

Fujitsu Storage ETERNUS DX60 S4, ETERNUS DX60 S3 Hybrid Storage Systems

Design Guide (Basic)



System configuration design

Table of Contents

| | |
|---|-----------|
| 1. Function | 12 |
| RAID Functions..... | 13 |
| Supported RAID..... | 13 |
| User Capacity (Logical Capacity)..... | 19 |
| RAID Group | 21 |
| Volume..... | 22 |
| Hot Spares | 24 |
| Data Protection..... | 26 |
| Data Block Guard..... | 26 |
| Disk Drive Patrol | 27 |
| Redundant Copy | 28 |
| Rebuild..... | 29 |
| Fast Recovery | 30 |
| Copyback/Copybackless..... | 31 |
| Protection (Shield)..... | 33 |
| Reverse Cabling..... | 34 |
| Operations Optimization (Virtualization)..... | 35 |
| Thin Provisioning | 35 |
| Optimization of Volume Configurations..... | 42 |
| RAID Migration | 43 |
| Logical Device Expansion..... | 45 |
| LUN Concatenation..... | 47 |
| Wide Striping | 49 |
| User Access Management | 50 |
| Account Management..... | 50 |
| User Authentication | 52 |
| Audit Log..... | 54 |
| Improving Host Connectivity..... | 55 |
| Host Affinity | 55 |
| iSCSI Security..... | 57 |
| Environmental Burden Reduction..... | 58 |
| Eco-mode | 58 |
| Power Consumption Visualization | 61 |

- Operation Management/Device Monitoring..... 62
 - Operation Management Interface..... 62
 - Performance Information Management..... 63
 - Event Notification 65
 - Device Time Synchronization 68
- Power Control 69
 - Power Synchronized Unit..... 69
 - Remote Power Operation (Wake On LAN) 70
- Backup (Advanced Copy) 71
 - Backup..... 72
- Performance Tuning 80
 - Striping Size Expansion 80
 - Assigned CMs 81
- Stable Operation via Load Control..... 82
 - Host Response 82
- Data Migration..... 83
 - Storage Migration 83
- Non-disruptive Storage Migration 85
- Server Linkage Functions 87
 - Oracle VM Linkage 87
 - VMware Linkage 88
 - Veeam Storage Integration 90
 - Microsoft Linkage..... 93
 - OpenStack Linkage 94
 - Logical Volume Manager (LVM) 95
- Smart Setup Wizard 96

2. Connection Configuration 101

- SAN Connection 101
 - Host Interface..... 101
 - Access Method..... 104
- LAN Connection 106
 - LAN for Operation Management (MNT Port) 106
 - LAN for Remote Support (RMT Port) 108
 - LAN Control (Master CM/Slave CM) 111

| | |
|--|------------|
| Network Communication Protocols | 113 |
| Power Supply Connection | 115 |
| Input Power Supply Lines..... | 115 |
| UPS Connection..... | 115 |
| Power Synchronized Connections..... | 116 |
| Power Synchronized Connections (PWC)..... | 116 |
| Power Synchronized Connections (Wake On LAN) | 119 |
| 3. Hardware Configurations | 120 |
| Configuration Schematics | 120 |
| Optional Product Installation Conditions..... | 127 |
| Controller Module..... | 127 |
| Drive Enclosures | 128 |
| I/O Module | 128 |
| Drives | 129 |
| Standard Installation Rules..... | 132 |
| Controller Module..... | 132 |
| Drive Enclosure | 132 |
| I/O Module | 133 |
| Drive | 133 |
| Recommended RAID Group Configurations | 135 |
| 4. Maintenance/Expansion | 138 |
| Hot Swap/Hot Expansion | 138 |
| User Expansion..... | 140 |
| Drive Sanitization..... | 140 |
| A. Function Specification List | 141 |
| List of Supported Protocols | 141 |
| Target Pool for Each Function/Volume List | 142 |
| Target RAID Groups/Pools of Each Function..... | 142 |
| Target Volumes of Each Function..... | 142 |
| Combinations of Functions That Are Available for Simultaneous Executions | 144 |
| Combinations of Functions That Are Available for Simultaneous Executions..... | 144 |
| Number of Processes That Can Be Executed Simultaneously | 145 |

Capacity That Can Be Processed Simultaneously 146

List of Figures

| | | |
|-----------|---|----|
| Figure 1 | RAID0 Concept | 14 |
| Figure 2 | RAID1 Concept | 14 |
| Figure 3 | RAID1+0 Concept | 15 |
| Figure 4 | RAID5 Concept..... | 15 |
| Figure 5 | RAID5+0 Concept | 16 |
| Figure 6 | RAID6 Concept..... | 17 |
| Figure 7 | RAID6-FR Concept..... | 18 |
| Figure 8 | Example of a RAID Group..... | 21 |
| Figure 9 | Volume Concept..... | 22 |
| Figure 10 | Hot Spares..... | 24 |
| Figure 11 | Data Block Guard..... | 26 |
| Figure 12 | Disk Drive Patrol..... | 27 |
| Figure 13 | Redundant Copy Function | 28 |
| Figure 14 | Rebuild..... | 29 |
| Figure 15 | Fast Recovery | 30 |
| Figure 16 | Copyback..... | 31 |
| Figure 17 | Copybackless..... | 32 |
| Figure 18 | Protection (Shield) | 33 |
| Figure 19 | Reverse Cabling..... | 34 |
| Figure 20 | Storage Capacity Virtualization..... | 35 |
| Figure 21 | TPV Balancing (When Allocating Disproportionate TPV Physical Capacity Evenly) | 39 |
| Figure 22 | TPV Balancing (When Distributing Host Accesses Evenly after TPP Expansion) | 39 |
| Figure 23 | TPV Capacity Optimization..... | 41 |
| Figure 24 | RAID Migration (When Data Is Migrated to a High Capacity Drive) | 43 |
| Figure 25 | RAID Migration (When a Volume Is Moved to a Different RAID Level)..... | 43 |
| Figure 26 | Logical Device Expansion (When Expanding the RAID Group Capacity)..... | 45 |
| Figure 27 | Logical Device Expansion (When Changing the RAID Level) | 45 |
| Figure 28 | LUN Concatenation | 47 |
| Figure 29 | LUN Concatenation (Volume Concatenation)..... | 48 |
| Figure 30 | Wide Striping..... | 49 |
| Figure 31 | Account Management | 50 |
| Figure 32 | Audit Log | 54 |
| Figure 33 | Host Affinity..... | 55 |
| Figure 34 | Associating Host Groups, CA Port Groups, and LUN Groups | 56 |
| Figure 35 | Eco-mode..... | 58 |
| Figure 36 | Power Consumption Visualization..... | 61 |
| Figure 37 | Event Notification..... | 65 |
| Figure 38 | Device Time Synchronization..... | 68 |
| Figure 39 | Power Synchronized Unit | 69 |
| Figure 40 | Wake On LAN..... | 70 |
| Figure 41 | Example of Advanced Copy..... | 71 |
| Figure 42 | Restore OPC..... | 74 |
| Figure 43 | EC Reverse..... | 74 |
| Figure 44 | Targets for the Multi-Copy Function..... | 75 |
| Figure 45 | Multi-Copy..... | 75 |
| Figure 46 | Multi-Copy (Including SnapOPC+) | 76 |
| Figure 47 | Cascade Copy..... | 77 |

| | | |
|-----------|---|-----|
| Figure 48 | Assigned CMs..... | 81 |
| Figure 49 | Host Response..... | 82 |
| Figure 50 | Storage Migration..... | 84 |
| Figure 51 | Non-disruptive Storage Migration..... | 85 |
| Figure 52 | Oracle VM Linkage..... | 87 |
| Figure 53 | VMware Linkage..... | 88 |
| Figure 54 | Veeam Storage Integration | 90 |
| Figure 55 | Microsoft Linkage | 93 |
| Figure 56 | Logical Volume Manager (LVM)..... | 95 |
| Figure 57 | RAID Configuration Example (When 12 SSDs Are Installed) | 98 |
| Figure 58 | RAID Configuration Example (When 15 SAS Disks Are Installed) | 100 |
| Figure 59 | Single Path Connection (When a SAN Connection Is Used — Direct Connection)..... | 104 |
| Figure 60 | Single Path Connection (When a SAN Connection Is Used — Switch Connection) | 104 |
| Figure 61 | Multipath Connection (When a SAN Connection Is Used — Basic Connection Configuration) .. | 105 |
| Figure 62 | Multipath Connection (When a SAN Connection Is Used — Switch Connection)..... | 105 |
| Figure 63 | Connection Example without a Dedicated Remote Support Port..... | 107 |
| Figure 64 | Connection Example When the IP Address of the Slave CM Is Set (and a Dedicated Remote Support Port Is Not Used) | 107 |
| Figure 65 | Overview of the AIS Connect Function..... | 108 |
| Figure 66 | Security Features | 109 |
| Figure 67 | Connection Example with a Dedicated Remote Support Port | 110 |
| Figure 68 | Connection Example When the IP Address of the Slave CM Is Set (and a Dedicated Remote Support Port Is Used)..... | 111 |
| Figure 69 | LAN Control (Switching of the Master CM)..... | 112 |
| Figure 70 | LAN Control (When the IP Address of the Slave CM Is Set) | 112 |
| Figure 71 | Power Supply Control Using a Power Synchronized Unit (When Connecting One or Two Servers)..... | 117 |
| Figure 72 | Power Supply Control Using a Power Synchronized Unit (When Connecting Three or More Servers)..... | 118 |
| Figure 73 | Power Supply Control Using Wake On LAN..... | 119 |
| Figure 74 | Minimum Configuration Diagram: ETERNUS DX60 S4 | 120 |
| Figure 75 | Minimum Configuration Diagram: ETERNUS DX60 S3 | 121 |
| Figure 76 | Maximum Configuration Diagram: ETERNUS DX60 S4 | 122 |
| Figure 77 | Maximum Configuration Diagram: ETERNUS DX60 S3 | 124 |
| Figure 78 | Enclosure Connection Path (When Only One Controller Is Installed)..... | 126 |
| Figure 79 | Enclosure Connection Path (When Two Controllers Are Installed) | 126 |
| Figure 80 | Controller Installation Order | 132 |
| Figure 81 | I/O Module Installation Order | 133 |
| Figure 82 | Installation Diagram for 2.5" Drives..... | 134 |
| Figure 83 | Installation Diagram for 3.5" Drives..... | 134 |
| Figure 84 | Drive Combination 1 | 135 |
| Figure 85 | Drive Combination 2..... | 135 |
| Figure 86 | Drive Combination 3..... | 136 |
| Figure 87 | Drive Combination 4..... | 137 |
| Figure 88 | Drive Combination 5..... | 137 |

List of Tables

| | | |
|----------|--|-----|
| Table 1 | Basic Functions..... | 12 |
| Table 2 | RAID Level Comparison | 18 |
| Table 3 | Formula for Calculating User Capacity for Each RAID Level..... | 19 |
| Table 4 | User Capacity per Drive | 20 |
| Table 5 | RAID Group Types and Usage | 21 |
| Table 6 | Number of Drives to Configure and the Recommended Number of Drives per RAID Group | 22 |
| Table 7 | Volumes That Can Be Created..... | 23 |
| Table 8 | Hot Spare Selection Criteria..... | 25 |
| Table 9 | TPP Maximum Number and Capacity | 36 |
| Table 10 | Chunk Size According to the Configured TPP Capacity..... | 36 |
| Table 11 | Levels and Configurations for a RAID Group That Can Be Registered in a TPP..... | 36 |
| Table 12 | TPP Thresholds..... | 37 |
| Table 13 | TPV Thresholds..... | 37 |
| Table 14 | Optimization of Volume Configurations..... | 42 |
| Table 15 | Available Functions for Default Roles..... | 51 |
| Table 16 | Client Public Key (SSH Authentication)..... | 52 |
| Table 17 | Eco-mode Specifications..... | 59 |
| Table 18 | Levels and Contents of Events That Are Notified..... | 65 |
| Table 19 | SNMP Specifications..... | 66 |
| Table 20 | Control Software (Advanced Copy)..... | 71 |
| Table 21 | List of Functions (Copy Methods) | 72 |
| Table 22 | Characteristics of SnapOPC/SnapOPC+ Operations with Each Type of Copy Destination Logical Volume | 73 |
| Table 23 | Available Cascade Copy Combinations (When a Cascade Copy Performs Session 1 Followed by Session 2)..... | 77 |
| Table 24 | Modified Data Size for Each QuickOPC Resolution..... | 78 |
| Table 25 | Available Cascade Copy Combinations (When a Cascade Copy Performs Session 2 Followed by Session 1) | 78 |
| Table 26 | Available Stripe Depth..... | 80 |
| Table 27 | Specifications for Paths and Volumes between the Local Storage System and the External Storage System..... | 85 |
| Table 28 | Volume Types That Can Be Used with Veeam Storage Integration | 92 |
| Table 29 | Guideline for the Number of Drives and User Capacities (When 1.92TB SSDs Are Installed) | 96 |
| Table 30 | Guideline for the Number of Drives and User Capacities (When 1.2TB SAS Disks Are Installed) | 98 |
| Table 31 | Ethernet Frame Capacity (Jumbo Frame Settings)..... | 102 |
| Table 32 | LAN Port Availability..... | 113 |
| Table 33 | Drive Characteristics | 131 |
| Table 34 | Number of Installable Drives..... | 131 |
| Table 35 | Hot Swap and Hot Expansion Availability for Components (ETERNUS DX60 S4) | 138 |
| Table 36 | Hot Swap and Hot Expansion Availability for Components (ETERNUS DX60 S3) | 139 |
| Table 37 | List of Supported Protocols..... | 141 |
| Table 38 | Combinations of Functions That Can Be Executed Simultaneously (1/2)..... | 144 |
| Table 39 | Combinations of Functions That Can Be Executed Simultaneously (2/2) | 144 |

Fujitsu would like to thank you for purchasing the Fujitsu Storage ETERNUS DX60 S4, ETERNUS DX60 S3 (hereinafter referred to as ETERNUS DX).

The ETERNUS DX is designed to be connected to Fujitsu servers (Fujitsu SPARC Servers, PRIME-QUEST, PRIMERGY, and other servers) or non-Fujitsu servers.

This manual provides the system design information for the ETERNUS DX storage systems.

This manual is intended for use of the ETERNUS DX in regions other than Japan.

This manual applies to the latest controller firmware version.

Twenty-First Edition

July 2023

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 field engineers or system administrators who design ETERNUS DX systems or use the ETERNUS DX.

Related Information and Documents

The latest version of this manual and the latest information for your model are available at:

<https://www.fujitsu.com/global/support/products/computing/storage/manuals-list.html>

Refer to the following manuals of your model as necessary:

"Overview"

"Site Planning Guide"

"Product List"

"Configuration Guide (Basic)"

"Configuration Guide -Server Connection-"

"Configuration Guide (Web GUI)"

"ETERNUS Web GUI User's Guide"

"ETERNUS CLI User's Guide"

Document Conventions

■ Third-Party Product Names

- Oracle Solaris may be referred to as "Solaris", "Solaris Operating System", or "Solaris OS".
- Microsoft® Windows Server® may be referred to as "Windows Server".

■ Notice Symbols

The following notice symbols are used in this manual:



Indicates information that you need to observe when using the ETERNUS storage system. Make sure to read the information.



Indicates information and suggestions that supplement the descriptions included in this manual.

Warning Signs

Warning signs are shown throughout this manual in order to prevent injury to the user and/or material damage. These signs are composed of a symbol and a message describing the recommended level of caution. The following explains the symbol, its level of caution, and its meaning as used in this manual.



This symbol indicates the possibility of serious or fatal injury if the ETERNUS DX is not used properly.



This symbol indicates the possibility of minor or moderate personal injury, as well as damage to the ETERNUS DX and/or to other users and their property, if the ETERNUS DX is not used properly.

IMPORTANT This symbol indicates IMPORTANT information for the user to note when using the ETERNUS DX.

The following symbols are used to indicate the type of warnings or cautions being described.

Electric Shock



△The triangle emphasizes the urgency of the WARNING and CAUTION contents. Inside the triangle and above it are details concerning the symbol (e.g. Electrical Shock).

No Disassembly



⊘The barred "Do Not..." circle warns against certain actions. The action which must be avoided is both illustrated inside the barred circle and written above it (e.g. No Disassembly).

Unplug



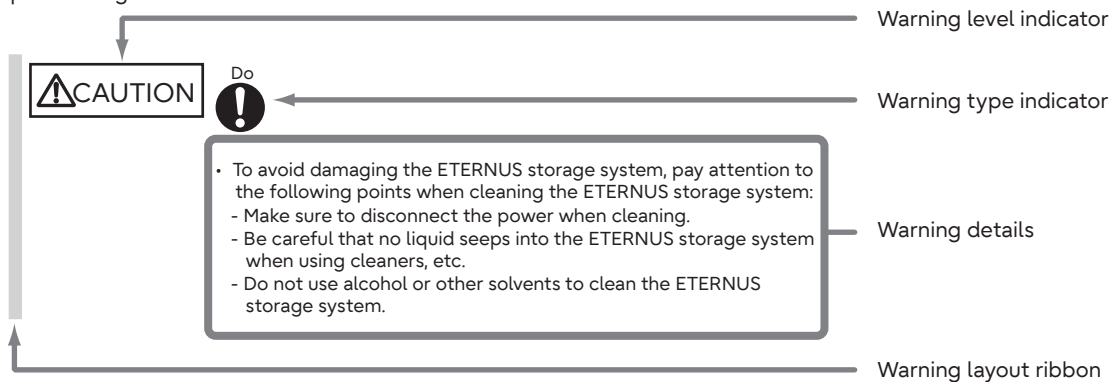
●The black "Must Do..." circle indicates actions that must be taken. The required action is both illustrated inside the black disk and written above it (e.g. Unplug).

How Warnings are Presented in This Manual

A message is written beside the symbol indicating the caution level. This message is marked with a vertical ribbon in the left margin, to distinguish this warning from ordinary descriptions.

A display example is shown here.

Example warning



1. Function

The ETERNUS DX provides various functions to ensure data integrity, enhance security, reduce cost, and optimize the overall performance of the system.

These functions enable to respond to problems from various situations.

Table 1 Basic Functions

| Overview | Function |
|---|--|
| Data protection Functions that ensure data integrity to improve data reliability. It is possible to detect and fix drive failures early. | "Data Block Guard" (page 26) "Disk Drive Patrol" (page 27) "Redundant Copy" (page 28) "Rebuild" (page 29) "Fast Recovery" (page 30) "Copyback/Copybackless" (page 31) "Protection (Shield)" (page 33) "Reverse Cabling" (page 34) |
| Resource utilization (virtualization) Functions that deliver effective resource utilization. | "Thin Provisioning" (page 35) |
| <ul style="list-style-type: none"> • Data capacity expansion Functions that expand or relocate a RAID group or a volume in order to flexibly meet any increases in the amount of data. • Guarantee of performance A function that creates a volume that is striped in multiple RAID groups in order to improve performance. | "RAID Migration" (page 43) "Logical Device Expansion" (page 45) "LUN Concatenation" (page 47) "Wide Striping" (page 49) |
| Security measures (user access management) Functions to prevent information leakage that are caused by a malicious access. | "Account Management" (page 50) "User Authentication" (page 52) "Audit Log" (page 54) |
| Security measures (unauthorized access prevention) Functions that prevent unintentional storage access. | "Host Affinity" (page 55) "iSCSI Security" (page 57) |
| Environmental burden reduction Functions that adjust the operating time and the environment of the installation location in order to reduce power consumption. | "Eco-mode" (page 58) "Power Consumption Visualization" (page 61) |
| Operation management (device monitoring) Function that reduce load on the system administrator, and that improve system stability and increase operating ratio of the system. | "Operation Management Interface" (page 62) "Performance Information Management" (page 63) "Event Notification" (page 65) "Device Time Synchronization" (page 68) |
| Power control Power control functions that are used to link power-on and power-off operations with servers and perform scheduled operations. | "Power Synchronized Unit" (page 69) "Remote Power Operation (Wake On LAN)" (page 70) |
| <ul style="list-style-type: none"> • High-speed backup • Continuous business Data can be duplicated at any point without affecting other operations. | "Backup" (page 72) |
| Performance Tuning Functions that can perform tuning in order to improve performance. | "Striping Size Expansion" (page 80) "Assigned CMs" (page 81) |
| Stable operation For stable operation of server connections, the appropriate response action can be specified for each server. | "Host Response" (page 82) |

| Overview | Function |
|--|--|
| Data relocation A function that migrates data between ETERNUS storage systems. | "Storage Migration" (page 83) |
| Non-disruptive data relocation A function that migrates data between ETERNUS storage systems without stopping the business server. | "Non-disruptive Storage Migration" (page 85) |
| Information linkage (function linkage with servers) Functions that cooperate with a server to improve performance in a virtualized environment. Beneficial effects such as centralized management of the entire storage system and a reduction of the load on servers can be realized. | "Oracle VM Linkage" (page 87) "VMware Linkage" (page 88) "Veeam Storage Integration" (page 90) "Microsoft Linkage" (page 93) "OpenStack Linkage" (page 94) "Logical Volume Manager (LVM)" (page 95) |
| Simple configuration A wizard that simplifies the configuration of Thin Provisioning. | "Smart Setup Wizard" (page 96) |

RAID Functions

This section explains the points to note before configuring a system using the ETERNUS DX.

Supported RAID

The ETERNUS DX supports the following RAID levels.

- RAID0 (striping)
- RAID1 (mirroring)
- RAID1+0 (striping of pairs of drives for mirroring)
- RAID5 (striping with distributed parity)
- RAID5+0 (double striping with distributed parity)
- RAID6 (striping with double distributed parity)
- RAID6-FR (provides the high speed rebuild function, and striping with double distributed parity)

IMPORTANT

Remember that a RAID0 configuration is not redundant. This means that if a RAID0 drive fails, the data will not be recoverable.

This section explains the concepts and purposes (RAID level selection criteria) of the supported RAID levels.

Caution

When Nearline SAS disks that have 6TB or more are used, the available RAID levels are RAID0, RAID1, RAID6, and RAID6-FR.

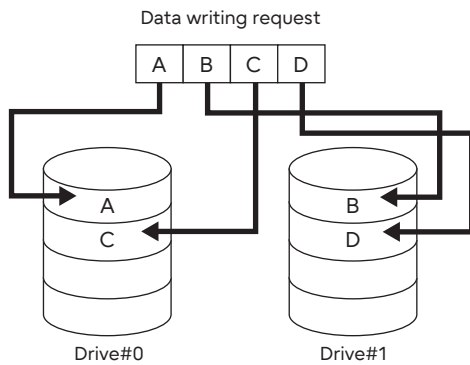
■ RAID Level Concept

A description of each RAID level is shown below.

● RAID0 (Striping)

Data is split in unit of blocks and stored across multiple drives.

Figure 1 RAID0 Concept

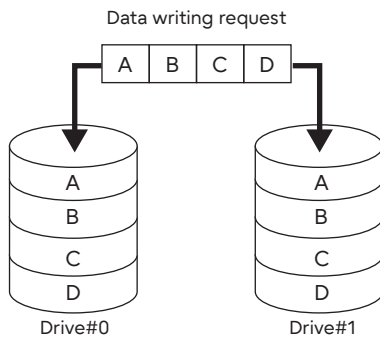


● RAID1 (Mirroring)

The data is stored on two duplicated drives at the same time.

If one drive fails, other drive continues operation.

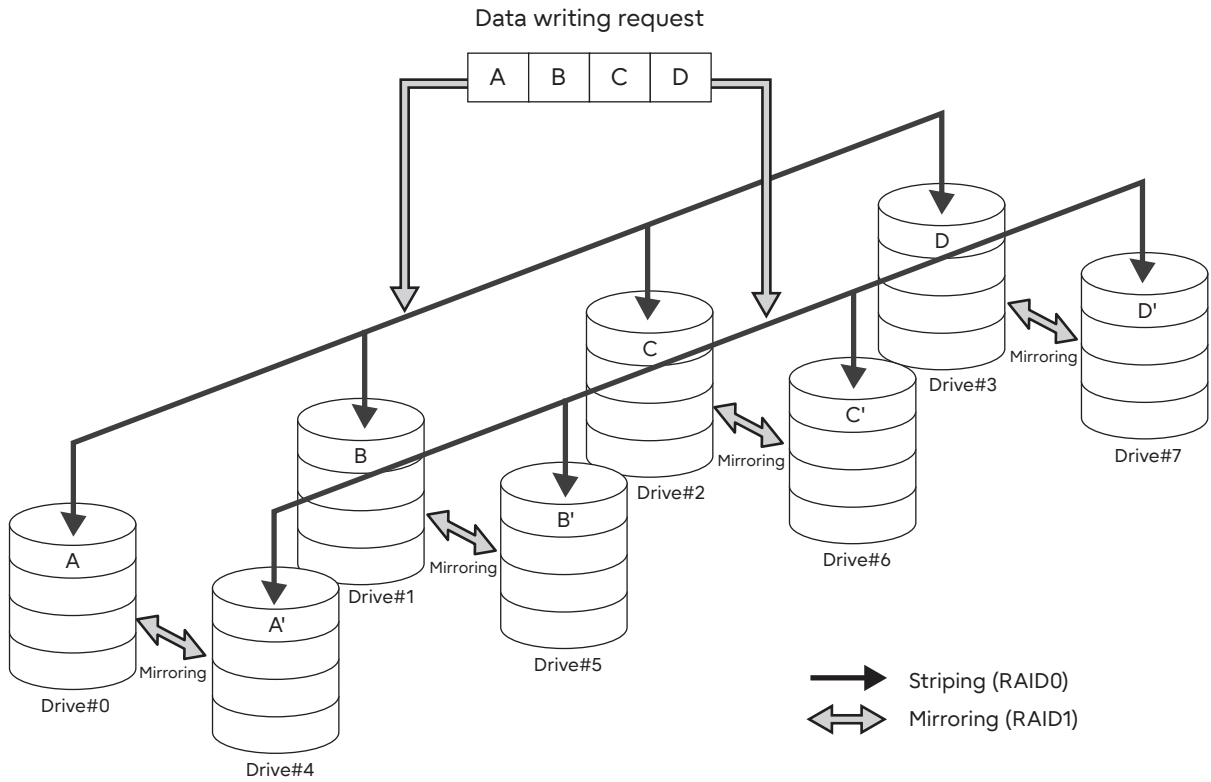
Figure 2 RAID1 Concept



● RAID1+0 (Striping of Pairs of Drives for Mirroring)

RAID1+0 combines the high I/O performance of RAID0 (striping) with the reliability of RAID1 (mirroring).

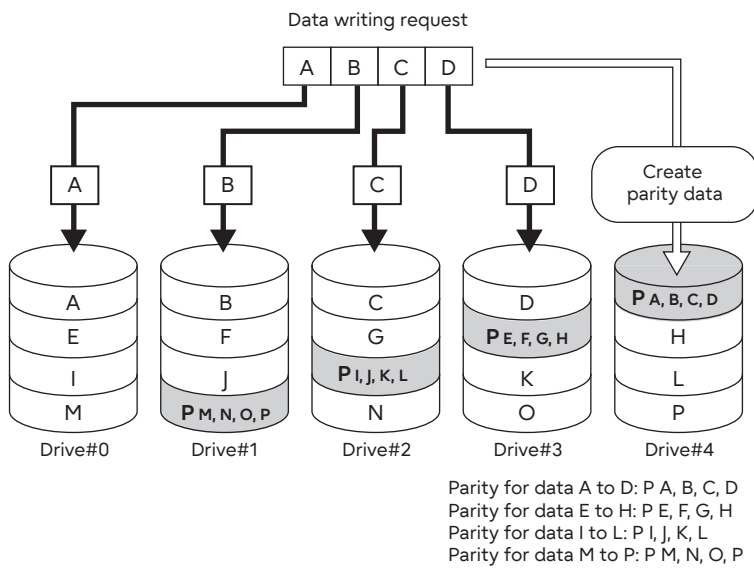
Figure 3 RAID1+0 Concept



● RAID5 (Striping with Distributed Parity)

Data is divided into blocks and allocated across multiple drives together with parity information created from the data in order to ensure the redundancy of the data.

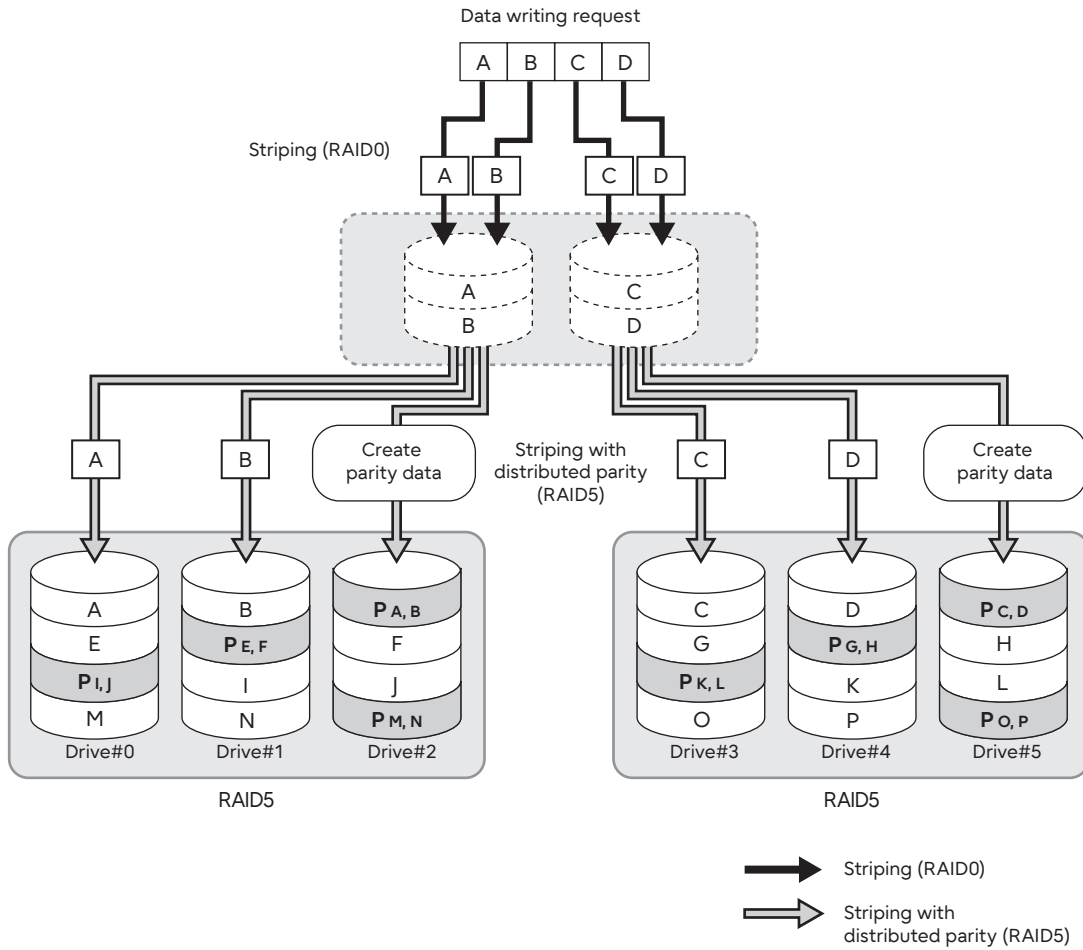
Figure 4 RAID5 Concept



● RAID5+0 (Double Striping with Distributed Parity)

Multiple RAID5 volumes are RAID0 striped. For large capacity configurations, RAID5+0 provides better performance, better reliability, and shorter rebuilding times than RAID5.

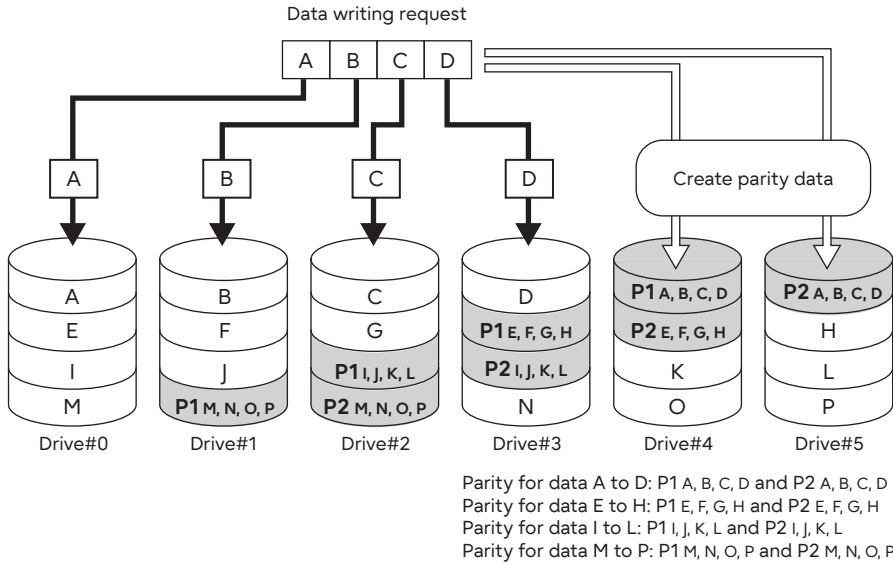
Figure 5 RAID5+0 Concept



● RAID6 (Striping with Double Distributed Parity)

Allocating two different parities on different drives (double parity) makes it possible to recover from up to two drive failures.

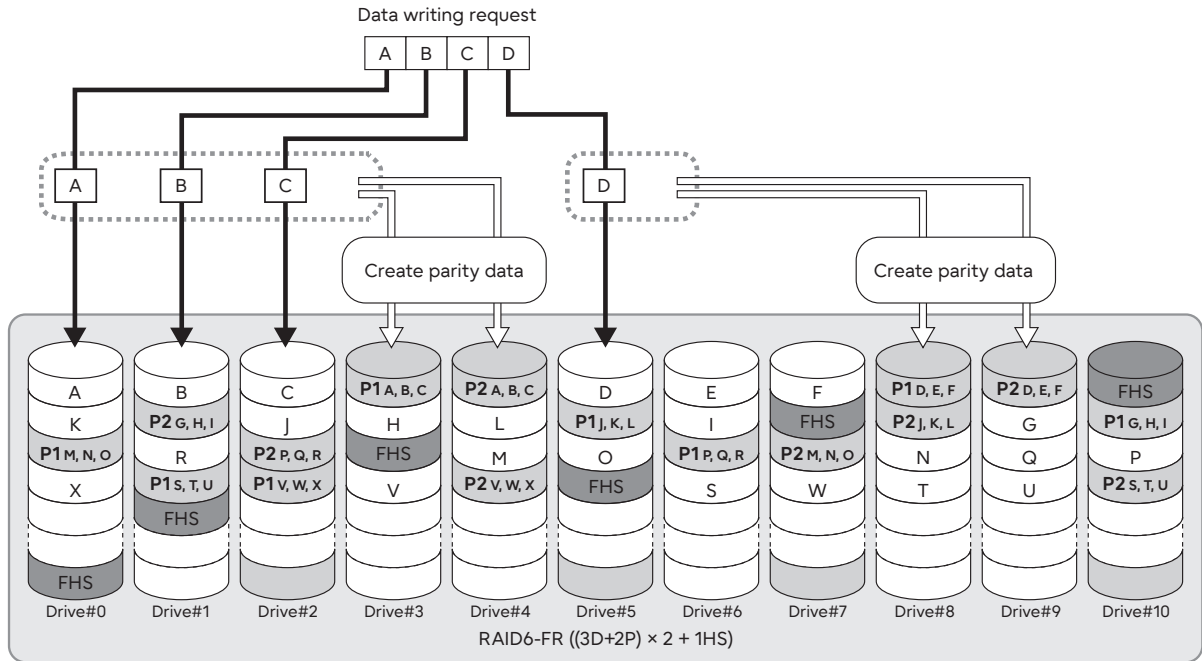
Figure 6 RAID6 Concept



● RAID6-FR (Provides the High Speed Rebuild Function, and Striping with Double Distributed Parity)

Distributing multiple data groups and reserved space equivalent to hot spares to the configuration drives makes it possible to recover from up to two drive failures. RAID6-FR requires less build time than RAID6.

Figure 7 RAID6-FR Concept



Parity for data A, B, C: P1 A, B, C and P2 A, B, C
 Parity for data D, E, F: P1 D, E, F and P2 D, E, F
 Parity for data G, H, I: P1 G, H, I and P2 G, H, I
 Parity for data J, K, L: P1 J, K, L and P2 J, K, L
 Parity for data M, N, O: P1 M, N, O and P2 M, N, O
 Parity for data P, Q, R: P1 P, Q, R and P2 P, Q, R
 Parity for data S, T, U: P1 S, T, U and P2 S, T, U
 Parity for data V, W, X: P1 V, W, X and P2 V, W, X
 :
 Fast recovery Hot Spare: FHS

■ Reliability, Performance, Capacity for Each RAID Level

Table 2 shows the comparison result of reliability, performance, capacity for each RAID level.

Table 2 RAID Level Comparison

| RAID level | Reliability | Performance (*1) | Capacity |
|------------|-------------|------------------|----------|
| RAID0 | × | ◎ | ◎ |
| RAID1 | ○ | ○ | △ |
| RAID1+0 | ○ | ◎ | △ |
| RAID5 | ○ | ○ | ○ |
| RAID5+0 | ○ | ○ | ○ |
| RAID6 | ◎ | ○ | ○ |
| RAID6-FR | ◎ | ○ | ○ |

◎: Very good ○: Good △: Reasonable ×: Poor

*1: Performance may differ according to the number of drives and the processing method from the host.

■ Recommended RAID Level

Select the appropriate RAID level according to the usage.

- Recommended RAID levels are RAID1, RAID1+0, RAID5, RAID5+0, RAID6, and RAID6-FR.
- When importance is placed upon read and write performance, a RAID1+0 configuration is recommended.
- For read only file servers and backup servers, RAID5, RAID5+0, RAID6, or RAID6-FR can also be used for higher efficiency. However, if the drive fails, note that data restoration from parities and rebuilding process may result in a loss in performance.
- For SSDs, a RAID5 configuration or a fault tolerant enhanced RAID6 configuration is recommended because SSDs operate much faster than other types of drive. For large capacity SSDs, using a RAID6-FR configuration, which provides excellent performance for the rebuild process, is recommended.
- Using a RAID6 or RAID6-FR configuration is recommended when Nearline SAS disks that have 6TB or more are used. For details on the RAID levels that can be configured with Nearline SAS disks that have 6TB or more, refer to "Supported RAID" (page 13).

User Capacity (Logical Capacity)

User Capacity for Each RAID Level

The user capacity depends on the capacity of drives that configure a RAID group and the RAID level.

Table 3 shows the formula for calculating the user capacity for each RAID level.

Table 3 Formula for Calculating User Capacity for Each RAID Level

| RAID level | Formula for user capacity computation |
|------------|--|
| RAID0 | Drive capacity × Number of drives |
| RAID1 | Drive capacity × Number of drives ÷ 2 |
| RAID1+0 | Drive capacity × Number of drives ÷ 2 |
| RAID5 | Drive capacity × (Number of drives - 1) |
| RAID5+0 | Drive capacity × (Number of drives - 2) |
| RAID6 | Drive capacity × (Number of drives - 2) |
| RAID6-FR | Drive capacity × (Number of drives - (2 × N) - Number of hot spares) (*1) |

*1: "N" is the number of RAID6 configuration sets. For example, if a RAID6 group is configured with "(3D+2P)×2+1HS", N is "2".

User Capacity of Drives

Table 4 shows the user capacity for each drive.

The supported drives vary between the ETERNUS DX60 S4 and the ETERNUS DX60 S3. For details about drives, refer to "Overview" of the currently used storage systems.

Table 4 User Capacity per Drive

| Product name (*1) | User capacity |
|-----------------------------|---------------|
| 200GB SSD | 186,624MB |
| 400GB SSD | 374,528MB |
| 800GB SSD | 750,080MB |
| 960GB SSD | 914,432MB |
| 1.6TB SSD | 1,501,440MB |
| 1.92TB SSD | 1,830,144MB |
| 3.84TB SSD | 3,661,568MB |
| 300GB SAS disk | 279,040MB |
| 600GB SAS disk | 559,104MB |
| 900GB SAS disk | 839,168MB |
| 1.2TB SAS disk | 1,119,232MB |
| 1.8TB SAS disk | 1,679,360MB |
| 2.4TB SAS disk | 2,239,744MB |
| 1TB Nearline SAS disk | 937,728MB |
| 2TB Nearline SAS disk | 1,866,240MB |
| 4TB Nearline SAS disk | 3,733,504MB |
| 6TB Nearline SAS disk (*2) | 5,601,024MB |
| 8TB Nearline SAS disk (*2) | 7,468,288MB |
| 10TB Nearline SAS disk (*2) | 9,341,696MB |
| 12TB Nearline SAS disk (*2) | 11,210,496MB |
| 14TB Nearline SAS disk (*2) | 13,079,296MB |

***1:** The capacity of the product names for the drives is based on the assumption that 1MB = 1,000² bytes, while the user capacity for each drive is based on the assumption that 1MB = 1,024² bytes. Furthermore, OS file management overhead will reduce the actual usable capacity.

The user capacity is constant regardless of the drive size (2.5"/3.5") or the SSD type (Value SSD and MLC SSD).

***2:** For details on the RAID levels that can be configured with Nearline SAS disks that have 6TB or more, refer to "Supported RAID" (page 13).

RAID Group

This section explains RAID groups.

A RAID group is a group of drives. It is a unit that configures RAID. Multiple RAID groups with the same RAID level or multiple RAID groups with different RAID levels can be set together in the ETERNUS DX. After a RAID group is created, RAID levels can be changed and drives can be added.

Table 5 RAID Group Types and Usage

| Type | Usage | Maximum capacity |
|-----------------------------------|---|--------------------------|
| RAID group | Areas to store normal data. Volumes (Standard, WSV, SDV, SDPV) for work and Advanced Copy can be created in a RAID group. | Approximately 324TB (*1) |
| Thin Provisioning Pool (TPP) (*2) | RAID groups that are used for Thin Provisioning in which the areas are managed as a Thin Provisioning Pool (TPP). Thin Provisioning Volumes (TPVs) can be created in a TPP. | 1,024TB |


*1: This value is for a 14TB Nearline SAS disk RAID6-FR([13D+2P]×2+1HS) configuration.

For details on the number of configuration drives for each RAID level and recommended configurations, refer to [Table 6](#).

*2: For details on the number of configuration drives for each RAID level and recommended configurations, refer to [Table 11](#).

SAS disks and Nearline SAS disks can exist together in the same RAID group. However, from a performance perspective, use the same type of disk (SAS disks or Nearline SAS disks) to configure RAID groups.

Figure 8 Example of a RAID Group



Note

- SAS disks and Nearline SAS disks can be installed together in the same group. Note that SAS disks and Nearline SAS disks cannot be installed with SSDs.
- Use drives that have the same size, capacity, rotational speed, and Advanced Format support to configure RAID groups.
 - If a RAID group is configured with drives that have different capacities, all the drives in the RAID group are recognized as having the same capacity as the drive with the smallest capacity in the RAID group and the rest of the capacity in the drives that have a larger capacity cannot be used.
 - If a RAID group is configured with drives that have different rotational speeds, the performance of all of the drives in the RAID group is reduced to that of the drive with the lowest rotational speed.
 - For details on the RAID levels that can be configured with Nearline SAS disks that have 6TB or more, refer to ["Supported RAID" \(page 13\)](#).

Table 6 shows the recommended number of drives that configure a RAID group.

Table 6 Number of Drives to Configure and the Recommended Number of Drives per RAID Group

| RAID level | Number of configuration drives | Recommended number of drives (*1) |
|------------|--------------------------------|--|
| RAID0 | 2 to 16 | — |
| RAID1 | 2 | 2(1D+1M) |
| RAID1+0 | 4 to 32 | 4(2D+2M), 6(3D+3M), 8(4D+4M), 10(5D+5M) |
| RAID5 | 3 to 16 | 3(2D+1P), 4(3D+1P), 5(4D+1P), 6(5D+1P) |
| RAID5+0 | 6 to 32 | 3(2D+1P) × 2, 4(3D+1P) × 2, 5(4D+1P) × 2, 6(5D+1P) × 2 |
| RAID6 | 5 to 16 | 5(3D+2P), 6(4D+2P), 7(5D+2P) |
| RAID6-FR | 11 to 31 | 17 ((6D+2P) × 2+1HS) |

*1: D = Data, M = Mirror, P = Parity, HS = Hot Spare

Note

- Sequential access performance hardly varies with the number of drives for the RAID group.
- Random access performance tends to be proportional to the number of drives for the RAID group.
- Use of higher capacity drives will increase the time required for the drive rebuild process to complete.
- For RAID5, RAID5+0, and RAID6, ensure that a single RAID group is not being configured with too many drives.
If the number of drives increases, the time to perform data restoration from parities and Rebuild/Copyback when a drive fails also increases.
For details on the recommended number of drives, refer to [Table 6](#).
- For details on the Thin Provisioning function and the RAID configurations that can be registered in Thin Provisioning Pools, refer to ["Storage Capacity Virtualization" \(page 35\)](#).

An assigned CM is allocated to each RAID group. For details, refer to ["Assigned CMs" \(page 81\)](#).
For the installation locations of the drives that configure the RAID group, refer to ["Recommended RAID Group Configurations" \(page 135\)](#).


Volume

This section explains volumes.

Logical drive areas in RAID groups are called volumes.

A volume is the basic RAID unit that can be recognized by the server.

Figure 9 Volume Concept



A volume may be up to 128TB. However, the maximum capacity of volume varies depending on the OS of the server.

The maximum number of volumes that can be created in the ETERNUS DX is 1,024. Volumes can be created until the combined total for each volume type reaches the maximum number of volumes.

A volume can be expanded or moved if required. Multiple volumes can be concatenated and treated as a single volume. For availability of expansion, displacement, and concatenation for each volume, refer to "Target Volumes of Each Function" (page 142).

The types of volumes that are listed in the table below can be created in the ETERNUS DX.

Table 7 Volumes That Can Be Created

| Type | Usage | Maximum capacity |
|--------------------------------|---|--|
| Standard (Open) | A standard volume is used for normal usage, such as file systems and databases. The server recognizes it as a single logical unit. "Standard" is displayed as the type for this volume in ETERNUS Web GUI/ETERNUS CLI and "Open" is displayed in ETERNUS SF software. | 128TB (*1) |
| Snap Data Volume (SDV) | This area is used as the copy destination for SnapOPC/SnapOPC+. There is a SDV for each copy destination. | 24 [MB] + copy source volume capacity × 0.1 [%] (*2) |
| Snap Data Pool Volume (SDPV) | This volume is used to configure the Snap Data Pool (SDP) area. The SDP capacity equals the total capacity of the SDPVs. A volume is supplied from a SDP when the amount of updates exceeds the capacity of the copy destination SDV. | 2TB |
| Thin Provisioning Volume (TPV) | This virtual volume is created in a Thin Provisioning Pool area. TPVs are used as the copy destination for SnapOPC+. | 128TB |
| Wide Striping Volume (WSV) | This volume is created by concatenating distributed areas in from 2 to 48 RAID groups. Processing speed is fast because data access is distributed. | 128TB |
| ODX Buffer volume | An ODX Buffer volume is a dedicated volume that is required to use the Offloaded Data Transfer (ODX) function of Windows Server 2012 or later. It is used to save the source data when data is updated while a copy is being processed. It can be created one per ETERNUS DX. Its volume type is Standard or TPV. | 1TB |

***1:** When multiple volumes are concatenated using the LUN Concatenation function, the maximum capacity is also 128TB.

***2:** The capacity differs depending on the copy source volume capacity.

After a volume is created, formatting automatically starts. A server can access the volume while it is being formatted. Wait for the format to complete if high performance access is required for the volume.


Caution

- In the ETERNUS DX, volumes have different stripe sizes that depend on the RAID level and the stripe depth parameter.
For details about the stripe sizes for each RAID level and the stripe depth parameter values, refer to "ETERNUS Web GUI User's Guide".
Note that the available user capacity can be fully utilized if an exact multiple of the stripe size is set for the volume size. If an exact multiple of the stripe size is not set for the volume size, the capacity is not fully utilized and some areas remain unused.
- When a Thin Provisioning Pool (TPP) is created, a control volume is created for each RAID group that configures the relevant TPP. Therefore, the maximum number of volumes that can be created in the ETERNUS DX decreases by the number of RAID groups that configure a TPP.

Hot Spares

Hot spares are used as spare drives for when drives in a RAID group fail, or when drives are in error status.

Figure 10 Hot Spares



Note

When the RAID level is RAID6-FR, data in a failed drive can be restored to a reserved space in a RAID group even when a drive error occurs because a RAID6-FR RAID group retains a reserved space for a whole drive in the RAID group. If the reserved area is in use and an error occurs in another drive (2nd) in the RAID group, then the hot spare is used as a spare.

■ Types of Hot Spares

The following two types of hot spare are available:

- Global Hot Spare
This is available for any RAID group. When multiple hot spares are installed, the most appropriate drive is automatically selected and incorporated into a RAID group.
- Dedicated Hot Spare
This is only available to the specified RAID group (one RAID group).
The Dedicated Hot Spare cannot be registered in a RAID group that is registered in TPPs.

Note

- Assign "Dedicated Hot Spares" to RAID groups that contain important data, in order to preferentially improve their access to hot spares.
- When an Advanced Format disk is used as the hot spare, the sector format information that can be confirmed from the server may be changed.

■ Number of Installable Hot Spares

Register one hot spare for each drive type.

■ Types of Drives

If a combination of SAS disks, Nearline SAS disks, and SSDs is installed in the ETERNUS DX, each different type of drive requires a corresponding hot spare.

There are two types of rotational speeds for SAS disks; 10,000rpm and 15,000rpm. If a drive error occurs and a hot spare is configured in a RAID group with different rotational speed drives, the performance of all the drives in the RAID group is determined by the drive with the slowest rotational speed. When using SAS disks with different rotational speeds, prepare hot spares that correspond to the different rotational speed drives if required. Even if a RAID group is configured with SAS disks that have different interface speeds, performance is not affected.

The capacity of each hot spare must be equal to the largest capacity of the same-type drives.

■ Selection Criteria

When multiple Global Hot Spares are installed, the following criteria are used to select which hot spare will replace a failed drive:

Table 8 Hot Spare Selection Criteria

| Selection order | Selection criteria |
|-----------------|---|
| 1 | A hot spare with the same type, same capacity, and same rotational speed as the failed drive |
| 2 | A hot spare with the same type and same rotational speed as the failed drive but with a larger capacity (*1) |
| 3 | A hot spare with the same type and same capacity as the failed drive but with a different rotational speed |
| 4 | A hot spare with the same type as the failed drive but with a larger capacity and a different rotational speed (*1) |

***1:** When there are multiple hot spares with a larger capacity than the failed drive, the hot spare with the smallest capacity among them is used first.

Data Protection

Data Block Guard

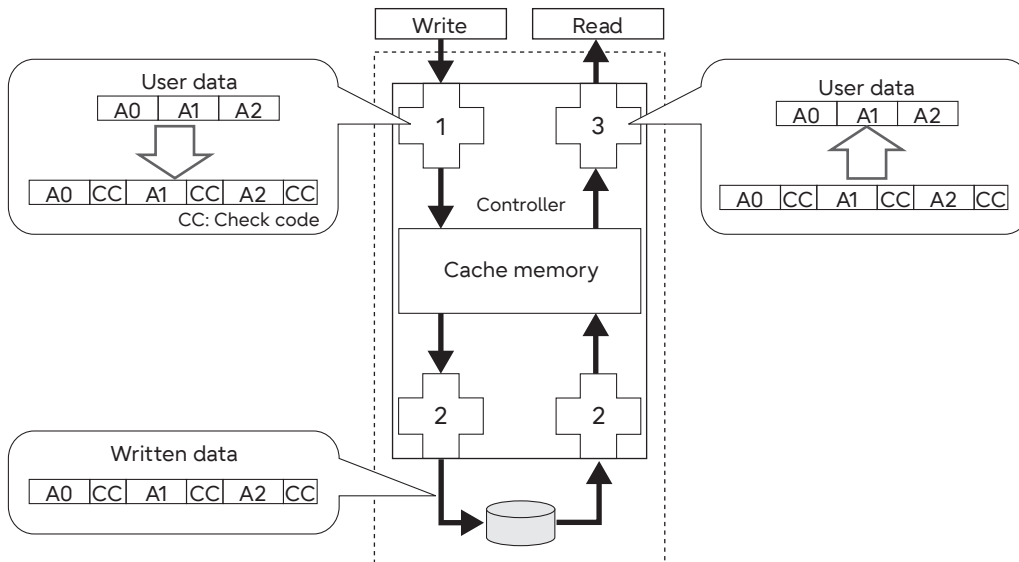
When a write request is issued by a server, the data block guard function adds check codes to all of the data that is to be stored. The data is verified at multiple checkpoints on the transmission paths to ensure data integrity.

When data is written from the server, the Data Block Guard function adds eight bytes check codes to each block (every 512 bytes) of the data and verifies the data at multiple checkpoints to ensure data consistency. This function can detect a data error when data is destroyed or data corruption occurs. When data is read from the server, the check codes are confirmed and then removed, ensuring that data consistency is verified in the whole storage system.

If an error is detected while data is being written to a drive, the data is read again from the data that is duplicated in the cache memory. This data is checked for consistency and then written.

If an error is detected while data is being read from a drive, the data is restored using RAID redundancy.

Figure 11 Data Block Guard



1. The check codes are added
2. The check codes are confirmed
3. The check codes are confirmed and removed

Also, the T10-Data Integrity Field (T10-DIF) function is supported. T10-DIF is a function that adds a check code to data that is to be transferred between the Oracle Linux server and the ETERNUS DX, and ensures data integrity at the SCSI level.

The server generates a check code for the user data in the host bus adapter (HBA), and verifies the check code when reading data in order to ensure data integrity.

The ETERNUS DX double-checks data by using the data block guard function and by using the supported T10-DIF to improve reliability.

Data is protected at the SCSI level on the path to the server. Therefore, data integrity can be ensured even if data is corrupted during a check code reassignment.

By linking the Data Integrity Extensions (DIX) function of Oracle DB, data integrity can be ensured in the entire system including the server.

The T10-DIF function can be used when connecting with HBAs that support T10-DIF with an FC interface.

The T10-DIF function can be enabled or disabled for each volume when the volumes are created. This function cannot be enabled or disabled after a volume has been created.

Caution

- The T10-DIF function can be enabled only in the Standard volume.
- LUN concatenation cannot be performed for volumes where the T10-DIF function is enabled.

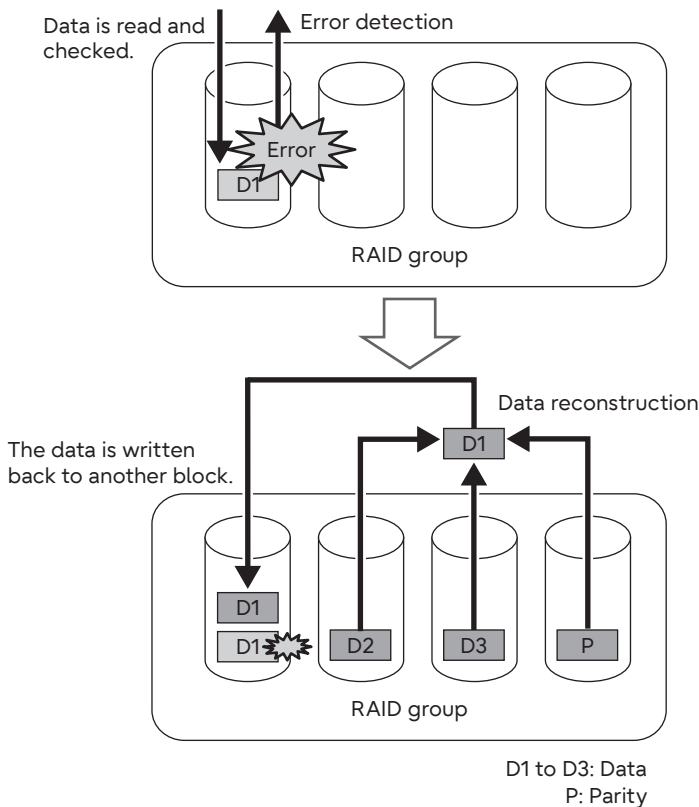
Disk Drive Patrol

In the ETERNUS DX, all of the drives are checked in order to detect drive errors early and to restore drives from errors or disconnect them.

The Disk Drive Patrol function regularly diagnoses and monitors the operational status of all drives that are installed in the ETERNUS DX. Drives are checked (read check) regularly as a background process.

For drive checking, read check is performed sequentially for a part of the data in all the drives. If an error is detected, data is restored using drives in the RAID group and the data is written back to another block of the drive in which the error occurred.

Figure 12 Disk Drive Patrol



Read checking is performed during the diagnosis.

These checks are performed in blocks (default 2MB) for each drive sequentially and are repeated until all the blocks for all the drives have been checked. Patrol checks are performed every second, 24 hours a day (default).

Note

- Drives that are stopped by Eco-mode are checked when the drives start running again.
- The Disk Drive Patrol function is suspended during a drive sanitization. After drive sanitization is completed for all the target drives in the ETERNUS DX, the Disk Drive Patrol function is resumed.

Caution

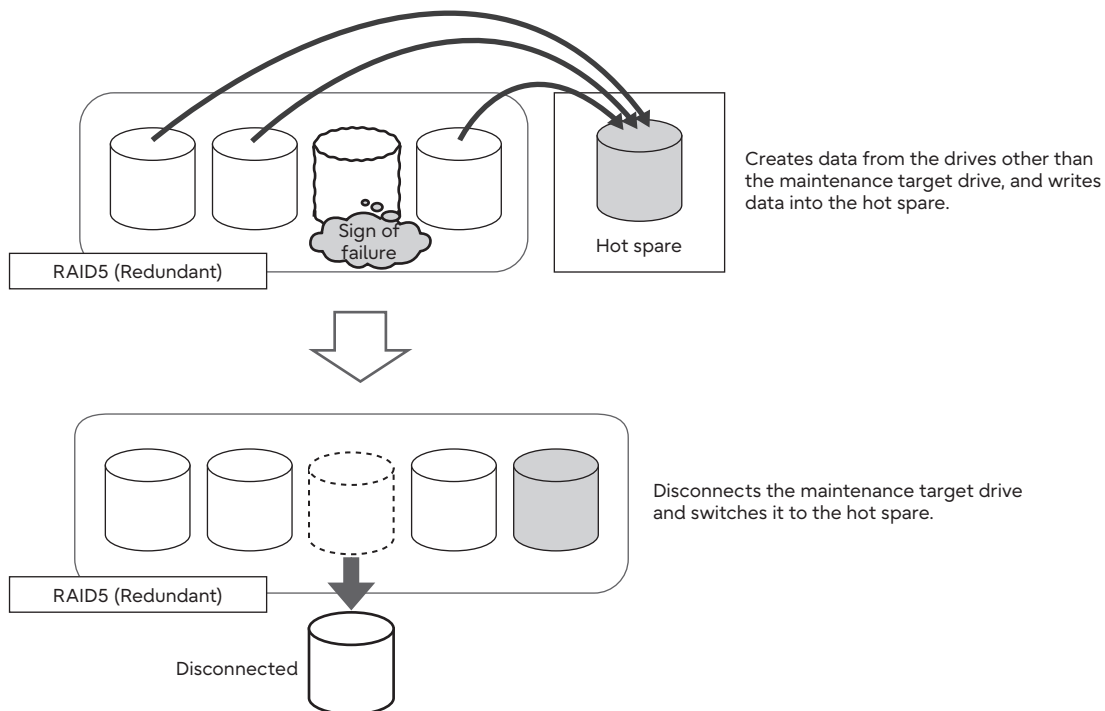
The Maintenance Operation privilege is required to set detailed parameters.

Redundant Copy

Redundant Copy is a function that copies the data of a drive that shows a possible sign of failure to a hot spare.

When the Disk Patrol function decides that preventative maintenance is required for a drive, the data of the maintenance target drive is re-created by the remaining drives and written to the hot spare. The Redundant Copy function enables data to be restored while maintaining data redundancy.

Figure 13 Redundant Copy Function



Note

If a bad sector is detected when a drive is checked, an alternate track is automatically assigned. This drive is not recognized as having a sign of drive failure during this process. However, the drive will be disconnected by the Redundant Copy function if the spare sector is insufficient and the problem cannot be solved by assigning an alternate track.

IMPORTANT

- Redundant Copy speed

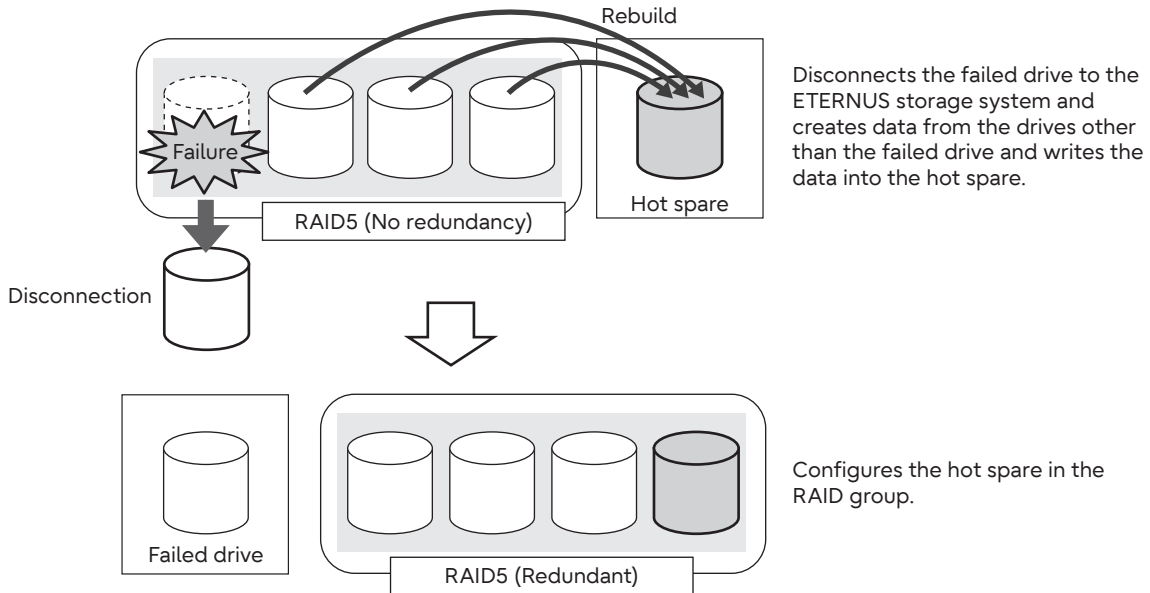
Giving priority to Redundant Copy over host access can be specified. By setting a higher Rebuild priority, the performance of Redundant Copy operations may improve.

However, it should be noted that when the priority is high and a Redundant Copy operation is performed for a RAID group, the performance (throughput) of this RAID group may be reduced.

Rebuild

Rebuild processes recover data in failed drives by using other drives. If a free hot spare is available when one of the RAID group drives has a problem, data of this drive is automatically replicated in the hot spare. This ensures data redundancy.

Figure 14 Rebuild



Note

When no hot spares are registered, rebuilding processes are only performed when a failed drive is replaced or when a hot spare is registered.

IMPORTANT

- Rebuild Speed

Giving priority to rebuilding over host access can be specified. By setting a higher rebuild priority, the performance of rebuild operations may improve.

However, it should be noted that when the priority is high and a rebuild operation is performed for a RAID group, the performance (throughput) of this RAID group may be reduced.

Fast Recovery

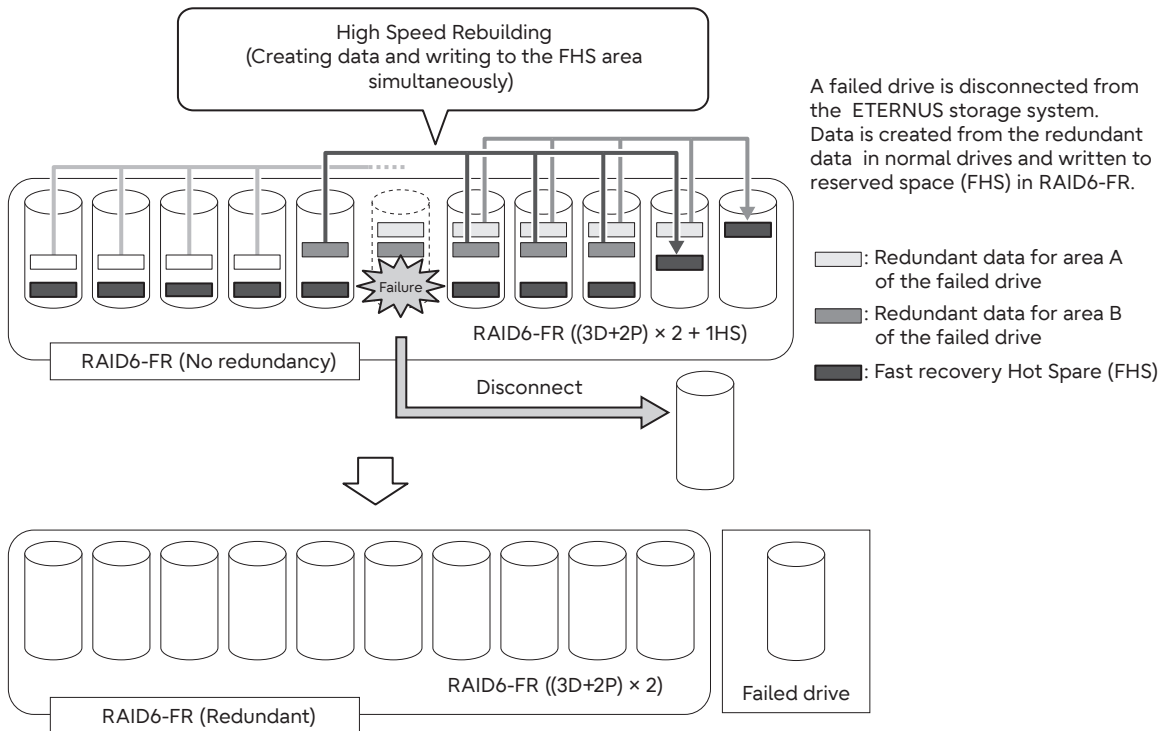
This function recovers data quickly by relocating data in the failed drive to the other remaining drives when a drive error is detected.

For a RAID group that is configured with RAID6-FR, Fast Recovery is performed for the reserved area that is equivalent to hot spares in the RAID group when a drive error occurs.

If a second drive fails when the reserved area is already used by the first failed drive, a normal rebuild (hot spare rebuild in the ETERNUS DX) is performed.

For data in a failed drive, redundant data and reserved space are allocated in different drives according to the area. A fast rebuild can be performed because multiple rebuild processes are performed for different areas simultaneously.

Figure 15 Fast Recovery



Caution

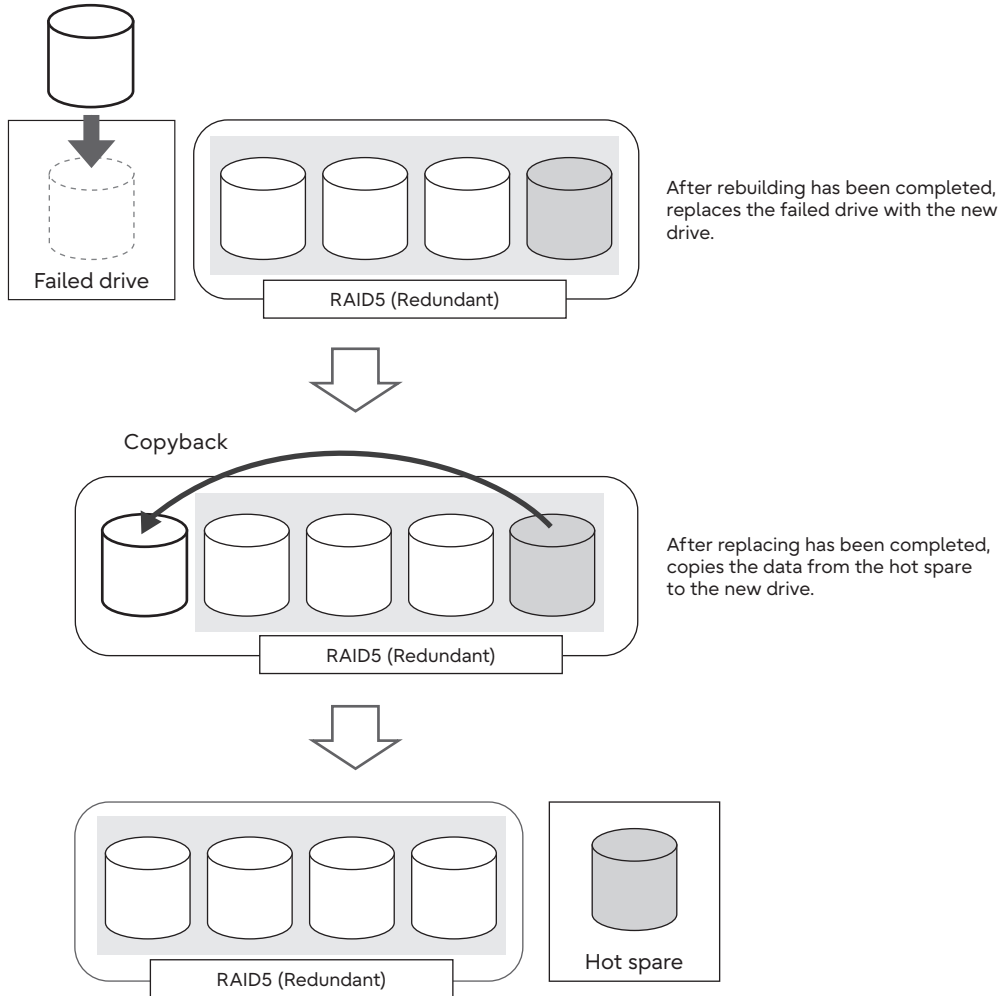
For the Fast Recovery function that is performed when the first drive fails, a copyback is performed after the failed drive is replaced even if the Copybackless function is enabled.

For a normal rebuild process that is performed when the reserved space is already being used and the second drive fails, a copyback is performed according to the settings of the Copybackless function.

Copyback/Copybackless

A Copyback process copies data in a hot spare to the new drive that is used to replace the failed drive.

Figure 16 Copyback



IMPORTANT

- Copyback speed

Giving priority to Copyback over host access can be specified. By setting a higher Rebuild priority, the performance of Copyback operations may improve.

However, it should be noted that when the priority is high and a Copyback operation is performed for a RAID group, the performance (throughput) of this RAID group may be reduced.

If copybackless is enabled, the drives that are registered in the hot spare become part of the RAID group configuration drives after a rebuild or a redundant copy is completed for the hot spare.

The failed drive is disconnected from the RAID group configuration drives and then registered as a hot spare. Copyback is not performed for the data even if the failed drive is replaced by a new drive because the failed drive is used as a hot spare.

A copyback operation is performed when the following conditions for the copybackless target drive (or hot spare) and the failed drive are the same.

- Drive type (SAS disks, Nearline SAS disks, and SSDs)

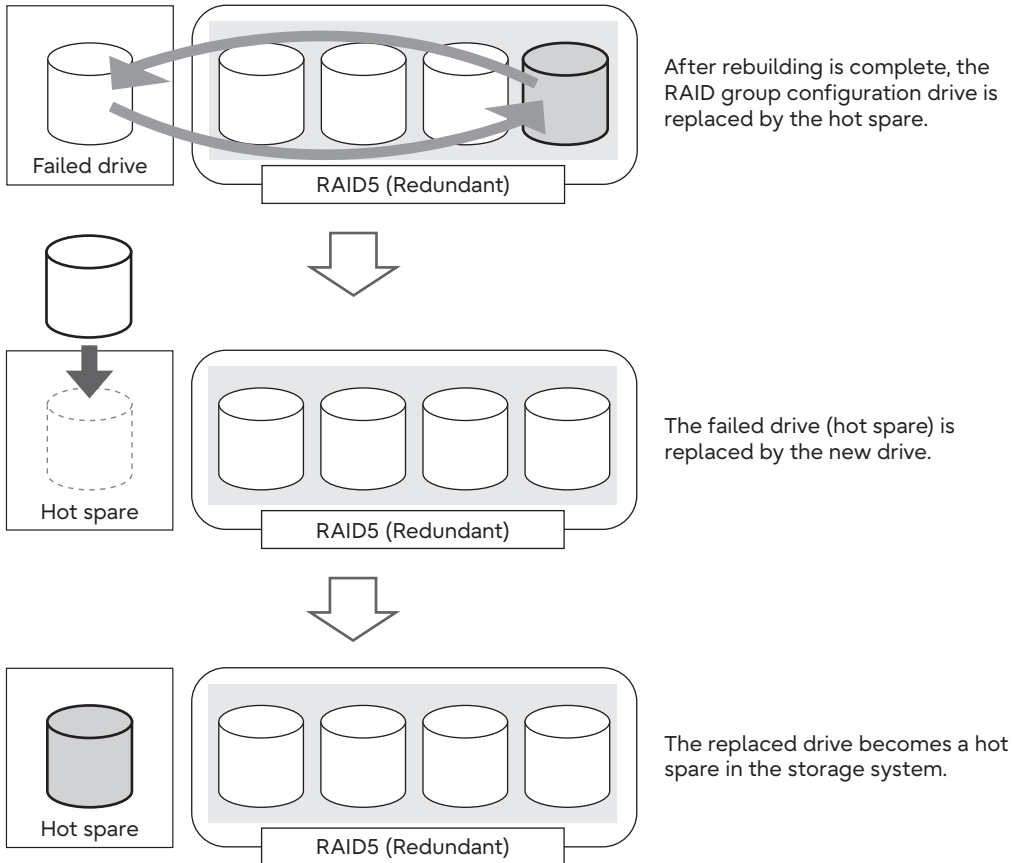
- Size (2.5" and 3.5")
- Capacity
- Rotational speed (15,000rpm, 10,000rpm, and 7,200rpm) (*1)

*1: For SAS disks or Nearline SAS disks only.

If different types of drives have been selected as the hot spare, copyback is performed after replacing the drives even when the Copybackless function is enabled.

The Copybackless function can be enabled or disabled. This function is enabled by default.

Figure 17 Copybackless



Caution

- To set the Copybackless function for each storage system, use the subsystem parameter settings. These settings can be performed with the Storage Management policy or the Maintenance Operation policy. After the settings are changed, the ETERNUS DX does not need to be turned off and on again.
- If the Copybackless function is enabled, the drive that is replaced with the failed drive cannot be installed in the prior RAID group configuration. This should be taken into consideration when enabling or disabling the Copybackless function.

Protection (Shield)

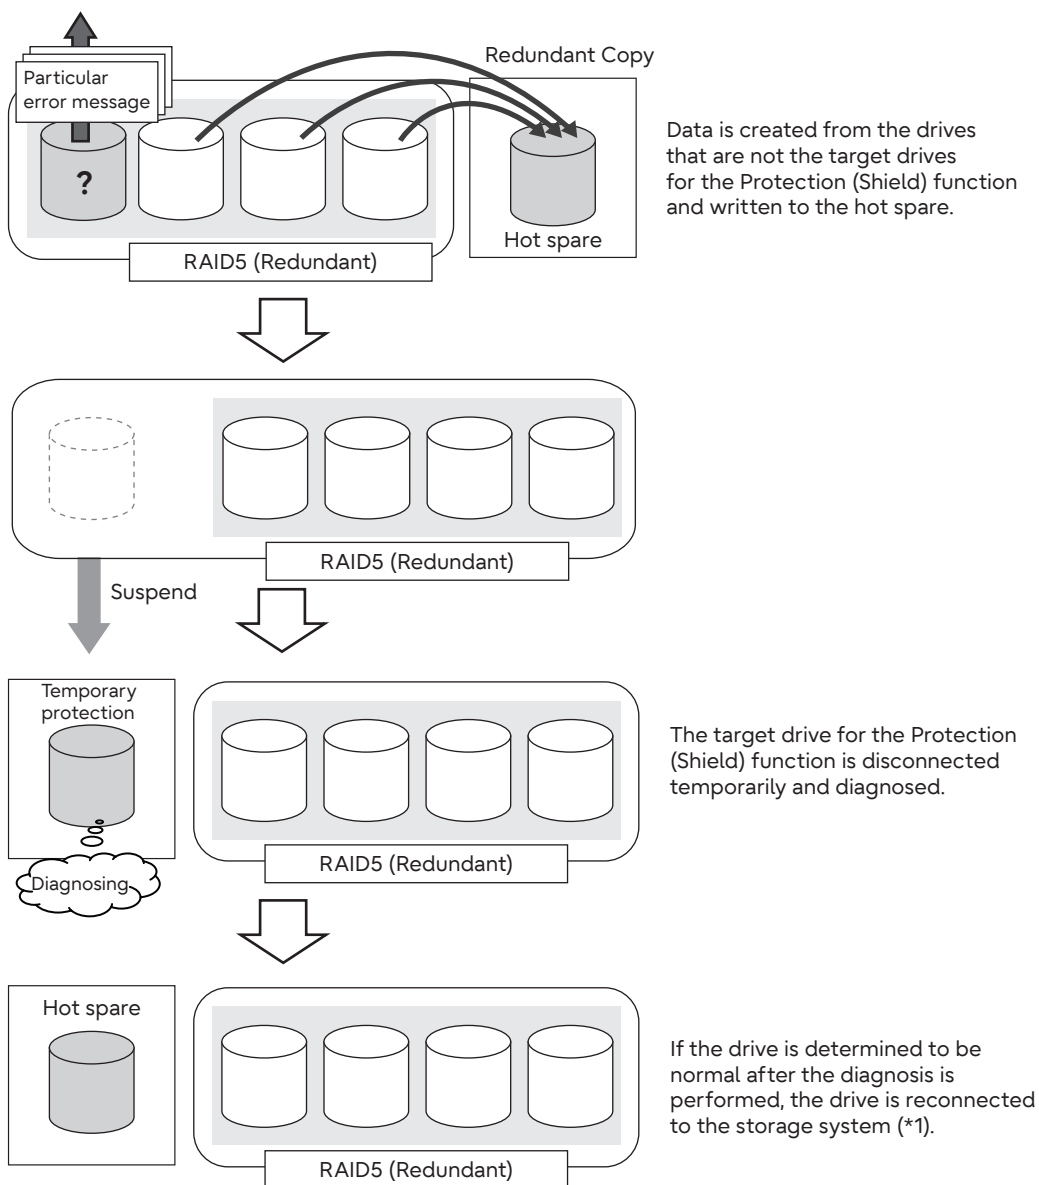
The Protection (Shield) function diagnoses temporary drive errors. A drive can continue to be used if it is determined to be normal. The target drive temporarily changes to diagnosis status when drive errors are detected by the Disk Drive Patrol function or error notifications.

For a drive that configures a RAID group, data is moved to a hot spare by a rebuild or redundant copy before the drive is diagnosed. For a drive that is disconnected from a RAID group, whether the drive has a permanent error or a temporary error is determined. The drive can be used again if it is determined that the drive has only a temporary error.

The target drives of the Protection (Shield) function are all the drives that are registered in RAID groups or registered as hot spares. Note that the Protection (Shield) function is not available for unused drives.

The Protection (Shield) function can be enabled or disabled. This function is enabled by default.

Figure 18 Protection (Shield)



*1: If copybackless is enabled, the drive is used as a hot spare disk. If copybackless is disabled, the drive is used as a RAID group configuration drive and copyback starts. The copybackless setting can be enabled or disabled until the drive is replaced.

Caution

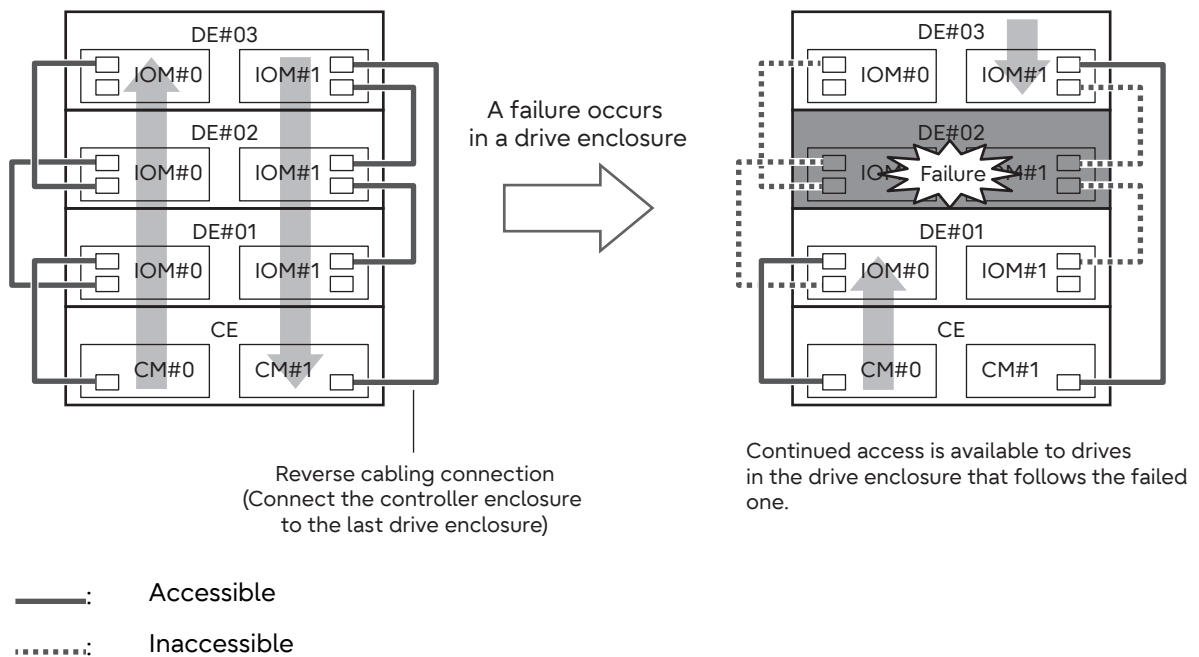
- The target drives are deactivated and then reactivated during temporary drive protection. Even though a system status error may be displayed during this period, this phenomenon is only temporary. The status returns to normal after the diagnosis is complete.
The following phenomenon may occur during temporary drive protection.
 - The Fault LEDs (amber) on the operation panel and the drive turn on
 - An error status is displayed by the ETERNUS Web GUI and the ETERNUS CLI
 - Error or Warning is displayed as the system status
 - Error, Warning, or Maintenance is displayed as the system status
- Target drives of the Protection (Shield) function only need to be replaced when drive reactivation fails.
If drive reactivation fails, a drive failure error is notified as an event notification message (such as SNMP/REMCS). When drive reactivation is successful, an error message is not notified. To notify this message, use the event notification settings.
- To set the Protection (Shield) function for each storage system, use the subsystem parameter settings. The maintenance operation privilege is required to perform this setting.
After the settings are changed, the ETERNUS DX does not need to be turned off and on again.

Reverse Cabling

Because the ETERNUS DX uses reverse cabling connections for data transfer paths between controllers and drives, continued access is ensured even if a failure occurs in a drive enclosure.

If a drive enclosure fails for any reason, access to drives that are connected after the failed drive can be maintained because normal access paths are secured by using reverse cabling.

Figure 19 Reverse Cabling



Operations Optimization (Virtualization)

IMPORTANT

A single controller configuration differs from a dual controller configuration in the following ways:

- The Thin Provisioning function cannot be used.

Thin Provisioning

The Thin Provisioning function has the following features:

- Storage Capacity Virtualization

The physical storage capacity can be reduced by allocating the virtual drives to a server, which allows efficient use of the storage capacity. The volumes more than the capacity of all the installed drives can be allocated by setting the capacity required for virtual volumes in the future.

- TPV Balancing

I/O access to the virtual volume can be distributed among the RAID groups in a pool, by relocating and balancing the physical allocation status of the virtual volume.

- TPV Capacity Optimization (Zero Reclamation)

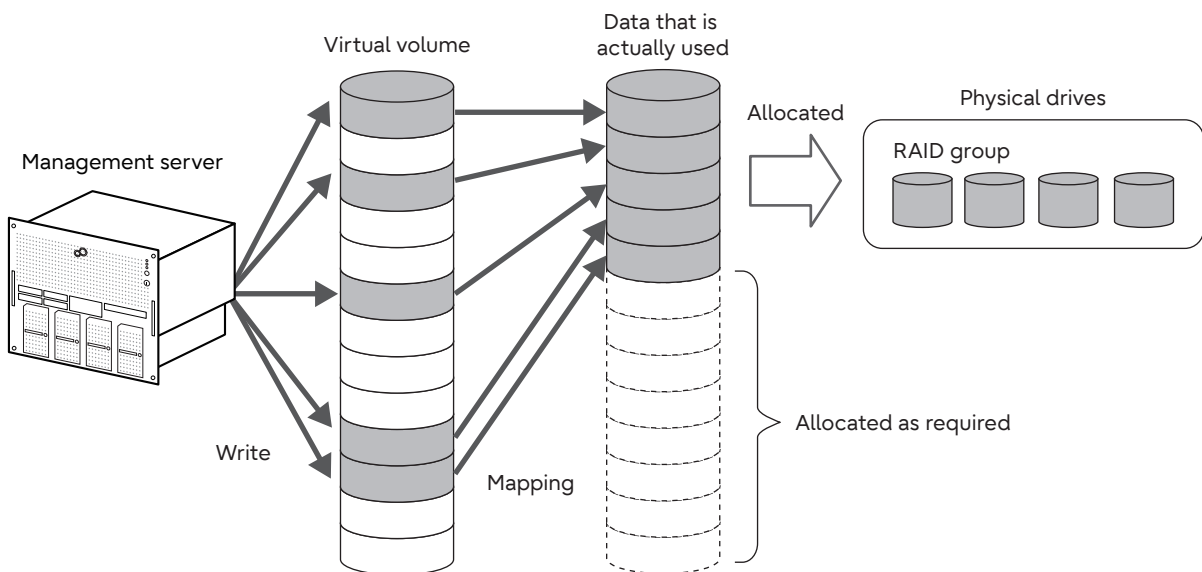
Data in physically allocated areas are checked in blocks and unnecessary areas (areas where 0 is allocated to all of the data in each block) are released to unallocated areas.

Storage Capacity Virtualization

Thin Provisioning improves the usability of the drives by managing the physical drives in a pool, and sharing the unused capacity among the virtual volumes in the pool. The volume capacity that is seen from the server is virtualized to allow the server to recognize a larger capacity than the physical volume capacity. Because a large capacity virtual volume can be defined, the drives can be used in a more efficient and flexible manner.

Initial cost can be reduced because less drive capacity is required even if the capacity requirements cannot be estimated. The power consumption requirements can also be reduced because a fewer number of drives are installed.

Figure 20 Storage Capacity Virtualization



In the Thin Provisioning function, the RAID group, which is configured with multiple drives, is managed as a Thin Provisioning Pool (TPP). When a Write request is issued, a physical area is allocated to the virtual volume. The free space in the TPP is shared among the virtual volumes which belong to the TPP, and a virtual volume, which is larger than the drive capacity in the ETERNUS DX, can be created. A virtual volume to be created in a TPP is referred to as a Thin Provisioning Volume (TPV).

- Thin Provisioning Pool (TPP)

A TPP is a physical drive pool which is configured with one or more RAID groups. TPP capacity can be expanded in the units of RAID groups. Add RAID groups with the same specifications (RAID level, drive type, and number of member drives) as those of the existing RAID groups.

The following table shows the maximum number and the maximum capacity of TPPs that can be registered in the ETERNUS DX.

Table 9 TPP Maximum Number and Capacity

| Item | ETERNUS DX60 S4/DX60 S3 |
|------------------------|-------------------------|
| Number of pools (max.) | 48 |
| Pool capacity (max.) | 1,024TB |

The following table shows the TPP chunk size that is applied when TPPs are created.

Table 10 Chunk Size According to the Configured TPP Capacity

| Setting value of the maximum pool capacity | Chunk size (*1) |
|--|-----------------|
| Up to 128TB | 21MB |
| Up to 256TB | 42MB |
| Up to 512TB | 84MB |
| Up to 1,024TB | 168MB |

*1: Chunk size is for delimiting data. The chunk size is automatically set according to the maximum pool capacity.

The following table shows the RAID configurations that can be registered in a TPP.

Table 11 Levels and Configurations for a RAID Group That Can Be Registered in a TPP

| RAID level | Number of configurable drives | Recommended configurations |
|------------|--|----------------------------|
| RAID0 | 4 (4D) | — |
| RAID1 | 2 (1D+1M) | 2 (1D+1M) |
| RAID1+0 | 4 (2D+2M), 8 (4D+4M), 16 (8D+8M), 24 (12D+12M) | 8 (4D+4M) |
| RAID5 | 4 (3D+1P), 5 (4D+1P), 7 (6D+1P), 8 (7D+1P), 9 (8D+1P), 13 (12D+1P) | 4 (3D+1P), 8 (7D+1P) |
| RAID6 | 6 (4D+2P), 8 (6D+2P), 9 (7D+2P), 10 (8D+2P) | 8 (6D+2P) |
| RAID6-FR | 13 ((4D+2P) ×2+1HS), 17 ((6D+2P) ×2+1HS), 31 ((8D+2P) ×3+1HS), 31 ((4D+2P) ×5+1HS) | 17 ((6D+2P) ×2+1HS) |

- Thin Provisioning Volume (TPV)

The maximum capacity of a TPV is 128TB. Note that the total TPV capacity must be smaller than the maximum capacity of the TPP.

When creating a TPV, the Allocation method can be selected.

- Thin

When data is written from the host to a TPV, a physical area is allocated to the created virtual volume. The capacity size (chunk size) that is applied is the same value as the chunk size of the TPP where the TPV is created. The physical storage capacity can be reduced by allocating a virtualized storage capacity.

- Thick

When creating a volume, the physical area is allocated to the entire volume area. This can be used for volumes in the system area to prevent a system stoppage due to a pool capacity shortage during operations.

In general, selecting "Thin" is recommended. The Allocation method can be changed after a TPV is created.

Perform a TPV capacity optimization if "Thick" has changed to "Thin". By optimizing the capacity, the area that was allocated to a TPV is released and the TPV becomes usable. If a TPV capacity optimization is not performed, the usage of the TPV does not change even after the Allocation method is changed.

The capacity of a TPV can be expanded after it is created.

For details on the number of TPVs that can be created, refer to ["Volume" \(page 22\)](#).

● **Threshold Monitoring of Used Capacity**

When the used capacity of a TPP reaches the threshold, a notification is sent to the notification destination, (SNMP Trap, e-mail, or Syslog) specified using the [Setup Event Notification] function. There are two types of thresholds: "Attention" and "Warning". A different value can be specified for each threshold type.

Also, ETERNUS SF Storage Cruiser can be used to monitor the used capacity.

• TPP Thresholds

There are two TPP usage thresholds: Attention and Warning.

Table 12 TPP Thresholds

| Threshold | Selectable range | Default | Setting conditions |
|-----------|------------------|---------|--|
| Attention | 5 (%) to 80 (%) | 75 (%) | Attention threshold ≤ Warning threshold The "Attention" threshold can be omitted. |
| Warning | 5 (%) to 99 (%) | 90 (%) | |

• TPV Thresholds

There is only one TPV usage threshold: Attention. When the physically allocated capacity of a TPV reaches the threshold, a response is sent to a host via a sense. The threshold is determined by the ratio of free space in the TPP and the unallocated TPV capacity.

Table 13 TPV Thresholds

| Threshold | Selectable range | Default |
|-----------|------------------|---------|
| Attention | 1 (%) to 100 (%) | 80 (%) |

 **Caution**

- Use of TPVs is also not recommended when the OS writes meta information to the whole LUN during file system creation.
- TPVs should be backed up of files as sets of their component files. While backing up a whole TPV is not difficult, unallocated areas will also be backed up as dummy data. If the TPV then needs to be restored from the backup, the dummy data is also "restored". This requires allocation of the physical drive area for the entire TPV capacity, which negates the effects of thin provisioning.
- For advanced performance tuning, use standard RAID groups.
- Refer to the applicable OS and file system documentation before dynamically expanding the volume capacity because expanded volumes may not be recognized by some types and versions of server-side platforms (OSs).
- If a TPP includes one or more RAID groups that are configured with Advanced Format drives, all TPVs created in the relevant TPP are treated as Advanced Format volumes. In this case, the write performance may be reduced when accessing the relevant TPV from an OS or an application that does not support Advanced Format.

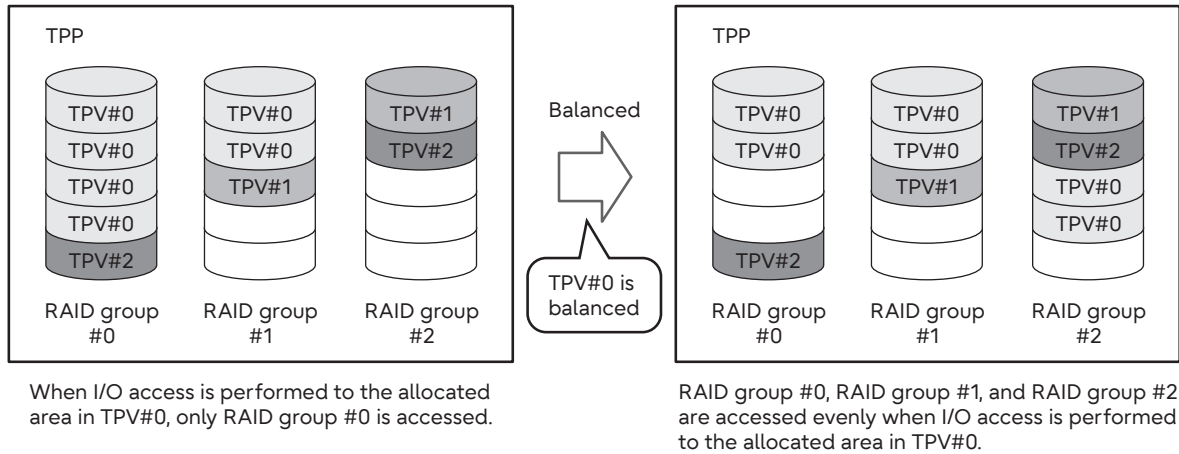
TPV Balancing

A drive is allocated when a write is issued to a virtual volume (TPV). Depending on the order and the frequency of writes, more drives in a specific RAID group may be allocated disproportionately. Also, the physical capacity is unevenly allocated among the newly added RAID group and the existing RAID groups when physical drives are added to expand the capacity.

Balancing of TPVs can disperse the I/O access to virtual volumes among the RAID groups in the Thin Provisioning Pool (TPP).


- When allocating disproportionate TPV physical capacity evenly

Figure 21 TPV Balancing (When Allocating Disproportionate TPV Physical Capacity Evenly)



- When distributing host accesses evenly after TPP expansion (after drives are added)

Figure 22 TPV Balancing (When Distributing Host Accesses Evenly after TPP Expansion)



Balance Thin Provisioning Volume is a function that evenly relocates the physically allocated capacity of TPVs among the RAID groups that configure the TPP.

Balancing TPV allocation can be performed for TPVs in the same TPP. TPV balancing cannot be performed at the same time as RAID Migration to a different TPP for which the target TPV does not belong.

When a write is issued to a virtual volume, a drive is allocated. When data is written to multiple TPVs in the TPP, physical areas are allocated by rotating the RAID groups that configure the TPP in the order that the TPVs were accessed. When using this method, depending on the write order or frequency, TPVs may be allocated unevenly to a specific RAID group. In addition, when the capacity of a TPP is expanded, the physical capacity is unevenly allocated among the newly added RAID group and the existing RAID groups.

● Balancing Level

The TPV balance status is displayed by three levels; "High", "Middle", and "Low". "High" indicates that the physical capacity of TPV is allocated evenly in the RAID groups registered in the TPP. "Low" indicates that the physical capacity is allocated unequally to a specific RAID group in the TPP.

TPV balancing may not be available when other functions are being used in the device or the target volume.

Refer to "[Combinations of Functions That Are Available for Simultaneous Executions](#)" (page 144) for details on the functions that can be executed simultaneously, the number of the process that can be processed simultaneously, and the capacity that can be processed concurrently.

Caution

- When a TPP has RAID groups unavailable for the balancing due to lack of free space, etc., the physical allocation capacity is balanced among the remaining RAID groups within the TPP. In this case, the balancing level after the balancing is completed may not be "High".
- By performing the TPV balancing, areas for working volumes (the migration destination TPVs with the same capacity as the migration source) are secured for the TPP to which the TPVs belong. If this causes the total logical capacity of the TPVs in all the TPPs that include these working volumes to exceed the maximum pool capacity, a TPV balancing cannot be performed.
In addition, this may cause a temporary alarm state ("Caution" or "Warning", which indicates that the threshold has been exceeded) in the TPP during a balancing execution. This alarm state is removed once balancing completes successfully.
- While TPV balancing is being performed, the balancing level may become lower than before balancing was performed if the capacity of the TPP to which the TPVs belong is expanded.

TPV Capacity Optimization

TPV capacity optimization can increase the unallocated areas in a pool (TPP) by changing the physical areas where 0 is allocated for all of the data to unallocated areas. This improves functional efficiency.

Once an area is physically allocated to a TPV, the area is never automatically released.

If operations are performed when all of the areas are physically allocated, the used areas that are recognized by a server and the areas that are actually allocated might have different sizes.

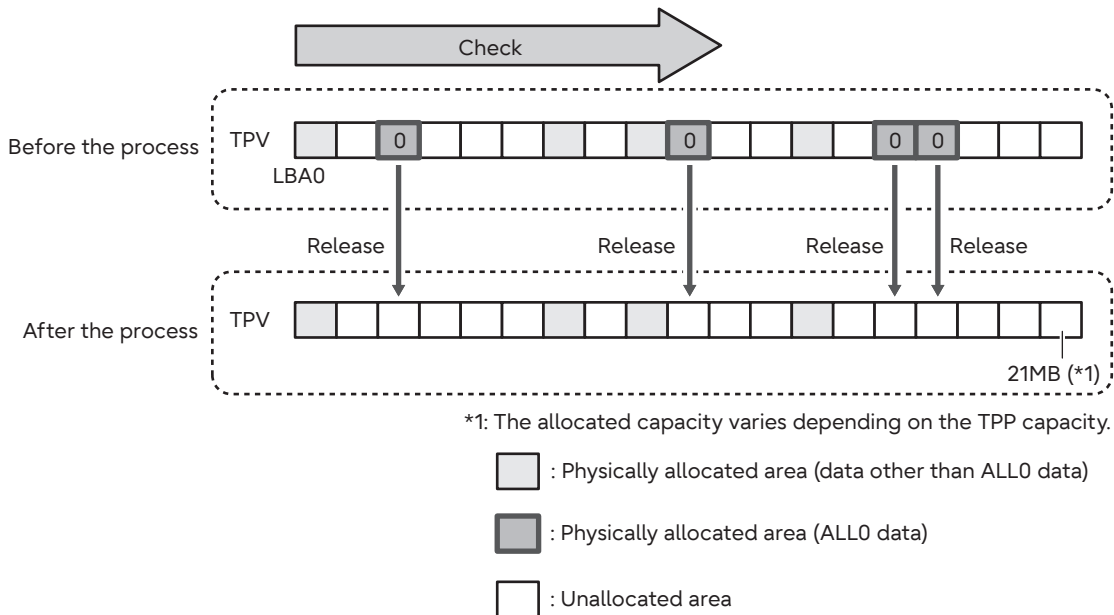
The following operations are examples of operations that create allocated physical areas with sequential data to which only 0 is allocated:

- Restoration of data for RAW image backup
- RAID Migration from Standard volumes to TPVs
- Creation of a file system in which writing is performed to the entire area

The TPV capacity optimization function belongs to Thin Provisioning. This function can be started after a target TPV is selected via ETERNUS Web GUI or ETERNUS CLI. This function is also available when the RAID Migration destination is a TPP.

TPV capacity optimization reads and checks the data in each allocated area for the Thin Provisioning function. This function releases the allocated physical areas to unallocated areas if data that contains all zeros is detected.

Figure 23 TPV Capacity Optimization



TPV capacity optimization may not be available when other functions are being used in the device or the target volume.

For details on the functions that can be executed simultaneously, refer to "[Combinations of Functions That Are Available for Simultaneous Executions](#)" (page 144).

Note

For details on the procedure of TPV capacity optimization, refer to "ETERNUS Web GUI User's Guide".

Optimization of Volume Configurations

The ETERNUS DX allows for the expansion of volumes and RAID group capacities, migration among RAID groups, and changing of RAID levels according to changes in the operation load and performance requirements. There are several expansion functions.

Table 14 Optimization of Volume Configurations

| Function/usage | Volume expansion | RAID group expansion | Migration among RAID groups | Changing the RAID level | Striping for RAID groups |
|--------------------------|---|----------------------|-----------------------------|---|--------------------------|
| RAID Migration | ○ (Adding capacity during migration) (*1) | × | ○ | ○ | × |
| Logical Device Expansion | × | ○ | × | ○ (Adding drives to existing RAID groups) | × |
| LUN Concatenation | ○ (Concatenating free spaces) | × | × | × | × |
| Wide Striping | × | × | × | × | ○ |

○: Possible, ×: Not possible

*1: For TPVs, the capacity cannot be expanded during a migration.

● Expansion of Volume Capacity

- RAID Migration (with increased migration destination capacity)

When volume capacity is insufficient, a volume can be moved to a RAID group that has enough free space. This function is recommended for use when the desired free space is available in the destination.

- LUN Concatenation

Adds areas of free space to an existing volume to expand its capacity. This uses free space from a RAID group to efficiently expand the volume.

● Expansion of RAID Group Capacity

- Logical Device Expansion

Adds new drives to an existing RAID group to expand the RAID group capacity. This is used to expand the existing RAID group capacity instead of adding a new RAID group to add the volumes.

● Migration among RAID Groups

- RAID Migration

The performance of the current RAID groups may not be satisfactory due to conflicting volumes after performance requirements have been changed. Use RAID Migration to improve the performance by redistributing the volumes amongst multiple RAID groups.

● Changing the RAID Level

- RAID Migration (to a RAID group with a different RAID level)

Migrating to a RAID group with a different RAID level changes the RAID level of volumes. This is used to convert a given volume to a different RAID level.

- Logical Device Expansion (and changing RAID levels when adding the new drives)
The RAID level for RAID groups can be changed. Adding drives while changing is also available. This is used to convert the RAID level of all the volumes belonging to a given RAID group.

- **Striping for Multiple RAID Groups**

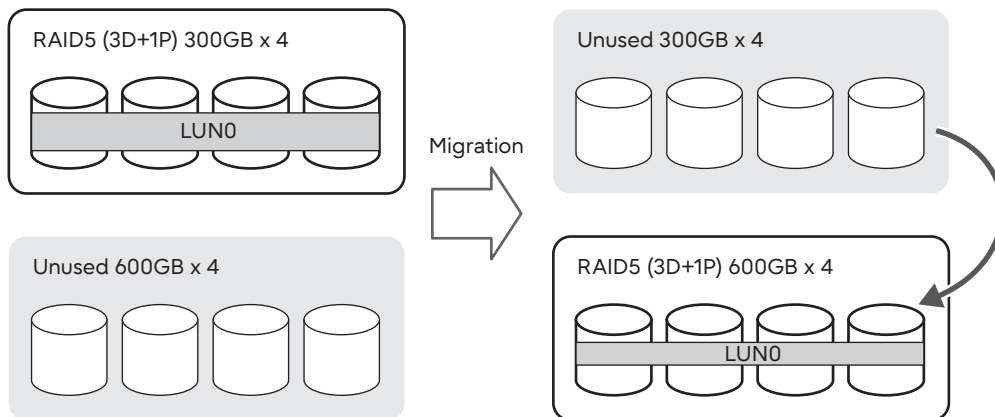
- Wide Striping
Distributing a single volume to multiple RAID groups makes I/O access from the server more efficient and improves the performance.

RAID Migration

RAID Migration is a function that moves a volume to a different RAID group with the data integrity being guaranteed. This allows easy redistribution of volumes among RAID groups in response to customer needs. RAID Migration can be carried out while the system is running, and may also be used to switch data to a different RAID level changing from RAID5 to RAID1+0, for example.

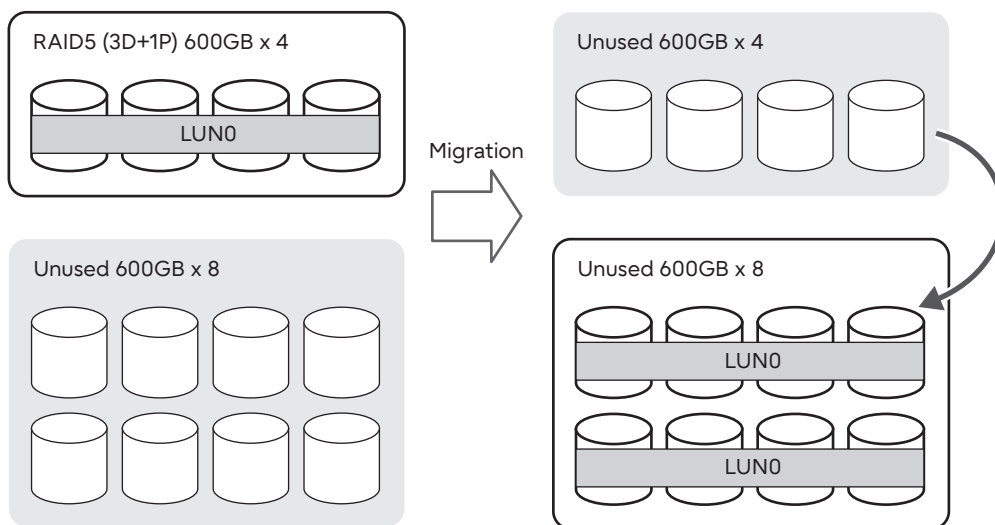
- Volumes moved from a 300GB drive configuration to a 600GB drive configuration

Figure 24 RAID Migration (When Data Is Migrated to a High Capacity Drive)



- Volumes moved to a different RAID level (RAID5 → RAID1+0)

Figure 25 RAID Migration (When a Volume Is Moved to a Different RAID Level)



The volume number (LUN) does not change before and after the migration. The host can access the volume without being affected by the volume number.

The following changes can be performed by RAID migration.

- Changing the volume type

A volume is changed to the appropriate type for the migration destination RAID groups or pools (TPP).

- Changing the number of concatenations and the Wide Stripe Size (for WSV)

The following processes can also be specified.

- Capacity expansion

When migration between RAID groups is performed, capacity expansion can also be performed at the same time. However, the capacity cannot be expanded for TPVs.

- TPV capacity optimization

When the migration destination is a pool (TPP), TPV capacity optimization after the migration can be set.

For details on the features of the TPV capacity optimization, refer to ["TPV Capacity Optimization" \(page 41\)](#).

Specify unused areas in the migration destination (RAID group or pool) with a capacity larger than the migration source volume.

RAID migration may not be available when other functions are being used in the ETERNUS DX or the target volume.

Refer to ["Combinations of Functions That Are Available for Simultaneous Executions" \(page 144\)](#) for details on the functions that can be executed simultaneously, the number of the process that can be processed simultaneously, and the capacity that can be processed concurrently.

 **Caution**

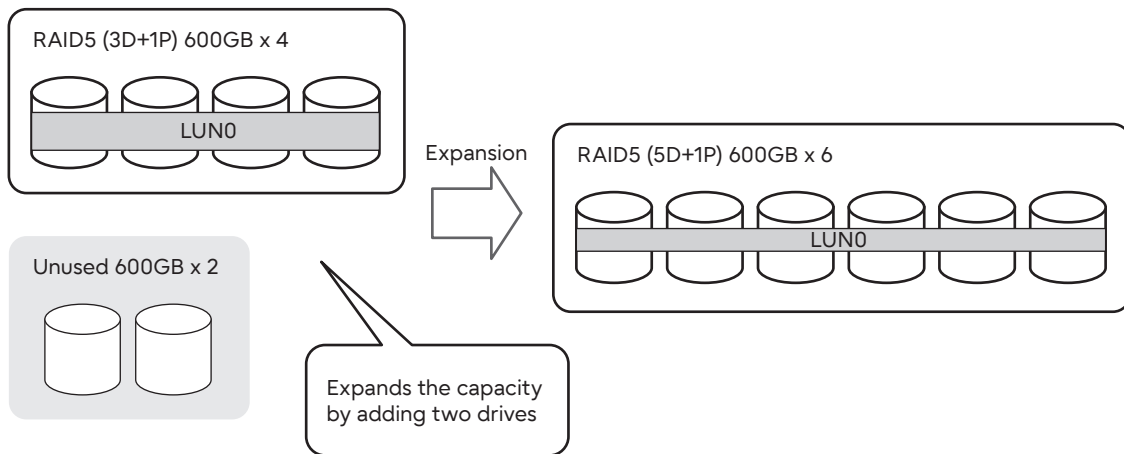
During RAID Migration, the access performance for the RAID groups that are specified as the RAID Migration source and RAID Migration destination may be reduced.

Logical Device Expansion

Logical Device Expansion (LDE) allows the capacity of an existing RAID group to be dynamically expanded by changing of the RAID level or the drive configuration of the RAID group. When this function is performed, drives can be also added at the same time. By using this LDE function to expand the capacity of an existing RAID group, a new volume can be added without having to add new RAID groups.

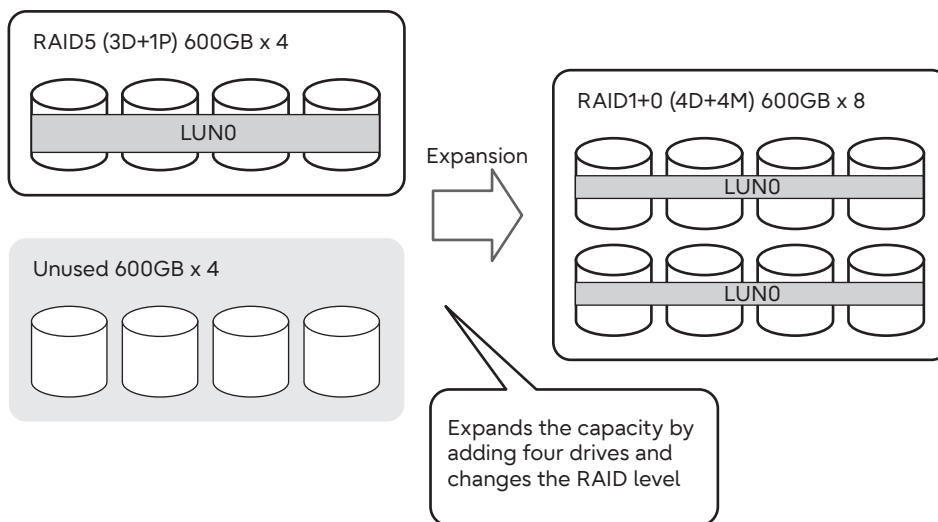
- Expand the RAID group capacity (from RAID5(3D+1P) → RAID5(5D+1P))

Figure 26 Logical Device Expansion (When Expanding the RAID Group Capacity)



- Change the RAID levels (from RAID5(3D+1P) → RAID1+0(4D+4M))

Figure 27 Logical Device Expansion (When Changing the RAID Level)



LDE works in terms of RAID group units. If a target RAID group contains multiple volumes, all of the data in the volumes is automatically redistributed when LDE is performed. Note that LDE cannot be performed if it causes the number of data drives to be reduced in the RAID group.

In addition, LDE cannot be performed for RAID groups in which the following conditions apply.

- RAID groups that belong to TPPs
- RAID groups in which volumes are not registered
- RAID groups in which WSVs are registered
- RAID groups that are configured with RAID5+0 or RAID6-FR

LDE may not be available when other functions are being used in the ETERNUS DX or the target RAID group.

For details on the functions that can be executed simultaneously and the number of the process that can be processed simultaneously, refer to "[Combinations of Functions That Are Available for Simultaneous Executions](#)" (page 144).

 **Caution**

- If drives of different capacities exist in a RAID group that is to be expanded while adding drives, the smallest capacity becomes the standard for the RAID group after expansion, and all other drives are regarded as having the same capacity as the smallest drive. In this case, the remaining drive space is not used.

If drives of different rotational speeds exist in a RAID group, the access performance of the RAID group is reduced by the slower drives.

- Since the data cannot be recovered after the failure of LDE, back up all the data of the volumes in the target RAID group to another area before performing LDE.
- If configuring RAID groups with Advanced Format drives, the write performance may be reduced when accessing volumes created in the relevant RAID group from an OS or an application that does not support Advanced Format.
- If the firmware version of the storage system is earlier than V10L88-6000, an error may occur due to a Flexible Write Through action while LDE is running (*1). Before executing LDE, set "Flexible Write Through" to "Disable" with the "Setup Subsystem Parameters" function.

***1 :** When using the controller firmware for the China region, this applies to controller firmware versions earlier than V10L88-9006.

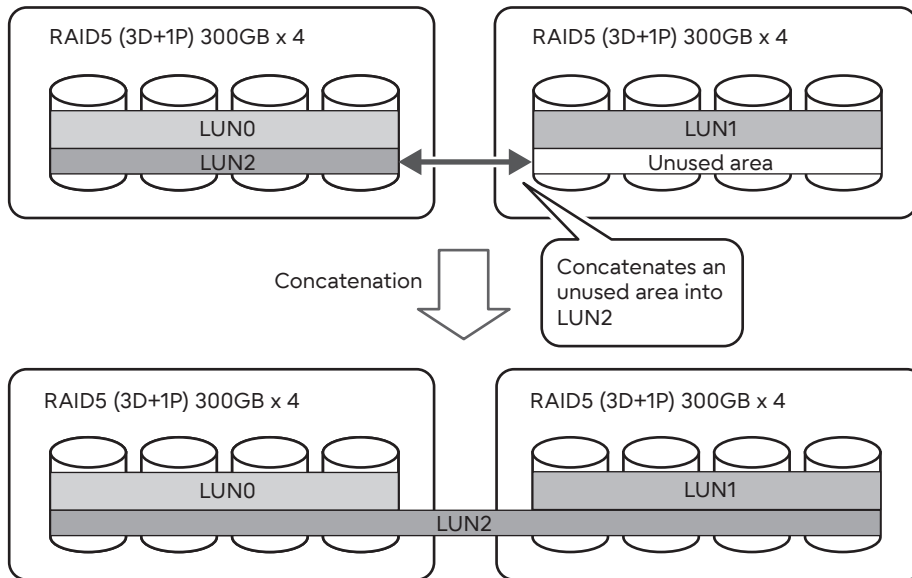
LUN Concatenation

LUN Concatenation is a function that is used to add new area to a volume and so expand the volume capacity available to the server. This function enables the reuse of leftover free area in a RAID group and can be used to solve capacity shortages.

Unused areas, which may be either part or all of a RAID group, are used to create new volumes that are then added together (concatenated) to form a single large volume.

The capacity can be expanded during an operation.

Figure 28 LUN Concatenation



LUN Concatenation is a function to expand a volume capacity by concatenating volumes.

Up to 16 volumes with a minimum capacity of 1GB can be concatenated.

When there are concatenation source volumes in SAS disks or Nearline SAS disks, concatenation can be performed with volumes in SAS disks or Nearline SAS disks.

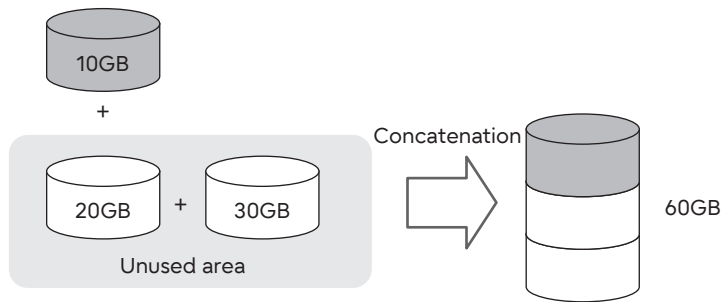
For SSDs, the drives for the concatenation source and destination volumes must be the same type (SSD).

From a performance perspective, using RAID groups with the same RAID level and the same drives (type, size, capacity, and rotational speed) is recommended as the concatenation source.

A concatenated volume can be used as an OPC, EC, or QuickOPC copy source or copy destination. It can also be used as a SnapOPC/SnapOPC+ copy source.

The LUN number stays the same before and after the concatenation. Because the server-side LUNs are not changed, an OS reboot is not required. Data can be accessed from the host in the same way regardless of the concatenation status (before, during, or after concatenation). However, the recognition methods of the volume capacity expansion vary depending on the OS types.

Figure 29 LUN Concatenation (Volume Concatenation)



Only Standard type volumes can be used for LUN Concatenation.

LUN Concatenation may not be available when other functions are being used in the device or the target volume.

For details on the functions that can be executed simultaneously, refer to "[Combinations of Functions That Are Available for Simultaneous Executions](#)" (page 144).

Caution

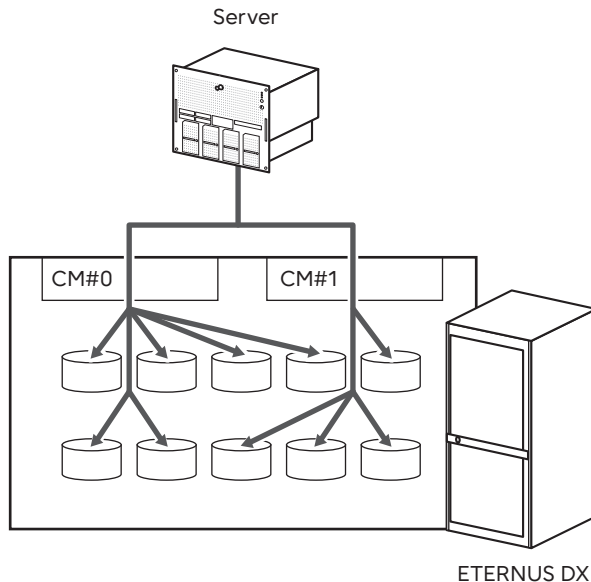
- It is recommended that the data on the volumes that are to be concatenated be backed up first.
- Refer to the applicable OS and file system documentation before dynamically expanding the volume capacity because expanded volumes may not be recognized by some types and versions of server-side platforms (OSs).
- When a volume that is using ETERNUS SF AdvancedCopy Manager to run backups is expanded via LUN Concatenation, the volume will need to be registered with ETERNUS SF Advanced-Copy Manager again.
- When specifying a volume in the RAID group configured with Advanced Format drives as a concatenation source or a concatenation destination to expand the capacity, the write performance may be reduced when accessing the expanded volumes from an OS or an application that does not support Advanced Format.

Wide Striping

Wide Striping is a function that concatenates multiple RAID groups by striping and uses many drives simultaneously to improve performance. This function is effective when high Random Write performance is required.

I/O accesses from the server are distributed to multiple drives by increasing the number of drives that configure a LUN, which improves the processing performance.

Figure 30 Wide Striping



Wide Striping creates a WSV that can be concatenated across 2 to 48 RAID groups.

The number of RAID groups that are to be concatenated is defined when creating a WSV. The number of concatenated RAID groups cannot be changed after a WSV is created. To change the number of concatenated groups or expand the group capacity, perform RAID Migration.

Other volumes (Standard, SDVs, SDPVs, or WSVs) can be created in the free area of a RAID group that is concatenated by Wide Striping.

WSVs cannot be created in RAID groups with the following conditions.

- RAID groups that belong to TPPs
- RAID groups with different stripe size values
- RAID groups that are configured with different types of drives
- RAID groups that are configured with RAID6-FR
- RAID groups without a continuous free area greater than the capacity of the volumes to be concatenated

Caution

- Capacity expansion cannot be performed (LDE is not possible) for RAID groups used to configure a WSV.
- If one or more RAID groups that are configured with Advanced Format drives exist in the RAID group that is to be concatenated by striping to create a WSV, the write performance may be reduced when accessing the created WSVs from an OS or an application that does not support Advanced Format.

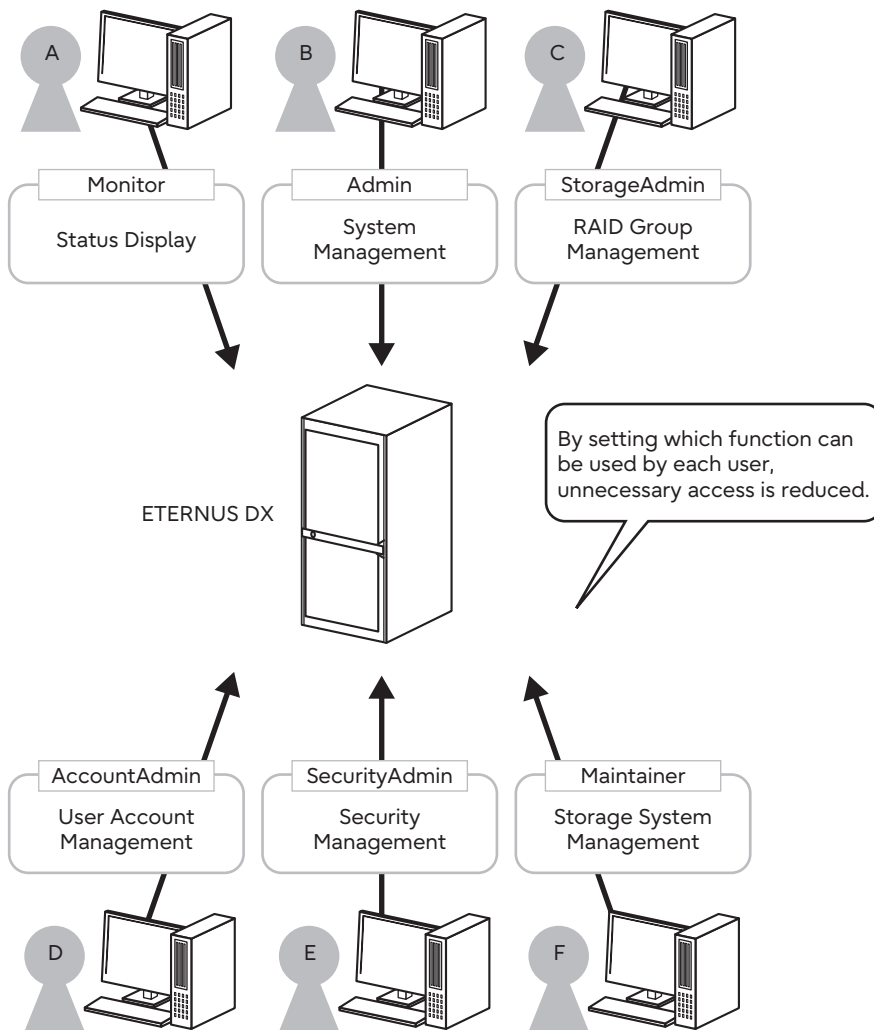
User Access Management

Account Management

The ETERNUS DX uses roles to allocate access authority when a user account is created, and sets which functions can be used depending on the user privileges.

Since the authorized functions of the storage administrator are classified according to the usage and only minimum privileges are given to the administrator, security is improved and operational mistakes and management hours can be reduced.

Figure 31 Account Management



Up to 60 user accounts can be set in the ETERNUS DX.

Up to 16 users can be logged in at the same time using ETERNUS Web GUI or ETERNUS CLI.

The menu that is displayed after logging on varies depending on the role that is added to a user account.

- Roles and available functions

Seven default roles are provided in the ETERNUS DX. The following table shows the roles and the available functions (categories).

Table 15 Available Functions for Default Roles

| Categories | Roles | | | | | | |
|------------------------------|---------|-------|---------------|---------------|----------------|------------|---------------|
| | Monitor | Admin | Storage Admin | Account Admin | Security Admin | Maintainer | Software (*1) |
| Status Display | ○ | ○ | ○ | × | ○ | ○ | × |
| RAID Group Management | × | ○ | ○ | × | × | ○ | × |
| Volume - Create / Modify | × | ○ | ○ | × | × | ○ | × |
| Volume - Delete / Format | × | ○ | ○ | × | × | ○ | × |
| Host Interface Management | × | ○ | ○ | × | × | ○ | × |
| Advanced Copy Management | × | ○ | ○ | × | × | ○ | × |
| Copy Session Management | × | ○ | ○ | × | × | ○ | × |
| Storage Migration Management | × | ○ | ○ | × | × | ○ | × |
| Storage Management | × | ○ | × | × | × | ○ | × |
| User Management | × | ○ | × | ○ | × | × | × |
| Authentication / Role | × | ○ | × | ○ | × | × | × |
| Security Setting | × | ○ | × | × | ○ | × | × |
| Maintenance Information | × | ○ | × | × | ○ | ○ | × |
| Firmware Management | × | ○ | × | × | × | ○ | × |
| Maintenance Operation | × | × | × | × | × | ○ | × |

○: Supported category ×: Not supported

***1:** This is the role that is used for external software. A user account with a "Software" role cannot be used with ETERNUS Web GUI or ETERNUS CLI.

 **Caution**

- The default roles cannot be deleted or edited.
- The function categories for the roles cannot be changed.
- A role must be assigned when creating a user account.

User Authentication

Internal Authentication and External Authentication are available as logon authentication methods. RADIUS authentication can be used for External Authentication.

The user authentication functions described in this section can be used when performing storage management and operation management, and when accessing the ETERNUS DX via operation management LAN.

● Internal Authentication

Internal Authentication is performed using the authentication function of the ETERNUS DX.

The following authentication functions are available when the ETERNUS DX is connected via a LAN using operation management software.

- User account authentication

User account authentication uses the user account information that is registered in the ETERNUS DX to verify user logins. Up to 60 user accounts can be set to access the ETERNUS DX.

- SSL authentication

ETERNUS Web GUI and SMI-S support HTTPS connections using SSL/TLS. Since data on the network is encrypted, security can be ensured. Server certifications that are required for connection are automatically created in the ETERNUS DX.

- SSH authentication

Since ETERNUS CLI supports SSH connections, data that is sent or received on the network can be encrypted. The server key for SSH varies depending on the ETERNUS DX. When the server certification is updated, the server key is updated as well.

Password authentication and client public key authentication are available as authentication methods for SSH connections.

The supported client public keys are shown below.

Table 16 Client Public Key (SSH Authentication)

| Type of public key | Complexity (bits) |
|---------------------------|----------------------|
| IETF style DSA for SSH v2 | 1024, 2048, and 4096 |
| IETF style RSA for SSH v2 | 1024, 2048, and 4096 |

● External Authentication

External Authentication uses the user account information (user name, password, and role name) that is registered on an external authentication server. RADIUS authentication supports ETERNUS Web GUI and the ETERNUS CLI login authentication for the ETERNUS DX, and authentication for connections to the ETERNUS DX through a LAN using operation management software.

- RADIUS authentication

RADIUS authentication uses the Remote Authentication Dial-In User Service (RADIUS) protocol to consolidate authentication information for remote access.

An authentication request is sent to the RADIUS authentication server that is outside the ETERNUS system network. The authentication method can be selected from CHAP and PAP. Two RADIUS authentication servers (the primary server and the secondary server) can be connected to balance user account information and to create a redundant configuration. When the primary RADIUS server failed to authenticate, the secondary RADIUS server attempts to authenticate.

User roles are specified in the Vendor Specific Attribute (VSA) of the Access-Accept response from the server. The following table shows the syntax of the VSA based account role on the RADIUS server.

| Item | Size (octets) | Value | Description |
|--------------------|---------------|------------------|---|
| Type | 1 | 26 | Attribute number for the Vendor Specific Attribute |
| Length | 1 | 7 or more | Attribute size (calculated by server) |
| Vendor-Id | 4 | 211 | Fujitsu Limited (SMI Private Enterprise Code) |
| Vendor type | 1 | 1 | Eternus-Auth-Role |
| Vendor length | 1 | 2 or more | Attribute size described after Vendor type (calculated by server) |
| Attribute-Specific | 1 or more | ASCII characters | One or more assignable role names for successfully authenticated users (*1) |

*1: The server-side role names must be identical to the role names of the ETERNUS DX. Match the letter case when entering the role names.

[Example] RoleName0

 **Caution**

- If RADIUS authentication fails when "Do not use Internal Authentication" has been selected for "Authentication Error Recovery" on ETERNUS Web GUI, ETERNUS CLI, or SMI-S, logging on to ETERNUS Web GUI or ETERNUS CLI will not be available.
When the setting to use Internal Authentication for errors caused by network problems is configured, Internal Authentication is performed if RADIUS authentication fails on both primary and secondary RADIUS servers, or at least one of these failures is due to network error.
- So long as there is no RADIUS authentication response the ETERNUS DX will keep retrying to authenticate the user for the entire "Timeout" period set on the "Set RADIUS Authentication (Initial)" menu. If authentication does not succeed before the "Timeout" period expires, RADIUS Authentication is considered to be a failure.
- When using RADIUS authentication, if the role that is received from the server is unknown (not set) for the device, RADIUS authentication fails.

Audit Log

The ETERNUS DX can send information such as access records by the administrator and setting changes as audit logs to the Syslog servers.


Audit logs are audit trail information that record operations that are executed for the ETERNUS DX and the response from the system. This information is required for auditing.

The audit log function enables monitoring of all operations and any unauthorized access that may affect the system.

Syslog protocols (RFC3164 and RFC5424) are supported for audit logs.

Information that is to be sent is not saved in the ETERNUS DX and the Syslog protocols are used to send out the information. Two Syslog servers can be set as the destination servers in addition to the Syslog server that is used for event notification.

Figure 32 Audit Log

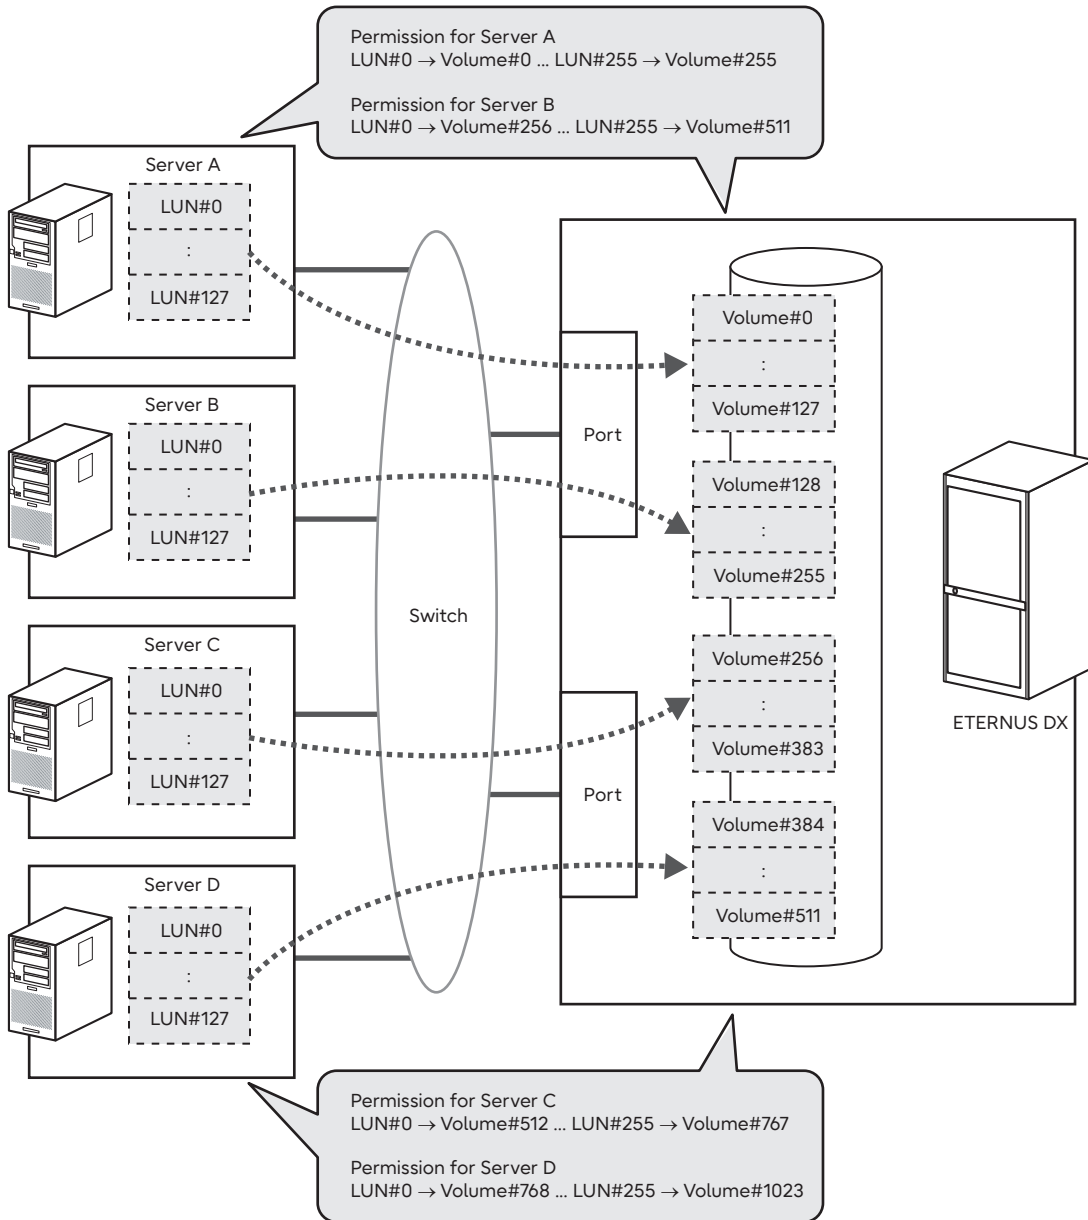


Improving Host Connectivity

Host Affinity

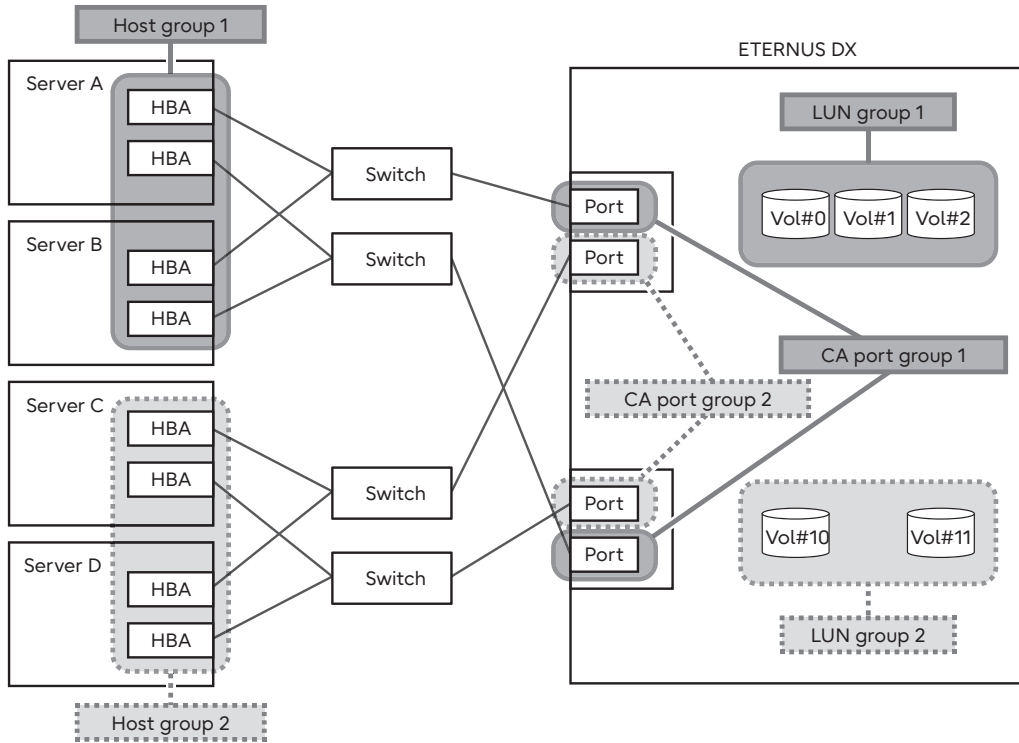
The host affinity function prevents data from being damaged due to inadvertent storage access. By defining a server that can access the volume, security can be ensured when multiple servers are connected.

Figure 33 Host Affinity



The host affinity can be set by associating "Host Groups", "CA Port Groups", and "LUN Groups".

Figure 34 Associating Host Groups, CA Port Groups, and LUN Groups



The host affinity can also be set by directly specifying the host and the CA port without creating host groups and CA port groups.

● **Host Group**

A host group is a group of hosts that have the same host interface type and that access the same LUN group. HBAs in multiple hosts can be configured in a single host group.

● **CA Port Group**

A CA port group is a group of the same CA type ports that are connected to a specific host group. A CA port group is configured with ports that access the same LUN group, such as ports that are used for multipath connection to the server or for connecting to the cluster configuring server. A single CA port group can be connected to multiple host groups.

● **LUN Group**

A LUN group is a group of LUNs that can be recognized by the host and the LUN group can be accessed from the same host group and CA port groups.

A LUN group is mapping information for LUNs and volumes.

Caution

- Host access must be prevented when changing or deleting already set host affinity settings. When adding a new LUN to the host affinity settings, it is not necessary to stop host access.
- When servers are duplicated and connected using a cluster configuration to share a single ETERNUS DX among multiple servers, cluster control software is required.

iSCSI Security

For an iSCSI interface, the iSCSI authentication function can be used when the initiator accesses the target. The iSCSI authentication function is available for host connections and remote copying.

The Challenge Handshake Authentication Protocol (CHAP) is supported for iSCSI authentication. For CHAP Authentication, unidirectional CHAP or bidirectional CHAP can be selected. When unidirectional CHAP is used, the target authenticates the initiator to prevent fraudulent access. When bidirectional CHAP is used, the target authenticates the initiator to prevent fraudulent access and the initiator authenticates the target to prevent impersonation.

Note that the Internet Storage Name Service (iSNS) is also supported as an iSCSI name resolution.

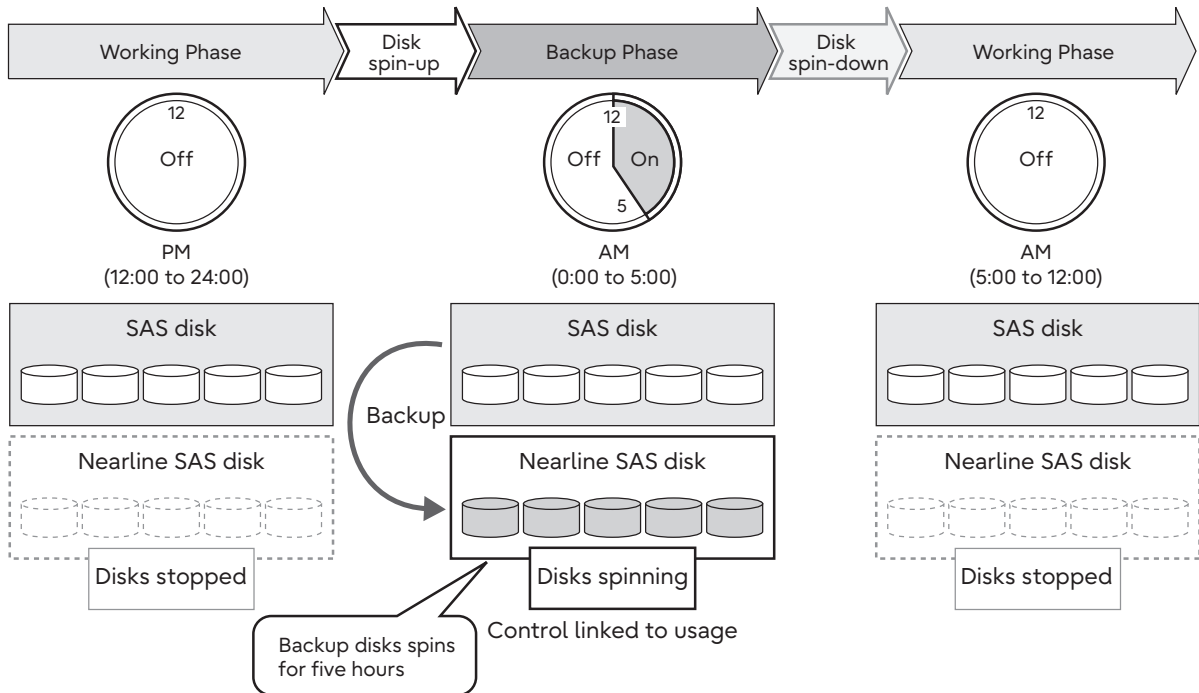
Environmental Burden Reduction

Eco-mode

Eco-mode is a function that reduces power consumption for limited access disks by stopping the disks rotation during specified periods or by powering off the disks.

Disk spin-up and spin-down schedules can be set for each RAID group or TPP. These schedules can also be set to allow backup operations.

Figure 35 Eco-mode



The Eco-mode of the ETERNUS DX is a function specialized for reducing power consumption attributed to Massive Arrays of Idle Disks (MAID). The operational state for stopping a disk can be selected from two modes: "stop motor" or "turn off drive power".

The disks to be controlled are SAS disks and Nearline SAS disks.

Eco-mode cannot be used for the following drives:

- Global Hot Spares (Dedicated Hot Spares are possible)
- SSDs
- Unused drives (that are not used by RAID groups)

The Eco-mode schedule cannot be specified for the following RAID groups or pools:

- No volumes are registered
- Configured with SSDs
- RAID groups to which the volume with Storage Migration path belongs

For RAID groups with the following conditions, the Eco-mode schedule can be set but the disks motor cannot be stopped or the power supply cannot be turned off:

- SDPVs are registered
- ODX Buffer volumes are registered

If disk access occurs while the disk motor is stopped, the disk is immediately spun up and can be accessed within one to five minutes.

The Eco-mode function can be used with the following methods:

- Schedule control
Controls the disk motors by configuring the Eco-mode schedule on ETERNUS Web GUI or ETERNUS CLI. The operation time schedule settings/management is performed for each RAID group and TPP.
- External application control (software interaction control)
Disk motor is controlled for each RAID group on ETERNUS SF Software.
The disk motors are controlled by interacting with applications installed on the server side and responding to instructions from the applications. Applications which can be interacted with are as follows:
 - ETERNUS SF Storage Cruiser
 - ETERNUS SF AdvancedCopy Manager

The following hierarchical storage management software can be also linked with Eco-mode.

When using the Eco-mode function with these products, an Eco-mode disk operating schedule does not need to be set. A drive in a stopped condition starts running when it is accessed.

- IBM Tivoli Storage Manager for Space Management
- IBM Tivoli Storage Manager HSM for Windows
- Symantec Veritas Storage Foundation Dynamic Storage Tiering (DST) function

The following table shows the specifications of Eco-mode.

Table 17 Eco-mode Specifications

| Item | Description | Remarks |
|--|---------------------------------|--|
| Number of registrable schedules | 64 | Up to 8 events (during disk operation) can be set for each schedule. |
| Host I/O Monitoring Interval (*1) | 30 minutes (default) | Monitoring time can be set from 10 to 60 minutes. This setting can only be changed by users with the Maintenance Operation policy. |
| Disk Motor Spin-down Limit Count (per day) | 25 (default) | The number of times the disk is stopped can be set from 1 to 25. When it exceeds the upper limit, Eco-mode becomes unavailable, and the disks keep running. This setting can only be changed by users with the Maintenance Operation policy. |
| Target drive | SAS disks Nearline SAS disks | SSD is not supported. |

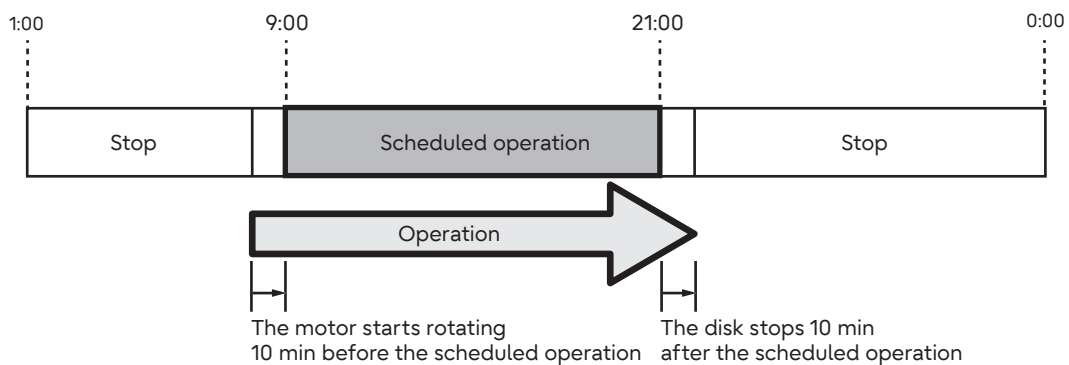
***1:** The monitoring time period to check if there is no access to a disk for a given length of time and stop the drive.

Caution

- To set Eco-mode schedule, use ETERNUS Web GUI, ETERNUS CLI, ETERNUS SF Storage Cruiser, or ETERNUS SF AdvancedCopy Manager. Note that schedules that are created by ETERNUS Web GUI or ETERNUS CLI and schedules that are created by ETERNUS SF Storage Cruiser or ETERNUS SF AdvancedCopy Manager cannot be shared. Make sure to use only one type of software to manage a RAID group.
- Use ETERNUS Web GUI or ETERNUS CLI to set Eco-mode for TPPs. ETERNUS SF Storage Cruiser or ETERNUS SF AdvancedCopy Manager cannot be used to set the Eco-mode for TPPs and FTRPs.
- Specify the same Eco-mode schedule for the RAID groups that configure a WSV. If different Eco-mode schedules are specified, activation of stopped disks when host access is performed occurs and the response time may increase.
- The operation time of disks varies depending on the Eco-mode schedule and the disk access.
 - Access to a stopped disk outside of the scheduled operation time period causes the motor of the stopped disk to be spun up, allowing normal access in about one to five minutes. When a set time elapses since the last access to a disk, the motor of the disk is stopped.
 - If a disk is activated from the stopped state more than a set amount of times in a day, the Eco-mode schedule is not applied and disk motors are not stopped by the Eco-mode.

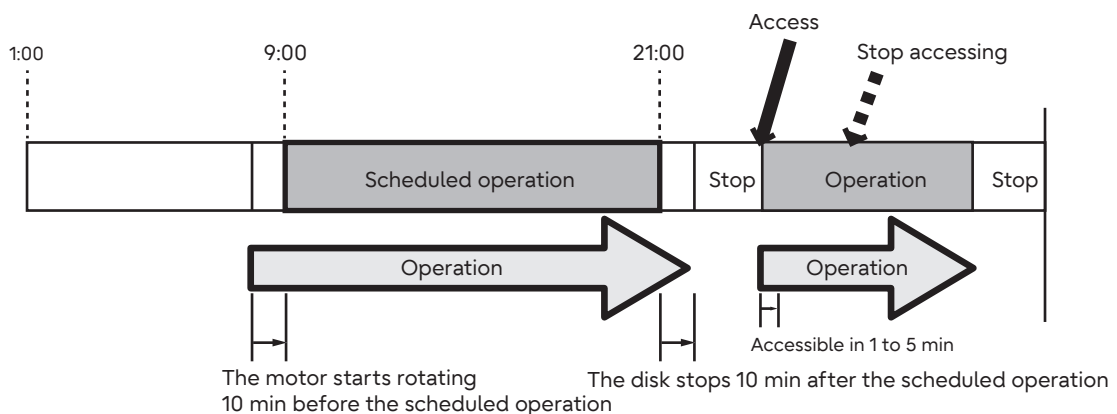
(Example 1) Setting the Eco-mode schedule via ETERNUS Web GUI

Operation schedule is set as 9:00 to 21:00 and there are no accesses outside of the scheduled period



(Example 2) Setting the Eco-mode schedule via ETERNUS Web GUI

Operation schedule is set as 9:00 to 21:00 and there are accesses outside of the scheduled period



Caution

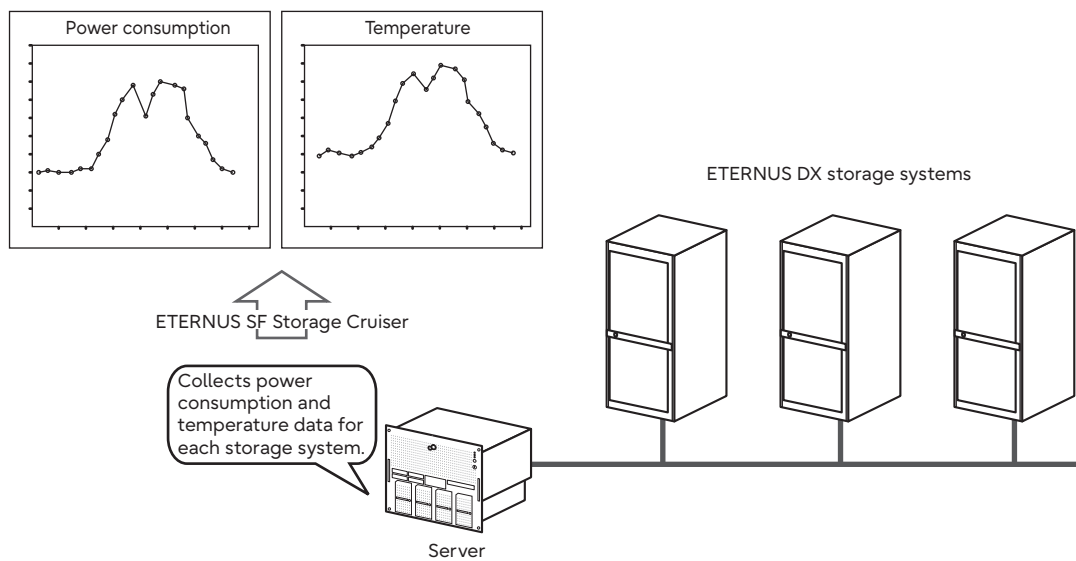
- Eco-mode schedules are executed according to the date and time that are set in the ETERNUS DX. To turn on and turn off the disk motors according to the schedule that is set, use the Network Time Protocol (NTP) server in the date and time setting in ETERNUS Web GUI to set automatic adjustment of the date and time.
- If the number of drives that are activated in a single drive enclosure is increased, the time for system activation may take longer (about 1 to 5 minutes). This is because all of the disks cannot be activated at the same time.
- Even if the disk motor is turned on and off repeatedly according to the Eco-mode schedule, the failure rate is not affected comparing to the case when the motor is always on.

Power Consumption Visualization

The power consumption and the temperature of the ETERNUS DX can be visualized with a graph by using the ETERNUS SF Storage Cruiser integrated management software in a storage system environment. The ETERNUS DX collects information on power consumption and the ambient temperature in the storage system. Collected information is notified using SNMP and graphically displayed on the screens by ETERNUS SF Storage Cruiser. Cooling efficiency can be improved by understanding local temperature rises in the data center and reviewing the location of air-conditioning.

Understanding the drives that have a specific time to be used from the access frequency to RAID groups enables the Eco-mode schedule to be adjusted accordingly.

Figure 36 Power Consumption Visualization



Operation Management/Device Monitoring

Operation Management Interface

Operation management software can be selected in the ETERNUS DX according to the environment of the user.

ETERNUS Web GUI and ETERNUS CLI are embedded in the ETERNUS DX controllers.

The setting and display functions can also be used with ETERNUS SF Web Console.

■ ETERNUS Web GUI

ETERNUS Web GUI is a program for settings and operation management that is embedded in the ETERNUS DX and accessed by using a web browser via http or https.

ETERNUS Web GUI has an easy-to-use design that makes intuitive operation possible.

The settings that are required for the ETERNUS DX initial installation can be easily performed by following the wizard and inputting the parameters for the displayed setting items.

SSL v3 and TLS are supported for https connections. However, when using https connections, it is required to register a server certification in advance or self-generate a server certification. Self-generated server certifications are not already certified with an official certification authority registered in web browsers. Therefore, some web browsers will display warnings. Once a server certification is installed in a web browser, the warning will not be displayed again.

When operations are managed using ETERNUS Web GUI, a Web browser must be prepared in the administration terminal.

For details on the operating environment of ETERNUS Web GUI, refer to "ETERNUS Web GUI User's Guide" or "Configuration Guide (Web GUI)".

When using ETERNUS Web GUI to connect the ETERNUS DX, the default port number is 80 for http.

■ ETERNUS CLI

ETERNUS CLI supports Telnet or SSH connections. The ETERNUS DX can be configured and monitored using commands and command scripts.

With the ETERNUS CLI, SSH v2 encrypted connections can be used. SSH server keys differ for each storage system, and must be generated by the SSH server before using SSH.

Password authentication and client public key authentication are supported as authentication methods for SSH.

For details on supported client public key types, refer to "[User Authentication](#)" (page 52).

■ ETERNUS SF

ETERNUS SF can manage a Fujitsu storage products centered storage environment. An easy-to-use interface enables complicated storage environment design and setting operations, which allows easy installation of a storage system without needing to have high level skills.

ETERNUS SF ensures stable operation by managing the entire storage environment.

■ SMI-S

Storage systems can be managed collectively using the general storage management application that supports Version 1.6 of Storage Management Initiative Specification (SMI-S). SMI-S is a storage management interface standard of the Storage Network Industry Association (SNIA). SMI-S can monitor the ETERNUS DX status and change configurations such as RAID groups, volumes, and Advanced Copy (EC/OPC/SnapOPC/SnapOPC+).

Performance Information Management

The ETERNUS DX supports a function that collects and displays the performance data of the storage system via ETERNUS Web GUI or ETERNUS CLI. The collected performance information shows the operation status and load status of the ETERNUS DX and can be used to optimize the system configuration.

ETERNUS SF Storage Cruiser can be used to easily understand the operation status and load status of the ETERNUS DX by graphically displaying the collected information on the GUI. ETERNUS SF Storage Cruiser can also monitor the performance threshold and retain performance information for the duration that a user specifies.

When performance monitoring is operated from ETERNUS SF Storage Cruiser, ETERNUS Web GUI, or ETERNUS CLI, performance information in each type is obtained during specified intervals (30 - 300 seconds) in the ETERNUS DX.

The performance information can be stored and exported in the text file format, as well as displayed, from ETERNUS Web GUI. The performance information, which can be obtained, are indicated as follows.

● Volume Performance Information for Host I/O

- Read IOPS (the read count per second)
- Write IOPS (the write count per second)
- Read Throughput (the amount of transferred data that is read per second)
- Write Throughput (the amount of transferred data that is written per second)
- Read Response Time (the average response time per host I/O during a read)
- Write Response Time (the average response time per host I/O during a write)
- Read Process Time (the average process time in the storage system per host I/O during a read)
- Write Process Time (the average process time in the storage system per host I/O during a write)
- Read Cache Hit Rate (cache hit rate for read)
- Write Cache Hit Rate (cache hit rate for write)
- Prefetch Cache Hit Rate (cache hit rate for prefetch)

● Volume Performance Information for the Advanced Copy Function

- Read IOPS (the read count per second)
- Write IOPS (the write count per second)
- Read Throughput (the amount of transferred data that is read per second)
- Write Throughput (the amount of transferred data that is written per second)
- Read Cache Hit Rate (cache hit rate for read)
- Write Cache Hit Rate (cache hit rate for write)
- Prefetch Cache Hit Rate (cache hit rate for prefetch)

- **Controller Performance Information**

- Busy Ratio (CPU usage)
- CPU core usage

- **CA Port Performance Information**

- Read IOPS (the read count per second)
- Write IOPS (the write count per second)
- Read Throughput (the amount of transferred data that is read per second)
- Write Throughput (the amount of transferred data that is written per second)

- **Drive Performance Information**

- Busy Ratio (drive usage)

 **Caution**

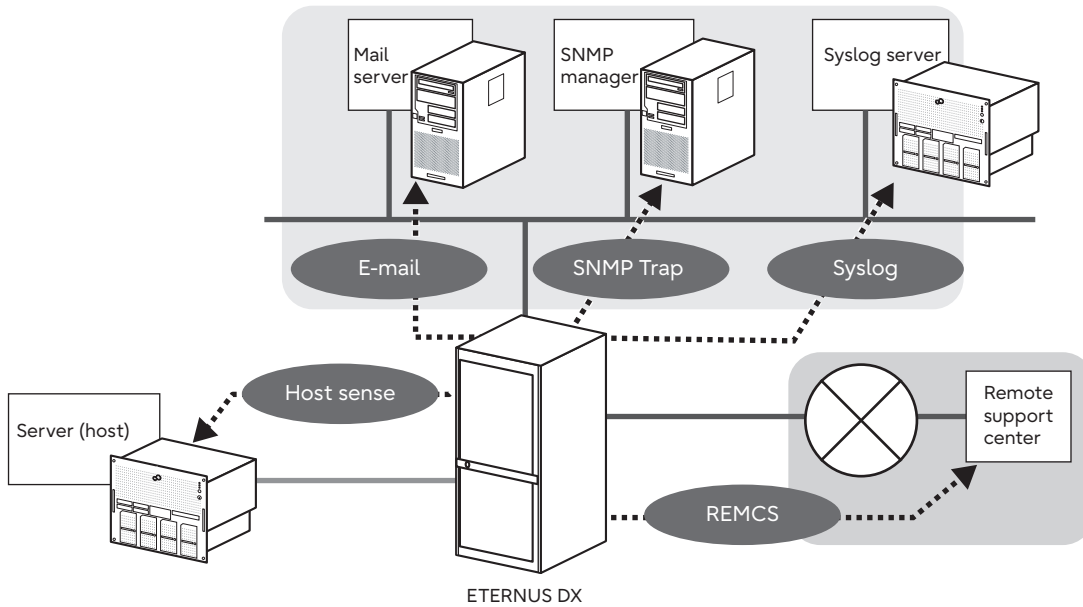
- When the ETERNUS DX is rebooted, the performance monitoring process is stopped.
 - If performance monitoring is started from ETERNUS SF Storage Cruiser, ETERNUS Web GUI or ETERNUS CLI cannot stop the process.
 - If performance monitoring is started from ETERNUS Web GUI or ETERNUS CLI, the process can be stopped from ETERNUS SF Storage Cruiser.
-

Event Notification

When an error occurs in the ETERNUS DX, the event notification function notifies the event information to the administrator. The administrator can be informed that an error occurred without monitoring the screen all the time.

The methods to notify an event are e-mail, SNMP Trap, syslog, remote support, and host sense.

Figure 37 Event Notification



The notification methods and levels can be set as required.

The following events are notified.

Table 18 Levels and Contents of Events That Are Notified

| Level | Level of importance | Event contents |
|----------------------------|-------------------------------------|---|
| Error | Maintenance is necessary | Component failure, temperature error, end of battery life (*1), rebuild/copy-back, etc. |
| Warning | Preventive maintenance is necessary | Module warning, battery life warning (*1), etc. |
| Notification (information) | Device information | Component restoration notification, user login/logout, RAID creation/deletion, storage system power on/off, firmware update, etc. |

*1: Battery related events are notified only for the ETERNUS DX60 S4.

● E-Mail

When an event occurs, an e-mail is sent to the specified e-mail address.

The ETERNUS DX supports "SMTP AUTH" and "SMTP over SSL" as user authentication. A method can be selected from CRAM-MD5, PLAIN, LOGIN, or AUTO which automatically selects one of these methods.

● Simple Network Management Protocol (SNMP)

Using the SNMP agent function, management information is sent to the SNMP manager (network management/monitoring server).

The ETERNUS DX supports the following SNMP specifications.

Table 19 SNMP Specifications

| Item | Specification | Remarks |
|--------------|-----------------------|---|
| SNMP version | SNMP v1, v2c, v3 | — |
| MIB | MIB II | Only the information managed by the ETERNUS DX can be sent with the GET command. The SET command send operation is not supported. |
| | FibreAlliance MIB 2.2 | This is a MIB which is defined for the purpose of FC base SAN management. Only the information managed by the ETERNUS DX can be sent with the GET command. The SET command send operation is not supported. |
| | Unique MIB | This is a MIB in regard to hardware configuration of the ETERNUS DX. |
| Trap | Unique Trap | A trap number is defined for each category (such as a component disconnection and a sensor error) and a message with a brief description of an event as additional information is provided. |

● Syslog

By registering the syslog destination server in the ETERNUS DX, various events that are detected by the ETERNUS DX are sent to the syslog server as event logs.

The ETERNUS DX supports the syslog protocol which conforms to RFC3164 and RFC5424.

● Remote Support

The errors that occur in the ETERNUS DX are notified to the remote support center. The ETERNUS DX sends additional information (logs and system configuration information) for checking the error. This shortens the time to collect information.

Remote support has the following maintenance functions.

- Failure notice

This function reports various failures, that occur in the ETERNUS DX, 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.

- Firmware download

The latest firmware in the remote support center is automatically registered in the ETERNUS DX. This function ensures that the latest firmware is registered in the ETERNUS DX, and prevents known errors from occurring. Firmware can also be registered manually.

● Host Sense

The ETERNUS DX returns host senses (sense codes) to notify specific status to the server. Detailed information such as error contents can be obtained from the sense code.

Caution

- Note that the ETERNUS DX cannot check whether the event log is successfully sent to the syslog server. Even if a communication error occurs between the ETERNUS DX and the syslog server, event logs are not sent again. When using the syslog function (enabling the syslog function) for the first time, confirm that the syslog server has successfully received the event log of the relevant operation.
- Using the ETERNUS Multipath Driver to monitor the storage system by host senses is recommended.

Sense codes that cannot be detected in a single configuration can also be reported.

Device Time Synchronization

The ETERNUS DX treats the time that is specified in the Master CM as the system standard time and distributes that time to other modules to synchronize the storage time. The ETERNUS DX also supports the time correction function by using the Network Time Protocol (NTP). The ETERNUS DX corrects the system time by obtaining the time information from the NTP server during regular time correction.

The ETERNUS DX has a clock function and manages time information of date/time and the time zone (the region in which the ETERNUS DX is installed). This time information is used for internal logs and for functions such as Eco-mode and remote support.

The automatic time correction by NTP is recommended to synchronize time in the whole system.

When using the NTP, specify the NTP server or the SNTP server. The ETERNUS DX supports NTP protocol v4. The time correction mode is Step mode (immediate correction). The time is regularly corrected every three hours once the NTP is set.


Caution

- If an error occurs in a system that has a different date and time for each device, analyzing the cause of this error may be difficult.
- Make sure to set the date and time correctly when using Eco-mode.

The stop and start process of the disk motors does not operate according to the Eco-mode schedule if the date and time in the ETERNUS DX are not correct.

Using NTP to synchronize the time in the ETERNUS DX and the servers is recommended.

Figure 38 Device Time Synchronization




Power Control

Power Synchronized Unit

A power synchronized unit detects changes in the AC power output of the Uninterruptible Power Supply (UPS) unit that is connected to the server and automatically turns on and off the ETERNUS DX.

Figure 39 Power Synchronized Unit



Remote Power Operation (Wake On LAN)

Wake On LAN is a function that turns on the ETERNUS DX via a network.


When "magic packet" data is sent from an administration terminal, the ETERNUS DX detects the packet and the power is turned on.

To perform Wake On LAN, utility software for Wake On LAN such as Systemwalker Runbook Automation is required and settings for Wake On LAN must be performed.

The MAC address for the ETERNUS DX can be checked on ETERNUS CLI.

ETERNUS Web GUI or ETERNUS CLI can be used to turn off the power of an ETERNUS DX remotely.

Figure 40 Wake On LAN




Backup (Advanced Copy)

The Advanced Copy function (high-speed copying function) enables data backup (data replication) at any point without stopping the operations of the ETERNUS DX.

For an ETERNUS DX backup operation, data can be replicated without placing a load on the business server. The replication process for large amounts of data can be performed by controlling the timing and business access so that data protection can be considered separate from operation processes.

An example of an Advanced Copy operation using ETERNUS SF AdvancedCopy Manager is shown below.

Figure 41 Example of Advanced Copy



Advanced Copy functions include One Point Copy (OPC), QuickOPC, SnapOPC, SnapOPC+, and Equivalent Copy (EC).

The following table shows ETERNUS related software for controlling the Advanced Copy function.

Table 20 Control Software (Advanced Copy)

| Control software | Feature |
|---------------------------------|--|
| ETERNUS Web GUI / ETERNUS CLI | The copy functions can be used without optional software. |
| ETERNUS SF AdvancedCopy Manager | ETERNUS SF AdvancedCopy Manager supports various OSs and ISV applications, and enables the use of all the Advanced Copy functions. This software can also be used for backups that interoperate with Oracle, SQL Server, Exchange Server, or Symfoware Server without stopping operations. |
| ETERNUS SF Express | ETERNUS SF Express allows easy management and backup of systems with a single product. |

Table 21 List of Functions (Copy Methods)

| Number of available sessions | Control software | | |
|------------------------------|-------------------------------|--|--------------------|
| | ETERNUS Web GUI / ETERNUS CLI | ETERNUS SF AdvancedCopy Manager | ETERNUS SF Express |
| 1,024 | SnapOPC+ | SnapOPC SnapOPC+ QuickOPC OPC EC | SnapOPC+ |

A copy is executed for each LUN. With ETERNUS SF AdvancedCopy Manager, a copy can also be executed for each logical disk (which is called a partition or a volume depending on the OS).

A copy cannot be executed if another function is running in the storage system or the target volume. For details on the functions that can be executed simultaneously, refer to ["Combinations of Functions That Are Available for Simultaneous Executions"](#) (page 144).

Backup

Type of Copy

The Advanced Copy functions offer the following copy methods: "Mirror Suspend", "Background Copy", and "Copy-on-Write". The function names that are given to each method are as follows: "EC" for the "Mirror Suspend" method, "OPC" for the "Background Copy" method, and "SnapOPC" for the "Copy-on-Write" method.

When a physical copy is performed for the same area after the initial copy, OPC offers "QuickOPC", which only performs a physical copy of the data that has been updated from the previous version. The SnapOPC+ function only copies data that is to be updated and performs generation management of the copy source volume.

● OPC

All of the data in a volume at a specific point in time is copied to another volume in the ETERNUS DX.

OPC is suitable for the following usages:

- Performing a backup
- Performing system test data replication
- Restoring backup data (restoration after replacing a drive when the copy source drive has failed)

● QuickOPC

QuickOPC copies all data as initial copy in the same way as OPC. After all of the data is copied, only updated data (differential data) is copied. QuickOPC is suitable for the following usages:

- Creating a backup of the data that is updated in small amounts
- Performing system test data replication
- Restoration from a backup

● SnapOPC/SnapOPC+ (*1)

As updates occur in the source data, SnapOPC/SnapOPC+ saves the data prior to change to the copy destination (SDV/TPV). The data, prior to changes in the updated area, is saved to an SDP/TPP. Create an SDPV for the SDP when performing SnapOPC/SnapOPC+ by specifying an SDV as the copy destination.

SnapOPC/SnapOPC+ is suitable for the following usages:

- Performing temporary backup for tape backup
- Performing a backup of the data that is updated in small amounts (generation management is available for SnapOPC+)
- SnapOPC/SnapOPC+ operations that use an SDV/TPV as the copy destination logical volume have the following characteristics. Check the characteristics of each volume type before selecting the volume type.

Table 22 Characteristics of SnapOPC/SnapOPC+ Operations with Each Type of Copy Destination Logical Volume

| Item to compare | SDV | TPV |
|------------------------------|---|--|
| Ease of operation settings | The operation setting is complex because a dedicated SDV and SDP must be set | The operation setting is easy because a dedicated SDV and SDP are not required |
| Usage efficiency of the pool | The usage efficiency of the pool is higher because the allocated size of the physical area is small (8KB) | The usage efficiency of the pool is lower because the allocated size of the physical area is large with a chunk size of 21MB / 42MB / 84MB / 168MB |

- *1: The difference between SnapOPC and SnapOPC+ is that SnapOPC+ manages the history of updated data as opposed to SnapOPC, which manages updated data for a single generation only. While SnapOPC manages updated data in units per session thus saving the same data redundantly, SnapOPC+ has updated data as history information which can provide multiple backups for multiple generations.

● EC

An EC creates data that is mirrored from the copy source to the copy destination beforehand, and then suspends the copy and handles each data independently.

When copying is resumed, only updated data in the copy source is copied to the copy destination. If the copy destination data has been modified, the copy source data is copied again in order to maintain equivalence between the copy source data and the copy destination data. EC is suitable for the following usages:

- Performing a backup
- Performing system test data replication

 **Caution**

- If the SDP capacity is insufficient, a copy cannot be performed. In order to avoid this situation, an operation that notifies the operation administrator of event information according to the remaining SDP capacity is recommended. For more details on event notification, refer to "[Event Notification](#)" (page 65).
- For EC, the data in the copy destination cannot be referenced or updated until the copy session is suspended. If the monitoring software (ServerView Agents) performs I/O access to the data in the copy destination, an I/O access error message is output to the server log message and other destinations. To prevent error messages from being output, consider using other monitoring methods.

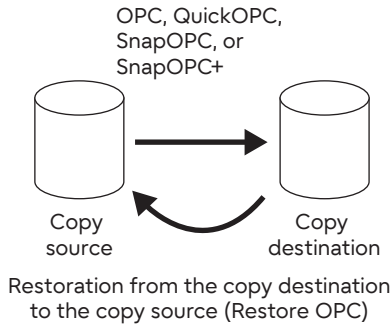
Available Advanced Copy Combinations

Different Advanced Copy types can be combined and used together.

- **Restore OPC**

For OPC, QuickOPC, SnapOPC, and SnapOPC+, restoration of the copy source from the copy destination is complete immediately upon request.

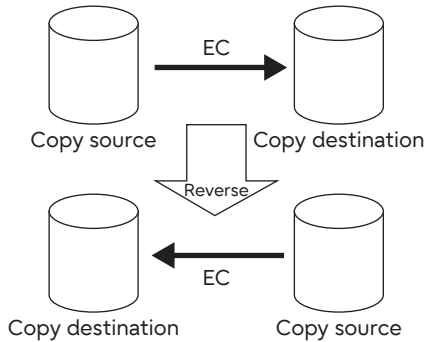
Figure 42 Restore OPC



- **EC Reverse**

Restoration can be performed by switching the copy source and destination of the EC.

Figure 43 EC Reverse



● Multi-Copy

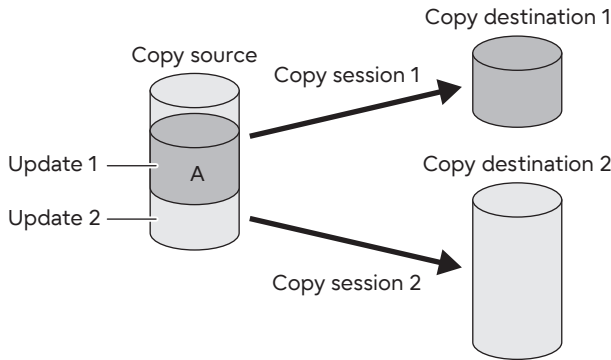
Multiple copy destinations can be set for a single copy source area to obtain multiple backups.

In the multi-copy shown in Figure 44, the entire range that is copied for copy session 1 will be the target for the multi-copy function.

When copy sessions 1 and 2 are EC, updates to area A in the copy source (update 1) are copied to both copy destination 1 and copy destination 2.

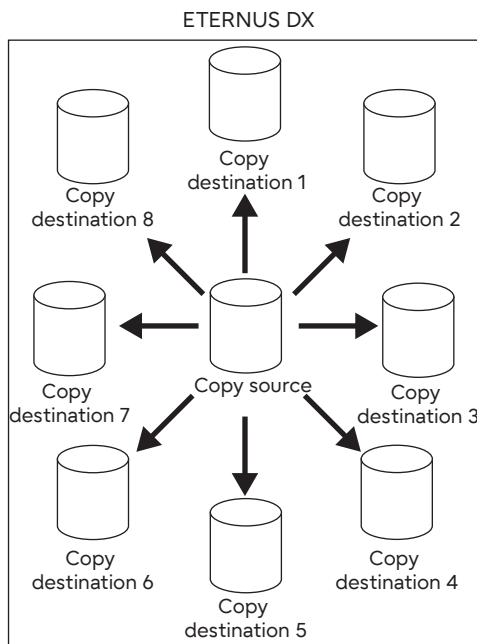
Updates to areas other than A in the copy source (update 2) are copied only to copy destination 2.

Figure 44 Targets for the Multi-Copy Function



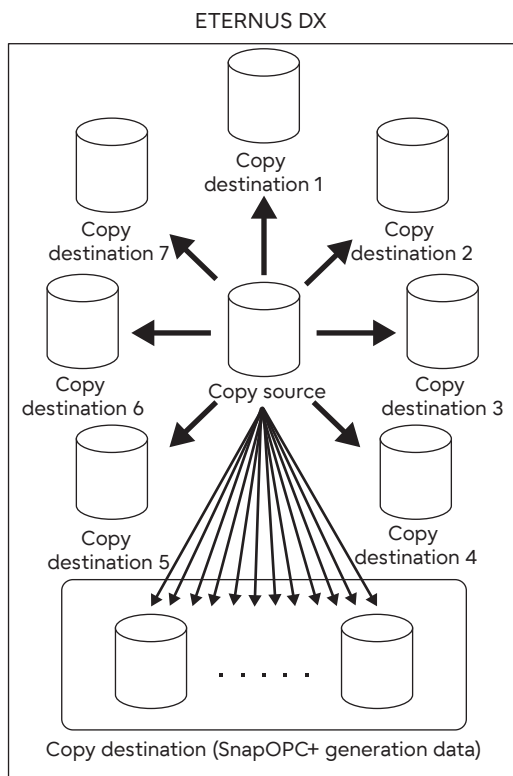
Up to eight OPC, QuickOPC, SnapOPC, or EC sessions can be set for a multi-copy.

Figure 45 Multi-Copy



For a SnapOPC+, the maximum number of SnapOPC+ copy session generations can be set for a single copy source area when seven or less multi-copy sessions are already set.

Figure 46 Multi-Copy (Including SnapOPC+)



● Cascade Copy

A copy destination with a copy session that is set can be used as the copy source of another copy session.

A Cascade Copy is performed by combining two copy sessions.

In Figure 47, "Copy session 1" refers to a copy session in which the copy destination area is also used as the copy source area of another copy session and "Copy session 2" refers to a copy session in which the copy source area is also used as the copy destination area of another copy session.

For a Cascade Copy, the copy destination area for copy session 1 and the copy source area for copy session 2 must be identical or the entire copy source area for copy session 2 must be included in the copy destination area for copy session 1.

A Cascade Copy can be performed when all of the target volumes are the same size or when the copy destination volume for copy session 2 is larger than the other volumes.

Figure 47 Cascade Copy

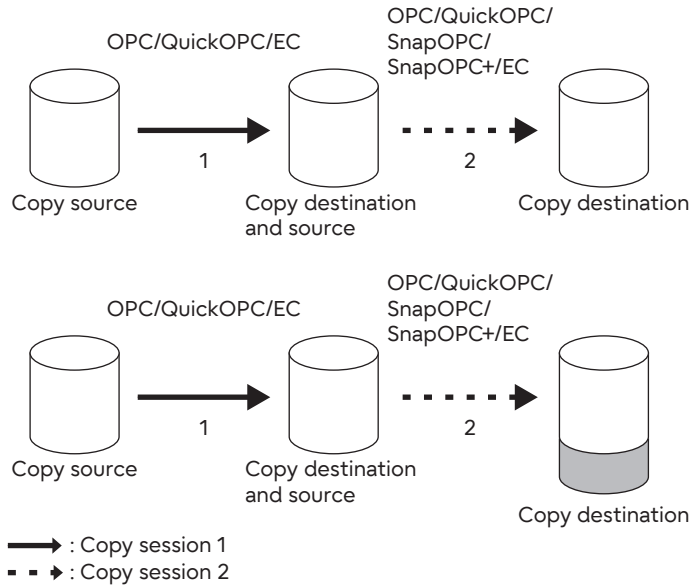


Table 23 shows the supported combinations when adding a copy session to a copy destination volume where a copy session has already been configured.

Table 23 Available Cascade Copy Combinations (When a Cascade Copy Performs Session 1 Followed by Session 2)

| Copy session 2 | Copy session 1 | | | | |
|----------------|----------------|-------------|---------|----------|----|
| | OPC | QuickOPC | SnapOPC | SnapOPC+ | EC |
| OPC | ○ (*1) | ○ (*1) | × | × | ○ |
| QuickOPC | ○ (*1) | ○ (*1) (*2) | × | × | ○ |
| SnapOPC | ○ (*1) | ○ (*1) | × | × | ○ |
| SnapOPC+ | ○ (*1) | ○ (*1) | × | × | ○ |
| EC | ○ | ○ | × | × | ○ |

○: Possible, ×: Not possible

*1: When copy session 2 is an OPC, QuickOPC, SnapOPC, or SnapOPC+ session, data in the copy destination of copy session 1 is backed up. Data is not backed up in the copy source of copy session 1.

*2: The supported modified data size varies depending on the controller firmware version.

- For controller firmware versions earlier than V10L82

This combination is supported only if the copy size in both the copy source volume and the copy destination volume is less than 2TB.

If the copy size is 2TB or larger, perform the following operations instead.

- When performing a temporary recovery
Use a Cascade Copy of QuickOPC (copy session 1) and OPC (copy session 2).
- When backing up two generations
Use a multi-copy that is configured with QuickOPC and QuickOPC.

- For controller firmware versions V10L82 or later

When issuing a re-copy request for copy session 1 (a request to copy only the data that has been updated after all data is copied), the modified data size should conform to the sizes shown in [Table 24](#).

For example, when a QuickOPC session is created at a resolution of 2, the modified data size should be less than 2TB. The "Resolution" for QuickOPC sessions can be checked on the [Advanced Copy (QuickOPC Sessions)] screen.

Table 24 Modified Data Size for Each QuickOPC Resolution

| Resolution | 1 | 2 | 4 | 8 | 16 | 32 | 64 |
|--------------------|---------------|---------------|---------------|---------------|----------------|----------------|----------------|
| Modified data size | Less than 1TB | Less than 2TB | Less than 4TB | Less than 8TB | Less than 16TB | Less than 32TB | Less than 64TB |

If a re-copy request is issued for a larger size, the command for the re-copy request may time out and fail. The current "Modified Data Size" for copy session 1 can be checked on the [Advanced Copy (QuickOPC Sessions)] screen. If an error occurs for copy session 1, perform one of the following operations instead.

- When performing a continuous Cascade Copy of QuickOPC
After the phase for copy session 2 becomes "Tracking", delete copy session 1. After changing the copy resolution of the ETERNUS DX, re-create copy session 1. For details on changing the resolution, refer to the "Modify Copy Table Size" section in "ETERNUS Web GUI User's Guide".
- When performing a temporary recovery
Use a Cascade Copy of QuickOPC (copy session 1) and OPC (copy session 2).
- When backing up two generations
Use a multi-copy that is configured with QuickOPC and QuickOPC.

[Table 25](#) shows the supported combinations when adding a copy session to a copy source volume where a copy session has already been configured.

Table 25 Available Cascade Copy Combinations (When a Cascade Copy Performs Session 2 Followed by Session 1)

| Copy session 1 | Copy session 2 | | | | |
|----------------|----------------|----------|---------|----------|----|
| | OPC | QuickOPC | SnapOPC | SnapOPC+ | EC |
| OPC | ○ | ○ | ○ | ○ | ○ |
| QuickOPC | ○ | ○ (*1) | ○ | ○ | ○ |
| SnapOPC | × | × | × | × | × |
| SnapOPC+ | × | × | × | × | × |
| EC | ○ | ○ | ○ | ○ | ○ |

○: Possible, ×: Not possible

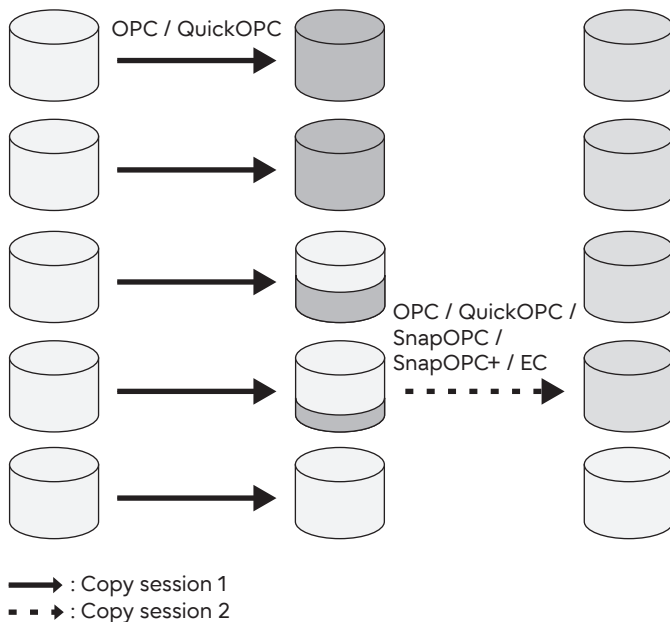
*1: The supported modified data size varies depending on the controller firmware version. Refer to footnote *2 of [Table 23](#) for details.

Caution

- To suspend a Cascade Copy where session 1 is performed before session 2 and session 2 is an EC session, perform the Suspend command after the physical copy for copy session 1 is complete.
- A Cascade Copy can be performed when the copy type for copy session 1 is XCOPYY or ODX. The copy destination area for XCOPYY or ODX and the copy source area for copy session 2 do not have to be completely identical. For example, a Cascade Copy can be performed when the copy source area for copy session 2 is only part of the copy destination area for copy session 1. XCOPYY or ODX cannot be set as the copy type for copy session 2 in a Cascade Copy.
- For more details on XCOPYY and ODX, refer to "Server Linkage Functions" (page 87).
- To acquire valid backup data in the copy destination for copy session 2, a physical copy must be completed or suspended in all of the copy sessions that configure the Cascade Copy. Check the copy status for copy sessions 1 and 2 when using the backup data.

However, if a Cascade Copy performs session 1 before session 2, and copy session 1 is an OPC or QuickOPC session and copy session 2 is an OPC, QuickOPC, SnapOPC, or SnapOPC+ session, the data in the copy destination for copy session 2 is available even during a physical copy.

- If copy session 1 is an EC session and copy session 2 is an OPC, QuickOPC, SnapOPC, or SnapOPC+ session, setting copy session 2 after setting copy session 1 to an equivalent or suspended state is recommended.
- When stopping an OPC or QuickOPC session for copy session 1 during a physical copy, stop copy session 2 in advance if copy session 2 is an OPC, QuickOPC, SnapOPC, or SnapOPC+ session.
- If copy session 2 is an EC session, copy session 2 does not transition to an equivalent state until the physical copy for copy session 1 is complete. For an EC session, a copy session cannot be suspended until the session transitions to an equivalent state.
- If a Cascade Copy performs session 1 before session 2, and copy session 1 is an OPC or QuickOPC session, the logical data in the intermediate volume when copy session 2 is started (the copy destination volume for copy session 1) is copied to the copy destination volume for copy session 2. A logical data copy is shown below.



Performance Tuning

Striping Size Expansion

Striping Size Expansion is a function to expand the stripe depth value by specifying the stripe depth when creating a RAID group.

Expansion of the stripe size enables advanced performance tuning. For normal operations, the default value does not need to be changed.

An expanded stripe depth reduces the number of drives that are accessed. A reduced number of commands to drives improves the access performance of the corresponding RAID1+0 RAID groups. However, it should be noted that an expanded stripe depth may reduce the sequential write performance for RAID5.

The stripe depth values that are available for each RAID level are shown below.

Table 26 Available Stripe Depth

| RAID level | Drive configuration (*1) | Available stripe depth |
|-----------------------------|--------------------------|--|
| Mirroring (RAID1) | 1D+1M | — |
| High performance (RAID1+0) | All drive configurations | 64KB, 128KB, 256KB, 512KB, and 1,024KB |
| Striping (RAID0) | | |
| High capacity (RAID5) | 2D+1P – 4D+1P | 64KB, 128KB, 256KB, and 512KB |
| | 5D+1P – 8D+1P | 64KB, 128KB, and 256KB |
| | 9D+1P – 15D+1P | 64KB and 128KB |
| Reliability (RAID5+0) | All drive configurations | 64KB |
| High reliability (RAID6) | | |
| High reliability (RAID6-FR) | | |

*1: D: Data, M: Mirror, P: Parity

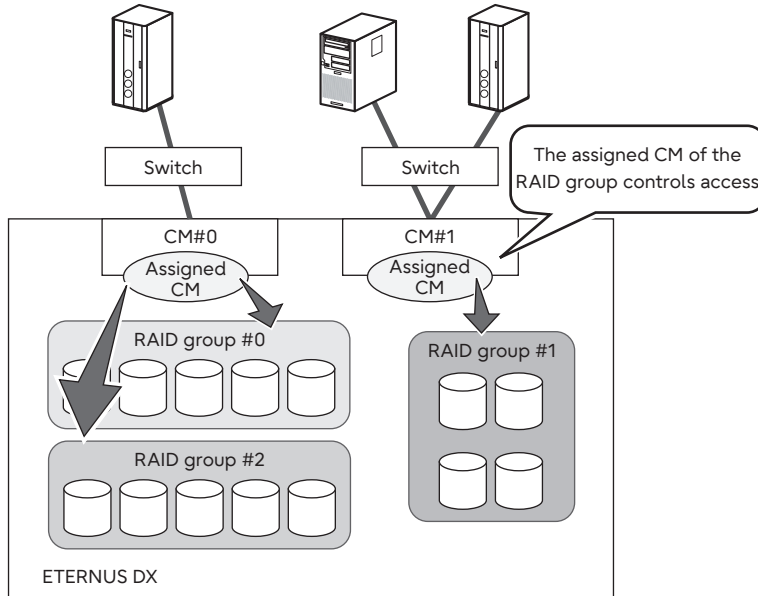
Caution

- The read/write performance of the random access can be enhanced by changing the setting, however, note that the performance can be degraded, depending on the used system.
- The following restrictions are applied to the RAID groups with expanded stripe sizes:
 - Logical Device Expansion cannot be performed on volumes that belong to the RAID group.
 - RAID groups configured with different stripe sizes cannot coexist in the same TPP pool.
 - A WSV cannot be configured by concatenating RAID groups with different stripe sizes.
- "Stripe Depth 512KB" cannot be specified for a "RAID5 (4D+1P)" configuration that is used for TPPs.
- "Stripe Depth 256KB" cannot be specified for a "RAID5 (8D+1P)" configuration that is used for TPPs.

Assigned CMs

A controller that controls access is assigned to each RAID group and manages the load balance in the ETERNUS DX. The controller that controls a RAID group is called an assigned CM.

Figure 48 Assigned CMs



When the load is unbalanced between the controllers, change the assigned CM. For more details, refer to "Change Controlling CM" in "ETERNUS Web GUI User's Guide".

If an assigned controller is disconnected for any reason, the assigned CM is replaced by another controller. After the disconnected controller is installed again and returns to normal status, this controller becomes the assigned CM again.

Stable Operation via Load Control


Host Response

The response from the ETERNUS DX can be optimized by switching the setup information of the host response for each connected server.

The server requirements of the supported functions, LUN addressing, and the method for command responses vary depending on the connection environments such as the server OS and the driver that will be used. A function that handles differences in server requirements is supported. This function can specify the appropriate operation mode for the connection environment and convert host responses that respond to the server in the ETERNUS DX.

The host response settings can be specified for the server or the port to which the server connects. For details on the settings, refer to "Configuration Guide -Server Connection-".

Figure 49 Host Response



▶ Caution

- If the host response settings are not set correctly, a volume may not be recognized or the desired performance may not be possible. Make sure to select appropriate host response settings.
- The number of LUNs that can be referenced from the host varies depending on "LUN Addressing" and "LUN Expand Mode (Peripheral Device Addressing)" of the host response settings.

Data Migration

▶ Caution

For the migration destination ETERNUS DX, select the same type and speed for the host interface as the migration source storage system.

To change the host interface, it is recommended that a validation is performed in advance to avoid operational impact after the migration.

For example, when data is migrated from an FC connected storage system to an ETERNUS DX that is connected using a different type of host interface, note that the performance may seem to be reduced due to a degraded I/O response.

Storage Migration

Storage Migration is a function that migrates the volume data from an old storage system to volumes in a new storage system without using a host in cases such as when replacing a storage system.

The migration source storage system and migration destination ETERNUS DX are connected using FC cables. Data read from the target volume in the migration source is written to the migration destination volume in the ETERNUS DX.

Since Storage Migration is controlled by ETERNUS DX controllers, no additional software is required.

The connection interface is FC. In addition, the direct connection and switch connection topologies are supported.

Online Storage Migration and offline Storage Migration are supported.


- Offline method

Stop the server during the data migration. Host access becomes available after the data migration to the migration destination volume is complete. Therefore, this method prevents host access from affecting the ETERNUS DX and can shorten the time of the migration. This method is suitable for cases requiring quick data migration.

- Online method

Host access becomes available after the data migration to the migration destination volume starts. Operations can be performed during the data migration. Therefore, this method can shorten the time for the stopped operation. This method is suitable for cases requiring continued host access during the data migration.

Figure 50 Storage Migration



The Storage Migration function migrates whole volumes at the block level. A data migration can be started by specifying a text file with migration information that is described in a dedicated format from ETERNUS Web GUI. The path between the migration source and the migration destination is called a migration path. The maximum number of migration volumes for each migration path is 512. Up to 16 migration source devices can be specified and up to eight migration paths can be created for each migration source device.

The capacity of a volume that is to be specified as the migration destination area must be larger than the migration source volume capacity.

Caution

- For online Storage Migration, the capacity of the migration destination volume must be the same as the migration source volume.
- For offline Storage Migration, stop server access to both the migration source volume and the migration destination volume during a migration.
For online Storage Migration, stop server access to the migration source volume and the migration destination volume before starting a migration. In addition, do not access the migration source volume from the server during a migration.
- Online storage migration can be manually resumed on the following volumes after the process (of deleting a copy session) is complete.
 - TPV capacity optimization is running
 - An Advanced Copy session exists
- For the migration destination device, the FC port mode needs to be switched to "Initiator" and the port parameter also needs to be set.
- Make sure to delete the migration path after Storage Migration is complete.

Non-disruptive Storage Migration

Non-disruptive Storage Migration is a function that migrates the volume data from an old storage system to volumes in a new storage system without stopping a business server in cases such as when replacing a storage system.

The connection interface between the migration source storage system (external storage system) and the migration destination storage system (local storage system) is only the FC cable. In addition, the direct connection and switch connection topologies are supported.

Figure 51 Non-disruptive Storage Migration




Table 27 Specifications for Paths and Volumes between the Local Storage System and the External Storage System

| Item | Quantity |
|---|----------|
| The maximum number of multipath connections between the local storage system and the external storage system (per external storage system) | 8 paths |
| The maximum number of ports in the external storage system that can be connected from the local storage system (per FC-Initiator port) | 32 ports |
| The maximum number of migration target volumes that can be imported to the local storage system (*1) | 512 |
| The maximum number of migration target volumes in the external storage system that can be imported simultaneously to the local storage system | 512 |

*1: The number of migration target volumes that are imported to the local storage system is added to the number of volumes in the local storage system.

Connect the external storage system to the local storage system ETERNUS DX using FC cables. After the connection is established, add multipath connections between the local storage system and the business server to prepare for the data migration.

After disconnecting the multipath connection between the external storage system and the business server, use RAID Migration to read data from the migration target volume in the external storage system and write data to the migration destination volume in the local storage system.

Data consistency is ensured because the data remains in the migration source volume for consolidated management during the data migration.

 **Caution**

- Only FC ports (connected in the FC-Initiator mode) are supported for connections with external storage systems.
 - The Non-disruptive Storage Migration License must be registered to use this function.
For details on the license, contact your sales representative.
 - Only data migrations from the external storage system to the local storage system is supported.
Data migrations from the local storage system to the external storage system or between external storage systems is not supported.
 - The local storage system volume returns the same information as the External Volume even after a migration is completed.
 - Do not set the copy operation suppression and the cache parameters for the External Volume.
 - The functions that can be used for the External Volumes are delete, rename, and RAID migration. Other functions cannot be used until the data migration is successfully completed.
 - Make sure to delete the Non-disruptive Storage Migration License after the Non-disruptive Storage Migration is complete.
-

Server Linkage Functions

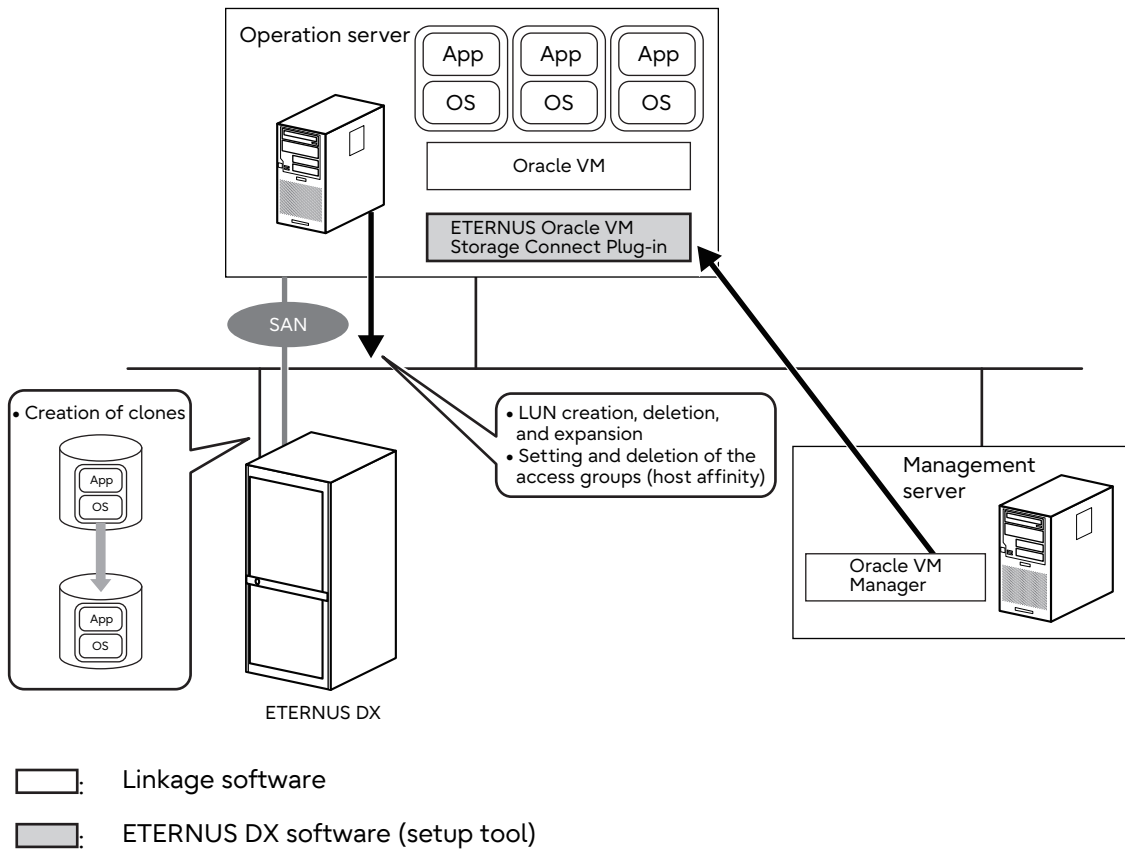
Oracle VM Linkage

"Oracle VM Manager", which is the user interface of the "Oracle VM" server environment virtualization software, can provision the ETERNUS DX.

"ETERNUS Oracle VM Storage Connect Plug-in" is required to use this function.

The Oracle VM Storage Connect framework enables Oracle VM Manager to directly use the resources and functions of the ETERNUS DX in an Oracle VM environment. Native storage services such as Logical Unit Number (LUN) creation, deletion, expansion, and snapshots are supported.

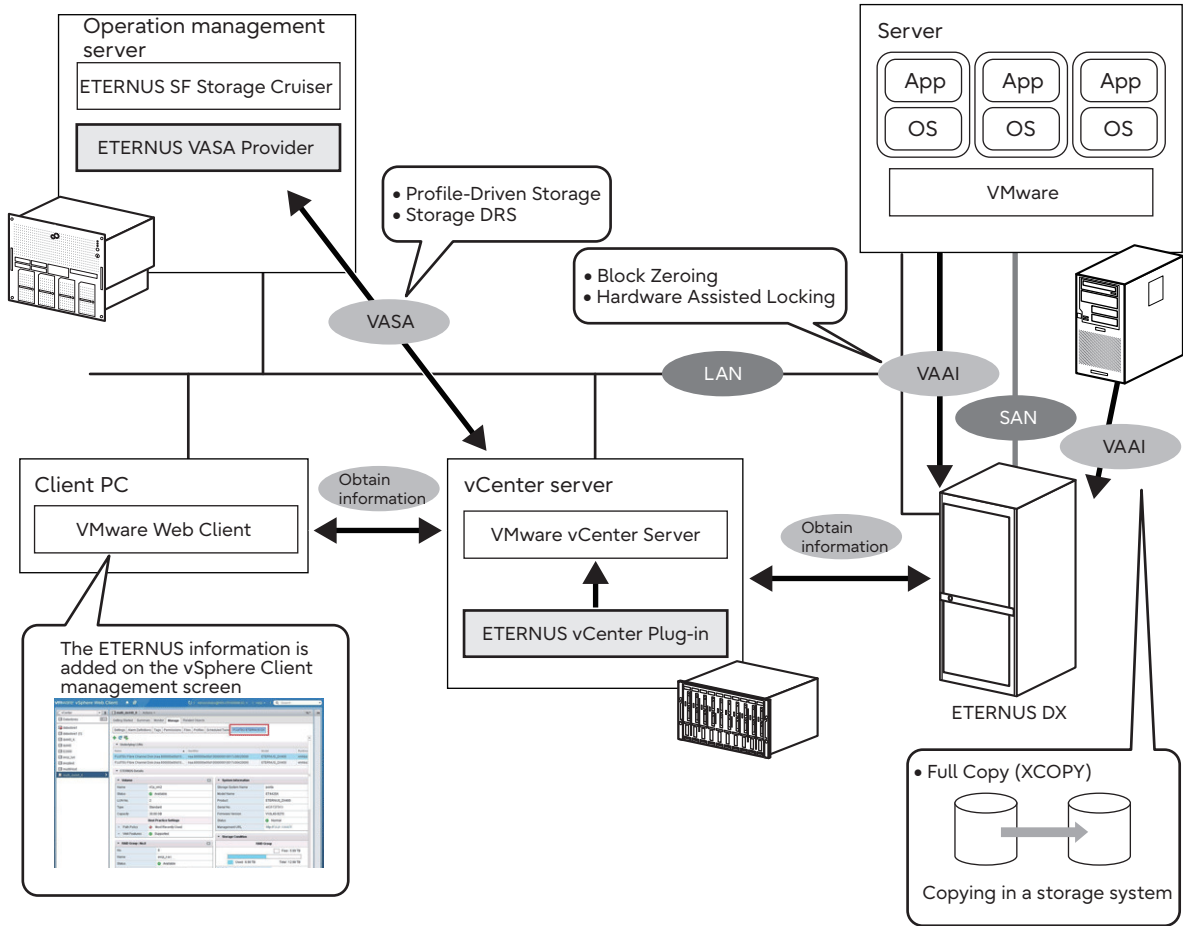
Figure 52 Oracle VM Linkage



VMware Linkage

By linking with "VMware vSphere" (which virtualizes platforms) and "VMware vCenter Server" (which supports integrated management of VMware vSphere), the resources of the ETERNUS DX can be effectively used and system performance can be improved.

Figure 53 VMware Linkage



- Linkage software
- ETERNUS DX software (setup tool)

■ VMware VASA

vStorage API for Storage Awareness (VASA) is an API that enables vCenter Server to link with the storage system and obtain storage system information. With VMware, VASA integrates the virtual infrastructure of the storage, and enhances the Distributed Resource Scheduling (DRS) function and the troubleshooting efficiency.

ETERNUS VASA Provider is required to use the VASA function.

ETERNUS VASA Provider obtains and monitors information from the ETERNUS DX by using functions of ETERNUS SF Storage Cruiser.

- Profile-Driven Storage

The Profile-Driven Storage function classifies volumes according to the service level in order to allocate virtual machines with the most suitable volumes.

- Distributed Resource Scheduler (Storage DRS)

The Storage DRS function moves original data in virtual machines to the most suitable storage area according to the access volume. Storage DRS balances the loads on multiple physical servers in order to eliminate the need for performance management on each virtual machine.

■ VMware VAAI

vStorage APIs for Array Integration (VAAI) are APIs that improve system performance and scalability by using the storage system resources more effectively.

The ETERNUS DX supports the following features.

- Full Copy (XCOPY)

Data copying processes can be performed in the ETERNUS DX without the use of a server such as when replicating or migrating the virtual machine. With Full Copy (XCOPY), the load on the servers is reduced and the system performance is improved.

- Block Zeroing

When allocating storage areas to create new virtual machines, it is necessary to zero out these storage areas for the initialization process. This process was previously performed on the server side. By performing this process on the ETERNUS DX side instead, the load on the servers is reduced and the dynamic capacity allocation (provisioning) of the virtual machines is accelerated.

- Hardware Assisted Locking

This control function enables the use of smaller blocks that are stored in the ETERNUS DX for exclusive control of specific storage areas.

Compared to LUN (logical volume) level control that is implemented in "VMware vSphere", enabling access control in block units minimizes the storage areas that have limited access using exclusive control and improves the operational efficiency of virtual machines.

■ VMware vCenter Server

- vCenter linkage

Various information of the ETERNUS DX can be displayed on vSphere Web Client by expanding the user interface of VMware Web Client. Because storage side information is more visualized, integrated management of the infrastructure under a virtual environment can be realized and usability can be improved.

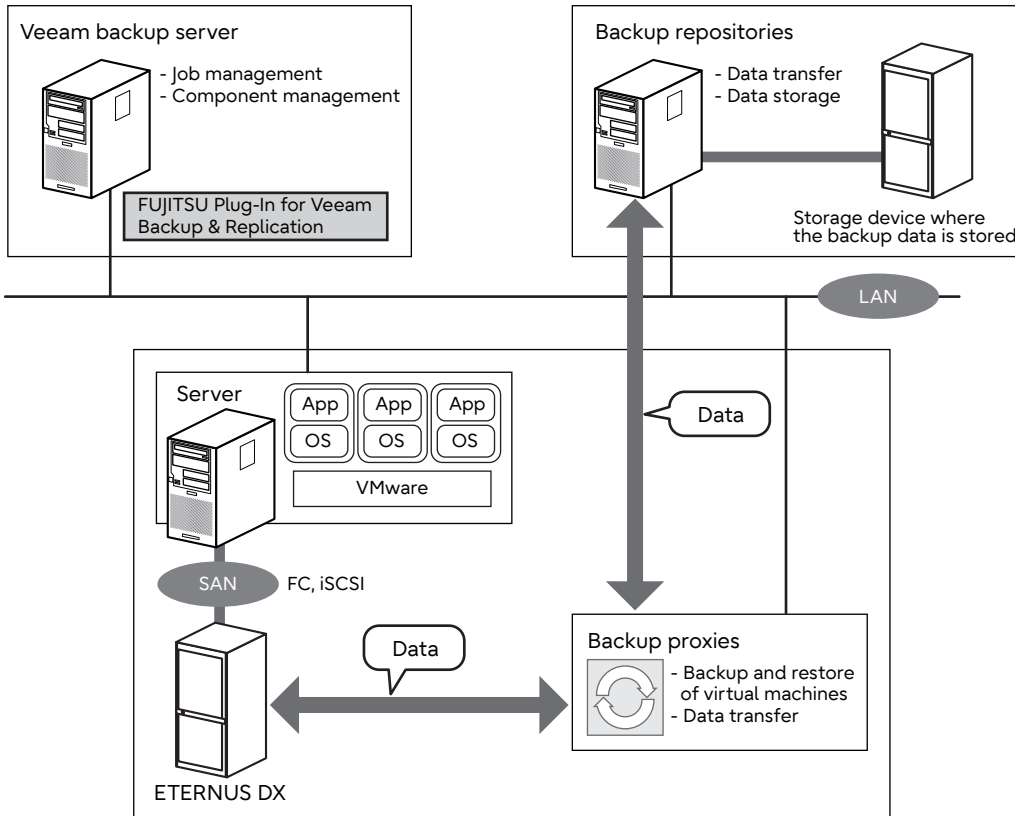
ETERNUS vCenter Plug-in is required to use this function.



Veeam Storage Integration

The operability and efficiency of Virtual Machine backups in virtual environments (VMware) are improved by using the ETERNUS DX storage snapshot integration with Veeam Backup & Replication provided by Veeam Software.

Veeam Storage Integration is available for the ETERNUS DX60 S4.

Figure 54 Veeam Storage Integration



- : Veeam Storage Integration software
- : ETERNUS DX software (setup tool)

 **Caution**

- The controller firmware version of the ETERNUS DX must be V10L86 or later.
- The Veeam Storage Integration license must be obtained and registered in the ETERNUS DX.
- iSCSI and FC host interfaces are supported in Veeam Storage Integration for the connection between backup proxies and the ETERNUS DX.
- To connect a Backup Proxy with the ETERNUS DX via an FC, the host affinity settings must be configured for the Backup Proxy using ETERNUS CLI. For more details, refer to "ETERNUS CLI User's Guide".
- For a restore process using a host that does not have any LUNs attached, host affinity settings for the host must be configured in advance. For more details, refer to "ETERNUS CLI User's Guide".
- To enable the ETERNUS DX storage snapshot integration with Veeam Backup & Replication, FUJITSU Plug-In for Veeam Backup & Replication must be installed to the Veeam backup server.
- If a volume has several snapshot generations and these snapshots have been created with different resolutions, only the oldest snapshot generation can be deleted.
- The following volumes cannot be managed or operated by Veeam Backup & Replication:
 - Volumes with Advanced Copy sessions except SnapOPC+ sessions
 - Volumes with SnapOPC+ sessions created by ETERNUS SF AdvancedCopy Manager
- Veeam Backup & Replication jobs or operations may fail during a RAID migration or a Thin Provisioning Volume balancing.
- SnapOPC+ is used for Veeam Storage Integration.

Thin Provisioning Volumes (TPVs) are used as SnapOPC+ copy destination volumes.

Configure an appropriate maximum pool capacity for the Thin Provisioning function by taking the total capacity of volumes used for Veeam Storage Integration and the number of snapshot generations into consideration.

For more details about the maximum pool capacity setting, refer to "Thin Provisioning Pool Management" in "ETERNUS Web GUI User's Guide".

Guidelines for the maximum pool capacity for the Thin Provisioning function:

Maximum pool capacity \geq total capacity of TPVs + total capacity of volumes for Veeam Storage Integration \times (number of snapshot generations + 1)

- It is not recommended to use multiple Veeam Backup & Replication for managing a single ETERNUS DX. In such configuration, an error might occur at the jobs that are in conflict with each other when being executed from multiple Veeam Backup & Replication.

Note

- Veeam Storage Integration supports the following volumes.

Table 28 Volume Types That Can Be Used with Veeam Storage Integration

| Volume type | Copy source | Copy destination |
|-------------|-------------|------------------|
| Standard | ○ | × |
| WSV | ○ | × |
| TPV | ○ | ○ |
| SDV | × | × |
| SDPV | × | × |

○: Supported ×: Not supported

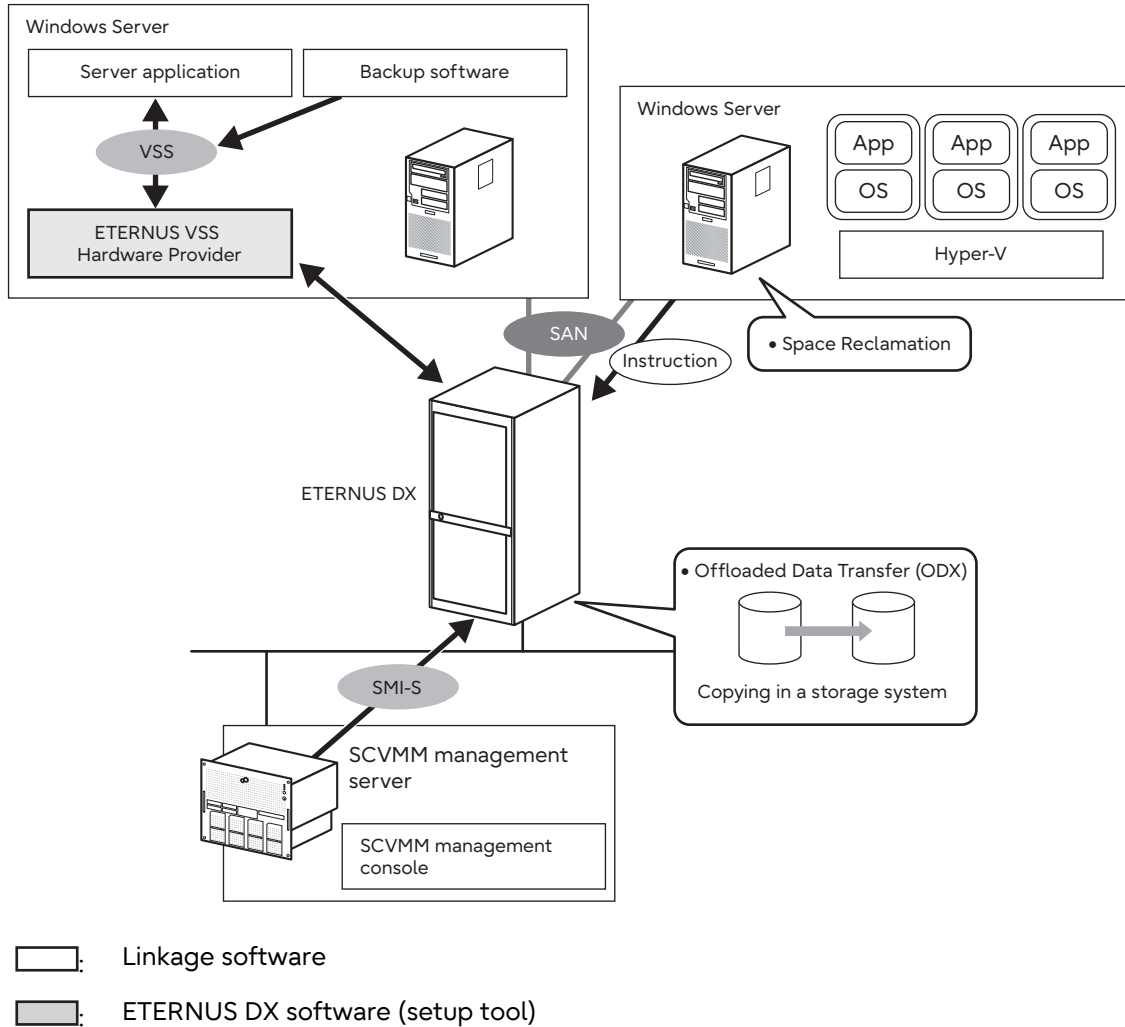
- Copy destination TPVs are automatically created when snapshots are created with Veeam Backup & Replication (*1).

- *1 :** Copy destination TPVs are automatically created in the following TPPs.
- If the copy source volume is a Standard volume or a WSV, the TPP with the smallest number in the ETERNUS DX
 - If the copy source volume is a TPV, the same TPP as the copy source
- Copy destination TPVs can also be created in any TPP using ETERNUS CLI.

Microsoft Linkage

The ETERNUS DX supports integrated management of virtualized platforms and cloud linkage by using functions in Windows Server and System Center.

Figure 55 Microsoft Linkage



■ Windows Server

The ETERNUS DX supports the following functions in Windows Server.

- Offloaded Data Transfer (ODX)

The ODX function of Windows Server 2012 or later offloads the processing load for copying and transferring files from the CPU of the server to the storage system.

- Thin Provisioning Space Reclamation

The Thin Provisioning Space Reclamation function of Windows Server 2012 or later automatically releases areas in the storage system that are no longer used by the OS or applications. A notification function for the host is provided when the amount of allocated blocks of the TPV reaches the threshold.

- Hyper-V

Hyper-V is virtualization software for Windows Server.

By using the Hyper-V virtualized Fibre Channel, direct access to the SAN environment from a guest OS can be performed. The volumes in the ETERNUS DX can be directly recognized and mounted from the guest OS.

- Volume Shadow Copy Service (VSS)

VSS is performed in combination with the backup software and the server applications that are compatible with Windows Server VSS while online backups are performed via the Advanced Copy function for the ETERNUS DX.

ETERNUS VSS Hardware Provider is required to use this function.

SnapOPC+ and QuickOPC can be used as the copy method.

▶ Caution

To use the ODX function, the controller firmware version of the ETERNUS DX must be V10L80-2000 or later.

■ System Center Virtual Machine Manager (SCVMM)

System Center is a platform to manage operations of data centers and clouds. This platform also provides an integrated tool set for the management of applications and services.

SCVMM is a component of System Center 2012 or 2016 that performs integrated management of virtualized environments. The ETERNUS DX can be managed from SCVMM by using the SMI-S functions of the ETERNUS DX.

OpenStack Linkage

ETERNUS OpenStack VolumeDriver is a program that supports linkage between the ETERNUS DX and OpenStack.

By using the VolumeDriver for the ETERNUS DX, the ETERNUS DX can be used as a Block Storage for cinder. Creating volumes in the ETERNUS DX and assigning created volumes to VM instances can be performed via an OpenStack standard interface (Horizon).

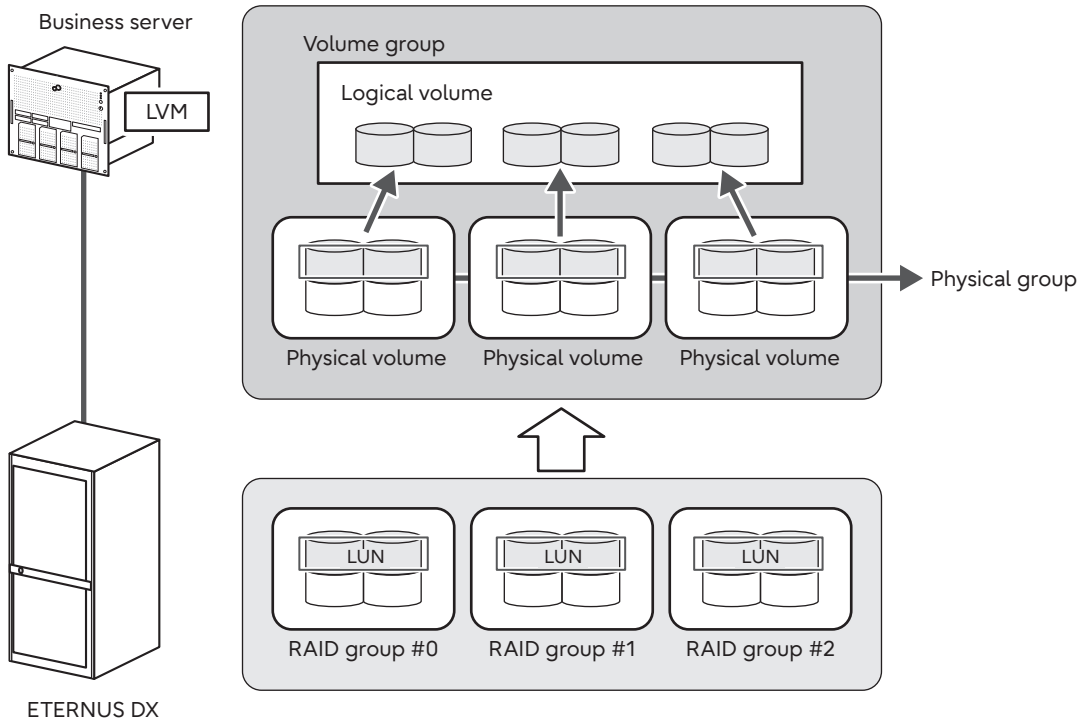
Logical Volume Manager (LVM)

The Logical Volume Manager is a management function that groups the save areas in multiple drives and partitions and manages these areas as one logical drive. Adding drives and expanding logical volumes can be performed without stopping the system. This function can be used on UNIX OSs (includes Linux).

LVM has a snapshot function. This function obtains any logical volume data as a snapshot and saves the snapshot as a different logical volume.

To use LUNs in the ETERNUS DX to configure an LVM, the LVM can be configured by registering LUNs in the ETERNUS DX as physical volumes.

Figure 56 Logical Volume Manager (LVM)



Smart Setup Wizard

The Smart Setup Wizard is a wizard that simplifies the creation of Thin Provisioning Pools and configuration of host affinity for configurations enabled with Thin Provisioning.

For the procedure on configuration using the Smart Setup Wizard, refer to "Configuration Guide (Basic)".

Note

If a Thin Provisioning Pool has not been created, the Thin Provisioning Pool configuration is automatically determined based on the type of drives and the number of drives installed in the ETERNUS DX.

- The priority for selecting drive types is as follows.

SSD > SAS > Nearline SAS

If multiple drive types exist, the drive type with the highest priority is selected to create a Thin Provisioning Pool.

To create another Thin Provisioning Pool with the unselected drive types, this wizard cannot be used. Use the dedicated function provided by this storage system to create a Thin Provisioning Pool.

- The RAID levels and the number of drives for RAID groups that configure the Thin Provisioning Pool are as follows.

| Drive type | RAID level | Number of drives |
|----------------------|------------|------------------|
| SSD | RAID5 | 5 to 48 |
| SAS and Nearline SAS | RAID6 | 7 to 48 |

- A Global Hot Spare is registered for each Thin Provisioning Pool.

The following shows an example of creating a Thin Provisioning Pool using the Smart Setup Wizard.

For SSDs

RAID groups are created with RAID5, which has high storage efficiency.

Table 29 shows a guideline for the number of drives and user capacities when 1.92TB SSDs are installed and Figure 57 shows an example RAID configuration.

Table 29 Guideline for the Number of Drives and User Capacities (When 1.92TB SSDs Are Installed)

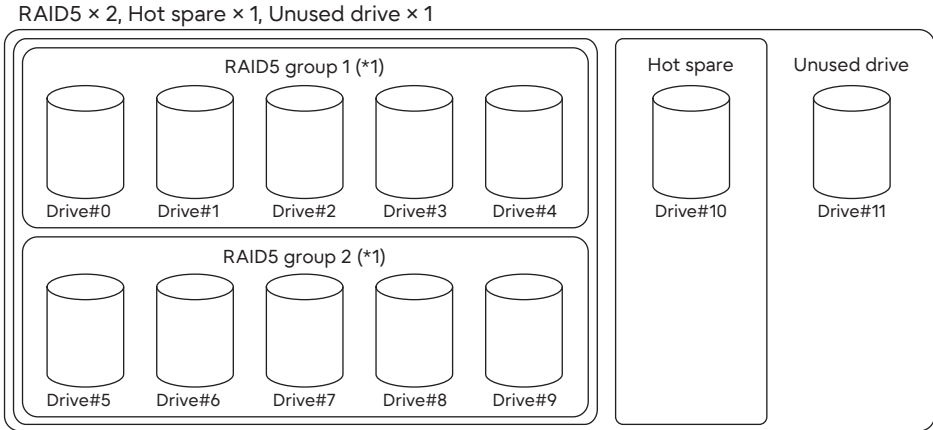
| Number of installed drives | RAID configuration that is to be created | | Hot spare | Unused drive | User capacity | |
|----------------------------|--|--|-----------|--------------|----------------|--------------------|
| | RAID group | Capacity of the user data area (equivalent number of drives) | | | Per RAID group | Per storage system |
| 4 or less | RAID groups cannot be created | - | - | - | - | - |
| 5 | RAID5 × 1 | 3 | 1 | 0 | Approx. 5.2TB | Approx. 5.2TB |
| 6 | RAID5 × 1 | 4 | 1 | 0 | Approx. 6.9TB | Approx. 6.9TB |
| 7 | RAID5 × 1 | 4 | 1 | 1 | Approx. 6.9TB | Approx. 6.9TB |
| 8 | RAID5 × 1 | 6 | 1 | 0 | Approx. 10.4TB | Approx. 10.4TB |
| 9 | RAID5 × 1 | 7 | 1 | 0 | Approx. 12.2TB | Approx. 12.2TB |

1. Function
Smart Setup Wizard

| Number of installed drives | RAID configuration that is to be created | | Hot spare | Unused drive | User capacity | |
|----------------------------|--|--|-----------|--------------|----------------|--------------------|
| | RAID group | Capacity of the user data area (equivalent number of drives) | | | Per RAID group | Per storage system |
| 10 | RAID5 × 1 | 8 | 1 | 0 | Approx. 13.9TB | Approx. 13.9TB |
| 11 | RAID5 × 2 | 8 | 1 | 0 | Approx. 6.9TB | Approx. 13.9TB |
| 12 | RAID5 × 2 | 8 | 1 | 1 | Approx. 6.9TB | Approx. 13.9TB |
| 13 | RAID5 × 3 | 9 | 1 | 0 | Approx. 5.2TB | Approx. 15.7TB |
| 14 | RAID5 × 3 | 9 | 1 | 1 | Approx. 5.2TB | Approx. 15.7TB |
| 15 | RAID5 × 2 | 12 | 1 | 0 | Approx. 10.4TB | Approx. 20.9TB |
| 16 | RAID5 × 2 | 12 | 1 | 1 | Approx. 10.4TB | Approx. 20.9TB |
| 17 | RAID5 × 2 | 14 | 1 | 0 | Approx. 12.2TB | Approx. 24.4TB |
| 18 | RAID5 × 2 | 14 | 1 | 1 | Approx. 12.2TB | Approx. 24.4TB |
| 19 | RAID5 × 2 | 16 | 1 | 0 | Approx. 13.9TB | Approx. 27.9TB |
| 20 | RAID5 × 2 | 16 | 1 | 1 | Approx. 13.9TB | Approx. 27.9TB |
| 21 | RAID5 × 4 | 16 | 1 | 0 | Approx. 6.9TB | Approx. 27.9TB |
| 22 | RAID5 × 3 | 18 | 1 | 0 | Approx. 10.4TB | Approx. 31.4TB |
| 23 | RAID5 × 3 | 18 | 1 | 1 | Approx. 10.4TB | Approx. 31.4TB |
| 24 | RAID5 × 3 | 18 | 1 | 2 | Approx. 10.4TB | Approx. 31.4TB |
| 25 | RAID5 × 3 | 21 | 1 | 0 | Approx. 12.2TB | Approx. 36.6TB |
| 26 | RAID5 × 3 | 21 | 1 | 1 | Approx. 12.2TB | Approx. 36.6TB |
| 27 | RAID5 × 3 | 21 | 1 | 2 | Approx. 12.2TB | Approx. 36.6TB |
| 28 | RAID5 × 3 | 24 | 1 | 0 | Approx. 13.9TB | Approx. 41.8TB |
| 29 | RAID5 × 4 | 24 | 1 | 0 | Approx. 10.4TB | Approx. 41.8TB |
| 30 | RAID5 × 4 | 24 | 1 | 1 | Approx. 10.4TB | Approx. 41.8TB |
| 31 | RAID5 × 6 | 24 | 1 | 0 | Approx. 6.9TB | Approx. 41.8TB |
| 32 | RAID5 × 6 | 24 | 1 | 1 | Approx. 6.9TB | Approx. 41.8TB |
| 33 | RAID5 × 4 | 28 | 1 | 0 | Approx. 12.2TB | Approx. 48.8TB |
| 34 | RAID5 × 4 | 28 | 1 | 1 | Approx. 12.2TB | Approx. 48.8TB |
| 35 | RAID5 × 4 | 28 | 1 | 2 | Approx. 12.2TB | Approx. 48.8TB |
| 36 | RAID5 × 5 | 30 | 1 | 0 | Approx. 10.4TB | Approx. 52.3TB |
| 37 | RAID5 × 4 | 32 | 1 | 0 | Approx. 13.9TB | Approx. 55.8TB |
| 38 | RAID5 × 4 | 32 | 1 | 1 | Approx. 13.9TB | Approx. 55.8TB |
| 39 | RAID5 × 4 | 32 | 1 | 2 | Approx. 13.9TB | Approx. 55.8TB |
| 40 | RAID5 × 4 | 32 | 1 | 3 | Approx. 13.9TB | Approx. 55.8TB |
| 41 | RAID5 × 5 | 35 | 1 | 0 | Approx. 12.2TB | Approx. 61.0TB |
| 42 | RAID5 × 5 | 35 | 1 | 1 | Approx. 12.2TB | Approx. 61.0TB |
| 43 | RAID5 × 6 | 36 | 1 | 0 | Approx. 10.4TB | Approx. 62.8TB |
| 44 | RAID5 × 6 | 36 | 1 | 1 | Approx. 10.4TB | Approx. 62.8TB |
| 45 | RAID5 × 6 | 36 | 1 | 2 | Approx. 10.4TB | Approx. 62.8TB |
| 46 | RAID5 × 5 | 40 | 1 | 0 | Approx. 13.9TB | Approx. 69.8TB |
| 47 | RAID5 × 5 | 40 | 1 | 1 | Approx. 13.9TB | Approx. 69.8TB |

| Number of installed drives | RAID configuration that is to be created | | Hot spare | Unused drive | User capacity | |
|----------------------------|--|--|-----------|--------------|----------------|--------------------|
| | RAID group | Capacity of the user data area (equivalent number of drives) | | | Per RAID group | Per storage system |
| 48 | RAID5 × 5 | 40 | 1 | 2 | Approx. 13.9TB | Approx. 69.8TB |

Figure 57 RAID Configuration Example (When 12 SSDs Are Installed)



*1: The capacity of the user data area is equivalent to four drives.

● For SAS Disks and Nearline SAS Disks

RAID groups are created with RAID6, which has high storage efficiency.

Table 30 shows a guideline for the number of drives and user capacities when 1.2TB SAS disks are installed and Figure 58 shows an example RAID configuration.

Table 30 Guideline for the Number of Drives and User Capacities (When 1.2TB SAS Disks Are Installed)

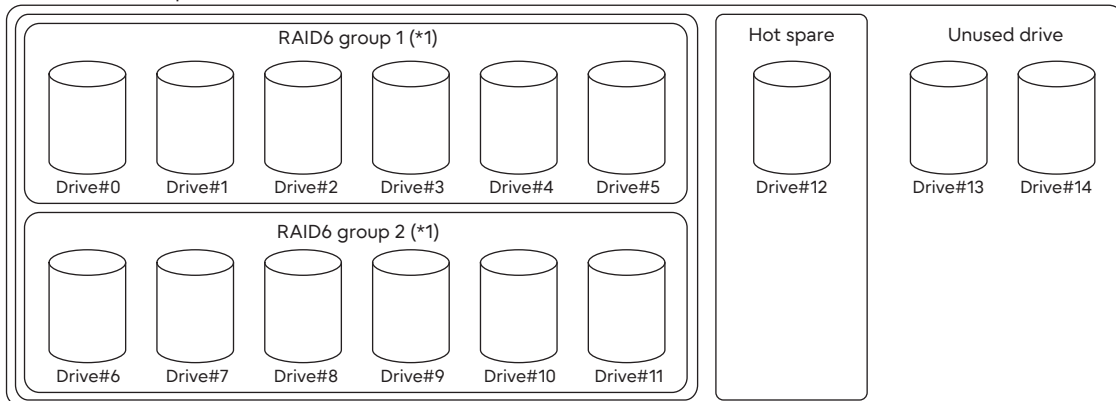
| Number of installed drives | RAID configuration that is to be created | | Hot spare | Unused drive | User capacity | |
|----------------------------|--|--|-----------|--------------|----------------|--------------------|
| | RAID group | Capacity of the user data area (equivalent number of drives) | | | Per RAID group | Per storage system |
| 6 or less | RAID groups cannot be created | - | - | - | - | - |
| 7 | RAID6 × 1 | 4 | 1 | 0 | Approx. 4.2TB | Approx. 4.2TB |
| 8 | RAID6 × 1 | 4 | 1 | 1 | Approx. 4.2TB | Approx. 4.2TB |
| 9 | RAID6 × 1 | 6 | 1 | 0 | Approx. 6.4TB | Approx. 6.4TB |
| 10 | RAID6 × 1 | 7 | 1 | 0 | Approx. 7.4TB | Approx. 7.4TB |
| 11 | RAID6 × 1 | 8 | 1 | 0 | Approx. 8.5TB | Approx. 8.5TB |
| 12 | RAID6 × 1 | 8 | 1 | 1 | Approx. 8.5TB | Approx. 8.5TB |
| 13 | RAID6 × 2 | 8 | 1 | 0 | Approx. 4.2TB | Approx. 8.5TB |
| 14 | RAID6 × 2 | 8 | 1 | 1 | Approx. 4.2TB | Approx. 8.5TB |
| 15 | RAID6 × 2 | 8 | 1 | 2 | Approx. 4.2TB | Approx. 8.5TB |

1. Function
Smart Setup Wizard

| Number of installed drives | RAID configuration that is to be created | | Hot spare | Unused drive | User capacity | |
|----------------------------|--|--|-----------|--------------|----------------|--------------------|
| | RAID group | Capacity of the user data area (equivalent number of drives) | | | Per RAID group | Per storage system |
| 16 | RAID6 × 2 | 8 | 1 | 3 | Approx. 4.2TB | Approx. 8.5TB |
| 17 | RAID6 × 2 | 12 | 1 | 0 | Approx. 6.4TB | Approx. 12.8TB |
| 18 | RAID6 × 2 | 12 | 1 | 1 | Approx. 6.4TB | Approx. 12.8TB |
| 19 | RAID6 × 2 | 14 | 1 | 0 | Approx. 7.4TB | Approx. 14.9TB |
| 20 | RAID6 × 2 | 14 | 1 | 1 | Approx. 7.4TB | Approx. 14.9TB |
| 21 | RAID6 × 2 | 16 | 1 | 0 | Approx. 8.5TB | Approx. 17.0TB |
| 22 | RAID6 × 2 | 16 | 1 | 1 | Approx. 8.5TB | Approx. 17.0TB |
| 23 | RAID6 × 2 | 16 | 1 | 2 | Approx. 8.5TB | Approx. 17.0TB |
| 24 | RAID6 × 2 | 16 | 1 | 3 | Approx. 8.5TB | Approx. 17.0TB |
| 25 | RAID6 × 3 | 18 | 1 | 0 | Approx. 6.4TB | Approx. 19.2TB |
| 26 | RAID6 × 3 | 18 | 1 | 1 | Approx. 6.4TB | Approx. 19.2TB |
| 27 | RAID6 × 3 | 18 | 1 | 2 | Approx. 6.4TB | Approx. 19.2TB |
| 28 | RAID6 × 3 | 21 | 1 | 0 | Approx. 7.4TB | Approx. 22.4TB |
| 29 | RAID6 × 3 | 21 | 1 | 1 | Approx. 7.4TB | Approx. 22.4TB |
| 30 | RAID6 × 3 | 21 | 1 | 2 | Approx. 7.4TB | Approx. 22.4TB |
| 31 | RAID6 × 3 | 24 | 1 | 0 | Approx. 8.5TB | Approx. 25.6TB |
| 32 | RAID6 × 3 | 24 | 1 | 1 | Approx. 8.5TB | Approx. 25.6TB |
| 33 | RAID6 × 4 | 24 | 1 | 0 | Approx. 6.4TB | Approx. 25.6TB |
| 34 | RAID6 × 4 | 24 | 1 | 1 | Approx. 6.4TB | Approx. 25.6TB |
| 35 | RAID6 × 4 | 24 | 1 | 2 | Approx. 6.4TB | Approx. 25.6TB |
| 36 | RAID6 × 4 | 24 | 1 | 3 | Approx. 6.4TB | Approx. 25.6TB |
| 37 | RAID6 × 4 | 28 | 1 | 0 | Approx. 7.4TB | Approx. 29.8TB |
| 38 | RAID6 × 4 | 28 | 1 | 1 | Approx. 7.4TB | Approx. 29.8TB |
| 39 | RAID6 × 4 | 28 | 1 | 2 | Approx. 7.4TB | Approx. 29.8TB |
| 40 | RAID6 × 4 | 28 | 1 | 3 | Approx. 7.4TB | Approx. 29.8TB |
| 41 | RAID6 × 4 | 32 | 1 | 0 | Approx. 8.5TB | Approx. 34.1TB |
| 42 | RAID6 × 4 | 32 | 1 | 1 | Approx. 8.5TB | Approx. 34.1TB |
| 43 | RAID6 × 4 | 32 | 1 | 2 | Approx. 8.5TB | Approx. 34.1TB |
| 44 | RAID6 × 4 | 32 | 1 | 3 | Approx. 8.5TB | Approx. 34.1TB |
| 45 | RAID6 × 4 | 32 | 1 | 4 | Approx. 8.5TB | Approx. 34.1TB |
| 46 | RAID6 × 5 | 35 | 1 | 0 | Approx. 7.4TB | Approx. 37.3TB |
| 47 | RAID6 × 5 | 35 | 1 | 1 | Approx. 7.4TB | Approx. 37.3TB |
| 48 | RAID6 × 5 | 35 | 1 | 2 | Approx. 7.4TB | Approx. 37.3TB |

Figure 58 RAID Configuration Example (When 15 SAS Disks Are Installed)

RAID6 × 2, Hot spare × 1, Unused drive × 2



***1:** The capacity of the user data area is equivalent to four drives.

2. Connection Configuration

This chapter explains the connection configuration of the ETERNUS DX.

SAN Connection

FC, iSCSI, and SAS are available as host interfaces. The server and the ETERNUS DX can be connected directly or via a switch.

Host Interface

This section describes each host interface.

The supported host interfaces vary between the ETERNUS DX60 S4 and the ETERNUS DX60 S3. For details about host interfaces, refer to "Overview" of the currently used storage systems.

When switches are used, zoning should be set for the switches to ensure the security of data.

■ Fibre Channel (FC)

The FC connection topologies that are supported are Fibre Channel Arbitrated Loop (FC-AL) and Fabric. Direct connections and switch connections to the servers are available.

The following types of host interfaces are available:

- FC 16Gbit/s

This host interface is supported only for the ETERNUS DX60 S4.

- FC 8Gbit/s

One of the following transfer rates can be specified:

- For FC 16Gbit/s
 - 16Gbit/s
 - 8Gbit/s
 - 4Gbit/s
- For FC 8Gbit/s
 - 8Gbit/s
 - 4Gbit/s

■ iSCSI

Direct connections and switch connections to servers are available.

The following types of host interfaces are available:

- iSCSI 10Gbit/s

One of the following transfer rates can be specified:

- 10Gbit/s
- 1Gbit/s

- iSCSI 1Gbit/s

In order to maintain iSCSI performance, the iSCSI network should be physically separated from other types of networks (such as networks for Internet access and file transfers).

- Operation Mode

The iSCSI 10Gbit/s operation mode is 10GBASE-T.

The iSCSI 1Gbit/s operation mode is 1000BASE-T Full Duplex (FULL).

- CHAP

CHAP authentication can prevent unauthorized access. The following CHAP authentication methods are supported:

- Unidirectional CHAP
- Bidirectional CHAP

- Tag VLAN

The tag VLAN function is supported. 16 tags (VLAN ID) can be used for each port.

- Jumbo Frame

Enabling Jumbo Frame makes data transfer more efficient by increasing the amount of data that can be transferred for each Frame.

Table 31 Ethernet Frame Capacity (Jumbo Frame Settings)

| Jumbo Frame settings | Ethernet frame capacity |
|----------------------|-------------------------|
| Enabled | Up to 9000 bytes |
| Disabled | Up to 1500 bytes |

▶ Caution

Server-side CPU load can be reduced by using Jumbo Frame. However, I/O performance may be reduced by 10% to 30%.

- Security Architecture for Internet Protocol (IPsec)

The IPsec function is not supported. Connect the server using a LAN switch that has the IPsec function as required.

- Internet Protocol

IPv4 and IPv6 are supported.

- Data Center Bridging (DCB)

iSCSI 10Gbit/s interfaces support the Data Center Bridging (DCB) function.

DCB is an enhanced function of traditional Ethernet and a standard for Fabric connections in data centers. The DCB function allows connections to Converged Enhanced Ethernet (CEE) environments.

■ SAS

Simple, cost effective, and high performance network storage environment can be configured. Direct connections and switch connections to servers are available.

The following types of host interfaces are available:

- SAS 12Gbit/s

This host interface is supported only for the ETERNUS DX60 S4.

- SAS 6Gbit/s

One of the following transfer rates can be specified:

- For SAS 12Gbit/s
 - 12Gbit/s
 - 6Gbit/s
 - 3Gbit/s
- For SAS 6Gbit/s
 - 6Gbit/s
 - 3Gbit/s
 - 1.5Gbit/s

Access Method

This section explains the connection configurations between server Host Bus Adapters (HBAs) and ETERNUS DX host interface ports.

■ Single Path Connection

A single path configuration connects the ETERNUS DX to a server via a single path.

The server cannot access an ETERNUS DX when a component (such as a controller, HBA, switch, or cable) on the path has a problem. The system must be stopped when a failed component on a path needs to be replaced or when the controller firmware needs to be updated.

In a single path connection configuration, the path failover and load balancing functions are not supported.

A multipath connection configuration is recommended to maintain availability when a problem occurs.

Figure 59 Single Path Connection (When a SAN Connection Is Used — Direct Connection)

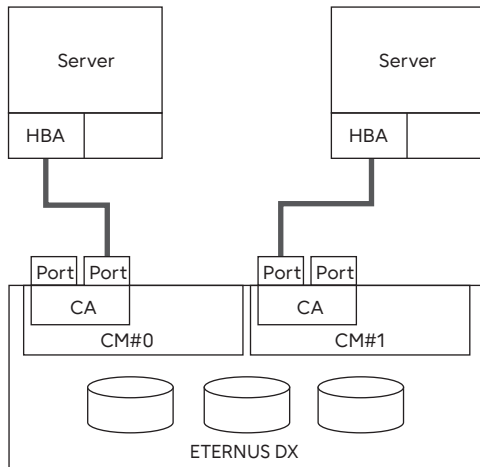



Figure 60 Single Path Connection (When a SAN Connection Is Used — Switch Connection)



■ Multipath Configuration

A multipath configuration connects the ETERNUS DX to a server via multiple paths (multipath). System reliability is improved due to the path redundancy.

If a path fails, access can continue by using the path failover function that switches access from the failed path to another path.

Figure 61 Multipath Connection (When a SAN Connection Is Used – Basic Connection Configuration)





Figure 62 Multipath Connection (When a SAN Connection Is Used – Switch Connection)



Caution

When configuring multipathing for reliability, make sure to configure a redundant connection for the controllers of the ETERNUS DX. Configure paths to connect to different controllers (CM#0 and CM#1).

■ Cluster Configuration

When servers are duplicated and connected using a cluster configuration to share a single ETERNUS DX among multiple servers, cluster control software is required.

LAN Connection

The ETERNUS DX requires a LAN connection for operation management.

In addition, information such as ETERNUS DX failures is notified to the remote support center.

Note

Make sure to connect each controller to the LAN for operation management.

Specifications for the LAN ports of the ETERNUS DX are shown below.

- Operation Mode
Ethernet (1000BASE-T/100BASE-TX/10BASE-T)
- Internet Protocol
IPv4 and IPv6 are supported.

■ IP Addresses for the ETERNUS DX

In order to connect to the LAN for operation management, an IP address for ETERNUS DX must be prepared in advance.

LAN for Operation Management (MNT Port)

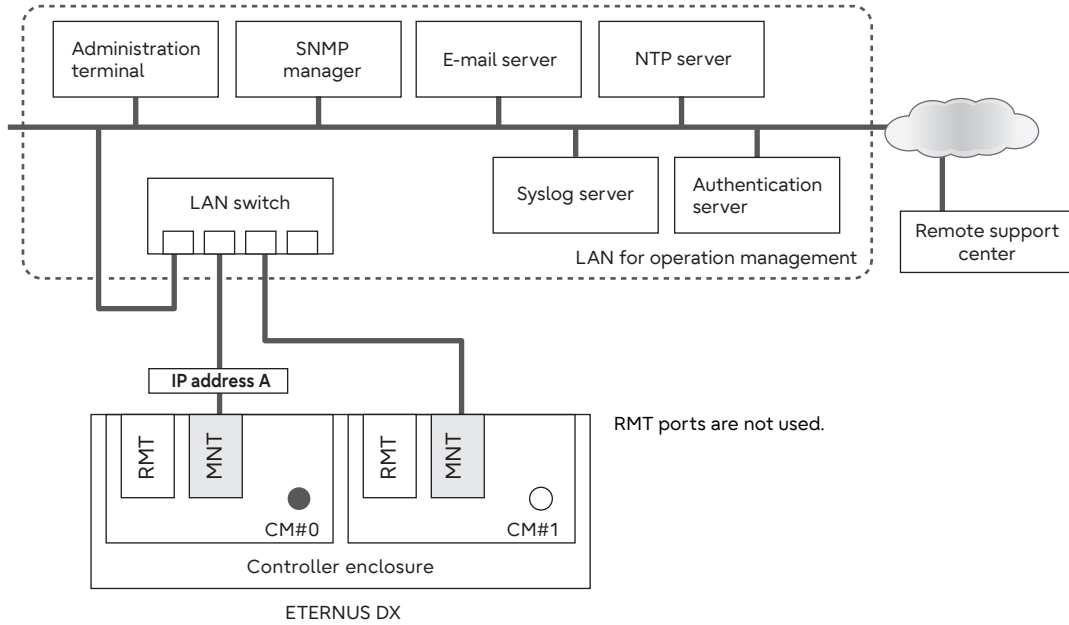
In an ETERNUS DX, the system administrator logs on to the ETERNUS DX via a LAN to set the RAID configuration, manage operations, and perform maintenance.

In addition, the failures that occur in the ETERNUS DX are notified to the remote support center. The remote support uses MNT ports for a network connection by default. In this situation, the network connection for the remote support is transferred via the LAN for operation management. When the network connection for the remote support needs to be separated from the LAN for operation management, refer to "[LAN for Remote Support \(RMT Port\)](#)" (page 108) and use the RMT ports to connect to the remote support center via a different network.

■ Connection Configurations

MNT ports are used for connecting a LAN for operation management.

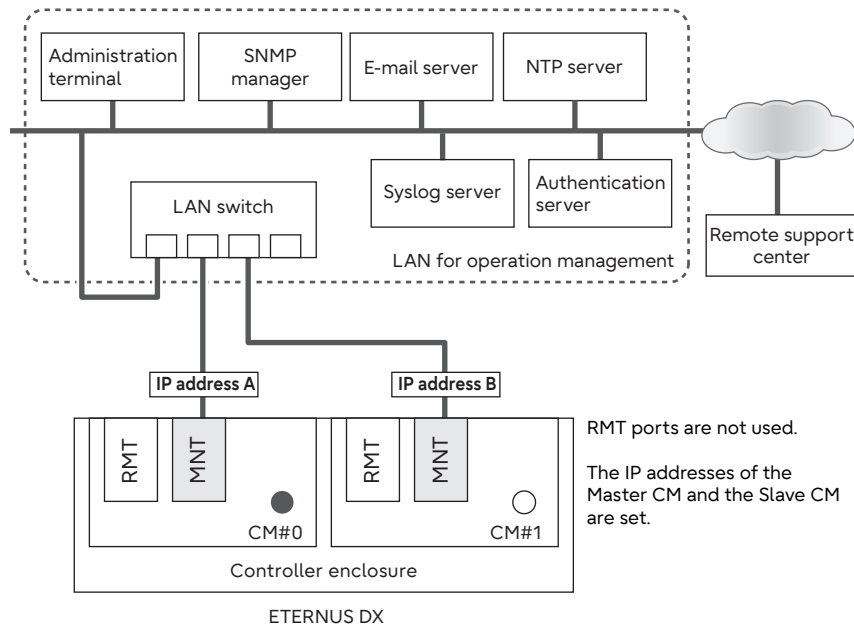
Figure 63 Connection Example without a Dedicated Remote Support Port



● Note

The following figure provides connection examples for when setting the IP address of the Slave CM.

Figure 64 Connection Example When the IP Address of the Slave CM Is Set (and a Dedicated Remote Support Port Is Not Used)



LAN for Remote Support (RMT Port)

When the network connection for the remote support needs to be separated from the company LAN, use the RMT ports to connect to the remote support center via a different network.

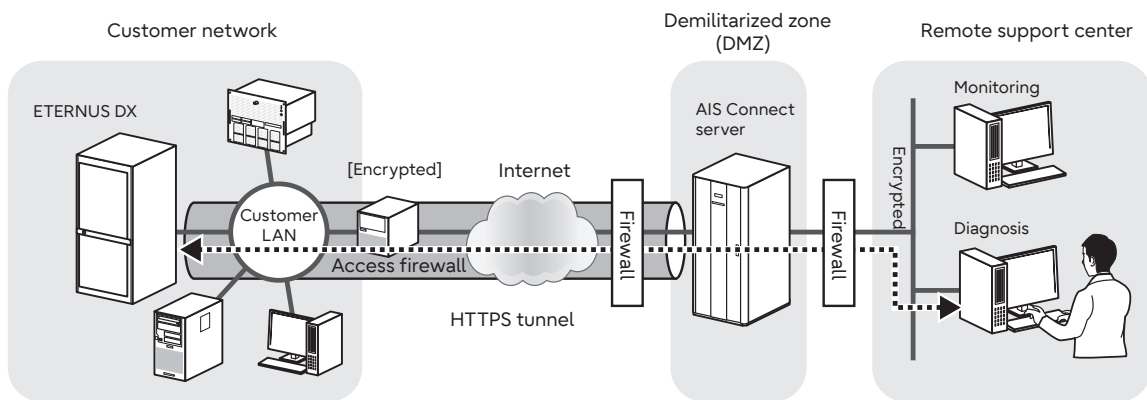
Caution

Service for AIS Connect was discontinued in December 2022.

AIS Connect Function

Overview of the AIS Connect Function

Figure 65 Overview of the AIS Connect Function



The diagram above describes the overview of the AIS Connect function. The left-hand side represents the customer and the right-hand side represents the service provider. The connection set-up initiative is always from the customer side based on regular AIS Connect agent contacts (Simple Object Access Protocol (SOAP) messages) with the AIS Connect server that can be reached via the Internet.

AIS Connect agent requests to the AIS Connect server can be handled directly and immediately. AIS Connect server requests to the AIS Connect agent (remote access) cannot be sent until the next contact has been set up. If remote access of the AIS Connect is enabled by the customer, the AIS Connect agent then executes the requests from customer operations such as setting up a tunnel for "remote access" or initiating a file transfer.


Contact set-up and request processing are performed via an HTTPS tunnel. Under certain circumstances, the firewall must be configured on the customer side to enable this type of tunnel to be set up. Likewise, proxies (plus ID and password) can be specified during Internet access configuration.

AIS Connect agent can perform the following actions:

- Notifying events (Information event, Warning event, or Error event) in the ETERNUS DX to the AIS Connect server
- Sending ETERNUS DX logs to the AIS Connect server
- Remote access from the AIS Connect server to the ETERNUS DX

● Security Features

Figure 66 Security Features



Connection from an AIS Connect agent to the AIS Connect server can be set up via a SOAP message that is based on HTTPS. Access can only be initiated by an AIS Connect agent at the ETERNUS DX site of the customer as illustrated by the diagram above. The AIS Connect server offers a certificate and the AIS Connect agent verifies this certificate for every connection setup. All transferred data is protected against spy-out and manipulation.

For setup procedure for remote support (by AIS Connect), refer to "Configuration Guide (Basic)".

▶ **Caution**

In some regions, the usage of AIS Connect is limited to contract customers.
Contact the Support Department for details.

■ REMCS

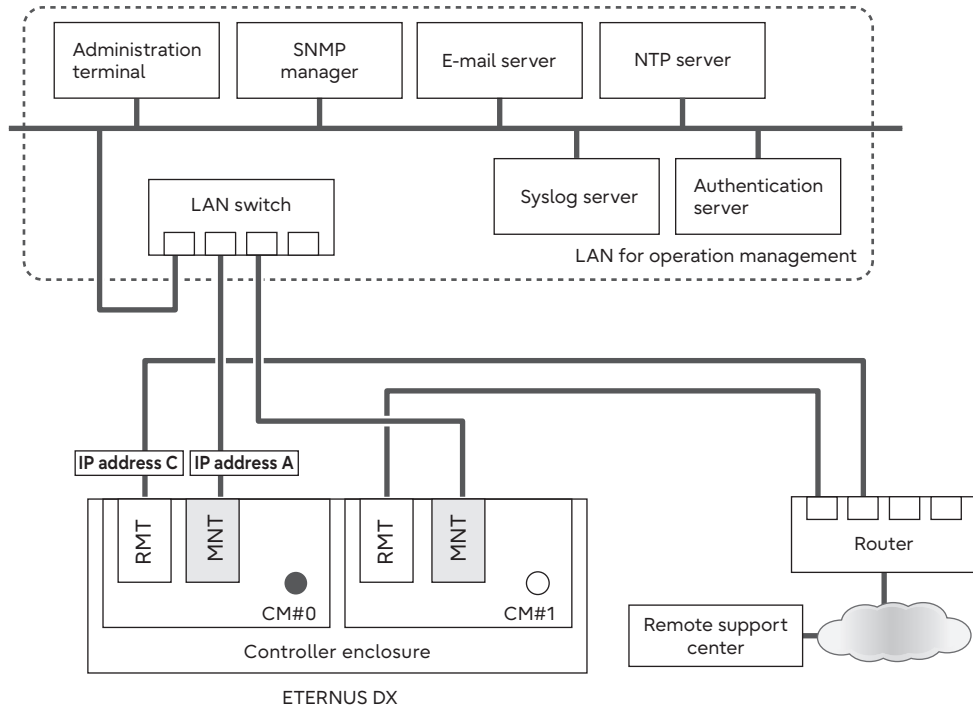
For setup procedure for remote support (by REMCS), refer to "Configuration Guide (Basic)".

The sections that are shown below explain how to set the ETERNUS DX for remote support. For details on the settings, refer to "Configuration Guide (Web GUI)".

■ Connection Configurations

For the ETERNUS DX, two IP addresses are required (one IP address for the MNT port and one IP address for the RMT port).

Figure 67 Connection Example with a Dedicated Remote Support Port



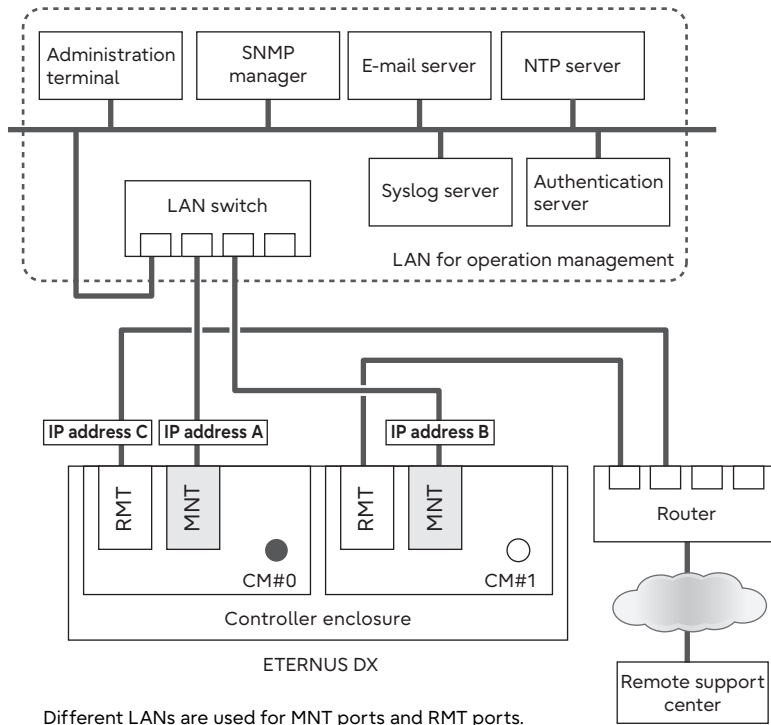
Different LANs are used for MNT ports and RMT ports.

Note

The following figure provides connection examples and explains the necessary preparation for when the IP address of the Slave CM is set.

For the ETERNUS DX, three IP addresses are required (two IP addresses for the MNT ports and one IP address for the RMT port).

Figure 68 Connection Example When the IP Address of the Slave CM Is Set (and a Dedicated Remote Support Port Is Used)



LAN Control (Master CM/Slave CM)

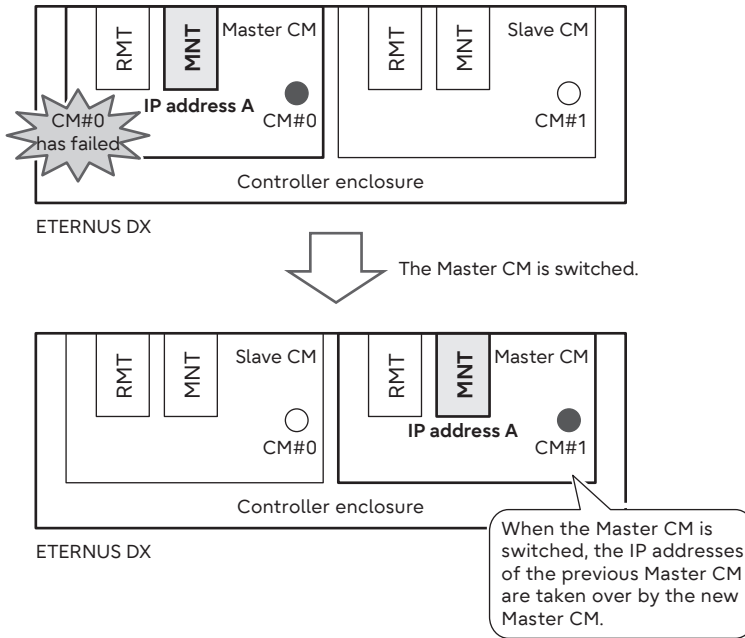
This section explains how the LAN control controller of the ETERNUS DX operates.

When an ETERNUS DX has two controllers, the controller (CM) that is given the authority to manage the LAN is called the Master CM and the other CM is called the Slave CM.

When an error occurs in the Master CM or LAN, the Master CM is switched automatically.

IP addresses of the LAN ports are not assigned to each CM. These IP addresses are assigned to the role of master or slave. If the Master CM is switched, the same IP addresses are reused. Therefore, even if the Master CM is switched and the physical port is changed, access can be maintained via the same IP addresses. The MAC address is not inherited.

Figure 69 LAN Control (Switching of the Master CM)



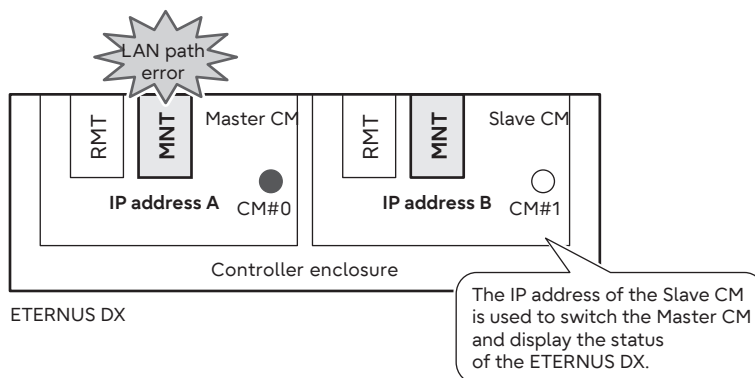
Note

- Each CM has an LED that lights up green to identify when it is the Master CM.
- Setting the IP address of the Slave CM ensures that ETERNUS Web GUI or ETERNUS CLI can be used from the Slave CM if an error occurs on the LAN path for the Master CM.

The Master CM and the Slave CM perform different functions. The Slave CM can only switch the Master CM and display the status of the ETERNUS DX.

The IP address of the Slave CM does not need to be set for normal operation.

Figure 70 LAN Control (When the IP Address of the Slave CM Is Set)



Network Communication Protocols

The usable LAN ports and functions are different depending on the usage and protocol.
The following table shows how the LAN ports may be used (by usage and protocol).

Table 32 LAN Port Availability

| Usage | Protocol | tcp / udp | Port number | Direc- tion | Master CM | | Slave CM | | Remarks | |
|-----------------|-------------------|--------------|----------------|----------------|-----------|--------|----------|--------|---|---|
| | | | | | MNT | RMT | MNT | RMT | | |
| ETERNUS Web GUI | http / https | tcp | 80 / 443 | from | ○ | ○ | △ (*1) | △ (*1) | Ac- cessed from a Web browser | |
| ETERNUS CLI | telnet / ssh | tcp | 23 / 22 | from | ○ | ○ | △ (*1) | △ (*1) | — | |
| | ftp (cli- ent) | tcp | 21 | to | ○ | ○ | △ (*1) | △ (*1) | — | |
| SNMP | agent | snmp | udp | 161 | from | ○ | ○ | ○ | ○ | — |
| | trap | snmp trap | udp | 162 (*2) | to | ○ (*3) | ○ (*3) | × | × | — |
| SMI-S | http / https | tcp | 5988 / 5989 | from | ○ | × | × | × | Used for SMI-S client commu- nication | |
| | http / https | tcp | — | to | ○ | × | × | × | Used for event commu- nications with the SMI-S listener, etc. | |
| | SLP | udp | 427 | to | ○ | × | × | × | Used for service inquiry commu- nication from the SMI-S client | |
| | SLP | tcp | 5993 | from | ○ | × | × | × | | |
| E-mail | smtp (client) | tcp | 25 (*2) | to | ○ (*3) | ○ (*3) | × | × | Used for failure notifica- tion, etc. | |
| NTP | NTP (cli- ent) | udp | 123 | to | ○ (*3) | ○ (*3) | × | × | — | |

| Usage | Protocol | tcp / udp | Port number | Direction | Master CM | | Slave CM | | Remarks |
|---|-----------------|-----------|-----------------|-----------|-----------|--------|----------|-----|-------------------------------------|
| | | | | | MNT | RMT | MNT | RMT | |
| REMCS (remote support) | smtp (client) | tcp | 25 (*2) | to | ○ (*3) | ○ (*3) | × | × | Used for failure notification, etc. |
| | pop3 (client) | tcp | 110 (*2) | to | ○ (*3) | ○ (*3) | × | × | Used for SMTP authentication |
| | http (client) | tcp | 80 or 8080 (*2) | to | ○ (*3) | ○ (*3) | × | × | Used for firmware download, etc. |
| AIS Connect (remote support) (*4) | https (client) | tcp | 443 | to | ○ (*3) | ○ (*3) | × | × | — |
| Syslog (event notification and audit log sending) | Syslog | udp | 514 (*2) | to | ○ (*3) | ○ (*3) | × | × | — |
| RADIUS | Radius | udp | 1812 (*2) | to | ○ (*3) | ○ (*3) | × | × | — |
| ping | ICMP | — | — | from | ○ (*3) | ○ (*3) | × | × | — |
| ETERNUS DX Discovery | Unique protocol | udp | 9686 | from | ○ | × | × | × | — |

○: Available / △: Available for some functions / ×: Not available

***1:** Only the following functions are available:

- Checking the ETERNUS DX status
- Switching the Master CM

***2:** Modifiable

***3:** May use either the MNT port or RMT port.

***4:** Service for AIS Connect was discontinued in December 2022.

For details on the port numbers for the ETERNUS SF Software, refer to the manual of each Storage Foundation Software ETERNUS SF.

Power Supply Connection

Connect the power cords (AC cables) of the ETERNUS DX to the power sockets, the UPS sockets, or the power control unit sockets.

For details about the types and the number of power sockets that can be used, refer to "Power Socket Specifications" in "Site Planning Guide".

Input Power Supply Lines

For details on input power supply lines, refer to "Site Planning Guide".

UPS Connection

It is recommended that an Uninterruptible Power Supply System (UPS) is used as the power supply source for the ETERNUS DX to cope with power outages and momentary voltage drops in the normal power supply.

Note that when connecting an ETERNUS DX to a single UPS, the total value of the power requirements for all the enclosures must not exceed the UPS output capacity.

When one of the power supply lines fails for redundantly configured UPS units with two power supply lines, all of the power for the ETERNUS DX must be supplied from the other line. Select a UPS that can supply sufficient power so that the total value of the power requirements does not exceed the UPS output capacity in case only one power supply line is available.

For details about the necessary UPS output capacity, refer to the specifications of the UPS that is used.

A UPS must satisfy the following conditions:

- **Rating Capacity**

Secure a sufficient rating capacity for the total value of the maximum power requirements for the enclosures that are to be installed.

To find the maximum power level requirements for each enclosure, refer to "Installation Specifications" in "Site Planning Guide".

- **Supply Time**

The total time for the server to shut down and for the ETERNUS DX to power off must be taken into consideration for the power supply time of the battery.

- **Switching Time during Power Failure**

The normal power supply must be switched to UPS output within 15ms after a power failure occurs.

- **Socket Type**

If the power plug type and the socket type of the UPS do not match, it is necessary to equip the UPS (AC output) with an appropriate type of socket. Request a qualified electrician to perform the necessary work to make terminal block connections available.

- **Power Supply Configuration**

If a UPS is used, make sure that it supplies power to all the enclosures.

Configurations where the controller enclosure is powered by the UPS while the drive enclosures are powered directly from AC are not supported.

Power Synchronized Connections

This section describes the connections to automatically control powering the ETERNUS DX on and off with a server.

In order to control powering the ETERNUS DX on and off with servers, the power control of the ETERNUS DX must be linked with all of the connected servers.

Power Synchronized Connections (PWC)

■ Power Synchronized Unit

A power synchronized unit enables the ETERNUS DX to be powered on and off with a server. The power synchronized unit detects changes in the AC power output of a UPS unit that is connected to a server (server UPS unit) and automatically turns on and off the ETERNUS DX. In addition to server UPS units, units that control the AC socket power output can also be connected. When three or more servers are connected, power can be synchronized by adding an AC sensor unit.

● Power Synchronization via a Server UPS Connection

The power synchronized unit detects the AC power output of the target devices for power synchronization and commands the ETERNUS DX to synchronize the power with the target devices.

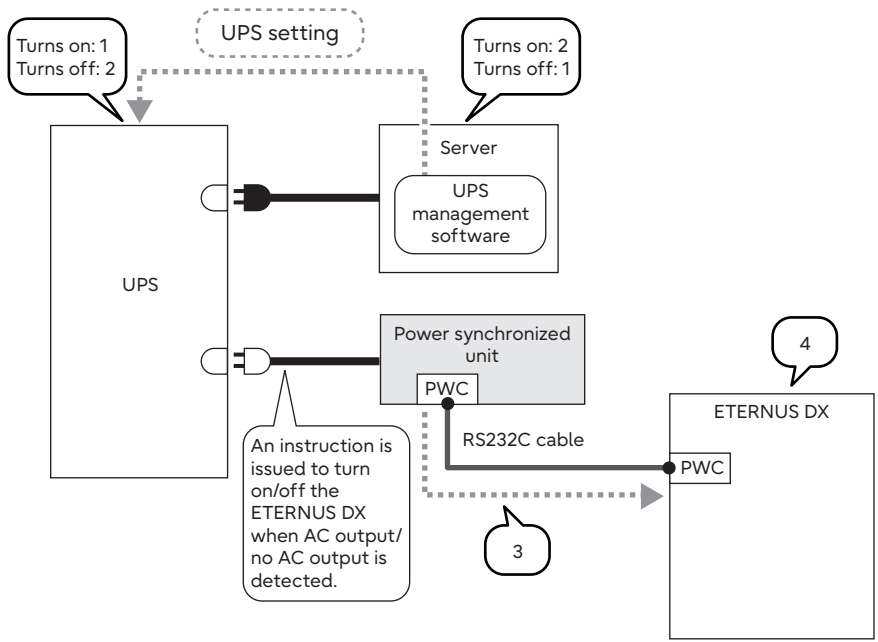
When the power synchronized unit detects the AC power output of any server UPS unit, the power synchronized unit commands the ETERNUS DX to turn on.

When the power synchronized unit does not detect AC power output in any of the server UPS units, the power synchronized unit commands the ETERNUS DX to turn off.

The server UPS unit must have a function in the management software that controls the AC power output according to when the server powers on and off. The server UPS unit must have one unused outlet to connect to the power synchronized unit.

- When connecting one or two servers

Figure 71 Power Supply Control Using a Power Synchronized Unit (When Connecting One or Two Servers)



Powering on sequence

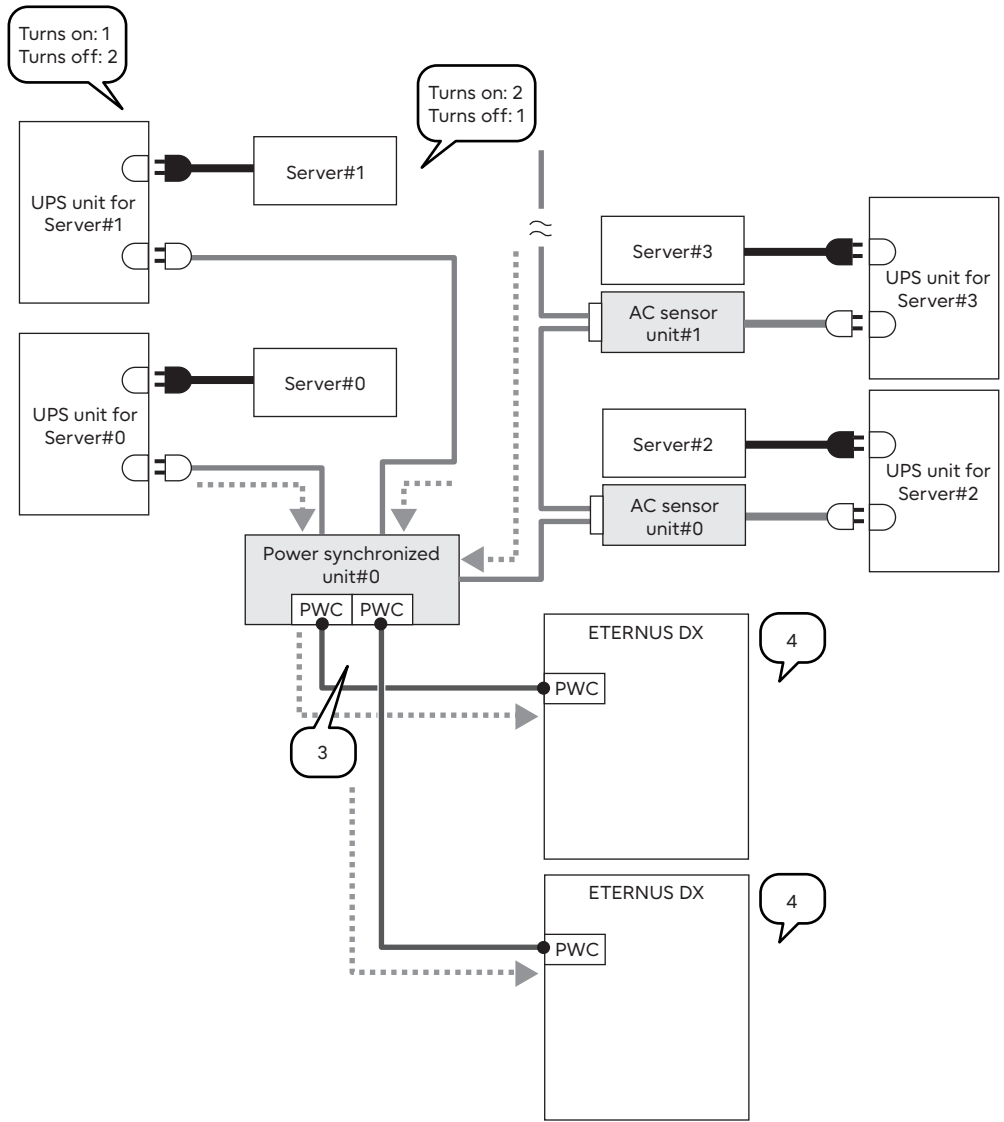
- 1 Power on of all the server UPS units
 - 2 Server startup
The server OS startup is suspended until the ETERNUS DX startup is complete (*1).
 - 3 Issuance of command to turn on the ETERNUS DX from the power synchronized unit
 - 4 ETERNUS DX startup
- *1: The server must be set to suspend server OS startup until the ETERNUS DX startup is complete.

Powering off sequence

- 1 Shutdown of all the servers
- 2 Shutdown of all the server UPS units
- 3 Issuance of command to turn off the ETERNUS DX from the power synchronized unit
- 4 ETERNUS DX shutdown

- When connecting three or more servers

Figure 72 Power Supply Control Using a Power Synchronized Unit (When Connecting Three or More Servers)



Powering on sequence

- 1 Power on of all the server UPS units
- 2 Server startup
The server OS startup is suspended until the ETERNUS DX startup is complete (*1).
- 3 Issuance of command to turn on the ETERNUS DX from the power synchronized unit
- 4 ETERNUS DX startup

*1: The server must be set to suspend server OS startup until the ETERNUS DX startup is complete.

Powering off sequence

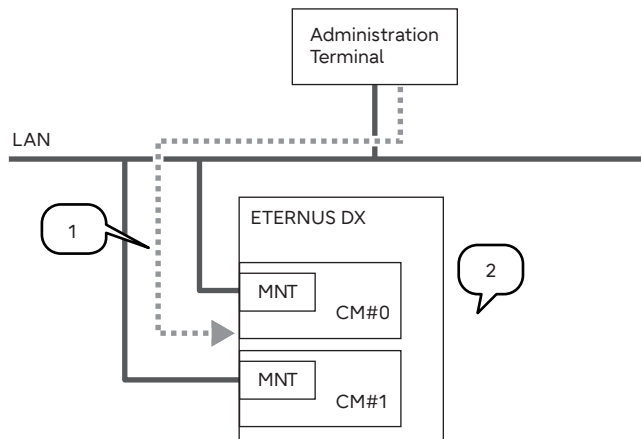
- 1 Shutdown of all the servers
- 2 Shutdown of all the server UPS units
- 3 Issuance of command to turn off the ETERNUS DX from the power synchronized unit
- 4 ETERNUS DX shutdown

Refer to the manual that is provided with a power synchronized unit for details about connection configurations with power synchronized units and required settings.

Power Synchronized Connections (Wake On LAN)

By using Wake On LAN, an instruction to power on the storage system can be issued via the LAN. "Magic packet" is sent by the utility software for Wake On LAN. The ETERNUS DX detects this packet and the power is turned on.

Figure 73 Power Supply Control Using Wake On LAN



Powering on sequence

- 1 A power on instruction is issued to the ETERNUS DX via the LAN.
- 2 ETERNUS DX startup

3. Hardware Configurations

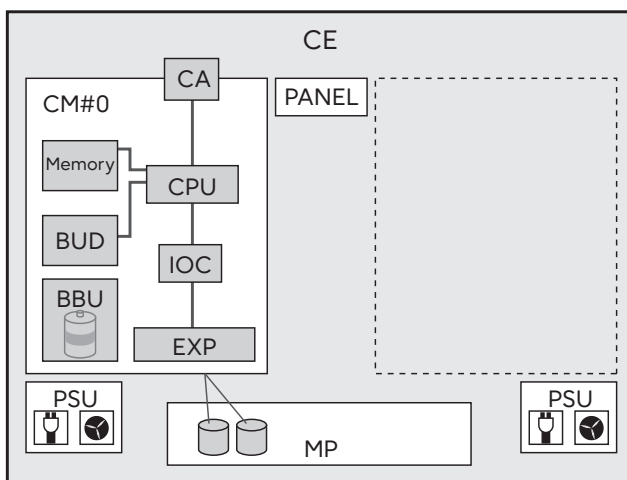
Multiple options can be selected in the ETERNUS DX according to the customer's requirements. This chapter describes the installation conditions and standard installation rules for each component.

Configuration Schematics

The following diagrams show minimum and maximum configurations for ETERNUS DX storage systems.

■ Minimum Configuration

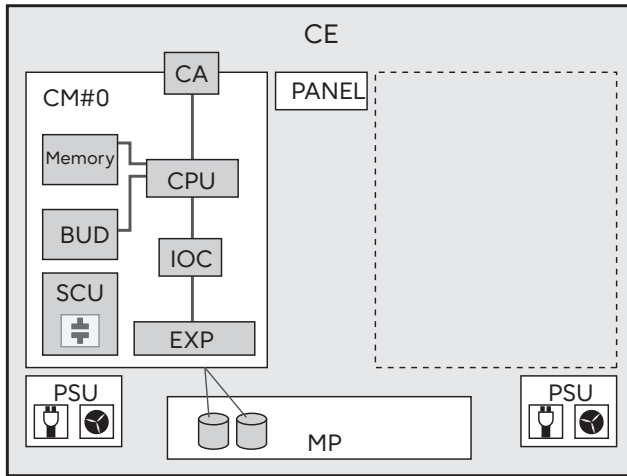
Figure 74 Minimum Configuration Diagram: ETERNUS DX60 S4



- PANEL: Operation Panel
- CE: Controller Enclosure
- CM: Controller Module
- CA: Channel Adapter
Host Interface
- Memory: System Memory
- BUD: Bootup and Utility Device
 - Backup area in case of power outage
 - Storage area for firmware
- BBU: Battery Backup Unit
Backup power source in case of power outage
- IOC: I/O Controller
Controller to control I/O
- EXP: SAS Expander
Expander chip for SAS connections
- MP: Mid Plane
Board that is located between the front of the enclosure and the rear of the enclosure (the controller (CM) or I/O module (IOM) side)

PSU: Power Supply Unit

Figure 75 Minimum Configuration Diagram: ETERNUS DX60 S3



PANEL: Operation Panel

CE: Controller Enclosure

CM: Controller Module

CA: Channel Adapter
Host Interface

Memory: Cache Memory

BUD: Bootup and Utility Device

- Backup area in case of power outage
- Storage area for firmware

SCU: System Capacitor Unit
Backup power source in case of power outage

IOC: I/O Controller
Controller to control I/O


EXP: SAS Expander
Expander chip for SAS connections

MP: Mid Plane
Board that is located between the front of the enclosure and the rear of the enclosure (the controller (CM) or I/O module (IOM) side)

PSU: Power Supply Unit

■ Maximum Configuration

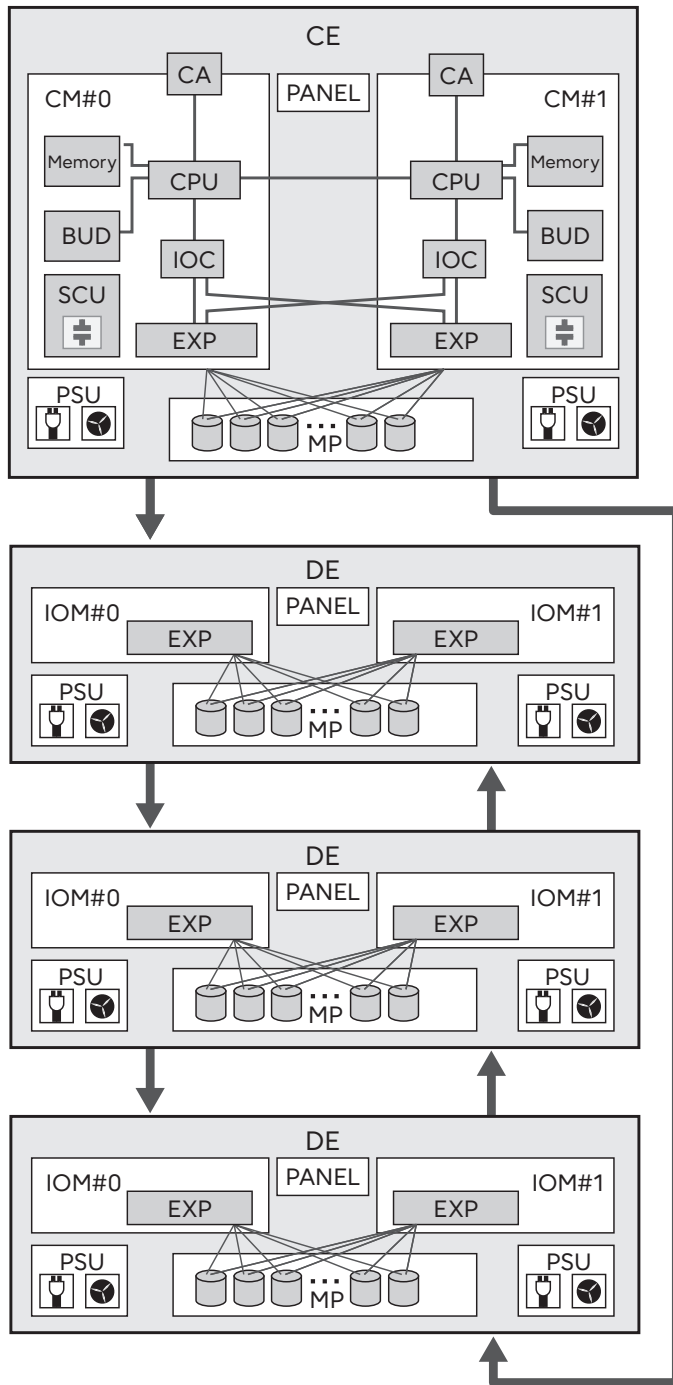
Figure 76 Maximum Configuration Diagram: ETERNUS DX60 S4



- PANEL: Operation Panel
- CE: Controller Enclosure
- CM: Controller Module
- CA: Channel Adapter
Host Interface

| | |
|---------|---|
| Memory: | System Memory |
| BUD: | Bootup and Utility Device <ul style="list-style-type: none">• Backup area in case of power outage• Storage area for firmware |
| BBU: | Battery Backup Unit Backup power source in case of power outage |
| IOC: | I/O Controller Controller to control I/O |
| EXP: | SAS Expander Expander chip for SAS connections |
| DE: | Drive Enclosure |
| IOM: | I/O Module Unit that controls I/O between controllers and drives |
| MP: | Mid Plane Board that is located between the front of the enclosure and the rear of the enclosure (the controller (CM) or I/O module (IOM) side) |
| PSU: | Power Supply Unit |

Figure 77 Maximum Configuration Diagram: ETERNUS DX60 S3



- PANEL: Operation Panel
- CE: Controller Enclosure
- CM: Controller Module
- CA: Channel Adapter
Host Interface
- Memory: Cache Memory

| | |
|------|---|
| BUD: | Bootup and Utility Device <ul style="list-style-type: none">• Backup area in case of power outage• Storage area for firmware |
| SCU: | System Capacitor Unit Backup power source in case of power outage |
| IOC: | I/O Controller Controller to control I/O |
| EXP: | SAS Expander Expander chip for SAS connections |
| DE: | Drive Enclosure |
| IOM: | I/O Module Unit that controls I/O between controllers and drives |
| MP: | Mid Plane Board that is located between the front of the enclosure and the rear of the enclosure (the controller (CM) or I/O module (IOM) side) |
| PSU: | Power Supply Unit |

■ Enclosure Connection Path


In the ETERNUS DX, multiple paths are used to connect a controller enclosure (CE) to drive enclosures (DE).

A drive enclosure has two independent drive interface ports. Path redundancy is maintained by connecting the drive enclosure to two controllers directly. This configuration allows operation to continue even if one of the connection paths fails.

Connect the controller enclosure to the drive enclosure as the following figures show.

Connection paths are not duplicated when only one controller is installed.

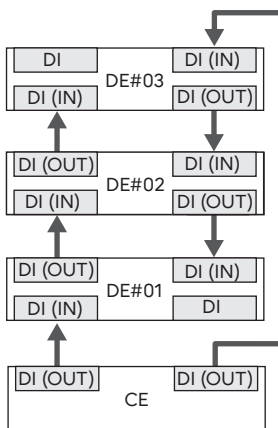
Figure 78 Enclosure Connection Path (When Only One Controller Is Installed)



DI (OUT): Drive interface (OUT) port

DI (IN): Drive interface (IN) port

Figure 79 Enclosure Connection Path (When Two Controllers Are Installed)



DI (OUT): Drive interface (OUT) port

DI (IN): Drive interface (IN) port

Optional Product Installation Conditions

This section provides the types and number of optional products and explains policies for optional product estimations.

Installation conditions for the following optional products are described in this section.

- Controller modules
- Drive enclosures
- I/O modules
- Drives

Note

To install the following optional products, a specific firmware version may be required for the ETERNUS DX.

For details, refer to the "Firmware Release Information" section in "Overview" or refer to "Product List".

- ETERNUS DX60 S4
 - Controller modules (SAS 12Gbit/s)
- ETERNUS DX60 S3
 - Controller modules (iSCSI 10Gbit/s)
 - Controller modules (FC 16Gbit/s)
 - Drive enclosures
 - SSDs

Controller Module

This section explains the installation conditions for controller modules.

■ Types of Controller Modules

Host interfaces are installed in the controller.

The following types of host interfaces are available:

● ETERNUS DX60 S4

- FC 16Gbit/s
- FC 8Gbit/s
- iSCSI 10Gbit/s
- iSCSI 1Gbit/s
- SAS 12Gbit/s
- SAS 6Gbit/s

● ETERNUS DX60 S3

- FC 8Gbit/s
- iSCSI 10Gbit/s
- iSCSI 1Gbit/s

- SAS 6Gbit/s

The features of each host interface are described below.

- **FC**

Fibre Channel (FC) enables high speed data transfer over long distances by using optical fibers and coaxial cables. FC is used for database servers where enhanced scalability and high performance are required.

- **iSCSI**

iSCSI is a communication protocol that transfers SCSI commands by encapsulating them in IP packets over Ethernet.

Since this interface can be installed at a lower cost and the network configuration is easier to change than FC, this interface is commonly used by divisions of large companies and by small and medium-sized companies where scalability and cost-effectiveness are valued over performance.

- **SAS**

Serial Attached SCSI (SAS) enables high speed data transfers by upgrading existing reliable SCSI connection standards to allow serial transfers. SAS is commonly used for small-sized systems where performance and cost-effectiveness are valued over scalability.

■ Number of Installable Controller Modules

Up to two controller modules can be installed in a controller enclosure.

Drive Enclosures

This section explains the installation conditions for drive enclosures.

■ Types of Drive Enclosures

There are two types of drive enclosures that correspond to available drive sizes (2.5" and 3.5").

Twenty-four 2.5" drives can be installed in a single 2.5" type drive enclosure.

Twelve 3.5" drives can be installed in a single 3.5" type drive enclosure.

■ Number of Installable Drive Enclosures

Up to three drive enclosures can be installed.

I/O Module

Up to two I/O modules can be installed in a drive enclosures.

Drives

This section explains the installation conditions for drives.

■ Types of Drives

The following drive types can be installed in the ETERNUS DX.

- SAS disks
- Nearline SAS disks
- SSDs

2.5" and 3.5" drive sizes are available.

Since 2.5" drives are lighter and require less power than 3.5" drives, the total weight and power consumption when 2.5" drives are installed is less than when the same number of 3.5" drives is installed.

When the data I/O count is compared based on the number of drives in an enclosure (2.5" drives: 24, 3.5" drives: 12), the Input Output Per Second (IOPS) performance for each enclosure in a 2.5" drive configuration is superior to a 3.5" drive configuration since more 2.5" drives can be installed in an enclosure than 3.5" drives.

● SAS Disks

SAS disks are high-performance and high-reliability disks. They are used to store high performance databases and other frequently accessed data.

▶ Caution

- The following disks are Advanced Format (512e) disks.
 - 1.8TB SAS disks
 - 2.4TB SAS disks
- When using Advanced Format (512e) disks, make sure that Advanced Format (512e) is supported by the server OS and the applications. If the server OS and applications do not support Advanced Format (512e), random write performance may be reduced.

● Nearline SAS Disks

Nearline SAS disks are high capacity cost effective disks for data backup and archive use. Nearline SAS disks can store information that requires a low access rate at a reasonable speed more cost effectively than SAS disks.

 **Caution**

- To use 10TB or larger Nearline SAS disks to configure RAID6-FR, the controller firmware version of the ETERNUS DX must be V10L86-4000 or later (*1).
 - Nearline SAS disks are used to store data that does not need the access performance of SAS disks. They are far more cost effective than SAS disks. It is recommended that SAS disks be used for data that is constantly accessed or when high performance/reliability is required.
 - If the ambient temperature exceeds the operating environment conditions, Nearline SAS disk performance may be reduced.
 - Nearline SAS disks can be used as Advanced Copy destinations and for the storage of archived data.
 - When Nearline SAS disks are used as an Advanced Copy destination, delayed access responses and slower copy speeds may be noticed, depending on the amount of I/O and the number of copy sessions.
 - The following disks are Advanced Format (512e) disks.
 - 2.5" Nearline SAS disks (2TB)
 - 3.5" Nearline SAS disks (6TB)
 - 3.5" Nearline SAS disks (8TB)
 - 3.5" Nearline SAS disks (10TB)
 - 3.5" Nearline SAS disks (12TB)
 - 3.5" Nearline SAS disks (14TB)
 - When using Advanced Format (512e) disks, make sure that Advanced Format (512e) is supported by the server OS and the applications. If the server OS and applications do not support Advanced Format (512e), random write performance may be reduced.
 - For details on the RAID levels that can be configured with Nearline SAS disks that have 6TB or more, refer to "[Supported RAID](#)" (page 13).
- *1 : When using the controller firmware for the China region, the controller firmware version must be V10L86-9006 or later.

● **SSDs**

SSDs are reliable drives with high performance. SSDs are used to store high performance databases and other frequently accessed data.

Flash memory as a storage media provides better random access performance than disks such as SAS disks and Nearline SAS disks. Containing no motors or other moving parts, they are highly resistant to impact.

The ETERNUS DX supports SSDs that have a high level wear leveling function and sufficient reserved space. Note that if the expected total write capacity is exceeded, the frequency of errors gradually increases, which may lead to a reduction in the write performance. The ETERNUS DX has a function that shows the remaining capacity as a percentage (health) in proportion to the expected total write capacity.

The number of rewrites may exceed the limit within the product warranty period if SSDs are used in a RAID1 configuration that has a high I/O access load. Using a RAID1+0, RAID5, RAID5+0, or RAID6 configuration is recommended.

The ETERNUS DX60 S4 supports Value SSDs. They are available at a lower cost per capacity than conventional SSDs by optimizing the guaranteed write endurance and the reserved space.

Note that the product's actual capacity of a Value SSD differs from the capacity that is displayed in ETERNUS Web GUI and ETERNUS CLI. For example, 2.00TB is displayed for a 1.92TB Value SSD.

Caution

- To use SSDs, the controller firmware version of the ETERNUS DX must be V10L90-4000 or later (*1).
 - If SSDs installed in the ETERNUS DX 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.
 - Turn on the ETERNUS DX and provide the SSDs with power for at least 10 days before 10 weeks pass since the SSDs had no power.
 - Back up the data in the SSDs regularly.
Backing up the current data to another area is recommended just before turning off the ETERNUS DX.
 - When storing SSDs without power, keep the ambient temperature of the ETERNUS DX below 40°C.
 - 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. For details, refer to "Drive Sanitization" (page 140).
- *1 : When using the controller firmware for the China region, the controller firmware version must be V10L90-9004 or later.

The drive characteristics of SAS disks, Nearline SAS disks, and SSDs are shown below.

Table 33 Drive Characteristics

| Type | Reliability | Performance | Price (*1) |
|--------------------|-------------|-------------|------------|
| SAS disks | ○ | ○ | ○ |
| Nearline SAS Disks | △ | △ | ◎ |
| SSDs | ◎ | ◎ | △ |

◎: Very good ○: Good △: Reasonable

*1: Bits are used to compare the prices for each drive type and the price goes up in the following order; ◎, ○, and △.

■ Number of Installable Drives

Up to 96 drives can be installed.

The following table shows the number of installable drives when the maximum number of drive enclosures are installed.

Table 34 Number of Installable Drives

| Type | ETERNUS DX60 S4/DX60 S3 |
|----------------------------|-------------------------|
| 2.5" type drive enclosures | 96 (*1) |
| 3.5" type drive enclosures | 48 (*2) |

*1: Number of drives when 2.5" controller enclosures are used.

*2: Number of drives when 3.5" controller enclosures are used.

For details on the number of required hot spares, refer to "■ Number of Installable Hot Spares" (page 25) in "Hot Spares" (page 24).

Standard Installation Rules

This section describes the standard installation rules for the following optional products.

- Controller modules
- Drive enclosures
- I/O modules
- Drives

Controller Module


This section describes the standard installation rules for controller modules.

By installing controller modules, the controller enclosure can connect with the drive enclosure even if drives are not installed.

■ Installation Order

Install CM#0 first, followed by CM#1.

Figure 80 Controller Installation Order



Rear view of a controller enclosure

Drive Enclosure

This section describes the installation order for drive enclosures.

2.5" type and 3.5" type drive enclosures can be installed together in the ETERNUS DX.

Drive enclosures can connect with other enclosures even if drives are not installed.

The installation priority order for drive enclosures varies depending on the controller enclosure type.

■ 2.5" Type Controller Enclosure

- 1 2.5" type drive enclosures
- 2 3.5" type drive enclosures

■ 3.5" Type Controller Enclosure

- 1 3.5" type drive enclosures
- 2 2.5" type drive enclosures

Drive enclosures are installed above the controller enclosure according to the priority order.

I/O Module


This section describes the standard installation rules for I/O modules.

By installing I/O modules, the drive enclosure can connect with the controller enclosure even if drives are not installed.

■ Installation Order

Install IOM#0 first, and then install IOM#1.

Figure 81 I/O Module Installation Order



Rear view of a drive enclosure

Drive

This section explains the installation rules for drives.

The supported drives vary between the ETERNUS DX60 S4 and the ETERNUS DX60 S3. For details about drives, refer to "Overview" of the currently used storage systems.

■ 2.5" Drives


The installation priority order for 2.5" drives is shown below.

- 1 2.5" SSDs (200GB)
- 2 2.5" SSDs (400GB) (MLC SSDs)
- 3 2.5" SSDs (400GB) (Value SSDs)
- 4 2.5" SSDs (800GB)
- 5 2.5" SSDs (960GB)
- 6 2.5" SSDs (1.6TB)
- 7 2.5" SSDs (1.92TB)
- 8 2.5" SSDs (3.84TB)
- 9 2.5" SAS disks (300GB/15krpm)
- 10 2.5" SAS disks (600GB/15krpm)
- 11 2.5" SAS disks (900GB/15krpm)
- 12 2.5" SAS disks (300GB/10krpm)
- 13 2.5" SAS disks (600GB/10krpm)
- 14 2.5" SAS disks (900GB/10krpm)
- 15 2.5" SAS disks (1.2TB/10krpm)
- 16 2.5" SAS disks (1.8TB/10krpm)

- 17 2.5" SAS disks (2.4TB/10krpm)
- 18 2.5" Nearline SAS disks (1TB/7.2krpm)
- 19 2.5" Nearline SAS disks (2TB/7.2krpm)

According to the installation order, install drives in the slots of a controller enclosure from Slot#0 to Slot#23 in ascending order.

Figure 82 Installation Diagram for 2.5" Drives




■ 3.5" Drives

The installation priority order for 3.5" drives is shown below.

- 1 3.5" SSDs (200GB)
- 2 3.5" SSDs (400GB) (MLC SSDs)
- 3 3.5" SSDs (400GB) (Value SSDs)
- 4 3.5" SSDs (800GB)
- 5 3.5" SSDs (1.6TB)
- 6 3.5" SAS disks (900GB/10krpm)
- 7 3.5" SAS disks (1.2TB/10krpm)
- 8 3.5" SAS disks (1.8TB/10krpm)
- 9 3.5" SAS disks (2.4TB/10krpm)
- 10 3.5" Nearline SAS disks (2TB/7.2krpm)
- 11 3.5" Nearline SAS disks (4TB/7.2krpm)
- 12 3.5" Nearline SAS disks (6TB/7.2krpm)
- 13 3.5" Nearline SAS disks (8TB/7.2krpm)
- 14 3.5" Nearline SAS disks (10TB/7.2krpm)
- 15 3.5" Nearline SAS disks (12TB/7.2krpm)
- 16 3.5" Nearline SAS disks (14TB/7.2krpm)

According to the installation order, install drives in the slots of a controller enclosure from Slot#0 to Slot#11 in ascending order. Then, install drives in the slots of a drive enclosure from Slot#0 to Slot#11 in ascending order.

Figure 83 Installation Diagram for 3.5" Drives



Recommended RAID Group Configurations

No restrictions on the installation location of drives apply if the same types of drives are used to create RAID groups.

To improve reliability, the installation location of drives that configure a RAID group must be considered.

Data reliability when enclosures fail can be improved by selecting configuration drives that are distributed to a larger number of enclosures when a RAID group is created.

For details on the recommended number of drives that configure a RAID group, refer to "RAID Group" (page 21).

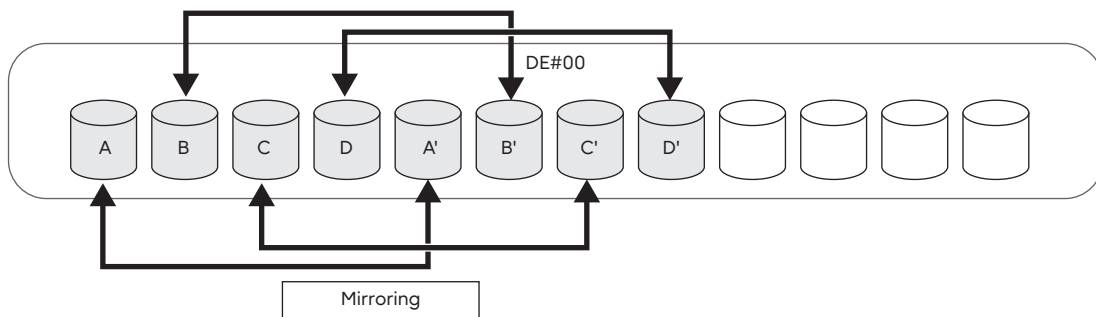
■ Mirroring Configuration

RAID1+0(4D+4M) is used in the following examples to explain how drives are installed to configure a mirroring RAID level.

The drive number is determined by the DE-ID of the enclosure and the slot number in which the drive is installed. Starting from the smallest drive number in the configuration, half of the drives are allocated into one group and the remaining drives are allocated into the other group. Each drive in the different groups are paired for mirroring.

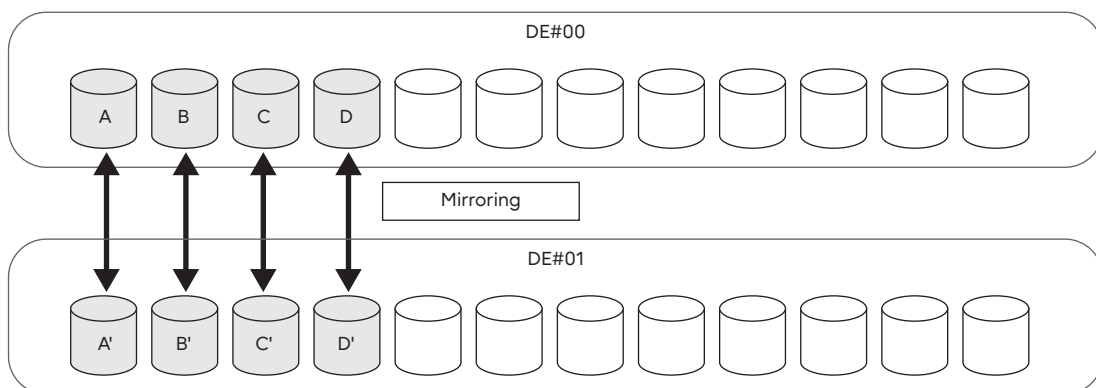
Example 1: All drives are installed in a single enclosure

Figure 84 Drive Combination 1




Example 2: Paired drives are installed in two different enclosures

Figure 85 Drive Combination 2



Example 3: Paired drives are installed in three different enclosures

Figure 86 Drive Combination 3



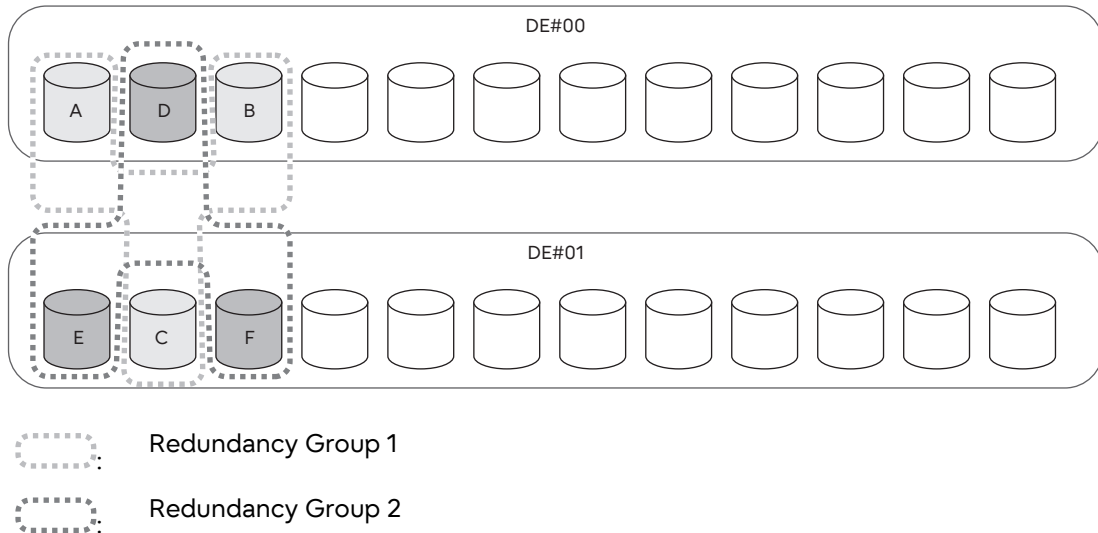
■ Double Striping Configuration with Distributed Parity

RAID5+0(2D+1P) ×2 is used in the following examples to explain how to install drives that are configured for double striping with distributed parity.

The drive number is determined by the DE-ID of the enclosure and the slot number in which the drive is installed. Drives are divided into two redundant set groups in ascending order of drive numbers.

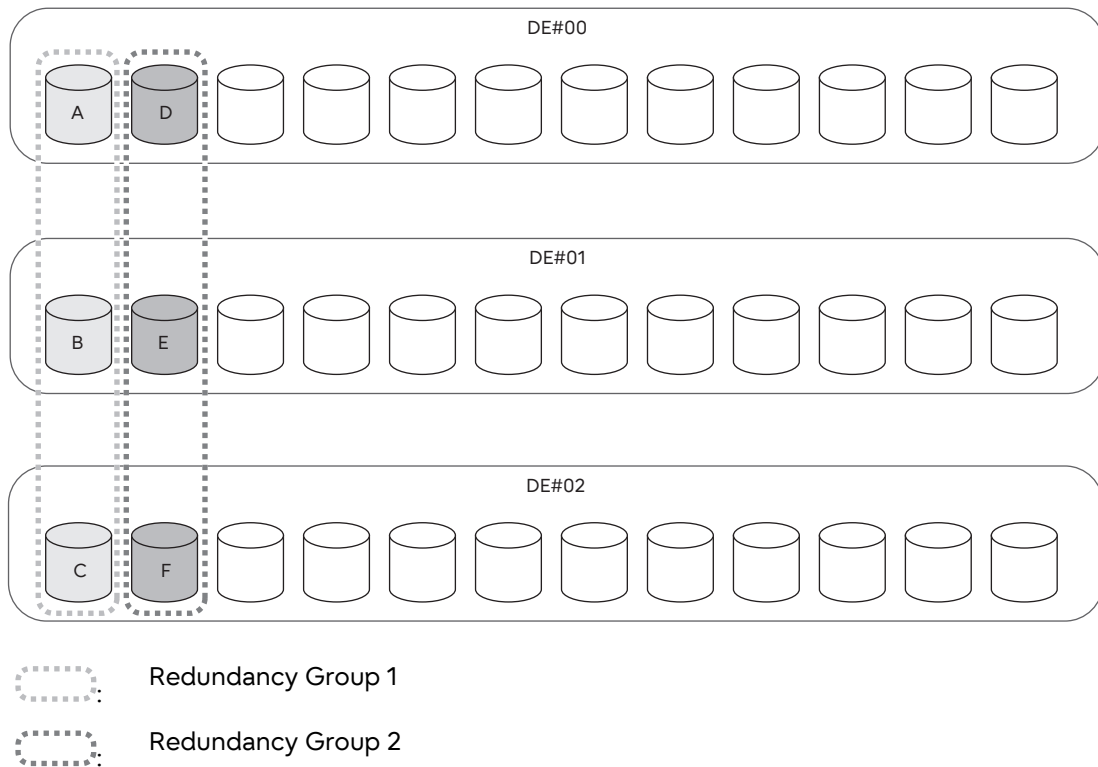
Example 4: Drives are installed in two different enclosures

Figure 87 Drive Combination 4



Example 5: Drives are installed in three different enclosures

Figure 88 Drive Combination 5



4. Maintenance/Expansion

Hot Swap/Hot Expansion

"Hot swap" allows components to be replaced, or allows the firmware to be updated while the system is running.

"Hot expansion" allows components to be added while the system is running.

■ Hot Swap/Hot Expansion (ETERNUS DX60 S4)

The table below shows whether hot swap or hot expansion for components of the ETERNUS DX60 S4 is possible.

Table 35 Hot Swap and Hot Expansion Availability for Components (ETERNUS DX60 S4)

| Component | Hot swap | Hot expansion | Remarks |
|---------------------------|-------------|---------------|---|
| Controller enclosure (CE) | × | — | Replace the controller enclosure (CE) when the MP (*1) fails. |
| Controller module (CM) | ○ (*2) | ○ (*3) | — |
| System memory | ○ (*2) | — | — |
| BBU | ○ (*2) | — | — |
| BUD | ○ (*2) | — | — |
| Controller firmware | ○ (*2) (*4) | — | — |
| SFP+ | ○ (*2) | ○ | — |
| iSCSI port | — | ○ | — |
| SAS port | — | ○ | — |
| Power supply unit (PSU) | ○ | — | — |
| Disk (HDD) | ○ | ○ | — |
| SSD | ○ | ○ | — |
| Operation panel (PANEL) | × (*5) | — | — |
| Disk firmware | ○ (*4) | — | — |
| Drive enclosure (DE) | ○ | ○ | Replace the drive enclosure (DE) when the MP (*1) fails. |
| Power supply unit (PSU) | ○ | — | — |
| Disk (HDD) | ○ | ○ | — |
| SSD | ○ | ○ | — |
| Operation panel (PANEL) | × (*5) | — | — |
| I/O module (IOM) | ○ | ○ (*3) | — |
| Disk firmware | ○ (*4) | — | — |

○: Allowed / ×: Not allowed (cold swap is possible) / —: Not applicable

***1:** Mid Plane. This is a board that is located between the front (drive side) and rear (controller (CM) or I/O module (IOM) side) of the ETERNUS DX.

***2:** When a multipath configuration is used, switch to the host paths of the CM that will not have maintenance performed to continue operation.

- *3: Hot expansion can only be performed when it is expanded together with a CM in a single-controller configuration.
- *4: Depending on the changes in the firmware, this function may require I/Os to be temporarily stopped.
- *5: Operation can be continued during a failure. The status of the ETERNUS DX can be monitored via ETERNUS Web GUI or ETERNUS CLI.

■ Hot Swap/Hot Expansion (ETERNUS DX60 S3)

The table below shows whether hot swap or hot expansion for components of the ETERNUS DX60 S3 is possible.

Table 36 Hot Swap and Hot Expansion Availability for Components (ETERNUS DX60 S3)

| Component | Hot swap | Hot expansion | Remarks |
|---------------------------|-------------|---------------|---|
| Controller enclosure (CE) | × | — | Replace the controller enclosure (CE) when the MP (*1) fails. |
| Controller module (CM) | ○ (*2) | ○ (*3) | — |
| Cache memory | ○ (*2) | — | — |
| Controller firmware | ○ (*2) (*4) | — | — |
| SFP+ | ○ (*2) | ○ | — |
| iSCSI port | — | ○ | — |
| SAS port | — | ○ | — |
| Power supply unit (PSU) | ○ | — | — |
| Disk (HDD) | ○ | ○ | — |
| SSD | ○ | ○ | — |
| Operation panel (PANEL) | × (*5) | — | — |
| Disk firmware | ○ (*4) | — | — |
| Drive enclosure (DE) | ○ | ○ | Replace the drive enclosure (DE) when the MP (*1) fails. |
| Power supply unit (PSU) | ○ | — | — |
| Disk (HDD) | ○ | ○ | — |
| SSD | ○ | ○ | — |
| Operation panel (PANEL) | × (*5) | — | — |
| I/O module (IOM) | ○ | ○ (*3) | — |
| Disk firmware | ○ (*4) | — | — |

○: Allowed / ×: Not allowed (cold swap is possible) / —: Not applicable

- *1: Mid Plane. This is a board that is located between the front (drive side) and rear (controller (CM) or I/O module (IOM) side) of the ETERNUS DX.
- *2: When a multipath configuration is used, switch to the host paths of the CM that will not have maintenance performed to continue operation.
- *3: Hot expansion can only be performed when it is expanded together with a CM in a single-controller configuration.
- *4: Depending on the changes in the firmware, this function may require I/Os to be temporarily stopped.

- *5: Operation can be continued during a failure. The status of the ETERNUS DX can be monitored via ETERNUS Web GUI or ETERNUS CLI.

User Expansion

Customers can expand (add) the following components.

To add components that are not described in this section, contact your sales representative or maintenance engineer.

- Drives
- Drive enclosures

Drive Sanitization

Drive sanitization is a function that deletes data in a drive using the sanitization function of the drives (SSDs, Online SAS disks, or Nearline SAS disks). The drive sanitization function can be used to delete user data when, for example, discarding drives.

Caution

- To use this function, the controller firmware version must be upgraded to V10L88 or later.
 - The Maintenance Operation policy is required to perform this function. Also, change the status of the ETERNUS DX to the maintenance status.
 - To use this function, run "Sanitize Drive" from ETERNUS Web GUI.
 - If the sanitization function is not supported by the drive, this function will terminate with an error.
-

A. Function Specification List

This appendix shows the combinations of functions that can be executed at the same time and the targets for each function.

List of Supported Protocols

Table 37 List of Supported Protocols

| Item | | LAN for Operation Management | iSCSI (SAN) |
|--|---------------------|------------------------------------|--------------------------|
| Operation mode | | 1000BASE-T/ 100BASE-TX/10BASE-T | 1000BASE-T/ 10GBASE-T |
| ETERNUS Web GUI | http | ○ | × |
| | https (SSL v3, TLS) | ○ | × |
| ETERNUS CLI | SSH v2 | ○ | × |
| | telnet | ○ | × |
| | ftp (client) | ○ | × |
| SMI-S | http / https | ○ | × |
| | SLP | ○ | × |
| NTP (time) | NTP v4 | ○ | × |
| E-mail | SMTP (Client) | ○ | × |
| SNMP | SNMP v1, v2c, v3 | ○ | × |
| Event notification and audit log sending | Syslog | ○ | × |
| Ping | ICMP | ○ | ○ |
| Network address | IPv4, IPv6 | ○ | ○ |
| Authentication | RADIUS | ○ | × |
| | CHAP | × | ○ |

○: Supported ×: Not supported

Target Pool for Each Function/Volume List

This section describes the functions which can be performed on RAID groups, pools, and volumes.

Target RAID Groups/Pools of Each Function

| Action | RAID group | SDP | TPP |
|--------------------------------|--------------------|-----------------------|----------------------------|
| Components | Standard, SDV, WSV | SDPV | TPV |
| Max. number | 48 | 1 | 48 |
| Create | ○ | — (*1) | ○ |
| Delete | ○ | — (*2) | ○ (*3) |
| Rename | ○ | × | ○ |
| Expand capacity | ○ (by LDE) | ○ (by adding an SDPV) | ○ (by adding a RAID group) |
| Migration | ○ | × | × |
| Logical Device Expansion (LDE) | ○ | ○ | × |
| Format (All area) | × | × | ○ |
| Format (Unformatted area) | × | × | ○ |
| Modify threshold | × | ○ | ○ |
| Eco-mode setup | ○ | × | ○ |
| Switch controlling CM | ○ | ○ | ○ |

○: Possible ×: Impossible - : N/A

*1: If an SDPV is created, an SDP is automatically created.

*2: If all of the SDPVs are deleted, SDPs are automatically deleted.

*3: When a volume exists in a pool, the pool cannot be deleted.

Target Volumes of Each Function

| Action | Standard | | SDV | SDPV | TPV | WSV | ODX Buffer volume (*1) |
|--------------------------------|----------|----------------|-----|------|--------|--------|------------------------|
| | Single | Con-caten-ated | | | | | |
| Create | ○ | ○ (*2) | ○ | ○ | ○ | ○ | ○ |
| Delete | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| Rename | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| Format | ○ | ○ | ○ | ○ | ○ | ○ | ○ |
| Eco-mode (*3) | ○ | ○ | ○ | × | ○ | × | × |
| TPV capacity expansion | × | × | × | × | ○ | × | ○ |
| RAID Migration | ○ | △ (*4) | × | × | ○ (*5) | ○ (*6) | ○ |
| Logical Device Expansion (LDE) | ○ | ○ | ○ | ○ | × | × | ○ |

| Action | Standard | | SDV | SDPV | TPV | WSV | ODX Buffer volume (*1) |
|--|---------------|----------------|---------|--------|---------|-----|------------------------|
| | Single | Con-caten-ated | | | | | |
| LUN Concatenation | ○ (*7)(*8) | ○ (*7)(*8) | × | × | × | × | ○ |
| Balancing | × | × | × | × | ○ | × | ○ |
| TPV Capacity Optimization | × | × | × | × | ○ | × | × |
| Modify threshold | × | × | × | △ (*9) | ○ | × | ○ |
| Advanced Copy function (Local copy) | ○ | ○ | ○ (*10) | × | ○ | ○ | ○ |
| Forbid Advanced Copy | ○ | ○ | × | × | ○ | ○ | × |
| Reserved and forced deletion | × | × | × | ○ | × | × | × |
| Release reservation | ○ | ○ | × | × | ○ | ○ | × |
| Performance monitoring | ○ | ○ | ○ | × | ○ | ○ | ○ |
| Modify cache parameters | ○ (*11) | ○ (*11) | ○ | ○ | ○ (*11) | ○ | ○ |
| Create a LUN while rebuild- ing | ○ | ○ | ○ | ○ | ○ | ○ | × |
| Storage Migration | ○ | ○ | × | × | ○ | ○ | × |

○: Possible, △: Partially possible, ×: Not possible

- *1 : An ODX Buffer volume is an exclusive volume that is required to perform ODX. Standard volumes or TPVs can be used for the ODX Buffer volume.
- *2 : A concatenated volumes cannot be created at the same time as a volume. Existing volumes can be concatenated by using LUN Concatenation.
- *3 : Eco-mode is configured for Standard volumes/SDVs in RAID group units that include the target volume, and TPVs in pool (TPP) units.
- *4 : If multiple volumes have been concatenated using the LUN Concatenation function, RAID Migration can be executed only on concatenation source volumes.
- *5 : Capacity expansion and RAID Migration cannot be performed at the same time.
- *6 : The volume capacity can be expanded by specifying a relatively large capacity for the destination volume when RAID Migration is performed.
- *7 : The maximum number of concatenated volumes is 16.
- *8 : If T10-DIF is enabled, this function cannot be executed.
- *9 : A threshold can be set to the pool which includes the target volume.
- *10: SDVs are used as the copy destination for SnapOPC/SnapOPC+. SDVs are also used as the copy source if the copy destination of SnapOPC/SnapOPC+ is also set as the copy source (Cascade Copy/restore).
- *11: This is part of the cache parameter settings, and cannot be configured in the units of volumes.

Combinations of Functions That Are Available for Simultaneous Executions

This section describes the availability of simultaneous execution with other functions, the number of processes that can be executed concurrently, and the capacity that can be processed concurrently for each function.

Combinations of Functions That Are Available for Simultaneous Executions

There are functions which cannot be performed concurrently when another function is being executed in the ETERNUS DX.

The following table shows which function combinations can be executed at the same time.

Table 38 Combinations of Functions That Can Be Executed Simultaneously (1/2)

| Process to be run | Already running process | | | | | | |
|---------------------------|-----------------------------------|---------------|---------------|----------------|--------------------------------|-------------------|--------------------|
| | Rebuild/ Copyback/ Redundant Copy | Fast Recovery | Format Volume | RAID Migration | Logical Device Expansion (LDE) | LUN Concatenation | Advanced Copy (*1) |
| Format Volume | ○ | ○ | ○ (*3) | ○ (*4) | ○ (*4) | ○ (*3) | ○ (*4) |
| RAID Migration | ○ | ○ | ○ (*4) | ○ (*4) | ○ (*4) | ○ (*5) | ○ (*6) |
| Logical Device Expansion | ○ (*7) | ○ | ○ (*7) | ○ (*7) | × | ○ (*7) | ○ |
| LUN Concatenation | ○ | ○ | ○ | ○ (*4) | ○ (*4) | ○ | ○ (*8) |
| Advanced Copy (*2) | ○ | ○ | ○ | ○ (*9) | ○ | ○ | ○ (*10) |
| Switch controlling CM | ○ (*7) | ○ (*7) | ○ (*7) | × | × | ○ (*7) | ○ |
| TPV capacity expansion | ○ | ○ | ○ | ○ | ○ | ○ | ○ (*8) |
| TPV Balancing | ○ | ○ | ○ (*4) | ○ (*4) | ○ | ○ | ○ (*6) |
| TPV Capacity Optimization | ○ | ○ | ○ | ○ (*4) | ○ | ○ | ○ |

○: Possible, ×: Not possible

Table 39 Combinations of Functions That Can Be Executed Simultaneously (2/2)

| Process to be run | Already running process | | | |
|--------------------------|-------------------------|-----------------------------|---------------|---------------------------|
| | Disk Patrol | Eco-mode (motor is stopped) | TPV Balancing | TPV Capacity Optimization |
| Format Volume | ○ | ○ (*11) | ○ (*4) | ○ |
| RAID Migration | ○ | ○ (*11) | ○ (*4) | ○ (*4) |
| Logical Device Expansion | ○ | ○ (*11) | ○ | ○ |

| Process to be run | Already running process | | | |
|---------------------------|-------------------------|-----------------------------|---------------|---------------------------|
| | Disk Patrol | Eco-mode (motor is stopped) | TPV Balancing | TPV Capacity Optimization |
| LUN Concatenation | ○ | ○ (*11) | ○ | ○ |
| Advanced Copy (*2) | ○ | ○ (*12) | ○ | ○ |
| Switch controlling CM | ○ | ○ | ○ | ○ |
| TPV capacity expansion | ○ | ○ (*11) | ○ (*4) | ○ |
| TPV Balancing | ○ | ○ (*11) | ○ (*4) | ○ (*4) |
| TPV Capacity Optimization | ○ | ○ (*11) | ○ (*4) | × |

○: Possible, ×: Not possible

- *1 : This indicates that the copy session is being set or the copy session is already set.
- *2 : This indicates the copy session setting operation.
- *3 : Using the "Format Volume" function will cause losing the entire data in volumes.
- *4 : The function cannot be executed when the same volume is set as the execution target.
- *5 : LUN concatenated volumes can be set as the RAID Migration source.
 Note that LUN concatenated volumes cannot be set as the RAID Migration destination.
- *6 : RAID Migration that expands the capacity cannot be performed on any volumes with copy sessions in LUN units.
- *7 : The function cannot be executed when the same RAID group is set as the execution target.
- *8 : When a copy session with a copy area that is specified in LUN units exists, volume capacity expansion cannot be performed because copying cannot be processed for the expanded part. When a copy session is specified for a logical disk, volume capacity expansion can be performed.
- *9 : Copy sessions in LUN units cannot be performed on a volume that is being expanded with RAID Migration.
- *10: Complies with Advanced Copy specifications (multi-copy/Cascade Copy).
- *11: The motor-off state of Eco-mode is stopped, and the drive motor is activated (spin-up).
- *12: EC can be performed (the motor-off state of Eco-mode will be stopped).
 OPC/QuickOPC/SnapOPC/SnapOPC+ cannot be performed (the ETERNUS DX returns an error).

Number of Processes That Can Be Executed Simultaneously

The following upper limits are applied to the number of processes that are to be executed.

- Only one process can be executed with Logical Device Expansion (LDE). Two or more processes cannot be executed simultaneously in the same device.
- A total number of 32 processes can be executed simultaneously for TPV balancing, RAID Migration, and Non-disruptive Storage Migration.

Capacity That Can Be Processed Simultaneously

The following upper limit is applied to the capacity of processes that are to be executed.

- A total capacity of 128TB can be processed simultaneously for TPV balancing, RAID Migration, and Non-disruptive Storage Migration.

Fujitsu Storage ETERNUS DX60 S4, ETERNUS DX60 S3 Hybrid Storage Systems
Design Guide (Basic)

P3AM-8992-21ENZO

Date of issuance: July 2023
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