide/cli-configure-profiles.html>

- EC2 Instances
- VPC
- Subnets
- Elastic Load Balancer (ELB)
- Internet Gateway
- NAT Gateway
- Elastic Block Storage (EBS) Volumes
- Security Groups
- Route Tables
- IAM Roles

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",
 "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-*"
 }
 }
}
]
}

```



#### 8.2.3.5.4 Fill out the Add Infrastructure Provider Form in the UI

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.

## 8.3 Applications

This section includes information on the applications you can deploy in DKP.

- [AppDeployment resources](#) (see page 561)
- [Manage an Application using the UI](#) (see page 564)
- [Platform Applications](#) (see page 570)

### 8.3.1 AppDeployment resources

**Use AppDeployments to deploy, and customize platform, DKP catalog, and custom applications in your environment**

An `AppDeployment` is a [Custom Resource](#)<sup>397</sup> 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 `HelmsRelease`. `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.

---

<sup>397</sup> <https://kubernetes.io/docs/concepts/extend-kubernetes/api-extension/custom-resources/>

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-40.0.0
 kind: ClusterApp
```

### 8.3.1.1 Customization

### 8.3.1.2 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.

### 8.3.1.3 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 584) 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-40.0.0
 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

```

### 8.3.1.4 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.

#### 8.3.1.4.1 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-40.0.0 host-cluster
kubecost kubecost-0.28.0 host-cluster
[...]

```

#### 8.3.1.4.2 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-40.0.0 host-cluster

```

#### 8.3.1.5 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 118](#)), [project-level](#) (see [page 619](#)), and [per-cluster deployment and customization of workspace applications](#) (see [page 584](#)).

#### 8.3.1.6 More Information

Refer to the [CLI documentation](#) (see [page 919](#)) for more information on how to `create`, or `get` an AppDeployment.

### 8.3.2 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 565](#))
  - [Enable an Application using the UI](#) (see [page 565](#))
  - [Customize an Application using the UI](#) (see [page 566](#))

- [Verify an Application using the UI \(see page 567\)](#)
- [Essential - Manage an Application using the UI \(see page 568\)](#)
  - [Enable an Application using the UI - Essential \(see page 568\)](#)
  - [Customize an Application using the UI - Essential \(see page 569\)](#)
  - [Verify an Application via UI - Essential \(see page 569\)](#)

### 8.3.2.1 Enterprise - Manage an Application using the UI

Enterprise

Gov Advanced



To use the CLI to deploy or uninstall applications, see [Application Deployment \(see page 570\)](#).

- [Enable an Application using the UI \(see page 565\)](#)
- [Customize an Application using the UI \(see page 566\)](#)
- [Verify an Application using the UI \(see page 567\)](#)

#### 8.3.2.1.1 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 575) documentation. Review them carefully prior to customizing to ensure that the applications are deployed successfully.

#### 8.3.2.1.1.1 Next topic:

[Customize an Application using the UI](#) (see page 566)

#### 8.3.2.1.2 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 565).

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**.

#### 8.3.2.1.2.1 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**.

#### 8.3.2.1.2.2 Next topic:

[Verify an App using the UI \(see page 567\)](#)

#### 8.3.2.1.3 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.

### 8.3.2.2 Essential - Manage an Application using the UI

- To use the CLI to deploy or uninstall applications, see [Application Deployment \(see page 570\)](#).

- [Enable an Application using the UI - Essential \(see page 568\)](#)
- [Customize an Application using the UI - Essential \(see page 569\)](#)
- [Verify an Application via UI - Essential \(see page 569\)](#)

#### 8.3.2.2.1 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 \(see page 575\)](#) documentation. Review them carefully prior to customizing to ensure that the applications are deployed successfully.



#### 8.3.2.2.1.1 Next topic:

[Customize an App using the UI - Essential \(see page 569\)](#)

#### 8.3.2.2.2 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 \(see page 568\)](#).

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**.

#### 8.3.2.2.2.1 Next topic:

[Verify an App via UI - Essential \(see page 569\)](#)

#### 8.3.2.2.3 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.

### 8.3.3 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 118) 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.

#### 8.3.3.1 Related Topics

- [Deploy Platform Applications via CLI](#) (see page 570)
- [Cluster-scoped Configuration for Existing AppDeployments](#) (see page 584)
- [Manage an Application using the UI](#) (see page 564)
- [Cluster-scoped Application Configuration via the UI](#) (see page 582)
- [Platform Application Dependencies](#) (see page 575)
- [Workspace Platform Application Defaults and Resource Requirements](#) (see page 118)

#### 8.3.3.2 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.

##### 8.3.3.2.1 Related Topics

- [Available Workspace Platform Applications](#) (see page 118)
- [Customize an existing AppDeployment](#) (see page 562)
- [Cluster-scoped configuration via CLI](#) (see page 584)
- [Cluster-scoped configuration via UI](#) (see page 582)

##### 8.3.3.2.2 Prerequisites

Before you begin, you must have:

- A running cluster with **Kommander installed** (see page 100).
- An **existing Kubernetes cluster attached to Kommander** (see page 678).
- 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 572).

### 8.3.3.2.3 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 561).

1. Obtain the **APP ID** and **Version** of the application from the [applications table in the latest Release Notes](#) (see page 1064). You will need them in the `<APP-ID>-<Version>` format, for example, `istio-1.15.3`.
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-1.15.3 --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`.

#### 8.3.3.2.4 Verify Applications

The applications are now enabled. Connect to the attached cluster and watch the `HelmsReleases` to verify the deployment. In this example, we are checking if `istio` got deployed correctly:

```
kubectll get helmreleases istio -n ${WORKSPACE_NAMESPACE} -w
```

You should eventually see the `HelmsRelease` 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 \(see page 575\)](#) for that table.

### 8.3.3.3 Upgrade Platform Applications

#### 8.3.3.3.1 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>
```

### 8.3.3.3.2 Upgrade Platform Applications from the CLI


The [DKP upgrade \(see page 1024\)](#) process deploys and upgrades Platform applications as a bundle for each cluster or workspace. For the Management Cluster or workspace, DKP upgrade handles all Platform applications; no other steps are necessary to upgrade the Platform application bundle. However, for managed or attached clusters or workspaces, you **MUST** manually upgrade the Platform applications bundle with the following command.

 If you are upgrading your Platform applications as part of the [DKP upgrade \(see page 1024\)](#), 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 \(see page 1056\)](#) of this release, and upgrading Kubernetes is part of the DKP Konvoy upgrade process.

#### 8.3.3.3.2.1 Prerequisites for upgrading clusters with NVIDIA GPU Nodes

 **NOTE:** The information in this subsection is applicable only to clusters with GPU nodes. If your cluster is not running GPU nodes, then skip these steps and proceed to the [Run DKP upgrade command subsection \(see page 574\)](#).

If your attached or managed cluster is running GPU nodes and has deployed the `nvidia` Platform Application, you must run additional steps **prior** to upgrading the workspace. In DKP 2.4, the `nvidia` application is replaced by `nvidia-gpu-operator` during upgrade, and the configuration for the new application must be created prior to the upgrade to ensure that the new application is deployed with the proper configuration, or it may not deploy successfully.

 **NOTE:** The following steps enable the application for all clusters in the workspace and add configuration overrides at the workspace level, which apply to all clusters in the workspace. If only a subset of clusters in the workspace are running GPU nodes, or some clusters require different configurations, refer to [Customize an Application per Cluster \(see page 588\)](#) on how to customize the application with different configuration overrides per cluster.

1. Create the ConfigMap with the necessary workspace-level configuration overrides to [set the correct Toolkit version \(see page 814\)](#). 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: ${WORKSPACE_NAMESPACE}
 name: nvidia-gpu-operator-ws-overrides
data:
 values.yaml: |
 toolkit:
 version: v1.10.0-centos7
EOF
```

2. Create an `AppDeployment` named `nvidia-gpu-operator` in the workspace namespace:

```
cat <<EOF | kubectl apply -f -
apiVersion: apps.kommander.d2iq.io/v1alpha3
kind: AppDeployment
metadata:
 name: nvidia-gpu-operator
 namespace: ${WORKSPACE_NAMESPACE}
spec:
 appRef:
 kind: ClusterApp
 configOverrides:
 name: nvidia-gpu-operator-ws-overrides
EOF
```

3. You may now proceed to the workspace upgrade steps.

### 8.3.3.3 Execute the DKP Upgrade Command

Use this command to upgrade all platform applications in the given 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}"
...
```

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}]'
```

### 8.3.3.4 Platform Application Dependencies

#### 8.3.3.4.1 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 570).

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.

#### 8.3.3.4.2 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](#)<sup>398</sup>: Automates TLS certificate management and issuance.
- [reloader](#)<sup>399</sup>: A controller that watches changes on ConfigMaps and Secrets, and automatically triggers updates on the dependent applications.

<sup>398</sup> <https://cert-manager.io/docs>

<sup>399</sup> <https://github.com/stakater/Reloader>

- [traefik](#)<sup>400</sup>: Provides an HTTP reverse proxy and load balancer. Requires cert-manager and reloader.

| Platform Application | Dependencies           |
|----------------------|------------------------|
| cert-manager         |                        |
| reloader             |                        |
| traefik              | cert-manager, reloader |

### 8.3.3.4.3 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](#)<sup>401</sup>: Open source and multi-platform log processor tool which aims to be a generic Swiss knife for logs processing and distribution.
- [grafana-logging](#)<sup>402</sup>: Logging dashboard used to view logs aggregated to Grafana Loki.
- [grafana-loki](#)<sup>403</sup>: A horizontally-scalable, highly-available, multi-tenant log aggregation system inspired by Prometheus.
- [logging-operator](#)<sup>404</sup>: Automates the deployment and configuration of a Kubernetes logging pipeline.
- [rook-ceph](#)<sup>405</sup> and [rook-ceph-cluster](#)<sup>406</sup>: A Kubernetes-native high performance object store with an S3-compatible API that supports deploying into private and public cloud infrastructures.

| Platform Application | Dependencies      |
|----------------------|-------------------|
| fluent-bit           |                   |
| grafana-logging      | grafana-loki      |
| grafana-loki         | rook-ceph-cluster |
| logging-operator     |                   |

400 <https://traefik.io/>

401 <https://docs.fluentbit.io/manual/>

402 <https://grafana.com/oss/grafana/>


403 <https://grafana.com/oss/loki/>

404 <https://banzaicloud.com/docs/one-eye/logging-operator/>

405 <https://rook.io/docs/rook/v1.10/Helm-Charts/operator-chart/>

406 <https://rook.io/docs/rook/v1.10/Helm-Charts/ceph-cluster-chart/>



| Platform Application | Dependencies                                                                                                                                                                                                                                                                                                                                                                                |
|----------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| rook-ceph            |                                                                                                                                                                                                                                                                                                                                                                                             |
| rook-ceph-cluster    | <ul style="list-style-type: none"> <li>• rook-ceph</li> <li>• kube-prometheus-stack</li> </ul> <div style="border: 1px solid purple; padding: 10px; margin-top: 10px;"> <p> kube-prometheus-stack is an optional dependency. Users can override the config to remove the dependency, as needed.</p> </div> |

#### 8.3.3.4.4 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/)<sup>407</sup>: 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/)<sup>408</sup>: 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)<sup>409</sup>: A stack of applications that collect metrics and provide visualization and alerting capabilities.  
**NOTE:** [Prometheus](https://prometheus.io/)<sup>410</sup>, [Prometheus Alertmanager](https://prometheus.io/docs/alerting/latest/alertmanager)<sup>411</sup> and [Grafana](https://grafana.com/)<sup>412</sup> are included in the bundled installation.
- [nvidia-gpu-operator](https://catalog.ngc.nvidia.com/orgs/nvidia/containers/gpu-operator)<sup>413</sup>: 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)<sup>414</sup>: Provides cluster metrics from Prometheus.

407 <https://kubecost.com/>

408 <https://kubernetes.io/docs/tasks/access-application-cluster/web-ui-dashboard/>

409 <https://github.com/prometheus-community/helm-charts/tree/main/charts/kube-prometheus-stack>

410 <https://prometheus.io/>

411 <https://prometheus.io/docs/alerting/latest/alertmanager>

412 <https://grafana.com/>

413 <https://catalog.ngc.nvidia.com/orgs/nvidia/containers/gpu-operator>

414 <https://github.com/DirectXMan12/k8s-prometheus-adapter>

| Platform Application  | Dependencies          |
|-----------------------|-----------------------|
| kubecost              | traefik               |
| kubernetes-dashboard  | traefik               |
| kube-prometheus-stack | traefik               |
| prometheus-adapter    | kube-prometheus-stack |
| nvidia-gpu-operator   |                       |

#### 8.3.3.4.5 Security

Allows management of security constraints and capabilities for the clusters and users.

- [gatekeeper](#)<sup>415</sup>: A policy Controller for Kubernetes.

| Platform Application | Dependencies |
|----------------------|--------------|
| gatekeeper           |              |

#### 8.3.3.4.6 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](#)<sup>416</sup>: A reverse proxy server that authenticates users using OIDC to Kubernetes API servers where OIDC authentication is not available.
- [traefik-forward-auth](#)<sup>417</sup>: Installs a forward authentication application providing Google OAuth based authentication for Traefik.

| Platform Application | Dependencies          |
|----------------------|-----------------------|
| kube-oidc-proxy      | cert-manager, traefik |

<sup>415</sup> <https://github.com/open-policy-agent/gatekeeper>

<sup>416</sup> <https://github.com/jetstack/kube-oidc-proxy>


<sup>417</sup> <https://github.com/thomaseddon/traefik-forward-auth>

| Platform Application | Dependencies |
|----------------------|--------------|
| traefik-forward-auth | traefik      |

### 8.3.3.4.7 Backup

This platform application assists you with backing up and restoring your environment.

- [velero](https://velero.io/)<sup>418</sup>: 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> |

### 8.3.3.4.8 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](https://istio.io/)<sup>419</sup>: Addresses the challenges developers and operators face with a distributed or microservices architecture.
- [jaeger](https://www.jaegertracing.io/)<sup>420</sup>: A distributed tracing system used for monitoring and troubleshooting microservices-based distributed systems.
- [kiali](https://kiali.io/)<sup>421</sup>: 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.

<sup>418</sup> <https://velero.io/>

<sup>419</sup> <https://istio.io/latest/about/service-mesh/>

<sup>420</sup> <https://www.jaegertracing.io/>

<sup>421</sup> <https://kiali.io/>

| Platform Application | Dependencies                                       |
|----------------------|----------------------------------------------------|
| istio                | kube-prometheus-stack                              |
| jaeger               | istio                                              |
| kiali                | istio<br>jaeger (optional for monitoring purposes) |
| knative              | istio                                              |

#### 8.3.3.4.9 Related Information

- [Kommander security architecture \(see page 780\)](#)
- [Traefik Ingress controller \(see page 781\)](#)
- [Logging and audits \(see page 751\)](#)

#### 8.3.3.5 Workspace Platform Application Resource Requirements

Refer to [Workspace Platform Application Defaults and Resource Requirements \(see page 118\)](#) for a list of all platform applications, their default deployment configuration, required resources, and storage minimums.

## 8.4 Workspaces

### Enterprise

#### Allow teams to manage their own clusters using workspaces

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 their own clusters and role-based permissions, while retaining an overall organization-level view of all clusters in operation.

#### 8.4.1 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.

## 8.4.2 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.

## 8.4.3 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.

## 8.4.4 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.

## 8.4.5 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 570)
- [Platform Applications](#) (see page 570)
- [Platform Application Dependencies](#) (see page 575)
- [Workspace Platform Application Defaults and Resource Requirements](#) (see page 118)

## 8.4.6 Workspace Applications

### Enterprise

This section documents the applications and application types that you can use with DKP.

Application types are:

- [Catalog Applications](#) (see page 593) 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 570) are applications integrated into DKP.
- [Cluster-scoped Application Configuration via the UI](#) (see page 582)
- [Cluster-scoped Configuration for Existing AppDeployments](#) (see page 584)
- [Workspace Catalog Applications](#) (see page 593)

### 8.4.6.1 Cluster-scoped Application Configuration via the UI

#### Enterprise

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.

#### 8.4.6.1.1 Prerequisites

Ensure that you've provisioned or attached clusters in one of the following environments:

- [Amazon Web Services \(AWS\)](#) (see page 225)

- [Amazon Elastic Kubernetes Service \(EKS\)](#) (see page 284)
- [Microsoft Azure](#) (see page 258)

Refer to the current list of Catalog and Platform Applications:

- [Workspace Catalog Applications](#) (see page 593)
- [Workspace Platform Applications](#) (see page 118)

#### 8.4.6.1.2 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.

#### 8.4.6.1.3 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 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.

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.

#### 8.4.6.1.4 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.

#### 8.4.6.2 Cluster-scoped Configuration for Existing AppDeployments

##### Enterprise

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 \(see page 933\)](#) section.



### 8.4.6.2.1 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 585)
- [Enable an Application per Cluster](#) (see page 586)
- [Customizations per Cluster](#) (see page 587)
- [Verify the Current Application Configuration](#) (see page 593)

### 8.4.6.2.2 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 570), [DKP catalog](#) (see page 593) and [Custom Applications](#) (see page 605) 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.

### 8.4.6.2.3 Prerequisites (per-cluster application configuration)

#### Enterprise

- 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 570), [workspace catalog application deployment](#) (see page 600)).
- 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>
```

#### 8.4.6.2.4 Enable an Application per Cluster

### Enterprise

##### 8.4.6.2.4.1 Prerequisites

Ensure you have reviewed the [prerequisites \(see page 585\)](#) section.

##### 8.4.6.2.4.2 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.

```
dkp create appdeployment kube-prometheus-stack --app kube-prometheus-stack-34.9.3 --workspace ${WORKSPACE_NAME} --clusters attached-cluster1,attached-cluster2
```

2. **Optional:** [Check the current status \(see page 593\)](#) of the `AppDeployment` to see the names of the clusters where the application is currently enabled.

##### 8.4.6.2.4.3 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. Once you have [enabled the application at the workspace level \(see page 0\)](#), the `spec.clusterSelector` field [\(see page 880\)](#) will be populated.

- 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 which clusters your application is currently deployed, refer to the [Print and Review the Current State of your AppDeployment](#) (see page 563) 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:

```
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-34.9.3
 kind: ClusterApp
 clusterSelector:
 matchExpressions:
 - key: kommander.d2iq.io/cluster-name
 operator: In
 values:
 - attached-cluster1
 - attached-cluster3-new
EOF
```

#### 8.4.6.2.4.4 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 563) section.

#### 8.4.6.2.5 Customizations per Cluster

Enable or disable a customization per cluster.

- [Customize an Application per Cluster](#) (see page 588)
- [Disable a Custom Configuration of an Application for a Cluster](#) (see page 591)

### 8.4.6.2.5.1 Customize an Application per Cluster

## Enterprise

#### Prerequisites

Ensure you have reviewed the [prerequisites](#) (see page 585) 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 588), 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-34.9.3
 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-34.9.3
 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 562) 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 563) section.

#### 8.4.6.2.5.2 Disable a Custom Configuration of an Application for a Cluster

### Enterprise

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
yaml
```

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-34.9.3
 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-34.9.3
 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 \(see page 563\)](#) section.

#### 8.4.6.2.6 Verify the Current Application Configuration

Enterprise

##### 8.4.6.2.6.1 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 \(see page 563\)](#) section.

#### 8.4.6.3 Workspace Catalog Applications

Enterprise

Catalog applications are any third-party or open source applications that appear in the Catalog.

D2iQ provides [DKP Catalog Applications \(see page 593\)](#) for use in your environment.

- [Workspace DKP Catalog Applications \(see page 593\)](#)
- [Deployment of Catalog Applications in Workspaces \(see page 600\)](#)
- [Workspace Catalog Application Upgrades \(see page 604\)](#)
- [Custom Applications \(see page 605\)](#)

##### 8.4.6.3.1 Workspace DKP Catalog Applications

Enterprise

Catalog applications are applications provided by D2iQ for use in your environment.

### 8.4.6.3.1.1 Prerequisites

- Ensure your clusters run on a supported Kubernetes version for [this release of DKP](#) (see page 1055), and that said Kubernetes version is also compatible with your [catalog application version](#) (see page 595).
- For customers with an [Enterprise license](#) (see page 104) 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.

#### Install the DKP Catalog via the CLI

Follow these steps to install the DKP catalog from the CLI.

1. Refer to [Install DKP in an Air-gapped Environment with Catalog Applications](#) (see page 511) instructions, 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: v2.4.2
 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}
```

**Workspace DKP Catalog Applications**


| Name                      | App ID             | Compatible Kubernetes versions |
|---------------------------|--------------------|--------------------------------|
| kafka-operator-0.20.0     | kafka-operator     | 1.21                           |
| kafka operator-0.20.2     | kafka-operator     | 1.21 - 1.24                    |
| spark-operator-1.1.6      | spark-operator     | 1.21                           |
| spark operator-1.1.17     | spark-operator     | 1.21 - 1.24                    |
| zookeeper-operator-0.2.13 | zookeeper-operator | 1.21 - 1.24                    |

**8.4.6.3.1.2 Kafka in a Workspace****Enterprise****Information about the Kafka Operator**

[Apache Kafka](#)<sup>422</sup> 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](#)<sup>423</sup> 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.

**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 625) custom resource documentation for more information on creating Kafka clusters.

 Only install the Kafka operator once per workspace.

<sup>422</sup> <https://kafka.apache.org/>

<sup>423</sup> <https://banzaicloud.com/docs/supertubes/kafka-operator/>

1. Follow the generic installation instructions for workspace catalog applications on the [Application Deployment \(see page 600\)](#) 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](#)<sup>424</sup> for details on custom configuration for the operator.

### 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`:

```
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
```

### Resources

- [Kafka Operator Documentation](#)<sup>425</sup>
- [Kafka Operator GitHub Repository](#)<sup>426</sup>
- [Sample Kafka Operator Custom Resources](#)<sup>427</sup>
- [Apache Kafka Documentation](#)<sup>428</sup>

<sup>424</sup> <https://github.com/banzaicloud/koperator/tree/master/charts/kafka-operator#configuration>

<sup>425</sup> <https://banzaicloud.com/docs/supertubes/kafka-operator/>

<sup>426</sup> <https://github.com/banzaicloud/koperator>

<sup>427</sup> <https://github.com/banzaicloud/koperator/tree/master/config/samples>

<sup>428</sup> <https://kafka.apache.org/documentation/>

### 8.4.6.3.1.3 Spark Operator in a Workspace

## Enterprise

### 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, please refer to the API Definition. For details on its design, please 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, please provide your override configmap to enable desired Spark Operator features.

### Install Spark Operator

You can find generic installation instructions for workspace catalog applications on the [Application Deployment](#) (see page 600) topic.

- Only install the Spark operator once per workspace.

For details on custom configuration for the operator, refer to the [Spark Operator Helm Chart documentation](#)<sup>429</sup>.

After you finish the installation, see [Spark Operator in a Project](#) (see page 629) custom resource documentation for more information about how to submit your Spark jobs.

### Sample Override Configuration File

- Ensure you configure the AppDeployment with the appropriate override configmap.

- Using UI

---

<sup>429</sup> <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/charts/spark-operator-chart/README.md>

```
podLabels:
 owner: john
 team: operations
```

- Using CLI

See [Application Deployment \(see page 600\)](#) 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
```

### 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. Please 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
```

#### Resources

Here are some resources to learn more about Spark Operator:

- [Spark Operator Documentation](#)<sup>430</sup>
- [Spark Operator Quick Start Guide](#)<sup>431</sup>
- [Spark Operator User Guide](#)<sup>432</sup>
- [Spark Documentation](#)<sup>433</sup>

#### 8.4.6.3.1.4 Zookeeper Operator in a Workspace

### Enterprise

[ZooKeeper](#)<sup>434</sup> is a centralized service for maintaining configuration information, naming, providing distributed synchronization, and providing group services. The [ZooKeeper operator](#)<sup>435</sup> 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.

#### 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 627) 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 600) page.

<sup>430</sup> <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/README.md>

<sup>431</sup> <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/docs/quick-start-guide.md>

<sup>432</sup> <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/docs/user-guide.md>

<sup>433</sup> <https://spark.apache.org/docs/latest/>

<sup>434</sup> <https://zookeeper.apache.org/index.html>

<sup>435</sup> <https://github.com/pravega/zookeeper-operator>

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, please refer to the [ZooKeeper operator Helm Chart documentation](#)<sup>436</sup>.

#### 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).

1. Delete all `ZookeeperClusters`, see [Deleting ZooKeeper Clusters](#) (see page 629).
2. Uninstall a ZooKeeper operator `AppDeployment`:

```
kubectl -n <workspace namespace> delete AppDeployment <name of AppDeployment>
```

3. Remove ZooKeeper CRDs:

**WARNING:** Once you remove the CRDs, all deployed `ZookeeperClusters` will be deleted!

```
kubectl delete crds zookeeperclusters.zookeeper.pravega.io
```

#### Resources

- [ZooKeeper Operator Documentation](#)<sup>437</sup>
- [ZooKeeper Cluster Helm Chart](#)<sup>438</sup>
- [ZooKeeper Documentation](#)<sup>439</sup>

### 8.4.6.3.2 Deployment of Catalog Applications in Workspaces

Enterprise

#### Deploy applications to attached clusters using the CLI

<sup>436</sup> <https://github.com/pravega/zookeeper-operator/tree/master/charts/zookeeper-operator#configuration>

<sup>437</sup> <https://github.com/pravega/zookeeper-operator>

<sup>438</sup> <https://github.com/pravega/zookeeper-operator/tree/master/charts/zookeeper>

<sup>439</sup> <https://zookeeper.apache.org/documentation>



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 584) section of the documentation.

#### 8.4.6.3.2.1 Prerequisites

Before you begin, you must have:

- A running cluster with [Kommander installed](#) (see page 475). The cluster should be on a supported Kubernetes version for [this release of DKP](#) (see page 1055), and also compatible with the [catalog application version](#) (see page 595) you wish to install.
- An [Attach an Existing Kubernetes Cluster](#) (see page 678) section of the documentation completed.

Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace the attached cluster exists in:

```
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 603).

#### 8.4.6.3.2.2 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](#)<sup>440</sup>.

Alternately, you can [use the CLI to enable your catalog applications](#) (see page 0).

#### 8.4.6.3.2.3 Enable the Application using the CLI

See [Workspace Catalog Applications](#) (see page 593) 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 678) with an `AppDeployment` resource.
2. Within the `AppDeployment`, define the `appRef` to specify which `App` to enable:

```
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.

#### 8.4.6.3.2.4 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:

<sup>440</sup> <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#dns-subdomain-names>

```
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.

#### 8.4.6.3.2.5 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 | spark-operator | True  | Release reconciliation succeeded |
| 7m3s                 |                |       |                                  |

### 8.4.6.3.3 Workspace Catalog Application Upgrades

#### Enterprise

#### Upgrade catalog applications using the CLI and UI

##### 8.4.6.3.3.1 Prerequisites

- Ensure your clusters run on a supported Kubernetes version for [this release of DKP \(see page 1055\)](#), and that said Kubernetes version is also compatible with your [catalog application version \(see page 0\)](#).
- For customers with an [Enterprise license \(see page 104\)](#) 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.

#### 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.

If you are running on an **air-gapped environment with catalog applications**, ensure you have updated your catalog repository before upgrading. The catalog repository should contain the Docker registry containing the [DKP catalog application \(see page 593\)](#) images and a charts bundle file containing [DKP catalog applications \(see page 0\)](#) charts.



Catalog applications must be upgraded to the latest version BEFORE upgrading the Kubernetes version (or Konvoy version for managed Konvoy clusters) on attached clusters, due to the previous versions' incompatibility with Kubernetes 1.22.

#### 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**.

## 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>
```

```
dkp upgrade catalogapp <appdeployment-name> --workspace=<my-workspace-name> --to-version=<version.number>
```

For example, the following command upgrades the Kafka Operator application, named `kafka-operator-abc`, in a workspace to version `0.20.2`:

```
dkp upgrade catalogapp kafka-operator-abc --workspace=my-workspace --to-version=0.20.2
```

ⓘ 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.

### 8.4.6.3.4 Custom Applications

#### Enterprise

**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 606)
- [Git Repository Structure](#) (see page 607)
- [Workspace Application Metadata](#) (see page 609)
- [Enable a Custom Application from the Workspace Catalog](#) (see page 611)

#### 8.4.6.3.4.1 Create a Git Repository

### Enterprise

#### 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 504) 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 |
|--------------|-------------------------------------------------------------------|-------|
| 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](#)<sup>441</sup>.

### 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
[...]
```

### Related Information

- [Flux](#)<sup>442</sup>
- [Flux docs](#)<sup>443</sup>

#### 8.4.6.3.4.2 Git Repository Structure

### Enterprise

**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

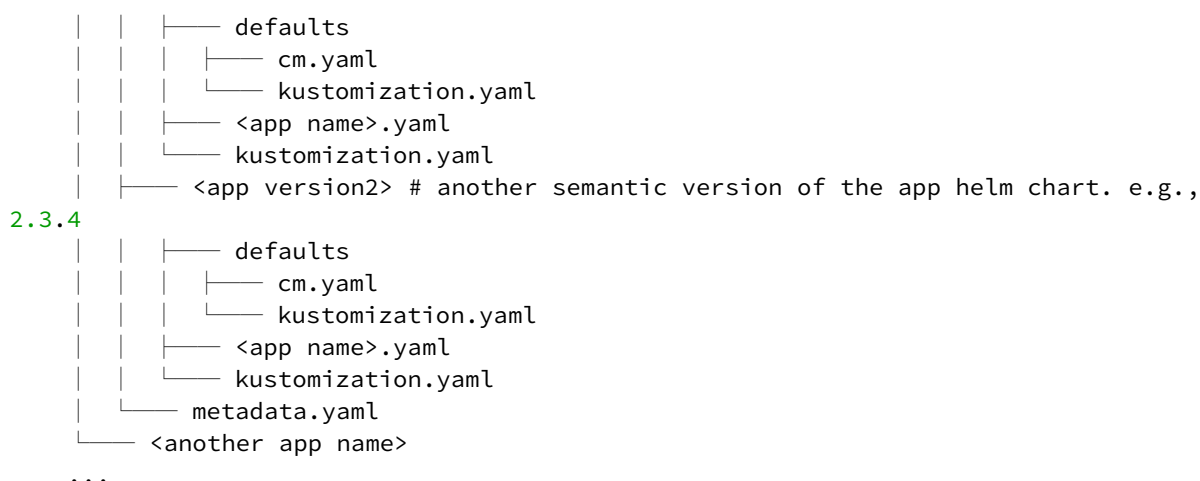
Use the following basic directory structure for your git repository:

```
├── 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
```

<sup>441</sup> <https://fluxcd.io/flux/components/source/gitrepositories/#secret-reference>

<sup>442</sup> <https://fluxcd.io/>

<sup>443</sup> <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](#)<sup>444</sup>, 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](#)<sup>445</sup>.
- 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 609).

See [the DKP Catalog repository](#)<sup>446</sup> 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

```

<sup>444</sup> <https://fluxcd.io/docs/components/helm/helmreleases/>

<sup>445</sup> <https://kubectl.docs.kubernetes.io/references/kustomize/kustomization/>

<sup>446</sup> <https://github.com/mesosphere/dkp-catalog-applications>



```
- ../../../../helm-repositories/<helm repository 1>
```

For more information, see the flux documentation about [HelmRepositories](#)<sup>447</sup>.

### Substitution Variables

Some [substitution variables](#)<sup>448</sup> 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.

### Related Information

- [Flux](#)<sup>449</sup>
- [Flux docs](#)<sup>450</sup>
- [Getting started with Flux guide](#)<sup>451</sup>

#### 8.4.6.3.4.3 Workspace Application Metadata

### Enterprise

**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. |

<sup>447</sup> <https://fluxcd.io/docs/components/source/helmrepositories/>

<sup>448</sup> <https://fluxcd.io/docs/components/kustomize/kustomization/#variable-substitution>

<sup>449</sup> <https://fluxcd.io/>

<sup>450</sup> <https://fluxcd.io/docs>

<sup>451</sup> <https://fluxcd.io/docs/get-started/>

| Field       | Default | Description                                                                                     |
|-------------|---------|-------------------------------------------------------------------------------------------------|
| description | ""      | Short description, should be a sentence or two, displayed in the UI on the application card.    |
| 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 xLLWJhY2tncm91bmQ6bmV3IDA gMCAzMdAgMzAwIiB4bWw6c3BhY2U9InByZXNlc nZlIj48cG  
F0aCBkPSJNMTUwIDUwQzk0LjggNTAgNTAgOTQuOCA1MCAxNTBzNDQuOCAxMDAgMTAwIDEwMCAxMDAtNDQuOCA  
xMDAtMTAwUzIwNS4yIDUwIDE1MCA1MHptMCAxODcuMmMtMTUuNyAwLTI4LjUtMTAuNS0yOC41LTIzLjRoNTYu  
OWMuMSAxMi45L TEyLjcgMjMuNC0yOC40IDIZLjR6bTQ3LTmxLjJoLTk0di0xN2g5NHYxN3ptLS4zLTI1LjloL  
TkzLjRjLS4zLS40LS42LS43LS45L TEuMS05LjYtMTEuNy0xMS45L TE3LjgtMTQuMS0yNCAwLS4yIDExLjcgMi  
40IDIwIDQuMyAwIDA gNC4zIDEgMTAuNSAyLjEtNi03LTkuNi0xNi05LjYtMjUuMSAwLTiWIDE1LjQtMzcuNiA  
5LjgtNTEuNyA1LjQuNCAxMS4yIDExLjQgMTEuNiAyOC41IDUuNy03LjkgOC4xLTIyLjQgOC4xLTMxLjMgMC05  
LjI gNi4xL TE5LjkgMTIuMS0yMC4yLTUuNCA4LjkgMS40IDE2LjUgNy40IDM1LjUgMi4zIDcuMSAyIDE5LjEgM  
y43IDI2LjcuNi0xNS44IDMuMy0zOC44IDEzLjMtNDYuNy00LjQgMTAgLjcgMjUuNSA0LjEgMjguNSA1LjYgOS  
43IDkgMTcuMSA5IDMxIDA gOS4zLTMuNCAxOC4xLTKuMyAyNSA2LjYtMS4yIDExLjItMi40IDExLjItMi40bDI  
xLjQtNC4yYy4xIDA tMyAxMi44L TE0LjkgMjUuMXoiIHNoeWxPSJmaWxs0iNmODQzMTEiLz48L3N2Zz4=

#### 8.4.6.3.4.4 Enable a Custom Application from the Workspace Catalog

Enterprise

##### 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 584) section.



From within a workspace, you can enable applications to deploy. Verify that an application has successfully deployed [via the CLI](#) (see page 0).

##### 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](#)<sup>452</sup>.

Alternately, you can [use the CLI to enable your custom applications](#) (see page 612).

#### 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>
```

#### 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:

---

<sup>452</sup> <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: my-custom-app
 namespace: ${WORKSPACE_NAMESPACE}
spec:
 appRef:
 name: custom-app-0.0.1
 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.

### 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.

### 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 ${WORKSPACE_NAMESPACE}
```

The output looks similar to this:

| NAMESPACE            | NAME          | READY | STATUS                           |
|----------------------|---------------|-------|----------------------------------|
| AGE                  |               |       |                                  |
| workspace-test-vjsfq | my-custom-app | True  | Release reconciliation succeeded |
| 7m3s                 |               |       |                                  |

## 8.4.7 Workspace Role Bindings

### Enterprise

**Workspace Role Bindings grant access to specified Workspace Roles for a specified group of people.**

### 8.4.7.1 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**.

## 8.5 Projects

### Enterprise

#### 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.

### 8.5.1 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.

### 8.5.2 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.

### 8.5.3 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

[+ Add Label](#)

**1 Cluster**

| Name     | Type    | Provider | Labels    |
|----------|---------|----------|-----------|
| cluster1 | Managed | aws      | dev: true |

### 8.5.4 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 617)
- [Project Deployments](#) (see page 641)
- [Project Role Bindings](#) (see page 657)
- [Project Roles](#) (see page 660)
- [Project ConfigMaps](#) (see page 663)
- [Project Secrets](#) (see page 665)
- [Project Quotas & Limit Ranges](#) (see page 666)
- [Project Network Policies](#) (see page 669)

## 8.5.5 Project Applications

### Enterprise

This section documents the applications and application types that you can utilize with DKP.

Application types are:

- [Catalog Applications](#) (see page 623) 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 624) are applications that are provided by D2iQ and added to the Catalog.
  - [Custom Applications](#) (see page 632) are applications you create and add to the Catalog.
- [Platform Applications](#) (see page 617) 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.

### 8.5.5.1 Project Platform Applications

#### Enterprise

#### 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 622) 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 \(see page 570\)](#).


### 8.5.5.1.1 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 \(see page 570\)](#)

 There may be dependencies between the applications, which are listed in [Project Platform Application Dependencies \(see page 621\)](#). Review them carefully prior to customizing to ensure that the applications are deployed successfully.

### 8.5.5.1.2 Platform Applications

| NAME                           | APP ID                  | Deployed by default |
|--------------------------------|-------------------------|---------------------|
| project-grafana-logging-6.38.1 | project-grafana-logging | False               |
| project-grafana-loki-0.48.6    | project-grafana-loki    | False               |
| project-logging-1.0.3          | project-logging         | False               |

### 8.5.5.1.3 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 [Upgrade Kommander](#) (see page 1026) for more information on how to upgrade these applications.

### 8.5.5.1.4 Deploy Project Platform Applications

#### Enterprise

#### 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 617).

See [Project Platform Applications](#) (see page 617) for a list of all applications and those that are enabled by default.

#### 8.5.5.1.4.1 Prerequisites

Before you begin, you must have:

- A running cluster with [Kommander installed](#) (see page 121).
- An [existing Kubernetes cluster attached to Kommander](#) (see page 678).

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_NAME} -n $
{WORKSPACE_NAMESPACE} -o jsonpath='{.status.namespaceRef.name}')
```

#### 8.5.5.1.4.2 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 \(see page 617\)](#).

1. Deploy one of the supported applications to [your existing attached cluster \(see page 678\)](#) 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.

#### 8.5.5.1.4.3 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_NAME}
```

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.

#### 8.5.5.1.4.4 Verify Applications

The applications are now enabled. Connect to the attached cluster and check the `HelmReleases` to verify the deployment:

```

kubectl get helmreleases -n ${PROJECT_NAMESPACE}

```

| NAMESPACE          | NAME                    | READY | STATUS                 |
|--------------------|-------------------------|-------|------------------------|
| AGE                |                         |       |                        |
| 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 \(see page 621\)](#) for that table.

#### 8.5.5.1.5 Project Platform Application Dependencies

Enterprise

##### 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 \(see page 619\)](#).

### 8.5.5.1.5.1 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.

### 8.5.5.1.5.2 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](#)<sup>453</sup>: Defines resources for the Logging Operator which uses them to direct the project's logs to its respective Grafana Loki application.
- [project-grafana-loki](#)<sup>454</sup>: A horizontally-scalable, highly-available, multi-tenant log aggregation system inspired by Prometheus.
- [project-grafana-logging](#)<sup>455</sup>: Logging dashboard used to view logs aggregated to Grafana Loki.



The project logging applications depend on the [workspace logging applications](#) (see page 753) being deployed.

| Application             | Dependencies                                                            |
|-------------------------|-------------------------------------------------------------------------|
| project-logging         | logging-operator (workspace)                                            |
| project-grafana-loki    | project-logging, grafana-loki (workspace), logging-operator (workspace) |
| project-grafana-logging | project-grafana-loki                                                    |

### 8.5.5.1.6 Project Platform Application Configuration Requirements

**ENTERPRISE**

<sup>453</sup> <https://grafana.com/oss/grafana/>

<sup>454</sup> <https://grafana.com/oss/loki/>

<sup>455</sup> <https://grafana.com/oss/grafana/>

### 8.5.5.1.6.1 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 |
|-------------------------|-----------------------------|-------------------------------------------------|---------------------|
| project-grafana-logging | cpu: 200m<br>memory: 100Mi  |                                                 | No                  |
| project-grafana-loki    |                             | # of PVs: 3<br>PV sizes: 10Gi x 3 (total: 30Gi) | No                  |
| project-logging         |                             |                                                 | No                  |

### 8.5.5.2 Project Catalog Applications

#### Enterprise

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 632) that can be used but are not supported by D2iQ.

- [Upgrade Project Catalog Applications](#) (see page 623)
- [Project-level DKP Applications](#) (see page 624)
- [Use Custom Resources with Workspace Catalog Applications](#) (see page 632)
- [Custom Project Applications](#) (see page 632)

#### 8.5.5.2.1 Upgrade Project Catalog Applications

#### Enterprise

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 Kubernetes version (or Konvoy version for managed Konvoy clusters) on attached clusters, due to the previous versions' incompatibility with Kubernetes 1.22.

#### 8.5.5.2.1.1 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**.

#### 8.5.5.2.1.2 Upgrade with the CLI

```
dkp upgrade catalogapp <appdeployment-name> --workspace=my-workspace --project=my-project --to-version=<version.number>
```

- [-] 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.

#### 8.5.5.2.2 Project-level DKP Applications

Enterprise

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.

#### 8.5.5.2.2.1 Kafka in a Project

### Enterprise


#### Deploying Kafka in a project

##### Get Started

To get started with creating and managing a Kafka Cluster in a project, you first need to deploy the [Kafka operator](#) (see page 595) and the [ZooKeeper operator](#) (see page 599) in the workspace where the project exists.

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](#)<sup>456</sup> for examples of the custom resources and their configurations.

##### 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 641) 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 first deploying a ZooKeeper cluster and then 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. 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's namespace:

```
kubectl apply -f - <<EOF
apiVersion: zookeeper.pravega.io/v1beta1
kind: ZookeeperCluster
metadata:
```

<sup>456</sup> <https://github.com/banzaicloud/koperator/tree/master/config/samples>

```

name: zookeeper
namespace: ${PROJECT_NAMESPACE}
spec:
 replicas: 1
EOF

```

3. Check the status of your ZooKeeper cluster using `kubectl` :

```
kubectl -n ${PROJECT_NAMESPACE} get zookeeperclusters
```

4. Download the [sample Kafka Cluster](#)<sup>457</sup> file.
5. Update the following attribute `zkAddresses` , replacing `zookeeper-client.zookeeper:2181` with `zookeeper-client.<project namespace>:2181` . You can use `sed` to update the file:

- MacOS sed

```

If you are using sed that comes installed on macOS
sed -i 's/zookeeper-client.zookeeper:2181/zookeeper-client.$
{PROJECT_NAMESPACE}:2181/g' simplekafkacluster.yaml

```

- GNU sed

```

If you are using GNU sed
sed -i "s/zookeeper-client.zookeeper:2181/zookeeper-client.$
{PROJECT_NAMESPACE}:2181/g" simplekafkacluster.yaml

```

6. Verify the file contains the following line:

```

...
zkAddresses:
 - "zookeeper-client.<project namespace>:2181"
...

```

7. Apply the `KafkaCluster` to your project's namespace:

```
kubectl apply -n ${PROJECT_NAMESPACE} -f simplekafkacluster.yaml
```

8. Check the status of your Kafka cluster using `kubectl` :

```
kubectl -n ${PROJECT_NAMESPACE} get kafkaclusters
```

<sup>457</sup> <https://raw.githubusercontent.com/banzaicloud/koperator/master/config/samples/simplekafkacluster.yaml>

With both the ZooKeeper cluster and Kafka cluster running in your project's namespace, refer to the [Kafka Operator documentation](#)<sup>458</sup> 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](#)<sup>459</sup>
- [Kafka Operator GitHub Repository](#)<sup>460</sup>
- [Sample Kafka Operator Custom Resources](#)<sup>461</sup>
- [Apache Kafka Documentation](#)<sup>462</sup>

#### 8.5.5.2.2.2 Zookeeper in a Project

## Enterprise

### 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 599) 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.

<sup>458</sup> <https://banzaicloud.com/docs/supertubes/kafka-operator/test/>

<sup>459</sup> <https://banzaicloud.com/docs/supertubes/kafka-operator/>

<sup>460</sup> <https://github.com/banzaicloud/koperator>

<sup>461</sup> <https://github.com/banzaicloud/koperator/tree/master/config/samples>

<sup>462</sup> <https://kafka.apache.org/documentation/>

A [Helm chart](#)<sup>463</sup> 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 641) 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        |          | 0.2.13  | 0.2.13  |         | 10.100.200 |
| .18:2181  | N/A      |          | 94s      |         |         |         |            |

<sup>463</sup> <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](#)<sup>464</sup>
- [ZooKeeper Cluster Helm Chart](#)<sup>465</sup>
- [ZooKeeper Documentation](#)<sup>466</sup>

**8.5.5.2.2.3 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](#)<sup>467</sup> 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 641) which enables you to leverage GitOps to deploy the resources. Otherwise, you will need to create the resources manually in each cluster.

<sup>464</sup> <https://github.com/pravega/zookeeper-operator>

<sup>465</sup> <https://github.com/pravega/zookeeper-operator/tree/master/charts/zookeeper>

<sup>466</sup> <https://zookeeper.apache.org/documentation>

<sup>467</sup> <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/docs/api-docs.md>

**Prerequisites**

Follow these steps:

1. Deploy your Spark Operator. See the [Spark Operator \(see page 597\)](#) 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](#)<sup>468</sup> 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
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
```

<sup>468</sup> <https://github.com/mesosphere/spark-on-k8s-operator/blob/d2iq-master/docs/quick-start-guide.md#about-the-spark-job-namespace>

```

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 615) 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 630) 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>
```

**8.5.5.2.3 Use Custom Resources with Workspace Catalog Applications****Enterprise**

Some workspace catalog applications provision some `CustomResourceDefinition`, which allow you to deploy Custom Resources. Refer to your workspace catalog application's documentation for instructions.

**8.5.5.2.4 Custom Project Applications****Enterprise**



### 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 \(see page 633\)](#)
- [Projects - Git repository structure \(see page 634\)](#)
- [Projects - Application Metadata \(see page 636\)](#)
- [Enable a Custom Application from the Project Catalog \(see page 638\)](#)

#### 8.5.5.2.4.1 Projects - Create a Git Repository

### Enterprise

#### 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 \(see page 509\)](#), 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                                                      | <a href="https://github.com/example-org/example-repo">https://github.com/example-org/example-repo</a> | True  |
| Fetched 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](#)<sup>469</sup>.

### 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
[...]
```

### Related Information

- [Flux](#)<sup>470</sup>
- [Flux docs](#)<sup>471</sup>

#### 8.5.5.2.4.2 Projects - Git repository structure

### Enterprise

**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

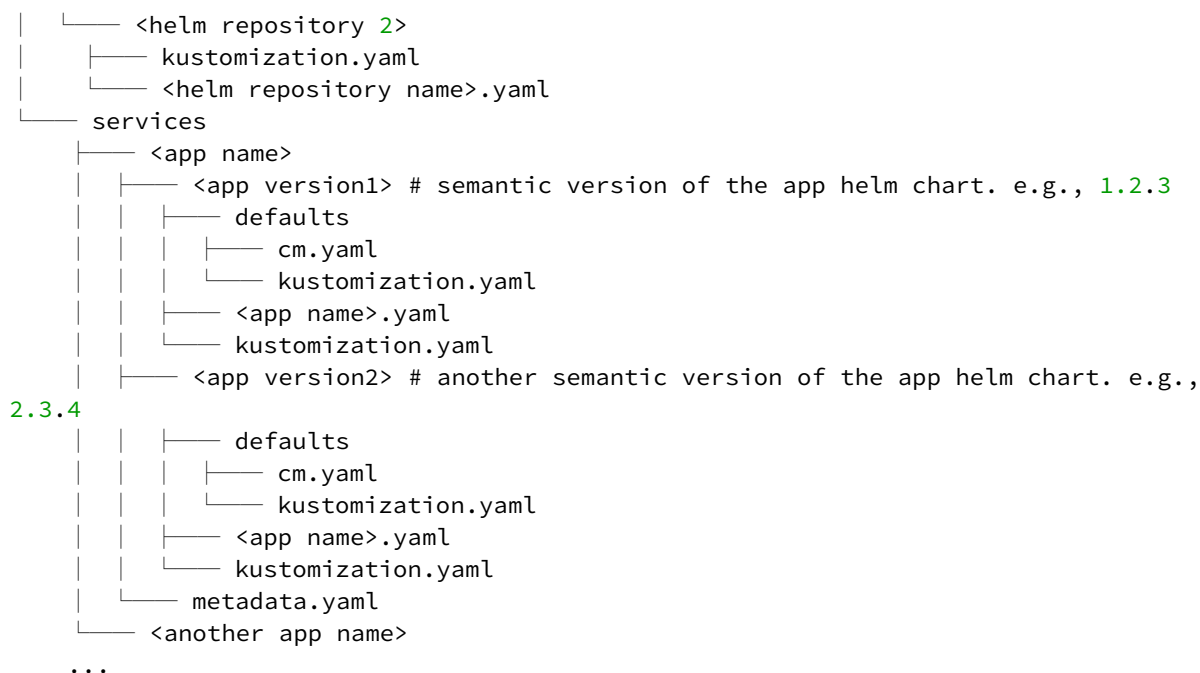
Use the following basic directory structure for your git repository:

```
├── helm-repositories
│ ├── <helm repository 1>
│ │ ├── kustomization.yaml
│ │ └── <helm repository name>.yaml
```

<sup>469</sup> <https://fluxcd.io/docs/components/source/gitrepositories/>

<sup>470</sup> <https://fluxcd.io/>

<sup>471</sup> <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](#)<sup>472</sup>, 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](#)<sup>473</sup>.
- 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 609).

See [the DKP Catalog repository](#)<sup>474</sup> 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:

<sup>472</sup> <https://fluxcd.io/docs/components/helm/helmreleases/>

<sup>473</sup> <https://kubectl.docs.kubernetes.io/references/kustomize/kustomization/>

<sup>474</sup> <https://github.com/mesosphere/dkp-catalog-applications>

```
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](#)<sup>475</sup>.

### Substitution Variables

Some [substitution variables](#)<sup>476</sup> 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](#)<sup>477</sup>
- [Flux docs](#)<sup>478</sup>
- [Getting started with Flux guide](#)<sup>479</sup>

#### 8.5.5.2.4.3 Projects - Application Metadata

### Enterprise

**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:

<sup>475</sup> <https://fluxcd.io/docs/components/source/helmrepositories/>

<sup>476</sup> <https://fluxcd.io/docs/components/kustomize/kustomization/#variable-substitution>

<sup>477</sup> <https://fluxcd.io/>

<sup>478</sup> <https://fluxcd.io/docs>

<sup>479</sup> <https://fluxcd.io/docs/get-started/>

| Field       | Default              | Description                                                                                     |
|-------------|----------------------|-------------------------------------------------------------------------------------------------|
| displayName | falls back to App ID | Display name of the application for the UI.                                                     |
| description | ""                   | Short description, should be a sentence or two, displayed in the UI on the application card.    |
| 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       | workspace            | List of scopes, can be workspace and/or project 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-iHZpZXdCb3g9IjAgMCAzMdAgMzAwIiBzdHlsZT0iZW5hYmxlLWJhY2tncm91bmQ6bmV3IDAgMCAzMdAgMzAwIiB4bWw6c3BhY2U9InByZXNlcnZlIj48cGF0aCBkPSJNMTUwIDUwQzZk0LjggNTAgNTAgOTQuOCA1MCAxNTBzNDQuOCAxMDAgMTAwIDEwMCAxMDAtNDQuOCAxMDAtMTAwUzIwNS4yIDUwIDE1MCA1MHptMCAxODcuMmMtMTUuNyAwLTI4LjUtMTAuNS0yOCA1LTIzLjRoNTYuOWMuMSAxMi45LTEyLjcgMjMuNC0yOCA4IDIZLjR6bTQ3LTMxLjJoLTk0di0xN2g5NHYxN3ptLS4zLTI1LjloLTKzLjRjLS4zLS40LS42LS43LS45LTEuMS05LjYtMTEuNy0xMS45LTE3LjgtMTQuMS0yNCAwLS4yIDExLjcgMj40IDIIwIDQuMyAwIDAgNC4zIDEGMTAuNSAyLjEtNi03LTkuNi0xNi05LjYtMjUuMSAwLTIwIDE1LjQtMzcuNiA5LjgtNTEuNyA1LjQuNCAxMS4yIDExLjQgMTEuNiAyOCA4IDUuNy03LjkgOCAxLTIyLjQgOCAxLTMxLjMgMCA05LjIjIj48LjE5LjkgMTEuMS0yMCA4LjUuNCA4LjkgMS40IDE2LjUgNy40IDM1LjUgMi4zIDcuMSAyIDE5LjEgMy43IDI2LjcuNi0xNS44IDMuMy0zOCA4IDEzLjMtNDYuNy00LjQgMTAgLjcgMjUuNSA0LjEgMjguNSA1LjYgOS43IDkgMTEuMSA5IDMxIDA0S4zLTMuNCAxOCA4LjUuNyAyNSA2LjYtMjUyIDExLjItMi40IDExLjItMi40bDIxLjQtNC4yYy4xIDAtMyAxMi44LTE0LjkgMjUuMXoiIHh0eWwlpSjmaWxsOiNmODQzMTEiLz48L3N2Zz4=

#### 8.5.5.2.4.4 Enable a Custom Application from the Project Catalog

### Enterprise

#### 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](#)<sup>480</sup>.

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
```

<sup>480</sup> <https://kubernetes.io/docs/concepts/overview/working-with-objects/names/#dns-subdomain-names>

```

metadata:
 name: my-custom-app
 namespace: ${PROJECT_NAMESPACE}
spec:
 appRef:
 name: custom-app-0.0.1
 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<br>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.

## 8.5.5.3 Project AppDeployments

### Enterprise

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 \(see page 561\)](#) documentation.

## 8.5.6 Project Deployments

### Enterprise

### 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? \(see page 642\)](#)
- [Continuous Delivery with GitOps \(see page 642\)](#)
- [Continuous Deployment \(see page 652\)](#)
- [Project Deployments Troubleshooting \(see page 656\)](#)
- [View Helm Releases \(see page 656\)](#)

### 8.5.6.1 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.

### 8.5.6.2 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 \(see page 643\)](#)
- [Deploy a Sample App from DKP GitOps \(see page 646\)](#)

### 8.5.6.2.1 Store secrets in GitOps repository using SealedSecrets

#### 8.5.6.2.1.1 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.

#### 8.5.6.2.1.2 Set up

These instructions are used as an example. For instructions on the latest release, see the [release page](#)<sup>481</sup>. For full documentation on SealedSecrets, see the [GitHub repo](#)<sup>482</sup>.

#### 8.5.6.2.1.3 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
```

#### 8.5.6.2.1.4 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
```

<sup>481</sup> <https://github.com/bitnami-labs/sealed-secrets/releases>

<sup>482</sup> <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.

### 8.5.6.2.1.5 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`<sup>483</sup> 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
```

<sup>483</sup> <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.

#### 8.5.6.2.1.6 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
```

#### 8.5.6.2.1.7 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
```

### 8.5.6.2.2 Deploy a Sample App from DKP GitOps

#### Enterprise

Use this procedure to deploy a sample `podinfo` application from DKP Enterprise GitOps.

#### 8.5.6.2.2.1 Prerequisites

- [Enterprise DKP installed](#) (see page 121)
- [Kommander installed](#) (see page 985)
- Github account and [personal access token](#)<sup>484</sup>
- [Add cluster to Kommander](#) (see page 678)
- [Setup Workspace and Projects](#) (see page 580)

#### 8.5.6.2.2.2 Deployment Steps

 This procedure was run on an AWS cluster with DKP 2.4.0 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.

<sup>484</sup> <https://docs.github.com/en/authentication/keeping-your-account-and-data-secure/creating-a-personal-access-token>

2. [Create a project \(see page 615\)](#).

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 |
|--------------------------|----------------|-------------------------|----------|-----------|
| <a href="#">pod info</a> | pod-info-xt2sz | No description provided |          |           |

3. [Optional] Create a secret in order to pull from the repository, for private repositories.
- Select the Secrets tab and set up your secret according to the [Continuous Deployment documentation \(see page 652\)](#).
  - 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 |
|-----------------------------|-------------------------------------|------|-----|
| <b>default</b> -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 |

- In the UI, select your project and then select the **Continuous Deployment (CD)** tab.
- 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.

- 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
```

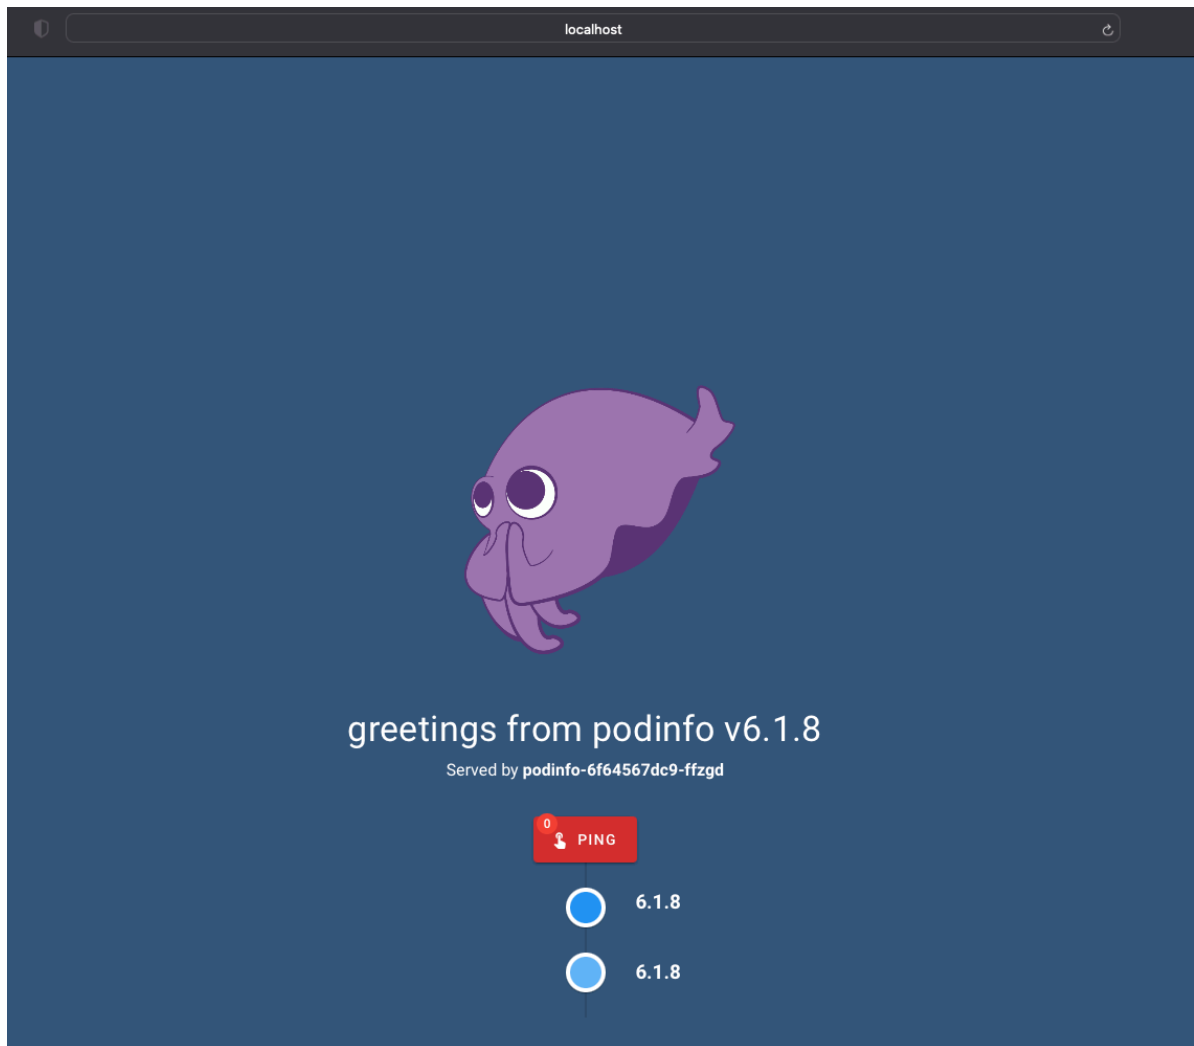
| NAME            | TYPE      | CLUSTER-IP    | EXTERNAL-IP | PORT(S)           |
|-----------------|-----------|---------------|-------------|-------------------|
| service/podinfo | ClusterIP | 10.99.239.120 | <none>      | 9898/TCP,9999/TCP |
| AGE             |           |               |             |                   |
| 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:



### 8.5.6.3 Continuous Deployment

#### Enterprise


After installing Kommander and [configuring your project and its clusters](#) (see page 674), 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.

#### 8.5.6.3.1 Prerequisites

You must have [Kommander installed](#) (see page 475) to run this procedure.

### 8.5.6.3.2 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 665).

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](#)<sup>485</sup>. 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.


<sup>485</sup> <https://docs.github.com/en/authentication/keeping-your-account-and-data-secure/creating-a-personal-access-token>

 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.

### 8.5.6.3.3 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
```

```
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.

#### 8.5.6.3.4 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](#)<sup>486</sup> and [Kustomize controller](#)<sup>487</sup> docs for advanced configuration and more examples.

### 8.5.6.4 Project Deployments Troubleshooting

#### Enterprise

- 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](#)<sup>488</sup> 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.
- Ensure your GitOps repository does not contain any `HelmRelease` and `Kustomization` resources that are targeting a different namespace than the project namespace.

### 8.5.6.5 View Helm Releases

#### Enterprise

#### 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.

<sup>486</sup> <https://fluxcd.io/docs/components/source/>

<sup>487</sup> <https://fluxcd.io/docs/components/kustomize/>

<sup>488</sup> <https://toolkit.fluxcd.io/components/>



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.

## 8.5.7 Project Role Bindings

Enterprise

**Project Role Bindings grant access to a specified Project Role for a specified group of people**

### 8.5.7.1 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**.

### 8.5.7.2 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
```

### 8.5.7.3 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-dbfj-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

```

#### 8.5.7.4 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 \(see page 899\)](#).

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`.

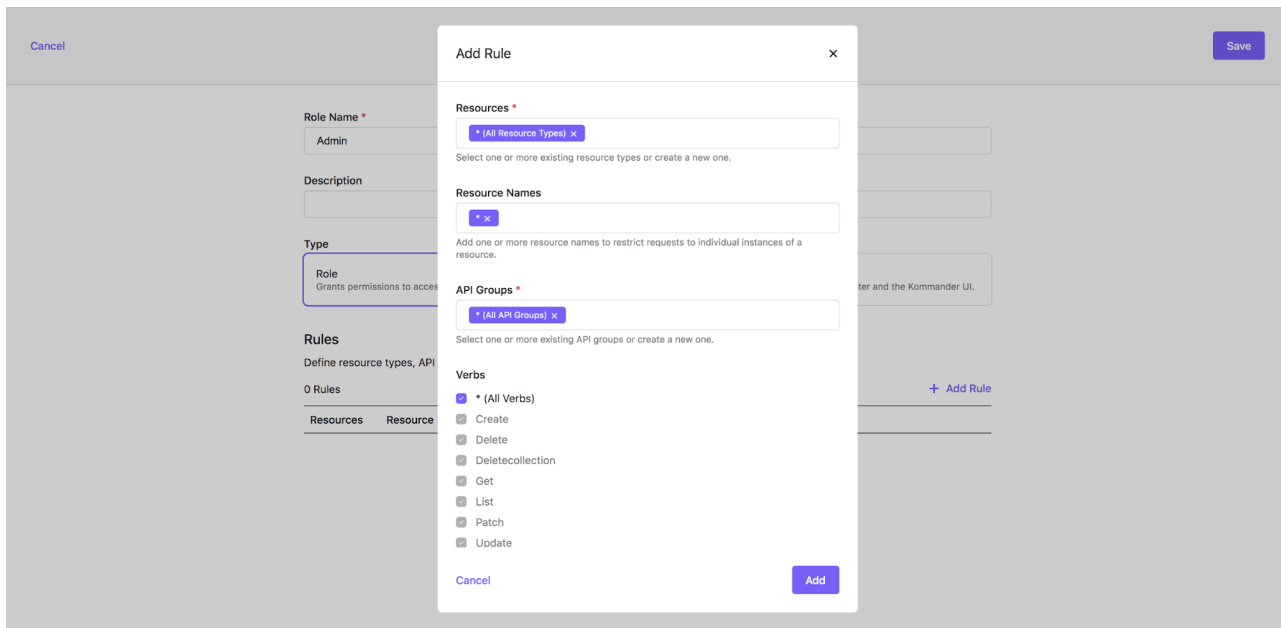
### 8.5.8 Project Roles

Enterprise

**Project Roles are used to define permissions at the namespace level.**

#### 8.5.8.1 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.

### 8.5.8.2 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:
 - '*'
```

## 8.5.9 Project ConfigMaps

Enterprise

### 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](https://kubernetes.io/docs/concepts/configuration/configmap/)<sup>489</sup>.

#### 8.5.9.1 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.

<sup>489</sup> <https://kubernetes.io/docs/concepts/configuration/configmap/>

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.

### 8.5.9.2 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
```

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
```

## 8.5.10 Project Secrets

### Enterprise

**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.**

### 8.5.10.1 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.


### 8.5.10.2 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
```

## 8.5.11 Project Quotas & Limit Ranges

### Enterprise

**Project Quotas and Limit Ranges can be set up to limit the number of resources the Project team uses.**

### 8.5.11.1 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.

### 8.5.11.2 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

```

## 8.5.12 Project Network Policies

### Enterprise

**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.

### 8.5.12.1 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.

### 8.5.12.2 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.

### 8.5.12.3 Creating Network Policies

You create network policies in three main parts:

- General information
- Ingress rules
- Egress rules

### 8.5.12.3.1 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.

### 8.5.12.3.2 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.

### 8.5.12.3 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.

### 8.5.12.4 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.

#### 8.5.12.4.1 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`

##### 8.5.12.4.1.1 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".

#### 8.5.12.4.1.2 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.

#### 8.5.12.4.1.3 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**.

#### 8.5.12.4.2 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`

##### 8.5.12.4.2.1 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”.

##### 8.5.12.4.2.2 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”.

#### 8.5.12.4.2.3 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**.

#### 8.5.12.4.3 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:

##### 8.5.12.4.3.1 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**.

##### 8.5.12.4.3.2 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**.

#### 8.5.12.4.4 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:

##### 8.5.12.4.4.1 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”.

#### 8.5.12.4.4.2 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**.

## 8.6 Manage 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.

- [Cluster Types](#) (see page 674)
- [Cluster Statuses](#) (see page 674)
- [Cluster Resources](#) (see page 676)
- [Platform Application](#) (see page 677)
- [Kubernetes Cluster Federation \(KubeFed\)](#) (see page 677)
- [Attach an Existing Kubernetes Cluster](#) (see page 678)
- [Advanced Creation of CLI Clusters](#) (see page 719)
- [Management Cluster](#) (see page 720)
- [Cluster Applications](#) (see page 721)
- [Custom Domains and Certificates Configuration](#) (see page 723)
- [Disconnect or Delete Clusters](#) (see page 729)

### 8.6.1 Cluster Types

There are several types of clusters that display in the Clusters tab. The cluster type appears in the cluster card just under the cluster name.

The type values include:

- **Attached:** An Attached cluster is one that was not created with Kommander. You cannot manage an Attached cluster's lifecycle, but you can monitor it.
- **Managed:** A Managed cluster is a Konvoy cluster that was created with Kommander. You can use Kommander to manage a Managed cluster's lifecycle.
- **Management:** This is the Konvoy cluster that hosts Kommander.

### 8.6.2 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.                             |
| 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.                                           |

| Status        | Description                                                                                                                                                                                                                                 |
|---------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 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 <a href="#">Manually attach a CLI-created cluster</a> (see page 717) page to resolve this status. |

\*These statuses only appear on Managed clusters.

### 8.6.3 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](#)<sup>490</sup> about resources.

## 8.6.4 Platform Application

Platform applications, are services that the management cluster installs. 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 118) 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 570) section.

## 8.6.5 Kubernetes Cluster Federation (KubeFed)

[Kubernetes Cluster Federation \(KubeFed\)](#)<sup>491</sup> 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.

In DKP, KubeFed is used 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. Once 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:

<sup>490</sup> <https://kubernetes.io/docs/concepts/configuration/manage-resources-containers/>

<sup>491</sup> <https://github.com/kubernetes-sigs/kubefed>

- <https://github.com/kubernetes-sigs/kubefed/blob/master/docs/concepts.md>
- <https://github.com/kubernetes-sigs/kubefed/blob/master/docs/userguide.md>

## 8.6.6 Attach an Existing Kubernetes Cluster

Enterprise

### Attach an existing Kubernetes cluster using kubeconfig

#### 8.6.6.1 Attach Kubernetes Cluster

You can attach an existing cluster directly to Kommander. 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](#) (see page 118) to ensure the attached cluster has sufficient resources. For more information on platform applications and customizing them, see [Workspace Applications](#) (see page 582).

If the cluster you want to attach was created using Amazon EKS or Google GKE, create a service account as described below. If you are attaching an Amazon EKS cluster to Kommander, [detailed instructions are available](#) (see page 708).

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](#) (see page 582).


- [Requirements for Attaching an Existing Cluster](#) (see page 678)
- [Create a kubeconfig File for your Cluster](#) (see page 680)
- [Generate a kubeconfig File](#) (see page 683)
- [Attach a Cluster with no Networking Restrictions](#) (see page 685)
- [Attach a Cluster with Networking Restrictions](#) (see page 686)
- [Finish Attaching the Existing Cluster](#) (see page 708)
- [Attach Amazon EKS Cluster](#) (see page 708)
- [Attach GKE Cluster](#) (see page 713)
- [Manually Attach a CLI-created Cluster](#) (see page 717)
- [Access a Managed Cluster](#) (see page 718)
- [Finalize Attaching your Cluster From the UI](#) (see page 719)

#### 8.6.6.2 Requirements for Attaching an Existing Cluster

Enterprise

### 8.6.6.2.1 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.

 DKP does not support attachment of K3s clusters.

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](#)<sup>492</sup>. You must ensure that the target cluster is running a compatible version.

### 8.6.6.2.2 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                 |
|-------------------------|-----------------|---------------|-----------------------------------|
| ebs-sc <b>(default)</b> | ebs.csi.aws.com | Delete        | WaitForFirstConsumer <b>false</b> |
| 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"
```

### 8.6.6.2.3 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 580) give you a logical way to represent your teams and specific configurations. [Projects](#) (see page 615) 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.

<sup>492</sup> <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29923823/2.3.0+Release+Notes#Supported-versions>

 Do not attach a cluster in the "Management Cluster Workspace" workspace. This workspace is reserved for your Application Management cluster only.

#### 8.6.6.2.4 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 570) that DKP provides by default.

#### 8.6.6.2.5 Attach Existing AWS, EKS and GKE clusters

Attaching an existing AWS cluster requires that the cluster be fully [configured and running](#) (see page 175). 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 284).

To attach an existing GKE cluster, refer to the specific information in [Attach GKE Cluster](#) (see page 713).

#### 8.6.6.2.6 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.

#### 8.6.6.3 Create a kubeconfig File for your Cluster

If you already have a kubeconfig file, go directly to [Attaching a cluster](#) (see page 680).



To get started, ensure you have [kubectl](#)<sup>493</sup> set up and configured with [ClusterAdmin](#)<sup>494</sup> 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](#)<sup>495</sup> 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"
 name: kommander-cluster-admin-sa-token
 namespace: default
 resourceVersion: "8554"
 uid: 72c2a4f0-636d-4a70-9f1c-55a75f15e520
type: kubernetes.io/service-account-token
```

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

494 <https://kubernetes.io/docs/concepts/cluster-administration/>

495 <https://eng.d2iq.com/blog/service-account-tokens-in-kubernetes-v1.24/#whats-changed-in-kubernetes-v124>

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 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
```

### 8.6.6.4 Generate a kubeconfig File

#### How to create a service account and generate a kubeconfig file for attaching an existing cluster

You should create a separate service account when attaching existing Amazon EKS, Azure AKS, or Google GKE Kubernetes clusters. This service account is needed because the kubeconfig files generated from those clusters are not usable out-of-the-box. They call CLI commands, such as `aws` or `gcloud`, and use locally-obtained authentication tokens. Having a separate service account also allows you to keep access to the cluster specifics and isolated to DKP Application Manager.

To get started, ensure you have `kubectl`<sup>496</sup> set up and configured with `ClusterAdmin`<sup>497</sup> for the cluster you want to connect.

1. Create the required service account using the command:

```
kubectl -n kube-system create serviceaccount kommander-cluster-admin
```

2. Configure the new service account for `cluster-admin` permissions. You can and paste this example ensuring that you use the service account created previously.

```
cat << EOF | kubectl apply -f -
apiVersion: rbac.authorization.k8s.io/v1
```

<sup>496</sup> <https://kubernetes.io/docs/tasks/tools/install-kubectl/>

<sup>497</sup> <https://kubernetes.io/docs/concepts/cluster-administration/cluster-administration-overview/>

```

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

```

3. Set up the following environment variables with access data needed for producing a new kubeconfig file:

```

export USER_TOKEN_NAME=$(kubectl -n kube-system get serviceaccount kommander-cluster-admin -o=jsonpath='{.secrets[0].name}')
export USER_TOKEN_VALUE=$(kubectl -n kube-system get secret/${USER_TOKEN_NAME} -o=go-template='{{.data.token}}' | base64 --decode)
export CURRENT_CONTEXT=$(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}'"}}{{index .cluster .server}} {{end}}')

```

4. Generate a kubeconfig file with these values:

```

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 procedure produces a file in your current working directory called, `kommander-cluster-admin-config`. The contents of this file are used in DKP Application Manager to attach the cluster.

Before importing this configuration, you can verify that it is functional by running the following command:

```
kubectl --kubeconfig $(pwd)/kommander-cluster-admin-config get all --all-namespaces
```

Then, you can use this kubeconfig to:

- Attach a cluster with [no additional networking restrictions](#) (see page 685)
- Attach a cluster that [has networking restrictions](#) (see page 686)



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.

### 8.6.6.5 Attach a Cluster with no Networking Restrictions

#### Enterprise

#### 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.
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 **Submit** to attach your cluster.

### 8.6.6.6 Attach a Cluster with Networking Restrictions

#### Enterprise

#### How to attach an existing cluster that has additional networking restrictions

Use this option when you want to attach a cluster that is in a DMZ, behind a NAT gateway, behind a proxy server or a firewall, or that requires additional access information. This procedure gathers the information required to create a kubeconfig file for the network tunnel between Kommander and the cluster you want to attach.

- If your cluster blocks public access, you may need to make the additional step of allowing certain authorized networks where Docker images are hosted for Konvoy to use your cluster, specifically `https://registry-1.docker.io/`

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.
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`.

**NOTE:** 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.

- If you have not attached this cluster before, you must create a new secret in the **Root CA Certificate** drop down menu. To do this in your Konvoy management cluster, view your base64 encoded Kubernetes secret values to copy and paste into the **Root CA Certificate** field:

```
echo $(kubectl get secret -n cert-manager kommander-ca -o go-template='{{index .data "tls.crt"}}')
```

Otherwise, select from the list of available Secrets.

- Add any **Extra Annotations** as needed.
- Select the **Save & Generate kubeconfig** button to generate the kubeconfig file for the network tunnel.

After the above is complete, finish [Attaching the Cluster](#) (see page 708).

As an alternative procedure, you can follow these instructions to [Use CLI to Add Managed Clusters to Kommander](#) (see page 717).

For information on TunnelGateway, review the [API documentation](#) (see page 876).

### 8.6.6.6.1 Attach Cluster Using Tunnel

#### Using the CLI to attach a Kubernetes Cluster using a Tunnel

Enterprise

##### 8.6.6.6.1.1 Identify the Management Cluster Endpoint

Obtain the hostname and CA certificate for the management cluster:

```
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"}}')
```

##### 8.6.6.6.1.2 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, if you wish to create a new workspace instead of using an existing workspace:

```
workspace=sample
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}
```

### 8.6.6.1.3 Create a Tunnel Gateway

Create a [tunnel gateway](#) (see page 0) on the management cluster to listen for tunnel agents on remote clusters:

**NOTE:** 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.

```

cacert_secret=kubetunnel-ca
gateway=sample-gateway

cat > gateway.yaml <<EOF
apiVersion: v1
kind: Secret
metadata:
 namespace: ${namespace}
 name: ${cacert_secret}
data:
 ca.crt:
 ${b64ca_cert}

apiVersion: kubetunnel.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@kubernetescd
EOF
kubectl apply -f gateway.yaml

```

You can verify the gateway exists using the command:

```
kubectl get tunnelgateway -n ${namespace} ${gateway}
```

#### 8.6.6.1.4 Connect a Remote Cluster

##### Create a tunnel connector

Create a [tunnel connector](#) (see page 704) on the management cluster for the remote cluster:

```

connector=sample-connector

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.

Verify the connector exists:

```
kubectl get tunnelconnector -n ${namespace} ${connector}
```

Wait for the tunnel connector to reach `Listening` state and then 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

kubectl get secret -n ${namespace} ${manifest} -o jsonpath='{.data.manifests\.yaml}'
| base64 -d > manifest.yaml

```

The `manifest.yaml` is applied successfully after the command completes.

Fetch the `manifest.yaml` to use it in the following section:

```

kubectl get secret -n ${namespace} ${manifest} -o jsonpath='{.data.manifests\.yaml}'
| base

```

- ⚠ 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 managed cluster

In the following commands, the `--kubeconfig` flag ensures that you set the context to the Attached or Managed cluster.

Copy the `manifest.yaml` file to the managed cluster and deploy the tunnel agent:

```

kubectl apply --kubeconfig=<managed_cluster_kubeconfig.conf> -f manifest.yaml

```

You can 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 | 0/1   | Running   | 0        | 14m |

#### Add the managed cluster into Kommander

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
```

Add the cluster into Kommander:

```
managed=private-cluster
display_name=${managed}

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
```

Wait for the managed cluster to join the attached cluster:

```
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}


kubectl get 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. 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.

#### 8.6.6.6.1.5 Use a Remote Cluster

To access services running on the remote cluster 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 managed 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.

**Managed cluster service**

These sections require a service to run on the managed 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 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
nginx.org.

Commercial support is available at
nginx.com.</p>
```

```
<p>Thank you for using nginx.</p>
</body>
</html>
```

### Use of kubeconfig file

This is primarily useful for running `kubectl` commands on the management cluster to monitor the managed cluster.

On the management cluster, a `kubeconfig` file for the 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 managed service resource, on the management cluster run:

```
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 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 managed cluster, demonstrating that the job on the management cluster accessed the service running on the 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='{.s
pec.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 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 managed cluster, demonstrating that the job on the management cluster accessed the service running on the managed cluster.

#### 8.6.6.6.1.6 API documentation (v1alpha1)

##### API Documentation (v1alpha1)

*This document is automatically generated from the API definition in the code.*

##### Page Contents

**TunnelGateway**

Provides an endpoint for remote clusters to connect to the management cluster.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>498</sup>	false
spec		<a href="#">TunnelGatewaySpec</a> (see page 703)	false

**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

---

<sup>498</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.24/#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

**TunnelGatewayList**

Contains a list of `TunnelGateway`.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>499</sup>	false
items		<a href="#">[]TunnelGateway</a> <sup>500</sup>	true

### 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> .	<a href="#">TunnelGatewayIngressSpec</a> (see page 701)	false

### KubeconfigWebhookStatus

Status of the kubeconfig webhook.

Field	Description	Scheme	Required
deploymentRef	A reference to the deployment for the kubeconfig webhook.	<code>corev1.LocalObjectReference</code>	false
serviceRef	A reference to the service for the kubeconfig webhook.	<code>corev1.LocalObjectReference</code>	false
ingressRef	A reference to the ingress for the kubeconfig webhook.	<code>corev1.LocalObjectReference</code>	false

<sup>499</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.24/#objectmeta-v1-meta>

<sup>500</sup> <https://docs.d2iq.com/dkp/kommander/2.2/clusters/tunnel-cli/api-reference/#tunnelgateway>

**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

**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		<a href="#">metav1.ObjectMeta</a> <sup>501</sup>	false
spec		<a href="#">TunnelConnectorSpec</a> (see page 705)	false
status		<a href="#">TunnelConnectorStatus</a> (see page 705)	false

**TunnelConnectorList**

Contains a list of `TunnelConnector`.

---

<sup>501</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.24/#objectmeta-v1-meta>



Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>502</sup>	false
items		<a href="#">[]TunnelConnector</a> (see <a href="#">page 704</a> )	true

**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

**TunnelConnectorStatus**


---

<sup>502</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.24/#objectmeta-v1-meta>

Field	Description	Scheme	Required
state	State of the tunnel connector: <b>Starting</b> - the initial state; <b>Listening</b> - the local tunnel server is waiting for the remote agent to connect; <b>Pending</b> - the remote agent has connected but the local proxy is not ready; <b>Connected</b> - the tunnel is configured and contact to the remote API server succeeded; <b>Disconnected</b> - the tunnel is configured but contact to the remote API server failed; <b>Failed</b> - an unexpected error occurred, such as not being able to parse the kubeconfig.	TunnelConnectorState	false
tunnelServer	Status of the tunnel server.	<a href="#">TunnelServerStatus</a> (see page 707)	false
kubeconfigWebhook	Status of the kubeconfig webhook.	<a href="#">KubeconfigWebhookStatus</a> (see page 703)	false
tunnelAgent	Status of the tunnel agent.	<a href="#">TunnelAgentStatus</a> (see page 704)	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

**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

### 8.6.6.7 Finish Attaching the Existing Cluster

#### Enterprise

#### How to apply the kubeconfig file to create the network tunnel to attach an existing cluster

Though the required file is now generated, you still need to apply it 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 686).
2. Copy the `kubectl apply . . .` command from the user interface and paste into your terminal session, substituting the actual name of the file for the variable. 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.
3. (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.

### 8.6.6.8 Attach Amazon EKS Cluster

#### Enterprise


### 8.6.6.8.1 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 674) this cluster. The following procedure shows how to attach an existing Amazon Elastic Kubernetes Service (EKS) cluster.

#### 8.6.6.8.2 Before you Begin

This procedure requires the following items and configurations:

- A fully configured and running Amazon [EKS](#)<sup>503</sup> cluster with administrative privileges.
- The current version DKP Enterprise is [installed](#) (see page 513) on your cluster.
- Ensure you have installed `kubectl` in your Management cluster.
- Install [aws-iam-authenticator](#)<sup>504</sup>. This binary is used to access your cluster using `kubectl`.

 This procedure assumes you have an existing and spun up Amazon EKS cluster(s) with administrative privileges. Refer to the Amazon [EKS](#)<sup>505</sup> for setup and configuration information.

### 8.6.6.8.3 Attach Amazon EKS Clusters

If you have an existing EKS cluster, you will use this section to attach your cluster. If you have a newly created DKP CLI cluster, you will use the [Manually Attach a CLI-created Cluster](#) (see page 717) document.

#### 8.6.6.8.3.1 Ensure you have access to 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
```


---

503 <https://aws.amazon.com/eks/>

504 <https://docs.aws.amazon.com/eks/latest/userguide/install-aws-iam-authenticator.html>

505 <https://aws.amazon.com/eks/>

### 8.6.6.8.3.2 Create a kubeconfig file for your EKS cluster

 If you already have a `kubeconfig` file, **SKIP this section**.

To get started, ensure you have `kubectl`<sup>506</sup> set up and configured with `ClusterAdmin`<sup>507</sup> 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](#)<sup>508</sup> 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:
```

<sup>506</sup> <https://kubernetes.io/docs/tasks/tools/#kubectl>

<sup>507</sup> <https://kubernetes.io/docs/concepts/cluster-administration/>

<sup>508</sup> <https://eng.d2iq.com/blog/service-account-tokens-in-kubernetes-v1.24/#whats-changed-in-kubernetes-v124>

```

annotations:
 kubernetes.io/service-account.name: kommander-cluster-admin
 kubernetes.io/service-account.uid: b22bc32e-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
```


#### 8.6.6.8.3.3 Finalize attaching 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 \(see page 686\)](#).
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.

#### 8.6.6.8.4 Related Information

For information on related topics or procedures, refer to the following:

- [Configuring and Running Amazon EKS Clusters](#)<sup>509</sup>
- [Installing and Configuring Konvoy v2.0 or above](#) (see page 259)
- [Installing and Configuring Kommander v2.0 or above](#) (see page 513)
- [Manage Clusters](#) (see page 674)

#### 8.6.6.9 Attach GKE Cluster

##### Enterprise


#### 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 674) this cluster. The following procedure shows how to attach an existing **standard** GKE cluster.

##### 8.6.6.9.1 Before you Begin

This procedure requires the following items and configurations:

- A fully configured and running [GKE cluster](#)<sup>510</sup> with a [supported Kubernetes version](#) (see page 1056), and administrative privileges.
- The current version of DKP Enterprise [installed](#) (see page 513) 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.

<sup>509</sup> <https://aws.amazon.com/eks/>

<sup>510</sup> <https://cloud.google.com/kubernetes-engine/docs/concepts/types-of-clusters>

## 8.6.6.9.2 Attach GKE Clusters

### 8.6.6.9.2.1 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
```

### 8.6.6.9.2.2 Configure a kubeconfig file



If you already have a `kubeconfig` file, **SKIP this section**.

To get started, ensure you have `kubect`<sup>511</sup> set up and configured with `ClusterAdmin`<sup>512</sup> 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
```

<sup>511</sup> <https://kubernetes.io/docs/tasks/tools/#kubectl>

<sup>512</sup> <https://kubernetes.io/docs/concepts/cluster-administration/>

```
type: kubernetes.io/service-account-token
EOF
```

For more information on Service Account Tokens, refer to [this article](#)<sup>513</sup> 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"
 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
```

<sup>513</sup> <https://eng.d2iq.com/blog/service-account-tokens-in-kubernetes-v1.24/#whats-changed-in-kubernetes-v124>

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}'"}}{{ index .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
```

### 8.6.6.9.2.3 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 \(see page 686\)](#).
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.

### 8.6.6.9.3 Related Information

For information on related topics or procedures, refer to the following:

- [Installing and Configuring Kommander \(see page 513\)](#)
- [Manage Clusters \(see page 674\)](#)

## 8.6.6.10 Manually Attach a CLI-created Cluster

Enterprise

**Manually attach a cluster that was created with the CLI**

When you create a cluster in a Workspace namespace using the DKP CLI, it does not attach automatically. To automatically attach a cluster, generate the cluster objects using the DKP CLI `--dry-run -o yaml` flags and create a cluster as stated in the [Advanced Creation of CLI Clusters \(see page 719\)](#) guide.

However, if you created the cluster using the CLI, you will still need to attach it. You will be able to see the new cluster in the UI while it is being provisioned. Once provisioning is completed, the status will change to Unattached.

#### 8.6.6.10.1 Manually Attach a Cluster that was Created with the CLI

From the CLI find out the `name` of the created `Cluster` so you can reference it later:

```
$ kubectl -n <workspace_namespace> get clusters
```

Attach the cluster by creating a `KommanderCluster` :

```
cat << EOF | kubectl apply -f -
apiVersion: kommander.mesosphere.io/v1beta1
kind: KommanderCluster
metadata:
 name: <cluster_name>
 namespace: <workspace_namespace>
spec:
 kubeconfigRef:
 name: <cluster_name>-kubeconfig
 clusterRef:
 capiCluster:
 name: <cluster_name>
EOF
```

#### 8.6.6.11 Access a Managed Cluster

Enterprise

**How to access an attached (managed) cluster with credentials**

##### 8.6.6.11.1 Accessing your Managed Clusters Using your UI Administrator Credentials

After the cluster is successfully attached, you can retrieve a custom kubeconfig file from the UI.

1. Select the Kommander username in the top right corner, and then select **Generate Token**.
2. Select the attached cluster name, and follow the instructions to assemble a kubeconfig for accessing its Kubernetes API. If Kommander prompts you to log in, use the credentials used to first login to 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 attached cluster name displays the instructions to assemble a kubeconfig for accessing its Kubernetes API.

### 8.6.6.12 Finalize Attaching your Cluster From the UI

Using the **Add Cluster** option, you can attach an existing Kubernetes or Konvoy cluster directly to Kommander. This enables you to access the multi-cluster management and monitoring benefits that Kommander provides, while keeping your existing cluster on its current provider and infrastructure.

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 \(see page 686\)](#).
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 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.

## 8.6.7 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.

### 8.6.7.1 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 \(see page 175\)](#) for an example on how to use the `--dry-run` and `--output` flags to create a set of cluster objects.

### 8.6.7.2 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 \(see page 615\)](#). 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.

## 8.6.8 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 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 \(see page 615\)](#) and enables the management of its [Platform Applications \(see page 617\)](#) 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.

### 8.6.8.1 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.



## 8.6.8.2 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.

## 8.6.9 Cluster Applications

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:

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 118) to ensure that the attached clusters have sufficient resources. For more information on applications and how to customize them, see [Workspace Applications](#) (see page 593).

### 8.6.9.1 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.

### 8.6.9.1.1 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.	

Key	Description	Required
<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
```

## 8.6.10 Custom Domains and Certificates Configuration

### Configure a custom domain and certificate for your Management or any Managed/Attached cluster

DKP supports configuring a custom domain name per cluster, so you can access the DKP UI and other platform services via that domain. Additionally, you can provide a custom certificate for the domain, or one can be issued automatically by Let's Encrypt (or other certificate authorities supporting the ACME protocol).

The configuration path is the same regardless of whether you are configuring a custom domain and certificate on the Management cluster, or a Managed/Attached cluster. However, you can choose to set up a customized domain and certificate for the Management cluster during DKP installation.

- Customize a domain or certificate in the [Management cluster, during installation \(see page 478\)](#).
- Customize a domain or certificate in the [Management or a Managed/Attached Cluster, after installation \(see page 725\)](#).

To customize the domain or certificate for a specific cluster after the installation of DKP, adapt or create an API YAML to specify the custom domain and certification details.

Here is an overview of the information provided in this section:

- [Why Should you set up a Custom Domain or Certificate? \(see page 724\)](#)

- [Configure Custom Domains or Custom Certificates](#) (see page 725)
- [Configuration Example with Let's Encrypt](#) (see page 726)
- [Verify and Troubleshoot Configuration Status](#) (see page 727)

## 8.6.10.1 Why Should you set up a Custom Domain or Certificate?

### 8.6.10.1.1 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.

### 8.6.10.1.2 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).



Using Let's Encrypt or other 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.


### 8.6.10.1.3 Related Topics

[Configure Custom Domains or Custom Certificates](#) (see page 725)

<https://d2iq.atlassian.net/wiki/spaces/DENT/pages/137593048/Configuration+Example+with+Let+s+Encrypt> (see page 726)

[Verify and Troubleshoot Configuration Status](#) (see page 727)

## 8.6.10.2 Configure Custom Domains or Custom Certificates

 Ensure your `dkp` configuration references the Management cluster of the environment where you want to customize the domain or certificate by setting the `KUBECONFIG=` environment variable, or using the `--kubeconfig` flag, [in accordance with Kubernetes conventions](#)<sup>514</sup>.

### 8.6.10.2.1 Edit the KommanderCluster Resource

To customize the domain or certificate of a cluster, alter the `spec` values of the `ingress` object in the `KommanderCluster` resource. Note that you can reference an issuer as an `issuerRef` **OR** a secret as a `certificateSecretRef`, as long as the object is created in the cluster where you want to customize the configuration.

### 8.6.10.2.2 Management, Managed or Attached cluster?

In the Management cluster, both the `KommanderCluster` and `issuerRef` or `certificateSecretRef` objects are on the same cluster. In Managed and Attached clusters, the `KommanderCluster` object is stored on the Management cluster, and the `issuerRef` or `certificateSecretRef` object is on the Managed or Attached cluster.

### 8.6.10.2.3 Configuration

Use the API YAML to customize the domain (via the `hostname` field), and the certificate (via the `issuerRef` or `certificateSecretRef` field).

You have two options to update and apply the `KommanderCluster` resource with the required ingress. Refer to the following examples:

- One option is to use a certificate that is managed automatically and supported by cert-manager like ACME (if you use Let's Encrypt, refer to the [example \(see page 726\)](#). For this, reference the **name of the Issuer** or `ClusterIssuer` that contains your ACME provider information in the `issuerRef` field, and enter the custom domain name in the `hostname` field of the target cluster:

```
cat <<EOF | kubectl -n <workspace_namespace> --kubeconfig
<management_cluster_kubeconfig> patch \
```

<sup>514</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

```
kommandercluster <cluster_name> --type='merge' --patch-file=/dev/stdin
spec:
 ingress:
 hostname: <cluster_hostname>
 issuerRef:
 name: <issuer_name>
 kind: ClusterIssuer # or Issuer depending on the issuer config
```

- Another option is to use a manually create a certificate that is **customized for your hostname**. Ensure the secret in the `certificateSecretRef` field and the custom domain name in the `hostname` field of the target cluster are provided:

```
kubectl create secret generic -n "${WORKSPACE_NAMESPACE}" domain-tls-certs \
 --from-file=ca.crt=$CERT_CA_PATH \
 --from-file=tls.crt=$CERT_PATH \
 --from-file=tls.key=$CERT_KEY_PATH \
 --type=kubernetes.io/tls
```



It is not possible to configure the namespace of the secret with a command. Ensure the secret is stored in the workspace namespace of the target cluster.

#### 8.6.10.2.4 Related topics

[Why Should you set up a Custom Domain or Certificate? \(see page 724\)](#)

[Configuration Example with Let's Encrypt \(see page 726\)](#)

[Verify and Troubleshoot Configuration Status \(see page 727\)](#)

### 8.6.10.3 Configuration Example with Let's Encrypt

#### 8.6.10.3.1 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: custom-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::YYYYYYYYYYYYY: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

```

### 8.6.10.3.2 Related Topics

[Why Should you set up a Custom Domain or Certificate? \(see page 724\)](#)

[Configure Custom Domains or Custom Certificates \(see page 725\)](#)

[Verify and Troubleshoot Configuration Status \(see page 727\)](#)

### 8.6.10.4 Verify and Troubleshoot Configuration Status

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` .

Inspect the modified `KommanderCluster` object:

```
kubectl describe kommandercluster -n <workspace_name> <cluster_name>
```

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
[...]
```

#### 8.6.10.4.1 Related Topics

[Why Should you set up a Custom Domain or Certificate? \(see page 724\)](#)

<https://d2iq.atlassian.net/wiki/spaces/DENT/pages/137593048/Configuration+Example+with+Let+s+Encrypt> (see page 726)

[Configure Custom Domains or Custom Certificates \(see page 725\)](#)




## 8.6.11 Disconnect or Delete Clusters

### 8.6.11.1 Disconnect or delete a cluster


#### 8.6.11.2 Disconnect vs. Delete

When you attach a cluster to your environment **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](#)<sup>515</sup>. 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. Deleting clusters via the CLI will not remove all DKP resources, in which case you will have to follow the [troubleshooting instructions](#) (see page 730) to finalize the detachment.

 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.

#### 8.6.11.2.1 Statuses

See [Statuses](#) (see page 674) for a list of possible states a cluster can have when it is getting disconnected or deleted.

<sup>515</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

### 8.6.11.3 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:

```
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.

## 8.7 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 using the following procedure.

- [Configure Velero with Rook Ceph Storage](#) (see page 731)
- [Configure Velero with Azure, or GCP storage](#) (see page 731)
- [Configure Velero with AWS](#) (see page 740)
- [Back up with Velero](#) (see page 746)

## 8.7.1 Configure Velero with Rook Ceph Storage

For default installations, DKP deploys [Velero](#)<sup>516</sup> integrated with [Rook Ceph](#)<sup>517</sup>, operating inside the same cluster.

For production use-cases, D2iQ advises to provide an *external* storage class to use with [Rook Ceph](#)<sup>518</sup>. See [Rook Ceph in DKP](#) (see page 858) for more information.

### 8.7.1.1 Cloud Provider Storage

If you do not want to use Rook Ceph for storing Velero backups and you can [configure Velero to use cloud provider storage](#) (see page 731).

## 8.7.2 Configure Velero with Azure, or GCP storage

### 8.7.2.1 Configure Velero to use Azure Blob Storage

**Prerequisites:** Create Azure related assets such as storage account, blob containers, resource group, roles, service principals prior to continuing.

1. Confirm that you have created your storage account, and your blob container [using these instructions](#)<sup>519</sup>.
2. Prep your credentials-velero file with the values. You will need to use the same credentials that you [created when creating the cluster](#) (see page 0). Please note that these credentials **should not** be Base64 encoded, as Velero will not read them properly. Export the `AZURE_BACKUP_RESOURCE_GROUP` that you created in the last step to be the `AZURE_RESOURCE_GROUP` in this step (in a later step, you will also use `AZURE_BACKUP_RESOURCE_GROUP`).

```
cat << EOF > ./credentials-velero
AZURE_SUBSCRIPTION_ID=${AZURE_SUBSCRIPTION_ID}
AZURE_TENANT_ID=${AZURE_TENANT_ID}
AZURE_CLIENT_ID=${AZURE_CLIENT_ID}
AZURE_CLIENT_SECRET=${AZURE_CLIENT_SECRET}
AZURE_RESOURCE_GROUP=${AZURE_RESOURCE_GROUP}
AZURE_CLOUD_NAME=AzurePublicCloud
EOF
```

<sup>516</sup> <https://velero.io/>

<sup>517</sup> <https://rook.io/>

<sup>518</sup> <https://rook.io/>

<sup>519</sup> <https://github.com/vmware-tanzu/velero-plugin-for-microsoft-azure#create-azure-storage-account-and-blob-container>

3. Use the `credentials-velero` file to create the secret (in this case, it is named as `azure-bsl-credentials`). Note that we used `--from-env-file` referencing the `credentials-velero` file. If you are backing up the Management cluster, the namespace is `kommander`.

```
kubectl create secret generic -n ${WORKSPACE_NAMESPACE} azure-bsl-credentials
--from-file=azure="credentials-velero" --kubeconfig=${CLUSTER_NAME}.conf
```

### 8.7.2.1.1 Configure Velero on Attached or Managed Clusters

1. Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace

```
export WORKSPACE_NAMESPACE=<your_workspace_namespace>
```

2. Create a ConfigMap to apply Azure to the Velero configuration

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
 namespace: ${WORKSPACE_NAMESPACE}
 name: velero-overrides
data:
 values.yaml: |
 initContainers:
 - name: velero-plugin-for-aws
 image: velero/velero-plugin-for-aws:v1.1.0
 imagePullPolicy: IfNotPresent
 volumeMounts:
 - mountPath: /target
 name: plugins
 - 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: azure-bsl-credentials
EOF
```

3. Patch the Velero AppDeployment by adding the `configOverrides` value. This applies the ConfigMap to your instance after the cluster has been configured.

```
cat << EOF | kubectl -n ${WORKSPACE_NAMESPACE} patch appdeployment velero --
type="merge" --patch-file=/dev/stdin
spec:
 configOverrides:
 name: velero-overrides
EOF
```

4. After patching the AppDeployment, you will see the ConfigMap on the `HelmRelease` object

```
kubectl wait --for=jsonpath='{.spec.valuesFrom[1].name}'=velero-overrides
HelmRelease/velero -n ${WORKSPACE_NAMESPACE}
```

5. Create the backup storage location via Velero CLI (note that this calls for the `BLOB_CONTAINER` and `AZURE_STORAGE_ACCOUNT_ID` variable that was used when creating the blob container in step 1, as well as the `AZURE_BACKUP_SUBSCRIPTION_ID` which will be the same as the `AZURE_SUBSCRIPTION_ID` set previously):

```
velero backup-location create azure -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=azure-bsl-credentials=azure --kubeconfig=${CLUSTER_NAME}.conf
```

6. Verify that the Azure backup location is created:

```
velero backup-location get -n ${WORKSPACE_NAMESPACE} --kubeconfig=${
CLUSTER_NAME}.conf
```

7. Check the Helm releases that the new Velero configuration has been applied:

```
kubectl get helmrelease -n ${WORKSPACE_NAMESPACE} --kubeconfig=${
CLUSTER_NAME}.conf
```

8. Verify that the Velero pod is running:

```
kubectl get pods -A --kubeconfig=${CLUSTER_NAME}.conf |grep velero
```

9. Create a test backup for Azure:

```
velero backup create azure-velero-testbackup -n ${WORKSPACE_NAMESPACE} --
kubeconfig=${CLUSTER_NAME}.conf --storage-location azure --snapshot-
volumes=false
```

10. View your backup:

```
velero backup describe azure-velero-testbackup
```

### 8.7.2.1.2 Configure Velero on the Management Cluster

1. Create the backup storage location via Velero CLI (note that this calls for the `BLOB_CONTAINER` and `AZURE_STORAGE_ACCOUNT_ID` variable that was used when creating the blob container in step 1, as well as the `AZURE_BACKUP_SUBSCRIPTION_ID` which will be the same as the `AZURE_SUBSCRIPTION_ID` set earlier):

```
velero backup-location create azure -n kommander \
--provider azure \
--bucket ${BLOB_CONTAINER} \
--config resourceGroup=${AZURE_BACKUP_RESOURCE_GROUP},storageAccount=${
AZURE_STORAGE_ACCOUNT_ID},subscriptionId=${AZURE_BACKUP_SUBSCRIPTION_ID} \
--credential=azure-bsl-credentials=azure --kubeconfig=${CLUSTER_NAME}.conf
```

2. Verify that the Azure backup location is created:

```
velero backup-location get -n kommander --kubeconfig=${CLUSTER_NAME}.conf
```

3. Output the Kommander configuration to `kommander.yaml` (or use your existing configuration file)

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

- a. Configure DKP to load the plugins and to include the secret in the `apps.velero` section:  
**NOTE:** This process has been tested to work with plugins for AWS v1.1.0 and Azure v1.5.1. Newer versions of these plugins can be used, but have not been tested by D2iQ.

```
...
velero:
 enabled: true
 values: |
 initContainers:
 - name: velero-plugin-for-aws
 image: velero/velero-plugin-for-aws:v1.1.0
```

```

 imagePullPolicy: IfNotPresent
 volumeMounts:
 - mountPath: /target
 name: plugins
 - 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: azure-bsl-credentials
 ...

```

4. Use the modified `kommander.yaml` configuration in install this Velero configuration:

```

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

```

5. Check the Helm releases that the new Velero configuration applied/loaded:

```

kubectl get helmrelease -n kommander --kubeconfig=${CLUSTER_NAME}.conf

```

6. Ensure that the Velero pod is running:

```

kubectl get pods -A --kubeconfig=${CLUSTER_NAME}.conf |grep velero

```

7. Create a test backup for Azure:

```

velero backup create azure-velero-testbackup -n kommander --kubeconfig=${CLUSTER_NAME}.conf --storage-location azure --snapshot-volumes=false

```

8. View your backup:

```

velero backup describe azure-velero-testbackup

```

### 8.7.2.2 Configure Velero to use Google Cloud Buckets

You can also store your backups in **Google Cloud/GCP**.

See the [official docs](#)<sup>520</sup> for details on how to use different types of authentication.

**Prerequisites:** Create GCP related assets such as GCS Bucket, GCP project, service accounts, and service account keys prior to continuing, and the `velero`, `gcloud`, and `gsutil` CLIs installed locally (`gsutil` is optional, you may buckets through the GCP web application).

1. Confirm that you have created your storage account, and your bucket, [using these instructions](#)<sup>521</sup>.
2. Prep your credentials-velero file with the values, using the `SERVICE_ACCOUNT_EMAIL` you used to [grant permissions to your bucket](#)<sup>522</sup>. This creates a `credentials-velero` file in your local directory.

```
gcloud iam service-accounts keys create credentials-velero \
 --iam-account $SERVICE_ACCOUNT_EMAIL
```

3. Use the `credentials-velero` file to create the secret (in this case, we named it `bsl-credentials`). Note that we used `--from-env-file` referencing the `credentials-velero` file. If you are backing up the Management cluster, the namespace is `kommander`.

```
kubectl create secret generic -n ${WORKSPACE_NAMESPACE} bsl-credentials --from-
file=gcp=credentials-velero --kubeconfig=${CLUSTER_NAME}.conf
```

### 8.7.2.2.1 Configuring Velero on Attached or Managed Clusters

1. Set the `WORKSPACE_NAMESPACE` environment variable to the name of the workspace's namespace

```
export WORKSPACE_NAMESPACE=<your_workspace_namespace>
```

2. Create the backup storage location via Velero CLI (note that this calls for the `BUCKET` variable that was used when creating the bucket container in step 1:

```
velero backup-location create gcp-backup -n ${WORKSPACE_NAMESPACE} \
 --provider gcp \
 --bucket $BUCKET \
 --credential=bsl-credentials=gcp
```

<sup>520</sup> <https://github.com/vmware-tanzu/velero-plugin-for-gcp>

<sup>521</sup> <https://github.com/vmware-tanzu/velero-plugin-for-gcp#create-an-gcs-bucket>

<sup>522</sup> <https://github.com/vmware-tanzu/velero-plugin-for-gcp#create-custom-role-with-permissions-for-the-velero-gsa>



3. Verify that the GCP backup location is created:

```
velero backup-location get -n ${WORKSPACE_NAMESPACE} --kubeconfig=${CLUSTER_NAME}.conf
```

4. Create a ConfigMap to apply GCP to the Velero configuration

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
kind: ConfigMap
metadata:
 namespace: ${WORKSPACE_NAMESPACE}
 name: velero-overrides
data:
 values.yaml: |
 initContainers:
 - name: velero-plugin-for-aws
 image: velero/velero-plugin-for-aws:v1.1.0
 imagePullPolicy: IfNotPresent
 volumeMounts:
 - mountPath: /target
 name: plugins
 - name: velero-plugin-for-gcp
 image: velero/velero-plugin-for-gcp:v1.5.0
 imagePullPolicy: IfNotPresent
 volumeMounts:
 - mountPath: /target
 name: plugins
 credentials:
 extraSecretRef: bsl-credentials
EOF
```

5. Patch the Velero AppDeployment by adding the configOverrides value. This applies the ConfigMap in this instance after the cluster has been configured.

```
cat << EOF | kubectl -n ${WORKSPACE_NAMESPACE} patch appdeployment velero --
type="merge" --patch-file=/dev/stdin
spec:
 configOverrides:
 name: velero-overrides
EOF
```

6. After patching the AppDeployment, you will see the ConfigMap on the `HelmlRelease` object:

```
kubectl wait --for=jsonpath='{.spec.valuesFrom[1].name}'=velero-overrides
HelmlRelease/velero -n ${WORKSPACE_NAMESPACE}
```

7. Check the Helm releases that the new Velero configuration applied/loaded:

```
kubectl get helmrelease -n ${WORKSPACE_NAMESPACE} --kubeconfig=${CLUSTER_NAME}.conf
```

8. Ensure that the Velero pod is running:

```
kubectl get pods -A --kubeconfig=${CLUSTER_NAME}.conf |grep velero
```

9. Create a test backup for GCP:

```
velero backup create gcp-velero-testbackup -n ${WORKSPACE_NAMESPACE} --
kubeconfig=${CLUSTER_NAME}.conf --storage-location gcp-backup --snapshot-
volumes=false
```

10. Please note: 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}
```

And then confirm your `backupstoragelocation` was configured correctly

```
kubectl get backupstoragelocations -n ${WORKSPACE_NAMESPACE}
```

If your backup storage location is "Available", repeat step 9 and proceed to the rest of the setup

NAME	PHASE	LAST VALIDATED	AGE	DEFAULT
gcp-backup	Available	38s	60m	

11. View your backup:

```
velero backup describe gcp-velero-testbackup
```

### 8.7.2.2.2 Configuring Velero on the Management Cluster

1. Create the backup storage location via Velero CLI (note that this calls for the `BUCKET` variable that was used when creating the bucket container in step 1):

```
velero backup-location create gcp-backup -n kommander \
```

```
--provider gcp \
--bucket $BUCKET \
--credential=bsl-credentials=gcp
```

2. Verify that the GCP backup location is created:

```
velero backup-location get -n kommander --kubeconfig=${CLUSTER_NAME}.conf
```

3. Output the Kommander configuration to `kommander.yaml` (or use your existing configuration file)

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

4. Configure DKP to load the plugins and to include the secret under the `apps.velero` section:

**NOTE:** This process has been tested to work with plugins for AWS v1.1.0 and GCP v1.5.0. Newer versions of these plugins can be used, but have not been tested by D2iQ.

```
...
velero:
 enabled: true
 values: |
 initContainers:
 - name: velero-plugin-for-aws
 image: velero/velero-plugin-for-aws:v1.1.0
 imagePullPolicy: IfNotPresent
 volumeMounts:
 - mountPath: /target
 name: plugins
 - name: velero-plugin-for-gcp
 image: velero/velero-plugin-for-gcp:v1.5.0
 imagePullPolicy: IfNotPresent
 volumeMounts:
 - mountPath: /target
 name: plugins
 credentials:
 extraSecretRef: bsl-credentials
...

```

5. Use the modified `kommander.yaml` configuration to install this Velero configuration:

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

6. Check the Helm releases that the new Velero configuration applied/loaded (this normally takes a few minutes to catch up after running the install command):

```
kubectl get helmrelease -n kommander --kubeconfig=${CLUSTER_NAME}.conf
```

7. Ensure that the Velero pod is running:

```
kubectl get pods -A --kubeconfig=${CLUSTER_NAME}.conf |grep velero
```

8. Create a test backup for GCP:

```
velero backup create gcp-velero-testbackup -n kommander --kubeconfig=${CLUSTER_NAME}.conf --storage-location gcp-backup --snapshot-volumes=false
```

9. Please note: 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 kommander
```

And then confirm your `backupstoragelocation` was configured correctly

```
kubectl get backupstoragelocations -n kommander
```

If your backup storage location is "Available", repeat step 8 and proceed to the rest of the setup

NAME	PHASE	LAST VALIDATED	AGE	DEFAULT
gcp-backup	Available	38s	60m	

10. View your backup:

```
velero backup describe gcp-velero-testbackup
```

## 8.7.3 Configure Velero with AWS

Configure Velero to use S3 buckets as storage for its backup operations.

### 8.7.3.1 Overview

Follow these procedures to set up Velero with AWS S3:

- [Velero with AWS S3 - Prepare your Environment \(see page 741\)](#)
- [Velero with AWS S3 - Configure Velero \(see page 742\)](#)
- [Velero with AWS S3 - Establish a Backup Location \(see page 745\)](#)

## 8.7.3.2 Velero with AWS S3 - Prepare your Environment

### 8.7.3.2.1 Prerequisites

- Ensure you have installed Velero (included in the default DKP installation).
- Ensure you have [installed the Velero CLI \(see page 747\)](#).
- You have [created an S3 bucket with AWS](#)<sup>523</sup>.

#### 8.7.3.2.1.1 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:

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

**i** This can be the `kommander` namespace for the Management cluster or any other additional workspace namespace for Attached or Managed clusters. To list all available workspaces and clusters, use the `dkp get clusters -A` (see page 982) 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>
```

### 8.7.3.2.2 Prepare your AWS Credentials

**📖** See the [official docs](#)<sup>524</sup> 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`.

<sup>523</sup> <https://docs.aws.amazon.com/AmazonS3/latest/userguide/creating-bucket.html>

<sup>524</sup> <https://github.com/vmware-tanzu/velero-plugin-for-aws>

```
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, a Managed or an Attached 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
```

### 8.7.3.2.3 Next Step:

[Velero with AWS S3 - Configure Velero \(see page 742\)](#)

### 8.7.3.3 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}
 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 588).

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

### 8.7.3.3.1 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
```

### 8.7.3.3.2 Next Step:

[Velero with AWS S3 - Establish a Backup Location](#) (see page 745)

## 8.7.3.4 Velero with AWS S3 - Establish a Backup Location


### 8.7.3.4.1 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}`

#### 8.7.3.4.2 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 \(see page 746\)](#)

### 8.7.4 Back up with Velero

DKP provides [Velero](#)<sup>525</sup> by default, to support backup and restore operations for your Kubernetes clusters and persistent volumes.

- [Install the Velero CLI \(see page 747\)](#)
- [Regular Backup Operations \(see page 747\)](#)
- [Restore a cluster from backup \(see page 749\)](#)

---

<sup>525</sup> <https://velero.io/>

- [Backup Service Diagnostics](#) (see page 751)

### 8.7.4.1 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](#)<sup>526</sup> for more information.

**IMPORTANT:** The version of Velero that is shipped in DKP v2.4 is **v1.9.2**. We recommend using the same version for the Velero CLI.

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
```

### 8.7.4.2 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 749)

#### 8.7.4.2.1 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

<sup>526</sup> <https://velero.io/docs/v1.5/basic-install/#install-the-cli>

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.

#### 8.7.4.2.1.1 Alternate backup schedules

The Velero CLI provides an easy way to create alternate backup schedules. For example, you can use a command similar to:

```
velero create schedule thrice-daily --schedule="@every 8h"
```

To change the default backup service settings:

1. Check the backup schedules currently configured for the cluster by running the following command:

```
velero get schedules
```

2. Delete the `velero-default` schedule by running the following command:

```
velero delete schedule velero-default
```

3. Replace the default schedule with your custom settings by running the following command:

```
velero create schedule velero-default --schedule="@every 24h"
```

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 system-critical --include-namespaces=kube-system,kube-public,k
ommander --schedule="@every 24h"
```

The Velero command line interface provides many more options worth exploring. You can also find tutorials for [disaster recovery](#)<sup>527</sup> and [cluster migration](#)<sup>528</sup> on the Velero community site.

#### 8.7.4.2.2 Back up on Demand

In some cases, you might find it necessary create a backup outside of the regularly-scheduled interval. For example, if you are preparing to upgrade any components or modify your cluster configuration, you should perform a backup immediately before taking that action.

Create a backup by running the following command:

<sup>527</sup> <https://velero.io/docs/v0.11.0/disaster-case>

<sup>528</sup> <https://velero.io/docs/v0.11.0/migration-case>

```
velero backup create <BACKUP-NAME> --snapshot-volumes=false
```

### 8.7.4.3 Restore a cluster from backup

#### 8.7.4.3.1 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.

#### 8.7.4.3.2 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 731)), 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
```

### 8.7.4.3.3 Restore your DKP Management Cluster

When restoring a backup to the Management Cluster, you must adjust configuration to avoid restore errors.

#### 8.7.4.3.3.1 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"
 }
]'
```

#### 8.7.4.3.2 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"
 }
]'
```

#### 8.7.4.4 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
```

## 8.8 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).

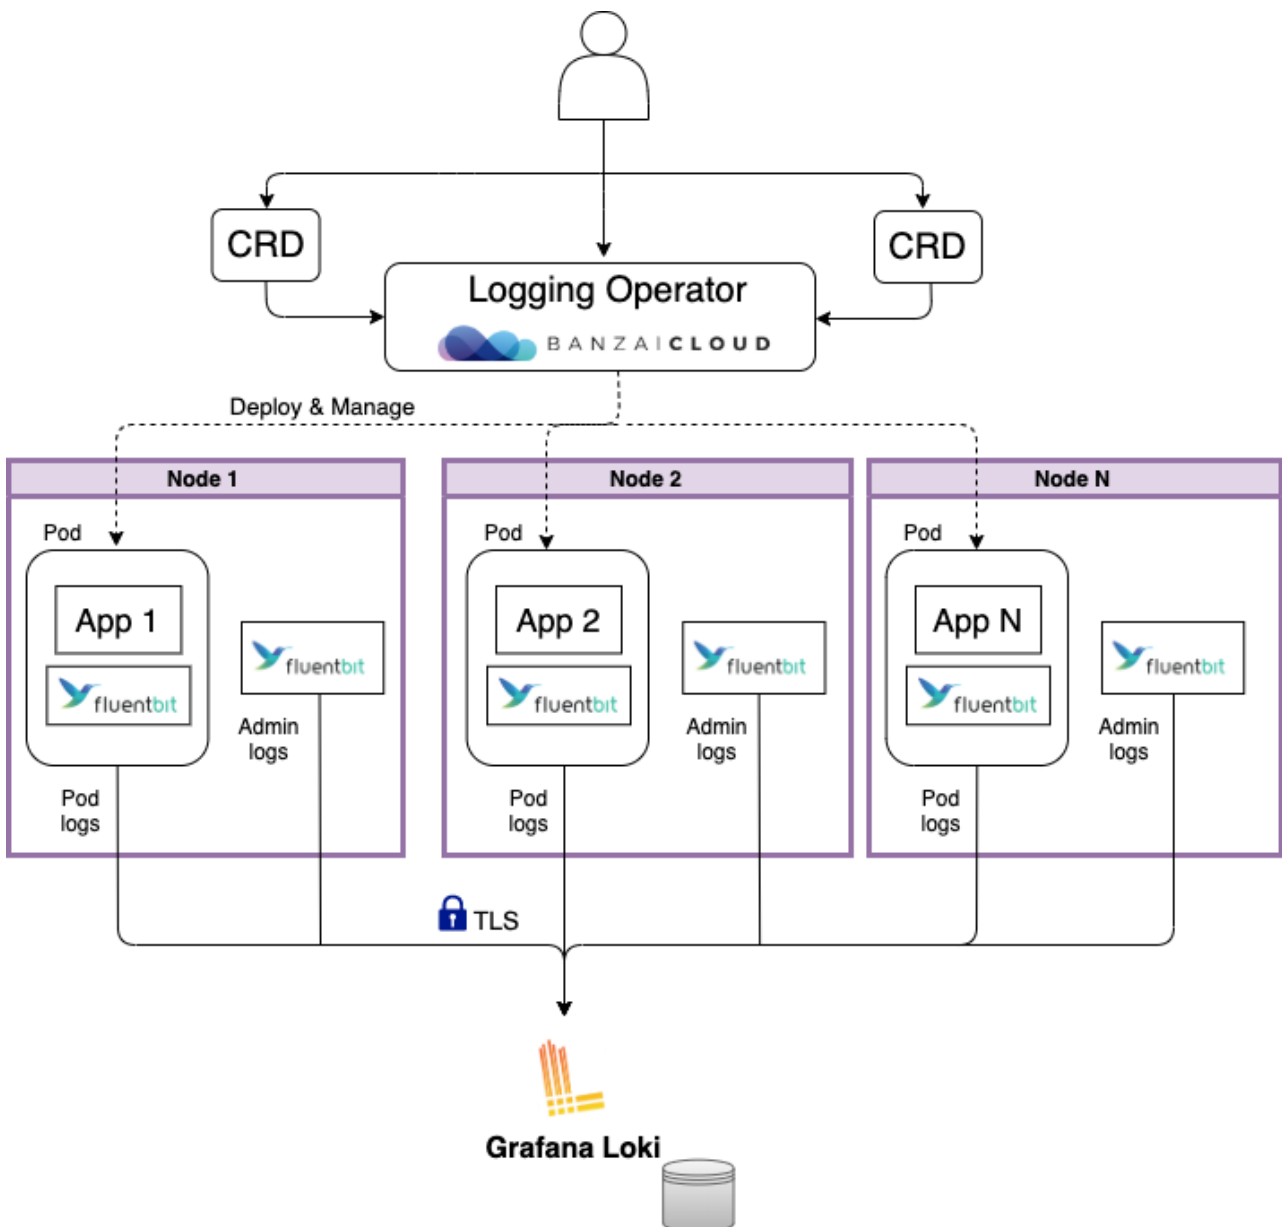
By default, logging is disabled on managed and attached clusters. You need to enable the logging stack applications explicitly on the workspace to make use of these capabilities.

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.

- [Admin-level logs](#) (see page 753)
- [Enable Workspace-level Logging](#) (see page 753)
- [Multi-Tenant Logging Overview](#) (see page 763)
- [Fluent Bit](#) (see page 773)
- [Configuring Loki to use AWS S3 Storage in DKP](#) (see page 778)

## 8.8.1 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 773) 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 858) for more details on how to configure the Ceph Cluster.

## 8.8.2 Enable Workspace-level Logging

### 8.8.2.1 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](#) (see page 754)
- [Enable Logging Applications through the UI](#) (see page 754)
- [Create AppDeployments to Enable Workspace Logging](#) (see page 755)
- [Override ConfigMap to Restrict Logging to Specific Namespaces](#) (see page 756)
- [Override ConfigMap to Modify the Storage Retention Period in Workspace Grafana Loki](#) (see page 758)
- [Verify Cluster Logging Stack Installation](#) (see page 760)
- [View Cluster Log Data](#) (see page 761)

## 8.8.2.2 Logging Prerequisites

Before you begin, you must:

- Be a cluster administrator with permissions to configure cluster-level platform services.
- Set a [default storage class](#) on each attached cluster for successful Loki deployment.


## 8.8.2.3 Enable Logging Applications through the UI

### 8.8.2.3.1 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 755).

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 721) page, or you can [verify the Cluster Logging Stack installation via the CLI](#) (see page 760).
9. Then, you can [view cluster log data](#) (see page 761).

 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 858) for more details on how to configure the Ceph Cluster.

## 8.8.2.4 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](#)<sup>529</sup> of the attached cluster for the following `kubectl` command.)

<sup>529</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

```
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-3.17.9 --
workspace ${WORKSPACE_NAME}
dkp create appdeployment rook-ceph --app rook-ceph-1.10.3 --workspace $
${WORKSPACE_NAME}
dkp create appdeployment rook-ceph-cluster --app rook-ceph-cluster-1.10.3 --
workspace ${WORKSPACE_NAME}
dkp create appdeployment grafana-loki --app grafana-loki-0.48.6 --workspace $
${WORKSPACE_NAME}
dkp create appdeployment grafana-logging --app grafana-logging-6.38.1 --
workspace ${WORKSPACE_NAME}
```

Then, you can [verify the cluster logging stack installation](#) (see page 760).



To deploy the applications to selected clusters within the workspace, refer to the [cluster-scoped configuration](#) (see page 584) 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 858) 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-0.20.9 --workspace $
${WORKSPACE_NAME}
```

### 8.8.2.5 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 \(see page 755\)](#) procedure.

### 8.8.2.5.1 Prerequisites

- Implement each of the steps listed in [Enable Workspace-level Logging \(see page 753\)](#).
- Ensure that log data is available before you execute this procedure.

### 8.8.2.5.2 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.

### 8.8.2.6 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](#)<sup>530</sup>. This retention policy can be customized.



The minimum retention period is 24h.

<sup>530</sup> <https://grafana.com/docs/loki/latest/operations/storage/boltdb-shipper/#compactor>

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](#)<sup>531</sup> 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
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
```

<sup>531</sup> <https://grafana.com/docs/loki/latest/operations/storage/retention/#retention-configuration>

```
kind: AppDeployment
metadata:
 name: grafana-loki
 namespace: ${WORKSPACE_NAMESPACE}
spec:
 appRef:
 name: grafana-loki-0.48.6
 kind: ClusterApp
 configOverrides:
 name: grafana-loki-custom-overrides
```

### 8.8.2.7 Verify Cluster Logging Stack Installation

 How to verify the cluster's logging stack installed successfully.

You must wait for the cluster's logging stack `HelmReleases` 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](#)<sup>532</sup> 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.

<sup>532</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>




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 761).

### 8.8.2.8 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:

```
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](#)<sup>533</sup> of the attached cluster for the following kubectl commands:

3. Get the Grafana URL:

<sup>533</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

```
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\''
```

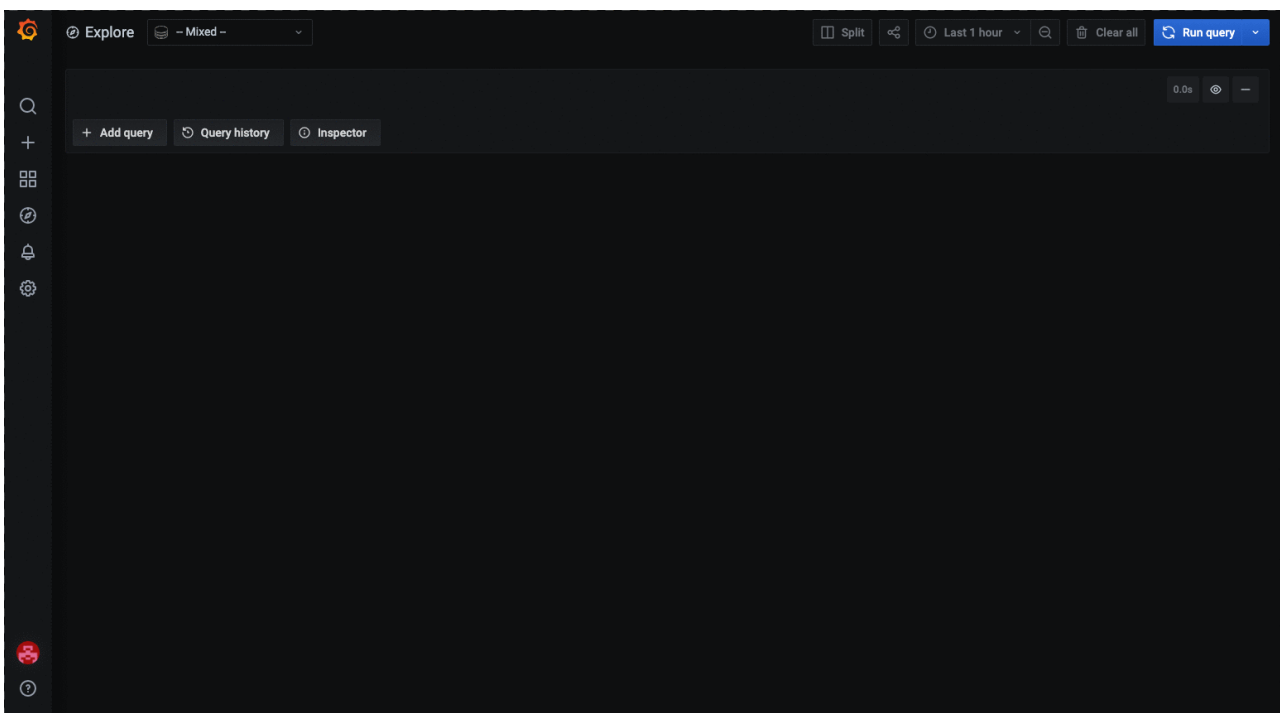
To view logs in Grafana:

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 781) for more information.
3. At the top of the page, change the datasource to `Loki`.

See the [Grafana Loki documentation](#)<sup>534</sup> for more on how to use the interface to view and query logs.



<sup>534</sup> <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.

### 8.8.3 Multi-Tenant Logging Overview

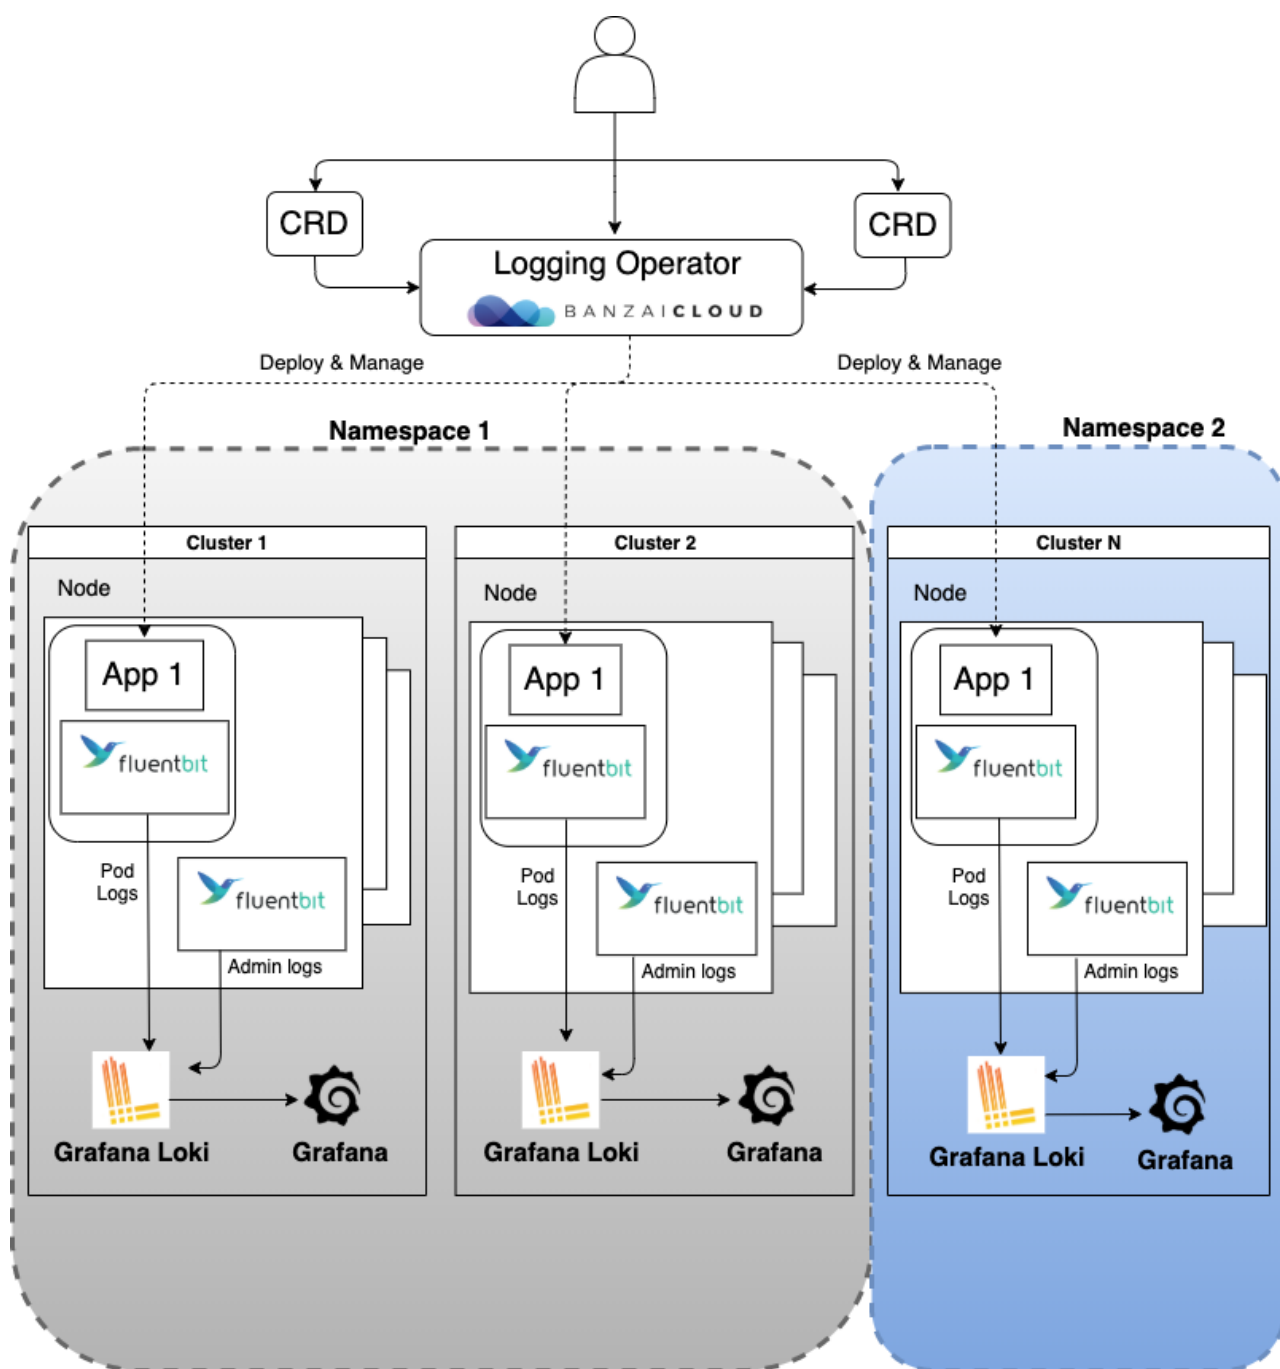
#### Enterprise

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 \(see page 753\)](#) 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 765)
- [Create a Project for Logging](#) (see page 765)
- [Create Project-level Logging AppDeployments](#) (see page 766)
- [Override ConfigMap to Modify the Storage Retention Period in Project Grafana Loki](#) (see page 767)
- [Verify Project Logging Stack Installation](#) (see page 769)
- [View Project Log Data](#) (see page 770)

### 8.8.3.1 Enable Multi-tenant Logging

Enterprise

#### 8.8.3.1.1 Prerequisites

Before you begin, you must:

- [Enable Workspace-level Logging](#) (see page 753) before you can configure multi-tenant logging.
- Be a cluster administrator with permissions to configure cluster level platform services.

#### 8.8.3.1.2 Multi-tenant Logging Enablement Process

The steps required to enable multi-tenant logging include:

1. [Create a Project](#). (see page 765)
2. [Create the required Project-level AppDeployments](#). (see page 766)
3. [Verify that the Project logging stack is installed](#). (see page 769)
4. [View a Project's log data](#). (see page 770)

Get started with multi-tenant logging by [Creating a Project](#) (see page 765).

### 8.8.3.2 Create a Project for Logging

Enterprise

To enable multi-tenant logging, you must first [Create a Project](#) (see page 615) 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 766).

### 8.8.3.3 Create Project-level Logging AppDeployments

#### Enterprise

#### 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>
```

- Copy these commands and execute them from a command line:

```
dkp create appdeployment project-grafana-loki --app project-grafana-loki-0.48.6
 --workspace ${WORKSPACE_NAME} --project ${PROJECT_NAME}
dkp create appdeployment project-grafana-logging --app project-grafana-
logging-6.38.1 --workspace ${WORKSPACE_NAME} --project ${PROJECT_NAME}
```

Then, you can [Verify the Project Logging Stack Installation](#) (see page 769) for multi-tenant logging.

### 8.8.3.4 Override ConfigMap to Modify the Storage Retention Period in Project Grafana Loki

#### Enterprise

DKP configures Project Grafana Loki to retain log metadata and logs for 1 week, using the [Compactor](#)<sup>535</sup>. 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:

- 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. This may NOT be identical to the Display Name of the `Workspace`.

- Set the `WORKSPACE_NAMESPACE` variable to the namespace copied in the previous step:

```
export WORKSPACE_NAMESPACE=<WORKSPACE_NAMESPACE>
```

- Get the namespace of your project:

```
kubectl get projects --namespace ${WORKSPACE_NAMESPACE}
```

<sup>535</sup> <https://grafana.com/docs/loki/latest/operations/storage/boltdb-shipper/#compactor>

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](#)<sup>536</sup> 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 -
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
```

<sup>536</sup> <https://grafana.com/docs/loki/latest/operations/storage/retention/#retention-configuration>



### 8.8.3.5 Verify Project Logging Stack Installation

#### Enterprise

#### How to verify the project logging stack installation for multi-tenant logging

You must wait for the Project's logging stack `HelMReleases` 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`<sup>537</sup> 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}
```

<sup>537</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

**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
<code>AGE</code>			
<code>\${PROJECT_NAMESPACE}</code>	<code>project-grafana-logging</code>	True	Release
	<code>reconciliation succeeded 15m</code>		
<code>\${PROJECT_NAMESPACE}</code>	<code>project-grafana-loki</code>	True	Release
	<code>reconciliation succeeded 11m</code>		
<code>\${PROJECT_NAMESPACE}</code>	<code>project-loki-object-bucket-claims</code>	True	Release
	<code>reconciliation succeeded 11m</code>		

Then, you can [View Project Log Data](#) (see page 770) within multi-tenant logging.

### 8.8.3.6 View Project Log Data

#### Enterprise

#### 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.

##### 8.8.3.6.1 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.

```
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](#)<sup>538</sup> 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"}}'
```

### 8.8.3.6.2 View Logs in Grafana

1. Go to the `Explore` tab:

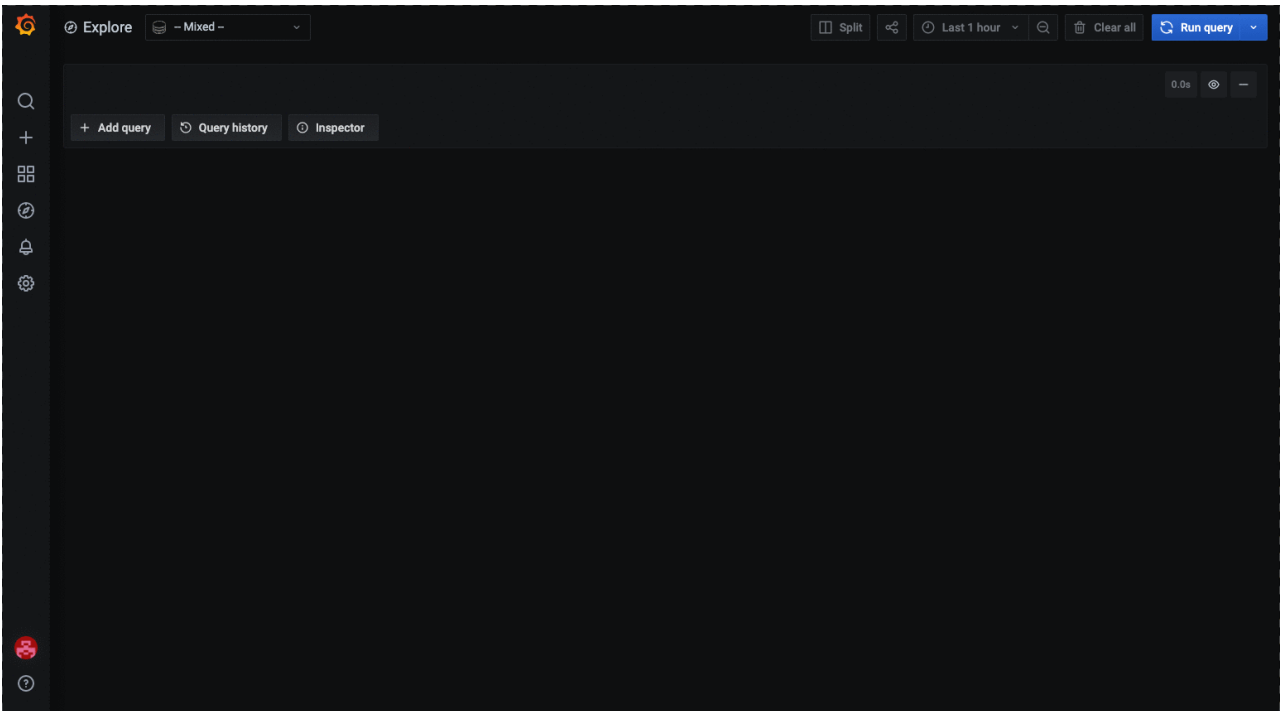
```
kubectl get ingress -n ${PROJECT_NAMESPACE} ${PROJECT_NAMESPACE}-project-
grafana-logging -o go-template='https://{{with
index .status.loadBalancer.ingress 0}}{{or .hostname .ip}}{{end}}{{with
index .spec.rules 0}}{{with index .http.paths 0}}{{.path }}{{end}}{{end}}/
explore{{"\n"}}'
```

2. You may be prompted to log in using the SSO flow. See [Authentication and Authorization](#) (see page 781) for more information.
3. At the top of the page, change the data source to `Loki`.

See the [Grafana Loki documentation](#)<sup>539</sup> for more on how to use the interface to view and query logs.

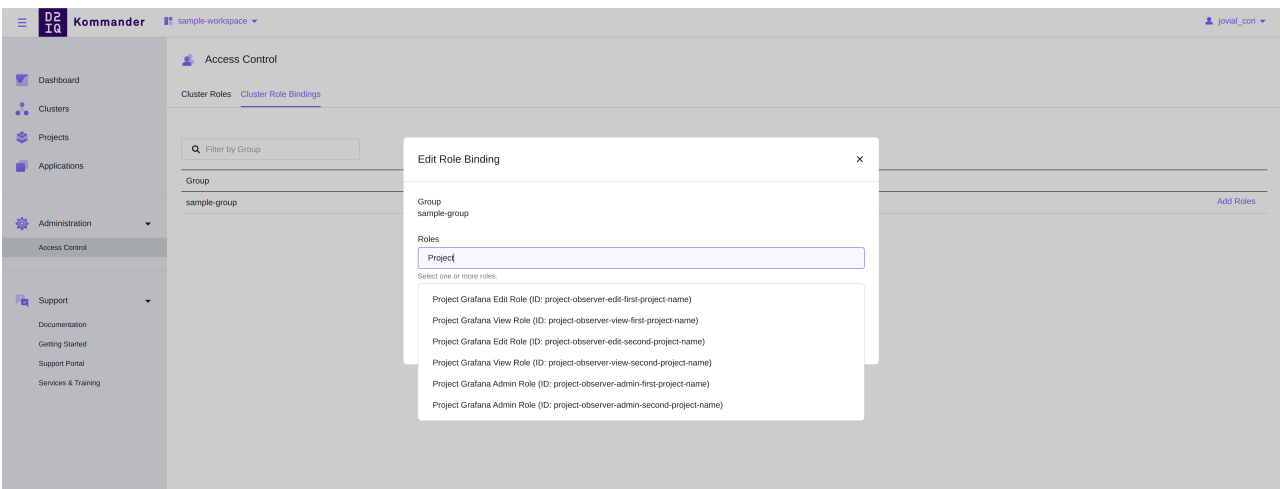
<sup>538</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

<sup>539</sup> <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.


You can [configure workspace policy \(see page 751\)](#) 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:



Select the correct role bindings for each group for a project at the workspace level.

## 8.8.4 Fluent Bit

[Fluent Bit](#)<sup>540</sup> 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 858](#)) for more details on how to configure the Rook Ceph Cluster.

### 8.8.4.1 Audit Log Collection

[Auditing](#)<sup>541</sup> 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](#)<sup>542</sup> configuration, you must modify the log retention settings by [Configuring the Control Plane](#) (see [page 466](#)) 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 773](#))

### 8.8.4.2 Related Information

For information on related topics or procedures, refer to the following:

- [Logging](#) (see [page 751](#))

### 8.8.4.3 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):

<sup>540</sup> <https://fluentbit.io/>

<sup>541</sup> <https://kubernetes.io/docs/tasks/debug-application-cluster/audit/>

<sup>542</sup> <https://kubernetes.io/docs/tasks/debug/debug-cluster/audit/#log-backend>

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 rescape 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 773) 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

```

13. 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.

## 8.8.5 Configuring Loki to use AWS S3 Storage in DKP

Follow the instructions on this page to configure Loki/Velero to use AWS S3 Storage in DKP

### 8.8.5.1 Configuring Loki

The easiest way to get started with using AWS S3 storage with Grafana Loki is to use a set of static AWS credentials.

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 can also be added to the installer configuration if you are configuring Grafana Loki on the Management Cluster.

```
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

```


5. 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

```

-  The reloader annotation only works in DKP 2.5.2+. Any changes to the AWS credential secret will not automatically reload some Loki pods. In this scenario, The ingester and compactor pods need to be manually restarted.
- For more information, refer to <https://github.com/grafana/helm-charts/issues/1905>.

## 8.9 Security

Security architecture.

- [Authentication and authorization architecture](#) (see page 781)
- [OpenID Connect \(OIDC\) Introduction](#) (see page 782)
- [SAML Connector](#) (see page 786)
- [Policy Controls](#) (see page 789)
- [Traefik-Forward-Authentication in DKP \(TFA\)](#) (see page 793)

## 8.9.1 Authentication and authorization architecture

### 8.9.1.1 Details on distributed authentication and authorization between clusters

### 8.9.1.2 Authentication

DKP UI comes with a pre-configured authentication [Dex](#)<sup>543</sup> 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](#)<sup>544</sup> 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](#)<sup>545</sup> platform service.



The `kube-oidc-proxy`<sup>546</sup> 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.

<sup>543</sup> <https://github.com/dexidp/dex>

<sup>544</sup> [https://openid.net/specs/openid-connect-core-1\\_0.html#IDToken](https://openid.net/specs/openid-connect-core-1_0.html#IDToken)

<sup>545</sup> <https://github.com/mesosphere/traefik-forward-auth>

<sup>546</sup> <https://github.com/jetstack/kube-oidc-proxy>

### 8.9.1.2.1 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 UI cluster Dex issued ID tokens.

When the `traefik-forward-auth` is used as a [Traefik Ingress authenticator](#)<sup>547</sup>, it checks if the user identity was issued by the Kommander cluster Dex service. An anonymous user is redirected to the UI cluster Dex service to authenticate and confirm their identity.

Never enter your own credentials on any of the attached clusters. On the UI cluster use the static admin credentials or an external identity provider (IDP).

### 8.9.1.3 Authorization

There is no centralized authorization component in Kommander. Each component and service makes its own authorization decisions based on user identity.

## 8.9.2 OpenID Connect (OIDC) Introduction

### 8.9.2.1 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.

### 8.9.2.2 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 use already use one or more of the following IdPs, you can configure Dex to use them:

---

<sup>547</sup> <https://docs.traefik.io/v2.4/providers/kubernetes-ingress/>

Name	Supports Refresh Tokens	Supports Groups Claim	Supports preferred_username Claim	Status	Notes
<a href="#">LDAP</a> <sup>548</sup>	yes	yes	yes	stable	
<a href="#">GitHub</a> <sup>549</sup>	yes	yes	yes	stable	
<a href="#">SAML 2.0</a> <sup>550</sup>	no	yes	no	stable	
<a href="#">GitLab</a> <sup>551</sup>	yes	yes	yes	beta	
<a href="#">OpenID Connect</a> <sup>552</sup>	yes	yes	yes	beta	Includes Salesforce, Azure, etc.
<a href="#">Google</a> <sup>553</sup>	yes	yes	yes	alpha	
<a href="#">LinkedIn</a> <sup>554</sup>	yes	no	no	beta	
<a href="#">Microsoft</a> <sup>555</sup>	yes	yes	no	beta	
<a href="#">AuthProxy</a> <sup>556</sup>	no	no	no	alpha	Authentication proxies such as Apache2 mod_auth, etc.
<a href="#">Bitbucket Cloud</a> <sup>557</sup>	yes	yes	no	alpha	

548 <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/ldap.md>

549 <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/github.md>

550 <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/saml.md>

551 <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/gitlab.md>

552 <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/oidc.md>

553 <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/google.md>

554 <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/linkedin.md>

555 <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/microsoft.md>

556 <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/authproxy.md>

557 <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/bitbucketcloud.md>

Name	Supports Refresh Tokens	Supports Groups Claim	Supports preferred_username Claim	Status	Notes
<a href="#">OpenShift</a> <sup>558</sup>	no	yes	no	stable	

**NOTE:** These are the Identity Providers supported by Dex [2.22.0](#)<sup>559</sup>, the version used by DKP.

### 8.9.2.3 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](#)<sup>560</sup> 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
 redirectURI: https://DKP_CLUSTER_DOMAIN/dex/callback
 scopes:
 - openid
 - profile
 - email
 insecureSkipEmailVerified: true
 insecureEnableGroups: true
 userIDKey: email
 userNameKey: email
```

[...]

<sup>558</sup> <https://github.com/dexidp/dex/blob/v2.22.0/Documentation/connectors/openshift.md>

<sup>559</sup> <https://github.com/dexidp/dex/blob/v2.22.0/README.md>

<sup>560</sup> <https://dexidp.io/docs/connectors/>



### 8.9.2.4 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](#)<sup>561</sup> 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"
[...]
```

### 8.9.2.5 Authentication

OpenID Connect is an extension of the [OAuth2 authentication protocol](#)<sup>562</sup>. 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.

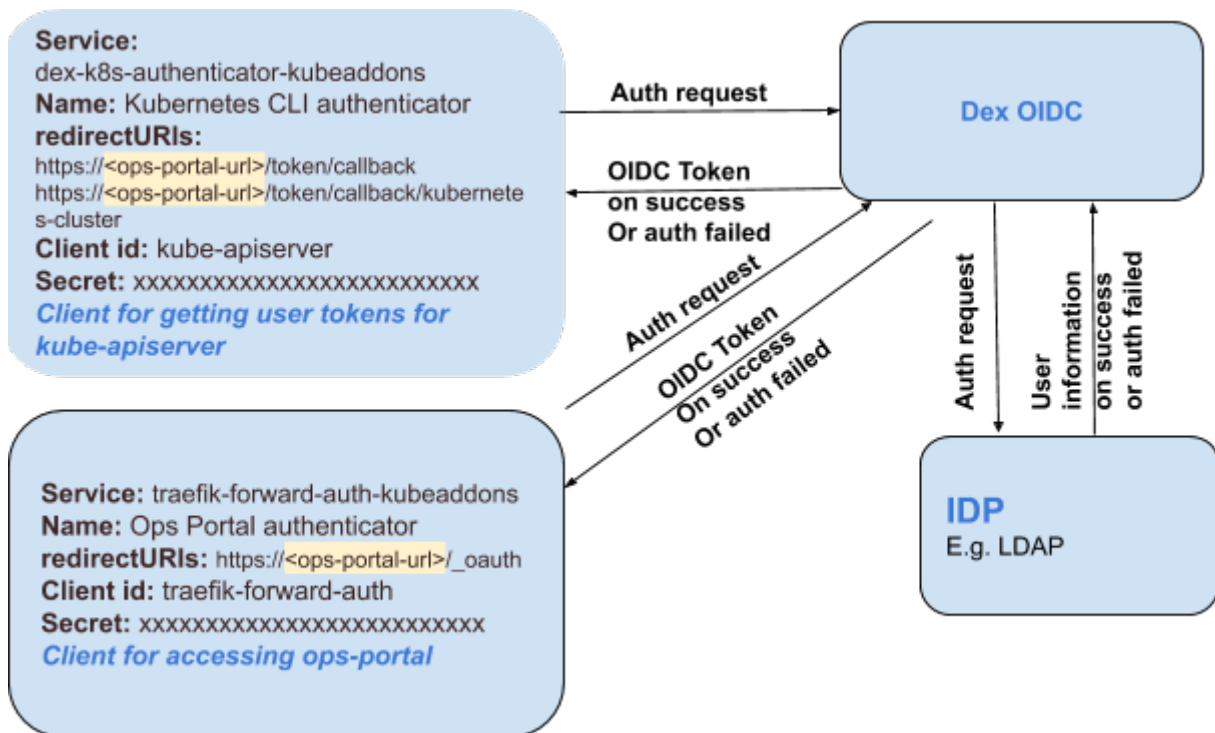
In Kommander, Dex comes pre-configured with a client for these access use cases. The clients talk to Dex for authentication. Dex talks to the configured Identity Provider, or IdP, (for example LDAP, SAML, etc) to perform the actual task of authenticating the user.

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:

<sup>561</sup> <https://dexidp.io/docs/id-tokens/#expiration-and-rotation-settings>

<sup>562</sup> <https://oauth.net/2/>



### 8.9.3 SAML Connector

#### 8.9.3.1 Connect your Kommander cluster to an IdP using SAML

#### 8.9.3.2 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 \(see page 121\)](#).
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 475) 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 529) 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.

## 8.9.4 Policy Controls

### 8.9.4.1 Workload Policy Controls

### 8.9.4.2 Enforce Policies Using Gatekeeper

#### 8.9.4.2.1 Learn how to enforce policies using Gatekeeper

[Gatekeeper](#)<sup>563</sup> is the policy controller for Kubernetes, allowing organizations to enforce configurable policies using the [Open Policy Agent](#)<sup>564</sup>, 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](#)<sup>565</sup>:

- [Prevent the running of privileged pods](#) (see page 0)
- [Prevent mounting host path volumes](#) (see page 0)

#### 8.9.4.2.2 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](#)<sup>566</sup> topic.
- If you install Kommander with a custom configuration, make sure you enabled Gatekeeper.

#### 8.9.4.2.3 Use Gatekeeper

Gatekeeper uses the [OPA Constraint Framework](#)<sup>567</sup> 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](#)<sup>568</sup> to replace Pod Security Policies which you will use in the following tutorials.

---

563 <https://github.com/open-policy-agent/gatekeeper>

564 <https://github.com/open-policy-agent/opa>

565 <https://kubernetes.io/docs/concepts/policy/pod-security-policy/>

566 <https://docs.d2iq.com/dkp/konvoy/2.2/choose-infrastructure/>

567 <https://github.com/open-policy-agent/frameworks/tree/master/constraint>

568 <https://github.com/open-policy-agent/gatekeeper-library/tree/master/library/pod-security-policy>

### 8.9.4.2.3.1 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:

```
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}
```

### 8.9.4.2.3.2 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
```

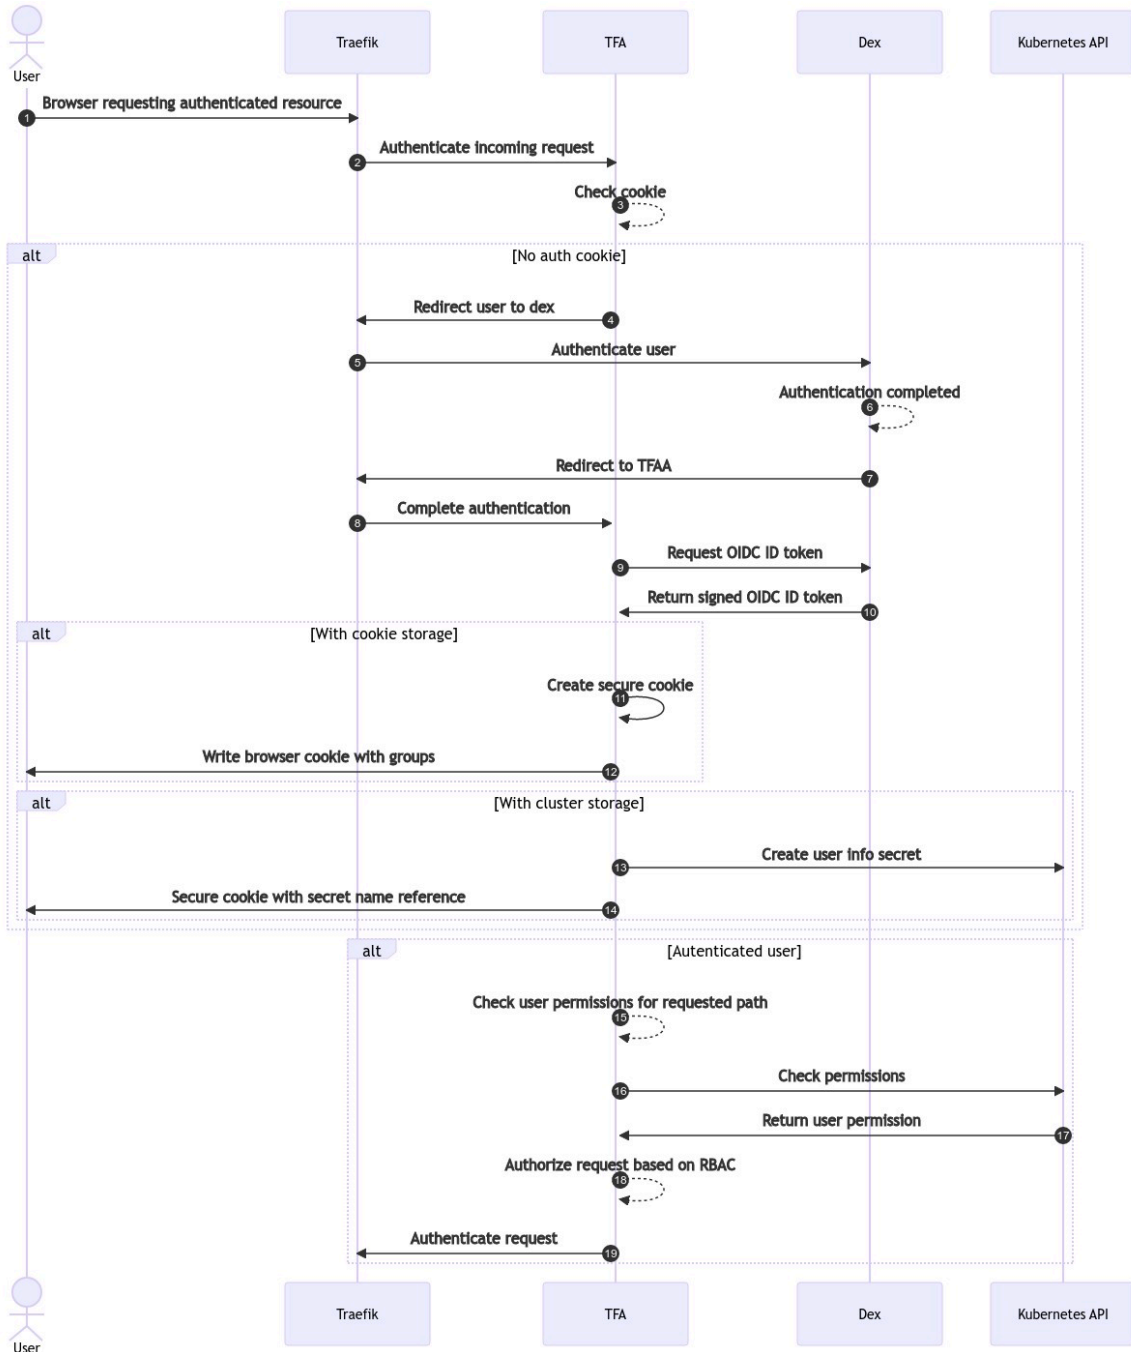
## 8.9.5 Traefik-Forward-Authentication in DKP (TFA)

### 8.9.5.1 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.

### 8.9.5.2 TFA Authentication Workflow



### 8.9.5.3 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.

### 8.9.5.4 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](#)<sup>569</sup>.

---

<sup>569</sup> <https://github.com/mesosphere/traefik-forward-auth>

## 8.10 Networking

### 8.10.1 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 796)
- [Required Domains](#) (see page 802)
- [Configure Ingress for load balancing](#) (see page 802)
- [Ingress](#) (see page 805)
- [Load Balancing](#) (see page 806)
- [External DNS](#) (see page 807)
- [Use Istio as a Microservice](#) (see page 808)

### 8.10.2 Networking Service

A [Service](#)<sup>570</sup> 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
 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.

---

<sup>570</sup> <https://kubernetes.io/docs/concepts/services-networking/service/#service-resource>

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.

### 8.10.2.1 Service Topology

[Service Topology](#)<sup>571</sup> 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"
```

- If the value of the `topologyKeys` field does not match any pattern, the traffic is rejected.

<sup>571</sup> <https://kubernetes.io/docs/concepts/services-networking/service-topology/#using-service-topology>

## 8.10.2.2 EndpointSlices

[EndpointSlices](#)<sup>572</sup> 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
```

## 8.10.2.3 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.

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`.

---

<sup>572</sup> <https://kubernetes.io/docs/concepts/services-networking/endpoint-slices/#endpointslice-resource>

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](#)<sup>573</sup>.

### 8.10.2.4 Ingress and Networking

[Ingress](#)<sup>574</sup> 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:
 paths:
 - path: /path
 pathType: Prefix
 backend:
 service:
 name: my-konvoy-service
 port:
```

<sup>573</sup> <https://kubernetes.io/docs/concepts/services-networking/dns-pod-service/>

<sup>574</sup> <https://kubernetes.io/docs/concepts/services-networking/ingress/#what-is-ingress>

number: 80

In Kommander, you can expose services to the outside world using Ingress objects.

#### 8.10.2.4.1 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](#)<sup>575</sup>:

- [HAProxy Ingress](#)<sup>576</sup> is a highly customizable community-driven ingress controller for HAProxy.
- NGINX offers support and maintenance for the [NGINX Ingress Controller](#)<sup>577</sup> for Kubernetes.
- [Traefik](#)<sup>578</sup> is a fully featured Ingress controller (Let's Encrypt, secrets, http2, websocket), and has commercial support.
- [Ambassador API Gateway](#)<sup>579</sup> EXPERIMENTAL is an Envoy based Ingress controller with community and commercial support.

In Kommander, `Traefik` deploys by default as a well-suited Ingress controller.

#### 8.10.2.5 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:

```
apiVersion: networking.k8s.io/v1
kind: NetworkPolicy
metadata:
 name: network-konvoy-policy
 namespace: default
```

<sup>575</sup> <https://kubernetes.io/docs/concepts/services-networking/ingress-controllers/#additional-controllers>

<sup>576</sup> <https://haproxy-ingress.github.io/>

<sup>577</sup> <https://www.nginx.com/products/nginx-ingress-controller>

<sup>578</sup> <https://github.com/containous/traefik>

<sup>579</sup> <https://www.getambassador.io/>



```

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).

### 8.10.2.5.1 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"
```

## 8.10.3 Required Domains

### 8.10.3.1 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.com](http://mcr.microsoft.com)<sup>580</sup>
- [nvcr.io](http://nvcr.io)
- <http://quay.io>
- <http://us.gcr.io>



In an air-gapped installation, these domains do not need to be accessible.

## 8.10.4 Configure Ingress for load balancing

### 8.10.4.1 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.

---

<sup>580</sup> <http://mcr.microsoft.com>

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.

This tutorial demonstrates how to expose an application running on the Konvoy cluster by configuring an Ingress for load balancing (layer-7).

### 8.10.4.2 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.

### 8.10.4.3 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 illustrates:

- 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 PORT(S)	TYPE	CLUSTER-IP AGE	EXTERNAL-IP
kommander-traefik	LoadBalancer	10.0.24.215	
		abf2e5bda6ca811e982140acb7ee21b7-37522315.us-west-2.elb.amazonaws.com	80:31169/TCP,443:32297/TCP,8080:31923/TCP
		4h22m	

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

```

## 8.10.5 Ingress

### 8.10.5.1 Traefik Ingress Controller

Kubernetes Ingress resources expose HTTP and HTTPS routes from outside the cluster to services within the cluster. In Kommander, the [Traefik](#)<sup>581</sup> 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.

### 8.10.5.2 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
- 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

---

<sup>581</sup> <https://landscape.cncf.io/card-mode?category=service-proxy&grouping=category&selected=traefik>

### 8.10.5.3 Related Information

For information on related topics or procedures, refer to the following:

- [List of Ingress controllers](#)<sup>582</sup>
- [Traefik v2.4 docs](#)<sup>583</sup>
- [Configure Ingress for load balancing](#) (see page 802)
- [Load Balancing](#) (see page 806)

## 8.10.6 Load Balancing

In a Kubernetes cluster, depending on the flow of traffic direction, there are two kinds of load balancing:

- Load balancing for the traffic within a Kubernetes cluster
- Load balancing for the traffic coming from outside the cluster

### 8.10.6.1 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`<sup>584</sup> rules to distribute direct traffic to real servers, whereas `IPVS`<sup>585</sup> leverages in-kernel transport-layer load-balancing. Read a [comparison between these two methods](#)<sup>586</sup>. 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.

### 8.10.6.2 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.



In DKP environments, the external load balancer must be configured without TLS termination.

<sup>582</sup> <https://kubernetes.io/docs/concepts/services-networking/ingress-controllers/>

<sup>583</sup> <https://doc.traefik.io/traefik/v2.4/>

<sup>584</sup> [https://www.linuxtopia.org/Linux\\_Firewall\\_iptables/x4013.html](https://www.linuxtopia.org/Linux_Firewall_iptables/x4013.html)

<sup>585</sup> [https://en.wikipedia.org/wiki/IP\\_Virtual\\_Server](https://en.wikipedia.org/wiki/IP_Virtual_Server)

<sup>586</sup> <https://www.projectcalico.org/comparing-kube-proxy-modes-iptables-or-ipvs/>

### 8.10.6.2.1 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.

### 8.10.6.2.2 On-premises

For an on-premises deployment, DKP ships with [MetalLB](#)<sup>587</sup>, which provides load-balancing services. The environments that use MetalLB are pre-provisioned and vSphere infrastructures.

For more information on how to configure MetalB for these environments, refer to the following pages:

- [Pre-provisioned Configure MetalLB \(see page 316\)](#)
- [Configure MetalLB for a vSphere infrastructure \(see page 352\)](#)

### 8.10.6.2.3 Custom Load Balancer for External Traffic

If you want to use a non-DKP load balancer for external traffic, see [Install Kommander with an External Load Balancer \(see page 501\)](#).

## 8.10.7 External DNS

**This section describes how to use `external-dns` to maintain your DNS records**

You require a DNS record when setting up a [custom domain for your cluster \(see page 723\)](#). 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.

The following example shows how to configure `external-dns` to manage DNS records in AWS Route 53 automatically:

1. Open the `kommander.yaml` file:
  - a. If you do not have a `kommander.yaml` file, [initialize the configuration file \(see page 475\)](#), 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. Adjust the `app` section of your `kommander.yaml` file to include these values:

---

<sup>587</sup> <https://metallb.universe.tf/concepts/>

```

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>

```

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>

```

Refer to the [external-dns documentation](#)<sup>588</sup> 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.

## 8.10.8 Use Istio as a Microservice

 EXPERIMENTAL

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.

### 8.10.8.1 Prerequisites

Before you begin, verify the following:

---

<sup>588</sup> <https://artifacthub.io/packages/helm/bitnami/external-dns>



- 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.
- Set the `WORKSPACE_NAME` environment variable to the name of the workspace where the cluster is attached:

```
export WORKSPACE_NAME=<workspace_name>
```

### 8.10.8.2 Deploy Istio Using DKP

#### EXPERIMENTAL

1. Review the [list of available applications \(see page 1064\)](#) to obtain the current *APP ID* and *version* for Istio and its [dependencies \(see page 579\)](#), as you will need this information to execute the following commands.
2. Install [Istio's dependencies \(see page 579\)](#) with an `AppDeployment` resource. Replace the `<APPID>` and `<APPID-version>` variables with the information of the application:

**Note:** 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. Enable the deployment of Istio to your managed or [attached cluster \(see page 678\)](#) with an `AppDeployment` resource.

```
dkp create appdeployment istio --app istio-1.15.3 --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.

### 8.10.8.2.1 Download the Istio Command Line Utility

1. Pull a copy of the corresponding Istio command line to your system:

```
curl -L https://istio.io/downloadIstio | ISTIO_VERSION=1.15.3 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)
```

### 8.10.8.2.2 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](#)<sup>589</sup>.

<sup>589</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

```
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:

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

3. 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
```

4. Follow the steps in the Istio [BookInfo Application](#)<sup>590</sup> documentation to understand the different Istio features.

For more information, see [Istio's official documentation](#)<sup>591</sup>.

## 8.11 GPUs

This section describes NVIDIA GPU support on Kommander. DKP supported [NVIDIA driver](#)<sup>592</sup> 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 214) section.

<sup>590</sup> <https://istio.io/docs/examples/bookinfo/>

<sup>591</sup> <https://istio.io/latest/docs/>

<sup>592</sup> <https://www.nvidia.com/Download/Find.aspx>

## 8.11.1 Kommander GPU Overview

GPU support on Kommander uses the [NVIDIA](#)<sup>593</sup> 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.
- [NVIDIA Device Plugin](#)<sup>594</sup>: 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](#)<sup>595</sup>: 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](#)<sup>596</sup> page for NVIDIA.

## 8.11.2 Kommander GPU configuration

### 8.11.2.1 Configure GPU for Kommander clusters

### 8.11.2.2 Prerequisites

Before you begin, you must:

- Ensure nodes provide an NVIDIA GPU.
- If you are using a public cloud service such as AWS, create an AMI with KIB using the instructions on the [KIB for GPU](#) (see page 427) page.
- If you are deploying in a pre-provisioned environment, ensure that you have created the appropriate secret for your GPU nodepool 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.


---

593 <https://github.com/NVIDIA/gpu-operator>

594 <https://github.com/NVIDIA/k8s-device-plugin>

595 <https://developer.nvidia.com/dcgm>

596 <https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/getting-started.html#chart-customization-options>

-  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 [Enable NVIDIA Platform Application on Kommander \(Management Cluster\)](#) (see page 813) section.
  - For instructions on enabling the NVIDIA platform application on attached or managed clusters, follow the instructions in the [Enable NVIDIA Platform Application on Attached or Managed Clusters](#) (see page 813) section.

Once `nvidia-gpu-operator` has been enabled depending on the cluster type, proceed to the [Select the correct Toolkit version for your NVIDIA GPU Operator](#) (see page 814) section.

### 8.11.2.3 Enable NVIDIA Platform Application on Kommander (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
```

4. Proceed to the [Select the correct Toolkit version for your NVIDIA GPU Operator](#) (see page 814) section.

### 8.11.2.4 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/143663151/2.4%2BPlatform%2BApplications#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 570) 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 586) 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.

### 8.11.2.5 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.

#### 8.11.2.5.1 Kommander (Management Cluster) Customization

1. Select the correct Toolkit version based on your OS:

##### **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:

```
kind: Installation
apps:
 nvidia-gpu-operator:
 values: |
 toolkit:
 version: v1.10.0-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 `install.yaml` to the following:

```
kind: Installation
apps:
 nvidia-gpu-operator:
```

```
values: |
 toolkit:
 version: v1.10.0-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 `install.yaml` to the following:

```
kind: Installation
apps:
 nvidia-gpu-operator:
 values: |
 toolkit:
 version: v1.11.0-ubuntu20.04
```

2. Install Kommander, using the configuration file you created:

```
dkp install kommander --installer-config ./install.yaml
```

### 8.11.2.5.2 Workspace (Attached and Managed clusters) Customization



Refer to [AppDeployment resources \(see page 561\)](#) 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 \(see page 588\)](#) 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.



### 8.11.2.6 Validate that the Application has Started Correctly

Run the following command to validate that your application has started correctly:

```
kubectll 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.

### 8.11.2.7 NVIDIA GPU Monitoring

Kommander uses the [NVIDIA Data Center GPU Manager](https://developer.nvidia.com/dcgm)<sup>597</sup> 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.

### 8.11.2.8 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.

<sup>597</sup> <https://developer.nvidia.com/dcgm>

**NOTE:** MIG is only available for the following NVIDIA devices: H100, A100, and A30.

### 8.11.2.8.1 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.

### 8.11.2.9 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.
|
```



```

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

```

## 8.11.2.10

### 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 317)

## 8.12 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 823)
- [Grafana Dashboards](#) (see page 825)
- [Cluster Metrics](#) (see page 827)


- [Configure Alerts Using AlertManager](#) (see page 828)
- [Centralized Monitoring](#) (see page 833)
- [Centralized Cost Monitoring](#) (see page 837)
- [Monitor Applications using Prometheus](#) (see page 839)
- [Set Storage Capacity for Prometheus](#) (see page 841)

## 8.12.1 Recommendations

### Enterprise

#### 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.

### 8.12.1.1 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>
-----	-----	----	--------------------------------------------------------------------------------------------------------------------------------------

## 8.12.2 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](#)<sup>598</sup>.

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:

```
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:

<sup>598</sup> <https://github.com/prometheus-community/helm-charts/tree/main/charts/kube-prometheus-stack/templates/grafana/dashboards-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 588) 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`.

### 8.12.2.1 Add Custom Dashboards

In Kommander you can define your own custom dashboards. You can use a few methods to [import dashboards](#)<sup>599</sup> to Grafana.

One method is to [use ConfigMaps to import dashboards](#)<sup>600</sup>. 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:

```
{
 "annotations": {
 "list": []
 },
 "description": "etcd sample Grafana dashboard with Prometheus",
 "editable": true,
 "gnetId": null,
 "hideControls": false,
 "id": 6,
 "links": [],
 "refresh": false,
 ...
}
```

<sup>599</sup> <https://github.com/grafana/helm-charts/tree/main/charts/grafana#import-dashboards>

<sup>600</sup> <https://github.com/grafana/helm-charts/tree/main/charts/grafana#sidecar-for-dashboards>

```
}

```

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](#)<sup>601</sup>:

```
kubectl apply -f etcd-custom-dashboard.yaml

```

### 8.12.3 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.

<sup>601</sup> <https://github.com/grafana/helm-charts/tree/main/charts/grafana#sidecar-for-dashboards>

- 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](#)<sup>602</sup>. 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.

## 8.12.4 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
- KubeAPIDown
- CoreDNSDown
- KubeVersionMismatch

A complete list with all the pre-defined alerts can be found [on GitHub](#)<sup>603</sup>.

### 8.12.4.1 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:

<sup>602</sup> <https://github.com/kubernetes/kube-state-metrics/tree/master/docs#exposed-metrics>

<sup>603</sup> <https://github.com/prometheus-community/helm-charts/tree/main/charts/kube-prometheus-stack/templates/prometheus/rules-1.14>

```
export WORKSPACE_NAMESPACE=<workspace_namespace>
```

### 8.12.4.1.1 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:

```
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 828) 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
```

- 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
```

- 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
```

- 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](#)<sup>604</sup>.

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`.

#### 8.12.4.1.2 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:
 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'

```

<sup>604</sup> <https://prometheus.io/docs/alerting/configuration/>

```

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 }}

```

- 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 -
```

## 8.12.5 Centralized Monitoring

### Enterprise

#### 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:

```
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}'
```

### 8.12.5.1 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](#)<sup>605</sup> 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 an attached cluster goes down, Kommander no longer collects or presents its metrics, including past data.

You can access the centralized Grafana UI at:

<sup>605</sup> <https://thanos.io/v0.5/components/query/#query>

```
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.

### 8.12.5.1.1 Adding Custom Dashboards

You can also define custom dashboards for centralized monitoring on Kommander. There are a few methods to [import dashboards](#)<sup>606</sup> 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: |
```

<sup>606</sup> <https://github.com/mesosphere/charts/tree/master/stable/grafana#import-dashboards>

```
{
 "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
```


### 8.12.5.2 Centralized Alerts

A centralized view of alerts, from attached clusters, is provided using an alert dashboard called [Karma](#)<sup>607</sup>. 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.

<sup>607</sup> <https://github.com/primitive/karma>

### 8.12.5.2.1 Federating Prometheus Alerting Rules

You can define additional [Prometheus alerting rules](#)<sup>608</sup> 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](#)<sup>609</sup>.

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
```

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](#)<sup>610</sup>.

```
kubectl get federatedprometheusrules kube-prometheus-stack-alertmanager.rules -n kommander -oyaml
```

608 [https://prometheus.io/docs/prometheus/latest/configuration/alerting\\_rules/](https://prometheus.io/docs/prometheus/latest/configuration/alerting_rules/)

609 <https://github.com/kubernetes-sigs/kubefed/blob/master/docs/installation.md#kubefedctl-cli>


610 <https://github.com/kubernetes-sigs/kubefed/blob/master/docs/userguide.md#propagation-status>

## 8.12.6 Centralized Cost Monitoring

### Enterprise

#### Monitoring costs of all attached clusters with Kubecost

[Kubecost](#)<sup>611</sup>, 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.

### 8.12.6.1 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](#)<sup>612</sup>.

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}'
```

<sup>611</sup> <https://kubecost.com/>

<sup>612</sup> <https://docs.kubecost.com/using-kubecost/getting-started/cost-allocation>

### 8.12.6.1.1 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.

#### 8.12.6.1.1.1 AWS

For more accurate AWS Spot pricing, follow [these steps](#)<sup>613</sup> to configure a data feed for the AWS Spot instances.

To allocate out of cluster costs for AWS, visit [this guide](#)<sup>614</sup>.

### 8.12.6.1.2 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`.

## 8.12.6.2 Related Information

For information on related topics or procedures, refer to the following:

- [Kubecost documentation](#)<sup>615</sup>

<sup>613</sup> <https://docs.kubecost.com/using-kubecost/getting-started/spot-checklist#implementing-spot-nodes-in-your-cluster>

<sup>614</sup> <https://docs.kubecost.com/install-and-configure/advanced-configuration/cloud-integration/aws-cloud-integrations/aws-out-of-cluster>

<sup>615</sup> <https://docs.kubecost.com/>

## 8.12.7 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`
- `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](#)<sup>616</sup>.

## 8.12.8 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:
 values.yaml: |

 prometheus:
 prometheusSpec:
 resources:
 limits:
 cpu: "4"
 memory: "8Gi"
 requests:
 cpu: "2"
 memory: "6Gi"
 storageSpec:
 volumeClaimTemplate:
 spec:
 resources:
 requests:
 storage: "100Gi"
```

## 8.13 DKP Troubleshooting

The following pages provide the steps to troubleshoot a Kubernetes cluster installation.

- [Generate a Support Bundle](#) (see page 841)
- [Custom Collectors](#) (see page 846)

### 8.13.1 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.

---

<sup>616</sup> <https://github.com/coreos/prometheus-operator/blob/master/Documentation/user-guides/getting-started.md#related-resources>


### 8.13.1.1 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 \(see page 919\)](#).

### 8.13.1.2 Create a 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 [on this public GitHub repository](#)<sup>617</sup>.

`dkp diagnose` supports [multiple support bundle collectors](#)<sup>618</sup> 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](#)<sup>619</sup> collects basic information about the cluster
- [clusterResources](#)<sup>620</sup> collects a subset of available resources in the cluster
- [configMap](#)<sup>621</sup> collects the values of Kubernetes ConfigMaps
- [secrets](#)<sup>622</sup> collects the values of Kubernetes ConfigMaps
- [execCopyFromHost](#) (see page 0) runs a container on each node on the cluster and copies the created data
- [allLogs](#) (see page 846) is capable of collecting logs from all containers on the cluster

617 <https://github.com/mesosphere/troubleshoot>

618 <https://troubleshoot.sh/docs/collect/all/>

619 <https://troubleshoot.sh/docs/collect/cluster-info/>

620 <https://troubleshoot.sh/docs/collect/cluster-resources/>

621 <https://troubleshoot.sh/docs/collect/configmap/>

622 <https://troubleshoot.sh/docs/collect/secret/>

### 8.13.1.3 Generate a Support Bundle

- 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 -xzf 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
```

#### 8.13.1.3.1 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
```

### 8.13.1.3.2 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`.

### 8.13.1.4 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](#)<sup>623</sup>.

 All other group definitions in the `inventory.yaml` file are ignored.

Refer to the following example file:

```
all:
```

<sup>623</sup> [https://docs.ansible.com/ansible/latest/user\\_guide/intro\\_inventory.html#connecting-to-hosts-behavioral-inventory-parameters](https://docs.ansible.com/ansible/latest/user_guide/intro_inventory.html#connecting-to-hosts-behavioral-inventory-parameters)

```
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 919)

## 8.13.2 Custom Collectors

For creating diagnostic bundles, D2iQ is using a customized version of `troubleshoot.sh` integrated into `dkp-diagnose`.

### 8.13.2.1 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](#)<sup>624</sup> of the changes.

#### 8.13.2.1.1 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 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:

---

<sup>624</sup> <https://github.com/mesosphere/troubleshoot/blob/v0.13-d2iq/README.d2iq.md>

```
<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
│ │ └── ...
│ ├── whoami_validate
│ ├── 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.

### 8.13.2.1.2 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
| | | |----- ...
| | | |----- whoami_validate
| | |----- 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.

### 8.13.2.1.3 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: [""]` ).

#### 8.13.2.1.4 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.

#### 8.13.2.1.5 ClusterResources Collector

Another customization is added to collect custom resource definitions and all custom resources in the cluster.

## 8.14 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:

#### 8.14.1 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.

## 8.14.2 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.

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).

### 8.14.2.1 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](#)<sup>625</sup> (for example, they can be mounted once 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.

### 8.14.2.2 Related Information

- [Kubernetes Storage](#)<sup>626</sup>
- [Kubernetes persistent storage design document](#)<sup>627</sup>

## 8.14.3 Provision a Static Local Volume

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

The `localvolumeprovisioner` component uses the [local volume static provisioner](#)<sup>628</sup> 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.

---

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

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

627 <https://github.com/kubernetes/community/blob/master/contributors/design-proposals/storage/persistent-storage.md>

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

- 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`.

### 8.14.3.1 Before you Begin

Before starting this tutorial, 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.

For this tutorial, you **do not deploy** using all the default settings as described in the [Quick start guide \(see page 47\)](#).

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.

### 8.14.3.2 Provision the Cluster and a Volume

1. Create a pre-provisioned cluster by following the steps outlined in the [Pre-provisioned Infrastructure \(see page 288\)](#) 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.

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
 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		

NAME	CAPACITY	ACCESS MODES	RECLAIM POLICY	STATUS
CLAIM	STORAGECLASS	REASON	AGE	
local-pv-4c7fc8ba	3986Mi	RWO	Delete	Bound
<b>default</b> /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.

## 8.14.4 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](#)<sup>629</sup> for [dynamic persistent volume \(PV\)](#)<sup>630</sup> creation. The table below lists the driver and default StorageClass for each supported cloud provisioner.

Cloud Provisioner	Driver	Default Storage Class
AWS	<a href="#">aws-ebs-csi-driver</a> <sup>631</sup>	ebs-sc
Azure	<a href="#">azuredisk-csi-driver</a> <sup>632</sup>	azuredisk-sc
Pre-provisioned	<a href="#">local-static-provisioner</a> <sup>633</sup>	localvolumeprovisioner
vSphere	<a href="#">vsphere-csi-driver</a> <sup>634</sup>	vsphere-raw-block-sc
GCP	<a href="#">gcp-filestore-csi-driver</a> <sup>635</sup>	csi-gce-pd

<sup>629</sup> <https://kubernetes.io/docs/concepts/storage/storage-classes/>

<sup>630</sup> <https://kubernetes.io/docs/concepts/storage/dynamic-provisioning/>


<sup>631</sup> <https://github.com/kubernetes-sigs/aws-ebs-csi-driver>

<sup>632</sup> <https://github.com/kubernetes-sigs/azuredisk-csi-driver>

<sup>633</sup> <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

<sup>634</sup> <https://github.com/kubernetes-sigs/vsphere-csi-driver>

<sup>635</sup> <https://github.com/kubernetes-sigs/gcp-filestore-csi-driver>

 DKP uses local static provisioner as the default storage provider. However, `localvolume-provisioner` is not suitable for production use. You should use a [Kubernetes CSI](#)<sup>636</sup> compatible storage that is suitable for production.

You can choose from any of the [storage options](#)<sup>637</sup> 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](#)<sup>638</sup>.

When a default StorageClass is specified, persistent volume claims (PVCs) can be created 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:
 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
```

#### 8.14.4.1 Multiple Storage Classes

The default `StorageClass` provisioned with DKP is acceptable for most workloads and offers a good cost to performance ratio. If your workload has different requirements, you can create additional `StorageClass` types with specific configurations.

<sup>636</sup> <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

<sup>637</sup> <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

<sup>638</sup> <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

In some instances you can change the default `StorageClass`. Refer to this procedure:

- [Changing the default storage class](#)<sup>639</sup>

## 8.14.4.2 Driver Information

All default drivers implement the [Container Storage Interface](#)<sup>640</sup> (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.

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

### 8.14.4.2.1 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](#)<sup>641</sup>
- Volume types and pricing: [volume types](#)<sup>642</sup>

<sup>639</sup> <https://kubernetes.io/docs/tasks/administer-cluster/change-default-storage-class/>

<sup>640</sup> <https://github.com/container-storage-interface/spec/blob/master/spec.md>

<sup>641</sup> <https://github.com/kubernetes-sigs/aws-ebs-csi-driver>

<sup>642</sup> <https://aws.amazon.com/ebs/features/>



#### 8.14.4.2.2 Azure CSI Driver

DKP deploys with StandardSSD\_LRS for Azure Virtual Disks.

- Driver documentation: [azuredisk-csi-driver](#)<sup>643</sup>
- Volume types and pricing: [volume types](#)<sup>644</sup>

#### 8.14.4.2.3 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>

#### 8.14.4.2.4 Pre-provisioned CSI Driver

In a [Pre-provisioned Infrastructure](#) (see page 288) environment (which includes on-premises, vSphere, AWS, Azure, and GCP), DKP will also deploy a CSI compatible driver and configure a default StorageClass - `localvolumeprovisioner`.

- Driver documentation: [local-static-provisioner](#)<sup>645</sup>



DKP uses [local-static-provisioner](#)<sup>646</sup> (`localvolumeprovisioner`) as the default storage provider for a pre-provisioned environment. However, [local-static-provisioner](#)<sup>647</sup> is not suitable for production use. You should use a different [Kubernetes CSI](#)<sup>648</sup> compatible storage that is suitable for production.

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: <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](#)<sup>649</sup> the default storage.

<sup>643</sup> <https://github.com/kubernetes-sigs/azuredisk-csi-driver>

<sup>644</sup> <https://azure.microsoft.com/en-us/pricing/details/storage/page-blobs/>

<sup>645</sup> <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

<sup>646</sup> <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

<sup>647</sup> <https://github.com/kubernetes-sigs/sig-storage-local-static-provisioner>

<sup>648</sup> <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

<sup>649</sup> <https://kubernetes.io/docs/concepts/storage/volumes/#volume-types>

Ceph can also be used as CSI storage. Refer to that section in the documentation for information on how to use [Rook Ceph in DKP](#) (see page 858).

#### 8.14.4.2.5 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](#)<sup>650</sup>
- Persistent volumes and dynamic provisioning: [volume types](#)<sup>651</sup>

#### 8.14.4.3 On-Premises and other Storage Options

In an on-premises environment, accessible storage can be used for PV and PVCs. Using the Kubernetes CSI and third party drivers, you can use your local volumes and other storage devices in your data center.


#### 8.14.4.4 Related Information

- [Kubernetes Storage](#)<sup>652</sup>
- [Kubernetes Local Persistent Volumes](#)<sup>653</sup>

### 8.14.5 Rook Ceph in DKP

DKP is shipped with a Rook Ceph cluster, which is used as the primary blob storage for various DKP components in 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.

 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 866).

- [Rook Ceph in DKP - Prerequisites](#) (see page 859)
- [Rook Ceph Configuration](#) (see page 860)
- [Rook Ceph Dashboard](#) (see page 865)
- [BYOS \(Bring Your Own Storage\) to DKP Clusters](#) (see page 866)

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

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

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

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

### 8.14.5.1 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 866).

If you do not plan on using any of the logging stack components such as `grafana-loki`, `project-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 731) then you do not need Ceph to be installed.

- = For more information about Ceph, refer to the following:

- [Intro to Ceph – Ceph Documentation](#)<sup>654</sup>
- [Rook – Rook Ceph Documentation](#)<sup>655</sup>

#### 8.14.5.1.1 Disabling Ceph during Upgrades (from DKP 2.3.x)


Up until DKP 2.3.x, `MinIO` was used as the default blob storage for DKP components. This has been replaced by Ceph from `2.4.x`, therefore when you upgrade from `2.3.x` -> `2.4.x`, Ceph is automatically installed if any of the following are true:


- `minio-operator` was installed (used by logging stack).
- `velero` was installed (used for backups).

However, if you did not have `minio-operator` installed in `2.3.x` and instead had configured `velero` to work with an [external cloud storage](#) (see page 731) such as Amazon S3 or Azure Blob Storage, then you do not need `rook-ceph` and `rook-ceph-cluster` post upgrade.

To explicitly disable installing these apps during upgrade, specify the following in the command line, then run the following command:

```
dkp upgrade kommander --disable-appdeployments rook-ceph,rook-ceph-cluster
```

 The command above fails if `minio-operator` is installed as you cannot disable Ceph installation if `minio-operator` is installed.

 See [Rook Ceph Configuration](#) (see page 860) for information on how you can configure Ceph for your Ceph Environment.

#### 8.14.5.2 Rook Ceph Configuration

**This page contains information about configuring Rook Ceph in your DKP Environment.**

<sup>654</sup> <https://docs.ceph.com/en/quincy/start/intro/>

<sup>655</sup> <https://rook.io/docs/rook/v1.10/Getting-Started/intro/>

**[-]** 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 866).

#### 8.14.5.2.1 Components of a Rook Ceph Cluster

Ceph supports creating clusters in different modes as listed in [CephCluster CRD - Rook Ceph Documentation](#)<sup>656</sup>. DKP, specifically is shipped with a PVC Cluster, as documented in [PVC Storage Cluster - Rook Ceph Documentation](#)<sup>657</sup>. It is recommended to use the PVC mode to keep the deployment and upgrades simple and agnostic to technicalities with node draining.

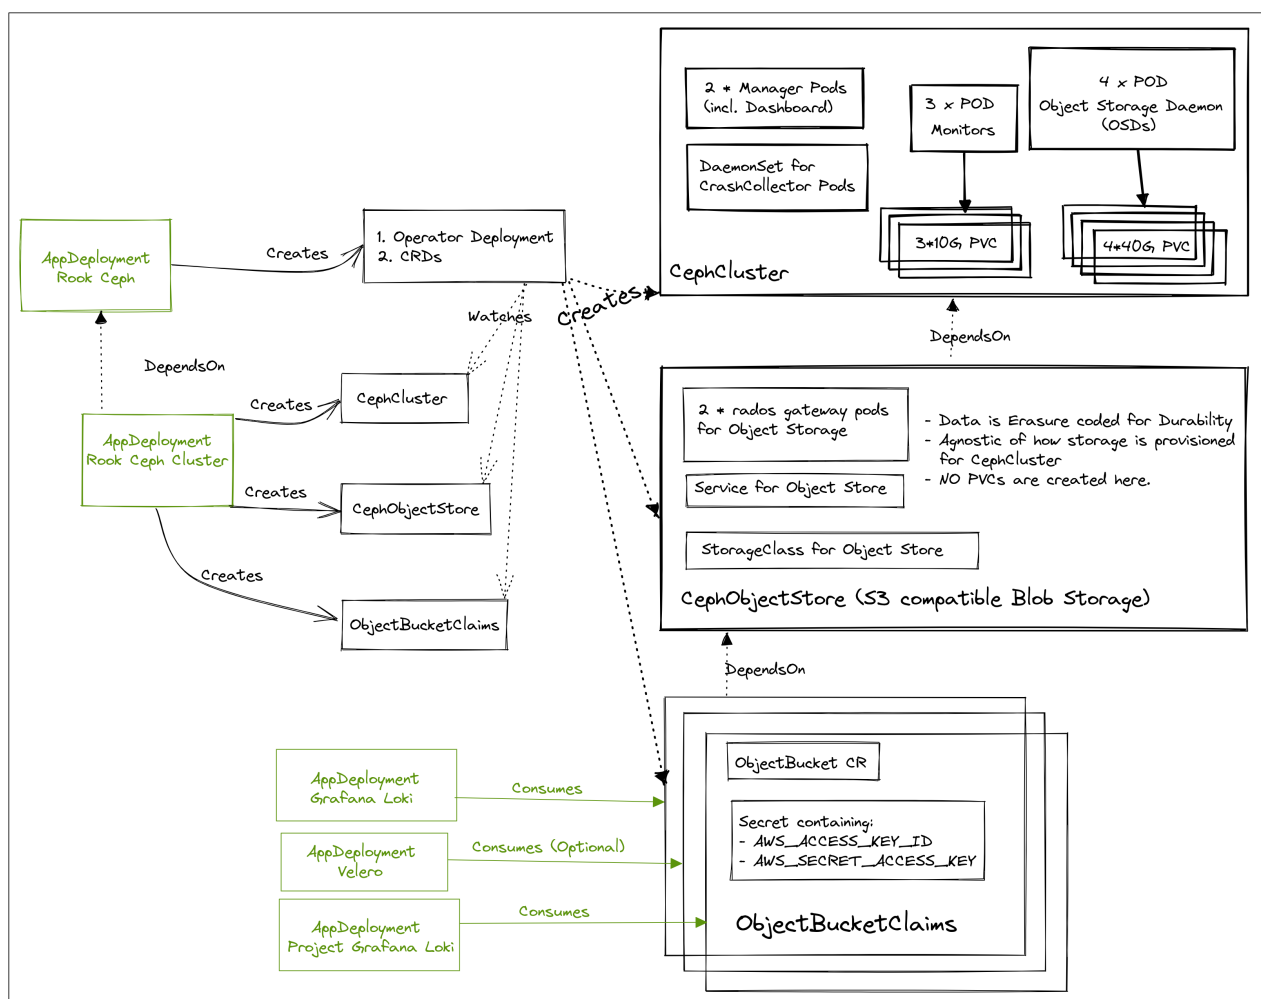
Ceph cannot be your CSI Provisioner when installing in PVC mode as Ceph relies on an existing CSI provisioner to bind the PVCs created by it. It is possible to use Ceph as your CSI provisioner, but that is outside the scope of this document. If you have an instance of Ceph that acts as the CSI Provisioner, then it is possible to reuse it for your DKP Storage needs. See [BYOS \(Bring Your Own Storage\) to DKP Clusters](#) (see page 866) 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:

---

<sup>656</sup> <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/>

<sup>657</sup> <https://rook.io/docs/rook/v1.10/CRDs/Cluster/pvc-cluster/#pvc-storage-only-for-monitors>



### 1 Rook Ceph Cluster Components

Items highlighted in green are user-facing and configurable.

**📖** Please refer to [Rook Ceph Storage Architecture](#)<sup>658</sup> and [Ceph Architecture](#)<sup>659</sup> 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](#)<sup>660</sup> page.

### 8.14.5.2.2 Resource Requirements

The following is a non-exhaustive list of the resource requirements for long running components of Ceph:

658 <https://rook.io/docs/rook/v1.10/Getting-Started/storage-architecture/>  
 659 <https://docs.ceph.com/en/quincy/architecture/>  
 660 <https://github.com/rook/rook/blob/release-1.10/design/ceph/data-model.md>

Type	Resources	Total
<b>CPUs</b>	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
<b>Memory</b>	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
<b>Disk</b>	4 x 40Gi PVCs with <code>Block</code> mode for <code>ObjectStorageDaemons</code> <sup>661</sup> 3 x 10Gi PVCs with <code>Block</code> or <code>FileSystem</code> mode for <code>Mons</code> <sup>662</sup>	190Gi

Your default `StorageClass` should support creation of `PersistentVolume`s that satisfy the `PersistentVolumeClaim`s created by Ceph with `volumeMode: Block`.

### 8.14.5.2.3 Ceph Storage Configuration

Ceph is highly configurable and can support Replication or Erasure Coding to ensure [data durability](#)<sup>663</sup>. 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.

#### 8.14.5.2.3.1 Replication<sup>664</sup>

- 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.

<sup>661</sup> <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#storage-selection-settings>

<sup>662</sup> <https://rook.io/docs/rook/v1.10/CRDs/Cluster/ceph-cluster-crd/#mon-settings>

<sup>663</sup> [https://en.wikipedia.org/wiki/Durability\\_\(database\\_systems\)](https://en.wikipedia.org/wiki/Durability_(database_systems))

<sup>664</sup> [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.

#### 8.14.5.2.3.2 Erasure Coding<sup>665</sup>

- 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 ( $\geq n$ ) 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)<sup>666</sup> and override the replication strategy according to your needs by referring to [CephObjectStore CRD](https://www.rook.io/docs/rook/v1.10/CRDs/Object-Storage/ceph-object-store-crd/)<sup>667</sup> documentation.

For more information about configuring storage in Rook Ceph, refer to the following pages:

<sup>665</sup> [https://en.wikipedia.org/wiki/Erasure\\_code](https://en.wikipedia.org/wiki/Erasure_code)

<sup>666</sup> <https://github.com/mesosphere/kommander-applications/blob/v2.4.0/services/rook-ceph-cluster/1.10.3/defaults/cm.yaml#L126-L175>

<sup>667</sup> <https://www.rook.io/docs/rook/v1.10/CRDs/Object-Storage/ceph-object-store-crd/>



- [Ceph OSD Management - Rook Ceph Documentation](https://www.rook.io/docs/rook/v1.11/Storage-Configuration/Advanced/ceph-osd-mgmt/)<sup>668</sup> - for general information on how to configure Object Storage Daemons (OSDs).
- [Ceph Configuration - Rook Ceph Documentation](https://www.rook.io/docs/rook/v1.11/Storage-Configuration/Advanced/ceph-configuration/?h=expa#auto-expansion-of-osds)<sup>669</sup> - for information on how to set up auto-expansion of OSDs.

### See Also:

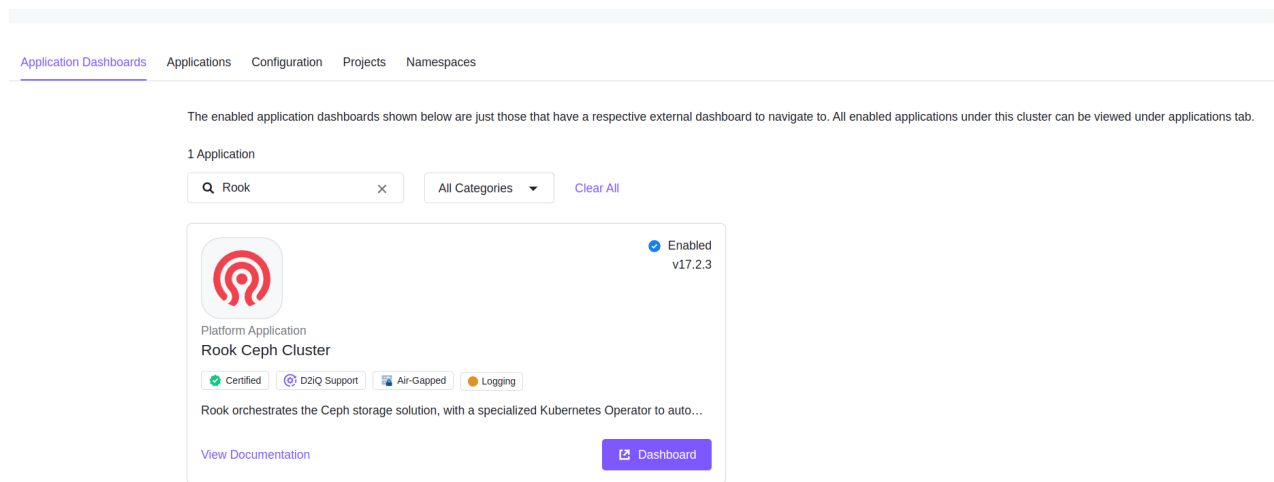
[Rook Ceph in DKP - Prerequisites \(see page 859\)](#)

[Rook Ceph Dashboard \(see page 865\)](#)

## 8.14.5.3 Rook Ceph Dashboard

### 8.14.5.3.1 Rook Ceph Dashboard Login

Rook Ceph Cluster provides a Dashboard which can be accessed from the UI in the Application Dashboards tab inside Kommander.



#### 2 Rook Ceph Dashboard Link in the UI

The dashboard can be used to view current cluster health and logs.

#### 8.14.5.3.1.1 How to Access the Rook Ceph Dashboard

1. Go to the applications dashboard
2. Select the Dashboard button
3. Username is `admin`

<sup>668</sup> <https://www.rook.io/docs/rook/v1.11/Storage-Configuration/Advanced/ceph-osd-mgmt/>

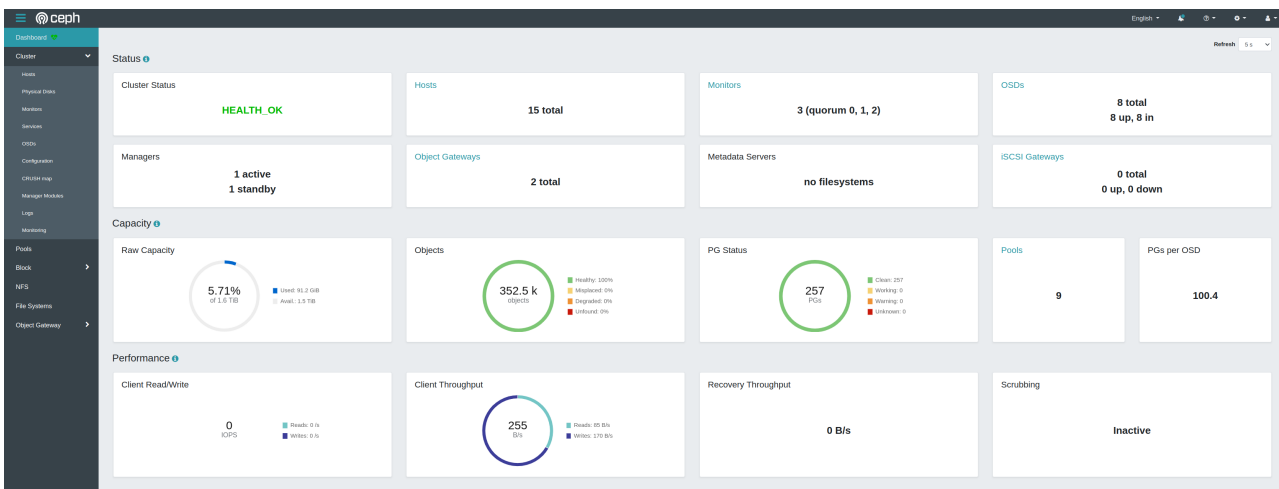
<sup>669</sup> <https://www.rook.io/docs/rook/v1.11/Storage-Configuration/Advanced/ceph-configuration/?h=expa#auto-expansion-of-osds>

- 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 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:



3 Rook Ceph Cluster Dashboard

#### 8.14.5.4 BYOS (Bring Your Own Storage) to DKP Clusters

Ceph can be used 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.



This guide assumes you have a Ceph cluster that is not managed by DKP. Refer to [Rook Ceph Configuration \(see page 860\)](#) for information on how to configure the Ceph instance installed by DKP for use by DKP platform applications.

##### 8.14.5.4.1 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.

#### 8.14.5.4.2 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](#)<sup>670</sup> 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.

##### 8.14.5.4.2.1 Creating `CephObjectStore`

There are two ways to install Ceph:

- Using [Helm Charts](#) (see page 867)
- [Directly applying Kubernetes manifests](#) (see page 869)

#### Helm Chart Values

This section is relevant if you have installed Ceph using `helm install` or some other managed Helm resource mechanism.

<sup>670</sup> <https://www.rook.io/docs/rook/v1.10/Getting-Started/quickstart/#create-a-ceph-cluster>

If you have applied any configuration overrides to your Rook Ceph operator, ensure it was deployed with `currentNamespaceOnly` set to `false`<sup>671</sup> (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`.

Ensure the following configuration<sup>672</sup> for `rook-ceph` helm chart<sup>673</sup>:

```
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<sup>674</sup> for the `rook-ceph-cluster` helm chart<sup>675</sup>:

```
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:
```

671 <https://www.rook.io/docs/rook/v1.10/Helm-Charts/operator-chart/>

672 <https://www.rook.io/docs/rook/v1.10/Helm-Charts/operator-chart/#configuration>

673 <https://www.rook.io/docs/rook/v1.10/Helm-Charts/operator-chart/#release>

674 <https://www.rook.io/docs/rook/v1.10/Helm-Charts/ceph-cluster-chart/#ceph-object-stores>

675 <https://www.rook.io/docs/rook/v1.10/Helm-Charts/ceph-cluster-chart/#release>

```

limits:
 cpu: "750m"
 memory: "1Gi"
requests:
 cpu: "250m"
 memory: "500Mi"
healthCheck:
 bucket:
 interval: 60s
storageClass:
 enabled: true
 name: dkp-object-store
 reclaimPolicy: Delete

```

### Managing resources directly

1. Set a variable to refer to the namespace the `AppDeployments` 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

```

#### 8.14.5.4.2.2 Creating ObjectBucketClaims

1. Once the Object Store is Connected, create the ObjectBucketClaim in the same namespace as velero and grafana-loki.

This results in the creation of ObjectBucket, which 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

### 8.14.5.4.3 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
```

### 8.14.5.4.4 Overriding `velero` and `grafana-loki` Configuration

Once 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:
 provider: "aws"
 backupStorageLocation:
 bucket: dkp-velero
 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.

#### 8.14.5.4.5 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.

## 9 CLI and API Tools

CLI commands and API tools.

- [API Documentation](#) (see page 876)
- [CLI Commands](#) (see page 919)

### 9.1 API Documentation

This document is automatically generated from the API definition in the code.

- [apps.kommander.mesosphere.io/v1alpha2](#) (see page 876)
- [apps.kommander.mesosphere.io/v1alpha3](#) (see page 880)
- [dispatch.d2iq.io/v1alpha2](#) (see page 887)
- [kommander.mesosphere.io/v1beta1](#) (see page 888)
- [workspaces.kommander.mesosphere.io/v1alpha1](#) (see page 901)

#### 9.1.1 apps.kommander.mesosphere.io/v1alpha2

##### 9.1.1.1 App

App is the Schema for a Service or Application in Kommander.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>676</sup>	false
spec		<a href="#">GenericAppSpec</a> (see page 878)	false
status		object	false

##### 9.1.1.2 AppDeployment

AppDeployment is the Schema for a concrete installation of a Service or Application in Kommander.

---

<sup>676</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>677</sup>	false
spec		<a href="#">AppDeploymentSpec</a> (see page 876)	false
status		object	false

### 9.1.1.3 AppDeploymentList

AppDeploymentList contains a list of AppDeployments.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>678</sup>	false
items		[...]AppDeployment (see page 876)	true

### 9.1.1.4 AppDeploymentSpec

AppDeploymentSpec defines an instance of an Application.

Field	Description	Scheme	Required
appRef	AppRef provides reference to a ClusterApp or App object.	<a href="#">TypedLocalObjectReference</a> (see page 879)	true
configOverrides	ConfigOverrides allows a user to define a ConfigMap that contains configuration overrides	<a href="#">corev1.LocalObjectReference</a> <sup>679</sup>	false

<sup>677</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>678</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>679</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

### 9.1.1.5 AppList

AppList contains a list of Apps.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>680</sup>	false
items		[...]App (see page 876)	true

### 9.1.1.6 ClusterApp

ClusterApp is the Schema for a Service or Application in Kommander.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>681</sup>	false
spec		<a href="#">GenericAppSpec</a> (see page 878)	false
status		object	false

### 9.1.1.7 ClusterAppList

ClusterAppList contains a list of ClusterApps.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>682</sup>	false
items		[...]ClusterApp (see page 877)	true

<sup>680</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>681</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>682</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

### 9.1.1.8 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

### 9.1.1.9 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	<a href="#">CrossNamespaceGitRepositoryReference</a> (see <a href="#">page 878</a> )	true
version	Version depicts the version of the service in the semantic versioning format.	string	true

### 9.1.1.10 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

## 9.1.2 apps.kommander.mesosphere.io/v1alpha3

### 9.1.2.1 App

App is the Schema for a Service or Application in Kommander.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>683</sup>	false
spec		<a href="#">GenericAppSpec</a> (see <a href="#">page 886</a> )	false
status		object	false

### 9.1.2.2 AppDeployment

AppDeployment is the Schema for a concrete installation of a Service or Application in Kommander.

---

<sup>683</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>684</sup>	false
spec		<a href="#">AppDeploymentSpec</a> (see page 882)	false
status		object	false

### 9.1.2.3 AppDeploymentClusterCondition

Field	Description	Scheme	Required
lastTransitionTime	LastTransitionTime is the last time the condition transitioned from one status to another.	<a href="#">metav1.Time</a> <sup>685</sup>	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

### 9.1.2.4 AppDeploymentList

AppDeploymentList contains a list of AppDeployments.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>686</sup>	false

<sup>684</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>685</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

<sup>686</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>



Field	Description	Scheme	Required
items		[...]AppDeployment (see page 880)	true

### 9.1.2.5 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 886)	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 883)	false
clusterSelector	ClusterSelector defines the label selectors to select clusters on which to enable the AppDeployment.	metav1.LabelSelector <sup>687</sup>	false
configOverrides	ConfigOverrides allows a user to define a ConfigMap that contains configuration overrides at the workspace level.	corev1.LocalObjectReference <sup>688</sup>	false

### 9.1.2.6 AppDeploymentStatus

AppDeploymentStatus defines the current state of the AppDeployment.

<sup>687</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#labelselector-v1-meta>

<sup>688</sup> <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 885)	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

### 9.1.2.7 AppList

AppList contains a list of Apps.

Field	Description	Scheme	Required
metadata		metav1.ListMeta <sup>689</sup>	false
items		[...]App (see page 880)	true

<sup>689</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

### 9.1.2.8 ClusterApp

ClusterApp is the Schema for a Service or Application in Kommander.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>690</sup>	false
spec		<a href="#">GenericAppSpec</a> (see page 886)	false
status		object	false

### 9.1.2.9 ClusterAppList

ClusterAppList contains a list of ClusterApps.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>691</sup>	false
items		[...]ClusterApp (see page 883)	true

### 9.1.2.10 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.	<a href="#">metav1.LabelSelector</a> <sup>692</sup>	false

<sup>690</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>691</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>692</sup> <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

### 9.1.2.11 ClusterDeploymentStatus

Field	Description	Scheme	Required
clusterConfigOverridesRef	ClusterConfigOverridesRef is set to reference the overrides ConfigMap in the management cluster	<a href="#">corev1.LocalObjectReference</a> <sup>693</sup>	false
conditions		<a href="#">[...]AppDeploymentClusterCondition</a> (see page 880)	false
name		string	true

### 9.1.2.12 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

<sup>693</sup> <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

### 9.1.2.13 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	<a href="#">CrossNamespaceGitRepositoryReference</a> (see <a href="#">page 885</a> )	true
version	Version depicts the version of the service in the semantic versioning format.	string	true

### 9.1.2.14 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

## 9.1.3 dispatch.d2iq.io/v1alpha2

### 9.1.3.1 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		<a href="#">metav1.ObjectMeta</a> <sup>694</sup>	false
spec		<a href="#">GitopsRepositorySpec</a> (see page 887)	true

### 9.1.3.2 GitopsRepositoryList

GitopsRepositoryList contains a list of Repository.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>695</sup>	false
items		[...] <a href="#">GitopsRepository</a> (see page 887)	true

### 9.1.3.3 GitopsRepositorySpec

GitopsRepositorySpec defines the desired state of GitopsRepository.

<sup>694</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>695</sup> <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.	<a href="#">GitopsRolloutTemplate</a> (see page 887)	false

#### 9.1.3.4 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.	<a href="#">sourcectrlv1beta1.GitRepositoryRef</a> <sup>696</sup>	false
suspend	Whether to suspend periodic or webhook-notified sync.	bool	false

<sup>696</sup> <https://fluxcd.io/flux/components/source/api/#source.toolkit.fluxcd.io/v1beta1.GitRepositoryRef>

## 9.1.4 kommander.mesosphere.io/v1beta1

### 9.1.4.1 CAPIClusterReference

Field	Description	Scheme	Required
name		string	true
namespace		string	false

### 9.1.4.2 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		<a href="#">CAPIClusterReference</a> (see page 888)	false

### 9.1.4.3 GenericClusterReference

GenericClusterReference sets the name of the cluster.

Field	Description	Scheme	Required
name		string	true

### 9.1.4.4 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.	<a href="#">corev1.LocalObjectReference</a> <sup>697</sup>	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.	<a href="#">IssuerReference</a> (see <a href="#">page 890</a> )	false

### 9.1.4.5 IngressStatus

IngressSpec holds information about the cluster's Kommander managed Ingress.

---

<sup>697</sup> <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

#### 9.1.4.6 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

### 9.1.4.7 KommanderCluster

KommanderCluster is the Schema for the kommander clusters API.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>698</sup>	false
spec		<a href="#">KommanderClusterSpec</a> (see page 892)	false
status		<a href="#">KommanderClusterStatus</a> (see page 893)	false

### 9.1.4.8 KommanderClusterCondition

Field	Description	Scheme	Required
lastTransitionTime		<a href="#">metav1.Time</a> <sup>699</sup>	false
message		string	false
reason		string	false
status		string	true
type		string	true

<sup>698</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>699</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

### 9.1.4.9 KommanderClusterList

KommanderClusterList contains a list of Cluster.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>700</sup>	false
items		[...]KommanderCluster ( <a href="#">see page 892</a> )	true

### 9.1.4.10 KommanderClusterSpec

KommanderClusterSpec defines the desired state of Cluster.

Field	Description	Scheme	Required
clusterRef		<a href="#">ClusterReference</a> ( <a href="#">see page 888</a> )	false
clusterTunnelConnectorRef	ClusterTunnelConnectorRef is a reference to TunnelConnector that should be used for connecting to cluster.	<a href="#">corev1.LocalObjectReference</a> <sup>701</sup>	false
ingress	Ingress configures the Kommander managed Ingress. If no value is provided, the default ingress will be provisioned.	<a href="#">IngressSpec</a> ( <a href="#">see page 888</a> )	false
kubeconfigRef		<a href="#">corev1.LocalObjectReference</a> <sup>702</sup>	false

<sup>700</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>701</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>702</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

### 9.1.4.11 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		[...]KommanderClusterCondition (see page 892)	false
dextfaclientRef	DexTFAClientRef holds a reference to a DexClient provisioned for Traefik Forward Auth running on managed cluster.	corev1.ObjectReference <sup>703</sup>	false
ingress	Ingress holds information about the cluster's Ingress	IngressStatus (see page 889)	false
kubefedclusterRef	KubefedClusterRef holds a reference to a kubefedcluster in the kubefed system namespace.	corev1.LocalObjectReference <sup>704</sup>	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

<sup>703</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectreference-v1-core>

<sup>704</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

Field	Description	Scheme	Required
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

### 9.1.4.12 License

License is the Schema for the licenses API.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>705</sup>	false
spec		<a href="#">LicenseSpec</a> (see page 896)	false
status		<a href="#">LicenseStatus</a> (see page 896)	false

### 9.1.4.13 LicenseCondition

Field	Description	Scheme	Required
lastTransitionTime		<a href="#">metav1.Time</a> <sup>706</sup>	false
message		string	false

<sup>705</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>706</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

Field	Description	Scheme	Required
reason		string	false
status		string	true
type		string	true

#### 9.1.4.14 LicenseExternalAWS

Field	Description	Scheme	Required
licenseArn	LicenseArn is a fully qualified AWS arn to a license for Kommander.	string	false

#### 9.1.4.15 LicenseExternalReference

Field	Description	Scheme	Required
awsLicense	AWSLicense holds a single reference to a license in AWS's License Manager	<a href="#">LicenseExternalAWS (see page 896)</a>	false
type	Type is the source of the external license referenced, i.e. AWS, GCP, Azure	string	true

#### 9.1.4.16 LicenseList

LicenseList contains a list of License.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>707</sup>	false
items		[...]License (see page 895)	true

### 9.1.4.17 LicenseSpec

LicenseSpec defines the desired state of License.

Field	Description	Scheme	Required
externalLicenseRef	ExternalLicenseRef holds a reference to an external license	<a href="#">LicenseExternalReference</a> (see page 896)	false
licenseRef	LicenseReference holds a single reference to the secret holding the license JWT	<a href="#">corev1.LocalObjectReference</a> <sup>708</sup>	false

### 9.1.4.18 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 895)	false

<sup>707</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>708</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>



Field	Description	Scheme	Required
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.	<a href="#">metav1.Time</a> <sup>709</sup>	false
licenseId	The license's ID as provided from Salesforce.	string	true
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.	<a href="#">metav1.Time</a> <sup>710</sup>	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

<sup>709</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

<sup>710</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

Field	Description	Scheme	Required
version	The license's version ID as provided when the license was created.	string	true

#### 9.1.4.19 PlacementSelector

PlacementSelector defines the fields to select clusters where to apply the project.

Field	Description	Scheme	Required
clusterSelector		<a href="#">metav1.LabelSelector</a> <sup>711</sup>	false
clusters		<a href="#">[...]GenericClusterReference</a> (see page 888)	false

#### 9.1.4.20 VirtualGroup

VirtualGroup is the Schema for the virtualgroups API.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>712</sup>	false
spec		<a href="#">VirtualGroupSpec</a> (see page 900)	false

#### 9.1.4.21 VirtualGroupClusterRoleBinding

VirtualGroupClusterRoleBinding is the Schema for the virtualgroupclusterrolebindings API.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>713</sup>	false

<sup>711</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#labelselector-v1-meta>

<sup>712</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>713</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

Field	Description	Scheme	Required
spec		<a href="#">VirtualGroupClusterRoleBindingSpec</a> (see page 900)	true

#### 9.1.4.22 VirtualGroupClusterRoleBindingList

VirtualGroupClusterRoleBindingList contains a list of VirtualGroupClusterRoleBinding.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>714</sup>	false
items		[...] <a href="#">VirtualGroupClusterRoleBinding</a> (see page 899)	true

#### 9.1.4.23 VirtualGroupClusterRoleBindingSpec

VirtualGroupClusterRoleBindingSpec defines the desired state of VirtualGroupClusterRoleBinding.

Field	Description	Scheme	Required
clusterRoleRef		<a href="#">corev1.LocalObjectReference</a> <sup>715</sup>	true
placement		<a href="#">PlacementSelector</a> (see page 899)	false
virtualGroupRef		<a href="#">corev1.LocalObjectReference</a> <sup>716</sup>	true

#### 9.1.4.24 VirtualGroupList

VirtualGroupList contains a list of VirtualGroup.

<sup>714</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>715</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>716</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>717</sup>	false
items		[...] <a href="#">VirtualGroup</a> (see page 899)	true

### 9.1.4.25 VirtualGroupSpec

VirtualGroupSpec defines the desired state of VirtualGroup.

Field	Description	Scheme	Required
subjects		[...] <a href="#">rbacv1.Subject</a> <sup>718</sup>	false

## 9.1.5 workspaces.kommander.mesosphere.io/v1alpha1

### 9.1.5.1 KommanderProjectRole

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>719</sup>	false
spec		<a href="#">KommanderProjectRole Spec</a> (see page 901)	false
status		<a href="#">KommanderProjectRole Status</a> (see page 901)	false

<sup>717</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>718</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#subject-v1-rbac-authorization-k8s-io>

<sup>719</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

### 9.1.5.2 KommanderProjectRoleList

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>720</sup>	false
items		[...] <a href="#">KommanderProjectRole</a> (see page 901)	true

### 9.1.5.3 KommanderProjectRoleSpec

Field	Description	Scheme	Required
projectObjectVerbs		[...]string	false
rules		[...] <a href="#">rbacv1.PolicyRule</a> <sup>721</sup>	false

### 9.1.5.4 KommanderProjectRoleStatus

Field	Description	Scheme	Required
roleInProjectRef		<a href="#">corev1.LocalObjectReference</a> <sup>722</sup>	false
roleInWorkspaceRef		<a href="#">corev1.LocalObjectReference</a> <sup>723</sup>	false

<sup>720</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>721</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#policyrule-v1-rbac-authorization-k8s-io>

<sup>722</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>723</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

### 9.1.5.5 KommanderWorkspaceRole

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>724</sup>	false
spec		<a href="#">KommanderWorkspaceRoleSpec</a> (see page 902)	false
status		<a href="#">KommanderWorkspaceRoleStatus</a> (see page 903)	false

### 9.1.5.6 KommanderWorkspaceRoleList

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>725</sup>	false
items		[...] <a href="#">KommanderWorkspaceRole</a> (see page 902)	true

### 9.1.5.7 KommanderWorkspaceRoleSpec

Field	Description	Scheme	Required
rules		[...] <a href="#">rbacv1.PolicyRule</a> <sup>726</sup>	false
workspaceObjectVerbs		[...]string	false

<sup>724</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>725</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>726</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#policyrule-v1-rbac-authorization-k8s-io>

### 9.1.5.8 KommanderWorkspaceRoleStatus

Field	Description	Scheme	Required
clusterRoleRef		<a href="#">corev1.LocalObjectReference</a> <sup>727</sup>	false
roleInWorkspaceRef		<a href="#">corev1.LocalObjectReference</a> <sup>728</sup>	false

### 9.1.5.9 Project

Project is a logical top-level container for a set of Kommander resources.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>729</sup>	false
spec		<a href="#">ProjectSpec</a> (see page 906)	false
status		<a href="#">ProjectStatus</a> (see page 906)	false

### 9.1.5.10 ProjectCondition

Field	Description	Scheme	Required
lastTransitionTime	Last time the condition transitioned from one status to another.	<a href="#">metav1.Time</a> <sup>730</sup>	false

<sup>727</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>728</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>729</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>730</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

Field	Description	Scheme	Required
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

### 9.1.5.11 ProjectList

ProjectList is a list of Project objects.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>731</sup>	false
items		[...] <a href="#">Project</a> (see page 903)	true

### 9.1.5.12 ProjectRole

ProjectRole is the Schema for the workspaces ProjectRole API.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>732</sup>	false

<sup>731</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>732</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>



Field	Description	Scheme	Required
spec		<a href="#">ProjectRoleSpec</a> (see page 906)	false
status		<a href="#">ProjectRoleStatus</a> (see page 906)	false

### 9.1.5.13 ProjectRoleList

ProjectRoleList contains a list of ProjectRole.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>733</sup>	false
items		[...] <a href="#">ProjectRole</a> (see page 905)	true

### 9.1.5.14 ProjectRoleSpec

Field	Description	Scheme	Required
rules		[...] <a href="#">rbacv1.PolicyRule</a> <sup>734</sup>	false

### 9.1.5.15 ProjectRoleStatus

Field	Description	Scheme	Required
federatedRoleRef		<a href="#">corev1.LocalObjectReference</a> <sup>735</sup>	false

<sup>733</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>734</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#policyrule-v1-rbac-authorization-k8s-io>

<sup>735</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

### 9.1.5.16 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		<a href="#">v1beta1.PlacementSelector</a> <sup>736</sup>	false
workspaceRef		<a href="#">corev1.LocalObjectReference</a> <sup>737</sup>	false

### 9.1.5.17 ProjectStatus

ProjectStatus is information about the current status of a Project.

Field	Description	Scheme	Required
conditions	Represents the latest available observations of a project's current state.	<a href="#">[...]ProjectCondition</a> (see <a href="#">page 903</a> )	false
namespaceRef		<a href="#">corev1.LocalObjectReference</a> <sup>738</sup>	false

### 9.1.5.18 VirtualGroupKommanderClusterRoleBinding

VirtualGroupKommanderClusterRoleBinding is the Schema for the VirtualGroupWorkspaceRoleBinding API.

<sup>736</sup> <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/110854693/kommander.mesosphere.io+v1beta1#PlacementSelector>

<sup>737</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>738</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>739</sup>	false
spec		<a href="#">VirtualGroupKommanderClusterRoleBindingSpec</a> (see page 907)	true
status		<a href="#">VirtualGroupKommanderClusterRoleBindingStatus</a> (see page 908)	false

### 9.1.5.19 VirtualGroupKommanderClusterRoleBindingList

VirtualGroupKommanderClusterRoleBindingList contains a list of VirtualGroupKommanderClusterRoleBinding.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>740</sup>	false
items		[...]VirtualGroupKommanderClusterRoleBinding (see page 906)	true

### 9.1.5.20 VirtualGroupKommanderClusterRoleBindingSpec

Field	Description	Scheme	Required
clusterRoleRef		<a href="#">corev1.LocalObjectReference</a> <sup>741</sup>	true
virtualGroupRef		<a href="#">corev1.LocalObjectReference</a> <sup>742</sup>	true

<sup>739</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>740</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>741</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>742</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

### 9.1.5.21 VirtualGroupKommanderClusterRoleBindingStatus

Field	Description	Scheme	Required
clusterRoleBindingRef		<a href="#">corev1.LocalObjectReference</a> <sup>743</sup>	false

### 9.1.5.22 VirtualGroupKommanderProjectRoleBinding

VirtualGroupKommanderProjectRoleBinding is the Schema to be used in the API.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>744</sup>	false
spec		<a href="#">VirtualGroupKommanderProjectRoleBindingSpec</a> (see page 909)	true
status		<a href="#">VirtualGroupKommanderProjectRoleBindingStatus</a> (see page 909)	false

### 9.1.5.23 VirtualGroupKommanderProjectRoleBindingList

VirtualGroupKommanderProjectRoleBindingList contains a list of VirtualGroupKommanderProjectRoleBinding.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>745</sup>	false
items		<a href="#">[...]VirtualGroupKommanderProjectRoleBinding</a> (see page 908)	true

<sup>743</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>744</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>745</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

### 9.1.5.24 VirtualGroupKommanderProjectRoleBindingSpec

Field	Description	Scheme	Required
kommanderProjectRoleRef		<a href="#">corev1.LocalObjectReference</a> <sup>746</sup>	true
virtualGroupRef		<a href="#">corev1.LocalObjectReference</a> <sup>747</sup>	true

### 9.1.5.25 VirtualGroupKommanderProjectRoleBindingStatus

Field	Description	Scheme	Required
roleBindingInProjectRef		<a href="#">corev1.LocalObjectReference</a> <sup>748</sup>	false
roleBindingInWorkspaceRef		<a href="#">corev1.LocalObjectReference</a> <sup>749</sup>	false

### 9.1.5.26 VirtualGroupKommanderWorkspaceRoleBinding

VirtualGroupKommanderWorkspaceRoleBinding is the Schema to be used in the API.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>750</sup>	false
spec		<a href="#">VirtualGroupKommanderWorkspaceRoleBindingSpec</a> (see page 911)	true

<sup>746</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>747</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>748</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>749</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>750</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

Field	Description	Scheme	Required
status		<a href="#">VirtualGroupKommanderWorkspaceRoleBindingStatus</a> (see page 911)	false

### 9.1.5.27 VirtualGroupKommanderWorkspaceRoleBindingList

VirtualGroupKommanderWorkspaceRoleBindingList contains a list of VirtualGroupKommanderWorkspaceRoleBinding.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>751</sup>	false
items		[...]VirtualGroupKommanderWorkspaceRoleBinding (see page 909)	true

### 9.1.5.28 VirtualGroupKommanderWorkspaceRoleBindingSpec

Field	Description	Scheme	Required
kommanderWorkspaceRoleRef		<a href="#">corev1.LocalObjectReference</a> <sup>752</sup>	true
virtualGroupRef		<a href="#">corev1.LocalObjectReference</a> <sup>753</sup>	true

<sup>751</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

<sup>752</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>753</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

### 9.1.5.29 VirtualGroupKommanderWorkspaceRoleBindingStatus

Field	Description	Scheme	Required
clusterRoleBindingRef		<a href="#">corev1.LocalObjectReference</a> <sup>754</sup>	false
roleBindingInWorkspaceRef		<a href="#">corev1.LocalObjectReference</a> <sup>755</sup>	false

### 9.1.5.30 VirtualGroupProjectRoleBinding

VirtualGroupProjectRoleBinding is the Schema for the VirtualGroupProjectRoleBinding API.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>756</sup>	false
spec		<a href="#">VirtualGroupProjectRoleBindingSpec</a> (see page 913)	true
status		<a href="#">VirtualGroupProjectRoleBindingStatus</a> (see page 913)	false

### 9.1.5.31 VirtualGroupProjectRoleBindingList

VirtualGroupProjectRoleBindingList contains a list of VirtualGroupProjectRoleBinding.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>757</sup>	false

<sup>754</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>755</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>756</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>757</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

Field	Description	Scheme	Required
items		[...]VirtualGroupProjectRoleBinding (see page 911)	true

### 9.1.5.32 VirtualGroupProjectRoleBindingSpec

Field	Description	Scheme	Required
projectRoleRef		<a href="#">corev1.LocalObjectReference</a> <sup>758</sup>	false
virtualGroupRef		<a href="#">corev1.LocalObjectReference</a> <sup>759</sup>	true
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.	<a href="#">corev1.LocalObjectReference</a> <sup>760</sup>	false

<sup>758</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>759</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>760</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>



### 9.1.5.33 VirtualGroupProjectRoleBindingStatus

Field	Description	Scheme	Required
federatedRoleBindingRef		<a href="#">corev1.LocalObjectReference</a> <sup>761</sup>	false

### 9.1.5.34 VirtualGroupWorkspaceRoleBinding

VirtualGroupWorkspaceRoleBinding is the Schema for the VirtualGroupWorkspaceRoleBinding API.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>762</sup>	false
spec		<a href="#">VirtualGroupWorkspaceRoleBindingSpec</a> (see page 914)	true
status		<a href="#">VirtualGroupWorkspaceRoleBindingStatus</a> (see page 914)	false

### 9.1.5.35 VirtualGroupWorkspaceRoleBindingList

VirtualGroupWorkspaceRoleBindingList contains a list of VirtualGroupWorkspaceRoleBinding.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>763</sup>	false
items		[...] <a href="#">VirtualGroupWorkspaceRoleBinding</a> (see page 913)	true

<sup>761</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>762</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>763</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

### 9.1.5.36 VirtualGroupWorkspaceRoleBindingSpec

Field	Description	Scheme	Required
placement		<a href="#">v1beta1.PlacementSelector</a> <sup>764</sup>	false
virtualGroupRef		<a href="#">corev1.LocalObjectReference</a> <sup>765</sup>	true
workspaceRoleRef		<a href="#">corev1.LocalObjectReference</a> <sup>766</sup>	true

### 9.1.5.37 VirtualGroupWorkspaceRoleBindingStatus

Field	Description	Scheme	Required
federatedClusterRoleBindingRef		<a href="#">corev1.LocalObjectReference</a> <sup>767</sup>	false

### 9.1.5.38 Workspace

Workspace is the Schema for the workspaces API.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ObjectMeta</a> <sup>768</sup>	false
spec		<a href="#">WorkspaceSpec</a> (see <a href="#">page 917</a> )	false

<sup>764</sup> <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/110854693/kommander.mesosphere.io+v1beta1#PlacementSelector>

<sup>765</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>766</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>767</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

<sup>768</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

Field	Description	Scheme	Required
status		<a href="#">WorkspaceStatus</a> (see page 918)	false

### 9.1.5.39 WorkspaceCondition

Field	Description	Scheme	Required
lastTransitionTime	Last time the condition transitioned from one status to another.	<a href="#">metav1.Time</a> <sup>769</sup>	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 workspace condition.	string	true

### 9.1.5.40 WorkspaceList

WorkspaceList contains a list of Workspace.

Field	Description	Scheme	Required
metadata		<a href="#">metav1.ListMeta</a> <sup>770</sup>	false

<sup>769</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#time-v1-meta>

<sup>770</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

Field	Description	Scheme	Required
items		[...]Workspace (see page 914)	true

### 9.1.5.41 WorkspaceRole

WorkspaceRole is the Schema for the workspaces API.

Field	Description	Scheme	Required
metadata		metav1.ObjectMeta <sup>771</sup>	false
spec		WorkspaceRoleSpec (see page 917)	false
status		WorkspaceRoleStatus (see page 917)	false

### 9.1.5.42 WorkspaceRoleList

WorkspaceRoleList contains a list of WorkspaceRole.

Field	Description	Scheme	Required
metadata		metav1.ListMeta <sup>772</sup>	false
items		[...]WorkspaceRole (see page 917)	true

<sup>771</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#objectmeta-v1-meta>

<sup>772</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#listmeta-v1-meta>

### 9.1.5.43 WorkspaceRoleSpec

Field	Description	Scheme	Required
aggregationRule		<a href="#">rbacv1.AggregationRule</a> <sup>773</sup>	false
rules		[...] <a href="#">rbacv1.PolicyRule</a> <sup>774</sup>	false

### 9.1.5.44 WorkspaceRoleStatus

Field	Description	Scheme	Required
federatedClusterRoleReference		<a href="#">corev1.LocalObjectReference</a> <sup>775</sup>	false

### 9.1.5.45 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

<sup>773</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#aggregationrule-v1-rbac-authorization-k8s-io>

<sup>774</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#policyrule-v1-rbac-authorization-k8s-io>

<sup>775</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

### 9.1.5.46 WorkspaceStatus

Field	Description	Scheme	Required
conditions	Represents the latest available observations of a workspace's current state.	[...] <a href="#">WorkspaceCondition</a> (see page 916)	false
namespaceRef		<a href="#">corev1.LocalObjectReference</a> <sup>776</sup>	false

## 9.2 CLI Commands

### 9.2.1 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 433) commands.

- [dkp attach](#) (see page 920)
- [dkp check](#) (see page 921)
- [dkp completion](#) (see page 923)
- [dkp config](#) (see page 928)
- [dkp create](#) (see page 931)
- [dkp delete](#) (see page 968)
- [dkp describe](#) (see page 973)
- [dkp detach](#) (see page 974)
- [dkp diagnose](#) (see page 975)
- [dkp get](#) (see page 978)
- [dkp import](#) (see page 983)
- [dkp install](#) (see page 984)
- [dkp move](#) (see page 986)
- [dkp open](#) (see page 987)
- [dkp push](#) (see page 989)
- [dkp scale](#) (see page 992)
- [dkp serve](#) (see page 993)
- [dkp update](#) (see page 994)
- [dkp upgrade](#) (see page 1009)
- [dkp edit](#) (see page 1020)

---

<sup>776</sup> <https://kubernetes.io/docs/reference/generated/kubernetes-api/v1.25/#localobjectreference-v1-core>

- [dkp version](#) (see page 1021)
- [dkp cluster](#) (see page 1022)

## 9.2.2 dkp attach

Attach one of [cluster]

### 9.2.2.1 Options

```

 --config string Config file to use (default "/root/.kommander/
config")
 --context string The name of the kubeconfig context to use
-h, --help help for attach
 --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.

```

### 9.2.2.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.2.3 SEE ALSO

- [dkp attach cluster](#) (see page 920) - Attach a cluster

### 9.2.2.4 dkp attach cluster

Attach a cluster

```
dkp attach cluster -n NAME --attached-kubeconfig FILENAME [flags]
```

#### 9.2.2.4.1 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

```

<code>-w, --workspace string</code>	Name of the workspace of the attached cluster
-------------------------------------	-----------------------------------------------

#### 9.2.2.4.2 Options inherited from parent commands

<code>--config string</code>	Config file to use ( <b>default</b> <code>"/root/.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 <b>for</b> 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

#### 9.2.2.4.3 SEE ALSO

- [dkp attach](#) (see page 920) - Attach one of [cluster]

### 9.2.3 dkp check

Check, one of [cluster]

#### 9.2.3.1 Options

<code>-h, --help</code>	help <b>for</b> check
-------------------------	-----------------------

#### 9.2.3.2 Options inherited from parent commands

<code>-v, --verbose int</code>	Output verbosity
--------------------------------	------------------

#### 9.2.3.3 SEE ALSO

- [dkp check cluster](#) (see page 921) - Check a cluster, one of [fips]

#### 9.2.3.4 dkp check cluster

Check a cluster, one of [fips]



### 9.2.3.4.1 Options

```
-h, --help help for cluster
```

### 9.2.3.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.3.4.3 SEE ALSO

- [dkp check \(see page 921\)](#) - Check, one of [cluster]
- [dkp check cluster fips \(see page 922\)](#) - Validate the components in your cluster are FIPS compliant

### 9.2.3.4.4 **dkp check cluster fips**

Validate the components in your cluster are FIPS compliant

#### 9.2.3.4.4.1 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]
```

#### 9.2.3.4.4.2 Options

```
-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)
```

#### 9.2.3.4.4.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.3.4.4.4 SEE ALSO

- [dkp check cluster](#) (see page 921) - Check a cluster, one of [fips]

## 9.2.4 dkp completion

Generate the autocompletion script for the specified shell

### 9.2.4.1 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.

### 9.2.4.2 Options

```
-h, --help help for completion
```

### 9.2.4.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.4.4 SEE ALSO

- [dkp completion bash \(see page 925\)](#) - Generate the autocompletion script for bash
- [dkp completion fish \(see page 924\)](#) - Generate the autocompletion script for fish
- [dkp completion powershell \(see page 927\)](#) - Generate the autocompletion script for powershell
- [dkp completion zsh \(see page 926\)](#) - Generate the autocompletion script for zsh

### 9.2.4.5 dkp completion fish

Generate the autocompletion script for fish

#### 9.2.4.5.1 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]
```

#### 9.2.4.5.2 Options

```
-h, --help help for fish
--no-descriptions disable completion descriptions
```

### 9.2.4.5.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.4.5.4 SEE ALSO

- [dkp completion](#) (see page 923) - Generate the autocompletion script for the specified shell

## 9.2.4.6 dkp completion bash

Generate the autocompletion script for bash

### 9.2.4.6.1 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:

#### 9.2.4.6.1.1 Linux:

```
dkp completion bash > /etc/bash_completion.d/dkp
```

#### 9.2.4.6.1.2 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
```

### 9.2.4.6.2 Options

```
-h, --help help for bash
--no-descriptions disable completion descriptions
```

### 9.2.4.6.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.4.6.4 SEE ALSO

- [dkp completion \(see page 923\)](#) - Generate the autocompletion script for the specified shell

## 9.2.4.7 dkp completion zsh

Generate the autocompletion script for zsh

### 9.2.4.7.1 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:

#### 9.2.4.7.1.1 Linux:

```
dkp completion zsh > "${fpath[1]}/_dkp"
```

#### 9.2.4.7.1.2 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]
```

#### 9.2.4.7.2 Options

```
-h, --help help for zsh
--no-descriptions disable completion descriptions
```

#### 9.2.4.7.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.4.7.4 SEE ALSO

- [dkp completion](#) (see page 923) - Generate the autocompletion script for the specified shell

### 9.2.4.8 dkp completion powershell

Generate the autocompletion script for powershell

#### 9.2.4.8.1 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]
```

### 9.2.4.8.2 Options

```
-h, --help help for powershell
--no-descriptions disable completion descriptions
```

### 9.2.4.8.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.4.8.4 SEE ALSO

- [dkp completion](#) (see page 923) - Generate the autocompletion script for the specified shell

## 9.2.5 dkp config

Manage DKP Kommander's configuration

### 9.2.5.1 Options

```
--config string Config file to use (default "/root/.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.
```

### 9.2.5.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.5.3 SEE ALSO

- [dkp config get](#) (see page 929) - Retrieve DKP Kommander's configuration
- [dkp config set](#) (see page 930) - Modify DKP Kommander's configuration

## 9.2.5.4 dkp config get

Retrieve DKP Kommander's configuration

### 9.2.5.4.1 Options

```
-h, --help help for get
```

### 9.2.5.4.2 Options inherited from parent commands

<code>--config string</code>	Config file to use ( <b>default</b> <code>"/root/.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 <b>for</b> 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

### 9.2.5.4.3 SEE ALSO

- [dkp config](#) (see page 928) - Manage DKP Kommander's configuration
- [dkp config get default-workspace](#) (see page 929) - Displays the name of the configured default Workspace

### 9.2.5.4.4 dkp config get default-workspace

Displays the name of the configured default Workspace

```
dkp config get default-workspace [flags]
```

#### 9.2.5.4.4.1 Options

```
-h, --help help for default-workspace
```



#### 9.2.5.4.2 Options inherited from parent commands

<code>--config string</code>	Config file to use ( <b>default</b> <code>"/root/.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 <b>for</b> 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

#### 9.2.5.4.3 SEE ALSO

- [dkp config get](#) (see page 929) - Retrieve DKP Kommander's configuration

### 9.2.5.5 dkp config set

Modify DKP Kommander's configuration

#### 9.2.5.5.1 Options

```
-h, --help help for set
```

#### 9.2.5.5.2 Options inherited from parent commands

<code>--config string</code>	Config file to use ( <b>default</b> <code>"/root/.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 <b>for</b> 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

#### 9.2.5.5.3 SEE ALSO

- [dkp config](#) (see page 928) - Manage DKP Kommander's configuration
- [dkp config set default-workspace](#) (see page 931) - Set the name of the default Workspace to use in all commands when none is provided

### 9.2.5.5.4 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]
```

#### 9.2.5.5.4.1 Options

```
-h, --help help for default-workspace
```

#### 9.2.5.5.4.2 Options inherited from parent commands

```

--config string Config file to use (default "/root/.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

```

#### 9.2.5.5.4.3 SEE ALSO

- [dkp config set \(see page 930\)](#) - Modify DKP Kommander's configuration

## 9.2.6 dkp create

Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

### 9.2.6.1 Options

```
-h, --help help for create
```

## 9.2.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

## 9.2.6.3 SEE ALSO

- [dkp create appdeployment](#) (see page 933) - Create an AppDeployment
- [dkp create bootstrap](#) (see page 936) - Create bootstrap cluster
- [dkp create capi-components](#) (see page 934) - Create the CAPI components in the cluster
- [dkp create chart-bundle](#) (see page 937) - Create charts bundle based on a catalog applications git repository
- [dkp create cluster](#) (see page 938) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]
- [dkp create image-bundle](#) (see page 935) - Create an image bundle
- [dkp create nodepool](#) (see page 956) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]
- [dkp create workspace](#) (see page 932) - Create a Workspace

## 9.2.6.4 dkp create workspace

Create a Workspace

```
dkp create workspace WORKSPACE_NAME [flags]
```

### 9.2.6.4.1 Options

<code>--config string</code>	Config file to use ( <b>default</b> <code>"/root/.kommander/config"</code> )
<code>--context string</code>	The name of the kubeconfig context to use
<code>--dry-run</code>	Export in YAML format to stdout
<code>-h, --help</code>	help <b>for</b> workspace
<code>--kubeconfig string</code>	Path to the kubeconfig file to use <b>for</b> CLI requests.
<code>-n, --namespace string</code>	Name of the Namespace to create <b>for</b> the workspace
<code>-o, --output string</code>	Output format. One of: yml json ( <b>default</b> <code>"yaml"</code> )
<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.

### 9.2.6.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.6.4.3 SEE ALSO

- [dkp create](#) (see page 931) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

## 9.2.6.5 dkp create appdeployment

Create an AppDeployment

### 9.2.6.5.1 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]
```

### 9.2.6.5.2 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 "/
root/.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.
-w, --workspace string Name of the workspace to create the
AppDeployment in

```

### 9.2.6.5.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.6.5.4 SEE ALSO

- [dkp create](#) (see page 931) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

## 9.2.6.6 dkp create capi-components

Create the CAPI components in the cluster

```
dkp create capi-components [flags]
```

### 9.2.6.6.1 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. 1s, 2m, 3h). (default 10m0s)
--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.

### 9.2.6.6.2 Options inherited from parent commands

`-v, --verbose` **int** Output verbosity

### 9.2.6.6.3 SEE ALSO

- [dkp create](#) (see page 931) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

## 9.2.6.7 dkp create image-bundle

Create an image bundle

```
dkp create image-bundle [flags]
```

### 9.2.6.7.1 Options

```
-h, --help help for image-bundle
--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])
```

### 9.2.6.7.2 Options inherited from parent commands

`-v, --verbose` **int** Output verbosity

### 9.2.6.7.3 SEE ALSO

- [dkp create](#) (see page 931) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

### 9.2.6.8 dkp create bootstrap

Create bootstrap cluster

```
dkp create bootstrap [flags]
```

#### 9.2.6.8.1 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:v2.4.2-rc.2")
--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. 1s, 2m, 3h). (default 10m0s)
--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.

```

#### 9.2.6.8.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.6.8.3 SEE ALSO

- [dkp create](#) (see page 931) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

### 9.2.6.9 dkp create chart-bundle

Create charts bundle based on a catalog applications git repository

```
dkp create chart-bundle [flags]
```

#### 9.2.6.9.1 Options

```

 --catalog-repository string Git repository containing catalog
application definitions
 --config string Config file to use (default "/"
root/.kommander/config")
 --context string The name of the kubeconfig context to use
 --dry-run 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 [])

```

#### 9.2.6.9.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```



### 9.2.6.9.3 SEE ALSO

- [dkp create](#) (see page 931) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]

### 9.2.6.10 dkp create cluster

Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

#### 9.2.6.10.1 Options

```
-h, --help help for cluster
```

#### 9.2.6.10.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.6.10.3 SEE ALSO

- [dkp create](#) (see page 931) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]
- [dkp create cluster aks](#) (see page 949) - Create a Konvoy cluster in AKS
- [dkp create cluster aws](#) (see page 940) - Create a Konvoy cluster in AWS
- [dkp create cluster azure](#) (see page 943) - Create a Konvoy cluster in Azure
- [dkp create cluster eks](#) (see page 938) - Create a Konvoy cluster in EKS
- [dkp create cluster gcp](#) (see page 954) - Create a Konvoy cluster in GCP
- [dkp create cluster preprovisioned](#) (see page 951) - Create a Konvoy cluster on pre-provisioned infrastructure
- [dkp create cluster vsphere](#) (see page 946) - Create a Konvoy cluster in vSphere

#### 9.2.6.10.4 dkp create cluster eks

Create a Konvoy cluster in EKS

```
dkp create cluster eks [flags]
```

## 9.2.6.10.4.1 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 stringWithString 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.
--etcd-image-repository string The image repository to use for
pulling the etcd image
--etcd-version string The version of etcd to use.
Overriding kubeadm's default value as etcd v3.5.x is not recommended for production
use. This default value will removed in a future release once etcd is fixed. (default
"3.4.13-0")
-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:v2.4.2-rc.2")
--kubeconfig string Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply.
--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 "1.23.12")
-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).
--region string AWS region to deploy cluster to
(default "us-west-2")
--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. 1s, 2m, 3h). (default 10m0s)
--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)

```

#### 9.2.6.10.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.6.10.4.3 SEE ALSO

- [dkp create cluster](#) (see page 938) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

#### 9.2.6.10.5 dkp create cluster aws

Create a Konvoy cluster in AWS

```
dkp create cluster aws [flags]
```

## 9.2.6.10.5.1 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 (default [])
--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 for Lookup search (ex.
'centos-7', 'ubuntu-18.04', 'ubuntu-20.04') (default "ubuntu-20.04")
--ami-format string Lookup Format string to generate
AMI search name from (default "capa-ami-{{.BaseOS}}-?{{.K8sVersion}}-*")
--ami-owner string Owner ID for AMI Lookup search (d
efault "258751437250")
--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. (default 0)
-c, --cluster-name name Name used to prefix the cluster
and all the created resources.
--control-plane-http-proxy string HTTP proxy for control plane
machines
--control-plane-https-proxy string HTTPS proxy for control plane
machines
--control-plane-iam-instance-profile string Name of the IAM instance profile
to assign to control plane machines. (default "control-plane.cluster-api-provider-
aws.sigs.k8s.io")
--control-plane-instance-type string Control Plane machine instance
type (default "m5.xlarge")
--control-plane-no-proxy strings No Proxy list for control plane
machines (default [])
--control-plane-replicas int Number of control plane replicas
(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.
Overriding kubeadm's default value as etcd v3.5.x is not recommended for production
use. This default value will removed in a future release once etcd is fixed. (default
"3.4.13-0")
--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:v2.4.2-rc.2")
--kubeconfig string Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply.
--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
"1.24.6")
-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).
--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 string 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. 1s, 2m, 3h). (default 10m0s)

```

<code>--vpc-id string</code>	Existing VPC ID to use <b>for</b> the cluster
<code>--wait</code>	If <b>true</b> , wait <b>for</b> operations to complete before returning. ( <b>default true</b> )
<code>--with-aws-bootstrap-credentials</code>	Set <b>true</b> to use AWS bootstrap credentials from your environment. When <b>false</b> , 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 <b>true</b> to use GCP bootstrap credentials from your environment. When <b>false</b> , the service account of the VM instance where the CAPG controller is scheduled on will be used instead.
<code>--worker-availability-zone string</code>	The AvailabilityZone in the region to deploy the worker nodes to, <b>if</b> not set a random one will be selected (ex. us-west-2a)
<code>--worker-http-proxy string</code>	HTTP proxy <b>for</b> nodes
<code>--worker-https-proxy string</code>	HTTPS proxy <b>for</b> nodes
<code>--worker-iam-instance-profile string</code>	Name of the IAM instance profile to assign to worker machines. ( <b>default "nodes.cluster-api-provider-aws.sigs.k8s.io"</b> )
<code>--worker-instance-type string</code>	Worker machine instance type ( <b>default "m5.2xlarge"</b> )
<code>--worker-no-proxy strings</code>	No Proxy list <b>for</b> nodes ( <b>default []</b> )
<code>--worker-replicas int</code>	Number of workers ( <b>default 4</b> )

#### 9.2.6.10.5.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.6.10.5.3 SEE ALSO

- [dkp create cluster](#) (see page 938) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

#### 9.2.6.10.6 dkp create cluster azure

Create a Konvoy cluster in Azure

```
dkp create cluster azure [flags]
```

##### 9.2.6.10.6.1 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 goyaml 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 replicas (defaul
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. Overriding
kubeadm's default value as etcd v3.5.x is not recommended for production use. This
default value will be removed in a future release once etcd is fixed. (default
"3.4.13-0")
--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:v2.4.2-rc.2")
--kubeconfig string Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply.
--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 "1.24.6")
--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 (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).
--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 string 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 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. 1s, 2m, 3h). (default 10m0s)
--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 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)

```



### 9.2.6.10.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.6.10.6.3 SEE ALSO

- [dkp create cluster](#) (see page 938) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.10.7 dkp create cluster vsphere

Create a Konvoy cluster in vSphere

```
dkp create cluster vsphere [flags]
```

#### 9.2.6.10.7.1 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.
--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 replicas (default
t 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. Overriding
kubeadm's default value as etcd v3.5.x is not recommended for production use. This
default value will be removed in a future release once etcd is fixed. (default
"3.4.13-0")
--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:v2.4.2-rc.2")
--kubeconfig string Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply.
--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 "1.24.6")
-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 (default
ult [])
-o, --output string Output format. One of: (json, yaml,
name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-
json, jsonpath-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 string 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 IP or FQDN.
--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. 1s, 2m, 3h). (default 10m0s)
--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. (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-cpus int The number of virtual processors in a
worker machine (default 8)
--worker-disk-size int The size of a worker machine's disk, in
GB (default 80)
--worker-http-proxy string HTTP proxy for nodes
--worker-https-proxy string HTTPS proxy for nodes
--worker-memory int The size of a worker machine's memory,
in GB (default 32)
--worker-no-proxy strings No Proxy list for nodes (default [])
--worker-replicas int Number of workers (default 4)

```

### 9.2.6.10.7.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.6.10.7.3 SEE ALSO

- [dkp create cluster](#) (see page 938) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.10.8 dkp create cluster aks

Create a Konvoy cluster in AKS

```
dkp create cluster aks [flags]
```

#### 9.2.6.10.8.1 Options

<code>--additional-tags</code> <code>stringToString</code>	Tags to apply to the provisioned infrastructure ( <b>default</b> [])
<code>--allow-missing-template-keys</code>	If <b>true</b> , 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. ( <b>default true</b> )
<code>--aws-service-endpoints</code> <code>string</code>	Custom AWS service endpoints in a semi-colon separated format: <code>\${SigningRegion1}:\${ServiceID1}=\${URL},\${ServiceID2}=\${URL};\${SigningRegion2}...</code>
<code>--certificate-renew-interval</code> <code>int</code>	The interval number of days Kubernetes managed PKI certificates are renewed. For example, an Interval value of <b>30</b> means the certificates will be refreshed every <b>30</b> days. A value of <b>0</b> disables the feature. ( <b>default 0</b> )
<code>-c, --cluster-name</code> <code>name</code>	Name used to prefix the cluster and all the created resources.
<code>--control-plane-http-proxy</code> <code>string</code>	HTTP proxy <b>for</b> control plane machines
<code>--control-plane-https-proxy</code> <code>string</code>	HTTPS proxy <b>for</b> control plane machines
<code>--control-plane-machine-size</code> <code>string</code>	Control Plane machine size ( <b>default "Standard_D4s_v3"</b> )
<code>--control-plane-no-proxy</code> <code>strings</code>	No Proxy list <b>for</b> control plane machines ( <b>default</b> [])
<code>--control-plane-replicas</code> <code>int</code>	Number of control plane replicas ( <b>default 3</b> )
<code>--dry-run</code>	Only print the objects that would be created, without creating them.
<code>--etcd-image-repository</code> <code>string</code>	The image repository to use <b>for</b> pulling the etcd image

```

--etcd-version string The version of etcd to use. Overriding
kubeadm's default value as etcd v3.5.x is not recommended for production use. This
default value will be removed in a future release once etcd is fixed. (default
"3.4.13-0")
--extra-sans strings A comma separated list of additional
Subject Alternative Names for the API Server signing cert (default [])
-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:v2.4.2-rc.2")
--kubeconfig string Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply.
--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. 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 v1.24.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 (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).
--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 string 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.
--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. 1s, 2m, 3h). (default 10m0s)

```

<code>--wait</code>	If <b>true</b> , wait <b>for</b> operations to complete before returning. ( <b>default true</b> )
<code>--with-aws-bootstrap-credentials</code>	Set <b>true</b> to use AWS bootstrap credentials from your environment. When <b>false</b> , 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 <b>true</b> to use GCP bootstrap credentials from your environment. When <b>false</b> , 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 <b>for</b> nodes
<code>--worker-https-proxy string</code>	HTTPS proxy <b>for</b> nodes
<code>--worker-machine-size string</code>	Worker machine size ( <b>default "Standard_D8s_v3"</b> )
<code>--worker-no-proxy strings</code>	No Proxy list <b>for</b> nodes ( <b>default []</b> )
<code>--worker-replicas int</code>	Number of workers ( <b>default 4</b> )

### 9.2.6.10.8.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

### 9.2.6.10.8.3 SEE ALSO

- [dkp create cluster](#) (see page 938) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.10.9 dkp create cluster preprovisioned

Create a Konvoy cluster on pre-provisioned infrastructure

```
dkp create cluster preprovisioned [flags]
```

#### 9.2.6.10.9.1 Options

<code>--additional-tags stringToString</code>	Tags to apply to the provisioned infrastructure ( <b>default []</b> )
<code>--allow-missing-template-keys</code>	If <b>true</b> , ignore any errors in templates when a field or map key is missing in the template. Only applies to goyaml and jsonpath output formats. ( <b>default true</b> )
<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>--certificate-renew-interval int</code>	The interval number of days Kubernetes managed PKI certificates are renewed. For example, an Interval value of <b>30</b>

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.

--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-no-proxy strings No Proxy list **for** control plane machines (**default** [])

--control-plane-replicas **int** Number of control plane replicas (**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. Overriding kubeadm's **default** value as etcd v3.5.x is not recommended **for** production use. This **default** value will be removed in a future release once etcd is fixed. (**default** "3.4.13-0")

--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:v2.4.2-rc.2")

--kubeconfig string Path to the kubeconfig **for** the management cluster. If unspecified, **default** discovery rules apply.

--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** "1.24.6")

-n, --namespace string If present, the namespace scope **for** **this** CLI request. (**default** "default")

--no-proxy strings No Proxy list **for** CAPI controllers (**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).

```

--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 string 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-private-key-file string Path to the private SSH key file used
by Konvoy to access the pre-provisioned hosts
--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. 1s, 2m, 3h). (default 10m0s)
--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.
--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-http-proxy string HTTP proxy for nodes
--worker-https-proxy string HTTPS proxy for nodes
--worker-no-proxy strings No Proxy list for nodes (default [])
--worker-replicas int Number of workers (default 4)

```



### 9.2.6.10.9.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.6.10.9.3 SEE ALSO

- [dkp create cluster](#) (see page 938) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.10.10 dkp create cluster gcp

Create a Konvoy cluster in GCP

```
dkp create cluster gcp [flags]
```

#### 9.2.6.10.10.1 Options

<code>--additional-tags</code> stringToString infrastructure ( <b>default</b> [])	Tags to apply to the provisioned
<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. ( <b>default true</b> )	If <b>true</b> , ignore any errors in
<code>--associate-public-ip-address</code> machines. When set to <b>false</b> the specified network must have Cloud NAT configured to provide internet access. ( <b>default true</b> )	Associate a <b>public</b> IP <b>for</b> all
<code>--aws-service-endpoints</code> string a semi-colon separated format: <code>\${SigningRegion1}:\${ServiceID1}=\${URL},\${ServiceID2}=\${URL};\${SigningRegion2}...</code>	Custom AWS service endpoints in
<code>--certificate-renew-interval</code> <b>int</b> Kubernetes managed PKI certificates are renewed. For example, an Interval value of <b>30</b> means the certificates will be refreshed every <b>30</b> days. A value of <b>0</b> disables the feature. ( <b>default 0</b> )	The interval number of days
<code>-c, --cluster-name</code> name and all the created resources.	Name used to prefix the cluster
<code>--control-plane-http-proxy</code> string machines	HTTP proxy <b>for</b> control plane
<code>--control-plane-https-proxy</code> string machines	HTTPS proxy <b>for</b> control plane
<code>--control-plane-instance-type</code> string type ( <b>default "n2-standard-4"</b> )	Control Plane machine instance
<code>--control-plane-no-proxy</code> strings machines ( <b>default</b> [])	No Proxy list <b>for</b> control plane

```

--control-plane-replicas int Number of control plane replicas
(default 3)
--control-plane-service-account-email string Control Plane Service Account
email address (default "default")
--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.
Overriding kubeadm's default value as etcd v3.5.x is not recommended for production
use. This default value will be removed in a future release once etcd is fixed. (default
"3.4.13-0")
--extra-sans strings A comma separated list of
additional Subject Alternative Names for the API Server signing cert (default [])
-h, --help help for gcp
--http-proxy string HTTP proxy for CAPI controllers
--https-proxy string HTTPS proxy for CAPI controllers
--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}}')
--kind-cluster-image string Kind node image for the
bootstrap cluster (default "mesosphere/konvoy-bootstrap:v2.4.2-rc.2")
--kubeconfig string Path to the kubeconfig for the
management cluster. If unspecified, default discovery rules apply.
--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
"1.24.6")
-n, --namespace string If present, the namespace scope
for this CLI request. (default "default")
--network string The GCP network name to deploy
the cluster to (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).
--project string The GCP project name to deploy
the cluster to
--region string GCP region to deploy cluster to
(default "us-west1")
--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

```

<code>--registry-mirror-url</code> string	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 <b>true</b> , 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 <b>false</b> , a management cluster is used. <b>(default false)</b>
<code>--show-managed-fields</code>	If <b>true</b> , keep the managedFields when printing objects in JSON or YAML format.
<code>--ssh-public-key-file</code> string	Path to the authorized SSH key <b>for</b> the user
<code>--ssh-username</code> string	Name of the user to create on the instance <b>(default "konvoy")</b>
<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 goolang templates [ <a href="http://golang.org/pkg/text/template/#pkg-overview">http://golang.org/pkg/text/template/#pkg-overview</a> ].
<code>--timeout</code> duration	The length of time to wait before giving up. Zero means wait forever (e.g. 1s, 2m, 3h). <b>(default 10m0s)</b>
<code>--wait</code>	If <b>true</b> , wait <b>for</b> operations to complete before returning. <b>(default true)</b>
<code>--with-aws-bootstrap-credentials</code>	Set <b>true</b> to use AWS bootstrap credentials from your environment. When <b>false</b> , 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 <b>true</b> to use GCP bootstrap credentials from your environment. When <b>false</b> , 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 <b>for</b> nodes
<code>--worker-https-proxy</code> string	HTTPS proxy <b>for</b> nodes
<code>--worker-instance-type</code> string	Worker machine instance type <b>(default "n2-standard-8")</b>
<code>--worker-no-proxy</code> strings	No Proxy list <b>for</b> nodes <b>(default [])</b>
<code>--worker-replicas</code> <b>int</b>	Number of workers <b>(default 4)</b>
<code>--worker-service-account-email</code> string	Worker machine Service Account email address <b>(default "default")</b>
<code>--worker-zone</code> string	Zone in the region to deploy the worker nodes to, <b>if</b> not set a random one will be selected (ex. us-west1-a)

#### 9.2.6.10.10.2 Options inherited from parent commands

`-v`, `--verbose` **int** Output verbosity

#### 9.2.6.10.10.3 SEE ALSO

- [dkp create cluster](#) (see page 938) - Create a Kubernetes cluster, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.11 dkp create nodepool

Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

#### 9.2.6.11.1 Options

```
-h, --help help for nodepool
```

#### 9.2.6.11.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.6.11.3 SEE ALSO

- [dkp create](#) (see page 931) - Create one of [appdeployment, bootstrap, capi-components, chart-bundle, cluster, image-bundle, nodepool, workspace]
- [dkp create nodepool aks](#) (see page 961) - Create a nodepool in AKS
- [dkp create nodepool aws](#) (see page 962) - Create a nodepool in AWS
- [dkp create nodepool azure](#) (see page 956) - Create a nodepool in Azure
- [dkp create nodepool eks](#) (see page 964) - Create a nodepool for EKS
- [dkp create nodepool gcp](#) (see page 965) - Create a nodepool in GCP
- [dkp create nodepool preprovisioned](#) (see page 967) - Create a nodepool for a Pre Provisioned Cluster.
- [dkp create nodepool vsphere](#) (see page 959) - Create a nodepool in vSphere

#### 9.2.6.11.4 dkp create nodepool azure

Create a nodepool in Azure

```
dkp create nodepool azure name [flags]
```

##### 9.2.6.11.4.1 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)
```

```

--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
--kubecfg string Path to the kubecfg for the management
cluster. If unspecified, default discovery rules apply.
--kubernetes-version string Kubernetes version (default "1.24.6")
--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).
--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 string 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. 1s, 2m, 3h). (default 30m0s)
--wait If true, wait for operations to complete
before returning.

```

### 9.2.6.11.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.6.11.4.3 SEE ALSO

- [dkp create nodepool](#) (see page 956) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.11.5 dkp create nodepool vsphere

Create a nodepool in vSphere

```
dkp create nodepool vsphere name [flags]
```

#### 9.2.6.11.5.1 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 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 "1.24.6")

```

<code>--memory</code> <b>int</b>	The size of a worker machine's memory, in GB ( <b>default</b> 32)
<code>-n, --namespace</code> string	If present, the namespace scope <b>for this</b> CLI request. ( <b>default</b> "default")
<code>--network</code> string	The vSphere network to deploy the workload cluster on.
<code>--no-proxy</code> strings	No Proxy list <b>for</b> nodes ( <b>default</b> [])
<code>-o, --output</code> string	Output format. One of: (json, yaml, name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json, jsonpath-file).
<code>--registry-mirror-cacert</code> file	CA certificate chain to use <b>while</b> communicating with the registry mirror using TLS
<code>--registry-mirror-password</code> string	Password to authenticate to the registry mirror with
<code>--registry-mirror-url</code> string	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> <b>int</b>	Number of replicas ( <b>default</b> 1)
<code>--resource-pool</code> string	The vSphere resource pool <b>for</b> the workload cluster's virtual machines.
<code>--server</code> string	The vCenter server IP or FQDN.
<code>--show-managed-fields</code>	If <b>true</b> , keep the managedFields when printing objects in JSON or YAML format.
<code>--ssh-<b>public</b>-key-file</code> string	Path to the authorized SSH key <b>for</b> the user
<code>--ssh-username</code> string	Name of the user to create on the instance ( <b>default</b> "konvoy")
<code>--storage-policy</code> string	This is the vSphere storage policy. Set it to "" <b>if</b> you don't want to use a storage policy.
<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 golang templates [ <a href="http://golang.org/pkg/text/template/#pkg-overview">http://golang.org/pkg/text/template/#pkg-overview</a> ].
<code>--timeout</code> duration	The length of time to wait before giving up. Zero means wait forever (e.g. 1s, 2m, 3h). ( <b>default</b> 30m0s)
<code>--tls-thumb-print</code> string	sha1 thumbprint of the vcenter certificate: <code>openssl x509 -sha1 -fingerprint -in ca.crt -noout</code>
<code>--vm-template</code> string	The virtual machine template to use <b>for</b> the workload cluster's virtual machines.
<code>--wait</code>	If <b>true</b> , wait <b>for</b> operations to complete before returning.

### 9.2.6.11.5.2 Options inherited from parent commands

`-v, --verbose` **int** Output verbosity

### 9.2.6.11.5.3 SEE ALSO

- [dkp create nodepool](#) (see page 956) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.11.6 dkp create nodepool aks

Create a nodepool in AKS

```
dkp create nodepool aks name [flags]
```

#### 9.2.6.11.6.1 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 goyaml 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 v1.24.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).
--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. 1s, 2m, 3h). (default 30m0s)
--wait If true, wait for operations to complete
before returning.

```

### 9.2.6.11.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.6.11.6.3 SEE ALSO

- [dkp create nodepool](#) (see page 956) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.11.7 dkp create nodepool aws

Create a nodepool in AWS

```
dkp create nodepool aws [flags]
```

#### 9.2.6.11.7.1 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 go lang
and jsonpath output formats. (default true)
--ami string AMI ID to use for machines
--ami-base-os string Base OS for Lookup search (ex.
'centos-7', 'ubuntu-18.04', 'ubuntu-20.04') (default "ubuntu-20.04")
--ami-format string Lookup Format string to generate AMI
search name from (default "capa-ami-{{.BaseOS}}-?{{.K8sVersion}}-*")
--ami-owner string Owner ID for AMI Lookup search (defau
lt "258751437250")
--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 <b>for</b> aws
--http-proxy string	HTTP proxy <b>for</b> nodes
--https-proxy string	HTTPS proxy <b>for</b> nodes
--iam-instance-profile string	Name of the IAM instance profile to assign to worker machines. ( <b>default</b> "nodes.cluster-api-provider-aws.sigs.k8s.io")
--instance-type string	Worker machine instance type ( <b>default</b> "m5.2xlarge")
--kubeconfig string	Path to the kubeconfig <b>for</b> the management cluster. If unspecified, <b>default</b> discovery rules apply.
--kubernetes-version string	Kubernetes version ( <b>default</b> "1.24.6")
-n, --namespace string	If present, the namespace scope <b>for</b> <b>this</b> CLI request. ( <b>default</b> "default")
--no-proxy strings	No Proxy list <b>for</b> nodes ( <b>default</b> [])
-o, --output string	Output format. One of: (json, yaml, name, go-template, go-template-file, template, templatefile, jsonpath, jsonpath-as-json, jsonpath-file).
--registry-mirror-cacert file	CA certificate chain to use <b>while</b> communicating with the registry mirror using TLS
--registry-mirror-password string	Password to authenticate to the registry mirror with
--registry-mirror-url string	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 <b>int</b>	Number of replicas ( <b>default</b> 1)
--show-managed-fields	If <b>true</b> , keep the managedFields when printing objects in JSON or YAML format.
--ssh- <b>public</b> -key-file string	Path to the authorized SSH key <b>for</b> the user
--ssh-username string	Name of the user to create on the instance ( <b>default</b> "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 [ <a href="http://golang.org/pkg/text/template/#pkg-overview">http://golang.org/pkg/text/template/#pkg-overview</a> ].
--timeout duration	The length of time to wait before giving up. Zero means wait forever (e.g. 1s, 2m, 3h). ( <b>default</b> 30m0s)
--wait	If <b>true</b> , wait <b>for</b> operations to complete before returning.

### 9.2.6.11.7.2 Options inherited from parent commands

-v, --verbose **int** Output verbosity

### 9.2.6.11.7.3 SEE ALSO

- [dkp create nodepool](#) (see page 956) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.11.8 dkp create nodepool eks

Create a nodepool for EKS

```
dkp create nodepool eks [flags]
```

#### 9.2.6.11.8.1 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)
 --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 "1.23.12")
 -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).
 --replicas int Number of replicas (default 1)
 --show-managed-fields 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 <b>for</b> the user
<code>--ssh-username</code> string	Name of the user to create on the instance ( <b>default</b> "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 [ <a href="http://golang.org/pkg/text/template/#pkg-overview">http://golang.org/pkg/text/template/#pkg-overview</a> ].
<code>--timeout</code> duration	The length of time to wait before giving up. Zero means wait forever (e.g. 1s, 2m, 3h). ( <b>default</b> 30m0s)
<code>--wait</code>	If <b>true</b> , wait <b>for</b> operations to complete before returning.

### 9.2.6.11.8.2 Options inherited from parent commands

`-v, --verbose` **int** Output verbosity

### 9.2.6.11.8.3 SEE ALSO

- [dkp create nodepool](#) (see page 956) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.11.9 dkp create nodepool gcp

Create a nodepool in GCP

```
dkp create nodepool gcp name [flags]
```

#### 9.2.6.11.9.1 Options

<code>--additional-tags</code> stringToString	Tags to apply to the provisioned infrastructure ( <b>default</b> [])
<code>--allow-missing-template-keys</code>	If <b>true</b> , 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. ( <b>default true</b> )
<code>--associate-public-ip-address</code>	Associate a <b>public</b> IP <b>for</b> all machines. When set to <b>false</b> the specified network must have Cloud NAT configured to provide internet access. ( <b>default true</b> )
<code>-c, --cluster-name</code> name	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 <b>for</b> gcp
<code>--http-proxy</code> string	HTTP proxy <b>for</b> 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 "1.24.6")
-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).
--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 string 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 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. 1s, 2m, 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)

```

### 9.2.6.11.9.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.6.11.9.3 SEE ALSO

- [dkp create nodepool](#) (see page 956) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.6.11.10 dkp create nodepool preprovisioned

Create a nodepool for a Pre Provisioned Cluster.

```
dkp create nodepool preprovisioned name [flags]
```

#### 9.2.6.11.10.1 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 "1.24.6")
 -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).
 --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 string 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. 1s, 2m, 3h). (default 30m0s)
--wait If true, wait for operations to complete
before returning.

```

#### 9.2.6.11.10.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.6.11.10.3 SEE ALSO

- [dkp create nodepool](#) (see page 956) - Create a nodepool, one of [aks, aws, azure, eks, gcp, preprovisioned, vsphere]

## 9.2.7 dkp delete

Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

### 9.2.7.1 Options

```
-h, --help help for delete
```

### 9.2.7.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.7.3 SEE ALSO

- [dkp delete bootstrap](#) (see page 969) - Delete bootstrap cluster
- [dkp delete capi-components](#) (see page 969) - Delete the CAPI components from the cluster
- [dkp delete chart](#) (see page 971) - Delete a chart from the repository
- [dkp delete cluster](#) (see page 970) - Delete a Kubernetes cluster
- [dkp delete nodepool](#) (see page 972) - Delete a nodepool for a given cluster

### 9.2.7.4 dkp delete bootstrap

Delete bootstrap cluster

```
dkp delete bootstrap [flags]
```

#### 9.2.7.4.1 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. 1s, 2m, 3h). (default 5m0s)
--wait If true, wait for operations to complete before
returning. (default true)
```

#### 9.2.7.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.7.4.3 SEE ALSO

- [dkp delete](#) (see page 968) - Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

### 9.2.7.5 dkp delete capi-components

Delete the CAPI components from the cluster



```
dkp delete capi-components [flags]
```

### 9.2.7.5.1 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. 1s, 2m, 3h). (default 5m0s)
--wait If true, wait for operations to complete before
returning. (default true)
```

### 9.2.7.5.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.7.5.3 SEE ALSO

- [dkp delete](#) (see page 968) - Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

## 9.2.7.6 dkp delete cluster

Delete a Kubernetes cluster

```
dkp delete cluster [flags]
```

### 9.2.7.6.1 Options

```
--aws-service-endpoints string Custom AWS service endpoints in a semi-colon
separated format: ${SigningRegion1}:${ServiceID1}=${URL},${ServiceID2}=${URL};$
${SigningRegion2}...
-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 (default "mesosphere/konvoy-bootstrap:v2.4.2-rc.2")
--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. 1s, 2m, 3h). (default 15m0s)
--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.

```

### 9.2.7.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.7.6.3 SEE ALSO

- [dkp delete](#) (see page 968) - Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

### 9.2.7.7 dkp delete chart

Delete a chart from the repository

```
dkp delete chart [chartName] [chartVersion] [flags]
```

#### 9.2.7.7.1 Options

```

--config string Config file to use (default "/root/.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.
```

### 9.2.7.7.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.7.7.3 SEE ALSO

- [dkp delete](#) (see page 968) - Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

## 9.2.7.8 dkp delete nodepool

Delete a nodepool for a given cluster

```
dkp delete nodepool name [flags]
```

### 9.2.7.8.1 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")
```

### 9.2.7.8.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.7.8.3 SEE ALSO

- [dkp delete](#) (see page 968) - Delete one of [bootstrap, capi-components, chart, cluster, nodepool]

## 9.2.8 dkp describe

Describe one of [cluster]

### 9.2.8.1 Options

```
-h, --help help for describe
```

### 9.2.8.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.8.3 SEE ALSO

- [dkp describe cluster \(see page 973\)](#) - Describe a Kubernetes cluster status

### 9.2.8.4 dkp describe cluster

Describe a Kubernetes cluster status

```
dkp describe cluster [flags]
```

#### 9.2.8.4.1 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. (de
fault "default")
```

#### 9.2.8.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.8.4.3 SEE ALSO

- [dkp describe](#) (see page 973) - Describe one of [cluster]

## 9.2.9 dkp detach

Detach one of [cluster]

### 9.2.9.1 Options

```

--config string Config file to use (default "/root/.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.
```

### 9.2.9.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.9.3 SEE ALSO

- [dkp detach cluster](#) (see page 974) - Detach a cluster

## 9.2.9.4 dkp detach cluster

Detach a cluster

```
dkp detach cluster CLUSTER_NAME [flags]
```

### 9.2.9.4.1 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

```

#### 9.2.9.4.2 Options inherited from parent commands

```

--config string Config file to use (default "/root/.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

```

#### 9.2.9.4.3 SEE ALSO

- [dkp detach](#) (see page 974) - Detach one of [cluster]

## 9.2.10 dkp diagnose

Generate a support bundle

### 9.2.10.1 Synopsis

A support bundle is an archive of files, output, metrics and state from a server that can be used to assist when troubleshooting a Kubernetes cluster.

```
dkp diagnose [flags]
```

### 9.2.10.2 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
--cache-dir string Default cache directory (default "/root/.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
-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

```

### 9.2.10.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.10.4 SEE ALSO

- [dkp diagnose default-config](#) (see page 976) - Prints the default configuration of the diagnostics bundle collectors
- [dkp diagnose ssh](#) (see page 976) - Collect node-level diagnostics data over SSH

### 9.2.10.5 dkp diagnose ssh

Collect node-level diagnostics data over SSH

```
dkp diagnose ssh path/to/inventory-file.yaml [flags]
```

#### 9.2.10.5.1 Options

```
-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)
```

#### 9.2.10.5.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.10.5.3 SEE ALSO

- [dkp diagnose](#) (see page 975) - Generate a support bundle

### 9.2.10.6 dkp diagnose default-config

Prints the default configuration of the diagnostics bundle collectors

```
dkp diagnose default-config [flags]
```

#### 9.2.10.6.1 Options

```
-h, --help help for default-config
```

#### 9.2.10.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```



### 9.2.10.6.3 SEE ALSO

- [dkp diagnose](#) (see page 975) - Generate a support bundle

## 9.2.11 dkp get

Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

### 9.2.11.1 Options

```
-h, --help help for get
```

### 9.2.11.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.11.3 SEE ALSO

- [dkp get appdeployments](#) (see page 979) - Get AppDeployments from a Workspace, Project, or all Workspaces and Projects
- [dkp get chart](#) (see page 981) - Obtain information about charts stored in the repository
- [dkp get clusters](#) (see page 982) - Get clusters from specified Workspace
- [dkp get kubeconfig](#) (see page 978) - Retrieve cluster kubeconfig and modify local kubeconfig file
- [dkp get nodepools](#) (see page 981) - Get nodepools for a given cluster
- [dkp get workspaces](#) (see page 980) - Get Workspaces

### 9.2.11.4 dkp get kubeconfig

Retrieve cluster kubeconfig and modify local kubeconfig file

```
dkp get kubeconfig [flags]
```

#### 9.2.11.4.1 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

```

#### 9.2.11.4.2 Options inherited from parent commands

```

-v, --verbose int Output verbosity

```

#### 9.2.11.4.3 SEE ALSO

- [dkp get \(see page 978\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

### 9.2.11.5 dkp get appdeployments

Get AppDeployments from a Workspace, Project, or all Workspaces and Projects

#### 9.2.11.5.1 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]

```

#### 9.2.11.5.2 Options

```

-A, --all-namespaces If present, list the requested object(s) across all
namespaces.
 --config string Config file to use (default "/root/.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.
-w, --workspace string Name of the workspace to show AppDeployments from

```

### 9.2.11.5.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.11.5.4 SEE ALSO

- [dkp get \(see page 978\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

## 9.2.11.6 dkp get workspaces

Get Workspaces

```
dkp get workspaces [flags]
```

### 9.2.11.6.1 Options

```

-A, --all-namespaces If present, list the requested object(s) across all
 namespaces.
 --config string Config file to use (default "/root/.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.

```

### 9.2.11.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.11.6.3 SEE ALSO

- [dkp get \(see page 978\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

### 9.2.11.7 dkp get nodepools

Get nodepools for a given cluster

```
dkp get nodepools [flags]
```

#### 9.2.11.7.1 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. (de
fault "default")
```

#### 9.2.11.7.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.11.7.3 SEE ALSO

- [dkp get \(see page 978\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

### 9.2.11.8 dkp get chart

Obtain information about charts stored in the repository

```
dkp get chart [chartName] [chartVersion] [flags]
```

### 9.2.11.8.1 Options

-A, --all-namespaces	If present, list the requested object(s) across all namespaces.
--config string	Config file to use ( <b>default</b> <code>"/root/.kommander/config"</code> )
--context string	The name of the kubeconfig context to use
-h, --help	help <b>for</b> chart
--kubeconfig string	Path to the kubeconfig file to use <b>for</b> 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.

### 9.2.11.8.2 Options inherited from parent commands

-v, --verbose **int** Output verbosity

### 9.2.11.8.3 SEE ALSO

- [dkp get \(see page 978\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

## 9.2.11.9 dkp get clusters

Get clusters from specified Workspace

```
dkp get clusters [CLUSTER_NAME] [flags]
```

### 9.2.11.9.1 Options

-A, --all-namespaces	If present, list the requested object(s) across all namespaces.
--config string	Config file to use ( <b>default</b> <code>"/root/.kommander/config"</code> )
--context string	The name of the kubeconfig context to use
-h, --help	help <b>for</b> clusters
--kubeconfig string	Path to the kubeconfig file to use <b>for</b> 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.
-w, --workspace string Name of the workspace to show clusters from

```

### 9.2.11.9.2 Options inherited from parent commands

```

-v, --verbose int Output verbosity

```

### 9.2.11.9.3 SEE ALSO

- [dkp get \(see page 978\)](#) - Get one of [appdeployments, chart, clusters, kubeconfig, nodepools, workspaces]

## 9.2.12 dkp import

Import images from an image bundle into Containerd

### 9.2.12.1 Options

```

-h, --help help for import

```

### 9.2.12.2 Options inherited from parent commands

```

-v, --verbose int Output verbosity

```

### 9.2.12.3 SEE ALSO

- [dkp import image-bundle \(see page 983\)](#) - Import images from image bundles into Containerd

### 9.2.12.4 dkp import image-bundle

Import images from image bundles into Containerd

```

dkp import image-bundle [flags]

```

### 9.2.12.4.1 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 [])

```

### 9.2.12.4.2 Options inherited from parent commands

```

-v, --verbose int Output verbosity

```

### 9.2.12.4.3 SEE ALSO

- [dkp import](#) (see page 983) - Import images from an image bundle into Containerd

## 9.2.13 dkp install

Install one of [kommander]

### 9.2.13.1 Options

```

 --config string Config file to use (default "/root/.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.

```

### 9.2.13.2 Options inherited from parent commands

```

-v, --verbose int Output verbosity

```

### 9.2.13.3 SEE ALSO

- [dkp install kommander \(see page 985\)](#) - Install kommander

### 9.2.13.4 dkp install kommander

Install kommander

```
dkp install kommander [flags]
```

#### 9.2.13.4.1 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 "v2.4.2-rc.2")
-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)

```

#### 9.2.13.4.2 Options inherited from parent commands

```

--config string Config file to use (default "/root/.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

```

### 9.2.13.4.3 SEE ALSO

- [dkp install](#) (see page 984) - Install one of [kommander]

## 9.2.14 dkp move

Move one of [capi-resources]

### 9.2.14.1 Synopsis

Command "move" is deprecated, use "dkp move capi-resources" instead Move one of [capi-resources]

```
dkp move [flags]
```

### 9.2.14.2 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)

```

### 9.2.14.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

## 9.2.14.4 SEE ALSO

- [dkp move capi-resources](#) (see page 987) - Move controllers and objects from one cluster to the other

## 9.2.14.5 dkp move capi-resources

Move controllers and objects from one cluster to the other

```
dkp move capi-resources [flags]
```

### 9.2.14.5.1 Options

```

--from-context string Context to be used within the from-cluster's
kubecfg file. If empty, current context will be used.
--from-kubecfg file Path to the kubecfg 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 kubecfg
file. If empty, current context will be used.
--to-kubecfg file Path to the kubecfg for pivot's destination cluster

```

### 9.2.14.5.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.14.5.3 SEE ALSO

- [dkp move](#) (see page 986) - Move one of [capi-resources]

## 9.2.15 dkp open

Open one of [dashboard]

### 9.2.15.1 Options

```

--config string Config file to use (default "/root/.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.

```

### 9.2.15.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.15.3 SEE ALSO

- [dkp open dashboard](#) (see page 988) - Open DKP UI in your browser

### 9.2.15.4 dkp open dashboard

Open DKP UI in your browser

```
dkp open dashboard [flags]
```

#### 9.2.15.4.1 Options

```
-h, --help help for dashboard
```

#### 9.2.15.4.2 Options inherited from parent commands

```

--config string Config file to use (default "/root/.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

```

### 9.2.15.4.3 SEE ALSO

- [dkp open](#) (see page 987) - Open one of [dashboard]

## 9.2.16 dkp push

Push one of [chart, chart-bundle, image-bundle]

### 9.2.16.1 Options

```

-h, --help help for push

```

### 9.2.16.2 Options inherited from parent commands

```

-v, --verbose int Output verbosity

```

### 9.2.16.3 SEE ALSO

- [dkp push chart](#) (see page 991) - Upload charts to the repository
- [dkp push chart-bundle](#) (see page 989) - Upload chart bundles to the repository
- [dkp push image-bundle](#) (see page 990) - Push images from an image bundle into an existing OCI registry

### 9.2.16.4 dkp push chart-bundle

Upload chart bundles to the repository

```

dkp push chart-bundle [chartsTarball]... [flags]

```

### 9.2.16.4.1 Options

```

 --config string Config file to use (default "/root/.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.

```

### 9.2.16.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.16.4.3 SEE ALSO

- [dkp push](#) (see page 989) - Push one of [chart, chart-bundle, image-bundle]

## 9.2.16.5 dkp push image-bundle

Push images from an image bundle into an existing OCI registry

```
dkp push image-bundle [flags]
```

### 9.2.16.5.1 Options

```

 --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 image-bundle
 --image-bundle strings Tarball containing list of images to
push. Can also be a glob pattern. (default [])
 --to-registry string Registry to push images to
 --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)

```

<code>--to-registry-password string</code>	Password to use to log in to destination registry
<code>--to-registry-username string</code>	Username to use to log in to destination registry

### 9.2.16.5.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

### 9.2.16.5.3 SEE ALSO

- [dkp push](#) (see page 989) - Push one of [chart, chart-bundle, image-bundle]

## 9.2.16.6 dkp push chart

Upload charts to the repository

```
dkp push chart [chartTarball]... [flags]
```

### 9.2.16.6.1 Options

<code>--config string</code>	Config file to use ( <b>default</b> <code>"/root/.kommander/config"</code> )
<code>--context string</code>	The name of the kubeconfig context to use
<code>-h, --help</code>	help <b>for</b> chart
<code>--kubeconfig string</code>	Path to the kubeconfig file to use <b>for</b> 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.

### 9.2.16.6.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

### 9.2.16.6.3 SEE ALSO

- [dkp push](#) (see page 989) - Push one of [chart, chart-bundle, image-bundle]

## 9.2.17 dkp scale

Scale one of [nodepool]

### 9.2.17.1 Options

```
-h, --help help for scale
```

### 9.2.17.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.17.3 SEE ALSO

- [dkp scale nodepool](#) (see page 992) - Scale a nodepool of a given cluster to the number of replicas

### 9.2.17.4 dkp scale nodepool

Scale a nodepool of a given cluster to the number of replicas

```
dkp scale nodepool name [flags]
```

#### 9.2.17.4.1 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.
```

#### 9.2.17.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.17.4.3 SEE ALSO

- [dkp scale](#) (see page 992) - Scale one of [nodepool]

### 9.2.18 dkp serve

Serve image or Helm chart bundles from an OCI registry

#### 9.2.18.1 Options

```
-h, --help help for serve
```

#### 9.2.18.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.18.3 SEE ALSO

- [dkp serve image-bundle](#) (see page 993) - Serve an OCI registry from image bundles

#### 9.2.18.4 dkp serve image-bundle

Serve an OCI registry from image bundles

```
dkp serve image-bundle [flags]
```

##### 9.2.18.4.1 Options

```
-h, --help help for image-bundle
```



<code>--image-bundle strings</code>	Tarball of images to serve. Can also be a glob pattern. ( <b>default</b> [])
<code>--listen-address string</code>	Address to listen on ( <b>default</b> "localhost")
<code>--listen-port uint16</code>	Port to listen on (0 means use any free port)
<code>--tls-cert-file string</code>	TLS certificate file
<code>--tls-<b>private</b>-key-file string</code>	TLS <b>private</b> key file

#### 9.2.18.4.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

#### 9.2.18.4.3 SEE ALSO

- [dkp serve](#) (see page 993) - Serve image or Helm chart bundles from an OCI registry

### 9.2.19 dkp update

Update one of [bootstrap (cluster), controlplane, nodepool]

#### 9.2.19.1 Options

`-h, --help` help **for** update

#### 9.2.19.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

#### 9.2.19.3 SEE ALSO

- [dkp update bootstrap](#) (see page 1005) - Update bootstrap cluster
- [dkp update controlplane](#) (see page 1000) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]
- [dkp update nodepool](#) (see page 994) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

#### 9.2.19.4 dkp update nodepool

Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.19.4.1 Options

```
-h, --help help for nodepool
```

### 9.2.19.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.19.4.3 SEE ALSO

- [dkp update](#) (see page 994) - Update one of [bootstrap (cluster), controlplane, nodepool]
- [dkp update nodepool aws](#) (see page 996) - Update a Konvoy cluster node pool in AWS
- [dkp update nodepool azure](#) (see page 995) - Update a Konvoy cluster node pool in Azure
- [dkp update nodepool eks](#) (see page 997) - Update a Konvoy cluster node pool in EKS
- [dkp update nodepool gcp](#) (see page 999) - Update a Konvoy cluster node pool in GCP
- [dkp update nodepool preprovisioned](#) (see page 997) - Update a Konvoy cluster node pool in Preprovisioned
- [dkp update nodepool vsphere](#) (see page 998) - Update a Konvoy cluster node pool in vSphere

### 9.2.19.4.4 dkp update nodepool azure

Update a Konvoy cluster node pool in Azure

```
dkp update nodepool azure [flags]
```

#### 9.2.19.4.4.1 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')
```

<pre>-n, --namespace string request. (default "default")   --use-context string   --wait returning. (default true)</pre>	<pre>If present, the namespace scope for this CLI Use a specific context in a kubeconfig file. If true, wait for operations to complete before</pre>
--------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------

#### 9.2.19.4.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.19.4.4.3 SEE ALSO

- [dkp update nodepool](#) (see page 994) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

#### 9.2.19.4.5 dkp update nodepool aws

Update a Konvoy cluster node pool in AWS

```
dkp update nodepool aws [flags]
```

##### 9.2.19.4.5.1 Options

<pre>--ami string -c, --cluster-name name created resources. -h, --help   --instance-type string   --kubeconfig string cluster. If unspecified, default discovery rules apply.   --kubernetes-version string -n, --namespace string request. (default "default")   --use-context string   --wait returning. (default true)</pre>	<pre>AMI id to use for node pool machines Name used to prefix the cluster and all the help for aws Instance type to use for node pool machines Path to the kubeconfig for the management Kubernetes version If present, the namespace scope for this CLI Use a specific context in a kubeconfig file. If true, wait for operations to complete before</pre>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

##### 9.2.19.4.5.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.19.4.5.3 SEE ALSO

- [dkp update nodepool](#) (see page 994) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.19.4.6 dkp update nodepool preprovisioned

Update a Konvoy cluster node pool in Preprovisioned

```
dkp update nodepool preprovisioned [flags]
```

#### 9.2.19.4.6.1 Options

-c, --cluster-name name	Name used to prefix the cluster and all the created resources.
-h, --help	help <b>for</b> preprovisioned
--kubeconfig string	Path to the kubeconfig <b>for</b> the management cluster. If unspecified, <b>default</b> discovery rules apply.
--kubernetes-version string	Kubernetes version
-n, --namespace string	If present, the namespace scope <b>for this</b> CLI request. ( <b>default</b> "default")
--use-context string	Use a specific context in a kubeconfig file.
--wait	If <b>true</b> , wait <b>for</b> operations to complete before returning. ( <b>default</b> true)

#### 9.2.19.4.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.19.4.6.3 SEE ALSO

- [dkp update nodepool](#) (see page 994) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.19.4.7 dkp update nodepool eks

Update a Konvoy cluster node pool in EKS

```
dkp update nodepool eks [flags]
```

### 9.2.19.4.7.1 Options

<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>-h, --help</code>	help <b>for</b> eks
<code>--instance-type string</code>	Instance type to use <b>for</b> node pool machines
<code>--kubeconfig string</code>	Path to the kubeconfig <b>for</b> the management cluster. If unspecified, <b>default</b> discovery rules apply.
<code>--kubernetes-version string</code>	Kubernetes version
<code>-n, --namespace string</code>	If present, the namespace scope <b>for this</b> CLI request. ( <b>default</b> "default")
<code>--use-context string</code>	Use a specific context in a kubeconfig file.
<code>--wait</code>	If <b>true</b> , wait <b>for</b> operations to complete before returning. ( <b>default true</b> )

### 9.2.19.4.7.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

### 9.2.19.4.7.3 SEE ALSO

- [dkp update nodepool](#) (see page 994) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

## 9.2.19.4.8 dkp update nodepool vsphere

Update a Konvoy cluster node pool in vSphere

```
dkp update nodepool vsphere [flags]
```

### 9.2.19.4.8.1 Options

<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>--cpus <b>int</b></code>	The number of virtual processors in a virtual machine.
<code>--disk-size <b>int</b></code>	The size of a virtual machine's disk, in GB.
<code>-h, --help</code>	help <b>for</b> vsphere
<code>--kubeconfig string</code>	Path to the kubeconfig <b>for</b> the management cluster. If unspecified, <b>default</b> discovery rules apply.

<code>--kubernetes-version</code> string	Kubernetes version
<code>--memory</code> <b>int</b>	The size of a virtual machine's memory, in GB.
<code>-n, --namespace</code> string	If present, the namespace scope <b>for this</b> CLI
request. ( <b>default</b> "default")	
<code>--use-context</code> string	Use a specific context in a kubeconfig file.
<code>--vm-template</code> string	The virtual machine template to use <b>for</b> a virtual
machine.	
<code>--wait</code>	If <b>true</b> , wait <b>for</b> operations to complete before
returning. ( <b>default</b> <b>true</b> )	

### 9.2.19.4.8.2 Options inherited from parent commands

`-v, --verbose` **int** Output verbosity

### 9.2.19.4.8.3 SEE ALSO

- [dkp update nodepool](#) (see page 994) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.19.4.9 dkp update nodepool gcp

Update a Konvoy cluster node pool in GCP

```
dkp update nodepool gcp [flags]
```

#### 9.2.19.4.9.1 Options

<code>-c, --cluster-name</code> name	Name used to prefix the cluster and all the
created resources.	
<code>-h, --help</code>	help <b>for</b> gcp
<code>--image</code> string	Full reference to an image to use <b>for</b> all nodes
(set either <b>this</b> 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 <b>for</b> all
nodes (set either <b>this</b> or <code>--image</code> ) (ex. 'projects/my-project/global/images/family/konvoy-ubuntu-2004-{{.K8sVersion}}')	
<code>--instance-type</code> string	Worker machine instance type (ex. "n2-standard-8")
<code>--kubeconfig</code> string	Path to the kubeconfig <b>for</b> the management
cluster. If unspecified, <b>default</b> discovery rules apply.	
<code>--kubernetes-version</code> string	Kubernetes version
<code>-n, --namespace</code> string	If present, the namespace scope <b>for this</b> CLI
request. ( <b>default</b> "default")	
<code>--use-context</code> string	Use a specific context in a kubeconfig file.

```
--wait If true, wait for operations to complete before
returning. (default true)
```

#### 9.2.19.4.9.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.19.4.9.3 SEE ALSO

- [dkp update nodepool](#) (see page 994) - Update a Kubernetes cluster node pool, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.19.5 dkp update controlplane

Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

#### 9.2.19.5.1 Options

```
-h, --help help for controlplane
```

#### 9.2.19.5.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.19.5.3 SEE ALSO

- [dkp update](#) (see page 994) - Update one of [bootstrap (cluster), controlplane, nodepool]
- [dkp update controlplane aws](#) (see page 1001) - Update a Konvoy cluster control plane in AWS
- [dkp update controlplane azure](#) (see page 1002) - Update a Konvoy cluster control plane in Azure
- [dkp update controlplane eks](#) (see page 1003) - Update a Konvoy cluster control plane in EKS
- [dkp update controlplane gcp](#) (see page 1004) - Update a Konvoy cluster control plane in GCP
- [dkp update controlplane preprovisioned](#) (see page 1000) - Update a Konvoy cluster control plane in Preprovisioned
- [dkp update controlplane vsphere](#) (see page 1004) - Update a Konvoy cluster control plane in vSphere

#### 9.2.19.5.4 dkp update controlplane preprovisioned

Update a Konvoy cluster control plane in Preprovisioned

```
dkp update controlplane preprovisioned [flags]
```

#### 9.2.19.5.4.1 Options

<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>-h, --help</code>	help <b>for</b> preprovisioned
<code>--kubeconfig string</code>	Path to the kubeconfig <b>for</b> the management cluster. If unspecified, <b>default</b> discovery rules apply.
<code>--kubernetes-version string</code>	Kubernetes version
<code>-n, --namespace string</code>	If present, the namespace scope <b>for this</b> CLI request. ( <b>default</b> "default")
<code>--use-context string</code>	Use a specific context in a kubeconfig file.
<code>--wait</code>	If <b>true</b> , wait <b>for</b> operations to complete before returning. ( <b>default true</b> )

#### 9.2.19.5.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.19.5.4.3 SEE ALSO

- [dkp update controlplane](#) (see page 1000) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

#### 9.2.19.5.5 dkp update controlplane aws

Update a Konvoy cluster control plane in AWS

```
dkp update controlplane aws [flags]
```

#### 9.2.19.5.5.1 Options

<code>--ami string</code>	AMI id to use <b>for</b> control plane machines
<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>-h, --help</code>	help <b>for</b> aws
<code>--instance-type string</code>	Instance type to use <b>for</b> 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)

```

### 9.2.19.5.5.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.19.5.5.3 SEE ALSO

- [dkp update controlplane](#) (see page 1000) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.19.5.6 dkp update controlplane azure

Update a Konvoy cluster control plane in Azure

```
dkp update controlplane azure [flags]
```

#### 9.2.19.5.6.1 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")
--use-context string Use a specific context in a kubeconfig file.
--wait If true, wait for operations to complete before
returning. (default true)

```

### 9.2.19.5.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.19.5.6.3 SEE ALSO

- [dkp update controlplane](#) (see page 1000) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

## 9.2.19.5.7 dkp update controlplane eks

Update a Konvoy cluster control plane in EKS

```
dkp update controlplane eks [flags]
```

### 9.2.19.5.7.1 Options

-c, --cluster-name name	Name used to prefix the cluster and all the created resources.
-h, --help	help <b>for</b> eks
--kubeconfig string	Path to the kubeconfig <b>for</b> the management cluster. If unspecified, <b>default</b> discovery rules apply.
--kubernetes-version string	Kubernetes version
-n, --namespace string	If present, the namespace scope <b>for this</b> CLI request. ( <b>default</b> "default")
--use-context string	Use a specific context in a kubeconfig file.
--wait	If <b>true</b> , wait <b>for</b> operations to complete before returning. ( <b>default</b> <b>true</b> )

### 9.2.19.5.7.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.19.5.7.3 SEE ALSO

- [dkp update controlplane](#) (see page 1000) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.19.5.8 dkp update controlplane vsphere

Update a Konvoy cluster control plane in vSphere

```
dkp update controlplane vsphere [flags]
```

#### 9.2.19.5.8.1 Options

-c, --cluster-name name	Name used to prefix the cluster and all the created resources.
--cpus <b>int</b>	The number of virtual processors in a virtual machine.
--disk-size <b>int</b>	The size of a virtual machine's disk, in GB.
-h, --help	help <b>for</b> vsphere
--kubeconfig string	Path to the kubeconfig <b>for</b> the management cluster. If unspecified, <b>default</b> discovery rules apply.
--kubernetes-version string	Kubernetes version
--memory <b>int</b>	The size of a virtual machine's memory, in GB.
-n, --namespace string	If present, the namespace scope <b>for this</b> CLI request. ( <b>default</b> "default")
--use-context string	Use a specific context in a kubeconfig file.
--vm-template string	The virtual machine template to use <b>for</b> a virtual machine.
--wait	If <b>true</b> , wait <b>for</b> operations to complete before returning. ( <b>default true</b> )

#### 9.2.19.5.8.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.19.5.8.3 SEE ALSO

- [dkp update controlplane](#) (see page 1000) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.19.5.9 dkp update controlplane gcp

Update a Konvoy cluster control plane in GCP

```
dkp update controlplane gcp [flags]
```

### 9.2.19.5.9.1 Options

<code>-c, --cluster-name name</code>	Name used to prefix the cluster and all the created resources.
<code>-h, --help</code>	help <b>for</b> gcp
<code>--image string</code>	Full reference to an image to use <b>for</b> all nodes (set either <b>this</b> or <code>--image-family</code> ) (ex. 'projects/my-project/global/images/konvoy-ubuntu-2004-1-99-99-1234567890')
<code>--image-family string</code>	Full reference to an image family to use <b>for</b> all nodes (set either <b>this</b> or <code>--image</code> ) (ex. 'projects/my-project/global/images/family/konvoy-ubuntu-2004-{{.K8sVersion}}')
<code>--instance-type string</code>	Control Plane machine instance type (ex. "n2-standard-4")
<code>--kubeconfig string</code>	Path to the kubeconfig <b>for</b> the management cluster. If unspecified, <b>default</b> discovery rules apply.
<code>--kubernetes-version string</code>	Kubernetes version
<code>-n, --namespace string</code>	If present, the namespace scope <b>for this</b> CLI request. ( <b>default</b> "default")
<code>--use-context string</code>	Use a specific context in a kubeconfig file.
<code>--wait</code>	If <b>true</b> , wait <b>for</b> operations to complete before returning. ( <b>default true</b> )

### 9.2.19.5.9.2 Options inherited from parent commands

`-v, --verbose int` Output verbosity

### 9.2.19.5.9.3 SEE ALSO

- [dkp update controlplane](#) (see page 1000) - Update a Kubernetes cluster control plane, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

## 9.2.19.6 dkp update bootstrap

Update bootstrap cluster

### 9.2.19.6.1 Options

`-h, --help` help **for** bootstrap

### 9.2.19.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.19.6.3 SEE ALSO

- [dkp update](#) (see page 994) - Update one of [bootstrap (cluster), controlplane, nodepool]
- [dkp update bootstrap credentials](#) (see page 1006) - Update credentials in the cluster

### 9.2.19.6.4 dkp update bootstrap credentials

Update credentials in the cluster

#### 9.2.19.6.4.1 Options

```
-h, --help help for credentials
```

#### 9.2.19.6.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.19.6.4.3 SEE ALSO

- [dkp update bootstrap](#) (see page 1005) - Update bootstrap cluster
- [dkp update bootstrap credentials aws](#) (see page 1006) - Update AWS credentials in the cluster and restart CAPA controllers
- [dkp update bootstrap credentials azure](#) (see page 1007) - Update Azure credentials in the cluster and restart CAPZ controllers
- [dkp update bootstrap credentials gcp](#) (see page 1008) - Update GCP credentials in the cluster and restart CAPG controllers
- [dkp update bootstrap credentials vsphere](#) (see page 1008) - Update VSphere credentials in the cluster and restart CAPV controllers

#### 9.2.19.6.4.4 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 1006) - Update credentials in the cluster

**9.2.19.6.4.5 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 1006) - Update credentials in the cluster

9.2.19.6.4.6 `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 1006) - Update credentials in the cluster

9.2.19.6.4.7 `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 \(see page 1006\)](#) - Update credentials in the cluster

## 9.2.20 dkp upgrade

Upgrade one of [addons, capi-components, catalogapp, kommander, workspace]

### 9.2.20.1 Options

```
-h, --help help for upgrade
```

### 9.2.20.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.20.3 SEE ALSO

- [dkp upgrade addons \(see page 1013\)](#) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]
- [dkp upgrade capi-components \(see page 1011\)](#) - Upgrade the CAPI components in the cluster
- [dkp upgrade catalogapp \(see page 1012\)](#) - Upgrade a Catalog Application to a newer version
- [dkp upgrade kommander \(see page 1009\)](#) - Upgrade the Kommander version of the targeted cluster
- [dkp upgrade workspace \(see page 1013\)](#) - Upgrade all platform applications in the given workspace and its projects to the same version as platform applications running on the management cluster

### 9.2.20.4 dkp upgrade kommander

Upgrade the Kommander version of the targeted cluster



### 9.2.20.4.1 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]
```

### 9.2.20.4.2 Options

<code>--charts-bundle</code> stringArray	Path to charts-bundle to upload to chartmuseum, apart from parsing the kommander applications repository ( <b>default</b> [])
<code>--config</code> string	Config file to use ( <b>default</b> " <code>root/.kommander/config</code> ")
<code>--context</code> string	The name of the kubeconfig context to use
<code>--core-app-timeout</code> duration	Timeout to wait <b>for</b> upgrade of each kommander core application ( <b>default</b> <code>20m0s</code> )
<code>--disable-appdeployments</code> strings	List of AppDeployments to be disabled during upgrade ( <b>default</b> [])
<code>--disallow-charts-download</code>	make CLI rely solely on provided chart bundles and <b>do not try</b> to download charts from the Internet
<code>--gitea-kommander-repository-name</code> string	gitea kommander repository name ( <b>default</b> " <code>kommander</code> ")
<code>-h, --help</code>	help <b>for</b> kommander
<code>--kommander-applications-repository</code> string	git repository with application definitions ( <b>default</b> " <code>v2.4.2-rc.2</code> ")
<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 <b>for</b> CLI requests.
<code>--platform-apps-timeout</code> duration	Timeout to wait <b>for</b> upgrade of the set of platform applications ( <b>default</b> <code>30m0s</code> )
<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. <code>1s</code> , <code>2m</code> , <code>3h</code> ). A value of zero means don't timeout requests.

### 9.2.20.4.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.20.4.4 SEE ALSO

- [dkp upgrade](#) (see page 1009) - Upgrade one of [addons, capi-components, catalogapp, kommander, workspace]

#### 9.2.20.5 dkp upgrade capi-components

Upgrade the CAPI components in the cluster

```
dkp upgrade capi-components [flags]
```

##### 9.2.20.5.1 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. 1s, 2m, 3h). (default 10m0s)
--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.

```

##### 9.2.20.5.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.20.5.3 SEE ALSO

- [dkp upgrade](#) (see page 1009) - Upgrade one of [addons, capi-components, catalogapp, kommander, workspace]

### 9.2.20.6 dkp upgrade catalogapp

Upgrade a Catalog Application to a newer version

```
dkp upgrade catalogapp CATALOGAPP_NAME --to-version VERSION [--workspace WORKSPACE |
--project PROJECT] [flags]
```

#### 9.2.20.6.1 Options

<code>--config</code> string	Config file to use ( <b>default</b> <code>"/root/.kommander/config"</code> )
<code>--context</code> string	The name of the kubeconfig context to use
<code>--core-app-timeout</code> duration	Timeout to wait <b>for</b> upgrade of each kommander core application ( <b>default</b> 20m0s)
<code>--disable-appdeployments</code> strings	List of AppDeployments to be disabled during upgrade ( <b>default</b> [])
<code>-h, --help</code>	help <b>for</b> catalogapp
<code>--kubeconfig</code> string	Path to the kubeconfig file to use <b>for</b> CLI requests.
<code>--platform-apps-timeout</code> duration	Timeout to wait <b>for</b> upgrade of the set of platform applications ( <b>default</b> 30m0s)
<code>--project</code> string	Name of the Project to upgrade the Catalog App in
<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>--to-version</code> string	Version the Catalog App should be upgraded to
<code>-w, --workspace</code> string	Name of the Workspace to upgrade the Catalog App in

#### 9.2.20.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.20.6.3 SEE ALSO

- [dkp upgrade](#) (see page 1009) - Upgrade one of [addons, capi-components, catalogapp, kommander, workspace]

### 9.2.20.7 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]
```

#### 9.2.20.7.1 Options

```

--config string Config file to use (default "/
root/.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.
```

#### 9.2.20.7.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.20.7.3 SEE ALSO

- [dkp upgrade](#) (see page 1009) - Upgrade one of [addons, capi-components, catalogapp, kommander, workspace]

## 9.2.20.8 dkp upgrade addons

Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.20.8.1 Options

```
-h, --help help for addons
```

### 9.2.20.8.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.20.8.3 SEE ALSO

- [dkp upgrade](#) (see page 1009) - Upgrade one of [addons, capi-components, catalogapp, kommander, workspace]
- [dkp upgrade addons aws](#) (see page 1013) - Upgrade the core Addons in a AWS cluster
- [dkp upgrade addons azure](#) (see page 1015) - Upgrade the core Addons in a Azure cluster
- [dkp upgrade addons eks](#) (see page 1017) - Upgrade the core Addons in a EKS cluster
- [dkp upgrade addons gcp](#) (see page 1019) - Upgrade the core Addons in a GCP cluster
- [dkp upgrade addons preprovisioned](#) (see page 1016) - Upgrade the core Addons in a Preprovisioned cluster
- [dkp upgrade addons vsphere](#) (see page 1018) - Upgrade the core Addons in a vSphere cluster

### 9.2.20.8.4 dkp upgrade addons aws

Upgrade the core Addons in a AWS cluster

```
dkp upgrade addons aws [flags]
```

#### 9.2.20.8.4.1 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 goyaml 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).
--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 go lang templates
[http://golang.org/pkg/text/template/#pkg-overview].

```

#### 9.2.20.8.4.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

#### 9.2.20.8.4.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1013) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

#### 9.2.20.8.5 dkp upgrade addons azure

Upgrade the core Addons in a Azure cluster

```
dkp upgrade addons azure [flags]
```

##### 9.2.20.8.5.1 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 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 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).
 --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 go lang templates
 [http://golang.org/pkg/text/template/#pkg-overview].

```

### 9.2.20.8.5.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.20.8.5.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1013) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.20.8.6 dkp upgrade addons preprovisioned

Upgrade the core Addons in a Preprovisioned cluster

```
dkp upgrade addons preprovisioned [flags]
```

#### 9.2.20.8.6.1 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 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 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).

```

```

--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 go lang templates
[http://golang.org/pkg/text/template/#pkg-overview].

```

### 9.2.20.8.6.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.20.8.6.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1013) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.20.8.7 dkp upgrade addons eks

Upgrade the core Addons in a EKS cluster

```
dkp upgrade addons eks [flags]
```

#### 9.2.20.8.7.1 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 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 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).
--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 go lang templates
[http://golang.org/pkg/text/template/#pkg-overview].

```



### 9.2.20.8.7.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.20.8.7.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1013) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.20.8.8 dkp upgrade addons vsphere

Upgrade the core Addons in a vSphere cluster

```
dkp upgrade addons vsphere [flags]
```

#### 9.2.20.8.8.1 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 goyaml 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).
--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 goyaml templates
[http://golang.org/pkg/text/template/#pkg-overview].

```

### 9.2.20.8.8.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.20.8.8.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1013) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

### 9.2.20.8.9 dkp upgrade addons gcp

Upgrade the core Addons in a GCP cluster

```
dkp upgrade addons gcp [flags]
```

#### 9.2.20.8.9.1 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 goyaml 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).
--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 goyaml templates
[http://golang.org/pkg/text/template/#pkg-overview].

```

### 9.2.20.8.9.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.20.8.9.3 SEE ALSO

- [dkp upgrade addons](#) (see page 1013) - Upgrade the core Addons in a cluster, one of [aws, azure, eks, gcp, preprovisioned, vsphere]

## 9.2.21 dkp edit

Edit a resource on the server

### 9.2.21.1 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)
```

### 9.2.21.2 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 goyaml and jsonpath
output formats. (default true)
--config string Config file to use (default "/root/.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")
 --windows-line-endings Defaults to the line ending native to your
platform.

```

### 9.2.21.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.22 dkp version

Print version information

### 9.2.22.1 Synopsis

Print version information

```
dkp version [flags]
```

### 9.2.22.2 Options

```
-h, --help help for version
--long If true, print additional version information.
-o, --output string One of 'yaml' or 'json'.
```

### 9.2.22.3 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

## 9.2.23 dkp cluster

Get cluster information

### 9.2.23.1 Options

```
--config string Config file to use (default "/root/.kommander/
config")
--context string 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.
```

### 9.2.23.2 Options inherited from parent commands

```
-v, --verbose int Output verbosity
```

### 9.2.23.3 SEE ALSO

- [dkp cluster type](#) (see page 1023) - Retrieve cluster type

### 9.2.23.4 dkp cluster type

Retrieve cluster type

```
dkp cluster type [flags]
```

#### 9.2.23.4.1 Options

```
-h, --help help for type
```

#### 9.2.23.4.2 Options inherited from parent commands

<code>--config</code> string	Config file to use ( <b>default</b> <code>"/root/.kommander/config"</code> )
<code>--context</code> string	The name of the kubeconfig context to use
<code>--kubeconfig</code> string	Path to the kubeconfig file to use <b>for</b> CLI requests.
<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> <b>int</b>	Output verbosity

#### 9.2.23.4.3 SEE ALSO

- [dkp cluster](#) (see page 1022) - Get cluster information

# 10 Upgrade DKP

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 1021).

## 10.2 Supported Upgrade Paths

Use this table to determine your correct upgrade path:

		Upgrading from Release...																	
		2.1.0	2.1.1	2.1.2	2.1.3	2.1.4	2.2.0	2.2.1	2.2.2	2.2.3	2.3.0	2.3.1	2.3.2	2.3.3	2.4.0	2.4.1	2.4.2	2.5.0	
... to R e l e a s e	2.2.0	Yes	Yes	No	No	No	No	No	No	No	No	No	No	No	No	No	No	No	No
	2.2.1	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No	No	No	No
	2.2.2	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No	No	No
	2.2.3	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No	No
	2.3.0	No	No	No	No	No	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No
	2.3.1	No	No	No	No	No	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No	No

2.3.1	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No	No
2.3.2	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No	No
2.3.3	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No	No	No	No
2.4.0	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes	No	No	No	No	No
2.4.1	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes	No	No	No	No
2.4.2	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
2.5.0	No	No	No	No	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	No	No
2.5.1	No	No	No	No	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes
2.5.2	No	No	No	No	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes	Yes	Yes

### 10.2.1 Understand the Upgrade Process

For this release, you perform the upgrade sequentially, beginning with the DKP UI and then moving to upgrading clusters and CAPI components.

When upgrading DKP, the process is different depending on whether you run a stand-alone Management Cluster, or a multi-cluster environment that includes a combination of a Management cluster and managed or attached workspace clusters.



Start with your Management Cluster in the UI, and then, if more than one exists, proceed workspace by workspace until complete. You can then move to upgrading Konvoy, cluster by cluster.

The overall process for upgrading to the latest version of DKP is done on each Workspace or cluster, with the following processes:



- Before upgrading, we **strongly recommend** reading the [release notes \(see page 1055\)](#) and verifying your current setup against any possible breaking changes.
- For this version of DKP, it is extremely important that you review the [list of major Kubernetes changes \(see page 1077\)](#) that may affect your system.

For **Kommander**, on your Management Cluster:

1. **Upgrade DKP UI** ([see page 1026](#)) and all Platform Applications.

If you do not have any managed or attached clusters, skip to upgrading Konvoy on your Management Cluster.

On your Workspaces (which include Management Cluster and managed or attached clusters):

1. **Upgrade your Workspaces** ([see page 1013](#)), which upgrades all Platform Applications on your managed or attached workspace clusters.
2. **Upgrade all DKP Catalog applications** ([see page 604](#)) deployed to Workspaces.
3. **Upgrade all DKP Catalog applications** ([see page 623](#)) deployed to Projects.
4. **Verify any Custom Catalog applications** ([see page 1053](#)) and ensure they are compatible with the Kubernetes version included in the [new release \(see page 1055\)](#).

For **Konvoy**, on your Management Cluster:

1. **Upgrade CAPI components** ([see page 1033](#)) (Essential or Enterprise section). This upgrades the CAPI controllers, which only run on the Management Cluster.
2. **Upgrade the Core Addons** ([see page 1033](#)) (Essential or Enterprise section). This upgrades multiple addons such as CSI, CNI, Cluster Autoscaler, and Node Feature Discovery.
3. **Upgrade the Kubernetes version** ([see page 1033](#)) (Essential or Enterprise section). This upgrades your cluster's control plane and node pools.

If you do not have any managed or attached clusters, you have finished the upgrade process and can start testing your environment. If you have managed or attached clusters, continue with the next section.

For **Konvoy**, on your Managed Clusters:

1. **Upgrade the Core Addons** ([see page 0](#)). This upgrades multiple addons such as CSI, CNI, Cluster Autoscaler, and Node Feature Discovery.
2. **Upgrade the Kubernetes version** ([see page 1037](#)). This upgrades your cluster's control plane and node pools. We recommend you upgrade your Kubernetes version on any attached clusters.

## 10.3 Upgrade Kommander

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.

### Page Contents

#### 10.3.1 Prerequisites

- **REQUIRED** Before upgrading, create an [on-demand backup](#) (see page 730) of your current configuration with Velero.
- [Download](#) (see page 100) and install the supported DKP CLI binary of [this release](#) (see page 1055) on your computer.
- Ensure you are attempting to run a [supported DKP upgrade](#) (see page 0).
- Ensure you are on Kubernetes version 1.23 or higher.
- If you have attached clusters, ensure they are on Kubernetes versions 1.22 or higher.
- Review the [DKP 2.4.1 Components and Applications](#) (see page 1081) versions that are part of this upgrade.
- For air-gapped environments **with** DKP Catalog Applications in a multi-cluster environment: [Load the Docker images into your Docker registry](#)<sup>777</sup>
- For air-gapped environments **without** DKP Catalog Applications: [Load the Docker images into your Docker registry](#)<sup>778</sup>



#### DKP Upgrade with Insights Installed


If your environment has Insights installed, you must uninstall Insights BEFORE upgrading DKP.

<sup>777</sup> <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29924336/Install+DKP+in+an+Air-gapped+Environment+with+Catalog+Applications#Load-the-Docker-images-into-your-Docker-registry>

<sup>778</sup> <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/89131679/%282.3%29+Install+DKP+in+an+Air-gapped+Environment+with+Catalog+Applications>

### 10.3.1.1 Prerequisites for Configurations with NVIDIA GPU Nodes

If your cluster is running GPU nodes and has deployed the `nvidia` Platform Application, you must run additional steps **prior** to upgrading Kommander. In DKP 2.4, the `nvidia` application is replaced by `nvidia-gpu-operator` during upgrade, and the configuration for the new application must be created prior to the upgrade to ensure that the new application is deployed with the proper configuration, or it may not deploy successfully. If your cluster is not running GPU nodes, then skip these steps and proceed directly to the [Upgrade Kommander](#) (see page 1029) section.

 The GPU nodes will experience downtime during the upgrade due to the new version of the Nvidia application.

1. Create the ConfigMap with the necessary configuration overrides to [set the correct Toolkit version](#) (see page 1026). 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
```

2. Create an `AppDeployment` named `nvidia-gpu-operator` in the `kommander` namespace:

```
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


3. You may now proceed to the Kommander upgrade steps.

### 10.3.2 Upgrade Kommander

Before running the following command, ensure that your `dkp` configuration **references the Kommander Management cluster**, otherwise it attempts to run the upgrade on the bootstrap cluster. You can do this by setting the `KUBECONFIG` environment variable [to the appropriate kubeconfig file's location](#)<sup>779</sup>.

In this version of DKP, pre-provisioned environments require a special upgrade path. Refer to [Upgrade Kommander in a Pre-provisioned Environment](#) (see page 1031) for more information.

 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.

 As stated earlier, an alternative to initializing the `KUBECONFIG` environment variable is to use the `-kubeconfig=cluster_name.conf` flag. This ensures that Kommander upgrades on the workload cluster.

1. Use the DKP CLI to upgrade Kommander and all the Platform Applications in the Management Cluster:
  - a. For **air-gapped** environments:

```
dkp upgrade kommander --charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.4.2.tar.gz --kommander-applications-repository ./application-repositories/kommander-applications-v2.4.2.tar.gz
```

- b. For **air-gapped** environments with **DKP Catalog Applications** in a multi-cluster environment:

```
dkp upgrade kommander --charts-bundle ./application-charts/dkp-kommander-charts-bundle-v2.4.2.tar.gz --charts-bundle ./application-charts/dkp-catalog-applications-charts-bundle-v2.4.2.tar.gz --kommander-applications-repository ./application-repositories/kommander-applications-v2.4.2.tar.gz
```

After the upgrade, if you have DKP Catalog Applications deployed, update the `GitRepository` for `dkp-catalog-applications`.

<sup>779</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

⚠️⚠️⚠️ 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. ⚠️⚠️⚠️

- i. In the `kommander.yaml` you are currently using for your environment, update the DKP Catalog Applications by setting the correct DKP version:

```
apiVersion: config.kommander.mesosphere.io/v1alpha1
kind: Installation
...
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-v2.4.2.tar.gz # modify this
 version to match the DKP upgrade version
```

- ii. Refresh the `kommander.yaml` to apply the updated tarball:

⚠️ 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 --installer-config kommander.yaml
```

- c. For **non-air-gapped** environments:

```
dkp upgrade kommander
```

After the upgrade, if you have DKP Catalog Applications deployed, update the GitRepository for `dkp-catalog-applications`.

- i. 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":"v2.4.2"}}}'
```

**i** This command updates the catalog application repositories for all workspaces.

- ii. 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>
```

- iii. Update the GitRepository for the additional workspace:

```
kubectl patch gitrepository -n ${WORKSPACE_NAMESPACE} dkp-catalog-
applications --type merge --patch '{"spec":{"ref":
{"tag":"v2.4.2"}}}'
```

2. For **Enterprise customers** (see page 104) (multi-cluster environment): Upgrade your additional **Workspaces** (see page 0) on a per-Workspace basis to upgrade the Platform Applications on other clusters than the Management Cluster.  
For **Essential customers** (see page 107) (single-cluster environment): Proceed with the **Konvoy Upgrade** (see page 1033).

### 10.3.3 Troubleshooting

If the upgrade fails, run the following command to get more information on the upgrade process:

```
dkp upgrade kommander -v 6
```

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}]'
```

You can always go back to the **DKP Upgrade overview** (see page 0), to review the next steps depending on your environment and license type.

## 10.3.4 Upgrade Kommander in a Pre-provisioned Environment

### 10.3.4.1 Prerequisites

- Ensure you meet all prerequisites stated in the [Kommander upgrade \(see page 1026\)](#) section.
- Ensure you have set up a minimum of 4 x 40 GB of [raw \(unformatted disk\) storage for each of the worker nodes \(see page 862\)](#) of your cluster. We recommend setting up 1 volume per node, or multiple volumes per node only if node pool is less than 4.

#### 10.3.4.1.1 Modify Rook Ceph's Deployment

1. Clone the `kommander-applications` repo to your local machine where you are running the upgrade:

```
git clone --depth 1 --branch v2.4.2 https://github.com/mesosphere/kommander-applications.git
```

2. Prepare `rook-ceph` for the upgrade by editing the storage specifications of the `rook-ceph-cluster`:

```
cd kommander-applications/services/rook-ceph-cluster/1.10.8/defaults/
vi cm.yaml
```

3. Adapt the `cephClusterSpec.storage` in `cm.yaml` to include the following values:

```
storageClassDeviceSets: []
useAllDevices: true
useAllNodes: true
```

4. **Delete** the following value in the same file:

```
onlyApplyOSDPlacement
```

#### 10.3.4.1.2 Upgrade the Kommander component of DKP

1. Ensure you are targeting the correct DKP version:

```
dkp version
```

The output should show the current binary version, in this case, DKP 2.4.0.

2. Execute the upgrade by specifying the repository you cloned and updated in the previous section:

```
dkp upgrade kommander --kommander-applications-repository ./kommander-
applications --platform-apps-timeout
```

## 10.4 Upgrade Konvoy

### 10.4.1 Steps to upgrade Konvoy via CLI

This section describes how to upgrade your DKP clusters to the latest Konvoy version. To prevent compatibility issues, review and ensure you are following the correct upgrade process for customers with a single-cluster or multi-cluster experience.

- [DKP Enterprise Upgrade](#) (see page 1033)
- [DKP Essential Upgrade](#) (see page 1043)
- [Upgrade Cluster Node Pools](#) (see page 1052)

### 10.4.2 DKP Enterprise Upgrade

Enterprise

Upgrade your Konvoy environment within the DKP Enterprise license.

#### 10.4.2.1 Prerequisites

- Create an [on-demand backup](#) (see page 730) of your current configuration with Velero.
- Follow the steps listed in the [DKP upgrade overview](#) (see page 0).
- Check what version of DKP you have downloaded currently using cli command [dkp version](#) (see page 1021).
- Ensure that all platform applications in the management cluster have been upgraded to avoid compatibility issues with the [Kubernetes version](#) (see page 1056) included in this release. This is done automatically when [upgrading Kommander](#) (see page 1026), so ensure that you upgrade Kommander prior to upgrading Konvoy.
- Set the appropriate environment variables:
  - For Air-gapped, download the required bundles at our [support site](#)<sup>780</sup>.
  - For Azure, set the required [environment variables](#) (see page 226).
  - For [AWS and EKS](#) (see page 151), set the required [environment variables](#) (see page 47).
  - For vSphere, set the required [environment variables](#) (see page 327).

<sup>780</sup> <https://support.d2iq.com/hc/en-us>



- For GCP, set the required [environment variables](#) (see page 382).
- **vSphere only:** If you want to resize your disk, ensure you have reviewed [Create a Base OS Image](#) (see page 430).

### 10.4.2.2 Overview

To upgrade Konvoy for DKP Enterprise:

1. Upgrade the Cluster API (CAPI) components.
2. Upgrade the core addons.
3. Upgrade the Kubernetes version.
4. Upgrade the Managed clusters.
5. Upgrade the Kubernetes version of Managed clusters.

Perform all three steps on the management cluster first. Then, execute the second and third steps on additional managed clusters one cluster at a time. For the managed clusters, you use the KUBECONFIG for the management cluster and specify the name of the managed cluster(s) to upgrade.

For a full list of DKP Enterprise features, see [DKP Enterprise](#) (see page 104).




- For pre-provisioned air-gapped environments, you must run `konvoy-image upload artifacts` before beginning Upgrade the Capi Components section below.
- You must maintain your attached clusters manually. Review the documentation from your cloud provider for more information.
- See this section for a full explanation of [DKP Enterprise](#) (see page 104).


### 10.4.2.3 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](#)<sup>781</sup>.

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 \(see page 938\)](#) for flag descriptions for `--with-aws-bootstrap-credentials` and `--aws-service-endpoints`
  - Refer to the HTTP section for details: [Universal Configurations for all Infrastructures \(see page 128\)](#)

```
dkp upgrade capi-components
```

The output resembles the following:

- ✓ Upgrading CAPI components
- ✓ Waiting **for** CAPI components to be upgraded
- ✓ Initializing **new** CAPI components
- ✓ Deleting Outdated Global ClusterResourceSets


If the upgrade fails, review the prerequisites section and ensure that you've followed the steps in the [DKP Upgrade \(see page 1024\)](#) overview. Furthermore, ensure you have adhered to the Prerequisites at the top of this page.

#### 10.4.2.4 Upgrade the Core Addons

To install the core addons, DKP relies on the `ClusterResourceSet` [Cluster API feature](#)<sup>782</sup>. In the CAPI component upgrade, we deleted the previous set of outdated global `ClusterResourceSets` because prior to DKP 2.4 some addons were installed using a global configuration. In order to support individual cluster upgrades, DKP 2.4 now installs all addons with a unique set of `ClusterResourceSet` and corresponding referenced resources, all named using the cluster's name as a suffix. For example: `calico-cni-installation-my-aws-cluster`.

<sup>781</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

<sup>782</sup> <https://cluster-api.sigs.k8s.io/>

 If you have modified any of the `clusterResourceSet` definitions, these changes will **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 remake 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.

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](#)<sup>783</sup>.
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.

#### Additional Considerations for EKS

- The Kubernetes version for EKS clusters supported on DKP 2.4 is now v1.23. As EKS has disabled support for in-tree EBS volume provisioning in favor of [CSI Volumes](#)<sup>784</sup>, DKP automatically deploys a new EBS CSI driver to your EKS cluster when you run the `dkp upgrade addons` command.
- Before running the `dkp upgrade addons` command when deploying EKS clusters, you must add the necessary IAM policy to your worker instances. If you are using the default IAM instance profile name, run the following command:

```
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
```

If you have customized your `AWSMachineTemplate` to use a different instance profile, review and add the policy to that profile.

Examples for upgrade core addons commands:

```
export CLUSTER_NAME=my-azure-cluster
dkp upgrade addons azure --cluster-name=${CLUSTER_NAME}
```

OR

<sup>783</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

<sup>784</sup> <https://kubernetes.io/blog/2019/12/09/kubernetes-1-17-feature-csi-migration-beta/>

```
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
configmap/nvidia-feature-discovery-my-aws-cluster upgraded
```

4. After addons are upgraded, you must remove deprecated addons. In DKP version 2.4, the `nvidia-feature-discovery` addon will no longer be shipped on new clusters, but rather will be handled by the `nvidia-gpu-operator` Kommander component. To perform removal, execute the following steps:

a. List all cluster resource sets by running:

```
kubectl get clusterresourcesets
```

NAME	AGE
aws-ebs-csi-my-aws-cluster	7m46s
calico-cni-installation-my-aws-cluster	7m46s
cluster-autoscaler-my-aws-cluster	7m46s
node-feature-discovery-my-aws-cluster	7m46
nvidia-feature-discovery-my-aws-cluster	7m46s

b. Delete the `ClusterResourceSet` for `nvidia-feature-discovery` by running:

```
kubectl delete clusterresourceset nvidia-feature-discovery-my-aws-cluster
```

c. Delete `ConfigMap ClusterResourceSet` referred to by running the following command always using named `nvidia-feature-discover-{CLUSTER_NAME}`. If there is no related `ConfigMap`, then you can move on to the next step.

```
kubectl delete configmap nvidia-feature-discovery-my-aws-cluster
```

d. Get the `kubeconfig` for the cluster by running:

```
dkp get kubeconfig -c my-aws-cluster >> my-aws-cluster.conf
```

e. Delete the corresponding `daemonset` on the remote cluster by running. If there is no related `daemonset`, then you can move on to the next step.

```
kubectl --kubeconfig=my-aws-cluster.conf delete daemonset nvidia-feature-discovery-gpu-feature-discovery -n node-feature-discovery
```

Now that you completed updating your core addons, begin upgrading the Kubernetes version in the section below.

#### 10.4.2.4.1 See Also

[DKP upgrade addons](#) (see page 1013) for more CLI command help.

### 10.4.2.5 Upgrade the Kubernetes Version

When upgrading the Kubernetes version of a cluster, first upgrade the control plane and then the node pools. If you have any additional managed or attached clusters, you need to upgrade the core addons and Kubernetes version for each one.

1. 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 414).
  - If an Azure Machine Image was specified for Azure, you must build a new one with [Konvoy Image Builder](#) (see page 420).
  - If a vSphere template Image was specified for vSphere, you must build a new one with [Konvoy Image Builder](#) (see page 430).
2. 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:** 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} --kubernetes-version=v1.24.6
```

#### Azure

```
dkp update controlplane azure --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6 --compute-gallery-id <Azure Compute Gallery built by KIB for
Kubernetes v1.24.6>
```

### vSphere

```
dkp update controlplane vsphere --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6 --vm-template <vSphere template built by KIB for Kubernetes v1.24.6>
```

### GCP

```
dkp update controlplane gcp --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6 --image=projects/${GCP_PROJECT}/global/images/<GCP image built by KIB
for Kubernetes v1.24.6>
```

### Pre-provisioned

```
dkp update controlplane preprovisioned --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6
```

### EKS

```
dkp update controlplane eks --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.23.12
```



Due to environment restrictions, EKS clusters will display different Kubernetes patch version numbers upon upgrade completion as EKS does not allow for patch version specification.

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|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.

### 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 controlplane aws --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6+fips.0 --ami=<ami-with-fips-id>
```

3. 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](#)<sup>785</sup>) for your critical applications. For more information, refer to [Update Cluster Nodepools](#) (see page 473) documentation.

- First, get a list of all node pools available in your cluster by running the following command:

```
dkp get nodepool --cluster-name ${CLUSTER_NAME}
```

- Select the nodepool you want to upgrade with the command below:

```
export NODEPOOL_NAME=my-nodepool
```

- 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: 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=v1.24.6
```

### Azure

<sup>785</sup> <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

```
dkp update nodepool azure ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --
kubernetes-version=v1.24.6 --compute-gallery-id <Azure Compute Gallery built by KIB
for Kubernetes v1.24.6>
```

### vSphere

```
dkp update nodepool vsphere ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --
kubernetes-version=v1.24.6 --vm-template <vSphere template built by KIB for
Kubernetes v1.24.6>
```

### GCP

```
dkp update nodepool gcp ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6 --image=projects/${GCP_PROJECT}/global/images/<GCP image built by KIB
for Kubernetes v1.24.6>
```

### Pre-provisioned

```
dkp update nodepool preprovisioned ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --
kubernetes-version=v1.24.6
```

### EKS

```
dkp update nodepool eks ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.23.12
```

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=v1.24.6+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 [DKP Upgrade \(see page 0\)](#) for more details.



### 10.4.2.6 Upgrade Managed Clusters

If you have managed clusters, follow these steps to upgrade each cluster:

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}
```

5. Check to see if you have any cluster resource sets that need to be cleaned up:

```
kubectl get clusterresourcesets -n ${CLUSTER_WORKSPACE}
```

6. Delete the ClusterResourceSet for nvidia-feature-discovery by running:

```
kubectl delete clusterresourceset nvidia-feature-discovery-my-aws-cluster -n ${CLUSTER_WORKSPACE}
```

7. Delete ConfigMap ClusterResourceSet referred to by running the following command, ensure you use `nvidia-feature-discover-${CLUSTER_NAME}`. If there is no related ConfigMap, then you can move on to the next step.

```
kubectl delete configmap nvidia-feature-discovery-my-aws-cluster -n ${CLUSTER_WORKSPACE}
```

8. Get the kubeconfig for the managed cluster by running:

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

9. Delete the corresponding daemonset on the remote cluster by running. If there is no related daemonset, then you can move on to the next step.

```
kubectl --kubeconfig=${CLUSTER_NAME}.conf delete daemonset nvidia-feature-
discovery-gpu-feature-discovery -n node-feature-discovery
```

#### 10.4.2.6.1 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 start upgrading the Kubernetes version:

```
dkp update controlplane aws --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6 -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=v1.24.6 -n ${CLUSTER_WORKSPACE}
```

### 10.4.3 DKP Essential Upgrade

Upgrade your Konvoy environment within the DKP Essential License.

### 10.4.3.1 Prerequisites

- Create an [on-demand backup](#) (see page 730) of your current configuration with Velero.
- Follow the steps listed in the [DKP upgrade overview](#) (see page 0).
- Check what version of DKP you have downloaded currently using cli command `dkp version` (see page 1021).
- Ensure that all platform applications in the management cluster are upgraded to avoid compatibility issues with the [Kubernetes version](#) (see page 1056) included in this release. This is done automatically when [upgrading Kommander](#) (see page 1026), so ensure that you upgrade Kommander prior to upgrading Konvoy.
- Set the appropriate environment variables:
  - For air-gapped: Download the required bundles either at our [support site](#)<sup>786</sup> or by using the CLI.
  - For Azure, set the required [environment variables](#) (see page 52).
  - For AWS, set the required [environment variables](#) (see page 47).
  - For vSphere, set the required [environment variables](#) (see page 327).
  - For GCP, set the required [environment variables](#) (see page 58).
- **vSphere only:** If you want to resize your disk, ensure you have reviewed [Create a Base OS Image](#) (see page 430).

### 10.4.3.2 Overview

To upgrade Konvoy for DKP Essential:

1. Upgrade the Cluster API (CAPI) components.
2. Upgrade the core addons.
3. Upgrade the Kubernetes version.

If you have more than one Essential cluster, repeat all of these steps for each Essential cluster (management cluster).




- For pre-provisioned air-gapped environments, you must run `konvoy-image upload artifacts` before beginning Upgrade the Capi Components section below.
- For a full list of DKP Essential features, see [DKP Essential](#) (see page 107).


<sup>786</sup> <https://support.d2iq.com/hc/en-us>

### 10.4.3.2.1 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.

 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](#)<sup>787</sup>.

Run the following 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](#) (see page 938) for flag descriptions for `--with-aws-bootstrap-credentials` and `--aws-service-endpoints`
  - Refer to the HTTP section for details: [Universal Configurations for all Infrastructures](#) (see page 128)

```
dkp upgrade capi-components
```

The command should output something similar to the following:

- ✓ Upgrading CAPI components
- ✓ Waiting **for** CAPI components to be upgraded
- ✓ Initializing **new** CAPI components
- ✓ Deleting Outdated Global ClusterResourceSets

If the upgrade fails, review the prerequisites section and ensure that you've followed the steps in the [DKP upgrade overview](#) (see page 0).

<sup>787</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

### 10.4.3.3 Upgrade the Core Addons

To install the core addons, DKP relies on the `ClusterResourceSet` [Cluster API feature](#)<sup>788</sup>. In the CAPI component upgrade, we deleted the previous set of outdated global `ClusterResourceSets` because prior to DKP 2.4 some addons were installed using a global configuration. In order to support individual cluster upgrades, DKP 2.4 now installs all addons with a unique set of `ClusterResourceSet` 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.



If you have more than one essential cluster, 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](#)<sup>789</sup>.

Perform the following steps to update your addons.

1. For 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](#)<sup>790</sup>.
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.

#### Additional Considerations for EKS

<sup>788</sup> <https://cluster-api.sigs.k8s.io/>

<sup>789</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

<sup>790</sup> <https://kubernetes.io/docs/tasks/access-application-cluster/configure-access-multiple-clusters/>

- The Kubernetes version for EKS clusters supported on DKP 2.4 is now v1.23. As EKS has disabled support for in-tree EBS volume provisioning in favor of [CSI Volumes](#)<sup>791</sup>, DKP automatically deploys a new EBS CSI driver to your EKS cluster when you run the `dkp upgrade addons` command.
- Before running the `dkp upgrade addons` command when deploying EKS clusters, you must add the necessary IAM policy to your worker instances. If you are using the default IAM instance profile name, run the following command:

```
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
```

If you have customized your AWSMachineTemplate to use a different instance profile, review and add the policy to that profile.

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
configmap/nvidia-feature-discovery-my-aws-cluster upgraded
```

<sup>791</sup> <https://kubernetes.io/blog/2019/12/09/kubernetes-1-17-feature-csi-migration-beta/>

4. After addons are upgraded, you must remove deprecated addons. In DKP version 2.4, the `nvidia-feature-discovery` addon is no longer be shipped on new clusters, but rather will be handled by the `nvidia-gpu-operator` Kommander component. To perform removal, execute the following steps:

- a. List all cluster resource sets by running:

```
kubectl get clusterresourcesets
```

NAME	AGE
aws-ebs-csi-my-aws-cluster	7m46s
calico-cni-installation-my-aws-cluster	7m46s
cluster-autoscaler-my-aws-cluster	7m46s
node-feature-discovery-my-aws-cluster	7m46
nvidia-feature-discovery-my-aws-cluster	7m46s

- b. Delete the `ClusterResourceSet` for `nvidia-feature-discovery` by running:

```
kubectl delete clusterresourceset nvidia-feature-discovery-my-aws-cluster
```

- c. Delete `ConfigMap ClusterResourceSet` referred to by running the following command always using named `nvidia-feature-discover- $\{\{CLUSTER\_NAME\}\}$`  :

```
kubectl delete configmap nvidia-feature-discovery-my-aws-cluster
```

- d. Get the `kubeconfig` for the cluster by running:

```
dkp get kubeconfig -c my-aws-cluster >> my-aws-cluster.conf
```

- e. Delete the corresponding `daemonset` on the remote cluster by running:

```
kubectl --kubeconfig=my-aws-cluster.conf delete daemonset nvidia-feature-discovery-gpu-feature-discovery -n node-feature-discovery
```

Now that you completed updating your core addons, begin upgrading the Kubernetes version in the section below.

**See also** [DKP upgrade addons \(see page 1013\)](#)

### 10.4.3.4 Upgrade the Kubernetes version

When upgrading the Kubernetes version of a cluster, first upgrade the control plane and then the node pools.

If you have any additional managed or attached clusters, you need to upgrade the core addons and Kubernetes version for each one.

1. 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 414).
  - If an Azure Machine Image was specified for Azure, you must build a new one with [Konvoy Image Builder](#)<sup>792</sup>.
  - If a vSphere template Image was specified for vSphere, you must build a new one with [Konvoy Image Builder](#) (see page 333).
2. 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:** 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} --kubernetes-
version=v1.24.6
```

### Azure

```
dkp update controlplane azure --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6 --compute-gallery-id <Azure Compute Gallery built by KIB for
Kubernetes v1.24.6>
```

### vSphere

```
dkp update controlplane vsphere --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6 --vm-template <vSphere template built by KIB for Kubernetes v1.24.6>
```

### GCP

```
dkp update controlplane gcp --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6 --image=projects/${GCP_PROJECT}/global/images/<GCP image built by KIB
for Kubernetes v1.24.6>
```

### Pre-provisioned

```
dkp update controlplane preprovisioned --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.24.6
```

<sup>792</sup> <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29918922/Create+a+custom+Azure+Image>



## EKS

```
dkp update controlplane eks --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.23.12
```

- Due to environment restrictions, EKS clusters will display different Kubernetes patch version numbers upon upgrade completion as EKS does not allow for patch version specification.

The output should be similar to the below example, with the provider name corresponding to the CLI you executed from the choices above:

- NOTE:** Some advanced options are available for various providers. An example would be regarding AMI instance type: `aws: --ami, --instance-type`. To see all the options for your particular provider, run this command `dkp update controlplane aws|vsphere|preprovisioned|azure --help` for more advance options.

The output should be similar to:

```
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
```

- If upgrading a FIPS cluster, there is a bug in the upgrade of `kube-proxy` `DaemonSet` in that it doesn't get automatically upgraded. After updating the control plane, execute the following command to finish the kube-proxy upgrade:

```
kubectl set image -n kube-system daemonset.v1.apps/kube-proxy kube-proxy=docker.io/
mesosphere/kube-proxy:v1.24.4_fips.0
```

3. Upgrade the Kubernetes version of your 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](https://kubernetes.io/docs/concepts/workloads/pods/disruptions/)<sup>793</sup>) for your critical applications. For more information, refer to [Update Cluster Nodepools](#) (see [page 473](#)) 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=v1.24.6
```

#### Azure

```
dkp update nodepool azure ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-version=v1.24.6 --compute-gallery-id <Azure Compute Gallery built by KIB for Kubernetes v1.24.6>
```

#### vSphere

```
dkp update nodepool vsphere ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-version=v1.24.6 --vm-template <vSphere template built by KIB for Kubernetes v1.24.6>
```

#### GCP

```
dkp update nodepool gcp ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-version=v1.24.6 --image=projects/${GCP_PROJECT}/global/images/<GCP image built by KIB for Kubernetes v1.24.6>
```

#### Pre-provisioned

```
dkp update nodepool preprovisioned ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-version=v1.24.6
```

<sup>793</sup> <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

## EKS

```
dkp update nodepool eks ${NODEPOOL_NAME} --cluster-name=${CLUSTER_NAME} --kubernetes-
version=v1.23.12
```

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
```

4. 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=v1.24.6+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 [DKP Upgrade \(see page 0\)](#) for more details.



For the overall process for upgrading to the latest version of DKP, refer back to [DKP Upgrade \(see page 0\)](#).

## 10.4.4 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](#)<sup>794</sup>) 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.

<sup>794</sup> <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>


### 10.4.4.1 Prerequisites:

- Deploy [Pod Disruption Budget](#)<sup>795</sup> (PDB)
- [Konvoy Image Builder](#) (see page 411) (KIB)

### 10.4.4.2 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.

 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 411).

---

<sup>795</sup> <https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>

## 10.5 Upgrade Custom Applications

### 10.5.1 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 \(see page 1055\)](#) 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.

Once 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.

# 11 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.com](mailto:sales@d2iq.com)<sup>796</sup> before attempting to download or install DKP.

### Release Notes for specific releases:

- [DKP 2.4.0 Release Notes](#) (see page 1055)
- [DKP 2.4.1 Release Notes](#) (see page 1078)

## 11.1 DKP 2.4.0 Release Notes

DKP® version 2.4.0 was released on November 21st, 2022.

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.com](mailto:sales@d2iq.com)<sup>797</sup> before attempting to download or install DKP.

### 11.1.1 Release Summary

Welcome to D2iQ Kubernetes Platform (DKP) 2.4! This release focuses on **MultiCluster - MultiCloud** capabilities. Also this release provides fixes to reported issues, integrates changes from previous releases, and maintains compatibility and support for other packages used in DKP. Below, you will find information about the following:

- [DKP 2.4.0 Supported Kubernetes Versions](#) (see page 1056)
- [DKP 2.4.0 Features and Enhancements](#) (see page 1057)
- [DKP 2.4.0 Security Enhancements](#) (see page 1062)
- [DKP 2.4.0 Components and Applications](#) (see page 1063)

---

<sup>796</sup> <mailto:sales@d2iq.com>

<sup>797</sup> <mailto:sales@d2iq.com>

- [DKP 2.4.0 Customer Incidents](#) (see page 1071)
- [DKP 2.4.0 Known Issues and Limitations](#) (see page 1072)
- [DKP 2.4.0 Deprecations](#) (see page 1076)
- [Kubernetes major updates and deprecations](#) (see page 1077)

### 11.1.2 Additional resources


- For more information about working with native Kubernetes, see the [Kubernetes documentation](#)<sup>798</sup>.
- For a full list of attributed 3rd party software, see <http://d2iq.com/legal/3rd>.

### 11.1.3 DKP 2.4.0 Supported Kubernetes Versions


#### 11.1.3.1 Deploying Clusters Versions

The newest and oldest Kubernetes versions used with DKP must be within one minor version.

- newest supported Kubernetes version is at **1.24.6 for deploying DKP** for all providers except EKS.
- Kubernetes versions are supported at **1.23.x for deploying DKP on EKS**.

 DKP 2.4 comes with support for Kubernetes 1.24.6, enabling you to benefit from the latest features and security fixes in upstream Kubernetes. This release comes with approximately 46 enhancements. To read more about major features in this release, visit <https://kubernetes.io/blog/2022/05/03/kubernetes-1-24-release-announcement/>.

#### 11.1.3.2 Attaching Clusters Versions

 When attempting to deploy a cluster with a lower Kubernetes version than the DKP default, you need to build and use an AMI with that lower version of Kubernetes. See [Konvoy Image Builder](#) (see page 411) for documentation on building images. Failure to do this could result in a failure to deploy error.

DKP 2.4.0 supports attaching clusters with the following Kubernetes versions:

---

<sup>798</sup> <https://kubernetes.io/docs/home/>

Product	Compatible Kubernetes Versions
EKS	1.23.x
AKS	1.24.x
GKE	1.24.x

## 11.1.4 DKP 2.4.0 Features and Enhancements

The following improvements are included in this release.

### 11.1.4.1 DKP Upgrades

DKP 2.4 supports upgrades of Kubernetes and the Platform Applications from DKP 2.3 For more information see [Supported Upgrade Paths \(see page 0\)](#).



For environments with **Kaptain** only: [Upgrade Kaptain](#)<sup>799</sup> to the latest version BEFORE upgrading DKP. This version of DKP removes a MinIO version required for Kaptain versions 2.1 and earlier.

### 11.1.4.2 NVIDIA GPUs

This feature will enable you who are using GPU nodes for their computational heavy applications, to run GPU nodes On-prem, and in Air-gapped environments, enabling our security conscious customers to use GPU workloads on DKP.

Konvoy Image Builder instructions have been added for building an image for GPU capable systems under the Konvoy Image Builder section called [KIB for GPU \(see page 427\)](#).

When upgrading a [nodepool \(see page 1037\)](#) running GPU workloads, users must drain the existing nodes in the nodepool and replace them with new nodes configured to run GPU workloads. With DKP 2.4, the nodes with NVIDIA GPUs are configured with `nvidia-gpu-operator` (<https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/overview.html>) and NVIDIA drivers to support the container runtime. Critical GPU workloads can be impacted while upgrading the nodepool.

Users may have to specify configuration overrides for the `nvidia-gpu-operator` app prior to upgrading the Kommander component of DKP and for any workspace upgrades that have attached clusters running GPU nodes. For information on how on how to do this, refer to <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/144214580/2.4%2BUgrade%2BKommander#Prerequisites-for-Configurations-with-NVIDIA-GPU-Nodes> (see page 1027).

<sup>799</sup> <https://d2iq.atlassian.net/wiki/spaces/DKAP/pages/125894830/Upgrade+Kaptain+and+Migrate+your+Data>



To ensure minimum downtime and maintain high availability of the critical application workloads during the upgrade process, we recommend deploying Pod Disruption Budget (see page 473) (<https://kubernetes.io/docs/concepts/workloads/pods/disruptions/>) 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.

### 11.1.4.3 RHEL 8.6 Support

DKP 2.4 now supports the RHEL 8.6 OS.

### 11.1.4.4 Installation Enhancements

The DKP Kommander install command will now wait for all enabled applications to become ready by default. This can be disabled by setting the wait flag to false i.e. `--wait=false`.

You can view the progress by setting the `-v 2` flag.

### 11.1.4.5 Cluster-scoped Application Configuration via the UI

DKP UI now allows for the custom configuration and deployment of applications on a per-cluster basis. This functionality allows you to use DKP in a multi-cluster scenario without restricting the management of multiple clusters from a single workspace. For instructions, refer to [Cluster-scoped Application Configuration via the UI](#) (see page 582).

### 11.1.4.6 Enhanced KIB Documentation

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. Inside the new [Konvoy Image Builder](#) (see page 411) section of documentation, you will find provider specific instructions for building an image compatible with a specific environment as well as a [version compatibility table](#) (see page 412). Any override files that are needed are found within that section as well. Plus, inside each [Advanced Configuration](#) (see page 128) for your provider, the KIB instructions are listed in the order in which they should be performed.

### 11.1.4.7 Complete DKP Air-gapped Bundle for Download

In DKP 2.4.0 for air-gapped environments, a new compressed archive is available for download which contains all the DKP components needed for an air-gapped installation, including the DKP CLI, Konvoy Image Builder and all air-gapped bundles. Navigate to the [Download](#) (see page 100) site for the new Complete Air-gapped Bundle.

The download site contains these sections:

- DKP CLI binary downloads
- Konvoy Image Builder downloads
- **NEW** Complete DKP Air-gapped Bundle - the complete bundle of all the components needed for an air-gapped install

- Air-gapped Components for Individual Download - in the event you want to customize and only download each component separately to create your own bundle rather than the new complete bundle

#### 11.1.4.8 Get Started and DKP Terms Section of Documentation

Our Documentation sections regarding installation have been reorganized for better understanding. In the [Get Started with DKP \(see page 113\)](#) section, you will find step by step instructions to go from [Download \(see page 100\)](#) all the way through log in to your UI dashboard in order. At the bottom of each page you will find a link to the next step or next relevant topic to help guide you. Other new sections such as [DKP Concepts and Terms \(see page 113\)](#) contain helpful terminology for understanding the intricacies of the product. Also, an [Advanced Kommander \(see page 475\)](#) section detailing installation was broken out and placed above the Day 2 configuration information.

#### 11.1.4.9 Ability to Upgrade GCP

With the release of DKP 2.4, you are now able to upgrade their DKP clusters provisioned on Google Cloud Provider (GCP) from DKP 2.3 to DKP 2.4. Refer to both the [DKP Enterprise Upgrade \(see page 1033\)](#) and [Upgrade DKP \(see page 1024\)](#) documentation for instructions.

#### 11.1.4.10 Provision DKP Clusters on Azure Using DKP UI

With the release of DKP 2.4, you are now able to provision DKP clusters on Microsoft Azure via the DKP UI. Refer to the [Create a new Azure Cluster via UI \(see page 258\)](#) documentation for instructions.

#### 11.1.4.11 Rook Ceph Default Storage

[Rook Ceph<sup>800</sup>](#) is a storage platform that is now included by default for usage in DKP Platform Applications. [This replaces MinIO \(see page 1076\)](#).

It is important to note that when you upgrade from `2.3.x` -> `2.4.x`, Ceph is automatically installed if any of the following are true:

- `minio-operator` was installed (used by logging stack).
- `velero` was installed (used for backups).


However, if you did not have `minio-operator` installed in `2.3.x` and instead had configured `velero` to work with an [external cloud storage \(see page 731\)](#) such as Amazon S3 or Azure Blob Storage, then you do not need `rook-ceph` and `rook-ceph-cluster` post upgrade.

In order to explicitly disable installing these apps during upgrade, specify the following in the command line, then run the command:

---

<sup>800</sup> <https://rook.io/docs/rook/v1.9/quickstart.html>

```
dkp upgrade kommander --disable-appdeployments rook-ceph,rook-ceph-cluster
```

 The command above fails if `minio-operator` is installed as you cannot disable ceph installation if minio-operator is installed.

Refer to [Rook Ceph in DKP](#) (see page 858) for additional information.

## 11.1.4.12 DKP Insights Enhancements

### 11.1.4.12.1 Release Summary

Welcome to DKP Insights 0.3.x! This release provides improved user experience, refined workflows, and maintains compatibility and support for other packages used in Insights.

For an overview of the features in Insights, see:



#### 11.1.4.12.2 Supported Kubernetes Versions

Insights supports the same [Kubernetes versions as the DKP platform](#)<sup>801802</sup>. For an overview of compatible DKP and Insights versions, see [\(2.4\) DKP and Insights Compatibility](#)<sup>803804</sup>.

#### 11.1.4.12.3 Expanded Cluster Details

With this release, we now provide the cluster name and project name in the Insight Alert Details Page, which enables you to get a more inclusive summary of each insight alert from the dashboard.

---

801 <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/144213313/2.4+DKP+2.4.0+Supported+Kubernetes+Versions>

802 <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/144213313/2.4+DKP+2.4.0+Supported+Kubernetes+Versions>

803 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/145261084/%282.4%29+DKP+and+Insights+Compatibility>

804 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/145261084/%282.4%29+DKP+and+Insights+Compatibility>

#### 11.1.4.12.4 Resolve or Snooze Insight Alerts

You can [clear Insight alerts from the open view](#)<sup>805806</sup> by selecting either *Mark as Resolved* or *Snooze*.

#### 11.1.4.12.5 Execute Single Actions on Multiple Insight Items

From the Insights table, you can now [select multiple insight items](#)<sup>807808</sup> and then select a single action for all items.

#### 11.1.4.12.6 Additional 3rd-party Scanning Tools

To assist with scanning configuration anomalies, we integrated these [additional third-party, open-source components](#)<sup>809810</sup> into the DKP Insights Engine:

- Trivy
- Pluto
- Nova

##### 11.1.4.12.6.1 Trivy

#### Runtime CVE scan for customer workloads

To keep the Kubernetes workload and cluster deployments always safe, secure, and operational, this release of Insights adds the functionality to perform [scans of CVE/CIS databases](#)<sup>811</sup> with [\(2.4\) Trivy](#)<sup>812813</sup>. When enabled, this feature automatically notifies the users about any vulnerabilities and security issues in all installed docker images and Kubernetes cluster deployments. Currently, this function is available in non-air-gapped environments only.

### 11.1.5 DKP 2.4.0 Security Enhancements

#### 11.1.5.1 Security Enhancements for DKP and Kommander roles

In DKP 2.4, four default user roles with edit and view rights were updated with stricter rules to improve security. These rules result in greater access control to resources and actions, limiting the scope of your users' *view* or *edit* rights.

---

805 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/145817647/2.4%2BDKP%2BInsight%2BAlerts#Resolve-or-Snooze-Alerts>

806 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/145817647/2.4%2BDKP%2BInsight%2BAlerts#Resolve-or-Snooze-Alerts>

807 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/145817647/2.4+DKP+Insight+Alerts#resolve-snooze-alerts>

808 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/145817647/2.4+DKP+Insight+Alerts#resolve-snooze-alerts>

809 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/145261037/2.4+Configuration+Anomalies>

810 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/145261037/2.4+Configuration+Anomalies>

811 <https://aquasecurity.github.io/trivy/v0.35/docs/vulnerability/detection/data-source/>

812 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/145261136/%282.4%29+Trivy>

813 <https://d2iq.atlassian.net/wiki/spaces/DINS/pages/145261136/%282.4%29+Trivy>

### 11.1.5.1.1 The following roles have been improved:

Role	Access Control	Description
dkp-kommander-edit	UI Access	Can edit fields and existing configurations via the UI. Cannot edit infrastructure providers or other user's roles.
dkp-kommander-view	UI Access	Can access the UI and view all workspaces and configurations. Cannot make modifications.
kommander-workspace-edit	Kommander Resource Access	Can edit workspace-specific resources for the workspace they are in.
kommander-workspace-view	Kommander Resource Access	Can view workspace-specific resources for the workspace they are in.

With DKP 2.4, some operations that were previously available to users with *view* and *edit* rights are no longer available. If you are using any of these roles to manage access to the UI (dkp-kommander roles) or to access Kommander resources and components (kommander-workspace roles), verify that your users still have the required rights to perform their usual tasks. If required, you can assign them additional roles to expand the scope of access.

## 11.1.6 DKP 2.4.0 Components and Applications

The following are component and application versions for DKP 2.4.0.

### 11.1.6.1 Components

Component Name	Version
Cluster API Core (CAPI)	1.1.3-d2iq.5
Cluster API AWS Infrastructure Provider (CAPA)	1.4.1
Cluster API Google Cloud Infrastructure Provider (CAPG)	1.1.0

Component Name	Version
Cluster API Pre-provisioned Infrastructure Provider (CAPPP)	0.12.1
Cluster API vSphere Infrastructure Provider (CAPV)	1.3.1
Cluster API Azure Infrastructure Provider (CAPZ)	1.3.2
Konvoy Image Builder (KIB)	<a href="#">1.24.2</a> <sup>814</sup>
containerd	1.4.13
etcd	3.4.13

### 11.1.6.2 Applications

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Centralized Grafana	centralized-grafana	40.0.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 40.0.0</a><sup>815</sup></li> <li>• prometheus-operator: 0.59.1</li> <li>• grafana: 9.1.6</li> </ul>	<a href="#">Link</a> <sup>816</sup>	<a href="#">Link</a> <sup>817</sup>
Centralized Kubecost	centralized-kubecost	0.28.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.28.0</a><sup>818</sup></li> <li>• kubecost: 1.97.0</li> </ul>	<a href="#">Link</a> <sup>819</sup>	<a href="#">Link</a> <sup>820</sup>

<sup>814</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.24.2>

<sup>815</sup> <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/40.0.0>

<sup>816</sup> <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/40.0.0?modal=values>

<sup>817</sup> <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/centralized-grafana/40.0.0/defaults/cm.yaml>

<sup>818</sup> <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.28.0>

<sup>819</sup> <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.28.0?modal=values>

<sup>820</sup> <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/centralized-kubecost/0.28.0/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Cert Manager	cert-manager	1.9.1	<ul style="list-style-type: none"> <li>chart: 1.9.1<sup>821</sup></li> <li>cert-manager: 1.9.1</li> </ul>	<a href="#">Link</a> <sup>822</sup>	<a href="#">Link</a> <sup>823</sup>
Chartmuseum	chartmuseum	3.9.0	<ul style="list-style-type: none"> <li>chart: 3.9.0<sup>824</sup></li> <li>chartmuseum: 3.9.0</li> </ul>	<a href="#">Link</a> <sup>825</sup>	<a href="#">Link</a> <sup>826</sup>
Dex	dex	2.10.0	<ul style="list-style-type: none"> <li>chart: 2.10.0<sup>827</sup></li> <li>dex: 2.31.0</li> </ul>	<a href="#">Link</a> <sup>828</sup>	<a href="#">Link</a> <sup>829</sup>
Dex K8s Authenticator	dex-k8s-authenticator	1.2.13	<ul style="list-style-type: none"> <li>chart: 1.2.13<sup>830</sup></li> <li>dex-k8s-authenticator: 1.2.4</li> </ul>	<a href="#">Link</a> <sup>831</sup>	<a href="#">Link</a> <sup>832</sup>
DKP Insights Management	dkp-insights-management	0.3.1	<ul style="list-style-type: none"> <li>chart: 0.3.1</li> <li>dkp-insights-management: 0.3.1</li> </ul>	N/A	<a href="#">Link</a> <sup>833</sup>
External DNS	external-dns	6.10.2	<ul style="list-style-type: none"> <li>chart: 6.10.2<sup>834</sup></li> <li>external-dns: 0.12.2</li> </ul>	<a href="#">Link</a> <sup>835</sup>	<a href="#">Link</a> <sup>836</sup>

821 <https://artifacthub.io/packages/helm/cert-manager/cert-manager/1.9.1>

822 <https://artifacthub.io/packages/helm/cert-manager/cert-manager/1.9.1?modal=values>

823 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/cert-manager/1.9.1/defaults/cm.yaml>

824 <https://artifacthub.io/packages/helm/chartmuseum/chartmuseum/3.9.0>

825 <https://artifacthub.io/packages/helm/chartmuseum/chartmuseum/3.9.0?modal=values>

826 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/chartmuseum/3.9.0/defaults/cm.yaml>

827 <https://artifacthub.io/packages/helm/mesosphere-stable/dex/2.10.0>

828 <https://artifacthub.io/packages/helm/mesosphere-stable/dex/2.10.0?modal=values>

829 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/dex/2.10.0/defaults/cm.yaml>

830 <https://artifacthub.io/packages/helm/mesosphere/dex-k8s-authenticator/1.2.13>

831 <https://artifacthub.io/packages/helm/mesosphere/dex-k8s-authenticator/1.2.13?modal=values>

832 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/dex-k8s-authenticator/1.2.13/defaults/cm.yaml>

833 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/dkp-insights-management/0.3.1/defaults/cm.yaml>

834 <https://artifacthub.io/packages/helm/bitnami/external-dns/6.10.2>

835 <https://artifacthub.io/packages/helm/bitnami/external-dns/6.10.2?modal=values>

836 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/external-dns/6.10.2/defaults/cm.yaml>



Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Fluent Bit	fluent-bit	0.20.9	<ul style="list-style-type: none"> <li>chart: 0.20.9<sup>837</sup></li> <li>fluent-bit: 1.9.9</li> </ul>	<a href="#">Link</a> <sup>838</sup>	<a href="#">Link</a> <sup>839</sup>
Gatekeeper	gatekeeper	3.10.0	<ul style="list-style-type: none"> <li>chart: 3.10.0<sup>840</sup></li> <li>gatekeeper: 3.10.0</li> </ul>	<a href="#">Link</a> <sup>841</sup>	<a href="#">Link</a> <sup>842</sup>
Gitea	gitea	6.0.1	<ul style="list-style-type: none"> <li>chart: 6.0.1<sup>843</sup></li> <li>gitea: 1.17.2</li> </ul>	<a href="#">Link</a> <sup>844</sup>	<a href="#">Link</a> <sup>845</sup>
Grafana Logging	grafana-logging	6.38.1	<ul style="list-style-type: none"> <li>chart: 6.38.1<sup>846</sup></li> <li>grafana: 9.1.5</li> </ul>	<a href="#">Link</a> <sup>847</sup>	<a href="#">Link</a> <sup>848</sup>
Grafana Loki	grafana-loki	0.48.6	<ul style="list-style-type: none"> <li>chart: 0.48.4<sup>849</sup></li> <li>loki: 2.5.0</li> </ul>	<a href="#">Link</a> <sup>850</sup>	<a href="#">Link</a> <sup>851</sup>
Istio	istio	1.15.3	<ul style="list-style-type: none"> <li>chart: 1.15.3<sup>852</sup></li> <li>istio: 1.15.3</li> </ul>	<a href="#">Link</a> <sup>853</sup>	<a href="#">Link</a> <sup>854</sup>

837 <https://artifacthub.io/packages/helm/fluent/fluent-bit/0.20.9>

838 <https://artifacthub.io/packages/helm/fluent/fluent-bit/0.20.9?modal=values>

839 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/fluent-bit/0.20.9/defaults/cm.yaml>

840 <https://artifacthub.io/packages/helm/gatekeeper/gatekeeper/3.10.0>

841 <https://artifacthub.io/packages/helm/gatekeeper/gatekeeper/3.10.0?modal=values>

842 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/gatekeeper/3.10.0/defaults/cm.yaml>

843 <https://artifacthub.io/packages/helm/gitea/gitea/6.0.1>

844 <https://artifacthub.io/packages/helm/gitea/gitea/6.0.1?modal=values>

845 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/gitea/6.0.1/defaults/cm.yaml>

846 <https://artifacthub.io/packages/helm/grafana/grafana/6.38.1>

847 <https://artifacthub.io/packages/helm/grafana/grafana/6.38.1?modal=values>

848 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/grafana-logging/6.38.1/defaults/cm.yaml>

849 <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.48.4>

850 <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.48.4?modal=values>

851 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/grafana-loki/0.48.6/defaults/cm.yaml>

852 <https://artifacthub.io/packages/helm/mesosphere/istio/1.15.3>

853 <https://artifacthub.io/packages/helm/mesosphere/istio/1.15.3?modal=values>

854 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/istio/1.15.3/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Jaeger	jaeger	2.36.0	<ul style="list-style-type: none"> <li>chart: 2.36.0<sup>855</sup></li> <li>jaeger: 1.38.0</li> </ul>	<a href="#">Link</a> <sup>856</sup>	<a href="#">Link</a> <sup>857</sup>
Karma	karma	2.0.1	<ul style="list-style-type: none"> <li>chart: 2.0.1<sup>858</sup></li> <li>karma: 0.70</li> </ul>	<a href="#">Link</a> <sup>859</sup>	<a href="#">Link</a> <sup>860</sup>
Kiali	kiali	1.57.1	<ul style="list-style-type: none"> <li>chart: 1.57.1<sup>861</sup></li> <li>kiali: 1.57.1</li> </ul>	<a href="#">Link</a> <sup>862</sup>	<a href="#">Link</a> <sup>863</sup>
Knative	knative	0.4.0	<ul style="list-style-type: none"> <li>chart: 0.4.0<sup>864</sup></li> <li>knative: 0.22.3</li> </ul>	<a href="#">Link</a> <sup>865</sup>	<a href="#">Link</a> <sup>866</sup>
Flux	kommander-flux	0.36.0	<ul style="list-style-type: none"> <li>chart: N/A</li> <li>flux: 0.36.0</li> </ul>	N/A	N/A
Kube OIDC Proxy	kube-oidc-proxy	0.3.2	<ul style="list-style-type: none"> <li>chart: 0.3.1<sup>867</sup></li> <li>kube-oidc-proxy: 0.3.0</li> </ul>	<a href="#">Link</a> <sup>868</sup>	<a href="#">Link</a> <sup>869</sup>

---

855 <https://artifacthub.io/packages/helm/jaegertracing/jaeger-operator/2.36.0>

856 <https://artifacthub.io/packages/helm/jaegertracing/jaeger-operator/2.36.0?modal=values>

857 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/jaeger/2.36.0/defaults/cm.yaml>

858 <https://artifacthub.io/packages/helm/mesosphere-stable/karma/2.0.1>

859 <https://artifacthub.io/packages/helm/mesosphere-stable/karma/2.0.1?modal=values>

860 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/karma/2.0.1/defaults/cm.yaml>

861 <https://artifacthub.io/packages/helm/kiali/kiali-operator/1.57.1>

862 <https://artifacthub.io/packages/helm/kiali/kiali-operator/1.57.1?modal=values>

863 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/kiali/1.57.1/defaults/cm.yaml>

864 <https://artifacthub.io/packages/helm/mesosphere/knative/0.4.0>

865 <https://artifacthub.io/packages/helm/mesosphere/knative/0.4.0?modal=values>

866 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/knative/0.4.0/defaults/cm.yaml>

867 <https://artifacthub.io/packages/helm/mesosphere/kube-oidc-proxy/0.3.1>

868 <https://artifacthub.io/packages/helm/mesosphere/kube-oidc-proxy/0.3.1?modal=values>

869 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/kube-oidc-proxy/0.3.2/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Kube Prometheus Stack	kube-prometheus-stack	40.0.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 40.0.0</a><sup>870</sup></li> <li>• prometheus-operator: 0.59.1</li> <li>• grafana: 9.1.6</li> <li>• prometheus: 2.38.0</li> <li>• prometheus-alertmanager: 0.24.0</li> </ul>	<a href="#">Link</a> <sup>871</sup>	<a href="#">Link</a> <sup>872</sup>
Kubecost	kubecost	0.28.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.28.0</a><sup>873</sup></li> <li>• kubecost: 1.97.0</li> </ul>	<a href="#">Link</a> <sup>874</sup>	<a href="#">Link</a> <sup>875</sup>
Kubefed	kubefed	0.10.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.10.0</a><sup>876</sup></li> <li>• kubefed: 0.10.0</li> </ul>	<a href="#">Link</a> <sup>877</sup>	<a href="#">Link</a> <sup>878</sup>
Kubernetes Dashboard	kubernetes-dashboard	5.11.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 5.11.0</a><sup>879</sup></li> <li>• kubernetes-dashboard: 2.7.0</li> </ul>	<a href="#">Link</a> <sup>880</sup>	<a href="#">Link</a> <sup>881</sup>
Kubetunnel	kubetunnel	0.0.15	<ul style="list-style-type: none"> <li>• chart: 0.0.15</li> <li>• kubetunnel: 0.0.15</li> </ul>	N/A	<a href="#">Link</a> <sup>882</sup>

870 <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/40.0.0>

871 <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/40.0.0?modal=values>

872 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/kube-prometheus-stack/40.0.0/defaults/cm.yaml>

873 <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.28.0>

874 <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.28.0?modal=values>

875 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/kubecost/0.28.0/defaults/cm.yaml>

876 <https://artifacthub.io/packages/helm/kubefed/kubefed/0.10.0>

877 <https://artifacthub.io/packages/helm/kubefed/kubefed/0.10.0?modal=values>

878 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/kubefed/0.10.0/defaults/cm.yaml>

879 <https://artifacthub.io/packages/helm/k8s-dashboard/kubernetes-dashboard/5.11.0>

880 <https://artifacthub.io/packages/helm/k8s-dashboard/kubernetes-dashboard/5.11.0?modal=values>

881 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/kubernetes-dashboard/5.11.0/defaults/cm.yaml>

882 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/kubetunnel/0.0.15/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Logging Operator	logging-operator	3.1 7.9	<ul style="list-style-type: none"> <li>• <a href="#">chart: 3.17.9</a><sup>883</sup></li> <li>• logging-operator: 3.17.9</li> <li>• logging-operator-logging: 3.17.9</li> </ul>	<a href="#">Link</a> <sup>884</sup>	<a href="#">Link</a> <sup>885</sup>
NFS Server Provisioner	nfs-server-provisioner	0.6. 0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.6.0</a><sup>886</sup></li> <li>• nfs-server-provisioner: 2.3.0</li> </ul>	<a href="#">Link</a> <sup>887</sup>	<a href="#">Link</a> <sup>888</sup>
NVIDIA GPU Operator	nvidia-gpu-operator	1.1 1.1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 1.11.1</a><sup>889</sup></li> <li>• nvidia-gpu-operator: 1.11.1</li> </ul>	<a href="#">Link</a> <sup>890</sup>	<a href="#">Link</a> <sup>891</sup>
Grafana (project)	project-grafana-logging	6.3 8.1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 6.38.1</a><sup>892</sup></li> <li>• grafana: 9.1.5</li> </ul>	<a href="#">Link</a> <sup>893</sup>	<a href="#">Link</a> <sup>894</sup>
Grafana Loki (project)	project-grafana-loki	0.4 8.6	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.48.4</a><sup>895</sup></li> <li>• loki: 2.5.0</li> </ul>	<a href="#">Link</a> <sup>896</sup>	<a href="#">Link</a> <sup>897</sup>

883 <https://artifacthub.io/packages/helm/banzaicloud-stable/logging-operator/3.17.9>

884 <https://artifacthub.io/packages/helm/banzaicloud-stable/logging-operator/3.17.9?modal=values>

885 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/logging-operator/3.17.9/defaults/cm.yaml>

886 <https://artifacthub.io/packages/helm/mesosphere/nfs-server-provisioner/0.6.0>

887 <https://artifacthub.io/packages/helm/mesosphere/nfs-server-provisioner/0.6.0?modal=values>

888 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/nfs-server-provisioner/0.6.0/defaults/cm.yaml>

889 <https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/getting-started.html#chart-customization-options>

890 <https://raw.githubusercontent.com/NVIDIA/gpu-operator/v1.11.1/deployments/gpu-operator/values.yaml>

891 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/nvidia-gpu-operator/1.11.1/defaults/cm.yaml>

892 <https://artifacthub.io/packages/helm/grafana/grafana/6.38.1>

893 <https://artifacthub.io/packages/helm/grafana/grafana/6.38.1?modal=values>

894 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/project-grafana-logging/6.38.1/defaults/cm.yaml>

895 <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.48.4>

896 <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.48.4?modal=values>

897 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/project-grafana-loki/0.48.6/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Prometheus Adapter	prometheus-adapter	3.4.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 3.4.0</a><sup>898</sup></li> <li>• prometheus-adapter: 0.10.0</li> </ul>	<a href="#">Link</a> <sup>899</sup>	<a href="#">Link</a> <sup>900</sup>
Reloader	reloader	0.0.121	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.0.121</a><sup>901</sup></li> <li>• reloader: 0.0.121</li> </ul>	<a href="#">Link</a> <sup>902</sup>	<a href="#">Link</a> <sup>903</sup>
Rook Ceph	rook-ceph	1.10.3	<ul style="list-style-type: none"> <li>• <a href="#">chart: 1.10.3</a><sup>904</sup></li> <li>• rook-ceph: 1.10.3</li> </ul>	<a href="#">Link</a> <sup>905</sup>	<a href="#">Link</a> <sup>906</sup>
Rook Ceph Cluster	rook-ceph-cluster	1.10.3	<ul style="list-style-type: none"> <li>• <a href="#">chart: 1.10.3</a><sup>907</sup></li> <li>• rook-ceph: 1.10.3</li> <li>• rook-ceph-cluster: 17.2.3</li> </ul>	<a href="#">Link</a> <sup>908</sup>	<a href="#">Link</a> <sup>909</sup>
Thanos	thanos	0.4.8	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.4.6</a><sup>910</sup></li> <li>• thanos: 0.17.1</li> </ul>	<a href="#">Link</a> <sup>911</sup>	<a href="#">Link</a> <sup>912</sup>

898 <https://artifacthub.io/packages/helm/prometheus-community/prometheus-adapter/3.4.0>

899 <https://artifacthub.io/packages/helm/prometheus-community/prometheus-adapter/3.4.0?modal=values>

900 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/prometheus-adapter/3.4.0/defaults/cm.yaml>

901 <https://github.com/stakater/Reloader/tree/v0.0.121/README.md#helm-charts>

902 <https://artifacthub.io/packages/helm/stakater/reloader/0.0.121?modal=values>

903 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/reloader/0.0.121/defaults/cm.yaml>

904 <https://github.com/rook/rook/tree/v1.10.3/Documentation/Helm-Charts/operator-chart.md>

905 <https://raw.githubusercontent.com/rook/rook/v1.10.3/deploy/charts/rook-ceph/values.yaml>

906 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/rook-ceph/1.10.3/defaults/cm.yaml>

907 <https://github.com/rook/rook/tree/v1.10.3/Documentation/Helm-Charts/ceph-cluster-chart.md>

908 <https://raw.githubusercontent.com/rook/rook/v1.10.3/deploy/charts/rook-ceph-cluster/values.yaml>

909 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/rook-ceph-cluster/1.10.3/defaults/cm.yaml>

910 <https://artifacthub.io/packages/helm/banzaicloud-stable/thanos/0.4.6>

911 <https://artifacthub.io/packages/helm/banzaicloud-stable/thanos/0.4.6?modal=values>

912 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/thanos/0.4.8/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Traefik	traefik	10.30.1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 10.30.1</a><sup>913</sup></li> <li>• traefik: 2.8.7</li> </ul>	<a href="#">Link</a> <sup>914</sup>	<a href="#">Link</a> <sup>915</sup>
Traefik ForwardAuth	traefik-forward-auth	0.3.8	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.3.8</a><sup>916</sup></li> <li>• traefik-forward-auth: 3.1.0</li> </ul>	<a href="#">Link</a> <sup>917</sup>	<a href="#">Link</a> <sup>918</sup>
Velero	velero	3.3.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 2.31.8</a><sup>919</sup></li> <li>• velero: 1.9.2</li> </ul>	<a href="#">Link</a> <sup>920</sup>	<a href="#">Link</a> <sup>921</sup>

### 11.1.7 DKP 2.4.0 Customer Incidents

The following resolved incidents have been included in this release.

#### 11.1.7.1 kubernetes/crictl cannot pull images from harbor registry

##### D2IQ-93950

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.

#### 11.1.7.2 DKP users can download cluster admin kubeconfigs

##### D2IQ-93695

An incident allowing for non-admin DKP UI users to download cluster admin kubeconfigs has been resolved, making it so only administrator-level roles are allowed to download these configurations.

<sup>913</sup> <https://artifacthub.io/packages/helm/traefik/traefik/10.30.1>

<sup>914</sup> <https://artifacthub.io/packages/helm/traefik/traefik/10.30.1?modal=values>

<sup>915</sup> <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/traefik/10.30.1/defaults/cm.yaml>

<sup>916</sup> <https://artifacthub.io/packages/helm/mesosphere/traefik-forward-auth/0.3.8>

<sup>917</sup> <https://artifacthub.io/packages/helm/mesosphere/traefik-forward-auth/0.3.8?modal=values>

<sup>918</sup> <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/traefik-forward-auth/0.3.8/defaults/cm.yaml>

<sup>919</sup> <https://artifacthub.io/packages/helm/vmware-tanzu/velero/2.31.8>

<sup>920</sup> <https://artifacthub.io/packages/helm/vmware-tanzu/velero/2.31.8?modal=values>

<sup>921</sup> <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.0/services/velero/3.3.0/defaults/cm.yaml>

### 11.1.7.3 nvidia-feature-discovery-gpu - error opening libnvidia-ml.so.1

#### D2IQ-93676

After creating a cluster, an incident arose where the nvidia-feature-discovery-gpu-feature-discovery produced a CrashLoopBackOff state, with the following error:

```
gpu-feature-discovery: 2022/10/14 08:52:32 Start running
gpu-feature-discovery: 2022/10/14 08:52:32 Warning: Error removing output file:
failed to remove output file: remove /etc/kubernetes/node-feature-discovery/
features.d/gfd: no such file or directory
gpu-feature-discovery: 2022/10/14 08:52:32 Exiting
gpu-feature-discovery: 2022/10/14 08:52:32 Error: error creating NVML labeler: failed
to initialize NVML: unexpected failure calling nvml.Init: error opening libnvidia-
ml.so.1: libnvidia-ml.so.1: cannot open shared object file: No such file or directory
```

This incident has been resolved by updating the override with appropriate registry details as part of this release.

### 11.1.7.4 Cannot specify custom Packer directory in KIB 1.5.0

#### D2IQ-93657

As part of trying to implement a custom CentOS image adhering to STIG requirements, an issue arose where Packer can not execute because `/tmp` is mounted with `noexec`.

This incident has been resolved by updating packer.json with the appropriate arguments as part of this release.

## 11.1.8 DKP 2.4.0 Known Issues and Limitations

### 11.1.8.1 Known issues and limitations

The following items are known issues with this release.

#### 11.1.8.1.1 Use Static Credentials to Provision an Azure Cluster

Only static credentials can be used when [provisioning an Azure cluster](#) (see page 235).

#### 11.1.8.1.2 EKS Upgrades do not use Declared Kubernetes Version

Upgrading an EKS cluster to DKP version 2.4 does not support specifying a particular Kubernetes patch version when upgrading EKS clusters. Instead, your cluster will be upgraded to the highest currently available patch version available in EKS for the MAJOR.MINOR version of Kubernetes you selected.

### 11.1.8.1.3 EKS Clusters use CSI-based EBS Drivers

EKS has disabled support for in-tree EBS volume provisioning in favor of [CSI Volumes](#)<sup>922</sup>. See the dropdown labeled “Additional Considerations for EKS” in the [\(2.4\) DKP Enterprise Upgrade | Upgrade the Core Addons](#) (see page 0) page for more details on what steps you must take when upgrading.

### 11.1.8.1.4 Intermittent Error Status when Creating EKS Clusters in the UI

When provisioning an EKS cluster through the UI, you may receive a brief error state because the EKS cluster may 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 will resolve, 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.

### 11.1.8.1.5 Installation and upgrade issue in pre-provisioned environments

DKP made the transition from Minio to Rook Ceph for cluster storage. An issue with Rook Ceph’s deployment prevents pre-provisioned environments from installing and upgrading to this DKP version. To solve this issue, you must set up 40 GB of raw storage for your worker nodes and customize your Rook Ceph installation as indicated in [Install Kommander in a Pre-provisioned Environment](#) (see page 515) or [Upgrade Kommander in a Pre-provisioned Environment](#) (see page 1031).

### 11.1.8.1.6 Cluster Roles' information not available in the Projects section of the UI

Selecting a role of the *Cluster Role* type in the **Projects > Roles** section of a workspace displays an error message in the UI.

#### 11.1.8.1.6.1 Workaround

You can still access a *Cluster Role*’s description and configuration from the UI. Take the following alternative path to *view* or *edit* the desired role:

1. Select your workspace from the top navigation bar.
2. Select **Administration > Access Control** from the sidebar.
3. A table appears that lists roles from all *Projects* in the selected workspace.
4. Select the **Name** or **ID** of the *Cluster Role* you want to access. A page opens that contains more information and configuration options for the role.

---

<sup>922</sup> <https://kubernetes.io/blog/2019/12/09/kubernetes-1-17-feature-csi-migration-beta/>



- To access a *Cluster Role*'s description and configuration page, ensure your user has sufficient cluster *view* or *edit* rights. You will not be able to select roles for which you do not have access rights.

### 11.1.8.1.7 Resolve issues with failed HelmReleases

An [issue with the Flux helm-controller](https://github.com/fluxcd/helm-controller/issues/149)<sup>923</sup> can cause HelmReleases to fail with the error message *Helm upgrade failed: another operation (install/upgrade/rollback) is in progress*. This can happen when the helm-controller is restarted while a HelmRelease is still upgrading, or installing.

#### 11.1.8.1.7.1 Workaround

To ensure the HelmRelease error was caused by the helm-controller restarting, first try to suspend/resume the HelmRelease:

```
kubectl -n <namespace> patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": true}]'
```

```
kubectl -n <namespace> patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": false}]'
```

This might resolve the issue. If not, continue with the following steps:

You should see the HelmRelease attempting to reconcile, and then it either succeeds (with status: *Release reconciliation succeeded*) or it fails with the same error as before.

If the HelmRelease is still in the failed state, it is likely related to the helm-controller restarting. For example, if the 'reloader' HelmRelease is the one that is stuck.

To resolve the issue, follow these steps:

- List secrets containing the affected HelmRelease name:

```
kubectl get secrets -n ${NAMESPACE} | grep reloader
```

The output should look like this:

```
kommander-reloader-reloader-token-9qd8b kubernetes.io/
service-account-token 3 171m
sh.helm.release.v1.kommander-reloader.v1
release.v1 1 171m
sh.helm.release.v1.kommander-reloader.v2
release.v1 1 117m
```

<sup>923</sup> <https://github.com/fluxcd/helm-controller/issues/149>

In this example, `sh.helm.release.v1.kommander-reloader.v2` is the most recent revision.

2. Find and delete the most recent revision secret, for example,

```
sh.helm.release.v1.*.<revision>:
```

```
kubectl delete secret -n <namespace> <most recent helm revision secret name>
```

3. Suspend and resume the HelmRelease to trigger a reconciliation:

```
kubectl -n <namespace> patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": true}]'
kubectl -n <namespace> patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": false}]'
```

You should see the HelmRelease is reconciled and eventually the upgrade and install succeed.

#### 11.1.8.1.7.2 Calico not updated during DKP upgrade on Flatcar

When upgrading a DKP cluster running on Flatcar OS, you may find that after the upgrade the Calico services were not updated. This occurs because the upgrade procedure is not correctly updating the Flatcar specific CNI ClusterResourceSet(CRS). This issue only impacts the Calico CRS.

Follow these steps to manually correct this issue:

1. Update the ConfigMap as follows:

```
cat <<EOF | kubectl apply -f -
apiVersion: v1
data:
 custom-resources.yaml: |+
 # This section includes base Calico installation configuration.
 # For more information, see: <https://docs.projectcalico.org/reference/
 installation/api>
 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: 192.168.0.0/16
 encapsulation: IPIP
 natOutgoing: Enabled
 nodeSelector: all()
 bgp: Enabled
```

```

nodeAddressAutodetectionV4:
 firstFound: true
 # FlexVolume path must be mounted under /opt on flatcar/coreos systems
 flexVolumePath: /opt/libexec/kubernetes/kubelet-plugins/volume/exec/
kind: ConfigMap
metadata:
 name: calico-cni-installation- $\$$ CLUSTER_NAME
EOF

```

2. Run these commands:

```
kubectl edit clusterresourceset calico-cni-installation- $\$$ CLUSTER_NAME
```

and update

```
spec.clusterSelector.matchLabels.konvoy.d2iq.io/osHint to konvoy.d2iq.io/osHint:
flatcar
```

#### 11.1.8.1.8 kube-oidc-proxy not ready after upgrade

If you installed or attached a cluster in 2.1, `kube-oidc-proxy` is not available after upgrading to 2.3 and 2.4. This application is required to access the Kubernetes API (with `kubectl`) using SSO. For affected customers, there are issues with the authentication via `kubectl`.

To make the application available, run the following command on each cluster that was installed, created or attached in 2.1, and is now on DKP version 2.4. Replace `<namespace>` with each cluster's workspace namespace:

```
kubectl -n <namespace> patch appdeployment kube-oidc-proxy --type=json -p
' [{"op": "remove", "path": "/spec/configOverrides"}]'
```

#### 11.1.8.1.9 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 created using OS images. The flags work by resizing the primary disk of the VM. When the VM boots, the root file system is expanded 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, please see [Create a Base OS Image \(see page 430\)](#).

### 11.1.9 DKP 2.4.0 Deprecations

The following items are deprecated or removed from this version of DKP.

### 11.1.9.1 MinIO

In previous versions of DKP, we deployed MinIO with velero, grafana-loki, and project-grafana loki, operating inside the same cluster. Beginning with DKP 2.4.0, MinIO is no longer deployed by default and will no longer be supported. However, we will not remove any MinIO instances that you have running on any version of DKP and were installed independently of the Kommander installation.

As a replacement for MinIO, we are integrating [Rook Ceph](#)<sup>924</sup> in this and future releases.

It is important to note that when you upgrade from `2.3.x` → `2.4.x`, Ceph is automatically installed if any of the following are true:

- `minio-operator` was installed (used by logging stack).
- `velero` was installed (used for backups).

However, if you did not have `minio-operator` installed in `2.3.x` and instead had configured `velero` to work with an [external cloud storage](#) (see page 731) such as Amazon S3 or Azure Blob Storage, then you do not need `rook-ceph` and `rook-ceph-cluster` post upgrade.

In order to explicitly disable installing these apps during upgrade, specify the following in the command line, then run the command:

```
dkp upgrade kommander --disable-appdeployments rook-ceph,rook-ceph-cluster
```

**NOTE:** The command above fails if `minio-operator` is installed as you cannot disable `ceph` installation if `minio-operator` is installed.


## 11.1.10 Kubernetes major updates and deprecations

### 11.1.10.1 LegacyServiceAccountTokenNoAutoGeneration Feature Gate

With Kubernetes 1.24 the **LegacyServiceAccountTokenNoAutoGeneration** feature gate is now enabled by default. This affects the attachment of existing Kubernetes clusters, for which you now must manually create the token. We include instructions on how to do this in the [Attach a GKE cluster](#) (see page 713) and [Attach an Amazon EKS cluster](#) (see page 708) pages.

Furthermore, upgrading to 1.24 may affect some other components of your environment:


<sup>924</sup> <https://docs.ceph.com/en/latest/radosgw/s3/>

 Before upgrading, we **strongly recommend** verifying your current setup against the information on this page and reading more about Kubernetes' new features in our blog <https://eng.d2iq.com/blog/service-account-tokens-in-kubernetes-v1.24/>, and in [Kubernetes urgent upgrade notes](#)<sup>925</sup>.

### 11.1.10.2 Control Plane Node Label and Taint

A new label has been implemented called `node-role.kubernetes.io/control-plane` to be added to control plane nodes.

In an effort to migrate Kubernetes away from the usage of the word `master` in labels and taints, for new clusters, the label `node-role.kubernetes.io/master` will no longer be added to control plane nodes. The new label `node-role.kubernetes.io/control-plane` will be added. For upgraded clusters, the label `node-role.kubernetes.io/master` will be removed from existing control-plane nodes. For new clusters, both the old taint `node-role.kubernetes.io/master:NoSchedule` and new taint `node-role.kubernetes.io/control-plane:NoSchedule` will be added to control plane nodes. For clusters that are being upgraded, the command will add the new taint `node-role.kubernetes.io/control-plane:NoSchedule` to existing control plane nodes.

 Before upgrading, modify your workloads that are currently relying on the label `node-role.kubernetes.io/master` to use `node-role.kubernetes.io/control-plane` instead for its `nodeSelector` and `affinity`. For any workloads that have a toleration for `node-role.kubernetes.io/master`, add a toleration for `node-role.kubernetes.io/control-plane`.

See the [Kubernetes urgent upgrade notes](#)<sup>926</sup> for a full list of changes.


## 11.2 DKP 2.4.1 Release Notes

**DKP® version 2.4.1 was released on May 22, 2023.**

[Download DKP](#)

<sup>925</sup> <https://github.com/kubernetes/kubernetes/blob/master/CHANGELOG/CHANGELOG-1.24.md#urgent-upgrade-notes>

<sup>926</sup> <https://github.com/kubernetes/kubernetes/blob/master/CHANGELOG/CHANGELOG-1.24.md#urgent-upgrade-notes>

-  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.com](mailto:sales@d2iq.com)<sup>927</sup> before attempting to download or install DKP.

## 11.2.1 Release Summary

Welcome to D2iQ Kubernetes Platform (DKP) 2.4.1! This release provides fixes to reported issues, integrates changes from previous releases, and maintains compatibility and support for other packages used in DKP. Below, you will find information about the following:

- [DKP 2.4.1 Enhancements](#) (see page 1079)
- [DKP 2.4.1 Fixes and Updates](#) (see page 1079)
- [DKP 2.4.1 Customer Incidents](#) (see page 1079)
- [DKP 2.4.1 Components and Applications](#) (see page 1081)
- [DKP 2.4.1 Supported Kubernetes Versions](#) (see page 1089)
- [DKP 2.4.1 Known Issues and Limitations](#) (see page 1090)

## 11.2.2 Additional resources

- For more information about working with native Kubernetes, see the [Kubernetes documentation](#)<sup>928</sup>.
- For a full list of attributed 3rd party software, see <http://d2iq.com/legal/3rd>.

## 11.2.3 DKP 2.4.1 Enhancements

The following improvements are included in this release.

### 11.2.3.1 Improved Installation Times for Kommander

This version of DKP includes an updated application deployment order for the Kommander component, which significantly decreases installation times.

### 11.2.3.2 External Load Balancer Support

If you want to use a [non-DKP load balancer](#) (see page 806) for external traffic, you can now [Install Kommander with an External Load Balancer](#) (see page 501).

---

<sup>927</sup> <mailto:sales@d2iq.com>

<sup>928</sup> <https://kubernetes.io/docs/home/>

## 11.2.4 DKP 2.4.1 Fixes and Updates

The following updates and fixes are included in this release.

### 11.2.4.1 kube-oidc-proxy issue was resolved

#### D2IQ-94629

If you installed or attached a cluster in 2.1, `kube-oidc-proxy` was not available after upgrading to 2.3. This prevented authentication via `kubectl` using SSO. This issue has been resolved.

## 11.2.5 DKP 2.4.1 Customer Incidents

The following issues are corrected or resolved in this release:

### 11.2.5.1 Improved workflow for Deploying a Lower Version of Kubernetes

#### D2IQ-95873

The DKP documentation now provides a workflow to deploy a lower version of Kubernetes.

### 11.2.5.2 KIB 1.24.x Fails to Create Ubuntu Images

#### D2IQ-95429

Attempting to create an Ubuntu 1804 or 2004 image for GCP with KIB 1.24.2 or 1.24.3 would fail with an error.

### 11.2.5.3 Traefik entry point TCP 9090 is tagged Velero-minio

#### D2IQ-95330

The entry point for Rook-ceph Object storage was tagged from previous versions of DKP when we previously used the Minlo application. The entry point is updated to Rook-Ceph.

### 11.2.5.4 Upgrading Azure Breaks Volume Attachments for some Pods

#### D2IQ-95191

When upgrading a DKP cluster on Azure from 2.3.0 to 2.4.0, some pods failed to start due to a bug in the Azure CSI driver.

### 11.2.5.5 Kommander Install and Upgrade fail on Pre-provisioned

#### D2IQ-94959

Upgrades and installations in 2.4 fail for pre-provisioned environments, requiring a workaround to complete the process. This is corrected and the workaround is no longer necessary.

### 11.2.5.6 Logging-operator fluentd Error

D2iQ-94914

The fluentd logging-operator was failing to capture logs from applications. This is corrected and the logs are now captured.

### 11.2.5.7 Calico Not Updated during DKP Upgrade on Flatcar

D2iQ-94862

When upgrading a cluster using Flatcar OS, the Calico resources were not properly upgraded due to a missing `osHint` label on some cluster resources. The appropriate `osHint` labels are now applied on cluster creation and upgrade.

### 11.2.5.8 Cannot Connect Cluster with Kommander installed to an existing Management Cluster

D2iQ-94859

Attempting to connect a cluster with Kommander already installed resulted in an error, as this was not a supported action. The [documentation is updated \(see page 678\)](#) to support this action.

### 11.2.5.9 Sizing Guidelines for Loki and Logging Operator

D2iQ-90805

The documentation was missing sizing guidelines for these two items.

### 11.2.5.10 License counts Control Plane Cores post 2.4 upgrade

D2iQ-95188

In 2.4.0, the licensing page was erroneously counting Control plane cores. The problem is corrected in this release.

### 11.2.5.11 Kommander Upgrade Fails with grafana-loki

D2iQ-94944

Rook-Ceph, introduced in 2.4.0, has additional storage requirements for pre-provisioned environments. [Please review the storage requirements \(see page 1031\)](#) before upgrading to 2.4.x.



## 11.2.6 DKP 2.4.1 Components and Applications

The following are component and application versions for DKP 2.4.1.

### 11.2.6.1 Components

Component Name	Version
Cluster API Core (CAPI)	1.2.4-d2iq.0
Cluster API AWS Infrastructure Provider (CAPA)	1.5.1
Cluster API Google Cloud Infrastructure Provider (CAPG)	1.1.0
Cluster API Pre-provisioned Infrastructure Provider (CAPPP)	0.12.3
Cluster API vSphere Infrastructure Provider (CAPV)	1.3.1
Cluster API Azure Infrastructure Provider (CAPZ)	1.5.2
Konvoy Image Builder (KIB)	<a href="https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.24.8">1.24.8</a> <sup>929</sup>
containerd	1.4.13
etcd	3.4.13

---

<sup>929</sup> <https://github.com/mesosphere/konvoy-image-builder/releases/tag/v1.24.8>

## 11.2.6.2 Applications

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Centralized Grafana	centralized-grafana	40.0.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 40.0.0</a><sup>930</sup></li> <li>• prometheus-operator: 0.59.1</li> <li>• grafana: 9.1.6</li> </ul>	<a href="#">Link</a> <sup>931</sup>	<a href="#">Link</a> <sup>932</sup>
Centralized Kubecost	centralized-kubecost	0.28.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.28.0</a><sup>933</sup></li> <li>• kubecost: 1.97.0</li> </ul>	<a href="#">Link</a> <sup>934</sup>	<a href="#">Link</a> <sup>935</sup>
Cert Manager	cert-manager	1.9.1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 1.9.1</a><sup>936</sup></li> <li>• cert-manager: 1.9.1</li> </ul>	<a href="#">Link</a> <sup>937</sup>	<a href="#">Link</a> <sup>938</sup>
Chartmuseum	chartmuseum	3.9.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 3.9.0</a><sup>939</sup></li> <li>• chartmuseum: 3.9.0</li> </ul>	<a href="#">Link</a> <sup>940</sup>	<a href="#">Link</a> <sup>941</sup>
Dex	dex	2.10.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 2.10.0</a><sup>942</sup></li> <li>• dex: 2.31.0</li> </ul>	<a href="#">Link</a> <sup>943</sup>	<a href="#">Link</a> <sup>944</sup>

930 <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/40.0.0>

931 <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/40.0.0?modal=values>

932 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/centralized-grafana/40.0.0/defaults/cm.yaml>

933 <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.28.0>

934 <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.28.0?modal=values>

935 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/centralized-kubecost/0.28.0/defaults/cm.yaml>

936 <https://artifacthub.io/packages/helm/cert-manager/cert-manager/1.9.1>

937 <https://artifacthub.io/packages/helm/cert-manager/cert-manager/1.9.1?modal=values>

938 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/cert-manager/1.9.1/defaults/cm.yaml>

939 <https://artifacthub.io/packages/helm/chartmuseum/chartmuseum/3.9.0>

940 <https://artifacthub.io/packages/helm/chartmuseum/chartmuseum/3.9.0?modal=values>

941 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/chartmuseum/3.9.0/defaults/cm.yaml>

942 <https://artifacthub.io/packages/helm/mesosphere-stable/dex/2.10.0>

943 <https://artifacthub.io/packages/helm/mesosphere-stable/dex/2.10.0?modal=values>

944 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/dex/2.10.0/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Dex K8s Authenticator	dex-k8s-authenticator	1.2.14	<ul style="list-style-type: none"> <li>chart: 1.2.14<sup>945</sup></li> <li>dex-k8s-authenticator: 1.2.4</li> </ul>	<a href="#">Link</a> <sup>946</sup>	<a href="#">Link</a> <sup>947</sup>
DKP Insights Management	dkp-insights-management	0.3.1	<ul style="list-style-type: none"> <li>chart: 0.3.1</li> <li>dkp-insights-management: 0.3.1</li> </ul>	N/A	<a href="#">Link</a> <sup>948</sup>
External DNS	external-dns	6.10.2	<ul style="list-style-type: none"> <li>chart: 6.10.2<sup>949</sup></li> <li>external-dns: 0.12.2</li> </ul>	<a href="#">Link</a> <sup>950</sup>	<a href="#">Link</a> <sup>951</sup>
Fluent Bit	fluent-bit	0.20.9	<ul style="list-style-type: none"> <li>chart: 0.20.9<sup>952</sup></li> <li>fluent-bit: 1.9.9</li> </ul>	<a href="#">Link</a> <sup>953</sup>	<a href="#">Link</a> <sup>954</sup>
Gatekeeper	gatekeeper	3.10.0	<ul style="list-style-type: none"> <li>chart: 3.10.0<sup>955</sup></li> <li>gatekeeper: 3.10.0</li> </ul>	<a href="#">Link</a> <sup>956</sup>	<a href="#">Link</a> <sup>957</sup>
Gitea	gitea	6.0.1	<ul style="list-style-type: none"> <li>chart: 6.0.1<sup>958</sup></li> <li>gitea: 1.17.2</li> </ul>	<a href="#">Link</a> <sup>959</sup>	<a href="#">Link</a> <sup>960</sup>

945 <https://artifacthub.io/packages/helm/mesosphere/dex-k8s-authenticator/1.2.14>

946 <https://artifacthub.io/packages/helm/mesosphere/dex-k8s-authenticator/1.2.14?modal=values>

947 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/dex-k8s-authenticator/1.2.14/defaults/cm.yaml>

948 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/dkp-insights-management/0.3.1/defaults/cm.yaml>

949 <https://artifacthub.io/packages/helm/bitnami/external-dns/6.10.2>

950 <https://artifacthub.io/packages/helm/bitnami/external-dns/6.10.2?modal=values>

951 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/external-dns/6.10.2/defaults/cm.yaml>

952 <https://artifacthub.io/packages/helm/fluent/fluent-bit/0.20.9>

953 <https://artifacthub.io/packages/helm/fluent/fluent-bit/0.20.9?modal=values>

954 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/fluent-bit/0.20.9/defaults/cm.yaml>

955 <https://artifacthub.io/packages/helm/gatekeeper/gatekeeper/3.10.0>

956 <https://artifacthub.io/packages/helm/gatekeeper/gatekeeper/3.10.0?modal=values>

957 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/gatekeeper/3.10.0/defaults/cm.yaml>

958 <https://artifacthub.io/packages/helm/gitea/gitea/6.0.1>

959 <https://artifacthub.io/packages/helm/gitea/gitea/6.0.1?modal=values>

960 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/gitea/6.0.1/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Grafana Logging	grafana-logging	6.3 8.1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 6.38.1</a><sup>961</sup></li> <li>• grafana: 9.1.5</li> </ul>	<a href="#">Link</a> <sup>962</sup>	<a href="#">Link</a> <sup>963</sup>
Grafana Loki	grafana-loki	0.4 8.7	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.48.4</a><sup>964</sup></li> <li>• loki: 2.5.0</li> </ul>	<a href="#">Link</a> <sup>965</sup>	<a href="#">Link</a> <sup>966</sup>
Istio	istio	1.1 5.3	<ul style="list-style-type: none"> <li>• <a href="#">chart: 1.15.3</a><sup>967</sup></li> <li>• istio: 1.15.3</li> </ul>	<a href="#">Link</a> <sup>968</sup>	<a href="#">Link</a> <sup>969</sup>
Jaeger	jaeger	2.3 6.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 2.36.0</a><sup>970</sup></li> <li>• jaeger: 1.38.0</li> </ul>	<a href="#">Link</a> <sup>971</sup>	<a href="#">Link</a> <sup>972</sup>
Karma	karma	2.0. 1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 2.0.1</a><sup>973</sup></li> <li>• karma: 0.70</li> </ul>	<a href="#">Link</a> <sup>974</sup>	<a href="#">Link</a> <sup>975</sup>
Kiali	kiali	1.5 7.1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 1.57.1</a><sup>976</sup></li> <li>• kiali: 1.57.1</li> </ul>	<a href="#">Link</a> <sup>977</sup>	<a href="#">Link</a> <sup>978</sup>

961 <https://artifacthub.io/packages/helm/grafana/grafana/6.38.1>

962 <https://artifacthub.io/packages/helm/grafana/grafana/6.38.1?modal=values>

963 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/grafana-logging/6.38.1/defaults/cm.yaml>

964 <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.48.4>

965 <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.48.4?modal=values>

966 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/grafana-loki/0.48.7/defaults/cm.yaml>

967 <https://artifacthub.io/packages/helm/mesosphere/istio/1.15.3>

968 <https://artifacthub.io/packages/helm/mesosphere/istio/1.15.3?modal=values>

969 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/istio/1.15.3/defaults/cm.yaml>

970 <https://artifacthub.io/packages/helm/jaegertracing/jaeger-operator/2.36.0>

971 <https://artifacthub.io/packages/helm/jaegertracing/jaeger-operator/2.36.0?modal=values>

972 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/jaeger/2.36.0/defaults/cm.yaml>

973 <https://artifacthub.io/packages/helm/mesosphere-stable/karma/2.0.1>

974 <https://artifacthub.io/packages/helm/mesosphere-stable/karma/2.0.1?modal=values>

975 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/karma/2.0.1/defaults/cm.yaml>

976 <https://artifacthub.io/packages/helm/kiali/kiali-operator/1.57.1>

977 <https://artifacthub.io/packages/helm/kiali/kiali-operator/1.57.1?modal=values>

978 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/kiali/1.57.1/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Knative	knative	0.4.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.4.0</a><sup>979</sup></li> <li>• knative: 0.22.3</li> </ul>	<a href="#">Link</a> <sup>980</sup>	<a href="#">Link</a> <sup>981</sup>
Flux	kommander-flux	0.37.0	<ul style="list-style-type: none"> <li>• chart: N/A</li> <li>• flux: 0.37.0</li> </ul>	N/A	N/A
Kube OIDC Proxy	kube-oidc-proxy	0.3.2	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.3.1</a><sup>982</sup></li> <li>• kube-oidc-proxy: 0.3.0</li> </ul>	<a href="#">Link</a> <sup>983</sup>	<a href="#">Link</a> <sup>984</sup>
Kube Prometheus Stack	kube-prometheus-stack	40.0.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 40.0.0</a><sup>985</sup></li> <li>• prometheus-operator: 0.59.1</li> <li>• grafana: 9.1.6</li> <li>• prometheus: 2.38.0</li> <li>• prometheus-alertmanager: 0.24.0</li> </ul>	<a href="#">Link</a> <sup>986</sup>	<a href="#">Link</a> <sup>987</sup>
Kubecost	kubecost	0.28.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.28.0</a><sup>988</sup></li> <li>• kubecost: 1.97.0</li> </ul>	<a href="#">Link</a> <sup>989</sup>	<a href="#">Link</a> <sup>990</sup>
Kubefed	kubefed	0.10.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.10.0</a><sup>991</sup></li> <li>• kubefed: 0.10.0</li> </ul>	<a href="#">Link</a> <sup>992</sup>	<a href="#">Link</a> <sup>993</sup>

979 <https://artifacthub.io/packages/helm/mesosphere/knative/0.4.0>

980 <https://artifacthub.io/packages/helm/mesosphere/knative/0.4.0?modal=values>

981 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/knative/0.4.0/defaults/cm.yaml>

982 <https://artifacthub.io/packages/helm/mesosphere/kube-oidc-proxy/0.3.1>

983 <https://artifacthub.io/packages/helm/mesosphere/kube-oidc-proxy/0.3.1?modal=values>

984 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/kube-oidc-proxy/0.3.2/defaults/cm.yaml>

985 <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/40.0.0>

986 <https://artifacthub.io/packages/helm/mesosphere/kube-prometheus-stack/40.0.0?modal=values>

987 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/kube-prometheus-stack/40.0.0/defaults/cm.yaml>

988 <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.28.0>

989 <https://artifacthub.io/packages/helm/mesosphere-stable/kubecost/0.28.0?modal=values>

990 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/kubecost/0.28.0/defaults/cm.yaml>

991 <https://artifacthub.io/packages/helm/kubefed/kubefed/0.10.0>

992 <https://artifacthub.io/packages/helm/kubefed/kubefed/0.10.0?modal=values>

993 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/kubefed/0.10.0/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Kubernetes Dashboard	kubernetes-dashboard	5.11.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 5.11.0</a><sup>994</sup></li> <li>• kubernetes-dashboard: 2.7.0</li> </ul>	<a href="#">Link</a> <sup>995</sup>	<a href="#">Link</a> <sup>996</sup>
Kubetunnel	kubetunnel	0.0.15	<ul style="list-style-type: none"> <li>• chart: 0.0.15</li> <li>• kubetunnel: 0.0.15</li> </ul>	N/A	<a href="#">Link</a> <sup>997</sup>
Logging Operator	logging-operator	3.17.9	<ul style="list-style-type: none"> <li>• <a href="#">chart: 3.17.9</a><sup>998</sup></li> <li>• logging-operator: 3.17.9</li> <li>• logging-operator-logging: 3.17.9</li> </ul>	<a href="#">Link</a> <sup>999</sup>	<a href="#">Link</a> <sup>1000</sup>
NFS Server Provisioner	nfs-server-provisioner	0.6.0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.6.0</a><sup>1001</sup></li> <li>• nfs-server-provisioner: 2.3.0</li> </ul>	<a href="#">Link</a> <sup>1002</sup>	<a href="#">Link</a> <sup>1003</sup>
NVIDIA GPU Operator	nvidia-gpu-operator	1.11.1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 1.11.1</a><sup>1004</sup></li> <li>• nvidia-gpu-operator: 1.11.1</li> </ul>	<a href="#">Link</a> <sup>1005</sup>	<a href="#">Link</a> <sup>1006</sup>

994 <https://artifacthub.io/packages/helm/k8s-dashboard/kubernetes-dashboard/5.11.0>

995 <https://artifacthub.io/packages/helm/k8s-dashboard/kubernetes-dashboard/5.11.0?modal=values>

996 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/kubernetes-dashboard/5.11.0/defaults/cm.yaml>

997 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/kubetunnel/0.0.15/defaults/cm.yaml>

998 <https://artifacthub.io/packages/helm/banzaicloud-stable/logging-operator/3.17.9>

999 <https://artifacthub.io/packages/helm/banzaicloud-stable/logging-operator/3.17.9?modal=values>

1000 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/logging-operator/3.17.9/defaults/cm.yaml>

1001 <https://artifacthub.io/packages/helm/mesosphere/nfs-server-provisioner/0.6.0>

1002 <https://artifacthub.io/packages/helm/mesosphere/nfs-server-provisioner/0.6.0?modal=values>

1003 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/nfs-server-provisioner/0.6.0/defaults/cm.yaml>

1004 <https://docs.nvidia.com/datacenter/cloud-native/gpu-operator/getting-started.html#chart-customization-options>

1005 <https://raw.githubusercontent.com/NVIDIA/gpu-operator/v1.11.1/deployments/gpu-operator/values.yaml>

1006 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/nvidia-gpu-operator/1.11.1/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Grafana (project)	project-grafana-logging	6.3 8.1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 6.38.1</a><sup>1007</sup></li> <li>• grafana: 9.1.5</li> </ul>	<a href="#">Link</a> <sup>1008</sup>	<a href="#">Link</a> <sup>1009</sup>
Grafana Loki (project)	project-grafana-loki	0.4 8.7	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.48.4</a><sup>1010</sup></li> <li>• loki: 2.5.0</li> </ul>	<a href="#">Link</a> <sup>1011</sup>	<a href="#">Link</a> <sup>1012</sup>
Prometheus Adapter	prometheus-adapter	3.4. 0	<ul style="list-style-type: none"> <li>• <a href="#">chart: 3.4.0</a><sup>1013</sup></li> <li>• prometheus-adapter: 0.10.0</li> </ul>	<a href="#">Link</a> <sup>1014</sup>	<a href="#">Link</a> <sup>1015</sup>
Reloader	reloader	0.0. 121	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.0.121</a><sup>1016</sup></li> <li>• reloader: 0.0.121</li> </ul>	<a href="#">Link</a> <sup>1017</sup>	<a href="#">Link</a> <sup>1018</sup>
Rook Ceph	rook-ceph	1.1 0.8	<ul style="list-style-type: none"> <li>• <a href="#">chart: 1.10.8</a><sup>1019</sup></li> <li>• rook-ceph: 1.10.8</li> </ul>	<a href="#">Link</a> <sup>1020</sup>	<a href="#">Link</a> <sup>1021</sup>

<sup>1007</sup> <https://artifacthub.io/packages/helm/grafana/grafana/6.38.1>

<sup>1008</sup> <https://artifacthub.io/packages/helm/grafana/grafana/6.38.1?modal=values>

<sup>1009</sup> <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/project-grafana-logging/6.38.1/defaults/cm.yaml>

<sup>1010</sup> <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.48.4>

<sup>1011</sup> <https://artifacthub.io/packages/helm/grafana/loki-distributed/0.48.4?modal=values>

<sup>1012</sup> <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/project-grafana-loki/0.48.7/defaults/cm.yaml>

<sup>1013</sup> <https://artifacthub.io/packages/helm/prometheus-community/prometheus-adapter/3.4.0>

<sup>1014</sup> <https://artifacthub.io/packages/helm/prometheus-community/prometheus-adapter/3.4.0?modal=values>

<sup>1015</sup> <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/prometheus-adapter/3.4.0/defaults/cm.yaml>

<sup>1016</sup> <https://github.com/stakater/Reloader/tree/v0.0.121/README.md#helm-charts>

<sup>1017</sup> <https://artifacthub.io/packages/helm/stakater/reloader/0.0.121?modal=values>

<sup>1018</sup> <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/reloader/0.0.121/defaults/cm.yaml>

<sup>1019</sup> <https://github.com/rook/rook/tree/v1.10.8/Documentation/Helm-Charts/operator-chart.md>

<sup>1020</sup> <https://raw.githubusercontent.com/rook/rook/v1.10.8/deploy/charts/rook-ceph/values.yaml>

<sup>1021</sup> <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/rook-ceph/1.10.8/defaults/cm.yaml>

Common Application Name	APP ID	Version	Component Versions	Helm Values	DKP Values
Rook Ceph Cluster	rook-ceph-cluster	1.10.8	<ul style="list-style-type: none"> <li>• <a href="#">chart: 1.10.8</a><sup>1022</sup></li> <li>• rook-ceph: 1.10.8</li> <li>• rook-ceph-cluster: 17.2.5</li> </ul>	<a href="#">Link</a> <sup>1023</sup>	<a href="#">Link</a> <sup>1024</sup>
Thanos	thanos	0.4.8	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.4.6</a><sup>1025</sup></li> <li>• thanos: 0.17.1</li> </ul>	<a href="#">Link</a> <sup>1026</sup>	<a href="#">Link</a> <sup>1027</sup>
Traefik	traefik	10.30.1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 10.30.1</a><sup>1028</sup></li> <li>• traefik: 2.8.7</li> </ul>	<a href="#">Link</a> <sup>1029</sup>	<a href="#">Link</a> <sup>1030</sup>
Traefik ForwardAuth	traefik-forward-auth	0.3.8	<ul style="list-style-type: none"> <li>• <a href="#">chart: 0.3.8</a><sup>1031</sup></li> <li>• traefik-forward-auth: 3.1.0</li> </ul>	<a href="#">Link</a> <sup>1032</sup>	<a href="#">Link</a> <sup>1033</sup>
Velero	velero	3.3.1	<ul style="list-style-type: none"> <li>• <a href="#">chart: 2.31.8</a><sup>1034</sup></li> <li>• velero: 1.9.2</li> </ul>	<a href="#">Link</a> <sup>1035</sup>	<a href="#">Link</a> <sup>1036</sup>

1022 <https://github.com/rook/rook/tree/v1.10.8/Documentation/Helm-Charts/ceph-cluster-chart.md>

1023 <https://raw.githubusercontent.com/rook/rook/v1.10.8/deploy/charts/rook-ceph-cluster/values.yaml>

1024 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/rook-ceph-cluster/1.10.8/defaults/cm.yaml>

1025 <https://artifacthub.io/packages/helm/bitnami/thanos/0.4.6>

1026 <https://artifacthub.io/packages/helm/bitnami/thanos/0.4.6?modal=values>

1027 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/thanos/0.4.8/defaults/cm.yaml>

1028 <https://artifacthub.io/packages/helm/traefik/traefik/10.30.1>

1029 <https://artifacthub.io/packages/helm/traefik/traefik/10.30.1?modal=values>

1030 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/traefik/10.30.1/defaults/cm.yaml>

1031 <https://artifacthub.io/packages/helm/mesosphere/traefik-forward-auth/0.3.8>

1032 <https://artifacthub.io/packages/helm/mesosphere/traefik-forward-auth/0.3.8?modal=values>

1033 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/traefik-forward-auth/0.3.8/defaults/cm.yaml>

1034 <https://artifacthub.io/packages/helm/vmware-tanzu/velero/2.31.8>

1035 <https://artifacthub.io/packages/helm/vmware-tanzu/velero/2.31.8?modal=values>

1036 <https://raw.githubusercontent.com/mesosphere/kommander-applications/v2.4.1/services/velero/3.3.1/defaults/cm.yaml>



## 11.2.7 DKP 2.4.1 Supported Kubernetes Versions

### 11.2.7.1 Deploying Clusters Versions

The newest and oldest Kubernetes versions used with DKP must be within one minor version.

- newest supported Kubernetes version is at **1.24.6 for deploying DKP for all providers except EKS.**
- Kubernetes versions are supported at **1.23.x for deploying DKP on EKS.**

**■** DKP 2.4 comes with support for Kubernetes 1.24.6, enabling you to benefit from the latest features and security fixes in upstream Kubernetes. This release comes with approximately 46 enhancements. To read more about major features in this release, visit <https://kubernetes.io/blog/2022/05/03/kubernetes-1-24-release-announcement/>.

### 11.2.7.2 Attaching Clusters Versions

**■** When attempting to deploy a cluster with a lower Kubernetes version than the DKP default, you need to build and use an AMI with that lower version of Kubernetes. See [Konvoy Image Builder \(see page 411\)](#) for documentation on building images. Failure to do this could result in a failure to deploy error.

DKP 2.4.1 supports attaching clusters with the following Kubernetes versions:

Product	Compatible Kubernetes Versions
EKS	1.23.x
AKS	1.24.x
GKE	1.24.x

## 11.2.8 DKP 2.4.1 Known Issues and Limitations

### 11.2.8.1 Known issues and limitations

The following items are known issues with this release.

#### 11.2.8.1.1 Use Static Credentials to Provision an Azure Cluster

Only static credentials can be used when [provisioning an Azure cluster](#) (see page 235).

#### 11.2.8.1.2 Workaround for missing `allowVolumeExpansion` on EBS StorageClass

To allow for volume expansion, perform the following steps:

1. Add the missing `allowVolumeExpansion` on EBS StorageClass by running command below to edit this storageclass YAML and add in `allowVolumeExpansion: true`:

```
kubectl patch storageclass ebs-sc --patch '{"allowVolumeExpansion": true}'
```

2. Enable CSI volume expansion by setting `allowVolumeExpansion` for:
  - \* [EBS CSI driver](#)<sup>1037</sup>
  - \* [AzureDisk CSI driver](#)<sup>1038</sup>
  - \* [GCP CSI driver](#)<sup>1039</sup>
  - \* [vSphere CSI driver](#)<sup>1040</sup>
3. Add permissions to the following areas for respective cloud provider:
  - AWS
  - Azure
  - GCP
  - vSphere

<sup>1037</sup> <https://github.com/kubernetes-sigs/aws-ebs-csi-driver/blob/8803fbf2d46de300891ac65b8b449783fd79ce02/examples/kubernetes/resizing/manifests/storageclass.yaml#L6>

<sup>1038</sup> <https://github.com/kubernetes-sigs/azuredisk-csi-driver/blob/master/deploy/example/resize/README.md#volume-online-resize>

<sup>1039</sup> <https://github.com/kubernetes-sigs/gcp-compute-persistent-disk-csi-driver/blob/master/docs/kubernetes/user-guides/resize.md#kubernetes-resize-user-guide>

<sup>1040</sup> <https://github.com/kubernetes-sigs/vsphere-csi-driver/blob/51dfa2a52b07eea1569c9a54cbd19461c2eaea0f/example/vanilla-k8s-RWO-file-system-volumes/example-sc-volume-expansion.yaml#L8>

### 11.2.8.1.3 EKS Upgrades do not use Declared Kubernetes Version

Upgrading an EKS cluster to DKP version 2.4 does not support specifying a particular Kubernetes patch version when upgrading EKS clusters. Instead, your cluster will be upgraded to the highest currently available patch version available in EKS for the MAJOR.MINOR version of Kubernetes you selected.

### 11.2.8.1.4 EKS Clusters use CSI-based EBS Drivers

EKS has disabled support for in-tree EBS volume provisioning in favor of [CSI Volumes](#)<sup>1041</sup>. See the dropdown labeled “Additional Considerations for EKS” in the [\(2.4\) DKP Enterprise Upgrade | Upgrade the Core Addons](#) (see page 0) page for more details on what steps you must take when upgrading.

### 11.2.8.1.5 Intermittent Error Status when Creating EKS Clusters in the UI

When provisioning an EKS cluster through the UI, you may receive a brief error state because the EKS cluster may 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 will resolve, 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.

### 11.2.8.1.6 Installation and upgrade issue in pre-provisioned environments

DKP made the transition from Minio to Rook Ceph for cluster storage. An issue with Rook Ceph’s deployment prevents pre-provisioned environments from installing and upgrading to this DKP version. To solve this issue, you must set up 40 GB of raw storage for your worker nodes and customize your Rook Ceph installation as indicated in [Install Kommander in a Pre-provisioned Environment](#) (see page 515) or [Upgrade Kommander in a Pre-provisioned Environment](#) (see page 1031).

### 11.2.8.1.7 Cluster Roles' information not available in the Projects section of the UI

Selecting a role of the *Cluster Role* type in the **Projects > Roles** section of a workspace displays an error message in the UI.

#### 11.2.8.1.7.1 Workaround

You can still access a *Cluster Role*’s description and configuration from the UI. Take the following alternative path to *view* or *edit* the desired role:

---

<sup>1041</sup> <https://kubernetes.io/blog/2019/12/09/kubernetes-1-17-feature-csi-migration-beta/>

1. Select your workspace from the top navigation bar.
2. Select **Administration > Access Control** from the sidebar.
3. A table appears that lists roles from all *Projects* in the selected workspace.
4. Select the **Name** or **ID** of the *Cluster Role* you want to access. A page opens that contains more information and configuration options for the role.



To access a *Cluster Role*'s description and configuration page, ensure your user has sufficient *cluster view* or *edit* rights. You will not be able to select roles for which you do not have access rights.

### 11.2.8.1.8 Resolve issues with failed HelmReleases

An [issue with the Flux helm-controller](#)<sup>1042</sup> can cause HelmReleases to fail with the error message *Helm upgrade failed: another operation (install/upgrade/rollback) is in progress*. This can happen when the helm-controller is restarted while a HelmRelease is still upgrading, or installing.

#### 11.2.8.1.8.1 Workaround

To ensure the HelmRelease error was caused by the helm-controller restarting, first try to suspend/resume the HelmRelease:

```
kubectl -n <namespace> patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": true}]'
```

```
kubectl -n <namespace> patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[{"op": "replace", "path": "/spec/suspend", "value": false}]'
```

This might resolve the issue. If not, continue with the following steps:

You should see the HelmRelease attempting to reconcile, and then it either succeeds (with status: *Release reconciliation succeeded*) or it fails with the same error as before.

If the HelmRelease is still in the failed state, it is likely related to the helm-controller restarting. For example, if the 'reloader' HelmRelease is the one that is stuck.

To resolve the issue, follow these steps:

1. List secrets containing the affected HelmRelease name:

```
kubectl get secrets -n ${NAMESPACE} | grep reloader
```

The output should look like this:

<sup>1042</sup> <https://github.com/fluxcd/helm-controller/issues/149>

```

kommander-reloader-reloader-token-9qd8b kubernetes.io/
service-account-token 3 171m
sh.helm.release.v1.kommander-reloader.v1 helm.sh/
release.v1 1 171m
sh.helm.release.v1.kommander-reloader.v2 helm.sh/
release.v1 1 117m

```

In this example, `sh.helm.release.v1.kommander-reloader.v2` is the most recent revision.

2. Find and delete the most recent revision secret, for example,

```
sh.helm.release.v1.*.<revision>:
```

```
kubectl delete secret -n <namespace> <most recent helm revision secret name>
```

3. Suspend and resume the HelmRelease to trigger a reconciliation:

```

kubectl -n <namespace> patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[
{"op": "replace", "path": "/spec/suspend", "value": true}]'
kubectl -n <namespace> patch helmrelease <HELMRELEASE_NAME> --type='json' -p='[
{"op": "replace", "path": "/spec/suspend", "value": false}]'

```

You should see the HelmRelease is reconciled and eventually the upgrade and install succeed.

#### 11.2.8.1.8.2 Calico not updated during DKP upgrade on Flatcar

When upgrading a DKP cluster running on Flatcar OS, you may find that after the upgrade the Calico services were not updated. This occurs because the upgrade procedure is not correctly updating the Flatcar specific CNI ClusterResourceSet(CRS). This issue only impacts the Calico CRS.

Follow these steps to manually correct this issue:

1. Update the ConfigMap as follows:

```

cat <<EOF | kubectl apply -f -
apiVersion: v1
data:
 custom-resources.yaml: |+
 # This section includes base Calico installation configuration.
 # For more information, see: <https://docs.projectcalico.org/reference/
installation/api>
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: 192.168.0.0/16
 encapsulation: IPIP
 natOutgoing: Enabled
 nodeSelector: all()
 bgp: Enabled
 nodeAddressAutodetectionV4:
 firstFound: true
 # FlexVolume path must be mounted under /opt on flatcar/coreos systems
 flexVolumePath: /opt/libexec/kubernetes/kubelet-plugins/volume/exec/
kind: ConfigMap
metadata:
 name: calico-cni-installation- $\$$ CLUSTER_NAME
EOF

```

2. Run these commands:

```
kubectl edit clusterresourceset calico-cni-installation- $\$$ CLUSTER_NAME
```

and update

```
spec.clusterSelector.matchLabels.konvoy.d2iq.io/osHint to konvoy.d2iq.io/osHint:
flatcar
```

### 11.2.8.1.9 Kube-oidc-proxy not ready after upgrade

If you installed or attached a cluster in 2.1, `kube-oidc-proxy` is not available after upgrading to 2.3 and 2.4. This application is required to access the Kubernetes API (with `kubectl`) using SSO. For affected customers, there are issues with the authentication via `kubectl`.

To make the application available, run the following command on each cluster that was installed, created or attached in 2.1, and is now on DKP version 2.4. Replace `<namespace>` with each cluster's workspace namespace:

```
kubectl -n <namespace> patch appdeployment kube-oidc-proxy --type=json -p
' [{"op": "remove", "path": "/spec/configOverrides"}]'
```

## 12 Access Documentation

The following sections describe how to access other versions of documentation:

### 12.1 Supported Documentation

You can access all n-2 supported documentation at [D2iQ Help Center](#)<sup>1043</sup>.

### 12.2 Archived Documentation

In accordance with our [version support policy](#)<sup>1044</sup><sup>1045</sup>, 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>1046</sup>.
- Konvoy, Kommander, Kaptain, DC/OS and Dispatch at <https://github.com/mesosphere/dcos-docs-site/tags>.

---

<sup>1043</sup> <http://docs.d2iq.com/>

<sup>1044</sup> <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29923883/Version+Support+Policy>

<sup>1045</sup> <https://d2iq.atlassian.net/wiki/spaces/DENT/pages/29923883/Version+Support+Policy>

<sup>1046</sup> <https://archive-docs.d2iq.com/>

## 13 Download Documentation

This PDF contains the entire DKP documentation space. The file name contains the version and date created.



DKP\_2.4\_9-1-23.pdf

[\(see page 1097\)](#)



## 14 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 2.4, then D2iQ supports all patch versions of DKP 2.3, and 2.2.

When the next version releases, support continues for 2.4, and 2.3. Support for DKP version 2.2.x expires. You should upgrade DKP with every new release to stay up-to-date with the latest features and bug fixes.

### 14.1 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 1055).

### 14.2 Supported Operating Systems

Details for supported operating systems on DKP can be found in [Supported Operating Systems](#) (see page 93).

### 14.3 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.

### 14.4 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.

## 14.5 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.

## 14.6 Support Definitions

### 14.6.1 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

Type	Scope Example	Support Offered
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
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

Type	Scope Example	Support Offered
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

## 14.7 Standard Level & Severity Definitions

To view our severity level and support terms refer to [D2iQ Support and Maintenance Terms](https://d2iq.com/legal/support-terms)<sup>1047</sup>.

---

<sup>1047</sup> <https://d2iq.com/legal/support-terms>

## 15 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>1048</sup> for additional details.

---

<sup>1048</sup> <https://d2iq.com/legal/3rd>