Best Practices for NetApp ASA

Sweksha Chaubey, Michael Peppers NetApp

July 2025 | TR-5009

Abstract

This technical report provides best practices configuration for SAN solution with NetApp ASA storage system. The document is focused on the new release of ASA and is less relevant to the original ASA release. This document also illustrates optimal SAN configurations at network, compute, storage, hypervisor (VMware) layers.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Executive Summary

NetApp ASA provides a robust, scalable, and secure infrastructure foundation designed to support enterprise application workloads with high performance, availability, and centralized management. The NetApp ASA framework leverages modular design principles, enabling seamless integration of compute, storage, and network components within a validated, interoperable ecosystem. When deployed with SAN (Storage Area Network), NetApp ASA decouples storage from compute, allowing shared, block-level access to storage resources via Fibre Channel or iSCSI, optimized for reliability and speed.

The architecture supports secure multi-tenancy through ONTAP Storage Virtual Machines (SVMs), ensures consistent data protection via technologies like SnapCenter, and integrates easily with leading virtualization platforms such as VMware. Best practices around SAN zoning, LUN masking, multipathing, and consistency groups (CGs) are enforced to guarantee operational resilience and data integrity. Additionally, compatibility with industry-standard hardware and adherence to NetApp 's Hardware and Software Universe (HWU) ensures supportability and lifecycle alignment. Overall, NetApp ASA delivers an efficient, enterprise-grade solution tailored for mission-critical applications and modern data center environments.

Supported SAN Protocols and Guidance

- NVMe/FC – Consistent sub-millisecond latency with higher IOPS
- FC & iSCSI – For broader compatibility depending on the network infrastructure and host support
- NVMe/TCP - For low latency and extreme flexibility

SAN Multipathing

In a Storage Area Network (SAN) environment, multipathing refers to the use of multiple physical paths between the host (initiator) and the storage array (target) to provide redundancy, load balancing, and increased throughput. If a path fails due to a hardware fault (HBA, cable, or switch), I/O operations can continue through an alternate path without interruption.

Multipathing is a critical component of enterprise-grade SAN environments and is especially beneficial when deploying all-flash storage arrays like NetApp ASA in conjunction with servers and switches.

Figure 1 SAN Multipathing

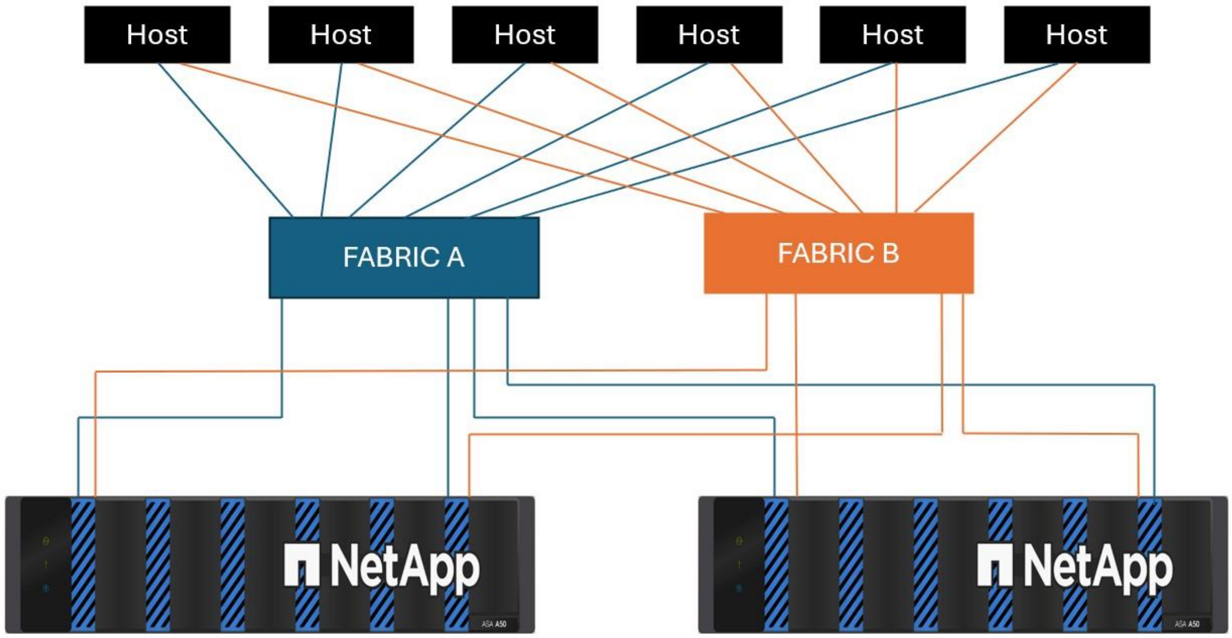

Switches: Switches configured in redundant fabric or network (Fabric A and Fabric B)

Storage: NetApp ASA all-flash arrays with dual controllers, each connected to both fabrics

Each host HBA connects to both fabrics, and the storage controllers are similarly connected to both. This forms a full-mesh topology for optimal redundancy.

If a link or switch fails, traffic is rerouted to another available path without interrupting application access

ONTAP's symmetric-active-active ensures both controllers can serve I/O for any LUN simultaneously, minimizing latency during path switch.

Best practice

- Ensure each host has at least 2 paths per SVM per node (minimum 4 total) for redundancy and load balancing.
- Always avoid any single points of failure (SPOF), this is where a single failure will make data inaccessible. Resiliency and high availability rely on fully redundant components and paths so that there are no single points of failure.

SAN Solution Scaling

Storage Area Network (SAN) provides a reliable and scalable infrastructure for deploying virtual infrastructures like VMware and business critical databases applications running on Microsoft SQL and Oracle databases. The building blocks for a SAN solution are compute, network, and storage components. As there are a variety of components to choose from, understanding your requirements thoroughly will be important for selecting the right solution, architecture and components.

Small-scale SAN

In a small business or branch office, a typical SAN configuration with NetApp ASA might consist of 2–4 physical or virtual servers (such as file servers, application servers, or a VMware host), connected to 1 NetApp storage system through a redundant network switch setup. For simplicity and cost-effectiveness, iSCSI or NVMe/TCP over Ethernet is often preferred. The NetApp ASA system hosts LUNs or NVMe namespaces which are presented to servers, allowing them to store virtual machines, databases, or application data centrally.

Key Components in Small-Scale SAN

- NetApp ASA: Provides high-performance block storage only.
- 2–4 Servers: Running Windows, Linux, or VMware.
- 2 or more Ethernet or FC Switches: For SAN connectivity and redundancy. While it is possible to use a single switch, this configuration lacks resiliency and has a single point of failure.
- SAN Protocols: iSCSI, NVMe/TCP, NVMe/FC or FC depending on your solution, existing infrastructure, skill sets, etc.

Medium-scale SAN

In a medium-scale SAN environment, it is deployed with dual-controller configurations, often backed by one or more disk shelves populated with high-performance SSDs. Connectivity is provided via Fibre Channel (FC) or Ethernet (depending on your solution) over a pair of redundant Ethernet or FC SAN switches, which connect the storage to 10–50 physical or virtual servers. These servers may host business-critical applications, file and database servers, or virtualized workloads on platforms like VMware, Hyper-V, or Linux Virtualization.

Medium-scale SANs also benefit from ONTAP’s advanced data services such as snapshot technology, Snap Mirror replication, QoS controls, and storage efficiency features like compression and deduplication. This allows businesses to scale storage capacity and performance while keeping management streamlined.

Key Features of Medium-Scale SAN

- Servers: 10-50 physical/virtual servers
- Switches: 2 SAN switches (FC or Ethernet) for redundancy
- Storage: 1-2 NetApp ASA HA pairs
- Capacity: 100 TB to 1 PB+ depending on use
- Protocols: FC, iSCSI, NVMe/FC, NVMe/TCP

Large-scale SAN

In large-scale SAN deployments, the SAN components are deployed across the environment to support vast amounts of data and client connections. The storage arrays are connected through redundant highspeed Fibre Channel or Ethernet fabrics, using enterprise-class switches and storage protocols. The system provides LUNs (for FC/iSCSI) or NVMe namespaces (for NVMe/FC and NVMe/TCP), which are mapped to hundreds of servers using initiator group(igroup) and subsystems to managed access.

Key Features of Large-Scale SAN

- Multi-petabyte capacity using dense SSDs
- Fibre Channel, iSCSI, NVMe/FC, or NVMe/TCP protocol support
- Hundreds to thousands of host connections
- High availability and failover using controller HA pairs and path redundancy
- Advanced data protection: Snapshots, SnapMirror, SnapCenter integration
- Performance: Designed for high IOPS workloads (databases, analytics, VDI, etc.)

Direct-attached storage vs. switch-attached storage

Direct Attached Storage (DAS) refers to a storage system that is directly connected to a single or small number of servers, without using any intermediate switches. It typically uses interfaces such as Ethernet

NIC ports to connect hard drives or SSDs to the host system. DAS provides block-level access, meaning the host operating system manages file systems and controls the storage directly. This type of connection does not scale well because it is limited by the number of physical ports on both host and storage. Furthermore, due to the port limitations, you may also face limits to the amount of resiliency and scalability this configuration can deliver.

Note: Directly connecting servers and storage using fibre channel is not supported because ONTAP's FC target requires NPIV services from an FC switch to support ONTAP's FC logical interfaces (LIFs). Directly connecting ONTAP storage as an initiator, usually for tape backups or FLI is supported because ONTAP's FC Initiator doesn't use LIFs.

A Storage Area Network (SAN) is a high-speed, networked storage solution that allows multiple servers to access centralized storage devices. It uses a dedicated storage network, typically built with Fibre Channel (Fibre Channel and NVMe/FC) or Ethernet (iSCSI and NVMe/TCP), and connects servers and storage arrays through SAN switches. Like DAS, SAN provides block-level access, but the storage is delivered over a network fabric. The addition of switches allows for easily scaling the size of the fabric as well as providing redundancy and therefore resiliency to our customer's SANs.

Table 1) Direct-attached storage vs. switch-attached storage

Features	Direct-attached storage	Switch-attached storage
Connection Type	Connected directly via interfaces like SATA, SAS, USB, and Ethernet.	Connected via switches using Fibre Channel or Ethernet.
Architecture	Point-to-point, each storage device is dedicated to one or a small number of servers.	Network based, shared storage accessible by multiple servers via a switching network/fabric.
Scalability	Limited, typically requires downtime to scale and has physical expansion limits.	Highly scalable, allows dynamic addition of storage without downtime.
Performance	Typically, high for the single server its connected to, but not optimized for concurrent access.	High performance with low-latency access across multiple servers, optimized for enterprise workloads.
Data sharing	Low, a small number of hosts can be connected to DAS since the number of physical ports on your storage will be your limiting factor	Enables multiple hosts to share and access the same storage volumes concurrently.
Cost	Lower initial cost, fewer components required.	Higher upfront investment due to switches and management infrastructure.
Management Complexity	Simple to configure and manage.	Requires SAN knowledge and tools for zoning, LUN masking, etc.
Use Cases	Local backups- Standalone servers, small business setups.	Enterprise databases- virtualization environments, mission-critical applications.
Fault Tolerance & Redundancy	Limited, often no built-in redundancy unless RAID is configured locally.	Built-in redundancy (multi-pathing, fabric failover) high availability solutions are common.

NetApp All-flash SAN Array

The NetApp ASA (All SAN Array) is a purpose-built, high-performance storage platform specifically designed for block-based SAN environments. NetApp ASA is an All-SAN Array and does not host either NAS or object protocols like S3 or their related features. It is optimized for enterprise SAN workloads that require low latency, high availability, and consistent performance, such as databases, virtual machines, and mission-critical applications. Unlike NetApp 's FAS or AFF systems, which support both file and block protocols, the NetApp ASA series is focused solely on block storage (LUNs and NVMe namespaces), making it a dedicated SAN solution.

Note: MCC is not supported for NetApp ASA systems.

The new NetApp ASA A-Series family, which is the focus of this TR, includes ASA A20, A30, A50, A70, A90, and A1K.

Table 2) Key NetApp ASA platforms technical specifications

Specifications	A1K	A90	A70	A50	A30	A20	C30
Form factor	2 x 2U	4U	4U	2U	2U	2U	2U
Drive count	Requires NS224	48	48	24	24	24	24
Max cluster size	12 Nodes	12 Nodes	12 Nodes	12 Nodes	8 Nodes	6 Nodes	
Max raw capacity per HA pair	2.67 PB	2.67 PB	2.67 PB	1.8 PB	1.1 PB	734 TB	
Max raw capacity per cluster	16 PB	16 PB	16 PB	11 PB	4.4 PB	2.2 PB	
PCIe expansion slots per HA pair	18	18	18	8	8	8	
Max FC speed	64 Gbps	64 Gbps	64 Gbps	64 Gbps	64 Gbps	64 Gbps	64 Gbps
Max Ethernet speed	200 Gbps	200 Gbps	200 Gbps	100 Gbps	100 Gbps	100 Gbps	100 Gbps
Minimum ONTAP version	9.16.0 GA	9.16.0 GA	9.16.0 GA	9.16.1	9.16.1	9.16.1	9.16.1

Note: Please see NetApp ASA datasheet and NetApp Hardware Universe for detailed technical specifications and the various supported configurations and cluster limits of these new NETAPP ASA platforms.

The new NetApp ASA A-Series storage systems deliver robust enterprise-class capabilities with the following key advantages:

- High Availability Architecture: Designed with a dual-controller, symmetric active-active architecture, the NetApp ASA A-Series ensures uninterrupted access to storage through multipath I/O. It supports non-disruptive firmware upgrades and delivers industry-leading reliability with 100% data availability guarantee.
- Scalable Performance Through Clustering: The NetApp ASA platform supports scale-out clustering for enhanced performance and capacity. System expansion supports up to 6 nodes for the NetApp ASA A20, up to 8 nodes for the NetApp ASA A30, and up to 12 nodes for higher-tier NetApp ASA models—enabling linear scalability across mid-range and high-end deployments.
- Always-On Data Efficiency: NetApp 's storage efficiency technologies—including inline data compression, deduplication, and compaction—maximize usable capacity. These capabilities support a

guaranteed 4:1 storage efficiency ratio, significantly reducing physical storage footprint without impacting performance.

- Comprehensive Security and Access Control: Advanced security features protect both data and administrative operations. Capabilities include multi-factor authentication (MFA), role-based access control (RBAC), multi-admin verification for critical actions, and secure multi-tenancy to isolate workloads. The platform also supports FIPS 140-2 compliant mode to help meet stringent regulatory requirements.
- Streamlined Deployment and Management: ONTAP System Manager provides an intuitive interface for provisioning storage units (SUs, also called LUNs and NVMe namespaces), configuring data protection, and performing system upgrades. Seamless integration with VMware environments via the NetApp ONTAP Tools for VMware vSphere plugin allows administrators to provision and manage storage directly from vCenter.

NetApp platforms

• Low-End Platform

It is the entry-level system in the NetApp ASA A-Series, designed for small IT environments, branch offices, or edge deployments where budget and footprint are key considerations. It features a compact 2U form factor with a dual-controller architecture and supports symmetric active-active access for uninterrupted I/O. With up to 6-node clustering, it provides basic scale-out capabilities, though its throughput and IOPS are optimized for modest workloads.

It supports SSDs and offers inline storage efficiency features like deduplication and compression, making it ideal for general-purpose SAN workloads, small-scale virtualization, and departmental file or block storage. The A20 is best suited for customers who need enterprise-grade resiliency and efficiency but with lighter performance and capacity demands.

Figure 2 Low-End Platform

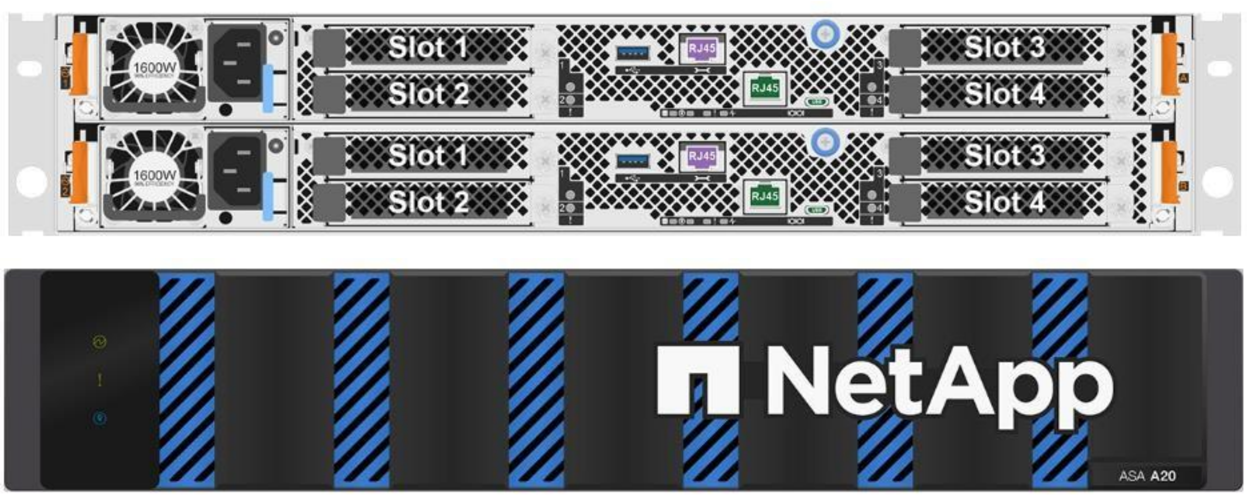

• Mid-Range Platform

Represent the midpoint in the A-Series lineup, designed for medium-sized enterprises requiring a balance of performance, capacity, and cost-efficiency. They support up to 12-node clusters, delivering strong linear scalability for both performance and capacity. With robust controller specs, native NVMe support, and expanded I/O capabilities, the NetApp ASA A50 is well-suited for business-critical applications including database workloads, enterprise virtualization, and VDI.

It’s always-on storage efficiency features ensure that capacity is used optimally, while ONTAP software provides seamless data protection, high availability, and intelligent storage tiering. The A50 is ideal for customers consolidating multiple workloads onto a single platform while maintaining predictable performance and simplified management.

Figure 3 Mid-End Platform

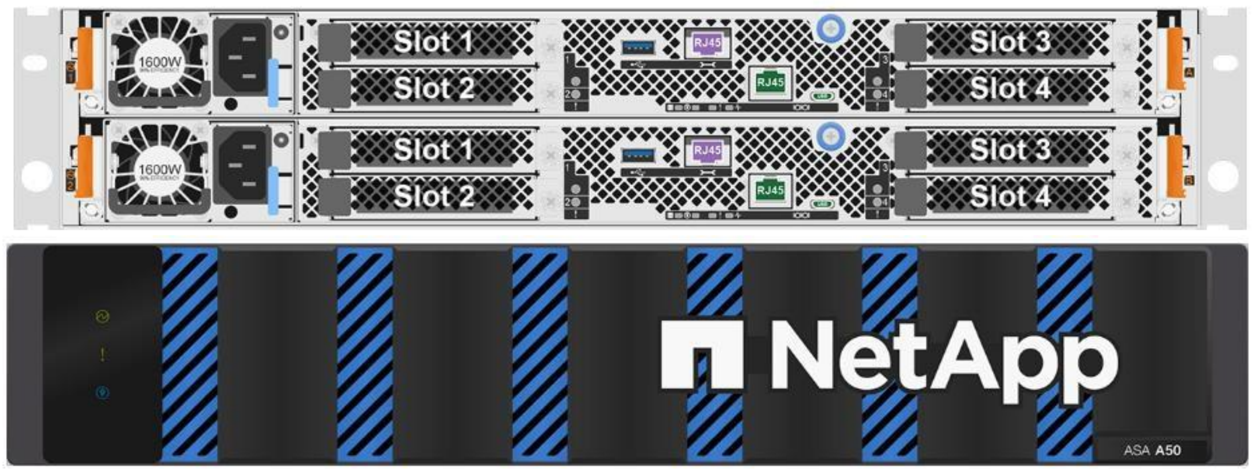

- **High-End Platform**

Sit at the top of the NetApp ASA A-Series portfolio as the flagship all-flash SAN solution. Built for the most demanding enterprise and service provider environments, it offers unmatched performance, scale, and reliability. It combines ultra-high-performance NVMe storage with dense compute and memory resources to drive massive IOPS and throughput at sub-millisecond latencies.

Supporting up to 12 nodes in a cluster, the A1K is designed for hyper-scale virtualization, large-scale AI/ML pipelines, high-end analytics, and consolidated enterprise workloads. With advanced data services, full-stack security, and deep integration with automation and orchestration platforms, the NetApp ASA A1K delivers mission-critical storage at data center scale.

Figure 4 High-End Platform

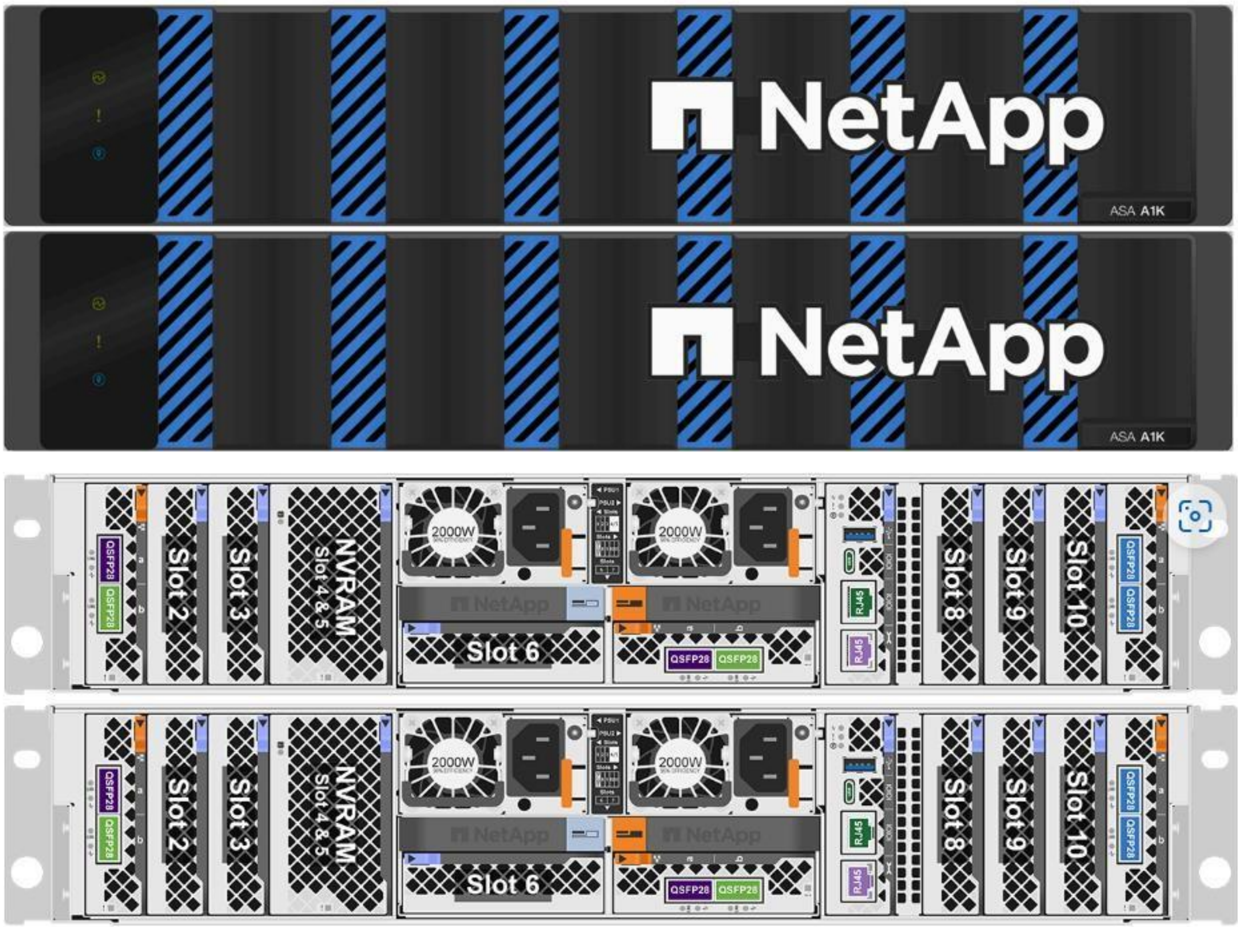

ONTAP Overview

NetApp ® ONTAP® is a robust enterprise-grade data management platform designed to unify, secure, and optimize data across SAN and hybrid cloud environments. It includes industry-leading storage efficiency technologies—such as inline deduplication, compression, and compaction—and supports scale-out clustering of up to 12 nodes for SAN deployments when leveraging mid-range and high-end NetApp ASA platforms. ONTAP integrates NetApp Snapshot™ technology, enabling space-efficient, instantaneous point-in-time copies of datasets without impacting performance. These snapshots can be used for rapid backup, cloning, and disaster recovery operations, ensuring data protection and operational continuity across critical enterprise applications.

NetApp storage platforms provide a hybrid cloud-ready foundation, powered by NetApp ’s intelligent data infrastructure. This enables seamless data movement from the edge, where data is generated, to the core, where it is processed, and to the cloud, where it can be analyzed at scale using elastic compute and AI/ML services. This approach supports dynamic data tiering and integrated cloud services to help customers extract business insights from their data while optimizing cost and performance across environments. For up-to-date information on supported features, platform scalability, and hardware capabilities, administrators should refer to the NetApp Hardware Universe and ASA documentation sites for the latest updates.

ONTAP System Manager is the embedded web-based GUI for ONTAP, offering centralized management of the storage system via a modern browser interface. It enables administrators to provision and manage storage objects—such as LUNs and NVMe namespaces—and define consistency groups for applicationaligned data protection. Consistency groups allow coordinated Snapshot creation across multiple storage volumes, ensuring crash-consistent backups for multi-volume applications. The interface also provides visibility into system alerts, capacity utilization, storage efficiency metrics, and network configuration. Additionally, System Manager streamlines ONTAP software upgrades and access control configuration, supporting simplified day-to-day operations and reduced administration overhead.

If you are already familiar with NetApp AFF, FAS, and prior NetApp ASA systems, the new NetApp ASA systems have the following main differences:

- When a cluster is created, a default data SVM, svm1, is created with the SAN protocols enabled. IP data LIFs support iSCSI and NVMe/TCP protocols and use the default-data-blocks service policy by default.
- Instead of user-managed aggregates and volumes, ONTAP manages the common pool of storage with storage availability zone (SAZ) and automatically creates or deletes the associated volume when a storage unit is created or deleted.
- Storage Units (SUs – a blanket term meaning LUN and/or namespace) are always thick provisioned (the space for them is completely reserved from storage) and the containing volume is thinly provisioned to its maximum size of 600TB
- Temperature-sensitive storage efficiency (TSSE) is not applicable

ONTAP Secure Multi- tenancy

ONTAP® secure multi-tenancy aligns with the architecture's principles of modularity, scalability, and centralized control.

ONTAP Secure Multi-Tenancy within the NetApp ASA is implemented using Storage Virtual Machines (SVMs), which logically partition a single physical NetApp storage system into multiple isolated environments. Each SVM operates independently with its own network configuration—including dedicated IPspaces and Logical Interfaces (LIFs) allowing overlapping IP addresses and complete network segmentation. This ensures strict tenant isolation at the network layer. Within each SVM, volumes, export policies, and protocol services (such iSCSI, and Fibre Channel) are managed separately, enforcing data access controls and namespace segregation. Authentication and authorization mechanisms, including role-based access control (RBAC) and integration with external identity services (LDAP, Kerberos), are applied at the SVM level, ensuring that administrative domains remain isolated. This architecture enables secure, scalable multi-tenant deployments on shared physical infrastructure, aligning with NetApp ASA goals of modularity, security, and centralized but segmented management.

What's new in ONTAP

- Platforms

The following NetApp ASA systems are available. These platforms deliver a unified hardware and software solution that creates a simplified experience specific to the needs of SAN-only customers.

  - ASA A1K
  - ASA A90
  - ASA A70

- System Manager

Simplified storage management: NetApp ASA systems introduce the use of storage units with consistency groups for simplified storage management.

A *storage unit* makes storage space available to your SAN hosts for data operations. A storage unit refers to a LUN for SCSI hosts or an NVMe namespace for NVMe hosts.

A consistency group is a collection of storage units that are managed as a single unit.

- Data security

Onboard key manager and dual-layer encryption: NetApp ASA systems support an onboard key manager and dual-layer (hardware and software) encryption.

The new capabilities available in ONTAP 9.16.1 for NetApp ASA.

- Platforms

The following NetApp ASA systems are supported beginning in ONTAP 9.16.1. These platforms deliver a unified hardware and software solution that creates a simplified experience specific to the needs of SANonly customers.

  - ASA A50
  - ASA A30
  - ASA A20
  - ASA C30 (new capacity-optimized option supporting both 15TB and 30TB QLC drives)

- Data protection

Support for encryption key migration between key managers: When you switch from the ONTAP onboard key manager to an external key manager at the cluster level, you can use the ONTAP command line interface (CLI) to easily migrate the encryption keys from one key manager to the other.

Support for hierarchical consistency groups: Hierarchical consistency groups allow you to create a parent consistency group that contains multiple child consistency groups. This simplifies data protection and management for complex data structures.

- Protocol support

NVMe support for symmetric Active/Active multipathing: NVMe/FC and NVMe/TCP now support symmetric active-active architecture for multipathing so that all paths between the hosts and storage are active/optimized. • Storage efficiency

Support for automatic rebalancing of storage units: ONTAP will automatically rebalance storage units across your storage availability zones for optimal performance and capacity utilization.

NVMe space deallocation enabled by default: Space deallocation (also called “hole punching” and “unmap”) is enabled for NVMe namespaces by default. Space deallocation allows a host to deallocate unused blocks from namespaces to reclaim space.

This greatly improves overall storage efficiency, especially with workloads that have high data turnover.

Modern SAN solution with ASA

SAN solution components

A SAN (Storage Area Network) solution with ASA using servers and switches involves several detailed components, each contributing to performance, availability, and scalability:

Server

Standard rack-mount or blade servers configured for enterprise workloads, each equipped with dual hotswappable power supplies connected to independent power sources for electrical redundancy. The servers feature redundant Host Bus Adapters (HBAs) ports for Fibre Channel or dual-port/high-throughput Network Interface Cards (NICs) for iSCSI (typically 10GbE or 25GbE) and ASA platforms are also offering 40 / 100 GbE connectivity, ensuring link redundancy and failover capability. Server operating systems include native multipath I/O (MPIO) stacks that balance traffic between multiple paths and maintain uninterrupted access to storage during path or component failures.

SAN Storage Array

A centralized enterprise-grade storage array providing block-level storage access via FC, iSCSI, NVMe/FC, and NVMe/TCP targets. ASA storage typically includes dual active-active controllers, highperformance SSDs, RAID protection, auto-tiering, automatic placement during provisioning, hardwareaccelerated snapshots, and native replication features. They support multiple front-end protocols and come with built-in failover mechanisms to maintain uptime and performance continuity.

SAN Switches

The enterprise-class Fibre Channel switches or Layer 2/3 Ethernet switches used in Ethernet SANs. These switches are deployed in a redundant dual-fabric topology (fabric or network A and fabric or network B) to ensure fault tolerance and load distribution. Fabric zoning is applied in Fibre Channel environments to logically segment traffic and isolate communication paths between initiators (servers) and targets (storage), enhancing both security and performance. Both FC and Ethernet SANs also use masking to actively identify which hosts can see and access SUs.

HBAs/NICs (Host Connectivity Interfaces)

Installed in each server, HBAs are used for Fibre Channel connections—typically supporting 8/16/32/64 Gbps speeds with SFP+ or QSFP+ transceivers, while NICs used for Ethernet can operate at 1/10/25/40/100Gbps over copper or optical fiber cabling. These interfaces connect to separate Ethernet switches in each fabric, ensuring path redundancy. FC fabrics also rely on NPIV services from the fibre channel switches in their environments to create logical interfaces (LIFs) to logical separate FC traffic.

Zoning and LUN Masking

Zoning (on SAN switches): Implemented using WWPN-based zoning to limit SAN fabric visibility, ensuring each server (initiator) communicates only with designated storage ports (targets). This reduces risk of crosstalk, misconfigurations, and enhances fabric security.

LUN Masking (on Storage Array): Configured to define which LUNs are visible to which servers, preventing unauthorized or accidental data access and enforcing data path isolation per host/application.

> **Note:** Best practices suggest either not using FabricPool with block storage (SAN) or only using it to tier snapshots to lower storage tiers.

SAN Solution Topology

The SAN topology under NetApp ASA uses a centralized storage architecture where servers are decoupled from storage, and communication between them occurs over a dedicated SAN network. The topology is designed for high availability, redundancy, scalability, and secure resource access, typically structured in a dual-fabric design.

Figure 5 SAN Solution Topology

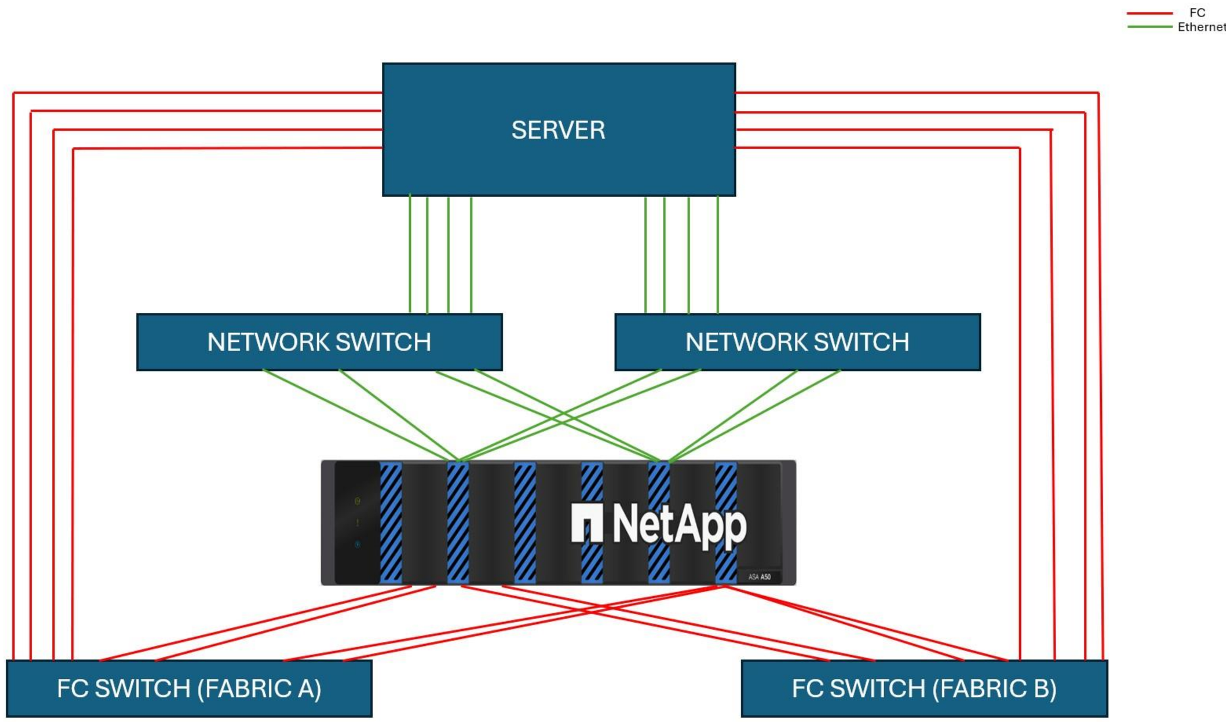

Topology for IP-based SAN

An IP-based SAN typically using the iSCSI protocol provides block-level storage access over a standard Ethernet network instead of dedicated Fibre Channel infrastructure. In the NetApp ASA, which emphasizes modularity, centralized resource control, and high availability, this topology supports scalability and cost-efficiency using standard Ethernet switches and generic x86 servers.

This IP-based SAN topology is fully compliant with ASA goals by offering modular design, centralized management, fault tolerance, and cost-effective deployment using open, vendor-neutral components.

Best practice

- Use a dedicated, redundant network with dual switches and NICs for iSCSI traffic, isolated via VLANs to ensure high availability and performance. Enable jumbo frames, QoS, and multipathing (MPIO) for throughput and fault tolerance. Secure access with CHAP authentication and implement LUN masking to control host visibility.

Topology for FC-based SAN

A Fibre Channel SAN provides high-speed, low-latency, and lossless block-level storage access, ideal for enterprise-grade applications such as databases, virtualization, and transaction-heavy workloads. Within NetApp ASA, which emphasizes modularity, fault tolerance, and centralized management, the FC-based SAN is deployed using dual-fabric architecture for resilience, utilizing standard rack servers and generic Fibre Channel switches.

This topology ensures a resilient, secure, and high-performing block storage environment, aligning with ASA goals of modular infrastructure, centralized resource control, and enterprise-grade availability.

Best practice

- Use dual-fabric architecture with redundant HBAs and Fibre Channel switches to ensure high availability and fault tolerance. Implement zoning for secure initiator-target mapping and configure multipathing (MPIO) for load balancing and failover. Regularly monitor performance and firmware versions for optimal stability and efficiency.

Workloads

This section includes recommendations on implementing LUN striping to optimize throughput and ensure balanced performance across the storage environment. Additionally, it outlines how ONTAP Quality of Service (QoS) can be used to manage resource utilization and maintain consistent performance levels across different workloads, minimizing the risk of any one workload adversely affecting others.

LUN Striping

To provide sufficient throughput for a given application it is strongly recommended to use striping of several LUNs at the host level to distribute the workload from a single application onto different LUNs. Often, LUN striping is configured by using logical volume managers (LVM), which are available natively in Linux and Unix OSs. LVMs combine 2 or more LUNs together and create a logical volume that stripes data across the LUNs the LVM was created from. This allows more throughput to be provided for an application beyond a single LUN limit and allows efficient utilization of the ONTAP based storage controllers. In combination with Quality of Service described within the next section, LUN stripping helps to provide the required throughput for a given application and to reduce noisy neighbors effects of applications generating large sequential write workloads such as SAP HANA.

Application specific recommendations can be found here:

- NetApp solutions: NetApp Solutions
- SAP and SAP HANA: SAP Business Application and SAP HANA Database Solutions
- Enterprise databases: ONTAP and enterprise applications

QoS

Quality of Service (QoS) is a crucial ONTAP feature designed to ensure predictable and consistent performance across SAN workloads. QoS enables administrators to define performance policies by setting thresholds (limits and floors) on IOPS (input/output operations per second) and throughput (MB/s) for specific volumes or LUNs. These limits and floors can be used to avoid contention for shared resources between the different workloads sharing those resources. This helps eliminate the "noisy neighbor" problem where one workload might affect the performance of others. NETAPP ASA supports both maximum and minimum IOPS thresholds—maximum IOPS to cap resource usage, and minimum IOPS to guarantee performance for critical applications. Additionally, adaptive QoS can automatically adjust policies based on volume size, ensuring consistent performance as data grows. By leveraging QoS, ASA systems deliver service-level objectives for performance-sensitive workloads, making them ideal for enterprise SAN environments.

Best practice

- The QoS burst percentage should be changed from 50% to 5% to limit the possibility of microbursts influencing the overall performance.
- Use a non-shared throughput ceiling QoS group-policy for each LUN where large sequential write workload is expected.

Enable client access to storage

For the ASA storage systems, it is very easy to provision storage units for client access. The process is similar whether you are using SCSI LUNs with iSCSI protocol or NVMe namespaces with NVMe/TCP protocol.

Storage-end configuration

1. Create IPspaces

An IPspace is a distinct IP address space in which storage VMs reside. When you create IPspaces, you enable your storage VMs to have their own secure storage, administration, and routing.

2. Create subnets

A subnet allows you to allocate specific blocks of IPv4 or IPv6 addresses to use when you create a LIF (network interface).

3. Create a storage VM

During cluster setup, your default data storage virtual machine (VM) is created. All new storage units are created inside your default data storage VM unless you create and select a different storage VM.

4. Create a LIF (network interface)

A LIF (network interface) is an IP address associated with a physical or logical port. Create LIFs on the ports you want to use to access data. Storage VMs serve data to clients through one or more LIFs.

Host-end configuration

1. Fabric Configuration In Fibre Channel SAN:

Identify each server's WWPNs from its HBAs.

Configure zoning on FC switches to allow communication between server HBAs (initiators) and storage ports (targets).

In iSCSI SAN:

Assign iSCSI qualified names (IQNs) to each client server.

Ensure servers can reach storage over dedicated VLANs/subnets.

2. LUN Masking and Access Control

Configure LUN masking on the storage array to expose only authorized LUNs to specific servers.

This prevents unintended access and enhances data security.

3. Multipath Configuration (MPIO)

Install and configure Multipath I/O (MPIO) software on each client server to enable:

Redundant data paths (via dual fabrics or networks).

Load balancing and automatic path failover.

4. Storage Protocol Setup*

For FC SAN: Ensure HBAs are properly installed, drivers updated, and the OS recognizes the FC targets.

For iSCSI SAN: Configure the iSCSI initiator with target IPs and authentication (e.g., CHAP) if required.

5. OS-Level Configuration*

On the client server

a. Rescan for new devices.

b. Partition and format the LUNs as needed.

c. Mount or assign drive letters for use by applications or OS.

Validation and Monitoring

Verify connectivity using tools like multipath -ll, esxcli, or OS-specific disk utilities.

Monitor paths, latency, and throughput using SAN management tools and server-side diagnostics.

Best practice

- An iSCSI LIF on a controller port can be configured to automatically failover to the controller's HA partner when the port goes down or when controller failover happens. LIF failover provides enhanced IO availability for these failure scenarios. The default IP LIF, which supports both iSCSI and NVMe/TCP protocols, does not support LIF failover. For the solution validation, we will configure iSCSI-only LIFs for iSCSI protocol usage as failover is supported for iSCSI LIFs. Due to utilizing different VLANs for traffic separation, we will also configure separate NVMe/TCP-only LIFs, which do not support failover, for NVMe/TCP protocol usage
- You must use the ONTAP command line interface (CLI) to create iSCSI-only LIF and enable failover for iSCSI-only LIF
- Create LUNs on the SAN storage system based on capacity and performance requirements.
- Assign the LUNs to specific initiators (client servers) or host groups.

Note: Please refer to https://www.NetApp .com/media/135828-flexpod-san-ucs-xdirect-NetApp ASAnva-deploy-guide.pdf for more information on client access to storage.

Ecosystem interoperability validation

Ecosystem interoperability validation in ASA ensures that all components within the architecture—including servers, storage, networking devices, and software—work seamlessly together to deliver a stable, reliable, and optimized infrastructure. This process involves rigorous testing and certification of hardware and software combinations to validate compatibility, performance, and supportability across the ecosystem.

The validation also helps maintain compliance with industry standards and ensures that all elements function cohesively within the modular and scalable design principles of ASA.

NetApp Interoperability Matrix Tool (IMT)

To check for interoperability information for ASA solution in NetApp Interoperability Matrix Tool (IMT): https://support.NetApp .com/matrix.

1. Go to NetApp IMT site.

1. Use the Advanced Search tool to search for ASA.

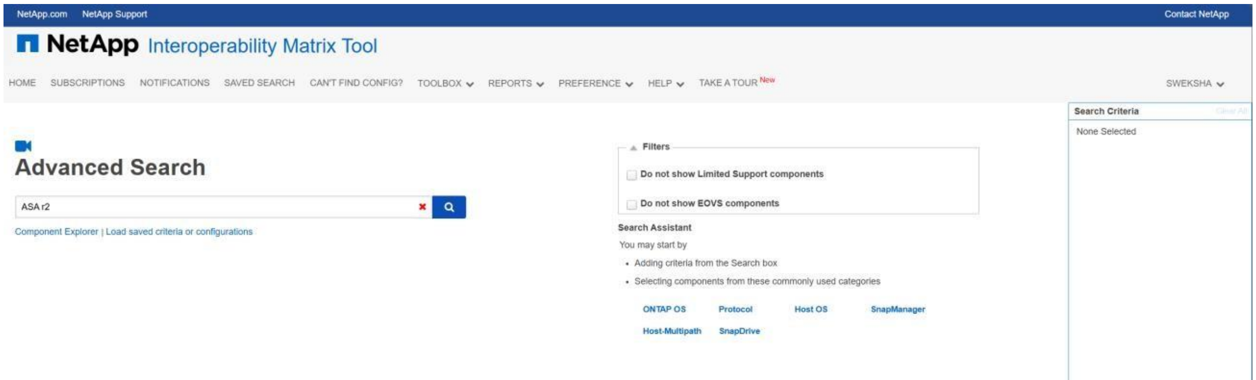

2. Highlight the ASA and click Add next to it to add it to the search criteria.

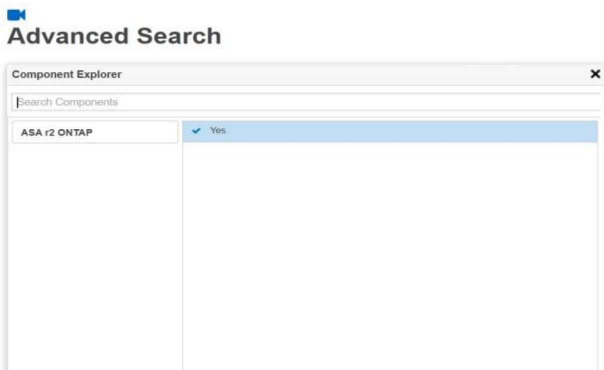

3. Click on Next >> Refine Search Criteria link in the middle to bring up the search refinement dialog.

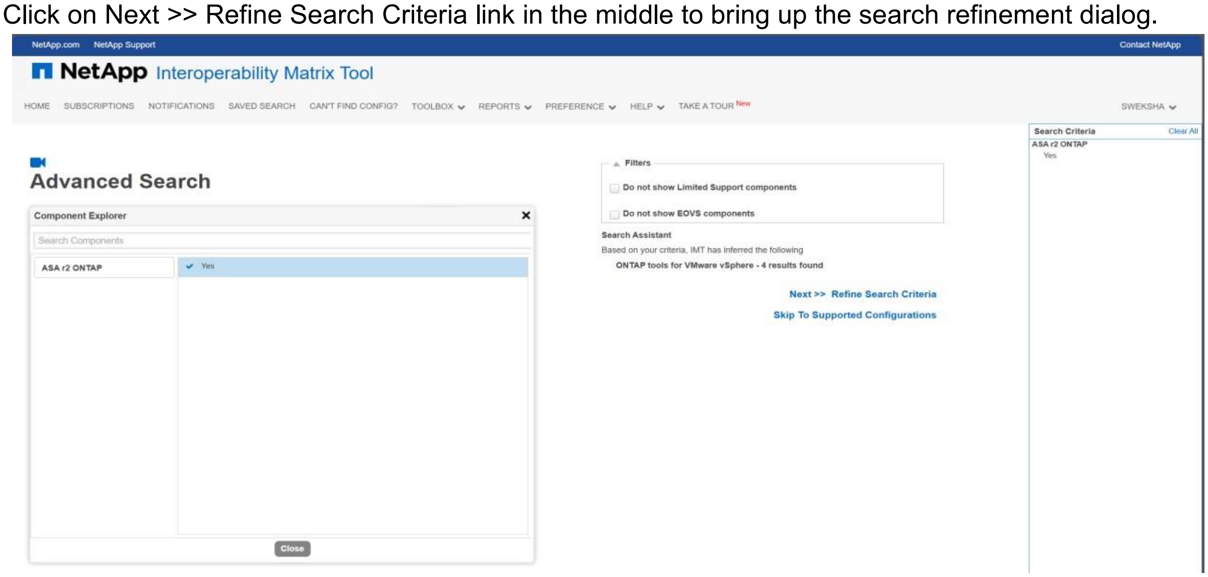

4. In the Refine Search Criteria window, select the desired ONTAP version, storage protocol, and host operating system for the ASA. Click the arrow on the right to view the supported configurations.

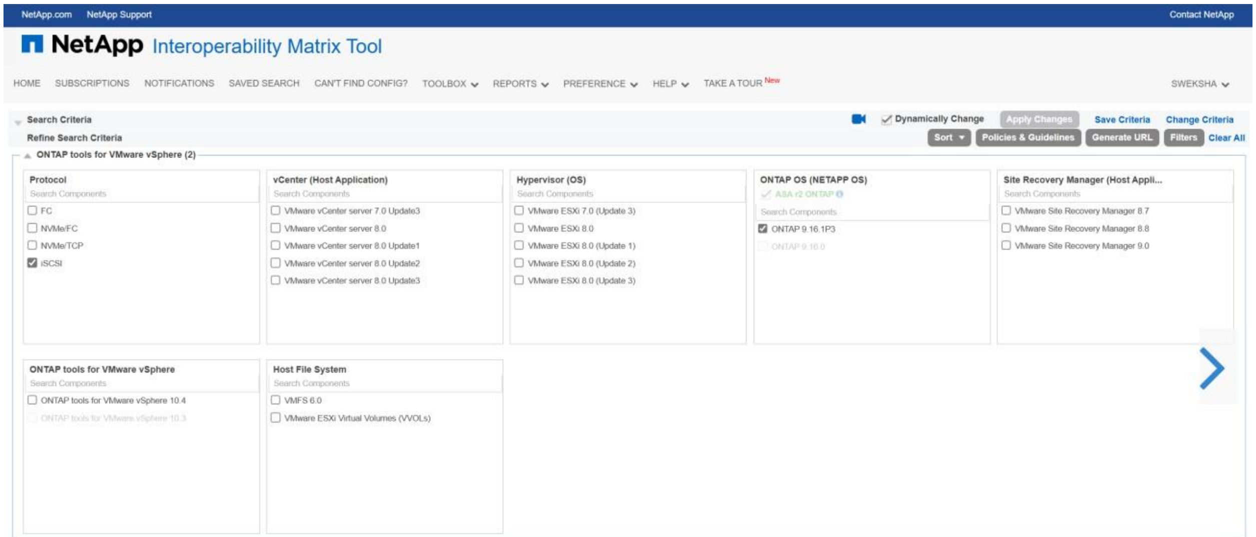

5. The supported configuration is shown along with notes related to the supported configuration.

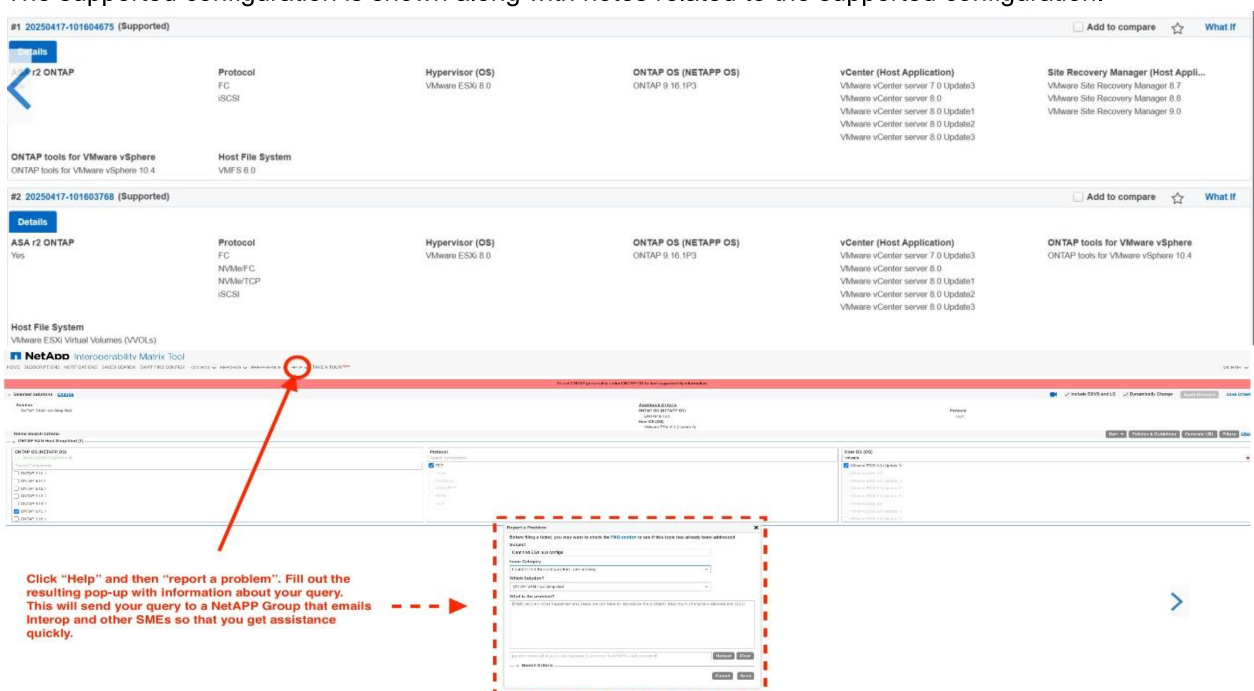

Note: Use “report a problem” from the help menu if you have questions about your query or need assistance finding a specific configuration

Broadcom compatibility guide (BCG)

To check for the compatibility of the VMware solution ecosystem, go to Broadcom Compatibility Guide (BCG) website: https://compatibilityguide.broadcom.com/ .

Using the compatibility guide helps organizations align their ASA infrastructure with VMware’s support matrix, avoiding potential integration issues and ensuring access to technical support and updates. It also supports planning for upgrades, ensuring that firmware, drivers, and VMware software versions remain in sync.

In the ASA when deploying VMware-based solutions, the Broadcom (formerly VMware) Compatibility Guide is a critical resource for ensuring hardware and software interoperability. This guide provides detailed validation for servers, I/O devices, storage arrays, and SAN components to ensure they are fully supported with specific VMware ESXi versions and features. Within NETAPP ASA, adherence to the Broadcom Compatibility Guide ensures that components—such as servers with supported CPUs, HBAs, NICs, and storage controllers—are certified to operate reliably in a virtualized VMware environment.

VMware virtual infrastructure

VMware virtual infrastructure provides a platform where multiple virtual machines (VMs) can run on a single physical server running the ESXi hypervisor, leveraging centralized storage and advanced networking capabilities for enterprise-grade performance and scalability.

SAN boot

Enables an ESXi host to boot its operating system directly from a SAN LUN rather than from local disks. This setup allows for diskless server architecture, which simplifies hardware replacement, enhances security by centralizing storage, and streamlines backup and disaster recovery processes. A second and possibly even bigger advantage of using SAN boot configurations is that administrators can perfect a “golden image” and then clone that as needed to create identical configurations for either scale or as a very easy way to replace or upgrade ESXi servers as needed.

SAN fabrics

This refers to the network topology and configuration of storage networks, typically built using Fibre Channel (FC) or Fibre Channel over Ethernet (FCoE). These fabrics include components such as HBAs (Host Bus Adapters), FC switches, and zoning configurations they provide multiple redundant paths between hosts and storage arrays to ensure high availability and load balancing.

Virtual switches

Virtual switches are software-defined network switches within the ESXi host that connect VMs to each other and to the physical network. They support VLAN tagging (IEEE 802.1Q), NIC teaming for failover and load balancing, traffic shaping, and other Layer 2 features to manage and secure VM traffic.

Virtual NICs

Virtual network interface cards (vNIC) assigned to VMs, allowing them to send and receive network traffic via the virtual switch. Each vNIC is mapped to a physical NIC (pNIC) on the host, depending on the vSwitch configuration, and can be configured with specific MAC addresses, VLANs, and traffic policies.

Jumbo frames

Jumbo frames are Ethernet frames with a payload size greater than the standard 1500 bytes—commonly configured up to 9000 bytes. Jumbo frames must be enabled end-to-end (on vNICs, vSwitches, pNICs, and network switches), if all intermediate devices don’t support and have jumbo frames, session setup and communication won’t work. Jumbo frames reduce protocol overhead and CPU utilization, significantly improving performance for large data transfers such as VMotion, NFS, and iSCSI traffic. These components together create a robust, flexible, and high-performing virtual infrastructure.

Solution availability and ONTAP life-cycle management

The NetApp ASA is purpose-built to deliver exceptional solution availability for enterprise environments where uptime and data integrity are critical. It features an active-active dual-controller architecture, which ensures that if one controller fails, the other seamlessly takes over, maintaining uninterrupted access to data. All hardware components—including power supplies, fans, NVRAM, and SAN ports—are fully redundant to eliminate single points of failure. In-flight data is protected using NVRAM, allowing data to be safely written even in the event of a power outage.

To enhance path resiliency and performance, the ASA uses ALUA (Asymmetric Logical Unit Access) to communicate path and path state changes to the host's MPIO stack. This works in conjunction with VMware's native multipathing or NetApp 's ONTAP Path Selection Plugin to provide automatic path failover and load balancing. The ASA is powered by NetApp 's ONTAP operating system, which includes powerful availability features like Snapshot copies for instant, point-in-time, space-efficient backups; SnapMirror for real-time or scheduled replication to a disaster recovery site; and SnapRestore for rapid full-volume recovery. Additionally, FlexClone technology enables fast, zero-impact cloning for testing or rollback purposes.

The ASA can also be deployed in high-availability configurations within the same data center or across sites using SnapMirror synchronous replication, ensuring business continuity even during site-level disruptions. Altogether, these features make NetApp ASA a robust and highly available solution for SANbased storage in VMware and other enterprise workloads.

ONTAP storage controller failover

Storage controller failover is used when you would like to add I/O adapters to the storage cluster or upgrade the storage cluster software one node at a time. It is a good test to confirm that you have multipathing configured properly throughout your SAN ecosystem.

To exercise storage failover, you can shut down or reboot one of the controller nodes from ONTAP System Manager using the steps below. We will confirm the SAN paths for a datastore LUN to validate SAN path availability and iSCSI LIF failover.

1. Login to vCenter.
2. Select one of the ESXi hosts in the data center Inventory view.
3. Go to Configure tab, select Storage Devices, select one of the storage devices, and then click on the Paths tab below the device list to see the paths.

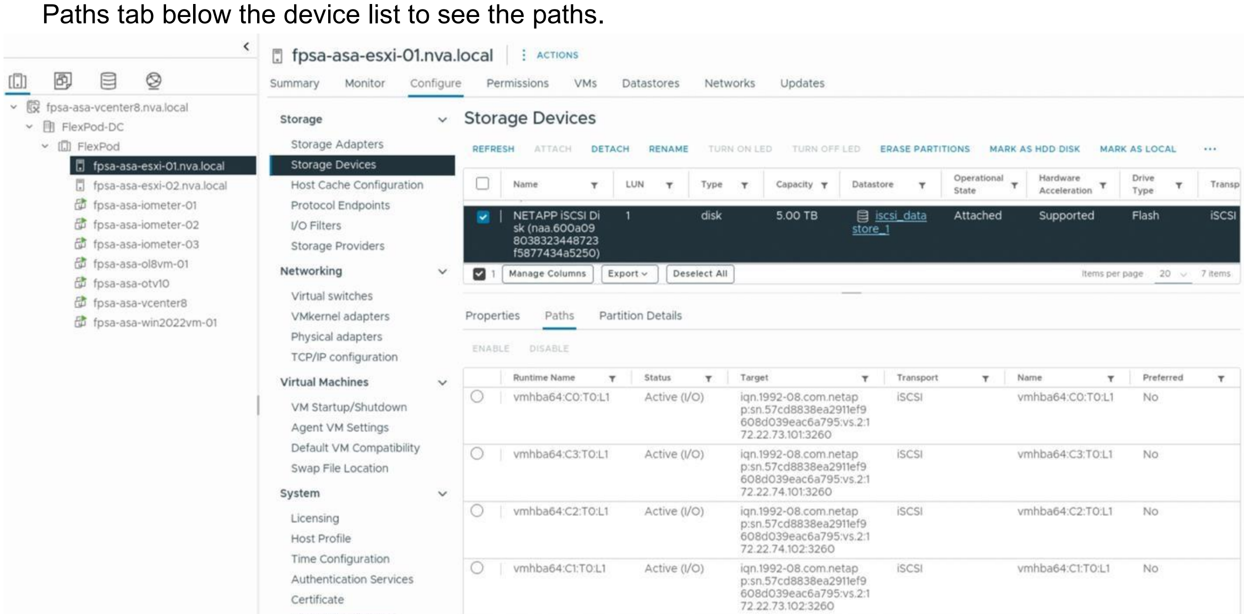

4. Confirm that all four paths to the LUN are showing Active(I/O) status in normal condition.
5. Login to ONTAP CLI.
6. Check iSCSI LIF information to confirm that each node has two iSCSI LIFs located in their respective VLAN ports as shown in the example below.

```
fpsa-a50-u0909::> net int show -lif iscsi*
  (network interface show)
```



```
             Logical    Status     Network            Current       Current Is
Vserver      Interface  Admin/Oper Address/Mask       Node          Port    Home ----
------- ---------- ---------- ------------------ ------------- ------- ---- svm1
             iscsi-lif-01a up/up   172.22.73.101/24   fpsa-a50-u0909-01 e2b-2273 true
iscsi-lif-01b up/up   172.22.74.101/24   fpsa-a50-u0909-01 e4b-2274 true
iscsi-lif-02a up/up   172.22.73.102/24   fpsa-a50-u0909-02 e2b-2273 true
iscsi-lif-02b up/up   172.22.74.102/24   fpsa-a50-u0909-02 e4b-2274 true 4 entries
were displayed.
```

7. Login to ONTAP System Manager.
8. Select Cluster > Overview from the left menu to navigate to the Cluster Overview page.
9. Scroll down to the Nodes section.
10. Click on the three dots next to a node and select either Shut down or Reboot from the menu.

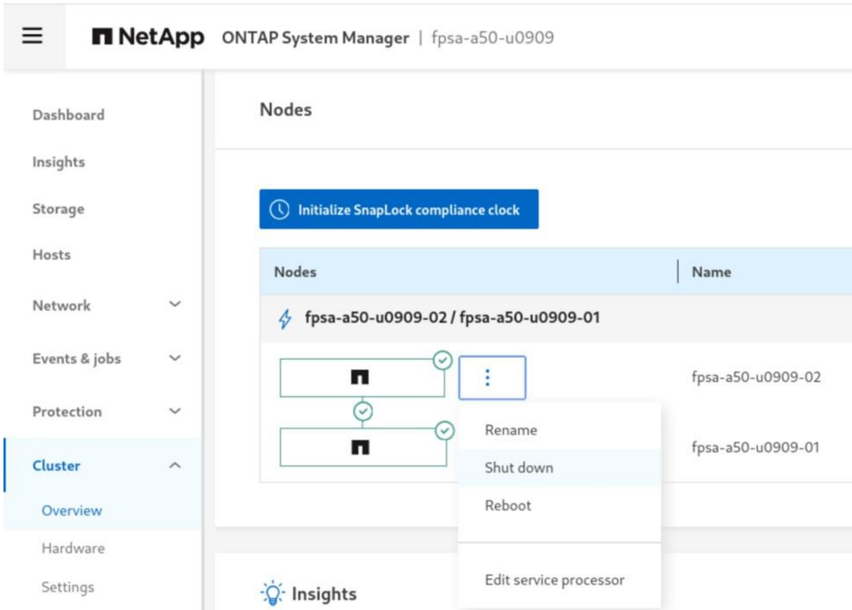

Note: When a node is rebooted or shut down, the partner controller will automatically takeover the workload previously served by the node which was getting rebooted or shut down.

11. Provide a reason for performing the operation. For example, shutting down a node to add an adapter for additional connectivity.

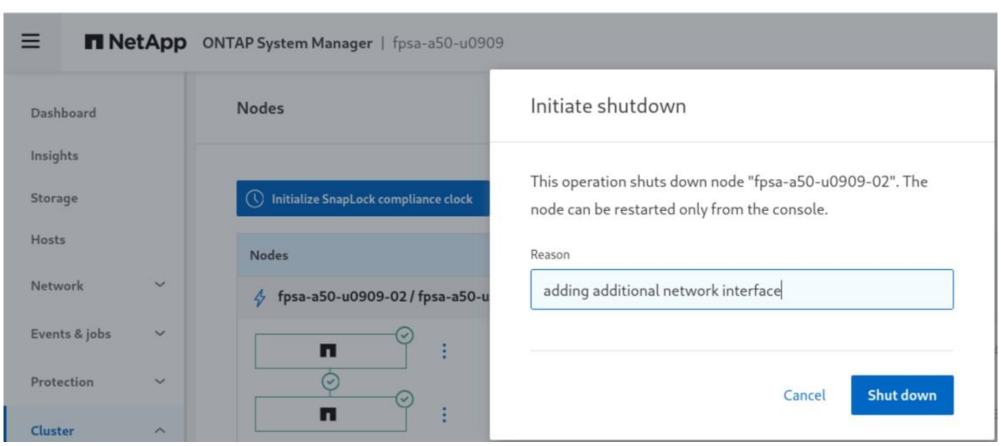

Note: After a shutdown, the node can be started only from the console or remote connection to the node's service processor.

12. ONTAP System Manager will indicate that a takeover happened. In this case, node 02 was shut down and node 01 has taken over node 02.

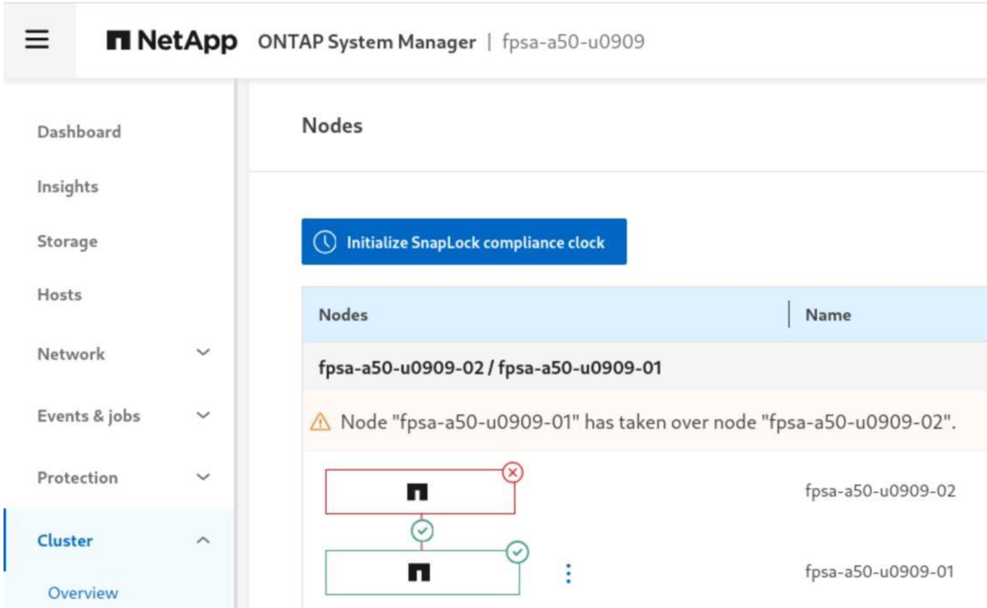

13. Go back to vCenter.
14. Select the same ESXi host from the data center Inventory view.
15. Go to Configure tab, select Storage Devices, select one of the storage devices, and then click on the Paths tab below the device list to see the paths.
16. Confirm that all paths to the same LUN are still showing Active(I/O) status which indicates that iSCSI LIF migration happened.

Note: If you see two of the paths showing Dead status, please go back to the iSCSI LIF configuration section for steps on properly configuring iSCSI-only LIFs with failover enabled.

17. Go back to ONTAP CLI.
18. Check iSCSI LIF information to confirm that the LIFs originally located in node 02 were migrated to node 01 as shown in the example below. The migrated LIFs should show false for the Is Home status.

```
 fpsa-a50-u0909::> net int show -lif iscsi*
(network interface show)
            Logical    Status     Network            Current       Current Is
Vserver     Interface  Admin/Oper Address/Mask       Node          Port    Home -------
---- ---------- ---------- ------------------ ------------- -------- ---- svm1
            iscsi-lif-01a up/up   172.22.73.101/24   fpsa-a50-u0909-01 e2b-2273 true
iscsi-lif-01b up/up   172.22.74.101/24   fpsa-a50-u0909-01 e4b-2274 true
iscsi-lif-02a up/up   172.22.73.102/24   fpsa-a50-u0909-01 e2b-2273 false
iscsi-lif-02b up/up   172.22.74.102/24   fpsa-a50-u0909-01 e4b-2274 false 4 entries
were displayed.
```

19. When you are ready to boot node 02 up again, connect to the console of node 02 and issue autoboot command from LOADER prompt.

```
LOADER-B> autoboot
```

20. After a few minutes, the cluster should be back to normal status again. You should check the paths to devices on the ESXi hosts to confirm that all paths are in Active(I/O) status. In addition, check iSCSI LIF information with ONTAP CLI to confirm that the migrated iSCSI LIFs are now back to their home node.

Note: For the use case of adding additional network interface cards into the storage cluster, you can then proceed to go through the process of shutting down the other node to add the interface card in that node before booting it up again.

ONTAP software upgrade

ONTAP storage cluster software and firmware upgrades utilize the storage failover mechanism to upgrade the nodes one at a time. Before performing an upgrade, be sure to confirm the ecosystem interoperability for the upgrade destination combination using the process discussed previously.

To perform ONTAP upgrade, follow the steps below.

1. Login to NetApp support site [https://support.NetApp .com](https://support.NetApp.com) with your credential.
2. Go to the [ONTAP 9 download link](#).
3. Click on the Download Latest Release link.

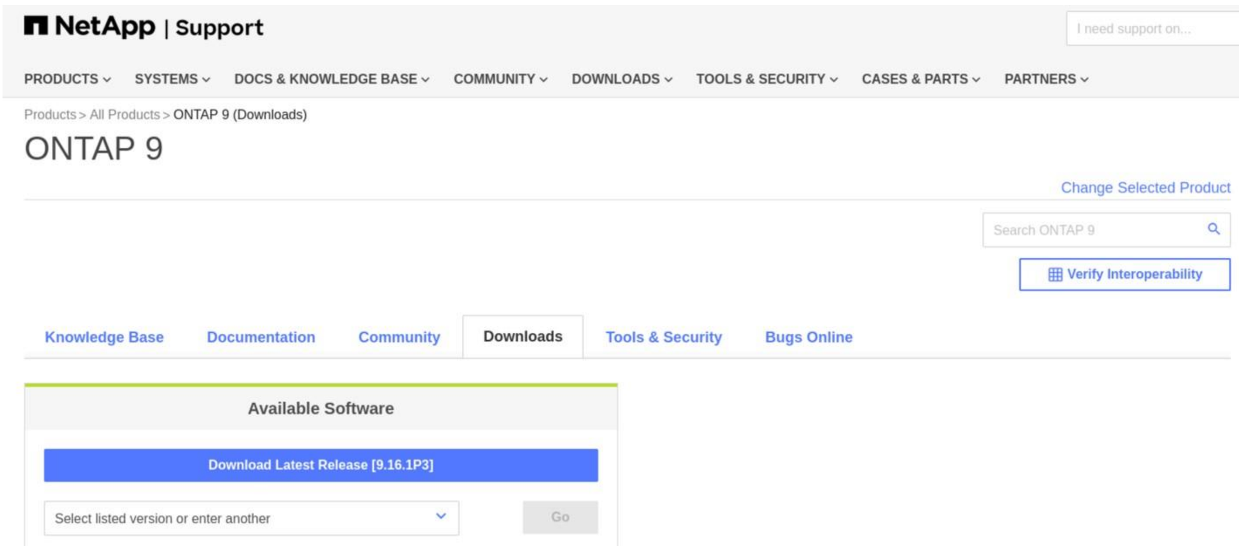

4. Review the Caution/Must Read information and End User License Agreement. Click on the box below the information to acknowledge and accept. Click ACCEPT & CONTINUE.
5. Review the information regarding Restrictions on Encryption Technology and select the appropriate ONTAP 9 image to download.
6. Login to ONTAP System Manager.
7. Expand the Cluster menu on the left and select Overview page.
8. Click on the Update button under ONTAP version or click on the three dots on the right and select ONTAP update.

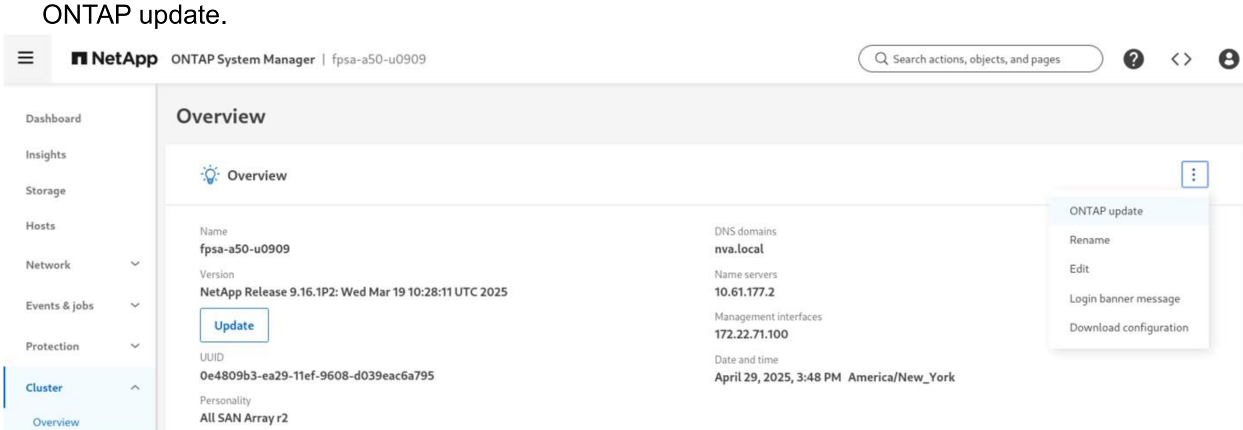

9. On the Software updates screen, click on +Add image and download the image from the server or browse to select the downloaded ONTAP image for upload from the local client.

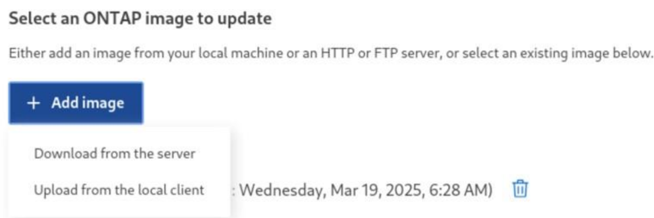

10. After the ONTAP image has been uploaded, click to select the uploaded image and click Update at the bottom of the page.

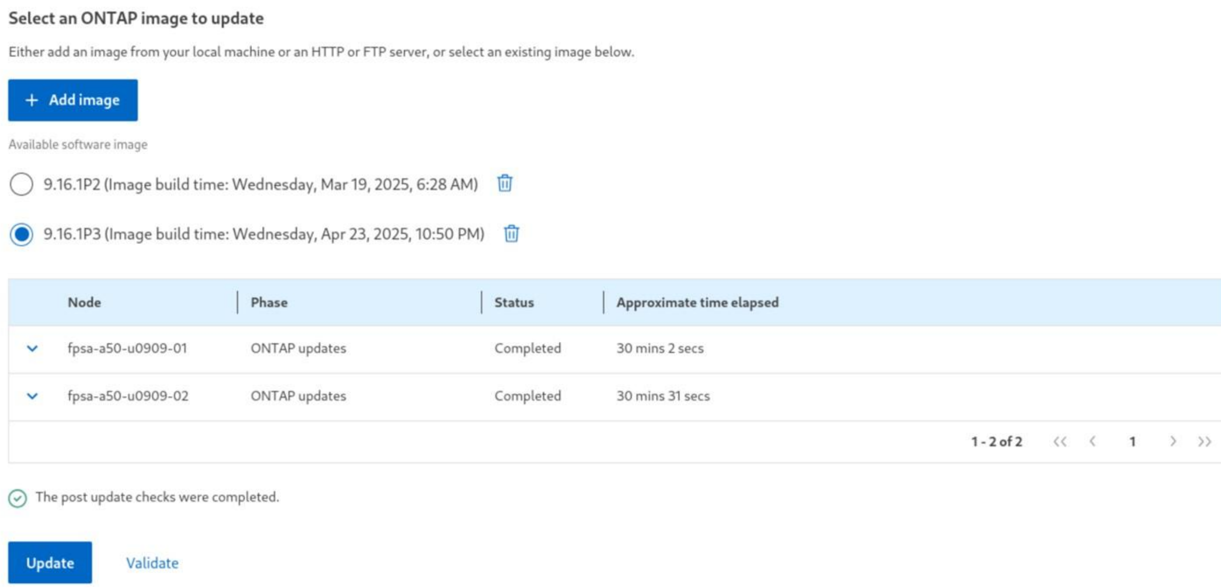

11. A series of pre-update checks will be performed. Afterwards, and review the resulting messages from the pre-update checks.

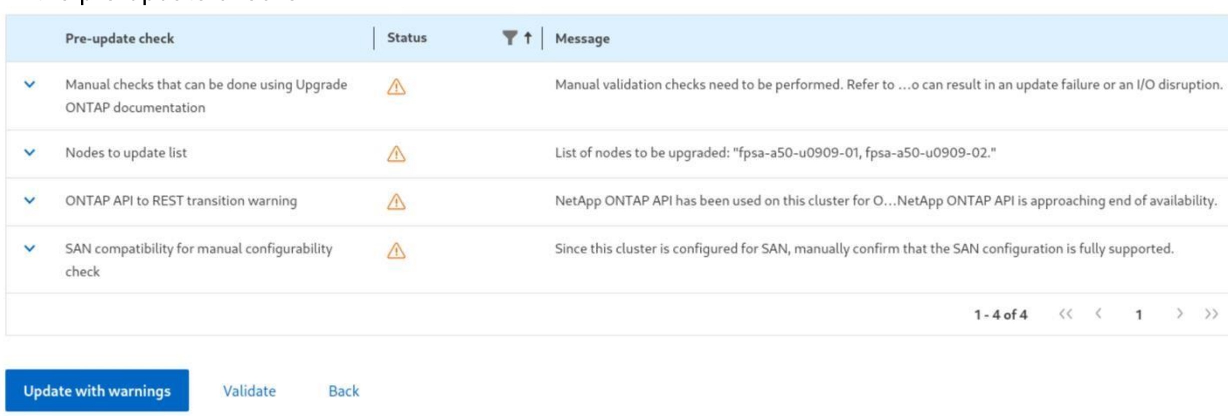

12. Address any errors or warnings and click Validate again if needed or click Update with warnings to proceed with update.
13. During the ONTAP update process, you can click on the down arrow in front of the controller node names to expand the details of the update status.

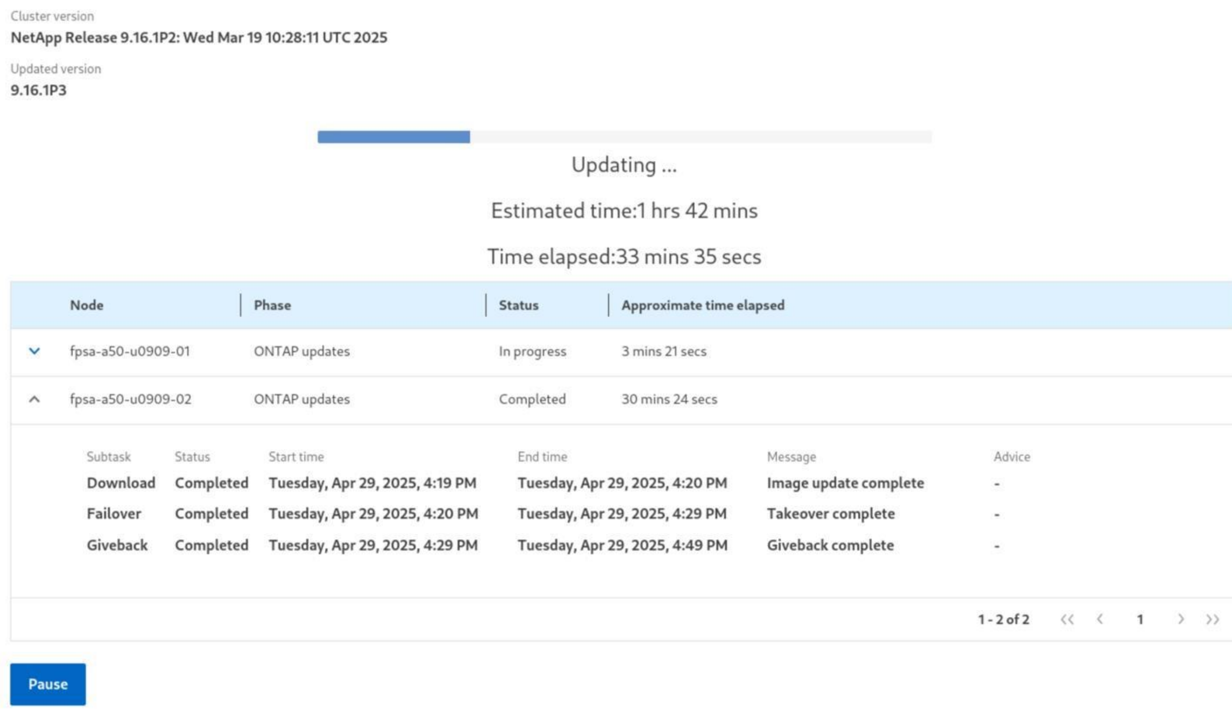

Note: For a single HA pair configuration, the updates for controllers happen sequentially for the nodes through a takeover and giveback process after completing image update for the new image to take effect. Client IO will continue during the non-disruptive ONTAP upgrade process, assuming SAN multipathing had been properly configured as described in this Technical report.

14. Wait for the upgrade process to complete and optionally delete the prior ONTAP image.

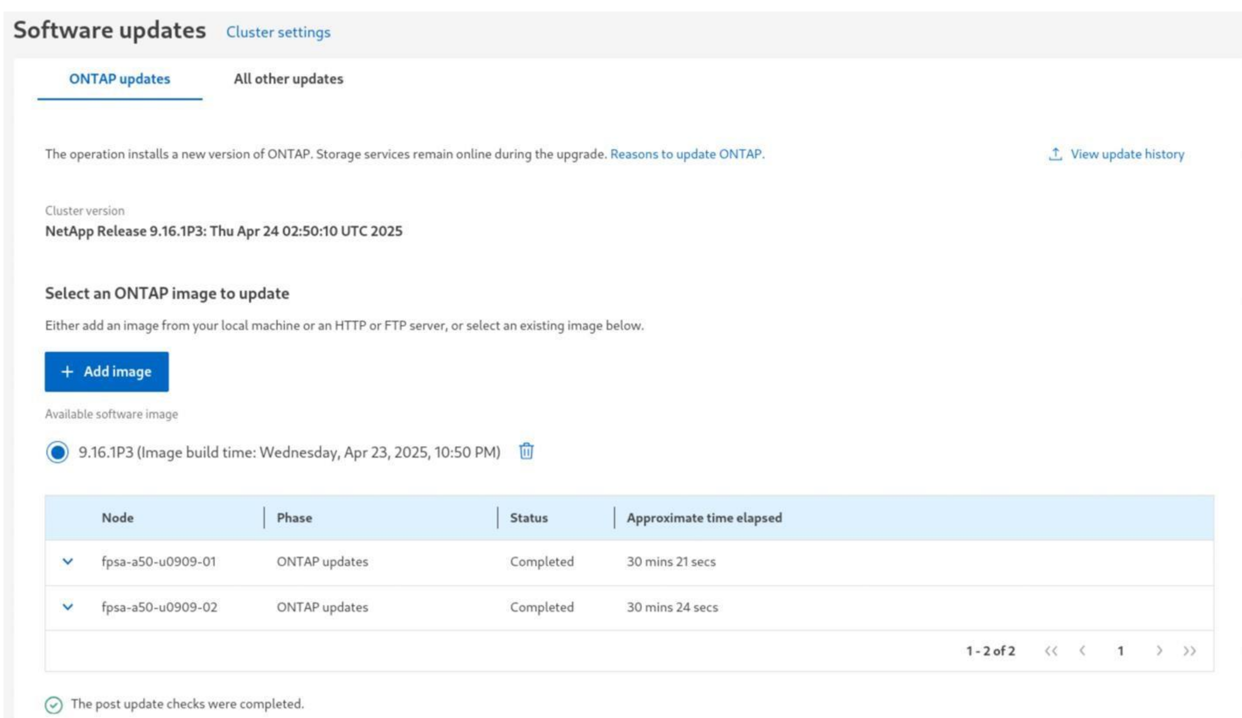

Data protection for SAN solutions

Data protection in ASA is achieved through a layered approach combining hardware-level resilience, SAN-based features, and integration with backup and disaster recovery solutions.

Snapshots

In ONTAP, a snapshot is a space-efficient, point-in-time image of a SAN-based storage volume or LUN, enabling rapid data protection and recovery without impacting performance. Snapshots are created directly on the storage array using copy-on-write or redirect-on-write mechanisms, which preserve the original data state while only storing changes made after the snapshot. This minimizes storage consumption and allows for near-instantaneous creation. Snapshots in ASA are fully integrated with SAN management tools and can be scheduled at regular intervals with automated retention policies, supporting stringent Recovery Point Objective/ Recovery Time Objective (RPO/RTO) requirements. They provide application-consistent backups when integrated with platforms like VMware, Microsoft SQL Server, or Oracle. Writable clones of snapshots are also used for testing and development without affecting production environments. Additionally, snapshots can be replicated to secondary storage systems as part of a broader disaster recovery strategy. Managed centrally within the ASA framework, snapshots enhance data resilience, ensure operational continuity, and support agile, low-impact data recovery operations in enterprise environments.

Consistency groups

A collection of Storage Units (SUs) that are managed as a single unit. CGs ease management, and protection of application workloads spanning multiple volumes. Each LUN / Namespace (NS) in a CG except for vVol metadata objects will be in a distinct FlexVol. New APIs to map/unmap all the LUNs/NSs of a CG is supported. Consistency groups align with ASA goals of modularity, high availability, and centralized control by supporting coordinated data protection, streamlined failover, and simplified recovery processes across enterprise workloads.

SnapCenter

In ONTAP, SnapCenter plays a critical role in delivering application-consistent data protection and centralized backup management across SAN-based infrastructures. SnapCenter is a NetApp data protection software offering designed to orchestrate backup, restore, and cloning operations using secondary storage systems. Support for ASA was introduced in SnapCenter starting with version 6.1. Within the ASA architecture which emphasizes modular, high-availability, and centrally managed solutions SnapCenter ensures efficient, policy-driven protection of enterprise applications such as Oracle, Microsoft SQL Server, SAP HANA, VMware and more.

SnapCenter Plugin for VMware vSphere

In the ASA, the SnapCenter Plug-in for VMware vSphere (SCV) provides a streamlined, efficient solution for protecting virtual machines through deep integration with both NetApp storage systems and VMware vCenter. Support for ASA was introduced in Snapcenter plugin for VMware vSphere starting with version 6.1. It leverages NetApp ONTAP's storage-based snapshot capabilities to perform fast, space-efficient, and crash-consistent backups of VMs hosted on SAN datastores, without impacting performance. The plug-in enables policy-driven automation for scheduling, retention, and replication, and supports granular recovery options—including full VM, individual VMDK, or file-level restores. SCV is fully integrated into the vSphere Web Client, allowing administrators to manage backup and recovery directly from the VMware interface, while maintaining centralized visibility and control through SnapCenter. It also supports application-consistent backups when used alongside other SnapCenter application plug-ins, ensuring data integrity for VMs running critical workloads. By aligning with ASA principles of centralized management, high availability, and modular scalability, the SnapCenter Plug-in for VMware vSphere plays a vital role in delivering robust, enterprise-grade virtual infrastructure protection.

SnapMirror Active Sync

SnapMirror Active Sync is a feature in NetApp ONTAP used primarily for synchronous data replication between two storage systems to provide zero data loss and automatic application failover in case of site failures. It's particularly designed for business-critical applications requiring high availability (HA) and low Recovery Time Objective (RTO) and Recovery Point Objective (RPO).

In the context of NetApp ASA (All SAN Array), which is optimized for block-based workloads (especially SAN environments using iSCSI and FC), SnapMirror Active Sync plays a crucial role in business continuity and disaster recovery (BCDR) strategies.

Note: SnapMirror is only supported between NetApp ASA New clusters

Appendix A: NetApp ONTAP tools for VMware vSphere deployment

ONTAP tools for VMware vSphere (OTV) is a set of tools for virtual machine lifecycle management. It is a collection of scalable, event-driven, microservices deployed as a virtual appliance and it integrates the VMware ecosystem and ONTAP storage to provide the following functionalities.

- Virtual machine functionality like basic protection and disaster recovery
- ASA Provider for VM granular management
- Storage policy-based management
- Storage Replication Adapter (SRA)

Appendix B: ONTAP tools for VMware vSphere 10.4 preinstallation considerations

Before deploying, please review the prerequisites information to determine the specific deployment configuration that suits your needs. For example:

- Select thin-provisioned or thick-provisioned storage
- Select a deployment type: non-HA (small / medium) or HA (small / medium / large)

See the table in the prerequisite page linked above for the CPU, memory, storage requirements for the different deployment types and the number of ESXi hosts and vVols they support. For this validation, we are using Non-HA Small deployment type. **Error! Reference source not found.** highlights the associated resources requirements and supported limits.

Table 3) Non-HA small OTV deployment resource requirements and limits

Resources	Requirement / Limit
CPU requirement	9
Memory requirement (GB)	18
Disk requirement (GB) thick provisioned	350
Number of vVols supported	~12,000

Number of ESXi hosts supported	32

To facilitate proper operations and communications for the deployment, please refer to the network port requirements listed on the prerequisites page under the port requirements section. Please ensure that the necessary network configurations are in place to permit the needed traffic within your network for the associated services to function correctly.

The pre-deployment check section highlights the need to already have vCenter deployed and the login information for deploying the tool into vCenter. In addition, the three IP addresses for load-balancer, Kubernetes control plane and worker node in the Non-HA Small deployment type need to be already configured in DNS.

Refer to https://docs.NetApp .com/us-en/ontap-tools-vmware-vsphere-10/deploy/ontap-tools-deployment.html for OTV installation. For detailed step-by-step installing refer this NVA https://www.NetApp .com/media/135828-flexpod-san-ucs-xdirect-NetApp_ASA-nva-deploy-guide.pdf

Appendix C: Install NetApp Windows Host Utilities

Install pre-requisites for NetApp Windows Host Utilities installation

NetApp recommends installing the **latest hotfixes, cumulative updates,** and security updates that are available from Microsoft.

1. Download the updates from the Microsoft Update Catalog 2022 and Microsoft Security Update links below.

   https://www.catalog.update.microsoft.com https://msrc.microsoft.com/update-guide/

2. Follow the instructions provided by Microsoft to install the hotfixes and updates.
3. Enable Windows Updates as needed from your host and install available updates.

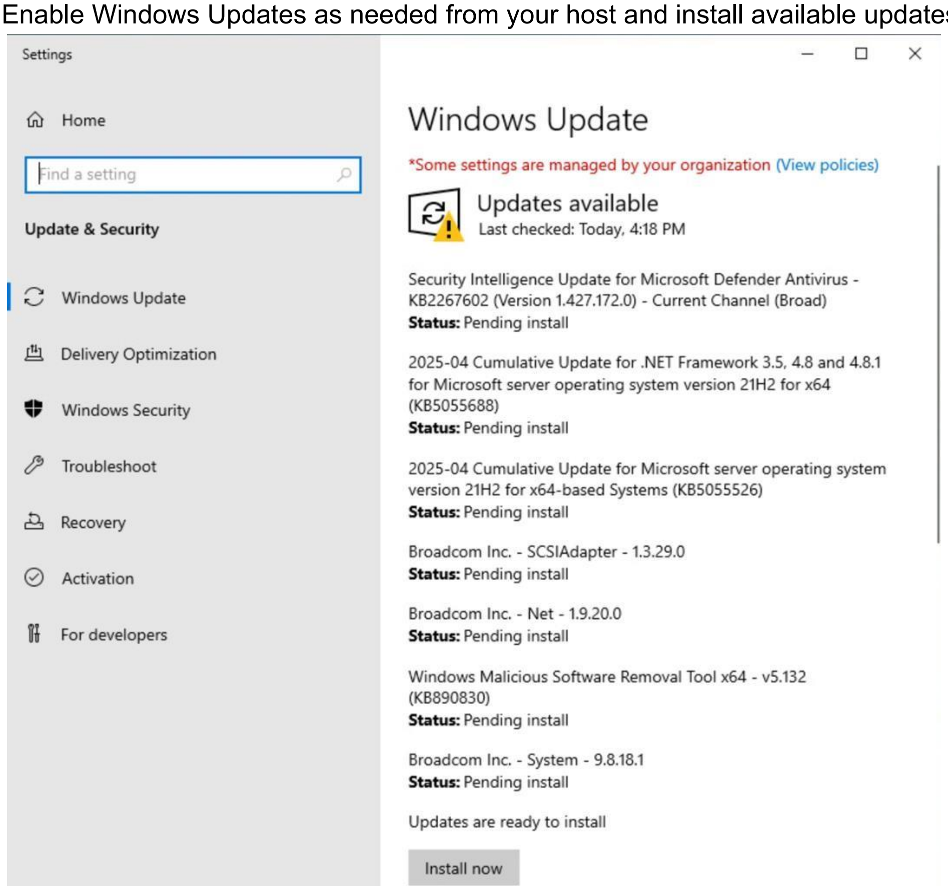

4. Many hotfixes and updates require a reboot of your Windows host. Reboot the host if instructed.

Before you install the Host Utilities, you should verify that your host and storage system configuration are supported. Please refer to the information in the following page. https://docs.NetApp .com/us-en/ontap-sanhost/hu_wuhu_72.html

Download NetApp Host Utilities Version 7.2 for Windows from NetApp support site download link below. https://mysupport.NetApp .com/site/products/all/details/hostutilities/downloadstab/download/61343/7.2/downloads

Appendix E: Install NetApp Linux Host Utilities

1. Download the NetApp Linux Unified Host Utilities 7.1 tool from https://support.NetApp .com.

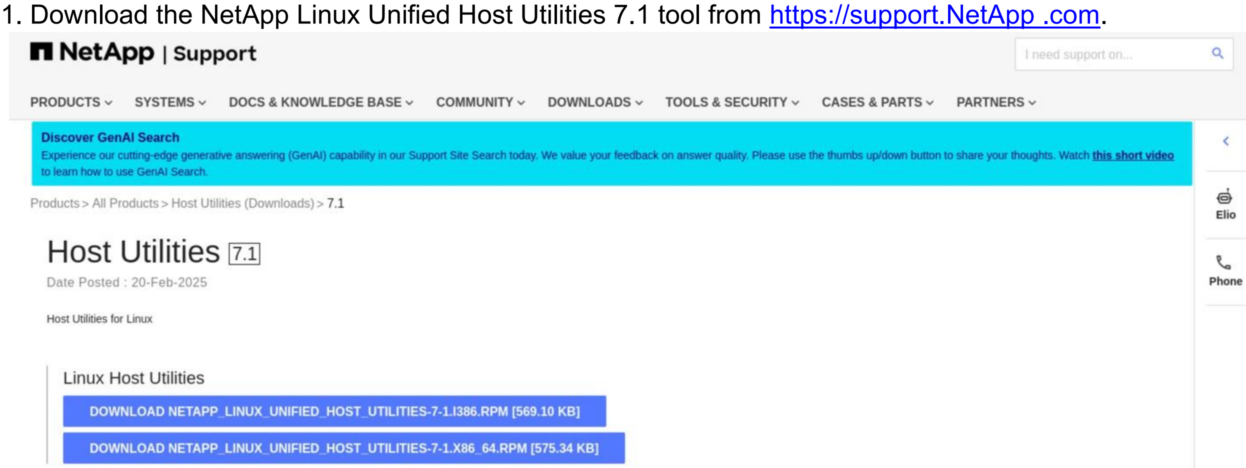

Note: You will need a NetApp support account and login with your username and password.

History

Version	Date	Document version history
Version 1.0	July 2025	Initial release

Where to find additional information

- NetApp ASA documentation: https://docs.NetApp .com/us-en/NetApp_ASA-r2/index.html
- NetApp ASA datasheet: https://www.NetApp .com/media/85736-ds-4254-NetApp_ASA.pdf
- TR-4080: Best practices for modern SAN: https://www.NetApp .com/media/10680-tr4080.pdf
- FlexPod with NetApp ASA NVA: https://www.NetApp .com/media/135828-flexpod-san-ucs-xdirectNetApp ASA-nva-deploy-guide.pdf
- FlexPod datacenter zero trust framework CVD: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/UCS_CVDs/flexpod_zero_trust.html

- FlexPod security hardening TR: https://www.NetApp .com/pdf.html?item=/media/99202tr_4984_flexpod_security_hardening.pdf
- NetApp support site: https://support.NetApp .com
- NetApp hardware universe: https://hwu.NetApp .com
- NetApp interoperability matrix tool: http://support.NetApp .com/matrix
- Broadcom compatibility guide: https://compatibilityguide.broadcom.com/
- NetApp solutions NetApp Solutions
- SAP and SAP HANA SAP Business Application and SAP HANA Database Solutions
- Enterprise databases ONTAP and enterprise applications

Acknowledgements

Marco Schoen, Jyh-shing Chen - NetApp

