

Fujitsu Storage ETERNUS DX8900 S4 – hardware architecture

Content

Management summary and general remarks	2
System architecture	3
Controller Module (CM)	5
Controller Enclosure (CE)	7
Front-end Router (FRT)	9
Front-end Enclosure (FE)	9
Drive Enclosure (DE)	10
System firmware	12
Remote service	13
Environmental sustainability	13
Summary	14
Appendix – technical specifications	14
Appendix – online resources	14

Management summary and general remarks

Based on a unique consistent system design, ETERNUS DX is a seamless family of hybrid disk storage systems ranging from economy, entry-level, and midrange systems up to the enterprise class. The major design principles of ETERNUS DX are compatibility and modularity. All models utilize the same disk types, disk shelves, racks, cables and other components, and use a uniform management software. Principally only the Controller Modules are model specific. The idea is to minimize operational, training and migration efforts and costs. Furthermore, this family concept allows upgrading a smaller model to a larger one by exchanging the control module. All other components can be reused. The uniform management system supports easy interaction between different model sizes: For example, ETERNUS DX100 S4 entry-level and DX500 S4 midrange models in remote locations can replicate data to any generation of ETERNUS DX and ETERNUS AF models. The comprehensive system configuration flexibility of ETERNUS DX makes it easy to balance user requirements in terms of speed, capacity and costs.

The latest member of the family, the ETERNUS DX8900 S4 enterprise hybrid storage system, is purpose-built for the requirements of behemoth data centers and delivers outstanding peak performance of up to 10 million* Input/Output Operations per Second (IOPS) and a enormous capacity of up to 140 petabytes (PB)** when used in full flash configuration, paired with enterprise-class resiliency and always-on availability. The systems incorporate a complete redesign of the system to achieve maximum modularity and configuration flexibility, while delivering highly scalable performance and operational robustness. This means that enterprises no longer need to precisely anticipate future capacity demand levels, nor select a specific high-end model with a fixed data-capacity range. The traditional, fixed-capacity, model-based approach typically led to a situation in which IT organizations invested in modules larger than they actually needed, just to be sure that future capacity could be covered and to avoid unplanned high migration costs. That's why Fujitsu based the ETERNUS DX8900 S4 on a completely new modular architecture that can be flexibly expanded according to business needs. With minimum initial cost, customers can start with a system having two controller modules. Maximum performance with up to 384 front-end ports for host connectivity can be achieved easily by scaling up to twenty-four controller modules in ETERNUS DX8900 S4. In order to boost performance, the system supports secondary cache accelerated by NVMe. Eight dual-port NVMe (PCIe) SSDs can be used in each CE (up to 96 for the full system) as secondary cache. Simultaneous and flexible connection to different network types is possible. FC and iSCSI host interfaces can be deployed in mixed configurations.

Fujitsu's Business-centric concept enables a higher level of automation for storage tiering, including support for SAS, Nearline SAS and SSD drive options and capacities that ensure an optimum service level in dynamic and fast-changing environments while also reducing the overall total cost of storage. The system makes use of advanced data reduction technology across storage tiers by offloading the compression process into the dedicated Storage Acceleration Engine (SAE), leading to improved capacity optimization with minimal performance impact.

The system architecture of the enterprise systems can cope with multiple component failures without any interruption to service. In conjunction with the unique transparent failover feature – ETERNUS Storage Cluster – non-stop operation despite a complete system or site failure is guaranteed. The challenges posed by data protection are addressed using innovative Self-Encrypting Drives (SED) or system-inherent data encryption methods.

This white paper describes the architecture and hardware features of the ETERNUS DX8900 S4 model. Software functionalities and features are the subject of other publications.

* 8 KB Random Read Miss

** Bigger capacities may apply in the future, depending on upcoming hard disk or SSD capacities.

System architecture

Figure 1 shows the basic architecture of the systems – Controller Modules (CM) interconnected via PCI Express lines over Front-end Routers (FRT). The Controller Modules, as a unit, have one or more host-side interfaces Channel Adapter (CA) and one or more drive-side interfaces installed. To increase the availability of the device, each Controller Module is connected to each Front-end Router so the interconnection of the Controller Modules is quadruplicated – what we call the **Quad Star Architecture**. In addition, a Drive Enclosure (DE) for drive installation has two independent SAS-3 interface ports that are directly connected to two Controller Modules in the Controller Enclosure (CE) for path redundancy. This means that Drive Enclosures that are not directly connected to the Controller Enclosure, can be accessed from any other Controller Module via the Front-end Routers in the Front-end Enclosure (FE). Thus access to any data located in any drive is possible from any Controller Module via redundant paths.

Figure 2 shows the schematic architecture of an ETERNUS DX8900 S4 Controller Module (CM). The basic functional blocks are the Intel Xeon processor which runs the firmware-based operating system, DRAM System Memory including the Data Cache and the Channel Adapters (CA) which realize the connection to the hosts. Each Controller Module connects to a maximum number of 15 Drive Enclosures (DE). The Drive Enclosures are connected via SAS-3 redundant paths to the SAS Expanders (EXP) in the Controller Modules. The four-ported PCIe 3.0 Interface (INF) realizes the high-bandwidth connectivity to the Front-end Routers.

Each CE can be optionally equipped with up to 24 (2.5 inch) dual-port SSDs or SAS drives.

The Controller Modules residing in pairs in a Controller Enclosure are interconnected via the Front-end Enclosures (FE) consisting of four redundant Front-end Routers (FRT) which enable PCIe 3.0 high-speed data interchange between the controllers (Figure 3). Each CE can additionally support up to 8 dual-port NVMe (PCIe) SSDs (and up to 96 NVMe SSDs in the total system) as secondary flash cache (NVMe Extreme Cache). Usage of NVMe Extreme Cache lowers costs while offering superior performance.

The minimal configuration of an ETERNUS DX8900 S4 model consists of 2 controllers housed in one Controller Enclosure (CE) and can be enhanced by adding pairs of Controller Modules up to the maximum configuration of 24 controllers in 12 Controller Enclosures. Both the performance and the capacity of the system scale with the number of controllers.

The combination of Front-end Router PCIe 3.0 and the Back-end SAS-3 drive interface is key to enable outstanding data throughput in any ETERNUS DX8900 S4 system.




Figure 1: Basic system architecture




Figure 2: Controller Module (CM)




Figure 3: Front-end Routers (FRT) enable PCIe 3.0 high-speed data interchange between the controllers.

Legend

FE	Front-end Enclosure	CM	Controller Module
FRT	Front-end Router	CA	Channel Adapter
CE	Controller Enclosure	DE	Drive Enclosure

The schematic configuration of the highest possible configuration is shown in Figure 4.




Figure 4: Schematic illustration of the highest possible configuration.

The maximum configuration of ETERNUS DX8900 S4 consists of 24 Controller Modules with 96 Channel Adapters and 180 Drive Enclosures holding up to 6,912 disk drives.

Mechanically the system populates one 19", 42U rack (minimum); ETERNUS DX8900 S4 can be expanded up to 12 fully populated racks (Figure 5).

This architecture is the basis for the extremely flexible configuration options of Fujitsu's high-end system. In addition, 3.5" and 2.5" disk drive shelves can be freely mixed in the system.

The number of controllers is related to the number of disks and vice versa. Thus the new architecture ensures that performance always scales out when higher capacity is needed.

Furthermore, the number of Solid State Drives (SSD) in the system is not limited. This means that all drive slots can be populated with SSDs, thereby running the DX8900 S4 in a flash-only configuration.

The Quad Star Architecture delivers big performance and high-availability advantages as the internal data paths are virtual point-to-point connections, thus avoiding latencies which are typical of switched or matrix-like back-end concepts.

In addition, it offers wide-ranging scaling capacity by simultaneously determining (without limits) which disk tier classes are utilized in line with the needs of specific applications.

The same wide choice applies to the host connectors, which are just modules within the controller and can be easily exchanged as fabric requirements change.




Figure 5: ETERNUS DX8900 S4 can be expanded up to 12 fully populated racks.

Controller Module (CM)

The Controller Module is the basic component of the ETERNUS DX8900 S4 system.

The block diagram below shows the components which are incorporated in the Controller Module.

The CM consists of four CAs, two SAS Controllers, one SAS EXP, one Direct memory access (DMA) engine, and one non-transparent bridge (NTB). The NTB is equipped with a four-port interface for directly connecting CMs with the PCIe cables. Up to four Channel Adapters (CA) per Controller Module realize the host connections.

All components within the Controller Module are connected as indicated in this table:

Connection	Technology
CPU-Memory	DDR4
CPU-CA	PCIe 3.0
CPU-INF	PCIe 3.0
CPU-IOC	PCIe 3.0
CPU-BUD	PCIe 3.0

Controller Module (CM)

Cache and main memory

The main memory or cache memory consists of six DIMM slots that can be populated with 8, 16, 32 or 128 gigabyte (GB) DDR4 DIMMs. The CM supports cache up to 128 GB DIMM. The minimum main memory for one Controller Module is 96 GB for a single DX8900 S4 Controller Module, so that a fully extended ETERNUS DX8900 S4 can offer up to 18 TB of cache memory.

System	ETERNUS DX8900 S4
min. Cache/CM	96 GB
min. Cache/CE	192 GB
max. Cache/CE	1,536 GB
max. Cache/System	18 TB

The combination of highly dynamic cache allocation mechanisms with a very lean operating system kernel supports very efficient and performance-optimized use of the main memory for caching purposes and not for system overhead.




Figure 6: Controller Module (CM)

Controller Module (CM)

Channel Adapter (CA)

The Channel Adapters (CA) realize the connectivity to the hosts via the storage network as well as the array-to-array connectivity for remote mirroring. Up to four Channel Adapters can be mounted on each Controller Module.

Currently four different types of Channel Adapters are available

- 2-port FC 16 Gb/s / 32 Gb/s
- 4-port FC 16 Gb/s / 32 Gb/s
- 2-port iSCSI 10 Gb/s
- 2-port iSCSI 1 Gb/s

Remote replication between arrays can be realized with both FC and iSCSI.

All types of Channel Adapters can be mixed on a Controller Module or in the complete system without limitations. This offers the highest flexibility in terms of connectivity and disaster recovery options during the entire lifecycle of the system. Even if new connectivity technologies or protocols come to market, the flexible CM/CA design enables quick adoption in existing systems, and thus offers additional investment protection.




Figure 7: All types of Channel Adapters can be mixed on a Controller Module.

Controller Enclosure (CE)

A Controller Enclosure is a 3U, standard 19" frame and always holds the following main components: 2 Controller Modules, 8 PCI Express Flash Modules, 2 Power Supply Units, 2 Battery Backup Units and 4 Boot up and Utility Device (BUD) units. It also incorporates the Service Controller, FPGA for accelerated processing, and the SAS-3 drive interfaces. Up to 24 (2.5 inch) additional dual port NVMe-SSDs or SAS drives can be installed in the CE, which are connected to both CMs. SAS interfaces between CMs are cross-connected to provide a redundant path to the drives on the CEs. In addition, a Drive Enclosure (DE) for drive installation has four independent interface ports that are directly connected to two CMs for path redundancy. Each CE can be connected with 15 DEs (three DEs for one path group and

four DEs for three path groups). Indirect access via FRT is used for DEs that are not directly connected to the CE. Up to 15 Drive Enclosures in a mixed configuration, Small Form Factor 2.5" (SFF), Large Form Factor 3.5" (LFF) and SAS 3.0 3.5" Drive High Dense DE, can be connected per Controller Enclosure.

The architecture was designed under the stipulation that all components must be non-disruptively upgradable or replaceable in the running system. Thus all components in the Controller Enclosures are hot-pluggable.

Starting with one Controller Enclosure, ETERNUS DX8900 S4 scales up to 12 Controller Enclosures/24 Controllers.





Figure 8: Schematic illustration of Controller Enclosures – all components in the enclosures are hot-pluggable.



Legend

- | | |
|--------------------------------|----------------------|
| BUD Boot up and Utility Device | CA Channel Adapter |
| BBU Battery Backup Unit | CM Controller Module |
| PSU Power Supply Unit | |

NVMe Extreme Cache

The system cache can be expanded using the Non-Volatile Memory Express (NVMe) Flash secondary cache (also called the NVMe Extreme Cache), which is located directly within the Controller Enclosure (CE). It can be accessed by both controller modules of the CE. Usage of NVMe Extreme Cache (which can directly connect to the PCIe bus) mounted on the CE improves performance while using fewer drives, thus optimizing system-wide costs and power consumption. ETERNUS DX8900 S4 supports up to 8 dual-port NVMe (PCIe) SSDs per CE, each with capacities of 1.6 TB or 3.2 TB. One CE can support up to a total of 25.6 TB NVMe Extreme cache capacity for read* operations.

A fully equipped ETERNUS DX8900 S4 system can support up to 96 NVMe Dual Port SSDs or 307.2 TB of NVMe Extreme Cache.

Boot up and Utility Device (BUD)

The Boot up and Utility Devices are very fast non-volatile SSD drives which have two major functions:

1. They contain the operating system software and the system configuration data
2. They provide sufficient storage capacity for flushing cached data, which has not yet been written to disk from cache, to the SSD in case of power faults.

A BUD consists of two such SSDs in a mirrored configuration.

Battery Backup Unit (BBU)

During a power outage or power failure, the Battery Backup Unit provides sufficient power to destage the data residing in the cache to the non-volatile vault drives realized by the Boot up and Utility Device. This protects the data in the cache memory without any limitations to the retention period, no matter how long the outage lasts.

Front-end Router (FRT)

The Front-end Routers are the central components of the Quad Star Architecture. They are high-performance PCIe switches and realize the any-to-any connections between the controllers. Four units which are hot-swappable in case of failures are always installed in the Front-end Enclosure.


One Controller Enclosure and the Drive Enclosures connected directly to that Controller Enclosure, collectively forms a "node." The Quad Star Architecture realizes the high-performance access of Drive Enclosures on a different node, accessing the Drive Enclosure indirectly through the Front-end Router using ultra-fast PCIe 3.0 connections. w

Front-end Enclosure (FE)


The Front-end Enclosure is the central component of ETERNUS DX8900 S4 storage systems. It realizes the controller-controller connection and communication. It is a 5U, standard 19" frame and holds the following main components: 4 Front-end Routers (FRTs), 2 Service Controllers (SVCs), 1 System Operation Panel (installed), 4 FAN units, and 4 Power Supply Units (FPSU). The FE inside is divided into two path groups, each of which consists of one SVC and two FRTs. The FE internal structure and SVC are duplicated and the FRT is quadruplicated for redundancy. The Service Controllers control the monitoring function for the entire ETERNUS DX system.

The integrated System Operation Panel provides basic information about the system status using LED indicators and messages on the LCD screen.

Front



Rear



Legend

- SVC Service Controllers
- PSU Power Supply Unit
- FRT Front-end Routers

Drive Enclosure (DE)




The utilized Drive Enclosure types are identical from ETERNUS DX60 S4 up to ETERNUS DX8900 S4. Drive Enclosures are 2U, standard 19" enclosures, and ETERNUS DX8900 S4 can be equipped with a maximum of 180 Drive Enclosures (DE) in a total of 12 cabinets.

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

The ETERNUS DX8900 S4 series also uses reverse cabling connections for the data transfer paths between the controllers and the Drive Enclosures. This ensures continued access even if a failure occurs in a Drive Enclosure.

Three different Drive Enclosure types are available:

- a Small Form Factor (SFF) Drive Enclosure holds up to 24 x 2.5" disk drives
- a Large Form Factor (LFF) Drive Enclosure holds up to 12 x 3.5" disk drives
- a Large Form Factor (LFF) High Dense Drive Enclosure holds up to 60 x SAS 3.0 3.5" SAS Drives



Legend
 IOM I/O Module
 PSU Power Supply Unit

Different disk types having the same form factor can be mixed within one Drive Enclosure without restriction. Up to 15 Drive Enclosures with a maximum of 576 disks in mixed configurations (SFF and LFF) can be connected per Controller Enclosure in addition to the disks residing in the CE.

As the operating system and the system configuration data resides on the backup disks in the controller enclosure, all disks in the drive enclosures can be used for the storage of customer data.

Disk drives

All disk drives in ETERNUS DX S4 systems are equipped with an SAS-3 12 Gb/s interface to realize the connection to the drive enclosures. Three different main types are available and realize different storage tiers in terms of performance and reliability:

- Very fast Solid State Drives (SSD) for the highest IOPS and response time requirements.
- High-performance SAS drives with both 15k and 10k RPM
- Low-spinning Nearline-SAS drives with 7.2 RPM for backup, archiving and low-access data

All these types are available with different capacities in both LFF and SFF, and optionally as Self-Encrypting Disks (SED).

ETERNUS DX8900 S4 thus offers flexibility with nearly unlimited disk combination options to perfectly match any kind of application or customer requirement.



System firmware

ETERNUS DX8900 S4 offers multifold firmware-based features and functionalities which complement the hardware design to improve the reliability and performance of the system.

Advanced RAID organization ensures data access even during a drive enclosure failure. Each disk drive in a RAID group is located in a different Drive Enclosure, and attached to different Controller Enclosures. This keeps the RAID group accessible even if one entire Drive Enclosure, Controller or Controller Enclosure becomes unusable. Available RAID levels are RAID 1, 1+0, 5, 5+0 and 6.

Fast recovery rebuilds data very quickly by simultaneously relocating data in the failed drive to space reserved on the remaining drives of the RAID.

Data block guard guarantees very high data integrity and improved system reliability. This function adds check codes for identification of all data written. It can also confirm/guarantee data integrity at multiple checkpoints along the data transmission path.

Disk drive patrol is a background process that constantly checks the disk drives to detect errors and write failures. Erroneous data is reconstructed and written to a new valid area on the disk drive. Thus disk write failures are avoided in a preemptive manner.

Controller based encryption prevents unauthorized data access LUN-wise, whereas Self-Encrypting Disks do not offer this kind of granularity. The selected encryption technology can either be the world-standard 128-bit AES technology or Fujitsu's unique encryption with high-process performance.

Encrypted communication between ETERNUS DX and user terminals via SSL/SSH guarantees secure transfer of data through a network and prevents unauthorized access and information theft.

Thin provisioning helps lower initial investments in disk drives while offering virtual pools of disk capacity.

Quality of Service options help administrators improve or maintain system performance for the most demanding and business-critical applications.

Storage compression improves capacity optimization by offloading the data reduction process into a dedicated Storage Acceleration Engine (SAE), thereby reducing the consumption of CPU resources. To optimize compression in database environments, the unit of the compression process adopted is 8 KB. The compression process is performed across multiple CMs (24 CMs in maximum configuration); thus the performance of the compression process improves in proportion to the number of CMs. The system also features support for large capacities of up to 48PB for each compression pool, along with support for the UNMAP command which improves the capacity efficiency in virtualized environments.

A more detailed description of all DX8900 S4 features and the even more sophisticated functionalities of the ETERNUS SF management suite (which complements and enhances ETERNUS DX systems) is available in the ETERNUS DX software white paper.



Remote service

Fujitsu's AIS Connect is an optional software feature that offers call-home functionality in case of alerts, incidents and other events. In addition, the solution allows server-less remote access to ETERNUS storage systems for remote services in customer environments.



It can either be leveraged as a software functionality of the firmware or with a separate device – the AIS Connect Remote Support Gateway. The device consists of a rack-mountable 1U, standard 19" enclosure.

Implementation of this real-time service reduces the potential downtimes of ETERNUS storage systems to a minimum. The call-home feature can alert the maintenance provider to issues long before a potential system failure becomes imminent. System specialists can remotely access the system for preemptive maintenance, fast analysis and troubleshooting to prevent incidents.

Environmental sustainability

The ETERNUS DX8900 S4 series was engineered to reduce the environmental load for both the manufacturing process and operations.

The main target in development and design was low power consumption, which has been achieved by the deployment of power-saving SFF and SSD hard drives, along with highly efficient power supplies and intelligent control of fan rotation.

In addition, power and cooling can be reduced by the ECO-mode feature which controls drive spindle rotation and disk "spin down" for idling disk drives.

Real-time monitoring of power consumption and temperature with ETERNUS SF management suite helps raise storage administrators' awareness for power savings.

The required battery capacity was significantly reduced thanks to the power-saving backup method realized with the Backup Disk implementation on the Controller Module. This was also a prerequisite for the deployment of NiMH batteries which save 50 percent in weight and 80 percent in volume while almost doubling the battery lifetime when compared to the lead batteries formerly used.

The avoidance of hazardous chemical substances is a key prerequisite throughout the entire manufacturing process and is complemented by the 3R design of systems and manufacturing – Recycle, Reuse, Reduce.



Figure 9: Real-time monitoring of power consumption and temperature with ETERNUS SF helps raise storage administrators' awareness for saving energy.

Summary

The Fujitsu Storage ETERNUS DX8900 S4 series provides non-disruptive capacity upgrades up to 6,912 disk drives, linearly scaling performance and capacity by adding Controller Modules and Drive Enclosures. Maximum performance is provided by configurations based on from 2 to 24 controller units, with maximal 384 front-end ports for host connectivity. Optional dual-port SSDs and SAS drives can be mounted on CEs along with additional NVMe Extreme Cache, which can significantly boost performance. For simultaneous connection to different network types, FC and iSCSI host interfaces can operate in mixed configurations. The system also features support for FC 32 Gbps.

Tiered storage optimization using a range of SAS, Nearline SAS and SSD drive options lowers the overall total cost of storage. The challenges of data protection are addressed by Self-Encrypting Drives (SED). Flexible management across

different quality of service profiles, using different drive types for each application scenario, ensures efficient assignment of system resources matched to the performance priorities of different applications.

The use of hardware accelerated data reduction across storage tiers by offloading the compression process into dedicated Storage Acceleration Engine (SAE) leads to efficient capacity optimization. System robustness with redundant interconnections to the storage array, internal redundant components and local data copy options ensure high availability for today's non-stop business. Additional extensive high-availability and disaster recovery features including remote replication and its native transparent fail-over capability make ETERNUS DX enterprise models the ideal storage systems for all business-critical data – Business-centric Storage.

Appendix – technical specifications

ETERNUS DX8900 S4	
Architecture	Scale-out
Storage controllers	2-24
Maximum storage capacity	141,558 TB
Maximum disk drives	6,912
Maximum cache memory (DRAM)	18 TB
Maximum second-level cache (SSD)	307.2 TB
Maximum connectable hosts	8,192

Appendix – online resources

- Fujitsu Global Portal: ETERNUS DX8900 S4 product page
- Datasheet: ETERNUS DX8900 S4
- Brochure: ETERNUS DX Family
- Whitepaper: ETERNUS AF/DX Feature Set

Published by
Fujitsu Limited

www.fujitsu.com

© 2022 Fujitsu. All rights reserved, including intellectual property rights. Technical data subject to modifications and delivery subject to availability. Any liability that the data and illustrations are complete, actual or correct is excluded. Designations may be trademarks and/or copyrights of the respective manufacturer, the use of which by third parties for their own purposes may infringe the rights of such owner.

For further information see www.fujitsu.com/eternus
FUJITSU-PUBLIC