



FlexPod Solutions

FlexPod

NetApp
March 25, 2024

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

Table of Contents

| | |
|---|-----|
| FlexPod Solutions | 1 |
| FlexPod Definition | 2 |
| FlexPod Express Technical Specifications | 2 |
| FlexPod Datacenter Technical Specifications | 26 |
| FlexPod Datacenter | 61 |
| FlexPod DataCenter with NetApp SnapMirror Business Continuity and ONTAP 9.10 | 61 |
| FlexPod Datacenter with VMware vSphere 7.0, Cisco VXLAN Single-Site Fabric, and NetApp ONTAP 9.7 - Design | 118 |
| FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7 - Deployment | 118 |
| FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Design | 119 |
| FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Deployment | 119 |
| FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Design | 120 |
| FlexPod Datacenter with VMware vSphere 6.7 U2, Cisco UCS fourth-generation Fabric and NetApp ONTAP 9.6 | 120 |
| FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS fourth-generation fabric, and NetApp AFF A-Series - Design | 121 |
| FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS fourth-generation fabric, and NetApp AFF A-Series | 121 |
| FlexPod Datacenter with Cisco ACI Multi-Pod, NetApp MetroCluster IP, and VMware vSphere 6.7 - Design | 122 |
| FlexPod Datacenter with Cisco ACI Multi-Pod with NetApp MetroCluster IP and VMware vSphere 6.7 - Deployment | 122 |
| Hybrid Cloud | 123 |
| FlexPod Hybrid Cloud with Cloud Volumes ONTAP for Epic | 123 |
| FlexPod Hybrid Cloud for Google Cloud Platform with NetApp Cloud Volumes ONTAP and Cisco Intersight | 158 |
| FlexPod hybrid cloud with NetApp Astra and Cisco Intersight for Red Hat OpenShift | 239 |
| NetApp Cloud Insights for FlexPod | 296 |
| FlexPod with FabricPool - Inactive Data Tiering to Amazon AWS S3 | 320 |
| FlexPod Datacenter with IBM Cloud Private | 341 |
| FlexPod Datacenter for Hybrid Cloud with Cisco CloudCenter and NetApp private storage - Design | 342 |
| FlexPod Datacenter for Multicloud with Cisco CloudCenter and NetApp Data Fabric | 342 |
| Enterprise Databases | 343 |
| SAP | 343 |
| Oracle | 349 |
| Microsoft SQL Server | 351 |
| Healthcare | 353 |
| FlexPod for Genomics | 353 |
| FlexPod for MEDITECH Directional Sizing Guide | 393 |
| FlexPod Datacenter for MEDITECH Deployment Guide | 404 |
| FlexPod for Medical Imaging | 434 |
| Virtual Desktop Infrastructure | 466 |
| FlexPod Datacenter with Citrix Virtual Apps & Desktops 1912 LTSR and VMware vSphere 7 for up to | |

| | |
|--|-----|
| 6000 seats | 466 |
| FlexPod Datacenter with VMware Horizon View 7.10, VMware vSphere 6.7 U2, Cisco UCS Manager 4.0, and NetApp ONTAP 9.6 for up to 6700 seats | 466 |
| 3D graphics visualization with Citrix and NVIDIA - White paper | 466 |
| FlexPod Datacenter with Citrix XenDesktop/XenApp 7.15 and VMware vSphere 6.5 Update 1 for 6000 seats | 467 |
| FlexPod Datacenter with VMware Horizon View 7.3 and VMware vSphere 6.5 Update 1 with Cisco UCS Manager 3.2 for 5000 seats | 467 |
| FlexPod Datacenter with VMware Horizon View 7.10, VMware vSphere 6.7 U2, Cisco UCS Manager 4.0, and NetApp ONTAP 9.6 for up to 6700 seats | 467 |
| Modern Apps | 469 |
| FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning - Design | 469 |
| Deploy NetApp Trident CSI plug-in on Cisco Container Platform with FlexPod | 469 |
| FlexPod Datacenter for OpenShift Container Platform 4 - Deployment | 469 |
| FlexPod Datacenter with Docker Enterprise Edition for Container Management | 470 |
| FlexPod Datacenter for OpenShift Container Platform 4 - Design | 470 |
| FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning - Deployment | 471 |
| 3D graphics visualization with VMware and NVIDIA on Cisco UCS - White paper | 471 |
| 3D graphics visualization with Citrix and NVIDIA - White paper | 471 |
| FlexPod Express | 472 |
| FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series Design Guide | 472 |
| FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series Deployment Guide | 483 |
| FlexPod Express with Cisco UCS C-Series and AFF A220 Series Design Guide | 577 |
| FlexPod Express with Cisco UCS C-Series and AFF A220 Series Deployment Guide | 587 |
| FlexPod Express with VMware vSphere 6.7U1 and NetApp AFF A220 with Direct-Attached IP-Based Storage | 667 |
| FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS - NVA - Deployment | 776 |
| FlexPod and Security | 777 |
| FlexPod, The Solution to Ransomware | 777 |
| FIPS 140-2 security-compliant FlexPod solution for healthcare | 796 |
| Cisco Intersight with NetApp ONTAP Storage | 821 |
| Cisco Intersight with NetApp Storage Quick Start Guide | 821 |
| What's new | 821 |
| Requirements | 825 |
| Before you begin | 826 |
| Configure AIQ UM proxy server for IMT service | 832 |
| Claim targets | 833 |
| Monitor NetApp storage from Cisco Intersight | 834 |
| Use cases | 837 |
| Infrastructure | 841 |
| End-to-End NVMe for FlexPod with Cisco UCSM, VMware vSphere 7.0, and NetApp ONTAP 9 | 841 |
| Legal notices | 852 |
| Copyright | 852 |
| Trademarks | 852 |

| | |
|----------------------|-----|
| Patents | 852 |
| Privacy policy | 852 |

FlexPod Solutions

FlexPod Definition

FlexPod Express Technical Specifications

TR-4293: FlexPod Express Technical Specifications

Karthick Radhakrishnan, Arvind Ramakrishnan, Lindsey Street, Savita Kumari, NetApp

FlexPod Express is a predesigned, best practice architecture that is built on the Cisco Unified Computing System (Cisco UCS) and the Cisco Nexus family of switches, and the storage layer is built by using the NetApp FAS or the NetApp E-Series storage. FlexPod Express is a suitable platform for running various virtualization hypervisors and bare metal operating systems (OSs) and enterprise workloads.

FlexPod Express delivers not only a baseline configuration, but also the flexibility to be sized and optimized to accommodate many different use cases and requirements. This document categorizes the FlexPod Express configurations based on the storage system used, FlexPod Express with NetApp FAS and FlexPod Express with E-Series.

FlexPod platforms

There are three FlexPod platforms:

- **FlexPod Datacenter.** This platform is a massively scalable virtual data center infrastructure suited for workload enterprise applications, virtualization, VDI, and public and private cloud. FlexPod Datacenter has its own specifications, which are documented in [TR-4036: FlexPod Datacenter Technical Specifications](#).
- **FlexPod Express.** This platform is a compact converged infrastructure that is targeted for remote office and edge use cases.

This document provides the technical specifications for the FlexPod Express platform.

FlexPod Rules

The FlexPod design allows a flexible infrastructure that encompasses many different components and software versions.

Use the rule sets as a guide for building or assembling a valid FlexPod configuration. The numbers and rules listed in this document are the minimum requirements for FlexPod; they can be expanded in the included product families as required for different environments and use cases.

Supported versus validated FlexPod configurations

The FlexPod architecture is defined by the set of rules described in this document. The hardware components and software configurations must be supported by the Cisco Hardware Compatibility List (HCL) and the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Each Cisco Validated Design (CVD) or NetApp Verified Architecture (NVA) is a possible FlexPod configuration. Cisco and NetApp document these configuration combinations and validate them with extensive end-to-end testing. The FlexPod deployments that deviate from these configurations are fully supported if they follow the

guidelines in this document and all the components are listed as compatible in the Cisco HCL and NetApp [IMT](#).

For example, adding additional storage controllers or Cisco UCS servers and upgrading software to newer versions is fully supported if the software, hardware, and configurations meet the guidelines defined in this document.

Storage Software

FlexPod Express supports storage systems that run NetApp ONTAP or SANtricity operating systems.

NetApp ONTAP

The NetApp ONTAP software is the operating system that runs on AFF and FAS storage systems. ONTAP provides a highly scalable storage architecture that enables nondisruptive operations, nondisruptive upgrades, and an agile data infrastructure.

For more information about ONTAP, see the [ONTAP product page](#).

E-Series SANtricity software

E-Series SANtricity software is the operating system that runs on E-Series storage systems. SANtricity provides a highly flexible system that meets varying application needs and offers built-in high availability and a wide variety of data protection features.

For more information, see the [SANtricity product page](#).

Minimum hardware requirements

This section describes the minimum hardware requirements for the different versions of FlexPod Express.

FlexPod Express with NetApp FAS

The hardware requirements for FlexPod Express solutions that use NetApp FAS controllers for underlying storage include the configurations describe in this section.

CIMC-based configuration (standalone rack servers)

The Cisco Integrated Management Controller (CIMC) configuration includes the following hardware components:

- Two 10Gbps standard Ethernet switches in a redundant configuration (Cisco Nexus 31108 is recommended, with Cisco Nexus 3000 and 9000 models supported)
- Cisco UCS C-Series standalone rack servers
- Two AFF C190, AFF A250, FAS2600, or FAS 2700 series controllers in a high-availability (HA) pair configuration deployed as a two-node cluster

Cisco UCS-managed configuration

The Cisco UCS-managed confirmation includes the following hardware components:

- Two 10Gbps standard Ethernet switches in a redundant configuration (Cisco Nexus 3524 is recommended)
- One Cisco UCS 5108 alternating current (AC) blade server chassis
- Two Cisco UCS 6324 fabric interconnects
- Cisco UCS B-Series servers (at least four Cisco UCS B200 M5 blade servers)
- Two AFF C190, AFF A250, FAS2750, or FAS2720 controllers in an HA pair configuration (requires two available unified target adapter 2 [UTA2] ports per controller)

FlexPod Express with E-Series

The hardware requirements for the FlexPod Express with E-Series starter configuration include:

- Two Cisco UCS 6324 fabric interconnects
- One Cisco UCS Mini chassis 5108 AC2 or DC2 (the Cisco UCS 6324 fabric interconnects are only supported in the AC2 and DC2 chassis)
- Cisco UCS B-Series servers (at least two Cisco UCS B200 M4 blade servers)
- One HA pair configuration of an E-Series E2824 storage system loaded with minimum 12 disk drives
- Two 10Gbps standard Ethernet switches in a redundant configuration (existing switches in the data center can be used)

These hardware components are required to build a starter configuration of the solution; additional blade servers and disk drives can be added as needed. The E-Series E2824 storage system can be replaced with a higher platform and can also be run as an all-flash system.

Minimum Software Requirements

This section describes the minimum software requirements for the different versions of FlexPod Express.

Software requirements for FlexPod Express with NetApp AFF or FAS

The software requirements for the FlexPod Express with NetApp FAS include:

- ONTAP 9.1 or later
- Cisco NX-OS version 7.0(3)I6(1) or later
- In the Cisco UCS- managed configuration, Cisco UCS Manager UCS 4.0(1b)

All software must be listed and supported in the [NetApp IMT](#). Certain software features might require more recent versions of code than the minimums listed in previous architectures.

Software requirements for FlexPod Express with E-Series

The software requirements for the FlexPod Express with E-Series include:

- E-Series SANtricity software 11.30 or higher
- Cisco UCS Manager 4.0(1b).

All software must be listed and supported in the [NetApp IMT](#).

Connectivity requirements

This section describes the connectivity requirements for the different versions of FlexPod Express.

Connectivity requirements for FlexPod Express with NetApp FAS

The connectivity requirements for FlexPod Express with NetApp FAS include:

- NetApp FAS storage controllers must be directly connected to the Cisco Nexus switches, except in the Cisco UCS-managed configuration, where storage controllers are connected to the fabric interconnects.
- No additional equipment can be placed inline between the core FlexPod components.
- Virtual port channels (vPCs) are required to connect the Cisco Nexus 3000/9000 series switches to the NetApp storage controllers.
- Although it is not required, enabling jumbo frame support is recommended throughout the environment.

Connectivity requirements for FlexPod Express with NetApp E-Series

The connectivity requirements for FlexPod Express with E-Series include:

- The E-Series storage controllers must be directly connected to the fabric interconnects.
- No additional equipment should be placed inline between the core FlexPod components.
- vPCs are required between the fabric interconnects and the Ethernet switches.

Connectivity requirements for FlexPod Express with NetApp AFF

The connectivity requirements for FlexPod Express with NetApp AFF include:

- NetApp AFF storage controllers must be directly connected to the Cisco Nexus switches, except in the Cisco UCS-managed configuration, where storage controllers are connected to the fabric interconnects.
- No additional equipment can be placed inline between the core FlexPod components.
- Virtual port channels (vPCs) are required to connect the Cisco Nexus 3000/9000 series switches to the NetApp storage controllers.
- Although it is not required, enabling jumbo frame support is recommended throughout the environment.

Other requirements

Additional requirements for FlexPod Express include the following:

- Valid support contracts are required for all equipment, including:
 - SMARTnet support for Cisco equipment
 - SupportEdge Advisor or SupportEdge Premium support for NetApp equipment
- All software components must be listed and supported in the [NetApp IMT](#).
- All NetApp hardware components must be listed and supported on [NetApp Hardware Universe](#).
- All Cisco hardware components must be listed and supported on [Cisco HCL](#).

Optional Features

This section describes the optional features for FlexPod Express.

iSCSI boot option

The FlexPod Express architecture uses iSCSI boot. The minimum requirements for the iSCSI boot option include:


- An iSCSI license/feature activated on the NetApp storage controller
- A two-port 10Gbps Ethernet adapter on each node in the NetApp storage controller HA pair
- An adapter in the Cisco UCS server that is capable of iSCSI boot

Configuration options

This section provides more information about the configuration required and validated in the FlexPod Express architecture.

FlexPod Express with Cisco UCS C-Series and AFF C190 Series


The following figure illustrates the FlexPod Express with Cisco UCS C-Series and AFF C190 series solution. This solution supports both 10GbE uplinks.



For more information about this configuration, see the FlexPod Express with VMware vSphere 6.7 and NetApp AFF C190 NVA Deployment Guide (in progress).

FlexPod Express with Cisco UCS Mini and AFF A220 and FAS 2750/2720

The following figure illustrates the FlexPod Express with Cisco UCS- managed configuration.



For more information about this configuration, see [FlexPod Express with VMware vSphere 6.7U1 and NetApp AFF A220 with Direct - Attached IP - Based Storage](#).

Cisco components

Cisco is a substantial contributor to the FlexPod Express design and architecture; it contributes the compute and networking layers of the solution. This section describes the Cisco UCS and Cisco Nexus components that are available for FlexPod Express.

Cisco UCS B-Series blade server options

Cisco UCS B-Series blades currently supported in the Cisco UCS Mini platform are B200 M5 and B420 M4. Other blades will be listed in the following table as they become supported in the Cisco UCS Mini platform.

| Cisco UCS B-Series server | Part number | Technical specifications |
|---------------------------|--------------|---|
| Cisco UCS B200 M5 | UCSB-B200-M5 | https://www.cisco.com/c/en/us/support/servers-unified-computing/ucs-b200-m5-blade-server/model.html |
| Cisco UCS B200 M4 | UCSB-B200-M4 | http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/b200m4-specsheet.pdf |
| Cisco UCS B420 M4 | UCSB-B420-M4 | http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/b420m4-spec-sheet.pdf |

Cisco UCS C-Series rack server options

Cisco UCS C-Series blades are available in one-rack and two-rack unit (RU) varieties, with various CPU, memory, and I/O options. The part numbers listed in the following table are for the base server; they do not include CPUs, memory, disk drives, PCIe cards, or the Cisco FEX. Multiple configuration options are available and supported in FlexPod.

| Cisco UCS C-Series rack server | Part number | Technical specifications |
|--------------------------------|---------------|---|
| Cisco UCS C220 M4 | UCSC-C220-M4S | http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c220m4-sff-spec-sheet.pdf |
| Cisco UCS C240 M4 | UCSC-C240-M4S | http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c240m4-sff-spec-sheet.pdf |
| Cisco UCS C460 M4 | UCSC-C460-M4 | http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c460m4_specsheets.pdf |

Cisco Nexus switches

Redundant switches are required for all FlexPod Express architectures.

The FlexPod Express with NetApp AFF or FAS architecture is built with the Cisco Nexus 31108 switch. FlexPod Express with the Cisco UCS Mini (Cisco UCS- managed) architecture is validated by using the Cisco Nexus 3524 switch. This configuration can also be deployed with a standard switch.

The FlexPod Express with E-Series can be deployed with a standard switch.

The following table lists the part numbers for the Cisco Nexus series chassis; they do not include additional SFP or add-on modules.

| Cisco Nexus Series switch | Part number | Technical specifications |
|---------------------------|-----------------|---|
| Cisco Nexus 3048 | N3K-C3048TP-1GE | http://www.cisco.com/c/en/us/products/collateral/switches/nexus-3000-series-switches/data_sheet_c78-685363.html |
| Cisco Nexus 31108 | N3K-C31108PC-V | http://www.cisco.com/c/en/us/products/switches/nexus-31108pc-v-switch/index.html |
| Cisco Nexus 9396 | N9K-C9396PX | http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-729405.html |
| Cisco Nexus 3172 | N3K-C3172 | https://www.cisco.com/c/en/us/products/collateral/switches/nexus-3000-series-switches/data_sheet_c78-729483.html |

Cisco Support licensing options

Valid SMARTnet support contracts are required on all Cisco equipment in the FlexPod Express architecture.



The licenses required and the part numbers for those licenses should be verified by your sales representative because they can differ for different products.

The following table lists the Cisco support licensing options.

| Cisco Support licensing | License guide |
|-------------------------|---|
| SMARTnet 24x7x4 | http://www.cisco.com/web/services/portfolio/product-technical-support/smartnet/index.html |

NetApp components

NetApp storage controllers provide the storage foundation in the FlexPod Express architecture for both boot and application data storage. This section lists the different NetApp options in the FlexPod Express architecture.

NetApp storage controller options

NetApp FAS

Redundant AFF C190, AFF A220, or FAS2750 series controllers are required in the FlexPod Express architecture. The controllers run ONTAP software. When ordering the storage controllers, the preferred software version can be preloaded on the controllers. For ONTAP, the cluster can be deployed either with a pair of cluster interconnect switches or in a switchless cluster configuration.

The part numbers listed in the following table are for an empty controller. Different options and configurations are available based on the storage platform selected. Consult your sales representative for details about these additional components.

| Storage controller | FAS part number | Technical specifications |
|--------------------|------------------------------------|---|
| FAS2750 | Based on individual options chosen | https://www.netapp.com/us/products/storage-systems/hybrid-flash-array/fas2700.aspx |
| FAS2720 | Based on individual options chosen | https://www.netapp.com/us/products/storage-systems/hybrid-flash-array/fas2700.aspx |
| AFF C190 | Based on individual options chosen | https://www.netapp.com/us/products/entry-level-aff.aspx |
| AFF A220 | Based on individual options chosen | https://www.netapp.com/us/documentation/all-flash-fas.aspx |
| FAS2620 | Based on individual options chosen | http://www.netapp.com/us/products/storage-systems/fas2600/fas2600-tech-specs.aspx |
| FAS2650 | Based on individual options chosen | http://www.netapp.com/us/products/storage-systems/fas2600/fas2600-tech-specs.aspx |

E-Series storage

An HA pair of NetApp E2800 series controllers is required in the FlexPod Express architecture. The controllers run the SANtricity OS.

The part numbers listed in the following table are for an empty controller. Different options and configurations are available based on the storage platform selected. Consult your sales representative for details about these additional components.

| Storage controller | Part number | Technical specifications |
|--------------------|------------------------------------|---|
| E2800 | Based on individual options chosen | http://www.netapp.com/us/products/storage-systems/e2800/e2800-tech-specs.aspx |

NetApp Ethernet expansion modules

NetApp FAS

The following table lists the NetApp FAS10GbE adapter options.

| Component | Part number | Technical specifications |
|---------------|-------------|---|
| NetApp X1117A | X1117A-R6 | https://library.netapp.com/ecm/ecm_download_file/ECMM1280307 |



The FAS2500 and 2600 series storage systems have onboard 10GbE ports.

The NetApp X1117A adapter is for FAS8020 storage systems.

E-Series storage

The following table lists the E-Series 10GbE adapter options.

| Component | Part number |
|----------------------------|-----------------|
| 10GbE iSCSI/16Gb FC 4-port | X-56025-00-0E-C |
| 10GbE iSCSI/16Gb FC 2-port | X-56024-00-0E-C |



The E2824 series storage systems have onboard 10GbE ports.

The 10GbE iSCSI/16Gb FC 4-port host interface card (HIC) can be used for additional port density.

The onboard ports and the HIC can function as iSCSI adapters or FC adapters depending on the feature activated in SANtricity OS.

For more information about supported adapter options, see the Adapter section of [NetApp Hardware Universe](#).

NetApp disk shelves and disks

NetApp FAS

A minimum of one NetApp disk shelf is required for storage controllers. The NetApp shelf type selected determines which drive types are available within that shelf.

The FAS2700 and FAS2600 series of controllers are offered as a configuration that includes dual storage controllers plus disks housed within the same chassis. This configuration is offered with SATA or SAS drives; therefore, additional external disk shelves are not needed unless performance or capacity requirements dictate more spindles.



All disk shelf part numbers are for the empty shelf with two AC PSUs. Consult your sales representative for additional part numbers.

Disk drive part numbers vary according to the size and form factor of the disk you intend to purchase. Consult your sales representative for additional part numbers.

The following table lists the NetApp disk shelf options, along with the drives supported for each shelf type, which can be found on NetApp Hardware Universe. Follow the Hardware Universe link, select the version of ONTAP you are using, then select the shelf type. Under the shelf image, click Supported Drives to see the drives supported for specific versions of ONTAP and the disk shelves.

| Disk shelf | Part number | Technical specifications |
|------------|-------------|---|
| DS212C | DS212C-0-12 | Disk Shelves and Storage Media Technical Specifications Supported Drives on NetApp Hardware Universe |
| DS224C | DS224C-0-24 | Disk Shelves and Storage Media Technical Specifications Supported Drives on NetApp Hardware Universe |

| Disk shelf | Part number | Technical specifications |
|------------|-------------------|--|
| DS460C | DS460C-0-60 | Disk Shelves and Storage Media Technical Specifications Supported Drives on NetApp Hardware Universe |
| DS2246 | X559A-R6 | Disk Shelves and Storage Media Technical Specifications Supported Drives on NetApp Hardware Universe |
| DS4246 | X24M-R6 | Disk Shelves and Storage Media Technical Specifications Supported Drives on NetApp Hardware Universe |
| DS4486 | DS4486-144TB-R5-C | Disk Shelves and Storage Media Technical Specifications Supported Drives on NetApp Hardware Universe |

E-Series storage

A minimum of one NetApp disk shelf is required for storage controllers that do not house any drives in their chassis. The NetApp shelf type selected determines which drive types are available within that shelf.

The E2800 series of controllers are offered as a configuration that includes dual storage controllers plus disks housed within a supported disk shelf. This configuration is offered with SSD or SAS drives.



Disk drive part numbers vary according to the size and form factor of the disk you intend to purchase. Consult your sales representative for additional part numbers.

The following table lists the NetApp disk shelf options and the drives supported for each shelf type, which can be found on NetApp Hardware Universe. Follow the Hardware Universe link, select the version of ONTAP you are using, then select the shelf type. Under the shelf image, click Supported Drives to see the drives supported for specific versions of ONTAP and the disk shelves.

| Disk shelf | Part number | Technical specifications |
|------------|------------------|--|
| DE460C | E-X5730A-DM-0E-C | Disk Shelves Technical Specifications Supported Drives on NetApp Hardware Universe |
| DE224C | E-X5721A-DM-0E-C | Disk Shelves Technical Specifications Supported Drives on NetApp Hardware Universe |
| DE212C | E-X5723A-DM-0E-C | Disk Shelves Technical Specifications Supported Drives on NetApp Hardware Universe |

NetApp software licensing options

NetApp FAS

The following table lists the NetApp FAS software licensing options.

| NetApp Software Licensing | Part Number | Technical Specifications |
|---------------------------|--|--------------------------|
| Base cluster license | Consult your NetApp sales team for more licensing information. | |

E-Series storage

The following table lists the E-Series software licensing options.

| NetApp software licensing | Part number | Technical specifications |
|---------------------------|--|--------------------------|
| Standard features | Consult your NetApp sales team for more licensing information. | |
| Premium features | | |

NetApp Support licensing options

SupportEdge Premium licenses are required, and the part numbers for those licenses vary based on the options selected in the FlexPod Express design.

NetApp FAS

The following table lists the NetApp support licensing options for NetApp FAS.

| NetApp Support licensing | Part number | Technical specifications |
|---|-------------|---|
| SupportEdge Premium4 hours onsite; months: 36 | CS-O2-4HR | https://www.netapp.com/pdf.html?item=/media/19784-ds-3873.pdf |

E-Series storage

The following table lists the NetApp support licensing options for E-Series storage.

| NetApp Support licensing | Part number | Technical specifications |
|---|-------------------|---|
| Hardware support Premium 4 hours onsite; months: 36 | SVC-O2-4HR-E | https://www.netapp.com/pdf.html?item=/media/19784-ds-3873.pdf |
| Software support | SW-SSP-O2-4HR-E | |
| Initial installation | SVC-INST-O2-4HR-E | |

Power and cabling requirements

This section describes the power and minimum cabling requirements for a FlexPod Express design.

Power requirements

The power requirements are based on U.S. specifications and assume the use of AC power. Other countries might have different power requirements. Direct current (DC) power options are also available for most components. For additional data about the maximum power required as well as other detailed power

information, consult the detailed technical specifications for each hardware component.

For detailed Cisco UCS power data, see the [Cisco UCS Power Calculator](#).

The following table lists the power ports required per device.

| Cisco Nexus switches | Power cables required |
|----------------------|---|
| Cisco Nexus 3048 | 2x C13/C14 power cables for each Cisco Nexus 3000 series switch |
| Cisco Nexus 3524 | 2x C13/C14 power cables for each Cisco Nexus 3000 series switch |
| Cisco Nexus 9396 | 2x C13/C14 power cables for each Cisco Nexus 9000 series switch |

| Cisco UCS chassis | Power cables required |
|-------------------|---|
| Cisco UCS 5108 | 2 CAB-US515P-C19-US/CAB-US520-C19-US for each Cisco UCS chassis |

| Cisco UCS B-Series servers | Power cables required |
|----------------------------|---|
| Cisco UCS B200 M4 | N/A; blade server is powered by chassis |
| Cisco UCS B420 M4 | N/A; blade server is powered by chassis |
| Cisco UCS B200 M5 | N/A; blade server is powered by chassis |
| Cisco UCS B480 M5 | N/A; blade server is powered by chassis |

| Cisco UCS C-Series servers | Power ports required |
|----------------------------|--|
| Cisco UCS C220 M4 | 2 x C13/C14 power cables for each Cisco UCS server |
| Cisco UCS C240 M4 | |
| Cisco UCS C460 M4 | |
| Cisco UCS C220 M5 | |
| Cisco UCS C240 M5 | |
| Cisco UCS C480 M5 | |

| NetApp FAS controllers | Power ports required (per HA pair) |
|------------------------|------------------------------------|
| FAS2554 | 2 x C13/C14 |
| FAS2552 | 2 x C13/C14 |
| FAS2520 | 2 x C13/C14 |
| FAS8020 | 2 x C13/C14 |

| E-Series controllers | Power ports required (per HA pair) |
|----------------------|------------------------------------|
| E2824 | 2 x C14/C20 |

| NetApp FAS disk shelves | Power ports required |
|-------------------------|----------------------|
| DS212C | 2 x C13/C14 |
| DS224C | 2 x C13/C14 |
| DS460C | 2 x C13/C14 |
| DS2246 | 2 x C13/C14 |
| DS4246 | 4 x C13/C14 |

| E-Series disk shelves | Power ports required |
|-----------------------|----------------------|
| DE460C | 2 x C14/C20 |
| DE224C | 2 x C14/C20 |
| DE212C | 2 x C14/C20 |

Minimum cable requirements

This section describes the minimum cable requirements for a FlexPod Express design. Most FlexPod implementations require additional cables, but the number varies based on the deployment size and scope.

The following table lists the minimum number of cables required for each device.

| Cisco Nexus 3000 Series switches | Cables required |
|----------------------------------|---|
| Cisco Nexus 31108 | At least two 10GbE fiber or Twinax cables per switch |
| Cisco Nexus 3172PQ | |
| Cisco Nexus 3048 | |
| Cisco Nexus 3524 | |
| Cisco Nexus 9396 | |
| DS212C | Number of SAS cables depends on specific configuration of disk shelves |
| DS2246 | |
| DS460C | |
| DS224C | |
| DS4246 | |
| E2800 | <ul style="list-style-type: none"> • At least one Gigabit Ethernet (1GbE) cable for management per controller • At least two 10GbE cables per controller (for iSCSI) or two FC cables matching speed requirements |
| DE460C | 2 x mini-SAS HD cables per disk shelf |
| DE224C | 2 x mini-SAS HD cables per disk shelf |
| DE212C | 2 x mini-SAS HD cables per disk shelf |

Technical Specifications and References

This section describes additional important technical specifications for each of the FlexPod Express components.

Cisco UCS B-Series blade servers

The following table lists the Cisco UCS B-Series blade server options.

| Component | Cisco UCS B200 M4 | Cisco UCS B420 M4 | Cisco UCS B200 M5 |
|-----------------------------|---------------------------------|-------------------------------|---|
| Processor support | Intel Xeon E5-2600 | Intel Xeon E5-4600 | Intel Xeon Scalable processors |
| Maximum memory capacity | 24 DIMMs for a maximum of 768GB | 48 DIMMs for a maximum of 3TB | 24 DIMMs for a maximum of 3072GB |
| Memory size and speed | 32GB DDR4; 2133MHz | 64GB DDR4; 2400MHz | 16GB, 32GB, 64GB, and 128GB DDR4; 2666MHz |
| SAN boot support | Yes | Yes | Yes |
| Mezzanine I/O adapter slots | 2 | 3 | 2, front and rear, including GPU support |
| I/O maximum throughput | 80Gbps | 160Gbps | 80Gbps |

Cisco UCS C-Series rack servers

The following table lists Cisco UCS C-Series rack server options.

| Component | Cisco UCS C220 M4 | Cisco UCS C240 M4 | Cisco UCS C460 M4 | Cisco UCS C220 M5 |
|-------------------------|-----------------------------|----------------------------------|---------------------------------------|---|
| Processor support | 1 or 2 Intel E5-2600 series | 1 or 2 Intel Xeon E5-2600 series | 2 or 4 Intel Xeon E7-4800/8800 series | Intel Xeon Scalable processors (1 or 2) |
| Maximum memory capacity | 1.5GB | 1.5TB | 6TB | 3072GB |
| PCIe slots | 2 | 6 | 10 | 2 |
| Form factor | 1RU | 2RU | 4RU | 1 RU |

The following table lists the datasheets for the Cisco UCS C-Series rack server options.

| Component | Cisco UCS datasheet |
|-------------------|---|
| Cisco UCS C220 M4 | http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c220m4-sff-spec-sheet.pdf |
| Cisco UCS C240 M4 | http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c240-m4-rack-server/datasheet-c78-732455.html |

| Component | Cisco UCS datasheet |
|-------------------|---|
| Cisco UCS C460 M4 | http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c460-m4-rack-server/datasheet-c78-730907.html |
| Cisco UCS C220 M5 | https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-c-series-rack-servers/c220m5-sff-specsheet.pdf |

Cisco Nexus 3000 Series switches

The following table lists the Cisco Nexus 3000 series switch options.

| Component | Cisco Nexus 3048 | Cisco Nexus 3524 | Cisco Nexus 31108 | Cisco Nexus 3172PQ |
|---------------------|------------------|------------------|--------------------|--|
| Form factor | 1RU | 1RU | 1RU | 1 RU |
| Maximum 1Gbps ports | 48 | 24 | 48 (10/40/100Gbps) | 72 1/10GbE ports, or 48 1/10GbE plus six 40GbE ports |
| Forwarding rate | 132Mbps | 360Mbps | 1.2Bpps | 1Bpps |
| Jumbo frame support | Yes | Yes | Yes | Yes |

The following table lists the datasheets for the Cisco Nexus 3000 series switch options.

| Component | Cisco Nexus Datasheet |
|-----------------------|---|
| Cisco Nexus 31108 | http://www.cisco.com/c/en/us/products/switches/nexus-31108pc-v-switch/index.html |
| Cisco Nexus 3172PQ | https://www.cisco.com/c/en/us/products/switches/nexus-3172pq-switch/index.html |
| Cisco Nexus 3048 | https://www.cisco.com/c/en/us/products/switches/nexus-3048-switch/index.html |
| Cisco Nexus 3172PQ-XL | https://www.cisco.com/c/en/us/products/switches/nexus-3172pq-switch/index.html |
| Cisco Nexus 3548 XL | https://www.cisco.com/c/en/us/products/switches/nexus-3548-x-switch/index.html |
| Cisco Nexus 3524 XL | https://www.cisco.com/c/en/us/products/switches/nexus-3524-x-switch/index.html |
| Cisco Nexus 3548 | https://www.cisco.com/c/en/us/products/switches/nexus-3548-x-switch/index.html |
| Cisco Nexus 3524 | https://www.cisco.com/c/en/us/products/switches/nexus-3524-x-switch/index.html |

The following table lists the Cisco Nexus 9000 series switch options.

| Component | Cisco Nexus 9396 | Cisco Nexus 9372 |
|--------------------------|------------------|------------------|
| Form factor | 2RU | 1RU |
| Maximum ports | 60 | 54 |
| 10Gbps SFP+ uplink ports | 48 | 48 |

The following table lists the Cisco Nexus 9000 series switch options datasheets.

| Component | Cisco Nexus datasheet |
|------------------|---|
| Cisco Nexus 9396 | http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736967.html |
| Cisco Nexus 9372 | http://www.cisco.com/c/en/us/products/collateral/switches/nexus-9000-series-switches/datasheet-c78-736967.html |
| Nexus 9396X | https://www.cisco.com/c/en/us/products/switches/nexus-9396px-switch/index.html?dtdid=ossdc000283 |

NetApp FAS storage controllers

The following table lists the current NetApp FAS storage controller options.

| Current component | FAS2620 | FAS2650 |
|--|--------------------------------|-------------------------------|
| Configuration | 2 controllers in a 2U chassis | 2 controllers in a 4U chassis |
| Maximum raw capacity | 1440TB | 1243TB |
| Internal drives | 12 | 24 |
| Maximum number of drives (internal plus external) | 144 | 144 |
| Maximum volume size | 100TB | |
| Maximum aggregate size | 4TB | |
| Maximum number of LUNs | 2,048 per controller | |
| Storage networking supported | iSCSI, FC, FCoE, NFS, and CIFS | |
| Maximum number of NetApp FlexVol volumes | 1,000 per controller. | |
| Maximum number of NetApp Snapshot copies | 255,000 per controller | |
| Maximum NetApp Flash Pool intelligent data caching | 24TB | |



For details about the FAS storage controller option, see the [FAS models](#) section of the Hardware Universe. For AFF, see [AFF models](#) section.

The following table lists the characteristics of a FAS8020 controller system.

| Component | FAS8020 |
|---|-------------------------------|
| Configuration | 2 controllers in a 3U chassis |
| Maximum raw capacity | 2880TB |
| Maximum number of drives | 480 |
| Maximum volume size | 70TB |
| Maximum aggregate size | 324TB |
| Maximum number of LUNs | 8,192 per controller |
| Storage networking supported | iSCSI, FC, NFS, and CIFS |
| Maximum number of FlexVol volumes | 1,000 per controller |
| Maximum number of Snapshot copies | 255,000 per controller |
| Maximum NetApp Flash Cache intelligent data caching | 3TB |
| Maximum Flash Pool data caching | 24TB |

The following table lists the datasheets for NetApp storage controllers.

| Component | Storage controller datasheet |
|----------------|---|
| FAS2600 series | http://www.netapp.com/us/products/storage-systems/fas2600/fas2600-tech-specs.aspx |
| FAS2500 series | http://www.netapp.com/us/products/storage-systems/fas2500/fas2500-tech-specs.aspx |
| FAS8000 series | http://www.netapp.com/us/products/storage-systems/fas8000/fas8000-tech-specs.aspx |

NetApp FAS Ethernet adapters

The following table lists NetApp FAS 10GbE adapters.

| Component | X1117A-R6 |
|--------------|-----------------|
| Port count | 2 |
| Adapter type | SFP+ with fibre |

The X1117A-R6 SFP+ adapter is supported on FAS8000 series controllers.

The FAS2600 and FAS2500 series storage systems have onboard 10GbE ports. For more information, see the [NetApp 10GbE adapter datasheet](#).



For more adapter details based on the AFF or FAS model, see the [Adapter section](#) in the Hardware Universe.

NetApp FAS disk shelves

The following table lists the current NetApp FAS disk shelf options.

| Component | DS460C | DS224C | DS212C | DS2246 | DS4246 |
|----------------------|------------------------|------------------------|------------------------|------------------------|------------------------|
| Form factor | 4RU | 2RU | 2RU | 2RU | 4RU |
| Drives per enclosure | 60 | 24 | 12 | 24 | 24 |
| Drive form factor | 3.5" large form factor | 2.5" small form factor | 3.5" large form factor | 2.5" small form factor | 3.5" large form factor |
| Shelf I/O modules | Dual IOM12 modules | Dual IOM12 modules | Dual IOM12 modules | Dual IOM6 modules | Dual IOM6 modules |

For more information, see the [NetApp disk shelves datasheet](#).



For more information about the disk shelves, see the NetApp Hardware Universe [Disk Shelves section](#).

NetApp FAS disk drives

The technical specifications for the NetApp disks include form factor size, disk capacity, disk RPM, supporting controllers, and Data ONTAP version requirements and are located in the Drives section on [NetApp Hardware Universe](#).

E-Series storage controllers

The following table lists the current E-Series storage controller options.

| Current Component | E2812 | E2824 | E2860 |
|---|-------------------------------|-------------------------------|-------------------------------|
| Configuration | 2 controllers in a 2U chassis | 2 controllers in a 2U chassis | 2 controllers in a 4U chassis |
| Maximum raw capacity | 1800TB | 1756.8TB | 1800TB |
| Internal drives | 12 | 24 | 60 |
| Maximum number of drives (internal plus external) | 180 | | |
| Maximum SSD | 120 | | |
| Maximum volume size for disk pool volume | 1024TB | | |
| Maximum disk pools | 20 | | |
| Storage networking supported | iSCSI and FC | | |
| Maximum number of volumes | 512 | | |

The following table lists the datasheets for the current E-Series storage controller.

| Component | Storage controller datasheet |
|-----------|---|
| E2800 | http://www.netapp.com/us/media/ds-3805.pdf |

E-Series adapters

The following table lists the E-Series adapters.

| Component | X-56023-00-0E-C | X-56025-00-0E-C | X-56027-00-0E-C | X-56024-00-0E-C | X-56026-00-0E-C |
|--------------|-----------------|------------------------|-----------------|------------------------|-----------------|
| Port count | 2 | 4 | 4 | 2 | 2 |
| Adapter type | 10Gb Base-T | 16G FC and 10GbE iSCSI | SAS | 16G FC and 10GbE iSCSI | SAS |

E-Series disk shelves

The following table lists the E-Series disk shelf options.

| Component | DE212C | DE224C | DE460C |
|----------------------|--------------------------------|--------|--------------------------------|
| Form factor | 2RU | 2RU | 4RU |
| Drives per enclosure | 12 | 24 | 60 |
| Drive form factor | 2.5" small form factor 3.5" | 2.5" | 2.5" small form factor 3.5" |
| Shelf I/O modules | IOM12 | IOM12 | IOM12 |

E-Series disk drives

The technical specifications for the NetApp disk drives include form factor size, disk capacity, disk RPM, supporting controllers, and SANtricity version requirements and are located in the Drives section on [NetApp Hardware Universe](#).

Previous architectures and equipment

FlexPod is a flexible solution allowing customers to use both existing and new equipment currently for sale by Cisco and NetApp. Occasionally, certain models of equipment from both Cisco and NetApp are designated end of life.

Even though these models of equipment are no longer available, customers who bought one of these models before the end-of-sale date can use that equipment in a FlexPod configuration.

Additionally, FlexPod Express architectures are periodically refreshed to introduce the latest hardware and software from Cisco and NetApp to the FlexPod Express solution. This section lists the previous FlexPod Express architectures and hardware used within them.

Previous FlexPod Express architectures


This section describes the previous FlexPod Express architectures.

FlexPod Express small and medium configurations


The FlexPod Express small and medium configurations include the following components:

- Two Cisco Nexus 3048 switches in a redundant configuration
- At least two Cisco UCS C-Series rack mount servers
- Two FAS2200 or FAS2500 series controllers in an HA pair configuration

The following figure illustrates the FlexPod Express small configuration.



The following figure illustrates the FlexPod Express medium configuration.

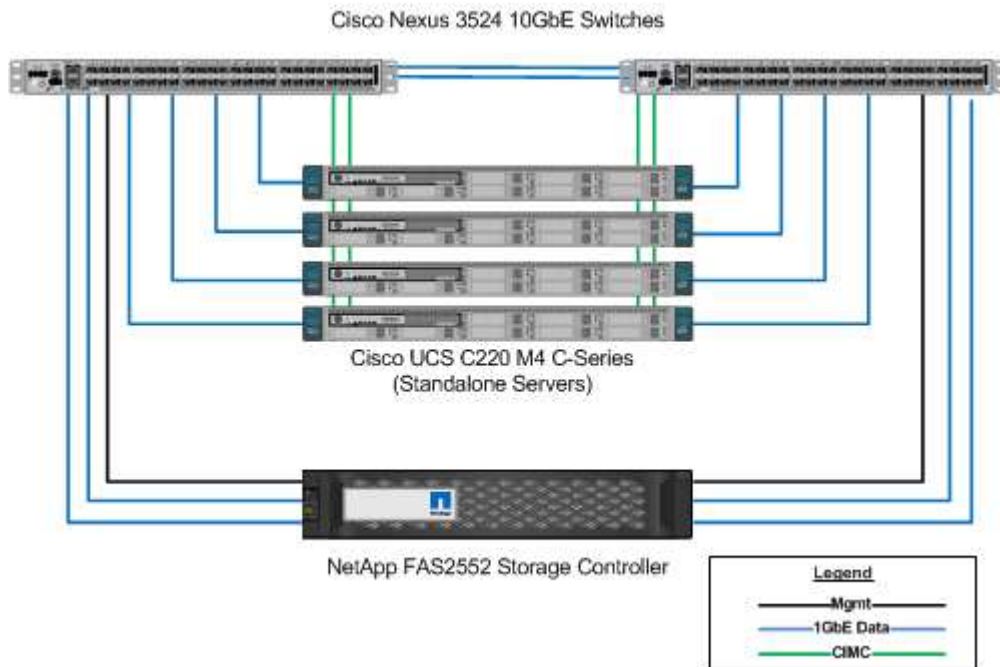


FlexPod Express large configuration

The FlexPod Express large configuration includes the following components:

- Two Cisco Nexus 3500 series or Cisco Nexus 9300 series switches in a redundant configuration
- At least two Cisco UCS C-Series rack mount servers
- Two FAS2552, FAS2554, or FAS8020 controllers in an HA pair configuration (requires two 10GbE ports per controller)
- One NetApp disk shelf with any supported disk type (when the FAS8020 is used)

The following figure illustrates the FlexPod Express large configuration.



Previous FlexPod Express verified architectures

Previous FlexPod Express verified architectures are still supported. The architecture and deployment documents include:

- [FlexPod Express with Cisco UCS C-Series and NetApp FAS2500 Series](#)
- [FlexPod Express with VMware vSphere 6.0: Small and Medium Configurations](#)
- [FlexPod Express with VMware vSphere 6.0: Large Configuration](#)
- [FlexPod Express with Microsoft Windows Server 2012 R2 Hyper-V: Small and Medium Configurations](#)
- [FlexPod Express with Microsoft Windows Server 2012 R2 Hyper-V: Large Configuration](#)

Previous hardware

The following table lists the hardware used in previous FlexPod Express architectures.

| Hardware used in previous architectures | Technical specifications (if available) |
|---|---|
| Cisco UCS C220 M3 | http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c220-m3-rack-server/data_sheet_c78-700626.html |
| Cisco UCS C24 M3 | http://www.cisco.com/en/US/prod/collateral/ps10265/ps10493/data_sheet_c78-706103.html |
| Cisco UCS C22 M3 | http://www.cisco.com/en/US/prod/collateral/ps10265/ps10493/data_sheet_c78-706101.html |
| Cisco UCS C240 M3 | http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-c240-m3-rack-server/data_sheet_c78-700629.html |
| Cisco UCS C260 M2 | http://www.cisco.com/en/US/prod/collateral/ps10265/ps10493/c260m2_specsheet.pdf |
| Cisco UCS C420 M3 | http://www.cisco.com/en/US/products/ps12770/index.html |
| Cisco UCS C460 M2 | http://www.cisco.com/en/US/prod/collateral/ps10265/ps10493/ps11587/spec_sheet_c17-662220.pdf |
| Cisco UCS B200 M3 | http://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b200-m3-blade-server/data_sheet_c78-700625.html |
| Cisco UCS B420 M3 | N/A |
| Cisco UCS B22 M3 | http://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/b22m3_specsheet.pdf |
| Cisco Nexus 3524 | http://www.cisco.com/c/en/us/products/switches/nexus-3524-switch/index.html |
| FAS2240 | |
| FAS2220 | http://www.netapp.com/us/products/storage-systems/fas2200/fas2200-tech-specs.aspx |
| DS4243 | N/A |

Legacy equipment

The following table lists the NetApp legacy storage controller options.

| Storage controller | FAS part number | Technical specifications |
|--------------------|------------------------------------|---|
| FAS2520 | Based on individual options chosen | http://www.netapp.com/us/products/storage-systems/fas2500/fas2500-tech-specs.aspx |
| FAS2552 | Based on individual options chosen | http://www.netapp.com/us/products/storage-systems/fas2500/fas2500-tech-specs.aspx |

| Storage controller | FAS part number | Technical specifications |
|--------------------|------------------------------------|---|
| FAS2554 | Based on individual options chosen | http://www.netapp.com/us/products/storage-systems/fas2500/fas2500-tech-specs.aspx |
| FAS8020 | Based on individual options chosen | http://www.netapp.com/us/products/storage-systems/fas8000/fas8000-tech-specs.aspx |

The following table lists the NetApp legacy disk shelf options for NetApp FAS.

| Disk shelf | Part number | Technical specifications |
|------------|---------------------|--|
| DE1600 | E-X5682A-DM-0E-R6-C | Disk Shelves Technical Specifications Supported Drives on NetApp Hardware Universe |
| DE5600 | E-X4041A-12-R6 | Disk Shelves Technical Specifications Supported Drives on NetApp Hardware Universe |
| DE6600 | X-48564-00-R6 | Disk Shelves Technical Specifications Supported Drives on NetApp Hardware Universe |

NetApp legacy FAS controllers

The following table lists the legacy NetApp FAS controller options.

| Current component | FAS2554 | FAS2552 | FAS2520 |
|---|--------------------------------|-------------------------------|-------------------------------|
| Configuration | 2 controllers in a 4U chassis | 2 controllers in a 2U chassis | 2 controllers in a 2U chassis |
| Maximum raw capacity | 576TB | 509TB | 336TB |
| Internal drives | 24 | 24 | 12 |
| Maximum number of drives (internal plus external) | 144 | 144 | 84 |
| Maximum volume size | 60TB | | |
| Maximum aggregate size | 120TB | | |
| Maximum number of LUNs | 2,048 per controller | | |
| Storage networking supported | iSCSI, FC, FCoE, NFS, and CIFS | | iSCSI, NFS, and CIFS |
| Maximum number of NetApp FlexVol volumes | 1,000 per controller | | |
| Maximum number of NetApp Snapshot copies | 255,000 per controller | | |



For more NetApp FAS models, see the [FAS models section](#) in the Hardware Universe.

Additional Information

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

- AFF and FAS System Documentation Center

<https://docs.netapp.com/platstor/index.jsp>

- AFF Documentation Resources page

<https://www.netapp.com/us/documentation/all-flash-fas.aspx>

- FAS Storage Systems Documentation Resources page

<https://www.netapp.com/us/documentation/fas-storage-systems.aspx>

- FlexPod

<https://flexpod.com/>

- NetApp documentation

<https://docs.netapp.com>

FlexPod Datacenter Technical Specifications

TR-4036: FlexPod Datacenter Technical Specifications

Arvind Ramakrishnan, and Jyh-shing Chen, NetApp

The FlexPod platform is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp storage controllers (AFF, ASA, or FAS systems).

FlexPod is a suitable platform for running a variety of virtualization hypervisors as well as bare-metal operating systems and enterprise workloads. FlexPod delivers not only a baseline configuration, but also the flexibility to be sized and optimized to accommodate many different use cases and requirements.



Before you order a complete FlexPod configuration, see the [FlexPod Converged Infrastructure](#) page on netapp.com for the latest version of these technical specifications.

Next: [FlexPod platforms](#).

FlexPod platforms

There are two FlexPod platforms:

- **FlexPod Datacenter.** This platform is a massively scalable virtual data center infrastructure that is suited

for workload enterprise applications; virtualization; virtual desktop infrastructure (VDI); and public, private, and hybrid cloud workloads.

- **FlexPod Express.** This platform is a compact converged infrastructure that is targeted to remote office and edge use cases. FlexPod Express has its own specifications that are documented in the [FlexPod Express Technical Specifications](#).

This document provides the technical specifications for the FlexPod Datacenter platform.

FlexPod rules

The FlexPod design enables a flexible infrastructure that encompasses many different components and software versions.

Use the rule sets as a guide for building or assembling a valid FlexPod configuration. The numbers and rules that are listed in this document are the minimum requirements for a FlexPod configuration. They can be expanded in the included product families as required for different environments and use cases.

Supported versus validated FlexPod configurations

The FlexPod architecture is defined by the set of rules that are described in this document. The hardware components and software configurations must be supported by the [Cisco UCS Hardware and Software Compatibility List](#) and the [NetApp Interoperability Matrix Tool \(IMT\)](#).

Each Cisco Validated Design (CVD) or NetApp Verified Architecture (NVA) is a possible FlexPod configuration. Cisco and NetApp document these configuration combinations and validate them with extensive end-to-end testing. The FlexPod deployments that deviate from these configurations are fully supported if they follow the guidelines in this document and if all the components are listed as compatible in the Cisco UCS Hardware and Software Compatibility List and the NetApp [IMT](#).

For example, adding more storage controllers or Cisco UCS Servers and upgrading software to newer versions are fully supported if the software, hardware, and configurations meet the guidelines that are defined in this document.

NetApp ONTAP

NetApp ONTAP software is installed on all NetApp FAS, AFF, and AFF All SAN Array (ASA) systems. FlexPod is validated with ONTAP software, providing a highly scalable storage architecture that enables nondisruptive operations, nondisruptive upgrades, and an agile data infrastructure.

For more information about ONTAP, see the [ONTAP Data Management Software](#) product page.

Cisco Nexus switching modes of operation

A variety of Cisco Nexus products can be used as the switching component of a given FlexPod deployment. Most of these options leverage the traditional Cisco Nexus OS or NX-OS software. The Cisco Nexus family of switches offers varying capabilities within its product lines. These capabilities are detailed later in this document.

Cisco's offering in the software-defined networking space is called Application Centric Infrastructure (ACI). The Cisco Nexus product line that supports the ACI mode, also called fabric mode, is the Cisco Nexus 9300 series. These switches can also be deployed in NX-OS or standalone mode.

Cisco ACI is targeted at data center deployments that focus on the requirements of a specific application. Applications are instantiated through a series of profiles and contracts that allow connectivity from the host or virtual machine (VM) all the way through the network to the storage.

FlexPod is validated with both modes of operation of the Cisco Nexus switches. For more information about the ACI and the NX-OS modes, see the following Cisco pages:

- [Cisco Application Centric Infrastructure](#)
- [Cisco NX-OS Software](#)

Minimum hardware requirements

A FlexPod Datacenter configuration has minimum hardware requirements, including, but not limited to, switches, fabric interconnects, servers, and NetApp storage controllers.

You must use Cisco UCS Servers. Both C-Series and B-Series Servers have been used in the validated designs. Cisco Nexus Fabric Extenders (FEXs) are optional with C-Series Servers.

A FlexPod configuration has the following minimum hardware requirements:

- Two Cisco Nexus switches in a redundant configuration. This configuration can consist of two redundant switches from the Cisco Nexus 5000, 7000, or 9000 Series. The two switches should be of the same model and should be configured in the same mode of operation.

If you are deploying an ACI architecture, you must observe the following additional requirements:

- Deploy the Cisco Nexus 9000 Series Switches in a leaf-spine topology.
- Use three Cisco Application Policy Infrastructure Controllers (APICs).
- Two Cisco UCS 6200, 6300, or 6400 Series Fabric Interconnects in a redundant configuration.
- Cisco UCS Servers:
 - If the solution uses B-Series Servers, one Cisco UCS 5108 B-Series Blade Server Chassis plus two Cisco UCS B-Series Blade Servers plus two 2104, 2204/8, 2408, or 2304 I/O modules (IOMs).
 - If the solution uses C-Series Servers, two Cisco UCS C-Series Rack Servers.

For larger deployments of Cisco UCS C-Series Rack Servers, you can choose a pair of 2232PP FEX modules. However, the 2232PP is not a hardware requirement.

- Two NetApp storage controllers in a high-availability (HA) pair configuration:

This configuration can consist of any supported NetApp FAS, AFF, or ASA-series storage controllers. See the [NetApp Hardware Universe](#) application for a current list of supported FAS, AFF, and ASA controller models.

- The HA configuration requires two redundant interfaces per controller for data access; the interfaces can be FCoE, FC, or 10/25/100Gb Ethernet (GbE).
- If the solution uses NetApp ONTAP, a cluster interconnect topology that is approved by NetApp is required. For more information, see the [Switches](#) tab of the NetApp Hardware Universe.

- If the solution uses ONTAP, at least two additional 10/25/100GbE ports per controller are required for data access.
- For ONTAP clusters with two nodes, you can configure a two-node switchless cluster.
- For ONTAP clusters with more than two nodes, a pair of cluster interconnect switches are required.
- One NetApp disk shelf with any supported disk type. See the Shelves tab of the [NetApp Hardware Universe](#) for a current list of supported disk shelf models.

Minimum software requirements

A FlexPod configuration has the following minimum software requirements:

- NetApp ONTAP:
 - ONTAP software version requires ONTAP 9.1 or later
- Cisco UCS Manager releases:
 - Cisco UCS 6200 Series Fabric Interconnect—2.2(8a)
 - Cisco UCS 6300 Series Fabric Interconnect—3.1(1e)
 - Cisco UCS 6400 Series Fabric Interconnect—4.0(1)
- Cisco Intersight Managed Mode:
 - Cisco UCS 6400 Series Fabric Interconnect – 4.1(2)
- For Cisco Nexus 5000 Series Switches, Cisco NX-OS software release 5.0(3)N1(1c) or later, including NX-OS 5.1.x
- For Cisco Nexus 7000 Series Switches:
 - The 4-slot chassis requires Cisco NX-OS software release 6.1(2) or later
 - The 9-slot chassis requires Cisco NX-OS software release 5.2 or later
 - The 10-slot chassis requires Cisco NX-OS software release 4.0 or later
 - The 18-slot chassis requires Cisco NX-OS software release 4.1 or later
- For Cisco Nexus 9000 Series Switches, Cisco NX-OS software release 6.1(2) or later



The software that is used in a FlexPod configuration must be listed and supported in the NetApp [IMT](#). Some features might require more recent releases of the software than the ones that are listed.

Connectivity requirements

A FlexPod configuration has the following connectivity requirements:

- A separate 100Mbps Ethernet/1Gb Ethernet out-of-band management network is required for all components.
- NetApp recommends that you enable jumbo frame support throughout the environment, but it is not required.
- The Cisco UCS Fabric Interconnect appliance ports are recommended only for iSCSI and NAS connections.
- No additional equipment can be placed in line between the core FlexPod components.

Uplink connections:

- The ports on the NetApp storage controllers must be connected to the Cisco Nexus 5000, 7000, or 9000 Series Switches to enable support for virtual port channels (vPCs).
- vPCs are required from the Cisco Nexus 5000, 7000, or 9000 Series Switches to the NetApp storage controllers.
- vPCs are required from the Cisco Nexus 5000, 7000, or 9000 Series Switches to the fabric interconnects.
- A minimum two connections are required for a vPC. The number of connections within a vPC can be increased based on the application load and performance requirements.

Direct connections:

- NetApp storage controller ports that are directly connected to the fabric interconnects can be grouped to enable a port channel. vPC is not supported for this configuration.
- FCoE port channels are recommended for end-to-end FCoE designs.

SAN boot:

- FlexPod solutions are designed around a SAN-boot architecture using iSCSI, FC, or FCoE protocols. Using boot-from-SAN technologies provides the most flexible configuration for the data center infrastructure and enables the rich features available within each infrastructure component. Although booting from SAN is the most efficient configuration, booting from local server storage is a valid and supported configuration.
- SAN boot over FC-NVME is not supported.

Other requirements

A FlexPod architecture has the following additional interoperability and support-related requirements:

- All hardware and software components must be listed and supported on the NetApp [IMT](#), the [Cisco UCS Hardware and Software Compatibility List](#), and the Cisco UCS Hardware and Software Interoperability Matrix Tool.
- Valid support contracts are required for all equipment, including:
 - Smart Net Total Care (SmartNet) support for Cisco equipment
 - SupportEdge Advisor or SupportEdge Premium support for NetApp equipment

For more information, see the NetApp [IMT](#).

Optional features


NetApp supports several optional components to further enhance FlexPod Datacenter architectures. Optional components are outlined in the following subsections.

MetroCluster


FlexPod supports both variants of the NetApp MetroCluster software for continuous availability, in either two- or four-node cluster configurations. MetroCluster provides synchronous replication for critical workloads. It requires a dual-site configuration that is connected with Cisco switching. The maximum supported distance between the sites is approximately 186 miles (300km) for MetroCluster FC and increases to approximately 435 miles (700km) for MetroCluster IP. The following figures illustrate a FlexPod Datacenter with NetApp

MetroCluster architecture and FlexPod Datacenter with NetApp MetroCluster IP architecture, respectively.

The following figure depicts FlexPod Datacenter with NetApp MetroCluster architecture.

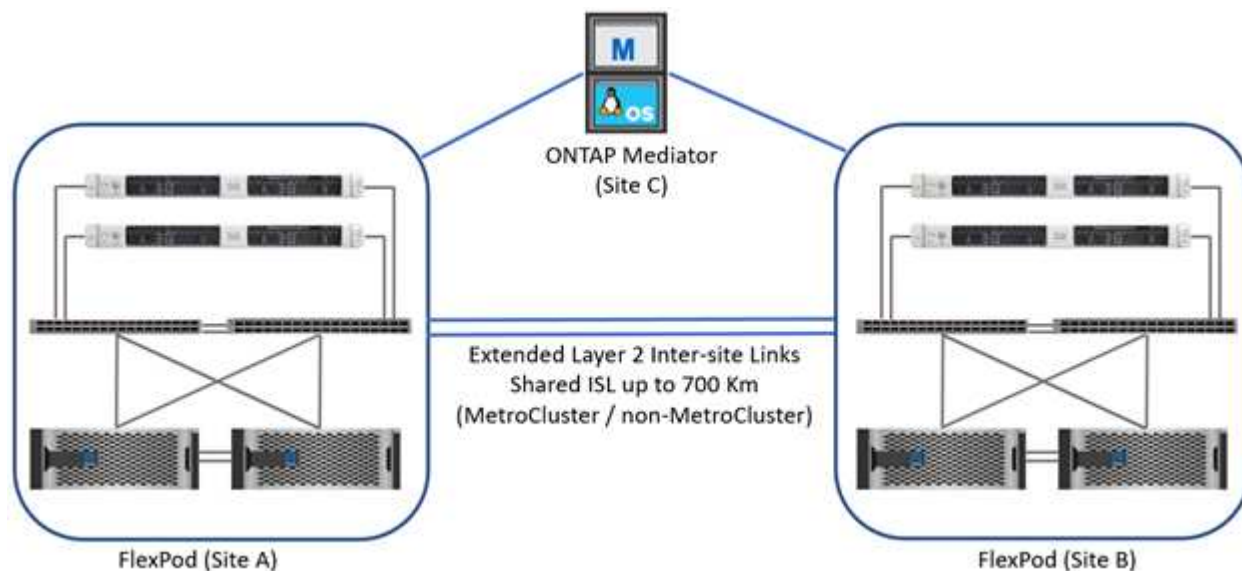


The following figure depicts the FlexPod Datacenter with NetApp MetroCluster IP architecture.



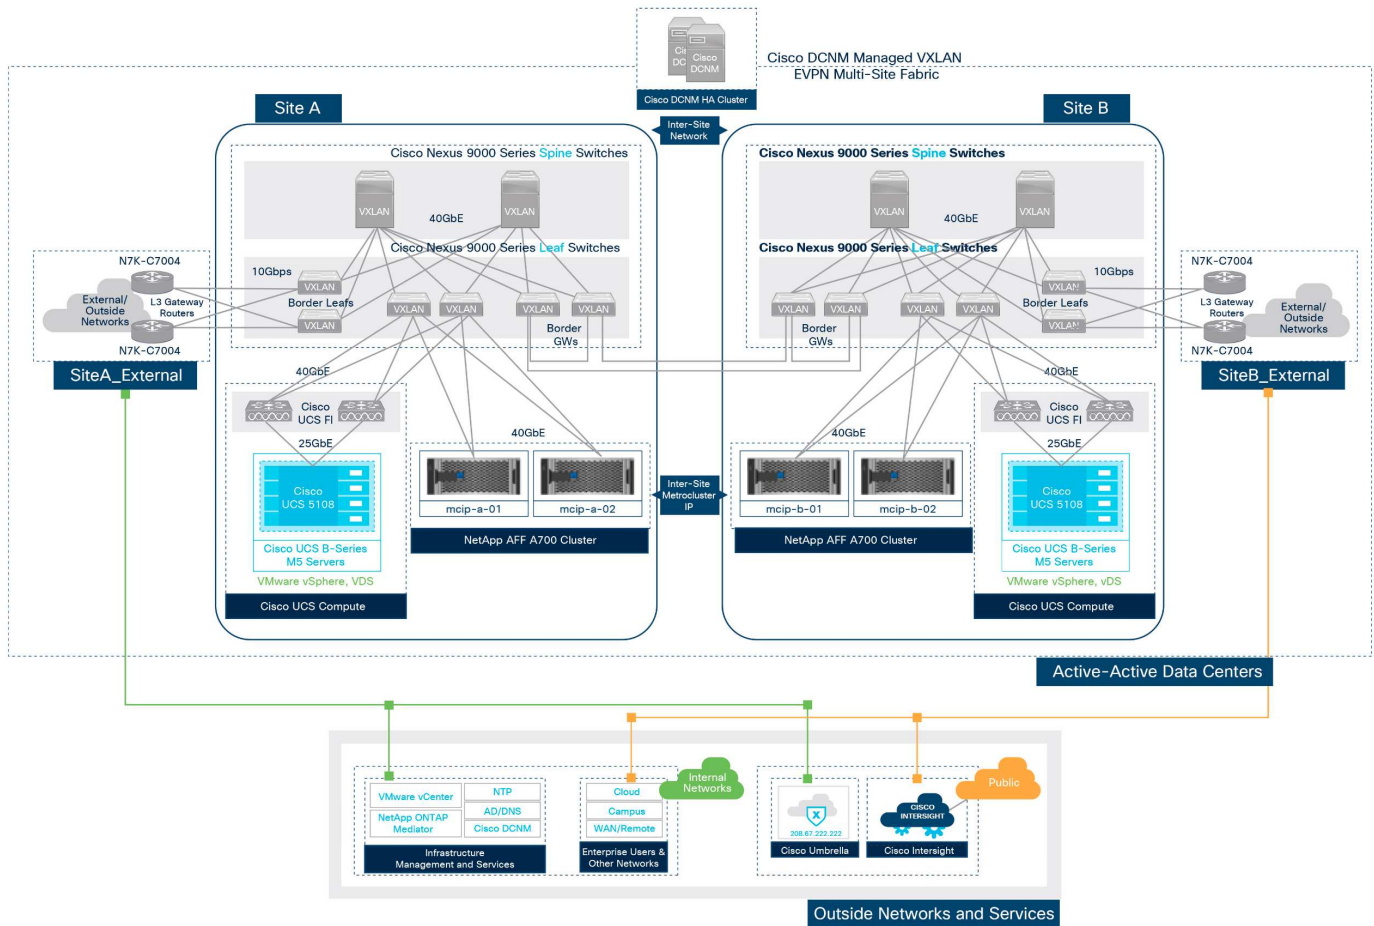
Starting with ONTAP 9.8, ONTAP Mediator can be deployed at a third site to monitor the MetroCluster IP solution and to facilitate automated unplanned switchover when a site disaster occurs.

For a FlexPod MetroCluster IP solution deployment with extended layer-2 site-to-site connectivity, you can achieve cost savings by sharing ISL and using FlexPod switches as compliant MetroCluster IP switches if the network bandwidth and switches meet the requirements as illustrated in the following figure, which depicts the FlexPod MetroCluster IP solution with ISL sharing and compliant switches.




The following two figures depict the VXLAN Multi-Site fabric and the MetroCluster IP storage fabric for a FlexPod MetroCluster IP solution with VXLAN Multi-Site fabric deployment.

- VXLAN Multi-Site fabric for FlexPod MetroCluster IP solution



- MetroCluster IP storage fabric for FlexPod MetroCluster IP solution



End-to-end FC-NVMe


An end-to-end FC-NVMe seamlessly extends a customer's existing SAN infrastructure for real-time applications while simultaneously delivering improved IOPS and throughput with reduced latency.

An existing 32G FC SAN transport can be used to simultaneously transport both NVMe and SCSI workloads.

The following figure illustrates the Flexpod Datacenter for FC with Cisco MDS.

More details about the FlexPod configurations and performance benefits, see [Introducing End-to-End NVMe for FlexPod White Paper](#).

For more information about ONTAP implementation, see [TR-4684: Implementing and Configuring Modern SANs with NVMe](#).



FC SAN boot through Cisco MDS

To provide increased scalability by using a dedicated SAN network, FlexPod supports FC through Cisco MDS switches and Nexus switches with FC support such as Cisco Nexus 93108TC-FX. The FC SAN boot option through Cisco MDS has the following licensing and hardware requirements:

- A minimum of two FC ports per NetApp storage controller; one port for each SAN fabric
- An FC license on each NetApp storage controller
- Cisco MDS switches and firmware versions that are supported on the NetApp [IMT](#)

For more guidance on an MDS-based design, see the CVD [FlexPod Datacenter with VMware vSphere 6.7U1 Fibre Channel and iSCSI Deployment Guide](#).


The following figures show an example of FlexPod Datacenter for FC with MDS connectivity and FlexPod Datacenter for FC with Cisco Nexus 93180YC-FX, respectively.

Cisco Unified Computing System
 Cisco UCS 6454 Fabric Interconnects,
 UCS B-Series Blade Servers with UCS VIC 1440, and
 UCS C-Series Rack Servers with UCS VIC 1457

Cisco Nexus 9336C-FX2

NetApp storage controllers AFF-A800

Cisco MDS 9148T or 9132T switch




Cisco Unified Computing System

Cisco UCS 6454 Fabric Interconnects, UCS 2408 Fabric Extenders, UCS B-Series Blade Servers with UCS VIC 1440, UCS C-Series Rack Servers with UCS VIC 1457, UCS C4200 Chassis, and UCS C125 Servers with UCS VIC 1455

Cisco Nexus 93180YC-FX

NetApp storage controllers AFF-A400



FC SAN boot with Cisco Nexus

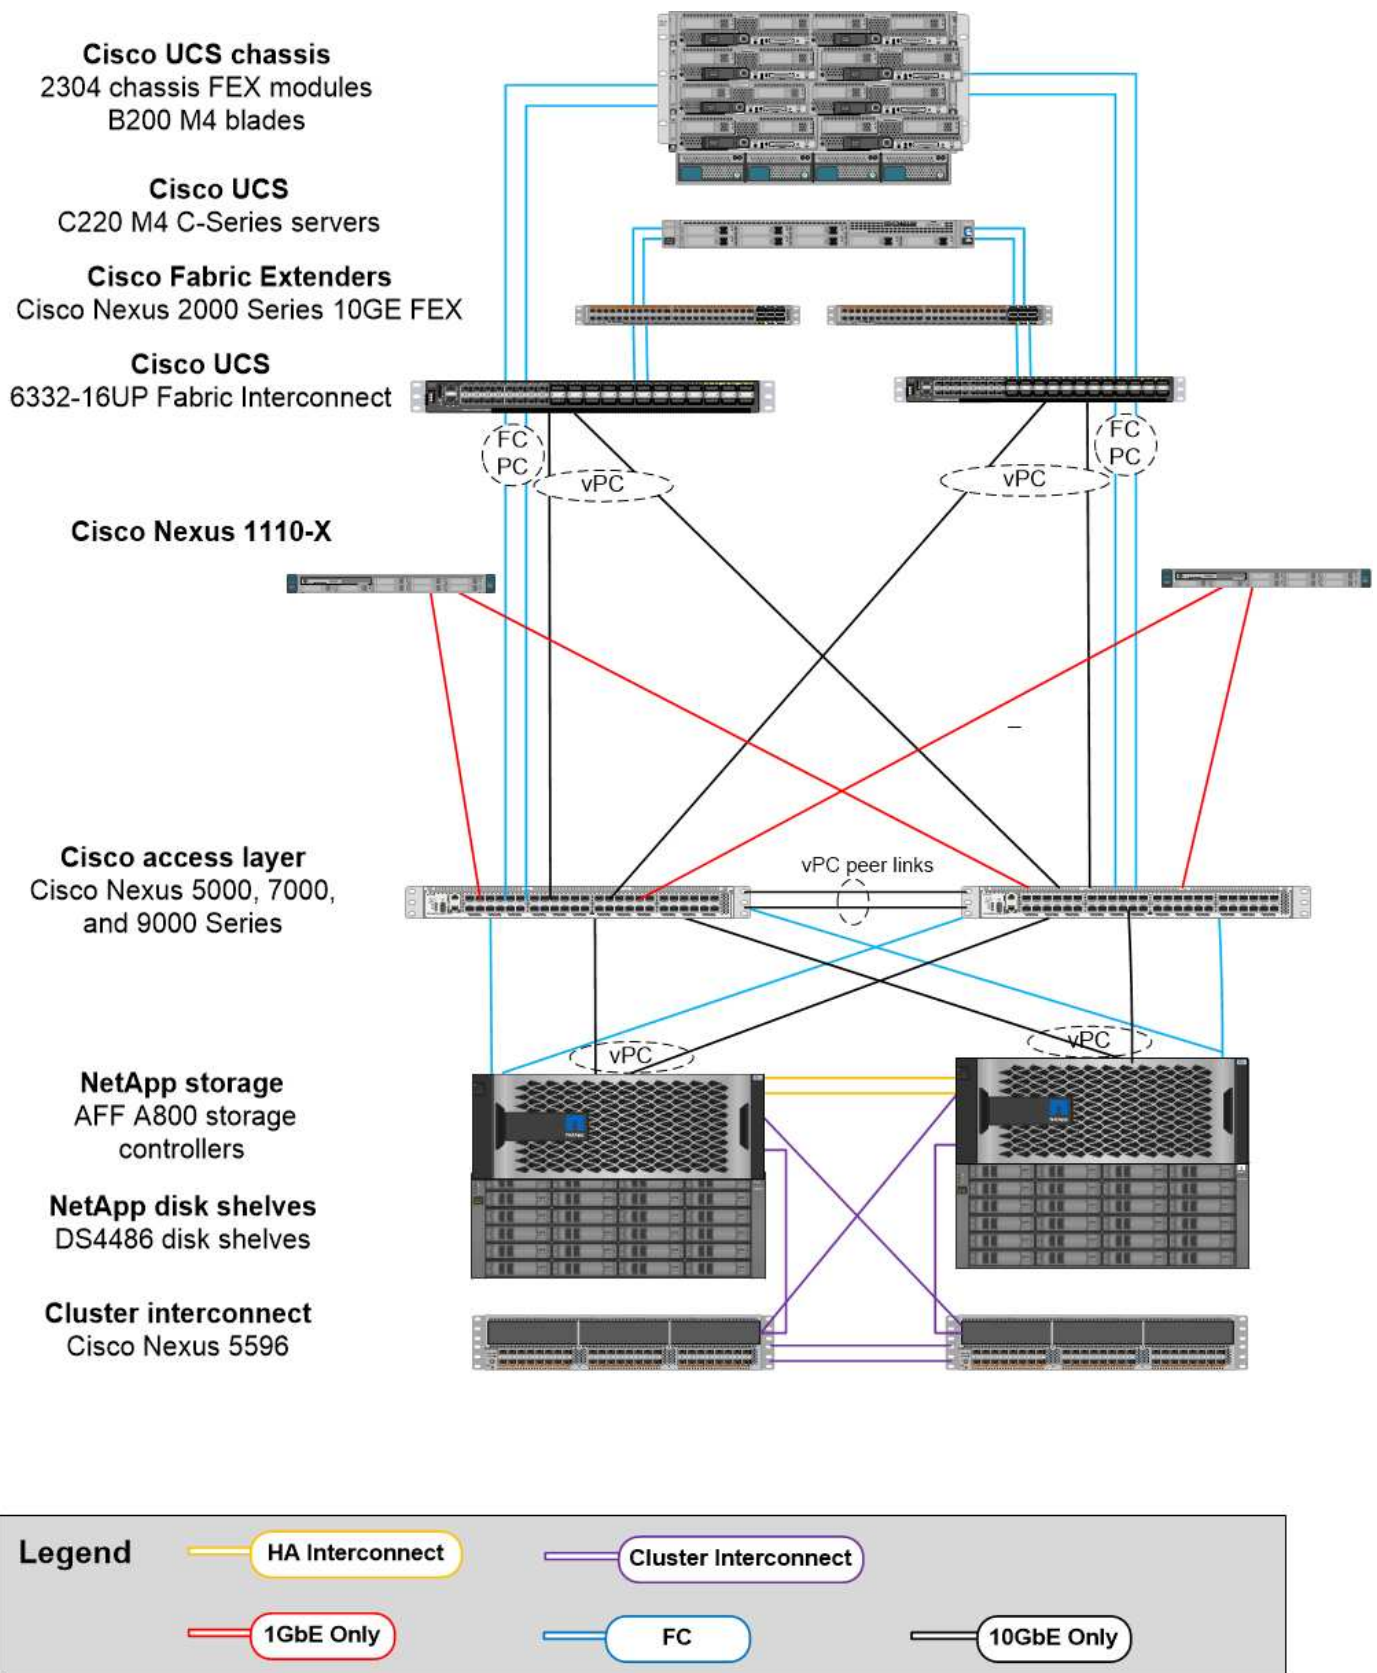
The classic FC SAN boot option has the following licensing and hardware requirements:

- When FC zoning is performed in the Cisco Nexus 5000 Series Switch, a Storage Protocols Service Package license for the Cisco Nexus 5000 Series Switches (FC_FEATURES_PKG) is required.
- When FC zoning is performed in the Cisco Nexus 5000 Series Switch, SAN links are required between the fabric interconnect and the Cisco Nexus 5000 Series Switch. For additional redundancy, SAN port channels are recommended between the links.
- The Cisco Nexus 5010, 5020, and 5548P Switches require a separate FC or universal port (UP) module for connectivity into the Cisco UCS Fabric Interconnect and into the NetApp storage controller.
- The Cisco Nexus 93180YC-FX requires an FC feature license for capabilities to enable FC.
- Each NetApp storage controller requires a minimum of two 8/16/32Gb FC ports for connectivity.
- An FC license on the NetApp storage controller is required.



The use of the Cisco Nexus 7000 or 9000 family of switches precludes the use of traditional FC unless FC zoning is performed in the fabric interconnect. In that case, SAN uplinks to the switch are not supported.

The following figure shows an FC connectivity configuration.



FCoE SAN boot option

The FCoE SAN boot option has the following licensing and hardware requirements:

- When FC zoning is performed in the switch, a Storage Protocols Service Package license for the Cisco

Nexus 5000 or 7000 Series Switches (FC_FEATURES_PKG) is required.

- When FC zoning is performed in the switch, FCoE uplinks are required between the fabric interconnect and the Cisco Nexus 5000 or 7000 Series Switches. For additional redundancy, FCoE port channels are also recommended between the links.
- Each NetApp storage controller requires at least one dual-port unified target adapter (UTA) add-on card for FCoE connectivity unless onboard unified target adapter 2 (UTA2) ports are present.
- This option requires an FC license on the NetApp storage controller.
- If you use the Cisco Nexus 7000 Series Switches and FC zoning is performed in the switch, a line card that is capable of supporting FCoE is required.



The use of the Cisco Nexus 9000 Series Switches precludes the use of FCoE unless FC zoning is performed in the fabric interconnect and storage is connected to the fabric interconnects with appliance ports. In that case, FCoE uplinks to the switch are not supported.

The following figure shows an FCoE boot scenario.

Cisco UCS chassis
2304 chassis FEX modules
B200 M4 blades

Cisco UCS
C220 M4 C-Series servers

Cisco Fabric Extenders
Cisco Nexus 2000 Series 10GE FEX


Cisco UCS
6332-16UP Fabric Interconnect

Cisco access layer
Cisco Nexus 5000, 7000,
and 9000 Series

NetApp storage
AFF A800 storage
controllers

NetApp disk shelves
DS4486 disk shelves

Cluster interconnect
Cisco Nexus 5596



Legend

HA Interconnect

Cluster Interconnect

FCoE Only

FCoE and 10GbE

10GbE Only

iSCSI boot option

The iSCSI boot option has the following licensing and hardware requirements:

- An iSCSI license on the NetApp storage controller is required.
- An adapter in the Cisco UCS Server that is capable of iSCSI boot is required.
- A two-port 10Gbps Ethernet adapter on the NetApp storage controller is required.

The following figure shows an Ethernet-only configuration that is booted by using iSCSI.

Cisco UCS chassis
2304 Chassis FEX modules
B200 M4 blades

Cisco UCS
C220 M4 C-Series servers

Cisco Fabric Extenders
Cisco Nexus 2000 Series 10GE FEX


Cisco UCS
6332-16UP Fabric Interconnect

Cisco access layer
Cisco Nexus 5000, 7000,
and 9000 Series

NetApp storage
AFF A800 storage
controllers

NetApp disk shelves
DS4486 Disk shelves

Cluster Interconnect
Cisco Nexus 5596



Legend

HA Interconnect

10GbE Only

Cluster Interconnect

FCoE

Cisco UCS direct connect with NetApp storage

NetApp AFF and FAS controllers can be directly connected to the Cisco UCS fabric interconnects without any upstream SAN switch.



Four Cisco UCS port types can be used to directly connect to NetApp storage:

- **Storage FC port.** Directly connect this port to an FC port on NetApp storage.
- **Storage FCoE port.** Directly connect this port to an FCoE port on NetApp storage.
- **Appliance port.** Directly connect this port to a 10GbE port on NetApp storage.
- **Unified storage port.** Directly connect this port to a NetApp UTA.

The licensing and hardware requirements are as follows:

- A protocol license on the NetApp storage controller is required.
- A Cisco UCS adapter (initiator) is required on the server. For a list of supported Cisco UCS adapters, see the NetApp [IMT](#).
- A target adapter on the NetApp storage controller is required.


The following figure shows an FC direct-connect configuration.



Notes:

- Cisco UCS is configured in FC switching mode.
- FCoE ports from the target to fabric interconnects are configured as FCoE storage ports.
- FC ports from the target to fabric interconnects are configured as FC storage ports.

The following figure shows an iSCSI/Unified IP direct-connect configuration.



Notes:

- Cisco UCS is configured in Ethernet switching mode.
- iSCSI ports from the target to fabric interconnects are configured as Ethernet storage ports for iSCSI data.
- Ethernet ports from the target to fabric interconnects are configured as Ethernet storage ports for CIFS/NFS data.

Cisco components

Cisco has contributed substantially to the FlexPod design and architecture, covering both the compute and networking layers of the solution. This section describes the Cisco UCS and Cisco Nexus options that are available for FlexPod. FlexPod supports both Cisco UCS B-Series and C-Series servers.

Cisco UCS fabric interconnect options

Redundant fabric interconnects are required in the FlexPod architecture. When you add multiple Cisco UCS chassis to a pair of fabric interconnects, remember that the maximum number of chassis in an environment is determined by both an architectural and a port limit.

The part numbers that are shown in the following table are for the base fabric interconnects. They do not include the power supply unit (PSU) or SFP+, QSFP+, or expansion modules. Additional fabric interconnects are supported; see the [NetApp IMT](#) for a complete list.

| Cisco UCS fabric interconnect | Part number | Technical specifications |
|-------------------------------|----------------|--|
| Cisco UCS 6332UP | UCS-FI-6332-UP | Cisco UCS 6332 Fabric Interconnect |

| Cisco UCS fabric interconnect | Part number | Technical specifications |
|-------------------------------|---------------|--|
| Cisco UCS 6454 | UCS-FI-6454-U | Cisco UCS 6454 Fabric Interconnect |

Cisco UCS 6454

The Cisco UCS 6454 Series offers line-rate, low-latency, lossless 10/25/40/100GbE Ethernet and FCoE connectivity, as well as unified ports that are capable of either Ethernet or FC operation. The 44 10/25Gbps ports can operate as 10Gbps or 25Gbps converged Ethernet, of which eight are unified ports capable of operating at 8/16/32Gbps for FC. Four ports operate at 1/10/25Gbps for legacy connectivity, and six QSFP ports serve as 40/100Gbps uplink ports or breakout ports. You can establish 100Gbps end-to-end network connectivity with NetApp storage controllers that support 100Gbps adapters. For adapters and platform support, see the [NetApp Hardware Universe](#).

For details about ports, see the [Cisco UCS 6454 Fabric Interconnect Datasheet](#).

For technical specifications about the 100Gb QSFP data modules, see the [Cisco 100GBASE QSFP Modules Datasheet](#).

Cisco UCS B-Series chassis option

To use Cisco UCS B-Series blades, you must have a Cisco UCS B-Series chassis. The table below describes the Cisco UCS B-Series chassis option.

| Cisco UCS B-Series chassis | Part number | Technical specifications |
|----------------------------|-------------|--|
| Cisco UCS 5108 | N20-C6508 | Cisco UCS 5100 Series Blade Server Chassis |

Each Cisco UCS 5108 blade chassis must have two Cisco UCS 2200/2300/2400 Series IOMs to provide redundant connectivity to the fabric interconnects.

Cisco UCS B-Series blade server options

Cisco UCS B-Series Blade Servers are available in half-width and full-width varieties, with various CPU, memory, and I/O options. The part numbers that are listed in the following table are for the base server. They do not include the CPU, memory, drives, or mezzanine adapter cards. Multiple configuration options are available and are supported in the FlexPod architecture.

| Cisco UCS B-Series blade | Part number | Technical specifications |
|--------------------------|--------------|--|
| Cisco UCS B200 M6 | UCSB-B200-M6 | Cisco UCS B200 M6 Blade Server |

Previous generations of Cisco UCS B-Series blades can be used in the FlexPod architecture, if they are supported on the [Cisco UCS Hardware and Software Compatibility List](#). The Cisco UCS B-Series Blade Servers must also have a valid SmartNet support contract.

Cisco UCS X-Series chassis option

To use Cisco UCS X-Series compute nodes, you must have a Cisco UCS X-Series chassis. The following table describes the Cisco UCS X-Series chassis option.

| Cisco UCS X-Series blade | Part number | Technical specifications |
|--------------------------|-------------|--|
| Cisco UCS 9508 M6 | UCSX-9508 | Cisco UCX9508 X-Series Chassis |

Each Cisco UCS 9508 chassis must have two Cisco UCS 9108 Intelligent Fabric Modules (IFMs) to provide redundant connectivity to the fabric interconnects.

Cisco UCS X-Series device options

Cisco UCS X-Series compute nodes are available with various CPU, memory, and I/O options. The part numbers listed in the following table are for the base node. They do not include the CPU, memory, drives, or mezzanine adapter cards. Multiple configuration options are available and are supported in the FlexPod architecture.

| Cisco UCS X-Series compute nodes | Part number | Technical specifications |
|----------------------------------|--------------|---|
| Cisco UCS X210c M6 | UCSX-210C-M6 | Cisco UCS X210c M6 Compute Node |

Cisco UCS C-Series rack server options

Cisco UCS C-Series Rack Servers are available in one and two rack-unit (RU) varieties, with various CPU, memory, and I/O options. The part numbers that are listed in the second table below are for the base server. They do not include CPUs, memory, drives, Peripheral Component Interconnect Express (PCIe) cards, or the Cisco Fabric Extender. Multiple configuration options are available and are supported in the FlexPod architecture.

The following table lists the Cisco UCS C-Series Rack Server options.

| Cisco UCS C-Series rack server | Part number | Technical specifications |
|--------------------------------|--------------|---|
| Cisco UCS C220 M6 | UCSC-C220-M6 | Cisco UCS C220 M6 Rack Server |
| Cisco UCS C225 M6 | UCSC-C225-M6 | Cisco UCS C225 M6 Rack Server |
| Cisco UCS C240 M6 | UCSC-C240-M6 | Cisco UCS C240 M6 Rack Server |
| Cisco UCS C245 M6 | UCSC-C245-M6 | Cisco UCS C245 M6 Rack Server |

Previous generations of Cisco UCS C-Series servers can be used in the FlexPod architecture, if they are supported on the [Cisco UCS Hardware and Software Compatibility List](#). The Cisco UCS C-Series servers must also have a valid SmartNet support contract.

Cisco Nexus 5000 Series switch options

Redundant Cisco Nexus 5000, 7000, or 9000 Series Switches are required in the FlexPod architecture. The part numbers that are listed in the table below are for the Cisco Nexus 5000 Series chassis; they do not include SFP modules, add-on FC, or Ethernet modules.

| Cisco Nexus 5000 Series switch | Part number | Technical specifications |
|--------------------------------|-----------------|--|
| Cisco Nexus 56128P | N5K-C56128P | Cisco Nexus 5600 Platform Switches |
| Cisco Nexus 5672UP-16G | N5K-C5672UP-16G | |

| Cisco Nexus 5000 Series switch | Part number | Technical specifications |
|--------------------------------|----------------|--|
| Cisco Nexus 5596UP | N5K-C5596UP-FA | Cisco Nexus 5548 and 5596 Switches |
| Cisco Nexus 5548UP | N5K-C5548UP-FA | |

Cisco Nexus 7000 series switch options

Redundant Cisco Nexus 5000, 7000, or 9000 Series Switches are required in the FlexPod architecture. The part numbers that are listed in the table below are for the Cisco Nexus 7000 Series chassis; they do not include SFP modules, line cards, or power supplies, but they do include fan trays.

| Cisco Nexus 7000 Series Switch | Part number | Technical specifications |
|--------------------------------|-------------|--|
| Cisco Nexus 7004 | N7K-C7004 | Cisco Nexus 7000 4-Slot Switch |
| Cisco Nexus 7009 | N7K-C7009 | Cisco Nexus 7000 9-Slot Switch |
| Cisco Nexus 7702 | N7K-C7702 | Cisco Nexus 7700 2-Slot Switch |
| Cisco Nexus 7706 | N77-C7706 | Cisco Nexus 7700 6-Slot Switch |

Cisco Nexus 9000 series switch options

Redundant Cisco Nexus 5000, 7000, or 9000 Series Switches are required in the FlexPod architecture. The part numbers that are listed in the table below are for the Cisco Nexus 9000 Series chassis; they do not include SFP modules or Ethernet modules.

| Cisco Nexus 9000 Series Switch | Part Number | Technical Specifications |
|--------------------------------|-----------------|--|
| Cisco Nexus 93180YC-FX | N9K-C93180YC-FX | Cisco Nexus 9300 Series Switches |
| Cisco Nexus 93180YC-EX | N9K-93180YC-EX | |
| Cisco Nexus 9336PQ ACI Spine | N9K-C9336PQ | |
| Cisco Nexus 9332PQ | N9K-C9332PQ | |
| Cisco Nexus 9336C-FX2 | N9K-C9336C-FX2 | |
| Cisco Nexus 92304QC | N9K-C92304QC | Cisco Nexus 9200 Series Switches |
| Cisco Nexus 9236C | N9K-9236C | |



Some Cisco Nexus 9000 Series Switches have additional variants. These variants are supported as part of the FlexPod solution. For the complete list of Cisco Nexus 9000 Series Switches, see [Cisco Nexus 9000 Series Switches](#) on the Cisco website.

Cisco APIC options

When deploying Cisco ACI, you must configure the three Cisco APICs in addition to the items in the section [Cisco Nexus 9000 Series Switches](#). For more information about the Cisco APIC sizes, see the [Cisco Application Centric Infrastructure Datasheet](#).

For more information about APIC product specifications refer to Table 1 through Table 3 on the [Cisco Application Policy Infrastructure Controller Datasheet](#).

Cisco Nexus fabric extender options

Redundant Cisco Nexus 2000 Series rack-mount FEXs are recommended for large FlexPod architectures that use C-Series servers. The table below describes a few Cisco Nexus FEX options. Alternate FEX models are also supported. For more information, see the [Cisco UCS Hardware and Software Compatibility List](#).

| Cisco Nexus rack-mount FEX | Part number | Technical specifications |
|--|--------------------------|--|
| Cisco Nexus 2232PP | N2K-C2232PP | Cisco Nexus 2000 Series Fabric Extenders |
| Cisco Nexus 2232TM-E | N2K-C2232TM-E | |
| Cisco Nexus 2348UPQ | N2K-C2348UPQ | Cisco Nexus 2300 Platform Fabric Extenders |
| Cisco Nexus 2348TQCisco Nexus 2348TQ-E | N2K-C2348TQN2K-C2348TQ-E | |

Cisco MDS options

Cisco MDS switches are an optional component in the FlexPod architecture. Redundant SAN switch fabrics are required when you implement the Cisco MDS switch for FC SAN. The table below lists the part numbers and details for a subset of the supported Cisco MDS switches. See the [NetApp IMT](#) and [Cisco Hardware and Software Compatibility List](#) for a complete list of supported SAN switches.

| Cisco MDS 9000 series switch | Part number | Description |
|------------------------------|----------------|--|
| Cisco MDS 9148T | DS-C9148T-24IK | Cisco MDS 9100 Series Switches |
| Cisco MDS 9132T | DS-C9132T-MEK9 | |
| Cisco MDS 9396S | DS-C9396S-K9 | Cisco MDS 9300 Series Switches |

Cisco software licensing options

Licenses are required to enable storage protocols on the Cisco Nexus switches. The Cisco Nexus 5000 and 7000 Series of switches all require a storage services license to enable the FC or FCoE protocol for SAN boot implementations. The Cisco Nexus 9000 Series Switches currently do not support FC or FCoE.

The required licenses and the part numbers for those licenses vary depending on the options that you select for each component of the FlexPod solution. For example, software license part numbers vary depending on the number of ports and which Cisco Nexus 5000 or 7000 Series Switches you choose. Consult your sales representative for the exact part numbers. The table below lists the Cisco software licensing options.

| Cisco software licensing | Part number | License information |
|--|--|---|
| Cisco Nexus 5500 Storage License, 8-, 48-, and 96-port | N55-8P-SSK9/N55-48P-SSK9/N55-96P-SSK9 | Licensing Cisco NX-OS Software Features |
| Cisco Nexus 5010/5020 Storage Protocols License | N5010-SSK9/N5020-SSK9 | |
| Cisco Nexus 5600 Storage Protocols License | N56-16p-SSK9/N5672-72P-SSK9/N56128-128P-SSK9 | |
| Cisco Nexus 7000 Storage Enterprise License | N7K-SAN1K9 | |
| Cisco Nexus 9000 Enterprise Services License | N95-LAN1K9/N93-LAN1K9 | |

Cisco support licensing options

Valid SmartNet support contracts are required on all Cisco equipment in the FlexPod architecture.

The required licenses and the part numbers for those licenses must be verified by your sales representative because they can vary for different products. The table below lists the Cisco support licensing options.

| Cisco Support licensing | License guide |
|-------------------------------------|--|
| Smart Net Total Care Onsite Premium | Cisco Smart Net Total Care Service |

NetApp components

NetApp storage controllers provide the storage foundation in the FlexPod architecture for both boot and application data storage. NetApp components include storage controllers, cluster interconnect switches, drives and disk shelves, and licensing options.

NetApp storage controller options

Redundant NetApp FAS, AFF, or AFF ASA controllers are required in the FlexPod architecture. The controllers run ONTAP software. When the storage controllers are ordered, the preferred software version can be preloaded on the controllers. For ONTAP, a complete cluster is ordered. A complete cluster includes a pair of storage controllers and a cluster interconnect (switch or switchless).

Different options and configurations are available, depending on the selected storage platform. Consult your sales representative for details about these additional components.

The controller families that are listed in the table below are appropriate for use in a FlexPod Datacenter solution because their connection to the Cisco Nexus switches is seamless. See the [NetApp Hardware Universe](#) for specific compatibility details on each controller model.

| Storage controller family | Technical specifications |
|---------------------------|---|
| AFF A-Series | AFF A-Series Documentation |
| AFF ASAA-Series | AFF ASAA-Series Documentation |
| FAS Series | FAS Series Documentation |

Cluster interconnect switch options

The following table lists the Nexus cluster interconnect switches that are available for FlexPod architectures. In addition, FlexPod supports all ONTAP supported cluster switches including non-Cisco switches, provided they are compatible with the version of ONTAP being deployed. See the [NetApp Hardware Universe](#) for additional compatibility details for specific switch models.

| Cluster interconnect switch | Technical specifications |
|-----------------------------|--|
| Cisco Nexus 3132Q-V | NetApp Documentation: Cisco Nexus 3132Q-V switches |
| Cisco Nexus 9336C-FX2 | NetApp Documentation: Cisco Nexus 9336C-FX2 switches |

NetApp disk shelf and drive options

A minimum of one NetApp disk shelf is required for all storage controllers.

The selected NetApp shelf type determines which drive types are available within that shelf.



For all disk shelves and disk part numbers, consult your sales representative.

For more information about the supported drives, click the NetApp Hardware Universe link in the following table and then select Supported Drives.

| Disk shelf | Technical specifications |
|------------|---|
| DS224C | Disk Shelves and Storage Media Supported Drives on NetApp Hardware Universe |
| DS212C | |
| DS460C | |
| NS224 | |

NetApp software licensing options

The following table lists the NetApp software licensing options that are available for the FlexPod Datacenter architecture. NetApp software is licensed at the FAS and AFF controller level.

| NetApp software licensing | Part number | Technical specifications |
|---------------------------------------|---------------------|-------------------------------------|
| SW, Complete BNDL (Controller), -C | SW-8XXX-COMP-BNDL-C | Product Library A–Z |
| SW, ONTAP Essentials (Controller), -C | SW-8XXX-ONTAP9-C | |

NetApp support licensing options

NetApp SupportEdge Premium licenses are required for the FlexPod architecture, but the part numbers for those licenses vary based on the options that you select in the FlexPod design. For example, software license part numbers are different depending on which FAS controller you choose. Consult your sales representative for information about the exact part numbers for individual support licenses. The table below shows an example of a SupportEdge license.

| NetApp support licensing | Part number | Technical specifications |
|--|-------------|--|
| SupportEdge Premium 4 hours on site—months: 36 | CS-O2-4HR | NetApp SupportEdge Premium |

Power and cabling requirements

A FlexPod design has minimum requirements for power and cabling.

Power requirements

Power requirements for FlexPod Datacenter differ based on the location where the FlexPod Datacenter configuration is installed.

For more data about the maximum power that is required and for other detailed power information, consult the technical specifications for each hardware component listed in the section [Technical Specifications and References: Hardware Components](#).

For detailed Cisco UCS power data, see the [Cisco UCS power calculator](#).

For NetApp storage controller power data, see the [NetApp Hardware Universe](#). Under Platforms, select the storage platform that you want to use in the configuration (FAS/V-Series or AFF). Select the ONTAP version and storage controller, and then click the Show Results button.

Minimum cable requirements

The number and type of cables and adapters that are required vary per FlexPod Datacenter deployment. The cable type, transceiver type, and number are determined during the design process based on your requirements. The table below lists the minimum number of cables required.

| Hardware | Model number | Cables required |
|-------------------|----------------|--|
| Cisco UCS chassis | Cisco UCS 5108 | At least two twinaxial cables per Cisco UCS 2104XP, 2204XP, or 2208XP module |

| Hardware | Model number | Cables required |
|---|-------------------------|--|
| Cisco UCS Fabric Interconnects | Cisco UCS 6248UP | <ul style="list-style-type: none"> Two Cat5e cables for management ports |
| | Cisco UCS 6296UP | <ul style="list-style-type: none"> Two Cat5e cables for the L1, L2 interconnects, per pair of fabric interconnects |
| | Cisco UCS 6332-16UP | <ul style="list-style-type: none"> At least four twinaxial cables per fabric interconnect |
| | Cisco UCS 6454 | <ul style="list-style-type: none"> At least four FC cables per fabric interconnect |
| | Cisco UCS 6332 | <ul style="list-style-type: none"> Two Cat5e cables for management ports Two Cat5e cables for the L1, L2 interconnects, per pair of fabric interconnects At least four twinaxial cables per fabric interconnect |
| | Cisco UCS 6324 | <ul style="list-style-type: none"> Two 10/100/1000Mbps management ports At least two twinaxial cables per fabric interconnect |
| Cisco Nexus 5000 and 7000 Series Switches | Cisco Nexus 5000 Series | <ul style="list-style-type: none"> At least two 10GbE fiber or twinaxial cables per switch At least two FC cables per switch (if FC/FCoE connectivity is required) |
| | Cisco Nexus 7000 Series | |
| Cisco Nexus 9000 Series Switches | Cisco Nexus 9000 Series | At least two 10GbE cables per switch |
| NetApp FAS controllers | AFF A-Series | <ul style="list-style-type: none"> A pair of SAS or SATA cables per storage controller At least two FC cables per controller, if using legacy FC At least two 10GbE cables per controller |
| | FAS Series | <ul style="list-style-type: none"> At least one GbE cable for management per controller For ONTAP, eight short twinaxial cables are required per pair of cluster interconnect switches |

| Hardware | Model number | Cables required |
|---------------------|--------------|--|
| NetApp disk shelves | DS212C | Two SAS, SATA, or FC cables per disk shelf |
| | DS224C | |
| | DS460C | |
| | NS224 | Two 100Gbps copper cables per disk shelf |

Technical specifications and references

Technical specifications provide details about the hardware components in a FlexPod solution, such as chassis, FEXs, servers, switches, and storage controllers.

Cisco UCS B-Series blade server chassis

The technical specifications for Cisco UCS B-Series Blade Server chassis, as shown in the table below, include the following components:

- Number of rack units
- Maximum number of blades
- Unified Fabric capability
- Midplane I/O bandwidth per server
- Number of I/O bays for FEXs

| Component | Cisco UCS 5100 Series blade server chassis |
|---------------------------|--|
| Rack units | 6 |
| Maximum full-width blades | 4 |
| Maximum half-width blades | 8 |
| Capable of Unified Fabric | Yes |
| Midplane I/O | Up to 80Gbps of I/O bandwidth per server |
| I/O bays for FEXs | Two bays for Cisco UCS 2104XP, 2204/8XP, 2408XP, and 2304 FEXs |

For more information, see the [Cisco UCS 5100 Series Blade Server Chassis Datasheet](#).

Cisco UCS B-Series blade servers

The technical specifications for Cisco UCS B-Series Blade Servers, as shown in the table below, include the following components:

- Number of processor sockets
- Processor support
- Memory capacity
- Size and speed

- SAN boot support
- Number of mezzanine adapter slots
- I/O maximum throughput
- Form factor
- Maximum number of servers per chassis

| Component | Cisco UCS datasheet |
|-------------------|--|
| Cisco UCS B200 M6 | Cisco UCS B200 M6 Blade Server |

Cisco UCS C-Series rack servers

The technical specifications for the Cisco UCS C-Series rack servers include processor support, maximum memory capacity, the number of PCIe slots, and the size of the form factor. For additional details on compatible UCS server models, see the [Cisco Hardware Compatibility](#) list. The following tables illustrate the C-Series Rack Server datasheets and Cisco UCS C-Series chassis option, respectively.

| Component | Cisco UCS datasheet |
|-------------------|---|
| Cisco UCS C220 M6 | Cisco UCS C220 M6 Rack Server |
| Cisco UCS C225 M6 | Cisco UCS C225 M6 Rack Server |
| Cisco UCS C240 M6 | Cisco UCS C240 M6 Rack Server |
| Cisco UCS C245 M6 | Cisco UCS C245 M6 Rack Server |

Cisco UCS X-Series chassis

The technical specifications for Cisco UCS X-Series chassis, as shown in the table below, include the following components:

- Number of rack units
- Maximum number of nodes
- Unified Fabric capability
- Number of I/O bays for IFMs

| Component | Cisco UCS 9508 X-Series compute node chassis |
|---------------------------|---|
| Rack units | 7 |
| Maximum number of nodes | 8 |
| Capable of Unified Fabric | Yes |
| I/O bays for IFMs | Two bays for Cisco UCS 9108 Intelligent Fabric Modules (IFMs) |

For more information, see the [Cisco UCS X9508 X-Series Chassis Datasheet](#).

Cisco UCS X-Series compute node

The technical specifications for Cisco UCS X-Series compute node, as shown in the following table below,

include the following components:

- Number of processor sockets
- Processor support
- Memory capacity
- Size and speed
- SAN boot support
- Number of mezzanine adapter slots
- I/O maximum throughput
- Form factor
- Maximum number of compute nodes per chassis

| Component | Cisco UCS datasheet |
|--------------------|---|
| Cisco UCS X210c M6 | Cisco UCS X210c M6 Compute Node |

GPU recommendation for FlexPod AI, ML, and DL

The Cisco UCS C-Series Rack Servers listed in the table below can be used in a FlexPod architecture for hosting AI, ML, and DL workloads. The Cisco UCS C480 ML M5 Servers are purpose built for AI, ML, and DL workloads and use NVIDIA's SXM2- based GPUs while the other servers use PCIe- based GPUs.

The table below also lists the recommended GPUs that can be used with these servers.

| Server | GPUs |
|-------------------|------------------------|
| Cisco UCS C220 M6 | NVIDIA T4 |
| Cisco UCS C225 M6 | NVIDIA T4 |
| Cisco UCS C240 M6 | NVIDIA TESLA A10, A100 |
| Cisco UCS C245 M6 | NVIDIA TESLA A10, A100 |

Cisco UCS VIC adapters for Cisco UCS B-Series blade servers

The technical specifications for Cisco UCS Virtual Interface Card (VIC) adapters for Cisco UCS B-Series Blade Servers include the following components:

- Number of uplink ports
- Performance per port (IOPS)
- Power
- Number of blade ports
- Hardware offload
- Single root input/output virtualization (SR-IOV) support

All currently validated FlexPod architectures use a Cisco UCS VIC. Other adapters are supported if they are listed on the NetApp [IMT](#) and are compatible with your deployment of FlexPod, but they might not deliver all the features that are outlined in corresponding reference architectures. The following table illustrates the Cisco UCS VIC adapter datasheets.

| Component | Cisco UCS datasheet |
|--------------------------------------|--|
| Cisco UCS Virtual Interface Adapters | Cisco UCS VIC Datasheets |

Cisco UCS fabric interconnects

The technical specifications for Cisco UCS fabric interconnects include form factor size, the total number of ports and expansion slots, and throughput capacity. The following table illustrates the Cisco UCS fabric interconnect datasheets.

| Component | Cisco UCS datasheet |
|------------------|--|
| Cisco UCS 6248UP | Cisco UCS 6200 Series Fabric Interconnects |
| Cisco UCS 6296UP | |
| Cisco UCS 6324 | Cisco UCS 6324 Fabric Interconnect |
| Cisco UCS 6300 | Cisco UCS 6300 Series Fabric Interconnects |
| Cisco UCS 6454 | Cisco UCS 6400 Series Fabric Interconnects |

Cisco Nexus 5000 Series switches

The technical specifications for Cisco Nexus 5000 Series Switches, including the form factor size, the total number of ports, and layer- 3 module and daughter card support, are contained in the datasheet for each model family. These datasheets can be found in the following table.

| Component | Cisco Nexus datasheet |
|-------------------------|---|
| Cisco Nexus 5548UP | Cisco Nexus 5548UP Switch |
| Cisco Nexus 5596UP (2U) | Cisco Nexus 5596UP Switch |
| Cisco Nexus 56128P | Cisco Nexus 56128P Switch |
| Cisco Nexus 5672UP | Cisco Nexus 5672UP Switch |

Cisco Nexus 7000 Series switches

The technical specifications for Cisco Nexus 7000 Series Switches, including the form factor size and the maximum number of ports, are contained in the datasheet for each model family. These datasheets can be found in the following table.

| Component | Cisco Nexus datasheet |
|------------------|--|
| Cisco Nexus 7004 | Cisco Nexus 7000 Series Switches |
| Cisco Nexus 7009 | |
| Cisco Nexus 7010 | |
| Cisco Nexus 7018 | |

| Component | Cisco Nexus datasheet |
|------------------|--|
| Cisco Nexus 7702 | Cisco Nexus 7700 Series Switches |
| Cisco Nexus 7706 | |
| Cisco Nexus 7710 | |
| Cisco Nexus 7718 | |

Cisco Nexus 9000 Series switches

The technical specifications for Cisco Nexus 9000 Series Switches are contained in the datasheet for each model. Specifications include the form factor size; the number of supervisors, fabric module, and line card slots; and the maximum number of ports. These datasheets can be found in the following table.

| Component | Cisco Nexus datasheet |
|-------------------------------------|---|
| Cisco Nexus 9000 Series | Cisco Nexus 9000 Series Switches |
| Cisco Nexus 9500 Series | Cisco Nexus 9500 Series Switches |
| Cisco Nexus 9300 Series | Cisco Nexus 9300 Series Switches |
| Cisco Nexus 9336PQ ACI Spine Switch | Cisco Nexus 9336PQ ACI Spine Switch |
| Cisco Nexus 9200 Series | Cisco Nexus 9200 Platform Switches |

Cisco Application Policy Infrastructure controller

When you deploy Cisco ACI, in addition to the items in the section [Cisco Nexus 9000 Series Switches](#), you must configure three Cisco APICs. The following table lists the Cisco APIC datasheet.

| Component | Cisco Application Policy Infrastructure datasheet |
|--|---|
| Cisco Application Policy Infrastructure Controller | Cisco APIC Datasheet |

Cisco Nexus fabric extender details

The technical specifications for the Cisco Nexus FEX include speed, the number of fixed ports and links, and form factor size.

The following table lists the Cisco Nexus 2000 Series FEX datasheet.

| Component | Cisco Nexus fabric extender datasheet |
|--|---|
| Cisco Nexus 2000 Series Fabric Extenders | Nexus 2000 Series FEX Datasheet |

SFP modules

For information about the SFP modules, review the following resources:

- For information about the Cisco 10Gb SFP, see [Cisco 10 Gigabit Modules](#).
- For information about the Cisco 25Gb SFP, see [Cisco 25 Gigabit Modules](#).
- For information about the Cisco QSFP module, see the [Cisco 40GBASE QSFP Modules datasheet](#).

- For information about the Cisco 100Gb SFP, see [Cisco 100 Gigabit Modules](#).
- For information about the Cisco FC SFP module, see the [Cisco MDS 9000 Family Pluggable Transceivers datasheet](#).
- For information about all supported Cisco SFP and transceiver modules, see [Cisco SFP and SFP+ Transceiver Module Installation Notes](#) and [Cisco Transceiver Modules](#).

NetApp storage controllers

The technical specifications for NetApp storage controllers include the following components:

- Chassis configuration
- Number of rack units
- Amount of memory
- NetApp FlashCache caching
- Aggregate size
- Volume size
- Number of LUNs
- Supported network storage
- Maximum number of NetApp FlexVol volumes
- Maximum number of supported SAN hosts
- Maximum number of Snapshot copies

FAS Series

All available models of FAS storage controllers are supported for use in a FlexPod Datacenter. Detailed specifications for all FAS series storage controllers are available in the [NetApp Hardware Universe](#). See the platform-specific documentation listed in the following table for detailed information about a specific FAS model.

| Component | FAS Series controller platform documentation |
|----------------|--|
| FAS9000 Series | FAS9000 Series Datasheet |
| FAS8700 Series | FAS8700 Series Datasheet |
| FAS8300 Series | FAS8300 Series Datasheet |
| FAS500f Series | FAS500f Series Datasheet |
| FAS2700 Series | FAS2700 Series Datasheet |

AFF A-Series

All current models of NetApp AFF A-Series storage controllers are supported for use in FlexPod. Additional information can be found in the [AFF Technical Specifications](#) datasheet and in the [NetApp Hardware Universe](#). See the platform-specific documentation listed in the following table for detailed information about a specific AFF Model.

| Component | AFF A-Series controller platform documentation |
|------------------|--|
| NetApp AFF A800 | AFF A800 Platform Documentation |
| NetApp AFF A700 | AFF A700 Platform Documentation |
| NetApp AFF A700s | AFF A700s Platform Documentation |
| NetApp AFF A400 | AFF A400 Platform Documentation |
| NetApp AFF A250 | AFF A250 Platform Documentation |

AFF ASA A-Series

All current models of NetApp AFF ASA A-Series storage controllers are supported for use in FlexPod. Additional information can be found in the All SAN Array documentation resources, ONTAP AFF All SAN Array System technical report, and in the NetApp Hardware Universe. See the platform-specific documentation listed in the following table for detailed information about a specific AFF Model.

| Component | AFF A-Series controller platform documentation |
|---------------------|---|
| NetApp AFF ASA A800 | AFF ASA A800 Platform Documentation |
| NetApp AFF ASA A700 | AFF ASA A700 Platform Documentation |
| NetApp AFF ASA A400 | AFF ASA A400 Platform Documentation |
| NetApp AFF ASA A250 | AFF ASA A250 Platform Documentation |
| NetApp AFF ASA A220 | AFF ASA A220 Platform Documentation |

NetApp disk shelves

The technical specifications for NetApp disk shelves include the form factor size, the number of drives per enclosure, and the shelf I/O modules; this documentation can be found in the following table. For more information, see the [NetApp Disk Shelves and Storage Media Technical Specifications](#) and the [NetApp Hardware Universe](#).

| Component | NetApp FAS/AFF disk shelf documentation |
|----------------------------------|---|
| NetApp DS212C Disk Shelf | DS212C Disk Shelf Documentation |
| NetApp DS224C Disk Shelf | DS224C Disk Shelf Documentation |
| NetApp DS460C Disk Shelf | DS460C Disk Shelf Documentation |
| NetApp NS224 NVMe-SSD Disk Shelf | NS224 Disk Shelf Documentation |

NetApp drives

The technical specifications for NetApp drives include the form factor size, disk capacity, disk RPM, supporting controllers, and ONTAP version requirements. These specifications can be found in the Drives section of the [NetApp Hardware Universe](#).

Legacy equipment

FlexPod is a flexible solution that enables you to use your existing equipment and new equipment that is currently for sale by Cisco and NetApp. Occasionally, certain models of

equipment from both Cisco and NetApp are designated as end of life (EOL).

Even though these equipment models are no longer available, if you purchased one of these models before the end-of-availability (EOA) date, you can use that equipment in a FlexPod configuration. A complete list of the legacy equipment models that are supported in FlexPod that are no longer for sale can be referenced on the [NetApp Service and Support Product Programs End of Availability Index](#).

For more information on legacy Cisco equipment, see the Cisco EOL and EOA notices for [Cisco UCS C-Series Rack Servers](#), [Cisco UCS B-Series Blade Servers](#), and [Nexus switches](#).

Legacy FC Fabric support includes the following:

- 2Gb Fabric
- 4Gb Fabric

Legacy software includes the following:

- NetApp Data ONTAP operating in 7-Mode, 7.3.5 and later
- ONTAP 8.1.x through 9.0.x
- Cisco UCS Manager 1.3 and later
- Cisco UCS Manager 2.1 through 2.2.7

Where to find additional Information

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

- NetApp Product Documentation
<https://docs.netapp.com/>
- NetApp Support Communications
<https://mysupport.netapp.com/info/communications/index.html>
- NetApp Interoperability Matrix Tool (IMT)
<https://mysupport.netapp.com/matrix/#welcome>
- NetApp Hardware Universe
<https://hwu.netapp.com/>
- NetApp Support
<https://mysupport.netapp.com/>

FlexPod Datacenter

FlexPod DataCenter with NetApp SnapMirror Business Continuity and ONTAP 9.10

TR-4920: FlexPod Datacenter with NetApp SnapMirror Business Continuity and ONTAP 9.10

Jyh-shing Chen, NetApp

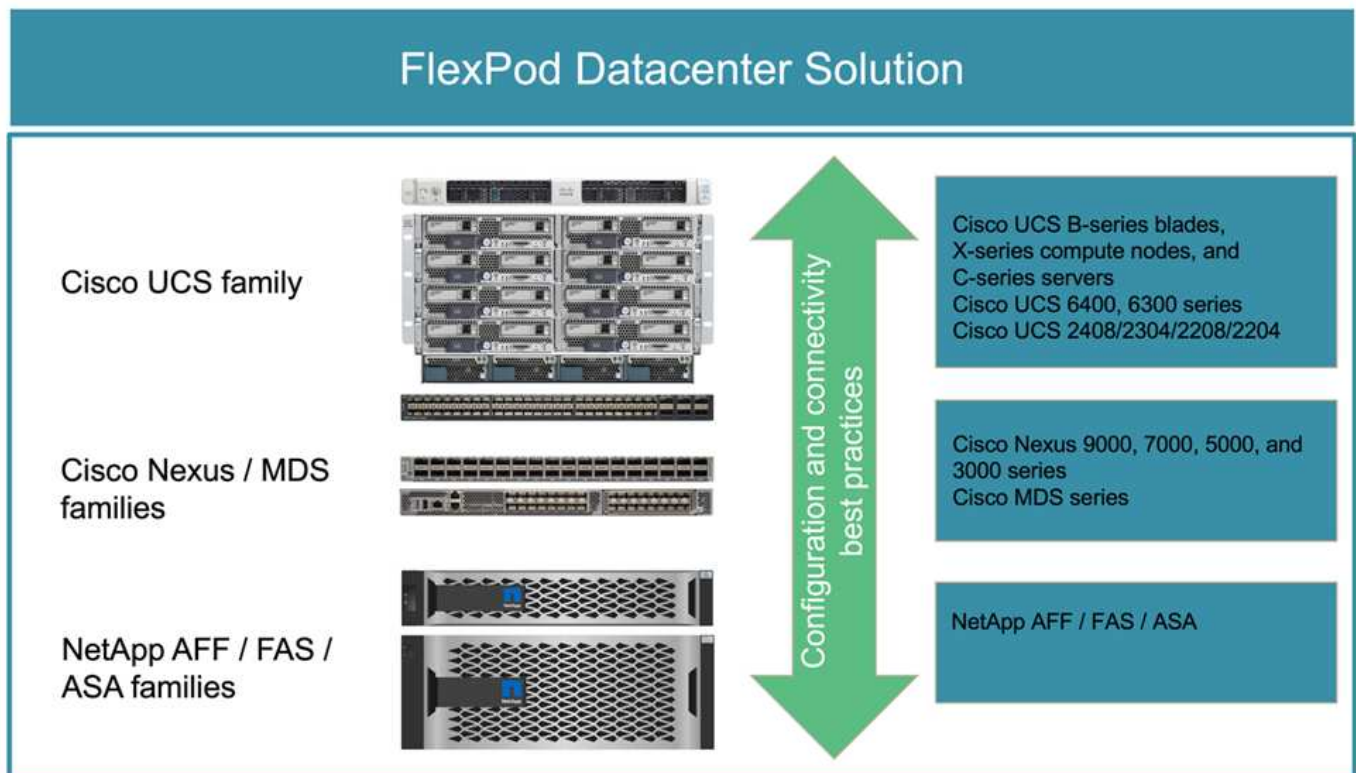
Introduction

FlexPod solution

FlexPod is a best-practice converged-infrastructure data center architecture that includes the following components from Cisco and NetApp:

- Cisco Unified Computing System (Cisco UCS)
- Cisco Nexus and MDS families of switches
- NetApp FAS, NetApp AFF, and NetApp All SAN Array (ASA) systems

The following figure depicts some of the components used for creating FlexPod solutions. These components are connected and configured according to the best practices of both Cisco and NetApp to provide an ideal platform for running a variety of enterprise workloads with confidence.



A large portfolio of Cisco Validated Designs (CVDs) and NetApp Verified Architectures (NVAs) are available. These CVDs and NVAs cover all major data center workloads and are the result of continued collaborations

and innovations between NetApp and Cisco on FlexPod solutions.

Incorporating extensive testing and validations in their creation process, FlexPod CVDs and NVAs provide reference solution architecture designs and step-by-step deployment guides to help partners and customers deploy and adopt FlexPod solutions. By using these CVDs and NVAs as guides for design and implementation, businesses can reduce risks; reduce solution downtime; and increase the availability, scalability, flexibility, and security of the FlexPod solutions they deploy.

Each of the FlexPod component families shown (Cisco UCS, Cisco Nexus/MDS switches, and NetApp storage) offers platform and resource options to scale the infrastructure up or down, while supporting the features and functionality that are required under the configuration and connectivity best practices of FlexPod. FlexPod can also scale out for environments that require multiple consistent deployments by rolling out additional FlexPod stacks.

Disaster recovery and business continuity

There are various methods that companies can adopt to make sure that they can quickly recover their application and data services from disasters. Having a disaster recovery (DR) and business continuity (BC) plan, implementing a solution which meets the business objectives, and performing regular testing of the disaster scenarios enables companies to recover from a disaster and continue critical business services after a disaster situation occurs.

Companies might have different DR and BC requirements for different types of application and data services. Some applications and data might not be needed during an emergency or disaster situation, while others might need to be continuously available to support business requirements.

For mission- critical application and data services that could disrupt your business when they are not available, a careful evaluation is needed to answer questions such as what kind of maintenance and disaster scenarios the business needs to consider, how much data the business can afford to lose in case of a disaster, and how quickly the recovery can and should take place.

For businesses that rely on data services for revenue generation, the data services might need to be protected by a solution that can withstand not only various single-point-of-failure scenarios but also a site outage disaster scenario to provide continuous business operations.

Recovery point objective and recovery time objective

The recovery point objective (RPO) measures how much data, in terms of time, you can afford to lose, or the point up to which you can recover your data. With a daily backup plan, a company might lose a day's worth of data because the changes made to the data since the last backup could be lost in a disaster. For business-critical and mission-critical data services, you might require a zero RPO and an associated plan and infrastructures to protect data without any data loss.

The recovery time objective (RTO) measures how much time you can afford to not have the data available, or how quickly data services must be brought back up. For example, a company might have a backup and recovery implementation that uses traditional tapes for certain data sets due to its size. As a result, to restore the data from the backup tapes, it might take several hours, or even days if there is an infrastructure failure. Time considerations must also include time to bring the infrastructure back up in addition to restoring data. For mission-critical data services, you might require a very low RTO and thus can only tolerate a failover time of seconds or minutes to quickly bring the data services back online for business continuity.

SM-BC

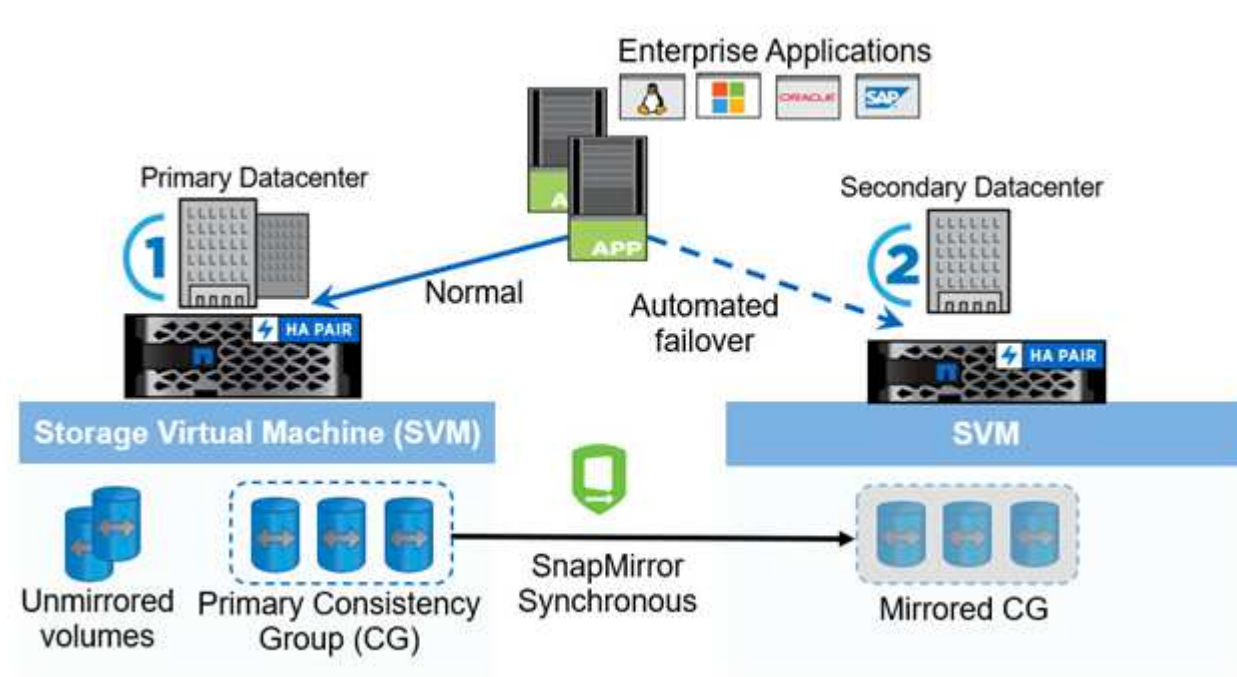
Beginning with ONTAP 9.8, you can protect SAN workloads for transparent application failover with NetApp SM-BC. You can create consistency group relationships between two AFF clusters or two ASA clusters for data

replication to achieve zero RPO and near zero RTO.

The SM-BC solution replicates data by using the SnapMirror Synchronous technology over an IP network. It provides application-level granularity and automatic failover to protect your business-critical data services such as Microsoft SQL Server, Oracle, and so on with iSCSI or FC protocol-based SAN LUNs. An ONTAP Mediator deployed at a third site monitors the SM-BC solution and enables automatic failover upon a site disaster.

A consistency group (CG) is a collection of FlexVol volumes that provides a write order consistency guarantee for the application workload which needs to be protected for business continuity. It enables simultaneous crash-consistent Snapshot copies of a collection of volumes at a point in time. A SnapMirror relationship, also known as a CG relationship, is established between a source CG and a destination CG. The group of volumes picked to be part of a CG can be mapped to an application instance, a group of applications instances, or for an entire solution. In addition, the SM-BC consistency group relationships can be created or deleted on demand based on business requirements and changes.

As illustrated in the following figure, the data in the consistency group is replicated to a second ONTAP cluster for disaster recovery and business continuity. The applications have connectivity to the LUNs in both ONTAP clusters. I/O is normally served by the primary cluster and automatically resumes from the secondary cluster if a disaster happens at the primary. When designing a SM-BC solution, the supported object counts for the CG relationships (for example, a maximum of 20 CGs and maximum of 200 endpoints) must be observed to avoid exceeding the supported limits.



[Next: FlexPod SM-BC solution.](#)

FlexPod SM-BC solution

[Previous: Introduction.](#)

Solution overview

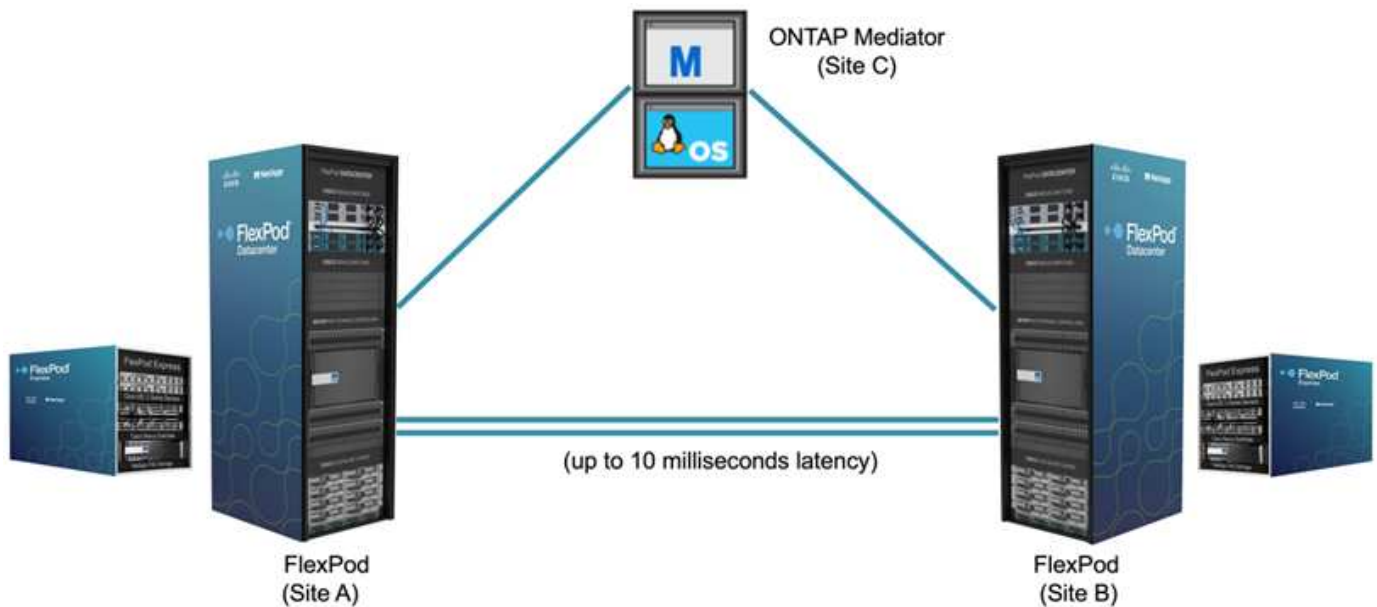
At a high level, a FlexPod SM-BC solution consists of two FlexPod systems, located at two sites separated by some distance, connected, and paired together to provide a highly available, highly flexible, and highly reliable data center solution that can provide business continuity despite a site failure.

In addition to deploying two new FlexPod infrastructures to create a FlexPod SM-BC solution, the solution can also be implemented on two existing FlexPod infrastructures that are compatible with SM-BC or by adding a new FlexPod to peer with an existing FlexPod.

The two FlexPod systems in a FlexPod SM-BC solution do not need to be identical in configurations. However, the two ONTAP clusters need to be of the same storage families, either two AFF or two ASA systems, but not necessarily the same hardware model. The SM-BC solution does not support FAS systems.

The two FlexPod sites require network connectivity which meets the solution bandwidth and quality-of-service requirements and has less than 10 milliseconds (10ms) round-trip latency between sites as required by the ONTAP SM-BC solution. For this FlexPod SM-BC solution validation, the two FlexPod sites are interconnected via extended layer-2 network in the same lab.

The NetApp ONTAP SM-BC solution provides synchronous replication between the two NetApp storage clusters for high availability and disaster recovery in a campus or metropolitan area. The ONTAP Mediator deployed at a third site monitors the solution and enables automated failover in case of a site disaster. The following figure provides a high-level view of the solution components.



With the FlexPod SM-BC solution, you can deploy a VMware vSphere-based private cloud on a distributed and yet integrated infrastructure. The integrated solution enables multiple sites to be coordinated as a single solution infrastructure to protect data services from a variety of single-point-of-failure scenarios and a complete site failure.

This technical report highlights some of the end-to-end design considerations of the FlexPod SM-BC solution. The practitioners are encouraged to reference information available in the various FlexPod CVDs and NVAs for additional FlexPod solution implementation details.

Although the solution was validated by deploying two FlexPod systems based on FlexPod best practices as documented in CVDs, it takes into accounts the requirements for the SM-BC solution. The deployed FlexPod SM-BC solution discussed in this report has been validated for resiliency and fault tolerance during various failure scenarios as well as a simulated site failure scenario.

Solution requirements

The FlexPod SM-BC solution is designed to address the following key requirements:

- Business continuity for business-critical applications and data services in the event of a complete data center (site) failure
- Flexible, distributed workload placement with workload mobility across data centers
- Site affinity where virtual machine data is accessed locally, from the same data center site, during normal operations
- Quick recovery with zero data loss when a site failure occurs


Solution components

Cisco compute components

The Cisco UCS is an integrated computing infrastructure to provide unified computing resources, unified fabric, and unified management. It enables companies to automate and accelerate deployment of applications, including virtualization and bare-metal workloads. The Cisco UCS supports a wide range of deployment use cases including remote and branch locations, data centers, and hybrid cloud use cases. Depending on the specific solution requirements, the FlexPod Cisco compute implementation can utilize a variety of components at different scales. The following subsections provide additional information on some of the UCS components.

UCS server and compute node

The following figure shows some examples of the UCS server components, including UCS C- Series rack servers, UCS 5108 chassis with B-Series blade servers, and the new UCS X9508 chassis with X-Series compute nodes. The Cisco UCS C-Series rack servers are available in one and two rack-unit (RU) form factor, Intel and AMD CPU based models, and with various CPU speeds and cores, memory, and I/O options. The Cisco UCS B-Series blade servers and the new X-Series compute nodes are also available with various CPU, memory, and I/O options, and they are all supported in the FlexPod architecture to meet the diverse business requirements.



In addition to the latest generation C220/C225/C240/C245 M6 rack servers, B200 M6 blade servers, and X210c compute nodes shown in this figure, prior generations of rack and blade servers can also be used if they are still supported.


I/O Module and Intelligent Fabric Module

The I/O Module (IOM)/Fabric Extender and Intelligent Fabric Module (IFM) provide unified fabric connectivity for the Cisco UCS 5108 blade server chassis and the Cisco UCS X9508 X-Series chassis, respectively.

The fourth generation UCS IOM 2408 has eight 25-G unified Ethernet ports for connecting the UCS 5108 chassis with Fabric Interconnect (FI). Each 2408 has four 10-G backplane Ethernet connectivity through the midplane to each blade server in the chassis.

The UCSX 9108 25G IFM has eight 25-G unified Ethernet ports for connecting the blade servers in the UCS X9508 chassis with fabric interconnects. Each 9108 has four 25-G connections towards each UCS X210c compute node in the X9108 chassis. The 9108 IFM also works in concert with the fabric interconnect to manage the chassis environment.

The following figure depicts the UCS 2408 and earlier IOM generations for the UCS 5108 chassis and the 9108 IFM for the X9508 chassis.



UCS Fabric Interconnects

The Cisco UCS Fabric Interconnects (FIs) provide connectivity and management for the entire Cisco UCS. Typically deployed as an active/active pair, the system's FIs integrate all components into a single, highly available management domain controlled by the Cisco UCS Manager or Cisco Intersight. Cisco UCS FIs provide a single unified fabric for the system with low-latency and lossless, cut-through switching that supports LAN, SAN, and management traffic using a single set of cables.

There are two variants for the fourth-generation Cisco UCS FIs: UCS FI 6454 and 64108. They include support for 10/25 Gbps Ethernet ports, 1/10/25-Gbps Ethernet ports, 40/100-Gbps Ethernet up-link ports, and unified ports that can support 10/25 Gigabit Ethernet or 8/16/32-Gbps Fibre Channel. The following figure shows the fourth-generation Cisco UCS FIs along with the third-generation models that are also supported.






To support the Cisco UCS X-Series chassis, fourth-generation fabric interconnects configured in Intersight Managed Mode (IMM) are required. However, the Cisco UCS 5108 B-series chassis can be supported both in IMM mode and in UCSM managed mode.



The UCS FI 6324 uses the IOM form factor and is embedded in a UCS Mini chassis for deployments that require only a small UCS domain.

UCS Virtual Interface Cards

Cisco UCS Virtual Interface Cards (VICs) unify system management and LAN and SAN connectivity for rack and blade servers. It supports up to 256 virtual devices, either as virtual Network Interface Cards (vNICs) or as virtual Host Bus Adapters (vHBAs) using the Cisco SingleConnect technology. As a result of virtualization, the VIC cards greatly simplify the network connectivity and reduce the number of network adapters, cables, and switch ports needed for solution deployment. The following figure shows some of the Cisco UCS VICs available for the B-Series and C-Series servers and the X-Series compute nodes.



The different adapter models support different blade and rack servers with different port counts, port speeds, and form factors of modular LAN on Motherboard (mLOM), mezzanine cards, and PCIe interfaces. The adapters can support some combinations of 10/25/40/100-G Ethernet and Fibre Channel over Ethernet (FCoE). They incorporate Cisco's Converged Network Adapter (CNA) technology, support a comprehensive feature set, and simplify adapter management and application deployment. For example, the VIC supports Cisco's Data Center Virtual Machine Fabric Extender (VM-FEX) technology, which extends the Cisco UCS fabric interconnect ports to virtual machines, thus simplifying server virtualization deployment.

With a combination of Cisco VIC in mLOM, mezzanine, and port expander and bridge card configurations, you can take full advantage of the bandwidth and connectivity available to the blade servers. For example, by using the two 25-G links on the VIC 14825 (mLOM) and 14425 (mezzanine) and the 14000 (bridge card) for the X210c compute node, the combined VIC bandwidth is $2 \times 50\text{-G} + 2 \times 50\text{-G}$, or 100G per fabric/IFM and 200G total per server with the dual IFM configuration.

For details on the Cisco UCS product families, technical specifications, and documentations, see the [Cisco UCS](#) web site for information.

Cisco switching components

Nexus switches

FlexPod uses Cisco Nexus Series switches to provide Ethernet switching fabric for communications between Cisco UCS and NetApp storage controllers. All currently supported Cisco Nexus switch models, including the Cisco Nexus 3000, 5000, 7000, and 9000 Series, are supported for FlexPod deployment.

When selecting a switch model for FlexPod deployment, there are many factors to consider, such as performance, port speed, port density, switching latency, and protocols such as ACI and VXLAN support, for your design objectives as well as the switches' support timespan.

The validation for many recent FlexPod CVDs uses Cisco Nexus 9000 series switches such as the Nexus 9336C-FX2 and the Nexus 93180YC-FX3, which deliver high performance 40/100G and 10//25G ports, low latency, and exceptional power efficiency in a compact 1U form factor. Additional speeds are supported via uplink ports and breakout cables. The following figure shows a few Cisco Nexus 9k and 3k switches, including the Nexus 9336C-FX2 and the Nexus 3232C used for this validation.

Nexus 9336C-FX2



Nexus 93180YC-FX3



Nexus 3232C



See [Cisco Data Center Switches](#) for more information on the available Nexus switches and their specifications and documentations.

MDS switches

The Cisco MDS 9100/9200/9300 Series Fabric switches are an optional component in the FlexPod architecture. These switches are highly reliable, highly flexible, secure, and can provide visibility into the traffic flow in the fabric. The following figure shows some example MDS switches that can be used to build redundant FC SAN fabrics for a FlexPod solution to meet application and business requirements.

MDS 9132T



MDS 9148T




MDS 9148S



MDS 9250i



MDS 9396T



Cisco MDS 9132T/9148T/9396T high performance 32G Multilayer Fabric Switches are cost effective and are highly reliable, flexible, and scalable. The advanced storage networking features and functions come with ease of management and are compatible with the entire Cisco MDS 9000 family portfolio for a reliable SAN implementation.

State-of-the-art SAN analytics and telemetry capabilities are built into this next-generation hardware platform. The telemetry data extracted from the inspection of the frame headers can be streamed to an analytics-visualization platform, including the Cisco Data Center Network Manager. The MDS switches supporting 16G FC, such as the MDS 9148S, are also supported in FlexPod. In addition, Multiservice MDS switches, such as MDS 9250i, which supports FCoE and FCIP protocols in addition to FC protocol, are also part of the FlexPod solution portfolio.

On semi-modular MDS switches such as 9132T and 9396T, additional port expansion module and port licenses can be added to support additional device connectivity. On the fixed switches such as 9148T, additional port licenses can be added as needed. This pay-as-you-grow flexibility provides an operational expenses component to help reduce the capital expenses for the implementation and operation of MDS switch-based SAN infrastructure.

See [Cisco MDS Fabric Switches](#) for more information on the available MDS Fabric switches and see the [NetApp IMT](#) and [Cisco Hardware and Software Compatibility List](#) for a complete list of supported SAN switches.


NetApp components

Redundant NetApp AFF or ASA controllers running ONTAP software 9.8, or later releases are required to create a FlexPod SM-BC solution. The latest ONTAP release, currently 9.10.1, is recommended for SM-BC deployment to take advantage of continued ONTAP innovations, performance, and quality improvements and the increased maximum object count for SM-BC support.

NetApp AFF and ASA controllers with industry-leading performance and innovations provide enterprise data protection and feature-rich data management capabilities. The AFF and ASA systems support end-to-end NVMe technologies, including NVMe-attached SSDs and NVMe over Fibre Channel (NVMe/FC) front-end host connectivity. You can improve your workload throughput and reduce I/O latency by adopting NVMe/FC-based SAN infrastructure. However, NVMe/FC-based datastores can currently only be used for workloads not protected by SM-BC, because SM-BC solution currently supports only iSCSI and FC protocols.

NetApp AFF and ASA storage controllers also provide a hybrid-cloud foundation for customers to take advantages of the seamless data mobility enabled by NetApp Data Fabric. With Data Fabric, you can easily get data from the edge where it is generated to the core where it is used and to the cloud to take advantage of the on-demand elastic compute and AI and ML capabilities to gain actionable business insights.

As shown in the following figure, NetApp offers a variety of storage controllers and disk shelves to meet your performance and capacity requirements. See the following table for links to product pages for information about the NetApp AFF and ASA controller capabilities and specifications.




| Product family | Technical specifications |
|----------------|--|
| AFF series | AFF series documentation |
| ASA series | ASA series documentation |

Consult the [NetApp disk shelves and storage media documentation](#) and [NetApp Hardware Universe](#) for details on the disk shelves and the supported disk shelves for each storage controller model.


Solution topologies

FlexPod solutions are flexible in topology and can be scaled up or scaled out to meet different solution requirements. A solution that requires business continuity protection and only minimum compute and storage resources can use a simple solution topology, as illustrated in the following figure. This simple topology uses the UCS C-Series rack servers and AFF/ASA controllers with SSDs in the controller without additional disk shelves.



The redundant compute, network, and storage components are interconnected with redundant connectivity between the components. This highly available design provides solution resiliency and enables it to withstand single-point-of-failure scenarios. The multi-site design and ONTAP SM-BC synchronous data replication relationships provide business-critical data services despite the potential for single-site storage failure.


An asymmetric deployment topology that could be used by companies between a data center and a branch office in a metropolitan area might look like the following figure. For this asymmetric design, the data center requires a higher performance FlexPod with more compute and storage resources. However, the branch office requirement is less and can be met by a much smaller FlexPod.



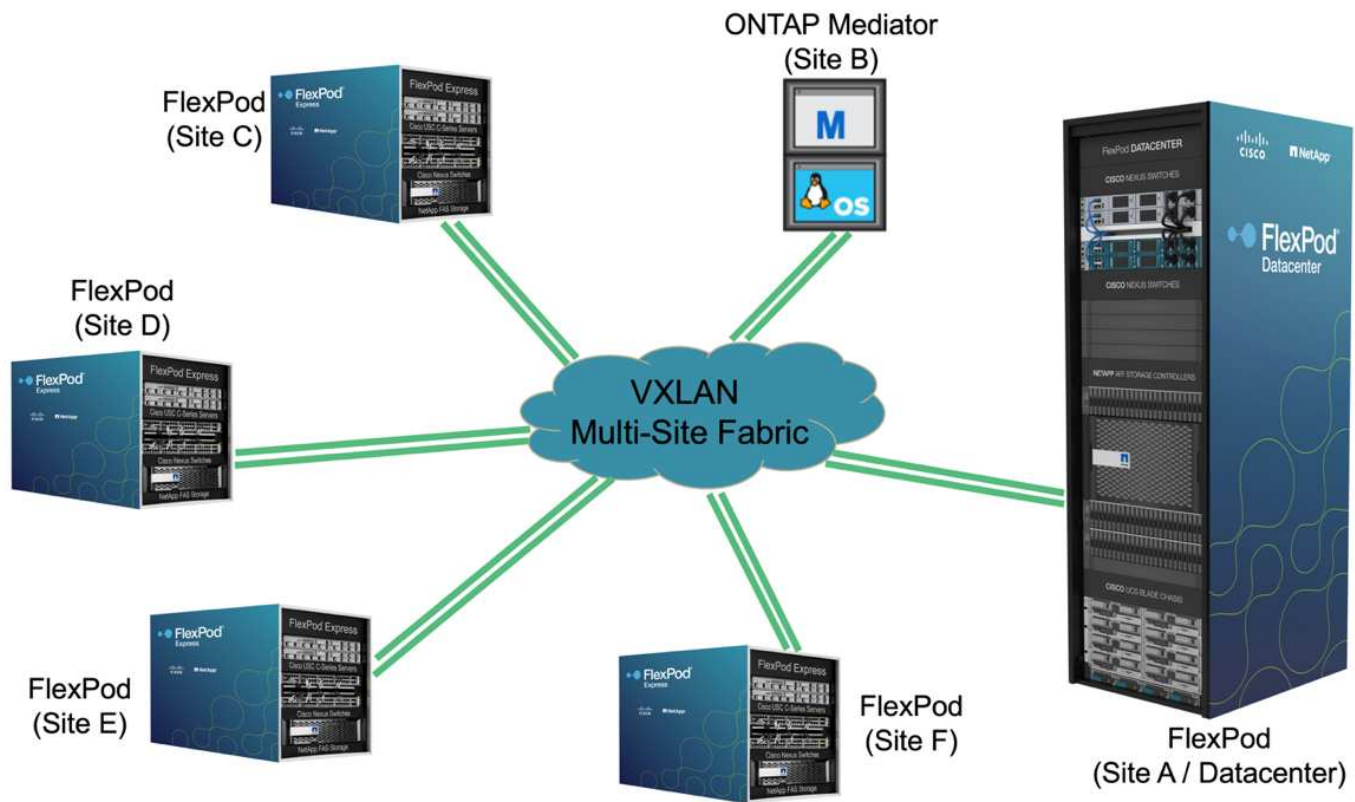
For companies with greater compute and storage resource requirements and multiple sites, a VXLAN-based multi-site fabric allows the multiple sites to have a seamless network fabric to facilitate application mobility so

an application can be served from any site.

There might be an existing FlexPod solution using the Cisco UCS 5108 chassis and B-Series blade servers that must be protected by a new FlexPod instance. The new FlexPod instance can use the latest UCS X9508 chassis with X210c compute nodes managed by Cisco Intersight, as shown in the following figure. In this case, the FlexPod systems at each site are connected to a larger data center fabric, and the sites are connected through an interconnect network to form a VXLAN multi-site fabric.



For companies that have a datacenter and several branch offices in a metro area that all need to be protected to provide business continuity, the FlexPod SM-BC deployment topology shown in the following figure can be implemented to protect critical application and data services to achieve zero RPO and near zero RTO objectives for all branch sites.



For this deployment model, each branch office establishes the SM-BC relationships and consistency groups it requires with the datacenter. You must take into account the supported SM-BC object limits, so the overall consistency group relationships and endpoint counts do not exceed the supported maximums at the datacenter.

[Next: Solution validation overview.](#)

Solution validation

Solution validation - Overview

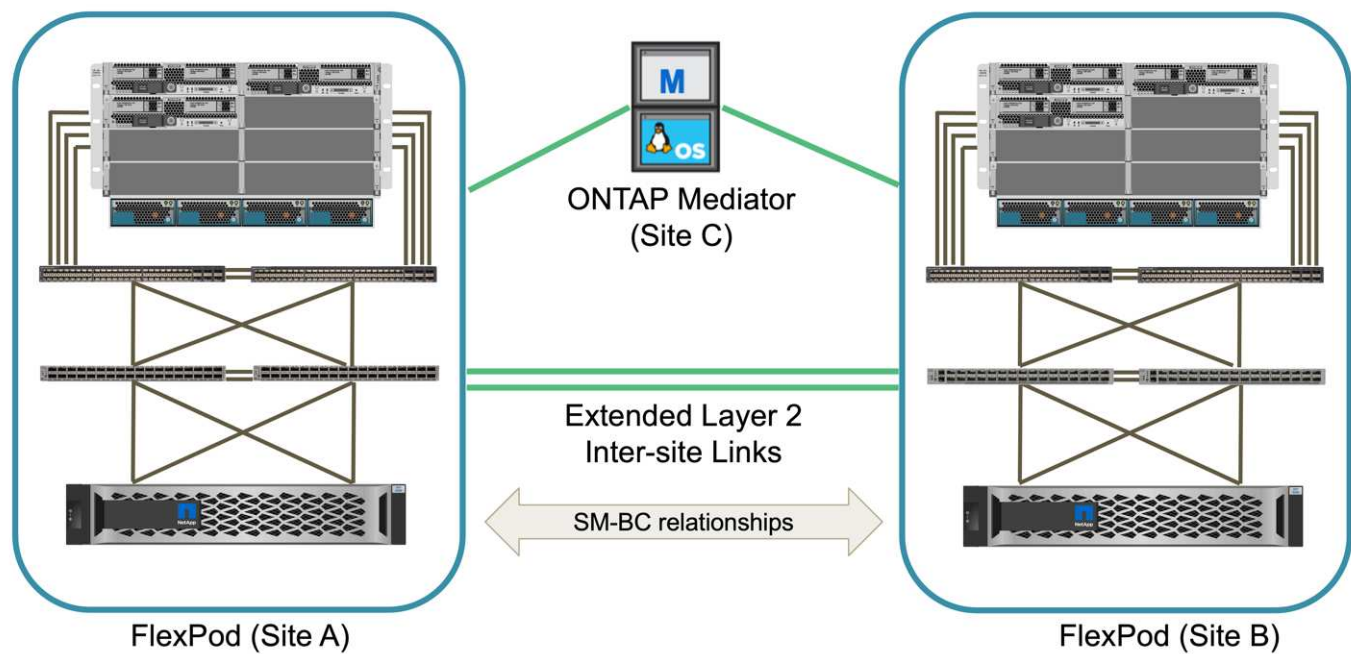
[Previous: FlexPod SM-BC solution.](#)

The FlexPod SM-BC solution design and implementation details depend on the specific FlexPod situation configuration and solution objectives. After the general business continuity requirements are defined, the FlexPod SM-BC solution can be created by implementing a completely new solution with two new FlexPod systems, adding a new FlexPod at another site to pair with an existing FlexPod, or by pairing two existing FlexPod systems together.

Since FlexPod solutions are flexible in nature in its configurations, all supported FlexPod configurations and components can potentially be used. The remainder of this section provides information for the implementation validations performed for a VMware-based virtual infrastructure solution. Except for the SM-BC related aspects, the implementation follows the standard FlexPod deployment processes. Please see the available FlexPod CVDs and NVAs appropriate for your specific configurations for general FlexPod implementation details.

Validation topology

For validation of the FlexPod SM-BC solution, supported technology components from NetApp, Cisco, and VMware are used. The solution features NetApp AFF A250 HA pairs running ONTAP 9.10.1, dual Cisco Nexus 9336C-FX2 switches at site A and dual Cisco Nexus 3232C switches at site B, Cisco UCS 6454 FIs at both sites, and three Cisco UCS B200 M5 servers at each site running VMware vSphere 7.0u2 and managed by UCS Manager and VMware vCenter server. The following figure shows the component-level solution validation topology with two FlexPod systems running at site A and site B connected by extended layer-2 inter-site links and ONTAP Mediator running at site C.



Hardware and software

The following table lists the hardware and software used for the solution validation. It is important to note that Cisco, NetApp, and VMware have interoperability matrixes used to determine support for any specific implementation of FlexPod:

- <http://support.netapp.com/matrix/>
- [Cisco UCS Hardware and Software Interoperability Tool](#)
- <http://www.vmware.com/resources/compatibility/search.php>

| Category | Component | Software version | Quantity |
|----------|--|------------------|----------------|
| Compute | Cisco UCS Fabric Interconnect 6454 | 4.2(1f) | 4 (2 per site) |
| | Cisco UCS B200 M5 servers | 4.2(1f) | 6 (3 per site) |
| | Cisco UCS IOM 2204XP | 4.2(1f) | 4 (2 per site) |
| | Cisco VIC 1440 (PID: UCSB-MLOM-40G-04) | 5.2(1a) | 2 (1 per site) |
| | Cisco VIC 1340 (PID: UCSB-MLOM-40G-03) | 4.5(1a) | 4 (2 per site) |

| Category | Component | Software version | Quantity |
|----------------|---|------------------|----------------|
| Network | Cisco Nexus 9336C-FX2 | 9.3(6) | 2 (site A) |
| | Cisco Nexus 3232C | 9.3(6) | 2 (site B) |
| Storage | NetApp AFF A250 | 9.10.1 | 4 (2 per site) |
| | NetApp System Manager | 9.10.1 | 2 (1 per site) |
| | NetApp Active IQ Unified Manager | 9.10 | 1 |
| | NetApp ONTAP Tools for VMware vSphere | 9.10 | 1 |
| | NetApp SnapCenter Plugin for VMware vSphere | 4.6 | 1 |
| | NetApp ONTAP Mediator | 1.3 | 1 |
| | NAbox | 3.0.2 | 1 |
| | NetApp Harvest | 21.11.1-1 | 1 |
| | | | |
| Virtualization | VMware ESXi | 7.0U2 | 6 (3 per site) |
| | VMware ESXi nenic Ethernet Driver | 1.0.35.0 | 6 (3 per site) |
| | VMware vCenter | 7.0U2 | 1 |
| | NetApp NFS Plug-in for VMware VAAI | 2.0 | 6 (3 per site) |
| | | | |
| Testing | Microsoft Windows | 2022 | 1 |
| | Microsoft SQL Server | 2019 | 1 |
| | Microsoft SQL Server Management Studio | 18.10 | 1 |
| | HammerDB | 4.3 | 1 |
| | Microsoft Windows | 10 | 6 (3 per site) |
| | IOMeter | 1.1.0 | 6 (3 per site) |

[Next: Solution validation - Compute.](#)


Solution validation - Compute

[Previous: Solution validation - Overview.](#)


The compute configuration for the FlexPod SM-BC solution follows typical FlexPod solution best practices. The following sections highlight some of the connectivity and configurations used for the validation. Some of the SM-BC-related considerations are also highlighted to provide implementation references and guidance.

Connectivity

The connectivity between the UCS B200 blade servers and the IOMs are provided by the UCS VIC card through the UCS 5108 chassis backplane connections. The UCS 2204XP Fabric Extenders used for the validation has sixteen 10G ports each to connect to the eight half-width blade servers, for example, two for each server. To increase server connectivity bandwidth, an additional mezzanine-based VIC can be added to connect the server to the alternative UCS 2408 IOM which provides four 10G connections to each server.



The connectivity between the UCS 5108 chassis and the UCS 6454 FIs used for the validation are provided by the IOM 2204XP which use four 10G connections. The FI ports 1 through 4 are configured as server ports for these connections. The FI ports 25 through 28 are configured as network uplink ports to the Nexus switch A and B at the local site. The following figure and table provide the connectivity diagram and port connection details for the UCS 6454 FIs to connect to the UCS 5108 chassis and the Nexus switches.



| Local device | Local port | Remote device | Remote port |
|---------------|------------|---------------|-------------|
| UCS 6454 FI A | 1 | IOM A | 1 |
| | 2 | | 2 |
| | 3 | | 3 |
| | 4 | | 4 |
| | 25 | Nexus A | 1/13/1 |
| | 26 | | 1/13/2 |
| | 27 | Nexus B | 1/13/3 |
| | 28 | | 1/13/4 |
| | L1 | UCS 6454 FI B | L1 |
| | L2 | | L2 |
| UCS 6454 FI B | 1 | IOM B | 1 |
| | 2 | | 2 |
| | 3 | | 3 |
| | 4 | | 4 |
| | 25 | Nexus A | 1/13/3 |
| | 26 | | 1/13/4 |
| | 27 | Nexus B | 1/13/1 |
| | 28 | | 1/13/2 |
| | L1 | UCS 6454 FI A | L1 |
| | L2 | | L2 |



The connections above are similar for both sites A and B, despite site A using Nexus 9336C-FX2 switches and site B using Nexus 3232C switches. 40G to 4x10G breakout cables are used for the Nexus to FI connections. The FI connections to Nexus utilizes port channel and virtual port channels are configured on the Nexus switches to aggregate the connections to each FI.



When using a different combination of IOM, FI, and Nexus switch components, be sure to use appropriate cables and port speed for the environment combination.



Additional bandwidth can be achieved by using components that support higher speed connections or more connections. Additional redundancy can be achieved by adding additional connections with components that support them.

Service profiles

A blade server chassis with fabric interconnects managed by UCS Manager (UCSM) or Cisco Intersight can abstract the servers by using service profiles available in UCSM and server profiles in Intersight. This validation uses UCSM and service profiles to simplify server management. With service profiles, replacing or upgrading a server can be done simply by associating the original service profile with the new hardware.

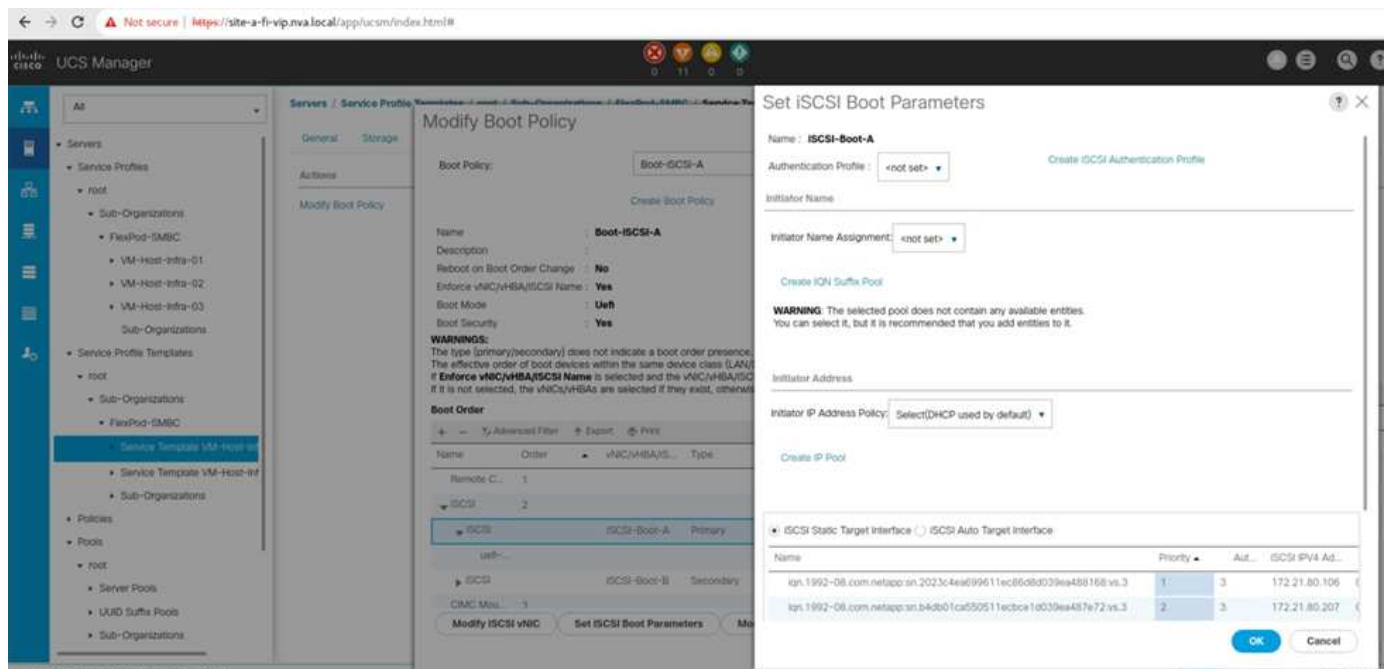
The created service profiles support the following for the VMware ESXi hosts:

- SAN boot from the AFF A250 storage at either site using iSCSI protocol.
- Six vNICs are created for the servers where:
 - Two redundant vNICs (vSwitch0-A and vSwitch0-B) carry in-band management traffic. Optionally, these vNICs can also be used by NFS protocol data that is not protected by SM-BC.
 - Two redundant vNICs (vDS-A and vDS-B) are used by the vSphere distributed switch to carry VMware vMotion and other application traffic.
 - iSCSI-A vNIC used by iSCSI-A vSwitch to provide access to iSCSI-A path.
 - iSCSI-B vNIC used by iSCSI-B vSwitch to provide access to iSCSI-B path.

SAN boot

For iSCSI SAN boot configuration, the iSCSI boot parameters are set to allow iSCSI boot from both iSCSI fabrics. To accommodate the SM-BC failover scenario in which an iSCSI SAN boot LUN is served from the secondary cluster when the primary cluster is not available, the iSCSI static target configuration should include targets from both site A and site B. In addition, to maximize boot LUN availability, configure the iSCSI boot parameter settings to boot from all storage controllers.

The iSCSI static target can be configured in the boot policy of service profile templates under the Set iSCSI Boot Parameter dialog as shown in the following figure. The recommended iSCSI boot parameter setting configuration is shown in the following table, which implements the boot strategy discussed above to achieve high availability.



| iSCSI fabric | Priority | iSCSI target | iSCSI LIF |
|--------------|----------|---------------------|---------------------------------|
| iSCSI A | 1 | Site A iSCSI target | Site A Controller 1 iSCSI A LIF |
| | 2 | Site B iSCSI target | Site B Controller 2 iSCSI A LIF |

| iSCSI fabric | Priority | iSCSI target | iSCSI LIF |
|--------------|----------|---------------------|---------------------------------|
| iSCSI B | 1 | Site B iSCSI target | Site B Controller 1 iSCSI B LIF |
| | 2 | Site A iSCSI target | Site A Controller 2 iSCSI B LIF |

[Next: Solution validation - Network.](#)

Solution validation - Network


[Previous: Solution validation - Compute.](#)

The network configuration for FlexPod SM-BC solution follows typical FlexPod solution best practices at each site. For inter-site connectivity, the solution validation configuration connects the FlexPod Nexus switches at the two sites together to provide inter-site connectivity that extends VLANs between the two sites. The following sections highlight some of the connectivity and configurations used for the validation.

Connectivity

The FlexPod Nexus switches at each site provides the local connectivity between the UCS compute and ONTAP storage in a highly available configuration. The redundant components and redundant connectivity provide the resiliency against single-point-of-failure scenarios.


The following diagram shows the Nexus switch local connectivity at each site. In addition to what is shown in the diagram, there are also console and management network connections for each component that are not shown. The 40G to 4 x 10G breakout cables are used to connect the Nexus switches to the UCS FIs and the ONTAP AFF A250 storage controllers. Alternatively, the 100G to 4 x 25G breakout cables can be used to increase the communication speed between the Nexus switches and the AFF A250 storage controllers. For simplicity, the two AFF A250 controllers are logically shown as side-by-side for cabling illustration. The two connections between the two storage controllers allow the storage to form a switchless cluster.



The following table shows the connectivity between Nexus switches and AFF A250 storage controllers at each site.

| Local device | Local port | Remote device | Remote port |
|--------------|------------|---------------|-------------|
| Nexus A | 1/10/1 | AFF A250 A | e1a |
| | 1/10/2 | | e1b |
| | 1/10/3 | AFF A250 B | e1a |
| | 1/10/4 | | e1b |
| Nexus B | 1/10/1 | AFF A250 A | e1c |
| | 1/10/2 | | e1d |
| | 1/10/3 | AFF A250 B | e1c |
| | 1/10/4 | | e1d |

The connectivity between the FlexPod switches at site A and site B is shown in the following figure with cabling details listed in the accompanying table. The connections between the two switches at each site are for the vPC peer links. On the other hand, the connections between the switches across sites provide the inter-site links. The links extend the VLANs across sites for intercluster communication, SM-BC data replication, in-band management, and data access for the remote site resources.



| Local device | Local port | Remote device | Remote port |
|-----------------|------------|-----------------|-------------|
| Site A switch A | 33 | Site B switch A | 31 |
| | 34 | | 32 |
| | 25 | Site A switch B | 25 |
| | 26 | | 26 |
| Site A switch B | 33 | Site B switch B | 31 |
| | 34 | | 32 |
| | 25 | Site A switch A | 25 |
| | 26 | | 26 |
| Site B switch A | 31 | Site A switch A | 33 |
| | 32 | | 34 |

| Local device | Local port | Remote device | Remote port |
|-----------------|------------|-----------------|-------------|
| | 25 | Site B switch B | 25 |
| | 26 | | 26 |
| Site B switch B | 31 | Site A switch B | 33 |
| | 32 | | 34 |
| | 25 | Site B switch A | 25 |
| | 26 | | 26 |



The table above lists connectivity from the perspectives of each FlexPod switch. As a result, the table contains duplicate information for readability.

Port channel and virtual port channel

Port channel enables link aggregation by using the Link Aggregation Control Protocol (LACP) for bandwidth aggregation and link failure resiliency. Virtual port channel (vPC) allows the port channel connections between two Nexus switches to logically appear as one. This further improves failure resiliency for scenarios such as a single link failure or a single switch failure.

The UCS server traffic to storage take the paths of IOM A to FI A and IOM B to FI B before reaching the Nexus switches. As the FI connections to Nexus switches utilize port channel on the FI side and virtual port channel on the Nexus switch side, the UCS server can effectively use paths through both Nexus switches and can survive single-point-of-failure scenarios. Between the two sites, the Nexus switches are inter-connected as illustrated in the previous figure. There are two links each to connect the switch pairs between the sites and they also use a port- channel configuration.

The in-band management, inter-cluster, and iSCSI / NFS data storage protocol connectivity is provided by interconnecting the storage controllers at each site to the local Nexus switches in a redundant configuration. Each storage controller is connected to two Nexus switches. The four connections are configured as part of an interface group on the storage for increased resiliency. On the Nexus switch side, those ports are also part of a vPC between switches.

The following table lists the port channel ID and usage at each site.

| Port channel ID | Usage |
|-----------------|-----------------------------|
| 10 | Local Nexus peer link |
| 15 | Fabric interconnect A links |
| 16 | Fabric interconnect B links |
| 27 | Storage controller A links |
| 28 | Storage controller B links |
| 100 | Inter-site switch A links |
| 200 | Inter-site switch B links |

VLANs

The following table lists VLANs configured for setting up the FlexPod SM-BC solution validation environment along with their usage.

| Name | VLAN ID | Usage |
|-------------------|---------|---|
| Native-VLAN | 2 | VLAN 2 used as native VLAN instead of default VLAN (1) |
| OOB-MGMT-VLAN | 3333 | Out-of-band management VLAN for devices |
| IB-MGMT-VLAN | 3334 | In-band management VLAN for ESXi hosts, VM management, etc. |
| NFS-VLAN | 3335 | Optional NFS VLAN for NFS traffic |
| iSCSI-A-VLAN | 3336 | iSCSI-A fabric VLAN for iSCSI traffic |
| iSCSI-B-VLAN | 3337 | iSCSI-B fabric VLAN for iSCSI traffic |
| vMotion-VLAN | 3338 | VMware vMotion traffic VLAN |
| VM-Traffic-VLAN | 3339 | VMware VM traffic VLAN |
| Intercluster-VLAN | 3340 | Intercluster VLAN for ONTAP cluster peer communications |



While SM-BC does not support NFS or CIFS protocols for business continuity, you can still use them for workloads that do not need to be protected for business continuity. NFS datastores were not created for this validation.

[Next: Solution validation - Storage.](#)


Solution validation - Storage

[Previous: Solution validation - Network.](#)

The storage configuration for FlexPod SM-BC solution follows typical FlexPod solution best practices at each site. For SM-BC cluster peering and data replication, they use the inter-site links established between the FlexPod switches at both sites. The following sections highlight some of the connectivity and configurations used for the validation.

Connectivity

The storage connectivity to the local UCS FIs and blade servers is provided by the Nexus switches at the local site. Through the Nexus switch connectivity between sites, the storage can also be accessed by the remote UCS blade servers. The following figure and table show the storage connectivity diagram and a list of connections for the storage controllers at each site.




| Local device | Local port | Remote device | Remote port |
|--------------|------------|---------------|-------------|
| AFF A250 A | e0c | AFF A250 B | e0c |
| | e0d | | e0d |
| | e1a | Nexus A | 1/10/1 |
| | e1b | | 1/10/2 |
| | e1c | Nexus B | 1/10/1 |
| | e1d | | 1/10/2 |
| AFF A250 B | e0c | AFF A250 A | e0c |
| | e0d | | e0d |
| | e1a | Nexus A | 1/10/3 |
| | e1b | | 1/10/4 |
| | e1c | Nexus B | 1/10/3 |
| | e1d | | 1/10/4 |

Connections and interfaces

Two physical ports on each storage controller are connected to each Nexus switches for bandwidth aggregation and redundancy for this validation. Those four connections participate in an interface group configuration on the storage. The corresponding ports on the Nexus switches participate in a vPC for link aggregation and resiliency.

The in-band management, inter-cluster, and NFS/iSCSI data storage protocols use VLANs. VLAN ports are created on the interface group to segregate the different types of traffic. Logical interfaces (LIFs) for the respective functions are created on top of the corresponding VLAN ports. The following figure shows the relationship between the physical connections, interface groups, VLAN ports, and logical interfaces.




SAN boot

NetApp recommends implementing SAN boot for the Cisco UCS servers in the FlexPod solution. Implementing SAN boot enables you to safely secure the operating system within the NetApp storage system, providing better performance and flexibility. For this solution, iSCSI SAN boot was validated.

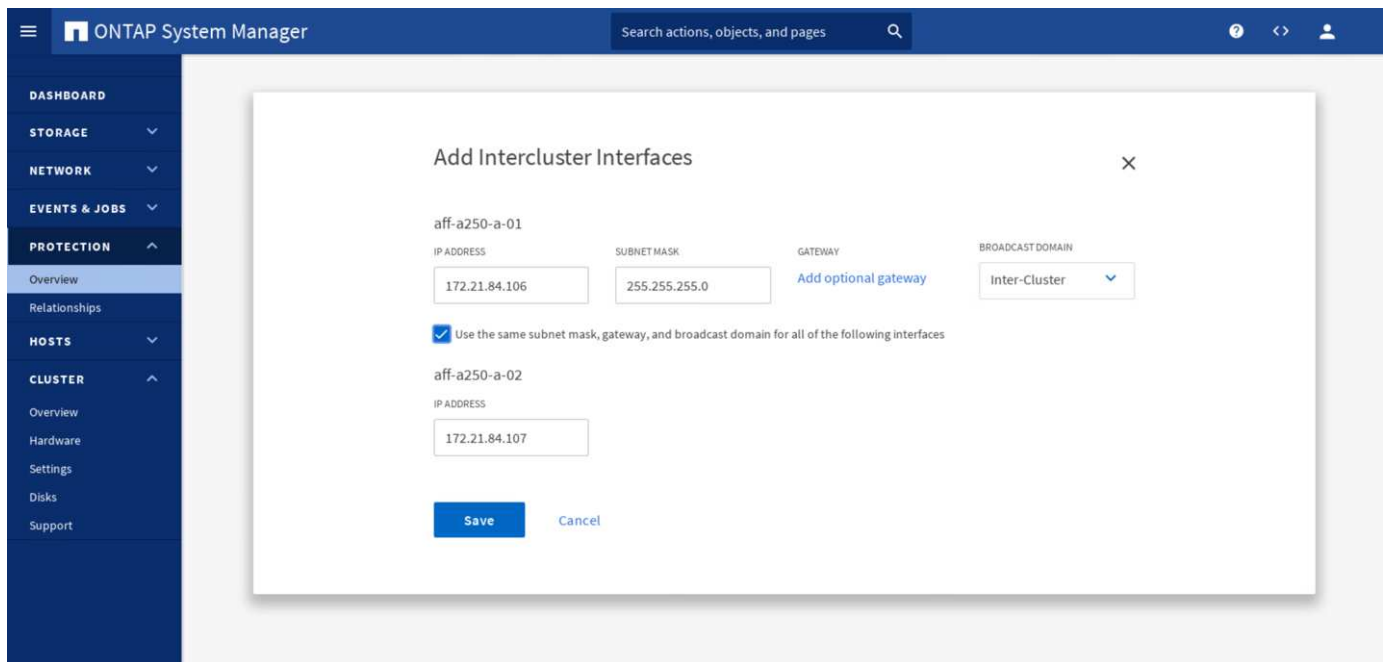
The following figure depicts the connectivity for iSCSI SAN boot of Cisco UCS server from NetApp Storage. In iSCSI SAN boot, each Cisco UCS server is assigned two iSCSI vNICs (one for each SAN fabric) that provide redundant connectivity from the server all the way to the storage. The 10/25-G Ethernet storage ports that are connected to the Nexus switches (in this example e1a, e1b, e1c, and e1d) are grouped together to form one interface group (ifgrp) (in this example, a0a). The iSCSI VLAN ports are created on the ifgrp and the iSCSI LIFs are created on the iSCSI VLAN ports.

Each iSCSI boot LUN is mapped to the server that boots from it through the iSCSI LIFs by associating the boot LUN with the server's iSCSI Qualified Names (IQNs) in its boot igroup. The server's boot igroup contains two IQNs, one for each vNIC / SAN fabric. This feature enables only the authorized server to have access to the boot LUN created specifically for that server.




Cluster peering

ONTAP cluster peers communicate via the intercluster LIFs. Using ONTAP System Manager for the two clusters, you can create the needed intercluster LIFs under the Protection > Overview pane.

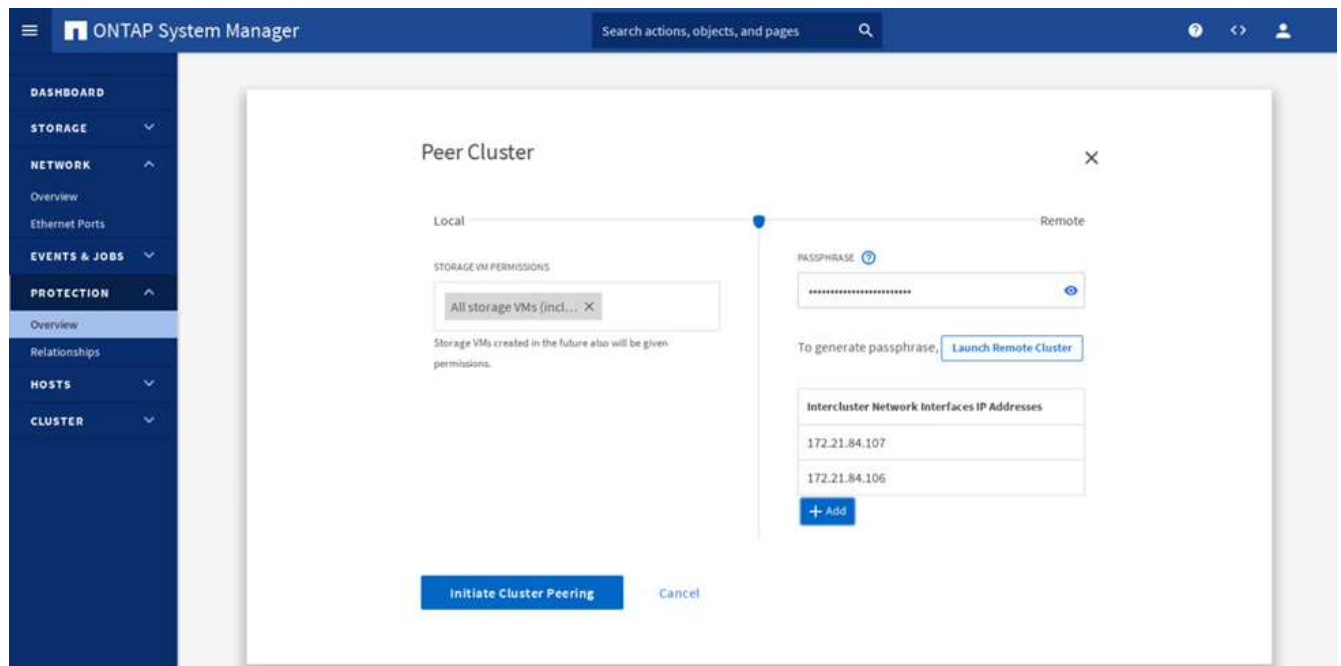


To peer the two clusters together, complete the following steps:


1. Generate cluster peering passphrase in the first cluster.



2. Invoke the Peer Cluster option in the second cluster and provide the passphrase and intercluster LIF information.



3. The System Manager Protection > Overview pane shows cluster peer information.




ONTAP Mediator installation and configuration


The ONTAP Mediator establishes a quorum for the ONTAP clusters in an SM-BC relationship. It coordinates automated failover when a failure is detected and helps to avoid split-brain scenarios when each cluster simultaneously tries to establish control as the primary cluster.

Before installing the ONTAP Mediator, check out the [Install or upgrade the ONTAP Mediator service](#) page for prerequisites, supported Linux versions, and the procedures for installing them on the various supported Linux operating systems.

After the ONTAP Mediator is installed, you can add the security certificate of the ONTAP Mediator to the ONTAP clusters and then configuring the ONTAP Mediator in the System Manager Protection > Overview pane. The following screenshot shows of the ONTAP Mediator configuration GUI.



After you provide the necessary information, the configured ONTAP Mediator then appears in the System Manager Protection > Overview pane.



SM-BC consistency group

A consistency group provides a write-order consistency guarantee for an application workload spanning a collection of specified volumes. For ONTAP 9.10.1, here are some of the important restrictions and limitations.

- The maximum number of SM-BC consistency group relationships in a cluster is 20.
- The maximum number of volumes supported per SM-BC relationship is 16.
- The maximum number of total source and destination endpoints in a cluster is 200.

For additional details, see the ONTAP SM-BC documentation on the [restrictions and limitations](#).

For the validation configuration, ONTAP System Manager was used to create the consistency groups to protect both the ESXi boot LUNs and the shared datastore LUNs for both sites. The consistency group creation dialog is accessible by going to Protection > Overview > Protect for Business Continuity > Protect Consistency Group. To create a consistency group, provide the needed source volumes, destination cluster, and destination storage virtual machine information for the creation.

The following table lists the four consistency groups that are created and the volumes that are included in each consistency group for the validation testing.

| System Manager | Consistency group | Volumes |
|----------------|----------------------|--|
| Site A | cg_esxi_a | esxi_a |
| Site A | cg_infra_datastore_a | infra_datastore_a_01 infra_datastore_a_02 |
| Site B | cg_esxi_b | esxi_b |
| Site B | cg_infra_datastore_b | infra_datastore_b_01 infra_datastore_b_02 |

After the consistency groups are created, they show up under the respective protection relationships in site A and site B.


This screenshot shows the consistency group relationships at site A.

| Source | Destination | Protection Policy | Relationship Health | State | Lag |
|-------------------------------------|--|-------------------|---------------------|---------|----------|
| Infra-SVM.1/cg/cg_infra_datastore_b | Infra-SVM-a/cg/cg_infra_datastore_b_dest | AutomatedFailOver | Healthy | In sync | 0 second |
| Infra-SVM.1/cg/cg_esxi_b | Infra-SVM-a/cg/cg_esxi_b_dest | AutomatedFailOver | Healthy | In sync | 0 second |

This screenshot shows the consistency group relationships at site B.

| Source | Destination | Protection Policy | Relationship Health | State | Lag |
|-------------------------------------|--|-------------------|---------------------|---------|----------|
| Infra-SVM.1/cg/cg_esxi_a | Infra-SVM-b/cg/cg_esxi_a_dest | AutomatedFailOver | Healthy | In sync | 0 second |
| Infra-SVM.1/cg/cg_infra_datastore_a | Infra-SVM-b/cg/cg_infra_datastore_a_dest | AutomatedFailOver | Healthy | In sync | 0 second |

This screenshot shows the consistency group relationship details for the cg_infra_datastore_b group.



Volumes, LUNs, and host mappings

After the consistency groups are created, SnapMirror synchronizes the source and the destination volumes so the data can always be in sync. The destination volumes at the remote site carries the volume names with the _dest ending. For example, for the esxi_a volume in site A cluster, there is a corresponding esxi_a_dest data protection (DP) volume in site B.

This screenshot shows the volume information for site A.

```
aff-a250-a::> vol show -vserver Infra-SVM-a
Vserver Volume Aggregate State Type Size Available Used%
-----
Infra-SVM-a esxi_a aggr1_aff_a250_a_01 online RW 320GB 315.9GB 1%
Infra-SVM-a esxi_b_dest aggr1_aff_a250_a_02 online DP 3.86GB 638.4MB 83%
Infra-SVM-a infra_datastore_a_01 aggr1_aff_a250_a_01 online RW 1TB 717.6GB 29%
Infra-SVM-a infra_datastore_a_02 aggr1_aff_a250_a_02 online RW 1TB 828.4GB 19%
Infra-SVM-a infra_svm_root aggr1_aff_a250_a_01 online RW 1GB 966.5MB 0%
Infra-SVM-a infra_svm_root_m01 aggr1_aff_a250_a_01 online LS 1GB 966.6MB 0%
Infra-SVM-a infra_svm_root_m02 aggr1_aff_a250_a_02 online LS 1GB 966.6MB 0%
Infra-SVM-a vol_infra_datastore_b_01_dest aggr1_aff_a250_a_01 online DP 138.7GB 31.52GB 76%
Infra-SVM-a vol_infra_datastore_b_02_dest aggr1_aff_a250_a_01 online DP 49.37GB 9.03GB 80%
9 entries were displayed.
```

This screenshot shows the volume information for site B.

```
aff-a250-b::> vol show -vserver Infra-SVM-b
Vserver Volume Aggregate State Type Size Available Used%
-----
Infra-SVM-b esxi_a_dest aggr1_aff_a250_b_02 online DP 4.10GB 768.2MB 80%
Infra-SVM-b esxi_b aggr1_aff_a250_b_01 online RW 320GB 315.8GB 1%
Infra-SVM-b infra_datastore_b_01 aggr1_aff_a250_b_01 online RW 1TB 911.9GB 10%
Infra-SVM-b infra_datastore_b_02 aggr1_aff_a250_b_02 online RW 1TB 964.0GB 5%
Infra-SVM-b infra_svm_root aggr1_aff_a250_b_01 online RW 1GB 966.9MB 0%
Infra-SVM-b infra_svm_root_m01 aggr1_aff_a250_b_01 online LS 1GB 967.0MB 0%
Infra-SVM-b infra_svm_root_m02 aggr1_aff_a250_b_02 online LS 1GB 967.0MB 0%
Infra-SVM-b vol_infra_datastore_a_01_dest aggr1_aff_a250_b_02 online DP 270.0GB 27.39GB 89%
Infra-SVM-b vol_infra_datastore_a_02_dest aggr1_aff_a250_b_02 online DP 202.8GB 28.20GB 85%
9 entries were displayed.
```

To facilitate transparent application failover, the mirrored SM-BC LUNs also need to be mapped to the hosts from the destination cluster. This allows the hosts to properly see paths to the LUNs from both the source and destination clusters. The `igroup show` and `lun show` outputs for both site A and site B are captured in the following two screenshots. With the created mappings, each ESXi host in the cluster sees its own SAN boot LUN as ID 0 and all the four shared iSCSI datastore LUNs.

This screenshot shows the host igroups and LUN mapping for site A cluster.

```

aff-a250-a:> igroup show
Vserver    Igroup      Protocol OS Type  Initiators
-----
Infra-SVM-a MGMT-Hosts  iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-a:1
                                     iqn.2010-11.com.flexpod:ucs-smbc-a:2
                                     iqn.2010-11.com.flexpod:ucs-smbc-a:3
                                     iqn.2010-11.com.flexpod:ucs-smbc-b:1
                                     iqn.2010-11.com.flexpod:ucs-smbc-b:2
                                     iqn.2010-11.com.flexpod:ucs-smbc-b:3
Infra-SVM-a VM-Host-Infra-a-01 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-a:1
Infra-SVM-a VM-Host-Infra-a-02 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-a:2
Infra-SVM-a VM-Host-Infra-a-03 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-a:3
Infra-SVM-a VM-Host-Infra-b-01 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-b:1
Infra-SVM-a VM-Host-Infra-b-02 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-b:2
Infra-SVM-a VM-Host-Infra-b-03 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-b:3
7 entries were displayed.

aff-a250-a:> lun show -m
Vserver    Path                                     Igroup    LUN ID  Protocol
-----
Infra-SVM-a /vol/esxi_a/VM-Host-Infra-a-01          VM-Host-Infra-a-01  0  iscsi
Infra-SVM-a /vol/esxi_a/VM-Host-Infra-a-02          VM-Host-Infra-a-02  0  iscsi
Infra-SVM-a /vol/esxi_a/VM-Host-Infra-a-03          VM-Host-Infra-a-03  0  iscsi
Infra-SVM-a /vol/esxi_a/swap_lun_a           MGMT-Hosts    13  iscsi
Infra-SVM-a /vol/esxi_b_dest/VM-Host-Infra-b-01      VM-Host-Infra-b-01  0  iscsi
Infra-SVM-a /vol/esxi_b_dest/VM-Host-Infra-b-02      VM-Host-Infra-b-02  0  iscsi
Infra-SVM-a /vol/esxi_b_dest/VM-Host-Infra-b-03      VM-Host-Infra-b-03  0  iscsi
Infra-SVM-a /vol/esxi_b_dest/swap_lun_b           MGMT-Hosts    23  iscsi
Infra-SVM-a /vol/infra_datastore_a_01/datastore_lun_a_01 MGMT-Hosts    11  iscsi
Infra-SVM-a /vol/infra_datastore_a_02/datastore_lun_a_02 MGMT-Hosts    12  iscsi
Infra-SVM-a /vol/vol_infra_datastore_b_01_dest/datastore_lun_b_01 MGMT-Hosts    21  iscsi
Infra-SVM-a /vol/vol_infra_datastore_b_02_dest/datastore_lun_b_02 MGMT-Hosts    22  iscsi
12 entries were displayed.

```

This screenshot shows the host igroups and LUN mapping for site B cluster.

```

aff-a250-b:> igroup show
Vserver    Igroup      Protocol OS Type  Initiators
-----
Infra-SVM-b MGMT-Hosts  iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-b:1
                               iqn.2010-11.com.flexpod:ucs-smbc-b:2
                               iqn.2010-11.com.flexpod:ucs-smbc-b:3
                               iqn.2010-11.com.flexpod:ucs-smbc-a:1
                               iqn.2010-11.com.flexpod:ucs-smbc-a:2
                               iqn.2010-11.com.flexpod:ucs-smbc-a:3
Infra-SVM-b VM-Host-Infra-a-01 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-a:1
Infra-SVM-b VM-Host-Infra-a-02 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-a:2
Infra-SVM-b VM-Host-Infra-a-03 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-a:3
Infra-SVM-b VM-Host-Infra-b-01 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-b:1
Infra-SVM-b VM-Host-Infra-b-02 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-b:2
Infra-SVM-b VM-Host-Infra-b-03 iscsi   vmware   iqn.2010-11.com.flexpod:ucs-smbc-b:3
7 entries were displayed.

aff-a250-b:> lun show -m
Vserver    Path                                     Igroup    LUN ID  Protocol
-----
Infra-SVM-b /vol/esxi_a_dest/VM-Host-Infra-a-01    VM-Host-Infra-a-01  0  iscsi
Infra-SVM-b /vol/esxi_a_dest/VM-Host-Infra-a-02    VM-Host-Infra-a-02  0  iscsi
Infra-SVM-b /vol/esxi_a_dest/VM-Host-Infra-a-03    VM-Host-Infra-a-03  0  iscsi
Infra-SVM-b /vol/esxi_a_dest/swap_lun_a          MGMT-Hosts    13  iscsi
Infra-SVM-b /vol/esxi_b/VM-Host-Infra-b-01        VM-Host-Infra-b-01  0  iscsi
Infra-SVM-b /vol/esxi_b/VM-Host-Infra-b-02        VM-Host-Infra-b-02  0  iscsi
Infra-SVM-b /vol/esxi_b/VM-Host-Infra-b-03        VM-Host-Infra-b-03  0  iscsi
Infra-SVM-b /vol/esxi_b/swap_lun_b              MGMT-Hosts    23  iscsi
Infra-SVM-b /vol/infra_datastore_b_01/datastore_lun_b_01 MGMT-Hosts    21  iscsi
Infra-SVM-b /vol/infra_datastore_b_02/datastore_lun_b_02 MGMT-Hosts    22  iscsi
Infra-SVM-b /vol/vol_infra_datastore_a_01_dest/datastore_lun_a_01 MGMT-Hosts    11  iscsi
Infra-SVM-b /vol/vol_infra_datastore_a_02_dest/datastore_lun_a_02 MGMT-Hosts    12  iscsi
12 entries were displayed.

```

[Next: Solution validation - Virtualization.](#)

Solution validation - Virtualization

[Previous: Solution validation - Storage.](#)

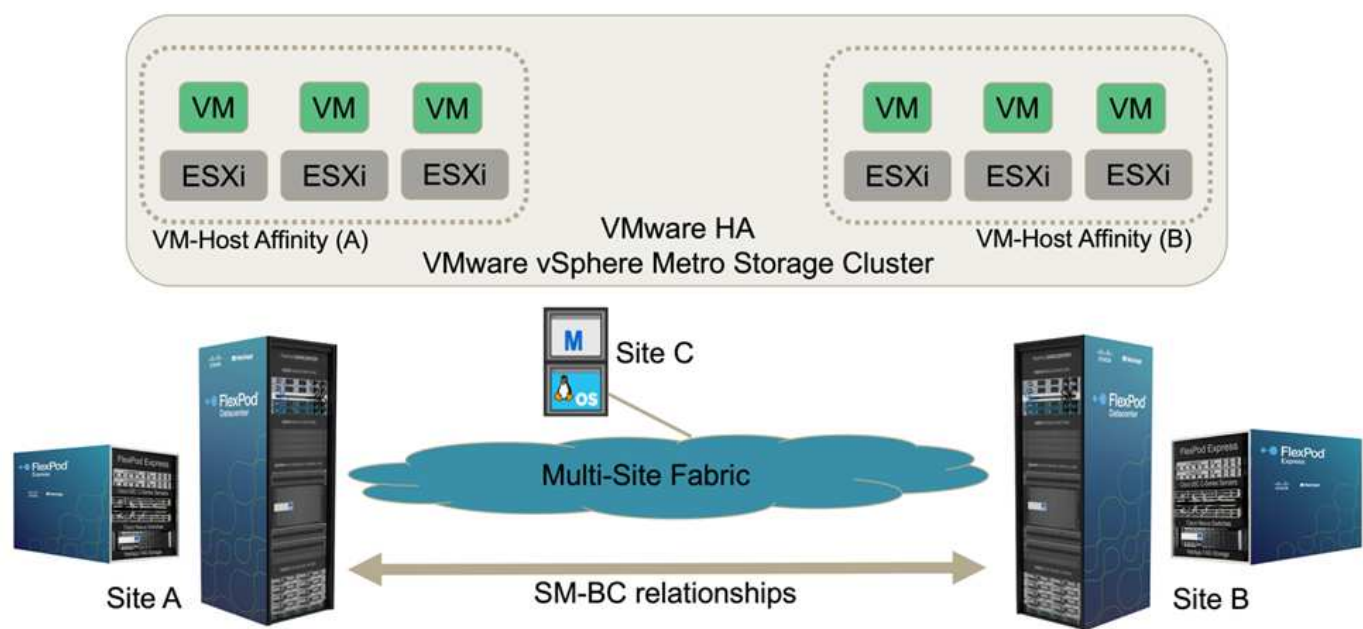
In the multi-site FlexPod SM-BC solution, a single VMware vCenter manages the virtual infrastructure resources for the entire solution. The hosts in both data centers participate in the single VMware HA cluster which spans both data centers. The hosts have access to the NetApp SM-BC solution where storage with defined SM-BC relationships can be accessed from both sites.

The SM-BC solution storage conforms to the uniform access model in the VMware vSphere Metro Storage Cluster (vMSC) feature to avoid disaster and downtime. For optimal virtual-machine performance, the virtual-machine disks should be hosted on the local NetApp AFF A250 systems to minimize latency and traffic across the WAN links under normal operation.

As part of the design implementation, the distribution of the virtual machines across the two sites must be determined. You can determine this virtual machine site affinity and application distribution across the two sites according to your site preferences and application requirements. The VMware cluster VM/Host Groups and VM/Host Rules are used to configure VM/Host affinity to make sure that VMs are running on hosts at the desired site.


However, configurations allowing the VMs to run at both sites will make sure that VMs can be restarted by VMware HA at remote-site hosts to provide solution resiliency. To accommodate virtual machines to run at both sites, all the iSCSI shared datastores must be mounted on all the ESXi hosts to ensure a smooth vMotion operation of virtual machines between sites.

The following figure shows a high-level FlexPod SM-BC solution virtualization view which includes both VMware HA and vMSC features to provide high availability for compute and storage services. The active-active datacenter solution architecture enables workload mobility between sites and provides DR/BC protection.



End-to-end network connectivity

The FlexPod SM-BC solution includes FlexPod infrastructures at each site, network connectivity between sites, and the ONTAP mediator deployed at a third site to meet the required RPO and RTO objectives. The following figure shows the end-to-end network connectivity between the Cisco UCS B200M5 servers at each site and the NetApp storage featuring SM-BC capabilities within a site and across sites.




The FlexPod deployment architecture is identical at each site for this solution validation. However, the solution supports asymmetric deployments and can also be added onto an existing FlexPod solutions if they meet the requirements.

Extended layer-2 architecture is used for a seamless multi-site data fabric that provides connectivity between port-channelled Cisco UCS compute and NetApp storage in each data center, as well as connectivity between data centers. Port channel configuration, and virtual port channel configuration where appropriate, is used for bandwidth aggregation and fault tolerance between the compute, network, and storage layers as well as for the cross-site links. As a result, The UCS blade servers have connectivity and multipath access to both local and remote NetApp storage.

Virtual networking

Each host in the cluster is deployed using identical virtual networking regardless of its location. The design separates the different traffic types using VMware virtual switches (vSwitch) and VMware Virtual Distributed Switches (vDS). The VMware vSwitch is used primarily for the FlexPod infrastructure networks and vDS for application networks, but it is not required.

The virtual switches (vSwitch, vDS) are deployed with two uplinks per virtual switch; the uplinks at the ESXi hypervisor level are referred to as vmnics and virtual NICs (vNICs) on Cisco UCS Software. The vNICs are created on the Cisco UCS VIC adapter in each server using Cisco UCS service profiles. Six vNICs are defined, two for vSwitch0, two for vDS0, two for vSwitch1, and two for the iSCSI uplinks as shown in the following figure.



vSwitch0 is defined during VMware ESXi host configuration, and it contains the FlexPod infrastructure management VLAN and the ESXi host VMkernel (VMK) ports for management. An infrastructure management virtual machine port group is also placed on vSwitch0 for any critical infrastructure management virtual machines that are needed.

It is important to place such management infrastructure virtual machines on vSwitch0 instead of the vDS because if the FlexPod infrastructure is shut down or power cycled and you attempt to activate that management virtual machine on a host other than the host on which it was originally running, it boots up fine on the network on vSwitch0. This process is particularly important if VMware vCenter is the management virtual machine. If vCenter were on the vDS and moved to another host and then booted, it would not be connected to the network after booting up.

Two iSCSI boot vSwitches are used in this design. Cisco UCS iSCSI boot requires separate vNICs for iSCSI boot. These vNICs use iSCSI VLAN of the appropriate fabric as the native VLAN and are attached to the appropriate iSCSI boot vSwitch. Optionally, you could also deploy iSCSI networks on vDS by deploying a new vDS or using an existing one.


VM-Host affinity groups and rules

To enable virtual machines to run on any ESXi host at both SM-BC sites, all ESXi hosts must mount the iSCSI datastores from both sites. If the datastores from both sites are properly mounted by all ESXi hosts, you can migrate a virtual machine between any hosts with vMotion and the VM still maintains access to all its virtual disks created from those datastores.


For a virtual machine that uses local datastores, its access to virtual disks becomes remote if it is migrated to a host at the remote site and thus increasing read operation latency due to the physical distance between the sites. Therefore, it is a best practice to keep virtual machines on the local hosts and utilize local storage at the site.

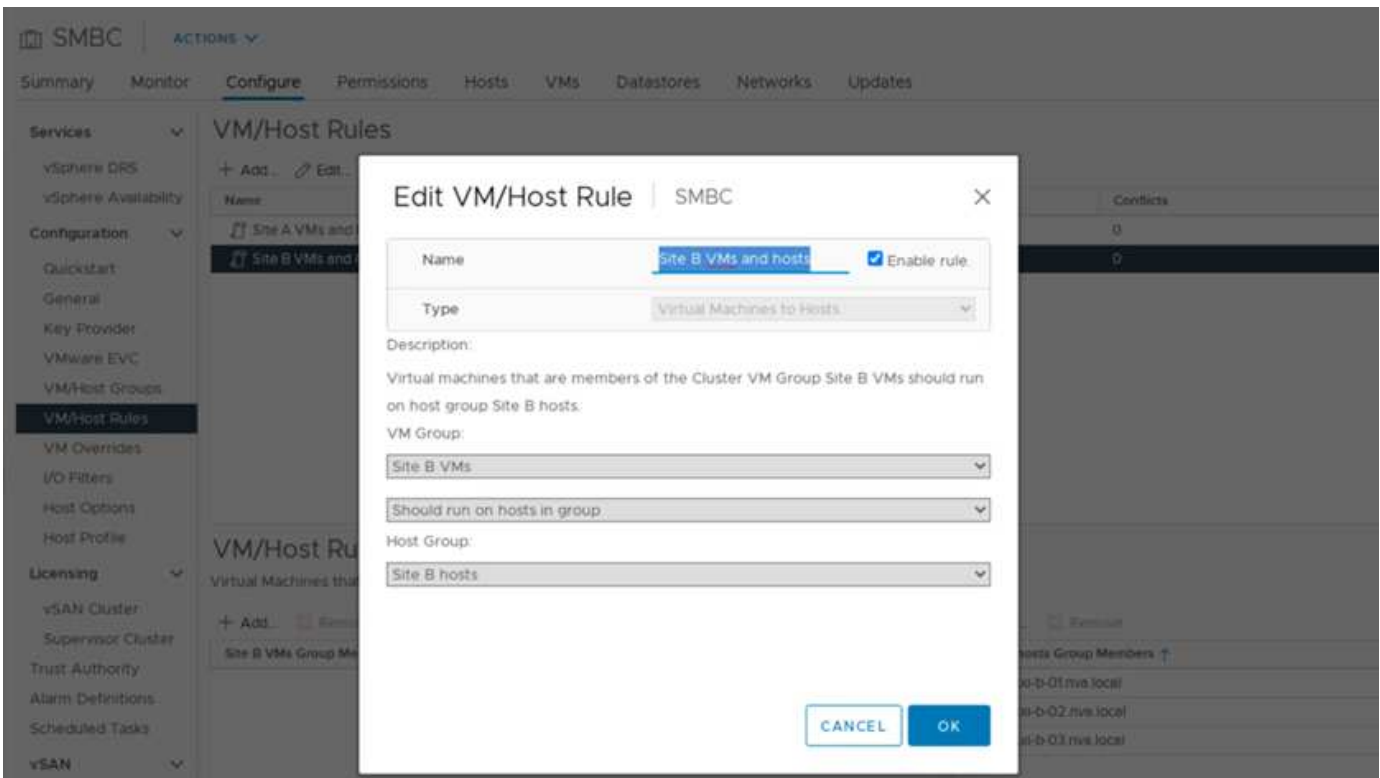
By using a VM/host affinity mechanism, you can use VM/Host Groups to create a VM group and a host group for virtual machines and hosts located at a particular site. Using VM/Host Rules, you can specify the policy for the VMs and hosts to follow. To allow virtual-machine migration across sites during a site maintenance or disaster scenario, use the “Should run on hosts in group” policy specification for that flexibility.

The following screenshot shows that two host groups and two VM groups are created for site A and site B hosts and VMs



In addition, the following two figures show the VM/Host rules that are created for site A and site B VMs to run on the hosts in their respective sites using the “Should run on hosts in group” policy.

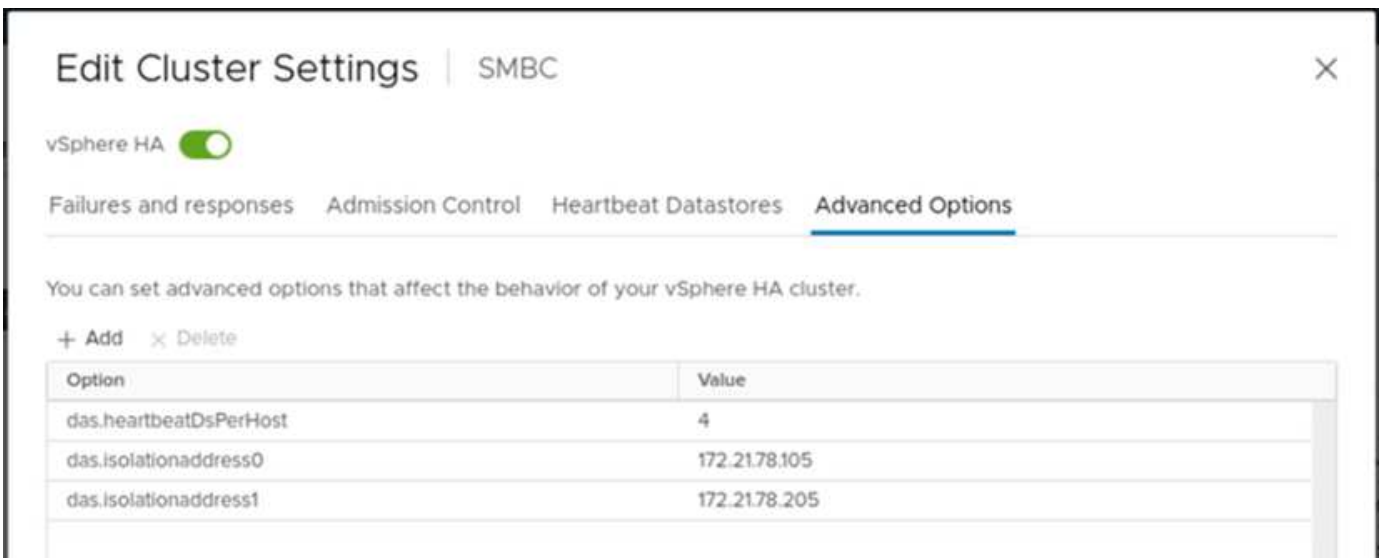




vSphere HA heartbeat

VMware vSphere HA has a heartbeat mechanism for host state validation. The primary heartbeat mechanism is through networking, and the secondary heartbeat mechanism is through the datastore. If heartbeats are not received, it then decides if it is isolated from the network by pinging the default gateway or the manually configured isolation addresses. For the datastore heartbeat, VMware recommends increasing the heartbeat datastores from the minimum of two to four for a stretched cluster.

For the solution validation, the two ONTAP cluster management IP addresses are used as the isolation address. In addition, the recommended vSphere HA advanced option `das.heartbeatDsPerHost` with a value of 4 was added as shown in the following figure.



For the heartbeat datastore, specify the four shared datastores from the cluster and complement automatically, as shown in the following figure.

Edit Cluster Settings
SMBC

vSphere HA
☒

Failures and responses
Admission Control
Heartbeat Datastores
Advanced Options

vSphere HA uses datastores to monitor hosts and virtual machines when the HA network has failed. vCenter Server selects 2 datastores for each host using the policy and datastore preferences specified below.

Heartbeat datastore selection policy:

☐ Automatically select datastores accessible from the hosts
☐ Use datastores only from the specified list
☒ Use datastores from the specified list and complement automatically if needed

Available heartbeat datastores

| | Name | Datastore Cluster | Hosts Mounting Datastore ↓ |
|-------------------------------------|----------------------|-------------------|----------------------------|
| <input type="checkbox"/> | infra_swap_a | N/A | 6 |
| <input type="checkbox"/> | infra_swap_b | N/A | 6 |
| <input checked="" type="checkbox"/> | infra_datastore_b_02 | N/A | 6 |
| <input checked="" type="checkbox"/> | infra_datastore_a_01 | N/A | 6 |
| <input checked="" type="checkbox"/> | infra_datastore_a_02 | N/A | 6 |
| <input checked="" type="checkbox"/> | infra_datastore_b_01 | N/A | 6 |

CANCEL
OK

For additional best practices and configurations for VMware HA Cluster and VMware vSphere Metro storage cluster, see [Creating and Using vSphere HA Clusters](#), [VMware vSphere Metro Storage Cluster \(vMSC\)](#) and the VMWare KB for [NetApp ONTAP with NetApp SnapMirror Business Continuity \(SM-BC\)](#) and [VMware vSphere Metro Storage Cluster \(vMSC\)](#).

[Next: Solution validation - Validated scenarios.](#)

Solution validation - Validated scenarios

[Previous: Solution validation - Virtualization.](#)

The FlexPod Datacenter SM-BC solution protects data services for a variety of single-point-of-failure scenarios as well as for a site disaster. The redundant design implemented at each site provides high availability, and the SM-BC implementation with synchronous data replication across sites protects data services from a sitewide disaster of one site. The deployed solution is validated for its desired solution functions and various failure scenarios for which the solution is designed to protect.

Solution functions validation

A variety of test cases are used to verify solution functions and simulate partial and complete site failure scenarios. To minimize duplication with the tests already performed in the existing FlexPod Datacenter solutions under Cisco Validated Design Program, the focus of this report is on the SM-BC related aspects of the solution. Some general FlexPod validations are included for practitioners to go through for their implementation validations.

For the solution validation, one Windows 10 virtual machine per ESXi host was created on all ESXi hosts at both sites. The IOMeter tool was installed and used to generate I/O to two virtual data disks that are mapped from the shared local iSCSI datastores. The IOMeter workload parameters configured were 8-KB I/O, 75% read, and 50% random, with 8 outstanding I/O commands for each data disk. For most of the test scenarios performed, the continuation of IOMeter I/O serves as an indication that the scenario did not cause a data service outage.

Since SM-BC is critical for business applications such as database servers, the Microsoft SQL server 2019 instance on a Windows server 2022 virtual machine was also included as part of the testing to confirm that the application continues to run when storage at its local site is not available and data service is resumed at the remote site storage without application disruptions.

ESXi Host iSCSI SAN boot test


The ESXi hosts in the solution are configured to boot from iSCSI SAN. Using SAN boot simplifies server management when replacing a server because the service profile of the server can be associated with a new server for it to boot up without making any additional configuration changes.

In addition to booting an ESXi host located at a site from its local iSCSI boot LUN, testing was also performed to boot the ESXi host when its local storage controller is in a takeover state or when its local storage cluster is completely unavailable. These validation scenarios make sure that the ESXi hosts are properly configured per design and can boot up during a storage maintenance or disaster scenario for disaster-recovery to provide business continuity.

Before the SM-BC consistency group relationship is configured, an iSCSI LUN hosted by a storage controller HA pair has four paths, two through each iSCSI fabric, based on the implementation of best practices. A host can get to the LUN through the two iSCSI VLANs/fabrics to the LUN hosting controller as well as through the high-availability partner of the controller.

After the SM-BC consistency group relationship is configured and the mirrored LUNs are properly mapped to the initiators, the path count for the LUN doubles. For this implementation, it goes from having two active/optimized paths and two active/non-optimized paths to having two active/optimized paths and six active/non-optimized paths.

The following figure illustrates the paths an ESXi host can take to access a LUN, for example, LUN 0. As the LUN is attached to the site A controller 01, only the two paths directly accessing the LUN via that controller are active/optimized and all the remaining six paths are active/nonoptimized.



The following screenshot of the storage-device-path information shows how the ESXi host sees the two types of device paths. The two active/optimized paths are shown as having `active (I/O)` path status, whereas the six active/non-optimized paths are shown only as `active`. Also note that the Target column shows the two iSCSI targets and the respective iSCSI LIF IP addresses to get to the targets.

esxi-a-01.nva.local

Summary Monitor **Configure** Permissions VMs Datastores Networks Updates

Storage

Storage Adapters

Host Cache Configuration

Protocol Endpoints

I/O Filters

Networking

Virtual switches

VMkernel adapters

Physical adapters

TCP/IP configuration

Virtual Machines

VM Startup/Shutdown

Agent VM Settings

Default VM Compatibility

Swap File Location

System

Licensing

Host Profile

Time Configuration

Authentication Services

Storage Adapters

+ Add Software Adapter Refresh Rescan Storage... Rescan Adapter Remove

| Adapter | Type | Status | Identifier | Targets | Devices | Paths |
|---------------------------------------|------------|---------|---|---------|---------|-------|
| Model: iSCSI Software Adapter | | | | | | |
| vmhba64 | iSCSI | Online | iscsi_urn(ign.2010-11.com:flexpod:ucs-smbc-a-1) | 8 | 7 | 56 |
| Model: Lewisburg SATA AHCI Controller | | | | | | |
| vmhba0 | Block SCSI | Unknown | - | 0 | 0 | 0 |

Copy All 2 items

Properties Devices **Paths** Dynamic Discovery Static Discovery Network Port Binding Advanced Options

Enable Disable

| Runtime Name | Target | LUN | Status |
|------------------|--|-----|--------------|
| vmhba64 C0:T0:L0 | ign.1992-08.com:netapp:sn.2023c4ee6996f1ec86d8d039ee488168 vs. 3.172.2180.106.3260 | 0 | Active (I/O) |
| vmhba64 C3:T0:L0 | ign.1992-08.com:netapp:sn.2023c4ee6996f1ec86d8d039ee488168 vs. 3.172.2180.107.3260 | 0 | Active |
| vmhba64 C2:T0:L0 | ign.1992-08.com:netapp:sn.2023c4ee6996f1ec86d8d039ee488168 vs. 3.172.2181106.3260 | 0 | Active (I/O) |
| vmhba64 C1:T0:L0 | ign.1992-08.com:netapp:sn.2023c4ee6996f1ec86d8d039ee488168 vs. 3.172.2181107.3260 | 0 | Active |
| vmhba64 C0:T1:L0 | ign.1992-08.com:netapp:sn.b4db01ca5505f1ecbce1d039ee487e72 vs. 3.172.2180.206.3260 | 0 | Active |
| vmhba64 C1:T1:L0 | ign.1992-08.com:netapp:sn.b4db01ca5505f1ecbce1d039ee487e72 vs. 3.172.2180.207.3260 | 0 | Active |
| vmhba64 C2:T1:L0 | ign.1992-08.com:netapp:sn.b4db01ca5505f1ecbce1d039ee487e72 vs. 3.172.2181206.3260 | 0 | Active |
| vmhba64 C3:T1:L0 | ign.1992-08.com:netapp:sn.b4db01ca5505f1ecbce1d039ee487e72 vs. 3.172.2181207.3260 | 0 | Active |

When one of the storage controllers goes down for maintenance or upgrade, the two paths that reach the down controller are no longer available and show up with a path status of `dead` instead.

If a failover of the consistency group occurs on the primary storage cluster, either due to manual failover testing or automatic disaster failover, the secondary storage cluster continues to provide data services for the LUNs in the SM-BC consistency group. Because the LUN identities are preserved and the data has been replicated

synchronously, all ESXi host boot LUNs protected by SM-BC consistency groups remain available from the remote storage cluster.


VMware vMotion and VM/host affinity test

Although a generic FlexPod VMware Datacenter solution supports multi-protocols such as FC, iSCSI, NVMe, and NFS, the FlexPod SM-BC solution feature supports FC and iSCSI SAN protocols typically used for business-critical solutions. This validation only uses iSCSI protocol- based datastores and iSCSI SAN boot.

To allow virtual machines to use storage services from either SM-BC site, the iSCSI datastores from both sites must be mounted by all the hosts in the cluster to enable migration of virtual machines between the two sites and for disaster failover scenarios.

For applications running on the virtual infrastructure that do not require the SM-BC consistency group protection across sites, NFS protocol and NFS datastores can also be used. In that case, caution must be observed when allocating storage for VMs so that the business-critical applications are properly using the SAN datastores protected by SM-BC consistency group to provide business continuity.

The following screenshot shows that hosts are configured to mount iSCSI datastores from both sites.



You have the option of migrating virtual-machine disks between available iSCSI datastores from both sites, as shown in the following figure. For performance considerations, it is optimal to have virtual machines using storage from their local storage cluster to reduce disk I/O latencies. This is especially true when the two sites are located at some distances apart due to the physical round-trip distance latency of roughly 1ms per 100Km distance.

Migrate | iometer-a-01

✓ 1 Select a migration type

2 Select storage

3 Ready to complete

Select storage

Select the destination storage for the virtual machine migration.

VM origin ⓘ

BATCH CONFIGURE

CONFIGURE PER DISK

CONFIGURE

| <input type="checkbox"/> | Virtual Machine | File | Storage | Disk format | VM Storage Policy |
|--------------------------|-----------------|------------------------|----------------------|------------------------|-------------------|
| <input type="checkbox"/> | iometer-a-01 | Configuration File | infra_datastore_a_01 | N/A | Datastore Default |
| <input type="checkbox"/> | iometer-a-01 | Hard disk 1 (64.00 GB) | infra_datastore_a_02 | Same format as sour... | Datastore Default |
| <input type="checkbox"/> | iometer-a-01 | Hard disk 2 (20.00 GB) | infra_datastore_b_01 | Same format as sour... | Datastore Default |
| <input type="checkbox"/> | iometer-a-01 | Hard disk 3 (20.00 GB) | infra_datastore_b_02 | Same format as sour... | Datastore Default |

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT


Tests of vMotion of virtual machines to a different host at the same site as well as across sites were performed and were successful. After manually migrating a virtual machine across sites, the VM/Host affinity rule activates and migrates the virtual machine back to the group where it belongs under the normal condition.

Planned storage failover


Planned storage failover operations should be performed on the solution after initial configuration to determine whether the solution is working properly after storage failover. The testing can help to identify any connectivity or configuration problems which might lead to I/O disruptions. Regularly testing and resolving any connectivity or configuration problems helps to provide uninterrupted data services when a real site disaster occurs. Planned storage failover can also be used before a scheduled storage maintenance activity so that data services can be served from the unaffected site.

To initiate a manual failover of site A storage data services to site B, you can use site B ONTAP System Manager to perform the action.

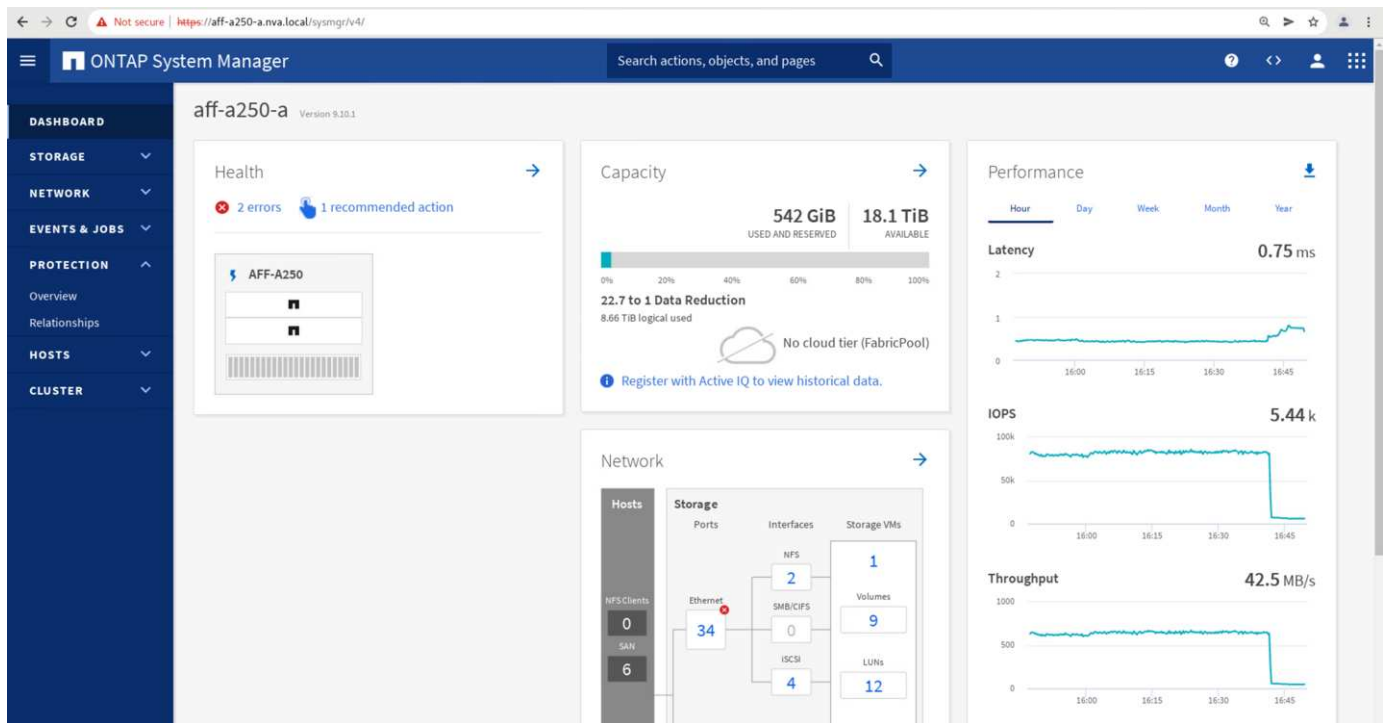
1. Navigate to the Protection > Relationships screen to confirm that the consistency group relationship state is In Sync. If it is still in the Synchronizing state, wait for the state to become In Sync before performing a failover.
2. Expand the dots next to the Source name and click Failover.




3. Confirm failover for the action to start.



Shortly after initiating the failover of the two consistency groups, `cg_esxi_a` and `cg_infra_datastore_a`, on the site B System Manager GUI, the site A I/O serving those two consistency groups moved over to site B. As a result, the I/O at site A reduced significantly as shown in the site A System Manager performance pane.



On the other hand, the Performance pane of the site B System Manager dashboard shows a significant increase in IOPs, due to serving additional I/O moved over from site A, to about 130K IOPs, and reached a throughput of approximately 1GB/s while maintaining an I/O latency of less than 1 millisecond.



With the I/O transparently migrated from site A to site B, the site A storage controllers can now be brought down for scheduled maintenance. After the maintenance work or testing is completed and site A storage cluster is brought back up and operational, check and wait for the consistency group protection state to change back to In sync before performing a failover to return the failover I/O from site B back to site A. Please note that the longer a site is taken down for maintenance or testing, the longer it takes before data are synchronized and the consistency group is returned to the In sync state.

Not secure | https://aff-a250-a.nva.local/symgr/v4/protection/relationships

ONTAP System Manager

Search actions, objects, and pages

Relationships

Protect


Search Download Show/Hide Filter

| Source | Destination | Protection Policy | Relationship Health | State | Lag |
|--|---|-------------------|---------------------|---------|----------|
| ▼ Infra-SVM.1:/cg/cg_infra_datastore_b | Infra-SVM-a:/cg/cg_infra_datastore_b_dest | AutomatedFailOver | Healthy | In sync | 0 second |
| ▼ Infra-SVM.1:/cg/cg_esxi_a_dest | Infra-SVM-a:/cg/cg_esxi_a | AutomatedFailOver | Healthy | In sync | 0 second |
| ▼ Infra-SVM.1:/cg/cg/ | Infra-SVM-a:/cg/cg_infra_datastore_a | AutomatedFailOver | Healthy | In sync | 0 second |
| ▼ Infra-SVM.1:/cg/cg/ | Infra-SVM-a:/cg/cg_esxi_b_dest | AutomatedFailOver | Healthy | In sync | 0 second |

Delete Update Failover

Unplanned storage failover

An unplanned storage failover can occur when a real disaster happens or during a disaster simulation. For example, see the following figure in which the storage system at site A experiences a power outage, an unplanned storage failover is triggered, and the data services for site A LUNs, which are protected by the SM-BC relationships, continue from site B.





To simulate a storage disaster at site A, both storage controllers at site A can be powered off by physically turning off the power switch to discontinue the supply of power to the controllers, or by using the storage controller service processors' system power management command to power off the controllers.

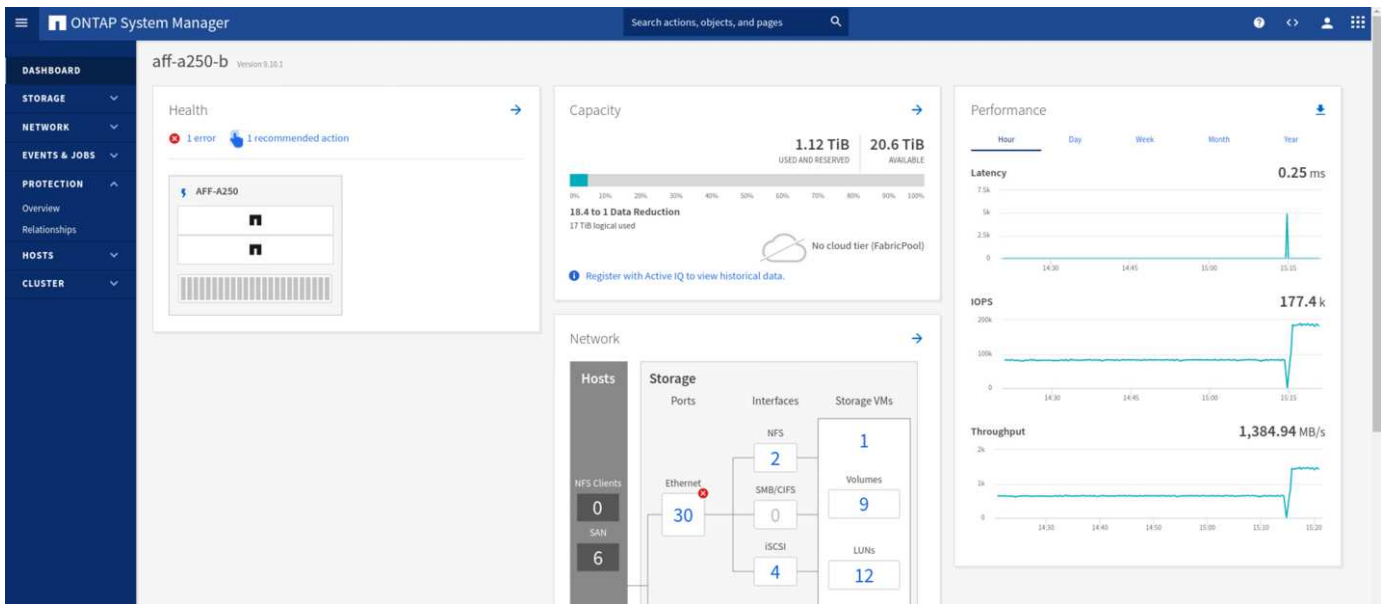
When the storage cluster at site A losses power, there is a sudden stop of the data services provided by the site A storage cluster. Then, the ONTAP Mediator, which monitors the SM-BC solution from a third site, detects the site A storage failure condition and enables the SM-BC solution to perform an automated unplanned failover. This allows site B storage controllers to continue data services for the LUNs configured in the SM-BC consistency group relationships with site A.

From the application perspective, the data services pause briefly while the operating system checks the path status for the LUNs and then resume I/O on the available paths to the surviving site B storage controllers.

During the validation testing, the IOMeter tool on the VMs at both sites generates I/O to their local datastores. After the site A cluster was powered off, I/O paused briefly and then resumed afterwards. See the following two figures for the dashboards of the storage cluster at site A and site B respectively before the disaster which show roughly 80k IOPS and 600 MB/s throughput at each site.



After powering off the storage controllers at site A, we can visually validate that site B storage controller I/O increased sharply to provide additional data services on behalf of site A (see the following figure). In addition, the GUI of the IOMeter VMs also showed that I/O continued despite site A storage cluster outage. Please note that if there are additional datastores backed by LUNs not protected by SM-BC relationships, those datastores will no longer be accessible when the storage disaster occurs. Therefore, it is important to evaluate the business needs of the various application data and properly place them in datastores protected by SM-BC relationships to provide business continuity.



While the site A cluster is down, the relationships of the consistent groups show Out of sync status as shown in the following figure. After the power is turned back on for the storage controllers at site A, the storage cluster boots up and the data synchronization between site A and site B happens automatically.

| Source | Destination | Protection Policy | Relationship Health | State | Lag |
|-------------------------------------|--|-------------------|---------------------|-------------|-----------------------------------|
| infra-SVM.1/cg/cg_esxi_a | infra-SVM-b/cg/cg_esxi_a_dest | AutomatedFailOver | Healthy | Out of sync | 1 hour, 22 minutes and 56 seconds |
| infra-SVM.1/cg/cg_infra_datastore_a | infra-SVM-b/cg/cg_infra_datastore_a_dest | AutomatedFailOver | Healthy | Out of sync | 1 hour, 29 minutes and 35 seconds |

Before returning data services from site B back to site A, you must check site A System Manager and make sure that the SM-BC relationships catches up and the status are back in sync. After confirming that the consistency groups are in sync, a manual failover operation can be initiated to return data services in the consistency group relationships back to site A.

| Source | Destination | Protection Policy | Relationship Health | State | Lag |
|--|--|-------------------|---------------------|---------|----------|
| infra-SVM.1/cg/cg_infra_datastore_b | infra-SVM-a/cg/cg_infra_datastore_b_dest | AutomatedFailOver | Healthy | In sync | 0 second |
| infra-SVM.1/cg/cg_esxi_a_dest | infra-SVM-a/cg/cg_esxi_a | AutomatedFailOver | Healthy | In sync | 0 second |
| infra-SVM.1/cg/cg_infra_datastore_a_dest | infra-SVM-a/cg/cg_infra_datastore_a | AutomatedFailOver | Healthy | In sync | 0 second |
| infra-SVM.1/cg/cg_esxi_b | infra-SVM-a/cg/cg_esxi_b_dest | AutomatedFailOver | Healthy | In sync | 0 second |

Complete site maintenance or site failure

A site might need site maintenance, experience power loss, or might be affected by a natural disaster such as

a hurricane or an earthquake. Therefore, it is crucial that you exercise planned and unplanned site failure scenarios to help ensure that your FlexPod SM-BC solution is properly configured to survive such failures for all your business-critical applications and data services. The following site-related scenarios were validated.

- Planned site maintenance scenario by migrating virtual machines and critical data services to the other site
- Unplanned site outage scenario by powering off servers and storage controllers for disaster simulation

To get a site ready for planned site maintenance, a combination of migrating affected virtual machines off the site with vMotion and a manual failover of the SM-BC consistency group relationships are needed to migrate virtual machines and critical data services to the alternative site. Testing was performed in two different orders: vMotion first followed by SM-BC failover and SM-BC failover first followed by vMotion, to confirm that virtual machines continue to run and data services are not interrupted.

Before performing the planned migration, update the VM/Host affinity rule so the VMs that are currently running on the site are automatically migrated off the site that is undergoing maintenance. The following screenshot shows an example of modifying the site A VM/Host affinity rule for the VMs to migrate from site A to site B automatically. Instead of specifying that the VMs now need to run on site B, you can also choose to disable the affinity rule temporarily so the VMs can be migrated manually.

Edit VM/Host Rule | SMBC ✕

| | | |
|------|-----------------------------|--|
| Name | Site A VMs and hosts | <input checked="" type="checkbox"/> Enable rule. |
| Type | Virtual Machines to Hosts ▼ | |

Description:

Virtual machines that are members of the Cluster VM Group Site A VMs must run on host group Site B hosts.

VM Group:

Site A VMs ▼

Must run on hosts in group ▼

Host Group:

Site B hosts ▼

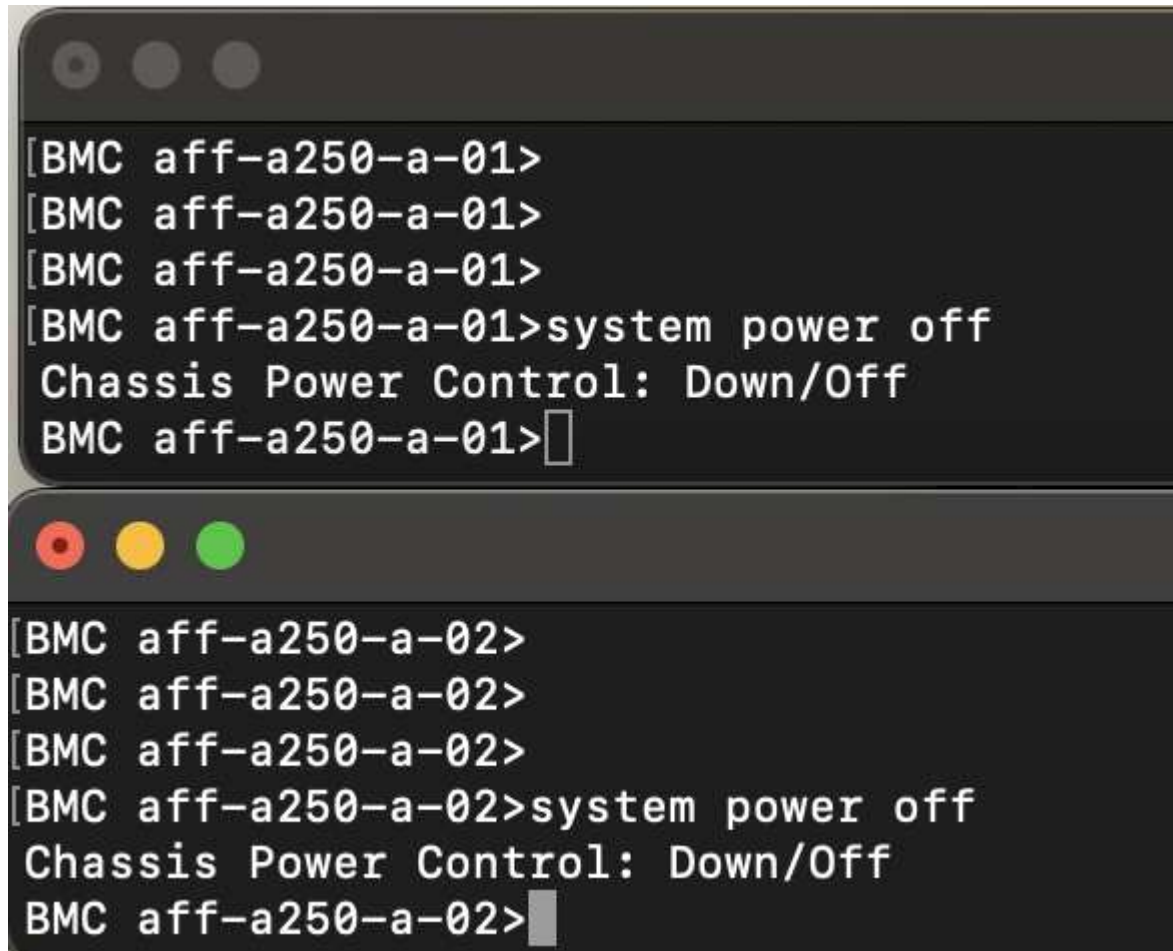
CANCEL OK

After virtual machines and storage services have been migrated, you can power off servers, storage controllers, disk shelves, and switches and perform the needed site maintenance activities. When site maintenance is completed and the FlexPod instance is brought back up, you can change the host group affinity for the VMs to return to their original site. Afterwards, you should change the “Must run on hosts in group” VM/Host site affinity rule back to “Should run on hosts in group” so virtual machines are allowed to run on hosts at the other site should a disaster happens. For the validation testing, all virtual machines were successfully migrated to the other site and the data services continued without problems after performing a failover for the

SM-BC relationships.

For the unplanned site disaster simulation, the servers and storage controllers were powered off to simulate a site disaster. The VMware HA feature detects the downed virtual machines and restarts those virtual machines on the surviving site. In addition, the ONTAP Mediator running at a third site detects the site failure and the surviving site initiates a failover and starts providing data services for the down site as expected.


The following screenshot shows that the storage controllers' service processor CLI were used to power off the site A cluster abruptly to simulate site A storage disaster.



```
[BMC aff-a250-a-01>
[BMC aff-a250-a-01>
[BMC aff-a250-a-01>
[BMC aff-a250-a-01>system power off
Chassis Power Control: Down/Off
BMC aff-a250-a-01>

[BMC aff-a250-a-02>
[BMC aff-a250-a-02>
[BMC aff-a250-a-02>
[BMC aff-a250-a-02>system power off
Chassis Power Control: Down/Off
BMC aff-a250-a-02>
```

The storage clusters' storage virtual machine dashboards as captured by the NetApp Harvest data collection tool and displayed in Grafana dashboard in the NAbbox monitoring tool are shown in the following two screenshots. As can be seen on the right-hand side of the IOPS and Throughputs graphs, the site B cluster picks up the cluster A storage workload right away after site A cluster goes down.



Microsoft SQL Server

Microsoft SQL Server is a widely adopted and deployed database platform for enterprise IT. The Microsoft SQL Server 2019 release brings a lot of new features and enhancements to its relational and analytical engines. It supports workloads with applications running on-premises, in the cloud, and in hybrid could using a combination of the two. In addition, it can be deployed on multiple platforms, including Windows, Linux, and containers.

As part of the business-critical workload validation for the FlexPod SM-BC solution, Microsoft SQL Server 2019 installed on a Windows Server 2022 VM is included along with the IOMeter VMs for SM-BC planned and unplanned storage failover testing. On the Windows Server 2022 VM, SQL Server Management Studio is installed to manage the SQL server. For testing, the HammerDB database tool is used to generate database transactions.

The HammerDB database testing tool was configured for testing with the Microsoft SQL Server TPROC-C workload. For the schema build configurations, the options were updated to use 100 warehouses with 10 virtual users as shown in the following screenshot.


The screenshot shows the 'Microsoft SQL Server TPROC-C Build Options' dialog box. It contains the following fields and options:

- SQL Server:** (local)
- TCP:** ☐
- SQL Server Port:** 1433
- Azure:** ☐
- SQL Server ODBC Driver:** ODBC Driver 17 for SQL Server
- Authentication:** ☒ Windows Authentication, ☐ SQL Server Authentication
- SQL Server User ID:** sa
- SQL Server User Password:** admin
- TPROC-C SQL Server Database:** tpcc
- In-Memory OLTP:** ☐
- In-Memory Hash Bucket Multiplier:** 1
- In-Memory Durability:** ☒ SCHEMA_AND_DATA, ☐ SCHEMA_ONLY
- Number of Warehouses:** 100
- Virtual Users to Build Schema:** 10

At the bottom are 'OK' and 'Cancel' buttons.


After the schema build options were updated, the schema build process was started. A few minutes later, an unplanned simulated site B storage cluster failure was introduced by powering off both nodes of the two node AFF A250 storage cluster at about the same time using system processor CLI commands.

After a brief pause of database transactions, the automated failover for the disaster remediation kicked in and the transactions resumed. The following screenshot shows the HammerDB Transaction Counter screenshot around that time. As the database for the Microsoft SQL Server normally resides on the site B storage cluster, the transaction paused briefly when storage at site B went down and then resumed after the automated failover happened.




The storage cluster metrics were captured by using the NAbbox tool with the NetApp Harvest monitoring tool installed. The results are displayed in the predefined Grafana dashboards for the storage virtual machine and other storage objects. The dashboard provides metrics for latency, throughput, IOPS, and additional details with read and write statistics separated for both site B and site A.

This screenshot shows the NAbbox Grafana performance dashboard for site B storage cluster.



The IOPS for the site B storage cluster was around 100K IOPS before the disaster was introduced. Then, the performance metrics showed a sharp drop down to zero at the right-hand side of the graphs due to the disaster. Since the site B storage cluster was down, nothing could be gathered from the site B cluster after the disaster was introduced.

On the other hand, the IOPS for the site A storage cluster picked up the additional workloads from site B after the automated failover. The additional workload can be easily seen on the right-hand side of the IOPS and Throughput graphs in the following screenshot, which shows the NAbot Grafana performance dashboard for site A storage cluster.



The storage disaster test scenario above confirmed that the Microsoft SQL Server workload can survive a complete storage cluster outage at site B where the database resides. The application transparently used the data services provided by the site A storage cluster after the disaster was detected and the failover happened.

At the compute layer, when the VMs running at a particular site suffers a host failure, the VMs are designed to automatically restart by the VMware HA feature. For a complete site compute outage, the VM/Host affinity rules allow VMs to be restarted at the surviving site. However, for a business-critical application to provide uninterrupted services, an application-based clustering such as Microsoft Failover Cluster or Kubernetes container-based application architecture is required to avoid application downtime. Please see the relevant document for the implementation of the application-based clustering, which is beyond the scope of this technical report.

[Next: Conclusion.](#)

Conclusion

[Previous: Solution validation - Validated scenarios.](#)

The FlexPod Datacenter with SM-BC uses an active-active data center design to provide business continuity and disaster recovery for business-critical workloads. The solution typically interconnects two data centers deployed in separate, geographically dispersed locations in a metro area. The NetApp SM-BC solution uses synchronous replication to protect business-critical data services against a site failure. The solution requires that the two FlexPod deployment sites have a round-trip network latency of less than 10 milliseconds.

The NetApp ONTAP Mediator deployed at a third site monitors the SM-BC solution and enables automated failover when a site disaster is detected. The VMware vCenter with VMware HA and stretched VMware vSphere Metro Storage Cluster configuration work seamlessly with NetApp SM-BC to enable the solution to meet the desired zero RPO and near zero RTO objectives.

The FlexPod SM-BC solution can also be deployed on existing FlexPod infrastructures if they meet the requirements or by adding an additional FlexPod solution to an existing FlexPod to achieve business continuity objectives. Additional management, monitoring, and automation tools, such as Cisco Intersight, Ansible, and HashiCorp Terraform- based automation, are available from NetApp and Cisco so you can easily monitor the solution, gain insights on its operations, and automate its deployment and operations.

From the perspectives of a business-critical application such as Microsoft SQL Server, a database that resides on a VMware datastore protected by an ONTAP SM-BC CG relationship continues to be available despite a site storage outage. As verified during the validation testing, after a power outage of the storage cluster where the database resides, a failover of the SM-BC CG relationship occurs, and the Microsoft SQL Server transactions resume without application disruption.

With application granular data protection, the ONTAP SM-BC CG relationships can be created for your business-critical applications to meet zero RPO and near zero RTO requirements. So that the VMware cluster on which the Microsoft SQL Server application is running can survive a site storage outage, the boot LUNs of the ESXi hosts at each site are also protected by a SM-BC CG relationship.

The flexibility and scalability of FlexPod enables you to start out with a right-sized infrastructure that can grow and evolve as your business requirements change. This validated design enables you to reliably deploy VMware vSphere-based private cloud on a distributed and integrated infrastructure, thereby delivering a solution that is resilient to many single-point-of-failure scenarios as well as a site failure to protect critical business data services.

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

Where to find additional information and version history

Previous: [Conclusion.](#)

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

FlexPod

- 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>

- Cisco Servers - Unified Computing System (UCS)

<https://www.cisco.com/c/en/us/products/servers-unified-computing/index.html>

- NetApp Product Documentation

<https://www.netapp.com/support-and-training/documentation/>

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

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_m6_esxi7u2_design.ht

ml

- FlexPod Datacenter with Cisco UCS 4.2(1) in UCS Managed Mode, VMware vSphere 7.0 U2, and NetApp ONTAP 9.9 Deployment Guide

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_m6_esxi7u2.html

- FlexPod Datacenter with Cisco UCS X-Series, VMware 7.0 U2, and NetApp ONTAP 9.9 Design Guide

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_xseries_esxi7u2_design.html

- FlexPod Datacenter with Cisco UCS X-Series, VMware 7.0 U2, and NetApp ONTAP 9.9 Deployment Guide

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_xseries_vmware_7u2.html

- FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS NVA Design Guide

<https://www.netapp.com/pdf.html?item=/media/22621-nva-1154-DESIGN.pdf>

- FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS NVA Deployment Guide

<https://www.netapp.com/pdf.html?item=/media/21938-nva-1154-DEPLOY.pdf>

- FlexPod MetroCluster IP with VXLAN Multi-Site Frontend Fabric

<https://www.cisco.com/c/dam/en/us/products/collateral/servers-unified-computing/flexpod-metrocluster-ip-vxlan-multi-site-wp.pdf>

- NAbbox

<https://nabox.org>

- NetApp Harvest

<https://github.com/NetApp/harvest/releases>

SM-BC

- SM-BC

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

- TR-4878: SnapMirror Business Continuity (SM-BC) ONTAP 9.8

<https://www.netapp.com/pdf.html?item=/media/21888-tr-4878.pdf>

- How to correctly delete a SnapMirror relationship ONTAP 9

https://kb.netapp.com/Advice_and_Troubleshooting/Data_Protection_and_Security/SnapMirror/How_to_correctly_delete_a_SnapMirror_relationship_ONTAP_9

- SnapMirror Synchronous disaster recovery basics

<https://docs.netapp.com/us-en/ontap/data-protection/snapmirror-synchronous-disaster-recovery-basics-concept.html>

- Asynchronous SnapMirror disaster recovery basics

<https://docs.netapp.com/us-en/ontap/data-protection/snapmirror-disaster-recovery-concept.html#data-protection-relationships>

- Data protection and disaster recovery

<https://docs.netapp.com/us-en/ontap/data-protection-disaster-recovery/index.html>

- Install or upgrade the ONTAP Mediator service

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

VMware vSphere HA and vSphere Metro Storage Cluster

- Creating and Using vSphere HA Clusters

<https://docs.vmware.com/en/VMware-vSphere/7.0/com.vmware.vsphere.avail.doc/GUID-5432CA24-14F1-44E3-87FB-61D937831CF6.html>

- VMware vSphere Metro Storage Cluster (vMSC)

<https://core.vmware.com/resource/vmware-vsphere-metro-storage-cluster-vmssc>

- VMware vSphere Metro Storage Cluster Recommended Practices

<https://core.vmware.com/resource/vmware-vsphere-metro-storage-cluster-recommended-practices>

- NetApp ONTAP with NetApp SnapMirror Business Continuity (SM-BC) with VMware vSphere Metro Storage Cluster (vMSC). (83370)

<https://kb.vmware.com/s/article/83370>

- Protect tier-1 applications and databases with VMware vSphere Metro Storage Cluster and ONTAP

<https://community.netapp.com/t5/Tech-ONTAP-Blogs/Protect-tier-1-applications-and-databases-with-VMware-vSphere-Metro-Storage/ba-p/171636>

Microsoft SQL and HammerDB

- Microsoft SQL Server 2019

<https://www.microsoft.com/en-us/sql-server/sql-server-2019>

- Architecting Microsoft SQL Server on VMware vSphere Best Practices Guide

<https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/solutions/sql-server-on-vmware-best-practices-guide.pdf>

- HammerDB website

<https://www.hammerdb.com>

Compatibility Matrix

- Cisco UCS Hardware Compatibility Matrix

<https://ucshcltool.cloudapps.cisco.com/public/>

- NetApp Interoperability Matrix Tool

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

- NetApp Hardware Universe

<https://hwu.netapp.com>

- VMware Compatibility Guide

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

Version history

| Version | Date | Document version history |
|-------------|------------|--------------------------|
| Version 1.0 | April 2022 | Initial release. |

FlexPod Datacenter with VMware vSphere 7.0, Cisco VXLAN Single-Site Fabric, and NetApp ONTAP 9.7 - Design

Ramesh Isaac, Cisco
Abhinav Singh, NetApp

Cisco Validated Designs (CVDs) consist of 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 customers. Cisco and NetApp have partnered to deliver FlexPod, which serves as the foundation for a variety of workloads, and deliver architectural designs that are robust, efficient, and scalable to address customer requirements. A FlexPod solution is a validated approach for deploying Cisco and NetApp technologies and products to build shared private and public cloud infrastructure.

[FlexPod Datacenter with VMware vSphere 7.0, Cisco VXLAN Single-Site Fabric, and NetApp ONTAP 9.7 - Design](#)

FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7 - Deployment

John George, Cisco
Sree Lakshmi Lanka, NetApp

This document describes the Cisco and NetApp FlexPod Datacenter with NetApp ONTAP

9.7 on NetApp AFF A400 All Flash storage system, Cisco UCS Manager unified software release 4.1(2) with second-generation Intel Xeon Scalable Processors and VMware vSphere 7.0. Cisco UCS Manager (UCSM) 4.1(2) provides consolidated support of the following:

- All current Cisco UCS Fabric Interconnect models: 6200, 6300, 6324 (Cisco UCS Mini)
- 6400
- 2200/2300/2400 series IOM
- Cisco UCS B-Series
- Cisco UCS C-Series

Also included are Cisco Intersight and NetApp Active IQ SaaS management platforms.

FlexPod Datacenter with NetApp ONTAP 9.7, Cisco UCS unified software release 4.1(2), and VMware vSphere 7.0 comprise a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus 9000 family of switches, MDS 9000 multilayer fabric switches, and NetApp AFF A-Series storage arrays running ONTAP 9.7 data management software.

[FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7 - Deployment](#)

FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Design

John George, Cisco
Scott Kovacs, NetApp

This document describes the Cisco and NetApp FlexPod solution, which is a validated approach for deploying Cisco and NetApp technologies as shared cloud infrastructure. This validated design provides a framework for deploying VMware vSphere, the most popular virtualization platform in enterprise-class data centers, on FlexPod.

[FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Design](#)

FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Deployment

John George, Cisco
Scott Kovacs, NetApp

The current industry trend in data center design is towards shared infrastructures. By using virtualization along with pre-validated IT platforms, enterprise customers have embarked on the journey to the cloud by moving away from application silos and toward shared infrastructure that can be quickly deployed, thereby increasing agility and reducing costs. Cisco and NetApp have partnered to deliver FlexPod, which uses best of breed storage, server, and network components to serve as the foundation for a variety of workloads, enabling efficient architectural designs that can be quickly and confidently deployed.

FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Design

John George, Cisco
Scott Kovacs, NetApp

This document describes a validated solution for deploying Cisco and NetApp technologies as a shared cloud infrastructure. This validated design provides a framework for deploying VMware vSphere, the most popular virtualization platform in enterprise-class data centers, on FlexPod.

FlexPod is a leading integrated infrastructure supporting a broad range of enterprise workloads and use cases. This solution enables customers to quickly and reliably deploy a VMware vSphere-based private cloud on an integrated infrastructure.

[FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Design](#)

FlexPod Datacenter with VMware vSphere 6.7 U2, Cisco UCS fourth-generation Fabric and NetApp ONTAP 9.6

John George, Cisco
Sree Lakshmi Lanka, NetApp

This document describes the Cisco and NetApp FlexPod Datacenter with NetApp ONTAP 9.6, Cisco UCS Manager unified software release 4.0(4) with second-generation Intel Xeon Scalable Processors, and VMware vSphere 6.7 U2. Cisco UCS Manager (UCSM) 4.0(4) provides consolidated support of the following:

- All current Cisco UCS Fabric Interconnect models: 6200, 6300, 6324 (Cisco UCS Mini)
- 6454
- 2200/2300/2400 series IOM
- Cisco UCS B-Series
- Cisco UCS C-Series.

FlexPod Datacenter with NetApp ONTAP 9.6, Cisco UCS unified software release 4.0(4), and VMware vSphere 6.7 U2 is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus 9000 family of switches, MDS 9000 multilayer fabric switches, and NetApp AFF A-Series storage arrays running ONTAP 9.

[FlexPod Datacenter with VMware vSphere 6.7 U2, Cisco UCS fourth-generation Fabric and NetApp ONTAP 9.6](#)

FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS fourth-generation fabric, and NetApp AFF A-Series - Design

John George, Cisco
Sree Lakshmi Lanka, NetApp

This document describes the Cisco and NetApp FlexPod solution, which is a validated approach for deploying Cisco and NetApp technologies as shared cloud infrastructure. This validated design provides a framework for deploying VMware vSphere, the most popular virtualization platform in enterprise class data centers, on FlexPod.

FlexPod is a leading integrated infrastructure supporting a broad range of enterprise workloads and use cases. This solution enables customers to quickly and reliably deploy VMware vSphere-based private cloud on integrated infrastructure.

The recommended solution architecture is built on Cisco Unified Computing System (Cisco UCS) using the unified software release to support the Cisco UCS hardware platforms, including Cisco UCS B-Series blade and C-Series rack servers, Cisco UCS 6454 Fabric Interconnects, Cisco Nexus 9000 Series switches, Cisco MDS Fibre channel switches, and NetApp All Flash series storage arrays. Additionally, it includes VMware vSphere 6.7 Update 1, which provides a number of new features for optimizing storage utilization and facilitating a private cloud.

[FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS fourth-generation fabric, and NetApp AFF A-Series - Design](#)

FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS fourth-generation fabric, and NetApp AFF A-Series

John George, Cisco
Scott Kovacs, NetApp

This document describes the Cisco and NetApp FlexPod Datacenter with Cisco UCS Manager unified software release 4.0(2) and VMware vSphere 6.7 U1. Cisco UCS Manager (UCSM) 4.0(2) provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 (Cisco UCS Mini)), 6454, 2200/2300 series IOM, Cisco UCS B-Series, and Cisco UCS C-Series. FlexPod Datacenter with Cisco UCS unified software release 4.0(2) and VMware vSphere 6.7 U1 is a predesigned, best-practice data center architecture built on the Cisco Unified Computing System (UCS), the Cisco Nexus 9000 family of switches, MDS 9000 multilayer fabric switches, and NetApp AFF A-Series storage arrays running ONTAP 9 storage OS.

[FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS fourth-generation fabric, and NetApp AFF A-Series](#)

FlexPod Datacenter with Cisco ACI Multi-Pod, NetApp MetroCluster IP, and VMware vSphere 6.7 - Design

Haseeb Niazi, Cisco

Arvind Ramakrishnan, NetApp

This document describes the integration of the Cisco ACI Multi-Pod and NetApp MetroCluster IP solution into the FlexPod Datacenter to provide a highly available multi-data center solution. The multi-data center architecture offers the ability to balance workloads between two data centers utilizing non-disruptive workload mobility, thereby enabling the migration of services between sites without the need for sustaining an outage.

The FlexPod with ACI Multi-Pod and NetApp MetroCluster IP solution provides the following benefits:

- Seamless workload mobility across data centers
- Consistent policies across the sites
- Layer-2 extension across geographically dispersed data centers
- Enhanced downtime avoidance during maintenance
- Disaster avoidance and recovery

[FlexPod Datacenter with Cisco ACI Multi-Pod, NetApp MetroCluster IP, and VMware vSphere 6.7 - Design](#)

FlexPod Datacenter with Cisco ACI Multi-Pod with NetApp MetroCluster IP and VMware vSphere 6.7 - Deployment

Haseeb Niazi, Cisco

Ramesh Issac, Cisco

Arvind Ramakrishnan, NetApp

Cisco and NetApp have partnered to deliver a series of FlexPod solutions that enable strategic data center platforms. The FlexPod solution delivers an integrated architecture that incorporates best design practices for computing, storage, and networking, thereby minimizing IT risks by validating the integrated architecture to ensure compatibility between various components. The solution also addresses IT pain points by providing documented design guidance, deployment guidance, and support that can be used in various stages (planning, designing and implementation) of a deployment.

[FlexPod Datacenter with Cisco ACI Multi-Pod with NetApp MetroCluster IP and VMware vSphere 6.7 - Deployment](#)

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

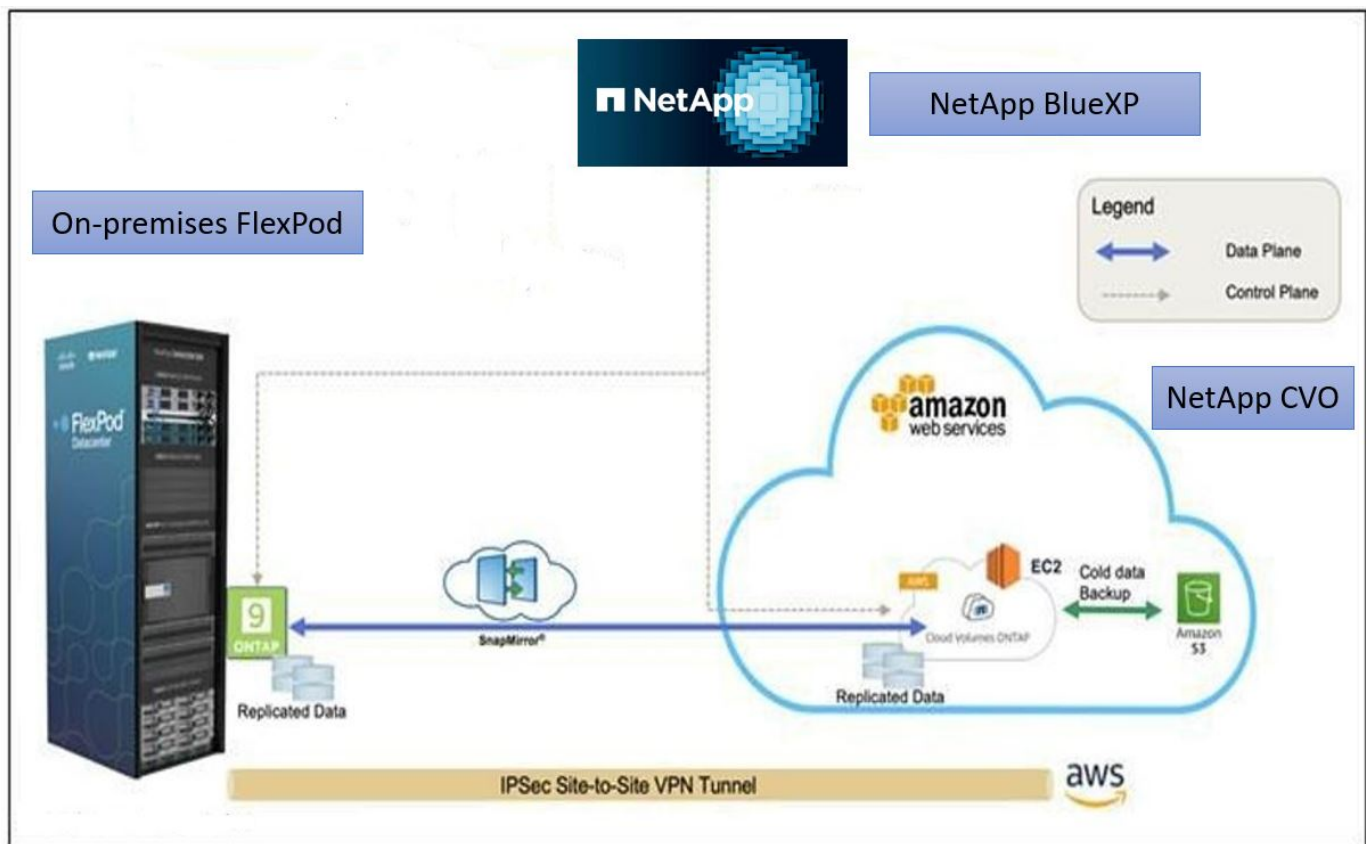
FlexPod Datacenter integrated with NetApp Cloud Volumes ONTAP offers the following benefits to healthcare workloads:

- **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 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.

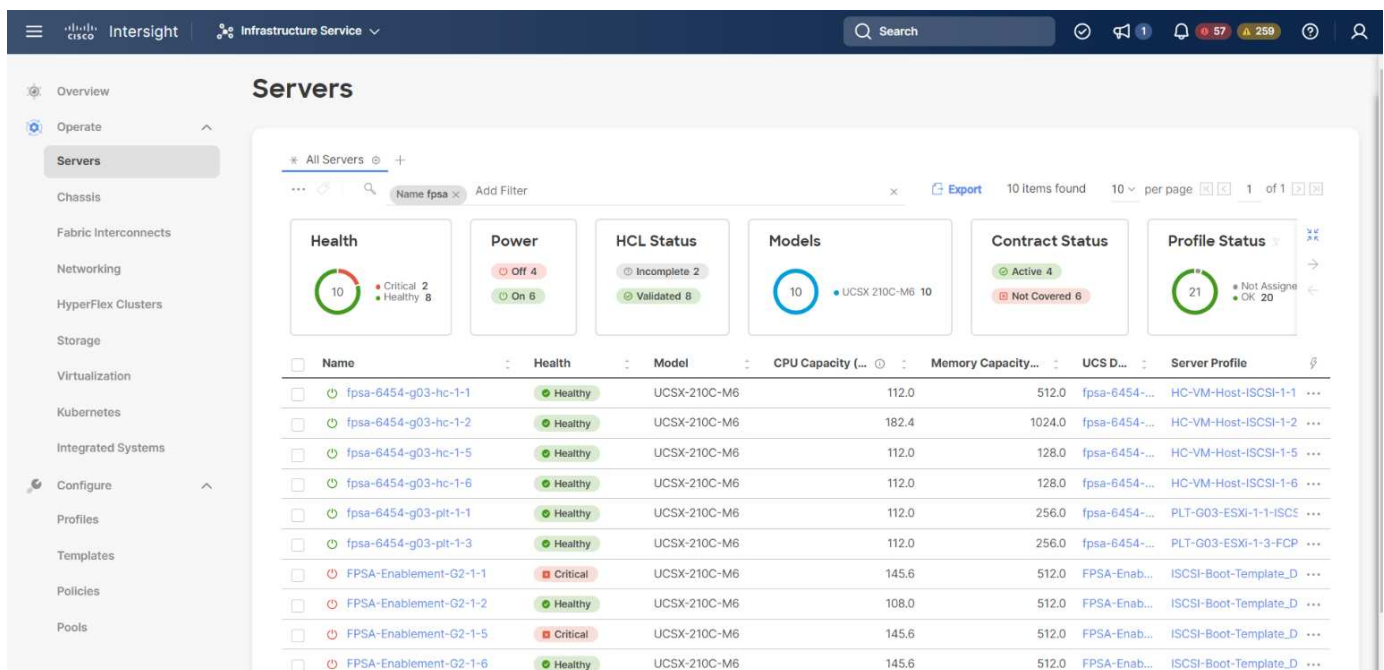
For more information, see [Active IQ Unified Manager](#).

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.

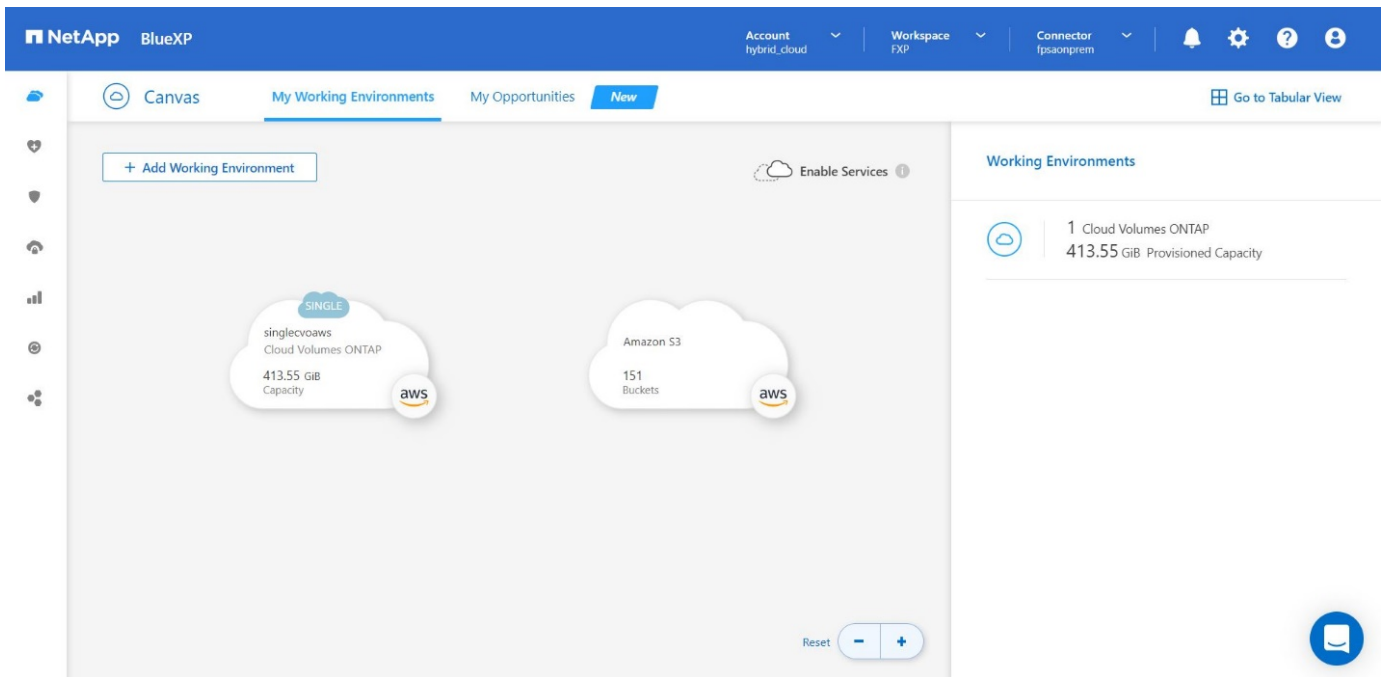
Before adding a Cloud Volume ONTAP working environment, you must deploy Connector. BlueXP prompts you if you try to create your first Cloud Volumes ONTAP working environment without Connector in place. To deploy Connector in AWS from BlueXP, see this [link](#).

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`.

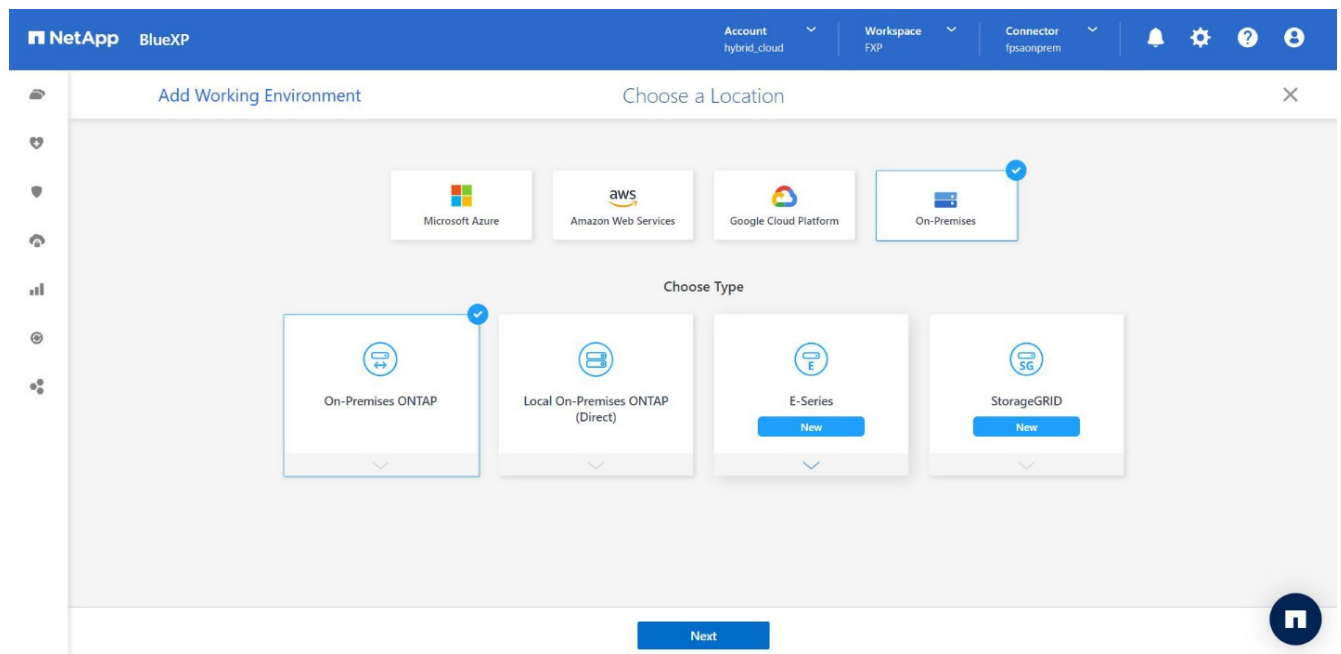
6. Map the LUN to the igroup.

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

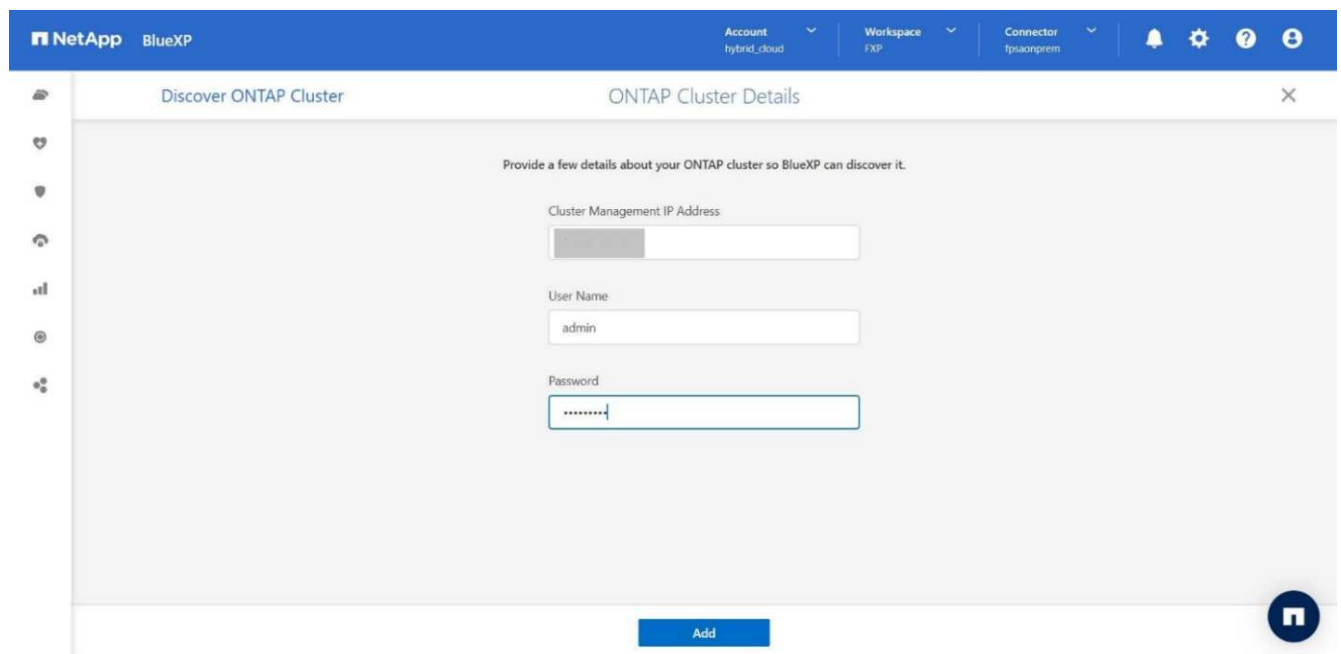
Add on-premises FlexPod storage to BlueXP

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**.

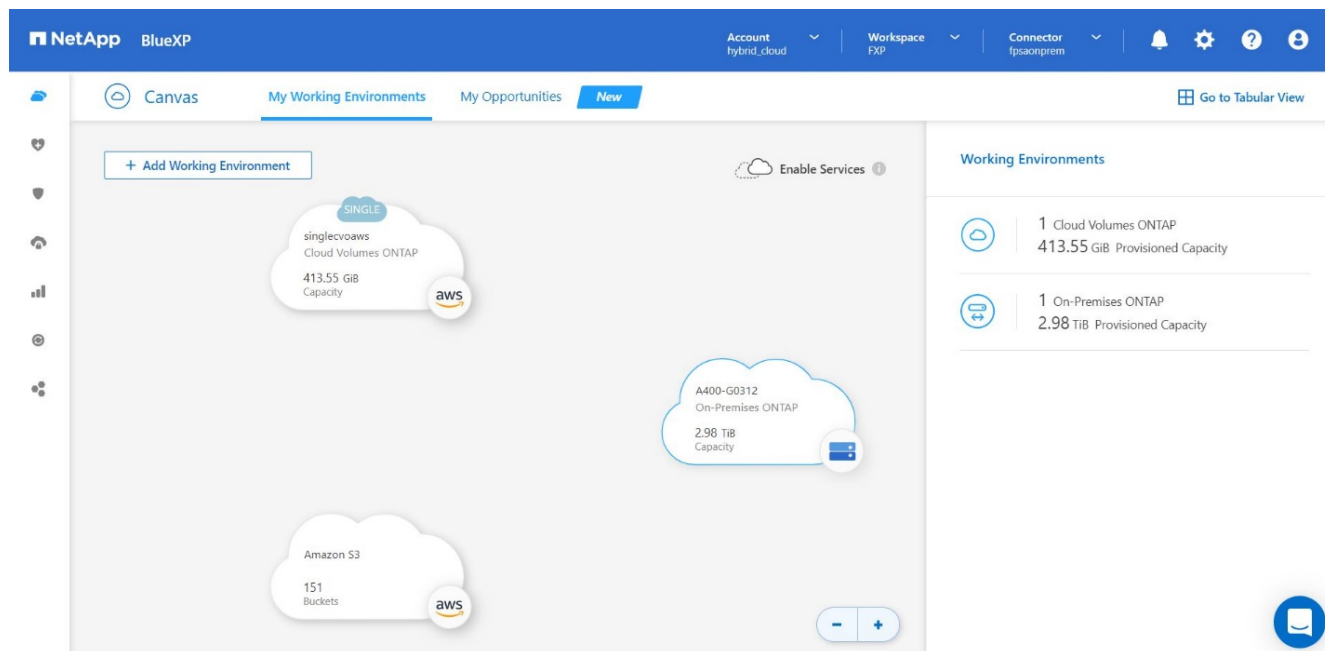


4. On the ONTAP Cluster Details page, enter the cluster management IP address and the password for the admin user account. Then click **Add**.



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 `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/data1v
meta-data=/dev/datavg/data1v      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/data1v /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-data1v | 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**.

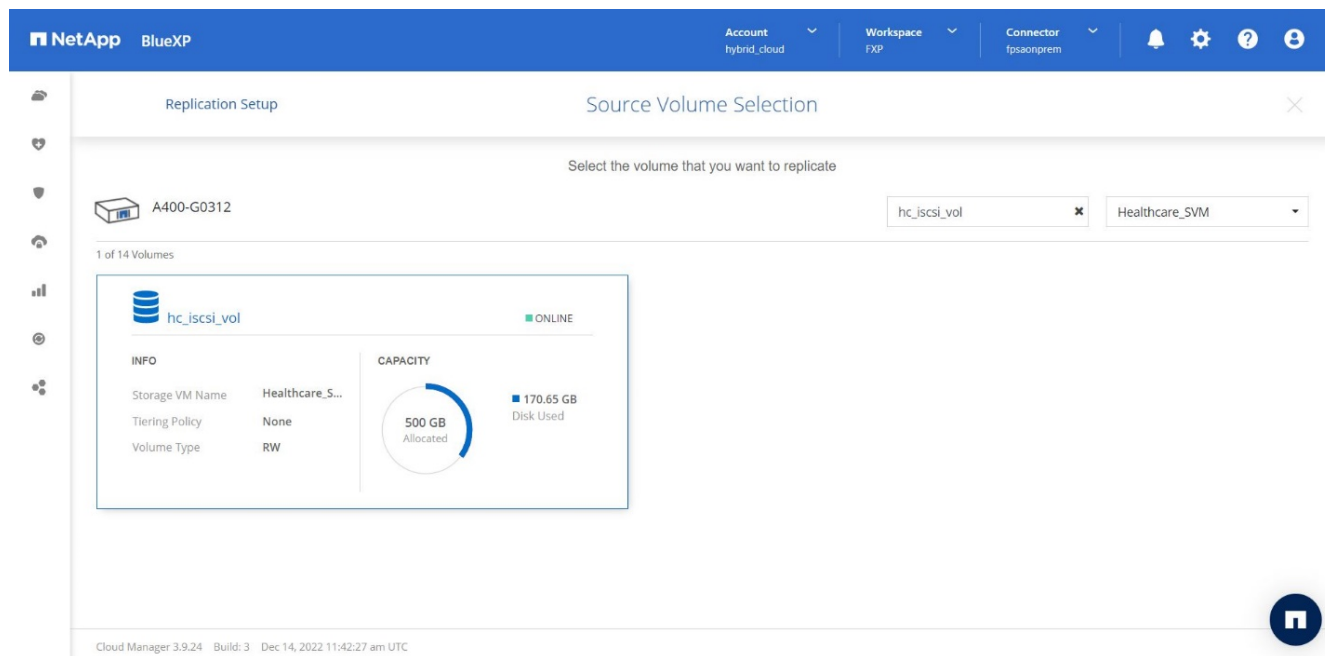


The remaining steps explain how to create a synchronous relationship between Cloud Volumes ONTAP and on-prem ONTAP clusters.


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)

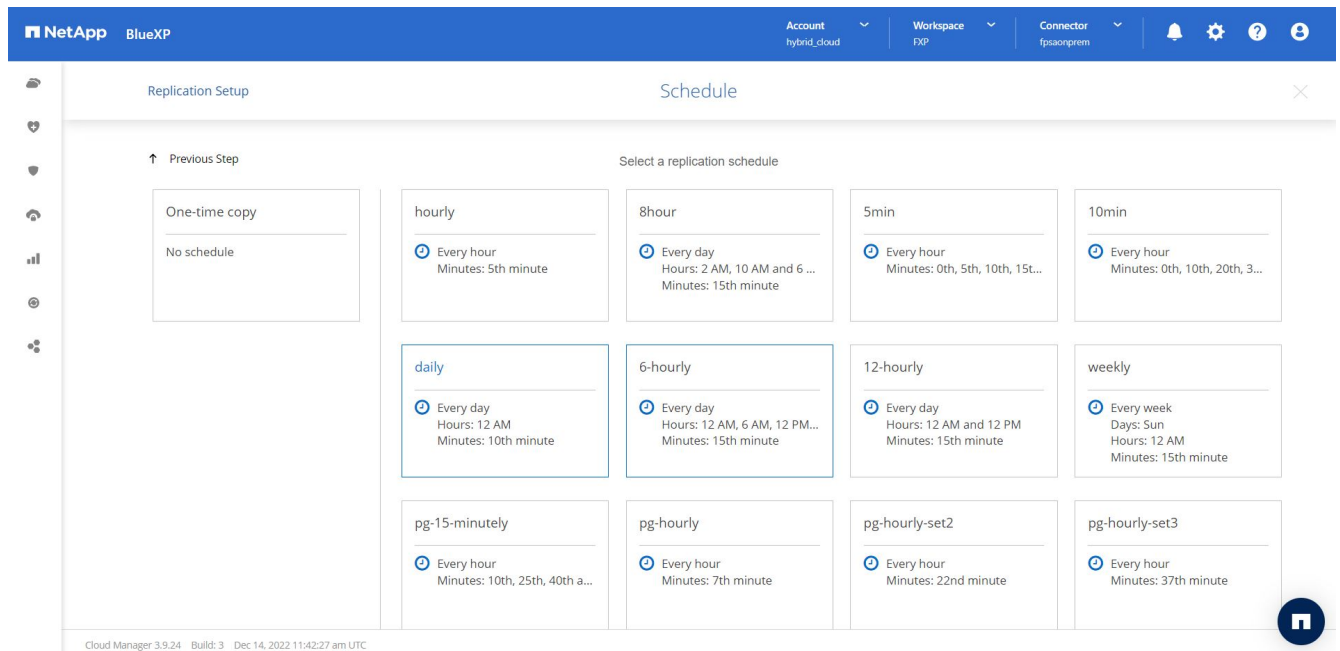
Continue

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**.

NetApp BlueXP Account hybrid_cloud Workspace FXP Connector fpxaonprem

Replication Setup Review & Approve

Previous Step

Source

A400-G0312

hc_iscsi_vol

Destination

singlecvoaws

hc_iscsi_vol_copy

Review your selection and start the replication process

☒ I understand that BlueXP will allocate the appropriate AWS resources to comply with my above requirements. [More information >](#)

| | | | |
|------------------------------------|---------------------------|-------------------------|---------------------------|
| Source Volume Allocated Size: | 500 GB | Destination Aggregate: | aggr3 (Automatically s... |
| Source Volume Used Size: | 170.65 GB | Destination Storage VM: | svm_singlecvoaws |
| Source Thin Provisioning: | Yes | Max Transfer Rate: | 100 MB/s |
| Destination Volume Allocated Size: | 500 GB | SnapMirror Policy: | Mirror |
| Destination Volume Disk Type: | General Purpose SSD (...) | Replication Schedule: | daily |
| Destination Thin Provisioning: | Yes | | |

Go

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

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.

NetApp BlueXP Account hybrid_cloud Workspace FXP Connector fpxaonprem

Canvas My Working Environments My Opportunities New Go to Tabular View

+ Add Working Environment

Enable Services

SINGLE

singlecvoaws

Cloud Volumes ONTAP

513.55 GiB Capacity

aws

Replication

A400-G0312

On-Premises ONTAP

3.08 TiB Capacity

Amazon S3

151 Buckets

Working Environments

- 1 Cloud Volumes ONTAP
513.55 GiB Provisioned Capacity
- 1 On-Premises ONTAP
3.08 TiB Provisioned Capacity

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 and workspace settings. 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 this, there are three sections: 'Transfer Info', 'Last Transfer Info', and 'Volume Info'. The 'Transfer Info' section contains a table with details like Status, Type, Total Size, Lag Duration, Priority, Max Transfer Rate, Total Transfer Time, Mirror State, Used Size / Used on Cloud, and Network Compression Ratio. The 'Last Transfer Info' section shows the last successful transfer date, size, duration, and type. The 'Volume Info' section lists source and destination availability zones and SVM names.

| 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 | snapirored | 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 | 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.

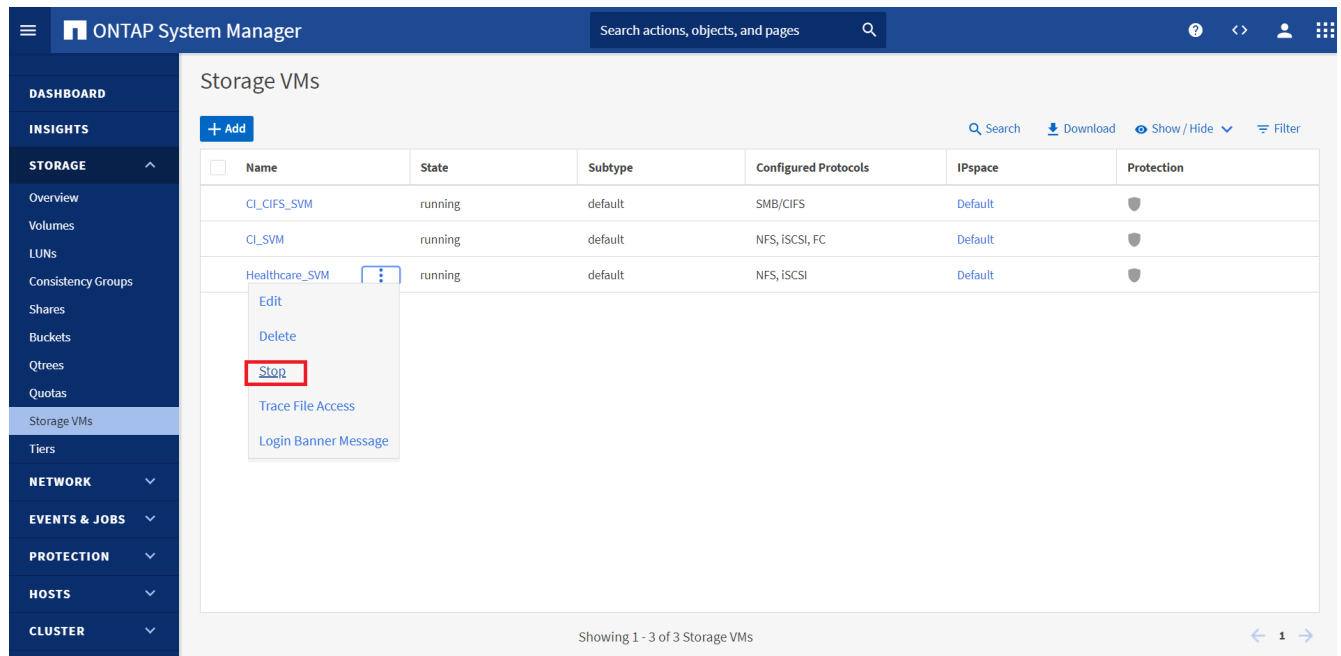
Disaster recovery

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.

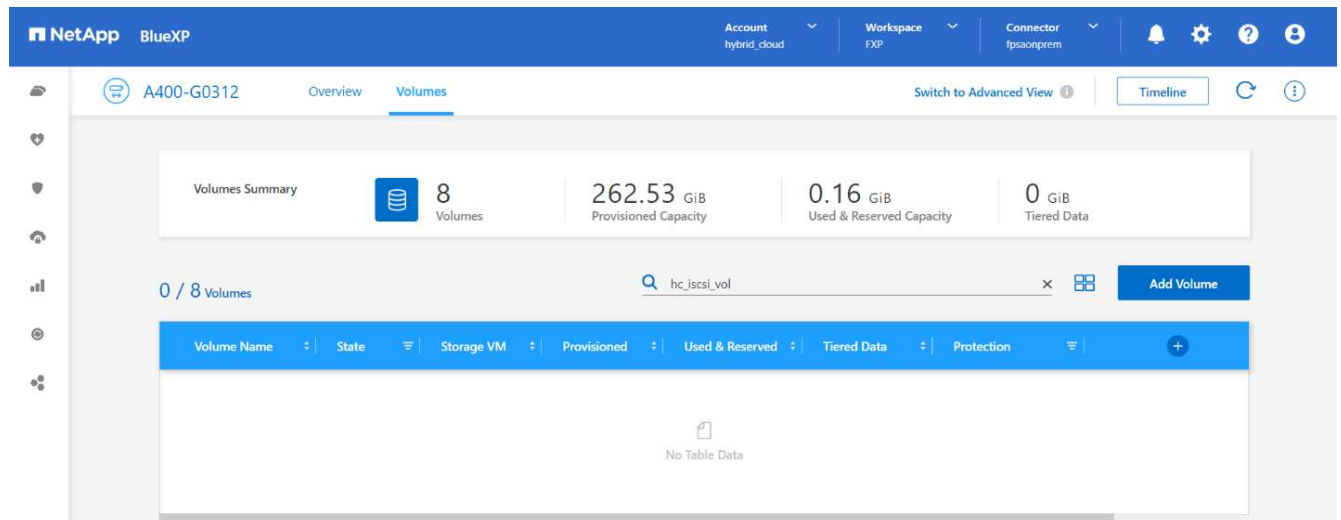
To perform the validation of a successful DR scenario, complete the following steps:

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 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_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.

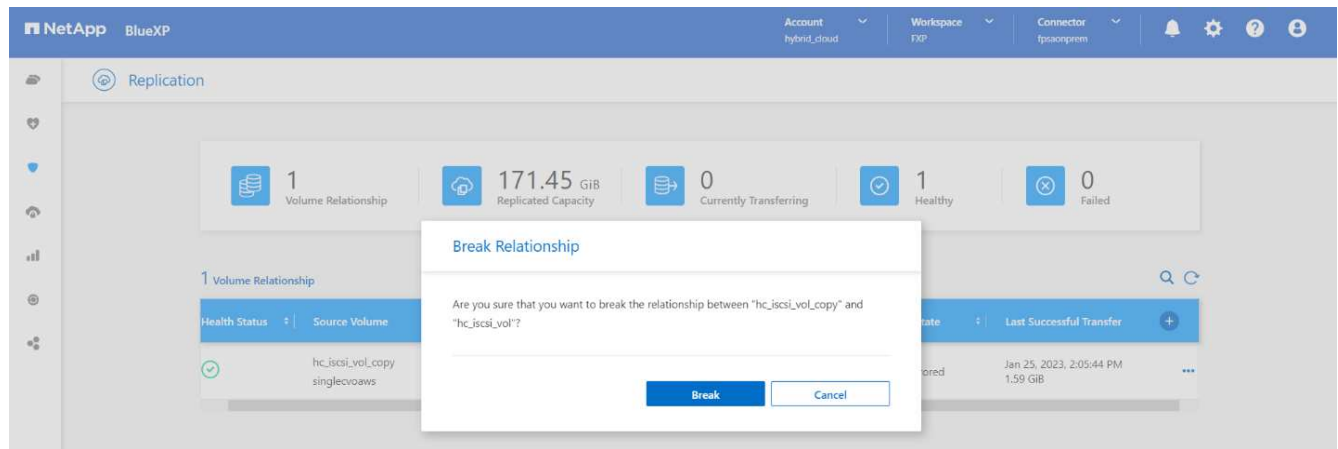
In this section, we talk about bringing up the source site, restoring the SnapMirror relationship between on-premises ONTAP and Cloud Volumes ONTAP, and finally performed a data integrity check on the source end.

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).



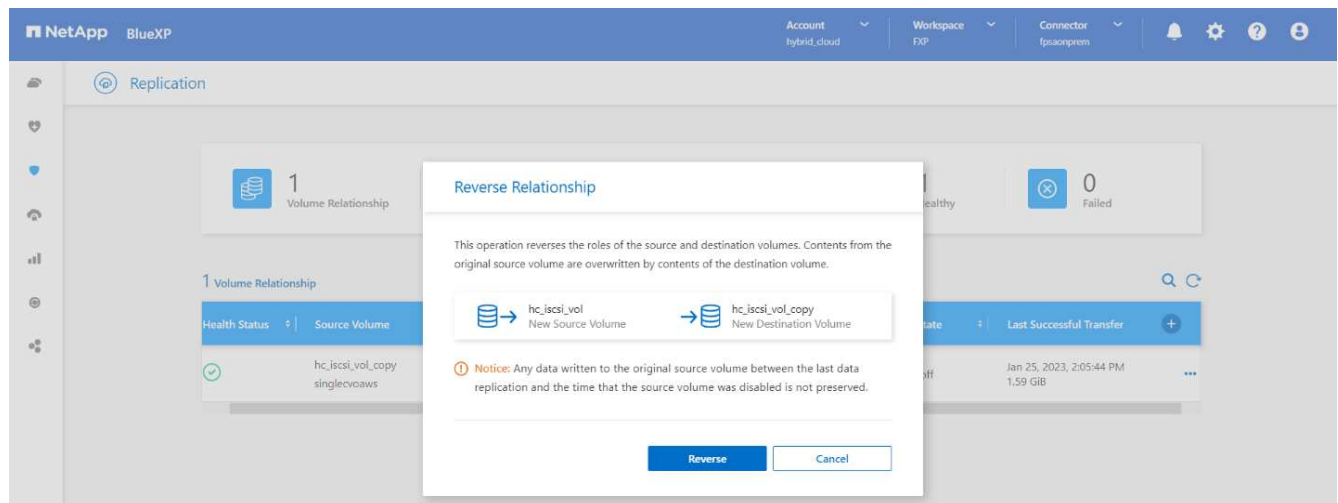
- 2. Break the SnapMirror replication relationship between Cloud Volumes ONTAP and on-premises ONTAP and promote the on-premises volume (hc_iscsi_vol) back to production.



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

- 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.

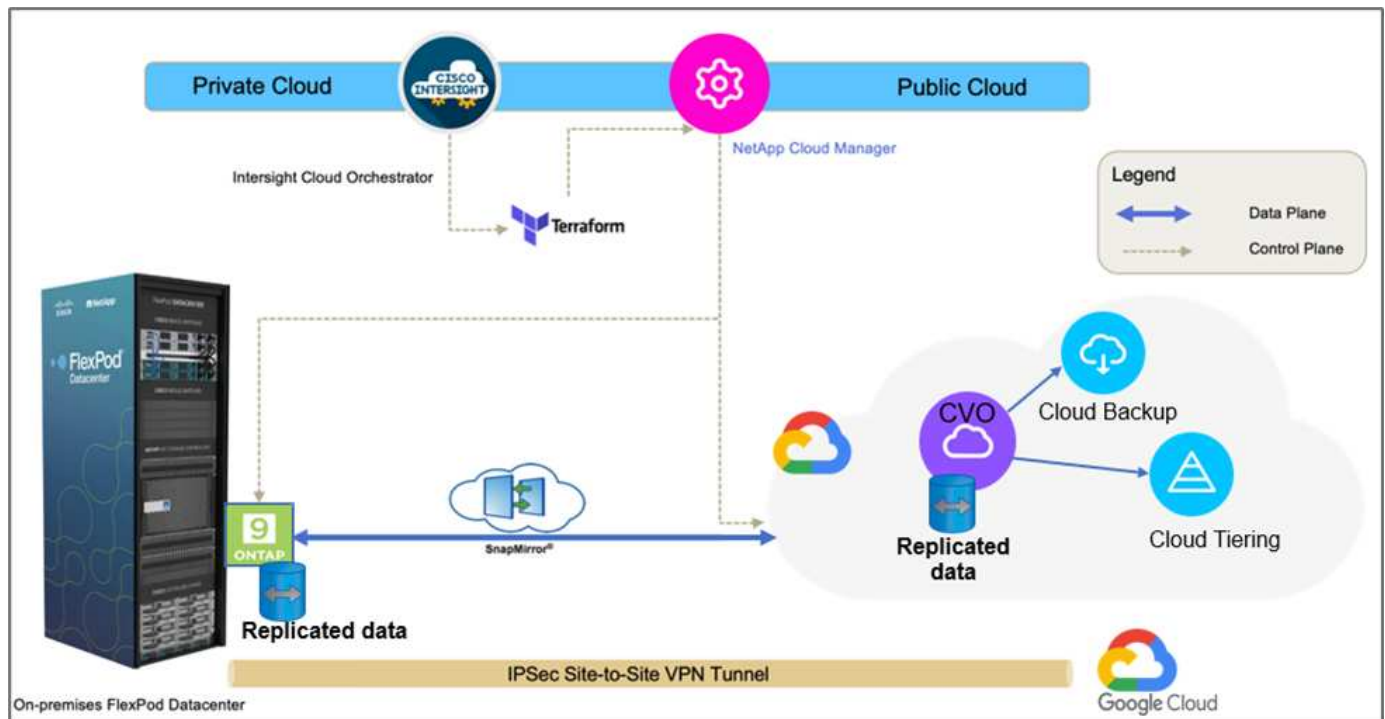
- **Increased IT productivity.** Using Intersight as the single secure, enterprise-grade platform for infrastructure and application lifecycle management simplifies configuration management and automation of manual tasks at scale for the solution.

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 of 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 |
| | Active IQ Unified Manager | 9.11 |
| Software | vSphere ESXi | 7.0(U3) |
| | VMware vCenter Appliance | 7.0.3 |

| Component | Product | Version |
|-----------|---|------------|
| | 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

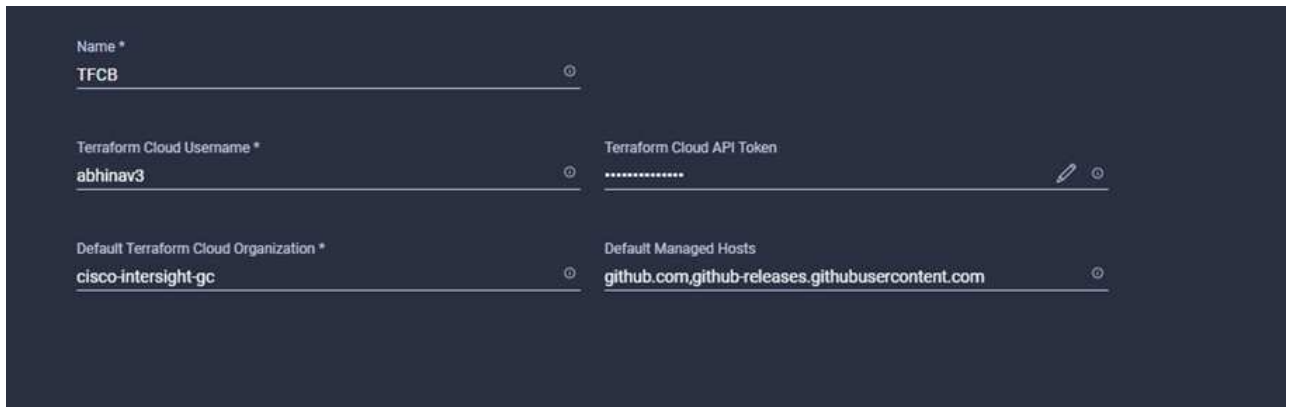
1. Log into Intersight with Account Administrator, Device Administrator, or Device Technician privileges.
2. Navigate to **ADMIN > Targets > Claim a New Target**.
3. In **Categories**, click **Cloud**.
4. Click **Terraform Cloud** and click **Start**.



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
 - github-releases.githubusercontent.com



The screenshot shows a configuration form for Terraform Cloud. It has four fields arranged in two rows. The first row has 'Name *' with the value 'TFCB' and 'Terraform Cloud API Token' with a masked value '*****'. The second row has 'Default Terraform Cloud Organization *' with the value 'cisco-intersight-gc' and 'Default Managed Hosts' with the value 'github.com,github-releases.githubusercontent.com'. Each field has a small circular icon to its right.

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.

The screenshot shows the 'Terraform Cloud target' configuration page. It includes the following fields:

- Name ***: flexpod-solution-terraform-agent
- Intersight Assist ***: g13-intersight-appliance.fpmc.sa
- Terraform Cloud Organization ***: cisco-intersight-gc
- Terraform Cloud Agent Pool Name ***: flexpod-solution-agent-pool

Below these fields is a section titled **Managed Hosts** containing a table:

| Hostname / IP Address / Subnets * | |
|---------------------------------------|--------------------------|
| github.com | [Info icon] [Trash icon] |
| github-releases.githubusercontent.com | [Info icon] [Trash icon] |

A plus sign (+) is located to the right of the table, indicating an option to add more hosts.



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.

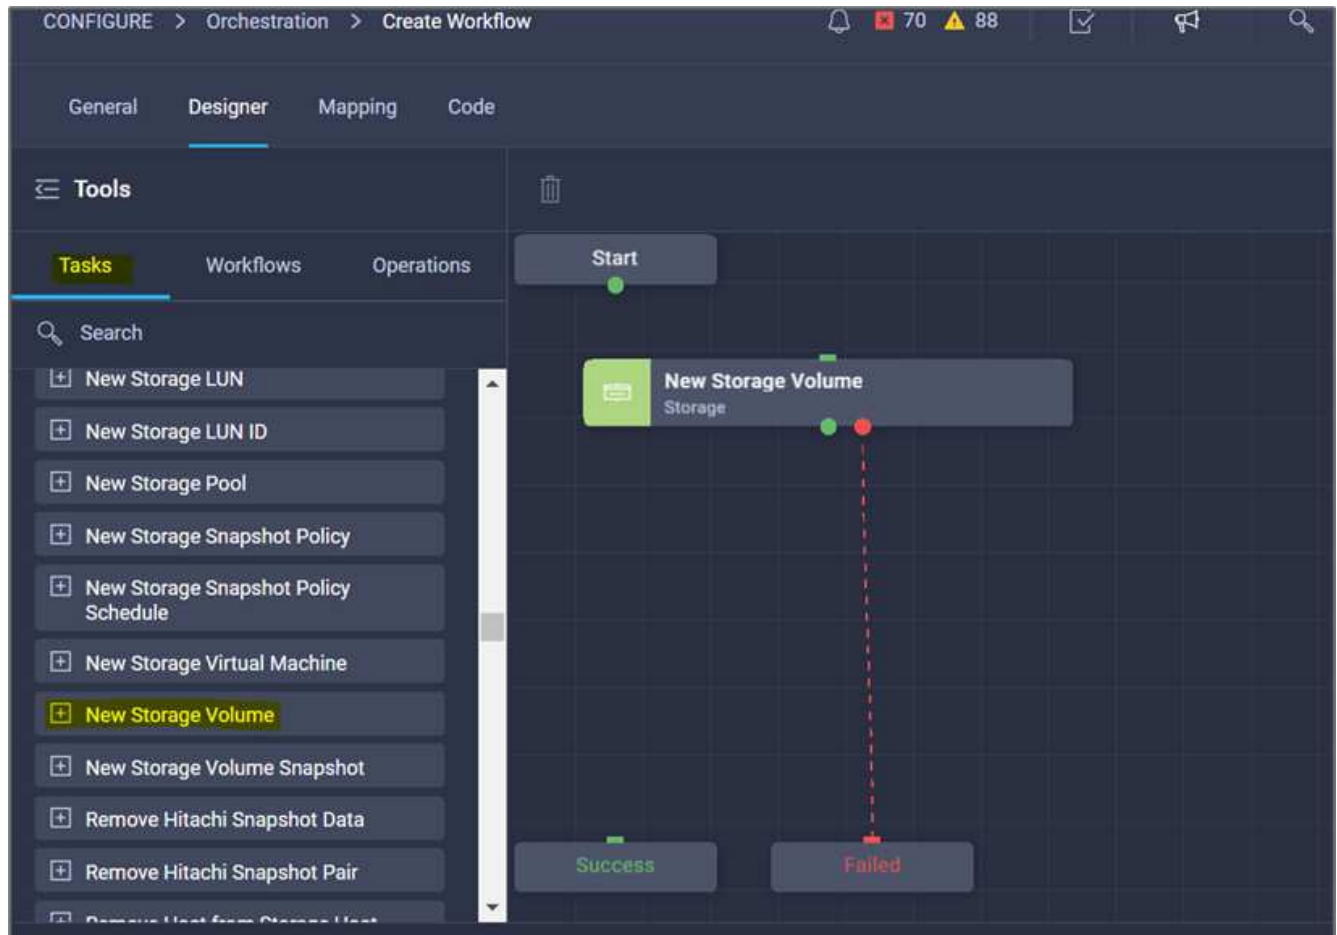
Procedure 1: Create the workflow

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.

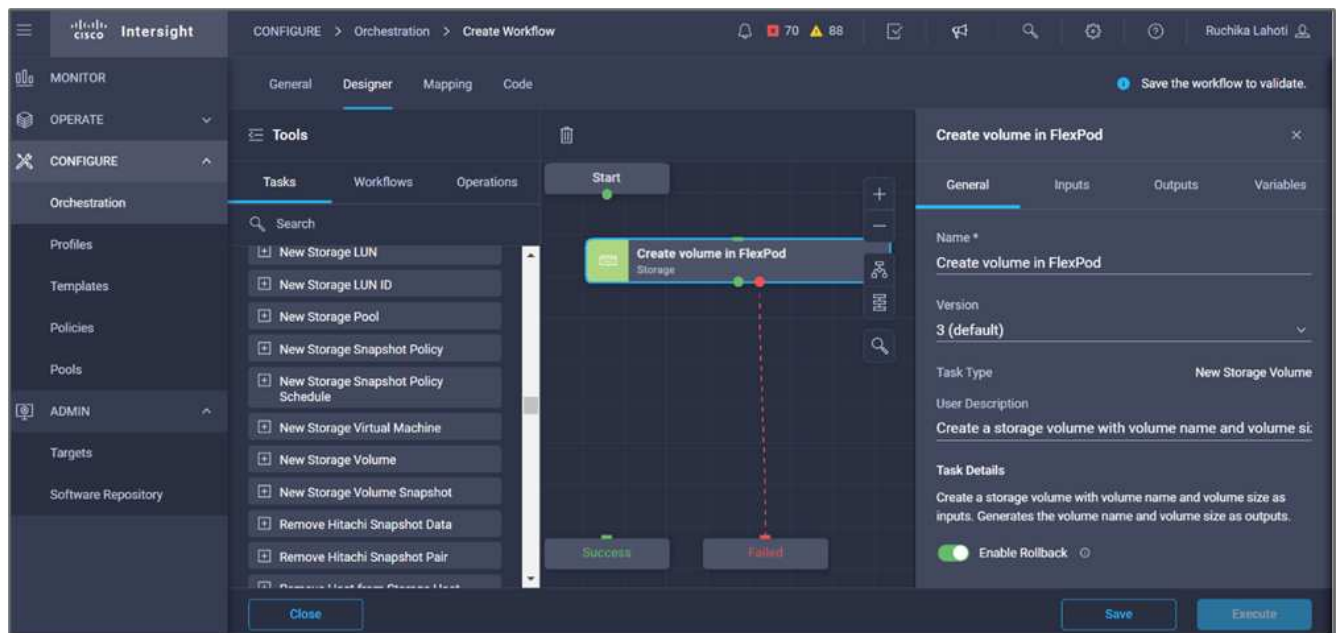
The screenshot shows the 'General' tab of a workflow creation interface. The 'Display Name' is 'Disaster Recovery Workflow' and the 'Reference Name' is 'DisasterRecoveryWorkflow'. The 'Organization' is set to 'default' and the 'Version' is '2 (default)'. The 'Description' is 'Workflow which creates and configures SnapMirror between FlexPod Storage and Cloud Volumes ONTAP'. Under 'Workflow Execution', there are three options: 'Failed/Terminated Actions' (checked), 'Enable Retry' (checked), and 'Enable Auto Rollback' (unchecked). There is also an 'Enable Debug Logs' option. At the bottom, there are tabs for 'Workflow Inputs', 'Workflow Variables', and 'Workflow Outputs', with 'Workflow Inputs' being the active tab. An 'Add Workflow Input' button is visible at the bottom left.

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**.




4. 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**.




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**.

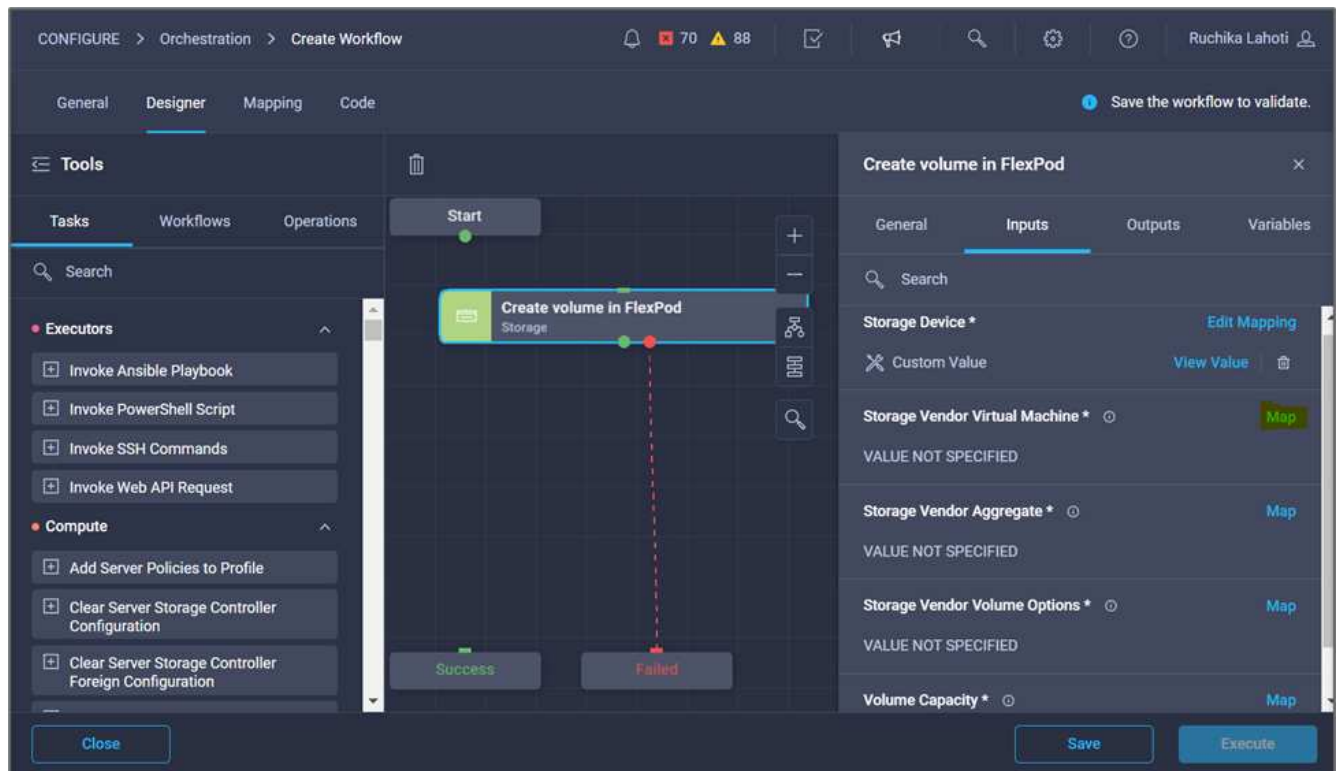
8. Click the storage target added and click **Select**.



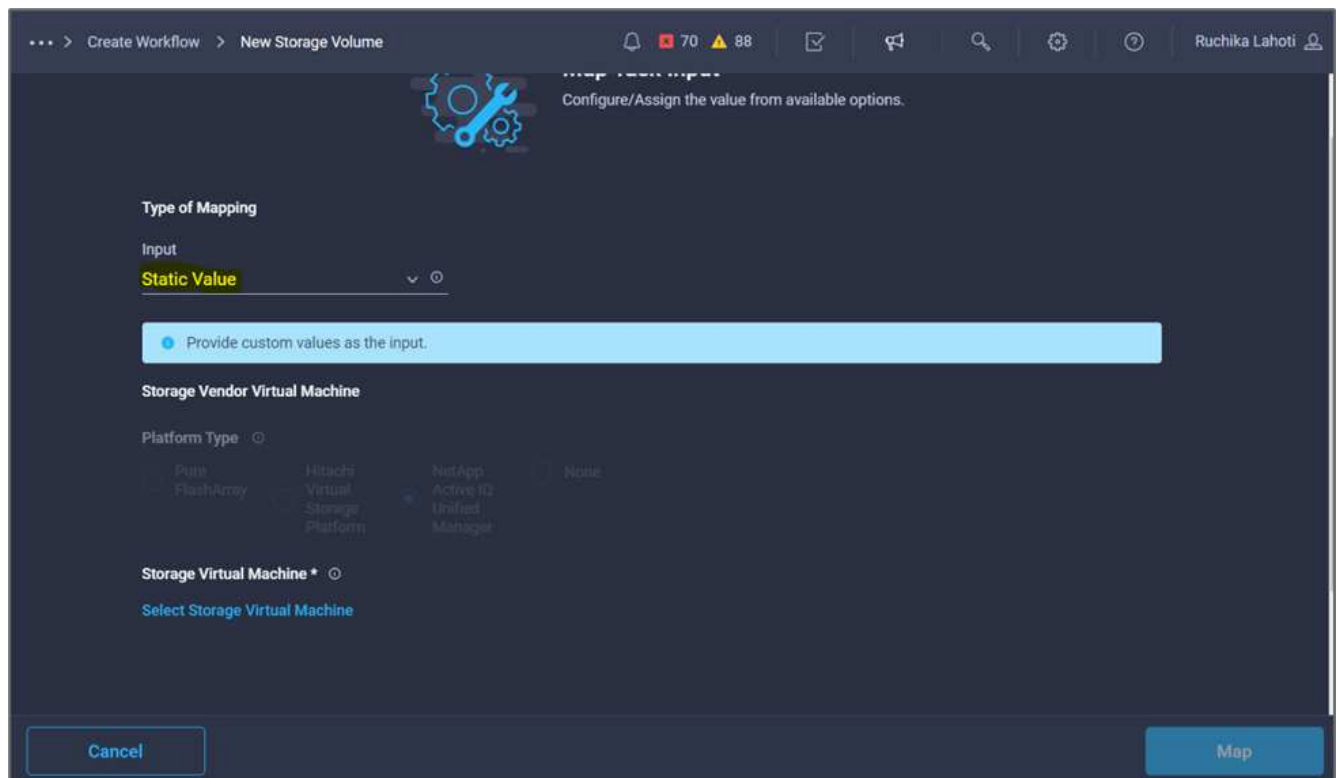
9. Click **Map**.



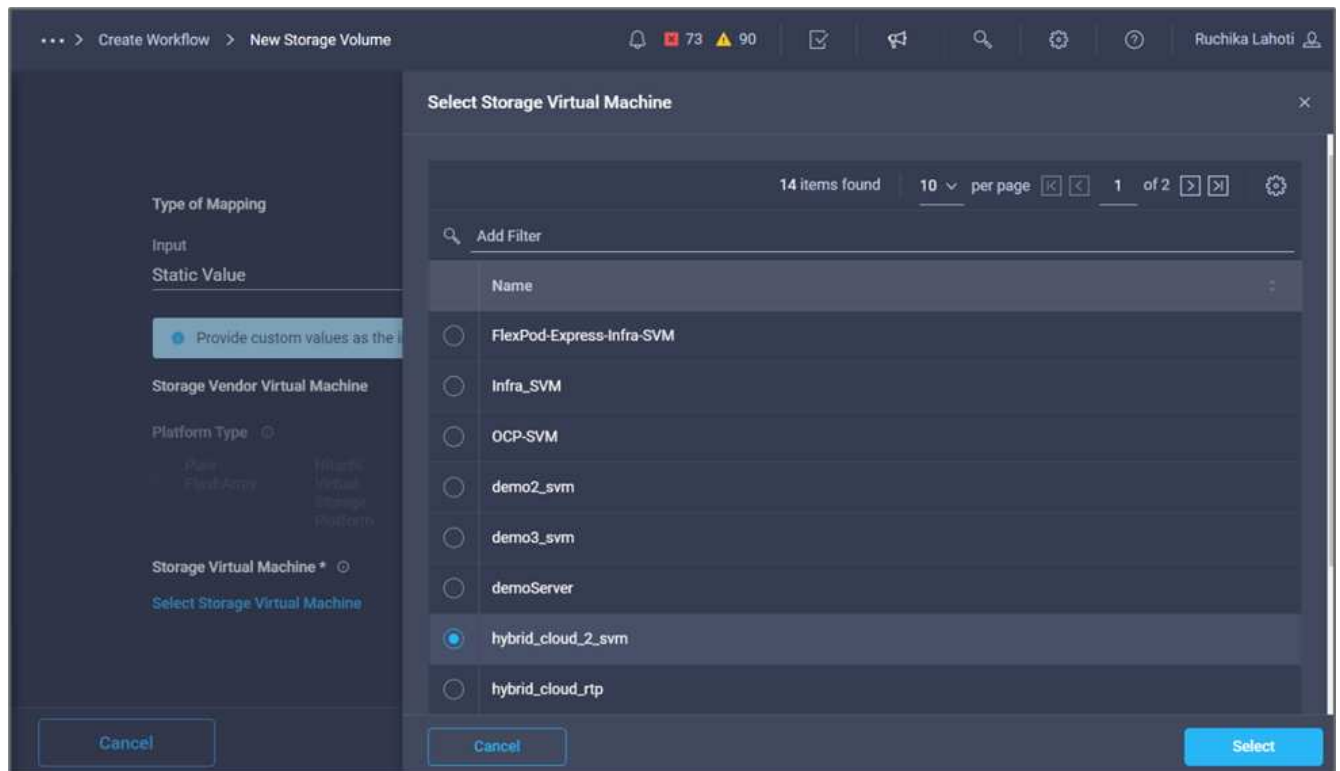
10. Click **Map** in the **Storage Vendor Virtual Machine** field.




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**.

... > Create Workflow > New Storage Volume

73 90

Map Task Input

Configure/Assign the value from available options.

Type of Mapping

Input

Direct Mapping v ⓘ

Map the workflow input, variable or any of the previous task's outputs to input.

Map to

Workflow Input v ⓘ

Input Name * v ⓘ

Add 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**.

Add Workflow Input

☒ Set Default Value ⓘ

☒ Allow User Override ⓘ

Default Values *

Storage Vendor Volume Options

Platform Type ⓘ

☐ Pure FlashArray

☐ Hitachi Virtual Storage Platform

☒ NetApp Active IQ Unified Manager

☐ None

Volume *

mssql_data_vol ⓘ

NFS Volume Option

☒ NFS ⓘ

Mount Path

/mssql_data_vol ⓘ

Cancel

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**.

... > Create Workflow > New Storage Volume > Volume Capacity

73 90

Ruchika Lahoti

Map Task Input

Configure/Assign the value from available options.

Type of Mapping

Input

Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.

Map to

Workflow Input


Input Name *

Add Workflow Input

Storage Vendor Volume Options


Cancel Map

24. In the Add Input wizard:
- Provide a display name and a reference name (optional).
 - Click **Required**.
 - For **Type**, select **Storage Capacity**.
 - Click **Set Default Value and Override**.
 - Provide a default value for the volume size and unit.
 - 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

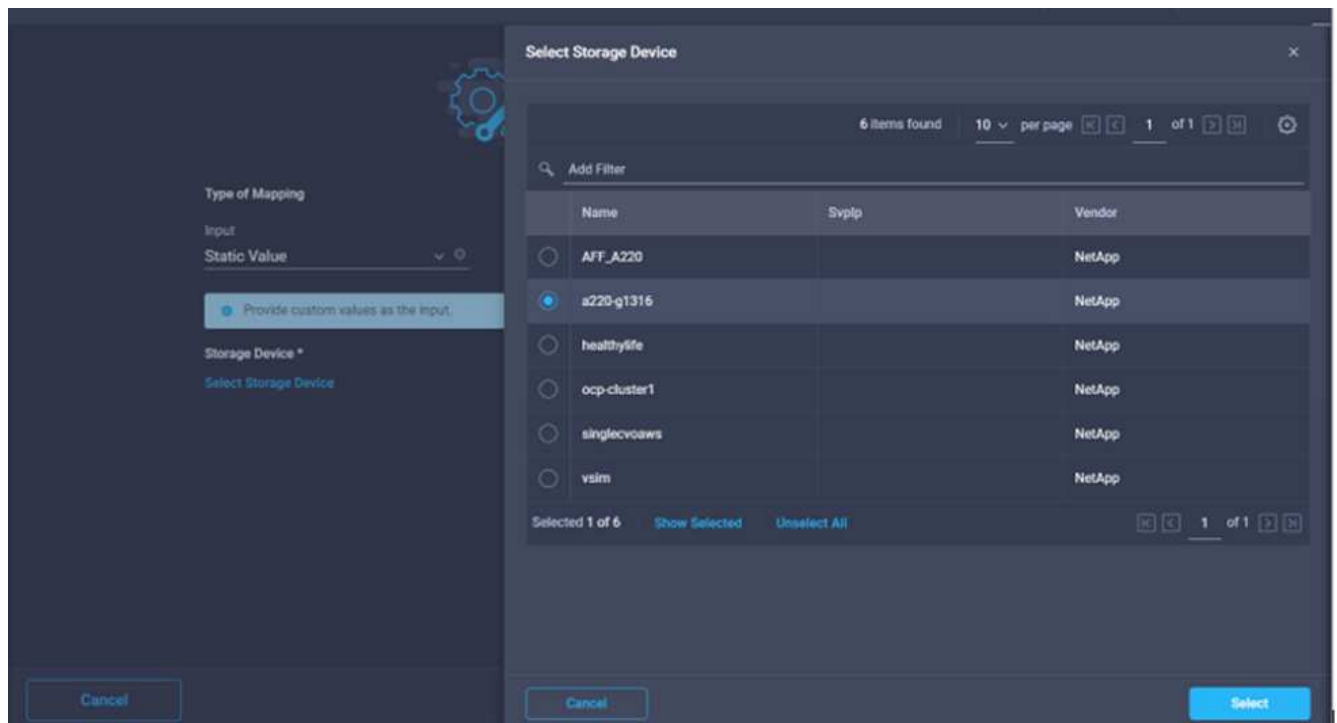
1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Storage > Add Storage Export Policy to Volume** task from the **Tools** section in the **Design** area.
3. Click **Add Storage Export Policy to Volume**. In the **Task Properties** area, click the **General** 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**.

The screenshot shows the NetApp Cloud Manager Designer interface. The top navigation bar includes 'CONFIGURE > Orchestration > Disaster recovery workflow > Edit'. The left sidebar shows the 'Tools' section with a search bar and a list of tasks. The main canvas displays a workflow with a 'Start' node, a 'Create volume in FlexPod' task, and an 'Add Storage Export Policy to Volume' task. The 'Add Storage Export Policy to Volume' task is selected, and its properties are shown on the right. The 'General' tab is active, showing the task name 'Add Storage Export Policy to Volume', version '1 (default)', task type 'Add Storage Export Policy to Volume', and user description 'Add an export policy to a volume with storage virtual machine'. The 'Task Details' section provides a description: 'Add an export policy to a volume with storage virtual machine name, volume name, export policy name as the inputs. On successful execution volume name and export policy added are generated as outputs.' The bottom of the interface shows a 'Last saved 7 minutes ago' status and 'Save' and 'Execute' buttons.


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.



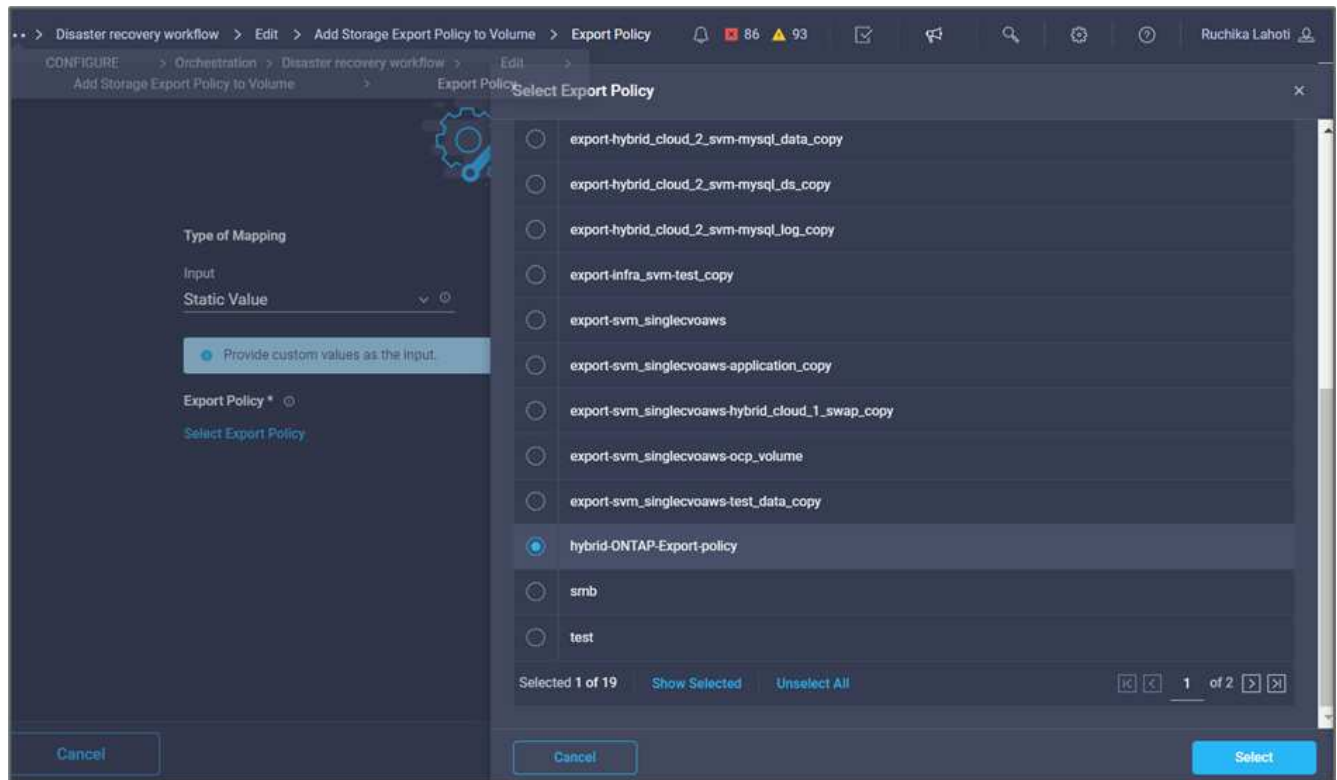
11. Click **Map**.
12. Click **Map** in the **Volume** field.
13. Click **Task Name** and then click **Create Volume in FlexPod**. Click **Output Name** and then **Volume**.



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.



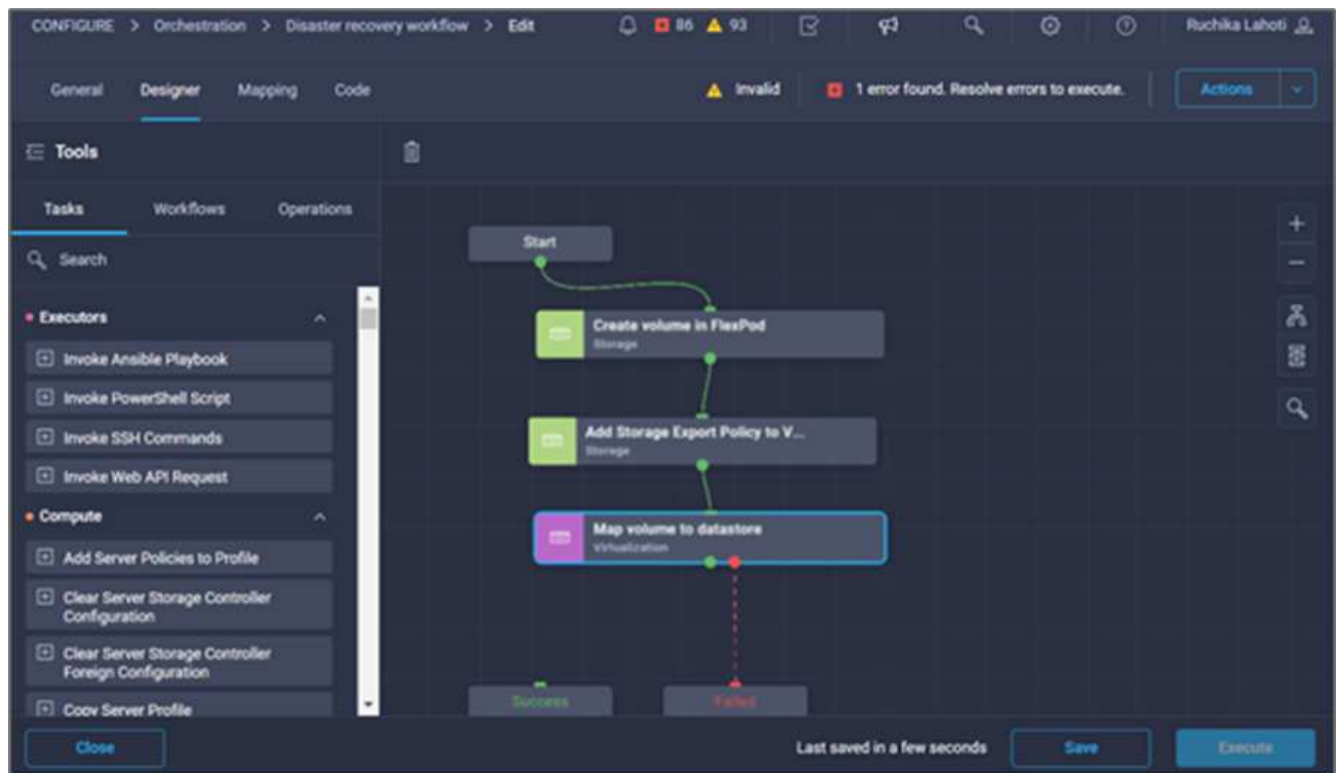
17. Click **Map** and then **Save**.




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

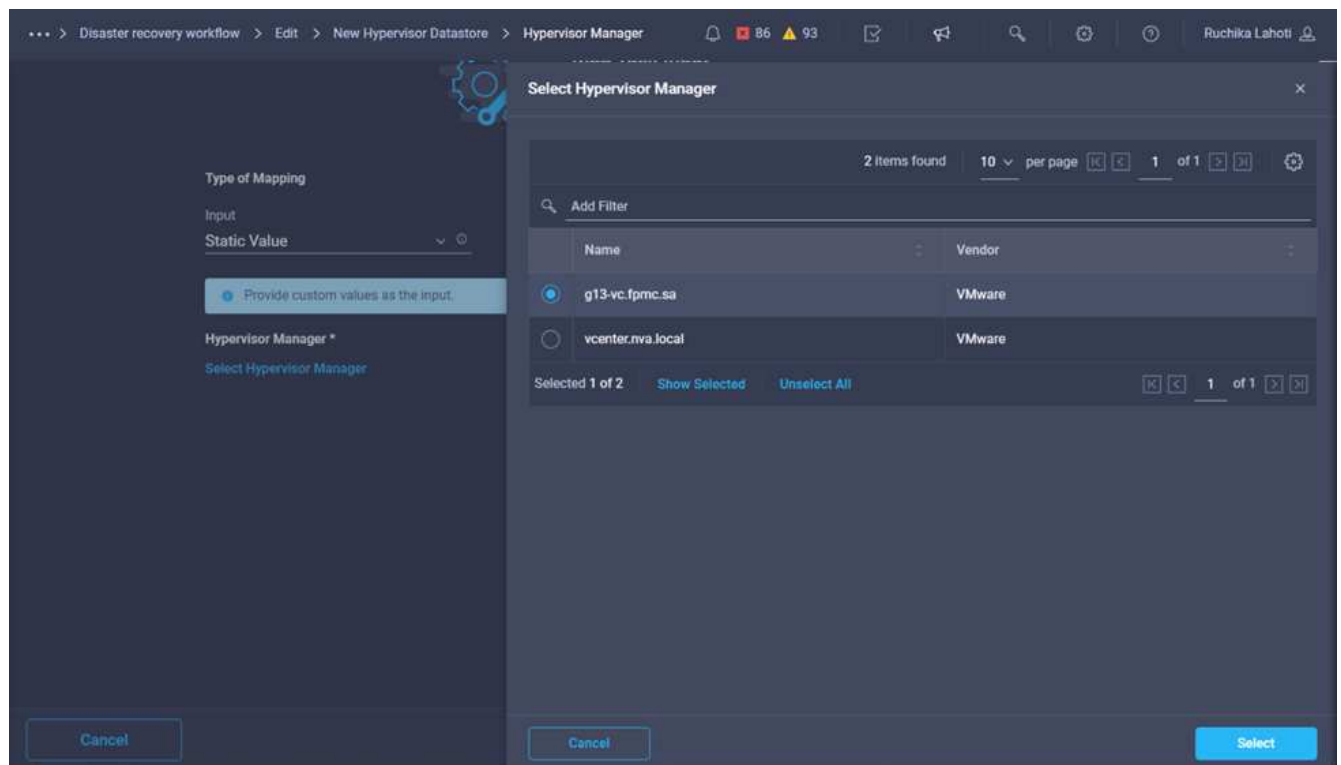
1. Go to the **Designer** tab and click **Tasks** from the **Tools** section.
2. Drag and drop the **Virtualization > New Hypervisor Datastore** task from the **Tools** section in the **Design** area.
3. Use Connector to make a connection between the **Add Storage Export Policy** and **New Hypervisor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




4. 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**.



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



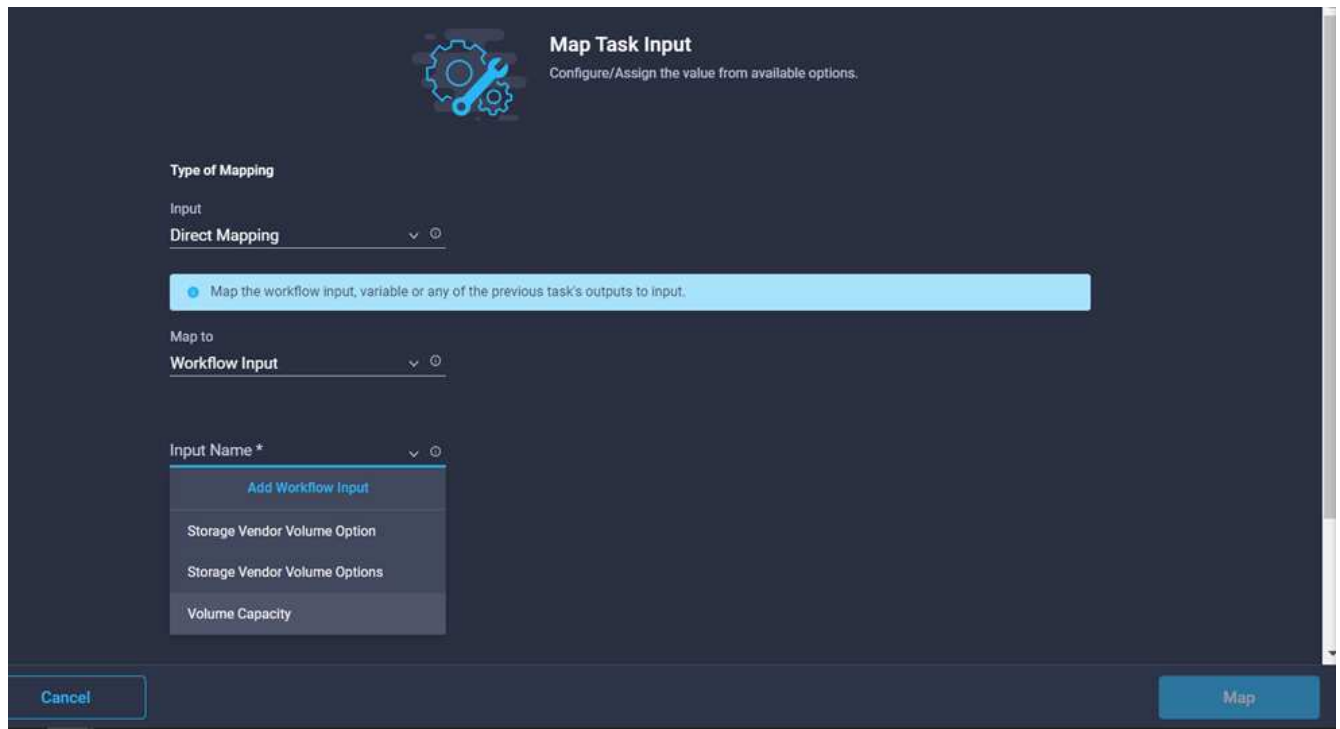
8. Click **Map**.



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**.



Map Task Input
Configure/Assign the value from available options.

Type of Mapping
Input
Direct Mapping

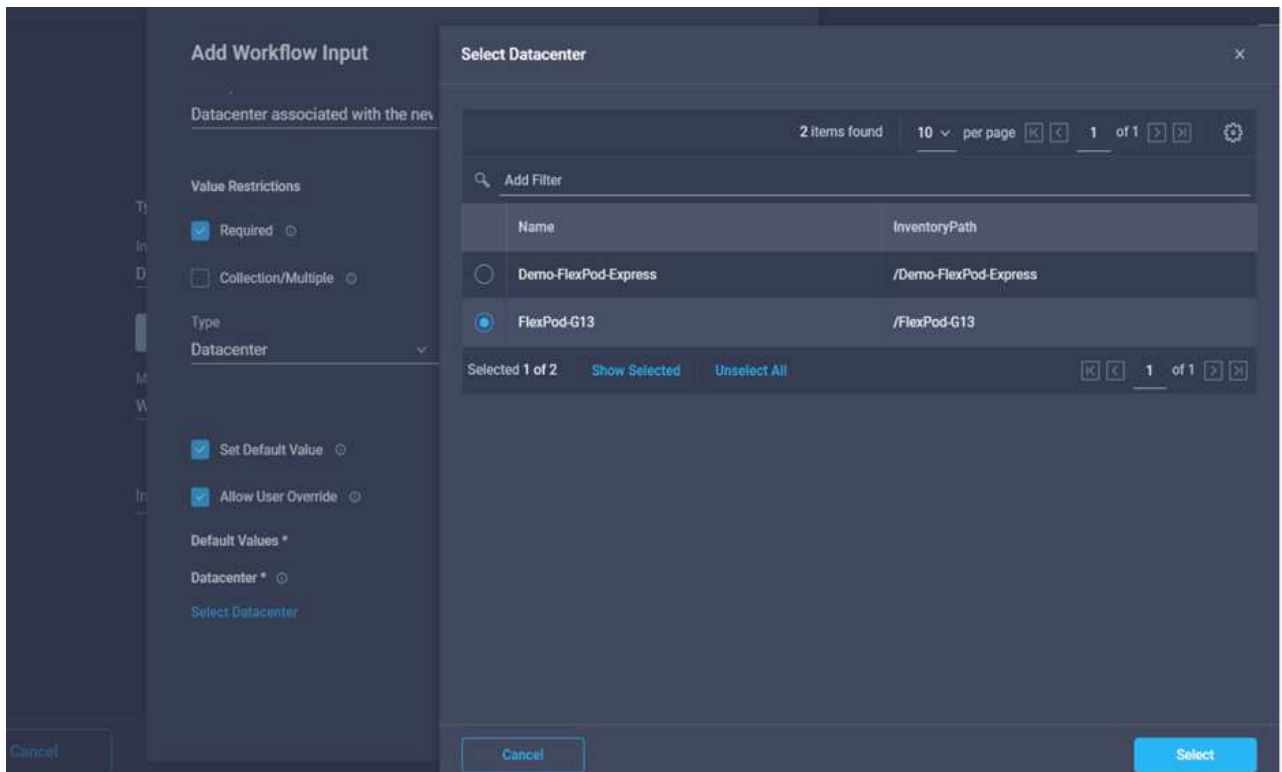
Map the workflow input, variable or any of the previous task's outputs to input.

Map to
Workflow Input

Input Name *
Add Workflow Input
Storage Vendor Volume Option
Storage Vendor Volume Options
Volume Capacity

Cancel Map

12. In the Add Input wizard, complete the following steps:
 - a. Provide a display name and reference name (optional).
 - b. Select **Datacenter** as the type.
 - c. Click **Set Default Value and Override**.
 - d. Click **Select Datacenter**.
 - e. Click the data center associated with the new datastore and then click **Select**.



Add Workflow Input

Datacenter associated with the new workflow input

Value Restrictions
☒ Required
☐ Collection/Multiple

Type
Datacenter

☒ Set Default Value
☒ Allow User Override

Default Values *
 Datacenter *
[Select Datacenter](#)

Select Datacenter

2 items found | 10 per page | 1 of 1

| Name | InventoryPath |
|----------------------|-----------------------|
| Demo-FlexPod-Express | /Demo-FlexPod-Express |
| FlexPod-G13 | /FlexPod-G13 |

Selected 1 of 2 | [Show Selected](#) | [Unselect All](#) | 1 of 1

Cancel Select

- Click **Add**.

13. Click **Map**.

14. Click **Map** in the **Cluster** field.

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

Map Task Input
Configure/Assign the value from available options.

Type of Mapping
Input
Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.


Map to
Workflow Input

Input Name *
Add Workflow Input
Datacenter
Storage Vendor Volume Option
Storage Vendor Volume Options
Volume Capacity

Cancel Map

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

- Provide a display name and reference name (optional).
- Click **Required**.
- Select Cluster as the type.
- Click **Set Default Value and Override**.
- Click **Select Cluster**.
- Click the cluster associated with the new datastore.
- Click **Select**.



h. Click **Add**.

17. Click **Map**.

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

Add Workflow Input

Cluster on which the datastore will

Value Restrictions

☐ Required

☐ Collection/Multiple

Type

Cluster

☒ Set Default Value

☒ Allow User Override

Default Values *

Cluster

Selected Cluster Washington

Cancel Add

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.

4 items found | 10 per page | 1 of 1

Add Filter

| Name |
|---|
| <input checked="" type="radio"/> 172.22.0.111 |
| <input type="radio"/> 172.22.0.112 |
| <input type="radio"/> esxi-01.nva.local |
| <input type="radio"/> esxi-02.nva.local |

Selected 1 of 4 | Show Selected | Unselect All | 1 of 1

Cancel | Select

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**.

Map Task Input

Configure/Assign the value from available options.

Type of Mapping

Input

Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.

Map to

Workflow Input

Input Name *

- Add Workflow Input
- Cluster
- Datacenter
- Storage Vendor Volume Option
- Storage Vendor Volume Options
- Volume Capacity

Cancel

24. In the Add Input wizard:
 - a. Provide a display name and reference name (optional).
 - b. Click **Required**.
 - c. Click **Set Default Value and Override**.
 - d. Provide a default value for the datastore and click **Add**.

Add Workflow Input

Type
String

Min 0 Max 0 Regex ^.{1,42}\$

☐ Secure

☒ Object Selector

☒ Set Default Value

☒ Allow User Override

Default Values *

Datastore *
hybrid-ds

Cancel 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**.

Type of Mapping

Input
Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.

Map to
Workflow Input

Input Name *
Add Workflow Input
Cluster
Datacenter
Datastore
Storage Vendor Volume Option
Storage Vendor Volume Options

Map

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 **Datastore** and click **Set Default Value and Override**.

Add Workflow Input

Display Name *
Type of Datastore

Reference Name *
DatastoreVersion

Description
Type and version of the new datast

Value Restrictions

☒ Required

☐ Collection/Multiple

Type
Types of Datastore


☒ Set Default Value

☒ Allow User Override

Default Values *
Type of Datastore

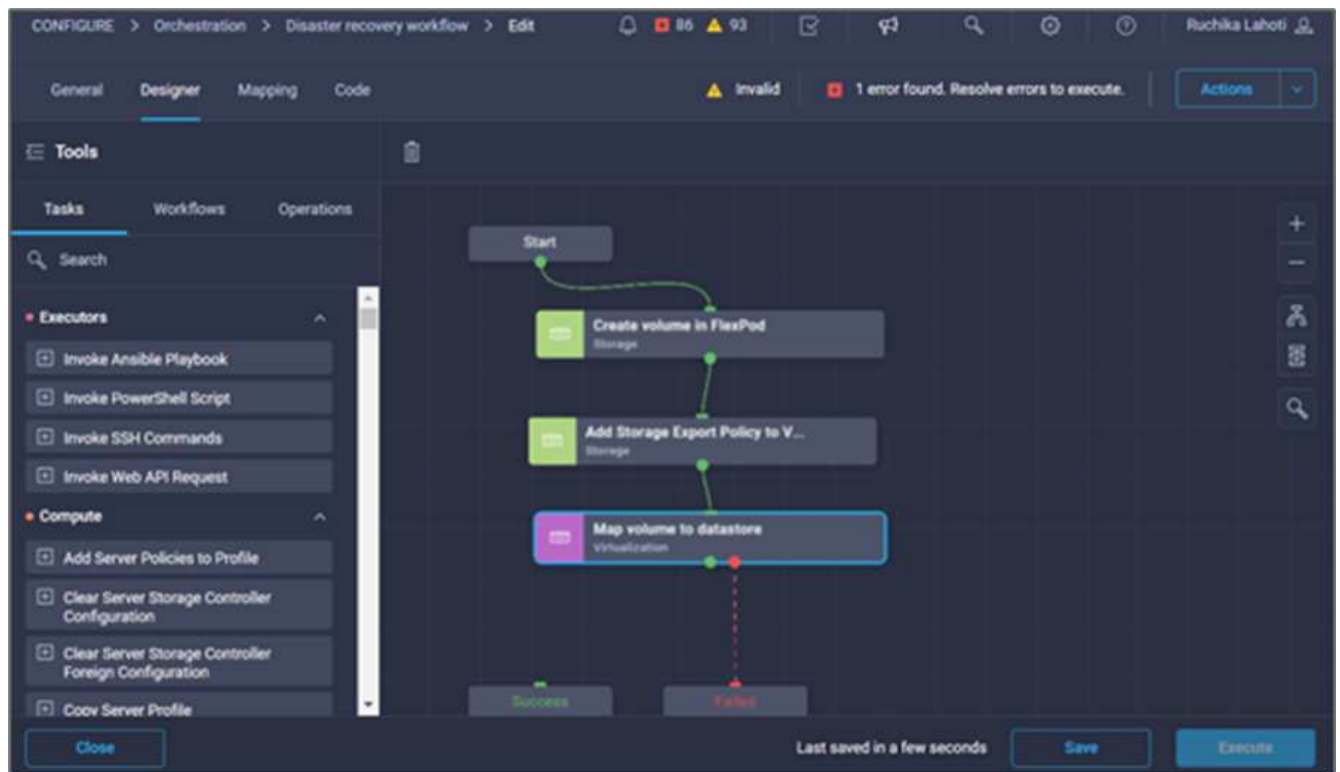
Cancel Add

- 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**.




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.](#)

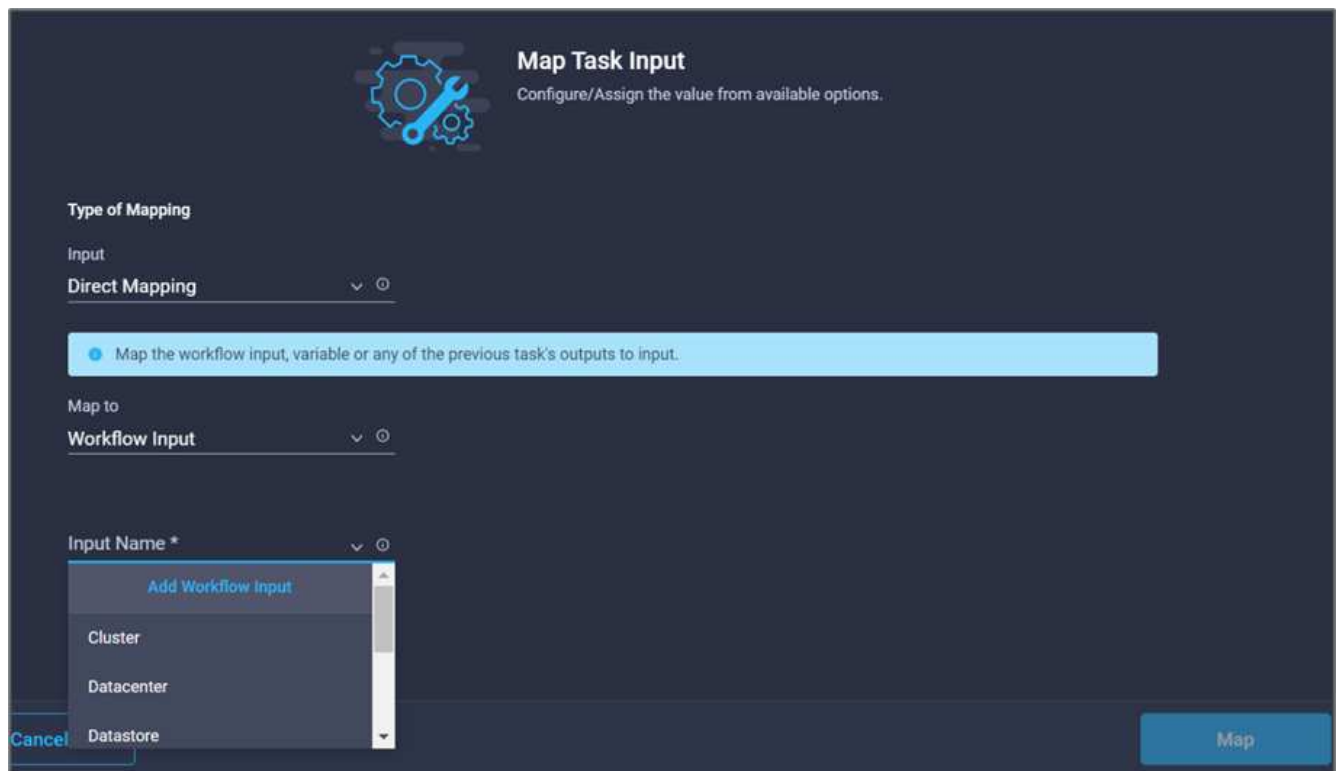


8. Click **Map**.

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**.



The image shows a 'Map Task Input' dialog box with a dark blue background. At the top left is a gear and wrench icon. The title 'Map Task Input' is at the top right, with the subtitle 'Configure/Assign the value from available options.' below it. The 'Type of Mapping' section has a dropdown menu set to 'Input', and the 'Direct Mapping' option is selected. A light blue instruction bar says 'Map the workflow input, variable or any of the previous task's outputs to input.' The 'Map to' section has a dropdown menu set to 'Workflow Input'. The 'Input Name *' section has a dropdown menu with options: 'Add Workflow Input' (highlighted in blue), 'Cluster', 'Datacenter', and 'Datastore'. At the bottom left is a 'Cancel' button, and at the bottom right is a 'Map' button.

Map Task Input
Configure/Assign the value from available options.

Type of Mapping
Input
Direct Mapping


Map the workflow input, variable or any of the previous task's outputs to input.

Map to
Workflow Input

Input Name *
Add Workflow Input
Cluster
Datacenter
Datastore

Cancel Map

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

Add Workflow Input ✕

Workspace Description ⓘ WorkspaceDescription ⓘ

Description
Description of the Terraform Work: ⓘ

Value Restrictions

☐ Required ⓘ

☐ Collection/Multiple ⓘ

Type
String ▼ ⓘ

Min 0 ⓘ Max 0 ⓘ Regex ⓘ

☐ Secure ⓘ

☒ Object Selector ⓘ

☒ Set Default Value ⓘ

☒ Allow User Override ⓘ

Cancel Add

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

Add Workflow Input

Value Restrictions

☐ Required ⓘ

☐ Collection/Multiple ⓘ

Type
String ▼ ⓘ

Min **0** ⓘ Max **0** ⓘ Regex ⓘ

☐ Secure ⓘ

☒ Object Selector ⓘ

☒ Set Default Value ⓘ

☒ Allow User Override ⓘ

Default Values *

Workspace Description
workspace to create CVO and configure SnapMirror ⓘ

Cancel Add

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

Type of Mapping

Input
 Static Value

Provide custom values as the input.

Execution Mode

ExecutionMode
 remote

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

Type of Mapping

Input
 Static Value

Provide custom values as the input.


Apply Method

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.


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




- In the **Workflow Properties** area, click **Inputs**.
- Click **Map** in the **Terraform Cloud Target** field.
- 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](#).



- Click **Map**.
- Click **Map** in the ***Terraform Organization Name *** field.
- 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**.

Add Workflow Input

Display Name *
Terraform Variable

Reference Name *
TerraformAddVariable

Description
Terraform Variable to be added

Value Restrictions

☒ Required

☐ Collection/Multiple

Type
String

Min
0

Max
0

Regex

☐ Secure

☐ Object Selector

Cancel Add

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 *

Edit Mapping

Custom Value

View Value

Workspace ID *

Edit Mapping

Task Output

WorkspaceId | Add Terraform Work...

Terraform Variable

Edit Mapping

Workflow Input

Terraform Variables

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.

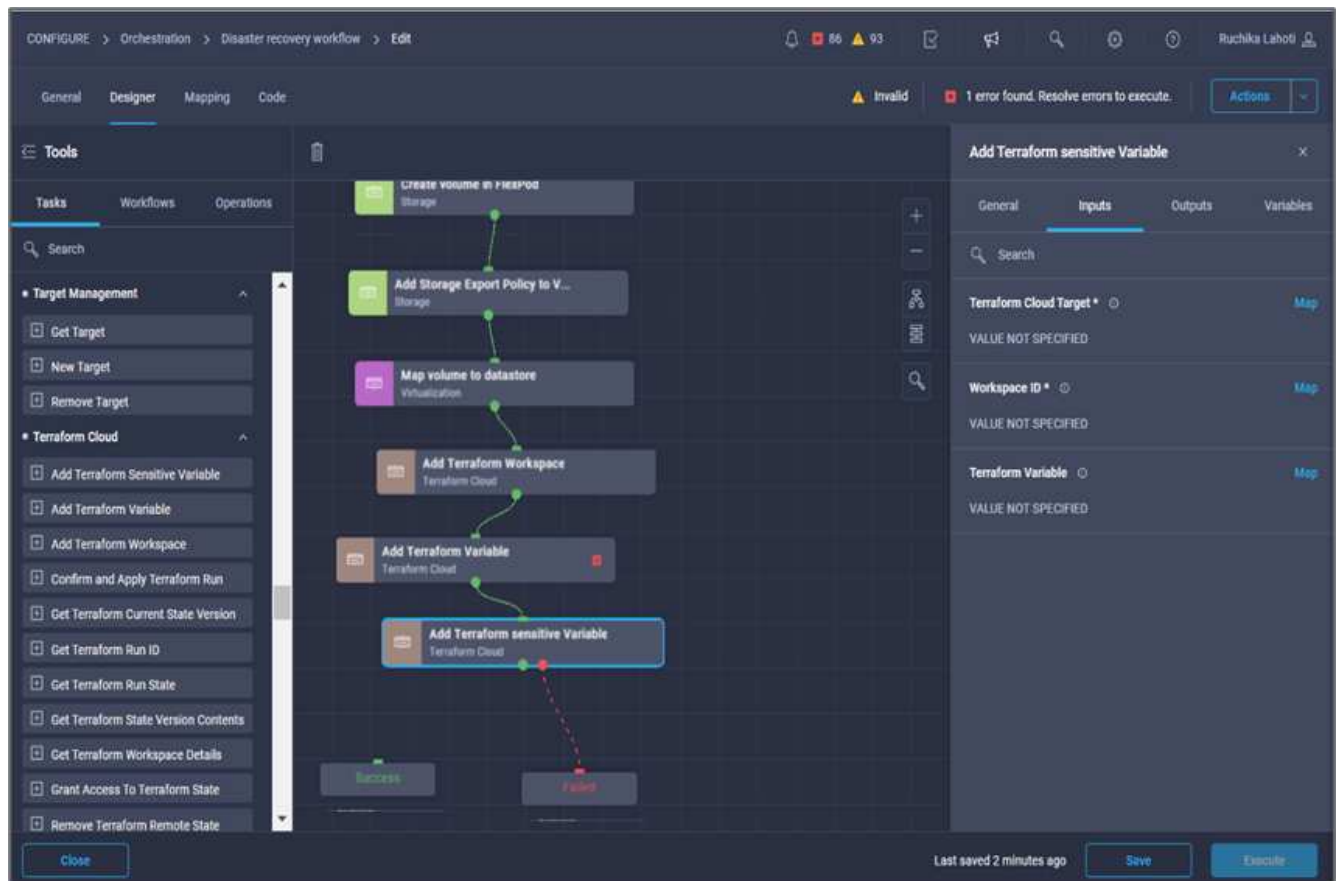
Procedure 7: Add sensitive variables to a workspace

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 *
terraformensitivevariable ⓘ

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

Key *

cloudmanager_refresh_token ⓘ

Value

ⓘ ⓘ

Description

cloudmanager refresh token ⓘ

☐ HCL ⓘ

+

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.

The screenshot displays the Cisco Intersight Designer interface. On the left, the 'Tools' panel shows a list of tasks under the 'Tasks' tab, with 'Start New Terraform Plan' highlighted. The central 'Design' area shows a workflow diagram with several tasks connected by arrows. The 'Start New Terraform Plan' task is the final step in the sequence. On the right, the 'Task Properties' panel for 'Start New Terraform Plan' is open, showing the 'General' tab. The 'Name' field is set to 'Start New Terraform Plan', the 'Version' is '1 (default)', and the 'Task Type' is 'Start New Terraform Plan'. The 'User Description' field contains the text: 'Starts a new plan or destroys a plan in the given Terraform Workspace'. At the bottom of the interface, there are buttons for 'Close', 'Save', and 'Execute', along with a status indicator 'Last saved 6 minutes ago'.

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**.

Map Task Input
Configure/Assign the value from available options.

Type of Mapping
Input
Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.

Map to
Task Output

Task Name *
Add Terraform Workspace

- Add Terraform sensitive Variable
- Add Terraform Variable
- Add Terraform Workspace**
- Map volume to datastore
- Add Storage Export Policy to Volume
- Create volume in FlexPod

Output Name *

Map

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**.

Add Workflow Input

☒ Required ⓘ

☐ Collection/Multiple ⓘ

Type
String ▼ ⓘ

Min **0** ⓘ Max **0** ⓘ Regex ⓘ

☐ Secure ⓘ

☐ Object Selector ⓘ

☒ Set Default Value ⓘ

☒ Allow User Override ⓘ


Default Values *

*Reason for starting plan **

terraform plan for replication between onprem volume and CVO ⓘ

Cancel 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**.

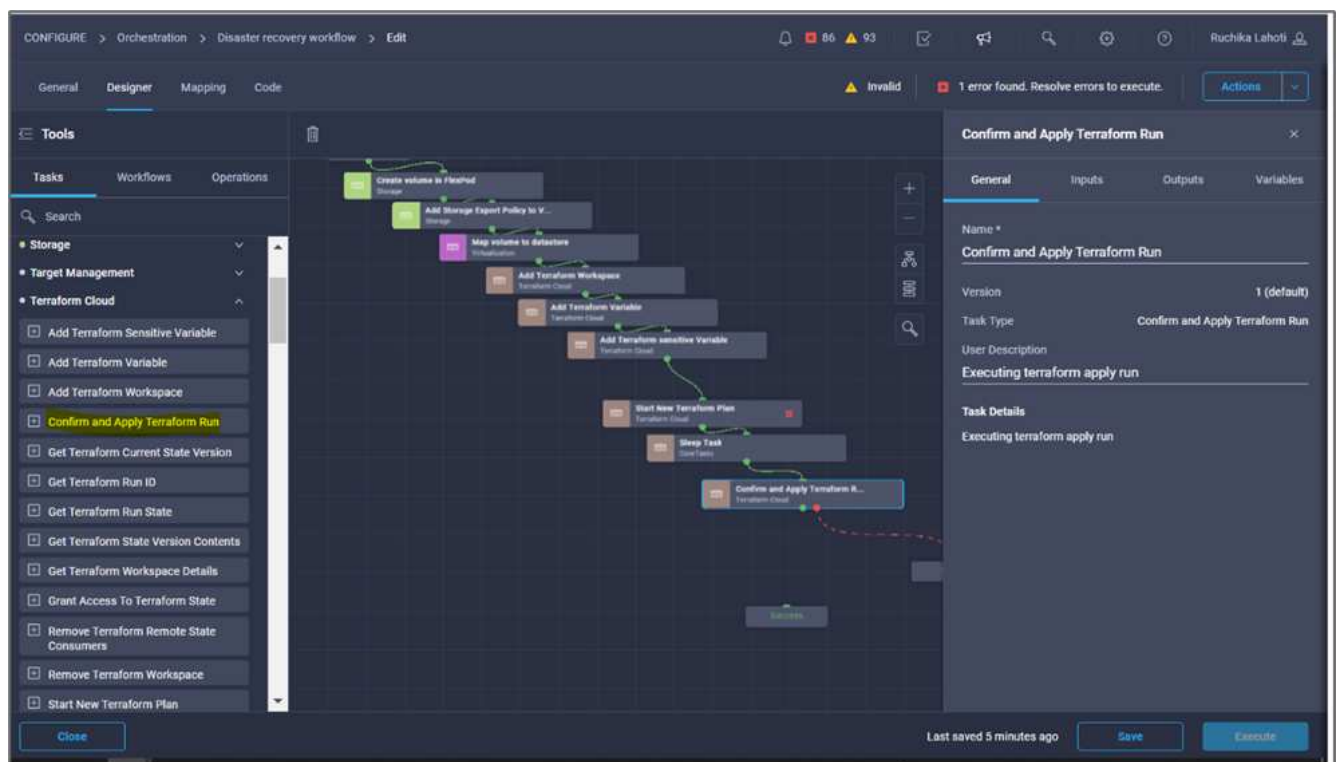


8. Click **Map**.
9. Click **Save**.

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**.

CONFIGURE > Orchestration > Disaster recovery workflow > Edit > Confirm and Apply Terraform Run > Run ID

86 93

Ruchika Lahoti

Map Task Input

Configure/Assign the value from available options.

Type of Mapping

Input
Direct Mapping

Map the workflow input, variable or any of the previous task's outputs to input.

Map to

Task Output

Task Name *
Start New Terraform Plan

Output Name *
Run ID

Cancel Map

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-prem FlexPod storage

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
☐ Pure FlashArray
 ☐ Hitachi Virtual Storage Platform
 ☒ NetApp Active IQ Unified Manager
 ☐ 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

| Name |
|--------------------|
| 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.

ONTAP System Manager

Search actions, objects, and pages

DASHBOARD

INSIGHTS

STORAGE

Overview

Volumes

LUNs

Consistency Groups

NVMe Namespaces

Shares

Qtrees

Quotas

Storage VMs

Tiers

Volumes

+ Add

More

| | Name | Storage VM | Status | Capacity |
|---|-----------------------|--------------------|--------|---|
| | Q | Q hybrid-cloud-svr | (All) | > |
| ✓ | application_copy | hybrid-cloud-svm | Online | 3.12 MiB used 19 GiB available 20 GiB |
| ✓ | audit_log_vol | hybrid-cloud-svm | Online | 32.7 MiB used 200 GiB available 200 GiB |
| ✓ | hybrid_cloud_svm_root | hybrid-cloud-svm | Online | 1.68 MiB used 971 MiB available 1 GiB |
| ✓ | test | hybrid-cloud-svm | Online | 648 KiB used 972 MiB available 1 GiB |
| ✓ | Test_Vol1 | hybrid-cloud-svm | Online | 10.6 MiB used 9.99 GiB available 10 GiB |

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

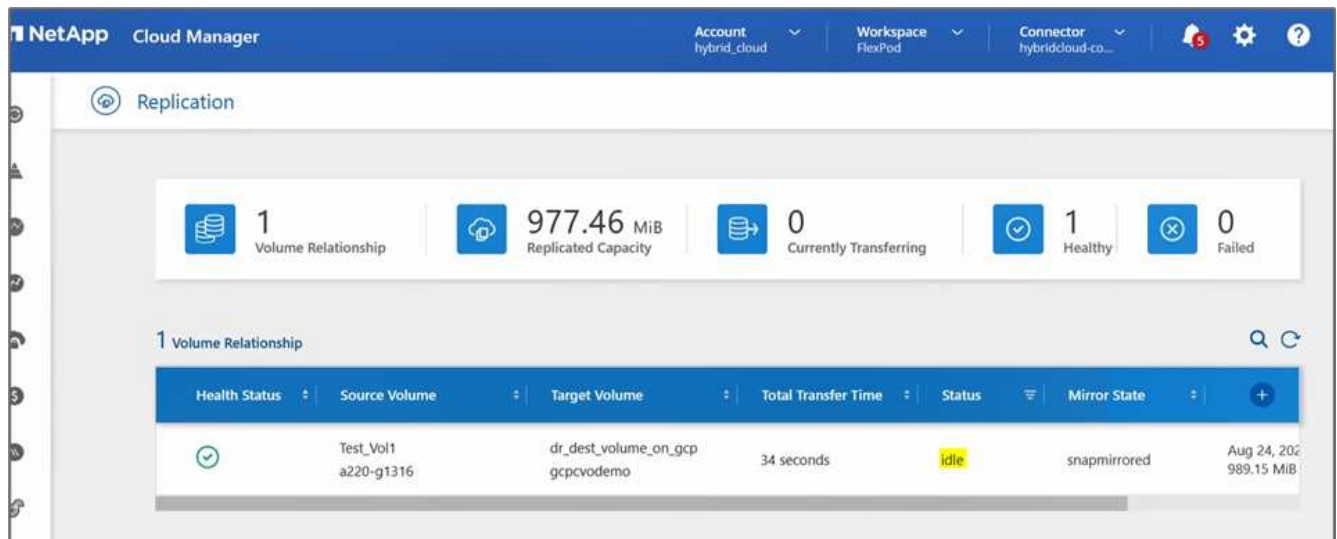
```

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
888a23c8495ad33fdf11a931ffc344c3643f15d5cefedbbf1326016e31ec5a59 test.zip
root@hybridcloudbackup:/snapmirror_demo#
root@hybridcloudbackup:/snapmirror_demo#

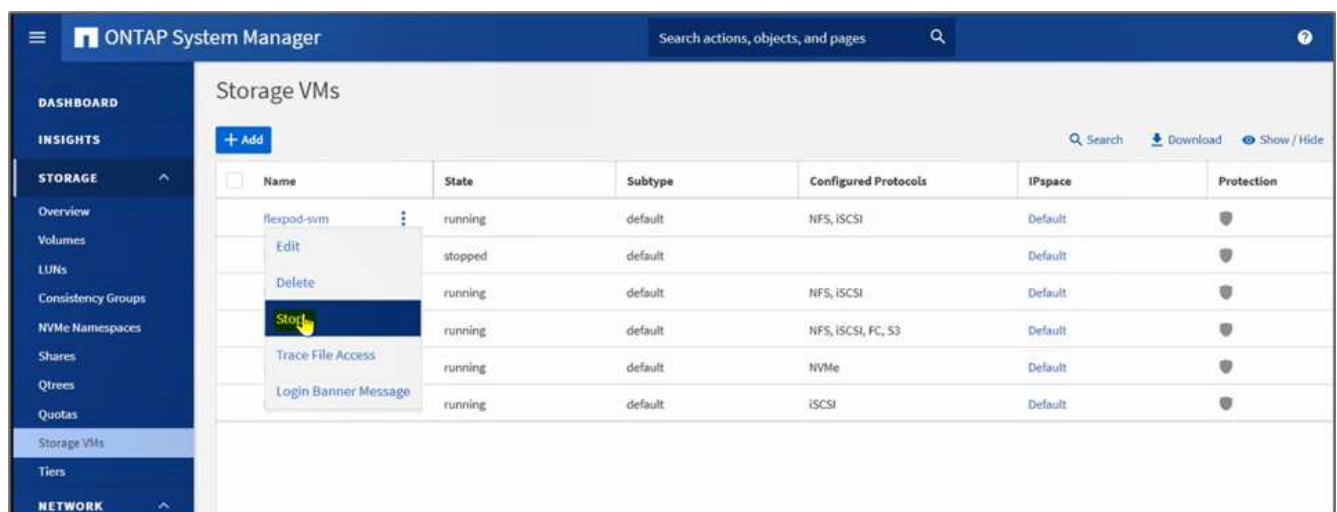
```

- 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**.


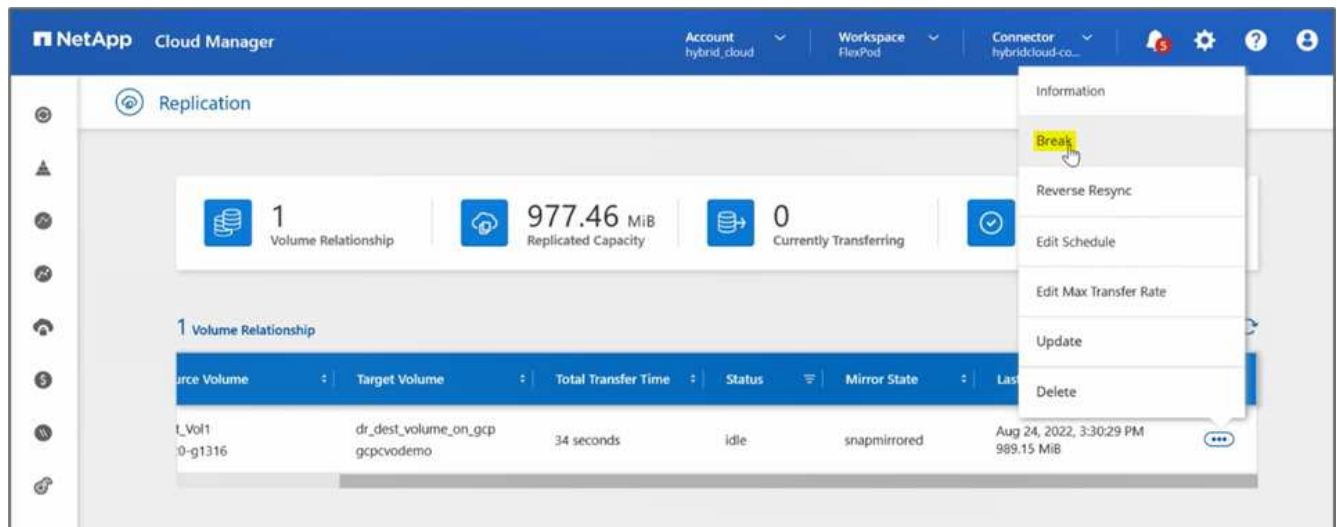


8. When the data transfer is complete, simulate a disaster on the source side by stopping the SVM that hosts the `Test_vol1` volume.


After the SVM has been stopped, the `Test_vol1` volume is not visible in the Cloud Manager.




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



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



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



↶ 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$
ruchikal_netapp_com@demo-nfs:/snapmirror_dest$ sha256sum test.zip
888a23c8495ad33fdf11a931ffc344c3643f15d5cefedbbf1326016e31ec5a59 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

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

- Cisco Intersight Service for HashiCorp Terraform Documentation

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

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

- 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

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 prerequisite

<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

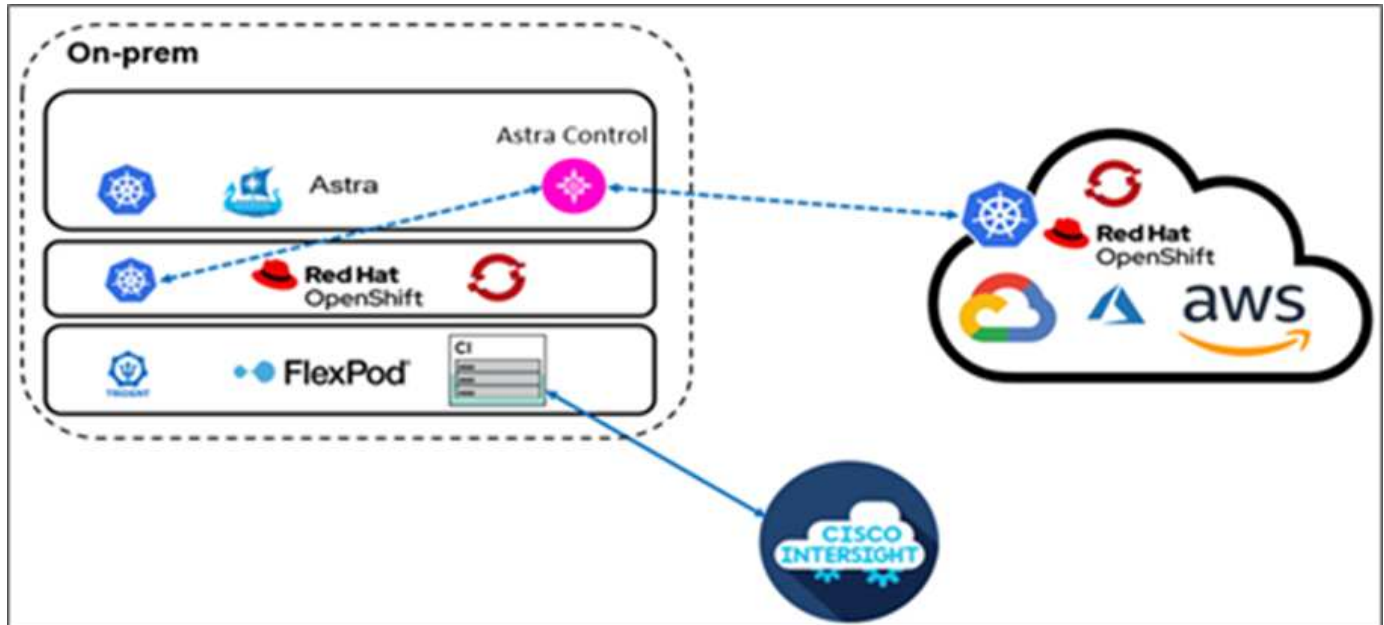
Abhinav Singh

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.

Solution components

[Previous: Solution overview.](#)

FlexPod

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

Astra Control

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

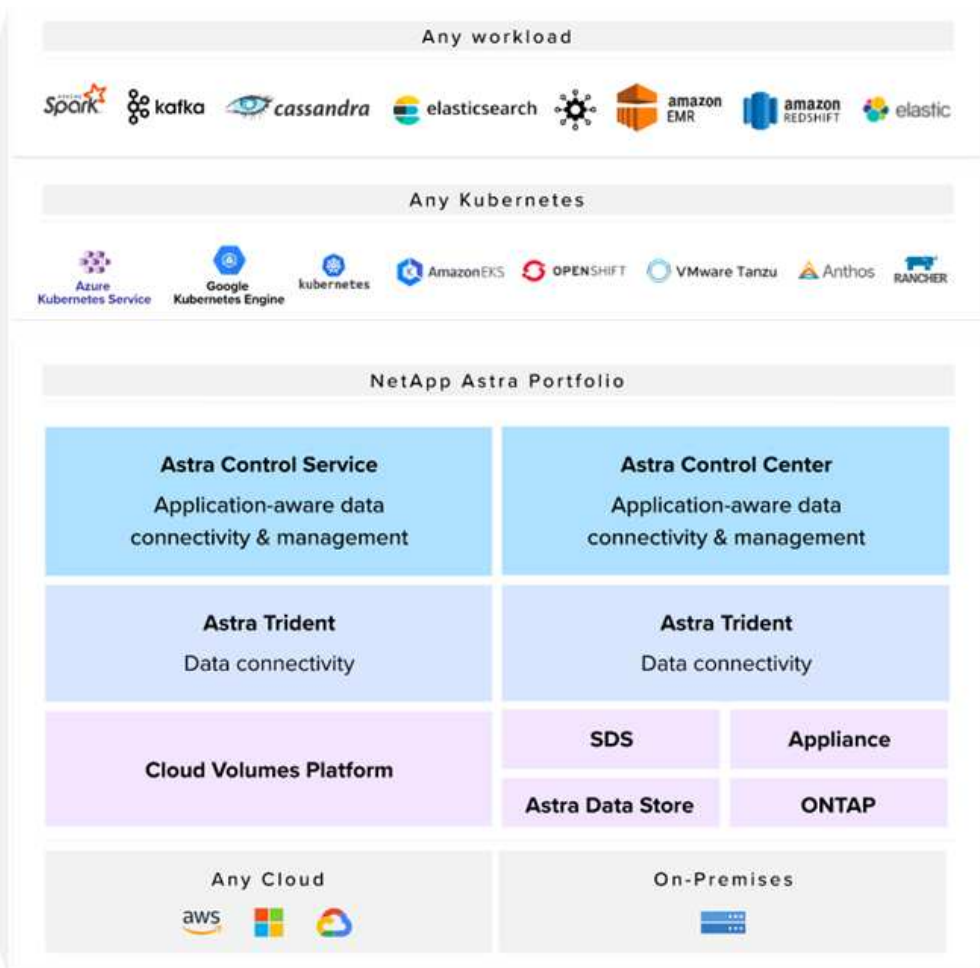
Features

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

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

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

The following image shows the Astra Control capabilities.




Astra Control Consumption models

Astra Control is available in two consumption models:

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

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

The following image shows the Astra Control architecture.



Astra Trident

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

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

Storage backend

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

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

functionality, such as NetApp Snapshot copies and mirroring.

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

NetApp Cloud Volumes ONTAP

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

Key benefits include:

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

Cloud Central

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

Cloud Manager

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

Connector

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

Connector is supported in the following locations:

- AWS
- Microsoft Azure
- Google Cloud
- On your premises

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

NetApp Cloud Insights

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

NetApp Active IQ Unified Manager

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

Cisco Intersight

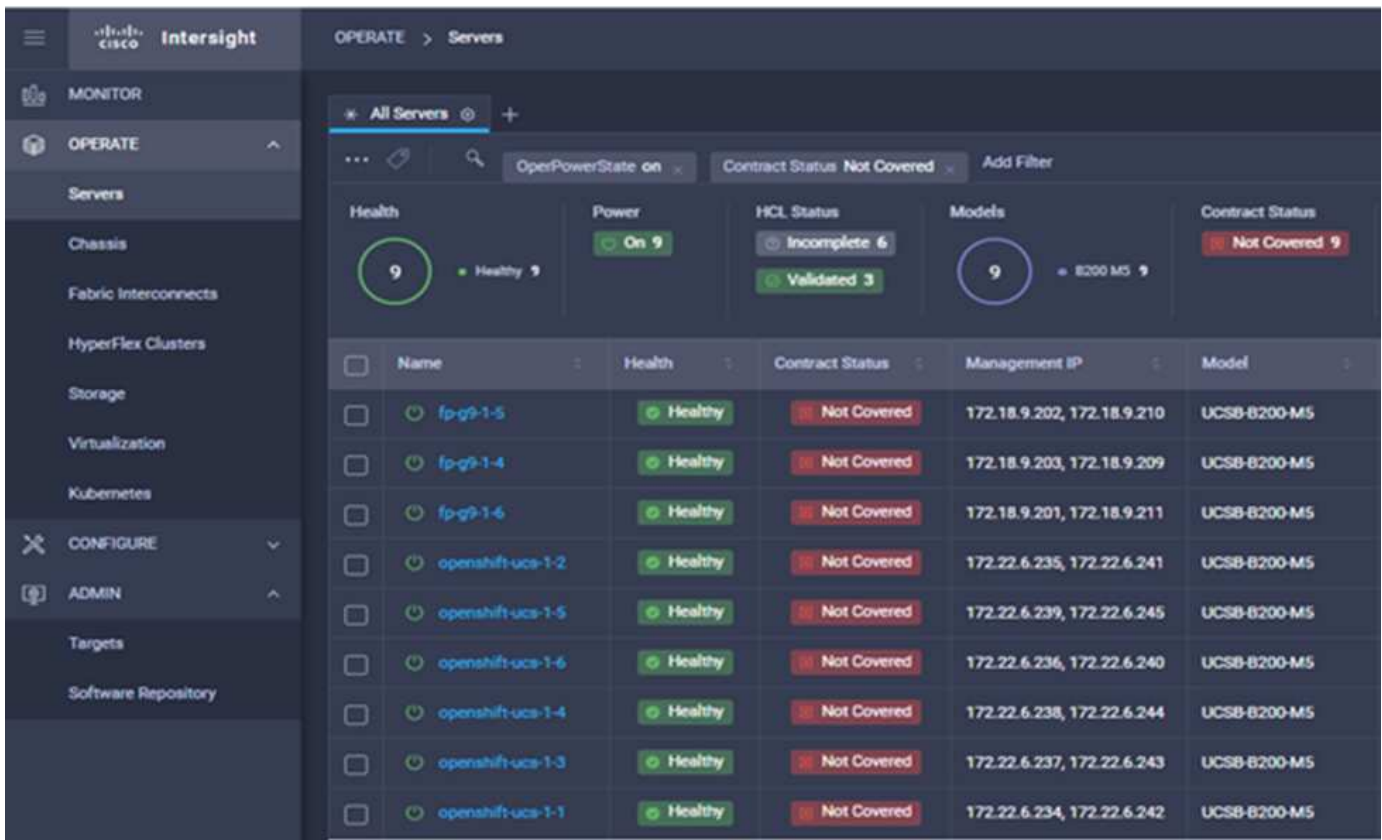
Cisco Intersight is a SaaS platform that delivers intelligent automation, observability, and optimization for traditional and cloud-native applications and infrastructure. The platform helps drive change with IT teams and delivers an operating model designed for hybrid cloud.

Cisco Intersight provides the following benefits:

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

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

The following image shows the Cisco Intersight dashboard.




Red Hat OpenShift Container Platform

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

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

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



Kubernetes infrastructure

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

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

Red Hat Core OS

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

RHCOS includes the following features:

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

VMware vSphere 7.0

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

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

VMware vSphere vCenter

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

Hardware and software revisions

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

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

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

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

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

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

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

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


Installation and configuration

FlexPod for OpenShift Container Platform 4 bare-metal installation

[Previous: Solution components.](#)

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

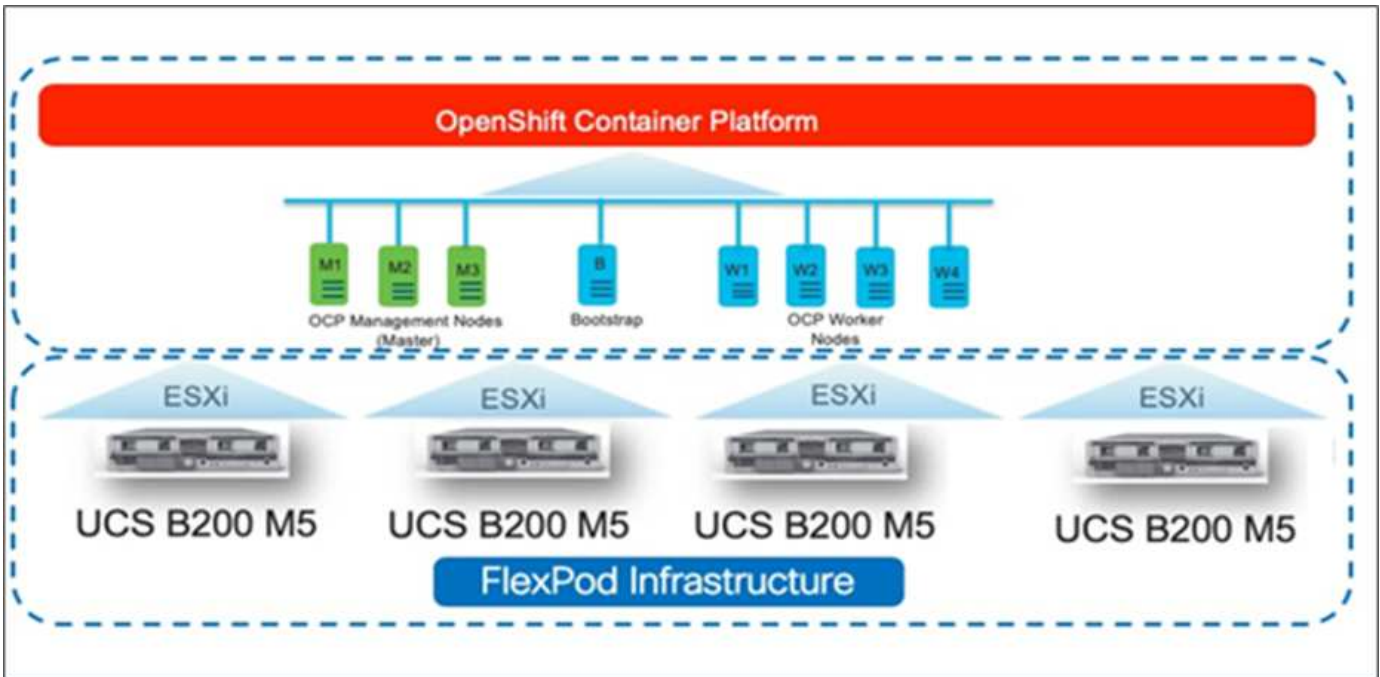
The following figure illustrates the OpenShift Container Platform 4 Bare Metal on FlexPod.



FlexPod for OpenShift Container Platform 4 on VMware installation

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

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



Next: Red Hat OpenShift on AWS.

Red Hat OpenShift on AWS

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

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

| Instances (6) Info | | | | | | | | |
|---|--------------------------------------|---------------------|----------------|---------------|-------------------|-----------------|----------|--|
| <input type="text" value="Search"/> | | | | | | | | |
| <input type="button" value="ocp"/> <input type="button" value="Clear filters"/> | | | | | | | | |
| <input type="checkbox"/> | Name | Instance ID | Instance state | Instance type | Availability Zone | Private IP a... | Key name | |
| <input type="checkbox"/> | ocpaws-v58kn-master-0 | i-0d2d81ca91a54276d | Running | m5.xlarge | us-east-1b | 172.30.165.160 | - | |
| <input type="checkbox"/> | ocpaws-v58kn-master-1 | i-0b161945421d2a23c | Running | m5.xlarge | us-east-1c | 172.30.166.162 | - | |
| <input type="checkbox"/> | ocpaws-v58kn-master-2 | i-0146a665e1060ea59 | Running | m5.xlarge | us-east-1a | 172.30.164.209 | - | |
| <input type="checkbox"/> | ocpaws-v58kn-worker-us-east-1a-zj8dj | i-05e6efa18d136c842 | Running | m5.large | us-east-1a | 172.30.164.128 | - | |
| <input type="checkbox"/> | ocpaws-v58kn-worker-us-east-1b-7nmbc | i-0879a088b50d2d966 | Running | m5.large | us-east-1b | 172.30.165.93 | - | |
| <input type="checkbox"/> | ocpaws-v58kn-worker-us-east-1c-96j6n | i-0c24ff3c2d701f82c | Running | m5.large | us-east-1c | 172.30.166.51 | - | |


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

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

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

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

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

NetApp Cloud Volumes ONTAP

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

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

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

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

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

Astra Control Center installation on OpenShift Container Platform

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

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

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

Environment requirements

- Astra Control Center supports multiple Kubernetes distributions; for Red Hat OpenShift, the supported

versions include Red Hat OpenShift Container Platform 4.8 or 4.9.

- Astra Control Center requires the following resources in addition to the environment's and the end-user's application resource requirements:

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

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



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

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



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

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

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



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

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

Center can manage applications on the K8s cluster on which it is running.

Application management requirements

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

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

Install Astra Control Center using OpenShift OperatorHub

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

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

```
openssl dgst -sha256 -verify astra-control-center[version].pub  
-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.

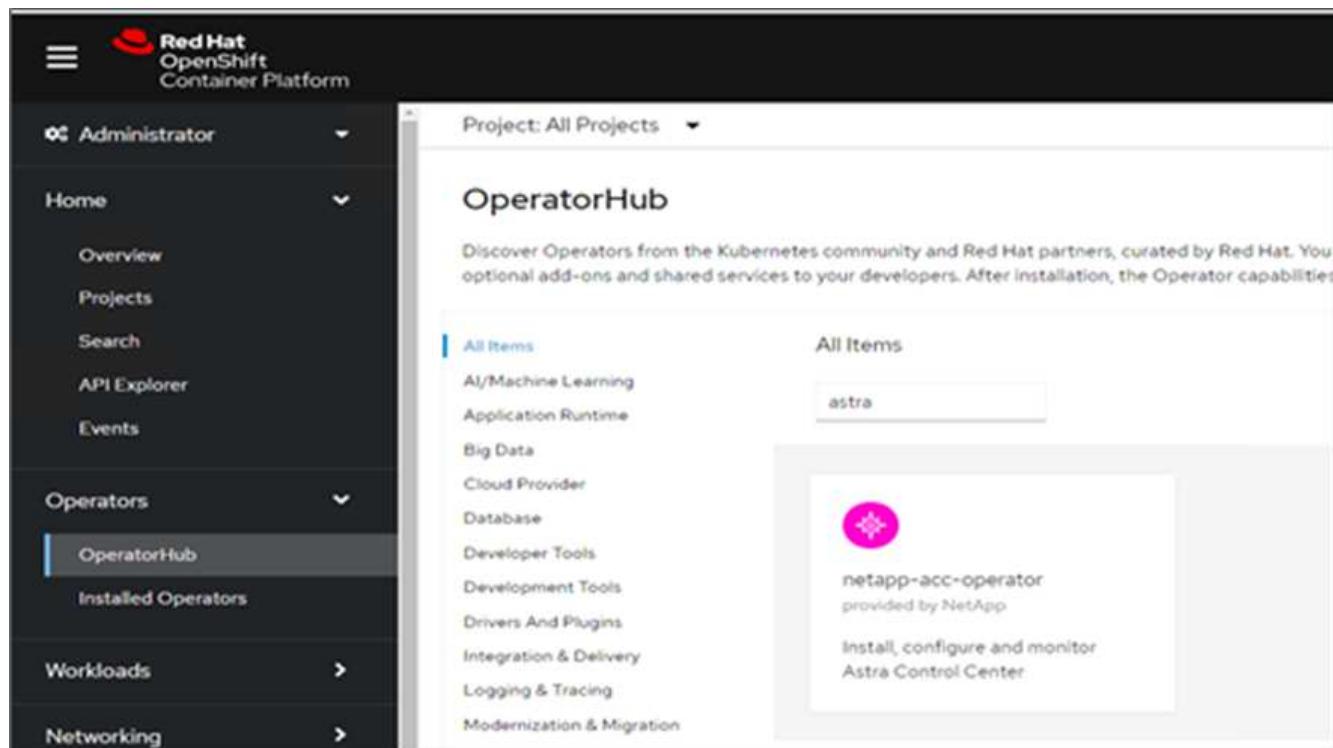
For Docker:

```
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:

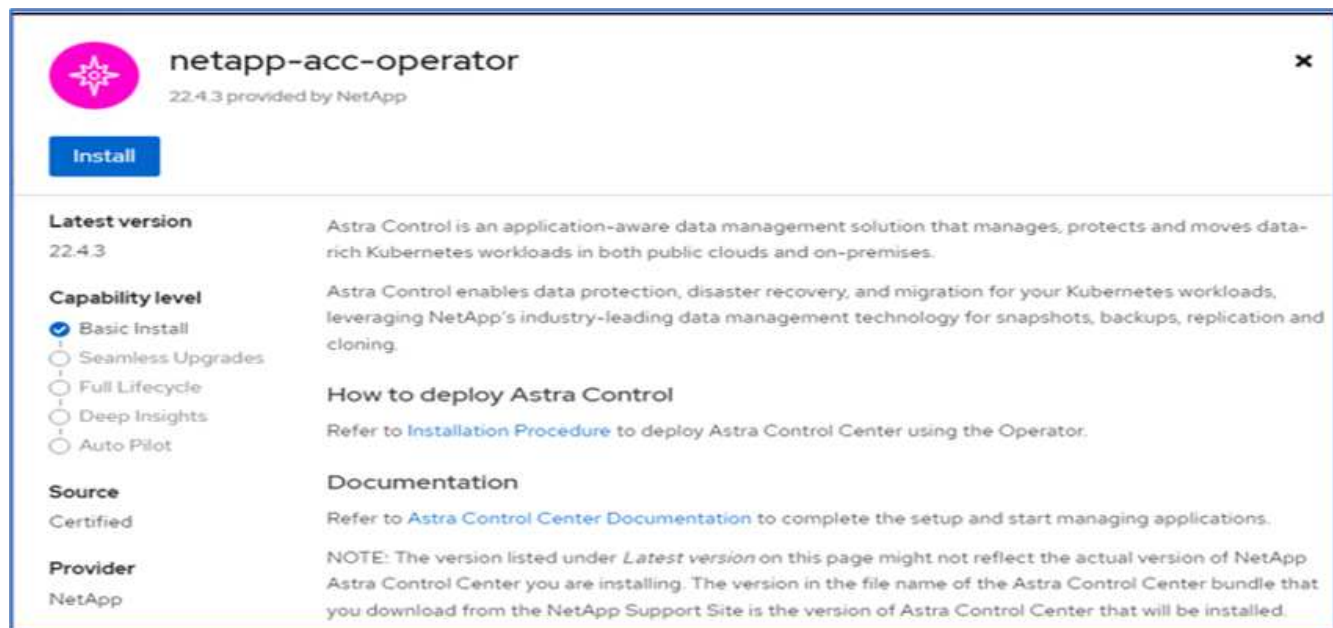
```
export REGISTRY=[Registry_path]
for astraImageFile in $(ls images/*.tar) ; do
    # Load to local cache. And store the name of the loaded image trimming
    the 'Loaded images: '
    astraImage=$(podman load --input ${astraImageFile} | sed 's/Loaded
image(s): //' )
    astraImage=$(echo ${astraImage}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



`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.



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 *

PR 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.

netapp-acc-operator
provided by NetApp

Provided APIs

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

Install **Cancel**

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

netapp-acc-operator
22.4.3 provided by NetApp

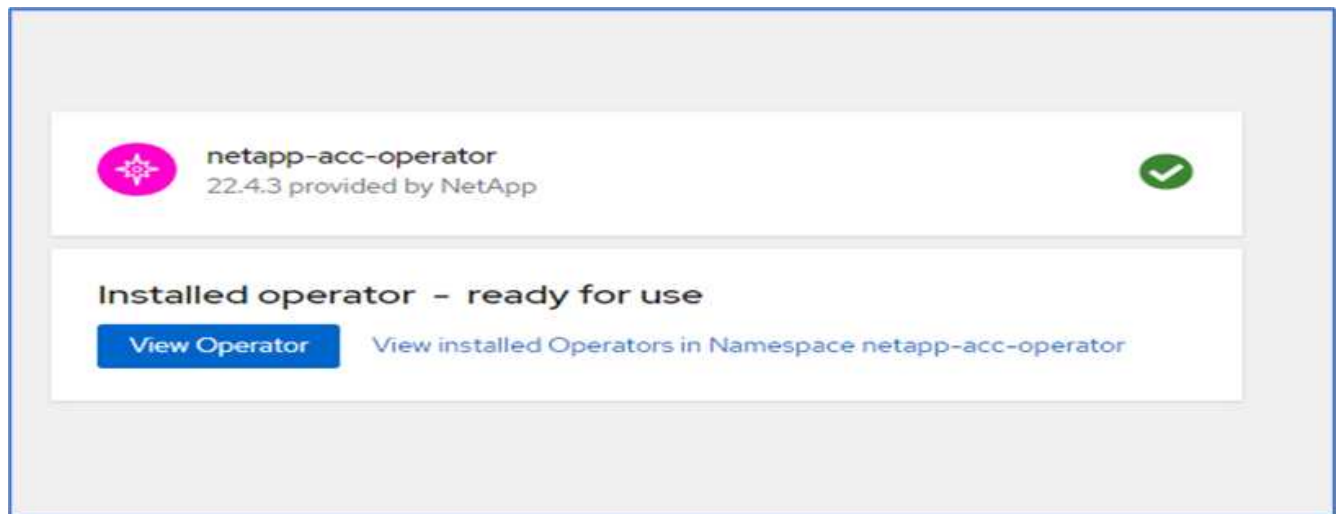
Manual approval required

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.

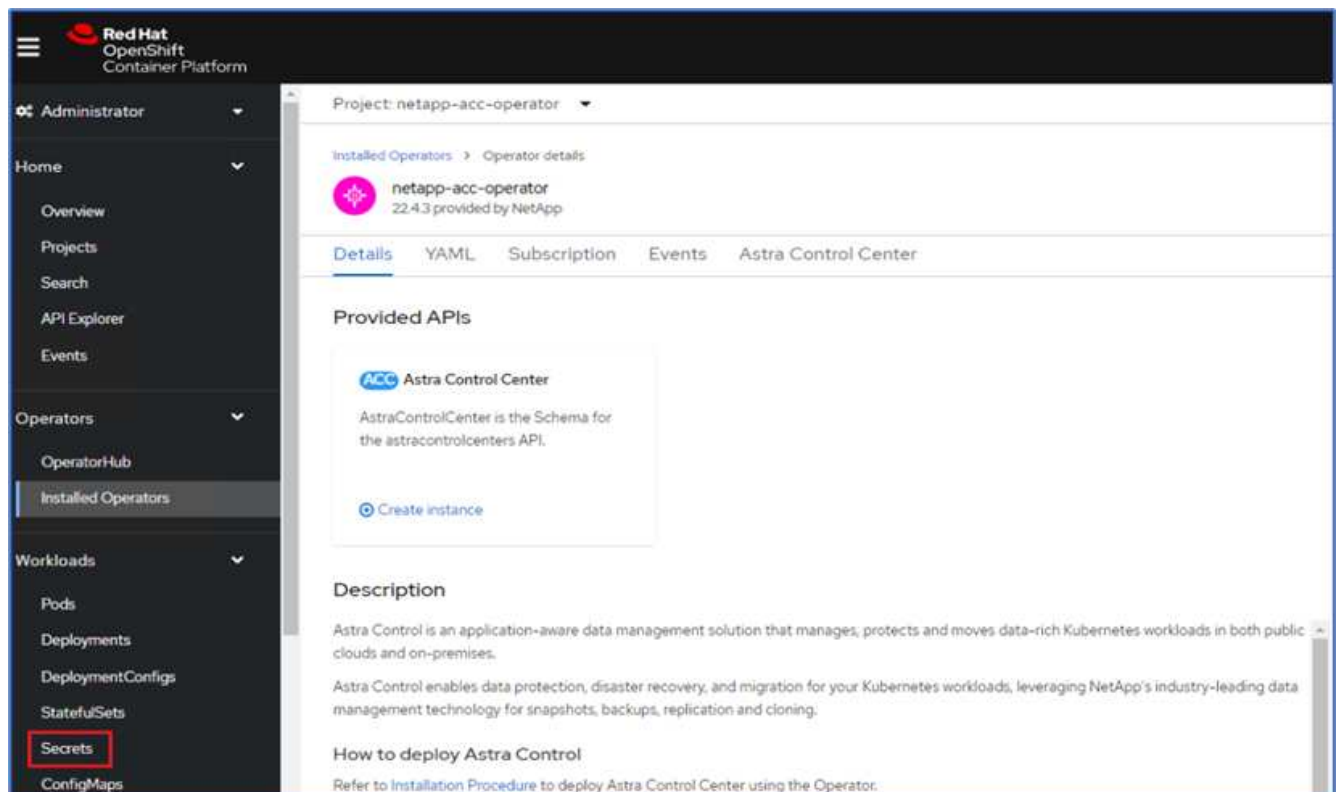
Approve **Deny**

[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 *

Unique name of the new secret.

Authentication type

Registry server address *

For example quay.io or docker.io

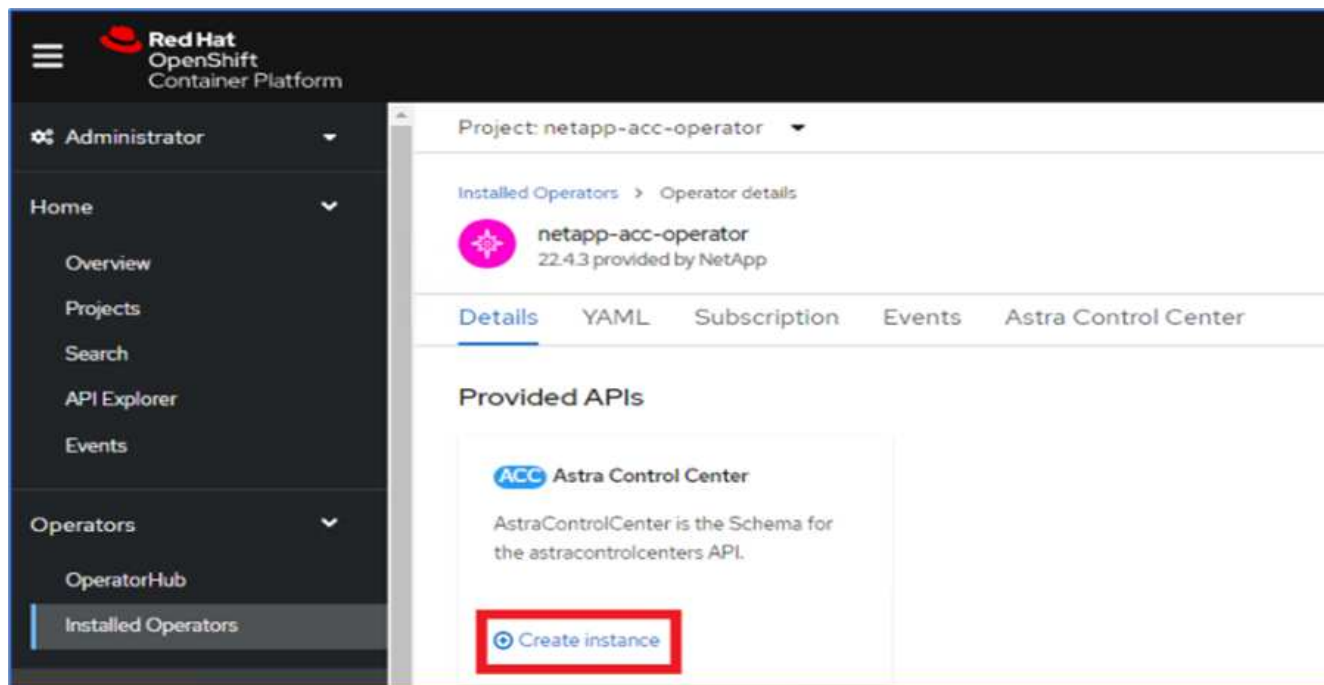
Username *

Password *

Email

[+ Add credentials](#)

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



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

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

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



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

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

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

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

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



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

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

Image Registry

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

First Name
Abhinav

The first name of the SRE supporting Astra

Astra Resources Scaler
Default

Scaling options for AstraControlCenter Resource limits.

Storage Class
ocp-nas-sc-gold

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

Crds

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

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

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

Project: netapp-acc-operator

Installed Operators > Operator details

netapp-acc-operator
22.43 provided by NetApp

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

AstraControlCenters [Create AstraControlCenter](#)

Name Search by name...

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

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

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

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

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

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

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

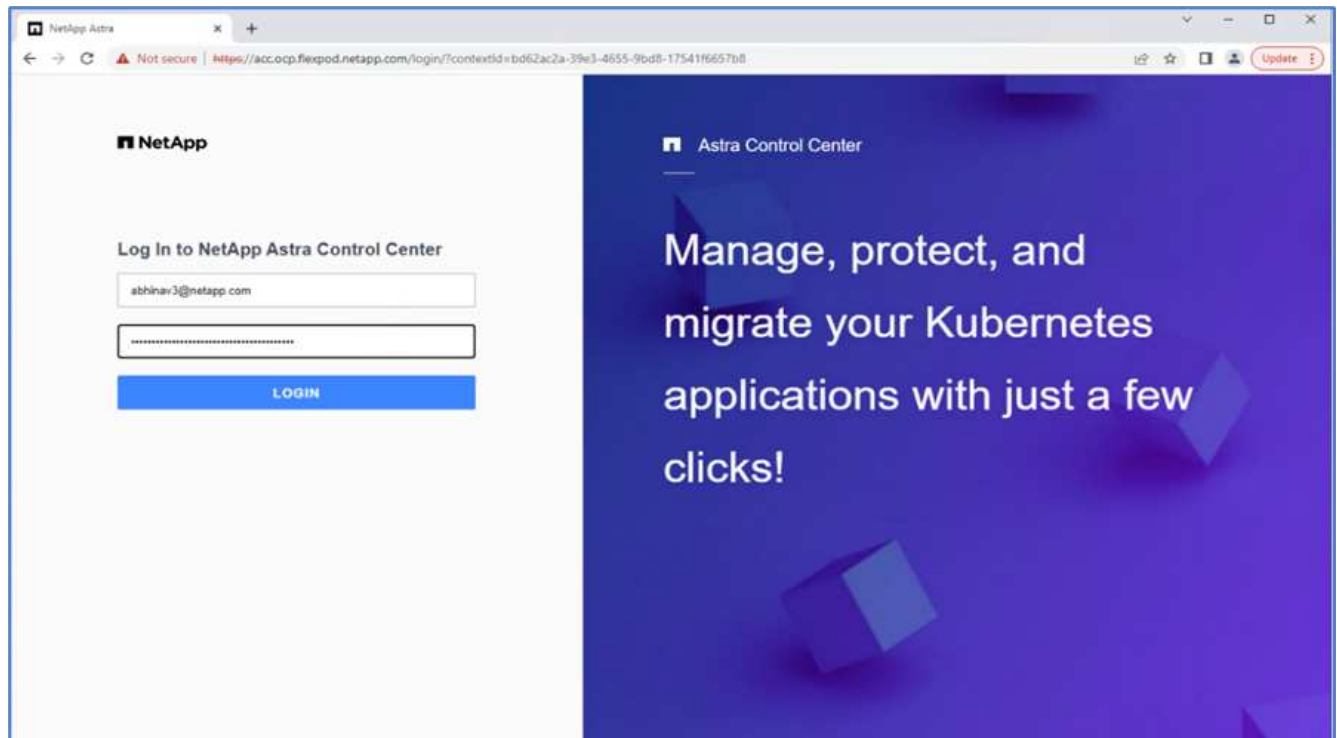


Each pod should have a status of Running. It might take several minutes before the system pods are deployed.

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

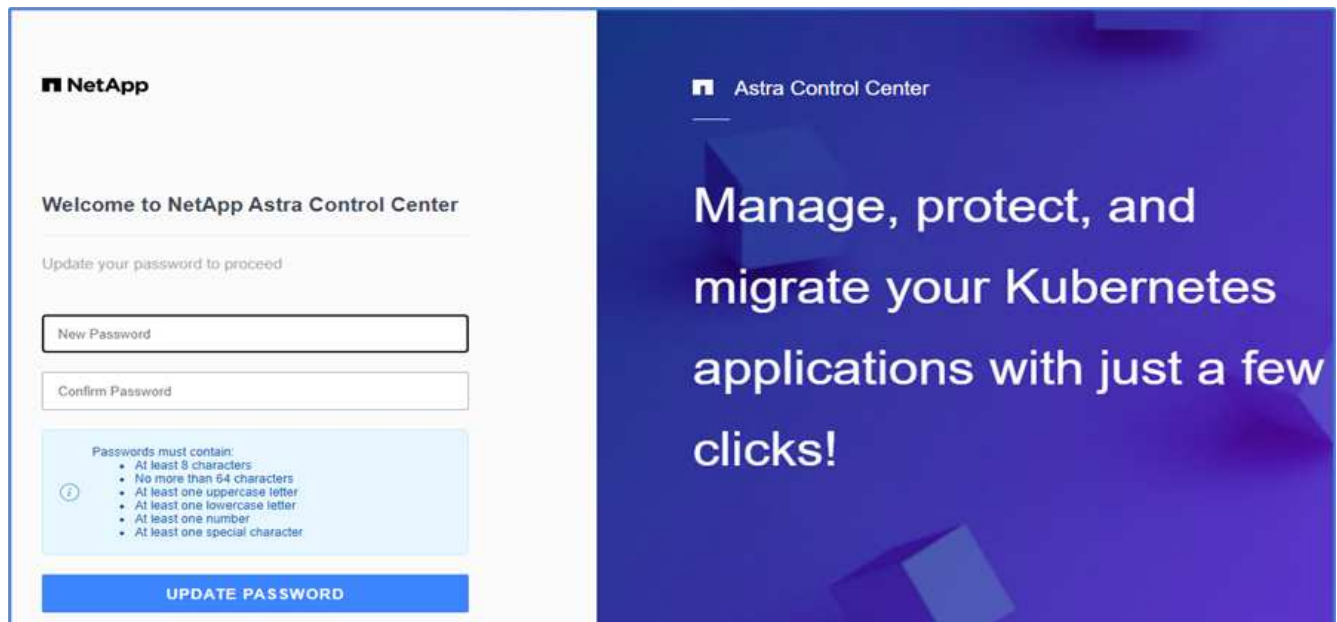
```
root@abhinav-ansible# oc get acc -o yaml -n netapp-acc-operator
```

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



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

23. Change the password and proceed.




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


Set up Astra Control Center

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


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




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



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



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




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



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

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



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

| Backends | | | | | | | |
|------------------------------|-------------|-------------------------------------|-------------------|--|----------------|--|---------|
| + Add | | <input type="text" value="Search"/> | | ★ Managed 🔍 Discovered | | 1-3 of 3 entries < > | |
| 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 | ⋮ |
| healthylife | ✓ Available | 5.16/106.42 TiB: 4.8% | Not available yet | ONTAP 9.11.1 | Not applicable | Not applicable | ⋮ |
| singlecvoaws | ✓ Available | 0.07/0.62 TiB: 11.9% | Not available yet | ONTAP 9.11.1 | Not applicable | Not applicable | ⋮ |

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

Dashboard

MANAGE YOUR APPLICATIONS

Applications

Clusters

MANAGE YOUR STORAGE

Backends

Buckets

MANAGE YOUR ACCOUNT

Account

Activity

Buckets

[+ Add](#)

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

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

Edit bucket

×

STORAGE BUCKET

Edit the access details of your existing object store bucket.

Type

Generic S3

Existing bucket name

acc-aws-bucket

Description (optional)

S3 server name or IP address

s3.us-east-1.amazonaws.com

☐ Make this bucket the default bucket for this cloud

SELECT CREDENTIALS

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

Add

Use existing

Access ID

Secret key

🔑

Credential name

EDITING STORAGE BUCKETS

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

Cancel

Update ✓



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

The Bucket state should be Healthy.

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

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

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

Next: [Solution validation overview](#).

Solution validation

Overview

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

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

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

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

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

Application recovery with remote backups

Previous: [Solution validation overview](#).

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


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

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

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



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



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

OperatorHub > Operator installation

Install Operator

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

Update channel * ⓘ

☒ alpha

Installation mode *

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

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

Installed Namespace *

NS argocd

Update approval * ⓘ

☒ Automatic

☐ Manual

Install **Cancel**

Argo CD
provided by Argo CD Community

Provided APIs

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

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

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

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

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

4. Go to the operator and click Create ArgoCD.

Project: argocd

Installed Operators > Operator details

Argo CD
0.3.0 provided by Argo CD Community

Actions

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

ArgoCDs **Create ArgoCD**

No operands found

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

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

Project: argocd ▾


[Argo CD](#) > Create ArgoCD

Create ArgoCD

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

Configure via: ☒ Form view ☐ YAML view

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



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

Name *

argocd-netapp

Labels

app=frontend

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

Project: argocd ▾

[Secrets](#) > Secret details

argocd-netapp-cluster

Managed by [ACD](#) argocd-netapp

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

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

Secret details

| | | | |
|--------------------|---|-------------|--------|
| Name | argocd-netapp-cluster | Type | Opaque |
| Namespace | argocd | | |
| Labels | Edit app.kubernetes.io/managed-by=argocd-netapp app.kubernetes.io/name=argocd-netapp-cluster app.kubernetes.io/part-of=argocd | | |
| Annotations | 0 annotations ✎ | | |
| Created at | 2 minutes ago | | |
| Owner | argocd-netapp | | |


Data

admin.password

.....

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

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



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

```

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

```

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

CREATE

CANCEL

EDIT AS YAML

GENERAL

Application Name

pricelist

Project

default

SYNC POLICY

Manual

SYNC OPTIONS

☐ SKIP SCHEMA VALIDATION

☒ AUTO-CREATE NAMESPACE

☐ PRUNE LAST

☐ APPLY OUT OF SYNC ONLY

☐ RESPECT IGNORE DIFFERENCES

PRUNE PROPAGATION POLICY: foreground

☐ REPLACE ⚠️

☐ RETRY

SOURCE

Repository URL

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

GIT ▼

Revision

main

Branches ▼

Path

pricelists/

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

DESTINATION

Cluster URL


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

URL ▼


Namespace

pricelist


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




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




13. Log into System Manager and verify the PVC.




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



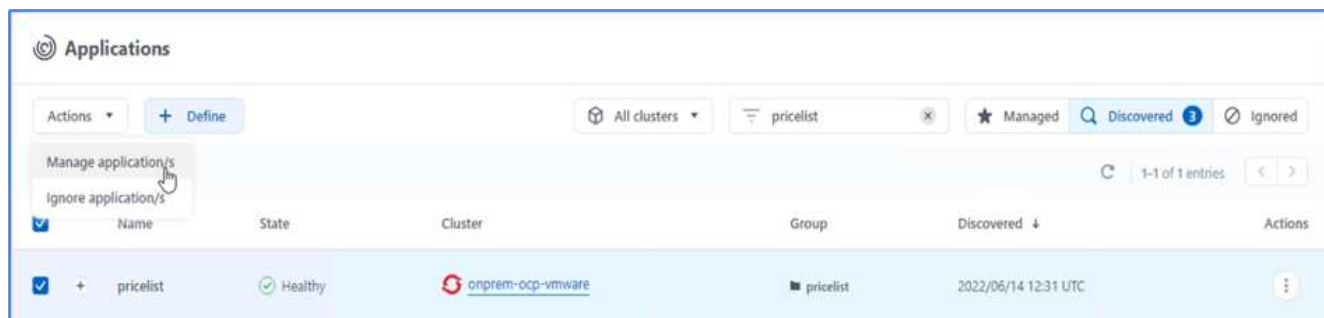
15. The Pricelist app homepage is displayed.




16. Create a few records on the web page.



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




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



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

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



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

Back up namespace application

STEP 1/2: DETAILS

✕

BACKUP DETAILS

Snapshot (optional)
pricelist-snapshot-20220614123756

Name
pricelist-backup-20220614123837

BACKUP DESTINATION

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

OVERVIEW

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

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

Cancel

Next

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

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

Copy S3 URI

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

Objects

Properties

Objects (5)
Objects are the fundamental entities stored in Amazon S3. You can use [Amazon S3 inventory](#) to get a list of all objects in your bucket. For others to access your objects, you'll need to explicitly grant them permissions. [Learn more](#)

Refresh

Copy S3 URI

Copy URL

Download

Open

Delete

Actions

Create folder

Upload

Find objects by prefix

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

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

ONTAP System Manager

Search actions, objects, and pages

Q

DASHBOARD
STORAGE

- Overview
- Volumes
- LUNs
- Consistency Groups

Storage VMs

+ Add

Infra

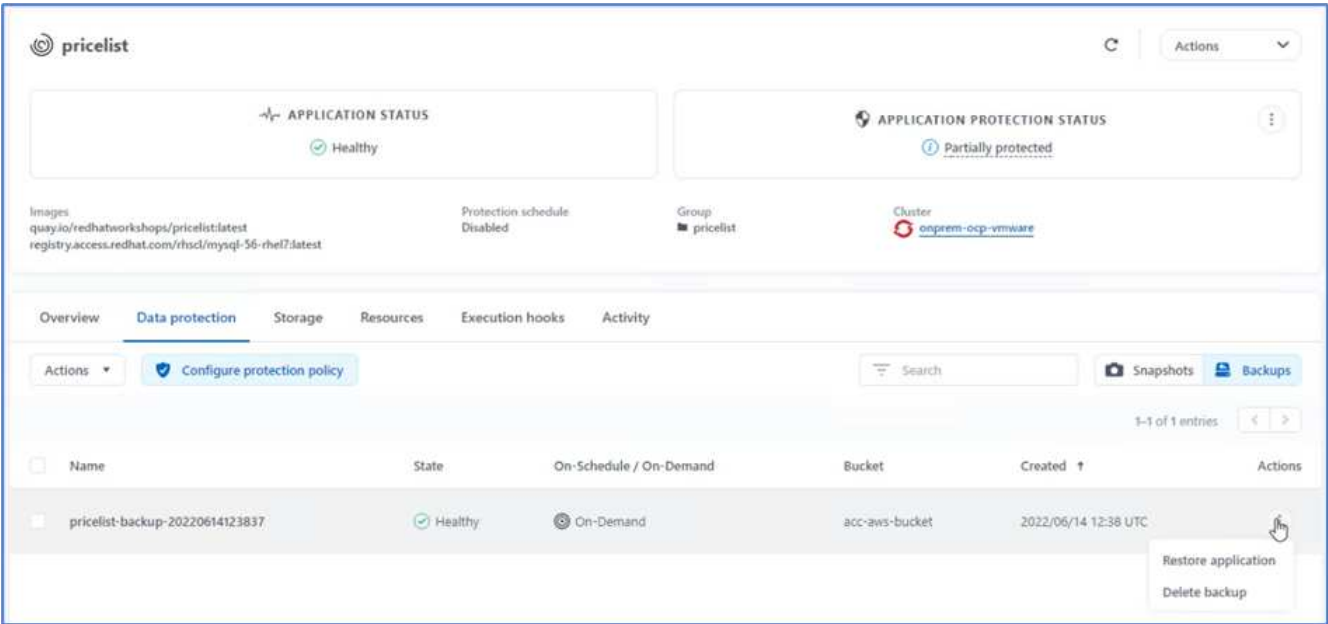
| <input type="checkbox"/> | Name | State | Subtype | Configured Protocols | IPspace |
|--------------------------|-----------|---------|---------|----------------------|---------|
| <input type="checkbox"/> | Infra_SVM | stopped | default | | Default |

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




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


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



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




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



27. Verify the same in the OpenShift web console.



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



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

Data protection with Snapshot copies and application mobility for DevTest


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

Part 1


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

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


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




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



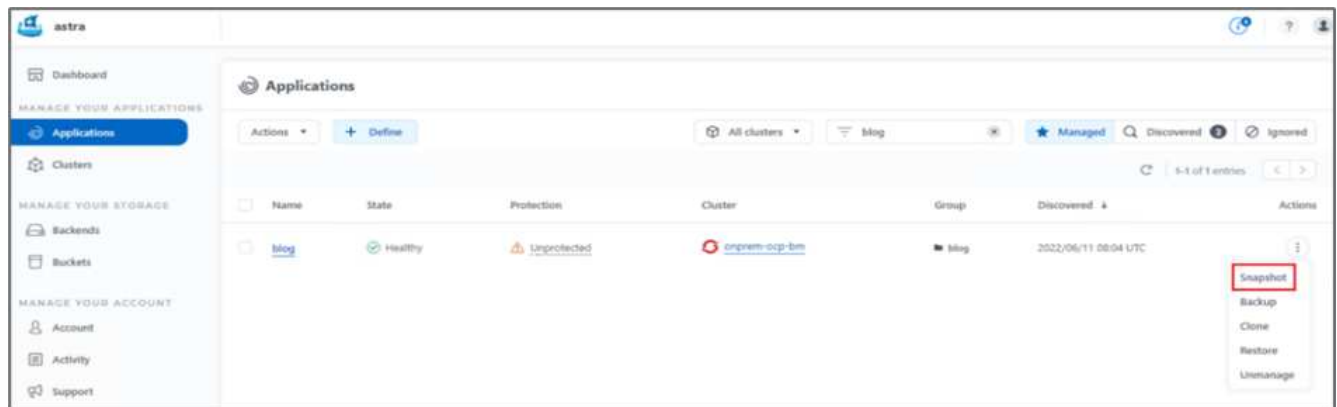
3. From the side menu, select Networking > Routes and click the URL.



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

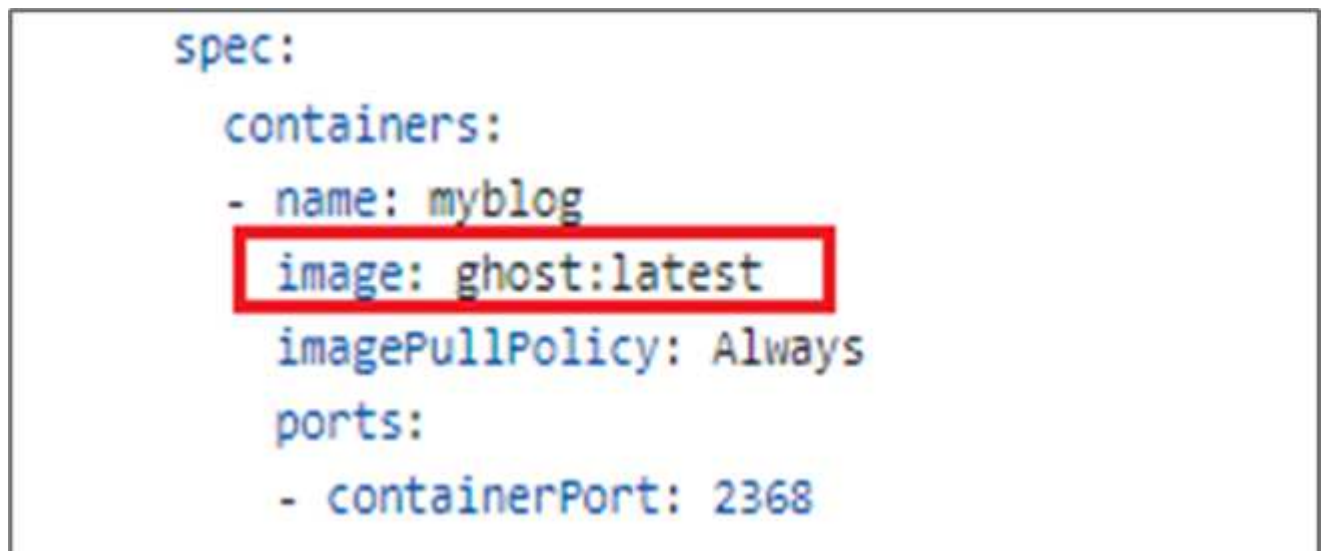


5. Go to Astra Control Center. First manage the app from the Discovered tab and then take a Snapshot copy.



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

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



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

Project: blog ▾

Pods ▸ Pod details

myblog-5f899f7b76-zv7rq CrashLoopBackOff

Details Metrics YAML Environment **Logs** Events Terminal

Log stream ended. myblog ▾ Current log ▾

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

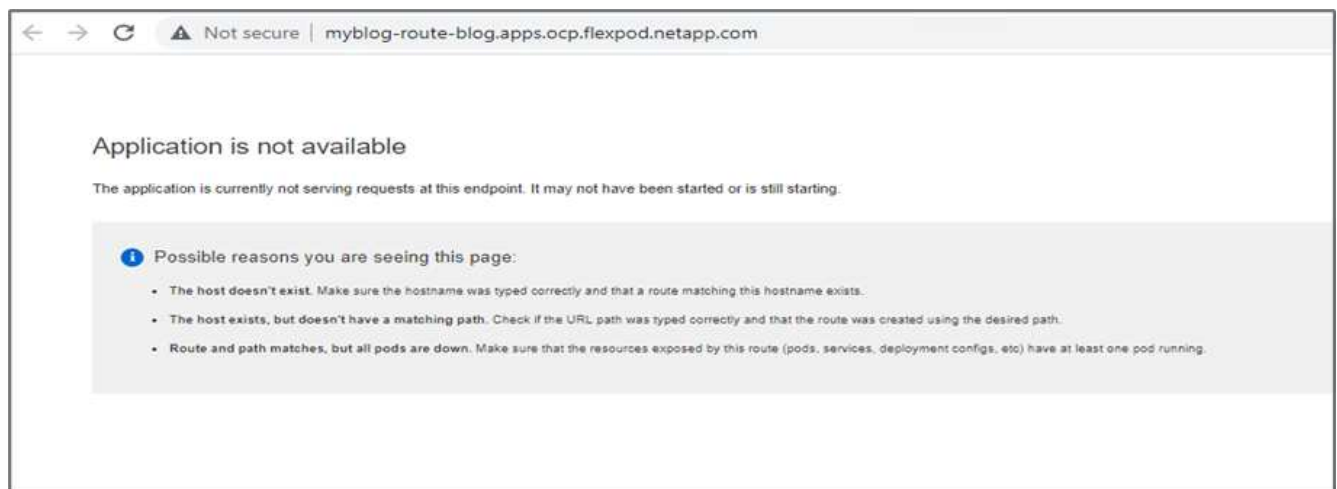
+{37m"}You must be on the latest v3.x to update across major versions - https://ghost.org/docs/update/"+"{39m
+{33m"}Run 'ghost update v3' to get the latest v3.x version, then run 'ghost update' to get to the latest.'"+"{39m

+{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
```

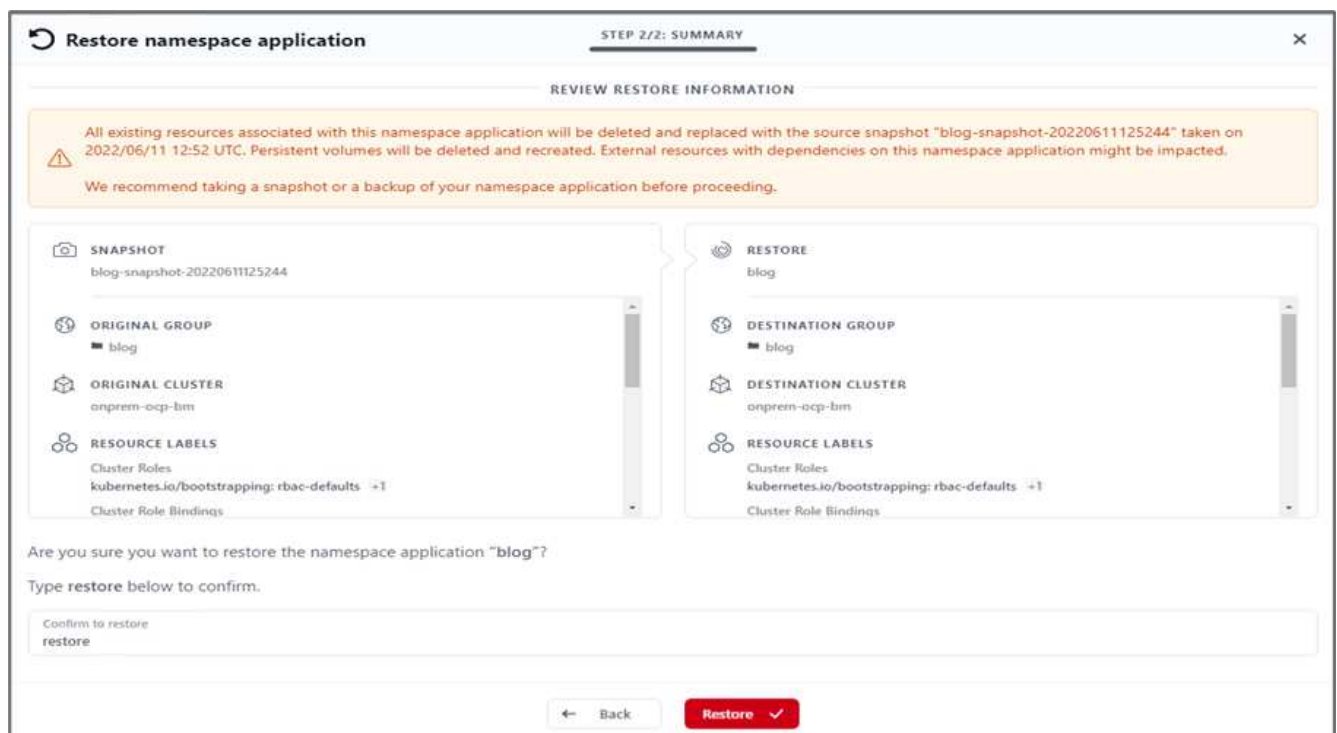
8. To confirm the unavailability of the blog site, refresh the URL.




9. Restore the app from the snapshot.



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



11. The app restore process starts immediately.



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

Applications

Actions

+ Define

All clusters

blog

Managed


Discovered

Ignored

1-1 of 1 entries

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

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



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

Part 2

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

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

Project: blog ▾


Pods > Pod details

myblog-55ffd9f658-tkbfq Running


Details Metrics YAML Environment Logs Events Terminal

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

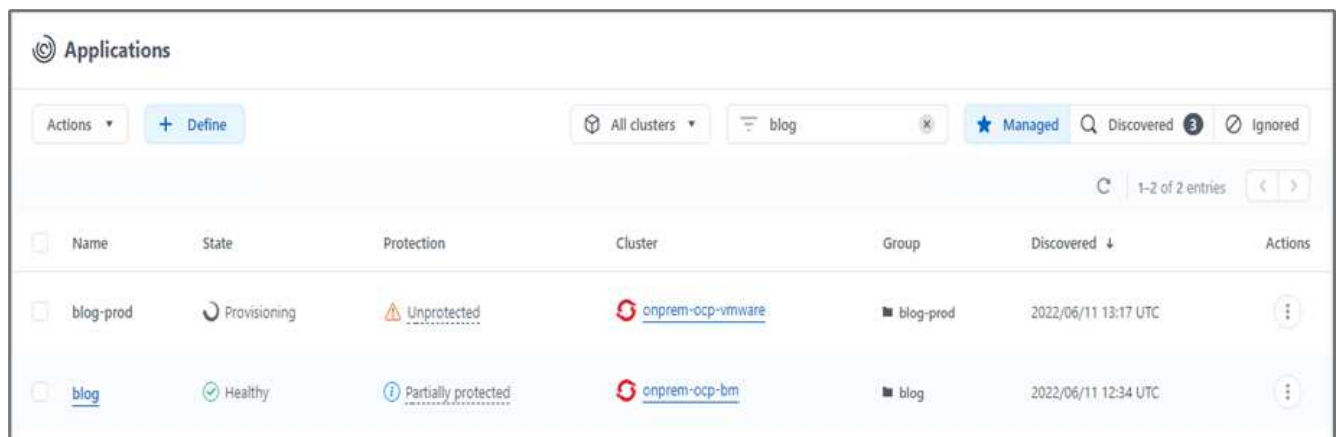
3. To verify the new theme, refresh the blog site.




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




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



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



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



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

[Next: Conclusion.](#)

Conclusion

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

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

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

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

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

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

Troubleshooting

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

Where to find additional information

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

- FlexPod Home Page

<https://www.flexpod.com>

- Cisco validated Design and deployment guides for FlexPod

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

- FlexPod deployment with Infrastructure as code for VMware using Ansible

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_m6_esxi7u2.html#AnsibleAutomationWorkflowandSolutionDeployment

- 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

- 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

- NetApp Cloud Volumes ONTAP

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:
 - a. Two CPUs
 - b. 8GB of RAM
 - c. 100GB of hard disk space
 - d. A network that can reach resources in the FlexPod Datacenter and the Cloud Insights server through an SSL connection on port 443.
 - e. 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 | | |
| > 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 | <input checked="" type="checkbox"/> | Connect... |
| > 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.

MANUAL PARTITIONING
RED HAT ENTERPRISE LINUX 7.6 INSTALLATION

Done

us

Help!

New Red Hat Enterprise Linux 7.6 Installation

DATA

/opt/netapp25 GiB>

rhel-opt_netapp

/var/log/netapp25 GiB

rhel-var_log_netapp

SYSTEM

/boot1024 MiB

sda1

/40 GiB

rhel-root

swap8064 MiB

rhel-swap

+

-

↺

AVAILABLE SPACE

1140.97 MiB

TOTAL SPACE

100 GiB

[1 storage device selected](#)

rhel-opt_netapp

Mount Point:

/opt/netapp

Device(s):

VMware Virtual disk (sda)

Desired Capacity:

25 GiB

Modify...

Device Type:

LVM

☐ Encrypt

File System:

xfs

☒ Reformat

Volume Group

rhel (4096 KiB free)

Modify...


Label:

Name:


opt_netapp

Reset All

- 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.



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



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.

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.


Installation Instructions [Need Help?](#)

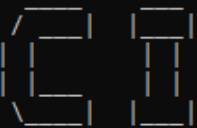
- 2 Paste the snippet into a bash shell to run the installer.
- 3 Please ensure you have copied and pasted the snippet into the bash shell.

[illegible]

309

```


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.

| | | | | | |
|---|--------------------|--------------|--------------------------------|------|-----------|
| NetApp PCS Sa... / Admin / Acquisition Units / NetApp-AU | | | | | Restart ▼ |
| Summary | | | | | |
| Name NetApp-AU | IP 10.1.156.115 | Status OK | Last Reported 9 minutes ago | Note | |

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


 Cloud Volumes ONTAP



 Data ONTAP 7-Mode



 ONTAP Data Management
 Software


 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.





NetApp
 ONTAP Data Management Software

Configure Collector

Add credentials and required settings [Need Help?](#)



 Configuration: Successfully pinged 192.168.156.50.
 Configuration: Successfully executed test command on device.

| | |
|--|---|
| Name ⓘ <input type="text" value="FlexPod Datacenter"/> | Acquisition Unit <input type="text" value="NetApp-AU"/> |
| NetApp Management IP Address <input type="text" value="192.168.156.50"/> | User Name <input type="text" value="admin"/> |
| Password <input type="password" value="*****"/> | |
| <input type="button" value="Complete Setup"/> <input type="button" value="Test Connection"/> | |

 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.

NetApp PCS Sa... / Admin / Data Collectors

Data Collectors 08Acquisition Units 08

Data Collectors (1)

+ Data Collector

Bulk Actions

FlexPod

| <input type="checkbox"/> | Name | Status | Type | Acquisition Unit | IP | Impact ↓ | Last Acquired |
|--------------------------|--------------------|----------------|---------------------------------------|------------------|----------------|----------|----------------|
| | FlexPod Datacenter | All successful | NetApp ONTAP Data Management Software | NetApp-AU | 192.168.156.50 | | 10 minutes ago |

- 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 FlexPod Datacenter | Type NetApp ONTAP Data Management Software | Types of Data Collected Inventory, Performance | Performance Recent Status Success | Note |
| 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)

Filter...

| Device ↑ | Name | IP |
|----------|-----------|----------------|
| 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... / **aa14-a800** Last 3 Hours Acquired 13 minutes ago, 12:51 PM Edit

Storage Summary

| | | | |
|--------------------------------------|---|---|---|
| Model: AFF-A800 | IP: 192.168.156.50 | IOPS - Total: 4,972.70 IO/s | Performance Policies: |
| 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 | |
| Serial Number: 1-80-000011 | Latency - Total: 0.05 ms | FC Fabrics Connected: 0 | |

User Data

Note: Testing annotations, Testing rules

Expert View

Latency - Total (ms)

Monday 10/19/2020 10:36:38 AM
aa14-a800: 0.04 ms


Resource

aa14-a800

Top Correlated

| | |
|-------------|-----|
| aa14-a800-2 | 79% |
| aa14-a800-1 | 23% |

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



- 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.

Active IQ Active IQ Digital Advisor Discovery Dashboard Asset Insights ...

Home > Cisco Systems Inc. > CISCO SYSTEMS - RTP - BUILDING 9 > aa14-a800

The Risk Acknowledgment feature has been migrated to Active IQ Digital Advisor. [Click here](#) to view and acknowledge risks.

Health Security Vulnerability Proactive Remediation Best Practices Performance System Health Storage Virtual Machine Health Health Trending

High Medium Low

| Ack | Node | Serial No | Impact Level | Public | Category | Risk | Details | Corrective Action |
|-----|-------------|--------------|--------------|--------|--------------|---|--|---------------------------------|
| | aa14-a800-2 | 941834000459 | High | No | ONTAP | A network interface (LIF) using a port on a X1116A, X1146A or X91146A NIC might not fail over to an alternate port. | 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". Potential Impact: Any network interface (LIF) using the port does not fail over to an alternate port in the event of failure. | Bug ID: 1322372 |
| | aa14-a800-2 | 941834000459 | High | Yes | FAS Hardware | On AFF A800 systems an erroneous 'Critical High' sensor reading can result in a system shutdown. | This AFF-A800 system is running BMC firmware 10.3 which is susceptible to bug 1279964. Potential Impact: System disruption caused by an erroneous 'Critical High' sensor reading. | Bug ID: 1279964 |
| | aa14-a800-2 | 941834000459 | High | Yes | ONTAP | AFF systems running an unfixed version of ONTAP with data compaction enabled and host services over FCP, iSCSI or NVMe can experience a disruption in service due to BUG 1273955. | 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. Potential Impact: The system may experience performance degradation and possible panic. | Bug ID: 1273955 |
| | aa14-a800-2 | 941834000459 | High | Yes | ONTAP | ONTAP 9.7 running on an All-Flash FAS (AFF) system having SAN workload might cause a storage controller disruption. | 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. Potential Impact: The system may experience a disruption. | KB ID: SU426 |
| | aa14-a800-1 | 941834000183 | High | No | ONTAP | A network interface (LIF) using a port on a X1116A, X1146A or X91146A NIC might not fail over to an alternate port. | 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". | Bug ID: 1322372 |

1 - 17 of 17 results

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.




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.




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.



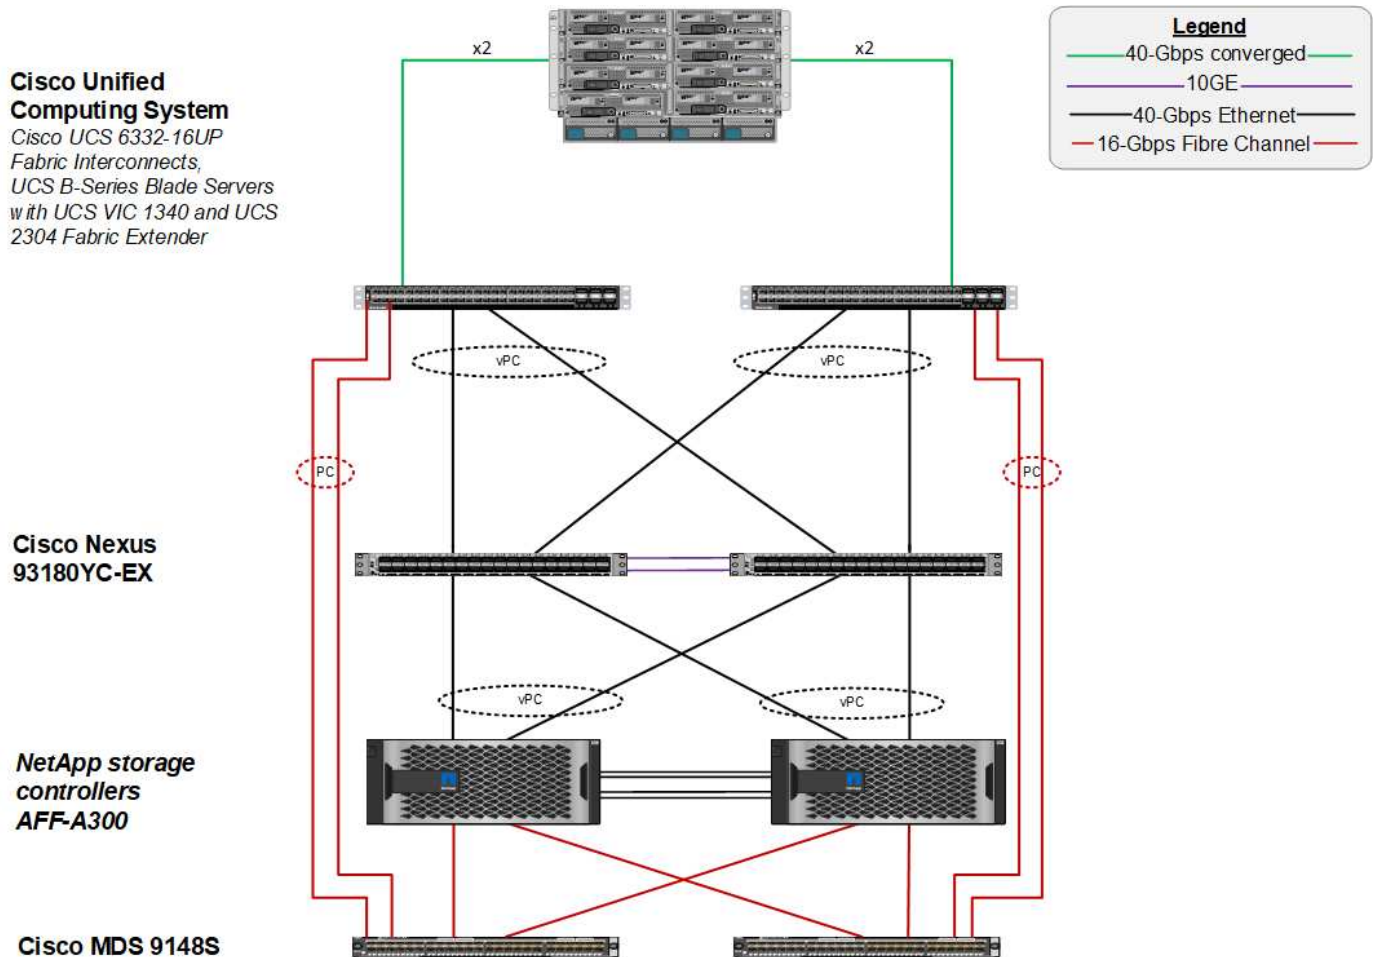
FlexPod architecture

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.

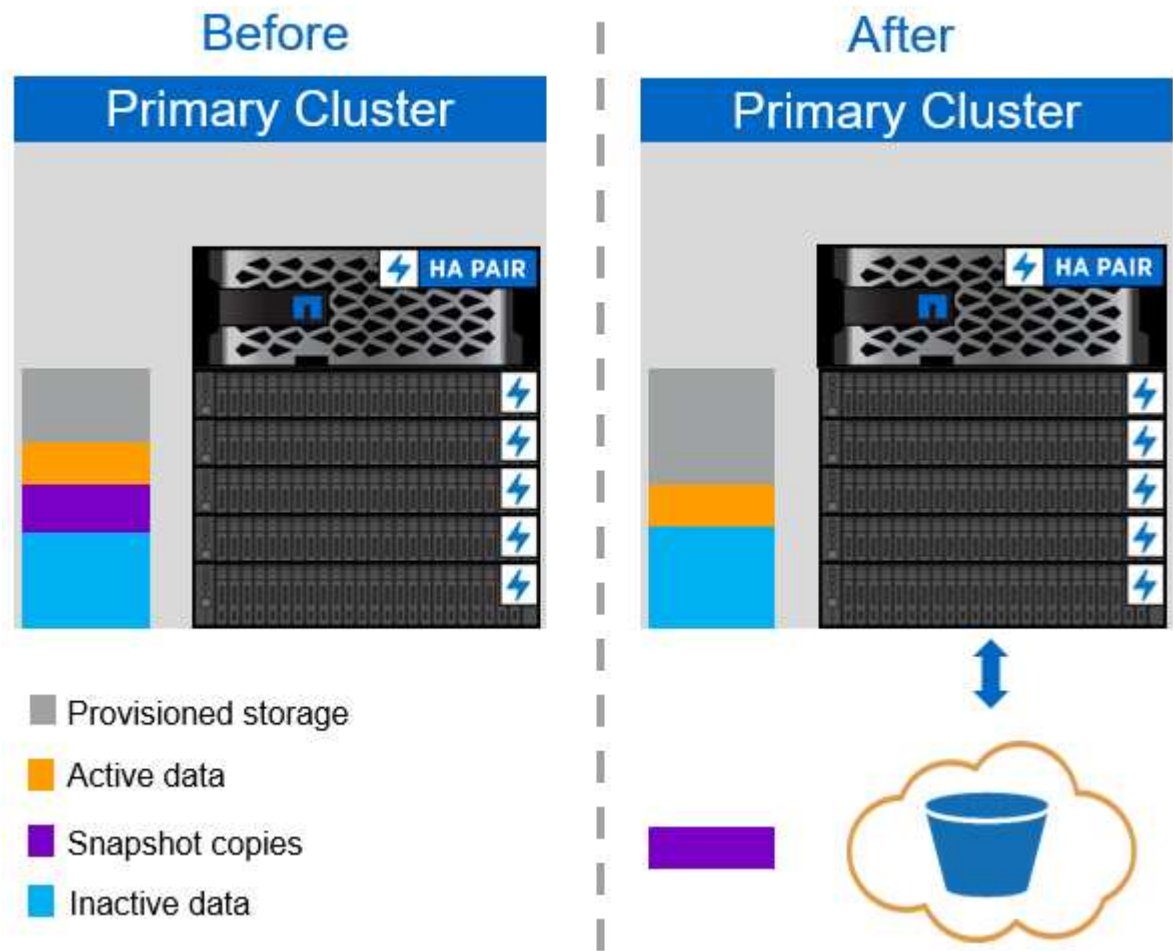
If performance tier utilization is at greater than 70% capacity, cold data is read directly from the cloud tier without being written back to the performance tier. By preventing cold data write-backs on heavily used aggregates, FabricPool preserves the aggregate for active data.

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.



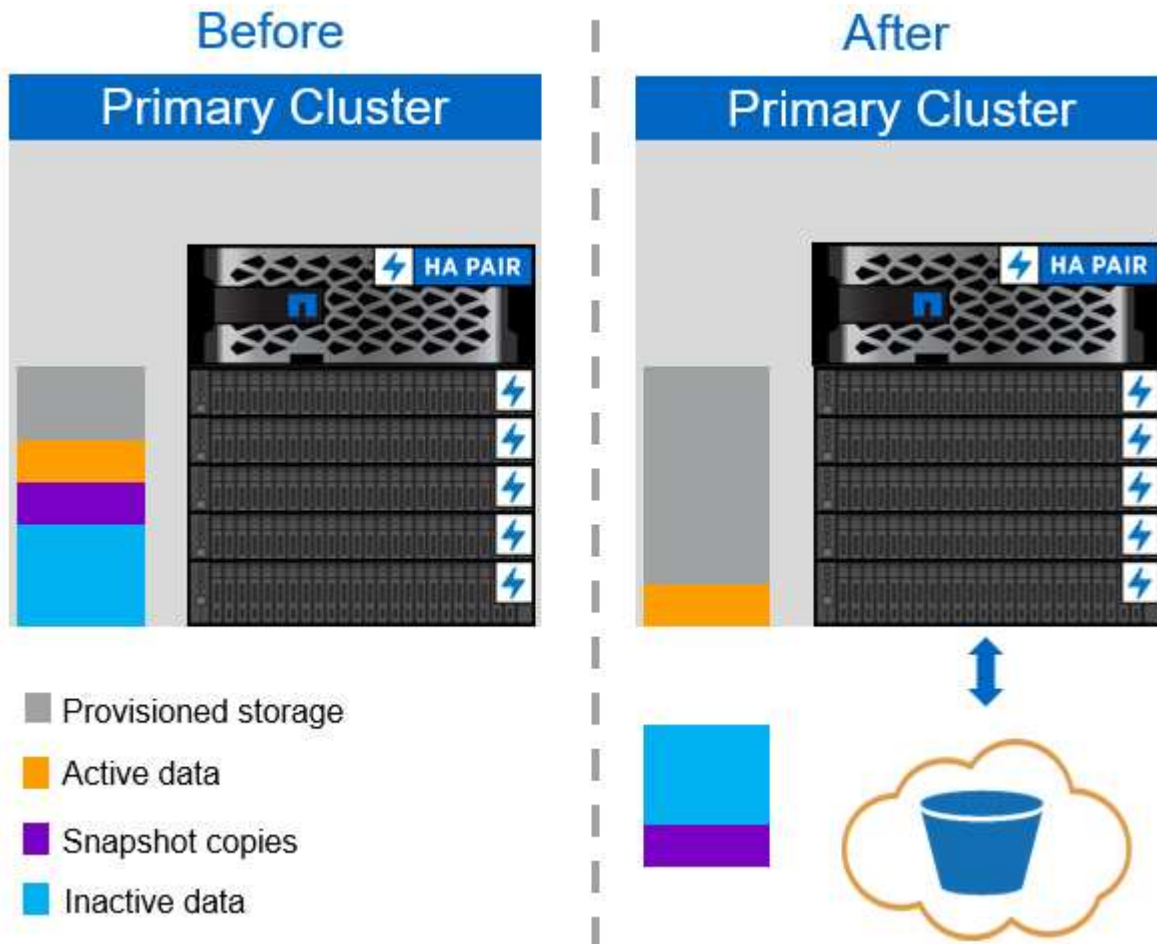
Auto tiering policy

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 | min | Latency (ms) max | avg | Throughput |
|-----|-------|--------|--------|-----|---------------------|------|------------|
| 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 6.7U2 | 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. The left sidebar has a menu with 'Configuration' and 'Licenses' highlighted with red boxes. The main area displays the 'Licenses' section with a table of installed licenses. A dialog box titled 'Add License Packages' is open, showing a text input field for license keys and a 'License Files' section with a 'Choose Files' button.

| Package | Entitlement Risk | Description |
|-----------------------------------|------------------|------------------------|
| (DEPRECATED)-Cluster Base License | -NA- | Installed on a cluster |
| Trusted Platform Module License | -NA- | No License Available |
| FabricPool License | -NA- | Installed on a cluster |
| NFS License | Medium risk | Medium risk |
| CIFS License | | |
| ISCSI License | | |
| FCP License | | |
| SnapRestore License | | |
| SnapMirror License | | |
| FlexClone License | | |
| SnapVault License | | |
| SnapLock License | | |

Add License Packages

Enter comma separated license keys

License Files

Browse to select a file... Choose Files

License files are required for features that use capacity based licenses. [Know more](#)

Add Cancel

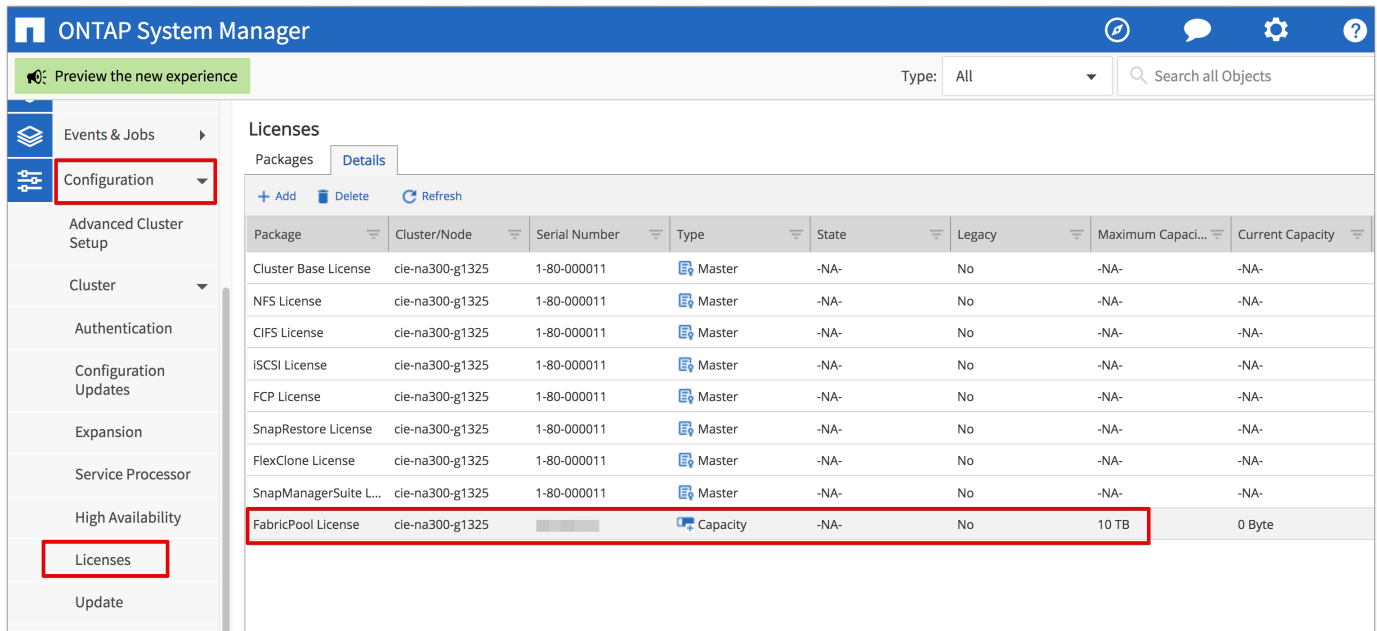
License capacity

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- |
| FabricPool License | cie-na300-g1325 | | Capacity | -NA- | No | 10 TB | 0 Byte |

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.

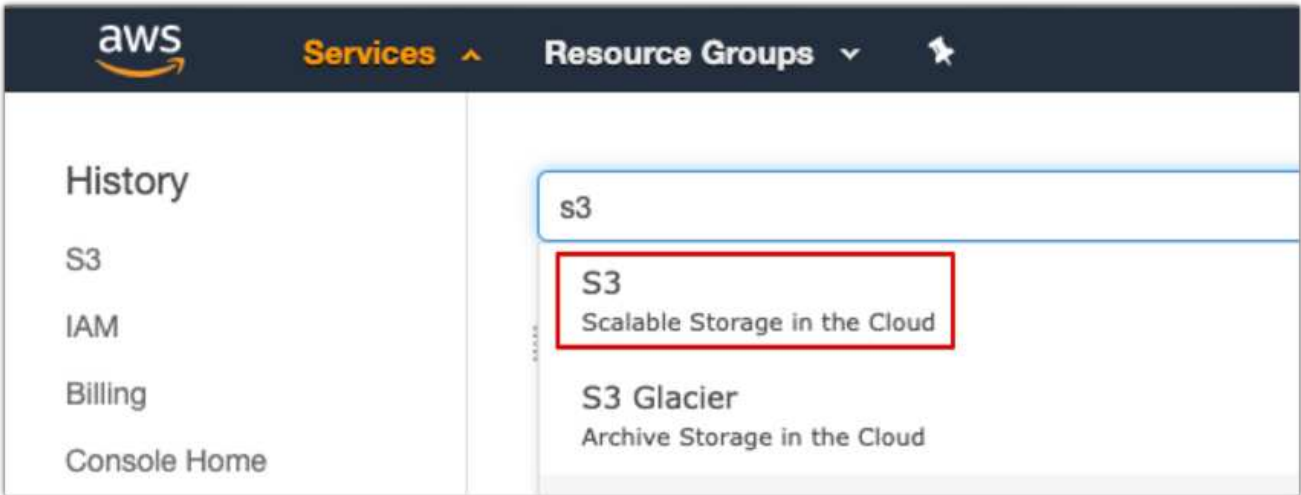
FabricPool supports the attachment of one bucket per aggregate, as illustrated in the following figure. A single bucket can be attached to a single aggregate, and a single bucket can be attached to multiple aggregates. However, a single aggregate cannot be attached to multiple buckets. Although a single 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.

The screenshot shows the 'Create bucket' wizard in the AWS Management Console. The title bar is blue with the text 'Create bucket' and a close button. Below the title bar is a progress bar with four steps: 1. Name and region (active), 2. Configure options, 3. Set permissions, and 4. Review. The main content area is dark blue. Under the 'Name and region' section, there is a 'Bucket name' field with the value 'flexpod-fp-bk-1' and a 'Region' dropdown menu set to 'US East (Ohio)'. Below these is a section titled 'Copy settings from an existing bucket' with a dropdown menu showing 'Select bucket (optional)' and '4 Buckets'. At the bottom of the form, there are three buttons: 'Create' (highlighted with a red box), 'Cancel', and 'Next'.

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

Name

Server Name (FQDN)

Access Key ID

Secret Key

 Container Name

 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.

CHANGE VOLUME TIERING POLICY

Select the tiering policy that you want to apply for the selected volume.

| Volume Name | Tiering Policy |
|---------------|----------------|
| affa3..._fp_1 | auto |

Tiering Policy

- snapshot-only
- none
- auto
- all

[Learn more about tiering policies.](#)

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

- NetApp Product Documentation

<https://docs.netapp.com>

- TR-4036: FlexPod Datacenter Technical Specification

<https://www.netapp.com/us/media/tr-4036.pdf>

FlexPod Datacenter with IBM Cloud Private

Sreenivasa Edula, Cisco

Thanachit Wichianchai, IBM

Jacky Ben-Bassat, IBM Global Alliance, NetApp

IBM Cloud Private (ICP) is an on-premises platform for developing and managing containerized applications for cloud-native and application-modernization use cases. It is an integrated environment built on Kubernetes as its container orchestration, and includes a private image repository for Docker containers, a management console, a

monitoring framework, many open source-based and IBM containerized applications, and more. Combining ICP with FlexPod, the converged infrastructure from Cisco and NetApp, can simplify the deployment and the management of your infrastructure. You can also benefit from improved storage efficiency, better data protection, lower risk, and the flexibility to scale this highly available enterprise-grade infrastructure stack to accommodate new business requirements and other changes over time.

[FlexPod Datacenter with IBM Cloud Private](#)

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](#)

FlexPod Datacenter for Multicloud with Cisco CloudCenter and NetApp Data Fabric

Haseeb Niazi, Cisco
David Arnette, NetApp

This document provides in-depth configuration and implementation guidelines for setting up FlexPod Datacenter for Hybrid Cloud. The following design elements distinguish this version of FlexPod from previous models:

- Integration of Cisco CloudCenter with FlexPod Datacenter with ACI as the private cloud
- Integration of Cisco CloudCenter with Amazon Web Services (AWS) and Microsoft Azure Resource Manager (MS Azure RM) public clouds
- Providing secure connectivity between the FlexPod data center and the public clouds for secure traffic between virtual machines (VMs)
- Providing secure connectivity between the FlexPod data center and NetApp Private Storage (NPS) for data replication traffic
- Ability to deploy application instances in either public or private clouds and to make up-to-date application data available to these instances through orchestration driven by Cisco CloudCenter
- Setting up, validating, and highlighting operational aspects of a development and test environment in this new hybrid cloud mode.

[FlexPod Datacenter for Multicloud with Cisco CloudCenter and NetApp Data Fabric](#)

Enterprise Databases

SAP

Introduction to SAP on FlexPod

The FlexPod platform is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp storage controllers.

FlexPod is a suitable platform for running SAP applications, and the solutions provided here allows you to quickly and reliably deploy SAP HANA with a model of tailored datacenter integration. FlexPod delivers not only a baseline configuration, but also the flexibility to be sized and optimized to accommodate many different use cases and requirements.

FlexPod Datacenter for SAP solution using FibreChannel SAN with Cisco UCS Manager 4.0 and NetApp ONTAP 9.7

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

This document describes the Cisco and NetApp FlexPod Datacenter with NetApp ONTAP 9.7 on NetApp AFF A400 storage and Cisco UCS Manager unified software release 4.1(1) with second-generation Intel Xeon Scalable Processors for SAP HANA in particular.

FlexPod Datacenter with NetApp ONTAP 9.7 and Cisco UCS unified software release 4.1(1) is a pre-designed, best-practice datacenter architecture built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus 9000 family of switches, MDS 9000 multilayer fabric switches, and NetApp AFF A-Series storage arrays running ONTAP 9.7 storage OS.

[FlexPod Datacenter for SAP solution using FibreChannel SAN with Cisco UCS Manager 4.0 and NetApp ONTAP 9.7](#)

SAP Non-HANA with SQL white paper - Design

The current IT industry is witnessing a dramatic transformation in data center solutions. In recent years, there has been a considerable interest in prevalidated and engineered data center solutions. Introduction of virtualization technology in critical areas has had a major impact on the design principles and architecture of these solutions. It has allowed many applications running on bare-metal systems to migrate to new virtualized integrated solutions. FlexPod is one such prevalidated and engineered data center solution designed to address the rapidly changing needs of IT departments. Cisco and NetApp have partnered to deliver FlexPod, which uses best-in-class computing, networking, and storage components as the foundation for a variety of enterprise workloads, including databases, enterprise resource planning (ERP), customer relationship management (CRM), and web applications.

The consolidation of IT applications, particularly databases, has generated considerable interest in recent years. The most widely adopted and deployed database platform over the past several years is Microsoft SQL

Server. SQL Server databases frequently have become subject to database sprawl, leading to IT challenges such as underutilized servers, incorrect licensing, security concerns, management concerns, and huge operational costs. Therefore, SQL Server databases are good candidates for consolidation on a more robust, flexible, and resilient platform. This document discusses a FlexPod reference architecture for deploying and consolidating SQL Server databases.

[SAP Non-HANA with SQL white paper - Design](#)

FlexPod Datacenter for SAP Solution with Cisco UCS third-generation fabric and NetApp AFF A-Series

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

This document describes the deployment methodology of Cisco and NetApp FlexPod Datacenter for SAP HANA based on second-generation Intel Xeon Scalable Processors supported Cisco UCS Computing System (Cisco UCS).

Cisco UCS Manager (UCSM) 4.0(4) provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 and 6454), 2200/2300 series IOM, Cisco UCS B-Series Blade, and Cisco UCS C-Series Rack Formfactor servers. FlexPod Datacenter with Cisco UCS unified software release 4.0(4d) and NetApp ONTAP 9.6, is a pre-designed, best-practice data center architecture built on the Cisco UCS, the Cisco Nexus 9000 family of switches, and NetApp AFF A-Series storage arrays.

[FlexPod Datacenter for SAP Solution with Cisco UCS third-generation Fabric and NetApp AFF A-Series](#)

FlexPod Datacenter for SAP solution using FibreChannel SAN with Cisco UCS Manager 4.0 and NetApp ONTAP 9.7 - Design

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

Cisco and NetApp have partnered to deliver a series of FlexPod solutions that enable strategic data center platforms. The FlexPod solution delivers an integrated architecture that incorporates computing, storage, and network design best practices, thereby minimizing IT risks by validating the integrated architecture to ensure compatibility between various components. The solution also addresses IT pain points by providing documented design guidance, deployment guidance, and support that can be used in various stages (planning, designing, and implementation) of a deployment.

[FlexPod Datacenter for SAP solution using FibreChannel SAN with Cisco UCS Manager 4.0 and NetApp ONTAP 9.7 - Design](#)

FlexPod Datacenter for SAP solution with Cisco ACI, Cisco UCS Manager 4.0, and NetApp AFF A-Series - Design

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

This document describes Cisco ACI integrated FlexPod solution as a validated approach for deploying SAP HANA Tailored Data Center Integration (TDI) environments. This

validated design provides guidelines and a framework for implementing SAP HANA with best practices from Cisco and NetApp.

The recommended solution architecture is built on the Cisco Unified Computing System (Cisco UCS) using a unified software release to support Cisco UCS hardware platforms that include the following components:

- Cisco UCS B-Series blade servers and Cisco UCS C-Series rack servers configurable with the Intel Optane Data Center Persistent Memory Module (DCPMM) option
- Cisco UCS 6400 series Fabric Interconnects
- Cisco Nexus 9000 Series Leaf and Spine switches
- NetApp All Flash series storage arrays

Additionally, this document provides validations for both Red Hat Enterprise Linux and SUSE Linux Enterprise Server for SAP HANA.

[FlexPod Datacenter for SAP solution with Cisco ACI, Cisco UCS Manager 4.0, and NetApp AFF A-Series - Design](#)

FlexPod Datacenter for SAP with Cisco ACI, Cisco UCS Manager 4.0, and NetApp AFF A-Series - Deployment

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

This document describes the architecture and deployment procedures for the SAP HANA Tailored DataCenter Integration option on FlexPod infrastructure, which is composed of:

- Cisco UCS Computing System (Cisco UCS) supported by second-generation Intel Xeon Scalable Processors.
- Switching products that leverage Cisco Application Centric Infrastructure (ACI).
- NetApp A-series AFF arrays.

The intent of this document is to show the detailed configuration steps for SAP HANA deployment

[FlexPod Datacenter for SAP with Cisco ACI, Cisco UCS Manager 4.0, and NetApp AFF A-Series - Deployment](#)

FlexPod Datacenter for SAP Solution with Cisco UCS Manager 4.0 and NetApp AFF A-Series - Design

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

This document describes the Cisco and NetApp FlexPod solution, which is a validated approach for deploying SAP HANA Tailored Data Center Integration (TDI) environments. This validated design provides guidelines and a framework for implementing SAP HANA with best practices from Cisco and NetApp.

FlexPod is a leading integrated infrastructure that supports a broad range of enterprise workloads and use cases. This solution allows you to quickly and reliably deploy SAP HANA with a model of a tailored data center

integration mode.

[FlexPod Datacenter for SAP Solution with Cisco UCS Manager 4.0 and NetApp AFF A-Series - Design](#)

FlexPod Datacenter for SAP solution with Cisco ACI on Cisco UCS M5 servers with SLES 12 SP3 and RHEL 7.4

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

This document describes the architecture and deployment procedures for SAP HANA Tailored DataCenter Integration option on FlexPod infrastructure composed of Cisco compute and switching products that leverage Cisco Application Centric Infrastructure (ACI) - the industry-leading software-defined networking solution (SDN) - along with NetApp A-series AFF arrays. The intent of this document is to show the design principles with the detailed configuration steps for SAP HANA deployment.

[FlexPod Datacenter for SAP solution with Cisco ACI on Cisco UCS M5 servers with SLES 12 SP3 and RHEL 7.4](#)

FlexPod Datacenter for SAP Solution with IP-based storage using NetApp AFF A-Series and Cisco UCS Manager 3.2

Shailendra Mruthunjaya, Cisco
Ralf Klahr, Cisco
Marco Schoen, NetApp

The reference architecture detailed in this document highlights the resiliency, cost benefit, and ease of deployment of an IP-based storage solution. A storage system capable of serving multiple protocols across a single interface allows for customer choice and investment protection because it truly is a wire-once architecture. The solution is designed to host scalable SAP HANA workloads.

[FlexPod Datacenter for SAP Solution with IP-based storage using NetApp AFF A-Series and Cisco UCS Manager 3.2](#)

FlexPod Datacenter for SAP solution using FibreChannel SAN with Cisco UCS Manager 4.0 and NetApp ONTAP 9.7

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

This document describes the Cisco and NetApp FlexPod Datacenter with NetApp ONTAP 9.7 on NetApp AFF A400 storage and Cisco UCS Manager unified software release 4.1(1) with second-generation Intel Xeon Scalable Processors for SAP HANA in particular.

FlexPod Datacenter with NetApp ONTAP 9.7 and Cisco UCS unified software release 4.1(1) is a pre-designed, best-practice datacenter architecture built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus 9000 family of switches, MDS 9000 multilayer fabric switches, and NetApp AFF A-Series storage arrays

running ONTAP 9.7 storage OS.

[FlexPod Datacenter for SAP solution using FibreChannel SAN with Cisco UCS Manager 4.0 and NetApp ONTAP 9.7](#)

Deploy SAP application servers on FlexPod with SQL

FlexPod is a pre-validated and engineered data center solution designed to address the rapidly changing needs of IT departments. Cisco and NetApp have partnered to deliver FlexPod, which uses best-in-class computing, networking, and storage components as the foundation for a variety of enterprise workloads, including databases, enterprise resource planning (ERP), customer relationship management (CRM), and web applications. The consolidation of IT applications, particularly databases, has generated considerable interest in recent years. The most widely adopted and deployed database platform over the past several years is Microsoft SQL Server. SQL Server databases frequently have become subject to database sprawl, leading to IT challenges such as underutilized servers, incorrect licensing, security concerns, management concerns, and huge operational costs. Therefore, SQL Server databases are good candidates for consolidation on a more robust, flexible, and resilient platform. This document discusses a FlexPod reference architecture for deploying and consolidating SQL Server databases.

[Deploy SAP application servers on FlexPod with SQL](#)

FlexPod Datacenter for SAP with Cisco ACI, Cisco UCS Manager 4.0, and NetApp AFF A-Series

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

This document describes the architecture and deployment procedures for the SAP HANA Tailored DataCenter Integration option on FlexPod infrastructure, which is composed of:

- Cisco UCS Computing System (Cisco UCS) supported by second-generation Intel Xeon Scalable Processors.
- Switching products that leverage Cisco Application Centric Infrastructure (ACI).
- NetApp A-series AFF arrays.

[FlexPod Datacenter for SAP with Cisco ACI, Cisco UCS Manager 4.0, and NetApp AFF A-Series](#)

FlexPod Datacenter for SAP solution with Cisco ACI, Cisco UCS Manager 4.0, and NetApp AFF A-Series - Design

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

This document describes Cisco ACI integrated FlexPod solution as a validated approach for deploying SAP HANA Tailored Data Center Integration (TDI) environments. This validated design provides guidelines and a framework for implementing SAP HANA with

best practices from Cisco and NetApp.

The recommended solution architecture is built on the Cisco Unified Computing System (Cisco UCS) using a unified software release to support Cisco UCS hardware platforms that include the following components:

- Cisco UCS B-Series blade servers and Cisco UCS C-Series rack servers configurable with Intel Optane Data Center Persistent Memory Module (DCPMM) option
- Cisco UCS 6400 series Fabric Interconnects
- Cisco Nexus 9000 Series Leaf and Spine switches
- NetApp All Flash series storage arrays

Additionally, this document provides validations for both Red Hat Enterprise Linux and SUSE Linux Enterprise Server for SAP HANA.

[FlexPod Datacenter for SAP solution with Cisco ACI, Cisco UCS Manager 4.0, and NetApp AFF A-Series - Design](#)

FlexPod Datacenter for SAP solution with Cisco UCS third-generation fabric and NetApp AFF A-Series

Shailendra Mruthunjaya, Cisco
Ralf Klahr, Cisco
Marco Schoen, NetApp

This document describes the deployment methodology of Cisco and NetApp FlexPod Datacenter for SAP HANA based on based on the Cisco UCS Computing System (Cisco UCS) supported by second-generation Intel Xeon Scalable Processors.

Cisco UCS Manager (UCSM) 4.0(4) provides consolidated support of all current Cisco UCS Fabric Interconnect models (6200, 6300, 6324 and 6454), 2200/2300 series IOM, Cisco UCS B-Series Blade, and Cisco UCS C-Series Rack Formfactor servers. FlexPod Datacenter with Cisco UCS unified software release 4.0(4d) and NetApp ONTAP 9.6 is a predesigned, best-practice data center architecture built on the Cisco UCS, the Cisco Nexus 9000 family of switches, and NetApp AFF A-Series storage arrays.

[FlexPod Datacenter for SAP solution with Cisco UCS third-generation fabric and NetApp AFF A-Series](#)

FlexPod Datacenter for SAP solution with Cisco UCS Manager 4.0 and NetApp AFF A-Series - Design

Pramod Ramamurthy, Cisco
Marco Schoen, NetApp

This document describes the Cisco and NetApp FlexPod solution, which is a validated approach for deploying SAP HANA Tailored Data Center Integration (TDI) environments. This validated design provides guidelines and a framework for implementing SAP HANA with best practices from Cisco and NetApp.

FlexPod is a leading integrated infrastructure that supports a broad range of enterprise workloads and use cases. This solution allows you to quickly and reliably deploy SAP HANA with a model of a tailored datacenter integration mode.

The recommended solution architecture is built on the Cisco Unified Computing System (Cisco UCS) using a unified software release to support Cisco UCS hardware platforms that include the following components:

- Cisco UCS B-Series blade servers, and Cisco UCS C-Series rack servers configurable with Intel Optane Data Center Persistent Memory Module (DCPMM) option
- Cisco UCS 6300 series Fabric Interconnects
- Cisco Nexus 9000 Series switches
- NetApp All Flash series storage arrays

In addition, this document provides validations for both Red Hat Enterprise Linux and SUSE Linux Enterprise Server for SAP HANA.

[FlexPod Datacenter for SAP solution with Cisco UCS Manager 4.0 and NetApp AFF A-Series - Design](#)

Oracle

FlexPod Datacenter with Oracle 19c RAC Databases on Cisco UCS and NetApp AFF with NVMe over FibreChannel

Tushar Patel, Cisco
Hardikkumar Vyas, Cisco

Cisco Validated Designs (CVDs) consist of systems and solutions that are designed, tested, and documented to facilitate and improve customer deployments. This CVD document describes the Cisco and NetApp FlexPod solution, which is a validated approach for deploying a highly available Oracle RAC Database environment. Cisco and NetApp have validated the reference architecture with various database workloads, like OLTP (Online Transactional Processing) and Data Warehouse in Cisco's UCS Datacenter lab. This document shows the hardware and software configuration of the components involved and the results of various tests. Additionally, the document offers a framework for implementing Oracle RAC Databases on NVMe/FC using Cisco UCS and NetApp Storage System.

[FlexPod Datacenter with Oracle 19c RAC Databases on Cisco UCS and NetApp AFF with NVMe over FibreChannel](#)

FlexPod Datacenter with Oracle RAC Databases on Cisco UCS and NetApp AFF A-Series

Tushar Patel, Cisco
Hardikkumar Vyas, Cisco

Cisco Validated Designs include 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 customers. Cisco and NetApp have partnered to deliver FlexPod, which serves as the foundation for a variety of workloads and enables efficient architectural designs that are based on customer requirements. A

FlexPod solution is a validated approach for deploying Cisco and NetApp technologies as a shared cloud infrastructure.

The FlexPod Datacenter with NetApp All Flash AFF system is a converged infrastructure platform that combines best-of-breed technologies from Cisco and NetApp into a powerful converged platform for enterprise applications. Cisco and NetApp work closely with Oracle to support the most demanding transactional and response-time-sensitive databases required by today's businesses.

This Cisco Validated Design (CVD) describes the reference FlexPod Datacenter architecture using Cisco UCS and NetApp All Flash AFF Storage for deploying a highly available Oracle RAC Database environment. This document shows the hardware and software configuration of the components involved and results of various tests. Also, this document offers implementation and best practices guidance using Cisco UCS Compute Servers, Cisco Fabric Interconnect Switches, Cisco MDS Switches, Cisco Nexus Switches, NetApp AFF Storage and Oracle RAC Database.

[FlexPod Datacenter with Oracle RAC Databases on Cisco UCS and NetApp AFF A-Series](#)

FlexPod Datacenter with Oracle RAC on Oracle Linux

Tushar Patel, Cisco
Niranjan Mohapatra, Cisco
John Elliott, NetApp

The Cisco Unified Computing System (Cisco UCS) is a next-generation data center platform that unites computing, network, storage access, and virtualization into a single cohesive system. Cisco UCS is an ideal platform for the architecture of mission-critical database workloads. The combination of Cisco UCS platform, NetApp storage, and Oracle Real Application Cluster (RAC) architecture can accelerate your IT transformation by enabling faster deployments, greater flexibility of choice, efficiency, and lower risk. This Cisco Validated Design (CVD) highlights a flexible, multitenant, high performance and resilient FlexPod reference architecture featuring the Oracle 12c RAC Database.

The FlexPod platform, developed by NetApp and Cisco, is a flexible, integrated infrastructure solution that delivers pre-validated storage, networking, and server technologies. It's designed to increase IT responsiveness to business demands while reducing the overall cost of computing. Think maximum uptime, minimal risk. FlexPod components are integrated and standardized to help you achieve timely, repeatable, consistent deployments. You can plan with accuracy the power, floor space, usable capacity, performance, and cost of each FlexPod deployment.

FlexPod embraces the latest technology and efficiently simplifies the data center workloads that redefine the way IT delivers value:

- Take advantage of the capability of NetApp FAS Hybrid Arrays with Flash Pool flash to provide the capability to deploy the precise proportion of flash to spinning media for your specific application or environment.
- Leverage a pre-validated platform to minimize business disruption and improve IT agility and reduce deployment time from months to weeks.
- Slash administration time and total cost of ownership (TCO) by 50 percent.
- Meet or exceed constantly expanding hardware performance demands for data center workloads.

[FlexPod Datacenter with Oracle RAC on Oracle Linux](#)

FlexPod Datacenter with Oracle RAC Databases on Cisco UCS and NetApp AFF A-Series

Tushar Patel, Cisco
Hardikkumar Vyas, Cisco

The FlexPod Datacenter with NetApp All Flash AFF system is a converged infrastructure platform that combines best-of-breed technologies from Cisco and NetApp into a powerful converged platform for enterprise applications. Cisco and NetApp work closely with Oracle to support the most demanding transactional and response-time-sensitive databases required by today's businesses.

This Cisco Validated Design (CVD) describes the reference FlexPod Datacenter architecture using Cisco UCS and NetApp All Flash AFF Storage for deploying a highly available Oracle RAC Database environment. This document shows the hardware and software configuration of the components involved and the results of various tests. Also, this document offers implementation and best practices guidance using Cisco UCS Compute Servers, Cisco Fabric Interconnect Switches, Cisco MDS Switches, Cisco Nexus Switches, NetApp AFF Storage and Oracle RAC Database.

[FlexPod Datacenter with Oracle RAC Databases on Cisco UCS and NetApp AFF A-Series](#)

Microsoft SQL Server

FlexPod Datacenter for Microsoft SQL Server 2019 and VMware vSphere 6.7

Gopu Narasimha Reddy, Cisco
Sanjeev Naldurgkar, Cisco
Atul Bhalodia, NetApp

This document describes a FlexPod reference architecture using the latest hardware and software products and provides deployment recommendations for hosting Microsoft SQL Server 2019 databases in VMware ESXi virtualized environments. This solution also uses Cisco Workload Optimization Manager (CWOM), which provides automated recommendations for optimal and efficient resource utilization for both SQL workloads and infrastructure.

The solution is built on Cisco Unified Computing System (Cisco UCS) using the unified software release 4.1.1c to support the Cisco UCS hardware platforms, including Cisco UCS B-Series Blade Servers, Cisco UCS 6400 Fabric Interconnects, Cisco Nexus 9000 Series Switches, and NetApp AFF Series Storage Arrays.

[FlexPod Datacenter for Microsoft SQL Server 2019 and VMware vSphere 6.7](#)

FlexPod Datacenter with Microsoft SQL Server 2016 and VMware vSphere 6.5

Gopu Narasimha Reddy, Cisco
Sanjeev Naldurgkar, Cisco
David Arnette, NetApp

This document discusses a FlexPod reference architecture using the latest hardware and software products and provides configuration recommendations for deploying Microsoft

SQL Server databases in a virtualized environment.

The recommended solution architecture is built on Cisco Unified Computing System (Cisco UCS) using the unified software release to support the Cisco UCS hardware platforms, including Cisco UCS B-Series Blade Servers, Cisco UCS 6300 Fabric Interconnects, Cisco Nexus 9000 Series Switches, and NetApp All Flash Series Storage Arrays. Additionally, this solution includes VMware vSphere 6.5, vSphere 6.5, providing a number of new features to optimize storage utilization and to facilitate a private cloud.

[FlexPod Datacenter with Microsoft SQL Server 2016 and VMware vSphere 6.5](#)

FlexPod Datacenter with Microsoft SQL Server 2017 on Linux VM running on VMware and Hyper-V

Gopu Narasimha Reddy, Cisco
Sanjeev Naldurgkar, Cisco
Atul Bhalodia, NetApp

This document discusses a FlexPod reference architecture using the latest hardware and software products and provides deployment recommendations for hosting Microsoft SQL Server databases in VMware ESXi and Microsoft Windows Hyper-V virtualized environments with Linux support enablement from Microsoft for SQL Server deployment.

The recommended solution architecture is built on Cisco Unified Computing System (Cisco UCS) using the unified software release 4.0.1c to support the Cisco UCS hardware platforms including Cisco UCS B-Series Blade Servers, Cisco UCS 6300 Fabric Interconnects, Cisco Nexus 9000 Series Switches, and NetApp AFF Series Storage Arrays.

[FlexPod Datacenter with Microsoft SQL Server 2017 on Linux VM running on VMware and Hyper-V](#)

FlexPod Datacenter with Microsoft SQL Server 2017 on Linux VM running on VMware and Hyper-V

Gopu Narasimha Reddy, Cisco
Sanjeev Naldurgkar, Cisco
Atul Bhalodia, NetApp

This document discusses a FlexPod reference architecture using the latest hardware and software products and provides deployment recommendations for hosting Microsoft SQL Server databases in VMware ESXi and Microsoft Windows Hyper-V virtualized environments with Linux support enablement from Microsoft for SQL Server deployment.

The recommended solution architecture is built on Cisco Unified Computing System (Cisco UCS) using the unified software release 4.0.1c to support the Cisco UCS hardware platforms, including Cisco UCS B-Series Blade Servers, Cisco UCS 6300 Fabric Interconnects, Cisco Nexus 9000 Series Switches, and NetApp AFF Series Storage Arrays.

[FlexPod Datacenter with Microsoft SQL Server 2017 on Linux VM running on VMware and Hyper-V](#)

Healthcare

FlexPod for Genomics

TR-4911: FlexPod Genomics

JayaKishore Esanakula, NetApp

There are few fields of medicine that are more important than genomics for healthcare and the life sciences, and genomics are fast becoming a key clinical tool for doctors and nurses. Genomics, when combined with medical imaging and digital pathology, help us understand how a patient's genes might be affected by treatment protocols. The success of genomics in healthcare increasingly depends on data interoperability at scale. The end goal is to make sense of the enormous volumes of genetic data and identify clinically relevant correlations and variants that improve diagnosis and make precision medicine a reality. Genomics help us understand the origin of disease outbreaks, how diseases evolve, and which treatments and strategies might be effective. Clearly, genomics has many benefits that span prevention, diagnosis, and treatment. Healthcare organizations are grappling with several challenges, including the following:

- Improved care quality
- Value-based care
- Data explosion
- Precision medicine
- Pandemics
- Wearables, remote monitoring, and care
- Cyber security

Standardized clinical pathways and clinical protocols are one of the critical components of modern medicine. One of the key aspects of standardization is interoperability between care providers, not just for medical records but also for genomic data. The big question is will healthcare organizations relinquish ownership of genomic data in lieu of patient ownership of their personal genomics data and related medical records?

Interoperable patient data is key for enabling precision medicine, one of the driving forces behind the recent explosion of data growth. The objective for precision medicine is to make health maintenance, disease prevention, diagnoses, and treatment solutions more effective and precise.

The rate of data growth has been exponential. In early February 2021, US laboratories sequenced approximately 8,000 COVID-19 strains per week. The number of genomes sequenced had increased to 29,000 per week by April 2021. Each fully sequenced human genome is around 125GB in size. Therefore, at a rate of 29,000 genomes sequenced per week, total genome storage at rest would be more than 180 petabytes per year. Various countries have committed resources to genomic epidemiology to improve genomic surveillance and prepare for the next wave of global health challenges.

The reduced cost of genomic research is driving genetic testing and research at an unprecedented rate. The three Ps are at an inflection point: computer power, privacy of data, and personalization of medicine. By 2025 researchers estimate that 100 million to as many as 2 billion human genomes will be sequenced. For genomics to be effective and a valuable proposition, genomics capabilities must be a seamless part of care workflows; it

should be easy to access and be actionable during a patient's visit. It is also equally important that patient electronic medical-record data be integrated with patient genomics data. With the advent of state-of-the-art converged infrastructure like FlexPod, organizations can bring their genomics capabilities into the everyday workflows of physicians, nurses, and clinic managers. For the latest FlexPod platform information, see this [FlexPod Datacenter with Cisco UCS X-Series White Paper](#).

For a physician, the true value of genomics includes precision medicine and personalized treatment plans based on the genomic data of a patient. There has never been such synergy between clinicians and data scientists in the past, and genomics is benefiting from the technological innovations in the recent past, and also real partnerships between healthcare organizations and technology leaders in the industry.

Academic medical centers and other healthcare and life science organizations are well on their way to establishing center of excellence (COE) in genome science. According to Dr. Charlie Gersbach, Dr. Greg Crawford, and Dr. Tim E Reddy from Duke University, "We know that genes aren't turned on or off by a simple binary switch, but instead it's a result of multiple gene regulatory switches that work together." They have also determined that "none of these parts of the genome work in isolation. The genome is a very complicated web that evolution has woven" ([ref](#)).

NetApp and Cisco have been hard at work implementing incremental improvements into the FlexPod platform for over 10 years. All customer feedback is heard, evaluated, and tied into the value streams and feature sets in FlexPod. It is this continuous loop of feedback, collaboration, improvement, and celebration that sets FlexPod apart as a trusted converged infrastructure platform the world over. It has been simplified and designed from the ground up to be the most reliable, robust, versatile, and agile platform for healthcare organizations.

Scope

The FlexPod converged infrastructure platform enables a healthcare organization to host one or more genomics workloads, along with other clinical and nonclinical healthcare applications. This technical report uses an open-source, industry-standard genomics tool called GATK during FlexPod platform validation. However, a deeper discussion of genomics or GATK is outside the scope of this document.

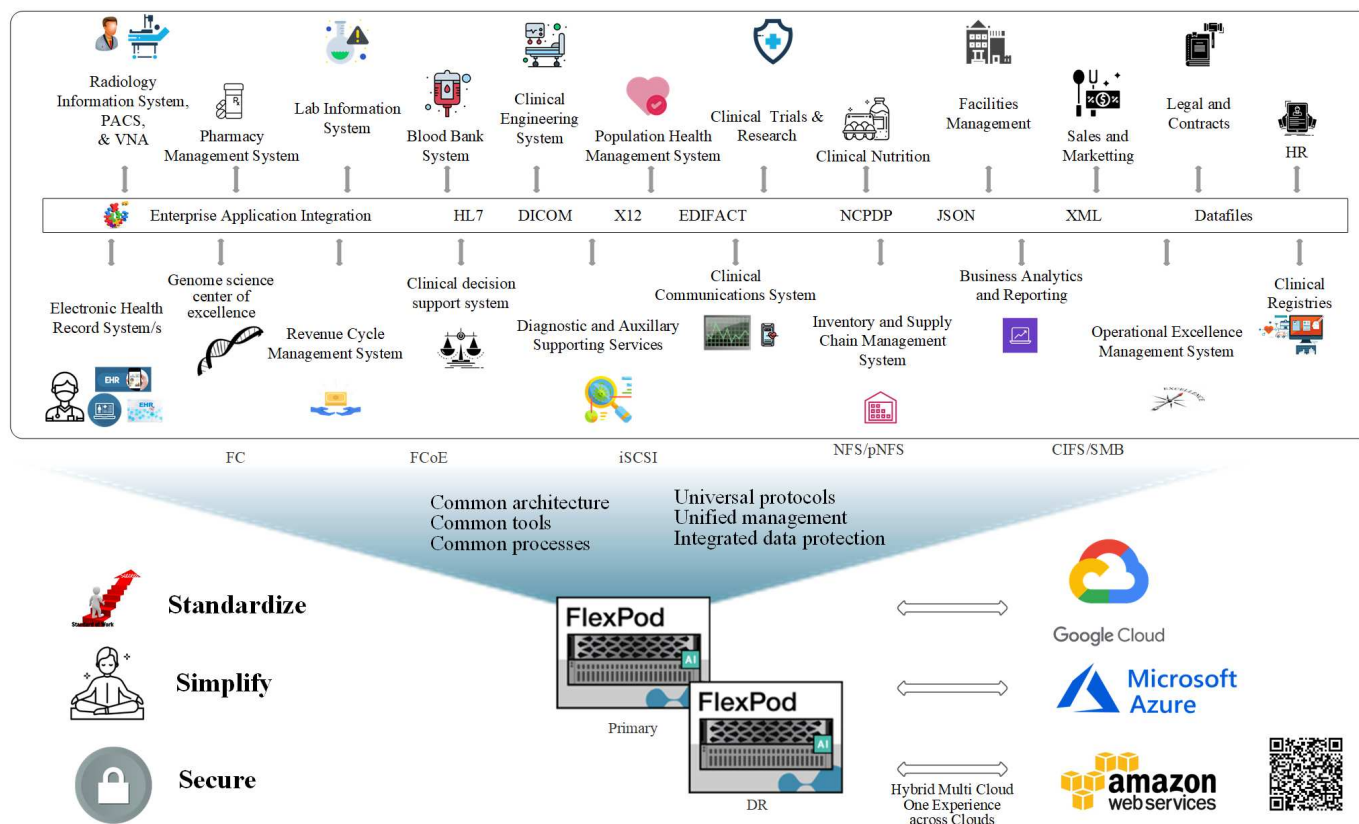
Audience

This document is intended for technical leaders in the healthcare industry and for Cisco and NetApp partner solutions engineers and professional services personnel. NetApp assumes that the reader has a good understanding of compute and storage sizing concepts as well as a technical familiarity with healthcare threats, healthcare security, healthcare IT systems, Cisco UCS, and NetApp storage systems.

Hospital capabilities deployed on FlexPod

A typical hospital has a diversified set of IT systems. The majority of such systems are purchased from a vendor, whereas very few are built by the hospital system in house. Therefore, the hospital system must manage a diverse infrastructure environment in their data centers. When hospitals unify their systems into a converged infrastructure platform like FlexPod, organizations can standardize their data center operations. With FlexPod, healthcare organizations can implement clinical and non-clinical systems on the same platform, thereby unifying data center operations.

Hospital capabilities deployed on a FlexPod



Next: [Benefits of deploying genomic workloads on FlexPod.](#)

Benefits of deploying genomic workloads on FlexPod

Previous: [Introduction.](#)

This section provides a brief list of benefits for running a genomics workload on a FlexPod converged infrastructure platform. Let's quickly describe the capabilities of a hospital. The following business architecture view shows a hospital's capabilities deployed on a hybrid-cloud-ready FlexPod converged infrastructure platform.

- Avoid siloes in healthcare.** Silos in healthcare are a very real concern. Departments are often siloed into their own set of hardware and software not by choice but organically by evolution. For example, radiology, cardiology, EHR, genomics, analytics, revenue cycle, and other departments end up with their individual set of dedicated software and hardware. Healthcare organizations maintain a limited set of IT professionals to manage their hardware and software assets. The inflection point comes when this set of individuals are expected to manage a very diversified set of hardware and software. Heterogeneity is made worse by an incongruent set of processes brought to the healthcare organization by vendors.
- Start small and grow.** The GATK tool kit is tuned for CPU execution, which best suites platforms like FlexPod. FlexPod enables independent scalability of network, compute, and storage. Start small and scale as your genomics capabilities and the environment grows. Healthcare organizations don't have to invest in specialized platforms to run genomic workloads. Instead, organizations can leverage versatile platforms like a FlexPod to run genomics and non-genomics workloads on the same platform. For example, if the pediatrics department wants to implement genomics capability, IT leadership can provision compute, storage, and networking on an existing FlexPod instance. As the genomics business unit grows, healthcare

organization can scale their FlexPod platform as needed.

- **Single control pane and unparalleled flexibility.** Cisco Intersight significantly simplifies IT operations by bridging applications with infrastructure, providing visibility and management from bare-metal servers and hypervisors to serverless applications, thereby reducing costs and mitigating risk. This unified SaaS platform uses a unified Open API design that natively integrates with third-party platforms and tools. Moreover, it allows management to occur from your data center operations team on site or from anywhere by using a mobile app.

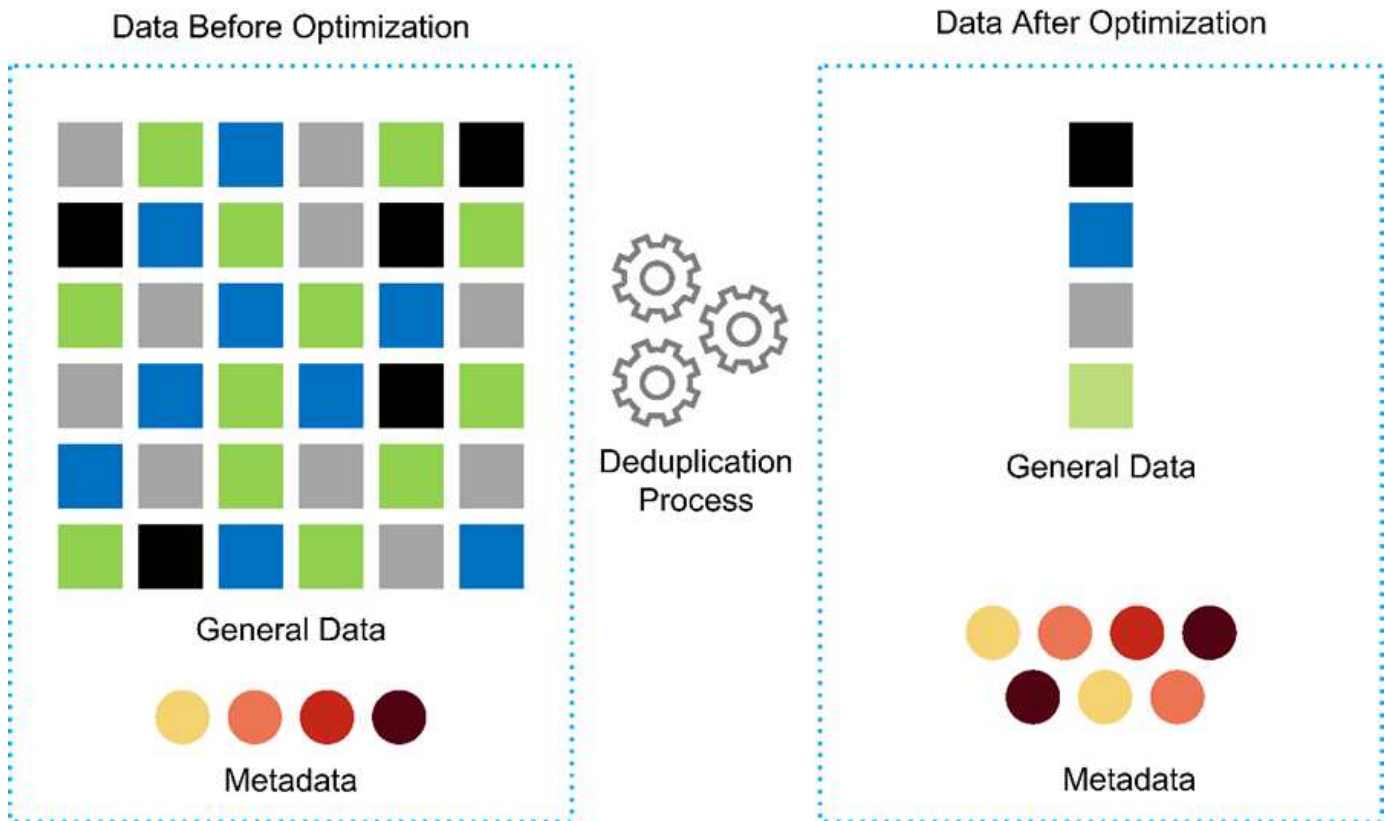
Users quickly unlock tangible value in their environment by leveraging Intersight as their management platform. Enabling automation for many daily manual tasks, Intersight removes errors and simplifies your daily operations. Moreover, advanced support capabilities facilitated by Intersight allow adopters to stay ahead of problems and accelerate issue resolution. Taken in combination, organizations spend far less time and money on their application infrastructure and more time on their core business development.

Leveraging Intersight management and FlexPod's easily scalable architecture enables organizations to run several genome workloads on a single FlexPod platform, increasing utilization and reducing total cost of ownership (TCO). FlexPod allows for flexible sizing, with choices starting with our small FlexPod Express and scaling into large FlexPod Datacenter implementations. With role-based access control capabilities built into Cisco Intersight, healthcare organizations can implement robust access control mechanisms, avoiding the need for separate infrastructure stacks. Multiple business units within the healthcare organization can leverage genomics as a key core competency.


Ultimately FlexPod helps simplify IT operations and lower operating costs, and it allows IT infrastructure admins to focus on tasks that help clinicians innovate rather than being relegated to keeping the lights on.

- **Validated design and guaranteed outcomes.** FlexPod design and deployment guides are validated to be repeatable, and they cover comprehensive configuration details and industry best practices that are needed to deploy a FlexPod with confidence. Cisco and NetApp validated design guides, deployment guides, and architectures help your healthcare or life science organization remove guesswork from the implementation of a validated and trusted platform from the beginning. With FlexPod, you can speed up deployment times and reduce cost, complexity, and risk. FlexPod validated designs and deployment guides establish FlexPod as the ideal platform for a variety of genomics workloads.
- **Innovation and agility.** FlexPod is recommended as an ideal platform by EHRs like Epic, Cerner, Meditech and imaging systems like Agfa, GE, Philips. For more information on [Epic honor roll](#) and target platform architecture, see the Epic userweb. Running genomics on [FlexPod](#) enables healthcare organizations to continue their journey of innovation with agility. With FlexPod, implementing organizational change comes naturally. When organizations standardize on a FlexPod platform, healthcare IT experts can provision their time, effort, and resources to innovate and thus be as agile as the ecosystem demands.
- **Data liberated.** With the FlexPod converged infrastructure platform and a NetApp ONTAP storage system, genomics data can be made available and accessible using a wide variety of protocols at scale from a single platform. FlexPod with NetApp ONTAP offers a simple, intuitive, and powerful hybrid cloud platform. Your data fabric powered by NetApp ONTAP weaves data together across sites, beyond physical boundaries, and across applications. Your data fabric is built for data-driven enterprises in a data-centric world. Data is created and used in multiple locations, and it often needs to be leveraged and shared with other locations, applications, and infrastructures. Therefore, you need a consistent and integrated way to manage it. FlexPod puts your IT team in control and simplifies ever-increasing IT complexity.
- **Secure multitenancy.** FlexPod uses FIPS 140-2 compliant cryptographic modules, hence enabling organizations to implement security as a foundational element, not an afterthought. FlexPod enables organizations implement secure multitenancy from a single converged infrastructure platform irrespective of the size of the platform. FlexPod with secured multitenancy and QoS help with workload separation and maximize utilization. This helps avoid capital being locked into specialized platforms that is potentially underutilized and requires a specialized skill set to manage.

- **Storage efficiency.** Genomics requires that the underlying storage have industry- leading storage efficiency capabilities. You can reduce storage costs with NetApp storage efficiency features such as deduplication (inline and on demand), data compression, and data compaction ([ref](#)). NetApp deduplication provides block-level deduplication in a FlexVol volume. Essentially, deduplication removes duplicate blocks, storing only unique blocks in the FlexVol volume. Deduplication works with a high degree of granularity and operates on the active file system of the FlexVol volume. The following figure shows an overview of how NetApp deduplication works. Deduplication is application transparent. Therefore, it can be used to deduplicate data originating from any application that uses the NetApp system. You can run volume deduplication as an inline process and as a background process. You can configure it to run automatically, to be scheduled, or to run manually through the CLI, NetApp ONTAP System Manager, or NetApp Active IQ Unified Manager.



- **Enable genomics interoperability.** ONTAP FlexCache is a remote caching capability that simplifies file distribution, reduces WAN latency, and lowers WAN bandwidth costs, ([ref](#)). One of the key activities during genomic variant identification and annotation is collaboration between clinicians. ONTAP FlexCache technology increases data throughput even when collaborating clinicians are in different geographic locations. Given the typical size of a *.BAM file (1GB to 100s of GB), it is critical that the underlying platform can make files available to clinicians in different geographic locations. FlexPod with ONTAP FlexCache makes genomic data and applications truly multisite ready, which makes collaboration between researchers located around the world seamless with low latency and high throughput. Healthcare organizations running genomics applications in a multisite setting can scale-out using the data fabric to balance manageability with cost and speed.
- **Intelligent use of storage platform.** FlexPod with ONTAP auto-tiering and NetApp Fabric Pool technology simplifies data management. FabricPool helps 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 the need to rearchitect the application infrastructure. FlexPod can benefit from the storage tiering capabilities of FabricPool to make more efficient use of ONTAP flash storage. For more information, see [FlexPod with FabricPool](#). The following diagram provides a high-level overview of FabricPool and its benefits.



Automatic tiering
 Zero-touch management
 Preserves file system
 Lower cost of ownership
 Choice of object tier locations



- **Faster variant analysis and annotation.** The FlexPod platform is faster to deploy and operationalize. The FlexPod platform enables clinician collaboration by making data available at scale with low latency and increased throughput. Increased interoperability enables innovation. Healthcare organizations can run their genomic and non-genomic workloads side by side, which means organizations do not need specialized platforms to start their genomics journey.

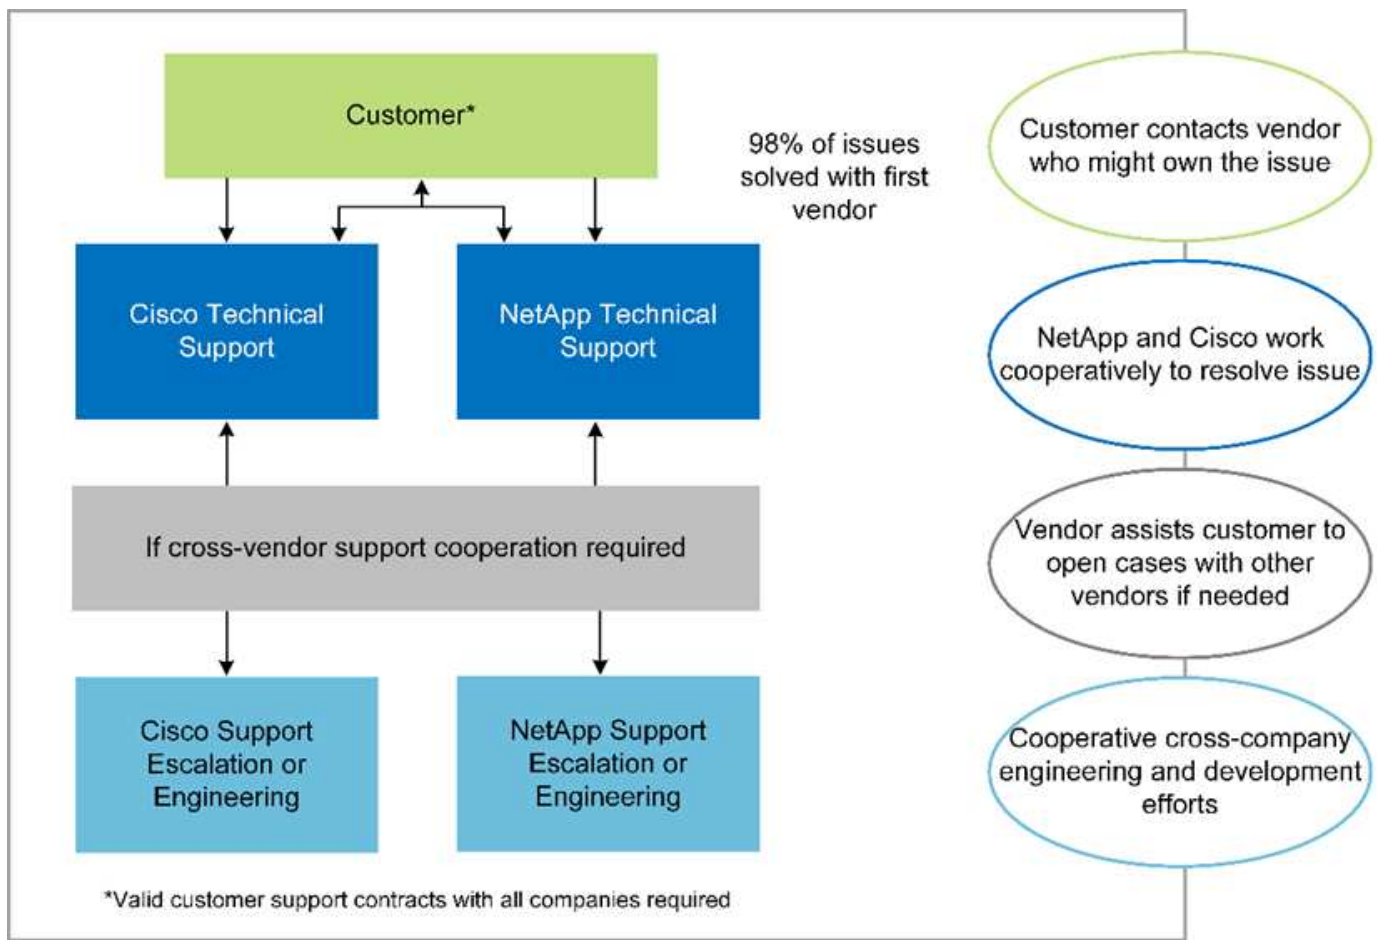
FlexPod ONTAP routinely adds cutting edge features to the storage platform. FlexPod Datacenter is the optimal shared infrastructure foundation for deploying FC- NVMe to allow high-performance storage access to applications that need it. As FC- NVMe evolves to include high availability, multipathing, and additional operating system support, FlexPod is well suited as the platform of choice, providing the scalability and reliability needed to support these capabilities. ONTAP with faster I/O with end-to-end NVMe allows genomics analyses to completed faster ([ref](#)).

Sequenced raw genome data produces large file sizes, and it is important that these files are made available to the variant analyzers to reduce the total time it takes from sample collection to variant annotation. NVMe (nonvolatile memory express) when used as a storage access and data transport protocol provides unprecedented levels of throughput and the fastest response times. FlexPod deploys the NVMe protocol while accessing flash storage via the PCI express bus (PCIe). PCIe enables implementation of tens of thousands of command queues, increasing parallelization and throughput. One single protocol from storage to memory makes data access fast.

- **Agility for clinical research from the ground up.** Flexible, expandable storage capacity and performance allows the healthcare research organizations to optimize the environment in an elastic or just-in-time (JIT) manner. By decoupling storage from compute and network infrastructure, FlexPod platform can be scaled up and out without disruption. Using Cisco Intersight, the FlexPod platform can be managed with both built-in and custom automated workflows. Cisco Intersight workflows enable healthcare organizations to reduce application life-cycle management times. When an academic medical center requires that patient data be anonymized and made available to their center for research informatics and/or center for quality, their IT organization can leverage Cisco Intersight FlexPod workflows to take secure data backups, clone, and the

restore in a matter of seconds, not hours. With NetApp Trident and Kubernetes, IT organizations can provision new data scientists and make clinical data available for model development in a matter of minutes, sometimes even in seconds.

- **Protecting genome data.** NetApp SnapLock provides a special-purpose volume in which files can be stored and committed to a non-erasable, non-rewritable state. The user's production data residing in a FlexVol volume can be mirrored or vaulted to a SnapLock volume through NetApp SnapMirror or SnapVault technology. The files in the SnapLock volume, the volume itself, and its hosting aggregate cannot be deleted until the end of the retention period. Using ONTAP FPolicy software organizations can prevent ransomware attacks by disallowing operations on files with specific extensions. An FPolicy event can be triggered for specific file operations. The event is tied to a policy, which calls out the engine it needs to use. You might configure a policy with a set of file extensions that could potentially contain ransomware. When a file with a disallowed extension tries to perform an unauthorized operation, FPolicy prevents that operation from executing ([ref](#)).
- **FlexPod Cooperative Support.** NetApp and Cisco have established FlexPod Cooperative Support, a strong, scalable, and flexible support model to meet the unique support requirements of the FlexPod converged infrastructure. This model uses the combined experience, resources, and technical support expertise of NetApp and Cisco to offer a streamlined process for identifying and resolving FlexPod support issues, regardless of where the problem resides. The following figure provides an overview of the FlexPod Cooperative Support model. The customer contacts the vendor who might own the issue, and both Cisco and NetApp work cooperatively to resolve it. Cisco and NetApp have cross-company engineering and development teams that work hand in hand to resolve issues. This support model reduces loss of information during translation, enables trust, and reduces downtime.

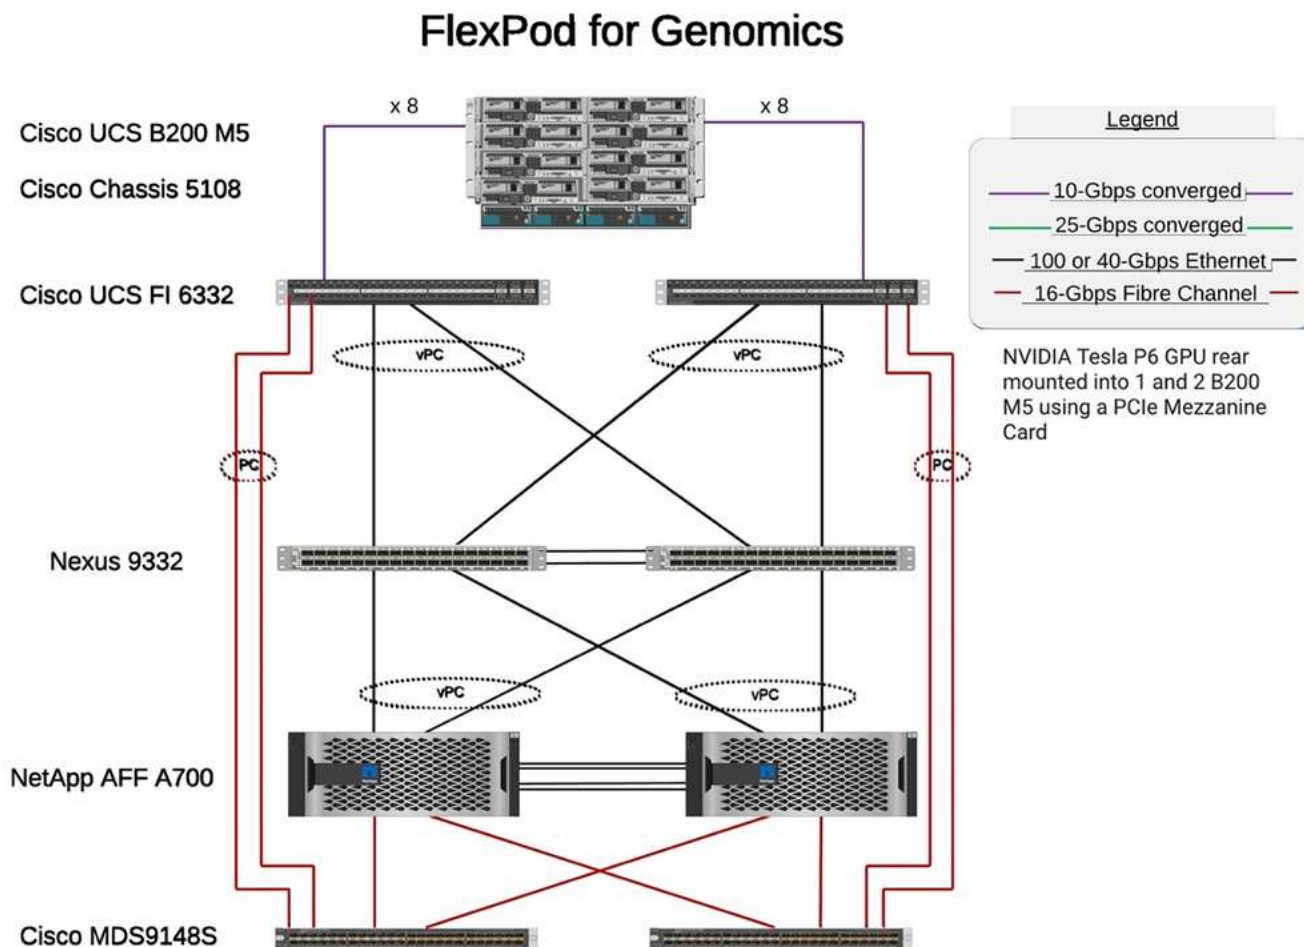


Next: Solution infrastructure hardware and software components.

Solution infrastructure hardware and software components


Previous: [Benefits of deploying genomic workloads on FlexPod.](#)

The following figure depicts the FlexPod system used for GATK setup and validation. We used [FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7 Cisco Validated Design \(CVD\)](#) during the setup process.



The following diagram depicts the FlexPod cabling details.

FlexPod for Genomics



The following table lists the hardware components used during the GATK testing enabling on a FlexPod. Here is the [NetApp Interoperability Matrix Tool \(IMT\)](#) and [Cisco Hardware Compatibility List \(HCL\)](#).

| Layer | Product family | Quantity and model | Details |
|---------|-------------------------|--------------------|--|
| Compute | Cisco UCS 5108 chassis | 1 or 2 | |
| | Cisco UCS blade servers | 6 B200 M5 | Each with 2x 20 or more cores, 2.7GHz, and 128-384GB RAM |

| Layer | Product family | Quantity and model | Details |
|-----------------|--|--------------------------------|--------------------------------|
| | Cisco UCS Virtual Interface Card (VIC) | Cisco UCS 1440 | See the |
| | 2x Cisco UCS Fabric Interconnects | 6332 | - |
| Network | Cisco Nexus switches | 2x Cisco Nexus 9332 | - |
| Storage network | IP network for storage access over SMB/CIFS, NFS, or iSCSI protocols | Same network switches as above | - |
| | Storage access over FC | 2x Cisco MDS 9148S | - |
| Storage | NetApp AFF A700 all-flash storage system | 1 Cluster | Cluster with two nodes |
| | Disk shelf | One DS224C or NS224 disk shelf | Fully populated with 24 drives |
| | SSD | 24, 1.2TB or larger capacity | - |

This table lists the infrastructure software.

| Software | Product family | Version or release | Details |
|------------|---|---|---------|
| Various | Linux | RHEL 8.3 | - |
| | Windows | Windows Server 2012 R2 (64 bit) | - |
| | NetApp ONTAP | ONTAP 9.8 or later | - |
| | Cisco UCS Fabric Interconnect | Cisco UCS Manager 4.1 or later | - |
| | Cisco Ethernet 3000 or 9000 series switches | For 9000 series, 7.0(3)I7(7) or later For 3000 series, 9.2(4) or later | - |
| | Cisco FC: Cisco MDS 9132T | 8.4(1a) or later | - |
| | Hypervisor | VMware vSphere ESXi 7.0 | - |
| Storage | Hypervisor management system | VMware vCenter Server 7.0 (vCSA) or later | - |
| Network | NetApp Virtual Storage Console (VSC) | VSC 9.7 or later | - |
| | NetApp SnapCenter | SnapCenter 4.3 or later | - |
| | Cisco UCS Manager | 4.1(3c) or later | |
| Hypervisor | ESXi | | |

| Software | Product family | Version or release | Details |
|------------|---|-------------------------|---------|
| Management | Hypervisor management systemVMware vCenter Server 7.0 (vCSA) or later | | |
| | NetApp Virtual Storage Console (VSC) | VSC 9.7 or later | |
| | NetApp SnapCenter | SnapCenter 4.3 or later | |
| | Cisco UCS Manager | 4.1(3c) or later | |

Next: [Genomics - GATK setup and execution](#).

Genomics - GATK setup and execution

Previous: [Solution infrastructure hardware and software components](#).

According to the National Human Genome Research Institute ([NHGRI](#)), “Genomics is the study of all of a person’s genes (the genome), including interactions of these genes with each other and with a person’s environment. ”

According to the [NHGRI](#), “Deoxyribonucleic acid (DNA) is the chemical compound that contains the instructions needed to develop and direct the activities of nearly all living organisms. DNA molecules are made of two twisting, paired strands, often referred to as a double helix.” “An organism’s complete set of DNA is called its genome.”


Sequencing is the process of determining the exact order of the bases in a strand of DNA. One of the most common types of sequencing used today is called sequencing by synthesis. This technique uses the emission of fluorescent signals to order the bases. Researchers can use DNA sequencing to search for genetic variations and any mutations that might play a role in the development or progression of a disease while a person is still in the embryonic stage.

From sample to variant identification, annotation, and prediction

At a high level, genomics can be classified into the following steps. This is not an exhaustive list:

1. Sample collection.
2. [Genome sequencing](#) using a sequencer to generate the raw data.
3. Preprocessing. For example, [deduplication](#) using [Picard](#).
4. Genomic analysis.
 - a. Mapping to a reference genome.
 - b. [Variant](#) identification and annotation typically performed using GATK and similar tools.
5. Integration into the electronic health record (EHR) system.
6. [Population stratification](#) and identification of genetic variation across geographical location and ethnic background.
7. [Predictive models](#) using significant single- nucleotide polymorphism.
8. [Validation](#).

The following figure shows the process from sampling to variant identification, annotation, and prediction.



The Human Genome project was completed in April 2003, and the project made a very high-quality simulation of the human genome sequence available in the public domain. This reference genome initiated an explosion in research and development of genomics capabilities. Virtually every human ailment has a signature in that human's genes. Up until recently, physicians were leveraging genes to predict and determine birth defects like sickle cell anemia, which is caused by a certain inheritance pattern caused by a change in a single gene. The treasure trove of data being made available by the human genome project led to the advent of the current state of genomics capabilities.

Genomics has a broad set of benefits. Here is a small set of benefits in the healthcare and life sciences domains:

- Better diagnosis at point of care
- Better prognosis
- Precision medicine
- Personalized treatment plans
- Better disease monitoring
- Reduction in adverse events
- Improved access to therapies
- Improved disease monitoring
- Effective clinical trial participation and better selection of patients for clinical trials based on genotypes.

Genomics is a "four-headed beast," because of the computational demands across the lifecycle of a dataset: acquisition, storage, distribution, and analysis.

Genome Analysis Toolkit (GATK)

GATK was developed as a data science platform at the [Broad Institute](#). GATK is a set of open-source tools that enable genome analysis, specifically variant discovery, identification, annotation, and genotyping. One of the benefits of GATK is that the set of tools and or commands can be chained to form a complete workflow. The primary challenges that Broad institute tackles are the following:

- Understand the root causes and biological mechanisms of diseases.
- Identify therapeutic interventions that act at the fundamental cause of a disease.
- Understand the line of sight from variants to function in human physiology.
- Create standards and policy [frameworks](#) for genome data representation, storage, analysis, security, and so on.
- Standardize and socialize interoperable genome aggregation databases (gnomAD).
- Genome-based monitoring, diagnosis, and treatment of patients with greater precision.
- Help implement tools that predict diseases well before symptoms appear.
- Create and empower a community of cross-disciplinary collaborators to help tackle the toughest and most important problems in biomedicine.

According to GATK and the Broad institute, genome sequencing should be treated as a protocol in a pathology lab; every task is well documented, optimized, reproducible, and consistent across samples and experiments. The following is a set of steps recommended by the Broad Institute, for more information, see the [GATK website](#).

FlexPod setup

Genomics workload validation includes a from-scratch setup of a FlexPod infrastructure platform. The FlexPod platform is highly available and can be independently scaled; for example, network, storage and compute can be scaled independently. We used the following Cisco validated design guide as the reference architecture document to set up the FlexPod environment: [FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7](#). See the following FlexPod platform set up highlights:

To perform FlexPod lab setup, complete the following steps:

1. FlexPod lab setup and validation uses the following IP4 reservations and VLANs.

IP Reservations

| VLAN | IP Range | Subnet Mask | Purpose |
|------|---------------|---------------|---------|
| 3281 | 172.21.25 /24 | 255.255.255.0 | IB-MGMT |
| 3282 | 172.21.26 /24 | 255.255.255.0 | vMotion |
| 3283 | 172.21.27 /24 | 255.255.255.0 | VM |
| 3284 | 172.21.28 /24 | 255.255.255.0 | NFS |
| 3285 | 172.21.29 /24 | 255.255.255.0 | iSCSI-A |
| 3286 | 172.21.30 /24 | 255.255.255.0 | iSCSI-B |

2. Configure iSCSI-based boot LUNs on the ONTAP SVM.

ONTAP System Manager

LUNs

[+ Add](#)

| | Name | Storage VM |
|---|-----------------|----------------|
| ▼ | ESXi_Boot_Lun_1 | Healthcare_SVM |
| ▼ | ESXi_Boot_Lun_2 | Healthcare_SVM |
| ▼ | ESXi_Boot_Lun_3 | Healthcare_SVM |
| ▼ | ESXi_Boot_Lun_4 | Healthcare_SVM |
| ▼ | ESXi_Boot_Lun_5 | Healthcare_SVM |
| ▼ | ESXi_Boot_Lun_6 | Healthcare_SVM |

3. Map LUNs to iSCSI initiator groups.

| | Name | Storage VM | Volume | Size | IOPS | Latency (ms) | Throughput (MB/s) |
|---|-----------------|----------------|---------------|-------|------|--------------|-------------------|
| ▲ | ESXi_Boot_Lun_1 | Healthcare_SVM | ESXi_Boot_Vol | 20 GB | 3 | 0.16 | 0.01 |

| STATUS | VOLUME | DESCRIPTION | SNAPSHOT COPIES (LOCAL) | SNAPMIRROR (LOCAL OR REMOTE) |
|--------|---------------|-------------|-------------------------|------------------------------|
| Online | ESXi_Boot_Vol | - | Protected | Unprotected |


| SERIAL NUMBER | QOS POLICY GROUP | MAPPED TO INITIATORS | ID | SNAPSHOT POLICY |
|---------------|------------------|---|----|-----------------|
| 80A4X+R8rAhP | - | GenomicsESXi_1 (1) iqn.1992-08.com.cisco:ucs-... | 0 | default |

| CAPACITY (AVAILABLE % TOTAL) | LUN FORMAT | PATH |
|--------------------------------|------------|------------------------------------|
| 95% 20 GB | VMware | /vol/ESXi_Boot_Vol/ESXi_Boot_Lun_1 |

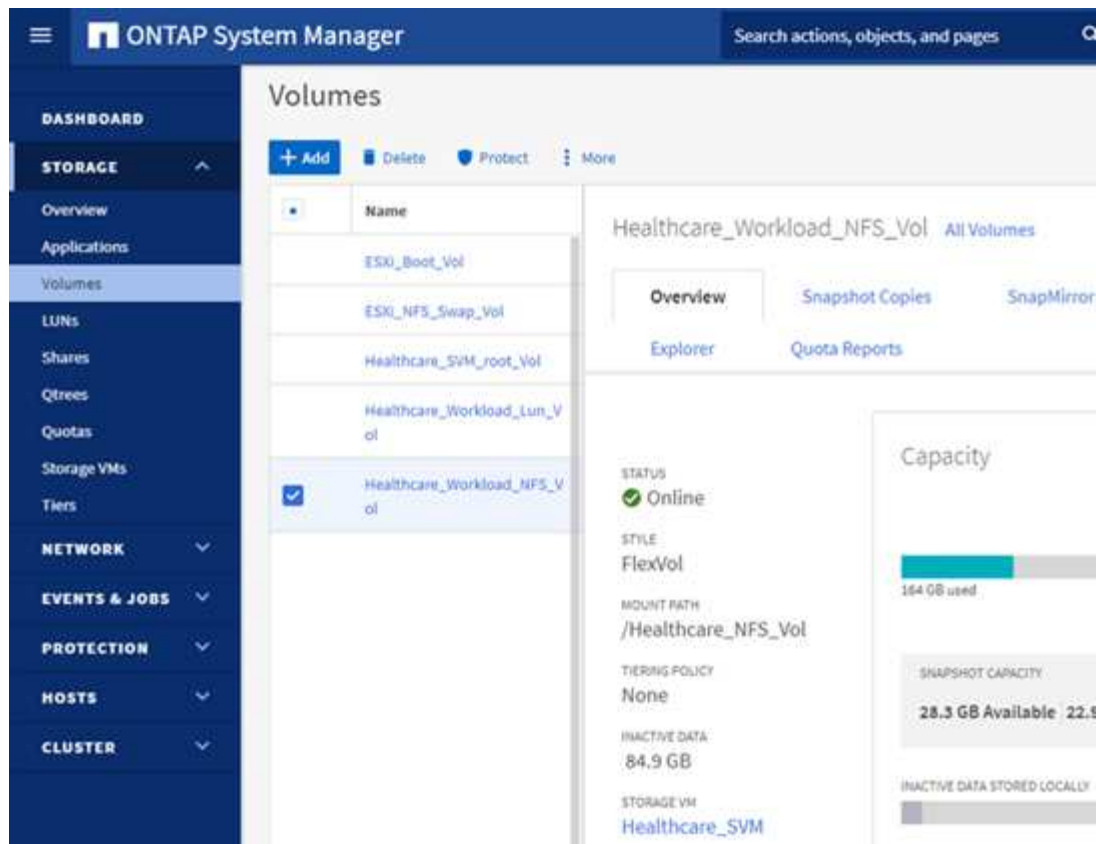
| | Name | Storage VM | Volume | Size | IOPS | Latency (ms) | Throughput (MB/s) |
|---|-----------------|----------------|---------------|-------|------|--------------|-------------------|
| ▼ | ESXi_Boot_Lun_1 | Healthcare_SVM | ESXi_Boot_Vol | 20 GB | 1 | 0.25 | 0.01 |
| ▲ | ESXi_Boot_Lun_2 | Healthcare_SVM | ESXi_Boot_Vol | 20 GB | 4 | 0.18 | 0.02 |

| | | | | |
|---|-------------------------|---|--|---|
| STATUS ✔ Online | VOLUME ESXi_Boot_Vol | DESCRIPTION - | SNAPSHOT COPIES (LOCAL) STATUS ✔ Protected | SNAPMIRROR (LOCAL OR REMOTE) STATUS ✔ Unprotected |
| SERIAL NUMBER 80A4X+R8rAhU | QOS POLICY GROUP - | MAPPED TO INITIATORS ^ GenomicsESXi_2 (1) iqn.1992-08.com.cisco:ucs-... | ID 0 | SNAPSHOT POLICY default |
| CAPACITY (AVAILABLE % TOTAL) 96% 20 GB | LUN FORMAT VMware | | | |


4. Install vSphere 7.0 with iSCSI boot.
5. Register ESXi hosts with the vCenter.



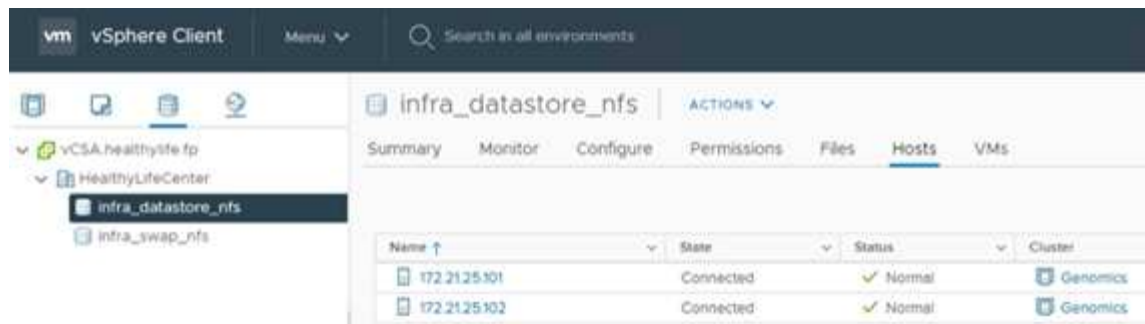
6. Provision an NFS datastore infra_datastore_nfs on the ONTAP storage.



7. Add the datastore to the vCenter.



8. Using vCenter, add an NFS datastore to the ESXi hosts.



9. Using the vCenter, create a Red Hat Enterprise Linux (RHEL) 8.3 VM to run GATK.
10. An NFS datastore is presented to the VM and mounted at `/mnt/genomics`, which is used to store GATK executables, scripts, Binary Alignment Map (BAM) files, reference files, index files, dictionary files, and out files for variant calling.

```
[root@genomics1 genomics]# df | grep genomics
/dev/sdb                308587328  5699492 287142812   2% /mnt/genomics
[root@genomics1 genomics]#
```

GATK setup and execution

Install the following prerequisites on the RedHat Enterprise 8.3 Linux VM:

- Java 8 or SDK 1.8 or later
- Download GATK 4.2.0.0 from the Broad Institute [GitHub site](#). Genome sequence data is generally stored in the form of a series of tab-delimited ASCII columns. However, ASCII takes too much space to store. Therefore, a new standard evolved called a BAM (*.bam) file. A BAM file stores the sequence data in a compressed, indexed, and binary form. We [downloaded](#) a set of publicly available BAM files for GATK execution from the [public domain](#). We also downloaded index files (*.bai), dictionary files (*.dict), and reference data files (*.fasta) from the same public domain.

After downloading, the GATK tool kit has a jar file and a set of support scripts.

- `gatk-package-4.2.0.0-local.jar` executable
- `gatk` script file.

We downloaded the BAM files and the corresponding index, dictionary, and reference genome files for a family that consisted of father, mother, and son *.bam files.

Cromwell engine

Cromwell is an open-source engine geared towards scientific workflows that enables workflow management. The Cromwell engine can be run in two [modes](#), Server mode or a single- workflow Run mode. The behavior of the Cromwell engine can be controlled using the [Cromwell engine configuration file](#).

- **Server mode.** Enables [RESTful](#) execution of workflows in Cromwell engine.
- **Run mode.** Run mode is best suited for executing single workflows in Cromwell, [ref](#) for a complete set of available options in Run mode.

We use the Cromwell engine to execute the workflows and pipelines at scale. The Cromwell engine uses a user-friendly [workflow description language](#) (WDL)-based scripting language. Cromwell also supports a second

workflow scripting standard called the common workflow language (CWL). Throughout this technical report, we used WDL. WDL was originally developed by the Broad Institute for genome analysis pipelines. Using the WDL workflows can be implemented using several strategies, including the following:

- **Linear chaining.** As the name suggests, output from task#1 is sent to task #2 as input.
- **Multi-in/out.** This is similar to linear chaining in that each task can have multiple outputs being sent as input to subsequent tasks.
- **Scatter-gather.** This is one of the most powerful enterprise application integration (EAI) strategies available, especially when used in event-driven architecture. Each task executes in a decoupled fashion, and the output for each task is consolidated into the final output.

There are three steps when WDL is used to run GATK in a standalone mode:

1. Validate syntax using `womtool.jar`.

```
[root@genomics1 ~]# java -jar womtool.jar validate ghplo.wdl
```

2. Generate inputs JSON.

```
[root@genomics1 ~]# java -jar womtool.jar inputs ghplo.wdl > ghplo.json
```

3. Run the workflow using the Cromwell engine and `Cromwell.jar`.

```
[root@genomics1 ~]# java -jar cromwell.jar run ghplo.wdl --inputs  
ghplo.json
```

The GATK can be executed by using several methods; this document explores three of these methods.

Execution of GATK using the jar file

Let's look at a single variant call pipeline execution using the Haplotype variant caller.

```
[root@genomics1 ~]# java -Dsamjdk.use_async_io_read_samtools=false \  
-Dsamjdk.use_async_io_write_samtools=true \  
-Dsamjdk.use_async_io_write_tribble=false \  
-Dsamjdk.compression_level=2 \  
-jar /mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-local.jar \  
HaplotypeCaller \  
--input /mnt/genomics/GATK/TEST\ DATA/bam/workshop_1906_2-  
germline_bams_father.bam \  
--output workshop_1906_2-germline_bams_father.validation.vcf \  
--reference /mnt/genomics/GATK/TEST\ DATA/ref/workshop_1906_2-  
germline_ref_ref.fasta
```


In this method of execution, we use the GATK local execution jar file, we use a single java command to invoke the jar file, and we pass several parameters to the command.

1. This parameter indicates that we are invoking the `HaplotypeCaller` variant caller pipeline.
2. `--input` specifies the input BAM file.
3. `--output` specifies the variant output file in variant call format (*.vcf) ([ref](#)).
4. With the `--reference` parameter, we are passing a reference genome.

Once executed, output details can be found in the section "[Output for execution of GATK using the jar file.](#)"

Execution of GATK using `./gatk` script

GATK tool kit can be executed using the `./gatk` script. Let's examine the following command:

```
[root@genomics1 execution]# ./gatk \
--java-options "-Xmx4G" \
HaplotypeCaller \
-I /mnt/genomics/GATK/TEST\ DATA/bam/workshop_1906_2-
germline_bams_father.bam \
-R /mnt/genomics/GATK/TEST\ DATA/ref/workshop_1906_2-
germline_ref_ref.fasta \
-O /mnt/genomics/GATK/TEST\ DATA/variants.vcf
```

We pass several parameters to the command.

- This parameter indicates that we are invoking the `HaplotypeCaller` variant caller pipeline.
- `-I` specifies the input BAM file.
- `-O` specifies the variant output file in variant call format (*.vcf) ([ref](#)).
- With the `-R` parameter, we are passing a reference genome.

Once executed, output details can be found in the section "[Output for execution of GATK using the `./gatk` script.](#)"

Execution of GATK using Cromwell engine

We use the Cromwell engine to manage GATK execution. Let's examine the command line and it's parameters.

```
[root@genomics1 genomics]# java -jar cromwell-65.jar \
run /mnt/genomics/GATK/seq/ghplo.wdl \
--inputs /mnt/genomics/GATK/seq/ghplo.json
```

Here, we invoke the Java command by passing the `-jar` parameter to indicate that we intend to execute a jar file, for example, `Cromwell-65.jar`. The next parameter passed (`run`) indicates that the Cromwell engine is running in Run mode, the other possible option is Server mode. The next parameter is `*.wdl` that the Run

mode should use to execute the pipelines. The next parameter is the set of input parameters to the workflows being executed.

Here's what the contents of the `ghplo.wdl` file look like:

```
[root@genomics1 seq]# cat ghplo.wdl
workflow helloHaplotypeCaller {
  call haplotypeCaller
}
task haplotypeCaller {
  File GATK
  File RefFasta
  File RefIndex
  File RefDict
  String sampleName
  File inputBAM
  File bamIndex
  command {
    java -jar ${GATK} \
      HaplotypeCaller \
      -R ${RefFasta} \
      -I ${inputBAM} \
      -O ${sampleName}.raw.indels.snps.vcf
  }
  output {
    File rawVCF = "${sampleName}.raw.indels.snps.vcf"
  }
}
[root@genomics1 seq]#
```

Here's the corresponding JSON file with the inputs to the Cromwell engine.

```
[root@genomics1 seq]# cat ghplo.json
{
  "helloHaplotypeCaller.haplotypeCaller.GATK": "/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-local.jar",
  "helloHaplotypeCaller.haplotypeCaller.RefFasta": "/mnt/genomics/GATK/TEST DATA/ref/workshop_1906_2-germline_ref_ref.fasta",
  "helloHaplotypeCaller.haplotypeCaller.RefIndex": "/mnt/genomics/GATK/TEST DATA/ref/workshop_1906_2-germline_ref_ref.fasta.fai",
  "helloHaplotypeCaller.haplotypeCaller.RefDict": "/mnt/genomics/GATK/TEST DATA/ref/workshop_1906_2-germline_ref_ref.dict",
  "helloHaplotypeCaller.haplotypeCaller.sampleName": "fatherbam",
  "helloHaplotypeCaller.haplotypeCaller.inputBAM": "/mnt/genomics/GATK/TEST DATA/bam/workshop_1906_2-germline_bams_father.bam",
  "helloHaplotypeCaller.haplotypeCaller.bamIndex": "/mnt/genomics/GATK/TEST DATA/bam/workshop_1906_2-germline_bams_father.bai"
}
[root@genomics1 seq]#
```

Please note that Cromwell uses an in-memory database for the execution. Once executed, the output log can be seen in the section ["Output for execution of GATK using the Cromwell engine."](#)

For a comprehensive set of steps on how to execute GATK, see the [GATK documentation](#).

[Next: Output for execution of GATK using the jar file.](#)

Output for execution of GATK using the jar file

[Previous: Genomics - GATK setup and execution.](#)

Execution of GATK using the jar file produced the following sample output.

```
[root@genomics1 execution]# java -Dsamjdk.use_async_io_read_samtools=false \
-Dsamjdk.use_async_io_write_samtools=true \
-Dsamjdk.use_async_io_write_tribble=false \
-Dsamjdk.compression_level=2 \
-jar /mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-local.jar \
HaplotypeCaller \
--input /mnt/genomics/GATK/TEST\ DATA/bam/workshop_1906_2-germline_bams_father.bam \
--output workshop_1906_2-germline_bams_father.validation.vcf \
--reference /mnt/genomics/GATK/TEST\ DATA/ref/workshop_1906_2-germline_ref_ref.fasta \
22:52:58.430 INFO NativeLibraryLoader - Loading libgkl_compression.so
from jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-local.jar!/com/intel/gkl/native/libgkl_compression.so
```

```

Aug 17, 2021 10:52:58 PM
shaded.cloud_nio.com.google.auth.oauth2.ComputeEngineCredentials
runningOnComputeEngine
INFO: Failed to detect whether we are running on Google Compute Engine.
22:52:58.541 INFO HaplotypeCaller -
-----
22:52:58.542 INFO HaplotypeCaller - The Genome Analysis Toolkit (GATK)
v4.2.0.0
22:52:58.542 INFO HaplotypeCaller - For support and documentation go to
https://software.broadinstitute.org/gatk/
22:52:58.542 INFO HaplotypeCaller - Executing as
root@genomics1.healthylife.fp on Linux v4.18.0-305.3.1.el8_4.x86_64 amd64
22:52:58.542 INFO HaplotypeCaller - Java runtime: OpenJDK 64-Bit Server
VM v1.8.0_302-b08
22:52:58.542 INFO HaplotypeCaller - Start Date/Time: August 17, 2021
10:52:58 PM EDT
22:52:58.542 INFO HaplotypeCaller -
-----
22:52:58.542 INFO HaplotypeCaller -
-----
22:52:58.542 INFO HaplotypeCaller - HTSJDK Version: 2.24.0
22:52:58.542 INFO HaplotypeCaller - Picard Version: 2.25.0
22:52:58.542 INFO HaplotypeCaller - Built for Spark Version: 2.4.5
22:52:58.542 INFO HaplotypeCaller - HTSJDK Defaults.COMPRESSION_LEVEL : 2
22:52:58.543 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_READ_FOR_SAMTOOLS : false
22:52:58.543 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_WRITE_FOR_SAMTOOLS : true
22:52:58.543 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_WRITE_FOR_TRIBBLE : false
22:52:58.543 INFO HaplotypeCaller - Deflater: IntelDeflater
22:52:58.543 INFO HaplotypeCaller - Inflater: IntelInflater
22:52:58.543 INFO HaplotypeCaller - GCS max retries/reopens: 20
22:52:58.543 INFO HaplotypeCaller - Requester pays: disabled
22:52:58.543 INFO HaplotypeCaller - Initializing engine
22:52:58.804 INFO HaplotypeCaller - Done initializing engine
22:52:58.809 INFO HaplotypeCallerEngine - Disabling physical phasing,
which is supported only for reference-model confidence output
22:52:58.820 INFO NativeLibraryLoader - Loading libgkl_utils.so from
jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar!/com/intel/gkl/native/libgkl_utils.so
22:52:58.821 INFO NativeLibraryLoader - Loading libgkl_pairhmm_omp.so
from jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar!/com/intel/gkl/native/libgkl_pairhmm_omp.so
22:52:58.854 INFO IntelPairHmm - Using CPU-supported AVX-512 instructions
22:52:58.854 INFO IntelPairHmm - Flush-to-zero (FTZ) is enabled when

```

```

running PairHMM
22:52:58.854 INFO   IntelPairHmm - Available threads: 16
22:52:58.854 INFO   IntelPairHmm - Requested threads: 4
22:52:58.854 INFO   PairHMM - Using the OpenMP multi-threaded AVX-
accelerated native PairHMM implementation
22:52:58.872 INFO   ProgressMeter - Starting traversal
22:52:58.873 INFO   ProgressMeter -           Current Locus   Elapsed Minutes
Regions Processed   Regions/Minute
22:53:00.733 WARN   InbreedingCoeff - InbreedingCoeff will not be
calculated at position 20:9999900 and possibly subsequent; at least 10
samples must have called genotypes
22:53:08.873 INFO   ProgressMeter -           20:17538652           0.2
58900              353400.0
22:53:17.681 INFO   HaplotypeCaller - 405 read(s) filtered by:
MappingQualityReadFilter
0 read(s) filtered by: MappingQualityAvailableReadFilter
0 read(s) filtered by: MappedReadFilter
0 read(s) filtered by: NotSecondaryAlignmentReadFilter
6628 read(s) filtered by: NotDuplicateReadFilter
0 read(s) filtered by: PassesVendorQualityCheckReadFilter
0 read(s) filtered by: NonZeroReferenceLengthAlignmentReadFilter
0 read(s) filtered by: GoodCigarReadFilter
0 read(s) filtered by: WellformedReadFilter
7033 total reads filtered
22:53:17.681 INFO   ProgressMeter -           20:63024652           0.3
210522             671592.9
22:53:17.681 INFO   ProgressMeter - Traversal complete. Processed 210522
total regions in 0.3 minutes.
22:53:17.687 INFO   VectorLoglessPairHMM - Time spent in setup for JNI call
: 0.010347438
22:53:17.687 INFO   PairHMM - Total compute time in PairHMM
computeLogLikelihoods() : 0.259172573
22:53:17.687 INFO   SmithWatermanAligner - Total compute time in java
Smith-Waterman : 1.27 sec
22:53:17.687 INFO   HaplotypeCaller - Shutting down engine
[August 17, 2021 10:53:17 PM EDT]
org.broadinstitute.hellbender.tools.walkers.haplotypecaller.HaplotypeCalle
r done. Elapsed time: 0.32 minutes.
Runtime.totalMemory()=5561122816
[root@genomics1 execution]#

```

Notice that the output file is located at the location specified after the execution.

Next: [Output for execution of GATK using the ./gatk script.](#)

Output for execution of GATK using the ./gatk script

[Previous: Output for execution of GATK using the jar file.](#)

The execution of GATK using the ./gatk script produced the following sample output.

```
[root@genomics1 gatk-4.2.0.0]# ./gatk --java-options "-Xmx4G" \
HaplotypeCaller \
-I /mnt/genomics/GATK/TEST\ DATA/bam/workshop_1906_2-
germline_bams_father.bam \
-R /mnt/genomics/GATK/TEST\ DATA/ref/workshop_1906_2-
germline_ref_ref.fasta \
-O /mnt/genomics/GATK/TEST\ DATA/variants.vcf
Using GATK jar /mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar
Running:
    java -Dsamjdk.use_async_io_read_samtools=false
-Dsamjdk.use_async_io_write_samtools=true
-Dsamjdk.use_async_io_write_tribble=false -Dsamjdk.compression_level=2
-Xmx4G -jar /mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-local.jar
HaplotypeCaller -I /mnt/genomics/GATK/TEST DATA/bam/workshop_1906_2-
germline_bams_father.bam -R /mnt/genomics/GATK/TEST
DATA/ref/workshop_1906_2-germline_ref_ref.fasta -O /mnt/genomics/GATK/TEST
DATA/variants.vcf
23:29:45.553 INFO NativeLibraryLoader - Loading libgkl_compression.so
from jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar!/com/intel/gkl/native/libgkl_compression.so
Aug 17, 2021 11:29:45 PM
shaded.cloud_nio.com.google.auth.oauth2.ComputeEngineCredentials
runningOnComputeEngine
INFO: Failed to detect whether we are running on Google Compute Engine.
23:29:45.686 INFO HaplotypeCaller -
-----
23:29:45.686 INFO HaplotypeCaller - The Genome Analysis Toolkit (GATK)
v4.2.0.0
23:29:45.686 INFO HaplotypeCaller - For support and documentation go to
https://software.broadinstitute.org/gatk/
23:29:45.687 INFO HaplotypeCaller - Executing as
root@genomics1.healthyliife.fp on Linux v4.18.0-305.3.1.el8_4.x86_64 amd64
23:29:45.687 INFO HaplotypeCaller - Java runtime: OpenJDK 64-Bit Server
VM v11.0.12+7-LTS
23:29:45.687 INFO HaplotypeCaller - Start Date/Time: August 17, 2021 at
11:29:45 PM EDT
23:29:45.687 INFO HaplotypeCaller -
-----
23:29:45.687 INFO HaplotypeCaller -
```

```

-----
23:29:45.687 INFO HaplotypeCaller - HTSJDK Version: 2.24.0
23:29:45.687 INFO HaplotypeCaller - Picard Version: 2.25.0
23:29:45.687 INFO HaplotypeCaller - Built for Spark Version: 2.4.5
23:29:45.688 INFO HaplotypeCaller - HTSJDK Defaults.COMPRESSION_LEVEL : 2
23:29:45.688 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_READ_FOR_SAMTOOLS : false
23:29:45.688 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_WRITE_FOR_SAMTOOLS : true
23:29:45.688 INFO HaplotypeCaller - HTSJDK
Defaults.USE_ASYNC_IO_WRITE_FOR_TRIBBLE : false
23:29:45.688 INFO HaplotypeCaller - Deflater: IntelDeflater
23:29:45.688 INFO HaplotypeCaller - Inflater: IntelInflater
23:29:45.688 INFO HaplotypeCaller - GCS max retries/reopens: 20
23:29:45.688 INFO HaplotypeCaller - Requester pays: disabled
23:29:45.688 INFO HaplotypeCaller - Initializing engine
23:29:45.804 INFO HaplotypeCaller - Done initializing engine
23:29:45.809 INFO HaplotypeCallerEngine - Disabling physical phasing,
which is supported only for reference-model confidence output
23:29:45.818 INFO NativeLibraryLoader - Loading libgkl_utils.so from
jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar!/com/intel/gkl/native/libgkl_utils.so
23:29:45.819 INFO NativeLibraryLoader - Loading libgkl_pairhmm_omp.so
from jar:file:/mnt/genomics/GATK/gatk-4.2.0.0/gatk-package-4.2.0.0-
local.jar!/com/intel/gkl/native/libgkl_pairhmm_omp.so
23:29:45.852 INFO IntelPairHmm - Using CPU-supported AVX-512 instructions
23:29:45.852 INFO IntelPairHmm - Flush-to-zero (FTZ) is enabled when
running PairHMM
23:29:45.852 INFO IntelPairHmm - Available threads: 16
23:29:45.852 INFO IntelPairHmm - Requested threads: 4
23:29:45.852 INFO PairHMM - Using the OpenMP multi-threaded AVX-
accelerated native PairHMM implementation
23:29:45.868 INFO ProgressMeter - Starting traversal
23:29:45.868 INFO ProgressMeter -          Current Locus  Elapsed Minutes
Regions Processed  Regions/Minute
23:29:47.772 WARN InbreedingCoeff - InbreedingCoeff will not be
calculated at position 20:9999900 and possibly subsequent; at least 10
samples must have called genotypes
23:29:55.868 INFO ProgressMeter -          20:18885652          0.2
63390          380340.0
23:30:04.389 INFO HaplotypeCaller - 405 read(s) filtered by:
MappingQualityReadFilter
0 read(s) filtered by: MappingQualityAvailableReadFilter
0 read(s) filtered by: MappedReadFilter
0 read(s) filtered by: NotSecondaryAlignmentReadFilter
6628 read(s) filtered by: NotDuplicateReadFilter

```

```

0 read(s) filtered by: PassesVendorQualityCheckReadFilter
0 read(s) filtered by: NonZeroReferenceLengthAlignmentReadFilter
0 read(s) filtered by: GoodCigarReadFilter
0 read(s) filtered by: WellformedReadFilter
7033 total reads filtered
23:30:04.389 INFO ProgressMeter - 20:63024652 0.3
210522 681999.9
23:30:04.389 INFO ProgressMeter - Traversal complete. Processed 210522
total regions in 0.3 minutes.
23:30:04.395 INFO VectorLoglessPairHMM - Time spent in setup for JNI call
: 0.012129203000000002
23:30:04.395 INFO PairHMM - Total compute time in PairHMM
computeLogLikelihoods() : 0.267345217
23:30:04.395 INFO SmithWatermanAligner - Total compute time in java
Smith-Waterman : 1.23 sec
23:30:04.395 INFO HaplotypeCaller - Shutting down engine
[August 17, 2021 at 11:30:04 PM EDT]
org.broadinstitute.hellbender.tools.walkers.haplotypecaller.HaplotypeCalle
r done. Elapsed time: 0.31 minutes.
Runtime.totalMemory()=2111832064
[root@genomics1 gatk-4.2.0.0]#

```

Notice that the output file is located at the location specified after the execution.

[Next: Output for execution of GATK using the Cromwell engine.](#)

Output for execution of GATK using the Cromwell engine

[Previous: Output for execution of GATK using the ./gatk script.](#)

The execution of GATK using the Cromwell engine produced the following sample output.

```

[root@genomics1 genomics]# java -jar cromwell-65.jar run
/mnt/genomics/GATK/seq/ghplo.wdl --inputs
/mnt/genomics/GATK/seq/ghplo.json
[2021-08-18 17:10:50,78] [info] Running with database db.url =
jdbc:hsqldb:mem:856a1f0d-9a0d-42e5-9199-
5e6c1d0f72dd;shutdown=false;hsqldb.tx=mvcc
[2021-08-18 17:10:57,74] [info] Running migration
RenameWorkflowOptionsInMetadata with a read batch size of 100000 and a
write batch size of 100000
[2021-08-18 17:10:57,75] [info] [RenameWorkflowOptionsInMetadata] 100%
[2021-08-18 17:10:57,83] [info] Running with database db.url =
jdbc:hsqldb:mem:6afe0252-2dc9-4e57-8674-
ce63c67aa142;shutdown=false;hsqldb.tx=mvcc
[2021-08-18 17:10:58,17] [info] Slf4jLogger started

```



```

[2021-08-18 17:10:58,33] [info] Workflow heartbeat configuration:
{
  "cromwellId" : "cromid-41b7e30",
  "heartbeatInterval" : "2 minutes",
  "ttl" : "10 minutes",
  "failureShutdownDuration" : "5 minutes",
  "writeBatchSize" : 10000,
  "writeThreshold" : 10000
}
[2021-08-18 17:10:58,38] [info] Metadata summary refreshing every 1
second.
[2021-08-18 17:10:58,38] [info] No metadata archiver defined in config
[2021-08-18 17:10:58,38] [info] No metadata deleter defined in config
[2021-08-18 17:10:58,40] [info] KvWriteActor configured to flush with
batch size 200 and process rate 5 seconds.
[2021-08-18 17:10:58,40] [info] WriteMetadataActor configured to flush
with batch size 200 and process rate 5 seconds.
[2021-08-18 17:10:58,44] [info] CallCacheWriteActor configured to flush
with batch size 100 and process rate 3 seconds.
[2021-08-18 17:10:58,44] [warn] 'docker.hash-lookup.gcr-api-queries-per-
100-seconds' is being deprecated, use 'docker.hash-lookup.gcr.throttle'
instead (see reference.conf)
[2021-08-18 17:10:58,54] [info] JobExecutionTokenDispenser - Distribution
rate: 50 per 1 seconds.
[2021-08-18 17:10:58,58] [info] SingleWorkflowRunnerActor: Version 65
[2021-08-18 17:10:58,58] [info] SingleWorkflowRunnerActor: Submitting
workflow
[2021-08-18 17:10:58,64] [info] Unspecified type (Unspecified version)
workflow 3e246147-b1a9-41dc-8679-319f81b7701e submitted
[2021-08-18 17:10:58,66] [info] SingleWorkflowRunnerActor: Workflow
submitted 3e246147-b1a9-41dc-8679-319f81b7701e
[2021-08-18 17:10:58,66] [info] 1 new workflows fetched by cromid-41b7e30:
3e246147-b1a9-41dc-8679-319f81b7701e
[2021-08-18 17:10:58,67] [info] WorkflowManagerActor: Starting workflow
3e246147-b1a9-41dc-8679-319f81b7701e
[2021-08-18 17:10:58,68] [info] WorkflowManagerActor: Successfully started
WorkflowActor-3e246147-b1a9-41dc-8679-319f81b7701e
[2021-08-18 17:10:58,68] [info] Retrieved 1 workflows from the
WorkflowStoreActor
[2021-08-18 17:10:58,70] [info] WorkflowStoreHeartbeatWriteActor
configured to flush with batch size 10000 and process rate 2 minutes.
[2021-08-18 17:10:58,76] [info] MaterializeWorkflowDescriptorActor
[3e246147]: Parsing workflow as WDL draft-2
[2021-08-18 17:10:59,34] [info] MaterializeWorkflowDescriptorActor
[3e246147]: Call-to-Backend assignments:
helloHaplotypeCaller.haplotypeCaller -> Local

```

```

[2021-08-18 17:11:00,54] [info] WorkflowExecutionActor-3e246147-b1a9-41dc-
8679-319f81b7701e [3e246147]: Starting
helloHaplotypeCaller.haplotypeCaller
[2021-08-18 17:11:01,56] [info] Assigned new job execution tokens to the
following groups: 3e246147: 1
[2021-08-18 17:11:01,70] [info] BackgroundConfigAsyncJobExecutionActor
[3e246147helloHaplotypeCaller.haplotypeCaller:NA:1]: java -jar
/mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-b1a9-41dc-
8679-319f81b7701e/call-haplotypeCaller/inputs/-179397211/gatk-package-
4.2.0.0-local.jar \
    HaplotypeCaller \
    -R /mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-
b1a9-41dc-8679-319f81b7701e/call-
haplotypeCaller/inputs/604632695/workshop_1906_2-germline_ref_ref.fasta \
    -I /mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-
b1a9-41dc-8679-319f81b7701e/call-
haplotypeCaller/inputs/604617202/workshop_1906_2-germline_bams_father.bam
\
    -O fatherbam.raw.indels.snps.vcf
[2021-08-18 17:11:01,72] [info] BackgroundConfigAsyncJobExecutionActor
[3e246147helloHaplotypeCaller.haplotypeCaller:NA:1]: executing: /bin/bash
/mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-b1a9-41dc-
8679-319f81b7701e/call-haplotypeCaller/execution/script
[2021-08-18 17:11:03,49] [info] BackgroundConfigAsyncJobExecutionActor
[3e246147helloHaplotypeCaller.haplotypeCaller:NA:1]: job id: 26867
[2021-08-18 17:11:03,53] [info] BackgroundConfigAsyncJobExecutionActor
[3e246147helloHaplotypeCaller.haplotypeCaller:NA:1]: Status change from -
to WaitingForReturnCode
[2021-08-18 17:11:03,54] [info] Not triggering log of token queue status.
Effective log interval = None
[2021-08-18 17:11:23,65] [info] BackgroundConfigAsyncJobExecutionActor
[3e246147helloHaplotypeCaller.haplotypeCaller:NA:1]: Status change from
WaitingForReturnCode to Done
[2021-08-18 17:11:25,04] [info] WorkflowExecutionActor-3e246147-b1a9-41dc-
8679-319f81b7701e [3e246147]: Workflow helloHaplotypeCaller complete.
Final Outputs:
{
    "helloHaplotypeCaller.haplotypeCaller.rawVCF": "/mnt/genomics/cromwell-
executions/helloHaplotypeCaller/3e246147-b1a9-41dc-8679-319f81b7701e/call-
haplotypeCaller/execution/fatherbam.raw.indels.snps.vcf"
}
[2021-08-18 17:11:28,43] [info] WorkflowManagerActor: Workflow actor for
3e246147-b1a9-41dc-8679-319f81b7701e completed with status 'Succeeded'.
The workflow will be removed from the workflow store.
[2021-08-18 17:11:32,24] [info] SingleWorkflowRunnerActor workflow
finished with status 'Succeeded'.

```

```

{
  "outputs": {
    "helloHaplotypeCaller.haplotypeCaller.rawVCF":
"/mnt/genomics/cromwell-executions/helloHaplotypeCaller/3e246147-b1a9-
41dc-8679-319f81b7701e/call-
haplotypeCaller/execution/fatherbam.raw.indels.snps.vcf"
  },
  "id": "3e246147-b1a9-41dc-8679-319f81b7701e"
}
[2021-08-18 17:11:33,45] [info] Workflow polling stopped
[2021-08-18 17:11:33,46] [info] 0 workflows released by cromid-41b7e30
[2021-08-18 17:11:33,46] [info] Shutting down WorkflowStoreActor - Timeout
= 5 seconds
[2021-08-18 17:11:33,46] [info] Shutting down WorkflowLogCopyRouter -
Timeout = 5 seconds
[2021-08-18 17:11:33,46] [info] Shutting down JobExecutionTokenDispenser -
Timeout = 5 seconds
[2021-08-18 17:11:33,46] [info] Aborting all running workflows.
[2021-08-18 17:11:33,46] [info] JobExecutionTokenDispenser stopped
[2021-08-18 17:11:33,46] [info] WorkflowStoreActor stopped
[2021-08-18 17:11:33,47] [info] WorkflowLogCopyRouter stopped
[2021-08-18 17:11:33,47] [info] Shutting down WorkflowManagerActor -
Timeout = 3600 seconds
[2021-08-18 17:11:33,47] [info] WorkflowManagerActor: All workflows
finished
[2021-08-18 17:11:33,47] [info] WorkflowManagerActor stopped
[2021-08-18 17:11:33,64] [info] Connection pools shut down
[2021-08-18 17:11:33,64] [info] Shutting down SubWorkflowStoreActor -
Timeout = 1800 seconds
[2021-08-18 17:11:33,64] [info] Shutting down JobStoreActor - Timeout =
1800 seconds
[2021-08-18 17:11:33,64] [info] Shutting down CallCacheWriteActor -
Timeout = 1800 seconds
[2021-08-18 17:11:33,64] [info] SubWorkflowStoreActor stopped
[2021-08-18 17:11:33,64] [info] Shutting down ServiceRegistryActor -
Timeout = 1800 seconds
[2021-08-18 17:11:33,64] [info] Shutting down DockerHashActor - Timeout =
1800 seconds
[2021-08-18 17:11:33,64] [info] Shutting down IoProxy - Timeout = 1800
seconds
[2021-08-18 17:11:33,64] [info] CallCacheWriteActor Shutting down: 0
queued messages to process
[2021-08-18 17:11:33,64] [info] JobStoreActor stopped
[2021-08-18 17:11:33,64] [info] CallCacheWriteActor stopped
[2021-08-18 17:11:33,64] [info] KvWriteActor Shutting down: 0 queued
messages to process

```

```
[2021-08-18 17:11:33,64] [info] IoProxy stopped
[2021-08-18 17:11:33,64] [info] WriteMetadataActor Shutting down: 0 queued
messages to process
[2021-08-18 17:11:33,65] [info] ServiceRegistryActor stopped
[2021-08-18 17:11:33,65] [info] DockerHashActor stopped
[2021-08-18 17:11:33,67] [info] Database closed
[2021-08-18 17:11:33,67] [info] Stream materializer shut down
[2021-08-18 17:11:33,67] [info] WDL HTTP import resolver closed
[root@genomics1 genomics]#
```

[Next: GPU setup.](#)

GPU setup

[Previous: Output for execution of GATK using the Cromwell engine.](#)

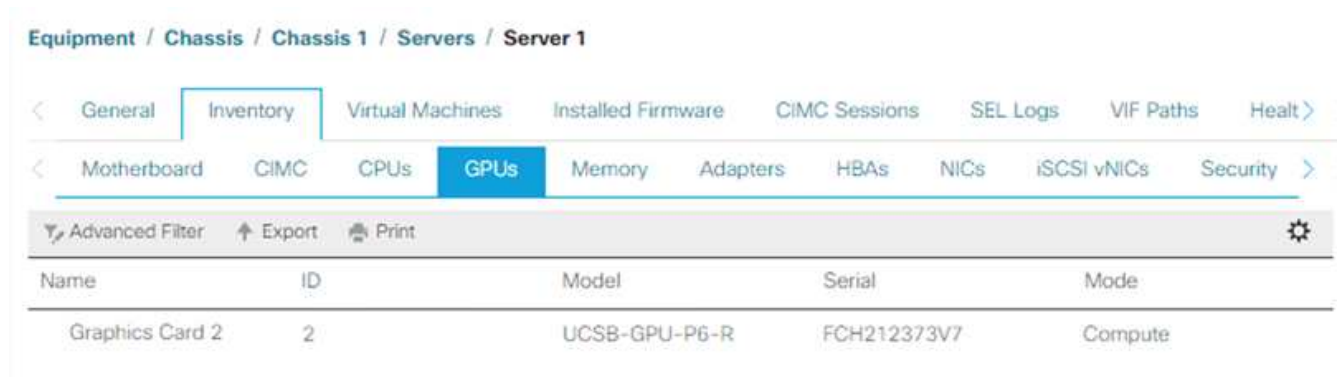
At the time of publication, the GATK tool does not have native support for GPU-based execution on premises. The following setup and guidance is provided to enable the readers understand how simple it is to use FlexPod with a rear-mounted NVIDIA Tesla P6 GPU using a PCIe mezzanine card for GATK.

We used the following Cisco-Validated Design (CVD) as the reference architecture and best-practice guide to set up the FlexPod environment so that we can run applications that use GPUs.

- [FlexPod Datacenter for AI/ML with Cisco UCS 480 ML for Deep Learning](#)

Here is a set of key takeaways during this setup:

1. We used a PCIe NVIDIA Tesla P6 GPU in a mezzanine slot in the UCS B200 M5 servers.



The screenshot shows the Cisco UCS Manager interface. The breadcrumb navigation at the top reads: Equipment / Chassis / Chassis 1 / Servers / Server 1. Below this, there are several tabs: General, Inventory, Virtual Machines, Installed Firmware, CIMC Sessions, SEL Logs, VIF Paths, and Health. The 'Inventory' tab is selected, and within it, the 'GPUs' sub-tab is active. Above the table, there are options for 'Advanced Filter', 'Export', and 'Print'. The table itself has five columns: Name, ID, Model, Serial, and Mode. One row is visible, representing a graphics card.

| Name | ID | Model | Serial | Mode |
|-----------------|----|---------------|-------------|---------|
| Graphics Card 2 | 2 | UCSB-GPU-P6-R | FCH212373V7 | Compute |

Equipment / Chassis / Chassis 1 / Servers / Server 2

< General **Inventory** Virtual Machines Installed Firmware CIMC Sessions SEL Logs VIF Paths Health >

< Motherboard CIMC CPUs **GPUs** Memory Adapters HBAs NICs iSCSI vNICs Security >

Advanced Filter Export Print

| Name | ID | Model | Serial | Mode |
|-----------------|----|---------------|-------------|---------|
| Graphics Card 2 | 2 | UCSB-GPU-P6-R | FCH212373Y1 | Compute |

- For this setup, we registered on the NVIDIA partner portal and obtained an evaluation license (also known as an entitlement) to be able to use the GPUs in compute mode.
- We downloaded the NVIDIA vGPU software required from the NVIDIA partner website.
- We downloaded the entitlement *.bin file from the NVIDIA partner website.
- We installed an NVIDIA vGPU license server and added the entitlements to the license server using the *.bin file downloaded from the NVIDIA partner site.
- Make sure to choose the correct NVIDIA vGPU software version for your deployment on the NVIDIA partner portal. For this setup we used driver version 460.73.02.
- This command installs the [NVIDIA vGPU Manager](#) in ESXi.

```
[root@localhost:~] esxcli software vib install -v
/vmfs/volumes/infra_datastore_nfs/nvidia/vib/NVIDIA_bootbank_NVIDIA-
VMware_ESXi_7.0_Host_Driver_460.73.02-1OEM.700.0.0.15525992.vib
Installation Result
Message: Operation finished successfully.
Reboot Required: false
VIBs Installed: NVIDIA_bootbank_NVIDIA-
VMware_ESXi_7.0_Host_Driver_460.73.02-1OEM.700.0.0.15525992
VIBs Removed:
VIBs Skipped:
```

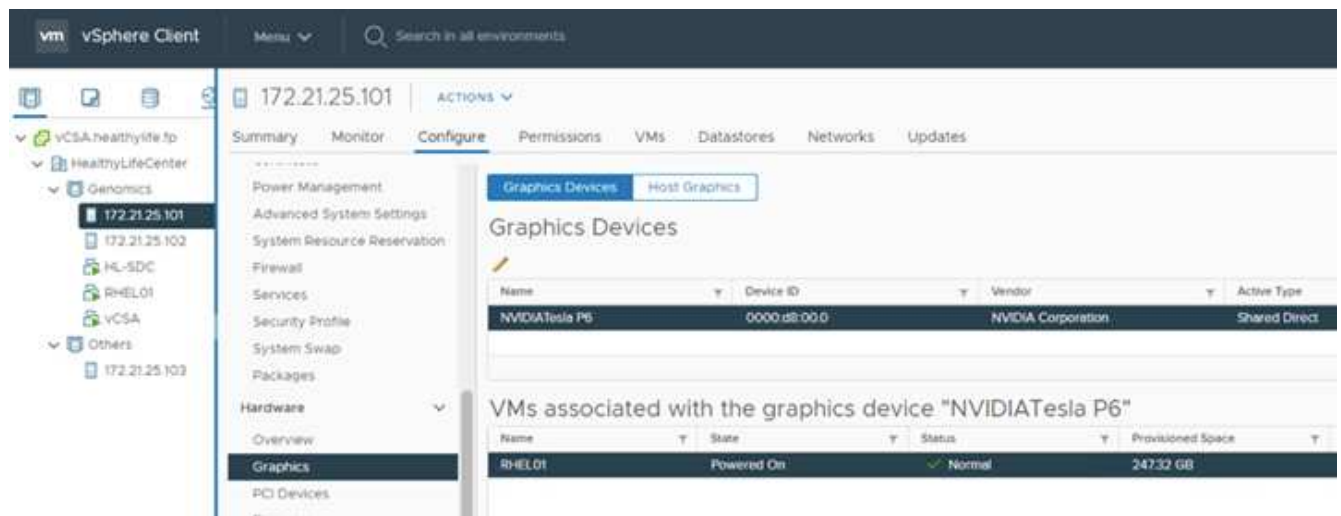
- After rebooting the ESXi server, run the following command to validate the installation and check the health of the GPUs.

```

[root@localhost:~] nvidia-smi
Wed Aug 18 21:37:19 2021
+-----+
+-----+
| NVIDIA-SMI 460.73.02      Driver Version: 460.73.02      CUDA Version: N/A
|
|-----+-----+
+-----+
| GPU   Name               Persistence-M| Bus-Id        Disp.A | Volatile
Uncorr. ECC |
| Fan   Temp   Perf   Pwr:Usage/Cap|      Memory-Usage | GPU-Util
Compute M. |
|
|
MIG M. |
|=====+=====+=====
=====|
|    0   Tesla P6             On      | 00000000:D8:00.0 Off |
0 |
| N/A    35C    P8      9W /  90W |  15208MiB / 15359MiB |      0%
Default |
|
|
N/A |
+-----+-----+
+-----+
+-----+
+-----+
+-----+
| Processes:
|
| GPU   GI    CI          PID    Type    Process name                  GPU
Memory |
|      ID    ID                 |          Usage
|
|=====+=====+=====
=====|
|    0   N/A   N/A     2812553      C+G     RHEL01
15168MiB |
+-----+-----+
+-----+
[root@localhost:~]

```

9. Using vCenter, [configure](#) the graphics device settings to “Shared Direct.”



10. Make sure that secure boot is disabled for the RedHat VM.
11. Make sure that the VM Boot Options firmware is set to EFI ([ref](#)).


| | |
|---------------------------------|---|
| > General Options | VM Name: RHEL01 |
| > VMware Remote Console Options | <input type="checkbox"/> Lock the guest operating system when the last remote user disconnects |
| > Encryption | Expand for encryption settings |
| > Power management | Expand for power management settings |
| > VMware Tools | Expand for VMware Tools settings |
| > Boot Options | |
| Firmware | EFI (recommended) ▼ |
| Secure Boot | <input type="checkbox"/> Enabled |
| Boot Delay | When powering on or resetting, delay boot order by 0 milliseconds |
| Force EFI setup | <input type="checkbox"/> During the next boot, force entry into the EFI setup screen |
| Failed Boot Recovery | <input type="checkbox"/> If the VM fails to find boot device, automatically retry after 10 seconds |
| > Advanced | Expand for advanced settings |
| > Fibre Channel NPIV | Expand for Fibre Channel NPIV settings |

CANCEL

OK

12. Make sure that the following PARAMS are added to the VM Options advanced Edit Configuration. The value of the `pciPassthru.64bitMMIOSizeGB` parameter depends on the GPU memory and number of GPUs assigned to the VM. For example:

- If a VM is assigned 4 x 32GB V100 GPUs, then this value should be 128.
- If a VM is assigned 4 x 16GB P6 GPUs, then this value should be 64.



15. On the RedHat Linux VM, NVIDIA drivers can be installed by running the following command:

```
[root@genomics1 genomics]# sh NVIDIA-Linux-x86_64-460.73.01-grid.run
```

16. Verify that the correct vGPU profile is being reported by running the following command:

```
[root@genomics1 genomics]# nvidia-smi -query-gpu=gpu_name  
-format=csv,noheader -id=0 | sed -e 's/ /-/g'  
GRID-P6-16C  
[root@genomics1 genomics]#
```

17. After reboot, verify that the correct NVIDIA vGPU are reported along with the driver versions.

```

[root@genomics1 genomics]# nvidia-smi
Wed Aug 18 20:30:56 2021
+-----+
+-----+
| NVIDIA-SMI 460.73.01      Driver Version: 460.73.01      CUDA Version:
11.2      |
|-----+-----+
+-----+
| GPU  Name                Persistence-M| Bus-Id        Disp.A | Volatile
Uncorr. ECC |
| Fan  Temp   Perf   Pwr:Usage/Cap|      Memory-Usage | GPU-Util
Compute M. |
|
MIG M. |
|=====+=====+=====
=====|
|   0  GRID P6-16C          On    | 00000000:02:02.0 Off |
N/A |
| N/A   N/A     P8      N/A /  N/A |   2205MiB / 16384MiB |      0%
Default |
|
N/A |
+-----+-----+
+-----+
+-----+
+-----+
+-----+
| Processes:
|
| GPU    GI    CI          PID    Type    Process name                        GPU
Memory |
|          ID    ID                                   Usage
|
|=====+=====+=====
=====|
|   0    N/A   N/A         8604      G    /usr/libexec/Xorg
13MiB |
+-----+-----+
+-----+
[root@genomics1 genomics]#

```

18. Make sure that the license server IP is configured on the VM in the vGPU grid configuration file.

a. Copy the template.


```
[root@genomics1 genomics]# cp /etc/nvidia/gridd.conf.template  
/etc/nvidia/gridd.conf
```

- b. Edit the file `/etc/nvidia/rid.conf`, add the license server IP address, and set the feature type to 1.

```
ServerAddress=192.168.169.10
```

```
FeatureType=1
```

19. After restarting the VM, you should see an entry under Licensed Clients in the license server as shown below.



20. Refer to the Solutions Setup section for more information on downloading the GATK and Cromwell software.
21. After GATK can use GPUs on premises, the workflow description language *.wdl has the runtime attributes as shown below.

```

task ValidateBAM {
  input {
    # Command parameters
    File input_bam
    String output_basename
    String? validation_mode
    String gatk_path
    # Runtime parameters
    String docker
    Int machine_mem_gb = 4
    Int additional_disk_space_gb = 50
  }
  Int disk_size = ceil(size(input_bam, "GB")) + additional_disk_space_gb
  String output_name = "${output_basename}_${validation_mode}.txt"
  command {
    ${gatk_path} \
      ValidateSamFile \
      --INPUT ${input_bam} \
      --OUTPUT ${output_name} \
      --MODE ${default="SUMMARY" validation_mode}
  }
  runtime {
    gpuCount: 1
    gpuType: "nvidia-tesla-p6"
    docker: docker
    memory: machine_mem_gb + " GB"
    disks: "local-disk " + disk_size + " HDD"
  }
  output {
    File validation_report = "${output_name}"
  }
}

```

[Next: Conclusion.](#)

Conclusion

[Previous: GPU setup.](#)

Many healthcare organizations around the world have standardized on FlexPod as a common platform. With FlexPod, you can deploy healthcare capabilities with confidence. FlexPod with NetApp ONTAP comes standard with the ability to implement an industry leading set of protocols out of the box. Irrespective of the origin of the request to run genomics of a given patient, interoperability, accessibility, availability, and scalability come standard with a FlexPod platform. When standardized on a FlexPod platform, the culture

of innovation becomes contagious.

Where to find additional information

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

- FlexPod Datacenter for AI/ML with Cisco UCS 480 ML for Deep Learning

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_480ml_aiml_deployment.pdf

- FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/fp_vmware_vsphere_7_0_ontap_9_7.html

- ONTAP 9 Documentation Center

<http://docs.netapp.com>

- Agile and efficient—how FlexPod drives data center modernization

<https://www.flexpod.com/idc-white-paper/>

- AI in healthcare

<https://www.netapp.com/us/media/na-369.pdf>

- FlexPod for healthcare Ease Your Transformation

<https://flexpod.com/solutions/verticals/healthcare/>

- FlexPod from Cisco and NetApp

<https://flexpod.com/>

- AI and Analytics for healthcare (NetApp)

<https://www.netapp.com/us/artificial-intelligence/healthcare-ai-analytics/index.aspx>

- AI in healthcare Smart infrastructure Choices Increase Success

<https://www.netapp.com/pdf.html?item=/media/7410-wp-7314.pdf>

- FlexPod Datacenter with ONTAP 9.8, ONTAP Storage Connector for Cisco Intersight, and Cisco Intersight Managed Mode.

<https://www.netapp.com/pdf.html?item=/media/25001-tr-4883.pdf>

- FlexPod Datacenter with Red Hat Enterprise Linux OpenStack Platform

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_openstack_osp6.html

Version history

| Version | Date | Document version history |
|-------------|---------------|--------------------------|
| Version 1.0 | November 2021 | Initial release. |

FlexPod for MEDITECH Directional Sizing Guide

TR-4774: FlexPod for MEDITECH Directional Sizing

Brandon Agee, John Duignan, NetApp
Mike Brennan, Jon Ebmeir, Cisco

In partnership with:



This report provides guidance for sizing FlexPod for a MEDITECH EHR application software environment.

Purpose

FlexPod systems can be deployed to host MEDITECH EXPANSE, 6.x, 5.x, and MAGIC services. FlexPod servers that host the MEDITECH application layer provide an integrated platform for a dependable, high-performance infrastructure. The FlexPod integrated platform is deployed rapidly by skilled FlexPod channel partners and is supported by Cisco and NetApp technical assistance centers.

Sizing is based on information in MEDITECH's hardware configuration proposal and the MEDITECH task document. The goal is to determine the optimal size for compute, network, and storage infrastructure components.

The [MEDITECH Workload Overview](#) section describes the types of compute and storage workloads that can be found in MEDITECH environments.

The [Technical Specifications for Small, Medium, and Large Architectures](#) section details a sample Bill of Materials for the different storage architectures described in the section. The configurations given are general guidelines only. Always size the systems using the sizers based on the workload and tune the configurations accordingly.

Overall solution benefits

Running a MEDITECH environment on the FlexPod architectural foundation can help healthcare organizations improve productivity and decrease capital and operating expenses. FlexPod provides a prevalidated, rigorously tested, converged infrastructure from the strategic partnership of Cisco and NetApp. It is engineered and designed specifically for delivering predictable low-latency system performance and high availability. This approach results in faster response time for users of the MEDITECH EHR system.

The FlexPod solution from Cisco and NetApp meets MEDITECH system requirements with a high performing, modular, prevalidated, converged, virtualized, efficient, scalable, and cost-effective platform. FlexPod Datacenter with MEDITECH delivers several benefits specific to the healthcare industry:

- **Modular architecture.** FlexPod addresses the various needs of the MEDITECH modular architecture with customized FlexPod systems for each specific workload. All components are connected through a clustered server and storage management fabric and use a cohesive management toolset.
- **Simplified operations and lowered costs.** You can eliminate the expense and complexity of legacy platforms by replacing them with a more efficient and scalable shared resource that can support clinicians wherever they are. This solution delivers better resource usage for greater return on investment (ROI).
- **Quicker deployment of infrastructure.** The integrated design of FlexPod Datacenter with MEDITECH enables customers to have the new infrastructure up and running quickly and easily for both on-site and remote data centers.
- **Scale-out architecture.** You can scale SAN and NAS from terabytes to tens of petabytes without reconfiguring running applications.
- **Nondisruptive operations.** You can perform storage maintenance, hardware lifecycle operations, and software upgrades without interrupting the business.
- **Secure multitenancy.** This benefit supports the increased needs of virtualized server and shared storage infrastructure, enabling secure multitenancy of facility-specific information. This benefit is important if you are hosting multiple instances of databases and software.
- **Pooled resource optimization.** This benefit can help reduce physical server and storage controller counts, load balance workload demands, boost utilization, and simultaneously improve performance.
- **Quality of service (QoS).** FlexPod offers quality of service (QoS) on the entire stack. Industry-leading QoS storage policies enable differentiated service levels in a shared environment. These policies enable optimal performance for workloads and help in isolating and controlling runaway applications.
- **Storage efficiency.** You can reduce storage costs with NetApp 7:1 storage efficiency.
- **Agility.** The industry-leading workflow automation, orchestration, and management tools offered by FlexPod systems allow IT to be far more responsive to business requests. These business requests can range from MEDITECH backup and provisioning of more testing and training environments to analytics database replications for population health management initiatives.
- **Productivity.** You can quickly deploy and scale this solution for optimal clinician end-user experiences.
- **Data Fabric.** The NetApp Data Fabric architecture weaves data together across sites, beyond physical boundaries, and across applications. The NetApp Data Fabric is built for data-driven enterprises in a data-centric world. Data is created and used in multiple locations, and is often shared with applications and infrastructures. Data Fabric provides a way to manage data that is consistent and integrated. It also offers IT more control of the data and simplifies ever-increasing IT complexity.

Scope

This document covers environments that use Cisco UCS and NetApp ONTAP based storage. It provides sample reference architectures for hosting MEDITECH.

It does not cover:

- Detailed sizing guidance using NetApp System Performance Modeler (SPM) or other NetApp sizing tools.
- Sizing for nonproduction workloads.

Audience

This document is intended for NetApp and partner systems engineers and NetApp Professional Services personnel. NetApp assumes that the reader has a good understanding of compute and storage sizing concepts as well as technical familiarity with Cisco UCS and NetApp storage systems.

Related Documents

The following technical reports and other documents are relevant to this Technical Report, and make up a complete set of documents required for sizing, designing, and deploying MEDITECH on FlexPod infrastructure.

- [TR-4753: FlexPod Datacenter for MEDITECH Deployment Guide](#)
- [TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#)
- [TR-4319: NetApp Deployment Guidelines for MEDITECH Environments](#)

 Login credentials for the NetApp Field Portal are required to access some of these reports.

MEDITECH Workload Overview

This section describes the types of compute and storage workloads that you might find in MEDITECH environments.

MEDITECH and backup workloads

When you size NetApp storage systems for MEDITECH environments, you must consider both the MEDITECH production workload and the backup workload.

| MEDITECH Host |
|---|
| A MEDITECH host is a database server. This host is also referred to as a MEDITECH file server (for the EXPANSE, 6.x or C/S 5.x platform) or a MAGIC machine (for the MAGIC platform). This document uses the term MEDITECH host to refer to a MEDITECH file server and a MAGIC machine. |

The following sections describe the I/O characteristics and performance requirements of these two workloads.

MEDITECH workload

In a MEDITECH environment, multiple servers that run MEDITECH software perform various tasks as an integrated system known as the MEDITECH system. For more information about the MEDITECH system, see the MEDITECH documentation:

- For production MEDITECH environments, consult the appropriate MEDITECH documentation to determine the number of MEDITECH hosts and the storage capacity that must be included as part of sizing the NetApp storage system.
- For new MEDITECH environments, consult the hardware configuration proposal document. For existing MEDITECH environments, consult the hardware evaluation task document. The hardware evaluation task is associated with a MEDITECH ticket. Customers can request either of these documents from MEDITECH.

You can scale the MEDITECH system to provide increased capacity and performance by adding hosts. Each host requires storage capacity for its database and application files. The storage available to each MEDITECH host must also support the I/O generated by the host. In a MEDITECH environment, a LUN is available for each host to support that host’s database and application storage requirements. The type of MEDITECH category and the type of platform that you deploy determines the workload characteristics of each MEDITECH host and, therefore, of the system as a whole.

MEDITECH Categories

MEDITECH associates the deployment size with a category number ranging from 1 to 6. Category 1 represents the smallest MEDITECH deployments; category 6 represents the largest. Examples of the MEDITECH application specification associated with each category include metrics such as:

- Number of hospital beds
- Inpatients per year
- Outpatients per year
- Emergency room visits per year
- Exams per year
- Inpatient prescriptions per day
- Outpatient prescriptions per day

For more information about MEDITECH categories, see the MEDITECH category reference sheet. You can obtain this sheet from MEDITECH through the customer or through the MEDITECH system installer.

MEDITECH Platforms

MEDITECH has four platforms:

- EXPANSE
- MEDITECH 6.x
- Client/Server 5.x (C/S 5.x)
- MAGIC

For the MEDITECH EXPANSE, 6.x and C/S 5.x platforms, the I/O characteristics of each host are defined as 100% random with a request size of 4,000. For the MEDITECH MAGIC platform, each host's I/O characteristics are defined as 100% random with a request size of either 8,000 or 16,000. According to MEDITECH, the request size for a typical MAGIC production deployment is either 8,000 or 16,000.

The ratio of reads and writes varies depending on the platform that is deployed. MEDITECH estimates the average mix of read and write and then expresses them as percentages. MEDITECH also estimates the average sustained IOPS value required for each MEDITECH host on a particular MEDITECH platform. The table below summarizes the platform-specific I/O characteristics that are provided by MEDITECH.

| MEDITECH Category | MEDITECH Platform | Average Random Read % | Average Random Write % | Average Sustained IOPS per MEDITECH Host |
|-------------------|-------------------|-----------------------|------------------------|--|
| 1 | EXPANSE, 6.x | 20 | 80 | 750 |
| 2-6 | EXPANSE | 20 | 80 | 750 |
| | 6.x | 20 | 80 | 750 |
| | C/S 5.x | 40 | 60 | 600 |
| | MAGIC | 90 | 10 | 400 |

In a MEDITECH system, the average IOPS level of each host must equal the IOPS values defined in the above table. To determine the correct storage sizing based on each platform, the IOPS values specified in the above table are used as part of the sizing methodology described in the [Technical Specifications for Small, Medium](#)

and [Large Architectures](#) section.

MEDITECH requires the average random write latency to stay below 1ms for each host. However, temporary increases of write latency up to 2ms during backup and reallocation jobs are considered acceptable. MEDITECH also requires the average random read latency to stay below 7ms for category 1 hosts and below 5ms for category 2 hosts. These latency requirements apply to every host regardless of which MEDITECH platform is being used.

The table below summarizes the I/O characteristics that you must consider when you size NetApp storage for MEDITECH workloads.

| Parameter | MEDITECH Category | EXPANSE | MEDITECH 6.x | C/S 5.x | MAGIC |
|------------------------------|-------------------|---------------------|---------------------|---------------------|---------------------|
| Request size | 1-6 | 4K | 4K | 4K | 8K or 16K |
| Random/sequential | | 100% random | 100% random | 100% random | 100% random |
| Average sustained IOPS | 1 | 750 | 750 | N/A | N/A |
| | 2-6 | 750 | 750 | 600 | 400 |
| Read/write ratio | 1-6 | 20% read, 80% write | 20% read, 80% write | 40% read, 60% write | 90% read, 10% write |
| Write latency | | <1ms | <1ms | <1ms | <1ms |
| Temporary peak write latency | 1-6 | <2ms | <2ms | <2ms | <2ms |
| Read latency | 1 | <7ms | <7ms | N/A | N/A |
| | 2-6 | <5ms | <5ms | <5ms | <5ms |



MEDITECH hosts in categories 3 through 6 have the same I/O characteristics as category 2. For MEDITECH categories 2 through 6, the number of hosts that are deployed in each category differs.

The NetApp storage system should be sized to satisfy the performance requirements described in previous sections. In addition to the MEDITECH production workload, the NetApp storage system must be able to maintain these MEDITECH performance targets during backup operations, as described in the following section.

Backup Workload Description

MEDITECH certified backup software backs up the LUN used by each MEDITECH host in a MEDITECH system. For the backups to be in an application-consistent state, the backup software quiesces the MEDITECH system and suspends I/O requests to disk. While the system is in a quiesced state, the backup software issues a command to the NetApp storage system to create a NetApp Snapshot copy of the volumes that contain the LUNs. The backup software later unquiesces the MEDITECH system, which enables production I/O requests to continue to the database. The software creates a NetApp FlexClone volume based on the Snapshot copy. This volume is used by the backup source while production I/O requests continue on the parent volumes that host the LUNs.

The workload that is generated by the backup software comes from the sequential reading of the LUNs that reside in the FlexClone volumes. The workload is defined as a 100% sequential read workload with a request

size of 64,000. For the MEDITECH production workload, the performance criterion is to maintain the required IOPS and the associated read/write latency levels. For the backup workload, however, the attention is shifted to the overall data throughput (MBps) that is generated during the backup operation. MEDITECH LUN backups are required to be completed in an eight-hour backup window, but NetApp recommends that the backup of all MEDITECH LUNs be completed in six hours or less. Aiming to complete the backup in less than six hours mitigates for events such as an unplanned increase in the MEDITECH workload, NetApp ONTAP background operations, or data growth over time. Any of these events might incur extra backup time. Regardless of the amount of application data stored, the backup software performs a full block-level backup of the entire LUN for each MEDITECH host.

Calculate the sequential read throughput that is required to complete the backup within this window as a function of the other factors involved:

- The desired backup duration
- The number of LUNs
- The size of each LUN to be backed up

For example, in a 50-host MEDITECH environment in which each host’s LUN size is 200GB, the total LUN capacity to backup is 10TB.

To back up 10TB of data in eight hours, the following throughput is required:

- = $(10 \times 10^6) \text{MB} (8 \times 3,600) \text{s}$
- = 347.2MBps

However, to account for unplanned events, a conservative backup window of 5.5 hours is selected to provide headroom beyond the six hours that is recommended.

To back up 10TB of data in eight hours, the following throughput is required:

- = $(10 \times 10^6) \text{MB} (5.5 \times 3,600) \text{s}$
- = 500MBps

At the throughput rate of 500MBps, the backup can complete within a 5.5-hour time frame, comfortably within the 8-hour backup requirement.

The table below summarizes the I/O characteristics of the backup workload to use when you size the storage system.

| Parameter | All Platforms |
|--------------------------|--|
| Request size | 64K |
| Random/sequential | 100% sequential |
| Read/write ratio | 100% read |
| Average throughput | Depends on the number of MEDITECH hosts and the size of each LUN: Backup must complete within 8 hours. |
| Required backup duration | 8 hours |

Cisco UCS Reference Architecture for MEDITECH

The architecture for MEDITECH on FlexPod is based on guidance from MEDITECH, Cisco, and NetApp and on partner experience in working with MEDITECH customers of all sizes. The architecture is adaptable and applies best practices for MEDITECH, depending on the customer's data center strategy: whether that is small or large, centralized, distributed, or multitenant.

When deploying MEDITECH, Cisco has designed Cisco UCS reference architectures that align directly with MEDITECH's best practices. Cisco UCS delivers a tightly integrated solution for high performance, high availability, reliability, and scalability to support physician practices and hospital systems with several thousand beds.

Technical specifications for small, medium and large architectures

This section discusses a sample Bill of Materials for different size storage architectures.

Bill of material for small, medium, and large architectures.

The FlexPod design is a flexible infrastructure that encompasses many different components and software versions. Use [TR-4036: FlexPod Technical Specifications](#) as a guide to assembling a valid FlexPod configuration. The configurations in the table below are the minimum requirements for FlexPod, and are just a sample. The configuration can be expanded for each product family as required for different environments and use cases.

For this sizing exercise small corresponds to a Category 3 MEDITECH environment, medium to a Category 5, and large to a Category 6.

| | Small | Medium | Large |
|---|--|-----------------------------|--|
| Platform | One NetApp AFF A220 all-flash storage system HA pair | One NetApp AFF A220 HA pair | One NetApp AFF A300 all-flash storage system HA pair |
| Disk shelves | 9TB x 3.8TB | 13TB x 3.8TB | 19TB x 3.8TB |
| MEDITECH database size | 3TB-12TB | 17TB | >30TB |
| MEDITECH IOPS | <22,000 IOPs | >25,000 IOPs | >32,000 IOPs |
| Total IOPS | 22000 | 27000 | 35000 |
| Raw | 34.2TB | 49.4TB | 68.4TB |
| Usable capacity | 18.53TiB | 27.96TiB | 33.82TiB |
| Effective capacity (2:1 storage efficiency) | 55.6TiB | 83.89TiB | 101.47TiB |



Some customer environments might have multiple MEDITECH production workloads running simultaneously or might have higher IOPS requirements. In such cases, work with the NetApp account team to size the storage systems according to the required IOPS and capacity. You should be able to determine the right platform to serve the workloads. For example, there are customers successfully running multiple MEDITECH environments on a NetApp AFF A700 all-flash storage system HA pair.

The following table shows the standard software required for MEDITECH configurations.

| Software | Product family | Version or release | Details |
|------------|--------------------------------------|-------------------------------------|-----------------------------|
| Storage | ONTAP | ONTAP 9.4 general availability (GA) | |
| Network | Cisco UCS fabric interconnects | Cisco UCSM 4.x | Current recommended release |
| | Cisco Nexus Ethernet switches | 7.0(3)I7(6) | Current recommended release |
| | Cisco FC: Cisco MDS 9132T | 8.3(2) | Current recommended release |
| Hypervisor | Hypervisor | VMware vSphere ESXi 6.7 | |
| | Virtual machines (VMs) | Windows 2016 | |
| Management | Hypervisor management system | VMware vCenter Server 6.7 U1 (VCSA) | |
| | NetApp Virtual Storage Console (VSC) | VSC 7.0P1 | |
| | NetApp SnapCenter | SnapCenter 4.0 | |
| | Cisco UCS Manager | 4.x | |

The following table shows an small (category 3) example configuration – infrastructure components.

| Layer | Product family | Quantity and model | Details |
|---------|---|--------------------|--|
| Compute | Cisco UCS 5108 Chassis | 1 | Supports up to eight half-width or four full-width blades. Add chassis as server requirement grows. |
| | Cisco Chassis I/O Modules | 2 x 2208 | 8GB x 10GB uplink ports |
| | Cisco UCS blade servers | 4 x B200 M5 | Each with 2 x 14 cores, 2.6GHz or higher clock speed, and 384GB BIOS 3.2(3#) |
| | Cisco UCS Virtual Interface Cards | 4 x UCS 1440 | VMware ESXi fNIC FC driver: 1.6.0.47 VMware ESXi eNIC Ethernet driver: 1.0.27.0 (See interoperability matrix: https://ucshcltool.cloudapps.cisco.com/public/) |
| | 2 x Cisco UCS Fabric Interconnects (FI) | 2 x UCS 6454 FI | 4th-generation fabric interconnects supporting 10/25/100GB Ethernet and 32GB FC |

| Layer | Product family | Quantity and model | Details |
|-----------------|--|---------------------|--|
| Network | Cisco Ethernet switches | 2 x Nexus 9336c-FX2 | 1GB, 10GB, 25GB, 40GB, 100GB |
| Storage network | IP Network Nexus 9k for BLOB storage | | FI and UCS chassis |
| | FC: Cisco MDS 9132T | | Two Cisco 9132T switches |
| Storage | NetApp AFF A300 all-flash storage system | 1 HA Pair | 2-node cluster for all MEDITECH workloads (File Server, Image Server, SQL Server, VMware, and so on) |
| | DS224C disk shelf | 1 DS224C disk shelf | |
| | Solid-state drive (SSD) | 9 x 3.8TB | |

The following table shows medium (category 5) example configuration – Infrastructure components

| Layer | Product family | Quantity and model | Details |
|-----------------|---|---------------------|---|
| Compute | Cisco UCS 5108 chassis | 1 | Supports up to eight half-width or four full-width blades. Add chassis as server requirement grows. |
| | Cisco chassis I/O modules | 2 x 2208 | 8GB x 10GB uplink ports |
| | Cisco UCS blade servers | 6 x B200 M5 | Each with 2 x 16 cores, 2.5GHz/or higher clock speed, and 384GB or more memory BIOS 3.2(3#) |
| | Cisco UCS virtual interface card (VIC) | 6 x UCS 1440 VICs | VMware ESXi fNIC FC driver: 1.6.0.47 VMware ESXi eNIC Ethernet driver: 1.0.27.0 (See interoperability matrix:) |
| | 2 x Cisco UCS Fabric Interconnects (FI) | 2 x UCS 6454 FI | 4th-generation fabric interconnects supporting 10GB/25GB/100GB Ethernet and 32GB FC |
| Network | Cisco Ethernet switches | 2 x Nexus 9336c-FX2 | 1GB, 10GB, 25GB, 40GB, 100GB |
| Storage network | IP Network Nexus 9k for BLOB storage | | |
| | FC: Cisco MDS 9132T | | Two Cisco 9132T switches |

| Layer | Product family | Quantity and model | Details |
|---------|--|-----------------------|--|
| Storage | NetApp AFF A220 all-flash storage system | 2 HA Pair | 2-node cluster for all MEDITECH workloads (File Server, Image Server, SQL Server, VMware, and so on) |
| | DS224C disk shelf | 1 x DS224C disk shelf | |
| | SSD | 13 x 3.8TB | |

The following table shows a large (category 6) example configuration – infrastructure components.

| Layer | Product family | Quantity and model | Details |
|-----------------|---|-------------------------|---|
| Compute | Cisco UCS 5108 chassis | 1 | |
| | Cisco chassis I/O modules | 2 x 2208 | 8 x 10GB uplink ports |
| | Cisco UCS blade servers | 8 x B200 M5 | Each with 2 x 24 cores, 2.7GHz and 768GB BIOS 3.2(3#) |
| | Cisco UCS virtual interface card (VIC) | 8 x UCS 1440 VICs | VMware ESXi fNIC FC driver: 1.6.0.47 VMware ESXi eNIC Ethernet driver: 1.0.27.0 (review interoperability matrix: https://ucshcltool.cloudapps.cisco.com/public/) |
| | 2 x Cisco UCS fabric interconnects (FI) | 2 x UCS 6454 FI | 4th-generation fabric interconnects supporting 10GB/25GB/100GB Ethernet and 32GB FC |
| Network | Cisco Ethernet switches | 2 x Nexus 9336c-FX2 | 2 x Cisco Nexus 9332PQ1, 10GB, 25GB, 40GB, 100GB |
| Storage network | IP Network N9k for BLOB storage | | |
| | FC: Cisco MDS 9132T | | Two Cisco 9132T switches |
| Storage | AFF A300 | 1 HA Pair | 2-node cluster for all MEDITECH workloads (File Server, Image Server, SQL Server, VMware, and so on) |
| | DS224C disk shelf | 1 x DS224C disk shelves | |
| | SSD | 19 x 3.8TB | |



These configurations provide a starting point for sizing guidance. Some customer environments might have multiple MEDITECH production and non-MEDITECH workloads running simultaneously, or they might have higher IOP requirements. You should work with the NetApp account team to size the storage systems based on the required IOPS, workloads, and capacity to determine the right platform to serve the workloads.

Additional Information

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

- FlexPod Datacenter with FC Cisco Validated Design.

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html

- NetApp Deployment Guidelines for MEDITECH Environments.

<https://fieldportal.netapp.com/content/248456> (NetApp login required)

- NetApp Sizing Guidelines for MEDITECH Environments.

www.netapp.com/us/media/tr-4190.pdf

- FlexPod Datacenter for Epic EHR Deployment

www.netapp.com/us/media/tr-4693.pdf

- FlexPod Design Zone

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

- FlexPod DC with FC Storage (MDS Switches) Using NetApp AFF, vSphere 6.5U1, and Cisco UCS Manager

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_esxi65u1_n9fc.html

- Cisco Healthcare

<https://www.cisco.com/c/en/us/solutions/industries/healthcare.html?dtid=osscdc000283>

Acknowledgments

The following people contributed to the writing and creation of this guide.

- Brandon Agee, Technical Marketing Engineer, NetApp
- John Duignan, Solutions Architect - Healthcare, NetApp
- Ketan Mota, Product Manager, NetApp
- Jon Ebmeier, Technical Solutions Architect, Cisco Systems, Inc
- Mike Brennan, Product Manager, Cisco Systems, Inc

FlexPod Datacenter for MEDITECH Deployment Guide

TR-4753: FlexPod Datacenter for MEDITECH Deployment Guide

Brandon Agee and John Duignan, NetApp
Mike Brennan and Jon Ebmeier, Cisco

In partnership with:



Overall solution benefits

By running a MEDITECH environment on the FlexPod architectural foundation, your healthcare organization can expect an improvement in staff productivity and a decrease in capital and operational expenditures. FlexPod Datacenter for MEDITECH delivers several benefits that are specific to the healthcare industry, including:

- **Simplified operations and lowered costs.** Eliminate the expense and complexity of legacy platforms by replacing them with a more efficient and scalable shared resource that can support clinicians wherever they are. This solution delivers higher resource utilization for greater return on investment (ROI).
- **Faster deployment of infrastructure.** Whether it's an existing data center or a remote location, with the integrated and tested design of FlexPod Datacenter, you can have your new infrastructure up and running in less time, with less effort.
- **Certified storage.** NetApp ONTAP data management software with MEDITECH gives you the superior reliability of a tested and certified storage vendor. MEDITECH does not certify other infrastructure components.
- **Scale-out architecture.** Scale SAN and NAS from terabytes (TB) to tens of petabytes (PB) without reconfiguring running applications.
- **Nondisruptive operations.** Perform storage maintenance, hardware lifecycle operations, and FlexPod upgrades without interrupting the business.
- **Secure multitenancy.** Support the increased needs of virtualized server and storage shared infrastructure, enabling secure multitenancy of facility-specific information, particularly if your system hosts multiple instances of databases and software.
- **Pooled resource optimization.** Help reduce physical server and storage controller counts, load-balance workload demands, and boost utilization while improving performance.
- **Quality of service (QoS).** FlexPod offers QoS on the entire stack. Industry-leading QoS network, compute, and storage policies enable differentiated service levels in a shared environment. These policies enable optimal performance for workloads and help in isolating and controlling runaway applications.
- **Storage efficiency.** Reduce storage costs with the [NetApp 7:1 storage efficiency guarantee](#).
- **Agility.** With the industry-leading workflow automation, orchestration, and management tools that FlexPod systems provide, your IT team can be far more responsive to business requests. These business requests can range from MEDITECH backup and provisioning of more test and training environments to analytics database replications for population health management initiatives.
- **Increased productivity.** Quickly deploy and scale this solution for optimal clinician end-user experiences.
- **NetApp Data Fabric.** The NetApp Data Fabric architecture weaves data together across sites, beyond

physical boundaries, and across applications. The NetApp Data Fabric is built for data-driven enterprises in a data-centric world. Data is created and is used in multiple locations, and often you need to leverage and to share it with other locations, applications, and infrastructures. You need a way to manage your data that is consistent and integrated. The Data Fabric provides a way to manage data that puts IT in control and that simplifies ever-increasing IT complexity.

FlexPod

New infrastructure approach for MEDITECH EHRs

Healthcare provider organizations like yours remain under pressure to maximize the benefits from substantial investments in industry-leading MEDITECH electronic health records (EHRs). For mission-critical applications, when customers design their data centers for MEDITECH solutions, they often identify the following goals for their data center architecture:

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

As MEDITECH users evolve their organizations to become accountable care organizations and adjust to tightened, bundled reimbursement models, the challenge becomes delivering the required MEDITECH infrastructure in a more efficient and agile IT delivery model.

Value of prevalidated converged infrastructure

Because of an overarching requirement to deliver predictable low-latency system performance and high availability, MEDITECH is prescriptive as to its customers' hardware requirements.

FlexPod is a prevalidated, rigorously tested converged infrastructure from the strategic partnership of Cisco and NetApp. It is engineered and designed specifically to deliver predictable low-latency system performance and high availability. This approach results in MEDITECH compliance and ultimately optimal response time for users of the MEDITECH system.

The FlexPod solution from Cisco and NetApp meets MEDITECH system requirements with a high-performing, modular, prevalidated, converged, virtualized, efficient, scalable, and cost-effective platform. It provides:

- **Modular architecture.** FlexPod meets the varied needs of the MEDITECH modular architecture with purpose-configured FlexPod platforms for each specific workload. All components are connected through a clustered server and a storage management fabric and a cohesive management toolset.
- **Industry-leading technology at each level of the converged stack.** Cisco, NetApp, VMware, and Microsoft Windows are all ranked as number 1 or number 2 by industry analysts in their respective categories of servers, networking, storage, and operating systems.
- **Investment protection with standardized, flexible IT.** The FlexPod reference architecture anticipates new product versions and updates, with rigorous ongoing interoperability testing to accommodate future technologies as they become available.

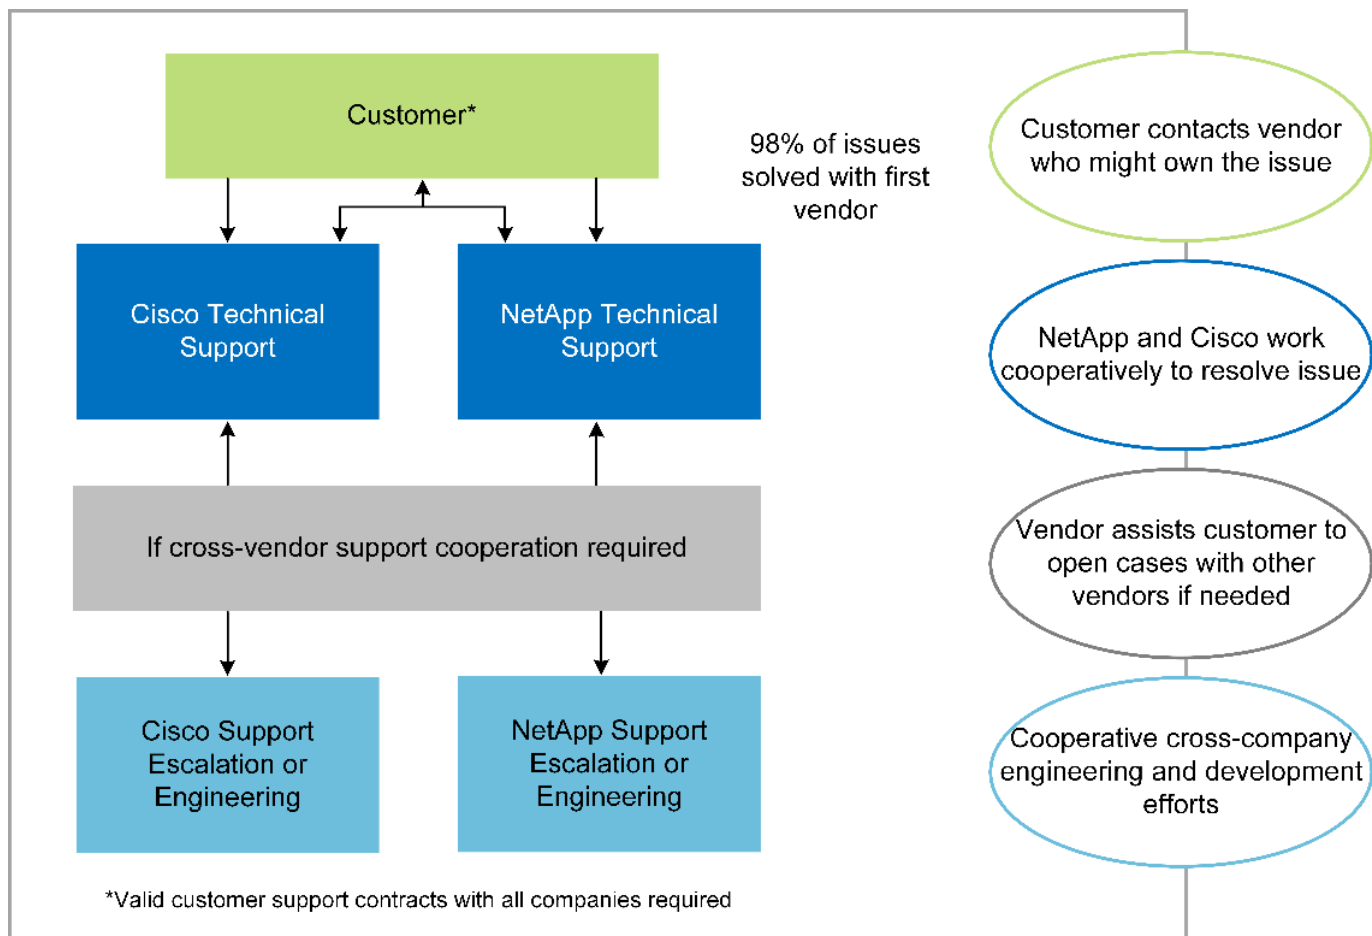
- **Proven deployment across a broad range of environments.** Pretested and jointly validated with popular hypervisors, operating systems, applications, and infrastructure software, FlexPod has been installed in multiple MEDITECH customer organizations.

Proven FlexPod architecture and cooperative support

FlexPod is a proven data center solution, offering a flexible, shared infrastructure that easily scales to support your growing workload demands without negatively affecting performance. By leveraging the FlexPod architecture, this solution delivers the full benefits of FlexPod, including:

- **Performance to meet the MEDITECH workload requirements.** Depending on your MEDITECH Hardware Configuration Proposal requirements, different ONTAP platforms can be deployed to meet your required I/O and latency requirements.
- **Scalability to easily accommodate clinical data growth.** Dynamically scale virtual machines (VMs), servers, and storage capacity on demand, without traditional limits.
- **Enhanced efficiency.** Reduce both administration time and TCO with a converged virtualized infrastructure, which is easier to manage and which stores data more efficiently while driving more performance from MEDITECH software.
- **Reduced risk.** Minimize business disruption with a prevalidated platform that is built on a defined architecture that eliminates deployment guesswork and accommodates ongoing workload optimization.
- **FlexPod Cooperative Support.** NetApp and Cisco have established Cooperative Support, a strong, scalable, and flexible support model to meet the unique support requirements of the FlexPod converged infrastructure. This model uses the combined experience, resources, and technical support expertise of NetApp and Cisco to provide a streamlined process for identifying and resolving your FlexPod support issue, regardless of where the problem resides. With the FlexPod Cooperative Support model, your FlexPod system operates efficiently and benefits from the most up-to-date technology, and you work with an experienced team to help you resolve integration issues.

FlexPod Cooperative Support is especially valuable to healthcare organizations that run business-critical applications such as MEDITECH on the FlexPod converged infrastructure. The following figure illustrates the FlexPod Cooperative Support model.



In addition to these benefits, each component of the FlexPod Datacenter stack with MEDITECH solution delivers specific benefits for MEDITECH EHR workflows.

Cisco Unified Computing System

A self-integrating, self-aware system, Cisco Unified Computing System (Cisco UCS) consists of a single management domain that is interconnected with a unified I/O infrastructure. So that the infrastructure can deliver critical patient information with maximum availability, Cisco UCS for MEDITECH environments has been aligned with MEDITECH infrastructure recommendations and best practices.

The foundation of MEDITECH on Cisco UCS architecture is Cisco UCS technology, with its integrated systems management, Intel Xeon processors, and server virtualization. These integrated technologies solve data center challenges and help you meet your goals for data center design for MEDITECH. Cisco UCS unifies LAN, SAN, and systems management into one simplified link for rack servers, blade servers, and VMs. Cisco UCS is an end-to-end I/O architecture that incorporates Cisco Unified Fabric and Cisco Fabric Extender Technology (FEX Technology) to connect every component in Cisco UCS with a single network fabric and a single network layer.

The system can be deployed as a single or multiple logical units that incorporate and scale across multiple blade chassis, rack servers, racks, and data centers. The system implements a radically simplified architecture that eliminates the multiple redundant devices that populate traditional blade server chassis and rack servers. In traditional systems, redundant devices such as Ethernet and FC adapters and chassis management modules result in layers of complexity. Cisco UCS consists of a redundant pair of Cisco UCS Fabric Interconnects (FIs) that provide a single point of management, and a single point of control, for all I/O traffic.

Cisco UCS uses service profiles to help ensure that virtual servers in the Cisco UCS infrastructure are configured correctly. Service profiles are composed of network, storage, and compute policies that are created once by subject-matter experts in each discipline. Service profiles include critical server information about the

server identity such as LAN and SAN addressing, I/O configurations, firmware versions, boot order, network virtual LAN (VLAN), physical port, and QoS policies. Service profiles can be dynamically created and associated with any physical server in the system in minutes, rather than in hours or days. The association of service profiles with physical servers is performed as a simple, single operation and enables migration of identities between servers in the environment without requiring any physical configuration changes. It facilitates rapid bare-metal provisioning of replacements for retired servers.

The use of service profiles helps ensure that servers are configured consistently throughout the enterprise. When multiple Cisco UCS management domains are employed, Cisco UCS Central can use global service profiles to synchronize configuration and policy information across domains. If maintenance needs to be performed in one domain, the virtual infrastructure can be migrated to another domain. This approach helps to ensure that even when a single domain is offline, applications continue to run with high availability.

To demonstrate that it meets the server configuration requirements, Cisco UCS has been extensively tested with MEDITECH over a multiyear period. Cisco UCS is a supported server platform, as listed on the MEDITECH Product Resources System Support site.

Cisco networking

Cisco Nexus switches and Cisco MDS multilayer directors provide enterprise-class connectivity and SAN consolidation. Cisco multiprotocol storage networking reduces business risk by providing flexibility and options: FC, Fibre Connection (FICON), FC over Ethernet (FCoE), SCSI over IP (iSCSI), and FC over IP (FCIP).

Cisco Nexus switches offer one of the most comprehensive data center network feature sets in a single platform. They deliver high performance and density for both data center and campus cores. They also offer a full feature set for data center aggregation, end-of-row, and data center interconnect deployments in a highly resilient modular platform.

Cisco UCS integrates computing resources with Cisco Nexus switches and a unified I/O fabric that identifies and handles different types of network traffic. This traffic includes storage I/O, streamed desktop traffic, management, and access to clinical and business applications. You get:

- **Infrastructure scalability.** Virtualization, efficient power and cooling, cloud scale with automation, high density, and high performance all support efficient data center growth.
- **Operational continuity.** The design integrates hardware, NX-OS software features, and management to support zero-downtime environments.
- **Network and computer QoS.** Cisco delivers policy-driven class of service (CoS) and QoS across the networking, storage, and compute fabric for optimal performance of mission-critical applications.
- **Transport flexibility.** Incrementally adopt new networking technologies with a cost-effective solution.

Together, Cisco UCS with Cisco Nexus switches and Cisco MDS multilayer directors provides an optimal compute, networking, and SAN connectivity solution for MEDITECH.

NetApp ONTAP

NetApp storage that runs ONTAP software reduces your overall storage costs while it delivers the low-latency read and write response times and IOPS that MEDITECH workloads need. ONTAP supports both all-flash and hybrid storage configurations to create an optimal storage platform that meets MEDITECH requirements. NetApp flash-accelerated systems have received MEDITECH's validation and certification, giving you as a MEDITECH customer the performance and responsiveness that are key to latency-sensitive MEDITECH operations. By creating multiple fault domains in a single cluster, NetApp systems can also isolate production from nonproduction. NetApp systems also reduce performance issues with a guaranteed performance level minimum for workloads with ONTAP QoS.

The scale-out architecture of the ONTAP software can flexibly adapt to various I/O workloads. To deliver the necessary throughput and low latency that clinical applications need while also providing a modular scale-out architecture, all-flash configurations are typically used in ONTAP architectures. NetApp AFF nodes can be combined in the same scale-out cluster with hybrid (HDD and flash) storage nodes that are suitable for storing large datasets with high throughput. Along with a MEDITECH-approved backup solution, you can clone, replicate, and back up your MEDITECH environment from expensive solid-state drive (SSD) storage to more economical HDD storage on other nodes. This approach meets or exceeds MEDITECH guidelines for SAN-based cloning and backup of production pools.


Many of the ONTAP features are especially useful in MEDITECH environments: simplifying management, increasing availability and automation, and reducing the total amount of storage needed. With these features, you get:

- **Outstanding performance.** The NetApp AFF solution shares the Unified Storage Architecture, ONTAP software, management interface, rich data services, and advanced feature set that the rest of the NetApp FAS product families have. This innovative combination of all-flash media with ONTAP delivers the consistent low latency and high IOPS of all-flash storage with the industry-leading quality of ONTAP software.
- **Storage efficiency.** Reduce total capacity requirements with deduplication, NetApp FlexClone data replication technology, inline compression, inline compaction, thin replication, thin provisioning, and aggregate deduplication.

NetApp deduplication provides block-level deduplication in a NetApp FlexVol volume or data constituent. Essentially, deduplication removes duplicate blocks, storing only unique blocks in the FlexVol volume or data constituent.


Deduplication works with a high degree of granularity and operates on the active file system of the FlexVol volume or data constituent. It is application transparent; therefore, you can use it to deduplicate data that originates from any application that uses the NetApp system. You can run volume deduplication as an inline process (starting in ONTAP 8.3.2). You can also run it as a background process that you can configure to run automatically, to be scheduled, or to run manually through the CLI, NetApp ONTAP System Manager, or NetApp Active IQ Unified Manager.

The following figure illustrates how NetApp deduplication works at the highest level.



- **Space-efficient cloning.** The FlexClone capability enables you to almost instantly create clones to support backup and testing environment refresh. These clones consume more storage only as changes are made.
- **NetApp Snapshot and SnapMirror technologies.** ONTAP can create space-efficient Snapshot copies of the logical unit numbers (LUNs) that the MEDITECH host uses. For dual-site deployments, you can implement SnapMirror software for more data replication and resiliency.
- **Integrated data protection.** Full data protection and disaster recovery features help you protect critical data assets and provide disaster recovery.
- **Nondisruptive operations.** You can perform upgrades and maintenance without taking data offline.
- **QoS and adaptive QoS (AQoS).** Storage QoS enables you to limit potential bully workloads. More important, QoS can guarantee a performance minimum for critical workloads such as MEDITECH production. By limiting contention, NetApp QoS can reduce performance-related issues. AQoS works with predefined policy groups, which you can apply directly to a volume. These policy groups can automatically scale a throughput ceiling or floor-to-volume size, maintaining the ratio of IOPS to terabytes and gigabytes as the size of the volume changes.
- **NetApp Data Fabric.** The NetApp Data Fabric simplifies and integrates data management across cloud and on-premises environments to accelerate digital transformation. It delivers consistent and integrated data management services and applications for data visibility and insights, data access and control, and data protection and security. NetApp is integrated with Amazon Web Services (AWS), Azure, Google Cloud Platform, and IBM Cloud clouds, giving you a wide breadth of choice.

The following figure illustrates the FlexPod architecture for MEDITECH workloads.



MEDITECH overview

Medical Information Technology, Inc., commonly known as MEDITECH, is a Massachusetts-based software company that provides information systems for healthcare organizations. MEDITECH provides an EHR system that is designed to store and to organize the latest patient data and provides the data to clinical staff. Patient data includes, but is not limited to, demographics; medical history; medication; laboratory test results; radiology images; and personal information such as age, height, and weight.

It is beyond the scope of this document to cover the wide span of functions that MEDITECH software supports. Appendix A provides more information about these broad sets of MEDITECH functions. MEDITECH applications require several VMs to support these functions. To deploy these applications, see the recommendations from MEDITECH.

For each deployment, from the storage system point of view, all MEDITECH software systems require a distributed patient-centric database. MEDITECH has its own proprietary database, which uses the Windows operating system.

BridgeHead and Commvault are the two backup software applications that are certified by both NetApp and MEDITECH. The scope of this document does not cover the deployment of these backup applications.

The primary focus of this document is to enable the FlexPod stack (servers and storage) to meet the performance-driven requirements for the MEDITECH database and the backup requirements in the EHR environment.

Purpose-built for specific MEDITECH workloads

MEDITECH does not resell server, network, or storage hardware, hypervisors, or operating systems; however, it has specific requirements for each component of the infrastructure stack. Therefore, Cisco and NetApp

worked together to test and to enable FlexPod Datacenter to be successfully configured, deployed, and supported to meet the MEDITECH production environment requirements of customers like you.

MEDITECH categories

MEDITECH associates the deployment size with a category number that ranges from 1 to 6. Category 1 represents the smallest MEDITECH deployments, and category 6 represents the largest MEDITECH deployments.

For information about the I/O characteristics and performance requirements for a MEDITECH host in each category, see NetApp [TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).

MEDITECH platform

The MEDITECH Expanse platform is the latest version of the company's EHR software. Earlier MEDITECH platforms are Client/Server 5.x and MAGIC. This section describes the MEDITECH platform (applicable to Expanse, 6.x, C/S 5.x, and MAGIC), pertaining to the MEDITECH host and its storage requirements.

For all the preceding MEDITECH platforms, multiple servers run MEDITECH software, performing various tasks. The previous figure depicts a typical MEDITECH system, including MEDITECH hosts serving as application database servers and other MEDITECH servers. Examples of other MEDITECH servers include the Data Repository application, the Scanning and Archiving application, and Background Job Clients. For the complete list of other MEDITECH servers, see the "Hardware Configuration Proposal" (for new deployments) and "Hardware Evaluation Task" (for existing deployments) documents. You can obtain these documents from MEDITECH through the MEDITECH system integrator or from your MEDITECH Technical Account Manager (TAM).

MEDITECH host

A MEDITECH host is a database server. This host is also referred to as a MEDITECH file server (for the Expanse, 6.x, or C/S 5.x platform) or as a MAGIC machine (for the MAGIC platform). This document uses the term MEDITECH host to refer to a MEDITECH file server or a MAGIC machine.

MEDITECH hosts can be physical servers or VMs that run on the Microsoft Windows Server operating system. Most commonly in the field, MEDITECH hosts are deployed as Windows VMs that run on a VMware ESXi server. As of this writing, VMware is the only hypervisor that MEDITECH supports. A MEDITECH host stores its program, dictionary, and data files on a Microsoft Windows drive (for example, drive E) on the Windows system.

In a virtual environment, a Windows E drive resides on a LUN that is attached to the VM by way of a raw device mapping (RDM) in physical compatibility mode. The use of Virtual Machine Disk (VMDK) files as a Windows E drive in this scenario is not supported by MEDITECH.

MEDITECH host workload I/O characteristic

The I/O characteristic of each MEDITECH host and the system as a whole depends on the MEDITECH platform that you deploy. All MEDITECH platforms (Expanse, 6.x, C/S 5.x, and MAGIC) generate workloads that are 100% random.

The MEDITECH Expanse platform generates the most demanding workload because it has the highest percentage of write operations and overall IOPS per host, followed by 6.x, C/S 5.x, and the MAGIC platforms.

For more details about the MEDITECH workload descriptions, see [TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).

Storage network

MEDITECH requires that the FC Protocol be used for data traffic between the NetApp FAS or AFF system and the MEDITECH hosts of all categories.

Storage presentation for a MEDITECH host

Each MEDITECH host uses two Windows drives:

- **Drive C.** This drive stores the Windows Server operating system and the MEDITECH host application files.
- **Drive E.** The MEDITECH host stores its program, dictionary, and data files on drive E of the Windows Server operating system. Drive E is a LUN that is mapped from the NetApp FAS or AFF system by using the FC Protocol. MEDITECH requires that the FC Protocol be used so that the MEDITECH host's IOPS and read and write latency requirements are met.

Volume and LUN naming convention

MEDITECH requires that a specific naming convention be used for all LUNs.

Before any storage deployment, verify the MEDITECH Hardware Configuration Proposal to confirm the naming convention for the LUNs. The MEDITECH backup process relies on the volume and LUN naming convention to properly identify the specific LUNs to back up.

Comprehensive management tools and automation capabilities

Cisco UCS with Cisco UCS Manager

Cisco focuses on three key elements to deliver a superior data center infrastructure: simplification, security, and scalability. The Cisco UCS Manager software combined with platform modularity provides a simplified, secure, and scalable desktop virtualization platform:

- **Simplified.** Cisco UCS provides a radical new approach to industry-standard computing and provides the core of the data center infrastructure for all workloads. Cisco UCS offers many features and benefits, including reduction in the number of servers that you need and reduction in the number of cables that are used per server. Another important feature is the capability to rapidly deploy or to reprovision servers through Cisco UCS service profiles. With fewer servers and cables to manage and with streamlined server and application workload provisioning, operations are simplified. Scores of blade and rack servers can be provisioned in minutes with Cisco UCS Manager service profiles. Cisco UCS service profiles eliminate server integration runbooks and eliminate configuration drift. This approach accelerates the time to productivity for end users, improves business agility, and allows IT resources to be allocated to other tasks.

Cisco UCS Manager automates many mundane, error-prone data center operations such as configuration and provisioning of server, network, and storage access infrastructure. In addition, Cisco UCS B-Series Blade Servers and C-Series Rack Servers with large memory footprints enable high application user density, which helps reduce server infrastructure requirements.

Simplification leads to a faster, more successful MEDITECH infrastructure deployment.

- **Secure.** Although VMs are inherently more secure than their physical predecessors, they introduce new security challenges. Mission-critical web and application servers that use a common infrastructure such as virtual desktops are now at a higher risk for security threats. Inter- VM traffic now poses an important security consideration that your IT managers must address, especially in dynamic environments in which VMs, using VMware vMotion, move across the server infrastructure.

Virtualization, therefore, significantly increases the need for VM- level awareness of policy and security,

especially given the dynamic and fluid nature of VM mobility across an extended computing infrastructure. The ease with which new virtual desktops can proliferate magnifies the importance of a virtualization-aware network and security infrastructure. Cisco data center infrastructure (Cisco UCS, Cisco MDS, and Cisco Nexus family solutions) for desktop virtualization provides strong data center, network, and desktop security, with comprehensive security from the desktop to the hypervisor. Security is enhanced with segmentation of virtual desktops, VM-aware policies and administration, and network security across the LAN and WAN infrastructure.

- **Scalable.** Growth of virtualization solutions is all but inevitable, so a solution must be able to scale, and to scale predictably, with that growth. The Cisco virtualization solutions support high VM density (VMs per server), and more servers scale with near-linear performance. Cisco data center infrastructure provides a flexible platform for growth and improves business agility. Cisco UCS Manager service profiles allow on-demand host provisioning and make it as easy to deploy hundreds of hosts as it is to deploy dozens.

Cisco UCS Servers provide near-linear performance and scale. Cisco UCS implements the patented Cisco Extended Memory Technology to offer large memory footprints with fewer sockets (with scalability of up to 1TB of memory with 2- and 4-socket servers). By using Unified Fabric technology as a building block, Cisco UCS Server aggregate bandwidth can scale up to 80Gbps per server, and the northbound Cisco UCS Fabric Interconnect can output 2Tbps at line rate. This capability helps prevent desktop virtualization I/O and memory bottlenecks. Cisco UCS, with its high-performance, low-latency Unified Fabric-based networking architecture, supports high volumes of virtual desktop traffic, including high-resolution video and communications traffic. In addition, ONTAP helps to maintain data availability and optimal performance during boot and login storms as part of the FlexPod virtualization solutions.

Cisco UCS, Cisco MDS, and Cisco Nexus data center infrastructure designs provide an excellent platform for growth. You get transparent scaling of server, network, and storage resources to support desktop virtualization, data center applications, and cloud computing.

VMware vCenter Server

VMware vCenter Server provides a centralized platform for managing MEDITECH environments so that your healthcare organization can automate and deliver a virtual infrastructure with confidence:

- **Simple deployment.** Quickly and easily deploy vCenter Server by using a virtual appliance.
- **Centralized control and visibility.** Administer the entire VMware vSphere infrastructure from a single location.
- **Proactive optimization.** Allocate and optimize resources for maximum efficiency.
- **Management.** Use powerful plug-ins and tools to simplify management and to extend control.

Virtual Storage Console for VMware vSphere

Virtual Storage Console (VSC), vSphere API for Storage Awareness (VASA) Provider, and VMware Storage Replication Adapter (SRA) for VMware vSphere from NetApp make up a single virtual appliance. The product suite includes SRA and VASA Provider as plug-ins to vCenter Server, which provides end-to-end lifecycle management for VMs in VMware environments that use NetApp storage systems.

The virtual appliance for VSC, VASA Provider, and SRA integrates smoothly with the VMware vSphere Web Client and enables you to use SSO services. In an environment with multiple VMware vCenter Server instances, each vCenter Server instance that you want to manage must have its own registered instance of VSC. The VSC dashboard page enables you to quickly check the overall status of your datastores and VMs.

By deploying the virtual appliance for VSC, VASA Provider, and SRA, you can perform the following tasks:

- **Use VSC to deploy and manage storage and to configure the ESXi host.** You can use VSC to add credentials, to remove credentials, to assign credentials, and to set up permissions for storage controllers in your VMware environment. In addition, you can manage ESXi servers that are connected to NetApp storage systems. With a couple clicks, you can set recommended best practice values for host timeouts, NAS, and multipathing for all the hosts. You can also view storage details and collect diagnostic information.
- **Use VASA Provider to create storage capability profiles and to set alarms.** VASA Provider for ONTAP is registered with VSC when you enable the VASA Provider extension. You can create and use storage capability profiles and virtual datastores. You can also set alarms to alert you when the thresholds for volumes and aggregates are almost full. You can monitor the performance of VMDKs and the VMs that are created on virtual datastores.
- **Use SRA for disaster recovery.** You can use SRA to configure protected and recovery sites in your environment for disaster recovery during failures.

NetApp OnCommand Insight and ONTAP

NetApp OnCommand Insight integrates infrastructure management into the MEDITECH service delivery chain. This approach gives your healthcare organization better control, automation, and analysis of your storage, network, and compute infrastructure. IT can optimize your current infrastructure for maximum benefit while simplifying the process of determining what and when to buy. It also mitigates the risks that are associated with complex technology migrations. Because it requires no agents, installation is straightforward and nondisruptive. Installed storage and SAN devices are continually discovered, and detailed information is collected for full visibility of your entire storage environment. You can quickly identify misused, misaligned, underused, or orphaned assets and reclaim them to fuel future expansion. OnCommand Insight helps you:

- **Optimize existing resources.** Identify misused, underused, or orphaned assets by using established best practices to avoid problems and to meet service levels.
- **Make better decisions.** Real-time data helps resolve capacity problems more quickly to accurately plan future purchases, to avoid overspending, and to defer capital expenditures.
- **Accelerate IT initiatives.** Better understand your virtual environments to help you manage risks, minimize downtime, and speed cloud deployment.

Design

The architecture of FlexPod for MEDITECH is based on guidance from MEDITECH, Cisco, and NetApp and from partner experience in working with MEDITECH customers of all sizes. The architecture is adaptable and applies best practices for MEDITECH, depending on your data center strategy; the size of your organization; and whether your system is centralized, distributed, or multitenant.

The correct storage architecture can be determined by the overall size with the total IOPS. Performance alone is not the only factor, and you might decide to use a larger node count based on additional customer requirements. The advantage of using NetApp storage is that you can easily and nondisruptively scale up the cluster as your requirements change. You can also nondisruptively remove nodes from the cluster to repurpose equipment or during equipment refreshes.

Here are some of the benefits of the NetApp ONTAP storage architecture:

- **Easy, nondisruptive scale-up and scale-out.** You can upgrade, add, or remove disks and nodes by using ONTAP nondisruptive operations. You can start with four nodes and move to six nodes or upgrade to larger controllers nondisruptively.

- **Storage efficiencies.** Reduce your total capacity requirements with deduplication, NetApp FlexClone, inline compression, inline compaction, thin replication, thin provisioning, and aggregate deduplication. The FlexClone capability enables you to almost instantly create clones to support backup and testing environment refreshes. These clones consume more storage only as changes are made.
- **Disaster recovery shadow database server.** The disaster recovery shadow database server is part of your business continuity strategy (used to support storage read-only functionality and potentially configured to be a storage read/write instance). Therefore, the placement and sizing of the third storage system are usually the same as in your production database storage system.
- **Database consistency (requires some consideration).** If you use NetApp SnapMirror backup copies in relation to business continuity, see [TR-3446: SnapMirror Async Overview and Best Practices Guide](#).

Storage layout

Dedicated aggregates for MEDITECH hosts

The first step toward meeting MEDITECH's high-performance and high-availability requirements is to properly design the storage layout for the MEDITECH environment to isolate the MEDITECH host production workload onto dedicated, high-performance storage.

One dedicated aggregate should be provisioned on each storage controller for storing the program, dictionary, and data files of the MEDITECH hosts. To eliminate the possibility of other workloads using the same disks and affecting performance, no other storage is provisioned from these aggregates.



Storage that you provision for the other MEDITECH servers should not be placed on the dedicated aggregate for the LUNs that are used by the MEDITECH hosts. You should place the storage for other MEDITECH servers on a separate aggregate. Storage requirements for other MEDITECH servers are available in the “Hardware Configuration Proposal” (for new deployments) and “Hardware Evaluation Task” (for existing deployments) documents. You can obtain these documents from MEDITECH through the MEDITECH system integrator or from your MEDITECH Technical Account Manager (TAM). NetApp solutions engineers might consult with the NetApp MEDITECH Independent Software Vendor (ISV) team to facilitate a proper and complete NetApp storage sizing configuration.

Spread MEDITECH host workload evenly across all storage controllers

NetApp FAS and AFF systems are deployed as one or more high-availability pairs. NetApp recommends that you spread the MEDITECH Expanse and 6.x workloads evenly across each storage controller to apply the compute, network, and caching resources on each storage controller.

Use the following guidelines to spread the MEDITECH workloads evenly across each storage controller:

- If you know the IOPS for each MEDITECH host, you can spread the MEDITECH Expanse and 6.x workloads evenly across all storage controllers by confirming that each controller services a similar number of IOPS from the MEDITECH hosts.
- If you do not know the IOPS for each MEDITECH host, you can still spread the MEDITECH Expanse and 6.x workloads evenly across all storage controllers. Complete this task by confirming that the capacity of the aggregates for the MEDITECH hosts is evenly distributed across all storage controllers. By doing so, the number of disks is the same across all data aggregates that are dedicated to the MEDITECH hosts.
- Use similar disk types and identical RAID groups to create the storage aggregates of both controllers for distributing the workloads equally. Before you create the storage aggregate, contact a NetApp Certified Integrator.



According to MEDITECH, two hosts in the MEDITECH system generate higher IOPS than the rest of the hosts. The LUNs for these two hosts should be placed on separate storage controllers. You should identify these two hosts with the assistance of the MEDITECH team before you deploy your system.

Storage Placement

Database storage for MEDITECH hosts

The database storage for a MEDITECH host is presented as a block device (that is, a LUN) from the NetApp FAS or AFF system. The LUN is typically mounted to the Windows operating system as the E drive.

Other storage

The MEDITECH host operating system and the database application normally generate a considerable amount of IOPS on the storage. Storage provisioning for the MEDITECH host VMs and their VMDK files, if necessary, is considered independent from the storage that is required to meet the MEDITECH performance thresholds.

Storage that is provisioned for the other MEDITECH servers should not be placed on the dedicated aggregate for the LUNs that the MEDITECH hosts use. Place the storage for other MEDITECH servers on a separate aggregate.

Storage controller configuration

High availability

To mitigate the effect of controller failure and to enable nondisruptive upgrades of the storage system, you should configure your storage system with controllers in a high-availability pair in the high-availability mode.

With the high-availability controller pair configuration, disk shelves should be connected to controllers by multiple paths. This connection increases storage resiliency by protecting against a single-path failure, and it improves performance consistency if a controller failover occurs.

Storage performance during storage controller failover

For storage systems that are configured with controllers in a high-availability pair, in the unlikely event of a controller failure, the partner controller takes over the failed controller's storage resources and workloads. It is important to consult the customer to determine the performance requirements that must be met if there is a controller failure and to size the system accordingly.

Hardware-assisted takeover

NetApp recommends that you turn on the hardware-assisted takeover feature on both storage controllers.

Hardware-assisted takeover is designed to minimize the storage controller failover time. It enables one controller's Remote LAN Module or Service Processor module to notify its partner about a controller failure faster than a heartbeat timeout trigger can, reducing the time that it takes to failover. The hardware-assisted takeover feature is enabled by default for storage controllers in a high-availability configuration.

For more information about hardware-assisted takeover, see the [ONTAP 9 Documentation Center](#).

Disk type

To support the low read latency requirement of MEDITECH workloads, NetApp recommends that you use a

high-performance SSD for aggregates on AFF systems that are dedicated for the MEDITECH hosts.

NetApp AFF

NetApp offers high-performance AFF arrays to address MEDITECH workloads that demand high throughput and that have random data access patterns and low- latency requirements. For MEDITECH workloads, AFF arrays offer performance advantages over systems that are based on HDDs. The combination of flash technology and enterprise data management delivers advantages in three major areas: performance, availability, and storage efficiency.

NetApp Support tools and services

NetApp offers a complete set of support tools and services. The NetApp AutoSupport tool should be enabled and configured on NetApp AFF/FAS systems to call home if there is a hardware failure or system misconfiguration. Calling home alerts the NetApp Support team to remediate any issues in a timely manner. NetApp Active IQ is a web based application that is based on AutoSupport information from your NetApp systems providing predictive and proactive insight to help improve availability, efficiency, and performance.

Deployment and configuration

Overview

The NetApp storage guidance for FlexPod deployment that is provided in this document covers:

- Environments that use ONTAP
- Environments that use Cisco UCS blade and rack-mount servers

This document does not cover:

- Detailed deployment of the FlexPod Datacenter environment

For more information, see [FlexPod Datacenter with FC Cisco Validated Design](#) (CVD).

- An overview of MEDITECH software environments, reference architectures, and integration best practices guidance.

For more information, see [TR-4300i: NetApp FAS and All-Flash Storage Systems for MEDITECH Environments Best Practices Guide](#) (NetApp login required).

- Quantitative performance requirements and sizing guidance.


For more information, see [TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).

- Use of NetApp SnapMirror technologies to meet backup and disaster recovery requirements.
- Generic NetApp storage deployment guidance.

This section provides an example configuration with infrastructure deployment best practices and lists the various infrastructure hardware and software components and the versions that you can use.

Cabling diagram

The following figure illustrates the 32Gb FC/40GbE topology diagram for a MEDITECH deployment.



Always use the [Interoperability Matrix Tool \(IMT\)](#) to validate that all versions of software and firmware are supported. The table in section "[MEDITECH modules and components](#)" lists the infrastructure hardware and software components that were used in the solution testing.

Next: [Base infrastructure Configuration.](#)

Base infrastructure configuration

Network connectivity

The following network connections must be in place before you configure the infrastructure:

- Link aggregation that uses port channels and virtual port channels (vPCs) is used throughout, enabling the design for higher bandwidth and high availability:
 - vPC is used between the Cisco FI and Cisco Nexus switches.
 - Each server has virtual network interface cards (vNICs) with redundant connectivity to the Unified Fabric. NIC failover is used between FIs for redundancy.
 - Each server has virtual host bus adapters (vHBAs) with redundant connectivity to the Unified Fabric.
- The Cisco UCS FI is configured in end- host mode as recommended, providing dynamic pinning of vNICs to uplink switches.

Storage connectivity

The following storage connections must be in place before you configure the infrastructure:

- Storage port interface groups (ifgroups, vPC)
- 10Gb link to switch N9K-A
- 10Gb link to switch N9K-B
- In- band management (active-passive bond):
 - 1Gb link to management switch N9K-A

- 1Gb link to management switch N9K-B
- 32Gb FC end-to-end connectivity through Cisco MDS switches; single initiator zoning configured
- FC SAN boot to fully achieve stateless computing; servers are booted from LUNs in the boot volume that is hosted on the AFF storage cluster
- All MEDITECH workloads are hosted on FC LUNs, which are spread across the storage controller nodes

Host software

The following software must be installed:

- ESXi installed on the Cisco UCS blades
- VMware vCenter installed and configured (with all the hosts registered in vCenter)
- VSC installed and registered in VMware vCenter
- NetApp cluster configured

[Next: Cisco UCS Blade Server and Switch Configuration.](#)

Cisco UCS blade server and switch configuration

The FlexPod for MEDITECH software is designed with fault tolerance at every level. There is no single point of failure in the system. For optimal performance, Cisco recommends the use of hot spare blade servers.

This document provides high-level guidance on the basic configuration of a FlexPod environment for MEDITECH software. In this section, we present high-level steps with some examples to prepare the Cisco UCS compute platform element of the FlexPod configuration. A prerequisite for this guidance is that the FlexPod configuration is racked, powered, and cabled per the instructions in the [FlexPod Datacenter with Fibre Channel Storage using VMware vSphere 6.5 Update 1, NetApp AFF A-series and Cisco UCS Manager 3.2 CVD](#).

Cisco Nexus switch configuration

A fault- tolerant pair of Cisco Nexus 9300 Series Ethernet switches is deployed for the solution. You should cable these switches as described in the [Cabling Diagram](#) section. The Cisco Nexus configuration helps ensure that Ethernet traffic flows are optimized for the MEDITECH application.

1. After you have completed the initial setup and licensing, run the following commands to set global configuration parameters on both switches:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
port-channel load-balance src-dst l4port
ntp server <global-ntp-server-ip> use-vrf management
ntp master 3
ip route 0.0.0.0/0 <ib-mgmt-vlan-gateway>
copy run start
```

2. Create the VLANs for the solution on each switch using the global configuration mode:

```
vlan <ib-mgmt-vlan-id>
name IB-MGMT-VLAN
vlan <native-vlan-id>
name Native-VLAN
vlan <vmotion-vlan-id>
name vMotion-VLAN
vlan <vm-traffic-vlan-id>
name VM-Traffic-VLAN
vlan <infra-nfs-vlan-id>
name Infra-NFS-VLAN
exit
copy run start
```

3. Create the Network Time Protocol (NTP) distribution interface, port channels, port channel parameters, and port descriptions for troubleshooting per [FlexPod Datacenter with Fibre Channel Storage using VMware vSphere 6.5 Update 1, NetApp AFF A-series and Cisco UCS Manager 3.2 CVD](#).

Cisco MDS 9132T configuration

The Cisco MDS 9100 Series FC switches provide redundant 32Gb FC connectivity between the NetApp AFF A200 or AFF A300 controllers and the Cisco UCS compute fabric. You should connect the cables as described in the [Cabling Diagram](#) section.

1. From the consoles on each MDS switch, run the following commands to enable the required features for the solution:

```
configure terminal
feature npiv
feature fport-channel-trunk
```

2. Configure individual ports, port channels, and descriptions as per the FlexPod Cisco MDS switch configuration section in [FlexPod Datacenter with FC Cisco Validated Design](#).
3. To create the necessary virtual SANs (VSANs) for the solution, complete the following steps while in global configuration mode:
 - a. For the Fabric-A MDS switch, run the following commands:

```
vsan database
vsan <vsan-a-id>
vsan <vsan-a-id> name Fabric-A
exit
zone smart-zoning enable vsan <vsan-a-id>
vsan database
vsan <vsan-a-id> interface fc1/1
vsan <vsan-a-id> interface fc1/2
vsan <vsan-a-id> interface port-channel110
vsan <vsan-a-id> interface port-channel112
```

The port channel numbers in the last two lines of the command were created when the individual ports, port channels, and descriptions were provisioned by using the reference document.

- b. For the Fabric-B MDS switch, run the following commands:

```
vsan database
vsan <vsan-b-id>
vsan <vsan-b-id> name Fabric-B
exit
zone smart-zoning enable vsan <vsan-b-id>
vsan database
vsan <vsan-b-id> interface fc1/1
vsan <vsan-b-id> interface fc1/2
vsan <vsan-b-id> interface port-channel111
vsan <vsan-b-id> interface port-channel113
```

The port channel numbers in the last two lines of the command were created when the individual ports, port channels, and descriptions were provisioned by using the reference document.

4. For each FC switch, create device alias names that make the identification of each device intuitive for ongoing operations by using the details in the reference document.
5. Finally, create the FC zones by using the device alias names that were created in step 4 for each MDS switch as follows:
 - a. For the Fabric-A MDS switch, run the following commands:

```

configure terminal
zone name VM-Host-Infra-01-A vsan <vsan-a-id>
member device-alias VM-Host-Infra-01-A init
member device-alias Infra-SVM-fcp_lif01a target
member device-alias Infra-SVM-fcp_lif02a target
exit
zone name VM-Host-Infra-02-A vsan <vsan-a-id>
member device-alias VM-Host-Infra-02-A init
member device-alias Infra-SVM-fcp_lif01a target
member device-alias Infra-SVM-fcp_lif02a target
exit
zoneset name Fabric-A vsan <vsan-a-id>
member VM-Host-Infra-01-A
member VM-Host-Infra-02-A
exit
zoneset activate name Fabric-A vsan <vsan-a-id>
exit
show zoneset active vsan <vsan-a-id>

```

b. For the Fabric-B MDS switch, run the following commands:

```

configure terminal
zone name VM-Host-Infra-01-B vsan <vsan-b-id>
member device-alias VM-Host-Infra-01-B init
member device-alias Infra-SVM-fcp_lif01b target
member device-alias Infra-SVM-fcp_lif02b target
exit
zone name VM-Host-Infra-02-B vsan <vsan-b-id>
member device-alias VM-Host-Infra-02-B init
member device-alias Infra-SVM-fcp_lif01b target
member device-alias Infra-SVM-fcp_lif02b target
exit
zoneset name Fabric-B vsan <vsan-b-id>
member VM-Host-Infra-01-B
member VM-Host-Infra-02-B
exit
zoneset activate name Fabric-B vsan <vsan-b-id>
exit
show zoneset active vsan <vsan-b-id>

```

Cisco UCS configuration guidance

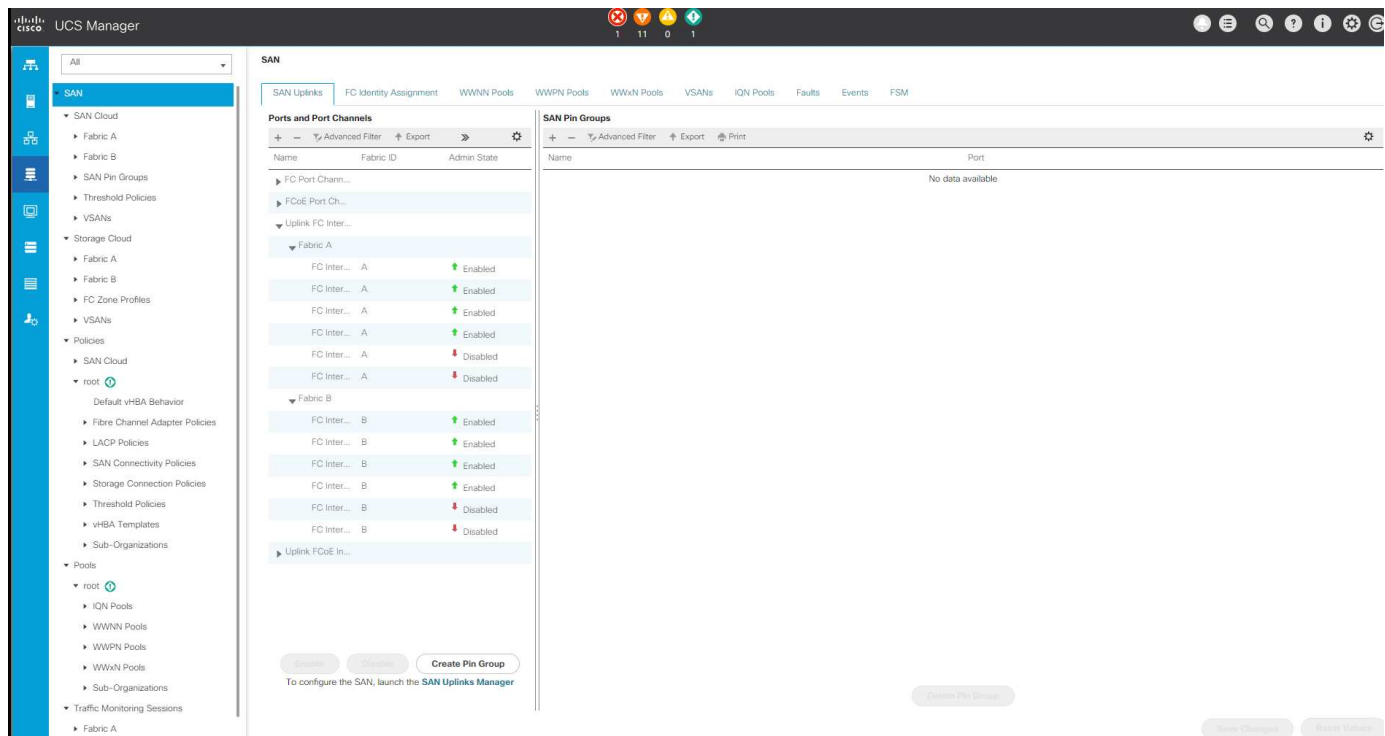
Cisco UCS enables you as a MEDITECH customer to leverage your subject-matter experts in network, storage, and compute to create policies and templates that tailor the environment to your specific needs. After

they are created, these policies and templates can be combined into service profiles that deliver consistent, repeatable, reliable, and fast deployments of Cisco blade and rack servers.

Cisco UCS provides three methods for managing a Cisco UCS system, called a domain:

- Cisco UCS Manager HTML5 GUI
- Cisco UCS CLI
- Cisco UCS Central for multidomain environments

The following figure shows a sample screenshot of the SAN node in Cisco UCS Manager.



In larger deployments, independent Cisco UCS domains can be built for more fault tolerance at the major MEDITECH functional component level.

In highly fault- tolerant designs with two or more data centers, Cisco UCS Central plays a key role in setting global policy and global service profiles for consistency between hosts throughout the enterprise.

To set up the Cisco UCS compute platform, complete the following procedures. Perform these procedures after the Cisco UCS B200 M5 Blade Servers are installed in the Cisco UCS 5108 AC blade chassis. Also, you must complete the cabling requirements as described in the [Cabling Diagram](#) section.

1. Upgrade the Cisco UCS Manager firmware to version 3.2(2f) or later.
2. Configure the reporting, Cisco call home features, and NTP settings for the domain.
3. Configure the server and uplink ports on each Fabric Interconnect.
4. Edit the chassis discovery policy.
5. Create the address pools for out- of- band management, universal unique identifiers (UUIDs), MAC address, servers, worldwide node name (WWNN), and worldwide port name (WWPN).
6. Create the Ethernet and FC uplink port channels and VSANs.
7. Create policies for SAN connectivity, network control, server pool qualification, power control, server BIOS,

and default maintenance.

8. Create vNIC and vHBA templates.
9. Create vMedia and FC boot policies.
10. Create service profile templates and service profiles for each MEDITECH platform element.
11. Associate the service profiles with the appropriate blade servers.

For the detailed steps to configure each key element of the Cisco UCS service profiles for FlexPod, see the [FlexPod Datacenter with Fibre Channel Storage using VMware vSphere 6.5 Update 1, NetApp AFF A-series and Cisco UCS Manager 3.2 CVD](#) document.

[Next: ESXi Configuration Best Practices.](#)

ESXi configuration best practices

For the ESXi host-side configuration, configure the VMware hosts as you would run any enterprise database workload:


- VSC for VMware vSphere checks and sets the ESXi host multipathing settings and HBA timeout settings that work best with NetApp storage systems. The values that VSC sets are based on rigorous internal testing by NetApp.
- For optimal storage performance, consider using storage hardware that supports VMware vStorage APIs - Array Integration (VAAI). The NetApp Plug-In for VAAI is a software library that integrates the VMware Virtual Disk Libraries that are installed on the ESXi host. The VMware VAAI package enables the offloading of certain tasks from the physical hosts to the storage array.

You can perform tasks such as thin provisioning and hardware acceleration at the array level to reduce the workload on the ESXi hosts. The copy offload feature and space reservation feature improve the performance of VSC operations. You can download the plug-in installation package and obtain the instructions for installing the plug-in from the NetApp Support site.

VSC sets ESXi host timeouts, multipath settings, and HBA timeout settings and other values for optimal performance and successful failover of the NetApp storage controllers. Follow these steps:

1. From the VMware vSphere Web Client home page, select vCenter > Hosts.
2. Right-click a host and then select Actions > NetApp VSC > Set Recommended Values.
3. In the NetApp Recommended Settings dialog box, select the values that work best with your system.

The standard recommended values are set by default.



4. Click OK.

[Next: NetApp Configuration.](#)

NetApp configuration

NetApp storage that is deployed for MEDITECH software environments uses storage controllers in a high-availability-pair configuration. Storage must be presented from both controllers to MEDITECH database servers over the FC Protocol. The configuration presents storage from both controllers to evenly balance the application load during normal operation.

ONTAP configuration

This section describes a sample deployment and provisioning procedures that use the relevant ONTAP commands. The emphasis is to show how storage is provisioned to implement the storage layout that NetApp recommends, which uses a high-availability controller pair. One of the major advantages with ONTAP is the ability to scale out without disturbing the existing high-availability pairs.

ONTAP licenses

After you have set up the storage controllers, apply licenses to enable the ONTAP features that NetApp recommends. The licenses for MEDITECH workloads are FC, CIFS, and NetApp Snapshot, SnapRestore, FlexClone, and SnapMirror technologies.

To configure licenses, open NetApp ONTAP System Manager, go to Configuration-Licenses, and then add the appropriate licenses.

Alternatively, run the following command to add licenses by using the CLI:

```
license add -license-code <code>
```

AutoSupport configuration

The NetApp AutoSupport tool sends summary support information to NetApp through HTTPS. To configure AutoSupport, run the following ONTAP commands:


```
autosupport modify -node * -state enable
autosupport modify -node * -mail-hosts <mailhost.customer.com>
autosupport modify -node prod1-01 -from prod1-01@customer.com
autosupport modify -node prod1-02 -from prod1-02@customer.com
autosupport modify -node * -to storageadmins@customer.com
autosupport modify -node * -support enable
autosupport modify -node * -transport https
autosupport modify -node * -hostnamesubj true
```

Hardware-assisted takeover configuration

On each node, enable hardware-assisted takeover to minimize the time that it takes to initiate a takeover in the unlikely event of a controller failure. To configure hardware-assisted takeover, complete the following steps:

1. Run the following ONTAP command to xxx.

Set the partner address option to the IP address of the management port for `prod1-01`.

```
MEDITECH::> storage failover modify -node prod1-01 -hwassist-partner-ip
<prod1-02-mgmt-ip>
```

2. Run the following ONTAP command to xxx:

Set the partner address option to the IP address of the management port for `cluster1-02`.

```
MEDITECH::> storage failover modify -node prod1-02 -hwassist-partner-ip
<prod1-01-mgmt-ip>
```

3. Run the following ONTAP command to enable hardware-assisted takeover on both the `prod1-01` and the `prod1-02` HA controller pair.

```
MEDITECH::> storage failover modify -node prod1-01 -hwassist true
MEDITECH::> storage failover modify -node prod1-02 -hwassist true
```

[Next: Aggregate Configuration.](#)

Aggregate configuration

NetApp RAID DP

NetApp recommends NetApp RAID DP technology as the RAID type for all aggregates in a NetApp FAS or AFF system, including regular NetApp Flash Pool aggregates. MEDITECH documentation might specify the use of RAID 10, but MEDITECH has approved the use of RAID DP.

RAID group size and number of RAID groups

The default RAID group size is 16. This size might or might not be optimal for the aggregates for the MEDITECH hosts at your specific site. For the number of disks that NetApp recommends that you use in a RAID group, see [NetApp TR-3838: Storage Subsystem Configuration Guide](#).

The RAID group size is important for storage expansion because NetApp recommends that you add disks to an aggregate with one or more groups of disks equal to the RAID group size. The number of RAID groups depends on the number of data disks and the RAID group size. To determine the number of data disks that you need, use the NetApp System Performance Modeler (SPM) sizing tool. After you determine the number of data disks, adjust the RAID group size to minimize the number of parity disks to within the recommended range for RAID group size per disk type.

For details on how to use the SPM sizing tool for MEDITECH environments, see [NetApp TR-4190: NetApp Sizing Guidelines for MEDITECH Environments](#).

Storage expansion considerations

When you expand aggregates with more disks, add the disks in groups that are equal to the aggregate RAID group size. Following this approach helps provide performance consistency throughout the aggregate.

For example, to add storage to an aggregate that was created with a RAID group size of 20, the number of disks that NetApp recommends adding is one or more 20-disk groups. So, you should add 20, 40, 60, and so on, disks.

After you expand aggregates, you can improve performance by running reallocation tasks on the affected volumes or aggregate to spread existing data stripes over the new disks. This action is helpful particularly if the existing aggregate was nearly full.



You should plan reallocation of schedules during nonproduction hours because it is a high-CPU and disk-intensive task.

For more information about using reallocation after an aggregate expansion, see [NetApp TR-3929: Reallocate Best Practices Guide](#).

Aggregate-level Snapshot copies

Set the aggregate-level NetApp Snapshot copy reserve to zero and disable the default aggregate Snapshot schedule. Delete any preexisting aggregate-level Snapshot copies if possible.

[Next: Storage Virtual Machine Configuration.](#)

Storage virtual machine configuration

This section pertains to deployment on ONTAP 8.3 and later versions.



A storage virtual machine (SVM) is also known as a Vserver in the ONTAP API and in the ONTAP CLI.

SVM for MEDITECH host LUNs

You should create one dedicated SVM per ONTAP storage cluster to own and to manage the aggregates that contain the LUNs for the MEDITECH hosts.

SVM language encoding setting

NetApp recommends that you set the language encoding for all SVMs. If no language encoding setting is specified at the time that the SVM is created, the default language encoding setting is used. The default language encoding setting is C.UTF-8 for ONTAP. After the language encoding has been set, you cannot modify the language of an SVM with Infinite Volume later.

The volumes that are associated with the SVM inherit the SVM language encoding setting unless you explicitly specify another setting when the volumes are created. To enable certain operations to work, you should use the language encoding setting consistently in all volumes for your site. For example, SnapMirror requires the source and destination SVM to have the same language encoding setting.

[Next: Volume Configuration.](#)

Volume configuration

Volume provisioning

MEDITECH volumes that are dedicated for MEDITECH hosts can be either thick or thin provisioned.

Default volume-level Snapshot copies

Snapshot copies are created as part of the backup workflow. Each Snapshot copy can be used to access the data stored in the MEDITECH LUNs at different times. The MEDITECH- approved backup solution creates thin-provisioned FlexClone volumes based on these Snapshot copies to provide point-in-time copies of the MEDITECH LUNs. The MEDITECH environment is integrated with an approved backup software solution. Therefore, NetApp recommends that you disable the default Snapshot copy schedule on each of the NetApp FlexVol volumes that make up the MEDITECH production database LUNs.

Important: FlexClone volumes share parent data volume space, so it is vital for the volume to have enough space for the MEDITECH data LUNs and the FlexClone volumes that the backup servers create. FlexClone volumes do not occupy more space the way that data volumes do. However, if there are huge deletions on the MEDITECH LUNs in a short time, the clone volumes might grow.

Number of volumes per aggregate

For a NetApp FAS system that uses Flash Pool or NetApp Flash Cache caching, NetApp recommends provisioning three or more volumes per aggregate that are dedicated for storing the MEDITECH program, dictionary, and data files.

For AFF systems, NetApp recommends dedicating four or more volumes per aggregate for storing the MEDITECH program, dictionary, and data files.

Volume-level reallocate schedule

The data layout of storage becomes less optimal over time, especially when it is used by write-intensive workloads such as the MEDITECH Expanse, 6.x, and C/S 5.x platforms. Over time, this situation might increase sequential read latency, resulting in a longer time to complete the backup. Bad data layout or fragmentation can also affect the write latency. You can use volume-level reallocation to optimize the layout of data on disk to improve write latencies and sequential read access. The improved storage layout helps to complete the backup within the allocated time window of 8 hours.

Best practice

At a minimum, NetApp recommends that you implement a weekly volume reallocation schedule to run reallocation operations during the allocated maintenance downtime or during off-peak hours on a production site.



NetApp highly recommends that you run the reallocation task on one volume at a time per controller.

For more information about determining an appropriate volume reallocation schedule for your production database storage, see section 3.12 in [NetApp TR-3929: Reallocate Best Practices Guide](#). That section also guides you on how to create a weekly reallocation schedule for a busy site.

[Next: LUN Configuration.](#)

LUN configuration

The number of MEDITECH hosts in your environment determines the number of LUNs that are created within the NetApp FAS or AFF system. The Hardware Configuration Proposal specifies the size of each LUN.

LUN provisioning

MEDITECH LUNs that are dedicated for MEDITECH hosts can be either thick or thin provisioned.

LUN operating system type

To properly align the LUNs that are created, you must correctly set the operating system type for the LUNs. Misaligned LUNs incur unnecessary write operation overhead, and it is costly to correct a misaligned LUN.

The MEDITECH host server typically runs in the virtualized Windows Server environment by using the VMware vSphere hypervisor. The host server can also run in the Windows Server environment on a bare-metal server. To determine the correct operating system type value to set, refer to the “LUN Create” section of [Clustered Data ONTAP 8.3 Commands: Manual Page Reference](#).

LUN size

To determine the LUN size for each MEDITECH host, see the Hardware Configuration Proposal (new deployment) or the Hardware Evaluation Task (existing deployment) document from MEDITECH.

LUN presentation

MEDITECH requires that storage for program, dictionary, and data files be presented to MEDITECH hosts as LUNs by using the FC Protocol. In the VMware virtual environment, the LUNs are presented to the VMware ESXi servers that host the MEDITECH hosts. Then each LUN that is presented to the VMware ESXi server is mapped to each MEDITECH host VM by using RDM in the physical compatibility mode.

You should present the LUNs to the MEDITECH hosts by using the proper LUN naming conventions. For example, for easy administration, you must present the LUN `MTFS01E` to the MEDITECH host `mt-host-01`.

Refer to the MEDITECH Hardware Configuration Proposal when you consult with the MEDITECH and backup system installer to devise a consistent naming convention for the LUNs that the MEDITECH hosts use.

An example of a MEDITECH LUN name is `MTFS05E`, in which:

- `MTFS` denotes the MEDITECH file server (for the MEDITECH host).
- `05` denotes host number 5.
- `E` denotes the Windows E drive.

[Next: Initiator Group Configuration.](#)

Initiator group configuration

When you use FC as the data network protocol, create two initiator groups (igroups) on each storage controller. The first igroup contains the WWPNs of the FC host interface cards on the VMware ESXi servers that host the MEDITECH host VMs (igroup for MEDITECH).

You must set the MEDITECH igroup operating system type according to the environment setup. For example:

- Use the igroup operating system type `Windows` for applications that are installed on bare-metal-server hardware in a Windows Server environment.
- Use the igroup operating system type `VMware` for applications that are virtualized by using the VMware vSphere hypervisor.



The operating system type for an igroup might be different from the operating system type for a LUN. As an example, for virtualized MEDITECH hosts, you should set the igroup operating system type to `VMware`. For the LUNs that are used by the virtualized MEDITECH hosts, you should set the operating system type to `Windows 2008 or later`. Use this setting because the MEDITECH host operating system is the Windows Server 2008 R2 64-bit Enterprise Edition.

To determine the correct value for the operating system type, see the sections “LUN Igroup Create” and “LUN Create” in the [Clustered Data ONTAP 8.2 Commands: Manual Page Reference](#).

[Next: LUN Mappings.](#)

LUN mappings

LUN mappings for the MEDITECH hosts are established when the LUNs are created.

MEDITECH modules and components

The MEDITECH application covers several modules and components. The following table lists the functions that are covered by these modules. For additional information about setting up and deploying these modules, see the MEDITECH documentation.

| Function | Type |
|---------------------------------|--|
| Connectivity | <ul style="list-style-type: none"> • Web server • Live application server (WI – Web Integration) • Test application server (WI) • SAML authentication server (WI) • SAML proxy server (WI) • Database server |
| Infrastructure | <ul style="list-style-type: none"> • File server • Background Job Client • Connection server • Transaction server |
| Scanning and archiving | <ul style="list-style-type: none"> • Image server |
| Data repository | <ul style="list-style-type: none"> • SQL Server |
| Business and clinical analytics | <ul style="list-style-type: none"> • Live intelligence server (BCA) • Test intelligence server (BCA) • Database server (BCA) |
| Home care | <ul style="list-style-type: none"> • Remote site solution • Connectivity • Infrastructure • Printing • Field devices • Scanning • Hosted site requirements • Firewall configuration |
| Support | <ul style="list-style-type: none"> • Background Job Client (CALs – Client Access License) |
| User devices | <ul style="list-style-type: none"> • Tablets • Fixed devices |
| Printing | <ul style="list-style-type: none"> • Live network print server (required; might already exist) • Test network print server (required; might already exist) |

| Function | Type |
|-------------------------|--|
| Third-party requirement | <ul style="list-style-type: none"> • First Databank (FDB) MedKnowledge Framework v4.3 |

Acknowledgments

The following people contributed to the creation of this guide.

- Brandon Agee, Technical Marketing Engineer, NetApp
- Atul Bhalodia, Technical Marketing Engineer, NetApp
- Ketan Mota, Senior Product Manager, NetApp
- John Duignan, Solutions Architect—Healthcare, NetApp
- Jon Ebmeier, Cisco
- Mike Brennan, Cisco

Where to find additional information

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

FlexPod design zone

- [FlexPod Design Zone](#)
- [FlexPod Data Center with FC Storage \(MDS Switches\) Using NetApp AFF, vSphere 6.5U1, and Cisco UCS Manager](#)

NetApp technical reports

- [TR-3929: Reallocate Best Practices Guide](#)
- [TR-3987: Snap Creator Framework Plug-In for InterSystems Caché](#)
- [TR-4300i: NetApp FAS and All-Flash Storage Systems for MEDITECH Environments Best Practices Guide](#)
- [TR-4017: FC SAN Best Practices](#)
- [TR-3446: SnapMirror Async Overview and Best Practices Guide](#)

ONTAP documentation

- [NetApp Product Documentation](#)
- [Virtual Storage Console \(VSC\) for vSphere documentation](#)
- [ONTAP 9 Documentation Center](#):
 - [FC Express Guide for ESXi](#)
- [All ONTAP 9.3 Documentation](#):
 - [Software Setup Guide](#)
 - [Disks and Aggregates Power Guide](#)

- [SAN Administration Guide](#)
- [SAN Configuration Guide](#)
- [FC Configuration for Windows Express Guide](#)
- [FC SAN Optimized AFF Setup Guide](#)
- [High-Availability Configuration Guide](#)
- [Logical Storage Management Guide](#)
- [Performance Management Power Guide](#)
- [SMB/CIFS Configuration Power Guide](#)
- [SMB/CIFS Reference](#)
- [Data Protection Power Guide](#)
- [Data Protection Tape Backup and Recovery Guide](#)
- [NetApp Encryption Power Guide](#)
- [Network Management Guide](#)
- [Commands: Manual Page Reference for ONTAP 9.3](#)

Cisco Nexus, MDS, Cisco UCS, and Cisco UCS Manager guides

- [Cisco UCS Servers Overview](#)
- [Cisco UCS Blade Servers Overview](#)
- [Cisco UCS B200 M5 Datasheet](#)
- [Cisco UCS Manager Overview](#)
- [Cisco UCS Manager 3.2\(3a\) Infrastructure Bundle](#) (requires Cisco.com authorization)
- [Cisco Nexus 9300 Platform Switches](#)
- [Cisco MDS 9132T FC Switch](#)

FlexPod for Medical Imaging

TR-4865: FlexPod for Medical Imaging

Jaya Kishore Esanakula and Atul Bhalodia, NetApp

Medical imaging accounts for 70% of all data that is generated by Healthcare organizations. As digital modalities continue to advance and new modalities emerge, the amount of data will continue to increase. For example, the transition from analog to digital pathology will dramatically increase image sizes at a rate that will challenge any data management strategies currently in place.

COVID-19 has clearly reshaped the digital transformation; according to a recent [report](#), COVID-19 has accelerated digital commerce by 5 years. The technological innovation driven by problem solvers is fundamentally changing the way that we go about our daily life. This technology-driven change will overhaul many critical aspects of our life, including healthcare.

Healthcare is poised to undergo a major change in the coming years. COVID is accelerating innovation in healthcare that will propel the industry by at least several years. At the heart of this change is the need to make

healthcare more flexible in handling pandemics by being more affordable, available, and accessible, without compromising reliability.

At the foundation of this healthcare change is a well-designed platform. One of the key metrics to measure the platform is the ease with which platform changes can be implemented. Speed is the new scale and data protection cannot be compromised. Some of the world's most critical data is being created and consumed by the clinical systems that support clinicians. NetApp has made critical data available for patient care where the clinicians need it, on premise, in the cloud, or in a hybrid setting. Hybrid multi- cloud environments are the current state of the art for IT architecture.

Healthcare as we know it revolves around providers (doctors, nurses, radiologists, medical device technicians, and so on) and patients. As we bring patients and providers closer together, making the geographic location a mere data point, it becomes even more important for the underlying platform to be available when providers and patients need it. The platform must be both efficient and cost-effective in the long term. In their efforts to drive patient care costs even lower, [Accountable Care Organizations](#) (ACOs) would be empowered by an efficient platform.

When it comes to health information systems used by healthcare organizations, the question of build versus purchase tends to have a single answer: purchase. This could be for many subjective reasons. Purchasing decisions made over many years can create heterogeneous information systems. Each system has a specific set of requirements for the platform that they are deployed on. The most significant issue is the large, diverse set of storage protocols and performance levels that information systems require, which makes platform standardization and optimal operational efficiency a significant challenge. Healthcare organizations cannot focus on mission critical issues because their attention is spread thin by trivial operational needs like the large set of platforms that require a diversified set of skills and thus SME retention.

The challenges can be classified into the following categories:

- Heterogeneous storage needs
- Departmental silos
- IT operational complexity
- Cloud connectivity
- Cybersecurity
- Artificial intelligence and deep learning

With FlexPod, you get a single platform that supports FC, FCoE, iSCSI, NFS/pNFS, SMB/CIFS and so on from a single platform. People, processes, and technology are part of the DNA that FlexPod is designed and built upon. FlexPod adaptive QoS helps to break down the departmental silos by supporting multiple mission critical clinical systems on the same underlying FlexPod platform. FlexPod is FedRAMP certified and FIPS 140-2 certified. Additionally, healthcare organizations are faced with opportunities such as artificial intelligence and deep learning. FlexPod and NetApp solve these challenges and make the data available where it is needed on premises or in a hybrid multi- cloud setting in a standardized platform. For more information and a series customer success stories, see [FlexPod Healthcare](#).

Typical medical imaging information and PACS systems have the following set of capabilities:

- Reception and registration
- Scheduling
- Imaging
- Transcription
- Management

- Data exchange
- Image archive
- Image viewing for image capturing and reading for technicians and image viewing for clinicians

Regarding imaging, the healthcare sector is trying to solve the following clinical challenges:

- Wider adoption of [natural language processing](#) (NLP)-based assistants by technicians and physicians for image reading. Radiology department can benefit from voice recognition to transcribe reports. NLP can be used to identify and anonymize a patient's record, specifically DICOM tags embedded in the DICOM image. NLP capabilities require high performing platforms with low latency response times for image processing. FlexPod QoS not only delivers and performance but also provides mature capacity projections for future growth.
- Wider adoption of standardized clinical pathways and protocols by ACOs and community health organizations. Historically, clinical pathways have been used as a static set of guidelines rather than an integrated workflow that guides clinical decisions. With advancements in NLP and image processing, DICOM tags in images can be integrated into clinical pathways as facts to drive clinical decisions. Therefore, these processes require high performance, low latency, and high throughput from the underlying infrastructure platform and storage systems.
- ML models that leverage convolutional neural networks enable automation of image- processing capabilities in real time and thus require infrastructure that is GPU-capable. FlexPod offers both CPU and GPU compute components built into the same system, and CPUs and GPUs can be scaled independently of each other.
- If DICOM tags are used as facts in clinical best-practice advisories, then the system must perform more reads of DICOM artifacts with low latency and high throughput.
- When evaluating images, real-time collaboration between radiologists across organizations requires high performance graphics processing in the end- user compute devices. NetApp provides industry- leading VDI solutions specifically designed and proven for high-end graphics use cases. More information can be found [here](#).
- Image and media management across ACO health organizations can uses a single platform, regardless of the system of record for the image, by using protocols such as Digital Imaging and Communications in Medicine ([DICOM](#)) and web access to DICOM-persistent objects ([WADO](#))
- Health information exchange ([HIE](#)) includes images embedded in messages.
- Mobile modalities, such as handheld, wireless scanning devices (for example, pocket handheld ultrasound scanners attached to a phone), require a robust network infrastructure with DoD-level security, reliability, and latency at the edge, the core, and in the cloud. [A data fabric enabled by NetApp](#) provide organizations with this capability at scale.
- Newer modalities have exponential storage needs; for example, CT and MRI require a few hundred MBs for each modality, but digital pathology images (including whole slide imaging) can be a few GBs in size. FlexPod is designed with [performance, reliability and scaling as foundational traits](#).

A well-architected medical imaging system platform is at the heart of innovation. The FlexPod architecture provides flexible compute and storage capabilities with industry-leading storage efficiency.

Overall solution benefits

By running an imaging application environment on a FlexPod architectural foundation, your healthcare organization can expect to see an improvement in staff productivity and a decrease in capital and operating expenses. FlexPod provides a rigorously tested, prevalidated, converged that is engineered and designed to deliver predictable low-latency system performance and high availability. This approach results in high comfort levels and, ultimately, optimal response times for users of the medical imaging system.

Different components of the imaging system might require the storage of data in SMB/CIFS, NFS, Ext4, or NTFS file systems. That requirement means that the infrastructure must provide data access over the NFS, SMB/CIFS, and SAN protocols. A single NetApp storage system can support the NFS, SMB/CIFS, and SAN protocols, thus eliminating the need for the legacy practice of protocol-specific storage systems.


The FlexPod infrastructure is a modular, converged, virtualized, scalable (scale-out and scale-up), and cost-effective platform. With the FlexPod platform, you can independently scale out compute, network, and storage to accelerate your application deployment. And the modular architecture enables nondisruptive operations even during system scale-out and upgrade activities.

FlexPod delivers several benefits that are specific to the medical imaging industry:

- **Low-latency system performance.** Radiologist time is a high-value resource, and efficient use of a radiologist's time is paramount. Waiting for images or videos to load can contribute to clinician burnout and can affect clinician's efficiency and patient safety.
- **Modular architecture.** FlexPod components are connected through a clustered server, a storage management fabric, and a cohesive management toolset. As imaging facilities grow year over year and the number of studies increase, there will be a need for the underlying infrastructure to scale accordingly. FlexPod can scale compute, storage, and network independently.
- **Quicker deployment of infrastructure.** Whether it is in an existing data center or a remote location, the integrated and tested design of FlexPod Datacenter with Medical Imaging enables you to get the new infrastructure up and running in less time, with less effort.
- **Accelerated application deployment.** A prevalidated architecture reduces implementation integration time and risk for any workload, and NetApp technology automates infrastructure deployment. Whether you use the solution for an initial rollout of medical imaging, a hardware refresh, or expansion, you can shift more resources to the business value of the project.
- **Simplified operations and lower costs.** You can eliminate the expense and complexity of legacy proprietary platforms by replacing them with a more efficient and scalable shared resource that can meet the dynamic needs of your workload. This solution delivers higher infrastructure resource utilization for greater return on investment (ROI).
- **Scale-out architecture.** You can scale SAN and NAS from terabytes to tens of petabytes without reconfiguring running applications.
- **Nondisruptive operations.** You can perform storage maintenance, hardware lifecycle operations, and software upgrades without interrupting your business.
- **Secure multitenancy.** This benefit supports the increased needs of virtualized server and storage shared infrastructure, enabling secure multitenancy of facility-specific information, particularly if you are hosting multiple instances of databases and software.
- **Pooled resource optimization.** This benefit can help you reduce physical server and storage controller counts, load-balance workload demands, and boost utilization while improving performance.
- **Quality of service (QoS).** FlexPod offers QoS on the entire stack. These industry-leading QoS storage policies enable differentiated service levels in a shared environment. These policies help optimize performance for workloads and help to isolate and control runaway applications.
- **Support for storage tier SLAs by using QoS.** You don't have to deploy different storage systems for the different storage tiers that a medical imaging environment typically requires. A single storage cluster with multiple NetApp FlexVol volumes with specific QoS policies for different tiers can serve that purpose. With this approach, storage infrastructure can be shared by dynamically accommodating the changing needs of a particular storage tier. NetApp AFF can support different SLAs for storage tiers by allowing QoS at the level of the FlexVol volume, thus eliminating the need for different storage systems for different storage tiers for the application.

- **Storage efficiency.** Medical images are typically pre-compressed by the imaging application to jpeg2k lossless compression which is around 2.5:1. However, this is imaging application and vendor specific. In larger imaging application environments (greater than 1PB), 5-10% storage savings are possible, and you can reduce storage costs with NetApp storage efficiency features. Work with your imaging application vendors and your NetApp subject matter expert to unlock potential storage efficiencies for your medical imaging system.
- **Agility.** With the industry-leading workflow automation, orchestration, and management tools that FlexPod systems offer, your IT team can be far more responsive to business requests. These business requests can range from medical imaging backup and provisioning of additional test and training environments to analytics database replications for population health- management initiatives.
- **Higher productivity.** You can quickly deploy and scale this solution for optimal clinician end-user experiences.
- **Data fabric.** Your data fabric powered by NetApp weaves data together across sites, beyond physical boundaries, and across applications. Your data fabric powered by NetApp is built for data-driven enterprises in a data-centric world. Data is created and used in multiple locations, and it often needs to be leveraged and shared with other locations, applications, and infrastructures. So, you want a consistent and integrated way to manage it. This solution provides a way to manage data that puts your IT team in control and that simplifies ever-increasing IT complexity.
- **FabricPool.** NetApp ONTAP FabricPool helps 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 the need to rearchitect the application infrastructure. FlexPod can benefit from the storage tiering capabilities of FabricPool to make more efficient use of ONTAP flash storage. For full information, see [FlexPod with FabricPool](#).
- **FlexPod security.** Security is at the very foundation of FlexPod. In the past few years, ransomware has become a significant and increasing threat. Ransomware is malware that is based on crypto virology, the use of cryptography to build malicious software. This malware can use both symmetric and asymmetric key encryption to lock a victim's data and demand a ransom to provide the key to decrypt the data. To learn how FlexPod helps mitigate threats like ransomware, see [The Solution to Ransomware](#). FlexPod infrastructure components are also Federal Information Processing Standard (FIPS) 140-2 compliant.
- **FlexPod Cooperative Support.** NetApp and Cisco have established FlexPod Cooperative Support, a strong, scalable, and flexible support model to meet the unique support requirements of the FlexPod converged infrastructure. This model uses the combined experience, resources, and technical support expertise of NetApp and Cisco to provide a streamlined process for identifying and resolving your FlexPod support issue, regardless of where the problem resides. The FlexPod Cooperative Support model helps confirm that your FlexPod system operates efficiently and benefits from the most up-to-date technology, while providing an experienced team to help resolve integration issues.

FlexPod Cooperative Support is especially valuable if your healthcare organization runs business-critical applications. The illustration below shows an overview of the FlexPod Cooperative Support model.



Scope

This document provides a technical overview of a Cisco Unified Computing System (Cisco UCS) and NetApp ONTAP-based FlexPod infrastructure for hosting this medical imaging solution.

Audience

This document is intended for technical leaders in the healthcare industry and for Cisco and NetApp partner solutions engineers and professional services personnel. NetApp assumes that the reader has a good understanding of compute and storage sizing concepts as well as technical familiarity with the medical imaging system, Cisco UCS, and NetApp storage systems.

Medical imaging application

A typical medical imaging application offers a suite of applications that together make an enterprise-grade imaging solution for small, medium, and large healthcare organizations.

At the heart of the product suite are the following clinical capabilities:

- Enterprise imaging repository
- Supports traditional image sources such as radiology and cardiology. Also supports other care areas like ophthalmology, dermatology, colonoscopy, and other medical imaging objects like photos and videos.
- [Picture archiving and communication system \(PACS\)](#), which is a computerized means of replacing the roles of conventional radiological film
- Enterprise Imaging Vendor Neutral Archive (VNA):

- Scalable consolidation of DICOM and non-DICOM documents
- Centralized Medical Imaging system
- Support for document synchronization and data integrity between multiple (PACSs) in the enterprise
- Document lifecycle management by a rules-based expert system that leverages document metadata, such as:
 - Modality type
 - Age of study
 - Patient age (current and at the time of image capture)
- Single point of integration within and outside (HIE) of the enterprise:
- Context- aware document linking
- Health Level Seven International (HL7), DICOM, and WADO
- Storage- agnostic archiving capability
- Integration with other health information systems that use HL7 and context-aware linking:
 - Enables EHRs to implement direct links to patient images from patient charts, imaging workflows, and so on.
 - Helps embed a patient's longitudinal care image history into EHRs.
- Radiology technologist workflows
- Enterprise zero footprint viewers for image viewing from anywhere on any capable device
- Analytical tools that leverage retrospective and real-time data:
 - Compliance reporting
 - Operational reports
 - Quality control and quality assurance reports

Size of the healthcare organization and platform sizing

Healthcare organizations can be broadly classified by using standards-based methods that help programs such as ACO. One such classification uses the concept of a clinical integrated network (CIN). A group of hospitals can be called a CIN if they collaborate and adhere to proven standard clinical protocols and pathways to improve the value of care and reduce patient costs. Hospitals within a CIN have controls and practices in place to onboard physicians who follow the core values of the CIN. Traditionally, an integrated delivery networks (IDN) has been limited to hospitals and physician groups. A CIN crosses traditional IDN boundaries, and a CIN can still be part of an ACO. Following the principles of a CIN, healthcare organizations can be classified into small, medium, and large.

Small healthcare organizations

A healthcare organization is small if it includes only a single hospital with ambulatory clinics and an inpatient department, but it is not part of a CIN. Physicians work as caregivers and coordinate patient care during a care continuum. These small organizations typically include physician-operated facilities. They might or might not offer emergency and trauma care as integrated care for the patient. Typically, a small-sized healthcare organization performs about 250,000 clinical imaging studies annually. Imaging centers are considered to be small healthcare organizations and they do provide imaging services. Some also provide radiology dictation services to other organizations.

Medium healthcare organizations

A healthcare organization considered to be of medium size if it includes multiple hospital systems with focused organizations, such as the following:

- Adult care clinics and adult inpatient hospitals
- Labor and delivery departments
- Childcare clinics and child inpatient hospitals
- A cancer treatment center
- Adult emergency departments
- Child emergency departments
- A family medicine and primary care office
- An adult trauma care center
- A child trauma care center

In a medium-sized healthcare organization, physicians follow the principles of a CIN and operate as a single unit. Hospitals have separate hospital, physician, and pharmacy billing functions. Hospitals might be associated with academic research institutes and perform interventional clinical research and trials. A medium healthcare organization performs as many as 500,000 clinical imaging studies annually.

Large healthcare organizations

A healthcare organization is considered to be large if it includes the traits of a medium- sized healthcare organization and offers the medium-sized clinical capabilities to the community in multiple geographical locations.

A large healthcare organization typically performs the following functions:

- Has a central office to manage the overall functions
- Participates in joint ventures with other hospitals
- Negotiates rates with payer organizations annually
- Negotiates payer rates by state and region
- Participates in Meaningful Use (MU) programs
- Performs advanced clinical research across population health cohorts by using standards-based population health management (PHM) tools
- Performs up to one million clinical imaging studies annually

Some large healthcare organizations that participate in a CIN also have AI- based imaging reading capabilities. These organizations typically perform one to two million clinical imaging studies annually.

Before you look into how these different-sized organizations translate into an optimally sized FlexPod system, you should understand the various FlexPod components and the different capabilities of a FlexPod system.

FlexPod

Cisco Unified Computing System

Cisco UCS consists of a single management domain that is interconnected with a unified I/O infrastructure. Cisco UCS for medical imaging environments has been aligned with NetApp medical imaging system

infrastructure recommendations and best practices so that the infrastructure can deliver critical patient information with maximum availability.

The compute foundation of enterprise medical imaging is Cisco UCS technology, with its integrated systems management, Intel Xeon processors, and server virtualization. These integrated technologies solve data center challenges and enable you to meet your goals for data center design with a typical medical imaging system. Cisco UCS unifies LAN, SAN, and systems management into one simplified link for rack servers, blade servers, and virtual machines (VMs). Cisco UCS consists of a redundant pair of Cisco UCS fabric interconnects that provide a single point of management and a single point of control for all I/O traffic.

Cisco UCS uses service profiles so that virtual servers in the Cisco UCS infrastructure are configured correctly and consistently. Service profiles include critical server information about the server identity, such as LAN and SAN addressing, I/O configurations, firmware versions, boot order, network virtual LAN (VLAN), physical port, and QoS policies. Service profiles can be dynamically created and associated with any physical server in the system in minutes rather than in hours or days. The association of service profiles with physical servers is performed as a single, simple operation that enables migration of identities between servers in the environment without requiring any physical configuration changes. It also facilitates rapid bare-metal provisioning of replacements for failed servers.

The use of service profiles helps confirm that servers are configured consistently throughout the enterprise. When using multiple Cisco UCS management domains, Cisco UCS Central can use global service profiles to synchronize configuration and policy information across domains. If maintenance must be performed in one domain, the virtual infrastructure can be migrated to another domain. With this approach, even when a single domain is offline, applications continue to run with high availability.

Cisco UCS is a next-generation solution for blade and rack server computing. The system integrates a low-latency, lossless, 40GbE unified network fabric with enterprise-class, x86-architecture servers. The system is an integrated, scalable, multi-chassis platform in which all resources participate in a unified management domain. Cisco UCS accelerates the delivery of new services simply, reliably, and securely through end-to-end provisioning and migration support for both virtualized and nonvirtualized systems. Cisco UCS provides the following features:

- Comprehensive management
- Radical simplification
- High performance

Cisco UCS consists of the following components:

- **Compute.** The system is based on an entirely new class of computing system that incorporates rack-mounted and blade servers based on the Intel Xeon scalable processor product family.
- **Network.** The system is integrated into a low-latency, lossless, 40Gbps unified network fabric. This network foundation consolidates LANs, SANs, and high-performance computing networks, which are separate networks today. The unified fabric lowers costs by reducing the number of network adapters, switches, and cables and also by decreasing power and cooling requirements.
- **Virtualization.** The system unleashes the full potential of virtualization by enhancing the scalability, performance, and operational control of virtual environments. Cisco security, policy enforcement, and diagnostic features are now extended into virtualized environments to better support changing business and IT requirements.
- **Storage access.** The system provides consolidated access to both SAN storage and NAS over the unified fabric. It is also an ideal system for software-defined storage. By combining the benefits of a single framework to manage both the compute and the storage servers in a single pane, QoS can be implemented if needed to inject I/O throttling in the system. And your server administrators can preassign storage-access policies to storage resources, which simplifies storage connectivity and management and

can help increase productivity. In addition to external storage, both rack and blade servers have internal storage that can be accessed through built-in hardware RAID controllers. By setting up the storage profile and disk configuration policy in Cisco UCS Manager, the storage needs of the host OS and application data are fulfilled by user-defined RAID groups. The result is high availability and better performance.

- **Management.** The system uniquely integrates all system components so that the entire solution can be managed as a single entity by Cisco UCS Manager. To manage all system configuration and operations, Cisco UCS Manager has an intuitive GUI, a CLI, and a powerful scripting library module for Microsoft Windows PowerShell that are built on a robust API.

Cisco Unified Computing System fuses access layer networking and servers. This high-performance, next-generation server system gives your data center a high degree of workload agility and scalability.

Cisco UCS Manager

Cisco UCS Manager provides unified, embedded management for all software and hardware components in Cisco UCS. By using single-connection technology, UCS Manager manages, controls, and administers multiple chassis for thousands of VMs. Through an intuitive GUI, a CLI, or an XML API, your administrators use the software to manage the entire Cisco UCS as a single logical entity. Cisco UCS Manager resides on a pair of Cisco UCS 6300 Series Fabric Interconnects that use clustered, active-standby configuration for high availability.

Cisco UCS Manager offers a unified embedded management interface that integrates your servers, network, and storage. Cisco UCS Manager performs auto discovery to detect the inventory of, to manage, and to provision system components that you add or change. It offers a comprehensive set of XML APIs for third-party integration, and it exposes 9,000 points of integration. It also facilitates custom development for automation, for orchestration, and to achieve new levels of system visibility and control.

Service profiles benefit both virtualized and nonvirtualized environments. They increase the mobility of nonvirtualized servers, such as when you move workloads from server to server or when you take a server offline for service or upgrade. You can also use profiles in conjunction with virtualization clusters to bring new resources online easily, complementing existing VM mobility.

For more information about Cisco UCS Manager, see the [Cisco UCS Manager product page](#).

Cisco UCS differentiators

Cisco Unified Computing System is revolutionizing the way that servers are managed in the data center. See the following unique differentiators of Cisco UCS and Cisco UCS Manager:

- **Embedded management.** In Cisco UCS, the servers are managed by the embedded firmware in the fabric interconnects, eliminating the need for any external physical or virtual devices to manage them.
- **Unified fabric.** In Cisco UCS, from blade server chassis or rack servers to fabric interconnects, a single Ethernet cable is used for LAN, SAN, and management traffic. This converged I/O reduces the number of cables, SFPs, and adapters that you need, in turn reducing your capital and operational expenses for the overall solution.
- **Autodiscovery.** By simply inserting the blade server in the chassis or by connecting rack servers to the fabric interconnects, discovery and inventory of compute resource occurs automatically without any management intervention. The combination of unified fabric and auto discovery enables the wire-once architecture of Cisco UCS, where its compute capability can be extended easily while keeping the existing external connectivity to LAN, SAN, and management networks.
- **Policy-based resource classification.** When a compute resource is discovered by Cisco UCS Manager, it can be automatically classified to a given resource pool based on the policies that you defined. This capability is useful in multitenant cloud computing.

- **Combined rack and blade server management.** Cisco UCS Manager can manage B-Series blade servers and C-Series rack servers under the same Cisco UCS domain. This feature, along with stateless computing, makes compute resources truly hardware form factor–agnostic.
- **Model-based management architecture.** The Cisco UCS Manager architecture and management database are model-based and data-driven. The open XML API that is provided to operate on the management model enables easy and scalable integration of Cisco UCS Manager with other management systems.
- **Policies, pools, and templates.** The management approach in Cisco UCS Manager is based on defining policies, pools, and templates instead of a cluttered configuration. It enables a simple, loosely coupled, data-driven approach in managing compute, network, and storage resources.
- **Loose referential integrity.** In Cisco UCS Manager, a service profile, a port profile, or policies can refer to other policies or to other logical resources with loose referential integrity. A referred policy cannot exist at the time of authoring the referring policy, but a referred policy can be deleted even though other policies are referring to it. This feature enables different subject-matter experts to work independently from each other. You gain great flexibility by enabling different experts from different domains—such as network, storage, security, server, and virtualization—to work together to accomplish a complex task.
- **Policy resolution.** In Cisco UCS Manager, you can create a tree structure of organizational unit hierarchy that mimics the real-life tenants and organizational relationships. You can define various policies, pools, and templates at different levels of your organizational hierarchy. A policy that refers to another policy by name is resolved in the organizational hierarchy with the closest policy match. If no policy with a specific name is found in the hierarchy of the root organization, then a special policy named “default” is searched. This policy resolution practice enables automation-friendly management APIs and provides great flexibility to the owners of the different organizations.
- **Service profiles and stateless computing.** A service profile is a logical representation of a server, carrying its various identities and policies. You can assign this logical server to any physical compute resource, as long as it meets the resource requirements. Stateless computing enables procurement of a server within minutes, which used to take days in legacy server management systems.
- **Built-in multitenancy support.** The combination of policies, pools, templates, a loose referential integrity, policy resolution in organizational hierarchy, and a service profiles- based approach to compute resources makes Cisco UCS Manager inherently friendly to multitenant environments that are typically observed in private and public clouds.
- **Extended memory.** The enterprise-class Cisco UCS B200 M5 Blade Server extends the capabilities of the Cisco Unified Computing System portfolio in a half-width blade form factor. The Cisco UCS B200 M5 harnesses the power of the latest Intel Xeon scalable- processor CPUs with up to 3TB of RAM. This feature enables the huge VM-to-physical- server ratio that many deployments need or enables certain architectures to support large memory operations, such as big data.
- **Virtualization- aware network.** Cisco Virtual Machine Fabric Extender (VM-FEX) technology makes the access network layer aware of host virtualization. This awareness prevents pollution of compute and network domains with virtualization when a virtual network is managed by port profiles that are defined by your network administrator team. VM-FEX also offloads hypervisor CPU by performing switching in the hardware, thus enabling the hypervisor CPU to perform more virtualization- related tasks. To simplify cloud management, VM-FEX technology is well integrated with VMware vCenter, Linux Kernel-Based Virtual Machine (KVM), and Microsoft Hyper-V SR-IOV.
- **Simplified QoS.** Even though FC and Ethernet are converged in the Cisco UCS, built-in support for QoS and lossless Ethernet make it seamless. By representing all system classes in one GUI panel, network QoS is simplified in Cisco UCS Manager.

Cisco Nexus IP and MDS switches

Cisco Nexus switches and Cisco MDS multilayer directors give you enterprise-class connectivity and SAN

consolidation. Cisco multiprotocol storage networking helps reduce your business risk by providing flexibility and options: FC, Fiber Connection (FICON), FC over Ethernet (FCoE), iSCSI, and FC over IP (FCIP).

Cisco Nexus switches offer one of the most comprehensive data center network feature sets in a single platform. They deliver high performance and density for both the data center and the campus core. They also offer a full feature set for data center aggregation, end-of-row, and data center interconnect deployments in a highly resilient modular platform.

Cisco UCS integrates compute resources with Cisco Nexus switches and a unified fabric that identifies and handles different types of network traffic. This traffic includes storage I/O, streamed desktop traffic, management, and access to clinical and business applications. You get the following capabilities:

- **Infrastructure scalability.** Virtualization, efficient power and cooling, cloud scale with automation, high density, and performance all support efficient data center growth.
- **Operational continuity.** The design integrates hardware, Cisco NX-OS software features, and management to support zero-downtime environments.
- **Transport flexibility.** You can incrementally adopt new networking technologies with this cost-effective solution.

Together, Cisco UCS with Cisco Nexus switches and MDS multilayer directors provide a compute, networking, and SAN connectivity solution for an enterprise medical Imaging system.

NetApp all-flash storage

NetApp storage that runs ONTAP software reduces your overall storage costs while delivering the low- latency read and write response times and high IOPS that medical imaging system workloads need. To create an optimal storage system that meets a typical medical imaging system requirement, ONTAP supports both all-flash and hybrid storage configurations. NetApp flash storage gives medical imaging system customers like you the key components of high performance and responsiveness to support latency-sensitive medical imaging system operations. By creating multiple fault domains in a single cluster, NetApp technology can also isolate your production environments from your nonproduction environments. And by guaranteeing that system performance do not drop below a certain level for workloads with ONTAP minimum QoS, NetApp reduces performance issues for your system.

The scale-out architecture of ONTAP software can flexibly adapt to your various I/O workloads. To deliver the necessary throughput and low latency that clinical applications need and to provide a modular scale-out architecture, all-flash configurations are typically used in ONTAP architectures. NetApp AFF nodes can be combined in the same scale-out cluster with hybrid (HDD and flash) storage nodes, suitable for storing large datasets with high throughput. You can clone, replicate, and back up your medical imaging system environment from expensive SSD storage to more economical HDD storage on other nodes. With NetApp cloud-enabled storage and a data fabric delivered by NetApp, you can back up to object storage on premises or in the cloud.

For medical imaging, ONTAP has been validated by most leading medical imaging systems. That means it has been tested to deliver fast and reliable performance for medical imaging. Additionally, the following features simplify management, increase availability and automation, and reduce the total amount of storage that you need.

- **Outstanding performance.** The NetApp AFF solution shares the same unified storage architecture, ONTAP software, management interface, rich data services, and advanced feature set as the rest of the NetApp FAS product families. This innovative combination of all-flash media with ONTAP gives you the consistent low latency and high IOPS of all-flash storage with industry- leading ONTAP software.
- **Storage efficiency.** You can reduce your total capacity requirements work with your NetApp SME to understand how this applied your specific medical imaging system.

- **Space-efficient cloning.** With the FlexClone capability, your system can almost instantly create clones to support backup and testing environment refresh. These clones consume additional storage only as changes are made.
- **Integrated data protection.** Full data protection and disaster recovery features help you protect your critical data assets and provide disaster recovery.
- **Nondisruptive operations.** You can perform upgrades and maintenance without taking data offline.
- **QoS.** Storage QoS helps you limit potential bulky workloads. More importantly, QoS creates a minimum performance guarantee that your system performance will not drop below a certain level for critical workloads such as a medical imaging system's production environment. And by limiting contention, NetApp QoS can also reduce performance-related issues.
- **Data fabric.** To accelerate digital transformation, your data fabric delivered by NetApp simplifies and integrates data management across cloud and on-premises environments. It delivers consistent and integrated data management services and applications for superior data visibility and insights, data access and control, and data protection and security. NetApp is integrated with large public clouds, such as AWS, Azure, Google Cloud, and IBM Cloud, giving you a wide breadth of choice.

Host virtualization — VMware vSphere

FlexPod architectures are validated with VMware vSphere 6.x, which is the industry- leading virtualization platform. VMware ESXi 6.x is used to deploy and run the VMs. vCenter Server Appliance 6.x is used to manage the ESXi hosts and VMs. Multiple ESXi hosts that run on Cisco UCS B200 M5 blades are used to form a VMware ESXi cluster. The VMware ESXi cluster pools the compute, memory, and network resources from all the cluster nodes and provides a resilient platform for the VMs that are running on the cluster. The VMware ESXi cluster features, vSphere high availability, and Distributed Resource Scheduler (DRS) all contribute to the vSphere cluster's tolerance to withstand failures, and they help distribute the resources across the VMware ESXi hosts.

The NetApp storage plug-in and the Cisco UCS plug-in integrate with VMware vCenter to enable operational workflows for your required storage and compute resources.

The VMware ESXi cluster and vCenter Server give you a centralized platform for deploying medical imaging environments in VMs. Your healthcare organization can realize all the benefits of an industry- leading virtual infrastructure with confidence, such as the following:


- **Simple deployment.** Quickly and easily deploy vCenter Server by using a virtual appliance.
- **Centralized control and visibility.** Administer the entire vSphere infrastructure from a single location.
- **Proactive optimization.** Allocate, optimize, and migrate resources for maximum efficiency.
- **Management.** Use powerful plug-ins and tools to simplify management and to extend control.

Architecture

The FlexPod architecture is designed to provide high availability if a component or a link fails in your entire compute, network, and storage stack. Multiple network paths for client access and storage access provide load balancing and optimal resource utilization.

The following figure illustrates the 16Gb FC/40Gb Ethernet (40GbE) topology for the medical imaging system solution deployment.

FlexPod Infrastructure for an Enterprise Medical Imaging System



Storage architecture

Use the storage architecture guidelines in this section to configure your storage infrastructure for an enterprise medical imaging system.

Storage tiers

A typical enterprise medical imaging environment consists of several different storage tiers. Each tier has specific performance and storage protocol requirements. NetApp storage supports various RAID technologies; more information can be found [here](#). Here is how NetApp AFF storage systems serve the needs of different storage tiers for the imaging system:

- **Performance Storage (tier 1).** This tier offers high performance and high redundancy for databases, OS drives, VMware Virtual Machine File System (VMFS) datastores, and so on. Block I/O moves over fiber to a shared storage array of SSD, as is configured in ONTAP. The minimum latency is 1ms to 3ms, with an occasional peak of 5ms. This storage tier is typically used for short-term storage cache, typically 6 to 12 months of image storage for fast access to online DICOM images. This tier offers high performance and high redundancy for image caches, database backup, and so on. NetApp all-flash arrays provide <1ms latency at a sustained bandwidth, which is far lower than the service times that are expected by a typical enterprise medical imaging environment. NetApp ONTAP supports both RAID-TEC (triple parity RAID to sustain three disk failures) and RAID DP (double-parity RAID to sustain two disk failures).
- **Archive storage (tier 2).** This tier is used for typical cost-optimized file access, RAID 5 or RAID 6 storage for larger volumes, and long-term lower-cost/performance archiving. NetApp ONTAP supports both RAID-

TEC (triple parity RAID to sustain three disk failures) and RAID DP (double-parity RAID to sustain two disk failures). NetApp FAS in FlexPod enables imaging application I/O over NFS/SMB to a SAS disk array. NetApp FAS systems provide ~10ms latency at sustained bandwidth, which is far lower than the service times that are expected for storage tier 2 in an enterprise medical imaging system environment.

Cloud-based archiving in a hybrid-cloud environment can be used for archiving to a public cloud storage provider using S3 or similar protocols. NetApp SnapMirror technology enables replication of imaging data from all-flash or FAS arrays to slower disk-based storage arrays or to Cloud Volumes ONTAP for AWS, Azure, or Google Cloud.

NetApp SnapMirror provides industry leading data replication capabilities that help protect your medical imaging system with unified data replication. Simplify data-protection management across the data fabric with cross-platform replication—from flash to disk to cloud:

- Transport data seamlessly and efficiently between NetApp storage systems to support both backup and disaster recovery with the same target volume and I/O stream.
- Failover to any secondary volume. Recover from any point-in-time Snapshot on the secondary storage.
- Safeguard your most critical workloads with available zero-data-loss synchronous replication (RPO=0).
- Cut network traffic. Shrink your storage footprint through efficient operations.
- Reduce network traffic by transporting only changed data blocks.
- Preserve storage-efficiency benefits on the primary storage during transport—including deduplication, compression, and compaction.
- Deliver additional inline efficiencies with network compression.

More information can be found [here](#).

The table below lists each tier that a typical medical imaging system requires for specific latency and the throughput performance characteristics.

| Storage tier | Requirements | NetApp recommendation |
|--------------|--|---|
| 1 | 1–5ms latency 35–500MBps throughput | AFF with <1ms latency AFF A300 high-availability (HA) pair with two disk shelves can handle throughput of up to ~1.6GBps |
| 2 | On premises archive | FAS with up to 30ms latency |
| | Archive to cloud | SnapMirror replication to Cloud Volumes ONTAP or backup archiving with NetApp StorageGRID software |

Storage network connectivity

FC fabric

- The FC fabric is for host OS I/O from compute to storage.
- Two FC fabrics (Fabric A and Fabric B) are connected to Cisco UCS Fabric A and UCS Fabric B, respectively.
- A storage virtual machine (SVM) with two FC logical interfaces (LIFs) is on each controller node. On each node, one LIF is connected to Fabric A and the other is connected to Fabric B.

- 16Gbps FC end-to-end connectivity is through Cisco MDS switches. A single initiator, multiple target ports, and zoning are all configured.
- FC SAN boot is used to create fully stateless computing. Servers are booted from LUNs in the boot volume that is hosted on the AFF storage cluster.

IP network for storage access over iSCSI, NFS, and SMB/CIFS

- Two iSCSI LIFs are in the SVM on each controller node. On each node, one LIF is connected to Fabric A and the second is connected to Fabric B.
- Two NAS data LIFs are in the SVM on each controller node. On each node, one LIF is connected to Fabric A and the second is connected to Fabric B.
- Storage port interface groups (virtual port channel [vPC]) for 10Gbps link to switch N9k-A and for 10Gbps link to switch N9k-B.
- Workload in Ext4 or NTFS file systems from VM to storage:
 - iSCSI protocol over IP.
- VMs hosted in NFS datastore:
 - VM OS I/O goes over multiple Ethernet paths through Nexus switches.

In-band management (active-passive bond)

- 1Gbps link to management switch N9k-A, and 1Gbps link to management switch N9k-B.

Backup and recovery

FlexPod Datacenter is built on a storage array that is managed by NetApp ONTAP data management software. ONTAP software has evolved over 20 years to provide many data management features for VMs, Oracle databases, SMB/CIFS file shares, and NFS. It also provides protection technology such as NetApp Snapshot technology, SnapMirror technology, and NetApp FlexClone data replication technology. NetApp SnapCenter software has a server and a GUI client to use ONTAP Snapshot, SnapRestore, and FlexClone features for VM, SMB/CIFS file shares, NFS, and Oracle database backup and recovery.

NetApp SnapCenter software employs [patented](#) Snapshot technology to create a backup of an entire VM or Oracle database on a NetApp storage volume instantaneously. Compared with Oracle Recovery Manager (RMAN), Snapshot copies do not require a full baseline backup copy, because they are not stored as physical copies of blocks. Snapshot copies are stored as pointers to the storage blocks as they existed in the ONTAP WAFL file system when the Snapshot copies were created. Because of this tight physical relationship, the Snapshot copies are maintained on the same storage array as the original data. Snapshot copies can also be created at the file level to give you more granular control for the backup.

Snapshot technology is based on a redirect-on-write technique. It initially contains only metadata pointers and does not consume much space until the first data change to a storage block. If an existing block is locked by a Snapshot copy, a new block is written by the ONTAP WAFL file system as an active copy. This approach avoids the double- writes that occur with the change-on-write technique.

For Oracle database backup, Snapshot copies yield incredible time savings. For example, a backup that took 26 hours to complete by using RMAN alone can take less than 2 minutes to complete by using SnapCenter software.

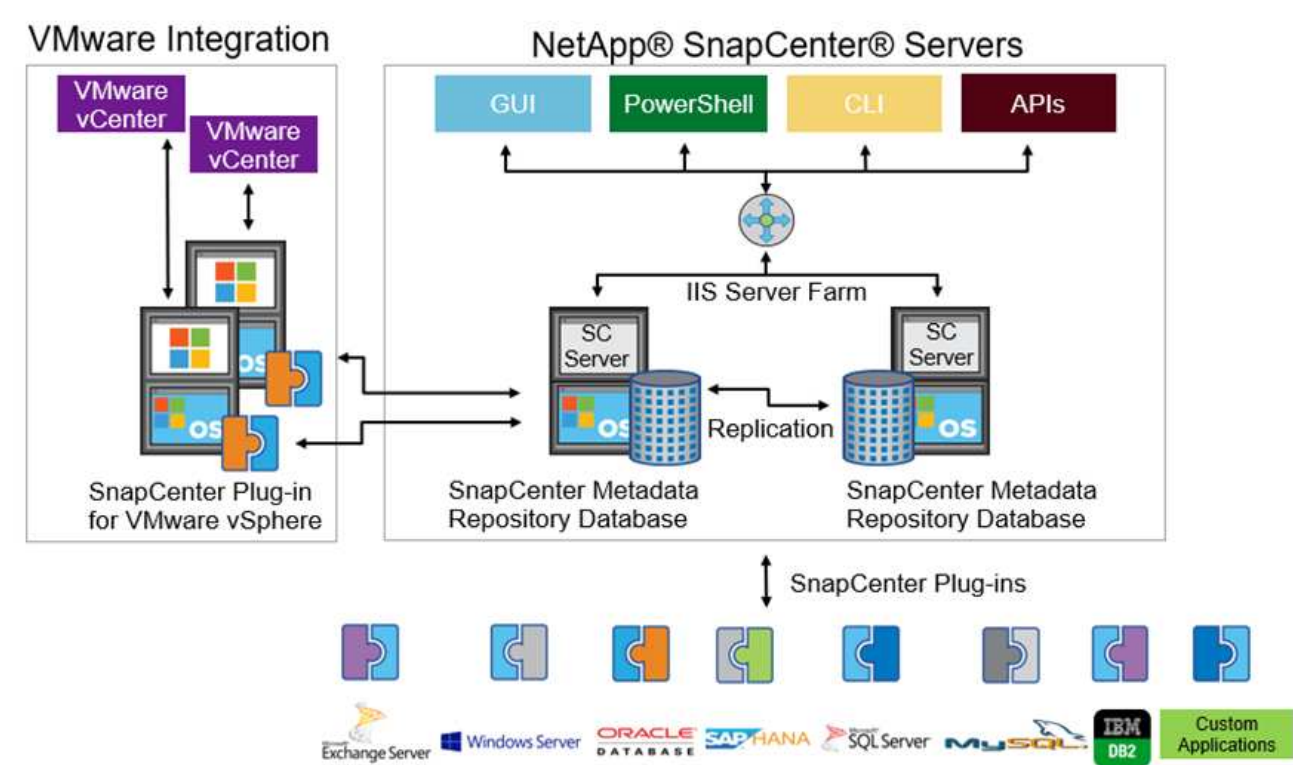
And because data restoration does not copy any data blocks but instead flips the pointers to the application-consistent Snapshot block images when the Snapshot copy was created, a Snapshot backup copy can be restored almost instantaneously. SnapCenter cloning creates a separate copy of metadata pointers to an

existing Snapshot copy and mounts the new copy to a target host. This process is also fast and storage efficient.

The following table summarizes the primary differences between Oracle RMAN and NetApp SnapCenter software.

| | Backup | Restore | Clone | Need Full Backup | Space usage | Off-site copy |
|------------|--------|---------|-------|------------------|-------------|---------------|
| RMAN | Slow | Slow | Slow | Yes | High | Yes |
| SnapCenter | Fast | Fast | Fast | No | Low | Yes |

The following figure presents the SnapCenter architecture.



NetApp MetroCluster configurations are used by thousands of enterprises worldwide for high availability (HA), zero data loss, and nondisruptive operations both within and beyond the data center. MetroCluster is a free feature of ONTAP software that synchronously mirrors data and configuration between two ONTAP clusters in separate locations or failure domains. MetroCluster provides continuously available storage for applications by automatically handling two objectives: Zero recovery point objective (RPO) by synchronously mirroring data written to the cluster. Near zero recovery time objective (RTO) by mirroring configuration and automating access to data at the second site MetroCluster provides simplicity with automatic mirroring of data and configuration between the two independent clusters located in the two sites. As storage is provisioned within one cluster, it is automatically mirrored to the second cluster at the second site. NetApp SyncMirror technology provides a complete copy of all data with a zero RPO. , Therefore, workloads from one site can switch over at any time to the opposite site and continue serving data without data loss. More information can be found [here](#).

Networking

A pair of Cisco Nexus switches provides redundant paths for IP traffic from compute to storage, and for

external clients of the medical imaging system image viewer:

- Link aggregation that uses port channels and vPCs is employed throughout, enabling the design for higher bandwidth and high availability:
 - vPC is used between the NetApp storage array and the Cisco Nexus switches.
 - vPC is used between the Cisco UCS fabric interconnect and the Cisco Nexus switches.
 - Each server has virtual network interface cards (vNICs) with redundant connectivity to the unified fabric. NIC failover is used between fabric interconnects for redundancy.
 - Each server has virtual host bus adapters (vHBAs) with redundant connectivity to the unified fabric.
- The Cisco UCS fabric interconnects are configured in end-host mode as recommended, providing dynamic pinning of vNICs to uplink switches.
- An FC storage network is provided by a pair of Cisco MDS switches.

Compute—Cisco Unified Computing System

Two Cisco UCS fabrics through different fabric interconnects provide two failure domains. Each fabric is connected to both IP networking switches and to different FC networking switches.

Identical service profiles for each Cisco UCS blade are created as per FlexPod best practices to run VMware ESXi. Each service profile should have the following components:

- Two vNICs (one on each fabric) to carry NFS, SMB/CIFS, and client or management traffic
- Additional required VLANs to the vNICs for NFS, SMB/CIFS, and client or management traffic
- Two vNICs (one on each fabric) to carry iSCSI traffic
- Two storage FC HBAs (one on each fabric) for FC traffic to storage
- SAN boot

Virtualization

The VMware ESXi host cluster runs workload VMs. The cluster comprises ESXi instances running on Cisco UCS blade servers.

Each ESXi host includes the following network components:

- SAN boot over FC or iSCSI
- Boot LUNs on NetApp storage (in a dedicated FlexVol for boot OS)
- Two VMNICs (Cisco UCS vNIC) for NFS, SMB/CIFS, or management traffic
- Two storage HBAs (Cisco UCS FC vHBA) for FC traffic to storage
- Standard switch or distributed virtual switch (as needed)
- NFS datastore for workload VMs
- Management, client traffic network, and storage network port groups for VMs
- Network adapter for management, client traffic, and storage access (NFS, iSCSI, or SMB/CIFS) for each VM
- VMware DRS enabled
- Native multipathing enabled for FC or iSCSI paths to storage

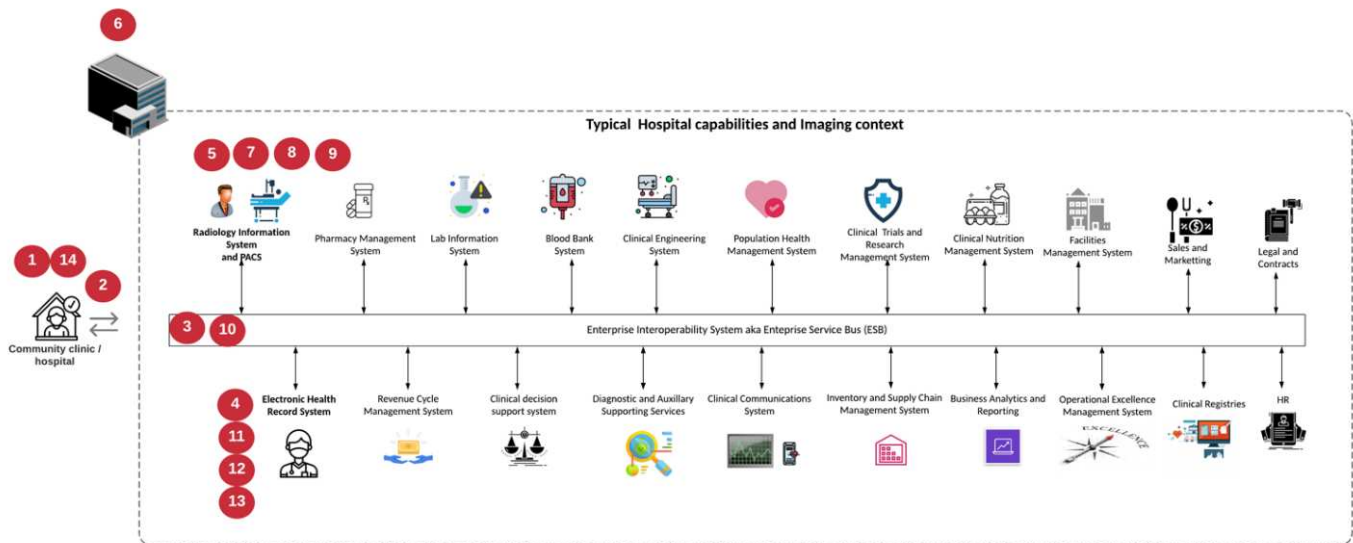
- VMware snapshots for VM turned off
- NetApp SnapCenter deployed for VMware for VM backups

Medical imaging system architecture

In healthcare organizations, medical imaging systems are critical applications and well-integrated into the clinical workflows that begin from patient registration and end with billing related activities in the revenue cycle.

The following diagram shows the various systems involved in a typical large hospital; this diagram is intended to provide architectural context to a medical imaging system before we zoom into the architectural components of a typical medical imaging system. Workflows vary widely and are hospital and use- case specific.

The figure below shows the medical imaging system in the context of a patient, a community clinic, and a large hospital.



1. The patient visits the community clinic with symptoms. During the consultation, the community physician places an imaging order that is sent to the larger hospital in the form of a HL7 order message.
2. The community physician's EHR system sends the HL7 order/ORD message to the large hospital.
3. The enterprise interoperability system (also known as the Enterprise Service Bus [ESB]) processes the order message and sends the order message to the EHR system.
4. The EHR processes the order message. If a patient record does not exist, a new patient record is created.
5. The EHR sends an imaging order to the medical imaging system.
6. The patient calls the large hospital for an imaging appointment.
7. The imaging reception and registration desk schedules patient for an imaging appointment using a radiology information or similar system.
8. The patient arrives for the imaging appointment, and the images or video is created and sent to the PACS.
9. The radiologist reads the images and annotates the images in the PACS using a high-end/GPU graphics-enabled diagnostic viewer. Certain imaging systems have artificial intelligence (AI)- enabled efficiency improvement capabilities built into the imaging workflows.
10. Image order results are sent to the EHR in the form of an order results HL7 ORU message via the ESB.
11. The EHR processes the order results into the patient's record, places thumbnail image with a context-aware link to the actual DICOM image. Physicians can launch the diagnostic viewer if a higher resolution

image is needed from within the EHR.

12. The physician reviews the image and enters physician notes into the patient's record. The physician could use the clinical decision support system to enhance the review process and aid in proper diagnosis for the patient.
13. The EHR system then sends the order results in the form of an order results message to the community hospital. At this point, if the community hospital could receive the complete image, then the image is sent either via WADO or DICOM.
14. The community physician completes the diagnosis and provides next steps to the patient.

A typical medical imaging system uses an N- tiered architecture. The core component of a medical imaging system is an application server to host various application components. Typical application servers are either Java runtime- based or C# .Net CLR- based. Most enterprise medical imaging solutions use an Oracle database Server or MS SQL Server or Sybase as the primary database. Additionally, some enterprise medical imaging systems also use databases for content acceleration and caching over a geographic region. Some enterprise medical imaging systems also use NoSQL databases like MongoDB, Redis, and so on in conjunction with enterprise integration servers for DICOM interfaces and or APIs.

A typical medical imaging system provides access to images for two distinct set of users: diagnostic user/radiologist, or the clinician or physician that ordered the imaging.

Radiologists typically use high- end, graphics- enabled diagnostic viewers that are running on high- end compute and graphics workstations that are either physical or part of a virtual desktop infrastructure. If you are about to start your virtual desktop infrastructure journey, more information can be found [here](#).

When hurricane Katrina destroyed two of Louisiana's major teaching hospitals, leaders came together and built a resilient electronic health record system that included over 3000 virtual desktops in record time. More information on use cases reference architecture and FlexPod reference bundles can be found [here](#).

Clinicians access images in two primary ways:

- **Web- based access.** which is typically used by EHR systems to embed PACS images as context- aware links into the electronic medical record (EMR) of the patient, and links that can be placed into imaging workflows, procedure workflows, progress notes workflows, and so on. Web based links are also use to provide image access to the patients via patient portals. Web based access uses a technology pattern called context aware links. Context aware links can either be static links/URLs to the DICOM media directly or dynamically generated links/URLs using custom macros.
- **Thick client.** Some enterprise medical systems also allow you to use a thick- client- based approach to view the images. You can launch a thick client from within the EMR of the patient or as a standalone application.


The medical imaging system can provide image access to a community of physicians or to CIN-participating physicians. Typical medical imaging systems include components that enable image interoperability with other health IT systems within and outside of your healthcare organization. Community physicians can either access images via a web-based application or leverage an image exchange platform for image interoperability. Image-exchange platforms typically use either WADO or DICOM as the underlying image exchange protocol.

Medical imaging systems can also support academic medical centers that need PACS or imaging systems for use in a classroom. To support academic activities, a typical medical imaging system can have the capabilities of a PACS system in a smaller footprint or a teaching- only imaging environment. Typical vendor- neutral archiving systems and some enterprise- class medical imaging systems offer DICOM image tag morphing capabilities to anonymize the images that are used for teaching purposes. Tag morphing enables healthcare organization to exchange DICOM images between different vendor medical imaging systems in a vendor- neutral fashion. Also, tag morphing enables medical imaging systems to implement an enterprise- wide,


vendor- neutral archival capability for medical images.

Medical imaging systems are starting to use GPU-based compute capabilities to enhance human workflows by preprocessing the images and thus improving efficiencies. Typical enterprise medical imaging systems take advantage of industry- leading NetApp storage efficiency capabilities. Enterprise medical imaging systems typically use RMAN for backup, recovery, and restore activities. For better performance and to reduce the time that it takes to create backups, Snapshot technology is available for backup operations and SnapMirror technology is available for replication.

The figure below shows the logical application components in a layered architectural view.



The figure below shows the physical application components.



The logical application components require that the infrastructure support a diverse set of protocols and file systems. NetApp ONTAP software supports an industry-leading set of protocols and file systems.

The table below lists the application components, storage protocol, and file system requirements.

| Application component | SAN/NAS | File system type | Storage tier | Replication type |
|-----------------------|---------------------|------------------------------|------------------------------|------------------|
| VMware host prod DB | local | SAN | VMFS | Tier 1 |
| Application | VMware host prod DB | REP | SAN | VMFS |
| Tier 1 | Application | VMware host prod application | local | SAN |
| VMFS | Tier 1 | Application | VMware host prod application | REP |

| Application component | SAN/NAS | File system type | Storage tier | Replication type |
|-----------------------|----------|------------------|--------------|----------------------------------|
| SAN | VMFS | Tier 1 | Application | Core database server |
| SAN | Ext4 | Tier 1 | Application | Backup database server |
| SAN | Ext4 | Tier 1 | None | Image cache server |
| NAS | SMB/CIFS | Tier 1 | None | Archive server |
| NAS | SMB/CIFS | Tier 2 | Application | Web server |
| NAS | SMB/CIFS | Tier 1 | None | WADO Server |
| SAN | NFS | Tier 1 | Application | Business intelligence server |
| SAN | NTFS | Tier 1 | Application | Business intelligence backup |
| SAN | NTFS | Tier 1 | Application | Interoperability server |
| SAN | Ext4 | Tier 1 | Application | Interoperability database server |

Solution infrastructure hardware and software components

The following tables list the hardware and software components, respectively, of the FlexPod infrastructure for the medical imaging system.

| Layer | Product family | Quantity and model | Details |
|---------|--|--|---|
| Compute | Cisco UCS 5108 chassis | 1 or 2 | Based on the number of blades required to support the number of annual studies |
| | Cisco UCS blade servers | B200 M5 | Number of blades based on the number of studies annually Each with 2 x 20 or more cores, 2.7GHz, and 128-384GB RAM |
| | Cisco UCS Virtual Interface Card (VIC) | Cisco UCS 1440 | See the |
| | 2 x Cisco UCS fabric interconnects | 6454 or later | — |
| Network | Cisco Nexus switches | 2 x Cisco Nexus 3000 Series or 9000 Series | — |

| Layer | Product family | Quantity and model | Details |
|-----------------|--|--|--------------------------------|
| Storage network | IP network for storage access over SMB/CIFS, NFS, or iSCSI protocols | Same network switches as above | — |
| | Storage access over FC | 2 x Cisco MDS 9132T | — |
| Storage | NetApp AFF A400 all-flash storage system | 1 or more HA pair | Cluster with two or more nodes |
| | Disk shelf | 1 or more DS224C or NS224 disk shelves | Fully populated with 24 drives |
| | SSD | >24, 1.2TB or larger capacity | — |

| Software | Product family | Version or release | Details |
|-----------------------------------|--|---|---------|
| Enterprise medical imaging system | | | |
| | MS SQL or Oracle Database Server | As suggested by the medical imaging system vendor | |
| | No SQL DBs like MongoDB Server | As suggested by the medical imaging system vendor | |
| | Application Servers | As suggested by the medical imaging system vendor | |
| | Integration Server (MS Biztalk, MuleSoft, Rhapsody, Tibco) | As suggested by the medical imaging system vendor | |
| | VMs | Linux (64 bit) | |
| | VMs | Windows Server (64 bit) | |
| Storage | ONTAP | ONTAP 9.7 or later | |
| Network | Cisco UCS Fabric Interconnect | Cisco UCS Manager 4.1 or later | |
| | Cisco Ethernet switches | 9.2(3)I7(2) or later | |
| | Cisco FC: Cisco MDS 9132T | 8.4(2) or later | |
| Hypervisor | Hypervisor | VMware vSphere ESXi 6.7 U2 or later | |
| Management | Hypervisor management system | VMware vCenter Server 6.7 U1 (vCSA) or later | |
| | NetApp Virtual Storage Console (VSC) | VSC 9.7 or later | |
| | SnapCenter | SnapCenter 4.3 or later | |

Solution sizing

Storage sizing

This section describes the number of studies and the corresponding infrastructure requirements.

The storage requirements that are listed in the following table assume that existing data is 1 year's worth plus projected growth for 1 year of study in the primary system (tier 1, 2). Additional storage needs for projected growth for 3 years beyond the first 2 years are listed separately.

| | Small | Medium | Large |
|--|---------------|-------------------|------------------------|
| Annual studies | <250K studies | 250K–500K studies | 500K–1 million studies |
| Tier 1 Storage | | | |
| IOPS (average) | 1.5K–5K | 5K–15K | 15K–40K |
| IOPS (peak) | 5K | 20K | 65K |
| Throughput | 50–100MBps | 50–150MBps | 100–300MBps |
| Capacity data center 1 (1 year of old data and 1 year of new study) | 70TB | 140TB | 260TB |
| Capacity data center 1 (additional need for 4 years for new study) | 25TB | 45TB | 80TB |
| Capacity data center 2 (1 year of old data and 1 year of new study) | 45TB | 110TB | 165TB |
| Capacity data center 2 (additional need for 4 years for new study) | 25TB | 45TB | 80TB |
| Tier 2 Storage | | | |
| IOPS (average) | 1K | 2K | 3K |
| Capacity data center 1 | 320TB | 800TB | 2000TB |

Compute sizing

The table below lists the compute requirements for small, medium, and large medical imaging systems.

| | Small | Medium | Large |
|--------------------------------|---------------|-------------------|------------------------|
| Annual studies | <250K studies | 250K–500K studies | 500K–1 million studies |
| Data Center 1 | | | |
| Number of VMs | 21 | 27 | 35 |
| Total virtual CPU (vCPU) count | 56 | 124 | 220 |
| Total memory requirement | 225GB | 450GB | 900GB |

| | Small | Medium | Large |
|--|--|--|---|
| Physical server (blades) specs (assume 1 vCPU = 1 core) | 4 x servers with 20 cores and 192GB RAM each | 8 x servers with 20 cores and 128GB RAM each | 14 x servers with 20 cores and 128GB RAM each |
| Data Center 2 | | | |
| Number of VMs | 15 | 17 | 22 |
| Total vCPU count | 42 | 72 | 140 |
| Total memory requirement | 179GB | 243GB | 513GB |
| Physical server (blades) specs (assume 1 vCPU = 1 core) | 3 x servers with 20 cores and 168GB RAM each | 6 x servers with 20 cores and 128GB RAM each | 8 x servers with 24 cores and 128GB RAM each |

Networking and Cisco UCS infrastructure sizing

The table below lists the networking and Cisco UCS infrastructure requirements for small, medium, and large medical imaging systems.

| | Small | Medium | Large |
|--|--|----------------|----------------|
| Data Center 1 | | | |
| Number of storage node ports | 2 converged network adapters (CNAs); 2 FCs | 2 CNAs; 2 FCs | 2 CNAs; 2 FCs |
| IP network switch ports (Cisco Nexus 9000) | 48-port switch | 48-port switch | 48-port switch |
| FC switch (Cisco MDS) | 32-port switch | 32-port switch | 48-port switch |
| Cisco UCS chassis count | 1 x 5108 | 1 x 5108 | 2 x 5108 |
| Cisco UCS Fabric Interconnect | 2 x 6332 | 2 x 6332 | 2 x 6332 |
| Data Center 2 | | | |
| Cisco UCS chassis count | 1 x 5108 | 1 x 5108 | 1 x 5108 |
| Cisco UCS Fabric Interconnect | 2 x 6332 | 2 x 6332 | 2 x 6332 |
| Number of storage node ports | 2 CNAs; 2 FCs | 2 CNAs; 2 FCs | 2 CNAs; 2 FCs |
| IP network switch ports (Cisco Nexus 9000) | 48-port switch | 48-port switch | 48-port switch |
| FC switch (Cisco MDS) | 32-port switch | 32-port switch | 48-port switch |

Best practices

Storage best practices

High availability

The NetApp storage cluster design provides high availability at every level:

- Cluster nodes
- Back-end storage connectivity
- RAID TEC that can sustain three disk failures
- RAID DP that can sustain two disk failures
- Physical connectivity to two physical networks from each node
- Multiple data paths to storage LUNs and volumes

Secure multitenancy

NetApp storage virtual machines (SVMs) provide a virtual storage array construct to separate your security domain, policies, and virtual networking. NetApp recommends that you create separate SVMs for each tenant organization that hosts data on the storage cluster.

NetApp storage best practices



Consider the following NetApp storage best practices:

- Always enable NetApp AutoSupport technology, which sends support summary information to NetApp through HTTPS.
- For maximum availability and mobility, make sure that a LIF is created for each SVM on each node in the NetApp ONTAP cluster. Asymmetric logical unit access (ALUA) is used to parse paths and to identify active optimized (direct) paths versus active nonoptimized paths. ALUA is used for both FC or FCoE and iSCSI.
- A volume that contains only LUNs does not need to be internally mounted, nor is a junction path required.
- If you use the Challenge-Handshake Authentication Protocol (CHAP) in ESXi for target authentication, you must also configure it in ONTAP. Use the CLI (`vserver iscsi security create`) or NetApp ONTAP System Manager (edit Initiator Security under Storage > SVMs > SVM Settings > Protocols > iSCSI).

SAN boot

NetApp recommends that you implement SAN boot for Cisco UCS Servers in the FlexPod Datacenter solution. This step enables the operating system to be safely secured by the NetApp AFF storage system, providing better performance. The design that is outlined in this solution uses iSCSI SAN boot.

In iSCSI SAN boot, each Cisco UCS Server is assigned two iSCSI vNICs (one for each SAN fabric), which provide redundant connectivity all the way to the storage. The storage ports in this example, e2a and e2e, which are connected to the Cisco Nexus switches, are grouped together to form one logical port called an interface group (ifgrp) (in this example, a0a). The iSCSI VLANs are created on the ifgroup, and the iSCSI LIFs are created on iSCSI port groups (in this example, a0a-`iSCSI-A-VLAN`). The iSCSI boot LUN is exposed to the servers through the iSCSI LIF by using ifgroups. This approach enables only the authorized server to have access to the boot LUN. For the port and LIF layout, see the figure below.



Unlike NAS network interfaces, the SAN network interfaces are not configured to fail over during a failure. Instead, if a network interface becomes unavailable, the host chooses a new optimized path to an available network interface. ALUA, a standard supported by NetApp, provides information about SCSI targets, which enables a host to identify the best path to the storage.

Storage efficiency and thin provisioning

NetApp has led the industry in storage efficiency innovation, such as with the first deduplication for primary workloads and with inline data compaction, which enhances compression and stores small files and I/Os efficiently. ONTAP supports both inline and background deduplication, as well as inline and background compression.

To realize the benefits of deduplication in a block environment, the LUNs must be thin-provisioned. Although the LUN is still seen by your VM administrator as taking the provisioned capacity, the deduplication savings are returned to the volume to be used for other needs. NetApp recommends that you deploy these LUNs in FlexVol volumes that are also thin-provisioned with a capacity that is two times the size of the LUN. When you deploy the LUN that way, the FlexVol volume acts merely as a quota. The storage that the LUN consumes is reported in the FlexVol volume and its containing aggregate.

For maximum deduplication savings, consider scheduling background deduplication. These processes use system resources when they're running, however. So, ideally, you should schedule them during less active times (such as weekends) or run them more frequently to reduce the amount of changed data to be processed. Automatic background deduplication on AFF systems has much less of an effect on foreground activities. Background compression (for hard disk-based systems) also consumes resources, so you should consider it only for secondary workloads with limited performance requirements.

Quality of service

Systems that run ONTAP software can use the ONTAP storage QoS feature to limit throughput in megabits per second (MBps) and to limit IOPS for different storage objects such as files, LUNs, volumes, or entire SVMs. Adaptive QoS is used to set an IOPS floor (QoS minimum) and ceiling (QoS maximum), which dynamically adjust based on the datastore capacity and used space.

Throughput limits are useful for controlling unknown or test workloads before a deployment to confirm that they don't affect other workloads. You might also use these limits to constrain a bully workload after it has been identified. Minimum levels of service based on IOPS are also supported to provide consistent performance for SAN objects in ONTAP.

With an NFS datastore, a QoS policy can be applied to the entire FlexVol volume or to individual Virtual Machine Disk (VMDK) files within it. With VMFS datastores (Cluster Shared Volumes [CSV] in Hyper-V) that use ONTAP LUNs, you can apply the QoS policies to the FlexVol volume that contains the LUNs or to the individual LUNs. However, because ONTAP has no awareness of the VMFS, you cannot apply the QoS policies to individual VMDK files. When you use VMware Virtual Volumes (VVols) with VSC 7.1 or later, you can set maximum QoS on individual VMs by using the storage capability profile.

To assign a QoS policy to a LUN, including VMFS or CSV, you can obtain the ONTAP SVM (displayed as `Vserver`), LUN path, and serial number from the Storage Systems menu on the VSC home page. Select the storage system (SVM), then Related Objects > SAN. Use this approach when you specify QoS by using one of the ONTAP tools.

You can set the QoS maximum throughput limit on an object in MBps and in IOPS. If you use both, the first limit that is reached is enforced by ONTAP. A workload can contain multiple objects, and a QoS policy can be applied to one or more workloads. When you apply a policy to multiple workloads, the workloads share the total limit of the policy. Nested objects are not supported (for example, for a file within a volume, they cannot each have their own policy). QoS minimums can be set only in IOPS.

Storage layout

This section provides best practices for layout of LUNs, volumes, and aggregates on storage.

Storage LUNs

For optimal performance, management, and backup, NetApp recommends the following LUN-design best practices:

- Create a separate LUN to store database data and log files.
- Create a separate LUN for each instance to store Oracle database log backups. The LUNs can be part of the same volume.
- Provision LUNs with thin provisioning (disable the Space Reservation option) for database files and log files.
- All imaging data is hosted in FC LUNs. Create these LUNs in FlexVol volumes that are spread across the aggregates that are owned by different storage controller nodes.

For placement of the LUNs in a storage volume, follow the guidelines in the next section.

Storage volumes

For optimal performance and management NetApp recommends the following volume design best practices:

- Isolate databases with I/O-intensive queries on separate storage volumes.
- The datafiles could be placed on a single LUN or a volume, but multiple volumes/LUNs are recommended for higher throughput.
- I/O parallelism can be attained by using any supported filesystem when multiple LUNs are used.
- Place database files and transaction logs on separate volumes to increase the recovery granularity.
- Consider using volume attributes like auto size, Snapshot reserve, QoS, and so on.

Aggregates

Aggregates are the primary storage containers for NetApp storage configurations and contain one or more RAID groups that consist of both data disks and parity disks.

NetApp performed various I/O workload characterization tests by using shared and dedicated aggregates with data files and transaction log files separated. The tests show that one large aggregate with more RAID groups and drives (HDDs or SSDs) optimizes and improves storage performance and is easier for administrators to manage for two reasons:

- One large aggregate makes the I/O abilities of all drives available to all files.
- One large aggregate enables the most efficient use of disk space.

For effective disaster recovery, NetApp recommends that you place the asynchronous replica on an aggregate that is part of a separate storage cluster in your disaster recovery site and use SnapMirror technology to replicate content.

For optimal storage performance, NetApp recommends that you have at least 10% free space available in an aggregate.

Storage aggregate layout guidance for AFF A300 systems (with two disk shelves with 24 drives) includes:

- Keep two spare drives.
- Use Advanced Disk Partitioning to create three partitions on each drive: root and data.
- Use a total of 20 data partitions and two parity partitions for each aggregate.

Backup best practices

NetApp SnapCenter is used for VM and database backups. NetApp recommends the following backup best practices:

- When SnapCenter is deployed to create Snapshot copies for backups, turn off the Snapshot schedule for the FlexVol that host VMs and application data.
- Create a dedicated FlexVol for host boot LUNs.
- Use a similar or a single backup policy for VMs that serve the same purpose.
- Use a similar or a single backup policy per workload type; for example, use a similar policy for all database workloads. Use different policies for databases, web servers, end-user virtual desktops, and so on.
- Enable verification of the backup in SnapCenter.
- Configure archiving of the backup Snapshot copies to the NetApp SnapVault backup solution.
- Configure retention of the backups on primary storage based on the archiving schedule.

Infrastructure best practices

Networking best practices

NetApp recommends the following networking best practices:

- Make sure that your system includes redundant physical NICs for production and storage traffic.
- Separate VLANs for iSCSI, NFS, and SMB/CIFS traffic between compute and storage.
- Make sure that your system includes a dedicated VLAN for client access to the medical imaging system.

You can find additional networking best practices in the FlexPod infrastructure design and deployment guides.

Compute best practices

NetApp recommends the following compute best practice:

- Make sure that each specified vCPU is supported by a physical core.

Virtualization best practices

NetApp recommends the following virtualization best practices:

- Use VMware vSphere 6 or later.
- Set the ESXi host server BIOS and OS layer to Custom Controlled–High Performance.
- Create backups during off-peak hours.

Medical imaging system best practices

See the following best practices and some requirements from a typical medical imaging system:

- Do not overcommit virtual memory.
- Make sure that the total number of vCPUs equals the number of physical CPUs.
- If you have a large environment, dedicated VLANs are required.
- Configure database VMs with dedicated HA clusters.
- Make sure that the VM OS VMDKs are hosted in fast tier 1 storage.
- Work with the medical imaging system vendor to identify the best approach to prepare VM templates for quick deployment and maintenance.
- Management, storage, and production networks require LAN segregation for the database, with isolated VLANs for VMware vMotion.
- Use the NetApp storage-array-based replication technology called [SnapMirror](#) instead of vSphere-based replication.
- Use backup technologies that leverage VMware APIs; backup windows should be outside the normal production hours.

Conclusion

By running a medical imaging environment on FlexPod, your healthcare organization can expect to see an improvement in staff productivity and a decrease in capital and operating expenses. FlexPod provides a prevalidated, rigorously tested converged infrastructure from the strategic partnership of Cisco and NetApp. It is engineered and designed specifically to deliver predictable low-latency system performance and high availability. This approach results in a superior user experience and optimal response time for users of the medical imaging system.

Different components of a medical imaging system require data storage in SMB/CIFS, NFS, Ext4, and NTFS file systems. Therefore, your infrastructure must provide data access over NFS, SMB/CIFS, and SAN protocols. NetApp storage systems support these protocols from a single storage array.

High availability, storage efficiency, Snapshot copy-based scheduled fast backups, fast restore operations, data replication for disaster recovery, and the FlexPod storage infrastructure capabilities all provide an industry-leading data storage and management system.

Additional information

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

- FlexPod Datacenter for AI/ML with Cisco UCS 480 ML for Deep Learning Design Guide

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_c480m5l_aiml_design.html

- FlexPod Datacenter Infrastructure with VMware vSphere 6.7 U1, Cisco UCS 4th Generation, and NetApp AFF A-Series

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_datacenter_vmware_netappaffa.html

- FlexPod Datacenter Oracle Database Backup with SnapCenter Solution Brief

<https://www.netapp.com/us/media/sb-3999.pdf>

- FlexPod Datacenter with Oracle RAC Databases on Cisco UCS and NetApp AFF A-Series

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_orc12cr2_affaseries.html

- FlexPod Datacenter with Oracle RAC on Oracle Linux

https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_orcrac_12c_bm.html

- FlexPod for Microsoft SQL Server

<https://flexpod.com/solutions/use-cases/microsoft-sql-server/>

- FlexPod from Cisco and NetApp

<https://flexpod.com/>

- [NetApp Solutions for MongoDB](#) Solution Brief (NetApp login required)

<https://fieldportal.netapp.com/content/734702>

- TR-4700: SnapCenter Plug-In for Oracle Database

<https://www.netapp.com/us/media/tr-4700.pdf>

- NetApp Product Documentation

<https://www.netapp.com/us/documentation/index.aspx>

- FlexPod for Virtual Desktop Infrastructure (VDI) Solutions

<https://flexpod.com/solutions/use-cases/virtual-desktop-infrastructure/>

Virtual Desktop Infrastructure

FlexPod Datacenter with Citrix Virtual Apps & Desktops 1912 LTSR and VMware vSphere 7 for up to 6000 seats

Jeff Nichols, Cisco
Suresh Thoppay, NetApp
Dre Jackson, NetApp

This document provides the architecture and design of a virtual desktop infrastructure for up to 6000 end-user computing users. The solution is virtualized on fifth-generation Cisco UCS B200 M5 blade servers, booting VMware vSphere 7.01 Update 1 through FC SAN from the AFF A400 storage array. The virtual desktops are powered using Citrix Provisioning Server 1912 LTSR and Citrix RDS/Citrix Virtual Apps & Desktops 1912 LTSR, with a mix of RDS-hosted shared desktops (6000), pooled and/or non-persistent-hosted virtual Windows 10 desktops (5000), and persistent-hosted virtual Windows 10 desktops provisioned with Citrix Machine Creation Services (5000) to support the user population. Where applicable, the document provides best-practice recommendations and sizing guidelines for customer deployments of this solution.

[FlexPod Datacenter with Citrix Virtual Apps & Desktops 1912 LTSR and VMware vSphere 7 for up to 6000 seats](#)

FlexPod Datacenter with VMware Horizon View 7.10, VMware vSphere 6.7 U2, Cisco UCS Manager 4.0, and NetApp ONTAP 9.6 for up to 6700 seats

Vadim Lebedev, Cisco
Suresh Thoppay, NetApp

This document provides a reference architecture and design guide for a 5000-seat to 6000-seat desktop workload, end-user computing environment on FlexPod Datacenter with Cisco UCS and NetApp AFF A300 and NetApp ONTAP data management software. The solution includes VMware Horizon server-based RDS Windows Server 2019 sessions, VMware Horizon persistent full clone Microsoft Windows 10 virtual desktops and VMware Horizon non-persistent, instant-clone Microsoft Windows 10 virtual desktops on VMware vSphere 6.7U2

[FlexPod Datacenter with VMware Horizon View 7.10, VMware vSphere 6.7 U2, Cisco UCS Manager 4.0, and NetApp ONTAP 9.6 for up to 6700 seats](#)

3D graphics visualization with Citrix and NVIDIA - White paper

This document describes the performance of Citrix XenDesktop on Citrix XenServer with

NVIDIA Tesla P4, P6, and P40 cards on Cisco UCS C240 M5 and B200 M5 servers with SPECviewperf 13.

[3D graphics visualization with Citrix and NVIDIA - White paper](#)

FlexPod Datacenter with Citrix XenDesktop/XenApp 7.15 and VMware vSphere 6.5 Update 1 for 6000 seats

Vadim Lebedev, Cisco
Chris Rodriguez, NetApp

This document provides a Reference Architecture for a virtual desktop and application design using Citrix XenApp/XenDesktop 7.15 built on Cisco UCS with a NetApp All Flash FAS (AFF) A300 storage and the VMware vSphere ESXi 6.5 hypervisor platform.

The landscape of desktop and application virtualization is changing constantly. The new M5 high-performance Cisco UCS Blade Servers and Cisco UCS Unified Fabric combined as part of the FlexPod proven Infrastructure, with the latest generation NetApp AFF storage result in a more compact, more powerful, more reliable and more efficient platform.

[FlexPod Datacenter with Citrix XenDesktop/XenApp 7.15 and VMware vSphere 6.5 Update 1 for 6000 seats](#)

FlexPod Datacenter with VMware Horizon View 7.3 and VMware vSphere 6.5 Update 1 with Cisco UCS Manager 3.2 for 5000 seats

Ramesh Guduru, Cisco
David Arnette, NetApp

This document provides a reference architecture, design guide, and deployment for up to a 5000-seat, mixed workload end-user computing environment on FlexPod Datacenter with Cisco UCS and NetApp All Flash FAS (AFF) A300 storage. The solution includes VMware Horizon server-based Remote Desktop Server Hosted sessions, VMware Horizon persistent Microsoft Windows 10 virtual desktops, and VMware Horizon non-persistent, Microsoft Windows 10 instant clone virtual desktops on VMware vSphere 6.5.

[FlexPod Datacenter with VMware Horizon View 7.3 and VMware vSphere 6.5 Update 1 with Cisco UCS Manager 3.2 for 5000 seats](#)

FlexPod Datacenter with VMware Horizon View 7.10, VMware vSphere 6.7 U2, Cisco UCS Manager 4.0, and NetApp ONTAP 9.6 for up to 6700 seats

Vadim Lebedev, Cisco
Suresh Thoppay, NetApp

This document provides a reference architecture and design guide for a 5000-seat to

6000-seat desktop workload end-user computing environment on FlexPod Datacenter with Cisco UCS and NetApp AFF A300 and NetApp ONTAP data management software. The solution includes VMware Horizon server-based RDS Windows Server 2019 sessions, VMware Horizon persistent, full clone Microsoft Windows 10 virtual desktops, and VMware Horizon non-persistent, instant-clone Microsoft Windows 10 virtual desktops on VMware vSphere 6.7 U2.

[FlexPod Datacenter with VMware Horizon View 7.10, VMware vSphere 6.7 U2, Cisco UCS Manager 4.0, and NetApp ONTAP 9.6 for up to 6700 seats](#)

Modern Apps

FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning - Design

Haseeb Niazi, Cisco
Arvind Ramakrishnan, NetApp

This document provides design details around the integration of the Cisco UCS C480 ML M5 platform into the FlexPod Datacenter solution to deliver a unified approach for providing AI and ML capabilities within the converged infrastructure. By providing customers the ability to manage the servers with combined AI and ML capabilities with the familiar tools they use to administer traditional FlexPod systems, the administrative overhead as well as the cost of deploying deep learning platform is greatly reduced. The design presented in this CVD also includes other Cisco UCS platforms such as the C220 M5 server with two NVIDIA T4 GPUs and the C240 M5 server equipped with two NVIDIA V100 32GB PCIe cards as additional options for handling concurrent AI and ML workloads.

[FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for Deep Learning - Design](#)

Deploy NetApp Trident CSI plug-in on Cisco Container Platform with FlexPod

This document provides step-by-step procedures for deploying the NetApp Trident Container Storage Interface (CSI) plug-in on a Cisco Container Platform Kubernetes tenant cluster in a FlexPod solution.

[Deploy NetApp Trident CSI plug-in on Cisco Container Platform with FlexPod](#)

FlexPod Datacenter for OpenShift Container Platform 4 - Deployment

Haseeb Niazi, Cisco
Alan Cowles, NetApp

Red Hat OpenShift is an enterprise-ready Kubernetes container platform to manage hybrid cloud and multi-cloud deployments. Red Hat OpenShift Container Platform includes everything needed for hybrid cloud, enterprise container, and Kubernetes development and deployments. It includes an enterprise-grade Linux operating system, container runtime, networking, monitoring, container registry, authentication, and authorization solutions.

Combining Red Hat OpenShift with the FlexPod Datacenter solution can simplify the deployment and the management of the container infrastructure. Customers can benefit from improved efficiency, better data protection, lower risk, and the flexibility to scale this highly available enterprise-grade infrastructure stack to

accommodate new business requirements. The pre-validated converged solution approach helps organizations achieve the speed, flexibility, and scale required for all of their application modernization and digital transformation initiatives.

[FlexPod Datacenter for OpenShift Container Platform 4 - Deployment](#)

FlexPod Datacenter with Docker Enterprise Edition for Container Management

Muhammad Afzal, Cisco
John George, Cisco
Amit Borulkar, NetApp
Uday Shetty, Docker

Docker is the world's leading software container platform for developers and IT operations to build, ship, and run distributed applications anywhere. With microservices architecture shaping the next generation of IT, enterprises with large investments in monolithic applications are finding ways to adopt Docker as a strategy for modernizing their application architectures and keeping the organization competitive and cost effective. Containerization provides the agility, control, and portability that developers and IT operations require to build and deploy applications across any infrastructure. The Docker platform allows distributed applications to be easily composed into a lightweight application container that can change dynamically yet non-disruptively. This capability makes the applications portable across development, test, and production environments running on physical or virtual machines locally, in data centres, and across the networks of different cloud service providers.

[FlexPod Datacenter with Docker Enterprise Edition for Container Management](#)

FlexPod Datacenter for OpenShift Container Platform 4 - Design

Haseeb Niazi, Cisco
Alan Cowles, NetApp

Cisco and NetApp have partnered to deliver a series of FlexPod solutions that enable strategic data center platforms. The FlexPod solution delivers an integrated architecture that incorporates best practices for computing, storage, and network design, thereby minimizing IT risks by validating the integrated architecture to ensure compatibility between various components. The solution also addresses IT pain points by providing documented design guidance, deployment guidance, and support that can be used in various stages (planning, designing and implementation) of a deployment.

[FlexPod Datacenter for OpenShift Container Platform 4 - Design](#)

FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning - Deployment

Haseeb Niazi, Cisco

Arvind Ramakrishnan, NetApp

This document provides deployment details and guidance around the integration of the Cisco UCS C480 ML M5 platform into the FlexPod data center solution to deliver a unified approach for providing AI and ML capabilities within the converged infrastructure. This document also explains the NVIDIA GPUs configuration on Cisco UCS C220 and C240 platforms. For a detailed design discussion about the platforms and technologies used in this solution, refer to the [FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning design](#).

[FlexPod Datacenter for combined AI and ML with Cisco UCS 480 ML for deep learning - Deployment](#)

3D graphics visualization with VMware and NVIDIA on Cisco UCS - White paper

This document describes the performance of the VMware ESXi hypervisor and VMware Horizon with NVIDIA Tesla P4, P6, and P40 solution on Cisco UCS C240 M5 Rack Servers and B200 M5 Blade Servers.

[3D graphics visualization with VMware and NVIDIA on Cisco UCS - White paper](#)

3D graphics visualization with Citrix and NVIDIA - White paper

This document describes the performance of Citrix XenDesktop on Citrix XenServer with NVIDIA Tesla P4, P6, and P40 cards on Cisco UCS C240 M5 and B200 M5 servers with SPECviewperf 13.

[3D graphics visualization with Citrix and NVIDIA - White paper](#)

FlexPod Express

FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series Design Guide

NVA-1139-DESIGN: FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series

Savita Kumari, NetApp

In partnership with:



Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices that uses the technology that they are familiar with in their data center.

FlexPod Express is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp AFF systems. The components of FlexPod Express are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

[Next: Program summary.](#)

Program summary

FlexPod Converged Infrastructure Portfolio

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or as NetApp Verified Architectures (NVAs). Deviations that are based on customer requirements from a given CVD or NVA are permitted if those variations do not result in the deployment of unsupported configurations.

As illustrated in the following figure, the FlexPod portfolio includes the following solutions: FlexPod Express and FlexPod Datacenter.

- **FlexPod Express** is an entry-level solution with technologies from Cisco and NetApp.
- **FlexPod Datacenter** delivers an optimal multipurpose foundation for various workloads and applications.

Expanded portfolio of platforms

FlexPod® Express

Departmental deployments and VAR velocity

Target: Primarily MSB, remote, and departmental deployments



Entry level: Cisco UCS, Cisco Nexus, and NetApp AFF and FAS systems

FlexPod Datacenter

Massively scalable, mission-critical workloads

Target: Enterprise/service provider



Cisco UCS, Cisco Nexus, and NetApp AFF and FAS systems

Distinct Architectures

Distinct Architectures

NetApp Verified Architecture program

The NetApp Verified Architecture program offers customers a verified architecture for NetApp solutions. An NVA solution has the following qualities:

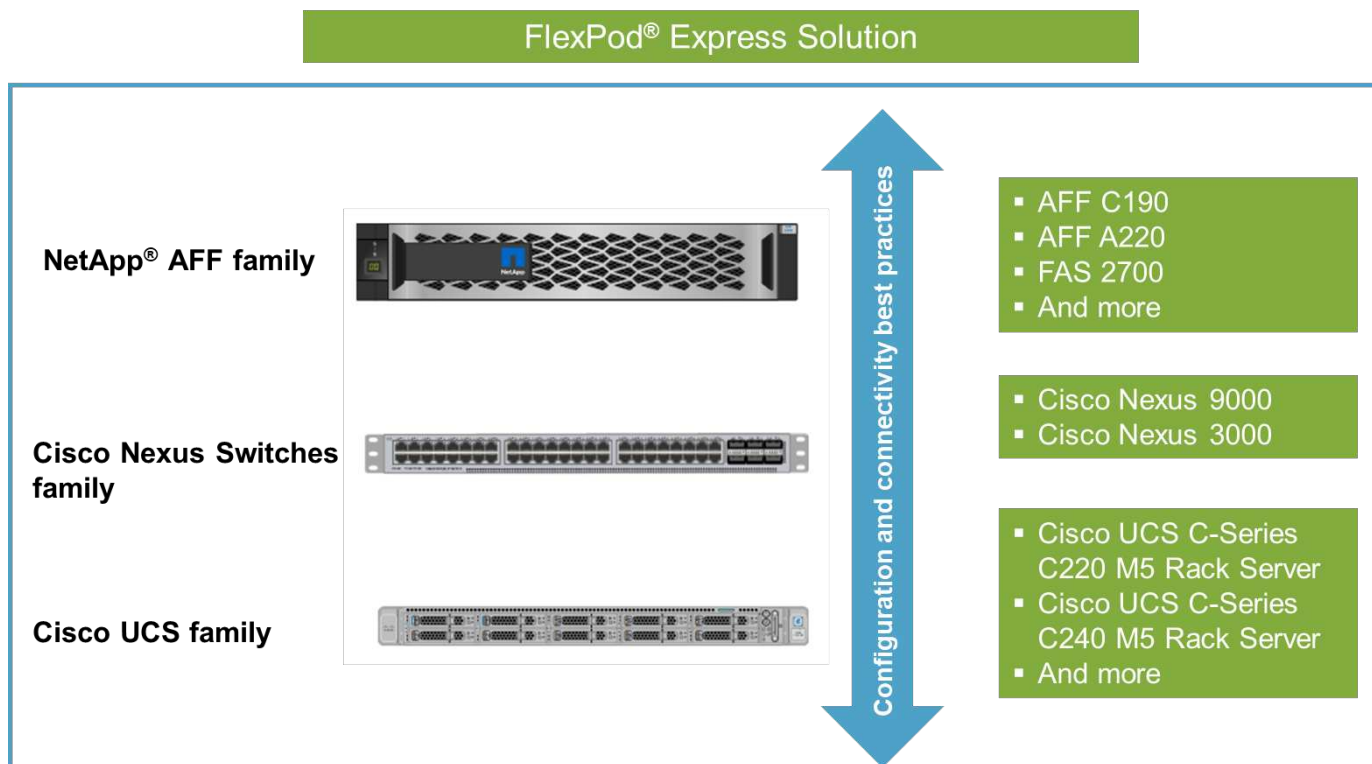
- Is thoroughly tested
- Is prescriptive in nature
- Minimizes deployment risks
- Accelerates time to market This guide details the design of FlexPod Express with VMware vSphere.

In addition, this design leverages the all-new AFF C190 system, which runs NetApp ONTAP 9.6 software, Cisco Nexus 31108 switches, and Cisco UCS C220 M5 servers as hypervisor nodes.

Solution overview

FlexPod Express is designed to run mixed virtualization workloads. It is targeted for remote and branch offices and for small to midsize businesses. It is also optimal for larger businesses that want to implement a dedicated solution for a specific purpose. This new solution for FlexPod Express adds new technologies such as NetApp ONTAP 9.6, NetApp AFF C190 system, and VMware vSphere 6.7U2.

The following figure shows the hardware components that are included in the FlexPod Express solution.




Target audience

This document is intended for people who want to take advantage of an infrastructure that is built to deliver IT efficiency and to enable IT innovation. The audience for this document includes, but is not limited to, sales engineers, field consultants, professional services personnel, IT managers, partner engineers, and customers.

Solution technology

This solution leverages the latest technologies from NetApp, Cisco, and VMware. It features the new NetApp AFF C190 system, which runs ONTAP 9.6 software, dual Cisco Nexus 31108 switches, and Cisco UCS C220 M5 rack servers that run VMware vSphere 6.7U2. This validated solution, illustrated in the following figure, uses 10 Gigabit Ethernet (10GbE) technology. Guidance is also provided on how to scale by adding two hypervisor nodes at a time so that the FlexPod Express architecture can adapt to an organization's evolving business needs.



Next: [Technology requirements.](#)

Technology requirements

FlexPod Express requires a combination of hardware and software components that depends on the selected hypervisor and network speed. In addition, FlexPod Express lays out the hardware components that are required to add hypervisor nodes to the system in units of two.

Hardware requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, you can use a different hypervisor on the same FlexPod Express hardware.

The following table lists the hardware components that are required for this FlexPod Express configuration and to implement this solution. The hardware components that are used in any implementation of the solution can vary based on customer requirements.

| Hardware | Quantity |
|--------------------------|----------|
| AFF C190 2-node cluster | 1 |
| Cisco UCS C220 M5 Server | 2 |
| Cisco Nexus 31108 Switch | 2 |

| Hardware | Quantity |
|---|----------|
| Cisco UCS Virtual Interface Card (VIC) 1457 for Cisco UCS C220 M5 rack server | 2 |

Software requirements

The following table lists the software components that are required to implement the architectures of the FlexPod Express solution.

| Software | Version | Details |
|---|-------------|---------------------------------|
| Cisco Integrated Management Controller (CIMC) | 4.0.4 | For C220 M5 rack servers |
| Cisco NX-OS | 7.0(3)I7(6) | For Cisco Nexus 31108 switches |
| NetApp ONTAP | 9.6 | For NetApp AFF C190 controllers |

The following table lists the software that is required for all VMware vSphere implementations on FlexPod Express.

| Software | Version |
|---------------------------------|---------|
| VMware vCenter Server Appliance | 6.7U2 |
| VMware vSphere ESXi | 6.7U2 |
| NetApp VAAI Plug-In for ESXi | 1.1.2 |
| NetApp Virtual Storage Console | 9.6 |

Next: [Design choices](#).

Design choices

The technologies listed in this section were chosen during the architectural design phase. Each technology serves a specific purpose in the FlexPod Express infrastructure solution.

NetApp AFF C190 Series with ONTAP 9.6

This solution leverages two of the newest NetApp products: NetApp AFF C190 system and ONTAP 9.6 software.

AFF C190 system

The target group is customers who want to modernize their IT infrastructure with all- flash technology at an affordable price. The AFF C190 system comes with the new ONTAP 9.6 and flash bundle licensing, which means that the following functions are on board:

- CIFS, NFS, iSCSI, and FCP
- NetApp SnapMirror data replication software, NetApp SnapVault backup software, NetApp SnapRestore data recovery software, NetApp SnapManager storage management software product suite, and NetApp SnapCenter software

- FlexVol technology
- Deduplication, compression, and compaction
- Thin provisioning
- Storage QoS
- NetApp RAID DP technology
- NetApp Snapshot technology
- FabricPool

The following figures show the two options for host connectivity.

The following figure illustrates UTA 2 ports where SFP+ module can be inserted.



The following figure illustrates 10GBASE-T ports for connection through conventional RJ-45 Ethernet cables.



For the 10GBASE-T port option, you must have a 10GBASE-T based uplink switch.

The AFF C190 system is offered exclusively with 960GB SSDs. There are four stages of expansions from which you can choose:

- 8x 960GB
- 12x 960GB
- 18x 960GB
- 24x 960GB

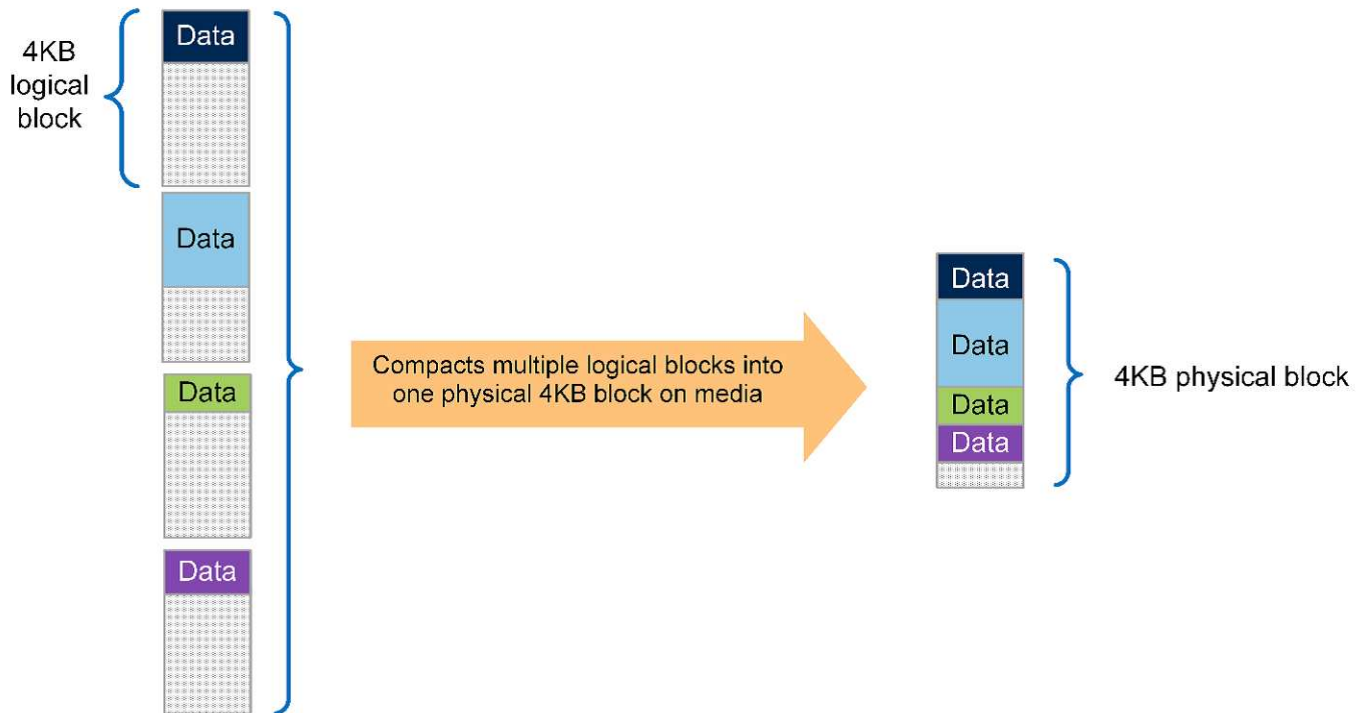
For full information about the AFF C190 hardware system, see the [NetApp AFF C190 All-Flash Array page](#).

ONTAP 9.6 software

NetApp AFF C190 systems use the new ONTAP 9.6 data management software. ONTAP 9.6 is the industry's leading enterprise data management software. It combines new levels of simplicity and flexibility with powerful

data management capabilities, storage efficiencies, and leading cloud integration.

ONTAP 9.6 has several features that are well suited for the FlexPod Express solution. Foremost is NetApp's commitment to storage efficiencies, which can be one of the most important features for small deployments. The hallmark NetApp storage efficiency features such as deduplication, compression, compaction, and thin provisioning are available in ONTAP 9.6. The NetApp WAFL system always writes 4KB blocks; therefore, compaction combines multiple blocks into a 4KB block when the blocks are not using their allocated space of 4KB. The following figure illustrates this process.



ONTAP 9.6 now supports an optional 512-byte block size for NVMe volumes. This capability works well with the VMware Virtual Machine File System (VMFS), which natively uses a 512-byte block. You can stay with the default 4K size or optionally set the 512-byte block size.

Other feature enhancements in ONTAP 9.6 include:

- **NetApp Aggregate Encryption (NAE).** NAE assigns keys at the aggregate level, thereby encrypting all volumes in the aggregate. This feature allows volumes to be encrypted and deduplicated at the aggregate level.
- **NetApp ONTAP FlexGroup volume enhancement.** In ONTAP 9.6, you can easily rename a FlexGroup volume. There's no need to create a new volume to migrate the data to. The volume size can also be reduced by using ONTAP System Manager or CLI.
- **FabricPool enhancement.** ONTAP 9.6 added additional support for object stores as cloud tiers. Support for Google Cloud and Alibaba Cloud Object Storage Service (OSS) was also added to the list. FabricPool supports multiple object stores, including AWS S3, Azure Blob, IBM Cloud object storage, and NetApp StorageGRID object-based storage software.
- **SnapMirror enhancement.** In ONTAP 9.6, a new volume replication relationship is encrypted by default before leaving the source array and is decrypted at the SnapMirror destination.

Cisco Nexus 3000 Series

The Cisco Nexus 31108PC-V is a 10Gbps SFP+ based top-of-rack (ToR) switch with 48 SFP+ ports and 6 QSFP28 ports. Each SFP+ port can operate in 100Mbps, 10Gbps, and each QSFP28 port can operate in native 100Gbps or 40Gbps mode or 4x 10Gbps mode, offering flexible migration options. This switch is a true PHY-less switch that is optimized for low latency and low power consumption.

The Cisco Nexus 31108PC-V specification includes the following components:

- 2.16Tbps switching capacity and forwarding rate of up to 1.2Tbps for 31108PC-V
- 48 SFP ports support 1 and 10 Gigabit Ethernet (10GbE); 6x QSFP28 ports support 4x 10GbE or 40GbE each or 100GbE

The following figure illustrates the Cisco Nexus 31108PC-V switch.



For more information about Cisco Nexus 31108PC-V switches, see [Cisco Nexus 3172PQ, 3172TQ, 3172TQ-32T, 3172PQ-XL, and 3172TQ-XL Switches Data Sheet](#).

Cisco UCS C-Series

The Cisco UCS C-Series rack server was chosen for FlexPod Express because its many configuration options allow it to be tailored for specific requirements in a FlexPod Express deployment.

Cisco UCS C-Series rack servers deliver unified computing in an industry-standard form factor to reduce TCO and to increase agility.

Cisco UCS C-Series rack servers offer the following benefits:

- A form-factor-agnostic entry point into Cisco UCS
- Simplified and fast deployment of applications
- Extension of unified computing innovations and benefits to rack servers
- Increased customer choice with unique benefits in a familiar rack package




The Cisco UCS C220 M5 rack server, shown in the above figure, is among the most versatile general-purpose enterprise infrastructure and application servers in the industry. It is a high-density two-socket rack server that delivers industry-leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications. Cisco UCS C-Series rack servers can be deployed as standalone servers or as part of Cisco UCS to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' TCO and increase their business agility.

For more information about C220 M5 servers, see [Cisco UCS C220 M5 Rack Server Data Sheet](#).

Cisco UCS VIC 1457 connectivity for C220 M5 rack servers

The Cisco UCS VIC 1457 adapter shown in the following figure is a quad-port small form-factor pluggable (SFP28) modular LAN on motherboard (mLOM) card designed for the M5 generation of Cisco UCS C-Series Servers. The card supports 10/25Gbps Ethernet or FCoE. The card can present PCIe standards-compliant interfaces to the host, and these can be dynamically configured as either NICs or HBAs.



For full information about the Cisco UCS VIC 1457 adapter, see [Cisco UCS Virtual Interface Card 1400 Series Data Sheet](#).

VMware vSphere 6.7U2

VMware vSphere 6.7U2 is one of the hypervisor options for use with FlexPod Express. VMware vSphere allows organizations to reduce their power and cooling footprint while confirming that the purchased compute capacity is used to its fullest. In addition, VMware vSphere allows hardware failure protection (VMware High Availability, or VMware HA) and compute resource load balancing across a cluster of vSphere hosts (VMware Distributed Resource Scheduler in maintenance mode, or VMware DRS-MM).

Because it restarts only the kernel, VMware vSphere 6.7U2 allows customers to quick boot, loading vSphere ESXi without restarting the hardware. The vSphere 6.7U2 vSphere client (HTML5-based client) has some new enhancements like Developer Center with Code Capture and API Explore. With Code Capture, you can record your actions in the vSphere client to deliver simple, usable code output. vSphere 6.7U2 also contains new features like DRS in maintenance mode (DRS-MM).

VMware vSphere 6.7U2 offers the following features:

- VMware is deprecating the external VMware Platform Services Controller (PSC) deployment model.



Starting with the next major vSphere release, external PSC will not be an available option.

- New protocol support for backing up and restoring a vCenter server appliance. Introducing NFS and SMB as supported protocol choices, up to 7 total (HTTP, HTTPS, FTP, FTPS, SCP, NFS, and SMB) when configuring a vCenter Server for file-based backup or restore operations.
- New functionality when using the content library. Syncing a native VM template between content libraries is now available when the vCenter Server is configured for enhanced linked mode.

- Update to the [Client Plug-Ins page](#).
- VMware vSphere Update Manager also adds enhancements to the vSphere client. You can perform attach-check compliance and remediate actions all from one screen.

For more information about VMware vSphere 6.7 U2, see the [VMware vSphere Blog page](#).

For more information about the VMware vCenter Server 6.7 U2 updates, see the [Release Notes](#).



Although this solution was validated with vSphere 6.7U2, it supports any vSphere version qualified with the other components by the [NetApp Interoperability Matrix Tool \(IMT\)](#). NetApp recommends that you deploy the next released version of vSphere for its fixes and enhanced features.

Boot architecture

The supported options for the FlexPod Express boot architecture include:


- iSCSI SAN LUN
- Cisco FlexFlash SD card
- Local disk

FlexPod Datacenter is booted from iSCSI LUNs; therefore, solution manageability is enhanced by using iSCSI boot for FlexPod Express as well.

ESXi Host Virtual Network Interface Card layout

Cisco UCS VIC 1457 has four physical ports. This solution validation includes these four physical ports in using the ESXi host. If you have a smaller or larger number of NICs, you might have different VMNIC numbers.

In an iSCSI boot implementation, iSCSI boot requires separate virtual network interface cards (vNICs) for iSCSI boot. These vNICs use the appropriate fabric's iSCSI VLAN as the native VLAN and are attached to the iSCSI boot vSwitches, as shown in the following figure.



Next: [Conclusion.](#)

Conclusion

The FlexPod Express validated design is a simple and effective solution that uses industry-leading components. By scaling and providing options for the hypervisor platform, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed for small to midsize businesses, remote and branch offices, and other businesses that require dedicated solutions.

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

Where to find additional information

To learn more about the information described in this document, see the following documents and websites:

- AFF and FAS System Documentation Center

<https://docs.netapp.com/platstor/index.jsp>

- AFF Documentation Resources page

<https://www.netapp.com/us/documentation/all-flash-fas.aspx>

- FlexPod Express with VMware vSphere 6.7 and NetApp AFF C190 Deployment Guide (in progress)

- NetApp documentation

<https://docs.netapp.com>

FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series Deployment Guide

NVA-1142-DEPLOY: FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series - NVA Deployment

Savita Kumari, NetApp

Industry trends indicate that a vast data center transformation is occurring toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices that uses technology that they are familiar with in their data center.

FlexPod® Express is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp® storage technologies. The components in a FlexPod Express system are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

FlexPod Datacenter and FlexPod Express deliver a baseline configuration and have the flexibility to be sized and optimized to accommodate many different use cases and requirements. Existing FlexPod Datacenter customers can manage their FlexPod Express system with the tools to which they are accustomed. New FlexPod Express customers can easily transition to managing FlexPod Datacenter as their environment grows.

FlexPod Express is an optimal infrastructure foundation for remote and branch offices and for small to midsize businesses. It is also an optimal solution for customers who want to provide infrastructure for a dedicated workload.

FlexPod Express provides an easy-to-manage infrastructure that is suitable for almost any workload.

Solution overview

This FlexPod Express solution is part of the FlexPod Converged Infrastructure Program.

FlexPod converged infrastructure program

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or NetApp Verified Architectures (NVAs). Deviations based on customer requirements from a given CVD or NVA are permitted if these variations do not create an unsupported configuration.

The FlexPod program includes two solutions: FlexPod Express and FlexPod Datacenter.

- **FlexPod Express.** Offers customers an entry-level solution with technologies from Cisco and NetApp.
- **FlexPod Datacenter.** Delivers an optimal multipurpose foundation for various workloads and applications.

The FlexPod Portfolio

A prevalidated, flexible platform that features



FlexPod® Express

Remote office or branch office, retail, small and midsize business, and edge



FlexPod Datacenter

Enterprise apps, unified infrastructure, and virtualization

11

NetApp Verified Architecture program

The NetApp Verified Architecture program offers customers a verified architecture for NetApp solutions. A NetApp Verified Architecture provides a NetApp solution architecture with the following qualities:


- Thoroughly tested
- Prescriptive in nature
- Minimized deployment risks
- Accelerated time to market

This guide details the design of FlexPod Express with VMware vSphere. In addition, this design uses the all-new AFF C190 system (running NetApp ONTAP® 9.6), the Cisco Nexus 31108, and Cisco UCS C-Series C220 M5 servers as hypervisor nodes.

Solution technology

This solution leverages the latest technologies from NetApp, Cisco, and VMware. This solution features the new NetApp AFF C190 running ONTAP 9.6, dual Cisco Nexus 31108 switches, and Cisco UCS C220 M5 rack

servers running VMware vSphere 6.7U2. This validated solution uses 10GbE technology. Guidance is also provided on how to scale compute capacity by adding two hypervisor nodes at a time so that the FlexPod Express architecture can adapt to an organization's evolving business needs.



To use the four physical 10GbE ports on the VIC 1457 efficiently, create two extra links from each server to the top rack switches.

Use case summary

The FlexPod Express solution can be applied to several use cases, including the following:

- Remote or branch offices
- Small and midsize businesses
- Environments that require a dedicated and cost-effective solution

FlexPod Express is best suited for virtualized and mixed workloads. Although this solution was validated with vSphere 6.7U2, it supports any vSphere version qualified with the other components by the NetApp Interoperability Matrix Tool. NetApp recommends deploying vSphere 6.7U2 because of its fixes and enhanced features, such as the following:

- New protocol support for backing up and restoring a vCenter server appliance, including HTTP, HTTPS, FTP, FTPS, SCP, NFS and SMB.
- New functionality when utilizing the content library. Syncing of native VM templates between content libraries is now available when vCenter Server is configured for enhanced linked mode.

- An updated Client Plug-In page.
- Added enhancements in the vSphere Update Manager (VUM) and the vSphere client. You can now perform the attach, check- compliance, and remediate actions, all from one screen.

For more information on this subject, see the [vSphere 6.7U2 page](#) and the [vCenter Server 6.7U2 Release Notes](#).

Technology requirements

A FlexPod Express system requires a combination of hardware and software components. FlexPod Express also describes the hardware components that are required to add hypervisor nodes to the system in units of two.

Hardware requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, you can use a different hypervisor on the same FlexPod Express hardware.

The following table lists the hardware components that are required for FlexPod Express configuration and implementation. The hardware components that are used in any implementation of the solution might vary based on customer requirements.

| Hardware | Quantity |
|---|----------|
| AFF C190 two-node cluster | 1 |
| Cisco C220 M5 server | 2 |
| Cisco Nexus 31108PC-V switch | 2 |
| Cisco UCS virtual interface card (VIC) 1457 for Cisco UCS C220 M5 rack server | 2 |

This table lists the hardware that is required in addition to the base configuration for implementing 10GbE.

| Hardware | Quantity |
|--------------------------|----------|
| Cisco UCS C220 M5 server | 2 |
| Cisco VIC 1457 | 2 |

Software requirements

The following table lists the software components that are required to implement the architectures of the FlexPod Express solutions.

| Software | Version | Details |
|---|----------|------------------------------------|
| Cisco Integrated Management Controller (CIMC) | 4.0.4 | For Cisco UCS C220 M5 rack servers |
| Cisco nenic driver | 1.0.0.29 | For VIC 1457 interface cards |

| Software | Version | Details |
|--------------|-------------|------------------------------------|
| Cisco NX-OS | 7.0(3)I7(6) | For Cisco Nexus 31108PC-V switches |
| NetApp ONTAP | 9.6 | For AFF C190 controllers |


This table lists the software that is required for all VMware vSphere implementations on FlexPod Express.

| Software | Version |
|---------------------------------|---------|
| VMware vCenter server appliance | 6.7U2 |
| VMware vSphere ESXi hypervisor | 6.7U2 |
| NetApp VAAI Plug-In for ESXi | 1.1.2 |
| NetApp VSC | 9.6 |

FlexPod Express cabling information

This reference validation is cabled as shown in the following figures and tables.

This figure shows the reference validation cabling.



The following table lists the cabling information for Cisco Nexus switch 31108PC-V-A.

| Local device | Local port | Remote device | Remote port |
|-----------------------------------|------------|---|--------------------|
| Cisco Nexus switch 31108PC-V A | Eth1/1 | NetApp AFF C190 storage controller A | e0c |
| | Eth1/2 | NetApp AFF C190 storage controller B | e0c |
| | Eth1/3 | Cisco UCS C220 C-Series standalone server A | MLOM0 |
| | Eth1/4 | Cisco UCS C220 C-Series standalone server B | MLOM0 |
| | Eth1/5 | Cisco UCS C220 C-Series standalone server A | MLOM1 |
| | Eth1/6 | Cisco UCS C220 C-Series standalone server B | MLOM1 |
| | Eth1/25 | Cisco Nexus switch 31108PC-V B | Eth1/25 |
| | Eth1/26 | Cisco Nexus switch 31108PC-V B | Eth1/26 |
| | Eth1/33 | NetApp AFF C190 storage controller A | e0M |
| | Eth1/34 | Cisco UCS C220 C-Series standalone server A | CIMC (FEX135/1/25) |

This table lists the cabling information for Cisco Nexus switch 31108PC-V- B.

| Local device | Local port | Remote device | Remote port |
|--------------------------------|------------|---|--------------------|
| Cisco Nexus switch 31108PC-V B | Eth1/1 | NetApp AFF C190 storage controller A | e0d |
| | Eth1/2 | NetApp AFF C190 storage controller B | e0d |
| | Eth1/3 | Cisco UCS C220 C-Series standalone server A | MLOM2 |
| | Eth1/4 | Cisco UCS C220 C-Series standalone server B | MLOM2 |
| | Eth1/5 | Cisco UCS C220 C-Series standalone server A | MLOM3 |
| | Eth1/6 | Cisco UCS C220 C-Series standalone server B | MLOM3 |
| | Eth1/25 | Cisco Nexus switch 31108 A | Eth1/25 |
| | Eth1/26 | Cisco Nexus switch 31108 A | Eth1/26 |
| | Eth1/33 | NetApp AFF C190 storage controller B | e0M |
| | Eth1/34 | Cisco UCS C220 C-Series standalone server B | CIMC (FEX135/1/26) |

This table lists the cabling information for NetApp AFF C190 storage controller A.

| Local device | Local Port | Remote device | Remote port |
|--------------------------------------|------------|--------------------------------------|-------------|
| NetApp AFF C190 storage controller A | e0a | NetApp AFF C190 storage controller B | e0a |
| | e0b | NetApp AFF C190 storage controller B | e0b |
| | e0c | Cisco Nexus switch 31108PC-V A | Eth1/1 |
| | e0d | Cisco Nexus switch 31108PC-V B | Eth1/1 |
| | e0M | Cisco Nexus switch 31108PC-V A | Eth1/33 |

This table lists the cabling information for NetApp AFF C190 storage controller B.

| Local device | Local port | Remote device | Remote port |
|--------------------------------------|------------|--------------------------------------|-------------|
| NetApp AFF C190 storage controller B | e0a | NetApp AFF C190 storage controller A | e0a |
| | e0b | NetApp AFF C190 storage controller A | e0b |
| | e0c | Cisco Nexus switch 31108PC-V A | Eth1/2 |
| | e0d | Cisco Nexus switch 31108PC-V B | Eth1/2 |
| | e0M | Cisco Nexus switch 31108PC-V B | Eth1/33 |

Deployment procedures

Overview

This document provides details for configuring a fully redundant, highly available FlexPod Express system. To reflect this redundancy, the components being configured in each step are referred to as either component A or component B. For example, controller A and controller B identify the two NetApp storage controllers that are provisioned in this document. Switch A and switch B identify a pair of Cisco Nexus switches.

In addition, this document describes steps for provisioning multiple Cisco UCS hosts, which are identified sequentially as server A, server B, and so on.

To indicate that you should include information pertinent to your environment in a step, `<<text>>` appears as part of the command structure. See the following example for the `vlan create` command:

```
Controller01> network port vlan create -node <<var_nodeA>> -vlan-name
<<var_vlan-name>>
```

This document enables you to fully configure the FlexPod Express environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and virtual local area network (VLAN) schemes. The following table describes the VLANs required for deployment, as outlined in this guide. This table can be completed based on the specific site variables and used to implement the document configuration steps.



If you use separate in-band and out-of-band management VLANs, you must create a layer-3 route between them. For this validation, a common management VLAN was used.

| VLAN name | VLAN purpose | VLAN ID | |
|-----------------|--------------------------------|---------|----------|
| Management VLAN | VLAN for management interfaces | 3437 | vSwitch0 |
| NFS VLAN | VLAN for NFS traffic | 3438 | vSwitch0 |

| VLAN name | VLAN purpose | VLAN ID | |
|---------------------|--|---------|------------------|
| VMware vMotion VLAN | VLAN designated for the movement of virtual machines (VMs) from one physical host to another | 3441 | vSwitch0 |
| VM traffic VLAN | VLAN for VM application traffic | 3442 | vSwitch0 |
| iSCSI-A-VLAN | VLAN for iSCSI traffic on fabric A | 3439 | iScsiBootvSwitch |
| iSCSI-B-VLAN | VLAN for iSCSI traffic on fabric B | 3440 | iScsiBootvSwitch |
| Native VLAN | VLAN to which untagged frames are assigned | 2 | |

The VLAN numbers are needed throughout the configuration of FlexPod Express. The VLANs are referred to as <<var_xxxx_vlan>>, where xxxx is the purpose of the VLAN (such as iSCSI-A).

There are two vSwitches created in this validation.

The following table lists the solution vSwitches.

| vSwitch name | Active adapters | Ports | MTU | Load balancing |
|------------------|-----------------|---------------|------|---|
| vSwitch0 | Vmnic2, vmnic4 | default (120) | 9000 | Route based on IP hash |
| iScsiBootvSwitch | Vmnic3, vmnic5 | default (120) | 9000 | Route based on the originating virtual port ID. |



The IP hash method of load balancing requires proper configuration for the underlying physical switch using SRC-DST-IP EtherChannel with a static (mode on) port-channel. In the event of intermittent connectivity due to possible switch misconfiguration, temporarily shut down one of the two associated uplink ports on the Cisco switch to restore communication to the ESXi management vmkernel port while troubleshooting the port-channel settings.

The following table lists the VMware VMs that are created.

| VM description | Host name |
|-------------------------|--------------|
| VMware vCenter Server | FlexPod-VCSA |
| Virtual Storage Console | FlexPod-VSC |

Deploy Cisco Nexus 31108PC-V

This section details the Cisco Nexus 331108PC-V switch configuration used in a FlexPod Express environment.

Initial Setup of Cisco Nexus 31108PC-V Switch

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod Express environment.



This procedure assumes that you are using a Cisco Nexus 31108PC-V running NX-OS software release 7.0(3)I7(6).

1. Upon initial boot and connection to the console port of the switch, the Cisco NX-OS setup automatically starts. This initial configuration addresses basic settings, such as the switch name, the mgmt0 interface configuration, and Secure Shell (SSH) setup.
2. The FlexPod Express management network can be configured in multiple ways. The mgmt0 interfaces on the 31108PC-V switches can be connected to an existing management network, or the mgmt0 interfaces of the 31108PC-V switches can be connected in a back-to-back configuration. However, this link cannot be used for external management access such as SSH traffic.



In this deployment guide, the FlexPod Express Cisco Nexus 31108PC-V switches are connected to an existing management network.

3. To configure the Cisco Nexus 31108PC-V switches, power on the switch and follow the on-screen prompts, as illustrated here for the initial setup of both the switches, substituting the appropriate values for the switch-specific information.

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

*Note: setup is mainly used for configuring the system initially, when no configuration is present. So setup always assumes system defaults and not the current system configuration values.

Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip the remaining dialogs.

Would you like to enter the basic configuration dialog (yes/no): y

Do you want to enforce secure password standard (yes/no) [y]: y

Create another login account (yes/no) [n]: n

Configure read-only SNMP community string (yes/no) [n]: n

Configure read-write SNMP community string (yes/no) [n]: n

Enter the switch name : 31108PC-V-B

Continue with Out-of-band (mgmt0) management configuration? (yes/no) [y]: y

Mgmt0 IPv4 address : <<var_switch_mgmt_ip>>

Mgmt0 IPv4 netmask : <<var_switch_mgmt_netmask>>

Configure the default gateway? (yes/no) [y]: y

IPv4 address of the default gateway : <<var_switch_mgmt_gateway>>

Configure advanced IP options? (yes/no) [n]: n

Enable the telnet service? (yes/no) [n]: n

Enable the ssh service? (yes/no) [y]: y

Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa

Number of rsa key bits <1024-2048> [1024]: <enter>

Configure the ntp server? (yes/no) [n]: y

NTP server IPv4 address : <<var_ntp_ip>>

Configure default interface layer (L3/L2) [L2]: <enter>

Configure default switchport interface state (shut/noshut) [noshut]: <enter>

Configure CoPP system profile (strict/moderate/lenient/dense) [strict]: <enter>

4. You then see a summary of your configuration, and you are asked if you would like to edit it. If your configuration is correct, enter n.

Would you like to edit the configuration? (yes/no) [n]: n

5. You are then asked if you would like to use this configuration and save it. If so, enter y.

Use this configuration and save it? (yes/no) [y]: Enter

6. Repeat this procedure for Cisco Nexus switch B.

Enable the advanced features

Certain advanced features must be enabled in Cisco NX-OS to provide additional configuration options. To enable the appropriate features on Cisco Nexus switch A and switch B, enter configuration mode using the command (config t) and run the following commands:

```
feature interface-vlan
feature lacp
feature vpc
```



The default port channel load-balancing hash uses the source and destination IP addresses to determine the load-balancing algorithm across the interfaces in the port channel. You can achieve better distribution across the members of the port channel by providing more inputs to the hash algorithm beyond the source and destination IP addresses. For the same reason, NetApp highly recommends adding the source and destination TCP ports to the hash algorithm.

From configuration mode (config t), enter the following commands to set the global port channel load-balancing configuration on Cisco Nexus switch A and switch B:

```
port-channel load-balance src-dst ip-l4port
```

Configure global spanning tree

The Cisco Nexus platform uses a new protection feature called bridge assurance. Bridge assurance helps protect against a unidirectional link or other software failure with a device that continues to forward data traffic when it is no longer running the spanning-tree algorithm. Ports can be placed in one of several states, including network or edge, depending on the platform.

NetApp recommends setting bridge assurance so that all ports are considered to be network ports by default. This setting forces the network administrator to review the configuration of each port. It also reveals the most common configuration errors, such as unidentified edge ports or a neighbor that does not have the bridge assurance feature enabled. In addition, it is safer to have the spanning tree block many ports rather than too few, which allows the default port state to enhance the overall stability of the network.

Pay close attention to the spanning-tree state when adding servers, storage, and uplink switches, especially if they do not support bridge assurance. In such cases, you might need to change the port type to make the ports active.

The Bridge Protocol Data Unit (BPDU) guard is enabled on edge ports by default as another layer of protection. To prevent loops in the network, this feature shuts down the port if BPDUs from another switch are seen on this interface.

From configuration mode (config t), run the following commands to configure the default spanning tree options, including the default port type and BPDU guard, on Cisco Nexus switch A and switch B:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
spanning-tree port type edge bpdufilter default
ntp server <<var_ntp_ip>> use-vrf management
ntp master 3
```

Define the VLANs

Before individual ports with different VLANs are configured, the layer- 2 VLANs must be defined on the switch. It is also a good practice to name the VLANs for easy troubleshooting in the future.

From configuration mode (config t), run the following commands to define and describe the layer- 2 VLANs on Cisco Nexus switch A and switch B:

```
vlan <<nfs_vlan_id>>
  name NFS-VLAN
vlan <<iSCSI_A_vlan_id>>
  name iSCSI-A-VLAN
vlan <<iSCSI_B_vlan_id>>
  name iSCSI-B-VLAN
vlan <<vmotion_vlan_id>>
  name vMotion-VLAN
vlan <<vmtraffic_vlan_id>>
  name VM-Traffic-VLAN
vlan <<mgmt_vlan_id>>
  name MGMT-VLAN
vlan <<native_vlan_id>>
  name NATIVE-VLAN
exit
```

Configure access and management port descriptions

As is the case with assigning names to the layer- 2 VLANs, setting descriptions for all the interfaces can help with both provisioning and troubleshooting.

From configuration mode (config t) in each of the switches, enter the following port descriptions for the FlexPod Express large configuration:

Cisco Nexus Switch A

```

int eth1/1
    description AFF C190-A e0c
int eth1/2
    description AFF C190-B e0c
int eth1/3
    description UCS-Server-A: MLOM port 0 vSwitch0
int eth1/4
    description UCS-Server-B: MLOM port 0 vSwitch0
int eth1/5
    description UCS-Server-A: MLOM port 1 iScsiBootvSwitch
int eth1/6
    description UCS-Server-B: MLOM port 1 iScsiBootvSwitch
int eth1/25
    description vPC peer-link 31108PC-V-B 1/25
int eth1/26
    description vPC peer-link 31108PC-V-B 1/26
int eth1/33
    description AFF C190-A e0M
int eth1/34
    description UCS Server A: CIMC

```

Cisco Nexus Switch B

```

int eth1/1
    description AFF C190-A e0d
int eth1/2
    description AFF C190-B e0d
int eth1/3
    description UCS-Server-A: MLOM port 2 vSwitch0
int eth1/4
description UCS-Server-B: MLOM port 2 vSwitch0
int eth1/5
    description UCS-Server-A: MLOM port 3 iScsiBootvSwitch
int eth1/6
    description UCS-Server-B: MLOM port 3 iScsiBootvSwitch
int eth1/25
    description vPC peer-link 31108PC-V-A 1/25
int eth1/26
    description vPC peer-link 31108PC-V-A 1/26
int eth1/33
    description AFF C190-B e0M
int eth1/34
    description UCS Server B: CIMC

```

Configure server and storage management interfaces

The management interfaces for both the server and the storage typically use only a single VLAN. Therefore, configure the management interface ports as access ports. Define the management VLAN for each switch and change the spanning-tree port type to edge.

From configuration mode (config t), enter the following commands to configure the port settings for the management interfaces of both the servers and the storage:

Cisco Nexus Switch A

```
int eth1/33-34
  switchport mode access
  switchport access vlan <<mgmt_vlan>>
  spanning-tree port type edge
  speed 1000
exit
```

Cisco Nexus Switch B

```
int eth1/33-34
  switchport mode access
  switchport access vlan <<mgmt_vlan>>
  spanning-tree port type edge
  speed 1000
exit
```

Perform the virtual port channel global configuration

A virtual port channel (vPC) enables links that are physically connected to two different Cisco Nexus switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device. A vPC can provide layer- 2 multipathing, which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

A vPC provides the following benefits:

- Enabling a single device to use a port channel across two upstream devices
- Eliminating spanning-tree- protocol blocked ports
- Providing a loop-free topology
- Using all available uplink bandwidth
- Providing fast convergence if either the link or a device fails
- Providing link-level resiliency
- Helping provide high availability

The vPC feature requires some initial setup between the two Cisco Nexus switches to function properly. If you use the back-to-back mgmt0 configuration, use the addresses defined on the interfaces and verify that they

can communicate by using the ping <<switch_A/B_mgmt0_ip_addr>>vrf management command.

From configuration mode (config t), run the following commands to configure the vPC global configuration for both switches:

Cisco Nexus Switch A

```
vpc domain 1
  role priority 10
  peer-keepalive destination <<switch_B_mgmt0_ip_addr>> source
<<switch_A_mgmt0_ip_addr>> vrf
management
peer-switch
peer-gateway
auto-recovery
delay restore 150
ip arp synchronize
int eth1/25-26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>,
<<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
  spanning-tree port type network
  vpc peer-link
  no shut
exit
copy run start
```

Cisco Nexus Switch B


```

vpc domain 1
  peer-switch
  role priority 20
  peer-keepalive destination <<switch_A_mgmt0_ip_addr>> source
<<switch_B_mgmt0_ip_addr>> vrf management
  peer-gateway
  auto-recovery
  delay-restore 150
  ip arp synchronize
int eth1/25-26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>,
<<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
  spanning-tree port type network
  vpc peer-link
no shut
exit
copy run start

```

Configure the storage port channels

The NetApp storage controllers allow an active-active connection to the network using the Link Aggregation Control Protocol (LACP). The use of LACP is preferred because it adds both negotiation and logging between the switches. Because the network is set up for vPC, this approach enables you to have active-active connections from the storage to separate physical switches. Each controller has two links to each of the switches. However, all four links are part of the same vPC and interface group (ifgrp).

From configuration mode (config t), run the following commands on each of the switches to configure the individual interfaces and the resulting port channel configuration for the ports connected to the NetApp AFF controller.

1. Run the following commands on switch A and switch B to configure the port channels for storage controller A:

```

int eth1/1
    channel-group 11 mode active
int Po11
    description vPC to Controller-A
    switchport
    switchport mode trunk
    switchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan
<<nfs_vlan_id>>,<<mgmt_vlan_id>>,<<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
    spanning-tree port type edge trunk
    mtu 9216
    vpc 11
    no shut

```

2. Run the following commands on switch A and switch B to configure the port channels for storage controller B:

```

int eth1/2
    channel-group 12 mode active
int Po12
    description vPC to Controller-B
    switchport
    switchport mode trunk
    switchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan <<nfs_vlan_id>>,<<mgmt_vlan_id>>,
<<iSCSI_A_vlan_id>>, <<iSCSI_B_vlan_id>>
    spanning-tree port type edge trunk
    mtu 9216
    vpc 12
    no shut
exit
copy run start

```

Configure the server connections

The Cisco UCS servers have a four-port virtual interface card, VIC1457, that is used for data traffic and booting of the ESXi operating system using iSCSI. These interfaces are configured to fail over to one another, providing additional redundancy beyond a single link. Spreading these links across multiple switches enables the server to survive even a complete switch failure.

From configuration mode (config t), run the following commands to configure the port settings for the interfaces connected to each server.

Cisco Nexus Switch A: Cisco UCS Server-A and Cisco UCS Server-B configuration

```
int eth1/5
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<iSCSI_A_vlan_id>>,<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_i
d>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  no shut
exit
copy run start
```

Cisco Nexus Switch B: Cisco UCS Server-A and Cisco UCS Server-B configuration

```
int eth1/6
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<iSCSI_B_vlan_id>>,<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_i
d>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  no shut
exit
copy run start
```

Configure the server port channels

Run the following commands on switch A and switch B to configure the port channels for Server-A:

```

int eth1/3
  channel-group 13 mode active
int Po13
  description vPC to Server-A
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_id>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  vpc 13
  no shut

```

Run the following commands on switch A and switch B to configure the port channels for Server-B:

```

int eth1/4
  channel-group 14 mode active
int Po14
  description vPC to Server-B
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_id>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  vpc 14
  no shut

```



An MTU of 9000 was used in this solution validation. However, you can configure an different value for the MTU appropriate for your application requirements. It is important to set the same MTU value across the FlexPod solution. Incorrect MTU configurations between components result in packets being dropped and these packets will need to be transmitted again, affecting the overall performance of the solution.



To scale the solution by adding additional Cisco UCS servers, run the previous commands with the switch ports that the newly added servers have been plugged into on switches A and B.

Uplink into an existing network infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 31108 switches included in the FlexPod environment into the infrastructure. The uplinks can be 10GbE uplinks for a 10GbE infrastructure solution or 1GbE for a 1GbE infrastructure solution if

required. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy start to save the configuration on each switch after the configuration is completed.

Next: [NetApp storage deployment procedure \(part 1\)](#).

NetApp storage deployment procedure (part 1)

This section describes the NetApp AFF storage deployment procedure.

NetApp storage controller AFF C190 Series installation

NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities.

Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install:

Access the [HWU](#) application to view the system configuration guides. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.

Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controller AFFC190 Series prerequisites

To plan the physical location of the storage systems, see the NetApp Hardware Universe. Refer to the following sections:

- Electrical Requirements
- Supported Power Cords
- Onboard Ports and Cables

Storage controllers

Follow the physical installation procedures for the controllers in the AFF [C190](#) Documentation.

NetApp ONTAP 9.6

Configuration worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the ONTAP 9.6 Software Setup Guide.



This system is set up in a two-node switchless cluster configuration.

The following table provides the ONTAP 9.6 installation and configuration information.

| Cluster detail | Cluster detail value |
|---|-----------------------------|
| Cluster node A IP address | <<var_nodeA_mgmt_ip>> |
| Cluster node A netmask | <<var_nodeA_mgmt_mask>> |
| Cluster node A gateway | <<var_nodeA_mgmt_gateway>> |
| Cluster node A name | <<var_nodeA>> |
| Cluster node B IP address | <<var_nodeB_mgmt_ip>> |
| Cluster node B netmask | <<var_nodeB_mgmt_mask>> |
| Cluster node B gateway | <<var_nodeB_mgmt_gateway>> |
| Cluster node B name | <<var_nodeB>> |
| ONTAP 9.6 URL | <<var_url_boot_software>> |
| Name for cluster | <<var_clustername>> |
| Cluster management IP address | <<var_clustermgmt_ip>> |
| Cluster B gateway | <<var_clustermgmt_gateway>> |
| Cluster B netmask | <<var_clustermgmt_mask>> |
| Domain name | <<var_domain_name>> |
| DNS server IP (you can enter more than one) | <var_dns_server_ip |
| NTP server IP (you can enter more than one) | <<var_ntp_server_ip>> |

Configure Node A

To configure node A, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

Allow the system to boot.

```
autoboot
```

2. Press Ctrl-C to enter the Boot menu.



If ONTAP 9.6 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.6 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

3. To install new software, select option 7.
4. Enter y to perform an upgrade.

5. Select e0M for the network port you want to use for the download.
6. Enter y to reboot now.
7. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeA_mgmt_ip>> <<var_nodeA_mgmt_mask>> <<var_nodeA_mgmt_gateway>>
```

8. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

9. Press Enter for the user name, indicating no user name.
10. Enter y to set the newly installed software as the default to be used for subsequent reboots.
11. Enter y to reboot the node.



When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

12. Press Ctrl-C to enter the Boot menu.
13. Select option 4 for Clean Configuration and Initialize All Disks.
14. Enter y to zero disks, reset config, and install a new file system.
15. Enter y to erase all the data on the disks.



The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the node B configuration while the disks for node A are zeroing.

While node A is initializing, begin configuring node B.

Configure Node B

To configure node B, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Press Ctrl-C to enter the Boot menu.

```
autoboot
```

3. Press Ctrl-C when prompted.



If ONTAP 9.6 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.6 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.A.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeB_mgmt_ip>> <<var_nodeB_mgmt_ip>><<var_nodeB_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press Enter for the user name, indicating no user name.
11. Enter y to set the newly installed software as the default to be used for subsequent reboots.
12. Enter y to reboot the node.



When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter y to zero disks, reset config, and install a new file system.
16. Enter y to erase all the data on the disks.



The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

Continuation of the node A configuration and cluster configuration

From a console port program attached to the storage controller A (node A) console port, run the node setup script. This script appears when ONTAP 9.6 boots on the node for the first time.



The node and cluster setup procedure has changed slightly in ONTAP 9.6. The cluster setup wizard is now used to configure the first node in a cluster, and NetApp ONTAP System Manager (formerly OnCommand® System Manager) is used to configure the cluster.

1. Follow the prompts to set up node A.

```
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
  "help" or "?" - if you want to have a question clarified,
  "back" - if you want to change previously answered questions, and
  "exit" or "quit" - if you want to quit the cluster setup wizard.
  Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.
Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/
Type yes to confirm and continue {yes}: yes
Enter the node management interface port [e0M]:
Enter the node management interface IP address: <<var_nodeA_mgmt_ip>>
Enter the node management interface netmask: <<var_nodeA_mgmt_mask>>
Enter the node management interface default gateway:
<<var_nodeA_mgmt_gateway>>
A node management interface on port e0M with IP address
<<var_nodeA_mgmt_ip>> has been created.
Use your web browser to complete cluster setup by accessing
https://<<var_nodeA_mgmt_ip>>
Otherwise, press Enter to complete cluster setup using the command line
interface:
```

2. Navigate to the IP address of the node's management interface.



Cluster setup can also be performed by using the CLI. This document describes cluster setup using System Manager guided setup.

3. Click Guided Setup to configure the cluster.

4. Enter <<var_clustername>> for the cluster name and <<var_nodeA>> and <<var_nodeB>> for each of the nodes that you are configuring. Enter the password that you would like to use for the storage system. Select Switchless Cluster for the cluster type. Enter the cluster base license.
5. You can also enter feature licenses for Cluster, NFS, and iSCSI.
6. You see a status message stating the cluster is being created. This status message cycles through several statuses. This process takes several minutes.
7. Configure the network.
 - a. Deselect the IP Address Range option.
 - b. Enter <<var_clustermgmt_ip>> in the Cluster Management IP Address field, <<var_clustermgmt_mask>> in the Netmask field, and <<var_clustermgmt_gateway>> in the Gateway field. Use the ... selector in the Port field to select e0M of node A.
 - c. The node management IP for node A is already populated. Enter <<var_nodeA_mgmt_ip>> for node B.
 - d. Enter <<var_domain_name>> in the DNS Domain Name field. Enter <<var_dns_server_ip>> in the DNS Server IP Address field.



You can enter multiple DNS server IP addresses.

- e. Enter 10.63.172.162 in the Primary NTP Server field.



You can also enter an alternate NTP server. The IP address 10.63.172.162 from <<var_ntp_server_ip>> is the Nexus Mgmt IP.

8. Configure the support information.
 - a. If your environment requires a proxy to access AutoSupport, enter the URL in Proxy URL.
 - b. Enter the SMTP mail host and email address for event notifications.



You must, at a minimum, set up the event notification method before you can proceed. You can select any of the methods.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



? AutoSupport ☒

? Proxy URL (Optional)

i Connection is verified after configuring AutoSupport on all nodes.

? Event Notifications

Notify me through:

| | | | |
|-------------------------------------|--------|-------------------------------------|---|
| <input checked="" type="checkbox"/> | Email | SMTP Mail Host <input type="text"/> | Email Addresses <input type="text" value="Separate email addresses with a comma..."/> |
| <hr/> | | | |
| <input type="checkbox"/> | SNMP | SNMP Trap Host <input type="text"/> | |
| <hr/> | | | |
| <input type="checkbox"/> | Syslog | Syslog Server <input type="text"/> | |
| <hr/> | | | |

Submit

When the system indicates that the cluster configuration has completed, click Manage Your Cluster to configure the storage.

Continuation of the storage cluster configuration

After the configuration of the storage nodes and base cluster, you can continue with the configuration of the storage cluster.

Zero all spare disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```

Set the on-board UTA2 ports personality

1. Verify the current mode and the current type for the ports by running the `ucadmin show` command.

```
AFF C190::> ucadmin show
```

| Node | Adapter | Current Mode | Current Type | Pending Mode | Pending Type | Admin Status |
|------------|---------|--------------|--------------|--------------|--------------|--------------|
| AFF C190_A | 0c | cna | target | - | - | online |
| AFF C190_A | 0d | cna | target | - | - | online |
| AFF C190_A | 0e | cna | target | - | - | online |
| AFF C190_A | 0f | cna | target | - | - | online |
| AFF C190_B | 0c | cna | target | - | - | online |
| AFF C190_B | 0d | cna | target | - | - | online |
| AFF C190_B | 0e | cna | target | - | - | online |
| AFF C190_B | 0f | cna | target | - | - | online |

8 entries were displayed.

2. Verify that the current mode of the ports that are in use is `cna` and that the current type is set to `target`. If not, change the port personality by using the following command:

```
ucadmin modify -node <home node of the port> -adapter <port name> -mode cna -type target
```



The ports must be offline to run the previous command. To take a port offline, run the following command:

```
network fcp adapter modify -node <home node of the port> -adapter <port name> -state down
```



If you changed the port personality, you must reboot each node for the change to take effect.

Rename the management logical interfaces

To rename the management logical interfaces (LIFs), complete the following steps:

1. Show the current management LIF names.

```
network interface show -vserver <<clustername>>
```

2. Rename the cluster management LIF.

```
network interface rename -vserver <<clustername>> -lif  
cluster_setup_cluster_mgmt_lif_1 -newname cluster_mgmt
```

3. Rename the node B management LIF.

```
network interface rename -vserver <<clustername>> -lif  
cluster_setup_node_mgmt_lif_AFF C190_B_1 -newname AFF C190-02_mgmt1
```

Set auto-revert on cluster management

Set the auto-revert parameter on the cluster management interface.

```
network interface modify -vserver <<clustername>> -lif cluster_mgmt -auto-  
revert true
```

Set up the service processor network interface

To assign a static IPv4 address to the service processor on each node, run the following commands:

```
system service-processor network modify -node <<var_nodeA>> -address  
-family IPv4 -enable true -dhcp none -ip-address <<var_nodeA_sp_ip>>  
-netmask <<var_nodeA_sp_mask>> -gateway <<var_nodeA_sp_gateway>>  
system service-processor network modify -node <<var_nodeB>> -address  
-family IPv4 -enable true -dhcp none -ip-address <<var_nodeB_sp_ip>>  
-netmask <<var_nodeB_sp_mask>> -gateway <<var_nodeB_sp_gateway>>
```



The service processor IP addresses should be in the same subnet as the node management IP addresses.

Enable storage failover in ONTAP

To confirm that storage failover is enabled, run the following commands in a failover pair:

1. Verify the status of storage failover.

```
storage failover show
```



Both <<var_nodeA>> and <<var_nodeB>> must be able to perform a takeover. Go to step 3 if the nodes can perform a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_nodeA>> -enabled true
```



Enabling failover on one node enables it for both nodes.

3. Verify the HA status of the two-node cluster.



This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

4. Go to step 6 if high availability is configured. If high availability is configured, you see the following message upon issuing the command:

```
High Availability Configured: true
```

5. Enable HA mode only for the two-node cluster.



Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true  
Do you want to continue? {y|n}: y
```

6. Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

```
storage failover hwassist show
```



The message `Keep Alive Status: Error:` indicates that one of the controllers did not receive `hwassist keep alive` alerts from its partner, indicating that hardware assist is not configured. Run the following commands to configure hardware assist.

```
storage failover modify -hwassist-partner-ip <<var_nodeB_mgmt_ip>> -node
<<var_nodeA>>
storage failover modify -hwassist-partner-ip <<var_nodeA_mgmt_ip>> -node
<<var_nodeB>>
```

Create a jumbo frame MTU broadcast domain in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Remove the data ports from the default broadcast domain

The 10GbE data ports are used for iSCSI/NFS traffic, and these ports should be removed from the default domain. Ports e0e and e0f are not used and should also be removed from the default domain.

To remove the ports from the broadcast domain, run the following command:

```
broadcast-domain remove-ports -broadcast-domain Default -ports
<<var_nodeA>>:e0c, <<var_nodeA>>:e0d, <<var_nodeA>>:e0e,
<<var_nodeA>>:e0f, <<var_nodeB>>:e0c, <<var_nodeB>>:e0d,
<<var_nodeA>>:e0e, <<var_nodeA>>:e0f
```

Disable flow control on UTA2 ports

It is a NetApp best practice to disable flow control on all UTA2 ports that are connected to external devices. To disable flow control, run the following command:

```

net port modify -node <<var_nodeA>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y

```

Configure the interface group LACP in ONTAP

This type of interface group requires two or more Ethernet interfaces and a switch that supports LACP. make sure it's configured based on the steps in this guide in section 5.1.

From the cluster prompt, complete the following steps:


```

ifgrp create -node <<var_nodeA>> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0d
ifgrp create -node << var_nodeB>> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0d

```

Configure the jumbo frames in ONTAP

To configure an ONTAP network port to use jumbo frames (usually with an MTU of 9,000 bytes), run the following commands from the cluster shell:

```

AFF C190::> network port modify -node node_A -port a0a -mtu 9000
Warning: This command will cause a several second interruption of service
on
        this network port.
Do you want to continue? {y|n}: y
AFF C190::> network port modify -node node_B -port a0a -mtu 9000
Warning: This command will cause a several second interruption of service
on
        this network port.
Do you want to continue? {y|n}: y

```

Create VLANs in ONTAP

To create VLANs in ONTAP, complete the following steps:

1. Create NFS VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_nfs_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_nfs_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_NFS -ports
<<var_nodeA>>:a0a-<<var_nfs_vlan_id>>, <<var_nodeB>>:a0a-
<<var_nfs_vlan_id>>

```

2. Create iSCSI VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_iscsi_vlan_B_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_iscsi_vlan_B_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports
<<var_nodeA>>:a0a-<<var_iscsi_vlan_A_id>>,<<var_nodeB>>:a0a-
<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports
<<var_nodeA>>:a0a-<<var_iscsi_vlan_B_id>>,<<var_nodeB>>:a0a-
<<var_iscsi_vlan_B_id>>

```

3. Create MGMT-VLAN ports.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<mgmt_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<mgmt_vlan_id>>

```

Create data aggregates in ONTAP

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it contains.

To create aggregates, run the following commands:

```

aggr create -aggregate aggr1_nodeA -node <<var_nodeA>> -diskcount
<<var_num_disks>>
aggr create -aggregate aggr1_nodeB -node <<var_nodeB>> -diskcount
<<var_num_disks>>

```



Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.



Start with five disks; you can add disks to an aggregate when additional storage is required.



The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display the aggregate creation status. Do not proceed until `aggr1_nodeA` is online.

Configure Time Zone in ONTAP

To configure time synchronization and to set the time zone on the cluster, run the following command:

```
timezone <<var_timezone>>
```



For example, in the eastern United States, the time zone is America/New_York. After you begin typing the time zone name, press the Tab key to see available options.

Configure SNMP in ONTAP

To configure the SNMP, complete the following steps:

1. Configure SNMP basic information, such as the location and contact. When polled, this information is visible as the `sysLocation` and `sysContact` variables in SNMP.

```
snmp contact <<var_snmp_contact>>
snmp location "<<var_snmp_location>>"
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts.

```
snmp traphost add <<var_snmp_server_fqdn>>
```

Configure SNMPv1 in ONTAP

To configure SNMPv1, set the shared secret plain-text password called a community.

```
snmp community add ro <<var_snmp_community>>
```



Use the `snmp community delete all` command with caution. If community strings are used for other monitoring products, this command removes them.

Configure SNMPv3 in ONTAP

SNMPv3 requires that you define and configure a user for authentication. To configure SNMPv3, complete the following steps:

1. Run the `security snmpusers` command to view the engine ID.
2. Create a user called `snmpv3user`.

```
security login create -username snmpv3user -authmethod usm -application snmp
```

3. Enter the authoritative entity's engine ID and select md5 as the authentication protocol.
4. Enter an eight-character minimum-length password for the authentication protocol when prompted.
5. Select des as the privacy protocol.
6. Enter an eight-character minimum-length password for the privacy protocol when prompted.

Configure AutoSupport HTTPS in ONTAP

The NetApp AutoSupport tool sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts <<var_mailhost>> -transport https -support enable -noteto <<var_storage_admin_email>>
```

Create a storage virtual machine

To create an infrastructure storage virtual machine (SVM), complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate aggr1_nodeA -rootvolume-security-style unix
```

2. Add the data aggregate to the infra-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Infra-SVM -aggr-list aggr1_nodeA,aggr1_nodeB
```

3. Remove the unused storage protocols from the SVM, leaving NFS and iSCSI.

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,ndmp,fc
```

4. Enable and run the NFS protocol in the infra-SVM SVM.

```
nfs create -vserver Infra-SVM -udp disabled
```

5. Turn on the SVM `vstorage` parameter for the NetApp NFS VAAI plug-in. Then, verify that NFS has been configured.

```
vserver nfs modify -vserver Infra-SVM -vstorage enabled
vserver nfs show
```



Commands are prefaced by `vserver` in the command line because SVMs were previously called Vservers.

Configure NFSv3 in ONTAP

The following table lists the information needed to complete this configuration.

| Detail | Detail value |
|----------------------------|---------------------------|
| ESXi host A NFS IP address | <<var_esxi_hostA_nfs_ip>> |
| ESXi host B NFS IP address | <<var_esxi_hostB_nfs_ip>> |

To configure NFS on the SVM, run the following commands:

1. Create a rule for each ESXi host in the default export policy.
2. For each ESXi host being created, assign a rule. Each host has its own rule index. Your first ESXi host has rule index 1, your second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default
-ruleindex 1 -protocol nfs -clientmatch <<var_esxi_hostA_nfs_ip>>
-rorule sys -rwrule sys -superuser sys -allow-suid false
vserver export-policy rule create -vserver Infra-SVM -policyname default
-ruleindex 2 -protocol nfs -clientmatch <<var_esxi_hostB_nfs_ip>>
-rorule sys -rwrule sys -superuser sys -allow-suid false
vserver export-policy rule show
```

3. Assign the export policy to the infrastructure SVM root volume.

```
volume modify -vserver Infra-SVM -volume rootvol -policy default
```



The NetApp VSC automatically handles export policies if you choose to install it after vSphere has been set up. If you do not install it, you must create export policy rules when additional Cisco UCS C-Series servers are added.

Create the iSCSI service in ONTAP

To create the iSCSI service on the SVM, run the following command. This command also starts the iSCSI service and sets the iSCSI IQN for the SVM. Verify that iSCSI has been configured.

```
iscsi create -vserver Infra-SVM
iscsi show
```

Create load-sharing mirror of SVM root volume in ONTAP

To create a load-sharing mirror of the SVM root volume in ONTAP, complete the following steps:

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver Infra_Vserver -volume rootvol_m01 -aggregate
aggr1_nodeA -size 1GB -type DP
volume create -vserver Infra_Vserver -volume rootvol_m02 -aggregate
aggr1_nodeB -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:rootvol -destination-path
Infra-SVM:rootvol_m01 -type LS -schedule 15min
snapmirror create -source-path Infra-SVM:rootvol -destination-path
Infra-SVM:rootvol_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship and verify that it has been created.

```
snapmirror initialize-ls-set -source-path Infra-SVM:rootvol
snapmirror show
```

Configure HTTPS access in ONTAP

To configure secure access to the storage controller, complete the following steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a certificate authority.



Deleting expired certificates before creating certificates is a best practice. Run the `security certificate delete` command to delete expired certificates. In the following command, use TAB completion to select and delete each default certificate.

```
security certificate delete [TAB] ...  
Example: security certificate delete -vserver Infra-SVM -common-name  
Infra-SVM -ca Infra-SVM -type server -serial 552429A6
```

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the infra-SVM and the cluster SVM. Again, use TAB completion to aid in completing these commands.

```
security certificate create [TAB] ...  
Example: security certificate create -common-name infra-svm.netapp.com  
-type server -size 2048 -country US -state "North Carolina" -locality  
"RTP" -organization "NetApp" -unit "FlexPod" -email-addr  
"abc@netapp.com" -expire-days 3650 -protocol SSL -hash-function SHA256  
-vserver Infra-SVM
```

5. To obtain the values for the parameters required in the following step, run the `security certificate show` command.
6. Enable each certificate that was just created using the `-server-enabled true` and `-client-enabled false` parameters. Again, use TAB completion.

```
security ssl modify [TAB] ...  
Example: security ssl modify -vserver Infra-SVM -server-enabled true  
-client-enabled false -ca infra-svm.netapp.com -serial 55243646 -common  
-name infra-svm.netapp.com
```

7. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -sslsv3-enabled true
Warning: Modifying the cluster configuration will cause pending web
service requests to be interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http
-vserver <<var_clustername>>
```



It is normal for some of these commands to return an error message stating that the entry does not exist.

8. Revert to the admin privilege level and create the setup to allow the SVM to be available by the web.

```
set -privilege admin
vserver services web modify -name spi -vserver * -enabled true
```

Create a NetApp FlexVol volume in ONTAP

To create a NetApp FlexVol® volume, enter the volume name, size, and the aggregate on which it exists. Create two VMware datastore volumes and a server boot volume.

```
volume create -vserver Infra-SVM -volume infra_datastore -aggregate
aggr1_nodeB -size 500GB -state online -policy default -junction-path
/infra_datastore -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_nodeA
-size 100GB -state online -policy default -junction-path /infra_swap
-space-guarantee none -percent-snapshot-space 0 -snapshot-policy none
-efficiency-policy none
volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_nodeA
-size 100GB -state online -policy default -space-guarantee none -percent
-snapshot-space 0
```

Create LUNs in ONTAP

To create two boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-A -size
15GB -ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-B -size
15GB -ostype vmware -space-reserve disabled
```



When adding an extra Cisco UCS C-Series server, you must create an extra boot LUN.

Create iSCSI LIFs in ONTAP

The following table lists the information needed to complete this configuration.

| Detail | Detail value |
|--|---------------------------------|
| Storage node A iSCSI LIF01A | <<var_nodeA_iscsi_lif01a_ip>> |
| Storage node A iSCSI LIF01A network mask | <<var_nodeA_iscsi_lif01a_mask>> |
| Storage node A iSCSI LIF01B | <<var_nodeA_iscsi_lif01b_ip>> |
| Storage node A iSCSI LIF01B network mask | <<var_nodeA_iscsi_lif01b_mask>> |
| Storage node B iSCSI LIF01A | <<var_nodeB_iscsi_lif01a_ip>> |
| Storage node B iSCSI LIF01A network mask | <<var_nodeB_iscsi_lif01a_mask>> |
| Storage node B iSCSI LIF01B | <<var_nodeB_iscsi_lif01b_ip>> |
| Storage node B iSCSI LIF01B network mask | <<var_nodeB_iscsi_lif01b_mask>> |

Create four iSCSI LIFs, two on each node.

```
network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data
-data-protocol iscsi -home-node <<var_nodeA>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_nodeA_iscsi_lif01a_ip>> -netmask
<<var_nodeA_iscsi_lif01a_mask>> -status-admin up -failover-policy disabled
-firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data
-data-protocol iscsi -home-node <<var_nodeA>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_nodeA_iscsi_lif01b_ip>> -netmask
<<var_nodeA_iscsi_lif01b_mask>> -status-admin up -failover-policy disabled
-firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data
-data-protocol iscsi -home-node <<var_nodeB>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_nodeB_iscsi_lif01a_ip>> -netmask
<<var_nodeB_iscsi_lif01a_mask>> -status-admin up -failover-policy disabled
-firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data
-data-protocol iscsi -home-node <<var_nodeB>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_nodeB_iscsi_lif01b_ip>> -netmask
<<var_nodeB_iscsi_lif01b_mask>> -status-admin up -failover-policy disabled
-firewall-policy data -auto-revert false
network interface show
```

Create NFS LIFs in ONTAP

The following table lists the information needed to complete this configuration.

| Detail | Detail value |
|--|-------------------------------|
| Storage node A NFS LIF 01 IP | <<var_nodeA_nfs_lif_01_ip>> |
| Storage node A NFS LIF 01 network mask | <<var_nodeA_nfs_lif_01_mask>> |
| Storage node B NFS LIF 02 IP | <<var_nodeB_nfs_lif_02_ip>> |
| Storage node B NFS LIF 02 network mask | <<var_nodeB_nfs_lif_02_mask>> |

Create an NFS LIF.

```
network interface create -vserver Infra-SVM -lif nfs_lif01 -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_nodeA_nfs_lif_01_ip>> -netmask <<
var_nodeA_nfs_lif_01_mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif02 -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_nodeB_nfs_lif_02_ip>> -netmask <<
var_nodeB_nfs_lif_02_mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
network interface show
```

Add an infrastructure SVM administrator

The following table lists the information needed to add an SVM administrator.

| Detail | Detail value |
|------------------------|--------------------------|
| Vsmgmt IP | <<var_svm_mgmt_ip>> |
| Vsmgmt network mask | <<var_svm_mgmt_mask>> |
| Vsmgmt default gateway | <<var_svm_mgmt_gateway>> |

To add the infrastructure SVM administrator and SVM administration logical interface to the management network, complete the following steps:

1. Run the following command:

```
network interface create -vserver Infra-SVM -lif vsmgmt -role data
-data-protocol none -home-node <<var_nodeB>> -home-port e0M -address
<<var_svm_mgmt_ip>> -netmask <<var_svm_mgmt_mask>> -status-admin up
-failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-
revert true
```



The SVM management IP here should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway
<<var_svm_mgmt_gateway>>
network route show
```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Infra-SVM
Enter a new password: <<var_password>>
Enter it again: <<var_password>>
security login unlock -username vsadmin -vserver Infra-SVM
```

[Next: Deploy Cisco UCS C-Series rack server.](#)

Deploy Cisco UCS C-Series rack server

This section provides a detailed procedure for configuring a Cisco UCS C-Series standalone rack server for use in the FlexPod Express configuration.

Perform the initial Cisco UCS C-Series standalone server setup for CIMC

Complete these steps for the initial setup of the CIMC interface for Cisco UCS C-Series standalone servers.

The following table lists the information needed to configure CIMC for each Cisco UCS C-Series standalone server.

| Detail | Detail value |
|----------------------|------------------|
| CIMC IP address | <<cimc_ip>> |
| CIMC subnet mask | \<<cimc_netmask |
| CIMC default gateway | <<cimc_gateway>> |



The CIMC version used in this validation is CIMC 4.0.(4).

All servers

1. Attach the Cisco keyboard, video, and mouse (KVM) dongle (provided with the server) to the KVM port on the front of the server. Plug a VGA monitor and USB keyboard into the appropriate KVM dongle ports.

Power on the server and press F8 when prompted to enter the CIMC configuration.



Copyright (c) 2019 Cisco Systems, Inc.

Press <F2> BIOS Setup : <F6> Boot Menu : <F7> Diagnostics
Press <F8> CIMC Setup : <F12> Network Boot
Bios Version : C220M5.4.0.4g.0.0712190011
Platform ID : C220M5

Processor(s) Intel(R) Xeon(R) Silver 4114 CPU @ 2.20GHz
Total Memory = 64 GB Effective Memory = 64 GB
Memory Operating Speed 2400 Mhz
M.2 SWRAID configuration is not detected. Switching to AHCI mode.

Cisco IMC IPv4 Address : 10.63.172.160
Cisco IMC MAC Address : 70:69:5A:B5:8D:68

Entering CIMC Configuration Utility ...

92

2. In the CIMC configuration utility, set the following options:

a. Network interface card (NIC) mode:

Dedicated ☒ [X]

b. IP (Basic):

IPV4: ☒ [X]

DHCP enabled: ☐ []

CIMC IP: <<cimc_ip>>

Prefix/Subnet: <<cimc_netmask>>

Gateway: <<cimc_gateway>>

c. VLAN (Advanced): Leave cleared to disable VLAN tagging.

NIC redundancy

None: ☒ [X]

```

Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
NIC Properties
NIC mode
Dedicated:      [X]          NIC redundancy
Shared LOM:     [ ]          None:                          [X]
Cisco Card:     [ ]          Active-standby:                 [ ]
Riser1:        [ ]          Active-active:                   [ ]
Riser2:        [ ]          VLAN (Advanced)
MLom:          [ ]          VLAN enabled:                     [ ]
Shared LOM Ext: [ ]          VLAN ID:                          1
Priority:                                0
IP (Basic)
IPv4:           [X]          IPv6:      [ ]
DHCP enabled    [ ]
CIMC IP:        10.63.172.160
Prefix/Subnet:  255.255.255.0
Gateway:        10.63.172.1
Pref DNS Server: 0.0.0.0
Smart Access USB
Enabled         [ ]
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F1>Additional settings

```

3. Press F1 to see the additional settings:

a. Common properties:

Host name: <<esxi_host_name>>

Dynamic DNS: []

Factory defaults: Leave cleared.

b. Default user (basic):

Default password: <<admin_password>>

Reenter password: <<admin_password>>

Port properties: Use default values.

Port profiles: Leave cleared.

4. Press F10 to save the CIMC interface configuration.

5. After the configuration is saved, press Esc to exit.

Configure Cisco UCS C-Series Servers iSCSI boot

In this FlexPod Express configuration, the VIC1457 is used for iSCSI boot.

The following table lists the information needed to configure iSCSI boot.



An italicized font indicates variables that are unique for each ESXi host.

| Detail | Detail value |
|-----------------------------------|----------------------------------|
| ESXi host initiator A name | <<var_ucs_initiator_name_A>> |
| ESXi host iSCSI-A IP | <<var_esxi_host_iscsiA_ip>> |
| ESXi host iSCSI-A network mask | <<var_esxi_host_iscsiA_mask>> |
| ESXi host iSCSI A default gateway | <<var_esxi_host_iscsiA_gateway>> |
| ESXi host initiator B name | <<var_ucs_initiator_name_B>> |
| ESXi host iSCSI-B IP | <<var_esxi_host_iscsiB_ip>> |
| ESXi host iSCSI-B network mask | <<var_esxi_host_iscsiB_mask>> |
| ESXi host iSCSI-B gateway | <<var_esxi_host_iscsiB_gateway>> |
| IP address iscsi_lif01a | <<var_iscsi_lif01a>> |
| IP address iscsi_lif02a | <<var_iscsi_lif02a>> |
| IP address iscsi_lif01b | <<var_iscsi_lif01b>> |
| IP address iscsi_lif02b | <<var_iscsi_lif02b>> |
| Infra_SVM IQN | <<var_SVM_IQN>> |

Boot order configuration

To set the boot order configuration, complete the following steps:

1. From the CIMC interface browser window, click the Compute tab and select BIOS.
2. Click Configure Boot Order and then click OK.

Cisco Integrated Management Controller

[Home](#) / [Compute](#) / [BIOS](#) ★

[BIOS](#)
[Remote Management](#)
[Troubleshooting](#)
[Power Policies](#)
[PID Catalog](#)

[Enter BIOS Setup](#) | [Clear BIOS CMOS](#) | [Restore Manufacturing Custom Settings](#) | [Restore Defaults](#)

[Configure BIOS](#)
[Configure Boot Order](#)
[Configure BIOS Profile](#)

BIOS Properties

Running Version

C220M5.4.0.4g.0.0712190011

UEFI Secure Boot

☐

Actual Boot Mode

Uefi

Configured Boot Mode

▼

Last Configured Boot Order Source

BIOS

Configured One time boot device

▼

Save Changes

▼ Configured Boot Devices

Basic

▶ ☒ Advanced

Actual Boot Devices

UEFI: Built-in EFI Shell (NonPolicyTarget)

UEFI: PXE IP4 Intel(R) Ethernet Controller X550 (NonPolicyTarget)

UEFI: PXE IP4 Intel(R) Ethernet Controller X550 (NonPolicyTarget)

Configure Boot Order

3. Configure the following devices by clicking the device under Add Boot Device and going to the Advanced tab:

a. Add Virtual Media:

Name: KVM-CD-DVD

Subtype: KVM MAPPED DVD

State: Enabled

Order: 1

b. Add iSCSI Boot:

Name: iSCSI-A

State: Enabled

Order: 2

Slot: MLOM

Port: 1

c. Click Add iSCSI Boot:

Name: iSCSI-B

State: Enabled

Order: 3

Slot: MLOM

Port: 3

4. Click Add Device.

5. Click Save Changes and then click Close.

Configure Boot Order

Configured Boot Level: Advanced

Basic Advanced

Add Boot Device

- Add Local HDD
- Add PXE Boot
- Add SAN Boot
- Add iSCSI Boot
- Add USB
- Add Virtual Media
- Add PCHStorage
- Add UEFISHELL
- Add SD Card
- Add NVME
- Add Local CDD

Advanced Boot Order Configuration

Selected 1 / Total 3

| | Name | Type | Order | State |
|-------------------------------------|----------------|--------|-------|---------|
| <input checked="" type="checkbox"/> | KVM-MAPPED-DVD | VMEDIA | 1 | Enabled |
| <input type="checkbox"/> | iSCSI-A | ISCSI | 2 | Enabled |
| <input type="checkbox"/> | iSCSI-B | ISCSI | 3 | Enabled |

Save Changes Reset Values Close

6. Reboot the server to boot with your new boot order.

Disable RAID controller (if present)

Complete the following steps if your C-Series server contains a RAID controller. A RAID controller is not needed in the boot from SAN configuration. Optionally, you can also physically remove the RAID controller from the server.

1. Under the Compute tab, click BIOS in the left navigation pane in CIMC.
2. Select Configure BIOS.
3. Scroll down to PCIe Slot:HBA Option ROM.
4. If the value is not already disabled, set it to disabled.

| | | | | | |
|------|-------------------|-----------------|----------------|-------------|-------------------|
| BIOS | Remote Management | Troubleshooting | Power Policies | PID Catalog | |
| I/O | Server Management | Security | Processor | Memory | Power/Performance |

Note: Default values are shown in bold.

Reboot Host Immediately: ☒

| | |
|---------------------------|---------------|
| Intel VT for directed IO: | Enabled ▼ |
| Intel VTD ATS support: | Enabled ▼ |
| LOM Port 1 OptionRom: | Enabled ▼ |
| Pcie Slot 1 OptionRom: | Disabled ▼ |
| MLOM OptionRom: | Enabled ▼ |
| Front NVME 1 OptionRom: | Enabled ▼ |
| MRAID Link Speed: | Auto ▼ |
| PCIe Slot 1 Link Speed: | Auto ▼ |
| Front NVME 1 Link Speed: | Auto ▼ |
| VGA Priority: | Onboard ▼ |
| P-SATA OptionROM: | LSI SW RAID ▼ |
| USB Port Rear: | Enabled ▼ |
| USB Port Internal: | Enabled ▼ |
| IPv6 PXE Support: | Disabled ▼ |

| | |
|------------------------------|------------|
| Legacy USB Support: | Enabled ▼ |
| Intel VTD coherency support: | Disabled ▼ |
| All Onboard LOM Ports: | Enabled ▼ |
| LOM Port 2 OptionRom: | Enabled ▼ |
| Pcie Slot 2 OptionRom: | Disabled ▼ |
| MRAID OptionRom: | Enabled ▼ |
| Front NVME 2 OptionRom: | Enabled ▼ |
| MLOM Link Speed: | Auto ▼ |
| PCIe Slot 2 Link Speed: | Auto ▼ |
| Front NVME 2 Link Speed: | Auto ▼ |
| M.2 SATA OptionROM: | AHCI ▼ |
| USB Port Front: | Enabled ▼ |
| USB Port KVM: | Enabled ▼ |
| USB Port:M.2 Storage: | Enabled ▼ |


Configure Cisco VIC1457 for iSCSI boot

The following configuration steps are for the Cisco VIC 1457 for iSCSI boot.



The default port-channeling between ports 0, 1, 2, and 3 must be turned off before the four individual ports can be configured. If port channeling is not turned off, only two ports appear for the VIC 1457. Complete the following steps to enable the port channel on the CIMC:

1. Under the networking tab, click the Adapter Card MLOM.
2. Under the General tab, uncheck the port channel.
3. Save the changes and reboot the CIMC.



Create iSCSI vNICs

To create iSCSI vNICs, complete the following steps:

1. Under the networking tab, click Adapter Card MLOM.
2. Click Add vNIC to create a vNIC.
3. In the Add vNIC section, enter the following settings:
 - Name: eth1
 - CDN Name: iSCSI-vNIC-A
 - MTU: 9000
 - Default VLAN: <<var_iscsi_vlan_a>>
 - VLAN Mode: TRUNK
 - Enable PXE boot: Check
4. Click Add vNIC and then click OK.
5. Repeat the process to add a second vNIC:
 - Name the vNIC eth3.
 - CDN Name: iSCSI-vNIC-B
 - Enter <<var_iscsi_vlan_b>> as the VLAN.
 - Set the uplink port to 3.

▼ General

Name:

CDN:

MTU: (1500 - 9000)

Uplink Port: ▼

MAC Address: ☐ Auto
☒

Class of Service: (0 - 6)

Trust Host CoS: ☐

PCI Order: (0 - 7)

Default VLAN: ☐ None
☒ ?

6. Select the vNIC eth1 on the left.

General External Ethernet Interfaces **vNICs** vHBAs

▼ vNICs

- eth0
- eth1**
- eth2
- eth3

► vNIC Properties

▼ iSCSI Boot Properties

► General

▼ Initiator

Name: (0 - 222) chars

IP Address:

Subnet Mask:

Gateway:

Primary DNS:

► Primary Target

► Secondary Target

Unconfigure iSCSI Boot

7. Under iSCSI Boot Properties, enter the initiator details:

- Name: <<var_ucsa_initiator_name_a>>
- IP address: <<var_esxi_hostA_iscsiA_ip>>
- Subnet mask: <<var_esxi_hostA_iscsiA_mask>>
- Gateway: <<var_esxi_hostA_iscsiA_gateway>>

The screenshot shows the vSphere configuration interface for iSCSI Boot Properties. On the left, a sidebar lists vNICs: eth0, eth1 (selected), eth2, and eth3. The main panel is titled 'vNIC Properties' and contains a section for 'iSCSI Boot Properties'. This section is divided into 'General' and 'Initiator' sub-sections. The 'Initiator' section has fields for Name, IP Address, Subnet Mask, Gateway, and Primary DNS. To the right of these fields are additional settings: Initiator Priority (set to 'primary'), Secondary DNS, TCP Timeout (set to 15), CHAP Name, and CHAP Secret. Below the 'Initiator' section are sections for 'Primary Target' and 'Secondary Target', each with fields for Name, IP Address, and TCP Port. At the bottom of the 'Initiator' section is a blue button labeled 'Unconfigure iSCSI Boot'.

8. Enter the primary target details:

- Name: IQN number of infra-SVM
- IP address: IP address of iscsi_lif01a
- Boot LUN: 0

9. Enter the secondary target details:

- Name: IQN number of infra-SVM
- IP address: IP address of iscsi_lif02a
- Boot LUN:0



You can obtain the storage IQN number by running the `vserver iscsi show` command.



Be sure to record the IQN names for each vNIC. You need them for a later step. In addition, the IQN names for initiators must be unique for each server and for the iSCSI vNIC.

10. Click Save Changes.

11. Select the vNIC eth3 and click the iSCSI Boot button located on the top of the Host Ethernet Interfaces section.

12. Repeat the process to configure eth3.

13. Enter the initiator details:

- Name: <<var_ucsa_initiator_name_b>>
- IP address: <<var_esxi_hostb_iscsib_ip>>
- Subnet mask: <<var_esxi_hostb_iscsib_mask>>
- Gateway: <<var_esxi_hostb_iscsib_gateway>>

14. Enter the primary target details:

- Name: IQN number of infra-SVM
- IP address: IP address of iscsi_lif01b
- Boot LUN: 0

15. Enter the secondary target details:

- Name: IQN number of infra-SVM
- IP address: IP address of iscsi_lif02b
- Boot LUN: 0



You can obtain the storage IQN number by using the `vserver iscsi show` command.



Be sure to record the IQN names for each vNIC. You need them for a later step.

16. Click Save Changes.

17. Repeat this process to configure iSCSI boot for Cisco UCS server B.

Configure vNICs for ESXi

To configure vNICs for ESXi, complete the following steps:

1. From the CIMC interface browser window, click Inventory and then click Cisco VIC adapters on the right pane.

2. Under Networking > Adapter Card MLOM, select vNICs tab and then select the vNICs underneath.
3. Select eth0 and click Properties.
4. Set the MTU to 9000. Click Save Changes.
5. Set the VLAN to native VLAN 2.

The screenshot shows the Cisco Integrated Management Controller (CIMC) interface. The breadcrumb path is "/ ... / Adapter Card MLOM / vNICs". The "vNICs" tab is selected. On the left, a list of vNICs includes eth0, eth1, eth2, and eth3. The "vNIC Properties" section for eth0 is expanded, showing the "General" tab. The configuration fields are as follows:

- Name: eth0
- CDN: VIC-MLOM-eth0
- MTU: 9000 (range 1500 - 9000)
- Uplink Port: 0
- MAC Address: ☐ Auto, ☒ F8:0F:6F:89:26:CE
- Class of Service: 0 (range 0 - 6)
- Trust Host CoS: ☐
- PCI Order: 0 (range 0 - 7)
- Default VLAN: ☐ None, ☒ 2

6. Repeat steps 3 and 4 for eth1, verifying that the uplink port is set to 1 for eth1.

The screenshot shows the Cisco Integrated Management Controller (CIMC) interface. The breadcrumb path is "/ ... / Adapter Card MLOM / vNICs". The "vNICs" tab is selected. On the left, a list of vNICs includes eth0, eth1, eth2, and eth3. The "Host Ethernet Interfaces" table is displayed, showing the configuration for all vNICs.

| Name | CDN | MAC Address | MTU | usNIC | Uplink Port | CoS | VLAN | VLAN Mode | ISCSI Boot | PXE Boot | Channel | Port Profile | Uplink Failover |
|-------------------------------|-------------|-------------------|------|-------|-------------|-----|------|-----------|------------|----------|---------|--------------|-----------------|
| <input type="checkbox"/> eth0 | VIC-MLO... | F8:0F:6F:89:26:CE | 9000 | 0 | 0 | 0 | 2 | TRUNK | disabled | enabled | N/A | N/A | N/A |
| <input type="checkbox"/> eth1 | VIC-ISCS... | F8:0F:6F:89:26:CF | 9000 | 0 | 1 | 0 | 3439 | TRUNK | enabled | enabled | N/A | N/A | N/A |
| <input type="checkbox"/> eth2 | VIC-MLO... | F8:0F:6F:89:26:D0 | 9000 | 0 | 2 | 0 | 2 | TRUNK | disabled | enabled | N/A | N/A | N/A |
| <input type="checkbox"/> eth3 | VIC-ISCS... | F8:0F:6F:89:26:D1 | 9000 | 0 | 3 | 0 | 3440 | TRUNK | enabled | enabled | N/A | N/A | N/A |



This procedure must be repeated for each initial Cisco UCS server node and each additional Cisco UCS server node added to the environment.

Next: [NetApp AFF storage deployment procedure \(part 2\)](#).

NetApp AFF storage deployment procedure (part 2)

Set up ONTAP SAN boot storage

Create iSCSI igroups



You need the iSCSI initiator IQNs from the server configuration for this step.

To create igroups, run the following commands from the cluster management node SSH connection. To view the three igroups created in this step, run the `igroup show` command.

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-A -protocol iscsi
-ostype vmware -initiator <<var_vm_host_infra_a_iSCSI-
A_vNIC_IQN>>,<<var_vm_host_infra_a_iSCSI-B_vNIC_IQN>>
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-B -protocol iscsi
-ostype vmware -initiator <<var_vm_host_infra_b_iSCSI-
A_vNIC_IQN>>,<<var_vm_host_infra_b_iSCSI-B_vNIC_IQN>>
```



This step must be completed when adding additional Cisco UCS C-Series servers.

Map boot LUNs to igroups

To map boot LUNs to igroups, run the following commands from the cluster management SSH connection:

```
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-A -igroup
VM-Host-Infra-A -lun-id 0
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-B -igroup
VM-Host-Infra-B -lun-id 0
```



This step must be completed when adding additional Cisco UCS C-Series servers.

[Next: VMware vSphere 6.7U2 deployment procedure.](#)

VMware vSphere 6.7U2 deployment procedure

This section provides detailed procedures for installing VMware ESXi 6.7U2 in a FlexPod Express configuration. The deployment procedures that follow are customized to include the environment variables described in previous sections.

Multiple methods exist for installing VMware ESXi in such an environment. This procedure uses the virtual KVM console and virtual media features of the CIMC interface for Cisco UCS C-Series servers to map remote installation media to each individual server.



This procedure must be completed for Cisco UCS server A and Cisco UCS server B.




This procedure must be completed for any additional nodes added to the cluster.

Log in to CIMC interface for Cisco UCS C-Series standalone servers

The following steps detail the method for logging in to the CIMC interface for Cisco UCS C-Series standalone servers. You must log in to the CIMC interface to run the virtual KVM, which enables the administrator to begin installation of the operating system through remote media.

All hosts

1. Navigate to a web browser and enter the IP address for the CIMC interface for the Cisco UCS C-Series. This step launches the CIMC GUI application.
2. Log in to the CIMC UI using the admin user name and credentials.
3. In the main menu, select the Server tab.
4. Click Launch KVM Console.



5. From the virtual KVM console, select the Virtual Media tab.
6. Select Map CD/DVD.



You might first need to click Activate Virtual Devices. Select Accept This Session if prompted.

7. Browse to the VMware ESXi 6.7U2 installer ISO image file and click Open. Click Map Device.
8. Select the Power menu and choose Power Cycle System (Cold Boot). Click Yes.

Install VMware ESXi

The following steps describe how to install VMware ESXi on each host.

Download ESXi 6.7U2 Cisco custom image

1. Navigate to the [VMware vSphere download page](#) for custom ISOs.
2. Click Go to Downloads next to the Cisco Custom Image for the ESXi 6.7U2 Install CD.
3. Download the Cisco Custom Image for the ESXi 6.7U2 Install CD (ISO).
4. When the system boots, the machine detects the presence of the VMware ESXi installation media.
5. Select the VMware ESXi installer from the menu that appears. The installer loads, which can take several minutes.
6. After the installer has finished loading, press Enter to continue with the installation.
7. After reading the end-user license agreement, accept it and continue with the installation by pressing F11.
8. Select the NetApp LUN that was previously set up as the installation disk for ESXi, and press Enter to continue with the installation.



9. Select the appropriate keyboard layout and press Enter.
10. Enter and confirm the root password and press Enter.
11. The installer warns you that existing partitions are removed on the volume. Continue with the installation by pressing F11. The server reboots after the installation of ESXi.

Set up VMware ESXi host management networking

The following steps describe how to add the management network for each VMware ESXi host.

All hosts

1. After the server has finished rebooting, enter the option to customize the system by pressing F2.
2. Log in with root as the login name and the root password previously entered during the installation process.
3. Select the Configure Management Network option.
4. Select Network Adapters and press Enter.
5. Select the desired ports for vSwitch0. Press Enter.
6. Select the ports that correspond to eth0 and eth1 in CIMC.

Network Adapters

Select the adapters for this host's default management network connection. Use two or more adapters for fault-tolerance and load-balancing.

| Device Name | Hardware Label (MAC Address) | Status |
|--|-------------------------------|-----------------|
| <input type="checkbox"/> vmnic0 | LOM Port 1 (...:5a:b5:8d:6e) | Connected |
| <input type="checkbox"/> vmnic1 | LOM Port 2 (...:5a:b5:8d:6f) | Disconnected |
| <input checked="" type="checkbox"/> vmnic2 | VIC-MLOM-eth0 (...:70:6c:cc) | Connected (...) |
| <input type="checkbox"/> vmnic3 | VIC-iSCSI-A (...:3c:70:6c:cd) | Connected (...) |
| <input checked="" type="checkbox"/> vmnic4 | VIC-MLOM-eth2 (...:70:6c:ce) | Connected (...) |
| <input type="checkbox"/> vmnic5 | VIC-iSCSI-B (...:3c:70:6c:cf) | Connected (...) |

<D> View Details <Space> Toggle Selected <Enter> OK <Esc> Cancel

7. Select VLAN (optional) and press Enter.
8. Enter the VLAN ID <<mgmt_vlan_id>>. Press Enter.
9. From the Configure Management Network menu, select IPv4 Configuration to configure the IP address of the management interface. Press Enter.
10. Use the arrow keys to highlight Set Static IPv4 Address and use the space bar to select this option.
11. Enter the IP address for managing the VMware ESXi host <<esxi_host_mgmt_ip>>.
12. Enter the subnet mask for the VMware ESXi host <<esxi_host_mgmt_netmask>>.
13. Enter the default gateway for the VMware ESXi host <<esxi_host_mgmt_gateway>>.
14. Press Enter to accept the changes to the IP configuration.
15. Enter the IPv6 configuration menu.
16. Use the space bar to disable IPv6 by unselecting the Enable IPv6 (restart required) option. Press Enter.
17. Enter the menu to configure the DNS settings.
18. Because the IP address is assigned manually, the DNS information must also be entered manually.
19. Enter the primary DNS server's IP address <<nameserver_ip>>.
20. (Optional) Enter the secondary DNS server's IP address.
21. Enter the FQDN for the VMware ESXi host name: <<esxi_host_fqdn>>.
22. Press Enter to accept the changes to the DNS configuration.
23. Exit the Configure Management Network submenu by pressing Esc.
24. Press Y to confirm the changes and reboot the server.
25. Select Troubleshooting Options, and then Enable ESXi Shell and SSH.



These troubleshooting options can be disabled after the validation pursuant to the customer's security policy.

26. Press Esc twice to return to the main console screen.
27. Click Alt-F1 from the CIMC Macros > Static Macros > Alt-F drop-down menu at the top of the screen.
28. Log in with the proper credentials for the ESXi host.
29. At the prompt, enter the following list of esxcli commands sequentially to enable network connectivity.

```
esxcli network vswitch standard policy failover set -v vSwitch0 -a
vmnic2,vmnic4 -l iphash
```

Configure ESXi host

Use the information in the following table to configure each ESXi host.

| Detail | Detail value |
|------------------------------|-------------------------------|
| ESXi host name | <<esxi_host_fqdn>> |
| ESXi host management IP | <<esxi_host_mgmt_ip>> |
| ESXi host management mask | <<esxi_host_mgmt_netmask>> |
| ESXi host management gateway | <<esxi_host_mgmt_gateway>> |
| ESXi host NFS IP | <<esxi_host_NFS_ip>> |
| ESXi host NFS mask | <<esxi_host_NFS_netmask>> |
| ESXi host NFS gateway | <<esxi_host_NFS_gateway>> |
| ESXi host vMotion IP | <<esxi_host_vMotion_ip>> |
| ESXi host vMotion mask | <<esxi_host_vMotion_netmask>> |
| ESXi host vMotion gateway | <<esxi_host_vMotion_gateway>> |
| ESXi host iSCSI-A IP | <<esxi_host_iSCSI-A_ip>> |
| ESXi host iSCSI-A mask | <<esxi_host_iSCSI-A_netmask>> |
| ESXi host iSCSI-A gateway | <<esxi_host_iSCSI-A_gateway>> |
| ESXi host iSCSI-B IP | <<esxi_host_iSCSI-B_ip>> |
| ESXi host iSCSI-B mask | <<esxi_host_iSCSI-B_netmask>> |
| ESXi host iSCSI-B gateway | <<esxi_host_iSCSI-B_gateway>> |

Log in to the ESXi host

To log in to the ESXi host, complete the following steps:

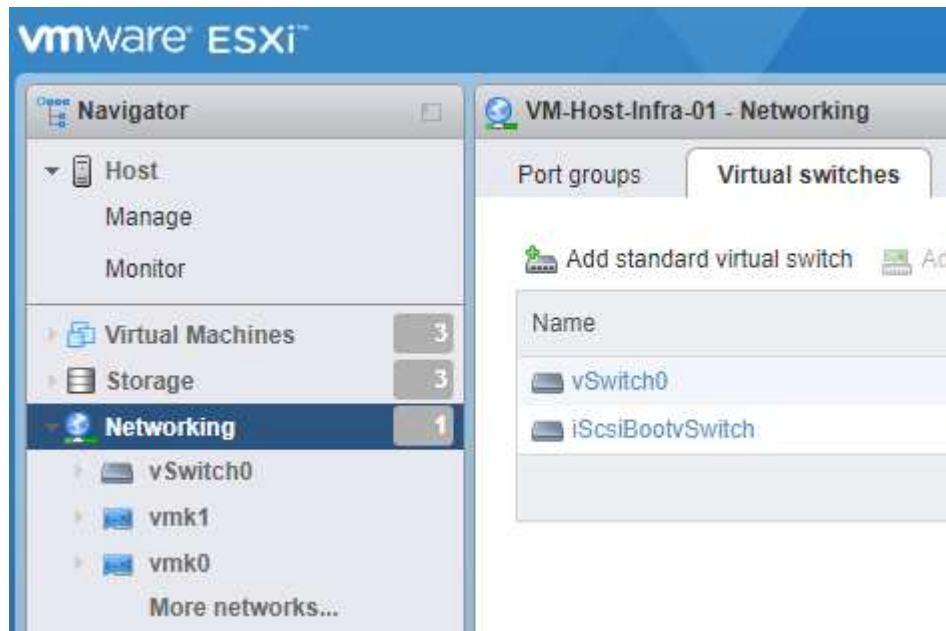
1. Open the host's management IP address in a web browser.
2. Log in to the ESXi host using the root account and the password you specified during the install process.

3. Read the statement about the VMware Customer Experience Improvement Program. After selecting the proper response, click OK.

Configure iSCSI boot

To configure iSCSI boot, complete the following steps:

1. Select Networking on the left.
2. On the right, select the Virtual Switches tab.



3. Click iScsiBootvSwitch.
4. Select Edit settings.
5. Change the MTU to 9000 and click Save.
6. Rename the iSCSIBootPG port to iSCSIBootPG-A.



Vmnic3 and vmnic5 are used for iSCSI boot in this configuration. If you have additional NICs in your ESXi host, you might have different vmnic numbers. To confirm which NICs are used for iSCSI boot, match the MAC addresses on the iSCSI vNICs in CIMC to the vmnics in ESXi.

7. In the center pane, select the VMkernel NICs tab.
8. Select Add VMkernel NIC.
 - a. Specify a new port group name of iScsiBootPG-B.
 - b. Select iScsiBootvSwitch for the virtual switch.
 - c. Enter <<iScsiB_vlan_id>> for the VLAN ID.
 - d. Change the MTU to 9000.
 - e. Expand IPv4 Settings.
 - f. Select Static Configuration.
 - g. Enter <<var_hosta_iScsiB_ip>> for Address.

h. Enter <<var_hosta_iscsib_mask>> for Subnet Mask.

i. Click Create.



Set the MTU to 9000 on iScsiBootPG-A.

9. To set the failover, complete the following steps:

- Click Edit Settings on iSCSIBootPG-A > Tiering and Failover > Failover Order > Vmnic3. Vmnic3 should be active and vmnic5 should be unused.
- Click Edit Settings on iSCSIBootPG-B > Teaming and Failover > Failover order > Vmnic5. Vmnic5 should be active and vmnic3 should be unused.

iScsiBootPG-A - Edit Settings

Properties

Security

Traffic shaping

Teaming and failover

Load balancing

Network failure detection

Notify switches

Failback

Failover order

☒ Override



Active adapters

vmnic3

Standby adapters

Unused adapters

vmnic5

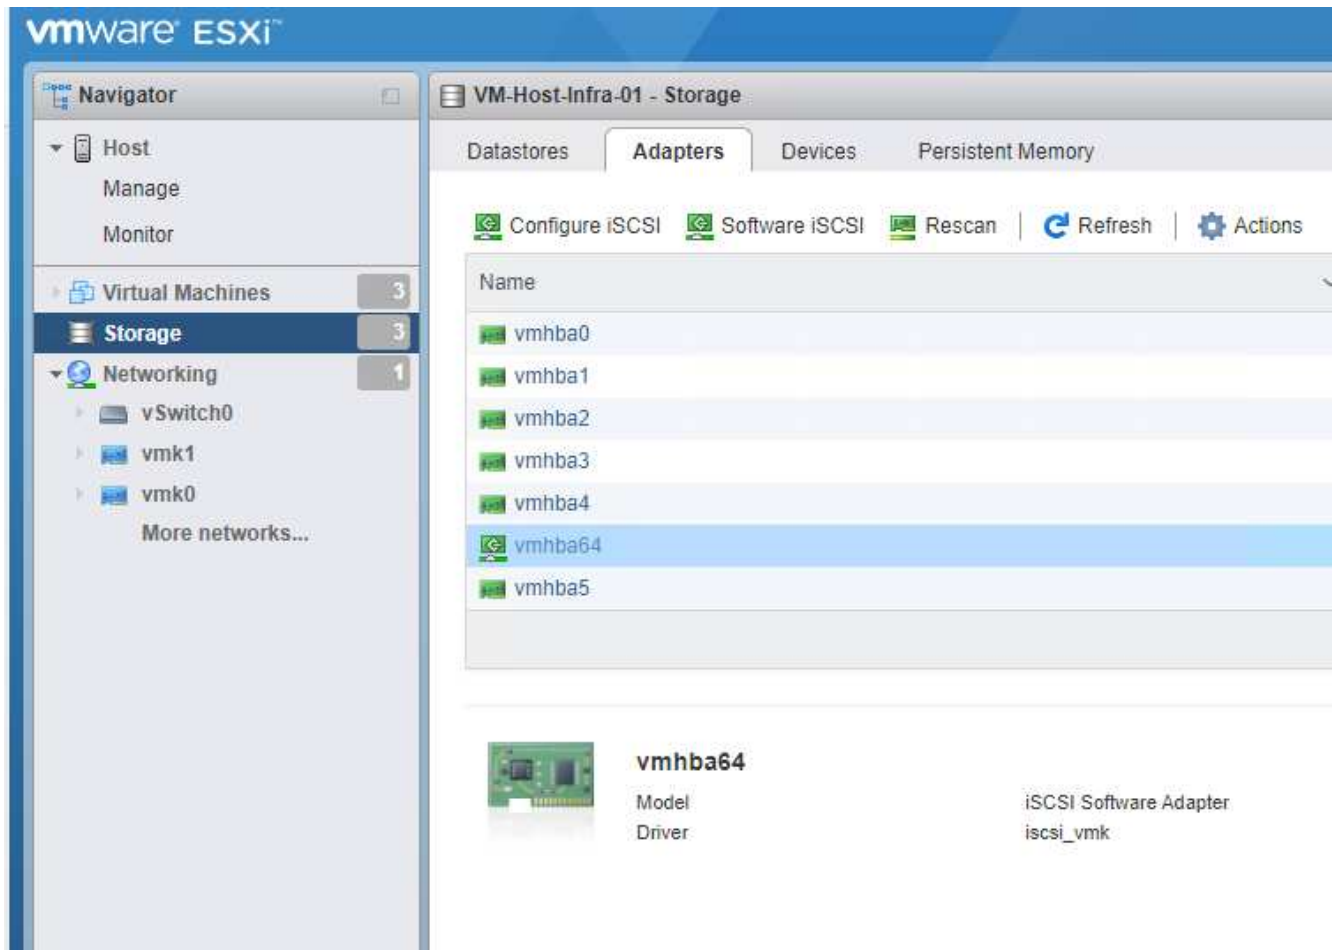
Select active and standby adapters

Configure iSCSI multipathing

To set up iSCSI multipathing on the ESXi hosts, complete the following steps:

- Select Storage in the left navigation pane. Click Adapters.

2. Select the iSCSI software adapter and click Configure iSCSI.



3. Under Dynamic Targets, click Add Dynamic Target.

Configure iSCSI - vmhba64

iSCSI enabled ☐ Disabled ☒ Enabled

▶ Name & alias `iqn.1992-01.com.cisco:ucsA-01`

▶ CHAP authentication Do not use CHAP

▶ Mutual CHAP authentication Do not use CHAP

▶ Advanced settings Click to expand

Network port bindings No port bindings

Static targets

Add static target Remove static target Edit settings Search

| Target | Address | Port |
|---|----------------|------|
| <code>iqn.1992-08.com.netapp.sn.e42fa6b2d2e011e9a68d00a098f...</code> | 172.21.183.105 | 3260 |
| <code>iqn.1992-08.com.netapp.sn.e42fa6b2d2e011e9a68d00a098f...</code> | 172.21.184.106 | 3260 |
| <code>iqn.1992-08.com.netapp.sn.e42fa6b2d2e011e9a68d00a098f...</code> | 172.21.183.106 | 3260 |
| <code>iqn.1992-08.com.netapp.sn.e42fa6b2d2e011e9a68d00a098f...</code> | 172.21.184.105 | 3260 |

Dynamic targets

Add dynamic target Remove dynamic target Edit settings Search

| Address | Port |
|----------------|------|
| 172.21.183.105 | 3260 |
| 172.21.184.105 | 3260 |
| 172.21.183.106 | 3260 |
| 172.21.184.106 | 3260 |

Save configuration Cancel

4. Enter the IP address `iscsi_lif01a`.
 - a. Repeat with the IP addresses `iscsi_lif01b`, `iscsi_lif02a`, and `iscsi_lif02b`.
 - b. Click Save Configuration.

Dynamic targets

Add dynamic target Remove dynamic target Edit settings Search

| Address | Port |
|----------------|------|
| 172.21.183.105 | 3260 |
| 172.21.184.105 | 3260 |
| 172.21.183.106 | 3260 |
| 172.21.184.106 | 3260 |

Save configuration Cancel



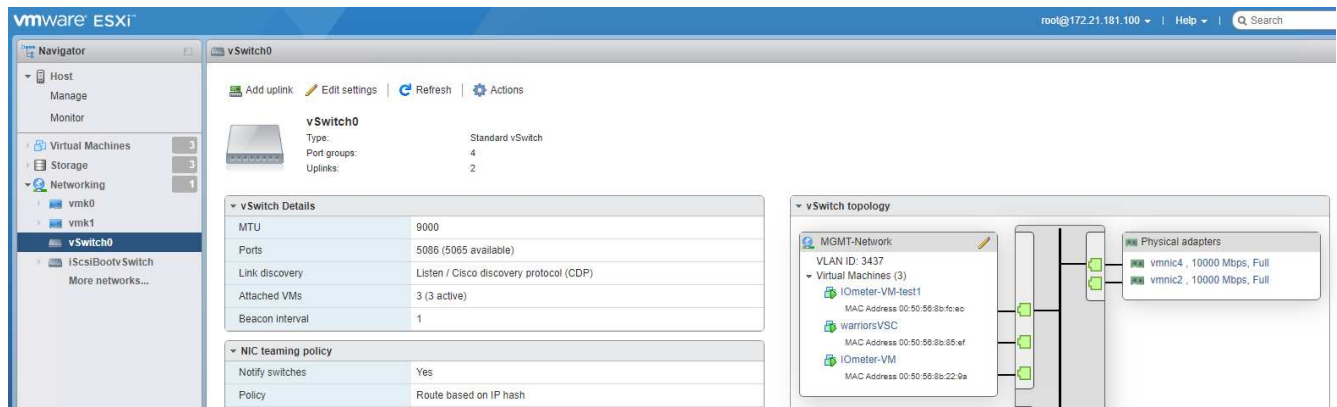
You can find the iSCSI LIF IP addresses by running the `network interface show` command on the NetApp cluster or by looking at the Network Interfaces tab in System Manager.

Configure the ESXi host

To configure ESXi boot, complete the following steps:

1. In the left navigation pane, select Networking.

2. Select vSwitch0.



3. Select Edit Settings.

4. Change the MTU to 9000.

5. Expand NIC Teaming and verify that both vmnic2 and vmnic4 are set to active and NIC Teaming and Failover is set to Route Based on IP Hash.



The IP hash method of load balancing requires the underlying physical switch to be properly configured using SRC-DST-IP EtherChannel with a static (mode- on) port channel. You might experience intermittent connectivity due to possible switch misconfiguration. If so, then temporarily shut down one of the two associated uplink ports on the Cisco switch to restore communication to the ESXi management vmkernel port while troubleshooting the port-channel settings.


Configure the port groups and VMkernel NICs

To configure the port groups and VMkernel NICs, complete the following steps:

1. In the left navigation pane, select Networking.
2. Right-click the Port Groups tab.



3. Right-click VM Network and select Edit. Change the VLAN ID to <<var_vm_traffic_vlan>>.
4. Click Add Port Group.
 - a. Name the port group MGMT-Network.
 - b. Enter <<mgmt_vlan>> for the VLAN ID.
 - c. Make sure that vSwitch0 is selected.
 - d. Click save.
5. Click the VMkernel NICs tab.



6. Select Add VMkernel NIC.
 - a. Select New Port Group.
 - b. Name the port group NFS-Network.
 - c. Enter <<nfs_vlan_id>> for the VLAN ID.
 - d. Change the MTU to 9000.
 - e. Expand IPv4 Settings.
 - f. Select Static Configuration.
 - g. Enter <<var_hosta_nfs_ip>> for Address.
 - h. Enter <<var_hosta_nfs_mask>> for Subnet Mask.
 - i. Click Create.
7. Repeat this process to create the vMotion VMkernel port.
8. Select Add VMkernel NIC.
 - a. Select New Port Group.
 - b. Name the port group vMotion.
 - c. Enter <<vmotion_vlan_id>> for the VLAN ID.
 - d. Change the MTU to 9000.
 - e. Expand IPv4 Settings.
 - f. Select Static Configuration.
 - g. Enter <<var_hosta_vmotion_ip>> for Address.
 - h. Enter <<var_hosta_vmotion_mask>> for Subnet Mask.
 - i. Make sure that the vMotion checkbox is selected after IPv4 Settings.

| | |
|-----------------|---|
| Virtual switch | vSwitch0 |
| VLAN ID | 3441 |
| MTU | 9000 |
| IP version | IPv4 only |
| ▼ IPv4 settings | |
| Configuration | <input type="radio"/> DHCP <input checked="" type="radio"/> Static |
| Address | 172.21.185.63 |
| Subnet mask | 255.255.255.0 |
| TCP/IP stack | Default TCP/IP stack |
| Services | <input checked="" type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication |

Create Cancel




There are many ways to configure ESXi networking, including by using the VMware vSphere distributed switch if your licensing allows it. Alternative network configurations are supported in FlexPod Express if they are required to meet business requirements.


Mount the first datastores

The first datastores to be mounted are the `infra_datastore` datastore for VMs and the `infra_swap` datastore for VM swap files.

1. Click Storage in the left navigation pane, and then click New Datastore.



2. Select Mount NFS Datastore.



3. Enter the following information in the Provide NFS Mount Details page:

- Name: `infra_datastore`
- NFS server: `<<var_nodea_nfs_lif>>`
- Share: `/infra_datastore`
- Make sure that NFS 3 is selected.

4. Click Finish. You can see the task completing in the Recent Tasks pane.

5. Repeat this process to mount the `infra_swap` datastore:


- Name: `infra_swap`
- NFS server: `<<var_nodea_nfs_lif>>`
- Share: `/infra_swap`

- Make sure that NFS 3 is selected.

Configure NTP

To configure NTP for an ESXi host, complete the following steps:


1. Click Manage in the left navigation pane. Select System in the right pane and then click Time & Date.
2. Select Use Network Time Protocol (Enable NTP Client).
3. Select Start and Stop with Host as the NTP service startup policy.
4. Enter <<var_ntp>> as the NTP server. You can set multiple NTP servers.
5. Click Save.



Move the VM swap file location

These steps provide details for moving the VM swap file location.

1. Click Manage in the left navigation pane. Select system in the right pane, then click Swap.



2. Click Edit Settings. Select infra_swap from the Datastore options.

| Configuration Option | Value |
|----------------------|---|
| Enabled | <input checked="" type="radio"/> Yes <input type="radio"/> No |
| Datastore | infra_swap |
| Local swap enabled | <input checked="" type="radio"/> Yes <input type="radio"/> No |
| Host cache enabled | <input checked="" type="radio"/> Yes <input type="radio"/> No |

Buttons: Save, Cancel

3. Click Save.

Next: [VMware vCenter Server 6.7U2 installation procedure.](#)

VMware vCenter Server 6.7U2 installation procedure

This section provides detailed procedures for installing VMware vCenter Server 6.7 in a FlexPod Express configuration.




FlexPod Express uses the VMware vCenter Server Appliance (VCSA).


Download the VMware vCenter Server Appliance

To download the VMware vCenter Server Appliance (VCSA), complete the following steps:

1. Download the VCSA. Access the download link by clicking the Get vCenter Server icon when managing the ESXi host.
2. Download the VCSA from the VMware site.
3. Although the Microsoft Windows vCenter Server installable is supported, VMware recommends the VCSA for new deployments.
4. Mount the ISO image.
5. Navigate to the `vcsa- ui-installer > win32` directory. Double-click `installer.exe`.
6. Click Install.
7. Click Next on the Introduction page.



8. Select Embedded Platform Services Controller as the deployment type.



If required, the External Platform Services Controller deployment is also supported as part of the FlexPod Express solution.

9. In the Appliance Deployment Target, enter the IP address of an ESXi host that you have deployed, the root user name, and the root password.

vCenter Server Appliance Installer

Installer

vm Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Appliance deployment target

Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.

| | | |
|----------------------------------|----------------|---|
| ESXi host or vCenter Server name | 172.21.181.100 | ? |
| HTTPS port | 443 | |
| User name | root | ? |
| Password | | |

CANCEL BACK NEXT

Activate Windows
Go to System in Settings

10. Set the appliance VM by entering VCSA as the VM name and the root password that you would like to use for the VCSA.

vCenter Server Appliance Installer

Installer

vm Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Set up appliance VM

Specify the VM settings for the appliance to be deployed.

VM name ⓘ


Set root password ⓘ

Confirm root password

CANCEL BACK NEXT

Activate Windows
Go to System in Centre

11. Select the deployment size that best fits your environment. Click Next.



12. Select the `infra_datastore` datastore. Click Next.
13. Enter the following information in the Configure network settings page and click Next.
 - a. Select MGMT-Network for Network.
 - b. Enter the FQDN or IP to be used for the VCSA.
 - c. Enter the IP address to be used.
 - d. Enter the subnet mask to be used.
 - e. Enter the default gateway.
 - f. Enter the DNS server.
14. On the Ready to Complete Stage 1 page, verify that the settings you have entered are correct. Click Finish.

Installer

vCenter Server Appliance Installer

vm Install - Stage 1: Deploy vCenter Server Appliance with an Embedded Platform Services Controller

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Configure network settings


Configure network settings for this appliance

| | | |
|------------------------------|-----------------------------|---|
| Network | MGMT-Network | ⓘ |
| IP version | IPv4 | |
| IP assignment | static | |
| FQDN | FlexPod-VCSA.cie.netapp.com | ⓘ |
| IP address | 172.21.181.105 | |
| Subnet mask or prefix length | 255.255.255.0 | ⓘ |
| Default gateway | 172.21.181.1 | |
| DNS servers | 10.61.184.251,10.61.184.252 | |
| Common Ports | | |
| HTTP | 80 | |
| HTTPS | 443 | |

CANCEL BACK NEXT


Activate Windows
Go to System in Control

15. Review your settings on stage 1 before starting the appliance deployment.




The VCSA installs now. This process takes several minutes.

- After stage 1 completes, a message appears stating that it has completed. Click Continue to begin stage 2 configuration.
- On the Stage 2 Introduction page, click Next.




18. Enter <<var_ntp_id>> for the NTP server address. You can enter multiple NTP IP addresses.
19. If you plan to use vCenter Server high availability (HA), make sure that SSH access is enabled.
20. Configure the SSO domain name, password, and site name. Click Next.




Record these values for your reference, especially if you deviate from the `vsphere.local` domain name.

21. Join the VMware Customer Experience Program if desired. Click Next.



22. View the summary of your settings. Click Finish or use the back button to edit settings.
23. A message appears stating that you will not be able to pause or stop the installation from completing after it has started. Click OK to continue.



The appliance setup continues. This takes several minutes.

A message appears indicating that the setup was successful.

24. The links that the installer provides to access vCenter Server are clickable.

[Next: VMware vCenter Server 6.7U2 and vSphere clustering configuration.](#)

VMware vCenter Server 6.7U2 and vSphere clustering configuration

To configure VMware vCenter Server 6.7 and vSphere clustering, complete the following steps:

1. Navigate to `https://<FQDN or IP of vCenter>/vsphere-client/`.
2. Click Launch vSphere Client.
3. Log in with the user name `administrator@vsphere.local` and the SSO password you entered during the VCSA setup process.
4. Right-click the vCenter name and select New Datacenter.
5. Enter a name for the data center and click OK.

Create a vSphere cluster

To create a vSphere cluster, complete the following steps:

1. Right-click the newly created data center and select New Cluster.
2. Enter a name for the cluster.
3. Enable DR and vSphere HA by selecting the checkboxes.
4. Click OK.

New Cluster | FlexPod-Datacenter

| | |
|------------|-------------------------------------|
| Name | FlexPod-Cluster |
| Location | FlexPod-Datacenter |
| DRS | <input checked="" type="checkbox"/> |
| vSphere HA | <input checked="" type="checkbox"/> |
| vSAN | <input type="checkbox"/> |


These services will have default settings - these can be changed later in the Cluster Quickstart workflow.

CANCEL **OK**

Add the ESXi hosts to the cluster

To add the ESXi hosts to the cluster, complete the following steps:

1. Right-click the cluster and select Add Host.



2. To add an ESXi host to the cluster, complete the following steps:
 - a. Enter the IP or FQDN of the host. Click Next.
 - b. Enter the root user name and password. Click Next.
 - c. Click Yes to replace the host's certificate with a certificate signed by the VMware certificate server.
 - d. Click Next on the Host Summary page.
 - e. Click the green + icon to add a license to the vSphere host.
3. This step can be completed later if desired.
 - a. Click Next to leave lockdown mode disabled.
 - b. Click Next at the VM location page.
 - c. Review the Ready to Complete page. Use the back button to make any changes or select Finish.
4. Repeat steps 1 and 2 for Cisco UCS host B.



This process must be completed for any additional hosts added to the FlexPod Express configuration.

Configure coredump on the ESXi hosts

To configure coredump on the ESXi hosts, complete the following steps:

1. Log into [https:// vCenter IP:5480/](https://vCenter IP:5480/), enter root for the user name, and enter the root password.
2. Click on services and select VMware vSphere ESXI Dump collector.
3. Start the VMware vSphere ESXI Dump collector service.

← → ↻ ⚠ Not secure | 172.21.181.105:5480/ui/services

vm Appliance Management
 Mon 10-28-2019 06:51 AM UTC

Summary
 Monitor
 Access
 Networking
 Firewall
 Time
 Services
 Update
 Administration
 Syslog
 Backup

RESTART START STOP

| | Name | |
|----------------------------------|---|--|
| <input type="radio"/> | vSAN health Service | |
| <input type="radio"/> | VMware vSphere Web Client | |
| <input type="radio"/> | VMware vSphere Update Manager | |
| <input type="radio"/> | VMware vSphere Profile-Driven Storage Service | |
| <input checked="" type="radio"/> | VMware vSphere ESXi Dump Collector | |
| <input type="radio"/> | VMware vSphere Client | |
| <input type="radio"/> | VMware vSphere Authentication Proxy | |
| <input type="radio"/> | VMware vService Manager | |
| <input type="radio"/> | VMware vSAN Data Protection Service | |
| <input type="radio"/> | VMware vCenter-Services | |
| <input type="radio"/> | VMware vCenter Server | |
| <input type="radio"/> | VMware vCenter High Availability | |
| <input type="radio"/> | VMware Topology Service | |

- Using SSH, connect to the management IP ESXi host, enter root for the user name, and enter the root password.
- Run the following commands:

```
esxcli system coredump network set -i ip_address_of_core_dump_collector
-v vmk0 -o 6500
esxcli system coredump network set --enable=true
esxcli system coredump network check
```

- The message Verified the configured netdump server is running appears after you enter the final command.

```
root@VM-Host-Infra-01:~] esxcli system coredump network set -i 172.21.181.105 -
vmk0 -o 6500
root@VM-Host-Infra-01:~]
root@VM-Host-Infra-01:~] esxcli system coredump network set --enable=true
root@VM-Host-Infra-01:~] esxcli system coredump network check
Verified the configured netdump server is running
```



This process must be completed for any additional hosts added to FlexPod Express.



`ip_address_of_core_dump_collector` in this validation is the vCenter IP.

Next: [NetApp Virtual Storage Console 9.6 deployment procedures](#).


NetApp Virtual Storage Console 9.6 deployment procedures

This section describes the deployment procedures for the NetApp Virtual Storage Console (VSC).

Install Virtual Storage Console 9.6

To install the VSC 9.6 software by using an Open Virtualization Format (OVF) deployment, follow these steps:

1. Go to vSphere Web Client > Host Cluster > Deploy OVF Template.
2. Browse to the VSC OVF file downloaded from the NetApp Support site.



3. Enter the VM name and select a datacenter or folder in which to deploy. Click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- 5 License agreements
- ✓ 6 Select storage
- 7 Select networks
- 8 Customize template

Select a name and folder

Specify a unique name and target location

Virtual machine name:

Select a location for the virtual machine.

- ▼  warriorsvcsa.cie.netapp.com
- >  FlexPod-Datacenter

4. Select the FlexPod-Cluster ESXi cluster and click Next.
5. Review the details and click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- 5 License agreements
- 6 Select storage
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

Review details

Verify the template details.

| | |
|---------------|---|
| Publisher | No certificate present |
| Product | Virtual Appliance - NetApp VSC, VASA Provider and SRA for ONTAP |
| Version | See appliance for version |
| Vendor | NetApp Inc. |
| Description | Virtual Appliance - NetApp VSC, VASA Provider, and SRA virtual appliance for NetApp storage systems. For more information or support please visit http://www.netapp.com/ |
| Download size | 1.0 GB |
| Size on disk | 2.1 GB (thin provisioned) |
| | 53.0 GB (thick provisioned) |

CANCEL

BACK

NEXT

6. Click Accept to accept the license and click Next.
7. Select the Thin Provision virtual disk format and one of the NFS datastores. Click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- 6 Select storage**
- 7 Select networks
- 8 Customize template
- 9 Ready to complete

Select storage

Select the storage for the configuration and disk files

☐ Encrypt this virtual machine (Requires Key Management Server)

Select virtual disk format: Thin Provision

VM Storage Policy: Datastore Default

| Name | Capacity | Provisioned | Free | Type |
|------------------|----------|-------------|-----------|------|
| infra_datastore | 75 GB | 360 KB | 75 GB | NF |
| infra_datastore1 | 475 GB | 639.9 GB | 276.86 GB | NF |
| infra_swap (1) | 100 GB | 4.98 GB | 95.02 GB | NF |

Compatibility

✓ Compatibility checks succeeded.

CANCEL

BACK

NEXT

8. From Select Networks, choose a destination network and click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- 7 Select networks**
- 8 Customize template
- 9 Ready to complete

Select networks

Select a destination network for each source network.

| Source Network | Destination Network |
|----------------|---------------------|
| nat | MGMT-Network |
| 1 items | |

IP Allocation Settings

IP allocation: Static - Manual

IP protocol: IPv4

CANCEL

BACK

NEXT

9. From Customize Template, enter the VSC administrator password, vCenter name or IP address, and other configuration details and click Next.

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 License agreements
- ✓ 6 Select storage
- ✓ 7 Select networks
- ✓ **8 Customize template**
- 9 Ready to complete

vCenter Server Address (*)
Specify the IP address/hostname of an existing vCenter to register to.
172.21.181.105

Port (*)
Specify the HTTPS port of an existing vCenter to register to.
443

Username (*)
Specify the username of an existing vCenter to register to.
administrator@vsphere.local

Password (*)
Specify the password of an existing vCenter to register to.
Password
Confirm Password

✓ Network Properties 8 settings

Host Name
Specify the hostname for the appliance. (Leave blank if DHCP is desired)

CANCEL

BACK

NEXT

10. Review the configuration details entered and click Finish to complete the deployment of NetApp-VSC VM.
11. Power on the NetApp-VSC VM and open the VM console.
12. During the NetApp-VSC VM boot process, you see a prompt to install VMware Tools. From vCenter, select NetApp-VSC VM > Guest OS > Install VMware Tools.

Booting VSC, VASA Provider, and SRA virtual appliance...Please wait...

VMware Tools OVF vCenter configuration not found.

VMware Tools OVF vCenter configuration not found.

VMware Tools OVF vCenter configuration not found.

VMware Tools installation

Before you can continue the VSC, VASA Provider, and SRA virtual appliance installation, you must install the VMware Tools:

1. Select VM > Guest OS > Install VMware Tools.


OR

Click on "Install VMware Tools" pop-up box on the vSphere Web Client.

2. Follow the prompts provided by the VMware Tools wizard.

Once you click on mount, the installation process will automatically continue.

13. Networking configuration and vCenter registration information was provided during OVF template customization. Therefore, after the NetApp-VSC VM is running, VSC, vSphere API for Storage Awareness (VASA), and VMware Storage Replication Adapter (SRA) are registered with vCenter.
14. Log out of the vCenter Client and log in again. From the Home menu, confirm that the NetApp VSC is installed.

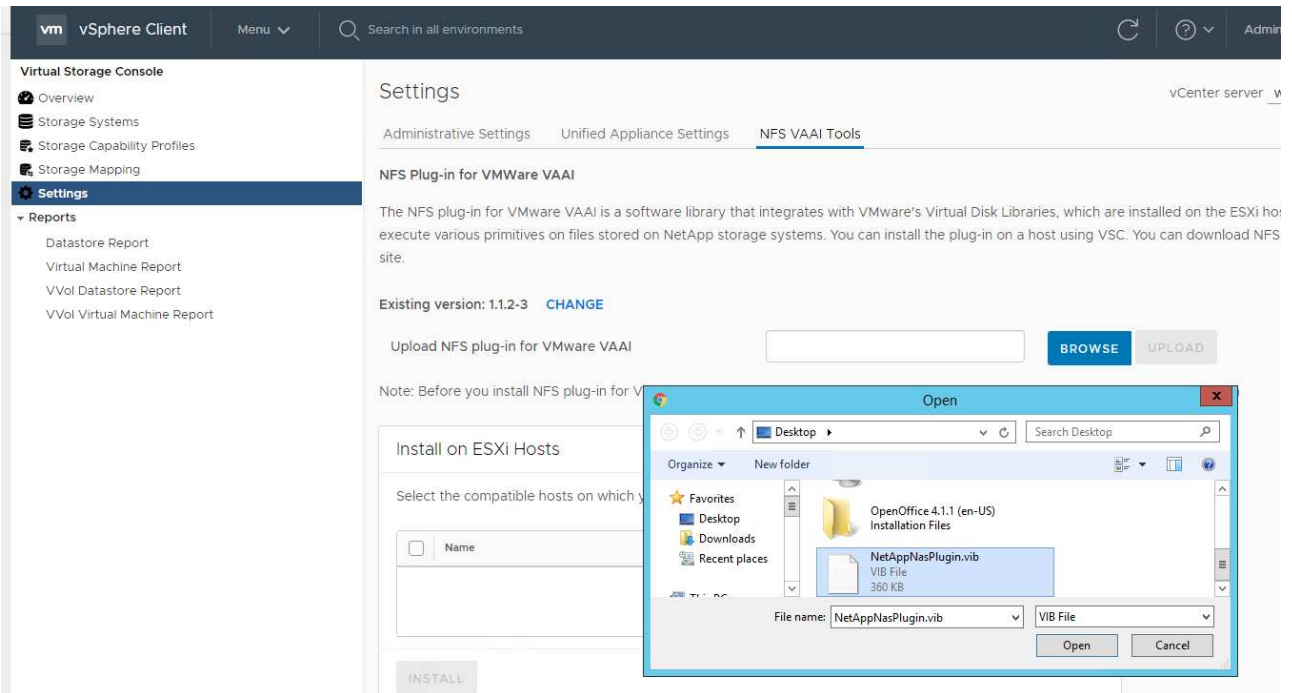


Download and install the NetApp NFS VAAI Plug-In

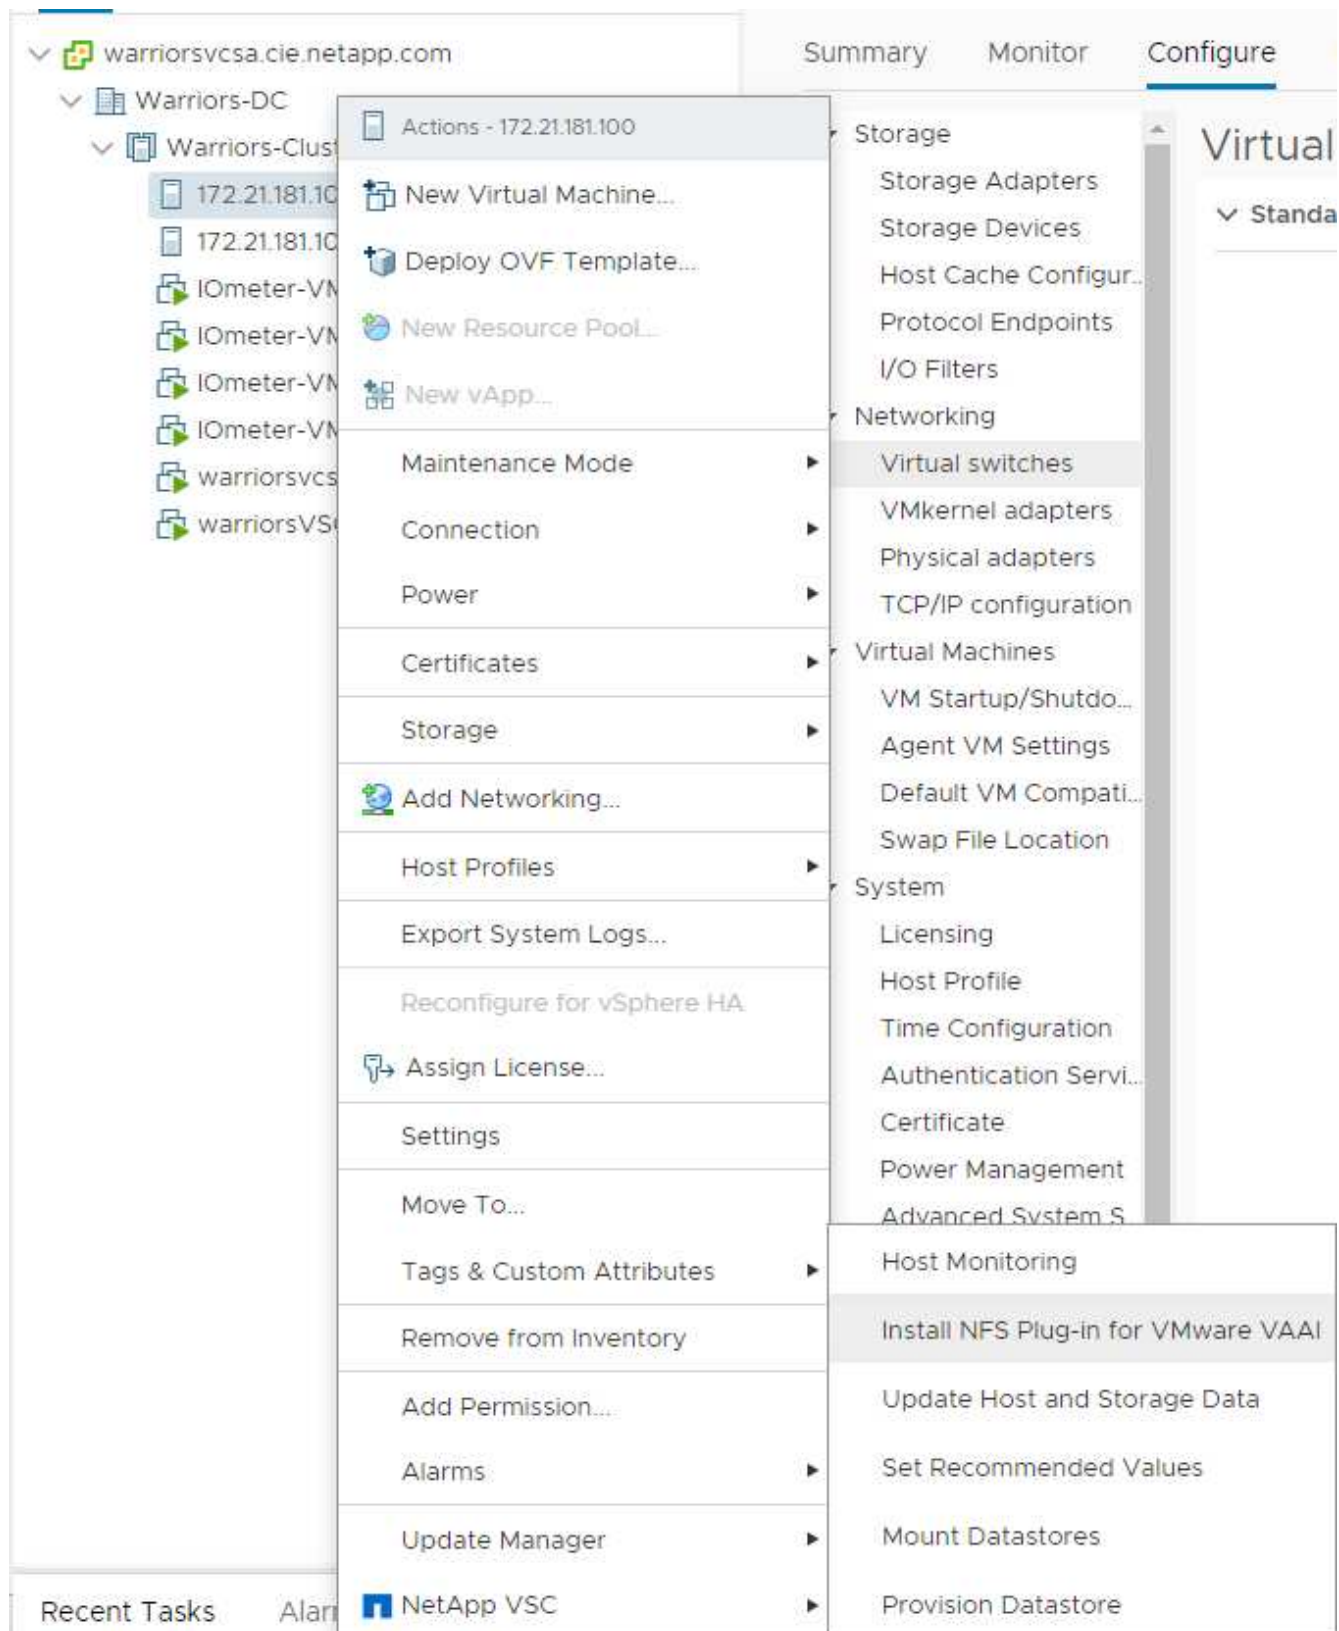
To download and install the NetApp NFS VAAI Plug-In, complete the following steps:

1. Download the NetApp NFS Plug-In 1.1.2 for VMware .vib file from the NFS Plugin Download page and save it to your local machine or admin host.
2. Download the NetApp NFS Plug-in for VMware VAAI:
 - a. Go to the [software download page](#).

- b. Scroll down and click NetApp NFS Plug-in for VMware VAAI.
- c. From the Home screen in the vSphere web client, select Virtual Storage Console.
- d. Under Virtual Storage Console > Settings > NFS VAAI Tools, upload the NFS Plug-in by choosing Select File and browsing to the location where the downloaded plug-in is stored.




3. Click Upload to transfer the plug-in to vCenter.
4. Select the host and then select NetApp VSC > Install NFS Plug-in for VMware VAAI.





Use the optimal storage settings for the ESXi hosts

VSC enables the automated configuration of storage-related settings for all ESXi hosts that are connected to NetApp storage controllers. To use these settings, complete the following steps:

1. From the Home screen, select vCenter > Hosts and Clusters. For each ESXi host, right-click and select NetApp VSC > Set Recommended Values.



2. Check the settings that you would like to apply to the selected vSphere hosts. Click OK to apply the settings.



3. Reboot the ESXi host after these settings are applied.

Conclusion

FlexPod Express provides a simple and effective solution by providing a validated design that uses industry-leading components. By scaling through the addition of components, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed for small to midsize businesses, ROBOs, and other businesses that require dedicated solutions.

Acknowledgments

The authors would like to acknowledge John George for his support and contribution to

this design.

Where to find additional information

To learn more about the information described in this document, refer to the following documents and/or websites:

NetApp Product Documentation

<http://docs.netapp.com>

FlexPod Express with Guide

NVA-1139-DESIGN: FlexPod Express with Cisco UCS C-Series and NetApp AFF C190 Series

<https://www.netapp.com/us/media/nva-1139-design.pdf>

Version history

| Version | Date | Document version history |
|-------------|---------------|--------------------------|
| Version 1.0 | November 2019 | Initial release. |

FlexPod Express with Cisco UCS C-Series and AFF A220 Series Design Guide

NVA-1125-DESIGN: FlexPod Express with Cisco UCS C-Series and AFF A220 Series

*
:hardbreaks:
:icons: font
:linkattrs:
:relative_path: ./express/
:imagesdir: /tmp/d20240325-6552-1p85s1v/source/./express/./../media/

Savita Kumari, NetApp
In partnership with:



Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices, leveraging the technology that they are familiar with in their data center.

FlexPod Express is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp AFF. The components in FlexPod Express are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are

optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

Next: [Program summary](#).

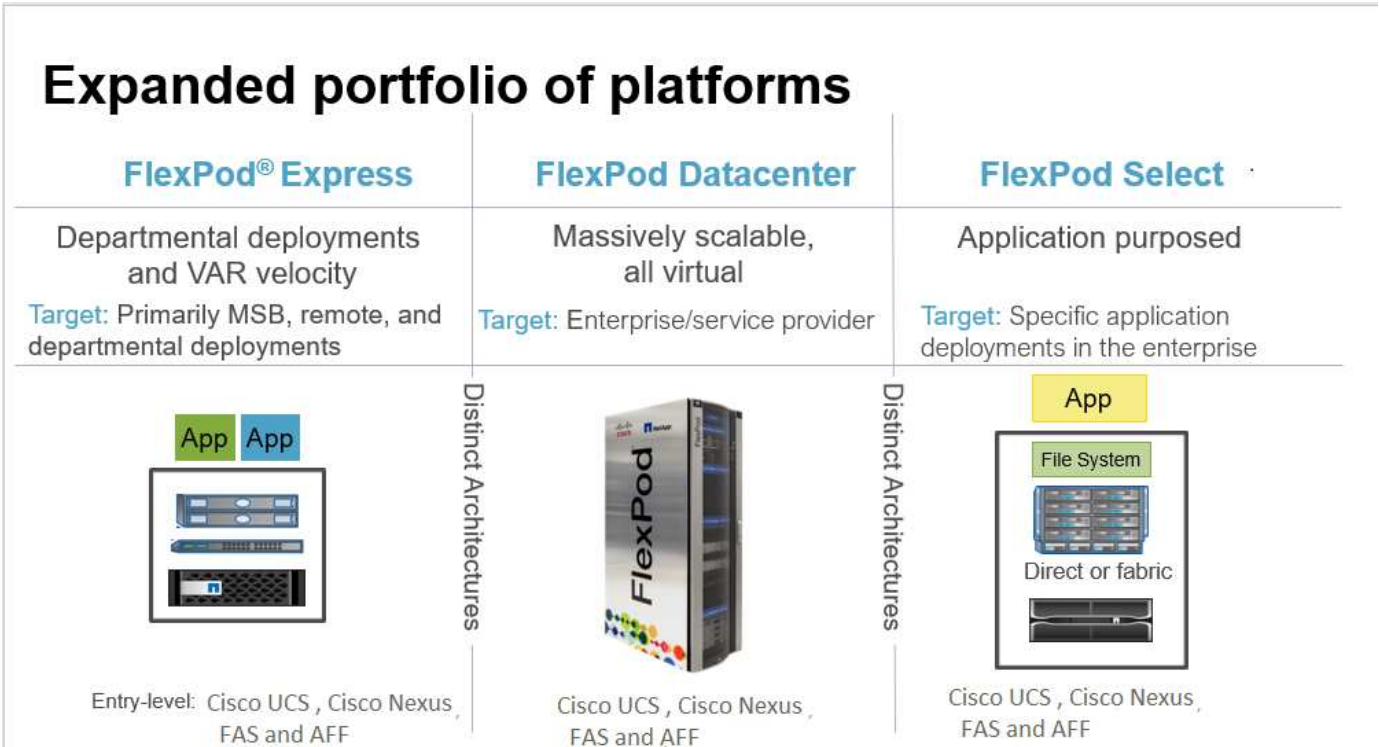
Program summary

FlexPod converged infrastructure portfolio

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or as NetApp Verified Architectures (NVAs). Deviations that are based on customer requirements from a given CVD or NVA are permitted if variations do not result in the deployment of unsupported configurations.

As depicted in the following figure, the FlexPod portfolio includes three solutions: FlexPod Express, FlexPod Datacenter, and FlexPod Select:

- **FlexPod Express.** Offers an entry-level solution that consists of technologies from Cisco and NetApp.
- **FlexPod Datacenter.** Delivers an optimal multipurpose foundation for various workloads and applications.
- **FlexPod Select.** Incorporates the best aspects of FlexPod Datacenter and tailors the infrastructure to a given application.



NetApp Verified Architecture program

The NVA program offers customers a verified architecture for NetApp solutions. An NVA means that the NetApp solution has the following qualities:

- Is thoroughly tested
- Is prescriptive in nature
- Minimizes deployment risks
- Accelerates time to market

This guide details the design of FlexPod Express with VMware vSphere. In addition, this design leverages the all-new AFF A220 system, which runs NetApp ONTAP 9.4 software, Cisco Nexus 3172P switches, and Cisco UCS C220 M5 servers as hypervisor nodes.


Although this document is validated for AFF A220, this solution also supports FAS2700.

[Next: Solution overview.](#)

Solution overview

FlexPod Express is designed to run mixed virtualization workloads. It is targeted for remote and branch offices and for small to midsize businesses. It is also optimal for larger businesses that want to implement a dedicated solution for a purpose. This new solution for FlexPod Express adds new technologies such as NetApp ONTAP 9.4, NetApp AFF A220, and VMware vSphere 6.7.

The following figure shows the hardware components that are included in the FlexPod Express solution.




Target audience

This document is intended for those who want to take advantage of an infrastructure that is built to deliver IT efficiency and enable IT innovation. The audience for this document includes, but is not limited to, sales engineers, field consultants, professional services personnel, IT managers, partner engineers, and customers.

Solution technology

This solution leverages the latest technologies from NetApp, Cisco, and VMware. This solution features the new NetApp AFF A220 system, which runs ONTAP 9.4 software, dual Cisco Nexus 3172P switches, and Cisco UCS C220 M5 Rack Servers that run VMware vSphere 6.7. This validated solution uses 10-Gigabit Ethernet (10GbE) technology. The following figure presents an overview. Guidance is also provided on how to scale by adding two hypervisor nodes at a time so that the FlexPod Express architecture can adapt to an organization's

evolving business needs.



40GbE is not validated, but it is a supported infrastructure.

[Next: Technology requirements.](#)

Technology requirements

FlexPod Express requires a combination of hardware and software components that depends on the selected hypervisor and network speed. In addition, FlexPod Express lays out the hardware components that are required to add hypervisor nodes to the system in units of two.

Hardware requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, either hypervisor can run on the same FlexPod Express hardware.

The following table lists the hardware components that are required for all FlexPod Express configurations and 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 |
|---|----------|
| AFF A220 two-node cluster | 1 |
| Cisco UCS C220 M5 server | 2 |
| Cisco Nexus 3172P switch | 2 |
| Cisco UCS Virtual Interface Card (VIC) 1387 for Cisco UCS C220 M5 Rack Server | 2 |
| Cisco CVR-QSFP-SFP10G adapter | 4 |

Software requirements

The following tables list the software components that are required to implement the architectures of the FlexPod Express solution.

The following table lists software requirements for the base FlexPod Express implementation.

| Software | Version | Details |
|---|---------------------|--------------------------------|
| Cisco Integrated Management Controller (CIMC) | 3.1.3 | For C220 M5 Rack Servers |
| Cisco NX-OS | nxos.7.0.3.17.5.bin | For Cisco Nexus 3172P switches |
| NetApp ONTAP | 9.4 | For AFF A220 controllers |

The following table lists the software that is required for all VMware vSphere implementations on FlexPod Express.

| Software | Version |
|---------------------------------|---------|
| VMware vCenter Server Appliance | 6.7 |
| VMware vSphere ESXi | 6.7 |
| NetApp VAAI Plug-In for ESXi | 1.1.2 |

[Next: Design choices.](#)

Design choices

The following technologies were chosen during the process of architecting this design. Each technology serves a specific purpose in the FlexPod Express infrastructure solution.

NetApp AFF A220 Series with ONTAP 9.4

This solution leverages two of the newest NetApp products: NetApp AFF A220 and ONTAP 9.4 software.

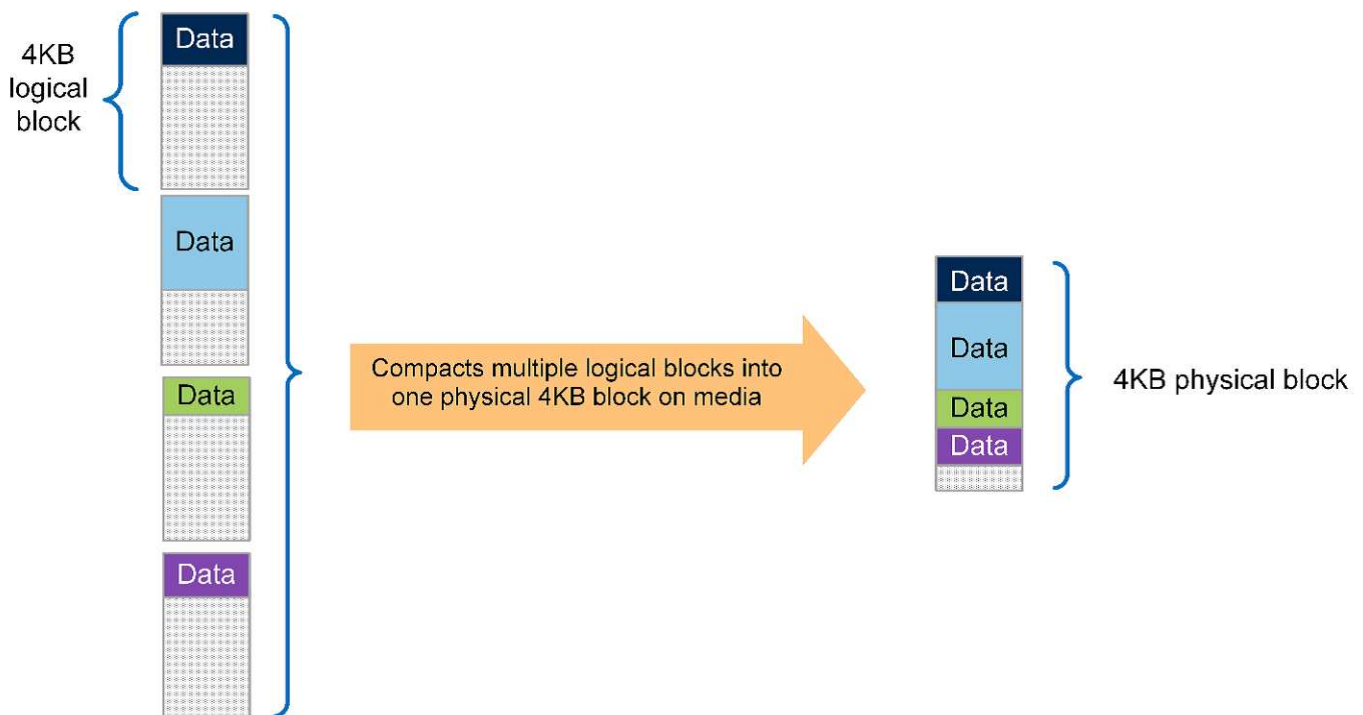
AFF A220 system

For more information about the AFF A220 hardware system, see the [AFF A-Series homepage](#).


ONTAP 9.4 software

NetApp AFF A220 systems use the new ONTAP 9.4 software. ONTAP 9.4 is the industry's leading enterprise data management software. It combines new levels of simplicity and flexibility with powerful data management capabilities, storage efficiencies, and leading cloud integration.

ONTAP 9.4 has several features that are well suited for the FlexPod Express solution. Foremost is NetApp's commitment to storage efficiencies, which can be one of the most important features for small deployments. The hallmark NetApp storage efficiency features such as deduplication, compression, and thin provisioning are available in ONTAP 9.4 with a new addition, compaction. Because the NetApp WAFL system always writes 4KB blocks, compaction combines multiple blocks into a 4KB block when the blocks are not using their allocated space of 4KB. The following figure illustrates this process.



Also, root-data partitioning can be leveraged on the AFF A220 system. This partitioning allows the root aggregate and two data aggregates to be striped across the disks in the system. Therefore, both controllers in a two-node AFF A220 cluster can leverage the performance of all the disks in the aggregate. See the following figure.



These are just a few key features that complement the FlexPod Express solution. For details about the additional features and functionality of ONTAP 9.4, see the [ONTAP 9 Data Management Software datasheet](#). Also, see the NetApp [ONTAP 9 Documentation Center](#), which has been updated to include ONTAP 9.4.

Cisco Nexus 3000 Series

The Cisco Nexus 3172P is a robust, cost-effective switch that offers 1/10/40/100Gbps switching. The Cisco Nexus 3172PQ switch, part of the Unified Fabric family, is a compact, 1-rack-unit (1RU) switch for top-of-rack data center deployments. (See the following figure.) It offers up to seventy-two 1/10GbE ports in 1RU or forty-eight 1/10GbE plus six 40GbE ports in 1RU. And for maximum physical layer flexibility, it also supports 1/10/40Gbps.

Because all the various Cisco Nexus series models run the same underlying operating system, NX-OS, multiple Cisco Nexus models are supported in the FlexPod Express and FlexPod Datacenter solutions.

Performance specifications include:

- Line-rate traffic throughput (both layers 2 and 3) on all ports
- Configurable maximum transmission units (MTUs) of up to 9216 bytes (jumbo frames)



For more information about Cisco Nexus 3172 switches, see the [Cisco Nexus 3172PQ, 3172TQ, 3172TQ-32T, 3172PQ-XL, and 3172TQ-XL switches data sheet](#).

Cisco UCS C-Series

The Cisco UCS C-Series rack server was chosen for FlexPod Express because its many configuration options allow it to be tailored for specific requirements in a FlexPod Express deployment.

Cisco UCS C-Series rack servers deliver unified computing in an industry-standard form factor to reduce TCO and to increase agility.

Cisco UCS C-Series rack servers provide the following benefits:

- A form-factor-agnostic entry point into Cisco UCS
- Simplified and fast deployment of applications
- Extension of unified computing innovations and benefits to rack servers
- Increased customer choice with unique benefits in a familiar rack package



The Cisco UCS C220 M5 rack server (in the previous figure) is among the most versatile general-purpose enterprise infrastructure and application servers in the industry. It is a high-density two-socket rack server that delivers industry-leading performance and efficiency for a wide range of workloads, including virtualization, collaboration, and bare-metal applications. Cisco UCS C-Series Rack Servers can be deployed as standalone servers or as part of Cisco UCS to take advantage of Cisco's standards-based unified computing innovations that help reduce customers' TCO and increase their business agility.

For more information about C220 M5 servers, see the [Cisco UCS C220 M5 Rack Server Data Sheet](#).

Connectivity options for C220 M5 rack servers

The connectivity options for the C220 M5 rack servers are as follows:

- **Cisco UCS VIC 1387**

The Cisco UCS VIC 1387 (in the following figure) offers dual-port enhanced QSFP+ 40GbE and FC over Ethernet (FCoE) in a modular-LAN-on-motherboard (mLOM) form factor. The mLOM slot can be used to install a Cisco VIC without consuming a Peripheral Component Interconnect Express (PCIe) slot, providing greater I/O expandability.



For more information about the Cisco UCS VIC 1387 adapter, see the [Cisco UCS Virtual Interface Card 1387](#) data sheet.

- **CVR-QSFP-SFP10G adapter**

The Cisco QSA Module converts a QSFP port into an SFP or SFP+ port. With this adapter, customers have the flexibility to use any SFP+ or SFP module or cable to connect to a lower-speed port on the other end of the network. This flexibility enables a cost-effective transition to 40GbE by maximizing the use of high-density 40GbE QSFP platforms. This adapter supports all SFP+ optics and cable reaches, and it supports several 1GbE SFP modules. Because this project has been validated by using 10GbE connectivity and because the VIC 1387 used is 40GbE, the CVR-QSFP-SFP10G adapter (in the following figure) is used for conversion.



VMware vSphere 6.7

VMware vSphere 6.7 is one hypervisor option for use with FlexPod Express. VMware vSphere allows organizations to reduce their power and cooling footprint while confirming that the purchased compute capacity is used to its fullest. In addition, VMware vSphere allows hardware failure protection (VMware High Availability, or VMware HA) and compute resource load balancing across a cluster of vSphere hosts (VMware Distributed Resource Scheduler, or VMware DRS).

Because it restarts only the kernel, VMware vSphere 6.7 allows customers to “quick boot” where it loads vSphere ESXi without restarting the hardware. This feature is available only with platforms and drivers that are

on the Quick Boot Whitelist. vSphere 6.7 extends the capabilities of the vSphere Client, which can do about 90% of what the vSphere Web Client can do.

In vSphere 6.7, VMware has extended this capability to enable customers to set Enhanced vMotion Compatibility (EVC) per virtual machine (VM) rather than per host basis. In vSphere 6.7, VMware has also exposed the APIs that can be used to create instant clones.

The following are some of the features of vSphere 6.7 U1:

- Fully featured HTML5 web-based vSphere Client
- vMotion for NVIDIA GRID vGPU VMs. Support for Intel FPGA.
- vCenter Server Converge Tool to move from external PSC to internal PCS.
- Enhancements for vSAN (HCI updates).
- Enhanced content library.

For details about vSphere 6.7 U1, see [What's New in vCenter Server 6.7 Update 1](#). Although this solution was validated with vSphere 6.7, it supports any vSphere version qualified with the other components by the NetApp Interoperability Matrix Tool. NetApp recommends deploying vSphere 6.7U1 for its fixes and enhanced features.

Boot architecture

Following are the supported options for the FlexPod Express boot architecture:

- iSCSI SAN LUN
- Cisco FlexFlash SD Card
- Local disk

Because FlexPod Datacenter is booted from iSCSI LUNs, solution manageability is enhanced by also using iSCSI boot for FlexPod Express.

[Next: Solution verification.](#)

Solution verification

Cisco and NetApp designed and built FlexPod Express to serve as a premier infrastructure platform for their customers. Because it was designed with industry-leading components, customers can trust FlexPod Express as their infrastructure foundation. In keeping with the fundamental principles of the FlexPod portfolio, the FlexPod Express architecture was thoroughly tested by Cisco and NetApp data center architects and engineers. From redundancy and availability to each individual feature, the entire FlexPod Express architecture is validated to instill confidence in our customers and to build trust in the design process.

VMware vSphere 6.7 was verified on the FlexPod Express infrastructure components. This validation included 10GbE uplink connectivity options for the hypervisor.

[Next: Conclusion.](#)

Conclusion

FlexPod Express offers a simple and effective solution by providing a validated design that uses industry-leading components. By scaling and by providing options for the hypervisor platform, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed keeping in mind small to midsize businesses, remote and branch offices, and other businesses that require dedicated solutions.

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

Where to find additional information

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

- NetApp documentation

<https://docs.netapp.com>

- FlexPod Express with VMware vSphere 6.7 and NetApp AFF A220 Deployment Guide

<https://www.netapp.com/us/media/nva-1123-deploy.pdf>

FlexPod Express with Cisco UCS C-Series and AFF A220 Series Deployment Guide

NVA-1123-DEPLOY: FlexPod Express with VMware vSphere 6.7 and NetApp AFF A220 deployment guide

Savita Kumari, NetApp

In partnership with:



Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices, leveraging the technology with which they are familiar in their data center.

FlexPod Express is a predesigned, best practice data center architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp storage technologies. The components in a FlexPod Express system are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal operating systems and enterprise workloads.

FlexPod Datacenter and FlexPod Express deliver a baseline configuration and have the flexibility to be sized

and optimized to accommodate many different use cases and requirements. Existing FlexPod Datacenter customers can manage their FlexPod Express system with the tools to which they are accustomed. New FlexPod Express customers can easily adapt to managing FlexPod Datacenter as their environment grows.

FlexPod Express is an optimal infrastructure foundation for remote and branch offices and for small to midsize businesses. It is also an optimal solution for customers who want to provide infrastructure for a dedicated workload.

FlexPod Express provides an easy-to-manage infrastructure that is suitable for almost any workload.

Solution overview


This FlexPod Express solution is part of the FlexPod Converged Infrastructure Program.

FlexPod Converged Infrastructure Program

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or NetApp Verified Architectures (NVAs). Deviations based on customer requirements from a given CVD or NVA are permitted if these variations do not create an unsupported configuration.

As depicted in the figure below, the FlexPod program includes three solutions: FlexPod Express, FlexPod Datacenter, and FlexPod Select:

- **FlexPod Express.** Offers customers an entry-level solution with technologies from Cisco and NetApp.
- **FlexPod Datacenter.** Delivers an optimal multipurpose foundation for various workloads and applications.
- **FlexPod Select.** Incorporates the best aspects of FlexPod Datacenter and tailors the infrastructure to a given application.



NetApp Verified Architecture Program

The NetApp Verified Architecture program offers customers a verified architecture for NetApp solutions. A NetApp Verified Architecture provides a NetApp solution architecture with the following qualities:

- Is thoroughly tested
- Is prescriptive in nature
- Minimizes deployment risks
- Accelerates time to market


This guide details the design of FlexPod Express with VMware vSphere. In addition, this design uses the all-new AFF A220 system, which runs NetApp ONTAP 9.4; the Cisco Nexus 3172P; and Cisco UCS C-Series C220 M5 servers as hypervisor nodes.

Solution technology

This solution leverages the latest technologies from NetApp, Cisco, and VMware. This solution features the new NetApp AFF A220 running ONTAP 9.4, dual Cisco Nexus 3172P switches, and Cisco UCS C220 M5 rack servers that run VMware vSphere 6.7. This validated solution uses 10GbE technology. Guidance is also provided on how to scale compute capacity by adding two hypervisor nodes at a time so that the FlexPod Express architecture can adapt to an organization's evolving business needs.

The following figure shows FlexPod Express with VMware vSphere 10GbE architecture.

FlexPod Express



This validation uses 10GbE connectivity and a Cisco UCS VIC 1387, which is 40GbE. To achieve 10GbE connectivity, the CVR-QSFP-SFP10G adapter is used.

Use case summary

The FlexPod Express solution can be applied to several use cases, including the following:

- Remote offices or branch offices
- Small and midsize businesses
- Environments that require a dedicated and cost-effective solution

FlexPod Express is best suited for virtualized and mixed workloads.



Although this solution was validated with vSphere 6.7, it supports any vSphere version qualified with the other components by the NetApp Interoperability Matrix Tool. NetApp recommends deploying vSphere 6.7U1 for its fixes and enhanced features.

Following are some features of vSphere 6.7 U1:

- Fully featured HTML5 web-based vSphere client
- vMotion for NVIDIA GRID vGPU VMs. Support for Intel FPGA
- vCenter Server Converge Tool to move from external PSC to internal PCS
- Enhancements for vSAN (HCI updates)
- Enhanced content library

For details about vSphere 6.7 U1, see [What's New in vCenter Server 6.7 Update 1](#).

Technology requirements

A FlexPod Express system requires a combination of hardware and software components. FlexPod Express also describes the hardware components that are required to add hypervisor nodes to the system in units of two.

Hardware requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, either hypervisor can run on the same FlexPod Express hardware.

The following table lists the hardware components required for all FlexPod Express configurations.

| Hardware | Quantity |
|--|----------|
| AFF A220 HA Pair | 1 |
| Cisco C220 M5 server | 2 |
| Cisco Nexus 3172P switch | 2 |
| Cisco UCS virtual interface card (VIC) 1387 for the C220 M5 server | 2 |
| CVR-QSFP-SFP10G adapter | 4 |

The following table lists the hardware required in addition to the base configuration for implementing 10GbE.

| Hardware | Quantity |
|--------------------------|----------|
| Cisco UCS C220 M5 server | 2 |
| Cisco VIC 1387 | 2 |
| CVR-QSFP-SFP10G adapter | 4 |

Software requirements

The following table lists the software components required to implement the architectures of the FlexPod Express solutions.

| Software | Version | Details |
|---|---------|------------------------------------|
| Cisco Integrated Management Controller (CIMC) | 3.1(3g) | For Cisco UCS C220 M5 rack servers |


| Software | Version | Details |
|--------------------|---------------------|--------------------------------|
| Cisco nenic driver | 1.0.25.0 | For VIC 1387 interface cards |
| Cisco NX-OS | nxos.7.0.3.17.5.bin | For Cisco Nexus 3172P switches |
| NetApp ONTAP | 9.4 | For AFF A220 controllers |

The following table lists the software required for all VMware vSphere implementations on FlexPod Express.

| Software | Version |
|---------------------------------|---------|
| VMware vCenter server appliance | 6.7 |
| VMware vSphere ESXi hypervisor | 6.7 |
| NetApp VAAI Plug-In for ESXi | 1.1.2 |

FlexPod Express cabling information

The following figure shows the reference validation cabling.



The following table shows cabling information for the Cisco Nexus switch 3172P A.

| Local device | Local port | Remote device | Remote port |
|----------------------------|------------|--------------------------------------|-------------|
| Cisco Nexus switch 3172P A | Eth1/1 | NetApp AFF A220 storage controller A | e0c |

| Local device | Local port | Remote device | Remote port |
|--------------|------------|---|------------------------------------|
| | Eth1/2 | NetApp AFF A220 storage controller B | e0c |
| | Eth1/3 | Cisco UCS C220 C-Series standalone server A | MLOM1 with CVR-QSFP-SFP10G adapter |
| | Eth1/4 | Cisco UCS C220 C-Series standalone server B | MLOM1 with CVR-QSFP-SFP10G adapter |
| | Eth1/25 | Cisco Nexus switch 3172P B | Eth1/25 |
| | Eth1/26 | Cisco Nexus switch 3172P B | Eth1/26 |
| | Eth1/33 | NetApp AFF A220 storage controller A | e0M |
| | Eth1/34 | Cisco UCS C220 C-Series standalone server A | CIMC |

The following table shows cabling information for Cisco Nexus switch 3172P B.

| Local device | Local port | Remote device | Remote port |
|----------------------------|------------|---|------------------------------------|
| Cisco Nexus switch 3172P B | Eth1/1 | NetApp AFF A220 storage controller A | e0d |
| | Eth1/2 | NetApp AFF A220 storage controller B | e0d |
| | Eth1/3 | Cisco UCS C220 C-Series standalone server A | MLOM2 with CVR-QSFP-SFP10G adapter |
| | Eth1/4 | Cisco UCS C220 C-Series standalone server B | MLOM2 with CVR-QSFP-SFP10G adapter |
| | Eth1/25 | Cisco Nexus switch 3172P A | Eth1/25 |
| | Eth1/26 | Cisco Nexus switch 3172P A | Eth1/26 |
| | Eth1/33 | NetApp AFF A220 storage controller B | e0M |
| | Eth1/34 | Cisco UCS C220 C-Series standalone server B | CIMC |

The following table shows the cabling information for NetApp AFF A220 storage controller A.

| Local device | Local port | Remote device | Remote port |
|--------------------------------------|------------|--------------------------------------|-------------|
| NetApp AFF A220 storage controller A | e0a | NetApp AFF A220 storage controller B | e0a |
| | e0b | NetApp AFF A220 storage controller B | e0b |

| Local device | Local port | Remote device | Remote port |
|--------------|------------|----------------------------|-------------|
| | e0c | Cisco Nexus switch 3172P A | Eth1/1 |
| | e0d | Cisco Nexus switch 3172P B | Eth1/1 |
| | e0M | Cisco Nexus switch 3172P A | Eth1/33 |

The following table shows cabling information for NetApp AFF A220 storage controller B.

| Local device | Local port | Remote device | Remote port |
|--------------------------------------|------------|--------------------------------------|-------------|
| NetApp AFF A220 storage controller B | e0a | NetApp AFF A220 storage controller A | e0a |
| | e0b | NetApp AFF A220 storage controller A | e0b |
| | e0c | Cisco Nexus switch 3172P A | Eth1/2 |
| | e0d | Cisco Nexus switch 3172P B | Eth1/2 |
| | e0M | Cisco Nexus switch 3172P B | Eth1/33 |

Deployment procedures

This document provides details for configuring a fully redundant, highly available FlexPod Express system. To reflect this redundancy, the components being configured in each step are referred to as either component A or component B. For example, controller A and controller B identify the two NetApp storage controllers that are provisioned in this document. Switch A and switch B identify a pair of Cisco Nexus switches.

In addition, this document describes steps for provisioning multiple Cisco UCS hosts, which are identified sequentially as server A, server B, and so on.

To indicate that you should include information pertinent to your environment in a step, `<<text>>` appears as part of the command structure. See the following example for the `vlan create` command:

```
Controller01>vlan create vif0 <<mgmt_vlan_id>>
```

This document enables you to fully configure the FlexPod Express environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and virtual local area network (VLAN) schemes. The table below describes the VLANs required for deployment, as outlined in this guide. This table can be completed based on the specific site variables and used to implement the document configuration steps.



If you use separate in-band and out-of-band management VLANs, you must create a layer-3 route between them. For this validation, a common management VLAN was used.

| AN Name | VLAN Purpose | ID Used in Validating This Document |
|------------------------------|--|-------------------------------------|
| Management VLAN | VLAN for management interfaces | 3437 |
| Native VLAN | VLAN to which untagged frames are assigned | 2 |
| NFS VLAN | VLAN for NFS traffic | 3438 |
| VMware vMotion VLAN | VLAN designated for the movement of virtual machines from one physical host to another | 3441 |
| Virtual machine traffic VLAN | VLAN for virtual machine application traffic | 3442 |
| iSCSI-A-VLAN | VLAN for iSCSI traffic on fabric A | 3439 |
| iSCSI-B-VLAN | VLAN for iSCSI traffic on fabric B | 3440 |

The VLAN numbers are needed throughout the configuration of FlexPod Express. The VLANs are referred to as `<<var_XXXX_vlan>>`, where `XXXX` is the purpose of the VLAN (such as iSCSI-A).

The table below lists the VMware virtual machines created.

| Virtual machine description | Host name |
|-----------------------------|-----------|
| VMware vCenter Server | |

Cisco Nexus 3172P deployment procedure

The following section details the Cisco Nexus 3172P switch configuration used in a FlexPod Express environment.

Initial setup of Cisco Nexus 3172P switch

The following procedures describe how to configure the Cisco Nexus switches for use in a base FlexPod Express environment.



This procedure assumes that you are using a Cisco Nexus 3172P running NX-OS software release 7.0(3)I7(5).

1. Upon initial boot and connection to the console port of the switch, the Cisco NX-OS setup automatically starts. This initial configuration addresses basic settings, such as the switch name, the mgmt0 interface configuration, and Secure Shell (SSH) setup.
2. The FlexPod Express management network can be configured in multiple ways. The mgmt0 interfaces on the 3172P switches can be connected to an existing management network, or the mgmt0 interfaces of the 3172P switches can be connected in a back-to-back configuration. However, this link cannot be used for external management access such as SSH traffic.

In this deployment guide, the FlexPod Express Cisco Nexus 3172P switches are connected to an existing

management network.

3. To configure the Cisco Nexus 3172P switches, power on the switch and follow the on- screen prompts, as illustrated here for the initial setup of both the switches, substituting the appropriate values for the switch-specific information.

```
This setup utility will guide you through the basic configuration of
the system. Setup configures only enough connectivity for management
of the system.
*Note: setup is mainly used for configuring the system initially,
when no configuration is present. So setup always assumes system
defaults and not the current system configuration values.
Press Enter at anytime to skip a dialog. Use ctrl-c at anytime
to skip the remaining dialogs.
Would you like to enter the basic configuration dialog (yes/no): y
Do you want to enforce secure password standard (yes/no) [y]: y
  Create another login account (yes/no) [n]: n
  Configure read-only SNMP community string (yes/no) [n]: n
  Configure read-write SNMP community string (yes/no) [n]: n
  Enter the switch name : 3172P-B
  Continue with Out-of-band (mgmt0) management configuration? (yes/no)
[y]: y
    Mgmt0 IPv4 address : <<var_switch_mgmt_ip>>
    Mgmt0 IPv4 netmask : <<var_switch_mgmt_netmask>>
  Configure the default gateway? (yes/no) [y]: y
    IPv4 address of the default gateway : <<var_switch_mgmt_gateway>>
  Configure advanced IP options? (yes/no) [n]: n
  Enable the telnet service? (yes/no) [n]: n
  Enable the ssh service? (yes/no) [y]: y
    Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
    Number of rsa key bits <1024-2048> [1024]: <enter>
  Configure the ntp server? (yes/no) [n]: y
    NTP server IPv4 address : <<var_ntp_ip>>
  Configure default interface layer (L3/L2) [L2]: <enter>
  Configure default switchport interface state (shut/noshut) [noshut]:
<enter>
  Configure CoPP system profile (strict/moderate/lenient/dense)
[strict]: <enter>
```

4. You then see a summary of your configuration, and you are asked if you would like to edit it. If your configuration is correct, enter n.

```
Would you like to edit the configuration? (yes/no) [n]: n
```

5. You are then asked if you would like to use this configuration and save it. If so, enter y.

```
Use this configuration and save it? (yes/no) [y]: Enter
```

6. Repeat this procedure for Cisco Nexus switch B.

Enable advanced features

Certain advanced features must be enabled in Cisco NX-OS to provide additional configuration options.



The `interface-vlan` feature is required only if you use the back-to-back `mgmt0` option described throughout this document. This feature allows you to assign an IP address to the interface VLAN (switch virtual interface), which enables in-band management communication to the switch (such as through SSH).

1. To enable the appropriate features on Cisco Nexus switch A and switch B, enter configuration mode using the command (`config t`) and run the following commands:

```
feature interface-vlan
feature lacp
feature vpc
```

The default port channel load-balancing hash uses the source and destination IP addresses to determine the load-balancing algorithm across the interfaces in the port channel. You can achieve better distribution across the members of the port channel by providing more inputs to the hash algorithm beyond the source and destination IP addresses. For the same reason, NetApp highly recommends adding the source and destination TCP ports to the hash algorithm.

2. From configuration mode (`config t`), enter the following commands to set the global port channel load-balancing configuration on Cisco Nexus switch A and switch B:

```
port-channel load-balance src-dst ip-l4port
```

Perform global spanning-tree configuration

The Cisco Nexus platform uses a new protection feature called bridge assurance. Bridge assurance helps protect against a unidirectional link or other software failure with a device that continues to forward data traffic when it is no longer running the spanning-tree algorithm. Ports can be placed in one of several states, including network or edge, depending on the platform.

NetApp recommends setting bridge assurance so that all ports are considered to be network ports by default. This setting forces the network administrator to review the configuration of each port. It also reveals the most common configuration errors, such as unidentified edge ports or a neighbor that does not have the bridge assurance feature enabled. In addition, it is safer to have the spanning tree block many ports rather than too few, which allows the default port state to enhance the overall stability of the network.

Pay close attention to the spanning-tree state when adding servers, storage, and uplink switches, especially if they do not support bridge assurance. In such cases, you might need to change the port type to make the ports active.

The Bridge Protocol Data Unit (BPDU) guard is enabled on edge ports by default as another layer of protection. To prevent loops in the network, this feature shuts down the port if BPDUs from another switch are seen on this interface.

From configuration mode (`config t`), run the following commands to configure the default spanning- tree options, including the default port type and BPDU guard, on Cisco Nexus switch A and switch B:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
```

Define VLANs

Before individual ports with different VLANs are configured, the layer 2 VLANs must be defined on the switch. It is also a good practice to name the VLANs for easy troubleshooting in the future.

From configuration mode (`config t`), run the following commands to define and describe the layer 2 VLANs on Cisco Nexus switch A and switch B:

```
vlan <<nfs_vlan_id>>
  name NFS-VLAN
vlan <<iSCSI_A_vlan_id>>
  name iSCSI-A-VLAN
vlan <<iSCSI_B_vlan_id>>
  name iSCSI-B-VLAN
vlan <<vmotion_vlan_id>>
  name vMotion-VLAN
vlan <<vmtraffic_vlan_id>>
  name VM-Traffic-VLAN
vlan <<mgmt_vlan_id>>
  name MGMT-VLAN
vlan <<native_vlan_id>>
  name NATIVE-VLAN
exit
```

Configure access and management port descriptions

As is the case with assigning names to the layer 2 VLANs, setting descriptions for all the interfaces can help with both provisioning and troubleshooting.

From configuration mode (`config t`) in each of the switches, enter the following port descriptions for the FlexPod Express large configuration:

Cisco Nexus Switch A

```

int eth1/1
    description AFF A220-A e0c
int eth1/2
    description AFF A220-B e0c
int eth1/3
    description UCS-Server-A: MLOM port 0
int eth1/4
    description UCS-Server-B: MLOM port 0
int eth1/25
    description vPC peer-link 3172P-B 1/25
int eth1/26
    description vPC peer-link 3172P-B 1/26
int eth1/33
    description AFF A220-A e0M
int eth1/34
    description UCS Server A: CIMC

```

Cisco Nexus Switch B

```

int eth1/1
    description AFF A220-A e0d
int eth1/2
    description AFF A220-B e0d
int eth1/3
    description UCS-Server-A: MLOM port 1
int eth1/4
    description UCS-Server-B: MLOM port 1
int eth1/25
    description vPC peer-link 3172P-A 1/25
int eth1/26
    description vPC peer-link 3172P-A 1/26
int eth1/33
    description AFF A220-B e0M
int eth1/34
    description UCS Server B: CIMC

```

Configure server and storage management interfaces

The management interfaces for both the server and the storage typically use only a single VLAN. Therefore, configure the management interface ports as access ports. Define the management VLAN for each switch and change the spanning-tree port type to edge.

From configuration mode (`config t`), enter the following commands to configure the port settings for the management interfaces of both the servers and the storage:

Cisco Nexus Switch A

```
int eth1/33-34
  switchport mode access
  switchport access vlan <<mgmt_vlan>>
  spanning-tree port type edge
  speed 1000
exit
```

Cisco Nexus Switch B

```
int eth1/33-34
  switchport mode access
  switchport access vlan <<mgmt_vlan>>
  spanning-tree port type edge
  speed 1000
exit
```

Perform virtual port channel global configuration

A virtual port channel (vPC) enables links that are physically connected to two different Cisco Nexus switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device. A vPC can provide layer-2 multipathing, which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

A vPC provides the following benefits:

- Enabling a single device to use a port channel across two upstream devices
- Eliminating spanning-tree protocol blocked ports
- Providing a loop-free topology
- Using all available uplink bandwidth
- Providing fast convergence if either the link or a device fails
- Providing link-level resiliency
- Helping provide high availability

The vPC feature requires some initial setup between the two Cisco Nexus switches to function properly. If you use the back-to-back mgmt0 configuration, use the addresses defined on the interfaces and verify that they can communicate by using the ping `[switch_A/B_mgmt0_ip_addr] vrf management` command.

From configuration mode (`config t`), run the following commands to configure the vPC global configuration for both switches:

Cisco Nexus Switch A

```

vpc domain 1
  role priority 10
  peer-keepalive destination <<switch_B_mgmt0_ip_addr>> source
<<switch_A_mgmt0_ip_addr>> vrf management
  peer-gateway
  auto-recovery
  ip arp synchronize
int eth1/25-26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>,
<<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
  spanning-tree port type network
  vpc peer-link
  no shut
exit
copy run start

```

Cisco Nexus Switch B

```

vpc domain 1
  peer-switch
  role priority 20
  peer-keepalive destination <<switch_A_mgmt0_ip_addr>> source
<<switch_B_mgmt0_ip_addr>> vrf management
  peer-gateway
  auto-recovery
  ip arp synchronize
int eth1/25- 26
  channel-group 10 mode active
int Po10
  description vPC peer-link
  switchport
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>,
<<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
  spanning-tree port type network
  vpc peer-link
no shut
exit
copy run start

```

Configure storage port channels

The NetApp storage controllers allow an active-active connection to the network using the Link Aggregation Control Protocol (LACP). The use of LACP is preferred because it adds both negotiation and logging between the switches. Because the network is set up for vPC, this approach enables you to have active-active connections from the storage to separate physical switches. Each controller has two links to each of the switches. However, all four links are part of the same vPC and interface group (IFGRP).

From configuration mode (`config t`), run the following commands on each of the switches to configure the individual interfaces and the resulting port channel configuration for the ports connected to the NetApp AFF controller.

1. Run the following commands on switch A and switch B to configure the port channels for storage controller A:


```

int eth1/1
    channel-group 11 mode active
int Po11
    description vPC to Controller-A
    switchport
    switchport mode trunk
    switchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan
<<nfs_vlan_id>>,<<mgmt_vlan_id>>,<<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>>
    spanning-tree port type edge trunk
    mtu 9216
    vpc 11
    no shut

```

2. Run the following commands on switch A and switch B to configure the port channels for storage controller B.

```

int eth1/2
    channel-group 12 mode active
int Po12
    description vPC to Controller-B
    switchport
    switchport mode trunk
    switchport trunk native vlan <<native_vlan_id>>
    switchport trunk allowed vlan <<nfs_vlan_id>>,<<mgmt_vlan_id>>,
<<iSCSI_A_vlan_id>>, <<iSCSI_B_vlan_id>>
    spanning-tree port type edge trunk
    mtu 9216
    vpc 12
    no shut
exit
copy run start

```



In this solution validation, an MTU of 9000 was used. However, based on application requirements, you can configure an appropriate value of MTU. It is important to set the same MTU value across the FlexPod solution. Incorrect MTU configurations between components will result in packets being dropped and these packets.

Configure server connections

The Cisco UCS servers have a two-port virtual interface card, VIC1387, that is used for data traffic and booting of the ESXi operating system using iSCSI. These interfaces are configured to fail over to one another, providing additional redundancy beyond a single link. Spreading these links across multiple switches enables the server to survive even a complete switch failure.

From configuration mode (config t), run the following commands to configure the port settings for the interfaces connected to each server.

Cisco Nexus Switch A: Cisco UCS Server-A and Cisco UCS Server-B configuration

```
int eth1/3-4
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<iSCSI_A_vlan_id>>,<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_i
d>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu9216
  no shut
exit
copy run start
```

Cisco Nexus Switch B: Cisco UCS Server-A and Cisco UCS Server-B configuration

```
int eth1/3-4
  switchport mode trunk
  switchport trunk native vlan <<native_vlan_id>>
  switchport trunk allowed vlan
<<iSCSI_B_vlan_id>>,<<nfs_vlan_id>>,<<vmotion_vlan_id>>,<<vmtraffic_vlan_i
d>>,<<mgmt_vlan_id>>
  spanning-tree port type edge trunk
  mtu 9216
  no shut
exit
copy run start
```

In this solution validation, an MTU of 9000 was used. However, based on application requirements, you can configure an appropriate value of MTU. It is important to set the same MTU value across the FlexPod solution. Incorrect MTU configurations between components will result in packets being dropped and these packets will need to be transmitted again. This will affect the overall performance of the solution.

To scale the solution by adding additional Cisco UCS servers, run the previous commands with the switch ports that the newly added servers have been plugged into on switches A and B.

Uplink into existing network infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 3172P switches included in the FlexPod environment into the infrastructure. The uplinks may be 10GbE uplinks for a 10GbE infrastructure solution or 1GbE for a 1GbE infrastructure solution if required. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run copy run start to save the configuration on each switch after the configuration is

completed.

[Next: NetApp Storage Deployment Procedure \(Part 1\)](#)

NetApp storage deployment procedure (part 1)

This section describes the NetApp AFF storage deployment procedure.

NetApp storage controller AFF2xx series installation

NetApp Hardware Universe

The NetApp Hardware Universe (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities.

Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install:

- 1. Access the [HWU](#) application to view the system configuration guides. Click the Controllers tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.
- 2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controller AFF2XX Series prerequisites

To plan the physical location of the storage systems, see the NetApp Hardware Universe. Refer to the following sections: Electrical Requirements, Supported Power Cords, and Onboard Ports and Cables.

Storage controllers

Follow the physical installation procedures for the controllers in the [AFF A220 Documentation](#).

NetApp ONTAP 9.4

Configuration worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the [ONTAP 9.4 Software Setup Guide](#).



This system is set up in a two-node switchless cluster configuration.

The following table shows ONTAP 9.4 installation and configuration information.

| Cluster detail | Cluster detail value |
|---------------------------|----------------------------|
| Cluster node A IP address | <<var_nodeA_mgmt_ip>> |
| Cluster node A netmask | <<var_nodeA_mgmt_mask>> |
| Cluster node A gateway | <<var_nodeA_mgmt_gateway>> |
| Cluster node A name | <<var_nodeA>> |
| Cluster node B IP address | <<var_nodeB_mgmt_ip>> |

| Cluster detail | Cluster detail value |
|---|-----------------------------|
| Cluster node B netmask | <<var_nodeB_mgmt_mask>> |
| Cluster node B gateway | <<var_nodeB_mgmt_gateway>> |
| Cluster node B name | <<var_nodeB>> |
| ONTAP 9.4 URL | <<var_url_boot_software>> |
| Name for cluster | <<var_clustername>> |
| Cluster management IP address | <<var_clustermgmt_ip>> |
| Cluster B gateway | <<var_clustermgmt_gateway>> |
| Cluster B netmask | <<var_clustermgmt_mask>> |
| Domain name | <<var_domain_name>> |
| DNS server IP (you can enter more than one) | <<var_dns_server_ip>> |
| NTP server IP (you can enter more than one) | <<var_ntp_server_ip>> |

Configure Node A

To configure node A, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Allow the system to boot.

```
autoboot
```

3. Press Ctrl-C to enter the Boot menu.

If ONTAP 9.4 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.4 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeA_mgmt_ip>> <<var_nodeA_mgmt_mask>> <<var_nodeA_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter `y` to zero disks, reset config, and install a new file system.
16. Enter `y` to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the node B configuration while the disks for node A are zeroing.

17. While node A is initializing, begin configuring node B.

Configure Node B

To configure node B, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Press Ctrl-C to enter the Boot menu.

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9.4 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.4 is the version being booted, select option 8 and `y` to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.

5. Enter `y` to perform an upgrade.
6. Select `e0M` for the network port you want to use for the download.
7. Enter `y` to reboot now.
8. Enter the IP address, netmask, and default gateway for `e0M` in their respective places.

```
<<var_nodeB_mgmt_ip>> <<var_nodeB_mgmt_ip>><<var_nodeB_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press Enter for the user name, indicating no user name.
11. Enter `y` to set the newly installed software as the default to be used for subsequent reboots.
12. Enter `y` to reboot the node.

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter `y` to zero disks, reset config, and install a new file system.
16. Enter `y` to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

Continuation of Node A configuration and cluster configuration

From a console port program attached to the storage controller A (node A) console port, run the node setup script. This script appears when ONTAP 9.4 boots on the node for the first time.



The node and cluster setup procedure has changed slightly in ONTAP 9.4. The cluster setup wizard is now used to configure the first node in a cluster, and System Manager is used to configure the cluster.

1. Follow the prompts to set up Node A.

```

Welcome to the cluster setup wizard.
You can enter the following commands at any time:
    "help" or "?" - if you want to have a question clarified,
    "back" - if you want to change previously answered questions, and
    "exit" or "quit" - if you want to quit the cluster setup wizard.
    Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
This system will send event messages and periodic reports to NetApp
Technical
Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.
Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/
Type yes to confirm and continue {yes}: yes
Enter the node management interface port [e0M]:
Enter the node management interface IP address: <<var_nodeA_mgmt_ip>>
Enter the node management interface netmask: <<var_nodeA_mgmt_mask>>
Enter the node management interface default gateway:
<<var_nodeA_mgmt_gateway>>
A node management interface on port e0M with IP address
<<var_nodeA_mgmt_ip>> has been created.
Use your web browser to complete cluster setup by accessing
https://<<var_nodeA_mgmt_ip>>
Otherwise, press Enter to complete cluster setup using the command line
interface:

```

2. Navigate to the IP address of the node's management interface.

Cluster setup can also be performed by using the CLI. This document describes cluster setup using NetApp System Manager guided setup.

3. Click Guided Setup to configure the cluster.

4. Enter <<var_clustername>> for the cluster name and <<var_nodeA>> and <<var_nodeB>> for each of the nodes that you are configuring. Enter the password that you would like to use for the storage system. Select Switchless Cluster for the cluster type. Enter the cluster base license.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:

Cluster Network Support Summary

Cluster Name

Nodes

Not sure all nodes have been discovered? [Refresh](#)

FAS2650 621630000092

HA-PAR

FAS2650 621630000093

Cluster Configuration: ☐ Switched Cluster ☐ Switchless Cluster

Username: admin

Password

Confirm Password

Cluster Base License (Optional)

For any queries related to licenses, contact mysupport.netapp.com

Feature Licenses (Optional)

Cluster Base License is mandatory to add Feature Licenses.

5. You can also enter feature licenses for Cluster, NFS, and iSCSI.
6. You see a status message stating the cluster is being created. This status message cycles through several statuses. This process takes several minutes.
7. Configure the network.
 - a. Deselect the IP Address Range option.
 - b. Enter `<<var_clustermgmt_ip>>` in the Cluster Management IP Address field, `<<var_clustermgmt_mask>>` in the Netmask field, and `<<var_clustermgmt_gateway>>` in the Gateway field. Use the ... selector in the Port field to select e0M of node A.
 - c. The node management IP for node A is already populated. Enter `<<var_nodeA_mgmt_ip>>` for node B.

- d. Enter <<var_domain_name>> in the DNS Domain Name field. Enter <<var_dns_server_ip>> in the DNS Server IP Address field.

You can enter multiple DNS server IP addresses.

- e. Enter <<var_ntp_server_ip>> in the Primary NTP Server field.

You can also enter an alternate NTP server.

8. Configure the support information.

- a. If your environment requires a proxy to access AutoSupport, enter the URL in Proxy URL.
- b. Enter the SMTP mail host and email address for event notifications.

You must, at a minimum, set up the event notification method before you can proceed. You can select any of the methods.

Guided Setup to Configure a Cluster

Provide the information required below to configure your cluster:



? AutoSupport ☒

? Proxy URL (Optional)

i Connection is verified after configuring AutoSupport on all nodes.

? Event Notifications

Notify me through:



Email

SMTP Mail Host

Email Addresses

Separate email addresses with a comma...



SNMP

SNMP Trap Host



Syslog

Syslog Server

Submit

- When indicated that the cluster configuration has completed, click Manage Your Cluster to configure the storage.

Continuation of storage cluster configuration

After the configuration of the storage nodes and base cluster, you can continue with the configuration of the

storage cluster.

Zero all spare disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```

Set on-board UTA2 ports personality

1. Verify the current mode and the current type of the ports by running the `ucadmin show` command.

```
AFF A220::> ucadmin show
```

| Node | Adapter | Current Mode | Current Type | Pending Mode | Pending Type | Admin Status |
|------------|---------|--------------|--------------|--------------|--------------|--------------|
| AFF A220_A | 0c | fc | target | - | - | online |
| AFF A220_A | 0d | fc | target | - | - | online |
| AFF A220_A | 0e | fc | target | - | - | online |
| AFF A220_A | 0f | fc | target | - | - | online |
| AFF A220_B | 0c | fc | target | - | - | online |
| AFF A220_B | 0d | fc | target | - | - | online |
| AFF A220_B | 0e | fc | target | - | - | online |
| AFF A220_B | 0f | fc | target | - | - | online |

8 entries were displayed.

2. Verify that the current mode of the ports that are in use is `cna` and that the current type is set to `target`. If not, change the port personality by using the following command:

```
ucadmin modify -node <home node of the port> -adapter <port name> -mode  
cna -type target
```

The ports must be offline to run the previous command. To take a port offline, run the following command:

```
`network fcp adapter modify -node <home node of the port> -adapter <port  
name> -state down`
```



If you changed the port personality, you must reboot each node for the change to take effect.

Rename management logical interfaces (LIFs)

To rename the management LIFs, complete the following steps:

1. Show the current management LIF names.

```
network interface show -vserver <<clustername>>
```

2. Rename the cluster management LIF.

```
network interface rename -vserver <<clustername>> -lif  
cluster_setup_cluster_mgmt_lif_1 -newname cluster_mgmt
```

3. Rename the node B management LIF.

```
network interface rename -vserver <<clustername>> -lif  
cluster_setup_node_mgmt_lif_AFF A220_B_1 -newname AFF A220-02_mgmt1
```

Set auto-revert on cluster management

Set the `auto-revert` parameter on the cluster management interface.

```
network interface modify -vserver <<clustername>> -lif cluster_mgmt -auto-  
revert true
```

Set up service processor network interface

To assign a static IPv4 address to the service processor on each node, run the following commands:

```
system service-processor network modify -node <<var_nodeA>> -address  
-family IPv4 -enable true -dhcp none -ip-address <<var_nodeA_sp_ip>>  
-netmask <<var_nodeA_sp_mask>> -gateway <<var_nodeA_sp_gateway>>  
system service-processor network modify -node <<var_nodeB>> -address  
-family IPv4 -enable true -dhcp none -ip-address <<var_nodeB_sp_ip>>  
-netmask <<var_nodeB_sp_mask>> -gateway <<var_nodeB_sp_gateway>>
```



The service processor IP addresses should be in the same subnet as the node management IP addresses.

Enable storage failover in ONTAP

To confirm that storage failover is enabled, run the following commands in a failover pair:

1. Verify the status of storage failover.

```
storage failover show
```

Both <<var_nodeA>> and <<var_nodeB>> must be able to perform a takeover. Go to step 3 if the nodes can perform a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_nodeA>> -enabled true
```

Enabling failover on one node enables it for both nodes.

3. Verify the HA status of the two-node cluster.

This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

4. Go to step 6 if high availability is configured. If high availability is configured, you see the following message upon issuing the command:

```
High Availability Configured: true
```

5. Enable HA mode only for the two-node cluster.



Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true  
Do you want to continue? {y|n}: y
```

6. Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

```
storage failover hwassist show
```

The message `Keep Alive Status : Error: did not receive hwassist keep alive alerts from partner` indicates that hardware assist is not configured. Run the following commands to configure hardware assist.

```
storage failover modify -hwassist-partner-ip <<var_nodeB_mgmt_ip>> -node
<<var_nodeA>>
storage failover modify -hwassist-partner-ip <<var_nodeA_mgmt_ip>> -node
<<var_nodeB>>
```

Create jumbo frame MTU broadcast domain in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Remove data ports from default broadcast domain

The 10GbE data ports are used for iSCSI/NFS traffic, and these ports should be removed from the default domain. Ports e0e and e0f are not used and should also be removed from the default domain.

To remove the ports from the broadcast domain, run the following command:

```
broadcast-domain remove-ports -broadcast-domain Default -ports
<<var_nodeA>>:e0c, <<var_nodeA>>:e0d, <<var_nodeA>>:e0e,
<<var_nodeA>>:e0f, <<var_nodeB>>:e0c, <<var_nodeB>>:e0d,
<<var_nodeA>>:e0e, <<var_nodeA>>:e0f
```

Disable flow control on UTA2 ports

It is a NetApp best practice to disable flow control on all UTA2 ports that are connected to external devices. To disable flow control, run the following command:

```

net port modify -node <<var_nodeA>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier.
Do you want to continue? {y|n}: y

```

Configure IFGRP LACP in ONTAP

This type of interface group requires two or more Ethernet interfaces and a switch that supports LACP. Make sure the switch is properly configured.

From the cluster prompt, complete the following steps.

```

ifgrp create -node <<var_nodeA>> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_nodeA>> -ifgrp a0a -port e0d
ifgrp create -node << var_nodeB>> -ifgrp a0a -distr-func port -mode
multimode_lacp
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0c
network port ifgrp add-port -node <<var_nodeB>> -ifgrp a0a -port e0d

```

Configure jumbo frames in NetApp ONTAP

To configure an ONTAP network port to use jumbo frames (that usually have an MTU of 9,000 bytes), run the following commands from the cluster shell:

```

AFF A220::> network port modify -node node_A -port a0a -mtu 9000
Warning: This command will cause a several second interruption of service
on
        this network port.
Do you want to continue? {y|n}: y
AFF A220::> network port modify -node node_B -port a0a -mtu 9000
Warning: This command will cause a several second interruption of service
on
        this network port.
Do you want to continue? {y|n}: y

```

Create VLANs in ONTAP

To create VLANs in ONTAP, complete the following steps:

1. Create NFS VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_nfs_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_nfs_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_NFS -ports
<<var_nodeA>>:a0a-<<var_nfs_vlan_id>>, <<var_nodeB>>:a0a-
<<var_nfs_vlan_id>>

```

2. Create iSCSI VLAN ports and add them to the data broadcast domain.


```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<var_iscsi_vlan_B_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<var_iscsi_vlan_B_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports
<<var_nodeA>>:a0a-<<var_iscsi_vlan_A_id>>, <<var_nodeB>>:a0a-
<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports
<<var_nodeA>>:a0a-<<var_iscsi_vlan_B_id>>, <<var_nodeB>>:a0a-
<<var_iscsi_vlan_B_id>>

```

3. Create MGMT-VLAN ports.

```

network port vlan create -node <<var_nodeA>> -vlan-name a0a-
<<mgmt_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name a0a-
<<mgmt_vlan_id>>

```

Create aggregates in ONTAP

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it contains.

To create aggregates, run the following commands:

```

aggr create -aggregate aggr1_nodeA -node <<var_nodeA>> -diskcount
<<var_num_disks>>
aggr create -aggregate aggr1_nodeB -node <<var_nodeB>> -diskcount
<<var_num_disks>>

```

Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.

Start with five disks; you can add disks to an aggregate when additional storage is required.

The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display the aggregate creation status. Do not proceed until `aggr1`_`nodeA` is online.

Configure time zone in ONTAP

To configure time synchronization and to set the time zone on the cluster, run the following command:

```
timezone <<var_timezone>>
```



For example, in the eastern United States, the time zone is `America/New York`. After you begin typing the time zone name, press the Tab key to see available options.

Configure SNMP in ONTAP

To configure the SNMP, complete the following steps:

1. Configure SNMP basic information, such as the location and contact. When polled, this information is visible as the `sysLocation` and `sysContact` variables in SNMP.

```
snmp contact <<var_snmp_contact>>
snmp location "<<var_snmp_location>>"
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts.

```
snmp traphost add <<var_snmp_server_fqdn>>
```

Configure SNMPv1 in ONTAP

To configure SNMPv1, set the shared secret plain-text password called a community.

```
snmp community add ro <<var_snmp_community>>
```



Use the `snmp community delete all` command with caution. If community strings are used for other monitoring products, this command removes them.

Configure SNMPv3 in ONTAP

SNMPv3 requires that you define and configure a user for authentication. To configure SNMPv3, complete the following steps:

1. Run the `security snmpusers` command to view the engine ID.
2. Create a user called `snmpv3user`.

```
security login create -username snmpv3user -authmethod usm -application snmp
```

3. Enter the authoritative entity's engine ID and select `md5` as the authentication protocol.
4. Enter an eight-character minimum-length password for the authentication protocol when prompted.
5. Select `des` as the privacy protocol.
6. Enter an eight-character minimum-length password for the privacy protocol when prompted.

Configure AutoSupport HTTPS in ONTAP

The NetApp AutoSupport tool sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts <<var_mailhost>> -transport https -support enable -noteto <<var_storage_admin_email>>
```

Create a storage virtual machine

To create an infrastructure storage virtual machine (SVM), complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate aggr1_nodeA -rootvolume-security-style unix
```

2. Add the data aggregate to the infra-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Infra-SVM -aggr-list aggr1_nodeA,aggr1_nodeB
```

3. Remove the unused storage protocols from the SVM, leaving NFS and iSCSI.

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,ndmp,fc
```

4. Enable and run the NFS protocol in the infra-SVM SVM.

```
`nfs create -vserver Infra-SVM -udp disabled`
```

5. Turn on the `SVM vstorage` parameter for the NetApp NFS VAAI plug-in. Then, verify that NFS has been configured.

```
`vserver nfs modify -vserver Infra-SVM -vstorage enabled`  
`vserver nfs show`
```



Commands are prefaced by `vserver` in the command line because storage virtual machines were previously called servers.

Configure NFSv3 in ONTAP

The following table lists the information needed to complete this configuration.

| Detail | Detail value |
|----------------------------|---------------------------|
| ESXi host A NFS IP address | <<var_esxi_hostA_nfs_ip>> |
| ESXi host B NFS IP address | <<var_esxi_hostB_nfs_ip>> |

To configure NFS on the SVM, run the following commands:

1. Create a rule for each ESXi host in the default export policy.
2. For each ESXi host being created, assign a rule. Each host has its own rule index. Your first ESXi host has rule index 1, your second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default  
-ruleindex 1 -protocol nfs -clientmatch <<var_esxi_hostA_nfs_ip>>  
-rorule sys -rwrule sys -superuser sys -allow-suid false  
vserver export-policy rule create -vserver Infra-SVM -policyname default  
-ruleindex 2 -protocol nfs -clientmatch <<var_esxi_hostB_nfs_ip>>  
-rorule sys -rwrule sys -superuser sys -allow-suid false  
vserver export-policy rule show
```

3. Assign the export policy to the infrastructure SVM root volume.

```
volume modify -vserver Infra-SVM -volume rootvol -policy default
```



The NetApp VSC automatically handles export policies if you choose to install it after vSphere has been set up. If you do not install it, you must create export policy rules when additional Cisco UCS C-Series servers are added.

Create iSCSI service in ONTAP

To create the iSCSI service, complete the following step:

1. Create the iSCSI service on the SVM. This command also starts the iSCSI service and sets the iSCSI IQN for the SVM. Verify that iSCSI has been configured.

```
iscsi create -vserver Infra-SVM
iscsi show
```

Create load-sharing mirror of SVM root volume in ONTAP

1. Create a volume to be the load- sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver Infra_Vserver -volume rootvol_m01 -aggregate
aggr1_nodeA -size 1GB -type DP
volume create -vserver Infra_Vserver -volume rootvol_m02 -aggregate
aggr1_nodeB -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:rootvol -destination-path
Infra-SVM:rootvol_m01 -type LS -schedule 15min
snapmirror create -source-path Infra-SVM:rootvol -destination-path
Infra-SVM:rootvol_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship and verify that it has been created.

```
snapmirror initialize-ls-set -source-path Infra-SVM:rootvol
snapmirror show
```

Configure HTTPS access in ONTAP

To configure secure access to the storage controller, complete the following steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS FQDN of the SVM. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a certificate authority.

Deleting expired certificates before creating certificates is a best practice. Run the `security certificate delete` command to delete expired certificates. In the following command, use TAB completion to select and delete each default certificate.

```
security certificate delete [TAB] ...  
Example: security certificate delete -vserver Infra-SVM -common-name  
Infra-SVM -ca Infra-SVM -type server -serial 552429A6
```

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the infra-SVM and the cluster SVM. Again, use TAB completion to aid in completing these commands.

```
security certificate create [TAB] ...  
Example: security certificate create -common-name infra-svm.netapp.com  
-type server -size 2048 -country US -state "North Carolina" -locality  
"RTP" -organization "NetApp" -unit "FlexPod" -email-addr  
"abc@netapp.com" -expire-days 365 -protocol SSL -hash-function SHA256  
-vserver Infra-SVM
```

5. To obtain the values for the parameters required in the following step, run the `security certificate show` command.
6. Enable each certificate that was just created using the `-server-enabled true` and `-client-enabled false` parameters. Again, use TAB completion.

```
security ssl modify [TAB] ...  
Example: security ssl modify -vserver Infra-SVM -server-enabled true  
-client-enabled false -ca infra-svm.netapp.com -serial 55243646 -common  
-name infra-svm.netapp.com
```

7. Configure and enable SSL and HTTPS access and disable HTTP access.

```

system services web modify -external true -sslv3-enabled true
Warning: Modifying the cluster configuration will cause pending web
service requests to be
        interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
system services firewall policy delete -policy mgmt -service http
-vserver <<var_clustername>>

```



It is normal for some of these commands to return an error message stating that the entry does not exist.

8. Revert to the admin privilege level and create the setup to allow SVM to be available by the web.

```

set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled
true

```

Create a NetApp FlexVol volume in ONTAP

To create a NetApp FlexVol volume, enter the volume name, size, and the aggregate on which it exists. Create two VMware datastore volumes and a server boot volume.

```

volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate
aggr1_nodeA -size 500GB -state online -policy default -junction-path
/infra_datastore_1 -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_nodeA
-size 100GB -state online -policy default -junction-path /infra_swap
-space-guarantee none -percent-snapshot-space 0 -snapshot-policy none
volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_nodeA
-size 100GB -state online -policy default -space-guarantee none -percent
-snapshot-space 0

```

Enable deduplication in ONTAP

To enable deduplication on appropriate volumes, run the following commands:

```

volume efficiency on -vserver Infra-SVM -volume infra_datastore_1
volume efficiency on -vserver Infra-SVM -volume esxi_boot

```

Create LUNs in ONTAP

To create two boot LUNs, run the following commands:

```
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-A -size
15GB -ostype vmware -space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-B -size
15GB -ostype vmware -space-reserve disabled
```



When adding an extra Cisco UCS C-Series server, an extra boot LUN must be created.

Create iSCSI LIFs in ONTAP

The following table lists the information needed to complete this configuration.

| Detail | Detail value |
|--|---------------------------------|
| Storage node A iSCSI LIF01A | <<var_nodeA_iscsi_lif01a_ip>> |
| Storage node A iSCSI LIF01A network mask | <<var_nodeA_iscsi_lif01a_mask>> |
| Storage node A iSCSI LIF01B | <<var_nodeA_iscsi_lif01b_ip>> |
| Storage node A iSCSI LIF01B network mask | <<var_nodeA_iscsi_lif01b_mask>> |
| Storage node B iSCSI LIF01A | <<var_nodeB_iscsi_lif01a_ip>> |
| Storage node B iSCSI LIF01A network mask | <<var_nodeB_iscsi_lif01a_mask>> |
| Storage node B iSCSI LIF01B | <<var_nodeB_iscsi_lif01b_ip>> |
| Storage node B iSCSI LIF01B network mask | <<var_nodeB_iscsi_lif01b_mask>> |

1. Create four iSCSI LIFs, two on each node.


```

network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data
-data-protocol iscsi -home-node <<var_nodeA>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_nodeA_iscsi_lif01a_ip>> -netmask
<<var_nodeA_iscsi_lif01a_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data
-data-protocol iscsi -home-node <<var_nodeA>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_nodeA_iscsi_lif01b_ip>> -netmask
<<var_nodeA_iscsi_lif01b_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data
-data-protocol iscsi -home-node <<var_nodeB>> -home-port a0a-
<<var_iscsi_vlan_A_id>> -address <<var_nodeB_iscsi_lif01a_ip>> -netmask
<<var_nodeB_iscsi_lif01a_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data
-data-protocol iscsi -home-node <<var_nodeB>> -home-port a0a-
<<var_iscsi_vlan_B_id>> -address <<var_nodeB_iscsi_lif01b_ip>> -netmask
<<var_nodeB_iscsi_lif01b_mask>> -status-admin up -failover-policy
disabled -firewall-policy data -auto-revert false
network interface show

```

Create NFS LIFs in ONTAP

The following table lists the information needed to complete this configuration.

| Detail | Detail Value |
|--|-------------------------------|
| Storage node A NFS LIF 01 IP | <<var_nodeA_nfs_lif_01_ip>> |
| Storage node A NFS LIF 01 network mask | <<var_nodeA_nfs_lif_01_mask>> |
| Storage node B NFS LIF 02 IP | <<var_nodeB_nfs_lif_02_ip>> |
| Storage node B NFS LIF 02 network mask | <<var_nodeB_nfs_lif_02_mask>> |

1. Create an NFS LIF.

```

network interface create -vserver Infra-SVM -lif nfs_lif01 -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_nodeA_nfs_lif_01_ip>> -netmask <<
var_nodeA_nfs_lif_01_mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif02 -role data
-data-protocol nfs -home-node <<var_nodeA>> -home-port a0a-
<<var_nfs_vlan_id>> -address <<var_nodeB_nfs_lif_02_ip>> -netmask <<
var_nodeB_nfs_lif_02_mask>> -status-admin up -failover-policy broadcast-
domain-wide -firewall-policy data -auto-revert true
network interface show

```

Add infrastructure SVM administrator

The following table lists the information needed to complete this configuration.

| Detail | Detail Value |
|------------------------|--------------------------|
| Vsmgmt IP | <<var_svm_mgmt_ip>> |
| Vsmgmt network mask | <<var_svm_mgmt_mask>> |
| Vsmgmt default gateway | <<var_svm_mgmt_gateway>> |

To add the infrastructure SVM administrator and SVM administration logical interface to the management network, complete the following steps:

1. Run the following command:

```

network interface create -vserver Infra-SVM -lif vsmgmt -role data
-data-protocol none -home-node <<var_nodeB>> -home-port e0M -address
<<var_svm_mgmt_ip>> -netmask <<var_svm_mgmt_mask>> -status-admin up
-failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-
revert true

```



The SVM management IP here should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```

network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway
<<var_svm_mgmt_gateway>>
network route show

```

3. Set a password for the SVM vsadmin user and unlock the user.

```
security login password -username vsadmin -vserver Infra-SVM
Enter a new password: <<var_password>>
Enter it again: <<var_password>>
security login unlock -username vsadmin -vserver Infra-SVM
```

[Next: Cisco UCS C-Series Rack Server Deployment Procedure](#)

Cisco UCS C-Series rack server deployment procedure

The following section provides a detailed procedure for configuring a Cisco UCS C-Series standalone rack server for use in the FlexPod Express configuration.

Perform initial Cisco UCS C-Series standalone server setup for Cisco Integrated Management Server

Complete these steps for the initial setup of the CIMC interface for Cisco UCS C-Series standalone servers.

The following table lists the information needed to configure CIMC for each Cisco UCS C-Series standalone server.


| Detail | Detail value |
|----------------------|------------------|
| CIMC IP address | <<cimc_ip>> |
| CIMC subnet mask | <<cimc_netmask>> |
| CIMC default gateway | <<cimc_gateway>> |



The CIMC version used in this validation is CIMC 3.1.3(g).

All servers

1. Attach the Cisco keyboard, video, and mouse (KVM) dongle (provided with the server) to the KVM port on the front of the server. Plug a VGA monitor and USB keyboard into the appropriate KVM dongle ports.
2. Power on the server and press F8 when prompted to enter the CIMC configuration.



3. In the CIMC configuration utility, set the following options:
- Network interface card (NIC) mode:
 - Dedicated ☒ [X]
 - IP (Basic):
 - IPV4: ☒ [X]
 - DHCP enabled: ☐ []
 - CIMC IP: <<cimc_ip>>
 - Prefix/Subnet: <<cimc_netmask>>
 - Gateway: <<cimc_gateway>>
 - VLAN (Advanced): Leave cleared to disable VLAN tagging.
 - NIC redundancy
 - None: ☒ [X]

```
Cisco IMC Configuration Utility Version 2.0 Cisco Systems, Inc.
*****
NIC Properties
NIC mode
Dedicated:      [X]
Shared LOM:     [ ]
Cisco Card:
  Riser1:       [ ]
  Riser2:       [ ]
  MLOm:         [ ]
Shared LOM Ext: [ ]
NIC redundancy
None:           [X]
Active-standby: [ ]
Active-active:  [ ]
VLAN (Advanced)
VLAN enabled:   [ ]
VLAN ID:        1
Priority:        0
IP (Basic)
IPv4:           [X]
DHCP enabled:   [ ]
CIMC IP:        10.61.185.215
Prefix/Subnet:  255.255.255.0
Gateway:        10.61.185.1
Pref DNS Server: 0.0.0.0
Smart Access USB
Enabled         [ ]
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F1>Additional settings
```

4. Press F1 to see additional settings.

- Common properties:
 - Host name: <<esxi_host_name>>
 - Dynamic DNS: []
 - Factory defaults: Leave cleared.
- Default user (basic):
 - Default password: <<admin_password>>
 - Reenter password: <<admin_password>>
 - Port properties: Use default values.
 - Port profiles: Leave cleared.

```

Cisco IMC Configuration Utility Version 2.0  Cisco Systems, Inc.
*****
Common Properties
  Hostname:      CIMC-Tiger-02
  Dynamic DNS:   [X]
  DDNS Domain:
FactoryDefaults
  Factory Default:      [ ]
Default User(Basic)
  Default password:      -
  Reenter password:
Port Properties
  Auto Negotiation:      [X]
                                Admin Mode      Operation Mode
  Speed[1000/100/10Mbps]:      Auto              1000
  Duplex mode[half/full]:      Auto              full
Port Profiles
  Reset:                  [ ]
  Name:
*****
<Up/Down>Selection  <F10>Save  <Space>Enable/Disable  <F5>Refresh  <ESC>Exit
<F2>PreviousPageettings

```

5. Press F10 to save the CIMC interface configuration.
6. After the configuration is saved, press Esc to exit.

Configure Cisco UCS C-Series servers iSCSI boot

In this FlexPod Express configuration, the VIC1387 is used for iSCSI boot.

The following table lists the information needed to configure iSCSI boot.



Italicized font indicates variables that are unique for each ESXi host.

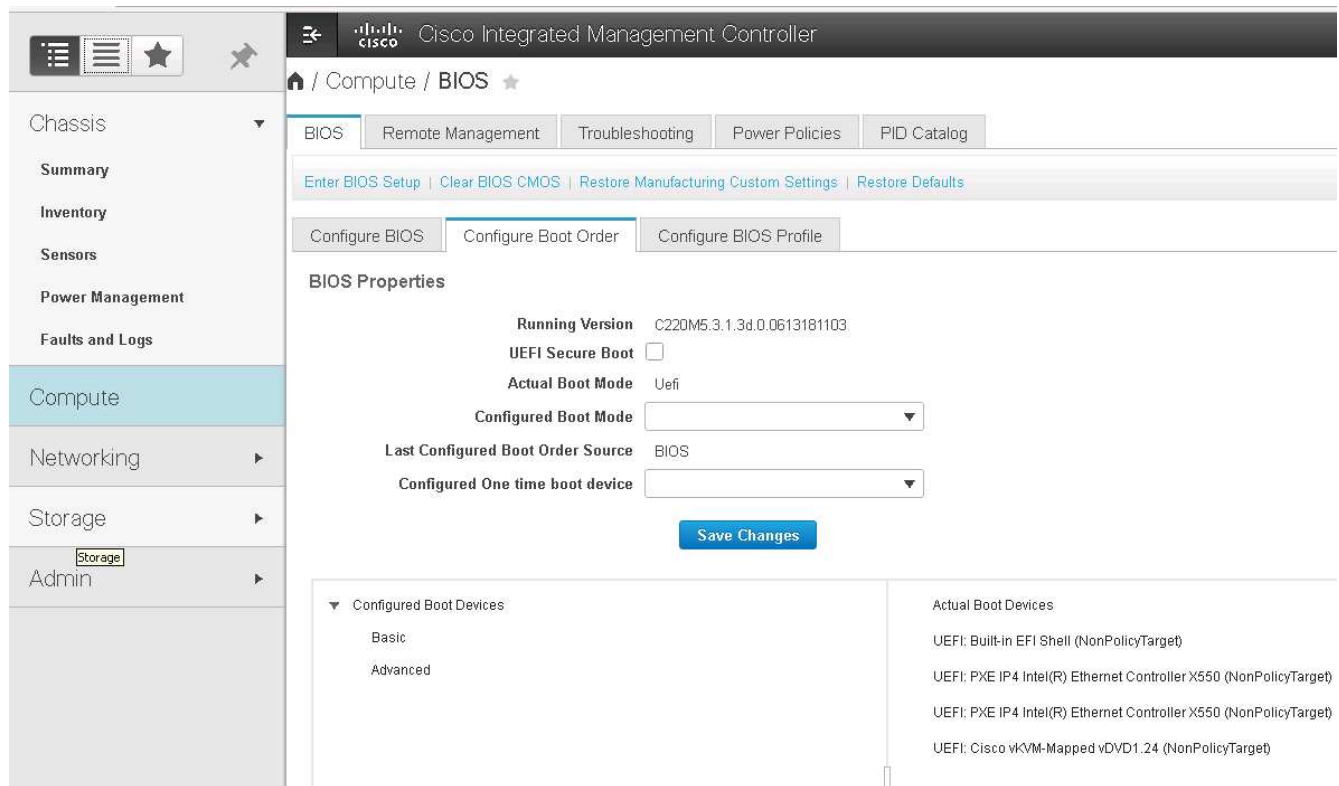
| Detail | Detail value |
|-----------------------------------|----------------------------------|
| ESXi host initiator A name | <<var_ucs_initiator_name_A>> |
| ESXi host iSCSI-A IP | <<var_esxi_host_iscsiA_ip>> |
| ESXi host iSCSI-A network mask | <<var_esxi_host_iscsiA_mask>> |
| ESXi host iSCSI A default gateway | <<var_esxi_host_iscsiA_gateway>> |
| ESXi host initiator B name | <<var_ucs_initiator_name_B>> |
| ESXi host iSCSI-B IP | <<var_esxi_host_iscsiB_ip>> |
| ESXi host iSCSI-B network mask | <<var_esxi_host_iscsiB_mask>> |
| ESXi host iSCSI-B gateway | <<var_esxi_host_iscsiB_gateway>> |

| Detail | Detail value |
|-------------------------|--------------|
| IP address iscsi_lif01a | |
| IP address iscsi_lif02a | |
| IP address iscsi_lif01b | |
| IP address iscsi_lif02b | |
| Infra_SVM IQN | |

Boot order configuration

To set the boot order configuration, complete the following steps:

1. From the CIMC interface browser window, click the Server tab and select BIOS.
2. Click Configure Boot Order and then click OK.



3. Configure the following devices by clicking the device under Add Boot Device, and going to the Advanced tab.
 - Add Virtual Media
 - Name: KVM-CD-DVD
 - Subtype: KVM MAPPED DVD
 - State: Enabled
 - Order: 1
 - Add iSCSI Boot.
 - Name: iSCSI-A

- State: Enabled
- Order: 2
- Slot: MLOM
- Port: 0
- Click Add iSCSI Boot.
 - Name: iSCSI-B
 - State: Enabled
 - Order: 3
 - Slot: MLOM
 - Port: 1

4. Click Add Device.

5. Click Save Changes and then click Close.

Configure Boot Order

Configured Boot Level: Advanced

Basic Advanced

Add Boot Device

- Add Local HDD
- Add PXE Boot
- Add SAN Boot
- Add iSCSI Boot
- Add USB
- Add Virtual Media
- Add PCHStorage
- Add UEFISHELL
- Add SD Card
- Add NVME
- Add Local CDD

Advanced Boot Order Configuration

Selected 1 / Total 3

| | Name | Type | Order | State |
|-------------------------------------|----------------|--------|-------|---------|
| <input checked="" type="checkbox"/> | KVM-MAPPED-DVD | VMEDIA | 1 | Enabled |
| <input type="checkbox"/> | iSCSI-A | ISCSI | 2 | Enabled |
| <input type="checkbox"/> | iSCSI-B | ISCSI | 3 | Enabled |

Save Changes Reset Values Close

6. Reboot the server to boot with your new boot order.

Disable RAID controller (if present)

Complete the following steps if your C-Series server contains a RAID controller. A RAID controller is not needed in the boot from SAN configuration. Optionally, you can also physically remove the RAID controller from the server.

1. Click BIOS on the left navigation pane in CIMC.
2. Select Configure BIOS.
3. Scroll down to PCIe Slot:HBA Option ROM.
4. If the value is not already disabled, set it to disabled.

| | | | | | |
|------|-------------------|-----------------|----------------|-------------|-------------------|
| BIOS | Remote Management | Troubleshooting | Power Policies | PID Catalog | |
| I/O | Server Management | Security | Processor | Memory | Power/Performance |

Note: Default values are shown in bold.

Reboot Host Immediately: ☒

| | |
|---------------------------|---------------|
| Intel VT for directed IO: | Enabled ▼ |
| Intel VTD ATS support: | Enabled ▼ |
| LOM Port 1 OptionRom: | Enabled ▼ |
| Pcie Slot 1 OptionRom: | Disabled ▼ |
| MLOM OptionRom: | Enabled ▼ |
| Front NVME 1 OptionRom: | Enabled ▼ |
| MRAID Link Speed: | Auto ▼ |
| PCIe Slot 1 Link Speed: | Auto ▼ |
| Front NVME 1 Link Speed: | Auto ▼ |
| VGA Priority: | Onboard ▼ |
| P-SATA OptionROM: | LSI SW RAID ▼ |
| USB Port Rear: | Enabled ▼ |
| USB Port Internal: | Enabled ▼ |
| IPv6 PXE Support: | Disabled ▼ |

| | |
|------------------------------|------------|
| Legacy USB Support: | Enabled ▼ |
| Intel VTD coherency support: | Disabled ▼ |
| All Onboard LOM Ports: | Enabled ▼ |
| LOM Port 2 OptionRom: | Enabled ▼ |
| Pcie Slot 2 OptionRom: | Disabled ▼ |
| MRAID OptionRom: | Enabled ▼ |
| Front NVME 2 OptionRom: | Enabled ▼ |
| MLOM Link Speed: | Auto ▼ |
| PCIe Slot 2 Link Speed: | Auto ▼ |
| Front NVME 2 Link Speed: | Auto ▼ |
| M.2 SATA OptionROM: | AHCI ▼ |
| USB Port Front: | Enabled ▼ |
| USB Port KVM: | Enabled ▼ |
| USB Port:M.2 Storage: | Enabled ▼ |

Configure Cisco VIC1387 for iSCSI boot

The following configuration steps are for the Cisco VIC 1387 for iSCSI boot.

Create iSCSI vNICs

1. Click Add to create a vNIC.
2. In the Add vNIC section, enter the following settings:
 - Name: iSCSI-vNIC-A
 - MTU: 9000
 - Default VLAN: <<var_iscsi_vlan_a>>
 - VLAN Mode: TRUNK
 - Enable PXE boot: Check

▼ vNIC Properties

▼ General

Name:

iSCSI-vNIC-A

CDN:

VIC-MLOM-iSCSI-vNIC-A

MTU:

9000

(1500 - 9000)

Uplink Port:

0 ▼

MAC Address:

☐ Auto
 ☒ 70:69:5A:C0:98:ED

Class of Service:

0

(0 - 6)

Trust Host CoS:

☒

PCI Order:

4

(0 - 5)

Default VLAN:

☐ None
 ☒ 3439

VLAN Mode:

Trunk ▼

Rate Limit:

☒ OFF
 ☐

Channel Number:

N/A

(1 - 1000)

PCI Link:

0

(0 - 1)

Enable NVGRE:

☐

Enable VXLAN:

☐

Advanced Filter:

☐

Port Profile:

N/A ▼

Enable PXE Boot:

☒

Enable VMQ:

☐

Enable aRFS:

☐

Enable Uplink Failover:

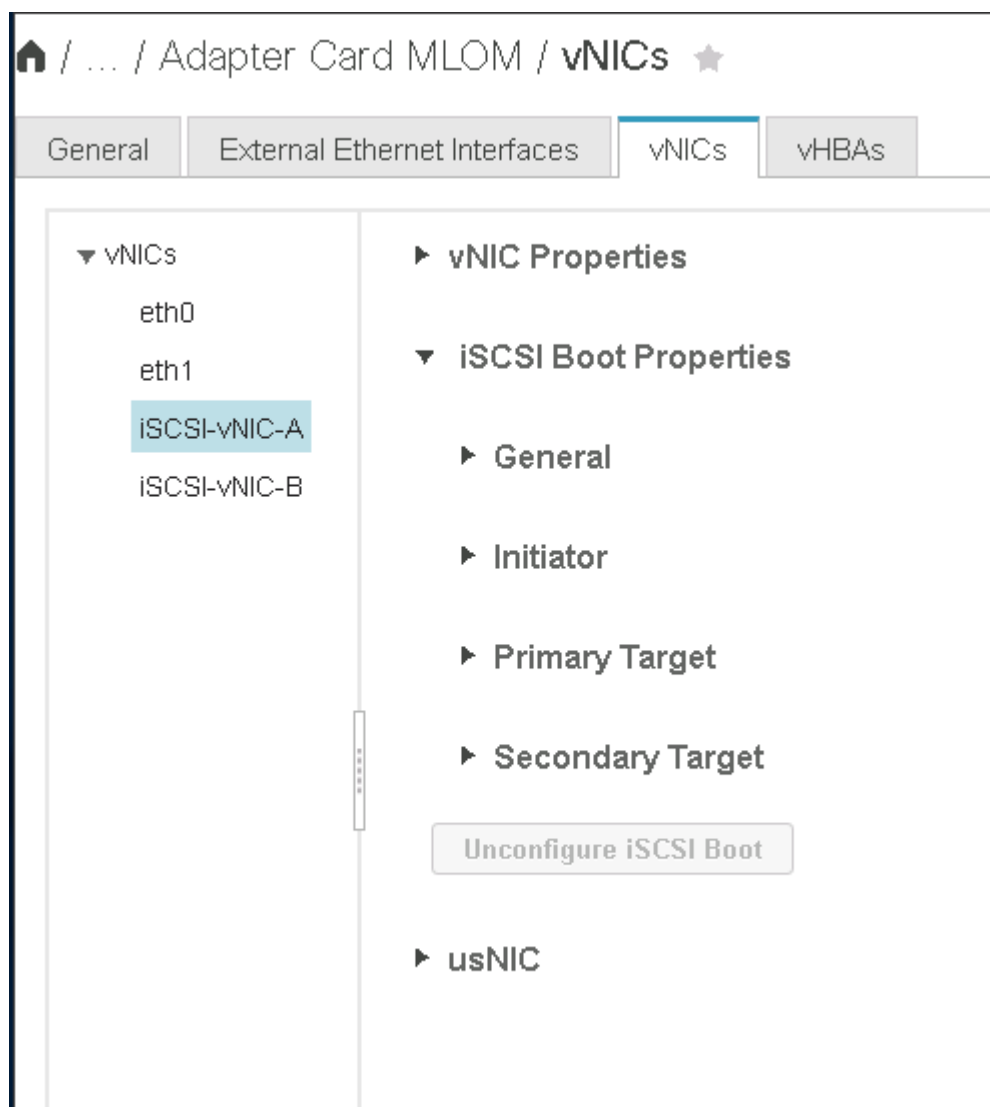
☐

Failback Timeout:

N/A

(0 - 600)

3. Click Add vNIC and then click OK.
4. Repeat the process to add a second vNIC.
 - a. Name the vNIC iSCSI-vNIC-B.
 - b. Enter <<var_iscsi_vlan_b>> as the VLAN.
 - c. Set the uplink port to 1.
5. Select the vNIC iSCSI-vNIC-A on the left.



6. Under iSCSI Boot Properties, enter the initiator details:
 - Name: <<var_ucsa_initiator_name_a>>
 - IP address: <<var_esxi_hostA_iscsiA_ip>>
 - Subnet mask: <<var_esxi_hostA_iscsiA_mask>>
 - Gateway: <<var_esxi_hostA_iscsiA_gateway>>

vNICs

eth0
eth1
ISCSI-v
ISCSI-v

ISCSI Boot Properties

General

Initiator

Name: (0 - 233) chars
Initiator Priority:

IP Address:
Secondary DNS:

Subnet Mask:
TCP Timeout:

Gateway:
CHAP Name:

Primary DNS:
CHAP Secret:

Primary Target

Secondary Target

7. Enter the primary target details.

- Name: IQN number of infra-SVM
- IP address: IP address of `iscsi_lif01a`
- Boot LUN: 0

8. Enter the secondary target details.

- Name: IQN number of infra-SVM
- IP address: IP address of `iscsi_lif02a`
- Boot LUN: 0

You can obtain the storage IQN number by running the `vserver iscsi show` command.



Be sure to record the IQN names for each vNIC. You need them for a later step.

General
External Ethernet Interfaces
vNICs
vHBAs

▼ vNICs
eth0
eth1
iSCSI-v
iSCSI-v

► Initiator

▼ Primary Target

Name: (0 - 233) chars
Boot LUN:

IP Address:
CHAP Name:

TCP Port: 3260
CHAP Secret:

▼ Secondary Target

Name: (0 - 233) chars
Boot LUN:

IP Address:
CHAP Name:

TCP Port: 3260
CHAP Secret:

Unconfigure iSCSI Boot

9. Click Configure iSCSI.

10. Select the vNIC `iSCSI-vNIC-B` and click the iSCSI Boot button located on the top of the Host Ethernet Interfaces section.

11. Repeat the process to configure `iSCSI-vNIC-B`.

12. Enter the initiator details.

- Name: `<<var_ucs_a_initiator_name_b>>`
- IP address: `<<var_esxi_hostb_iscsib_ip>>`
- Subnet mask: `<<var_esxi_hostb_iscsib_mask>>`
- Gateway: `<<var_esxi_hostb_iscsib_gateway>>`

13. Enter the primary target details.

- Name: IQN number of infra-SVM
- IP address: IP address of `iscsi_lif01b`
- Boot LUN: 0

14. Enter the secondary target details.

- Name: IQN number of infra-SVM
- IP address: IP address of `iscsi_lif02b`
- Boot LUN: 0

You can obtain the storage IQN number by using the `vserver iscsi show` command.



Be sure to record the IQN names for each vNIC. You need them for a later step.

15. Click Configure iSCSI.

16. Repeat this process to configure iSCSI boot for Cisco UCS server B.

Configure vNICs for ESXi

- 1. From the CIMC interface browser window, click Inventory and then click Cisco VIC adapters on the right pane.
- 2. Under Adapter Cards, select Cisco UCS VIC 1387 and then select the vNICs underneath.

Home / ... / Adapter Card

MLOM / vNICs ★

Refresh | Host Power | Launch KVM | Ping | CIMC Reboot | Locat

General | External Ethernet Interfaces | vNICs | vHBAs

▼ vNICs

eth0

eth1

iSCSI-v

iSCSI-v

Host Ethernet Interfaces

Selected 0

Add vNIC | Clone vNIC | Delete vNICs

| | Name | CDN | MAC Address | MTU | usNIC | Uplink Port | CoS | VLAN | VLAN Mode |
|--------------------------|------------|------------|-------------------|------|-------|-------------|-----|------|-----------|
| <input type="checkbox"/> | eth0 | VIC-MLO... | 70:69:5A:C0:98:49 | 1500 | 0 | 0 | 0 | NONE | TRUNK |
| <input type="checkbox"/> | eth1 | VIC-MLO... | 70:69:5A:C0:98:4A | 1500 | 0 | 1 | 0 | NONE | TRUNK |
| <input type="checkbox"/> | iSCSI-v... | VIC-MLO... | 70:69:5A:C0:98:4D | 9000 | 0 | 0 | 0 | 3439 | TRUNK |
| <input type="checkbox"/> | iSCSI-v... | VIC-MLO... | 70:69:5A:C0:98:4E | 9000 | 0 | 1 | 0 | 3440 | TRUNK |

- 3. Select eth0 and click Properties.
- 4. Set the MTU to 9000. Click Save Changes.

GeneralExternal Ethernet InterfacesvNICsvHBAs

▼ vNICs

eth0

eth1

ISCSI-v

ISCSI-v

Name: eth0

CDN: VIC-MLOM-eth0

MTU: 9000 (1500 - 9000)

Uplink Port: 0 ▼

MAC Address: ☐ Auto ☒ 70:69:5A:C0:98:49

Class of Service: 0 (0 - 6)

Trust Host CoS: ☐

PCI Order: 0 (0 - 5)

Default VLAN: ☒ None ☐ ?

5. Repeat steps 3 and 4 for eth1, verifying that the uplink port is set to 1 for eth1.

[/ ... / Adapter Card MLOM / vNICs](#) ★

GeneralExternal Ethernet InterfacesvNICsvHBAs

▼ vNICs

eth0

eth1

ISCSI-vNIC-A

ISCSI-vNIC-B

Host Ethernet Interfaces

Add vNICClone vNICDelete vNICs

| | Name | CDN | MAC Address | MTU | usNIC | Uplink Port |
|--------------------------|------------|------------|-------------------|------|-------|-------------|
| <input type="checkbox"/> | eth0 | VIC-MLO... | 70:69:5A:C0:98:49 | 9000 | 0 | 0 |
| <input type="checkbox"/> | eth1 | VIC-MLO... | 70:69:5A:C0:98:4A | 9000 | 0 | 1 |
| <input type="checkbox"/> | iSCSI-v... | VIC-MLO... | 70:69:5A:C0:98:4D | 9000 | 0 | 0 |
| <input type="checkbox"/> | iSCSI-v... | VIC-MLO... | 70:69:5A:C0:98:4E | 9000 | 0 | 1 |



This procedure must be repeated for each initial Cisco UCS Server node and each additional Cisco UCS Server node added to the environment.

NetApp AFF Storage Deployment Procedure (Part 2)

ONTAP SAN boot storage setup

Create iSCSI igroups

To create igroups, complete the following step:

You need the iSCSI initiator IQNs from the server configuration for this step.

1. From the cluster management node SSH connection, run the following commands. To view the three igroups created in this step, run the `igroup show` command.

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-A -protocol iscsi
-ostype vmware -initiator <<var_vm_host_infra_a_iSCSI-A_vNIC_IQN>>,
<<var_vm_host_infra_a_iSCSI-B_vNIC_IQN>>
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-B -protocol iscsi
-ostype vmware -initiator <<var_vm_host_infra_b_iSCSI-A_vNIC_IQN>>,
<<var_vm_host_infra_b_iSCSI-B_vNIC_IQN>>
```



This step must be completed when adding additional Cisco UCS C- Series servers.

Map boot LUNs to igroups

To map boot LUNs to igroups, run the following commands from the cluster management SSH connection:

```
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra- A -igroup
VM-Host-Infra- A -lun-id 0
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra- B -igroup
VM-Host-Infra- B -lun-id 0
```



This step must be completed when adding additional Cisco UCS C-Series servers.

[Next: VMware vSphere 6.7 Deployment Procedure.](#)

VMware vSphere 6.7 deployment procedure

This section provides detailed procedures for installing VMware ESXi 6.7 in a FlexPod Express configuration. The deployment procedures that follow are customized to include the environment variables described in previous sections.

Multiple methods exist for installing VMware ESXi in such an environment. This procedure uses the virtual KVM console and virtual media features of the CIMC interface for Cisco UCS C-Series servers to map remote installation media to each individual server.



This procedure must be completed for Cisco UCS server A and Cisco UCS server B.


This procedure must be completed for any additional nodes added to the cluster.

Log in to CIMC interface for Cisco UCS C-Series standalone servers

The following steps detail the method for logging in to the CIMC interface for Cisco UCS C-Series standalone servers. You must log in to the CIMC interface to run the virtual KVM, which enables the administrator to begin installation of the operating system through remote media.

All hosts

1. Navigate to a web browser and enter the IP address for the CIMC interface for the Cisco UCS C-Series. This step launches the CIMC GUI application.
2. Log in to the CIMC UI using the admin user name and credentials.
3. In the main menu, select the Server tab.
4. Click Launch KVM Console.



5. From the virtual KVM console, select the Virtual Media tab.
6. Select Map CD/DVD.



You might first need to click Activate Virtual Devices. Select Accept This Session if prompted.

7. Browse to the VMware ESXi 6.7 installer ISO image file and click Open. Click Map Device.
8. Select the Power menu and choose Power Cycle System (Cold Boot). Click Yes.

Install VMware ESXi

The following steps describe how to install VMware ESXi on each host.

Download ESXi 6.7 Cisco custom image

1. Navigate to the [VMware vSphere download page](#) for custom ISOs.
2. Click Go to Downloads next to the Cisco Custom Image for ESXi 6.7 GA Install CD.
3. Download the Cisco Custom Image for ESXi 6.7 GA Install CD (ISO).

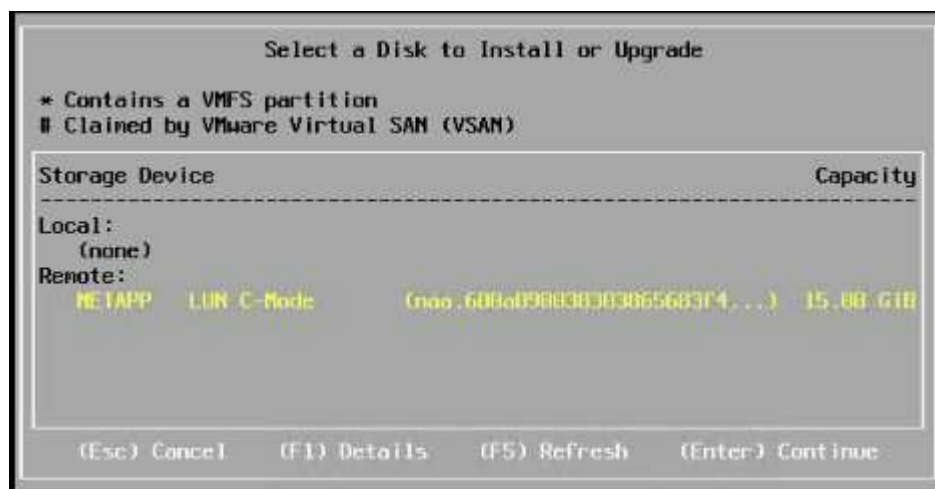
All hosts

1. When the system boots, the machine detects the presence of the VMware ESXi installation media.
2. Select the VMware ESXi installer from the menu that appears.

The installer loads. This takes several minutes.

3. After the installer has finished loading, press Enter to continue with the installation.

4. After reading the end-user license agreement, accept it and continue with the installation by pressing F11.
5. Select the NetApp LUN that was previously set up as the installation disk for ESXi, and press Enter to continue with the installation.



6. Select the appropriate keyboard layout and press Enter.
7. Enter and confirm the root password and press Enter.
8. The installer warns you that existing partitions are removed on the volume. Continue with the installation by pressing F11. The server reboots after the installation of ESXi.

Set up VMware ESXi host management networking


The following steps describe how to add the management network for each VMware ESXi host.

All hosts

1. After the server has finished rebooting, enter the option to customize the system by pressing F2.
2. Log in with root as the login name and the root password previously entered during the installation process.
3. Select the Configure Management Network option.
4. Select Network Adapters and press Enter.
5. Select the desired ports for vSwitch0. Press Enter.



Select the ports that correspond to eth0 and eth1 in CIMC.



6. Select VLAN (optional) and press Enter.
7. Enter the VLAN ID <<mgmt_vlan_id>>. Press Enter.
8. From the Configure Management Network menu, select IPv4 Configuration to configure the IP address of the management interface. Press Enter.
9. Use the arrow keys to highlight Set Static IPv4 address and use the space bar to select this option.
10. Enter the IP address for managing the VMware ESXi host <<esxi_host_mgmt_ip>>.
11. Enter the subnet mask for the VMware ESXi host <<esxi_host_mgmt_netmask>>.
12. Enter the default gateway for the VMware ESXi host <<esxi_host_mgmt_gateway>>.
13. Press Enter to accept the changes to the IP configuration.
14. Enter the IPv6 configuration menu.
15. Use the space bar to disable IPv6 by unselecting the Enable IPv6 (restart required) option. Press Enter.
16. Enter the menu to configure the DNS settings.
17. Because the IP address is assigned manually, the DNS information must also be entered manually.
18. Enter the primary DNS server's IP address [nameserver_ip].
19. (Optional) Enter the secondary DNS server's IP address.
20. Enter the FQDN for the VMware ESXi host name: [esxi_host_fqdn].
21. Press Enter to accept the changes to the DNS configuration.
22. Exit the Configure Management Network submenu by pressing Esc.
23. Press Y to confirm the changes and reboot the server.
24. Log out of the VMware Console by pressing Esc.

Configure ESXi host

You need the information in the following table to configure each ESXi host.

| Detail | Value |
|----------------|-------|
| ESXi host name | |


| Detail | Value |
|------------------------------|-------|
| ESXi host management IP | |
| ESXi host management mask | |
| ESXi host management gateway | |
| ESXi host NFS IP | |
| ESXi host NFS mask | |
| ESXi host NFS gateway | |
| ESXi host vMotion IP | |
| ESXi host vMotion mask | |
| ESXi host vMotion gateway | |
| ESXi host iSCSI-A IP | |
| ESXi host iSCSI-A mask | |
| ESXi host iSCSI-A gateway | |
| ESXi host iSCSI-B IP | |
| ESXi host iSCSI-B mask | |
| ESXi host iSCSI-B gateway | |

Log in to ESXi host

1. Open the host's management IP address in a web browser.
2. Log in to the ESXi host using the root account and the password you specified during the install process.
3. Read the statement about the VMware Customer Experience Improvement Program. After selecting the proper response, click OK.

Configure iSCSI boot

1. Select Networking on the left.
2. On the right, select the Virtual Switches tab.



3. Click iScsiBootvSwitch.
4. Select Edit settings.
5. Change the MTU to 9000 and click Save.
6. Click Networking in the left navigation pane to return to the Virtual Switches tab.
7. Click Add Standard Virtual Switch.
8. Provide the name iScsiBootvSwitch-B for the vSwitch name.
 - Set the MTU to 9000.
 - Select vmnic3 from the Uplink 1 options.
 - Click Add.



Vmnic2 and vmnic3 are used for iSCSI boot in this configuration. If you have additional NICs in your ESXi host, you might have different vmnic numbers. To confirm which NICs are used for iSCSI boot, match the MAC addresses on the iSCSI vNICs in CIMC to the vmnics in ESXi.

9. In the center pane, select the VMkernel NICs tab.
10. Select Add VMkernel NIC.
 - Specify a new port group name of iScsiBootPG-B.
 - Select iScsiBootvSwitch-B for the virtual switch.
 - Enter <<iscsib_vlan_id>> for the VLAN ID.
 - Change the MTU to 9000.
 - Expand IPv4 Settings.
 - Select Static Configuration.
 - Enter <<var_hosta_iscsib_ip>> for Address.
 - Enter <<var_hosta_iscsib_mask>> for Subnet Mask.
 - Click Create.

| | |
|-----------------|--|
| Port group | New port group ▼ |
| New port group | iScsiBootPG-B |
| Virtual switch | iScsiBootvSwitch-B ▼ |
| VLAN ID | 3440 |
| MTU | 9000 |
| IP version | IPv4 only ▼ |
| ▼ IPv4 settings | |
| Configuration | <input type="radio"/> DHCP <input checked="" type="radio"/> Static |
| Address | 172.21.184.63 |
| Subnet mask | 255.255.255.0 |
| TCP/IP stack | Default TCP/IP stack ▼ |
| Services | <input type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication |

Create Cancel




Set the MTU to 9000 on iScsiBootPG- A.

Configure iSCSI multipathing

To set up iSCSI multipathing on the ESXi hosts, complete the following steps:

1. Select Storage in the left navigation pane. Click Adapters.
2. Select the iSCSI software adapter and click Configure iSCSI.



3. Under Dynamic Targets, click Add Dynamic Target.

Configure iSCSI - vmhba64

| iSCSI enabled | <input type="radio"/> Disabled <input checked="" type="radio"/> Enabled | | | | | | | | |
|--|--|--------------|--|--------------|------------|--------------------|--|---------------|------|
| ▶ Name & alias | iqn.1992-08.com.cisco:ucsaiscsia | | | | | | | | |
| ▶ CHAP authentication | Do not use CHAP ▼ | | | | | | | | |
| ▶ Mutual CHAP authentication | Do not use CHAP ▼ | | | | | | | | |
| ▶ Advanced settings | Click to expand | | | | | | | | |
| Network port bindings | <div> Add port binding Remove port binding </div> <table border="1"> <thead> <tr> <th>VMkernel NIC</th> <th>Port group</th> <th>IPv4 address</th> </tr> </thead> <tbody> <tr> <td colspan="3">No port bindings</td> </tr> </tbody> </table> | | | VMkernel NIC | Port group | IPv4 address | No port bindings | | |
| VMkernel NIC | Port group | IPv4 address | | | | | | | |
| No port bindings | | | | | | | | | |
| Static targets | <div> Add static target Remove static target Edit settings <input type="text" value="Search"/> </div> <table border="1"> <thead> <tr> <th>Target</th> <th>Address</th> <th>Port</th> </tr> </thead> <tbody> <tr> <td>iqn.1992-08.com.netapp:sn.09591199033811e78eb...</td> <td>172.21.183.34</td> <td>3260</td> </tr> </tbody> </table> | | | Target | Address | Port | iqn.1992-08.com.netapp:sn.09591199033811e78eb... | 172.21.183.34 | 3260 |
| Target | Address | Port | | | | | | | |
| iqn.1992-08.com.netapp:sn.09591199033811e78eb... | 172.21.183.34 | 3260 | | | | | | | |
| Dynamic targets | <div> Add dynamic target Remove dynamic target Edit settings <input type="text" value="Search"/> </div> <table border="1"> <thead> <tr> <th>Address</th> <th>Port</th> </tr> </thead> <tbody> <tr> <td colspan="2">No dynamic targets</td> </tr> </tbody> </table> | | | Address | Port | No dynamic targets | | | |
| Address | Port | | | | | | | | |
| No dynamic targets | | | | | | | | | |

4. Enter the IP address `iscsi_lif01a`.
 - Repeat with the IP addresses `iscsi_lif01b`, `iscsi_lif02a`, and `iscsi_lif02b`.
 - Click Save Configuration.


| Dynamic targets | Add dynamic target Remove dynamic target Edit settings |
|-----------------|--|
| Address | Port |
| 172.21.183.33 | 3260 |
| 172.21.183.34 | 3260 |
| 172.21.184.33 | 3260 |
| 172.21.184.34 | 3260 |



You can find the iSCSI LIF IP addresses by running the ``network interface show`` command on the NetApp cluster or by looking at the Network Interfaces tab in OnCommand System Manager.

Configure ESXi host

1. In the left navigation pane, select Networking.
2. Select vSwitch0.



3. Select Edit Settings.
4. Change the MTU to 9000.
5. Expand NIC Teaming and verify that both vmnic0 and vmnic1 are set to active.


Configure port groups and VMkernel NICs

1. In the left navigation pane, select Networking.
2. Right-click the Port Groups tab.



3. Right-click VM Network and select Edit. Change the VLAN ID to <<var_vm_traffic_vlan>>.
4. Click Add Port Group.
 - Name the port group MGMT-Network.
 - Enter <<mgmt_vlan>> for the VLAN ID.
 - Make sure that vSwitch0 is selected.
 - Click Add.

5. Click the VMkernel NICs tab.



6. Select Add VMkernel NIC.
- Select New Port Group.
 - Name the port group NFS-Network.
 - Enter <<nfs_vlan_id>> for the VLAN ID.
 - Change the MTU to 9000.
 - Expand IPv4 Settings.
 - Select Static Configuration.
 - Enter <<var_hosta_nfs_ip>> for Address.
 - Enter <<var_hosta_nfs_mask>> for Subnet Mask.
 - Click Create.

| | |
|-----------------|--|
| Port group | New port group ▼ |
| New port group | NFS-Network |
| Virtual switch | vSwitch0 ▼ |
| VLAN ID | 3438 |
| MTU | 9000 |
| IP version | IPv4 only ▼ |
| ▼ IPv4 settings | |
| Configuration | <input type="radio"/> DHCP <input checked="" type="radio"/> Static |
| Address | 172.21.182.63 |
| Subnet mask | 255.255.255.0 |
| TCP/IP stack | Default TCP/IP stack ▼ |

Create Cancel

7. Repeat this process to create the vMotion VMkernel port.
8. Select Add VMkernel NIC.
 - a. Select New Port Group.
 - b. Name the port group vMotion.
 - c. Enter <<vmotion_vlan_id>> for the VLAN ID.
 - d. Change the MTU to 9000.
 - e. Expand IPv4 Settings.
 - f. Select Static Configuration.
 - g. Enter <<var_hosta_vmotion_ip>> for Address.
 - h. Enter <<var_hosta_vmotion_mask>> for Subnet Mask.
 - i. Make sure that the vMotion checkbox is selected after IPv4 Settings.

| | |
|-----------------|---|
| Virtual switch | vSwitch0 |
| VLAN ID | 3441 |
| MTU | 9000 |
| IP version | IPv4 only |
| ▼ IPv4 settings | |
| Configuration | <input type="radio"/> DHCP <input checked="" type="radio"/> Static |
| Address | 172.21.185.63 |
| Subnet mask | 255.255.255.0 |
| TCP/IP stack | Default TCP/IP stack |
| Services | <input checked="" type="checkbox"/> vMotion <input type="checkbox"/> Provisioning <input type="checkbox"/> Fault tolerance logging <input type="checkbox"/> Management <input type="checkbox"/> Replication <input type="checkbox"/> NFC replication |

Create Cancel



There are many ways to configure ESXi networking, including by using the VMware vSphere distributed switch if your licensing allows it. Alternative network configurations are supported in FlexPod Express if they are required to meet business requirements.


Mount first datastores

The first datastores to be mounted are the infra_datastore_1 datastore for virtual machines and the infra_swap datastore for virtual machine swap files.

1. Click Storage in the left navigation pane, and then click New Datastore.



2. Select Mount NFS Datastore.



3. Next, enter the following information in the Provide NFS Mount Details page:

- Name: `infra_datastore_1`
- NFS server: `<<var_nodea_nfs_lif>>`
- Share: `/infra_datastore_1`
- Make sure that NFS 3 is selected.

4. Click Finish. You can see the task completing in the Recent Tasks pane.

5. Repeat this process to mount the `infra_swap` datastore:

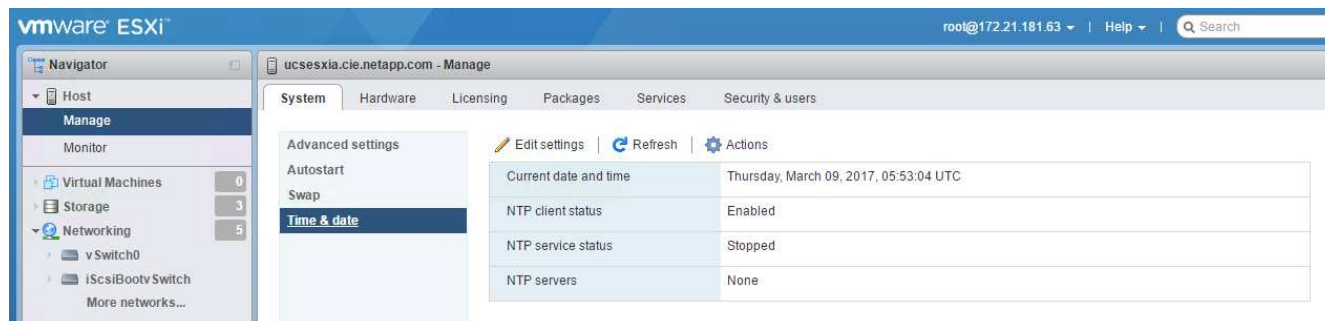
- Name: `infra_swap`
- NFS server: `<<var_nodea_nfs_lif>>`
- Share: `/infra_swap`

- Make sure that NFS 3 is selected.

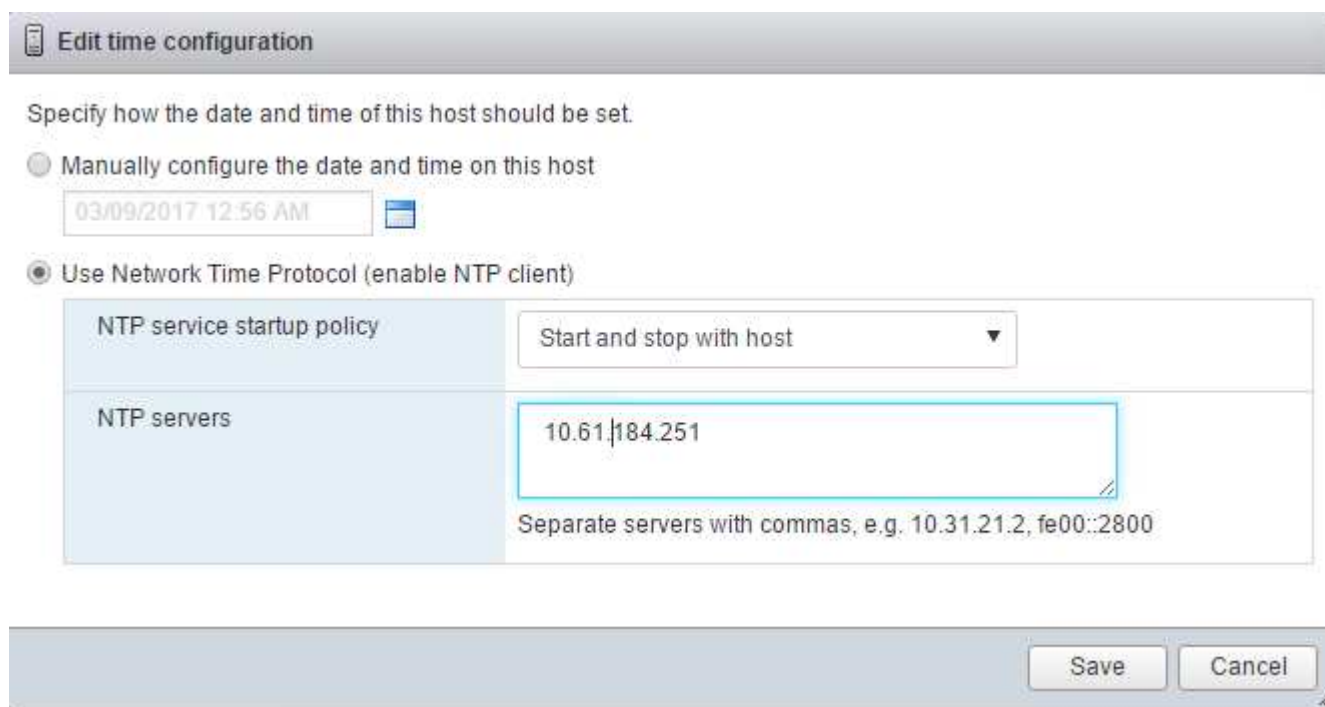
Configure NTP

To configure NTP for an ESXi host, complete the following steps:

1. Click Manage in the left navigation pane. Select System in the right pane and then click Time & Date.




2. Select Use Network Time Protocol (Enable NTP Client).
3. Select Start and Stop with Host as the NTP service startup policy.
4. Enter <<var_ntp>> as the NTP server. You can set multiple NTP servers.
5. Click Save.



Move the virtual machine swap-file location

These steps provide details for moving the virtual machine swap-file location.

1. Click Manage in the left navigation pane. Select system in the right pane, then click Swap.



2. Click Edit Settings. Select infra_swap from the Datastore options.



3. Click Save.

Install the NetApp NFS Plug-in 1.0.20 for VMware VAAI

To install the NetApp NFS Plug-in 1.0.20 for VMware VAAI, complete the following steps.

1. Enter the following commands to verify that VAAI is enabled:

```
esxcfg-advcfg -g /DataMover/HardwareAcceleratedMove
esxcfg-advcfg -g /DataMover/HardwareAcceleratedInit
```

If VAAI is enabled, these commands produce the following output:

```
~ # esxcfg-advcfg -g /DataMover/HardwareAcceleratedMove
Value of HardwareAcceleratedMove is 1
~ # esxcfg-advcfg -g /DataMover/HardwareAcceleratedInit
Value of HardwareAcceleratedInit is 1
```

2. If VAAI is not enabled, enter the following commands to enable VAAI:

```
esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedInit
esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedMove
```

These commands produce the following output:

```
~ # esxcfg-advcfg -s 1 /Data Mover/HardwareAcceleratedInit
Value of HardwareAcceleratedInit is 1
~ # esxcfg-advcfg -s 1 /DataMover/HardwareAcceleratedMove
Value of HardwareAcceleratedMove is 1
```

3. Download the NetApp NFS Plug-in for VMware VAAI:
 - a. Go to the [software download page](#).
 - b. Scroll down and click NetApp NFS Plug-in for VMware VAAI.
 - c. Select the ESXi platform.
 - d. Download either the offline bundle (.zip) or online bundle (.vib) of the most recent plug-in.
4. Install the plug-in on the ESXi host by using the ESX CLI.
5. Reboot the ESXi host.

```
[root@vm-host-infra-04:~] ls /vmfs/volumes/datastore1/NetAppNasPlugin.vib
/vmfs/volumes/datastore1/NetAppNasPlugin.vib
[root@vm-host-infra-04:~] esxcli software vib install -v /vmfs/volumes/datastore1/NetAppNasPlugin.vib
Installation Result
  Message: The update completed successfully, but the system needs to be rebooted for the changes to be effective.
  Reboot Required: true
  VIBs Installed: NetApp_bootbank_NetAppNasPlugin_1.1.2-3
  VIBs Removed:
  VIBs Skipped:
[root@vm-host-infra-04:~] █
```

[Next: Install VMware vCenter Server 6.7](#)

Install VMware vCenter Server 6.7


This section provides detailed procedures for installing VMware vCenter Server 6.7 in a FlexPod Express configuration.



FlexPod Express uses the VMware vCenter Server Appliance (VCSA).

Download the VMware vCenter server appliance

1. Download the VCSA. Access the download link by clicking the Get vCenter Server icon when managing the ESXi host.



2. Download the VCSA from the VMware site.



Although the Microsoft Windows vCenter Server installable is supported, VMware recommends the VCSA for new deployments.

3. Mount the ISO image.
4. Navigate to the vcsa-ui-installer> win32 directory. Double-click installer.exe.
5. Click Install.
6. Click Next on the Introduction page.
7. Accept the end-user license agreement.
8. Select Embedded Platform Services Controller as the deployment type.

Installer

vm Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Appliance deployment target

Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.

| | | |
|----------------------------------|---------------|---|
| ESXi host or vCenter Server name | 172.21.246.25 | i |
| HTTPS port | 443 | |
| User name | root | i |
| Password | ***** | |

CANCEL

BACK

NEXT

10. Set the appliance VM by entering `VCSA` as the VM name and the root password you would like to use for the VCSA.

12. Select the infra_datastore_1 datastore. Click Next.

vm

Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

1 Introduction

2 End user license agreement

3 Select deployment type

4 Appliance deployment target

5 Set up appliance VM

6 Select deployment size

7 Select datastore

8 Configure network settings

9 Ready to complete stage 1

Select datastore

Select the storage location for this appliance

☒ Install on an existing datastore accessible from the target host

| Name | Type | Capacity | Free | Provisioned | Thin Provisioning |
|-------------------|------|----------|-----------|-------------|-------------------|
| infra_datastore_1 | NFS | 500 GB | 499.98 GB | 18.38 MB | Supported |
| infra_swap | NFS | 100 GB | 99.99 GB | 10.95 MB | Supported |

2 items

☒ Enable Thin Disk Mode ⓘ

☐ Install on a new vSAN cluster containing the target host ⓘ

CANCEL

BACK

NEXT

13. Enter the following information in the Configure network settings page and click Next.

- Select MGMT-Network for Network.
- Enter the FQDN or IP to be used for the VCSA.
- Enter the IP address to be used.
- Enter the subnet mask to be used.
- Enter the default gateway.
- Enter the DNS server.

14. On the Ready to Complete Stage 1 page, verify that the settings you have entered are correct. Click Finish.

vCenter Server Appliance Installer
Installer

vm Install - Stage 1: Deploy vCenter Server with an Embedded Platform Services Controller

- 1 Introduction
- 2 End user license agreement
- 3 Select deployment type
- 4 Appliance deployment target
- 5 Set up appliance VM
- 6 Select deployment size
- 7 Select datastore
- 8 Configure network settings**
- 9 Ready to complete stage 1


Configure network settings

| | | |
|------------------------------|-----------------------------|---|
| IP version | IPv4 | |
| IP assignment | static | |
| FQDN | tigervcsa.cle.netapp.com | i |
| IP address | 172.21.246.41 | |
| Subnet mask or prefix length | 255.255.255.0 | i |
| Default gateway | 172.21.246.1 | |
| DNS servers | 10.61.184.251,10.61.184.252 | |
| Common Ports | | |
| HTTP | 80 | |
| HTTPS | 443 | |

CANCEL BACK NEXT

The VCSA installs now. This process takes several minutes.

- After stage 1 completes, a message appears stating that it has completed. Click Continue to begin stage 2 configuration.
- On the Stage 2 Introduction page, click Next.



17. Enter `<<var_ntp_id>>` for the NTP server address. You can enter multiple NTP IP addresses.

If you plan to use vCenter Server high availability (HA), make sure that SSH access is enabled.

18. Configure the SSO domain name, password, and site name. Click Next.

Record these values for your reference, especially if you deviate from the `vsphere.local` domain name.

19. Join the VMware Customer Experience Program if desired. Click Next.

20. View the summary of your settings. Click Finish or use the back button to edit settings.

21. A message appears stating that you will not be able to pause or stop the installation from completing after it has started. Click OK to continue.

The appliance setup continues. This takes several minutes.

A message appears indicating that the setup was successful.

The links that the installer provides to access vCenter Server are clickable.

[Next: Configure VMware vCenter Server 6.7 and vSphere clustering.](#)

Configure VMware vCenter Server 6.7 and vSphere clustering

To configure VMware vCenter Server 6.7 and vSphere clustering, complete the following steps:

1. Navigate to <https://<<FQDN or IP of vCenter>>/vsphere-client/>.
2. Click Launch vSphere Client.
3. Log in with the user name administrator@vsphere.local and the SSO password you entered during the VCSA setup process.
4. Right-click the vCenter name and select New Datacenter.
5. Enter a name for the data center and click OK.

Create vSphere cluster

Complete the following steps to create a vSphere cluster:

1. Right-click the newly created data center and select New Cluster.
2. Enter a name for the cluster.
3. Enable DR and vSphere HA by selecting the checkboxes.
4. Click OK.


The screenshot shows the 'New Cluster' dialog box in the vSphere Client. The dialog has a title bar with 'New Cluster' and 'FlexPod' on the left, and a close button (X) on the right. The main content area contains several fields and checkboxes:

- Name:** A text field containing 'Tiger3'.
- Location:** A dropdown menu showing 'FlexPod' with a building icon.
- DRS:** A checkbox labeled 'Turn ON' which is checked.
- vSphere HA:** A checkbox labeled 'Turn ON' which is checked.
- EVC:** A dropdown menu showing 'Disable'.

At the bottom right of the dialog are two buttons: 'CANCEL' and 'OK'.

Add ESXi hosts to cluster

1. Right-click the cluster and select Add Host.



2. To add an ESXi host to the cluster, complete the following steps:
 - a. Enter the IP or FQDN of the host. Click Next.
 - b. Enter the root user name and password. Click Next.
 - c. Click Yes to replace the host's certificate with a certificate signed by the VMware certificate server.
 - d. Click Next on the Host Summary page.
 - e. Click the green + icon to add a license to the vSphere host.



This step can be completed later if desired.

- f. Click Next to leave lockdown mode disabled.
 - g. Click Next at the VM location page.
 - h. Review the Ready to Complete page. Use the back button to make any changes or select Finish.
3. Repeat steps 1 and 2 for Cisco UCS host B. This process must be completed for any additional hosts added to the FlexPod Express configuration.

Configure coredump on ESXi hosts

1. Using SSH, connect to the management IP ESXi host, enter root for the user name, and enter the root password.
2. Run the following commands:

```
esxcli system coredump network set -i ip_address_of_core_dump_collector  
-v vmk0 -o 6500  
esxcli system coredump network set --enable=true  
esxcli system coredump network check
```

3. The message `Verified the configured netdump server is running` appears after you enter the final command.

This process must be completed for any additional hosts added to FlexPod Express.

Conclusion

FlexPod Express provides a simple and effective solution by providing a validated design that uses industry-leading components. By scaling through the addition of additional components, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed by keeping in mind small to midsize businesses, ROBOs, and other businesses that require dedicated solutions.

Where to find additional information

To learn more about the information described in this document, refer to the following documents and/or websites:

- NetApp product documentation

<http://docs.netapp.com>

- FlexPod Express with VMware vSphere 6.7 and NetApp AFF A220 Design Guide

<https://www.netapp.com/us/media/nva-1125-design.pdf>

FlexPod Express with VMware vSphere 6.7U1 and NetApp AFF A220 with Direct-Attached IP-Based Storage

NVA-1131-DEPLOY: FlexPod Express with VMware vSphere 6.7U1 and NetApp AFF A220 with Direct-Attached IP-Based Storage

Sree Lakshmi Lanka, NetApp

Industry trends indicate a vast data center transformation toward shared infrastructure and cloud computing. In addition, organizations seek a simple and effective solution for remote and branch offices, leveraging the technology with which they are familiar in their data center.

FlexPod Express is a predesigned, best practice architecture that is built on the Cisco Unified Computing System (Cisco UCS), the Cisco Nexus family of switches, and NetApp storage technologies. The components in a FlexPod Express system are like their FlexPod Datacenter counterparts, enabling management synergies across the complete IT infrastructure environment on a smaller scale. FlexPod Datacenter and FlexPod Express are optimal platforms for virtualization and for bare-metal OSs and enterprise workloads.

FlexPod Datacenter and FlexPod Express deliver a baseline configuration and have the versatility to be sized and optimized to accommodate many different use cases and requirements. Existing FlexPod Datacenter customers can manage their FlexPod Express system with the tools to which they are accustomed. New FlexPod Express customers can easily adapt to managing FlexPod Datacenter as their environment grows.

FlexPod Express is an optimal infrastructure foundation for remote offices and branch offices (ROBOs) and for small to midsize businesses. It is also an optimal solution for customers who want to provide infrastructure for a dedicated workload.

FlexPod Express provides an easy-to-manage infrastructure that is suitable for almost any workload.

Solution Overview

This FlexPod Express solution is part of the FlexPod converged infrastructure program.


FlexPod Converged Infrastructure Program

FlexPod reference architectures are delivered as Cisco Validated Designs (CVDs) or NetApp Verified Architectures (NVAs). Deviations based on customer requirements from a given CVD or NVA are permitted if these variations do not create an unsupported configuration.

As depicted in the figure below, the FlexPod program includes three solutions: FlexPod Express, FlexPod Datacenter, and FlexPod Select:

- **FlexPod Express** offers customers an entry-level solution with technologies from Cisco and NetApp.
- **FlexPod Datacenter** delivers an optimal multipurpose foundation for various workloads and applications.
- **FlexPod Select** incorporates the best aspects of FlexPod Datacenter and tailors the infrastructure to a given application.

The following figure shows the technical components of the solution.



NetApp Verified Architecture Program

The NVA program offers customers a verified architecture for NetApp solutions. An NVA provides a NetApp solution architecture with the following qualities:

- Is thoroughly tested
- Is prescriptive in nature
- Minimizes deployment risks
- Accelerates time to market

This guide details the design of FlexPod Express with direct- attached NetApp storage. The following sections list the components used for the design of this solution.

Hardware components

- NetApp AFF A220
- Cisco UCS Mini
- Cisco UCS B200 M5
- Cisco UCS VIC 1440/1480.
- Cisco Nexus 3000 Series Switches


Software components

- NetApp ONTAP 9. 5
- VMWare vSphere 6.7U1
- Cisco UCS Manager 4.0(1b)
- Cisco NXOS Firmware 7.0(3)I6(1)

Solution technology

This solution leverages the latest technologies from NetApp, Cisco, and VMware. It features the new NetApp AFF A220 running ONTAP 9.5, dual Cisco Nexus 31108PCV switches, and Cisco UCS B200 M5 servers that run VMware vSphere 6.7U1. This validated solution uses Direct Connect IP storage over 10GbE technology.

The following figure illustrates FlexPod Express with VMware vSphere 6.7U1 IP-Based Direct Connect architecture.



Use case summary

The FlexPod Express solution can be applied to several use cases, including the following:

- ROBOs
- Small and midsize businesses
- Environments that require a dedicated and cost-effective solution

FlexPod Express is best suited for virtualized and mixed workloads.

Technology requirements

A FlexPod Express system requires a combination of hardware and software

components. FlexPod Express also describes the hardware components that are required to add hypervisor nodes to the system in units of two.

Hardware requirements

Regardless of the hypervisor chosen, all FlexPod Express configurations use the same hardware. Therefore, even if business requirements change, either hypervisor can run on the same FlexPod Express hardware.

The following table lists the hardware components that are required for all FlexPod Express configurations.

| Hardware | Quantity |
|--|----------|
| AFF A220 HA Pair | 1 |
| Cisco UCS B200 M5 server | 2 |
| Cisco Nexus 31108PCV switch | 2 |
| Cisco UCS Virtual Interface Card (VIC) 1440 for the Cisco UCS B200 M5 server | 2 |
| Cisco UCS Mini with two Integrated UCS-FI-M-6324 fabric interconnects | 1 |

Software requirements

The following table lists the software components that are required to implement the architectures of the FlexPod Express solutions.

| Software | Version | Details |
|----------------------|-------------|---|
| Cisco UCS Manager | 4.0(1b) | For Cisco UCS Fabric Interconnect FI-6324UP |
| Cisco Blade software | 4.0(1b) | For Cisco UCS B200 M5 servers |
| Cisco nenic driver | 1.0.25.0 | For Cisco VIC 1440 interface cards |
| Cisco NX-OS | 7.0(3)I6(1) | For Cisco Nexus 31108PCV switches |
| NetApp ONTAP | 9.5 | For AFF A220 controllers |

The following table lists the software that is required for all VMware vSphere implementations on FlexPod Express.

| Software | Version |
|---------------------------------|---------|
| VMware vCenter Server Appliance | 6.7U1 |
| VMware vSphere ESXi hypervisor | 6.7U1 |

FlexPod Express Cabling Information

The reference validation cabling is documented in the following tables.

The following table lists cabling information for Cisco Nexus switch 31108PCV A.

| Local device | Local port | Remote device | Remote port |
|-------------------------------|------------|--------------------------------------|-------------|
| Cisco Nexus switch 31108PCV A | Eth1/1 | NetApp AFF A220 storage controller A | e0M |
| | Eth1/2 | Cisco UCS-mini FI-A | mgmt0 |
| | Eth1/3 | Cisco UCS-mini FI-A | Eth1/1 |
| | Eth 1/4 | Cisco UCS-mini FI-B | Eth1/1 |
| | Eth 1/13 | Cisco NX 31108PCV B | Eth 1/13 |
| | Eth 1/14 | Cisco NX 31108PCV B | Eth 1/14 |

The following table lists the cabling information for Cisco Nexus switch 31108PCV B.

| Local device | Local port | Remote device | Remote port |
|-------------------------------|------------|--------------------------------------|-------------|
| Cisco Nexus switch 31108PCV B | Eth1/1 | NetApp AFF A220 storage controller B | e0M |
| | Eth1/2 | Cisco UCS-mini FI-B | mgmt0 |
| | Eth1/3 | Cisco UCS-mini FI-A | Eth1/2 |
| | Eth 1/4 | Cisco UCS-mini FI-B | Eth1/2 |
| | Eth 1/13 | Cisco NX 31108PCV A | Eth 1/13 |
| | Eth 1/14 | Cisco NX 31108PCV A | Eth 1/14 |

The following table lists cabling information for NetApp AFF A220 storage controller A.

| Local device | Local port | Remote device | Remote port |
|--------------------------------------|------------|--------------------------------------|-------------|
| NetApp AFF A220 storage controller A | e0a | NetApp AFF A220 storage controller B | e0a |
| | e0b | NetApp AFF A220 storage controller B | e0b |
| | e0e | Cisco UCS-mini FI-A | Eth1/3 |
| | e0f | Cisco UCS-mini FI-B | Eth1/3 |
| | e0M | Cisco NX 31108PCV A | Eth1/1 |

The following table lists cabling information for NetApp AFF A220 storage controller B.

| Local device | Local port | Remote device | Remote port |
|--------------------------------------|------------|--------------------------------------|-------------|
| NetApp AFF A220 storage controller B | e0a | NetApp AFF A220 storage controller B | e0a |
| | e0b | NetApp AFF A220 storage controller B | e0b |
| | e0e | Cisco UCS-mini FI-A | Eth1/4 |
| | e0f | Cisco UCS-mini FI-B | Eth1/4 |
| | e0M | Cisco NX 31108PCV B | Eth1/1 |

The following table lists cabling information for Cisco UCS Fabric Interconnect A.

| Local device | Local port | Remote device | Remote port |
|---------------------------------|------------|--------------------------------------|-------------|
| Cisco UCS Fabric Interconnect A | Eth1/1 | Cisco NX 31108PCV A | Eth1/3 |
| | Eth1/2 | Cisco NX 31108PCV B | Eth1/3 |
| | Eth1/3 | NetApp AFF A220 storage controller A | e0e |
| | Eth1/4 | NetApp AFF A220 storage controller B | e0e |
| | mgmt0 | Cisco NX 31108PCV A | Eth1/2 |

The following table lists cabling information for Cisco UCS Fabric Interconnect B.

| Local device | Local port | Remote device | Remote port |
|---------------------------------|------------|--------------------------------------|-------------|
| Cisco UCS Fabric Interconnect B | Eth1/1 | Cisco NX 31108PCV A | Eth1/4 |
| | Eth1/2 | Cisco NX 31108PCV B | Eth1/4 |
| | Eth1/3 | NetApp AFF A220 storage controller A | e0f |
| | Eth1/4 | NetApp AFF A220 storage controller B | e0f |
| | mgmt0 | Cisco NX 31108PCV B | Eth1/2 |

Deployment Procedures

This document provides details for configuring a fully redundant, highly available FlexPod Express system. To reflect this redundancy, the components being configured in each step are referred to as either component A or component B. For example, controller A and controller B identify the two NetApp storage controllers that are provisioned in this document. Switch A and switch B identify a pair of Cisco Nexus switches. Fabric Interconnect A and Fabric Interconnect B are the two Integrated Nexus Fabric Interconnects.


In addition, this document describes steps for provisioning multiple Cisco UCS hosts, which are identified

sequentially as server A, server B, and so on.

To indicate that you should include information pertinent to your environment in a step, <<text>> appears as part of the command structure. See the following example for the `vlan create` command:

```
Controller01>vlan create vif0 <<mgmt_vlan_id>>
```

This document enables you to fully configure the FlexPod Express environment. In this process, various steps require you to insert customer-specific naming conventions, IP addresses, and virtual local area network (VLAN) schemes. The table below describes the VLANs required for deployment, as outlined in this guide. This table can be completed based on the specific site variables and used to implement the document configuration steps.



If you use separate in-band and out-of-band management VLANs, you must create a layer 3 route between them. For this validation, a common management VLAN was used.

| VLAN name | VLAN purpose | ID used in validating this document |
|---------------------|--|-------------------------------------|
| Management VLAN | VLAN for management interfaces | 18 |
| Native VLAN | VLAN to which untagged frames are assigned | 2 |
| NFS VLAN | VLAN for NFS traffic | 104 |
| VMware vMotion VLAN | VLAN designated for the movement of virtual machines (VMs) from one physical host to another | 103 |
| VM traffic VLAN | VLAN for VM application traffic | 102 |
| iSCSI-A-VLAN | VLAN for iSCSI traffic on fabric A | 124 |
| iSCSI-B-VLAN | VLAN for iSCSI traffic on fabric B | 125 |

The VLAN numbers are needed throughout the configuration of FlexPod Express. The VLANs are referred to as <<var_xxxx_vlan>>, where `xxxx` is the purpose of the VLAN (such as `iSCSI-A`).

The following table lists the VMware VMs created.

| VM Description | Host Name |
|-----------------------|------------------------------|
| VMware vCenter Server | Seahawks-vcsa.cie.netapp.com |

Cisco Nexus 31108PCV deployment procedure

This section details the Cisco Nexus 31308PCV switch configuration used in a FlexPod Express environment.

Initial setup of Cisco Nexus 31108PCV switch

This procedures describes how to configure the Cisco Nexus switches for use in a base FlexPod Express environment.



This procedure assumes that you are using a Cisco Nexus 31108PCV running NX-OS software release 7.0(3)I6(1).

1. Upon initial boot and connection to the console port of the switch, the Cisco NX-OS setup automatically starts. This initial configuration addresses basic settings, such as the switch name, the mgmt0 interface configuration, and Secure Shell (SSH) setup.
2. The FlexPod Express management network can be configured in multiple ways. The mgmt0 interfaces on the 31108PCV switches can be connected to an existing management network, or the mgmt0 interfaces of the 31108PCV switches can be connected in a back-to-back configuration. However, this link cannot be used for external management access such as SSH traffic.

In this deployment guide, the FlexPod Express Cisco Nexus 31108PCV switches are connected to an existing management network.

3. To configure the Cisco Nexus 31108PCV switches, power on the switch and follow the on-screen prompts, as illustrated here for the initial setup of both the switches, substituting the appropriate values for the switch-specific information.

This setup utility will guide you through the basic configuration of the system. Setup configures only enough connectivity for management of the system.

```

*Note: setup is mainly used for configuring the system initially, when
no configuration is present. So setup always assumes system defaults and
not the current system configuration values.
Press Enter at anytime to skip a dialog. Use ctrl-c at anytime to skip
the remaining dialogs.
Would you like to enter the basic configuration dialog (yes/no): y
Do you want to enforce secure password standard (yes/no) [y]: y
Create another login account (yes/no) [n]: n
Configure read-only SNMP community string (yes/no) [n]: n
Configure read-write SNMP community string (yes/no) [n]: n
Enter the switch name : 31108PCV-A
Continue with Out-of-band (mgmt0) management configuration? (yes/no)
[y]: y
Mgmt0 IPv4 address : <<var_switch_mgmt_ip>>
Mgmt0 IPv4 netmask : <<var_switch_mgmt_netmask>>
Configure the default gateway? (yes/no) [y]: y
IPv4 address of the default gateway : <<var_switch_mgmt_gateway>>
Configure advanced IP options? (yes/no) [n]: n
Enable the telnet service? (yes/no) [n]: n
Enable the ssh service? (yes/no) [y]: y
Type of ssh key you would like to generate (dsa/rsa) [rsa]: rsa
Number of rsa key bits <1024-2048> [1024]: <enter>
Configure the ntp server? (yes/no) [n]: y
NTP server IPv4 address : <<var_ntp_ip>>
Configure default interface layer (L3/L2) [L2]: <enter>
Configure default switchport interface state (shut/noshut) [noshut]:
<enter>
Configure CoPP system profile (strict/moderate/lenient/dense) [strict]:
<enter>

```

4. A summary of your configuration is displayed and you are asked if you would like to edit the configuration. If your configuration is correct, enter n.

```

Would you like to edit the configuration? (yes/no) [n]: no

```

5. You are then asked if you would like to use this configuration and save it. If so, enter y.

```

Use this configuration and save it? (yes/no) [y]: Enter

```

6. Repeat steps 1 through 5 for Cisco Nexus switch B.

Enable advanced features

Certain advanced features must be enabled in Cisco NX-OS to provide additional configuration options.

1. To enable the appropriate features on Cisco Nexus switch A and switch B, enter configuration mode by using the command (`config t`) and run the following commands:

```
feature interface-vlan
feature lacp
feature vpc
```



The default port channel load-balancing hash uses the source and destination IP addresses to determine the load-balancing algorithm across the interfaces in the port channel. You can achieve better distribution across the members of the port channel by providing more inputs to the hash algorithm beyond the source and destination IP addresses. For the same reason, NetApp highly recommends adding the source and destination TCP ports to the hash algorithm.

2. From configuration mode (`config t`), run the following commands to set the global port channel load-balancing configuration on Cisco Nexus switch A and switch B:

```
port-channel load-balance src-dst ip-l4port
```

Perform global spanning-tree configuration

The Cisco Nexus platform uses a new protection feature called bridge assurance. Bridge assurance helps protect against a unidirectional link or other software failure with a device that continues to forward data traffic when it is no longer running the spanning-tree algorithm. Ports can be placed in one of several states, including network or edge, depending on the platform.

NetApp recommends setting bridge assurance so that all ports are considered to be network ports by default. This setting forces the network administrator to review the configuration of each port. It also reveals the most common configuration errors, such as unidentified edge ports or a neighbor that does not have the bridge assurance feature enabled. In addition, it is safer to have the spanning tree block many ports rather than too few, which allows the default port state to enhance the overall stability of the network.

Pay close attention to the spanning-tree state when adding servers, storage, and uplink switches, especially if they do not support bridge assurance. In such cases, you might need to change the port type to make the ports active.

The Bridge Protocol Data Unit (BPDU) guard is enabled on edge ports by default as another layer of protection. To prevent loops in the network, this feature shuts down the port if BPDUs from another switch are seen on this interface.

From configuration mode (`config t`), run the following commands to configure the default spanning-tree options, including the default port type and BPDU guard, on Cisco Nexus switch A and switch B:

```
spanning-tree port type network default
spanning-tree port type edge bpduguard default
```

Define VLANs

Before individual ports with different VLANs are configured, the layer-2 VLANs must be defined on the switch. It is also a good practice to name the VLANs for easy troubleshooting in the future.

From configuration mode (`config t`), run the following commands to define and describe the layer 2 VLANs on Cisco Nexus switch A and switch B:

```
vlan <<nfs_vlan_id>>
  name NFS-VLAN
vlan <<iSCSI_A_vlan_id>>
  name iSCSI-A-VLAN
vlan <<iSCSI_B_vlan_id>>
  name iSCSI-B-VLAN
vlan <<vmotion_vlan_id>>
  name vMotion-VLAN
vlan <<vmtraffic_vlan_id>>
  name VM-Traffic-VLAN
vlan <<mgmt_vlan_id>>
  name MGMT-VLAN
vlan <<native_vlan_id>>
  name NATIVE-VLAN
exit
```

Configure access and management port descriptions

As is the case with assigning names to the layer-2 VLANs, setting descriptions for all the interfaces can help with both provisioning and troubleshooting.

From configuration mode (`config t`) in each of the switches, enter the following port descriptions for the FlexPod Express large configuration:

Cisco Nexus switch A

```
int eth1/1
  description AFF A220-A e0M
int eth1/2
  description Cisco UCS FI-A mgmt0
int eth1/3
  description Cisco UCS FI-A eth1/1
int eth1/4
  description Cisco UCS FI-B eth1/1
int eth1/13
  description vPC peer-link 31108PVC-B 1/13
int eth1/14
  description vPC peer-link 31108PVC-B 1/14
```

Cisco Nexus switch B

```
int eth1/1
    description AFF A220-B e0M
int eth1/2
    description Cisco UCS FI-B mgmt0
int eth1/3
    description Cisco UCS FI-A eth1/2
int eth1/4
    description Cisco UCS FI-B eth1/2
int eth1/13
    description vPC peer-link 31108PVC-B 1/13
int eth1/14
    description vPC peer-link 31108PVC-B 1/14
```

Configure server and storage management interfaces

The management interfaces for both the server and the storage typically use only a single VLAN. Therefore, configure the management interface ports as access ports. Define the management VLAN for each switch and change the spanning-tree port type to edge.

From configuration mode (`config t`), run the following commands to configure the port settings for the management interfaces of both the servers and the storage:

Cisco Nexus switch A

```
int eth1/1-2
    switchport mode access
    switchport access vlan <<mgmt_vlan>>
    spanning-tree port type edge
    speed 1000
exit
```

Cisco Nexus switch B

```
int eth1/1-2
    switchport mode access
    switchport access vlan <<mgmt_vlan>>
    spanning-tree port type edge
    speed 1000
exit
```

Add NTP distribution interface

Cisco Nexus switch A

From the global configuration mode, execute the following commands.

```
interface Vlan<ib-mgmt-vlan-id>  
ip address <switch-a-ntp-ip>/<ib-mgmt-vlan-netmask-length>  
no shutdown  
exitntp peer <switch-b-ntp-ip> use-vrf default
```

Cisco Nexus switch B

From the global configuration mode, execute the following commands.

```
interface Vlan<ib-mgmt-vlan-id>  
ip address <switch- b-ntp-ip>/<ib-mgmt-vlan-netmask-length>  
no shutdown  
exitntp peer <switch-a-ntp-ip> use-vrf default
```

Perform virtual port channel global configuration

A virtual port channel (vPC) enables links that are physically connected to two different Cisco Nexus switches to appear as a single port channel to a third device. The third device can be a switch, server, or any other networking device. A vPC can provide layer-2 multipathing, which allows you to create redundancy by increasing bandwidth, enabling multiple parallel paths between nodes, and load-balancing traffic where alternative paths exist.

A vPC provides the following benefits:

- Enabling a single device to use a port channel across two upstream devices
- Eliminating spanning-tree protocol blocked ports
- Providing a loop-free topology
- Using all available uplink bandwidth
- Providing fast convergence if either the link or a device fails
- Providing link-level resiliency
- Helping provide high availability

The vPC feature requires some initial setup between the two Cisco Nexus switches to function properly. If you use the back-to-back mgmt0 configuration, use the addresses defined on the interfaces and verify that they can communicate by using the ping <<switch_A/B_mgmt0_ip_addr>>vrf management command.

From configuration mode (`config t`), run the following commands to configure the vPC global configuration for both switches:

Cisco Nexus switch A

```

vpc domain 1
  role priority 10
peer-keepalive destination <<switch_B_mgmt0_ip_addr>> source
<<switch_A_mgmt0_ip_addr>> vrf management
  peer-gateway
  auto-recovery
  ip arp synchronize
  int eth1/13-14
  channel-group 10 mode active
int Po10description vPC peer-link
switchport
switchport mode trunkswitchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>,
<<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>> spanning-tree port type network
vpc peer-link
no shut
exit
int Po13
description vPC ucs-FI-A
switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan <<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
<<mgmt_vlan>> spanning-tree port type network
mtu 9216
vpc 13
no shut
exit
int eth1/3
  channel-group 13 mode active
int Po14
description vPC ucs-FI-B
switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan <<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
<<mgmt_vlan>> spanning-tree port type network
mtu 9216
vpc 14
no shut
exit
int eth1/4
  channel-group 14 mode active
copy run start

```

Cisco Nexus switch B

```
vpc domain 1
peer-switch
role priority 20
peer-keepalive destination <<switch_A_mgmt0_ip_addr>> source
<<switch_B_mgmt0_ip_addr>> vrf management
    peer-gateway
    auto-recovery
    ip arp synchronize
    int eth1/13-14
    channel-group 10 mode active
int Po10
description vPC peer-link
switchport
switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan <<nfs_vlan_id>>,<<vmotion_vlan_id>>,
<<vmtraffic_vlan_id>>, <<mgmt_vlan>>, <<iSCSI_A_vlan_id>>,
<<iSCSI_B_vlan_id>> spanning-tree port type network
vpc peer-link
no shut
exit
int Po13
description vPC ucs-FI-A
switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan <<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
<<mgmt_vlan>> spanning-tree port type network
mtu 9216
vpc 13
no shut
exit
int eth1/3
    channel-group 13 mode active
int Po14
description vPC ucs-FI-B
switchport mode trunk
switchport trunk native vlan <<native_vlan_id>>
switchport trunk allowed vlan <<vmotion_vlan_id>>, <<vmtraffic_vlan_id>>,
<<mgmt_vlan>> spanning-tree port type network
mtu 9216
vpc 14
no shut
exit
int eth1/4
```



```
channel-group 14 mode active
copy run start
```



In this solution validation, a maximum transmission unit (MTU) of 9000 was used. However, based on application requirements, you can configure an appropriate value of MTU. It is important to set the same MTU value across the FlexPod solution. Incorrect MTU configurations between components result in packets being dropped.

Uplink into existing network infrastructure

Depending on the available network infrastructure, several methods and features can be used to uplink the FlexPod environment. If an existing Cisco Nexus environment is present, NetApp recommends using vPCs to uplink the Cisco Nexus 31108PVC switches included in the FlexPod environment into the infrastructure. The uplinks can be 10GbE uplinks for a 10GbE infrastructure solution or 1GbE for a 1GbE infrastructure solution if required. The previously described procedures can be used to create an uplink vPC to the existing environment. Make sure to run `copy run start` to save the configuration on each switch after the configuration is completed.

NetApp storage deployment procedure (part 1)

This section describes the NetApp AFF storage deployment procedure.

NetApp Storage Controller AFF2xx Series Installation

NetApp Hardware Universe

The [NetApp Hardware Universe](#) (HWU) application provides supported hardware and software components for any specific ONTAP version. It provides configuration information for all the NetApp storage appliances currently supported by ONTAP software. It also provides a table of component compatibilities.

Confirm that the hardware and software components that you would like to use are supported with the version of ONTAP that you plan to install:

1. Access the [HWU](#) application to view the system configuration guides. Select the Compare Storage Systems tab to view the compatibility between different version of the ONTAP software and the NetApp storage appliances with your desired specifications.
2. Alternatively, to compare components by storage appliance, click Compare Storage Systems.

Controller AFF2XX Series prerequisites

To plan the physical location of the storage systems, see the the following sections:

Electrical requirements
Supported power cords
Onboard ports and cables

Storage controllers

Follow the physical installation procedures for the controllers in the [AFF A220 Documentation](#).

NetApp ONTAP 9.5

Configuration worksheet

Before running the setup script, complete the configuration worksheet from the product manual. The configuration worksheet is available in the [ONTAP 9.5 Software Setup Guide](#) (available in the [ONTAP 9 Documentation Center](#)). The table below illustrates ONTAP 9.5 installation and configuration information.



This system is set up in a two-node switchless cluster configuration.

| Cluster Detail | Cluster Detail Value |
|---|-----------------------------|
| Cluster node A IP address | <<var_nodeA_mgmt_ip>> |
| Cluster node A netmask | <<var_nodeA_mgmt_mask>> |
| Cluster node A gateway | <<var_nodeA_mgmt_gateway>> |
| Cluster node A name | <<var_nodeA>> |
| Cluster node B IP address | <<var_nodeB_mgmt_ip>> |
| Cluster node B netmask | <<var_nodeB_mgmt_mask>> |
| Cluster node B gateway | <<var_nodeB_mgmt_gateway>> |
| Cluster node B name | <<var_nodeB>> |
| ONTAP 9.5 URL | <<var_url_boot_software>> |
| Name for cluster | <<var_clustername>> |
| Cluster management IP address | <<var_clustermgmt_ip>> |
| Cluster B gateway | <<var_clustermgmt_gateway>> |
| Cluster B netmask | <<var_clustermgmt_mask>> |
| Domain name | <<var_domain_name>> |
| DNS server IP (you can enter more than one) | <<var_dns_server_ip>> |
| NTP server A IP | << switch-a-ntp-ip >> |
| NTP server B IP | << switch-b-ntp-ip >> |

Configure node A

To configure node A, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl- C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Allow the system to boot.

```
autoboot
```

3. Press Ctrl- C to enter the Boot menu.

If ONTAP 9. 5 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9. 5 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeA_mgmt_ip>> <<var_nodeA_mgmt_mask>> <<var_nodeA_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

10. Press Enter for the user name, indicating no user name.
11. Enter y to set the newly installed software as the default to be used for subsequent reboots.
12. Enter y to reboot the node.

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl- C to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter y to zero disks, reset config, and install a new file system.
16. Enter y to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize. You can continue with the node B configuration while the disks for node A are zeroing.

17. While node A is initializing, begin configuring node B.

Configure node B

To configure node B, complete the following steps:

1. Connect to the storage system console port. You should see a Loader-A prompt. However, if the storage system is in a reboot loop, press Ctrl-C to exit the autoboot loop when you see this message:

```
Starting AUTOBOOT press Ctrl-C to abort...
```

2. Press Ctrl-C to enter the Boot menu.

```
autoboot
```

3. Press Ctrl-C when prompted.

If ONTAP 9. 5 is not the version of software being booted, continue with the following steps to install new software. If ONTAP 9.4 is the version being booted, select option 8 and y to reboot the node. Then, continue with step 14.

4. To install new software, select option 7.
5. Enter y to perform an upgrade.
6. Select e0M for the network port you want to use for the download.
7. Enter y to reboot now.
8. Enter the IP address, netmask, and default gateway for e0M in their respective places.

```
<<var_nodeB_mgmt_ip>> <<var_nodeB_mgmt_ip>><<var_nodeB_mgmt_gateway>>
```

9. Enter the URL where the software can be found.



This web server must be pingable.

```
<<var_url_boot_software>>
```

10. Press Enter for the user name, indicating no user name
11. Enter y to set the newly installed software as the default to be used for subsequent reboots.
12. Enter y to reboot the node.

When installing new software, the system might perform firmware upgrades to the BIOS and adapter cards, causing reboots and possible stops at the Loader-A prompt. If these actions occur, the system might deviate from this procedure.

13. Press Ctrl-C to enter the Boot menu.
14. Select option 4 for Clean Configuration and Initialize All Disks.
15. Enter y to zero disks, reset config, and install a new file system.
16. Enter y to erase all the data on the disks.

The initialization and creation of the root aggregate can take 90 minutes or more to complete, depending on the number and type of disks attached. When initialization is complete, the storage system reboots. Note that SSDs take considerably less time to initialize.

Continuation node A configuration and cluster configuration

From a console port program attached to the storage controller A (node A) console port, run the node setup script. This script appears when ONTAP 9.5 boots on the node for the first time.

The node and cluster setup procedure has changed slightly in ONTAP 9.5. The cluster setup wizard is now used to configure the first node in a cluster, and System Manager is used to configure the cluster.

1. Follow the prompts to set up node A.

```
Welcome to the cluster setup wizard.
You can enter the following commands at any time:
    "help" or "?" - if you want to have a question clarified,
    "back" - if you want to change previously answered questions, and
    "exit" or "quit" - if you want to quit the cluster setup wizard.
    Any changes you made before quitting will be saved.
You can return to cluster setup at any time by typing "cluster setup".
To accept a default or omit a question, do not enter a value.
This system will send event messages and periodic reports to NetApp
Technical Support. To disable this feature, enter
autosupport modify -support disable
within 24 hours.
Enabling AutoSupport can significantly speed problem determination and
resolution should a problem occur on your system.
For further information on AutoSupport, see:
http://support.netapp.com/autosupport/
Type yes to confirm and continue {yes}: yes
Enter the node management interface port [e0M]:
Enter the node management interface IP address: <<var_nodeA_mgmt_ip>>
Enter the node management interface netmask: <<var_nodeA_mgmt_mask>>
Enter the node management interface default gateway:
<<var_nodeA_mgmt_gateway>>
A node management interface on port e0M with IP address
<<var_nodeA_mgmt_ip>> has been created.
Use your web browser to complete cluster setup by accessing
https://<<var_nodeA_mgmt_ip>>
Otherwise, press Enter to complete cluster setup using the command line
interface:
```

2. Navigate to the IP address of the node's management interface.



Cluster setup can also be performed by using the CLI. This document describes cluster setup using NetApp System Manager guided setup.

3. Click Guided Setup to configure the cluster.
4. Enter <<var_clustername>> for the cluster name and <<var_nodeA>> and <<var_nodeB>> for each of the nodes that you are configuring. Enter the password that you would like to use for the storage system.

Select Switchless Cluster for the cluster type. Enter the cluster base license.

5. You can also enter feature licenses for Cluster, NFS, and iSCSI.
6. You see a status message stating the cluster is being created. This status message cycles through several statuses. This process takes several minutes.
7. Configure the network.
 - a. Deselect the IP Address Range option.
 - b. Enter <<var_clustermgmt_ip>> in the Cluster Management IP Address field, <<var_clustermgmt_mask>> in the Netmask field, and <<var_clustermgmt_gateway>> in the Gateway field. Use the ... selector in the Port field to select e0M of node A.
 - c. The node management IP for node A is already populated. Enter <<var_nodeA_mgmt_ip>> for node B.
 - d. Enter <<var_domain_name>> in the DNS Domain Name field. Enter <<var_dns_server_ip>> in the DNS Server IP Address field.

You can enter multiple DNS server IP addresses.

- e. Enter <<switch-a-ntp-ip>> in the Primary NTP Server field.

You can also enter an alternate NTP server as <<switch- b-ntp-ip>>.

8. Configure the support information.
 - a. If your environment requires a proxy to access AutoSupport, enter the URL in Proxy URL.
 - b. Enter the SMTP mail host and email address for event notifications.

You must, at a minimum, set up the event notification method before you can proceed. You can select any of the methods.

9. When indicated that the cluster configuration has completed, click Manage Your Cluster to configure the storage.

Continuation of storage cluster configuration

After the configuration of the storage nodes and base cluster, you can continue with the configuration of the storage cluster.

Zero all spare disks

To zero all spare disks in the cluster, run the following command:

```
disk zerospares
```

Set on-board UTA2 ports personality

1. Verify the current mode and the current type of the ports by running the `ucadmin show` command.

```
AFFA220-Clus::> ucadmin show
```

| Node | Adapter | Current Mode | Current Type | Pending Mode | Pending Type | Admin Status |
|-----------------|---------|--------------|--------------|--------------|--------------|--------------|
| AFFA220-Clus-01 | 0c | cna | target | - | - | offline |
| AFFA220-Clus-01 | 0d | cna | target | - | - | offline |
| AFFA220-Clus-01 | 0e | cna | target | - | - | offline |
| AFFA220-Clus-01 | 0f | cna | target | - | - | offline |
| AFFA220-Clus-02 | 0c | cna | target | - | - | offline |
| AFFA220-Clus-02 | 0d | cna | target | - | - | offline |
| AFFA220-Clus-02 | 0e | cna | target | - | - | offline |
| AFFA220-Clus-02 | 0f | cna | target | - | - | offline |

```
8 entries were displayed.
```

2. Verify that the current mode of the ports that are in use is `cna` and that the current type is set to `target`. If not, change the port personality by running the following command:

```
ucadmin modify -node <home node of the port> -adapter <port name> -mode  
cna -type target
```

The ports must be offline to run the previous command. To take a port offline, run the following command:

```
network fcp adapter modify -node <home node of the port> -adapter <port  
name> -state down
```



If you changed the port personality, you must reboot each node for the change to take effect.

Enable Cisco Discovery Protocol

To enable the Cisco Discovery Protocol (CDP) on the NetApp storage controllers, run the following command:

```
node run -node * options cdpd.enable on
```

Enable Link-layer Discovery Protocol on all Ethernet ports

Enable the exchange of Link-layer Discovery Protocol (LLDP) neighbor information between the storage and network switches by running the following command. This command enables LLDP on all ports of all nodes in the cluster.

```
node run * options lldp.enable on
```

Rename management logical interfaces

To rename the management logical interfaces (LIFs), complete the following steps:

1. Show the current management LIF names.

```
network interface show -vserver <<clustername>>
```

2. Rename the cluster management LIF.

```
network interface rename -vserver <<clustername>> -lif  
cluster_setup_cluster_mgmt_lif_1 -newname cluster_mgmt
```

3. Rename the node B management LIF.

```
network interface rename -vserver <<clustername>> -lif  
cluster_setup_node_mgmt_lif_AFF A220_A_1 - newname AFF A220-01_mgmt1
```

Set auto-revert on cluster management

Set the `auto-revert` parameter on the cluster management interface.

```
network interface modify -vserver <<clustername>> -lif cluster_mgmt -auto-  
revert true
```

Set up service processor network interface

To assign a static IPv4 address to the service processor on each node, run the following commands:


```
system service-processor network modify -node <<var_nodeA>> -address  
-family IPv4 -enable true - dhcp none -ip-address <<var_nodeA_sp_ip>>  
-netmask <<var_nodeA_sp_mask>> -gateway <<var_nodeA_sp_gateway>>  
system service-processor network modify -node <<var_nodeB>> -address  
-family IPv4 -enable true - dhcp none -ip-address <<var_nodeB_sp_ip>>  
-netmask <<var_nodeB_sp_mask>> -gateway <<var_nodeB_sp_gateway>>
```



The service processor IP addresses should be in the same subnet as the node management IP addresses.

Enable storage failover in ONTAP

To confirm that storage failover is enabled, run the following commands in a failover pair:

1. Verify the status of storage failover.

```
storage failover show
```

Both <<var_nodeA>> and <<var_nodeB>> must be able to perform a takeover. Go to step 3 if the nodes can perform a takeover.

2. Enable failover on one of the two nodes.

```
storage failover modify -node <<var_nodeA>> -enabled true
```

3. Verify the HA status of the two-node cluster.



This step is not applicable for clusters with more than two nodes.

```
cluster ha show
```

4. Go to step 6 if high availability is configured. If high availability is configured, you see the following message upon issuing the command:

```
High Availability Configured: true
```

5. Enable HA mode only for the two-node cluster.

Do not run this command for clusters with more than two nodes because it causes problems with failover.

```
cluster ha modify -configured true  
Do you want to continue? {y|n}: y
```

6. Verify that hardware assist is correctly configured and, if needed, modify the partner IP address.

```
storage failover hwassist show
```

The message `Keep Alive Status : Error: did not receive hwassist keep alive alerts from partner` indicates that hardware assist is not configured. Run the following commands to configure hardware assist.

```
storage failover modify -hwassist-partner-ip <<var_nodeB_mgmt_ip>> -node <<var_nodeA>>
storage failover modify -hwassist-partner-ip <<var_nodeA_mgmt_ip>> -node <<var_nodeB>>
```

Create jumbo frame MTU broadcast domain in ONTAP

To create a data broadcast domain with an MTU of 9000, run the following commands:

```
broadcast-domain create -broadcast-domain Infra_NFS -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-A -mtu 9000
broadcast-domain create -broadcast-domain Infra_iSCSI-B -mtu 9000
```

Remove data ports from default broadcast domain

The 10GbE data ports are used for iSCSI/NFS traffic, and these ports should be removed from the default domain. Ports e0e and e0f are not used and should also be removed from the default domain.

To remove the ports from the broadcast domain, run the following command:

```
broadcast-domain remove-ports -broadcast-domain Default -ports
<<var_nodeA>>:e0c, <<var_nodeA>>:e0d, <<var_nodeA>>:e0e,
<<var_nodeA>>:e0f, <<var_nodeB>>:e0c, <<var_nodeB>>:e0d,
<<var_nodeA>>:e0e, <<var_nodeA>>:e0f
```

Disable flow control on UTA2 ports

It is a NetApp best practice to disable flow control on all UTA2 ports that are connected to external devices. To disable flow control, run the following commands:

```

net port modify -node <<var_nodeA>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeA>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0c -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0d -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0e -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y
net port modify -node <<var_nodeB>> -port e0f -flowcontrol-admin none
Warning: Changing the network port settings will cause a several second
interruption in carrier. Do you want to continue? {y|n}: y

```



The Cisco UCS Mini direct connection to ONTAP does not support LACP.

Configure jumbo frames in NetApp ONTAP

To configure an ONTAP network port to use jumbo frames (that usually have an MTU of 9,000 bytes), run the following commands from the cluster shell:

```

AFF A220::> network port modify -node node_A -port e0e -mtu 9000
Warning: This command will cause a several second interruption of service
on this network port.
Do you want to continue? {y|n}: y
AFF A220::> network port modify -node node_B -port e0e -mtu 9000
Warning: This command will cause a several second interruption of service
on this network port.
Do you want to continue? {y|n}: y
AFF A220::> network port modify -node node_A -port e0f -mtu 9000
Warning: This command will cause a several second interruption of service
on this network port.
Do you want to continue? {y|n}: y
AFF A220::> network port modify -node node_B -port e0f -mtu 9000
Warning: This command will cause a several second interruption of service
on this network port.
Do you want to continue? {y|n}: y

```

Create VLANs in ONTAP

To create VLANs in ONTAP, complete the following steps:

1. Create NFS VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name e0e-
<<var_nfs_vlan_id>>
network port vlan create -node <<var_nodeA>> -vlan-name e0f-
<<var_nfs_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name e0e-
<<var_nfs_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name e0f-
<<var_nfs_vlan_id>>
broadcast-domain add-ports -broadcast-domain Infra_NFS -ports
<<var_nodeA>>: e0e- <<var_nfs_vlan_id>>, <<var_nodeB>>: e0e-
<<var_nfs_vlan_id>> , <<var_nodeA>>:e0f- <<var_nfs_vlan_id>>,
<<var_nodeB>>:e0f-<<var_nfs_vlan_id>>

```

2. Create iSCSI VLAN ports and add them to the data broadcast domain.

```

network port vlan create -node <<var_nodeA>> -vlan-name e0e-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeA>> -vlan-name e0f-
<<var_iscsi_vlan_B_id>>
network port vlan create -node <<var_nodeB>> -vlan-name e0e-
<<var_iscsi_vlan_A_id>>
network port vlan create -node <<var_nodeB>> -vlan-name e0f-
<<var_iscsi_vlan_B_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-A -ports
<<var_nodeA>>: e0e- <<var_iscsi_vlan_A_id>>,<<var_nodeB>>: e0e-
<<var_iscsi_vlan_A_id>>
broadcast-domain add-ports -broadcast-domain Infra_iSCSI-B -ports
<<var_nodeA>>: e0f- <<var_iscsi_vlan_B_id>>,<<var_nodeB>>: e0f-
<<var_iscsi_vlan_B_id>>

```

3. Create MGMT-VLAN ports.

```

network port vlan create -node <<var_nodeA>> -vlan-name e0m-
<<mgmt_vlan_id>>
network port vlan create -node <<var_nodeB>> -vlan-name e0m-
<<mgmt_vlan_id>>

```

Create aggregates in ONTAP

An aggregate containing the root volume is created during the ONTAP setup process. To create additional aggregates, determine the aggregate name, the node on which to create it, and the number of disks it contains.

To create aggregates, run the following commands:

```

aggr create -aggregate aggr1_nodeA -node <<var_nodeA>> -diskcount
<<var_num_disks>>
aggr create -aggregate aggr1_nodeB -node <<var_nodeB>> -diskcount
<<var_num_disks>>

```

Retain at least one disk (select the largest disk) in the configuration as a spare. A best practice is to have at least one spare for each disk type and size.

Start with five disks; you can add disks to an aggregate when additional storage is required.

The aggregate cannot be created until disk zeroing completes. Run the `aggr show` command to display the aggregate creation status. Do not proceed until `aggr1_nodeA` is online.

Configure time zone in ONTAP

To configure time synchronization and to set the time zone on the cluster, run the following command:

```
timezone <<var_timezone>>
```



For example, in the eastern United States, the time zone is `America/New_York`. After you begin typing the time zone name, press the Tab key to see available options.

Configure SNMP in ONTAP

To configure the SNMP, complete the following steps:

1. Configure SNMP basic information, such as the location and contact. When polled, this information is visible as the `sysLocation` and `sysContact` variables in SNMP.

```
snmp contact <<var_snmp_contact>>
snmp location "<<var_snmp_location>>"
snmp init 1
options snmp.enable on
```

2. Configure SNMP traps to send to remote hosts.

```
snmp traphost add <<var_snmp_server_fqdn>>
```

Configure SNMPv1 in ONTAP

To configure SNMPv1, set the shared secret plain-text password called a community.

```
snmp community add ro <<var_snmp_community>>
```



Use the `snmp community delete all` command with caution. If community strings are used for other monitoring products, this command removes them.

Configure SNMPv3 in ONTAP

SNMPv3 requires that you define and configure a user for authentication. To configure SNMPv3, complete the following steps:

1. Run the `security snmpusers` command to view the engine ID.
2. Create a user called `snmpv3user`.

```
security login create -username snmpv3user -authmethod usm -application snmp
```

3. Enter the authoritative entity's engine ID and select `md5` as the authentication protocol.
4. Enter an eight-character minimum-length password for the authentication protocol when prompted.
5. Select `des` as the privacy protocol.
6. Enter an eight-character minimum-length password for the privacy protocol when prompted.

Configure AutoSupport HTTPS in ONTAP

The NetApp AutoSupport tool sends support summary information to NetApp through HTTPS. To configure AutoSupport, run the following command:

```
system node autosupport modify -node * -state enable -mail-hosts  
<<var_mailhost>> -transport https -support enable -noteto  
<<var_storage_admin_email>>
```

Create a storage virtual machine

To create an infrastructure storage virtual machine (SVM), complete the following steps:

1. Run the `vserver create` command.

```
vserver create -vserver Infra-SVM -rootvolume rootvol -aggregate  
aggr1_nodeA -rootvolume- security-style unix
```

2. Add the data aggregate to the infra-SVM aggregate list for the NetApp VSC.

```
vserver modify -vserver Infra-SVM -aggr-list aggr1_nodeA,aggr1_nodeB
```

3. Remove the unused storage protocols from the SVM, leaving NFS and iSCSI.

```
vserver remove-protocols -vserver Infra-SVM -protocols cifs,ndmp,fc
```

4. Enable and run the NFS protocol in the infra-SVM SVM.

```
nfs create -vserver Infra-SVM -udp disabled
```

5. Turn on the `SVM vstorage` parameter for the NetApp NFS VAAI plug-in. Then, verify that NFS has been configured.

```
vserver nfs modify -vserver Infra-SVM -vstorage enabled
vserver nfs show
```



Commands are prefaced by `vserver` in the command line because SVMs were previously called servers

Configure NFSv3 in ONTAP

The table below lists the information needed to complete this configuration.

| Detail | Detail Value |
|----------------------------|---------------------------|
| ESXi host A NFS IP address | <<var_esxi_hostA_nfs_ip>> |
| ESXi host B NFS IP address | <<var_esxi_hostB_nfs_ip>> |

To configure NFS on the SVM, run the following commands:

1. Create a rule for each ESXi host in the default export policy.
2. For each ESXi host being created, assign a rule. Each host has its own rule index. Your first ESXi host has rule index 1, your second ESXi host has rule index 2, and so on.

```
vserver export-policy rule create -vserver Infra-SVM -policyname default
-ruleindex 1 -protocol nfs -clientmatch <<var_esxi_hostA_nfs_ip>>
-rorule sys -rwrule sys -superuser sys -allow-suid false
vserver export-policy rule create -vserver Infra-SVM -policyname default -ruleindex 2
-protocol nfs -clientmatch <<var_esxi_hostB_nfs_ip>> -rorule sys -rwrule
sys -superuser sys -allow-suid false
vserver export-policy rule show
```

3. Assign the export policy to the infrastructure SVM root volume.

```
volume modify -vserver Infra-SVM -volume rootvol -policy default
```



The NetApp VSC automatically handles export policies if you choose to install it after vSphere has been set up. If you do not install it, you must create export policy rules when additional Cisco UCS B-Series servers are added.

Create iSCSI service in ONTAP

To create the iSCSI service, complete the following step:

1. Create the iSCSI service on the SVM. This command also starts the iSCSI service and sets the iSCSI Qualified Name (IQN) for the SVM. Verify that iSCSI has been configured.


```
iscsi create -vserver Infra-SVM
iscsi show
```

Create load-sharing mirror of SVM root volume in ONTAP

To create a load-sharing mirror of the SVM root volume in ONTAP, complete the following steps:

1. Create a volume to be the load-sharing mirror of the infrastructure SVM root volume on each node.

```
volume create -vserver Infra_Vserver -volume rootvol_m01 -aggregate
aggr1_nodeA -size 1GB -type DPvolume create -vserver Infra_Vserver
-volume rootvol_m02 -aggregate aggr1_nodeB -size 1GB -type DP
```

2. Create a job schedule to update the root volume mirror relationships every 15 minutes.

```
job schedule interval create -name 15min -minutes 15
```

3. Create the mirroring relationships.

```
snapmirror create -source-path Infra-SVM:rootvol -destination-path
Infra-SVM:rootvol_m01 -type LS -schedule 15min
snapmirror create -source-path Infra-SVM:rootvol -destination-path
Infra-SVM:rootvol_m02 -type LS -schedule 15min
```

4. Initialize the mirroring relationship and verify that it has been created.

```
snapmirror initialize-ls-set -source-path Infra-SVM:rootvol snapmirror
show
```

Configure HTTPS access in ONTAP

To configure secure access to the storage controller, complete the following steps:

1. Increase the privilege level to access the certificate commands.

```
set -privilege diag
Do you want to continue? {y|n}: y
```

2. Generally, a self-signed certificate is already in place. Verify the certificate by running the following command:

```
security certificate show
```

3. For each SVM shown, the certificate common name should match the DNS fully qualified domain name (FQDN) of the SVM. The four default certificates should be deleted and replaced by either self-signed certificates or certificates from a certificate authority.

Deleting expired certificates before creating certificates is a best practice. Run the `security certificate delete` command to delete expired certificates. In the following command, use TAB completion to select and delete each default certificate.

```
security certificate delete [TAB] ...  
Example: security certificate delete -vserver Infra-SVM -common-name  
Infra-SVM -ca Infra-SVM - type server -serial 552429A6
```

4. To generate and install self-signed certificates, run the following commands as one-time commands. Generate a server certificate for the infra-SVM and the cluster SVM. Again, use TAB completion to aid in completing these commands.

```
security certificate create [TAB] ...  
Example: security certificate create -common-name infra-svm.netapp.com  
-type server -size 2048 - country US -state "North Carolina" -locality  
"RTP" -organization "NetApp" -unit "FlexPod" -email- addr  
"abc@netapp.com" -expire-days 365 -protocol SSL -hash-function SHA256  
-vserver Infra-SVM
```

5. To obtain the values for the parameters required in the following step, run the `security certificate show` command.
6. Enable each certificate that was just created using the `-server-enabled true` and `-client-enabled false` parameters. Again, use TAB completion.

```
security ssl modify [TAB] ...  
Example: security ssl modify -vserver Infra-SVM -server-enabled true  
-client-enabled false -ca infra-svm.netapp.com -serial 55243646 -common  
-name infra-svm.netapp.com
```

7. Configure and enable SSL and HTTPS access and disable HTTP access.

```
system services web modify -external true -ssl3-enabled true
Warning: Modifying the cluster configuration will cause pending web
service requests to be interrupted as the web servers are restarted.
Do you want to continue {y|n}: y
System services firewall policy delete -policy mgmt -service http
-vserver <<var_clustername>>
```



It is normal for some of these commands to return an error message stating that the entry does not exist.

8. Revert to the admin privilege level and create the setup to allow SVM to be available by the web.

```
set -privilege admin
vserver services web modify -name spi|ontapi|compat -vserver * -enabled
true
```

Create a NetApp FlexVol volume in ONTAP

To create a NetApp FlexVol® volume, enter the volume name, size, and the aggregate on which it exists. Create two VMware datastore volumes and a server boot volume.

```
volume create -vserver Infra-SVM -volume infra_datastore_1 -aggregate
aggr1_nodeA -size 500GB -state online -policy default -junction-path
/infra_datastore_1 -space-guarantee none -percent-snapshot-space 0
volume create -vserver Infra-SVM -volume infra_datastore_2 -aggregate
aggr1_nodeB -size 500GB -state online -policy default -junction-path
/infra_datastore_2 -space-guarantee none -percent-snapshot-space 0
```

```
volume create -vserver Infra-SVM -volume infra_swap -aggregate aggr1_nodeA
-size 100GB -state online -policy default -junction-path /infra_swap -space
-guarantee none -percent-snapshot-space 0 -snapshot-policy none
volume create -vserver Infra-SVM -volume esxi_boot -aggregate aggr1_nodeA
-size 100GB -state online -policy default -space-guarantee none -percent
-snapshot-space 0
```

Enable deduplication in ONTAP

To enable deduplication on appropriate volumes once a day, run the following commands:

```

volume efficiency modify -vserver Infra-SVM -volume esxi_boot -schedule
sun-sat@0
volume efficiency modify -vserver Infra-SVM -volume infra_datastore_1
-schedule sun-sat@0
volume efficiency modify -vserver Infra-SVM -volume infra_datastore_2
-schedule sun-sat@0

```

Create LUNs in ONTAP

To create two boot logical unit numbers (LUNs), run the following commands:

```

lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-A -size
15GB -ostype vmware - space-reserve disabled
lun create -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra-B -size
15GB -ostype vmware - space-reserve disabled

```



When adding an extra Cisco UCS C-Series server, an extra boot LUN must be created.

Create iSCSI LIFs in ONTAP

The table below lists the information needed to complete this configuration.

| Detail | Detail Value |
|--|---------------------------------|
| Storage node A iSCSI LIF01A | <<var_nodeA_iscsi_lif01a_ip>> |
| Storage node A iSCSI LIF01A network mask | <<var_nodeA_iscsi_lif01a_mask>> |
| Storage node A iSCSI LIF01B | <<var_nodeA_iscsi_lif01b_ip>> |
| Storage node A iSCSI LIF01B network mask | <<var_nodeA_iscsi_lif01b_mask>> |
| Storage node B iSCSI LIF01A | <<var_nodeB_iscsi_lif01a_ip>> |
| Storage node B iSCSI LIF01A network mask | <<var_nodeB_iscsi_lif01a_mask>> |
| Storage node B iSCSI LIF01B | <<var_nodeB_iscsi_lif01b_ip>> |
| Storage node B iSCSI LIF01B network mask | <<var_nodeB_iscsi_lif01b_mask>> |

1. Create four iSCSI LIFs, two on each node.

```

network interface create -vserver Infra-SVM -lif iscsi_lif01a -role data
-data-protocol iscsi - home-node <<var_nodeA>> -home-port e0e-
<<var_iscsi_vlan_A_id>> -address <<var_nodeA_iscsi_lif01a_ip>> -netmask
<<var_nodeA_iscsi_lif01a_mask>> -status-admin up - failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif01b -role data
-data-protocol iscsi - home-node <<var_nodeA>> -home-port e0f-
<<var_iscsi_vlan_B_id>> -address <<var_nodeA_iscsi_lif01b_ip>> -netmask
<<var_nodeA_iscsi_lif01b_mask>> -status-admin up - failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02a -role data
-data-protocol iscsi - home-node <<var_nodeB>> -home-port e0e-
<<var_iscsi_vlan_A_id>> -address <<var_nodeB_iscsi_lif01a_ip>> -netmask
<<var_nodeB_iscsi_lif01a_mask>> -status-admin up - failover-policy
disabled -firewall-policy data -auto-revert false
network interface create -vserver Infra-SVM -lif iscsi_lif02b -role data
-data-protocol iscsi - home-node <<var_nodeB>> -home-port e0f-
<<var_iscsi_vlan_B_id>> -address <<var_nodeB_iscsi_lif01b_ip>> -netmask
<<var_nodeB_iscsi_lif01b_mask>> -status-admin up - failover-policy
disabled -firewall-policy data -auto-revert false
network interface show

```

Create NFS LIFs in ONTAP

The following table lists the information needed to complete this configuration.

| Detail | Detail value |
|--|---------------------------------|
| Storage node A NFS LIF 01 a IP | <<var_nodeA_nfs_lif_01_a_ip>> |
| Storage node A NFS LIF 01 a network mask | <<var_nodeA_nfs_lif_01_a_mask>> |
| Storage node A NFS LIF 01 b IP | <<var_nodeA_nfs_lif_01_b_ip>> |
| Storage node A NFS LIF 01 b network mask | <<var_nodeA_nfs_lif_01_b_mask>> |
| Storage node B NFS LIF 02 a IP | <<var_nodeB_nfs_lif_02_a_ip>> |
| Storage node B NFS LIF 02 a network mask | <<var_nodeB_nfs_lif_02_a_mask>> |
| Storage node B NFS LIF 02 b IP | <<var_nodeB_nfs_lif_02_b_ip>> |
| Storage node B NFS LIF 02 b network mask | <<var_nodeB_nfs_lif_02_b_mask>> |

1. Create an NFS LIF.

```

network interface create -vserver Infra-SVM -lif nfs_lif01_a -role data
-data-protocol nfs -home- node <<var_nodeA>> -home-port e0e-
<<var_nfs_vlan_id>> -address <<var_nodeA_nfs_lif_01_a_ip>> - netmask <<
var_nodeA_nfs_lif_01_a_mask>> -status-admin up -failover-policy
broadcast-domain-wide - firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif01_b -role data
-data-protocol nfs -home- node <<var_nodeA>> -home-port e0f-
<<var_nfs_vlan_id>> -address <<var_nodeA_nfs_lif_01_b_ip>> - netmask <<
var_nodeA_nfs_lif_01_b_mask>> -status-admin up -failover-policy
broadcast-domain-wide - firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif02_a -role data
-data-protocol nfs -home- node <<var_nodeB>> -home-port e0e-
<<var_nfs_vlan_id>> -address <<var_nodeB_nfs_lif_02_a_ip>> - netmask <<
var_nodeB_nfs_lif_02_a_mask>> -status-admin up -failover-policy
broadcast-domain-wide - firewall-policy data -auto-revert true
network interface create -vserver Infra-SVM -lif nfs_lif02_b -role data
-data-protocol nfs -home- node <<var_nodeB>> -home-port e0f-
<<var_nfs_vlan_id>> -address <<var_nodeB_nfs_lif_02_b_ip>> - netmask <<
var_nodeB_nfs_lif_02_b_mask>> -status-admin up -failover-policy
broadcast-domain-wide - firewall-policy data -auto-revert true
network interface show

```

Add infrastructure SVM administrator

The following table lists the information needed to complete this configuration.

| Detail | Detail value |
|------------------------|--------------------------|
| Vsmgmt IP | <<var_svm_mgmt_ip>> |
| Vsmgmt network mask | <<var_svm_mgmt_mask>> |
| Vsmgmt default gateway | <<var_svm_mgmt_gateway>> |

To add the infrastructure SVM administrator and SVM administration LIF to the management network, complete the following steps:

1. Run the following command:

```

network interface create -vserver Infra-SVM -lif vsmgmt -role data
-data-protocol none -home-node <<var_nodeB>> -home-port e0M -address
<<var_svm_mgmt_ip>> -netmask <<var_svm_mgmt_mask>> - status-admin up
-failover-policy broadcast-domain-wide -firewall-policy mgmt -auto-
revert true

```



The SVM management IP here should be in the same subnet as the storage cluster management IP.

2. Create a default route to allow the SVM management interface to reach the outside world.

```
network route create -vserver Infra-SVM -destination 0.0.0.0/0 -gateway  
<<var_svm_mgmt_gateway>> network route show
```

3. Set a password for the SVM `vsadmin` user and unlock the user.

```
security login password -username vsadmin -vserver Infra-SVM  
Enter a new password: <<var_password>>  
Enter it again: <<var_password>>  
security login unlock -username vsadmin -vserver
```

Cisco UCS server configuration

FlexPod Cisco UCS base

Perform Initial Setup of Cisco UCS 6324 Fabric Interconnect for FlexPod Environments.

This section provides detailed procedures to configure Cisco UCS for use in a FlexPod ROBO environment by using Cisco UCS Manager.

Cisco UCS fabric interconnect 6324 A

Cisco UCS uses access layer networking and servers. This high-performance, next-generation server system provides a data center with a high degree of workload agility and scalability.

Cisco UCS Manager 4.0(1b) supports the 6324 Fabric Interconnect that integrates the Fabric Interconnect into the Cisco UCS Chassis and provides an integrated solution for a smaller deployment environment. Cisco UCS Mini simplifies the system management and saves cost for the low scale deployments.

The hardware and software components support Cisco's unified fabric, which runs multiple types of data center traffic over a single converged network adapter.

Initial system setup

The first time when you access a fabric interconnect in a Cisco UCS domain, a setup wizard prompts you for the following information required to configure the system:

- Installation method (GUI or CLI)
- Setup mode (restore from full system backup or initial setup)
- System configuration type (standalone or cluster configuration)
- System name
- Admin password
- Management port IPv4 address and subnet mask, or IPv6 address and prefix

- Default gateway IPv4 or IPv6 address
- DNS Server IPv4 or IPv6 address
- Default domain name

The following table lists the information needed to complete the Cisco UCS initial configuration on Fabric Interconnect A

| Detail | Detail/value |
|--|---------------------------|
| System Name | <<var_ucs_clustername>> |
| Admin Password | <<var_password>> |
| Management IP Address: Fabric Interconnect A | <<var_ucsa_mgmt_ip>> |
| Management netmask: Fabric Interconnect A | <<var_ucsa_mgmt_mask>> |
| Default gateway: Fabric Interconnect A | <<var_ucsa_mgmt_gateway>> |
| Cluster IP address | <<var_ucs_cluster_ip>> |
| DNS server IP address | <<var_nameserver_ip>> |
| Domain name | <<var_domain_name>> |

To configure the Cisco UCS for use in a FlexPod environment, complete the following steps:

1. Connect to the console port on the first Cisco UCS 6324 Fabric Interconnect A.

Enter the configuration method. (console/gui) ? console

Enter the setup mode; setup newly or restore from backup.
(setup/restore) ? setup

You have chosen to setup a new Fabric interconnect. Continue? (y/n): y

Enforce strong password? (y/n) [y]: Enter

Enter the password for "admin":<<var_password>>
Confirm the password for "admin":<<var_password>>

Is this Fabric interconnect part of a cluster(select 'no' for standalone)? (yes/no) [n]: yes

Enter the switch fabric (A/B) []: A

Enter the system name: <<var_ucs_clustername>>

Physical Switch Mgmt0 IP address : <<var_ucsa_mgmt_ip>>

Physical Switch Mgmt0 IPv4 netmask : <<var_ucsa_mgmt_mask>>

IPv4 address of the default gateway : <<var_ucsa_mgmt_gateway>>

Cluster IPv4 address : <<var_ucs_cluster_ip>>

Configure the DNS Server IP address? (yes/no) [n]: y

DNS IP address : <<var_nameserver_ip>>

Configure the default domain name? (yes/no) [n]: y
Default domain name: <<var_domain_name>>

Join centralized management environment (UCS Central)? (yes/no) [n]:
no

NOTE: Cluster IP will be configured only after both Fabric Interconnects are initialized. UCSM will be functional only after peer FI is configured in clustering mode.

Apply and save the configuration (select 'no' if you want to re-enter)? (yes/no): yes

Applying configuration. Please wait.

Configuration file - Ok

2. Review the settings displayed on the console. If they are correct, answer yes to apply and save the configuration.
3. Wait for the login prompt to verify that the configuration has been saved.

The following table lists the information needed to complete the Cisco UCS initial configuration on Fabric Interconnect B.

| Detail | Detail/value |
|----------------------------|---------------------------|
| System Name | <<var_ucs_clustername>> |
| Admin Password | <<var_password>> |
| Management IP Address-FI B | <<var_ucsb_mgmt_ip>> |
| Management Netmask-FI B | <<var_ucsb_mgmt_mask>> |
| Default Gateway-FI B | <<var_ucsb_mgmt_gateway>> |
| Cluster IP Address | <<var_ucs_cluster_ip>> |
| DNS Server IP address | <<var_nameserver_ip>> |
| Domain Name | <<var_domain_name>> |

1. Connect to the console port on the second Cisco UCS 6324 Fabric Interconnect B.

```
Enter the configuration method. (console/gui) ? console

Installer has detected the presence of a peer Fabric interconnect.
This Fabric interconnect will be added to the cluster. Continue (y/n) ?
y

Enter the admin password of the peer Fabric
interconnect:<<var_password>>
Connecting to peer Fabric interconnect... done
Retrieving config from peer Fabric interconnect... done
Peer Fabric interconnect Mgmt0 IPv4 Address: <<var_ucsb_mgmt_ip>>
Peer Fabric interconnect Mgmt0 IPv4 Netmask: <<var_ucsb_mgmt_mask>>
Cluster IPv4 address: <<var_ucs_cluster_address>>

Peer FI is IPv4 Cluster enabled. Please Provide Local Fabric
Interconnect Mgmt0 IPv4 Address

Physical Switch Mgmt0 IP address : <<var_ucsb_mgmt_ip>>

Apply and save the configuration (select 'no' if you want to re-
enter)? (yes/no): yes
Applying configuration. Please wait.

Configuration file - Ok
```

2. Wait for the login prompt to confirm that the configuration has been saved.

Log into Cisco UCS Manager

To log into the Cisco Unified Computing System (UCS) environment, complete the following steps:

1. Open a web browser and navigate to the Cisco UCS Fabric Interconnect cluster address.

You may need to wait at least 5 minutes after configuring the second fabric interconnect for Cisco UCS Manager to come up.

2. Click the Launch UCS Manager link to launch Cisco UCS Manager.
3. Accept the necessary security certificates.
4. When prompted, enter admin as the user name and enter the administrator password.
5. Click Login to log in to Cisco UCS Manager.

Cisco UCS Manager software version 4.0(1b)

This document assumes the use of Cisco UCS Manager Software version 4.0(1b). To upgrade the Cisco UCS Manager software and the Cisco UCS 6324 Fabric Interconnect software refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Configure Cisco UCS Call Home

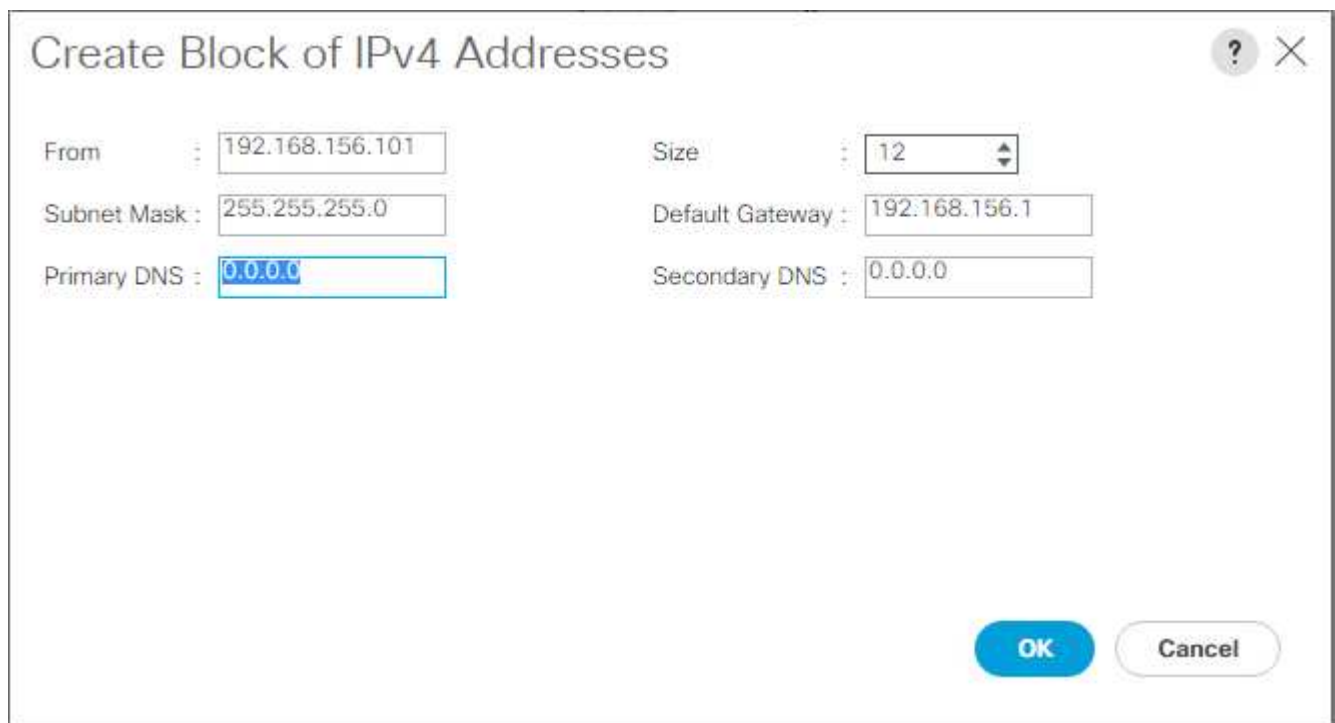
Cisco highly recommends that you configure Call Home in Cisco UCS Manager. Configuring Call Home accelerates the resolution of support cases. To configure Call Home, complete the following steps:

1. In Cisco UCS Manager, click Admin on the left.
2. Select All > Communication Management > Call Home.
3. Change the State to On.
4. Fill in all the fields according to your Management preferences and click Save Changes and OK to complete configuring Call Home.

Add block of IP addresses for keyboard, video, mouse access

To create a block of IP addresses for in band server keyboard, video, mouse (KVM) access in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Expand Pools > root > IP Pools.
3. Right-click IP Pool ext-mgmt and select Create Block of IPv4 Addresses.
4. Enter the starting IP address of the block, number of IP addresses required, and the subnet mask and gateway information.



Create Block of IPv4 Addresses

| | | | |
|---------------|--|-------------------|--|
| From : | <input type="text" value="192.168.156.101"/> | Size : | <input type="text" value="12"/> |
| Subnet Mask : | <input type="text" value="255.255.255.0"/> | Default Gateway : | <input type="text" value="192.168.156.1"/> |
| Primary DNS : | <input type="text" value="0.0.0.0"/> | Secondary DNS : | <input type="text" value="0.0.0.0"/> |

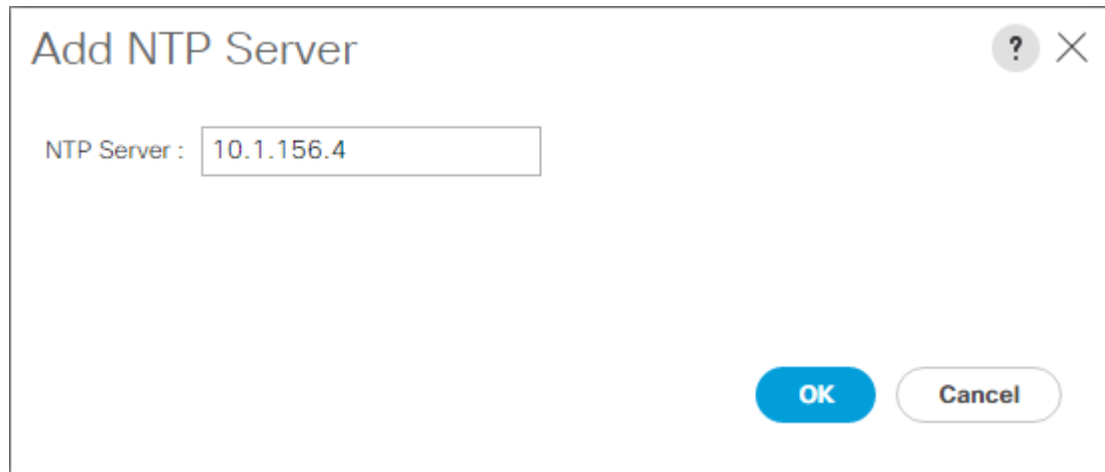
OK Cancel

5. Click OK to create the block.
6. Click OK in the confirmation message.

Synchronize Cisco UCS to NTP

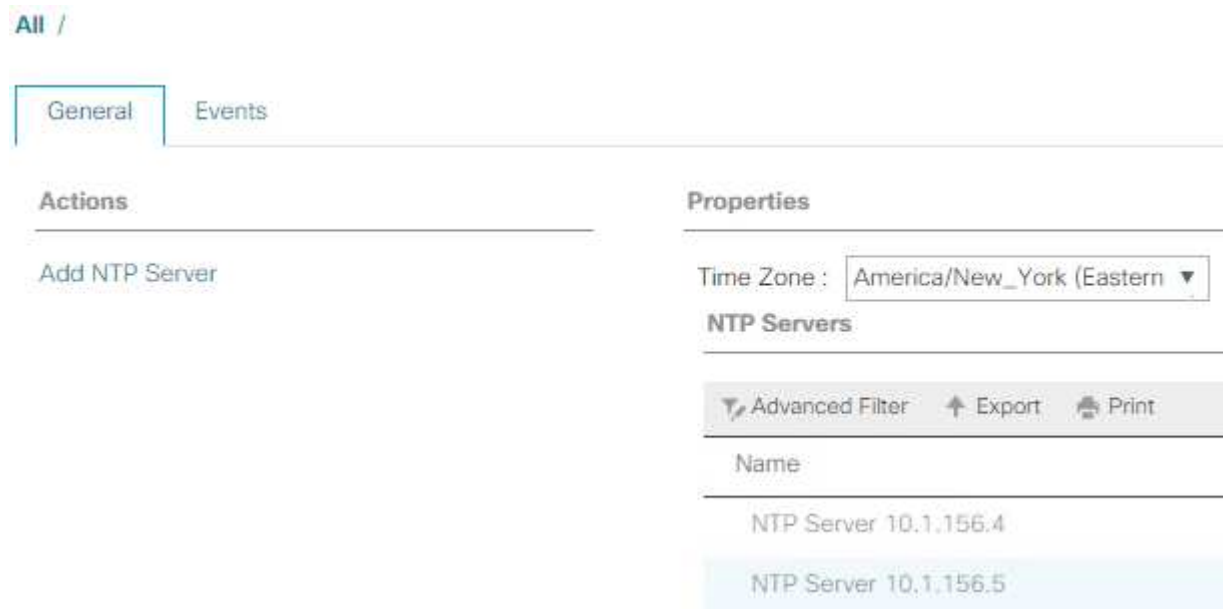
To synchronize the Cisco UCS environment to the NTP servers in the Nexus switches, complete the following steps:

1. In Cisco UCS Manager, click Admin on the left.
2. Expand All > Time Zone Management.
3. Select Time Zone.
4. In the Properties pane, select the appropriate time zone in the Time Zone menu.
5. Click Save Changes and click OK.
6. Click Add NTP Server.
7. Enter <switch-a-ntp-ip> or <Nexus-A-mgmt-IP> and click OK. Click OK.



The image shows a dialog box titled "Add NTP Server". It has a question mark icon and a close button (X) in the top right corner. Inside the dialog, there is a label "NTP Server :" followed by a text input field containing the IP address "10.1.156.4". At the bottom right of the dialog, there are two buttons: "OK" (a blue button) and "Cancel" (a white button with a grey border).

8. Click Add NTP Server.
9. Enter <switch-b-ntp-ip> or <Nexus-B-mgmt-IP> and click OK. Click OK on the confirmation.



The image shows a screenshot of the Cisco UCS Manager interface. At the top, there is a breadcrumb "All /". Below it, there are two tabs: "General" (which is selected) and "Events". The main content area is divided into two panes. The left pane, titled "Actions", contains a link "Add NTP Server". The right pane, titled "Properties", contains a "Time Zone" dropdown menu set to "America/New_York (Eastern)". Below this, there is a section titled "NTP Servers" which contains a table. The table has a header "Name" and two rows of data: "NTP Server 10.1.156.4" and "NTP Server 10.1.156.5". Above the table, there are three buttons: "Advanced Filter", "Export", and "Print".

Edit chassis discovery policy

Setting the discovery policy simplifies the addition of Cisco UCS B-Series chassis and of additional fabric extenders for further Cisco UCS C-Series connectivity. To modify the chassis discovery policy, complete the following steps:


1. In Cisco UCS Manager, click Equipment on the left and select Equipment in the second list.
2. In the right pane, select the Policies tab.
3. Under Global Policies, set the Chassis/FEX Discovery Policy to match the minimum number of uplink ports that are cabled between the chassis or fabric extenders (FEXes) and the fabric interconnects.
4. Set the Link Grouping Preference to Port Channel. If the environment being setup contains a large amount of multicast traffic, set the Multicast Hardware Hash setting to Enabled.
5. Click Save Changes.
6. Click OK.

Enable server, uplink, and storage ports

To enable server and uplink ports, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, select the Equipment tab.
2. Expand Equipment > Fabric Interconnects > Fabric Interconnect A > Fixed Module.
3. Expand Ethernet Ports.
4. Select ports 1 and 2 that are connected to the Cisco Nexus 31108 switches, right-click, and select Configure as Uplink Port.
5. Click Yes to confirm the uplink ports and click OK.
6. Select ports 3 and 4 that are connected to the NetApp Storage Controllers, right-click, and select Configure as Appliance Port.
7. Click Yes to confirm the appliance ports.
8. On the Configure as Appliance Port window, click OK.
9. Click OK to confirm.
10. In the left pane, select Fixed Module under Fabric Interconnect A.
11. From the Ethernet Ports tab, confirm that ports have been configured correctly in the If Role column. If any port C-Series servers were configured on the Scalability port, click on it to verify port connectivity there.

Equipment / Fabric Interconnects / Fabric Interconnect A (subordinate) / Fixed Module

| General Ethernet Ports FC Ports Faults Events | | | | | | | | |
|--|---------------|---------|-------------------|-------------------|----------|-----------------|-------------|------|
| Advanced Filter Export Print <input checked="" type="checkbox"/> All <input checked="" type="checkbox"/> Unconfigured <input checked="" type="checkbox"/> Network <input checked="" type="checkbox"/> Server <input checked="" type="checkbox"/> FCoE Uplink <input checked="" type="checkbox"/> Unified Uplink <input checked="" type="checkbox"/> Appliance Storage <input checked="" type="checkbox"/> FCoE Storage <input checked="" type="checkbox"/> Unified Storage <input checked="" type="checkbox"/> Monitor  | | | | | | | | |
| Slot | Aggr. Port ID | Port ID | MAC | If Role | If Type | Overall Status | Admin State | Peer |
| 1 | 0 | 1 | 00:DE:FB:30:36:88 | Network | Physical | Up | Enabled | |
| 1 | 0 | 2 | 00:DE:FB:30:36:89 | Network | Physical | Up | Enabled | |
| 1 | 0 | 3 | 00:DE:FB:30:36:8A | Appliance Storage | Physical | Up | Enabled | |
| 1 | 0 | 4 | 00:DE:FB:30:36:8B | Appliance Storage | Physical | Up | Enabled | |
| 1 | 5 | 1 | 00:DE:FB:30:36:8C | Unconfigured | Physical | Sfp Not Present | Disabled | |
| 1 | 5 | 2 | 00:DE:FB:30:36:8D | Unconfigured | Physical | Sfp Not Present | Disabled | |
| 1 | 5 | 3 | 00:DE:FB:30:36:8E | Unconfigured | Physical | Sfp Not Present | Disabled | |
| 1 | 5 | 4 | 00:DE:FB:30:36:8F | Unconfigured | Physical | Sfp Not Present | Disabled | |

12. Expand Equipment > Fabric Interconnects > Fabric Interconnect B > Fixed Module.
13. Expand Ethernet Ports.
14. Select Ethernet ports 1 and 2 that are connected to the Cisco Nexus 31108 switches, right-click, and select Configure as Uplink Port.
15. Click Yes to confirm the uplink ports and click OK.
16. Select ports 3 and 4 that are connected to the NetApp Storage Controllers, right-click, and select Configure as Appliance Port.
17. Click Yes to confirm the appliance ports.
18. On the Configure as Appliance Port window, click OK.
19. Click OK to confirm.
20. In the left pane, select Fixed Module under Fabric Interconnect B.
21. From the Ethernet Ports tab, confirm that ports have been configured correctly in the If Role column. If any port C-Series servers were configured on the Scalability port, click it to verify port connectivity there.

Equipment / Fabric Interconnects / Fabric Interconnect B (primar... / Fixed Module / Ethernet Ports

Ethernet Ports

| <input type="checkbox"/> Advanced Filter <input type="button" value="Export"/> <input type="button" value="Print"/> <input checked="" type="checkbox"/> All <input type="checkbox"/> Unconfigured <input checked="" type="checkbox"/> Network <input checked="" type="checkbox"/> Server <input checked="" type="checkbox"/> FCoE Uplink <input checked="" type="checkbox"/> Unified Uplink <input checked="" type="checkbox"/> Appliance Storage <input checked="" type="checkbox"/> FCoE Storage <input checked="" type="checkbox"/> Unified Storage <input checked="" type="checkbox"/> Monitor | | | | | | | | | |
|--|---------------|---------|-------------------|-------------------|----------|-----------------|-------------|------|--|
| Slot | Aggr. Port ID | Port ID | MAC | If Role | If Type | Overall Status | Admin State | Peer | |
| 1 | 0 | 1 | 00:DE:FB:30:3A:C8 | Network | Physical | Up | Enabled | | |
| 1 | 0 | 2 | 00:DE:FB:30:3A:C9 | Network | Physical | Up | Enabled | | |
| 1 | 0 | 3 | 00:DE:FB:30:3A:CA | Appliance Storage | Physical | Up | Enabled | | |
| 1 | 0 | 4 | 00:DE:FB:30:3A:CB | Appliance Storage | Physical | Up | Enabled | | |
| 1 | 5 | 1 | 00:DE:FB:30:3A:CC | Unconfigured | Physical | Sfp Not Present | Disabled | | |
| 1 | 5 | 2 | 00:DE:FB:30:3A:CD | Unconfigured | Physical | Sfp Not Present | Disabled | | |
| 1 | 5 | 3 | 00:DE:FB:30:3A:CE | Unconfigured | Physical | Sfp Not Present | Disabled | | |
| 1 | 5 | 4 | 00:DE:FB:30:3A:CF | Unconfigured | Physical | Sfp Not Present | Disabled | | |

Create uplink port channels to Cisco Nexus 31108 switches

To configure the necessary port channels in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, select the LAN tab in the navigation pane.



In this procedure, two port channels are created: one from Fabric A to both Cisco Nexus 31108 switches and one from Fabric B to both Cisco Nexus 31108 switches. If you are using standard switches, modify this procedure accordingly. If you are using 1 Gigabit Ethernet (1GbE) switches and GLC-T SFPs on the Fabric Interconnects, the interface speeds of Ethernet ports 1/1 and 1/2 in the Fabric Interconnects must be set to 1Gbps.

2. Under LAN > LAN Cloud, expand the Fabric A tree.
3. Right-click Port Channels.
4. Select Create Port Channel.
5. Enter 13 as the unique ID of the port channel.

6. Enter vPC-13-Nexus as the name of the port channel.
7. Click Next.

The screenshot shows a 'Create Port Channel' window. On the left, a blue sidebar contains two numbered steps: '1 Set Port Channel Name' and '2 Add Ports'. The main content area has two input fields: 'ID' with the value '1' and 'Name' with the value 'vPC-13-Nexus'. At the bottom right, there are four buttons: 'Previous' (disabled), 'Next >' (active), 'Finish' (disabled), and 'Cancel' (disabled). A help icon (?) and a close icon (X) are in the top right corner.

8. Select the following ports to be added to the port channel:
 - a. Slot ID 1 and port 1
 - b. Slot ID 1 and port 2
9. Click >> to add the ports to the port channel.
10. Click Finish to create the port channel. Click OK.
11. Under Port Channels, select the newly created port channel.

The port channel should have an Overall Status of Up.

12. In the navigation pane, under LAN > LAN Cloud, expand the Fabric B tree.
13. Right-click Port Channels.
14. Select Create Port Channel.
15. Enter 14 as the unique ID of the port channel.
16. Enter vPC-14-Nexus as the name of the port channel. Click Next.
17. Select the following ports to be added to the port channel:
 - a. Slot ID 1 and port 1
 - b. Slot ID 1 and port 2
18. Click >> to add the ports to the port channel.
19. Click Finish to create the port channel. Click OK.

20. Under Port Channels, select the newly created port-channel.

21. The port channel should have an Overall Status of Up.

Create an organization (optional)

Organizations are used to organizing resources and restricting access to various groups within the IT organization, thereby enabling multitenancy of the compute resources.



Although this document does not assume the use of organizations, this procedure provides instructions for creating one.

To configure an organization in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, from the New menu in the toolbar at the top of the window, select Create Organization.
2. Enter a name for the organization.
3. Optional: Enter a description for the organization. Click OK.
4. Click OK in the confirmation message.

Configure storage appliance ports and storage VLANs

To configure the storage appliance ports and storage VLANs, complete the following steps:

1. In the Cisco UCS Manager, select the LAN tab.
2. Expand the Appliances cloud.
3. Right-click VLANs under Appliances Cloud.
4. Select Create VLANs.
5. Enter NFS-VLAN as the name for the Infrastructure NFS VLAN.
6. Leave Common/Global selected.
7. Enter <<var_nfs_vlan_id>> for the VLAN ID.
8. Leave Sharing Type set to None.

Create VLANs

Create VLANs

VLAN Name/Prefix : NFS-VLAN

☒ Common/Global
 ☐ Fabric A
 ☐ Fabric B
 ☐ Both Fabrics Configured Differently

You are creating global VLANs that map to the same VLAN IDs in all available fabrics.
Enter the range of VLAN IDs.(e.g. "2009-2019", "29,35,40-45", "23", "23,34-45")

VLAN IDs : 3170

Sharing Type :
 ☒ None
 ☐ Primary
 ☐ Isolated
 ☐ Community

Check Overlap Ok Cancel

9. Click OK, and then click OK again to create the VLAN.
10. Right-click VLANs under Appliances Cloud.
11. Select Create VLANs.
12. Enter iSCSI-A-VLAN as the name for the Infrastructure iSCSI Fabric A VLAN.
13. Leave Common/Global selected.
14. Enter <<var_iscsi-a_vlan_id>> for the VLAN ID.
15. Click OK, and then click OK again to create the VLAN.
16. Right-click VLANs under Appliances Cloud.
17. Select Create VLANs.
18. Enter iSCSI-B-VLAN as the name for the Infrastructure iSCSI Fabric B VLAN.
19. Leave Common/Global selected.
20. Enter <<var_iscsi-b_vlan_id>> for the VLAN ID.
21. Click OK, and then click OK again to create the VLAN.
22. Right-click VLANs under Appliances Cloud.

23. Select Create VLANs.
24. Enter Native-VLAN as the name for the Native VLAN.
25. Leave Common/Global selected.
26. Enter <<var_native_vlan_id>> for the VLAN ID.
27. Click OK, and then click OK again to create the VLAN.

LAN / LAN Cloud / VLANs

VLANs

Advanced Filter Expand Print

| Name | ID | Type | Transport | Native | VLAN Sharing | Primary VLAN Name | Multicast Policy Name |
|-------------------------|-----|------|-----------|--------|--------------|-------------------|-----------------------|
| VLAN default (1) | 1 | Lan | Ether | Yes | None | | |
| VLAN 0002-Native (2) | 2 | Lan | Ether | No | None | | |
| VLAN public (18) | 18 | Lan | Ether | No | None | | |
| VLAN 0101-IB-MGMT (101) | 101 | Lan | Ether | No | None | | |
| VLAN 0102-VM (102) | 102 | Lan | Ether | No | None | | |
| VLAN 0103-yMotion (103) | 103 | Lan | Ether | No | None | | |
| VLAN 0104-NFS (104) | 104 | Lan | Ether | No | None | | |
| VLAN 0120-SCSI-A (120) | 120 | Lan | Ether | No | None | | |
| VLAN 0121-SCSI-B (121) | 121 | Lan | Ether | No | None | | |

28. In the navigation pane, under LAN > Policies, expand Appliances and right-click Network Control Policies.
29. Select Create Network Control Policy.
30. Name the policy Enable_CDP_LLPD and select Enabled next to CDP.
31. Enable the Transmit and Receive features for LLDP.

Properties for: Enable_CDP

General Events

Actions

Delete

Show Policy Usage

Use Global

Properties

Name: Enable_CDP

Description:

Owner: Local

CDP: ☐ Disabled ☒ Enabled

MAC Register Mode: ☒ Only Native Vlan ☐ All Host Vlans

Action on Uplink Fail: ☒ Link Down ☐ Warning

MAC Security

Forge: ☒ Allow ☐ Deny

LLDP

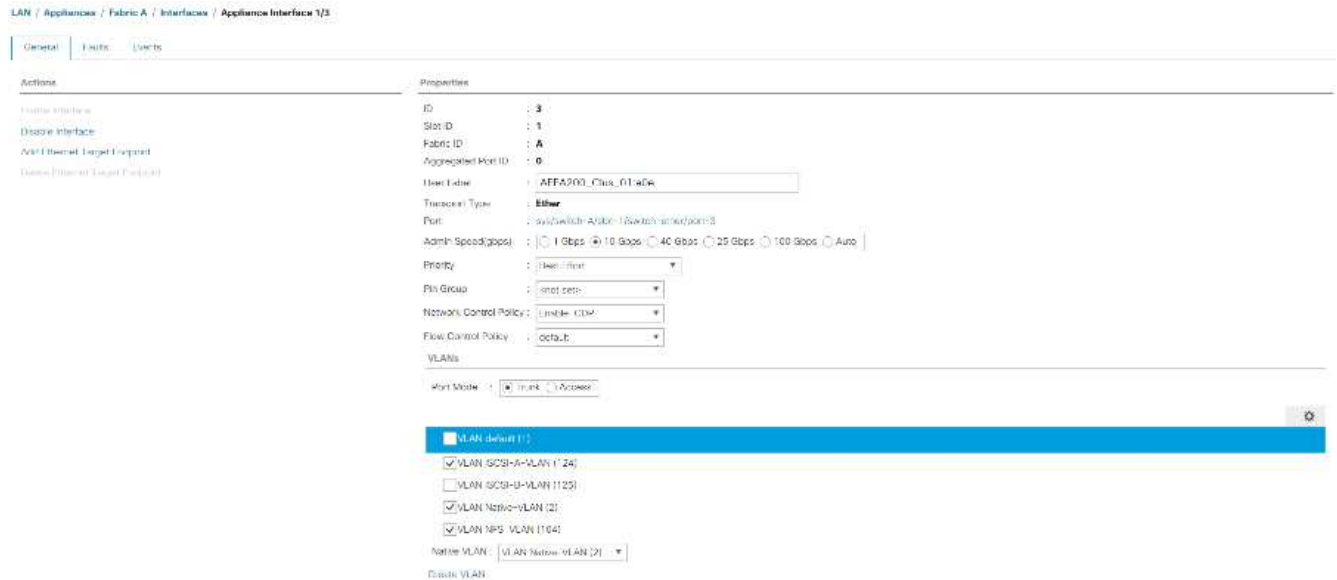
Transmit: ☐ Disabled ☒ Enabled

Receive: ☐ Disabled ☒ Enabled

OK Apply Cancel Help

32. Click OK and then click OK again to create the policy.

33. In the navigation pane, under LAN > Appliances Cloud, expand the Fabric A tree.
34. Expand Interfaces.
35. Select Appliance Interface 1/3.
36. In the User Label field, put in information indicating the storage controller port, such as <storage_controller_01_name>:e0e. Click Save Changes and OK.
37. Select the Enable_CDP Network Control Policy and select Save Changes and OK.
38. Under VLANs, select the iSCSI-A-VLAN, NFS VLAN, and Native VLAN. Set the Native-VLAN as the Native VLAN. Clear the default VLAN selection.
39. Click Save Changes and OK.



LAN / Appliances / Fabric A / Interfaces / Appliance Interface 1/3

General | Faults | Events

Actions

- Create Interface
- Discover Interface
- Add Ethernet Target Endpoint
- Remove Ethernet Target Endpoint

Properties

ID: 3
Slot ID: 1
Fabric ID: A
Aggregated Port ID: 0
User Label: AFFA200_Clus_01-e0e
Trunking Type: Ether
Port: 2x25Gbit-A/30e-1/25Gbit-A/30e/30e/3
Admin Speed(gbps): 1 Gbps 10 Gbps 40 Gbps 25 Gbps 100 Gbps Auto
Priority: 1 (High / Low)
Pin Group: 1 (High / Low)
Network Control Policy: Enable_CDP
Flow Control Policy: default
VLANs

Port Mode: 1 (High / Low)

☒ VLAN default (1)
☒ VLAN iSCSI-A-VLAN (124)
☐ VLAN iSCSI-B-VLAN (125)
☒ VLAN NFS-VLAN (2)
☒ VLAN NFS-VLAN (104)
Native VLAN: VLAN Native-VLAN (2)
Delete VLAN

40. Select Appliance Interface 1/4 under Fabric A.
41. In the User Label field, put in information indicating the storage controller port, such as <storage_controller_02_name>:e0e. Click Save Changes and OK.
42. Select the Enable_CDP Network Control Policy and select Save Changes and OK.
43. Under VLANs, select the iSCSI-A-VLAN, NFS VLAN, and Native VLAN.
44. Set the Native-VLAN as the Native VLAN.
45. Clear the default VLAN selection.
46. Click Save Changes and OK.
47. In the navigation pane, under LAN > Appliances Cloud, expand the Fabric B tree.
48. Expand Interfaces.
49. Select Appliance Interface 1/3.
50. In the User Label field, put in information indicating the storage controller port, such as <storage_controller_01_name>:e0f. Click Save Changes and OK.
51. Select the Enable_CDP Network Control Policy and select Save Changes and OK.
52. Under VLANs, select the iSCSI-B-VLAN, NFS VLAN, and Native VLAN. Set the Native-VLAN as the Native VLAN. Unselect the default VLAN.

General Faults Events

Actions

- Enable Interface
- Disable Interface
- Act Ethernet Target Endpoint
- Delete Ethernet Target Endpoint

Properties

ID : 3

Slot ID : 1

Fabric ID : B

Aggregated Port ID : 0

User Label : AFFA200_Clus_01:e0f

Transport Type : Ether

Port : sys/switch-B/slot-1/switch-ether/port-3

Admin Speed(gbps) : ☐ 1 Gbps ☒ 10 Gbps ☐ 40 Gbps ☐ 25 Gbps ☐ 100 Gbps ☐ Auto

Priority : Best Effort

Pin Group : <not set>

Network Control Policy : Enable_CDP

Flow Control Policy : default

VLANs

Port Mode : ☒ Trunk ☐ Access

☐ VLAN default (1)

☐ VLAN iSCSI-A-VLAN (124)

☒ VLAN iSCSI-B-VLAN (125)

☒ VLAN Native-VLAN (2)

☒ VLAN NFS_VLAN (104)

Native VLAN : VLAN Native-VLAN (2)

Create VLAN

53. Click Save Changes and OK.

54. Select Appliance Interface 1/4 under Fabric B.

55. In the User Label field, put in information indicating the storage controller port, such as <storage_controller_02_name>:e0f. Click Save Changes and OK.

56. Select the Enable_CDP Network Control Policy and select Save Changes and OK.

57. Under VLANs, select the iSCSI-B-VLAN, NFS VLAN, and Native VLAN. Set the Native-VLAN as the Native VLAN. Unselect the default VLAN.

58. Click Save Changes and OK.

Set jumbo frames in Cisco UCS fabric

To configure jumbo frames and enable quality of service in the Cisco UCS fabric, complete the following steps:

1. In Cisco UCS Manager, in the navigation pane, click the LAN tab.
2. Select LAN > LAN Cloud > QoS System Class.
3. In the right pane, click the General tab.
4. On the Best Effort row, enter 9216 in the box under the MTU column.

LAN / LAN Cloud / QoS System Class

General Events FSM

Actions Use Global Properties Owner: Local

| Priority | Enabled | CoS | Packet Drop | Weight | Weight (%) | MTU | Multicast Optimized |
|---------------|-------------------------------------|-----|-------------------------------------|--------|------------|--------|--------------------------|
| Platinum | <input type="checkbox"/> | 5 | <input type="checkbox"/> | 10 | N/A | normal | <input type="checkbox"/> |
| Gold | <input type="checkbox"/> | 4 | <input checked="" type="checkbox"/> | 9 | N/A | normal | <input type="checkbox"/> |
| Silver | <input type="checkbox"/> | 2 | <input checked="" type="checkbox"/> | 8 | N/A | normal | <input type="checkbox"/> |
| Bronze | <input type="checkbox"/> | 1 | <input checked="" type="checkbox"/> | 7 | N/A | normal | <input type="checkbox"/> |
| Best Effort | <input checked="" type="checkbox"/> | Any | <input checked="" type="checkbox"/> | 5 | 50 | 9216 | <input type="checkbox"/> |
| Fibre Channel | <input checked="" type="checkbox"/> | 3 | <input type="checkbox"/> | 5 | 50 | 10 | N/A |

5. Click Save Changes.

6. Click OK.

Acknowledge Cisco UCS chassis

To acknowledge all Cisco UCS chassis, complete the following steps:

1. In Cisco UCS Manager, select the Equipment tab, then Expand the Equipment tab on the right.
2. Expand Equipment > Chassis.
3. In the Actions for Chassis 1, select Acknowledge Chassis.
4. Click OK and then click OK to complete acknowledging the chassis.
5. Click Close to close the Properties window.

Load Cisco UCS 4.0(1b) firmware images

To upgrade the Cisco UCS Manager software and the Cisco UCS Fabric Interconnect software to version 4.0(1b) refer to [Cisco UCS Manager Install and Upgrade Guides](#).

Create host firmware package

Firmware management policies allow the administrator to select the corresponding packages for a given server configuration. These policies often include packages for adapter, BIOS, board controller, FC adapters, host bus adapter (HBA) option ROM, and storage controller properties.

To create a firmware management policy for a given server configuration in the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Expand Host Firmware Packages.
4. Select default.
5. In the Actions pane, select Modify Package Versions.
6. Select the version 4.0(1b) for both the Blade Packages.

Modify Package Versions

Blade Package : 4.0(1b)B

Rack Package : <not set>

Service Pack :

The images from Service Pack will take precedence over the images from Blade or Rack Package

Excluded Components:

- ☐ Adapter
- ☐ BIOS
- ☐ Board Controller
- ☐ CIMC
- ☐ FC Adapters
- ☐ Flex Flash Controller
- ☐ GPUs
- ☐ HBA Option ROM
- ☐ Host NIC
- ☐ Host NIC Option ROM
- ☒ Local Disk
- ☐ NVME Mswitch Firmware
- ☐ PSU
- ☐ SAS Expander

OK Apply Cancel Help

7. Click OK then OK again to modify the host firmware package.

Create MAC address pools

To configure the necessary MAC address pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Pools > root.


In this procedure, two MAC address pools are created, one for each switching fabric.

3. Right-click MAC Pools under the root organization.
4. Select Create MAC Pool to create the MAC address pool.
5. Enter MAC-Pool-A as the name of the MAC pool.
6. Optional: Enter a description for the MAC pool.
7. Select Sequential as the option for Assignment Order. Click Next.
8. Click Add.
9. Specify a starting MAC address.



For the FlexPod solution, the recommendation is to place 0A in the next-to-last octet of the starting MAC address to identify all of the MAC addresses as fabric A addresses. In our example, we have carried forward the example of also embedding the Cisco UCS domain number information giving us 00:25:B5:32:0A:00 as our first MAC address.

10. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources. Click OK.



11. Click Finish.
12. In the confirmation message, click OK.
13. Right-click MAC Pools under the root organization.
14. Select Create MAC Pool to create the MAC address pool.
15. Enter MAC-Pool-B as the name of the MAC pool.
16. Optional: Enter a description for the MAC pool.
17. Select Sequential as the option for Assignment Order. Click Next.
18. Click Add.
19. Specify a starting MAC address.



For the FlexPod solution, it is recommended to place 0B in the next to last octet of the starting MAC address to identify all the MAC addresses in this pool as fabric B addresses. Once again, we have carried forward in our example of also embedding the Cisco UCS domain number information giving us 00:25:B5:32:0B:00 as our first MAC address.

20. Specify a size for the MAC address pool that is sufficient to support the available blade or server resources. Click OK.
21. Click Finish.
22. In the confirmation message, click OK.

Create iSCSI IQN pool


To configure the necessary IQN pools for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click SAN on the left.
2. Select Pools > root.
3. Right- click IQN Pools.
4. Select Create IQN Suffix Pool to create the IQN pool.
5. Enter IQN-Pool for the name of the IQN pool.
6. Optional: Enter a description for the IQN pool.
7. Enter `iqn.1992-08.com.cisco` as the prefix.
8. Select Sequential for Assignment Order. Click Next.
9. Click Add.
10. Enter `ucs-host` as the suffix.



If multiple Cisco UCS domains are being used, a more specific IQN suffix might need to be used.

11. Enter 1 in the From field.
12. Specify the size of the IQN block sufficient to support the available server resources. Click OK.



13. Click Finish.

Create iSCSI initiator IP address pools

To configure the necessary IP pools iSCSI boot for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Pools > root.
3. Right-click IP Pools.
4. Select Create IP Pool.
5. Enter iSCSI-IP-Pool-A as the name of IP pool.
6. Optional: Enter a description for the IP pool.
7. Select Sequential for the assignment order. Click Next.
8. Click Add to add a block of IP address.
9. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
10. Set the size to enough addresses to accommodate the servers. Click OK.
11. Click Next.
12. Click Finish.

13. Right-click IP Pools.
14. Select Create IP Pool.
15. Enter iSCSI-IP-Pool-B as the name of IP pool.
16. Optional: Enter a description for the IP pool.
17. Select Sequential for the assignment order. Click Next.
18. Click Add to add a block of IP address.
19. In the From field, enter the beginning of the range to assign as iSCSI IP addresses.
20. Set the size to enough addresses to accommodate the servers. Click OK.
21. Click Next.
22. Click Finish.

Create UUID suffix pool

To configure the necessary universally unique identifier (UUID) suffix pool for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click UUID Suffix Pools.
4. Select Create UUID Suffix Pool.
5. Enter UUID-Pool as the name of the UUID suffix pool.
6. Optional: Enter a description for the UUID suffix pool.
7. Keep the prefix at the derived option.
8. Select Sequential for the Assignment Order.
9. Click Next.
10. Click Add to add a block of UUIDs.
11. Keep the From field at the default setting.
12. Specify a size for the UUID block that is sufficient to support the available blade or server resources. Click OK.
13. Click Finish.
14. Click OK.

Create server pool

To configure the necessary server pool for the Cisco UCS environment, complete the following steps:



Consider creating unique server pools to achieve the granularity that is required in your environment.

1. In Cisco UCS Manager, click Servers on the left.
2. Select Pools > root.
3. Right-click Server Pools.

4. Select Create Server Pool.
5. Enter `Infra-Pool` as the name of the server pool.
6. Optional: Enter a description for the server pool. Click Next.
7. Select two (or more) servers to be used for the VMware management cluster and click >> to add them to the `Infra-Pool` server pool.
8. Click Finish.
9. Click OK.

Create Network Control Policy for Cisco Discovery Protocol and Link Layer Discovery Protocol

To create a Network Control Policy for Cisco Discovery Protocol (CDP) and Link Layer Discovery Protocol (LLDP), complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click Network Control Policies.
4. Select Create Network Control Policy.
5. Enter Enable-CDP-LLDP policy name.
6. For CDP, select the Enabled option.
7. For LLDP, scroll down and select Enabled for both Transmit and Receive.
8. Click OK to create the network control policy. Click OK.

Create Network Control Policy

CDP : ☐ Disabled ☒ Enabled

MAC Register Mode : ☒ Only Native Vlan ☐ All Host Vlans

Action on Uplink Fail : ☒ Link Down ☐ Warning

MAC Security

Forge : ☒ Allow ☐ Deny

LLDP

Transmit : ☐ Disabled ☒ Enabled

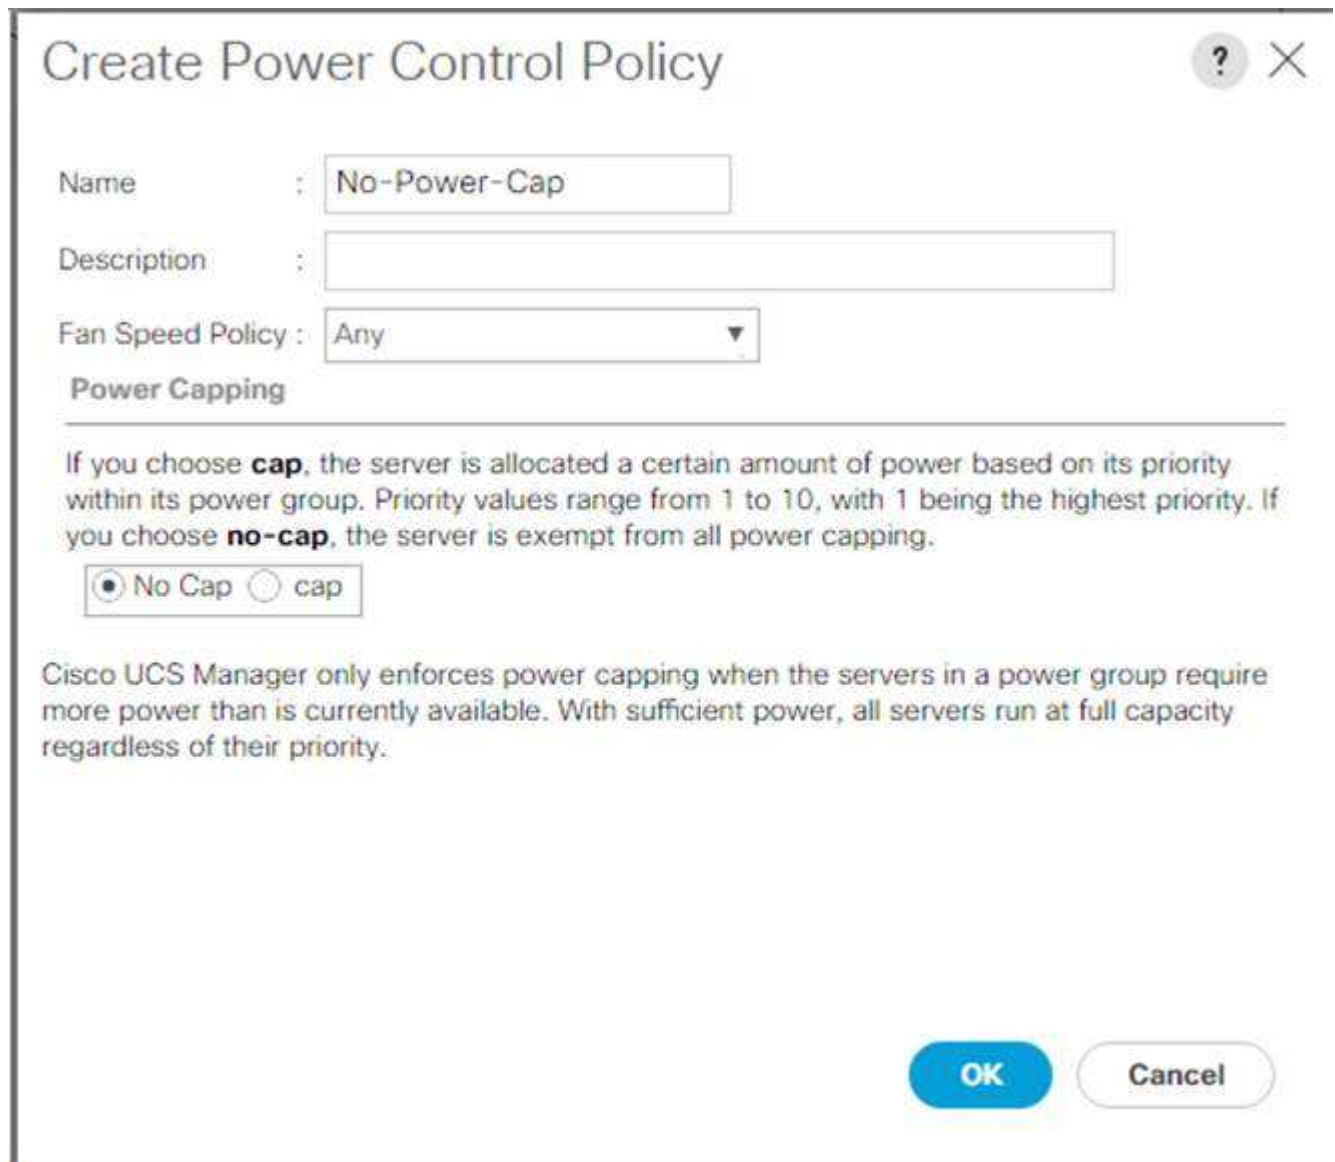
Receive : ☐ Disabled ☒ Enabled

OK **Cancel**

Create power control policy

To create a power control policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers tab on the left.
2. Select Policies > root.
3. Right-click Power Control Policies.
4. Select Create Power Control Policy.
5. Enter No-Power-Cap as the power control policy name.
6. Change the power capping setting to No Cap.
7. Click OK to create the power control policy. Click OK.



The screenshot shows a dialog box titled "Create Power Control Policy" with a question mark icon and a close button (X) in the top right corner. The dialog contains the following fields and options:

- Name:** A text box containing "No-Power-Cap".
- Description:** An empty text box.
- Fan Speed Policy:** A dropdown menu currently set to "Any".
- Power Capping:** A section with explanatory text and two radio buttons.
 - Text: "If you choose **cap**, the server is allocated a certain amount of power based on its priority within its power group. Priority values range from 1 to 10, with 1 being the highest priority. If you choose **no-cap**, the server is exempt from all power capping."
 - Radio buttons: "No Cap" (selected) and "cap".
- Footer text:** "Cisco UCS Manager only enforces power capping when the servers in a power group require more power than is currently available. With sufficient power, all servers run at full capacity regardless of their priority."
- Buttons:** "OK" (blue) and "Cancel" (white with blue border) at the bottom right.


Create server pool qualification policy (Optional)

To create an optional server pool qualification policy for the Cisco UCS environment, complete the following steps:



This example creates a policy for Cisco UCS B-Series servers with the Intel E2660 v4 Xeon Broadwell processors.

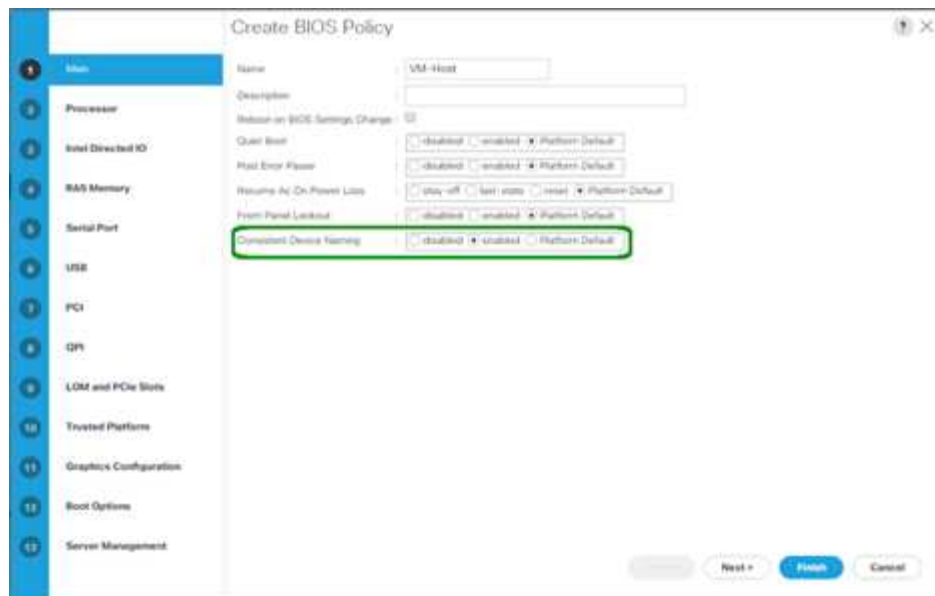
1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Select Server Pool Policy Qualifications.
4. Select Create Server Pool Policy Qualification or Add.
5. Name the policy Intel.
6. Select Create CPU/Cores Qualifications.
7. Select Xeon for the Processor/Architecture.
8. Enter <UCS-CPU- PID> as the process ID (PID).
9. Click OK to create the CPU/Core qualification.
10. Click OK to create the policy, and then click OK for the confirmation.



Create server BIOS policy

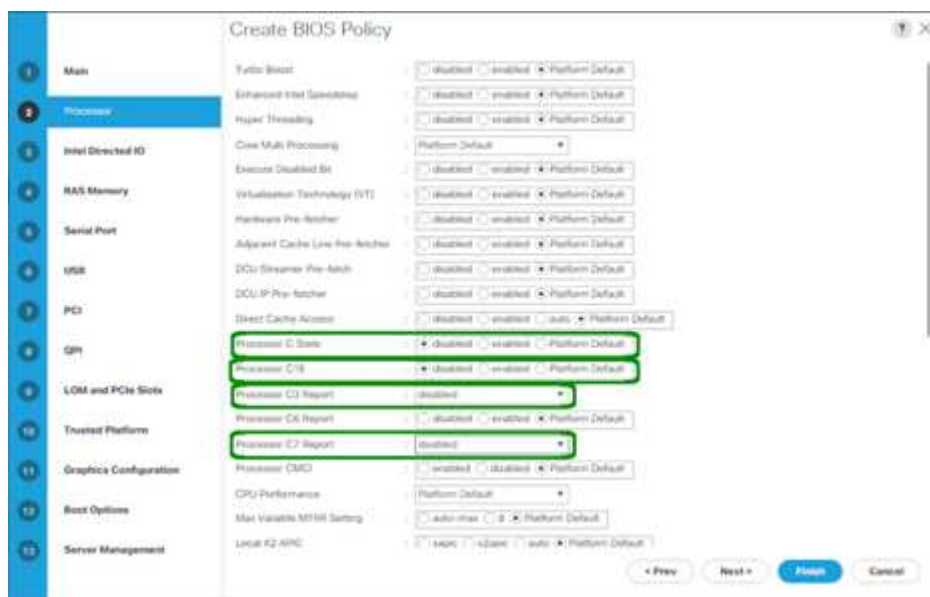
To create a server BIOS policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click BIOS Policies.
4. Select Create BIOS Policy.
5. Enter VM-Host as the BIOS policy name.
6. Change the Quiet Boot setting to disabled.
7. Change Consistent Device Naming to enabled.



8. Select the Processor tab and set the following parameters:

- Processor C State: disabled
- Processor C1E: disabled
- Processor C3 Report: disabled
- Processor C7 Report: disabled



9. Scroll down to the remaining Processor options and set the following parameters:

- Energy Performance: performance
- Frequency Floor Override: enabled
- DRAM Clock Throttling: performance



10. Click RAS Memory and set the following parameters:

- LV DDR Mode: performance mode



11. Click Finish to create the BIOS policy.

12. Click OK.

Update the default maintenance policy

To update the default Maintenance Policy, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Select Maintenance Policies > default.
4. Change the Reboot Policy to User Ack.
5. Select On Next Boot to delegate maintenance windows to server administrators.

Servers / Policies / root / Maintenance Poli... / default

General Events

Actions

Cancel

Show Policy Usage

Use Global

Properties

Name : default

Description :

Owner : Local

Soft Shutdown Timer : 150 Secs

Reboot Policy : ☐ Immediate ☒ User Ack ☐ Timer Automatic

☒ On Next Boot (Apply pending changes at next reboot.)

6. Click Save Changes.
7. Click OK to accept the change.

Create vNIC templates

To create multiple virtual network interface card (vNIC) templates for the Cisco UCS environment, complete the procedures described in this section.



A total of four vNIC templates are created.

Create infrastructure vNICs

To create an infrastructure vNIC, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter Site-XX-vNIC_A as the vNIC template name.
6. Select updating-template as the Template Type.
7. For Fabric ID, select Fabric A.
8. Ensure that the Enable Failover option is not selected.
9. Select Primary Template for Redundancy Type.
10. Leave the Peer Redundancy Template set to <not set>.
11. Under Target, make sure that only the Adapter option is selected.
12. Set Native-VLAN as the native VLAN.
13. Select vNIC Name for the CDN Source.
14. For MTU, enter 9000.
15. Under Permitted VLANs, select Native-VLAN, Site-XX-IB-MGMT, Site-XX-NFS, Site-XX-VM-Traffic, and Site-XX-vMotion. Use the Ctrl key to make this multiple selection.
16. Click Select. These VLANs should now appear under Selected VLANs.
17. In the MAC Pool list, select MAC_Pool_A.

18. In the Network Control Policy list, select Pool-A.
19. In the Network Control Policy list, select Enable-CDP-LLDP.
20. Click OK to create the vNIC template.
21. Click OK.

LAN > Policies > root > vNIC Templates > vNIC_Template_A

General | VLANs | VLAN Groups | Tags | Events

Actions

- Modify vNICs
- Modify vNIC Groups
- Delete
- Show Policy Usage
- Use Clone

Properties

Name: **vNIC_Template_A**

Description:

Owner: **Local**

Fabric ID: ☒ Fabric A ☐ Fabric B ☒ Enable Failover

Redundancy

Redundancy Type: ☐ No Redundancy ☒ Primary Template ☐ Secondary Template

Peer Redundancy Template: **vNIC_Template_B** [Create vNIC Template](#)

Target

☒ vNIC ☐ vNIC Group

Template Type: **Initial Template** [Updating Template](#)

QPV Source: **vNIC Name** [User Defined](#)

VPI: **8000**

Policies

MAC Policy: **MAC_Pool_Access**

QoS Policy: **vnic def**

Network Control Policy: **Enable_CDP**

Pre Queue: **vnic def**

State Threshold Policy: **default**

Connection Policies

☒ Dynamic vNIC ☐ vNIC ☐ VNIC

Dynamic vNIC Connection Policy: **vnic def**

To create the secondary redundancy template Infra-B, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter `Site-XX-vNIC_B` as the vNIC template name.
6. Select updating-template as the Template Type.
7. For Fabric ID, select Fabric B.
8. Select the Enable Failover option.



Selecting Failover is a critical step to improve link failover time by handling it at the hardware level, and to guard against any potential for NIC failure not being detected by the virtual switch.

9. Select Primary Template for Redundancy Type.
10. Leave the Peer Redundancy Template set to vNIC_Template_A.
11. Under Target, make sure that only the Adapter option is selected.
12. Set Native-VLAN as the native VLAN.
13. Select vNIC Name for the CDN Source.
14. For MTU, enter 9000.
15. Under Permitted VLANs, select Native-VLAN, Site-XX-IB-MGMT, Site-XX-NFS, Site-XX-VM-Traffic, and Site-XX-vMotion. Use the Ctrl key to make this multiple selection.
16. Click Select. These VLANs should now appear under Selected VLANs.
17. In the MAC Pool list, select MAC_Pool_B.
18. In the Network Control Policy list, select Pool-B.
19. In the Network Control Policy list, select Enable-CDP-LLDP.
20. Click OK to create the vNIC template.
21. Click OK.

LAN / Policies / root / vNIC Template / vNIC Template vNIC_Template_B

General VLANs VLAN Groups Tags Profiles

Actions

- Modify vNICs
- Modify VLAN Groups
- Delete
- Show Policy Usage
- Use Manual

Properties

Name: vNIC_Template_B

Description:

Owner: Local

Fabric ID: ☐ Fabric A ☒ Fabric B ☒ Enable External

Redundancy

Redundancy Type: ☐ No Redundancy ☐ Primary Template ☒ Secondary Template

Peer Redundancy Template: vNIC_Template_A [Create vNIC Template](#)

Target

☒ Adapter ☐ VM

Template Type: ☐ Native Template ☒ Updating Template

CDN Source: ☒ vNIC Name ☐ User Defined

MTU: 9000

Policies

MAC Pool: MAC_Pool_B(58/64)

QoS Policy: ☐ null add +

Network Control Policy: Enable_CDP

RH Group: ☐ null add +

Stats Threshold Policy: default

Connection Policies

☒ Dynamic vNIC ☐ usNIC ☐ VMQ

Dynamic vNIC Connection Policy: ☐ null add +

Create iSCSI vNICs

To create iSCSI vNICs, complete the following steps:

1. Select LAN on the left.

2. Select Policies > root.
3. Right-click vNIC Templates.
4. Select Create vNIC Template.
5. Enter Site- 01-iSCSI_A as the vNIC template name.
6. Select Fabric A. Do not select the Enable Failover option.
7. Leave Redundancy Type set at No Redundancy.
8. Under Target, make sure that only the Adapter option is selected.
9. Select Updating Template for Template Type.
10. Under VLANs, select only Site- 01-iSCSI_A_VLAN.
11. Select Site- 01-iSCSI_A_VLAN as the native VLAN.
12. Leave vNIC Name set for the CDN Source.
13. Under MTU, enter 9000.
14. From the MAC Pool list, select MAC-Pool-A.
15. From the Network Control Policy list, select Enable-CDP-LLDP.
16. Click OK to complete creating the vNIC template.
17. Click OK.

LAN / Policies / root / vNIC Templates / vNIC Template Site_01_iSCSI-A

General VLANs VLAN Groups Faults Events

Actions

- Modify VLANs
- Modify VLAN Groups
- Delete
- Show Policy Usage
- Use Global

Properties

Name : Site_01_iSCSI-A

Description :

Owner : Local

Fabric ID : ☒ Fabric A ☐ Fabric B ☐ Enable Failover

Redundancy :

Redundancy Type : ☒ No Redundancy ☐ Primary Template ☐ Secondary Template

Target

☒ Adapter ☐ VM

Template Type : ☐ Initial Template ☒ Updating Template

CDN Source : ☒ vNIC Name ☐ User Defined

MTU : 9000

Policies

MAC Pool : MAC_Pool_A(56/64)

QoS Policy : <not set>

Network Control Policy : Enable_CDP

Pin Group : <not set>

Stats Threshold Policy : default

Connection Policies

☒ Dynamic vNIC ☐ usNIC ☐ VMQ

Dynamic vNIC Connection Policy : <not set>

18. Select LAN on the left.
19. Select Policies > root.
20. Right-click vNIC Templates.
21. Select Create vNIC Template.
22. Enter Site- 01-iSCSI_B as the vNIC template name.
23. Select Fabric B. Do not select the Enable Failover option.
24. Leave Redundancy Type set at No Redundancy.
25. Under Target, make sure that only the Adapter option is selected.
26. Select Updating Template for Template Type.
27. Under VLANs, select only Site- 01-iSCSI_B_VLAN.
28. Select Site- 01-iSCSI_B_VLAN as the native VLAN.
29. Leave vNIC Name set for the CDN Source.
30. Under MTU, enter 9000.
31. From the MAC Pool list, select MAC-Pool-B.
32. From the Network Control Policy list, select Enable-CDP-LLDP.
33. Click OK to complete creating the vNIC template.
34. Click OK.

| General | VLANs | VLAN Groups | Faults | Events |
|---|-------|-------------|--------|--------|
| Actions Modify VNICs Modify VLAN Groups Delete Show Policy Usage Link Critical | | | | |
| Properties Name : Site_01_ISCSI-B Description : Owner : Local Fabric ID : <input type="radio"/> Fabric A <input checked="" type="radio"/> Fabric B <input type="checkbox"/> Enable Failover Redundancy Redundancy Type : <input checked="" type="radio"/> No Redundancy <input type="radio"/> Primary Template <input type="radio"/> Secondary Template Target <input checked="" type="checkbox"/> Adaptor <input type="checkbox"/> VM Template Type : <input type="radio"/> Initial Template <input checked="" type="radio"/> Updating Template CDN Source : <input checked="" type="radio"/> vNIC Name <input type="radio"/> User Defined MTU : 9000 Policies MAC Pool : MAC_Pool_B(50/64) QoS Policy : <not set> Network Control Policy : Enable_CDP Pin Group : <not set> Stats Threshold Policy : default Connection Policies <input checked="" type="radio"/> Dynamic vNIC <input type="radio"/> usNIC <input type="radio"/> VMQ Dynamic vNIC Connection Policy : <not set> | | | | |

Create LAN connectivity policy for iSCSI boot

This procedure applies to a Cisco UCS environment in which two iSCSI LIFs are on cluster node 1 (iscsi_lif01a and iscsi_lif01b) and two iSCSI LIFs are on cluster node 2 (iscsi_lif02a and iscsi_lif02b). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS 6324 A) and the B LIFs are connected to Fabric B (Cisco UCS 6324 B).

To configure the necessary Infrastructure LAN Connectivity Policy, complete the following steps:

1. In Cisco UCS Manager, click LAN on the left.
2. Select LAN > Policies > root.
3. Right-click LAN Connectivity Policies.
4. Select Create LAN Connectivity Policy.
5. Enter Site-XX-Fabric-A as the name of the policy.
6. Click the upper Add option to add a vNIC.
7. In the Create vNIC dialog box, enter Site-01-vNIC-A as the name of the vNIC.
8. Select the Use vNIC Template option.
9. In the vNIC Template list, select vNIC_Template_A.

10. From the Adapter Policy drop-down list, select VMWare.
11. Click OK to add this vNIC to the policy.

Modify vNIC

Name : **Site-01-vNIC-A**

Use vNIC Template : ☒

[Create vNIC Template](#)

vNIC Template : vNIC_Template_A ▼

Adapter Performance Profile

Adapter Policy : VMWare ▼

[Create Ethernet Adapter Policy](#)

[Create QoS Policy](#)

[Create Network Control Policy](#)

Connection Policies

☒ Dynamic vNIC ☐ usNIC ☐ VMQ

OK **Cancel**

12. Click the upper Add option to add a vNIC.
13. In the Create vNIC dialog box, enter Site-01-vNIC-B as the name of the vNIC.
14. Select the Use vNIC Template option.
15. In the vNIC Template list, select vNIC_Template_B.
16. From the Adapter Policy drop-down list, select VMWare.
17. Click OK to add this vNIC to the policy.
18. Click the upper Add option to add a vNIC.
19. In the Create vNIC dialog box, enter Site-01- iSCSI-A as the name of the vNIC.
20. Select the Use vNIC Template option.
21. In the vNIC Template list, select Site-01-iSCSI-A.
22. From the Adapter Policy drop-down list, select VMWare.
23. Click OK to add this vNIC to the policy.
24. Click the upper Add option to add a vNIC.

25. In the Create vNIC dialog box, enter `Site-01-iSCSI-B` as the name of the vNIC.
26. Select the Use vNIC Template option.
27. In the vNIC Template list, select `Site-01-iSCSI-B`.
28. From the Adapter Policy drop-down list, select VMWare.
29. Click OK to add this vNIC to the policy.
30. Expand the Add iSCSI vNICs option.
31. Click the Lower Add option in the Add iSCSI vNICs space to add the iSCSI vNIC.
32. In the Create iSCSI vNIC dialog box, enter `Site-01-iSCSI-A` as the name of the vNIC.
33. Select the Overlay vNIC as `Site-01-iSCSI-A`.
34. Leave the iSCSI Adapter Policy option to Not Set.
35. Select the VLAN as `Site-01-iSCSI-Site-A (native)`.
36. Select None (used by default) as the MAC address assignment.
37. Click OK to add the iSCSI vNIC to the policy.

Modify iSCSI vNIC ? ×

Name : **Site-01-ISCSI-A**

Overlay vNIC : Site-01-ISCSI-A ▼

iSCSI Adapter Policy : <not set> ▼ [Create iSCSI Adapter Policy](#)

VLAN : Site_01_ISCSI-A (native) ▼

iSCSI MAC Address

MAC Address Assignment: Select(None used by default)

[Create MAC Pool](#)

OK **Cancel**

38. Click the Lower Add option in the Add iSCSI vNICs space to add the iSCSI vNIC.
39. In the Create iSCSI vNIC dialog box, enter `Site-01-iSCSI-B` as the name of the vNIC.
40. Select the Overlay vNIC as `Site-01-iSCSI-B`.
41. Leave the iSCSI Adapter Policy option to Not Set.
42. Select the VLAN as `Site-01-iSCSI-Site-B (native)`.
43. Select None(used by default) as the MAC Address Assignment.
44. Click OK to add the iSCSI vNIC to the policy.
45. Click Save Changes.

LAN / Policies / root / LAN Connectivity Policies / Site01-SCSIBoot

General Events

Actions:
 Delete
 Show Policy Usage
 Live Details

Name: Site01-SCSIBoot
 Description:
 Owner: Local
 Click Add to specify one or more vNICs that the server should use to connect to the LAN.

| Name | MAC Address | Native VLAN |
|---------------------|-------------|-------------|
| vNIC Site-01-SCSI-A | Derived | |
| vNIC Site-01-SCSI-B | Derived | |
| vNIC Site-01-vNIC-A | Derived | |
| vNIC Site-01-vNIC-B | Derived | |

Filter: Add Modify

Add SCSI vNICs

| Name | Overlay vNIC Name | SCSI Adapter Policy | MAC Address |
|--------------------------|-------------------|---------------------|-------------|
| SCSI vNIC Site-01-SCSI-A | Site-01-SCSI-A | | Derived |
| SCSI vNIC Site-01-SCSI-B | Site-01-SCSI-B | | Derived |

Add Delete Modify

Create vMedia policy for VMware ESXi 6.7U1 install boot

In the NetApp Data ONTAP setup steps an HTTP web server is required, which is used for hosting NetApp Data ONTAP as well as VMware software. The vMedia Policy created here maps the VMware ESXi 6.7U1 ISO to the Cisco UCS server in order to boot the ESXi installation. To create this policy, complete the following steps:

1. In Cisco UCS Manager, select Servers on the left.
2. Select Policies > root.
3. Select vMedia Policies.
4. Click Add to create new vMedia Policy.
5. Name the policy ESXi-6.7U1-HTTP.
6. Enter Mounts ISO for ESXi 6.7U1 in the Description field.
7. Select Yes for Retry on Mount failure.
8. Click Add.
9. Name the mount ESXi-6.7U1-HTTP.
10. Select the CDD Device Type.
11. Select the HTTP Protocol.
12. Enter the IP Address of the web server.



The DNS server IPs were not entered into the KVM IP earlier, therefore, it is necessary to enter the IP of the web server instead of the hostname.

13. Enter VMware-VMvisor-Installer-6.7.0.update01-10302608.x86_64.iso as the Remote File name.

This VMware ESXi 6.7U1 ISO can be downloaded from [VMware Downloads](#).

14. Enter the web server path to the ISO file in the Remote Path field.

15. Click OK to create the vMedia Mount.
16. Click OK then OK again to complete creating the vMedia Policy.

For any new servers added to the Cisco UCS environment the vMedia service profile template can be used to install the ESXi host. On first boot, the host boots into the ESXi installer since the SAN mounted disk is empty. After ESXi is installed, the vMedia is not referenced as long as the boot disk is accessible.

The screenshot shows two overlapping dialog boxes in the Cisco UCS Manager interface. The background dialog is 'Create vMedia Policy' with fields for Name (ESXi-6.7U1-HTTP), Description (Mounts ISO for ESXi 6.7U1), and Retry on Mount Failure (Yes). The foreground dialog is 'Create vMedia Mount' with fields for Name (ESXi-6.7U1-HTTP), Description, Device Type (CDD), Protocol (HTTP), Hostname/IP Address (172.18.7.30), Image Name Variable (None), Remote File (VMware-VMvisor-Installer-6.7.0.update01-103026), Remote Path (http://172.18.7.30/seahawks/vSphere/), Username, Password, and Remap on Eject. Both dialogs have OK and Cancel buttons.

Create iSCSI boot policy

The procedure in this section applies to a Cisco UCS environment in which two iSCSI logical interfaces (LIFs) are on cluster node 1 (`iscsi_lif01a` and `iscsi_lif01b`) and two iSCSI LIFs are on cluster node 2 (`iscsi_lif02a` and `iscsi_lif02b`). Also, it is assumed that the A LIFs are connected to Fabric A (Cisco UCS Fabric Interconnect A) and the B LIFs are connected to Fabric B (Cisco UCS Fabric Interconnect B).




One boot policy is configured in this procedure. The policy configures the primary target to be `iscsi_lif01a`.

To create a boot policy for the Cisco UCS environment, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Policies > root.
3. Right-click Boot Policies.
4. Select Create Boot Policy.

5. Enter Site-01-Fabric-A as the name of the boot policy.
6. Optional: Enter a description for the boot policy.
7. Keep the Reboot on Boot Order Change option cleared.
8. Boot Mode is Legacy.
9. Expand the Local Devices drop-down menu and select Add Remote CD/DVD.
10. Expand the iSCSI vNICs drop-down menu and select Add iSCSI Boot.
11. In the Add iSCSI Boot dialog box, enter Site-01-iSCSI-A. Click OK.
12. Select Add iSCSI Boot.
13. In the Add iSCSI Boot dialog box, enter Site-01-iSCSI-B. Click OK.
14. Click OK to create the policy.



Create service profile template

In this procedure, one service profile template for Infrastructure ESXi hosts is created for Fabric A boot.

To create the service profile template, complete the following steps:

1. In Cisco UCS Manager, click Servers on the left.
2. Select Service Profile Templates > root.
3. Right-click root.
4. Select Create Service Profile Template to open the Create Service Profile Template wizard.
5. Enter VM-Host-Infra-iSCSI-A as the name of the service profile template. This service profile template is configured to boot from storage node 1 on fabric A.
6. Select the Updating Template option.
7. Under UUID, select UUID_Pool as the UUID pool. Click Next.

Create Service Profile Template

You must enter a name for the service profile template and specify the template type. You can also specify how a UUID will be assigned to the template and enter a description.

Name:

The template will be created in the following organization. Its name must be unique within this organization.
Where:

The template will be created in the following organization. Its name must be unique within this organization.
Type:

Specify how the UUID will be assigned to the server associated with the service generated by the template.
UUID Assignment:

The UUID will be assigned from the selected pool.
The available/total UUIDs are displayed after the pool name.

Optionally enter a description for the profile. The description can contain information about when and where the service profile should be used.

Configure storage provisioning

To configure storage provisioning, complete the following steps:

1. If you have servers with no physical disks, click Local Disk Configuration Policy and select the SAN Boot Local Storage Policy. Otherwise, select the default Local Storage Policy.
2. Click Next.

Configure networking options

To configure the networking options, complete the following steps:

1. Keep the default setting for Dynamic vNIC Connection Policy.
2. Select the Use Connectivity Policy option to configure the LAN connectivity.
3. Select iSCSI-Boot from the LAN Connectivity Policy drop-down menu.
4. Select IQN_Pool in Initiator Name Assignment. Click Next.

Create Service Profile Template

Optionally specify LAN configuration information.

Dynamic vNIC Connection Policy: Select a Policy to use (no Dynamic vNIC Policy by default) ▼

[Create Dynamic vNIC Connection Policy](#)

How would you like to configure LAN connectivity?

☐ Simple
 ☐ Expert
 ☐ No vNICs
 ☒ Use Connectivity Policy

LAN Connectivity Policy: Site01 - iSCSIBoot ▼ [Create LAN Connectivity Policy](#)

Initiator Name

Initiator Name Assignment: IQN Pool(60/64) ▼

Initiator Name: [Create IQN Suffix Pool](#)

The IQN will be assigned from the selected pool.
The available/total IQNs are displayed after the pool name.

< Prev Next > **Finish** Cancel

Configure SAN connectivity

To configure SAN connectivity, complete the following steps:

1. For the vHBAs, select No for the How Would you Like to Configure SAN Connectivity? option.
2. Click Next.

Configure zoning

To configure zoning, simply click Next.

Configure vNIC/HBA placement

To configure vNIC/HBA placement, complete the following steps:

1. From the Select Placement drop-down list, leave the placement policy as Let System Perform Placement.
2. Click Next.

Configure vMedia policy

To configure the vMedia policy, complete the following steps:

1. Do not select a vMedia Policy.
2. Click Next.

Configure server boot order

To configure the server boot order, complete the following steps:

1. Select Boot-Fabric-A for Boot Policy.

Create Service Profile Template

Optionally specify the boot policy for this service profile template.

Select a boot policy.

Boot Policy: **Site-01-Fabric-A** [Create Boot Policy](#)

Name: **Site-01-Fabric-A**
Description:
Reboot on Boot Order Change: **No**
Enforce vNIC/vHBA/iSCSI Name: **Yes**
Boot Mode: **Legacy**

WARNINGS:
The type (primary/secondary) does not indicate a boot order presence.
The effective order of boot devices within the same device class (LAN/Storage/iSCSI) is determined by PCIe bus scan order.
If **Enforce vNIC/vHBA/iSCSI Name** is selected and the vNIC/vHBA/iSCSI does not exist, a config error will be reported.
If it is not selected, the vNICs/vHBAs are selected if they exist, otherwise the vNIC/vHBA with the lowest PCIe bus scan order is used.

Boot Order

| Name | Order | vNIC/vHBA/iSCSI | vNIC | Type | LUN Na... | WWN | Slot Nu... | Boot Na... | Boot Path | Descripti... |
|-----------|-------|-----------------|------|-----------|-----------|-----|------------|------------|-----------|--------------|
| Reboot... | 1 | | | | | | | | | |
| ▼ iSCSI | 2 | | | | | | | | | |
| iSCSI | | Site-01-iSCSI-A | | Primary | | | | | | |
| iSCSI | | Site-01-iSCSI-B | | Second... | | | | | | |

[Set iSCSI Boot Parameters](#) [Set iSCSI Boot Parameters](#) [Set iSCSI Boot Parameters](#)

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

2. In the Boot order, select Site-01- iSCSI-A.
3. Click Set iSCSI Boot Parameters.
4. In the Set iSCSI Boot Parameters dialog box, leave the Authentication Profile option to Not Set unless you have independently created one appropriate for your environment.
5. Leave the Initiator Name Assignment dialog box Not Set to use the single Service Profile Initiator Name defined in the previous steps.
6. Set iSCSI_IP_Pool_A as the Initiator IP address Policy.
7. Select iSCSI Static Target Interface option.
8. Click Add.
9. Enter the iSCSI target name. To get the iSCSI target name of Infra-SVM, log in into storage cluster management interface and run the `iscsi show` command.

```
bb04-aff300::> iscsi show
Target                Target                Status
Vserver Name          Alias                Admin
-----
Infra-SVM iqn.1992-08.com.netapp:sn.b5acab9ef1c811e68d9d00a098a9fec2:vs.3
                        Infra-SVM                up
```

10. Enter the IP address of `iscsi_lif_02a` for the IPv4 Address field.

Create iSCSI Static Target ? ×

iSCSI Target Name :

Priority :

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

11. Click OK to add the iSCSI static target.
12. Click Add.
13. Enter the iSCSI target name.
14. Enter the IP address of `iscsi_lif_01a` for the IPv4 Address field.

Create iSCSI Static Target ? ×

iSCSI Target Name :

Priority :

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

15. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters

Name : **iSCSI-A-vNIC**

Authentication Profile : **<not set>** [Create iSCSI Authentication Profile](#)

Initiator Name

Initiator Name Assignment : **<not set>**

[Create IQN Suffix Pool](#)

WARNING: The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy : **iSCSI_IP_Pool_A(12/16)**

IPv4 Address : **0.0.0.0**
 Subnet Mask : **255.255.255.0**
 Default Gateway : **0.0.0.0**
 Primary DNS : **0.0.0.0**
 Secondary DNS : **0.0.0.0**

[Create IP Pool](#)
[Reset Initiator Address](#)
 The IP address will be automatically assigned from the selected pool.

☒ iSCSI Static Target Interface ☐ iSCSI Auto Target Interface

| Name | Priority | Port | Authentication Pro. | iSCSI IPv4 Address | LUN id |
|------------------|----------|------|---------------------|--------------------|--------|
| iqn.1992-08.c... | 1 | 3260 | | 192.168.10.62 | 0 |
| iqn.1992-08.c... | 2 | 3260 | | 192.168.10.61 | 0 |

OK **Cancel**



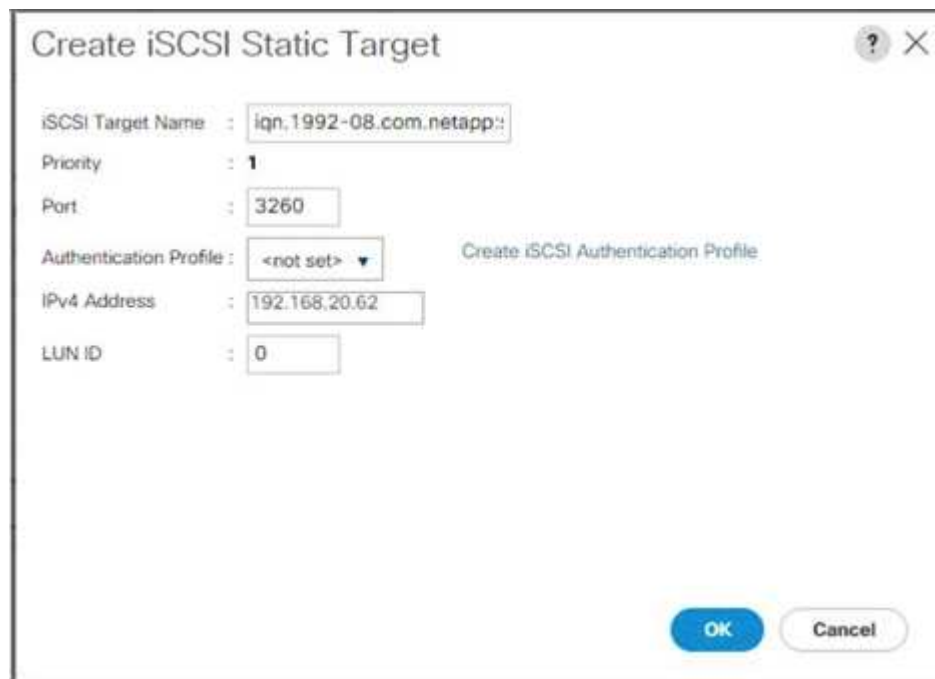
The target IPs were put in with the storage node 02 IP first and the storage node 01 IP second. This is assuming the boot LUN is on node 01. The host boots by using the path to node 01 if the order in this procedure is used.

16. In the Boot order, select iSCSI-B-vNIC.
17. Click Set iSCSI Boot Parameters.
18. In the Set iSCSI Boot Parameters dialog box, leave the Authentication Profile option as Not Set unless you have independently created one appropriate to your environment.
19. Leave the Initiator Name Assignment dialog box Not Set to use the single Service Profile Initiator Name defined in the previous steps.
20. Set `iSCSI_IP_Pool_B` as the initiator IP address policy.
21. Select the iSCSI Static Target Interface option.
22. Click Add.
23. Enter the iSCSI target name. To get the iSCSI target name of Infra-SVM, log in into storage cluster management interface and run the `iscsi show` command.

```
bb04-aff300::> iscsi show
```

| Vserver | Target Name | Target Alias | Status Admin |
|-----------|---|--------------|--------------|
| Infra-SVM | iqn.1992-08.com.netapp:sn.b5acab9ef1c811e68d9d00a098a9fec2:vs.3 | Infra-SVM | up |

24. Enter the IP address of `iscsi_lif_02b` for the IPv4 Address field.



Create iSCSI Static Target

iSCSI Target Name :

Priority :

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

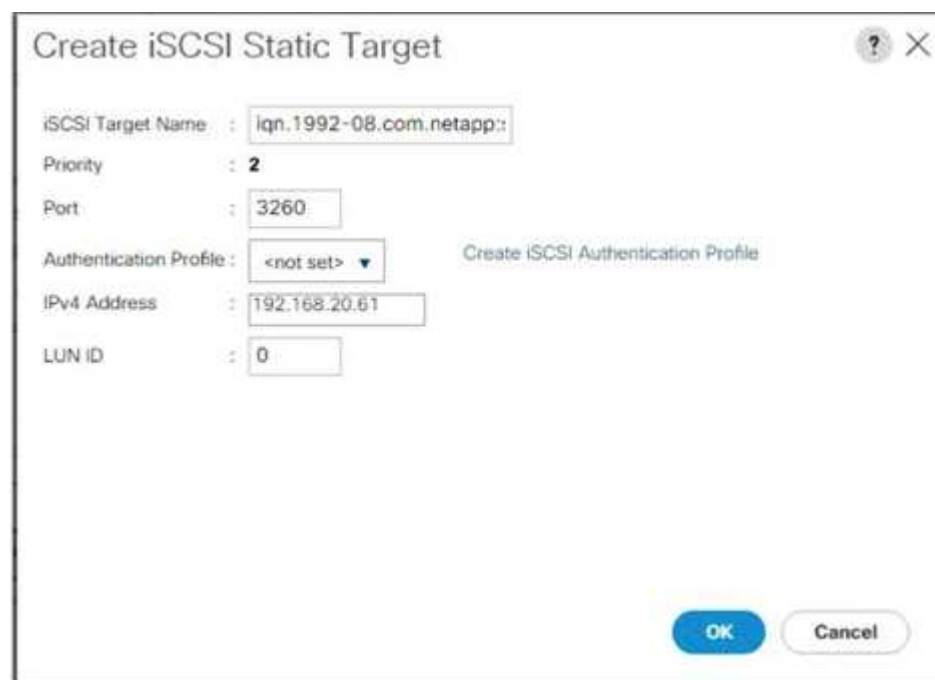
LUN ID :

25. Click OK to add the iSCSI static target.

26. Click Add.

27. Enter the iSCSI target name.

28. Enter the IP address of `iscsi_lif_01b` for the IPv4 Address field.



Create iSCSI Static Target

iSCSI Target Name :

Priority :

Port :

Authentication Profile : [Create iSCSI Authentication Profile](#)

IPv4 Address :

LUN ID :

29. Click OK to add the iSCSI static target.

Set iSCSI Boot Parameters

Create IQN Suffix Pool

WARNING: The selected pool does not contain any available entities. You can select it, but it is recommended that you add entities to it.

Initiator Address

Initiator IP Address Policy: iSCSI_IP_Pool_B(12/16)

IPv4 Address : 0.0.0.0

Subnet Mask : 255.255.255.0

Default Gateway : 0.0.0.0

Primary DNS : 0.0.0.0

Secondary DNS : 0.0.0.0

Create IP Pool

Reset Initiator Address

The IP address will be automatically assigned from the selected pool.

☒ iSCSI Static Target Interface

☐ iSCSI Auto Target Interface

| Name | Priority | Port | Authentication Pro.. | iSCSI IPv4 Address | LUN Id |
|------------------|----------|------|----------------------|--------------------|--------|
| iqn.1992-08.c... | 1 | 3260 | | 192.168.20.62 | 0 |
| iqn.1992-08.c... | 2 | 3260 | | 192.168.20.61 | 0 |

Add

Delete

Info

Minimum one instance of iSCSI Static Target Interface and maximum two are allowed.

OK

Cancel

30. Click Next.

Configure maintenance policy

To configure the maintenance policy, complete the following steps:

- 1. Change the maintenance policy to default.

Create Service Profile Template

Specify how disruptive changes such as reboots, network interruptions, and firmware upgrades should be applied to the server associated with this service profile.

Maintenance Policy

Select a maintenance policy to include with this service profile or create a new maintenance policy that will be accessible to all service profiles.

Maintenance Policy: default [Create Maintenance Policy](#)

| | |
|---------------------|----------|
| Name | default |
| Description | |
| Soft Shutdown Timer | 150 Secs |
| Reboot Policy | User Ack |

[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

2. Click Next.

Configure server assignment

To configure the server assignment, complete the following steps:

1. In the Pool Assignment list, select Infra-Pool.
2. Select Down as the power state to be applied when the profile is associated with the server.
3. Expand Firmware Management at the bottom of the page and select the default policy.

Create Service Profile Template

Optionally specify a server pool for this service profile template.

You can select a server pool you want to associate with this service profile template.

Pool Assignment: Infra-Pool [Create Server Pool](#)

Select the power state to be applied when this profile is associated with the server.

☐ Up ☒ Down

The service profile template will be associated with one of the servers in the selected pool. If desired, you can specify an additional server pool policy qualification that the selected server must meet. To do so, select the qualification from the list.

Server Pool Qualification: <not set>

Restrict Migration: ☒

Firmware Management (BIOS, Disk Controller, Adapter)

If you select a host firmware policy for this service profile, the profile will update the firmware on the server that it is associated with. Otherwise the system uses the firmware already installed on the associated server.

Host Firmware Package: default [Create Host Firmware Package](#)

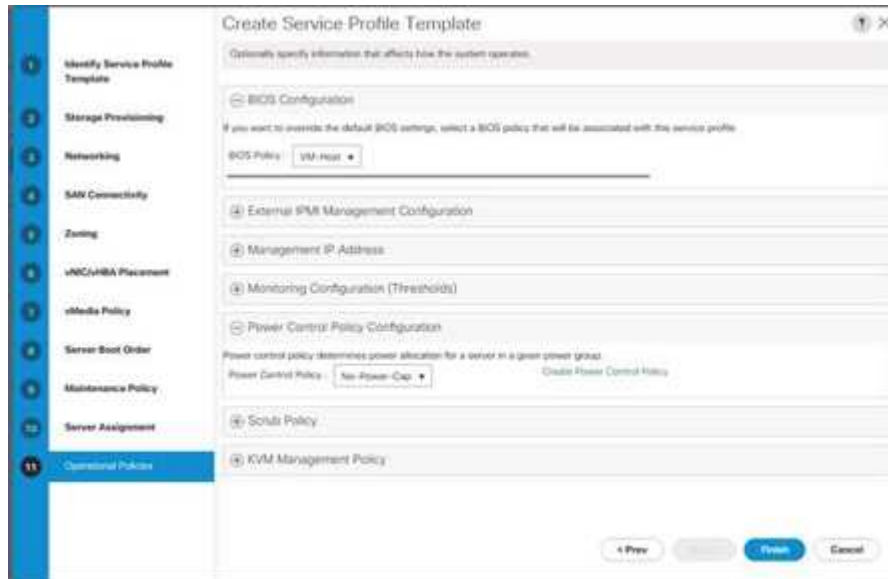
[< Prev](#) [Next >](#) [Finish](#) [Cancel](#)

4. Click Next.

Configure operational policies

To configure the operational policies, complete the following steps:

1. From the BIOS Policy drop-down list, select VM-Host.
2. Expand Power Control Policy Configuration and select No-Power-Cap from the Power Control Policy drop-down list.



3. Click Finish to create the service profile template.
4. Click OK in the confirmation message.

Create vMedia-enabled service profile template

To create a service profile template with vMedia enabled, complete the following steps:

1. Connect to UCS Manager and click Servers on the left.
2. Select Service Profile Templates > root > Service Template VM-Host-Infra-iSCSI-A.
3. Right-click VM-Host-Infra-iSCSI-A and select Create a Clone.
4. Name the clone VM-Host-Infra-iSCSI-A-vM.
5. Select the newly created VM-Host-Infra-iSCSI-A-vM and select the vMedia Policy tab on the right.
6. Click Modify vMedia Policy.
7. Select the ESXi-6.7U1-HTTP vMedia Policy and click OK.
8. Click OK to confirm.

Create service profiles

To create service profiles from the service profile template, complete the following steps:

1. Connect to Cisco UCS Manager and click Servers on the left.
2. Expand Servers > Service Profile Templates > root > Service Template <name>.

3. In Actions, click Create Service Profile from Template and complete the following steps:
 - a. Enter Site- 01-Infra-0 as the naming prefix.
 - b. Enter 2 as the number of instances to create.
 - c. Select root as the org.
 - d. Click OK to create the service profiles.



4. Click OK in the confirmation message.
5. Verify that the service profiles Site-01-Infra-01 and Site-01-Infra-02 have been created.



The service profiles are automatically associated with the servers in their assigned server pools.

Storage configuration part 2: boot LUNs and initiator groups

ONTAP boot storage setup

Create initiator groups

To create initiator groups (igroups), complete the following steps:

1. Run the following commands from the cluster management node SSH connection:

```
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-01 -protocol
iscsi -ostype vmware -initiator <vm-host-infra-01-iqn>
igroup create -vserver Infra-SVM -igroup VM-Host-Infra-02 -protocol
iscsi -ostype vmware -initiator <vm-host-infra-02-iqn>
igroup create -vserver Infra-SVM -igroup MGMT-Hosts -protocol iscsi
-ostype vmware -initiator <vm-host-infra-01-iqn>, <vm-host-infra-02-iqn>
```



Use the values listed in Table 1 and Table 2 for the IQN information.

2. To view the three igroups just created, run the `igroup show` command.

Map boot LUNs to igroups

To map boot LUNs to igroups, complete the following step:

1. From the storage cluster management SSH connection, run the following commands:

```
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra- A  
-igroup VM-Host-Infra-01 -lun-id 0  
lun map -vserver Infra-SVM -volume esxi_boot -lun VM-Host-Infra- B -igroup VM-Host-Infra-02 -lun-id 0
```

VMware vSphere 6.7U1 deployment procedure

This section provides detailed procedures for installing VMware ESXi 6.7U1 in a FlexPod Express configuration. After the procedures are completed, two booted ESXi hosts are provisioned.

Several methods exist for installing ESXi in a VMware environment. These procedures focus on how to use the built-in KVM console and virtual media features in Cisco UCS Manager to map remote installation media to individual servers and connect to their boot LUNs.

Download Cisco custom image for ESXi 6.7U1

If the VMware ESXi custom image has not been downloaded, complete the following steps to complete the download:

1. Click the following xref:./express/ [VMware vSphere Hypervisor \(ESXi\) 6.7U1](#).
2. You need a user ID and password on [vmware.com](#) to download this software.
3. Download the `.iso` file.

Cisco UCS Manager

The Cisco UCS IP KVM enables the administrator to begin the installation of the OS through remote media. It is necessary to log in to the Cisco UCS environment to run the IP KVM.

To log in to the Cisco UCS environment, complete the following steps:

1. Open a web browser and enter the IP address for the Cisco UCS cluster address. This step launches the Cisco UCS Manager application.
2. Click the Launch UCS Manager link under HTML to launch the HTML 5 UCS Manager GUI.
3. If prompted to accept security certificates, accept as necessary.
4. When prompted, enter `admin` as the user name and enter the administrative password.
5. To log in to Cisco UCS Manager, click Login.
6. From the main menu, click Servers on the left.
7. Select Servers > Service Profiles > root > VM-Host-Infra-01.
8. Right-click VM-Host-Infra-01 and select KVM Console.

9. Follow the prompts to launch the Java-based KVM console.
10. Select Servers > Service Profiles > root > VM-Host-Infra-02.
11. Right-click VM-Host-Infra-02. and select KVM Console.
12. Follow the prompts to launch the Java-based KVM console.

Set up VMware ESXi installation

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To prepare the server for the OS installation, complete the following steps on each ESXi host:

1. In the KVM window, click Virtual Media.
2. Click Activate Virtual Devices.
3. If prompted to accept an Unencrypted KVM session, accept as necessary.
4. Click Virtual Media and select Map CD/DVD.
5. Browse to the ESXi installer ISO image file and click Open.
6. Click Map Device.
7. Click the KVM tab to monitor the server boot.

Install ESXi

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To install VMware ESXi to the iSCSI-bootable LUN of the hosts, complete the following steps on each host:

1. Boot the server by selecting Boot Server and clicking OK. Then click OK again.
2. On reboot, the machine detects the presence of the ESXi installation media. Select the ESXi installer from the boot menu that is displayed.
3. After the installer is finished loading, press Enter to continue with the installation.
4. Read and accept the end-user license agreement (EULA). Press F11 to accept and continue.
5. Select the LUN that was previously set up as the installation disk for ESXi and press Enter to continue with the installation.
6. Select the appropriate keyboard layout and press Enter.
7. Enter and confirm the root password and press Enter.
8. The installer issues a warning that the selected disk will be repartitioned. Press F11 to continue with the installation.
9. After the installation is complete, select the Virtual Media tab and clear the P mark next to the ESXi installation media. Click Yes.



The ESXi installation image must be unmapped to make sure that the server reboots into ESXi and not into the installer.

10. After the installation is complete, press Enter to reboot the server.
11. In Cisco UCS Manager, bind the current service profile to the non-vMedia service profile template to prevent mounting the ESXi installation iso over HTTP.


Set up management networking for ESXi hosts

Adding a management network for each VMware host is necessary for managing the host. To add a management network for the VMware hosts, complete the following steps on each ESXi host:

ESXi Host VM-Host-Infra-01 and VM-Host-Infra-02

To configure each ESXi host with access to the management network, complete the following steps:

1. After the server has finished rebooting, press F2 to customize the system.
2. Log in as `root`, enter the corresponding password, and press Enter to log in.
3. Select Troubleshooting Options and press Enter.
4. Select Enable ESXi Shell and press Enter.
5. Select Enable SSH and press Enter.
6. Press Esc to exit the Troubleshooting Options menu.
7. Select the Configure Management Network option and press Enter.
8. Select Network Adapters and press Enter.
9. Verify that the numbers in the Hardware Label field match the numbers in the Device Name field.
10. Press Enter.



11. Select the VLAN (Optional) option and press Enter.
12. Enter the `<ib-mgmt-vlan-id>` and press Enter.
13. Select IPv4 Configuration and press Enter.
14. Select the Set Static IPv4 Address and Network Configuration option by using the space bar.
15. Enter the IP address for managing the first ESXi host.
16. Enter the subnet mask for the first ESXi host.

17. Enter the default gateway for the first ESXi host.
18. Press Enter to accept the changes to the IP configuration.
19. Select the DNS Configuration option and press Enter.



Because the IP address is assigned manually, the DNS information must also be entered manually.

20. Enter the IP address of the primary DNS server.
21. Optional: Enter the IP address of the secondary DNS server.
22. Enter the FQDN for the first ESXi host.
23. Press Enter to accept the changes to the DNS configuration.
24. Press Esc to exit the Configure Management Network menu.
25. Select Test Management Network to verify that the management network is set up correctly and press Enter.
26. Press Enter to run the test, press Enter again once the test has completed, review environment if there is a failure.
27. Select the Configure Management Network again and press Enter.
28. Select the IPv6 Configuration option and press Enter.
29. Using the spacebar, select Disable IPv6 (restart required) and press Enter.
30. Press Esc to exit the Configure Management Network submenu.
31. Press Y to confirm the changes and reboot the ESXi host.

Reset VMware ESXi host VMkernel port vmk0 MAC address (optional)

ESXi Host VM-Host-Infra-01 and VM-Host-Infra-02

By default, the MAC address of the management VMkernel port vmk0 is the same as the MAC address of the Ethernet port on which it is placed. If the ESXi host's boot LUN is remapped to a different server with different MAC addresses, a MAC address conflict will occur because vmk0 retains the assigned MAC address unless the ESXi system configuration is reset. To reset the MAC address of vmk0 to a random VMware-assigned MAC address, complete the following steps:

1. From the ESXi console menu main screen, press Ctrl-Alt-F1 to access the VMware console command line interface. In the UCSM KVM, Ctrl-Alt-F1 appears in the list of static macros.
2. Log in as root.
3. Type `esxcfg-vmknic -l` to get a detailed listing of interface vmk0. vmk0 should be a part of the Management Network port group. Note the IP address and netmask of vmk0.
4. To remove vmk0, enter the following command:

```
esxcfg-vmknic -d "Management Network"
```

5. To add vmk0 again with a random MAC address, enter the following command:

```
esxcfg-vmknics -a -i <vmk0-ip> -n <vmk0-netmask> "Management Network".
```

6. Verify that vmk0 has been added again with a random MAC address

```
esxcfg-vmknics -l
```

7. Type `exit` to log out of the command line interface.
8. Press Ctrl-Alt-F2 to return to the ESXi console menu interface.

Log into VMware ESXi hosts with VMware host client

ESXi Host VM-Host-Infra-01

To log in to the VM-Host-Infra-01 ESXi host by using the VMware Host Client, complete the following steps:

1. Open a web browser on the management workstation and navigate to the `VM-Host-Infra-01` management IP address.
2. Click Open the VMware Host Client.
3. Enter `root` for the user name.
4. Enter the root password.
5. Click Login to connect.
6. Repeat this process to log in to `VM-Host-Infra-02` in a separate browser tab or window.

Install VMware drivers for the Cisco Virtual Interface Card (VIC)

Download and extract the offline bundle for the following VMware VIC driver to the Management workstation:

- `nenic Driver version 1.0.25.0`

ESXi hosts VM-Host-Infra-01 and VM-Host-Infra-02

To install VMware VIC Drivers on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02, complete the following steps:

1. From each Host Client, select Storage.
2. Right-click `datastore1` and select Browse.
3. In the Datastore browser, click Upload.
4. Navigate to the saved location for the downloaded VIC drivers and select `VMW-ESX-6.7.0-nenic-1.0.25.0-offline_bundle-11271332.zip`.
5. In the Datastore browser, click Upload.
6. Click Open to upload the file to `datastore1`.
7. Make sure the file has been uploaded to both ESXi hosts.
8. Place each host into Maintenance mode if it isn't already.
9. Connect to each ESXi host through ssh from a shell connection or putty terminal.

10. Log in as root with the root password.
11. Run the following commands on each host:

```
esxcli software vib update -d /vmfs/volumes/datastore1/VMW-ESX-6.7.0-
nenic-1.0.25.0-offline_bundle-11271332.zip
reboot
```

12. Log into the Host Client on each host once reboot is complete and exit Maintenance Mode.

Set up VMkernel ports and virtual switch

ESXi Host VM-Host-Infra-01 and VM-Host-Infra-02

To set up the VMkernel ports and the virtual switches on the ESXi hosts, complete the following steps:

1. From the Host Client, select Networking on the left.
2. In the center pane, select the Virtual switches tab.
3. Select vSwitch0.
4. Select Edit settings.
5. Change the MTU to 9000.
6. Expand NIC teaming.
7. In the Failover order section, select vmnic1 and click Mark active.
8. Verify that vmnic1 now has a status of Active.
9. Click Save.
10. Select Networking on the left.
11. In the center pane, select the Virtual switches tab.
12. Select iScsiBootvSwitch.
13. Select Edit settings.
14. Change the MTU to 9000
15. Click Save.
16. Select the VMkernel NICs tab.
17. Select vmk1 iScsiBootPG.
18. Select Edit settings.
19. Change the MTU to 9000.
20. Expand IPv4 settings and change the IP address to an address outside of the UCS iSCSI-IP-Pool-A.



To avoid IP address conflicts if the Cisco UCS iSCSI IP Pool addresses should get reassigned, it is recommended to use different IP addresses in the same subnet for the iSCSI VMkernel ports.

21. Click Save.
22. Select the Virtual switches tab.

23. Select the Add standard virtual switch.
24. Provide a name of iScsiBootvSwitch-B for the vSwitch Name.
25. Set the MTU to 9000.
26. Select vmnic3 from the Uplink 1 drop-down menu.
27. Click Add.
28. In the center pane, select the VMkernel NICs tab.
29. Select Add VMkernel NIC
30. Specify a New port group name of iScsiBootPG-B.
31. Select iScsiBootvSwitch-B for Virtual switch.
32. Set the MTU to 9000. Do not enter a VLAN ID.
33. Select Static for the IPv4 settings and expand the option to provide the Address and Subnet Mask within the Configuration.



To avoid IP address conflicts, if the Cisco UCS iSCSI IP Pool addresses should get reassigned, it is recommended to use different IP addresses in the same subnet for the iSCSI VMkernel ports.

34. Click Create.
35. On the left, select Networking, then select the Port groups tab.
36. In the center pane, right-click VM Network and select Remove.
37. Click Remove to complete removing the port group.
38. In the center pane, select Add port group.
39. Name the port group Management Network and enter <ib-mgmt-vlan-id> in the VLAN ID field, and make sure Virtual switch vSwitch0 is selected.
40. Click Add to finalize the edits for the IB-MGMT Network.
41. At the top, select the VMkernel NICs tab.
42. Click Add VMkernel NIC.
43. For New port group, enter VMotion.
44. For Virtual switch, select vSwitch0 selected.
45. Enter <vmotion-vlan-id> for the VLAN ID.
46. Change the MTU to 9000.
47. Select Static IPv4 settings and expand IPv4 settings.
48. Enter the ESXi host vMotion IP address and netmask.
49. Select the vMotion stack TCP/IP stack.
50. Select vMotion under Services.
51. Click Create.
52. Click Add VMkernel NIC.
53. For New port group, enter NFS_Share.
54. For Virtual switch, select vSwitch0 selected.

55. Enter <infra-nfs-vlan-id> for the VLAN ID
56. Change the MTU to 9000.
57. Select Static IPv4 settings and expand IPv4 settings.
58. Enter the ESXi host Infrastructure NFS IP address and netmask.
59. Do not select any of the Services.
60. Click Create.
61. Select the Virtual Switches tab, then select vSwitch0. The properties for vSwitch0 VMkernel NICs should be similar to the following example:

vSwitch0

Type: Standard vSwitch
Port groups: 4
Uplinks: 2

| vSwitch Details | |
|-----------------|---|
| MTU | 9000 |
| Ports | 8816 (8798 available) |
| Link discovery | Listen + Cisco discovery protocol (CDP) |
| Attached VMs | 2 (1 active) |
| Beacon interval | 1 |

| NIC teaming policy | |
|--------------------|------------------------------------|
| Notify switches | Yes |
| Policy | Route based on originating port ID |
| Reverse policy | Yes |
| Fallback | Yes |

| Security policy | |
|------------------------|-----|
| Allow promiscuous mode | No |
| Allow forged transmits | Yes |
| Allow MAC changes | Yes |

| Shapping policy | |
|-----------------|----|
| Enabled | No |

vSwitch topology

- VM Network** (VLAN ID: 18)
 - Virtual Machines (2): vCenterServerApp-01 (MAC Address 00:0c:29:27:48:61), Linux-VM
- VMotion** (VLAN ID: 103)
 - VMkernel ports (1): vmk4: 192.168.103.208
- NFS_Share** (VLAN ID: 104)
 - VMkernel ports (1): vmk3: 192.168.104.208
- Management network** (VLAN ID: 18)
 - VMkernel ports (1): vmk2: 172.18.7.208
- Physical adapters**
 - vmnic1: 10000 Mbps, Full
 - vmnic0: 10000 Mbps, Full

62. Select the VMkernel NICs tab to confirm the configured virtual adapters. The adapters listed should be similar to the following example:

localhost.localdomain - Networking

Port groups

Virtual switches

Physical NICs

VMkernel NICs

TCP/IP stacks

Firewall rules

Add VMkernel NIC

Edit settings

Refresh

Actions

Search

| Name | Portgroup | TCP/IP stack | Services | IPv4 ad... | IPv6 addresses |
|------|--------------------|----------------------|------------|--------------|-----------------------------|
| vmk0 | Management Network | Default TCP/IP stack | Management | 172.18.7.... | fe80::225:b5ff:fe00:a2e/64 |
| vmk1 | iScsiBootPG | Default TCP/IP stack | | 192.168.... | fe80::225:b5ff:fe00:a3e/64 |
| vmk2 | iScsiBootPG-B | Default TCP/IP stack | | 192.168.... | fe80::250:56ff:fe64:1248... |
| vmk3 | NFS_Share | Default TCP/IP stack | | 192.168.... | fe80::250:56ff:fe65:29a4... |
| vmk4 | VMotion | Default TCP/IP stack | vMotion | 192.168.... | fe80::250:56ff:fe6c:2650... |

5 items

Setup iSCSI multipathing

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To set up the iSCSI multipathing on the ESXi host VM-Host-Infra-01 and VM-Host-Infra-02, complete the following steps:

1. From each Host Client, select Storage on the left.
2. In the center pane, click Adapters.
3. Select the iSCSI software adapter and click Configure iSCSI.


localhost.localdomain - Storage


Datastores


Adapters


Devices


Persistent Memory


 Configure iSCSI



 Software iSCSI

 Rescan

 Refresh

 Actions

 Search

| Name | Model | Status | Driver |
|---|--------------------------------|---------|-----------|
|  vmhba0 | Lewisburg SATA AHCI Controller | Unknown | vmw_ahci |
|  vmhba64 | iSCSI Software Adapter | Online | iscsi_vmk |

2 Items



vmhba64

Model
Driver

iSCSI Software Adapter
iscsi_vmk

4. Under Dynamic targets, click Add dynamic target.
5. Enter the IP Address of `iscsi_lif01a`.
6. Repeat entering these IP addresses: `iscsi_lif01b`, `iscsi_lif02a`, and `iscsi_lif02b`.
7. Click Save Configuration.

Configure iSCSI - vmhba64

iSCSI enabled: ☐ Disabled ☒ Enabled

Name & alias: iqn.1992-08.com.cisco:ucs-host:3

CHAP authentication: Do not use CHAP

Mutual CHAP authentication: Do not use CHAP

Advanced settings: Click to expand

Network port bindings:

Add port binding Remove port binding

VMkernel NIC Port group IPv4 address

No port bindings

Static targets:

Add static target Remove static target Edit settings Search

| Target | Address | Port |
|---------------------------------------|---------------|------|
| iqn.1992-08.com.netapp:sn.aff300:vs.3 | 192.168.124.3 | 3260 |
| iqn.1992-08.com.netapp:sn.aff300:vs.3 | 192.168.124.1 | 3260 |
| iqn.1992-08.com.netapp:sn.aff300:vs.3 | 192.168.125.3 | 3260 |
| iqn.1992-08.com.netapp:sn.aff300:vs.3 | 192.168.125.1 | 3260 |

Dynamic targets:

Add dynamic target Remove dynamic target Edit settings Search

| Address | Port |
|---------------|------|
| 192.168.124.1 | 3260 |
| 192.168.125.1 | 3260 |
| 192.168.125.3 | 3260 |

Save configuration Cancel

To obtain all of the `iscsi_lif` IP addresses, log in to NetApp storage cluster management interface and run the `network interface show` command.




The host automatically rescans the storage adapter and the targets are added to static targets.

Mount required datastores


ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To mount the required datastores, complete the following steps on each ESXi host:

1. From the Host Client, select Storage on the left.
2. In the center pane, select Datastores.
3. In the center pane, select New Datastore to add a new datastore.
4. In the New datastore dialog box, select Mount NFS datastore and click Next.

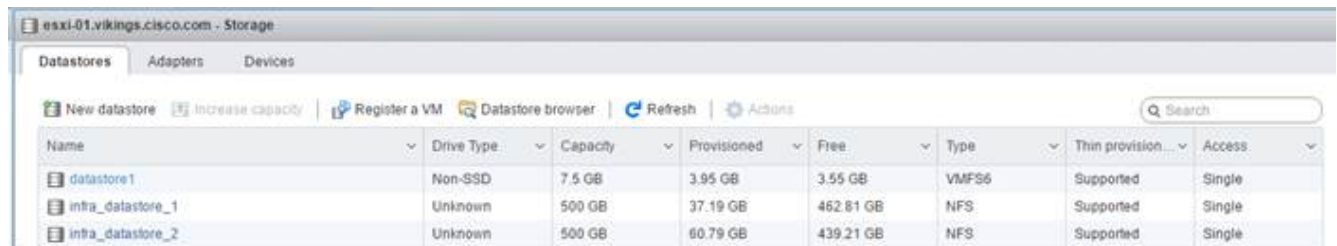


5. On the provide NFS Mount Details page, complete these steps:
 - a. Enter `infra_datastore_1` for the datastore name.
 - b. Enter the IP address for the `nfs_lif01_a` LIF for the NFS server.
 - c. Enter `/infra_datastore_1` for the NFS share.
 - d. Leave the NFS version set at NFS 3.
 - e. Click Next.



6. Click Finish. The datastore should now appear in the datastore list.
7. In the center pane, select New Datastore to add a new datastore.
8. In the New Datastore dialog box, select Mount NFS Datastore and click Next.
9. On the provide NFS Mount Details page, complete these steps:

- a. Enter `infra_datastore_2` for the datastore name.
 - b. Enter the IP address for the `nfs_lif02_a` LIF for the NFS server.
 - c. Enter `/infra_datastore_2` for the NFS share.
 - d. Leave the NFS version set at NFS 3.
 - e. Click Next.
10. Click Finish. The datastore should now appear in the datastore list.



| Name | Drive Type | Capacity | Provisioned | Free | Type | Thin provision... | Access |
|-------------------|------------|----------|-------------|-----------|-------|-------------------|--------|
| datastore1 | Non-SSD | 7.5 GB | 3.95 GB | 3.55 GB | VMFS6 | Supported | Single |
| infra_datastore_1 | Unknown | 500 GB | 37.19 GB | 462.81 GB | NFS | Supported | Single |
| infra_datastore_2 | Unknown | 500 GB | 60.79 GB | 439.21 GB | NFS | Supported | Single |

11. Mount both datastores on both ESXi hosts.

Configure NTP on ESXi hosts

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To configure NTP on the ESXi hosts, complete the following steps on each host:

1. From the Host Client, select Manage on the left.
2. In the center pane, select the Time & Date tab.
3. Click Edit Settings.
4. Make sure Use Network Time Protocol (enable NTP client) is selected.
5. Use the drop-down menu to select Start and Stop with Host.
6. Enter the two Nexus switch NTP addresses in the NTP servers box separated by a comma.

Edit time configuration

Specify how the date and time of this host should be set.

☐ Manually configure the date and time on this host

10/13/2016 4:09 PM

☒ Use Network Time Protocol (enable NTP client)

NTP service startup policy: Start and stop with host

NTP servers: 10.1.156.4,10.1.156.5

Separate servers with commas, e.g. 10.31.21.2, fe00::2800

Save Cancel

7. Click Save to save the configuration changes.
8. Select Actions > NTP service > Start.
9. Verify that NTP service is now running and the clock is now set to approximately the correct time



The NTP server time might vary slightly from the host time.

Configure ESXi host swap

ESXi Hosts VM-Host-Infra-01 and VM-Host-Infra-02

To configure host swap on the ESXi hosts, follow these steps on each host:

1. Click Manage in the left navigation pane. Select System in the right pane and click Swap.

vmware ESXi™

Navigator

- Host
 - Manage
 - Monitor
- Virtual Machines 0
- Storage 3
- Networking 5
 - v Switch0
 - iScsiBootvSwitch
 - More networks...

ucsesxia.cie.netapp.com - Manage

System Hardware Licensing Packages Services Security

Advanced settings

Autostart

Swap

Time & date

Edit settings Refresh

| | |
|------------|-----|
| Enabled | Yes |
| Datastore | No |
| Host cache | Yes |
| Local swap | Yes |

2. Click Edit Settings. Select `infra_swap` from the Datastore options.



| Edit swap configuration | |
|-------------------------|---|
| Enabled | <input checked="" type="radio"/> Yes <input type="radio"/> No |
| Datastore | infra_swap ▼ |
| Local swap enabled | <input checked="" type="radio"/> Yes <input type="radio"/> No |
| Host cache enabled | <input checked="" type="radio"/> Yes <input type="radio"/> No |
| <div>Save Cancel</div> | |

3. Click Save.

Install the NetApp NFS Plug-in 1.1.2 for VMware VAAI

To install the NetApp NFS Plug-in 1. 1.2 for VMware VAAI, complete the following steps.

1. Download the NetApp NFS Plug-in for VMware VAAI:
 - a. Go to the [NetApp software download page](#).
 - b. Scroll down and click NetApp NFS Plug-in for VMware VAAI.
 - c. Select the ESXi platform.
 - d. Download either the offline bundle (.zip) or online bundle (.vib) of the most recent plug-in.
2. The NetApp NFS plug-in for VMware VAAI is pending IMT qualification with ONTAP 9.5 and interoperability details will be posted to the NetApp IMT soon.
3. Install the plug-in on the ESXi host by using the ESX CLI.
4. Reboot the ESXI host.

Install VMware vCenter Server 6.7

This section provides detailed procedures for installing VMware vCenter Server 6.7 in a FlexPod Express configuration.

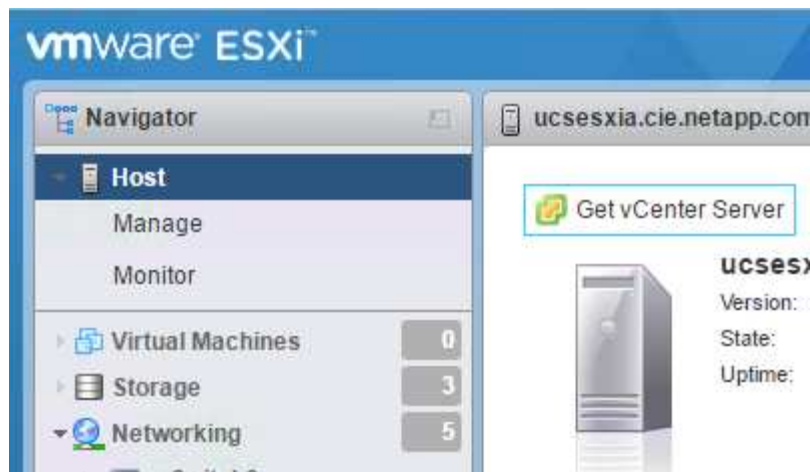


FlexPod Express uses the VMware vCenter Server Appliance (VCSA).

Install VMware vCenter server appliance

To install VCSA, complete the following steps:

1. Download the VCSA. Access the download link by clicking the Get vCenter Server icon when managing the ESXi host.



2. Download the VCSA from the VMware site.



Although the Microsoft Windows vCenter Server installable is supported, VMware recommends the VCSA for new deployments.

3. Mount the ISO image.


4. Navigate to the `vcsa-ui-installer > win32` directory. Double-click `installer.exe`.

5. Click Install.

6. Click Next on the Introduction page.

7. Accept the EULA.

8. Select Embedded Platform Services Controller as the deployment type.



If required, the External Platform Services Controller deployment is also supported as part of the FlexPod Express solution.

9. On the Appliance Deployment Target page, enter the IP address of an ESXi host you have deployed, the root user name, and the root password. Click Next.

Installer vCenter Server Appliance Installer

vm Install - Stage 1: Deploy appliance

1 Introduction
2 End user license agreement
3 Select deployment type
4 Appliance deployment target
5 Set up appliance VM
6 Select deployment size
7 Select datastore
8 Configure network settings
9 Ready to complete stage 1


Appliance deployment target

Specify the appliance deployment target settings. The target is the ESXi host or vCenter Server instance on which the appliance will be deployed.

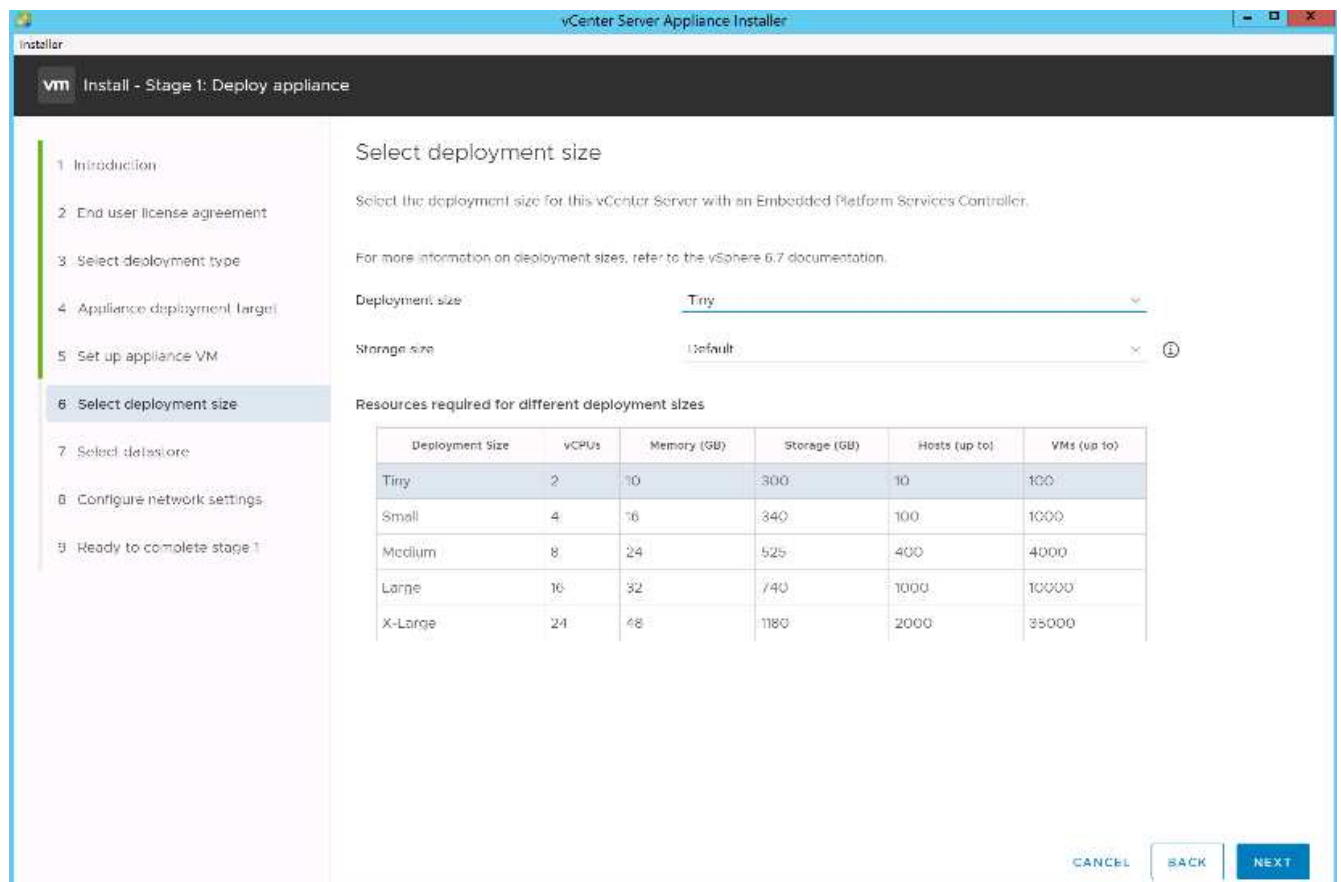
| | |
|----------------------------------|----------------|
| ESXi host or vCenter Server name | 172.18.7.208 ⓘ |
| HTTPS port | 443 |
| User name | root ⓘ |
| Password | ***** |

CANCEL BACK NEXT


10. Set the appliance VM by entering VCSA as the VM name and the root password you would like to use for the VCSA. Click Next.



11. Select the deployment size that best fits your environment. Click Next.




12. Select the `infra_datastore_1` datastore. Click Next.



13. Enter the following information on the Configure Network Settings page and click Next.


- Select MGMT-Network as your network.
- Enter the FQDN or IP to be used for the VCSA.
- Enter the IP address to be used.
- Enter the subnet mask to be used.
- Enter the default gateway.
- Enter the DNS server.



14. On the Ready to Complete Stage 1 page, verify that the settings you have entered are correct. Click Finish.

The VCSA installs now. This process takes several minutes.

15. After stage 1 completes, a message appears stating that it has completed. Click Continue to begin stage 2 configuration.



16. On the Stage 2 Introduction page, click Next.
17. Enter `<<var_ntp_id>>` for the NTP server address. You can enter multiple NTP IP addresses.

If you plan to use vCenter Server high availability, make sure that SSH access is enabled.

18. Configure the SSO domain name, password, and site name. Click Next.

Record these values for your reference, especially if you deviate from the `vsphere.local` domain name.

19. Join the VMware Customer Experience Program if desired. Click Next.
20. View the summary of your settings. Click Finish or use the back button to edit settings.
21. A message appears stating that you are not able to pause or stop the installation from completing after it has started. Click OK to continue.

The appliance setup continues. This takes several minutes.

A message appears indicating that the setup was successful.



The links that the installer provides to access vCenter Server are clickable.

Configure VMware vCenter Server 6.7 and vSphere clustering

To configure VMware vCenter Server 6.7 and vSphere clustering, complete the following steps:

1. Navigate to `https://<<FQDN or IP of vCenter>>/vsphere-client/`.
2. Click Launch vSphere Client.
3. Log in with the user name `administrator@vsphere.local` and the SSO password you entered during the VCSA setup process.
4. Right-click the vCenter name and select New Datacenter.
5. Enter a name for the data center and click OK.

Create vSphere Cluster.

To create a vSphere cluster, complete the following steps:

1. Right-click the newly created data center and select New Cluster.
2. Enter a name for the cluster.
3. Select and enable DRS and vSphere HA options.
4. Click OK.

New Cluster

Flexpod_SeaHawks

×

| | |
|------------|-------------------------------------|
| Name | Express |
| Location | Flexpod_SeaHawks |
| DRS | <input checked="" type="checkbox"/> |
| vSphere HA | <input checked="" type="checkbox"/> |
| vSAN | <input type="checkbox"/> |

These services will have default settings - these can be changed later in the Cluster Quickstart workflow.

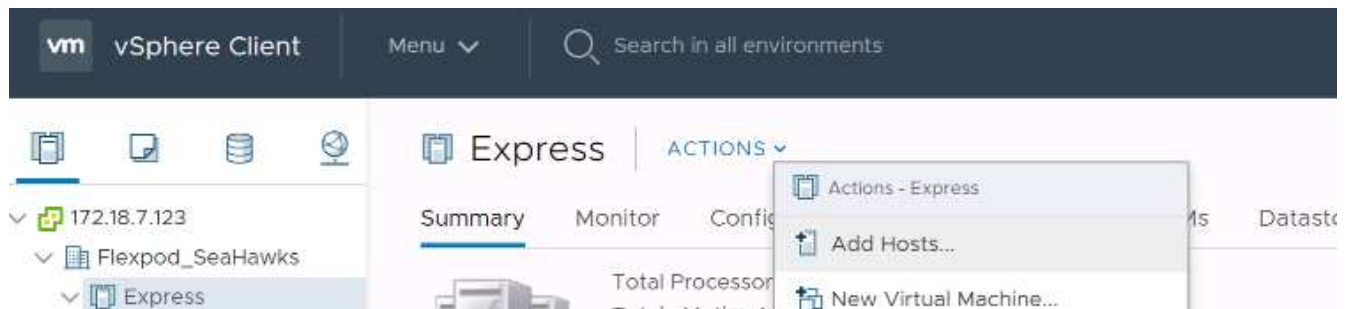
CANCEL

OK


Add ESXi Hosts to Cluster

To add ESXi hosts to the cluster, complete the following steps:

1. Select Add Host in the Actions menu of the cluster.



2. To add an ESXi host to the cluster, complete the following steps:
 - a. Enter the IP or FQDN of the host. Click Next.
 - b. Enter the root user name and password. Click Next.
 - c. Click Yes to replace the host's certificate with a certificate signed by the VMware certificate server.
 - d. Click Next on the Host Summary page.
 - e. Click the green + icon to add a license to the vSphere host.

 This step can be completed later if desired.

- f. Click Next to leave lockdown mode disabled.
- g. Click Next at the VM location page.
- h. Review the Ready to Complete page. Use the back button to make any changes or select Finish.

3. Repeat steps 1 and 2 for Cisco UCS host B.

This process must be completed for any additional hosts added to the FlexPod Express configuration.

Configure coredump on ESXi hosts

ESXi Dump Collector Setup for iSCSI-Booted Hosts

ESXi hosts booted with iSCSI using the VMware iSCSI software initiator need to be configured to do core dumps to the ESXi Dump Collector that is part of vCenter. The Dump Collector is not enabled by default on the vCenter Appliance. This procedure should be run at the end of the vCenter deployment section. To setup the ESXi Dump Collector, follow these steps:

1. Log in to the vSphere Web Client as administrator@vsphere.local and select Home.
2. In the center pane, click System Configuration.
3. In the left pane, select Services.
4. Under Services, click VMware vSphere ESXi Dump Collector.
5. In the center pane, click the green start icon to start the service.
6. In the Actions menu, click Edit Startup Type.
7. Select Automatic.
8. Click OK.
9. Connect to each ESXi host using ssh as root.
10. Run the following commands:

```
esxcli system coredump network set -v vmk0 -j <vcenter-ip>
esxcli system coredump network set -e true
esxcli system coredump network check
```

The message `Verified the configured netdump server is running` appears after you run the final command.



This process must be completed for any additional hosts added to FlexPod Express.

Conclusion

FlexPod Express provides a simple and effective solution by providing a validated design that uses industry-leading components. By scaling through the addition of additional components, FlexPod Express can be tailored for specific business needs. FlexPod Express was designed by keeping in mind small to midsize businesses, ROBOs, and other businesses that require dedicated solutions.

Additional Information

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

- NVA- 1130-DESIGN: FlexPod Express with VMware vSphere 6.7U1 and NetApp AFF A220 with Direct-Attached IP=Based Storage NVA Design

<https://www.netapp.com/us/media/nva-1130-design.pdf>

- AFF and FAS Systems Documentation Center

<http://docs.netapp.com/platstor/index.jsp>

- ONTAP 9 Documentation Center

<http://docs.netapp.com/ontap-9/index.jsp>

- NetApp Product Documentation

<https://docs.netapp.com>

FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS - NVA - Deployment

Jyh-shing Chen, NetApp

The FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS solution leverages Cisco UCS Mini with B200 M5 blade servers, Cisco UCS 6324 in-chassis Fabric Interconnects, Cisco Nexus 31108PC-V switches, or other compliant switches, and NetApp AFF A220, C190, or the FAS2700 series controller HA pair, which runs NetApp ONTAP 9.7 data management software. This NetApp Verified Architecture (NVA) deployment document provides the detailed steps needed to configure the infrastructure components and to deploy VMware vSphere 7.0 and the associated tools to create a highly reliable and highly available FlexPod Express-based virtual infrastructure.

[FlexPod Express for VMware vSphere 7.0 with Cisco UCS Mini and NetApp AFF/FAS - NVA - Deployment](#)

FlexPod and Security

FlexPod, The Solution to Ransomware

TR-4802: FlexPod, The Solution to Ransomware

Arvind Ramakrishnan, NetApp

In partnership with:



To understand ransomware, it is necessary to first understand a few key points about cryptography. Cryptographical methods enable the encryption of data with a shared secret key (symmetric key encryption) or a pair of keys (asymmetric key encryption). One of these keys is a widely available public key and the other is an undisclosed private key.

Ransomware is a type of malware that is based on cryptovirology, which is the use of cryptography to build malicious software. This malware can make use of both symmetric and asymmetric key encryption to lock a victim's data and demand a ransom to provide the key to decrypt the victim's data.

How does ransomware work?

The following steps describe how ransomware uses cryptography to encrypt the victim's data without any scope for decryption or recovery by the victim:

1. The attacker generates a key pair as in asymmetric key encryption. The public key that is generated is placed within the malware, and the malware is then released.
2. After the malware has entered the victim's computer or system, it generates a random symmetric key by using a pseudorandom number generator (PRNG) or any other viable random number- generating algorithm.
3. The malware uses this symmetric key to encrypt the victim's data. It eventually encrypts the symmetric key by using the attacker's public key that was embedded in the malware. The output of this step is an asymmetric ciphertext of the encrypted symmetric key and the symmetric ciphertext of the victim's data.
4. The malware zeroizes (erases) the victim's data and the symmetric key that was used to encrypt the data, thus leaving no scope for recovery.
5. The victim is now shown the asymmetric ciphertext of the symmetric key and a ransom value that must be paid in order to obtain the symmetric key that was used to encrypt the data.
6. The victim pays the ransom and shares the asymmetric ciphertext with the attacker. The attacker decrypts the ciphertext with his or her private key, which results in the symmetric key.
7. The attacker shares this symmetric key with the victim, which can be used to decrypt all the data and thus recover from the attack.

Challenges

Individuals and organizations face the following challenges when they are attacked by ransomware:

- The most important challenge is that it takes an immediate toll on the productivity of the organization or the individual. It takes time to return to a state of normalcy, because all the important files must be regained, and the systems must be secured.
- It could lead to a data breach that contains sensitive and confidential information that belongs to clients or customers and leads to a crisis situation that an organization would clearly want to avoid.
- There is a very good chance of data getting into the wrong hands or being erased completely, which leads to a point of no return that could be disastrous for organizations and individuals.
- After paying the ransom, there is no guarantee that the attacker will provide the key to restore the data.
- There is no assurance that the attacker will refrain from broadcasting the sensitive data in spite of paying the ransom.
- In large enterprises, identifying the loophole that led to a ransomware attack is a tedious task, and securing all the systems involves a lot of effort.

Who is at risk?

Anyone can be attacked by ransomware, including individuals and large organizations. Organizations that do not implement well-defined security measures and practices are even more vulnerable to such attacks. The effect of the attack on a large organization can be several times larger than what an individual might endure.

Ransomware accounts for approximately 28% of all malware attacks. In other words, more than one in four malware incidents is a ransomware attack. Ransomware can spread automatically and indiscriminately through the internet, and, when there is a security lapse, it can enter into the victim's systems and continue to spread to other connected systems. Attackers tend to target people or organizations that perform a lot of file sharing, have a lot of sensitive and critical data, or maintain inadequate protection against attacks.

Attackers tend to focus on the following potential targets:

- Universities and student communities
- Government offices and agencies
- Hospitals
- Banks

This is not an exhaustive list of targets. You cannot consider yourself safe from attacks if you fall outside of one of these categories.

How does ransomware enter a system or spread?

There are several ways in which ransomware can enter a system or spread to other systems. In today's world, almost all systems are connected to one another other through the internet, LANs, WANs, and so on. The amount of data that is being generated and exchanged between these systems is only increasing.

Some of the most common ways by which ransomware can spread include methods that we use on a daily basis to share or access data:

- Email
- P2P networks
- File downloads
- Social networking
- Mobile devices

- Connecting to insecure public networks
- Accessing web URLs

Consequences of data loss

The consequences or effects of data loss can reach more widely than organizations might anticipate. The effects can vary depending on the duration of downtime or the time period during which an organization doesn't have access to its data. The longer the attack endures, the bigger the effect on the organization's revenue, brand, and reputation. An organization can also face legal issues and a steep decline in productivity.

As these issues continue to persist over time, they begin to magnify and might end up changing an organization's culture, depending on how it responds to the attack. In today's world, information spreads at a rapid rate and negative news about an organization could cause permanent damage to its reputation. An organization could face huge penalties for data loss, which could eventually lead to the closure of a business.

Financial effects

According to a recent [McAfee report](#), the global costs incurred due to cybercrime are roughly \$600 billion, which is approximately 0.8% of global GDP. When this amount is compared against the growing worldwide internet economy of \$4.2 trillion, it equates to a 14% tax on growth.

Ransomware takes a significant share of this financial cost. In 2018, the costs incurred due to ransomware attacks were approximately \$8 billion—an amount predicted to reach \$11.5 billion in 2019.

What is the solution?

Recovering from a ransomware attack with minimal downtime is only possible by implementing a proactive disaster recovery plan. Having the ability to recover from an attack is good, but preventing an attack altogether is ideal.


Although there are several fronts that you must review and fix to prevent an attack, the core component that allows you to prevent or recover from an attack is the data center.

The data center design and the features it provides to secure the network, compute, and storage end-points play a critical role in building a secure environment for day-to-day operations. This document shows how the features of a FlexPod hybrid cloud infrastructure can help in quick data recovery in the event of an attack and can also help to prevent attacks altogether.

FlexPod Overview

FlexPod is a predesigned, integrated, and validated architecture that combines Cisco Unified Computing System (Cisco UCS) servers, the Cisco Nexus family of switches, Cisco MDS fabric switches, and NetApp storage arrays into a single, flexible architecture. FlexPod solutions are designed for high availability with no single points of failure, while maintaining cost-effectiveness and design flexibility to support a wide variety of workloads. A FlexPod design can support different hypervisors and bare metal servers and can also be sized and optimized based on customer workload requirements.

The figure below illustrates the FlexPod architecture and clearly highlights the high availability across all the layers of the stack. The infrastructure components of storage, network, and compute are configured in such a way that the operations can instantaneously fail over to the surviving partner in case one of the components fail.



A major advantage for a FlexPod system is that it is predesigned, integrated, and validated for several workloads. Detailed design and deployment guides are published for every solution validation. These documents include the best practices that you must employ for workloads to run seamlessly on FlexPod. These solutions are built with the best-in-class compute, network, and storage products and a host of features that focus on security and hardening of the entire infrastructure.

IBM's [X-Force Threat Intelligence Index](#) states, "Human error responsible for two-thirds of compromised records including historic 424% jump in misconfigured cloud infrastructure."

With a FlexPod system, you can avoid misconfiguring your infrastructure by using automation through Ansible playbooks that perform an end-to-end setup of the infrastructure according to the best practices described in Cisco Validated Designs (CVDs) and NetApp Verified Architectures (NVAs).

Ransomware protection measures

This section discusses the key features of NetApp ONTAP data management software and the tools for Cisco UCS and Cisco Nexus that you can use to effectively protect and recover from ransomware attacks.

Storage: NetApp ONTAP

ONTAP software provides many features useful for data protection, most of which are free of charge to customers who have an ONTAP system. You can use the following features at all times to safeguard data from attacks:

- **NetApp Snapshot technology.** A Snapshot copy is a read-only image of a volume that captures the state of a file system at a point in time. These copies help protect data with no effect on system performance and, at the same time, do not occupy a lot of storage space. NetApp recommends that you create a schedule for the creation of Snapshot copies. You should also maintain a long retention time because some malware can go dormant and then reactivate weeks or months after an infection. In the event of an attack, the volume can be rolled back using a Snapshot copy that was taken before the infection.
- **NetApp SnapRestore technology.** SnapRestore data recovery software is extremely useful to recover from data corruption or to revert only the file contents. SnapRestore does not revert the attributes of a volume; it is much faster than what an administrator can achieve by copying files from the Snapshot copy to the active file system. The speed at which data can be recovered is helpful when many files must be recovered as quickly as possible. In the event of an attack, this highly efficient recovery process helps to get business back online quickly.
- **NetApp SnapCenter technology.** SnapCenter software uses NetApp storage-based backup and replication functions to provide application- consistent data protection. This software integrates with enterprise applications and provides application- specific and database- specific workflows to meet the needs of application, database, and virtual infrastructure administrators. SnapCenter provides an easy-to-use enterprise platform to securely coordinate and manage data protection across applications, databases, and file systems. Its ability to provide application- consistent data protection is critical during data recovery because it makes it easy to restore applications to a consistent state more quickly.
- **NetApp SnapLock technology.** SnapLock provides a special purpose volume in which files can be stored and committed to a nonerasable, nonrewritable state. The user's production data residing in a FlexVol volume can be mirrored or vaulted to a SnapLock volume through NetApp SnapMirror or SnapVault technology, respectively. The files in the SnapLock volume, the volume itself, and its hosting aggregate cannot be deleted until the end of the retention period.
- **NetApp FPolicy technology.** Use FPolicy software to prevent attacks by disallowing operations on files with specific extensions. An FPolicy event can be triggered for specific file operations. The event is tied to a policy, which calls out the engine it needs to use. You might configure a policy with a set of file extensions that could potentially contain ransomware. When a file with a disallowed extension tries to perform an unauthorized operation, FPolicy prevents that operation from executing.

Network: Cisco Nexus

Cisco NX OS software supports the NetFlow feature that enables enhanced detection of network anomalies and security. NetFlow captures the metadata of every conversation on the network, the parties involved in the communication, the protocol being used, and the duration of the transaction. After the information is aggregated and analyzed, it can provide insight into normal behavior.

The collected data also allows identification of questionable patterns of activity, such as malware spreading across the network, which might otherwise go unnoticed.

NetFlow uses flows to provide statistics for network monitoring. A flow is a unidirectional stream of packets that arrives on a source interface (or VLAN) and has the same values for the keys. A key is an identified value for a field within the packet. You create a flow using a flow record to define the unique keys for your flow. You can export the data that NetFlow gathers for your flows by using a flow exporter to a remote NetFlow collector, such as Cisco Stealthwatch. Stealthwatch uses this information for continuous monitoring of the network and provides real-time threat detection and incident response forensics if a ransomware outbreak occurs.

Compute: Cisco UCS

Cisco UCS is the compute endpoint in a FlexPod architecture. You can use several Cisco products that can help to secure this layer of the stack at the operating system level.

You can implement the following key products in the compute or application layer:

- **Cisco Advanced Malware Protection (AMP) for Endpoints.** Supported on Microsoft Windows and Linux operating systems, this solution integrates prevention, detection, and response capabilities. This security software prevents breaches, blocks malware at the point of entry, and continuously monitors and analyzes file and process activity to rapidly detect, contain, and remediate threats that can evade front-line defenses.

The Malicious Activity Protection (MAP) component of AMP continually monitors all endpoint activity and provides run-time detection and blocking of abnormal behavior of a running program on the endpoint. For example, when endpoint behavior indicates ransomware, the offending processes are terminated, preventing endpoint encryption and stopping the attack.

- **Cisco Advanced Malware Protection for Email Security.** Emails have become the prime vehicle to spread malware and to carry out cyber-attacks. On average, approximately 100 billion emails are exchanged in a single day, which provides attackers with an excellent penetration vector into user's systems. Therefore, it is absolutely essential to defend against this line of attack.

AMP analyzes emails for threats such as zero-day exploits and stealthy malware hidden in malicious attachments. It also uses industry-leading URL intelligence to combat malicious links. It gives users advanced protection against spear phishing, ransomware, and other sophisticated attacks.

- **Next-Generation Intrusion Prevention System (NGIPS).** Cisco Firepower NGIPS can be deployed as a physical appliance in the datacenter or as a virtual appliance on VMware (NGIPSv for VMware). This highly effective intrusion prevention system provides reliable performance and a low total cost of ownership. Threat protection can be expanded with optional subscription licenses to provide AMP, application visibility and control, and URL filtering capabilities. Virtualized NGIPS inspects traffic between virtual machines (VMs) and make it easier to deploy and manage NGIPS solutions at sites with limited resources, increasing protection for both physical and virtual assets.

Protect and recover data on FlexPod

This section describes how an end user's data can be recovered in the event of an attack and how attacks can be prevented by using a FlexPod system.

Testbed overview

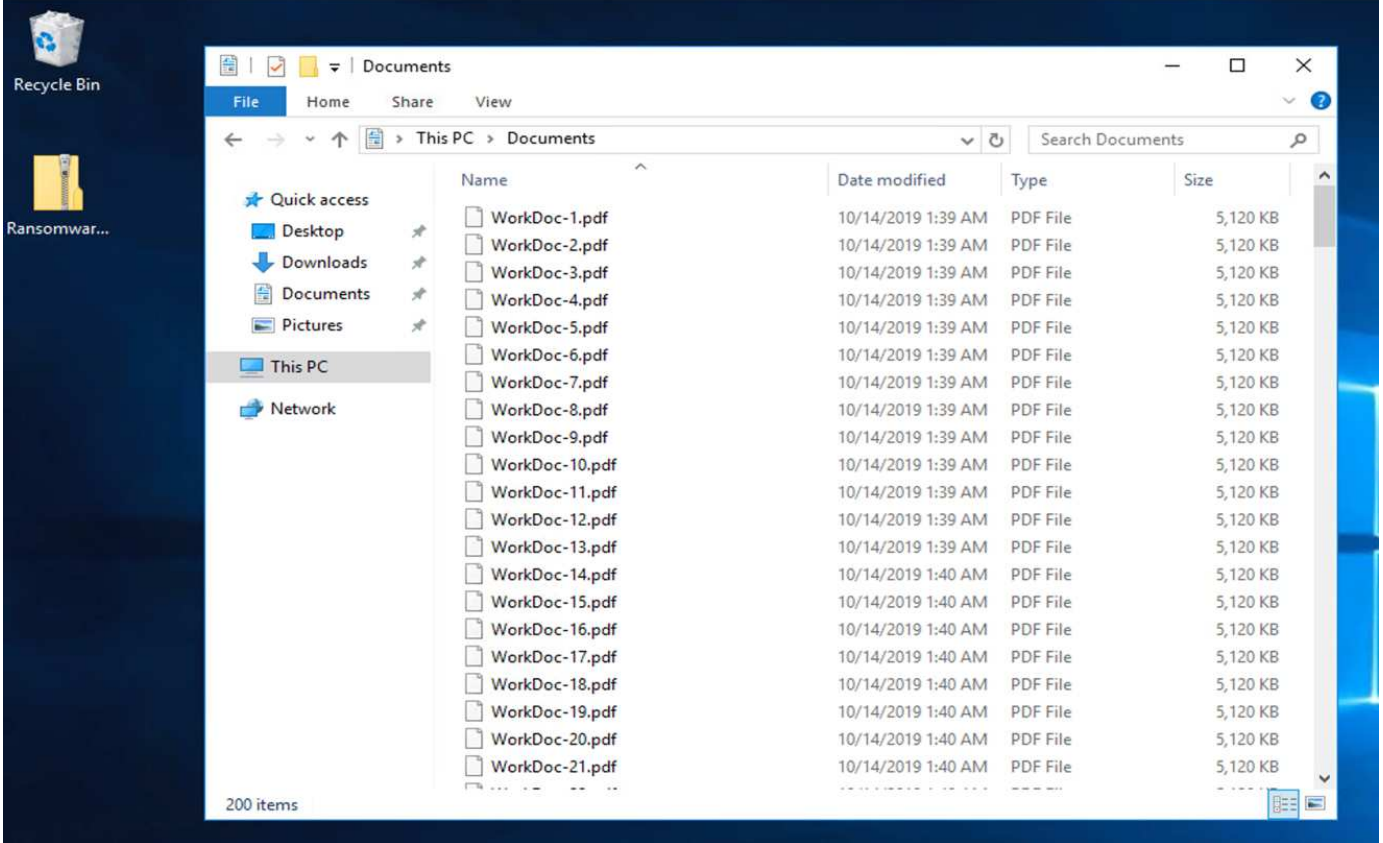
To showcase FlexPod detection, remediation, and prevention, a testbed was built based on the guidelines that are specified in the latest platform CVD available at the time this document was authored: [FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS 4th Generation, and NetApp AFF A-Series CVD](#).

A Windows 2016 VM, which provided a CIFS share from NetApp ONTAP software, was deployed in the VMware vSphere infrastructure. Then NetApp FPolicy was configured on the CIFS share to prevent the execution of files with certain extension types. NetApp SnapCenter software was also deployed to manage the Snapshot copies of the VMs in the infrastructure to provide application- consistent Snapshot copies.


State of VM and its files prior to an attack

This section provides shows the state of the files prior to an attack on the VM and the CIFS share that was mapped to it.


The Documents folder of the VM had a set of PDF files that have not yet been encrypted by the WannaCry malware.



The following screenshot shows the CIFS share that was mapped to the VM.



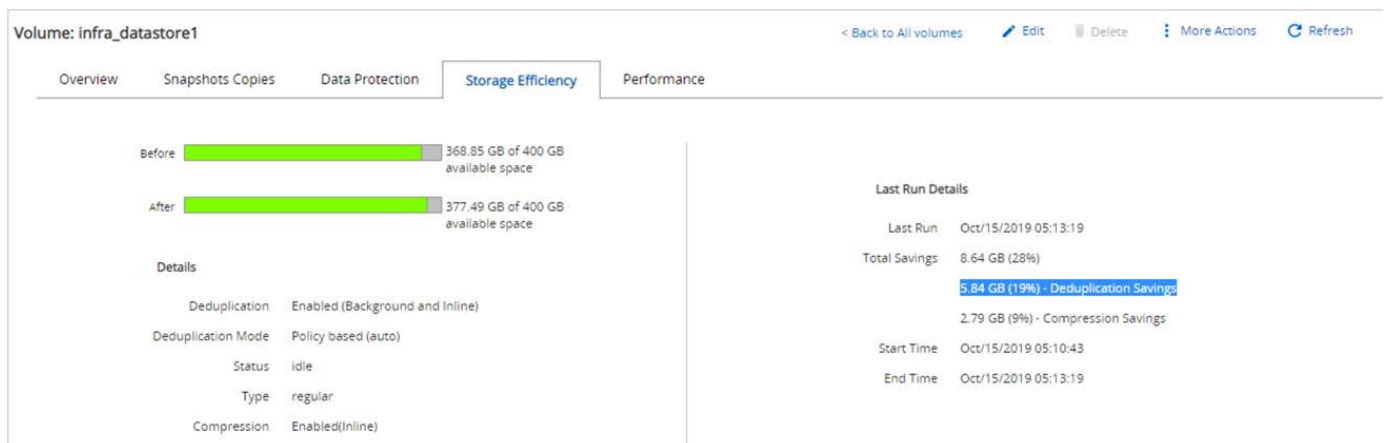
The following screenshot shows the files on the CIFS share `fpolicy_share` that have not yet been encrypted by the WannaCry malware.



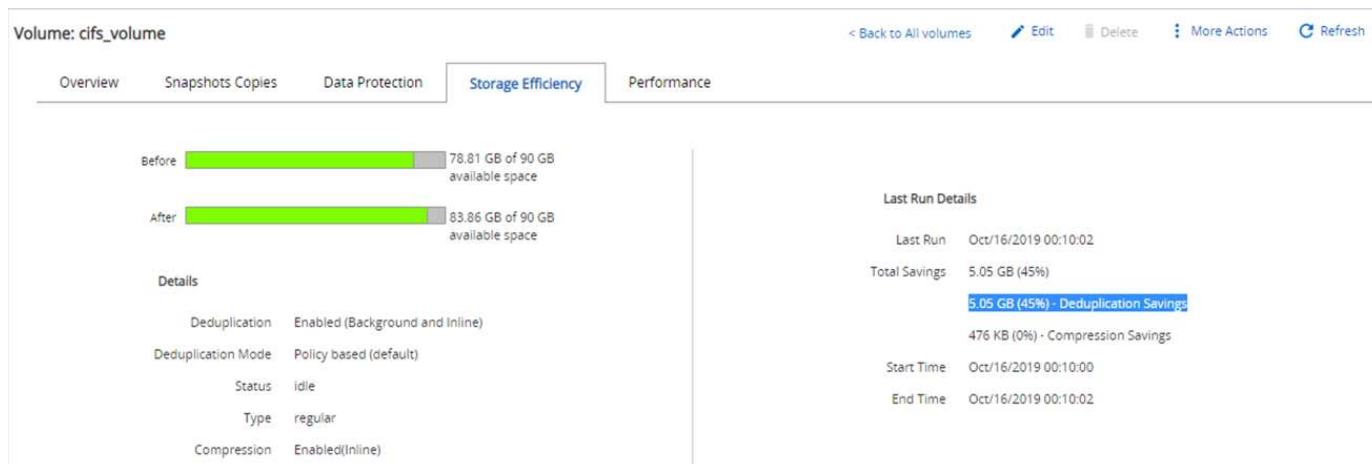
Deduplication and Snapshot information before an attack

The storage efficiency details and size of the Snapshot copy prior to an attack are indicated and used as a reference during the detection phase.

Storage savings of 19% were achieved with deduplication on the volume hosting the VM.



Storage savings of 45% were achieved with deduplication on the CIFS share `fpolicy_share`.



A Snapshot copy size of 456KB was observed for the volume hosting the VM.

Volume: infra_datastore1

< Back to All volumes Edit Delete More Actions Refresh

Overview **Snapshots Copies** Data Protection Storage Efficiency Performance

+ Create Configuration Settings More Actions Delete Refresh

| Status | State | Snapshot Name | Date Time | Total Size | Application Dependency |
|--------|-------|---------------|----------------------|------------|------------------------|
| Normal | -NA- | before_attack | Oct/18/2019 01:44:26 | 456 KB | None |

A Snapshot copy size of 160KB was observed for the CIFS share fpolicy_share.

Volume: cifs_volume

< Back to All volumes Edit Delete More Actions Refresh

Overview **Snapshots Copies** Data Protection Storage Efficiency Performance

+ Create Configuration Settings More Actions Delete Refresh

| Status | State | Snapshot Name | Date Time | Total Size | Application Dependency |
|--------|-------|--------------------|----------------------|------------|------------------------|
| Normal | -NA- | before_attack_cifs | Oct/18/2019 01:45:26 | 160 KB | None |

WannaCry infection on VM and CIFS share


In this section, we show how the WannaCry malware was introduced into the FlexPod environment and the subsequent changes to the system that were observed.

The following steps demonstrate how the WannaCry malware binary was introduced into the VM:

1. The secured malware was extracted.



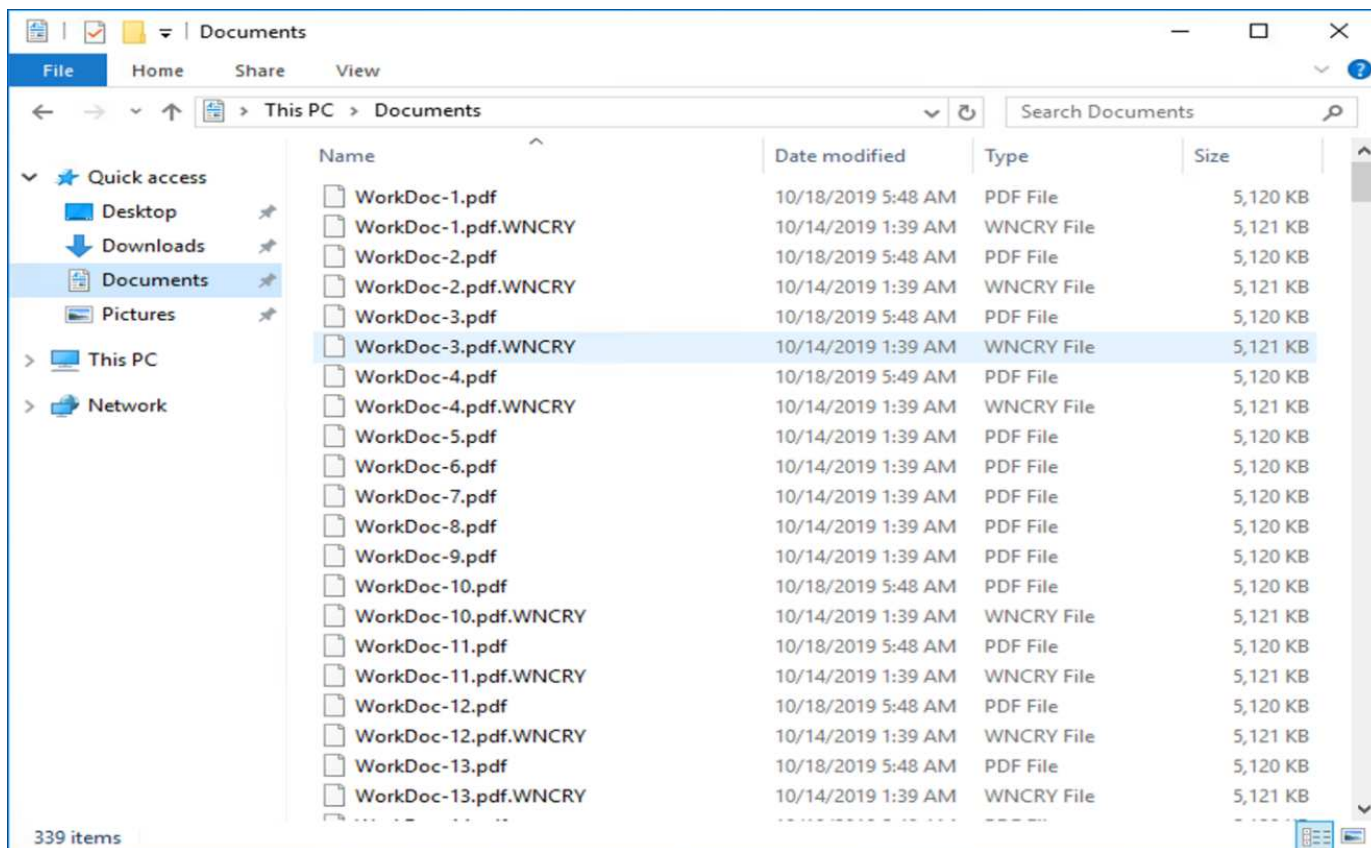
2. The binary was executed.




Case 1: WannaCry encrypts the file system within the VM and mapped CIFS share

The local file system and the mapped CIFS share were encrypted by the WannaCry malware.

Malware starts to encrypt files with WNCRY extensions.



The malware encrypts all the files in the local VM and the mapped share.



Detection

From the moment the malware started to encrypt the files, it triggered an exponential increase in the size of the Snapshot copies and an exponential decrease in the storage efficiency percentage.

We detected a dramatic increase in the Snapshot size to 820.98MB for the volume hosting the CIFS share during the attack.

Volume: cifs_volume

< Back to All volumes Edit Delete More Actions Refresh

Overview **Snapshots Copies** Data Protection Storage Efficiency Performance

+ Create Configuration Settings More Actions Delete Refresh

| Status | State | Snapshot Name | Date Time | Total Size | Application Dependency |
|--------|-------|--------------------|----------------------|------------|------------------------|
| Normal | -NA- | before_attack_cifs | Oct/18/2019 01:45:26 | 820.98 MB | None |

We detected an increase in the Snapshot copy size to 404.3MB for the volume hosting the VM.

Volume: infra_datastore1

< Back to All volumes Edit Delete More Actions Refresh

Overview **Snapshots Copies** Data Protection Storage Efficiency Performance

+ Create Configuration Settings More Actions Delete Refresh

| Status | State | Snapshot Name | Date Time | Total Size | Application Dependency |
|--------|-------|---------------|----------------------|------------|------------------------|
| Normal | -NA- | before_attack | Oct/18/2019 01:44:26 | 404.3 MB | None |

The storage efficiency for the volume hosting the CIFS share decreased to 34%.

Volume: cifs_volume

< Back to All volumes Edit Delete More Actions Refresh

Overview Snapshots Copies Data Protection **Storage Efficiency** Performance

Before 75.21 GB of 90 GB available space

After 80.21 GB of 90 GB available space

Details

| | |
|--------------------|---------------------------------|
| Deduplication | Enabled (Background and inline) |
| Deduplication Mode | Policy based (default) |
| Status | idle |
| Type | regular |
| Compression | Enabled(inline) |

Last Run Details

| | |
|---------------|------------------------------------|
| Last Run | Oct/16/2019 00:10:02 |
| Total Savings | 5 GB (34%) |
| | 5 GB (34%) - Deduplication Savings |
| | 180 KB (0%) - Compression Savings |
| Start Time | Oct/16/2019 00:10:00 |
| End Time | Oct/16/2019 00:10:02 |


Remediation

Restore the VM and mapped CIFS share by using a clean Snapshot copy create prior to the attack.


Restore VM

To restore the VM, complete the following steps:

1. Use the Snapshot copy you created with SnapCenter to restore the VM.



2. Select the desired VMware- consistent Snapshot copy for restore.



3. The entire VM is restored and restarted.

The screenshot shows the 'Restore' wizard window. On the left, a sidebar lists four steps: '1. Select backup' (checked), '2. Select scope' (active and highlighted in blue), '3. Select location', and '4. Summary'. The main area contains the following configuration options:

| | |
|------------------|-------------------------------------|
| Restore scope | Entire virtual machine |
| Restored VM name | WannaCry-2016-1 |
| ESXi host name | 172.21.211.10 |
| Restart VM | <input checked="" type="checkbox"/> |

At the bottom right, there are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

4. Click Finish to start the restore process.


The screenshot shows the 'Restore' wizard window at the '4. Summary' step. The sidebar now highlights '4. Summary' in blue, and the previous steps are marked with green checkmarks. The main area displays a summary of the restore operation:

| | |
|--|-------------------------------------|
| Virtual machine to be restored | WannaCry-2016-1 |
| Backup name | SnapCenter_10-18-2019_01.30.35.0093 |
| Restart virtual machine | Yes |
| ESXi host to be used to mount the backup | 172.21.211.10 |

Below the summary table, there is a yellow warning icon and the text: 'This virtual machine will be powered down during the process.'

At the bottom right, there are four buttons: 'Back', 'Next', 'Finish', and 'Cancel'.

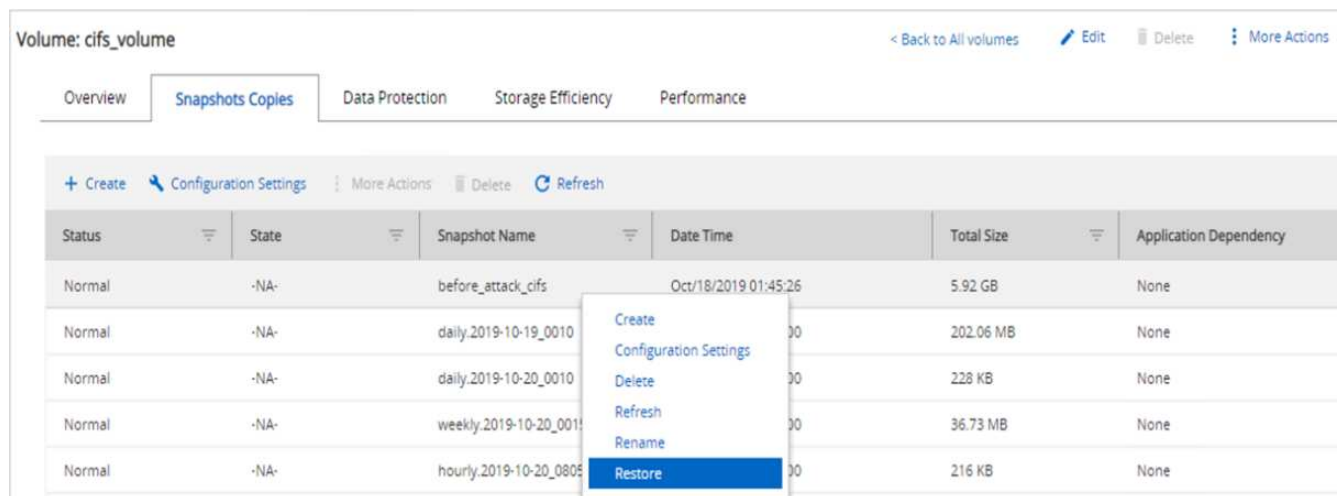
5. The VM and its files are restored.




Restore CIFS Share

To restore the CIFS share, complete the following steps:


1. Use the Snapshot copy of the volume taken prior to the attack to restore the share.



2. Click OK to initiate the restore operation.



3. View the CIFS share after the restore.



Case 2: WannaCry encrypts file system within the VM and tries to encrypt the mapped CIFS share that is protected through FPolicy

Prevention


Configure FPolicy

To configure FPolicy on the CIFS share, run the following commands on the ONTAP cluster:


```
vserver fpolicy policy event create -vserver infra_svm -event-name
Ransomware_event -protocol cifs -file-operations create,rename,write,open
vserver fpolicy policy create -vserver infra_svm -policy-name
Ransomware_policy -events Ransomware_event -engine native
vserver fpolicy policy scope create -vserver infra_svm -policy-name
Ransomware_policy -shares-to-include fpolicy_share -file-extensions-to
-include WNCRY,Locky,ad4c
vserver fpolicy enable -vserver infra_svm -policy-name Ransomware_policy
-sequence-number 1
```

With this policy, files with extensions WNCRY, Locky, and ad4c are not allowed to perform the file operations create, rename, write, or open.

View the status of files prior to attack—they are unencrypted and in a clean system.



The files on the VM are encrypted. The WannaCry malware tries to encrypt the files in the CIFS share, but FPolicy prevents it from affecting the files.



Continue business operations without paying ransom

The NetApp capabilities described in this document help you restore data within minutes after an attack and prevent attacks in the first place so that you can continue business operations unhindered.

A Snapshot copy schedule can be set to meet the desired recovery point objective (RPO). Snapshot copy-based restore operations are very quick; therefore, a very low recovery time objective (RTO) can be achieved.

Above all, you do not have to pay any ransom as a result of an attack, and you can quickly get back to regular operations.

Conclusion

Ransomware is a product of organized crime, and the attackers do not operate with ethics. They can refrain from providing the key for decryption even after receiving the ransom. The victim not only loses their data but also a substantial amount of money and will face consequences associated with the loss of production data.

According to a [Forbes article](#), only 19% of ransomware victims get their data back after paying the ransom. Therefore, the authors recommend not paying a ransom in the event of an attack because doing so reinforces the attacker's faith in their business model.

Data backup and restore operations play an important part of ransomware recovery. Therefore, they must be included as an integral part of business planning. The implementation of these operations should be budgeted for so that there is no compromise on recovery capabilities in the event of an attack.

The key is to select the correct technology partner in this journey, and FlexPod provides most of the needed capabilities natively with no additional cost in an all-flash FAS system.

Acknowledgements

The author would like to thank the following people for their support in the creation of this document:

- Jorge Gomez Navarrete, NetApp
- Ganesh Kamath, NetApp

Additional information

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

- NetApp Snapshot software

<https://www.netapp.com/us/products/platform-os/snapshot.aspx>

- SnapCenter Backup Management

<https://www.netapp.com/us/products/backup-recovery/snapcenter-backup-management.aspx>

- SnapLock Data Compliance

<https://www.netapp.com/us/products/backup-recovery/snaplock-compliance.aspx>

- NetApp Product Documentation

<https://www.netapp.com/us/documentation/index.aspx>

- Cisco Advanced Malware Protection (AMP)

<https://www.cisco.com/c/en/us/products/security/advanced-malware-protection/index.html>

- Cisco Stealthwatch

https://www.cisco.com/c/en_in/products/security/stealthwatch/index.html

FIPS 140-2 security-compliant FlexPod solution for healthcare

TR-4892: FIPS 140-2 security-compliant FlexPod solution for healthcare

JayaKishore Esanakula, NetApp
John McAbel, Cisco

The Health Information Technology for Economic and Clinical Health Act (HITECH) requires Federal Information Processing Standard (FIPS) 140-2-validated encryption of electronic Protected Health Information (ePHI). Health information technology (HIT) applications and software are required to be compliant with FIPS 140-2 for obtaining the Promoting Interoperability Program (formerly, Meaningful Use Incentive Program) certification. Eligible providers and hospitals are required to use a FIPS 140-2 (level 1)

compliant HIT for receiving Medicare and Medicaid incentives and for avoiding reimbursement penalties from the Center for Medicare and Medicaid (CMS). FIPS 140-2 certified encryption algorithms qualify as technical safeguards that are required as per the [Security Rule](#) of the Health Information Portability and Accountability Act (HIPAA).

FIPS 140-2 is a U.S. government standard that sets security requirements for cryptographic modules in hardware, software, and firmware that protect sensitive information. Compliance with the standard is mandated for use by U.S. government agencies, and it is also often used in such regulated industries as financial services and healthcare. This technical report helps the reader to understand the FIPS 140-2 security standard at a high level. It also helps the audience understand various threats faced by healthcare organizations. Finally, the technical report helps one to understand how a FIPS 140-2 compliant FlexPod system can help secure healthcare assets when deployed on a FlexPod converged infrastructure.

Scope

This document is a technical overview of a Cisco Unified Computing System (Cisco UCS), Cisco Nexus, Cisco MDS and NetApp ONTAP-based FlexPod infrastructure for hosting one or more healthcare IT applications or solutions that require FIPS 140-2 security compliance.

Audience

This document is intended for technical leaders in the healthcare industry and for Cisco and NetApp partner solutions engineers and professional services personnel. NetApp assumes that the reader has a good understanding of compute and storage sizing concepts as well as a technical familiarity with healthcare threats, healthcare security, healthcare IT systems, Cisco UCS, and NetApp storage systems.

[Next: Cybersecurity threats in healthcare.](#)

Cybersecurity threats in healthcare

[Previous: Introduction.](#)

Every problem presents a new opportunity—an example of one such opportunity is presented by the COVID pandemic. According to a [report](#) by the Department of Health and Human Services (HHS) Cybersecurity Program, the COVID response has resulted in an increased number of ransomware attacks. There were 6,000 new internet domains registered just in the third week of March 2020. More than 50% of the domains hosted malware. Ransomware attacks were responsible for almost 50% of all healthcare data breaches in 2020 affecting more than 630 healthcare organizations and approximately 29 million healthcare records. Nineteen leakers/sites doubled the extortion. At 24.5%, the healthcare industry saw the highest number of data breaches in 2020.


Malicious agents attempted to breach security and privacy of Protected Health Information (PHI) by selling the information or by threatening to destroy or expose it. Targeted and mass-broadcast attempts are frequently made to gain unauthorized access to ePHI. Approximately 75% of the exposed patient records in the second half of 2020 were due to compromised business associates.

The following list of healthcare organizations were targeted by the malicious agents:

- Hospital systems
- Life science labs

- Research labs
- Rehabilitation facilities
- Community hospitals and clinics

The diversity of applications that constitute a healthcare organization is undeniable and increasingly growing in complexity. Information security offices are challenged to provide governance for the vast array of IT systems and assets. The following figure depicts the clinical capabilities of a typical hospital system.



Patient data is at the heart of this image. The loss of patient data and the stigma associated with sensitive medical conditions is very real. Other sensitive issues include the risk of social exclusion, blackmail, profiling, vulnerability to targeted marketing, exploitation, and potential financial liability toward payers about medical information beyond the payer's privileges.

Threats to healthcare are multidimensional in nature and in impact. Governments worldwide have enacted various provisions to secure ePHI. The detrimental effects and the evolving nature of the threats to healthcare make it difficult for healthcare organizations to defend all threats.

Here is a list of common threats identified in healthcare:

- Ransomware attacks
- Loss or theft of equipment or data with sensitive information
- Phishing attacks
- Attacks against connected medical devices that can affect patient safety
- E-mail phishing attacks
- Loss or theft of equipment or data
- Remote desktop protocol compromise
- Software vulnerability

Healthcare organizations operate in a legal and regulatory environment that is as complicated as their digital ecosystems. This environment includes, but is not limited to, the following:

- Office of the National Coordinator (for Healthcare Technology) ONC Certified Electronic Health Information Technology interoperability standards
- Medicare access and the children's Health Insurance Program Reauthorization Act (MACRA)/Meaningful

Use

- Multiple obligations under the Food and Drug Administration (FDA)
- The Joint Commission accreditation processes
- HIPAA requirements
- HITECH requirements
- Minimum Acceptable Risk Standards for payers
- State privacy and security rules
- Federal Information Security Modernization Act requirements as incorporated into federal contracts and research grants through agencies such as the National Institutes of Health
- Payment Card Industry Data Security Standard (PCI-DSS)
- Substance Abuse and Mental Health Services Administration (SAMHSA) requirements
- The Gramm-Leach-Bliley Act for financial processing
- The Stark Law as it relates to providing services to affiliated organizations
- Family Educational Rights and Privacy Act (FERPA) for institutions that participate in higher education
- Genetic Information Nondiscrimination Act (GINA)
- The new General Data Protection Regulation (GDPR) in the European Union

Security architecture standards are fast evolving to stop the malicious actors from impacting healthcare information systems. One such standard is FIPS 140-2, defined by the National Institute of Standards and Technology (NIST). FIPS publication 140-2 details the U.S. government requirements for a cryptographic module. The security requirements cover areas related to a secure design and implementation of a cryptographic module and can be applied to HIT. Well-defined cryptographic boundaries allow for easier security management while staying current with the cryptographic modules. These boundaries help prevent weak crypto modules that can be easily exploited by malicious actors. They can also help prevent human errors when managing standard cryptographic modules.

NIST along with the Communications Security Establishment (CSE) have established the Cryptographic Module Validation Program (CMVP) to certify cryptographic modules for FIPS 140-2 validation levels. Using a FIPS 140-2 certified module, federal organizations are required to protect sensitive or valuable data while at-rest as well as while in motion. Due to its success in protecting sensitive or valuable information, many healthcare systems have chosen to encrypt ePHI by using FIPS 140-2 cryptographic modules beyond the legally required minimum level of security.

Leveraging and implementing the FlexPod FIPS 140-2 capabilities only takes hours (not days). Becoming FIPS compliant is within reach for most healthcare organizations, regardless of size. With clearly defined cryptographic boundaries and well-documented and simple implementation steps, a FIPS 140-2 compliant FlexPod architecture can set a solid security foundation for infrastructure and allow for simple enhancements to further increase protection for security threats.

[Next: Overview of FIPS 140-2.](#)

Overview of FIPS 140-2

[Previous: Cybersecurity threats in healthcare.](#)

FIPS 140-2 specifies the security requirements for a cryptographic module used within a security system that protects sensitive information in computer and telecommunication

systems. A cryptographic module should be a set of hardware, software, firmware, or a combination. FIPS applies to the cryptographic algorithms, key generation and key managers contained within a cryptographic boundary. It is important to understand that FIPS 140-2 applies specifically to the cryptographic module, not the product, architecture, data, or ecosystem. The cryptographic module, which is defined in the key terms later in this document, is the specific component (whether it's hardware, software, and/or firmware) that implements approved security functions. In addition, FIPS 140-2 specifies four levels. Approved cryptographic algorithms are common to all levels. Key elements and requirements of each security level include:

- **Security level 1**

- Specifies basic security requirements for a cryptographic module (at least one approved algorithm or security function is required).
- No specified physical security mechanisms are required for level 1 beyond the basic requirements for production-grade components.

- **Security level 2**

- Enhances the physical security mechanisms by adding the requirement for tamper-evidence by using tamper-evident solutions such as coatings or seals, locks on removable covers or doors of the cryptographic modules.
- Requires, at minimum, role-based access control (RBAC) in which the cryptographic module authenticates the authorization of an operator or administrator to assume a specific role and perform a corresponding set of functions.

- **Security level 3**

- Builds on the tamper-evident requirements of level 2 and attempts to prevent further access to critical security parameters (CSPs) within the cryptographic module.
- Physical security mechanisms required at level 3 are intended to have a high probability to detect and respond to attempts at physical access, or any use or modification of the cryptographic module. Examples might include strong enclosures, tamper detection, and response circuitry that zeros all plaintext CSPs when a removable cover on the cryptographic module is opened.
- Requires identity-based authentication mechanisms to enhance the security of the RBAC mechanisms specified in level 2. A cryptographic module authenticates the identity of an operator and verifies that the operator is authorized to use a role and perform the functions of the role.

- **Security level 4**

- The highest level of security in FIPS 140-2.
- The most useful level for operations in physically unprotected environments.
- At this level, the physical security mechanisms are intended to provide complete protection around the cryptographic module with the responsibility of detecting and responding to any unauthorized attempts at physical access.
- Penetration or exposure of the cryptographic module should have a high probability of detection and result in the immediate zeroization of all unsecure or plaintext CSPs.

Next: [Control plane versus data plane.](#)

Control plane versus data plane

[Previous: Overview of FIPS 140-2.](#)

When implementing a FIPS 140-2 strategy, it is important to understand what is being protected. This can easily be broken down into two areas: control plane and data plane. A control plane refers to the aspects that affect the control and operation of the components within the FlexPod system: for example, administrative access to the NetApp storage controllers, Cisco Nexus switches, and Cisco UCS servers. Protection at this layer is provided by limiting the protocols and cryptographic cyphers that administrators can use to connect to devices and make changes. A data plane refers to the actual information, such as the PHI, within the FlexPod system. This is protected by encrypting data at rest and again for FIPS, ensuring that the cryptographic modules in use meet the standards.

[Next: FlexPod Cisco UCS compute and FIPS 140-2.](#)

FlexPod Cisco UCS compute and FIPS 140-2

[Previous: Control plane versus data plane.](#)

A FlexPod architecture can be designed with a Cisco UCS server that is FIPS 140-2 compliant. In accordance with the U. S. NIST, Cisco UCS server can operate in FIPS 140-2 level 1 compliance mode. For a complete list of FIPS-compliant Cisco components, see [Cisco's FIPS 140 page](#). Cisco UCS Manager is FIPS 140-2 validated.

Cisco UCS and Fabric Interconnect

Cisco UCS Manager is deployed and runs from the Cisco Fabric Interconnects (FIs).

For more information about Cisco UCS and how to enable FIPS, see the [Cisco UCS Manager documentation](#).

To enable FIPS mode on the Cisco fabric interconnect on each fabric A and B, run the following commands:

```
fp-health-fabric-A# connect local-mgmt
fp-health-fabric-A(local-mgmt)# enable fips-mode
FIPS mode is enabled
```



To replace an FI in a cluster on Cisco UCS Manager Release 3.2(3) with an FI on a release earlier than Cisco UCS Manager Release 3.2(3), disable FIPS mode (disable `fips-mode`) on the existing FI before adding the replacement FI to the cluster. After the cluster is formed, as part of the Cisco UCS Manager boot up, FIPS mode is automatically enabled.

Cisco offers the following key products that can be implemented in the compute or application layer:

- **Cisco Advanced Malware Protection (AMP) for endpoints.** Supported on Microsoft Windows and Linux operating systems, this solution integrates prevention, detection, and response capabilities. This security software prevents breaches, blocks malware at the point of entry, and continuously monitors and analyzes file and process activity to rapidly detect, contain, and remediate threats that can evade front-line defenses. The Malicious Activity Protection (MAP) component of AMP continually monitors all endpoint activity and

provides run-time detection and blocking of abnormal behavior of a running program on the endpoint. For example, when endpoint behavior indicates ransomware, the offending processes are terminated, preventing endpoint encryption and stopping the attack.

- **AMP for email security.** Emails have become the prime vehicle to spread malware and to carry out cyberattacks. On average, approximately 100 billion emails are exchanged in a single day, which provides attackers with an excellent penetration vector into user's systems. Therefore, it is absolutely essential to defend against this line of attack. AMP analyzes emails for threats such as zero-day exploits and stealthy malware hidden in malicious attachments. It also uses industry-leading URL intelligence to combat malicious links. It gives users advanced protection against spear phishing, ransomware, and other sophisticated attacks.
- **Next- Generation Intrusion Prevention System (NGIPS).** Cisco Firepower NGIPS can be deployed as a physical appliance in the data center or as a virtual appliance on VMware (NGIPSv for VMware). This highly effective intrusion prevention system provides reliable performance and a low total cost of ownership. Threat protection can be expanded with optional subscription licenses to provide AMP, application visibility and control, and URL filtering capabilities. Virtualized NGIPS inspects traffic between virtual machines (VMs) and makes it easier to deploy and manage NGIPS solutions at sites with limited resources, increasing protection for both physical and virtual assets.

[Next: FlexPod Cisco networking and FIPS 140-2.](#)

FlexPod Cisco networking and FIPS 140-2

[Previous: FlexPod Cisco UCS compute and FIPS 140-2.](#)

Cisco MDS

Cisco MDS 9000 series platform with software 8.4.x is [FIPS 140-2 compliant](#). Cisco MDS implements cryptographic modules and the following services for SNMPv3 and SSH.

- Session establishment supporting each service
- All underlying cryptographic algorithms supporting each services key derivation functions
- Hashing for each service
- Symmetric encryption for each service

Before you enable FIPS mode, complete the following tasks on the MDS switch:

1. Make your passwords a minimum of eight characters in length.
2. Disable Telnet. Users should log in using SSH only.
3. Disable remote authentication through RADIUS/TACACS+. Only users local to the switch can be authenticated.
4. Disable SNMP v1 and v2. Any existing user accounts on the switch that have been configured for SNMPv3 should be configured only with SHA for authentication and AES/3DES for privacy.
5. Disable VRRP.
6. Delete all IKE policies that either have MD5 for authentication or DES for encryption. Modify the policies so they use SHA for authentication and 3DES/AES for encryption.
7. Delete all SSH Server RSA1 keypairs.

To enable FIPS mode and to display FIPS status on the MDS switch, complete the following steps:

1. Show the FIPS status.

```
MDSSwitch# show fips status
FIPS mode is disabled
MDSSwitch# conf
Enter configuration commands, one per line.  End with CNTL/Z.
```

2. Set up the 2048 bits SSH key.

```
MDSSwitch(config)# no feature ssh
XML interface to system may become unavailable since ssh is disabled
MDSSwitch(config)# no ssh key
MDSSwitch(config)# show ssh key
*****
could not retrieve rsa key information
bitcount: 0
*****
could not retrieve dsa key information
bitcount: 0
*****
no ssh keys present. you will have to generate them
*****
MDSSwitch(config)# ssh key
dsa    rsa
MDSSwitch(config)# ssh key rsa 2048 force
generating rsa key(2048 bits).....
...
generated rsa key
```

3. Enable FIPS mode.

```
MDSSwitch(config)# fips mode enable
FIPS mode is enabled
System reboot is required after saving the configuration for the system
to be in FIPS mode
Warning: As per NIST requirements in 6.X, the minimum RSA Key Size has
to be 2048
```

4. Show the FIPS status.

```
MDSSwitch(config)# show fips status
FIPS mode is enabled
MDSSwitch(config)# feature ssh
MDSSwitch(config)# show feature | grep ssh
sshServer          1          enabled
```

5. Save the configuration to the running configuration.

```
MDSSwitch(config)# copy ru st
[#####] 100%
exitCopy complete.
MDSSwitch(config)# exit
```

6. Restart MDS switch

```
MDSSwitch# reload
This command will reboot the system. (y/n)? [n] y
```

7. Show the FIPS status.

```
Switch(config)# fips mode enable
Switch(config)# show fips status
```

For more information, see [Enabling FIPS Mode](#).

Cisco Nexus

Cisco Nexus 9000 series switches (version 9.3) are [FIPS 140-2 compliant](#). Cisco Nexus implements cryptographic modules and the following services for SNMPv3 and SSH.

- Session establishment supporting each service
- All underlying cryptographic algorithms supporting each services key derivation functions
- Hashing for each service
- Symmetric encryption for each service

Before you enable FIPS mode, complete the following tasks on the Cisco Nexus switch:

1. Disable Telnet. Users should log in using Secure Shell (SSH) only.
2. Disable SNMPv1 and v2. Any existing user accounts on the device that have been configured for SNMPv3 should be configured only with SHA for authentication and AES/3DES for privacy.
3. Delete all SSH server RSA1 key-pairs.
4. Enable HMAC-SHA1 message integrity checking (MIC) to use during the Cisco TrustSec Security

Association Protocol (SAP) negotiation. To do so, enter the sap hash-algorithm HMAC-SHA-1 command from the `cts-manual` or `cts-dot1x` mode.

To enable FIPS mode on the Nexus switch, complete the following steps:

1. Set up 2048 bits SSH key.

```
NexusSwitch# show fips status
FIPS mode is disabled
NexusSwitch# conf
Enter configuration commands, one per line.  End with CNTL/Z.
```

2. Set up the 2048 bits SSH key.

```
NexusSwitch(config)# no feature ssh
XML interface to system may become unavailable since ssh is disabled
NexusSwitch(config)# no ssh key
NexusSwitch(config)# show ssh key
*****
could not retrieve rsa key information
bitcount: 0
*****
could not retrieve dsa key information
bitcount: 0
*****
no ssh keys present. you will have to generate them
*****
NexusSwitch(config)# ssh key
dsa    rsa
NexusSwitch(config)# ssh key rsa 2048 force
generating rsa key(2048 bits).....
...
generated rsa key
```

3. Enable FIPS mode.

```

NexusSwitch(config)# fips mode enable
FIPS mode is enabled
System reboot is required after saving the configuration for the system
to be in FIPS mode
Warning: As per NIST requirements in 6.X, the minimum RSA Key Size has
to be 2048
Show fips status
NexusSwitch(config)# show fips status
FIPS mode is enabled
NexusSwitch(config)# feature ssh
NexusSwitch(config)# show feature | grep ssh
sshServer          1          enabled
Save configuration to the running configuration
NexusSwitch(config)# copy ru st
[#####] 100%
exitCopy complete.
NexusSwitch(config)# exit

```

4. Restart the Nexus switch.

```

NexusSwitch# reload
This command will reboot the system. (y/n)? [n] y

```

5. Show the FIPS status.

```

NexusSwitch(config)# fips mode enable
NexusSwitch(config)# show fips status

```

Additionally, Cisco NX OS software supports the NetFlow feature that enables enhanced detection of network anomalies and security. NetFlow captures the metadata of every conversation on the network, the parties involved in the communication, the protocol being used, and the duration of the transaction. After the information is aggregated and analyzed, it can provide insight into normal behavior. The collected data also allows identification of questionable patterns of activity, such as malware spreading across the network, which might otherwise go unnoticed. NetFlow uses flows to provide statistics for network monitoring. A flow is a unidirectional stream of packets that arrives on a source interface (or VLAN) and has the same values for the keys. A key is an identified value for a field within the packet. You create a flow using a flow record to define the unique keys for your flow. You can export the data that NetFlow gathers for your flows by using a flow exporter to a remote NetFlow collector, such as Cisco Stealthwatch. Stealthwatch uses this information for continuous monitoring of the network and provides real-time threat detection and incident response forensics if a ransomware outbreak occurs.

[Next: FlexPod NetApp ONTAP storage and FIPS 140-2.](#)

FlexPod NetApp ONTAP storage and FIPS 140-2

[Previous: FlexPod Cisco networking and FIPS 140-2.](#)

NetApp offers a variety of hardware, software, and services, which can include various components of the cryptographic modules validated under the standard. Therefore, NetApp uses a variety of approaches for FIPS 140-2 compliance for the control plane and data plane:

- NetApp includes cryptographic modules that have achieved level 1 validation for data-in-transit and data-at-rest encryption.
- NetApp acquires both hardware and software modules that have been FIPS 140-2 validated by the suppliers of those components. For example, the NetApp Storage Encryption solution leverages FIPS level 2 validated drives.
- NetApp products can use a validated module in a way that complies with the standard even though the product or feature is not within the boundary of the validation. For example, NetApp Volume Encryption (NVE) is FIPS 140-2 compliant. Although not separately validated, it leverages the NetApp cryptographic module, which is level 1 validated. To understand the specifics of compliance for your version of ONTAP, contact your FlexPod SME.

NetApp Cryptographic modules are FIPS 140-2 level 1 validated

- The NetApp Cryptographic Security Module (NCSM) is FIPS 140-2 level 1 validated.

NetApp self-encrypting drives are FIPS 140-2 level 2 validated

NetApp purchases self-encrypting drives (SEDs) that have been FIPS 140-2 validated by the original equipment manufacturer (OEM); customers seeking these drives must specify them when ordering. Drives are validated at level 2. The following NetApp products can leverage validated SEDs:

- AFF A-Series and FAS storage systems
- E-Series and EF-Series storage systems


NetApp Aggregate Encryption and NetApp Volume Encryption

NVE and NetApp Aggregate Encryption (NAE) technologies enable encryption of data at the volume and aggregate level respectively, making the solution agnostic to the physical drive.

NVE is a software-based, data-at-rest encryption solution available starting with ONTAP 9.1, and it has been FIPS 140-2 compliant since ONTAP 9.2. NVE allows ONTAP to encrypt data for each volume for granularity. NAE, available with ONTAP 9.6, is an outgrowth of NVE; it allows ONTAP to encrypt data for each volume, and the volumes can share keys across the aggregate. Both NVE and NAE use AES 256-bit encryption. Data can also be stored on disk without SEDs. NVE and NAE enable you to use storage efficiency features even when encryption is enabled. An application- layer- only encryption defeats all benefits of storage efficiency. With NVE and NAE, storage efficiencies are maintained because the data comes in from the network through NetApp WAFL to the RAID layer, which determines whether the data should be encrypted. For greater storage efficiency, you can use aggregate deduplication with NAE. NVE volumes and NAE volumes can coexist on the same NAE aggregate. NAE aggregates do not support unencrypted volumes.

Here's how the process works: When data is encrypted, it is sent to the cryptographic module which is FIPS 140-2 level 1 validated. The cryptographic module encrypts the data and sends it back to the RAID layer. The encrypted data is then sent to the disk. Therefore, with the combination of NVE and NAE, the data is already encrypted on the way to the disk. Reads follow the reverse path. In other words, the data leaves the disk

encrypted, is sent to RAID, is decrypted by the cryptographic module, and is then sent up the rest of the stack, as shown in the following figure.



NVE uses a software cryptographic module which is FIPS 140-2 level 1 validated.

For more information about NVE, see the [NVE Datasheet](#).

NVE protects data in the cloud. Cloud Volumes ONTAP and Azure NetApp Files are capable of providing FIPS 140-2 compliant data encryption at rest.

Starting with ONTAP 9.7, newly created aggregates and volumes are encrypted by default when you have the NVE license and onboard or external key management. Starting with ONTAP 9.6, you can use aggregate-level encryption to assign keys to the containing aggregate for the volumes to be encrypted. Volumes you create in the aggregate are encrypted by default. You can override the default when you encrypt the volume.

ONTAP NAE CLI commands

Before you run the following CLI commands, make sure the cluster has the required NVE license.

To create an aggregate and encrypt it, run the following command (when run on an ONTAP 9.6 and later cluster CLI):

```
fp-health::> storage aggregate create -aggregate aggregatename -encrypt  
-with-aggr-key true
```

To convert a non-NAE aggregate to an NAE an aggregate, run the following command (when run on an ONTAP 9.6 and later cluster CLI):

```
fp-health::> storage aggregate modify -aggregate aggregatename -node  
svmname -encrypt-with-aggr-key true
```

To convert an NAE aggregate to a non-NAE aggregate, run the following command (when run on an ONTAP 9.6 and later cluster CLI):

```
fp-health::> storage aggregate modify -aggregate aggregatename -node  
svmname -encrypt-with-aggr-key false
```

ONTAP NVE CLI commands

Starting with ONTAP 9.6, you can use aggregate-level encryption to assign keys to the containing aggregate for the volumes to be encrypted. Volumes you create in the aggregate are encrypted by default.

To create a volume on an aggregate that is NAE enabled, run the following command (when run on an ONTAP 9.6 and later cluster CLI):

```
fp-health::> volume create -vserver svmname -volume volumename -aggregate  
aggregatename -encrypt true
```

To enable encryption of an existing volume “inplace” without a volume move, run the following command (when run on an ONTAP 9.6 and later cluster CLI):

```
fp-health::> volume encryption conversion start -vserver svmname -volume  
volumename
```


To verify that volumes are enabled for encryption, run the following CLI command:

```
fp-health::> volume show -is-encrypted true
```

NSE

NSE uses SEDs to perform the data encryption through a hardware-accelerated mechanism.

NSE is configured to use FIPS 140-2 level 2 self-encrypting drives to facilitate compliance and spares return by enabling the protection of data at rest through AES 256-bit transparent disk encryption. The drives perform all of the data encryption operations internally, as depicted in the following figure, including encryption key generation. To prevent unauthorized access to the data, the storage system must authenticate itself with the drive using an authentication key that is established the first time the drive is used.



NSE uses hardware encryption on each drive, which is FIPS 140-2 level 2 validated.

For more information about NSE, see the [NSE datasheet](#).

Key management

The FIPS 140-2 standard applies to the cryptographic module as defined by the boundary, as shown in the following figure.

2.1.1 Cryptographic Boundary

The logical cryptographic boundary of the CryptoMod module is the `cryptomod_fips.ko` component of ONTAP OS kernel. The logical boundary is depicted in the block diagram below. The Approved DRBG is used to supply the module's cryptographic keys. The physical boundary for the module is the enclosure of the NetApp controller.




Figure 1 - Block Diagram

Key manager keeps track of all the encryption keys used by ONTAP. NSE SEDs use the key manager to set the authentication keys for NSE SEDs. When using the key manager, the combined NVE and NAE solution is composed of a software cryptographic module, encryption keys, and a key manager. For each volume, NVE uses a unique XTS-AES 256 data encryption key, which the key manager stores. The key used for a data volume is unique to the data volume in that cluster and is generated when the encrypted volume is created. Similarly, an NAE volume uses unique XTS-AES 256 data encryption keys per aggregate, which the key manager also stores. NAE keys are generated when the encrypted aggregate is created. ONTAP does not pregenerate keys, reuse them, or display them in plain text—they are stored and protected by the key manager.

Support for external key manager

Beginning with ONTAP 9.3, external key managers are supported in both NVE and NSE solutions. The FIPS 140-2 standard applies to the cryptographic module used in the specific vendor's implementation. Most often, FlexPod and ONTAP customers use one of the following validated (per the [NetApp Interoperability Matrix](#)) key managers:

- Gemalto or SafeNet AT

- Vormetric (Thales)
- IBM SKLM
- Utimaco (formerly Microfocus, HPE)

NSE and NVMe SED authentication key is backed up to an external key manager by using the industry-standard OASIS Key Management Interoperability Protocol (KMIP). Only the storage system, drive, and key manager have access to the key, and the drive cannot be unlocked if it is moved outside the security domain, thus preventing data leakage. The external key manager also stores NVE volume encryption keys and NAE aggregate encryption keys. If the controller and disks are moved and no longer have access to the external key manager, the NVE and NAE volumes won't be accessible and cannot be decrypted.

The following example command adds two key management servers to the list of servers used by the external key manager for store virtual machine (SVM) `svmname1`.

```
fp-health::> security key-manager external add-servers -vserver svmname1
-key-servers 10.0.0.20:15690, 10.0.0.21:15691
```

When a FlexPod Datacenter is being used in a multitenancy scenario, ONTAP enables users by providing tenancy separation for security reasons at the SVM level.

To verify list of external key managers, run the following CLI command:

```
fp-health::> security key-manager external show
```

Combine encryption for double encryption (layered defense)

If you need to segregate access to data and make sure that data is protected all the time, NSE SEDs can be combined with network- or fabric-level encryption. NSE SEDs act like a backstop if an administrator forgets to configure or misconfigures higher-level encryption. For two distinct layers of encryption, you can combine NSE SEDs with NVE and NAE.

NetApp ONTAP cluster-wide control plane FIPS mode

NetApp ONTAP data management software has a FIPS mode configuration that instantiates an added level of security for the customer. This FIPS mode only applies to the control plane. When FIPS mode is enabled, in accordance with key elements of FIPS 140-2, Transport Layer Security v1 (TLSv1) and SSLv3 are disabled, and only TLS v1.1 and TLS v1.2 remain enabled.



ONTAP cluster-wide control pane in FIPS mode is FIPS 140-2 level 1 compliant. Cluster-wide FIPS mode uses a software-based cryptographic module provided by NCSM.

FIPS 140-2 compliance mode for cluster-wide control plane secures all control Interfaces of ONTAP. By default, the FIPS 140-2 only mode is disabled; however you can enable this mode by setting the `is-fips-enabled` parameter to `true` for the `security config modify` command.

To enable FIPS mode on the ONTAP cluster, run the following command:

```
fp-health::~> security config modify -interface SSL -is-fips-enabled true
```

When SSL FIPS mode is enabled, SSL communication from ONTAP to the external client or server components outside of ONTAP will use FIPS compliant cryptographic for SSL.

To show the FIPS status for the entire cluster, run the following commands:

```
fp-health::~> set advanced
fp-health::*> security config modify -interface SSL -is-fips-enabled true
```

[Next: Solution benefits of FlexPod converged infrastructure.](#)

Solution benefits of FlexPod converged infrastructure


[Previous: FlexPod NetApp ONTAP storage and FIPS 140-2.](#)

Healthcare organizations have several mission-critical systems. Two of the most critical systems are the electronic health record (EHR) systems and medical imaging systems. To demonstrate the FIPS setup on a FlexPod system, we used an open-source EHR and an open-source picture archiving and communication system (PACS) system for the lab setup and workload validation on the FlexPod system. For a complete list of EHR capabilities, EHR logical application components, and how EHR systems benefit when implemented on a FlexPod system see [TR-4881: FlexPod for Electronic Health Record Systems](#). For a complete list of a medical imaging system capabilities, logical application components, and how medical imaging systems benefit when implemented on FlexPod, see [TR-4865: FlexPod for Medical Imaging](#).

During the FIPS setup and workload validation, we exercised workload characteristics that were representative of a typical healthcare organization. For example, we exercised an open-source EHR system to include realistic patient data access and change scenarios. Additionally, we exercised medical imaging workloads that included digital imaging and communications in medicine (DICOM) objects in a *.dcm file format. DICOM objects with metadata were stored on both the file and block storage. Additionally, we implemented multipathing capabilities from within a virtualized RedHat Enterprise Linux (RHEL) server. We stored DICOM objects on an NFS, mounted LUNs using iSCSI, and mounted LUNs using FC. During the FIPS setup and validation, we observed that the FlexPod converged infrastructure exceeded our expectations and performed seamlessly.

The following figure depicts the FlexPod system used for FIPS setup and validation. We leveraged the [FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7 Cisco Validated Design \(CVD\)](#) during the setup process.

FIPS 140-2 security compliant FlexPod for Healthcare



Solution infrastructure hardware and software components

The following two figures list the hardware and software components respectively used during the FIPS testing enabling on a FlexPod. The recommendations in these tables are examples; you should work with your NetApp SME to make sure that the components are suitable for your organization. Also, make sure that the components and versions are supported in the [NetApp Interoperability Matrix Tool \(IMT\)](#) and [Cisco Hardware Compatibility List \(HCL\)](#).

| Layer | Product family | Quantity and model | Details |
|-----------------|--|--------------------------------|--|
| Compute | Cisco UCS 5108 chassis | 1 or 2 | |
| | Cisco UCS blade servers | 3 B200 M5 | Each with 2x 20 or more cores, 2.7GHz, and 128-384GB RAM |
| | Cisco UCS Virtual Interface Card (VIC) | Cisco UCS 1440 | See the |
| | 2x Cisco UCS Fabric Interconnects | 6332 | - |
| Network | Cisco Nexus switches | 2x Cisco Nexus 9332 | - |
| Storage network | IP network for storage access over SMB/CIFS, NFS, or iSCSI protocols | Same network switches as above | - |
| | Storage access over FC | 2x Cisco MDS 9148S | - |

| Layer | Product family | Quantity and model | Details |
|---------|--|--------------------------------|--------------------------------|
| Storage | NetApp AFF A700 all-flash storage system | 1 Cluster | Cluster with two nodes |
| | Disk shelf | One DS224C or NS224 disk shelf | Fully populated with 24 drives |
| | SSD | >24, 1.2TB or larger capacity | - |

| Software | Product family | Version or release | Details |
|------------|--|---|---------|
| Various | Linux | RHEL 7.X | - |
| | Windows | Windows Server 2012 R2 (64 bit) | - |
| | NetApp ONTAP | ONTAP 9.7 or later | - |
| | Cisco UCS Fabric Interconnect | Cisco UCS Manager 4.1 or later | - |
| | Cisco Ethernet 3000 or 9000 series switches | For 9000 series, 7.0(3)I7(7) or later For 3000 series, 9.2(4) or later | - |
| | Cisco FC: Cisco MDS 9132T | 8.4(1a) or later | - |
| | Hypervisor | VMware vSphere ESXi 6.7 U2 or later | - |
| Storage | Hypervisor management system | VMware vCenter Server 6.7 U3 (vCSA) or later | - |
| Network | NetApp Virtual Storage Console (VSC) | VSC 9.7 or later | - |
| | NetApp SnapCenter | SnapCenter 4.3 or later | - |
| | Cisco UCS Manager | 4.1(1c) or later | |
| Hypervisor | ESXi | | |
| Management | Hypervisor management system VMware vCenter Server 6.7 U3 (vCSA) or later | | |
| | NetApp Virtual Storage Console (VSC) | VSC 9.7 or later | |
| | NetApp SnapCenter | SnapCenter 4.3 or later | |
| | Cisco UCS Manager | 4.1(1c) or later | |

Next: [Additional FlexPod security considerations.](#)

Additional FlexPod security considerations

[Previous: Solution benefits of FlexPod converged infrastructure.](#)

The FlexPod infrastructure is a modular, converged, optionally virtualized, scalable (scale out and scale up), and cost-effective platform. With the FlexPod platform, you can independently scale out compute, network, and storage to accelerate your application deployment. And the modular architecture enables nondisruptive operations even during your system scale-out and upgrade activities.

Different components of an HIT system require data to be stored in SMB/CIFS, NFS, Ext4, and NTFS file systems. This requirement means that the infrastructure must provide data access over the NFS, CIFS, and SAN protocols. A single NetApp storage system can support all these protocols, eliminating the need for the legacy practice of protocol-specific storage systems. Additionally, a single NetApp storage system can support multiple HIT workloads such as EHRs, PACS or VNA, genomics, VDI, and more, with guaranteed and configurable performance levels.

When deployed in a FlexPod system, HIT delivers several benefits that are specific to the healthcare industry. The following list is a high-level description of these benefits:

- **FlexPod security.** Security is at the very foundation of a FlexPod system. In the past few years, ransomware has become a threat. Ransomware is a type of malware that is based on cryptovirology, the use of cryptography to build malicious software. This malware can use both symmetric and asymmetric key encryption to lock a victim's data and demand a ransom to provide the key to decrypt the data. To learn how the FlexPod solution helps mitigate threats like ransomware, see [TR-4802: The Solution to Ransomware](#). FlexPod infrastructure components are also [FIPS 140-2-compliant](#).
- **Cisco Intersight.** Cisco Intersight is an innovative, cloud-based, management-as-a-service platform that provides a single pane of glass for full-stack FlexPod management and orchestration. The Intersight platform uses FIPS 140-2 security-compliant cryptographic modules. The platform's out-of-band management architecture makes it out of scope for some standards or audits such as HIPAA. No individual identifiable health information on the network is ever sent to the Intersight portal.
- **NetApp FPolicy technology.** NetApp FPolicy (an evolution of the name file policy) is a file-access notification framework for monitoring and to managing file access over the NFS or SMB/CIFS protocols. This technology has been part of the ONTAP data management software for more than a decade—it is useful in helping detect ransomware. This Zero Trust engine provides extra security measures beyond permissions in access control lists (ACLs). FPolicy has two modes of operation: native and external:
 - Native mode provides both blacklisting and whitelisting of file extensions.
 - External mode has the same capabilities as native mode, but it also integrates with an FPolicy server that runs externally to the ONTAP system as well as a security information and event management (SIEM) system. For more information about how to fight ransomware, see the [Fighting Ransomware: Part Three – ONTAP FPolicy, Another Powerful Native \(aka Free\) Tool](#) blog.
- **Data at rest.** ONTAP 9 and later has three FIPS 140-2-compliant, data-at-rest encryption solutions:
 - NSE is a hardware solution that uses self-encrypting drives.
 - NVE is a software solution that enables encryption of any data volume on any drive type where it is enabled with a unique key for each volume.
 - NAE is a software solution that enables encryption of any data volume on any drive type where it is enabled with unique keys for each aggregate.



Starting with ONTAP 9.7, NAE and NVE are enabled by default if the NetApp NVE license package with name VE is in place.

- **Data in flight.** Starting with ONTAP 9.8, Internet Protocol security (IPsec) provides end-to-end encryption support for all IP traffic between a client and an ONTAP SVM. IPsec data encryption for all IP traffic includes NFS, iSCSI, and SMB/CIFS protocols. IPsec provides the only encryption in flight option for iSCSI traffic.
- **End-to-end data encryption across a hybrid, multicloud data fabric.** Customers who use data-at-rest encryption technologies such as NSE or NVE and Cluster Peering Encryption (CPE) for data replication traffic can now use end-to-end encryption between client and storage across their hybrid multicloud data fabric by upgrading to ONTAP 9.8 or later and using IPsec. Beginning with ONTAP 9, you can enable the FIPS 140-2 compliance mode for cluster-wide control plane interfaces. By default, the FIPS 140-2-only mode is disabled. Starting with ONTAP 9.6, CPE provides TLS 1.2 AES-256 GCM encryption support for ONTAP data replication features such as NetApp SnapMirror, NetApp SnapVault, and NetApp FlexCache technologies. Encryption is setup by way of a pre-shared key (PSK) between two cluster peers.
- **Secure multitenancy.** Supports the increased needs of virtualized server and storage shared infrastructure, enabling secure multitenancy of facility-specific information, particularly when hosting multiple instances of databases and software.

Next: [Conclusion](#).

Conclusion

Previous: [Additional FlexPod security considerations](#).

By running your healthcare application on a FlexPod platform, your healthcare organization is better protected by a FIPS 140-2-enabled platform. FlexPod offers multilayered protection at every single component: compute, network and storage. FlexPod data protection capabilities protect data at rest or in flight, and keep backups safe and ready when needed.

Avoid human errors by leveraging the FlexPod prevalidated designs that are rigorously tested converged infrastructures from the strategic partnership of Cisco and NetApp. A FlexPod system engineered and designed to deliver predictable, low-latency system performance and high availability with little impact, even when FIPS 140-2 is enabled in the compute, networking, and storage layers. This approach results in a superior user experience and optimal response time for users of your HIT system.

Next: [Acknowledgements, version history, and where to find additional information](#).

Acknowledgements, version history, and where to find additional information

Previous: [Conclusion](#).

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

- Cisco MDS 9000 Family NX-OS Security Configuration Guide

https://www.cisco.com/c/en/us/td/docs/switches/datacenter/mds9000/sw/8_x/config/security/cisco_mds9000_security_config_guide_8x/configuring_fips.html#task_1188151

- Cisco Nexus 9000 Series NX-OS Security Configuration Guide, Release 9.3(x)
<https://www.cisco.com/c/en/us/td/docs/switches/datacenter/nexus9000/sw/93x/security/configuration/guide/b-cisco-nexus-9000-nx-os-security-configuration-guide-93x/m-configuring-fips.html>
- NetApp and Federal Information Processing Standard (FIPS) Publication 140-2
<https://www.netapp.com/company/trust-center/compliance/fips-140-2/>
- FIPS 140-2
<https://fieldportal.netapp.com/content/902303>
- NetApp ONTAP 9 Hardening Guide
<https://www.netapp.com/us/media/tr-4569.pdf>
- NetApp Encryption Power Guide
<https://docs.netapp.com/ontap-9/index.jsp?topic=%2Fcom.netapp.doc.pow-nve%2Fhome.html>
- NVE and NAE Datasheet
<https://www.netapp.com/us/media/ds-3899.pdf>
- NSE Datasheet
<https://www.netapp.com/us/media/ds-3213-en.pdf>
- ONTAP 9 Documentation Center
<http://docs.netapp.com>
- NetApp and Federal Information Processing Standard (FIPS) Publication 140-2
<https://www.netapp.com/company/trust-center/compliance/fips-140-2/>
- Cisco and FIPS 140-2 Compliance
<https://www.cisco.com/c/en/us/solutions/industries/government/global-government-certifications/fips-140.html>
- NetApp Cryptographic Security Module
<https://csrc.nist.gov/csrc/media/projects/cryptographic-module-validation-program/documents/security-policies/140sp2648.pdf>
- Cybersecurity practices for medium and large healthcare organizations
<https://www.phe.gov/Preparedness/planning/405d/Documents/tech-vol2-508.pdf>
- Cisco and Cryptographic Module Validation Program (CMVP)
<https://csrc.nist.gov/projects/cryptographic-module-validation-program/validated-modules/search?SearchMode=Basic&Vendor=cisco&CertificateStatus=Active&ValidationYear=0>

- NetApp Storage Encryption, NVMe Self-Encrypting Drives, NetApp Volume Encryption, and NetApp Aggregate Encryption
<https://www.netapp.com/pdf.html?item=/media/17073-ds-3898.pdf>
- NetApp Volume Encryption and NetApp Aggregate Encryption
<https://www.netapp.com/pdf.html?item=/media/17070-ds-3899.pdf>
- NetApp Storage Encryption
<https://www.netapp.com/pdf.html?item=/media/7563-ds-3213-en.pdf>
- FlexPod for Electronic Health Record Systems
<https://www.netapp.com/pdf.html?item=/media/22199-tr-4881.pdf>
- Data Now: Improving Performance in Epic EHR Environments with Cloud-Connected Flash Technology
<https://www.netapp.com/media/10809-cloud-connected-flash-wp.pdf>
- FlexPod Datacenter for Epic EHR Infrastructure
<https://www.netapp.com/pdf.html?item=/media/17061-ds-3683.pdf>
- FlexPod Datacenter for Epic EHR Deployment Guide
<https://www.netapp.com/media/10658-tr-4693.pdf>
- FlexPod Datacenter Infrastructure for MEDITECH Software
<https://www.netapp.com/media/8552-flexpod-for-meditech-software.pdf>
- The FlexPod Standard Extends to MEDITECH Software
<https://blog.netapp.com/the-flexpod-standard-extends-to-meditech-software/>
- FlexPod for MEDITECH Directional Sizing Guide
<https://www.netapp.com/pdf.html?item=/media/12429-tr4774.pdf>
- FlexPod for medical imaging
<https://www.netapp.com/media/19793-tr-4865.pdf>
- AI in Healthcare
<https://www.netapp.com/us/media/na-369.pdf>
- FlexPod for healthcare Ease Your Transformation
<https://flexpod.com/solutions/verticals/healthcare/>
- FlexPod from Cisco and NetApp
<https://flexpod.com/>

Acknowledgements

- Abhinav Singh, Technical Marketing Engineer, NetApp
- Brian O'Mahony, Solution Architect Healthcare (Epic), NetApp
- Brian Pruitt, Pursuit Business Development Manager, NetApp
- Arvind Ramakrishnan, Senior Solutions Architect, NetApp
- Michael Hommer, FlexPod Global Field CTO, NetApp

Version History

| Version | Date | Document version history |
|-------------|------------|--------------------------|
| Version 1.0 | April 2021 | Initial release |

Cisco Intersight with NetApp ONTAP Storage

Cisco Intersight with NetApp Storage Quick Start Guide

In partnership with:



Introduction

NetApp and Cisco have partnered to provide Cisco Intersight, a single-pane view of the FlexPod ecosystem. This simplified integration creates a unified management platform for all components in the FlexPod infrastructure and FlexPod solution. Cisco Intersight allows you to monitor NetApp storage, Cisco compute, and VMware inventory. It also allows you to orchestrate or automate workflows to accomplish storage and virtualization tasks in tandem.

Related information

To learn more, see the following documents and websites:

[TR 4883: FlexPod Datacenter with ONTAP 9.8, ONTAP Storage Connector for Cisco Intersight, and Cisco Intersight Managed Mode](#)

[Cisco Intersight help center](#)

[Cisco Intersight Getting Started Overview](#)

[Intersight Appliance Install and Upgrade Guide](#)

What's new

This section lists new features and functionality available for Cisco Intersight with NetApp ONTAP storage.

January 2024

- NetApp storage orchestration using reference workflows now available for download in GitHub through the [FlexPod Intersight Workflow repository](#). For more information on the new reference workflows in GitHub, see [Use case 2: NetApp storage orchestration using Reference Workflows](#).

November 2023

- Added NVMe Namespaces page under the Inventory section of the user interface.

August 2023



An upgrade to NetApp Active IQ Unified Manager 9.13GA is required to ensure compatibility and full functionality with the latest release.

- Improved the New NetApp Smart LUN task to clearly indicate the availability of selection options for either creating a new initiator group or selecting an existing initiator group. When users now select the box to create a new initiator group, the parameter for choosing an existing initiator group is no longer available. If users deselect the box to create a new initiator group, the existing initiator group parameter then becomes available.
- Enhanced the New NetApp LUN Map and Remove NetApp LUN Map tasks. The new relationship between the LUN and initiator group is now updated. The UI inventory is immediately updated for both the LUN and the initiator group upon task execution.
- The Checks page now loads properly the first-time users log in and no longer requires a refresh.

July 2023



An upgrade to NetApp Active IQ Unified Manager 9.13GA is required to ensure compatibility and full functionality with the latest release.

- Updated names for NetApp storage tasks. See Use case 3 Custom workflows using designer-free form for the complete list of renamed tasks.
- NFS Interface IP address was added as an output of the New NetApp NAS Smart Volume task.
- A check that ASUP transport is HTTPS was added to the Checks tab.
- The correct tier type for all tiers now displays properly under the Tiers user interface.
- All compliant licenses now display properly under the Licenses page.
- Accurate value for CIFS shares without or without a home directory now display on the Shares page.
- Sorting and filtering now enabled for the Mapped column on the LUNS page.
- Sorting and filtering now enabled the Authentication Enabled column on the NTP Servers page.
- Added new checks and the following corresponding categories to the Checks tab.
 - Security
 - Anti-Ransomware
 - Availability
 - Other
- Under the Inventory detail view, report now used instead of physical used capacity.

June 2023



An upgrade to NetApp Active IQ Unified Manager 9.13RC1 is required to ensure compatibility and full functionality with the latest release.

- Updated names for NetApp storage tasks. See [Use case 3 Custom workflows using designer-free form](#) for the complete list of renamed tasks.

April 2023

- Added Protection Policies (SnapMirror) and Snapshot Policies tabs under the Policies page within the Inventory section of the user interface.
- Added NFS Clients page under the Inventory section of the user interface.

- Added Protected column in the Storage VMs page under the Inventory section of the user interface.
- Modified how Data Reduction information is reported and displayed.
- Added Local Tier and Cloud Tier tabs under the Tiers page within the Inventory section of the user interface.
- Node column now displays after the Name column under the Ports page within the Inventory section of the user interface.

January 2023



An upgrade to NetApp Active IQ Unified Manager 9.12 GA is required to ensure compatibility and full functionality with the latest release. For a list of known issues related to this release, see [Known Issues](#).

- Intersight interoperability checks can now distinguish between UCSM and IMM firmware modes when carrying out compatibility checks.
- Protection Relationships will not display in Intersight for ONTAP 9.7. This issue was fixed in ONTAP 9.8RC1.

August 2022



An upgrade to NetApp Active IQ Unified Manager 9.11 GA is required to ensure compatibility and full functionality with the latest release. For a list of known issues related to this release, see [Known Issues](#).

- Updated cluster available capacity calculation to match System Manager
- Updated cluster General page to hide the performance metrics summary until performance data is populated
- Fixed cluster General page UI issue that occasionally caused the page to hang
- Added CIFS shares, CIFS services, Qtrees, and SVM SnapMirror policies to backend inventory.
- Added Shares and Qtrees to the UI navigation menu under the Logical inventory section
- Added Shares as a tab from a selected Storage VM
- Added CIFS Service information on the Storage VM General tab if the Storage VM is CIFS enabled
- Added a cluster Checks page that allow users to validate the configuration of NetApp storage systems adhere to best practices

July 2022

- Improved visuals for Cluster Data Reduction ratio now available under the Capacity Widget
- Added FC Interfaces tab to the Network Interfaces page
- Creating a new volume using the generic “New Storage Volume” task now sets volume space guarantee to none and snapshot reserve percent to 0%
- Comment field under the Edit Snapshot Policy task now optional and no longer mandatory
- Improved UI inventory and orchestration consistency
- Intersight capacity information under Cluster Capacity now consistent with System Manager

- Added checkbox under New Storage Virtual Machine task to display all parameters when creating a new management interface to improve usability
- Moved Protocols below Client Match, now consistent with System Manager
- Export policy general page now displaying Access Protocol(s)
- igroup removal now conditionally logged
- Added “Failover Policy” and “autorevert” parameters for NAS under New Storage NAS Data Interface and New Storage iSCSI Data Interface
- Rollback for New Storage NAS Smart Volume task now removes export policy if no other volumes are attached
- Made enhancements for Smart Volume and Smart LUN tasks

April 2022



To ensure compatibility and complete functionality with future releases, it is recommended that you upgrade your NetApp Active IQ Unified Manager to version 9.10P1.

- Added Broadcast Domain to Ethernet Port Detail page
- Changed the term “Aggregate” to “Tier” for the Aggregate and SVM within the user interface
- Changed the term “Cluster Status” to “Array Status”
- MTU filter now works for <, >, =, <=, >= characters
- Added Network Interface Page to Cluster Inventory
- Added AutoSupport to Cluster Inventory
- Added `cdpd.enable` option to node
- Added an object for CDP neighbor
- Added NetApp workflow storage tasks within Cisco Intersight. See [Use case 3 Custom workflows using designer-free form](#) for a complete list of NetApp storage tasks.

January 2022

- Added event-based Intersight alarms for NetApp Active IQ Unified Manager 9.10 or above.



To ensure compatibility and complete functionality with future releases, it is recommended that you upgrade your NetApp Active IQ Unified Manager to version 9.10.

- Explicitly set each protocol enabled (true or false) for Storage Virtual Machine
- Mapped clusterHealthStatus state ok-with-suppressed to OK
- Renamed Health column to Cluster Status column under the Cluster list page
- Showing storage array “Unreachable” if the cluster is down or otherwise unreachable
- Renamed Health column to Array Status column under the Cluster General page
- SVM now has a “Volumes” tab that shows all the volumes for the SVM
- Volume has a snapshot capacity section
- Licenses now display correctly

October 2021

- Updated list of NetApp storage tasks available within Cisco Intersight. See [Use case 3 Custom workflows using designer-free form](#) for a complete list of NetApp storage tasks.
- Added Health column under the Cluster list page.
- Expanded details now available under the General page for a selected cluster.
- NTP Server table now accessible through the navigation pane.
- Added a new Sensors tab containing the General page for the Storage Virtual Machine.
- VLAN and link aggregation group summary now available under the Port General page.
- Total Data Capacity column added under the Volume Total Capacity table.
- Latency, IOPS, and Throughput columns added under Average Volume Statistics, Average LUN Statistics, Average Aggregate Statistics, Average Storage VM Statistics, and Average Node Statistics tables



The above performance metrics are only available for storage arrays monitored through NetApp Active IQ Unified Manager 9.9 or above.

Known Issues

- If you are using a version of AIQUM 9.11 or earlier, a discrepancy will occur between the displayed values on the Storage List page and capacity bar chart on the Storage general page. To resolve this issue, upgrade to AIQUM 9.12 or greater to ensure the accuracy of the displayed capacity values.
- If you are using AIQUM 9.11 or earlier, any checks performed by the "Interoperability" tab under the "Integrated Systems" page will fail to distinguish IMM and UCSM Cisco components accurately. To resolve this issue, upgrade to AIQUM 9.12 to ensure all components are properly identified.
- To ensure Intersight storage inventory data is unaffected during the data collection process, any unsupported ONTAP clusters (i.e., versions below ONTAP 9.7P1) must be removed from the Active IQ Unified Manager (AIQUM).
- All claimed targets require a minimum AIQUM version of 9.11 for FlexPod Integrated System Interoperability queries to complete successfully.
- The Storage Inventory Checks page will not populate if the ONTAP cluster is added to AIQUM using an FQDN. Users must add ONTAP clusters to AIQUM using an IP address.

Requirements

Verify that you meet the hardware, software, and licensing requirements for NetApp ONTAP storage integration with Cisco Intersight.

Hardware and software requirements

These are the minimum hardware and software components required to implement the solution. The components that are used in any particular implementation of the solution might vary based on customer requirements.

| Component | Requirement details |
|----------------------------------|--|
| NetApp ONTAP | ONTAP 9.7P1 and later |
| NetApp Active IQ Unified Manager | Latest version of NetApp Active IQ Unified Manager is required (currently 9.14RC1) |
| NetApp Storage Array | All ONTAP ASA, AFF, and FAS storage arrays supported for ONTAP 9.7P1 and later |
| Virtualization Hypervisor | vSphere 7.0 and later |



Refer to [Cisco Intersight supported systems](#) for the minimum requirements of Cisco UCS Compute Components and UCSM version.

Cisco Intersight licensing requirements

Cisco Intersight offers services such as Infrastructure Service and Cloud Orchestrator service to manage, automate, and optimize physical storage (NetApp storage). You can use these services to manage Cisco UCS server and Cisco HyperFlex system. The Infrastructure Service and Cloud Orchestrator service use a subscription-based licensing model with multiple tiers. You can choose the required Cisco UCS Server volume tier for the selected subscription term.

Licensing Model

The Cisco Intersight Infrastructure Services licensing model has been simplified and now offers the following two tiers:

- **Cisco Intersight Infrastructure Services Essentials** - The Essentials license tier offers server management including global health monitoring functionality, inventory, proactive support through Cisco TAC integration, multi-factor authentication, along with providing SDK and API access.
- **Cisco Intersight Infrastructure Services Advantage** - The Advantage license tier offers advanced server management with extended visibility, ecosystem integration, automation of Cisco and third-party hardware and software, along with providing multi-domain solutions.

For more information about the features covered by various licensing tiers, go to [Infrastructure Services license](#).

Before you begin

To monitor and orchestrate NetApp storage from Cisco Intersight, you need NetApp Active IQ Unified Manager and Cisco Intersight Assist Virtual Appliance installed in the vCenter environment.

Install or Upgrade NetApp Active IQ Unified Manager

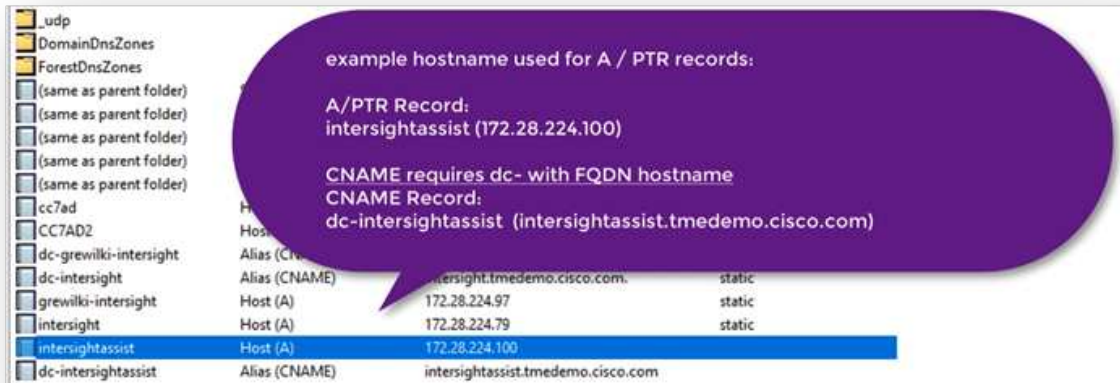
Install or upgrade to Active IQ Unified Manager (latest version is required, currently 9.14RC1) if you have not done so. For instructions, go to the [NetApp Active IQ Unified Manager Documentation](#).

Install Cisco Intersight Assist Virtual Appliance

Ensure that you meet the [Cisco Intersight Virtual Appliance Licensing, System, and Network requirements](#).

Steps

1. Create a Cisco Intersight Account.
Visit <https://intersight.com/> to create your Intersight account. You must have a valid Cisco ID to create a Cisco Intersight account.
2. Download the Intersight Virtual Appliance at software.cisco.com. For more information, go to the [Intersight Appliance Install and Upgrade Guide](#).
3. Deploy the OVA. DNS and NTP are required to deploy the OVA.
 - a. Configure DNS with A/PTR and CNAME Alias records prior to deploying the OVA. See the example below.



example hostname used for A / PTR records:

A/PTR Record:
intersightassist (172.28.224.100)

CNAME requires dc- with FQDN hostname
CNAME Record:
dc-intersightassist (intersightassist.tmedemo.cisco.com)

| | | | |
|-------------------------|---------------|------------------------------------|--------|
| _udp | | | |
| DomainDnsZones | | | |
| ForestDnsZones | | | |
| (same as parent folder) | | | |
| (same as parent folder) | | | |
| (same as parent folder) | | | |
| (same as parent folder) | | | |
| (same as parent folder) | | | |
| cc7ad | Host (A) | 172.28.224.97 | static |
| CC7AD2 | Host (A) | 172.28.224.79 | static |
| dc-grewilki-intersight | Alias (CNAME) | intersight.tmedemo.cisco.com. | static |
| dc-intersight | Alias (CNAME) | intersight.tmedemo.cisco.com. | static |
| grewilki-intersight | Host (A) | 172.28.224.97 | static |
| intersight | Host (A) | 172.28.224.79 | static |
| intersightassist | Host (A) | 172.28.224.100 | static |
| dc-intersightassist | Alias (CNAME) | intersightassist.tmedemo.cisco.com | static |

- b. Choose the appropriate configuration size (Tiny, Small, or Medium) based on your OVA deployment requirements for Intersight Virtual Appliance.

TIP: For a two-node ONTAP cluster with a large number of storage objects, NetApp recommends that you use the Small (16 vCPU, 32 Gi RAM) option.

Deploy OVF Template

✓ 1 Select an OVF template

✓ 2 Select a name and folder

✓ 3 Select a compute resource

✓ 4 Review details

5 Configuration

6 Select storage

7 Select networks

8 Customize template

9 Ready to complete

Configuration

Select a deployment configuration

| | |
|--|---|
| <input type="radio"/> Small(16 vCPU, 32 Gi RAM) | Description Deployment size supports Intersight Assist only |
| <input type="radio"/> Medium(24 vCPU, 64 Gi RAM) | |
| <input checked="" type="radio"/> Tiny(8 vCPU, 16 Gi RAM) | |

3 items

CANCEL

BACK

NEXT

- c. On the **Customize Template** page, customize the deployment properties of the OVF template. The administrator password is used for the local users: admin(webUI/cli/ssh).

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 Configuration
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Customize template**
- 9 Ready to complete

Customize template

Customize the deployment properties of this software solution.

✓ All properties have valid values

| Uncategorized | 8 settings |
|-----------------|---|
| Enable DHCP | Use DHCP for networking. All static params will be ignored. <input type="checkbox"/> |
| IP Address | IPv4 address (Must have PTR record in your DNS) <input type="text"/> |
| Net Mask | IPv4 Network Mask <input type="text" value="255.255.255.0"/> |
| Default Gateway | IPv4 Default Gateway <input type="text"/> |
| DNS Domain | DNS Search Domain <input type="text"/> |
| DNS Servers | Comma-separated list of DNS servers <input type="text"/> |

CANCEL

BACK

NEXT

Deploy OVF Template

- ✓ 1 Select an OVF template
- ✓ 2 Select a name and folder
- ✓ 3 Select a compute resource
- ✓ 4 Review details
- ✓ 5 Configuration
- ✓ 6 Select storage
- ✓ 7 Select networks
- 8 Customize template**
- 9 Ready to complete

| | |
|------------------------|---|
| Net Mask | IPv4 Network Mask |
| | 255.255.255.0 |
| Default Gateway | IPv4 Default Gateway |
| | |
| DNS Domain | DNS Search Domain |
| | |
| DNS Servers | Comma-separated list of DNS servers |
| | |
| Administrator password | Password for local admin account |
| | Password |
| | Confirm Password |
| NTP Server | Comma-separated list of NTP servers. If no servers are provided, NIST servers will be configured. |
| | |

CANCEL
BACK
NEXT

d. Click **Next**.

4. Post-deploy the Intersight Assist Appliance.

- a. Navigate to <https://FQDN-of-your-appliance> to complete the post-install set-up of your appliance.

The installation process automatically begins. Installation can take up to one hour depending on bandwidth to Intersight.com. It can also take several seconds for the secure site to be operational after the VM powers on.

- b. During the post-deployment process, select the following option:

- **Intersight Assist.** This deployment enables SaaS model to connect to Cisco Intersight.



When selecting Intersight Assist, take note of the device ID and claim code before you continue.

What would you like to Install ?

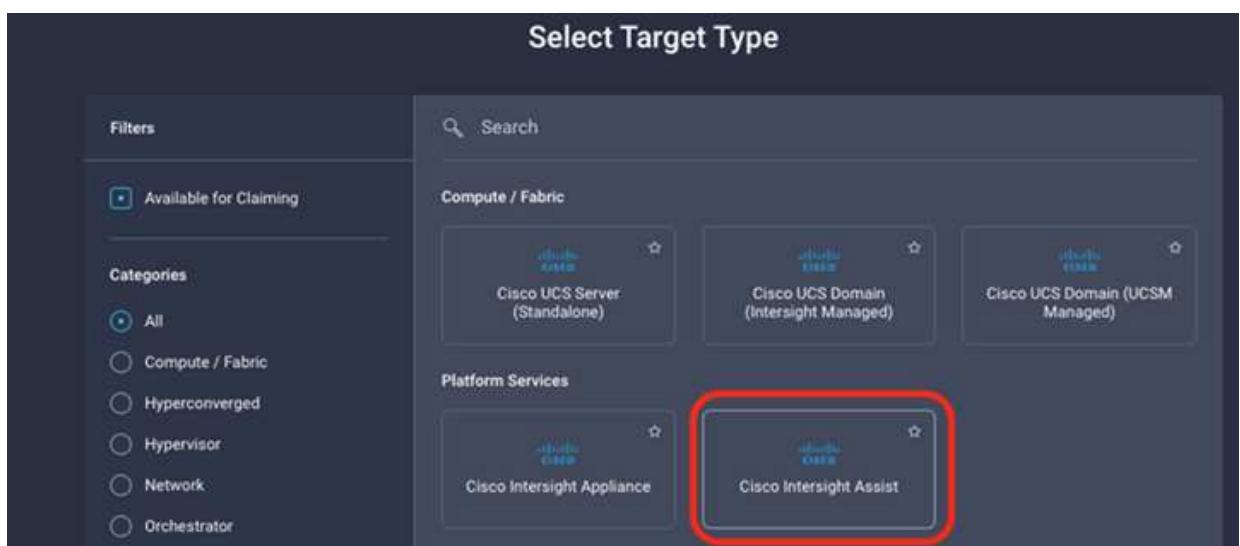
☐ Intersight Connected Virtual Appliance

☐ Intersight Private Virtual Appliance

☐ Intersight Assist

[Recover from backup](#) Proceed

- c. Click **Proceed**.
- d. Select **Intersight Assist** and complete the following steps:
 - i. Navigate to your SaaS Intersight account at <https://intersight.com>.
 - ii. Click **Targets**, **Cisco Intersight Assist**, and then **Start**.
 - iii. Claim the **Cisco Intersight Assist** appliance by copying and pasting the device ID and claim code from your newly deployed Intersight Assist virtual appliance.



- iv. Return to the **Cisco Intersight Assist** appliance and click **Continue**. You might need to refresh the browser.

The download and installation process begins. The binaries are transferred from Intersight Cloud to your on-prem appliance. Completion time varies depending on your bandwidth to the Intersight Cloud.

Configure AIQ UM proxy server for IMT service

If you are using a proxy server with AIQ UM for Cisco Intersight with NetApp ONTAP storage, you must configure the setup through the command line interface (CLI) to utilize the interoperability matrix tool service (IMT). The IMT service is available under the **Interoperability** tab of the **Integrated Systems** page. You must use the Active IQ Unified Manager Virtual Machine (OVA) Diag shell to configure the AIQ UM proxy server settings.



For information on how to access the AIQ UM Diag shell, see [How to access Active IQ Unified Manager Virtual Machine \(OVA\) DIAG shell](#)

Steps

1. Log into the AIQ UM terminal and run the following command to log into um.

```
um cli login -u <um maintenance user name>
```

Example

```
um cli login -u admin
```

2. Set the `imt_proxy_host` and `imt_proxy_port` by running the following commands.



The IMT proxy is a separate configuration from AutoSupport (ASUP) proxy configurations.

```
um option set imt.https.proxy.host=<IMT_PROXY_HOST>
um option set imt.https.proxy.port=<IMT_PROXY_PORT>
```

Example

```
um option set imt.https.proxy.host=example-proxy.cls.eng.com
um option set imt.https.proxy.port=8200
```



IMT proxy server configurations do not support authentication.

3. View the IMT proxy details to verify the `proxy_host` and `proxy_port` settings through the following command.

```
um option list |grep imt
```


Claim targets

After Cisco Intersight Assist is installed, you can claim your NetApp storage and virtualization devices. Return to the **Intersight Targets** page and add your vCenter and NetApp Active IQ Unified Manager targets. For more information about the claim process, watch the video [Claim a Target via Cisco Intersight Assist](#).



Make sure that the NetApp Active IQ Unified Manager (AIQ UM) API Gateway is enabled.


From the NetApp IQ Unified Manager, navigate to **Settings > General > Feature Settings**.



The following example shows the NetApp AIQ UM target being claimed from Cisco Intersight.



When you claim the NetApp AIQ UM target, all clusters managed by Active IQ Unified Manager are automatically added to Intersight.



NetApp Active IQ Unified Manager

To claim any on-premises target an Intersight Assist Appliance is required. Deploy and claim an Assist Appliance if needed before claiming the target

This target is intended for the functionality of Intersight Orchestrator

Intersight Assist *
isassist.cie.netapp.com

Hostname/IP Address *
NTAPAIQUM.fp.netapp.com

Username *
admin

Password *

☒ Secure

Monitor NetApp storage from Cisco Intersight

After targets are claimed, NetApp storage widgets, storage inventory, and virtualization tabs become available if you have an Advantage tier license. Orchestration tabs are available if you have a Premier tier license.

Storage inventory overview

The following screenshot displays the **Operate > Storage** screen.

OPERATE > Storage

The Trial period for Intersight is active. During the Trial period, the Premier tier features of Intersight are available. [Go to Licensing](#)

* All Storage

Export 8 items found 10 per page 1 of 1

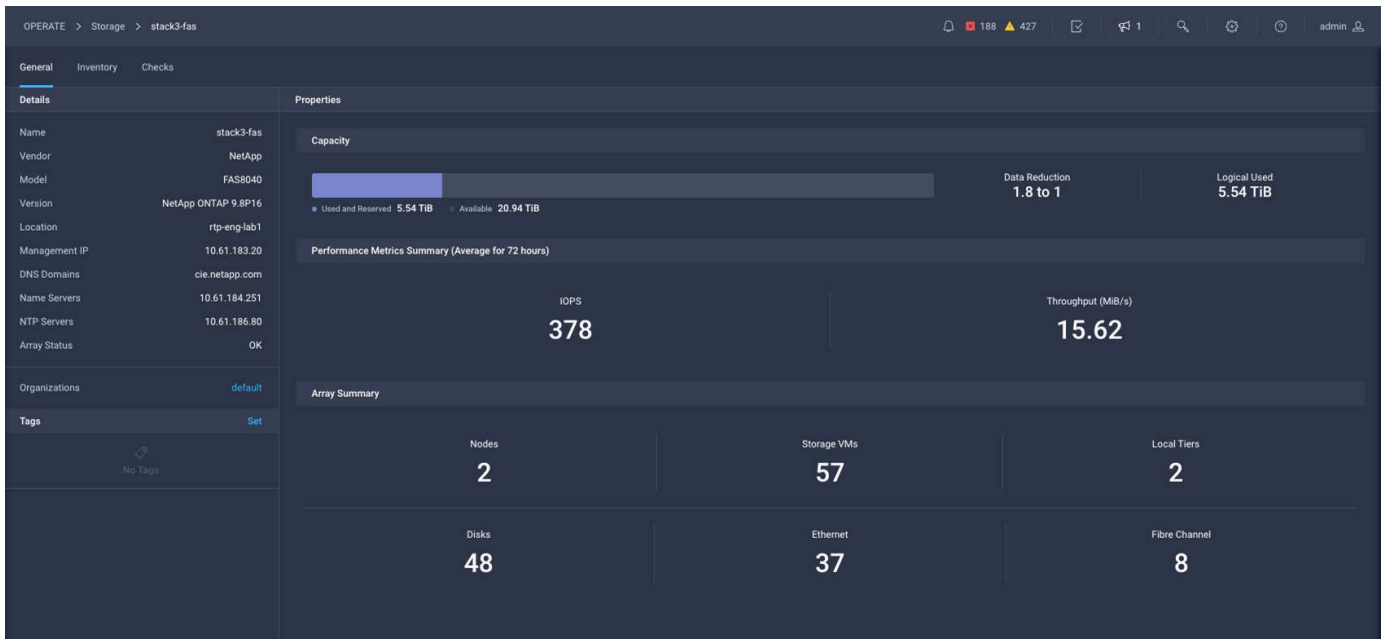
| Name | Vendor | Model | Version | Capacity | Capacity Utilization |
|------------------|--------|----------|----------------------|------------|----------------------|
| stack1-fas | NetApp | FAS2552 | NetApp ONTAP 9.7P8 | 27.61 TiB | 98.5% |
| aaron | NetApp | FAS8020 | NetApp ONTAP 9.8X28 | 1.76 TiB | 46.7% |
| cle-na2750-g1344 | NetApp | FAS2750 | NetApp ONTAP 9.7P8 | 104.34 TiB | 98.8% |
| stack3-fas | NetApp | FAS8040 | NetApp ONTAP 9.7P8 | 38.73 TiB | 40.6% |
| AFF8060-51-130 | NetApp | AFF8060 | NetApp ONTAP 9.8X22 | 3.77 TiB | 0.1% |
| niefas2650 | NetApp | FAS2650 | NetApp ONTAP 9.7P8 | 3.24 TiB | 0.0% |
| a220-f0234 | NetApp | AFF-A220 | NetApp ONTAP 9.9.1P1 | 5.77 TiB | 7.1% |
| rajeshcluster-1 | NetApp | SIMBOX | NetApp ONTAP 9.8.0 | 9.93 GiB | 0.1% |

1 of 1

The following screenshot shows the storage cluster overview.



The following performance metric summary information will only display if the storage array is monitored through NetApp Active IQ Unified Manager 9.9 or above.





Storage widgets

To view storage widgets, navigate to **Monitoring > Dashboards > View NetApp storage widgets**.

- The following screenshot shows the Storage Version Summary widget.



- This screenshot shows the Top 5 Storage Arrays by Capacity Utilization widget.

| Top 5 Storage Arrays by Capacity Utilization   | | | | |
|--|---------------------|--------|-----------|------------------------------|
| # | Name | Vendor | Capacity | Utilization |
| 1 | Warriors_Controller | NetApp | 13.83 TiB | <div><div></div></div> 89.4% |
| 2 | stack3-fas | NetApp | 8.95 TiB | <div><div></div></div> 66.2% |
| 3 | aaron | NetApp | 4.71 TiB | <div><div></div></div> 44.1% |
| 4 | aff-a400 | NetApp | 40.62 TiB | <div><div></div></div> 0.2% |

- This screenshot shows the Top 5 Storage Volumes by Capacity Utilization widget.

| Top 5 Storage Volumes by Capacity Utilization | | | | |
|---|------------------|--------|------------|------------------------------|
| # | Name | Vendor | Capacity | Utilization |
| 1 | test_1_vol | NetApp | 10.31 GiB | <div><div></div></div> 98.6% |
| 2 | test_lun_vol | NetApp | 10.31 GiB | <div><div></div></div> 97.9% |
| 3 | vmware_server_1 | NetApp | 50.00 GiB | <div><div></div></div> 95.0% |
| 4 | vmware_server_2 | NetApp | 50.00 GiB | <div><div></div></div> 82.3% |
| 5 | VM_Datastore_vol | NetApp | 150.00 GiB | <div><div></div></div> 67.0% |

Use cases

These are a few use case examples for monitoring and orchestration of NetApp storage from Cisco Intersight.

Use case 1: Monitoring NetApp storage inventory and widgets

When the NetApp storage environment is available in Cisco Intersight, you can monitor NetApp storage objects in detail from storage inventory and get an overview from storage widgets.

1. Deploy Intersight Assist OVA (OnPrem task in vCenter Environment).
2. Add NetApp AIQ UM devices in Intersight Assist.
3. Go to **Storage** and navigate through NetApp storage inventory.
4. Add **Widgets** for NetApp storage to your **Monitor Dashboard**.

Here is a [link](#) to the video showing NetApp ONTAP Storage Monitoring Features from Cisco Intersight.

Use case 2: NetApp storage orchestration using reference workflows

When NetApp storage and vCenter environments are available in Cisco Intersight, you can use end-to-end reference workflows available in GitHub through the [FlexPod Intersight Workflow repository](#).

The reference workflows include storage and virtualization tasks. The README file for the repository provides the prerequisites needed for executing workflows, links to helpful resources (including documentation on how to import a workflow), and documentation links for each reference workflow.

Each workflow has a folder in the repository containing two files:

- The JSON file to download and import into Intersight,
- A documentation file that provides a view of the tasks in the workflow, workflow inputs, and an example execution of the workflow.

Perform the following to import and use a reference workflow:

1. Deploy Intersight Assist OVA (OnPrem task in vCenter Environment).
2. Add NetApp AIQ UM devices in Intersight Assist.
3. Add the vCenter target to Intersight via Intersight Assist.
4. Download the JSON file for a reference workflow from the FlexPod-Intersight-Workflow repository.
5. Import the workflow into Intersight, then execute the workflow.

Here is a list of workflows available in the GitHub FlexPod-Intersight-Workflow repository:

- Add Initiators to NetApp Initiator Group
- New Export Policy for NetApp Volume
- New NAS Datastore Using NetApp Smart Volume
- New NetApp FC Data Interface
- New NetApp Initiator Group

- New NetApp iSCSI Data Interface
- New NetApp NAS Data Interface
- New NetApp Storage Virtual Machine
- New VMFS Datastore Using NetApp Smart LUN
- Remove Initiators from NetApp Initiator Group
- Remove NAS Datastore Using NetApp Smart Volume
- Remove NetApp Export Policy
- Remove NetApp Initiator Group
- Remove VMFS Datastore Using NetApp Smart LUN
- Update NAS Datastore Using NetApp Smart Volume
- Update VMFS Datastore Using NetApp Smart LUN

Use case 3: Custom workflows using designer-free form

When the NetApp Storage and vCenter environments are available in Cisco Intersight, you can build custom workflows using the NetApp storage and virtualization tasks.

1. Deploy Intersight Assist OVA (OnPrem task in vCenter Environment)
2. Add NetApp AIQ UM devices in Intersight Assist.
3. Add vCenter target to Intersight via Intersight Assist.
4. Navigate to the **Orchestration** tab in Intersight.
5. Select **Create Workflow**.
6. Add storage and virtualization tasks to your workflows.

Here are the NetApp storage tasks that are available from Cisco Intersight:

- Add ACL to NetApp CIFS Share
- Add Client Match to NetApp Export Policy Rule
- Add Export Policy to NetApp Volume
- Add Initiators to NetApp Initiator Group
- Add Rule to NetApp Export Policy
- Add Schedule to NetApp Snapshot Policy
- Confirm NetApp License Status
- Confirm NetApp Storage Virtual Machine FCP Protocol Status
- Edit NetApp Aggregates for Storage Virtual Machine
- Edit NetApp Asynchronous SnapMirror Policy
- Edit NetApp CIFS Share ACL Permission
- Edit NetApp Export Policy Rule
- Edit NetApp Snapshot Policy
- Edit NetApp Snapshot Policy Schedule
- Edit NetApp Volume Security Style

- Edit NetApp Volume Snapshot Policy
- Enable NetApp CIFS Services
- Expand NetApp LUN
- New NetApp Asynchronous SnapMirror Policy
- New NetApp CIFS Server
- New NetApp CIFS Share
- Find NetApp Initiator Group LUN Map
- Find NetApp LUN by ID
- Find NetApp Volume by ID
- New NetApp Export Policy
- New NetApp FC Data Interface
- New NetApp Initiator Group
- New NetApp iSCSI Data Interface
- New NetApp Load-Sharing Mirrors for SVM Root Volume
- New NetApp LUN
- New NetApp LUN Map
- New NetApp NAS Data Interface
- New NetApp NAS Smart Volume
- New NetApp Smart LUN
- New NetApp SnapMirror Relationship for Volume
- New NetApp Snapshot Policy
- New NetApp Storage Virtual Machine
- New NetApp Volume
- New NetApp Volume Snapshot
- Register DNS for NetApp Storage Virtual Machine
- Remove ACL from NetApp CIFS Share
- Remove Client Match from NetApp Export Policy Rule
- Remove Export Policy from NetApp Volume
- Remove Initiator from NetApp Initiator Group
- Remove NetApp CIFS Server
- Remove NetApp CIFS Share
- Remove NetApp Export Policy
- Remove NetApp FC Data Interface
- Remove NetApp Initiator Group
- Remove NetApp IP Interface
- Remove NetApp Load-Sharing Mirrors for SVM Root Volume
- Remove NetApp LUN

- Remove NetApp LUN Map
- Remove NetApp NAS Smart Volume
- Remove NetApp Smart LUN
- Remove NetApp Snapmirror Relationship for Volume
- Remove NetApp Snapmirror Policy
- Remove NetApp Snapshot Policy
- Remove NetApp Storage Virtual Machine
- Remove NetApp Volume
- Remove NetApp Volume Snapshot
- Remove Rule from NetApp Export Policy
- Remove Schedule from NetApp Snapshot Policy
- Rename NetApp Volume Snapshot
- Update NetApp Load-Sharing Mirrors for SVM Root Volume
- Update NetApp Volume Capacity

To learn more about customizing workflows with NetApp storage and virtualization tasks, watch the video [NetApp ONTAP Storage Orchestration in Cisco Intersight](#).

Infrastructure

End-to-End NVMe for FlexPod with Cisco UCSM, VMware vSphere 7.0, and NetApp ONTAP 9

TR-4914: End-to-End NVMe for FlexPod with Cisco UCSM, VMware vSphere 7.0, and NetApp ONTAP 9

Chris Schmitt and Kamini Singh, NetApp

In partnership with:



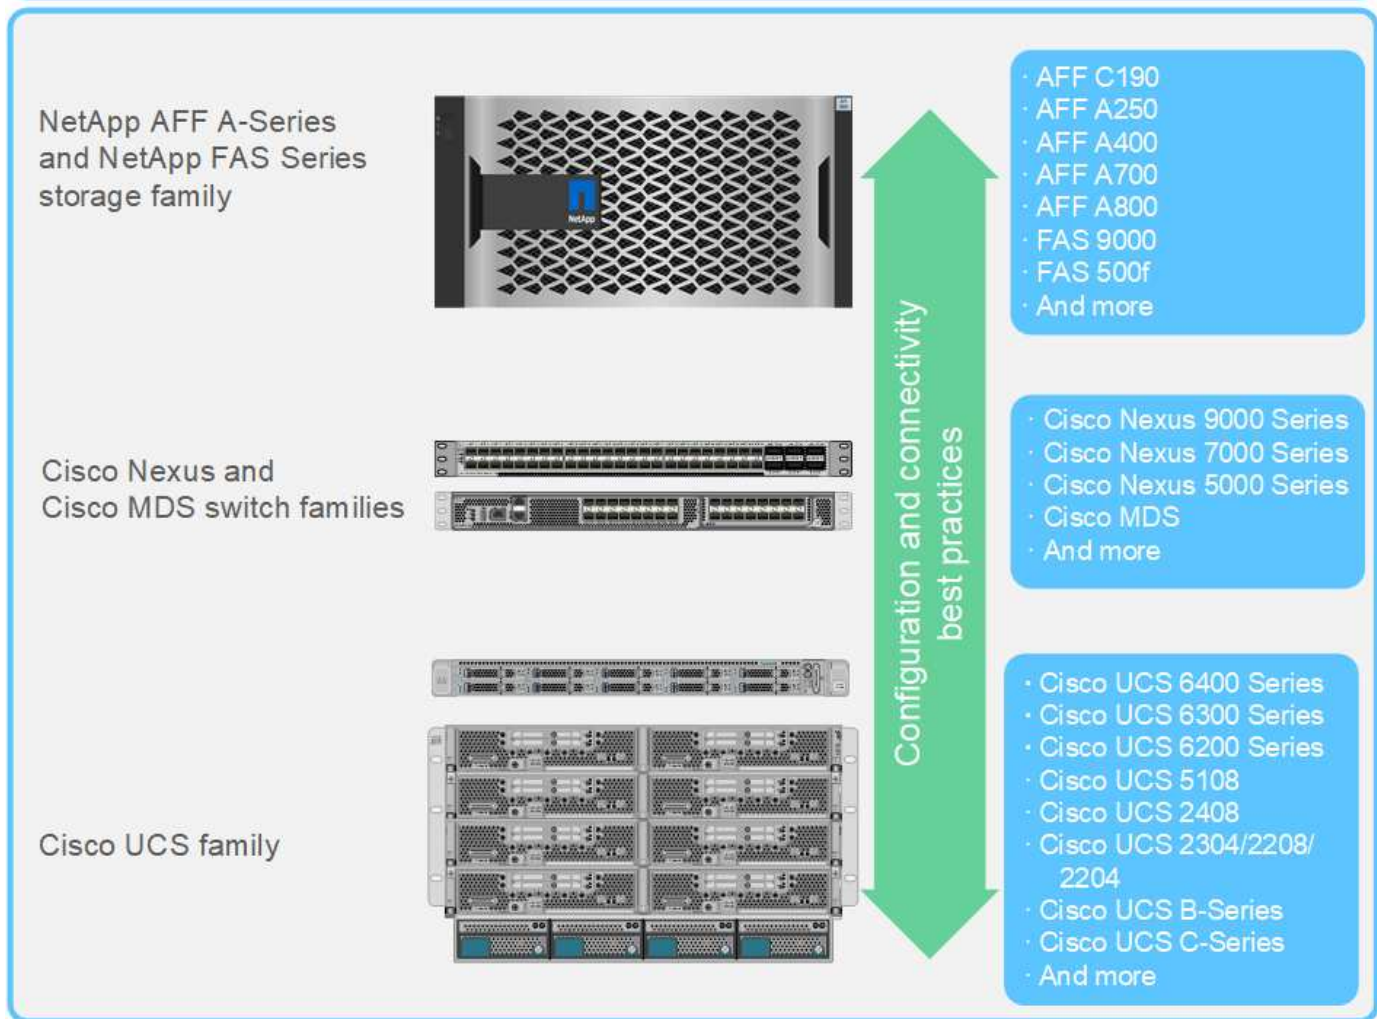
The NVMe data-storage standard, an emerging core technology, is transforming enterprise data storage access and transport by delivering very high bandwidth and very low latency storage access for current and future memory technologies. NVMe replaces the SCSI command set with the NVMe command set.

NVMe was designed to work with nonvolatile flash drives, multicore CPUs, and gigabytes of memory. It also takes advantage of the significant advances in computer science since the 1970s, enabling streamlined command sets that more efficiently parse and manipulate data. An end-to-end NVMe architecture also enables data center administrators to rethink the extent to which they can push their virtualized and containerized environments and the amount of scalability that their transaction-oriented databases can support.

FlexPod is a best-practice data center architecture that includes the Cisco Unified Computing System (Cisco UCS), Cisco Nexus switches, Cisco MDS switches, and NetApp AFF systems. These components are connected and configured according to the best practices of both Cisco and NetApp to provide an excellent platform for running a variety of enterprise workloads with confidence. FlexPod can scale up for greater performance and capacity (adding compute, network, or storage resources individually as needed), or it can scale out for environments that require multiple consistent deployments (such as rolling out of additional FlexPod stacks).

The following figure presents the FlexPod component families.

FlexPod Datacenter solution



FlexPod is the ideal platform for introducing FC-NVMe. It can be supported with the addition of the Cisco UCS VIC 1400 Series and Port Expander in existing Cisco UCS B200 M5 or M6 servers or Cisco UCS C-Series M5 or M6 Rack Servers and simple, nondisruptive software upgrades to the Cisco UCS system, the Cisco MDS 32Gbps switches, and the NetApp AFF storage arrays. After the supported hardware and software are in place, the configuration of FC-NVMe is similar to the FCP configuration.

NetApp ONTAP 9.5 and later provides a complete FC-NVMe solution. A nondisruptive ONTAP software update for AFF A300, AFF A400, AFF A700, AFF A700s, and AFF A800 arrays allow these devices to support an end-to-end NVMe storage stack. Therefore, servers with sixth-generation host bus adapters (HBAs) and NVMe driver support can communicate with these arrays using native NVMe.

Objective

This solution provides a high-level summary of the FC-NVMe performance with VMware vSphere 7 on FlexPod. The solution was verified to successfully pass FC-NVMe traffic, and performance metrics were captured for FC-NVMe with various data block sizes.

Solution benefits

End-to-end NVMe for FlexPod delivers exceptional value for customers with the following solution benefits:

- NVMe relies on PCIe, a high-speed and high-bandwidth hardware protocol that is substantially faster than older standards such as SCSI, SAS, and SATA. High-bandwidth, ultra-low latency connectivity between the Cisco UCS Server and NetApp storage array for most of the demanding applications.
- An FC-NVMe solution is lossless and can handle the scalability requirements of next-generation applications. These new technologies include artificial intelligence (AI), machine learning (ML), deep learning (DL), real-time analytics, and other mission-critical applications.
- Reduces the cost of IT by efficiently using all resources throughout the stack.
- Dramatically reduces response times and boosts application performance, which corresponds to improved IOPS and throughput with reduced latency. The solution delivers ~60% more performance and reduces latency by ~50% for existing workloads.
- FC-NVMe is a streamlined protocol with excellent queuing capabilities, especially in situations with more I/O operations per second (IOPS; that is, more transactions) and parallel activities.
- Offers nondisruptive software upgrades to the FlexPod components such as Cisco UCS, Cisco MDS, and the NetApp AFF storage arrays. Requires no modification to applications.

[Next: Testing approach.](#)

Testing approach

[Previous: Introduction.](#)

This section provides a high-level summary of the FC-NVMe on FlexPod validation testing. It includes both the test environment/configuration and the test plan adopted to perform the workload testing with respect to FC-NVMe for FlexPod with VMware vSphere 7.


Test environment

The Cisco Nexus 9000 Series Switches support two modes of operation:

- NX-OS standalone mode, using Cisco NX-OS software
- ACI fabric mode, using the Cisco Application Centric Infrastructure (Cisco ACI) platform

In standalone mode, the switch performs like a typical Cisco Nexus switch, with increased port density, low latency, and 40GbE and 100GbE connectivity.

FlexPod with NX-OS is designed to be fully redundant in the computing, network, and storage layers. There is no single point of failure from a device or traffic path perspective. The figure below shows the connection of the various elements of the latest FlexPod design used in this validation of FC-NVMe.



From an FC SAN perspective, this design uses the latest fourth-generation Cisco UCS 6454 fabric interconnects and the Cisco UCS VICs 1400 platform with port expander in the servers. The Cisco UCS B200 M6 Blade Servers in the Cisco UCS chassis use the Cisco UCS VIC 1440 with Port Expander connected to the Cisco UCS 2408 Fabric Extender IOM, and each Fibre Channel over Ethernet (FCoE) virtual host bus adapter (vHBA) has a speed of 40Gbps. The Cisco UCS C220 M5 Rack Servers managed by Cisco UCS use the Cisco UCS VIC 1457 with two 25Gbps interfaces to each Fabric Interconnect. Each C220 M5 FCoE vHBA has a speed of 50Gbps.

The fabric interconnects connect through 32Gbps SAN port channels to the latest-generation Cisco MDS 9148T or 9132T FC switches. The connectivity between the Cisco MDS switches and the NetApp AFF A800 storage cluster is also 32Gbps FC. This configuration supports 32Gbps FC, for Fibre Channel Protocol (FCP), and FC-NVMe storage between the storage cluster and Cisco UCS. For this validation, four FC connections to each storage controller are used. On each storage controller, the four FC ports are used for both FCP and FC-NVMe protocols.

Connectivity between the Cisco Nexus switches and the latest-generation NetApp AFF A800 storage cluster is also 100Gbps with port channels on the storage controllers and vPCs on the switches. The NetApp AFF A800 storage controllers are equipped with NVMe disks on the higher-speed Peripheral Connect Interface Express (PCIe) bus.

The FlexPod implementation used in this validation is based on [FlexPod Datacenter with Cisco UCS 4.2\(1\) in UCS Managed Mode](#), [VMware vSphere 7.0U2](#), and [NetApp ONTAP 9.9](#).

Validated hardware and software

The following table lists the hardware and software versions used during the solution validation process. Note that Cisco and NetApp have interoperability matrixes that should be referenced to determine support for any specific implementation of FlexPod. For more information, see the following resources:

- [NetApp Interoperability Matrix Tool](#)
- [Cisco UCS Hardware and Software Interoperability Tool](#)

| Layer | Device | Image | Comments |
|-----------------|---|----------------------|---|
| Computing | <ul style="list-style-type: none">• Two Cisco UCS 6454 Fabric Interconnects• One Cisco UCS 5108 blade chassis with two Cisco UCS 2408 I/O modules• Four Cisco UCS B200 M6 blades, each with one Cisco UCS VIC 1440 adapter and port expander card | Release 4.2(1f) | Includes Cisco UCS Manager, Cisco UCS VIC 1440, and port expander |
| CPU | Two Intel Xeon Gold 6330 CPUs at 2.0 GHz, with 42-MB Layer 3 cache and 28 cores per CPU | — | — |
| Memory | 1024GB (16x 64GB DIMMS operating at 3200MHz) | — | — |
| Network | Two Cisco Nexus 9336C-FX2 switches in NX-OS standalone mode | Release 9.3(8) | — |
| Storage network | Two Cisco MDS 9132T 32Gbps 32-port FC switches | Release 8.4(2c) | Supports FC-NVMe SAN analytics |
| Storage | Two NetApp AFF A800 storage controllers with 24x 1.8TB NVMe SSDs | NetApp ONTAP 9.9.1P1 | — |
| Software | Cisco UCS Manager | Release 4.2(1f) | — |
| | VMware vSphere | 7.0U2 | — |
| | VMware ESXi | 7.0.2 | — |
| | VMware ESXi native Fibre Channel NIC driver (NFNIC) | 5.0.0.12 | Supports FC-NVMe on VMware |

| Layer | Device | Image | Comments |
|--------------|--|----------|----------|
| | VMware ESXi native Ethernet NIC driver (NENIC) | 1.0.35.0 | — |
| Testing tool | FIO | 3.19 | — |

Test plan

We developed a performance test plan to validate NVMe on FlexPod using a synthetic workload. This workload allowed us to execute 8KB random reads and writes as well as 64KB reads and writes. We used VMware ESXi hosts to run our test cases against the AFF A800 storage.

We used FIO, an open-source synthetic I/O tool that can be used for performance measurement, to generate our synthetic workload.

To complete our performance testing, we conducted several configuration steps on both the storage and servers. Below are the detailed steps for the implementation:

1. On the storage side, we created four storage virtual machines (SVMs, formerly known as Vservers), eight volumes per SVM, and one namespace per volume. We created 1TB volumes and 960GB namespaces. We created four LIFs per SVM as well as one subsystem per SVM. The SVM LIFs were evenly spread across the eight available FC ports on the cluster.
2. On the server side, we created a single virtual machine (VM) on each of our ESXi hosts, for a total of four VMs. We installed FIO on our servers to run the synthetic workloads.
3. After the storage and the VMs were configured, we were able to connect to the storage namespaces from the ESXi hosts. This allowed us to create datastores based on our namespace and then create Virtual Machine Disks (VMDKs) based on those datastores.


[Next: Test results.](#)

Test results

[Previous: Testing approach.](#)


Testing consisted of running the FIO workloads to measure the FC-NVMe performance in terms of IOPS and latency.

The following graph illustrates our findings when running a 100% random read workload using 8KB block sizes.




In our testing, we found that the system achieved over 1.2M IOPS while maintaining just under 0.35ms of server-side latency.

The following graph illustrates our findings when running a 100% random write workload using 8KB block sizes.




In our testing, we found that the system achieved close to 300k IOPS while maintaining just under 1ms of server-side latency.

For 8KB block size with 80% random reads and 20% writes, we observed the following results:




In our testing, we found that the system achieved over 1M IOPS while maintaining just under 1ms of server-side latency.

For 64KB block size and 100% sequential reads, we observed the following results:



In our testing, we found that the system achieved around 250k IOPS while maintaining just under 1ms of server-side latency.

For 64KB block size and 100% sequential writes, we observed the following results:



In our testing, we found that the system achieved around 120k IOPS while maintaining under 1ms of server-side latency.

[Next: Conclusion.](#)

Conclusion

[Previous: Test results.](#)

The observed throughput for this solution was 14GBps and 220k IOPS for a sequential read workload under 1ms latency. For random read workloads, we reached a throughput of 9.5GBps and 1.25M IOPS. The ability of FlexPod to provide this performance with FC-NVMe can address the needs of any mission-critical applications.

FlexPod Datacenter with VMware vSphere 7.0 U2 is the optimal shared infrastructure foundation to deploy FC-NVMe for a variety of IT workloads thereby providing high-performance storage access to applications that require it. As FC-NVMe evolves to include high availability, multipathing, and additional operating system support, FlexPod is well suited as the platform of choice, providing the scalability and reliability needed to support these capabilities.

With FlexPod, Cisco and NetApp have created a platform that is both flexible and scalable for multiple use cases and applications. With FC-NVMe, FlexPod adds another feature to help organizations efficiently and effectively support business-critical applications running simultaneously from the same shared infrastructure. The flexibility and scalability of FlexPod also enables customers to start with a right-sized infrastructure that can grow with and adapt to their evolving business requirements.

Additional information

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

- Cisco Unified Computing System (UCS)
<http://www.cisco.com/en/US/products/ps10265/index.html>
- Cisco UCS 6400 Series Fabric Interconnects Data Sheet
<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/datasheet-c78-741116.html>
- Cisco UCS 5100 Series Blade Server Chassis
<http://www.cisco.com/en/US/products/ps10279/index.html>
- Cisco UCS B-Series Blade Servers
<http://www.cisco.com/en/US/partner/products/ps10280/index.html>
- Cisco UCS C-Series Rack Servers
<http://www.cisco.com/c/en/us/products/servers-unified-computing/ucs-c-series-rack-servers/index.html>
- Cisco Unified Computing System Adapters
http://www.cisco.com/en/US/products/ps10277/prod_module_series_home.html
- Cisco UCS Manager
<http://www.cisco.com/en/US/products/ps10281/index.html>
- Cisco Nexus 9000 Series Switches
<http://www.cisco.com/c/en/us/products/switches/nexus-9000-series-switches/index.html>
- Cisco MDS 9000 Multilayer Fabric Switches
<http://www.cisco.com/c/en/us/products/storage-networking/mds-9000-series-multilayer-switches/index.html>
- Cisco MDS 9132T 32-Gbps 32-Port Fibre Channel Switch
<https://www.cisco.com/c/en/us/products/collateral/storage-networking/mds-9100-series-multilayer-fabric-switches/datasheet-c78-739613.html>
- NetApp ONTAP 9
<http://www.netapp.com/us/products/platform-os/ontap/index.aspx>
- NetApp AFF A-Series
<http://www.netapp.com/us/products/storage-systems/all-flash-array/aff-a-series.aspx>
- VMware vSphere
<https://www.vmware.com/products/vsphere>
- VMware vCenter Server
<http://www.vmware.com/products/vcenter-server/overview.html>

- Best Practices for modern SAN

<https://www.netapp.com/us/media/tr-4080.pdf>

- Introducing End-to-End NVMe for FlexPod

<https://www.cisco.com/c/en/us/products/collateral/servers-unified-computing/ucs-b-series-blade-servers/whitepaper-c11-741907.html>

Interoperability matrixes

- NetApp Interoperability Matrix Tool

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

- Cisco UCS Hardware Compatibility Matrix

<https://ucshcltool.cloudapps.cisco.com/public/>

- VMware Compatibility Guide

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

Acknowledgements

The authors would like to thank John George from Cisco and Scott Lane and Bobby Oommen from NetApp for the assistance and guidance offered during this project execution.

Legal notices

Legal notices provide access to copyright statements, trademarks, patents, and more.

Copyright

<http://www.netapp.com/us/legal/copyright.aspx>

Trademarks

NETAPP, the NETAPP logo, and the marks listed on the NetApp Trademarks page are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

<http://www.netapp.com/us/legal/netapptmlist.aspx>

Patents

A current list of NetApp owned patents can be found at:

<https://www.netapp.com/us/media/patents-page.pdf>

Privacy policy

<https://www.netapp.com/us/legal/privacypolicy/index.aspx>

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.