



Hybrid Cloud with Self-Managed Components (On- premises/AWS/GCP/Azure)

NetApp Solutions

NetApp
May 17, 2024

Table of Contents

- NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads 1
 - Overview 1
 - NetApp Solution with Red Hat OpenShift Container platform workloads in Hybrid Cloud 3
 - Deploy and configure the Red Hat OpenShift Container platform on AWS 5
 - Data protection using Astra Control Center 14
 - Data migration using Astra Control Center 17

NetApp Hybrid Multicloud solutions for Red Hat OpenShift Container workloads

NetApp is seeing a significant increase in customers modernizing their legacy enterprise applications and building new applications using containers and orchestration platforms built around Kubernetes. Red Hat OpenShift Container Platform is one example that we see adopted by many of our customers.

Overview

As more and more customers begin adopting containers within their enterprises, NetApp is perfectly positioned to help serve the persistent storage needs of their stateful applications and classic data management needs such as data protection, data security, and data migration. However, these needs are met using different strategies, tools, and methods.

NetApp ONTAP based storage options listed below, deliver security, data protection, reliability, and flexibility for containers and Kubernetes deployments.

- Self-managed storage in on-premises:
 - NetApp Fabric Attached Storage (FAS), NetApp All Flash FAS Arrays (AFF), NetApp All SAN Array (ASA) and ONTAP Select
- Provider-managed storage in on-premises:
 - NetApp Keystone provides Storage as a Service (STaaS)
- Self-managed storage in the cloud:
 - NetApp Cloud Volumes ONTAP(CVO) provide self managed storage in the hyperscalers
- Provider-managed storage in the cloud:
 - Cloud Volumes Service for Google Cloud (CVS), Azure NetApp Files (ANF), Amazon FSx for NetApp ONTAP offer fully managed storage in the hyperscalers

ONTAP feature highlights



<p>Storage Administration</p> <ul style="list-style-type: none"> Multi-tenancy FlexVol & FlexGroup LUN Quotas ONTAP CLI & API System Manager & BlueXP 	<p>Performance & Scalability</p> <ul style="list-style-type: none"> FlexCache FlexClone nconnect, session trunking, multipathing Scale-out clusters
<p>Availability & Resilience</p> <ul style="list-style-type: none"> Multi-AZ HA deployment (MetroCluster) SnapShot & SnapRestore SnapMirror SnapMirror Business Continuity (MetroCluster) SnapMirror Cloud 	<p>Access Protocols</p> <ul style="list-style-type: none"> NFS –v3, v4, v4.1, v4.2 SMB – v2, v3 iSCSI Multi-protocol access
<p>Storage Efficiency</p> <ul style="list-style-type: none"> Deduplication & Compression Compaction Thin provisioning Data Tiering (Fabric Pool) 	<p>Security & Compliance</p> <ul style="list-style-type: none"> Fpolicy & Vscan Active Directory integration LDAP & Kerberos Certificate based authentication

NetApp BlueXP enables you to manage all of your storage and data assets from a single control plane/interface.

You can use BlueXP to create and administer cloud storage (for example, Cloud Volumes ONTAP and Azure NetApp Files), to move, protect, and analyze data, and to control many on-prem and edge storage devices.

NetApp Astra Trident is a CSI Compliant Storage Orchestrator that enable quick and easy consumption of persistent storage backed by a variety of the above-mentioned NetApp storage options. It is an open-source software maintained and supported by NetApp.

Astra Trident CSI feature highlights



<p>CSI specific</p> <ul style="list-style-type: none"> CSI NetApp® Snapshot™ copies and volume creation from CSI Snapshot copies CSI topology Volume expansion 	<p>Security</p> <ul style="list-style-type: none"> Dynamic-export policy management iSCSI initiator-groups dynamic management iSCSI bidirectional CHAP
<p>Control</p> <ul style="list-style-type: none"> Storage and performance consumption Monitoring Volume Import Cross Namespace Volume Access 	<p>Installation methods</p> <ul style="list-style-type: none"> Binary Helm chart Operator GitOps
<p>Choose your access mode</p> <ul style="list-style-type: none"> RWO (ReadWriteOnce, i.e 1↔1) RWX (ReadWriteMany, i.e 1↔n) ROX (ReadOnlyMany) RWOP (ReadWriteOnce POD) 	<p>Choose your protocol</p> <ul style="list-style-type: none"> NFS SMB iSCSI

Business critical container workloads need more than just persistent volumes. Their data management requirements require protection and migration of the application kubernetes objects as well.



Application data includes kubernetes objects in addition to the user data: Some examples are as follows:

- kubernetes objects such as pods specs, PVCs, deployments, services
- custom config objects such as config maps and secrets
- persistent data such as Snapshot copies, backups, clones
- custom resources such as CRs and CRDs

NetApp Astra Control, available as both fully-managed and self-managed software, provides orchestration for robust application data management. Refer to the [Astra documentation](#) for additional details on the Astra family of products.

This reference documentation provides validation of migration and protection of container-based applications, deployed on RedHat OpenShift container platform, using NetApp Astra Control Center. In addition, the solution provides high-level details for the deployment and the use of Red Hat Advanced Cluster Management (ACM) for managing the container platforms. The document also highlights the details for the integration of NetApp storage with Red Hat OpenShift container platforms using Astra Trident CSI provisioner. Astra Control Center is deployed on the hub cluster and is used to manage the container applications and their persistent storage lifecycle. Finally, it provides a solution for replication and failover and fail-back for container workloads on managed Red Hat OpenShift clusters in AWS (ROSA) using Amazon FSx for NetApp ONTAP (FSxN) as persistent storage.

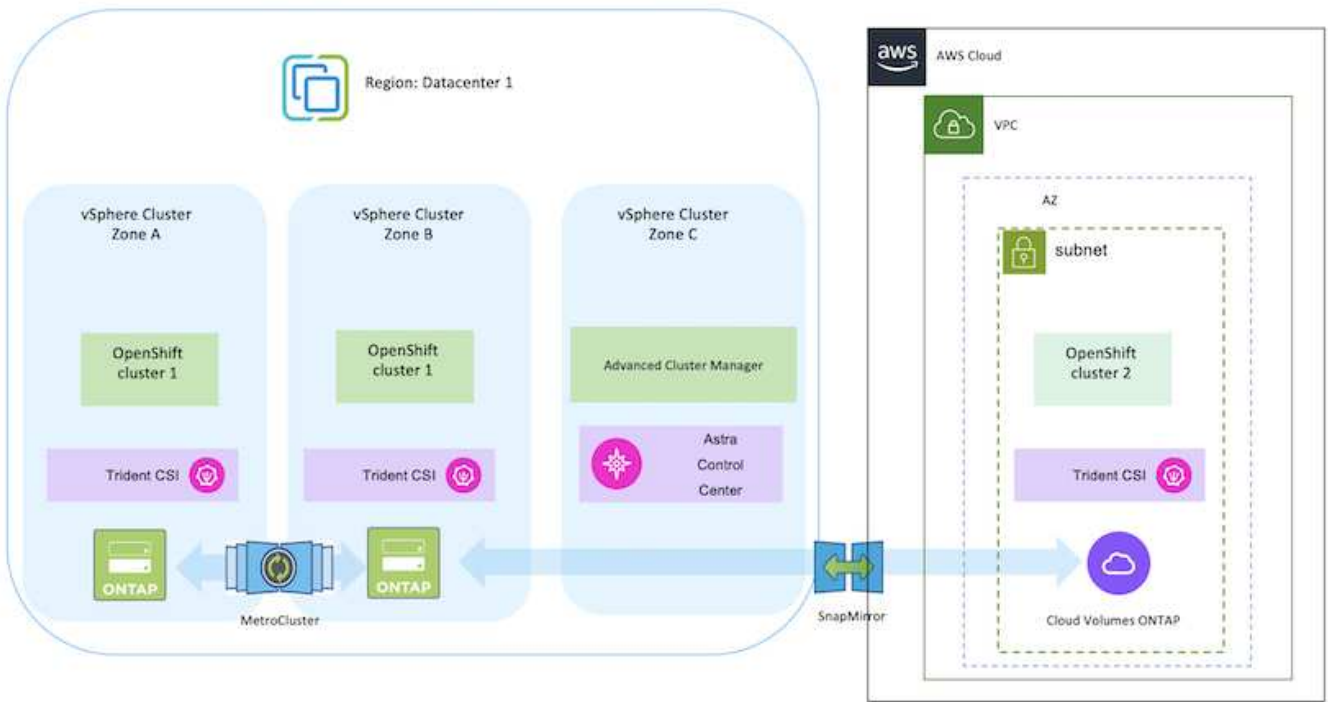
NetApp Solution with Red Hat OpenShift Container platform workloads in Hybrid Cloud

Customers may be at a point in their modernization journey when they are ready to move some select workloads or all workloads from their data centers to the cloud. They may choose to use self-managed OpenShift containers and self-managed NetApp storage in the cloud for various reasons. They should plan and deploy the Red Hat OpenShift container platform (OCP) in the cloud for a successful production-ready environment for migrating their container workloads from their data centers. Their OCP clusters can be deployed on VMware or Bare Metal in their data centers and on AWS, Azure or Google Cloud in the cloud environment.

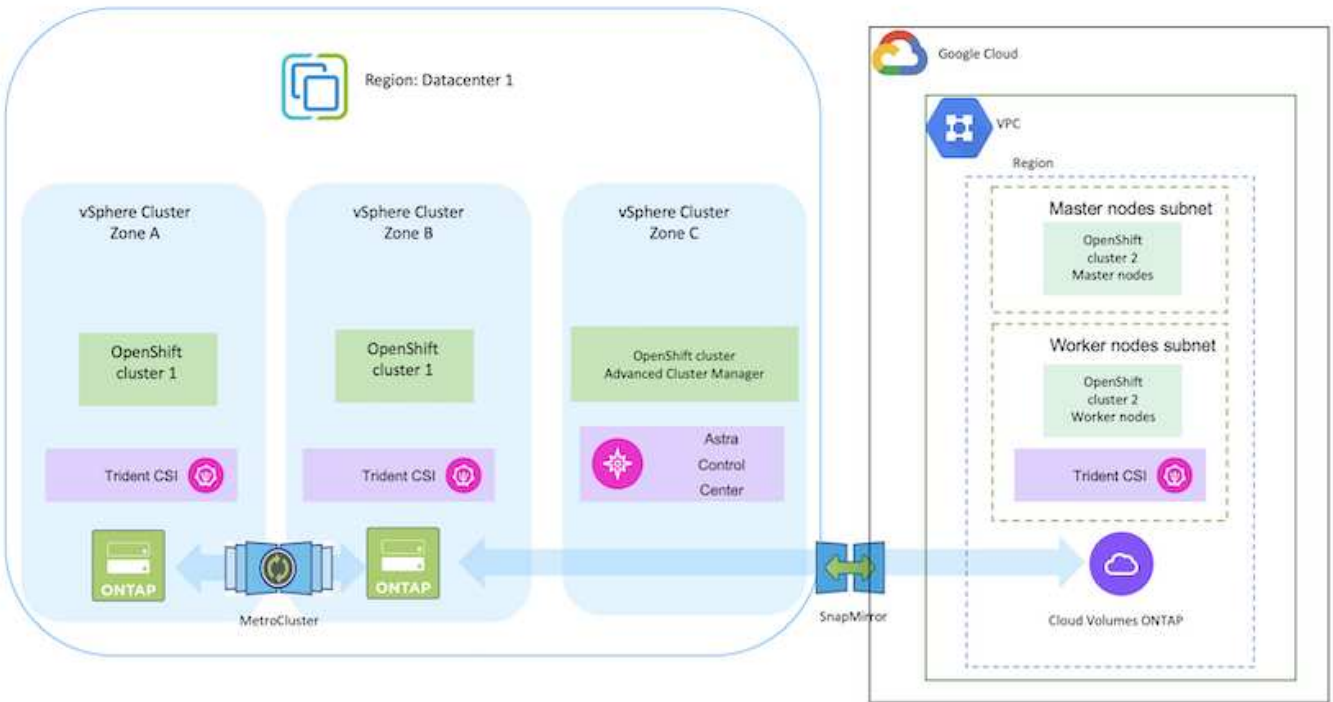
NetApp Cloud Volumes ONTAP storage delivers data protection, reliability, and flexibility for container deployments in AWS, Azure and in Google Cloud. Astra Trident serves as the dynamic storage provisioner to consume the persistent Cloud Volumes ONTAP storage for customers' stateful applications. Astra Control Center can be used to orchestrate the many data management requirements of stateful applications such as data protection, migration, and business continuity.

Data protection and migration solution for OpenShift Container workloads in a hybrid cloud using Astra Control Center


On-premises and AWS



On-premises and Google Cloud



On-premises and Azure Cloud




Deploy and configure the Red Hat OpenShift Container platform on AWS

This section describes a high-level workflow of how to set up and manage OpenShift Clusters in AWS and deploy stateful applications on them. It shows the use of NetApp Cloud Volumes ONTAP storage with the help of Astra Trident to provide persistent volumes. Details are provided about the use of Astra Control Center to perform data protection and migration activities for the stateful applications.



There are several ways of deploying Red Hat OpenShift Container platform clusters on AWS. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the [resources section](#).

Here is a diagram that depicts the clusters deployed on AWS and connected to the data center using a VPN.



The setup process can be broken down into the following steps:

Install an OCP cluster on AWS from the Advanced Cluster Management.

- Create a VPC with a site-to-site VPN connection (using pfsense) to connect to the on-premises network.
- On-premises network has internet connectivity.
- Create 3 private subnets in 3 different AZs.
- Create a Route 53 private hosted zone and a DNS resolver for the VPC.

Create OpenShift Cluster on AWS from the Advanced Cluster Management (ACM) Wizard. Refer to instructions [here](#).



You can also create the cluster in AWS from the OpenShift Hybrid Cloud console. Refer [here](#) for instructions.



When creating the cluster using the ACM, you have the ability to customize the installation by editing the yaml file after filling in the details in the form view. After the cluster is created, you can ssh login to the nodes of the cluster for troubleshooting or additional manual configuration. Use the ssh key you provided during installation and the username core to login.

Deploy Cloud Volumes ONTAP in AWS using BlueXP.

- Install the connector in on-premises VMware environment. Refer to instructions [here](#).
- Deploy a CVO instance in AWS using the connector. Refer to instructions [here](#).



The connector can also be installed in the cloud environment. Refer [here](#) for additional information.

Install Astra Trident in the OCP Cluster

- Deploy Trident Operator using Helm. Refer to instructions [here](#)
- Create a backend and a storage class. Refer to instructions [here](#).

Add the OCP cluster on AWS to the Astra Control Center.

Add the OCP cluster in AWS to Astra Control Center.

Using CSI Topology feature of Trident for multi-zone architectures

Cloud providers, today, enable Kubernetes/OpenShift cluster administrators to spawn nodes of the clusters that are zone based. Nodes can be located in different availability zones within a region, or across various regions. To facilitate the provisioning of volumes for workloads in a multi-zone architecture, Astra Trident uses CSI Topology. Using the CSI Topology feature, access to volumes can be limited to a subset of nodes, based on regions and availability zones. Refer [here](#) for additional details.



Kubernetes supports two volume binding modes:

- When **VolumeBindingMode is set to Immediate** (default), Astra Trident creates the volume without any topology awareness. Persistent Volumes are created without having any dependency on the requesting pod's scheduling requirements.
- When **VolumeBindingMode set to WaitForFirstConsumer**, the creation and binding of a Persistent Volume for a PVC is delayed until a pod that uses the PVC is scheduled and created. This way, volumes are created to meet the scheduling constraints that are enforced by topology requirements.

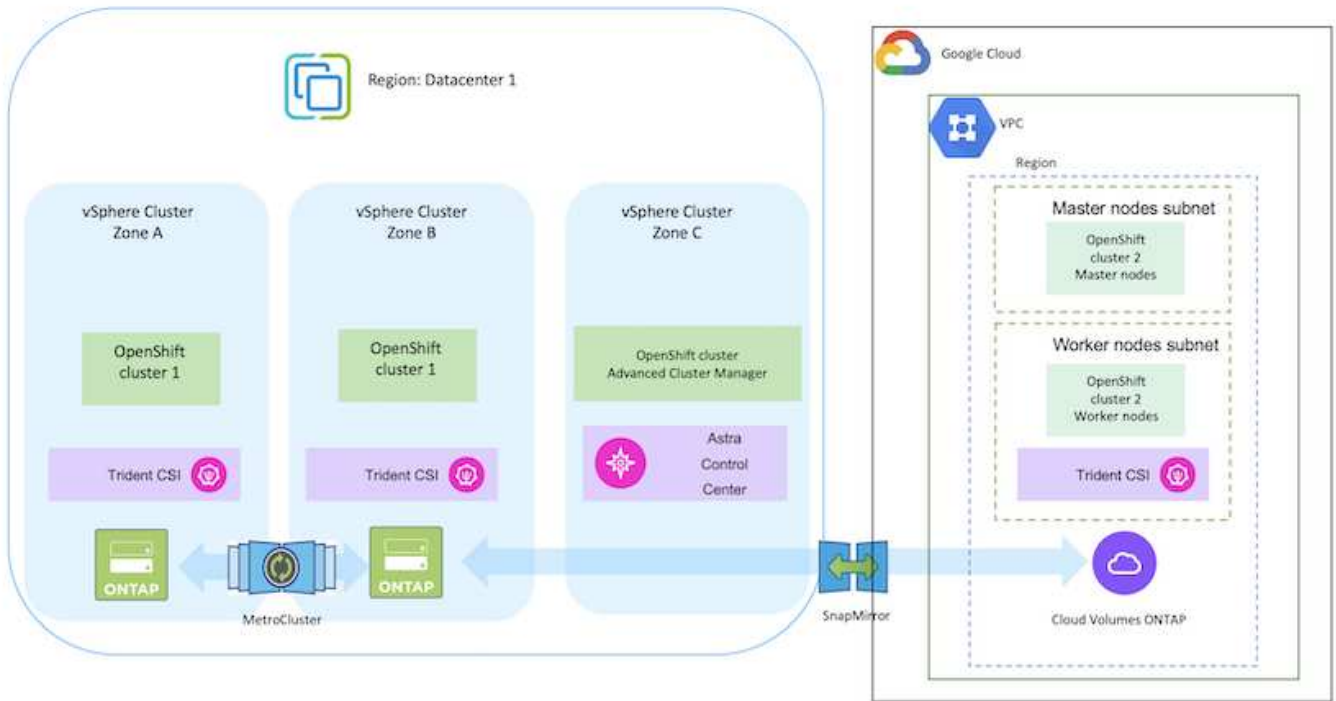
Astra Trident storage backends can be designed to selectively provision volumes based on availability zones (Topology-aware backend). For StorageClasses that make use of such a backend, a volume would only be created if requested by an application that is scheduled in a supported region/zone. (Topology-aware StorageClass)

Refer [here](#) for additional details.

Deploy and configure the Red Hat OpenShift Container platform on GCP

This section describes a high-level workflow of how to set up and manage OpenShift Clusters in GCP and deploy stateful applications on them. It shows the use of NetApp Cloud Volumes ONTAP storage with the help of Astra Trident to provide persistent volumes. Details are provided about the use of Astra Control Center to perform data protection and migration activities for the stateful applications.

Here is a diagram that shows the clusters deployed on GCP and connected to the data center using a VPN.



There are several ways of deploying Red Hat OpenShift Container platform clusters in GCP. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the [resources section](#).

The setup process can be broken down into the following steps:

Install an OCP cluster on GCP from the CLI.

- Ensure that you have met all the prerequisites stated [here](#).
- For the VPN connectivity between on-premises and GCP, a pfsense VM was created and configured. For instructions, see [here](#).
 - The remote gateway address in pfsense can be configured only after you have created a VPN gateway in Google Cloud Platform.
 - The remote network IP addresses for the Phase 2 can be configured only after the OpenShift cluster installation program runs and creates the infrastructure components for the cluster.
 - The VPN in Google Cloud can only be configured after the infrastructure components for the cluster are created by the installation program.
- Now install the OpenShift cluster on GCP.
 - Obtain the installation program and the pull secret and deploy the cluster following the steps provided in the documentation [here](#).
 - The installation creates a VPC network in Google Cloud Platform. It also creates a private zone in Cloud DNS and adds A records.
 - Use the CIDR block address of the VPC network to configure the pfsense and establish the VPN connection. Ensure firewalls are setup correctly.
 - Add A records in the DNS of the on-premises environment using the IP address in the A records of the Google Cloud DNS.
 - The installation of the cluster completes and will provide a kubeconfig file and username and password to login to the console of the cluster.

Deploy Cloud Volumes ONTAP in GCP using BlueXP.

- Install a connector in Google Cloud. Refer to instructions [here](#).
- Deploy a CVO instance in Google Cloud using the connector. Refer to instructions [here](https://docs.netapp.com/us-en/bluexp-cloud-volumes-ontap/task-getting-started-gcp.html).

Install Astra Trident in the OCP Cluster in GCP

- There are many methods to deploy Astra Trident as shown [here](#).
- For this project, Astra Trident was installed by deploying Astra Trident Operator manually using the instructions [here](#).
- Create backend and a storage classes. Refer to instructions [here](#).

Add the OCP cluster on GCP to the Astra Control Center.

- Create a separate KubeConfig file with a cluster role that contains the minimum permissions necessary for a cluster to be managed by Astra Control. The instructions can be found [here](#).
- Add the cluster to Astra Control Center following the instructions [here](#)

Using CSI Topology feature of Trident for multi-zone architectures

Cloud providers, today, enable Kubernetes/OpenShift cluster administrators to spawn nodes of the clusters that are zone based. Nodes can be located in different availability zones within a region, or across various regions. To facilitate the provisioning of volumes for workloads in a multi-zone architecture, Astra Trident uses CSI Topology. Using the CSI Topology feature, access to volumes can be limited to a subset of nodes, based on regions and availability zones. Refer [here](#) for additional details.



Kubernetes supports two volume binding modes:

- When **VolumeBindingMode is set to Immediate** (default), Astra Trident creates the volume without any topology awareness. Persistent Volumes are created without having any dependency on the requesting pod's scheduling requirements.
- When **VolumeBindingMode set to WaitForFirstConsumer**, the creation and binding of a Persistent Volume for a PVC is delayed until a pod that uses the PVC is scheduled and created. This way, volumes are created to meet the scheduling constraints that are enforced by topology requirements.

Astra Trident storage backends can be designed to selectively provision volumes based on availability zones (Topology-aware backend). For StorageClasses that make use of such a backend, a volume would only be created if requested by an application that is scheduled in a supported region/zone. (Topology-aware StorageClass)

Refer [here](#) for additional details.

Demonstration Video


[OpenShift Cluster installation on Google Cloud Platform](#)

[Importing OpenShift clusters into Astra Control Center](#)

Deploy and configure the Red Hat OpenShift Container platform on Azure

This section describes a high-level workflow of how to set up and manage OpenShift Clusters in Azure and deploy stateful applications on them. It shows the use of NetApp Cloud Volumes ONTAP storage with the help of Astra Trident/Astra Control Provisioner to provide persistent volumes. Details are provided about the use of Astra Control Center to perform data protection and migration activities for the stateful applications.

Here is a diagram that shows the clusters deployed on Azure and connected to the data center using a VPN.



There are several ways of deploying Red Hat OpenShift Container platform clusters in Azure. This high-level description of the setup provides documentation links for the specific method that was used. You can refer to the other methods in the relevant links provided in the [resources section](#).

The setup process can be broken down into the following steps:

Install an OCP cluster on Azure from the CLI.

- Ensure that you have met all the prerequisites stated [here](#).
- Create a VPN, subnets and network security groups and a private DNS zone. Create VPN gateway and site-to-site VPN Connection.
- For the VPN connectivity between on-premises and Azure, a pfsense VM was created and configured. For instructions, see [here](#).
- Obtain the installation program and the pull secret and deploy the cluster following the steps provided in the documentation [here](#).
- The installation of the cluster completes and will provide a kubeconfig file and username and password to login to the console of the cluster.

A sample install-config.yaml file is given below.

```
apiVersion: v1
baseDomain: sddc.netapp.com
compute:
- architecture: amd64
  hyperthreading: Enabled
  name: worker
  platform:
    azure:
      encryptionAtHost: false
      osDisk:
        diskSizeGB: 512
        diskType: "StandardSSD_LRS"
      type: Standard_D2s_v3
      ultraSSDCapability: Disabled
      #zones:
      #- "1"
      #- "2"
      #- "3"
  replicas: 3
controlPlane:
  architecture: amd64
  hyperthreading: Enabled
  name: master
  platform:
    azure:
      encryptionAtHost: false
      osDisk:
        diskSizeGB: 1024
        diskType: Premium_LRS
      type: Standard_D8s_v3
      ultraSSDCapability: Disabled
  replicas: 3
```

```

metadata:
  creationTimestamp: null
  name: azure-cluster
networking:
  clusterNetwork:
  - cidr: 10.128.0.0/14
    hostPrefix: 23
  machineNetwork:
  - cidr: 10.0.0.0/16
  networkType: OVNKubernetes
  serviceNetwork:
  - 172.30.0.0/16
platform:
  azure:
    baseDomainResourceGroupName: ocp-base-domain-rg
    cloudName: AzurePublicCloud
    computeSubnet: ocp-subnet2
    controlPlaneSubnet: ocp-subnet1
    defaultMachinePlatform:
      osDisk:
        diskSizeGB: 1024
        diskType: "StandardSSD_LRS"
        ultraSSDCapability: Disabled
    networkResourceGroupName: ocp-nc-us-rg
    #outboundType: UserDefinedRouting
    region: northcentralus
    resourceGroupName: ocp-cluster-ncusrg
    virtualNetwork: ocp_vnet_ncus
publish: Internal
pullSecret:

```

Deploy Cloud Volumes ONTAP in Azure using BlueXP.

- Install a connector in in Azure. Refer to instructions [here](#).
- Deploy a CVO instance in Azure using the connector. Refer to instructions [link:https://docs.netapp.com/us-en/bluexp-cloud-volumes-ontap/task-getting-started-azure.html](https://docs.netapp.com/us-en/bluexp-cloud-volumes-ontap/task-getting-started-azure.html) [here.]

Install Astra Control Provisioner in the OCP Cluster in Azure

- For this project, Astra Control Provisioner (ACP) was installed on all the clusters (on-prem cluster, on-prem cluster where Astra Control Center is deployed and the cluster in Azure). Learn more about the Astra Control Provisioner [here](#).
- Create backend and a storage classes. Refer to instructions [here](#).

Add the OCP cluster on Azure to the Astra Control Center.

- Create a separate KubeConfig file with a cluster role that contains the minimum permissions necessary for a cluster to be managed by Astra Control. The instructions can be found [here](#).
- Add the cluster to Astra Control Center following the instructions [here](#)

Using CSI Topology feature of Trident for multi-zone architectures

Cloud providers, today, enable Kubernetes/OpenShift cluster administrators to spawn nodes of the clusters that are zone based. Nodes can be located in different availability zones within a region, or across various regions. To facilitate the provisioning of volumes for workloads in a multi-zone architecture, Astra Trident uses CSI Topology. Using the CSI Topology feature, access to volumes can be limited to a subset of nodes, based on regions and availability zones. Refer [here](#) for additional details.



Kubernetes supports two volume binding modes:

- When **VolumeBindingMode is set to Immediate** (default), Astra Trident creates the volume without any topology awareness. Persistent Volumes are created without having any dependency on the requesting pod's scheduling requirements.
- When **VolumeBindingMode set to WaitForFirstConsumer**, the creation and binding of a Persistent Volume for a PVC is delayed until a pod that uses the PVC is scheduled and created. This way, volumes are created to meet the scheduling constraints that are enforced by topology requirements.

Astra Trident storage backends can be designed to selectively provision volumes based on availability zones (Topology-aware backend). For StorageClasses that make use of such a backend, a volume would only be created if requested by an application that is scheduled in a supported region/zone. (Topology-aware StorageClass)
Refer [here](#) for additional details.

Demonstration Video

[Using Astra Control for Failover and Failback of applications](#)

Data protection using Astra Control Center

This page shows the data protection options for Red Hat OpenShift Container based applications running on VMware vSphere or in the cloud using Astra Control Center (ACC).

As users take their journey of modernizing their applications with Red Hat OpenShift, a data protection strategy should be in place to protect them from accidental deletion or any other human errors. Often a protection strategy is also required for regulatory or compliance purposes to protect their data from a disaster.

The requirements of data protection varies from reverting back to a point in time copy to automatically failing over to a different fault domain without any human intervention. Many customers pick ONTAP as their preferred storage platform for their Kubernetes applications because of its rich features like multitenancy, multi-protocol, high performance and capacity offerings, replication and caching for multi-site locations, security and flexibility.


Customers may have a cloud environment setup as their data center extension, so that they can leverage the benefits of the cloud as well as be well positioned to move their workloads at a future time. For such customers, backing up of their OpenShift applications and their data to the cloud environment becomes an

inevitable choice. They can then restore the applications and the associated data either to an OpenShift cluster in the cloud or in their data center.

Backup and Restore with ACC

Application owners can review and update the applications discovered by ACC. ACC can take Snapshot copies using CSI and perform backup using the point in time Snapshot copy. Backup destination can be an object store in the cloud environment. Protection policy can be configured for scheduled backups and the number of backup versions to keep. The minimum RPO is one hour.

Restoring an application from a backup using ACC



Application specific execution hooks

Even though storage array level data protection features are available, often additional steps are needed to make backups and restores application consistent. The app-specific additional steps could be:

- before or after a Snapshot copy is created.
- before or after a backup is created.
- after restoring from a Snapshot copy or backup.

Astra Control can execute these app-specific steps coded as custom scripts called execution hooks.

NetApp's [open source project Verda](#) provides execution hooks for popular cloud-native applications to make protecting applications straightforward, robust, and easy to orchestrate. Feel free to contribute to that project if you have enough information for an application that is not in the repository.

Sample execution hook for pre-Snapshot of a redis application.

Edit execution hook
✕

HOOK DETAILS ?

Operation
 Pre-snapshot

Hook arguments (optional)
 1 pre ✕ ?
 Enter hook arguments

Hook name
 redis-pre-snapshot

EXECUTION HOOKS

Execution hooks allow Astra Control to execute your own custom scripts before or after a snapshot.

Read more in [Manage application execution hooks](#)

CONTAINER IMAGES ?

Apply to all container images

Use a regular expression to target container images for the hook.

Container image names to match:
 redis

SCRIPT ?

+ Add
Search

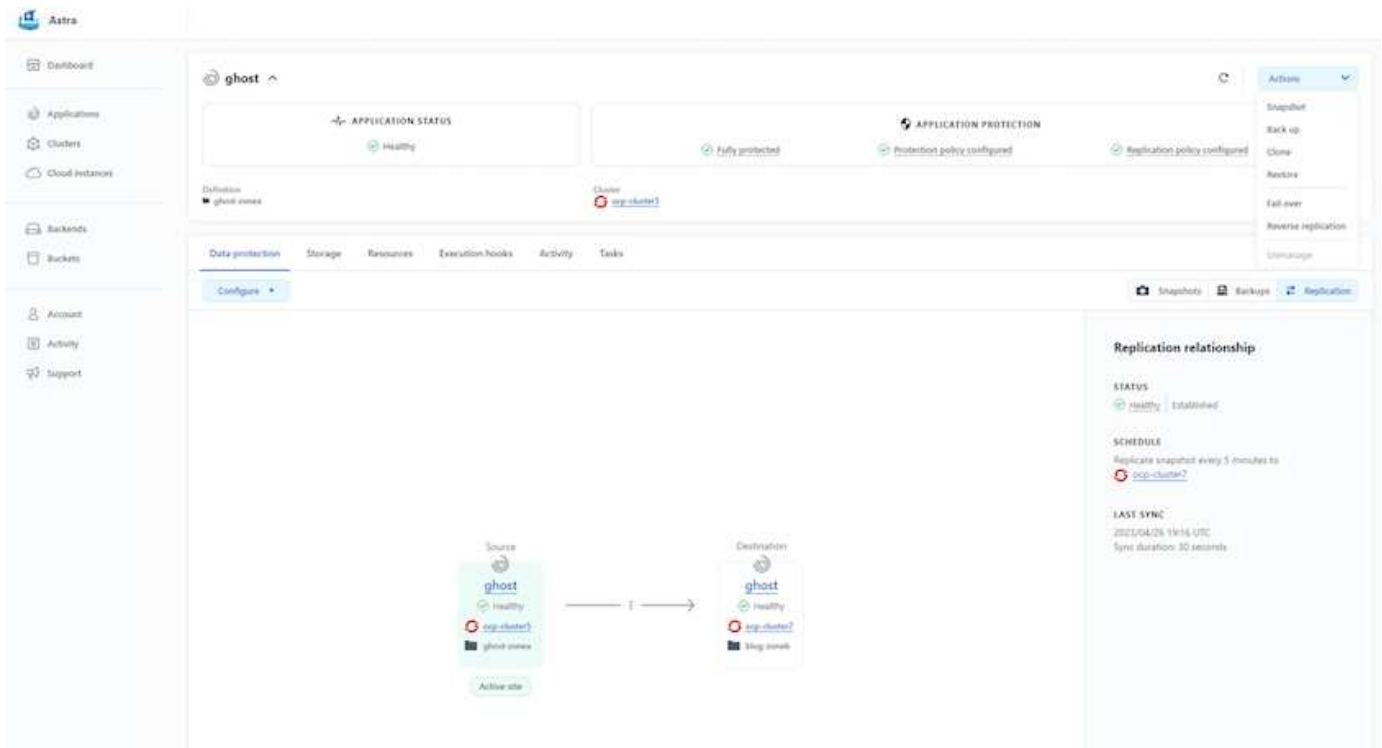
Name ↓
<input type="radio"/> mariadb_mysql.sh
<input type="radio"/> postgresql.sh
<input checked="" type="radio"/> redis_hook.sh

Cancel
Save ✓

Replication with ACC

For regional protection or for a low RPO and RTO solution, an application can be replicated to another Kubernetes instance running at a different site, preferably in another region. ACC utilizes ONTAP async SnapMirror with RPO as low as 5 minutes. Refer [here](#) for SnapMirror setup instructions.

SnapMirror with ACC



san-economy and nas-economy storage drivers do not support replication feature. Refer [here](#) for additional details.

Demo video:

[Demonstration video of disaster recovery with Astra Control Center](#)

[Data protection with Astra Control Center](#)

Details on Astra Control Center Data Protection features are available [here](#)

Disaster recovery (Failover and Failback using replication) with ACC

[Using Astra Control for Failover and Failback of applications](#)

Data migration using Astra Control Center

This page shows the data migration options for container workloads on Red Hat OpenShift clusters with Astra Control Center (ACC). Specifically, customers can use ACC to

- move some selected workloads or all workloads from their on-premises data centers to the cloud
- clone their apps to the cloud either for testing purposes or move from the data center to the cloud

Data Migration

To migrate application from one environment to another, you can use one of the following features of ACC:

- replication
- backup and restore
- clone

Refer to the [data protection section](#) for the **replication and backup and restore** options.

Refer [here](#) for additional details about **cloning**.



Astra Replication feature is only supported with Trident Container Storage Interface (CSI). However, replication is not supported by nas-economy & san-economy drivers.

Performing data replication using ACC

The screenshot displays the Astra management console interface for configuring a replication relationship. The main content area is titled 'ghost' and shows the 'APPLICATION STATUS' as 'Healthy'. Under 'APPLICATION PROTECTION', it indicates 'Fully protected', 'Protection policy configured', and 'Replication policy configured'. The 'Definition' is 'ghost-vmware' and the 'Cluster' is 'oci-cluster1'.

The 'Data protection' tab is active, showing a 'Configure' button. Below this, a diagram illustrates the replication relationship between a 'Source' and a 'Destination', both labeled 'ghost'. The source is associated with 'oci-cluster1' and 'ghost-vmware', while the destination is associated with 'oci-cluster2' and 'vmware'. An arrow points from the source to the destination.

On the right side, the 'Replication relationship' details are shown:

- STATUS:** healthy | Established
- SCHEDULE:** Replicate snapshot every 5 minutes to [oci-cluster2](#)
- LAST SYNC:** 2023-04-25 19:16 UTC, Sync duration: 30 seconds

The left sidebar contains navigation options: Dashboard, Applications, Clusters, Cloud instances, Backends, Buckets, Account, Activity, and Support.

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.