

Fujitsu ETERNUS DX Series

Note: The Fujitsu ETERNUS DX Product Analysis is no longer receiving updates as of September, 2023.

The ETERNUS DX is a family of storage systems that utilize common software and provide block storage and, with some models, file access. The current hardware models are designated by a generation number, with S5 being the latest generation for most models.

Fujitsu's ETERNUS DX storage systems span entry level, midrange, and high-end storage segments and are targeted for tier-1 applications and storage consolidation. The Fujitsu ETERNUS DX family of storage systems are designed for enterprise and distributed environments requiring high reliability, availability and serviceability features. The ETERNUS DX S5 and S4 systems are sold worldwide via Fujitsu companies and channel partners.

Products and vendors that compete with the ETERNUS DX systems include the Dell PowerMax, the VSP Hitachi VSP, HPE Alletra 9000, and IBM DS8900 systems.

Highlights

- Characteristics
 - Performance – Fujitsu publishes results for both SPC-1 and SPC-2 benchmarks. The StoragePerformance.org website should be consulted for the latest information. Historically, Fujitsu block storage systems have excellent performance and very good price performance.
 - Availability – Dual controllers provide non-disruptive HA. Multiple controller models work in a paired controller configuration.
 - Replication for BC/DR – Both asynchronous and synchronous remote replication is available to other Fujitsu DX systems
 - Support for VMware VAAI and SRM with an SRA.
 - Compression and deduplication selectable by volume
 - Multiple connection options depending on model.
 - Cache extension with SSDs.
 - Support for RAID0, 1, 5, 5+0, 6 and RAID1+0 with global hot spare disks
- Applications
 - Fujitsu's DX systems are targeted at enterprise applications either physical or virtual requiring high performance and capacity, or consolidating several smaller workloads
 - Web and client based GUI along with thin provisioning are included at no cost.
 - Additional cost software includes thin provisioning, snapshots and replication
- System environments



- The DX systems provide block storage using direct Fibre Channel (FCP), iSCSI, or FCoE with standard host drivers. File storage is supported as a NAS system with Ethernet interfaces for some models.
- Deployment and Administration
 - For customers without the required knowledge, Fujitsu and certified resellers offer rapid deployment services.
 - Several management consoles are available, including a web browser GUI and Storage Express and for multiple nodes
 - Eternus SF integrated management platform available for all models.
 - Remote Call Home support is also available and is enabled through a service processor, which may be deployed as a virtual machine.

Overview of System

The ETERNUS DX systems vary in scaling, capacity, performance, and connectivity. Heterogeneous server attachment is supported (individual server operating system support is listed in the Connectivity section of this document).

The ETERNUS DX series consists of models that address entry level, midrange, and high-end storage needs. The entry level systems consist of the DX60, DX100, and DX200. Midrange models include the DX500, DX600, and DX900. The DX8900 is Fujitsu’s high end ETERNUS system.

The entry level and most midrange models support both block and file storage personas. The high-end DX8900 and the top-level midrange model, the DX900, offer block support only.

Model	DX60 S5	DX100 S5	DX200 S5	DX500 S5	DX600 S5	DX900 S5	DX8900 S4
# of Controllers	1/2	1/2	1/2	2	2	2-4	2-24
Maximum SSD Capacity - Raw	31 TB	4,424 TB	8,110 TB	17,695 TB	32440 TB	70779 TB	141,558 TB
Maximum HDD Capacity - Raw	672 TB	2,016 TB	3,696 TB	7,786 TB	14,506 TB	31,699 TB	93,427 TB
Maximum Cache Capacity - DRAM	16 GB	32 GB	128 GB	512 Gb	768 GB	3 TB	18 TB
Extreme Cache-SSDs	N/A	1.6 TB	1.6 TB	25.6 TB	25.6 TB	51.2 TB	307.2 TB
Host Interfaces	8/16 Gbps FC, 1/10 Gbps iSCSI, 12 Gbps SAS	1/16/32 Gbps FC, 1/10 Gbps iSCSI, 12 Gbps SAS, 1/10 Gbps Ethernet	1/16/32 Gbps FC, 1/10 Gbps iSCSI, 12 Gbps SAS, 1/10 Gbps Ethernet	8/16/32 Gbps FC, 1/10 Gbps iSCSI, 1/19 Gbps Ethernet	8/16/32 Gbps FC, 1/10 Gbps iSCSI, 1/19 Gbps Ethernet	8/16/32 Gbps FC, 1/10 Gbps iSCSI	8/16/32 Gbps FC, 1/10 Gbps iSCSI

Table 1: ETERNUS DX Models



Figure 1: ETERNUS DX Systems (Source: Fujitsu)

Hardware Architecture

The Fujitsu ETERNUS DX family consists of a wide range of models to suit many different storage needs. All DX models offer a mixture of SSDs and HDDs to create a hybrid storage system. All Flash ETERNUS models are available under the ETERNUS AF series.

The DX8900 S4 is somewhat unique compared to the other models due to its scale out architecture which allows it to scale from 2 to 24 controllers representing both very large capacity and performance capabilities. The DX900 S5 has a more limited scaling capability of 2 to 4 controllers. PCIe routers are used for interconnection in scale out architectures. The remainder of the models are dual controller systems.

The DX systems use Intel Xeon Processors with multiple cores and a mixture of SSDs and HDDs may be included in the system. NVMe SSD devices can be used for caching.

Architecture diagrams can be seen for many of the models below.

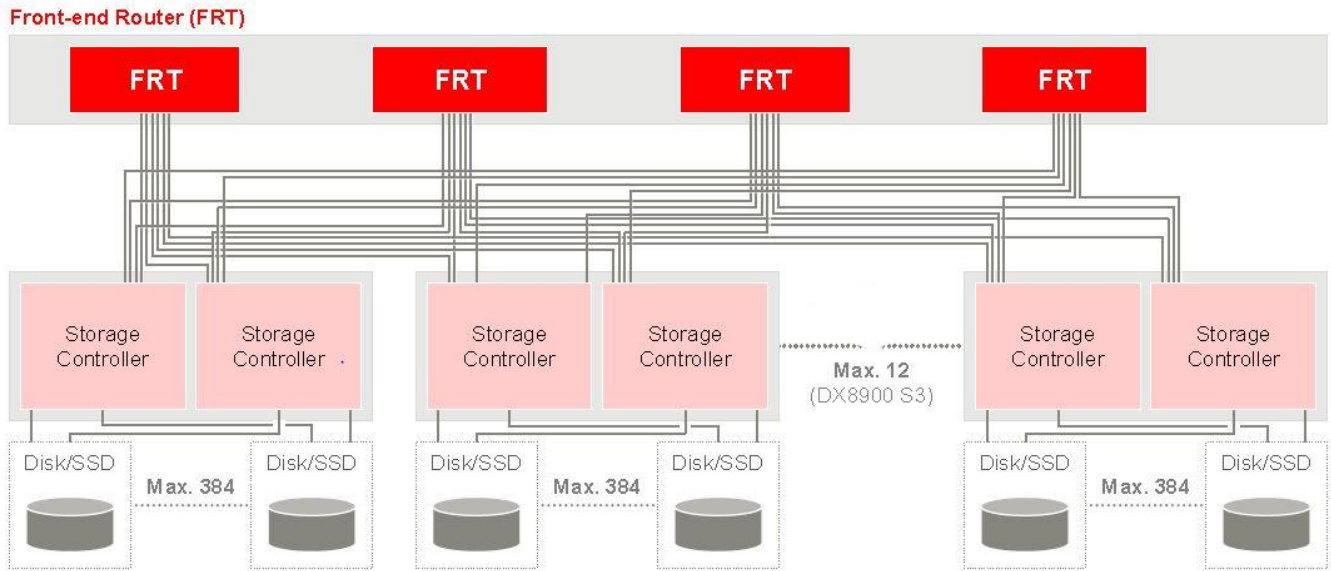


Figure 2: ETERNUS DX8900 S4 Architecture (Source: Fujitsu)

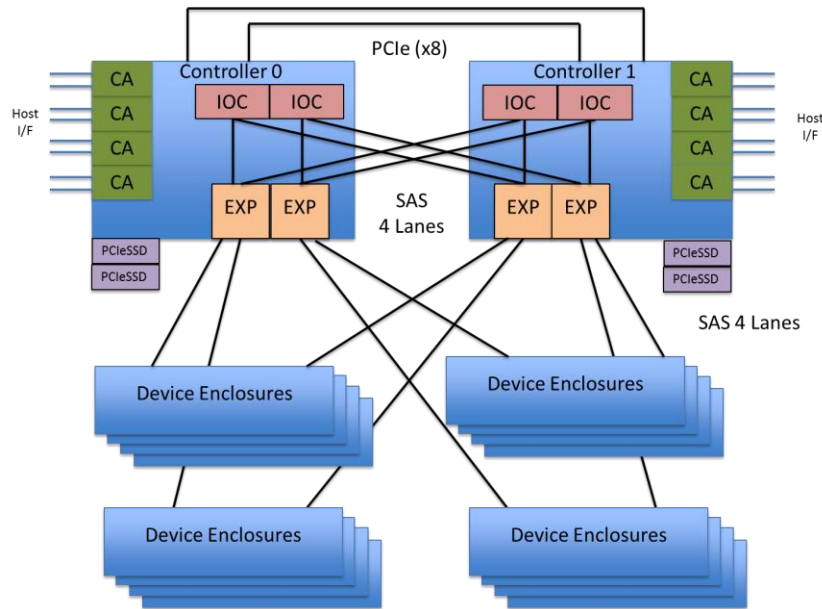


Figure 3: ETERNUS DX600 Architecture (Source : Fujitsu)

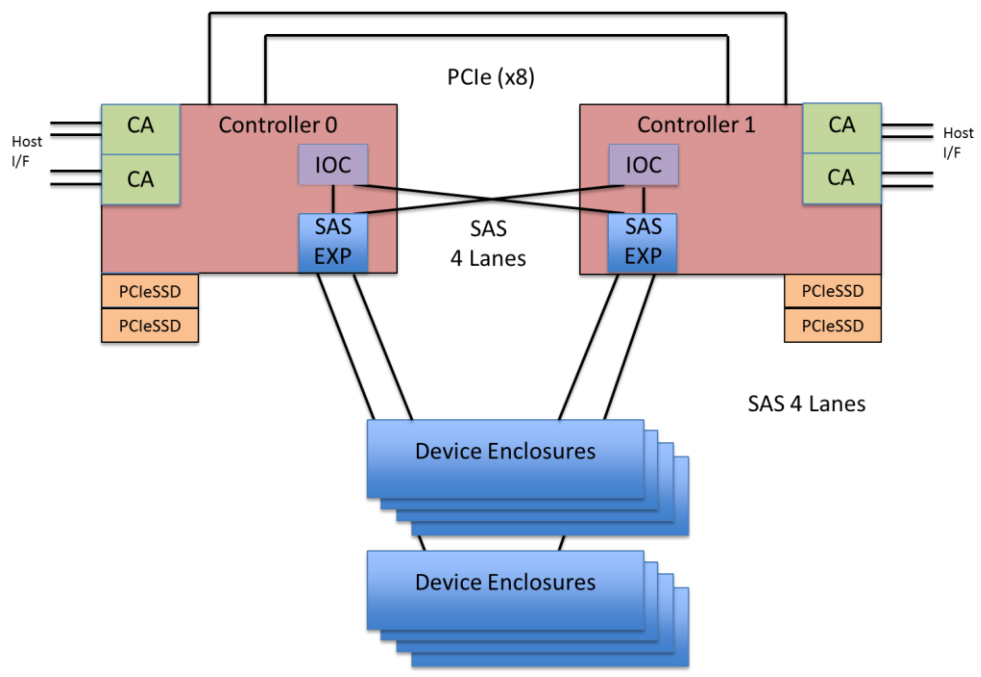


Figure 4: ETERNUS DX500 Architecture (Source: Fujitsu)

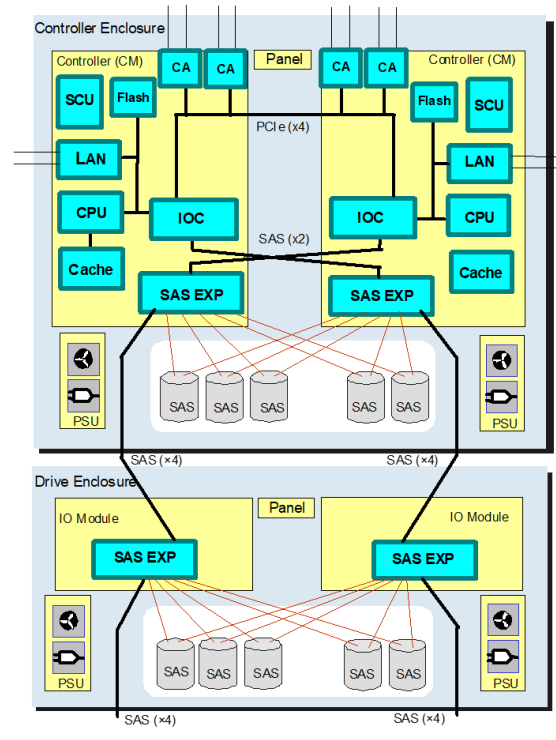


Figure 5: ETERNUS DX100 Architecture, dual controller (source: Fujitsu)

Controllers

The controller uses a PCI-Express v3 bus architecture for inter-processor communication between the controllers and for the attachment of the channel adapters (to the host) with front-end routers. The controllers execute the software that runs the storage systems. The paired controllers can run active-passive or in parallel (active-active) for faster performance.

The ETERNUS DX models use Intel multi-core processors. The cores are dynamically assigned for processing, encryption, and parity creation for RAID5 and RAID6.

Host interface connectivity for the system can be Fibre Channel, iSCSI, FCoE or a mixture.

Cache is utilized for multiple purposes, including for read and write operations, as well as maintaining information required for data protection and replication functions such as Fujitsu’s array-based replication and point in time copy data protection software. Cache is allocated dynamically. In addition to the DRAM cache, the system uses NAND flash on PCIe attached SSDs plugged into the controller PCIe bus for extended cache called Extreme Cache.

Channel Adapters

The Channel Adapters (CAs) provide Fibre Channel, iSCSI, and FCoE interfaces to the host (server) system. They process commands from the host and manage accesses to the cache. The channel adapters plug into a PCI-Express front-end routers housed in a front-end enclosure.

The channel adapters support Fibre Channel connectivity at 2, 4, 8 and 16 Gb/s. Channel adapters support 1Gbps and 10Gbps iSCSI as well as 10Gbps FCoE.

Table 2 below summarizes channel adapter connectivity across the DX series models.

Interface	DX8900	DX900	DX600	DX500	DX200	DX100	DX60
FC (16Gb/s) ports	4 – 384	4 – 64	4 – 16	4 – 8	2 – 8	2 – 8	2 – 4
FC (32Gb/s) ports	4 – 384	4 – 64	4 – 32	4 – 16	2 – 8	2 – 8	N/A
iSCSI (1GigE) ports	4 – 192	4 – 32	4 – 16	4 – 8	2 – 8	2 – 8	2 – 4
iSCSI (10GigE) ports	4 – 192	4 – 32	4 – 16	4 – 8	2 – 8	2 – 8	2 – 4
FCoE (10GigE) ports	N/A	N/A	4 – 16	4 – 8	4 – 8	4 – 8	N/A

Table 2: ETERNUS DX Models Host Connectivity

Device Interfaces

Integrated with the controller module, the Device Interface adapters (DIs) connect the disk drives to the controller modules using a 12Gb/s SAS connection. The system has four SAS connections per controller of 4 lanes each.

Device Enclosures

The Fujitsu ETERNUS DX8900 system support these enclosures:

- 2.5" device enclosure, 24 slot – 2 I/O modules
- 3.5" device enclosure, 12 slot – 2 I/O modules

The Fujitsu ETERNUS DX900 system support these enclosures:

- 2.5" device enclosure, 24 slot – 2 I/O modules
- 3.5" device enclosure, 12 slot – 2 I/O modules

The Fujitsu ETERNUS DX600 systems support three types of enclosures:

- 2.5" device enclosure, 24 slot – 2 I/O modules
- 3.5" device enclosure, 12 slot – 2 I/O modules
- 2.5" high density device enclosure, 60 slot – 2 I/O modules

The Fujitsu ETERNUS DX500 systems support three types of enclosures:

- 2.5" device enclosure, 24 slot – 2 I/O modules
- 3.5" device enclosure, 12 slot – 2 I/O modules
- 2.5" high density device enclosure, 60 slot – 2 I/O modules

The Fujitsu ETERNUS DX200 supports four types of enclosures:

- 2.5" drive enclosure, 24 slot – 1 I/O module
- 2.5" drive enclosure, 24 slot – 2 I/O modules
- 3.5" drive enclosure, 12 slot – 1 I/O module
- 3.5" drive enclosure, 12 slot – 2 I/O modules
- 2.5" drive enclosure, 60 slot – 2 I/O modules

The Fujitsu ETERNUS DX100 supports four types of enclosures:

- 2.5" drive enclosure, 24 slot – 1 I/O module
- 2.5" drive enclosure, 24 slot – 2 I/O modules
- 3.5" drive enclosure, 12 slot – 1 I/O module

- 3.5" drive enclosure, 12 slot – 2 I/O modules
- 2.5" drive enclosure, 60 slot – 2 I/O modules

RAID Groups

Fujitsu supports several RAID group configurations for both traditional and SSD disk devices. The configurations are shown below. The maximum number of RAID groups is dependent upon the number of disk drives in the system, and the configuration guidelines provided below.

The Fujitsu ETERNUS DX systems utilize RAID groups as the first unit of grouping and data protection. Host accessible volumes (or LUNs) are created by aggregating or splitting up RAID groups. Logical volumes are made available to servers as LUNs (Logical Unit Numbers).

EGI Comment: Striping and sparing are done on a per RAID group basis. That is, data may only be striped across the disks within a RAID group. This is a limitation that can impact performance due to the limited wide-striping capability.

The ETERNUS DX systems support multiple RAID levels and disk combinations providing different levels of protection and performance. The RAID levels supported include: RAID0, RAID1, RAID1+0 (data is mirrored on multiple disks and then striped across 4 to 32 disks), RAID5 with rotating distributed parity on multiple disks, RAID5+0 (which combines block striping of RAID0 with RAID5 rotating parity) and RAID6 which offers dual parity per RAID group.

Listed in the following table are the allowable and recommended RAID configurations.

RAID	Drives / RAID Group	Recommended Drives Configuration (Grouping)
RAID0	2 – 16	Not recommended due to no protection
RAID1	2	2(1D+1M)
RAID1+0	4 – 32	4(2D+2M), 6(3D+3M), 8(4D+4M), 10(5D+5M)
RAID5	3 – 16	3(2D+1P), 4(3D+1P), 5(4D+1P), 6(5D+1P)
RAID5+0	6 – 32	6[(2D+1P)x2], 8[(3D+1P)x2], 10[(4D+1P)x2], 12[(5D+1P)x2]
RAID6	5 – 16	5(3D+2P), 6(4D+2P), 7(5D+2P),

Table 3: Supported RAID Groups and Recommended Configurations

Hot Spare

The ETERNUS DX has two types of hot spare disk drives: global hot spares and dedicated hot spares. As the name implies, global spares that may be used by any RAID group with access to the hot spare and dedicated hot spare disk drives are reserved for the use of specific RAID groups. The hot spares along with the Redundant Copy feature allows rebuilding to start automatically when a disk drive fails, or before failure, if a high drive error rate is detected. After replacing a failed drive, the data is copied to the replaced drive and the original spare returned to spare status. If using different capacity drives in the system, the hot spare must be equal to or greater than the largest capacity drive.

When invoking hot sparing, a hot spare disk with the same capacity as the failed disk will be used first; if one does not exist, the largest capacity hot spare disk will be used. There is no limit to the number of hot spares that may be defined for the system.

An additional feature is the ability to use a Nearline 7.2k drive in place of a 15k SAS drive when both drives are installed together and a 15k hot spare is not available.

Evaluator Group Comment: The ability to utilize any drive as a spare provides administrators with the option to maintain high availability for a RAID group by sacrificing performance. This allows administrators the ability to choose to maintain a redundant system or maintain performance. However, sharing a spare among multiple RAID groups can lower overall system availability, as only one data rebuild operation can occur before a drive swap is required.

Disk Drives

The ETERNUS DX systems support multiple drive types, capacities and physical sizes. The disk drive types correspond to different workloads; IO centric for SSD, mixed performance capacity for enterprise 15k and 10k rpm drives, and capacity centric for the nearline 7.2k rpm drives. Both SFF (2.5”) and LFF (3.5”) form factor drives are supported on the system.

All drives utilize a Serial Attached SCSI (SAS) 12 Gb/s interface with dual ports for high availability.

Evaluator Group Comment: Fujitsu has chosen to support both 3.5" and 2.5" drives, providing significant choice for customers as the industry transitions to the newer 2.5" drive form factor. Additionally, the ETERNUS DX systems support SSD drives in both small and large form factor. Fujitsu is one of the few vendors offering SSD drives in this price class, which provides users with more choices when configuring for performance and capacity centric applications.

MAID - ECO-Mode

Fujitsu utilizes Massive Array of Idle Disks (MAID) technology for infrequently accessed data in order to reduce power and cooling requirements. This feature is referred to as 'ETERNUS ECO-Mode'. By allowing RAID groups that are infrequently accessed to spin down when they are not used, the total power consumption may be reduced.

The ECO-Mode feature controls the disk drive rotation within a RAID group, affecting all drives that form a RAID group. This feature is managed through ETERNUSmgr. ECO Mode is disabled when system activity requires access to the RAID group, including for data protection operations, mirror copying, RAID group rebuild operations and other system-required activity.

ECO Mode utilizes a setting, in conjunction with RAID group activity to determine if a RAID group should be spun down. If there is no activity, and the setting indicates the group should be spun down, then ECO Mode takes effect. If either of these two conditions is not met, the group will remain spinning. ECO Mode utilizes time frames for a specific RAID group. RAID groups may be enabled for the following time frames:

- Everyday - The mode is enabled every day from the specified start time to the specified end time.
- Every week - The mode is enabled every day from the specified start time to the specified end time.
- Specific day – The day or date of the month, along with start, and number of days is specified
- Specific week – A day of the week within the month, along with the start and ending day of the week and the starting and ending time are specified
- Last Access Time – A specific amount of time has elapsed since a RAID group was last accessed, this may be set to 10, 20, 30, 40, 50 or 60 minutes (default is 30 minutes)

Evaluator Group Comments: The lower power consumption of ECO-Mode can provide significant power saving opportunities for larger configuration. While not a unique feature in the midrange market segment, ECO-mode provides customers with flexibility to match their operational performance and cost objectives with business requirements.

LUN Management

Fujitsu refers to their LUN masking and LUN management features as 'Host Affinity.' The ETERNUS GUI configuration software controls LUN access, or Host Affinity, which is included with the system. Host Affinity controls access, which provides a security mechanism at the LUN level. A Host Affinity group is defined as a set of LUNs that are made available to a specific set of servers.

LUNs can be assigned to multiple servers, with one LUN used in different Host Affinity groups allowing for multi-host data access for cluster or other configurations using host software to coordinate data sharing.

The maximum number of LUNs per RAID group is 128. A LUN may be comprised of parts of one or more RAID groups, or a partial RAID group.

RAID Migration

Fujitsu refers to dynamic migration of a LUN to a different RAID group as RAID Migration. The capacity of the LUN can be extended during the migration. The target RAID group may be the same or different RAID level. This feature provides the ability to manage changing requirements for capacity, performance or reliability of a specific volume.

Virtual / Thin Provisioning

Fujitsu has added support for virtual capacity of volumes, which is known as “Thin Provisioning.” Additionally, Fujitsu supports expansion of a LUN through RAID Migration.

Evaluator Group Comments: Thin provisioning has become an increasingly important checklist item feature. The effectiveness of thin provisioning in general is application dependent, with some uses such as virtual server and VDI implementations providing some of the best uses.

Logical Device Expansion

Logical Device Expansion (LDE) supports changing the RAID level, or adding capacity to an existing RAID group. The capacity of an existing RAID group can be dynamically extended by adding extra disk drives. A RAID group’s RAID level can be dynamically changed. The capacity of the new RAID group must be at least as large as the existing RAID group.

LUN Concatenation

LUN concatenation allows several volumes to be aggregated together to provide a larger volume. Capacity for the concatenated LUN may be taken from unused areas in an existing RAID group, or the unmapped capacity in another RAID group. LUN concatenation creates a new volume with the obtained area, and concatenates the original volume and the new volumes together into a larger capacity volume.

LUN concatenation may be used both for expanding an existing volume and to combine multiple areas into a larger LUN.

Evaluator Group Comments: The requirement to concatenate LUNS to expand their capacity is an architectural limitation of the way RAID groups are utilized to enable storage pooling. Although this provides LUN expansion, data is not re-striped across the additional capacity.

Data Encryption

The ETERNUS DX systems support encryption of data on the devices.

The DX systems support the use of self-encrypting drives (SED), for drive-based encryption. With SED based encryption, AES-256 bit encryption is utilized. This method has no impact on controller performance, although it does limit the drive choices available.

Additionally, the ETERNUS supports encryption during replication using standard IPsec security (AES encryption with HMAC-MD5 authentication). This feature is enabled through the use of the optional iSCSI Remote Adapter for remote replication.

Evaluator Group Comments: Encryption of data is important for some environments. The encryption may support regulatory compliance when coupled with data preservation required for specific regulations. Fujitsu's encryption with included key management provides a simple and effective method of protecting data.

Disk Traffic Control (DTC)

DTC is a performance optimization that optimizes the back end disk drive by segmenting sequential and random accesses into separate extents when they both occur simultaneously. The two I/O streams (sequential and random) are processed alternately on a time interval basis to provide optimum usage of the disk drive resource.

Other Features

Fujitsu also supports two modes for specific environments:

- Oracle HARD
- Video Streaming Mode

The Oracle HARD feature provides end-to-end data fidelity. This prevents potential corrupt data from being written to the disk.

The Video streaming mode allows forced prefetching of data when sequential read or write operations are the predominant workload. This can improve performance for video applications.

Software Architecture

Block storage access is supported with the ETERNUS DX system.

The following table lists supported servers and operating systems:

Host Vendors	Supported Operating Systems
Fujitsu	Windows Server 2003 for Itanium-based Systems, Windows Server 2008 for Itanium-based Systems, Red Hat Enterprise Linux AS (V.4 for Itanium), Red Hat Enterprise Linux(V.5 for Itanium)
	Windows Server 2003, 2008, 2012, Red Hat Enterprise Linux 5, Red Hat Enterprise Linux 6, VMware vSphere 4, VMware vSphere 4.1
	Solaris 10 Operating System
	Solaris 8, 9, 10, 11 and beyond Operating System
	Windows 2003, 2012 x64 Editions, Red Hat Enterprise Linux AS v.3, Red Hat Enterprise Linux AS v.4, RedHat Enterprise Linux ES v.3, Red Hat Enterprise Linux ES v.4, Red Hat Enterprise Linux 5 (for x86), Red Hat Enterprise Linux 5 (for Intel64), SUSE Linux Enterprise Server 9 for x86, SUSE Linux Enterprise Server 10 for EM64T, SuSE Linux Enterprise Server 11, x86, SUSE Linux Enterprise Server 11 for EM64T, VMware vSphere 4, VMware vSphere 4.1
Oracle / Sun	Solaris Operating System
HP	HP-UX 11iv2, HP-UX 11iv3
IBM	AIX 5.2, AIX 5.3 and later
Others	Windows Server, Red Hat Enterprise Linux AS v.3, Red Hat Enterprise Linux AS v.4, RedHat Enterprise Linux ES v.3, Red Hat Enterprise Linux ES v.4, Red Hat Enterprise Linux 5 (for x86), Red Hat Enterprise Linux 5 (for Intel64), SUSE Linux Enterprise Server 9 for x86, SUSE Linux Enterprise Server 9 for EM64T, SuSE Linux Enterprise Server 10, x86, SUSE Linux Enterprise Server 10 for EM64T, VMware ESX/ESXi

Table 4: ETERNUS DX8900 Supported Servers and Operating Systems

Multi-Path IO

Fujitsu provides a multi pathing driver, referred to generically as MPD for the operating platforms supported by the ETERNUS DX line. For each platform there is a specific driver package known as ETERNUSmpd.

The ETERNUS DX supports either a fixed path access on a per LUN basis, or dynamic multi-path access on a per LUN basis. This feature is different than active/active controllers, and different still from multi-path I/O. The ETERNUS DX series permits asymmetric LUN access. This is a mode Fujitsu refers to as "Assigned CM."

The ETERNUS DX systems assign a controller module (CM) to a specific RA ID Group. The path(s) connecting to the controller normally used will have an active status, while the path(s) connecting to other controllers will have standby status.

Paths connected to the assigned CM are in active status, while paths connected to the non-assigned CM are in standby status. Load-balancing is performed on the paths in active status (called active paths). Paths in standby status (called standby paths) are normally not used for access.

Thus, the ETERNUS DX provides active/active controller paired use with multi-pathing, and asymmetric access to a LUN through a specific array controller (or control module - CM). The diagram below illustrates the configuration for AIX, although this is also the case for other OS environments. The figure below illustrates asymmetric access. Scale out is support with multiple paired controllers connected through a front-end router

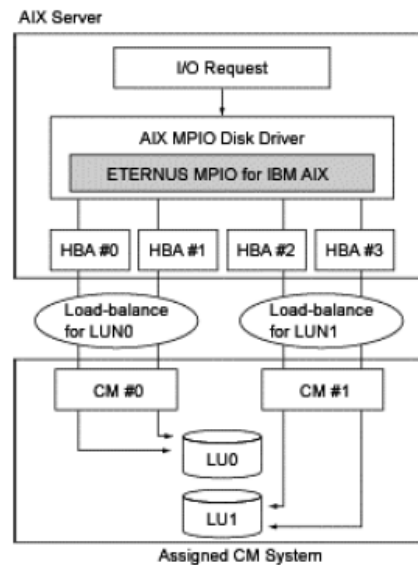


Figure 6: Asymmetric MPIO Access for ETERNUS DX

Solaris

Fujitsu provides a device driver for Solaris platforms supporting MPIO, utilizing Fujitsu’s multi-path driver or MPD.

Windows

Fujitsu provides a device driver for Windows platforms supporting MPIO, utilizing Fujitsu’s multi-path driver or MPD.

Linux

Fujitsu provides a device driver for Linux platforms supporting MPIO, utilizing Fujitsu’s multi-path driver or Linux DM driver.

IBM AIX

For the AIX platform, Fujitsu’s MPD is an interface to AIX’s native multi-pathing I/O capabilities. The ETERNUS MPIO for AIX is an IBM Path Control Module (PCM) that utilizes AIX’s native multi pathing capabilities, through the MPIO interface. ETERNUS systems utilizing this software may be accessed as an MPIO device by AIX. This allows active/active asymmetric logical unit access via AIX.

ETERNUS MPIO for IBM AIX performs load-balancing over the usable paths to improve the access performance.

iSCSI Connectivity

The ETERNUS DX systems support iSCSI connectivity. There are several options for iSCSI, which may be implemented. The following are supported by the ETERNUS systems:

- Jumbo frames are supported
- iSNS (internet Storage Name Services)
- Changing TCP/IP settings
- CHAP authentication
- Host affinity mode
- Cancel reservation
- Header Digest
- Data Digest

Specification	Supported Number
iSCSI Hosts	8,192
Affinity Groups	8,192
Affinity Groups per Port	256

Table 5: ETERNUS DX8900 iSCSI Host Limitations

Software Features and Functions

The ETERNUS SF software suite includes a number of offerings that provide options for system and SAN management, data protection and disaster recovery, backup and archiving, local replication and path failover.

Data Protection

Fujitsu has termed their point in time copy feature for the ETERNUS series, One Point Copy (OPC). There is local mirroring (Equivalent Copy), along with three different versions of One Point Copy:

- One Point Copy (OPC)
- Quick OPC
- Snap OPC

	Equivalent Copy	One Point Copy (OPC)		
		One Point Copy	QuickOPC	SnapOPC
Copy Type	Mirror	Clone	Clone	PIT
Architecture	Full Copy	Bit Map	Bit Map	Bit Map
Max # Copies	512/volume Maximum concurrent EC, OPC, QOPC or SnapOPC copy sessions: 8,192			
Re-sync	Yes	No	Yes	No
Potential Data Exposure	Synchronous—none	Exposure is back to time OPC was activated	Exposure is back to time QuickOPC was activated	Exposure is back to time when snap was taken
Availability for Restore	Mirror available after completion of initial mirror	Available for restore immediately, or able to replace volume after copy complete	Available for restore, or able to replace volume immediately	Available for restore only after command completes
DR, Fail-over Capability	Yes On Mirror suspend or break	Yes Able to replace volume after copy complete	Yes Able to replace volume as soon as command completes	No Not for DR or fail-over, depends on source data
Default Action	Mirror. Copy is a true mirror until it is split	Copy when command issued	Copy when command issued	Copy on Write
Additional Capacity	100%	100%	100%	Changes only
Access to Copy	R/W after split	R/W immediate	R/W immediate	R/O
Performance Impact	Some impact	Small Impact	Minimal	Minimal

Table 6: Fujitsu ETERNUS DX8900 Data Protection Options

Note: The maximum number of copy sessions is based on cache available for management of the copy sessions or a maximum of 8,192. Cache usage is determined by the size of the volumes being replicated, the number of volumes being replicated and the resolution of the bitmap. The maximum total copy sessions or 512 copy sessions per volume are not available for all configurations.

One Point Copy (OPC)

The copy is a “clone” copy where an exact original copy is made to another physical area. Entire volumes are copied using OPC.

OPC has two copy phases; the logical phase and the physical phase. The logical phase starts after the OPC command is initiated. A bitmap that represents the copy is created in cache memory. That process takes a few seconds and after completion, the copy is available for use. The physical phase then starts, which is the process of the storage system copying the data to the copy location. The physical copy is completed in the background. During the physical phase, host access to data is available. The copied volume is available almost instantly after OPC initiates, upon completion of the logical phase. This is accomplished as follows:

- Read to Target – If the requested block has been previously copied to the target, a normal read is done to the target volume. If the requested block has not been copied, the system will copy the block from the source to the target and then honor the read from the target volume.
- Write to Source – If the requested block has been copied to the target, a normal write operation is done to the source. If the requested block has not been copied yet, the original block is copied to the target and the write is then done to the source.
- Write to Target – If the requested block has been copied to the target, a normal write operation is done. If the block has not been copied, the bitmap is changed to indicate the copy has been done and the normal write is done.

The process of the physical copying persists through a power down or reboot.

There are three speed settings (pacing parameters) to control the amount of storage system resources consumed for the copy operation. The speed or priority settings are controlled via the management software, ETERNUSmgr. The three speed options are **low**, **high** and **automatic**. The automatic setting changes the priority dynamically in response to the system workload.

One Point Copy consumes cache memory for the bitmap and the additional data being copied.

QuickOPC

The first copy is the same as One Point Copy, except a bitmap is used to track all changed blocks from the time the logical copy completed. The design leverages the use of a bitmap, which tracks changed blocks for the volume. Subsequent copies are optimized by only copying the blocks changed since the previous copy.

SnapOPC

SnapOPC operates similar to other Copy on Write point in time copy operations. Rather than copying the new data to the copy volume, the old data is copied to the duplicate volume. Space is saved by only saving the original blocks of data, when new data is written to the SnapOPC volume. By copying only a portion of the source volume, it minimizes copy volume capacity.

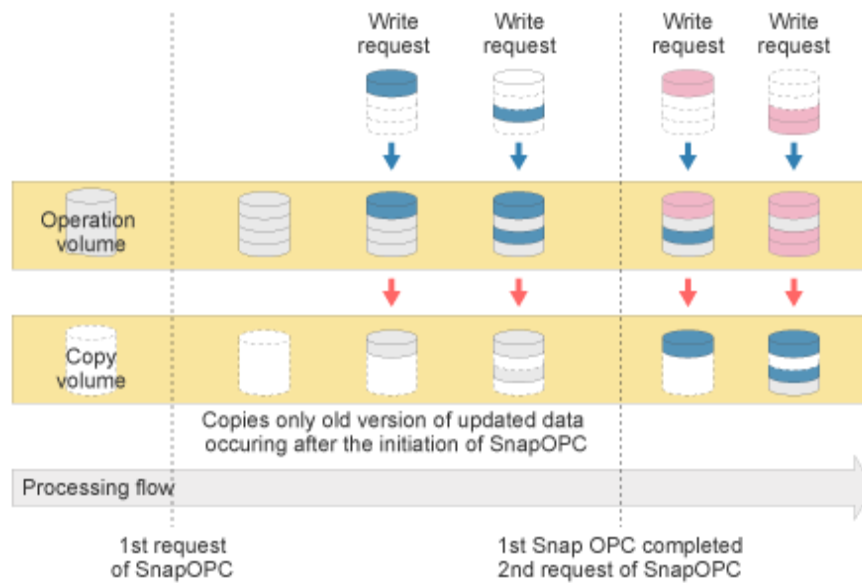


Figure 7: Example of SnapOPC for ETERNUS S4

SnapOPC+

Similar to SnapOPC, the SnapOPC+ feature backs up the previous data areas to a copy destination area. As a result, multiple generations of data may be preserved.

Equivalent Copy (EC)

Fujitsu local mirroring and replication product is known as Equivalent Copy (EC). The mirror is an exact image copy made to another physical area. Entire volumes are copied as initiated by the Fujitsu management software.

The mirror volume must be a local volume, thus the ETERNUS DX supports local, in-system mirroring. Like the One Point Copy feature, space for the copy must be reserved previously and be at least the same size. The RAID configuration may be different for the source and target.

The copied volume is not available until the initial copy is complete and either a “detach” (split or break) or a “suspend” command has been issued. During suspend, updates to source and target are tracked. A “resume” command causes resynchronization. Resynchronization copies the incremental updates from

the source and overwrites any changed made to the target. The target volume is not available during a resume. If a “detach” has been done, the target volume becomes independent at that time.

The mirror persists through a power down or reboot. As in One Point Copy, Equivalent Copy consumes cache memory for the bitmap and the additional data being copied (and the additional data being copied). To use EC, at least 1GB of cache for each controller module must be installed. While very similar to One Point Copy, the differences are that Equivalent Copy supports the Suspend/Resume function and the mirror volume is not available until the copy is complete.

Remote Equivalent Copy

Fujitsu’s remote replication for the ETERNUS DX is known as Remote Equivalent Copy (REC), which supports replication between multiple ETERNUS or other compatible Fujitsu DX storage systems. When referencing remote mirroring specifically, this feature is also known as Remote Equivalent Copy (REC). Fujitsu’s RAC product supports both synchronous and asynchronous remote replication.

Remote replication is supported over both FC and IP connections, utilizing the FCP or iSCSI protocols respectively. Utilizing iSCSI replication provides a cost effective method that utilizes TCP/IP connections without requiring an FCIP converter as is the typically the case when using FC protocols.

Fujitsu ETERNUS Remote Equivalent Copy provides the following features:

- Maintain an exact byte-for-byte copy of production data in a remote location
- Write order consistency is supported for both Sync and Asynchronous replication
- Consistency groups are supported
- Failover to secondary sites for disaster recovery
- Integration with OPC for secondary site backup operations
- Management of replication pairs using ETERNUS Advanced Copy Management
- Choose between synchronous and asynchronous mirroring for ETERNUS systems

The figure below illustrates the basic example of an asynchronous replication scenario between two ETERNUS DX storage subsystems.

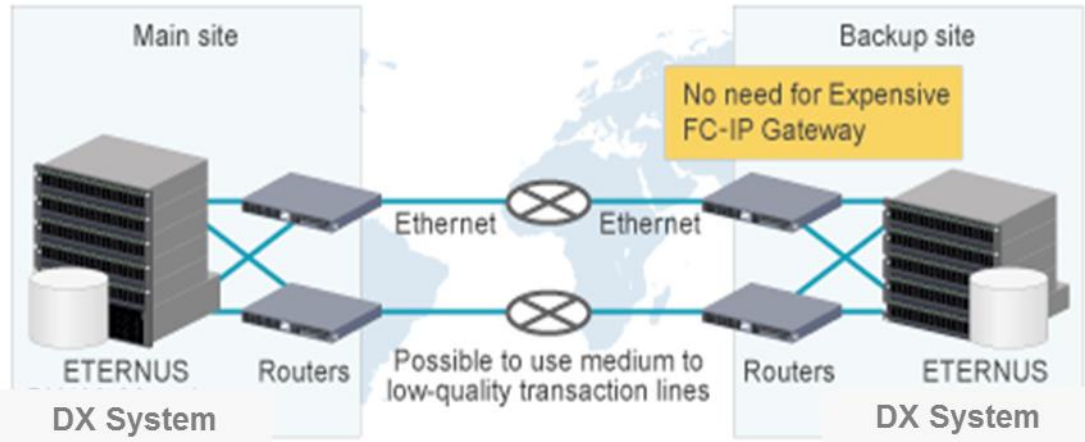


Figure 8: Basic Example of Asynchronous Replication between Two Systems

iSCSI Remote Adapter (optional hardware)

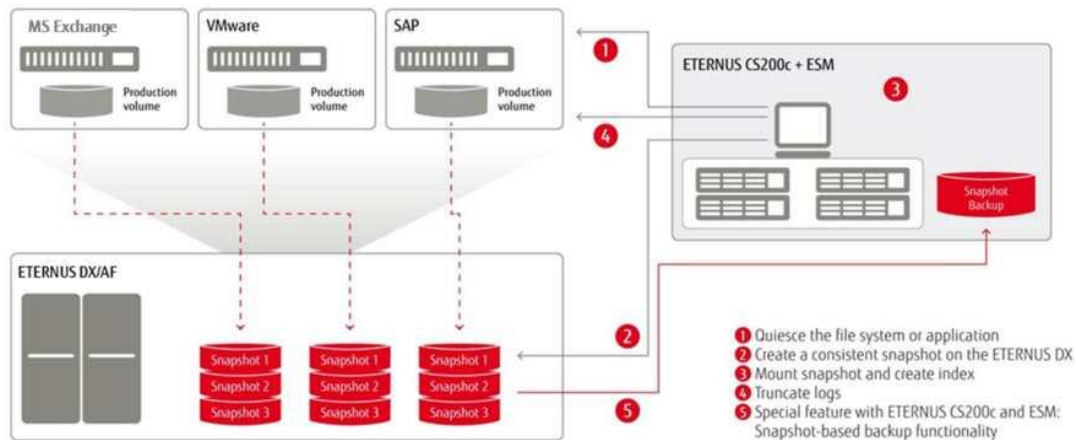
The iSCSI Remote Adapter (iSCSI-RA) provides a number of features for the ETERNUS DX systems.

- Optimized TCP acceleration
 1. Provides optimization for all TCP protocols (FTP, iSCSI, etc.)
 2. High bandwidth using large TCP window size
 1. Supports up to 300 Mbps effective TCP transfer rate per card
 2. Supports 500 ms RTT TCP connections
- Encrypted TCP Communication
 3. Utilizes AES encryption with HMAC-MD5 authentication
 4. Supports standard IPsec security
- Bandwidth Control
 5. Traffic shaping and QoS
 6. Automatic TCP window size adjustment with 50 micro-second precision
- Embedded WAN acceleration
 7. All functions supported on one chip on an AMC (Advanced Mezzanine Card)

Evaluator Group Comments: The Fujitsu iSCSI-RA cards may provide a significant advantage for customers looking to deploy a remote replication setup without requiring external FCIP or other WAN access hardware. The use of iSCSI leverages TCP/IP, removing the need for FCIP converters commonly found with FC based replication. Additionally, the encryption provides a level of security mandated by many companies either for security requirements or best practices guidelines. Finally, the inclusion a TCP accelerator can eliminate the need for a separate WAN accelerator in many cases. As a result, it is possible for customers to lower their network costs using the Fujitsu iSCSI-RA cards rather than external encryption and WAN accelerators.

ETERNUS Snapshot Manager

ETERNUS support creating and cataloging snapshots through the ETERNUS Snapshot Manager (ESM) GUI with the ability to replicate snapshots to another ETERNUS system for data protection. ESM is integrated with application software and VM systems for application-aware snapshots. Granular snapshots and recovery enable near real-time data protection. ESM is separately licensed per array. CommVault IntelliSnap technology is the basis for the ESM implementation.



How ESM works

Figure 9: ETERNUS Snapshot Manager (source: Fujitsu)

NAS Support

The ETERNUS DX systems, with the exception of the DX8900 and DX900, can be configured to support block, file, or both types of access simultaneously. When NAS functionality is enabled, remote protocols of NFS and CIFS/SMB allow access to file systems over the Ethernet interfaces.

File systems are configured as NAS volumes where a volume is allocated as a thin provisioned volume and assigned specifically for NAS usage. The number of filesystems supported varies by the DX model but a maximum of 128TB of capacity may be allocated.

Model	Maximum Number of Filesystems
ETERNUS DX100	1
ETERNUS DX200	2
ETERNUS DX500	4
ETERNUS DX600	8

Table 1: Maximum Filesystems by Model

Block sizes for the file system when the volume is allocated may be set based on the prominent application usage. File systems with different block sizes may be configured in the DX NAS system at the same time. The block size does have an effect on the maximum file system capacity.

Block Size	Maximum File Size	Maximum Volume Capacity
8KB	1TB	4TB
32KB	4TB	16TB
256KB	32TB	128TB

Table 2: Block Size Settings

Protocol Support

CIFS/SMB

- Version 3.0 through SAMBA 4.1
- SMB Encryption of data transferred
- Large MTU – from 64KB to 1MB
- Symbolic links supported
- Active Directory integration with support for Windows ACLs

NFS

- V2 and V3 supported
- LDAP support with NFS ACLs

File Sharing

When files are shared between NFS and CIFS/SMB, the NFS ACLs are utilized with a mapping from Windows to NFS.

Snapshots

A snapshot of an entire filesystem may be taken. The block-based snapshot mechanism is used to take the snapshot of the volume used for the filesystem.

Quotas

Quotas are supported by the ETERNUS NAS system.

NAS Backup Volumes

NAS volumes can be backed up to other volumes that are specifically configured in the DX system as NAS Backup Volumes. Restores may be done of entire volumes or individual files after the backup volume has been mounted.

Reliability, Availability, Serviceability Features

ETERNUS storage systems maintain data integrity and provide data availability via the following mechanisms:

- RAID protection of disk devices
- Redundancy of active components
- Dual controllers
- Drive enclosures with dual I/O modules (optional)
- Dual power sources
- Capacitor protected cache with flash memory termed Cache Protector
- Controller based data integrity checking with Fujitsu Data Block Guard
- Activities supported with multiple controllers active, non-degraded
 - Configuration changes
 - Adding, removing or changing hot spare disk drives
 - Changes supported while the system is active with a single controller:
 - Maintenance on controllers
 - Firmware upgrades
- Data Encryption

Device Connectivity

The connectivity to from the controllers to the devices is through SAS connections. Each enclosure connects to another via a 4X SAS connector, and is then connected to a SAS expander, with a direct connection to each drive. Similar to switched FC connectivity, there is a direct connection to each device, thereby removing any impact of a faulty connection or device on the overall system.

Data Block Guard

Fujitsu has implemented integrity checking of data within the controller via a mechanism referred to as Data Block Guard. This feature is also included on the ETERNUS DX600 systems. Data Block Guard adds an extra eight bytes of information on every 512-byte data block as it enters the controller. These eight bytes are retained with the data stored within the system. The eight bytes are composed of a four-byte cyclical redundancy check (CRC) and four bytes for the block ID. This information is checked whenever the block is stored on disk and is also checked when data is retrieved by the array controller. The Data Block Guard information is stripped whenever data is sent to the host by the controller (also known as Channel Adapter by Fujitsu).

Error Checking and Correction (ECC)

Data integrity is also a part of the cache function. Data in cache is error-correction code (ECC) protected to be able to detect and correct double bit, or any odd number of bits of cache memory errors.

Redundant Copy

All disks are monitored for hard and soft errors. When a disk is identified for preventive replacement (before hard failure), the redundant copy function builds the identified disk's data from the other disks in the RAID group to a hot spare. After the drive rebuild is complete, the drive identified is removed from the RAID group to ensure integrity of the RAID group.

Mirrored Cache

The controllers each contain cache. Only write data is mirrored from one controller to another controller. With a varying amount of write data, the amount of cache available between the controllers is also variable.

The write data is mirrored between the two controllers over a shared PCI-Express bus link. The bandwidth of the cache on a controller is not specified by Fujitsu.

Cache Protector

The ETERNUS DX systems utilize a capacitor known as the System Capacitor Unit (SCU) in case of a power failure to provide power while saving write data in the controller's cache memory. The SCU and flash memory configuration provides protection of cache data for an unlimited amount of time.

Evaluator Group Comment: SCU and Flash memory feature provides a unique combination that delivers unlimited protection time for power loss scenarios. The amount of time data is protected is even more important in entry-level systems due to the potential amount of time a system may be down before the situation is recognized.

Redundant Power System

The power system is based on redundant power supplies in the modular enclosures with dual line cord inputs. All power system components are replaceable non-disruptively.

RAID6 and RAID6-FR

With the use of dual parity drives in a RAID6 configuration, the RAID group provides better protection and reduced overhead than many alternative RAID configurations. A RAID6 group provides protection against a double disk failure, while minimizing the overhead associated with a RAID50 or similar configuration. RAID6-FR is an implementation that does a fast rebuild of a failed device to minimize the window for a second failure to occur.

Evaluator Group Comment: Protection against double disk failures is an important consideration with the new generation of high capacity disk drives. As disk capacities grow, the rebuild time also grows, thereby increasing the possibility of a double disk failure prior to data reconstruction. This issue is particularly important for high capacity drives over 600 GB. Due to the lengthy rebuild times, protection against double disk failures is an important consideration for all disk drive types.

Firmware Updates

Firmware updates may be accomplished non-disruptively on the ETERNUS DX storage systems via the common approach of failing over LUNs to the secondary controller, updating firmware, then failing LUNs to the other controller and updating firmware on that controller.

Concurrent Maintenance

A modular architecture is based on modules being able to be taken out of service while others take over the workload. The ETERNUS DX storage systems operate in a mode where some components can be replaced without failing over to another controller while others require a failover to occur for service. Fujitsu states the following components are FRUs (Field replaceable units) and may be replaced while providing data access.

- Array Controller, contains - Controllers, cache and channel adapters
- Power Supplies with integrated cooling fans
- Disk drives

Remote Call Home

The ETERNUS DX system supports remote call home with several remote support features including:

- Send Communication Log
- Automatic sending of system Log
- Manual system log sending

Evaluator Group Comment: The Remote Call Home feature is important to businesses of all sizes. Having remote call home allows small IT shops with limited expertise the ability to deliver high availability with the assistance of Fujitsu support.

Management Session Encryption

ETERNUS DX supports the use of SSL/SSH, protocols for encrypting data and sending encrypted data on the network. This prevents malicious accesses with unauthorized access via Web browsers or CLIs.

Storage Management

SMI-S Interface

The SNIA along with the DTMF have created a set of storage management interface specifications, known as “SMI-S” interface.

Evaluator Group Comment: Prior models of the ETERNUS DX systems have supported and been certified interoperable with SMI-S version 1.2 frameworks. Check with the vendor to confirm the status of current models.

ETERNUSmgr (Included)

ETERNUSmgr is a device management application that is included with ETERNUS storage systems, including ETERNUS DX8900 systems. The backend program is embedded in the ETERNUS systems and is accessed by a web browser, the ETERNUSmgr frontend software or other operation management software. This browser based management application is used to configure, manage and monitor (including email notification) one or more ETERNUS systems.

ETERNUSmgr Aid (Optional)

ETERNUSmgr Aid adds management value to ETERNUSmgr with its simplified GUI interface and additional management features. Customers will find it easier to manage their ETERNUS storage systems. ETERNUSmgr Aid is optional software that is installed on an administrative server.

ETERNUS SF Storage Cruiser (Optional)

Storage Cruiser is a SAN management application that provides configuration, specification, and monitoring of a SAN environment. It may also be used to manage the relationship between the storage as viewed from the host server and the physical storage, allowing the whole storage system's configuration and failure conditions to be managed. It will manage the ETERNUS family and also supports other selected vendors' products.

ETERNUS SF Storage Cruiser consists of two modules, with an optional third module:

- The Agent feature is installed on each server node being managed
- The Manager feature is installed on the selected administrative server
- (Optional) A remote Client may be installed on any computer used to perform management functions

ETERNUS SF Advanced Copy Manager (Optional)

The ETERNUS SF Advanced Copy Manager provides backup and replication with flexible options and works in conjunction with the optional Advance Copy functions of ETERNUS systems. It allows non-stop high-speed backup of major databases such as Oracle and SQL Server. The methods used are Equivalent Copy (EC), One-Point Copy (OPC), differential backup/replication (QuickOPC) and SnapOPC.

Performance

Performance and price performance levels of the ETERNUS DX family have historically been excellent as measured by the SPC benchmarks. The transaction centric SPC-1 indicates performance on I/O latency sensitive applications and the SPC-2 benchmark indicates performance for applications that drive large amounts of data. Fujitsu claims 10 million IOPs for the DX8900.

Fujitsu regularly conducts SPC benchmark tests and the website www.storageperformance.org should be consulted for the latest information.

Evaluator Group Comments

The ETERNUS DX is comprised of entry level, midrange, and high end systems with enterprise features, large performance, and high capacity. The many offerings within the ETERNUS DX family give plenty of options in terms of performance, capacity, and scalability. The ability to scale very large with a scale-out implementation in the high end DX8900 system is unique in the industry at this time. Performance and capacity scaling allow Fujitsu to offer customers consolidation systems that can grow to meet demand – both performance and capacity.

The systems utilize common architectural elements across the DX series. These systems share a common set of management, data protection and replication software.

Fujitsu has continued to enhance its platform and server based software features and interoperability. The high-end ETERNUS DX8900 systems are a good choice for enterprises wanting to consolidate to a central system or for deploying storage for several applications. The systems have a wide variety of data protection options and an excellent warranty.

Strengths:

The ETERNUS DX systems are a proven design with the features and characteristics expected of enterprise storage systems. The proven reliability of earlier generations is continued with the new DX models. The system is a low-risk selection for customers.

The scale-out capabilities for high end models are extremely valuable for IT environments, allowing customers to increase capacity and performance as needed without reducing the performance as more capacity is added.

The use of PCIe attached SSDs for extended cache and standard SSDs as high performance devices is a unique combination and should provide performance advantages compared to other systems.

The host level connectivity is very broad with the top-level data rates and the widest selection compared to other mid-range systems. The broad selection of interfaces allows the DX systems to serve many different market usages.

Perceived Challenges:

While having very powerful systems, Fujitsu is not well distributed in North America. Such great systems should be made more widely available. Europe and Asia are the strength markets for Fujitsu with the ETERNUS systems but there is opportunity in North America as well.

The storage pooling used in the ETERNUS DX S4 systems is based on the earlier S2 architectures with a mapping to the physical RAID protection across devices. The use of protection at the chunk level in the storage pool would give the system greater flexibility in rebuilds, expansion, and sparing.

Additionally, while the scale out functionality is a valuable feature present at the high end, the midrange and entry level models are lacking this functionality, offering a less scalable solution.



More detailed information is available at <http://evaluatorgroup.com>

Copyright 2023 Evaluator Group, Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying and recording, or stored in a database or retrieval system for any purpose without the express written consent of Evaluator Group Inc. The information contained in this document is subject to change without notice. Evaluator Group assumes no responsibility for errors or omissions. Evaluator Group makes no expressed or implied warranties in this document relating to the use or operation of the products described herein. In no event shall Evaluator Group be liable for any indirect, special, inconsequential or incidental damages arising out of or associated with any aspect of this publication, even if advised of the possibility of such damages. The Evaluator Series is a trademark of Evaluator Group, Inc. All other trademarks are the property of their respective companies.