



FlexPod Datacenter

FlexPod

NetApp
January 21, 2025

This PDF was generated from <https://docs.netapp.com/us-en/flexpod/flexpod-dc/sm-bcs-introduction.html> on January 21, 2025. Always check docs.netapp.com for the latest.

Table of Contents

- FlexPod Datacenter 1
 - FlexPod DataCenter with NetApp SnapMirror Business Continuity and ONTAP 9.10. 1
 - FlexPod Datacenter with VMware vSphere 7.0, Cisco VXLAN Single-Site Fabric, and NetApp ONTAP 9.7 - Design 58
 - FlexPod Datacenter with VMware vSphere 7.0 and NetApp ONTAP 9.7 - Deployment 58
 - FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Design. 59
 - FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Deployment. 59
 - FlexPod Datacenter with Cisco Intersight and NetApp ONTAP 9.7 - Design. 60
 - FlexPod Datacenter with VMware vSphere 6.7 U2, Cisco UCS fourth-generation Fabric and NetApp ONTAP 9.6 60
 - FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS fourth-generation fabric, and NetApp AFF A-Series - Design 61
 - FlexPod Datacenter with VMware vSphere 6.7 U1, Cisco UCS fourth-generation fabric, and NetApp AFF A-Series 61
 - FlexPod Datacenter with Cisco ACI Multi-Pod, NetApp MetroCluster IP, and VMware vSphere 6.7 - Design 62
 - FlexPod Datacenter with Cisco ACI Multi-Pod with NetApp MetroCluster IP and VMware vSphere 6.7 - Deployment 62

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

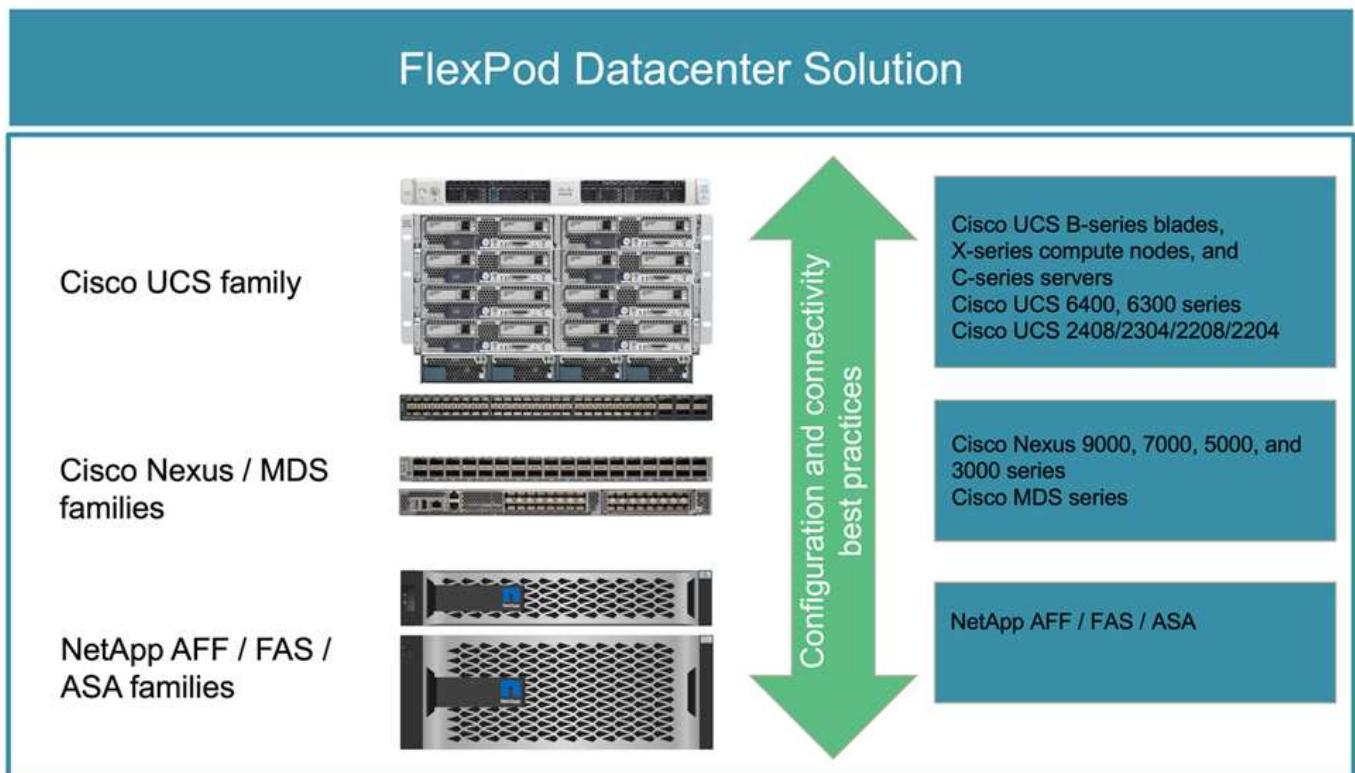
Introduction

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

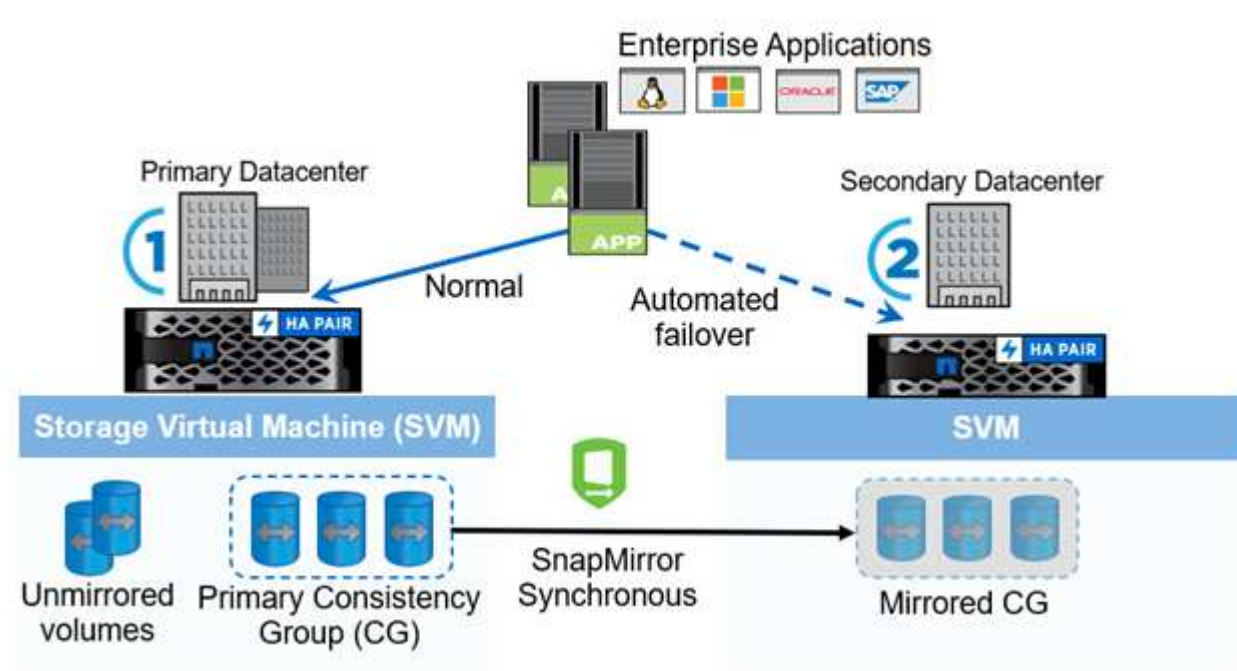
Beginning with ONTAP 9.8, you can protect SAN workloads for transparent application failover with NetApp SM-BC. You can create consistency group relationships between two AFF clusters or two ASA clusters for data

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.

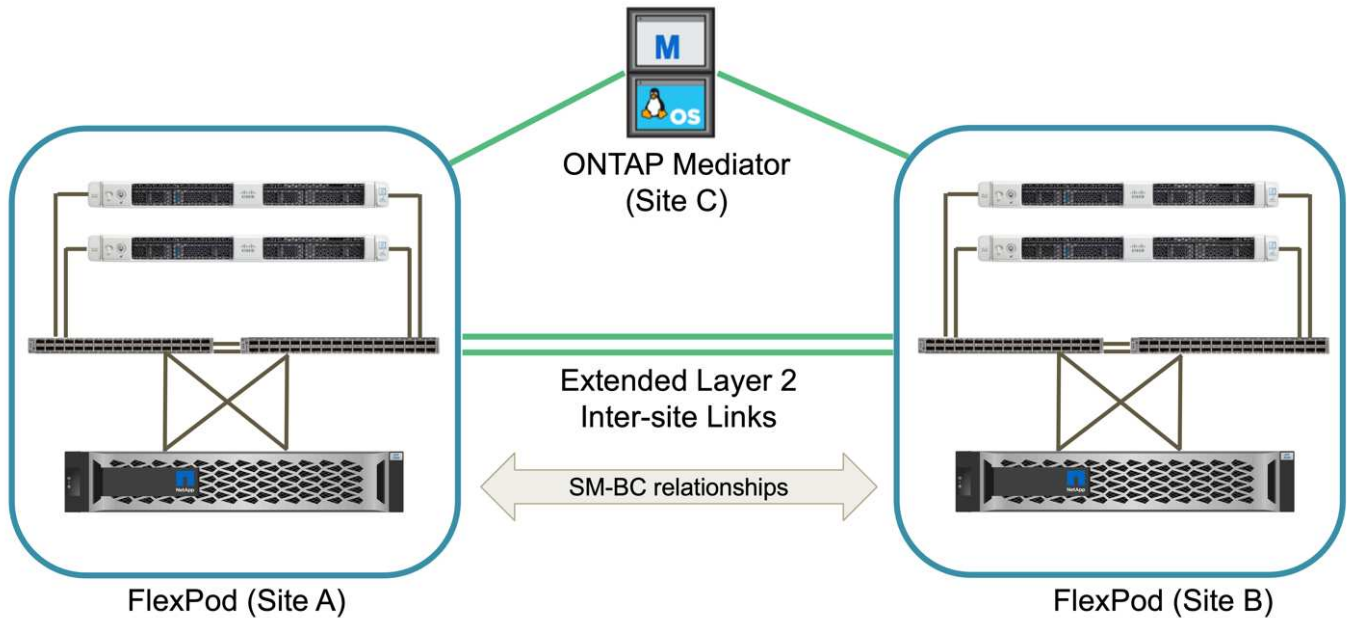


Product family	Technical specifications
AFF series	AFF series documentation
ASA series	ASA series documentation

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.

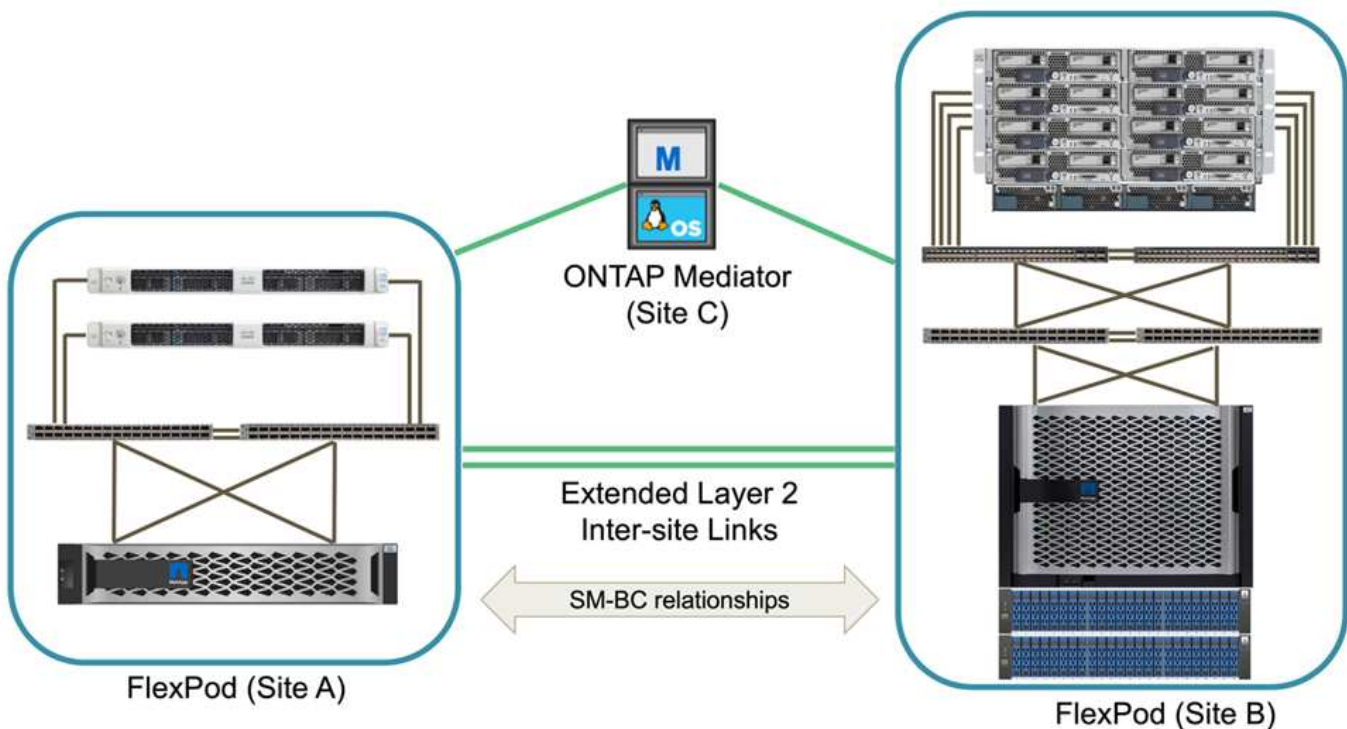
Solution topologies

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.

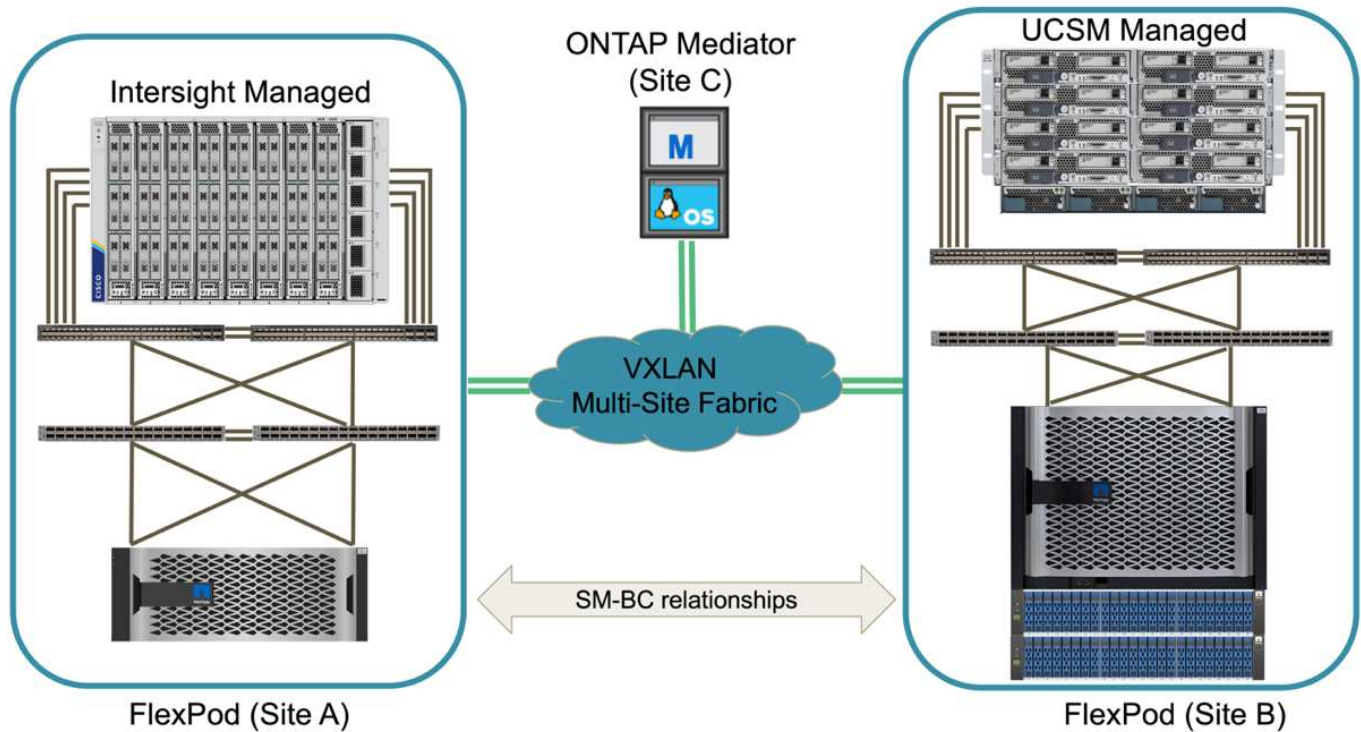
An asymmetric deployment topology that could be used by companies between a data center and a branch office in a metropolitan area might look like the following figure. For this asymmetric design, the data center requires a higher performance FlexPod with more compute and storage resources. However, 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.](#)

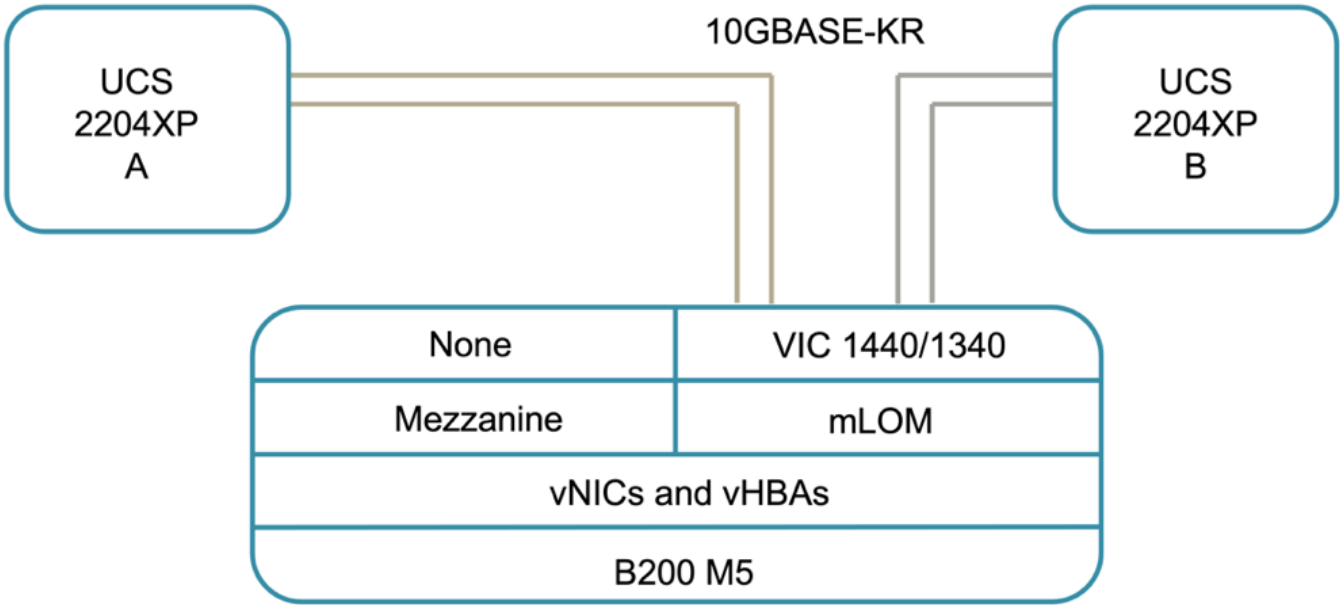
Solution validation - Compute

[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.

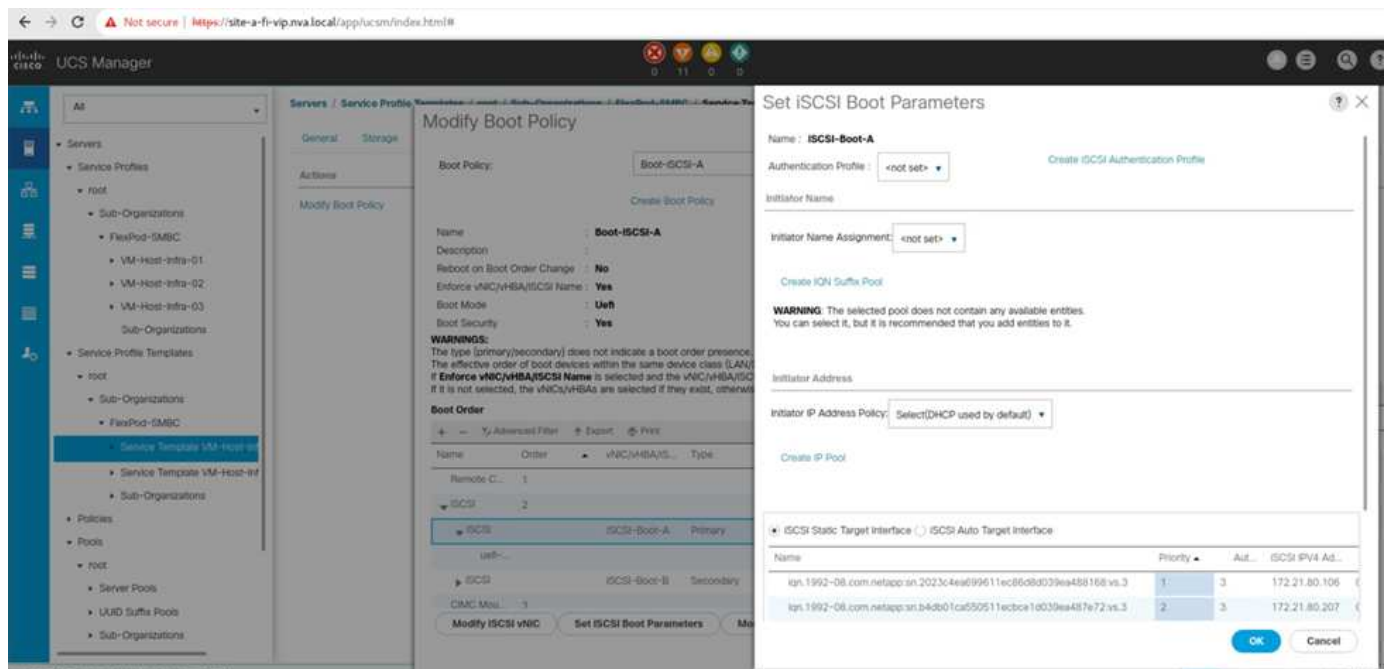
The created service profiles support the following for the VMware ESXi hosts:

- 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.

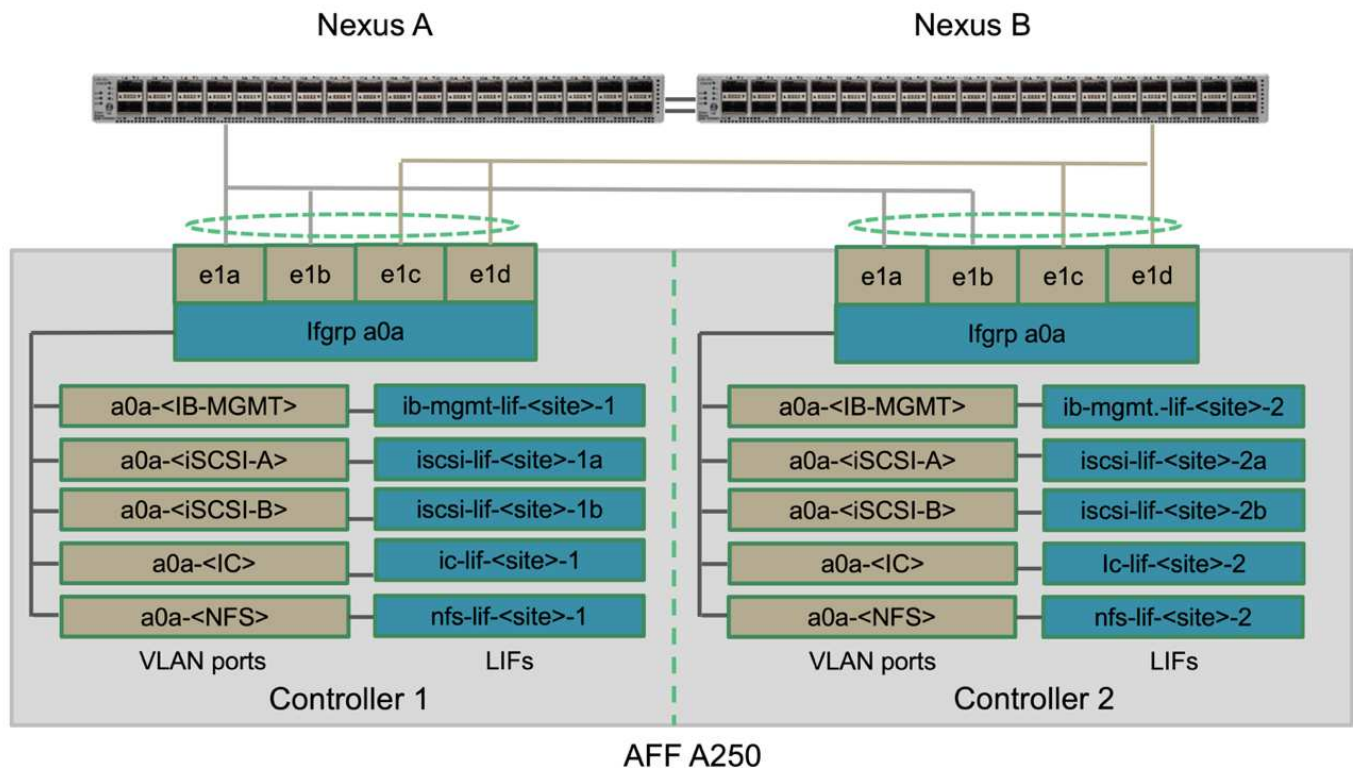


Local device	Local port	Remote device	Remote port
AFF A250 A	e0c	AFF A250 B	e0c
	e0d		e0d
	e1a	Nexus A	1/10/1
	e1b		1/10/2
	e1c	Nexus B	1/10/1
	e1d		1/10/2
AFF A250 B	e0c	AFF A250 A	e0c
	e0d		e0d
	e1a	Nexus A	1/10/3
	e1b		1/10/4
	e1c	Nexus B	1/10/3
	e1d		1/10/4

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.

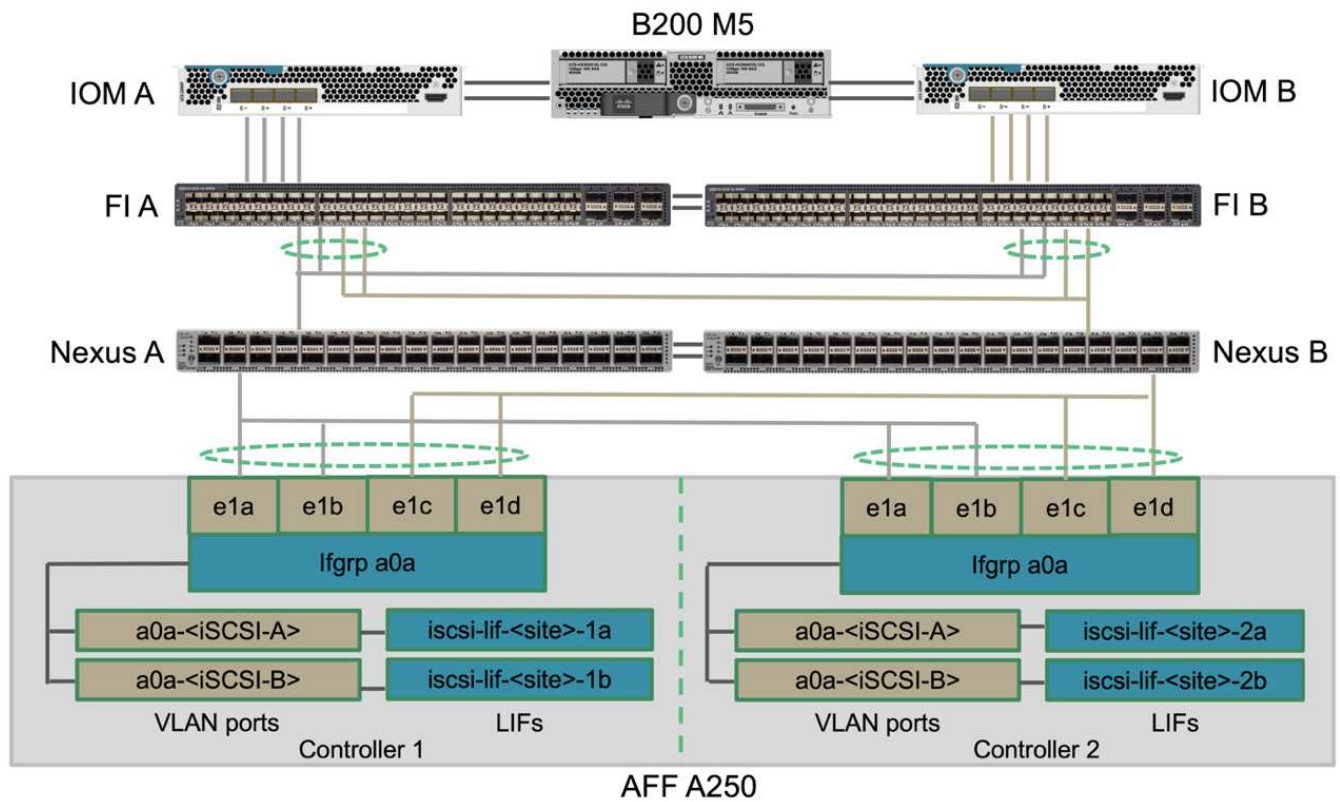


SAN boot

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.

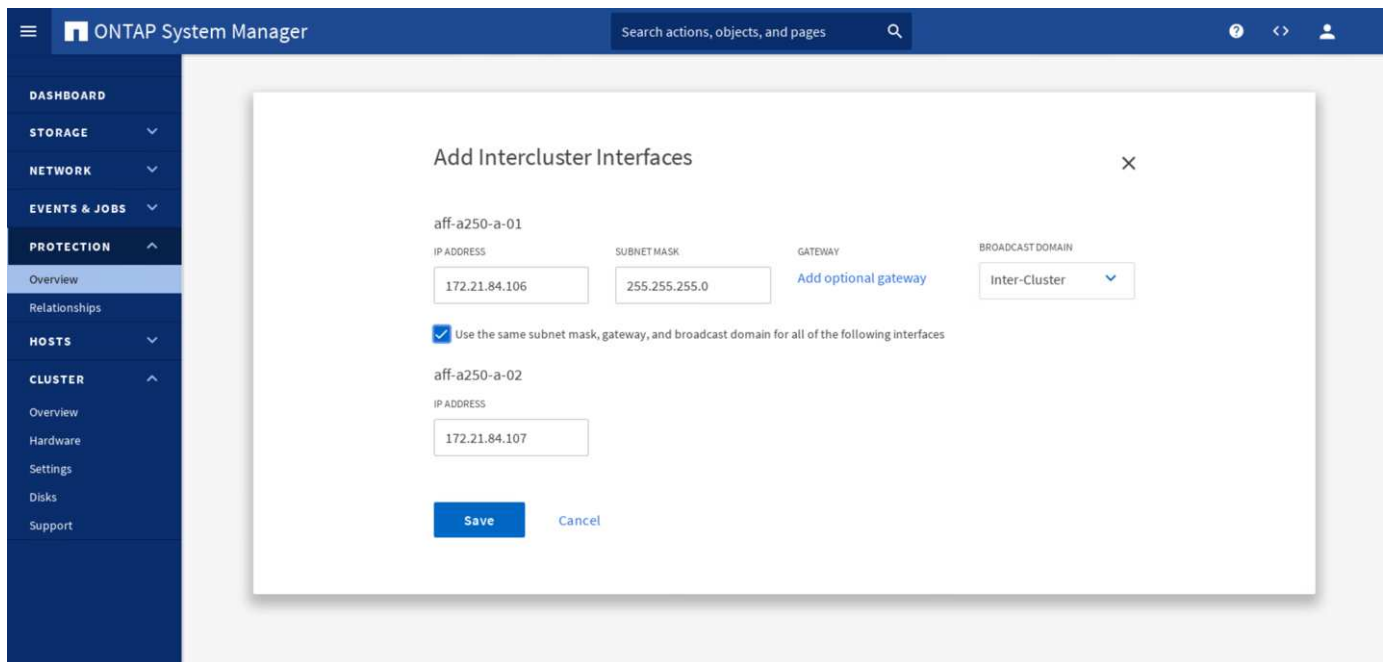
The following figure depicts the connectivity for iSCSI SAN boot of Cisco UCS server from NetApp Storage. In iSCSI SAN boot, each Cisco UCS server is assigned two iSCSI vNICs (one for each SAN fabric) that provide redundant connectivity from the server all the way to the storage. The 10/25-G Ethernet 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:

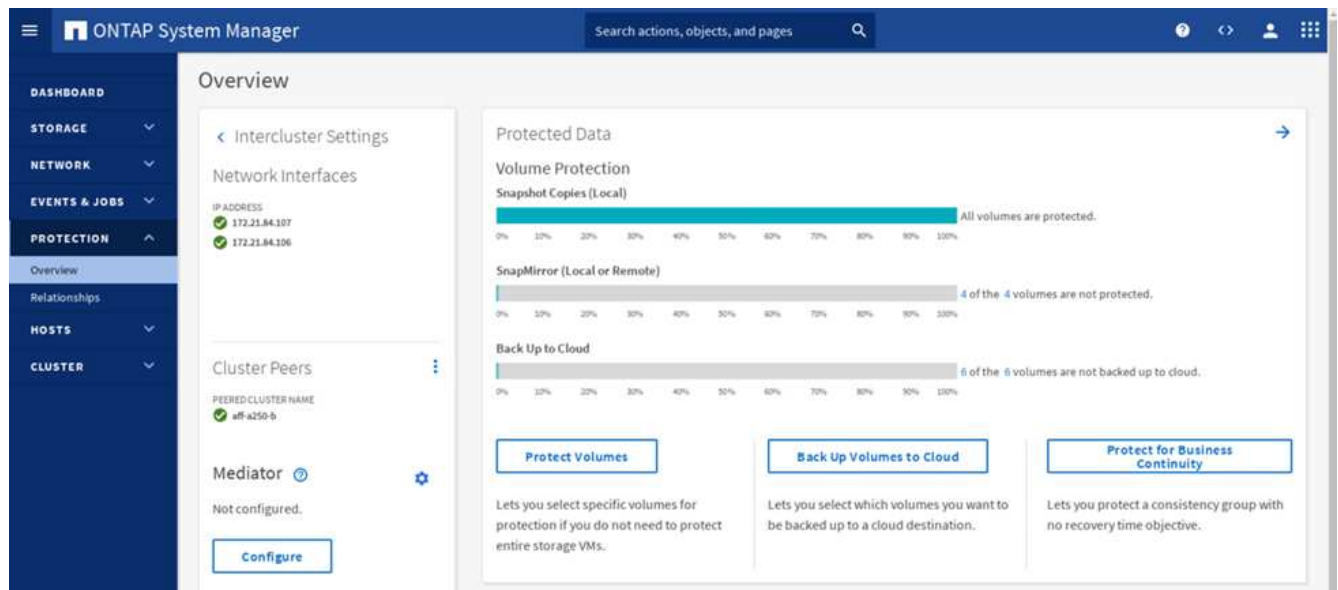
1. Generate cluster peering passphrase in the first cluster.



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.

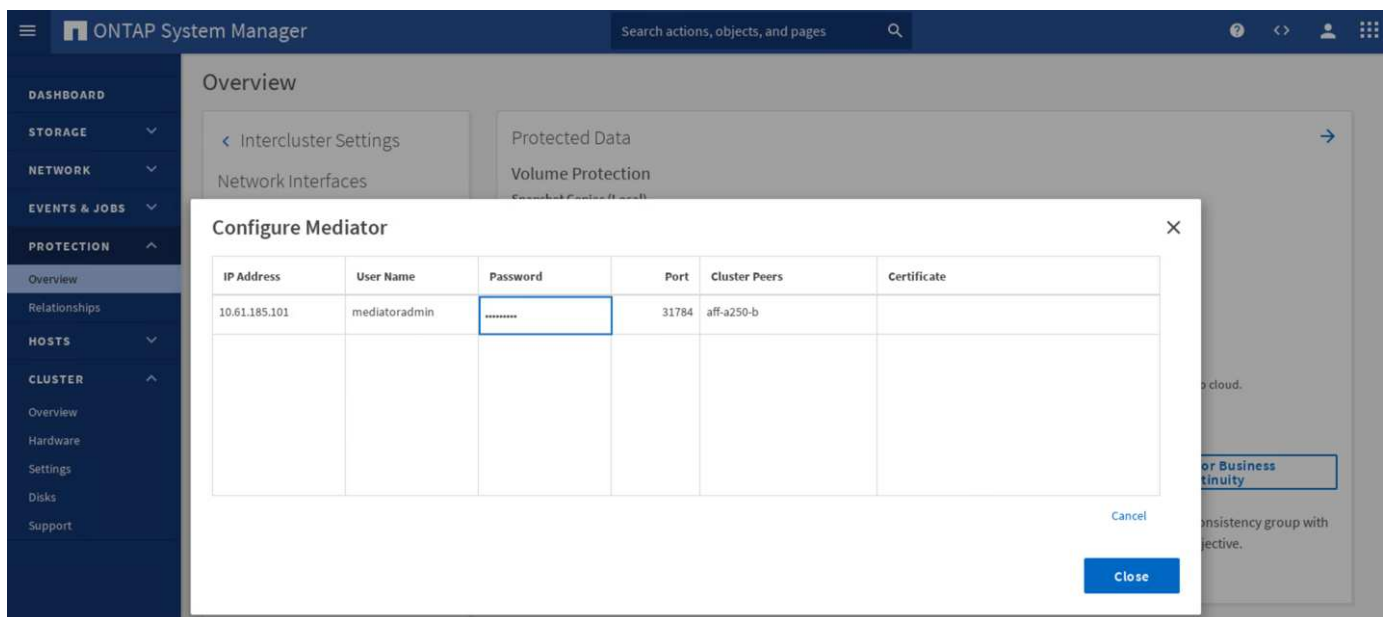


ONTAP Mediator installation and configuration

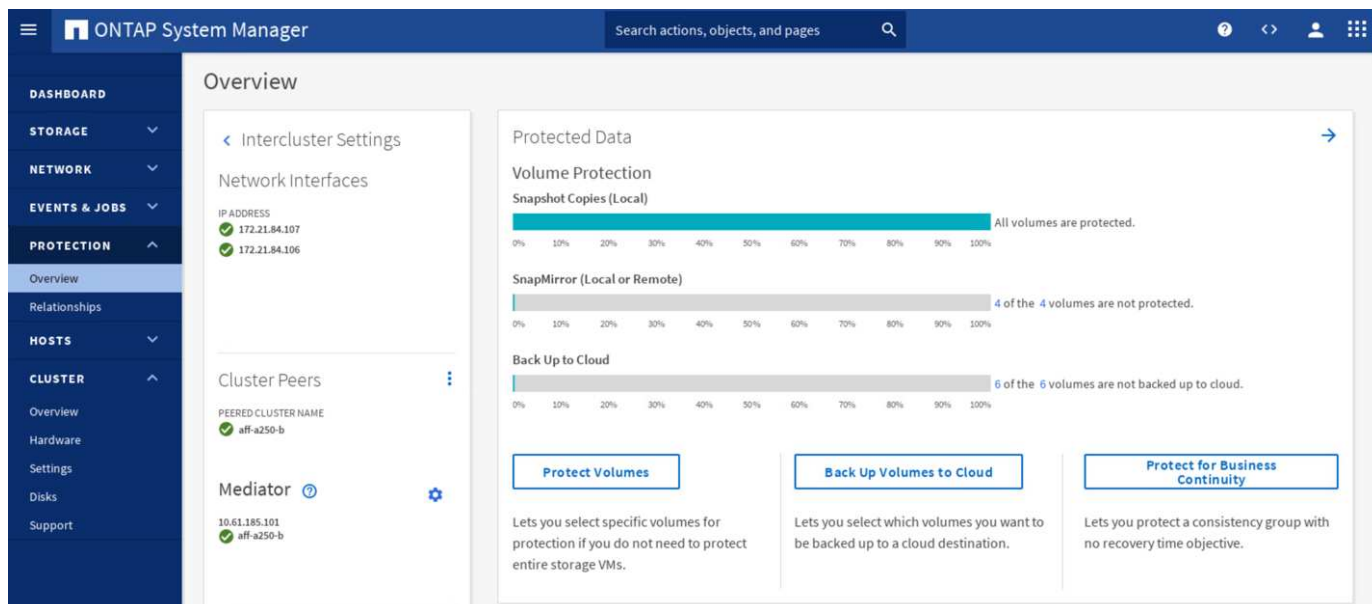
The ONTAP Mediator establishes a quorum for the ONTAP clusters in an SM-BC relationship. It coordinates automated failover when a failure is detected and helps to avoid split-brain scenarios when each cluster simultaneously tries to establish control as the primary cluster.

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.

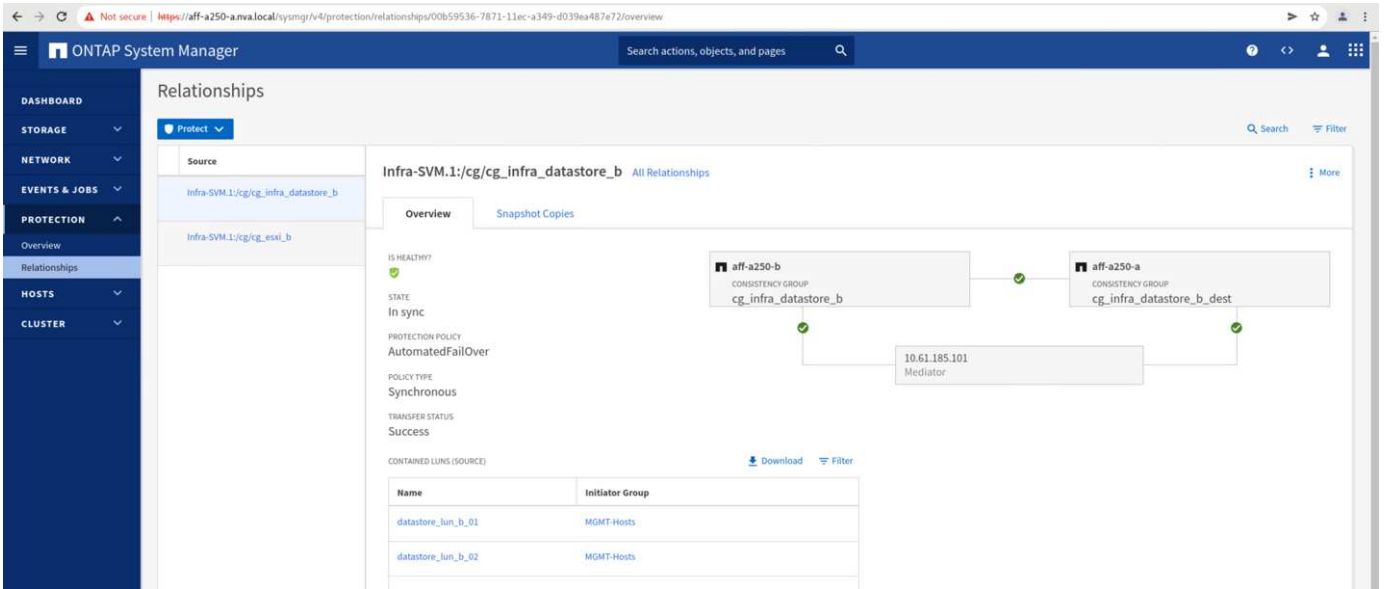
This screenshot shows the consistency group relationships at site A.

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.](#)

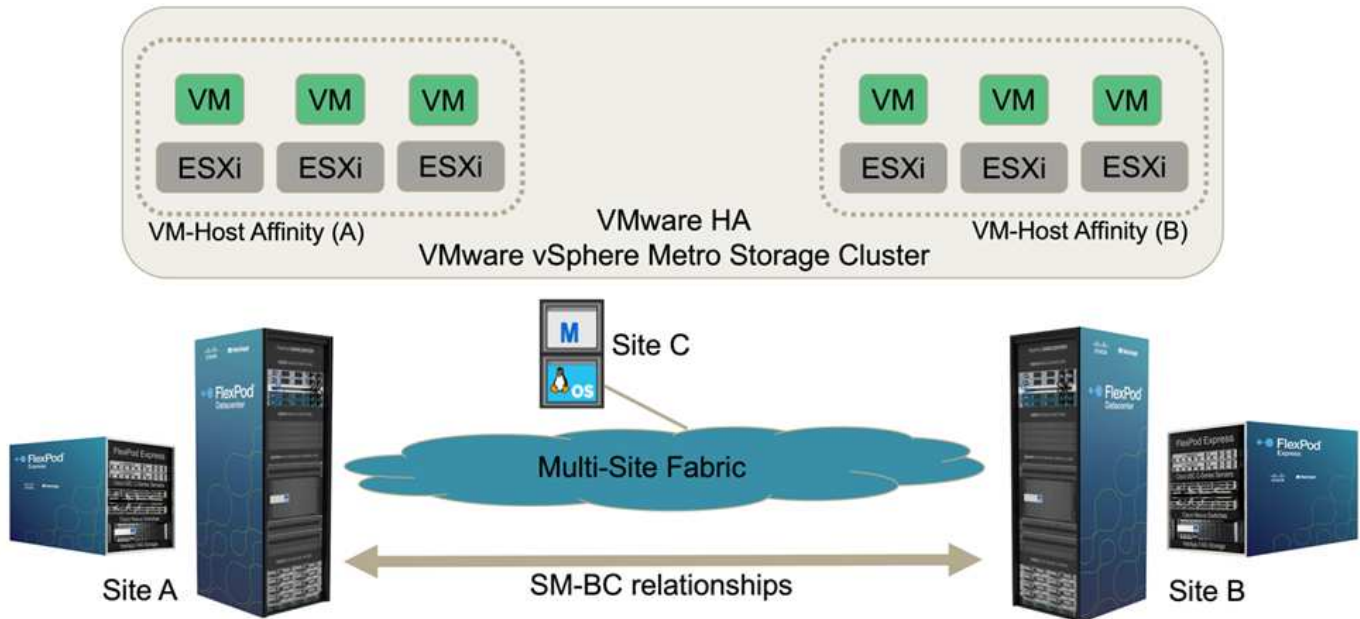
In the multi-site FlexPod SM-BC solution, a single VMware vCenter manages the virtual infrastructure resources for the entire solution. The hosts in both data centers participate in the single VMware HA cluster which spans both data centers. The hosts have access to the NetApp SM-BC solution where storage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

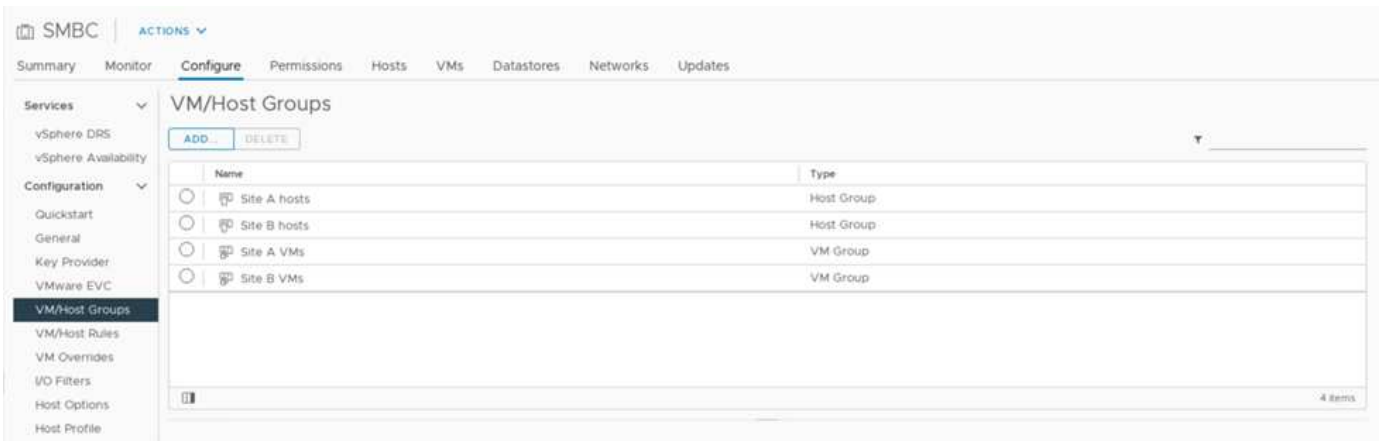
VM-Host affinity groups and rules

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.

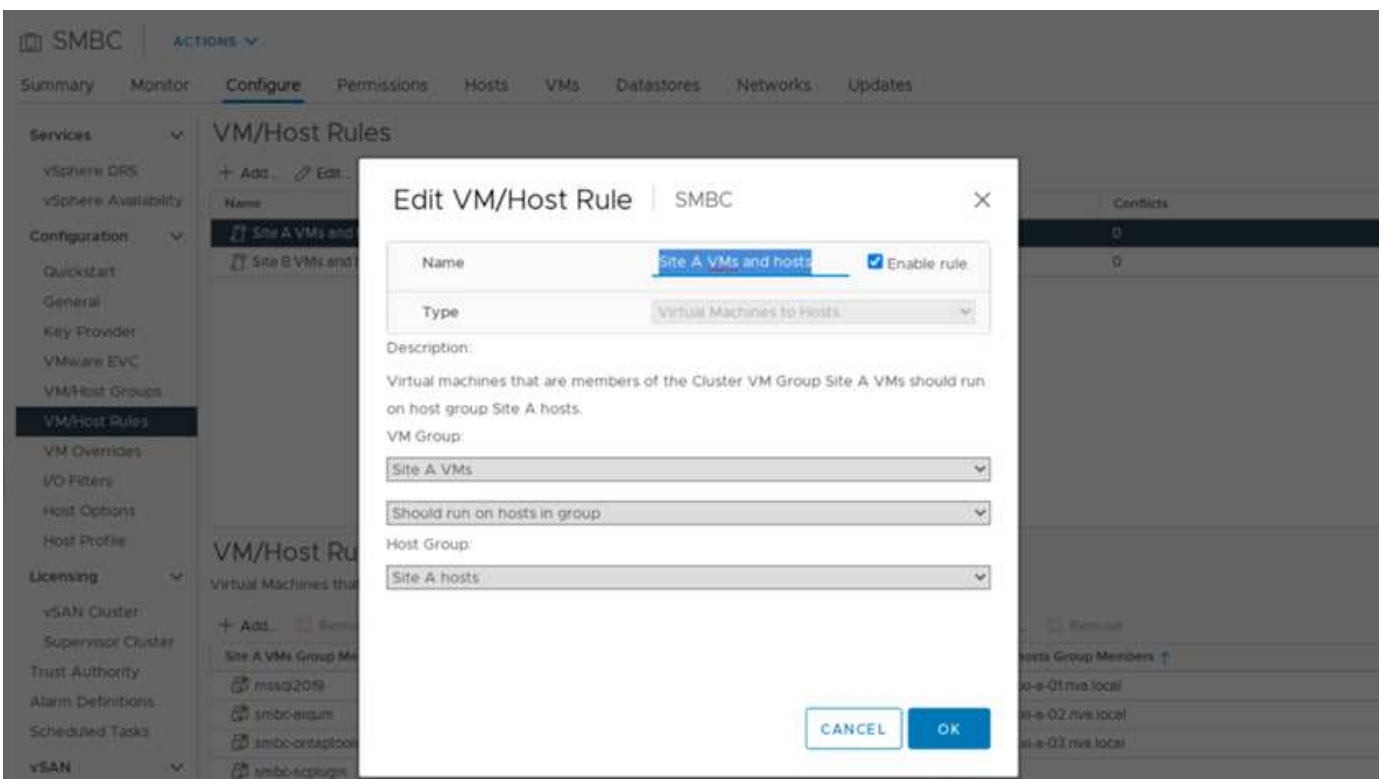
For a virtual machine that uses local datastores, its access to virtual disks becomes remote if it is migrated to a host at the remote site and thus increasing read operation latency due to the physical distance between the sites. Therefore, it is a best practice to keep virtual machines on the local 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	iSCSI	Online	iscsi_urn(ign.2010-11.com:flexpod:ucs-smbc-a-1)	8	7	56
Model: LSI MegaRAID SATA 09:07:07:00:00:00	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. 3.172.2181.106.3260	0	Active (I/O)
vmhba64 C1:T0:L0	ign.1992-08.com:netapp:sn.2023c4ee6996f1ec86d8d039ee488168 vs. 3.172.2181.107.3260	0	Active
vmhba64 C0:T1:L0	ign.1992-08.com:netapp:sn.b4db01ca5505f1ecb0e1d039ee487e72 vs. 3.172.2180.206.3260	0	Active
vmhba64 C1:T1:L0	ign.1992-08.com:netapp:sn.b4db01ca5505f1ecb0e1d039ee487e72 vs. 3.172.2180.207.3260	0	Active
vmhba64 C2:T1:L0	ign.1992-08.com:netapp:sn.b4db01ca5505f1ecb0e1d039ee487e72 vs. 3.172.2181.206.3260	0	Active
vmhba64 C3:T1:L0	ign.1992-08.com:netapp:sn.b4db01ca5505f1ecb0e1d039ee487e72 vs. 3.172.2181.207.3260	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.

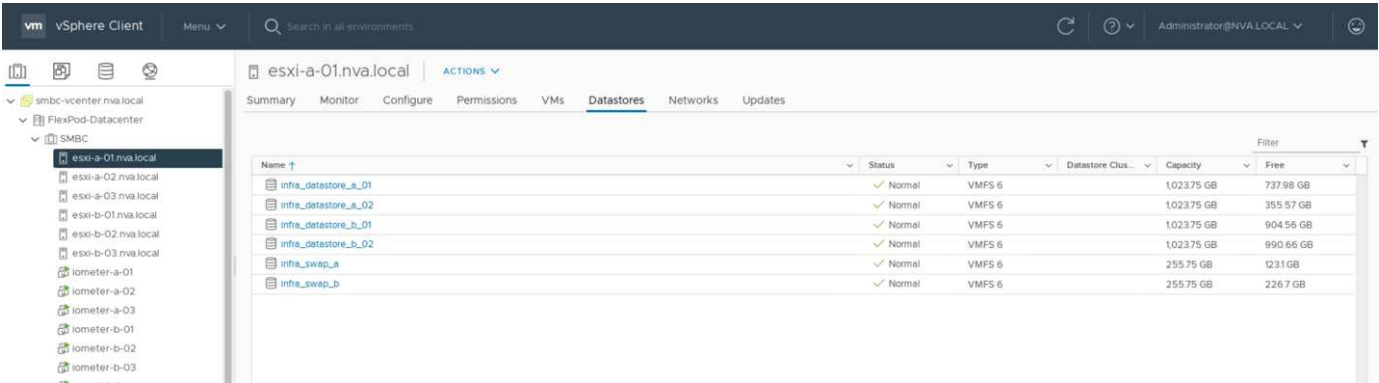
VMware vMotion and VM/host affinity test

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

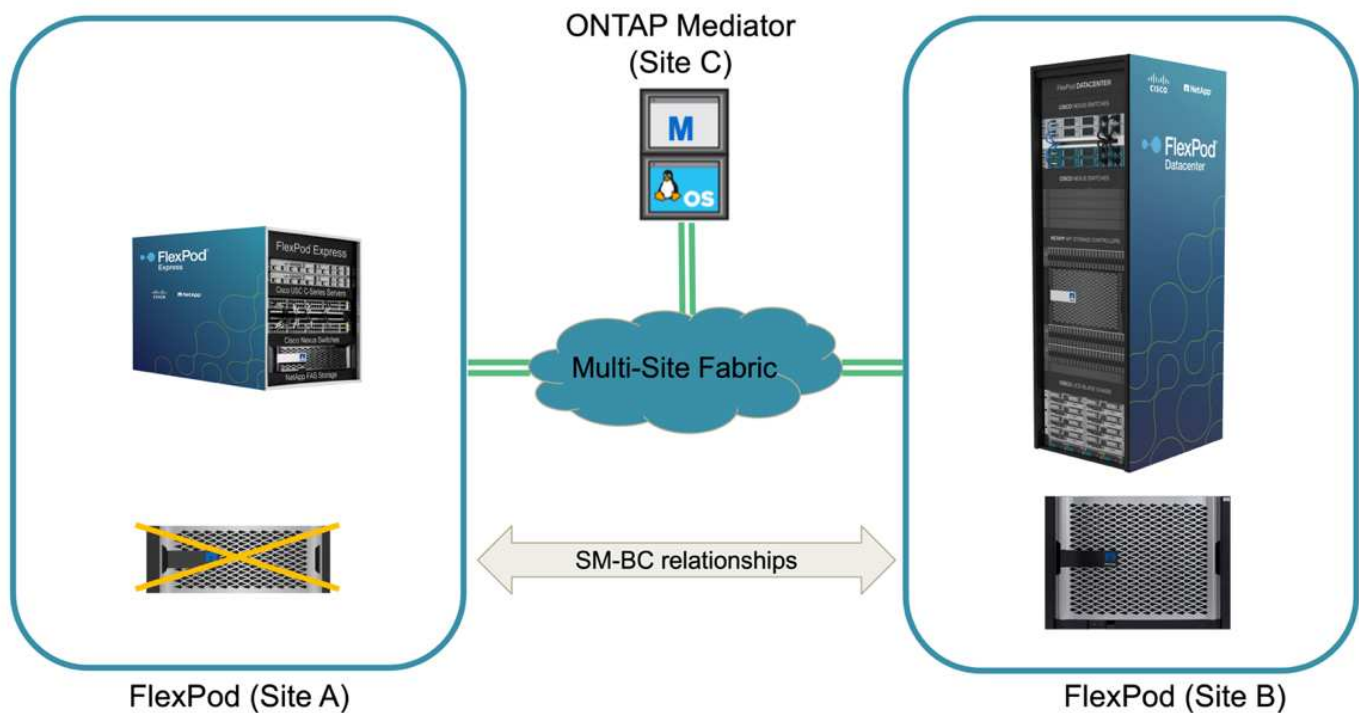
Protect

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

Actions: 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.

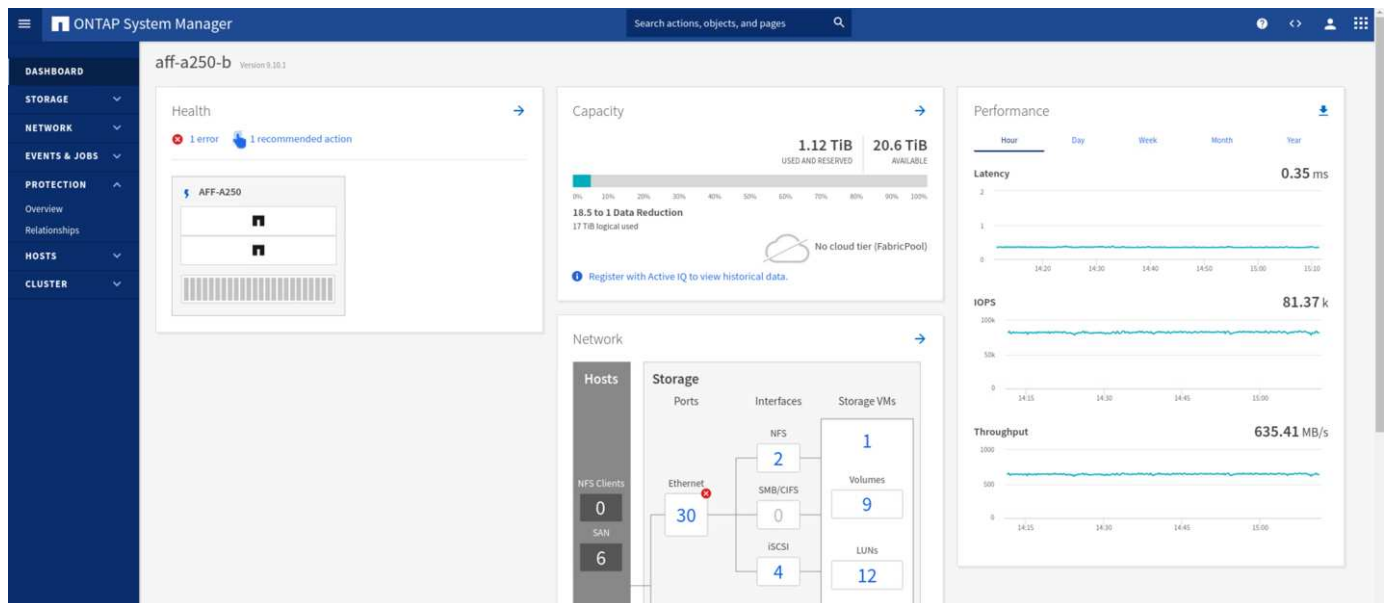
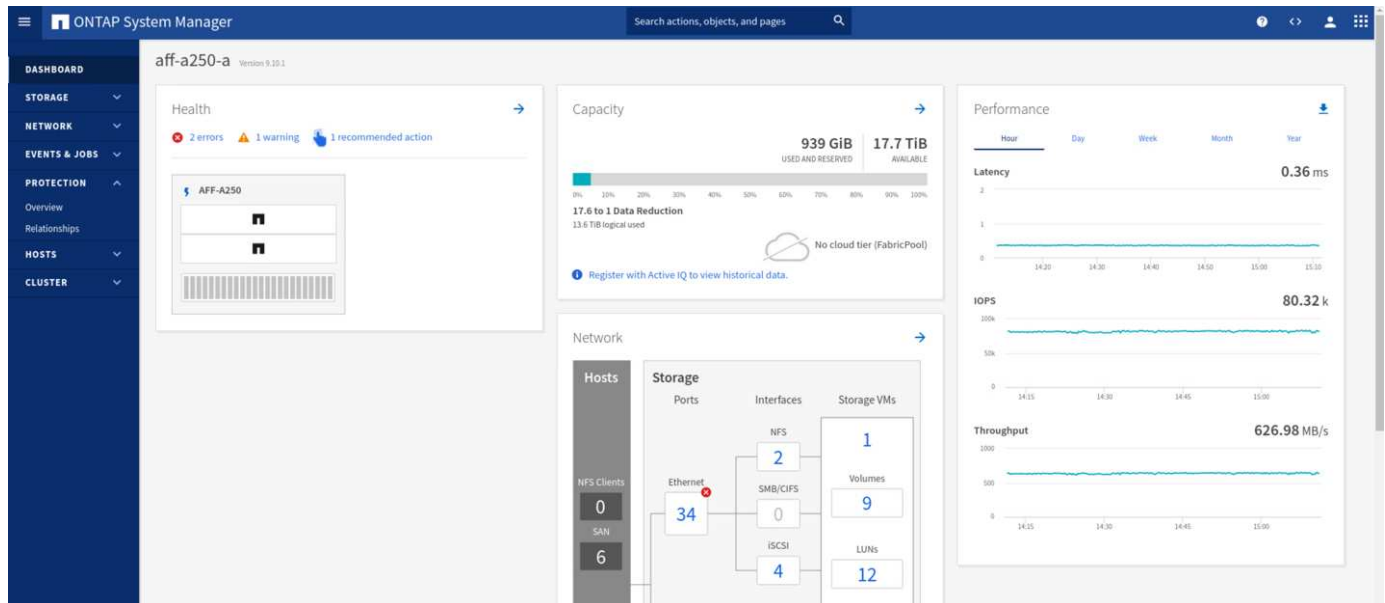


To simulate a storage disaster at site A, both storage controllers at site A can be powered off by physically turning off the power switch to discontinue the supply of power to the controllers, or by using the storage controller service processors' system power management command to power off the controllers.


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.

The Relationships page shows the following data:

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.

The Relationships page shows the following data:

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 ▼

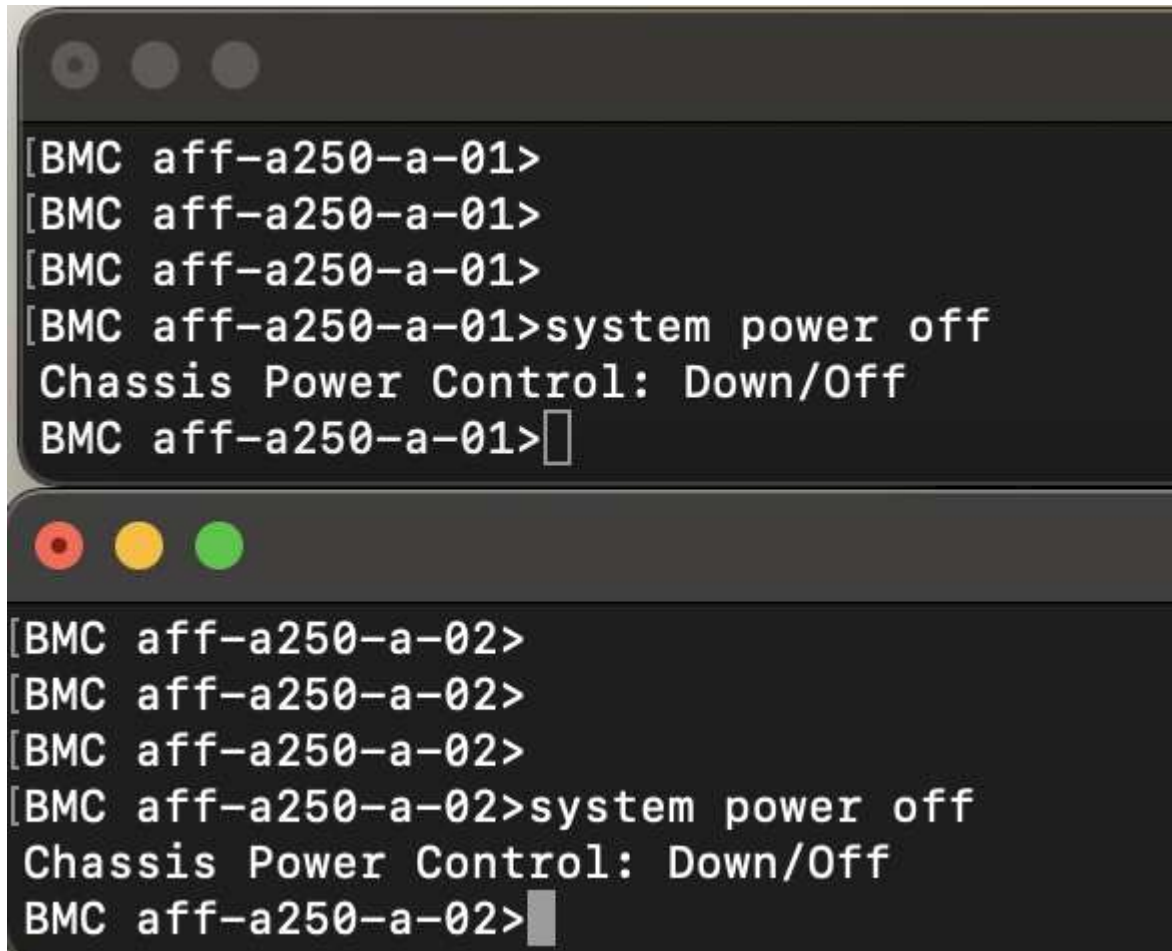
CANCEL OK

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.

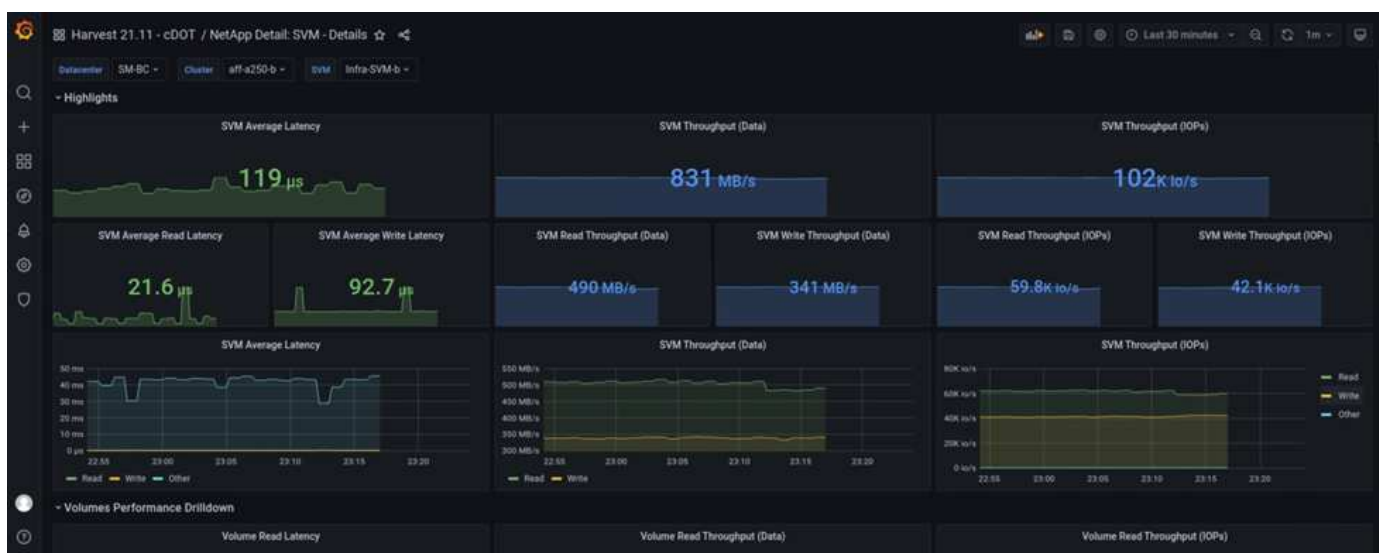
After the schema build options were updated, the schema build process was started. A few minutes later, an unplanned simulated site B storage cluster failure was introduced by powering off both nodes of the two node AFF A250 storage cluster at about the same time using system processor CLI commands.

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.



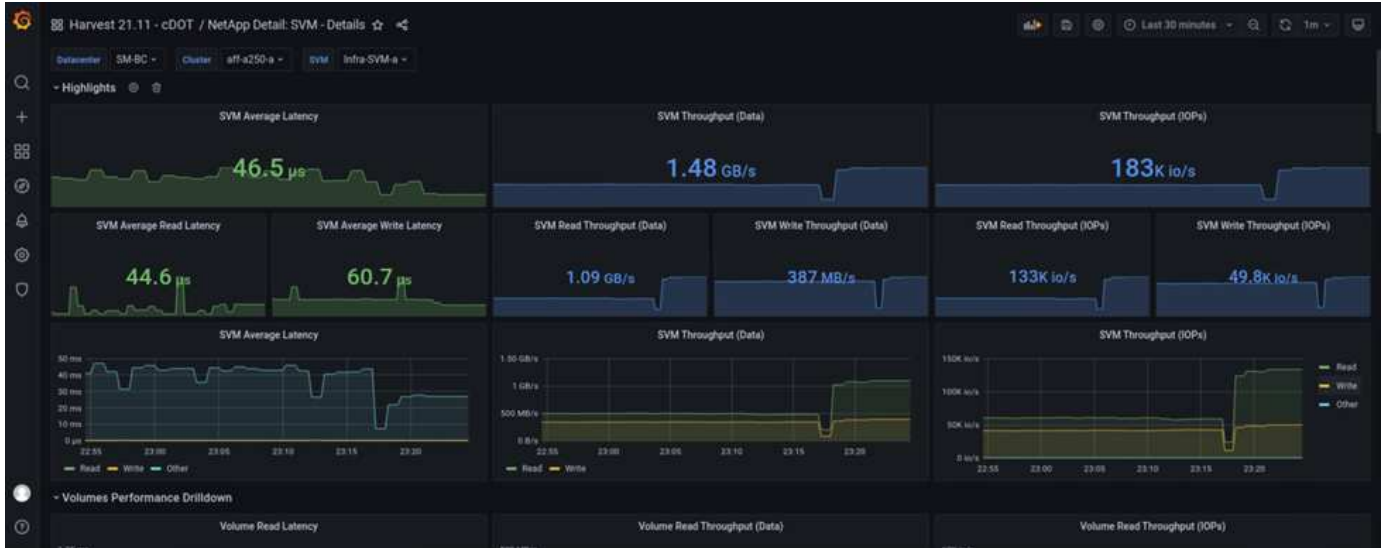
The storage cluster metrics were captured by using the NAbbox tool with the NetApp Harvest monitoring tool installed. The results are displayed in the predefined Grafana dashboards for the storage virtual machine and other storage objects. The dashboard provides metrics for latency, throughput, IOPS, 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](#)

Copyright information

Copyright © 2025 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

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

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

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

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

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

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

Trademark information

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