

INTRODUCTION TO THE EMC VNX SERIES

VNX5100, VNX5300, VNX5500, VNX5700, & VNX7500

A Detailed Review



Abstract

This white paper introduces the EMC® VNX™ series unified platform. It discusses the different models, new and improved features, and key benefits.

December 2013

Copyright © 2013 EMC Corporation. All rights reserved.

EMC believes the information in this publication is accurate of its publication date. The information is subject to change without notice.

The information in this publication is provided “as is.” EMC Corporation makes no representations or warranties of any kind with respect to the information in this publication, and specifically disclaims implied warranties of merchantability or fitness for a particular purpose.

Use, copying, and distribution of any EMC software described in this publication requires an applicable software license.

For the most up-to-date listing of EMC product names, see the EMC corporation trademarks page at EMC.com.

VMware is a registered trademark of VMware, Inc. All other trademarks used herein are the property of their respective owners.

Part Number h8217.3

Table of Contents

Executive summary	5
Audience	5
Terminology	5
VNX family overview	7
VNX hardware components	7
Block, file, and unified configurations	9
Gateway configurations	11
VNX VG2	11
VNX VG8	11
Components and features in the VNX	12
Xeon 5600 processor makes VNX faster and greener	12
Lead-free Generation 3 power supplies are 7 to 10 % more efficient	12
High-capacity and small form factor drives	13
VNX VG2 and VG8 gateway support for VMAXe and FCoE	13
DAEs, DPEs, and drives provide increased density	13
6 Gb/s x 4 lanes SAS back end improves performance and availability	26
Expanded UltraFlex I/O using PCI-E Gen 2 increases bandwidth	27
VNX models	28
VNX5100	30
VNX5300	33
VNX5500	34
VNX5700	35
VNX7500	38
I/O modules	39
I/O modules for the storage processor	39
I/O modules for the Data Mover	44
Control Station	48
Standby power supply	49
Software on the VNX series	50
Reserved area on VNX series storage systems	50
Unified management	50
Software for efficiency	51
Fully Automated Storage Tiering for Virtual Pools (FAST VP)	51

FAST Cache	52
Data deduplication	53
Thin Provisioning	55
Software for protection	56
Replicator	56
RecoverPoint integration.....	57
MirrorView/A	57
MirrorView/S	58
SAN Copy.....	58
Business continuance.....	58
SnapView Snapshots.....	58
Checkpoints.....	60
Clones	60
Security.....	60
Anti-virus software	60
Quota management and auditing.....	60
File-Level Retention.....	61
Advanced Data Services Support for file	62
Serviceability, availability, and performance	63
File-side control LUN protection	63
Unified software upgrade.....	63
Unified Remote Support.....	64
Block to Unified Plug and Play	64
In-Family Data in place Conversion	65
VNX scalability and performance.....	67
Conclusion	67
Appendix A: System level, block, and file components.....	68
Appendix B: Storage processor UltraFlex I/O slot availability.....	73

Executive summary

Scalability, efficiency, and simplicity are three challenges that are a constant in the storage industry. Storage systems need to provide scalability, extreme efficiency, and a simple, intuitive user interface.

The VNX series is designed for a wide range of environments that include midtier through enterprise. VNX provides offerings that include file only, block only, and unified (block and file) implementations. The VNX series is managed through a simple and intuitive user interface called Unisphere®; this single pane of glass completes the unified experience.

This white paper discusses the VNX series and its hardware and software, limitations, I/O Modules, etc. This paper also describes the serviceability and availability features that have been added to the VNX, such as LUN protection, Advanced Data Services Support for file, Unified Remote Support, and the Unisphere Service Manager.

Audience

This white paper is intended for IT architects, administrators, and others who are interested in VNX series arrays. You should be familiar with storage array concepts, general hardware, and the software services provided by the arrays.

Terminology

- **Automatic Volume Management (AVM)**—Feature of VNX for File that creates and manages volumes automatically without manual volume management by an administrator. AVM organizes volumes into storage pools that can be allocated to file systems.
- **Converged Network Adapter (CNA)**—A host adapter that allows a host to process Fibre Channel and Ethernet traffic through a single type of card and connection, decreasing infrastructure costs over time.
- **Disk array enclosure (DAE)**—Shelf in VNX that includes an enclosure, either 15 or 25 disk modules, two Fibre Channel link control cards (LCCs), and two power supplies, but does not contain storage processors (SPs).
- **Disk-processor enclosure (DPE)**—Shelf in VNX that includes an enclosure, disk modules, storage processors (SPs), two Fibre Channel link control cards (LCCs), two power supplies, and four fan packs. A DPE supports DAEs in addition to its own disk modules. This 3U form factor is used in the lower-end VNX models (VNX5100™, VNX5300™, and VNX5500™) and supports a maximum of 75, 125, and 250 drives, respectively.
- **Fibre Channel (FC)**—Nominally 1 Gb/s data transfer interface technology, although the specification allows data transfer rates from 133 Mb/s up to 4.25 Gb/s. Data can be transmitted and received simultaneously. Common transport protocols, such as Internet Protocol (IP) and Small Computer Systems Interface (SCSI), run over Fibre Channel. Consequently, a single connectivity technology can support high-speed I/O and networking.

- **Fibre Channel over Ethernet (FCoE)**—Allows Fibre Channel frames to be encapsulated via Ethernet. With Converged Network Adapter cards, a host may use one type of adapter and cable for Ethernet and Fibre Channel traffic.
- **Flash drive** – Flash drives provide extremely high performance and consume very little power. These drives enable a VNX to provide an *Extreme Performance* level of storage for critical workloads. Flash drives can also be used to extend the array's cache using [FAST Cache technology](#).
- **Logical unit number (LUN)**—Identifying number of a SCSI or iSCSI object that processes SCSI commands. The LUN is the last part of the SCSI address for a SCSI object. The LUN is an ID for the logical unit, but the term is often used to refer to the logical unit itself.
- **Near-Line Serial Attached SCSI (NL-SAS) drives**—Enterprise SATA drives with a SAS interface head, media, and rotational speed of Enterprise SATA drives with a fully capable SAS interface. These drives provide similar performance and capacity to SATA drives, utilizing a SAS interface for I/O.
- **Serial-attached SCSI (SAS)**—A communication protocol used to move data to and from computer storage devices. SAS is a point-to-point serial protocol that replaces parallel SCSI bus technology. Expanders facilitate the connection of multiple SAS end devices to a single initiator port.
- **I/O Module** – A hardware component that provides frontend and backend connectivity between the storage processors disk-array enclosures, and hosts.
- **Storage-processor enclosure (SPE)**— –Shelf in VNX for block that includes an enclosure, storage processors (SPs), Fibre Channel link control cards (LCCs), two power supplies, and fan packs. This 2U form factor is used for the higher-end VNX VNX5700™ and VNX7500™ models, and supports a maximum of 500 and 1,000 drives, respectively.
- **Storage pool**—A single repository of homogeneous or heterogeneous physical disks from which LUNs may be created. This concept enables ease of use and technologies such as FAST VP.

VNX family overview

The VNX family includes the *VNXe*® and *VNX* series arrays. The VNXe series is designed for small-to-midsize environments, while the VNX series is designed for midtier-to-enterprise storage environments.

The VNXe series provides high-availability, scalability, and application-aware, wizard-driven disk provisioning in a small, dense, and energy-efficient form factor. For more information about the VNXe series, see the white paper, *EMC VNXe Series Storage Systems* located on powerlink.

The VNX series is EMC's next generation of midrange-to-enterprise products. The VNX series unifies EMC's file-based and block-based offerings into a single product that you can manage with one easy to use GUI. This new VNX software environment offers significant advancements in efficiency, simplicity, and performance.

Benefits of the VNX series include:

- State-of-the-art hardware:
 - PCI-E Gen 2 I/O modules
 - 6 Gb/s x 4 lanes SAS back-end infrastructure
 - More memory
 - Intel's latest Xeon 5600 multi-core processors
 - Expanded UltraFlex I/O, delivered with denser packaging
- Block functionality that supports the FC, iSCSI, and FCoE protocols
- File functionality that supports the NFS, CIFS, and pNFS protocols
- Support for up to 1,000 drives
- Automated tiering with Fully Automated Storage Tiering for Virtual Pools (FAST VP) for block and file data
- Unified replication using EMC RecoverPoint for block and file data
- Updated unified management with Unisphere delivers a cohesive, unified user experience

VNX hardware components

VNX storage systems include the following components:

- **Storage processors (SPs)** support block data with UltraFlex I/O technology that supports Fibre Channel, iSCSI, and FCoE protocols. The SPs provide access for all external hosts and the file side of the VNX array.
- The **storage-processor enclosure (SPE)** is 2U in size and houses each storage processor. This form factor is used for the higher-end VNX VNX5700™ and VNX7500™ models, and supports a maximum of 500 and 1,000 drives, respectively.

- The **disk-processor enclosure (DPE)** is 3U in size and houses each storage processor and the first tray of disks. This form factor is used in the lower-end VNX models (VNX5100™, VNX5300™, and VNX5500™) and supports a maximum of 75, 125, and 250 drives, respectively.
- **Blade** (or *Data Move*) access data from the back end and provide host access using the same UltraFlex I/O technology that supports the NFS, CIFS, and pNFS protocols. The Blades in each array are scalable and provide redundancy.
- The **Data Mover enclosure (DME)** is 2U in size and houses the Data Movers (Blades). The DME is similar in form to the SPE and is used on all VNX models that support file.
- **Standby power supplies (SPSs)** are 1U in size and provide enough power to each storage processor to ensure that any data in flight is de-staged to the vault area in the event of a power failure. This ensures that no writes are lost. When the array is restarted, the pending writes are reconciled and persisted.
- **Control Stations** are 1U in size and provide management functions to the file-side components referred to as Blades. The Control Station is responsible for Blade failover. The Control Station may optionally be configured with a matching secondary Control Station to ensure redundancy on the VNX array.
- **Disk-array enclosures (DAE)** house the drives used in the array. They enable the expansion of each array's drive count to provide the storage needed for the expanding needs over time of each implementation. The 15-drive DAE is 3U in size, the 25-drive is 2U in size, and the 60-drive DAE is 4U in size.
- The SPE/DPE and DME use **dual power supplies** that are integrated into the chassis, and provide back-up electricity if a single power supply fails.
- For NAS services (i.e., Data Mover Blade I/O options):
 - 4-port 1 GbaseT Ethernet copper module
 - 2-port 1 GBaseT Ethernet copper plus 2-port 1 Gb Ethernet optical module
 - 2-port 10 Gb Ethernet optical/twin-ax module
 - 2-port 10 GBaseT Ethernet copper module
- For Block Services to host (i.e., Storage Processor I/O options):
 - 4-port 8 Gb optical Fibre Channel UltraFlex I/O module
 - 4-port 1 GBaseT Ethernet iSCSI/TOE module
 - 2-port 10 Gb Ethernet optical/twin-ax iSCSI/TOE module
 - 2-port 10 GBaseT Ethernet copper iSCSI/TOE module
 - 2-port FCoE optical/twin-ax module

Block, file, and unified configurations

The hardware components in a VNX depend on the overall configuration and options ordered. The three basic configurations are:

- Block—Supports block data
- File—Supports file data
- Unified—Supports block and file data

Block configuration – supports block data (shown in Figure 1)

This configuration supports the FC, iSCSi, and FCoE protocols and consists of:

- 3U DPE (disk-processor enclosure) for VNX5100, VNX5300, and VNX5500 *or* 2U SPE (storage-processor enclosure) for VNX5700, and VNX7500 + at least one DAE for Vault Drives.
- [optional] 2U 25-drive DAE, *or* 3U 15-drive DAE, *or* 4U 60-drive DAE (in deep rack) up to the system maximum
- 1U SPS



Figure 1: Block configurations for VNX5500 (left) and VNX5700 (right)

File configuration – supports file data (shown in Figure 2)

This configuration supports the NFS, CIFS, and pNFS protocols and consists of:

- 3U DPE (disk-processor enclosure) for VNX5100, VNX5300, and VNX5500 *or* 2U SPE (storage-processor enclosure) for VNX5700, and VNX7500 + at least one DAE for Vault Drives
- [optional] 2U 25-drive DAE (disk-array enclosure) *or* 3U 15-drive DAE *or* 4U 60-drive DAE (in deep rack) up to the system maximum
- At least one 2U Blade (or Data Mover) enclosure
- One *or* two 1U Control Stations
- 1U (SPS)



Figure 2. File configurations for VNX5300 (left) and VNX5700 (right)

Unified configuration – supports file and block data (shown in Figure 3)

The unified configuration supports the FC, iSCSi, and FCoE protocols for block data and the NFS, CIFS, and pNFS protocols for file data. A unified configuration has the same components as a file configuration.



Figure 3. Unified configurations for VNX5300 (left) and VNX5700 (right)

Gateway configurations

Each VNX Series Gateway product—the VG2 or VG8—is a dedicated network server optimized for file access and advanced functionality in a scalable, easy-to-use package. The gateway systems can connect to, boot from, and work with Symmetrix, the VNX series, and CLARiiON back-end array technologies, as well as multiple systems.

VNX VG2

- One or two 2U Blade (or Data Movers)
- Four core 2.40 Ghz Intel Xeon 5600 processors with 6 GB memory
- 256 TB usable capacity per Blade (256 TB per system)

VNX VG8

- Two to eight 2U Blades (or Data Movers)
- Six core 2.83 Ghz Intel Xeon 5600 processors with 24 GB memory
- 256 TB usable capacity per Blade (1,792 TB per system)

Components and features in the VNX

VNX series arrays are comprised of advanced components that have the following features:

- Xeon 5600 processor makes VNX faster and greener
- Lead-free Generation 3 power supplies are from 7 to 10% more efficient
- High- capacity and small form factor drives
- VNX VG2 and VG8 gateway support for VMAXe and FCoE
- Increased density from new DAEs, DPEs, and drives
- 6 Gb/s x 4 lanes SAS backend improves performance and availability
- Expanded UltraFlex I/O using PCI-E Gen 2 increases bandwidth

These components and the features each provides are discussed in more detail below.

Xeon 5600 processor makes VNX faster and greener

The VNX series uses Intel's Xeon 5600 processor. This new 32 nm processor architecture provides multi-core capability that translates into significant scalability for each VNX model. The processor speeds range from 2.13 GHz to 2.8 GHz and have four to six cores per socket. As a result, these processors help you accomplish more in a smaller footprint, consume less power, and keep systems cooler.

Lead-free Generation 3 power supplies are 7 to 10 % more efficient

The VNX series is more power-efficient than previous storage systems. The power supply provides the information necessary for Unisphere to monitor and display the ambient temperature and power consumption of the power supply. The Generation 3 power supplies:

- Are 7 % to 10 % more efficient than Gen 2 power supplies used on CX4/NS arrays
- Use a lead-free construction
- Meet the EPA Gold Standard for efficiency targets (90%)
- Provide adaptive cooling, in which an array adjusts the power supply fan speeds to spin only as fast as needed to ensure sufficient cooling. This optimizes blower speeds and energy consumption. The power supplies have an onboard sensor that provides energy information reporting. The reports include input power and air inlet temperature for each enclosure on the array, as well as input power reporting at an overall array level. The values are based on a rolling average for each hour that is reported on in the Unisphere GUI, as shown in Figure 4.




Figure 4. Energy information provided in the Unisphere Environment tab

High-capacity and small form factor drives

The VNX series supports high-capacity 3 TB drives, along with small form factor 2.5" flash drives. These drives are supported with VNX arrays running VNX OE for File v7.0.3X and VNX OE for Block v5.31.000.5.XXX.

VNX VG2 and VG8 gateway support for VMAXe and FCoE

The VNX series gateway models (VG2 and VG8) now support FCoE I/O modules on the storage processors, along with VMAXe backend connectivity. VMAXe is the Symmetrix business class that complements the higher-end line VMAX.

DAEs, DPEs, and drives provide increased density

The VNX series uses updated components that make it significantly denser than earlier models:

- **25-drive DAEs and DPEs paired with 2.5" SAS drives**—Increased performance and spindle density are provided by 25-drive DAEs and DPEs paired with 2.5" SAS drives.
- **15-drive DAEs paired with 3.5" 2 TB 7,200 rpm NL-SAS drives**—In terms of drive capacity, maximum density is achieved by using 15-drive DAEs paired with 3.5" 2 TB 7,200 rpm NL-SAS drives. This combination allows more spindles and storage capacity per square foot than in any other previous array using the 15-drive enclosures.
- **60-drive high-capacity DAE paired with 3.5" 7.2 K rpm NL-SAS drives and 2.5" 10 K rpm drives using a 3.5"-carrier**—In terms of drive capacity, maximum density is achieved by using 3.5" 3TB 7.2K rpm NL-SAS drives. This combination allows more spindles and storage capacity per square foot than in any other VNX-compatible

DAE. Increased performance may be achieved by using 2.5" and 3.5" flash and 10 K rpm drives with the use of a 3.5"-carrier. This DAE is only available on arrays running VNX-OE for Block R 31.5 and/or VNX-OE for File 7.0.3. 2.5" 300 GB 10 K drives are not supported in this enclosure. This DAE requires a deep rack to house it.





Figure 5: Block dense configuration example



• **Figure 6: Unified dense configuration example**

- **3U DPE for lower-end models**—Improved rack space utilization in the VNX5100, VNX5300, and VNX5500 is achieved with an updated DPE. This 3U DPE houses storage processors and the first shelf of disks in a 3U unit. This is a dramatic improvement in rack-space utilization compared to previous generations.

The 25-drive 2.5" SAS-drive DAE

Figure 5 and Figure 6 show the front and back of the 25-drive DAE that is housing 2.5" SAS drives.




Figure 5. Front of the 25-drive DAE housing 2.5" SAS drives



Figure 6. Back of the 25-drive DAE housing 2.5" SAS drives

Figure 7 shows the back of the 25-drive DAE. In this figure, you can see the primary/extension SAS ports; SPS ports; and the LEDs for power, fault, SAS lane status, bus number, and enclosure number.




Figure 7. A close-up of the back of the 25-drive DAE

The 15-drive DAE and DPE supports 2.5" SAS, 3.5" NL-SAS, and 3.5" flash drives

The 15-drive DAE, shown in Figure 8 and Figure 9, may be populated with 2.5" drives in 3.5" carriers. The 15-drive DAE may also be populated with any combination of 3.5" EFD, SAS, and NL-SAS drives.

These options enable you to configure your system for optimal efficiency and performance. The flash drives provide extreme performance, the mid-tier SAS drives provide a good balance of price and performance (reaching speeds of up to 15k rpm), and the cost-effective NL-SAS drives have a capacity of up to 2 TB.

These 15-drive DAEs, while similar in appearance to previous-generation CX™ and NS models, are not backward-compatible with CX and NS arrays¹. Also, because the DAEs use SAS as the backend bus architecture, the VNX series arrays do not accept previous-generation DAEs.

¹ To learn how to move data from a NS or CX model to a VNX model, please refer to the white papers [Migrating Data from an EMC Celerra Array to a VNX Platform using Celerra Replicator](#) and [Migrating Data from an EMC CLARiiON Array to a VNX Platform using SAN Copy](#).




Figure 8. Front of the 15-drive DAE



Figure 9. Back of the 15-drive DAE

As you can see in Figure 9, there is no bus-reset button on the back of the DAE; this is not needed, because VNX uses SAS technology for the backend.

Figure 10 shows the back of the 15-drive DAE. In this figure, you can see the primary/expansion SAS ports; and the LEDs for power, fault, SAS lane status, bus number, and enclosure number.




Figure 10. A close-up of the back of a 15-drive DAE

The 60-drive high-capacity DAE supports 2.5" SAS, 3.5" NL-SAS, and 3.5" flash drives

The 60-drive DAE can be populated with 2.5" SAS drives using a 3.5"-carrier. The 60-drive will hold up to 60 rotating or SSD-type drives in 3.5" (EFD, 7200 rpm and 10 K rpm) and 2.5" (EFD, 7200 rpm, 10 K rpm) form factors. For the 3.5" drive types, 15 K rpm drives are not supported in this enclosure.

These options allow you to configure your system for optimal efficiency and performance. The flash drives provide extreme performance in SFF (small form factor), 2.5" and LFF (large form factor), 3.5" sizes. The mid-tier SAS drives provide a good balance of price and performance (reaching speeds of up to 10 K rpm), and the cost-effective NL-SAS drives have a capacity of up to 3 TB.

The enclosure employs a slide-out drawer with access to the drives from the top. The drive matrix consists of 5 rows (banks A-E) of 12 disks (slots 0-11). These disks are addressed and notated using a combination of letters and numbers, such as A1 and B4, to uniquely identify a single drive in the enclosure. The enclosure itself is labeled clearly to ensure proper drive identification.

The dimensions of the 60-drive high-capacity enclosure are 7" (height) (4U) x 35" (depth, chassis only) with a maximum weight of 260 lbs. Because the design of this industry-leading enclosure provides far more density per square foot of rack space

than any other enclosure available, a special dimensionally enhanced rack (44" deep) is required for this enclosure. The deep rack provides three power zones (A, B, and C), each with their own power drops to power each zone. Each 60-drive high-capacity DAE requires four power connections for proper operation. Each generation of three power supplies are FRUs. These power supplies may be removed without detaching the bus cable by using the cable management arm (CMA) quick disconnects labeled in orange.

The 60-drive high-capacity enclosure may be the first DAE and contain the vault drives (row A, positions 0-3) for the VNX5700 and VNX7500 models. If there is a partially filled row in the DAE, the remaining empty slots will be populated with fillers.

The 60-drive DAE contains four hot swappable LCCs:

- Two ICM LCCs are located at the rear of the enclosure. They house the SAS ports and provide connectivity to the storage processors.
- Two internal LCCs are centrally located and provide connectivity to all of the disks in the enclosure.

All LCCs on this enclosure are FRUs. There are also three cooling modules located on the front of the enclosure that are FRUs.

Figure 11 shows the front of the 60-drive DAE.



Figure 11. Front of the 60-drive DAE

Figure 12 shows the top of the 60-drive DAE extended with the top panel open, which reveals the drives and cooling modules.




Figure 12. Top of the 60-drive DAE

Figure 13 shows the internal LCC cards that insert in the middle of the 60-drive DAE from the top of the DAE.




Figure 13. The internal LCC cards

Figure 14 shows one of the three cooling modules removed from the 60-drive DAE.




Figure 14. Cooling module in the 60-drive DAE

Figure 15 shows the cable management arms at the rear of the 60-drive DAE.




Figure 15. Cable management arms in the rear of the 60-drive DAE

Figure 16 shows a power supply at the rear of the 60-drive DAE, the ICM LCC (which contains the bus and enclosure LEDs), and SAS expander ports.




Figure 16. Power supply, ICM LCC, and SAS expander ports

6 Gb/s x 4 lanes SAS back end improves performance and availability

The backend of the VNX series uses an architecture that dramatically improves performance and availability.

The 6 Gb/s x 4 lanes SAS backend provides speeds of up to 3 GB/s. SAS is a significant improvement over previous-generation FC backends that provide speeds of 0.25 GB/s (4Gbps). This point-to-point architecture, in which all drives connect directly to an expanding switching matrix, allows direct communication from the controller to the drive. The data frames do not have to travel throughout the entire bus accessing each of the drives, as was the case in previous models that used a Fibre Channel arbitrated loop as the back-end architecture. This provides much higher bandwidth across the VNX line than previous models. In addition, because SAS uses acknowledgements (ACK/NAK), in case there is a failure, an immediate response is received, unlike Fibre Channel where the need to wait for a time-out is necessary.

The SAS backend increases availability by offering improved backend fault isolation, and VNX series systems benefit from this dramatic improvement. Because data frames no longer traverse the entire bus and do not pass through every drive, a corrupt frame is much easier to isolate to a specific drive or SAS cable, especially when compared to

Fibre Channel, where locating the exact drive or cable that corrupted a frame requires significant time and resources.

The white papers, [Migrating Data from an EMC Celerra Array to a VNX Platform using Celerra Replicator](#) and [Migrating Data from an EMC CLARiiON Array to a VNX Platform using SAN Copy](#), discuss data migration options.

Expanded UltraFlex I/O using PCI-E Gen 2 increases bandwidth

The VNX series expands the UltraFlex I/O technology of previous-generation midrange storage arrays. The UltraFlex I/O uses PCI-E Gen 2, which offers up to two times the bandwidth of the UltraFlex cards used in previous-generation models. This increase in bandwidth is an important factor in the enhanced performance of VNX arrays over previous generations.

VNX models

VNX series models are comprised of 5000 and 7000 class systems. The available models are VNX5100, VNX5300, VNX5500, VNX5700, and VNX7500. The hardware and connectivity options scale with each model, providing more power and options throughout the model range.



Figure 17. VNX series storage systems

All VNX models support 2.5” and 3.5” SAS drives using either the 15-drive DAE or the 25-drive DAE, and utilize the 6 Gb/s SAS backend. FC drives are no longer supported. SAS drives, NL-SAS rotational drives, and flash drives will be used moving forward.

The following table shows the types of models available in the VNX series. The sections after the table give more detailed information about each model.

Table 1. Model comparison chart

Unified platform model	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Industry position	Entry	Midrange/ Entry	Midrange/ Midtier	High-end/ Mid-capacity	High-end/ large-capacity
EMC predecessor	AX4-5	NS-120 and CX4-120	CX4-240	NS-480 and CX4-480	NS-960 and CX4-960
Management software	Unisphere	Unisphere	Unisphere	Unisphere	Unisphere
Drive count	75	125	250	500	1000
Drive types	Flash, SAS, NL-SAS	Flash, SAS, NL-SAS	Flash, SAS, NL-SAS	Flash, SAS, NL-SAS	Flash, SAS, NL-SAS
Block: Protocols	FC	FC, iSCSi, FCoE	FC, iSCSi, FCoE	FC, iSCSi, FCoE	FC, iSCSi, FCoE
Standby power supplies	1/2	1/2	2	2	2
File: Protocols	n/a	NFS, CIFS, pNFS	NFS, CIFS, pNFS	NFS, CIFS, pNFS	NFS, CIFS, pNFS
File: number of Blades	none	1-2	1-3	2-4	2-8
Array enclosure/ SP count	DPE/2 SP	DPE/2 SP	DPE/2 SP	SPE/2 SP	SPE/2 SP
CS count	NA	1-2	1-2	1-2	1-2

VNX5100

The VNX5100 is designed for the entry-to-mid-tier space. This model provides block-only services, supports the Fibre Channel protocol for host connectivity, and utilizes a DPE form factor.

This model uses a 1.6 GHz, two-core Xeon 5600 processor with 4 GB of RAM and a maximum of 75 drives with the following host connectivity option: FC. This model does not accept any UltraFlex I/O modules and, as such, only supports four Fibre Channel onboard ports for host connectivity. Note the systems ship with only 2 SFPs per SP in mezzanine ports 2-3. Seen here are optionally added SFPs.

This model uses a DPE that is available in 15 x 3.5" drive or 25 x 2.5" drive form factors. It includes four onboard 8 Gb/s Fibre Channel ports and two 6 Gb/s SAS ports for backend connectivity on each storage processor. A micro-DB9 port and service LAN port are available. EMC service personnel use these ports to connect to the VNX.

A LAN connection is provided on each SP for array management. Each SP in the enclosure also has a power supply module and two UltraFlex I/O module slots that are not available for expansion on the VNX5100. These I/O module slots will be populated with blanks (to ensure proper airflow).



Figure 18. Back view of the DPE with SP A (on the right) and SP B (on the left)

As you can see in Figure 18, the System Information tag is located on the back of the SP A side of the DPE (on the right side above). This tag contains serial number and part number information and should be handled carefully when racking the system.

Figure 19, Figure 20, and Figure 21 show the backs of the DPE-based storage processor.



Figure 19. A close-up of the back of the DPE-based storage processor




Figure 20. Power, fault, activity, link, and status LEDs




Figure 21. Storage processor onboard ports

The front of the DPE houses the first tray of drives. This may range from 15 to 25 disks, depending on which DPE has been selected for the model. At the top of the enclosure, from left to right, there is the enclosure fault LED (amber if there is a fault), and the enclosure power LED (blue). The drive LED's are shown at the bottom which will also switch to amber if there is an issue. This is shown in Figure 22.



Figure 22. Front of a 25-drive DPE with 2.5" drives

Because the VNX5100 is block-only, this system does not use Blades.

VNX5300

The VNX5300, the introductory model for the VNX unified platform, is designed for the mid-range entry space. This model provides either block and file services, file only services, or block only services, and uses a DPE.

The VNX5300 uses a 1.6 GHz, four-core Xeon 5600 processor with 8 GB RAM and a maximum of 125 drives with the following block-based host connectivity options: FC, iSCSI, and FCoE.

This model uses a DPE and is available in 15 x 3.5" drive or 25 x 2.5" drive form factors. It includes four onboard 8 Gb/s Fibre Channel ports and two 6 Gb/s SAS ports for backend connectivity on each storage processor. There is a micro-DB9 port and service LAN port available, which are used as a backup method of connectivity for system maintenance.

A LAN connection is provided on each SP for array management. Each SP in the enclosure also has a power supply module and two UltraFlex I/O module slots. Both I/O module slots may be populated on this model. Any slots without I/O modules will be populated with blanks (to ensure proper airflow).

The front of the DPE houses the first tray of drives. The number can range from 15 to 25 disks, depending on which DPE has been selected for the model.

The File/Unified model uses a 2.13 GHz, four-core Xeon 5600 processor with 6 GB RAM per blade, with a maximum storage capacity of 200 TB per blade. The model can have one or two blades, and each blade has redundant power supplies located in the front of the blade. The File/Unified model supports the following NAS protocol options: NFS, CIFS, and pNFS. The CPU module is also accessible via the front of each blade after removing the power supplies. The rear of each blade houses the slots used for internal/external network connectivity, backend fibre connectivity, and frontend connectivity options.

The File/Unified model uses a DME (Data Mover Enclosure) that appears similar in appearance to the SPE of the VNX5700 and VNX7500. The DME can house two blades per 2U enclosure, and each blade can accept a maximum of two I/O modules for frontend connectivity. Pictures of the DME are provided in Figure 23, Figure 24, and Figure 25.



Figure 23. Front of a DME



Figure 24. Back of a DME




Figure 25. Close-up of the back of a DME

VNX5500

The VNX5500 is designed for the mid-tier space. This model provides either block and file services, file only, or block only (without blades and Control Stations).

The block only model uses a 2.13 GHz, four-core Xeon 5600 processor with 12 GB RAM and a maximum of 250 drives with the following block-based host connectivity options: FC, iSCSI, and FCoE.

The VNX5500 uses a DPE that is available in 15 x 3.5" drive or 25 x 2.5" drive form factors. The DPE includes four onboard 8 Gb/s Fibre Channel ports and two 6 Gb/s SAS ports for backend connectivity on each storage processor. A micro-DB9 port and service LAN port are available. EMC service personnel use these ports to connect to the VNX.

A LAN connection is provided on each SP for array management. Each SP in the enclosure also has a power supply module and two UltraFlex I/O module slots. Both I/O module

slots may be populated on this model. Any slots without I/O modules will be populated with blanks (to ensure proper airflow).

The front of the DPE houses the first tray of drives. The number can range from 15 to 25 disks, depending on which DPE has been selected for the model.

The blades have a 2.13 GHz, four-core Xeon 5600 processor with 12 GB RAM per blade, with a maximum storage capacity of 256 TB per blade. The VNX5500 can have up to three blades and provides the following NAS protocol options: NFS, CIFS, and pNFS. Each blade has redundant power supplies located in the front of each blade. The CPU module is also accessible via the front of each blade after removing the power supplies. The rear of each blade houses the slots used for internal/external network connectivity, backend fibre connectivity, and frontend connectivity options.

The File/Unified model uses a DME that appears similar in appearance to the SPE. The DME can house two blades per 2U enclosure, and each blade on this model may accept a maximum of three I/O modules for frontend connectivity.

VNX5700

The VNX5700 is designed for the high-end, mid-capacity space. This model provides either block and file services, file only, or block only), and uses a SPE form factor.

Each SP in the SPE uses a 2.4 GHz, four-core Xeon 5600 processor with 18 GB RAM and a maximum of 500 drives with the following host connectivity options: FC, iSCSI, and FCoE.

The VNX5700 has a SPE that uses UltraFlex I/O slots for all connectivity. The first slot houses the internal network management switch, which includes a mini-serial port and service LAN port. EMC service personnel use these ports to connect to the VNX.

A LAN connection is provided on each SP for array management. Each SP in the enclosure also has two power supply modules. Each SP has five I/O module slots. Any slots without I/O modules will be populated with blanks (to ensure proper airflow).

The fault/status light of the SPE form factor is located on the front of the unit. This form factor requires an additional DAE to provide a minimum of four drives for the VNX Array Operating Environment. This model's form factor favors scalability, rather than density, requiring slightly more rack space for a minimal configuration.

The back of the SPE is shown in Figure 26. In this figure, you can see the various components including I/O modules. Figure 27 provides a close-up view of the back of the SPE.



Figure 26. Back of the SPE




Figure 27. A close-up of the SPE

The front of the SPE, shown in Figure 28, houses the power supplies, cooling fans, and CPU modules, along with the fault/status LED.



Figure 28. Front of the SPE

The blades in this model use a 2.4 GHz, four-core Xeon 5600 processor with 12 GB RAM per blade with a maximum storage capacity of 256 TB per blade with a maximum of four blades. They support the following NAS protocol options: NFS, CIFS, and pNFS. Each

blade has redundant power supplies located in the front of each blade. The CPU module is also accessible via the front of each blade after removing the power supplies. The rear of each blade houses the slots used for internal/external network connectivity, backend fibre connectivity, and frontend connectivity options.

This model uses a DME that is similar in appearance to the SPE. The DME can house two blades per 2U enclosure, and each blade on this model may accept a maximum of three I/O modules for frontend connectivity.

Figure 29, Figure 30, and Figure 31 show the front and back of the DME.



Figure 29. Front of a DME



Figure 30. Back of a DME




Figure 31. Close-up of the back of a DME

VNX7500

The VNX7500 is designed for the enterprise space. This model provides either block and file services, file only, or block only and uses a SPE form factor.

This model uses a 2.8 GHz, six-core Xeon 5600 processor with 24 or 48 GB RAM per blade and a maximum of 1,000 drives with the following host connectivity options: FC, iSCSI, and FCoE. There is also an optional upgrade to 48 GB of RAM if you have purchased a VNX or have an existing system with 24 GB of RAM per SP.

The VNX7500 has a SPE that uses UltraFlex I/O slots for all connectivity. The first slot houses the internal network management switch, which includes a mini-serial port and service LAN port. EMC service personnel use these ports to connect to the VNX.

A LAN connection is provided on each SP for array management. Each SP in the enclosure also has two power supply modules. Each SP has a maximum of five I/O module slots. Any slots without I/O modules will be populated with blanks (to ensure proper airflow).

The fault/status light of the SPE form factor is located on the front of the unit. This form factor requires an additional DAE to provide a minimum of four drives for the VNX Array Operating Environment. This model's form factor favors scalability, rather than density, requiring slightly more rack space for a minimal configuration.

The front of the SPE houses the power supplies, cooling fans, and CPU modules, along with the fault/status LED.

The blades in this model use a 2.8 GHz, six-core Xeon 5600 processor with 24 GB RAM per blade, with a maximum storage capacity of 256 TB per blade with a maximum of eight blades, and supports the following NAS protocol options: NFS, CIFS, and pNFS.

Each blade has redundant power supplies located in the front of each blade. The CPU module is also accessible via the front of each blade after removing the power supplies. The rear of each blade houses the slots used for internal/external network connectivity, backend fibre connectivity, and frontend connectivity options.

This model uses a DME that appears similar in appearance to the SPE of the VNX5700 and VNX7500. The DME can house two blades per 2U enclosure, and each blade on this model may accept a maximum of four I/O modules for frontend connectivity.

I/O modules

I/O modules for the storage processor

VNX series arrays support a variety of UltraFlex I/O modules on the storage processors, which are discussed in this section.



Figure 32. The quad-port 8 Gb/s FC optical I/O module

The quad-port 8 Gb/s FC optical I/O module:

- Auto-negotiates 2 Gb/s, 4 Gb/s, or 8 Gb/s
- Is supported on storage processors, except in the VNX5100
- Is used for frontend connectivity from storage processors to hosts
- Is used for backend connectivity from blades to storage processors on the VNX array




Figure 33. The quad-port 1 Gbase-T iSCSI/TOE module

The quad-port 1 Gbase-T iSCSI/TOE module:

- Is supported on storage processors, except in the VNX5100
- Is used for frontend connectivity to hosts
- Operates at 1 Gb/s
- Provides four copper ports




Figure 34. The dual-port 10 Gb/s Ethernet/FCoE module

The dual-port 10 Gb/s Ethernet/FCoE module:

- Is supported on storage processors
- Is used for frontend connectivity to hosts
- Operates at 10 Gb/s
- Supports SFP+ and Active TwinAx cables only (Passive TwinAx cables not supported)




Figure 35. The dual-port 10 GbE iSCSI (optical/TwinAx) I/O module

The dual-port 10 GbE iSCSI (optical/TwinAx) I/O module:

- Is supported on storage processors, except in the VNX5100
- Is used for frontend connectivity to hosts
- Operates at 10 Gb/s
- Provides two optical ports




Figure 36. The quad-port 6 Gb x 4 lane SAS I/O module

The quad-port 6 Gb x 4 lane SAS I/O module:

- Is supported on VNX5500, VNX5700 and VNX7500 storage processors for backend connections to DAEs
- Provides four lanes of 6 Gb/s SAS per port



Figure 377. The dual-port 10G Base -T copper I/O module

The dual-port 10G Base-T iSCSI copper I/O module:

- Is supported on VNX5300, VNX5500, VNX5700 and VNX7500 storage processors for backend connection to DAEs
- Operates at 1G/10G auto
- Provides two copper ports

I/O modules for the Data Mover

VNX series arrays support a variety of UltraFlex I/O modules on the blades, which are discussed in this section.




Figure 38. The quad-port 1 Gbase-T IP module

The quad-port 1 Gbase-T IP module:

- Enables connectivity to NAS clients
- Operates at 10/100/1 Gb/s
- Provides four copper ports




Figure 39. The dual-port 1 Gbase-T plus dual-port 1 GbE optical IP module

The dual-port 1 Gbase-T plus dual-port 1 GbE optical IP module:

- Enables connectivity to NAS clients
- Operates at 10/100/1 Gb/s copper and 10 Gb/s optical
- Provides two copper and two optical connectors



Figure 40. The dual-port 10G Base-T copper I/O module Gen 2

The dual-port 10G Base-T copper I/O module:

- Is supported on blades
- Enables connectivity to NAS clients
- Operates at 10 Mb/100 Mb/1G/10G auto
- Provides two copper ports




Figure 41. The dual-port 10G optical I/O module

The dual-port 10G optical I/O module:

- Is supported on blades
- Enabled optical connectivity to NAS clients
- Operates at 10G only
- Provides two optical ports

Table 2. Ultra-Flex I/O module compatibility chart

UltraFlex I/O module	Supported on Blades	Supported on storage processors
Quad-port 8 Gb FC optical	Yes	Yes
Quad-port 1G Base-T IP	Yes	No
Dual-port 1G Base-T plus dual-port 1GbE optical IP	Yes	No
Quad-port 1G Base-T iSCSI/TOE	No	Yes
Dual-port 10 Gb Ethernet/FCoE	No	Yes
Dual-port 10 GbE iSCSI (optical/TwinAx) ¹	No	Yes
Quad-port 6 Gb x 4 lane SAS ³	No	Yes
Dual-port 10G Base-T (copper)	Yes	No
Dual-port 10G Base-T iSCSI	No	Yes
Dual-port 10Gb Ethernet (optical/twin-ax)	Yes	No

¹ Model numbers: VSPMXGI2A (SFP)/VSPMXGI2TWA (Active TwinAx)

² Model numbers: VDMMXG2OPA (SFP)/VDMMXG2TWA (Active TwinAx)


³ Used for backend connectivity only; not supported for host connectivity. VNX5300 does not support additional SAS modules

Control Station

The Control Station is a 1U self-contained server, and is used in unified and file configurations. The Control Station provides administrative access to the blades; it also monitors the blades and facilitates failover in the event of an blade runtime issue.

The Control Station provides network communication to each storage processor. It uses Proxy ARP to enable communication over the Control Station management Ethernet port to each storage processor's own IP address.

An optional secondary Control Station is available that acts as a standby unit to provide redundancy for the primary Control Station.



1	AC power in connector	7	DB-9 serial console plug connector
2	PS/2 connector (mouse)—not used	8	Two USB 2.0 connectors
3	DB-9 serial modem plug connector	9	RJ-45 Ethernet NIC port (labeled A)
4	RJ-45 Ethernet NIC port (labeled CS ^a)	10	POST diagnostic LEDs ^b
5	RJ-45 Ethernet NIC port (labeled B)	11	DB-15 Video (VGA) socket connector—not used
6	RJ-45 Ethernet NIC port (labeled MGMT)	12	PS/2 connector (keyboard)—not used

Figure 40. Control Station

The RJ-45 Ethernet NIC port (number 5 in Figure 40) uses an Intelligent Platform Management Interface (IPMI) cable to connect to an optional secondary Control Station. This IPMI port is labeled CS on the Control Station.

Standby power supply

The standby power supply is a 1U unit that provides battery power to the DPE in the VNX5100, VNX5300, and VNX5500. It provides battery power to the DAE (bus 0, enclosure 0) and SPE and in a VNX5700 and VNX7500. This battery power allows the storage processors to de-stage data in-flight to the vault area of the reserved space in the event of a power failure. Once power is restored to the array, any writes that were de-staged are reconciled and persisted to the target backend disks to ensure that no data is lost. When using the 60-Drive DAE to hold the vault drives, the SPS is 2U 2.2KW.




Figure 41. Back of an SPS

Software on the VNX series

Reserved area on VNX series storage systems

The first four drives in enclosure 0 bus 0 (disks 0 through 3 in either the DPE for VNX5100/5300/5500 or the first DAE in VNX5700/7500, also known as *The Vault*) house the array's operating environment, boot image, and file control LUNs. This reserved area consumes 192 GB per disk and provides the landing area for data in-flight that is de-staged from the cache in the event of a power failure. NL-SAS drives can be used for vault drives for the VNX5100/5300/5500 only.

Unified management

EMC Unisphere provides a flexible, integrated experience for managing existing VNX storage systems.

Unisphere provides simplicity, flexibility, and automation—all key requirements for optimal storage management. Unisphere's ease of use is reflected in its intuitive task-based controls, customizable dashboards, and single-click access to real-time support tools and online customer communities. Unisphere's wizards help you provision and manage your storage while automatically implementing best practices for your configuration.

Unisphere is completely web-enabled for remote management of your storage environment. Unisphere Management Server runs on the SPs and the Control Station; it can be launched by pointing the browser to the IP address of either SP or the Control Station.

Unisphere has all the existing features and functionality of the previous interfaces, such as VMware awareness, LDAP integration, Analyzer, and Quality of Service Manager. Unisphere also adds many new features like dashboards, systems dashboard, task-based navigation and online support tools. For more information on Unisphere, refer to the white paper, [EMC Unisphere: Unified Storage Management Solution](#).

Figure 42 displays the new system dashboard. You can customize the view-blocks (also referred to as *panels*) in the dashboard.




Figure 42. Unisphere dashboard has customizable panels (view-blocks)

Software for efficiency

Fully Automated Storage Tiering for Virtual Pools (FAST VP)

EMC leverages FAST VP to migrate data to high-performance drives or high-capacity drives, depending on end-user access. Customers require fewer drives and receive the best ROI from those that are configured.




Figure 43. The FAST VP process

FAST VP functionality has been enhanced for the VNX series. FAST VP, previously available for block data and also known as sub-LUN tiering, is now available for file data, too. FAST VP optimizes storage utilization by automatically moving data between

storage tiers (for example, EFD, SAS, and NL-SAS). With the 7.1 release, data also moves within tiers.

The data is moved in 1 GB chunks based on the access patterns of the I/O. The storage pools on the left in Figure 43 show the initial storage configuration. After implementing FAST VP (shown on the right), the system proactively optimizes the storage pool by moving the 1 GB chunks of sub-LUN data to the most effective drive. This ensures that the appropriate data is housed on the right tier at the right time, which significantly increases efficiency and performance.

You can set policies and schedules to help determine how and when the data is moved. The [EMC FAST VP for Unified Storage Systems](#) white paper provides more information.

FAST Cache




Figure 44. FAST Cache process

FAST Cache remains an industry-leading feature for EMC storage arrays. The VNX series uses this feature to extend the array's read-write cache and ensure that unpredictable I/O spikes are serviced at flash speeds, which benefits all applications. FAST Cache

perfectly complements FAST VP, as it works at a more granular level by copying 64 KB chunks onto flash drives reserved for FAST Cache, depending upon the I/O characteristics.

Repeated access of the same 64 KB chunk of data causes the policy engine to promote that data to FAST Cache. FAST Cache also works in quicker cycles than FAST VP; it has a response time in the order of milliseconds to microseconds. FAST Cache reacts to I/O spikes and maintains access levels by acting as an extension to onboard memory.

FAST Cache is most appropriate for workloads with a high locality of reference, for example, applications that access a small area of storage with very high frequency, such as database indices and reference tables. FAST Cache is not beneficial for very large I/O streams that are sequential, such as backups, because each 64 KB chunk is accessed only once. SP Cache, however, uses various algorithms to optimize sequential I/O. FAST Cache is most useful for handling spikes in I/O, and it benefits applications that may have unpredictable I/O profiles. For more information, refer to the [EMC CLARiiON, Celerra Unified, and VNX FAST Cache – A Detailed Review](#) white paper.

The VNX5100 cannot use Thin Provisioning™ if FAST Cache is enabled, and vice versa. However, Thin Provisioning and FAST Cache usage at the same time is available in all of the other VNX models.

Data deduplication




Figure 45. Data deduplication

Data deduplication is a file-side feature that includes file-level compression and file-level single instancing. This asynchronous feature works in the background, scanning for inactive files that may contain duplicate data. If there is more than one instance of a file, it will be single-instanced, along with being compressed.

This feature helps increase the storage efficiency of file storage by eliminating redundant data within file systems, thereby reducing storage costs. It has user-configurable settings to filter certain files from being processed. It may be turned on or off per file system basis at any time. For more information, refer to the [EMC Data Compression—A Detailed Review](#) white paper.

Thin Provisioning




Figure 46. Thin Provisioning

EMC's Thin Provisioning feature is easy to set up and monitor. It allows storage administrators to allocate storage on demand. It presents a host with the total amount of storage that has been requested; however, it only allocates storage on the array that is actually being used. For example, on the file side, a 100 GB file system that has been thin provisioned will be seen by hosts and users as 100 GB. If only 50 percent of the file system is actually in use (contains data), only 50 GB will be used on the array. This feature prevents overprovisioning storage that is not used.


Thin provisioning increases storage utilization, bringing down the costs associated with maintaining the storage (such as power and cooling), and lowering acquisition costs of new storage hardware. When using this on the block side, an administrator creates thin LUNs as opposed to thick LUNs. Thin LUNs will allocate 3 GB at LUN creation; however, this will not reserve the space necessary for the entire size of the newly created LUN. Additional space is assigned 1 GB at a time; however, only 8 KB chunks are reserved as needed.

Administrators should understand the growth rate of their pools and file systems to know what percentage is practical and allow enough time to react and address potential oversubscription issues. For more information on Thin Provisioning, refer to the [EMC CLARiiON Thin Provisioning – Applied Technology](#) white paper.

The VNX5100 cannot use Thin Provisioning™ if FAST Cache is enabled, and vice versa. However, Thin Provisioning and FAST Cache usage at the same time is available in all of the other VNX models.

Software for protection

Replicator



EMC VNX Replicator is an asynchronous file-based replication solution. This feature is used for file-system-level replication, and provides point-in-time views of the source file systems on the replica file systems. VNX Replicator uses an adaptive scheduling technology that responds to changing network conditions by monitoring two key pieces of information:

- The amount of change (delta) that has built up in the production object since the last transfer
- The average transfer rate that was observed during the last transfer

This adaptive-scheduling technology dynamically allocates resources to ensure that objects are replicated in compliance with their customer-configurable recovery point objectives (RPOs). An RPO is a measurement of how much data may be lost before it negatively impacts the business. Each file system may have a different RPO as the importance of each file system's contents may differ.

VNX Replicator also allows bandwidth throttling based on a schedule, e.g., change the throughput during specific times throughout the week. This is very useful in environments where VNX Replicator shares network resources with other applications.

RecoverPoint integration




Figure 47. How RecoverPoint works

Unisphere has been designed to accept plug-ins that extend its management capabilities. VNX series arrays are managed with the latest version of Unisphere that includes RecoverPoint integration for replication of block and file data. RecoverPoint provides a synchronous/asynchronous replication solution to allow file and block cabinet-level business continuation on a remote VNX series array in the event of a disaster.

This support includes active/passive, as well as active/active, using continuous remote replication (CRR). Because RecoverPoint uses a single consistency group to provide array-level replication for files and does not provide point-in-time views, it is most appropriate for control LUNs and critical file systems. It is, therefore, still recommended that you use EMC VNX Replicator for file system-level replication.

MirrorView/A




Figure 48. How MirrorView/A works

EMC VNX MirrorView™/A is a block feature for remote mirroring that works asynchronously to periodically copy information on a LUN from one storage system to another. As an asynchronous process, it provides replication over long distances in the order of hundreds to thousands of miles and has an RPO from 30 minutes to hours. MirrorView/A is optimized for low network bandwidth.

MirrorView/S




Figure 49. How MirrorView/S works

EMC VNX MirrorView/S is a block feature for remote mirroring that works synchronously to copy information on a LUN from one storage system to another. The synchronous nature of this feature means that for every write to a LUN on the primary storage system, the same write is copied to the secondary storage system before the write acknowledgement is sent to the host.

With Dense Wavelength Division Multiplexing (DWDM), MirrorView/S provides replication over distances of up to 200 km. Without DWDM, MirrorView/S provides replication over distances of up to 60 km. In each case, the RPO is zero seconds, because both copies of data are identical.

SAN Copy

EMC VNX SAN Copy™ is a block feature that copies data between EMC and qualified third-party storage systems. SAN Copy copies data directly from a source LUN on one storage system to destination LUNs on other systems without using any host resources. SAN Copy may be used to create full and incremental copies of a source LUN.

Business continuance

VNX Snapshots





VNX Snapshots is a new feature created to improve snapshot capability for VNX Block. VNX Snapshots are point-in-time views of a LUN, which can be made accessible to another host, or be held as a copy for possible restoration. VNX Snapshots use a redirect-on-write algorithm, and are limited to pool based provisioned LUNs (i.e. not RAID Group LUNs). VNX Snapshots support 256 writable snaps per pool LUN. Branching, or *snap of a snap*, is also supported. There are no restrictions to the number of branches, as long as the entire Snapshot Family is within 256 members. Consistency Groups are also introduced, meaning that several pool LUNs can be combined into a Consistency Group and snapped at the same time. For more information see the White paper on VNX Snapshots on powerlink.

SnapView Snapshots

EMC VNX SnapView snapshots are point-in-time views of a LUN, which can be made accessible to another host, or be held as a copy for possible restoration. SnapView

snapshots use a pointer-and-copy-based algorithm. A memory map keeps track of *chunks* (blocks) of data.

Before chunks of data are written to the source LUN, they are copied to a reserved area in private space, and the memory map is updated with the new location of these chunks. This process is referred to as *Copy on First Write*. The source LUN, reserved LUN, and memory map work together to create the snapshot.



Checkpoints

EMC VNX checkpoints are point-in-time snapshots for file systems. This feature allows creation of multiple non-disaster recoverable copies of production data. Checkpoints enable end users to restore their files by integrating with Microsoft Shadow Copy Services. UNIX and Windows clients may also access checkpoints to restore individual files through a hidden directory.

The Unisphere GUI gives the storage administrator a quick and easy way to restore entire file systems from checkpoints. SnapSure™ is the underlying technology that supports checkpoints. SnapSure uses a very storage-efficient Copy-on-First-Write algorithm to maintain checkpoint views.

Clones

SnapView clones are fully populated point-in-time copies of LUNs that allow incremental synchronization between source and destination LUNs. Unlike snapshots that provide point-in-time views of data, clones provide fully populated point-in-time copies that maximize the flexibility of the storage environment. These point-in-time copies allow you to perform additional storage management tasks with minimal impact to production data. These tasks include backup/recovery, application testing, warehousing, and data movement.

Security

Anti-virus software

The VNX Event Enabler (VEE) allows VNX to integrate with industry-leading anti-virus applications, such as Symantec, McAfee, Computer Associates, Sophos, Kaspersky, and Trend Micro.

The Anti-Virus feature works with an external system that houses one of the supported anti-virus applications. Once configured, files on a file system are scanned before they are delivered to the requesting client. The file is either validated and sent to the client or denied if any issues are identified during the scan. A scan is also made after a file is modified and closed. If no modifications are performed on the file, it is not rescanned when closed. This can be configured to either scan files on write (default) or scan files on read. There is the capability to scan all files in the file system with a separate command as well. This feature helps protect files from viruses and malware while still allowing you to benefit from the advanced features of the VNX array, such as high availability and performance.

Quota management and auditing

The VEE allows the VNX to integrate seamlessly with industry-leading content/quota management software applications (CQM applications). This enhanced functionality extends the built-in management capability to help improve content and quota management along with auditing on file systems. This feature is comprised of a host

framework and an agent that have been validated to support several of the leading CQM application vendors.

File-Level Retention

The File-Level Retention feature protects files from deletion or modification until a user-specified retention date elapses. This feature prevents users from deleting or modifying files that are locked and protected. There are two levels of FLR protection that are enabled at the file system level upon creation:

- **FLR-C File-Level Retention (Compliance level)**—Protects data from changes made by users through CIFS, NFS, and FTP, including administrator actions. This level of FLR also meets the requirements of SEC Rule 17a-4(f). A file system that contains files that are locked with FLR-C cannot be deleted.
- **FLR-E File-Level Retention (Enterprise level)**—Protects data from changes made by users through CIFS, NFS, and FTP, but not including administrator actions. This means that only a VNX administrator with the appropriate authorization can delete an FLR-E file system that contains FLR-E protected files. However, administrators still cannot target individually locked files for deletion or modification.

Figure 50 illustrates the workflow for File-Level Retention.




Figure 50. Workflow for FLR

Advanced Data Services Support for file

This feature allows file to use pool-based LUNs, create file objects, and use FAST VP. When a file object (such as a file system) is created upon a block object (such as a LUN), configuring and changing the LUN could adversely affect the file system it supports. To prevent this from happening, Unisphere has the Advanced Data Service Support for file feature.

This feature allows the VNX system to take into account the relationship between these two sides (file and block) when certain functions are performed. The system warns the user, and even prevents some actions when these actions could cause undesired results. For example, VNX displays an error and prevents the action when the end user tries to change the mirror status of a storage-pool LUN (from mirrored to non-mirrored or vice versa), and this change would result in a pool having mirrored and non-mirrored LUNs.

VNX displays a warning (but still allows the action) when:

- LUNs underneath the same file system are spanning across multiple AVM pools.
- A new storage pool is discovered that contains mirrored and non-mirrored LUNs. Only one type of LUN will be mapped over to file; the rest will be skipped.
- A LUN is added to the “~filestorage” storage group from a pool with “mixed” data services. Data services include:
 - Tiering Policy (Start High then Auto-Tier/Auto-Tier/Highest/Lowest/No Movement)
 - Thick/thin LUN
 - Compression

A warning is displayed, because the system cannot guarantee file system performance when supported by LUNs with different data services. Although not guaranteed, the system uses a “best effort” algorithm to find LUNs in the pool that have matching data services during file system creation and extension. The ability for the system to find LUNs with matching data services is dependent on LUN availability.

Serviceability, availability, and performance

File-side control LUN protection

To ensure the safety and stability of the unified system, the control LUNs for file have been moved to the VNX Private Space. Control LUNs are added automatically to a special system storage group named “~filestorage.” This storage group is visible in the Unisphere GUI. The system protects these LUNs by disallowing them from being deleted or removed from the “~filestorage” group in the GUI. As further protection, navisecli commands cannot be used to manipulate these Control LUNs.

Unified software upgrade

The VNX series provides software upgrades as a single process utilizing Unisphere Service Manager. The Celerra Upgrade Tool and CLARiiON Software Assistant have been merged in to one utility that allows you to download updates, perform health checks, and install the new updates for both block and file. This is shown in Figure 51.




Figure 51. The System tab in Unisphere Service Manager for handling updates

Unified Remote Support

The VNX series fully supports Unified Remote Support by leveraging EMC Secure Remote Services (EMC Secure Remote Support Gateway). This secure IP-based system allows dial home, as well as remote access, to the block and file components. ESRS allows organizations complete control over policies that dictate who may access their VNX and when that access may occur. ESRS is a separate offering from EMC that is supported by the VNX series.

Block to Unified Plug and Play

Block to Unified adds file services to a block-only deployment. The upgrade will include installing the necessary hardware to enable file services, install the necessary operating systems (DART/FLARE), and configure file services within your environment.

Assumptions

- DART/FLARE compatible revisions are loaded on the private LUNs as a precondition to the B2U.
- The upgrade will provide uninterruptable block data I/O during the file services upgrade and uninterruptable management access.
- Block data services, MirrorView, SnapView and replication via Recover Point will not be interrupted.
- Dual control station configurations are supported as part of the block to file upgrade.

User Experience

The Block to Unified(B2U) upgrade is divided into four phases:

1. Readiness Check - the customer would use a tool to verify and prepare the Block system for the Unified upgrade.
2. Rack & Cable—The customer would rack and cable the newly acquired file hardware according to the provided instructions and then power-on the hardware.
3. Plug-n-Play—The customer would monitor the automated File software installation and set up the Unified system to be ready for configuration.
4. Configure—The customer would configure the system for production.

In addition, the installation and recovery for File/Unified systems is enhanced to improve TCE by reusing the technology developed for the Block to Unified upgrade.

In-Family Data in place Conversion

The In-Family Upgrade/Conversion features allow the customer to preserve system identity while moving to a higher model or upgrading memory of an existing model.

The Upgrade Readiness Checker is used before an in-family upgrade can be performed.

Source Platform (Model/Memory)	VNX5100/ 4GB	VNX5300/ 8GB	VNX5500/ 12GB	VNX5700/ 18GB	VNX7500/ 24GB ^[1]	VNX7500/ 48GB
VNX5100/4GB	N/A	Yes	Yes	Yes ^[2]	N/A	Yes ^[2]
VNX5300/8GB	N/A	N/A	Yes	Yes ^[2]	N/A	Yes ^[2]
VNX5500/12GB	N/A	N/A	N/A	Yes ^[2]	N/A	Yes ^[2]
VNX5700/18GB	N/A	N/A	N/A	N/A	N/A	Yes
VNX7500/24GB	N/A	N/A	N/A	N/A	N/A	Yes ^[3]
VNX7500/48 GB	N/A	N/A	N/A	N/A	N/A	N/A

Figure 52: In-Family Data in place Conversion options

- 1 The VNX7500/24GB is an invalid target. The VNX7500/48GB will be the standard target
- 2 Several conversion paths are chassis swap conversions which involve more hardware and process steps
- 3 This path is considered a memory upgrade, though it uses the in-family conversion mechanism.

Source	Target			
	VNX5300 Block 8GB	VNX5500 Block 12GB	VNX5700 Block 18GB	VNX7500 Block 48GB
	VNX5100 Block 4GB	SP Swap	SP Swap	Chassis Swap
	VNX5300 Block 8GB	N/A	SP Swap	Chassis Swap
	VNX5500 Block 12GB	N/A	Chassis Swap	Chassis Swap
	VNX5700 Block 18GB	N/A	N/A	SP Swap
	VNX7500 Block 24GB	N/A	N/A	Memory Add
	VNX7500 Block 48GB	N/A	N/A	N/A

Figure 53: Block In-Family Conversions

Source	Target			
	VNX5300 File/Unified 8GB	VNX5500 File/Unified 12GB	VNX5700 File/Unified 18GB	VNX7500 File/Unified 48GB
	VNX5100 File/Unified 4GB	N/A	N/A	N/A
	VNX5300 File/Unified 8GB	SP Swap Blade Swap	Chassis Swap Blade Swap	Chassis Swap Blade Swap
	VNX5500 File/Unified 12GB	N/A	Chassis Swap Blade Swap	Chassis Swap Blade Swap
	VNX5700 File/Unified 18GB	N/A	N/A	SP Swap Blade Swap
	VNX7500 File/Unified 24GB	N/A	N/A	Memory Add No Blade Swap
	VNX7500 File/Unified 48GB	N/A	N/A	N/A

Figure 54: Unified or File In-Family Conversions

VNX scalability and performance

Improved overall performance

The EMC VNX series array is highly scalable and provides significantly more performance than previous array models, such as the CX4 and NS. The VNX systems are optimized to take advantage of the multi-core Intel Xeon 5600 processors. In addition, other hardware improvements, such as the 6 Gb/s x 4 lanes SAS back end, PCIe Gen 2 I/O modules, and increased memory, make the VNX series two to three times faster than previous-generation CLARiiON and Celerra systems.

Increased degraded system performance

The VNX series systems show significantly improved performance when running on a single SP (also known as *degraded mode*). All VNX models perform better compared to its predecessors (CX4 or NS systems).

For example, a CX4-960 populated with 240 active drives running IOPS-intensive workloads would suffer a significant performance hit in the event of a single SP failure. However, a VNX7500 (with a similar configuration of 250 drives, that is, running IOPS-intensive workloads) does not suffer anywhere near the performance hit during a single-SP failure; in fact, it provides *more* performance than the CX4 running with *both* storage processors. This increase in robustness is directly attributable to the hardware advances within the array.

Optimized for flash

VNX series systems also take much better advantage of flash drives than previous-generation systems. This is because of the efficient use of the multi-core processors and the more powerful storage processors in the VNX series. VNX series arrays have two to three times more IOPS using flash drives with Exchange and file-serving workloads than previous-generation models.

Conclusion

The VNX series offers increased scalability, efficiency, and performance in a unified form factor that you can manage using a single pane of glass, offering an unparalleled unified experience. The advances in hardware, from the Intel Xeon 5600 processors, to the 6 Gb/s x 4 lanes SAS backend, to the optimizations made using Flash drives, place the VNX series in the forefront of EMC's midrange offerings.

The unification of file and block is another step in the continued evolution of our product lines. By unifying the hardware and creating a new unified management experience (through the conceptual merger of file and block inside Unisphere), EMC has provided a holistic way to view and manage an entire storage domain with the utmost simplicity.

Appendix A: System level, block, and file components

Table 3. System level, block, and file components

Components Configuration	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
SYSTEM LEVEL COMPONENTS					
Max FAST Cache memory (GB)	100	500 ¹	1000	1500 ¹	2100 ¹
Max FAST Cache drives	2	10	20	30	42
Max drives per array	75	125	250	500	1000
Min drives per storage system	4	4	4	4	4
Minimum configuration rack space	4U	7U ²	7U ²	8U ^{2, 3}	8U ^{2, 3}
3.5"–100 GB/200 GB Flash drives	Yes	Yes	Yes	Yes	Yes
2.5 –100 GB/200 GB Flash drives	Yes	Yes	Yes	Yes	Yes
3.5"–300 GB 15k & 10k drives	Yes	Yes	Yes	Yes	Yes
3.5"–600 GB 15k & 10k drives	Yes	Yes	Yes	Yes	Yes
3.5"–900 GB 10k drives	Yes	Yes	Yes	Yes	Yes
3.5"–1 TB 7.2k drives	Yes	Yes	Yes	Yes	Yes
3.5"–2 TB 7.2k drives	Yes	Yes	Yes	Yes	Yes
3.5"–3 TB 7.2k drives	Yes	Yes	Yes	Yes	Yes
2.5"–300 GB 10k drives	Yes	Yes	Yes	Yes	Yes
2.5"–600 GB 10k	Yes	Yes	Yes	Yes	Yes

Components Configuration	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
drives					
2.5"–900 GB 10k drives	Yes	Yes	Yes	Yes	Yes
BLOCK COMPONENTS					
Processor clock speed/number of cores/architecture (per SP)	1.6 GHz, 2-core Xeon 5600	1.6 GHz, 4-core Xeon 5600	2.13 GHz, 4-core Xeon 5600	2.4 GHz, 4-core Xeon 5600	2.8 GHz, 6-core Xeon 5600
Physical memory per SP (GB)	4	8	12	18	24 or 48
Max FC ports per SP (FE ports)—onboard only	4	4	4	0	0
Max FC ports per SP (FE ports)—I/O module add-On	0	4	4	12	16
Max 1 Gb/s iSCSI ports per SP (FE only)	0	4	8	8	8
Max 10 Gb/s iSCSI ports per SP (FE only)	0	4	4	6	6
Max FCoE ports per SP (FE only)	0	4	4	6	8
Max initiators per FC port	256	512	512	512	1024
Max initiators per 1 Gb/s iSCSI port	0	512	512	512	512
Max initiators per 10 Gb/s iSCSI port	0	512	1024	2048	2048
Max initiators per FCoE port	0	512	512	512	1024
Max VLANs per 1 Gb/s iSCSI port	0	8	8	8	8
Max VLANs per 10 Gb/s iSCSI port	0	8	8	8	8

Components Configuration	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Max initiators per SP	256	1024	2048	2048	4096
Max FE ports per SP (all types)	4	12	12	16	16
Max initiators per array	512	2048	4096	4096	8192
Max UltraFlex I/O module slots (per SP)	0	2	2	5	5
Max usable UltraFlex I/O module slots (front-end) (per SP)	Please see Table 5. Block configuration on page 73				
Embedded I/O ports per (per SP)	4 FC 2 back-end SAS	4 FC 2 back-end SAS	4 FC 2 back-end SAS	0	0
Available 6 Gb SAS ports for DAE connectivity	2	2	2 or 6	4	4 or 8
Max LUNs per storage system	512	2048	4096	4096	8192
Max user-visible LUNs ⁵	896	2816	4864	5376	10496
LUN ID range	0-1023	0-4095	0-8191	0-8191	0-24575
Max LUNs per RAID group	256	256	256	256	256
Max LUNs per storage group	256	512	512	1024	1024
RAID options	0,1,1/0,3,5,6	0,1,1/0,3,5,6	0,1,1/0,3,5,6	0,1,1/0,3,5,6	0,1,1/0,3,5,6
Max RAID groups per storage system	75	125	250	500	1000
Max drives per RAID group	16	16	16	16	16
DPE/SPE rack space	3U (DPE)	3U (DPE)	3U (DPE)	2U (SPE)	2U (SPE)
Max SnapView sessions(LUNs) per source LUN	8	8	8	8	8
Max SnapView snapshot source	128	256	256	512	512

Components Configuration	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
LUNs					
Max SnapView snapshot sessions per array	256	512	512	1024	2048
Max SnapView snapshot source LUNs in a Consistent Start	32	32	32	64	64
Max clones	128	256	512	1024	2048
Max clone source LUNs	64	128	256	512	1024
Max clones in a Consistent Facture	32	32	32	64	64
MirrorView/S objects	128	128	256	512	1024
MirrorView/A objects	128	256	256	256	256
SAN Copy	Yes	Yes	Yes	Yes	Yes
Max VNX Snapshots per source	256	256	256	256	256
Max VNX Snapshots per array	2048	4096	8192	16384	32768
Max VNX Snapshot consistency groups	128	128	128	256	256
Max source LUNs in consistency group	64	64	64	64	64
Max VNX Snapshot LUNs (Host Visible)	128	512	1024	2048	2048 4096 w/ 48G memory
FILE COMPONENTS					
Processor clock speed / number of cores / architecture (per blade)	N/A	2.13 GHz, 4-core Xeon 5600	2.13 GHz, 4-core Xeon 5600	2.4 GHz, 4-core Xeon 5600	2.8 GHz, 6-core Xeon 5600

Components Configuration	VNX5100	VNX5300	VNX5500	VNX5700	VNX7500
Physical memory per blade (GB)	N/A	6	12	12	24
Data Movers (min-max)	N/A	1 -2	1 -3	2 -4	2 - 8
Capacity per blade (TB)	N/A	200 ⁷	256	256	256
Max # of file systems per DM ⁶	N/A	2048	2048	2048	2048
Max # of file systems per array	N/A	4096	4096	4096	4096
Max file system size (TB)	N/A	16	16	16	16
Number of snaps per file system (SnapSure)	N/A	96	96	96	96
Number of replication file systems (Replicator)	N/A	1024	1024	1024	1024
Host connectivity	N/A	NFS, CIFS, pNFS	NFS, CIFS, pNFS	NFS, CIFS, pNFS	NFS, CIFS, pNFS
Control Stations	N/A	1- 2	1- 2	1- 2	1- 2
Max usable UltraFlex I/O module slots (front-end)	N/A	2	3	3	4
(per blade) ⁴	N/A	1,1/0,5,6	1,1/0,5,6	1,1/0,5,6	1,1/0,5,6
RAID options					

¹ Attainable using 100 GB Flash drives

² Single Control Station configuration

³ Using 2.5" 25-drive DAE

⁴ Each blade requires one I/O module slot for backend (fibre) connectivity. This has been deducted from the numbers shown in the chart.

⁵ These numbers include FLUs + pool LUNs + snapshot LUNs (SV and Adv) + MV/A LUNs

⁶ These numbers are without any snaps or replications

⁷ Limited by the number of total drives per array multiplied by the per drive capacity

Appendix B: Storage processor UltraFlex I/O slot availability

The following table shows the number of various components in a Unified configuration.

Table 4. Unified configuration

	VNX5300	VNX5500		VNX5700	VNX7500			
SAS back-end buses	1–2	1–2	3–6	1–4	1–4		5–8	
Number of blades	1–2	1–3	1–3	2–4	2–4	5–8	2–4	5–8
Total number of configurable host I/O connectivity slots	2	2	1	3	3	2	2	1

Table 5 and Table 6 pertain to components delivered in an EMC rack. I/O module slots are only reserved for upgrade from block-only to unified if the array is delivered in an EMC rack. The VNX5100 does not support the use of UltraFlex I/O modules, so it has been omitted in these tables.

Table 5. Block configuration

	VNX5300		VNX5500		VNX5700		VNX7500			
SAS back-end buses	1–2		1–2	3–6	1–4		1–4		5–8	
Space reserved for upgrade to unified?*	N	Y	N	Y	N	Y	N	Y	N	Y
Total number of configurable host I/O connectivity slots	2		2	1	4	3	4	3	3	2

Table 6. File configuration

	VNX5300	VNX5500		VNX5700	VNX7500			
SAS back-end buses	1–2	1–2	3–6	1–4	1–4		5–8	
Number of blades	1–2	1–3	1–3	2–4	2–4	5–8	2–4	5–8
Total number of configurable host I/O connectivity slots*	1	1	1	1	1	1	1	1

There is only one configurable slot available to account for the RecoverPoint Fibre Channel option. As a file-only configuration, no other host connectivity is allowed from the storage processors.