P3AG-5552-25ENZ0

Fujitsu Storage
ETERNUS AX series All-Flash Arrays,
ETERNUS AC series All-Flash Arrays,
ETERNUS HX series Hybrid Arrays

Notes and Cautions When Using This Product

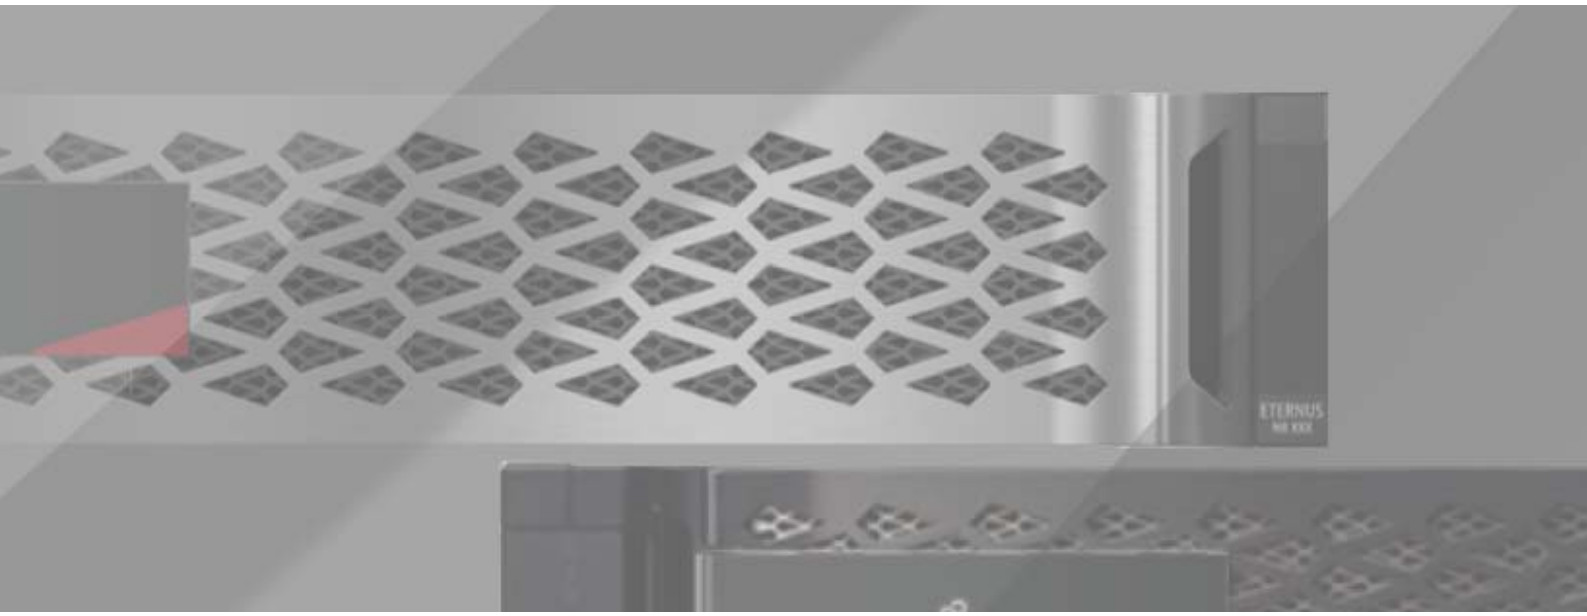

FUJITSU

Table of Contents

List of Figures

List of Tables

Preface

The Fujitsu Storage ETERNUS AX1100/AX1200, ETERNUS AX2100/AX2200, ETERNUS AX4100 All-Flash Arrays, the ETERNUS AC2100 All-Flash Arrays, and the ETERNUS HX2100/HX2200/HX2300, ETERNUS HX6100 Hybrid Arrays (hereinafter referred to as ETERNUS AX/AC/HX) are designed to be connected to Fujitsu servers (Fujitsu SPARC Servers, PRIMEQUEST, PRIMERGY, and other servers) or non-Fujitsu servers.

This manual provides notes and cautions for customers who use this product.

This manual is intended for use of the ETERNUS AX/AC/HX in regions other than Japan.

Twenty-fifth Edition
May 2025

Trademarks

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

About This Manual

Intended Audience

This manual is intended for customers, sales representatives, and field engineers.

Related Information and Documents

The latest information for your model 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/AC/HX. Make sure to read the information.
Note	Indicates information and suggestions that supplement the descriptions included in this manual.

Model Name Notation

In this manual, model names may be noted as follows.

Target model	Notation
ETERNUS AX1100/AX1200	ETERNUS AX1x00
ETERNUS AX2100/AX2200	ETERNUS AX2x00
ETERNUS HX2100/HX2200/HX2300	ETERNUS HX2x00

1. Product Organization

This chapter describes the base unit and the optional products that are configured in the ETERNUS AX/AC/HX.
The base units are categorized into All-Flash Arrays and Hybrid Arrays. These categories include the following products.

Table 1 List of ETERNUS AX/AC/HX Products

Category	Class	
	Entry model (lower model)	Midrange model (upper model)
All-Flash Arrays	ETERNUS AX1100 ETERNUS AX1200 ETERNUS AX2100 ETERNUS AX2200	ETERNUS AX4100
	ETERNUS AC2100	-
Hybrid Arrays	ETERNUS HX2100 ETERNUS HX2200 ETERNUS HX2300	ETERNUS HX6100

ETERNUS AX series (All-Flash Arrays)

- ETERNUS AX1100/AX1200/AX2100/AX2200, ETERNUS AX4100
  Refer to this website for the specifications.

Figure 1 ETERNUS AX Series Product Lineup

ETERNUS AC series (All-Flash Arrays)

Refer to this website for the specifications.

Figure 2 ETERNUS AC Series

ETERNUS HX series (Hybrid Arrays)

- ETERNUS HX2100/HX2200/HX2300, ETERNUS HX6100
  Refer to this website for the specifications.

Figure 3 ETERNUS HX Series Product Lineup

2. Notes and Cautions When Using This Product

This chapter describes notes and cautions when using this product.

Important Note

Notes on Netlogon Vulnerability (CVE-2022-38023)

The versions that have implemented the security update for Netlogon vulnerability (CVE-2022-38023) are as follows.

- ONTAP 9.7 P22
- ONTAP 9.8 P18
- ONTAP 9.9.1 P15
- ONTAP 9.10.1 P12
- ONTAP 9.11.1 P8
- ONTAP 9.12.1 P2 (*1)

*1: For ONTAP 9.12.1, upgrade to ONTAP 9.12.1 P4 or later.

Remote Support

Remote support services for notifying the errors that occur in the ETERNUS AX/AC/HX.
Contact your maintenance engineer for Remote support details.

REMCS

Remote support has the following maintenance functions.

- Failure notice
  This function reports various failures, that occur in the ETERNUS AX/AC/HX, to the remote support center.
  The maintenance engineer is notified of a failure immediately.
- Information transfer
  This function sends information such as logs and configuration information to be used when checking a failure. This shortens the time to collect the information that is necessary to check errors.

Events notified to the remote support center are Error events and Warning events mostly.

Notes Related to the Network Settings

About the Management LIF Setting Port

When collecting logs during maintenance work, confirm that e0M is the Management LIF port and collect it from e0M. However, if a port other than e0M is used, cooperation from the customer may be required.

About the Segment of SP/BMC (e0M) and Data LIF

If the segment of SP/BMC (e0M) and Data LIF are the same, the data that should flow in Data LIF flows in SP/BMC (e0M) instead and may become a performance bottleneck. If the operation environment can be operated in a segment that differs from the IP address of SP/BMC (e0M), set Data LIF with a segment IP address that is different from SP/BMC (e0M).

In addition, if the data network segment has a high load in the ETERNUS AX/AC/HX and the data network and e0M are in the same segment, the data that should flow in the data network instead flows through the e0M and causes SP/BMC to malfunction, which in turn may cause high-speed rotation of fans, detection of sensor malfunctions, and finally, a shutdown (takeover). If a shutdown occurs simultaneously on both nodes, the system may fail. To avoid this, separate the data network segment from the SP/BMC (e0M) segment.

About the IP Address Setting of SP/BMC (e0M)

For the NIC of SP/BMC (e0M), do not set the network address within the range of "10.1.1.0/24". If it was set within that range, SP/BMC may continually reboot, the fans may rotate at high-speed, and the volume may become loud.

About the Ports That Are Set as Cluster LIFs

In a cluster network, the controllers are connected to each other using the cluster LIFs, and the ports used for the cluster LIFs are defined for each model. If a port used for the cluster LIFs is changed to a port that is not defined in advance, the redundancy or the performance may degrade. Do not change the ports used for the cluster LIFs to ports that are not defined in advance.

Table 2 Ports Defined as Cluster LIFs in Advance

Model	Port defined in advance
ETERNUS AX1100/AX1200/AX2100, ETERNUS HX2100/HX2200/HX2300	e0a and e0b
ETERNUS AX2200, ETERNUS AC2100	e0c and e0d
ETERNUS AX4100	e3a and e3b
ETERNUS HX6100	e0c and e0d

Notes When Connecting Tape Devices

Because the tape alias setting uses the serial number of the drive/changer, if the drive/changer is replaced, the tape alias must be set again on the ETERNUS AX/AC/HX.

Notes and Cautions for SSDs

Notes on Using SSDs

SSDs have reserved spaces to handle stress on SSDs, which increases as the number of write operations on NAND blocks increases. These reserved spaces are also used for garbage collection to aggregate discrete data. The reserved space may decrease depending on the amount of write operations, so a decline in write performance and a delay in response may occur.
The ETERNUS AX/AC/HX constantly monitors the components life and performance, and performs preventive maintenance (Redundant Copy) before the SSDs affect operations.

Cautions When Storing SSDs for an Extended Period of Time

If SSDs installed in the ETERNUS AX/AC/HX are left without power for 12 weeks or longer, they may lose data or become unusable. This is due to the characteristics of the NAND storage media.

To store SSDs without power for an extended period of time, perform the following actions.

- Back up the data if the SSDs are to be stored for at least 14 days.
  Backing up the current data to another area is recommended just before turning off the ETERNUS AX/AC/HX.
- Turn on the ETERNUS AX/AC/HX and provide the SSDs with power for at least 10 days before 10 weeks pass since the SSDs had no power.
- Delete all data before storing the SSDs for at least one month.

To store SSDs for an extended period of time and then reuse them without using their data, delete the RAID groups that are configured with those SSDs and then delete the data before cutting the power.
To reuse the SSDs that have been stored without power for an extended period of time, make sure to turn on the power when there is enough time to restore the backup data or to replace parts if the SSDs are unusable.

Notes and Cautions When Starting the ETERNUS AX/AC/HX

When starting the system, turn the power on in order from the drive shelf and then the controller.
Turn on the drive shelf and then after the Shelf LED turns green, turn on the controller.
If the above order is not followed, the drives may not be properly recognized.

Starting the ETERNUS AX/AC/HX

Cautions When Configuring an External Security Key Management Server

If security key management is configured using an external security key management server, start the server before the ETERNUS AX/AC/HX. The ETERNUS AX/AC/HX cannot be started if the security key management server is not started.

Procedure ▶▶▶

1 Make sure that the breaker of the power distribution units connected to the ETERNUS AX/AC/HX is in the OFF position and that the AC switch of the Power Fan Canister/Power Canister (hereinafter referred to as "power supply unit") of each shelf is in the OFF position.

Note

An AC switch is not installed on the power supply units of the ETERNUS AX2200, ETERNUS AX4100, ETERNUS AC2100, or ETERNUS HX6100 controller shelf or NS224 drive shelf.
When using an ETERNUS AX4100 or ETERNUS HX6100, make sure that the four power cords are disconnected. When using an ETERNUS AX2200, ETERNUS AC2100, or NS224 drive shelf, make sure that the two power cords are disconnected.

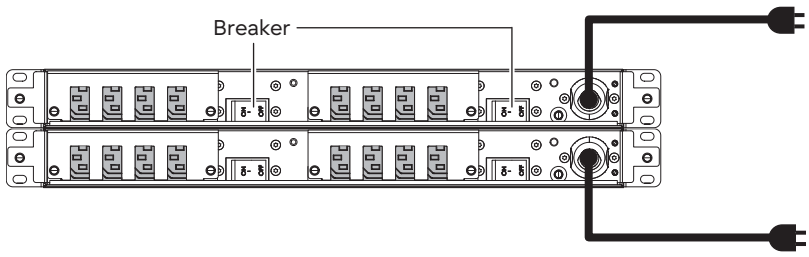

2 Turn on the AC power according to the following procedure.

- Cluster Switch

(1) Turn the breaker of the power distribution unit connected to the cluster switch to the ON position (marked "|").
If the power distribution unit option is not used, connect the power plug to an outlet (an AC switch is not installed).

(2) Make sure that the status LED is green.

■ Drive Shelf

(1) Turn the AC switch of all the power supply units to the ON position (marked "|").
An AC switch is not installed on the power supply units of the NS224 drive shelves. Connect the power supply units with two power cords.

(2) Turn the breaker of the power supply units connected to the drive shelf to the ON position (marked "|").
This step is unnecessary if the power distribution unit option is not used. The drive shelf is turned on in Step (1).

(3) Wait approximately two minutes.

(4) Make sure that the Shelf Power LED on the front left of the drive shelf is green.

(5) Make sure that the Activity LED of the installed drives is green.

■ Controller Shelf

(1) Turn the AC switch of all the power supply units to the ON position (marked "|").
An AC switch is not installed on the ETERNUS AX4100 or ETERNUS HX6100 controller shelf. Connect the power supply units with four power cords.

(2) Turn the breaker of the power supply units connected to the controller shelf to the ON position (marked "|").
This step is unnecessary if the power distribution unit option is not used. The controller shelf is turned on in Step (1).

(3) Wait approximately three minutes.

(4) Make sure that the Power LED of the power supply units is green.

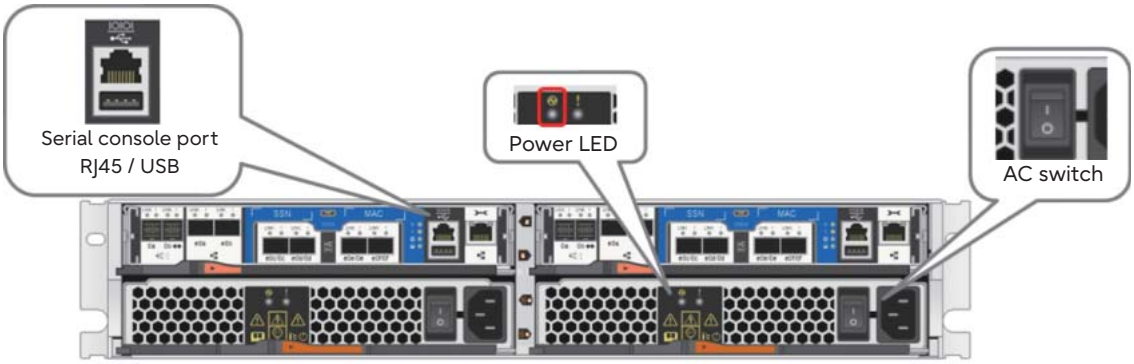

3 Make sure that the Shelf Power LED is green and the Fault LED is off on the front of the controller shelf.

4 Connect the serial console cable to the serial console port of the controller and then start a terminal emulator (such as TeraTerm) on the administration PC or server.

Set the terminal emulator with the following settings:

Baud rate:	115200 bps
Data length:	8 bits
Stop bit:	1 bit
Parity:	None
Flow control:	OFF
Buffer size (HyperTerminal):	Maximum

5 Check the cluster system.

5-1 On the console, press the Enter key.
A "login:" prompt appears.

5-2 Enter the login name and then the password.
The prompt "cluster1::>" appears.
To connect to the console, the username and password that have been set are required.
In the following example, the "cluster1::>" string shows the name that has been set for the cluster.

```
login: admin                                          ← Enter the login name
Password: password0                                   ← Enter the password
****************************************************
* This is a serial console session. Output from this *
* session is mirrored on the SP/RLM console session. *
****************************************************
cluster1::>
```

5-3 Execute the "cluster ha show" command.
Make sure that High Availability Configured is enabled (true).

```
cluster1::> cluster ha show
High Availability Configured: true
```

5-4 Check the cluster configuration.

(1) If the cluster consists of two nodes, execute the "cluster show" command.
Make sure that Health and Eligibility are enabled (true).

```
cluster1::> cluster show
Node                  Health  Eligibility
--------------------- ------- ------------
cluster1-01           true    true
cluster1-02           true    true
2 entries were displayed.
cluster1::>
```

(2) If the cluster consists of four or more nodes, execute the "storage failover show" command.
Make sure that Takeover Possible is enabled (true).

```
cluster1::> storage failover show
                                 Takeover
Node             Partner         Possible   State Description
---------------  --------------  --------   ------------------------------------
cluster1-01      cluster1-02     true       Connected to cluster1-02
cluster1-02      cluster1-01     true       Connected to cluster1-01
cluster1-03      cluster1-04     true       Connected to cluster1-04
cluster1-04      cluster1-03     true       Connected to cluster1-03
4 entries were displayed.
cluster1::>
```

Checking the LIF Status after Rebooting the ETERNUS AX/AC/HX

After the ETERNUS AX/AC/HX is rebooted, LIFs my be migrated to failover ports. Confirm that the LIFs are on the home ports.

Input

```
network interface show
```

Example

```
cluster1::> network interface show -vserver vs0
            Logical     Status      Network             Current       Current  Is
Vserver     Interface   Admin/Oper  Address/Mask        Node          Port     Home
----------- ----------  ----------  ------------------  ------------  -------  ----
vs0
            data001     down/down   192.0.2.120/24      node0         e0e      true
            data002     down/down   192.0.2.121/24      node0         e0f      true
            data003     down/down   192.0.2.122/24      node0         e2a      true
            data004     down/down   192.0.2.123/24      node0         e2b      true
            data005     down/down   192.0.2.124/24      node0         e0e      false
            data006     down/down   192.0.2.125/24      node0         e0f      false
            data007     down/down   192.0.2.126/24      node0         e2a      false
            data008     down/down   192.0.2.127/24      node0         e2b      false
8 entries were displayed.
```

If there is a LIF whose "Status Admin" is "down" or whose "Is Home" is "false", perform the following procedure.

Procedure ▶▶▶

1 Enable data LIFs.

Input

```
network interface modify {-role data} -status-admin up
```

Example

```
cluster1::> network interface modify {-role data} -status-admin up
8 entries were modified.
```

2 Revert the LIFs to their home ports.

Input

```
network interface revert *
```

Example

```
cluster1::> network interface revert *
8 entries were acted on.
```

3 Confirm that the LIFs are on their home ports.

Input

```
network interface show
```

Example

```
cluster1::> network interface show -vserver vs0
            Logical     Status      Network              Current        Current  Is
Vserver     Interface   Admin/Oper  Address/Mask         Node           Port     Home
----------- ----------  ----------  -------------------  -------------  -------  ----
vs0
            data001     up/up       192.0.2.120/24       node0          e0e      true
            data002     up/up       192.0.2.121/24       node0          e0f      true
            data003     up/up       192.0.2.122/24       node0          e2a      true
            data004     up/up       192.0.2.123/24       node0          e2b      true
            data005     up/up       192.0.2.124/24       node1          e0e      true
            data006     up/up       192.0.2.125/24       node1          e0f      true
            data007     up/up       192.0.2.126/24       node1          e2a      true
            data008     up/up       192.0.2.127/24       node1          e2b      true
8 entries were displayed.
```

◀◀◀

Stopping the ETERNUS AX/AC/HX

Cautions When Configuring an External Security Key Management Server

If security key management is configured using an external security key management server, stop the server after the ETERNUS AX/AC/HX has stopped.

Cautions When Performing Maintenance of the ETERNUS AX/AC/HX during a Power Outage

Before turning off the ETERNUS AX/AC/HX, execute the command to stop the dedicated OS of the ETERNUS AX/AC/HX ("halt" command).
Failure to do this operation causes unwritten data to remain in the cache memory and then causes the battery backup function to operate, thereby significantly shortening the battery life. Be sure to stop the ETERNUS AX/AC/HX before performing maintenance of the ETERNUS AX/AC/HX during a power outage.

Procedure ▶▶▶

1 Connect the serial console cable to the serial console port of controller A and then start a terminal emulator (such as TeraTerm) on the administration PC or server.

2 On the console, press the Enter key.

A "login:" prompt appears.

3 Enter the login name and then the password.

The prompt "cluster1::>" appears.
To connect to the console, the username and password that have been set are required.
In the following example, the "cluster1::>" string shows the name that has been set for the cluster.

```
login: admin                                          ← Enter the login name
Password: password0                                   ← Enter the password
****************************************************
* This is a serial console session. Output from this *
* session is mirrored on the SP/RLM console session. *
****************************************************
cluster1::>
```

4 Stop the ONTAP OS.

4-1 If the cluster consists of two nodes, perform the following procedure.

(1) Execute the following command:
system halt -node * -inhibit-takeover true -skip-lif-migration-before-shutdown true

```
cluster1::> system halt -node * -inhibit-takeover true -skip-lif-migration-
before-shutdown true
```

4-2 If the cluster consists of four nodes, perform the following procedure.

(1) Execute the following command:
system halt -node * -inhibit-takeover true -ignore-quorum-warnings true -skip-lif-migration-before-shutdown true

(2) When the 'Warning: Are you sure you want to halt node "cluster1-0x"? {y|n}:' prompt appears, enter "y".

(3) If the ETERNUS AX/AC/HX has multiple controllers, this prompt appears for each controller. Enter "y" each time the prompt appears.

(4) Make sure that "LOADER-A>" appears.

```
cluster1::> system halt -node * -inhibit-takeover true -ignore-quorum-warn-
ings true -skip-lif-migration-before-shutdown true

Warning: Are you sure you want to halt node "cluster1-01"? {y|n}: y

Warning: Are you sure you want to halt node "cluster1-02"? {y|n}: y
2 entries were acted on.
                                    (snip)

CPU Type: Intel(R) Xeon(R) CPU      E5240 @ 3.00GHz
LOADER-A>                                   ← Make sure that "LOADER-A>" appears
```

5 Connect the serial console cable from the serial console ports of controller B and subsequent controllers to the administration PC or server.

6 On the console screen, make sure that "LOADER-A>" or "LOADER-B>" appears.

7 Turn the AC switch of the power supply units of the controller shelf to the OFF position.

An AC switch is not installed on the power supply units of the ETERNUS AX4100 or ETERNUS HX6100 controller shelf or NS224 drive shelf. Disconnect the power cords of the shelf.

8 Wait approximately two minutes.

9 Make sure that the LED on the rear of the controller shelf is off.

10 Turn the AC switch of the power supply units of the drive shelf to the OFF position.

11 Wait approximately two minutes.

12 Make sure that the LED on the rear of the drive shelf is off.

13 Turn the breaker of the controller shelf, the drive shelf, and the power distribution unit that is connected to the cluster switch to the OFF position.

Notes When Performing Field Expansion

Notes When Placing an Order

Placing an order for a field expansion drive or license requires the serial numbers of the controllers. The serial number can be found on the System Manager administration screen, by using the command line, or on the label attached on the rear of the controller shelf. Refer to Online Help for details on how to check the serial number.
To perform a field expansion, provide the serial numbers to your Fujitsu or partner sales representative.

- Example of finding the serial number on the label
  On the rear of the controller shelf, the label is attached to the controller area. The label has a 12-digit number described after the string "SSN", which indicates a serial number. A unique serial number is assigned to each of the two controllers and both serial numbers are required.
- Example of finding the serial number using the command line
  Enter "run local sysconfig" on the system management console. The "System Serial Number:" string in the return value is followed by a 12-digit number, which indicates the serial number. A unique serial number is assigned to each of the two controllers and both serial numbers are required.

Notes When Adding a Drive

Be sure to apply the Disk Qualification Package (DQP) before adding a new capacity drive.

Notes Related to Failures

Drive Failure

Within the ETERNUS AX/AC/HX, installation of hot spare drives is mandatory and installation of two or more hot spare drives is recommended.

If a drive failure occurs or a spare drive has been prepared, make sure that the reconstruction process completed successfully. Usually, the amber LED that indicates an error turns on when a drive failure occurs. If there are no spare drives, RAID degradation continues. If another drive (*1) in the same RAID group fails in this state, the data in the volumes that belong to the relevant RAID group will be completely lost. Therefore, use a spare drive that is in the normal state to immediately replace the failed drive.
For spare drives, make sure to prepare ones with the same capacity as the data drives. If there are no drives with the same capacity as the failed drives, drives with a larger capacity than the failed drives are used for reconstruction. However, the larger drive is logically recognized as having the same capacity as the failed drive. This drive is treated as having the same capacity as the failed drive even after the failed drive is replaced and, as a result, spare drives with large capacity will become insufficient. In this case, use the Rapid RAID Recovery function to copyback the data to the replaced drive, restore the capacity of the large drive, and revert it back to a spare drive with a large capacity.

*1: For RAID-DP, two more drives have failed. For RAID-TEC, three more drives have failed.

Note

Creating an ADP with a minimum drive configuration (*2) assigns four drives (*3) to each node's root aggregate. This drive assignment does not allow for a hot spare drive due to a limited number of drives.
Without a hot spare drive, the fault tolerance of the RAID drives is reduced, potentially triggering a warning from ASUP.

*2: For the ETERNUS AX/AC, this configuration includes eight SSDs. For the ETERNUS HX, it includes four SSDs and eight HDDs.

*3: The drive type used for the ETERNUS AX/AC is SSD and HDD for the ETERNUS HX.

Handling Failed Drives

Remove the failed drives from the drive shelf the next time the system stopped or when hot swapping is performed. If the system is rebooted without removing the failed drive, the failed drive may be recognized as a normal drive during the initial diagnosis while the system is booting up and the system may configure the failed drive as a hot spare drive.
To prevent this, use the "disk fail" command to write "BROKEN" to the drive label of the failed drive. Drives labeled as "BROKEN" are not reused.

RAID Degradation

RAID groups become degraded when the parity disks fail. The file access service is managed normally and continues even in the degraded mode.
If degraded mode cannot be recovered for the ETERNUS AX/AC/HX without spare drives or for the ETERNUS AX/AC/HX where all the spare drives are used for the data recovery process, a message urging the user to replace the drive is repeated every hour and the system automatically shuts down after a certain period of time to protect data.

Notes When Replacing the Controllers

If a failure occurs, the controllers may need replacing.

- When controllers are replaced, the MAC address of the LAN port assigned to the controller is changed. Ask your system administrator to change the settings.
- When controllers are replaced, the World Wide Node Name (WWNN) and the World Wide Port Name (WWPN) of the host interface installed in the controller are changed. Ask your system administrator to change the zoning settings and the security settings.
- After the controllers are replaced, the license code must be changed at a later date.
  Because the license code is managed according to the system serial number, the license code must be changed if the controller is replaced.
  The license code associated with the new system serial number is provided within 30 days after the replacement. The license code is changed by our maintenance engineer at a later date.
- If all the controllers used to configure a cluster are replaced (the license correspondence is lost) before the license code is changed to a new one, the license configuration cannot be changed.
  - Example
    For CIFS licenses, Share cannot be changed.

Cautions When Using the Drive Shelf

Cautions about Copyback Operations When the Drive Shelf Is Connected

When a drive installed in the controller shelf is replaced by your maintenance engineer due to a failure, the new drive may be automatically assigned to a partner node by the copyback function.
To disable the copyback function, execute the following commands.

- Command to check the setting value

```
cluster1::>options raid.lm.disk.copyback.enable
```

- Command to disable the function

```
cluster1::>options raid.lm.disk.copyback.enable off
```

Cautions When Deploying a Spare Drive in the Drive Shelf

Even if a spare drive in the drive shelf is initialized and then manually partitioned during the setup, ONTAP may automatically release the partition during operation. For example, in a configuration where a spare drive installed on the drive shelf is partitioned and each partition is assigned to a different node, ONTAP may automatically release the partitions when a takeover occurs.
To allow automatic partitioning when this event occurs and a spare drive is required, enable the AutoAssign function.

Notes Related to MetroCluster IP

Supported Model Combinations

MetroCluster IP is supported by an ETERNUS AX2100 (excluding the ASA model) and ETERNUS HX2200 running ONTAP 9.7 or by all models (excluding the ETERNUS AX1100 and the ETERNUS HX2100) running ONTAP 9.8 or later. Note that the models used for the cluster that are to be installed must be the same model.
The supported combinations are noted in Table 3. Unlisted combinations are not supported.

Table 3 List of Supported Combinations for MetroCluster IP

Site A	Site B	Supported
AX1200	AX1200	Yes
AX2100	AX2100	Yes
AX2200	AX2200	Yes
AX4100	AX4100	Yes
AC2100	AC2100	Yes
HX2200	HX2200	Yes
HX6100	HX6100	Yes

Notes When Using an Adapter Card for MetroCluster IP

If an adapter card (HXBCPA or AXDCPA) for MetroCluster IP is used and its speed is changed from 40GbE to 100GbE, the node must be restarted.

Notes on Designing FabricPool Configurations

FabricPool relocates data blocks to the cloud tier according to the specified policy, and you cannot use only cloud tier data blocks to perform a recovery. Design a FabricPool configuration that includes a backup operation.
Note that BlueXP only supports the parts related to CVO/FabricPool.

Caution

For the tiering data in Windows environments that is accessed via SAN protocols, Fujitsu recommends leveraging private cloud services, such as ONTAP S3 and StorageGRID, due to connectivity considerations.

Notes on Subscription Models

If a subscription model contract is signed, some information must be registered in the SPF portal after an aggregate is created. After signing the subscription model contract, provide the information on the pool, storage system IP address, and login account to the SPF portal staff in charge of the onsite setup.

Notes on the Manageability Software Development Kit (SDK)

The Manageability Software Development Kit (SDK) is used to collect information when investigating the failures of various software.
Although this SDK is also a development kit required to create applications that monitor and manage ONTAP storage systems, Fujitsu shall not support the creation of applications using this SDK nor support applications created using this SDK.
Do not use this SDK for purposes other than to collect information when investigating software failures.

Notes on Monitoring the Storage System

For storage system monitoring performed through a network such as with SNMP traps, transmission of a failure notification may not be possible due to a network port failure. Additional monitoring such as checking for network traffic or checking the network status using System Manager is recommended.
Be especially mindful of this for operations where only one node of an HA is assumed to be running operations and the partner node is assumed to be a standby for a switchover in case of a failure. If a

failure on the partner node is left unnoticed, a switchover to the partner node may not be possible which can lead to a service shutdown.

NVMe/TCP (Supported in ONTAP 9.10.1 and Later)

To use NVMe/TCP in a SAN environment of the ETERNUS AX1100/AX2100 or ETERNUS HX series running ONTAP 9.10.1 or later, contact your Fujitsu or partner sales representative.

Cautions about CIFS

Disconnected Sessions after a CIFS Session Timeout

CIFS sessions are disconnected after a CIFS session timeout because of an idle timeout that occurs on Linux or other SMB clients.

About CIFS Access from Non-Windows Hosts

Normally, CIFS access from non-Windows hosts is not supported.
Since NFS connection is recommended from Linux hosts, support for CIFS connection is on a best-effort basis. When considering the use of CIFS access from non-Windows hosts, contact your sales representative or partner sales representative.

Cautions When the Access-Based Enumeration (ABE) Function Is Enabled

When the ABE function is enabled in a CIFS environment, the process to check the access permissions of all folders and files takes a while. The more users, folders, and files there are, the longer it may take. If the length of the process time is a concern, disable the ABE function.

Cautions When Using System-Level Diagnosis

Do not perform a system-level diagnosis (the diagnosis function executed using the "boot_diags" command on the LOADER prompt) if the version of ONTAP in the ETERNUS AX4100/AX4100 ASA or the ETERNUS HX6100 is earlier than 9.7 P12.
Doing so may corrupt the data on the boot device and prevent the system from booting.
If this occurs, the boot device must be replaced.

Cautions about Encrypting the Root Volume of a Node

In ONTAP 9.8 and later, Volume Encryption can be used to protect the root volume of a node, but the root volume cannot be decrypted once encrypted.
To revert an encrypted root volume to an unencrypted root volume, the ETERNUS AX/AC/HX needs to be set up again.

Restrictions on Autonomous Ransomware Protection

Autonomous Ransomware Protection is not supported for the following configurations.

- Unsupported System Configuration
  - SAN environment
  - ONTAP S3 environment
  - VMDKs in NFS
  - Cloud Volumes Service for AWS / Google Cloud
  - Azure NetApp Files
  - Amazon FSx for ONTAP

- Unsupported Volume Type
  - Offline volume
  - Restricted volume
  - SnapLock volume
  - FlexGroup volume
  - FlexCache volume (*1)
  - SAN-only volume
  - Volumes in a suspended Storage VM
  - Storage VM root volume
  - Data protection volume

*1: Autonomous Ransomware Protection is supported for transmission source FlexVol volumes, but not cache volumes.

Cautions about ADP Environments

- In ADP environments, do not change the state of the container drive assignment from the default state. If only half the container drives that are installed in the controller shelf (*1) are assigned to each node, the "ADP DISABLED" alert may be reported and some actions may be required the next time the ETERNUS AX/AC/HX is rebooted.
- If ADP is used with the Autoassign setting disabled (enabled by default), "Spare low" is reported because the spare partition, which is created when a spare drive is used during an HDD failure, is not automatically assigned to a node. Therefore, keep the Autoassign setting enabled when ADP is used.

  *1: Equivalent to the number of drives in a single drive shelf for midrange and high-end models.

  - Checking and configuring the Autoassign setting

```
cluster::> storage disk option show -fields autoassign
cluster::> storage disk option modify -autoassign on -node *
```

Cautions about NFSv4

In a volume whose LANG setting has not been set to UTF-8, filenames that contain non-ASCII characters are rejected.

Cautions about the UTA Port

For the UTA port, use the following command:

```
# system node hardware unified-connect show
```

Do not use the following command:

```
# run -node node_name system hardware unified-connect show
```

Cautions When Using SnapMirror

Cautions When Using SnapMirror Synchronous

If SnapMirror Synchronous is configured, an inconsistency between the ONTAP version of the source and that of the destination may result in the following problems. Configure SnapMirror Synchronous with the same version of ONTAP.

- An operation related to SnapMirror fails
- The status display is out of sync

Cautions When Multiple Clients Access SnapMirror Destination Volumes

For environments where the version of ONTAP is ONTAP 9.14.1 and multiple clients concurrently access SnapMirror destination volumes, disable oplocks for those volumes.
If oplocks is enabled, the node may stop.

- Command to disable oplocks

```
cifs share properties remove -vserver <vserver_name> -share-name <share_name>
-share-properties oplocks
```

Note

This problem has been resolved in ONTAP 9.14.1 P4.

Cautions When Using FlexGroup Rebalancing

Even if FlexGroup Rebalancing is executed with the default settings, an imbalance may not be resolved. For example, in configurations that have many files selected as targets for rebalancing but the size of those files are smaller than the minimum file size (default of 100MB), an imbalance may persist even after the process is completed due to a large number of files that remain unmoved between constituents.

If the result of the rebalancing process is not what is expected, change the rebalancing setting with the "volume rebalance modify" command and try again. The minimum file size selected as the target for rebalancing can be specified with the "-min-file-size" option.

If a value smaller than the default minimum file size is specified, the number of files that are moved is more than when the default setting is used. Note that this may affect the read/write I/O performance.

Cautions When the Maximum Capacity Value Is Set for SVMs

The maximum capacity value of SVMs can now be set from ONTAP 9.13.1. When the maximum capacity value of SVMs is set, the following events occur.

- A capacity expansion of FlexGroups may fail
  When the capacity of a FlexGroup is expanded, even if the maximum capacity of SVMs has not been exceeded, the capacity expansion of the FlexGroup may fail.
  If this occurs, expand the capacity after releasing the maximum capacity setting of SVMs.

  Note

  This problem has been resolved in ONTAP 9.14.1.

Notes When Using the robocopy Command

The robocopy command updates the ctime after the initial copy, resulting in unmodified files being marked as "modified".

For files copied to a volume using the CIFS protocol, an attempt to perform a differential backup using the NDMP protocol may create an incomplete backup because the backup operation is skipped.
For details, refer to "ONTAP 9.x Setup Guide" (*1).

*1: The "x" in "ONTAP 9.x" indicates the minor version number. A free Fujitsu ID is required to read this manual.

Cautions When an ETERNUS AX1200 Is Configured with Eight Drives

For an ETERNUS AX1200 configured with the minimum number of eight drives, the Root-Data-Data (RD2) configuration of ADP has changed in ONTAP 9.13.1 P8 and later, and ONTAP 9.14.1 P1 and later.

For the ETERNUS AX1200, the RD2 configuration of ADP is as follows.

- For ONTAP 9.13.1 P7 and earlier, and ONTAP 9.14.1 GA
  1D + 2P + 1S
- For ONTAP 9.13.1 P8 and later, and ONTAP 9.14.1 P1 and later
  2D + 2P + 0S

Upgrading the ONTAP OS to a version where the change occurred does not change the RD2 configuration of ADP. The configuration remains the same as that before the update.
After the ONTAP OS is upgraded to that version, the RD2 configuration of ADP is changed by initializing ADP.

Notes When Using ServerProtect for NetApp

In environments using ServerProtect for NetApp 5.8, upgrading the virus search engine (VSAPI) to version 22.610 causes character corruption in subsequent virus detection logs.
Therefore, consider applying the following patch.

- Service Pack 1 Patch 2 Critical Patch build 1326

  For details, refer to the Trend Micro website.

Notes When Using the Multi-Admin Verification Function

If "volume create" is specified for [Operation], an approval is requested for CLI commands only.
An approval is not requested when a volume is created in ONTAP System Manager.

3. Notes and Cautions When Using ONTAP

This chapter describes notes and cautions when using ONTAP.

Notes and Cautions When Updating

Upgrading from ONTAP 9.7

If ONTAP 9.7 is configured with an Event Management System (EMS) message before being upgraded to ONTAP 9.8 or a later version, the EMS message may not be sent after the upgrade. If this event occurs, perform the following procedure after the upgrade to delete the management-ems LIF service from the service policies whose EMS messages do not reach the destination.

Procedure ▶▶▶

1. Identify the service policy associated with an LIF that can be used to send EMS messages.
   Example:

```
cluster1::> network interface show -fields service-policy-services
management-ems
```

2. Check if a connection to the EMS message destination is established for each LIF (do this on each node).
   Example:

```
cluster1::> network ping -lif node1-mgmt -vserver cluster-1 -destination
10.10.10.10
```

3. Switch to advanced privilege.
   Example:

```
cluster1::> set advanced
```

4. Delete the management-ems LIF service from the service policy associated with the LIF where a connection to the EMS message destination failed.
   Example:

```
cluster1::*> network interface service-policy remove-service -vserver svm_name
-policy service_policy_name -service management-ems
```

5. Check if the management-ems LIF service is only associated with LIFs whose EMS messages can reach the destination.
   Example:

```
cluster1::*> network interface show -fields service-policy -services
management-ems
```

◀◀◀

Upgrading from ONTAP 9.8 P12

When upgrading ONTAP 9.8 P12 or later to ONTAP 9.9.1 or ONTAP 9.10.1, upgrade it to ONTAP 9.9.1 P14 or later, or ONTAP 9.10.1 P10 or later.

Upgrading from a Version of ONTAP Earlier than 9.10.1 to ONTAP 9.10.1 or Later

When updating from a version of ONTAP earlier than 9.10.1 to ONTAP 9.10.1 or later in an environment that is using adaptive compression, upgrade to the following versions.

- ONTAP 9.10.1: P12 or later
- ONTAP 9.11.1: P8 or later
- ONTAP 9.12.1: P2 or later

If an upgrade is performed to a version older than the above, an error may occur when adaptive compressed data is read after that.

Upgrading to ONTAP 9.10.1

During an upgrade to ONTAP 9.10.1, a write to a cloud service that interoperates with FabricPool may result in a node being stopped. Stop access to the volumes targeted by FabricPool during an upgrade to ONTAP 9.10.1.

Upgrading to ONTAP 9.12.1

- Because the node may go offline due to a read process of the aggregate to which the drive with the media error belongs, make sure to upgrade to ONTAP 9.12.1 P4 or later.
- An upgrade to ONTAP 9.12.1 fails if the data protection (DP) mode is set for a SnapMirror relationship. Before the upgrade, convert the mode to the extended data protection (XDP) mode.
- ONTAP 9.11.1 and later releases do not support the ssh-ed25519 SSH public key type. Existing SSH public key accounts without a supported key algorithm must be reconfigured with a supported public key type before the upgrade. Without a reconfiguration, authentication of the administrator fails.

Upgrading to ONTAP 9.12.1 P4 or ONTAP 9.13.1

After the upgrade to ONTAP 9.12.1 P4 or ONTAP 9.13.1, the shelf firmware fails to be automatically updated so the update process may repeat.

- Phenomenon
  - If an automatic upgrade of the shelf firmware is attempted multiple times, the following events are repeatedly recorded in the EMS.

```
[cluster-01: dsa_disc: sfu.firmwareDownrev.shelf:error]: Shelf 0a.shelf0 has
downrev firmware.
[cluster-01: dsa_disc: sfu.firmwareDownrev.shelf:error]: Shelf 0a.shelf1 has
downrev firmware.
[cluster-01: dsa_sfu: sfu.firmwareDownrev:error]: Disk shelf firmware needs to be
updated on 2 disk shelves.
[cluster-01: dsa_sfu: sfu.downloadStarted:info]: Update of disk shelf firmware
started on 2 shelves.
[cluster-01: dsa_sfu: sfu.downloadSummary:info]: Shelf firmware updated on 2
shelves.
```

  - If there is a node being updated due to repeated attempts of a shelf firmware update, "running" is displayed as follows.
    Input:

```
set advanced
storage shelf firmware show-update-status
```

    Example:

```
cluster::> set advanced
Warning: These advanced commands are potentially dangerous; use them only when
         directed to do so by Fujitsu personnel.
Do you want to continue? {y|n}: y

cluster::> storage shelf firmware show-update-status
                   Update  In-Progress
Node               Status  Count
------------------ ------- -----------
cluster-01         running          10
cluster-02         idle              -
2 entries were displayed.
```

- The old version is displayed when the shelf firmware version is checked.
  The current version of the shelf firmware is displayed as the value for "FW Rev." when the "storage shelf show -module" command is used.
  Input:

```
storage shelf show -module
```

  Example:

```
cluster::> storage shelf show -module
              Shelf Name: 1.0
                Stack ID: 1
                Shelf ID: 0
               Shelf UID: 50:0a:09:80:0b:51:85:81
           Serial Number: SHFFG1925000070
             Module Type: IOM12E
                   Model: DS224-12
            Shelf Vendor: NETAPP
              Disk Count: 8
         Connection Type: SAS
             Shelf State: Online
                  Status: Normal

Modules:                                          Module is
                                 Monitor    Is    Reporting          Latest   Swap Operational
ID  Part No.     ES Serial No.   is Active Master Element   FW Rev. FW Rev. Count Status
--- ------------ --------------- --------- ------ --------- ------- ------- ----- -----------
  A 111-03964+D0 021920015102    true      true   true      0240    -       0     normal
  B 111-03964+D0 021919048534    true      false  false     0240    -       0     normal
Errors:
-------
```

  In addition, for the "run -node <node_name> -command sysconfig -a" command, the current version is displayed as follows.
  Input:

```
run -node <node_name> -command sysconfig -a
```

  Example:

```
cluster::>run -node cluster1-01 -command sysconfig -a
(omitted)
Shelf 0: DS224-12 Firmware rev. IOM12E A: 0240 IOM12E B: 0240
(omitted)
```

- Action

  If an automatic upgrade of the shelf firmware is not completed after multiple attempts, reboot each node to update it successfully.
  For the procedure to reboot each node, refer to "ONTAP 9.13.1 Setup Guide" (*1).

  *1: A free Fujitsu ID is required to read this manual.

Cautions after Upgrading ONTAP

If ONTAP is upgraded to the following versions or later, delete the ONTAP image used for the upgrade.

- ONTA 9.11.1 P8
- ONTAP 9.12.1 P2
- ONTAP 9.13.1 GA

In the upgrade status confirmation (performed in the API/CLI) process after upgrading to the above versions, the ONTAP image used for the upgrade may be mounted multiple times in systemshell and may cause nodes to stop.
The following procedure can be performed to determine whether multiple mounts are occurring.

Procedure ▶▶▶

1 Change to the Diag mode.

```
set -priv diag
```

2 Execute df in systemshell.

```
system node systemshell -node nodename -command "df -ik"
```

◀◀◀

Note

- If multiple mounts have occurred, multiple mount points will be displayed under /mroot/etc/NDU/mnt/store.
- The increase in multiple mounts can be stopped by deleting the ONTAP image used for the upgrade.
  Example:
  `cluster image package delete -version 9.12.1P2`
- Multiple mounts can be deleted using the takeover/giveback operation on the relevant node.

Notes on Reverting the Firmware

- To revert the firmware, the general rule is to make sure Fujitsu Support is contacted.
  Refer to "ETERNUS AX/HX series Upgrade and Revert/Downgrade Guide" for more details.
- The ONTAP version cannot be reverted to ONTAP 9.10.1 or earlier if a file is set using SnapLock with a retention time later than January 19, 2071 at 8:44:07 am.

Cautions When Using ONTAP System Manager

- In ONTAP 9.12.1, "Show Cables" has been added to the Cluster > Hardware menu. If "Show Cables" is clicked, the cable connections using the cluster switch and the cable connections between shelves can be displayed, but they may not be displayed depending on the configuration.
- In ONTAP System Manager 9.7 and 9.8, use of ONTAP REST API Reference is restricted because the references are not accessible.
  When using ONTAP REST API, download "FUJITSU Storage ETERNUS AX/HX series REST API Developers Guide" from Fujitsu's manual hosting site and execute the API.
- In ONTAP System Manager 9.8, the "Disk Shelves I/O Modules" screen, which appears by selecting CLUSTER > Disks on the menu, contains the "View Disk Qualification Package" link. Clicking this link may download the Disk Qualification Package, which is not supported by Fujitsu. Therefore, do not click the link.
- In ONTAP System Manager 9.7, the link for Help/Documentation is in the "?" button on the top bar. Currently, the link destination is Fujitsu's support site.
  Documents related to ONTAP System Manager are at Fujitsu's manual hosting site. The site can be accessed from the listed URL or by performing the following procedure from Fujitsu's support site.

  - Fujitsu's support site

    1 In the following order, select the links that are displayed in the top bar of the support site.
      Support > Product

    2 The Support for Products page is displayed. In the following order, select the links in the Support category that are displayed to the left of the Support for Products page.
      Products > Computing Products > Storage

    3 The Storage page is displayed. Select "Manuals" in this page.
      The site for the manuals can also be accessed directly.
      https://www.fujitsu.com/global/support/products/computing/storage/manuals-list.html

- In ONTAP System Manager 9.8, 9.9.1, and 9.10.1, when you click "+Add Cloud Tier" on the "Tiers" screen which appears by selecting STORAGE > Tiers on the menu, ONTAP S3 and StorageGRID should become enabled but may appear disabled (dimmed).
  For that reason, add ONTAP S3 or StorageGRID as a cloud tier using the command line interface (CLI). For details about adding a cloud tier using CLI, refer to the following document at Fujitsu's manual hosting site.

  - Fujitsu's manual hosting site

    https://www.fujitsu.com/global/support/products/computing/storage/manuals-list.html
    "ETERNUS AX/HX Series ONTAP® 9 Managing Storage Tiers By Using FabricPool"

  Note

  This phenomenon has been fixed in ONTAP System Manager 9.8 P12 and later, 9.9.1 P9 and later, and 9.10.1 P3 and later.

- ONTAP System Manager does not necessarily display the latest content because the content is not automatically updated. Refresh the page with the F5 key or by selecting the desired item again after selecting a different item from the left pane.

Cautions for the ONTAP Version

Cautions When Using ONTAP 9.7 P5

ONTAP 9.7 P5 has an issue with switchless cluster configurations that incorrectly disables the switchless cluster option when the option to automatically detect switchless clusters is enabled.
If the switchless cluster option is disabled in a switchless cluster configuration, the cluster does not operate correctly. In this situation, because the controllers cannot communicate with each other correctly, a controller feedback failure occurs.

Therefore, note the following when using ONTAP 9.7 P5:

(1) Do not change the setting of the option to automatically detect switchless clusters.
Specifically, do not run the following command:

```
network options detect-switchless-cluster modify [-enabled] true
```

(2) Disable the switchless cluster option after adding a cluster switch.
To disables the switchless cluster option, run the following command:

```
network options switchless-cluster modify [-enabled] false
```

Note

- This command requires advanced privileges to run.
- To check the setting of the switchless cluster option, run the command shown below. This command also requires advanced privileges to run.

```
network options switchless-cluster show
```

- If a cluster switch is added and later removed, a workaround must be performed for the issue of the switchless cluster option setting after the cluster switch is removed and the cluster configuration become switchless. Contact your Fujitsu sales representative, partner sales representative, or Fujitsu maintenance engineer for the workaround if a cluster switch is to be added and then later removed.

(3) Contact your Fujitsu maintenance engineer before initializing or setting up a cluster.
When a cluster is initialized or set up while ONTAP 9.7 P5 is running, a workaround must be performed for the issue regarding the switchless cluster option setting. Contact your Fujitsu maintenance engineer for the workaround.

(4) After updating ONTAP to a version without the issue of the switchless cluster option setting, enable the option to automatically detect switchless clusters.
To enable the option to automatically detect switchless clusters, run the following command:

```
network options switchless-cluster modify [-enabled] true
```

Note

- This command requires advanced privileges to run.
- To check the setting of the switchless cluster option, run the command shown below. This command also requires advanced privileges to run.

```
network options detect-switchless-cluster show
```

- The ONTAP versions without the issue of the switchless cluster option setting are ONTAP 9.7 P12 or later, or ONTAP 9.8 or later.
- If the ONTAP version is reverted to 9.7 P5, a workaround must be performed for the issue of the switchless cluster option setting after the reversion. Contact your Fujitsu maintenance engineer for the workaround before reverting the ONTAP version to 9.7 P5.

Cautions When Using ONTAP 9.8/ONTAP 9.9.1

If a node name is changed in ONTAP 9.8/ONTAP 9.9.1, the management gateway may stop working after an upgrade to ONTAP 9.10.1 or later.
If a node name is changed on ONTAP 9.8 or ONTAP 9.9.1, rename the node to the original name before the upgrade.

Cautions When Using ONTAP 9.10.1

In ONTAP 9.10.1, if the following string is included in the volume name, the volume may not be displayed in ONTAP System Manager.

Example: __1234 (*1)

*1: Two underscores "_" followed by a four digit number

In ONTAP 9.10.1, do not use the above string as part of the volume name.

Cautions When Using ONTAP 9.10.1 P11

When the Autonomous Ransomware Protection (ARP) function is used together with SVM-DR, the snapshots created for ARP on the source SVM may not be deleted automatically. This may cause a capacity shortage on the source SVM. Therefore, if ARP is used together with SVM-DR, monitor the source SVM capacity regularly and delete any snapshots for ARP that are no longer needed.

Notes When Using ONTAP 9.11.1

- Setting the timeout period during an inactive state

  A function to set the timeout period during an inactive state is supported from ONTAP 9.11.1.
  A pop-up appears to confirm whether to log out when one minute remains.
  If the specified value is too small, GUI operations may be affected because the pop-up appears immediately.
  To change the inactivity timeout period to a shorter time, set the value at least five minutes or longer.

Cautions When Using ONTAP Versions 9.11.1 to 9.11.1 P2

- In ONTAP versions 9.11.1 to 9.11.1 P2, the following problems may occur when the FabricPool/Cloud Volumes ONTAP (CVO) function is used.
  Use the FabricPool/CVO function in ONTAP version 9.11.1 P3 or later. ONTAP 9.11.1 P3 is planned for release at a later time.
  - A FabricPool operation fails
  - Creating an aggregate fails in Cloud Manager
  - A failure occurs in replications that use SnapMirror between CVOs
  - A Multi-Disk Panic (MDP) occurs in an environment deployed on Azure
  - The Tiering dashboard of Cloud Manager displays an error
- In ONTAP versions 9.11.1 to 9.11.1 P2, disable cluster peer encryption for newly created cluster peers.
  Example:

```
::> cluster peer modify -cluster cluster1 -encryption-protocol-proposed none
-auth-status-admin use-authentication
```

Cautions When Using ONTAP 9.12.1

- ONTAP 9.12.1 has a command option that can force automatic switchover even when MetroCluster is configured and the NVRAM is not synchronized (metrocluster modify -allow-auto-forced-switchover).
  If this option is used when the unsynchronized area in the NVRAM contains write data, the write data between clusters will become inconsistent. Therefore, do not use this option if a risk of corruption or loss of write data is not acceptable.
- In ONTAP 9.12.1, do not set the default retention period (-default-retention-period) of SnapLock volumes to "unspecified".

Cautions When Using ONTAP 9.12.1/ONTAP 9.13.1

- ONTAP 9.12.1 and ONTAP 9.13.1 do not support conversions to FlexGroups or FlexGroup rebalancing for the following volumes:
  - Volumes that are migrated from 7-Mode and have non-Unicode names
- ONTAP 9.13.1 does not support NDMP restores to a primary cluster volume (*1) in the following environment:
  - SnapMirror-BC environment or SnapMirror-Sync environment
  - The ONTAP version of the primary cluster is ONTAP 9.13.1
  - The ONTAP version of the secondary cluster is ONTAP 9.13.1 P2 or later

  *1: NDMP restore is supported if ONTAP 9.13.1, or ONTAP 9.13.1 P2 or later version is running on both the primary cluster and secondary cluster.
- Do not create a file in a directory that has a level of 500 or lower for the following ONTAP versions. Creating such a file may cause a node to stop.
  - ONTAP version: 9.12.1 P4 to 9.12.1 P7
  - ONTAP version: 9.13.1 to 9.13.1 P5

Cautions When Using ONTAP 9.16.1

In ONTAP 9.16.1 iSCSI/FCP environments, do not configure QoS for LUNs that are currently in use, or for any SVM, volume, or file associated with those LUNs. Applying QoS in these scenarios can lead to availability issues (such as inaccessible systems, unresponsive systems, or hidden LUNs).
Additionally, be aware that System Manager may automatically configure QoS. Use the following commands to check whether QoS policies are currently in effect.

- SVM

```
vserver show -vserver svm1 -fields qos-policy-group
```

- Vol

```
vol show -vserver <vserver_name> -volume <volume_name> -fields qos-policy-group
```

- LUN

```
lun show -vserver <vserver_name> -path <lun_path> -fields qos-policy-group
```

- File

```
qos workload show -vserver <vserver_name> -volume <volume_name> -fields
file,policy-group
```

Caution for the ONTAP Command

About the Parameters of the storage failover Command

Although operating the "auto-giveback" and "auto-giveback-after-panic" parameters of the "storage failover" command with "false" is recommended, if these parameters are set to "true" due to the customer's requirements, these parameters are set to "false" after the ONTAP upgrade. Therefore, if these parameters have been set to "true", set these parameters to "true" again after the ONTAP upgrade.

Cautions When Applying ONTAP One to a Storage System That Does Not Support the S3 SnapMirror License

If the ONTAP version is 9.11.1 or earlier and NLF is applied using ONTAP System Manager, an error appears. The following error appears if NLF (ONTAP One) is applied to a storage system that does not support the S3 SnapMirror license and the ONTAP version is 9.11.1 or earlier.

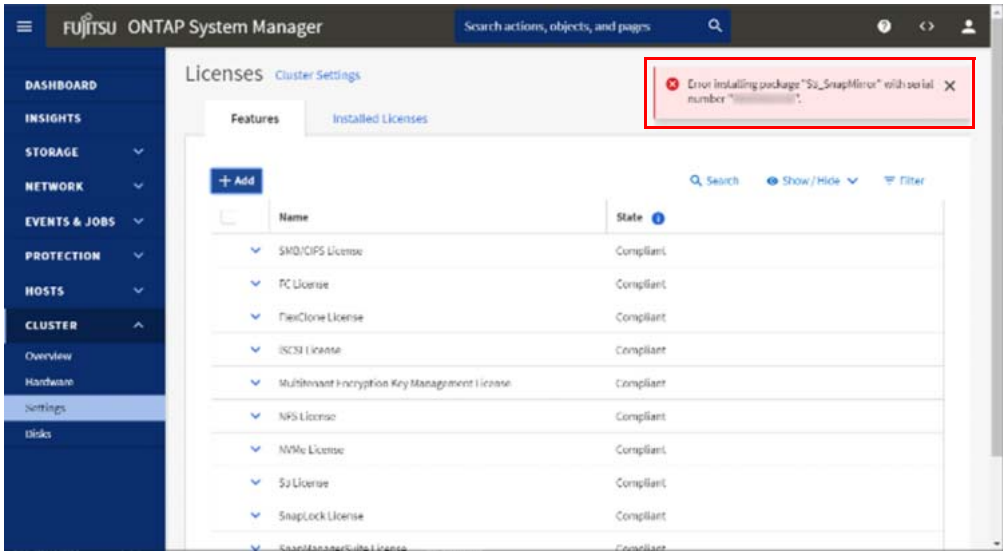

To use S3 SnapMirror, upgrade ONTAP to 9.12.1 or later and try again.

4. Notes and Cautions for Optional Software

This chapter describes notes and cautions for optional software.

Cautions When Using SnapCenter

When Using SnapCenter Version 4.3.1

For SnapCenter Version 4.3.1, when the SVM of the ETERNUS AX/AC/HX is registered using the "Add Storage System" menu, there are cases when the SVM of the ETERNUS AX/AC/HX cannot be registered.
If this issue occurs, perform the following action and register the SVM of the ETERNUS AX/AC/HX again using the "Add Storage System" menu of SnapCenter.

- Set the SVM to be registered in SnapCenter with Admin LIF using ONTAP System Manager or CLI.

Note

This issue is fixed in SnapCenter Version 4.4.

When Upgrading SnapCenter 4.4 to 4.6

During a SnapCenter Plug-in for Oracle (SCO) Plug-in upgrade, the services (spl.service and scc.service) are restarted on the host.

When Using SnapCenter and SnapCenter Plugin for VMware vSphere (SCV)

To register SVMs from SnapCenter and SCV, management LIFs must exist in those SVMs. In addition, an SVM name and the IP address of an SVM management LIF must be resolvable for both forward and reverse lookups.

When Upgrading to SnapCenter 5.0

- SVMs with Only FC Enabled Disappear from SnapCenter

  In an environment that is upgraded to SnapCenter 5.0, SVMs enabled with only FC disappear from the cluster and the following error message may appear.

  ```
  Create CG snapshot'<SNAPSHOT_NAME>'failed:Error:The credentials for the storage system<SVM_NAME>is not configured.
  ```

  After adding management LIFs in the SVM, perform a re-detection from SnapCenter. Note that enabling management LIFs is not required.

- An Error Occurs during the Installation of the Upgrade

  The following error message appears during the installation of the upgrade and the installation may fail.
  - Communication error

    ```
    ERROR: Failed to create SnapCenter Schema. Exception: Bad handshake
    ```

  - Access error

    ```
    ERROR: failed to create SnapCenter Schema. Exception: Access denied for user
    'root'@'localhost' (using password NO)
    ```

  Reboot the storage system and try again.

When Using SnapCenter 5.0

- VMs with Only FC Enabled Are Not Detected from SnapCenter

  SVMs with only FC enabled may not be detected from SnapCenter. If this occurs, add management LIFs in the SVM and perform a re-detection from SnapCenter.
  Note that enabling management LIFs is not required.

When Using SnapCenter 6.0

- Adding a Windows Host Fails

  Adding a Windows host in SnapCenter GUI may fail with the error message "The required Java version is not found on <HostName>. The SC plugins are compatible with Java 11 version." regardless of whether or not Java 11 is installed.
  If this error message appears, confirm that Java 11 is installed in the Windows host and add the host after selecting [Skip optional preinstall checks] under [More Options].

Notes When Using SnapCenter Plug-in for VMware vSphere (SCV)

When a vCenter Task of an SCV Backup Job Is Stuck at 2% after the Job Is Completed

- Phenomenon

  For environments that are using SCV 4.3 to 4.6, the vCenter task of an SCV backup job may be stuck at 2% even after the job is completed.

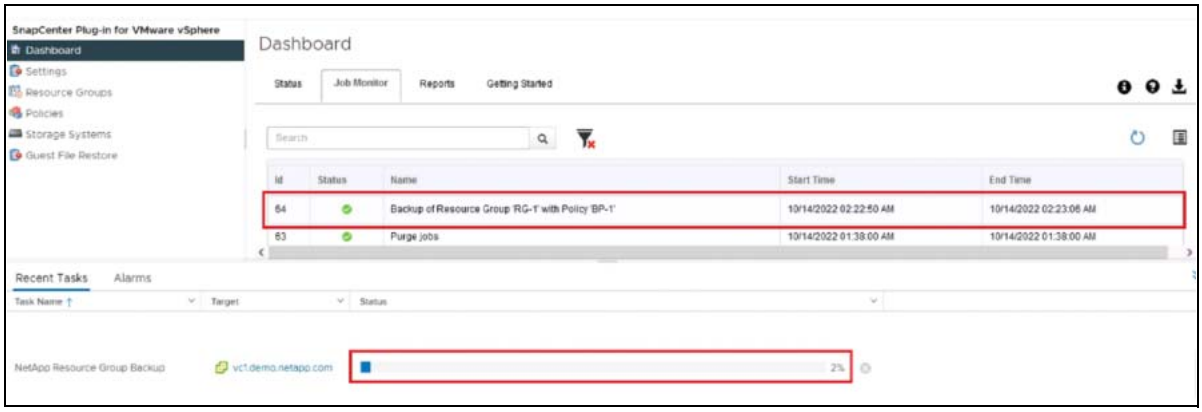

- Action

  This phenomenon has been fixed in SCV 4.7, so perform the following actions.

  (1) Upgrade to SCV 4.7.

  (2) Restart the vmware-vpxd (vpxd) service on the vCenter server, and then clear the previously stuck vCenter tasks.

When Adding an SVM Using SVM Management LIF in SCV

To add an SVM using SVM management LIF, the target aggregate must be added to [List of Aggregates Assigned]. If the aggregate is not included in the list (aggr-list), the following error is displayed and the SCV restore process fails.

Example:

Reason: aggregate <Aggregate_Name> is not in aggr-list of Vserver <Vserver_Name>

If this error occurs, use the "vserver modify" command to assign the aggregate specified/displayed in the error message to the target SVM.

When Storage Systems Cannot Be Registered with Certificate Authentication

If certificate authentication is used when registering a storage system, the following error may occur and the storage system cannot be registered in the SCV.
If this error continues, change the authentication method to credentials and then register the storage system.

```
Failed to add/update storage system information:  Failed to connect to storage system <SVM_NAME_OR_IP_ADDR>.
This instance has already started one or more requests. Properties can only be modified before sending the first request.
```

Registration or Update of a Storage System Fails on SCV 5.0 or Later

- If a Storage System Is Deleted and then Registered Again, An Error Occurs

  When a storage system is deleted and immediately registered again, the following error may occur and the registration fails.

  ```
  Failed to add/update storage system information: Failed to connect to storage system <SVM_NAME_OR_IP_ADDR>.
  This instance has already started one or more requests. Properties can only be modified before sending the first request.
  ```

  Wait a few minutes and register (or update) the storage system again. If the registration fails even after waiting for a while, contact Fujitsu Support.

- Registration of a Storage System Fails due to the ONTAP User Not Having Sufficient Privileges

  If the ONTAP user who is specified when a storage system is registered cannot access HTTP or ONTAPI type applications, the registration may fail with the error shown below. In addition, connections to SVMs of the registered storage system fail and the volume list may not be displayed.

  ```
  Failed to add/update storage system information: Failed to connect to storage system <SVM_NAME_OR_IP_ADDR>.
  Invalid credentials. Please check the username or password
  ```

  For SCV 5.0 and later, HTTP and ONTAPI type applications must be added as user login methods for ONTAP users who are allowed to access the SCV.
  Use the "security login show" command to confirm that users can access HTTP and ONTAPI type applications.

Notes on Using SnapManager

Due to a SnapManager limitation, FlexGroup is not supported.

Cautions about Active IQ Config Advisor or Active IQ OneCollect

Cautions and Limitations When Using Active IQ Config Advisor or Active IQ OneCollect

- The Collection Target Command May Reply with Some Errors

  When information is collected using Active IQ Config Advisor or Active IQ OneCollect, the collection target command may reply with some errors due to compatibility and environment dependencies. In this case, the data is analyzed based on the collected information.

- Part of the File Exported from [Job Monitor] Is Corrupt

  When the Excel file exported from [Job Monitor] in Active IQ Config Advisor is opened, the following message appears: "We found a problem with some content in '*filename*'. Do you want us to recover as much as we can? If you trust the source of this workbook, click Yes.". Delete the "Nic Cards" sheet because the recovery process fails to create the sheet.

- The Report File Exported from [Results] May Not Be Created Correctly

  In Active IQ Config Advisor 6.7.1, the report file exported from [Results] may not contain the following information. If the report lacks the necessary information, extract the information from the Excel file exported from [Job Monitor].
  - Summary report (profiled controller information and cluster overview)
  - Cluster information (such as the network interfaces, LUN information, and volume information)
  - Individual node information (such as shelf/drive information, hardware information, and license information)

- The [Discover] Panel Cannot Be Used to Create a Data Collection Job

  In Active IQ OneCollect, a data collection job cannot be created for a device that is detected using the [Discover] panel in the right frame after the appropriate panel is selected for the purpose of collecting data. Manually create a data collection job by adding a device input field using [Device Type] and [Device Category] in the left frame and then entering the necessary device information.

  Note

  This problem has been resolved in Active IQ OneCollect 2.7.2.

Cautions When Upgrading the Software

When upgrading the Active IQ Config Advisor software or Active IQ OneCollect software, do not delete the database.
The software must be upgraded by uninstalling the existing version of the software and then installing the latest version of the software.
In addition, the GUI screen must be shut down before the software is upgraded or uninstalled.

Caution

- When the message "Do you want to delete the Database for a clean installation?" appears, be sure to select "No". Selecting "Yes" clears the already registered information of the ETERNUS AX/AC/HX.
- When updating the software, uninstall the existing software version and delete the following ConfigAdvisorAIDE folder before installing the latest version of the software.
  C: \ Users \ <Username> \ ConfigAdvisorAIDE
  The collected data archives are stored in the ConfigAdvisorAIDE folder. To keep the ConfigAdvisorAIDE folder undeleted, rename it.

Notes When Using ONTAP Antivirus Connector

Using Data LIFs

In environments in which ONTAP Antivirus (AV) Connector 1.0.4, 1.0.4 P1, or 1.0.5 is being used, if Data LIF is used, ONTAP connections fail when Trend Micro ServerProtect uses AV Connector.
In addition, although the connection is re-established when ONTAP AV Connector is restarted, the problem will reoccur over time. Because of this, perform one of the following actions.

- Use Management LIF instead of Data LIF.
- As this phenomenon has been fixed in ONTAP AV Connector 1.0.5P1 and later, perform an upgrade.

Connecting to ServerProtect for Storage on NetApp

An SVM data LIF cannot be used to connect the ETERNUS AX/AC/HX to Trend Micro's ServerProtect for Storage on NetApp using ONTAP Antivirus Connector 1.0.5. Use an SVM management LIF or a cluster management LIF instead.

Note

This problem has been resolved in ONTAP Antivirus Connector 1.0.5 P1.

Cautions When Using Virtual Storage Console

If the ETERNUS AX2100 (ASA) model or the ETERNUS AX4100 (ASA) model is added in the "Storage System" menu of Virtual Storage Console 9.7, "Unknown" is displayed because the ETERNUS AX2100 (ASA) or the ETERNUS AX4100 (ASA) is not recognized.
Therefore, the ETERNUS AX2100 (ASA) model or the ETERNUS AX4100 (ASA) model cannot be used with Virtual Storage Console 9.7.

Cautions When Using RcfFileGenerator for MetroCluster IP v1.2

Do not generate an Rcf file in RcfFileGenerator for MetroCluster IP version 1.2 using "Configure switches for 4-8-4 Node Tech Refresh". This is because version 1.2 can generate Rcf files for models that do not support "Configure switches for 4-8-4 Node Tech Refresh". Use version 1.3 if an Rcf file needs to be generated using "Configure switches for 4-8-4 Node Tech Refresh".

Optional Software in the DVD Kit

The DVD kit for each ONTAP version may contain optional software of a version that is not supported by that ONTAP version.
The latest version of the software can be checked at the following public website (*1):

https://support.ts.fujitsu.com/IndexMySupport.asp?_gl=1*1x4v1nm*_ga*MTA0ODM5NDIxNS4xNjk0MDQ4MDA3*_ga_GNHKR21PZP*MTY5Nzc4NzEyMC42LjEuMTY5Nzc4NzIwNS4wLjAuMA

*1: A free Fujitsu ID is required to obtain the software.

Cautions about ONTAP tools for VMware vSphere (OTV)

When Using OTV 9.12

When the ETERNUS AX1200 ASA model is added to OTV 9.12, OTV 9.12 cannot recognize the model name and displays the model as "AFF A150". The model name is just displayed incorrectly and the functions are not affected.

When Upgrading to OTV 9.13P1

- The OTV Appliance Is Started in the GRUB Rescue Mode

  The following error may occur when upgrading to OTV 9.13P1 and the OTV appliance may start up in the GRUB Rescue mode.

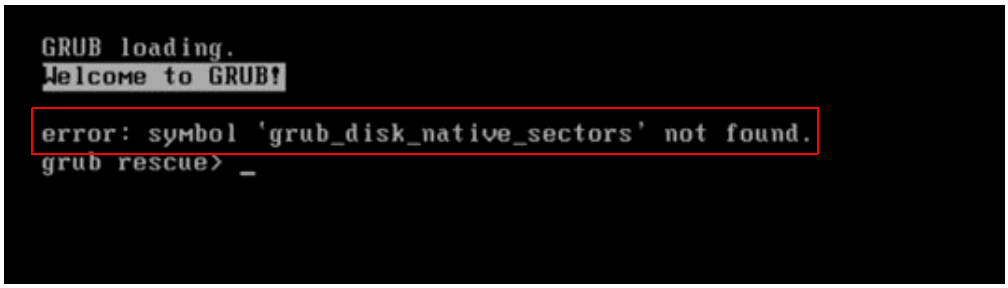

  Make sure to power off the OTV appliance and create a VM snapshot before upgrading to OTV 9.13P1. If the upgrade process fails, restore the VM from the VM snapshot and then try the OTV upgrade again.
  If this phenomenon occurs while upgrading to OTV9.13P1 without creating a VM snapshot, contact the Fujitsu Support.

- The Platform Information in the Storage Capability Profile Becomes Empty after the Upgrade

  After upgrading to OTV9.13P1, the platform information in the storage capability profile that is created by the user may become empty. If this phenomenon occurs, delete the storage capability profile that is created by the user and then create it again.

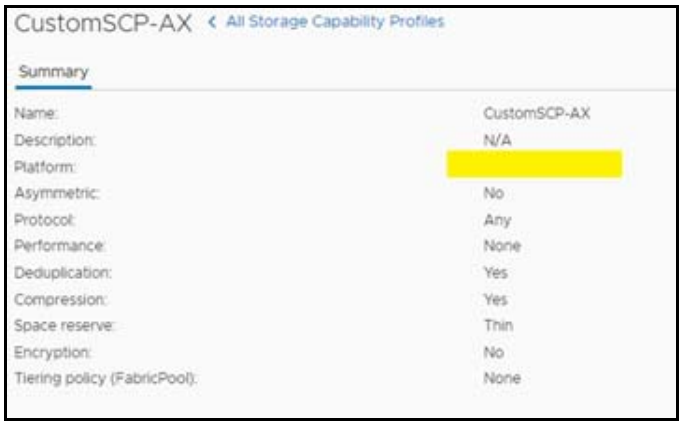

Cautions When Using Active IQ Unified Manager

When Using Active IQ Unified Manager 9.14

- Workload Provisioning Fails if [STORAGE EFFICIENCY POLICY (OPTIONAL)] Is Clicked

  For the ETERNUS AX/AC series, a proper provisioning is not performed if the [STORAGE EFFICIENCY POLICY (OPTIONAL)] pull-down list on the [Provision Workload] page is clicked. Specify each item without clicking [STORAGE EFFICIENCY POLICY (OPTIONAL)].

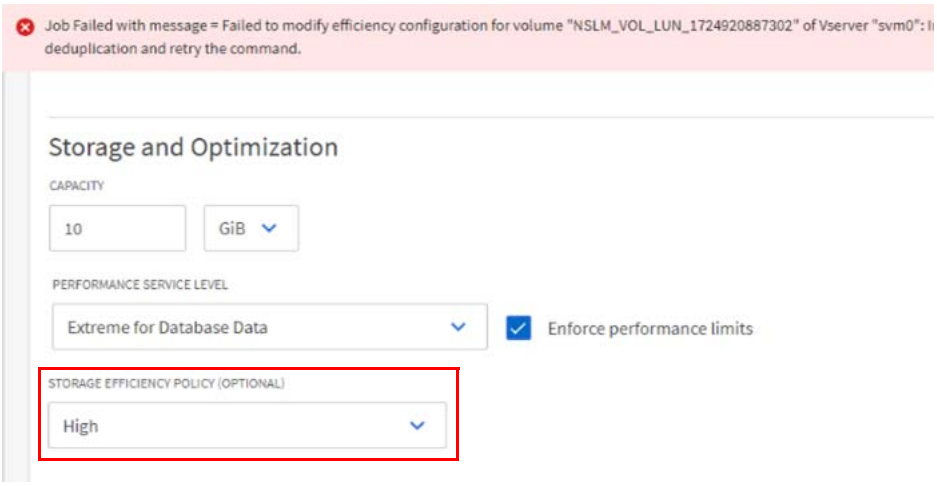

- Adding a Cluster May Fail with the Following Error

Unable to add cluster data source. This can occur if the clocks on the systems are not synchronized and the Active IQ Unified Manager HTTPS certificate start date is later than the date on the cluster, or if the cluster has reached the maximum number of EMS notification destinations.

If this error occurs, set the enable.cloudagent setting to false by performing the following procedure:

Procedure ▶▶▶

1. Open server.properties with a text editor.
   - For Windows
     C: \ Program Files \ netapp \ essentials \ conf \ server.properties
   - For Linux
     /opt/netapp/essentials/conf/server.properties
2. Change "enable.cloudagent=true" to "enable.cloudagent=false" and save the file.

3 Restart the service.

Note

For the procedure to restart the service, refer to the following website.

- For Windows
  https://storage-system.fujitsu.com/manual/en/active-iq-unified-manager/install-windows/task_restart_unified_manager_win.html
- For Linux
  https://storage-system.fujitsu.com/manual/en/active-iq-unified-manager/install-linux/task_restart_unified_manager.html

- The Root Aggregate of a Node May Appear in the List of Aggregates

  Active IQ Unified Manager 9.14 does not support monitoring of the root aggregate of nodes that use ONTAP 9.14 or later, but it may show the root aggregate in the list of aggregates. Do not perform any settings on the root aggregate of those nodes.

5. Contact Info

Contact your Fujitsu or partner sales representative.

A. FAQ

Question:

Can I disable unused ports?

Answer:

Yes. Execute the "network port modify" command with the "-up admin false" option to disable the specified port. Even if the port is disabled, the SFP information of the target port is displayed normally with the "run -node (nodename) sysconfig" command.
This operation is available for all ONTAP versions and ETERNUS AX/AC/HX models.

Fujitsu Storage
ETERNUS AX series All-Flash Arrays,
ETERNUS AC series All-Flash Arrays,
ETERNUS HX series Hybrid Arrays
Notes and Cautions When Using This Product

P3AG-5552-25ENZ0

Date of issuance: May 2025
Issuance responsibility: Fujitsu Limited

- The content of this manual is subject to change without notice.
- This manual was prepared with the utmost attention to detail.
  However, Fujitsu shall assume no responsibility for any operational problems as the result of errors, omissions, or the use of information in this manual.
- Fujitsu assumes no liability for damages to third party copyrights or other rights arising from the use of any information in this manual.
- The content of this manual may not be reproduced or distributed in part or in its entirety without prior permission from Fujitsu.


FUJITSU