P3AG-6342-01ENZ0

FUJITSU Storage
ETERNUS AX series All-Flash Arrays

Best Practices for Business-Critical Workloads

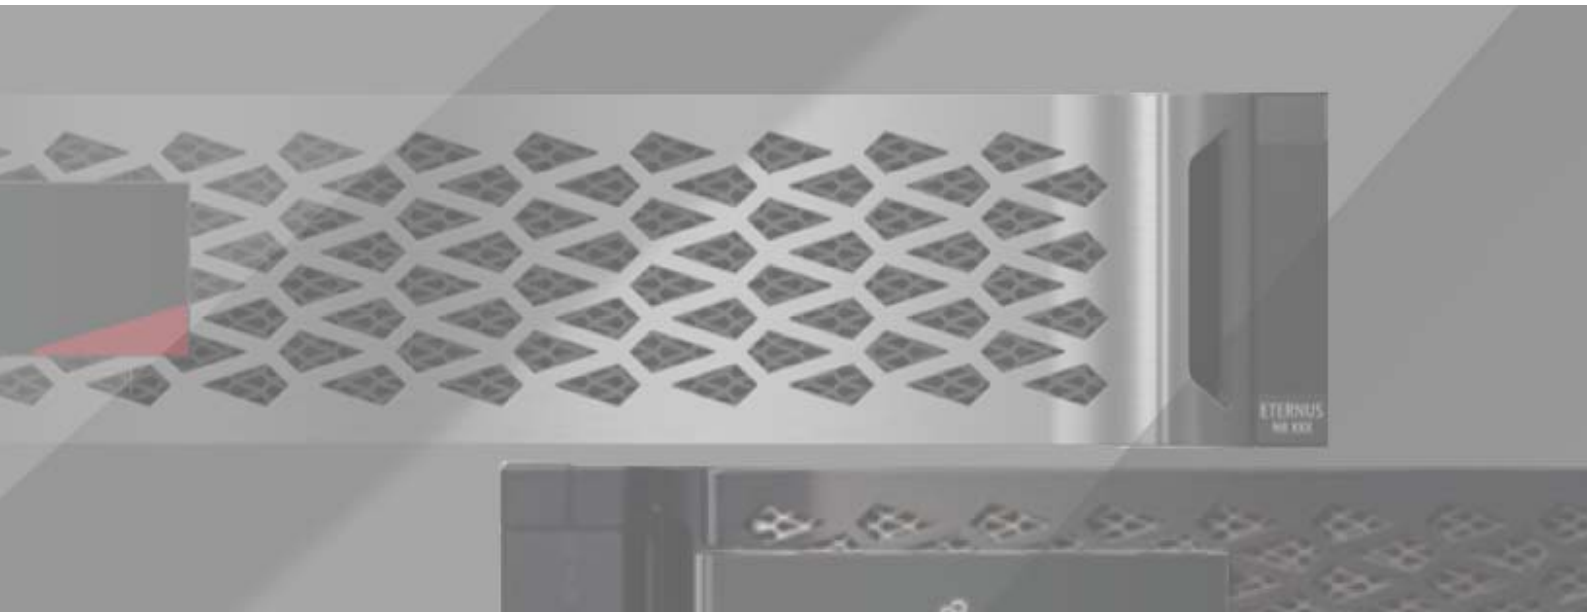

FUJITSU

Table of Contents

List of Figures

List of Tables

Preface

This document introduces ETERNUS AX series All SAN Array (ASA) systems and covers implementation and best practices recommendations for always-available, business-critical SAN configurations. This document corresponds to ONTAP 9.9.1.

First Edition
December 2021

Trademarks

Third-party trademark information related to this product is available at:
https://www.fujitsu.com/global/products/computing/storage/eternus/trademarks.html

Trademark symbols such as ™ and ® are omitted in this document.

About This Manual

Intended Audience

This manual is intended for system administrators who configure and manage operations of the ETERNUS AX, or field engineers who perform maintenance. Refer to this manual as required.

Related Information and Documents

The latest information for the ETERNUS AX is available at:
https://www.fujitsu.com/global/support/products/computing/storage/manuals-list.html

Document Conventions

Notice Symbols

The following notice symbols are used in this manual:

Caution	Indicates information that you need to observe when using the ETERNUS AX. Make sure to read the information.
Note	Indicates information and suggestions that supplement the descriptions included in this manual.

1. ETERNUS AX series Systems Introduction

ASA systems are built on ETERNUS AX series systems, which deliver industry-leading performance and reliability. ETERNUS AX series systems provide an enterprise-class SAN solution for customers who want to consolidate and to share storage resources for multiple workloads.

ETERNUS AX series SAN systems deliver:

- Industry-leading >99.9999% availability
- Massive scale clusters, which scale both up and out
- Industry-leading storage efficiency
- Support for the cloud environment
- Cost-effective seamless data protection

ASA systems build on the all-flash system to deliver continuous SAN availability for enterprises that run mission-critical applications. These systems provide uninterrupted access to data during a planned or unplanned storage failover and deliver streamlined implementation, configuration, and management through a solution that's dedicated only to running tier 1 SAN workloads.

Fujitsu recommends ASA systems when your requirements include:

- Mission-critical workloads such as databases that must have symmetric active-active paths from hosts to storage. All paths between the host and storage are active and optimized across high-availability (HA) partners in this design.
- Preference for a dedicated system to isolate some or all SAN workloads from all others.

ETERNUS AX series systems remain the preferred choice for customers who:

- Need to scale out SAN clusters to up to 12 nodes.
- Need asymmetric access to storage from hosts that are matched with the application requirement.
- Prefer a cluster that supports unified protocols and mixed NAS and SAN workloads.

2. ETERNUS AX series Systems Overview

This document is a detailed guide for storage architects who intend to run business-critical tier 1 workloads on ETERNUS AX series. It details an ASA storage configuration to validate its ability to provide consistent low-latency performance, high throughput, uninterrupted availability, and resiliency. It also discusses best practices for configuring, installing, validating, deploying, and monitoring tier 1 modern SAN storage environments.

This document and its prescriptions are the product of extensive performance testing to identify and to qualify a baseline configuration for consistent performance. It describes this configuration and makes conservative recommendations that are designed to optimize consistent performance. The ASA systems were designed to eliminate All Paths Down (APD) client disruptions that result from a storage failover and to eliminate variability in storage latency and performance, even during storage failover transitions. ASA systems offer uninterrupted availability while maintaining the industry-leading performance. And by concentrating on SAN protocols and features and by excluding NAS protocols and NAS-only features, they also reduce complexity.

Every organization has its own preferences for allocating and for clustering workloads, and for segregating or for integrating their SAN and NAS estates. There's no one best solution; it depends entirely on each company's business objectives, skillset, and technology roadmap. This document presents requirements and recommendations that will enable your IT organization to build systems that maximize performance while maintaining consistent low-latency operations, even during storage disruptions like with controller takeovers and givebacks.

The ASA configuration is optimized for symmetric active-active access and for consistent high performance with low latency.

The guidelines, requirements, and sample results that are enumerated in this document are calculated based on the workload and performance characteristics.

> **Note**
>
> In this document, we use the term tier 1 to refer to mission-critical workloads that can't accept any loss of access to data. Some use the term tier 0 to describe these same critical workloads.

3. Introducing the ETERNUS AX series

Before the release of ONTAP 9.7, all ONTAP controllers featured the architecture that is shown in Figure 1. This architecture advertised routes directly to the controller that hosted the LUN as active-optimized (AO) paths, with all other paths (indirect paths) advertised as active-non-optimized (ANO) paths. Active nonoptimized paths are not preferred and are not used unless no active optimized paths exist.

Figure 1 Unified ONTAP paths

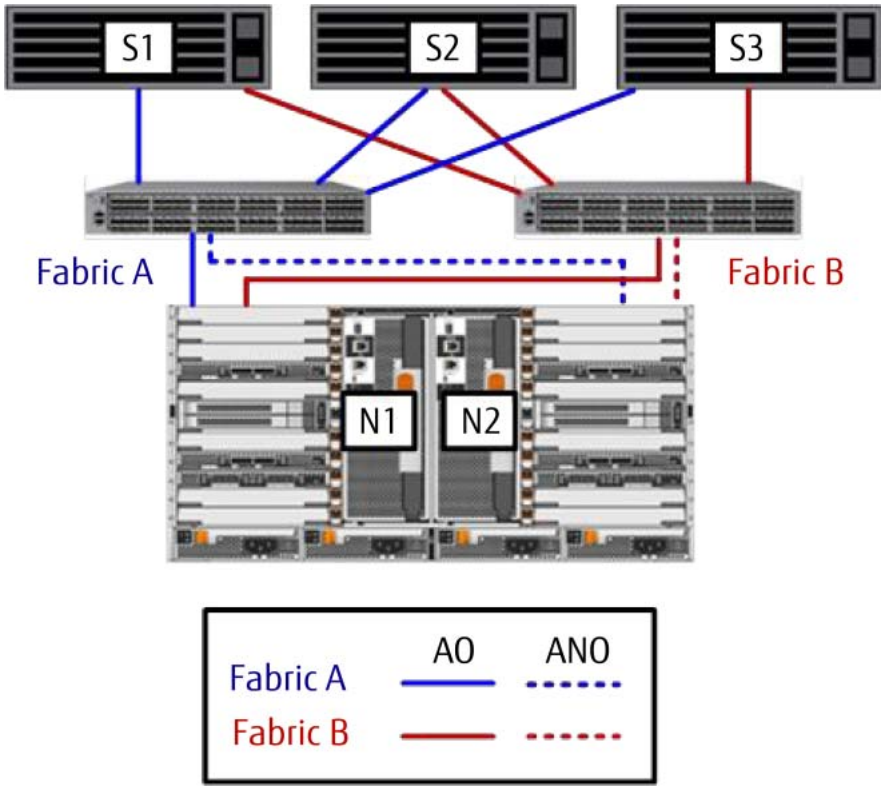

With ONTAP 9.7, Fujitsu introduced ETERNUS AX series ASA systems, which feature symmetric active-active topology, as shown in Figure 2. The ASA supports SAN (block protocols) only and is built on a single HA pair. It currently supports the FC and iSCSI protocols, and ONTAP 9.9.1 supports the NVMe-oF protocol in the ETERNUS AX ASA system.

The defining features of ASA systems include:

- Symmetric active-active operations, which means that all paths are active preferred paths to all LUNs. ASA advertises all paths as AO, which means that there are always active paths to all LUNs, even if a storage failover (SFO, also called a takeover or giveback) occurs. The practical effect is that hosts always have active paths and don't need to query for new paths if an SFO occurs. This feature reduces the impact of an SFO to times that match those on frame-style arrays. Unified clusters advertise both AO and active-non-optimized (ANO) paths.

  Note

  Hosts that connect to a unified cluster see both AO paths (preferred) and ANO paths (not preferred). If the host loses all AO paths and doesn't receive updates that advertise new AO paths, it changes the ANO paths that it still has to a LUN to AO or preferred paths. However, this process can take some time for the host to make those adjustments to its storage map.

- A SAN-only experience that's simplified by the absence of any variables and options that are related to NAS (file) protocols. This feature reduces the skillset that you need to configure, to provision, and to manage the ASA.

- Support in ONTAP 9.7, which also includes a complete streamlined ONTAP System Manager GUI. All aspects of provisioning, configuring, and managing of ONTAP SANs have been significantly simplified.

Figure 2 ETERNUS AX4100 ASA active-active pathing

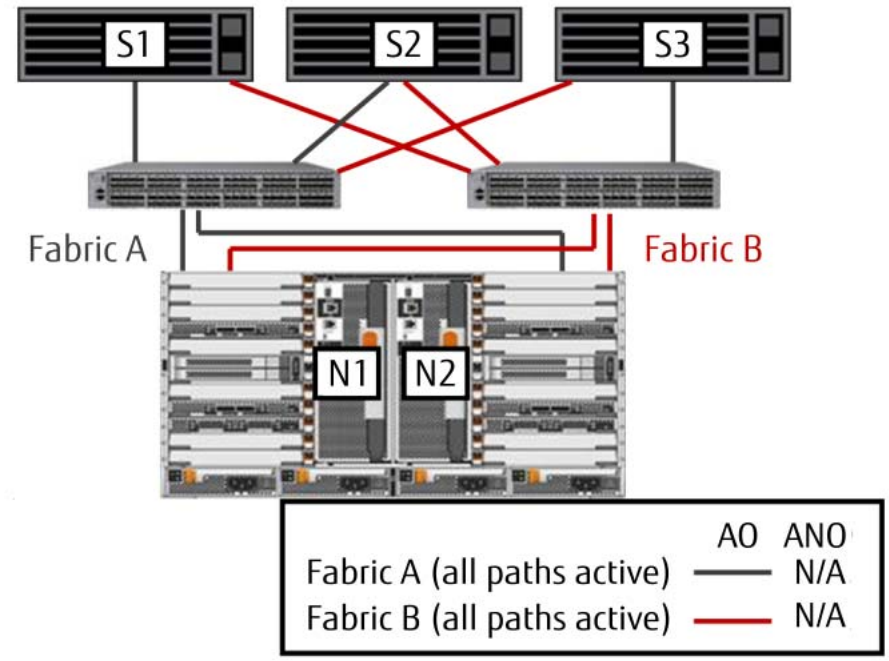

4. ASA Compatibility Guidelines

An ASA that's provisioned to serve business-critical applications can grow with the data storage requirements of your applications.

Applications and storage requirements that fit the following guidelines are an excellent fit for ASAs running current versions of ONTAP:

- Application architects should consider ETERNUS AX series for workloads in which continuous availability and consistent low-latency performance are more important than attaining the maximum possible steady-state throughput. For a discussion about performance optimization and consistent low latency, review the section titled "Steady-state Storage Utilization" (page 20) of this document.
- The ASA symmetric active-active architecture neutralizes the impact of planned and unplanned storage failovers or other component failures. In particular, because of the symmetric access that ASA provides to all LUNs, even with a path, fabric, network, or other failure, a well-designed and managed ASA still provides continuous, consistent, low-latency data access.

ASA Commitments and Service-Level Objectives

ASA service-level objectives (SLOs) are geared toward reducing failover times to an absolute minimum. By changing the ONTAP block architecture to all paths active and by using all controllers, ASA can offer symmetric active-active access to data with no client disruptions from APD. ASA also provides virtually instantaneous and nondisruptive failovers.

When comparing recovery times, measurement protocols matter. Fujitsu testing focuses on I/O resume times from the host's point of view (I/O resume time, or IORT). It is inadequate to measure recovery time by measuring transition times on the partner node. To really quantify the impact and disruption that an SFO causes, you must measure I/O resume time at the operating system or application level.

With ASA symmetric paths, we found no outages when storage failovers occurred, because hosts always have active paths to the LUNs to which they read and write data. With non-ASAs, testing showed variations in different host OS I/O stacks. The length of those disruption windows varied based on the OS, applications, and specific OS or application settings.

In fact, takeover and outage windows are primarily affected by the host OS. The takeover and pathing performance of many OSs can be improved to more quickly react to a loss of active paths by adjusting host I/O timeout thresholds-most of these tweaks were added to OS defaults on the most modern versions of those OSs. You can discover and review many of these configuration tweaks by reviewing the host utilities documentation associated with the OS you are interested in.

5. ONTAP 9.8: New SAN Features

ONTAP 9.8 added a number of new features, some of these are available on both unified and ASA platforms. Others are initially being introduced on the SAN-only ASA. The expectation is that these will likely be added to unified ONTAP clusters in upcoming releases of ONTAP. The new ONTAP SAN features include:

- Virtual machine identifier (VMID)-a virtual machine (VM) telemetry enhancement

VMID

VMware generates a globally unique identifier for each of the VMs it is hosting. It writes these identifiers to the header field to tag each FC frame so that it can be correlated to a specific VM. This feature allows administrators to identify and follow I/Os from each VM using a shared FC LUN-backed datastore. Prior to the VMID feature, the highest level of granularity possible was the ability to identify the datastore with which the I/O was associated. With VMID, administrators can identify and track I/O characteristics for each of the VMs sharing a datastore individually. This functionality allows for end-to-end QoS, which allows for significantly better insight into traffic patterns, workload characteristics, substantially enhanced troubleshooting, and more detailed VM traffic analysis and reporting. This functionality is initially supported with Brocade switches (Gen 6 and later).

Figure 3 VMID

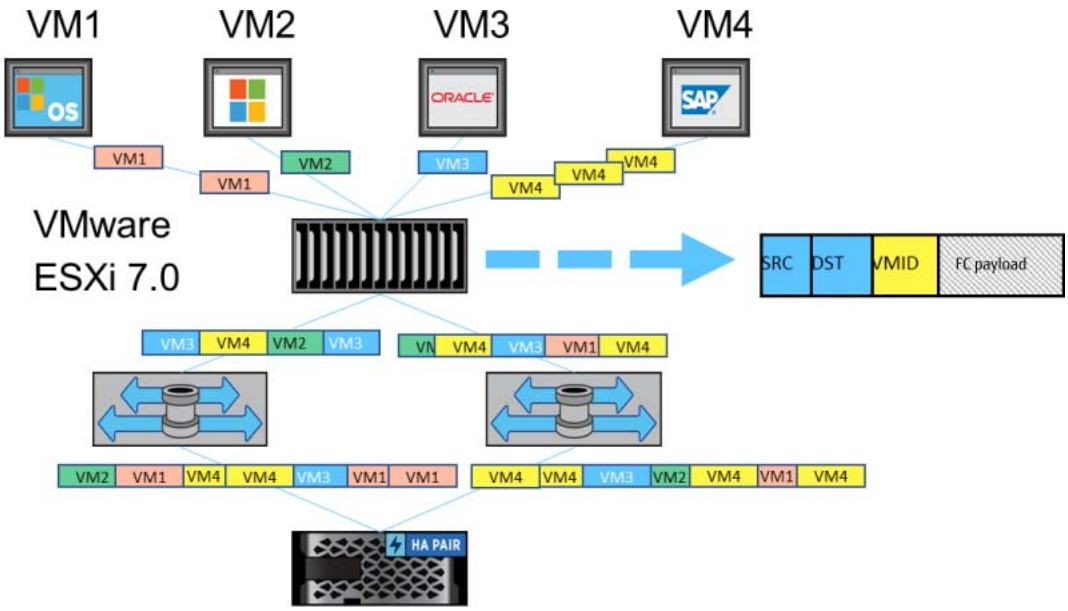

Tracking I/O from each VM in a shared datastore through an FC fabric includes these steps:

Procedure ▶▶▶

1. Hypervisor assigns a globally unique ID to each VM.
2. The VMID is tagged to each frame from the VM.
3. The switch and storage nodes propagate and reflect each frame and VMID.

◀◀◀

6. ONTAP 9.8: ASA-Specific New Features

ONTAP 9.8 introduced two new SAN features on the ASA only. These features will likely be added to unified ONTAP in an upcoming release. The ASA-only features are:

- An increase in maximum LUN size
- Persistent ports

Increase in Maximum LUN Size

ONTAP 9.8 increased the maximum LUN size from 16TB to 128TB. This increase also has a corresponding increase in size for volumes from 100TB to 300TB. The expectation is that the most popular use of those larger LUNs is going to be for LUNs backing hypervisor data stores.

Persistent Ports

Persistent ports reduce the impact of takeovers. They are able to do this by creating a shadow LIF on the corresponding physical port on the HA partner. When a node is taken over, the shadow LIF on the corresponding partner node assumes the identity of the original LIF, including the WWPN.

Persistent ports work because they are able to advertise the shadow LIF as active optimized to the host Multipath I/O (MPIO) stack before it changes the previous path status to down (faulty).

The host MPIO stack shifts I/O to the next active optimized path (formerly the shadow LIF) so that any I/O disruption is minimized. The host always sees the same number of paths to the target regardless of the target's state (steady state or in takeover).

Persistent ports are introduced with ONTAP 9.8 on ASA. When a node is upgraded to ONTAP 9.8, this feature is enabled by default.

Best practices for persistent ports require that FCP port characteristics must be identical within the HA pair:

- FCP port counts
- FCP port names
- FCP port speeds
- FCP LIF WWPN-based zoning
- Both the active and shadow LIF need to be in the same zone with the initiator

If any of these best practices are violated, an EMS message is generated with the following body:

```
EMS : scsiblade.lif.persistent.ports.fcp.init.error
```

The persistent ports feature is available with FC but not iSCSI. It requires that zone membership be identified by WWPN, since the WWPN is essentially spoofed on the shadow LIF also.

Figure 4 shows the persistent ports.

Figure 4 Persistent ports

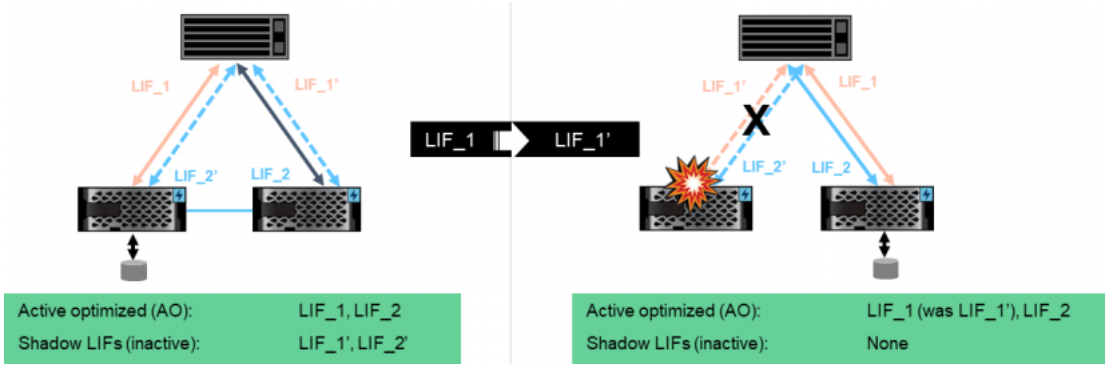

Figure 5 shows a persistent ports zoning example.

Figure 5 Persistent ports zoning example

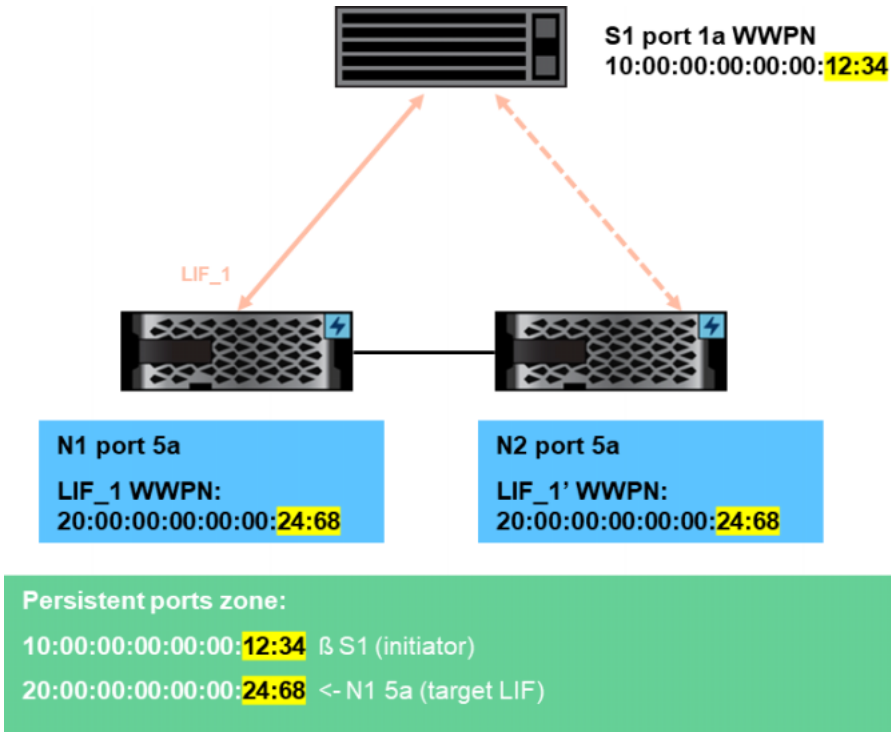

7. ONTAP 9.9.1: ASA New Features

ONTAP 9.9.1 adds two new ASA enhancements:

- SnapMirror Business Continuity (SM-BC)
- ASA can now be expanded to 12-node clusters
- NVMe/FC protocol support

SnapMirror Business Continuity

ONTAP 9.9.1 introduced SM-BC, which uses SnapMirror Synchronous to synchronously replicate applications using application consistency groups to manage and replicate all application objects between the two sites. SM-BC enables automated failovers between two synchronously replicated sites. This reduces any outage durations and significantly lowers administrative costs associated with maintaining

both mirrors and managing automated failovers. Figure 6 is a visualization of this topology. For more information, see ETERNUS AX/HX Series SnapMirror Business Continuity in the Fujitsu manual site.

Figure 6 SMBC

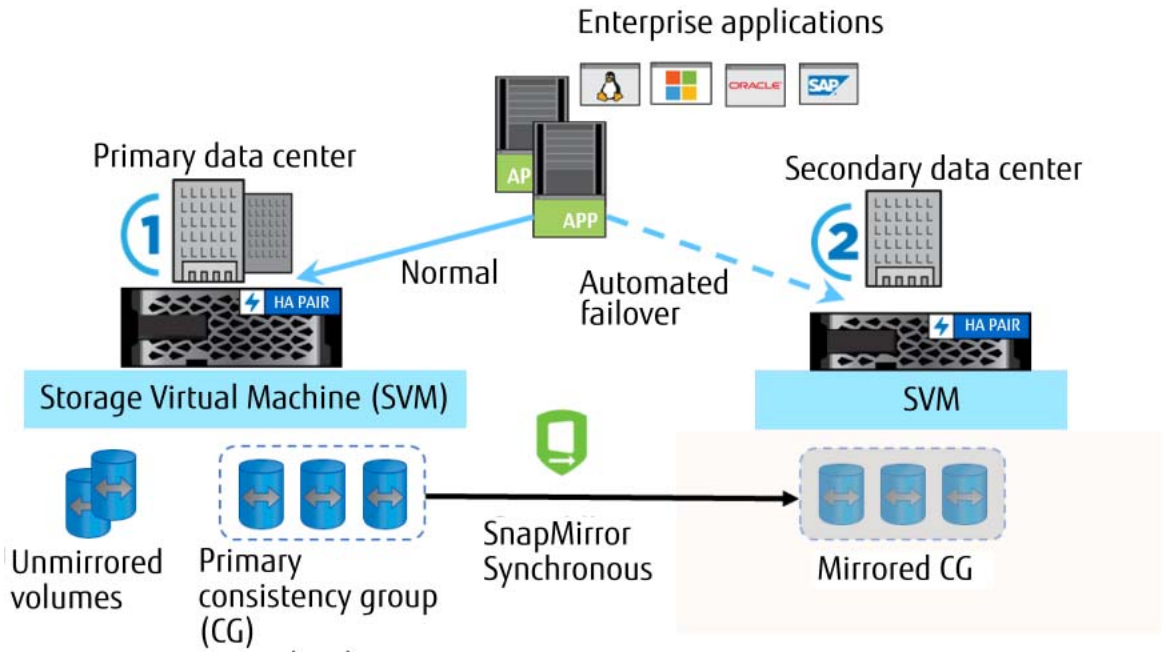

ASA Maximum Cluster Size Grows from a Single HA Pair to 12 Nodes, or Six HA Pairs

In configurations larger than a single HA pair, it is important to understand that the ASA cluster is essentially a federation of ASA HA pairs that are clustered together and share a single management plan. This allows all of the usual NDO operations and other features of ONTAP clusters, but with the exception that there are no NAS protocols or features supported on those nodes and certain other ASA specifics.
Because the primary feature that defines an ASA is that it offers symmetric active-active access, it is important to understand how that access and pathing works when ASAs are combined into a larger cluster. ASA within a single HA pair advertise all paths through both nodes as Active Optimized (AO or preferred paths) and therefore the host MPIO stack will use all AO paths. For ASA cluster that are larger than a single HA pair, each HA pair advertises all paths to LUNs hosted by that pair as AO. However, any paths through other controllers not part of the hosting HA pair would advertise paths as Active Non-Optimized (ANO or nonpreferred paths) these paths would not be used by host MPIO stacks. Additionally by default, ONTAP has a feature called Selective LUN Map (SLM) that only advertises paths through the hosting HA pair, therefore hosts would not be aware of other non-optimized or less preferred paths unless SLM was configured to advertise additional paths. Figure 7 illustrates a 12-node ASA.

For more information on SLM, see the Selective LUN Mapping section in ETERNUS AX/HX series Best Practices for ONTAP SAN Configurations in the Fujitsu manual site.

Figure 7 Scale ETERNUS AX series ASA up to 12-nodes in ONTAP 9.9.1

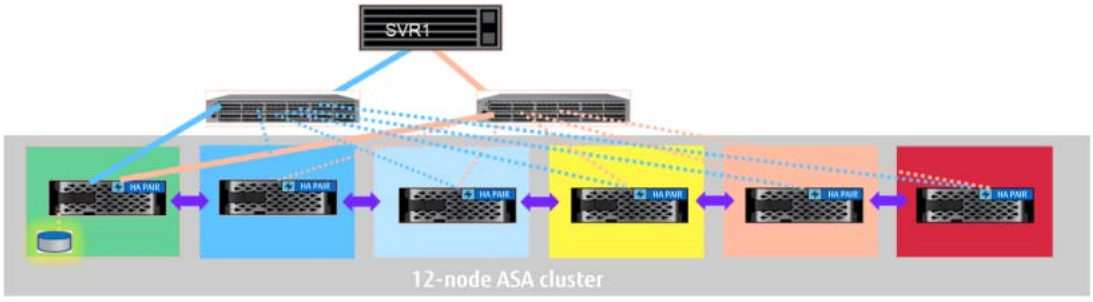

- ASA HA pairs are symmetric active-active within the HA pair (AO)
- Pathing is asymmetric between HA pairs (ANO)
- All nodes are managed as a single cluster
- All nondisruptive operations (NDO), nondisruptive upgrades (NDU), and vol/LUN move and copies are supported.

Legend
Blue Active optimized (AO) advertised path (SLM)
Peach Active optimized (AO) advertised path (SLM)
Blue Active nonoptimized (ANO) nonadvertised path (SLM)
Peach Active nonoptimized (ANO) nonadvertised path (SLM)

ASA Adds NVMe/FC Support

ONTAP 9.9.1 adds NVMe/FC as an additional block protocol. Unlike either FC or iSCSI, NVMe/FC on ASA will continue to be asymmetric (active optimized/active nonoptimized). This is due to differences in how NVMe-oF works with remote versus local paths.

8. ASA Configuration Requirements

This section details the requirements to implement and to maintain an ASA configuration. To validate an ASA configuration, you must fulfill the following requirements when you provision storage for applications.

You can check the configuration requirements and maximums by downloading and running Active IQ Config Advisor. To confirm that the storage system continues to conform to ASA configuration requirements, you should run Config Advisor after initial setup and provisioning and whenever you make significant changes to the configuration and workloads. To maintain consistent performance and to meet storage SLOs, if Active IQ Config Advisor discovers any inconsistencies with the baseline configuration, you must remediate them. Config Advisor queries the configuration and maximums and identifies any nonconforming items so that you can remediate them to maintain the rapid failover times that are critical for ASA performance.

Required Hardware and Software Components for ASA Configurations

All ASA configurations have the following mandatory components:

- As of this writing, the models that support ASA are as follows.
  - ETERNUS AX2100
  - ETERNUS AX2200
  - ETERNUS AX4100
- ONTAP 9.7 or later version
- Active IQ Unified Manager 9.7 or later version
- Active IQ Config Advisor

Tools for ASA

This section describes multiple tools that you can use with ASA to greatly ease ASA management. Use these tools for the following functions:

- Confirm supported configurations
- Gather, parse, and display customer storage estate configuration details
- Manage thresholds, reporting, alerts, and performance

ASA administrators should add all the tools described in this section to their tool chests. These tools can greatly ease configuration, administration, and management of ASAs and unified ONTAP platforms.
These tools are free to any Fujitsu customer or partner to use.

Active IQ OneCollect Data Collection Tool

Active IQ OneCollect is a data collection tool that gathers data from storage, hosts, and switches. The collected data is used for troubleshooting, solution validation, migration, and upgrade assessments. NetApp Active IQ OneCollect is available to NetApp customers, channel partners, and internal users.

You can use the OneCollect tool to gather all the necessary data about an existing configuration.

For more information about OneCollect, see the latest ETERNUS Software Active IQ OneCollect Installation and Setup Guide in the Fujitsu manual site.

Active IQ Unified Manager

Active IQ Unified Manager enables:

- Health monitoring
- Performance monitoring and analysis
- Utilization and usage reporting
- Thresholding and alerting

Active IQ Unified Manager, shown in Figure 8, provides complete ONTAP estate monitoring for all ONTAP clusters from a single pane. It's available as a Windows or Linux installation or as a VMware-based virtual appliance.

Figure 8 Active IQ Unified Manager

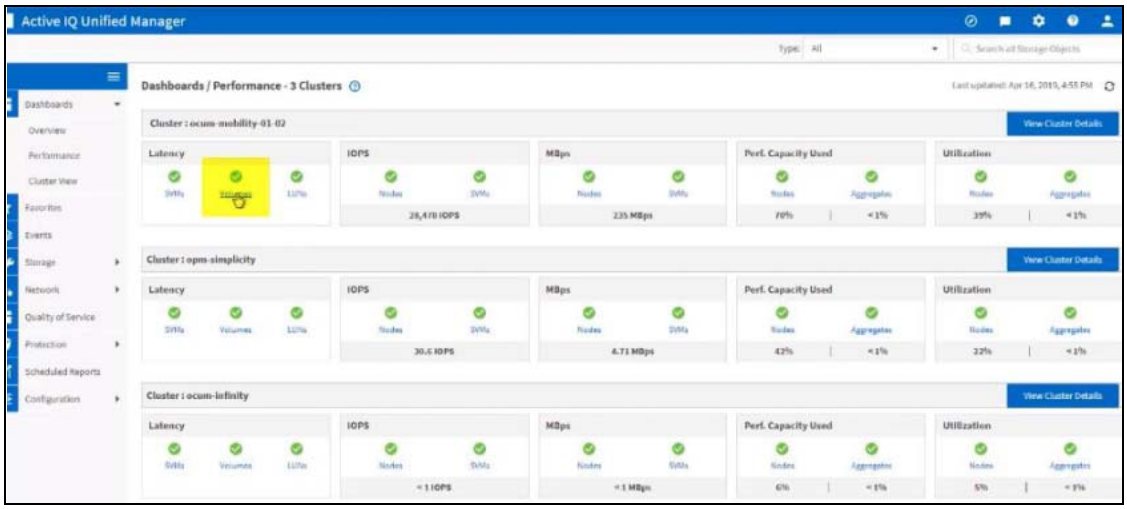

SAN Environmental Requirements

All ASA environments are assumed to have been architected to follow general SAN best practices: redundant fabrics and the use of dedicated high-speed storage networks that are segregated from general Ethernet communications networks. For details about best practices, see the ETERNUS AX/HX series Best Practices for ONTAP SAN Configurations in the Fujitsu manual site.

Hardware Configuration

ETERNUS AX series systems are introduced with ONTAP 9.7 as a single cluster that contains a single HA pair. The ASA of ONTAP 9.7/9.8 supports only single HA pairs.

Storage Controllers

Table 1 lists the Fujitsu storage controllers that support ASAs.

The baseline configuration was tested and qualified with a particular storage layout when running an ETERNUS AX series storage system. ETERNUS AX series nodes in a business-processing cluster must meet the storage subsystem hardware requirements that are described in Table 1.

Table 1 ETERNUS AX series cluster limits

Limits	ETERNUS AX series	Notes
Aggregate type	SSD only	-
Advanced Disk Partitioning (ADP)	Yes	Advanced Disk Partitioning 2 (ADP2): one root, two data partitions. Each disk has three partitions, with a data partition per controller, up to the first 48 disks. The remaining disks are partitioned normally.
Maximum storage devices/node	240	-
Data aggregates	1-10	-
Drives/RAID group	11-28	-
Maximum volumes/node	200	-
Maximum LUNs/node	12,288	-
ONTAP Snapshot copies/volume	40	-
Data aggregate space utilization	>75%	-
Controller utilization	50% performance capacity	For CPU/disk utilization: Use the Active IQ Unified Manager headroom tool. CPU utilization of <=50% applies to steady-state only. In takeover, CPU utilization can go higher than 50% due to load from the other node.

Steady-state Storage Utilization

Fujitsu recommends that you size ASA nodes to less than 50% of performance capacity per node. This recommendation helps prevent an impact on performance if a failover occurs, where one controller hosts both controllers' workloads. Fujitsu sizing tools are tuned to size ASA systems based on this recommendation. Although this recommendation does not allow both controllers to be optimized for steady-state operations, there is no performance variance when a failover occurs.

After the system is in operation, if workloads grow beyond the recommended maximum per node, Fujitsu suggests that ASA administrators balance these workloads back to below 50% per node. This rebalancing prevents performance impacts if a failover occurs. Neither ONTAP nor ASA specifically stops storage managers from provisioning beyond 50% performance capacity per node. The impact on takeover performance is correlated to the amount of performance capacity that is over 50%.

Fujitsu recommends the use of performance capacity to optimize performance while maintaining consistently low latency.

Software Configuration

The software configuration that's specific to a storage cluster running within the baseline configuration is meant to change over time as workloads and applications are added and removed. The software configuration section outlines the range of configuration values and settings that are included in the ASA configuration. To validate them automatically, you can use the Config Advisor tool. For more information about this tool and how to use it to validate a storage cluster's settings, see "Validate the ASA Configuration" (page 24).

Aggregate Thresholds

You can set a fullness threshold for aggregates so that when the total percentage of used space in the aggregate exceeds the threshold, an event is generated. This event can then be forwarded to an SNMP-based monitoring tool.

To increase warning times and reaction windows, you should set the ETERNUS AX series ASA controllers' nearly full threshold to 70% and their full threshold to 75%. By lowering both thresholds, storage administrators have ample opportunity to act well before an aggregate is completely filled, despite the smaller storage space that is commonly available when compared with storage controllers that use spinning media.

Host OS Configuration and Settings

Fujitsu publishes host utilities for the following host OS families:

- IBM AIX
- Microsoft Windows
- Linux
- Oracle Solaris

The host utilities software comes with:

- Documentation that's specific to the OS for which it's designed.
- Recommendations for configuration setting and tuning to optimize the OS for ONTAP SAN.
- The SANLUN utility, which provides several queries that are very helpful when documenting or troubleshooting host and ONTAP SAN interactions. These queries include listing paths, worldwide port name (WWPNs), iSCSI Qualified Names (IQNs), LUNs found, adapter settings, and so on.

> **Note**
>
> There are no differences in the host OS settings between hosts that connect to ONTAP unified controllers versus ASA settings.

ASA Specific Limits

To accelerate storage failover transition times, ASA configurations have lower maximum values for some parameters. Table 2 summarizes the differences between ETERNUS AX series systems and ASA systems at the time of their introduction. The limits of ASA will likely change over time because the workload and performance characteristics identify the maximum object limit that enables the ASA to minimize the failover transition time (takeover or giveback).

The virtually instantaneous transition time causes no impact because there are still active paths to all LUNs. I/Os are fenced while controllers are actively transitioning, then they are responded to after the storage transition is complete.

Table 2 ETERNUS AX series versus ASA maximums

Objects per node	ETERNUS AX series cluster maximums	ASA cluster maximums
Maximum volumes	1,000	200
Data Protection Optimized (DPO) volumes	1,000	Not applicable; Fujitsu does not recommend DPO volumes on ASA
LUNs: ETERNUS AX2100 ETERNUS AX2200 ETERNUS AX4100	 8,192 8,192 12,288	 8,192 8,192 12,288

Protocol Support

ASA supports block protocols exclusively and currently supports FC and iSCSI. In addition, NVMe over Fabrics (NVMe-oF) protocols are supported. Neither NAS protocols nor NAS-only features are supported on ASA.

Snapshot Scheduling and Policy

Although Snapshot copies are supported on ASA systems, in most cases, Fujitsu recommends that you disable Snapshot policies. There are two reasons to disable Snapshot copies:

- Snapshot copies should be managed by a storage management tool, for instance, the SnapCenter products, or should be application initiated to validate that they are application consistent.
- By disabling the Snapshot policy, your storage managers can also better manage the number of Snapshot copies and the amount of space that's consumed.

Use ONTAP System Manager to edit, delete, or disable Snapshot policies. See Figure 9.

Figure 9 Use ONTAP System Manager to manage Snapshot policies

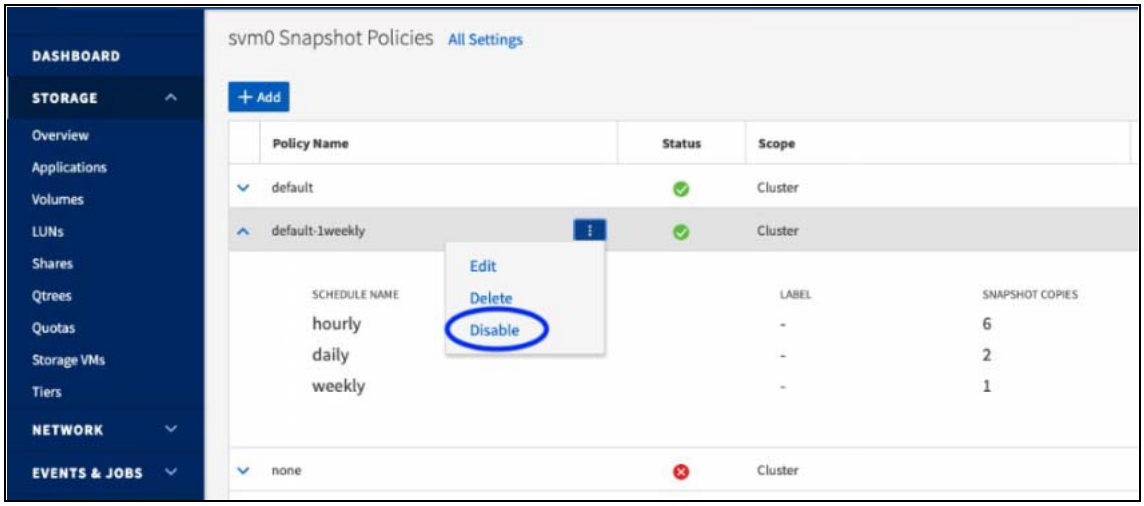

Thin Provisioning

ONTAP uses the WAFL file system, which does not preallocate storage on disk before consuming it. This storage allocation policy is known as thin provisioning or dynamic provisioning. You can set space reserves to subtract free space from a volume, an aggregate, or a LUN and to hold it in reserve for future write operations. This approach is called thick provisioning. When space reserves are turned off and LUNs are created that, when fully written, could consume more space than is immediately available in a volume or an aggregate, the policy is known as storage overcommitment.

Storage overcommitment requires that free space be continuously monitored to meet the needs of hosted applications. This policy also requires an action plan for increasing the free space that's available (either through non-disruptive data mobility operations or by expanding aggregate sizes). Therefore, the most conservative option is to fully provision storage, but at the cost of additional storage capacity that might not be required.

If you use thin provisioning, a strategy or action plan must be documented and in place to mitigate low-space scenarios. It is also a best practice to leave >25% free space in the hosting aggregate and to adjust free space thresholds for those aggregates. This recommendation is made to give storage managers enough time to react to low-space situations. For more information, see "Aggregate Thresholds" (page 21).

LUN Space Allocation

The `space allocation` option on LUNs is disabled by default; you should not enable it. The space allocation setting determines whether a LUN supports SCSI unmap/space reclamation.

Space Reclamation (T10 hole punching/unmap)

Space reclamation can be extremely processor intensive, potentially long-running, and can cause transient host-side performance impacts and latency spikes. It is therefore a best practice to leave unmap disabled. If any LUN that has this option enabled is replicated or migrated into the ASA, you should disable the option before allowing the LUN to be discovered by a host system. Not disabling this option could lead to potentially long-running performance impacts while unmapping scans are running on hosts and then communicated back to the ASA. This hole-punching is triggered by low-space or number-of-deleted-block thresholds that can be triggered during peak production times.

Note

When space reclamation was introduced, almost all storage was comprised of hard disk drives and most storage efficiency technologies like dedupe, zero block detection, compression, and compaction were still mostly in development. Because of those innovations in storage efficiency, space reclamation is unlikely to have the same benefits seen in a pre-storage efficiency environment.

Validate the ASA Configuration

You can validate the ASA configuration with the Active IQ Config Advisor tool, as shown in Figure 10. Config Advisor examines an ASA cluster's current configuration and compares it with the baseline configuration, as detailed in this document. Table 3 contains a list of the checks performed by Config Advisor. Fujitsu recommends that you keep the resulting list of warnings for archival purposes and use it as a list of items to be remediated (see Figure 11).

The resulting output details any areas where the storage cluster's current configuration differs from the baseline configuration. You should schedule remediation actions to reestablish compliance for any configuration details that do not conform to the baseline configuration.

Figure 10 Config Advisor with Managed ONTAP SAN plug-in

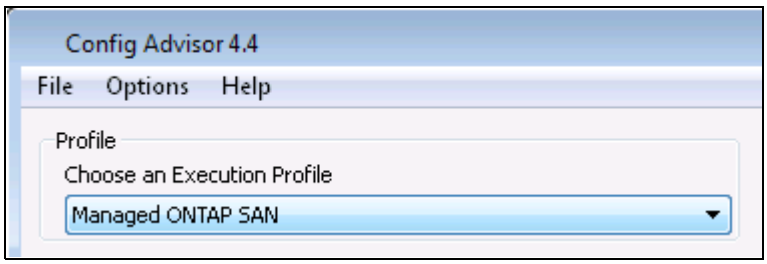

Figure 11 Configuration verification by Config Advisor with Managed ONTAP SAN plug-in

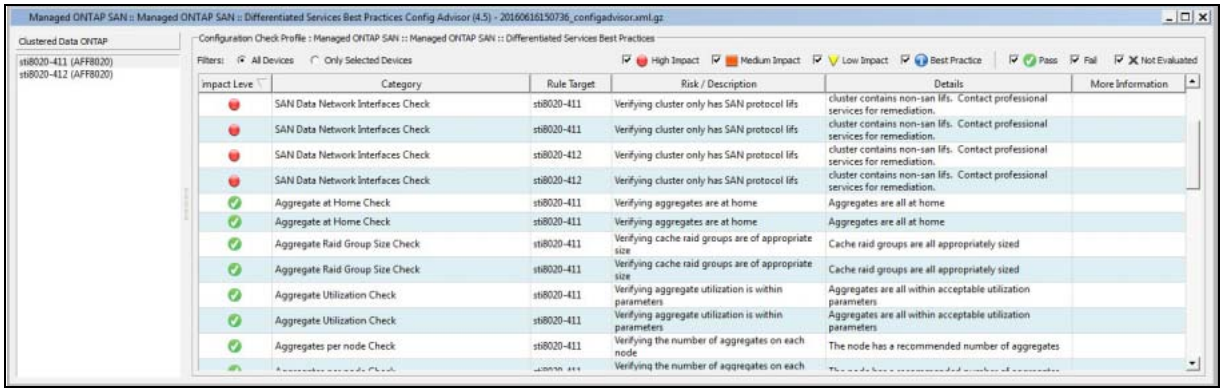

Table 3 Configuration checks performed by Active IQ Config Advisor

Check name	Description
Node health check	Verifies that nodes are healthy and can be queried for information
Model check	Verifies that all nodes are ASA supported controllers
Network interfaces check	Verifies that only SAN LIFs exist on the cluster
Aggregates per node check	Verifies that from 1 through 10 data aggregates are on each node
Aggregates at home check	Verifies that all aggregates are currently being serviced by their owning node
Aggregates utilization check	Verifies that no aggregates exceed 75% utilization
Volumes per node check	Verifies that no nodes own more than 200 volumes
Snapshot copies per volume check	Verifies that no volumes have more than 40 Snapshot copies
SFO check	Verifies that all nodes have SFO enabled
SAN SVM (formerly Vserver) quality of service (QoS) check	Verifies that QoS is enabled on all storage virtual machines (SVMs)
LUN space allocation check	Verifies that space allocation is disabled for all LUNs

9. Performance Capacity, CPU Utilization, Storage Utilization, and Performance Capacity Planning

To determine the optimal solution sizing, an initial sizing must be performed. After the initial sizing, Fujitsu recommends that you base all incremental performance sizing, monitoring, capacity planning, and workload placement on the Active IQ Unified Manager performance capacity determination.

Fujitsu best practice for sizing ASA systems is to use performance capacity to size each node to less than 50% of the performance capacity on each controller. By sizing this way, you can maintain acceptable low latency if a takeover occurs. The cost of this approach is that you sacrifice a little of the steady-state top-line performance.

10. ASA Service Offering Lifecycle

The ASA service lifecycle describes how to size, configure, and validate an ASA implementation. It includes a number of checklists and task lists that should be performed to put a new ASA into production.

Size an ASA Cluster

Most ASA clusters need to grow over time. The cluster's initial controller models, disks, and shelves must be determined. This determination can be made with OS and application vendor sizing tools, or with the deployment guide that's associated with the applications that the cluster hosts. For other sizing guides that are appropriate to particular applications, visit the Fujitsu manual site and see "Steady-state Storage Utilization" (page 20), in this document. Storage managers need to manage additional workload growth by rebalancing it across the ASA and even possibly between the current and additional ASAs.

Initial Setup and Prevalidation

Before you begin qualification and acceptance testing of a new ASA system, you should perform several steps after basic hardware installation of the cluster nodes. These steps are shown in the following checklists and validation guidelines.

Initial Hardware Setup Checklist

Install all the cluster nodes, including shelves, cluster network switches, and cabling, according to their installation guides. Table 4 shows the checklist items.

Table 4 Hardware setup checklist

	Checklist item
	All the ASA cluster's hardware components are operational.
	The nodes and network switches have no faults.
	All power supply units and system fans are operational.
	No shelf modules or SSDs display faults.
	The FCP and iSCSI licenses are enabled, as appropriate.
	The cluster's disks, cluster network, and HA failover cabling are correct and have been validated by the Config Advisor tool.

Initial Hardware Setup Validation

To validate the initial hardware setup checklist that is shown in Table 4, use the validation method from the corresponding checklist item in Table 5.

Table 5 Hardware checklist validation methods

	Checklist validation method
	Validate according to data center policies and guidelines: • Visually inspect cluster hardware for fault lights or other indicators. • Review storage controller environmental sensor readouts. • Review the cluster dashboard by using ONTAP System Manager.
	Review disk and shelf status values by using ONTAP System Manager. Under the cluster menu, review the overview and disks menu. The dashboard also has alerts for any problem components.
	Review Config Advisor output.
	Review the licenses that are currently installed by reviewing ONTAP System Manager, Cluster > Settings. The License tile displays licensed protocols and features; you can also enable any additional licenses that are supported on the ASA from that tile.

Configuration Tool Setup Checklist

For the list of configuration tools that are part of an ASA environment, see Table 6.

Table 6 ASA configuration tools

Configuration tool	Version	Schedule	Functionality
OneCollect	Latest	When configuration changes	Checks and preserves end-to-end configuration details
Config Advisor	Latest	When cluster configuration changes	Checks cabling and HA properties of storage systems

Predeployment Validation Tasks

Table 7 provides a checklist of predeployment validation tasks.

Table 7 Predeployment validation task checklist

	Prevalidation task	Desired result
	Identify the hosts, fabrics, and networks that connect to the ASA, including hosts used during validation phases and when the ASA is serving applications in a production role.	You have validated the hardware and software in your environment are supported in an ASA environment, including hosts, networks, and fabrics.
	Connect hosts to the ASA cluster by using the iSCSI, FC, or NVMe over Fabric Protocol.	LUNs provided by the ASA cluster that are suitable for testing are mounted on hosts in the ASA application environment.

For a description of SAN topologies and host setup details, see the ETERNUS AX/HX Series SAN Configuration Guide in the Fujitsu manual site.

Validation Testing

You must monitor and test your ASA configuration.

Table 8 Application validation test items

Number	Validation test	Desired result
1	Cable pull and port shutdown to cause path failure: • From the storage controller to the fabric or Ethernet switch • From the host to the fabric or Ethernet switch	Path faults are detected by Active IQ Unified Manager; storage volume performance is still within ASA parameters.
2	Planned takeover and giveback of storage controllers	Storage I/O is not disrupted; storage performance is unaffected; alerts are sent out by using Active IQ Unified Manager and AutoSupport.
3	Unplanned takeover and giveback of storage controllers	Storage I/O is not disrupted; storage performance is unaffected; alerts are sent out by using Active IQ Unified Manager and AutoSupport.

Manage and Schedule Operations that Help Increase System Utilization

There are several operations that a storage administrator can run that can increase CPU and disk utilization temporarily while the operations are being run.

Some of these operations include nondisruptive volume and LUN move operations, such as a volume move or LUN move, large Snapshot deletes, and SnapMirror initializations or rebaselines. As commonsense guidance, Fujitsu recommends that, where possible, you schedule these operations during nonpeak or lower-utilization periods.

Fujitsu also recommends that you reduce the number of concurrent operations that you run. For example, don't move 20 volumes at a time; such operations will reduce performance. By following these guidelines, you can achieve higher performance. In addition, operations such as volume moves complete more rapidly, which has the added benefit of reducing the amount of time that your controllers are subject to the utilization costs of these types of operations.

A. Configure Active Directory Domain Controller Access

Before an Active Directory account can access the SVM, you must configure Active Directory domain controller access to the cluster or SVM. Because a CIFS volume is not present on the ASA, you can create a computer account for the SVM on the Active Directory domain.

You have two options for configuring Active Directory domain controller authentication:

- **Configure an authentication tunnel**
  If you have already configured a CIFS server for a data SVM, you can use the `security login domain-tunnel create` command to configure the SVM as a gateway, or tunnel, for Active Directory access to the cluster.
- **Create an SVM computer account on the domain**
  If you have not configured a CIFS server for a data SVM, you can use the `vserver active-directory create` command to create a computer account for the SVM on the domain.

FUJITSU Storage
ETERNUS AX series All-Flash Arrays
Best Practices for Business-Critical Workloads

P3AG-6342-01ENZ0

Date of issuance: December 2021
Issuance responsibility: FUJITSU LIMITED


FUJITSU