

EMC VNXe SERIES STORAGE SYSTEMS

A Detailed Review

Abstract

This white paper introduces the architecture and functionality of EMC® VNXe™ storage systems, including the VNXe3100™ and VNXe3300™. This paper also discusses the software functionalities of these storage systems.

March 2012

Copyright © 2012 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 EMC Corporation Trademarks on EMC.com.

VMware, ESX, and VMware vCenter are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. All other trademarks used herein are the property of their respective owners. All other trademarks used herein are the property of their respective owners.

Part Number h8178.3

Table of Contents

Executive summary	5
Audience.....	5
Terminology	6
VNX unified storage family	7
VNXe series	8
Hardware overview	11
Software overview	15
Connection Utility and Configuration wizard	15
Unisphere	16
Storage provisioning	17
Exchange storage.....	18
Shared folders.....	19
Generic iSCSI storage	20
VMware storage	21
Hyper-V storage.....	21
Storage pool characteristics	22
Data protection	24
Snapshots.....	25
Replication.....	27
Backup and restore	28
CLI support.....	29
Advanced features	29
Link aggregation.....	29
Fail Safe Networking.....	30
Thin provisioning.....	30
File-Level Retention (enterprise)	31
Deduplication.....	32
Virus protection.....	32
User management	33
LDAP integration.....	33
High availability	33
Virtual integration with VMware vSphere	35
Advanced Storage Access (ASA)	35
vStorage API for Array Integration (VAAI).....	35
Remote monitoring	36
EMC Secure Remote Support (ESRS).....	36
Unisphere Remote.....	36

The VNXe series support ecosystem	37
Conclusion	37
References	38

Executive summary

Businesses continue to experience exponential growth in storage-capacity demand as well as require 24x7 availability. Research shows that companies are experiencing, on average, greater than 50 percent annual information growth. Most of this growth is happening with unstructured data, such as files.

Increased complexity is common as IT departments continue to add additional servers and applications, and move to a more virtualized environment. Solutions that help organizations better utilize server virtualization are compelling. Traditional storage management is often too complicated for businesses. Allocating and managing storage often requires an understanding of complicated storage terminology and specialized training for storage administrators. While the complexity of IT environments continues to increase, the service level demanded by IT users remains high. The increased cost of equipment, people, and facilities continues to be a challenge for IT generalists.

The EMC® VNXe™ series is a new unified storage solution for SMB and lower mid-market organizations and remote/branch offices (ROBO) in enterprise organizations. It addresses the challenges mentioned above. Designed for IT generalists with limited storage expertise, the VNXe series facilitates complete storage consolidation with advanced file and block functionality as well as a simple, application-driven approach to managing shared storage. The VNXe series delivers significant advancements in efficiency, simplicity, and affordability. Major benefits of VNXe series storage systems include:

- Support for file (CIFS and NFS) and block (iSCSI) storage
- Snapshots and replication (local and remote)
- Application-focused provisioning, and built-in best practices
- VMware and Hyper-V server integration
- A Unisphere™ interface for simplified management

This white paper provides an overview for VNXe series storage systems.

Audience

This white paper is intended for EMC employees, partners, IT planners, storage architects, administrators, and anyone involved in evaluating, acquiring, managing, operating, or designing an EMC networked storage environment using VNXe storage systems.

Terminology

- **Command line interface (CLI)** – An interface that allows you to do storage-related tasks using commands typed into a command line.
- **Common Internet File System (CIFS)** – An access protocol that allows users to access files and folders from Windows hosts located on a network.
- **Deduplication** – The process used to compress redundant data, allowing space to be saved on a file system. When multiple files have identical data, the file system stores only one copy of the data, and shares that data between the multiple files.
- **Link aggregation** – A high-availability feature that allows Ethernet ports that are associated with the same switch and have similar characteristics to combine into a single virtual device/link. The aggregated link may have a single IP address or multiple IP addresses.
- **Local replication** – Replication between two storage servers within the same VNXe system.
- **iSCSI protocol** – The Internet Small Computer System Interface (iSCSI) protocol provides a mechanism for accessing raw block-level data storage over network connections. The iSCSI protocol is based on a network-standard client/server model with iSCSI initiators (hosts) acting as storage clients and iSCSI targets acting as storage servers. Once a connection is established between an iSCSI host and the iSCSI server, the host can request storage resources and services from the server.
- **iSCSI server** – A VNXe server that uses the iSCSI protocol to manage Microsoft Exchange storage groups, generic storage virtual disks, Hyper-V datastores, and VMFS-based VMware® datastores.
- **Network File System (NFS)** – An access protocol that allows users to access files and folders from Linux/UNIX hosts located on a network.
- **Network Data Management Protocol (NDMP)** – A standard for backing up file servers on a network. It allows centralized applications to back up file systems running in a customer environment.
- **Remote replication** – The replication of stored data from one VNXe system to another (remote) VNXe system or VNX™/Celerra® system.
- **Server Message Block (SMB)** – The underlying protocol used by CIFS to request files, print, and communicate with a server over a network through TCP ports. VNXe storage systems support Microsoft's Server Message Block (SMB) 2.0 for CIFS shares. The SMB 2.0 protocol supports all characters from the Unicode 3.0 standard on systems running Microsoft Windows Server 2008 and Microsoft Windows 7.
- **Share** – A named, mountable instance of shared-folder storage, accessible through a shared folder or VMware VMFS datastore. Each share is accessible through the protocol (NFS or CIFS) defined for the shared folder where it resides.

- **Shared folder** – A VNXe storage resource that provides access to individual file systems for sharing files and folders. Shared folders contain either Windows shares (which transfer data according to the CIFS protocol) or NFS shares (which transfer data according to the NFS protocol). NFS shares are sometimes referred to as NFS exports.
- **Shared-folder server** – A VNXe server that uses either the CIFS or NFS protocol to catalog, organize, and transfer files within designated shares. A shared-folder server is required to create shared folders that contain CIFS or NFS shares, or NFS VMware datastores.
- **Snapshot** – A read-only, point-in-time copy of data stored on the storage system. You can recover files from snapshots or restore a storage element to a snapshot.
- **Storage pool** – A storage pool is a collection of disk drives configured with a particular storage profile. The storage profile defines the type of disks used to provide storage, as well as the type of RAID configured on the disks. The storage pool's configuration defines the number of disks and quantity of storage associated with the pool.
- **Storage processor (SP)** – A hardware component that provides the processing resources for performing storage operations such as creating, managing, and monitoring storage resources.
- **Unisphere** – A web-based management environment for creating storage resources, configuring and scheduling protection for stored data, and managing and monitoring other storage operations.
- **Virtual Provisioning™** – A VNXe storage feature used to allocate storage on an as-needed basis from a larger reserved resource. The portion of the total resource that the storage element initially uses is called its *initial allocation*.

VNX unified storage family

The VNX family is the next generation of EMC's midrange and low-end storage platforms. The VNX family (Figure 1) includes the VNX series and the VNXe series.




Figure 1. Models in the VNX family

The entry-level VNXe series includes two models: the VNXe3100™ and the VNXe3300™. These entry-level offerings provide storage over an IP network. They support the iSCSI, CIFS, and NFS protocols. The small footprint and affordable price of the VNXe make it well suited for small to medium businesses as well as remote-branch offices interested in consolidating storage of Exchange, VMware, Hyper-V, and file servers. VNXe is designed for the IT generalist with no specific storage knowledge. The newly designed management interface, Unisphere, is intuitive and enables application administrators to perform storage provisioning, management, and monitoring at the application level. Unisphere provides a centralized GUI that makes VNXe storage systems easy to manage, and it integrates with VMware and Hyper-V servers in the environment.

The VNX series is designed for midsize to enterprise environments. It begins with the VNX5100™, which is specialized with Fibre Channel-only front-end connectivity. The VNX5300™ through VNX7500™ can be configured for file, block, or unified (which includes file and block) environments. Optimized for virtual environments, the VNX models are simple, efficient, and powerful. The EMC VG2 and VG8 platforms should be used for those environments that require gateway access to a single or multiple Symmetrix®, CLARiiON®, or VNX arrays, or a combination thereof. For more details about the VNX series, refer to the [Introduction to EMC VNX Series Storage Systems](#) white paper available on EMC Powerlink®.

VNXe series

The VNXe series storage systems provide a consolidated platform for provisioning, managing, and monitoring application data storage within organizations and networks of all sizes. Each storage system provides highly efficient iSCSI block storage and highly accessible file-based storage for enabling users to work with files and data over network connections.

VNXe provides access to storage resources through the following protocols:

- **iSCSI** – Block-level data storage resources accessible through the iSCSI protocol over a network connection. Optional support for CHAP client/server authentication and iSNS discovery and management is also available.
- **Network-attached storage (NAS)** – File-based storage for a wide range of clients and applications that access storage over network connections. Protocol-specific file systems are located and managed on the storage system, which transfers data to hosts over TCP/IP using the CIFS and NFS file sharing protocols.

VNXe storage systems provide storage resources suited to the needs of specific applications, host operating systems, and user requirements. Unisphere can be used to create, configure, and monitor the following types of storage resources:

- **Microsoft Exchange** – Provides a resource for storing Microsoft Exchange databases, log files, and public folders based on simple parameters such as the number of users and the average user mailbox size.
- **Shared folder** – Allows hosts to store data and easily access shared folders and shares that integrate seamlessly into:
 - Windows environments that use the CIFS protocol for file sharing, Microsoft Active Directory for authentication, and Windows directory access for folder permissions
 - Linux/UNIX environments that use the NFS protocol for file sharing and POSIX access control lists for folder permissions
- **Generic iSCSI** – Provides generic block-level storage for hosts and applications that use the iSCSI protocol to access storage in the form of virtual disks.
- **VMware** – Provides storage for VMware virtual machines through datastores that are accessible through the NFS protocol or VMFS (over iSCSI) protocol.
- **Hyper-V** – Provides storage for Hyper-V virtual machines, including their virtual hard drives, configuration files, and snapshots, through datastores that are accessible to Windows Server 2008 hosts using the iSCSI protocol.

VNXe series storage systems support high-performance SAS drives and balanced-performance/capacity Near-Line SAS (NL-SAS) drives. Storage pools can be created from these drives to assign storage to hosts and applications.



Figure 2. Models in the VNXe series

The VNXe is an ideal platform for businesses with physical server infrastructures, as well as those making the move to server virtualization to drive consolidation and greater efficiency. The VNXe3100 is the entry-level VNXe storage system that provides high value for smaller user configurations and lighter performance needs. It is designed to support smaller businesses and applications with up to 300 users (dual SP configuration) or up to 150 users (single SP configuration). The VNXe3300 includes all the ease-of-use and application-driven management features of the VNXe3100, along with increased performance, scalability, and I/O expandability for use with larger configurations with higher performance demands. It is designed to support larger businesses and applications with up to 1,000 users. Both systems share a comprehensive set of features including increased capacity utilization, data protection, and availability solutions, and advanced support capabilities.

Some of the major benefits of VNXe series storage systems include:

- **Storage consolidation:** VNXe systems align storage management with applications by embedding best practices into the user interface and providing a faster, simpler user experience in completing everyday administrative tasks.
- **Simple and efficient, unified storage:** VNXe systems deliver unified IP storage for NAS and iSCSI environments. Application-centric management and provisioning wizards result in immediate familiarity for users, while integration of snapshot and replication with storage management workflows results in streamlined operations and uniform data protection coverage.
- **Compact storage platform:** VNXe systems provide a highly available design, where dual controller designs fit in only 2U or 3U of rack space.

Other features of the VNXe series are explained in the following sections.

Hardware overview

VNXe3100 systems can scale from one SP supporting up to 48 disks to two storage processors supporting up to 96 disks. In single storage-processor models of VNXe3100, a cache protection module in the empty storage-processor slot maintains a copy of write cache content. You can upgrade a single SP VNXe3100 system to a dual SP VNXe3100 system while keeping the original data in place; however you will not have access to that data during the upgrade.

VNXe3300 systems have two SP that support up to 120 disks. In both models, you can add additional disks by daisy-chaining disk array enclosures (DAEs) on the single SAS back-end bus. Table 1 compares the hardware features of these two storage-system models.

Table 1. VNXe3100 and VNXe3300 details

	VNXe3100		VNXe3300
	Single SP	Dual SP	
Number of SPs	1	2	2
Form factor (min)	2U	2U	3U
Max number of DAEs	3	7	7
Min / max number of drives	6 / 48	6 / 96	6 / 120
Drive options	300 GB 15k rpm SAS 600 GB 15k rpm SAS 1 TB 7.2k rpm NL-SAS 2 TB 7.2k rpm NL-SAS		100 GB Flash* 300 GB 15k rpm SAS 600 GB 15k rpm SAS 1 TB 7.2k rpm NL-SAS* 2 TB 7.2k rpm NL-SAS
Drive enclosure options	12 x 3.5" SAS / NL-SAS drives		15 x 3.5" Flash* / SAS / NL-SAS drives
CPU	1 x Intel Xeon Dual Core	2 x Intel Xeon Dual Core	2 x Intel Xeon Quad Core
Memory per system	4 GB	8 GB or 16 GB*	24 GB
Back-end connection / number of back-end ports	6 Gb SAS / 1		6 Gb SAS / 1
Embedded I/O ports per SP	2 x 1 GbE		4 x 1 GbE
Configurable I/O slots per SP	1		1
Port options per Flex I/O modules	4 x 1 GbE		4 x 1 GbE, 2 x 10 GbE
RAID options	1/0, 5, 6		1/0, 5, 6
Management	LAN 2 x 10/100/1000 Copper GbE		LAN 2 x 10/100/1000 Copper GbE
Limits and Support			
Maximum number of	128	256	512

supported LUNs		
Maximum LUN size	2 TB	2 TB
Maximum file system size	16 TB	16 TB
Max raw capacity	192 TB	240 TB
Maximum Initiators	256	256

* The VNXe system needs to be running code version 2.1.0.xxxx or later.

The total addressable capacity per SP depends on the raw capacity that has been configured on the storage system. There is no limit per SP.

Comparison of drive types

All VNXe storage systems support SAS and NL-SAS drives. VNXe3300 storage systems also support Flash drives. Flash drives are recommended for performance-intensive applications or parts of applications with very low response-time and high-throughput requirements. SAS drives should continue to be the choice for environments with large-capacity high-performance requirements. Flash and SAS drives are appropriate for applications such as database applications that require frequent read and write operations. NL-SAS drives are recommended for modest-performance and high-capacity environments. In addition, NL-SAS drives can provide energy efficient bulk storage capacity at low cost. NL-SAS drives are appropriate for storing large amounts of data that is less frequently used (such as video files, audio files, and images) and for users with applications that do not have strict performance requirements.

For the VNXe series, disks are sold in disk packs. A disk pack is a set of drives that share the same capacity, type, and speed. The VNXe3100 disk packs are available in the following combinations:

- **SAS disk packs** - Sold in groups of five to make 4+1 RAID 5 groups, or groups of six to make 3+3 RAID 1/0 groups. A hot spare disk is recommended for every 30 SAS disks. When ordering a system with SAS disk packs, it is highly recommended to order the base model with SAS drives. This ensures that the system drives (drives 0 – 3 in the first DAE) are build on faster performing SAS drives.
- **NL-SAS disk packs** - Sold in groups of six to make 4+2 RAID 6 groups.
- Individual hot spare drives for SAS and NL-SAS drives.

The VNXe3300 disk packs are available in the following combinations:

- **Flash drive packs** – Sold in groups of five to make 4+1 RAID 5 groups. A hot spare disk is recommended for every 30 Flash drives.

- **SAS disk packs** - Sold in groups of seven to make 6+1 RAID 5 groups, or groups of six to make 3+3 RAID 1/0 groups. A hot spare disk is recommended for every 30 SAS disks. When ordering a system with SAS disk packs, it is highly recommended to order the base model with SAS drives. This ensures that the system drives (drives 0 – 3 in the first DAE) are build on faster performing SAS drives.
- **NL-SAS disk packs** - Sold in groups of six to make 4+2 RAID 6 groups.
- Individual hot Spare drives for Flash, SAS and NL-SAS drives.

RAID types

When the VNXe system is initially configured, or when drives are added, the disk packs are configured into storage pools. Storage pools are an easy-to-use and economical alternative to traditional storage-system provisioning.

Each storage pool has a RAID group with the desired capacity and drive characteristics. The supported RAID levels are 5, 6, and 1/0. It is important to note that once a storage pool is configured, you cannot change the RAID type for the pool.

Each VNXe system has a disk pack in the first disk enclosure; these disks store the operating environment for the VNXe system. If a RAID 5 disk pack is chosen for this disk pack, a hot spare is automatically included.

RAID 5 is best suited for transaction processing, and is often used for general-purpose storage, relational databases, and enterprise resource systems. This RAID level provides a fairly low cost per megabyte while still retaining redundancy. RAID 5 stripes data at a block level across several disks and distributes parity among the disks. No single disk is devoted to parity. Distributed parity requires all disks, but only one to be present to operate. If a disk fails, it will reduce storage performance and should be replaced immediately, but data loss does not occur as a result of a single disk failure.

RAID 6 storage pools are appropriate for the same types of applications as RAID 5, but in situations where providing increased fault tolerance is important. RAID 6 is similar to RAID 5, but it uses double parity that is distributed across different disks, so it offers extremely high fault and disk-failure tolerance. RAID 6 provides block-level striping with parity data distributed across all disks. In this RAID configuration, storage systems continue to operate with up to two failed disks. RAID 6 gives time to rebuild the storage pool without the data being at risk if a single additional disk fails before the rebuild is complete.

RAID 1/0 storage pools provide both high performance and reliability at medium cost, while providing lower usable capacity per disk. RAID 1/0 may be more appropriate for applications with fast or high processing requirements, such as enterprise servers and moderate-sized database systems. RAID 1/0 storage pools in VNXe systems require a minimum of six physical disks, where three mirrored sets in a striped set together to provide fault tolerance. Although mirroring provides fault tolerance, if any

disk is lost it must immediately be replaced and the pool rebuilt because this configuration cannot handle the loss of more than one disk in the same mirror pair.

Software overview

This section discusses the software features available in VNXe systems.

Connection Utility and Configuration wizard

Connection Utility is a software tool that you use to establish an IP address for managing the VNXe storage system. You have to do this before you can configure the VNXe storage system or create storage resources. You can download Connection Utility directly from the product support page.

When you run Connection Utility from a computer in the same subnet as the VNXe storage system, Connection Utility automatically discovers any unconfigured VNXe systems. If you run the Connection Utility on a different subnet, you can save the configuration to a USB drive and transfer it to the VNXe system.

After you have configured the VNXe system, you can connect to the VNXe system through a web browser using the management IP address in Connection Utility.

The first time that you connect to a VNXe storage system, the VNXe Configuration wizard automatically starts. The Configuration wizard lets you set up the initial configuration of the VNXe system so that you can start to create storage resources. You set up the following parameters in the Configuration wizard:

- Login password for the user “admin”
- Service password for the “service” account
- DNS and NTP servers
- Shared folder and iSCSI server (optional)
- Disk pools (optional)
- SMTP server for product support options (optional)
- EMC online support credentials (optional)
- EMC Secure Remote Support (ESRS)

After running the Configuration wizard, you can download and install the license file for the software and features supported on your VNXe system. For additional information about downloading and installing VNXe license files, refer to the EMC Unisphere for VNXe: Next-Generation Storage Management white paper on EMC Online Support (<https://support.emc.com>) › VNXe Product Page.

Unisphere

Unisphere is the graphical user interface that you use to monitor and manage VNXe series storage systems. It provides tools for:

- Creating, configuring, and managing storage resources for data for Microsoft Exchange servers, VMware servers, databases, and network hosts and users that use shared folders and shares
- Setting up access to storage resources for users, applications, and hosts, using Active Directory, NFS, and iSCSI controls
- Monitoring storage operations and systems status through detailed graphical reporting services
- Monitoring storage system performance
- Protecting files and data by setting up automated schedules for recording point-in-time snapshots of stored data
- Protecting files and data by setting up local and remote replication with another storage system
- Recovering lost or unavailable data quickly and easily from recorded snapshot images
- Accessing VNXe [ecosystem](#) resources like FAQ documents, how-to videos, online discussion forums, and chat sessions
- Downloading and updating the software code on the VNXe series storage systems using a non-disruptive-upgrade (NDU) process

Unisphere's Home page (Figure 3) provides a convenient starting point for setting up network-accessible data storage, scheduling data protection, and monitoring system health status and performance.




Figure 3. Unisphere Home page

On the Home page you can:

- Monitor system health, capacity, and performance information
- Configure network settings, user access, host connections, software, and licenses
- Set system preferences for the management interface by specifying the language and password settings
- Obtain support information and access information resources

For additional information about Unisphere for VNXe series storage systems, refer to the EMC Unisphere for VNXe: Next-Generation Storage Management white paper on EMC Online Support (<https://support.emc.com>) > VNXe Product Page.

Storage provisioning

The types of storage the VNXe system supports depend on the types of storage servers configured to run on it: iSCSI servers, shared-folder servers, or both. You can configure VNXe storage server settings through the VNXe Configuration wizard or use the Settings page in Unisphere to add storage servers as needed.

Table 2 shows the supported VNXe storage types, the kinds of hosts that can access the storage, and the type of VNXe storage servers required to support that storage.

Table 2. VNXe storage details

Storage Type	Hosts	Managed Objects	Storage Server
Microsoft Exchange	Microsoft Exchange servers running on Windows hosts	Exchange storage groups, virtual disks	iSCSI server
Shared folders	Linux/UNIX hosts	Shared folders and shares (NFS)	Shared-folder server with NFS enabled
	Windows users/groups	Shared folders and shares	Shared-folder server with CIFS enabled
Generic iSCSI	Windows or Linux/UNIX hosts that use iSCSI for connection to storage	iSCSI virtual disks	iSCSI server
VMware	VMware hosts	NFS datastores	Shared-folder server (NFS-enabled) for datastores accessed via NFS
		VMFS datastores	iSCSI server for datastores accessed via iSCSI
Hyper-V	Windows Server 2008 hosts	Hyper-V datastores	iSCSI server

Exchange storage

VNXe Exchange storage resources provide storage optimized for supporting Microsoft Exchange 2003/2007 or Exchange 2010 server operation. When you configure storage for an Exchange server, VNXe creates and configures the necessary storage groups in accordance with the Exchange server's requirements and Microsoft best practice recommendations. When the Exchange server connects to the Exchange storage groups, it can access the storage as if it were a local resource.

An Exchange storage resource is a *container* for iSCSI storage groups that are associated with a specific Exchange server. Each Exchange container can contain one or more storage groups, depending on the number of mailboxes that the Exchange server supports; the average mailbox size; and the size of the database, logs, or public folders it uses.

Each storage group is a managed resource that stores Exchange server data. Each storage group contains virtual disks associated with the server database, log files, and (optionally) public folders.




Figure 4. Microsoft Exchange Wizard

It is important to remember that the Unisphere’s Exchange Details view (shown in Figure 4) shows the number of mailboxes supported based on the average mailbox size you specified during the creation process. This number does not indicate the actual number of mailboxes on the Exchange server.

Shared folders

Shared folders and shares provide file-based storage resources that Windows and Linux/UNIX hosts can access over network connections. A shared folder is a file-based storage resource that is associated with a particular quantity of storage and file access protocols (CIFS or NFS). A share is an access point through which hosts can access shared folders.

When file-based storage is enabled on the system, a shared-folder server maintains and manages file systems for shared folders, and transfers data to and from hosts using the CIFS or NFS file-sharing protocols. The type of access to VNXe shares varies according to whether it is a CIFS share or an NFS share.

In Unisphere, when you create the shared folder server on a VNXe system, you can choose to enable the CIFS and NFS protocols. This determines the type of shared folders that you can create using this shared folder server. However, even if you enable both the CIFS and NFS protocols on a shared-folder server, you can create only single-protocol shared folders on VNXe systems, which is illustrated in Figure 5. For example, if you choose the CIFS protocol when creating a shared folder, you can create only CIFS shares in that shared folder. You can create a new shared folder and enable NFS, allowing you to create NFS shares in the VNXe system.




Figure 5. Shared Folder Wizard

Shared-folder servers can be created as domain-joined or standalone. Domain-joined servers are configured using domain-based Kerberos authentication; they maintain their own identity in the domain, and leverage domain information to locate services such as domain controllers. Standalone servers do not have access to a Windows domain; these servers are useful in test and development environments.

Unisphere allows you to choose a replication option for your shared folder in the Results window of the Shared Folder Wizard. By default, replication for the shared folder is disabled in the wizard.

You can choose to replicate locally if you want to create a replica of the shared folder on the local system. You can also choose to replicate to a remote system if you want to create a replica of the shared folder on another system on the network. If you choose to create a local or remote replica of this shared folder, you can do this when you close the Shared Folder Wizard and the Replication Session Wizard appears. If you choose not to create a replica of this shared-folder storage resource, you can still initiate replication at a later time.

Generic iSCSI storage

Generic iSCSI storage resources provide hosts with access to general purpose block-level storage through network-based iSCSI connections. With iSCSI storage you can manage addressable partitions of block storage resources so that host systems can mount and use these resources (virtual disks) over IP connections. After a host connects to a Generic iSCSI virtual disk, it can use the virtual disk like a local storage drive.

For example, assume a database application requires a virtual disk for a database and a separate virtual disk for logs files. You can represent them under the same generic storage application in Unisphere by adding **Virtual Disk** to the Generic iSCSI application. Figure 6 shows an iSCSI application with two virtual disks.




Figure 6. iSCSI storage application with two virtual disks

Unisphere allows you to create, view, manage, and delete virtual disks associated with Generic iSCSI storage. You can also configure protection storage through the iSCSI storage wizard. To support third-party backup applications that may require storage resources, VNXe automatically reserves a minimum percentage (5%) of its allocated storage for protection, even when protection is not enabled. After you create a Generic iSCSI storage resource, you can increase but not reduce the quantity of primary and protection storage allocated for the resource.

Protection storage will be covered later in this whitepaper (see *Data Protection* section for more details).

VMware storage

Unisphere allows you to create storage resources optimized for use by VMware vCenter™ servers and ESX® hosts. A VMware datastore is associated with a specific quantity of storage and a type of storage access — VMFS or NFS. You can configure VMware datastores using the VNXe VMware Storage wizard, which configures datastore settings.

VMFS datastores require an iSCSI shared-folder server on the storage system. NFS datastores similarly require a shared-folder server to be available. The restrictions that apply for Generic iSCSI storage in a VNXe system also apply to VMware VMFS datastores. That is, after you create a VMFS datastore, you can increase but not reduce the quantity of primary and protection storage allocated for the resource.

The window under the Hosts tab in Unisphere allows you to discover and add VMware ESX server and vCenter server information from your network into Unisphere. Once the VMware host information has been imported into Unisphere, information about virtual machines deployed on those servers is automatically displayed and updated. Furthermore, when VMware datastores are created through Unisphere, the new storage is automatically scanned and becomes available to the ESX server.

Hyper-V storage

Hyper-V datastores provide hosts with access to general-purpose block-level storage through network-based iSCSI connections. Using a Windows Server host with Hyper-V, you can mount and format these resources (virtual disks) over an IP connection, and then use the virtual disk like a local storage drive to provide storage to Hyper-V virtual machines. The Hyper-V storage wizard in Unisphere requires an iSCSI storage server to be configured on the storage system.

The restrictions that apply for Generic iSCSI storage in a VNXe system also apply to Hyper-V storage. Here again, after you create a Hyper-V datastore, you can increase but not reduce the quantity of primary and protection storage allocated for the resource.

Storage pool characteristics

VNXe storage is contained in storage pools. A storage pool is a group of disks of similar type and speed. When storage is allocated for application use, you have the option of choosing the type of storage pool. Depending on the VNXe model and its disk configuration, different storage pools are available for new storage. Storage pools can be created automatically by the system depending on disk types present, or they can be custom created via a wizard.

Table 3 shows what types of storage pools you can create on VNXe systems.

Table 3. VNXe storage pool characteristics

Types of storage pools	VNXe3100	VNXe3300
Extreme Performance pool (Default)	N/A	4+1 RAID 5 (Flash)
Performance pool (Default)	4+1 RAID 5 (SAS)	6+1 RAID 5 (SAS)
Capacity pool (Default)	4+2 RAID 6 (NL-SAS)	4+2 RAID 6 (NL-SAS)
Custom pool	4+1 RAID 5 (SAS) 3+3 RAID 1/0 (SAS) 4+2 RAID 6 (NL-SAS)	6+1 RAID 5 (SAS) 3+3 RAID 1/0 (SAS) 4+2 RAID 6 (NL-SAS)

When you enable automatic disk configuration, the system allocates existing disks into a capacity, and/or performance and/or extreme performance pool, depending on the number and type of available disks:

- NL-SAS disks are allocated in multiples of six in RAID 6 (4+2) groups with no assigned spare disks. For example if 45 NL-SAS disks are available, the capacity pool uses 42 of the disks (in seven six-disk groups), does not allocate any spare disks, and leaves three disks unassigned. You can manually create a hot spare using NL-SAS disks, if needed.
- In a VNXe3300 system, SAS disks are assigned in multiples of seven in RAID 5 (6+1) groups. A spare disk is assigned for the first 0-30 disks, and another spare disk is assigned for each additional group of 30 disks. For example if 45 SAS

disks are available, the performance pool uses 42 of the disks (in six seven-disk groups), allocates two spare disks, and leaves one disk unassigned.

- In a VNXe3300 system, Flash drives are assigned in multiples of five in RAID 5 (4+1) groups. A spare disk is assigned if there are left over drives. For example if 11 Flash drives are available, the extreme performance pool uses 10 disks (in two five-disk groups), and allocates one spare disk. If you had 10 Flash drives, the extreme performance pool uses 10 disks and no hot spares.
- In a VNXe3100 system, SAS disks are assigned in multiples of five in RAID 5 (4+1) groups. A spare disk is assigned for the first 0-30 disks, and another spare disk is assigned for each additional group of 30 disks. For example if 45 SAS disks are available, the performance pool uses 40 disks (in eight five-disk groups), allocates two spare disks, and leaves three disks unassigned.

Instead of configuring the storage pools automatically, you can configure custom storage pools to provide storage that is optimized for specific applications and usage patterns, and to achieve a specific level of performance, capacity, or cost-efficiency. Unisphere's Disk Configuration wizard provides the following custom pool configuration options:

- **Optimizing for application usage** - You can select a storage pool profile based on the type of application that will use the storage and the anticipated usage of the storage (for example, general purpose, backup, or database). Unisphere gives the following application-profile options: Microsoft Exchange, Generic Storage, Shared Folders, VMware, and Hyper-V.
- **Optimizing for general characteristics** - You can create a storage pool based on disk type and general storage characteristics. The options include Balanced Performance/Capacity, High Performance, High Capacity and Best Performance.

When creating a storage pool for a specific application, different options are presented to the user depending on the quantity of storage in the system that is not assigned to a pool. The options depend on the number and types of available disks, and the RAID level associated with the application profile. For example, a profile that uses SAS disks in a RAID 1/0 (3+3) configuration requires disk assignment in groups of six. The options are rated to show the general effectiveness of each disk type for the chosen application type and its intended usage.

After a storage pool is created, you can add new disks to increase pool capacity. The storage type, RAID level, and configuration determine how many disks you can allocate to the pool in each grouping. For example, NL-SAS disks using RAID 6 are configured in 4+2 configurations; therefore you must allocate them in groups of six.




Figure 7. Disk Configuration Wizard

Unisphere's Disk Configuration Wizard (shown in Figure 7) can also configure hot spare disks for the VNXe system. A spare disk is used to replace a failed or faulted device in a storage pool. The spare disk is not used to actively store data but can be used to automatically replace a failed disk. If a disk fails and there are no spare disks, the system may run in degraded mode until a spare is provided or the failed disk is replaced. During this time, the data is at greater risk if an additional failure occurs. You can select the number of storage disks to be configured as hot spare disks by using the Disk Configuration Wizard.

Data protection

A small to medium organization's data is one of its most valuable assets. Therefore, the company's highest priorities must include safeguarding the data. EMC VNXe series provides integrated features that meet customers' goals of business continuity and data protection. In this white paper, data protection for VNXe systems is summarized in two categories: [snapshots](#) and [replication](#). For additional information about these features, refer to the *EMC VNXe Data Protection* white paper on EMC Online Support (<https://support.emc.com>) > VNXe Product Page.

Snapshots

A snapshot is a virtual point-in-time image of the data stored on a VNXe system. Snapshots provide a record of the content in the source storage at a particular date and time, but are not mirror copies of the data. VNXe snapshots are read-only. *Promoted* snapshots are for backup and restore purposes only. Periodically creating snapshots of the file systems and virtual devices provides an effective technique for meeting data protection and recovery requirements.

Unisphere automatically provisions extra storage for protection whenever a storage resource is created. This is in addition to the *primary storage* that is used for storing host, user, or application data. Primary storage can include Microsoft Exchange server data, files within shared folders and shares, or data.

Protection storage is automatically provisioned to store snapshot (and third-party replication) data, and to restore or recover primary storage. Every storage resource has a minimum amount of protection storage to support manual snapshots of the primary data if necessary.

You can manage primary and protection storage separately for storage resources. For example, you can increase the protection storage for a resource without affecting its primary storage. For Generic iSCSI storage resources, VMFS datastores, and Hyper-V storage, you can set separate host access permissions for primary and protection storage; however you cannot reduce the size of the protection storage once it has been created. Unisphere allows you to auto-adjust the size of protection storage, depending on the size of primary storage. When enabled, VNXe adjusts the protection storage size proportionately to any change to the size of the primary storage.

VNXe storage systems have built-in tools for protecting stored data by using snapshot schedules to create point-in-time copies of the data. You can choose to use standard snapshot schedules, create custom schedules, or take manual snapshots to create *storage checkpoints*. These checkpoints can be used to restore or recover data. Based on the importance and volatility of the source data, you can define recurring schedules that specifically determine the time when a snapshot is taken and how long that snapshot is kept in the system.

VNXe provides tools for performing manual (on-demand) snapshots, configuring snapshot schedules, and performing snapshot restore and recovery operations for the following storage types:

- CIFS shared folders
- NFS shared folders
- VMware datastores (using NFS)
- VMware datastores (using iSCSI)
- Generic iSCSI storage
- Hyper-V storage

Shared-folder snapshots are performed on the entire shared folder – you cannot take a snapshot of just a portion of the shared folder. Snapshots for the iSCSI storage resources mentioned above, when collected through Unisphere, are crash-consistent and not application-consistent, unless the applications are quiesced before taking the snapshot.

Replication Manager creates and manages application-consistent snapshots for Exchange¹, VMware datastores, Hyper-V datastores, and Generic iSCSI storage. Replication Manager runs on a separate Windows server. Allocating storage for the application and protection is configured with the Create Storage wizard within Unisphere.

An administrator can access the snapshots from a host mounted (NFS) or mapped (CIFS) to the share by adding `\.ckpt` to the end of the share path to go to the hidden snapshot directory. For a CIFS shared folder only, an administrator has the option to access the snapshot through Microsoft Explorer under the Previous Versions tab in the share's Properties window. Snapshot restores of shared folders can also be performed within Unisphere to retrieve the files or data recorded on a specific date and time. Snapshots that were created after the snapshot that you use to restore are destroyed and cannot be recovered in iSCSI storage resources (for example, Generic iSCSI storage, VMware VMFS datastores, Exchange storage, and Hyper-V datastores.) However, for shared folders and VMware NFS datastores, all snapshots remain on the system.

When you schedule snapshot creation through a protection schedule, if the protection storage runs out, the oldest snapshot is deleted to make space for new snapshots.

The VNXe system allows you to choose or customize snapshot schedules that specify regular times to perform snapshot operations (automatic snapshot creation and deletion). Administrators can select the intervals, times, days, and dates at which snapshot operations occur. A collection of rules within the schedule specify the interval, frequency, and time that snapshots are taken. The VNXe calculates the amount of storage protection space needed based on the complexity of the protection schedule chosen by the administrator. This value can be adjusted if desired. It is important to note that if snapshot schedules are created via Unisphere for applications that reside on iSCSI volumes, these snapshots should not be used for backup purposes. Further, if multiple Generic iSCSI storage instances are created via Unisphere and presented to an application, the snapshots are not crash consistent.

The VNXe system also allows you to manually create snapshots that contain on-demand images of specific storage resources at a particular time. This feature is useful in various situations, such as taking snapshots of data that will be unavailable for a period of time, that will be subject to unusual circumstances or risks, or that is used for testing or debugging purposes.

¹ Scheduling snapshots for Exchange storage must be performed using Replication Manager.

Replication

Storage replication is a process in which storage data is duplicated either locally or to a remote network device. Replication produces a read-only, point-in-time copy of source storage data and periodically updates the copy, keeping it consistent with the source data. Storage replication provides an enhanced level of redundancy in case the main storage backup system fails; this minimizes the downtime-associated costs of a system failure and simplifies the recovery process for a natural or human-caused disaster. VNXe replication provides easy-to-use asynchronous replication tools for creating point-in-time copies of storage resources. This copy is a complete copy of the data, unlike a snapshot of the data, in which only the changes to the data are saved.

There are two types of replication:

- **Local replication** – Local replication occurs within the same VNXe system.
- **Remote replication** – Remote replication occurs between the source VNXe system and the remote VNXe, VNX, or Celerra system.

File-level replication, such as CIFS and NFS shared folders or VMware NFS datastores, is managed entirely within the VNXe environment.

Block-level replication, such as Hyper-V and Exchange, leverages EMC Replication Manager to create, schedule, and manage iSCSI replications. Replication Manager manages EMC point-in-time replication technologies and coordinates the entire data replication process from discovery and configuration to the management of multiple disk-based replicas.

As of VNXe OE 2.2.x, replication sessions for Generic iSCSI storage resources can be established within Unisphere. After the replication session has been created, you will need to discover the session and set up replication from EMC Replication Manager.

Replication Manager provides a GUI for managing the replication of iSCSI LUNs. Before creating a replica, Replication Manager ensures that applications are in a quiescent state and that the cache is flushed so that the replica is consistent from the point of view of client applications.

In a VNXe system, you can control the frequency at which the source and destination systems are synchronized. The value can be set between 5 and 1,440 minutes (24 hours). The default value is 60 minutes. Settings greater than 6 hours result in a warning to the user that they should increase their Protection Reserve storage amount to a size equal to 10 percent of the application's size. The default 5 percent Protection Reserve will otherwise be sufficient for settings of 6 hours or less.

Once a replication connection has been established between the source and destination system, the following operations can be executed on the replicated data sets:

- **Replication failover** – In a failover scenario, the production site becomes unavailable and inaccessible, typically as the result of a disaster or unexpected outage. In response, the user can execute a failover from the destination system. The execution of a replication failover is not an automatic occurrence—it requires

manual execution. A failover operation is asynchronous and results in data loss if the data has not been synchronized between the source and destination prior to executing the failover.

If possible, run the Sync Now function to synchronize production with the destination before executing failover to avoid losing data. During the failover process, read/write access is provided to the destination application. When the production application becomes reachable, only read operations are allowed. After the production site has been restored to service, the original replication session can be restarted by executing a failback on the production application.

- **Replication switchover** – The switchover feature is typically used for disaster-scenario testing or data migrations. The switchover differs from the failover in that it synchronizes the production application with the destination using the differential snapshot from the source. Switchover stops the replication session and does not restart the replication session. This option should be used if the source site is available and you want to activate the destination file system as read/write. This command should only be executed from the source system. This command pauses replication, mounts the source object as read-only, and mounts the destination object as read/write so that it can act as the new source object.
- **Replication failback** – Following a switchover or failover, you can execute a failback to restore the source as the read/write production application, and to resume the replication session. Failback can only be initiated from the production side. The failback is initiated under the Replication tab on the production application. Failback resumes the replication session without losing the data written to the destination while in the failed-over or switched-over state. It does this by synchronizing the applications as part of the failover process.

Backup and restore

The VNXe series simplifies backup and restore operations. In a typical data center, backups need to be performed on multiple application servers. VNXe storage systems allow the application data to be stored at a central location, where the data can be backed up, and restored if required.

VNXe supports the following backup products:

- EMC NetWorker®
- EMC NetWorker® FastStart
- EMC Avamar 6.0
- EMC DataDomain 5.0
- CommVault Simpana
- Symantec Backup Exec
- Symantec NetBackup with NDMP

For additional information, refer to the *EMC Simple Support Matrix for EMC VNXe Series* on EMC Online Support (<https://support.emc.com>) > VNXe Product Page.

VNXe storage systems support NDMP v2-v4 over the network; direct-attach NDMP is not supported. This means that the tape drives need to be connected to a media server, and the VNXe system communicates with the media server over the network. NDMP has an advantage when using multiprotocol file systems because it backs up the Windows ACLs as well as the UNIX security information. Also, when you are using deduplication on a file system, the files are backed up in their compressed state when NDMP is used.

CLI support

The CLI in the VNXe system enables you to do scripting for some of the most commonly performed tasks. You can use the VNXe management IP address in CLI from a Windows or Linux/UNIX console to log in and execute commands on VNXe storage systems.

To use CLI, you have to install a CLI client on your host machine. You can use the same client to manage multiple storage systems. Details about the CLI format and options can be found in the product support page for VNXe systems.

Advanced features

A VNXe storage system has certain advanced features and network configuration options, which are discussed here.

Link aggregation

Both models in the VNXe series support link aggregation. With link aggregation, up to four Ethernet ports can be combined into one logical link. Each storage processor must have the same type and number of Ethernet ports, and the cabling on SPA must be identical to the cabling on SPB, otherwise link aggregation cannot be configured. Link aggregation provides the following advantages:

- **High availability of network paths to and from VNXe** – If one physical port of an aggregated port fails, the VNXe system does not lose connectivity.
- **Possible increased overall throughput** – This is because multiple physical ports are bonded into one logical port. If you are working with iSCSI protocol, you also have the option of using multi-pathing to increase throughput.

For example, if you combine two physical ports into one logical port and one of the links fails, the other link carries the network traffic without disrupting host access. Once the link comes back, network traffic flows over both healthy links again.

In order to configure link aggregation, both linked ports must be cabled to the same switch or to the same logical switch if the switches support stack interconnects that allow for cross-stack Link Aggregation, and the switch must be configured to use link aggregation using the Link Aggregation Control Protocol (LACP).

VNXe allows you to combine the first port from the on-board I/O ports (or the Flex I/O module) to any other port on the same card.

Fail Safe Networking

Fail Safe Networking (FSN) is configured on the VNXe storage systems by default. If one of the ports in the storage system fails, FSN will automatically reroute the I/O internally to the corresponding physical port on the peer storage processor.

For example, you can have port 'eth2' on SPA plugged into switch A; and port 'eth2' on SPB plugged into switch B (both of these switches must be on the same subnet). If a host application has data going over port 'eth2' on SPA and switch A fails, FSN will route the I/O traffic over port 'eth2' on SPB using switch B. Fail Safe Networking therefore helps in providing switch level redundancy in your environment.

Thin provisioning

Thin provisioning is the ability to present a server with more capacity than is actually allocated within the storage system, essentially giving the host the illusion it has capacity that is not physically allocated from the storage system.

When you enable thin provisioning for a storage resource, the amount of storage requested is not immediately allocated to the resource. Instead, Unisphere allocates a smaller quantity of storage to the storage resource. When the amount of storage consumed within the storage resource approaches the limit of the current allocation, the system allocates additional storage to the storage resource from the pool.

Thin provisioning allows multiple storage resources to subscribe to common storage capacity within a pool, while the system allocates only a portion of the physical capacity requested by each storage resource. The remaining storage is available for other storage resources to use. Figure 8 compares pool usage between standard provisioning and thin provisioning.




Figure 8. Comparison between standard provisioning and thin provisioning

With thin provisioning, you can improve storage efficiency while reducing the time and effort required to monitor and rebalance existing pool resources. Organizations can purchase less storage capacity upfront, and increase available disk capacity (by adding disks) as needed and according to actual storage usage, instead of basing disk requirements on the requests or predictions for connected hosts.

Thin provisioning can be enabled, via a checkbox, in the Storage Allocation wizard and is supported on all storage types in a VNXe storage system. When storage is provisioned, the storage system allocates a small amount of storage for that application. This is called the *initial allocation*. The initial allocation for file-based storage types is 10 GB and the initial allocation for iSCSI-based storage types is 1 GB. If the storage has been provisioned on an iSCSI storage server, thin provisioning can be enabled only in the Storage Allocation wizard, and this setting cannot be changed later. With shared folders and NFS VMware datastores, thin provisioning can be enabled or disabled after the storage resource has been created.

VNXe systems allow the storage pools to be *oversubscribed*. Oversubscription is a storage provisioning method that allows storage administrators to provision more capacity than may be physically available in a particular storage pool. When groups of thin-provisioned storage resources are associated with a storage pool, they can request (*or subscribe to*) more storage capacity than is actually in the pool. However, storage from the pool is allocated incrementally to each storage resource based on the amount of used storage within the resource. When the collective amount of storage used within the pool approaches the pool's physical capacity, Unisphere generates notification messages that more storage may soon be required.

The administrator can then add more storage disks to the system or assign more available disks to the pool. Hosts connected to a thinly provisioned storage resource see the *requested* space rather than the *allocated* space in the pool. Unisphere allows you to monitor storage pool usage and subscription levels so that you can assess the current status of pool subscription, and predict the need for additional disks. You can also specify the threshold at which the system generates pool usage notifications by specifying the percentage of pool utilization at which notifications occur. The default setting is 85%, but you can set this value anywhere between 50% and 85%.

File-Level Retention (enterprise)

File-Level Retention (FLR) in VNXe storage systems provides a way to set file-based permissions on a CIFS or NFS shared folder to limit write access to the files for a specific period of time. FLR ensures the integrity of data during that period by creating an unalterable set of files and directories. You can only enable FLR when you create the shared folder. After a shared folder is created and enabled with FLR, protection is applied on a per-file basis.

A file in an FLR-enabled shared folder is always in one of the following four states:

- **Not Locked** – A file that is not locked is treated exactly as a file in a file system that is not enabled for FLR; it can be renamed, modified, or deleted.

- **Locked** – A file in locked state can be managed by setting retention dates that, until the dates have passed, prevent the files from being modified or deleted.
- **Append** – A file in append state cannot be deleted or renamed. Existing data in this file cannot be modified, but new data can be added.
- **Expired** – A file in expired state cannot be renamed, modified, or appended to. These files can be deleted or relocked, if needed.

Deduplication

VNXe storage systems have integrated deduplication support for file-based storage (shared folders and NFS datastores). This optimizes storage efficiency by eliminating redundant data from the stored files, thereby saving storage space. From an end-user perspective, the NAS clients are not aware that they are accessing deduplicated data.

Deduplication operates on whole files that are stored in the file system. For example, if there are 10 unique files in a file system that is being deduplicated, 10 unique files will still exist but the data will be compressed, yielding a space savings of up to 50 percent. On the other hand, if there are 10 identical copies of a file, 10 files will still exist but they will share the same file data. The one instance of the shared file is also compressed, providing further space savings.

During deduplication, each deduplication-enabled file system is scanned for files that match specific criteria such as a particular file type or a modification time older than a certain date. When a file is found that matches the criteria, the file data is deduplicated and compressed if appropriate. Since file metadata is not affected by deduplication, different instances of the file can have different names, security attributes, or timestamps, but they are still deduplicated. The policy engine that controls the deduplication process is throttled according to the storage processor utilization. If the SP utilization is above a certain value, the deduplication process is paused and is started again when the utilization value drops below the threshold.

Virus protection

The VNXe storage system supports third-party anti-virus servers that perform virus scans and reports back to the VNXe. For example, when a shared-folder client creates, moves, or modifies a file, the VNXe storage system invokes the anti-virus server to scan the file for known viruses. If the file does not contain a virus, it is written to the VNXe. If the file is infected, corrective action is taken as defined by the antivirus server.

VNXe supports the following antivirus servers:

- Symantec SAV for NAS
- Symantec endpoint protection
- McAfee VirusScan
- Computer Associated eTrust
- Sophos Anti-Virus

- Kaspersky Anti-Virus
- Trend Micro ServerProtect

User management

VNXe storage systems provide tools for creating user accounts for managers and administrators who configure and monitor VNXe systems. When users access Unisphere, they are prompted to log in with account-based credentials before they access the system. Unisphere accounts combine a unique username and password with a specific role for each account. The role determines the types of actions that the user can perform after logging in.

The following user roles are available in Unisphere:

- **Operator** – This role can view Unisphere system and storage status information but cannot change system settings.
- **Storage administrator** – This role can view VNXe storage system data, edit Unisphere settings, use Unisphere tools, and create and delete storage resources and host configurations. However, this user cannot add user accounts or host configurations, perform initial configuration of the system, modify network settings, create or delete storage servers, or upgrade system software.
- **Administrator** – This role can view storage-system data, configure Unisphere system settings, and perform all other tasks accessible through Unisphere.

LDAP integration

Lightweight Directory Access Protocol (LDAP) is an application protocol for querying and modifying directory services running on TCP/IP networks. LDAP provides a management service for network authentication and authorization operations by centralizing user and group management across the network. Integrating Unisphere users into an existing LDAP environment provides a way to control management access to the VNXe system based on established user and group accounts within the LDAP directory. In Unisphere you can configure LDAP settings under the Directory Services tab of the Manage Administration window.

You should note that LDAP settings are used only for controlling access to Unisphere, not for controlling access to VNXe storage resources. Also, the Active Directory services should be installed in your environment before integrating the VNXe system into it.

High availability

The VNXe series of storage systems offer several built-in high-availability features. This high availability is provided through redundant components. If one component fails, the other one is available to back it up. The redundant components include SPs, cooling fans, AC power cords, power supplies, I/O modules, and Link Controller Cards (LCCs). Network high availability is provided through link aggregation.

For network high-availability features to work, the cable on each SP needs to have the same connectivity. If Port 0 on SPA is plugged in to Subnet X, Port 0 on SPB must also be plugged in to Subnet X. This is necessary for both server and network failover. If a VNXe server is configured to use a port that is not connected on the peer SP, an alert is generated. Unisphere does not verify if they are plugged in to the same subnet, but they should be, for proper failover. If you configure a server on a port that has no cable or connectivity, the traffic is routed over an SP interconnect path to the same port on the peer SP (just a single network connection for the entire system is not recommended).

For additional information about high availability in VNXe storage systems, refer to the *EMC VNXe High Availability* whitepaper on EMC Online Support (<https://support.emc.com>) > VNXe Product Page.

Some of the examples of high availability are:

- **SPs** – When an SP experiences hardware or software failure, reboots, or a user places it in Service Mode, a failover occurs. Storage servers that use the out-of-service SP fail over to the other SP if it is available, with minimal disruption between the VNXe system and connected hosts.

The type of disruption depends on the protocol used to access storage. The NFS protocol keeps retrying the active I/O indefinitely; the CIFS protocol fails the active I/O but retries each time a new I/O is attempted from the storage volume. In the case of iSCSI storage volumes, it is the iSCSI host's responsibility to decide how long the active I/O would be retried; the default value is 5 seconds, which is recommended to be changed to 600 seconds.

After the SP is available again, the storage servers fail back to the original SP. The failback policy can be configured through Unisphere.

- **Power supplies** – Each SP complex has its own power supply. Each VNXe disk enclosure (DPE or DAE) can continue to run with one failed power supply in the enclosure. The SP power supply also contains a battery backup unit (BBU) that provides enough backup power to write cache data to non-volatile media in the case of power interruption.
- **Network links** – If a network link fails, link aggregation provides an alternate path.
- **Network paths** – The VNXe supports network pass-through to provide network path redundancy. If a network path becomes unavailable due to a failed NIC, switch port, or bad cable, network traffic is re-routed through the peer SP using an inter-SP network, and all the network connections remain active.
- **Disk paths** – Dual paths are provided to all disks.

- **RAID protection** – The data on the disks is protected with RAID 5, RAID 6, or RAID 1/0. With RAID 5, data is still available after a single physical drive failure. With RAID 6, data is available in the event of a double disk failure. RAID 1/0 uses mirroring for availability and striping for performance; if a disk in the mirrored pair fails, the mirror provides continued access to data.
- **Hot spares** – VNXe supports hot spare technology, wherein a failing drive can be replaced automatically with a drive that has been configured as the hot spare for the system.

Virtual integration with VMware vSphere

The VNXe storage system offers tight integration with VMware vSphere. VMware administrators can take advantage of this functionality when managing their virtual environment.

Advanced Storage Access (ASA)

ASA allows VMware administrators to manage appropriately configured VNXe host configurations by taking advantage of advanced file operations that optimize NFS storage utilization. Once ASA is enabled on your shared folder storage or VMware datastore (configured at the server level), EMC's VSI Unified Storage Management tool can be utilized for the following:

- Simplifying the process of creating NFS datastores
- Compressing virtual machines in NFS datastores
- Reducing the amount of storage consumed by virtual machines by using compression and Fast Clone technologies. The cloning functions include:
 - FAST clones (thin copy/snaps) of Virtual Machine Disk (VMDK) files
 - FULL clones (full copy) of Virtual Machine Disk (VMDK) files
 - FULL clones (copies)

For additional information about VSI Storage Management, refer to the *EMC VSI for VMware vSphere: Unified Storage Management* product guide available on Powerlink.

vStorage API for Array Integration (VAAI)

VAAI² improves resource utilization of the ESXi host by offloading related tasks to the VNXe. NFS storage tasks are processed by the storage system, reducing the host processor and storage adapter resources required to perform select storage operations.

An example that illustrates VAAI is an operation such as provisioning full clones from one a template VM for your virtual environment. With VAAI enabled, the ESXi host streams write commands to the VNXe target. The VNXe storage system processes these requests internally, performs the write operations on the given SP and returns an update to the host when they are completed.

While the load is shifted to the storage system, the impact on host resources and front end ports of the VNXe system are significantly reduced.

² VAAI for file requires VMware vSphere 5.0

Remote monitoring

The status of a VNXe system can be monitored remotely by EMC personnel or customers themselves.

EMC Secure Remote Support (ESRS)

The EMC Secure Remote Support (ESRS) feature provides your authorized EMC service provider with remote access capabilities to your VNXe system using a secure and encrypted tunnel. For outbound access, the VNXe management IP network must allow outbound and inbound HTTPS traffic. The secure tunnel that ESRS establishes between the VNXe device and authorized systems on the EMC network can also be used to transfer files out to the VNXe system or transfer files back to EMC's network.

Unisphere Remote

Unisphere Remote is a centralized, easy to use, network application that proves administrators with a way to centrally monitor their VNXe storage systems.

Unisphere Remote enables you to:

- Monitor up to 3,000 VNXe systems from a single interface
- View aggregated alerts, health, and capacity and CPU usage for multiple systems
- Control access to the monitoring interface by setting up local Unisphere Remote users or integrating existing LDAP enabled users and groups
- Organize views of the VNXe nodes in logical ways, including by location, type, department, and so on
- Launch Unisphere for VNXe management interface for individual VNXe systems from Unisphere Remote

Note: Unisphere Remote does not provide active management capability of the VNXe systems in your environment. Instead, it leverages these capabilities in Unisphere by link-and-launching to the specific VNXe that you would like to manage.

The Unisphere Remote environment consists of a Unisphere Remote server running in a VMware virtualized environment, multiple VNXe systems, and a remote system to access the Unisphere Remote sever.

The VNXe series support ecosystem

VNXe series systems are customer-installable and customer-maintainable. When there is a hardware or software fault in the VNXe system, you are informed through the VNXe alert mechanism. The pop-up alert has links to context-sensitive knowledgebase articles that help you correct the fault. For example, if there is a faulted disk in the VNXe system, an alert message pops up that specifies which disk has faulted. Additionally, there are links to parts of the knowledgebase document that provide instructions to order a new drive, and replace the faulted disk in the VNXe system.

The Support window in Unisphere provides links to resources for learning about and getting assistance with your VNXe storage system. It provides the following features:

- **How to videos** – Videos to learn about your storage system; for example, instructions on how to replace a failed component.
- **Online documentation** – Online documents that provide the latest product documentation. These are routinely updated to ensure the most current information.
- **Online training** – Videos, slideshows, and other educational material to learn about VNXe systems.
- **Knowledgebase articles** – Articles, white papers, and other information regarding known issues, and solutions related to system installation, configuration, and operation.
- **Community access** – A way to interact with other VNXe customers and read, contribute, or ask questions regarding the storage system.
- **Live chat** – Quickly contact and chat with support personnel who can help assist you in real time. Note that this option is available only to those customers who are under EMC warranty.
- **Service center** – Access to information about your open service requests.

Conclusion

The VNXe storage system provides a consolidated platform for provisioning, managing, and monitoring application data storage within organizations and networks of small to medium sizes. Each storage system provides highly efficient iSCSI block storage and highly accessible file-based storage that enables users to work with files and data over network connections. Because these storage systems support both IP-based SAN (iSCSI) and NAS (network-attached storage) storage access, they provide storage resources for clients and hosts running a variety of operating systems and applications. VNXe provides a true multiprotocol data storage device.

Since VNXe software automatically implements best practices when you configure storage, you do not need to have detailed knowledge about storage and application technologies. VNXe systems are designed so they are easily installed, configured, and maintained by an IT generalist. Comprehensive, easy-to-understand guidance is also available through the VNXe support ecosystem.

From a hardware point of view, users can start with a small configuration based on their current requirements, and easily scale up as their requirements change. Advanced features like link aggregation, File-Level Retention, deduplication, and VMware vCenter integration are also available.

References

The following can be found on Powerlink or EMC Online Support (<https://support.emc.com>) › VNXe Product Page:

- EMC Unisphere for VNXe: Next-Generation Storage Management – A Detailed Review
- EMC VNXe Data Protection - Overview
- EMC VNXe Storage Systems – A Detailed Review
- EMC Simple Support Matrix for EMC VNXe Series
- EMC VSI for VMware vSphere: Unified Storage Management