



# **FlexPod hybrid cloud with NetApp Astra and Cisco Intersight for Red Hat OpenShift**

## **FlexPod**

NetApp  
March 25, 2024

This PDF was generated from <https://docs.netapp.com/us-en/flexpod/hybrid-cloud/flexpod-rho-cvo-solution-overview.html> on March 25, 2024. Always check docs.netapp.com for the latest.

# Table of Contents

- FlexPod hybrid cloud with NetApp Astra and Cisco Intersight for Red Hat OpenShift. . . . . 1
  - TR-4936: FlexPod hybrid cloud with NetApp Astra and Cisco Intersight for Red Hat OpenShift. . . . . 1
  - Solution components . . . . . 4
  - Installation and configuration . . . . . 12
  - Solution validation . . . . . 35
  - Conclusion . . . . . 56

# FlexPod hybrid cloud with NetApp Astra and Cisco Intersight for Red Hat OpenShift

## TR-4936: FlexPod hybrid cloud with NetApp Astra and Cisco Intersight for Red Hat OpenShift

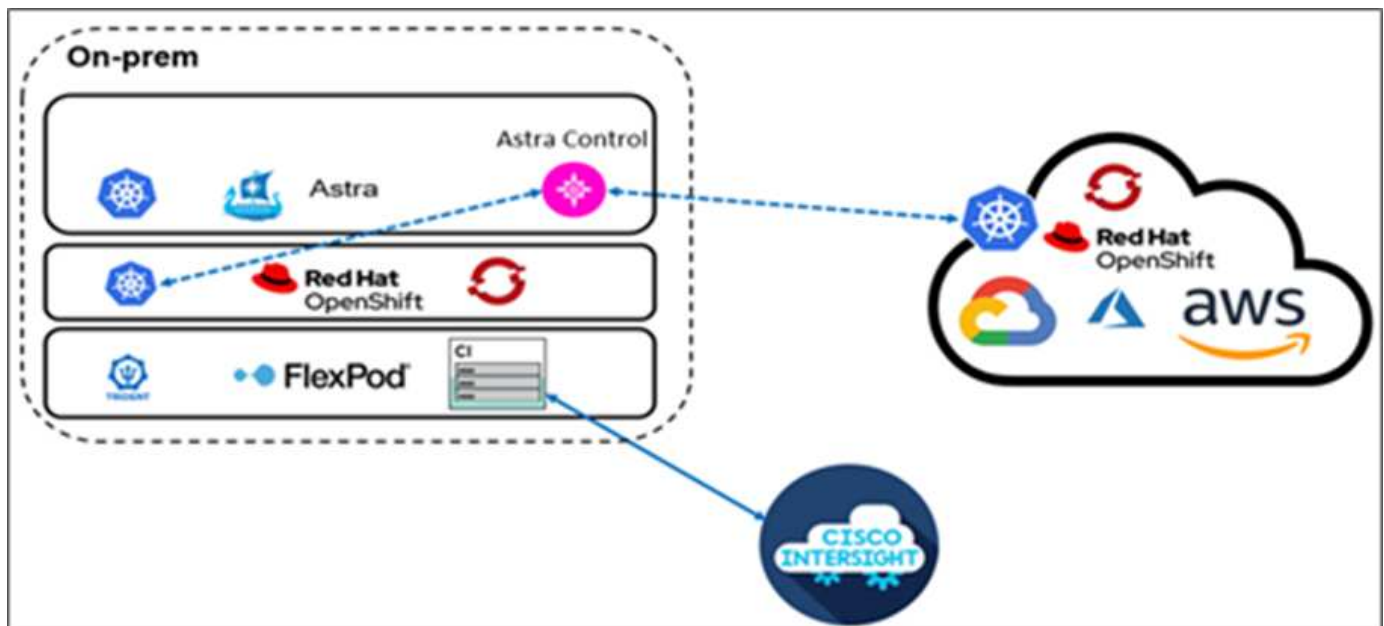
Abhinav Singh

### Introduction

As containers and Kubernetes become the de facto choice for developing, deploying, running, managing, and scaling containerized apps, enterprises are increasingly running business-critical applications on them. Business-critical applications are heavily dependent on state. A stateful application has associated state, data, and configuration information and depends on previous data transactions to execute its business logic. Business-critical applications while running on Kubernetes continue to have availability and business continuity requirements like traditional applications. A service outage can seriously affect a loss of revenue, productivity, and reputation of the company. Therefore, it's very critical to protect, recover, and move Kubernetes workloads quickly and easily within and across clusters, on-premises data centers, and Hybrid cloud environments. Enterprises have seen the benefits of shifting their business to a hybrid cloud model and modernizing their applications to a cloud-native form factor is high on their list.

This technical report brings together NetApp Astra Control Center with Red Hat OpenShift Container Platform on a FlexPod converged infrastructure solution and extends to Amazon Web Services (AWS) to form a hybrid cloud data center. Building on the familiarity with [FlexPod and Red Hat OpenShift](#), this document discusses NetApp Astra Control Center, starting from installation, configuration, application protection workflows, and application migration between on-premises and cloud. It also discusses the advantages of application-aware data management features (such as backup and recovery, business continuity) when using NetApp Astra Control Center for containerized applications running on Red Hat OpenShift.

The following figure illustrates the solution overview.



## Audience

The intended audience of this document includes chief technology officers (CTOs), application developers, cloud solution architects, site reliability engineers (SREs), DevOps Engineers, ITOps, and professional services teams that are focused on designing, hosting, and managing containerized applications.

## NetApp Astra Control – Key use cases

NetApp Astra Control aims at simplifying application protection for customers who deal with cloud native microservices:

- **Point-in-time (PiT) application representation with snapshots.** With Astra Control you can take end-to-end snapshots of your containerized applications that include the configuration details of the application running on Kubernetes and the associated persistent storage. In case of an incident, applications can be restored to a known good state in button click.
- **Full copy application backup.** With Astra Control you can take a full application backup on a predefined schedule which can be used to restore the application to the same K8s cluster or to a different K8s cluster on-demand in an automated fashion.
- **Application portability and migration with clones.** With Astra Control you can clone an entire application along with its data from one Kubernetes cluster to another or within the same K8s cluster. This feature also helps in porting or migrating an application across K8s clusters no matter where the clusters are located (simply delete the source application instance after cloning).
- **Customize application consistency.** With Astra Control you can take control of defining application quiesce states by leveraging the execution hooks. Drop the 'pre' and 'post' execution hooks to the snapshot and backup workflows, your applications will be quiesced in your own way before a snapshot or backup is taken.
- **Automate application-level disaster recovery (DR).** With Astra Control you can configure a business continuity disaster recovery (BCDR) plan for your containerized applications. NetApp SnapMirror is used in the backend and the complete implementation of the DR workflow is automated.

## Solution topology

This section describes the logical topology of the solution.

The following illustration represents the solution topology comprising the FlexPod on-premises environment running OpenShift Container Platform clusters, and a self-managed OpenShift Container Platform cluster on AWS with NetApp Cloud Volumes ONTAP, Cisco Intersight, and NetApp Cloud Manager SaaS platform.



# Solution components

[Previous: Solution overview.](#)

## FlexPod

FlexPod is a defined set of hardware and software that forms an integrated foundation for both virtualized and nonvirtualized solutions. FlexPod includes NetApp ONTAP storage, Cisco Nexus networking, Cisco MDS storage networking, Cisco Unified Computing System (Cisco UCS). The design is flexible enough that the networking, computing, and storage can fit into one data center rack, or it can be deployed according to a customer's data center design. Port density allows the networking components to accommodate multiple configurations.

## Astra Control

Astra Control offers application-aware data protection services for cloud-native applications that are hosted in both public clouds and on-premises. Astra Control delivers data protection, disaster recovery, and migration capabilities for your containerized application running on Kubernetes.


### Features

Astra Control offers critical capabilities for Kubernetes application data lifecycle management:

- Automatically manage persistent storage
- Create application consistent, on-demand snapshots and backups
- Automated policy-driven snapshot and backup operations
- Migrate applications and associated data from one Kubernetes cluster to another in a hybrid cloud setup
- Clone an application to the same K8s cluster or to another K8s cluster
- Visualize application protection status
- Provides a Graphical user interface and an exhaustive list of REST APIs to implement all protection workflows from existing in-house tools.

Astra Control provides a single pane of glass visualization for your containerized applications that includes an insight into their associated resources created on the Kubernetes cluster. You can view all your clusters, all your apps, in all clouds or in all data centers using one portal. You can use the Astra Control APIs across all environments (on-premises or public clouds) to implement your data management workflows.

The following image shows the Astra Control capabilities.



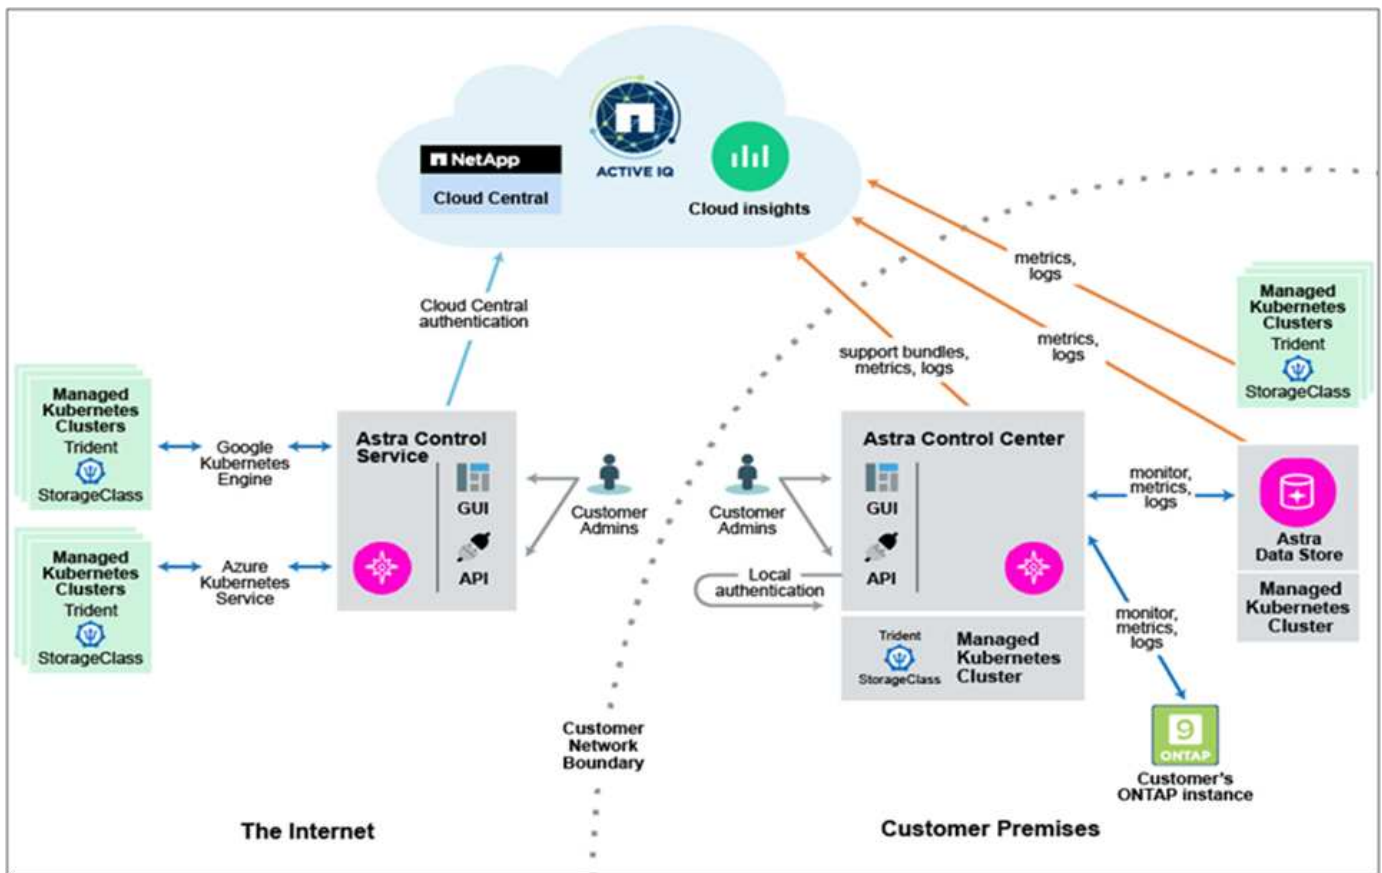
## Astra Control Consumption models

Astra Control is available in two consumption models:

- **Astra Control Service.** A fully managed service hosted by NetApp that provides application-aware data management of Kubernetes clusters in Google Kubernetes Engine (GKE), Azure Kubernetes Service (AKS).
- **Astra Control Center.** Self-managed software that provides application-aware data management of Kubernetes clusters running in your on-premises and hybrid cloud environment.

This technical report leverages Astra Control Center for the management of cloud-native applications running on Kubernetes.

The following image shows the Astra Control architecture.



## Astra Trident

Astra Trident is an open-source, fully supported storage orchestrator for containers and Kubernetes distributions. It was designed from the beginning to help you meet your containerized applications' persistence demands using industry-standard interfaces, such as the [Container Storage Interface \(CSI\)](#). With Astra Trident, microservices and containerized applications can take advantage of enterprise-class storage services provided by the NetApp portfolio of storage systems.

Astra Trident is deployed on Kubernetes clusters as pods and provides dynamic storage orchestration services for your Kubernetes workloads. It enables your containerized applications to consume persistent storage quickly and easily from NetApp's broad portfolio, which includes NetApp ONTAP (NetApp AFF, NetApp FAS, NetApp ONTAP Select, Cloud, and Amazon FSx for NetApp ONTAP), NetApp Element software (NetApp SolidFire), as well as the Azure NetApp Files service, Cloud Volume Service on Google Cloud, and the Cloud Volume Service on AWS. In a FlexPod environment, Astra Trident is used to dynamically provision and manage persistent volumes for containers that are backed by NetApp FlexVol volumes and LUNs hosted on an ONTAP storage platform such as NetApp AFF and FAS systems and Cloud Volumes ONTAP. Trident also plays a key role in the implementation of application protection schemes delivered by Astra Control. For more information about Astra Trident, see the [Astra Trident documentation](#).

## Storage backend

To use Astra Trident, you need supported storage backend. A Trident backend defines the relationship between Trident and a storage system. It tells Trident how to communicate with that storage system and how Trident should provision volumes from it. Trident will automatically offer up storage pools from backends that together match the requirements defined by a storage class.

- ONTAP AFF and FAS storage backend. As a storage software and hardware platform, ONTAP provides core storage services, support for multiple storage access protocols, and storage management



functionality, such as NetApp Snapshot copies and mirroring.

- Cloud Volumes ONTAP storage backend
- [Astra Data Store](#) storage backend

## NetApp Cloud Volumes ONTAP

NetApp Cloud Volumes ONTAP is a software-defined storage offering that delivers advanced data management for file and block workloads. With Cloud Volumes ONTAP, you can optimize your cloud storage costs and increase application performance while enhancing data protection, security, and compliance.

Key benefits include:

- Leverage built-in data deduplication, data compression, thin provisioning, and cloning to minimize storage costs.
- Ensure enterprise reliability and continuous operations in case of failures in your cloud environment.
- Cloud Volumes ONTAP leverages SnapMirror, NetApp's industry-leading replication technology, to replicate on-premises data to the cloud so it's easy to have secondary copies available for multiple use cases.
- Cloud Volumes ONTAP also integrates with Cloud Backup service to deliver backup and restore capabilities for protection, and long-term archive of your cloud data.
- Switch between high and low-performance storage pools on-demand without taking applications offline.
- Ensure consistency of Snapshot copies using NetApp SnapCenter.
- Cloud Volumes ONTAP supports data encryption and provides protection against viruses and ransomware.
- Integration with Cloud Data Sense helps you understand data context and identify sensitive data.

## Cloud Central

Cloud Central provides a centralized location to access and manage NetApp cloud data services. These services enable you to run critical applications in the cloud, create automated DR sites, back up your data, and effectively migrate and control data across multiple clouds. For more information, see [Cloud Central](#).

## Cloud Manager

Cloud Manager is an enterprise-class, SaaS-based management platform that enables IT experts and cloud architects to centrally manage their hybrid multi-cloud infrastructure using NetApp's cloud solutions. It provides a centralized system for viewing and managing your on-premises and cloud storage, supporting hybrid, multiple cloud providers and accounts. For more information, see [Cloud Manager](#).

## Connector

Connector is an instance that enables Cloud Manager to manage resources and processes within public cloud environment. A Connector is required to use many features that Cloud Manager provides. A Connector can be deployed in the cloud or on-premises network.

Connector is supported in the following locations:

- AWS
- Microsoft Azure
- Google Cloud

- On your premises

To learn more about Connector, see [this link](#).

## NetApp Cloud Insights

A NetApp cloud infrastructure monitoring tool, Cloud Insights enables you to monitor performance and utilization for your Kubernetes clusters managed by Astra Control Center. Cloud Insights correlates storage usage to workloads. When you enable the Cloud Insights connection in Astra Control Center, telemetry information shows in Astra Control Center UI pages.

## NetApp Active IQ Unified Manager

NetApp Active IQ Unified Manager allows you to monitor your ONTAP storage clusters from a single redesigned, intuitive interface that delivers intelligence from community wisdom and AI analytics. It provides comprehensive operational, performance, and proactive insights into the storage environment and the virtual machines (VMs) running on it. When an issue occurs with the storage infrastructure, Unified Manager can notify you about the details of the issue to help with identifying the root cause. The VM dashboard gives you a view into the performance statistics for the VM so that you can investigate the entire I/O path from the VMware vSphere host down through the network and finally to the storage. Some events also provide remedial actions that can be taken to rectify the issue. You can configure custom alerts for events so that when issues occur, you are notified through email and SNMP Traps. Active IQ Unified Manager enables planning for the storage requirements of your users by forecasting capacity and usage trends to proactively act before issues arise, preventing reactive short-term decisions that can lead to additional problems in the long term.

## Cisco Intersight


Cisco Intersight is a SaaS platform that delivers intelligent automation, observability, and optimization for traditional and cloud-native applications and infrastructure. The platform helps drive change with IT teams and delivers an operating model designed for hybrid cloud.

Cisco Intersight provides the following benefits:

- **Faster delivery.** Delivered as a service from the cloud or in the customer's data center with frequent updates and continued innovation, due to an agile-based software development model. This way, customer can just focus on accelerating delivery for line-of-business.
- **Simplified operations.** Simplify operations by using a single secure SaaS-delivered tool with common inventory, authentication, and APIs to work across full stack and all locations, eliminating silos across teams. From managing physical servers and hypervisors on-premises, to VMs, K8s, serverless, automation, optimization, and cost control across both on-premises and public clouds.
- **Continuous optimization.** Continuously optimize your environment by using intelligence provided by Cisco Intersight across every layer, as well as Cisco TAC. This intelligence is converted into recommended and automatable actions so you can adapt real-time to every change: from moving workloads and monitoring health of physical servers to auto sizing K8s clusters, to cost reduction recommendations the public clouds you work with.

There are two modes of management operations possible with Cisco Intersight: UCSM Managed Mode (UMM) and Intersight Managed Mode (IMM). You can select the native UMM or IMM for the fabric-attached Cisco UCS Systems during initial setup of the Fabric Interconnects. In this solution, native UMM is used.

The following image shows the Cisco Intersight dashboard.




## Red Hat OpenShift Container Platform

The Red Hat OpenShift Container Platform is a container application platform that brings together CRI-O and Kubernetes and provides an API and web interface to manage these services. CRI-O is an implementation of the Kubernetes Container Runtime Interface (CRI) to enable using Open Container Initiative (OCI) compatible runtimes. It is a lightweight alternative to using Docker as the runtime for Kubernetes.

OpenShift Container Platform allows customers to create and manage containers. Containers are standalone processes that run within their own environment, independent of operating system and the underlying infrastructure. OpenShift Container Platform helps develop, deploy, and manage container-based applications. It provides a self-service platform to create, modify, and deploy applications on demand, thus enabling faster development and release life cycles. OpenShift Container Platform has a microservices-based architecture of smaller, decoupled units that work together. It runs on top of a Kubernetes cluster, with data about the objects stored in etcd, a reliable clustered key-value store.

The following image is an overview of the Red Hat OpenShift Container platform.



## Kubernetes infrastructure

Within OpenShift Container Platform, Kubernetes manages containerized applications across a set of CRI-O runtime hosts and provides mechanisms for deployment, maintenance, and application-scaling. The CRI-O service packages, instantiates, and runs containerized applications.

A Kubernetes cluster consists of one or more masters and a set of worker nodes. This solution design includes high availability (HA) functionality at the hardware as well as the software stack. A Kubernetes cluster is designed to run in HA mode with three master nodes and a minimum of two worker nodes to help ensure that the cluster has no single point of failure.

## Red Hat Core OS

OpenShift Container Platform uses Red Hat Enterprise Linux CoreOS (RHCOS), a container-oriented operating system that combines some of the best features and functions of the CoreOS and Red Hat Atomic Host operating systems. RHCOS is specifically designed for running containerized applications from OpenShift Container Platform and works with new tools to provide fast installation, operator-based management, and simplified upgrades.

RHCOS includes the following features:

- Ignition, which OpenShift Container Platform uses as a first boot system configuration for initially bringing up and configuring machines.
- CRI-O, a Kubernetes native container runtime implementation that integrates closely with the operating system to deliver an efficient and optimized Kubernetes experience. CRI-O provides facilities for running, stopping, and restarting containers. It fully replaces the Docker Container Engine, which was used in OpenShift Container Platform 3.
- Kubelet, the primary node agent for Kubernetes, is responsible for launching and monitoring containers.

## VMware vSphere 7.0

VMware vSphere is a virtualization platform for holistically managing large collections of infrastructures (resources including CPUs, storage, and networking) as a seamless, versatile, and dynamic operating environment. Unlike traditional operating systems that manage an individual machine, VMware vSphere aggregates the infrastructure of an entire data center to create a single powerhouse with resources that can be allocated quickly and dynamically to any application in need.

For more information, see [VMware vSphere](#).

### VMware vSphere vCenter

VMware vCenter Server provides unified management of all hosts and VMs from a single console and aggregates performance monitoring of clusters, hosts, and VMs. VMware vCenter Server gives administrators a deep insight into the status and configuration of compute clusters, hosts, VMs, storage, the guest OS, and other critical components of a virtual infrastructure. VMware vCenter manages the rich set of features available in a VMware vSphere environment.

## Hardware and software revisions

This solution can be extended to any FlexPod environment that is running supported versions of software, firmware, and hardware as defined in the [NetApp Interoperability Matrix Tool](#) and [Cisco UCS Hardware Compatibility List](#). The OpenShift cluster is installed on FlexPod in a Bare Metal fashion as well as on VMware vSphere.

Only a single instance of Astra Control Center is required to manage multiple OpenShift (k8s) clusters, while Trident CSI is installed on each OpenShift cluster. Astra Control Center can be installed on any of these OpenShift cluster. In this solution, Astra Control Center is installed on the OpenShift bare-metal cluster.

The following table lists the FlexPod hardware and software revisions for OpenShift.

Component	Product	Version
Compute	Cisco UCS Fabric Interconnects 6454	4.1(3c)
	Cisco UCS B200 M5 Servers	4.1(3c)
Network	Cisco Nexus 9336C-FX2 NX-OS	9.3(8)
Storage	NetApp AFF A700	9.11.1
	NetApp Astra Control Center	22.04.0
	NetApp Astra Trident CSI Plugin	22.04.0
	NetApp Active IQ Unified Manager	9.11
Software	VMware ESXi nenic Ethernet Driver	1.0.35.0
	vSphere ESXi	7.0(U2)
	VMware vCenter Appliance	7.0 U2b
	Cisco Intersight Assist Virtual Appliance	1.0.9-342
	OpenShift Container Platform	4.9

Component	Product	Version
	OpenShift Container Platform Master Node	RHCOS 4.9
	OpenShift Container Platform Worker Node	RHCOS 4.9

The following table lists the software versions for OpenShift on AWS.

Component	Product	Version
Compute	Master Instance Type: m5.xlarge	n/a
	Worker Instance Type: m5.large	n/a
Network	Virtual Private Cloud Transit Gateway	n/a
Storage	NetApp Cloud Volumes ONTAP	9.11.1
	NetApp Astra Trident CSI Plugin	22.04.0
Software	OpenShift Container Platform	4.9
	OpenShift Container Platform Master Node	RHCOS 4.9
	OpenShift Container Platform Worker Node	RHCOS 4.9

[Next: FlexPod for OpenShift Container Platform 4 bare-metal installation.](#)

## Installation and configuration

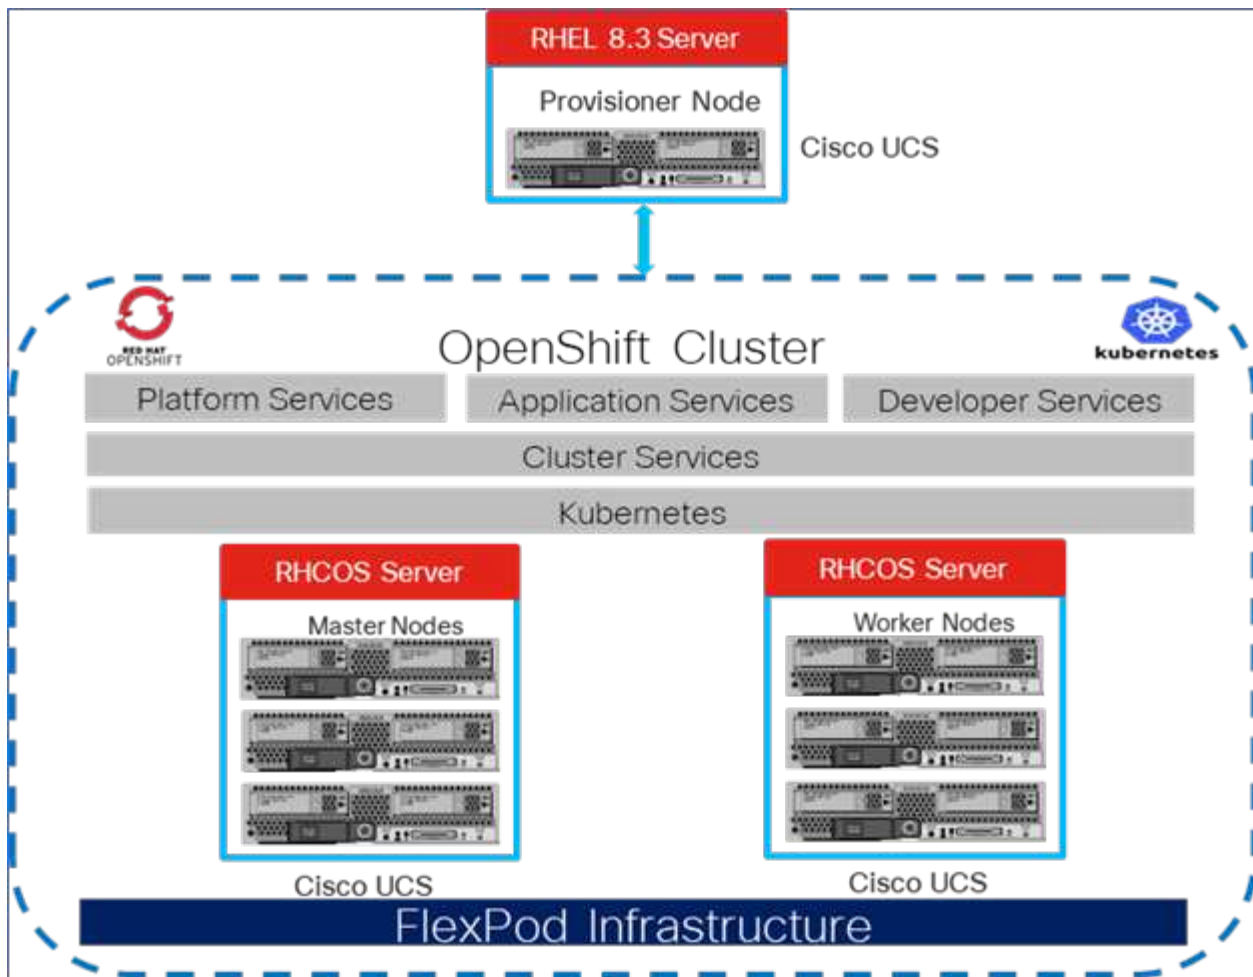
### FlexPod for OpenShift Container Platform 4 bare-metal installation

[Previous: Solution components.](#)

To understand FlexPod for OpenShift Container Platform 4 bare-metal design, deployment details, and the NetApp Astra Trident installation and configuration, see [FlexPod with OpenShift Cisco Validated Design and Deployment guide \(CVD\)](#). This CVD covers FlexPod and OpenShift Container Platform deployment using Ansible. The CVD also provide detailed information about preparing worker nodes, Astra Trident installation, storage backend, and storage class configurations, which are the few prerequisites for deploying and configuring Astra Control Center.

The following figure illustrates the OpenShift Container Platform 4 Bare Metal on FlexPod.

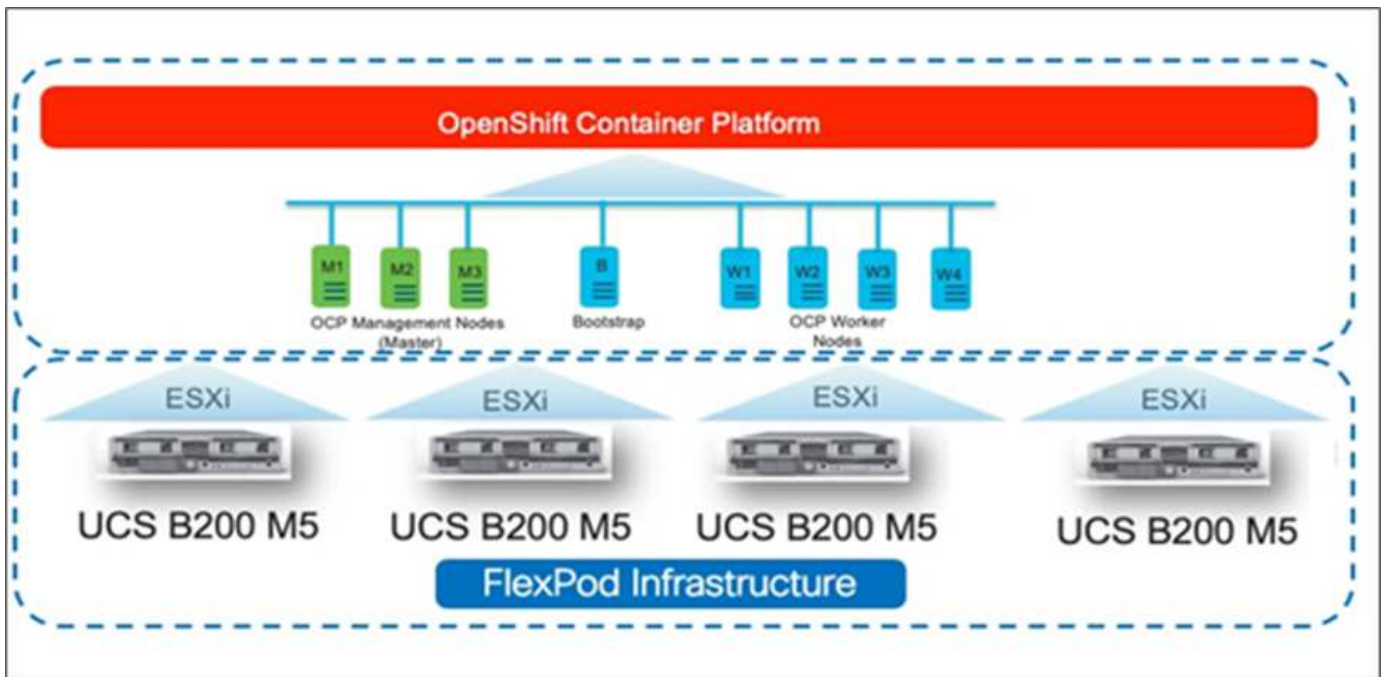




### FlexPod for OpenShift Container Platform 4 on VMware installation

For more information about deploying Red Hat OpenShift Container Platform 4 on FlexPod running VMware vSphere, see [FlexPod Datacenter for OpenShift Container Platform 4](#).

The following figure illustrates FlexPod for OpenShift Container Platform 4 on vSphere.



Next: Red Hat OpenShift on AWS.

## Red Hat OpenShift on AWS

Previous: FlexPod for OpenShift Container Platform 4 bare-metal installation.

A separate self-managed OpenShift Container Platform 4 cluster is deployed on AWS as a DR site. The master and worker nodes span across three availability zones for high availability.

Instances (6) [Info](#)

Q

Search

ocp

X

Clear filters

<input type="checkbox"/>	Name	Instance ID	Instance state	Instance type	Availability Zone	Private IP a...	Key name
<input type="checkbox"/>	ocpaws-v58kn-master-0	i-0d2d81ca91a54276d	<span>Running</span>	m5.xlarge	us-east-1b	172.30.165.160	-
<input type="checkbox"/>	ocpaws-v58kn-master-1	i-0b161945421d2a23c	<span>Running</span>	m5.xlarge	us-east-1c	172.30.166.162	-
<input type="checkbox"/>	ocpaws-v58kn-master-2	i-0146a665e1060ea59	<span>Running</span>	m5.xlarge	us-east-1a	172.30.164.209	-
<input type="checkbox"/>	ocpaws-v58kn-worker-us-east-1a-zj8dj	i-05e6efa18d136c842	<span>Running</span>	m5.large	us-east-1a	172.30.164.128	-
<input type="checkbox"/>	ocpaws-v58kn-worker-us-east-1b-7nmbc	i-0879a088b50d2d966	<span>Running</span>	m5.large	us-east-1b	172.30.165.93	-
<input type="checkbox"/>	ocpaws-v58kn-worker-us-east-1c-96j6n	i-0c24ff3c2d701f82c	<span>Running</span>	m5.large	us-east-1c	172.30.166.51	-



```
[ec2-user@ip-172-30-164-92 ~]$ oc get nodes
```

NAME	STATUS	ROLES	AGE	VERSION
ip-172-30-164-128.ec2.internal	Ready	worker	29m	v1.22.8+f34b40c
ip-172-30-164-209.ec2.internal	Ready	master	36m	v1.22.8+f34b40c
ip-172-30-165-160.ec2.internal	Ready	master	33m	v1.22.8+f34b40c
ip-172-30-165-93.ec2.internal	Ready	worker	30m	v1.22.8+f34b40c
ip-172-30-166-162.ec2.internal	Ready	master	36m	v1.22.8+f34b40c
ip-172-30-166-51.ec2.internal	Ready	worker	28m	v1.22.8+f34b40c

OpenShift is deployed as a [private cluster](#) into an existing VPC on AWS. A private OpenShift Container Platform cluster does not expose external endpoints and is accessible from only an internal network and is not visible to the internet. A single-node NetApp Cloud Volumes ONTAP is deployed using NetApp Cloud Manager, which provides a storage backend to Astra Trident.

For more information about installing OpenShift on AWS, see [OpenShift documentation](#).

[Next: NetApp Cloud Volumes ONTAP.](#)

## NetApp Cloud Volumes ONTAP

[Previous: Red Hat OpenShift on AWS.](#)

The NetApp Cloud Volumes ONTAP instance is deployed on AWS, and it serves as backend storage to Astra Trident. Before adding a Cloud Volumes ONTAP working environment, a Connector must be deployed. The Cloud Manager prompts you if you try to create your first Cloud Volumes ONTAP working environment without a Connector in place. To deploy a Connector in AWS, see [Create a Connector](#).

To deploy Cloud Volumes ONTAP on AWS, see [Quick Start for AWS](#).

After Cloud Volumes ONTAP is deployed, you can install Astra Trident and configure the storage backend and snapshot class on the OpenShift Container Platform cluster.

[Next: Astra Control Center installation on OpenShift Container Platform.](#)

## Astra Control Center installation on OpenShift Container Platform

[Previous: NetApp Cloud Volumes ONTAP.](#)

You can install Astra Control Center either on OpenShift cluster running on FlexPod or on AWS with a Cloud Volumes ONTAP storage backend. In this solution, Astra Control Center is deployed on the OpenShift bare-metal cluster.

Astra Control Center can be installed using the standard process described [here](#) or from the Red Hat OpenShift OperatorHub. Astra Control Operator is a Red Hat certified operator. In this solution, Astra Control Center is installed using the Red Hat OperatorHub.

## Environment requirements

- Astra Control Center supports multiple Kubernetes distributions; for Red Hat OpenShift, the supported versions include Red Hat OpenShift Container Platform 4.8 or 4.9.
- Astra Control Center requires the following resources in addition to the environment's and the end-user's application resource requirements:

Components	Requirement
Storage backend capacity	At least 500GB available
Worker nodes	At least 3 worker nodes, with 4 CPU cores and 12GB RAM each
Fully qualified domain name (FQDN) address	An FQDN address for Astra Control Center
Astra Trident	Astra Trident 21.04 or newer installed and configured
Ingress controller or load balancer	Configure the ingress controller to expose Astra Control Center with a URL or load balancer to provide IP address which will resolve to the FQDN

- You must have an existing private image registry to which you can push the Astra Control Center build images. You need to provide the URL of the image registry where you upload the images.



Some images are pulled while executing certain workflows, and containers are created and destroyed when necessary.

- Astra Control Center requires that a storage class be created and set as the default storage class. Astra Control Center supports the following ONTAP drivers provided by Astra Trident:
  - ontap-nas
  - ontap-nas-flexgroup
  - ontap-san
  - ontap-san-economy



We assume that the deployed OpenShift clusters have Astra Trident installed and configured with an ONTAP backend, and a default storage class is also defined.

- For application cloning in OpenShift environments, Astra Control Center needs to allow OpenShift to mount volumes and change the ownership of files. To modify the ONTAP export policy to allow these operations, run the following commands:

```
export-policy rule modify -vserver <storage virtual machine name>  
-policyname <policy name> -ruleindex 1 -superuser sys  
export-policy rule modify -vserver <storage virtual machine name>  
-policyname <policy name> -ruleindex 1 -anon 65534
```



To add a second OpenShift operational environment as a managed compute resource, make sure that the Astra Trident Volume snapshot feature is enabled. To enable and test volume snapshots with Astra Trident, see the official [Astra Trident instructions](#).

- A [VolumeSnapClass](#) should be configured on all Kubernetes clusters from where the applications is managed. This could also include the K8s cluster on which Astra Control Center is installed. Astra Control Center can manage applications on the K8s cluster on which it is running.

### Application management requirements

- **Licensing.** To manage applications using Astra Control Center, you need an Astra Control Center license.
- **Namespaces.** A namespace is the largest entity that can be managed as an application by Astra Control Center. You can choose to filter out components based on the application labels and custom labels in an existing namespace and manage a subset of resources as an application.
- **StorageClass.** If you install an application with a StorageClass explicitly set and you need to clone the application, the target cluster for the clone operation must have the originally specified StorageClass. Cloning an application with an explicitly set StorageClass to a cluster that does not have the same StorageClass fails.
- **Kubernetes resources.** Applications that use Kubernetes resources not captured by Astra Control might not have full application data management capabilities. Astra Control can capture the following Kubernetes resources:

Kubernetes resources		
ClusterRole	ClusterRoleBinding	ConfigMap
CustomResourceDefinition	CustomResource	CronJob
DaemonSet	HorizontalPodAutoscaler	Ingress
DeploymentConfig	MutatingWebhook	PersistentVolumeClaim
Pod	PodDisruptionBudget	PodTemplate
NetworkPolicy	ReplicaSet	Role
RoleBinding	Route	Secret
ValidatingWebhook		

### Install Astra Control Center using OpenShift OperatorHub

The following procedure installs Astra Control Center using Red Hat OperatorHub. In this solution, Astra Control Center is installed on a bare-metal OpenShift cluster running on FlexPod.

1. Download the Astra Control Center bundle (`astra-control-center-[version].tar.gz`) from the [NetApp Support site](#).
2. Download the .zip file for the Astra Control Center certificates and keys from the [NetApp Support site](#).
3. Verify the signature of the bundle.

```
openssl dgst -sha256 -verify astra-control-center[version].pub
-signture <astra-control-center[version].sig astra-control-
center[version].tar.gz
```

4. Extract the Astra images.

```
tar -vxzf astra-control-center-[version].tar.gz
```

5. Change to the Astra directory.

```
cd astra-control-center-[version]
```

6. Add the images to your local registry.

```
For Docker:
docker login [your_registry_path]OR
For Podman:
podman login [your_registry_path]
```

7. Use the appropriate script to load the images, tag the images, and push them to your local registry.

For Docker:

```
export REGISTRY=[Docker_registry_path]
for astraImageFile in $(ls images/*.tar) ; do
    # Load to local cache. And store the name of the loaded image trimming
    the 'Loaded images: '
    astraImage=$(docker load --input ${astraImageFile} | sed 's/Loaded
image: //' )
    astraImage=$(echo ${astraImage} | sed 's!localhost/!!')
    # Tag with local image repo.
    docker tag ${astraImage} ${REGISTRY}/${astraImage}
    # Push to the local repo.
    docker push ${REGISTRY}/${astraImage}
done
```


For Podman:

```

export REGISTRY=[Registry_path]
for astraImageFile in $(ls images/*.tar) ; do
    # Load to local cache. And store the name of the loaded image trimming
    the 'Loaded images: '
    astraImage=$(podman load --input ${astraImageFile} | sed 's/Loaded
image(s): //'')
    astraImage=$(echo ${astraImage} | sed 's!localhost/!!!')
    # Tag with local image repo.
    podman tag ${astraImage} ${REGISTRY}/${astraImage}
    # Push to the local repo.
    podman push ${REGISTRY}/${astraImage}
done


```

- Log into the bare-metal OpenShift cluster web console. From the side menu, select Operators > OperatorHub. Enter astra to list the netapp-acc-operator.



netapp-acc-operator is a certified Red Hat OpenShift Operator and is listed under the OperatorHub catalogue.

- Select netapp-acc-operator and click Install.



**netapp-acc-operator**  
 22.4.3 provided by NetApp

Install

**Latest version**  
 22.4.3

**Capability level**  
☒ Basic Install  
☐ Seamless Upgrades  
☐ Full Lifecycle  
☐ Deep Insights  
☐ Auto Pilot

**Source**  
 Certified

**Provider**  
 NetApp

Astra Control is an application-aware data management solution that manages, protects and moves data-rich Kubernetes workloads in both public clouds and on-premises.

Astra Control enables data protection, disaster recovery, and migration for your Kubernetes workloads, leveraging NetApp's industry-leading data management technology for snapshots, backups, replication and cloning.

**How to deploy Astra Control**

Refer to [Installation Procedure](#) to deploy Astra Control Center using the Operator.

**Documentation**

Refer to [Astra Control Center Documentation](#) to complete the setup and start managing applications.

**NOTE:** The version listed under *Latest version* on this page might not reflect the actual version of NetApp Astra Control Center you are installing. The version in the file name of the Astra Control Center bundle that you download from the NetApp Support Site is the version of Astra Control Center that will be installed.

10. Select the appropriate options and click Install.

OperatorHub > Operator Installation

## Install Operator

Install your Operator by subscribing to one of the update channels to keep the Operator up to date. The strategy determines either manual or automatic updates.

**Update channel \*** ⓘ  
☐ alpha  
☒ stable


**Installation mode \***  
☒ All namespaces on the cluster (default)  
 Operator will be available in all Namespaces.  
☐ A specific namespace on the cluster  
 This mode is not supported by this Operator

**Installed Namespace \***


**Update approval \*** ⓘ  
☐ Automatic  
☒ Manual

**Namespace creation**  
 Namespace **netapp-acc-operator** does not exist and will be created.

**Manual approval applies to all operators in a namespace**  
 Installing an operator with manual approval causes all operators installed in namespace **netapp-acc-operator** to function as manual approval strategy. To allow automatic approval, all operators installed in the namespace must use automatic approval strategy.


**netapp-acc-operator**  
 provided by NetApp


**Provided APIs**  


**Astra Control Center**  
 AstraControlCenter is the Schema for the astracontrolcenters API.


Install

Cancel


11. Approve the installation and wait for the operator to be installed.



- At this stage, the operator is installed successfully and ready for use. Click View Operator to start the installation of Astra Control Center.



- Before installing Astra Control Center, create the pull secret to download Astra images from the Docker registry that you pushed earlier.



14. To pull the Astra Control Center images from your Docker private repo, create a secret in the `netapp-acc-operator` namespace. This secret name is provided in the Astra Control Center YAML manifest in a later step.



Project: netapp-acc-operator ▼

## Create image pull secret

Image pull secrets let you authenticate against a private image registry.

**Secret name \***

Unique name of the new secret.

**Authentication type**

**Registry server address \***

For example quay.io or docker.io

**Username \***


**Password \***

**Email**

[+ Add credentials](#)

CreateCancel

- From the side menu, select Operators > Installed Operators and click Create Instance under the provided APIs section.



16. Complete the Create AstraControlCenter form. Provide the name, Astra address, and Astra version.

The screenshot shows the 'Create AstraControlCenter' form in the Red Hat OpenShift Container Platform interface. The left navigation menu is the same as in the previous image. The main content area is titled 'Project: netapp-acc-operator' and shows 'netapp-acc-operator > Create AstraControlCenter'. The form title is 'Create AstraControlCenter' with a subtitle 'Create by completing the form. Default values may be provided by the Operator authors.' Below the title are tabs for 'Form view' (selected) and 'YAML view'. A note states: 'Note: Some fields may not be represented in this form view. Please select "YAML view" for full control.' The form fields are:
 

- Name \***: acc
- Labels**: app=frontend
- Auto Support \***: A section with a right arrow and text: 'AutoSupport indicates willingness to participate in NetApp's proactive support application, NetApp Active IQ. An internet connection is required (port 442) and all support data is anonymized. The default election is true and indicates no support data will be sent to NetApp. An empty or blank election is the same as a default election. Air gapped installations should enter false.'
- Astra Address \***: acc.ocp.flexpod.netapp.com. Below this is a description: 'AstraAddress defines how Astra will be found in the data center. This IP address and/or DNS A record must be created prior to provisioning Astra Control Center. Example - "astra.example.com" The A record and its IP address must be allocated prior to provisioning Astra Control Center.'
- Astra Version \***: 22.04.0. Below this is a description: 'Version of AstraControlCenter to deploy. You are provided a Helm repository with a corresponding version. Example - 1.5.2, 1.4.2-patch.'



Under Astra Address, provide the FQDN address for Astra Control Center. This address is used to access the Astra Control Center Web console. The FQDN should also resolve to a reachable IP network and should be configured in the DNS.

17. Enter an account name, email address, administrator last name, and retain the default volume reclaim policy. If you are using a load balancer, set the Ingress Type to AccTraefik. Otherwise, select Generic for

Ingress.Controller. Under Image Registry, enter the container image registry path and secret.

The screenshot shows the configuration page for the 'netapp-acc-operator' project in Astra Control Center. The left sidebar contains navigation links: Administrator, Home, Operators, OperatorHub, Installed Operators, Workloads, Networking, Storage, Builds, Observe, Compute, User Management, and Administration. The main form area is titled 'Project: netapp-acc-operator' and contains the following fields:

- Account Name \***: ocp (Astra Control Center account name)
- Email \***: abhinav3@netapp.com (EmailAddress will be notified by Astra as events warrant.)
- Last Name**: Singh (The last name of the SRE supporting Astra.)
- Volume Reclaim Policy**: Retain (Reclaim policy to be set for persistent volumes)
- Ingress Type**: AccTraefik (IngressType The type of ingress to that ACC should be configured for)
- Astra Kube Config Secret**: (AstraKubeConfigSecret if present and secret exists operator will attempt to add KubeConfig to Managed Clusters.)
- Image Registry**:
  - Name**: (The name of the image registry. For example "example.registry/astra". Do not prefix with protocol.)
  - Secret**: astra-registry-cred (The name of the Kubernetes secret that will authenticate with the image registry.)



In this solution, the Metallb load balancer is used. Therefore, the ingress type is AccTraefik. This exposes the Astra Control Center traefik gateway as a Kubernetes service of type LoadBalancer.

18. Enter the admin first name, configure the resource scaling, and provide the storage class. Click Create.

**Image Registry**

The container image registry that is hosting the Astra application images, ACC Operator and ACC Helm Repository.

**First Name**  
Abhinav

The first name of the SRE supporting Astra

**Astra Resources Scaler**  
Default

Scaling options for AstraControlCenter Resource limits.

**Storage Class**  
ocp-nas-sc-gold

The storage class to be used for PVCs. If not set, default storage class will be used.

**Crds**

Options for how ACC should handle CRDs. Options for how ACC should handle CRDs. Options for how ACC should handle CRDs. Options for how ACC should handle CRDs.

[Create](#) [Cancel](#)

The status of the Astra Control Center instance should change from Deploying to Ready.

Project: netapp-acc-operator

Installed Operators > Operator details

netapp-acc-operator  
22.43 provided by NetApp

Details | YAML | Subscription | Events | **Astra Control Center**

**AstraControlCenters** [Create AstraControlCenter](#)

Name Search by name...

Name	Kind	Status	Labels	Last updated
ACC-acc	AstraControlCenter	Conditions: Ready, PostinstallComplete, Deployed	appacc	8 minutes ago

19. Verify that all system components have been installed successfully and that all pods are running.

```
root@abhinav-ansible# oc get pods -n netapp-acc-operator
```

NAME	READY	STATUS
acc-helm-repo-77745b49b5-7zg2v	1/1	Running
acc-operator-controller-manager-5c656c44c6-tqnmn	2/2	Running
activity-589c6d59f4-x2sfs	1/1	Running

6m4s			
api-token-authentication-4q5lj	1/1	Running	0
5m26s			
api-token-authentication-pzptd	1/1	Running	0
5m27s			
api-token-authentication-tbtg6	1/1	Running	0
5m27s			
asup-669df8d49-qps54	1/1	Running	0
5m26s			
authentication-5867c5f56f-dnpp2	1/1	Running	0
3m54s			
bucket-service-85495bc475-5zcc5	1/1	Running	0
5m55s			
cert-manager-67f486bbc6-txhh6	1/1	Running	0
9m5s			
cert-manager-cainjector-75959db744-4l5p5	1/1	Running	0
9m6s			
cert-manager-webhook-765556b869-g6wdf	1/1	Running	0
9m6s			
cloud-extension-5d595f85f-txrfl	1/1	Running	0
5m27s			
cloud-insights-service-674649567b-5s4wd	1/1	Running	0
5m49s			
composite-compute-6b58d48c69-46vhc	1/1	Running	0
6m11s			
composite-volume-6d447fd959-chnrt	1/1	Running	0
5m27s			
credentials-66668f8ddd-8qc5b	1/1	Running	0
7m20s			
entitlement-fd6fc5c58-wxnmh	1/1	Running	0
6m20s			
features-756bbb7c7c-rgcrm	1/1	Running	0
5m26s			
fluent-bit-ds-278pg	1/1	Running	0
3m35s			
fluent-bit-ds-5pqc6	1/1	Running	0
3m35s			
fluent-bit-ds-8l7cq	1/1	Running	0
3m35s			
fluent-bit-ds-9qbft	1/1	Running	0
3m35s			
fluent-bit-ds-nj475	1/1	Running	0
3m35s			
fluent-bit-ds-x9pd8	1/1	Running	0
3m35s			
graphql-server-698d6f4bf-kftwc	1/1	Running	0

3m20s			
identity-5d4f4c87c9-wjz6c	1/1	Running	0
6m27s			
influxdb2-0	1/1	Running	0
9m33s			
krakend-657d44bf54-8cb56	1/1	Running	0
3m21s			
license-594bbdc-rghdg	1/1	Running	0
6m28s			
login-ui-6c65fbbbd4-jg8wz	1/1	Running	0
3m17s			
loki-0	1/1	Running	0
9m30s			
metrics-facade-75575f69d7-hnlk6	1/1	Running	0
6m10s			
monitoring-operator-65dff79cfb-z78vk	2/2	Running	0
3m47s			
nats-0	1/1	Running	0
10m			
nats-1	1/1	Running	0
9m43s			
nats-2	1/1	Running	0
9m23s			
nautilus-7bb469f857-4hlc6	1/1	Running	0
6m3s			
nautilus-7bb469f857-vz94m	1/1	Running	0
4m42s			
openapi-8586db4bcd-gwwvf	1/1	Running	0
5m41s			
packages-6bdb949cfb-nrq8l	1/1	Running	0
6m35s			
polaris-consul-consul-server-0	1/1	Running	0
9m22s			
polaris-consul-consul-server-1	1/1	Running	0
9m22s			
polaris-consul-consul-server-2	1/1	Running	0
9m22s			
polaris-mongodb-0	2/2	Running	0
9m22s			
polaris-mongodb-1	2/2	Running	0
8m58s			
polaris-mongodb-2	2/2	Running	0
8m34s			
polaris-ui-5df7687dbd-trcnf	1/1	Running	0
3m18s			
polaris-vault-0	1/1	Running	0

9m18s			
polaris-vault-1	1/1	Running	0
9m18s			
polaris-vault-2	1/1	Running	0
9m18s			
public-metrics-7b96476f64-j88bw	1/1	Running	0
5m48s			
storage-backend-metrics-5fd6d7cd9c-vc4j	1/1	Running	0
5m59s			
storage-provider-bb85ff965-m7qrq	1/1	Running	0
5m25s			
telegraf-ds-4zqgz	1/1	Running	0
3m36s			
telegraf-ds-cp9x4	1/1	Running	0
3m36s			
telegraf-ds-h4n59	1/1	Running	0
3m36s			
telegraf-ds-jnp2q	1/1	Running	0
3m36s			
telegraf-ds-pdz5j	1/1	Running	0
3m36s			
telegraf-ds-znqtp	1/1	Running	0
3m36s			
telegraf-rs-rt64j	1/1	Running	0
3m36s			
telemetry-service-7dd9c74bfc-sfkzt	1/1	Running	0
6m19s			
tenancy-d878b7fb6-wf8x9	1/1	Running	0
6m37s			
traefik-6548496576-5v2g6	1/1	Running	0
98s			
traefik-6548496576-g82pq	1/1	Running	0
3m8s			
traefik-6548496576-psn49	1/1	Running	0
38s			
traefik-6548496576-qrkfd	1/1	Running	0
2m53s			
traefik-6548496576-srs6r	1/1	Running	0
98s			
trident-svc-679856c67-78kbt	1/1	Running	0
5m27s			
vault-controller-747d664964-xmn6c	1/1	Running	0
7m37s			




Each pod should have a status of Running. It might take several minutes before the system pods are deployed.

20. When all pods are running, run the following command to retrieve the one-time password. In the YAML version of the output, check the `status.deploymentState` field for the deployed value, and then copy the `status.uuid` value. The password is ACC- followed by the UUID value. (ACC-[UUID]).

```
root@abhinav-ansible# oc get acc -o yaml -n netapp-acc-operator
```


21. In a browser, navigate to the URL by using the FQDN that you had provided.
22. Log in using the default user name, which is the email address provided during the installation and the one-time password ACC-[UUID].



If you enter an incorrect password three times, then the administrator account is locked for 15 minutes.

23. Change the password and proceed.






For more information about the Astra Control Center installation, see the [Astra Control Center Installation overview](#) page.


## Set up Astra Control Center

After you install Astra Control Center, log into the UI, upload the license, add clusters, manage storage, and add buckets.


1. On the home page under Account, go to the License tab and select Add License to upload the Astra license.




2. Before adding the OpenShift cluster, create an Astra Trident Volume snapshot class from the OpenShift web console. The Volume snapshot class is configured with the `csi.trident.netapp.io` driver.



3. To add the Kubernetes cluster, go to Clusters on the home page and click Add Kubernetes Cluster. Then upload the kubeconfig file for the cluster and provide a credential name. Click Next.



4. The existing storage classes are discovered automatically. Select the default storage class, click Next, and then click Add cluster.




- The cluster is added in few minutes. To add additional OpenShift Container Platform clusters, repeat steps 1–4.



To add an additional OpenShift operational environment as a managed compute resource, make sure that the Astra Trident [VolumeSnapshotClass](#) objects are defined.

- To manage the storage, go to Backends, click the three dots under Actions against the backend that you would like to manage. Click Manage.



- Provide the ONTAP credentials and click Next. Review the information and click Managed. The backends should look like the following example.

Backends							
+ Add		Search		★ Managed 🔍 Discovered		1-3 of 3 entries < >	
Name ↓	State	Capacity	Throughput	Type	Cluster	Cloud	Actions
<a href="#">c190-cluster</a>	✓ Available	0.4/10.64 TiB: 3.8%	Not available yet	ONTAP 9.11.1	Not applicable	Not applicable	⋮
<a href="#">healthylife</a>	✓ Available	5.16/106.42 TiB: 4.8%	Not available yet	ONTAP 9.11.1	Not applicable	Not applicable	⋮
<a href="#">singlecvoaws</a>	✓ Available	0.07/0.62 TiB: 11.9%	Not available yet	ONTAP 9.11.1	Not applicable	Not applicable	⋮

8. To add a bucket to Astra Control, select Buckets and click Add.

Dashboard

MANAGE YOUR APPLICATIONS

Applications

Clusters

MANAGE YOUR STORAGE

Backends

**Buckets**

MANAGE YOUR ACCOUNT

Account

Activity

astradb

Buckets

+ Add

Name ↓	Description	State	Type
--------	-------------	-------	------

9. Select the bucket type and provide the bucket name, S3 server name, or IP address and S3 credential. Click Update.

Edit bucket

×

STORAGE BUCKET

Edit the access details of your existing object store bucket.

Type

Generic S3

Existing bucket name

acc-aws-bucket

Description (optional)

S3 server name or IP address

s3.us-east-1.amazonaws.com

☐

Make this bucket the default bucket for this cloud

SELECT CREDENTIALS

Astra Control requires S3 access credentials with the roles necessary to facilitate Kubernetes application data management.

Add

Use existing

Access ID

Secret key

Credential name

EDITING STORAGE BUCKETS

Edit your existing object store bucket. If the selected bucket is not currently defined as the default bucket for the cloud, you can replace the currently defined default bucket. Read more in [Storage buckets](#).

Cancel

Update ✓



In this solution, AWS S3 and ONTAP S3 buckets are both used. You can also use StorageGRID.

The Bucket state should be Healthy.

Name	Description	State	Type	Actions
acc-aws-bucket		Healthy	Generic S3	
astra-bucket	On Prem S3 Bucket	Healthy	NetApp ONTAP S3	

As a part of Kubernetes cluster registration with Astra Control Center for application-aware data management, Astra Control automatically creates role bindings and a NetApp monitoring namespace to collect metrics and logs from the application pods and worker nodes. Make one of the supported ONTAP-based storage classes the default.

After you [add a cluster to Astra Control management](#), you can install apps on the cluster (outside of Astra Control) and then go to the Apps page in Astra Control to manage the apps and their resources. For more information about managing apps with Astra, see the [App management requirements](#).

Next: [Solution validation overview](#).

## Solution validation

### Overview

Previous: [Astra Control Center installation on OpenShift Container Platform](#).

In this section, we revisit the solution with some use cases:

- Restoring a stateful application from a remote backup to another OpenShift cluster running in the cloud.
- Restoring a stateful application to the same namespace in the OpenShift cluster.
- Application mobility by cloning from one FlexPod system (OpenShift Container Platform Bare Metal) to another FlexPod system (OpenShift Container Platform on VMware).

Notably, only a few use cases are validated in this solution. This validation does not in any way represent the entire functionality of Astra Control Center.

Next: [Application recovery with remote backups](#).

### Application recovery with remote backups

Previous: [Solution validation overview](#).

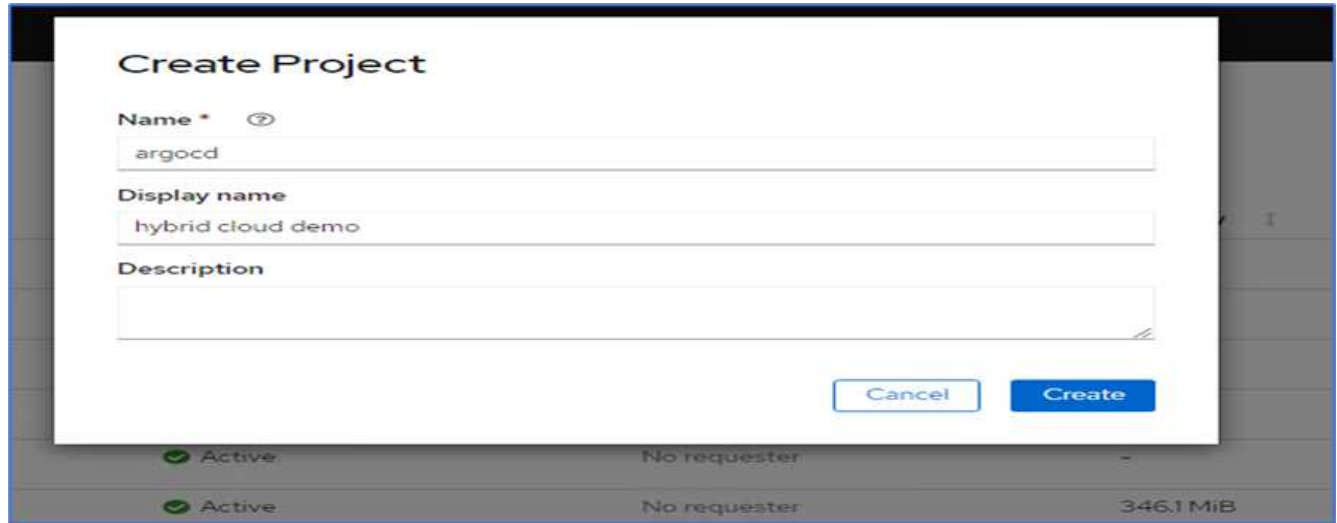
With Astra, you can take a full application-consistent backup that can be used to restore your application with its data to a different Kubernetes cluster running in an on-premises

data center or in a public cloud.


To validate a successful application recovery, simulate an on-premises failure of an application running on the FlexPod system and restore the application to a K8s cluster running in the cloud by using a remote backup.

The sample application is a pricelist application that uses MySQL for the database. To automate the deployment, we used the [Argo CD](#) tool. Argo CD is a declarative, GitOps, continuous delivery tool for Kubernetes.

1. Log into the on-premises OpenShift cluster and create a new project with the name `argocd`.



2. In the OperatorHub, search for `argocd` and select Argo CD operator.



3. Install the operator in the `argocd` namespace.

OperatorHub > Operator installation

## Install Operator

Install your Operator by subscribing to one of the update channels to keep the Operator up to date. The strategy determines either manual or automatic updates.

**Update channel \*** ⓘ

☒ alpha

**Installation mode \***

☐ All namespaces on the cluster (default)  
Operator will be available in all Namespaces.

☒ A specific namespace on the cluster  
Operator will be available in a single Namespace only.

**Installed Namespace \***

**NS** argocd

**Update approval \*** ⓘ

☒ Automatic

☐ Manual

**Install** **Cancel**

**Argo CD**  
provided by Argo CD Community

**Provided APIs**

**A** **Application**  
An Application is a group of Kubernetes resources as defined by a manifest.

**AS** **ApplicationSet**  
An ApplicationSet is a group or set of Application resources.

**AP** **AppProject**  
An AppProject is a logical grouping of Argo CD Applications.

**ACDE** **Argo CDEExport**  
ArgoCDEExport is the Schema for the argocdexports API

**ACD** **Argo CD**  
ArgoCD is the Schema for the argocds API

4. Go to the operator and click Create ArgoCD.

Project: argocd

Installed Operators > Operator details

**Argo CD**  
0.3.0 provided by Argo CD Community

Actions

Details YAML Subscription Events All instances Application ApplicationSet AppProject Argo CDEExport **Argo CD**

**ArgoCDs** **Create ArgoCD**

No operands found

Operands are declarative components used to define the behavior of the application.

5. To deploy the Argo CD instance in the argocd project, provide a name and click Create.

Project: argocd ▾

[Argo CD](#) > Create ArgoCD

## Create ArgoCD

Create by completing the form. Default values may be provided by the Operator authors.

Configure via: ☒ Form view ☐ YAML view

**Note:** Some fields may not be represented in this form view. Please select "YAML view" for full control.



**Argo CD**  
provided by Argo CD Community  
ArgoCD is the Schema for the argocds API

**Name \***

argocd-netapp

**Labels**

app=frontend

6. To log in to Argo CD, the default user is admin and the password is in a secret file with the name argocd-netapp-cluster.

Project: argocd ▾

[Secrets](#) > Secret details

## argocd-netapp-cluster

Managed by [ACD](#) argocd-netapp

[Add Secret to workload](#) [Actions](#) ▾

[Details](#) [YAML](#)

### Secret details

<b>Name</b>	argocd-netapp-cluster	<b>Type</b>	Opaque
<b>Namespace</b>	argocd		
<b>Labels</b>	<a href="#">Edit</a> app.kubernetes.io/managed-by=argocd-netapp app.kubernetes.io/name=argocd-netapp-cluster app.kubernetes.io/part-of=argocd		
<b>Annotations</b>	0 annotations <a href="#">✎</a>		
<b>Created at</b>	2 minutes ago		
<b>Owner</b>	argocd-netapp		

**Data**


admin.password

.....

[Reveal values](#) [Copied](#)

7. From the side menu, select Routes > Location and click the URL for the argocd routes. Enter the user name and password.





8. Add the on-premises OpenShift cluster to Argo CD through the CLI.

```

####Login to Argo CD####
abhinav3@abhinav-ansible$ argocd-linux-amd64 login argocd-netapp-server-
argocd.apps.ocp.flexpod.netapp.com --insecure
Username: admin
Password:
'admin:login' logged in successfully
Context'argocd-netapp-server-argocd.apps.ocp.flexpod.netapp.com' updated
####List the On-Premises OpenShift cluster####
abhinav3@abhinav-ansible$ argocd-linux-amd64 cluster add
ERRO[0000] Choose a context name from:
CURRENT  NAME
CLUSTER          SERVER
*          default/api-ocp-flexpod-netapp-com:6443/abhinav3
api-ocp-flexpod-netapp-com:6443
https://api.ocp.flexpod.netapp.com:6443
          default/api-ocp1-flexpod-netapp-com:6443/abhinav3
api-ocp1-flexpod-netapp-com:6443
https://api.ocp1.flexpod.netapp.com:6443
####Add On-Premises OpenShift cluster###
abhinav3@abhinav-ansible$ argocd-linux-amd64 cluster add default/api-
ocp1-flexpod-netapp-com:6443/abhinav3
WARNING: This will create a service account `argocd-manager` on the
cluster referenced by context `default/api-ocp1-flexpod-netapp-
com:6443/abhinav3` with full cluster level admin privileges. Do you want
to continue [y/N]? y
INFO[0002] ServiceAccount "argocd-manager" already exists in namespace
"kube-system"
INFO[0002] ClusterRole "argocd-manager-role" updated
INFO[0002] ClusterRoleBinding "argocd-manager-role-binding" updated
Cluster 'https://api.ocp1.flexpod.netapp.com:6443' added

```

9. In the ArgoCD UI, click NEW APP and enter the details about the app name and code repository.

CREATE

CANCEL

EDIT AS YAML

GENERAL

Application Name

pricelist

Project

default

SYNC POLICY

Manual

SYNC OPTIONS

☐ SKIP SCHEMA VALIDATION

☒ AUTO-CREATE NAMESPACE

☐ PRUNE LAST

☐ APPLY OUT OF SYNC ONLY

☐ RESPECT IGNORE DIFFERENCES

PRUNE PROPAGATION POLICY: foreground

☐ REPLACE ⚠️

☐ RETRY

SOURCE

Repository URL

https://github.com/netapp-abhinav/demo/

GIT

Revision

main

Branches

Path

pricelists/

10. Enter the OpenShift cluster where the app will be deployed along with the namespace.

DESTINATION

Cluster URL


https://api.ocp1.flexpod.netapp.com:6443

URL


Namespace

pricelist


11. To deploy the app on the on-premises OpenShift cluster, click SYNC.




12. In the OpenShift Container Platform console, go to Project Pricelist, and, under Storage, verify the name and size of the PVC.




13. Log into System Manager and verify the PVC.




14. After the Pods are running, select Networking > Routes from the side menu, and click the URL under Location.




15. The Pricelist app homepage is displayed.




16. Create a few records on the web page.



17. The app is discovered in Astra Control Center. To manage the app, go to Applications > Discovered, select the Pricelist app, and click Manage Applications under Actions.




18. Click the Pricelist app and select Data Protection. At this point, there should be no snapshots or backups. Click Create Snapshot to create an on-demand snapshot.



NetApp Astra Control Center supports both on-demand and scheduled snapshots and backups.

19. After the snapshot is created and the State is healthy, create a remote backup using that snapshot. This backup is stored in the S3 bucket.



20. Select the AWS S3 bucket and initiate the backup operation.

**Back up namespace application**

STEP 1/2: DETAILS

✕

**BACKUP DETAILS**

Snapshot (optional)  
pricelist-snapshot-20220614123756

Name  
pricelist-backup-20220614123837

**BACKUP DESTINATION**

Bucket  
acc-aws-bucket - AWS S3 bucket for ACC Available Default

**OVERVIEW**

**Application backups**  
Astra Control can take a backup of your application configuration and persistent storage. Persistent storage backups are transferred to your object store. Enter a backup name to get started.

- Namespace application pricelist
- Namespace pricelist
- Cluster onprem-ocp-vmware

Cancel

Next

21. The backup operation should create a folder with multiple objects in the AWS S3 bucket.

Amazon S3 > Buckets > acc-aws-bucket > 04330ccb-f13e-4eef-8f52-755f56aa3a3f/

Copy S3 URI

04330ccb-f13e-4eef-8f52-755f56aa3a3f/

Objects

Properties

**Objects (5)**

Objects are the fundamental entities stored in Amazon S3. You can use [Amazon S3 inventory](#) to get a list of all objects in your bucket. For others to access your objects, you'll need to explicitly grant them permissions. [Learn more](#)

↻

Copy S3 URI

Copy URL

Download

Open

Delete

Actions

Create folder

Upload

Find objects by prefix

<input type="checkbox"/>	Name	Type	Last modified	Size	Storage class
<input type="checkbox"/>	config	-	June 14, 2022, 05:39:19 (UTC-07:00)	155.0 B	Standard
<input type="checkbox"/>	data/	Folder	-	-	-
<input type="checkbox"/>	index/	Folder	-	-	-
<input type="checkbox"/>	keys/	Folder	-	-	-
<input type="checkbox"/>	snapshots/	Folder	-	-	-

22. When the remote backup is complete, simulate a disaster on the on-premises by stopping the storage virtual machine (SVM) that hosts the backing volume for the PV.

**ONTAP System Manager**

Search actions, objects, and pages

**DASHBOARD**

**STORAGE**

Overview

Volumes

LUNs

Consistency Groups

**Storage VMs**

+ Add

Infra


<input type="checkbox"/>	Name	State	Subtype	Configured Protocols	IPspace
<input type="checkbox"/>	Infra_SVM	stopped	default		Default

23. Refresh the webpage to confirm the outage. The webpage is unavailable.




As expected, the website is down, so let's quickly recover the app from the remote backup by using Astra to the OpenShift cluster running in AWS.

24. In Astra Control Center, click the Pricelist app and select Data Protection > Backups. Select the backup, and click Restore Application under Action.




25. Select `ocp-aws` as the destination cluster and give a name to the namespace. Click the on-demand backup, Next, and then Restore.






26. A new app with the name `pricelist-app` is provisioned on the OpenShift cluster running in AWS.



27. Verify the same in the OpenShift web console.



28. After all the pods under the `pricelist-aws` project are running, go to Routes and click the URL to launch the web page.



This process validates that the pricelist application has been successfully restored and that data integrity has been maintained on the OpenShift cluster running seamlessly on AWS with the help of Astra Control Center.

## Data protection with Snapshot copies and application mobility for DevTest


This use case consists of two parts, as described the following sections.

### Part 1


With Astra Control Center, you can take application-aware snapshots for local data protection. If you accidentally delete or corrupt your data, you can revert your applications and associated data to a known good state using a previously recorded snapshot.

In this scenario, a development and testing (DevTest) team deploys a sample stateful application (blog site) that is a Ghost blog application, adds some content, and upgrades the app to the latest version available. The Ghost application uses SQLite for the database. Before upgrading the application, a snapshot (on-demand) is taken using Astra Control Center for data protection. The detailed steps are as follows:


1. Deploy the sample blogging app and sync it from ArgoCD.




2. Log into the first OpenShift cluster, go to Project, and enter Blog in the search bar.




3. From the side menu, select Networking > Routes and click the URL.



4. The blog home page is displayed. Add some content to the blog site and publish it.



5. Go to Astra Control Center. First manage the app from the Discovered tab and then take a Snapshot copy.



You can also protect your apps by creating snapshots, backups, or both at a defined schedule. For more information, see [Protect apps with snapshots and backups](#).

6. After the On-Demand snapshot is created successfully, upgrade the app to the latest version. The current image version is `ghost: 3.6-alpine` and the target version is `ghost:latest`. To upgrade the app, make changes directly to the Git repository and sync them to Argo CD.



7. You can see that the direct upgrade to the latest version is not supported due to the blog site being down and the entire application being corrupted.

Project: blog ▾

Pods ▸ Pod details

**myblog-5f899f7b76-zv7rq** CrashLoopBackOff

Details Metrics YAML Environment **Logs** Events Terminal

Log stream ended. myblog ▾ Current log ▾

```
34 lines
[2022-06-11 12:54:05] +[36mINFO+[39m Creating database backup
[2022-06-11 12:54:05] +[36mINFO+[39m Database backup written to: /var/lib/ghost/content/data/astra.ghost.2022-06-11-12-54-05.json
[2022-06-11 12:54:05] +[36mINFO+[39m Running migrations.
[2022-06-11 12:54:06] +[36mINFO+[39m Rolling back: Unable to run migrations.
[2022-06-11 12:54:06] +[36mINFO+[39m Rollback was successful.
[2022-06-11 12:54:06] +[31mERROR+[39m Unable to run migrations
+[[31m
+[[31mUnable to run migrations+[[39m

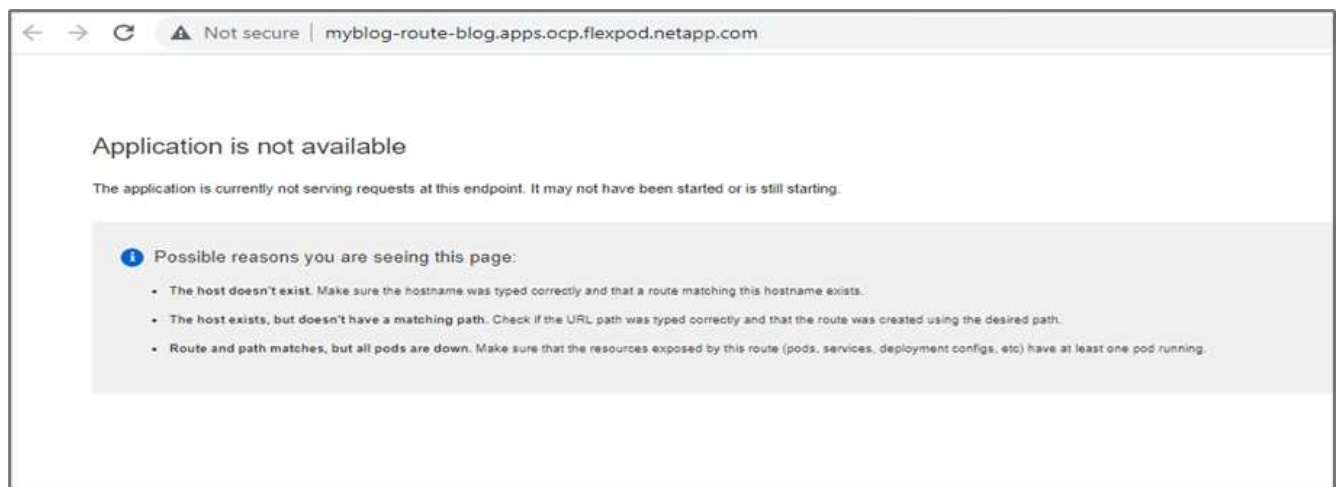
+[[37mYou must be on the latest v3.x to update across major versions - https://ghost.org/docs/update/" +[[39m
+[[33mRun 'ghost update v3' to get the latest v3.x version, then run 'ghost update' to get to the latest." +[[39m

+[[1m+[[37mError ID: +[[39m+[[22m
+[[90m93b99ce0-e985-11ec-9301-7d29b2c73999+[[39m


+[[90m-----+[[39m

+[[90mInternalServerError: Unable to run migrations
at /var/lib/ghost/versions/5.2.2/node_modules/knex-migrator/lib/index.js:1032:19
at up (/var/lib/ghost/versions/5.2.2/core/server/data/migrations/utils/migrations.js:118:19)
at Object.up (/var/lib/ghost/versions/5.2.2/core/server/data/migrations/utils/migrations.js:54:19)
at /var/lib/ghost/versions/5.2.2/node_modules/knex-migrator/lib/index.js:982:33
at /var/lib/ghost/versions/5.2.2/node_modules/knex/lib/execution/transaction.js:221:22+[[39m
+[[39m
[2022-06-11 12:54:06] +[[35mWARN+[[39m Ghost is shutting down
[2022-06-11 12:54:06] +[[35mWARN+[[39m Ghost has shut down
[2022-06-11 12:54:06] +[[35mWARN+[[39m Your site is now offline
[2022-06-11 12:54:06] +[[35mWARN+[[39m Ghost was running for a few seconds
```

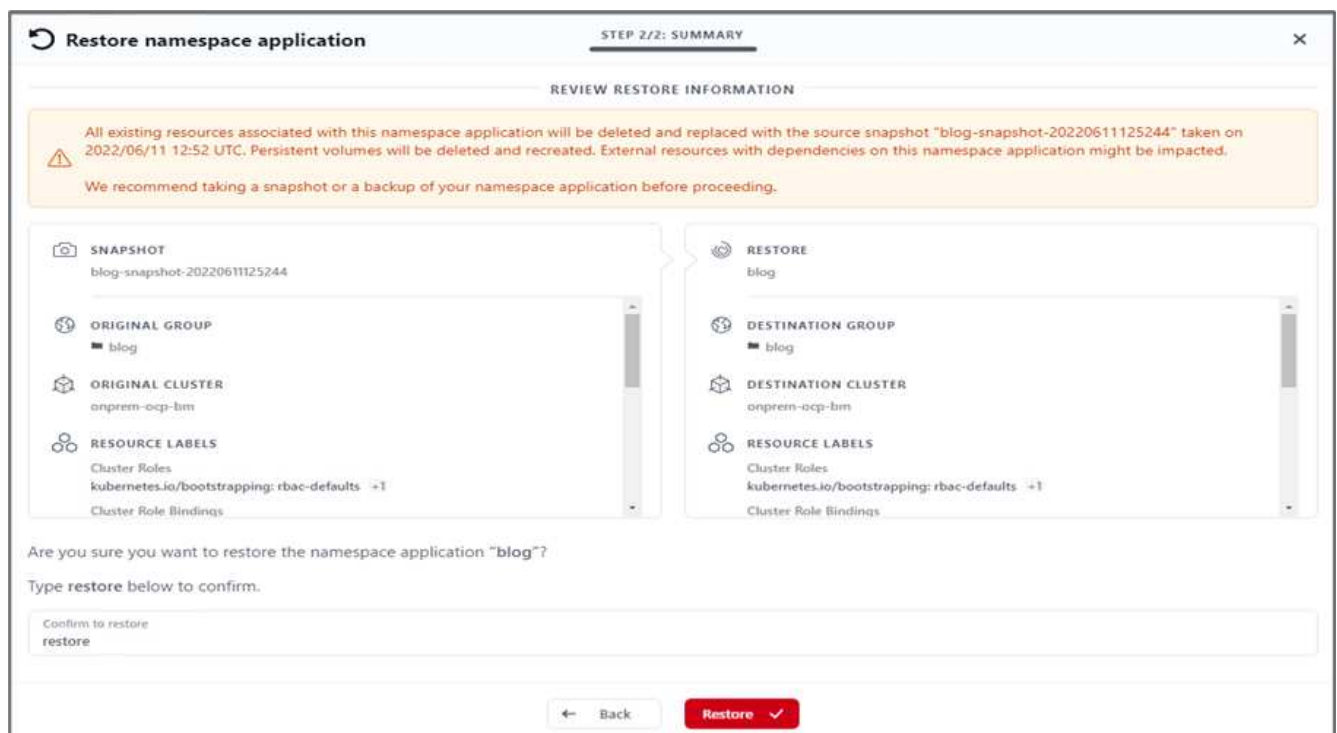
8. To confirm the unavailability of the blog site, refresh the URL.




9. Restore the app from the snapshot.



10. The app is restored on the same OpenShift cluster.



11. The app restore process starts immediately.



12. In few minutes, the app is restored successfully from the available snapshot.

Applications

Actions

+ Define

All clusters

blog

Managed

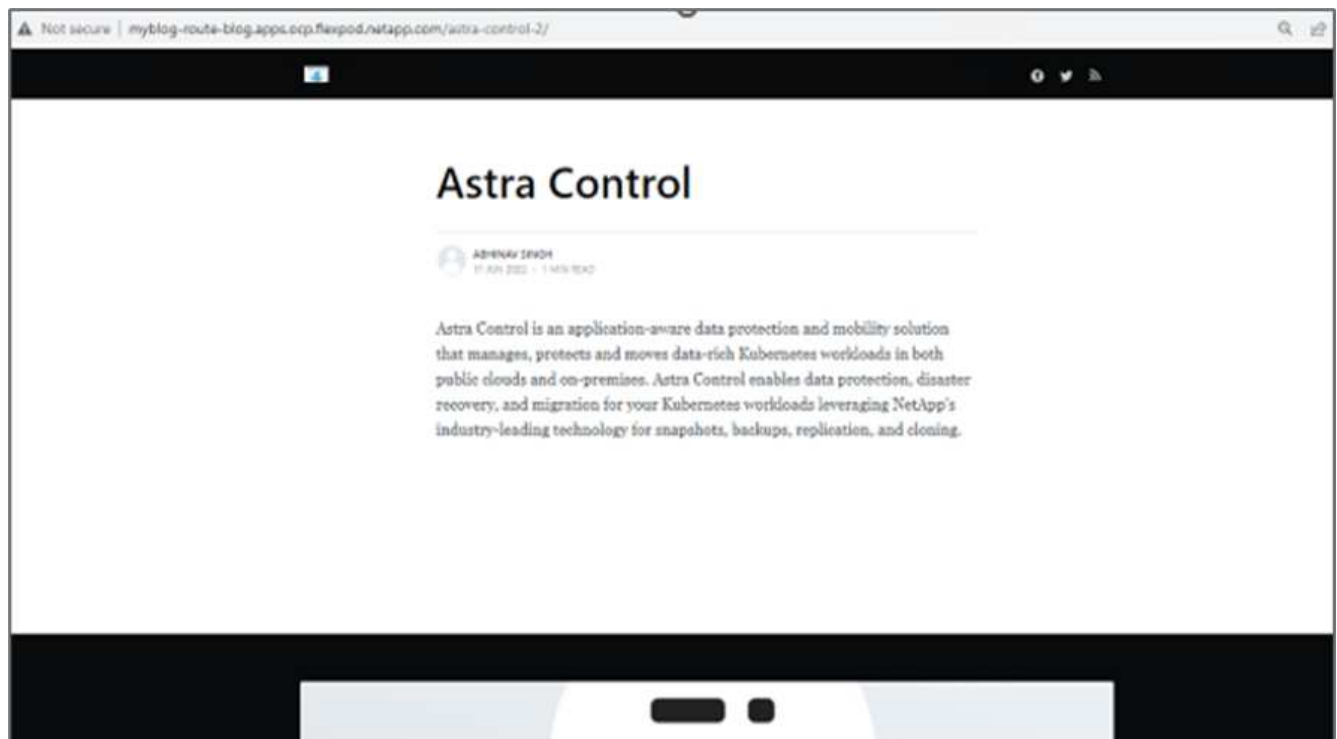
Discovered 3

Ignored

1-1 of 1 entries

<input type="checkbox"/>	Name	State	Protection	Cluster	Group	Discovered	Actions
<input type="checkbox"/>	<a href="#">blog</a>	<span>Healthy</span>	<span>Partially protected</span>	<span>onprem-ocp-bm</span>	blog	2022/06/11 12:34 UTC	<div></div>

13. To see whether the webpage is available, refresh the URL.



With the help of Astra Control Center, a DevTest team can successfully recover a blog site app and its associated data using the snapshot.

## Part 2

With Astra Control Center, you can move an entire application along with its data from one Kubernetes cluster to another, no matter where the clusters are located (on-premises or in the cloud).

1. The DevTest team initially upgrades the app to the supported version (`ghost-4.6-alpine`) before upgrading to the final version (`ghost-latest`) to make it production ready. They then post an upgrade the app that is cloned to the production OpenShift cluster running on a different FlexPod system.
2. At this point, the app is upgraded to the latest version and ready to be cloned to the production cluster.



Project: blog ▾


Pods > Pod details

**myblog-55ffd9f658-tkbfq** Running

Details Metrics YAML Environment Logs Events Terminal


```
180     ports:
181     - containerPort: 2368
182       protocol: TCP
183     imagePullPolicy: Always
184     volumeMounts:
185     - name: content
186       mountPath: /var/lib/ghost/content
187     - name: kube-api-access-t2sdz
188       readOnly: true
189       mountPath: /var/run/secrets/kubernetes.io/serviceaccount
190     terminationMessagePolicy: File
191     image: 'ghost:latest'
192   serviceAccount: default
193   volumes:
194   - name: content
195     persistentVolumeClaim:
196       claimName: blog-content
```

3. To verify the new theme, refresh the blog site.

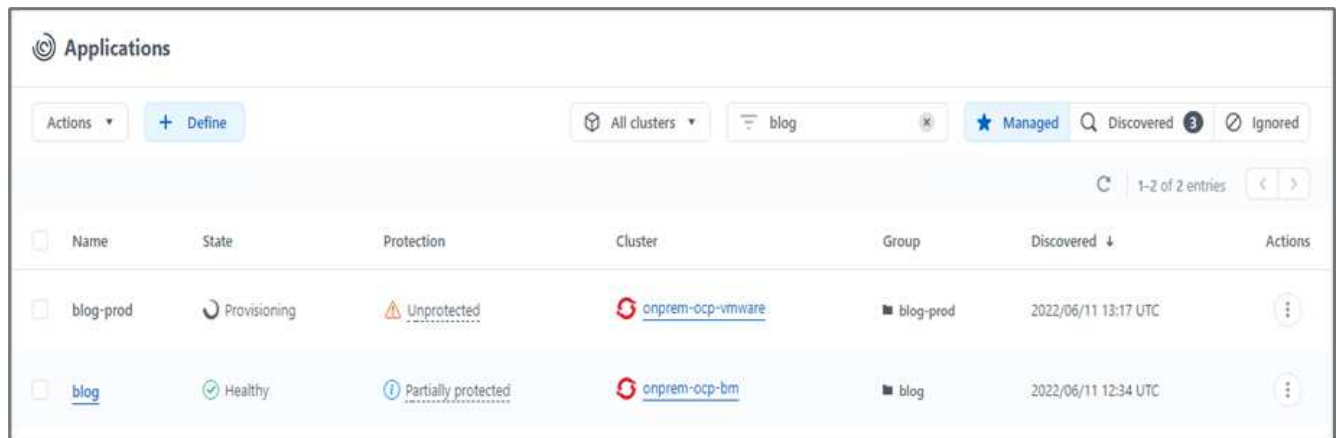


4. From Astra Control Center, clone the app to the other production OpenShift cluster running on VMware vSphere.






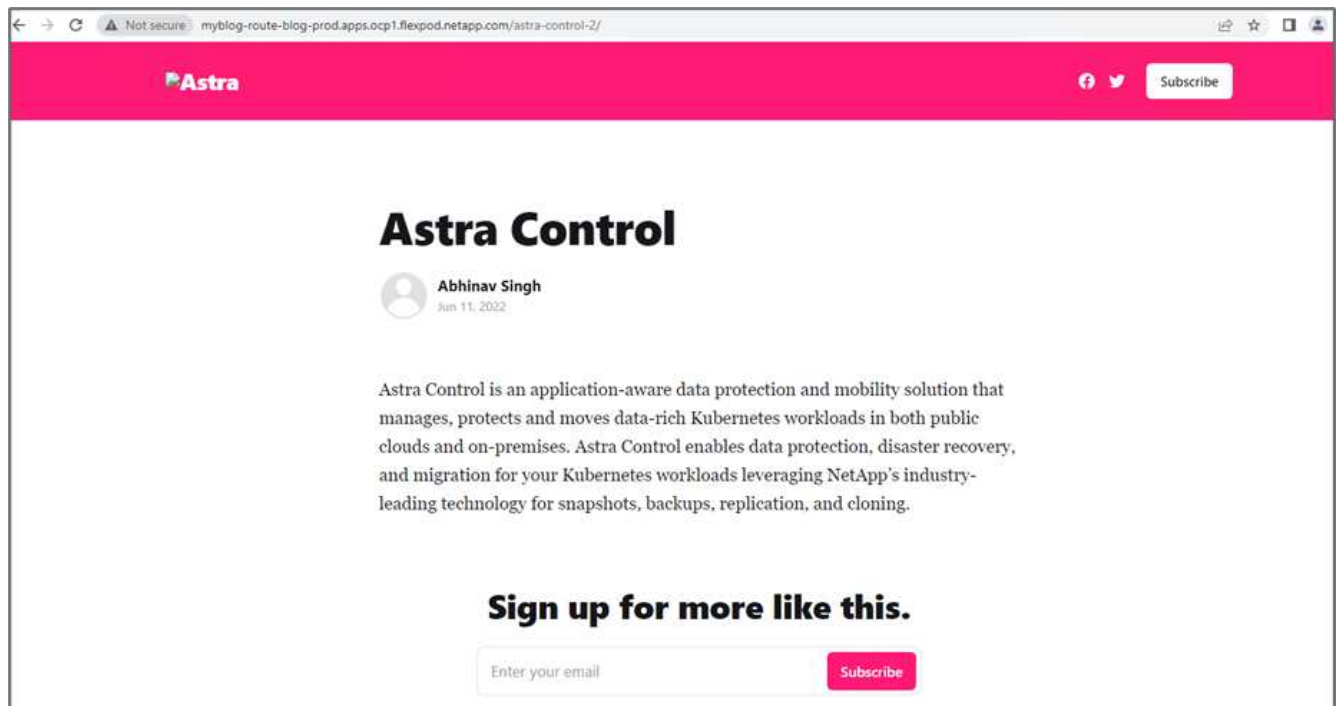
A new application clone is now provisioned in the production OpenShift cluster.



5. Log into the production OpenShift cluster and search for the project blog.



6. From the side menu, select Networking > Routes and click the URL under Location. The same homepage with the content is displayed.



This concludes the Astra Control Center solution validation. You can now clone an entire application and its data from one Kubernetes cluster to another no matter where the Kubernetes cluster is located.

[Next: Conclusion.](#)

## Conclusion

[Previous: Application recovery with remote backups.](#)

In this solution, we implemented a protection plan for containerized applications running on FlexPod and AWS using the NetApp Astra portfolio. NetApp Astra Control Center and Astra Trident, along with Cloud Volumes ONTAP, Red Hat OpenShift, and the FlexPod infrastructure, formed the core components of this solution.

We demonstrated the protection of applications by capturing snapshots, and we executed full-copy backups to restore apps across different K8s clusters running in the cloud and on-premises environments.

We also demonstrated the cloning of applications across K8s clusters, thereby enabling customers to migrate their apps to their choice of K8s clusters at their desired locations.

FlexPod has constantly evolved so that its customers can modernize their applications and business delivery processes. With this solution, FlexPod customers can confidently build their BCDR plan for their cloud-native apps with the public cloud as a location for a transient or full-time DR plan while keeping the cost of the solution low.

Astra Control enables you to move an entire application along with its data from one Kubernetes cluster to another, no matter where the clusters are located. It can also help you accelerate deployment, operations, and protection for your cloud-native applications.

## Troubleshooting

For troubleshooting guidance, see the [online documentation](#).

## Where to find additional information

To learn more about the information that is described in this document, review the following documents and/or websites:

- FlexPod Home Page

<https://www.flexpod.com>

- Cisco validated Design and deployment guides for FlexPod

<https://www.cisco.com/c/en/us/solutions/design-zone/data-center-design-guides/flexpod-design-guides.html>

- FlexPod deployment with Infrastructure as code for VMware using Ansible

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_m6\\_esxi7u2.html#AnsibleAutomationWorkflowandSolutionDeployment](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_m6_esxi7u2.html#AnsibleAutomationWorkflowandSolutionDeployment)

- FlexPod deployment with Infrastructure as code for Red Hat OpenShift Bare Metal using Ansible

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_iac\\_redhat\\_openshift.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_iac_redhat_openshift.html)

- Cisco UCS Hardware and Software Interoperability Tool

<http://www.cisco.com/web/techdoc/ucs/interoperability/matrix/matrix.html>

- Cisco Intersight Data Sheet

<https://intersight.com/help/saas/home>

- NetApp Astra documentation

<https://docs.netapp.com/us-en/astra-control-center/index.html>

- NetApp Astra Control Center

<https://docs.netapp.com/us-en/astra-control-center/index.html>

- NetApp Astra Trident

<https://docs.netapp.com/us-en/trident/index.html>

- NetApp Cloud Manager

[https://docs.netapp.com/us-en/occm/concept\\_overview.html](https://docs.netapp.com/us-en/occm/concept_overview.html)

- NetApp Cloud Volumes ONTAP

[https://docs.netapp.com/us-en/occm/task\\_getting\\_started\\_aws.html](https://docs.netapp.com/us-en/occm/task_getting_started_aws.html)

- Red Hat OpenShift

<https://www.openshift.com/>

- NetApp Interoperability Matrix Tool

<http://support.netapp.com/matrix/>

## Version history

Version	Date	Document version history
Version 1.0	July 2022	Release for ACC 22.04.0.

## Copyright information

Copyright © 2024 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

LIMITED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (b)(3) of the Rights in Technical Data -Noncommercial Items at DFARS 252.227-7013 (FEB 2014) and FAR 52.227-19 (DEC 2007).

Data contained herein pertains to a commercial product and/or commercial service (as defined in FAR 2.101) and is proprietary to NetApp, Inc. All NetApp technical data and computer software provided under this Agreement is commercial in nature and developed solely at private expense. The U.S. Government has a non-exclusive, non-transferrable, nonsublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b) (FEB 2014).

## Trademark information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.