FUJITSU

Technical White Paper
ETERNUS DX Feature Set

This white paper provides an overview of the main features supported by the Fujitsu ETERNUS DX storage systems. It highlights their benefits and use cases and briefly describes each feature.

Introduction

As data is one of the most important assets in any company, storage systems play a vital role in the IT infrastructures of all enterprises. IT administrators need to ensure that data is stored on reliable, highly available platforms that can scale efficiently in order to handle ongoing business changes.

New IT trends are also imposing new challenges on storage systems. The increasing use of business analytics and data warehousing, the hype around big data, server and desktop virtualization, the enormous growth of unstructured data are just a few of the examples.

These trends not only require massive storage capacity, but also – and more importantly – storage performance. Parameters, such as IOPS (Input & Output Operations per Second), latency and bandwidth are gaining in significance. Many enterprises are purchasing additional storage systems, just to provide the required performance. The result is storage capacity overprovisioning as well as higher operational and capital expenditure.

These are the reasons why Fujitsu has designed the ETERNUS DX family in such a manner – **Business-centric Storage.** Built on powerful performance architecture with unified block and file data access, it ensures that there are no more performance bottlenecks, and efficiently consolidates storage resources while keeping investments down to a minimum.

High processor performance, large system caches, fast disk interfaces and networking connectivity are just some of the highlights within the ETERNUS DX S3 family hardware.

The ETERNUS SF Management Suite complements the family with its simple setup and administration functionalities. This unified management system ensures that an IT administrator, once trained to operate a specific system in the ETERNUS DX family, will also be able to manage the other models within the family.
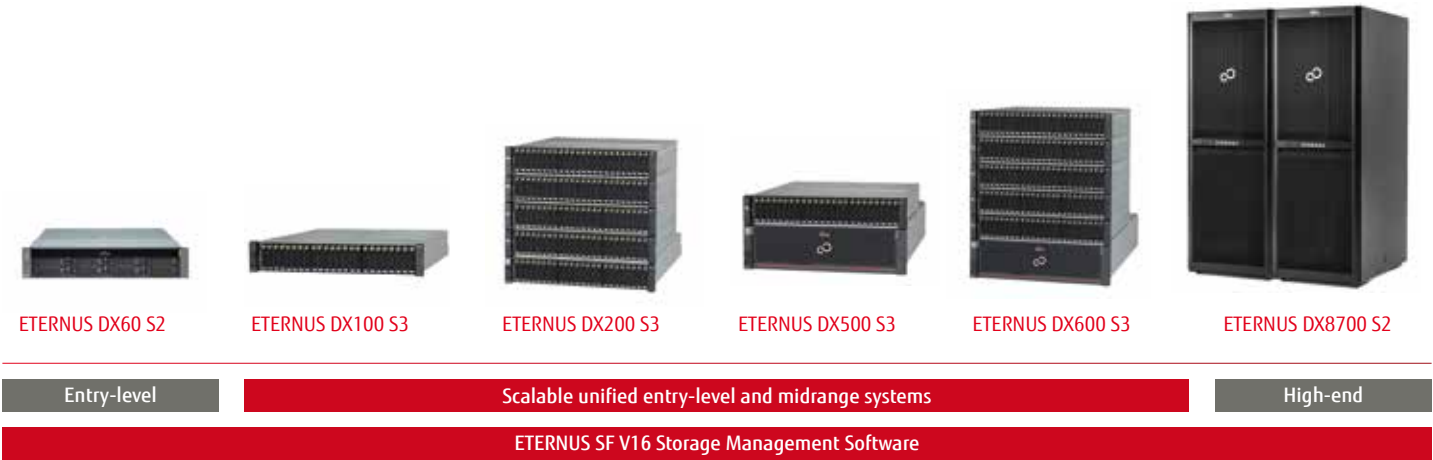
Advanced software features, particularly for automated service management quality, help to fulfill service levels and deliver predictable operations as well as stable response times for business-critical applications– at the right cost and with the right quality levels.

Hardware Architecture

Based on a unique and consistent system design, ETERNUS DX is a seamless family of disk storage systems ranging from entry to enterprise class. The main design principles of ETERNUS DX are scalability and flexibility. In order to provide an economic storage solution that satisfies the needs of smaller businesses which have tight budgets, the ETERNUS DX60 S2 supports all the basic functionalities. In the highly demanding enterprise segment, the ETERNUS DX8700 S2 has been enhanced to provide peak loads of up to 1 million IOPS. Its superior performance has been proven in various customer deployments and SPC benchmarks.

The scalable entry and mid-range segments have ETERNUS DX S3 models which are based on enhanced performance architecture. They provide unified block and files access. All models within the same generation use the same types of disks, disk shelves, racks, cables and other components. Only the controller module is unique to each model. The aim is to minimize the amount of work and costs involved in operations, training and migration. Furthermore, this family concept enables easy upgrading - from a smaller model to a bigger one – as you just have to change the controller module. All the other components can still be used.

ETERNUS DX60 S2	ETERNUS DX100 S3	ETERNUS DX200 S3	ETERNUS DX500 S3	ETERNUS DX600 S3	ETERNUS DX8700 S2

Entry-level	Scalable unified entry-level and midrange systems	High-end

**ETERNUS SF V16 Storage Management Software**

Scalable Entry Models

The scalable entry models – ETERNUS DX100 S3 and ETERNUS DX200 S3 – are designed in a compact style to ensure an optimized footprint. They comprise the following hardware:

Controller Enclosure (CE): The CE contains the Controller Modules (CM) and Power Supply Units (PSU) in the rear as well as the disk drives installed in the front. The controller enclosure can be equipped with either one controller module (single controller model) or with two modules (dual controller model). A single controller model can be easily upgraded to a dual controller model by just installing an extra controller.

Drive Enclosures (DE): They contain disk drives installed in the front as well as power supplies and expanders in the rear. The SAS expander is a unit that controls the connection to the controller modules. The disk enclosures can be equipped with one or two expanders depending on the number of controller modules. The various types of disk drives are described in detail in the following sections.

The basic architecture is shown below

CA: Channel Adapter
CE: Controller Enclosure
CM: Controller Module
EXP: SAS Expander
DE: Drive Enclosure
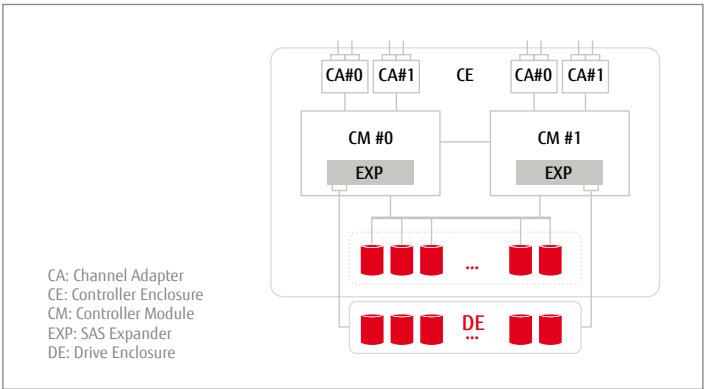
The Controller Module (CM) runs the firmware-based operating system and contains host interface ports, drive interface ports, cache memory, a System Capacitor Unit (SCU), RAID control processors, RAID control firmware, and Ethernet ports. Channel Adapters (CA), providing the connection to the host, are installed in each controller module and drive-side interfaces are installed onboard. The 2 CMs are interconnected via PCI Express. SAS interfaces between the CMs are cross-connected to provide a redundant path to the disks installed on CEs. A Drive Enclosure (DE) also has two independent interface ports that are directly connected to two CMs for path redundancy.

Mid-Range Models

The mid-range models of the ETERNUS DX S3 (ETERNUS DX500 S3 and ETERNUS DX600 S3) are designed to provide the highest possible capacity, performance and reliability levels. They contain the following hardware:

Controller Enclosure (CE): The CE contains the Controller Modules (CM), Channel Adapters (CA), PCI Express Flash Module (PFM), Battery Back-up Unit (BBU) and Power Supply Units (PSU). Unlike the entry-level models, the controller enclosure does not contain disk drives.

Drive Enclosures (DE): The minimum configuration consists of one DE in addition to the CE. Expansion drive enclosures can be added according to the capacity required. The same drive enclosures can be used in any ETERNUS DX S3 model, ensuring a smooth upgrade and full investment protection. Disk enclosures supporting 2.5″ and 3.5″ disk drives are available.

The architecture of the ETERNUS DX S3 mid-range models is shown below
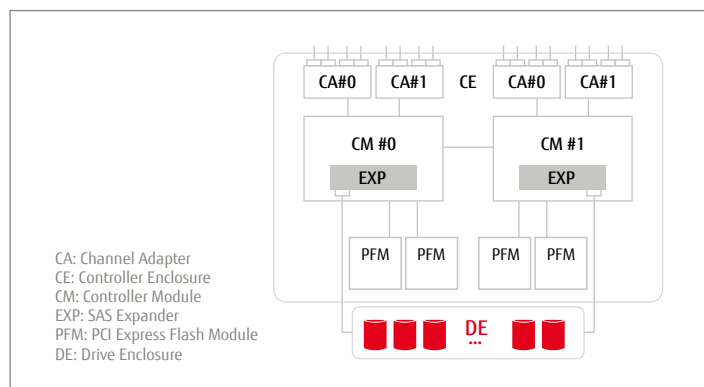
The basic architecture of the system consists of two interconnected CMs (via PCI Express), each having two Channel Adapters (CA), and two drive-side interfaces installed. It also supports high-speed PCI Express Flash Modules (PFM) as large capacity cache – the Extreme Cache (EXC) – which provide higher performance with less number of disk drives, optimizing costs and power consumption. The EXC supports hot-swapping. The system can host up to four modules providing up to 5.6 TB cache capacity for read operations. Unlike many implementations, the PFM modules are mounted in the controllers – not in the drive enclosures – providing some benefits, such as:
■ Short latency due to direct PCI Express connection to CPU
■ Better utilization of the drive enclosures, as no DE slots need to be dedicated for EXC
■ I/O performance of the drive enclosures are not impacted

High-End Model

The schematic architecture of the ETERNUS DX8700 S2 high-end model of the ETERNUS DX family is shown below as it differs slightly from other models. While most of the basic functional blocks are the same, the connection between the controller module and the disk drives located in the drive enclosures is realized using the back-end routers (BRT), which are not part of the controller, but rather separate modules installed according to the number of disk enclosures.
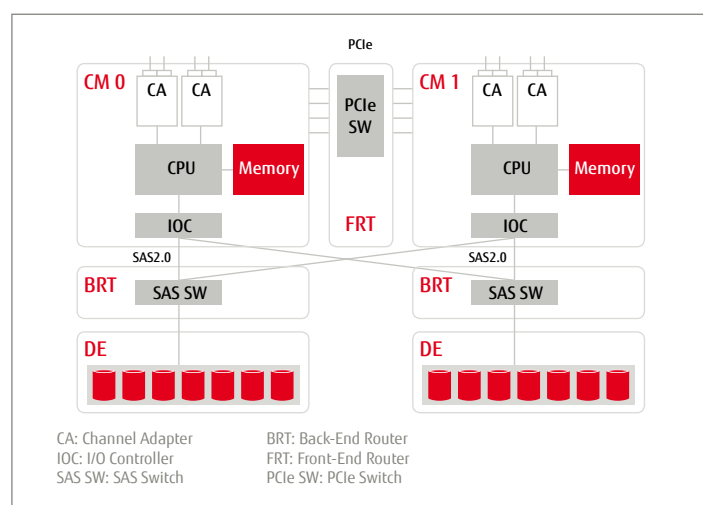

In order to ensure reliable operation and high availability, the ETERNUS disk storage systems have a fully redundant configuration and hot-maintenance capability. The interconnection of the CMs is duplicated. A drive enclosure (DE) has also two independent interface ports that are directly connected to two CMs for path redundancy.

The use of redundant components and multiple controller-to-drive connections ensures exceptional reliability. Hardware components supporting hot-exchange and hot-expansion features not only enable firmware upgrades during system operation, but also provide capacity expansion together with the LUN expansion feature where DEs or HDDs are added as needed.

Reverse cabling

In order to provide higher availability, the controller enclosure is connected to the drive enclosure via reverse cabling. The connection of the primary path is implemented in ascending order while the secondary path is connected in descending order, as shown in the following figure.

DE#09
DE#08
DE#07
DE#06
DE#05
DE#04
DE#03
DE#02
DE#01
CE
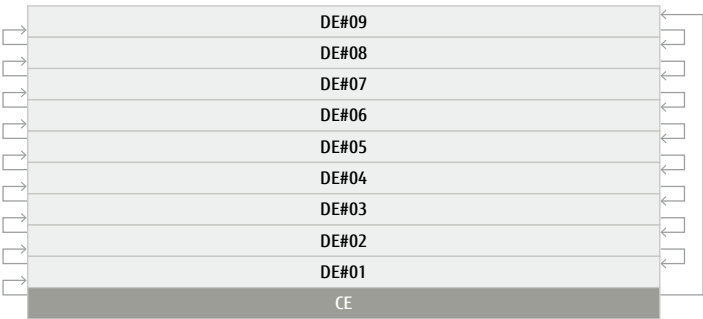
In the event of a DE failure, only the affected DE is disconnected. All other DEs remain accessible.

Symmetric Multi-Processing and 64-bit Operating System

ETERNUS DX S3 systems use powerful Intel® multi-core processors to provide leading performance. Symmetric multi-processing is a technology that allows controlling multiple processor cores by one operating system. Combined with using the latest 64-bit operating system, it offers many benefits, such as:

- Memory resources are better utilized
- Performance is dramatically increased as the core sufficiency ratio is improved by sorting required transactions flexibly even when cores are increased
- Better resource optimization as resources are not bound to cores
- Higher flexibility and processing performance as 64-bit operating system CPU can handle high-volume data

The figure below shows the benefit of symmetric multi-processing, as many different tasks can be executed in parallel, resulting in significant performance increase.
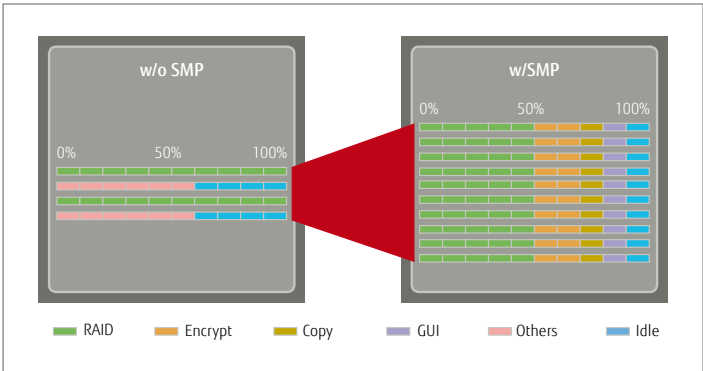
Interfaces

The ETERNUS DX S3 provides different types of host interfaces, offering customers full flexibility in selecting the most appropriate data center infrastructure. Fibre channel is the most widely used storage networking technology as it is highly reliable, efficient and secure. Up to 16Gbit/s bandwidth is supported, offering the highest possible bandwidth.

1Gbit/s & 10Gbit/s iSCSI are also supported for connecting to IP networks. They are simple to operate and hence preferred by many customers. 10Gbit/s FCoE (Fibre Channel over Ethernet) is a new networking technology that encapsulates FC frames over Ethernet networks, which allow using the FC protocol over existing Ethernet networks combining the benefits of both technologies. The backend interfaces are realized by 12Gbit/s SAS.

Types of disk drives

As the storage requirements differ according to the type of data and the frequency of its usage, various types of disk drives need to be supported in order to allocate the right disks for each type of data. Some data is mission-critical; it has to be accessed immediately in order to avoid revenue loss or productivity degradation. This data must be stored on drives with very high performance, such as SSD (Solid State Drives). On the other hand, some types of data do not require very high performance, but need to be stored for longer periods; it can thus be stored on more cost-efficient, high-capacity disks, such as Nearline SAS disk drives, enabling the customer to balance speed, capacity and costs.

In addition, 2.5″ and 3.5″ enclosures can be mixed in the same storage configuration. SSD, SAS and Nearline SAS disks can be mixed in the same 2.5″ drive enclosures. 3.5″ enclosures can host Nearline disk drives, which provide up to 4TB capacity.

In addition to manual tiering of data into the appropriate type of disk drives, the ETERNUS DX product family supports Automated Storage Tiering (AST) as described in the chapter "ETERNUS SF Management Suite" of this white paper.
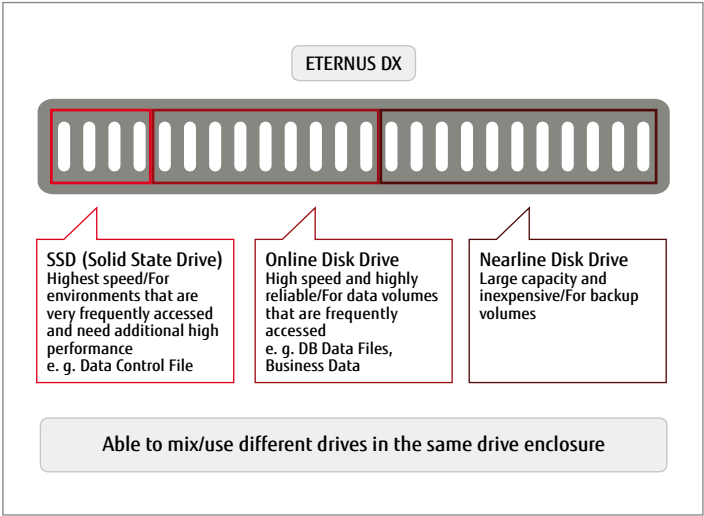
Low-power, highly reliable and high-performance Solid State Drives (SSD)

Solid State Drives (SSD) use semiconductor memory to store data. They contain no motors or moving parts, and thus have a much higher read/write access speed and reduced power consumption.

They benefit those applications with high random access requirements, such as databases. In addition, with no motors or moving parts, they are more reliable than disk drives. These features make SSD ideal for Tier 0 storage use. The SSD used in ETERNUS DX have enterprise-class performance and reliability. While maintaining compatibility with traditional disk drives, they support low-power consumption and high-speed operation. If a power failure occurs, SSD in ETERNUS DX can move any data in the volatile memory of the SSD (high-speed DRAM cache) to the flash memory and ensure a safe system halt.

Online SAS disk drives

For data volumes that are frequently accessed, but still do not require the very high performance of SSD, SAS disk drives are used, providing a balanced mix of performance and capacity, while keeping costs at a moderate level.

Cost-efficient, high-capacity Nearline SAS disk drives

Storing infrequently accessed data on high-performance storage devices generates unnecessary costs. To meet the growing demand for cost-effective storage of less frequently accessed data, Fujitsu provides high-capacity, highly reliable, yet cost-optimized Nearline disk drives in its ETERNUS DX disk storage systems.

This combination of online disk drives and Nearline disk drives in the same drive enclosure enables ETERNUS DX disk storage systems to support cost-effective operations, such as disk-to-disk backup and long-term storage of reference data.

High-density drive enclosures

High-density drive enclosures are available in order to meet the growing demand for more capacity. They house up to 60 3.5″ disk drives in one DE, which can provide up to 240 TB in 4 rack units.

## Energy savings – Eco-mode

Power consumption is one of the biggest contributors to the OPEX of any company. It is also a heavy burden on the environment. The disk drives used in all the ETERNUS DX products have Eco-mode support using MAID technology. MAID is an acronym for Massive Array of Idle Disks. This technology extends the life of the disk drives and reduces power consumption by stopping spindle rotation on less frequently accessed disk drives.

Eco-mode reduces power consumption because the disks are powered off during periods of inactivity. This time-controlled mode, based on the scheduled use of specific disks, can be set up for individual RAID groups and backup operations.

Rotations stop when scheduled disks are not accessed for specific periods and then restart when a data access order is received. Full rotation is restored within seconds.

Furthermore, two options are supported to control the power-off of idle disk drives:

- Scheduled control: disk drives that are inactive for 3 minutes will be powered off
- Software control: disk drives are synchronously controlled with the accessing server applications

Control in conjunction with operation
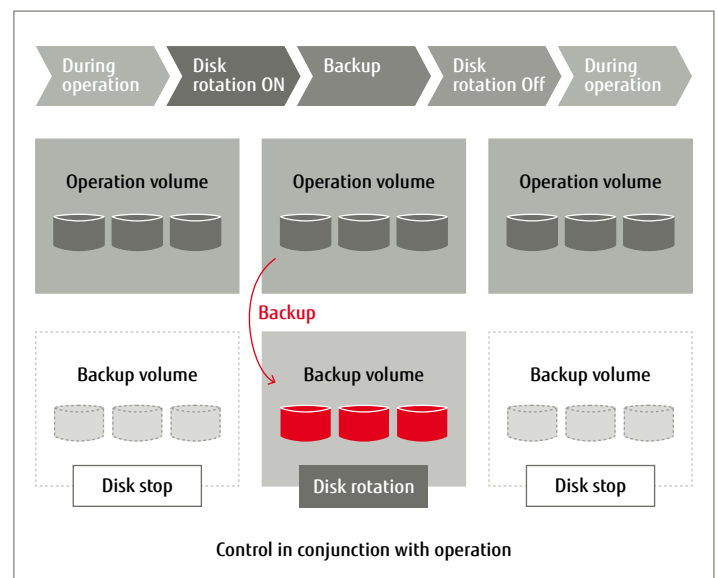
Data Integrity

Data errors can occur for different reasons. They result in data corruption, which in turn can lead to a loss of important company information. The ETERNUS DX storage systems support the following techniques which ensure data integrity:

Data Block Guard

The Data Block Guard function adds check codes for data stored during write operations. While verifying the codes for the read/write operations, it guarantees data integrity at multiple checkpoints along the data transmission route.

Oracle Database Data Guard

ETERNUS DX disk storage systems check data integrity using Data Block Guard technology. While this is very important, it still does not cover those situations where data corruption occurs in the interfaces between systems. This is because Data Block Guard only verifies data after it has reached the storage device.

Fujitsu also uses another data protection mechanism called Database Data Guard by Oracle. This combination of data security measures enables ETERNUS DX disk storage systems to provide very robust data integrity.

When data is written to the disk storage system, the database adds check codes. The disk storage system knows the logic of these check codes and where the codes are placed, thus enabling it to verify the data via the check codes. When an ETERNUS DX disk storage system identifies any data corruption, it stops further operations and notifies the administrator, thus preventing the use of data which is known to be corrupt.
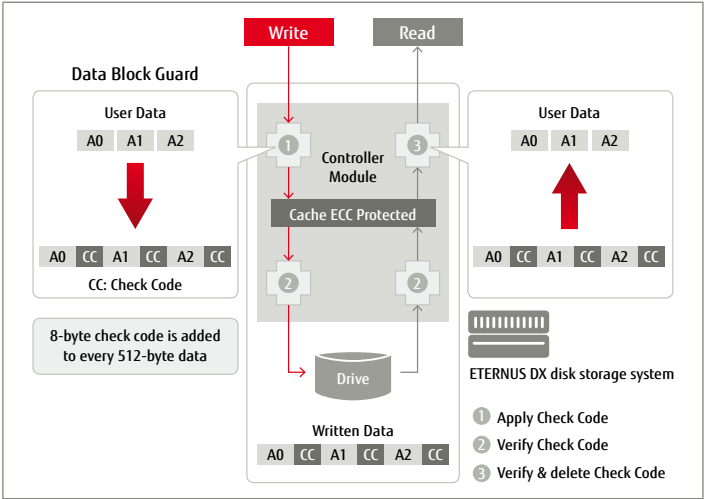
Disk Drive Patrol

Data on the ETERNUS DX disk storage systems is protected via a disk drive patrol function. The controller regularly checks the disk drives in order to detect errors and write failures. This process also ensures data consistency within the volume group.

Data on each disk drive is read, and if an error is detected, data is reconstructed via the redundant information contained within the volume group. The corrected data is then written to a new valid area on the disk drive.
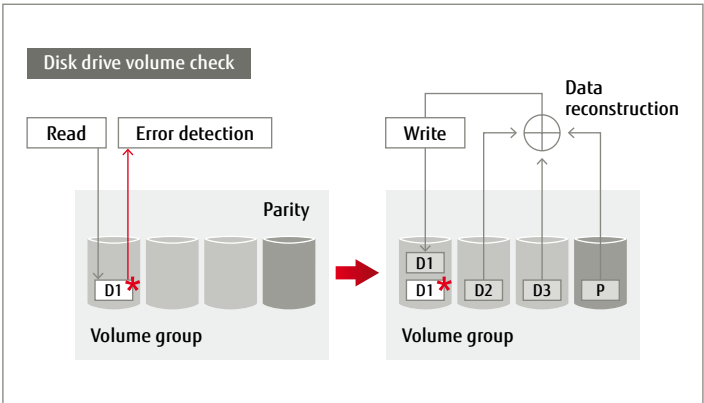
Benefit:

■ Higher data reliability as data errors are quickly found and corrected (by reconstruction) and disk write failures are avoided.
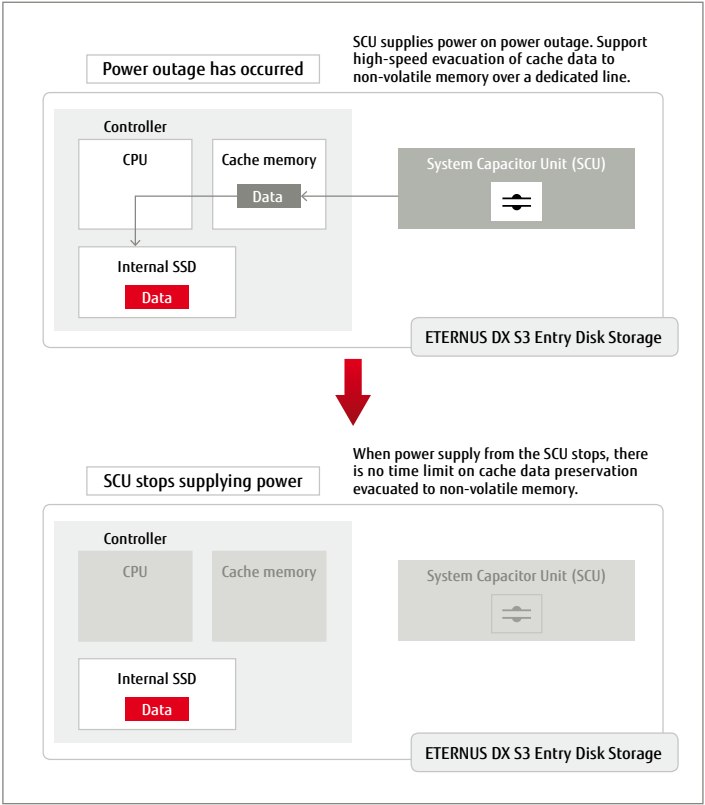

Data Block Guard
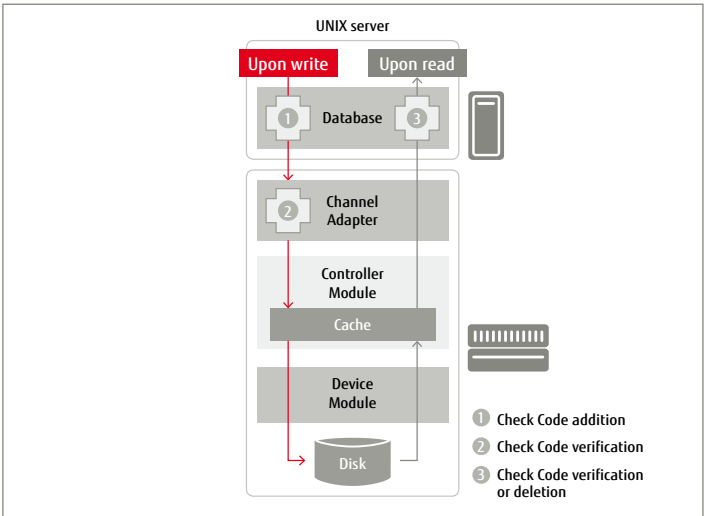
Database Data Guard

Disk drive patrol

Cache Protector

ETERNUS DX S3 scalable entry-level models guarantee data security even in the event of cache failure because cache is redundantly configured and constantly mirrored.

If the power supply fails, the controller cache is automatically evacuated and data is placed in an internal SSD. A system capacitor unit (SCU) provides sufficient power to always ensure that all data is successfully rescued. The internal SSD protects the data indefinitely.

The use of capacitors has some advantages over batteries; they shrink system size and weight because capacitors are smaller and lighter than batteries. Toxic waste is also reduced by using a permanent SCU instead of periodically replaceable batteries.

SCU supplies power on power outage. Support high-speed evacuation of cache data to non-volatile memory over a dedicated line.

Power outage has occurred

Controller
CPU
Cache memory
Data
System Capacitor Unit (SCU)
Internal SSD
Data
ETERNUS DX S3 Entry Disk Storage

When power supply from the SCU stops, there is no time limit on cache data preservation evacuated to non-volatile memory.

SCU stops supplying power

Controller
CPU
Cache memory
System Capacitor Unit (SCU)
Internal SSD
Data
ETERNUS DX S3 Entry Disk Storage

Cache Guard

The use of a super capacitor as a power supply is suitable for the entry-level models. The mid-range and high-end models, due to their high amounts of data, require batteries to power the mechanism.

Redundantly configured batteries inside the ETERNUS DX disk storage systems allow data in cache memory to be moved to non-volatile memory or to physical disk drives in the event of a power failure. This secured data can then be maintained in that state indefinitely.

Power outage occurred

Controller
CPU
Cache memory
Data
Battery
Non-volatile memory (Internal SSD)
Data
ETERNUS DX S3 mid-range and high-end models

Benefit:

- Cached data remains secure during any power outage regardless of the duration.
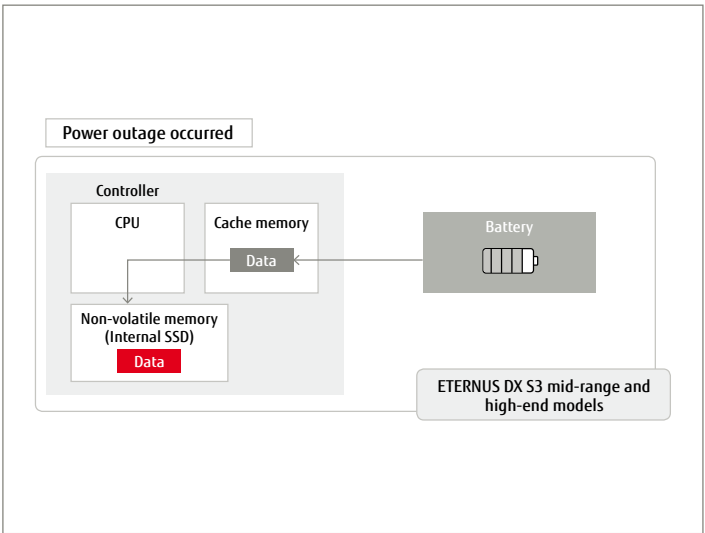
RAID: Technology that improves performance and prevents data loss due to disk failure

Damage to a company caused by disk failure is a steadily growing risk as data volumes and disk capacities increase. Storage system downtime can result in companies failing to take full advantage of business opportunities due to the management overheads involved in securing important data.

RAID technology not only prevents such data loss, but also enhances business performance.

RAID is the use of multiple disks to manage HDD data using a range of different techniques. These are divided into various levels. They all differ in terms of data deployment and the type of redundancy offered. It has also become popular to mix and match the various RAID level technologies in order to provide more specific cost reductions and performance enhancements. This document only concentrates on the main RAID levels that are widely used.
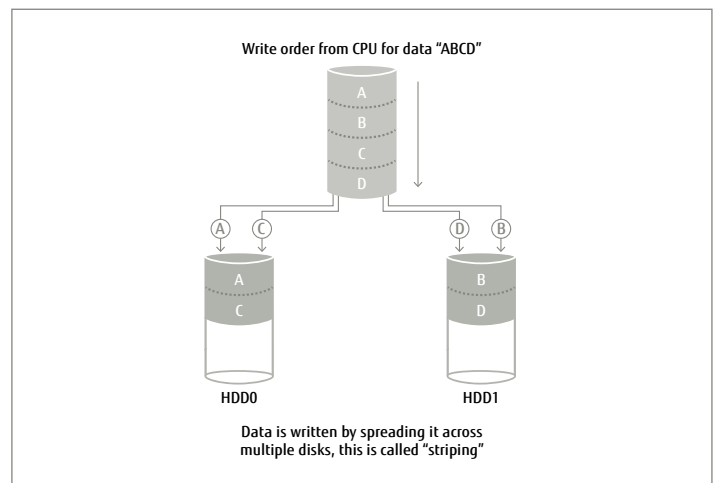
RAID 0

RAID 0 divides data into block units and writes them in a dispersed manner across multiple disks. As data is placed on every disk, it is also called "striping." This process enables high performance, as parallel access to the data on different disks improves the speed of retrieval. However, no recovery feature is provided if a disk failure occurs. If one disk fails, it affects both reads and writes, and as more disks are added to the array, the chance of a disk failure occurring is higher.
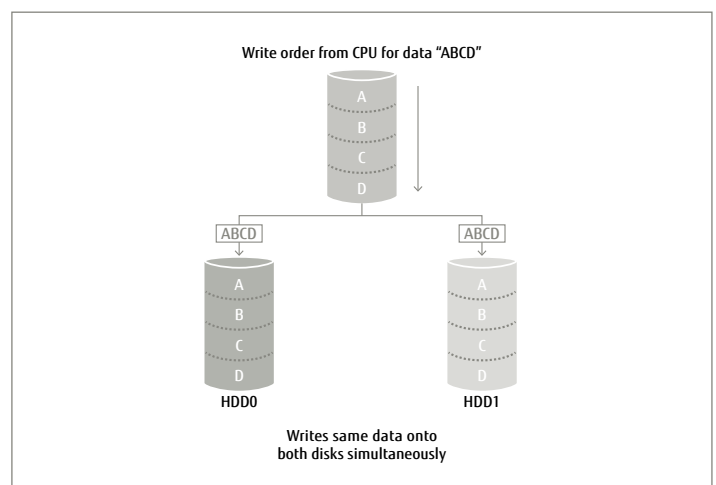
RAID 1

This level is called "mirroring" as it copies data to two disk drives simultaneously. Although there is no enhancement in access speeds, the automatic duplication of the data means there is little likelihood of data loss or system downtime. RAID 1 provides failure tolerance. If one disk fails, the other automatically takes over and continuous operation is maintained. There is no storage cost performance improvement because duplicating all data means only half of the total disk capacity is available for storage.
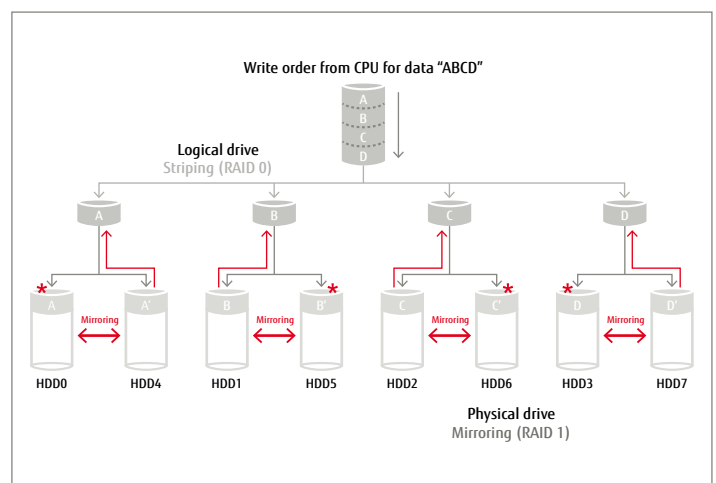
RAID 1+0

RAID 1+0 combines the benefits of RAID 0 and RAID 1. By configuring both technologies in a single array, both data duplication and improved access speed can be provided. Although this combination makes installation more expensive compared to other technologies, both reliability and high I/O performance can be guaranteed. What's more, RAID 1+0 on Fujitsu disk storage systems provides extra protection. This is because a single disk failure does not prevent striping to other disks.

RAID 0

RAID 1

RAID 1+0

RAID 5

RAID 5 is the most commonly used RAID technology today. It is based on a technique that avoids concentration of I/O on a dedicated parity disk as with RAID 4. RAID 5 divides the data and creates parity information, but the parity data is written separately across multiple disks. It enables multiple write orders to be implemented concurrently because updated parity data is dispersed across the multiple disks. This feature ensures higher performance compared to RAID 4.

RAID 5+0

RAID 5+0 stripes data across multiple RAID 5 groups using a front-end RAID 0 method. Such multiple RAID 5 striping enables one disk per group to be saved in the event of disk failure. This provides higher reliability in large-capacity configuration systems compared to a single RAID5 group. Furthermore, the rebuilding of transactions, which takes an increasingly longer time as disk capacity grows, can be executed much faster with RAID5+0 as the amount of data in each RAID group is smaller.

RAID 6

RAID 6 deploys two parity records to different disk drives (double parity), enabling two simultaneous disk drive failures in the same RAID group to be recovered. It is thus able to execute multiple write orders at the same time. This feature ensures higher performance. Furthermore, the ETERNUS DX8700 S2 is able to deploy disk drives where RAID 6 and 5 arrays are deployed across separate drive enclosures (DE) for improved reliability.

The following table summarizes the different RAID groups and highlights their main benefits:

|	Reliability	Data efficiency	Write performance
RAID 1	Good	OK	Good
RAID 1+0	Good	OK	Very good
RAID 5	Good	Good	Good
RAID 5+0	Good	Good	Good
RAID 6	Very good	Good	Good

Advanced RAID organization

In ETERNUS DX disk storage systems, each member disk drive of a RAID group is located in a different drive enclosure. This keeps the RAID group accessible even when one complete drive enclosure becomes unusable. Furthermore, with support of RAID 6 (double parity), two disk drive failures occurring in the same RAID group can be recovered. This ensures exceptional reliability.

Redundant copy

The redundant copy function enables preventive disk replacement while ensuring data redundancy. When a disk is diagnosed for preventive replacement, this function creates data using other disks in the same RAID group and writes that data to a hot spare. Once the write is complete, the hot spare takes over its functions and the disk at fault is detached. ETERNUS DX disk storage systems support two hot-spare types: global hot spare and dedicated hot spare.
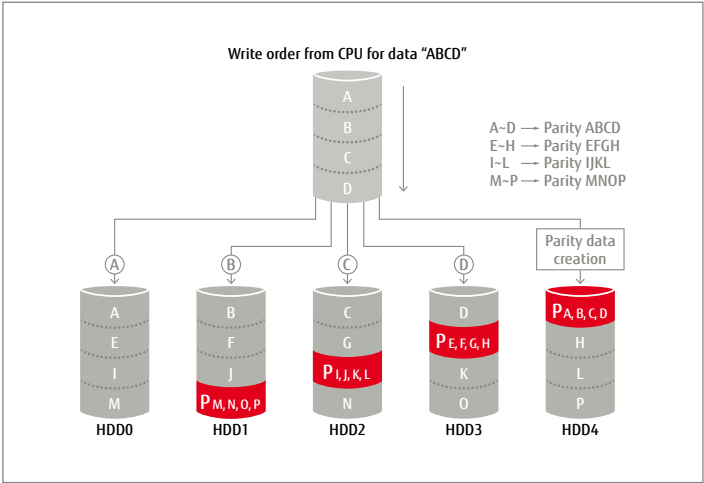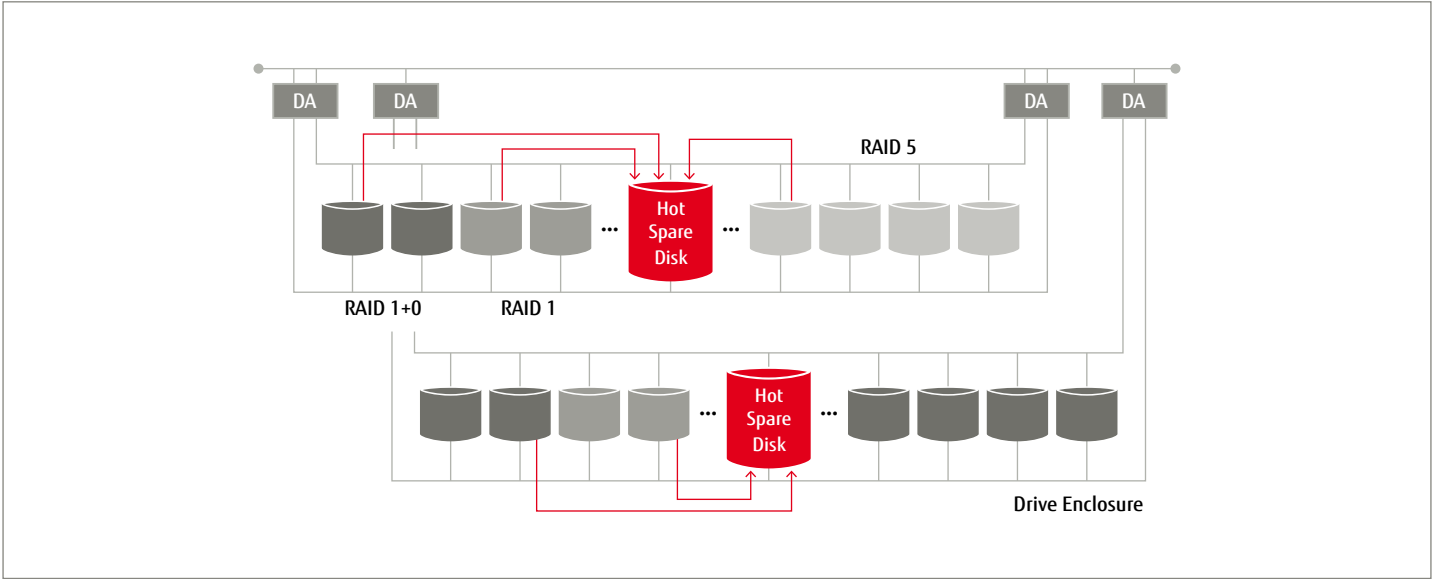

RAID 5
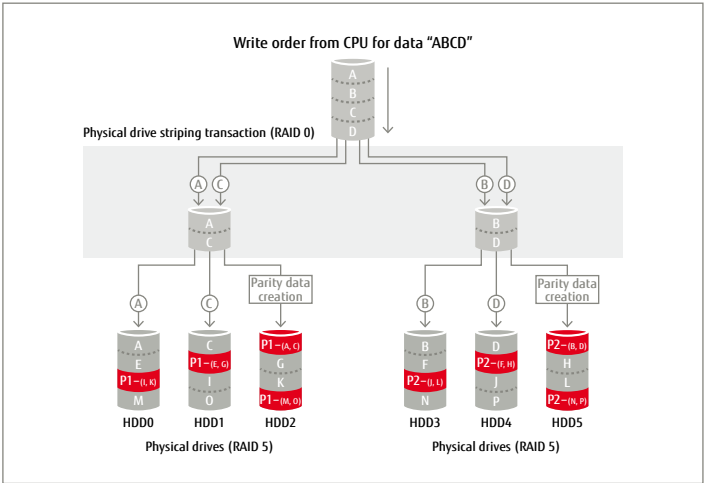
RAID 5+0
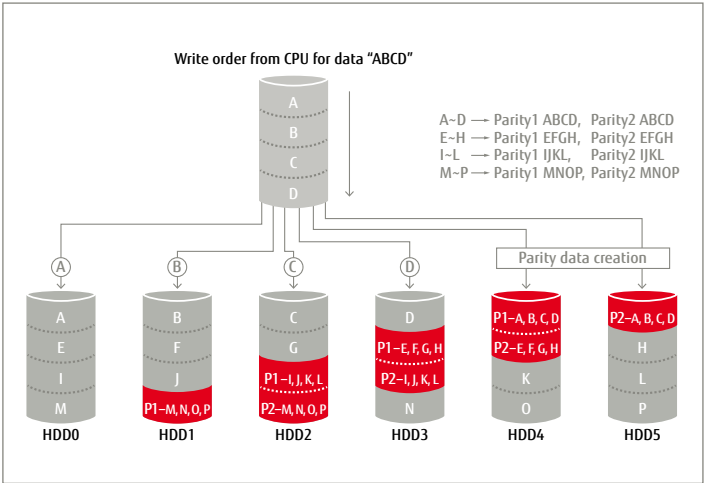
RAID 6

Global hot spare

Hot-spare disks are preparatory disk drives that are kept on active standby for use when a disk drive fails. This global hot spare function enables hot-spare disks to be used for any RAID group. When a disk drive in a RAID group fails, data on the disk drive is automatically reconstructed on the hot-spare disk in background.

Dedicated hot spare

Unlike global hot spare, which can be used with any RAID groups, dedicated hot spare can only be used with a specific RAID group. When a disk drive fails in the RAID group which has a dedicated hot spare, data on the disk drive is automatically reconstructed on the dedicated hot-spare disk. Dedicated hot spare is a special case of global hot spare for use in dedicated mission-critical environments where sharing does not offer sufficient security.
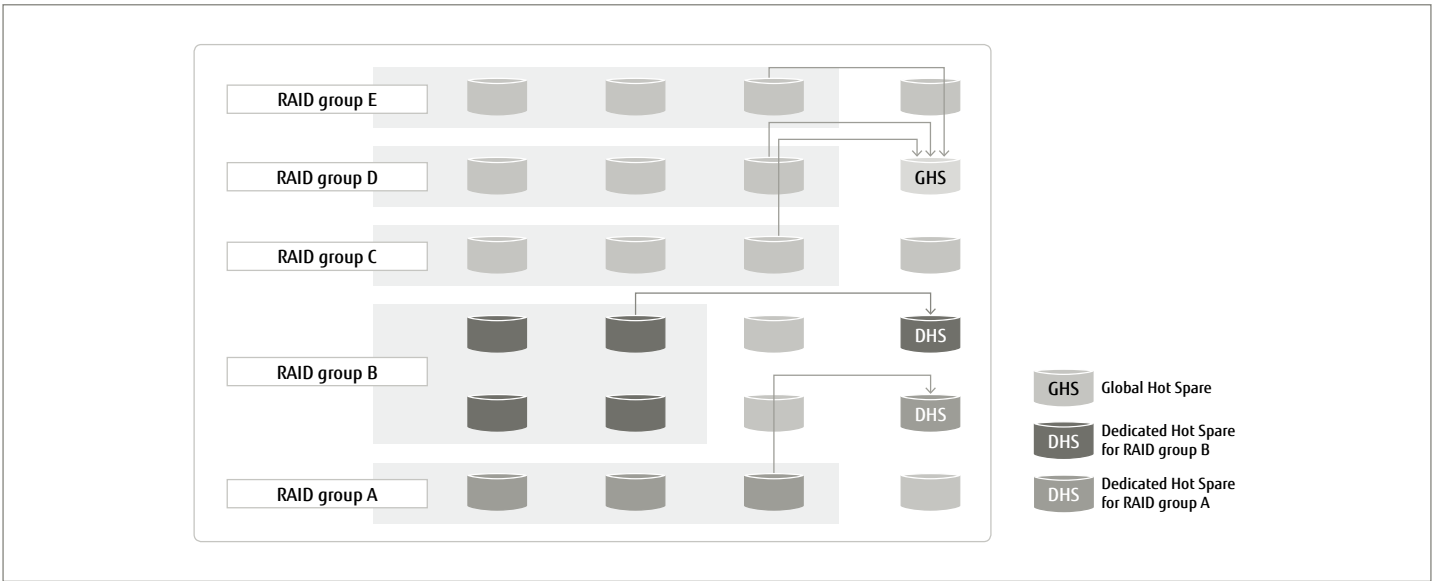
Copy-back and copy-back-less operation

After the failed disk has been replaced with a new disk, data on the hot-spare disk is retrieved (copied back) to the new disk. The copy-back function restores the RAID group while maintaining redundancy after the rebuild has been performed for the hot-spare in the RAID group.

Copy-back-less function is a feature which builds hot-spare disks into a RAID configuration after the completion of rebuild or redundant copy and the internal RAID configuration of the failed disk is changed to hot-spare disk. Immediately after maintenance and replacement, it can start working as a hot-spare disk. This feature means that the copy-back process is no longer required.

Example:

Rebuild: 4 hours, copy-back: 4 hours, disk replacement: 0.5 hours

	Working time for replacement of disk	Time until HS becomes available
Copy-back	Max 8.5 hours (Rebuild+Replacing Disk+Copyback)	8.5 hours (Rebuild+Replacing Disk+Copyback)
Copy-back-less	Max 4.5 hours (Rebuild+Replacing Disk)	4.5 hours (Rebuild+Replacing Disk)

The administrator can choose between copy-back and copy-back-less operation via the GUI.

Benefits:

- Disk replacement time is drastically reduced
- RAID availability is improved as the hot-spare downtime is reduced
- I/O performance degradation during the copy-back process can be avoided


Copy-back

Copy-back-less
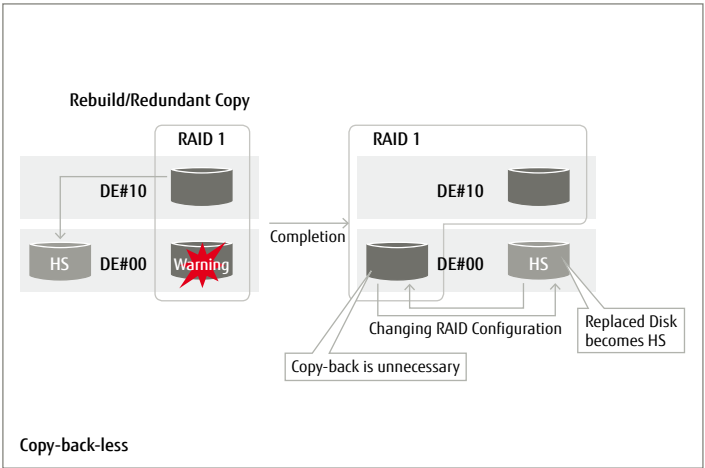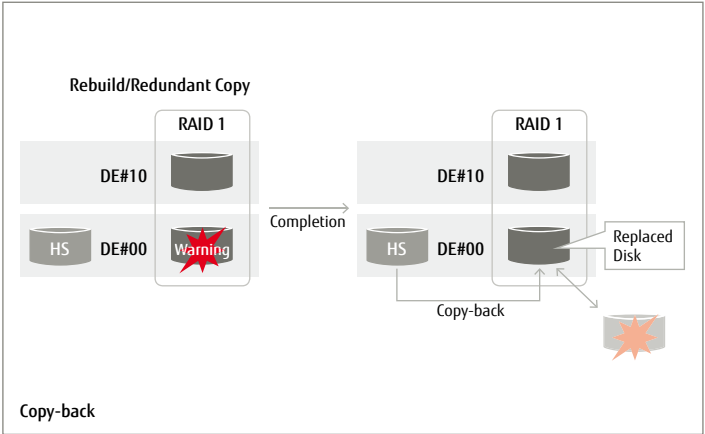
Fast rebuild

Rebuild is a process that restores the contents of the failed disk from the remaining normal disks. If a hot spare disk (HS) is installed when a disk failure occurs, the rebuild process is performed from the hot spare and redundancy is restored.

Fast rebuild is a technology to shorten the rebuild time. No exclusive hot spare disk is available, but a reserved area for rebuild is prepared in each disk in the RAID group, as shown in the figure below.

When a disk failure occurs, rebuild on reserved areas in multiple disks is performed simultaneously, unlike the traditional rebuild function where it is carried out disk by disk. This reduces the rebuild time for a 1TB disk to 90 minutes instead of 9 hours with traditional rebuild. RAID-6 is supported. Fast rebuild requires copy-back after a disk has been replaced. This feature will be available in the second version of the S3 release.
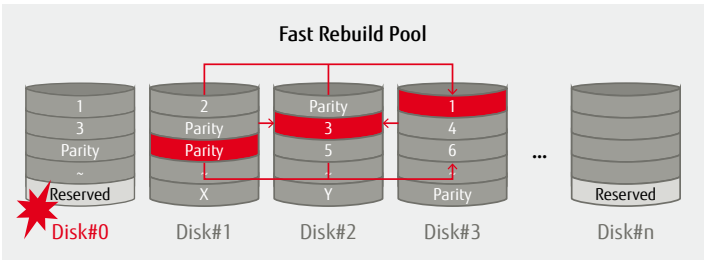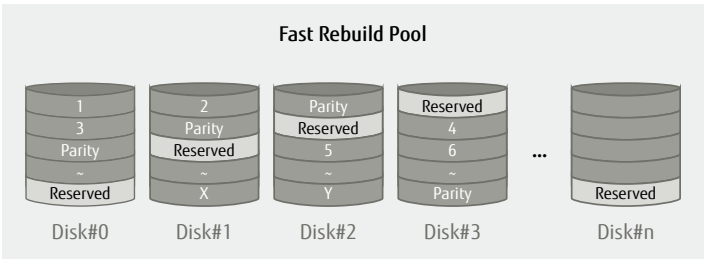
Data encryption

Due to various data protection laws, enterprise information and the security involved has become much more important from a corporate social responsibility standpoint. Laws and internal guidelines require that access to relevant stored data is restricted only to authorized users and that sensitive information is protected against unauthorized or accidental access. ETERNUS DX disk storage systems provide data encryption functions to address such requirements.

Controller-based encryption

Data can be automatically encrypted inside disk storage systems using high-security 128-bit AES technology and Fujitsu Original Encryption. This not only ensures that data is protected during use – it also ensures security during data transfer to off-site archive facilities.

Fujitsu Original Encryption is a unique encryption scheme that encrypts drive data in ETERNUS DX. Encryption is on a LUN basis. It comes at no extra cost and provides some key benefits in comparison with 128-bit AES encryption, such as:

- Less performance degradation
- Closed technology ensuring higher security

Robust security using SSL / SSH

The ETERNUS DX S3 series supports SSL (Secure Socket Layer)/SSH (Secure Shell) for encryption and secure transfer of data over a network. Normal data transfer without encryption bears the risk of possible unauthorized accesses from malicious web browsers and CLI that appear authorized yet are attempting to steal or manipulate data.

SSL enables a secure transfer of important data using SSL server certification (public key and secret key) on both the browser and web servers. SSH encrypts data using common key encryption mechanisms (DES, AES) when it is forwarded from one computer to another via a TCP/IP network. SSH achieves high data security by also hiding the common key using public key encryption mechanisms. Encrypted communication between ETERNUS DX systems and user terminals equipped with these technologies prevents the manipulation and theft of important information.
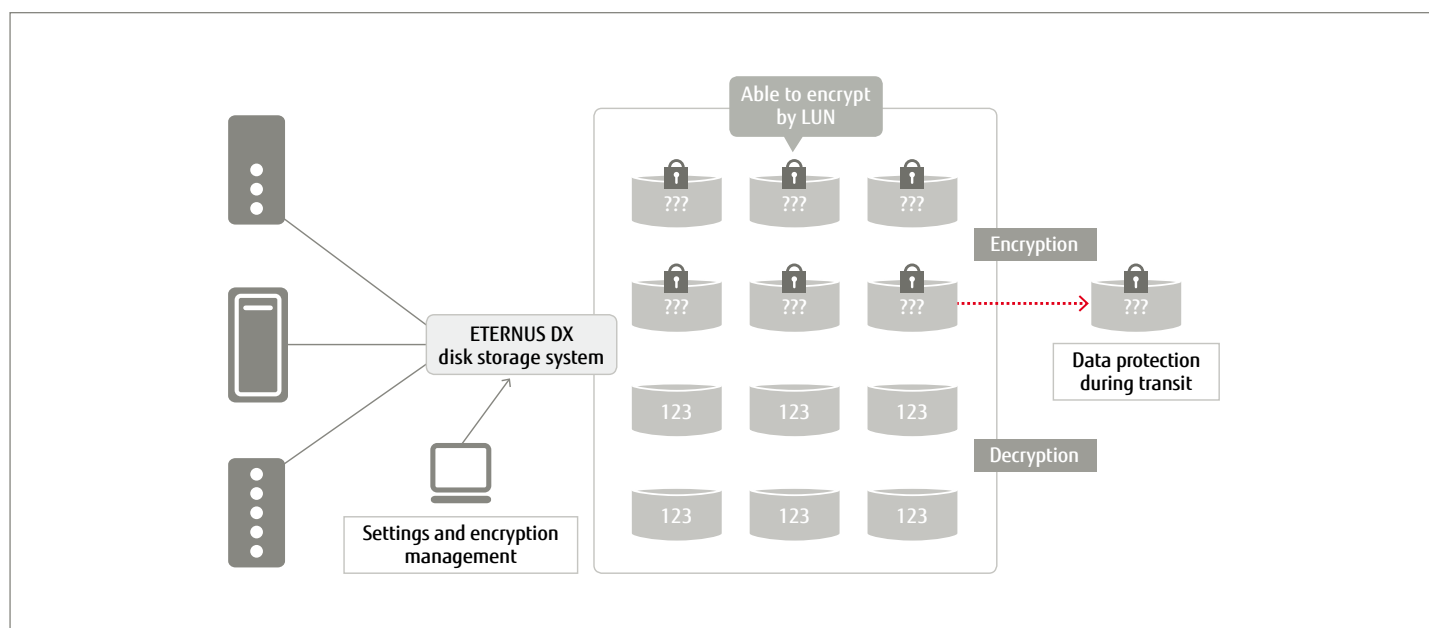
Self-encrypting drives (SED)

In order to ensure full data security, the ETERNUS DX family supports self-encrypting drives (SED). Self-encryption means that all the data transferred to the storage medium is automatically encrypted internally before the data is written – and vice versa. When the data is read from the storage medium, the data is automatically decrypted into plain text. All data passing the interface between the host controller and the disk drives interface is in plain text. The internal encryption process is transparent for the host. All read/write operations for the host are business as usual. The encryption is via a process where plain text is encrypted in order to hide its meaning. The plain text is encrypted (cipher text) when it is written to the disk and decrypted (deciphered) back to the original text when it is read from the disk. The encryption and decryption engines use the same secret internal data encryption key for this process.

The SED uses two methods for the encryption/decryption process:

■ **The internal data encryption key**
Each SED generates an internal data encryption key in the factory, which is embedded in the drive and cannot be read out or deleted. The encryption key can be modified to destroy or delete the data.
■ **The algorithm of the encryption/decryption engine**
The algorithm is a standard known as the Advanced Encryption Standard (AES), which is recommended by the US government. There are two versions of this standard: AES-128 and AES-256. The numbers 128 and 256 refer to the bit size of the encryption key used by the algorithm.
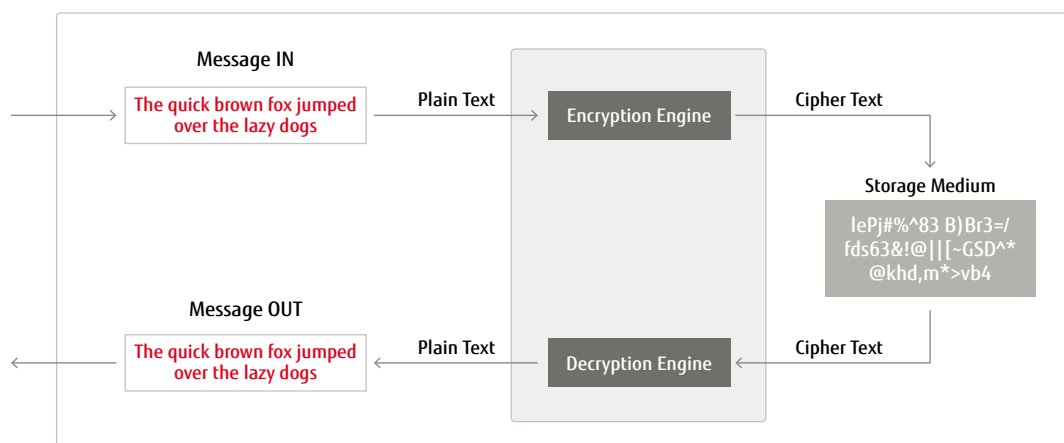
Offline storage migration

When, for example, a storage system is replaced, storage migration allows logical volume data to be moved from one ETERNUS DX storage system to a new system without involving the host. In this process, the new ETERNUS DX storage system (migration destination) connects directly to the existing ETERNUS DX storage system (migration source) in order to copy the data in the logical volume on a block level basis. Access from the host is suspended during data copying. No additional costly software or licenses are needed for the storage migration.

Storage migration can be performed just by changing the operating mode of the migration destination channel adapter (CA) port from normal CA mode to initiator mode. The destination can thus obtain data from the source. The path between the migration destination and source can be direct or via switch. Path redundancy is also supported in order to ensure higher reliability. The progress of data migration can be monitored from the GUI. Functions, such as pause, suspension and resume, are also available.

Compare functions exist in order to verify that the data migration has been completed without any errors:

■ **Quick compare:**
compares only several data blocks from the top of a volume
■ **Full compare:**
compares all data blocks in a volume

Having completed the data migration process, the operating mode of the destination CA can be changed back to CA mode and the host is connected to the new storage system.
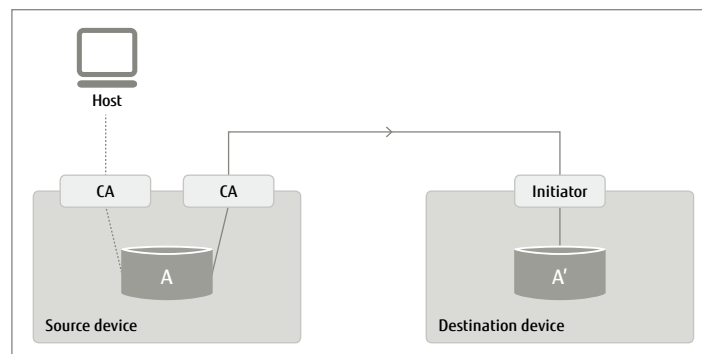
Thin Provisioning

Storage system growth continues year on year. Due to concerns about having sufficient storage capacity, users tend to deploy more physical storage than they actually need – "just to be safe." However, in practice the allocated capacity is often underutilized. Industry research organizations have even stated that in some cases only 20% to 30% of the provided capacity is actually used.

Thin provisioning technology has thus been developed to enable effective use of available storage capacity for better investment utilization. It reduces physical storage deployment by using virtual storage techniques that maximize available capacities.

Thin provisioning only assigns the total overall user capacity as virtual storage. The actual physical disk capacity is allocated as and when needed. All physical disks are managed as a single disk pool and allocated according to the amount of data written to the virtual volumes. This reduces the amount of unused physical disk capacity and supports much more effective storage operations. Furthermore, predefined thresholds avoid storage capacity shortages by issuing a warning that additional physical disks need to be added.

Example: A user requests 10 TB of resource allocation from the server administrator. While 10 TB of physical storage capacity may eventually be needed, current usage suggests that 2 TB of storage is sufficient. The system administrator therefore prepares 2 TB of physical storage, but allocates a 10 TB virtual volume to the server. This means that the

server can start using the existing physical disk pool which is only around 1/5 of the virtual volume. This "start small" approach enables more effective use of storage capacity. As more physical capacity is required to support the virtual volume (as shown in the diagram), existing physical volume capacity is consumed. In order to avoid a capacity shortage, the physical disk pool is monitored using a predefined usage threshold. For example, by defining 80% of the entire disk pool as the threshold, an alarm tells the administrator to expand the number of physical disks when that amount of 8 TB in our example is reached. This means that the new drives can be added without stopping the system, ensuring continuous system operation.

Benefits:

- Lowers initial investment by using storage capacity very efficiently (start small)
- Does not require any changes to storage capacity settings for changes on demand
- Reduces operational costs by integrating storage with virtualization
- Reduces overall power consumption via reductions in over-provisioning
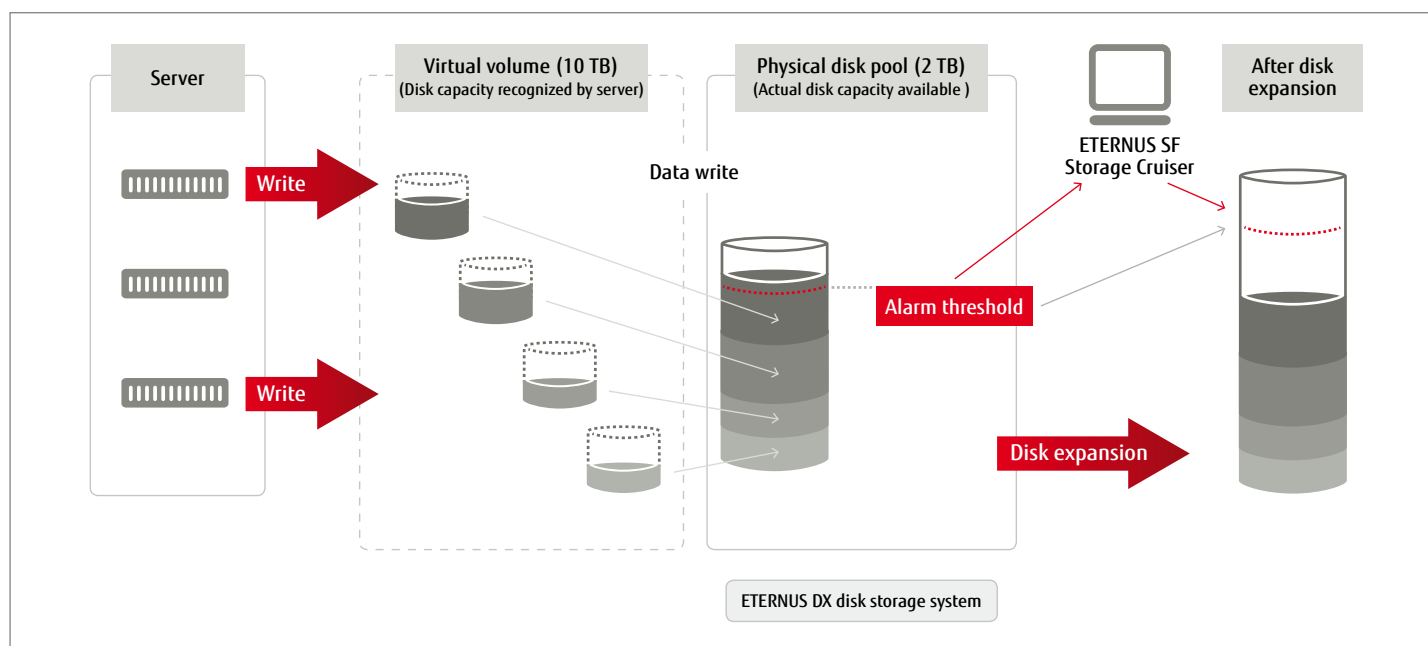
ETERNUS SF Management Suite

The ETERNUS SF Storage Management Software is the perfect fit for ETERNUS DX disk storage systems. It reduces the total cost of ownership, simplifies monitoring and management and helps to achieve business continuity. The following ETERNUS SF functionalities are available to meet individual requirements:

■ **ETERNUS SF Express:** to simplify the setup and administration of ETERNUS DX entry-level models
■ **ETERNUS SF Storage Cruiser:** For advanced management, stable operation of all ETERNUS DX storage systems and SAN management
■ **ETERNUS SF AdvancedCopy Manager:** To leverage the high-speed internal and external volume copy functions of ETERNUS DX

ETERNUS SF optimizes storage resources in complex IT environments. Its flexible and transparent license model guarantees that customers only pay for the functions they require. Furthermore, controller-based licensing ensures that customers can add capacity without extra license costs. The license key protection function for ETERNUS SF is structured according to functional and upgradable modules for each individual customer requirement. Each module and additional dedicated features can be activated by entering a license key in the ETERNUS SF integrated manager. When adding a license to an existing software installation,

the current environment can be used without any modifications. This framework eliminates the need for separate and isolated tools in order to manage each class of storage device.

Easy configuration and operation of the storage system
A high degree of automation and the easy-to-learn, easy-to-handle graphical user interface facilitate the management tasks. Storage resource optimization (including the implementation of policies for enhancing storage integration and operation, error discovery, health monitoring, capacity management, provisioning, cloning, backup, restore, and disaster recovery) are displayed with a consistent, user-friendly look and feel.

The user interface is based on a three-step operational concept: **Select** the type of action you want to perform from a well-structured menu, **check** the status and **execute** the task.

Helpful wizards, system data visualization and automated routine administration tasks reduce the monitoring and management workload. The supplied web-based startup wizard is all that is required for initial installation. Ongoing changes and system expansion are supported by equally intuitive wizard-based operations.

1 Choose from menu

2 Check status

3 Define action
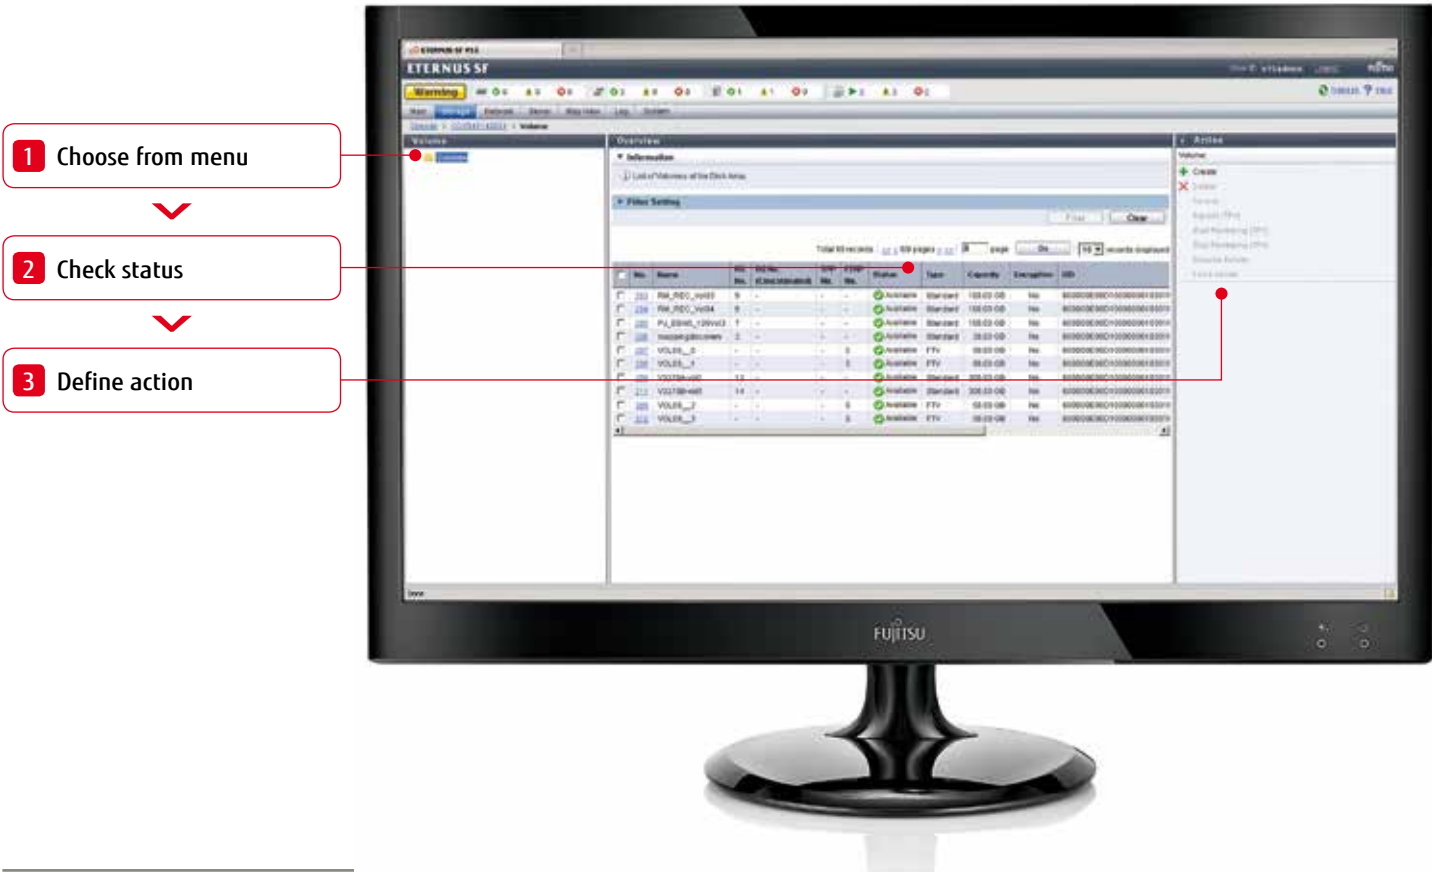
ETERNUS SF Express – Simplified management for entry level

ETERNUS SF Express is bundled with ETERNUS DX entry-level disk storage systems for simplified storage system management and maintenance. For those users installing storage systems for the first time, Express provides an easy-to-use wizard setup panel. This ensures that all the parameters and settings, including the connection of disk storage systems with all the application servers, are simple and straightforward.

Easy upgrade to the full feature set

Simply entering a license key for the ETERNUS SF software activates a shift to enterprise storage management software, which comprises a more comprehensive set of functionalities for storage resource management and advanced replication capabilities. This supports the cut-off procedures to reconstruct the environment while protecting storage investments.

ETERNUS SF Storage Cruiser

ETERNUS SF Storage Cruiser simplifies the setup of disk storage systems and the management of ETERNUS DX storage environments. Its ability to support complex storage configurations and settings via the user-friendly graphical interface enables administrators to implement storage environments easily.

Once operational, ETERNUS SF Storage Cruiser supports stable system operation by providing centralized management for the entire storage environment, including servers and Fibre Channel switches. ETERNUS SF Storage Cruiser also provides all storage system lifecycle functions, from configuration to maintenance. The resulting improved availability and stable operation reduce operation management costs, help support Green IT and enable the implementation of virtualized IT systems.

ETERNUS SF AdvancedCopy Manager

ETERNUS SF AdvancedCopy Manager (ACM) enables high-speed and high-reliability backup/restore and replication operations. To achieve continuous 24 x 7 business continuity, ETERNUS SF ACM takes over all the tasks involved in implementing a disaster-resilient storage infrastructure.

ETERNUS SF ACM thus significantly reduces operational downtime, especially in large-volume, transaction-intensive environments. The software can be deployed in conjunction with a wide range of mainstream database environments, including Oracle Database, SQL Server, DB2 and Exchange Server, thus providing reliable high-speed data copying without going offline and without downtime.

Data center managers can benefit from:

- High-speed backup operations
- Fast backup and restore, regardless of data volume
- Disaster recovery with remote copy
- Fast recovery of business operations
- Database backups without interruptions or heavy loads
- VMware environment backup/restore
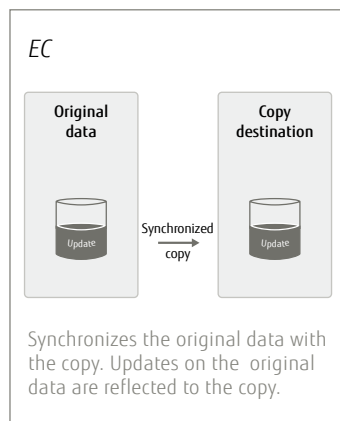- Hyper-V guest online backup

Advanced Copy Functions

The advanced copy functions allow the disk storage system to carry out high-speed copy operations without any need to draw on server CPU resources. Advanced Copy functions are used to copy a business data volume to a separate copy volume at any point in time, quickly and within the disk storage system. Once the copy is complete, the copy volume can be separated from the business volume in order to ensure that no further updates to the business volume are applied to the copy volume. This allows the copy volume data to be backed up to a tape device as a point-in-time copy of the business data while normal operations continue.

ETERNUS DX systems support the two distinct data copy modes: **Synchronous high-speed copy** and **Snapshot high-speed copy.**

- Synchronous high-speed copy maintains the equivalent status for a transaction volume and backup volume. The two copy types available are: EC (Equivalent Copy) and REC (Remote Equivalent Copy)
- Snapshot high-speed copy creates a snapshot of data. The copy types available with this function are: OPC (One Point Copy), QuickOPC, SnapOPC and SnapOPC+

Equivalent Copy – EC

Equivalent Copy creates and maintains a copy volume (mirror) synchronized to the business data volume, until they are "detached" (mirror suspend or break) so as to enable the start of a backup operation.

As the detached copy volume contains the same data as the business volume up to the time synchronization stops, it can be used as a point-in-time copy for

EC

Original data → Synchronized copy → Copy destination

Synchronizes the original data with the copy. Updates on the original data are reflected to the copy.

backup to a tape device, while business operations continue on the original business volume. Two methods of breaking the mirror are provided. If a complete break is made, the copy volume is fully detached and any subsequent use of EC copies all the operational data again to a new copy volume before maintaining the mirror in synchronization. If the Suspend/ Resume functions are used (mirror suspend), the same copy can be resumed. In this case only the differences between the business and the copy volume are copied until synchronization is once again reached and then maintained. The method of re-establishing synchronization depends on how long the copy volume has been detached. If the suspension time is relatively short, the suspend resume method is the quickest.
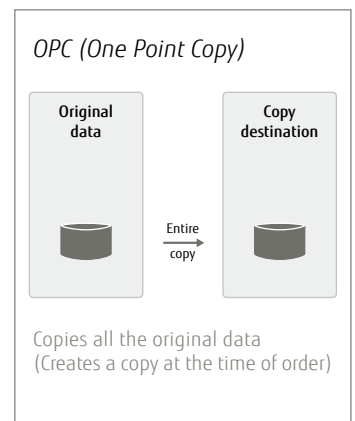
One Point Copy – OPC

One Point Copy (OPC) enables the creation of a high-speed copy of an entire business data volume at any specific point in time.

Unlike EC, with its data synchronization (mirroring) capability, the copy volume created by OPC is always separated from the business volume (a point-in-time copy) and never reflects ongoing updates to the business data.

OPC (One Point Copy)

Original data → Entire copy → Copy destination

Copies all the original data (Creates a copy at the time of order)

This means the copy is a snapshot of the business data at the time the OPC request is issued, and as such can be backed up to a tape device in parallel with the ongoing business operation. However, for a subsequent backup, OPC requires that all the data is copied once again. QuickOPC is provided to copy only the updated data.

Benefits:

- Ensures significant reductions in time by enabling backup data to be created and maintained while normal business operations are carried out in parallel.
- Enables copy data to be detached and used for other processes, data discovery, batch enquiries, system testing, at any point in time without impacting operational processes.
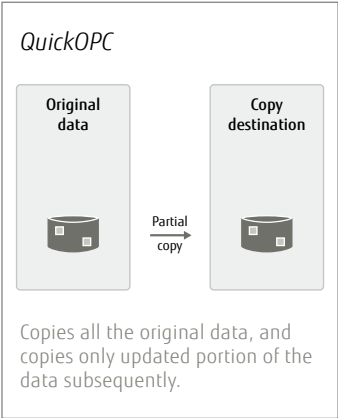
Benefit:

- High-speed backup and near real-time disk-to-disk-to-tape backup extends business operation hours and availability.
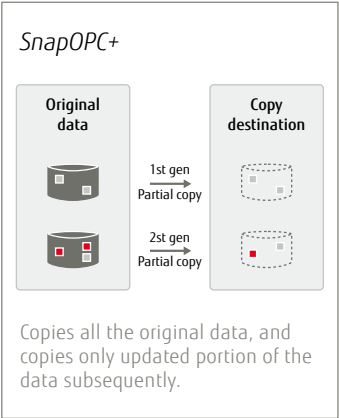
Quick One Point Copy – QuickOPC

QuickOPC initially copies all the business data volume to a copy volume. Subsequently, it only copies updates that occur on the business volume (the differences).

This reduces copying time, particularly with large databases, enabling high-speed backup operations. Both the business volume and the copy volume have the same data size. This process is particularly suitable for backup operations on mission-critical databases where robust data security is essential.

QuickOPC

| Original data | Copy destination |

Partial copy →

Copies all the original data, and copies only updated portion of the data subsequently.

Partial copy function – SnapOPC/SnapOPC+

SnapOPC only copies the "before" image of data that is being updated to the copy volume. By only copying data subject to change, the copy volume capacity can be significantly reduced in comparison to the original business data volume.

Furthermore, SnapOPC+ provides generation management of the updated data. The difference between SnapOPC+ and SnapOPC is that SnapOPC+ updates data only as history information, while SnapOPC stores the data redundantly. Logging as history information can provide disk-based generation backup using a smaller copy volume capacity.

SnapOPC+

| Original data | Copy destination |

1st gen Partial copy
2st gen Partial copy

Copies all the original data, and copies only updated portion of the data subsequently.

Benefit:

- Physical copying after the initial full copy requires significantly less time.

Benefit:

- Enables efficient copy retention because the overall copy volume capacity is much smaller.

Disaster Recovery

A loss of data due to human error or a natural disaster, such as earthquakes or fire, poses serious risks for IT administrators. Data has to reside at different geographical locations in order to ensure smooth recovery should disaster strike. ETERNUS DX storage systems support a number of features that guarantee reliable operation in disaster scenarios.

Remote Copy using Fibre Channel Interface

REC (Remote Equivalent Copy) provides a server-less remote mirroring function, which ensures fast recovery if the primary disk storage system site is not operational.

Remote Advanced Copy for Storage Area Networks (SAN)
By using Fibre Channel interfaces, Remote Advanced Copy can provide low-cost remote site support between a primary storage device and a secondary device.

Extended Remote Advanced Copy for Wide Area Networks (WAN)
Extended Remote Advanced Copy uses a combination of a Fibre Channel switch and WAN converter to cover very long distances over WAN. Replicated data can be located at a remote site hundreds of miles away from the primary site. This provides high security for the protection of critical data from any kind of disaster.
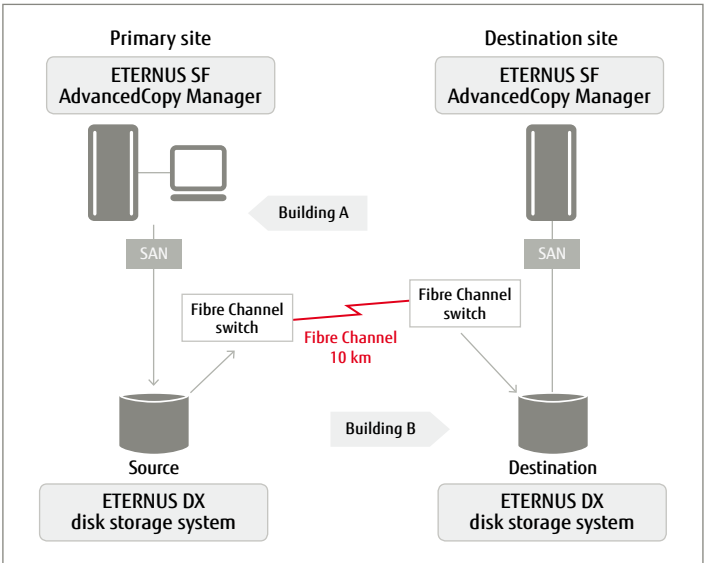
Furthermore, ETERNUS DX S3 supports replication to existing models as well as N:1 integrated backup. These capabilities enable flexible system configuration matched to customer requirements. Remote copy using iSCSI interface is supported as well.

Remote Copy Disk Buffered REC over low-bandwidth networks

A consistency model is provided in order to provide remote copying over low-bandwidth networks. It uses part of the cache memory as a buffer (the REC buffer). Data is then copied to the destination device and compiled on a block basis after accumulating the I/O from multiple REC sessions in the REC buffer for a specific period. Use of this mode provides the control to maintain transfer integrity even when data is transferred out of sequence, due to transfer delays as it travels to the destination device via WAN.

Furthermore, disk-buffered REC can be used if the cache memory capacity becomes insufficient due to instabilities in the link or increased traffic. This supports temporary increases in updated data using the larger buffering capacity of hard disks.

Transparent Failover

As some data is mission-critical, it must always be accessible. In order to ensure data availability, even in the event of a system or site failure, transparent failover will be introduced in the second version of the S3 release based on REC (Remote Equivalent Copy) functionality.

The basic concept is constructed around deploying a secondary storage system and a monitoring server. As long as the primary storage system is running, data is transferred from it to the secondary system via a synchronous REC function. The monitoring server continuously checks the status of the primary storage. If a failure is detected, it runs the failover logic and the primary storage information (e.g. LUN ID) is taken over to the secondary storage in order to recognize the volume transparently by the I/O server. Hence, operations run smoothly ensuring business continuity.

There are two scenarios that can be implemented:

Primary storage system failure

This scenario is targeted to overcome the situation should the primary storage system fail. The secondary system and the monitoring server are deployed at the same site as the primary storage.
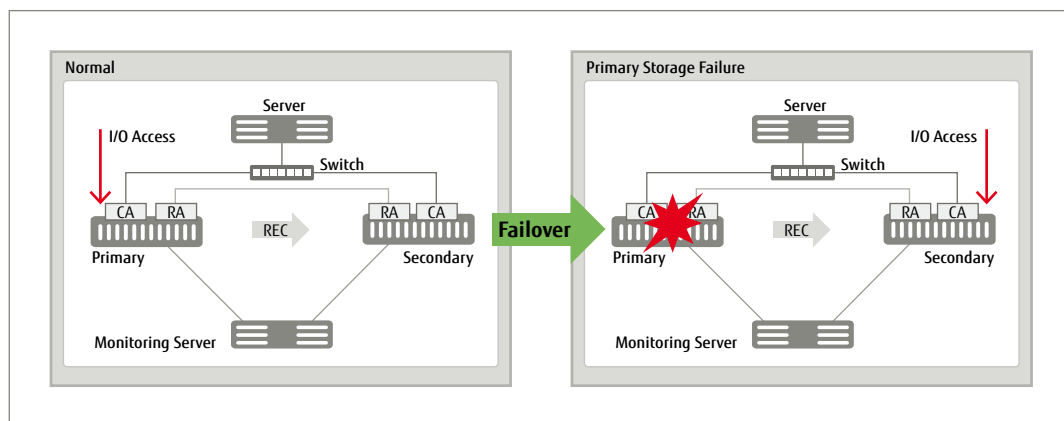
Primary site failure

As the first scenario does not cover a possible failure of the primary site, e.g. due to a natural disaster or man-made error, this scenario is supported in order to provide higher reliability and data availability. The secondary storage system is deployed at a different site, which can be up to 100 km away from the primary site. The monitoring server is located at a different site as well.
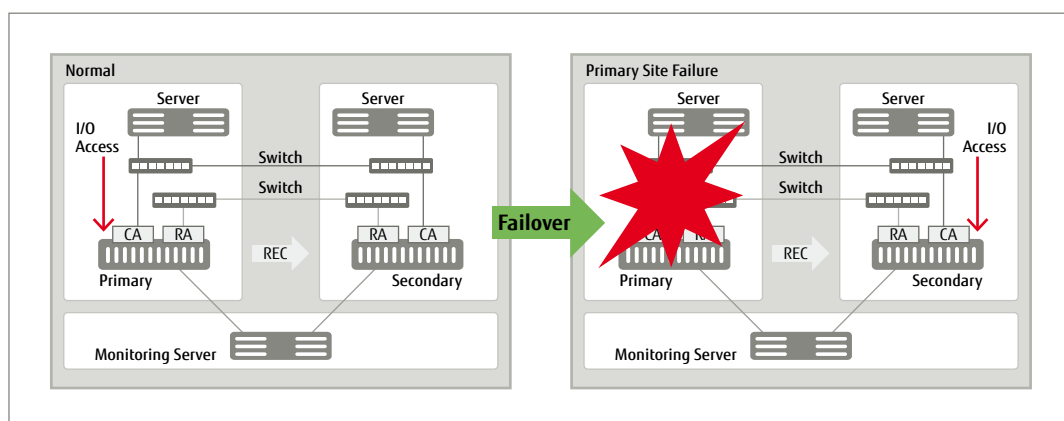
Storage efficiency features
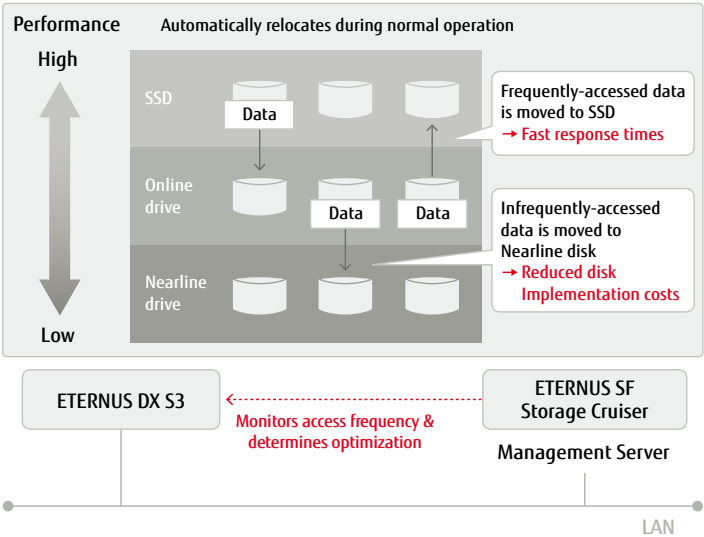
Automated Storage Tiering (AST)

The use of long-term data storage within organizations is increasing due to many laws and government regulations governing data retention, not to mention internal data audit requirements. One problem is the access frequency of older information, which typically decreases over time. But because of difficulties in managing such access, information which is well-suited for low-cost long-term storage is often left to reside in more expensive high-performance storage systems.

Automated Storage Tiering (AST) is a feature that monitors data access frequency in mixed environments that contain different storage classes and disk types. The storage administrator does not need to classify data or define policies. Once the tiers are configured, the ETERNUS DX storage system does all the work, enabling the storage administrator to focus on other storage-related responsibilities. The automation of tiered storage means that multiple storage tiers can be managed as a single entity. It helps ensure that the right data is in the right place at the right time.

The ETERNUS SF Storage Cruiser controls the destination and the arrangement of data, monitors its access frequency, and automatically relocates the data between drives to the most appropriate storage devices. This storage hierarchy control offers significant investment optimization and reduces storage costs by matching storage system capabilities and applications sensitivity to performance, availability, price and functionality. Infrequently used data and non-essential copies of primary application data, e.g. point-in-time snapshots, replication copies and data mining are located on Nearline drives, which have large capacity but are less expensive. For high priority applications, the best performance and response times for important information are improved by locating frequently accessed data on high-performance SSD. The overall arrangement of data on the different drive types is thus optimized regarding costs. The relocation of data is completely transparent to servers and applications and is carried out without any changes in server settings.

Data can be moved in 252 MB chunks providing high efficiency, as less data with low performance requirements would unnecessarily be moved to faster, more expensive disk drives. On the other hand, it guarantees that data demanding high performance will be moved to the fastest disk drives.

Calendar-based scheduling enables the exclusion of off-day performance, such as weekends and public holidays, from the tuning process.
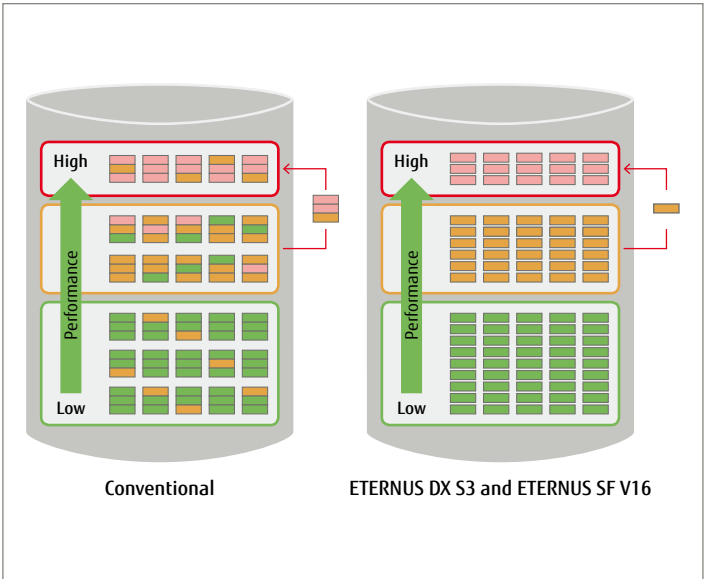
Benefits:

- Reduces data management time and costs due to automated operations
- Provides optimal performance while reducing costs
- Operational data reallocation policies can be flexibly set to meet requirements
- Reallocations are performed without changes in server settings

| Conventional | ETERNUS DX S3 and ETERNUS SF V16 |

Application I/O prioritization – Quality of Service

A prerequisite for any storage consolidation strategy is the ability to host multiple applications on a single storage platform without allowing the actions of one set of users to affect the I/O performance of others.

Potential problem areas for shared storage access include:

- Workloads with I/O and cache conflicts, such as online transaction processing (OLTP) and data warehousing
- Tiered storage access restrictions, such as development and production applications
- Peak processing demands for critical applications versus maintenance activities, such as backup or database reorganization

The ETERNUS DX Quality of Service feature with application I/O prioritization resolves these issues and enables the consolidation of multiple tiers of applications in a single storage system.

It sets performance limits for each connected server according to its priority. By prioritizing data access and dynamically managing any I/O conflict, high performance can be guaranteed for high-priority applications, and at the same time capacity is used more efficiently, thus increasing storage utilization without sacrificing performance. The QoS policies allow the user to specify the expected I/O patterns of each application (random, sequential, read or write-based, and mixed).

An example is shown in the figure below. Two servers are connected to an ETERNUS DX storage system. Server B is granted a higher priority than server A. Accordingly, limits for I/O requests from both servers are set and server B has a higher limit than server A. In the event of

Benefits:

- Mapping application Service Level Agreements (SLA) to storage infrastructure
- Increased storage utilization by combining different workload profiles
- Allows service providers to guarantee a specific QoS and charge accordingly

increased workloads on the low-priority server A, the system limits the I/O performance at the predefined level and the performance of the high-priority server B is not affected. Thus the required I/O performance is guaranteed regardless of the workloads on other servers with lower priority.

The Quality of Service functionality provides higher degree of automation to ensure simpler and more intuitive settings.

The ETERNUS SF Storage Cruiser Quality of Service management option sets values based on performance requirements and dynamically adjusts the values along with the result of performance monitoring.

This feature makes it easier for the user to start the settings. Furthermore, the automatic tuning ensures that the values used are more accurate, resulting in better service level fulfillment.
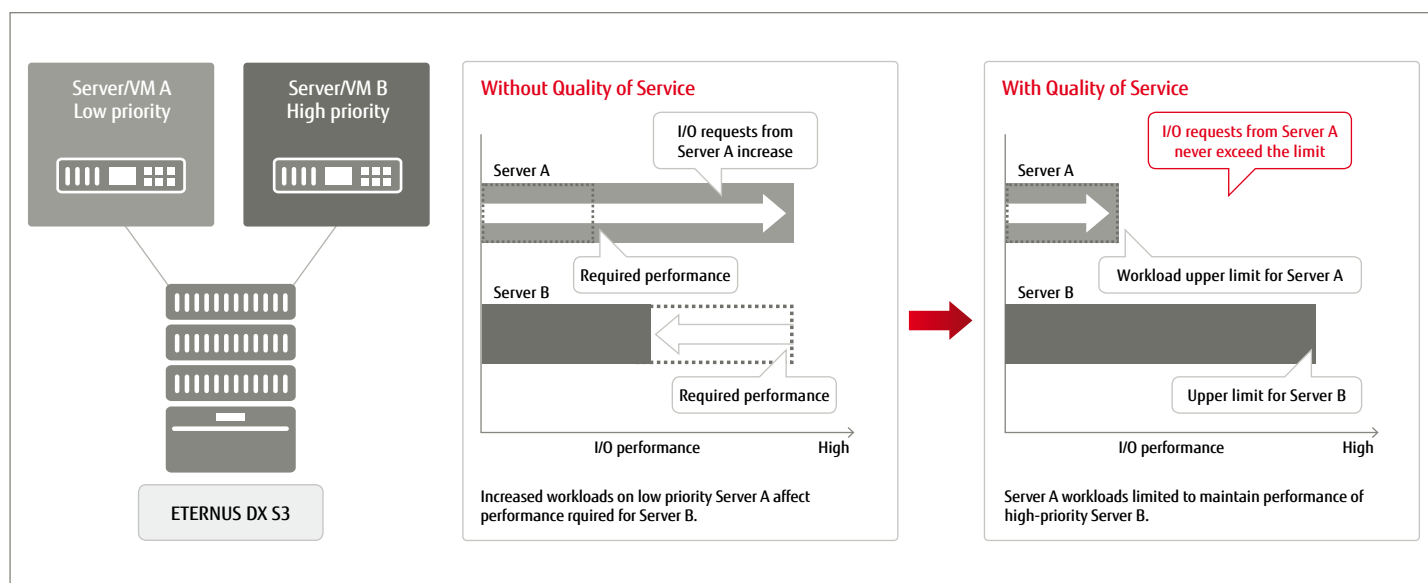
Unified Block and File Data Access

Reporting

Another important aspect of consolidation is the ability to provide reports on the storage usage of different applications, as it simplifies the billing procedure done by the storage administrator.

From ETERNUS SF Storage Cruiser V16 on, reporting functionality is provided for ETERNUS DX S2 & S3 models. Reports of storage space assignment to servers are automatically generated periodically (monthly). Administrators can generate reports and manually set parameters such as start- and end-time.

The reports are generated in XML format enabling customers to convert the data according to their preference. The reports can then be exported for billing calculation.

The next version supports CSV format.

Data is currently growing at a yearly average rate of 50%. Some industries report even up to 100%. Unstructured data, i.e. files, are a major factor in these percentages. In contrast to structured data in databases and OLTP (Online Transaction Processing) applications, it is much more difficult to anticipate the necessary storage capacity for files as their size can vary greatly.

File and block storage have been traditionally separated on dedicated systems. File storage, utilizing NAS (Network Attached Storage) architecture, has been preferred in many cases due to its deployment simplicity. On the other hand, block level storage systems with SAN (Storage Area Network) architecture, provide unmatched flexibility. They can fully concentrate on data handling whereas higher level data management, such as file services, runs outside, typically on servers.

In recent years, the boundaries between both types have started to diminish and now combine the benefits of both. Unified storage systems, which include file access on top of block access, provide higher consolidation potential by hosting structured data from databases plus unstructured data in the form of files.

The main benefits of such a unified approach are:

- Less operational and training expenditure
- Use of the same software features for both, block and file operation
- Higher storage resource consolidation

Fujitsu ETERNUS DX S3 is offering unified storage systems for the scalable entry and mid-range segments. Customers can choose between SAN-only and unified systems. Moreover, SAN-only models can be easily upgraded to the unified models.

| ETERNUS DX100 S3 | ETERNUS DX200 S3 | ETERNUS DX500 S3 | ETERNUS DX600 S3 |

Scalable unified entry-level and midrange systems

Architecture

When deploying the ETERNUS DX S3 systems for unified storage, the architecture is slightly adapted in order to handle both SAN and NAS functionalities.

In addition to the conventional SAN control, the same controller module can also process the NAS, as shown in the next figure.

Distributed file system is used to control the file management capability in the ETERNUS DX S3. Network File System (NFS) and Common Internet File System (CIFS) are supported as file sharing protocols. In the second version of S3 release more protocols will be supported, such as HTTP, HTTPS, FTP, SCP, SFTP and WebDAV.

The file system is installed on each controller module and performs shared control for a single NAS volume (Shared VOL). Hence, active-active connections can be implemented as cluster systems in CM#0/CM#1 from CIFS/NFS client.

The main features of NAS volumes are as follows:

- Thin Provisioning Volumes (TPVs) are used as NAS volumes
- Formatting a single TPV is performed with a single file system
- The Thin Provisioning Pool (TPP), to which the TPV belongs, can be selected
- SAN and NAS volumes can be mixed in the same TPV ensuring efficient operation
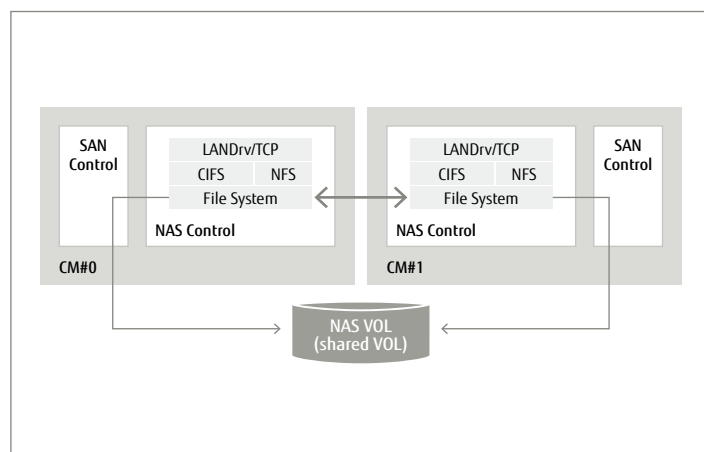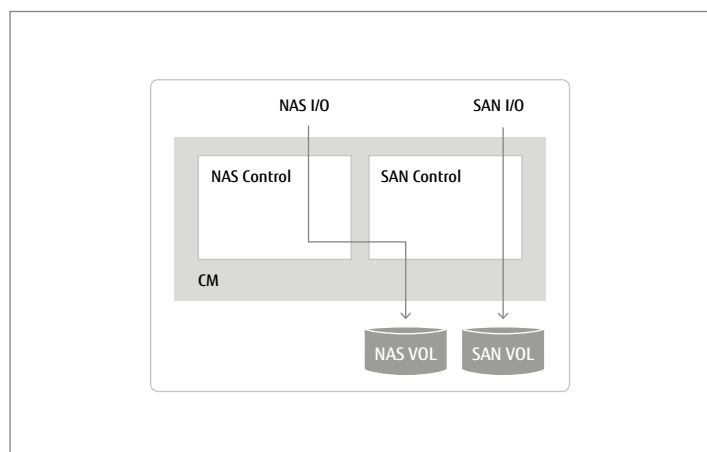
Topology

ETERNUS DX S3 supports three different topologies for the connection to NAS volumes:

Single Path Connection

This topology has the client access the NAS volume via a single path. Each NAS volume is connected via only one controller module. In the event of a CM failure, failover is not possible and the system is down.

In order to provide higher availability and reliability, the following multi-path connection topologies are also supported:


Active-Active Multi-Path Connection

In this case the NAS volume can be accessed by the client via two parallel paths using both controller modules. If one CM fails, information is taken over by the other CM and the failover operation is carried out. After maintenance for the failed hardware has been completed, the connection topology can return to the previous state.

This topology allows load balancing between both CMs.

Active-Standby Multi-Path Connection

In this topology the client can access the NAS volume only via the active path (CM#0 in the figure below). The other path (CM#1) is in standby mode. If the active connection fails, the client detects the failure and the other CM takes over. Failover operation is then carried out. After maintenance for the failed hardware has been completed, the connection topology can return to the previous state. This topology prevents performance degradation, even when one CM is not available.
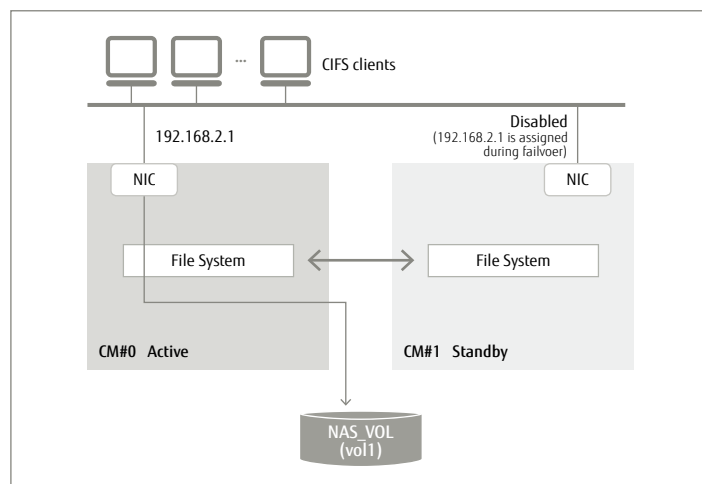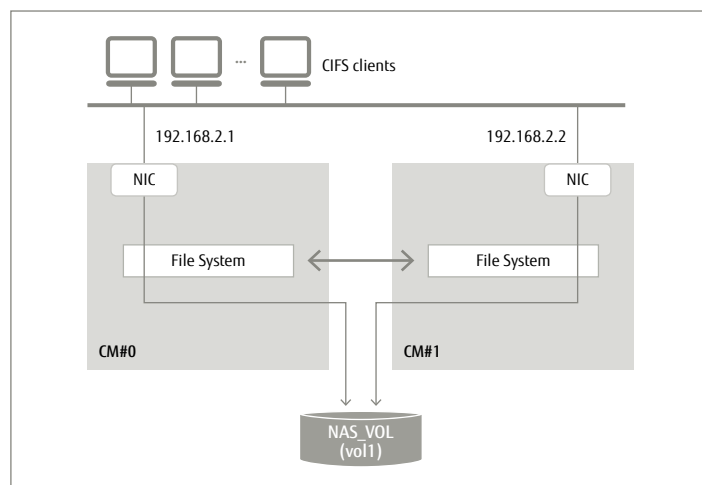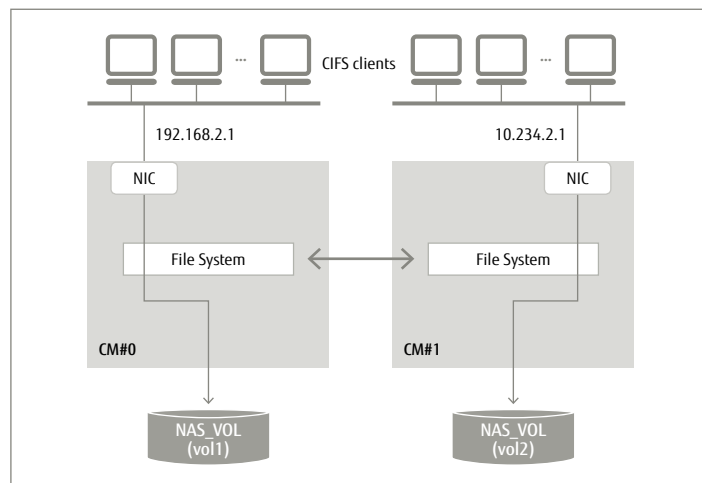
NAS File System Cache

As the file management consumes additional system resources, the unified storage systems have to provide higher performance in order to avoid quality degradation. The total system cache is thus increased for unified models and dedicated cache is assigned for the file management.

The cache area dedicated for SAN control ensures data writing on a moment-to-moment basis regardless of control from the application. On the other hand, the cache area for NAS file system writes data only when a Sync request is received from the client, according to RFC3530.

Data integrity function control for NAS I/O

In the same way as SAN control, data integrity function (DIF) control is performed for NAS I/O to improve data reliability. Data copying as well as DIF generation, removal, and checking are carried out between operating systems using hardware mechanisms without creating any additional CPU load.

If a DIF error occurs, the information is sent to the access host. For READ I/O, re-reading from the disk is performed internally.

Quota management

Quota limits are set for the number of files that can be created and the total file data size for each file owner and ownership group. The client is notified when the specified value is exceeded.

This function will be supported in the next version of the S3 release.

Security features

In addition to traditional data management, a unified storage system also performs file management functions, which requires enhanced security functions.

Authentication

In order to ensure that only authorized users can access the system, the following authentication methods are implemented depending on the protocol used.

When CIFS protocol is used to access the device and user management is performed for directories and files, authentication by an ActiveDirectory authentication server (Kerberos authentication method) must be performed.

If NFS protocol is used to access the device and user management for directories and files, authentication by an LDAP server (such as an OpenLDAP server) is required.

NAS data protection

This function allows data restoration should a data volume be lost. It collects and saves a copy of the entire directory in a specified volume. The stored data can then be used in the case of data loss. There are two options for this functionality:

Administrator function

Volumes used as NAS are cloned and copy source volumes are then restored. The whole operation is done independently of the operating systems and applications without affecting the server.

User function

Prompt data recovery of each file or directory is provided if a file or directory is deleted or updated by mistake. This function can be configured for individual users or for each application. This function will be supported in the next version of the S3 release.

NAS Data Protection functions require the use of ETERNUS SF Advanced-Copy Manager (ACM) software.

Link aggregation

Link aggregation is supported in order to improve fault tolerance and increase overall performance. It combines multiple ports and aggregates them together in order to provide link resiliency or higher throughput. Load balancing is supported as well.

Summary and conclusion

The ETERNUS DX product series is the perfect choice for customers looking for a flexible storage solution that adapts according to their business changes. Built on powerful performance hardware architecture with unified block and file data access, it ensures efficient storage consolidation while delivering the desired performance.

The uniform management software, ETERNUS SF, reduces the total cost of ownership, simplifies monitoring and management, and helps achieve business continuity.

Furthermore, innovative advanced software functions, particularly in the area automated Quality of Service management, allow maintaining stringent Service Level Agreements, thus making the ETERNUS DX storage systems – The Business-centric Storage – the better alternative for storage solutions.

Appendix:
ETERNUS DX Series – Hardware Specs

ETERNUS DX Online Storage Family						
	ETERNUS DX60 S2	ETERNUS DX100 S3	ETERNUS DX200 S3	ETERNUS DX500 S3	ETERNUS DX600 S3	ETERNUS DX8700 S2
Architecture	Flexible and seamless family design with uniform storage management					
Segment	Entry-level	Scalable unified entry-level and midrange systems				High-end
Maximum Storage Capacity	96 TB	576 TB	1,056 TB	2,112 TB	4,224 TB	6,144 TB
Maximum Disk Drives	24	144	264	528	1,056	1,536 (3.5 inch) 3,072 (2.5 inch)
Storage Controllers	1/2	1/2	1/2	2	2	2/4/6/8
Maximum Cache Memory	2 GB	8 + 8 GB (Block+File)	16 + 16 GB (Block+File)	64 + 32 GB (Block+File)	128 + 64 GB (Block+File)	up to 1 TB
Host Interface Type	4 Gbit/s FC 1 Gbit/s iSCSI 3 Gbit/s SAS	8/16 Gbit/s FC 10 Gbit/s FCoE 1/10 Gbit/s iSCSI 6 Gbit/s SAS 1/10 Gbit/s Ethernet		8/16 Gbit/s FC 10 Gbit/s FCoE 1/10 Gbit/s iSCSI 1/10 Gbit/s Ethernet		8/16 Gbit/s FC 10 Gbit/s FCoE 1/10 Gbit/s iSCSI
Storage Management	ETERNUS SF V16 Software Suite					
Quality of Service	Automated management of service level agreements (SLA management)					
Expandability		Extensive scalability with seamless upgrade options				
Efficiency		Thin Provisioning with zero space reclamation				
Replication			Remote Equivalent Copy (REC)			
Data Integrity	Reliability/RAID protection					
Encryption		Based on system controller and/or self-encrypting disks				
	Controller					
Redundancy	Dual controller/redundant components					