



# **Hybrid Cloud**

## FlexPod

NetApp  
November 04, 2025

This PDF was generated from <https://docs.netapp.com/us-en/flexpod/hybrid-cloud/fhc-cvoe-solution-overview.html> on November 04, 2025. Always check [docs.netapp.com](https://docs.netapp.com) for the latest.

# Table of Contents

Hybrid Cloud . . . . .	1
FlexPod Hybrid Cloud with Cloud Volumes ONTAP for Epic . . . . .	1
TR-4960: FlexPod Hybrid Cloud with Cloud Volumes ONTAP for Epic . . . . .	1
Solution components . . . . .	3
Installation and configuration . . . . .	7
SAN configuration . . . . .	11
Solution validation . . . . .	17
Conclusion . . . . .	25
Where to find additional information . . . . .	25
FlexPod Hybrid Cloud for Google Cloud Platform with NetApp Cloud Volumes ONTAP and Cisco Intersight . . . . .	26
TR-4939: FlexPod Hybrid Cloud for Google Cloud Platform with NetApp Cloud Volumes ONTAP and Cisco Intersight . . . . .	27
Solution components . . . . .	29
Installation and configuration . . . . .	34
Solution validation . . . . .	97
Conclusion . . . . .	105
FlexPod hybrid cloud with NetApp Astra and Cisco Intersight for Red Hat OpenShift . . . . .	107
TR-4936: FlexPod hybrid cloud with NetApp Astra and Cisco Intersight for Red Hat OpenShift . . . . .	108
Solution components . . . . .	111
Installation and configuration . . . . .	118
Solution validation . . . . .	140
Conclusion . . . . .	161
NetApp Cloud Insights for FlexPod . . . . .	163
TR-4868: NetApp Cloud Insights for FlexPod . . . . .	163
Use cases . . . . .	163
Architecture . . . . .	164
Design considerations . . . . .	166
Deploy Cloud Insights for FlexPod . . . . .	166
Use cases . . . . .	178
Videos and demos . . . . .	186
Additional information . . . . .	186
FlexPod with FabricPool - Inactive Data Tiering to Amazon AWS S3 . . . . .	187
TR-4801: FlexPod with FabricPool - Inactive Data Tiering to Amazon AWS S3 . . . . .	187
FlexPod overview and architecture . . . . .	187
FabricPool . . . . .	189
FabricPool requirements . . . . .	193
Configuration . . . . .	197
Performance considerations . . . . .	207
Cost of ownership . . . . .	208
Conclusion . . . . .	208
Where to find additional information . . . . .	208
FlexPod Datacenter for Hybrid Cloud with Cisco CloudCenter and NetApp private storage - Design . . . . .	208

# Hybrid Cloud

## FlexPod Hybrid Cloud with Cloud Volumes ONTAP for Epic

### TR-4960: FlexPod Hybrid Cloud with Cloud Volumes ONTAP for Epic

In partnership with:



Kamini Singh, NetApp

The key to making a digital transformation is simply doing more with data. Hospitals generate and require large amounts of data to run their organization and serve their patients effectively. Information is collected and processed when treating patients and managing staff schedules and medical resources.

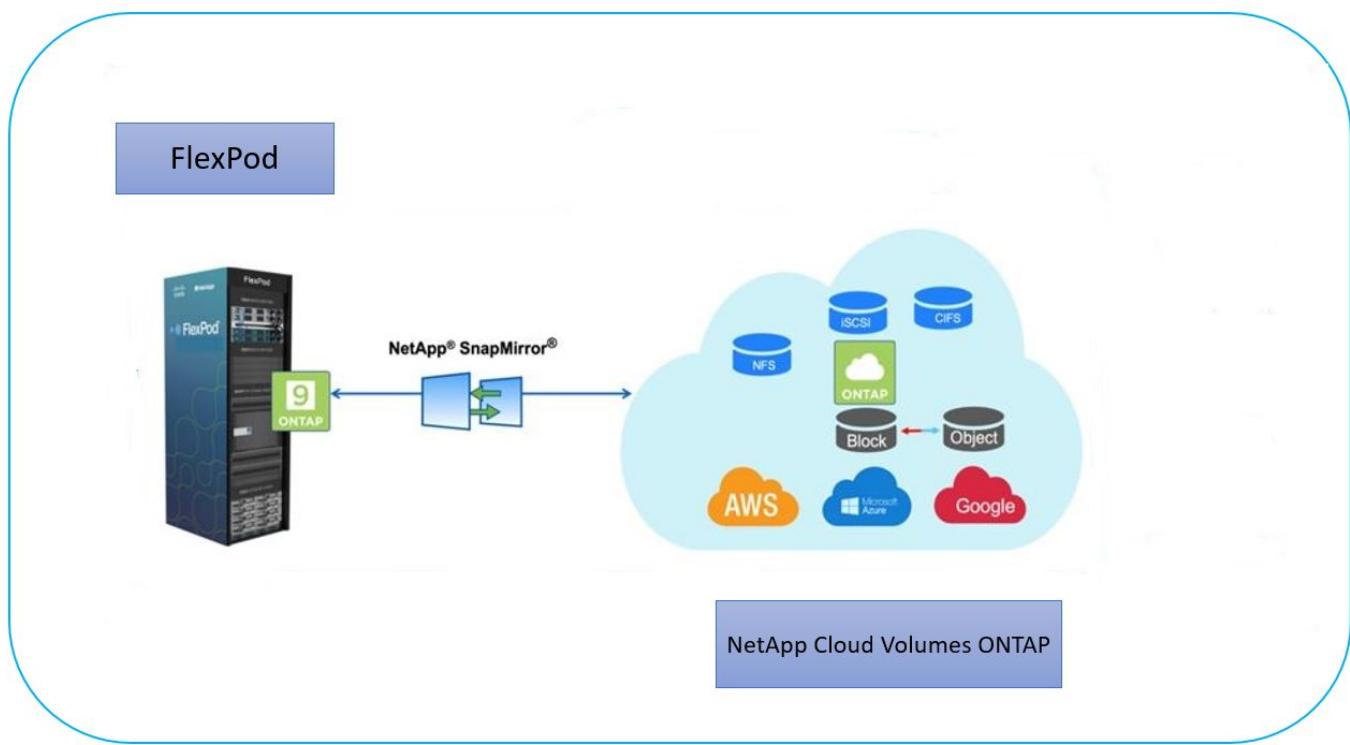
The ever-increasing size of healthcare data and the valuable insights that this data can provide make healthcare data services and data protection both critical and challenging. First, healthcare data must be both available and protected to meet data recovery, medical business continuity, or compliance requirements.

Second, healthcare data must be made readily available for analysis. Often this analysis uses artificial intelligence (AI)- and machine learning (ML)-based approaches to help medical businesses improve their solutions and create business values.

Third, the data service infrastructures and the data protection methodologies must accommodate the growth of healthcare data as a medical business grows. In addition, data mobility is increasingly becoming critical due to the need to move data from the edge where it is created to the core and cloud to use resources available there for data analysis or archival purposes.

NetApp offers a single data management solution for enterprise applications, including healthcare, and we are able to guide hospitals through their journey toward digital transformation. NetApp Cloud Volumes ONTAP delivers a solution for healthcare data management in which data can be efficiently replicated from a FlexPod Datacenter to Cloud Volumes ONTAP deployed on a public cloud like AWS.

By leveraging cost-effective and secure public cloud resources, Cloud Volumes ONTAP enhances cloud-based disaster recovery (DR) with highly efficient data replication, built-in storage efficiencies, and simple DR testing. These systems are managed with unified control and drag-and-drop simplicity, which provides cost-effective and bullet-proof protection against any kind of error, failure, or disaster. Cloud Volumes ONTAP provides NetApp SnapMirror technology as a solution for block-level data replication that keeps the destination up to date through incremental updates.



## Audience

This document is intended for NetApp and partner solutions engineers (SEs) and professional services personnel. NetApp assumes that the reader has the following background knowledge:

- A solid understanding of SAN and NAS concepts
- Technical familiarity with NetApp ONTAP storage systems
- Technical familiarity with the configuration and administration of ONTAP software

## Solution benefits

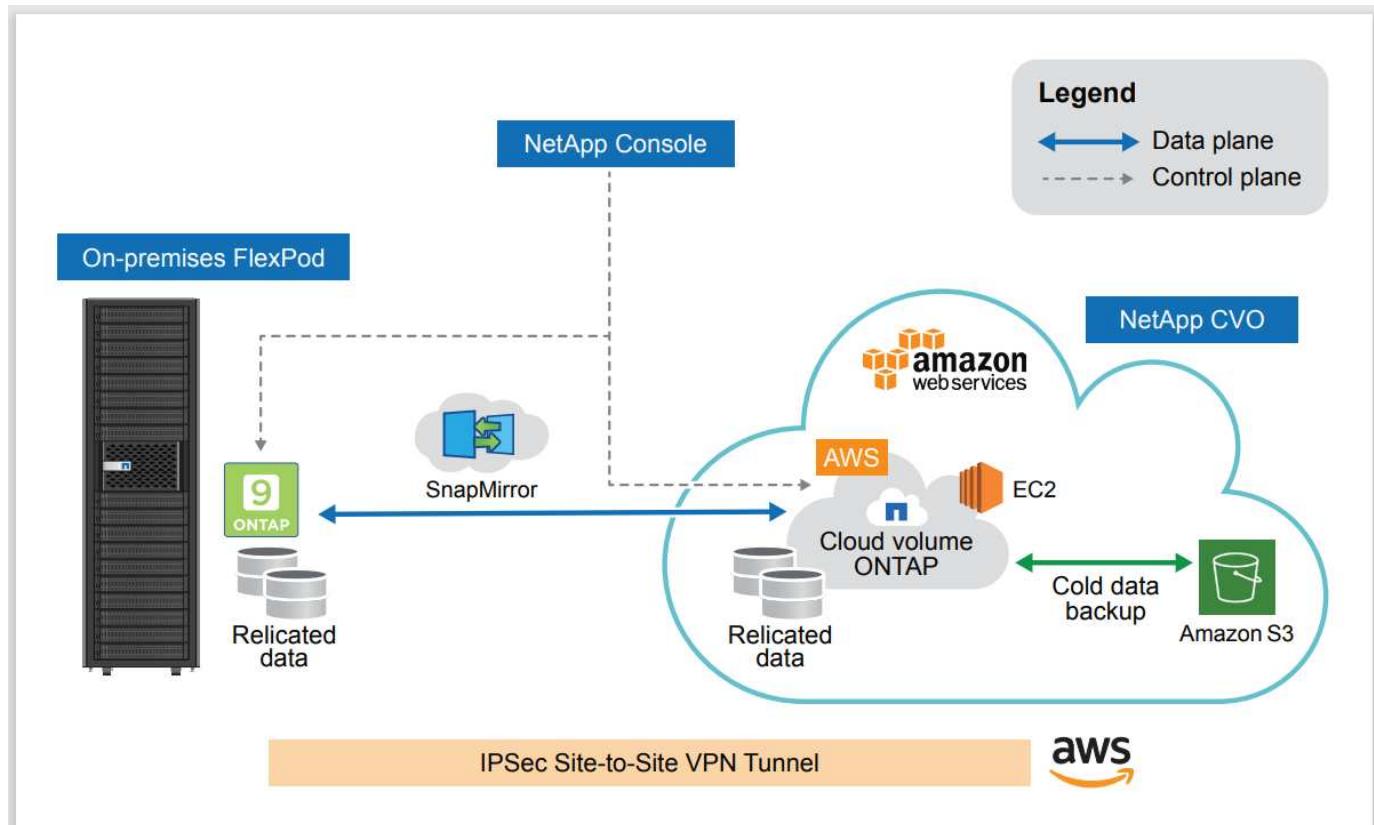
FlexPod Datacenter integrated with NetApp Cloud Volumes ONTAP offers the following benefits to healthcare workloads:

- **Customized protection.** Cloud Volumes ONTAP provides block-level data replication from ONTAP to the cloud that keeps the destination up to date through incremental updates. Users can specify a synchronization schedule to determine when changes at the source are transferred over. This provides customized protection for all sorts of healthcare data.
- **Failover and Fallback.** When a disaster occurs, storage administrators can quickly set failover to the cloud volumes. When the primary site is recovered, the new data created in the DR environment is synchronized back to the source volumes enabling the secondary data replication to be re-established. In this way, healthcare data can be easily recovered without disruption.
- **Efficiency.** The storage space and costs for the secondary cloud copy are optimized using data compression, thin provisioning, and deduplication. Healthcare data is transferred at the block-level in a compressed and deduplicated form, improving the speed of the transfers. Data is also automatically tiered to low-cost object storage and only brought back to high-performance storage when accessed, such as in a DR scenario. This significantly reduces ongoing storage costs.
- **Ransomware Protection.** NetApp Console ransomware protection scans data sources across on-premises and cloud environments, detects security vulnerabilities, and provides their current security status

and risk scoring. It then provides actionable recommendations that you can further investigate and follow to remediate. This enables you to protect your critical healthcare data from ransomware attacks.

## Solution topology

This section describes the logical topology of the solution. The following figure represents the solution topology composed of the FlexPod on-premises environment, NetApp Cloud Volumes ONTAP (CVO) running on Amazon Web Services (AWS), and the NetApp Console SaaS platform.



The control planes and data planes are clearly indicated between the endpoints. The data plane runs between the ONTAP instance running on all-flash FAS in FlexPod and the NetApp CVO instance in AWS by leveraging a secure site-to-site VPN connection. The replication of healthcare workload data from the on-premises FlexPod Datacenter to NetApp Cloud Volumes ONTAP is handled by NetApp SnapMirror replication. An optional backup and tiering of the cold data residing in the NetApp CVO instance to AWS S3 is also supported with this solution.

[Next: Solution components.](#)

## Solution components

[Previous: Solution Overview.](#)

### FlexPod

FlexPod is a defined set of hardware and software that forms an integrated foundation for both virtualized and non-virtualized solutions. FlexPod includes NetApp ONTAP storage, Cisco Nexus networking, Cisco MDS storage networking, and the Cisco Unified Computing System (Cisco UCS).

Healthcare organizations are looking for a solution to ease their digital transformation and improve patient

experiences and outcomes. With FlexPod, you get a secure, scalable platform that drives efficiency and empowers your staff to make more informed decisions faster so that they can provide better patient care.

FlexPod is the ideal platform for healthcare workload needs because it provides the following benefits:

- Optimization of operations to get faster insights and better patient outcomes.
- Streamlining imaging apps with scalable, reliable infrastructure.
- Deploying quickly and efficiently with a proven approach for healthcare-specific apps such as EHR.

## **EHR**

Electronic Health Records (EHRs) makes software for midsize and large medical groups, hospitals, and integrated healthcare organizations. Customers also include community hospitals, academic facilities, children's organizations, safety net providers, and multi-hospital systems. EHR-integrated software spans clinical, access, and revenue functions and extends into the home.

Healthcare provider organizations remain under pressure to maximize the benefits of their substantial investments in industry-leading EHRs. When customers design their data centers for EHR solutions and mission-critical applications, they often identify the following goals for their data center architecture:

- High availability of the EHR applications
- High performance
- Ease of implementing EHR in the data center
- Agility and scalability to enable growth with new EHR releases or applications
- Cost effectiveness
- Manageability, stability, and ease of support
- Robust data protection, backup, recovery, and business continuance

FlexPod is EHR certified and supports a platform containing Cisco Cisco UCS with Intel Xeon processors, Red Hat Enterprise Linux (RHEL), and virtualization with VMware ESXi. This platform, coupled with EHR's High Comfort Level ranking for NetApp storage running ONTAP, enables you to run your healthcare applications in a fully managed private cloud through FlexPod that can also be connected to any of the public cloud providers.

## **NetApp Console**

NetApp Console 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 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 [NetApp Console documentation](#).

## **Console agent**

A Console agent instance enables the Console to manage resources and processes within a public cloud environment. A Console agent is required for many of the features provided by the Console, and it can be deployed in the cloud or in the on-premises network.

A Console agent is supported in the following locations:

- Amazon Web Services
- Microsoft Azure

- Google Cloud
- On premises

[Learn more about Console agents.](#)

## NetApp Cloud Volumes ONTAP

NetApp Cloud Volumes ONTAP is a software-defined storage offering that runs ONTAP data management software in the cloud to deliver 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 the following:

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

For more detailed information, see [Cloud Volumes ONTAP](#).

## NetApp Active IQ Unified Manager

NetApp Active IQ Unified Manager allows the monitoring of your ONTAP storage clusters from a single, redesigned, and intuitive interface that delivers intelligence from community wisdom and AI analytics. It provides comprehensive operational, performant, and proactive insights into the storage environment and the virtual machines running on it. When an issue occurs with the storage infrastructure, Unified Manager can notify you about the details of the issue to help identify the root cause. The virtual machine dashboard gives you a view into the performance statistics for the VM so that you can investigate the entire I/O path from the 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 allows you to plan for the storage requirements of your users by forecasting capacity and usage trends so that you can act before issues arise, preventing reactive short-term decisions that can lead to additional problems in the long term.

For more information, see [Active IQ Unified Manager](#).

## 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 to drive change with IT teams and delivers an operating model designed for hybrid cloud. Cisco Intersight provides the following benefits:

- **Faster delivery.** Intersight is 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. In this way, the customer can focus on supporting critical business needs.
- **Simplified operations.** Intersight simplifies operations by using a single, secure SaaS-delivered tool with common inventory, authentication, and APIs to work across the full stack and all locations, eliminating silos across teams. This allows you to manage physical servers and hypervisors on-premises, to VMs, K8s, serverless, automation, optimization, and cost control both on-premises and in public clouds.
- **Continuous optimization.** You can continuously optimize your environment by using intelligence provided by Cisco Intersight across every layer, as well as by Cisco TAC. This intelligence is converted into recommended and automatable actions so that you can adapt in real-time to any changes: from moving workloads and monitoring the health of physical servers to cost reduction recommendations for the public clouds that 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 UCSM Managed Mode (UMM) or Intersight Managed Mode (IMM) for fabric-attached Cisco UCS systems during the initial setup of the fabric Interconnects. In this solution, native IMM is used. The following figure shows the Cisco Intersight Dashboard.

The screenshot shows the Cisco Intersight dashboard with the 'Servers' page selected. The left sidebar has 'Operate' selected, with 'Servers' highlighted. The main area shows a summary dashboard with metrics like Health (10), Power (Off 4, On 6), HCL Status (Incomplete 2, Validated 8), Models (UCSX 210C-M6 10), Contract Status (Active 4, Not Covered 6), and Profile Status (21, Not Assigned 20). Below this is a table of server details:

Name	Health	Model	CPU Capacity	Memory Capacity	UCS D...	Server Profile
fpsa-6454-g03-hc-1-1	Healthy	UCSX-210C-M6	112.0	512.0	fpsa-6454...	HC-VM-Host-ISC...
fpsa-6454-g03-hc-1-2	Healthy	UCSX-210C-M6	182.4	1024.0	fpsa-6454...	HC-VM-Host-ISC...
fpsa-6454-g03-hc-1-5	Healthy	UCSX-210C-M6	112.0	128.0	fpsa-6454...	HC-VM-Host-ISC...
fpsa-6454-g03-hc-1-6	Healthy	UCSX-210C-M6	112.0	128.0	fpsa-6454...	HC-VM-Host-ISC...
fpsa-6454-g03-pit-1-1	Healthy	UCSX-210C-M6	112.0	256.0	fpsa-6454...	PLT-G03-ESX...
fpsa-6454-g03-pit-1-3	Healthy	UCSX-210C-M6	112.0	256.0	fpsa-6454...	PLT-G03-ESX...
FPSA-Enablement-G2-1-1	Critical	UCSX-210C-M6	145.6	512.0	FPSA-Enab...	ISCSI-Boot-Templ...
FPSA-Enablement-G2-1-2	Healthy	UCSX-210C-M6	108.0	512.0	FPSA-Enab...	ISCSI-Boot-Templ...
FPSA-Enablement-G2-1-5	Critical	UCSX-210C-M6	145.6	512.0	FPSA-Enab...	ISCSI-Boot-Templ...
FPSA-Enablement-G2-1-6	Healthy	UCSX-210C-M6	145.6	512.0	FPSA-Enab...	ISCSI-Boot-Templ...

## VMware vSphere 7.0

VMware vSphere is a virtualization platform for holistically managing large collections of infrastructure (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 datacenter to create a single powerhouse with resources that can be allocated quickly and dynamically to any application in need.

For more information about VMware vSphere and its components, see [VMware vSphere](#).

## VMware vCenter Server

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.

For detailed information, see [VMware vCenter](#).

## Hardware and software revisions

This hybrid cloud 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](#), [UCS Hardware and Software Compatibility](#), and [VMware Compatibility Guide](#).

The following table shows the on-premises FlexPod hardware and software revisions.

Component	Product	Version
Compute	Cisco UCS X210c M6	5.0(1b)
	Cisco UCS Fabric Interconnects 6454	4.2(2a)
Network	Cisco Nexus 9336C-FX2 NX-OS	9.3(9)
Storage	NetApp AFF A400	ONTAP 9.11.1P2
	NetApp ONTAP Tools for VMware vSphere	9.11
	NetApp NFS Plug-in for VMware VAAI	2.0
	NetApp Active IQ Unified Manager	9.11P1
Software	VMware vSphere	7.0(U3)
	VMware ESXi nenic Ethernet Driver	1.0.35.0
	VMware vCenter Appliance	7.0.3
	Cisco Intersight Assist Virtual Appliance	1.0.9-342

The following table shows the Console and Cloud Volumes ONTAP versions.

Vendor	Product	Version
NetApp	Console	3.9.24
	Cloud Volumes ONTAP	ONTAP 9.11

[Next: Installation and configuration](#).

## Installation and configuration

[Previous: Solution components](#).

## NetApp Cloud Volumes ONTAP deployment

Complete the following steps to configure your Cloud Volumes ONTAP instance:

1. Prepare the public cloud service provider environment.

You must capture the environment details of your public cloud service provider for the solution configuration. For example, for Amazon Web Services (AWS) environment preparation, you need the AWS access key, the AWS secret key, and other network details like region, VPC, subnet, and so on.

2. Configure the VPC endpoint gateway.

A VPC endpoint gateway is required to enable the connection between the VPC and the AWS S3 service. This is used to enable the backup on CVO, an endpoint with the Gateway type.

3. Access the NetApp Console.

To access the Console and other cloud services, you need to sign up on [NetApp Console](#). For setting up workspaces and users in the Console account, see [NetApp Console setup and administration](#). You need an account that has permission to deploy the Console agent in your cloud provider directly from the Console. To obtain the permissions you need, refer to [Permissions summary for NetApp Console](#).

4. Deploy Console agent.

Before adding a Cloud Volume ONTAP system, you must deploy a Console agent. The Console prompts you if you try to create your first Cloud Volumes ONTAP system without a Console agent in place. To deploy a Console agent in AWS from the Console, see the [Console agent installation options in AWS](#).

5. Launch Cloud Volumes ONTAP in AWS.

You can launch Cloud Volumes ONTAP in a single-system configuration or as an HA pair in AWS. [Read the step-by-step instructions](#).

For detailed information about these steps, see the [Quick start guide for Cloud Volumes ONTAP in AWS](#).

In this solution, we have deployed a single-node Cloud Volumes ONTAP system in AWS.

## On-premises FlexPod Deployment

To understand FlexPod with UCS X-Series, VMware, and NetApp ONTAP design details, see the [FlexPod Datacenter with Cisco UCS X-Series](#) design guide. This document provides design guidance for incorporating the Cisco Intersight-managed UCS X-Series platform within the FlexPod Datacenter infrastructure.

For deploying the on-premises FlexPod instance, see [this deployment guide](#).

This document provides deployment guidance for incorporating the Cisco Intersight-managed UCS X-Series platform within a FlexPod Datacenter infrastructure. The document covers both configurations and best practices for a successful deployment.

FlexPod can be deployed in both UCS Managed Mode and Cisco Intersight Managed Mode (IMM). If you are deploying FlexPod in UCS Managed Mode, see this [design guide](#) and this [deployment guide](#).

FlexPod deployment can be automated with Infrastructure as code using Ansible. Below are the links to the GitHub repositories for End-to-End FlexPod deployment:

- Ansible configuration of FlexPod with Cisco UCS in UCS Managed Mode, NetApp ONTAP, and VMware vSphere can be seen [here](#).
- Ansible configuration of FlexPod with Cisco UCS in IMM, NetApp ONTAP, and VMware vSphere can be seen [here](#).

## On-premises ONTAP storage configuration

This section describes some of the important ONTAP configuration steps that are specific to this solution.

- Configure an SVM with the iSCSI service running.

```

1. vserver create -vserver Healthcare_SVM -rootvolume
Healthcare_SVM_root -aggregate aggr1_A400_G0312_01 -rootvolume-security
-style unix
2. vserver add-protocols -vserver Healthcare_SVM -protocols iscsi
3. vserver iscsi create -vserver Healthcare_SVM

```

To verify:

```

A400-G0312::> vserver iscsi show -vserver Healthcare_SVM
Vserver: Healthcare_SVM
Target Name:
iqn.1992-08.com.netapp:sn.1fbf00f438c111ed866cd039ea91fb56:vs.3
Target Alias: Healthcare_SVM
Administrative Status: up

```

If the iSCSI license was not installed during cluster configuration, make sure to install the license before creating the iSCSI service.

- Create a FlexVol volume.

```

1. volume create -vserver Healthcare_SVM -volume hc_iscsi_vol -aggregate
aggr1_A400_G0312_01 -size 500GB -state online -policy default -space
guarantee none

```

- Add interfaces for iSCSI access.

```
1. network interface create -vserver Healthcare_SVM -lif iscsi-lif-01a  
-service-policy default-data-iscsi -home-node <st-node01> -home-port  
a0a-<infra-iscsi-a-vlan-id> -address <st-node01-infra-iscsi-a-ip>  
-netmask <infra-iscsi-a-mask> -status-admin up  
2. network interface create -vserver Healthcare_SVM -lif iscsi-lif-01b  
-service-policy default-data-iscsi -home-node <st-node01> -home-port  
a0a-<infra-iscsi-b-vlan-id> -address <st-node01-infra-iscsi-b-ip>  
-netmask <infra-iscsi-b-mask> -status-admin up  
3. network interface create -vserver Healthcare_SVM -lif iscsi-lif-02a  
-service-policy default-data-iscsi -home-node <st-node02> -home-port  
a0a-<infra-iscsi-a-vlan-id> -address <st-node02-infra-iscsi-a-ip>  
-netmask <infra-iscsi-a-mask> -status-admin up  
4. network interface create -vserver Healthcare_SVM -lif iscsi-lif-02b  
-service-policy default-data-iscsi -home-node <st-node02> -home-port  
a0a-<infra-iscsi-b-vlan-id> -address <st-node02-infra-iscsi-b-ip>  
-netmask <infra-iscsi-b-mask> -status-admin up
```

In this solution, we created four iSCSI logical interfaces (LIFs), two on each node.

After the FlexPod instance is up and running with vCenter deployed and all ESXi hosts added to it, we need to deploy a Linux VM that acts as a server that connects to and accesses the NetApp ONTAP storage. In this solution, we have installed a CentOS 8 instance in vCenter.

#### 4. Create a LUN.

```
1. lun create -vserver Healthcare_SVM -path /vol/hc_iscsi_vol/iscsi_lun1  
-size 200GB -ostype linux -space-reserve disabled
```

For an EHR operational database (ODB), a journal, and application workloads, EHR recommends presenting storage to servers as iSCSI LUNs. NetApp also supports using FCP and NVMe/FC if you have versions of AIX and the RHEL operating systems that are capable, which enhances performance. FCP and NVMe/FC can coexist on the same fabric.

#### 5. Create an igroup.

```
1. igrup create -vserver Healthcare_SVM -igroup ehr -protocol iscsi  
-ostype linux -initiator iqn.1994-05.com.redhat:8e91e9769336
```

Igroups are used to allow server access to LUNs. For Linux host, the server IQN can be found in the file /etc/iscsi/initiatorname.iscsi.

#### 6. Map the LUN to the igroup.

```
1. lun mapping create -vserver Healthcare_SVM -path  
/vol/hc_iscsi_vol/iscsi_lun1 -igroup ehr -lun-id 0
```

## Add on-premises FlexPod storage to the NetApp Console

Complete the following steps to add your FlexPod storage to the system using the Console.

1. From the navigation menu, select **Storage > Systems**.
2. On the Systems page, click **Add System** and select **On-Premises**.
3. Select **On-Premises ONTAP**. Click **Next**.
4. On the ONTAP Cluster Details page, enter the cluster management IP address and the password for the admin user account. Then click **Add**.
5. On the Details and Credentials page, enter a name and description for the working environment, and then click **Go**.

The Console discovers the ONTAP cluster and adds it as a system on the Systems page.

For detailed information, see the page [Discover on-premises ONTAP clusters](#).

Next: [SAN configuration](#).

## SAN configuration

[Previous: Installation and configuration](#).

This section describes the host-side configuration required by EHR to enable the software to best integrate with NetApp storage. In this segment, we specifically discuss the host integration for Linux operating systems. Use the [NetApp Interoperability Matrix Tool \(IMT\)](#) to validate all versions of software and firmware.



The following configuration steps are specific to the CentOS 8 host that was used in this solution.

### NetApp Host Utility Kit

NetApp recommends installing the NetApp Host Utility Kit (Host Utilities) on the operating systems of hosts that are connected to and accessing NetApp storage systems. Native Microsoft Multipath I/O (MPIO) is supported. The OS must be asymmetric logical unit access (ALUA)-capable for multipathing. Installing the Host Utilities configures the host bus adapter (HBA) settings for NetApp storage.

NetApp Host Utilities can be downloaded [here](#). In this solution, we have installed Linux Host Utilities 7.1 on the host.

```
[root@hc-cloud-secure-1 ~]# rpm -ivh netapp_linux_unified_host_utilities-  
7-1.x86_64.rpm
```

## Discover ONTAP storage

Make sure the iSCSI service is running when the log-ins are supposed to occur. To set the log-in mode for a specific portal on a target or for all the portals on a target, use the `iscsiadm` command.

```
[root@hc-cloud-secure-1 ~]# rescan-scsi-bus.sh
[root@hc-cloud-secure-1 ~]# iscsiadm -m discovery -t sendtargets -p
<iscsi-lif-ip>
[root@hc-cloud-secure-1 ~]# iscsiadm -m node -L all
```

Now you can use `sanlun` to display information about the LUNs connected to the host. Make sure that you are logged in as root on the host.

```
[root@hc-cloud-secure-1 ~]# sanlun lun show
controller(7mode/E-Series) /
                               device      host      lun
vserver(cDOT/FlashRay) lun-pathname filename  adapter protocol size
product
-----
---
Healthcare_SVM           /dev/sdb  host33  iSCSI    200g
cDOT
                  /vol/hc_iscsi_vol/iscsi_lun1

Healthcare_SVM           /dev/sdc  host34  iSCSI    200g
cDOT
                  /vol/hc_iscsi_vol/iscsi_lun1
```

## Configure multipathing

Device Mapper Multipathing (DM-Multipath) is a native multipathing utility in Linux. It can be used for redundancy and to improve performance. It aggregates or combines the multiple I/O paths between servers and storage, so it creates a single device at the OS Level.

1. Before setting up DM-Multipath on your system, make sure that that your system has been updated and includes the `device-mapper-multipath` package.

```
[root@hc-cloud-secure-1 ~]# rpm -qa | grep multipath
device-mapper-multipath-libs-0.8.4-31.el8.x86_64
device-mapper-multipath-0.8.4-31.el8.x86_64
```

2. The configuration file is the `/etc/multipath.conf` file. Update the configuration file as shown below.

```
[root@hc-cloud-secure-1 ~]# cat /etc/multipath.conf
defaults {
    path_checker      readsector0
    no_path_retry     fail
}
devices {
    device {
        vendor        "NETAPP  "
        product       "LUN.*"
        no_path_retry  queue
        path_checker   tur
    }
}
```

### 3. Enable and start the multipath services.

```
[root@hc-cloud-secure-1 ~]# systemctl enable multipathd.service
[root@hc-cloud-secure-1 ~]# systemctl start multipathd.service
```

### 4. Add the loadable kernel module `dm-multipath` and restart the multipath service. Finally, check the multipathing status.

```
[root@hc-cloud-secure-1 ~]# modprobe -v dm-multipath
insmod /lib/modules/4.18.0-408.el8.x86_64/kernel/drivers/md/dm-
multipath.ko.xz

[root@hc-cloud-secure-1 ~]# systemctl restart multipathd.service

[root@hc-cloud-secure-1 ~]# multipath -ll
3600a09803831494c372b545a4d786278 dm-2 NETAPP,LUN C-Mode
size=200G features='3 queue_if_no_path pg_init_retries 50' hwhandler='1
alua' wp=rw
|-- policy='service-time 0' prio=50 status=active
|  `-- 33:0:0:0 sdb 8:16 active ready running
`-- policy='service-time 0' prio=10 status=enabled
  `-- 34:0:0:0 sdc 8:32 active ready running
```



For detailed information about these steps, see [here](#).

## Create physical volume

Use the `pvcreate` command to initialize a block device to be used as a physical volume. Initialization is analogous to formatting a file system.

```
[root@hc-cloud-secure-1 ~]# pvcreate /dev/sdb
Physical volume "/dev/sdb" successfully created.
```

## Create volume group

To create a volume group from one or more physical volumes, use the `vgcreate` command. This command creates a new volume group by name and adds at least one physical volume to it.

```
[root@hc-cloud-secure-1 ~]# vgcreate datavg /dev/sdb
Volume group "datavg" successfully created.
```

The `vgdisplay` command can be used to display volume group properties (such as size, extents, number of physical volumes, and so on) in a fixed form.

```
[root@hc-cloud-secure-1 ~]# vgdisplay datavg
--- Volume group ---
VG Name          datavg
System ID
Format          lvm2
Metadata Areas  1
Metadata Sequence No  1
VG Access        read/write
VG Status        resizable
MAX LV          0
Cur LV          0
Open LV          0
Max PV          0
Cur PV          1
Act PV          1
VG Size          <200.00 GiB
PE Size          4.00 MiB
Total PE        51199
Alloc PE / Size 0 / 0
Free  PE / Size 51199 / <200.00 GiB
VG UUID          C7jmI0-J0SS-Cq91-t6b4-A9xw-nTfi-RXcy28
```

## Create logical volume

When you create a logical volume, the logical volume is carved from a volume group using the free extents on the physical volumes that make up the volume group.

```
[root@hc-cloud-secure-1 ~]# lvcreate -l 100%FREE -n datalv datavg
Logical volume "datalv" created.
```

This command creates a logical volume called `data1v` that uses all of the unallocated space in the volume group `datavg`.

## Create file system

```
[root@hc-cloud-secure-1 ~]# mkfs.xfs -K /dev/datavg/datalv
meta-data=/dev/datavg/datalv      isize=512      agcount=4, agsize=13106944
blks
          =          sectsz=4096  attr=2, projid32bit=1
          =          crc=1      finobt=1, sparse=1, rmapbt=0
          =          reflink=1  bigtime=0 inobtcount=0
data      =          bsize=4096   blocks=52427776, imaxpct=25
          =          sunit=0    swidth=0 blks
naming    =version 2      bsize=4096  ascii-ci=0, ftype=1
log       =internal log   bsize=4096   blocks=25599, version=2
          =          sectsz=4096  sunit=1 blks, lazy-count=1
realtime  =none          extsz=4096   blocks=0, rtextents=0
```

## Make folder to mount

```
[root@hc-cloud-secure-1 ~]# mkdir /file1
```

## Mount the file system

```
[root@hc-cloud-secure-1 ~]# mount -t xfs /dev/datavg/datalv /file1

[root@hc-cloud-secure-1 ~]# df -k
Filesystem      1K-blocks  Used  Available  Use%  Mounted on
devtmpfs          8072804    0  8072804  0%  /dev
tmpfs             8103272    0  8103272  0%  /dev/shm
tmpfs             8103272  9404  8093868  1%  /run
tmpfs             8103272    0  8103272  0%  /sys/fs/cgroup
/dev/mapper/cs-root  45496624 5642104 39854520 13%  /
/dev/sda2          1038336 258712  779624 25%  /boot
/dev/sda1           613184  7416  605768  2%  /boot/efi
tmpfs              1620652   12  1620640  1%  /run/user/42
tmpfs              1620652    0  1620652  0%  /run/user/0
/dev/mapper/datavg-datalv 209608708 1494520 208114188 1%  /file1
```

For detailed information about these tasks, see the page [LVM Administration with CLI Commands](#).

## Data generation

Dgen.pl is a perl script data generator for EHR's I/O simulator (GenerateIO). Data inside the LUNs are

generated with the EHR Dgen.pl script. The script is designed to create data similar to the data in an EHR database.

```
[root@hc-cloud-secure-1 ~]# cd GenerateIO-1.17.3/  
  
[root@hc-cloud-secure-1 GenerateIO-1.17.3]# ./dgen.pl --directory /file1  
--jobs 80  
  
[root@hc-cloud-secure-1 ~]# cd /file1/  
[root@hc-cloud-secure-1 file1]# ls  
dir01  dir05  dir09  dir13  dir17  dir21  dir25  dir29  dir33  dir37  
dir41  dir45  dir49  dir53  dir57  dir61  dir65  dir69  dir73  dir77  
dir02  dir06  dir10  dir14  dir18  dir22  dir26  dir30  dir34  dir38  
dir42  dir46  dir50  dir54  dir58  dir62  dir66  dir70  dir74  dir78  
dir03  dir07  dir11  dir15  dir19  dir23  dir27  dir31  dir35  dir39  
dir43  dir47  dir51  dir55  dir59  dir63  dir67  dir71  dir75  dir79  
dir04  dir08  dir12  dir16  dir20  dir24  dir28  dir32  dir36  dir40  
dir44  dir48  dir52  dir56  dir60  dir64  dir68  dir72  dir76  dir80  
  
[root@hc-cloud-secure-1 file1]# df -k .  
Filesystem          1K-blocks  Used  Available  Use%  Mounted  
on  
/dev/mapper/datavg-datalv  209608708  178167156  31441552  85%  /file1
```

While running, the Dgen.pl script uses 85% of the file system for data generation by default.

## Configure SnapMirror replication between on-premises ONTAP and Cloud Volumes ONTAP

NetApp SnapMirror replicates data at high speeds over LAN or WAN, so you get high data availability and fast data replication in both virtual and traditional environments. When you replicate data to NetApp storage systems and continually update the secondary data, your data is kept current and remains available whenever you need it. No external replication servers are required.

Complete the following steps to configure SnapMirror replication between your on-premises ONTAP system and CVO.

1. From the navigation menu, select **Storage > Systems**.
2. In Systems, select the system that contains the source volume, drag it to the system to which you want to replicate the volume, and then select **Replication**.  
The remaining steps explain how to create a synchronous relationship between Cloud Volumes ONTAP and on-prem ONTAP clusters.
3. **Source and destination peering setup.** If this page appears, select all the intercluster LIFs for the cluster peer relationship.
4. **Source Volume Selection.** Select the volume that you want to replicate.
5. **Destination disk type and tiering.** If the target is a Cloud Volumes ONTAP system, select the destination disk type and choose whether you want to enable data tiering.

6. **Destination volume name:** Specify the destination volume name and choose the destination aggregate. If the destination is an ONTAP cluster, you must also specify the destination storage VM.
7. **Max transfer rate.** Specify the maximum rate (in megabytes per second) at which data can be transferred.
8. **Replication policy.** Choose a default policy or click **Additional Policies**, and then select one of the advanced policies. For help, [learn about replication policies](#).
9. **Schedule.** Choose a one-time copy or a recurring schedule. Several default schedules are available. If you want a different schedule, you must create a new schedule on the destination cluster using System Manager.
10. **Review.** Review your selections and click **Go**.

For detailed information about these configuration steps, see [here](#).

The Console starts the data replication process. At this stage, you can see the **Replication** service that was established between your on-premises ONTAP system and Cloud Volumes ONTAP.

In the Cloud Volumes ONTAP cluster, you can see the newly created volume.

You can also verify that the SnapMirror relationship is established between the on-premises volume and the cloud volume.

More information on the replication task can be found under the **Replication** tab.

[Next: Solution validation.](#)

## Solution validation

[Previous: SAN configuration.](#)

In this section, we review some solution use cases.

- One of the primary use cases for SnapMirror is data backup. SnapMirror can be used as a primary backup tool by replicating data within the same cluster or to remote targets.
- Using the DR environment to run application development testing (dev/test).
- DR in the event of a disaster in production.
- Data distribution and remote data access.

Notably, the relatively few use cases validated in this solution do not represent the entire functionality of SnapMirror replication.

### Application development and testing (dev/test)

To accelerate application development, you can quickly clone replicated data at the DR site and use it to dev/test applications. The colocation of DR and dev/test environments can significantly improve the utilization of backup or DR facilities, and on-demand dev/test clones provide as many data copies as you need to get to production more quickly.

NetApp FlexClone technology can be used to quickly create a read-write copy of a SnapMirror destination FlexVol volume in case you want to have read-write access of the secondary copy to confirm if all the production data is available.

Complete the following steps to use the DR environment to perform application dev/test:

1. Make a copy of production data. To do so, perform an application snapshot of an on-premises volume. Application snapshot creation consist of three steps: Lock, Snap, and Unlock.
  - a. Quiesce the file system so that I/O is suspended and applications maintain consistency. Any application writes hitting the filesystem stay in a wait state until the unquiesce command is issued in step c. Steps a, b, and c are executed through a process or a workflow that is transparent and does not affect the application SLA.

```
[root@hc-cloud-secure-1 ~]# fsfreeze -f /file1
```

This option requests the specified filesystem to be frozen from new modifications. Any process attempting to write to the frozen filesystem is blocked until the filesystem is unfrozen.

1. Create a snapshot of the on-prem volume.

```
A400-G0312::> snapshot create -vserver Healthcare_SVM -volume hc_iscsi_vol -snapshot kamini
```

1. Unquiesce the file system to restart I/O.

```
[root@hc-cloud-secure-1 ~]# fsfreeze -u /file1
```

This option is used to un-freeze the filesystem and allow operations to continue. Any filesystem modifications that were blocked by the freeze are unblocked and allowed to complete.

Application-consistent snapshot can also be performed using NetApp SnapCenter, which has the complete orchestration of the workflow outlined above as part of SnapCenter. For detailed information, see [here](#).

2. Perform a SnapMirror update operation to keep the production and DR systems in sync.

```
singlecvoaws::> snapmirror update -destination-path
svm_singlecvoaws:hc_iscsi_vol_copy -source-path
Healthcare_SVM:hc_iscsi_vol
```

```
Operation is queued: snapmirror update of destination
"svm_singlecvoaws:hc_iscsi_vol_copy".
```

A SnapMirror update can also be performed through the NetApp Console GUI under the **Replication** tab.

3. Create a FlexClone instance based on the application snapshot that was taken earlier.

```
singlecvoaws::> volume clone create -flexclone kamini_clone -type RW  
-parent-vserver svm_singlecvoaws -parent-volume hc_iscsi_vol_copy  
-junction-active true -foreground true -parent-snapshot kamini  
  
[Job 996] Job succeeded: Successful
```

For the previous task, a new snapshot can also be created, but you must follow the same steps as above to ensure application consistency.

4. Activate a FlexClone volume to bring up the EHR instance in the cloud.

```
singlecvoaws::> lun mapping create -vserver svm_singlecvoaws -path  
/vol/kamini_clone/iscsi_lun1 -igroup ehr-igroup -lun-id 0  
  
singlecvoaws::> lun mapping show  
Vserver      Path          Igroup      LUN ID  
Protocol  
-----  
-----  
svm_singlecvoaws  
/vol/kamini_clone/iscsi_lun1      ehr-igroup      0      iscsi
```

5. Execute the following commands on the EHR instance in the cloud to access the data or filesystem.
  - a. Discover ONTAP storage. Check the multipathing status.

```
sudo rescan-scsi-bus.sh
sudo iscsiaadm -m discovery -t sendtargets -p <iscsi-lif-ip>
sudo iscsiaadm -m node -L all
sudo sanlun lun show
```

Output:

controller(7mode/E-Series) / vserver(cDOT/FlashRay)	device	host	lun		
product	lun-pathname	filename	adapter	protocol	size
-----	-----	-----	-----	-----	-----
svm_singlecvoaws cDOT	/dev/sda	host2	iSCSI	200g	
	/vol/kamini_clone/iscsi_lun1				

```
sudo multipath -ll
```

Output:

```
3600a09806631755a452b543041313053 dm-0 NETAPP, LUN C-Mode
size=200G features='3 queue_if_no_path pg_init_retries 50'
hwhandler='1 alua' wp=rw
`-- policy='service-time 0' prio=50 status=active
`- 2:0:0:0 sda 8:0 active ready running
```

b. Activate the volume group.

```
sudo vgchange -ay datavg
Output:
1 logical volume(s) in volume group "datavg" now active
```

c. Mount the file system and display the summary of filesystem information.

```
sudo mount -t xfs /dev/datavg/datalv /file1
cd /file1
df -k .
Output:
Filesystem           1K-blocks   Used   Available   Use%
Mounted on
/dev/mapper/datavg-datalv  209608708 183987096  25621612   88%
/file1
```

This validates that you can use the DR environment for application dev/test. Performing application dev/test on your DR storage allows you to get more use out of resources that might otherwise sit idle much of the time.

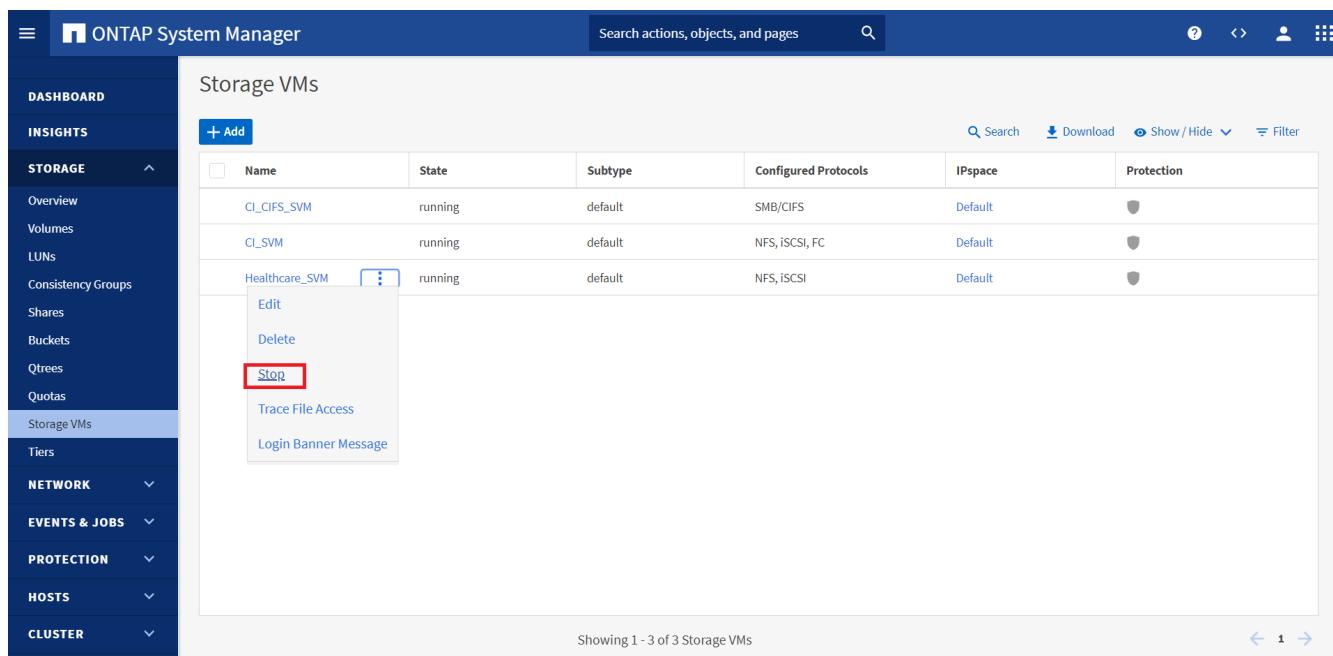
## Disaster recovery

SnapMirror technology is also used as a part of DR plans. If critical data is replicated to a different physical location, a serious disaster does not have to cause extended periods of data unavailability for business-critical applications. Clients can access replicated data across the network until the recovery of the production site from corruption, accidental deletion, natural disaster, and so on.

In the case of failback to the primary site, SnapMirror provides an efficient means of resynchronizing the DR site with the primary site, transferring only changed or new data back to the primary site from the DR site by simply reversing the SnapMirror relationship. After the primary production site resumes normal application operations, SnapMirror continues the transfer to the DR site without requiring another baseline transfer.

To perform the validation of a successful DR scenario, complete the following steps:

1. Simulate a disaster on the source (production) side by stopping the SVM that hosts the on-premises ONTAP volume (`hc_iscsi_vol`).



The screenshot shows the ONTAP System Manager interface. The left sidebar is the navigation menu with sections like DASHBOARD, INSIGHTS, STORAGE, NETWORK, EVENTS & JOBS, PROTECTION, HOSTS, and CLUSTER. The 'STORAGE' section is expanded, showing Overview, Volumes, LUNs, Consistency Groups, Shares, Buckets, Qtrees, Quotas, and Storage VMs. 'Storage VMs' is selected and highlighted in blue. The main content area is titled 'Storage VMs' and lists three VMs: CL\_CIFS\_SVM, CL\_SVM, and Healthcare\_SVM. The 'Healthcare\_SVM' row is selected, indicated by a blue border. A context menu is open over this row, with the 'Stop' option highlighted with a red box. Other options in the menu include Edit, Delete, Trace File Access, and Login Banner Message. The top navigation bar includes a search bar, a help icon, a download icon, a 'Show/Hide' dropdown, and a filter icon. The bottom of the screen shows a message 'Showing 1 - 3 of 3 Storage VMs'.

<input type="checkbox"/>	Name	State	Subtype	Configured Protocols	IPSpace	Protection
	CL_CIFS_SVM	running	default	SMB/CIFS	Default	🛡️
	CL_SVM	running	default	NFS, iSCSI, FC	Default	🛡️
	Healthcare_SVM	running	default	NFS, iSCSI	Default	🛡️

Make sure that SnapMirror replication is already set up between the on-premises ONTAP in FlexPod instance and Cloud Volumes ONTAP in AWS, so that you can create frequent application snapshots.

After the SVM has been stopped, the `hc_iscsi_vol` volume is not visible in the Console.

2. Activate DR in CVO.

- a. Break the SnapMirror replication relationship between on-prem ONTAP and Cloud Volumes ONTAP and promote the CVO destination volume (`hc_iscsi_vol_copy`) to production.

After the SnapMirror relationship is broken, the destination volume type changes from data protection (DP) to read/write (RW).

```
singlecvoaws::> volume show -volume hc_iscsi_vol_copy -fields type
server          volume          type
-----
svm_singlecvoaws hc_iscsi_vol_copy RW
```

b. Activate the destination volume in Cloud Volumes ONTAP to bring up the EHR instance on an EC2 instance in the cloud.

```
singlecvoaws::> lun mapping create -vserver svm_singlecvoaws -path
/vol/hc_iscsi_vol_copy/iscsi_lun1 -igroup ehr-igroup -lun-id 0

singlecvoaws::> lun mapping show
Vserver          Path          Igroup      LUN ID
Protocol
-----
-----
svm_singlecvoaws
/vol/hc_iscsi_vol_copy/iscsi_lun1 ehr-igroup 0      iscsi
```

c. To access the data and filesystem on the EHR instance in the cloud, first discover the ONTAP storage and verify multipathing status.

```
sudo rescan-scsi-bus.sh
sudo iscsiadm -m discovery -t sendtargets -p <iscsi-lif-ip>
sudo iscsiadm -m node -L all
sudo sanlun lun show
Output:
controller(7mode/E-Series) /           device      host      lun
vserver(cDOT/FlashRay) lun-pathname filename  adapter protocol size
product
-----
-----
svm_singlecvoaws           /dev/sda  host2  iSCSI  200g
cDOT
/vol/hc_iscsi_vol_copy/iscsi_lun1
sudo multipath -ll
Output:
3600a09806631755a452b543041313051 dm-0 NETAPP,LUN C-Mode
size=200G features='3 queue_if_no_path pg_init_retries 50'
hwandler='1 alua' wp=rw
`-+- policy='service-time 0' prio=50 status=active
`- 2:0:0:0 sda 8:0 active ready running
```

d. Then activate the volume group.

```
sudo vgchange -ay datavg
Output:
1 logical volume(s) in volume group "datavg" now active
```

e. Finally, mount the file system and display the filesystem information.

```
sudo mount -t xfs /dev/datavg/datalv /file1

cd /file1
df -k .
Output:
Filesystem           1K-blocks   Used   Available   Use%
Mounted on
/dev/mapper/datavg-datalv  209608708  183987096  25621612   88%
/file1
```

This output shows that users can access replicated data across the network until the recovery of the production site from disaster.

f. Reverse the SnapMirror relationship. This operation reverses the roles of the source and destination volumes.

When this operation is performed, the contents from the original source volume are overwritten by the contents of the destination volume. This is helpful when you want to reactivate a source volume that went offline.

Now the CVO volume (hc\_iscsi\_vol\_copy) becomes the source volume, and the on-premises volume (hc\_iscsi\_vol) becomes the destination volume.

Any data written to the original source volume between the last data replication and the time that the source volume was disabled is not preserved.

g. To verify write access to the CVO volume, create a new file on the EHR instance in the cloud.

```
cd /file1/
sudo touch newfile
```

When the production site is down, clients can still access the data and also perform writes to the Cloud Volumes ONTAP volume, which is now the source volume.

In the case of failback to the primary site, SnapMirror provides an efficient means of resynchronizing the DR site with the primary site, transferring only changed or new data back to the primary site from the DR site by simply reversing the SnapMirror relationship. After the primary production site resumes normal application operations, SnapMirror continues the transfer to the DR site without requiring another baseline transfer.

This section illustrates the successful resolution of a DR scenario when the production site is hit by disaster. Data can now be safely consumed by applications that can now serve the clients while the source site goes through restoration.

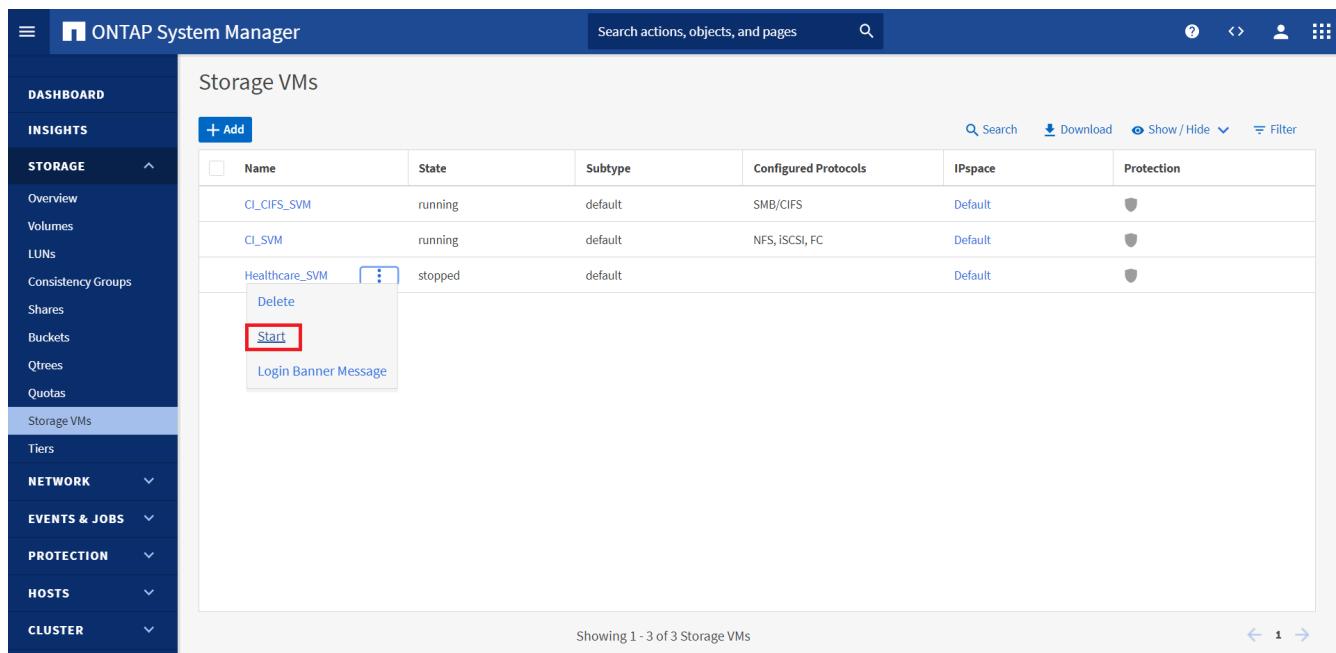
## Verification of data on the production site

After the production site is restored, you must make sure that the original configuration is restored and clients are able to access the data from the source site.

In this section, we talk about bringing up the source site, restoring the SnapMirror relationship between on-premises ONTAP and Cloud Volumes ONTAP, and finally performed a data integrity check on the source end.

The following procedure can be used for the verification of data on the production site:

1. Make sure that the source site is now up. To do so, start the SVM that hosts the on-premises ONTAP volume (`hc_iscsi_vol`).



The screenshot shows the ONTAP System Manager interface. The left sidebar is collapsed. The main area is titled 'Storage VMs' and contains a table with three rows. The first two rows are 'CLCIFS\_SVM' and 'CL\_SVM', both in 'running' state. The third row is 'Healthcare\_SVM', which is 'stopped'. Below the table, there is a context menu with options: 'Delete', 'Start' (which is highlighted with a red box), and 'Login Banner Message'. The table has columns: Name, State, Subtype, Configured Protocols, IPspace, and Protection. The 'Configured Protocols' column for 'Healthcare\_SVM' shows 'NFS, iSCSI, FC'. The 'Protection' column for 'Healthcare\_SVM' shows a shield icon. The bottom of the screen shows a message 'Showing 1 - 3 of 3 Storage VMs' and navigation arrows.

2. Break the SnapMirror replication relationship between Cloud Volumes ONTAP and on-premises ONTAP and promote the on-premises volume (`hc_iscsi_vol`) back to production.

After the SnapMirror relationship is broken, the on-premises volume type changes from data protection (DP) to read/write (RW).

```
A400-G0312::> volume show -volume hc_iscsi_vol -fields type
vserver          volume          type
-----
Healthcare_SVM  hc_iscsi_vol  RW
```

3. Reverse the SnapMirror relationship. Now, the on-premises ONTAP volume (`hc_iscsi_vol`) becomes the source volume as it was earlier, and the Cloud Volumes ONTAP volume (`hc_iscsi_vol_copy`) becomes the destination volume.

By following these steps, we have successfully restored the original configuration.

4. Reboot the on-premises EHR instance. Mount the filesystem and verify that the `newfile` that you created on the EHR instance in the cloud when production was down now exists here as well.

```
[root@hc-cloud-secure-1 ~]# mount -t xfs /dev/datavg/datalv /file1
[root@hc-cloud-secure-1 ~]# cd /file1/
[root@hc-cloud-secure-1 file1]# ls
dir01  dir05  dir09  dir13  dir17  dir21  dir25  dir29  dir33  dir37  dir41  dir45  dir49  dir53  dir57  dir61  dir65  dir69  dir73  dir77  kamini
dir02  dir06  dir10  dir14  dir18  dir22  dir26  dir30  dir34  dir38  dir42  dir46  dir50  dir54  dir58  dir62  dir66  dir70  dir74  dir78  latest file
dir03  dir07  dir11  dir15  dir19  dir23  dir27  dir31  dir35  dir39  dir43  dir47  dir51  dir55  dir59  dir63  dir67  dir71  dir75  dir79
dir04  dir08  dir12  dir16  dir20  dir24  dir28  dir32  dir36  dir40  dir44  dir48  dir52  dir56  dir60  dir64  dir68  dir72  dir76  dir80
```

We can infer that the data replication from the source to the destination has been completed successfully and that data integrity has been maintained. This completes the verification of data on the production site.

[Next: Conclusion.](#)

## Conclusion

[Previous: Solution validation.](#)

Building a hybrid cloud is a goal for most healthcare organizations to provide data availability at any time. In this solution, we implemented a FlexPod hybrid cloud solution with Cloud Volumes ONTAP, utilizing NetApp SnapMirror replication technology to validate some use cases to back up and recover healthcare applications and workloads.

FlexPod, a rigorously tested and prevalidated converged infrastructure from the strategic partnership of Cisco and NetApp is designed to deliver predictable low-latency system performance and high availability. This approach results in EHR high comfort levels and ultimately the best response time for users of the EHR system.

With NetApp, you can run EHR production, disaster recovery, backup, or tiering in the cloud just like you would run NetApp storage features in an on-premises datacenter. With NetApp Cloud Volumes ONTAP, NetApp provides the enterprise-class capabilities and the performance required to effectively run EHR in the cloud. NetApp cloud options provide block-over-iSCSI and file-over-NFS or SMB.

This solution caters to the need of healthcare organizations and enables them to take a step towards their digital transformation. It can also help them manage their applications and workloads in an efficient manner.

[Next: Where to find additional information.](#)

## Where to find additional information

[Previous: Conclusion.](#)

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](#)

- NetApp Console

<https://console.netapp.com/>

- NetApp Cloud Volumes ONTAP

<https://docs.netapp.com/us-en/cloud-manager-cloud-volumes-ontap/concept-overview-cvo.html>

- Quick start for Cloud Volumes ONTAP in AWS

<https://docs.netapp.com/us-en/cloud-manager-cloud-volumes-ontap/task-getting-started-aws.html>

- SnapMirror Replication

<https://docs.netapp.com/us-en/cloud-manager-replication/concept-replication.html>

- TR-3928: NetApp best practices for Epic

<https://www.netapp.com/pdf.html?item=/media/17137-tr3928pdf.pdf>

- TR-4693: FlexPod Datacenter for Epic EHR Deployment Guide

<https://www.netapp.com/media/10658-tr-4693.pdf>

- FlexPod for Epic

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_xseries\\_vmw\\_epic.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_xseries_vmw_epic.html)

- NetApp Interoperability Matrix Tool

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

- Cisco UCS Hardware and Software Interoperability Tool

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

- VMware Compatibility Guide

<http://www.vmware.com/resources/compatibility/search.php>

## Version history

Version	Date	Document version history
Version 1.0	March 2023	Initial version

# FlexPod Hybrid Cloud for Google Cloud Platform with NetApp Cloud Volumes ONTAP and Cisco Intersight

# TR-4939: FlexPod Hybrid Cloud for Google Cloud Platform with NetApp Cloud Volumes ONTAP and Cisco Intersight

Ruchika Lahoti, NetApp

## Introduction

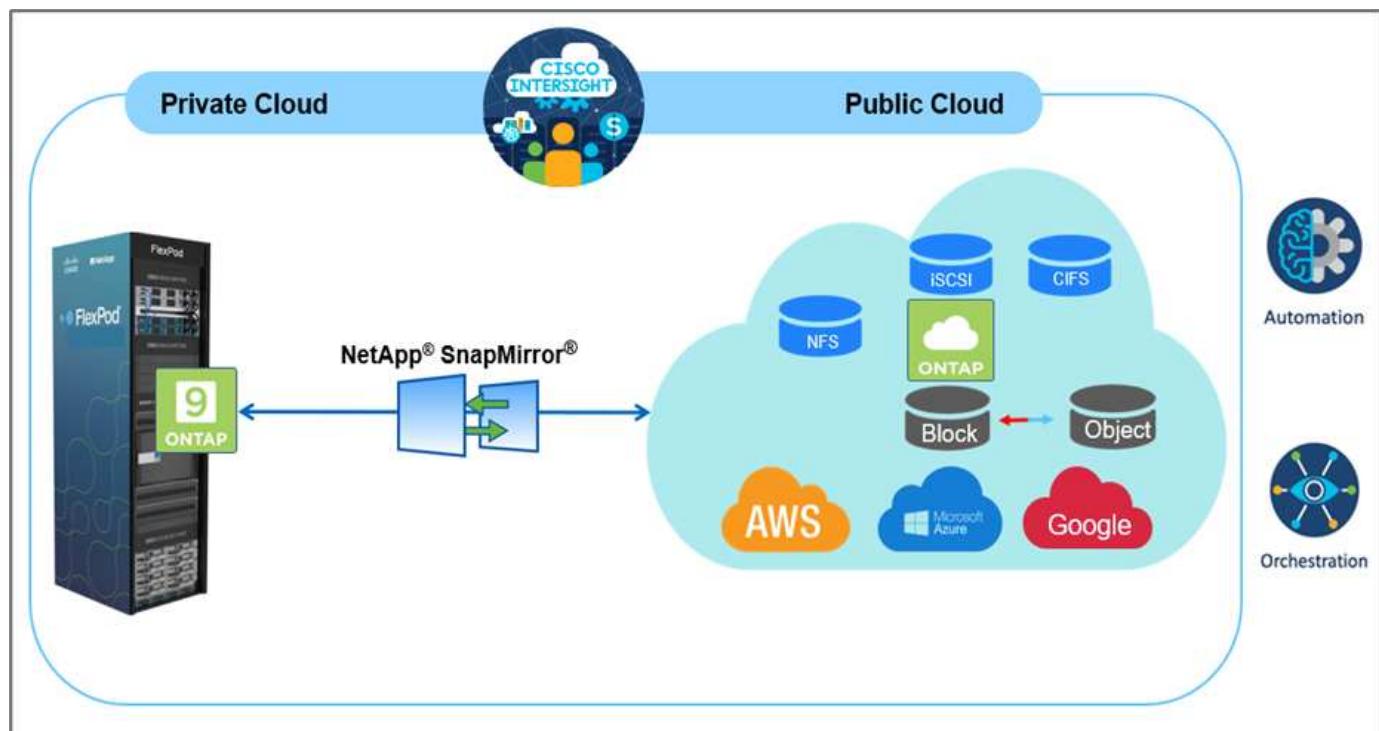
Protecting data with disaster recovery (DR) is a critical goal for businesses continuity. DR allows organizations to failover their business operations to a secondary location and later recover and failback to the primary site efficiently and reliably. Multiple concerns like natural disaster, network failures, software vulnerabilities, and human error make developing a DR strategy a top IT priority.

For DR, all workloads running on the primary site must be faithfully reproduced on the DR site. An organization must also have an up-to-date copy of all enterprise data, including database, file services, NFS and iSCSI storage, and so on. Because data in the production environment is constantly updated, changes must be transferred to the DR site on a regular basis.

Deploying DR environments is challenging for most organizations because of the requirement for infrastructure and site independence. The number of resources needed and the costs of setting up, testing, and maintaining a secondary data center can be very high, typically approaching the cost of the entire production environment. It is challenging to keep a minimal data footprint with adequate protection, while continuously synchronizing data and establishing seamless failover and failback. After building out the DR site, the challenge then becomes to replicate data from the production environment and to keep it synchronized going forward.

This technical report brings together the FlexPod converged infrastructure solution, NetApp Cloud Volumes ONTAP on Google Cloud, and Cisco Intersight to form a hybrid cloud data center for DR. In this solution we discuss designing and executing an on-premises ONTAP workflow using Cisco Intersight Cloud Orchestrator. We also discuss deploying NetApp Cloud Volumes ONTAP and orchestrating and automating data replication and DR between FlexPod and Cloud Volumes ONTAP using the Cisco Intersight Service for HashiCorp Terraform.

The following figure provide a solution overview.



This solution provides multiple advantages, including:

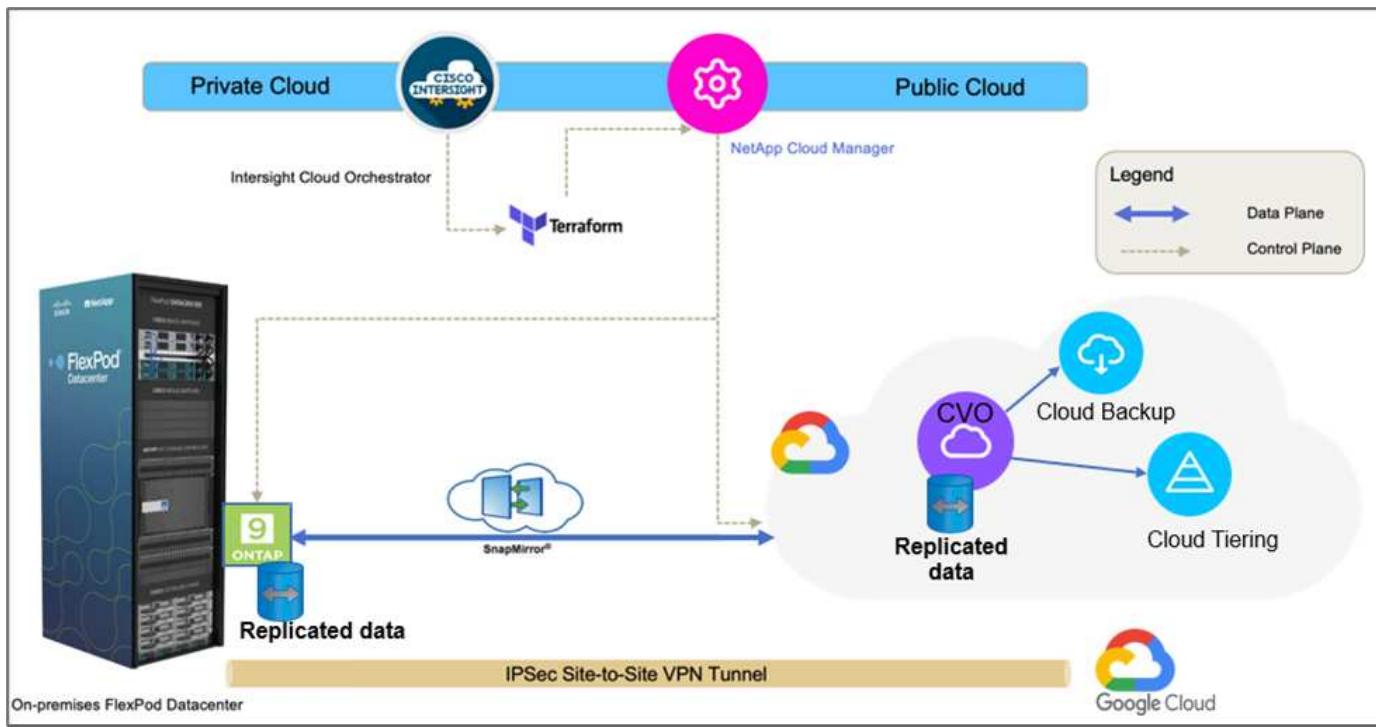
- **Orchestration and automation.** Cisco Intersight simplifies the day-to-day operations of FlexPod hybrid cloud infrastructure by providing consistent orchestration frameworks that are delivered via automation.
- **Customized Protection.** Cloud Volumes ONTAP provides block-level data replication from ONTAP to the cloud that keeps the destination up to date through incremental updates. Users can specify a synchronization schedule of every 5 minutes or every hour, for example, based on changes at the source that are transferred over.
- **Seamless failover and failback.** When a disaster occurs, storage administrators can quickly failover to the cloud volumes. When the primary site is recovered, the new data created in the DR environment is synchronized back to the source volumes, re-establishing secondary data replication.
- **Efficiency:** The storage space and costs for the secondary cloud copy are optimized through the use of data compression, thin provisioning, and deduplication. Data is transferred at the block-level in a compressed and deduplicated form, improving transfer speed. Data is also automatically tiered to low-cost object storage and only brought back to high-performance storage when accessed, such as in a DR scenario. This significantly reduces ongoing storage costs.
- **Increased IT productivity.** Using Intersight as the single secure, enterprise-grade platform for infrastructure and application lifecycle management simplifies configuration management and automation of manual tasks at scale for the solution.

## Audience

The audience for this document includes, but is not limited to; sales engineers, field consultants, professional services, IT managers, partner engineers, site reliability engineers, cloud architects, cloud engineers, and customers who want to take advantage of an infrastructure built to deliver IT efficiency and enable IT innovation.

## Solution topology

This section describes the logical topology of the solution. The following figure represents the solution topology of the on-premises FlexPod environment, NetApp Cloud Volumes ONTAP running on Google Cloud, Cisco Intersight, and NetApp Cloud Manager.



The control planes and data planes are clearly indicated between the endpoints. The data plane uses a secure site-to-site VPN connection to connect the ONTAP instance running on FlexPod All Flash FAS to the NetApp Cloud Volumes ONTAP instance on Google Cloud.

The replication of workload data from FlexPod to NetApp Cloud Volumes ONTAP is handled by NetApp SnapMirror, and the overall process is orchestrated using Cisco Intersight Cloud Orchestrator for both the on-premises and cloud environments. Cisco Intersight Cloud Orchestrator consumes Terraform Resource Providers for NetApp Cloud Manager to carry out operations related to NetApp Cloud Volumes ONTAP deployment and establish data replication relationships.



The optional backup and tiering of cold data residing in the NetApp Cloud Volumes ONTAP instance to Google Cloud Storage is also supported with this solution.

[Next: Solution components.](#)

## Solution components

[Previous: Solution overview.](#)

### FlexPod

FlexPod is a defined set of hardware and software that forms an integrated foundation for both virtualized and non-virtualized solutions. FlexPod includes NetApp ONTAP storage, Cisco Nexus networking, Cisco MDS storage networking, and 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.

### 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 to 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 the customer can 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 the 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 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 native UMM or IMM for fabric-attached Cisco UCS systems during initial setup of fabric interconnects. In this solution, native IMM is used.

### Cisco Intersight licensing

Cisco Intersight uses a subscription-based license with multiple tiers.

Cisco Intersight license tiers are as follows:

- **Cisco Intersight Essentials.** Includes all base functionality plus the following features:
  - Cisco UCS Central
  - Cisco IMC Supervisor entitlement
  - Policy-based configuration with Server Profiles
  - Firmware management
  - Valuation of compatibility with the Hardware Compatibility List (HCL)
- **Cisco Intersight Advantage.** Includes all the features and functionality of the Essentials tier plus the following features:
  - Widgets, inventory, capacity, utilization features, and cross-domain inventory correlation across physical compute, network, storage, VMware virtualization, and AWS public cloud.
  - The Cisco Security Advisory service where customers can receive important security alerts and field notices about impacted endpoint devices.
- **Cisco Intersight Premier.** In addition to the capabilities provided in the Advantage tier, Cisco Intersight Premier offers the following:
  - Intersight Cloud Orchestrator (ICO) for Cisco and third-party compute, network, storage, integrated systems, virtualization, container, and public-cloud platforms
  - Full subscription entitlement for Cisco UCS Director at no additional cost.

More information about Intersight Licensing and features supported in each licensing can be found [here](#).



In this solution, we use Intersight Cloud Orchestrator and Intersight Service for HashiCorp Terraform. These features are available for users with the Intersight Premier license, so this licensing tier must be enabled.

## Terraform Cloud Integration with ICO

You can use Cisco Intersight Cloud Orchestrator (ICO) to create and execute workflows that call Terraform Cloud (TFC) APIs. The Invoke Web API Request task supports Terraform Cloud as a target, and it can be configured with Terraform Cloud APIs using HTTP methods. So, the workflow can have a combination of tasks that calls multiple Terraform Cloud APIs using generic API tasks and other operations. You need a Premier license to use the ICO feature.

### Cisco Intersight Assist

Cisco Intersight Assist helps you add endpoint devices to Cisco Intersight. A data center could have multiple devices that do not connect directly with Cisco Intersight. Any device that is supported by Cisco Intersight but does not connect directly to it requires a connection mechanism. Cisco Intersight Assist provides that connection mechanism and helps you add devices into Cisco Intersight.

Cisco Intersight Assist is available within the Cisco Intersight Virtual Appliance, which is distributed as a deployable virtual machine contained within an Open Virtual Appliance (OVA) file format. You can install the appliance on an ESXi server. For more information, see the [Cisco Intersight Virtual Appliance Getting Started Guide](#).

After claiming Intersight Assist into Intersight, you can claim endpoint devices using the Claim Through Intersight Assist option. For more information, see [Getting Started](#).

### NetApp Cloud Volumes ONTAP

- Leveraging built-in data deduplication, data compression, thin provisioning, and cloning to minimize storage costs.
- Providing enterprise reliability and continuous operations in case of failures in your cloud environment.
- Cloud Volumes ONTAP uses NetApp SnapMirror, 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 the Cloud Backup service to deliver backup and restore capabilities for protection and long-term archiving of your cloud data.
- Switching between high and low-performance storage pools on-demand without taking applications offline.
- Providing 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 SaaS 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 cloud solutions. It provides a centralized system for viewing and managing your on-premises and cloud storage to support multiple hybrid-cloud providers and accounts. For more information, see [Cloud Manager](#).

## Connector

Connector enables Cloud Manager to manage resources and processes within a public cloud environment. A Connector instance is required to use many features provided by Cloud Manager and can be deployed in the cloud or on-premises network. Connector is supported in the following locations:

- AWS
- Microsoft Azure
- Google Cloud
- On premises

## 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 running on it. When an issue occurs with the storage infrastructure, Unified Manager can notify you about the details of the issue to help identify the root cause. The virtual machine dashboard gives you a view into the performance statistics for the VM so that you can investigate the entire I/O path from the vSphere host down through the network and finally to the storage.

Some events also provide remedial actions that you can take 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.

## VMware vSphere

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 about VMware vSphere, follow [this link](#).

## 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 versions

This hybrid cloud 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 the Cisco UCS Hardware Compatibility List.

The FlexPod solution that is used as a baseline platform in our on-premises environment was deployed according to the guidelines and specifications described [here](#).

The network within this environment is ACI- based. For more information, see [here](#).

- See the following links for more information:
- [NetApp Interoperability Matrix Tool](#)
- [VMware Compatibility Guide](#)
- [Cisco UCS Hardware and Software Interoperability Tool](#)

The following table shows the FlexPod hardware and software revisions.

Component	Product	Version
Compute	Cisco UCS X210C-M6	5.0(1b)
	Cisco UCS Fabric Interconnects 6454	4.2(2a)
Network	Cisco Nexus 9332C (Spine)	14.2(7s)
	Cisco Nexus 9336C-FX2 (Leaf)	14.2(7s)
	Cisco ACI	4.2(7s)
Storage	NetApp AFF A220	9.11.1
	NetApp ONTAP Tools for VMware vSphere	9.10
	NetApp NFS Plugin for VMware VAAI	2.0-15
Software	Active IQ Unified Manager	9.11
	vSphere ESXi	7.0(U3)
	VMware vCenter Appliance	7.0.3
	Cisco Intersight Assist Virtual Appliance	1.0.11-306

The execution of Terraform configurations happens on the Terraform Cloud for Business account. Terraform configuration uses the Terraform provider for NetApp Cloud Manager.

The following table lists the vendors, products, and versions.

Component	Product	Version
HashiCorp	Terraform	1.2.7

The following table shows the Cloud Manager and Cloud Volumes ONTAP versions.

Component	Product	Version
NetApp	Cloud Volumes ONTAP	9.11
	Cloud Manager	3.9.21

Next: [Installation and configuration - Deploy FlexPod](#).

## Installation and configuration

### Deploy FlexPod

[Previous: Solution components.](#)

To understand the FlexPod design and deployment details, including the configuration of various elements of design and associated best practices, see [Cisco Validated Designs for FlexPod](#).

FlexPod can be deployed in both UCS Managed Mode and Cisco Intersight Managed Mode. If you are deploying FlexPod in UCS Managed Mode, the latest Cisco Validated Design can be found [here](#).

Cisco Unified Compute System (Cisco UCS) X-Series is a brand new modular compute system, configured and managed from the cloud. It is designed to meet the needs of modern applications and to improve operational efficiency, agility, and scale through an adaptable, future-ready, modular design. The design guidance around incorporating the Cisco Intersight- managed UCS X-Series platform within FlexPod infrastructure can be found [here](#).

FlexPod with Cisco ACI deployment can be found [here](#).

[Next: Cisco Intersight configuration.](#)

### Cisco Intersight configuration

[Previous: Deploy FlexPod.](#)

To configure Cisco Intersight and Intersight Assist, see the Cisco Validated Designs for FlexPod found [here](#).

[Next: Terraform Cloud Integration with ICO prerequisite.](#)

### Terraform Cloud Integration with ICO prerequisite

[Previous: Cisco Intersight configuration.](#)

#### Procedure 1: Connect Cisco Intersight and Terraform Cloud

1. Claim or create a Terraform cloud target by providing the relevant Terraform Cloud account details.
2. Create a Terraform Cloud Agent target for private clouds so that customers can install the agent in the data center and enable communication with Terraform Cloud.

For more information, follow [this link](#).

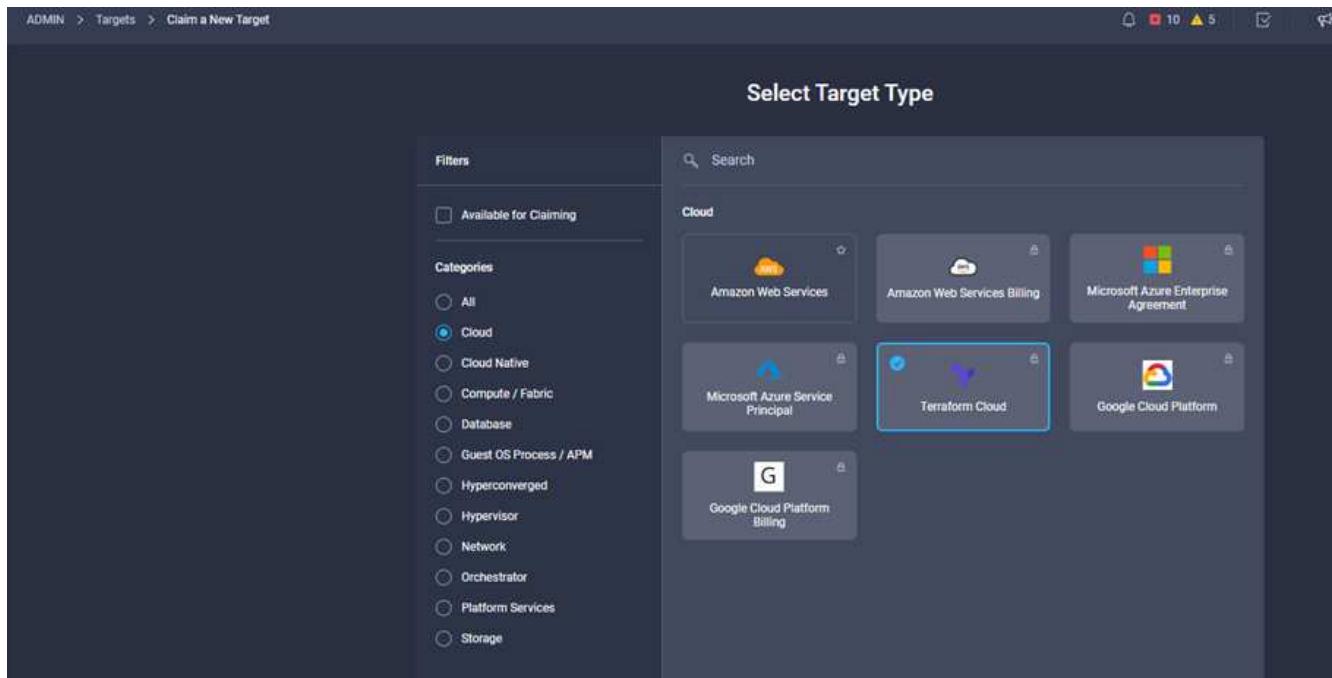
#### Procedure 2: Generate user token

As a part of adding a target for Terraform Cloud, you must provide the username and API token from the Terraform Cloud settings page.

1. Login to Terraform Cloud and go to **User Tokens**: <https://app.terraform.io/app/settings/tokens>.
2. Click **Create a new API token**.
3. Assign a name to remember and save the token in a secure place.

### Procedure 3: Claim Terraform Cloud Target

1. Log into Intersight with Account Administrator, Device Administrator, or Device Technician privileges.
2. Navigate to **ADMIN > Targets > Claim a New Target**.
3. In **Categories**, click **Cloud**.
4. Click **Terraform Cloud** and click **Start**.



5. Enter a name for the target, your username for the Terraform Cloud, the API token, and a default organization in Terraform Cloud as displayed in the following image.
6. In the **Default Managed Hosts** field, make sure to add the following links along with other managed hosts:
  - [github.com](https://github.com)
  - [github-releases.githubusercontent.com](https://github-releases.githubusercontent.com)

A screenshot of the 'Claim Terraform Cloud Target' form. It includes fields for 'Name' (TFCB), 'Terraform Cloud Username' (abhinav3), 'Terraform Cloud API Token' (redacted), 'Default Terraform Cloud Organization' (cisco-intersight-gc), and 'Default Managed Hosts' (github.com,github-releases.githubusercontent.com).

If everything is correctly entered, you will see your Terraform Cloud target displayed in the **Intersight Targets** section.

#### Procedure 4: Add Terraform Cloud agents

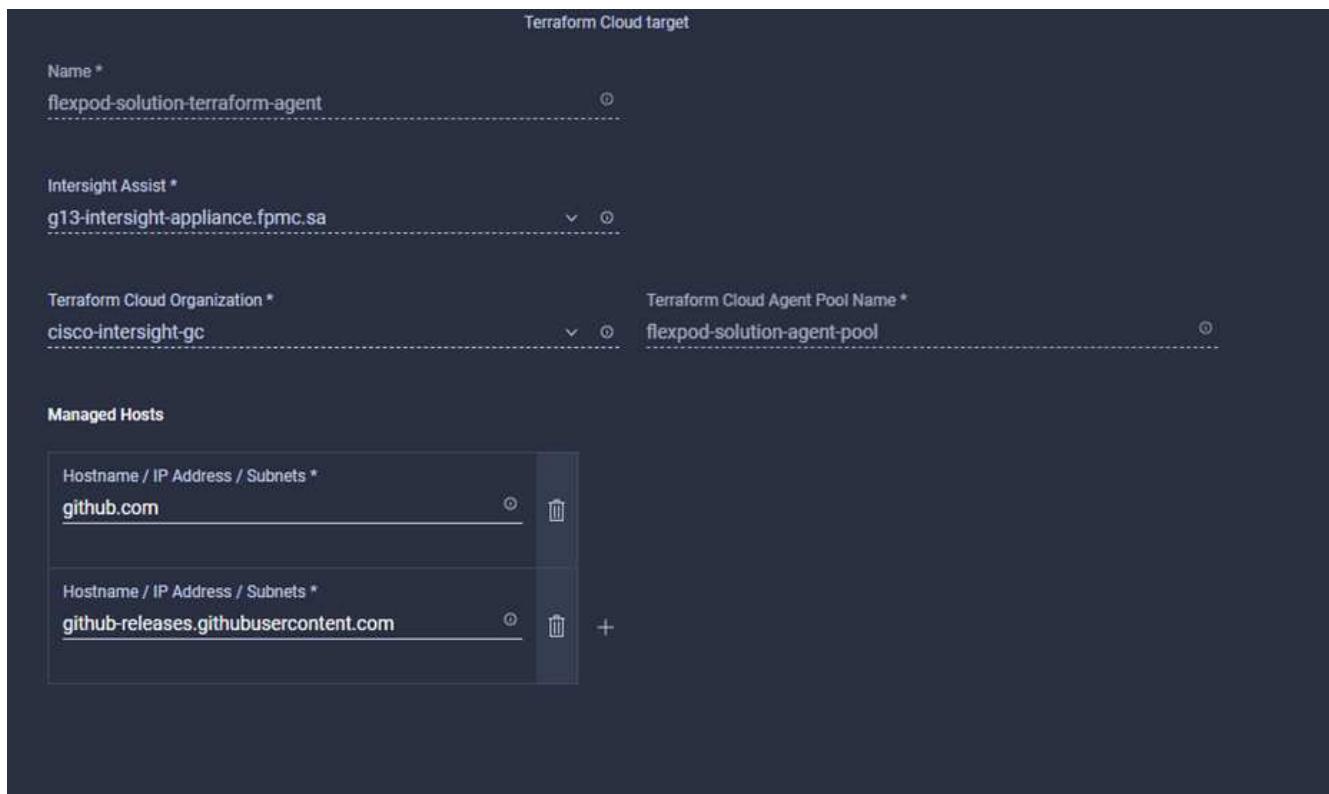
Prerequisites:

- Terraform Cloud target.
- Claimed Intersight Assist into Intersight before deploying the Terraform Cloud Agent.

 You can only claim five agents for each Assist.

 After you have created the connection to Terraform, you must spin up a Terraform Agent to execute the Terraform code.

1. Click **Claim Terraform Cloud Agent** from the drop-down list of your Terraform Cloud target.
2. Enter the details for the Terraform Cloud agent. The following screenshot shows the configuration details for Terraform agent.



 You can update any Terraform Agent property. If the target is in the **Not Connected** state and has never been in the **Connected** state, then a token has not been generated for the Terraform agent.

After the agent validation succeeds and an agent token is generated, you are unable to reconfigure the Organization and/or Agent Pool. Successful deployment of a Terraform agent is indicated by a status of **Connected**.

After you have enabled and claimed the Terraform Cloud integration, you can deploy one or more Terraform Cloud agents in Cisco Intersight Assist. The Terraform Cloud agent is modelled as a child target of the Terraform Cloud target. When you claim the agent target, you see a message to indicate that the target claim is in progress.

After a few seconds, the target is moved to the **Connected** state, and the Intersight platform routes HTTPS packets from the agent to the Terraform Cloud gateway.

Your Terraform Agent should be correctly claimed and should show up under targets as **Connected**.

Next: [Configure Public Cloud Service provider](#).

## Configure Public Cloud Service provider

Previous: [Terraform Cloud Integration with ICO prerequisite](#).

### Procedure 1: Access NetApp Cloud Manager

To access NetApp Cloud Manager and other cloud services, you need to sign up on [NetApp Cloud Central](#).



For setting up workspaces and users in the Cloud Central account, click [here](#).

### Procedure 2: Deploy Connector

To deploy Connector in Google Cloud, see this [link](#).

Next: [Automated deployment of Hybrid Cloud NetApp Storage](#).

## Automated deployment of Hybrid Cloud NetApp Storage

Previous: [Configure Public Cloud Service provider](#).

### Google Cloud

You must first enable APIs and create a service account that provides Cloud Manager with permissions to deploy and manage Cloud Volumes ONTAP systems that are in the same project as the Connector or in different projects.

Before you deploy a connector in a Google Cloud project, make sure that the connector isn't running on your premises or in a different cloud provider.

Two sets of permissions must be in place before you deploy a Connector directly from Cloud Manager:

- You need to deploy Connector using a Google account that has permissions to launch the Connector VM instance from Cloud Manager.
- When deploying Connector, you are prompted to select the VM instance. Cloud Manager gets permissions from the service account to create and manage Cloud Volumes ONTAP systems on your behalf. Permissions are provided by attaching a custom role to the service account. You need to set up two YAML files that include the required permissions for the user and the service account. Learn how to use [the YAML files to set up permissions](#) here.

See [this detailed video](#) for all required prerequisites.

## Cloud Volumes ONTAP deployment modes and architecture

Cloud Volumes ONTAP is available in Google Cloud as a single- node system and as a high-availability (HA) pair of nodes. Based on the requirements, we can choose the Cloud Volumes ONTAP deployment mode. Upgrading a single node system to an HA pair is not supported. If you want to switch between a single- node system and an HA pair, then you must deploy a new system and replicate data from the existing system to the

new system.

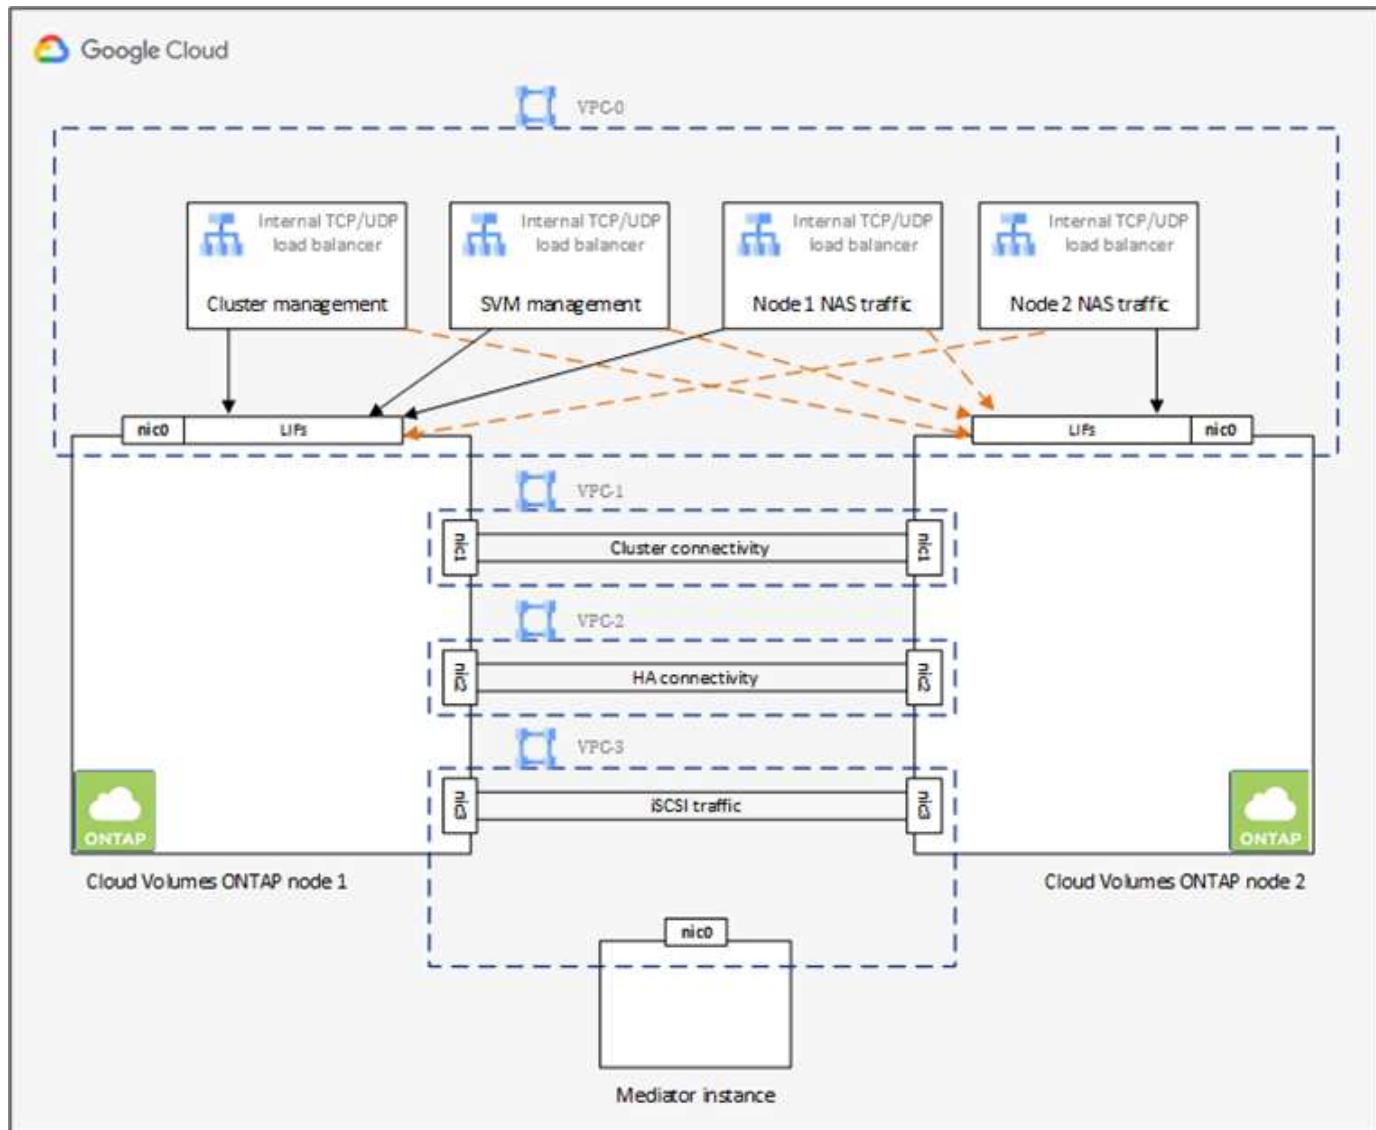
## Highly available Cloud Volumes ONTAP in Google Cloud

Google Cloud supports deployment of resources across multiple geographical regions and multiple zones within a region. The HA deployment consists of two ONTAP nodes that use powerful n1-standard or n2-standard machine types available in Google Cloud. Data is synchronously replicated between the two Cloud Volumes ONTAP nodes to provide availability in the event of a failure. HA deployment of Cloud Volumes ONTAP requires four VPCs and a private subnet in each VPC. The subnets in the four VPCs should be provisioned with non-overlapping CIDR ranges.

The four VPCs are used for the following purposes:

- VPC 0 enables inbound communication to data and Cloud Volumes ONTAP nodes.
- VPC 1 provides cluster connectivity between Cloud Volumes ONTAP nodes.
- VPC 2 allows for non-volatile ram (NVRAM) replication between nodes.
- VPC 3 is used for connectivity to the HA mediator instance and disk replication traffic for node rebuilds.

The following image shows a highly available Cloud Volumes ONTAP in Goggle Cloud.



For details, see [this link](#).

For networking requirements for Cloud Volumes ONTAP in Google Cloud, see [this link](#).

For details about data tiering, see [this link](#).

## Set up environment prerequisites

The automated creation of Cloud Volumes ONTAP clusters, SnapMirror configuration between an on-premises volume and a Cloud volume, creating a cloud volume, and so on are performed using Terraform configuration. These Terraform configurations are hosted on a Terraform Cloud for Business account. Using Intersight Cloud Orchestrator, you orchestrate tasks like creating a workspace in a Terraform Cloud for Business account, add all required variables to the workspace, execute a Terraform Plan, and so on.

For these automation and orchestration tasks, there are a few requirements and data needed, as is described in the following sections.

### GitHub repository

You need a GitHub account to host your Terraform code. Intersight Orchestrator creates a new workspace in the Terraform Cloud for Business account. This workspace is configured with a version control workflow. For this purpose, you need to keep the Terraform configuration in a GitHub repository and provide it as an input while creating the workspace.

[This GitHub link](#) provides the Terraform configuration with various resources. You can fork this repository and make a copy in your GitHub account.

In this repository, `provider.tf` has the definition for the required Terraform provider. Terraform provider for NetApp Cloud Manager is used.

`variables.tf` has all the variable declarations. The value for these variables is input as the Intersight Cloud Orchestrator's workflow input. This provides a convenient way to pass values to a workspace and execute the Terraform configuration.

`resources.tf` defines the various resources needed to add an on-premises ONTAP to the working environment, create a single node Cloud Volumes ONTAP cluster on Google Cloud, establish a SnapMirror relationship between on-premises and Cloud Volumes ONTAP, create a cloud volume on Cloud Volumes ONTAP, and so on.

In this repository:

- `provider.tf` has NetApp Cloud Manager as a definition for the required Terraform provider.
- `variables.tf` has the variable declarations that are used as input for the Intersight Cloud Orchestrator workflow. This provides a convenient way to pass values to workspace and execute Terraform configuration.
- `resources.tf` defines various resources to add an on-premises ONTAP to the working environment, create a single- node Cloud Volumes ONTAP cluster on Google Cloud, establish a SnapMirror relationship between on-premises and Cloud Volumes ONTAP, create a cloud volume on Cloud Volumes ONTAP, and so on.

You can add an additional resource block to create multiple volumes on Cloud Volumes ONTAP or use `count` or `for_each` Terraform constructs.

To connect Terraform workspaces, modules, and policy sets to git repositories containing Terraform configurations, Terraform Cloud needs access to your GitHub repo.

Add a client, and the OAuth Token ID of the client is used as one of the Intersight Cloud Orchestrator's workflow input.

1. Log in to your Terraform Cloud for Business account. Navigate to **Settings > Providers**.
2. Click **Add a VCS provider**.
3. Select your version.
4. Follow the steps under **Set up provider**.
5. You see the added client in **VCS Providers**. Make a note of the OAuth Token ID.

### Refresh token for NetApp Cloud Manager API operations

In addition to the web browser interface, Cloud Manager has a REST API that provides software developers with direct access to the Cloud Manager functionality through the SaaS interface. The Cloud Manager service consists of several distinct components that collectively form an extensible development platform. The refresh token enables you to generate access tokens that you add to the Authorization header for each API call.

Without calling an API directly, the netapp-cloudmanager provider uses a refresh token and translates the Terraform resources into corresponding API calls. You need to generate a refresh token for NetApp Cloud Manager API operations from [NetApp Cloud Central](#).

You need the client ID of the Cloud Manager Connector to create resources on Cloud Manager such as creating a Cloud Volumes ONTAP cluster, configuring SnapMirror, and so on.

1. Log into Cloud Manager: <https://cloudmanager.netapp.com/>.
2. Click **Connector**.
3. Click **Manage Connectors**.
4. Click the ellipses and copy the Connector ID.

### Develop Cisco Intersight Cloud Orchestrator workflow

Cisco Intersight Cloud Orchestrator is available in Cisco Intersight if:

- You have installed the Intersight Premier license.
- You are either an account administrator, storage administrator, virtualization administrator, or server administrator and have a minimum of one server assigned to you.

### Workflow Designer

The Workflow Designer helps you create new workflows (as well as tasks and data types) and edit existing workflows to manage targets in Cisco Intersight.

To launch the Workflow Designer, go to **Orchestration > Workflows**. A dashboard displays the following details under the tabs **My Workflows**, **Sample Workflows**, and **All Workflows**:

- Validation Status
- Last Execution Status
- Top Workflows by Execution Count

- Top Workflow Categories
- Number of System Defined Workflows
- Top Workflows by Targets

Using the dashboard, you can create, edit, clone, or delete a tab. To create your own custom view tab, click **+**, specify a name, and then select the required parameters that need to be displayed in the columns, tag columns, and widgets. You can rename a tab if it doesn't have a **Lock** icon.

Under the dashboard is a tabular list of workflows displaying the following information:

- Display Name
- Description
- System Defined
- Default Version
- Executions
- Last Execution Status
- Validation Status
- Last Update
- Organization

The Actions column allows you to perform the following actions:

- **Execute.** Executes the workflow.
- **History.** Displays workflow execution history.
- **Manage Versions.** Create and manage versions for workflows.
- **Delete.** Delete a workflow.
- **Retry.** Retry a failed workflow.

## Workflow

Create a workflow that consists of the following steps:

- **Defining a workflow.** Specify the display name, description, and other important attributes.
- **Define workflow inputs and workflow outputs.** Specify which input parameters are mandatory for the workflow execution, and the outputs generated on successful execution
- **Add workflow tasks.** Add one or more workflow tasks in the Workflow Designer that are needed for the workflow to carry out its function.
- \*Validate the workflow. \*Validate a workflow to ensure that there are no errors in connecting task inputs and outputs.

### Create workflows for on-premises FlexPod storage

To configure a workflow for on-premises FlexPod storage, see [this link](#).

Next: [DR workflow](#).

## DR workflow

Previous: [Automated deployment of Hybrid Cloud NetApp Storage.](#)

The sequence of steps are as follows:

1. Define the workflow.
  - Create a short, user-friendly name for the workflow, such as Disaster Recovery Workflow.
2. Define the workflow input. The inputs we take for this workflow include the following:
  - Volume options (volume name, mount path)
  - Volume capacity
  - Data center associated with the new datastore
  - Cluster on which the datastore is hosted
  - Name for the new datastore to create in vCenter
  - Type and version of the new datastore
  - Name of the Terraform organization
  - Terraform workspace
  - Description of the Terraform workspace
  - Variables (sensitive and nonsensitive) required to execute Terraform configuration
  - Reason for starting the plan
3. Add the workflow tasks.

The tasks related to operations in FlexPod include the following:

- Create volume in FlexPod.
- Add storage export policy to the created volume.
- Map the newly created volume to a datastore in VMware vCenter.

The tasks related to creating Cloud Volumes ONTAP cluster:

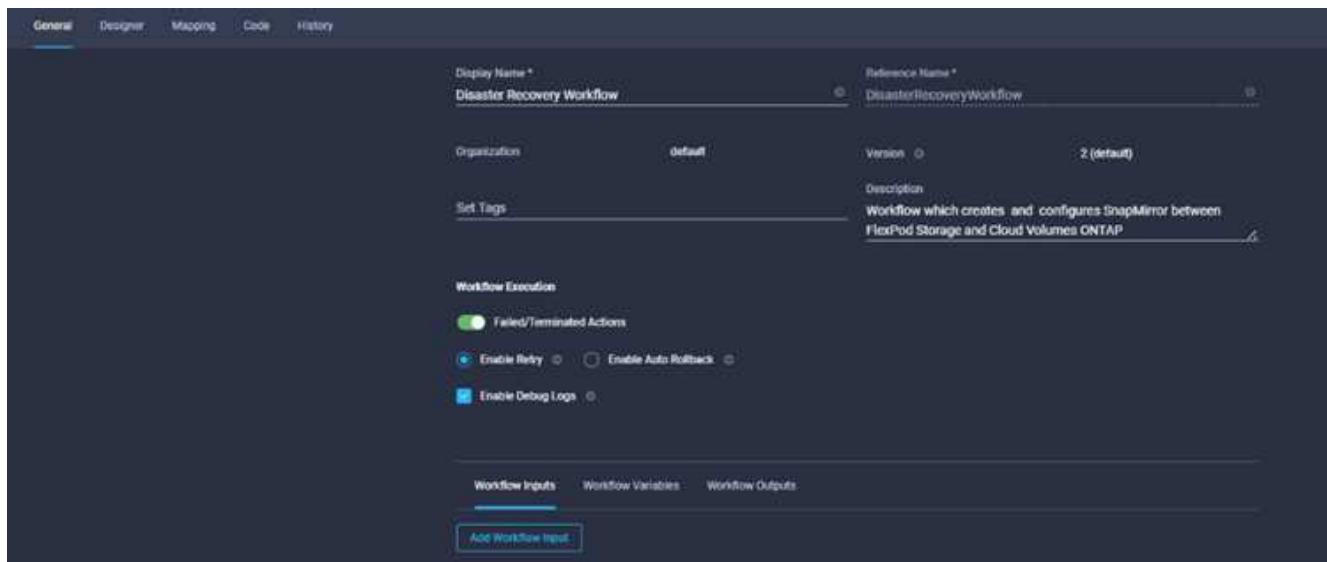
- Add Terraform workspace
- Add Terraform variables
- Add Terraform sensitive variables
- Start new Terraform plan
- Confirm Terraform run

4. Validate the workflow.

### Procedure 1: Create the workflow

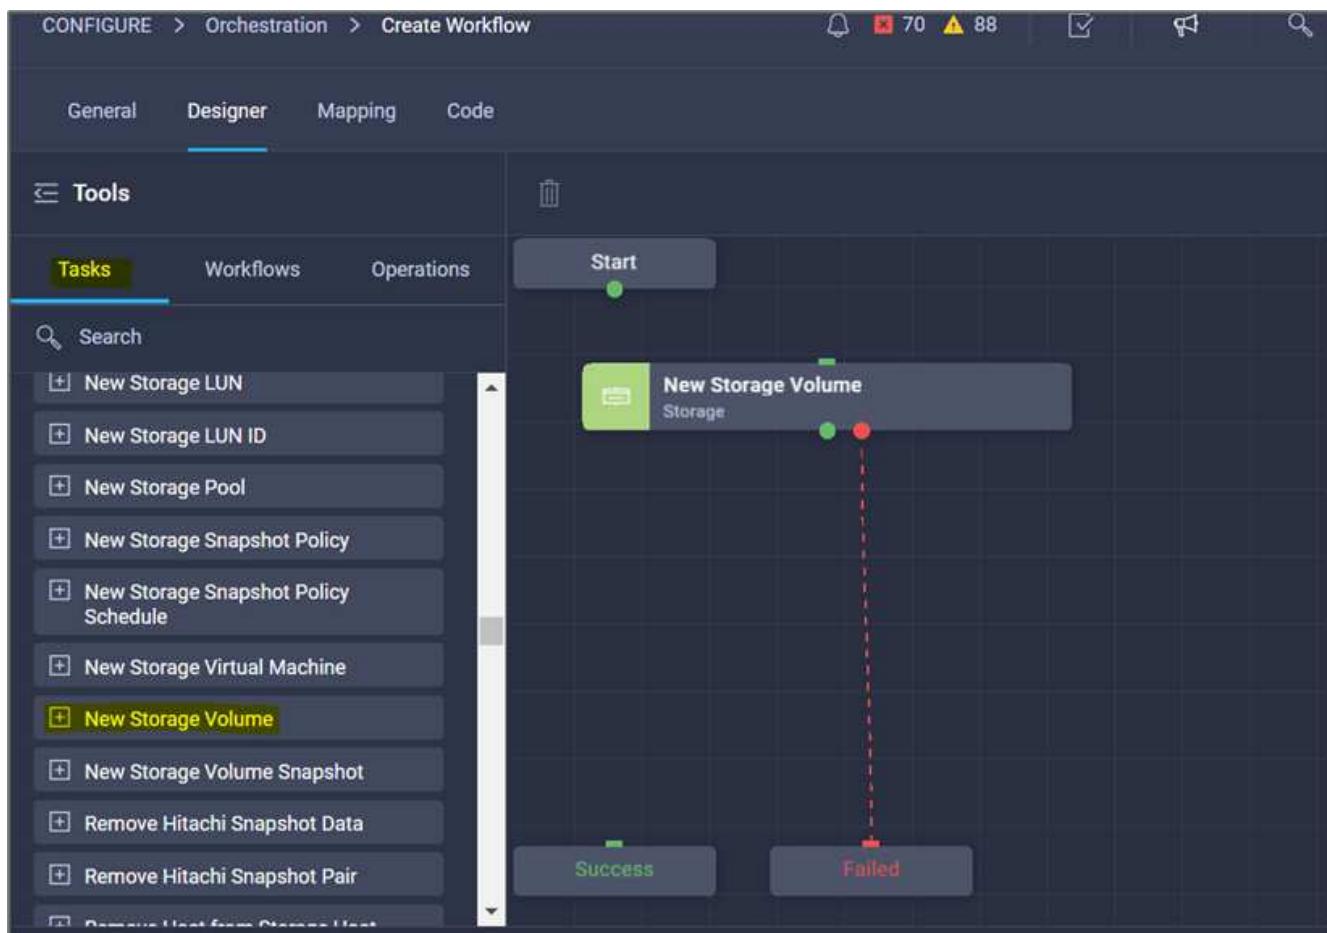
1. Click **Orchestration** from the left navigation pane and click **Create Workflow**.
2. In the **General** tab:
  - a. Provide the display name (Disaster Recovery Workflow).
  - b. Select the organization, set tags, and provide a description.

3. Click Save.

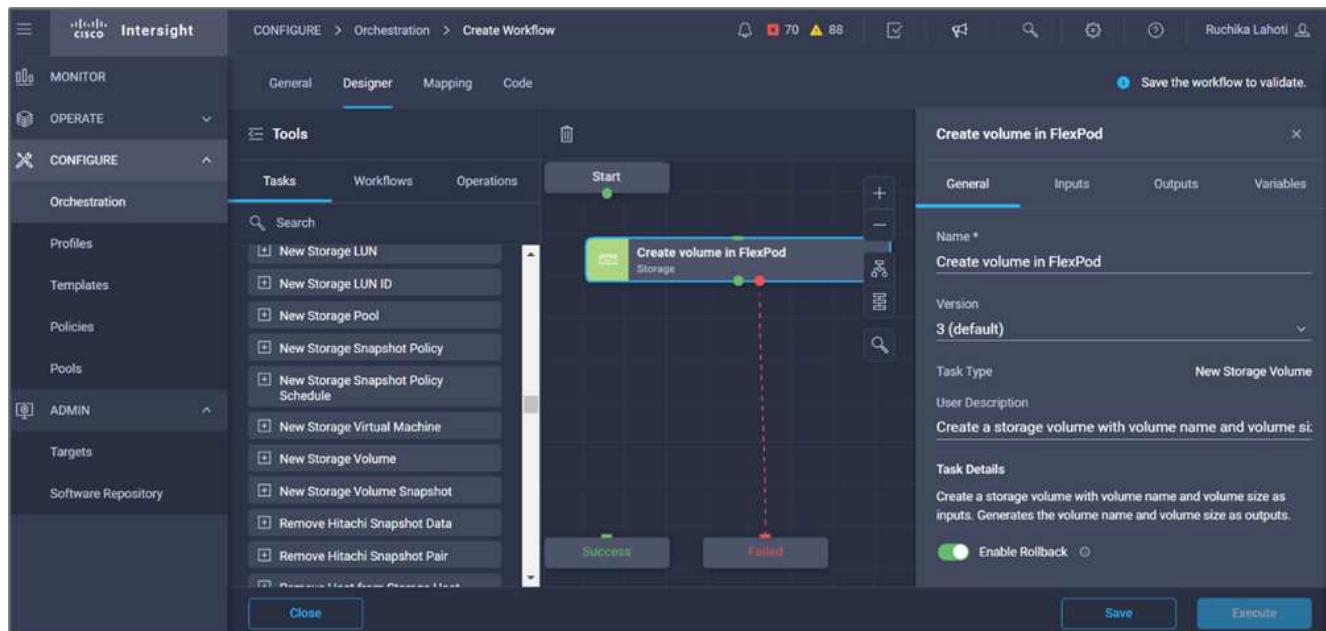


#### Procedure 2. Create a new volume in FlexPod

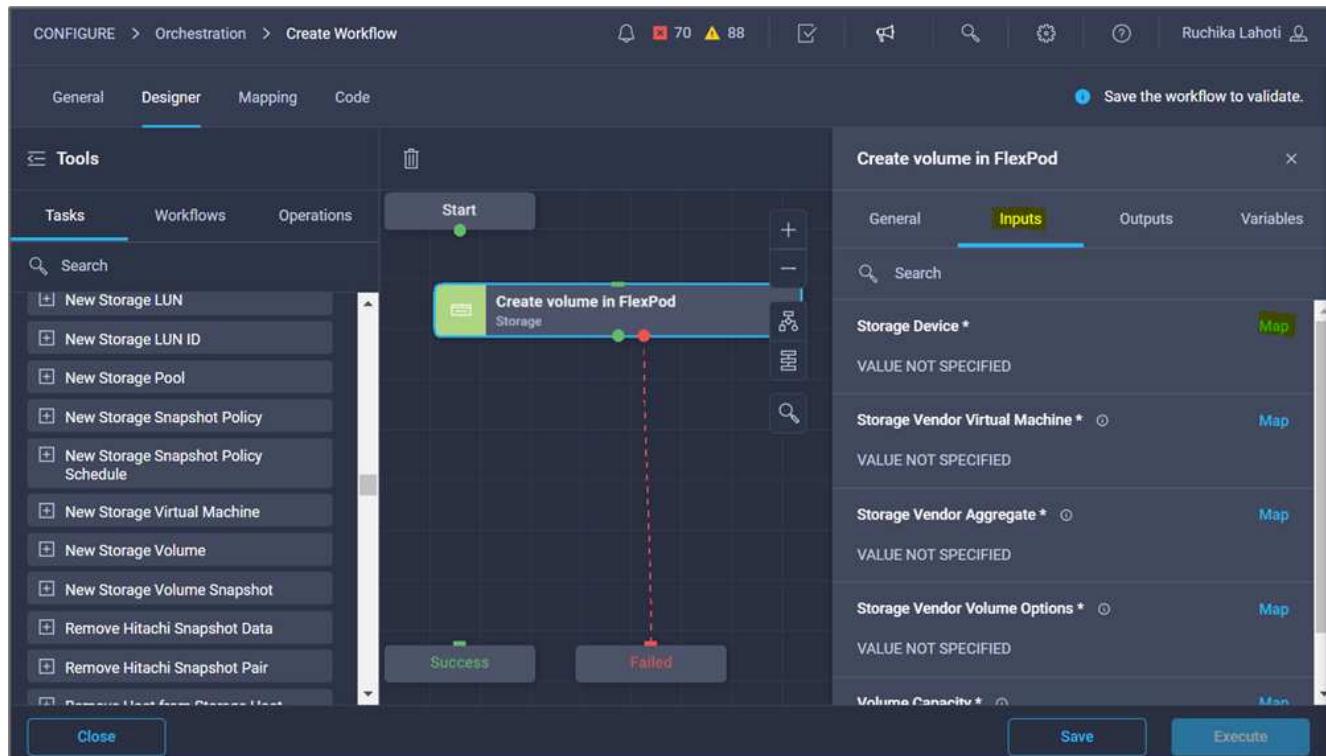
1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Storage > New Storage Volume** task from the **Tools** section into the **Design** area.
3. Click **New Storage Volume**.



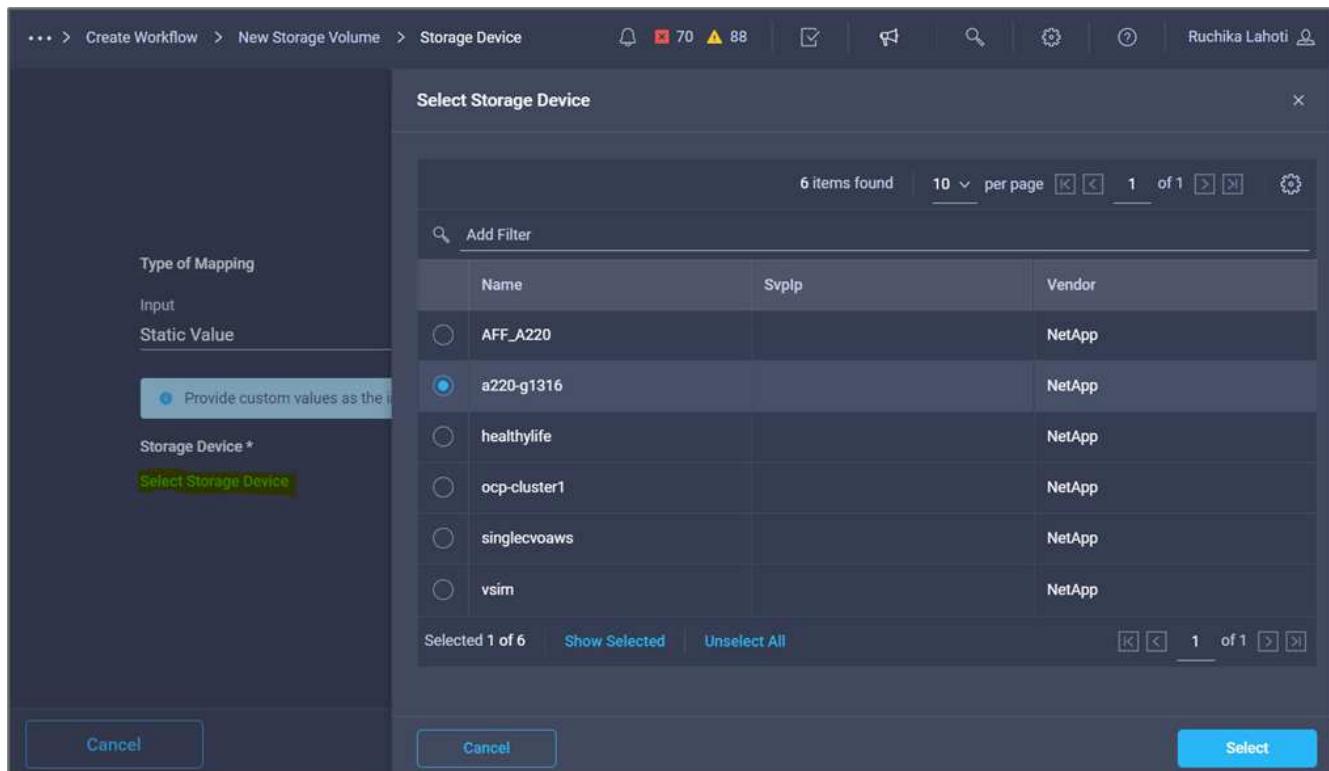
- In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task. In this example, the name of the task is **Create Volume in FlexPod**.



- In the **Task Properties** area, click **Inputs**.
- Click **Map** in the **Storage Device** field.



- Choose **Static Value** and click **Select Storage Device**.
- Click the storage target added and click **Select**.



... > Create Workflow > New Storage Volume > Storage Device

Select Storage Device

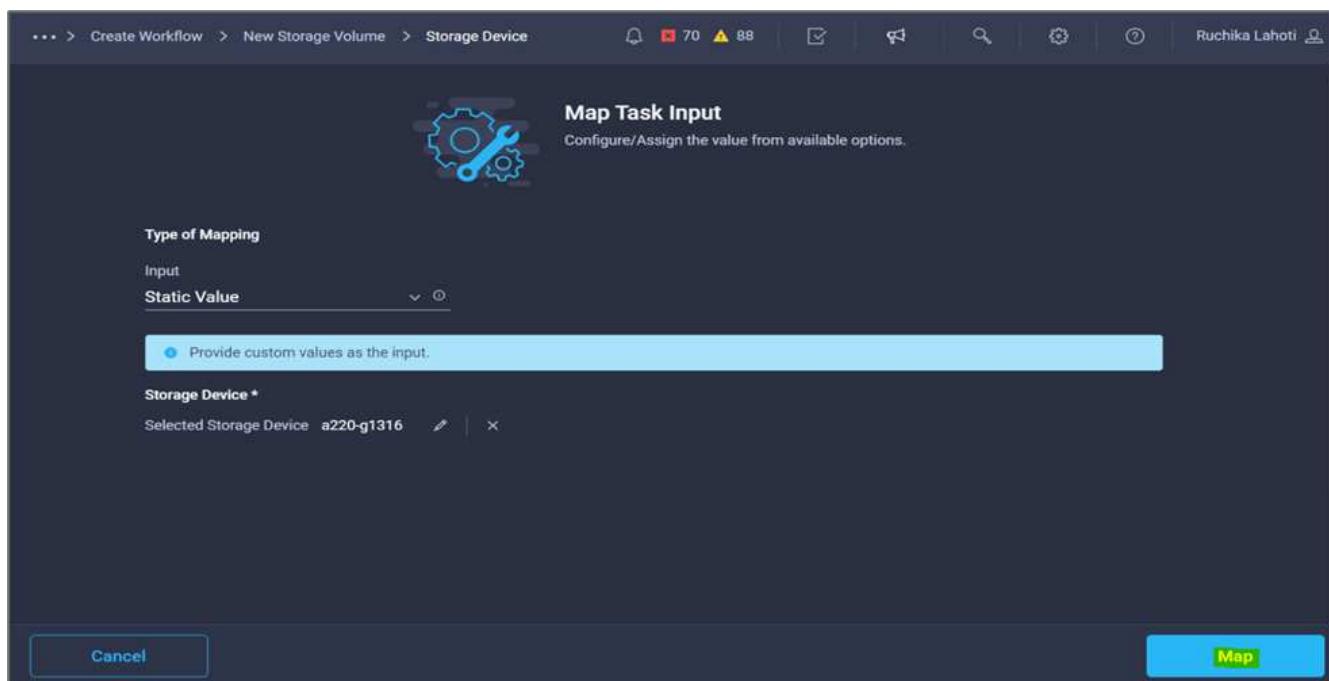
6 items found 10 per page 1 of 1

	Name	Svlp	Vendor
<input type="radio"/>	AFF_A220		NetApp
<input checked="" type="radio"/>	a220-g1316		NetApp
<input type="radio"/>	healthylife		NetApp
<input type="radio"/>	ocp-cluster1		NetApp
<input type="radio"/>	singlecvoaws		NetApp
<input type="radio"/>	vsim		NetApp

Selected 1 of 6 Show Selected Unselect All 1 of 1

Cancel Select

9. Click **Map**.



... > Create Workflow > New Storage Volume > Storage Device

Map Task Input

Configure/Assign the value from available options.

Type of Mapping

Input

Static Value

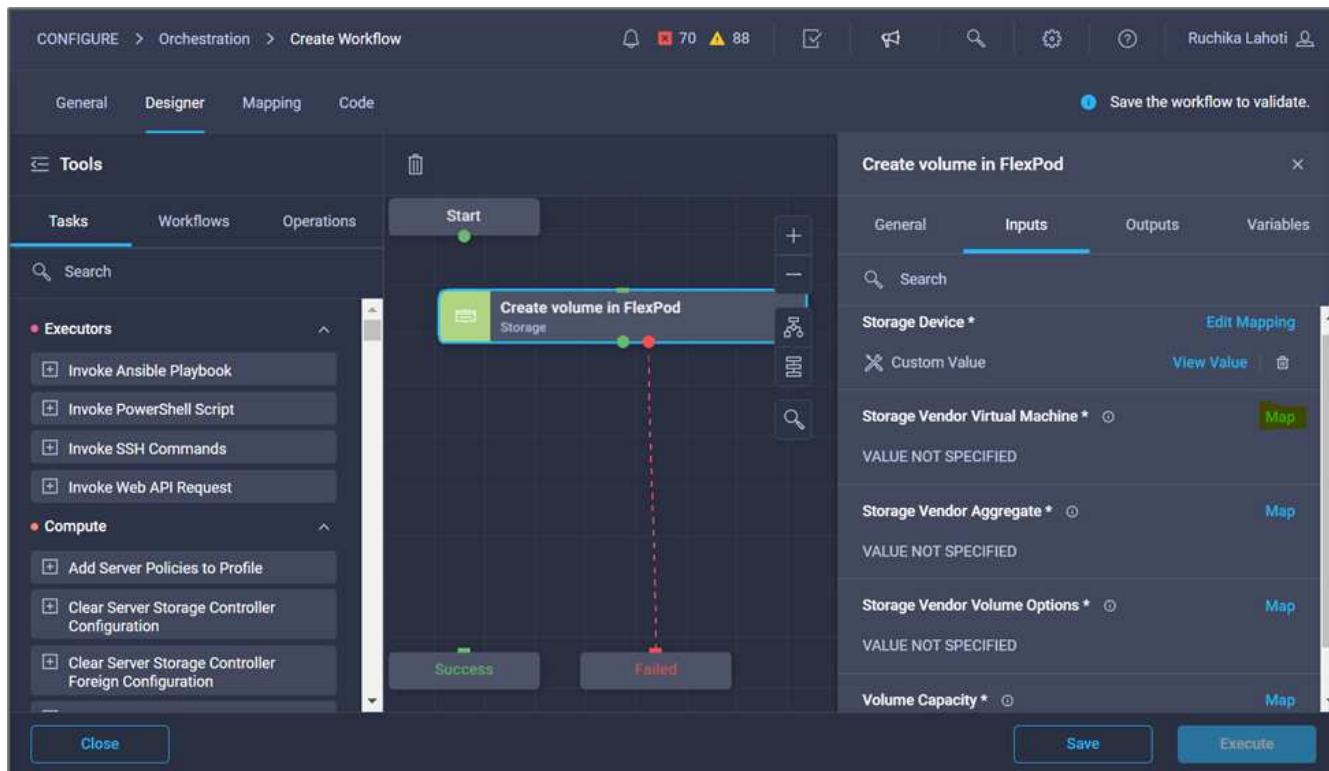
Provide custom values as the input.

Storage Device \*

Selected Storage Device a220-g1316

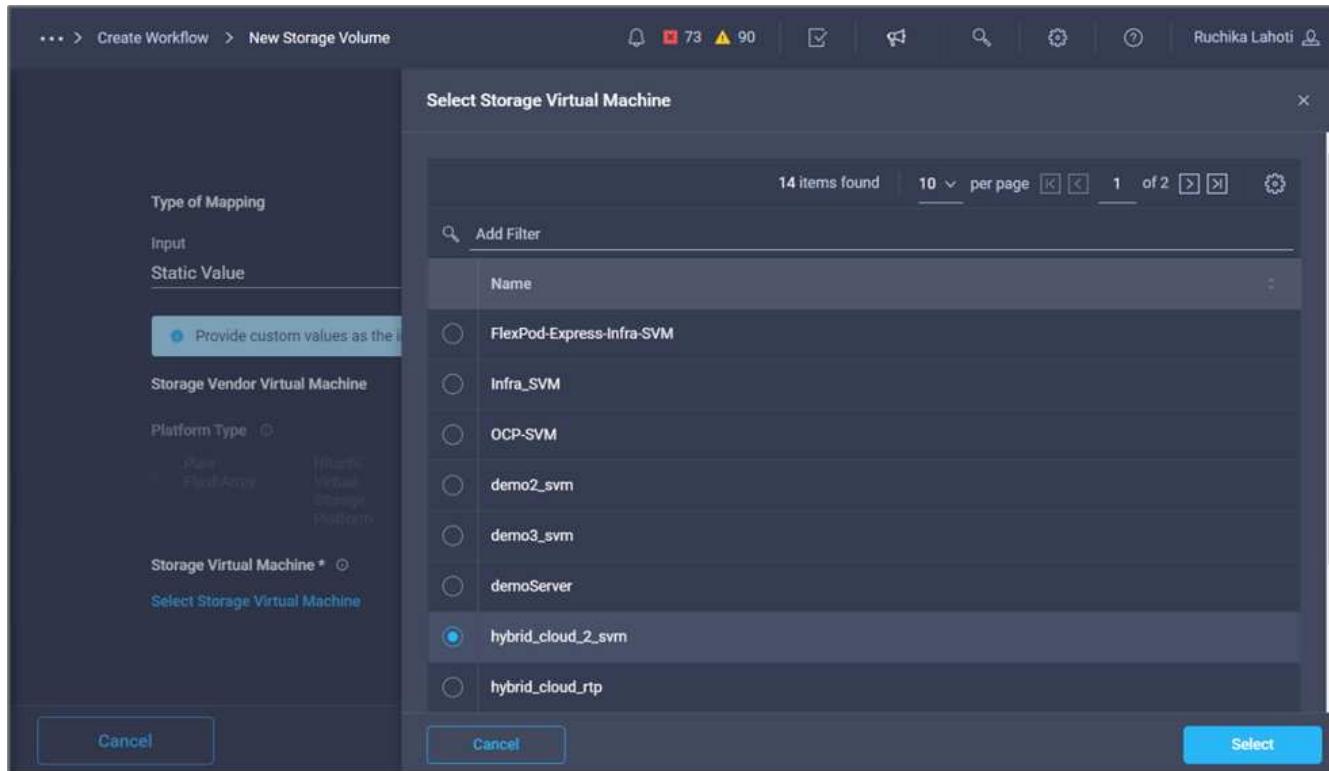
Cancel Map

10. Click **Map** in the **Storage Vendor Virtual Machine** field.

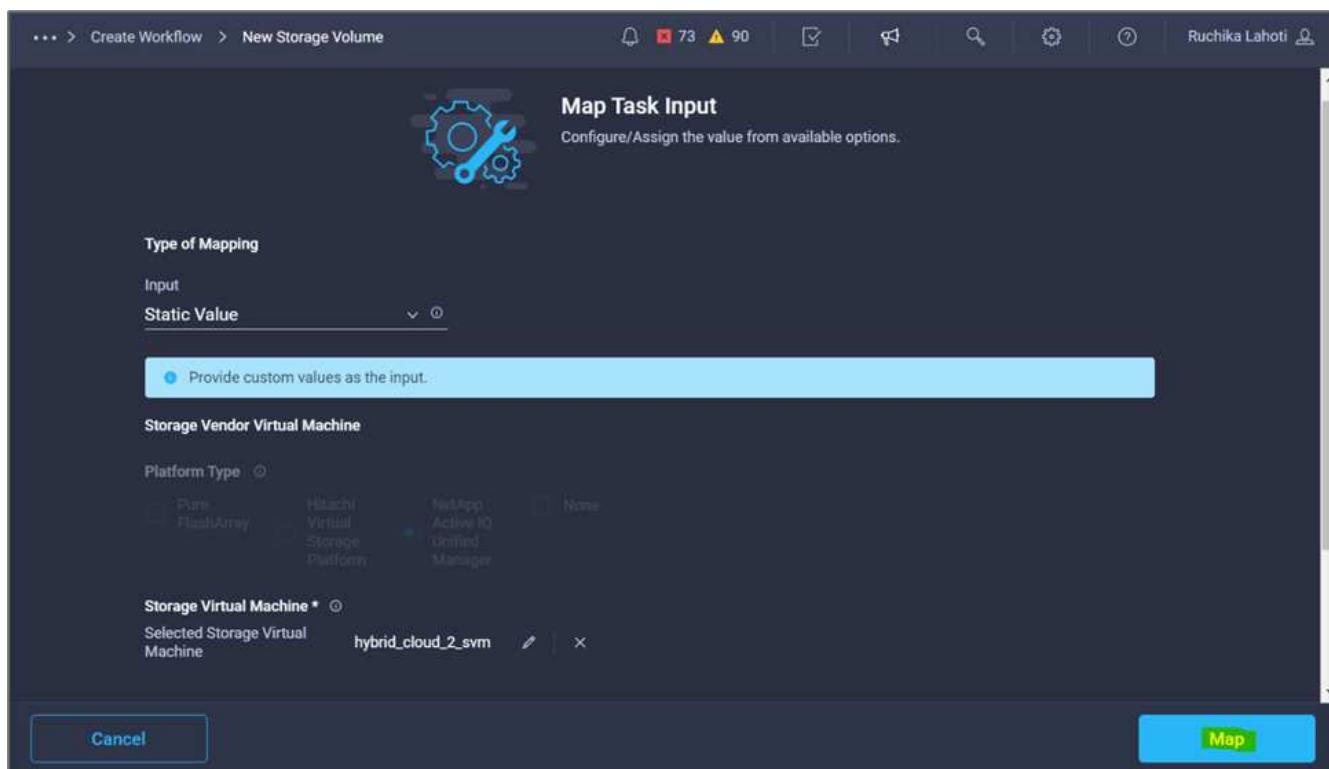


11. Choose **Static Value** and click **Select Storage Virtual Machine**.

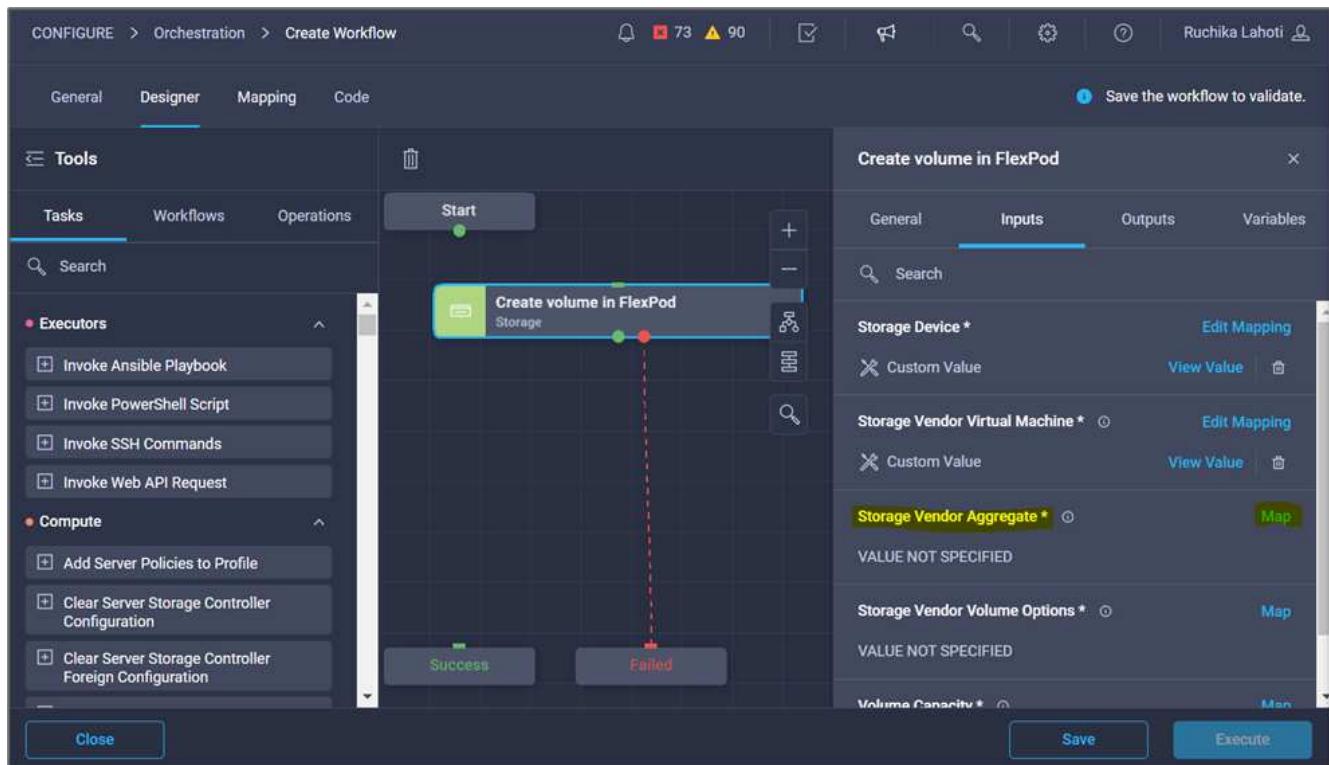
12. Select the storage virtual machine where the volume needs to be created and click **Select**.



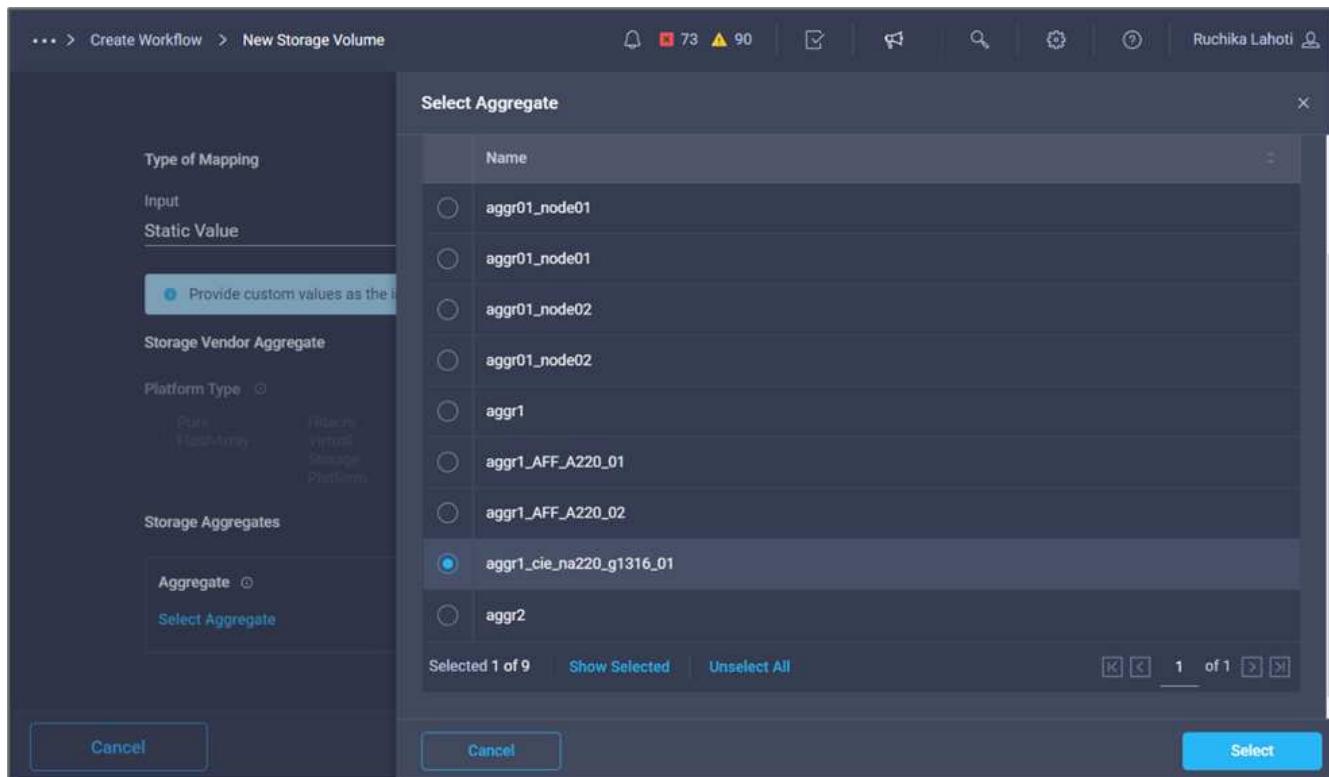
13. Click **Map**.



14. Click **Map** in the **Storage Vendor Aggregate** field.



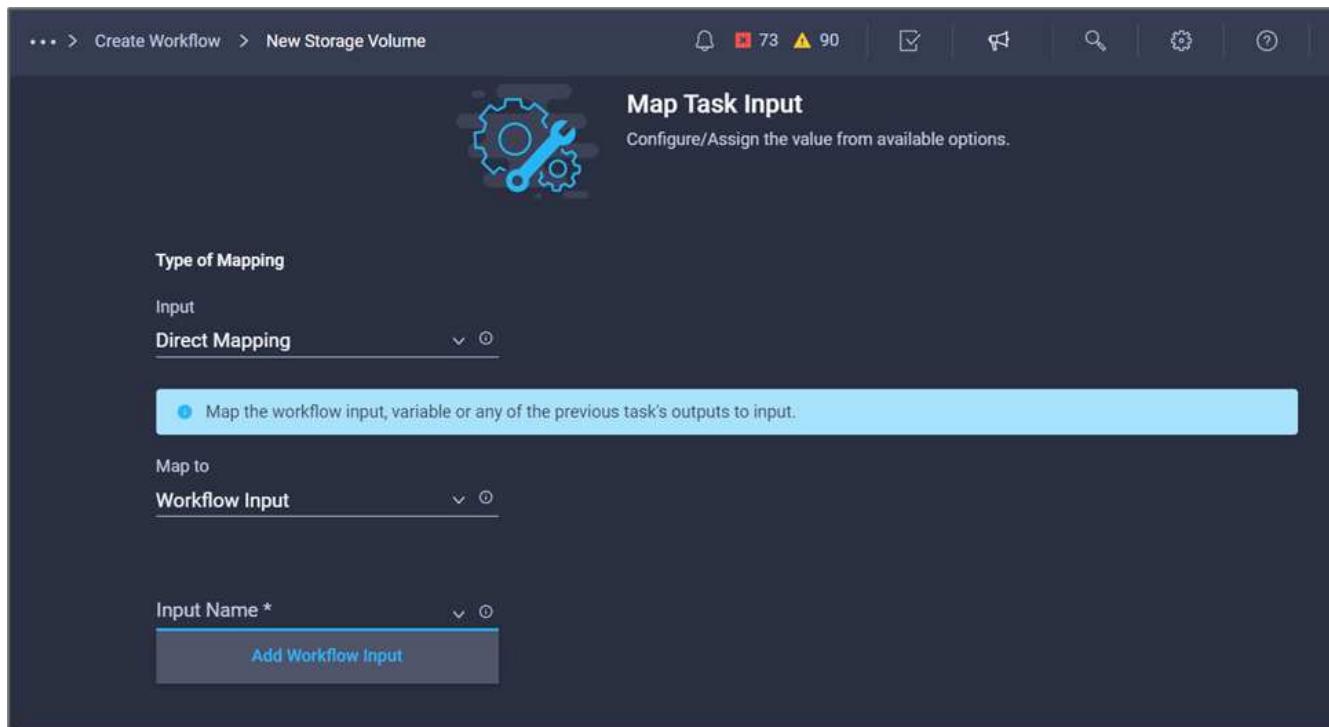
15. Choose **Static Value** and click **Select Storage Aggregate**. Choose the aggregate and click **Select**.



16. Click **Map**.

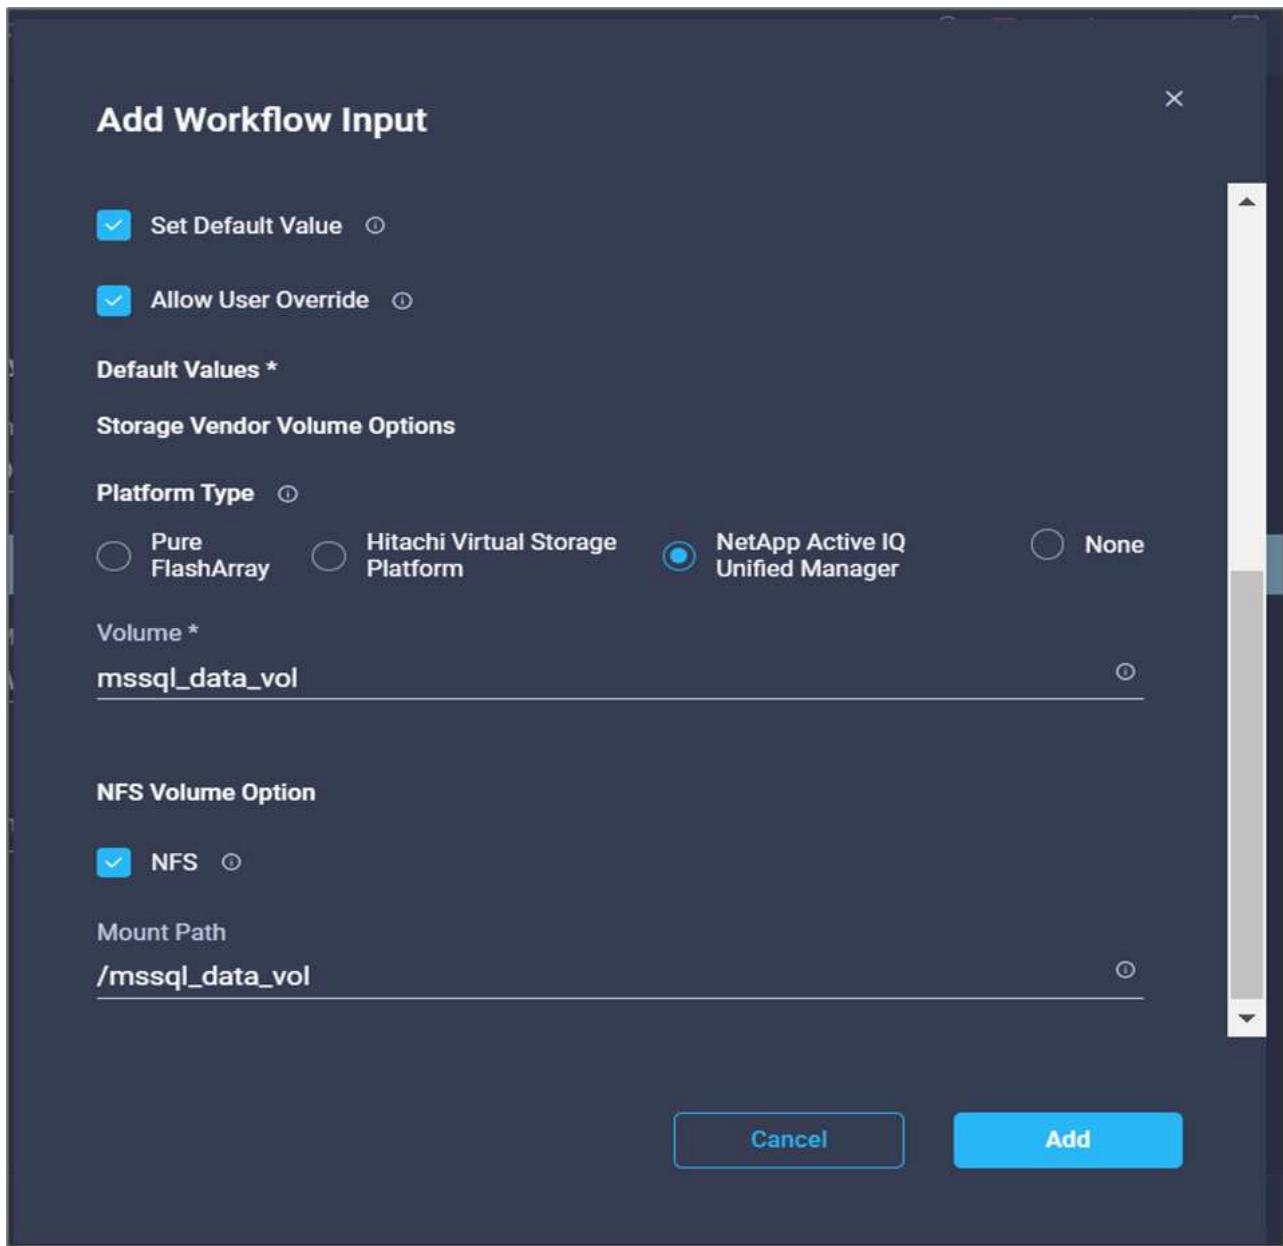
17. Click **Map** in the **Storage Vendor Volume Options** field.

18. Choose **Direct Mapping** and click **Workflow Input**.

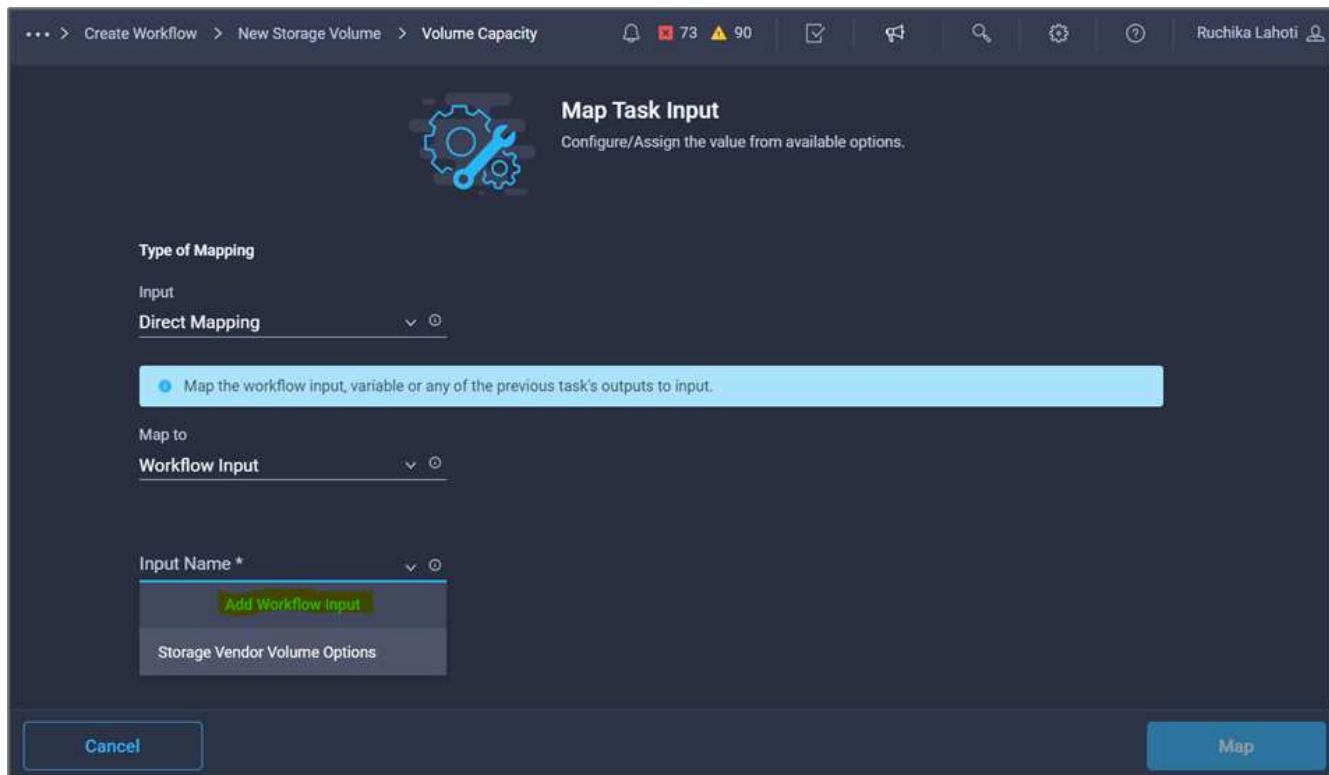


19. In the Add Input wizard, complete the following steps:

- Provide a display name and reference name (optional).
- Make sure that **Storage Vendor Volume Options** is selected for the **Type**.
- Click **Set Default Value and Override**.
- Click **Required**.
- Set the **Platform Type** to **NetApp Active IQ Unified Manager**.
- Provide a default value for the created volume under **Volume**.
- Click **NFS**. If NFS is set, an NFS volume is created. If this value is set to false, a SAN volume is created.
- Provide a mount path and click **Add**.

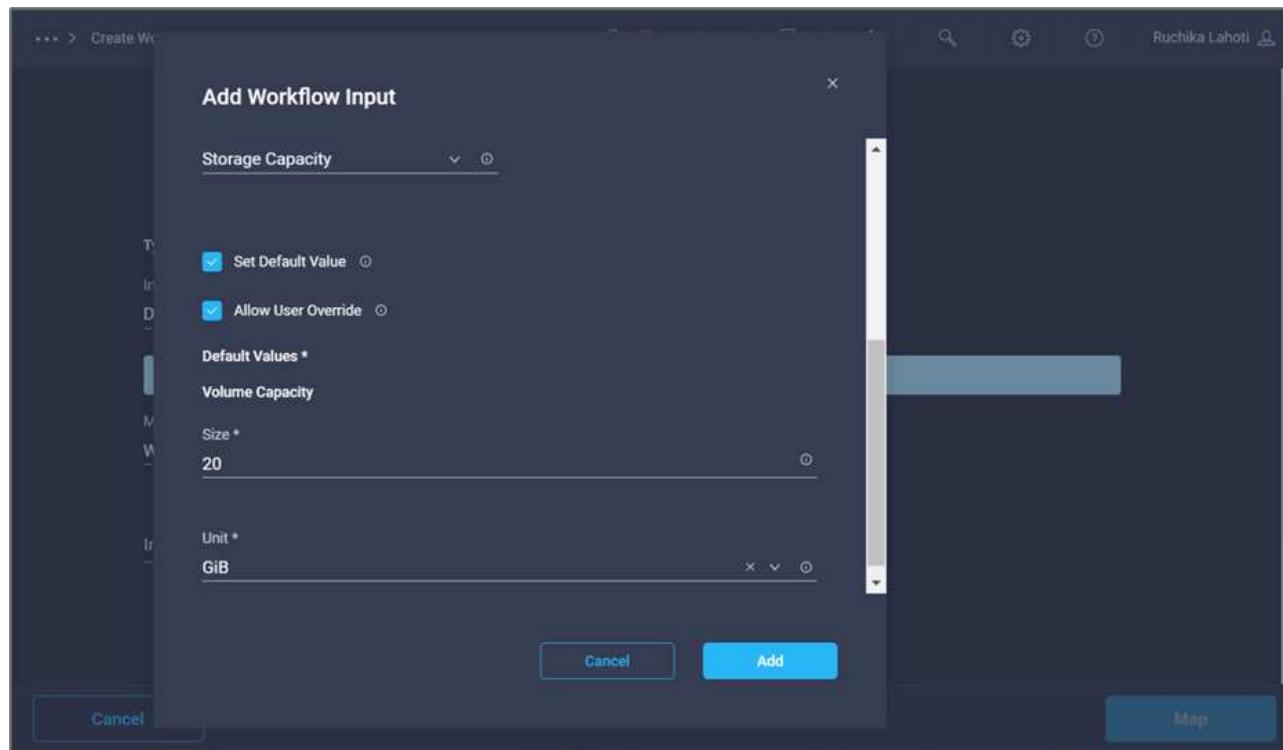


20. Click **Map**.
21. Click **Map** in the **Volume Capacity** field.
22. Choose **Direct Mapping** and click **Workflow Input**.
23. Click **Input Name** and **Create Workflow Input**.



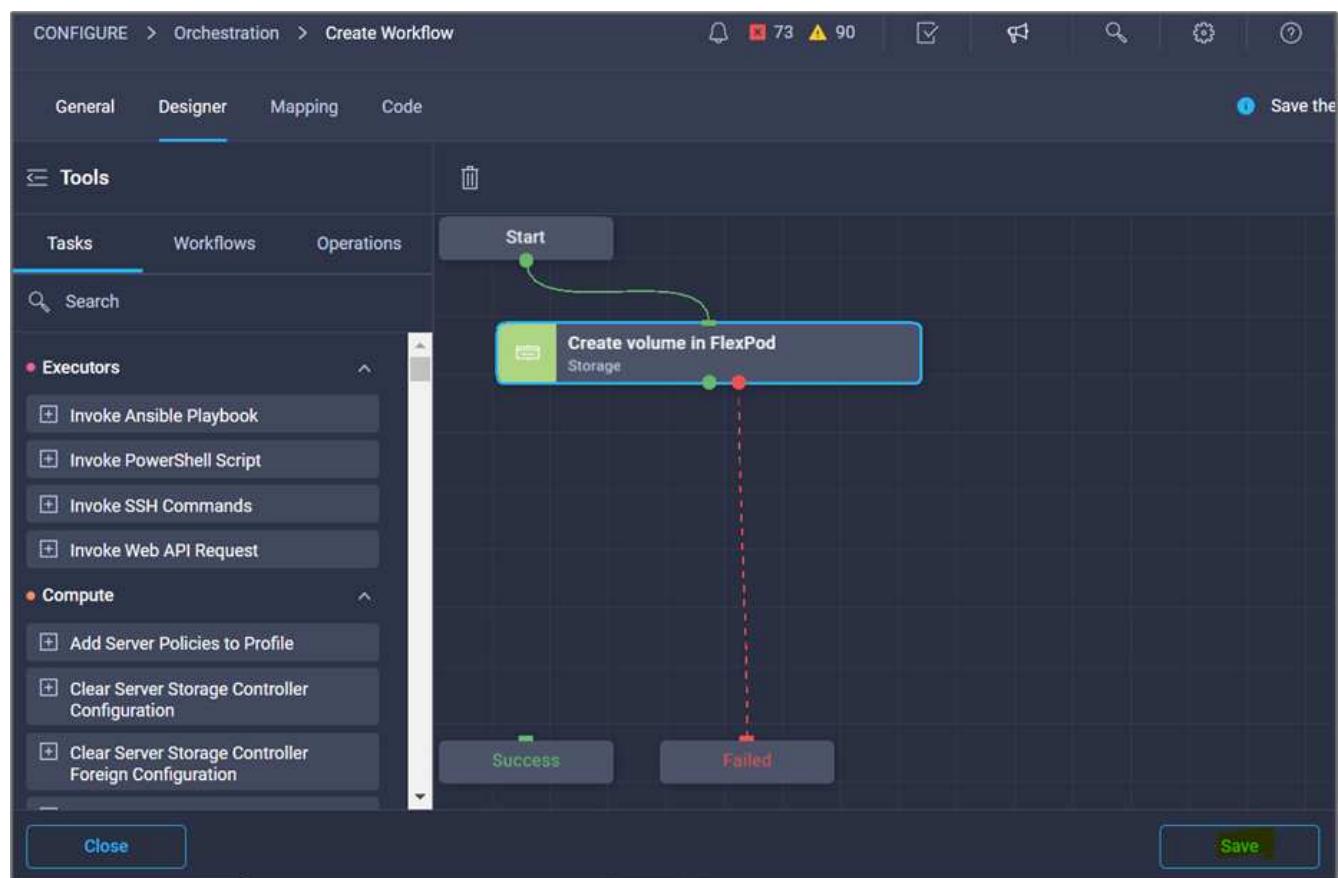
24. In the Add Input wizard:

- a. Provide a display name and a reference name (optional).
- b. Click **Required**.
- c. For **Type**, select **Storage Capacity**.
- d. Click **Set Default Value and Override**.
- e. Provide a default value for the volume size and unit.
- f. Click **Add**.



25. Click **Map**.

26. With Connector, create a connection between the **Start** and **Create Volume in FlexPod** tasks, and click **Save**.

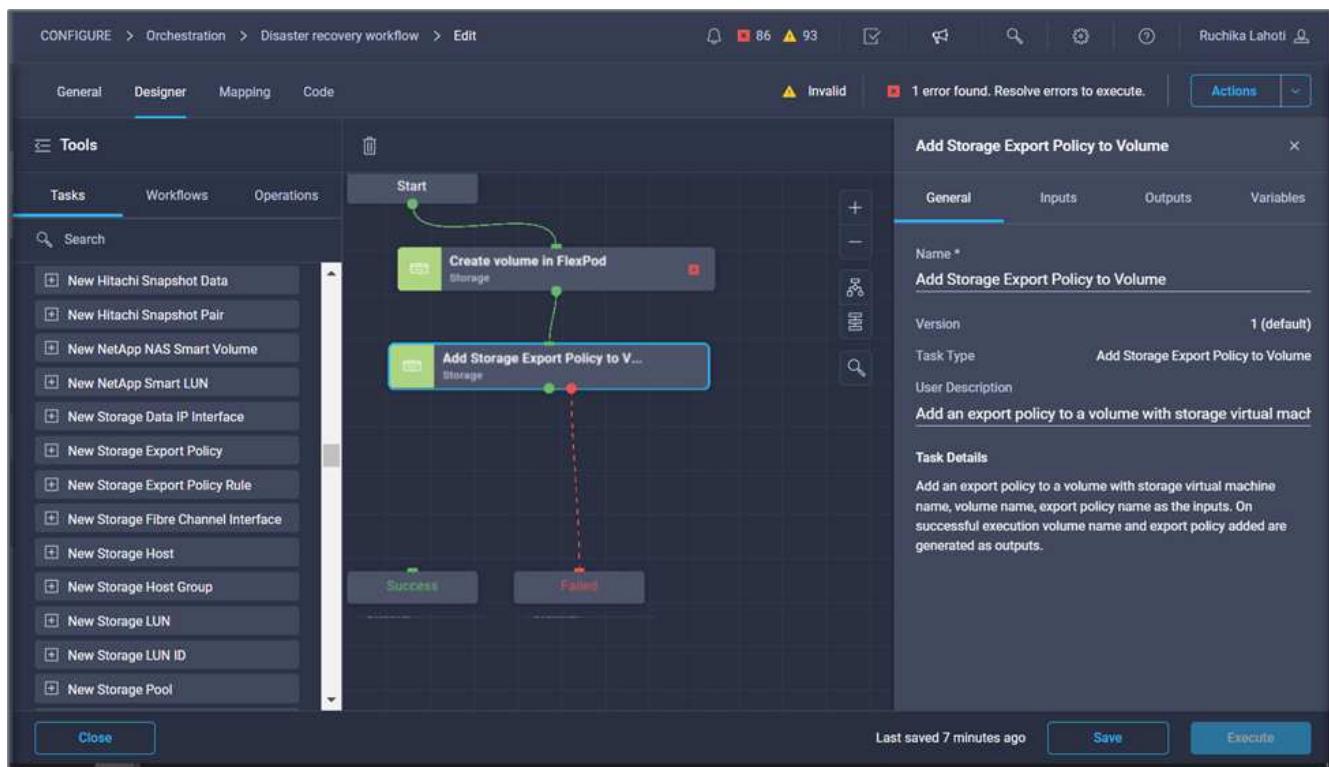




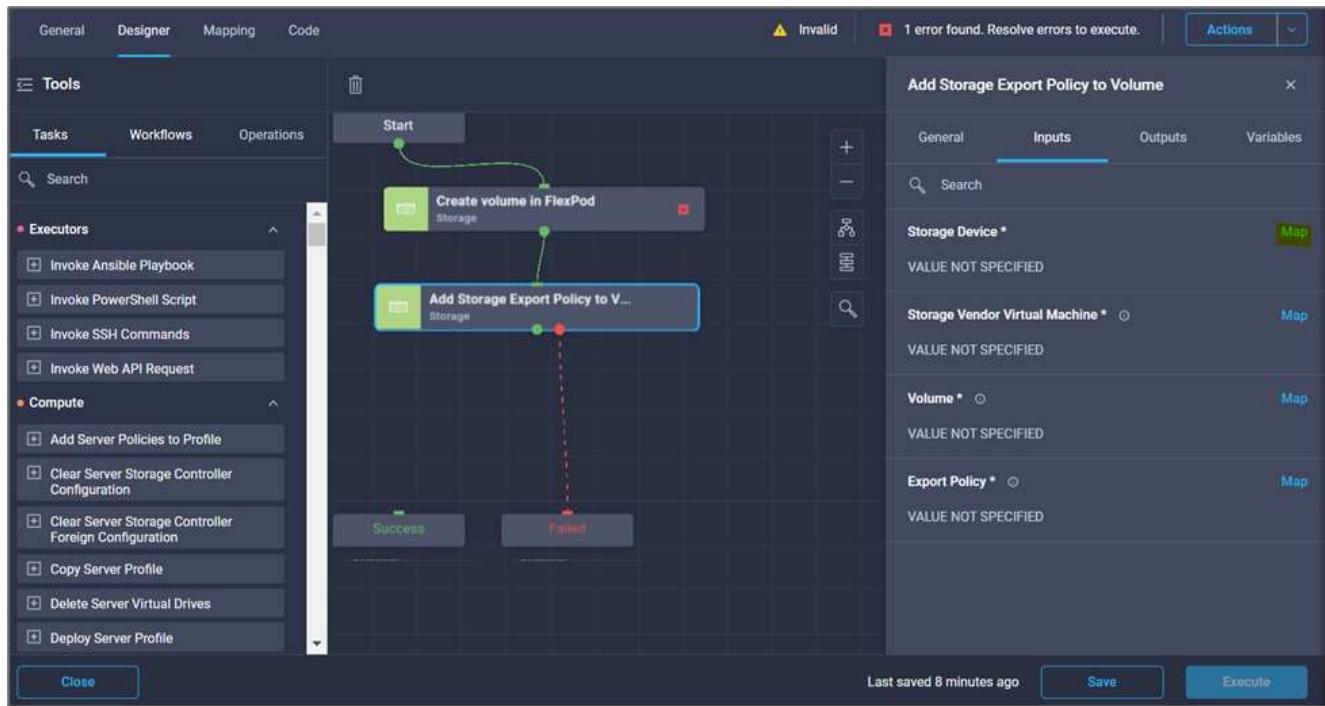
Ignore the error for now. This error displays because there is no connectivity between the tasks **Create Volume in FlexPod** and **Success** which is required to specify the successful transition.

#### Procedure 3: Add storage export policy

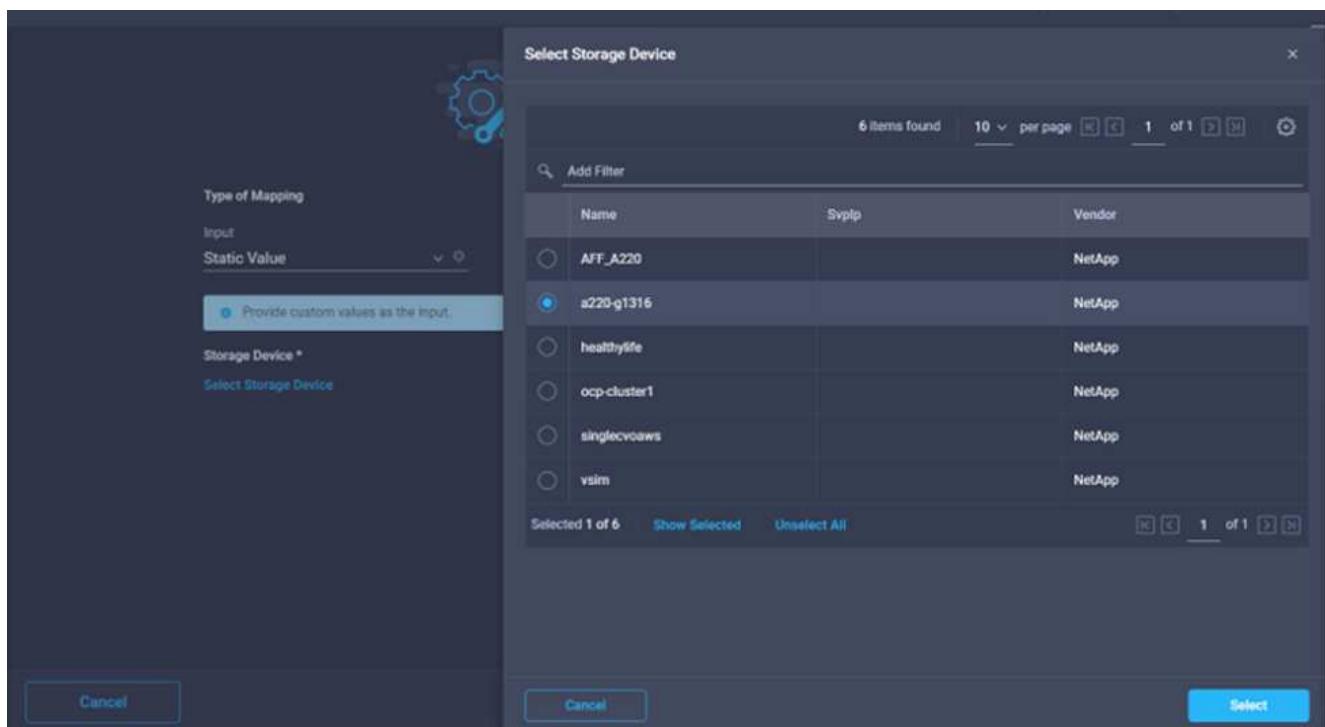
1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Storage > Add Storage Export Policy to Volume** task from the **Tools** section in the **Design** area.
3. Click **Add Storage Export Policy to Volume**. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task. In this example, the name of the task is **Add Storage Export Policy**.
4. Use Connector to make a connection between the tasks **Create Volume in FlexPod** and **Add Storage Export Policy**. Click **Save**.



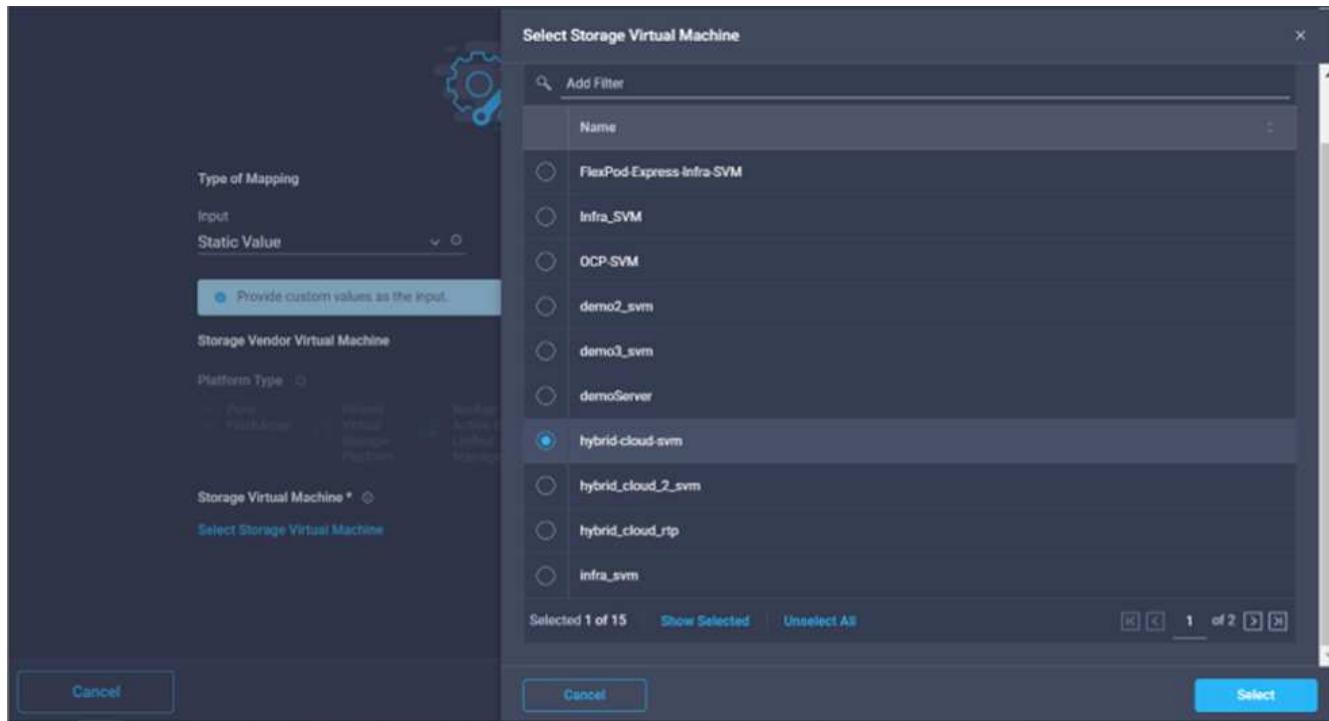
5. In the **Task Properties** area, click **Inputs**.
6. Click **Map** in the **Storage Device** field.



7. Choose **Static Value** and click **Select Storage Device**. Select the same storage target added while creating the previous task of creating a new storage volume.
8. Click **Map**.



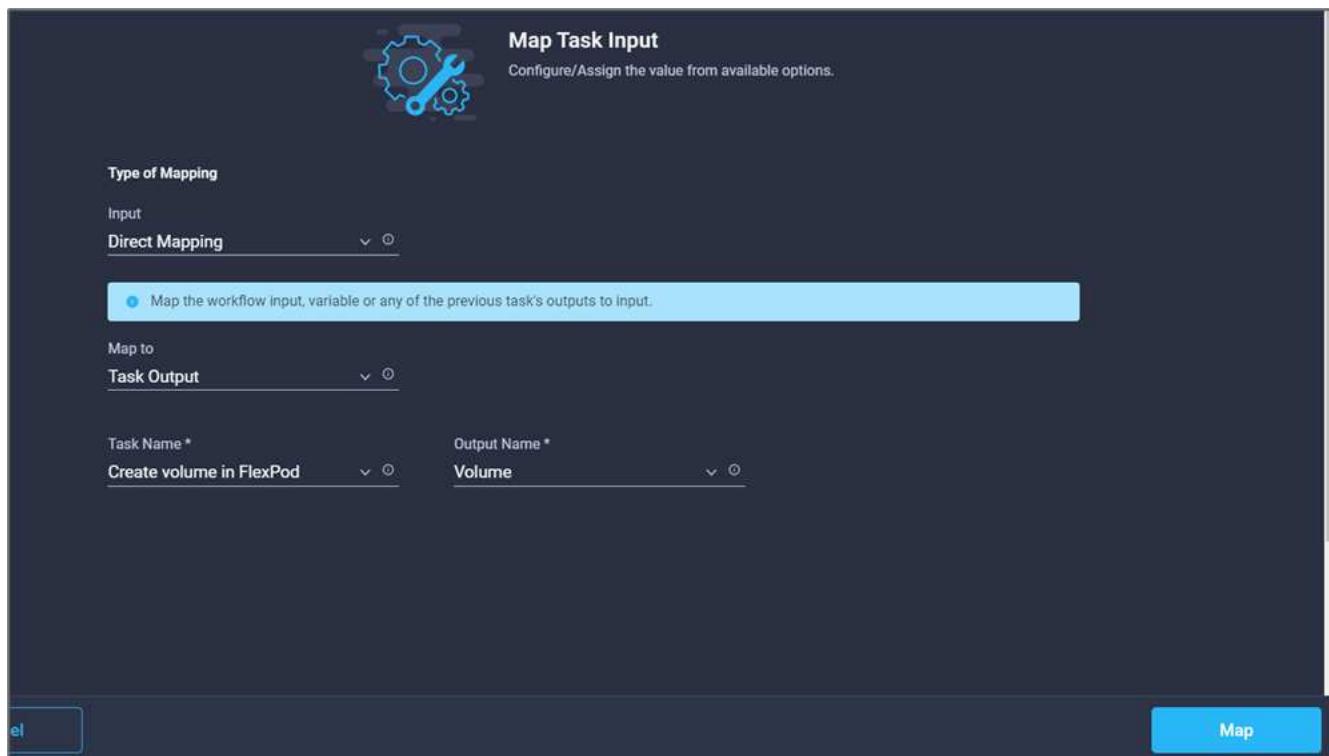
9. Click **Map** in the **Storage Vendor Virtual Machine** field.
10. Choose **Static Value** and click **Select Storage Virtual Machine**. Select the same storage virtual machine added while creating the previous task of creating a new storage volume.



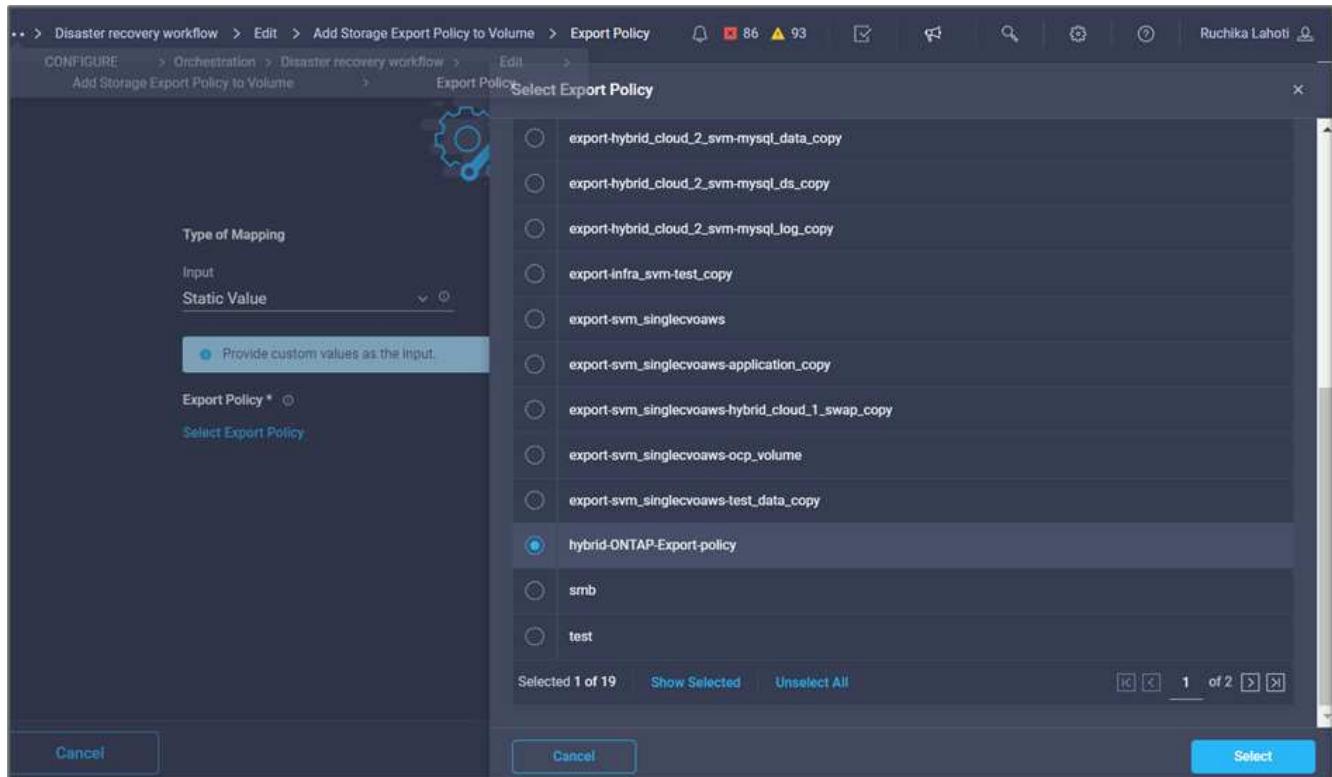
11. Click **Map**.
12. Click **Map** in the **Volume** field.
13. Click **Task Name** and then click **Create Volume in FlexPod**. Click **Output Name** and then **Volume**.



In Cisco Intersight Cloud Orchestrator, you can provide the output of a previous task as the input for a new task. In this example, the **Volume** details were provided from the **Create Volume in FlexPod** task as an input for the task **Add Storage Export Policy**.



14. Click **Map**.
15. Click **Map** in the **Export Policy** field.
16. Choose **Static Value** and click **Select Export Policy**. Select the export policy created.



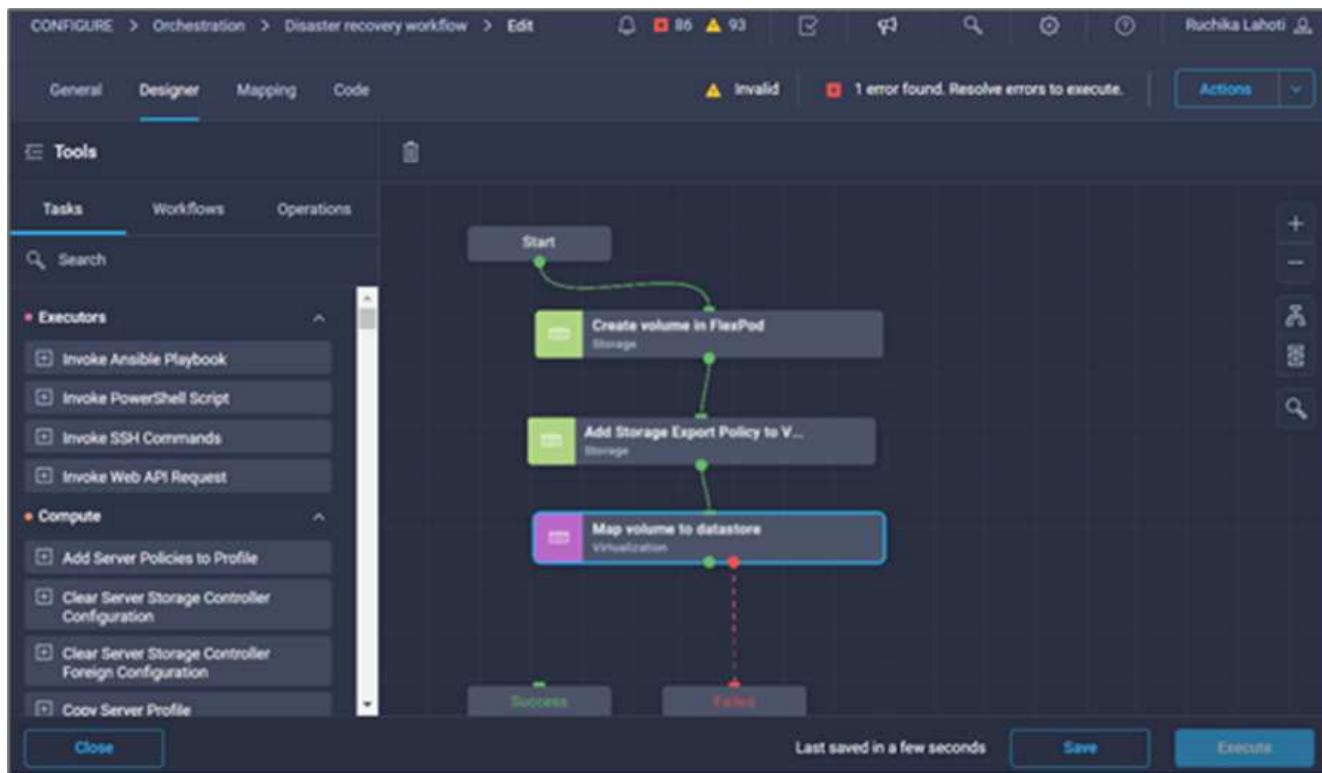
17. Click **Map** and then **Save**.



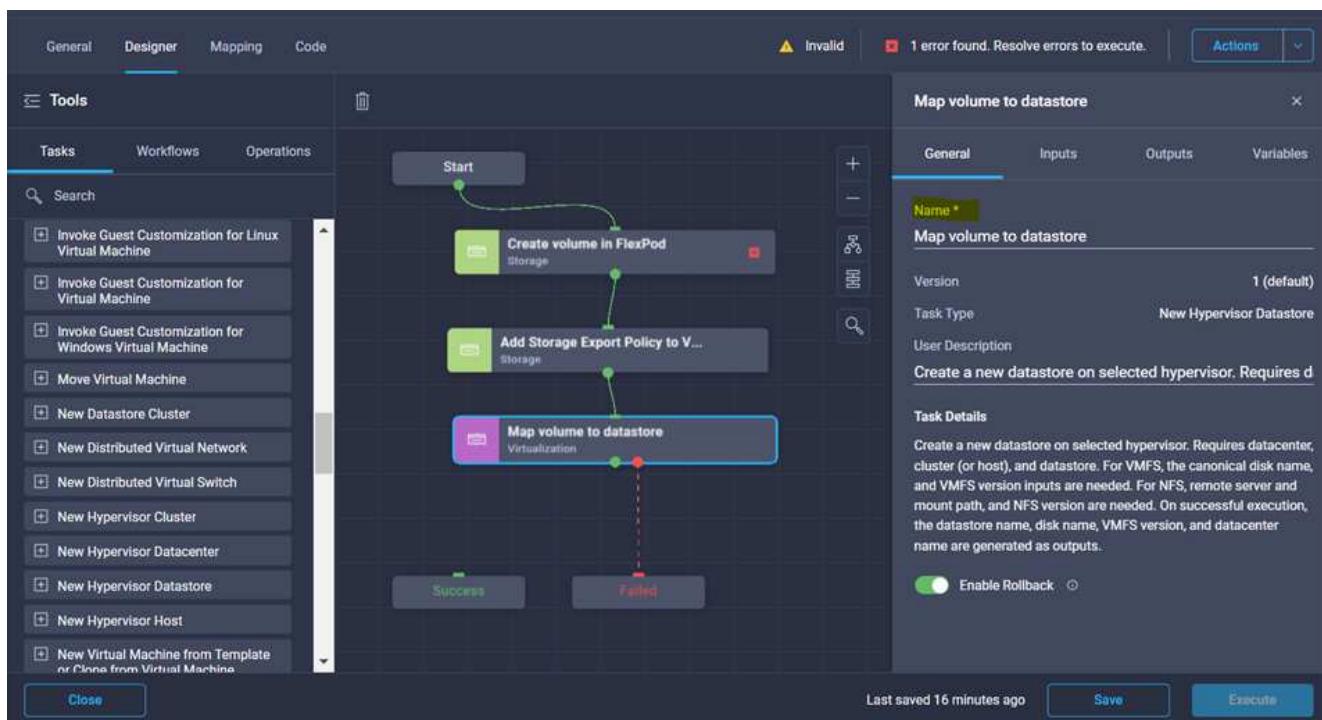
This completes addition of an export policy to the volume. Next, you create a new datastore mapping the created volume.

#### Procedure 4: Map FlexPod volume to datastore

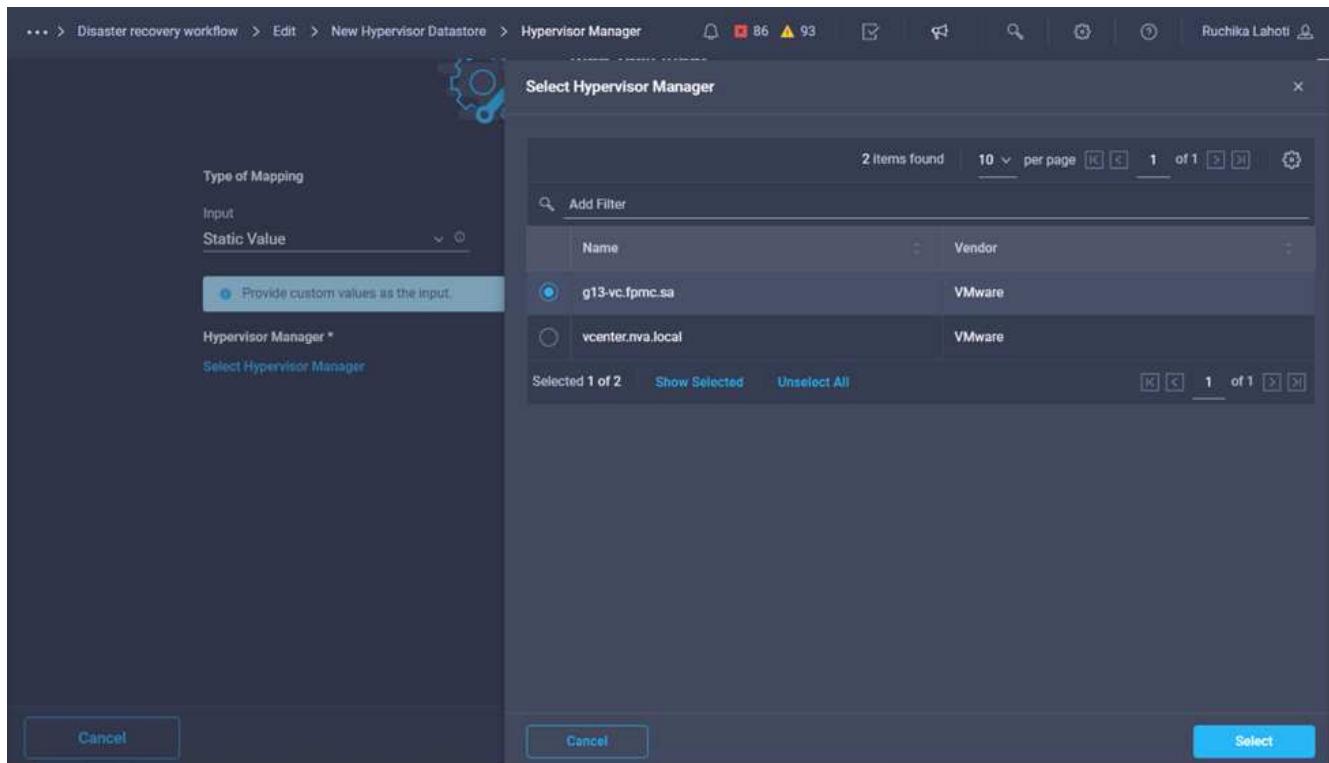
1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Virtualization > New Hypervisor Datastore** task from the **Tools** section in the **Design** area.
3. Use Connector to make a connection between the **Add Storage Export Policy** and **New Hypervisor Datastore** tasks. Click **Save**.



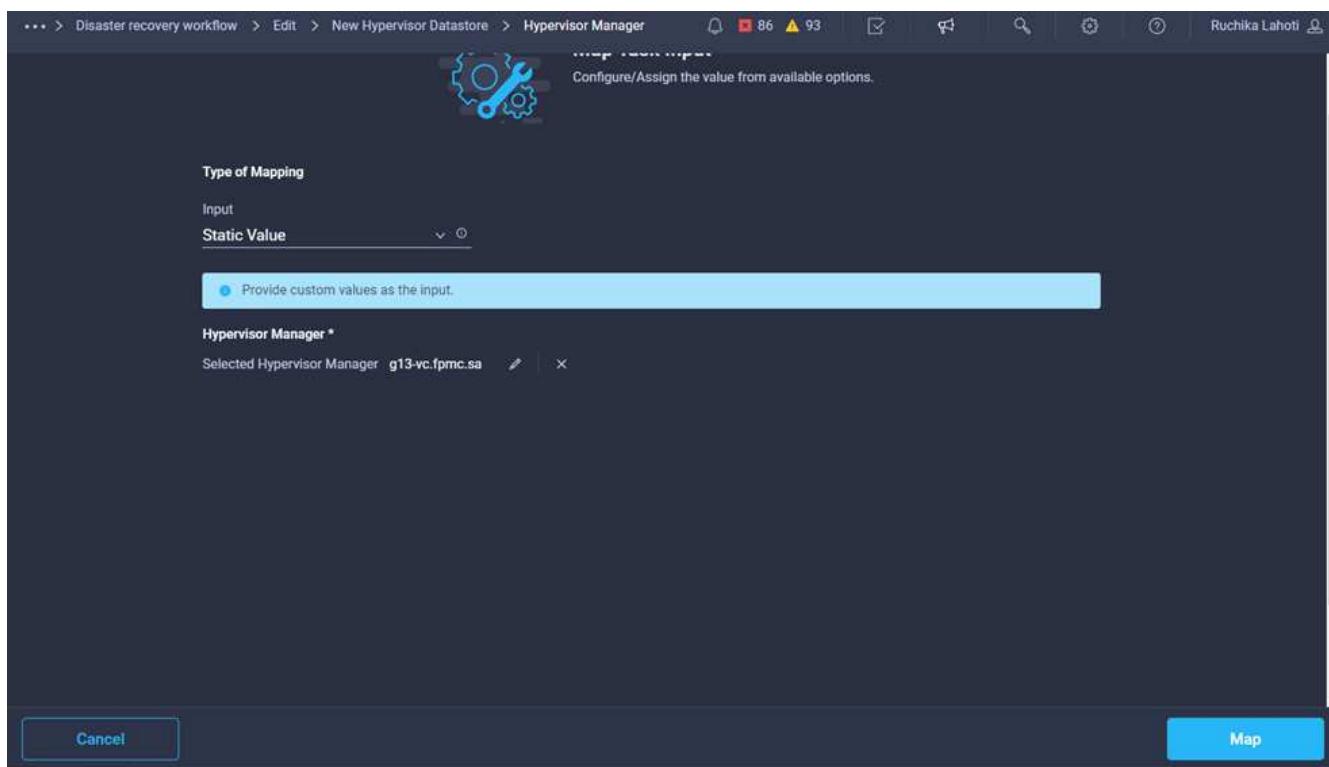
- Click **New Hypervisor Datastore**. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task. In this example, the name of the task is **Map volume to Datastore**.



- In the **Task Properties** area, click **Inputs**.
- Click **Map** in the **Hypervisor Manager** field.
- Choose **Static Value** and click **Select Hypervisor Manager**. Click the VMware vCenter target.



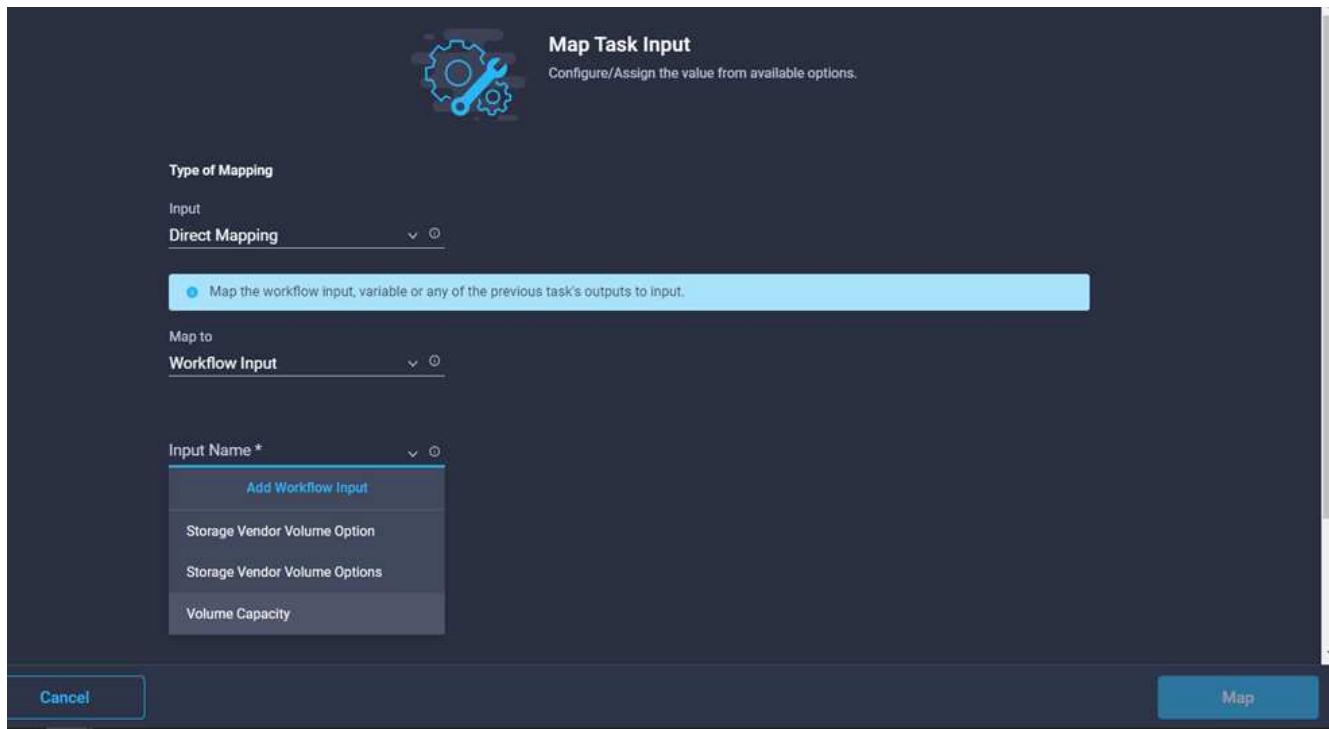
8. Click **Map**.



9. Click **Map** in the **Data center** field. This is the data center associated with the new datastore.

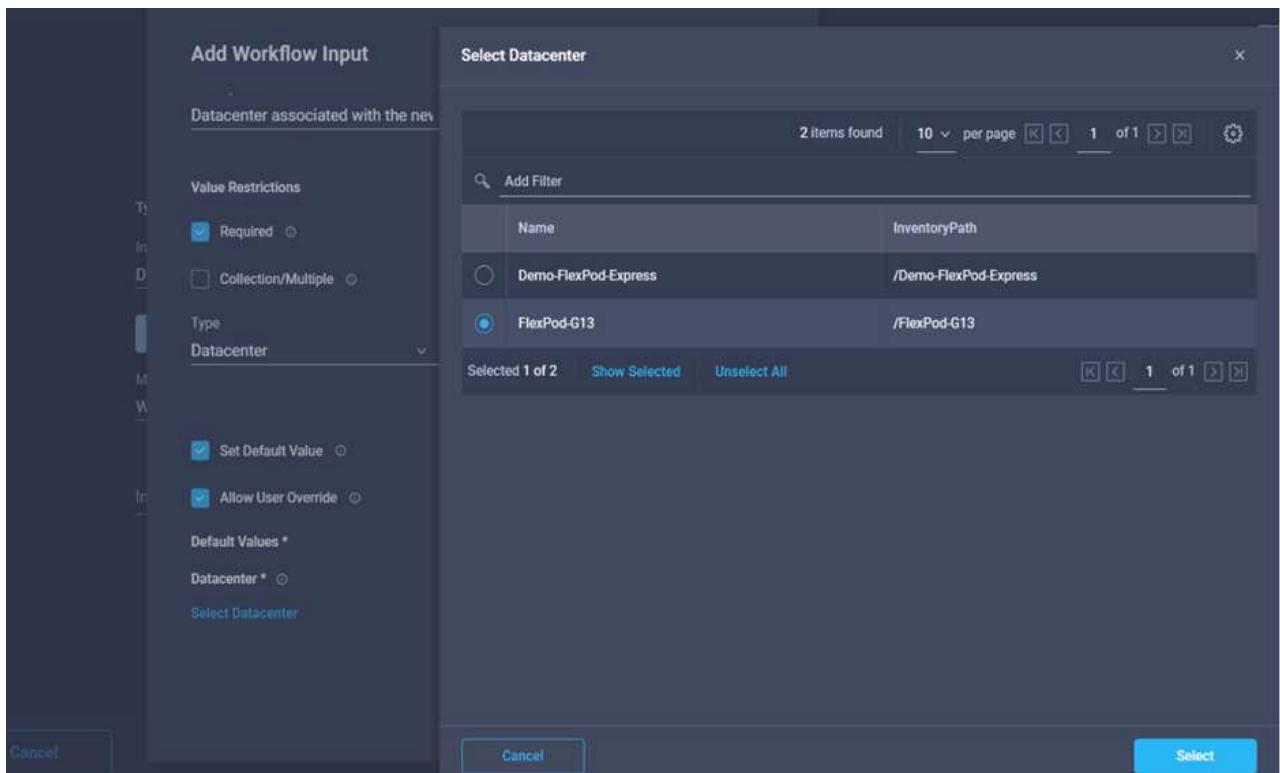
10. Choose **Direct Mapping** and click **Workflow Input**.

11. Click **Input Name** and then **Create Workflow Input**.



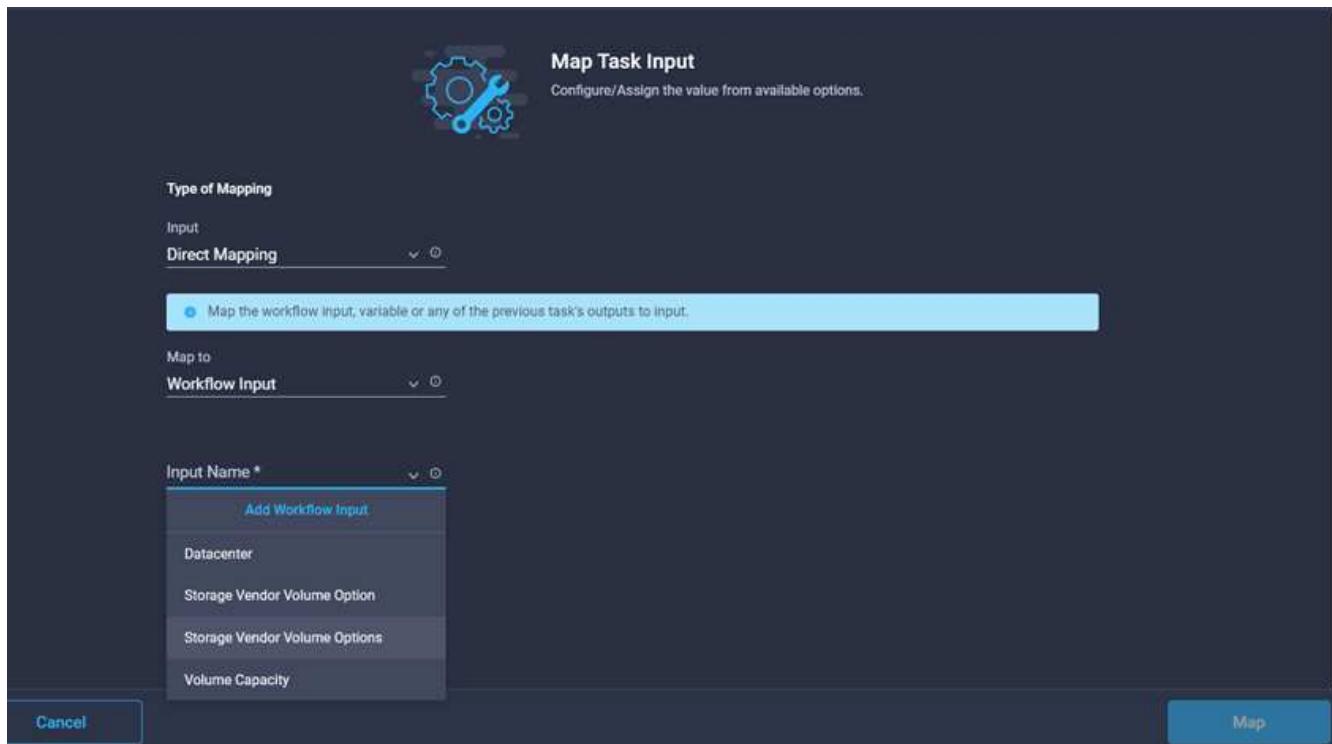
12. In the Add Input wizard, complete the following steps:

- Provide a display name and reference name (optional).
- Select **Datacenter** as the type.
- Click **Set Default Value and Override**.
- Click **Select Datacenter**.
- Click the data center associated with the new datastore and then click **Select**.

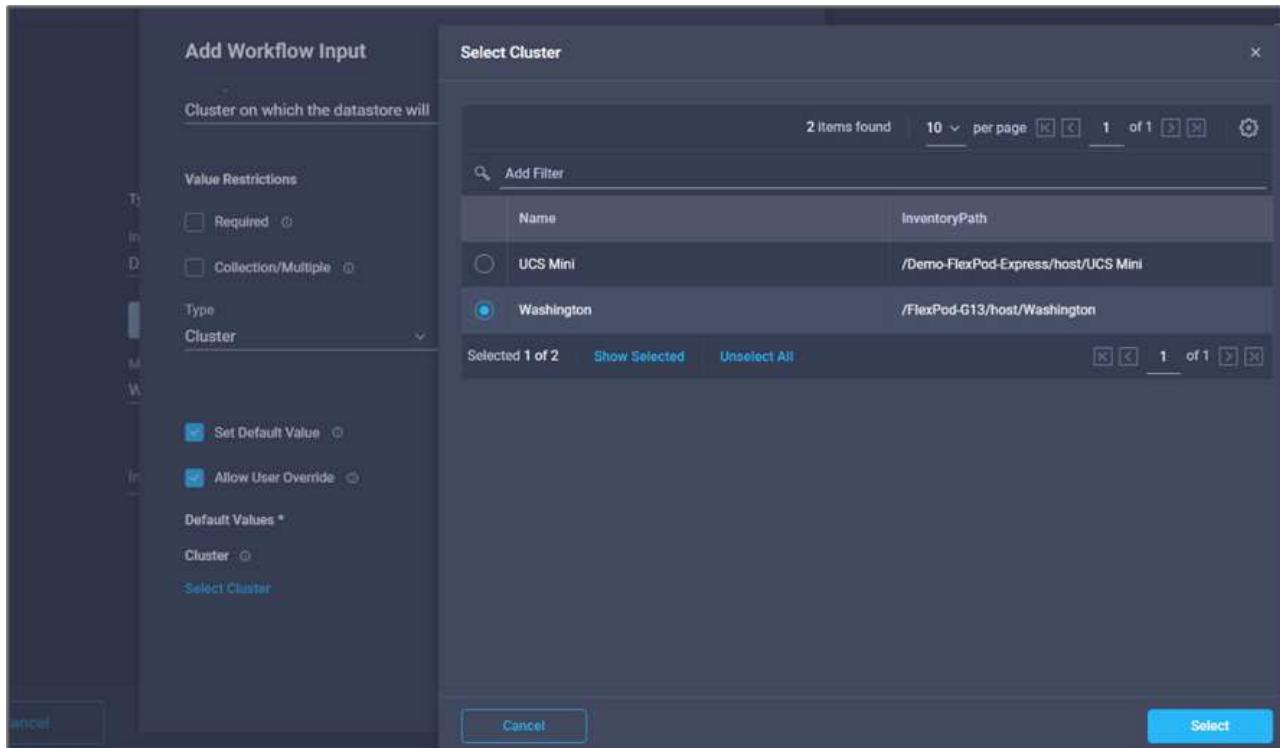


- Click **Add**.

13. Click **Map**.
14. Click **Map** in the **Cluster** field.
15. Choose **Direct Mapping** and click **Workflow Input**.



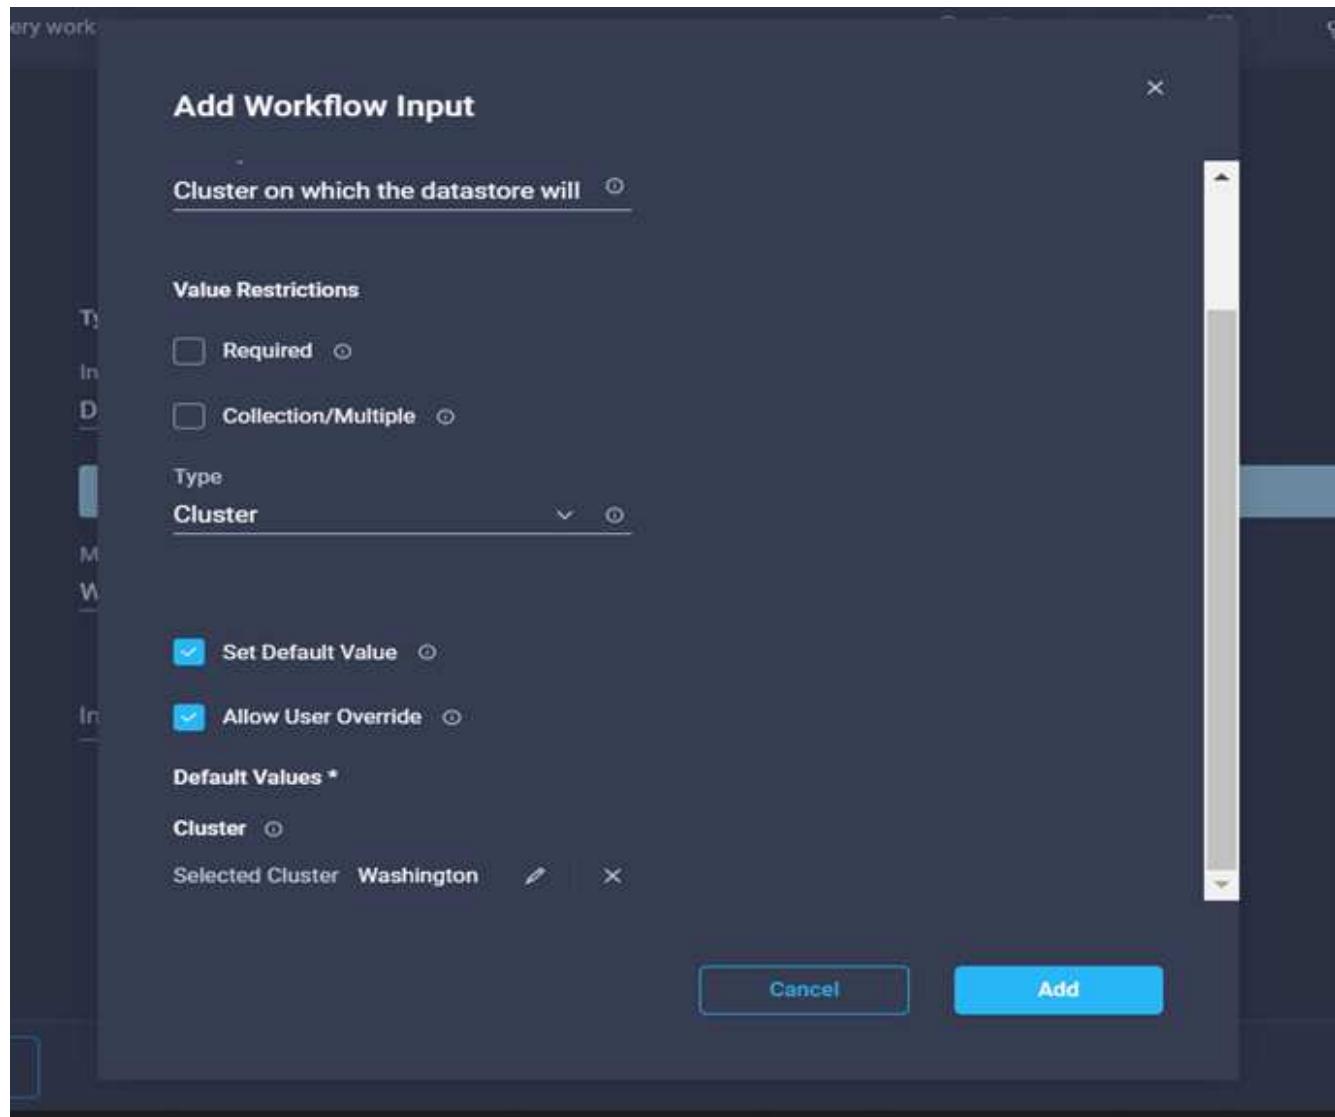
16. In the Add Input wizard, complete the following steps:
  - a. Provide a display name and reference name (optional).
  - b. Click **Required**.
  - c. Select Cluster as the type.
  - d. Click **Set Default Value and Override**.
  - e. Click **Select Cluster**.
  - f. Click the cluster associated with the new datastore.
  - g. Click **Select**.



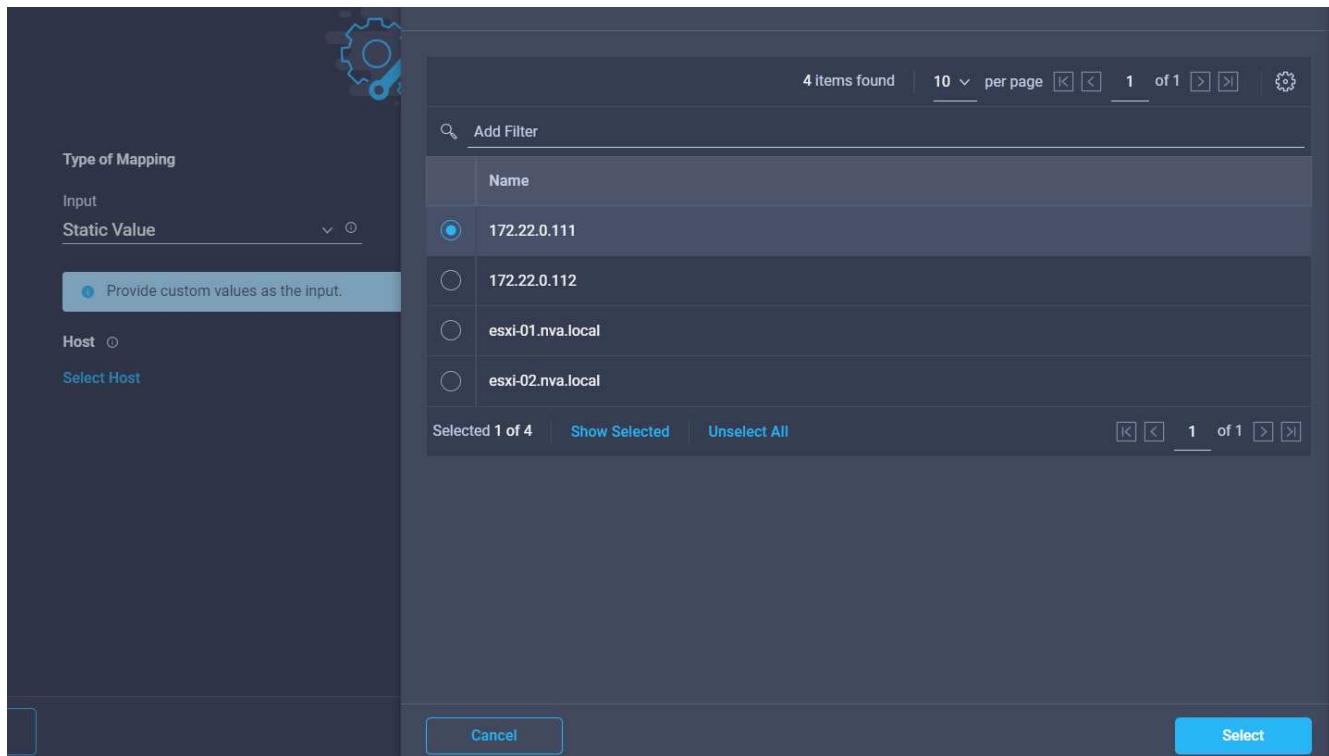
h. Click **Add**.

17. Click **Map**.

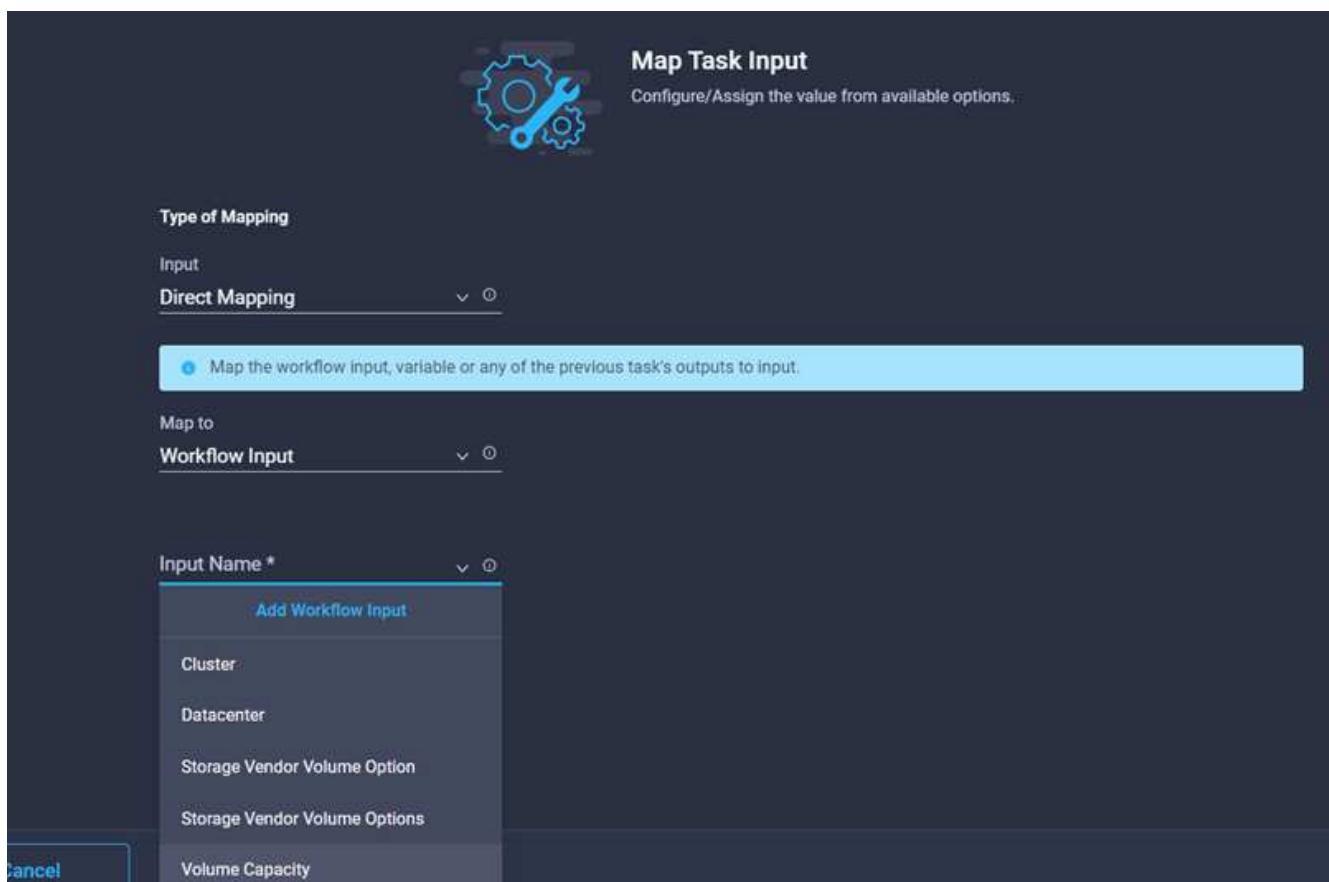
18. Click **Map** in the **Host** field.



19. Choose **Static Value** and click the host on which the datastore will be hosted. If a cluster is specified, then the host is ignored.

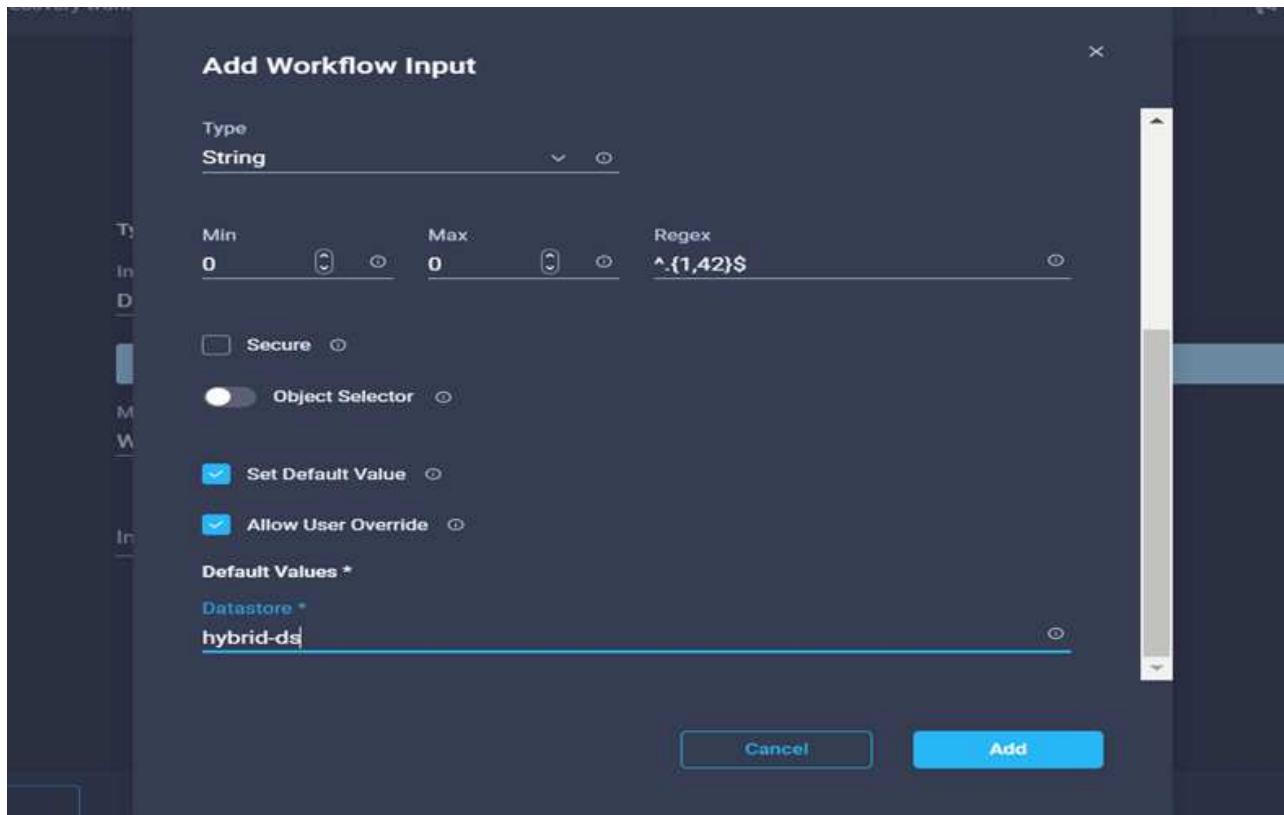


20. Click **Select and Map**.
21. Click **Map** in the **Datastore** field.
22. Choose **Direct Mapping** and click **Workflow Input**.
23. Click **Input Name** and **Create Workflow Input**.



24. In the Add Input wizard:

- Provide a display name and reference name (optional).
- Click **Required**.
- Click **Set Default Value and Override**.
- Provide a default value for the datastore and click **Add**.

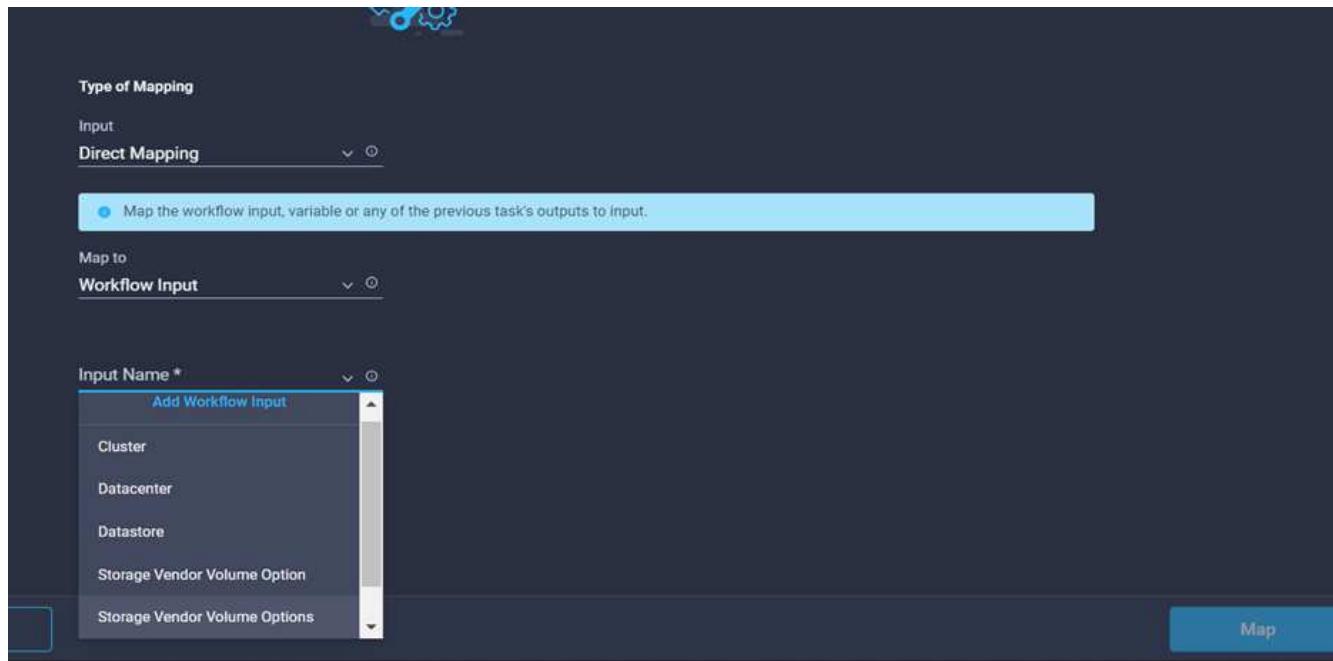


25. Click **Map**.

26. Click **Map** in the input field **Type of Datastore**.

27. Choose **Direct Mapping** and click **Workflow Input**.

28. Click **Input Name** and **Create Workflow Input**.



29. In the Add Input wizard, complete the following steps:

- Provide a display name and reference name (optional) and click **Required**.
- Make sure to select the type **Types of Datastore** and click **Set Default Value and Override**.

ter rr 93

### Add Workflow Input

Display Name \*

Type of Datastore

Description

Value Restrictions

Required

Collection/Multiple

Type

Set Default Value

Allow User Override

Default Values \*

Type of Datastore

Cancel Add

- c. Provide the Remote Path. This is the remote path of the NFS mount point.
- d. Provide the host names or IP addresses of remote NFS server in NFS Server Address.
- e. Click the **Access Mode**. The Access mode is for the NFS server. Click read-only if volumes are exported as read-only. Click **Add**.

Edit > Add Workflow Input

**Default Values \***

**Type of Datastore**

**Type of Datastore**  NFS3  VMFS-6  VMFS-5  NFS4.1

**Remote Path \***

/mssql\_data\_vol

**NFS Server Address \***

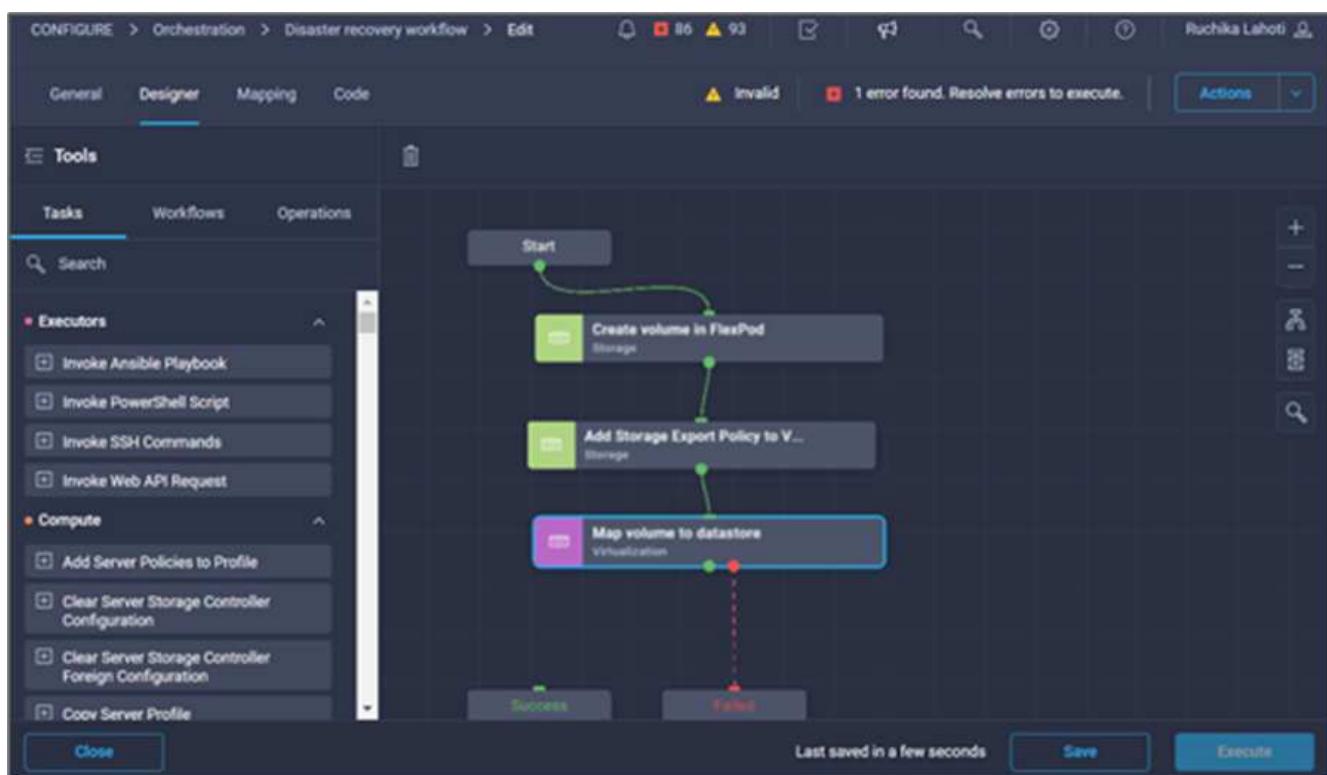
172.22.4.155

**Access Mode**  Read Write  Read Only

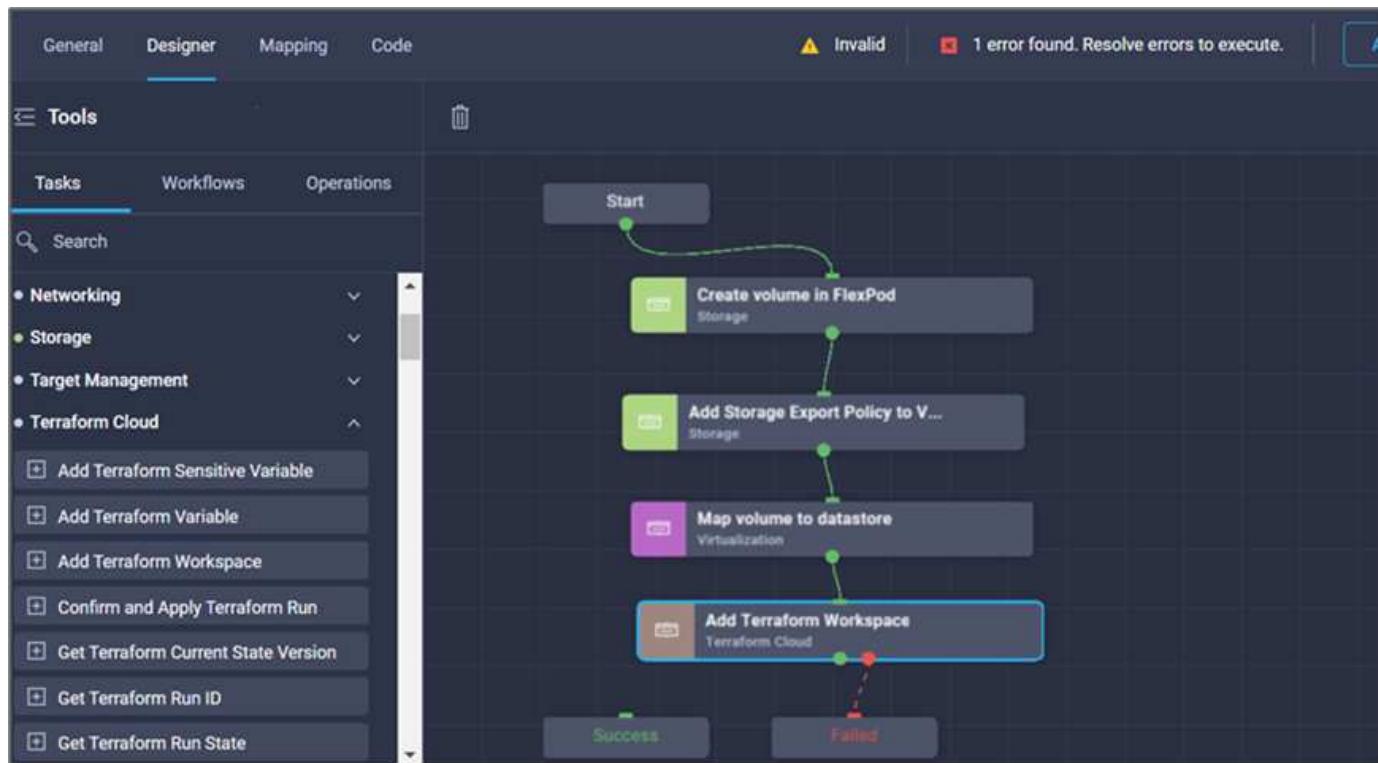
**Add**

30. Click **Map**.

31. Click **Save**.

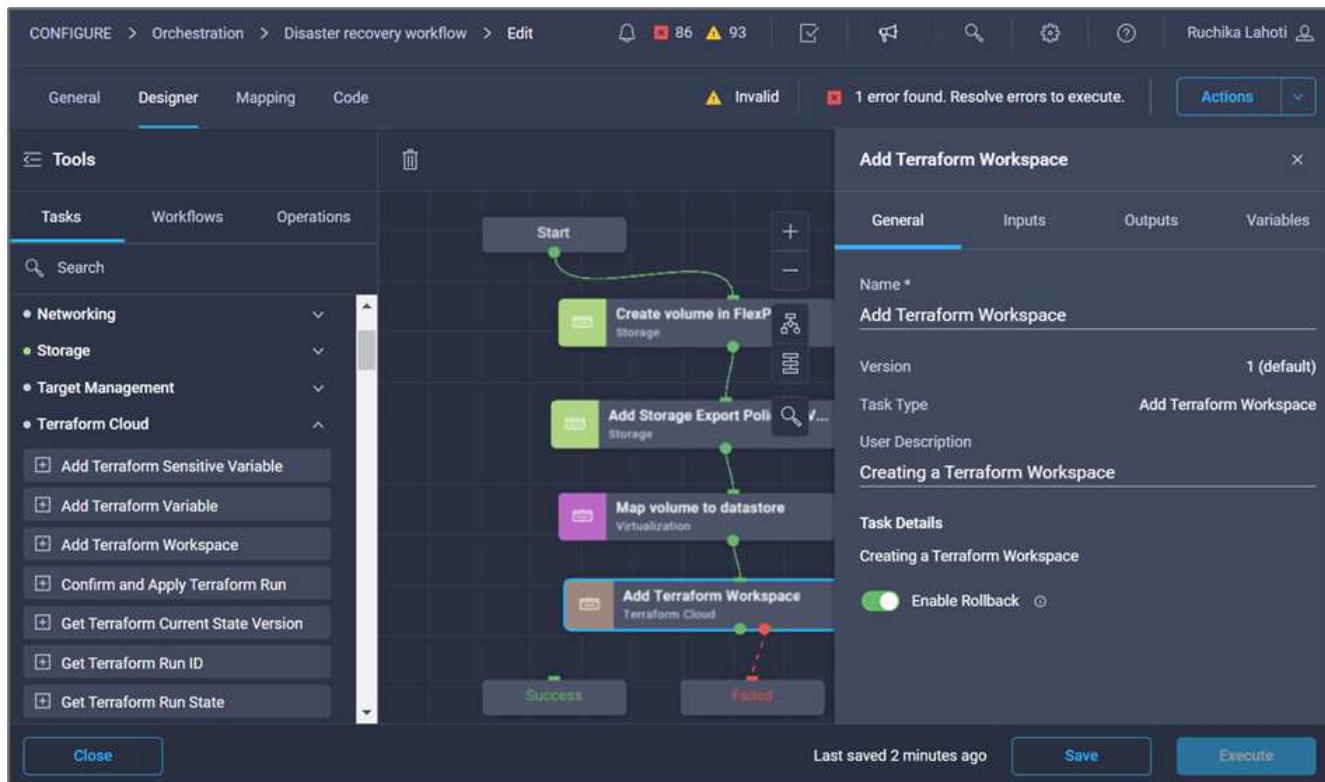


This completes the task of creating the datastore. All the tasks performed in the on-premises FlexPod Datacenter are completed.

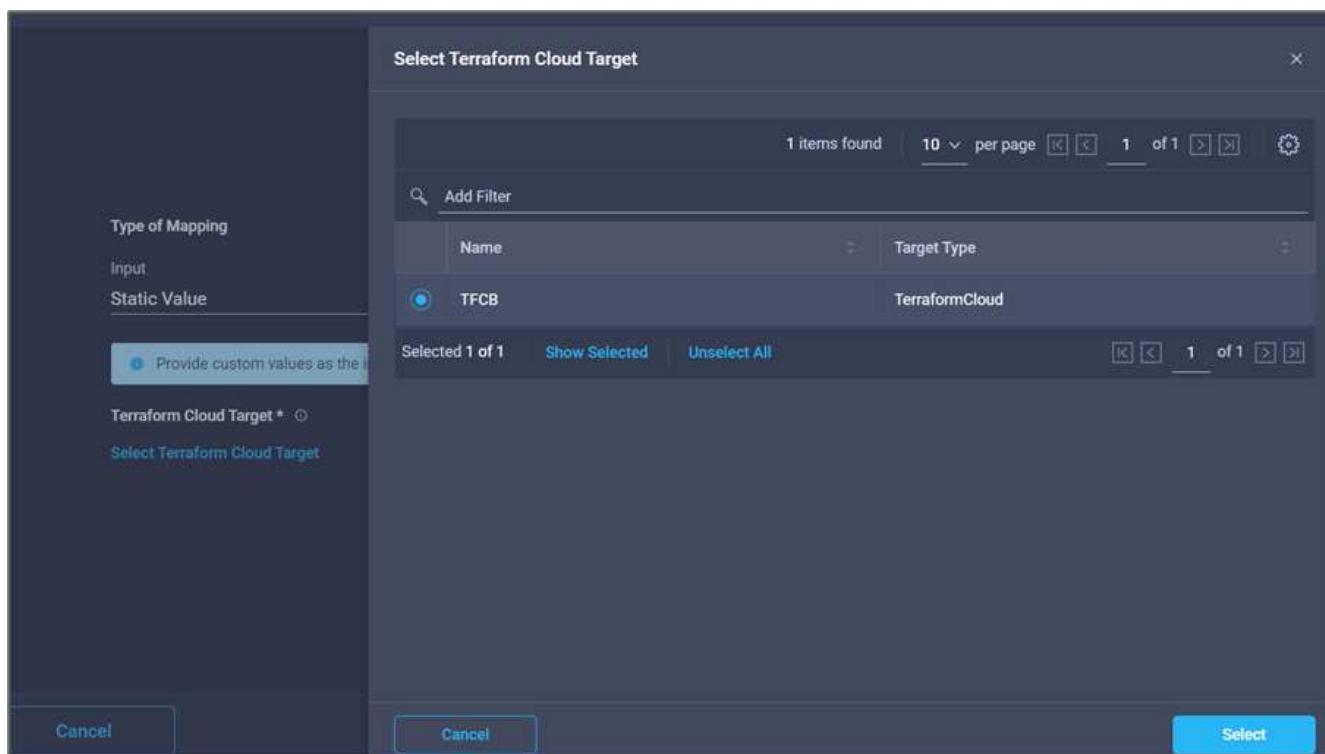


#### Procedure 5: Add a new Terraform workspace

1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Terraform Cloud > Add Terraform Workspace** task from the Tools section in the Design area.
3. Use Connector to connect the **Map volume to Datastore** and **Add Terraform Workspace** tasks and click **Save**.
4. Click **Add Terraform Workspace**. In the Task Properties area, click the **General** tab. Optionally, you can change the Name and Description for this task.

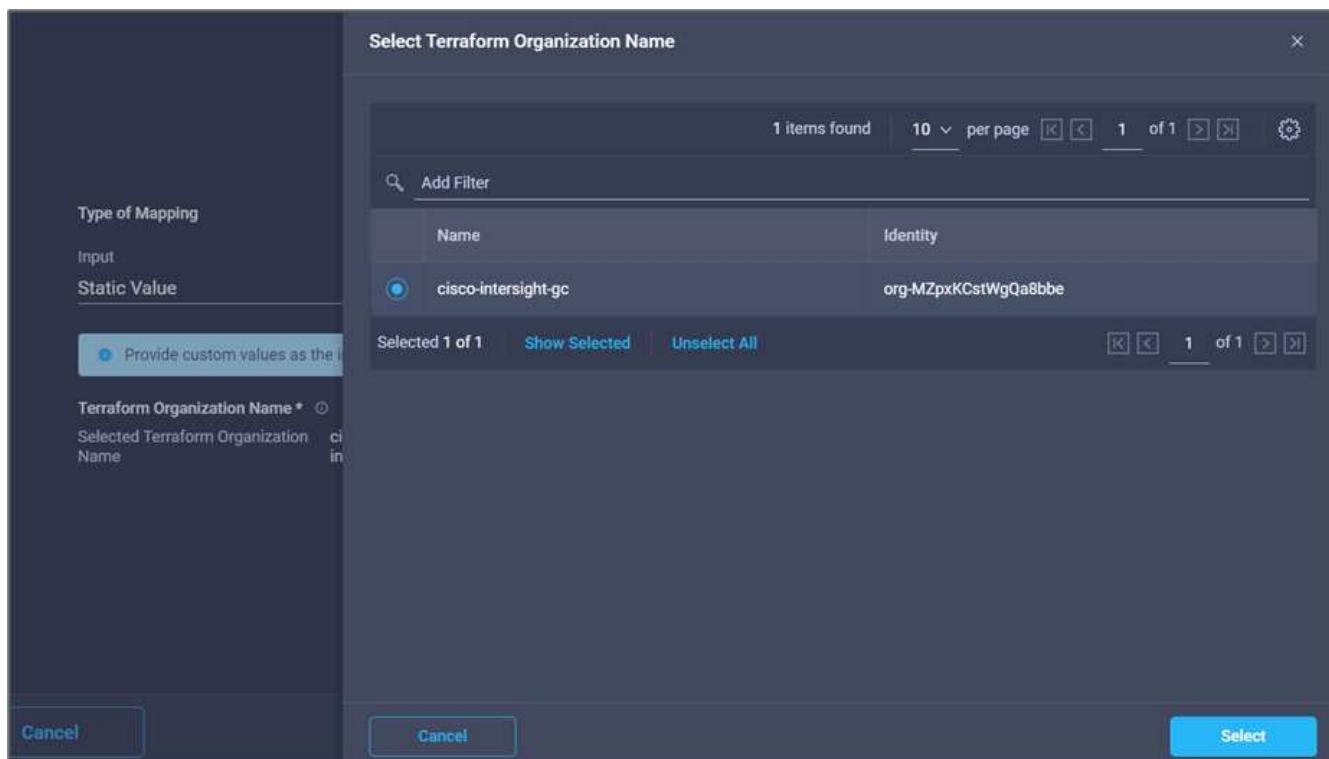


5. In the Task Properties area, click **Inputs**.
6. Click **Map** in the input field **Terraform Cloud Target**.
7. Choose **Static Value** and click **Select Terraform Cloud Target**. Select the Terraform Cloud for Business account that was added as explained in [Configure Cisco Intersight Service for HashiCorp Terraform](#).

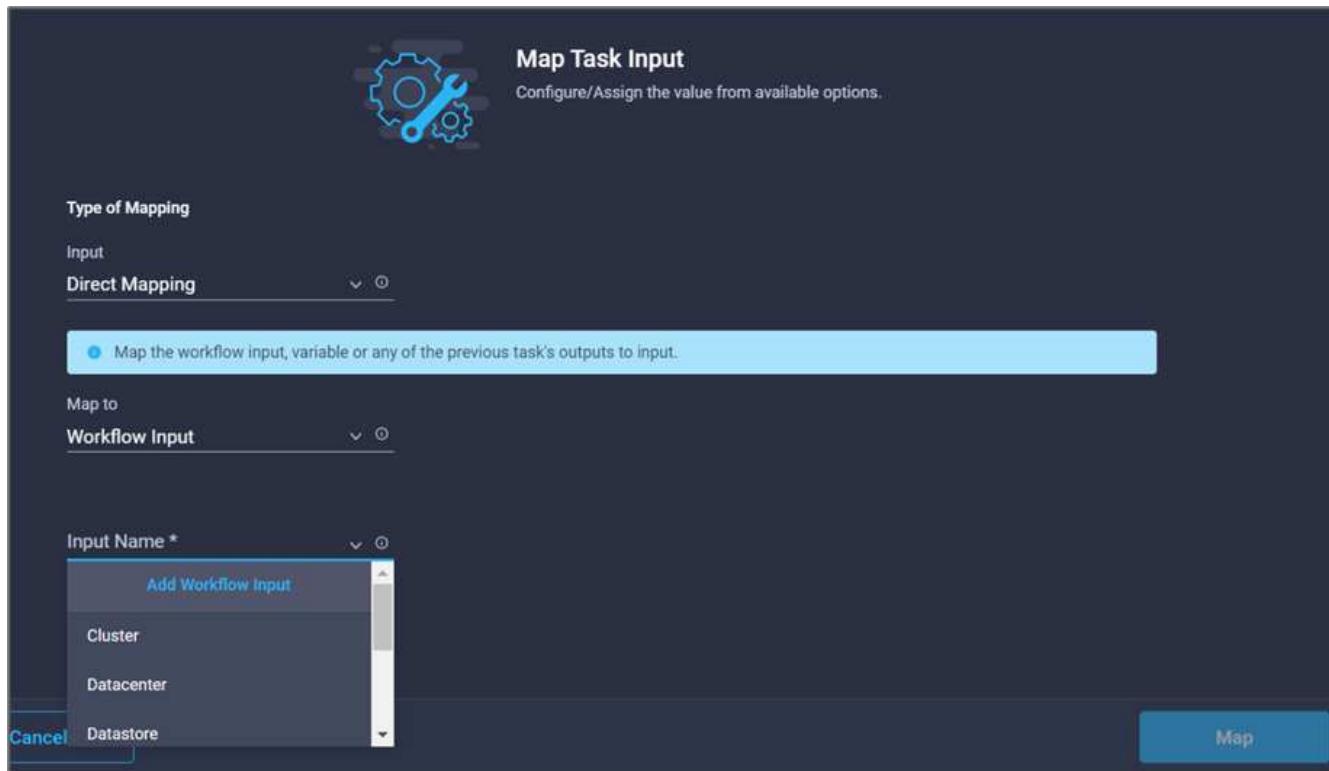


8. Click **Map**.

9. Click **Map** in the input field **Terraform Organization Name**.
10. Choose **Static Value** and then click **Select Terraform Organization**. Select the name of the Terraform Organization that you are part of in your Terraform Cloud for Business account.

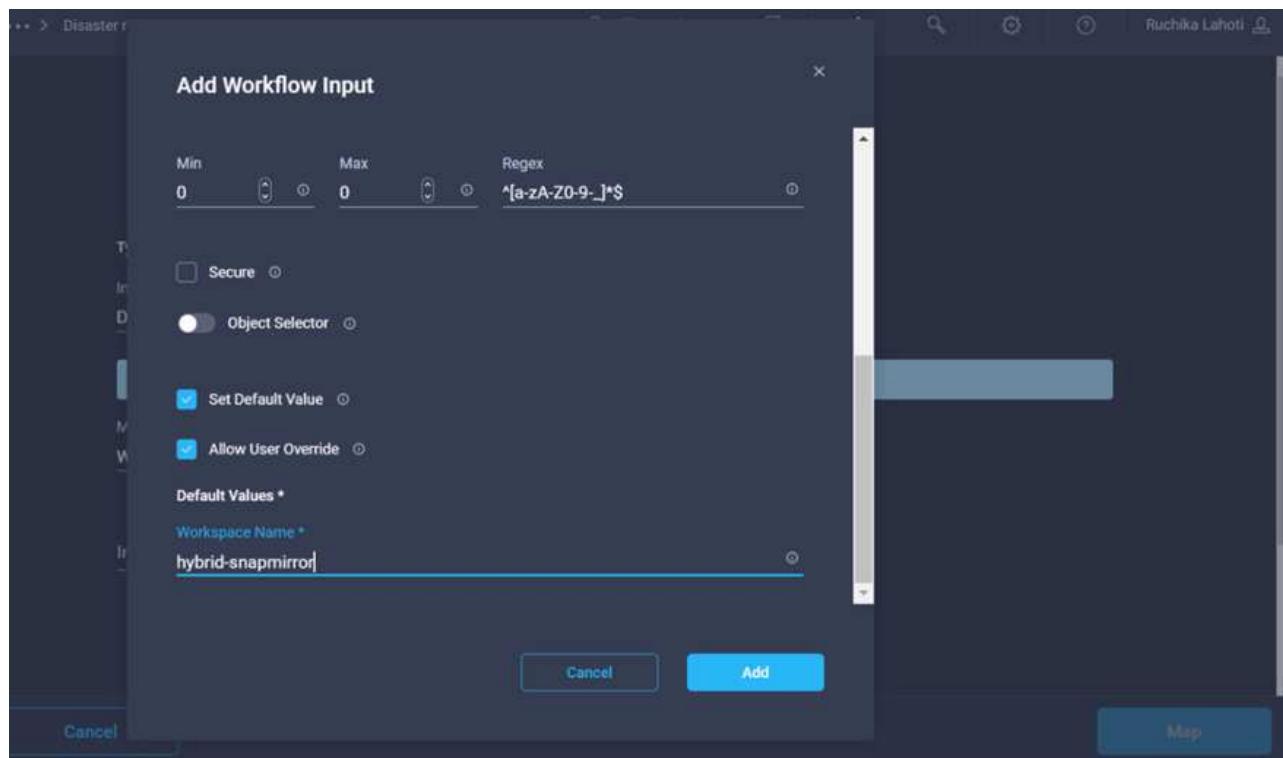


11. Click **Map**.
12. Click **Map** in the **Terraform Workspace Name** field. This is the new workspace in the Terraform Cloud for Business account.
13. Choose **Direct Mapping** and click **Workflow Input**.
14. Click **Input Name** and **Create Workflow Input**.

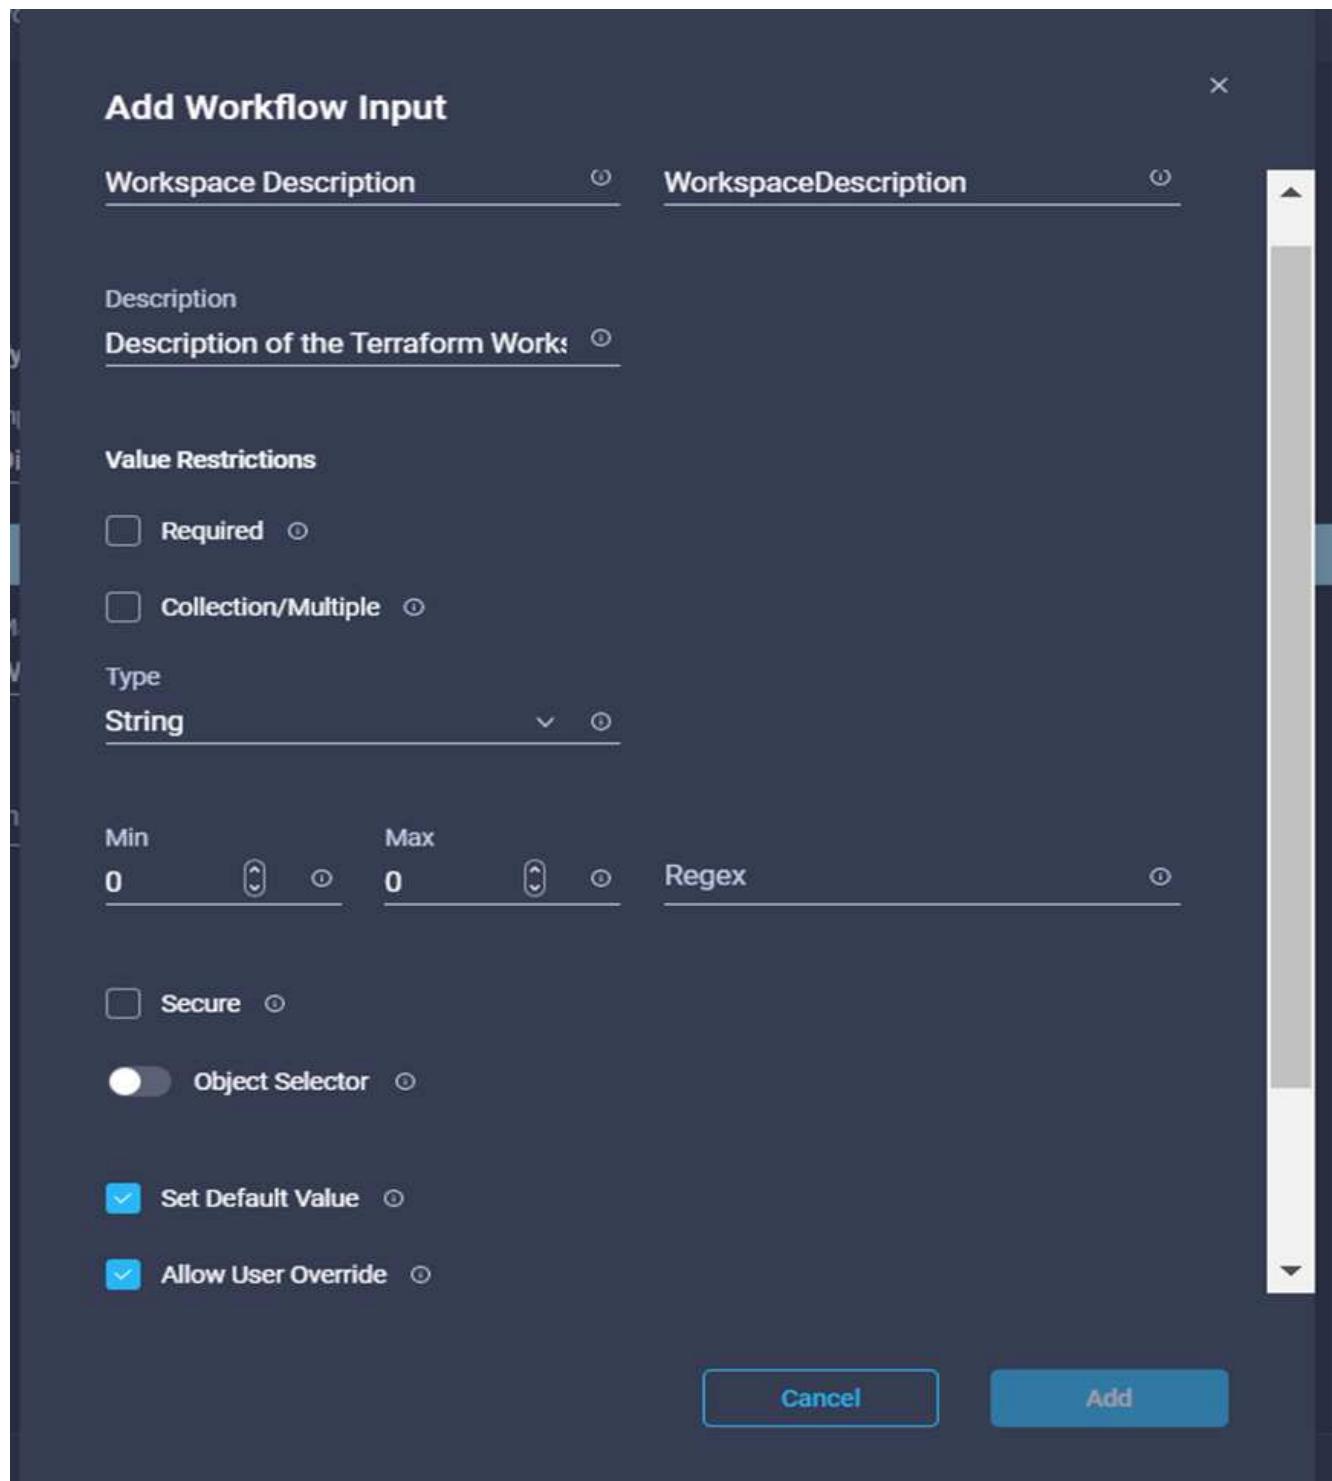


15. In the Add Input wizard, complete the following steps:

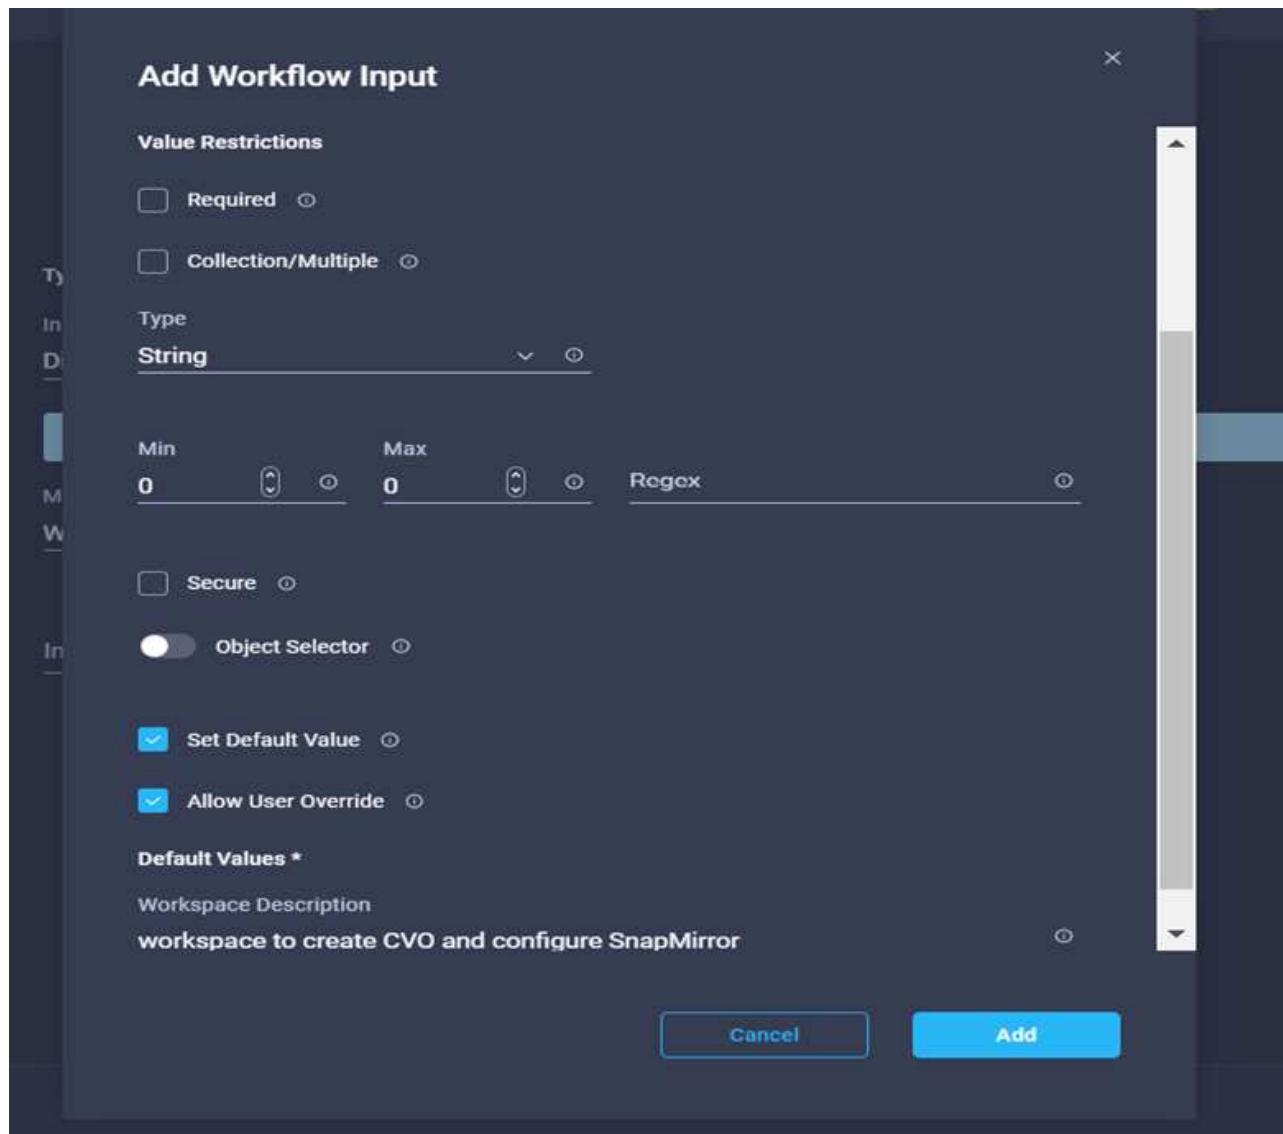
- Provide a display name and reference name (optional).
- Click **Required**.
- Make sure to select **String** for **Type**.
- Click **Set Default Value and Override**.
- Provide a default name for workspace.
- Click **Add**.



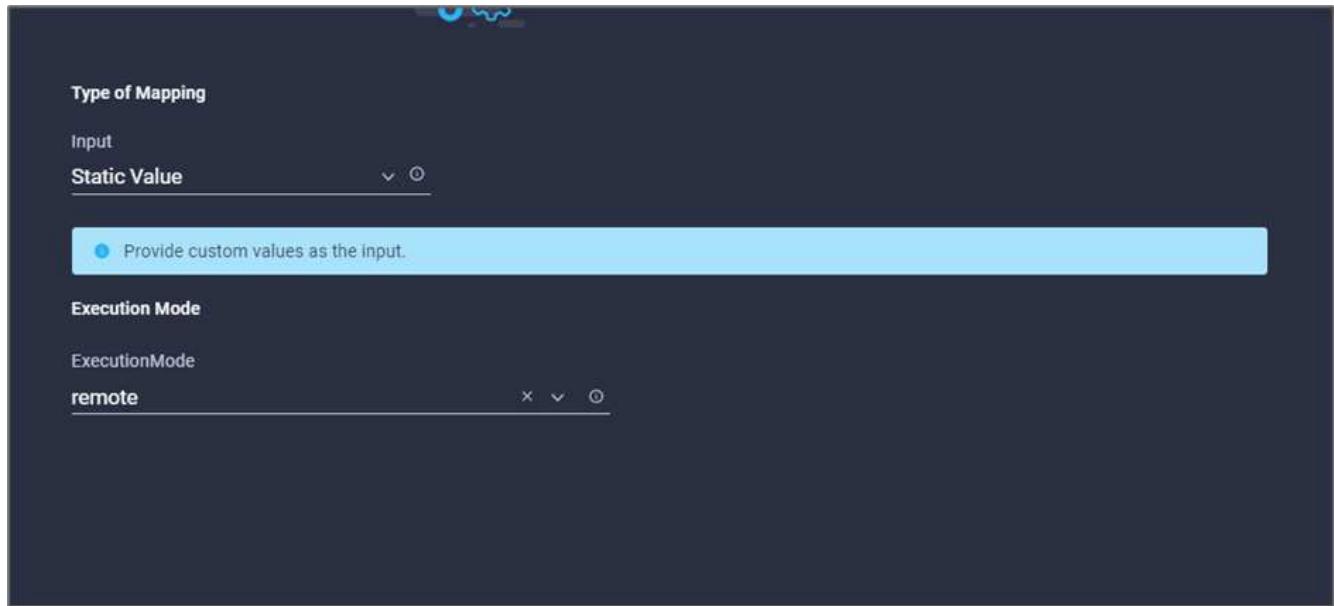
16. Click **Map**.
17. Click **Map** in the **Workspace Description** field.
18. Choose **Direct Mapping** and click **Workflow Input**.
19. Click **Input Name** and **Create Workflow Input**.



20. In the Add Input wizard, complete the following steps:
  - a. Provide a display name and reference name (optional).
  - b. Make sure to select **String** for **Type**.
  - c. Click **Set Default Value and Override**.
  - d. Provide a workspace description and click **Add**.



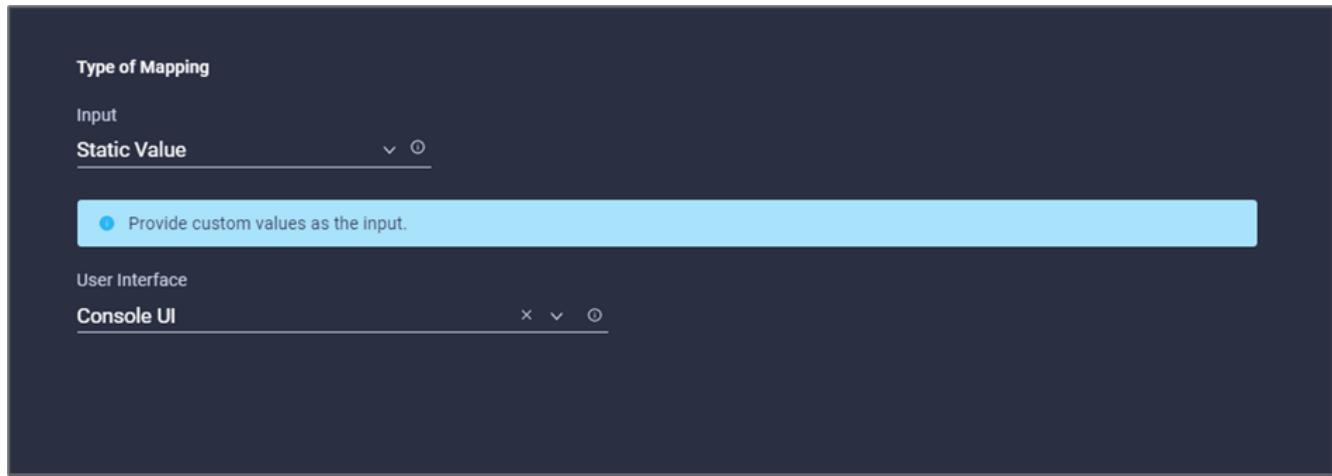
21. Click **Map**.
22. Click **Map** in the **Execution Mode** field.
23. Choose **Static Value**, click **Execution Mode**, and then click **remote**.



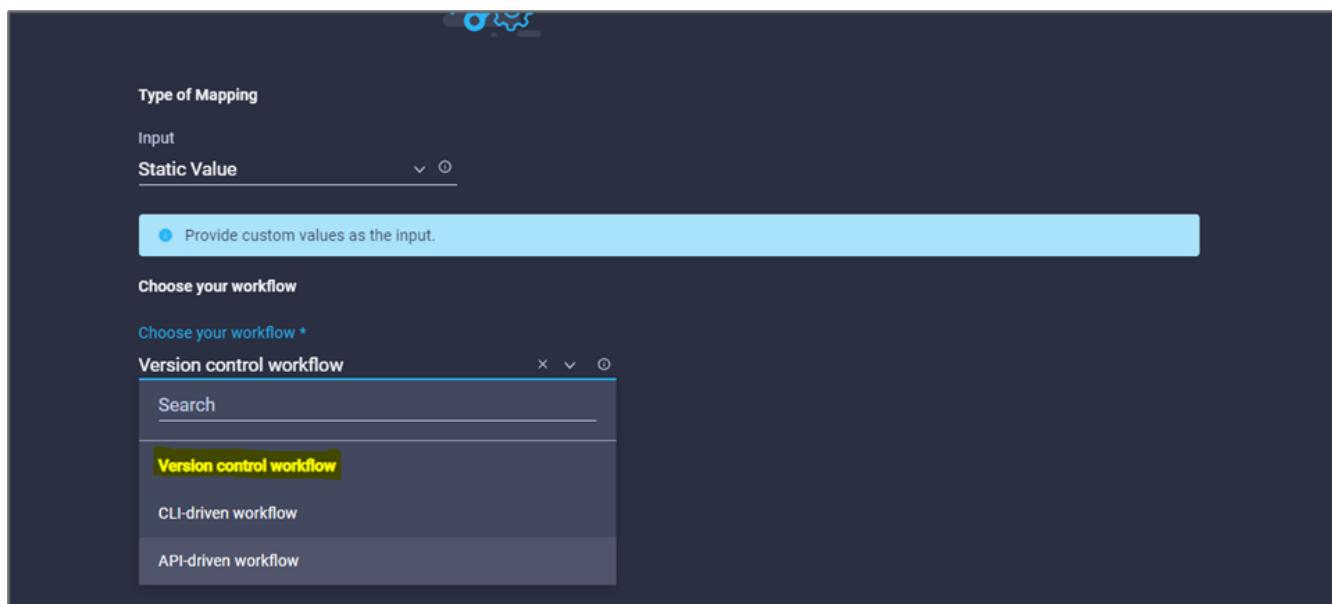
24. Click **Map**.
25. Click **Map** in the **Apply Method** field.
26. Choose **Static Value** and click **Apply Method**. Click **Manual Apply**.



27. Click **Map**.
28. Click **Map** in the **User Interface** field.
29. Choose **Static Value** and click **User Interface**. Click **Console UI**.



30. Click **Map**.
31. Click **Map** in the input field and select your workflow.
32. Select **Static Value**, and click **Choose Your Workflow**. Click **Version Control Workflow**.



33. Provide the following GitHub repository details:
  - a. In **Repository Name**, enter the name of the repository detailed in the section "["Set up environment prerequisites"](#)".
  - b. Provide the OAuth Token ID as detailed in the section "["Set up environment prerequisites"](#)".
  - c. Select the **Automatic Run Triggering** option.

**Type of Mapping**

Input

**Static Value**

v ⓘ

ⓘ Provide custom values as the input.

**Choose your workflow**

Choose your workflow \*

**Version control workflow**

x v ⓘ

**Choose repository and configure settings**

Repository Name \*

**NetApp-Automation/FlexPod-hybrid-cloud-for-GCP-wit** ⓘ

Oauth Token ID \*

[REDACTED]

ⓘ

**Terraform Working Directory**

ⓘ

**Automatic Run Triggering**

Automatic Run Triggering Options

**Always Trigger Runs**

x v ⓘ

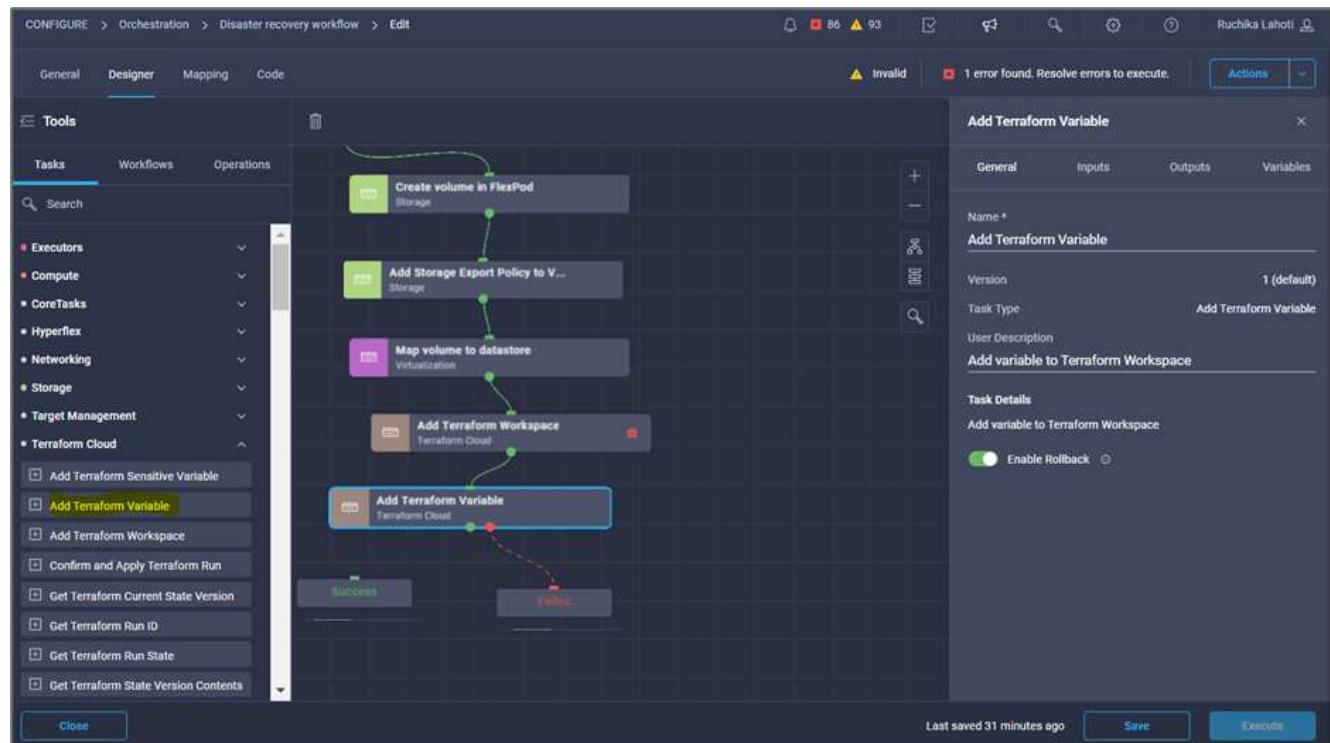
34. Click **Map**.35. Click **Save**.

This completes the task of creating a workspace in a Terraform Cloud for Business account.

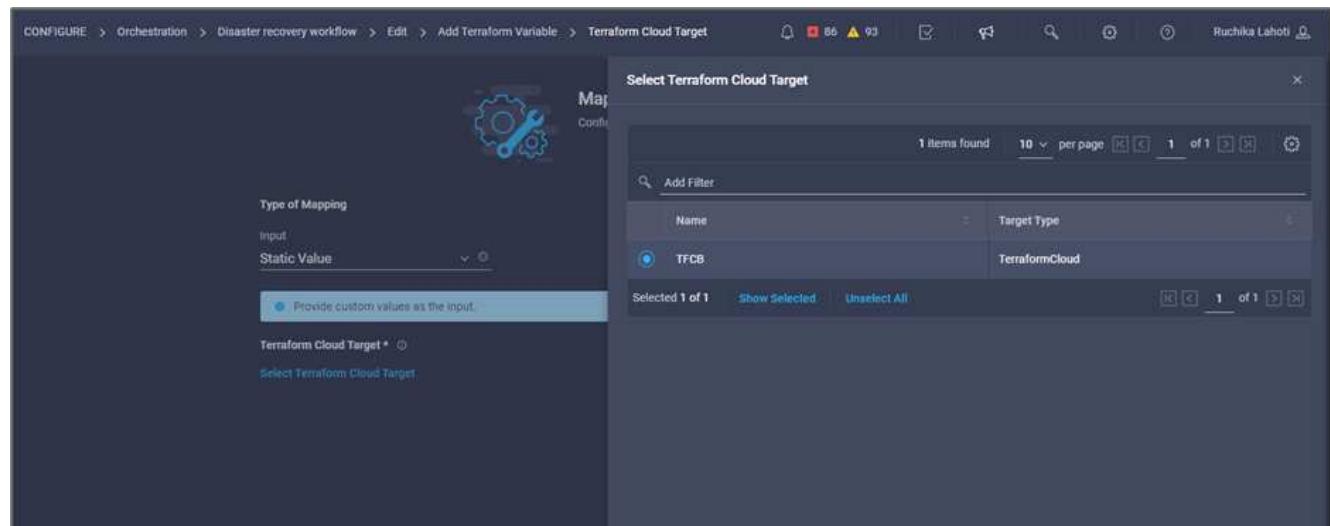
**Procedure 6: Add non-sensitive variables to workspace**

1. Go to the **Designer** tab and click the **Workflows from Tools** section.
2. Drag and drop the **Terraform > Add Terraform Variables** workflow from the **Tools** section in the **Design** area.
3. Use Connector to connect the **Add Terraform Workspace** and **Add Terraform Variables** tasks. Click **Save**.

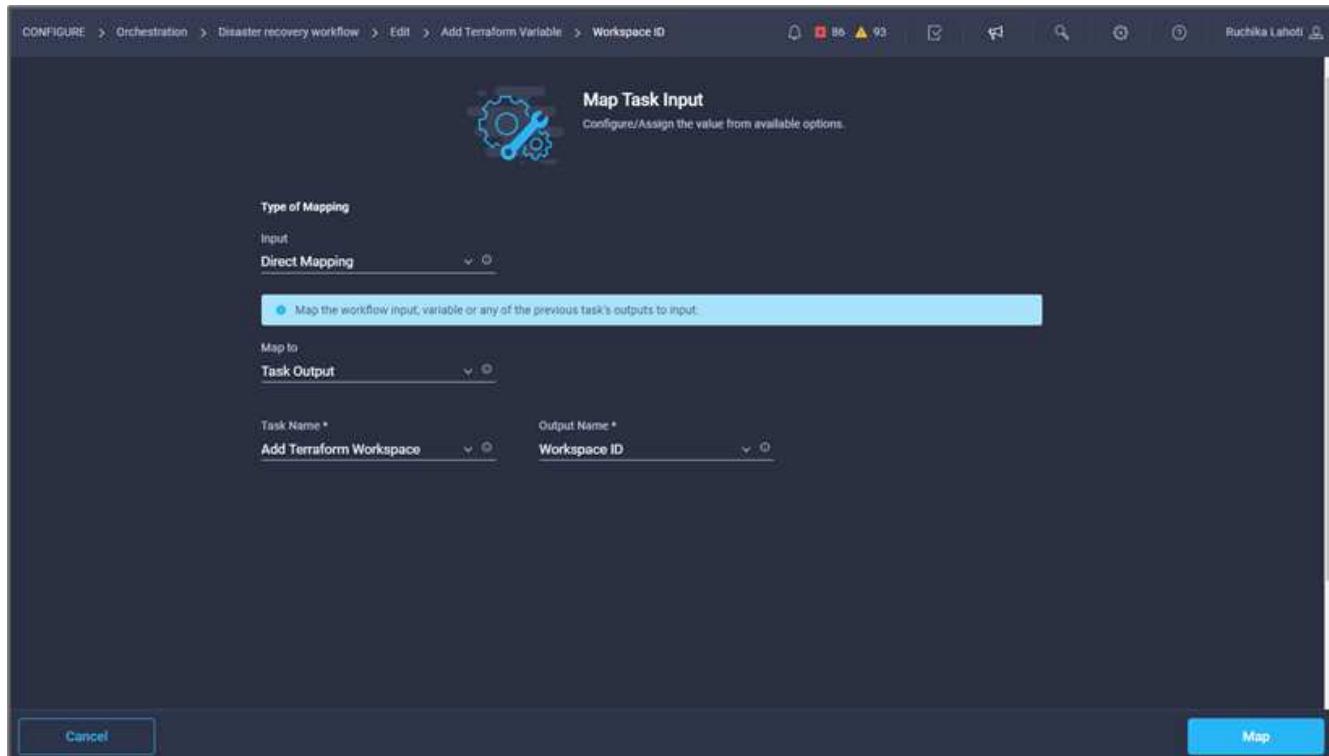
4. Click **Add Terraform Variables**. In the **Workflow Properties** area, click the **General** tab. Optionally, you can change the name and description for this task.



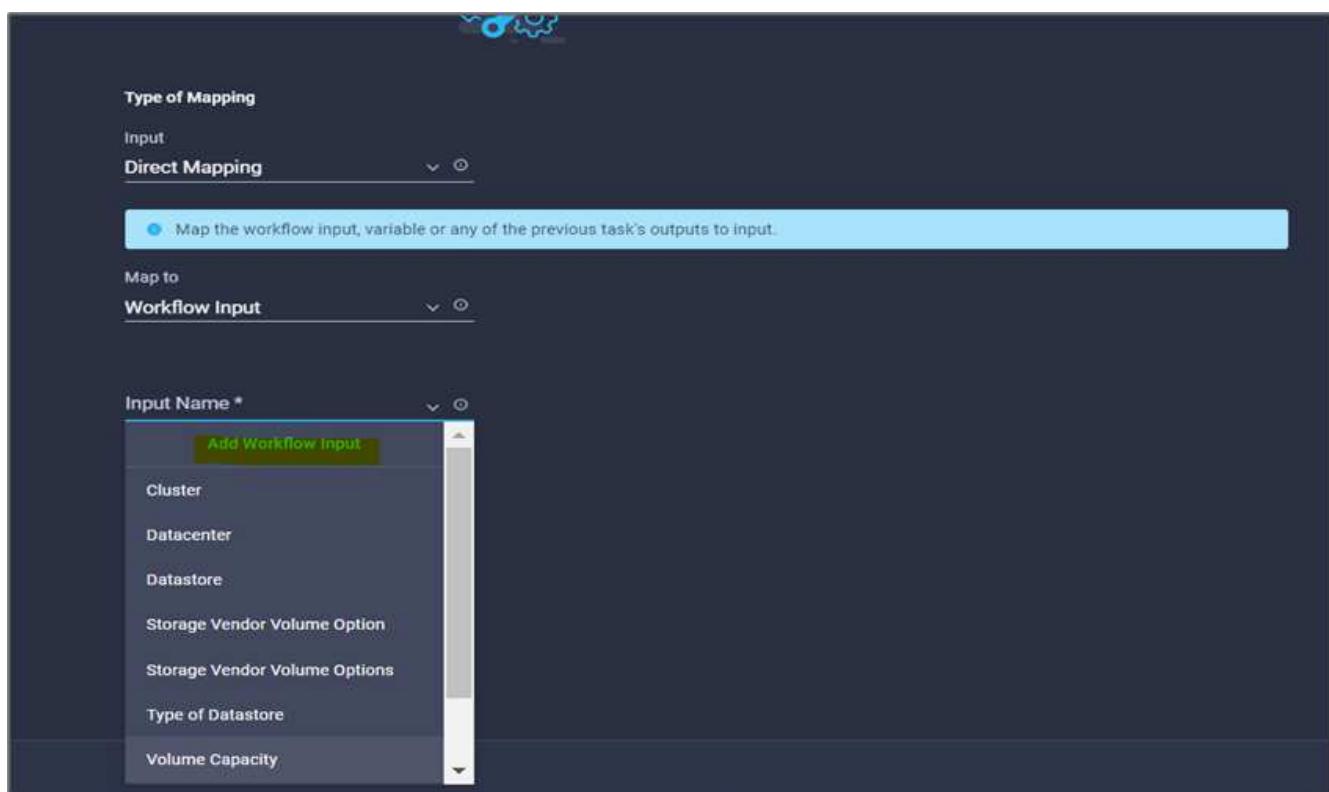
5. In the **Workflow Properties** area, click **Inputs**.
6. Click **Map** in the **Terraform Cloud Target** field.
7. Choose **Static Value** and click **Select Terraform Cloud Target**. Select the Terraform Cloud for Business account that was added as explained in [Configure Cisco Intersight Service for HashiCorp Terraform](#).



8. Click **Map**.
9. Click **Map** in the **\*Terraform Organization Name** field.
10. Choose **Static Value** and click **Select Terraform Organization**. Select the name of the Terraform Organization that you are part of in your Terraform Cloud for Business account.



11. Click **Map**.
12. Click **Map** in the **Terraform Workspace Name** field.
13. Choose **Direct Mapping** and click **Task Output**.
14. Click **Task Name** and click **Add Terraform Workspace**.



15. Click **Output Name** and click **Workspace Name**.

16. Click **Map**.
17. Click **Map** in the **Add Variables Options** field.
18. Choose **Direct Mapping** and click **Workflow Input**.
19. Click **Input Name** and **Create Workflow Input**.

The screenshot shows the 'Add Workflow Input' dialog box. The 'Display Name' field is set to 'Terraform Variable'. The 'Reference Name' field is set to 'TerraformAddVariable'. The 'Description' field contains the text 'Terraform Variable to be added'. Under 'Value Restrictions', the 'Required' checkbox is checked, and the 'Collection/Multiple' checkbox is unchecked. The 'Type' is set to 'String'. The 'Min' and 'Max' values are both set to 0. There are 'Secure' and 'Object Selector' options, both of which are unchecked. At the bottom, there are 'Cancel' and 'Add' buttons.

20. In the Add Input wizard, complete the following steps:
  - a. Provide a display name and reference name (Optional).
  - b. Make sure to select **String** for the **Type**.
  - c. Click **Set Default Value and Override**.
  - d. Click **Variable Type** and then click **Non-Sensitive Variables**.

21. In the **Add Terraform Variables** section, provide the following information:

- **Key.** name\_of\_on-prem-ontap
- **Value.** Provide the name of on-premises ONTAP.
- **Description.** Name of the on-premises ONTAP.

22. Click **+** to add additional variables.

Set Default Value

Allow User Override

**Default Values \***

**Terraform Variable**

Key \*

name\_of\_on-prem-ontap

Value

Provide the name of On-premise ONTAP added in section Deploying

Description

Name of the On-premise ONTAP

HCL

**Cancel** **Add**

23. Add all the Terraform Variables as shown in the following table. You can also provide a default value.

Terraform variable name	Description
name_of_on-prem-ontap	Name of the on-premises ONTAP (FlexPod)

Terraform variable name	Description
on-prem-ontap_cluster_ip	The IP address of the storage cluster management interface
on-prem-ontap_user_name	Admin username for the storage cluster
Zone	GCP region where the working environment will be created
subnet_id	GCP subnet id where the working environment will be created
vpc_id	The VPC ID where the working environment will be created
capacity_package_name	The type of license to use
source_volume	The name of the source volume
source_storage_vm_name	The name of the source SVM
destination_volume	Name of volume on Cloud Volumes ONTAP
schedule_of_replication	The default is 1 hour
name_of_volume_to_create_on_cvo	Name of the cloud volume
workspace_id	The workspace_id where the working environment will be created
Project_id	The project_id where the working environment will be created
name_of_cvo_cluster	The name of the Cloud Volumes ONTAP working environment
gcp_service_account	gcp_service_account of Cloud Volumes ONTAP working environment

24. Click **Map** and then **Save**.

## Add Terraform Variable

General Inputs Outputs Variables

Search

Terraform Cloud Target *	○	<a href="#">Edit Mapping</a>
Custom Value		<a href="#">View Value</a> <a href="#">Edit</a>
Workspace ID *	○	<a href="#">Edit Mapping</a>
Task Output		WorkspaceId   Add Terraform Work... <a href="#">Edit</a>
Terraform Variable	○	<a href="#">Edit Mapping</a>
Workflow Input		Terraform Variables <a href="#">Edit</a>

Last saved an hour ago [Save](#) [Execute](#)

This completes the task of adding the required Terraform variables to the workspace. Next, add the required sensitive Terraform variables to the workspace. You can also combine both into a single task.

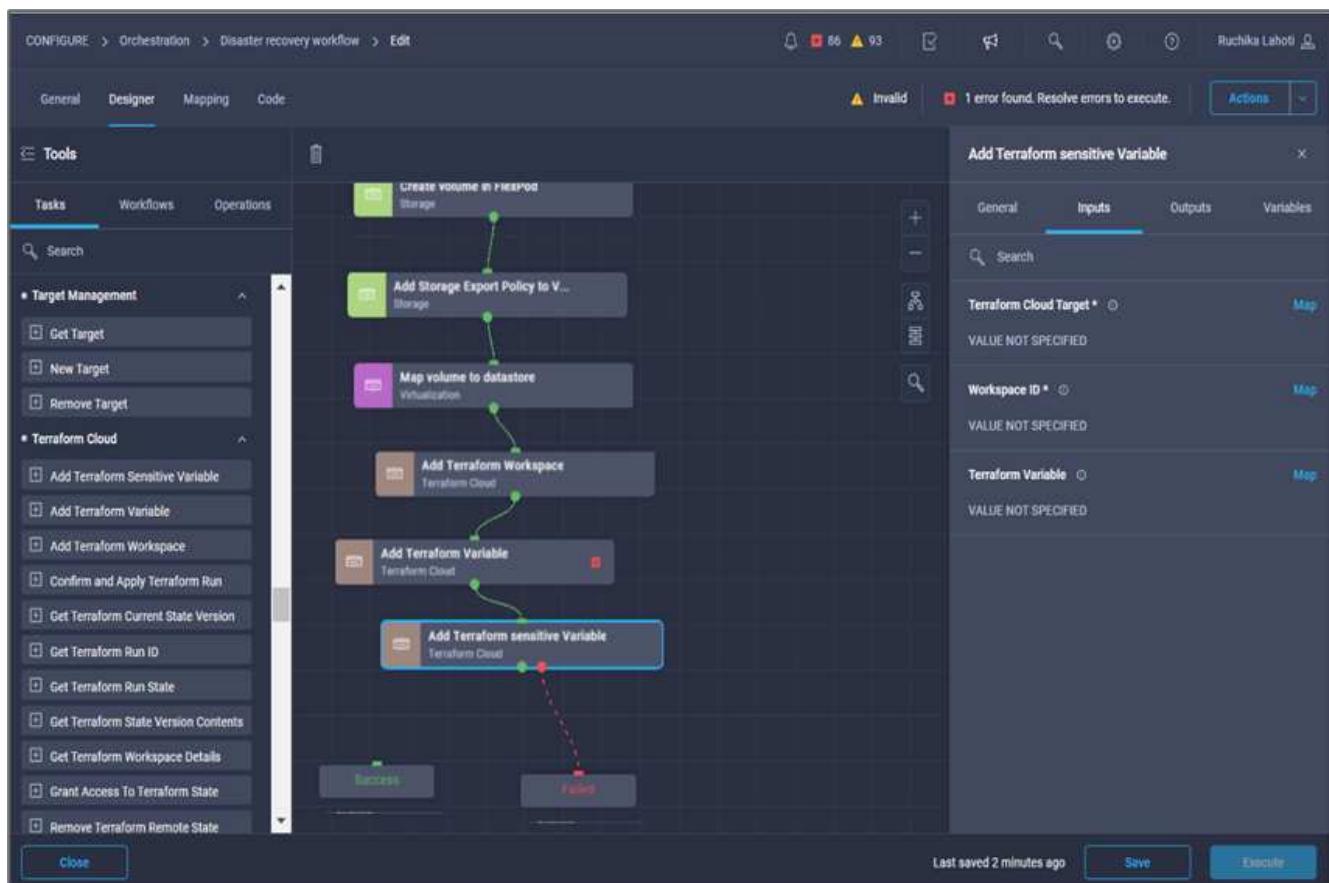
#### Procedure 7: Add sensitive variables to a workspace

1. Go to the **Designer** tab and click **Workflows** from the **Tools** section.
2. Drag and drop the **Terraform > Add Terraform Variables** workflow from the **Tools** section in the **Design** area.
3. Use Connector to connect the two **Add Terraform Workspace** tasks. Click **Save**.



A warning appears indicating that the two tasks have the same name. Ignore the error for now because you change the task name in the next step.

4. Click **Add Terraform Variables**. In the **Workflow Properties** area, click the **General** tab. Change the name to **Add Terraform Sensitive Variables**.



5. In the **Workflow Properties** area, click **Inputs**.
6. Click **Map** in the **Terraform Cloud Target** field.
7. Choose **Static Value** and click **Select Terraform Cloud Target**. Select the Terraform Cloud for Business account that was added in the section [Configure Cisco Intersight Service for HashiCorp Terraform](#).
8. Click **Map**.
9. Click **Map** in the **Terraform Organization Name** field.
10. Choose **Static Value** and click **Select Terraform Organization**. Select the name of the Terraform Organization that you are part of in your Terraform Cloud for Business account.
11. Click **Map**.
12. Click **Map** in the **Terraform Workspace Name** field.

13. Choose **Direct Mapping** and click **Task Output**.
14. Click **Task Name** and then click **Add Terraform Workspace**.
15. Click **Output Name** and click the output **Workspace Name**.
16. Click **Map**.
17. Click **Map** in the **Add Variables Options** field.
18. Choose **Direct Mapping** and then click **Workflow Input**.
19. Click **Input Name** and **Create Workflow Input**.
20. In the Add Input wizard, complete the following steps:
  - a. Provide a display name and reference name (optional).
  - b. Make sure to select **Terraform Add Variables Options** for the type.
  - c. Click **Set Default Value**.
  - d. Click **Variable Type** and then click **Sensitive Variables**.
  - e. Click **Add**.

### Add Workflow Input

Display Name \* **terraform sensitive variable**

Reference Name \* **terraformsensitivevariable**

Description **Add Variables**

Value Restrictions

Required

Collection/Multiple

Type **Terraform Add Variables Option**

Set Default Value

Allow User Override

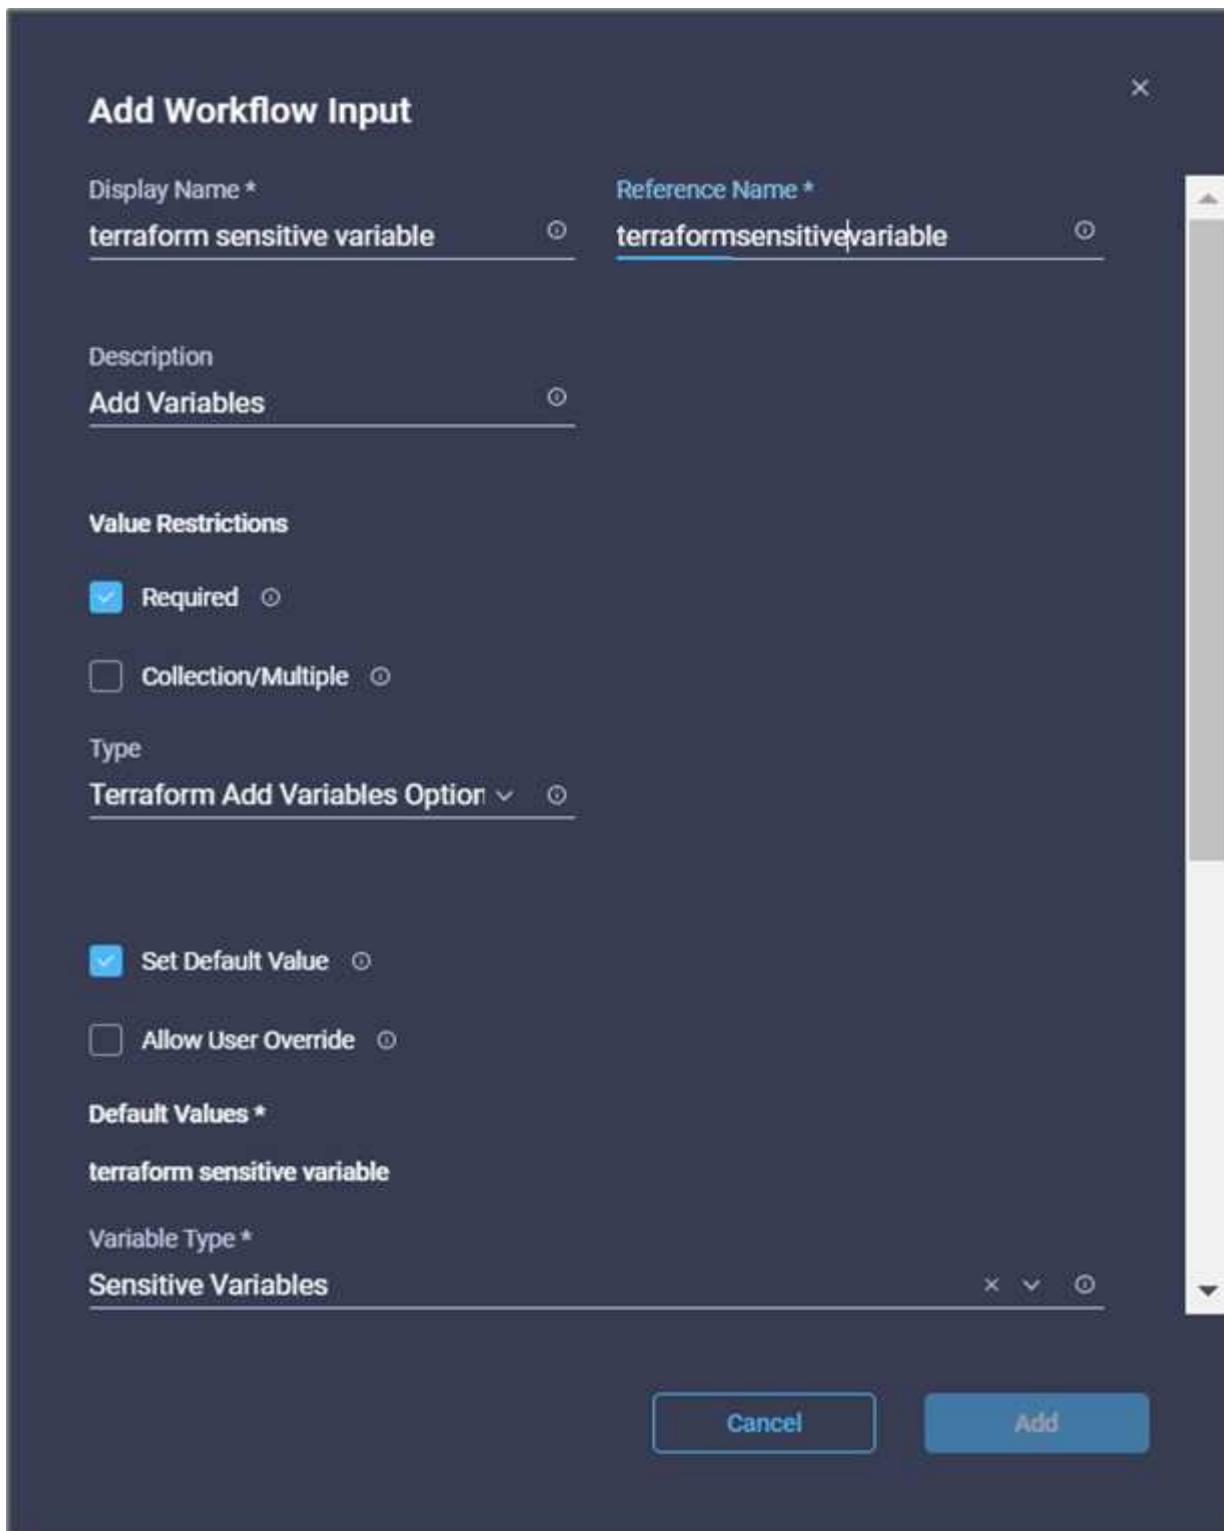
Default Values \*

**terraform sensitive variable**

Variable Type \*

**Sensitive Variables**

**Cancel** **Add**



21. In the **Add Terraform Variables** section, provide the following information:

- **Key.** `cloudmanager_refresh_token`.
- **Value.** Input the refresh token for NetApp Cloud Manager API operations.
- **Description.** Refresh token.



For more information about obtaining a refresh token for the NetApp Cloud Manager API operations, see the section [“Set up environment prerequisites.”](#)

### Add Workflow Input

Set Default Value

Allow User Override

**Default Values \***

**terraform sensitive variable**

**Variable Type \***

**Sensitive Variables**

**Add Sensitive Terraform Variables**

<b>Key *</b>	<b>Value</b>	<b>Description</b>	<b>HCL</b>
cloudmanager_refresh_token		cloudmanager refresh token	<input type="checkbox"/>

**Cancel** **Add**

22. Add all the Terraform sensitive variables as shown in the table below. You can also provide a default value.

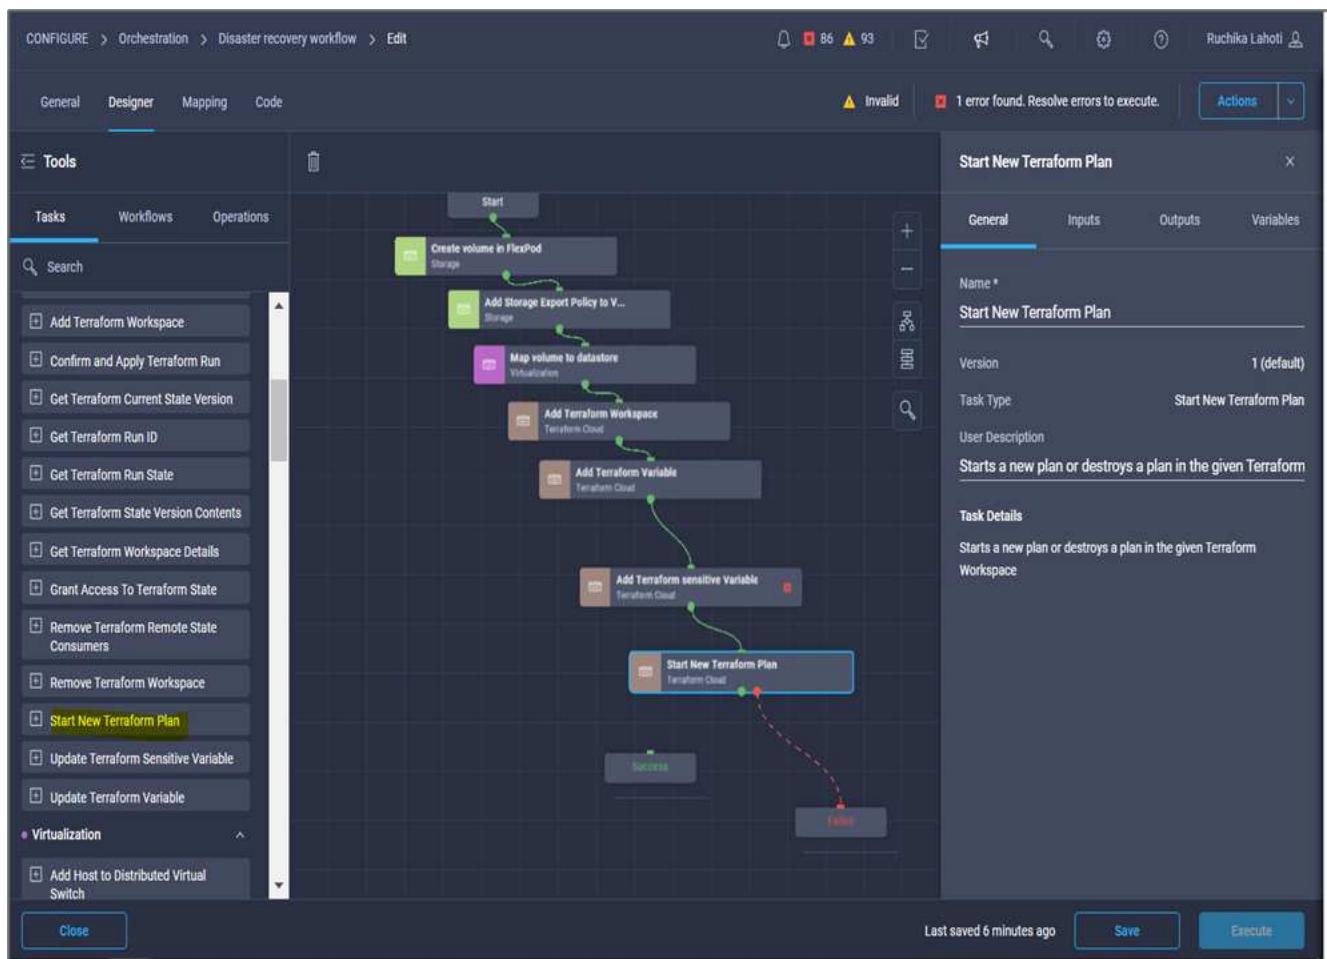
Terraform sensitive variable name	Description
cloudmanager_refresh_token	Refresh token. Obtain it from:
connector_id	The client ID of the Cloud Manager Connector. Obtain it from
cvo_admin_password	The admin password for Cloud Volumes ONTAP

Terraform sensitive variable name	Description
on-prem-ontap_user_password	Admin password for the storage cluster

23. Click **Map**. This completes the task of adding the required Terraform sensitive variables to workspace. Next, start a new Terraform plan in the configured workspace.

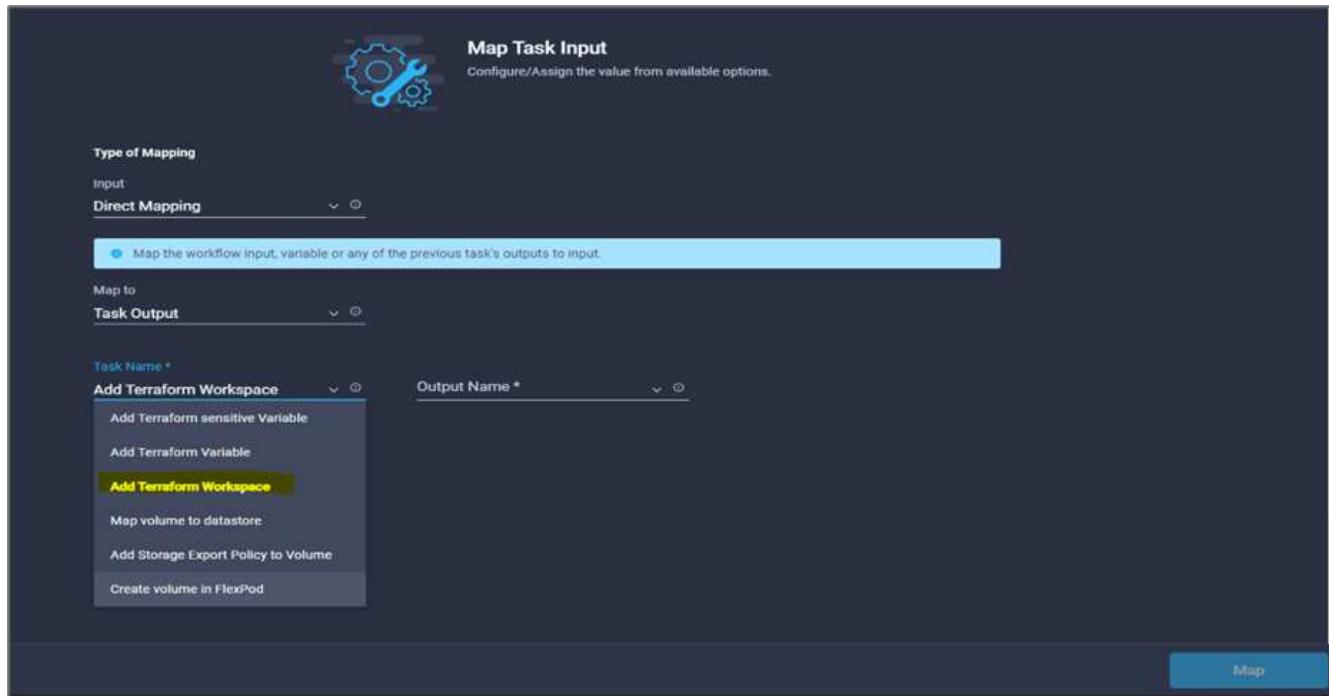
#### Procedure 8: Start a new Terraform plan

1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Terraform Cloud > Start New Terraform Plan** task from the **Tools** section on the **Design** area.
3. Use Connector to connect between the tasks **Add Terraform Sensitive Variables** and **Start New Terraform Plan tasks**. Click **Save**.
4. Click **Start New Terraform Plan**. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task.

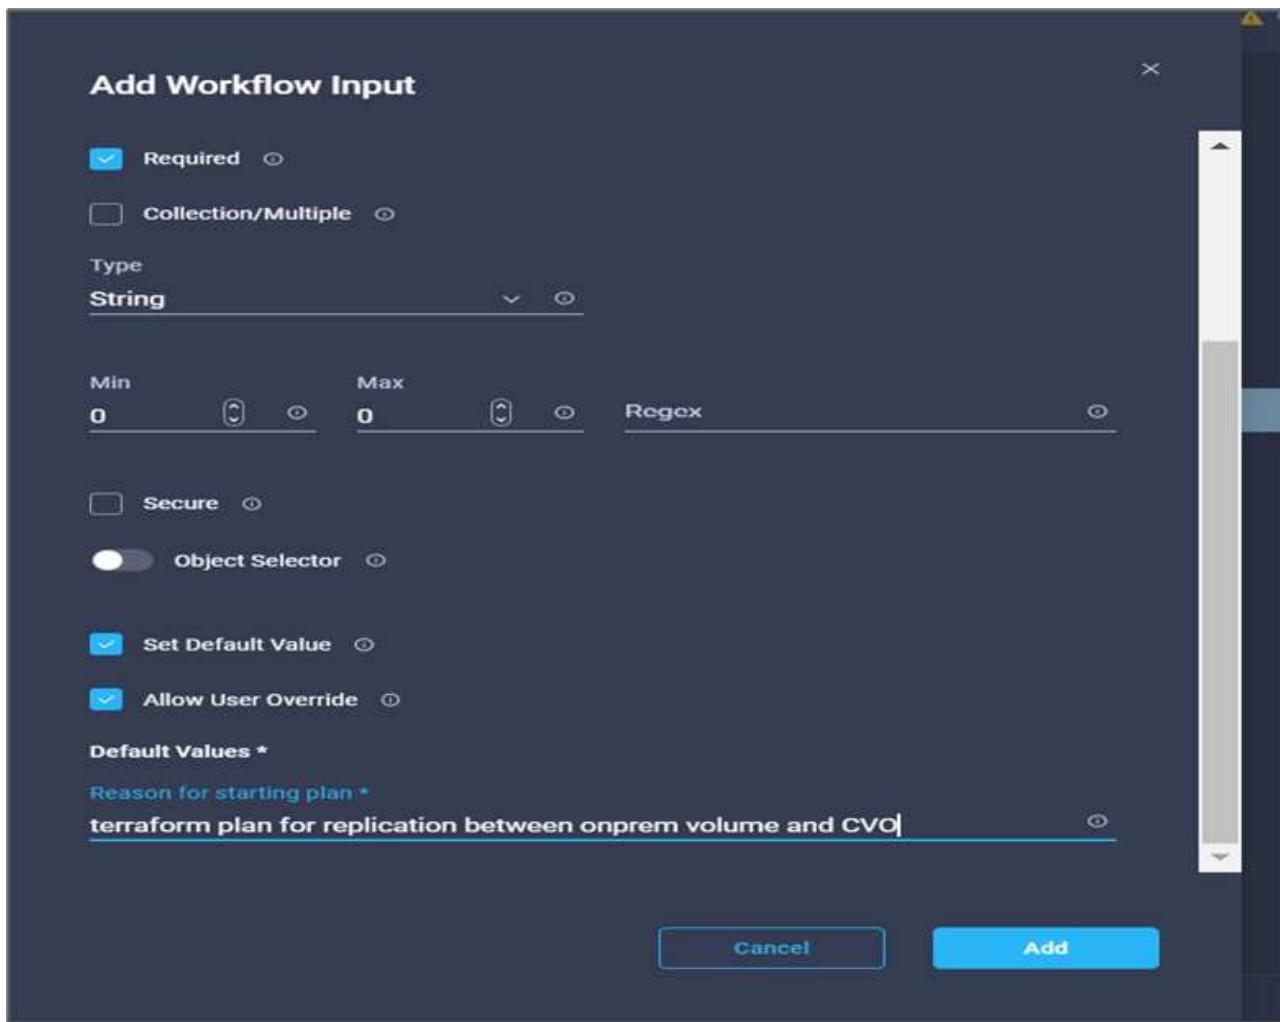


5. In the **Task Properties** area, click **Inputs**.
6. Click **Map** in the **Terraform Cloud Target** field.
7. Choose **Static Value** and click **Select Terraform Cloud Target**. Select the Terraform Cloud for Business account that was added in the section “Configuring Cisco Intersight Service for HashiCorp Terraform.”
8. Click **Map**.

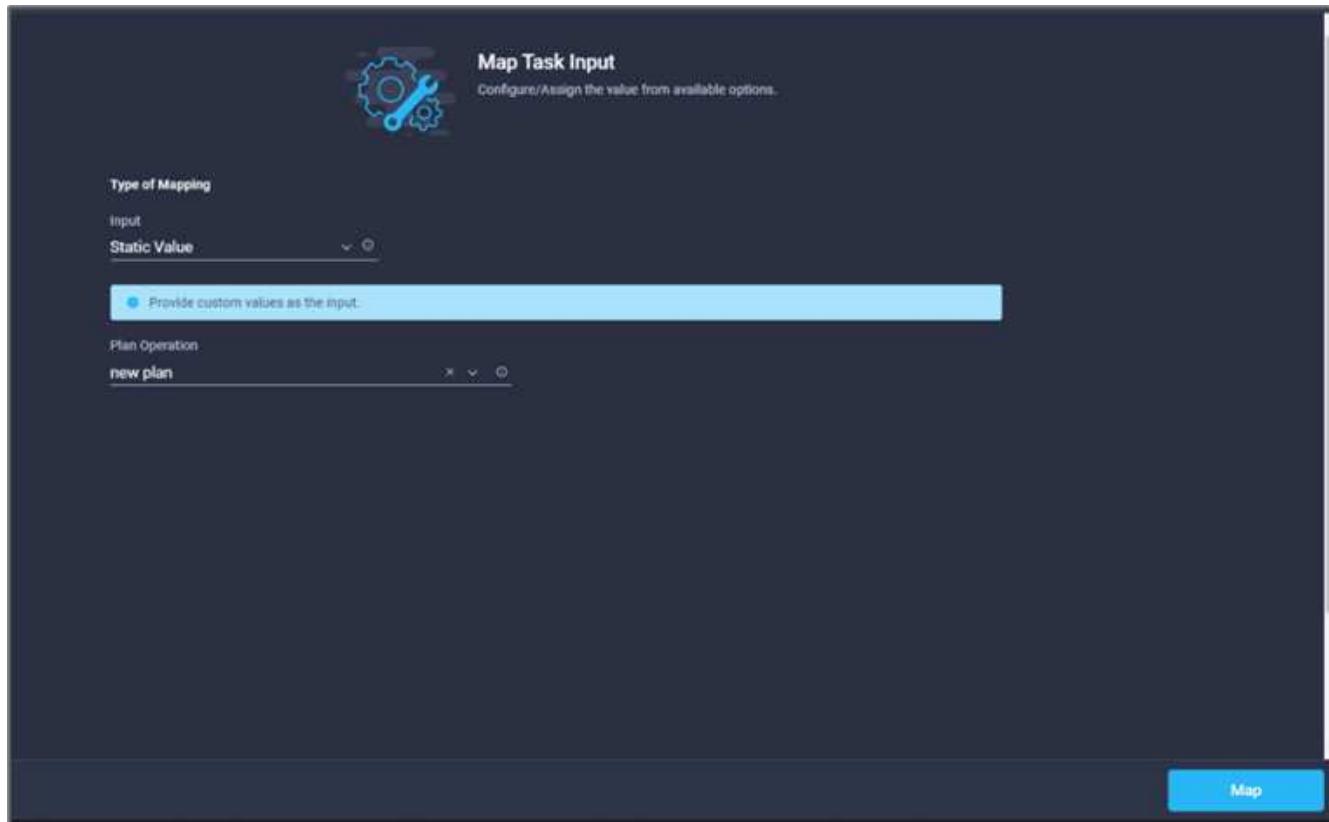
9. Click **Map** in the **Workspace ID** field.
10. Choose **Direct Mapping** and click **Task Output**.
11. Click **Task Name** and then click **Add Terraform Workspace**.



12. Click **Output Name**, **Workspace ID**, and then **Map**.
13. Click **Map** in the **Reason for starting plan** field.
14. Choose **Direct Mapping** and then click **Workflow Input**.
15. Click **Input Name** and then **Create Workflow Input**.
16. In the Add Input wizard, complete the following steps:
  - a. Provide a display name and reference name (optional).
  - b. Make sure to select **String** for the **Type**.
  - c. Click **Set Default Value and Override**.
  - d. Input a default value for **Reason for starting plan** and click **Add**.



17. Click **Map**.
18. Click **Map** in the **Plan Operation** field.
19. Choose **Static Value** and click **Plan Operation**. Click **new plan**.



20. Click **Map**.

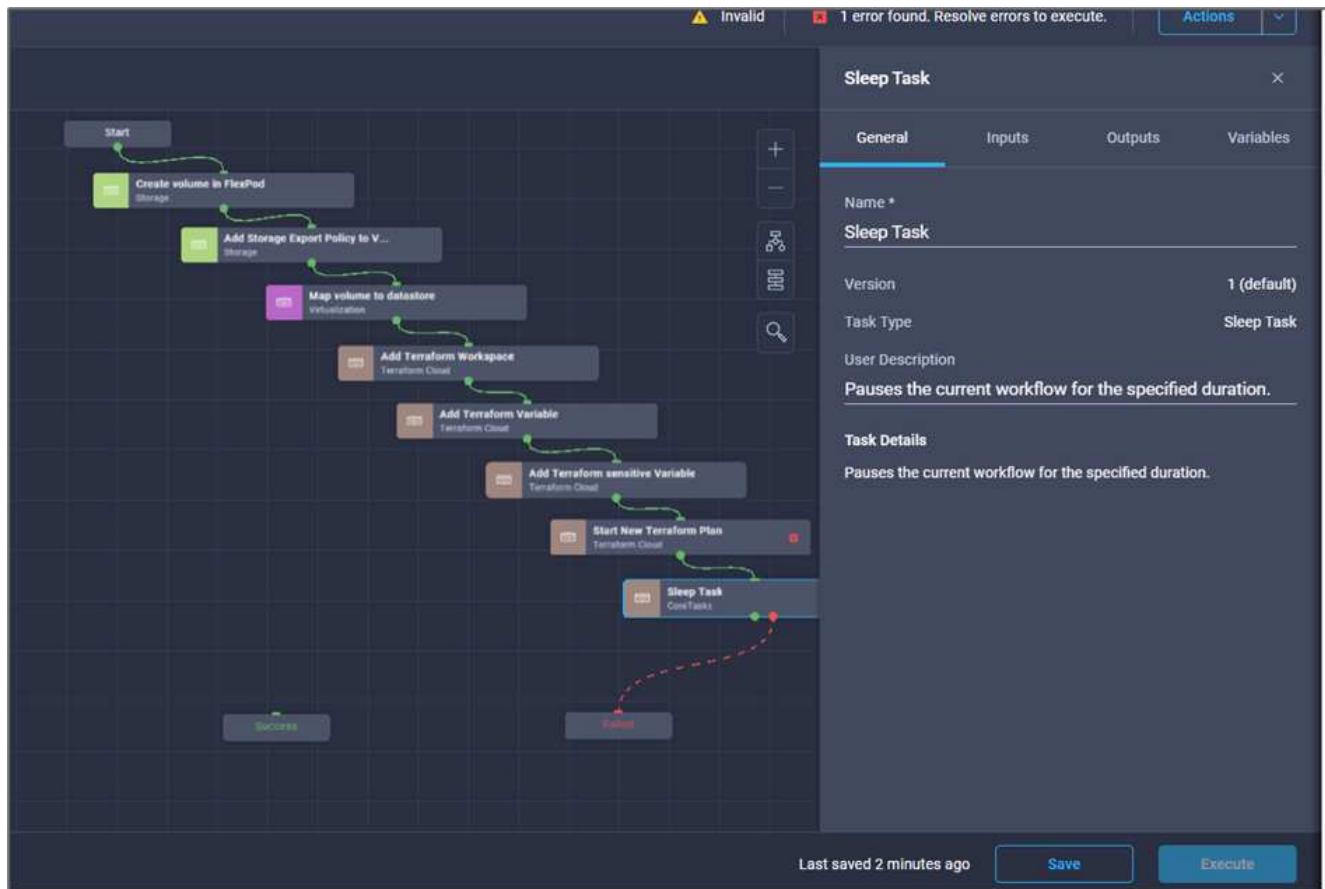
21. Click **Save**.

This completes the task of adding a Terraform Plan in Terraform Cloud for Business account. Next, create a sleep task for a few seconds.

#### Procedure 9: Sleep task for synchronization

Terraform Apply requires RunID, which is generated as a part of the Terraform Plan task. Waiting a few seconds between the Terraform Plan and Terraform Apply actions avoids timing issues.

1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Core Tasks > Sleep Task** from the **Tools** section in the **Design** area.
3. Use Connector to connect the tasks **Start New Terraform Plan** and **Sleep Task**. Click **Save**.



4. Click **Sleep Task**. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task. In this example, the name of the task is **Synchronize**.
5. In the **Task Properties** area, click **Inputs**.
6. Click **Map** in the **Sleep Time in Seconds** field.
7. Choose **Static Value** and input **15** in for the **Sleep Time in Seconds**.

Provide custom values as the input.

Sleep Time in Seconds \*

15

1 - 600

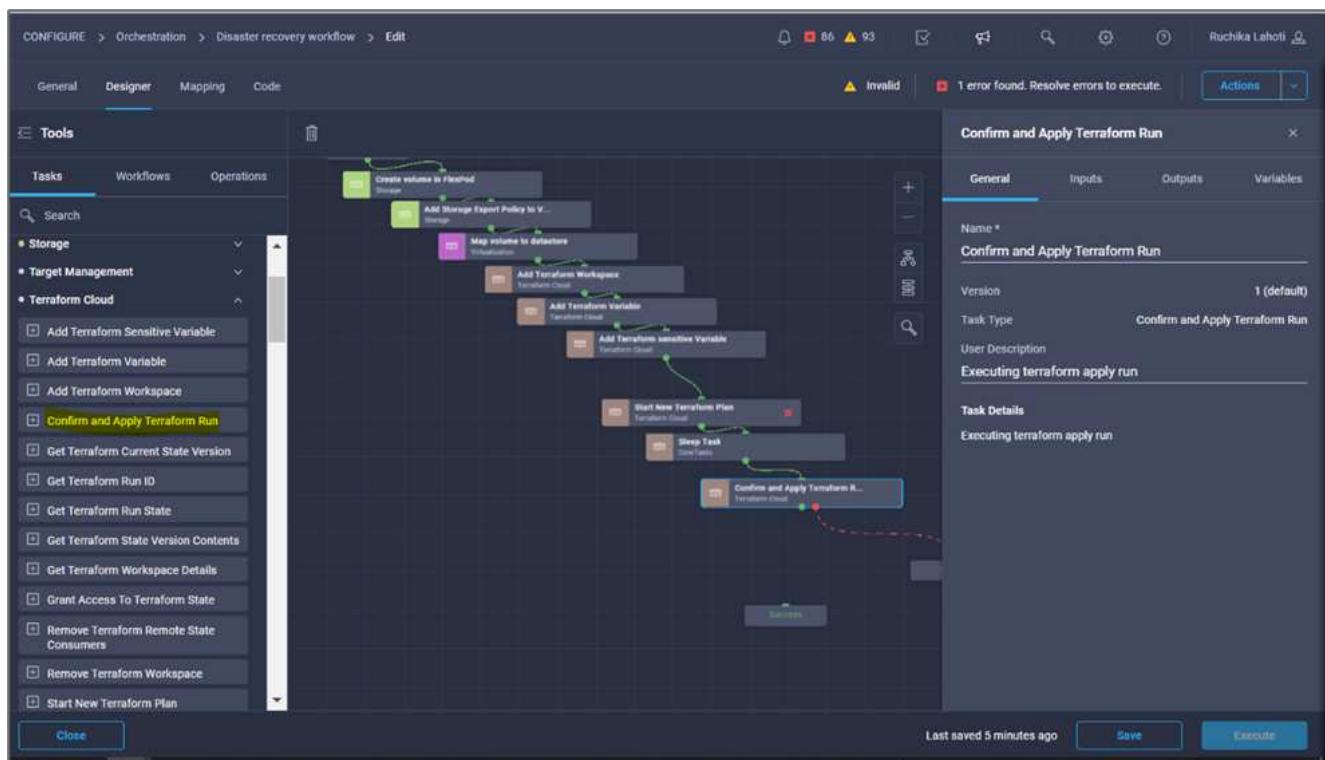
8. Click **Map**.

9. Click **Save**.

This completes the sleep task. Next, create the last task of this workflow, confirming and applying the Terraform Run.

#### Procedure 10: Confirm and apply Terraform Run

1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Terraform Cloud > Confirm and Apply Terraform Run** task from the **Tools** section in the **Design** area.
3. Use connector to connect the tasks **Synchronize** and **Confirm and Apply Terraform Run**. Click **Save**.
4. Click **Confirm and Apply Terraform Run**. In the **Task Properties** area, click the **General** tab. Optionally, you can change the name and description for this task.



5. In the **Task Properties** area, click **Inputs**.

6. Click **Map** in the **Terraform Cloud Target** field.

7. Choose **Static Value** and click **Select Terraform Cloud Target**. Select the Terraform Cloud for Business account that was added in [Configure Cisco Intersight Service for HashiCorp Terraform](#).

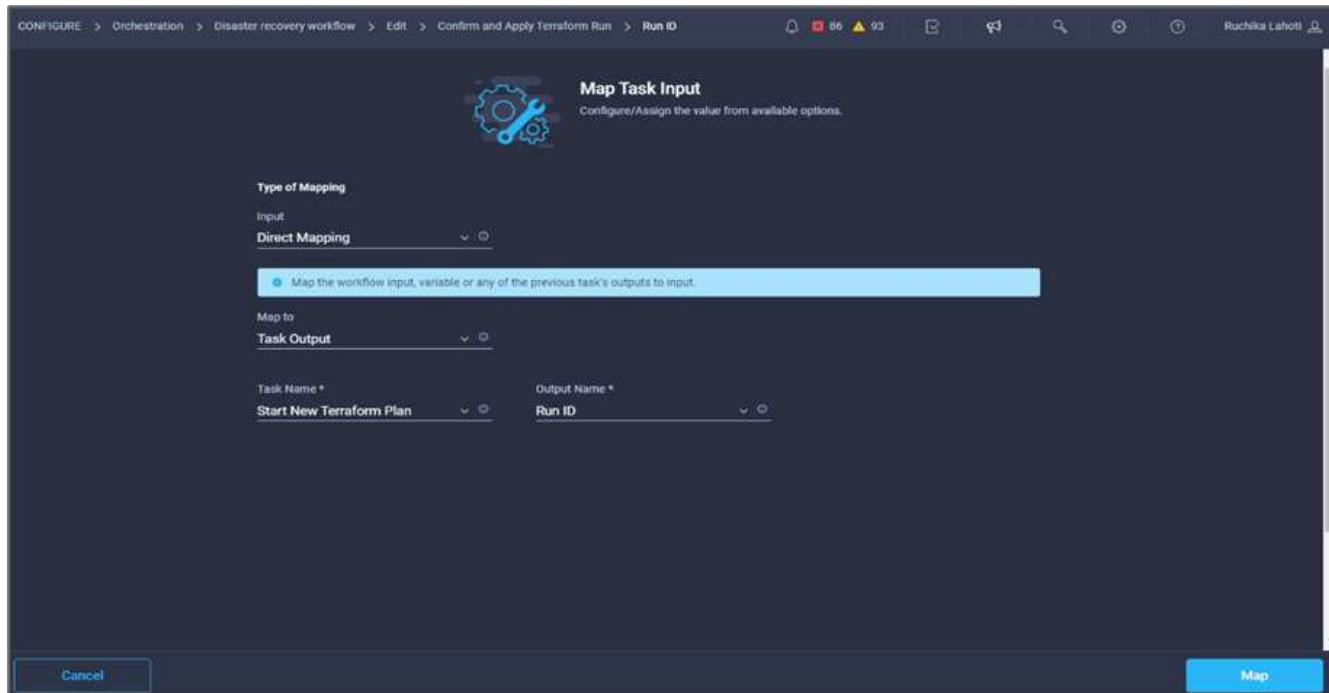
8. Click **Map**.

9. Click **Map** in the **Run ID** field.

10. Choose **Direct Mapping** and click **Task Output**.

11. Click **Task Name** and click **Start New Terraform Plan**.

12. Click **Output Name** and then click **Run ID**.



13. Click **Map**.
14. Click **Save**.
15. Click **Auto Align Workflow** so that all tasks are aligned. Click **Save**.



This completes the Confirm and Apply Terraform Run task. Use Connector to connect between the **Confirm and Apply Terraform Run** task and the **Success** and **Failed** tasks.

#### Procedure 11: Import a Cisco-built workflow

Cisco Intersight Cloud Orchestrator enables you to export workflows from a Cisco Intersight account to your system and then import them to another account. A JSON file was created by exporting the built workflow that can be imported to your account.

A JSON file for the workflow component is available in the [GitHub repository](#).

Next: Terraform execution from controller.

## Terraform execution from controller

Previous: DR workflow.

We can execute the Terraform plan using a controller. You can skip this section if you have already executed your Terraform plan using an ICO workflow.

### Prerequisites

Setup of the solution begins with a management workstation that has access to the Internet and with a working installation of Terraform.

A guide for installing Terraform can be found [here](#).

### Clone GitHub repo

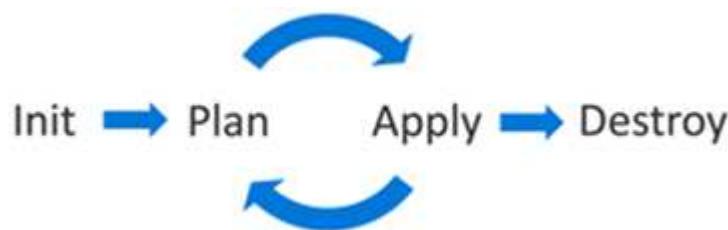
The first step in the process is to clone the GitHub repo to a new empty folder on the management workstation. To clone the GitHub repository, complete the following steps:

1. From the management workstation, create a new folder for the project. Create a new folder inside this folder named `/root/snapmirror-cvo` and Clone the GitHub repo into it.
2. Open a command-line or console interface on the management workstation and change directories to the new folder just created.
3. Clone the GitHub collection using the following command:

```
Git clone https://github.com/NetApp-Automation/FlexPod-hybrid-cloud-for-GCP-with-Intersight-and-CVO
```

1. Change directories to the new folder named `snapmirror-cvo`.

## Terraform execution



- **Init.** Initialize the (local) Terraform environment. Usually executed only once per session.
- **Plan.** Compare the terraform state with the as-in state in the cloud and build and display an execution plan. This does not change the deployment (read-only).
- **Apply.** Apply the plan from the plan phase. This potentially changes the deployment (read and write).
- **Destroy.** All resources that are governed by this specific terraform environment.

For details, see [here](#).

[Next: Solution validation.](#)

## Solution validation

[Previous: Terraform execution from controller.](#)

In this section, we revisit the solution with a sample data-replication workflow and take a few measurements to verify the integrity of data replication from the NetApp ONTAP instance running in FlexPod to NetApp Cloud Volumes ONTAP running on Google Cloud.

We used the Cisco Intersight workflow orchestrator in this solution and will continue to use this for our use case.

Notably, the limited set of Cisco Intersight workflows used in this solution do not represent the full set of workflows that Cisco Intersight is equipped with. You can create custom workflows based on your specific requirements and have them triggered from Cisco Intersight.

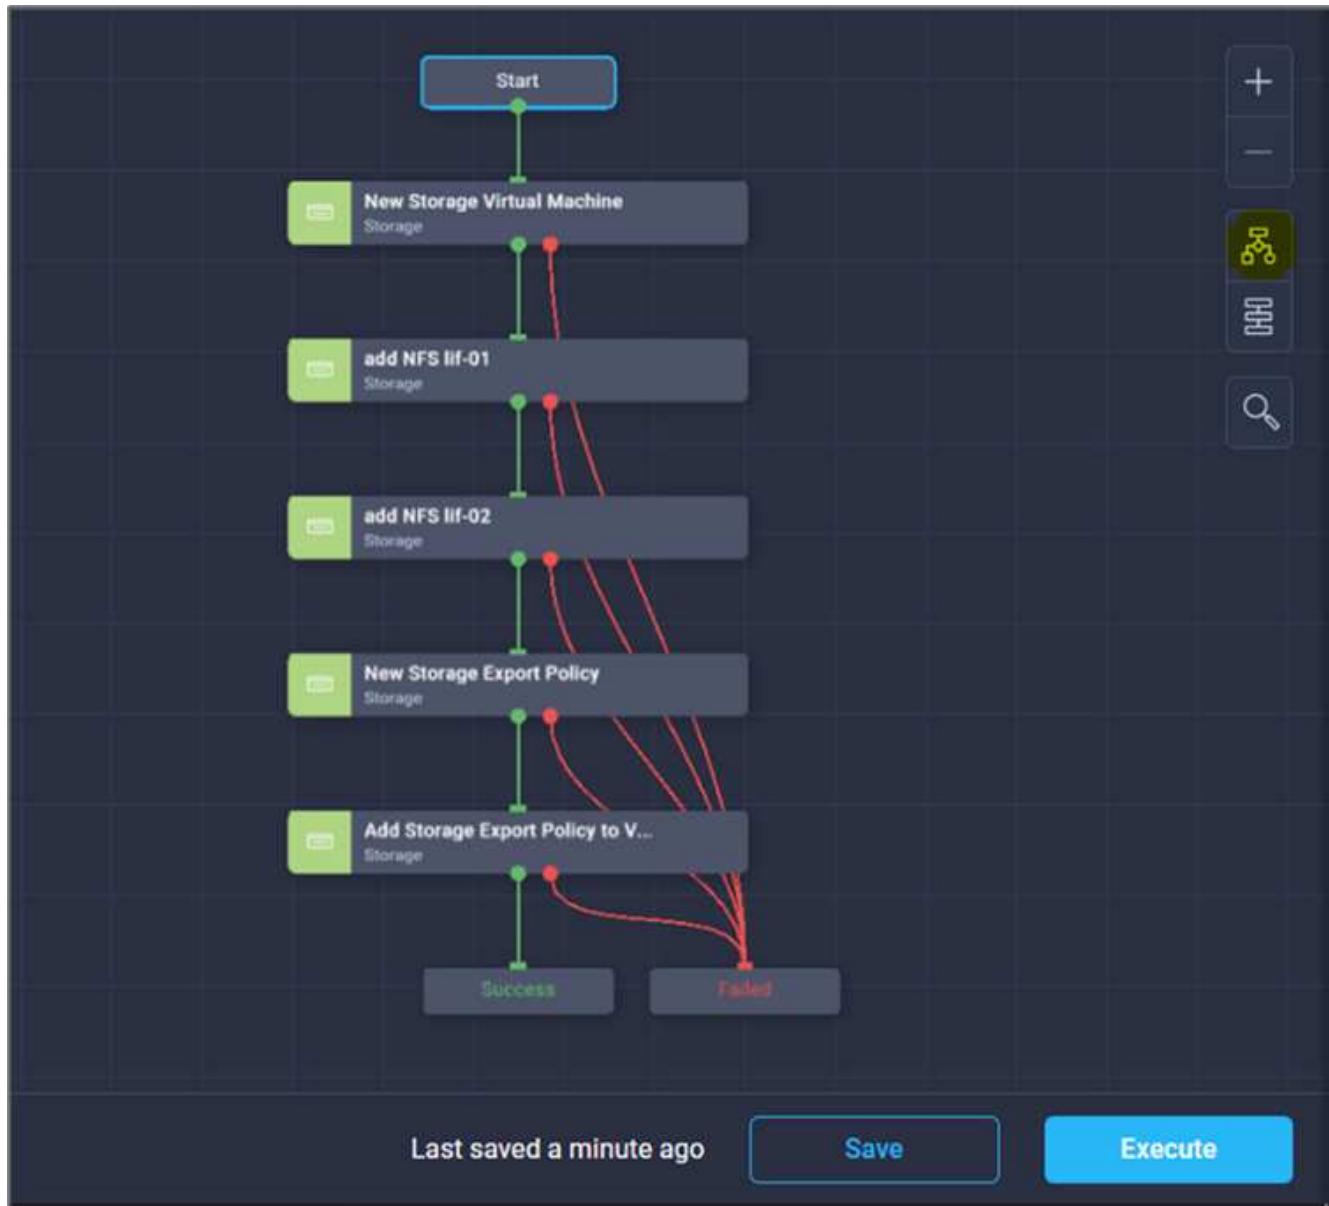
To perform the validation of a successful DR scenario, first move data from a volume in ONTAP that is part of FlexPod to Cloud Volumes ONTAP using SnapMirror. Then you can attempt to access the data from the Google cloud compute instance followed by a data integrity check.

The following high-level steps are used to verify the success criteria of this solution:

1. Generate an SHA256 checksum on the sample dataset that is present in an ONTAP volume in FlexPod.
2. Set up a volume SnapMirror relationship between ONTAP in FlexPod and Cloud Volumes ONTAP.
3. Replicate the sample dataset from FlexPod to Cloud Volumes ONTAP.
4. Break the SnapMirror relationship and promote the volume in Cloud Volumes ONTAP to production.
5. Map the Cloud Volumes ONTAP volume with the dataset to a compute instance in Google Cloud.
6. Generate an SHA256 checksum on the sample dataset in Cloud Volumes ONTAP.
7. Compare the checksum on the source and destination; presumably, the checksums on both sides match.

To execute the on-premises workflow, complete the following steps:

1. Create a workflow in Intersight for on-premises FlexPod.



2. Provide the required inputs and execute the workflow.

Execute Workflow: Configure-on-premFlexpodstorage

Execute Workflow

Fill Attributes

General

Organization \* default

Workflow Instance Name Configure on-prem FlexPod storage

Workflow Inputs

Storage Virtual Machine \* flexpod-svm

Storage Vendor Virtual Machine Options

Platform Type:  NetApp Active IQ Unified Manager  Pure FlashArray  Hitachi Virtual Storage Platform  None

NetApp Virtual Machine Options

Storage VM Protocols \* NFS

Storage VM Protocols \* iSCSI

Manage Administrator Account: vsadmin

Route Destination IPv4 Gateway 10.61.183.1

Execute

3. Verify the newly created SVM in the system manager.

☰ ONTAP System Manager

Search actions, objects, and pages

DASHBOARD

INSIGHTS

STORAGE

Overview

Volumes

LUNs

Consistency Groups

NVMe Namespaces

Shares

Qtrees

Quotas

Storage VMs

Tiers

Storage VMs

+ Add

More

<input type="checkbox"/> Name
<input checked="" type="checkbox"/> flexpod-svm
hybrid-cloud-svm
hybrid_cloud_2_svm
infra_svm
nvme1
terraform-demo-svm

flexpod-svm All Storage VMs

Overview Settings Snap

Security

Certificates

4. Create and execute another disaster recovery workflow to create a volume in on-prem FlexPod and establish a SnapMirror relationship between this volume in FlexPod and Cloud Volumes ONTAP.



5. Verify the newly created volume in ONTAP system manager.

The screenshot shows the ONTAP System Manager interface. The left sidebar is titled 'DASHBOARD' and includes sections for 'INSIGHTS', 'STORAGE' (with 'Overview', 'Volumes', 'LUNs', 'Consistency Groups', 'NVMe Namespaces', 'Shares', 'Qtrees', 'Quotas', 'Storage VMs', and 'Tier'), and 'Tiare'. The main content area is titled 'Volumes' and shows a table with columns: 'Name', 'Storage VM', 'Status', and 'Capacity'. The table lists several volumes, including 'application\_copy', 'audit\_log\_vol', 'hybrid\_cloud\_svm\_root', 'test', and 'Test\_Vol1' (which is highlighted with a yellow box). Each volume row includes a checkbox, a search icon, and a dropdown for status.

6. Mount the same NFS volume to an on-premises virtual machine, then copy the sample dataset and perform the checksum.

```
root@hybridcloudbackup:/snapmirror_demo# mount -t nfs 172.22.4.157:/Test_Vol1 /snapmirror_demo
root@hybridcloudbackup:/snapmirror_demo# df -kh
Filesystem      Size  Used  Avail Use% Mounted on
udev            1.9G   0     1.9G  0% /dev
tmpfs           394M  1.1M  393M  1% /run
/dev/sda2        16G   11G   4.2G  72% /
tmpfs           2.0G   0     2.0G  0% /dev/shm
tmpfs           5.0M   0     5.0M  0% /run/lock
tmpfs           2.0G   0     2.0G  0% /sys/fs/cgroup
/dev/loop1        55M   55M   0    100% /snap/core18/1705
/dev/loop2        69M   69M   0    100% /snap/lxd/14804
/dev/loop0        28M   28M   0    100% /snap/snapd/7264
172.22.4.157:/Test_Vol1  10G  512K  10G  1% /snapmirror_demo
root@hybridcloudbackup:/snapmirror_demo#
```

```
root@hybridcloudbackup:/snapmirror_demo#
root@hybridcloudbackup:/snapmirror_demo# sha256sum test.zip
888a23c8495ad33fdf11a931ffc344c3643f15d5cefdbbf1326016e31ec5a59  test.zip
root@hybridcloudbackup:/snapmirror_demo#
root@hybridcloudbackup:/snapmirror_demo#
```

7. Check the replication status in Cloud Manager. The data transfer can take few minutes based on the size of the data. After it is completed, you can see the SnapMirror status as **Idle**.

NetApp Cloud Manager Replication dashboard summary:

- 1 Volume Relationship
- 977.46 MiB Replicated Capacity
- 0 Currently Transferring
- 1 Healthy
- 0 Failed

1 Volume Relationship table:

Health Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Created
Green checkmark	Test_Vol1 a220-g1316	dr_dest_volume_on_gcp gpcvodemo	34 seconds	Idle	snapmirrored	Aug 24, 2022 989.15 MiB

- When the data transfer is complete, simulate a disaster on the source side by stopping the SVM that hosts the `Test_voll` volume.

After the SVM has been stopped, the `Test_voll` volume is not visible in the Cloud Manager.

ONTAP System Manager Storage VMs list:

Name	State	Subtype	Configured Protocols	IPspace	Protection
flexpod-svm	running	default	NFS, iSCSI	Default	🛡️
flexpod-svm	stopped	default		Default	🛡️
flexpod-svm	running	default	NFS, iSCSI	Default	🛡️
flexpod-svm	running	default	NFS, iSCSI, FC, S3	Default	🛡️
flexpod-svm	running	default	NVMe	Default	🛡️
flexpod-svm	running	default	iSCSI	Default	🛡️

Context menu for the stopped SVM:

- Edit
- Delete
- Start (highlighted)
- Trace File Access
- Login Banner Message

- Break the replication relationship and promote the Cloud Volumes ONTAP destination volume to production.

NetApp Cloud Manager - Replication

Account: hybrid\_cloud | Workspace: FlexPod | Connector: hybridcloud-co...

1 Volume Relationship

Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Success
Test_Vol1 a220-g1316	dr_dest_volume_on_gcp gpcvcdemo	34 seconds	idle	snapmirrored	Aug 24, 2022, 3:30:29 PM 989.15 MiB

Information

- Break (highlighted)
- Reverse Resync
- Edit Schedule
- Edit Max Transfer Rate
- Update
- Delete

1 Volume Relationship

Status	Source Volume	Target Volume	Total Transfer Time	Status	Mirror State	Last Success
broken-off	Test_Vol1 a220-g1316	dr_dest_volume_on_gcp gpcvcdemo	34 seconds	idle	broken-off	Aug 24, 2022, 3:30:29 PM 989.15 MiB

10. Edit the volume and enable client access by associating it with an export policy.

Edit volume dr\_dest\_volume\_on\_gcp

Protocol: NFS

Access control:

No access to the volume

Custom export policy

172.30.116.0/22

Advanced options

Protection

Snapshot Policy:

none

Update Cancel

11. Obtain the ready-to-use mount command for the volume.

The screenshot shows the NetApp Cloud Manager interface. In the top navigation bar, 'NetApp Cloud Manager' is selected. The top right shows 'Account: hybrid\_cloud', 'Workspace: FlexPod', and 'Connector: hybridcloud-co...'. There are also icons for notifications, settings, and help. The main area shows a project named 'gcpcvodemo'. Under 'Volumes', there are two volumes: 'dr\_dest\_volume\_on\_gcp' and 'test\_cvo\_volume'. The 'dr\_dest\_volume\_on\_gcp' volume has a context menu open with options: 'Clone', 'Mount Command' (which is highlighted in yellow), 'Restore from Snapshot copy', and 'Create a Snapshot copy'. The 'test\_cvo\_volume' volume is listed with 'ONLINE' status, 'PD-BALANCED' disk type, 'None' tiering policy, and 1 GB Allocated, 328 KB Disk Used.

## Mount Volume dr\_dest\_volume\_on\_gcp

Go to your Linux machine and enter this mount command

```
mount 172.30.116.153:/dr_dest_volume_on_gcp <dest...>
```

Copy

12. Mount the volume to a compute instance, verify that the data is present in the destination volume, and generate the SHA256 checksum of the sample\_dataset\_2GB file.

```
drwxr-xr-x 21 root      root          4096 Aug 24 10:20 ../
-rwxr-xr-x  1 nobody    4294967294 1015306240 Aug 24 09:59 test.zip*
ruchikal_netapp_com@demo-nfs:/snapmirror_dest$ sha256sum test.zip
888a23c8495ad33fdf11a931ffc344c3643f15d5cefdbbf1326016e31ec5a59  test.zip
ruchikal_netapp_com@demo-nfs:/snapmirror_dest$
```

13. Compare the checksum values at both the source (FlexPod) and the destination (Cloud Volumes ONTAP).
14. The checksums match to the source and destination.

You can confirm that the data replication from the source to the destination was completed successfully and the data integrity was maintained. This data can now be safely consumed by the applications to serve clients while the source site goes through restoration.

[Next: Conclusion.](#)

## Conclusion

Previous: [Solution validation](#).

In this solution, the NetApp Cloud Data service, Cloud Volumes ONTAP, and FlexPod Datacenter infrastructure were used to build a DR solution with a public cloud powered by the Cisco Intersight Cloud Orchestrator. The FlexPod solution has constantly evolved to enable customers to modernize their applications and business-delivery processes. With this solution, you can build a BCDR plan with the public cloud as your go-to location for a transient or full-time DR plan while keeping the cost of the DR solution low.

Data replication between on-premises FlexPod and NetApp Cloud Volumes ONTAP was handled by proven SnapMirror technology, but you can also select other NetApp data- transfer and synchronization tools like Cloud Sync for your data mobility requirements. Security of the data in-flight provided by built-in encryption technologies based on TLS/AES.

Whether you have a temporary DR plan for an application or a full-time DR plan for a business, the portfolio of products used in this solution can meet both requirements at scale. Powered by Cisco Intersight Workflow Orchestrator, the same can be automated with prebuilt workflows that not just eliminate the need to rebuild processes but also accelerate the implementation of a BCDR plan.

The solution enables the management of FlexPod on-premises and data replication across a hybrid cloud in a very easy and convenient manner with automation and orchestration provided by Cisco Intersight Cloud Orchestrator.

### Where to find additional information

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

#### GitHub

- All Terraform Configurations used

<https://github.com/NetApp-Automation/FlexPod-hybrid-cloud-for-GCP-with-Intersight-and-CVO>

- JSON files for importing workflows

[https://github.com/ucs-compute-solutions/FlexPod\\_DR\\_Workflows](https://github.com/ucs-compute-solutions/FlexPod_DR_Workflows)

#### Cisco Intersight

- Cisco Intersight Help Center

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

- Cisco Intersight Cloud Orchestrator Documentation:

[https://intersight.com/help/saas/features/orchestration/configure#intersight\\_cloud\\_orchestrator](https://intersight.com/help/saas/features/orchestration/configure#intersight_cloud_orchestrator)

- Cisco Intersight Service for HashiCorp Terraform Documentation

[https://intersight.com/help/saas/features/terraform\\_cloud/admin](https://intersight.com/help/saas/features/terraform_cloud/admin)

- Cisco Intersight Data Sheet

<https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/intersight/intersight-ds.html>

- Cisco Intersight Cloud Orchestrator Data Sheet

<https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/intersight/nb-06-intersight-cloud-orch-aag-cte-en.html>

- Cisco Intersight Service for HashiCorp Terraform Data Sheet

<https://www.cisco.com/c/en/us/products/collateral/cloud-systems-management/intersight/nb-06-intersight-terraf-ser-aag-cte-en.html>

## FlexPod

- FlexPod Home Page

<https://www.flexpod.com>

- Cisco Validated Design and deployment guides for FlexPod

[FlexPod Datacenter with Cisco UCS 4.2\(1\) in UCS Managed Mode, VMware vSphere 7.0 U2, and NetApp ONTAP 9.9 Design Guide](#)

- FlexPod Datacenter with Cisco UCS X-Series

[https://www.cisco.com/c/en/us/td/docs/unified\\_computing/ucs/UCS\\_CVDs/flexpod\\_xseries\\_esxi7u2\\_design.html](https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_xseries_esxi7u2_design.html)

## Interoperability

- NetApp Interoperability Matrix Tool

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

- Cisco UCS Hardware and Software Interoperability Tool

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

- VMware Compatibility Guide

<http://www.vmware.com/resources/compatibility/search.php>

## NetApp Cloud Volumes ONTAP reference documents

- NetApp Cloud Manager

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

- Cloud Volumes ONTAP

<https://docs.netapp.com/us-en/cloud-manager-cloud-volumes-ontap/task-getting-started-gcp.html>

- Cloud Volumes ONTAP TCO Calculator  
<https://cloud.netapp.com/google-cloud-calculator>
- Cloud Volumes ONTAP Sizer  
<https://cloud.netapp.com/cvo-sizer>
- Cloud Assessment Tool  
<https://cloud.netapp.com/assessments>
- NetApp Hybrid Cloud  
<https://cloud.netapp.com/hybrid-cloud>
- Cloud Manager API documentation  
[https://docs.netapp.com/us-en/occm/reference\\_infrastructure\\_as\\_code.html](https://docs.netapp.com/us-en/occm/reference_infrastructure_as_code.html)

#### **Troubleshooting issues**

[https://kb.netapp.com/Advice\\_and\\_Troubleshooting/Cloud\\_Services/Cloud\\_Volumes\\_ONTAP\\_\(CVO\)](https://kb.netapp.com/Advice_and_Troubleshooting/Cloud_Services/Cloud_Volumes_ONTAP_(CVO))

#### **Terraform**

- Terraform Cloud  
<https://www.terraform.io/cloud>
- Terraform Documentation  
<https://www.terraform.io/docs/>
- NetApp Cloud Manager Registry  
<https://registry.terraform.io/providers/NetApp/netapp-cloudmanager/latest>

#### **GCP**

- ONTAP High Availability for GCP  
<https://cloud.netapp.com/blog/gcp-cvo-blg-what-makes-cloud-volumes-ontap-high-availability-for-gcp-tick>
- GCP perquisite  
<https://netapp.hosted.panopto.com/Panopto/Pages/Viewer.aspx?id=f3d0368b-7165-4d43-a76e-ae01011853d6>

## **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

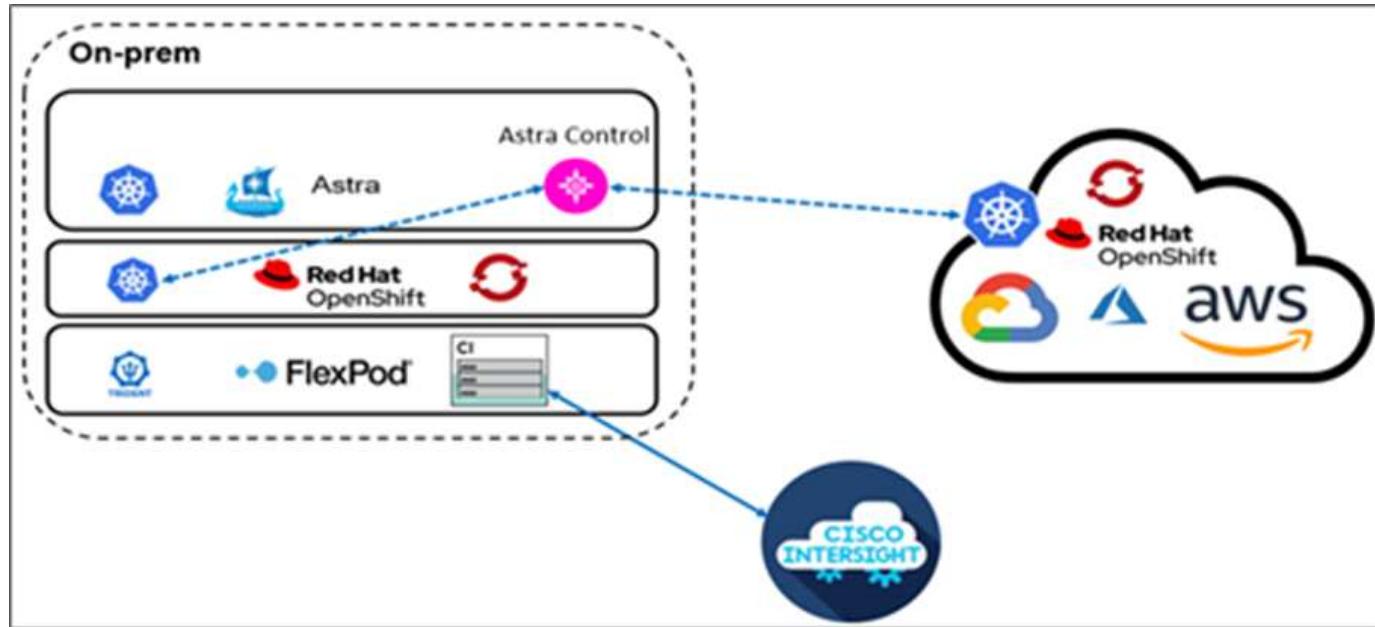
Abhinav Singh

## 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

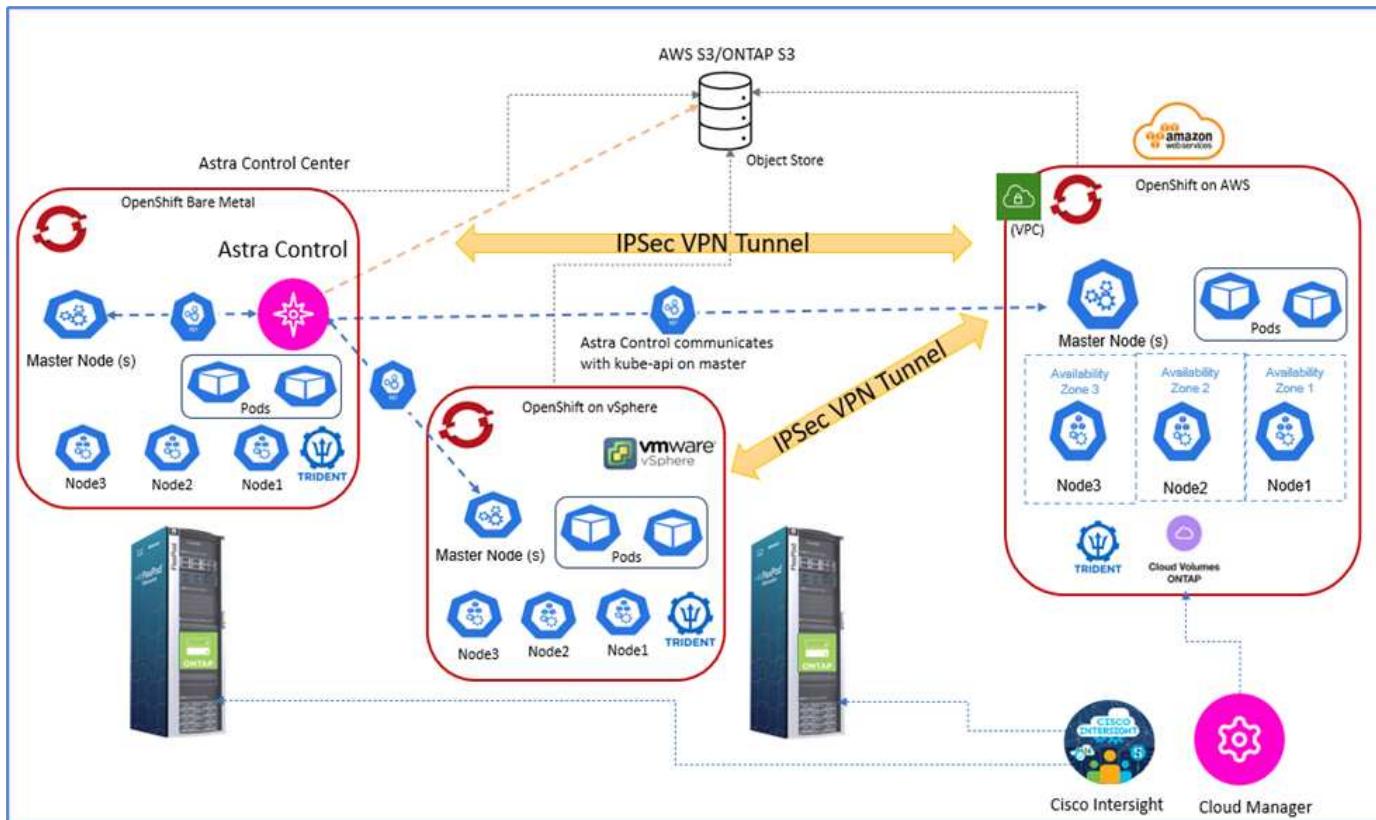
NetApp Astra Control aims at simplifying application protection for customers who deal with cloud native microservices:

- **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.



The first OpenShift Container Platform cluster is a bare-metal installation on FlexPod, the second OpenShift Container Platform cluster is deployed on VMware vSphere running on FlexPod, and the third OpenShift Container Platform cluster is deployed as a [private cluster](#) into an existing virtual private cloud (VPC) on AWS as a self-managed infrastructure.

In this solution, FlexPod is connected to AWS through a site-to-site VPN, however, customers can also use the direct connect implementations to extend to a hybrid cloud. Cisco Intersight is used to manage the FlexPod infrastructure components.

In this solution, Astra Control Center manages the containerized application hosted on the OpenShift Container Platform cluster running on FlexPod and on AWS. Astra Control Center is installed on the OpenShift bare-metal instance running on FlexPod. Astra Control communicates with the kube-api on the master node and continually watches the Kubernetes cluster for changes. Any new applications added to the K8s cluster are automatically discovered and made available for management.

PiT representations of containerized applications can be captured as snapshots using Astra Control Center. Application snapshots can be triggered through a scheduled protection policy or on demand. For applications that Astra supports, the snapshot is crash consistent. An application snapshot constitutes a snapshot of the application data in the persistent volumes as well as the application metadata of the various Kubernetes resources associated with that application.

A full copy backup of an application can be created by using Astra Control using a predefined backup schedule or on demand. An object storage is used to store the backup of the application data. NetApp ONTAP S3, NetApp StorageGRID, and any generic S3 implementation can be used as an object store.

[Next: Solution components.](#)

## 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.

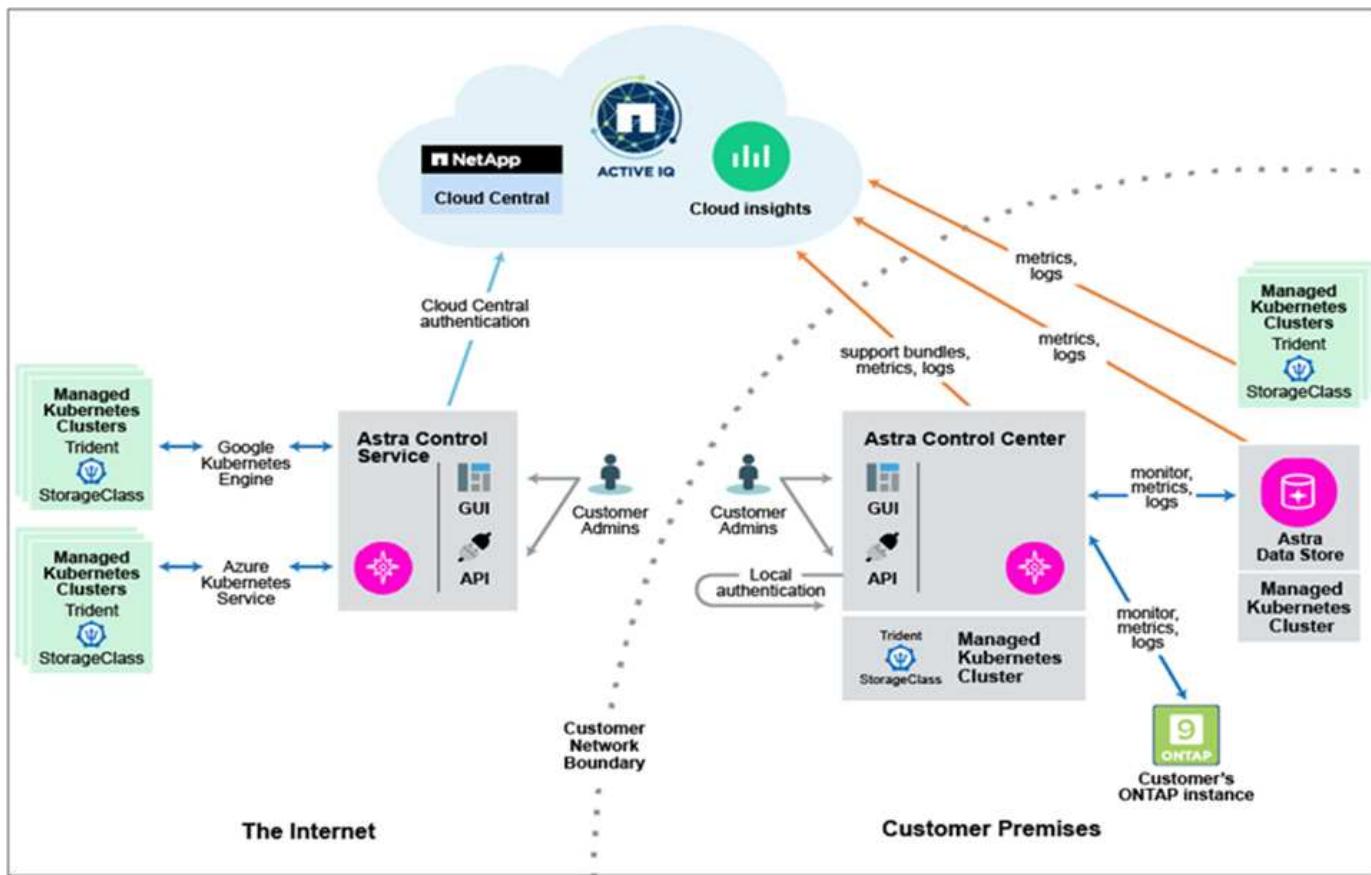
#### Astra Control Consumption models

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. 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.

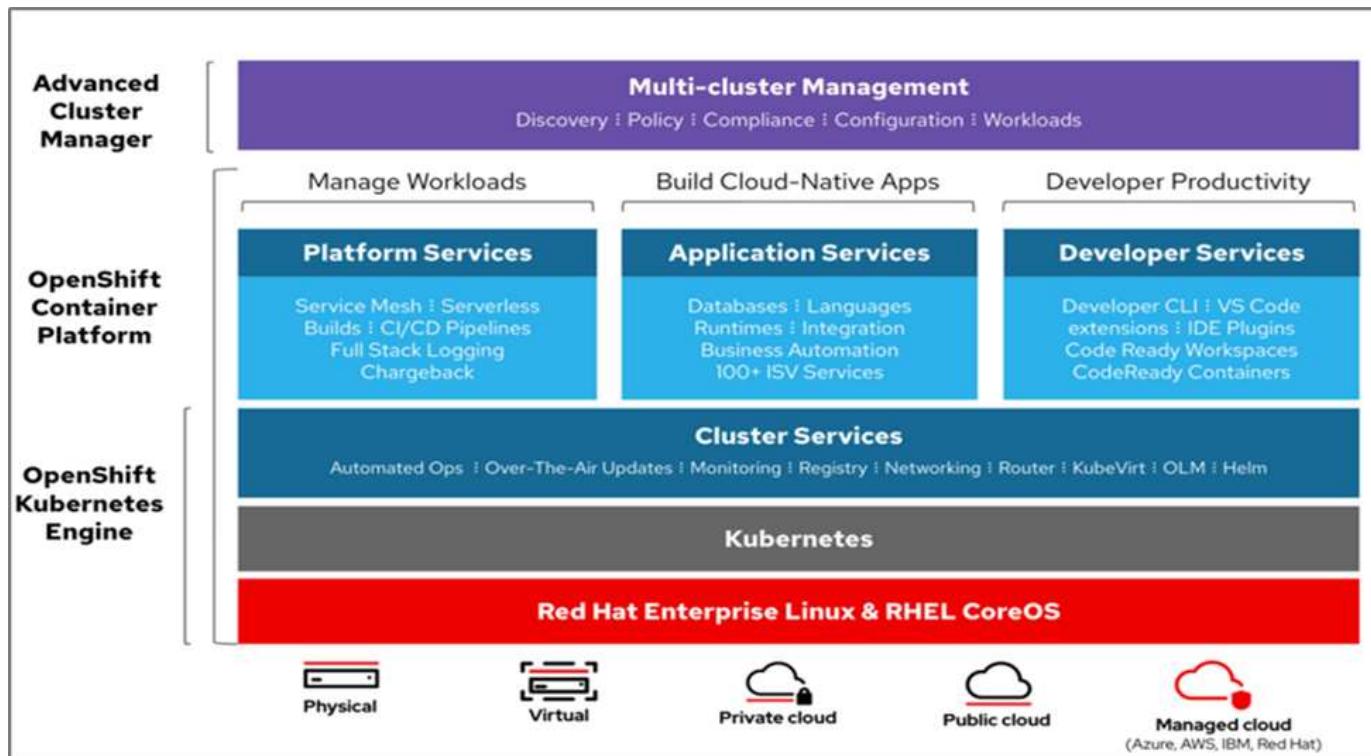
Name	Health	Contract Status	Management IP	Model
fp-g9-1-5	Healthy	Not Covered	172.18.9.202, 172.18.9.210	UCSB-B200-M5
fp-g9-1-4	Healthy	Not Covered	172.18.9.203, 172.18.9.209	UCSB-B200-M5
fp-g9-1-6	Healthy	Not Covered	172.18.9.201, 172.18.9.211	UCSB-B200-M5
openshift-ucs-1-2	Healthy	Not Covered	172.22.6.235, 172.22.6.241	UCSB-B200-M5
openshift-ucs-1-5	Healthy	Not Covered	172.22.6.239, 172.22.6.245	UCSB-B200-M5
openshift-ucs-1-6	Healthy	Not Covered	172.22.6.236, 172.22.6.240	UCSB-B200-M5
openshift-ucs-1-4	Healthy	Not Covered	172.22.6.238, 172.22.6.244	UCSB-B200-M5
openshift-ucs-1-3	Healthy	Not Covered	172.22.6.237, 172.22.6.243	UCSB-B200-M5
openshift-ucs-1-1	Healthy	Not Covered	172.22.6.234, 172.22.6.242	UCSB-B200-M5

## 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
	VMware ESXi nenic Ethernet Driver	1.0.35.0
Software	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.](#)

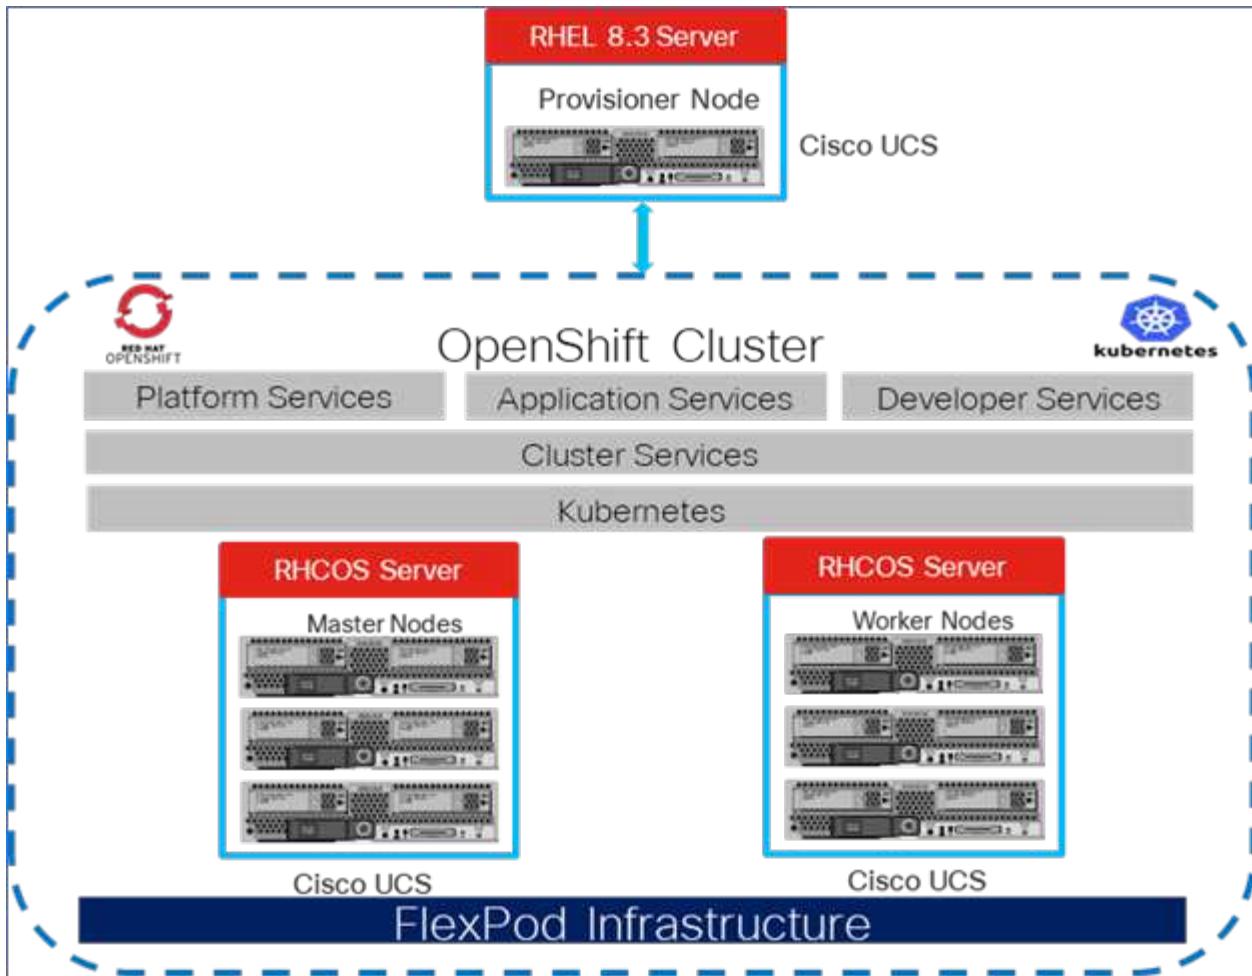
## Installation and configuration

### [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.

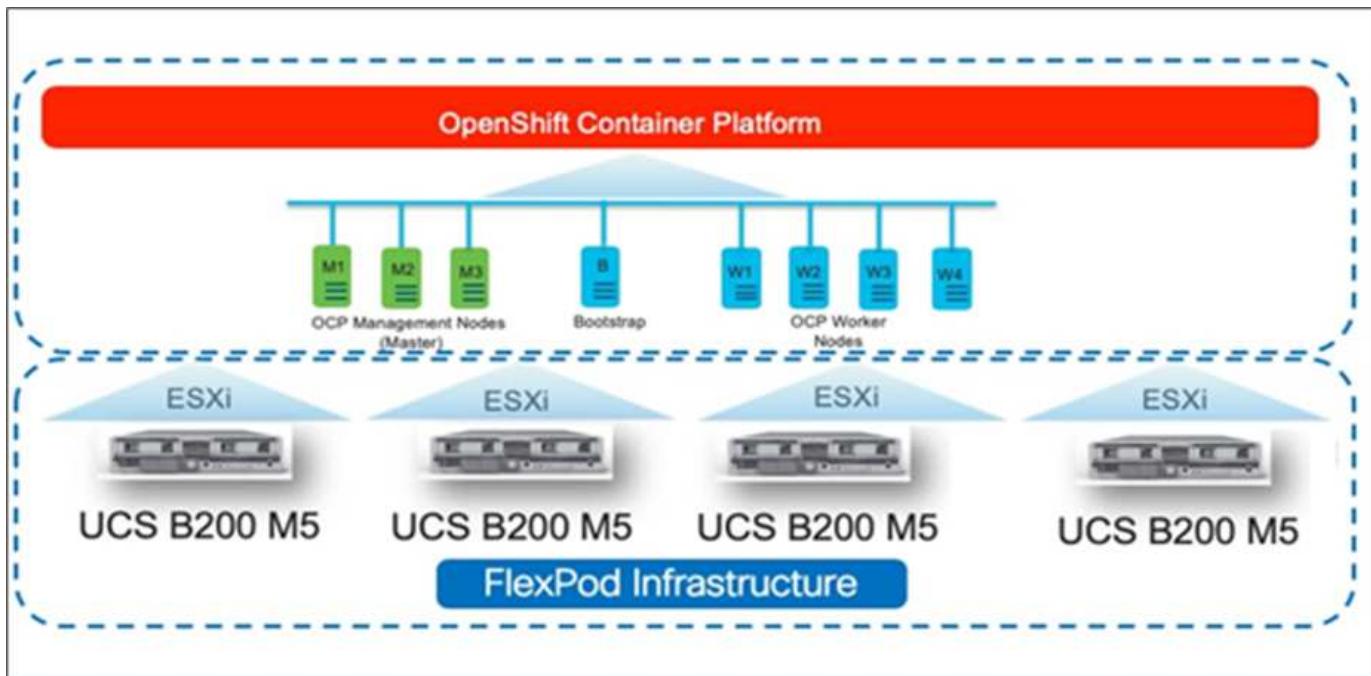
The following figure illustrates the OpenShift Container Platform 4 Bare Metal on FlexPod.



#### 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) <a href="#">Info</a>												
<input type="checkbox"/> <a href="#">Search</a>		<a href="#">C</a>										
<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>	<span>🕒🕒</span>	m5.xlarge		us-east-1b		172.30.165.160		-
<input type="checkbox"/>	ocpaws-v58kn-master-1		i-0b161945421d2a23c	<span>Running</span>	<span>🕒🕒</span>	m5.xlarge		us-east-1c		172.30.166.162		-
<input type="checkbox"/>	ocpaws-v58kn-master-2		i-0146a665e1060ea59	<span>Running</span>	<span>🕒🕒</span>	m5.xlarge		us-east-1a		172.30.164.209		-
<input type="checkbox"/>	ocpaws-v58kn-worker-us-east-1a-z8dj		i-05e6efa18d136c842	<span>Running</span>	<span>🕒🕒</span>	m5.large		us-east-1a		172.30.164.128		-
<input type="checkbox"/>	ocpaws-v58kn-worker-us-east-1b-7nmhc		i-0879a088b50d2d966	<span>Running</span>	<span>🕒🕒</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>	<span>🕒🕒</span>	m5.large		us-east-1c		172.30.166.51		-

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  
-signature <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.

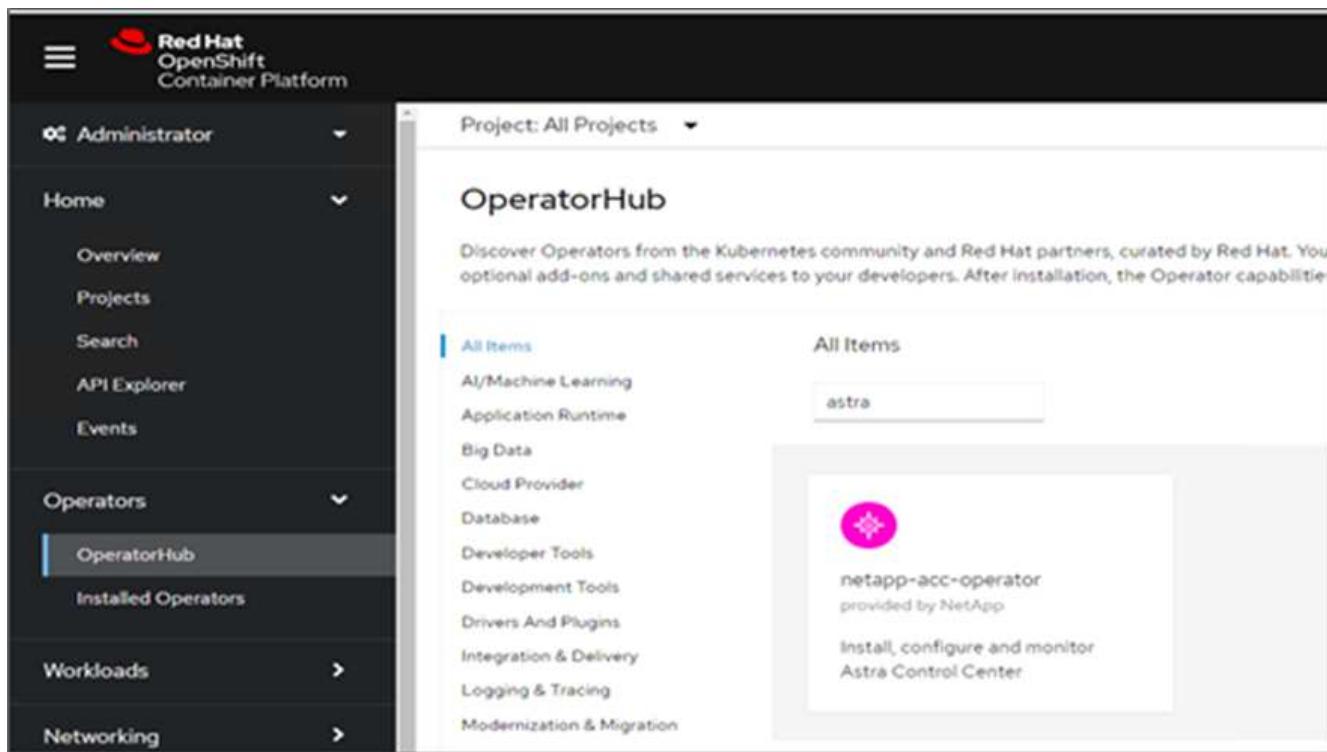
For Docker:

```
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
```

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

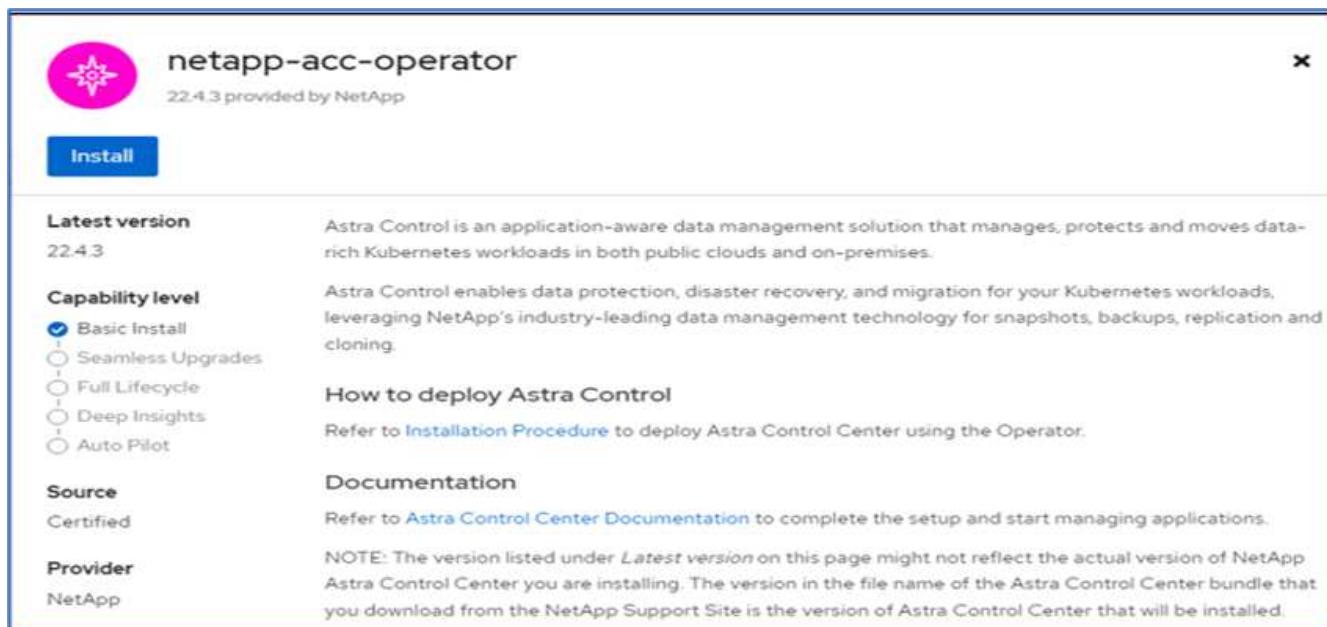


The screenshot shows the Red Hat OpenShift Container Platform web console. The left sidebar has a 'Operators' section with 'OperatorHub' selected. The main content area is titled 'OperatorHub' and shows a search bar with 'astra'. A search result for 'netapp-acc-operator' is displayed, categorized under 'Database'. The result includes a purple icon, the name 'netapp-acc-operator', 'provided by NetApp', and a description: 'Install, configure and monitor Astra Control Center'.



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

9. Select netapp-acc-operator and click Install.



The screenshot shows the 'netapp-acc-operator' details page in the OperatorHub. It includes a summary section with the operator name, version (22.4.3), and provider (NetApp). Below this are sections for 'Latest version' (22.4.3), 'Capability level' (with 'Basic Install' selected), 'How to deploy Astra Control' (link to Installation Procedure), 'Documentation' (link to Astra Control Center Documentation), and a note about version discrepancies. A large 'Install' button is prominently displayed.

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 \***

 **netapp-acc-operator** (Operator recommended)

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

**Update approval \*** 

Automatic  
 Manual

**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.

**Install** **Cancel**

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

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

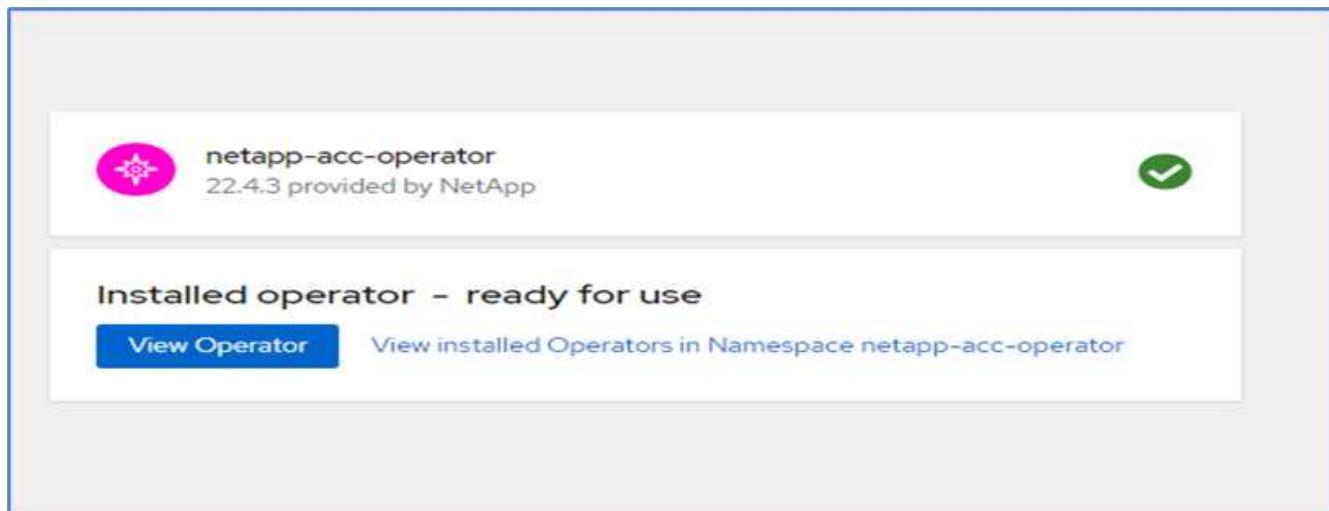
**Manual approval required**

Review the **manual install plan** for operators **acc-operator.v22.4.3**. Once approved, the following resources will be created in order to satisfy the requirements for the components specified in the plan. Click the resource name to view the resource in detail.

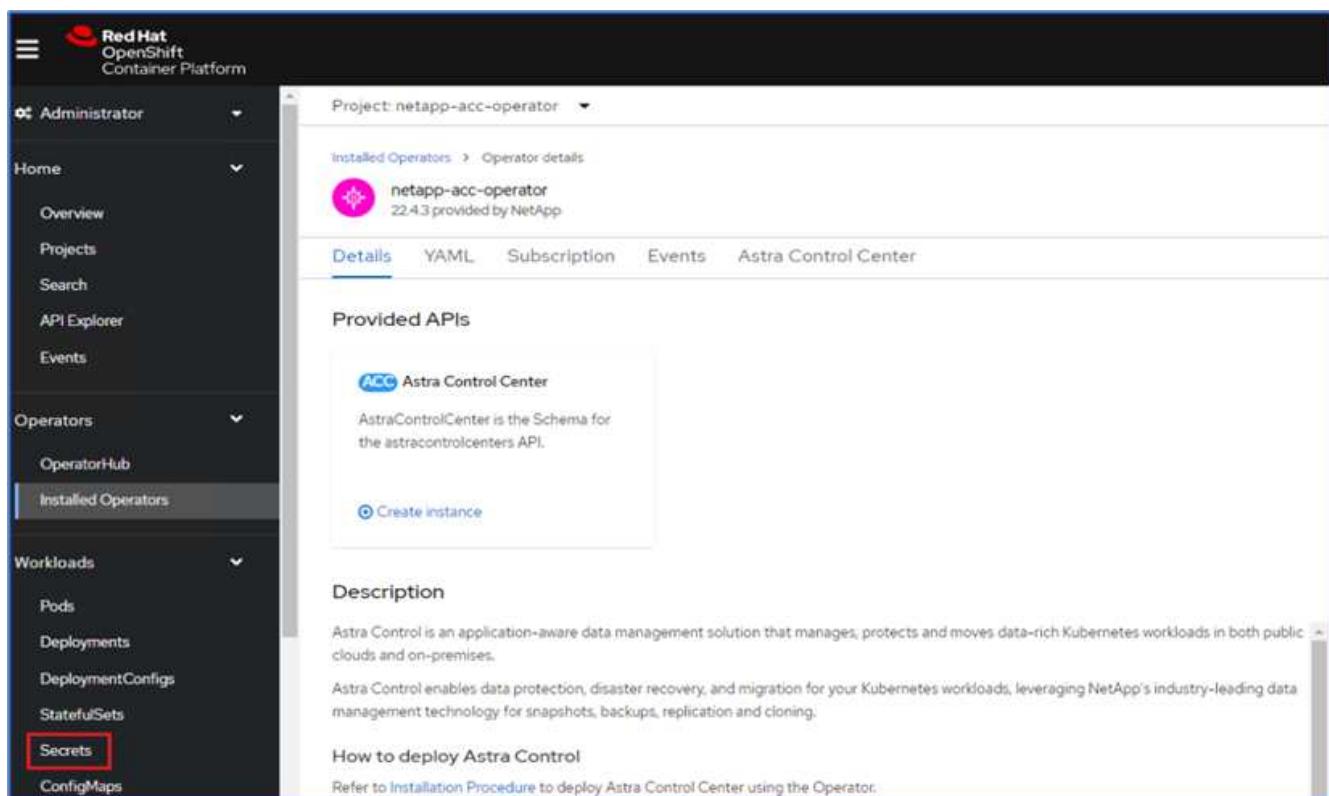
**Approve** **Deny**

[View installed Operators in Namespace netapp-acc-operator](#)

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



13. 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 \***  
astra-registry-cred

Unique name of the new secret.

**Authentication type**  
Image registry credentials

**Registry server address \***  
[REDACTED]  
For example quay.io or docker.io

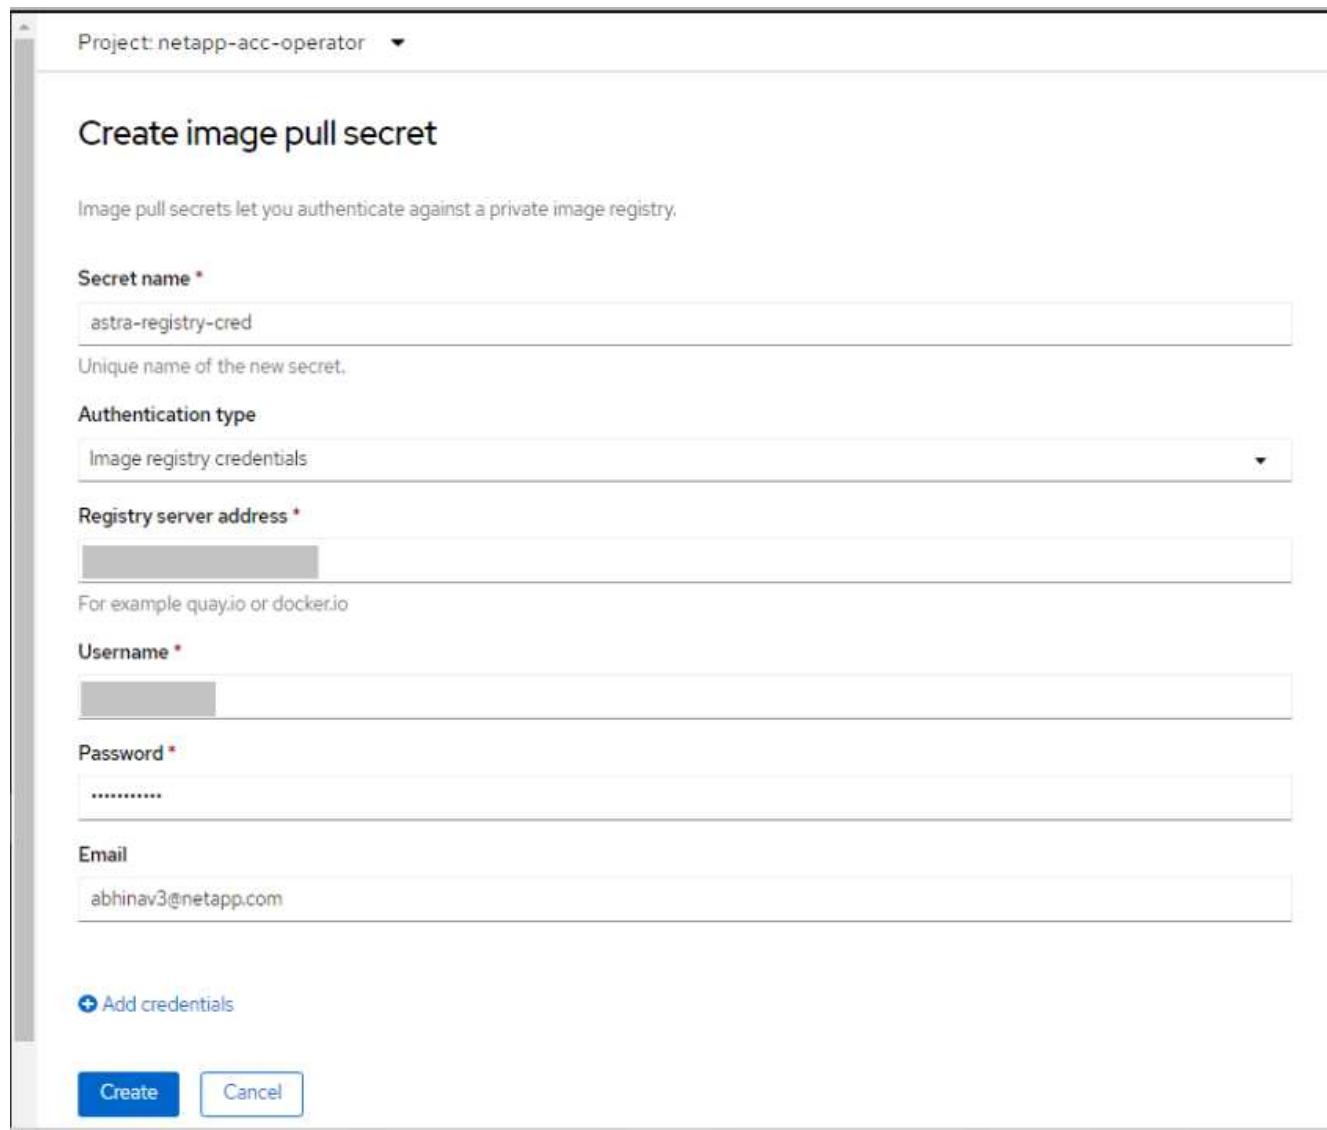
**Username \***  
[REDACTED]

**Password \***  
[REDACTED]

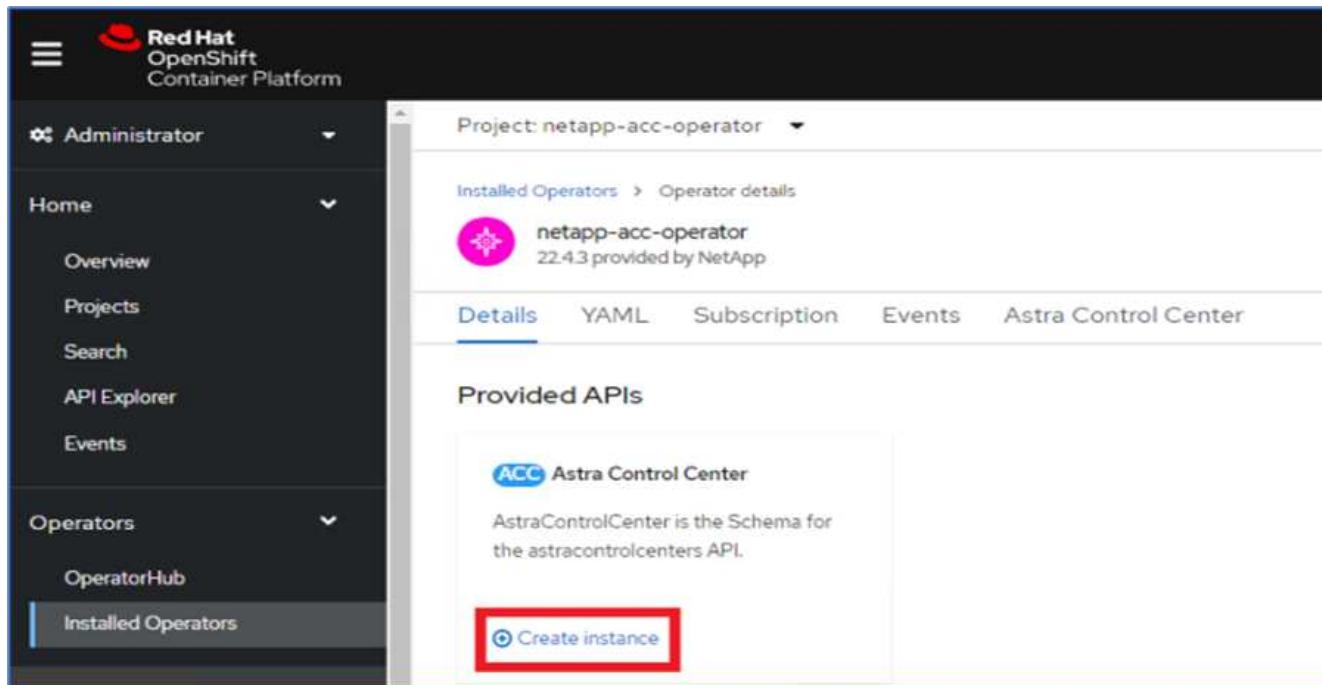
**Email**  
abhinav3@netapp.com

[+ Add credentials](#)

**Create** **Cancel**

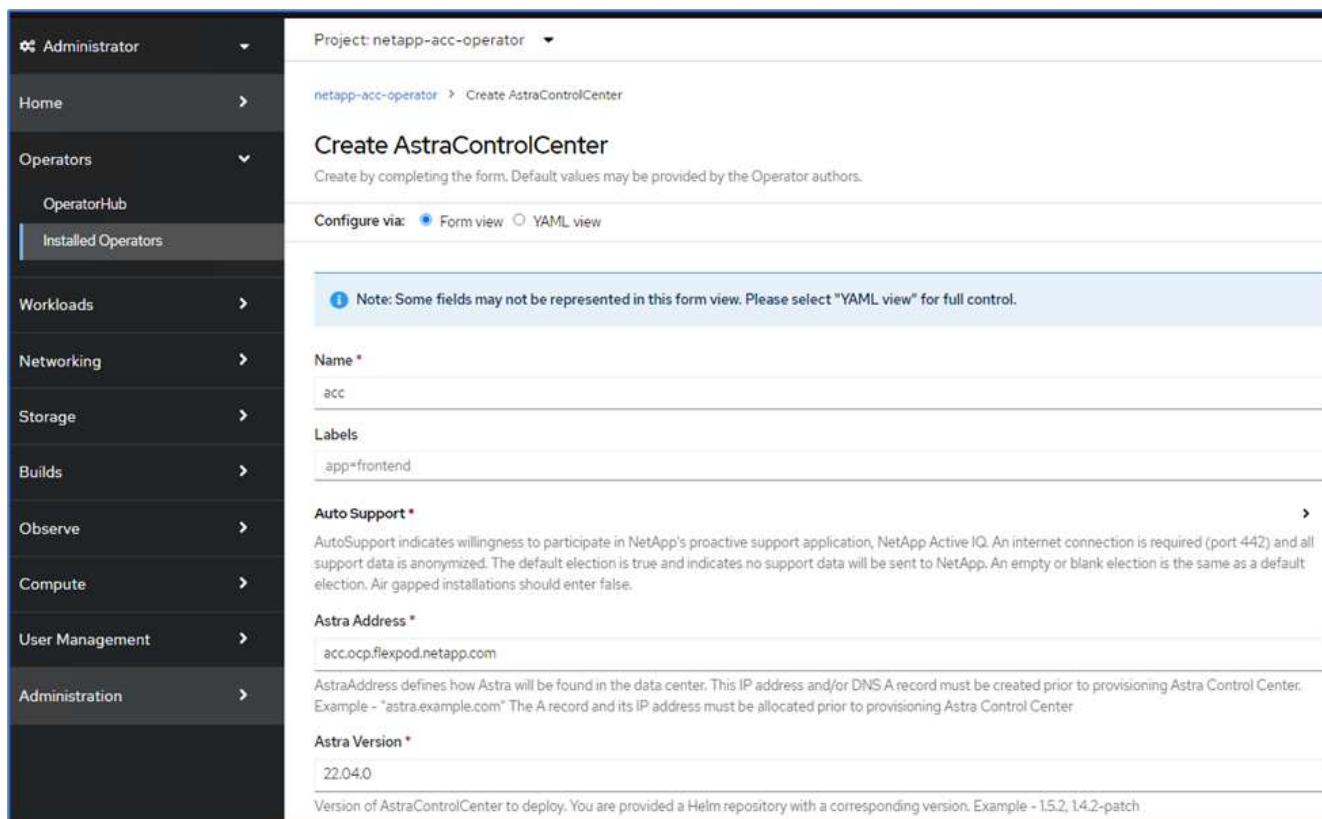


15. From the side menu, select Operators > Installed Operators and click Create Instance under the provided APIs section.



The screenshot shows the Red Hat OpenShift Container Platform web interface. The left sidebar is titled 'Administrator' and includes 'Home', 'Overview', 'Projects', 'Search', 'API Explorer', 'Events', 'Operators', 'OperatorHub', and 'Installed Operators'. The 'Installed Operators' option is selected. The main content area is titled 'Project: netapp-acc-operator' and shows 'Installed Operators > Operator details' for 'netapp-acc-operator' (version 22.4.3 provided by NetApp). It includes tabs for 'Details' (selected), 'YAML', 'Subscription', 'Events', and 'Astra Control Center'. Below the tabs, a section titled 'Provided APIs' contains a box for 'Astra Control Center' with the description 'AstraControlCenter is the Schema for the astracontrolcenters API.' and a 'Create instance' button, which is highlighted with a red box.

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



The screenshot shows the 'Create AstraControlCenter' form. The left sidebar is the same as the previous screenshot. The main form has a title 'Create AstraControlCenter' and a note: 'Create by completing the form. Default values may be provided by the Operator authors.' It includes a 'Configure via:' section with 'Form view' (selected) and 'YAML view' radio buttons. A note below says: 'Note: Some fields may not be represented in this form view. Please select "YAML view" for full control.' The form fields are: 'Name \*' (value: 'acc'), 'Labels' (value: 'app=frontend'), 'Auto Support \*' (note: '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 \*' (value: 'acc.ocp.flexpod.netapp.com'), 'Astra Address' (note: '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'), and 'Astra Version \*' (value: '22.04.0'). A note at the bottom says: 'Version of AstraControlCenter to deploy. You are provided a Helm repository with a corresponding version. Example - 1.5.2, 14.2-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.

Project: netapp-acc-operator

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**

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

Name  
[Redacted]

The name of the image registry. For example "example.registry/astra". Do not prefix with protocol.

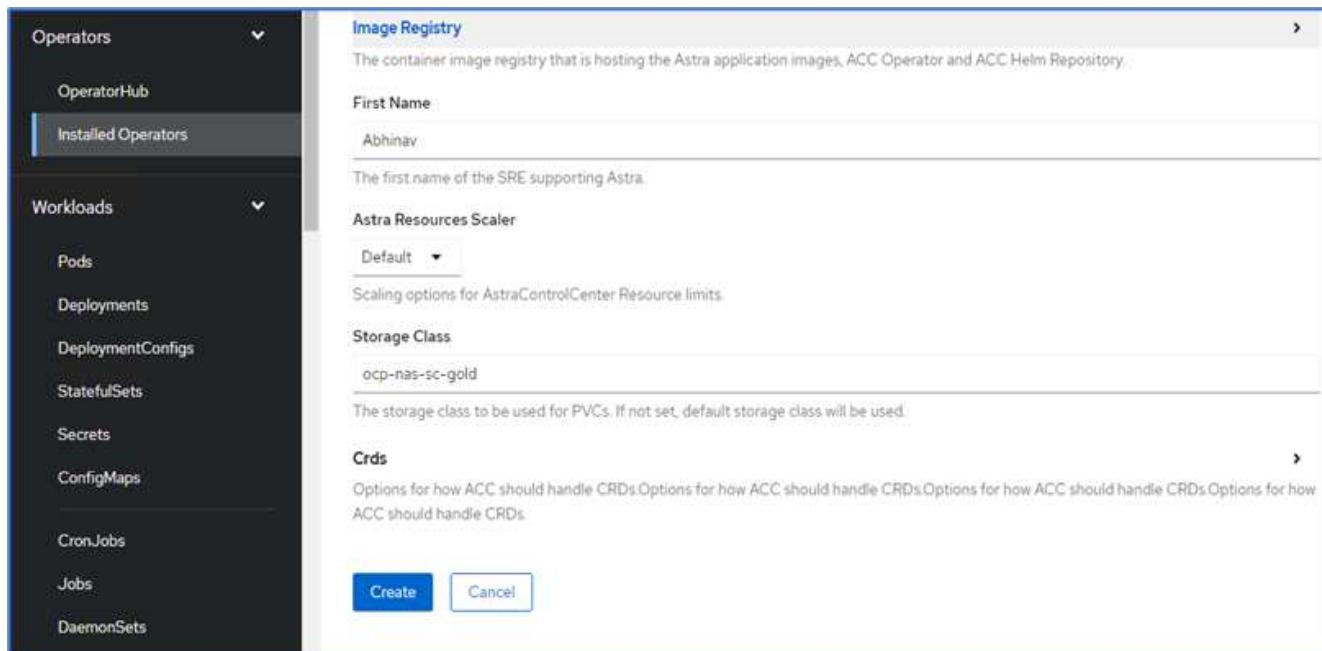
Secret  
astra-registry-cred

The name of the Kubernetes secret that will authenticate with the image registry.

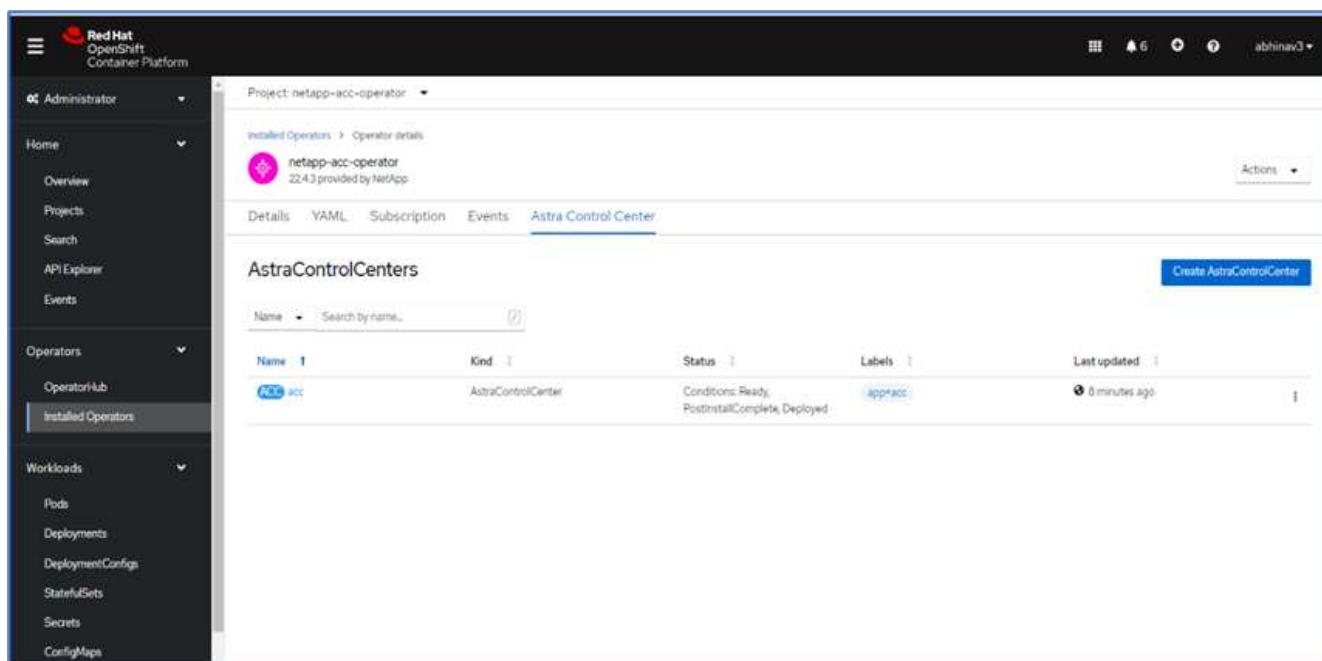


In this solution, the Metallb load balancer is used. Therefore, the ingress type is AccTraefik. This exposes the Astra Control Center traefik gateway as a Kubernetes service of type LoadBalancer.

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



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



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   0
acc-operator-controller-manager-5c656c44c6-tqnmn   2/2    Running   0
activity-589c6d59f4-x2sfs        1/1    Running   0
```

6m4s				
api-token-authentication-4q51j	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				
bucketservice-85495bc475-5zcc5	1/1	Running	0	
5m55s				
cert-manager-67f486bbc6-txhh6	1/1	Running	0	
9m5s				
cert-manager-cainjector-75959db744-415p5	1/1	Running	0	
9m6s				
cert-manager-webhook-765556b869-g6wdf	1/1	Running	0	
9m6s				
cloud-extension-5d595f85f-txrf1	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-817cq	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-6c65fbdd4-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-vcb4j	1/1	Running	0	
5m59s				
storage-provider-bb85ff965-m7qrg	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				

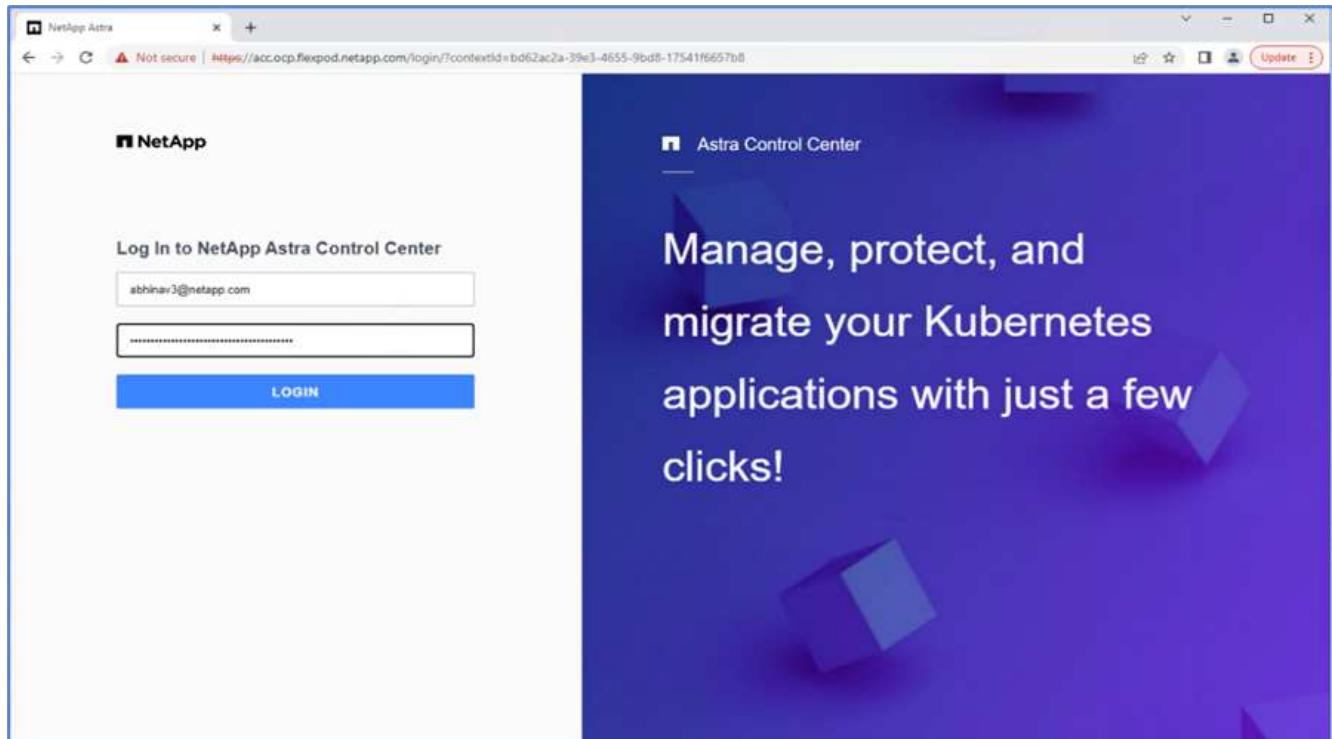


Each pod should have a status of Running. It might take several minutes before the system pods are deployed.

- 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
```

- In a browser, navigate to the URL by using the FQDN that you had provided.
- 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.

- Change the password and proceed.

NetApp

Welcome to NetApp Astra Control Center

Update your password to proceed

New Password

Confirm Password

>Passwords must contain:  

- At least 8 characters
- No more than 64 characters
- At least one uppercase letter
- At least one lowercase letter
- At least one number
- At least one special character

UPDATE PASSWORD

Astra Control Center

Manage, protect, and migrate your Kubernetes applications with just a few clicks!

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.

ocp

Dashboard

Account

Users Credentials Notifications License Packages Connections

ASTRA CONTROL CENTER LICENSE OVERVIEW

You have no active Astra Control Center license

To get started with Astra Control Center, use your account ID below to begin the license process. When you receive your license, select Add license to manually upload the file. [More information](#)

Astra Control Center account ID: 98338fa8-353b-4091-9b09-57694b3f815b

Have an evaluation license?

Select Add license to manually upload your evaluation license file. [More information](#)

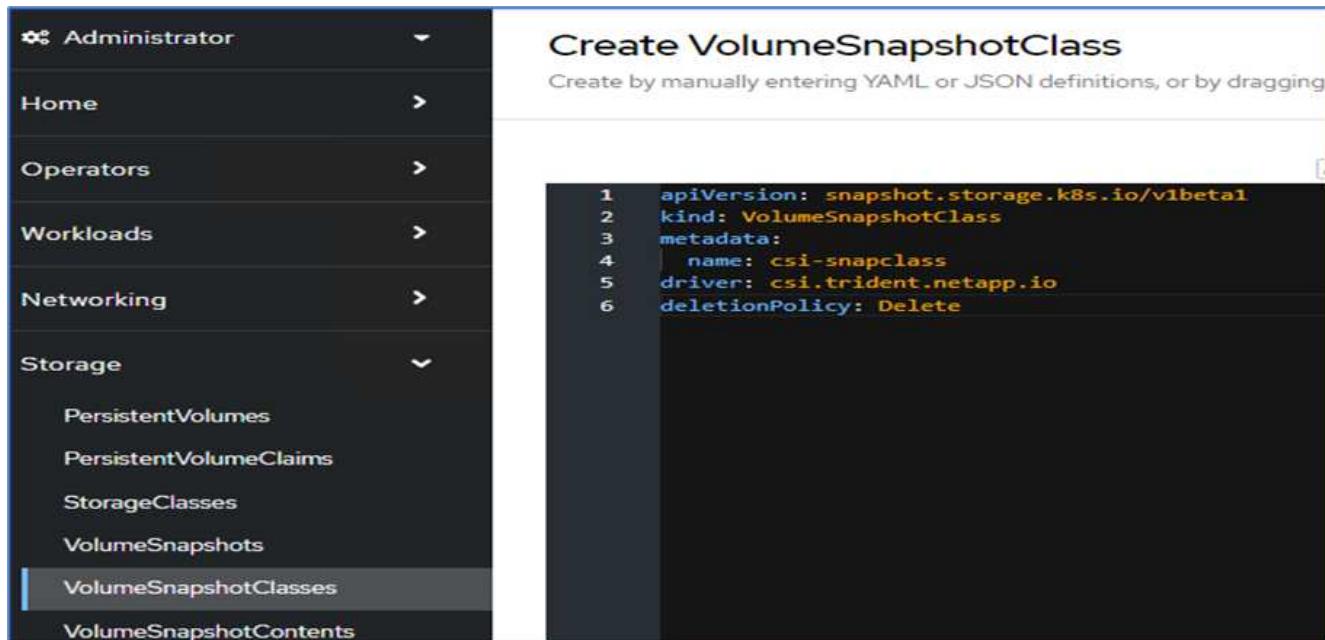
Add license

Astra Data Store licenses

+ Manually add license

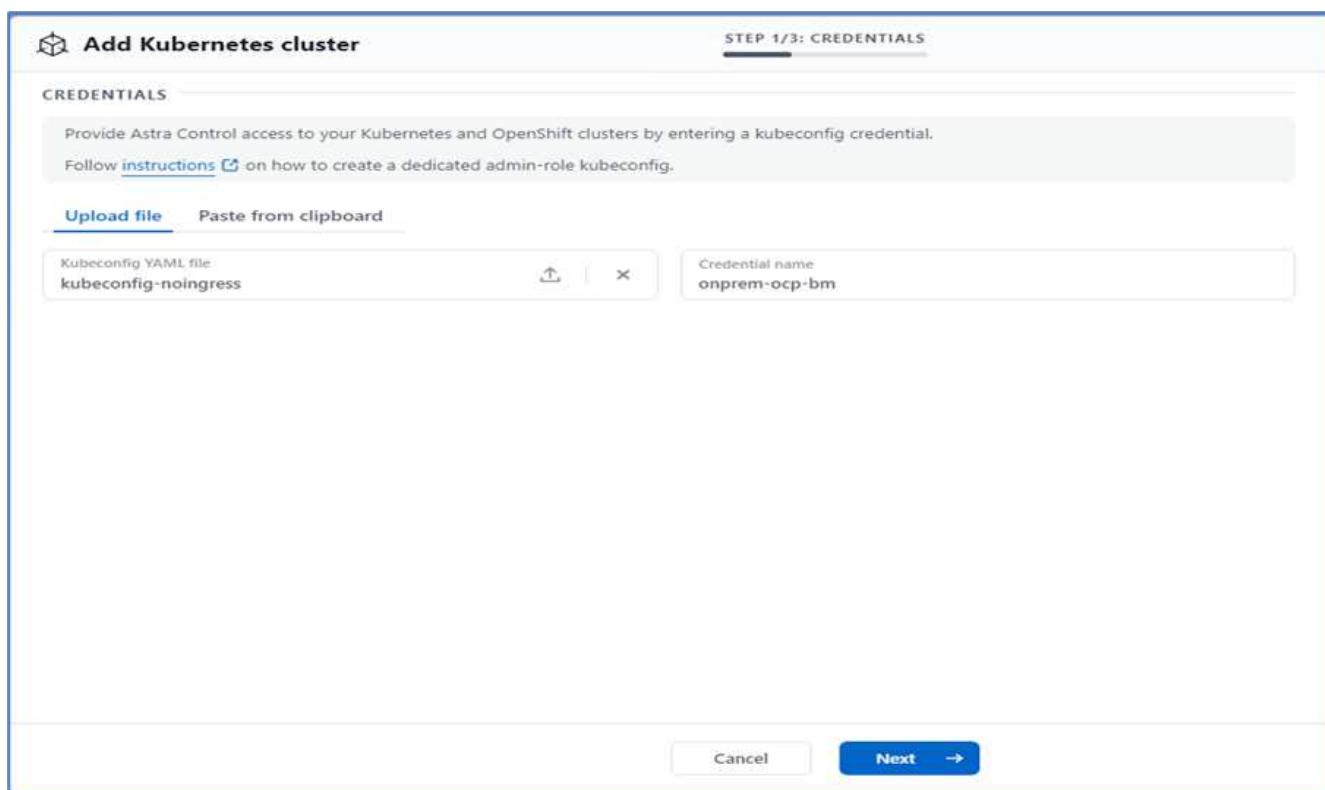
Search

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.



```
1  apiVersion: snapshot.storage.k8s.io/v1beta1
2  kind: VolumeSnapshotClass
3  metadata:
4    name: csi-snapclass
5  driver: csi.trident.netapp.io
6  deletionPolicy: Delete
```

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.

Set default	Storage class	Storage provisioner	Reclaim policy	Binding mode	Eligible
•	ocp-nas-sc-gold	csi.trident.netapp.io	Delete	Immediate	<input checked="" type="checkbox"/>

5. 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.

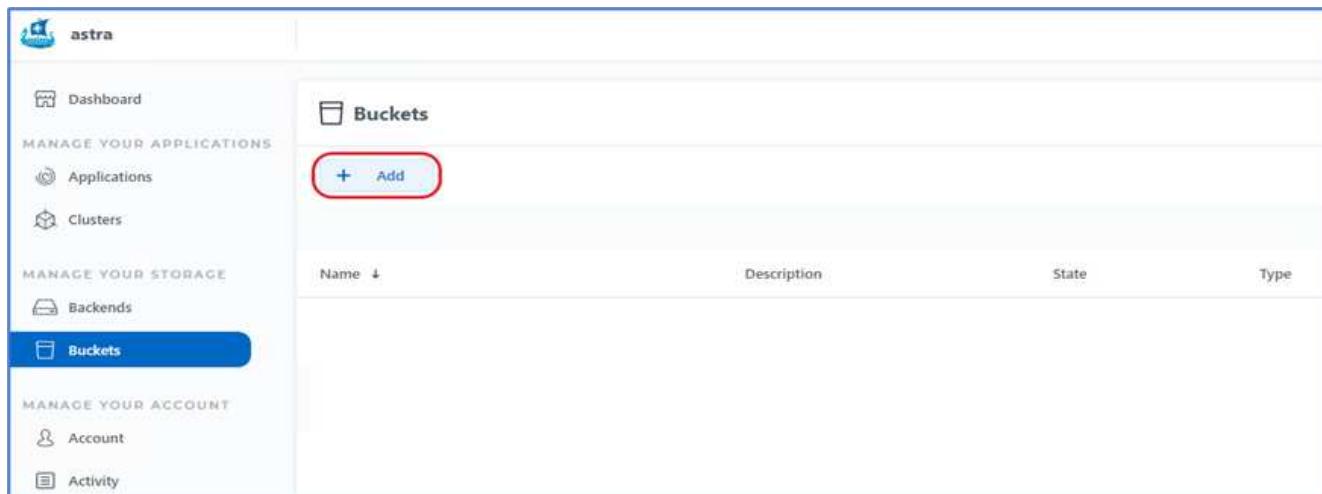
6. To manage the storage, go to Backends, click the three dots under Actions against the backend that you would like to manage. Click Manage.

Name	State	Capacity	Throughput	Type	Cluster	Cloud	Actions
ct90-cluster	Discovered	Not available yet	Not available yet	ONTAP 9.11, t	Not applicable	Not applicable	<a href="#">Manage</a> <a href="#">Remove</a>
healthylife	Discovered	Not available yet	Not available yet	ONTAP 9.11, t	Not applicable	Not applicable	<a href="#">Manage</a> <a href="#">Remove</a>
singlecvoaws	Discovered	Not available yet	Not available yet	ONTAP 9.11, t	Not applicable	Not applicable	<a href="#">Manage</a> <a href="#">Remove</a>

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

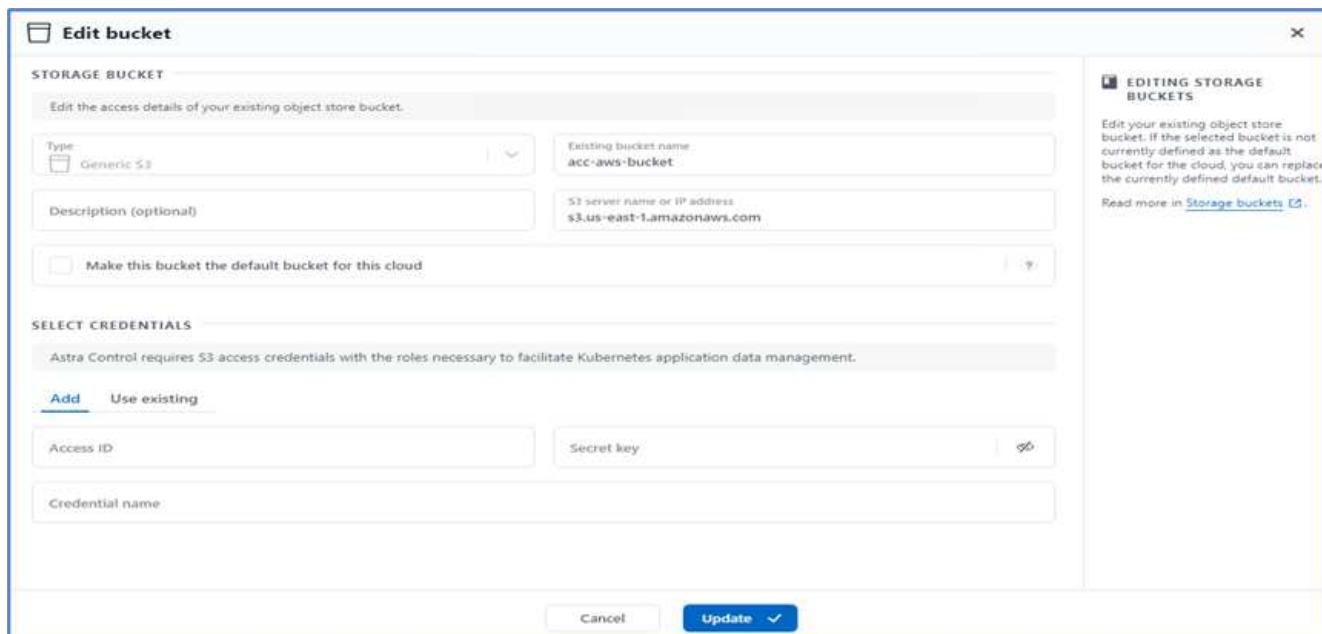
Backends							
Actions		Search		Managed		Discovered	
Name	State	Capacity	Throughput	Type	Cluster	Cloud	Actions
c190-cluster	Available	0.4/10.64 TiB: 3.8%	Not available yet	ONTAP 9.11.1	Not applicable	Not applicable	<input type="button" value="⋮"/>
healthylife	Available	5.16/106.42 TiB: 4.8%	Not available yet	ONTAP 9.11.1	Not applicable	Not applicable	<input type="button" value="⋮"/>
singlecvoaws	Available	0.07/0.62 TiB: 11.9%	Not available yet	ONTAP 9.11.1	Not applicable	Not applicable	<input type="button" value="⋮"/>

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



The screenshot shows the Astra Control web interface. On the left, there's a sidebar with 'Dashboard', 'MANAGE YOUR APPLICATIONS' (Applications, Clusters), 'MANAGE YOUR STORAGE' (Backends, **Buckets**), and 'MANAGE YOUR ACCOUNT' (Account, Activity). The 'Buckets' section is currently selected. The main area is titled 'Buckets' and contains a table with columns: Name, Description, State, and Type. A red box highlights the 'Add' button at the top of the table area.

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

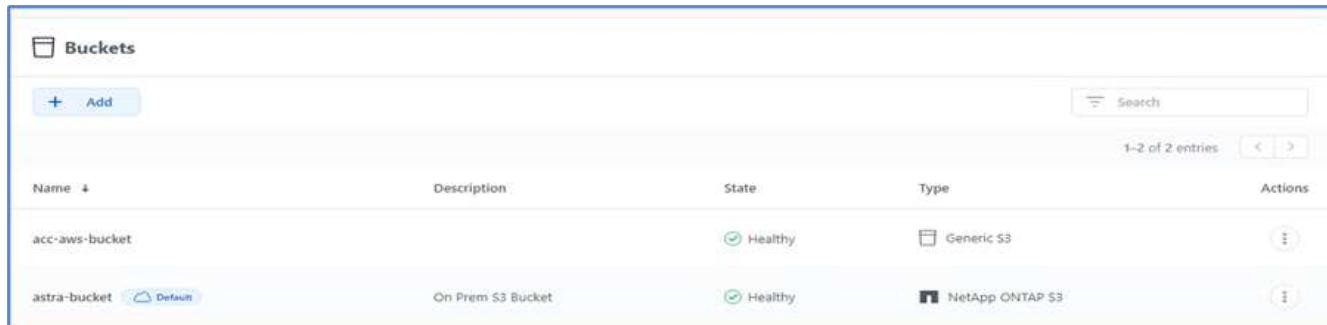


The screenshot shows the 'Edit bucket' dialog. Under 'STORAGE BUCKET', it says 'Edit the access details of your existing object store bucket.' It has a 'Type' dropdown set to 'Generic S3', an 'Existing bucket name' field with 'acc-aws-bucket', an 'S3 server name or IP address' field with 's3.us-east-1.amazonaws.com', and an 'Description (optional)' field that is empty. There's a checkbox 'Make this bucket the default bucket for this cloud' which is unchecked. Under 'SELECT CREDENTIALS', it says 'Astra Control requires S3-access credentials with the roles necessary to facilitate Kubernetes application data management.' It has tabs for 'Add' (selected) and 'Use existing', with fields for 'Access ID' and 'Secret key', and a 'Credential name' dropdown. A note on the right says 'Editing storage buckets' with a note about replacing the default bucket.



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.

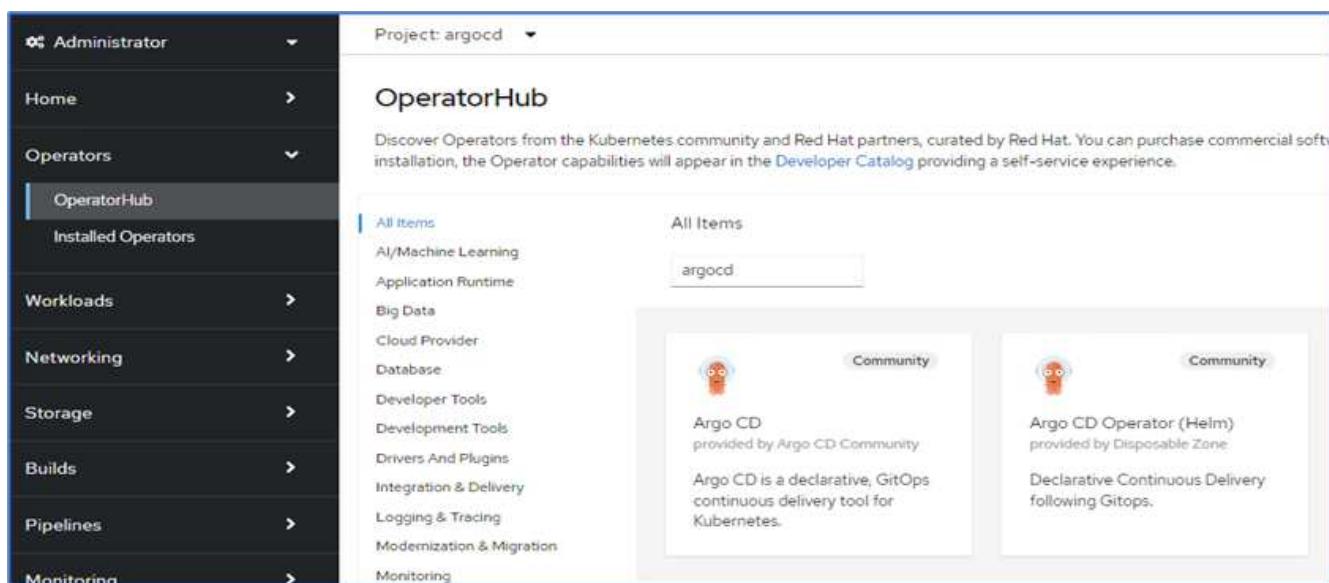
To validate a successful application recovery, simulate an on-premises failure of an application running on the FlexPod system and restore the application to a K8s cluster running in the cloud by using a remote backup.

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 \***

argocd

**Update approval \*** [?](#)

Automatic

Manual

**Provided APIs**

<b>Application</b>	<b>ApplicationSet</b>
An Application is a group of Kubernetes resources as defined by a manifest.	An ApplicationSet is a group or set of Application resources.
<b>AppProject</b>	<b>Argo CDExport</b>
An AppProject is a logical grouping of Argo CD Applications.	ArgoCDExport is the Schema for the argocdexports API
<b>Argo CD</b>	
ArgoCD is the Schema for the argocds API	

**Actions**

**Install** **Cancel**

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 CDExport **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.

**Name \***  
argocd-netapp

**Labels**  
app=frontend

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

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 argocd-netapp

**Details** **YAML**

**Secret details**

**Name** argocd-netapp-cluster **Type** Opaque

**Namespace** argocd

**Labels**  
app.kubernetes.io/managed-by=argocd-netapp app.kubernetes.io/name=argocd-netapp-cluster app.kubernetes.io/part-of=argocd

**Annotations**  
Annotations

**Created at** 2 minutes ago

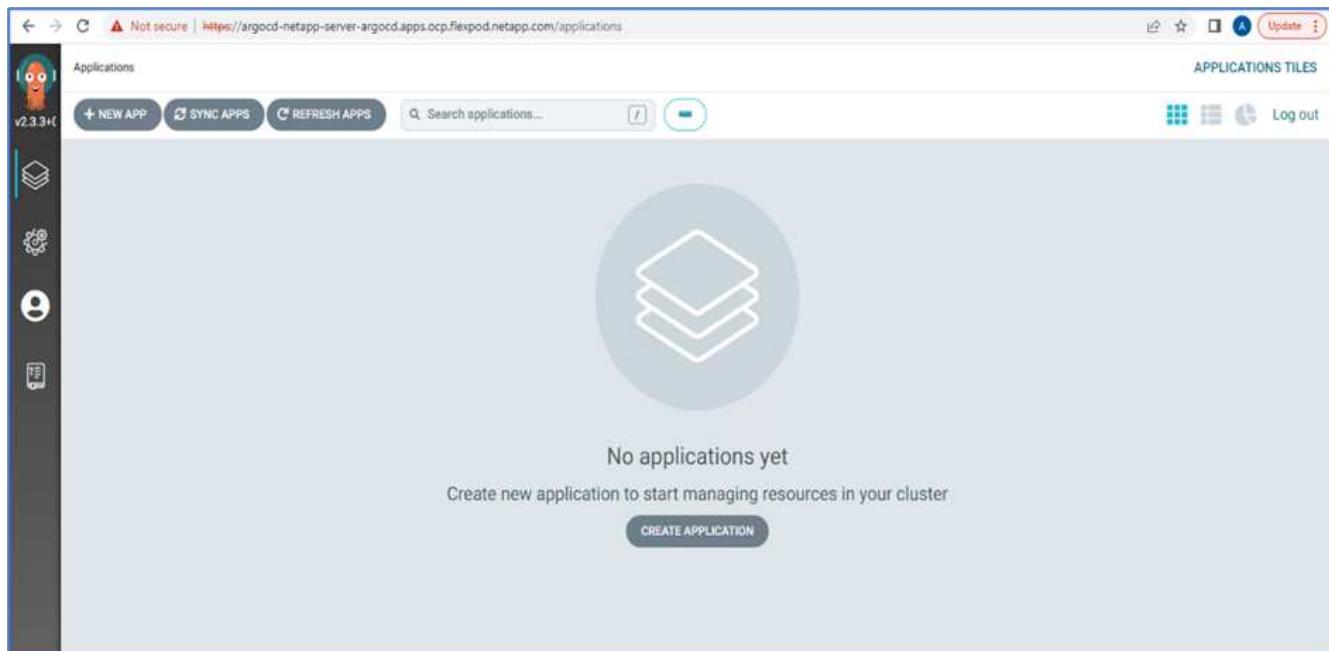
**Owner** argocd-netapp

**Data**

admin.password  
.....

**Copy**

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 X

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

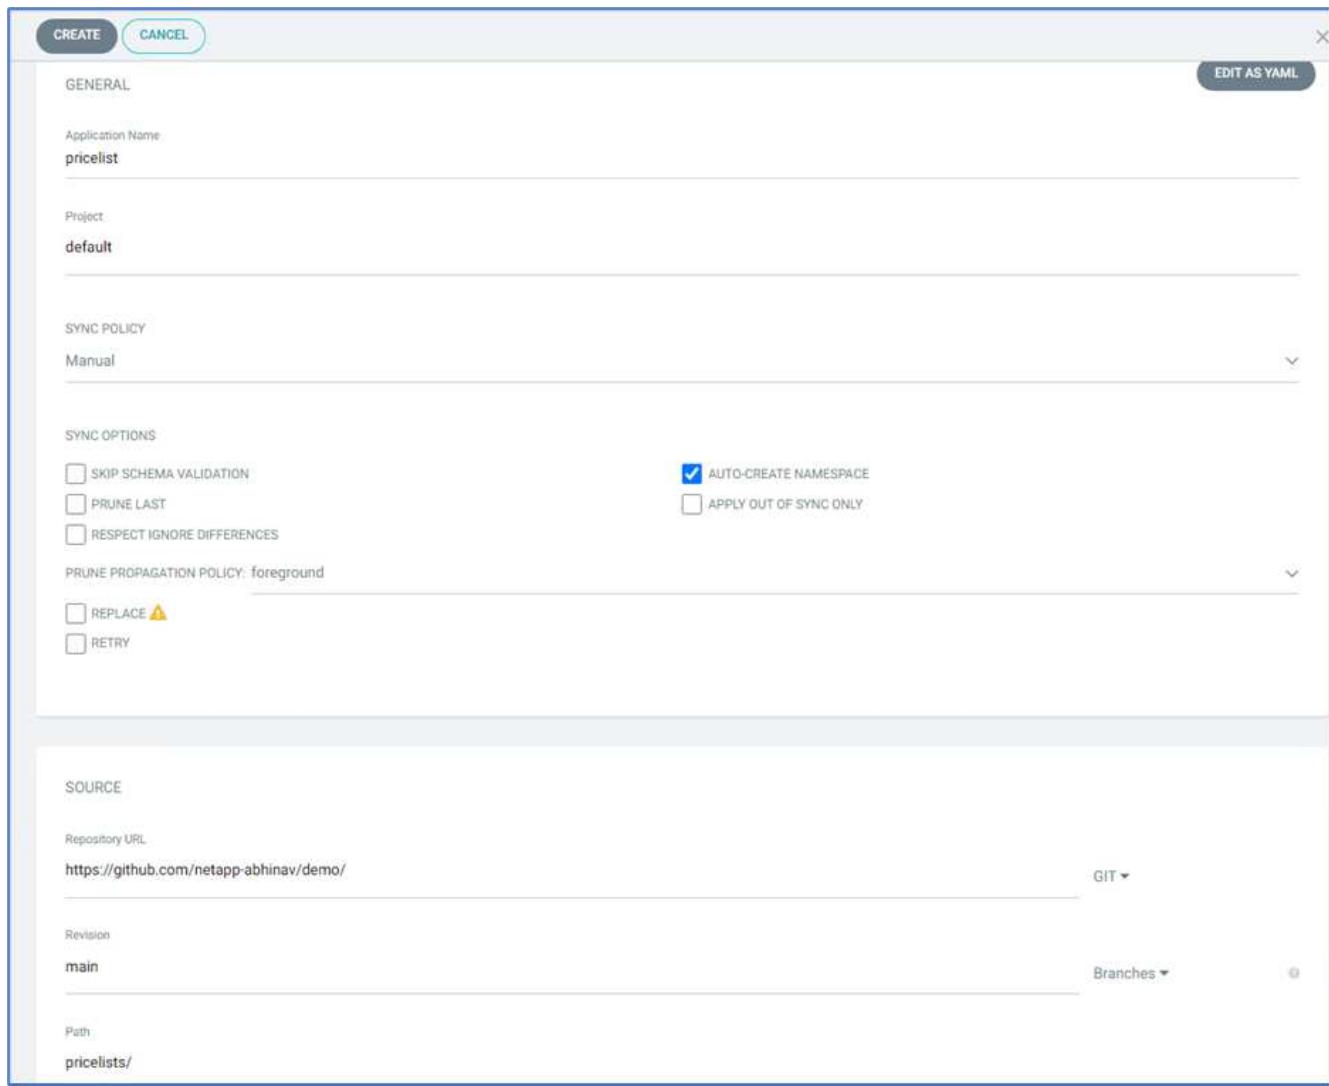
REPLACE  RETRY

SOURCE

Repository URL: <https://github.com/netapp-abhinav/demo/> GIT

Revision: main Branches

Path: pricelists/

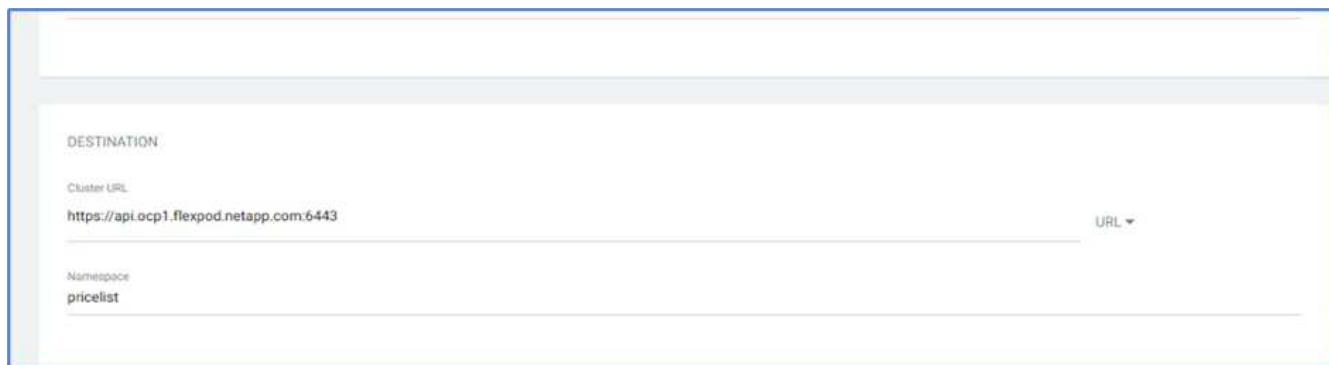


10. Enter the OpenShift cluster where the app will be deployed along with the namespace.

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.

Applications

+ NEW APP   SYNC APPS   REFRESH APPS   Search applications...

FILTERS

FAVORITES ONLY

SYNC STATUS

- Unknown 0
- Synced 0
- OutOfSync 1

HEALTH STATUS

- Unknown 0
- Progressing 0
- Suspended 0

pricelist

Project: default

Labels: missing

Status: OutOfSync

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

Target Revision: main

Path: pricelists/

Destination: default/api-ocp1-flexpod-netapp-com:6443/abhinav3

Namespace: pricelist

SYNC   REFRESH   DELETE

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

Project: pricelist

PersistentVolumeClaims

Create PersistentVolumeClaim

Name	Status	PersistentVolumes	Capacity	Used	StorageClass
pvc-64ef51a3-1fd7-489d-906d-3d368128a6e9	Bound	pvc-64ef51a3-1fd7-489d-906d-3d368128a6e9	1GiB	320 kB	ocp-nas-sc-gold

13. Log into System Manager and verify the PVC.

Volumes

+ Add   More

Name	Storage VM	Status	Capacity	IOPS	Latency (ms)
trident_pvc_64ef51a3_1fd7_489d_906d_3d368128a6e9	Infra_SVM	Online	5.21 MiB Used   0.995 GiB available   1.11 GiB	0	0

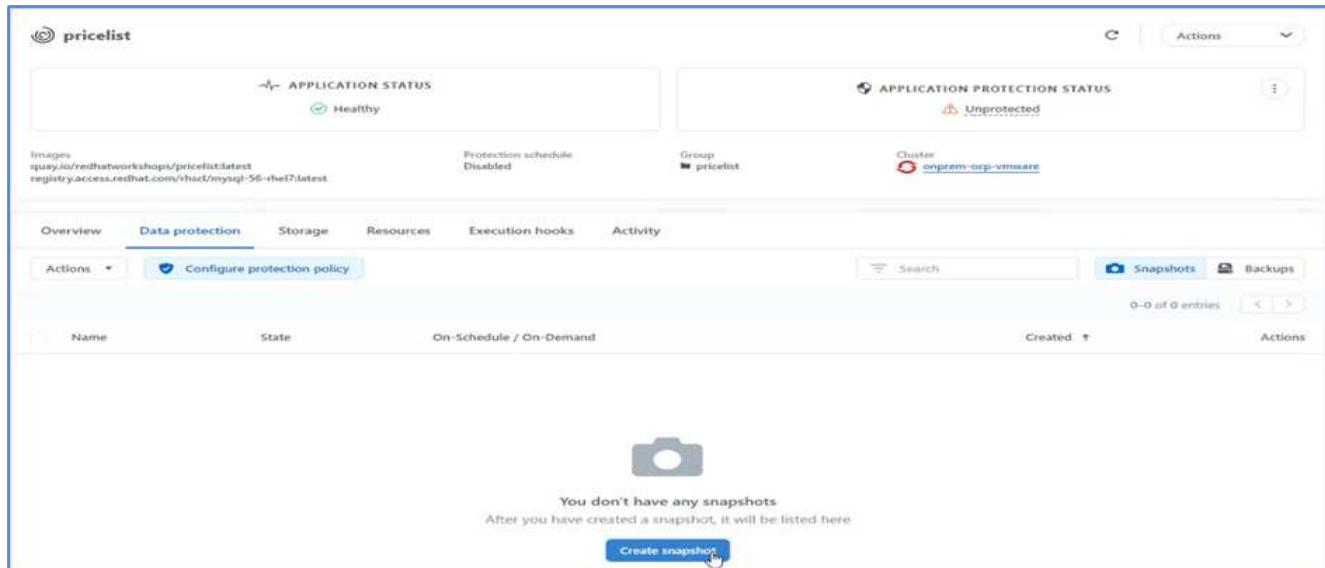
14. After the Pods are running, select Networking > Routes from the side menu, and click the URL under Location.

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.



APPLICATION STATUS  
Healthy

APPLICATION PROTECTION STATUS  
Unprotected

Images: quay.io/redhatworkshop/pricelist:latest  
registry.access.redhat.com/rhsc1/mysql-56-rhel7:latest

Protection schedule: Disabled  
Group: pricelist  
Cluster: onprem-ocp-vmware

Overview Data protection Storage Resources Execution hooks Activity

Actions Configure protection policy

Search:  Snapshots Backups

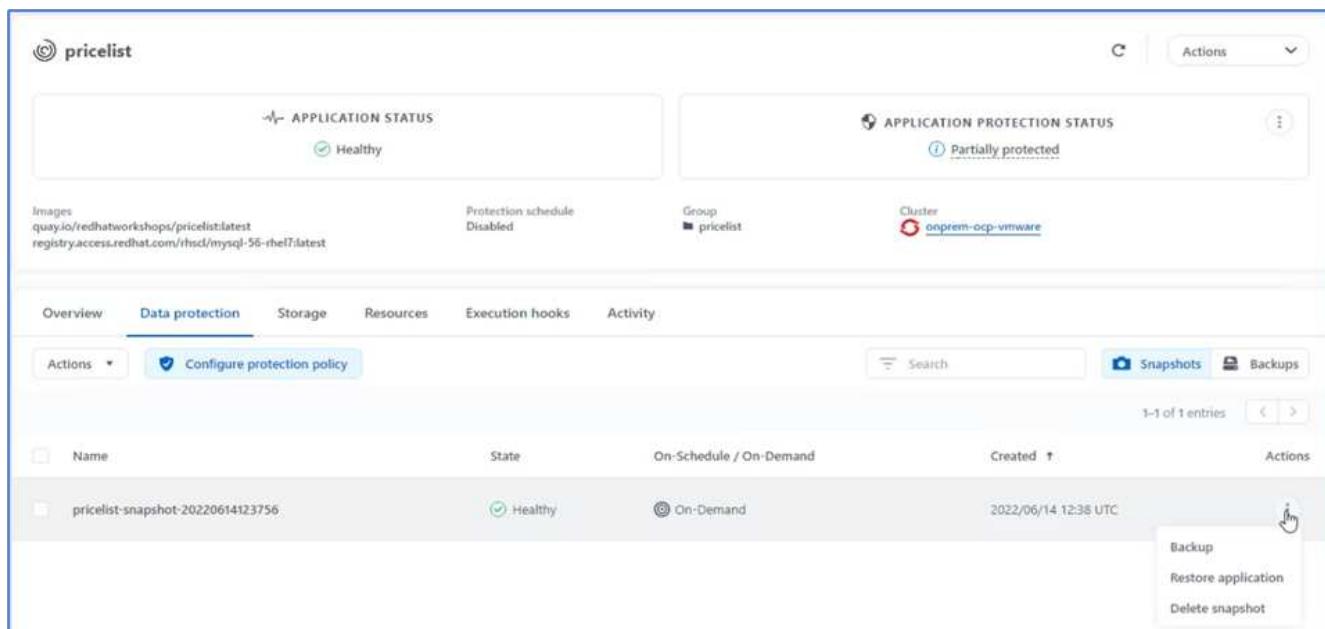
0-0 of 0 entries

Name	State	On-Schedule / On-Demand	Created	Actions
 You don't have any snapshots After you have created a snapshot, it will be listed here <a href="#">Create snapshot</a>				



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.



APPLICATION STATUS  
Healthy

APPLICATION PROTECTION STATUS  
Partially protected

Images: quay.io/redhatworkshop/pricelist:latest  
registry.access.redhat.com/rhsc1/mysql-56-rhel7:latest

Protection schedule: Disabled  
Group: pricelist  
Cluster: onprem-ocp-vmware

Overview Data protection Storage Resources Execution hooks Activity

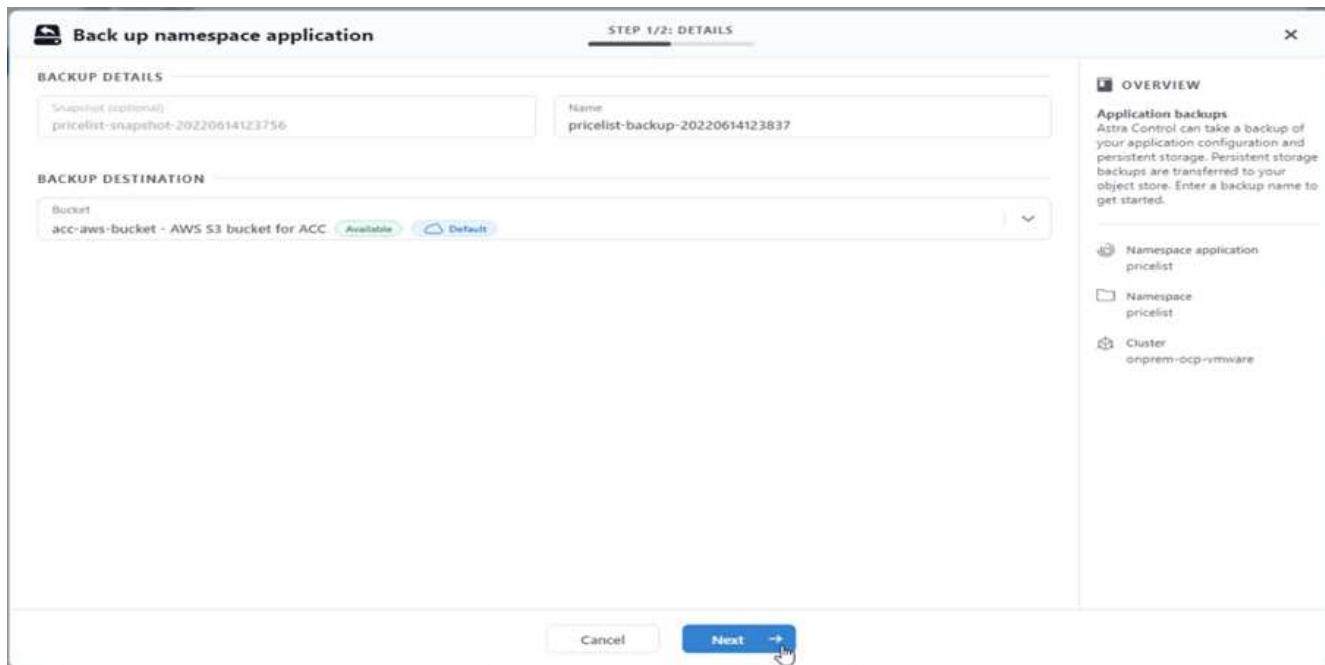
Actions Configure protection policy

Search:  Snapshots Backups

1-1 of 1 entries

Name	State	On-Schedule / On-Demand	Created	Actions
pricelist-snapshot-20220614123756	Healthy	On-Demand	2022/06/14 12:38 UTC	<a href="#">Backup</a> <a href="#">Restore application</a> <a href="#">Delete snapshot</a>

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



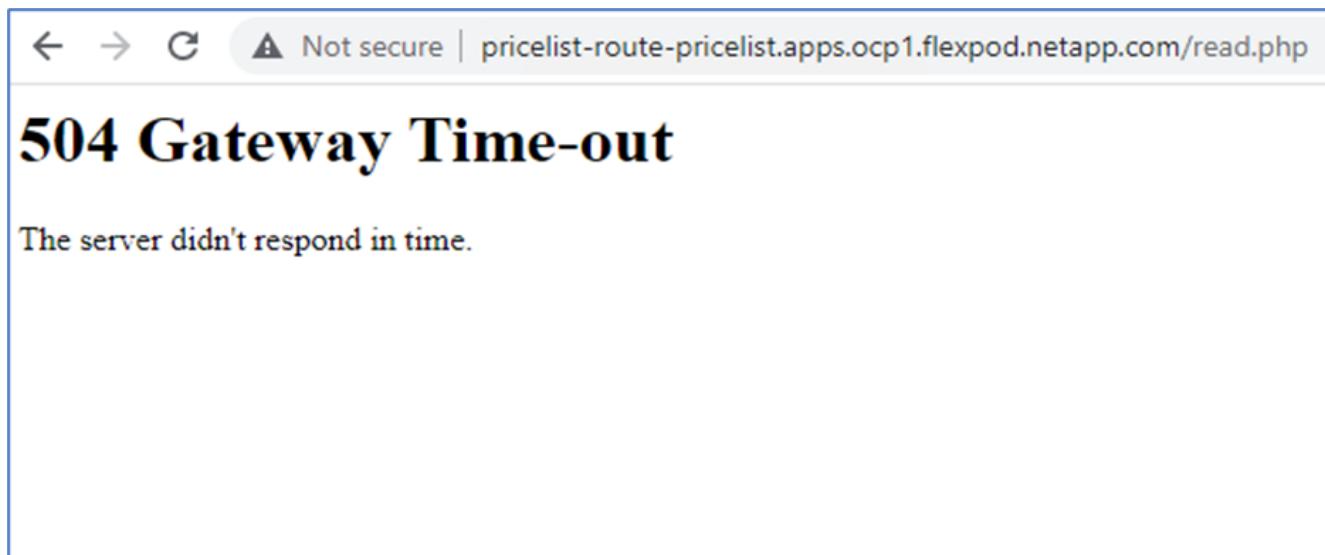
21. The backup operation should create a folder with multiple objects in the AWS S3 bucket.

Name	Type	Last modified	Size	Storage class
config	Folder	June 14, 2022, 05:39:19 (UTC-07:00)	155.0 B	Standard
data/	Folder	-	-	-
index/	Folder	-	-	-
keys/	Folder	-	-	-
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.

Name	State	Subtype	Configured Protocols	IPspace
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.

Name	State	On-Schedule / On-Demand	Bucket	Created	Actions
pricelist-backup-20220614123837	Healthy	On-Demand	acc-aws-bucket	2022/06/14 12:38 UTC	<a href="#">Restore application</a> <a href="#">Delete backup</a>

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 in the following sections.

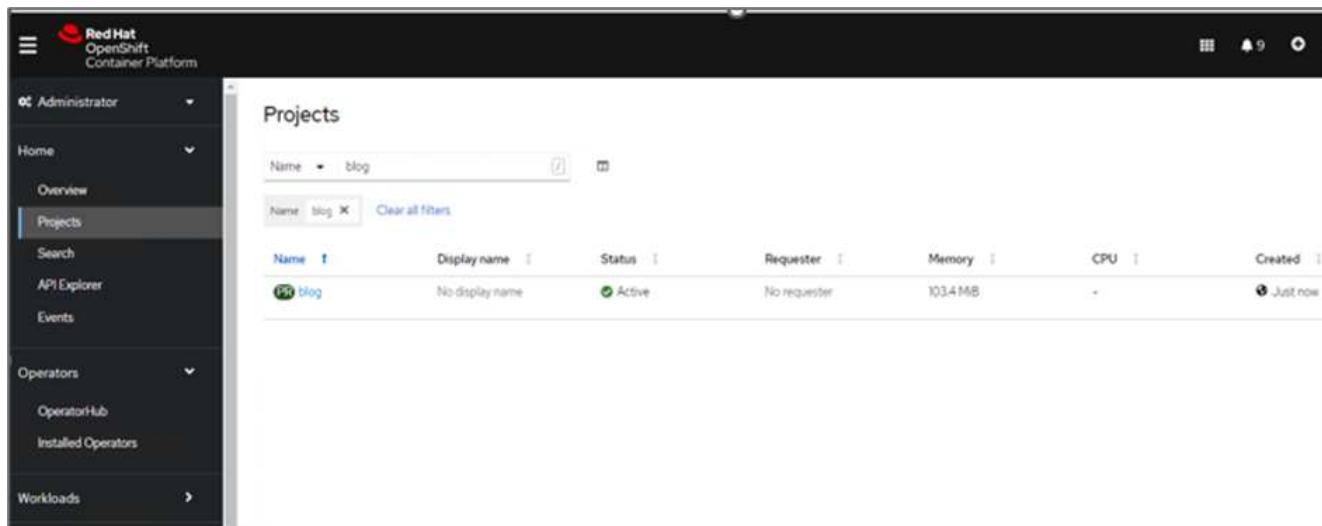
##### Part 1

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:

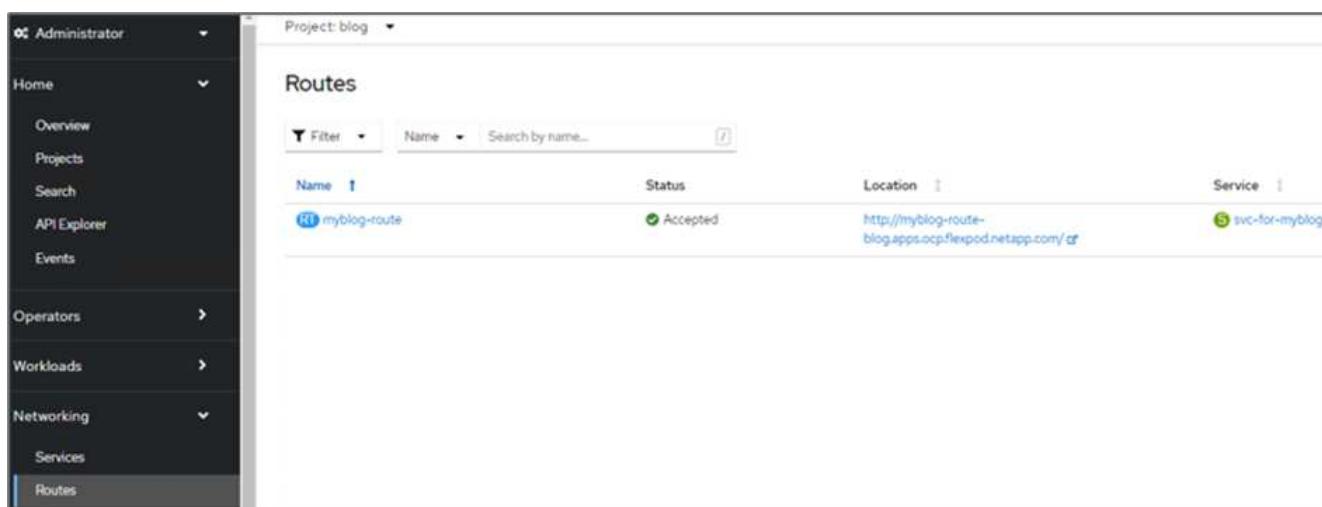
1. Deploy the sample blogging app and sync it from ArgoCD.

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



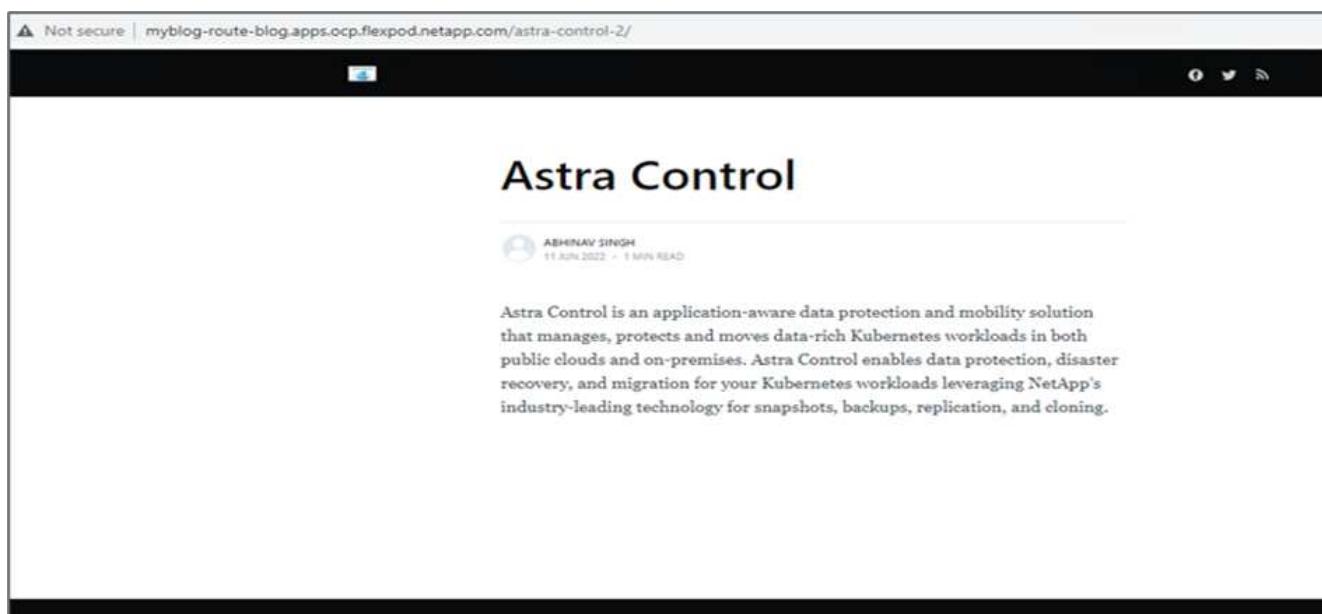
Name	Display name	Status	Requester	Memory	CPU	Created
blog	No display name	Active	No requester	103.4 MB	0	Just now

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



Name	Status	Location	Service
myblog-route	Accepted	http://myblog-route-blog.apps.ocp.flexpod.netapp.com/or	svc-for-myblog

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

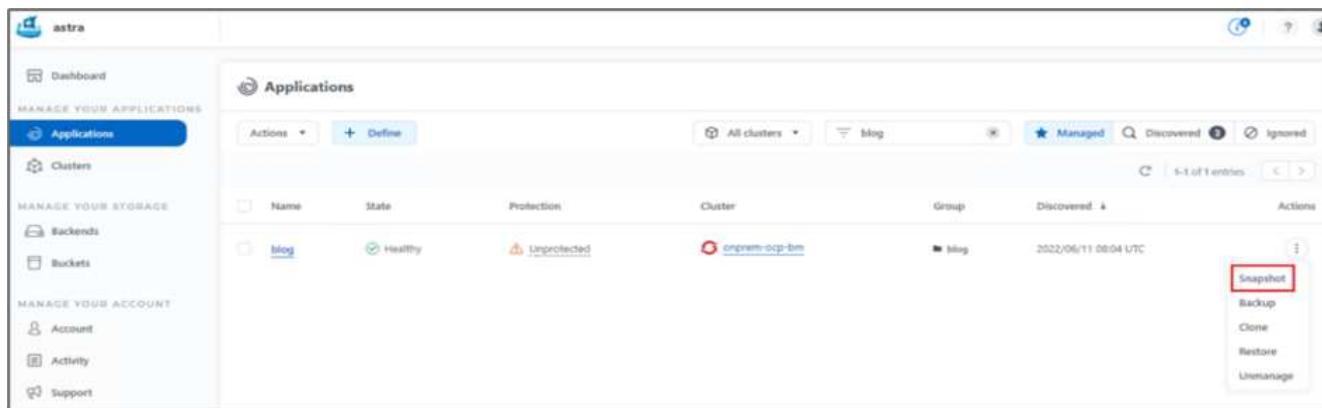


Astra Control

ABHINAV SINGH 11 JUN 2022 • 1 MIN READ

Astra Control is an application-aware data protection and mobility 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 technology for snapshots, backups, replication, and cloning.

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



Name	State	Protection	Cluster	Group	Discovered	Actions
Blog	Healthy	Unprotected	ongrem-ocp-bm	Blog	2022/06/11 08:04 UTC	<a href="#">Snapshot</a> <a href="#">Backup</a> <a href="#">Clone</a> <a href="#">Restore</a> <a href="#">Unmanage</a>



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.

```
spec:  
  containers:  
    - name: myblog  
      image: ghost:latest  
      imagePullPolicy: Always  
    ports:  
      - containerPort: 2368
```

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

**P** 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

+[37m"You must be on the latest v3.x to update across major versions - https://ghost.org/docs/update/"+[39m
+[33m"Run '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.

← → ⌂ Not secure | myblog-route-blog.apps.ocp.flexpod.netapp.com

Application is not available

The application is currently not serving requests at this endpoint. It may not have been started or is still starting.

**i** Possible reasons you are seeing this page:

- The host doesn't exist. Make sure the hostname was typed correctly and that a route matching this hostname exists.
- The host exists, but doesn't have a matching path. Check if the URL path was typed correctly and that the route was created using the desired path.
- Route and path matches, but all pods are down. Make sure that the resources exposed by this route (pods, services, deployment configs, etc) have at least one pod running.

9. Restore the app from the snapshot.

The screenshot shows the OpenShift Data Protection interface for the 'blog' application. The top navigation bar includes 'Actions' and a dropdown. The 'APPLICATION STATUS' section shows 'Healthy'. The 'APPLICATION PROTECTION STATUS' section shows 'Partially protected'. The 'Images' section lists 'ghost:3.6-alpine' and 'ghosttest'. The 'Protection schedule' is 'Disabled'. The 'Group' is 'blog' and the 'Cluster' is 'onprem-ocp-bm'. The 'Data protection' tab is selected, showing a table with one entry: 'blog-snapshot-20220611125244'. The table columns are 'Name', 'State', 'On-Schedule / On-Demand', 'Created', and 'Actions'. The 'Actions' column for the snapshot shows 'Backup' (disabled), 'Restore application' (highlighted with a red box), and 'Delete snapshot'. A search bar and buttons for 'Snapshots' and 'Backups' are also present.

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

The screenshot shows the 'Restore namespace application' wizard, Step 2/2: SUMMARY. The title is 'Restore namespace application'. The summary section contains a warning message: 'All existing resources associated with this namespace application will be deleted and replaced with the source snapshot "blog-snapshot-20220611125244" taken on 2022/06/11 12:52 UTC. Persistent volumes will be deleted and recreated. External resources with dependencies on this namespace application might be impacted.' It also says 'We recommend taking a snapshot or a backup of your namespace application before proceeding.' The left panel shows the 'SNAPSHOT' section with 'blog-snapshot-20220611125244'. The right panel shows the 'RESTORE' section with 'blog'. The 'DESTINATION GROUP' and 'DESTINATION CLUSTER' are both 'blog' and 'onprem-ocp-bm'. The 'RESOURCE LABELS' section shows 'Cluster Roles' and 'Cluster Role Bindings'. A confirmation message at the bottom asks 'Are you sure you want to restore the namespace application "blog"?'. A text input field 'Type restore below to confirm.' contains 'restore'. A 'Confirm to restore' button is present. At the bottom are 'Back' and 'Restore' buttons.

11. The app restore process starts immediately.

The screenshot shows the 'Applications' interface. The top navigation bar includes 'Actions', '+ Define', and a search bar with 'All clusters' dropdown, 'blog' search term, and filters for 'Managed', 'Discovered', and 'Ignored'. The table shows one entry: 'blog' with 'Restoring' state, 'Partially protected' protection, 'onprem-ocp-bm' cluster, 'blog' group, and 'Discovered' status. A search bar and buttons for 'Managed', 'Discovered', and 'Ignored' are also present.

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

The screenshot shows the Astra Control Center interface under the 'Applications' tab. At the top, there are buttons for 'Actions' (dropdown), '+ Define', and search fields for 'All clusters' (set to 'blog') and 'Managed' status. Below the search is a table with the following columns: Name, State, Protection, Cluster, Group, Discovered, and Actions. A single row is listed for the application 'blog', which is 'Healthy' and 'Partially protected' in the 'Protection' column. The 'Cluster' column shows 'onprem-ocp-bm'. The 'Group' column shows 'blog'. The 'Discovered' column shows '2022/06/11 12:34 UTC'. The 'Actions' column has a three-dot menu icon.

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

The screenshot shows a web browser window with the URL 'Not secure | myblog-route-blog.apps.ocp.flexpod.netapp.com/astra-control-2/'. The page title is 'Astra Control'. Below the title, there is a user profile for 'ASHWANI DEVKH' with a timestamp '17 JUN 2022 - 1 MIN READ'. A text block describes Astra Control as an application-aware data protection and mobility solution that manages, protects and moves data-rich Kubernetes workloads in both public clouds and on-premises. It mentions that Astra Control enables data protection, disaster recovery, and migration for Kubernetes workloads leveraging NetApp's industry-leading technology for snapshots, backups, replication, and cloning.

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 to 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

**P myblog-55ffd9f658-tkbfq** Running

Details Metrics **YAML** Environment Logs Events Terminal

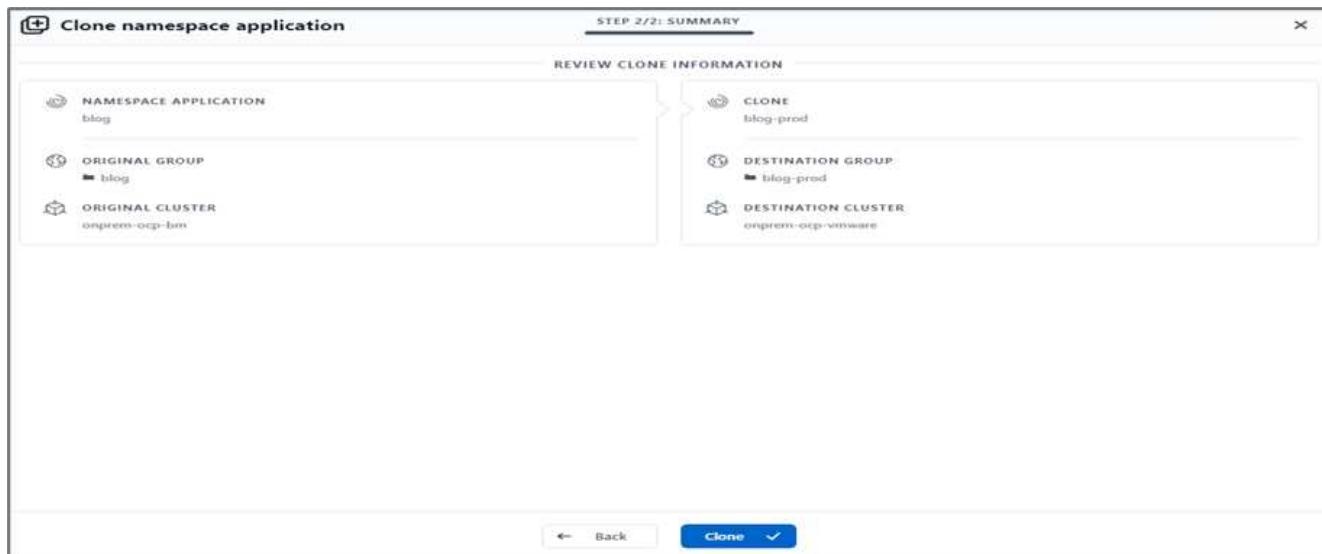
```
180
181   ports:
182     - containerPort: 2368
183     protocol: TCP
184   imagePullPolicy: Always
185   volumeMounts:
186     - name: content
187       mountPath: /var/lib/ghost/content
188     - name: kube-api-access-t2sdz
189       readOnly: true
190       mountPath: /var/run/secrets/kubernetes.io/serviceaccount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

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

The screenshot shows a web browser window with the following details:

- Address bar: Not secure: myblog-route-blog.apps.ocp.flexpod.netapp.com/astra-control-2/
- Header: Astra (with a logo), a user profile icon, and a "Subscribe" button.
- Section: Astra Control
- Author: Abhinav Singh (with a profile picture), Jun 11, 2022.
- Text: Astra Control is an application-aware data protection and mobility 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 technology for snapshots, backups, replication, and cloning.
- Call-to-action: Sign up for more like this. (with an input field for email and a "Subscribe" button).

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



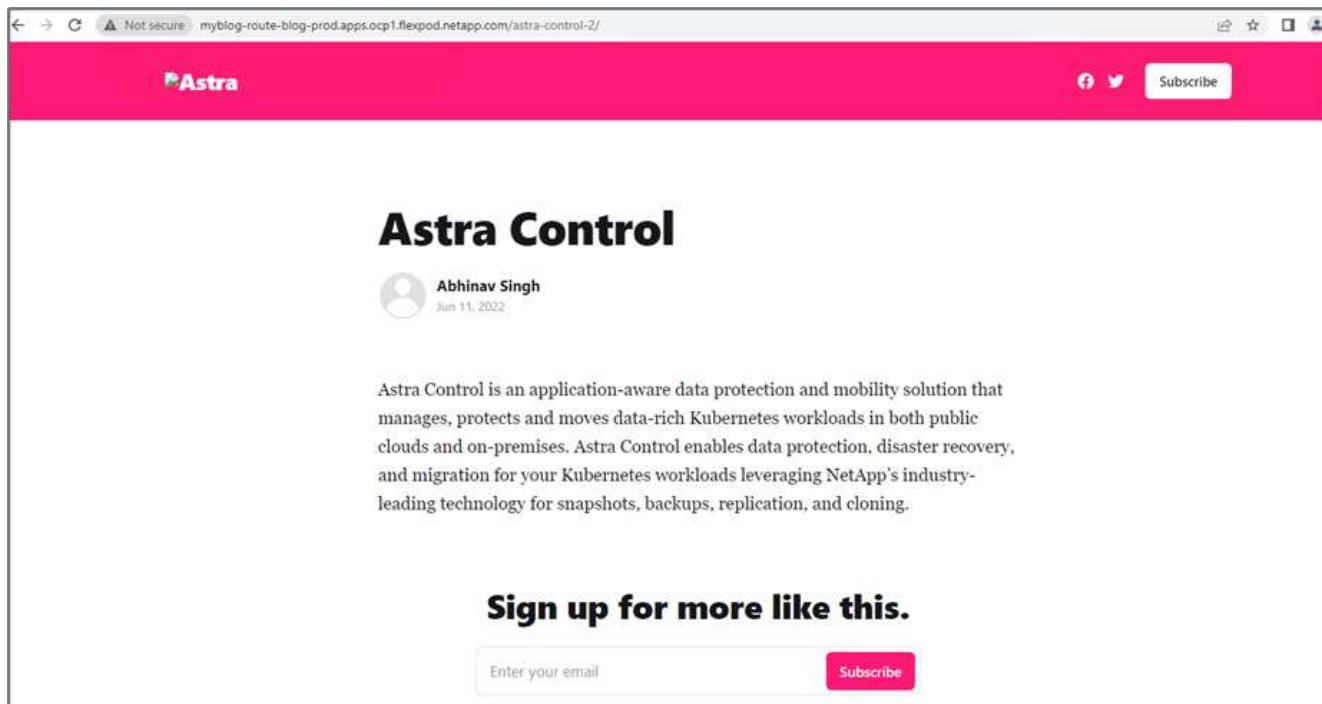
A new application clone is now provisioned in the production OpenShift cluster.

Name	State	Protection	Cluster	Group	Discovered	Actions
blog-prod	Provisioning	Unprotected	onprem-ocp-vmware	blog-prod	2022/06/11 13:17 UTC	⋮
blog	Healthy	Partially protected	onprem-ocp-bm	blog	2022/06/11 12:34 UTC	⋮

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

Name	Display name	Status	Requester	Memory	CPU	Created
PR blog-prod	No display name	Active	No requester	-	-	Just now

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



The screenshot shows a web browser window with the URL <https://myblog-route-blog-prod.apps.ocp1.flexpod.netapp.com/astra-control-2/>. The page has a pink header with the Astra logo and navigation icons. The main content area features a large title 'Astra Control' in bold black text. Below it is a user profile for 'Abhinav Singh' with a small profile picture, followed by the date 'Jun 11, 2022'. A detailed description of Astra Control follows, mentioning its application-aware data protection and mobility features. At the bottom, there is a call-to-action section with the text 'Sign up for more like this.' and a 'Subscribe' button.

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.

# NetApp Cloud Insights for FlexPod

## TR-4868: NetApp Cloud Insights for FlexPod

Alan Cowles, NetApp

In partnership with:



The solution detailed in this technical report is the configuration of the NetApp Cloud Insights service to monitor the NetApp AFF A800 storage system running NetApp ONTAP, which is deployed as a part of a FlexPod Datacenter solution.

### Customer value

The solution detailed here provides value to customers who are interested in a fully-featured monitoring solution for their hybrid cloud environments, where ONTAP is deployed as the primary storage system. This includes FlexPod environments that use NetApp AFF and FAS storage systems.

### Use cases

This solution applies to the following use cases:

- Organizations that want to monitor various resources and utilization in their ONTAP storage system deployed as part of a FlexPod solution.
- Organizations that want to troubleshoot issues and shorten resolution time for incidents that occur in their FlexPod solution with their AFF or FAS systems.
- Organizations interested in cost optimization projections, including customized dashboards to provide detailed information about wasted resources, and where cost savings can be realized in their FlexPod environment, including ONTAP.

## Target audience

The target audience for the solution includes the following groups:

- IT executives and those concerned with cost optimization and business continuity.
- Solutions architects with an interest in data center or hybrid cloud design and management.
- Technical support engineers responsible for troubleshooting and incident resolution.

You can configure Cloud Insights to provide several useful types of data that you can use to assist with planning, troubleshooting, maintenance, and ensuring business continuity. By monitoring the FlexPod Datacenter solution with Cloud Insights and presenting the aggregated data in easily digestible customized dashboards; it is not only possible to predict when resources in a deployment might need to be scaled to meet demands, but also to identify specific applications or storage volumes that are causing problems within the system. This helps to ensure that the infrastructure being monitored is predictable and performs according to expectations, allowing an organization to deliver on defined SLA's and to scale infrastructure as needed, eliminating waste and additional costs.

## Architecture

In this section, we review the architecture of a FlexPod Datacenter converged infrastructure, including a NetApp AFF A800 system that is monitored by Cloud Insights.

### Solution technology

A FlexPod Datacenter solution consists of the following minimum components to provide a highly available, easily scalable, validated, and supported converged infrastructure environment.

- Two NetApp ONTAP storage nodes (one HA pair)
- Two Cisco Nexus data center network switches
- Two Cisco MDS fabric switches (optional for FC deployments)
- Two Cisco UCS fabric interconnects
- One Cisco UCS blade chassis with two Cisco UCS B-series blade servers

Or

- Two Cisco UCS C-Series rackmount servers

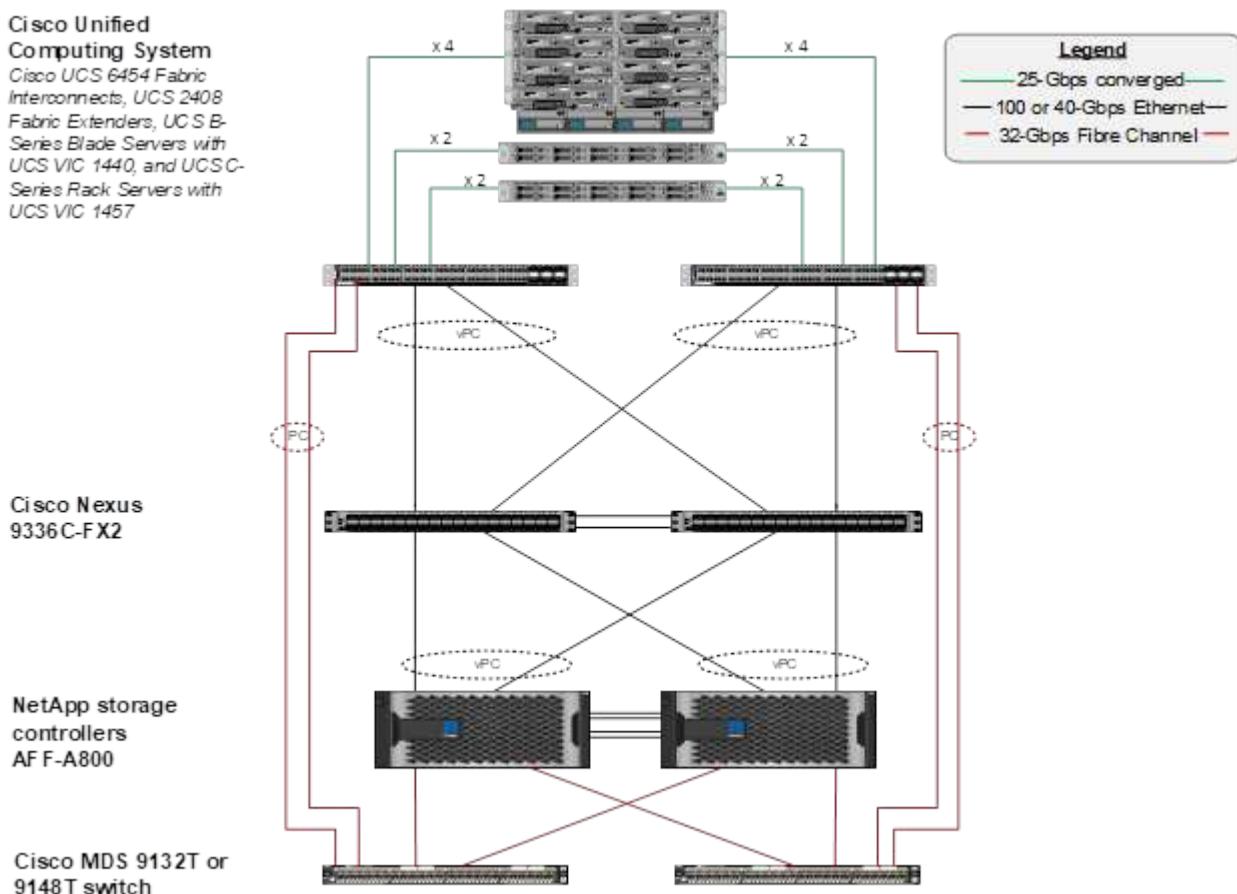
For Cloud Insights to collect data, an organization must deploy an Acquisition Unit as a virtual or physical machine either within their FlexPod Datacenter environment, or in a location where it can contact the components from which it is collecting data. You can install the Acquisition Unit software on a system running several supported Windows or Linux operating systems. The following table lists solution components for this software.

Operating system	Version
Microsoft Windows	10
Microsoft Windows Server	2012, 2012 R2, 2016, 2019
Red Hat Enterprise Linux	7.2 – 7.6
CentOS	7.2 – 7.6

Operating system	Version
Oracle Enterprise Linux	7.5
Debian	9
Ubuntu	18.04 LTS

## Architectural diagram

The following figure shows the solution architecture.



## Hardware requirements

The following table lists the hardware components that are required to implement the solution. The hardware components that are used in any particular implementation of the solution might vary based on customer requirements.

Hardware	Quantity
Cisco Nexus 9336C-FX2	2
Cisco UCS 6454 Fabric Interconnect	2
Cisco UCS 5108 Blade Chassis	1
Cisco UCS 2408 Fabric Extenders	2
Cisco UCS B200 M5 Blades	2

Hardware	Quantity
NetApp AFF A800	2

## Software requirements

The following table lists the software components that are required to implement the solution. The software components that are used in any particular implementation of the solution might vary based on customer requirements.

Software	Version
Cisco Nexus Firmware	9.3(5)
Cisco UCS Version	4.1(2a)
NetApp ONTAP Version	9.7
NetApp Cloud Insights Version	September 2020, Basic
Red Hat Enterprise Linux	7.6
VMware vSphere	6.7U3

## Use case details

This solution applies to the following use cases:

- Analyzing the environment with data provided to NetApp Active IQ digital advisor for assessment of storage system risks and recommendations for storage optimization.
- Troubleshooting problems in the ONTAP storage system deployed in a FlexPod Datacenter solution by examining system statistics in real-time.
- Generating customized dashboards to easily monitor specific points of interest for ONTAP storage systems deployed in a FlexPod Datacenter converged infrastructure.

## Design considerations

The FlexPod Datacenter solution is a converged infrastructure designed by Cisco and NetApp to provide a dynamic, highly available, and scalable data center environment for the running of enterprise workloads. Compute and networking resources in the solution are provided by Cisco UCS and Nexus products, and the storage resources are provided by the ONTAP storage system. The solution design is enhanced on a regular basis, when updated hardware models or software and firmware versions become available. These details, along with best practices for solution design and deployment, are captured in Cisco Validated Design (CVD) or NetApp Verified Architecture (NVA) documents and published regularly.

The latest CVD document detailing the FlexPod Datacenter solution design is available [here](#).

## Deploy Cloud Insights for FlexPod

To deploy the solution, you must complete the following tasks:

1. Sign up for the Cloud Insights service
2. Create a VMware virtual machine (VM) to configure as an Acquisition Unit
3. Install the Red Hat Enterprise Linux (RHEL) host
4. Create an Acquisition Unit instance in the Cloud Insights Portal and install the software
5. Add the monitored storage system from the FlexPod Datacenter to Cloud Insights.

## Sign up for the NetApp Cloud Insights service

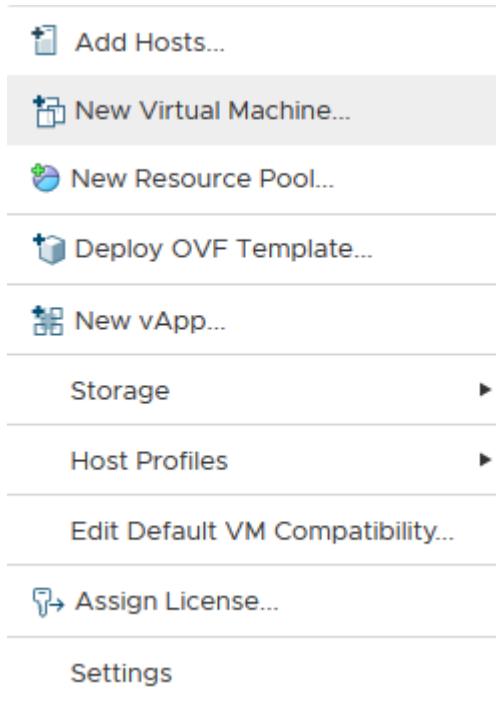
To sign up for the NetApp Cloud Insights Service, complete the following steps:

1. Go to <https://cloud.netapp.com/cloud-insights>
2. Click the button in the center of the screen to start the 14-day free trial, or the link in the upper right corner to sign up or log in with an existing NetApp Cloud Central account.

## Create a VMware virtual machine to configure as an acquisition unit

To create a VMware VM to configure as an acquisition unit, complete the following steps:

1. Launch a web browser and log in to VMware vSphere and select the cluster you want to host a VM.
2. Right-click that cluster and select Create A Virtual Machine from the menu.



3. In the New Virtual Machine wizard, click Next.
4. Specify the name of the VM and select the data center that you want to install it to, then click Next.
5. On the following page, select the cluster, nodes, or resource group you would like to install the VM to, then click Next.
6. Select the shared datastore that hosts your VMs and click Next.
7. Confirm the compatibility mode for the VM is set to ESXi 6.7 or later and click Next.

8. Select Guest OS Family Linux, Guest OS Version: Red Hat Enterprise Linux 7 (64-bit).

**Select a guest OS**

Choose the guest OS that will be installed on the virtual machine

Identifying the guest operating system here allows the wizard to provide the appropriate defaults for the operating system installation.

Guest OS Family:

Guest OS Version:

Compatibility: ESXi 6.7 and later (VM version 14)

[CANCEL](#)

[BACK](#)

[NEXT](#)

9. The next page allows for the customization of hardware resources on the VM. The Cloud Insights Acquisition Unit requires the following resources. After the resources are selected, click Next:

- Two CPUs
- 8GB of RAM
- 100GB of hard disk space
- A network that can reach resources in the FlexPod Datacenter and the Cloud Insights server through an SSL connection on port 443.
- An ISO image of the chosen Linux distribution (Red Hat Enterprise Linux) to boot from.

## Customize hardware

Configure the virtual machine hardware

Virtual Hardware

VM Options

[ADD NEW DEVICE](#)

> CPU *	2	▼	<a href="#">i</a>
> Memory *	8	▼	GB ▼
> New Hard disk *	100	GB	▼
> New SCSI controller *	VMware Paravirtual		
> New Network *	VM_Network	▼	<input checked="" type="checkbox"/> Connect...
> New CD/DVD Drive *	Datastore ISO File		
> Video card *	Specify custom settings		
VMCI device	Device on the virtual machine PCI bus that provides support for the virtual machine communication interface		

Compatibility: ESXi 6.7 and later (VM version 14)

[CANCEL](#)

[BACK](#)

[NEXT](#)

10. To create the VM, on the Ready to Complete page, review the settings and click Finish.

## Install Red Hat Enterprise Linux

To install Red Hat Enterprise Linux, complete the following steps:

1. Power on the VM, click the window to launch the virtual console, and then select the option to Install Red Hat Enterprise Linux 7.6.

## Red Hat Enterprise Linux 7.6

Install Red Hat Enterprise Linux 7.6

Test this media & install Red Hat Enterprise Linux 7.6

Troubleshooting >

Press Tab for full configuration options on menu items.

2. Select the preferred language and click Continue.

The next page is Installation Summary. The default settings should be acceptable for most of these options.

3. You must customize the storage layout by performing the following options:
  - a. To customize the partitioning for the server, click Installation Destination.
  - b. Confirm that the VMware Virtual Disk of 100GiB is selected with a black check mark and select the I Will Configure Partitioning radio button.

## Device Selection

Select the device(s) you'd like to install to. They will be left untouched until you click on the main menu's "Begin Installation" button.

### Local Standard Disks

100 GiB



VMware Virtual disk

sda / 100 GiB free

*Disks left unselected here will not be touched.*

### Specialized & Network Disks



Add a disk...

*Disks left unselected here will not be touched.*

## Other Storage Options

### Partitioning

Automatically configure partitioning.  I will configure partitioning.

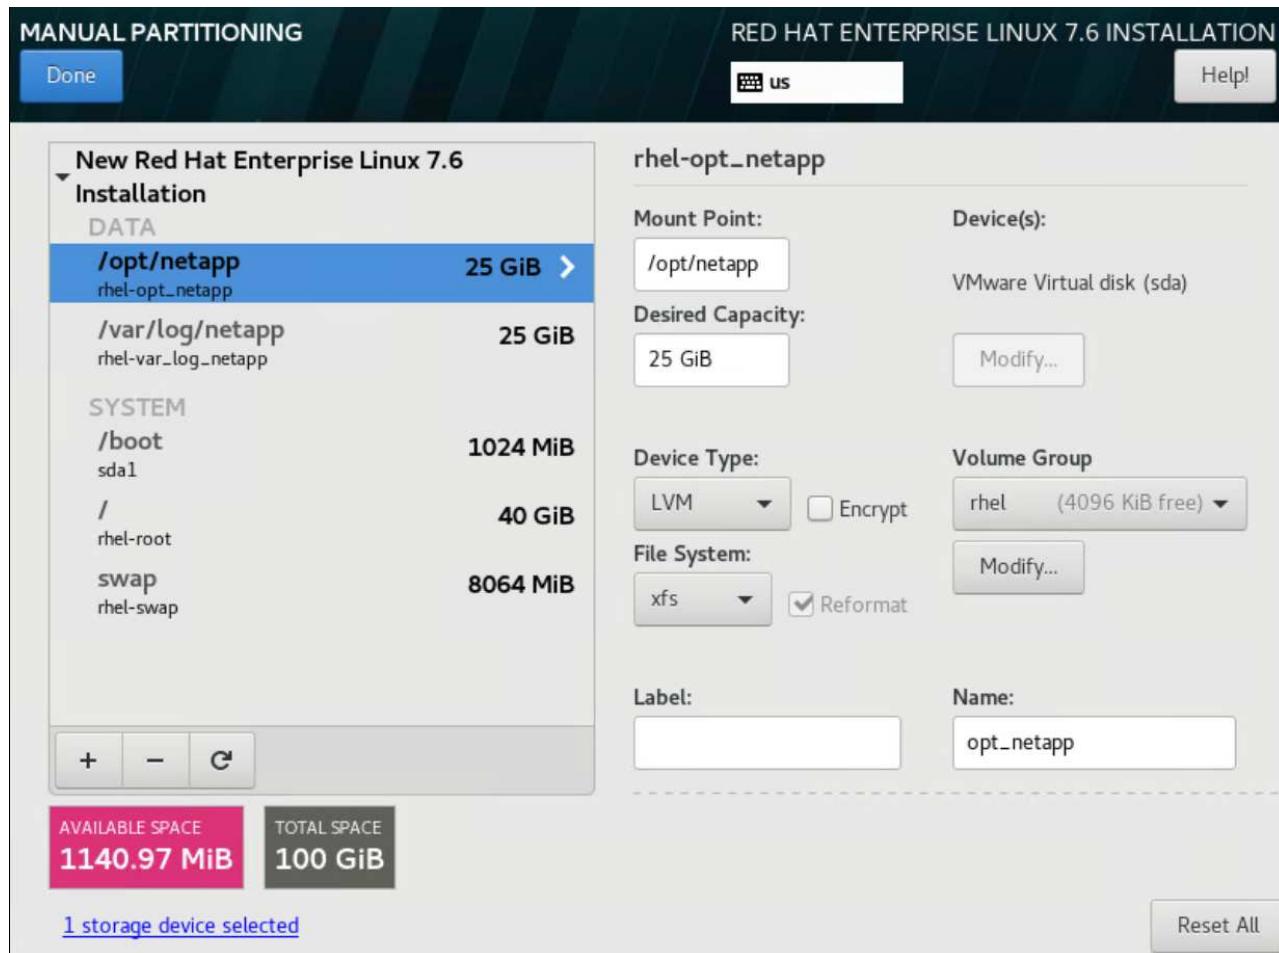
I would like to make additional space available.

[Full disk summary and boot loader...](#)

1 disk selected; 100 GiB capacity; 100 GiB free [Refresh...](#)

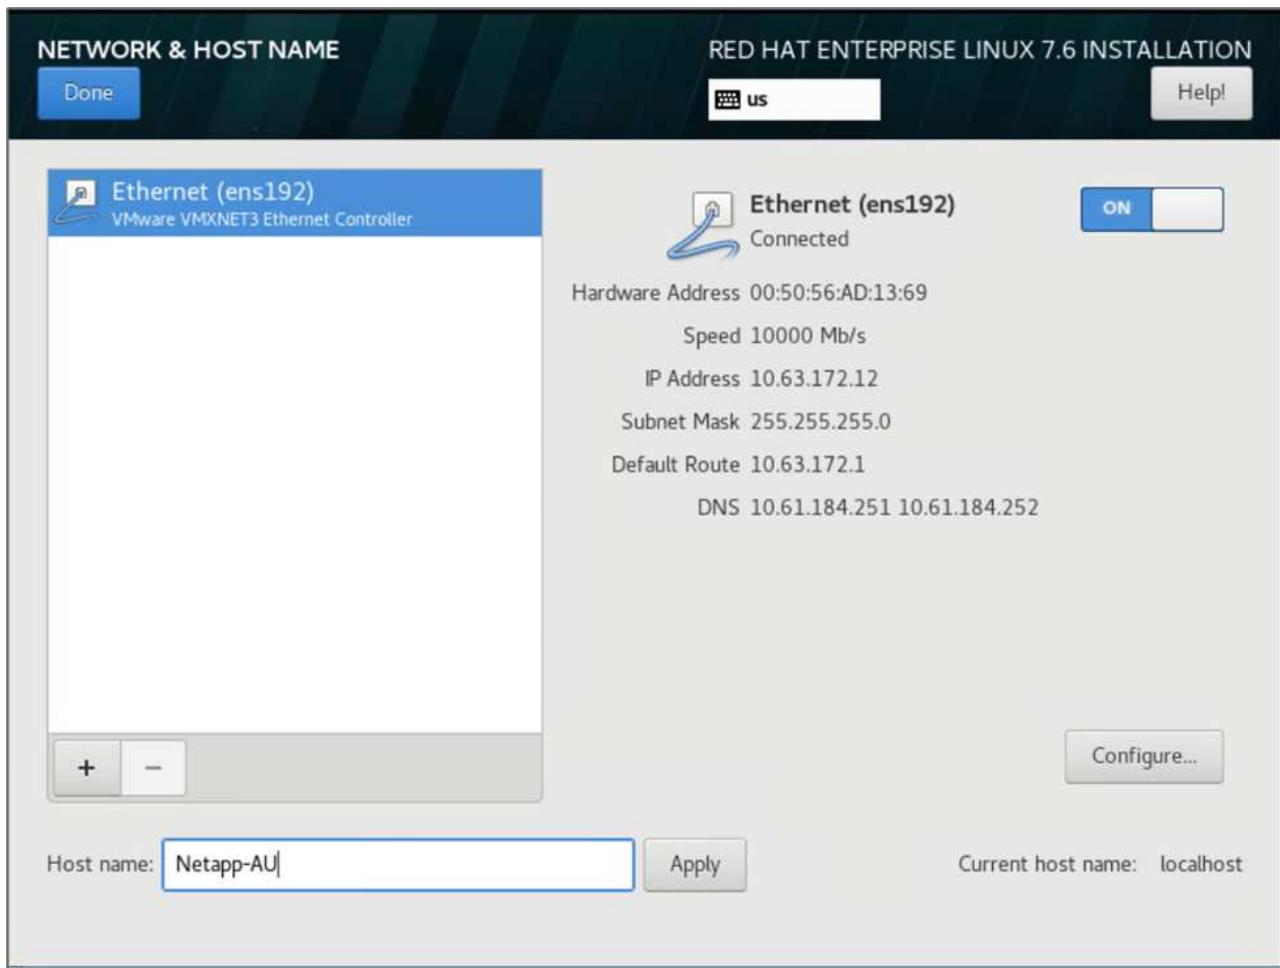
c. Click Done.

A new menu displays enabling you to customize the partition table. Dedicate 25 GB each to /opt/netapp and /var/log/netapp. You can automatically allocate the rest of the storage to the system.



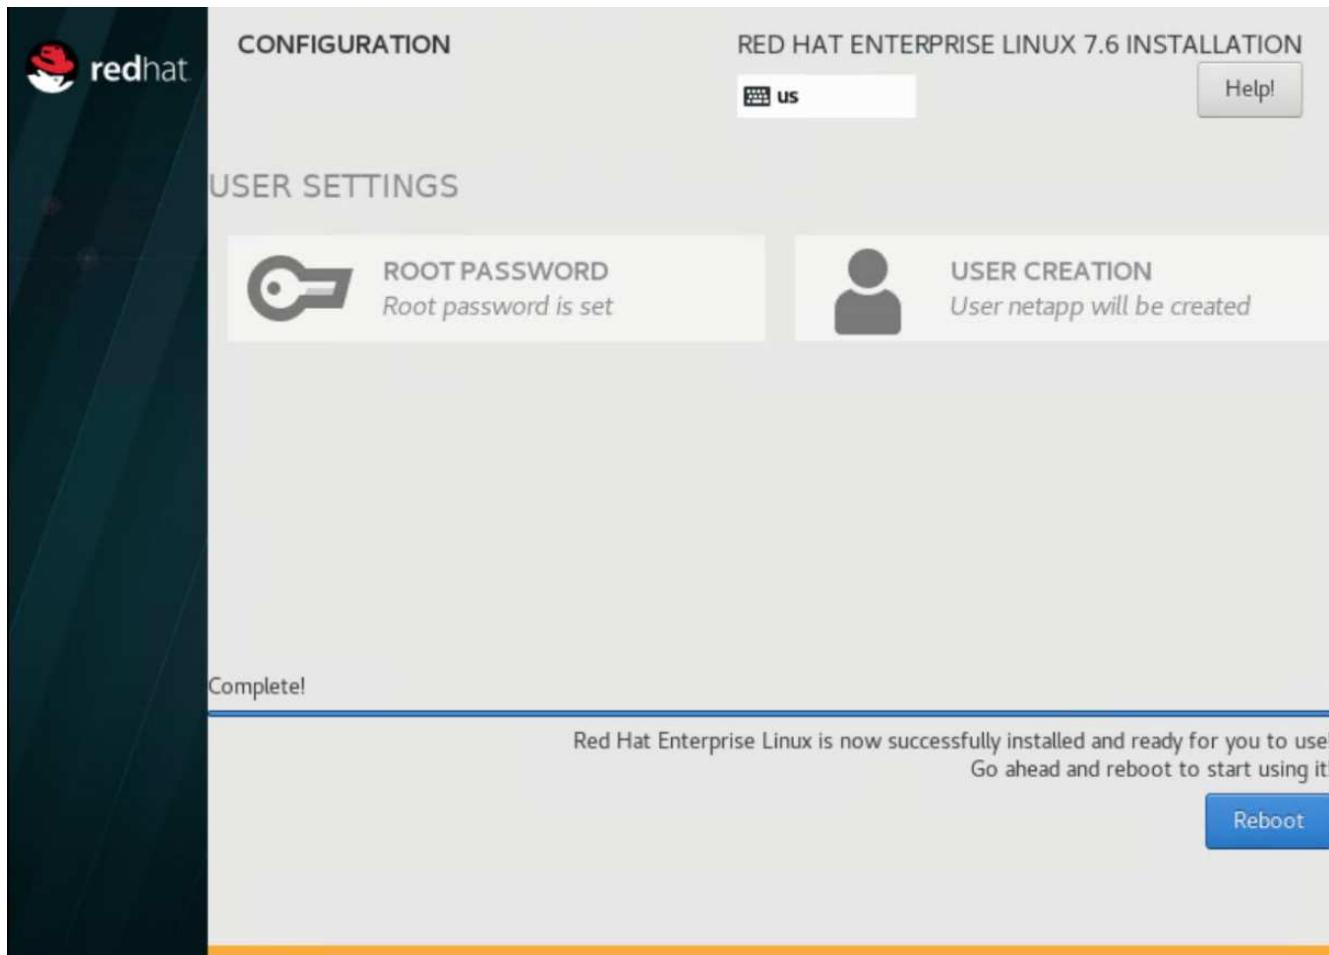
- d. To return to Installation Summary, click Done.

4. Click Network and Host Name.
  - a. Enter a host name for the server.
  - b. Turn on the network adapter by clicking the slider button. If Dynamic Host Configuration Protocol (DHCP) is configured on your network, you will receive an IP address. If it is not, click Configure, and manually assign an address.



c. . Click Done to return to Installation Summary.

5. On the Installation Summary page, click Begin Installation.
6. On the Installation Progress page, you can set the root password or create a local user account. When the installation finishes, click Reboot to restart the server.



7. After the system has rebooted, log in to your server and register it with Red Hat Subscription Manager.

```
[root@Netapp-AU ~]# subscription-manager register
Registering to: subscription.rhsm.redhat.com:443/subscription
Username: alan.cowles@netapp.com
Password:
The system has been registered with ID: a47f2e7b-81cd-4757-85c7-eb1818c2c2a1
The registered system name is: Netapp-AU
[root@Netapp-AU ~]#
```

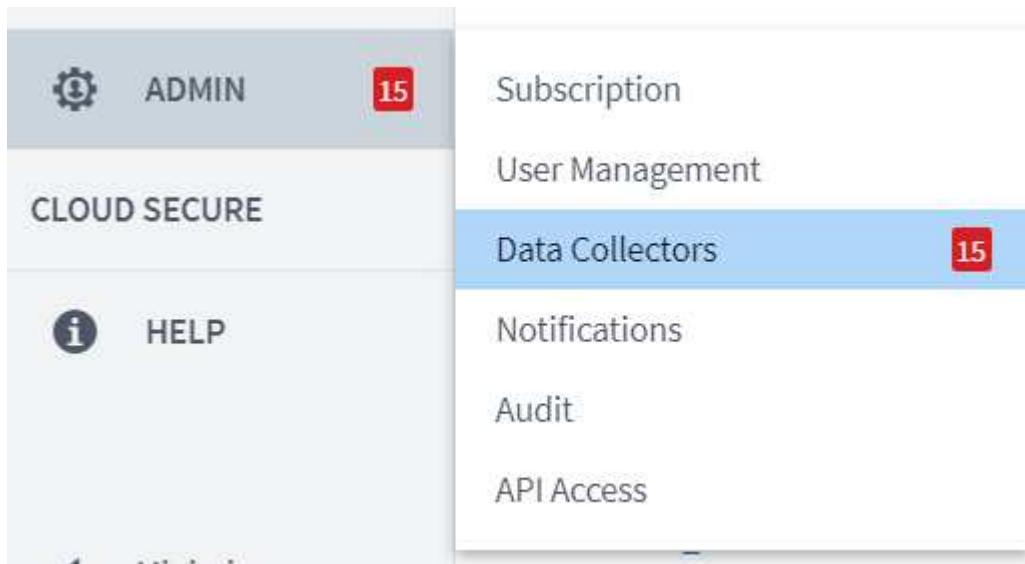
8. Attach an available subscription for Red Hat Enterprise Linux.

```
[root@Netapp-AU ~]# subscription-manager attach --pool=8a85f99b710f3b1901713b90b9e154cf
Successfully attached a subscription for: Red Hat Enterprise Linux, Standard Support (128 Sockets, NFR, Partner Only)
[root@Netapp-AU ~]#
```

#### Create an acquisition unit instance in the Cloud Insights portal and install the software

To create an acquisition unit instance in the Cloud Insights portal and install the software, complete the following steps:

1. From the home page of Cloud Insights, hover over the Admin entry in the main menu to the left and select Data Collectors from the menu.



2. In the top center of the Data Collectors page, click the link for Acquisition Units.

[Data Collectors ! 9](#) [Acquisition Units ! 7](#)

3. To create a new Acquisition Unit, click the button on the right.

[+ Acquisition Unit](#)

Filter...

4. Select the operating system that you want to use to host your Acquisition Unit and follow the steps to copy the installation script from the web page.

In this example, it is a Linux server, which provides a snippet and a token to paste into the CLI on our host. The web page waits for the Acquisition Unit to connect.

## Install Acquisition Unit

Cloud Insights collects device data via one or more Acquisition Units installed on local servers. Each Acquisition Unit can host multiple Data Collectors, which send device metrics to Cloud Insights for analysis.

## What Operating System or Platform Are You Using?

## Installation Instructions

Need Help?

1 Copy Installer Snippet

*This snippet has a unique key valid for 24 hours for this Acquisition Unit only.*

 Reveal Installer Snippet

token=eyJ0eXAiOiJKV1QiLCJhbGciOiJIUzI4NCT9.eyJhdWxvZ2luVXJsIjoiaHR0cHM6Ly9hdWxvZ2luLmMwMS5jbG91ZGlucl2lnaHRzLm5ldGFWcC5jb20iLCJvbW0aW1lVG9rZW5jZC1i6jYdKYZi1ZWNjLwU5MjctNDQ4YS05NmVhLtlNDjKyliWnOGI2MyIsInjvbGViJpbImFjCvXpC2l0aw9uX3NpZ25lci1dLCjZXJZ2XJZCvcmwiOjIdHrwczovLzhkNDE5NWE2LWVbIYjtgNGFkMC1hNmVLTQzXMaYzQwYjVhZi5jMDEUY2xvdwRpbnnPz2h0cy5uZXRhchauY29Ti1xiawXNzIjoib2NpIiwiZWhxjw0nNjAyMdk5Mja2LCJsb2pdBi6jMfjCvXpC2l0aw9uJjAzOTI0MD1yLTg2YzQntNDJmMC1hZDjmlWfHjYME1ZDEzMzhMiiSm1hdC16MTWjMjAxMjic0niwidXVpZC16jJjAzOTI0MD1yLTg2YzQntNDjmc1

2 Paste the snippet into a bash shell to run the installer.

3 Please ensure you have copied and pasted the snippet into the bash shell.

5. Paste the snippet into the CLI of the Red Hat Enterprise Linux machine that was provisioned and click Enter.

```
[root@Netapp-AU ~]# token=eyJ0eXAiOiJKV1QiLCJhbGciOiJIUzIwMjAyMDIyNjIzGluc2lnaHRzLm5ldGFwcC5jb20iLCJvbVmV0aW1lVG9rZW5jZC1iJq5ZTY0MGM5LTY5MTItNDQ4Yi04YmI4LTlTwNGY20TQ2YzY1YStsInJvbGVzIjpbImFjcXVpc2l0aw9uX3NpZ251ciJdLCjzZXJ2ZXJvcmwi0iJodHrwczovLzhkNDE5NWE2LWlV1jgtNGFkMC1hNmVmlTQxMzAyMzQwvJvhZi5jMDEuY2xvdWRpbnNpZ2h0cy5uZXRhcaHwY29tLiwiob2NpIiwiZxHwIjoxNjAyMtgz2LCJsb2dpbiI6ImFjcXVpc2l0aw9uLjci4MT1zGI3LTk50WqtNGNiYs05YmU1LTmWzTcxJzK00DRiZCIsImlhIdC16MTYwMjA5NjMyNiwidXvpZC1iJc4MT1zGI3LTk50WqtNGNiYs05YmU1LTmWzTcxJzk00DRiZCIsInRlwmFudC16IjhkNDE5NWE2LWlV1jgtNGFkMC1hNmVmlTQxMzAyMzQwvJvhZiIsInRlwmFudF1YmRvbWFpbIi6InBzMTMyNsJ9.RvWLR3wH1_k6fI0Ci0_h-Wok2STFPDj7VksmXqW-GZ-JqSiE8SZE4v3DuWrWM6 domainUrl=https://8d4195a6-ebb8-4d04-a6ef-41302340b5af.c01.cloudinsights.netapp.com/rest/v1/au version=1.253.0 bootstrap=cloudinsights-au-install-bootstrap.sh && curl $proxy_auth_scheme -H "Authorization: Bearer $token" -o $bootstrap $domainUrl/installerBootstrap && sudo chmod 755 $bootstrap && sudo /bin/bash -c "TOKEN=$token HTTPS_PROXY=$https_proxy PROXY_AUTH_SCHEME=$proxy_auth_scheme AU_VERSION=$version INSTALLER_NAME=cloudinsights-linux-au-installer-$version INSTALLER_URL=$domainUrl/installers/linux/$version ./bootstrap"
```

The installation program downloads a compressed package and begins the installation. When the installation is complete, you receive a message stating that the Acquisition Unit has been registered with NetApp Cloud Insights.

```

Welcome to CloudInsights (R) ...
Acquisition Unit

NetApp (R)
Installation: /opt/netapp/cloudinsights
Logs:          /opt/netapp/cloudinsights/logs -> /var/log/netapp/cloudinsights

To control the CloudInsights service:
  sudo cloudinsights-service.sh --help
To uninstall:
  sudo cloudinsights-uninstall.sh --help

1/8 Acquisition Unit Starting
2/8 Connecting to Cloud Insights
3/8 Sending Certificate-Signing Request..
4/8 Logging in to Cloud Insights
5/8 Updating Security Settings..
6/8 Downloading Data Collection Modules
7/8 Registering to Cloud Insights
8/8 Acquisition Unit Ready

Acquisition Unit has been installed successfully.
[root@Netapp-AU ~]#

```

## Add the monitored storage system from the FlexPod Datacenter to Cloud Insights

To add the ONTAP storage system from a FlexPod deployment, complete the following steps:

1. Return to the Acquisition Units page on Cloud Insights portal and find the listed newly registered unit. To display a summary of the unit, click the unit.

Name	IP	Status	Last Reported	Note
NetApp-AU	10.1.156.115	OK	9 minutes ago	

2. To start a wizard to add the storage system, on the Summary page, click the button for creating a data collector. The first page displays all the systems from which data can be collected. Use the search bar to search for ONTAP.

Choose a Data Collector to Monitor

NetApp  
Cloud Volumes ONTAP

NetApp  
Data ONTAP 7-Mode

NetApp  
ONTAP Data Management Software

NetApp  
ONTAP Select

### 3. Select ONTAP Data Management Software.

A page displays that enables you to name your deployment and select the Acquisition Unit that you want to use. You can provide the connectivity information and credentials for the ONTAP system and test the connection to confirm.

Configure Collector

Add credentials and required settings

NetApp Management IP Address: 192.168.156.50

User Name: admin

Acquisition Unit: NetApp-AU

NetApp Management IP Address: 192.168.156.50

User Name: admin

Advanced Configuration

### 4. Click Complete Setup.

The portal returns to the Data Collectors page and the Data Collector begins its first poll to collect data from the ONTAP storage system in the FlexPod Datacenter.

FlexPod Datacenter	All stand-by	NetApp ONTAP Data Management Software	NetApp-AU	192.168.156.50	Polling...	⋮
--------------------	--------------	---------------------------------------	-----------	----------------	------------	---

## Use cases

With Cloud Insights set up and configured to monitor your FlexPod Datacenter solution,

we can explore some of the tasks that you can perform on the dashboard to assess and monitor your environment. In this section, we highlight five primary use cases for Cloud Insights:

- Active IQ integration
- Exploring real-time dashboards
- Creating custom dashboards
- Advanced troubleshooting
- Storage optimization

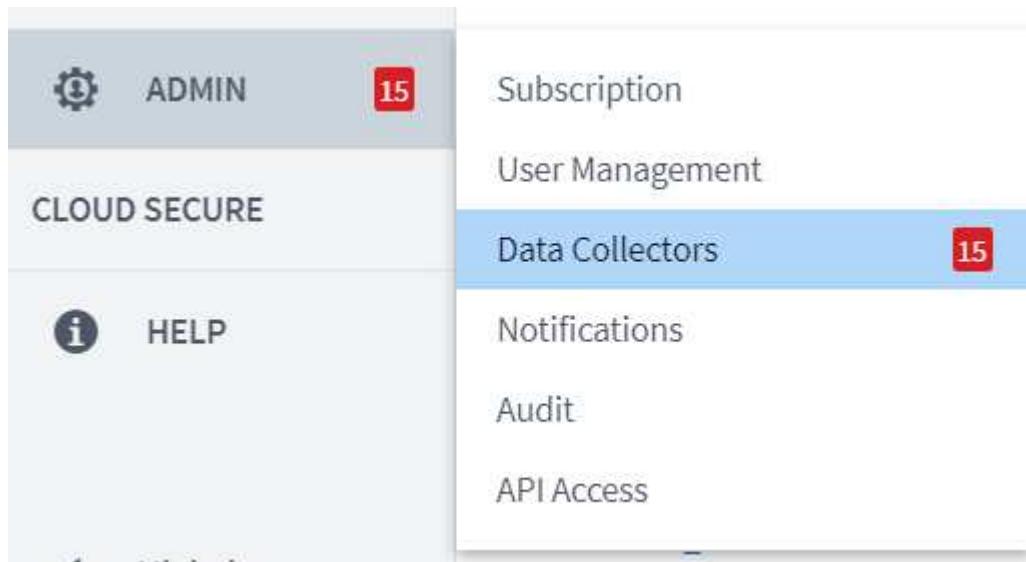
## Active IQ integration

Cloud Insights is fully integrated into the Active IQ storage monitoring platform. An ONTAP system, deployed as a part of a FlexPod Datacenter solution, is automatically configured to send information back to NetApp through the AutoSupport function, which is built into each system. These reports are generated on a scheduled basis, or dynamically whenever a fault is detected in the system. The data communicated through AutoSupport is aggregated and displayed in easily accessible dashboards under the Active IQ menu in Cloud Insights.

### Access Active IQ information through the Cloud Insights dashboard

To access the Active IQ information through the Cloud Insights dashboard, complete the following steps:

1. Click the Data Collector option under the Admin menu on the left.



2. Filter for the specific Data Collector in your environment. In this example, we filter by the term FlexPod.

A screenshot of the Data Collectors list page. The URL in the address bar is 'NetApp PCS Sa... / Admin / Data Collectors'. The page header shows 'Data Collectors 0/8' and 'Acquisition Units 0/8'. A search bar at the top right contains the text 'FlexPod'. The main table has a header: 'Data Collectors (1)' and columns: 'Name', 'Status', 'Type', 'Acquisition Unit', 'IP', 'Impact ↓', 'Last Acquired', and a 'More' (three dots) column. There is one entry in the table: 'FlexPod Datacenter' with status 'All successful', type 'NetApp ONTAP Data Management Software', acquisition unit 'NetApp-AU', IP '192.168.156.50', impact 'Low', last acquired '10 minutes ago', and a 'More' icon.

3. Click the Data Collector to get a summary of the environment and devices that are being monitored by that collector.

NetApp PCS Sa... / Admin / Data Collectors / Installed / **FlexPod Datacenter**

**Edit** ▾

**Summary**

Name	Type	Types of Data Collected	Performance Recent Status	Note
FlexPod Datacenter	NetApp ONTAP Data Management Software	Inventory, Performance	Success	
Acquisition Unit	NetApp-AU	Inventory Recent Status	Success	

**Event Timeline (Last 3 Weeks)**

Inventory	Performance
3 Weeks Ago	2 Weeks Ago
1 Week Ago	

**Inventory** 10/15/2020 1:51:42 PM - 10/19/2020 11:42:15 AM

**Devices Reported by This Collector (1)**

Device ↑	Name	IP
<span style="color: orange;">S</span> Storage	aa14-a800	192.168.156.50

**Show Recent Changes**

Under the device list near the bottom, click on the name of the ONTAP storage system being monitored. This displays a dashboard of information collected about the system, including the following details:

- Model
- Family
- ONTAP Version
- Raw Capacity
- Average IOPS
- Average Latency
- Average Throughput

NetApp PCS Sa... / S aa14-a800

Last 3 Hours

Acquired 13 minutes ago, 12:51 PM

**Storage Summary**

Model: AFF-A800	IP: 192.168.156.50	IOPS - Total: 4,972.70 IO/s	Performance Policies: 5m
Vendor: NetApp	Microcode Version: 9.7.0P1 clustered Data ONTAP	Throughput - Total: 7.98 MB/s	Risks: 35 risks detected by Active IQ
Family: AFF	Raw Capacity: 43,594.6 GB	Management: HTTPS://192.168.156.50:443	FC Fabrics Connected: 0
Serial Number: 1-80-000011	Latency - Total: 0.05 ms		

**User Data**

Note: Testing annotations  
Testing rules

**Expert View**

Latency - Total (ms)

Display Metrics ▾

**Resource**

S aa14-a800

**Top Correlated**

<span style="color: orange;">S</span> aa14-a800-2	79%
<span style="color: orange;">S</span> aa14-a800-1	23%

Also, on this page under the Performance Policies section, you can find a link to NetApp Active IQ.

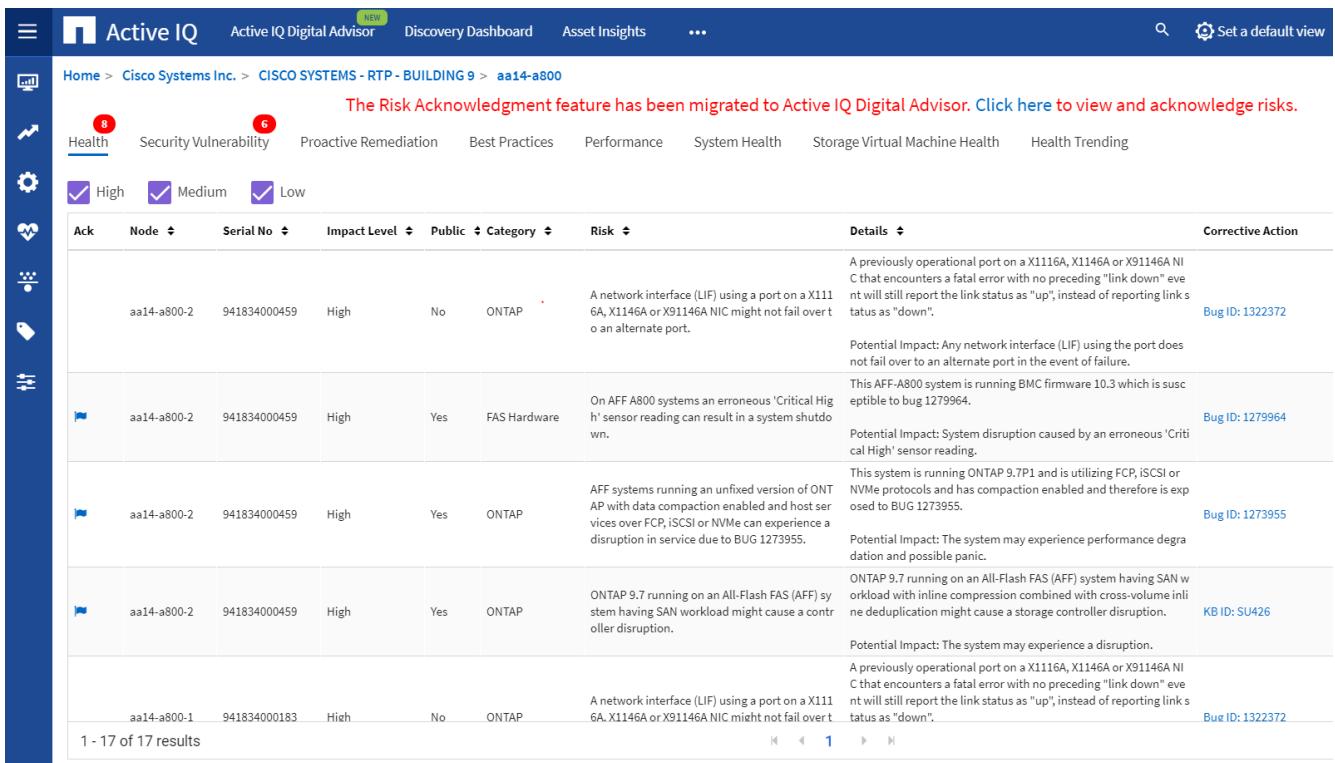
⌚ 5m

## Performance Policies:

### Risks:

❗ 35 risks detected  
by  Active IQ 

4. To open a new browser tab and take you to the risk mitigation page, which shows which nodes are affected, how critical the risks are, and what the appropriate action is that needs to be taken to correct the identified issues, click the link for Active IQ.



The screenshot shows the Active IQ Digital Advisor interface. The top navigation bar includes 'Active IQ Digital Advisor' (with a 'NEW' badge), 'Discovery Dashboard', 'Asset Insights', and a '...' button. On the far right are a search icon and a 'Set a default view' button. The main content area shows a breadcrumb path: 'Home > Cisco Systems Inc. > CISCO SYSTEMS - RTP - BUILDING 9 > aa14-a800'. A red message at the top states: 'The Risk Acknowledgment feature has been migrated to Active IQ Digital Advisor. Click here to view and acknowledge risks.' Below this, there are tabs for 'Health' (8), 'Security Vulnerability' (6), 'Proactive Remediation', 'Best Practices', 'Performance', 'System Health', 'Storage Virtual Machine Health', and 'Health Trending'. Under the 'Health' tab, there are filters for 'High', 'Medium', and 'Low' risk levels. The table below lists 17 risks for node 'aa14-a800-2'. Each row includes columns for 'Ack', 'Node', 'Serial No', 'Impact Level', 'Public', 'Category', 'Risk', 'Details', and 'Corrective Action'. The 'Details' column provides a brief description of the risk, and the 'Corrective Action' column includes a 'Bug ID' link. The table footer shows '1 - 17 of 17 results'.

Ack	Node	Serial No	Impact Level	Public	Category	Risk	Details	Corrective Action
	aa14-a800-2	941834000459	High	No	ONTAP		A previously operational port on a X1116A, X1146A or X91146A NIC that encounters a fatal error with no preceding "link down" event will still report the link status as "up", instead of reporting link status as "down".	<a href="#">Bug ID: 1322372</a>
	aa14-a800-2	941834000459	High	Yes	FAS Hardware		Potential Impact: Any network interface (LIF) using the port does not fail over to an alternate port in the event of failure.	<a href="#">Bug ID: 1279964</a>
	aa14-a800-2	941834000459	High	Yes	ONTAP		This AFF A800 system is running BMC firmware 10.3 which is susceptible to bug 1279964.	<a href="#">Bug ID: 1279964</a>
	aa14-a800-2	941834000459	High	Yes	ONTAP		Potential Impact: System disruption caused by an erroneous 'Critical High' sensor reading.	<a href="#">Bug ID: 1273955</a>
	aa14-a800-2	941834000459	High	Yes	ONTAP		This system is running ONTAP 9.7P1 and is utilizing FCP, iSCSI or NVMe protocols and has compaction enabled and therefore is exposed to BUG 1273955.	<a href="#">Bug ID: 1273955</a>
	aa14-a800-2	941834000459	High	Yes	ONTAP		Potential Impact: The system may experience performance degradation and possible panic.	<a href="#">Bug ID: SU426</a>
	aa14-a800-1	941834000183	High	No	ONTAP		ONTAP 9.7 running on an All-Flash FAS (AFF) system having SAN workload with inline compression combined with cross-volume inline deduplication might cause a storage controller disruption.	<a href="#">Bug ID: SU426</a>
	aa14-a800-1	941834000183	High	No	ONTAP		Potential Impact: The system may experience a disruption.	
	aa14-a800-1	941834000183	High	No	ONTAP		A previously operational port on a X1116A, X1146A or X91146A NIC that encounters a fatal error with no preceding "link down" event will still report the link status as "up", instead of reporting link status as "down".	<a href="#">Bug ID: 1322372</a>

### Explore real-time dashboards

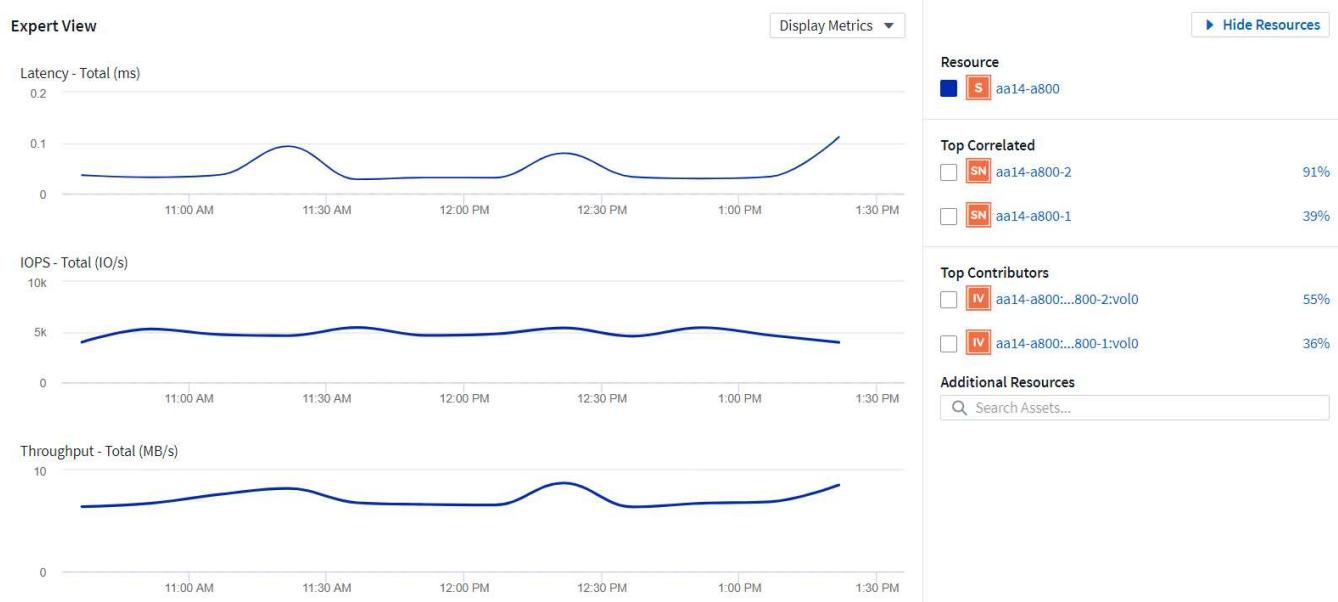
Cloud Insights can display real-time dashboards of the information that has been polled from the ONTAP storage system deployed in a FlexPod Datacenter solution. The Cloud Insights Acquisition Unit collects data in regular intervals and populates the default storage system dashboard with the information collected.

## Access real-time graphs through the Cloud Insights dashboard

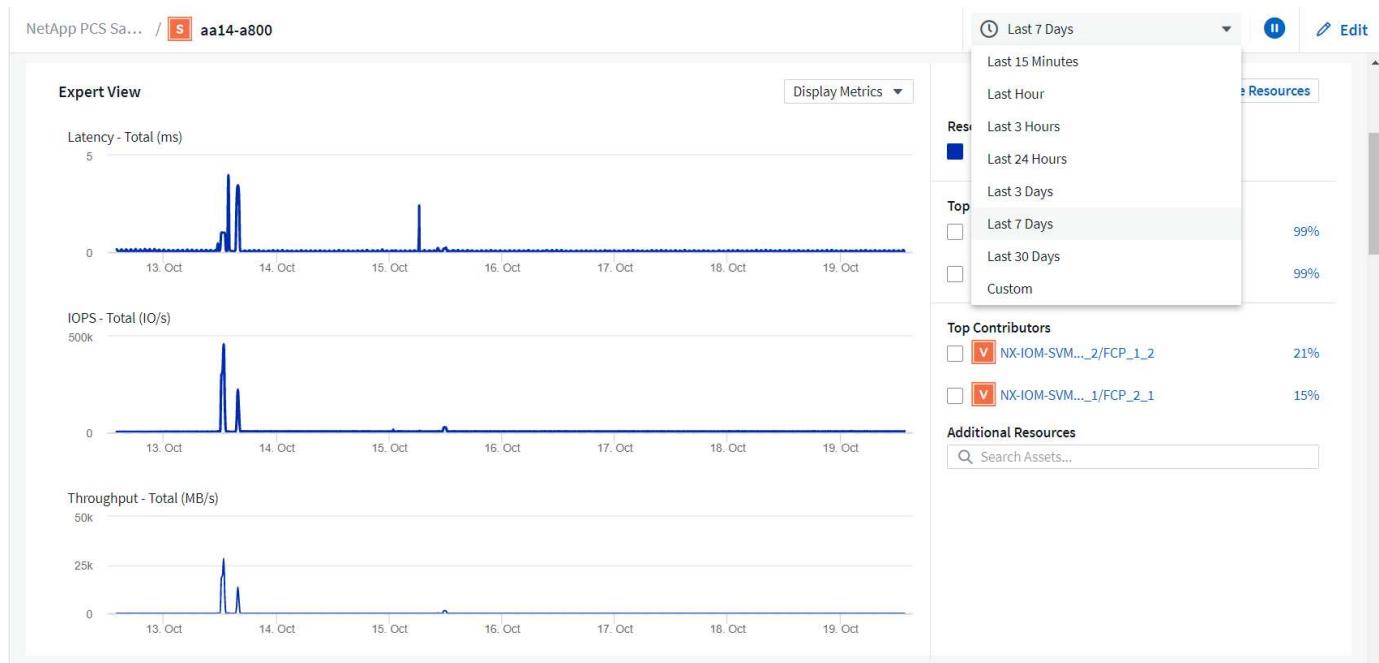
From the storage system dashboard, you can see the last time that the Data Collector updated the information. An example of this is shown in the figure below.

Acquired 3 minutes ago, 1:21 PM		
Details	X	
Data Collector	Status	Last Acquired
FlexPod Datacenter	All successful	3 minutes ago, 1:21 PM

By default, the storage system dashboard displays several interactive graphs that show system-wide metrics from the storage system being polled, or from each individual node, including: Latency, IOPS, and Throughput, in the Expert View section. Examples of these default graphs are shown in the figure below.



By default, the graphs show information from the last three hours, but you can set this to a number of differing values or a custom value from the dropdown list near the top right of the storage system dashboard. This is shown in the figure below.



## Create custom dashboards

In addition to making use of the default dashboards that display system-wide information, you can use Cloud Insights to create fully customized dashboards that enable you to focus on resource use for specific storage volumes in the FlexPod Datacenter solution, and thus the applications deployed in the converged infrastructure that depend on those volumes to run effectively. Doing so can help you to create a better visualization of specific applications and the resources they consume in the data center environment.

## Create a customized dashboard to assess storage resources

To create a customized dashboard to assess storage resources, complete the following steps:

1. To create a customized dashboard, hover over Dashboards on the Cloud Insights main menu and click + New Dashboard in the dropdown list.



MONITOR & OPTIMIZE

NetApp PCS Sa... / Admin / Da

HOME

DASHBOARDS

QUERIES

ALERTS

Summary

Show All Dashboards (1835)

+ New Dashboard

Kubernetes Explorer

The New Dashboard window opens.

2. Name the dashboard and select the type of widget used to display the data. You can select from a number of graph types or even notes or table types to present the collected data.

NetApp PCS Sa... / Dashboards / **New Dashboard**

Last 7 Days

Add Variable

Add Widget

Save

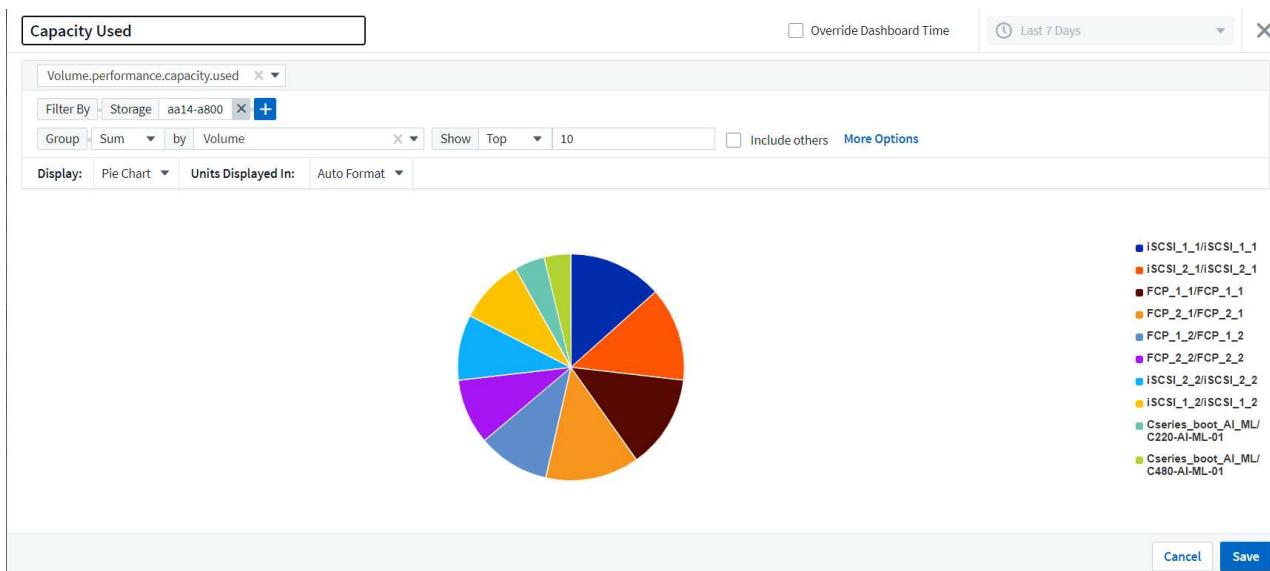
Choose Widget Type:

Line Chart	Spline Chart	Area Chart	Stacked Area Chart	Box Plot	Scatter Plot
Single Value	Solid Gauge	Bullet Gauge	Bar Chart	Column Chart	Pie Chart
Note	Table	Violations Table			

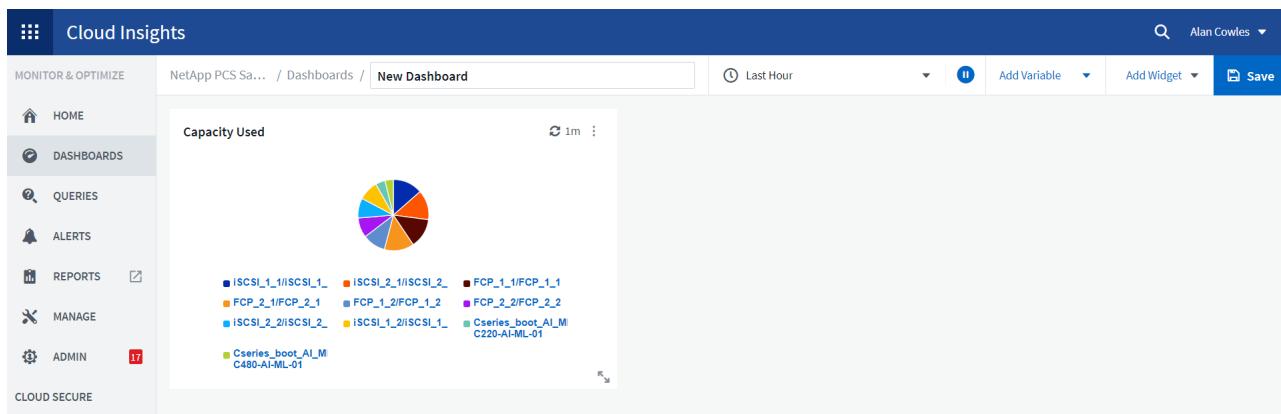
3. Choose customized variables from the Add Variable menu.

This enables the data presented to be focused to display more specific or specialized factors.

4. To create a custom dashboard, select the widget type you would like to use, for example, a pie chart to display storage utilization by volume:
  - a. Select the Pie Chart widget from the Add Widget dropdown list.
  - b. Name the widget with a descriptive identifier, such as Capacity Used.
  - c. Select the object you want to display. For example, you can search by the key term volume and select volume.performance.capacity.used.
  - d. To filter by storage systems, use the filter and type in the name of the storage system in the FlexPod Datacenter solution.
  - e. Customize the information to be displayed. By default, this selection shows ONTAP data volumes and lists the top 10.
  - f. To save the customized dashboard, click the Save.



After saving the custom widget, the browser returns to the New Dashboard page where it displays the newly created widget and allows for interactive action to be taken, such as modifying the data polling period.



## Advanced troubleshooting

Cloud Insights enables advanced troubleshooting methods to be applied to any storage environment in a FlexPod Datacenter converged infrastructure. Using components of each of the features mentioned above: Active IQ integration, default dashboards with real-time statistics, and customized dashboards, issues that might arise are detected early and solved rapidly. Using the list of risks in Active IQ, a customer can find reported configuration errors that could lead to issue or discover bugs that have been reported and patched versions of code that can remedy them. Observing the real-time dashboards on the Cloud Insights home page can help to discover patterns in system performance that could be an early indicator of a problem on the rise and help to resolve it expediently. Lastly, being able to create customized dashboards enables customers to focus on the most important assets in their infrastructure and monitor those directly to ensure that they can meet their business continuity objectives.

## Storage optimization

In addition to troubleshooting, it is possible to use the data collected by Cloud Insights to optimize the ONTAP storage system deployed in a FlexPod Datacenter converged infrastructure solution. If a volume shows a high latency, perhaps because several VMs with high performance demands are sharing the same datastore, that information is displayed on the Cloud Insights dashboard. With this information, a storage administrator can choose to migrate one or more VMs either to other volumes, migrate storage volumes between tiers of aggregates, or between nodes in the ONTAP storage system, resulting in a performance optimized environment. The information gleaned from the Active IQ integration with Cloud Insights can highlight configuration issues that lead to poorer than expected performance, and provide the recommended corrective action that if implemented, can remediate any issues, and ensure an optimally tuned storage system.

## Videos and demos

You can see a video demonstration of using NetApp Cloud Insights to assess the resources in an on-premises environment [here](#).

You can see a video demonstration of using NetApp Cloud Insights to monitor infrastructure and set alert thresholds for infrastructure [here](#).

You can see a video demonstration of using NetApp Cloud Insights to assess individual applications in the environment [here](#).

## Additional information

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

- Cisco Product Documentation  
<https://www.cisco.com/c/en/us/support/index.html>
- FlexPod Datacenter  
<https://www.flexpod.com>
- NetApp Cloud Insights  
<https://cloud.netapp.com/cloud-insights>
- NetApp Product Documentation  
<https://docs.netapp.com>

## FlexPod with FabricPool - Inactive Data Tiering to Amazon AWS S3

### TR-4801: FlexPod with FabricPool - Inactive Data Tiering to Amazon AWS S3

Scott Kovacs, NetApp

Flash storage prices continue to fall, making it available to workloads and applications that were not previously considered candidates for flash storage. However, making the most efficient use of the storage investment is still critically important for IT managers. IT departments continue to be pressed to deliver higher-performing services with little or no budget increase. To help address these needs, NetApp FabricPool allows you to leverage cloud economics by moving infrequently used data off of expensive on-premises flash storage to a more cost-effective storage tier in the public cloud. Moving infrequently accessed data to the cloud frees up valuable flash storage space on AFF or FAS systems to deliver more capacity for business-critical workloads to the high-performance flash tier.

This technical report reviews the FabricPool data- tiering feature of NetApp ONTAP in the context of a FlexPod converged infrastructure architecture from NetApp and Cisco. You should be familiar with the FlexPod Datacenter converged infrastructure architecture and the ONTAP storage software to fully benefit from the concepts discussed in this technical report. Building on familiarity with FlexPod and ONTAP, we discuss FabricPool, how it works, and how it can be used to achieve more efficient use of on-premises flash storage. Much of the content in this report is covered in greater detail in [TR-4598 FabricPool Best Practices](#) and other ONTAP product documentation. The content has been condensed for a FlexPod infrastructure and does not completely cover all use cases for FabricPool. All features and concepts examined are available in ONTAP 9.6.

Additional information about FlexPod is available in [TR-4036 FlexPod Datacenter Technical Specifications](#).

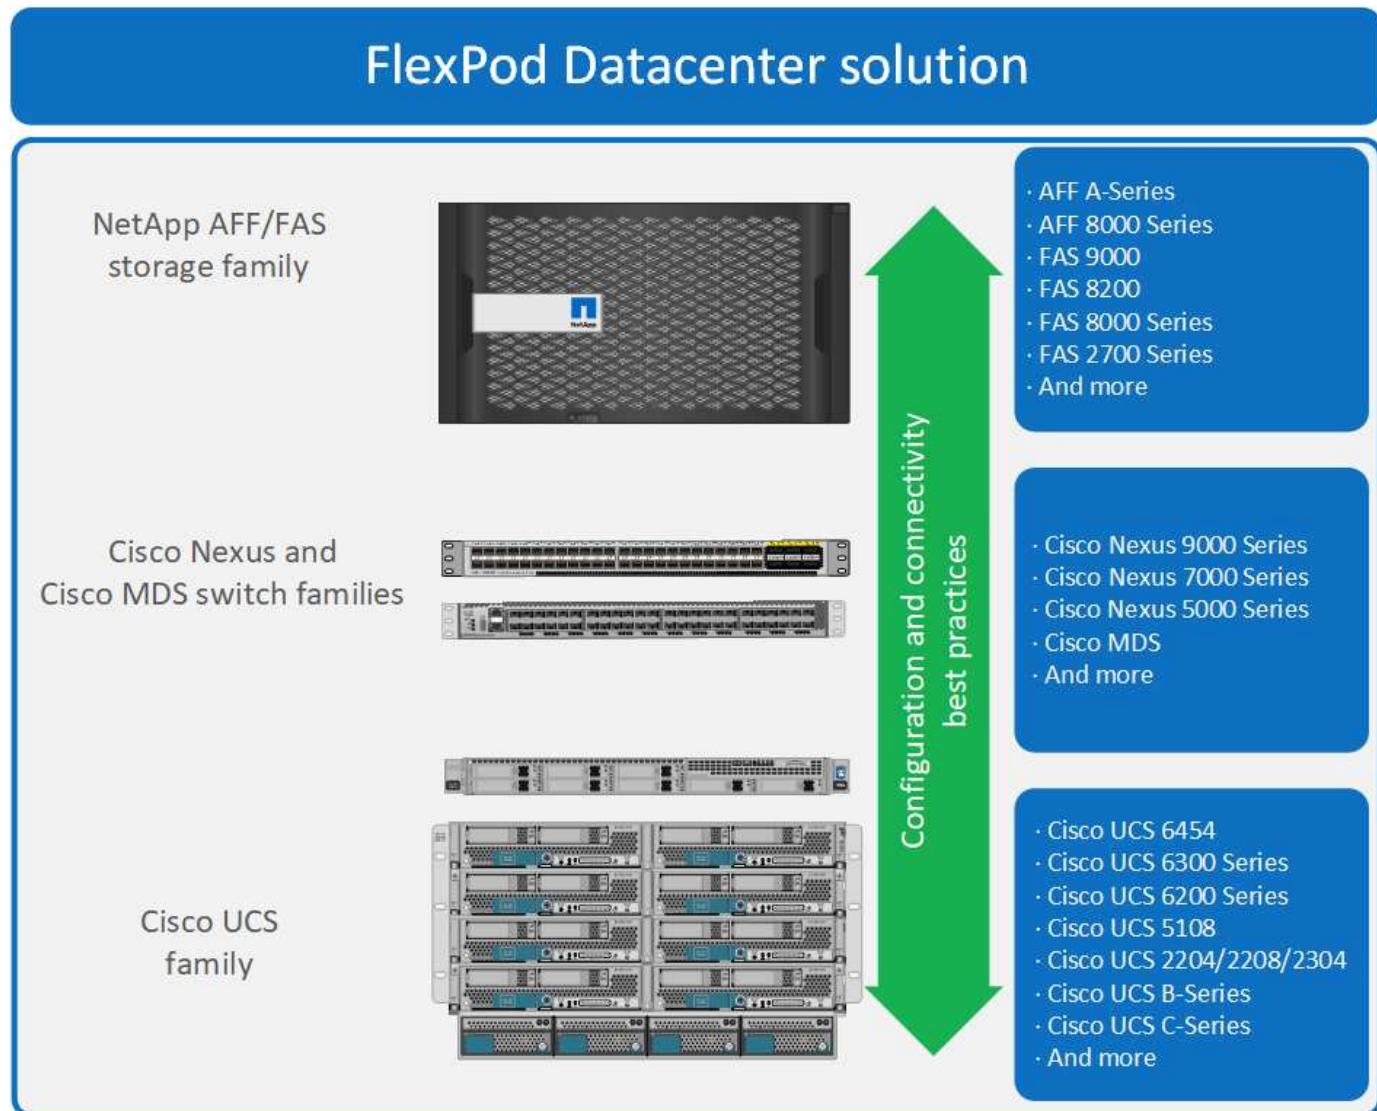
## FlexPod overview and architecture

### FlexPod overview

FlexPod is a defined set of hardware and software that forms an integrated foundation for both virtualized and nonvirtualized solutions. FlexPod includes NetApp AFF storage, Cisco Nexus networking, Cisco MDS storage networking, the Cisco Unified Computing System (Cisco UCS), and VMware vSphere software in a single

package. 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.

One benefit of the FlexPod architecture is the ability to customize, or flex, the environment to suit a customer's requirements. A FlexPod unit can easily be scaled as requirements and demand change. A unit can be scaled both up (adding resources to a FlexPod unit) and out (adding more FlexPod units). The FlexPod reference architecture highlights the resiliency, cost benefit, and ease of deployment of a Fibre Channel and IP-based storage solution. A storage system that is capable of serving multiple protocols across a single interface gives customers a choice and protects their investment because it is truly a wire-once architecture. The following figure shows many of the hardware components of FlexPod.



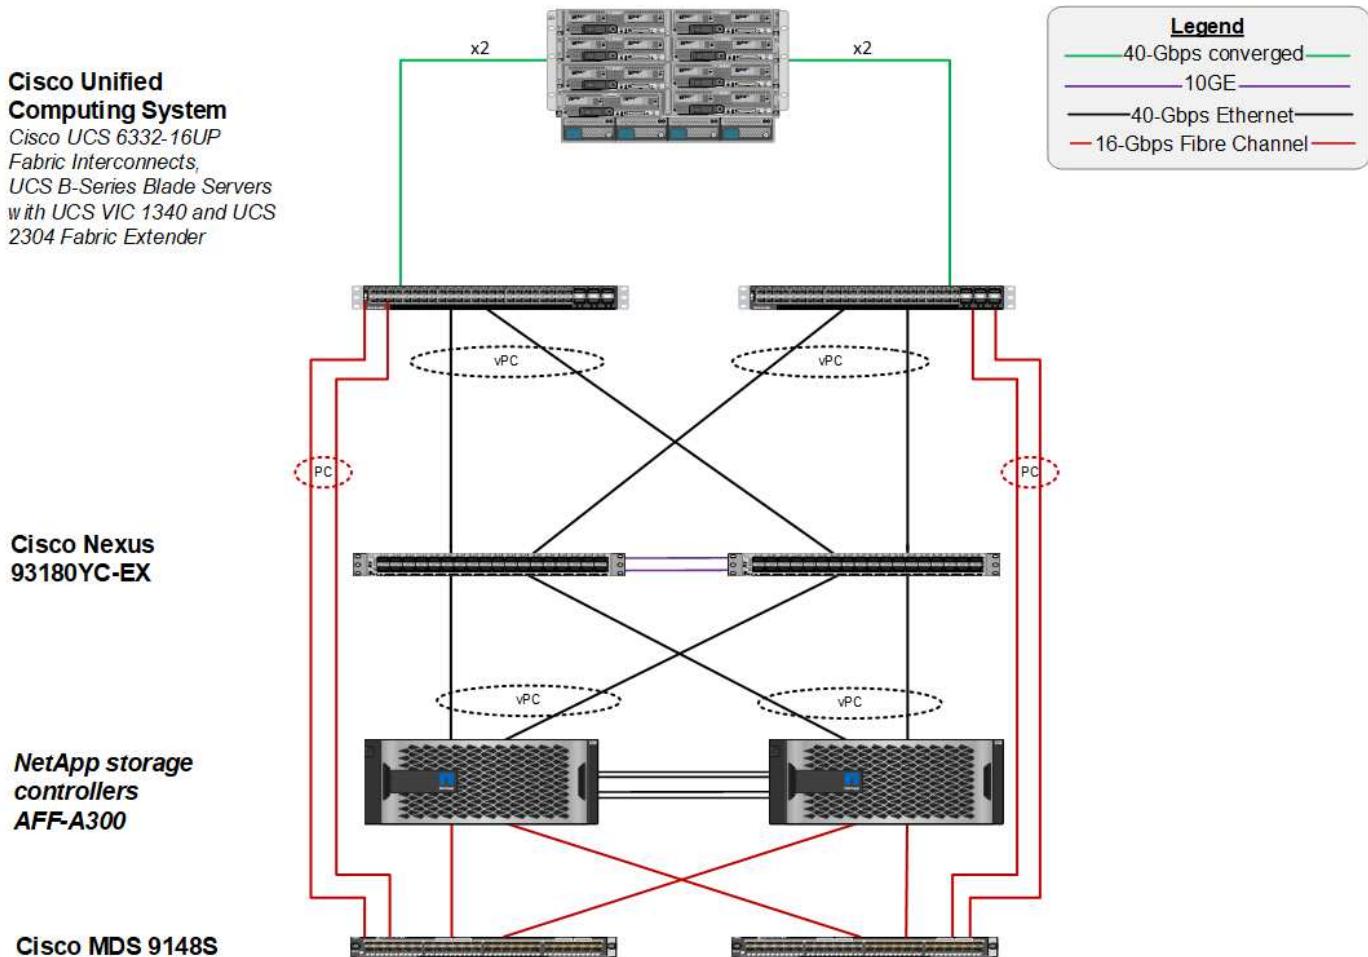
### FlexPod architecture

The following figure shows the components of a VMware vSphere and FlexPod solution and the network connections needed for Cisco UCS 6454 fabric interconnects. This design has the following components:

- Port-channelled 40Gb Ethernet connections between the Cisco UCS 5108 blade chassis and the Cisco UCS fabric interconnects
- 40Gb Ethernet connections between the Cisco UCS fabric interconnect and the Cisco Nexus 9000

- 40Gb Ethernet connections between the Cisco Nexus 9000 and the NetApp AFF A300 storage array

These infrastructure options expanded with the introduction of Cisco MDS switches sitting between the Cisco UCS fabric interconnect and the NetApp AFF A300. This configuration provides FC-booted hosts with 16Gb FC block-level access to shared storage. The reference architecture reinforces the wire-once strategy, because, as additional storage is added to the architecture, no recabling is required from the hosts to the Cisco UCS fabric interconnect.



## FabricPool

### FabricPool overview

FabricPool is a hybrid storage solution in ONTAP that uses an all-flash (SSD) aggregate as a performance tier and an object store in a public cloud service as a cloud tier. This configuration enables policy-based data movement, depending on whether or not data is frequently accessed. FabricPool is supported in ONTAP for both AFF and all-SSD aggregates on FAS platforms. Data processing is performed at the block level, with frequently accessed data blocks in the all-flash performance tier tagged as hot and infrequently accessed blocks tagged as cold.

Using FabricPool helps to reduce storage costs without compromising performance, efficiency, security, or protection. FabricPool is transparent to enterprise applications and capitalizes on cloud efficiencies by lowering storage TCO without having to rearchitect the application infrastructure.

FlexPod can benefit from the storage tiering capabilities of FabricPool to make more efficient use of ONTAP flash storage. Inactive virtual machines (VMs), infrequently used VM templates, and VM backups from NetApp

SnapCenter for vSphere can consume valuable space in the datastore volume. Moving cold data to the cloud tier frees space and resources for high-performance, mission-critical applications hosted on the FlexPod infrastructure.

 Fibre Channel and iSCSI protocols generally take longer before experiencing a timeout (60 to 120 seconds), but they do not retry to establish a connection in the same way that NAS protocols do. If a SAN protocol times out, the application must be restarted. Even a short disruption could be disastrous to production applications using SAN protocols because there is no way to guarantee connectivity to public clouds. To avoid this issue, NetApp recommends using private clouds when tiering data that is accessed by SAN protocols.

In ONTAP 9.6, FabricPool integrates with all the major public cloud providers: Alibaba Cloud Object Storage Service, Amazon AWS S3, Google Cloud Storage, IBM Cloud Object Storage, and Microsoft Azure Blob Storage. This report focuses on Amazon AWS S3 storage as the cloud object tier of choice.

### The composite aggregate

A FabricPool instance is created by associating an ONTAP flash aggregate with a cloud object store, such as an AWS S3 bucket, to create a composite aggregate. When volumes are created inside the composite aggregate, they can take advantage of the tiering capabilities of FabricPool. When data is written to the volume, ONTAP assigns a temperature to each of the data blocks. When the block is first written, it is assigned a temperature of hot. As time passes, if the data is not accessed, it undergoes a cooling process until it is finally assigned a cold status. These infrequently accessed data blocks are then tiered off the performance SSD aggregate and into the cloud object store.

The period of time between when a block is designated as cold and when it is moved to cloud object storage is modified by the volume tiering policy in ONTAP. Further granularity is achieved by modifying ONTAP settings that control the number of days required for a block to become cold. Candidates for data tiering are traditional volume snapshots, SnapCenter for vSphere VM backups and other NetApp Snapshot-based backups, and any infrequently used blocks in a vSphere datastore, such as VM templates and infrequently accessed VM data.

### Inactive data reporting

Inactive data reporting (IDR) is available in ONTAP to help evaluate the amount of cold data that can be tiered from an aggregate. IDR is enabled by default in ONTAP 9.6 and uses a default 31-day cooling policy to determine which data in the volume is inactive.

 The amount of cold data that is tiered depends on the tiering policies set on the volume. This amount may be different than the amount of cold data detected by IDR using the default 31-day cooling period.

### Object creation and data movement

FabricPool works at the NetApp WAFL block level, cooling blocks, concatenating them into storage objects, and migrating those objects to a cloud tier. Each FabricPool object is 4MB and is composed of 1,024 4KB blocks. The object size is fixed at 4MB based on performance recommendations from leading cloud providers and cannot be changed. If cold blocks are read and made hot, only the requested blocks in the 4MB object are fetched and moved back to the performance tier. Neither the entire object nor the entire file is migrated back. Only the necessary blocks are migrated.

 If ONTAP detects an opportunity for sequential readaheads, it requests blocks from the cloud tier before they are read to improve performance.

By default, data is moved to the cloud tier only when the performance aggregate is greater than 50% utilized. This threshold can be set to a lower percentage to allow a smaller amount of data storage on the performance flash tier to be moved to the cloud. This might be useful if the tiering strategy is to move cold data only when the aggregate is nearing capacity.

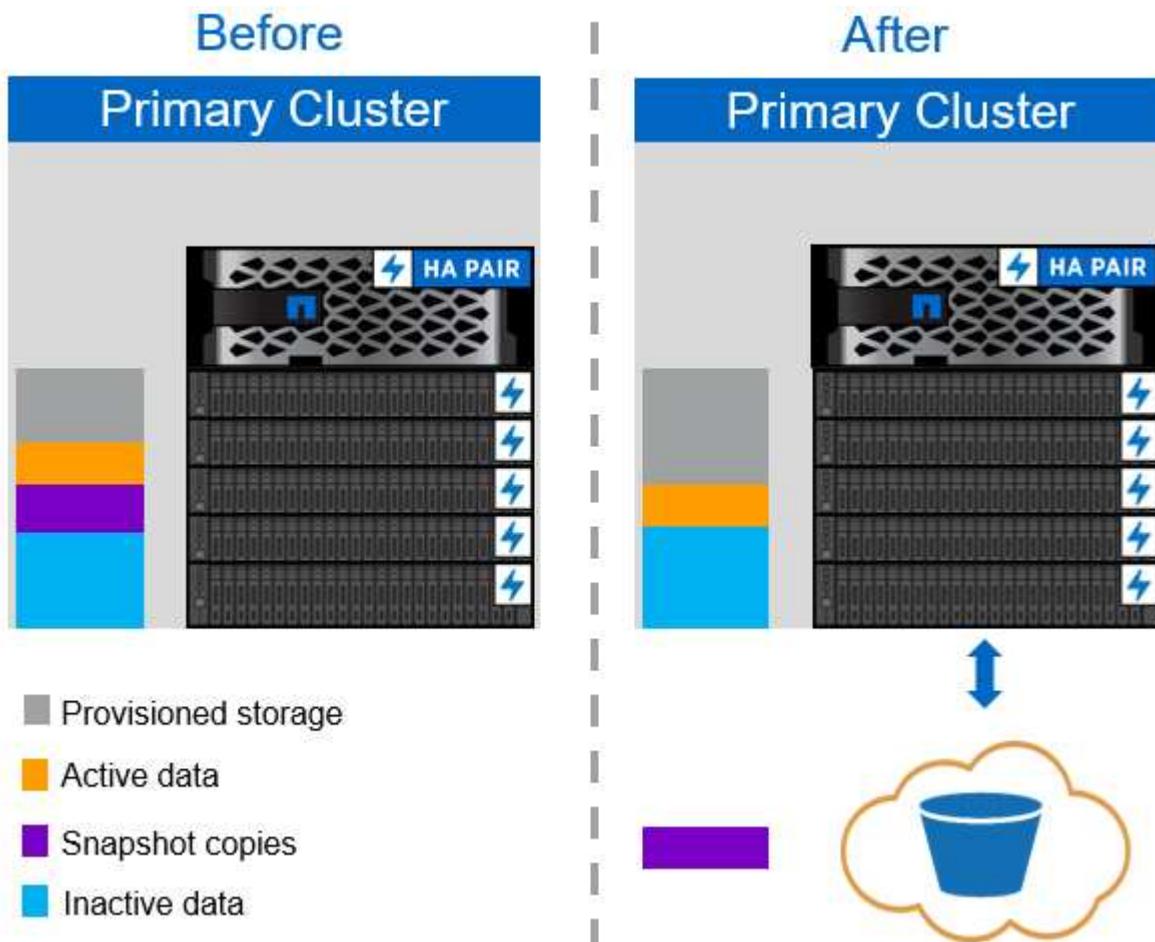
If performance tier utilization is at greater than 70% capacity, cold data is read directly from the cloud tier without being written back to the performance tier. By preventing cold data write-backs on heavily used aggregates, FabricPool preserves the aggregate for active data.

## Reclaim performance tier space

As previously discussed, the primary use case for FabricPool is to facilitate the most efficient use of high-performance on-premises flash storage. Cold data in the form of volume snapshots and VM backups of the FlexPod virtual infrastructure can occupy a significant amount of expensive flash storage. Valuable performance- tier storage can be freed by implementing one of two tiering policies: Snapshot-Only or Auto.

### Snapshot-Only tiering policy

The Snapshot-Only tiering policy, illustrated in the following figure, moves cold volume snapshot data and SnapCenter for vSphere backups of VMs that are occupying space but are not sharing blocks with the active file system into a cloud object store. The Snapshot-Only tiering policy moves cold data blocks to the cloud tier. If a restore is required, cold blocks in the cloud are made hot and moved back to the performance flash tier on the premises.



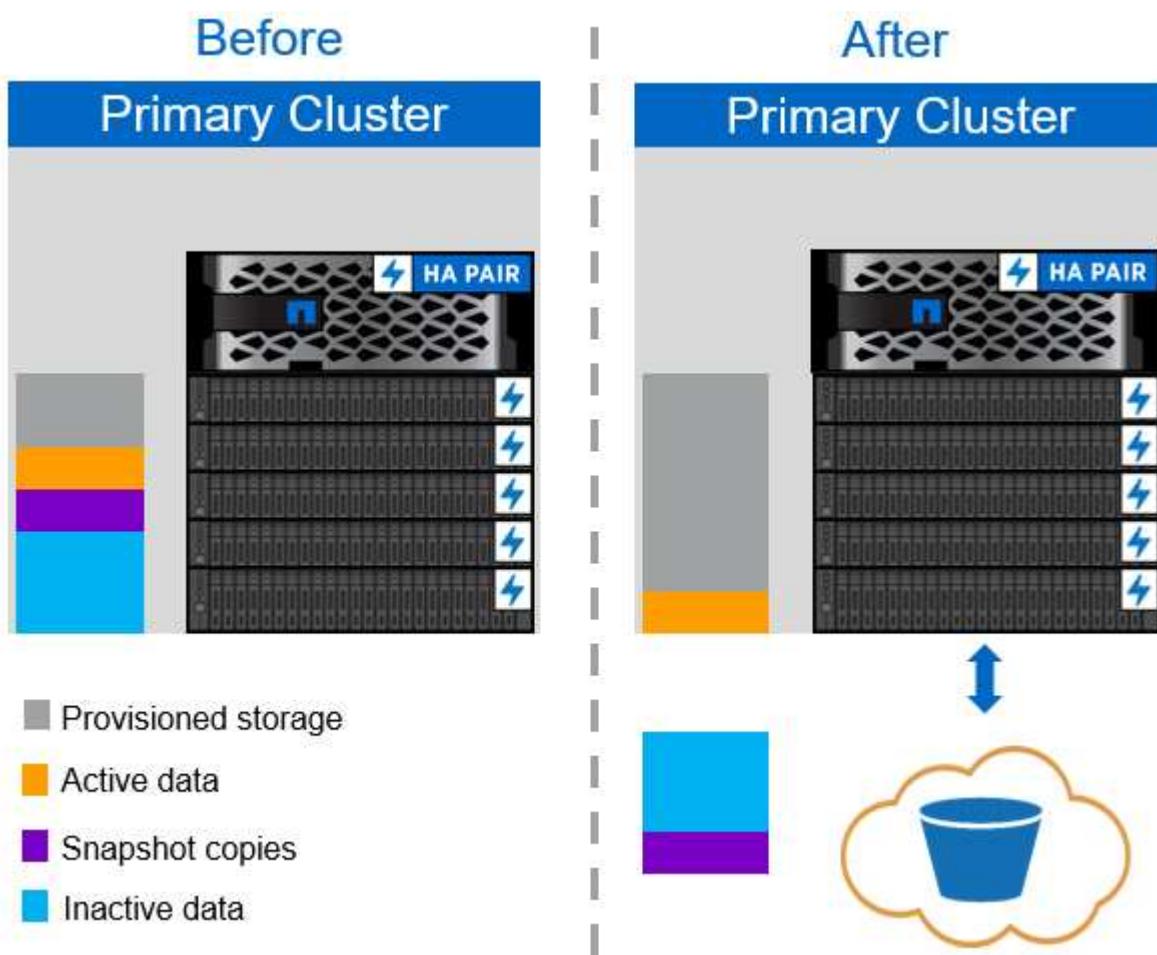
## Auto tiering policy

The FabricPool Auto tiering policy, illustrated in the following figure, not only moves cold snapshot data blocks to the cloud, it also moves any cold blocks in the active file system. This can include VM templates and any unused VM data in the datastore volume. Which cold blocks are moved is controlled by the tiering-minimum-cooling-days setting for the volume. If cold blocks in the cloud tier are randomly read by an application, those blocks are made hot and brought back to the performance tier. However, if cold blocks are read by a sequential process such as an antivirus scanner, the blocks remain cold and persist in the cloud object store; they are not moved back to the performance tier.

When using the Auto tiering policy, infrequently accessed blocks that are made hot are pulled back from the cloud tier at the speed of cloud connectivity. This may affect VM performance if the application is latency sensitive, which should be considered before using the Auto tiering policy on the datastore. NetApp recommends placing Intercluster LIFs on ports with a speed of 10GbE for adequate performance.



The object store profiler should be used to test latency and throughput to the object store before attaching it to a FabricPool aggregate.



## All tiering policy

Unlike the Auto and Snapshot-only policies, the All tiering policy moves entire volumes of data immediately into the cloud tier. This policy is best suited to secondary data protection or archival volumes for which data must be kept for historical or regulatory purposes but is rarely accessed. The All policy is not recommended for VMware datastore volumes because any data written to the datastore is immediately moved to the cloud tier. Subsequent read operations are performed from the cloud and could potentially introduce performance issues.

for VMs and applications residing in the datastore volume.

## Security

Security is a central concern for the cloud and for FabricPool. All the native security features of ONTAP are supported in the performance tier, and the movement of data is secured as it is transferred to the cloud tier. FabricPool uses the [AES-256-GCM](#) encryption algorithm on the performance tier and maintains this encryption end to end into the cloud tier. Data blocks that are moved to the cloud object store are secured with transport layer security (TLS) v1.2 to maintain data confidentiality and integrity between storage tiers.



Communicating with the cloud object store over an unencrypted connection is supported but not recommended by NetApp.

### Data encryption

Data encryption is vital to the protection of intellectual property, trade information, and personally identifiable customer information. FabricPool fully supports both NetApp Volume Encryption (NVE) and NetApp Storage Encryption (NSE) to maintain existing data protection strategies. All encrypted data on the performance tier remains encrypted when moved to the cloud tier. Client-side encryption keys are owned by ONTAP and the server-side object store encryption keys are owned by the respective cloud object store. Any data not encrypted with NVE is encrypted with the AES-256-GCM algorithm. No other AES-256 ciphers are supported.



The use of NSE or NVE is optional and not required to use FabricPool.

## FabricPool requirements

FabricPool requires ONTAP 9.2 or later and the use of SSD aggregates on any of the platforms listed in this section. Additional FabricPool requirements depend on the cloud tier being attached. For entry-level AFF platforms that have a fixed, relatively small capacity such as the NetApp AFF C190, FabricPool can be highly effective for moving inactive data to the cloud tier.

### Platforms

FabricPool is supported on the following platforms:

- NetApp AFF
  - A800
  - A700S, A700
  - A320, A300
  - A220, A200
  - C190
  - AFF8080, AFF8060, and AFF8040
- NetApp FAS
  - FAS9000
  - FAS8200
  - FAS8080, FAS8060, and FAS8040

- FAS2750, FAS2720
- FAS2650, FAS2620



Only SSD aggregates on FAS platforms can use FabricPool.

- Cloud tiers
  - Alibaba Cloud Object Storage Service (Standard, Infrequent Access)
  - Amazon S3 (Standard, Standard-IA, One Zone-IA, Intelligent-Tiering)
  - Amazon Commercial Cloud Services (C2S)
  - Google Cloud Storage (Multi-Regional, Regional, Nearline, Coldline)
  - IBM Cloud Object Storage (Standard, Vault, Cold Vault, Flex)
  - Microsoft Azure Blob Storage (Hot and Cool)

## Intercluster LIFs

Cluster high-availability (HA) pairs that use FabricPool require two intercluster logical interfaces (LIFs) to communicate with the cloud tier. NetApp recommends creating an intercluster LIF on additional HA pairs to seamlessly attach cloud tiers to aggregates on those nodes as well.

The LIF that ONTAP uses to connect with the AWS S3 object store must be on a 10Gbps port.

If more than one Intercluster LIF is used on a node with different routing, NetApp recommends placing them in different IPspaces. During configuration, FabricPool can select from multiple IPspaces, but it is not able to select specific intercluster LIFs within an IPspace.



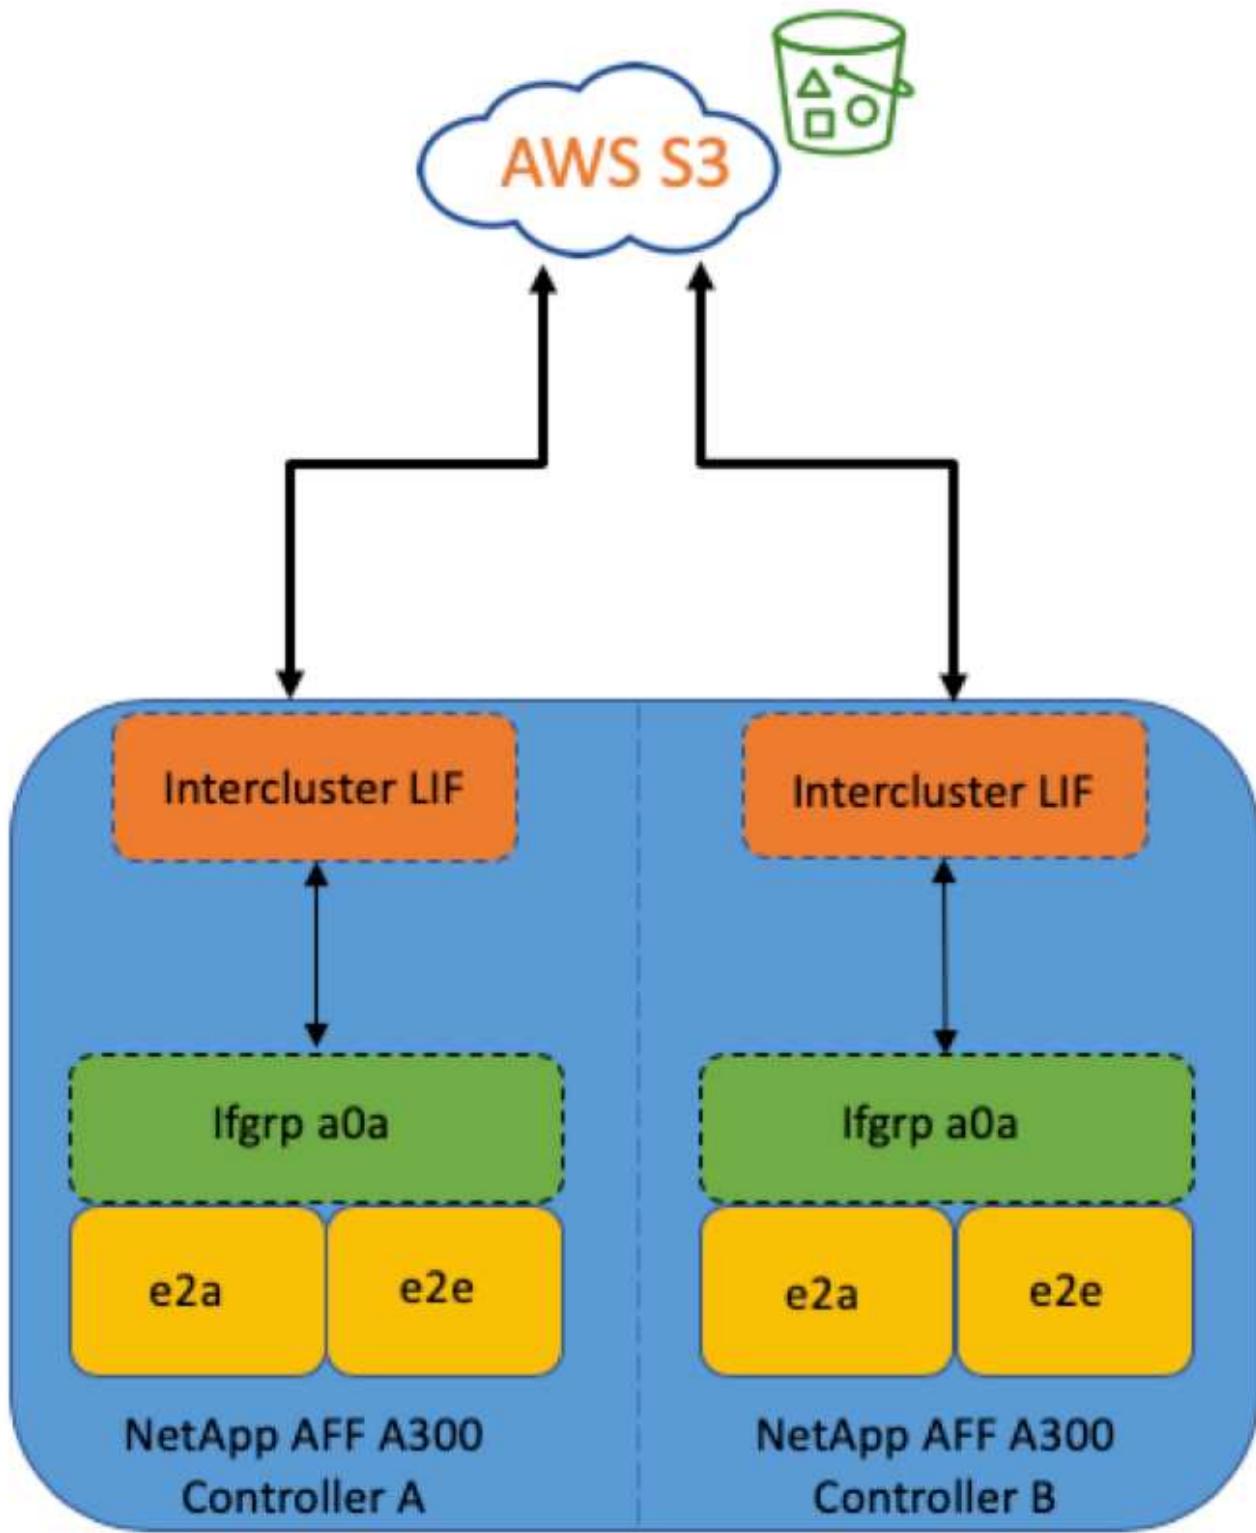
Disabling or deleting an intercluster LIF interrupts communication to the cloud tier.

## Connectivity

FabricPool read latency is a function of connectivity to the cloud tier. Intercluster LIFs using 10Gbps ports, illustrated in the following figure, provide adequate performance. NetApp recommends validating the latency and throughput of the specific network environment to determine the effect it has on FabricPool performance.



When using FabricPool in low-performance environments, minimum performance requirements for client applications must continue to be met, and recovery time objectives should be adjusted accordingly.



#### Object store profiler

The object store profiler, an example of which is shown below and is available through the ONTAP CLI, tests the latency and throughput performance of object stores before they are attached to a FabricPool aggregate.



The cloud tier must be added to ONTAP before it can be used with the object store profiler.

Start the object store profiler from the advanced privilege mode in ONTAP with the following command:

```
storage aggregate object-store profiler start -object-store-name <name>
-node <name>
```

To view the results, run the following command:

```
storage aggregate object-store profiler show
```

Cloud tiers do not provide performance similar to that found on the performance tier (typically GB per second). Although FabricPool aggregates can easily provide SATA-like performance, they can also tolerate latencies as high as 10 seconds and low throughput for tiering solutions that do not require SATA-like performance.

```
bb09-a300-2::*> storage aggregate object-store profiler show
Object store config name: aws_infra_fp_bk_1
Node name: bb09-a300-2-1
Status: Active. Issuing GETs
Start time: 10/3/2019 12:37:24

Op      Size     Total     Failed          Latency (ms)          Throughput
          min       max       avg
-----
```

Op	Size	Total	Failed	Latency (ms)	Throughput		
				min	max	avg	
PUT	4MB	1084	0	336	5951	2817	69.55MB
GET	4KB	158636	0	27	1132	41	32.22MB
GET	8KB	0	0	0	0	0	0B
GET	32KB	0	0	0	0	0	0B
GET	256KB	0	0	0	0	0	0B

5 entries were displayed.

## Volumes

Storage thin provisioning is a standard practice for the FlexPod virtual infrastructure administrator. NetApp Virtual Storage Console (VSC) provisions storage volumes for VMware datastores without any space guarantee (thin provisioning) and with optimized storage efficiency settings per NetApp best practices. If VSC is used to create VMware datastores, no additional action is required, because no space guarantee should be assigned to the datastore volume.



FabricPool cannot attach a cloud tier to an aggregate that contains volumes using a space guarantee other than None (for example, Volume).

```
volume modify -space-guarantee none
```

Setting the `space-guarantee none` parameter provides thin provisioning for the volume. The amount of space consumed by volumes with this guarantee type grows as data is added instead of being determined by the initial volume size. This approach is essential for FabricPool because the volume must support cloud tier data that becomes hot and is brought back to the performance tier.

## Licensing

FabricPool requires a capacity-based license when attaching third-party object storage providers (such as Amazon S3) as cloud tiers for AFF and FAS hybrid flash systems.

FabricPool licenses are available in perpetual or term-based (1-year or 3-year) format.

Tiering to the cloud tier stops when the amount of data (used capacity) stored on the cloud tier reaches the licensed capacity. Additional data, including SnapMirror copies to volumes using the All tiering policy, cannot be tiered until the license capacity is increased. Although tiering stops, data is still accessible from the cloud tier. Additional cold data remains on SSDs until the licensed capacity is increased.

A free 10TB capacity, term-based FabricPool license comes with the purchase of any new ONTAP 9.5 or later cluster, although additional support costs might apply. FabricPool licenses (including additional capacity for existing licenses) can be purchased in 1TB increments.

A FabricPool license can only be deleted from a cluster that contains no FabricPool aggregates.



FabricPool licenses are cluster-wide. You should have the UUID available when purchasing a license (`cluster identify show`). For additional licensing information, refer to the [NetApp Knowledgebase](#).

## Configuration

### Software revisions

The following table illustrates validated hardware and software versions.

Layer	Device	Image	Comments
Storage	NetApp AFF A300	ONTAP 9.6P2	
Compute	Cisco UCS B200 M5 blade servers with Cisco UCS VIC 1340	Release 4.0(4b)	
Network	Cisco Nexus 6332-16UP fabric interconnect	Release 4.0(4b)	
	Cisco Nexus 93180YC-EX switch in NX-OS standalone mode	Release 7.0(3)I7(6)	
Storage network	Cisco MDS 9148S	Release 8.3(2)	
Hypervisor		VMware vSphere ESXi 6.7.0,13006603 6.7U2	ESXi 6.7.0,13006603
		VMware vCenter Server	vCenter server 6.7.0.30000 Build 13639309
Cloud provider		Amazon AWS S3	Standard S3 bucket with default options

The basic requirements for FabricPool are outlined in [FabricPool Requirements](#). After all the basic requirements have been met, complete the following steps to configure FabricPool:

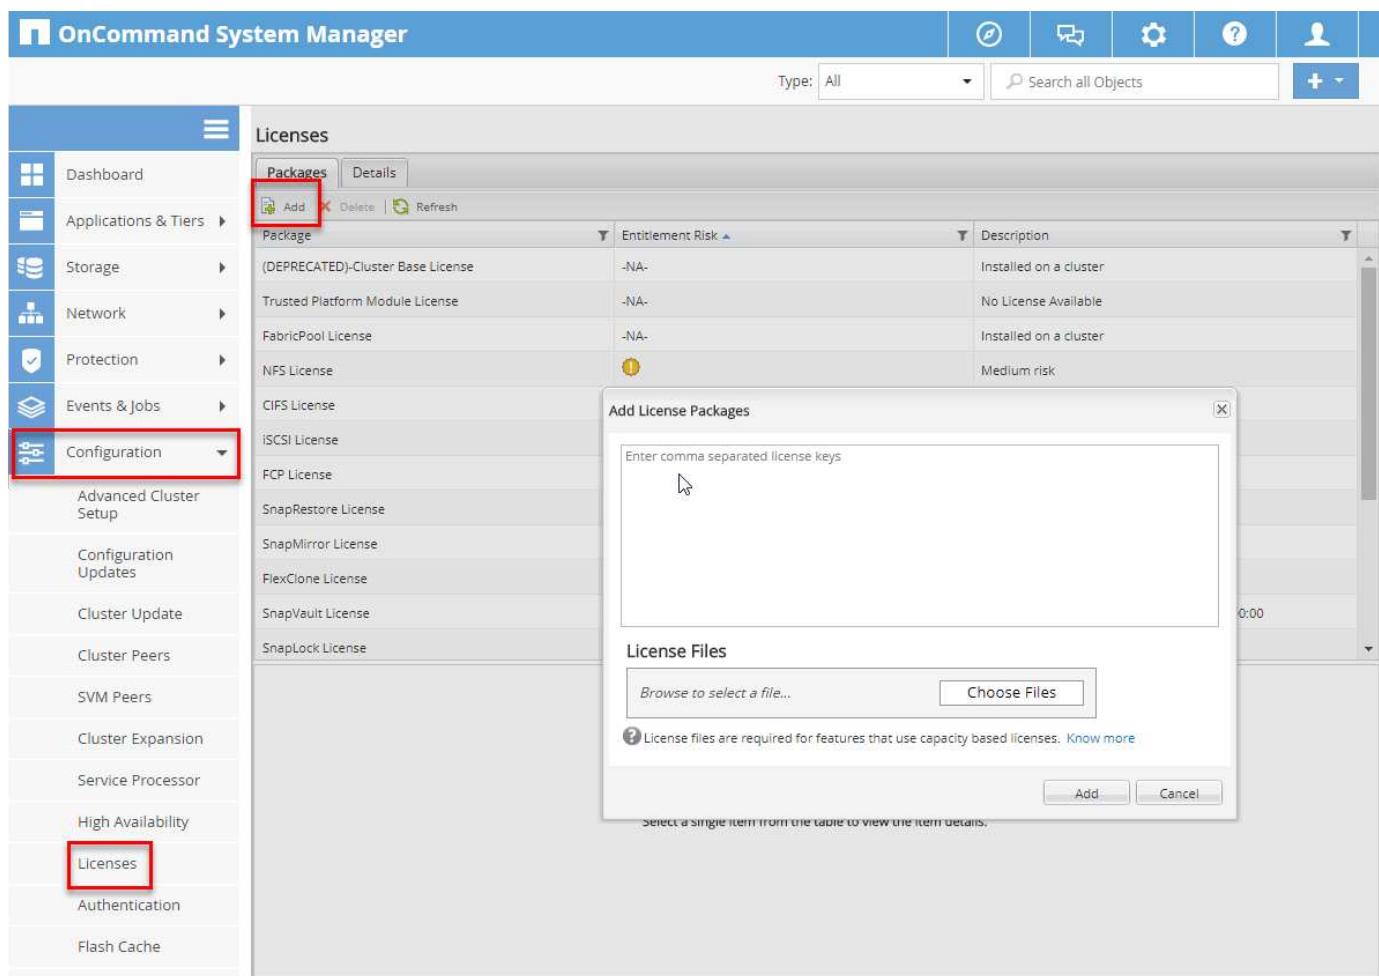
1. Install a FabricPool license.
2. Create an AWS S3 object store bucket.
3. Add a cloud tier to ONTAP.
4. Attach the cloud tier to an aggregate.
5. Set the volume tiering policy.

[Next: Install FabricPool license.](#)

### Install FabricPool license

After you acquire a NetApp license file, you can install it with OnCommand System Manager. To install the license file, complete the following steps:

1. Click Configurations.
2. Click Cluster.
3. Click Licenses.
4. Click Add.
5. Click Choose Files to browse and select a file.
6. Click Add.



The screenshot shows the OnCommand System Manager interface with the 'Licenses' screen selected. The left sidebar is expanded to show 'Configuration' and 'Licenses', with 'Licenses' highlighted by a red box. The main table lists various license packages, including '(DEPRECATED)-Cluster Base License', 'Trusted Platform Module License', 'FabricPool License', 'NFS License', and 'CIFS License'. The 'FabricPool License' row has a yellow warning icon. An 'Add' dialog box is overlaid on the screen, also with a red box around the 'Licenses' tab. The dialog has two main sections: 'Add License Packages' (containing a text input field 'Enter comma separated license keys') and 'License Files' (containing a 'Browse to select a file...' button and a 'Choose Files' button). A note at the bottom of the dialog states: 'License files are required for features that use capacity based licenses. [Know more](#)'.

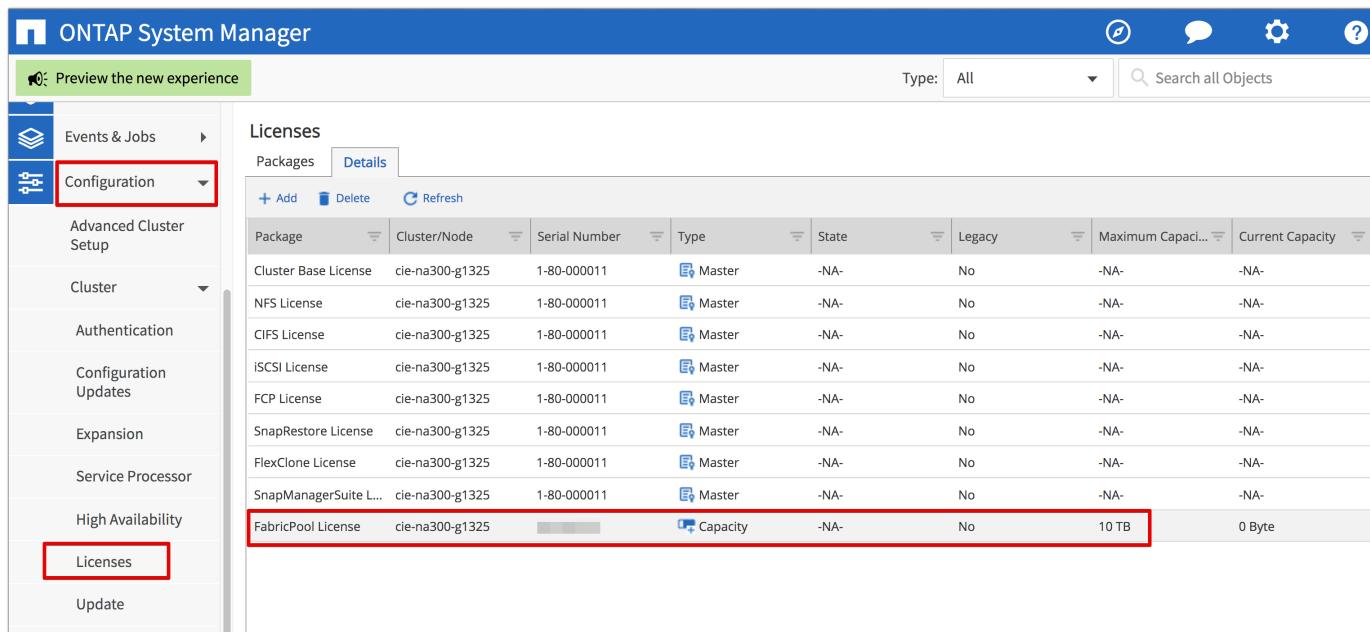
## License capacity

You can view the license capacity by using either the ONTAP CLI or OnCommand System Manager. To see the licensed capacity, run the following command in the ONTAP CLI:

```
system license show-status
```

In OnCommand System Manager, complete the following steps:

1. Click Configurations.
2. Click Licenses.
3. Click the Details tab.



Package	Cluster/Node	Serial Number	Type	State	Legacy	Maximum Capacity	Current Capacity
Cluster Base License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
NFS License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
CIFS License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
iSCSI License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
FCP License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
SnapRestore License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
FlexClone License	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
SnapManagerSuite L...	cie-na300-g1325	1-80-000011	Master	-NA-	No	-NA-	-NA-
<b>FabricPool License</b>	cie-na300-g1325		<b>Capacity</b>	<b>-NA-</b>	<b>No</b>	<b>10 TB</b>	<b>0 Byte</b>

Maximum capacity and current capacity are listed on the FabricPool License row.

[Next: Create AWS S3 bucket.](#)

## Create AWS S3 bucket

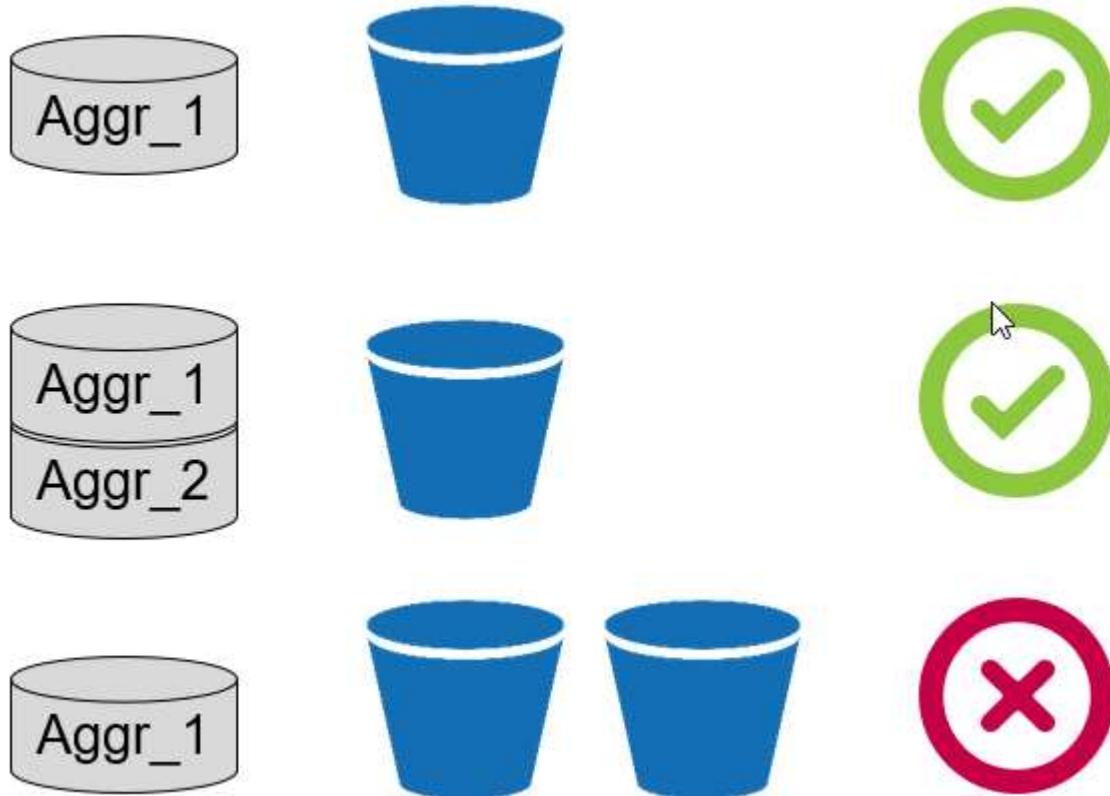
Buckets are object store containers that hold data. You must provide the name and location of the bucket in which data is stored before it can be added to an aggregate as a cloud tier.



Buckets cannot be created using OnCommand System Manager, OnCommand Unified Manager, or ONTAP.

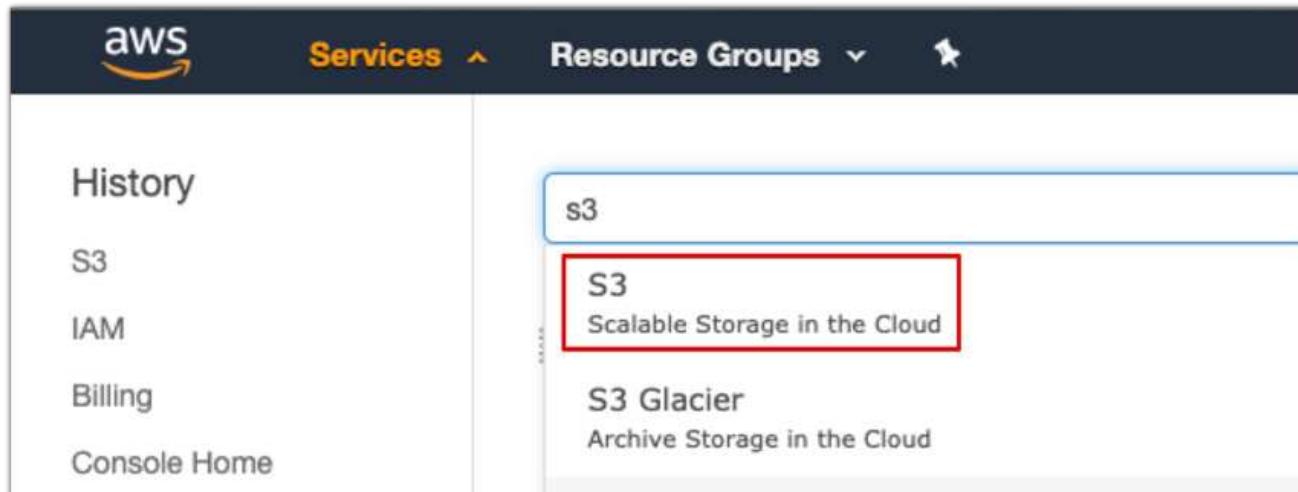
FabricPool supports the attachment of one bucket per aggregate, as illustrated in the following figure. A single bucket can be attached to a single aggregate, and a single bucket can be attached to multiple aggregates. However, a single aggregate cannot be attached to multiple buckets. Although a single bucket can be attached to multiple aggregates in a cluster, NetApp does not recommend attaching a single bucket to aggregates in multiple clusters.

When planning a storage architecture, consider how the bucket-to-aggregate relationship might affect performance. Many object store providers set a maximum number of supported IOPS at the bucket or container level. Environments that require maximum performance should use multiple buckets to reduce the possibility that object-store IOPS limitations might affect performance across multiple FabricPool aggregates. Attaching a single bucket or container to all FabricPool aggregates in a cluster might be more beneficial to environments that value manageability over cloud-tier performance.

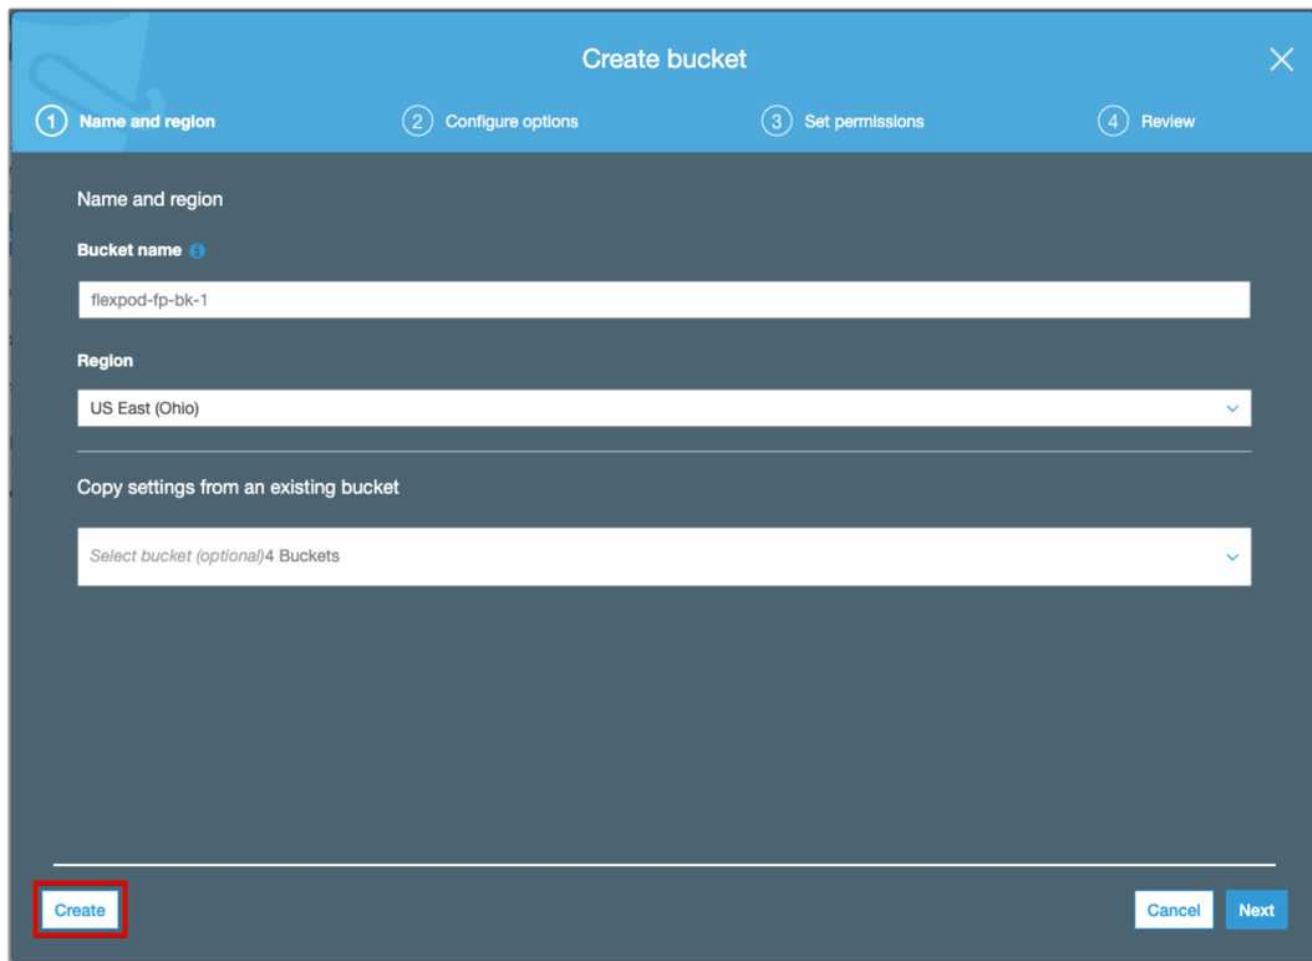


#### Create an S3 bucket

1. In the AWS management console from the home page, enter S3 in the search bar.
2. Select S3 Scalable Storage in the Cloud.



3. On the S3 home page, select Create Bucket.
4. Enter a DNS-compliant name and choose the region to create the bucket.



5. Click Create to create the object store bucket.

[Next: Add a cloud tier to ONTAP](#)

### Add a cloud tier to ONTAP

Before an object store can be attached to an aggregate, it must be added to and identified by ONTAP. This task can be completed with either OnCommand System Manager or the ONTAP CLI.

FabricPool supports Amazon S3, IBM Object Cloud Storage, and Microsoft Azure Blob Storage object stores as cloud tiers.

You need the following information:

- Server name (FQDN); for example, `s3.amazonaws.com`
- Access key ID
- Secret key
- Container name (bucket name)

## OnCommand System Manager

To add a cloud tier with OnCommand System Manager, complete the following steps:

1. Launch OnCommand System Manager.
2. Click Storage.
3. Click Aggregates & Disks.
4. Click Cloud Tiers.
5. Select an object store provider.
6. Complete the text fields as required for the object store provider.

In the Container Name field, enter the object store's bucket or container name.

7. Click Save and Attach Aggregates.

Add Cloud Tier 

Cloud tiers/ object stores are used to store infrequently-accessed data. [Learn more](#)

Cloud Tier Provider  Amazon S3

Type	<input type="text" value="Amazon S3"/>
Name	<input type="text" value="aws_infra_fp_bk_1"/>
Server Name (FQDN)	<input type="text" value="s3.amazonaws.com"/>
Access Key ID	<input type="text"/>
Secret Key	<input type="text"/>
Container Name	<input type="text" value="flexpod-fp-bkt-1"/>
Encryption	Enabled

## ONTAP CLI

To add a cloud tier with the ONTAP CLI, enter the following commands:

```
object-store config create
-object-store-name <name>
-provider-type <AWS>
-port <443/8082> (AWS)
-server <name>
-container-name <bucket-name>
-access-key <string>
-secret-password <string>
-ssl-enabled true
-ipspace default
```

Next: [Attach a cloud tier to an ONTAP aggregate.](#)

### Attach a cloud tier to an ONTAP aggregate

After an object store has been added to and identified by ONTAP, it must be attached to an aggregate to create a FabricPool. This task can be completed by using either OnCommand System Manager or the ONTAP CLI.

More than one type of object store can be connected to a cluster, but only one type of object store can be attached to each aggregate. For example, one aggregate can use Google Cloud, and another aggregate can use Amazon S3, but one aggregate cannot be attached to both.

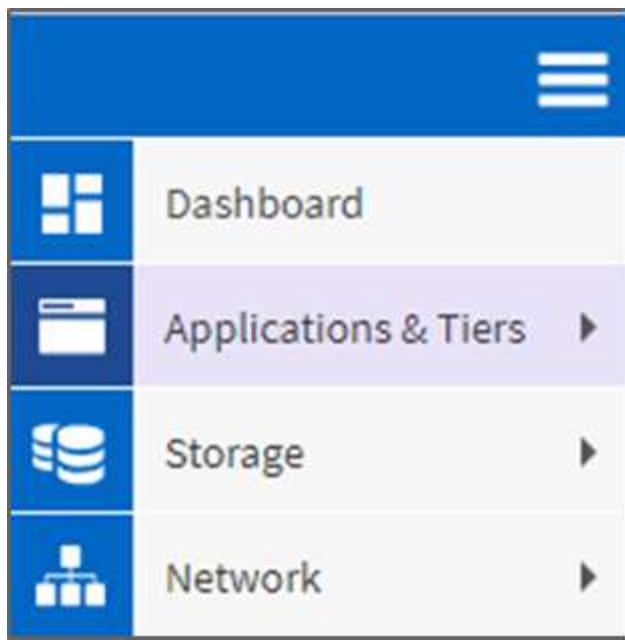


Attaching a cloud tier to an aggregate is a permanent action. A cloud tier cannot be unattached from an aggregate that it has been attached to.

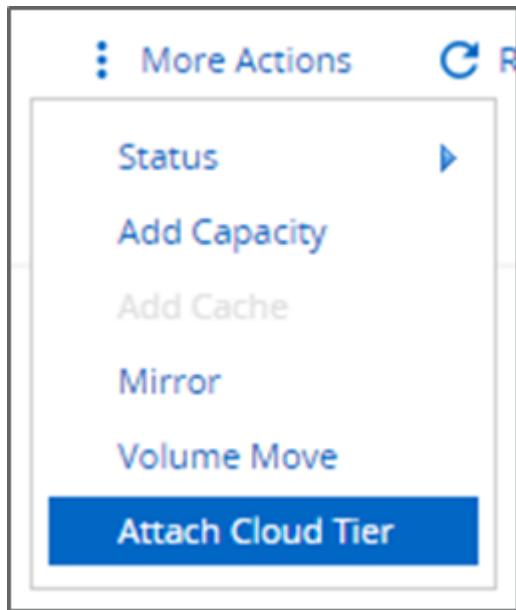
#### OnCommand System Manager

To attach a cloud tier to an aggregate by using OnCommand System Manager, complete the following steps:

1. Launch OnCommand System Manager.
2. Click Applications & Tiers.



3. Click Storage Tiers.
4. Click an aggregate.
5. Click Actions and select Attach Cloud Tier.



6. Select a cloud tier.
7. View and update the tiering policies for the volumes on the aggregate (optional). By default, the volume tiering policy is set as Snapshot-Only.
8. Click Save.

#### ONTAP CLI

To attach a cloud tier to an aggregate by using the ONTAP CLI, run the following commands:

```
storage aggregate object-store attach
-aggregate <name>
-object-store-name <name>
```

Example:

```
storage aggregate object-store attach -aggregate aggr1 -object-store-name
- aws_infra_fp_bk_1
```

[Next: Set volume tiering policy.](#)

### Set volume tiering policy

By default, volumes use the None volume tiering policy. After volume creation, the volume tiering policy can be changed by using OnCommand System Manager or the ONTAP CLI.

When used with FlexPod, FabricPool provides three volume tiering policies, Auto, Snapshot-Only, and None.

- **Auto**

- All cold blocks in the volume are moved to the cloud tier. Assuming that the aggregate is more than 50% utilized, it takes approximately 31 days for inactive blocks to become cold. The Auto cooling period is adjustable between 2 days and 63 days by using the `tiering-minimum-cooling-days` setting.
- When cold blocks in a volume with a tiering policy set to Auto are read randomly, they are made hot and written to the performance tier.
- When cold blocks in a volume with a tiering policy set to Auto are read sequentially, they stay cold and remain on the cloud tier. They are not written to the performance tier.

- **Snapshot-Only**

- Cold snapshot blocks in the volume that are not shared with the active file system are moved to the cloud tier. Assuming that the aggregate is more than 50% utilized, it takes approximately 2 days for inactive snapshot blocks to become cold. The Snapshot-Only cooling period is adjustable from 2 to 63 days by using the `tiering-minimum-cooling-days` setting.
- When cold blocks in a volume with a tiering policy set to Snapshot-Only are read, they are made hot and written to the performance tier.

- **None (Default)**

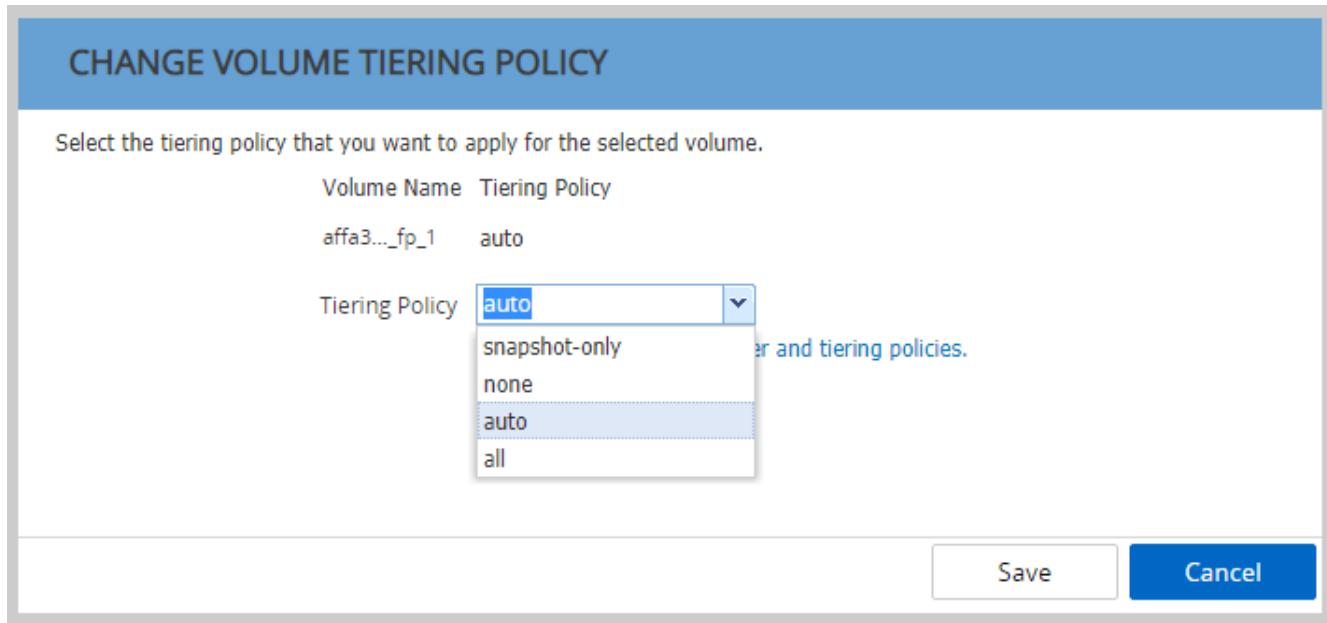
- Volumes set to use None as their tiering policy do not tier cold data to the cloud tier.
- Setting the tiering policy to None prevents new tiering.
- Volume data that has previously been moved to the cloud tier remains in the cloud tier until it becomes hot and is automatically moved back to the performance tier.

### OnCommand System Manager

To change a volume's tiering policy by using OnCommand System Manager, complete the following steps:

1. Launch OnCommand System Manager.

2. Select a volume.
3. Click More Actions and select Change Tiering Policy.
4. Select the tiering policy to apply to the volume.
5. Click Save.



## ONTAP CLI

To change a volume's tiering policy by using the ONTAP CLI, run the following command:

```
volume modify -vserver <svm_name> -volume <volume_name>
-tiering-policy <auto|snapshot-only|all|none>
```

Next: [Set volume tiering minimum cooling days](#).

### Set volume tiering minimum cooling days

The `tiering-minimum-cooling-days` setting determines how many days must pass before inactive data in a volume using the Auto or Snapshot-Only policy is considered cold and eligible for tiering.

#### Auto

The default `tiering-minimum-cooling-days` setting for the Auto tiering policy is 31 days.

Because reads keep block temperatures hot, increasing this value might reduce the amount of data that is eligible to be tiered and increase the amount of data kept on the performance tier.

If you would like to reduce this value from the default 31 days, be aware that data should no longer be active before being marked as cold. For example, if a multiday workload is expected to perform a significant number of writes on day 7, the volume's `tiering-minimum-cooling-days` setting should be set no lower than 8 days.



Object storage is not transactional like file or block storage. Making changes to files that are stored as objects in volumes with overly aggressive minimum cooling days can result in the creation of new objects, the fragmentation of existing objects, and the addition of storage inefficiencies.

### Snapshot-Only

The default `tiering-minimum-cooling-days` setting for the Snapshot-Only tiering policy is 2 days. A 2-day minimum gives additional time for background processes to provide maximum storage efficiency and prevents daily data-protection processes from having to read data from the cloud tier.

### ONTAP CLI

To change a volume's `tiering-minimum-cooling-days` setting by using the ONTAP CLI, run the following command:

```
volume modify -vserver <svm_name> -volume <volume_name> -tiering-minimum-cooling-days <2-63>
```

The advanced privilege level is required.



Changing the tiering policy between Auto and Snapshot-Only (or vice versa) resets the inactivity period of blocks on the performance tier. For example, a volume using the Auto volume tiering policy with data on the performance tier that has been inactive for 20 days will have the performance tier data inactivity reset to 0 days if the tiering policy is set to Snapshot-Only.

## Performance considerations

### Size the performance tier

When considering sizing, keep in mind that the performance tier should be capable of the following tasks:

- Supporting hot data
- Supporting cold data until the tiering scan moves the data to the cloud tier
- Supporting cloud tier data that becomes hot and is written back to the performance tier
- Supporting WAFL metadata associated with the attached cloud tier

For most environments, a 1:10 performance-to-capacity ratio on FabricPool aggregates is extremely conservative, while providing significant storage savings. For example, if the intent is to tier 200TB to the cloud tier, then the performance tier aggregate should be 20TB at a minimum.



Writes from the cloud tier to the performance tier are disabled if performance tier capacity is greater than 70%. If this occurs, blocks are read directly from the cloud tier.

### Size the cloud tier

When considering sizing, the object store acting as the cloud tier should be capable of the following tasks:

- Supporting reads of existing cold data

- Supporting writes of new cold data
- Supporting object deletion and defragmentation

## Cost of ownership

The [FabricPool Economic Calculator](#) is available through the independent IT analyst firm Evaluator Group to help project the cost savings between on premises and the cloud for cold data storage. The calculator provides a simple interface to determine the cost of storing infrequently accessed data on a performance tier versus sending it to a cloud tier for the remainder of the data lifecycle. Based on a 5-year calculation, the four key factors—source capacity, data growth, snapshot capacity, and the percentage of cold data—are used to determine storage costs over the time period.

## Conclusion

The journey to the cloud varies between organizations, between business units, and even between business units within organizations. Some choose a fast adoption, while others take a more conservative approach. FabricPool fits into the cloud strategy of organizations no matter their size and regardless of their cloud adoption speed, further demonstrating the efficiency and scalability benefits of a FlexPod infrastructure.

## Where to find additional information

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

- FabricPool Best Practices  
[www.netapp.com/us/media/tr-4598.pdf](http://www.netapp.com/us/media/tr-4598.pdf)
- NetApp Product Documentation  
<https://docs.netapp.com>
- TR-4036: FlexPod Datacenter Technical Specification  
<https://www.netapp.com/us/media/tr-4036.pdf>

## FlexPod Datacenter for Hybrid Cloud with Cisco CloudCenter and NetApp private storage - Design

Haseeb Niazi, Cisco  
David Arnette, NetApp

Cisco Validated Designs (CVDs) deliver systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. These designs incorporate a wide range of technologies and products into a portfolio of solutions that have been developed to address the business needs of the customers and to guide them

from design to deployment.

[FlexPod Datacenter for Hybrid Cloud with Cisco CloudCenter and NetApp private storage - Design](#)

## Copyright information

Copyright © 2025 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.