



# **FlexPod Hybrid Cloud with Cloud Volumes ONTAP for Epic**

**FlexPod**

NetApp  
October 22, 2024

This PDF was generated from <https://docs.netapp.com/us-en/flexpod/hybrid-cloud/fhc-cvoe-solution-overview.html> on October 22, 2024. Always check docs.netapp.com for the latest.

# Table of Contents

- 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 . . . . . 8
  - SAN configuration . . . . . 13
  - Solution validation . . . . . 25
  - Conclusion . . . . . 35
  - Where to find additional information. . . . . 35

# 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.

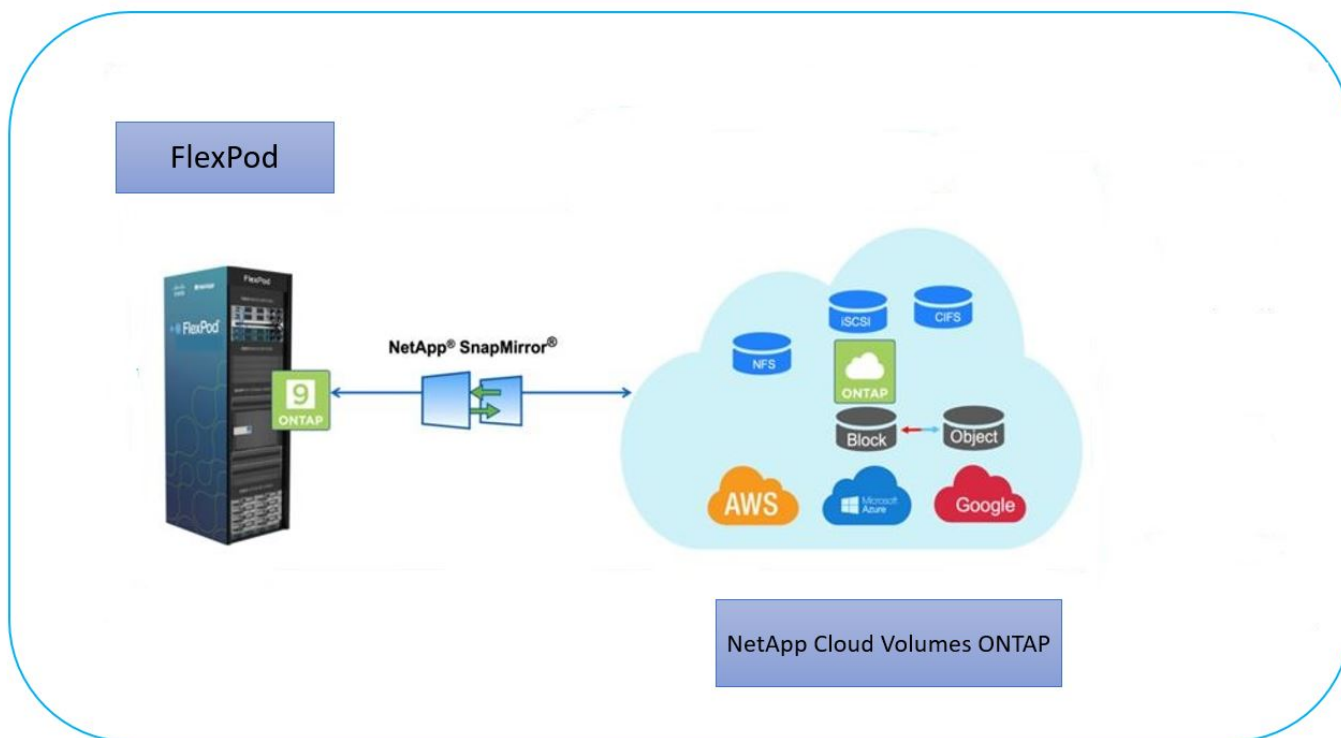
The ever-increasing size of healthcare data and the valuable insights that this data can provide make healthcare data services and data protection both critical and challenging. First, healthcare data must be both available and protected to meet data recovery, medical business continuity, or compliance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

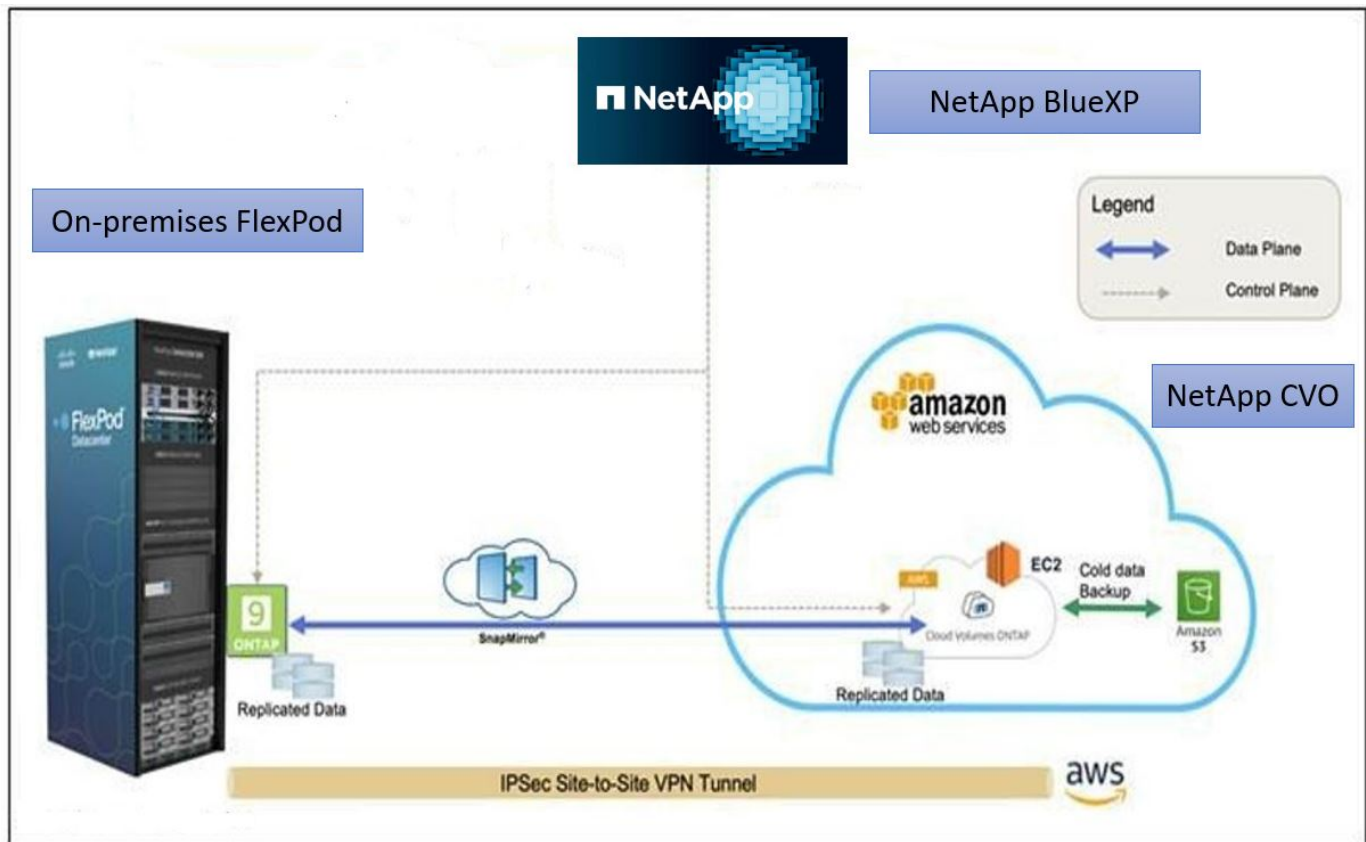
FlexPod Datacenter integrated with NetApp Cloud Volumes ONTAP offers the following benefits to healthcare workloads:

- **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 Failback.** 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 BlueXP 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. In this way, you can 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 BlueXP 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 validated and supports a platform containing 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, gives customers the confidence to run their healthcare applications in a fully managed private cloud through FlexPod that can also be connected to any of the public cloud providers.

## **NetApp BlueXP**

BlueXP (formerly NetApp 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, supporting hybrid, multiple cloud providers and accounts. For more information, see [BlueXP](#).

## **Connector**

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

Connector is supported in the following locations:

- Amazon Web Services
- Microsoft Azure
- Google Cloud
- On premises

To learn more about Connector, see the [Connector page](#).

## 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.

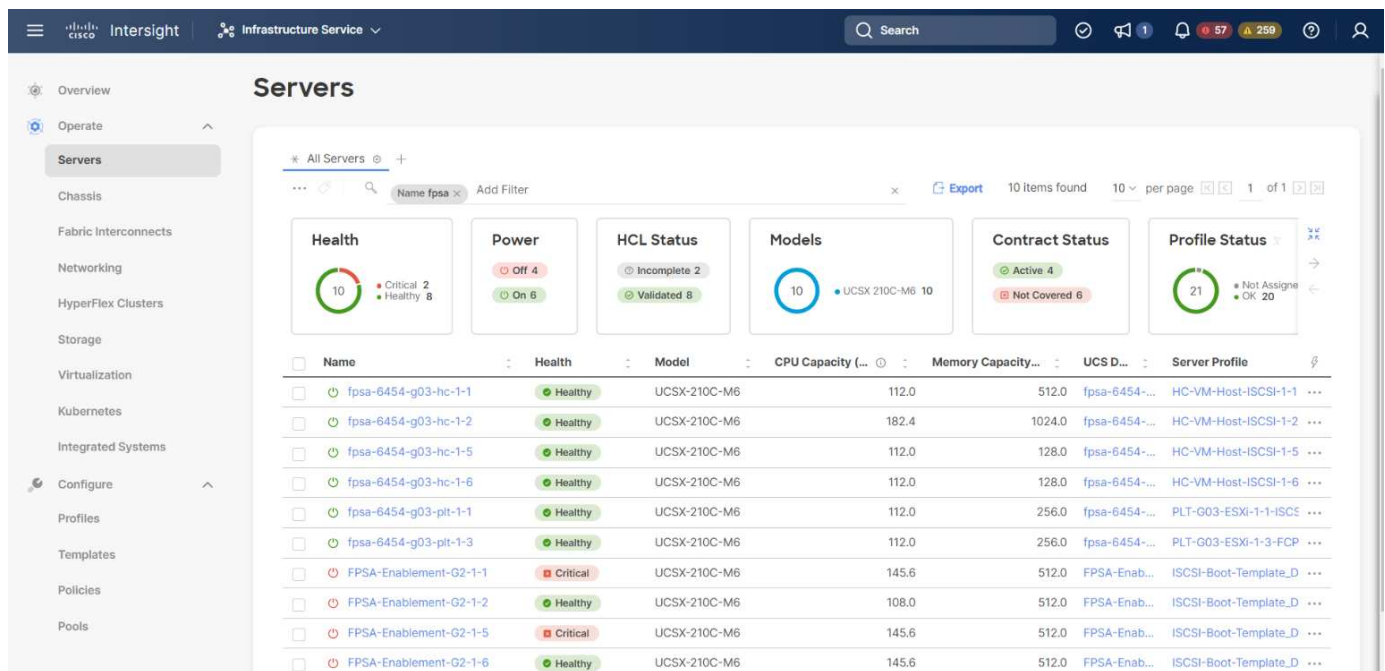
For more information, see [Active IQ Unified Manager](#).

## 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.



## 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 NetApp BlueXP and Cloud Volumes ONTAP versions.

Vendor	Product	Version
NetApp	BlueXP	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 NetApp BlueXP.

To access the NetApp BlueXP and other cloud services, you need to sign up on [NetApp BlueXP](#). For setting up workspaces and users in the BlueXP account, click [here](#). You need an account that has permission to deploy the Connector in your cloud provider directly from BlueXP. You can download the BlueXP policy from [here](#).

4. Deploy Connector.


Before adding a Cloud Volume ONTAP working environment, you must deploy Connector. BlueXP prompts you if you try to create your first Cloud Volumes ONTAP working environment without Connector in place. To deploy Connector in AWS from BlueXP, see this [link](#).

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. The following figure depicts the NetApp BlueXP Dashboard with single-node CVO instance.



## 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.

1. Configure an SVM with the iSCSI service running.

```

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

```

To verify:

```

A400-G0312::> vservers iscsi show -vservers 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.

## 2. Create a FlexVol volume.

```

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

```

## 3. Add interfaces for iSCSI access.

```

1. network interface create -vservers 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 -vservers 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 -vservers 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 -vservers 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. igroup 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`.

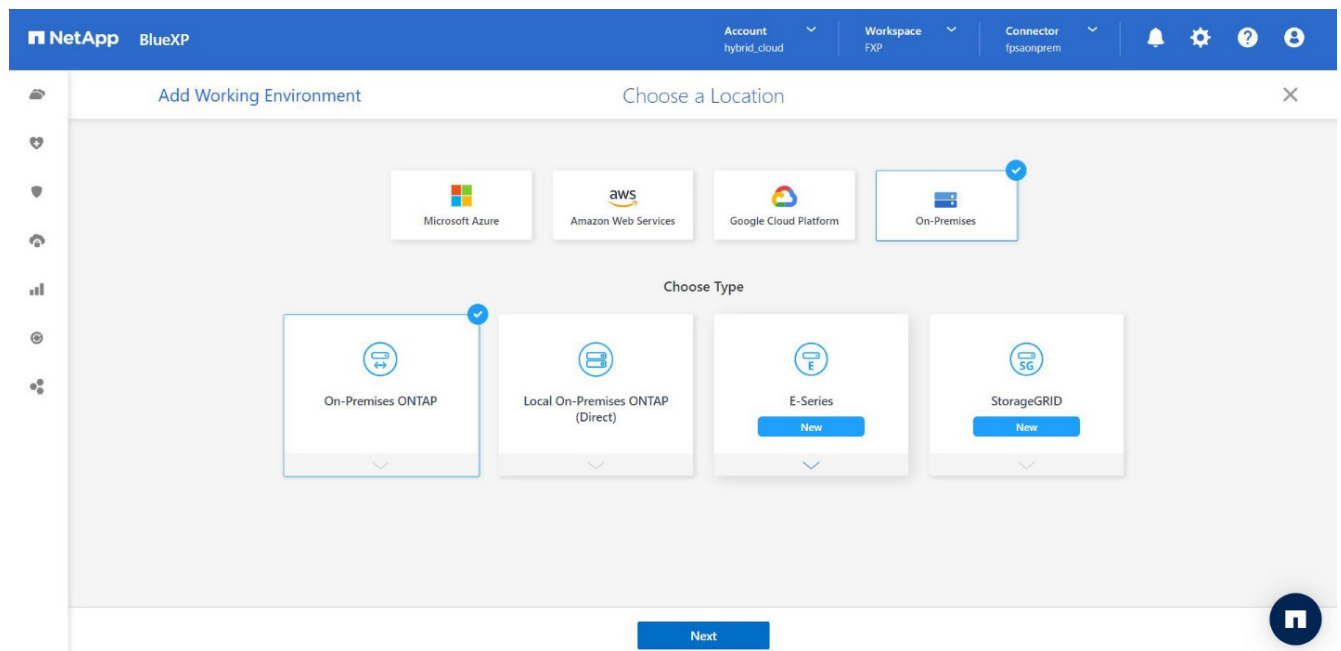
#### 6. Map the LUN to the igroup.

```
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 BlueXP

Complete the following steps to add your FlexPod storage to the working environment using NetApp BlueXP.

1. From the navigation menu, select **Storage > Canvas**.
2. On the Canvas page, click **Add Working Environment** 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**.

BlueXP discovers the ONTAP cluster and adds it as a working environment on the Canvas.



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 `datalv` 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 what would be found inside 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.

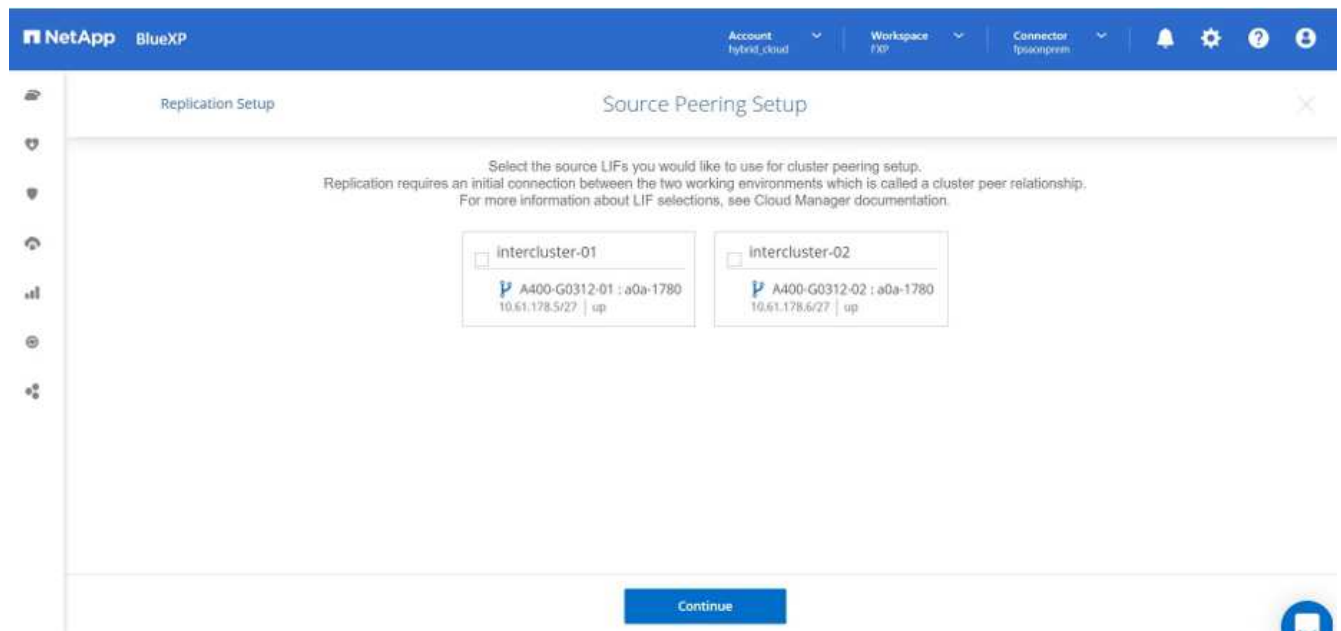
Complete the following steps to configure SnapMirror replication between your on-premises ONTAP system and CVO.

1. From the navigation menu, select **Storage > Canvas**.
2. In Canvas, select the working environment that contains the source volume, drag it to the working environment 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.

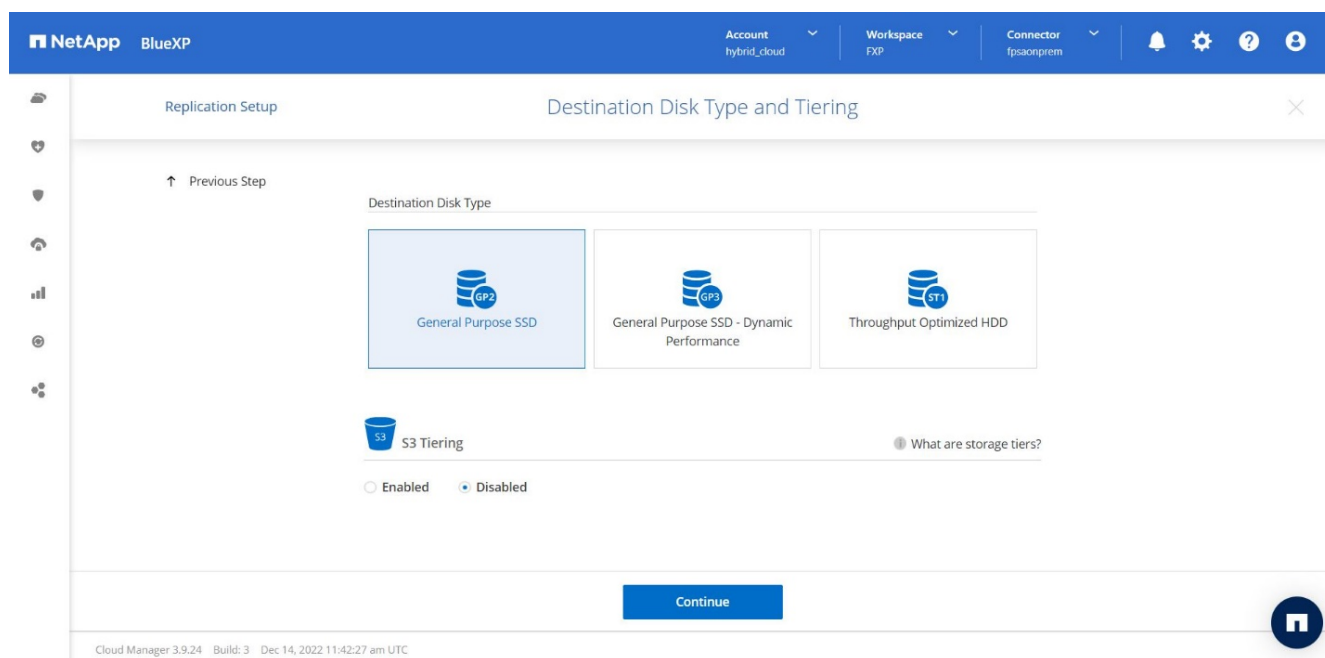
3. **Source and destination peering setup.** If this page appears, select all the intercluster LIFs for the cluster peer relationship.



4. **Source Volume Selection.** Select the volume that you want to replicate.



- 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.



- 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.

NetApp BlueXP

Account hybrid\_cloud Workspace FXP Connector fipsaonprem

Replication Setup Destination Volume Name

↑ Previous Step

Destination Volume Name  
hc\_iscsi\_vol\_copy

Destination Aggregate  
Automatically select the best aggregate

Continue

Cloud Manager 3.9.24 Build: 3 Dec 14, 2022 11:42:27 am UTC

7. **Max transfer rate.** Specify the maximum rate (in megabytes per second) at which data can be transferred.

NetApp BlueXP

Account hybrid\_cloud Workspace FXP Connector fipsaonprem

Replication Setup Max Transfer Rate

↑ Previous Step

You should limit the transfer rate. An unlimited rate might negatively impact the performance of other applications and it might impact your Internet performance.


☒ Limited to: 100 MB/s

☐ Unlimited (recommended for DR only machines)


Continue

Cloud Manager 3.9.24 Build: 3 Dec 14, 2022 11:42:27 am UTC

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](#).

BlueXP starts the data replication process. Now, 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.

NetApp BlueXP Account: hybrid\_cloud Workspace: FXP Connector: fpsaonprem

singlecvoaws Switch to Advanced View AWS AWS Managed Encryption

Volumes Cost Replications

Volumes hc\_iscsi Add Volume

★ New version available Upgrade now

1 of 21 Volumes | 500 GB Allocated | 170.02 GB Total Used (511.70 GB in EBS, 0 KB in S3)

hc\_iscsi\_vol\_copy ONLINE

INFO

Disk Type	GP2
Tiering Policy	None
Backup	OFF

CAPACITY

500 GB Allocated

170.02 GB EBS Used

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

NetApp BlueXP Account: hybrid\_cloud Workspace: FXP Connector: fpsaonprem

singlecvoaws Switch to Advanced View AWS AWS Managed Encryption

Volumes Cost Replications

1 Volume Relationships 170.26 GB Replicated Capacity 0 Currently Transferring 1 Healthy 0 Failed

Search 1 relationship Refresh Add / Remove columns

Source	Target	Lag Duration	Relationship Health	Status	Mirror State	Last Successful Transfer	Policy	Schedule
hc_iscsi_vol A400-G0312	hc_iscsi_vol_copy singlecvoaws	An hour	✓ Healthy	idle	snapmirrored	Dec 21, 2022 05:05:00 ... 0 Byte	Mirror	daily

Cloud Manager 3.9.24 Build: 3 Dec 14, 2022 11:42:27 am UTC

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

The screenshot shows the NetApp BlueXP Replication dashboard. At the top, there's a navigation bar with 'NetApp BlueXP' and various account, workspace, and connector dropdowns. The main content area is titled 'Replication' and displays three components: 'hc\_iscsi\_vol (A400-G0312)' as the Source Volume, 'hc\_iscsi\_vol\_copy (singlecvoaws)' as the Target Volume, and a 'Healthy' Replication Health status. Below these, there are three sections: 'Transfer Info', 'Last Transfer Info', and 'Volume Info'. Each section contains a table of details.

Transfer Info				
idle	N/A	101.48 GiB	6 hours 19 minutes 24 secon...	N/A
Status	Type	Total Size	Lag Duration	Priority
100 MiB/s	34 minutes 9 seconds	snappiromed	170.01 GiB / 0 B	1:1
Max Transfer Rate	Total Transfer Time	Mirror State	Used Size / Used on Cloud	Network Compression Ratio

Last Transfer Info			
Jan 19, 2023, 5:40:04 AM	25.63 KiB	2 seconds	update
Last Successful	Size	Duration	Type

Volume Info			
Source Availability Zone	Healthcare_SVM	us-east-1a	svm_singlecvoaws
	Source SVM Name	Destination Availability Zone	Destination SVM Name

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.

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

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

- c. 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 BlueXP 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 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/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`).




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 BlueXP.



## 2. Activate DR in CVO.

- 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 typev
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_singleecvoaws                    /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'
hwhandler='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.

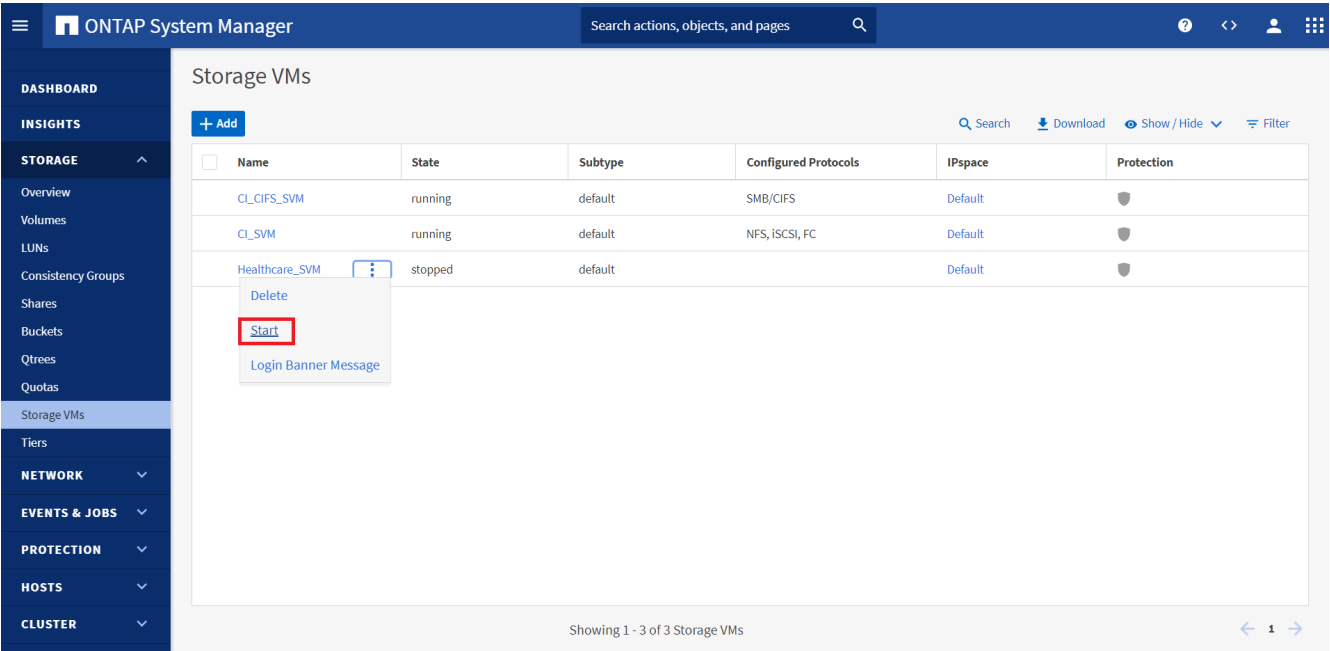
## Verification of data on the production site

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.

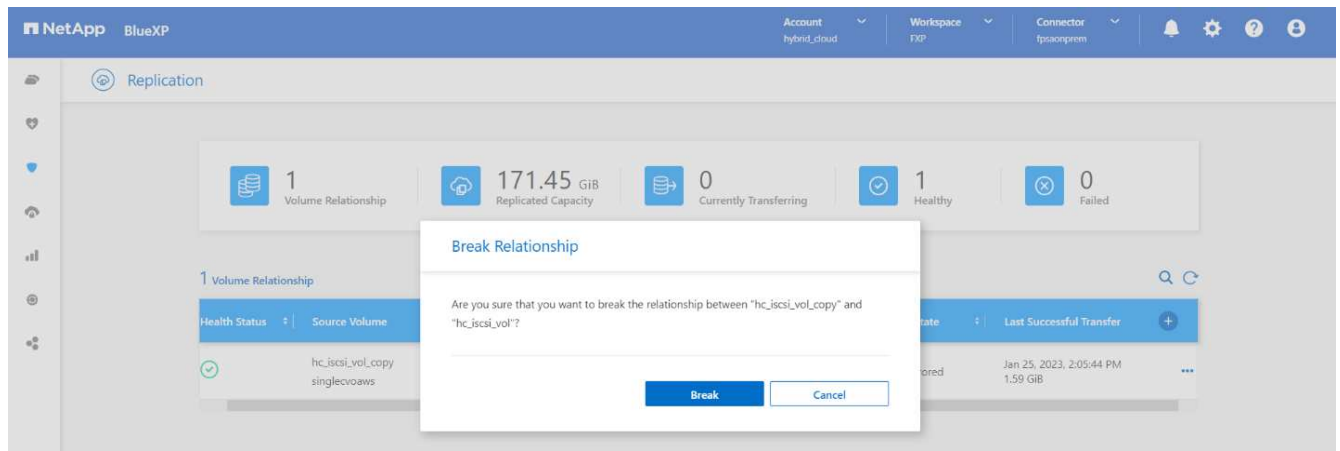
In this section, we talk about bringing up the source site, restoring the SnapMirror relationship between on-premises ONTAP and Cloud Volumes ONTAP, and finally performed a data integrity check on the source end.

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).



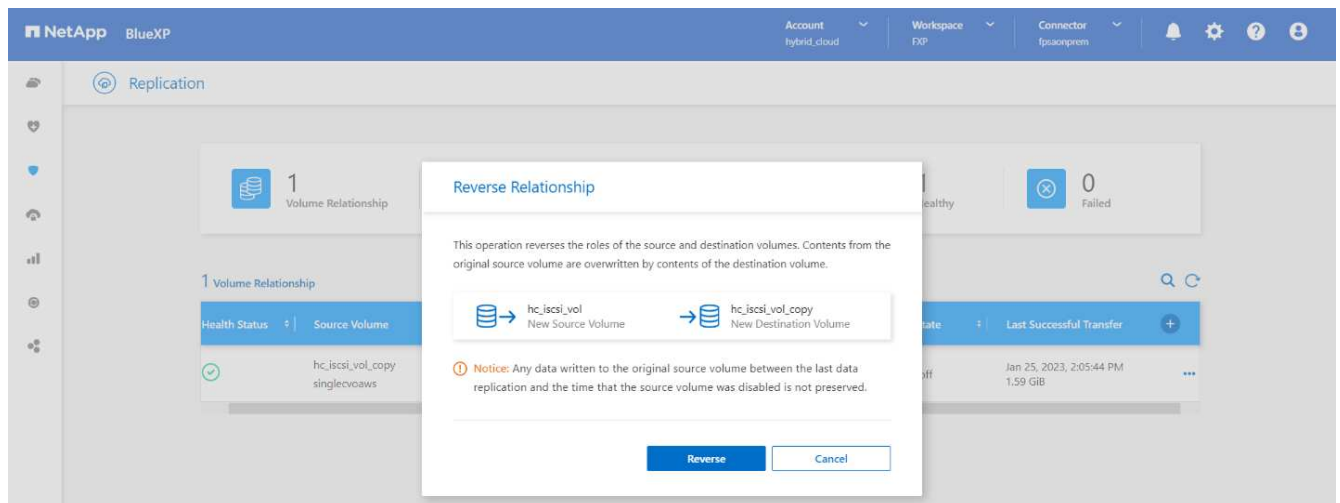
2. Break the SnapMirror replication relationship between Cloud Volumes ONTAP and on-premises ONTAP and promote the on-premises volume (hc\_iscsi\_vol) back to production.



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
```

- 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.

- 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/data1v /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 latest file
dir02 dir06 dir10 dir14 dir18 dir22 dir26 dir30 dir34 dir38 dir42 dir46 dir50 dir54 dir58 dir62 dir66 dir70 dir74 dir78 newfile
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 BlueXP

<https://bluexp.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

## Copyright information

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

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

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

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

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

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

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

## Trademark information

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