



EMC VSI for VMware vSphere: Unified Storage Management

Version 5.6.3

Product Guide

P/N 300-012-100

REV 15

EMC²

Copyright ©2010–2014 EMC Corporation. All rights reserved. Published in the USA.

Published May 2014

EMC believes the information in this publication is accurate as 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.

EMC², EMC, and the EMC logo are registered trademarks or trademarks of EMC Corporation in the United States and other countries. All other trademarks used herein are the property of their respective owners.

For the most up-to-date regulatory document for your product line, go to EMC Online Support (<https://support.emc.com>).

CONTENTS

Chapter 1	Overview	
	Executive summary	6
	Product overview	7
	Audience.....	8
	Scope	8
	Key terms.....	9
	Related documents	10
Chapter 2	Prerequisites and System Requirements	
	Downloading Unified Storage Management.....	12
	Setting up the environment.....	12
	Environment and system requirements	13
	Celerra system requirements.....	14
	CLARiiON system requirements	14
	VMAX system requirements.....	15
	VNX system requirements	16
	VNXe system requirements.....	17
	VPLEX system requirements	17
	XtremIO system requirements	17
	VMware environment requirements.....	18
	Citrix XenDesktop environment requirements.....	18
	Loading Unified Storage Management.....	18
	Removing quotes from system path.....	19
	DHSM/ASA setup instructions.....	19
Chapter 3	Installation	
	Installing Unified Storage Management.....	24
	Installing the EMC Unified Storage Access Control Utility	24
	Common installation errors	25
Chapter 4	Procedures	
	Managing access profiles in the EMC Unified Storage Access Control Utility	28
	Adding an access profile	28
	Removing an access profile	31
	Editing an access profile	31
	Exporting an access profile.....	31
	Managing storage systems and connection brokers.....	32
	Adding storage systems or connection brokers.....	32
	Removing storage systems or connection brokers	37
	Editing credentials	37
	Provisioning storage.....	37
	Provisioning storage for an NFS datastore	37
	Using CAVA with VMware vSphere	40
	Provisioning storage on an existing NFS export.....	41
	Provisioning storage for a VMFS datastore or RDM volume on CLARiiON, VNX, VNXe, or VMAX.....	42

	Provisioning storage for a VMFS datastore or RDM volume on VPLEX	44
	Provisioning storage for a VMFS datastore or RDM volume on XtremIO	47
	Extending storage	47
	Compressing and decompressing file storage system objects.....	48
	Compressing file storage system objects.....	49
	Decompressing file storage system objects	51
	Compressing and decompressing VNX block storage objects.....	51
	Compressing and decompressing datastores	52
	Compressing and decompressing virtual machines	52
	Enabling and disabling block deduplication on VNX systems.....	53
	Prerequisites.....	54
	Enabling deduplication	54
	Disabling deduplication	55
	Cloning virtual machines.....	55
	Creating Fast Clones.....	56
	Creating Full Clones	59
	Creating Native Clones	61
	Integrating clones with VMware Horizon View.....	63
	Integrating clones with Citrix XenDesktop.....	65
	Refreshing desktops	66
	Uninstalling Unified Storage Management.....	66
	Uninstalling the EMC Unified Storage Access Control Utility	66
Chapter 5	Troubleshooting	
	Known problems and limitations.....	68
	Logs.....	68
	Technical notes	68
	EMC Sales and Customer Service contacts	69
Appendix A	Configuration File Parameters	
	Configuration file name and location.....	72
	Parameters within the configuration file	72

CHAPTER 1

Overview

This chapter presents these topics:

- ◆ Executive summary 6
- ◆ Product overview 7
- ◆ Key terms 9
- ◆ Related documents 10

Executive summary

Unified Storage Management is a feature of EMC® Virtual Storage Integrator (VSI) for VMware vSphere. It is designed to simplify administration of the following storage systems:

- ◆ EMC Celerra® network-attached storage (NAS)
- ◆ EMC CLARiiON® block
- ◆ EMC Symmetrix® VMAX®
- ◆ EMC VNX® and EMC VNXe®
- ◆ EMC VPLEX®
- ◆ EMC XtremIO™

The Unified Storage Management feature enables VMware administrators to provision new Network File System (NFS) and Virtual Machine File System (VMFS) datastores and Raw Device Mapping (RDM) volumes directly from vSphere Client.

The Unified Storage Management feature is accessed and run from within VSI.

For NFS datastores on NAS, use the feature to rapidly provision new virtual machines with Full Clones or space-efficient Fast Clones. To improve storage utilization, access deduplication features that reduce storage consumption of virtual machines. VMware administrators who manage shared NFS storage can:

- ◆ Simplify the process of creating NFS datastores in accordance with best practices
- ◆ Automatically mount the NFS datastores to one or more VMware ESX/ESXi hosts
- ◆ Reduce the amount of storage consumed by virtual machines by using compression and Fast Clone technologies
- ◆ Reduce the duration of virtual machine copies by using Full Clone technology

For VMFS datastores and RDM volumes on block storage, use the feature to provision and mount new storage based on storage pools or RAID groups and to select a tiering policy for the new storage. The feature also supports array-based compression in VMFS datastores and block compression for thin LUNs on VNX block storage systems.

For VMFS datastores on next-generation VNX systems, you can use the feature to enable block deduplication, which can dramatically reduce storage cost and improve storage efficiency.

With the Unified Storage Management feature, VMware administrators who manage shared block storage can:

- ◆ Simplify the process of creating VMFS datastores and RDM volumes in accordance with best practices
- ◆ Automatically mount the VMFS datastores and RDM volumes to one or more ESX/ESXi hosts
- ◆ Compress objects in VNX block storage systems

The EMC Unified Storage Access Control Utility (ACU) is included with the Unified Storage Management feature. The ACU is a standalone tool that allows storage administrators to create access profiles for VMware administrators and specify the storage systems, storage pools, and RAID groups that the VMware administrator can access.

Product overview

The EMC VSI for VMware vSphere: Unified Storage Management feature can provision NFS datastores on NAS and VMFS datastores, and provision RDM volumes on block storage. The feature also performs array-based compression and array-based cloning of virtual machines in NFS datastores and array-based compression in VMFS datastores and RDM volumes. The cloning functions include Full Clones (copies), Fast Clones (snaps), and Native Clones of Virtual Machine Disk (VMDK) files. The feature allows VMware administrators to manage NAS and block storage in VMware environments using the existing vSphere Client user interface.

VMware administrators can use the feature to:

- ◆ Provision new NFS and VMFS datastores, and RDM volumes
- ◆ Enable asyncmtime NFS mount option in next-generation VNX systems to reduce the latency of write operations
- ◆ Create either a striped or concatenated meta volume when provisioning large VMAX storage
- ◆ Extend existing NFS and VMFS datastores
- ◆ Compress virtual machines in NFS datastores
- ◆ Compress objects on VNX block storage systems
- ◆ Enable deduplication on next-generation VNX block systems
- ◆ Clone virtual machines in NFS and VMFS datastores:
 - Fast Clone support (limited to the same file system)
 - Full Clone support (limited to file systems on the same Data Mover)
 - Native Clone support (limited to VMFS datastores)
- ◆ Integrate with VMware Horizon View
- ◆ Integrate with Citrix XenDesktop

To provision new NFS datastores, the feature:

- ◆ Creates a file system with automatic file system extension and EMC Virtual Provisioning™
- ◆ Exports the file system using NFS
- ◆ Provides one or more ESX/ESXi servers root access to the export
- ◆ Creates an NFS datastore on the newly created NAS file system

To provision new VMFS datastores, the feature:

- ◆ Binds a Fibre Channel (FC) or iSCSI LUN in a user-specified RAID group or storage pool

- ◆ Adds the newly bound LUN to the storage group associated with the selected ESX/ESXi hosts
- ◆ Creates a VMFS datastore on the newly created block LUN

To provision new RDM volumes, the feature:

- ◆ Binds an FC or iSCSI LUN in a user-specified RAID group or storage pool
- ◆ Does not add the newly bound LUN to any storage group, so the user is able to assign it to any virtual machine

The EMC Unified Storage ACU allows storage administrators to create access profiles for the VMware administrator. To create an access profile with the ACU, the storage administrator specifies:

- ◆ The storage systems that the VMware administrator will be allowed to access
- ◆ Whether the VMware administrator will have Distributed Hierarchical Storage Management (DHSM)/Advanced Storage Access (ASA) on file storage systems
- ◆ The storage pools and RAID groups that will be available to the VMware administrator

Audience

This feature was designed for VMware administrators who manage shared NFS or VMFS storage through vSphere Client. After the storage administrator has planned and implemented the storage environment, the VMware administrator sets up the VMware virtual environment and loads EMC VSI for VMware vSphere: Unified Storage Management. Using this feature, VMware administrators can reduce storage consumption for virtual machines as they grow, allocate, or expand NFS and VMFS datastores, and provision virtual machines without the direct involvement of the storage administrator.

Scope

This guide is intended to give readers an overview of the EMC VSI for VMware vSphere: Unified Storage Management feature and its use in their environment. Topics covered in this guide include:

- ◆ Unified Storage Management overview
- ◆ VMware vSphere prerequisites and system requirements
- ◆ EMC Celerra, CLARiiON, VMAX, VNX, VNXe, VPLEX, and XtremIO prerequisites and system requirements
- ◆ Downloading and installing Unified Storage Management
- ◆ Common installation errors
- ◆ Procedures and best practices
- ◆ Troubleshooting

Key terms

- ◆ **Cluster:** A group of hosts that share resources and have a common management interface. Changes made at the cluster level impact all the hosts within the cluster.
- ◆ **Compression:** An operation performed on an object so that it uses less space on disk than it would normally occupy. In the context of Unified Storage Management, if the compressed field value is **Yes**, compression has been applied.
- ◆ **Data Mover:** In a Celerra Network Server, a cabinet component that runs its own operating system to retrieve data from a storage device and make it available to a network client. A Data Mover is also referred to as a blade.
- ◆ **Deduplication:** The process of eliminating duplicate data in block storage or file system storage. In the context of Unified Storage Management, deduplication refers to block deduplication that can detect duplicate data at the storage pool level. As a result, 8 KB data blocks within pool slices are shared by multiple LUNs or within a single LUN.
- ◆ **ESX/ESXi:** The virtualization applications that abstract processor, memory, storage, and networking resources into multiple virtual machines running side-by-side on the same physical server.
- ◆ **Fast Clone:** A method of making a thin copy of a virtual machine using Celerra, VNX, or VNXe NFS-based snapshots.
- ◆ **Full Clone:** A method of making a full copy of a virtual machine using native Celerra, VNX, or VNXe functionality.
- ◆ **LUN:** Logical Unit Number. A SCSI identifier used by Fibre Channel and iSCSI to identify a disk, a subset of a disk, or an array of disks that is layered beneath a VMFS datastore or represented by an RDM volume.
- ◆ **Native Clone:** A method of cloning virtual machines across VMFS datastores using the VMware vSphere clone API. If VAAI is supported and enabled by the storage array, the clone operations are offloaded to a back-end storage array and benefit from certain hardware acceleration.
- ◆ **NFS:** Network file system protocol, which allows a user on a client computer to access files over a network as easily as if the network devices were attached to the client computer's local disks.
- ◆ **Path:** The NFS path to the virtual machine.
- ◆ **RAID group:** A group of disks that presents itself as a single unit on which to bind LUNs. Each LUN stored in a RAID group is distributed equally among the disks in the RAID group.
- ◆ **Snap:** Also known as a snapshot, a point-in-time image of a file system that does not mirror the data on the file system, thereby using less disk space than a standard backup.
- ◆ **Space savings:** The amount of space savings realized if a volume is compressed.
- ◆ **Storage pool:** An aggregation of disk storage from which datastores can be provisioned.

- ◆ **Storage processor (SP):** In a CLARiiON array, a cabinet component that runs its own operating system and provides access to the data stored on the array. Each SP has its own IP address.
- ◆ **Tiering policy:** The method used by the array to balance cost and performance by moving data between different drive types within a storage pool. Tiering is not available for RAID group-based LUNs.
- ◆ **Virtual machine:** A software computer that, like a physical computer, runs an operating system and applications. An operating system installed on a virtual machine is called a guest operating system.
- ◆ **VMware Virtual Machine File System (VMFS):** A virtual machine file system used in VMware ESX Server software to store files in a virtualized environment.
- ◆ **VMware vStorage APIs for Array Integration (VAAI):** A set of APIs and SCSI commands that offload certain I/O-intensive functions from the ESXi host to the storage array for more efficient performance.

Related documents

The following related documents are available at EMC Online Support:

- ◆ *EMC VSI for VMware vSphere: Unified Storage Management Release Notes*
- ◆ *Using EMC Celerra Storage with VMware vSphere and VMware Infrastructure* (TechBook)
- ◆ *Using EMC CLARiiON Storage with VMware vSphere and VMware Infrastructure* (TechBook)
- ◆ *Using EMC Symmetrix Storage with VMware vSphere* (TechBook)
- ◆ *Using EMC VNX Storage with VMware vSphere* (TechBook)
- ◆ *EMC CLARiiON Integration with VMware ESX* (Applied Technology White Paper)

CHAPTER 2

Prerequisites and System Requirements

This chapter presents these topics:

- ◆ [Downloading Unified Storage Management.....](#) 12
- ◆ [Setting up the environment.....](#) 12
- ◆ [Environment and system requirements](#) 13
- ◆ [Loading Unified Storage Management.....](#) 19

Downloading Unified Storage Management

EMC VSI for VMware vSphere: Unified Storage Management is distributed as a Zip file containing a single-file installer that is available for download from <https://support.emc.com>. This software is available with EMC Celerra, CLARiiON, VNX, VNXe, VMAX, VPLEX, and XtremIO systems at no additional cost.

Setting up the environment

Install Unified Storage Management on a workstation that has vSphere Client. [“Installing Unified Storage Management” on page 24](#) provides specific instructions. After you install the feature, ensure that network connectivity exists between the following:

- ◆ NAS Data Movers and ESX/ESXi servers
- ◆ NAS Control Station and Data Movers and vSphere Client
- ◆ One or more block storage processors and the ESX/ESXi servers
- ◆ One or more block storage processors and vSphere Client
- ◆ VMware vCenter Server and vSphere Client

Figure 1 depicts the reference architecture.

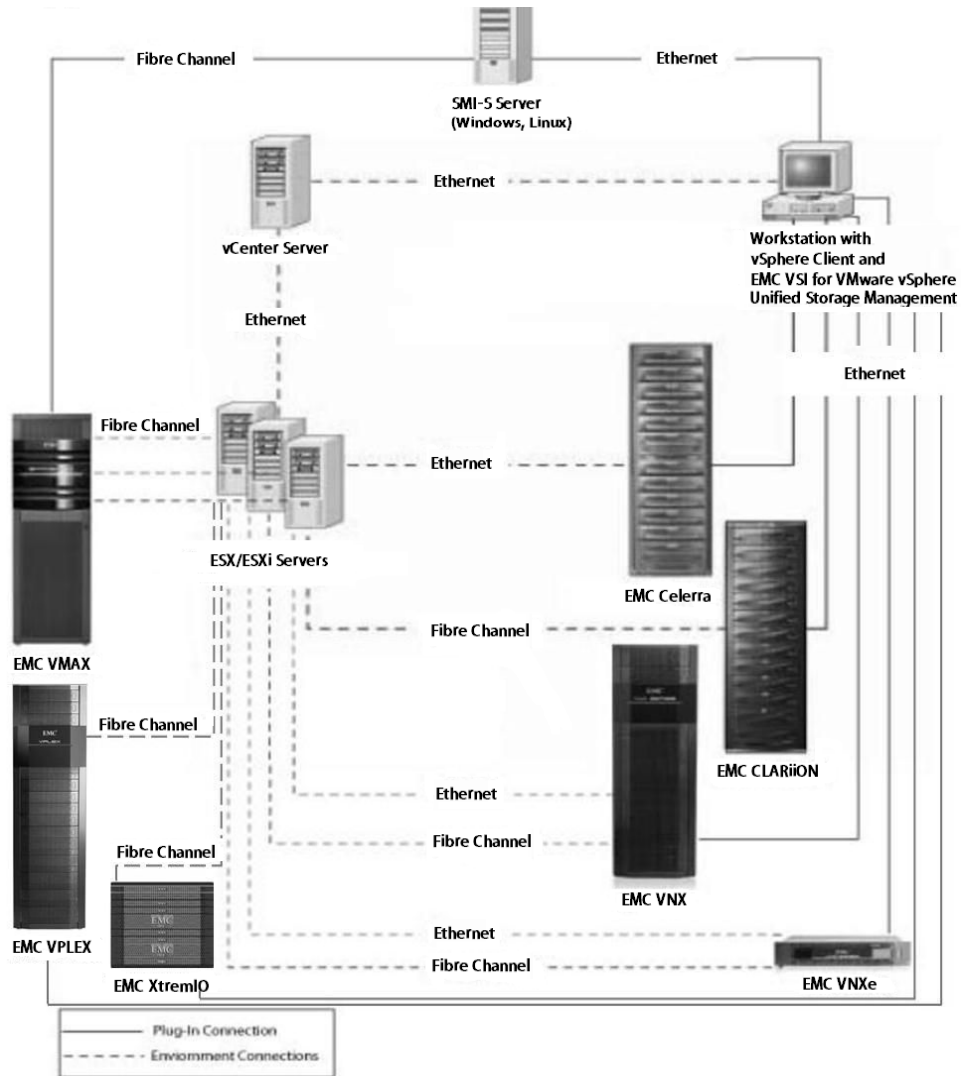


Figure 1 EMC VSI for VMware vSphere: Unified Storage Management reference architecture diagram

Environment and system requirements

Minimum system requirements for vSphere Client are specified in the *ESX/ESXi and vCenter Server Installation Guide*.

Before installing EMC VSI for VMware vSphere: Unified Storage Management, uninstall any previous installations of the EMC Celerra Plug-in for VMware or EMC Unified Storage Plug-in for VMware earlier than 4.0.0.45.

Celerra system requirements

- ◆ Celerra Network Server software version 5.6.48.7 or later is installed.
- ◆ All Celerra Data Movers have Celerra Network Server software version 5.6.48.701 or later installed. To obtain the required software, contact EMC Customer Support.
- ◆ NFS license is enabled.
- ◆ One or more storage volumes are created.
- ◆ Network ports are configured.
- ◆ Network connectivity exists between one or more Data Movers and ESX/ESXi servers.
- ◆ Network connectivity exists between the Control Station and vSphere Client on port 22 (required to provision NFS storage).
- ◆ Network connectivity exists between one or more Data Movers and vSphere Client on port 5080 (required to compress and clone virtual machines).
- ◆ Network connectivity exists between the vCenter Server and vSphere Client.
- ◆ DHSM must be set up on the Celerra to use the compression and cloning features.
- ◆ Celerra maintains the following limit: Maximum file system size = 16 TB.

CLARiiON system requirements

- ◆ FLARE[®] software version 04.29 or later is installed. To obtain the required software, contact EMC Customer Support.
- ◆ Navisphere[®] Secure CLI version 7.32 or later is installed on the vSphere Client host.

Note: If Navisphere Secure CLI is not installed on the vSphere Client host, download Navisphere CLI for Windows from EMC Online Support. The Navisphere CLI package for Windows includes Navisphere Secure CLI.

- ◆ One or more storage volumes or RAID groups are created.
- ◆ Network ports are configured.
- ◆ Network connectivity exists between one or more storage processors and ESX/ESXi servers.
- ◆ Network connectivity exists between one or more storage processors and vSphere Client on port 443 (required for provisioning VMFS datastores and RDM volumes and for setting tiering policies).
- ◆ Network connectivity exists between the vCenter Server and vSphere Client.
- ◆ Each ESX/ESXi host that will access block storage is connected to a storage group on the array.

VMAX system requirements

Note: Unified Storage Management currently supports only FC connectivity between the host and the VMAX storage system.

- ◆ EMC SMI-S Provider 4.6 is installed. SMI-S Provider 4.6 is available on EMC Online Support.
- ◆ Bidirectional connectivity requirements:
 - vCenter Server and vSphere Client
 - SMI-S server and vSphere Client (ports 5988, 5989, 5985, and 5986, configurable)
 - SMI-S server and VMAX arrays (VMAX arrays must be preconfigured in SMI-S Provider views)
 - VMAX arrays and ESXi hosts (ESXi hosts must be preconfigured in VMAX masking views)
- ◆ Storage systems:
 - VMAX: Enginuity™ 5875 or later
 - SMI-S server: EMC SMI-S Provider 4.6 or later (includes EMC Solutions Enabler 7.6)
SMI-S Provider runs on a variety of Microsoft Windows and Linux releases. For more information, see *EMC SMI-S Provider V4.6 Release Notes*.
- ◆ Symmetrix VMAX models supported:
 - VMAX 10K
 - VMAX 20K
 - VMAX 40K
- ◆ Symmetrix gatekeepers (on SMI-S server):
 - Configure six gatekeepers for each Symmetrix array accessed by the provider.
 - Only set up these gatekeepers for the host on which the SMI-S Provider is running.
- ◆ When started, the SMI-S Provider automatically discovers all Symmetrix storage arrays connected to the host on which the array provider is running. No other action is required. (You do not have to run the `symcfg discover` command.)
- ◆ Administrator authentication is required to query the EMC CIM Server.
 - Setting up authentication:
 1. Go to the URL `https://<ipaddress>:5989/ecomconfig`, and log in using the username **admin** and the password **#1Password**.
 2. Click **Add User** and create a user with the role of Administrator.
Use this newly created username to obtain access to the SMI-S Provider.

VNX system requirements

VNX for file:

- ◆ VNX for file Operating Environment (OE) software version 7.0 or later is installed.
- ◆ VNX for block OE (5.32) or later is required for Virtual Data Mover (VDM) support.
- ◆ All VNX Data Movers have VNX for file OE software version 7.0 or later installed. To obtain the required software, contact EMC Customer Support.
- ◆ NFS license is enabled.
- ◆ Storage pool is created.
- ◆ Network ports are configured.
- ◆ Network connectivity exists between one or more Data Movers and ESX/ESXi servers.
- ◆ Network connectivity exists between the Control Station and vSphere Client on port 22 (required to provision NFS storage).
- ◆ Network connectivity exists between one or more Data Movers and vSphere Client on port 5080 (required to compress and clone virtual machines).
- ◆ Network connectivity exists between the vCenter Server and vSphere Client.
- ◆ DHSM is set up to use the compression and cloning features.
- ◆ VNX for file maintains the following limit: Maximum file system size = 16 TB.

VNX for block:

- ◆ VNX for block OE software version 05.31 or later is installed. To obtain the required software, contact EMC Customer Support.
- ◆ Navisphere Secure CLI version 7.32 or later is installed on the vSphere Client host.

Note: If Navisphere Secure CLI is not installed on the vSphere Client host, download Navisphere CLI for Windows from the EMC online support website. The Navisphere CLI package for Windows includes Navisphere Secure CLI.

- ◆ Storage pool or RAID group is created.
- ◆ Network ports are configured.
- ◆ Network connectivity exists between one or more storage processors and ESX/ESXi servers.
- ◆ Network connectivity exists between one or more storage processors and vSphere Client on port 443 (required for provisioning VMFS datastores and RDM volumes and for setting tiering policies).
- ◆ Network connectivity exists between the vCenter Server and vSphere Client.
- ◆ Each ESX/ESXi host that will access block storage is connected to a storage group on the array.

VNXe system requirements

- ◆ VNXe OE software version 2.0.3 or later is installed. To obtain the required software, contact EMC Customer Support.

For ASA and the cloning and compression capabilities on VNXe, VNXe OE software version 2.2 is required.

- ◆ UEM CLI version 1.5 or later is installed on the vSphere Client host.

Notes:

- If Unisphere® CLI is not installed on the vSphere Client host, download Unisphere CLI for Windows from the EMC Online Support website.
 - Microsoft Windows 2012 does not support Unisphere CLI; therefore, the Unified Storage Management feature does not support the VNXe platform on Windows 2012.
-

- ◆ NFS license is enabled (to provision file storage).
- ◆ iSCSI license is enabled (to provision block storage).
- ◆ Network ports are configured.
- ◆ Network connectivity exists between the VNXe and one or more ESX/ESXi servers.
- ◆ Network connectivity exists between the VNXe and vSphere Client on port 443 (required for provisioning NFS and VMFS datastores and RDM volumes).
- ◆ Network connectivity exists between the vCenter Server and vSphere Client.

VPLEX system requirements

- ◆ EMC GeoSynchrony® 4.0 or later is installed.
- ◆ One or more storage volumes are created.
- ◆ One or more consistency groups are created.
- ◆ One or more storage views are created.
- ◆ Network ports are configured.
- ◆ Network connectivity exists between one or more storage processors and ESX/ESXi servers.
- ◆ Each ESX/ESXi host that will access block storage is configured on a storage view.

XtremIO system requirements

- ◆ The XtremIO system is version 2.4 or later.
- ◆ One or more storage volumes are created.
- ◆ One or more initiator groups are created.
- ◆ One or more storage views are created.
- ◆ Network ports are configured.

- ◆ Network connectivity exists between one or more storage processors and ESX/ESXi servers.

Note: Unified Storage Management currently supports only FC connectivity between the host and the XtremIO storage system.

- ◆ Each ESX/ESXi host that will access block storage is configured on a storage view.

VMware environment requirements

Note: EMC VSI for VMware vSphere does not support VMware vCenter Linked Mode.

- ◆ VMware ESXi Server 5.0, 5.1, or 5.5 is installed.
- ◆ All ESX/ESXi servers belong to a DNS domain.
- ◆ VMware vCenter Server 5.0, 5.1, or 5.5 is installed.
- ◆ VMware vSphere Client 5.0, 5.1, or 5.5 is installed.
- ◆ VMware Horizon View 5.1 or 5.2 is installed.
- ◆ The preferred locale of the Windows client that runs vSphere Client is **en_US**. However, if you have set a locale other than **en_US**, then input/output must be in ASCII characters.
- ◆ No previous installations of the EMC Celerra Plug-in for VMware or the EMC Unified Storage Plug-in for VMware earlier than 4.0.0.45 exist.
- ◆ Versions prior to 4.0.0.45 are uninstalled. (For versions 4.0.0.45 to 4.1, the installer will perform an upgrade.)
- ◆ EMC recommends uninstalling the EMC CLARiiON Plug-in for VMware if it is installed.

Citrix XenDesktop environment requirements

- ◆ Citrix XenDesktop Controller 5.6 or 7.0 is installed.
- ◆ Windows Remote Management (WinRM) is installed on the XenDesktop Controller host and on the vSphere Client host.
- ◆ WinRM service is started and configured to listen to incoming service requests.
- ◆ Windows PowerShell 2.0 is installed on the XenDesktop Controller host and on the vSphere Client host.
- ◆ A minimum of one hypervisor connection is configured on XenDesktop with an address that points to the vCenter Server where vSphere Client is registered.
- ◆ The XenDesktop Controller and vSphere Client are in the same Active Directory domain.
- ◆ Virtual Desktop Agent is installed on the virtual machine images to be cloned.

Loading Unified Storage Management

After you install Unified Storage Management on the workstation with vSphere Client, you must complete the following procedures on the workstation to load the feature and make it operational.

Removing quotes from system path

1. Right-click **My Computer**.
2. Select **Properties**.
3. Click the **Advanced** tab.
4. Click **Environment Variables**.
5. Edit the PATH variable. Change any entries that contain quotes to short notation (use the command `dir /x` to find the short notation of a folder).

Example:

```
"C:\Program Files\Common Files\emc"
```

Changes to:

```
C:\progra~1\common~1\emc
```

DHSM/ASA setup instructions

You can set up DHSM on EMC Celerra or EMC VNX in three ways. In the Unified Storage Management feature, you can run the DHSM setup script or manually enter the commands in the Celerra or VNX CLI.

On EMC VNXe storage systems, the DHSM functionality is called Advanced Storage Access (ASA). Configure ASA when adding VNXe systems to the feature.

Set up DHSM/ASA in the feature

You can set up DHSM/ASA when you are adding a Celerra, VNX, or VNXe to the feature. Refer to [“Add Celerra” on page 33](#) or [“Add VMAX” on page 34](#) for instructions on setting up DHSM; refer to [“Add VNXe” on page 35](#) for instructions on setting up ASA.

Set up DHSM using the setup script

Included with the feature package is a Perl script (`dhsmsetup.pl`) that automatically sets up DHSM on the EMC Celerra or EMC VNX. The script is usable only after it is copied to the Celerra or VNX Control Station.

1. Copy the script by downloading `pscp.exe` from the following website: <http://www.chiark.greenend.org.uk/~sgtatham/putty/download.html>
2. Run `pscp.exe` on your local machine by typing the following command:

```
pscp dhsmsetup.pl <User>@<Hostname>:<Path>
```

The variables in this executable are:

- *User* is the Control Station user.
- *Hostname* is the Control Station hostname or IP address.

- *Path* is the path on the Control Station where the script should be saved.

The Control Station prompts you for your password. This process transfers the dhsmsetup.pl script to the path that you provide on the Control Station.

Example:

```
pscp dhsmsetup.pl nasadmin@celerra.emc.com:/home/nasadmin
```

3. After the script is successfully copied, log in to the Control Station and navigate to the path where you saved the script. The script is executed with the following command:

```
perl dhsmsetup.pl <Data Mover> <User> <Host IP>
```

The variables in this script are:

- *Data Mover* is the name of the Celerra or VNX Data Mover for which you would like to configure DHSM.
- *User* is the name of the Celerra or VNX user to create that will be accessing DHSM.
- *Host IP* is the IP address of the host that will be accessing DHSM.

The system prompts you for a User ID, Group ID, Home Directory, and password for the new user. The User and Group IDs must be numeric. If you attempt to use an existing User ID, you are prompted to retry your input. The Home Directory field can be left blank if you choose. The password for the new user must be at least 6 characters.

After completing these tasks, the script displays the DHSM configuration and a message that the setup is successful.

Note: You must run the script once for each Data Mover that you intend to use with the feature. The feature requires that Data Movers in a Celerra or VNX cabinet have the same DHSM username and password.

Set up ASA when adding a VNXe to the feature

1. After you enter the platform credentials, select **Configure Advanced Storage Access**.
2. In the **Password** field, type the ASA password.
3. Click **Finish**.

Manually set up DHSM on Celerra or VNX

This procedure sets up DHSM on all Data Movers. The manual setup is a procedure for users who choose not to use the DHSM setup script or to set up DHSM in the feature GUI.

1. Create a new DHSM user:

```
/nas/sbin/server_user <server name> -add -md5 -passwd <username>
```

Example:

```
/nas/sbin/server_user server_2 -add -md5 -passwd dhsm_user
```

2. Enable digest authentication:

```
/nas/bin/server_http <server name> -modify dhsm -authentication  
digest -users <username>
```

Example:

```
/nas/bin/server_http server_2 -modify dhsm -authentication digest  
-users dhsm_user
```

3. Start DHSM service on the Data Mover:

```
/nas/bin/server_http <server name> -service dhsm -start
```

4. Add your vSphere Client IP(s) to the DHSM access list:

```
/nas/bin/server_http <server name> -append dhsm -hosts  
<clientIP1,clientIP2,clientIP3>
```

Example:

```
/nas/bin/server_http server_2 -append dhsm -hosts  
100.100.100.1,100.100.100.2
```

Note: Failure to set up DHSM on the Celerra or VNX or failure to set up ASA on VNXe disables compression and creation of Full Clones or Fast Clones.

CHAPTER 3

Installation

This chapter presents these topics:

- ◆ Installing Unified Storage Management 24
- ◆ Installing the EMC Unified Storage Access Control Utility 24
- ◆ Common installation errors 25


Installing Unified Storage Management

Note: During installation, Unified Storage Management checks to see if EMC Virtual Storage Integrator (VSI) is installed. If VSI is not detected, it is automatically installed as part of the Unified Storage Management feature installation.

1. Unzip the Zip file.
2. Double-click **emc-vsi-usm-5.6.3.x-vmware-vsphere-WINDOWS-x86.exe**.
3. Click **Next**.

Note: Navisphere CLI is not required to provision for VMAX only.

4. Select **I accept the terms in the license agreement** and click **Next**.
5. Click **Install**.
6. Click **Finish** to exit the installer.
7. Launch vSphere Client and connect to the vCenter Server.



Installing the EMC Unified Storage Access Control Utility

1. Unzip the Zip file.
2. Double-click **emc-acu-5.6.3.x-WINDOWS-x86.exe**.

3. Click **Next**.
4. Select **I accept the terms in the license agreement** and click **Next**.
5. Click **Install**.
6. Click **Finish** to exit the installer.
7. Launch the EMC Unified Storage Access Control Utility.

Common installation errors

Symptom: I receive a DHSM error when trying to connect to Celerra or VNX.

- ◆ Check to make sure the DHSM service is started.

Example:

```
/nas/bin/server_http server_http server_2 -info
```

Confirm that the **Active** field is set to **True**.

- ◆ Check to make sure the vSphere Client IP address is added to the DHSM access list.

Example:

```
/nas/bin/server_http server_http server_2 -info
```

Confirm that the **Allowed IPs** field contains the correct client IP addresses.

Symptom: The feature does not load in vSphere Client.

- ◆ Remove quotes from the Windows Environment path, as shown in the section [“Removing quotes from system path” on page 19](#). Restart the VMware vSphere Client.
- ◆ Select **Plug-ins > Manage Plug-ins** and verify that **EMC Virtual Storage Integrator (VSI) for vSphere** is enabled.

CHAPTER 4

Procedures

This chapter presents these topics:

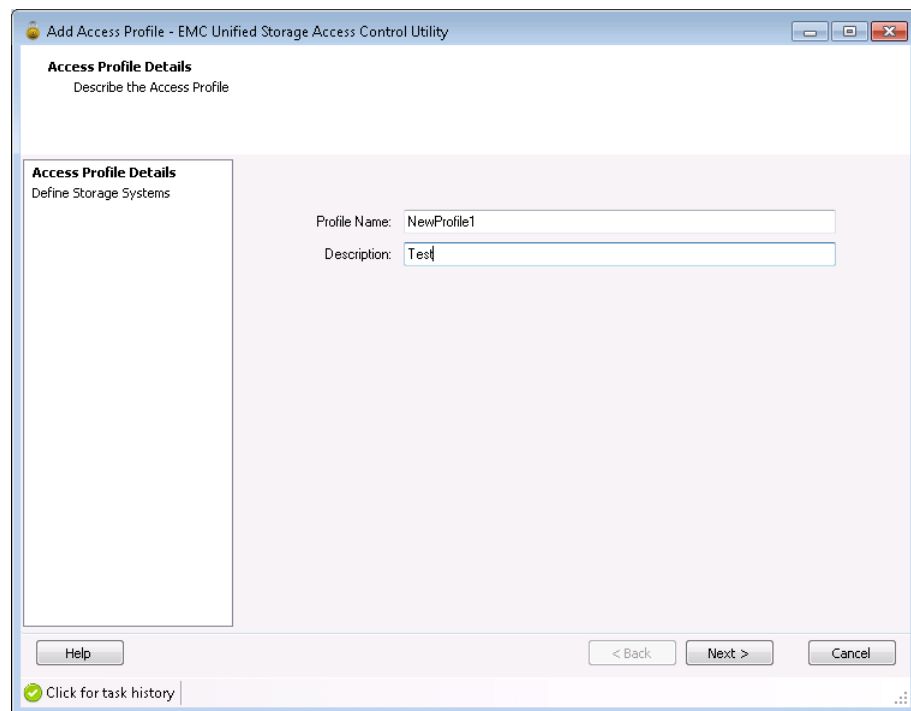
◆ Managing access profiles in the EMC Unified Storage Access Control Utility.....	28
◆ Managing storage systems and connection brokers.....	32
◆ Provisioning storage.....	37
◆ Extending storage	47
◆ Compressing and decompressing file storage system objects.....	48
◆ Compressing and decompressing VNX block storage objects.....	51
◆ Enabling and disabling block deduplication on VNX systems.....	53
◆ Cloning virtual machines	55
◆ Integrating clones with VMware Horizon View.....	63
◆ Integrating clones with Citrix XenDesktop.....	65
◆ Refreshing desktops	66
◆ Uninstalling Unified Storage Management.....	66
◆ Uninstalling the EMC Unified Storage Access Control Utility	66

Managing access profiles in the EMC Unified Storage Access Control Utility

Double-click the **Access Control Utility** icon to start the EMC Unified Storage Access Control Utility.

Adding an access profile

1. Click **Add**.
The **Add Access Profile** wizard appears.
2. In the **Profile Name** field, type a name for the access profile.
3. In the **Description** field, optionally type a description for the access profile.



4. Click **Next**.
5. Click **Add** and follow the applicable instructions to add a storage system:
 - ["Add Celerra" on page 33](#)
 - ["Add CLARiiON" on page 33](#)
 - ["Add VMAX" on page 34](#)
 - ["Add VNX" on page 34](#)
 - ["Add VNXe" on page 35](#)
 - ["Add VPLEX" on page 35](#)
 - ["Add XtremIO" on page 36](#)

Note: For EMC Celerra and EMC VNX for file, VMware administrators cannot compress, decompress, or clone datastores without DHSM access. VMware administrators must have ASA access to compress, decompress, or clone datastores on EMC VNXe.

6. At **Select Access Volumes Privileges**, select the level of storage volume access to allow on the storage system:
 - **All storage volumes**—Allows the VMware administrator to have unlimited access to all the storage volumes on the system. With this setting, the VMware administrator also will have access to any storage volumes that are created at a later time.
 - **No storage volumes**—Does not allow the VMware administrator any access to the storage volumes on the system. With this setting, the VMware administrator also will not have access to any storage volumes that are created at a later time.
 - **Selected storage volumes**—Allows the VMware administrator to have access to a specified list of storage volumes. With this setting, the VMware administrator will not have access to any storage volumes that are created at a later time.
7. Click **Finish** if you selected **All storage volumes** or **No storage volumes**, or click **Next** if you selected **Selected storage volumes**.
8. If you selected **Selected storage volumes**, choose storage volumes from the **Available Storage Volumes** table and click **Select**.

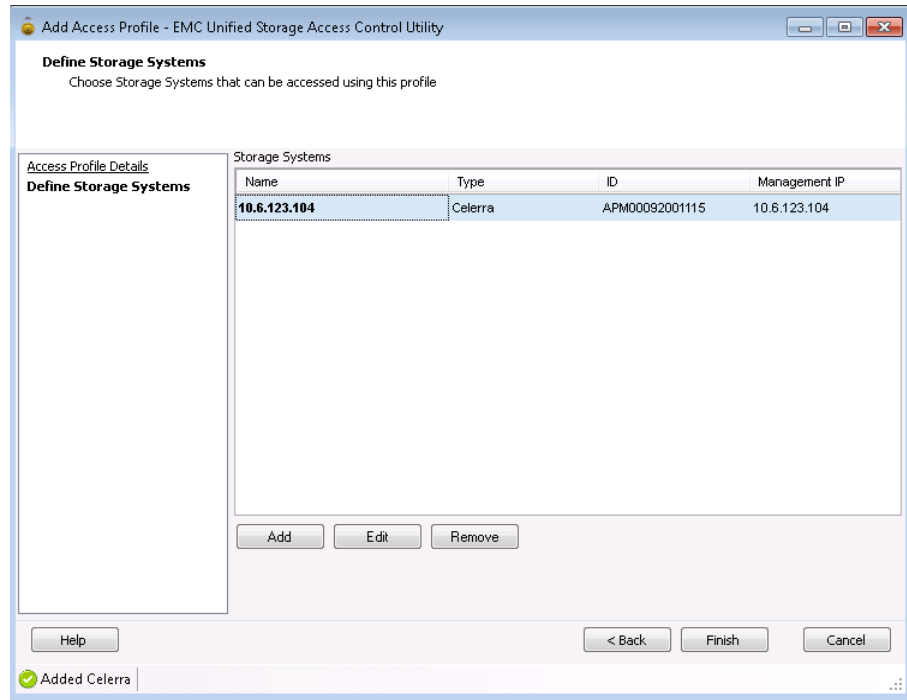
To deselect a storage volume, select it from the table and click **Remove**.

9. Click **Finish**.


The storage system now appears in the **Storage Systems** table.

- To edit a storage system, select it in the **Storage Systems** table, and click **Edit**.
- To remove a storage system, select it in the **Storage Systems** table, and click **Remove**.

- Click **Add** to add another storage system to the profile, or click **Finish** to complete the profile.



The new profile appears in the **Access Profiles** table.



Removing an access profile

1. In the **Access Profiles** table, select an access profile.
2. Click **Remove**.

Editing an access profile

To edit the credentials of a storage system or View Manager server already in the feature, follow these steps:

1. In the **Access Profiles** table, select an access profile and click **Edit**.
2. Edit the fields in the **Enter Credentials** dialog box.

Note: The **Password** field is not populated. If you do not know the password, you cannot edit the storage system.

3. Click **Next**.
4. Select the level of storage volume access to allow on the storage system:
 - **All storage volumes** allows the VMware administrator to have unlimited access to all the storage volumes on the system. With this setting, the VMware administrator will also have access to any storage volumes that are created at a later time.
 - **No storage volumes** does not allow the VMware administrator access to the storage volumes on the system. With this setting, the VMware administrator also does not have access to any storage volumes that are created at a later time.
 - **Selected storage volumes** allows the VMware administrator to have access to a specified list of storage volumes. With this setting, the VMware administrator will not have access to any storage volumes that are created at a later time.
5. Click **Finish** if you selected **All storage volumes** or **No storage volumes**, or click **Next** if you selected **Selected storage volumes**.
6. If you selected **Selected storage volumes**, choose storage volumes from the **Available Storage Volumes** table and click **Select**.

To deselect a storage volume, select the storage volume from the **Selected Storage Volume** table and click **Remove**.

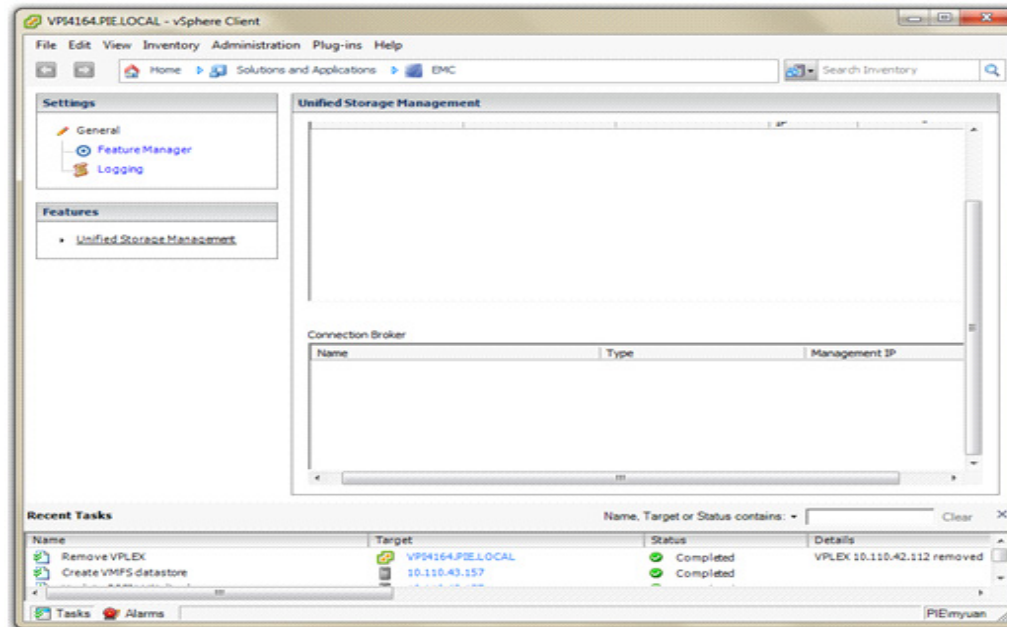
7. Click **Finish**.

Exporting an access profile

1. In the **Access Profiles** table, select an access profile.
2. Click **Export**.
3. In the **Save as** field, type a name for the file.
4. Click **...** (the browse button) to save the file in a location other than the default.
5. In the **Passphrase** field, type a passphrase for the file.
6. Click **OK**.

Managing storage systems and connection brokers

Navigate to **Home** > **Solutions and Applications** > **EMC**, and click **Unified Storage Management** in the **Features** list to manage storage systems and connection brokers.



Adding storage systems or connection brokers

1. From the EMC Unified Storage interface, click **Add** to add a storage system or connection broker to the feature.
2. Select **Enter Storage System or Connection Broker Credentials** or **Import Access Profile**, and then click **Next**.

To enter storage system credentials, continue with the next step. To import an access profile, skip to [“Import an access profile” on page 49](#).

3. Select a storage system, or select **View Manager server** or **XenDesktop controller**.
4. Click **Next**.
5. Follow the applicable instructions to add a storage system or connection broker:
 - ["Add Celerra" on page 33](#)
 - ["Add CLARiiON" on page 33](#)
 - ["Add VMAX" on page 34](#)
 - [“Add VNX” on page 34](#)
 - ["Add VNXe" on page 35](#)
 - ["Add VPLEX" on page 35](#)
 - ["Add XtremIO" on page 36](#)
 - [“Add VMware Horizon View Manager” on page 36](#)
 - [“Add Citrix XenDesktop Controller” on page 36](#)

Add Celerra

1. In the **Celerra Control Station Hostname/IP Address** field, type the hostname or IP address of the Control Station.
2. In the **Celerra Control Station Username** field, type the Control Station username.
To use LDAP authentication, type the username in this format:
`<username>@<domain_name>`
3. In the **Celerra Control Station Password** field, type the Control Station password.
4. To configure DHSM, select **Configure DHSM** and click **Next**.
 - To create a new DHSM user, select **Create New DHSM User**.
 - a. In the **DHSM Username** field, type a DHSM username.
 - b. In the **DHSM Password** field, type a DHSM password.
 - c. Select **Use next available User ID and Group ID**, or manually fill in the **User ID** and **Group ID** fields.

Note: If you manually select a User ID and Group ID that are already in use, the system will become unresponsive. Reboot the Control Station and add the storage system again.
 - d. Select **Allow access from all hosts** or **Allow access from specific hosts**.

To allow access from specific hosts, in the **Client IP Addresses** field, type the IP addresses of one or more vSphere Client hosts that will access the Celerra. Type each IP address on a separate line.
 - To use an existing DHSM user, select **Use Existing DHSM User**.
 - a. In the **DHSM Username** list box, select the DHSM username.
 - b. In the **DHSM Password** field, type the DHSM password.
 - c. Select **Allow access from all hosts** or **Allow access from specific hosts**.

To allow access from specific hosts, in the **Client IP Addresses** field, type the IP addresses of one or more vSphere Client hosts that will access the Celerra. Type each IP address on a separate line.
5. Click **Finish**.

Add CLARiiON

1. In the **EMC CLARiiON SP A** field, type the IP address of SP A.
2. In the **EMC CLARiiON SP B** field, type the IP address of SP B.
3. In the **Username** field, type the CLARiiON username.
4. In the **Password** field, type the CLARiiON password.
5. In the **Scope** list box, select a scope for CLARiiON access.

To use LDAP authentication, type the username and select **LDAP** from the list box. Do not type the domain name as part of the username.
6. Click **Finish**.

Add VMAX

Note: The SMI-S Provider is a proxy for VMAX 10K, VMAX 20K, and VMAX 40K.

1. In the **SMI-S Provider IP** field, type the SMI-S Provider IP address.
2. In the **Username** field, type the SMI-S Provider username.
3. In the **Password** field, type the SMI-S Provider password.
4. Click **Next**.
5. Select the VMAX system from the list.

Note: Multiple entries in the list might have the same Management IP. You can select only one system at a time.

6. Click **Finish**.

Add VNX

1. Select **Block**, **File**, or both.
2. Click **Next**.

To add VNX block:

1. In the **Storage Processor IP** field, type the storage processor IP address.
2. In the **Username** field, type the VNX block username.
3. In the **Password** field, type the VNX block password.
4. In the **Scope** list box, select a scope for VNX block access.

To use LDAP authentication, type the username and select **LDAP** from the list box. Do not type the domain name as part of the username.

5. Click **Finish**.

Note: To add both VNX block and VNX file, click **Next** to add the VNX file credentials.

To add VNX file:

1. In the **VNX Control Station Hostname/IP Address** field, type the Control Station hostname or IP address.
2. In the **VNX Control Station Username** field, type the Control Station username.

To use LDAP authentication, type the username in this format:

`<username>@<domain_name>`

3. In the **VNX Control Station Password** field, type the Control Station password.
4. To configure DHSM, select the **Configure DHSM** check box and click **Next**.

- To create a new DHSM user, select **Create New DHSM User**.
 - a. In the **DHSM Username** field, type a DHSM username.

- b. In the **DHSM Password** field, type a DHSM password.
- c. Select **Use next available User ID and Group ID**, or manually fill in the **User ID** and **Group ID** fields.

Note: If you manually select a User ID and Group ID that are already in use, the system will become unresponsive. Reboot the Control Station and add the storage system again.

- d. Select **Allow access from all hosts** or **Allow access from specific hosts**.

To allow access from specific hosts, in the **Client IP Addresses** field, type the IP addresses of one or more vSphere Client hosts that will access the VNX. Type each IP address on a separate line.

- e. Click **Finish**.

- To use an existing DHSM user, select **Use Existing DHSM User**.

- a. In the **DHSM Username** list box, select a DHSM username.

- b. In the **DHSM Password** field, type the DHSM password.

- c. Select **Allow access from all hosts** or **Allow access from specific hosts**.

To allow access from specific hosts, in the **Client IP Addresses** field, type the IP addresses of one or more vSphere Client hosts that will access the VNX. Type each IP address on a separate line.

5. Click **Finish**.

Add VNXe

1. In the **Management Address** field, type the management IP address.
2. In the **User** field, type the username.

To use LDAP authentication, type the username in this format:

```
<domain_name>/<username>
```

3. In the **Password** field, type the password.
4. Select **Configure Advanced Storage Access** to configure ASA.
5. In the **ASA Password** field, type the password.
6. Click **Finish**.

Add VPLEX

1. In the **Enter VPLEX Credentials** window, type the requested information:
 - VPLEX IP Address/Hostname
 - VPLEX service account
 - Password
2. Click **Finish**.

Add XtremIO

1. In the **XMS IP/Host name** field, type the IP address of the XMS server.
2. In the **XMCLI Username** field, type the XMCLI username.
3. In the **XMCLI Password** field, type the XMCLI password.
4. Click **Finish**.

Add VMware Horizon View Manager

1. In the **VMware View Manager Server** field, type the View Manager server IP address.
2. In the **Username** field, type the View Manager server domain and username in *domain\username* format.
3. In the **Password** field, type the View Manager server password.
4. Click **Finish**.

Add Citrix XenDesktop Controller

1. In the **XenDesktop Controller name/IP** field, type the name or IP address of the XenDesktop Controller.
2. In the **Domain Admin User** field, type the domain and administrator username for the XenDesktop controller in *domain\username* format.
3. In the **Password** field, type the password for the XenDesktop Controller.
4. Click **Finish**.

Note: If the system where vSphere Client is installed is not part of the same domain as the XenDesktop Controller, the add operation will fail.

Import an access profile

To import an access profile, follow these steps:

1. Click **...** (the browse button) to browse to the location where the access profile file is saved.
2. Double-click the access profile .xml file, or click **Open**.
3. In the **Passphrase** field, type the passphrase.
4. Click **Finish**.

If any storage systems specified in the access profile have already been added to the Unified Storage Management feature, the credentials and permissions from the access profile overwrite the existing credentials.

5. Click **OK** to overwrite any existing credentials, or click **Cancel** to stop the import operation.

Removing storage systems or connection brokers

1. From the EMC Unified Storage interface, select a storage system or connection broker from the list and click **Remove**.
2. Click **Yes**.

Editing credentials

1. From the EMC Unified Storage interface, select a storage system, or select **View Manager server** or **XenDesktop controller**, and click **Edit**.
2. Edit the fields in the **Add Credentials** dialog box.

Note: The **Password** field is not populated. If you do not know the password, you cannot edit the storage system.

3. Click **Finish**.

Provisioning storage

Storage provisioning prepares a NAS NFS, block VMFS file system, or block RDM volume for use by the ESX/ESXi servers. If you choose to provision storage on a cluster, folder, or datacenter, then all hosts within the selected object will mount the newly provisioned NFS datastore, VMFS datastore, or RDM volume.

Provisioning storage for an NFS datastore

To use the feature to provision NFS datastores, you can create a new NFS export or use an existing NFS export. VMware administrators who do not have the need or necessary privileges to create NAS file systems and NFS exports can use the provision storage feature to attach ESX/ESXi hosts to existing NFS exports.

NFS datastores can be provisioned on EMC Celerra, EMC VNX, and EMC VNXe platforms.

Provision storage on a new NFS export on EMC Celerra, EMC VNX, or EMC VNXe

Note: Do not create iSCSI LUNs on file systems that are exported over NFS and mounted by ESX/ESXi servers.

1. Right-click the object and select **EMC > Storage Manager > Provision Storage**.

The object can be a host, cluster, folder, or datacenter. If you choose a cluster, folder, or datacenter, then all ESX/ESXi hosts within the object will be attached to the newly provisioned storage. The **Provision Storage** wizard appears.

2. Select **Network File System**.
3. Click **Next**.
4. Select a storage system from the table.

If no storage systems are listed, click **Add**. The **Add Credentials** wizard appears. Refer to [“Add EMC Celerra” on page 39](#), [“Add EMC VNX” on page 41](#), or [“Add EMC VNXe” on page 44](#) to add a storage system.

5. Click **Next**.
6. In the **Datastore Name** field, type a name for the datastore.

Note: The /\% characters, names that contain spaces, and names that consist only of numeric characters are not allowed.

7. Click **Next**.

Continue to [step 8](#) to provision an NFS datastore on EMC VNXe. Skip to [step 16](#) to provision an NFS datastore on EMC Celerra or EMC VNX.

8. Select **Create New NFS Export**.
9. Click **Next**.
10. Select a storage volume from the table.

Note: If the storage system was added by importing an access profile from the EMC Unified Storage Access Control Utility, only storage volumes specified by the storage administrator are available for provisioning.

11. Click **Next**.
12. In the **Shared Folder Server** list box, select a shared folder server.
13. In the **Size** field, type the size for the NFS export and then select the unit of measure.
14. Select the **Enabled** checkbox if you want to enable thin provisioning. Click **Finish**.
15. Skip to [step 27 on page 40](#) to view the new NFS datastore on VNXe.

To provision an NFS datastore on EMC Celerra or EMC VNX, proceed with the next step.

16. In the **Data Mover Name** list box, select a Data Mover.
17. In the **Data Mover Interfaces** list box, select a Data Mover interface.
18. Click **Next**.
19. Select **Create New NFS Export**.
20. Click **Next**.
21. In the **Storage Pool** list box, select a storage volume.

Note: If the storage system was added by importing an access profile from the EMC Unified Storage Access Control Utility, only storage volumes specified by the storage administrator will be available for provisioning.

22. In the **Initial Capacity** field, type an initial capacity for the NFS export and then select the unit of measure from the list box to the right.
23. For thin provisioning, select **Thin Enabled**.

Note: Thin provisioning, a best practice with VMware on Celerra and VNX, is enabled by default. If you choose to use Virtual Provisioning, you must type a maximum capacity in the **Max Capacity** field.

24. If **Thin Enabled** is selected, in the **Max Capacity** field type a maximum capacity for the NFS export and then select the unit of measure from the list box to the right.
25. Click **Advanced** to see the advanced features available. You can select or deselect the following options.

Note: Advanced options are not available for VNXe systems.

- **Export Path** allows the user to export to a custom path. The feature will export the default file system path.
- **High Water Mark** specifies the file system usage threshold at which to initiate automatic file system extension. The threshold is an integer in the range of 50 to 99 percent. The default is 90 percent.
- **Direct Writes Enabled** (not applicable to next-generation VNX systems) enhances write performance to the Celerra or VNX over the NFS protocol. This mechanism allows well-formed writes to be sent directly to the disk without being cached on the server. This option is enabled by default.
- **Mtime Async Enabled**, an option that is available only if you are provisioning storage on a next-generation VNX system, reduces the latency of write operations by updating the mtime file in memory only when a data block is overwritten. The mtime file is flushed to persistent storage every sync cycle interval. This option is enabled by default.

Note: All file metadata (inode, directory, and indirect block) updates remain transactional. In addition, the mtime file update is transactional when the file operation involves other metadata changes such as file size, number of blocks, indirect blocks, and so on.

- **No Prefetch** turns the prefetch mechanism off. The prefetch mechanism performs read ahead processing for file systems. This mechanism is designed to optimize read operations of large files. Turning this mechanism off may affect performance. This option is disabled by default.
- **Virus Checking Enabled** turns on the Celerra AntiVirus Agent (CAVA). As a best practice, this check box is cleared. By disabling AV scanning when provisioning storage for virtual machines, the overall performance can be improved.
- **Export to Subnet** allows the user to export the file system to an entire subnet. By default, the feature exports the file system to each ESX/ESXi host VMkernel IP address. This option is disabled by default.
- **Set Timeout Settings** sets the NFS heartbeat settings on ESX to the best practices for VMware with Celerra or VNX. This option is enabled by default.

26. Click **Finish**.

After you click **Finish**, the Unified Storage Management feature:

- Creates a file system on the selected storage volume.
- Mounts the newly created file system on the selected Celerra or VNX Data Mover.
- Exports the newly created file system over NFS and provides root and access privileges to the ESX/ESXi hosts that will mount the NFS datastore.

- Creates the NFS datastore on the selected ESX/ESXi hosts.
- Updates the selected NFS options on the chosen ESX/ESXi hosts.

27. Click the **Summary** tab to see the newly provisioned storage.

Known issues for provisioning storage on a new NFS export

- ◆ If the user who is logged in to the NAS platform does not have sufficient privileges to create file systems or NFS exports, the provision storage operation will fail.
- ◆ If the user who is logged in to the vCenter Server does not have sufficient privileges to create NFS datastores, the provision storage operation will fail.

Using CAVA with VMware vSphere

Celerra AntiVirus Agent (CAVA) provides an antivirus solution to file-based clients using an EMC Celerra Network Server or EMC VNX. It uses the Common Internet File System (CIFS) protocol in a Microsoft Windows Server 2003, Windows 2000, or Windows NT domain. CAVA uses third-party antivirus software to identify and eliminate known viruses before they infect files on the storage system. CAVA provides benefits such as scan-on-first-read, scan-on-write, and automatic updating of virus definition files to ensure that infected files are not stored in the Celerra or VNX shared storage.

The *Using Celerra AntiVirus Agent* technical module on EMC Online Support provides further information about CAVA. The antivirus solution is only for clients running the CIFS protocol. If NFS or FTP protocols are used to move or modify files, the files are not scanned for viruses. Therefore, files accessed by ESX as part of the virtual machine deployment (as in files virtualized in virtual disks) are not scanned for viruses. Furthermore, since CAVA is a file-based solution, block-level storage that is presented to ESX from Celerra, or VNX is not scanned for viruses either.

However, files accessed by Windows virtual machines through the CIFS protocol, by using mapped network shares from Celerra or VNX, are scanned for viruses.

CAVA is most suitable for virtual machine user data that is accessed using CIFS, such as home directories and network shares. This permits you to architect a centralized solution for virus scanning, avoiding the need to scan these files locally on each virtual machine.

To ensure that the third-party antivirus software that is configured as part of CAVA will not attempt to scan virtual disk files, complete one of the following processes:

- ◆ For NFS file systems that are presented to ESX, mount the file system on the Celerra or VNX Data Mover with the `noscan` option, which instructs CAVA not to scan the file system. This is the optimal alternative, as it has CAVA focus solely on the file systems with files that should be scanned.
- ◆ Alternatively, if a file system is presented to ESX using NFS and simultaneously also to virtual machines using CIFS, then you can configure CAVA to exclude all file types that are used for file encapsulation of a virtual machine. This involves using the `excl=` parameter in the `viruschecker.conf` configuration file.

The latter alternative is less favorable, mainly because typically virtual machines should not be granted access to the encapsulated files of the virtual machines. The remainder of this section focuses on the first alternative.

Disable CAVA virus scanning on new NAS datastores

EMC recommends disabling CAVA scanning when provisioning storage for virtual machines.

In the **Advanced Options** dialog box, clear the CAVA scanning check box to disable CAVA scanning.

Testing shows that disabling antivirus scanning can improve performance. To verify the antivirus scanning settings as part of the provision storage operation, **Provision Storage > Advanced Options** is selected.

Provisioning storage on an existing NFS export

Note: Do not create iSCSI LUNs on file systems that are exported over NFS and mounted by ESX/ESXi server(s).

1. Right-click the object, which can be a host, cluster, folder, or datacenter.

Note: If you choose a cluster, folder, or datacenter, then all ESX/ESXi hosts within the object will be attached to the newly provisioned storage.

2. Select **EMC > Storage Manager**.
3. Select **Provision Storage**. The **Provision Storage** wizard appears.
4. Select **Network File System**.
5. Click **Next**.
6. Select a storage system from the table.

If no storage systems are listed, click **Add**. The **Add Credentials** wizard appears. Refer to [“Add EMC Celerra” on page 39](#), [“Add EMC VNX” on page 41](#), or [“Add EMC VNXe” on page 44](#) to add a storage system.

7. Click **Next**.
8. In the **Datastore Name** field, type a name for the datastore.
9. Click **Next**.

Continue with [step 10](#) to provision an NFS datastore on EMC VNXe. Skip to [step 14 on page 42](#) to provision an NFS datastore on EMC Celerra or EMC VNX.

10. Select **Use Existing NFS Export**.
11. Click **Next**.
12. Select an NFS export from the table.

Note: If the storage system was added by importing an access profile from the EMC Unified Storage Access Control Utility, only NFS exports on storage volumes specified by the storage administrator will be available for provisioning.

13. Click **Finish**.

To provision an NFS datastore on EMC Celerra or EMC VNX, proceed with the next step.

14. In the **Data Mover Name** list box, select a Data Mover.
15. In the **Data Mover Interfaces** list box, select a Data Mover interface.
16. Click **Next**.
17. Select **Use Existing NFS Export**.
18. Click **Next**.
19. In the **NFS Export Name** list box, select an NFS export.

Note: If the storage system was added by importing an access profile from the EMC Unified Storage Access Control Utility, only NFS exports on storage pools specified by the storage administrator will be available for provisioning.

20. Select **Edit Advanced Settings** to see the advanced features available. Here, you can select or clear the option to **Set Timeout Settings**.

Set Timeout Settings is selected by default.

Note: Advanced options are not available for an EMC VNXe.

21. Click **Finish**.

Known issues for provisioning storage on an existing NFS export

- ◆ Each ESX/ESXi host can create only one NFS datastore on each NFS export.
- ◆ When you are mounting an existing NFS export, best practices such as uncached and prefetch that are configured automatically by the **Create New NFS Export** option must be configured manually on the Celerra or VNX.

Provisioning storage for a VMFS datastore or RDM volume on CLARiiON, VNX, VNXe, or VMAX

1. Right-click the object, which can be a host, cluster, folder, or datacenter.

Note: If you choose a cluster, folder, or datacenter, then all ESX/ESXi hosts within the object will be attached to the newly provisioned storage.

2. Select **EMC > Storage Manager**.
3. Select **Provision Storage**.

The **Provision Storage** wizard appears.

4. Select **Disk/LUN**.
5. Click **Next**.
6. Select a storage array from the table.

If there are no storage systems listed, click **Add**. The **Add Credentials** wizard appears. [“Add EMC Celerra” on page 39](#), [“Add EMC VNX” on page 41](#), [“Add EMC VNXe” on page 44](#), or [“Add EMC VMAX” on page 45](#) provide instructions for adding a storage system.

7. From the table, select the storage volume where the new LUN will reside.

For VMAX, multiple entries in the table might have the same Management IP address.

Note: If the storage system was added by importing an access profile from the EMC Unified Storage Access Control Utility, only storage volumes specified by the storage administrator will be available for provisioning.

8. Click **Next**.

If you are provisioning storage on an EMC CLARiiON or EMC VNX system, skip to [step 20 on page 44](#).

9. For VNXe systems, in the **iSCSI Node** list box, select an iSCSI node and then click **Next**.
10. Choose the masking view for VMAX.
11. If prompted, select **VMFS-5** or **VMFS-3**.

Note: Do not select **VMFS-5** if the datastore is going to be accessed by ESXi hosts with an ESX version older than 5.0.

12. For VMFS-3 datastores, in the **Maximum File Size** list box, select a maximum file size.
13. Click **Next**.
14. Select **VMFS Datastore** or **RDM Volume**.

Note: Unlike VMFS datastores, which can be shared across multiple virtual machines, RDM volumes are bound to a single virtual machine and cannot be shared across multiple virtual machines. EMC recommends using VMFS datastores unless a one-to-one mapping between physical and virtual storage is required.

15. For VMFS datastores:
 - a. In the **Datastore Name** field, type a name for the datastore.
 - b. In the **Maximum File Size** list box, select a maximum file size.
16. In the **LUN Number** list box, select a LUN number.

Note: The **LUN Number** list box does not appear when provisioning block storage on VNXe.

17. In the **LUN Ownership** list box, select which storage processor will own the LUN.
18. In the **Capacity** field, type an initial capacity for the datastore and select the unit of measure from the list box.

For VMAX only, if the provisioned capacity is equal to or smaller than 240 GB, the device will be created as a TDEV device; if the provisioned capacity is larger than 240 GB, the device will be created as a meta device.

19. For VMAX systems with a provisioned capacity larger than 240 GB, select **Concatenated** or **Striped** for **Meta volume type**.

20. For CLARiiON and VNX systems, click the **Advanced** button to set a tiering policy for the selected LUN.

Note: Advanced options are not available for an EMC VNXe or EMC VMAX.

Choose from the following tiering policy options:

- **Start High Then Auto-Tier:** Sets the initial data placement to the highest available tier, with subsequent data movement controlled by auto-tier.

Note: For VNX systems, this option is supported only on block OE 5.32 or later.

- **Auto-Tier:** The software distributes the initial data placement across all drive types in the pool, to maximize spindle usage for the LUN. Subsequent data relocation is based on the LUN performance statistics such that data is relocated among tiers according to I/O activity.
- **Highest Available Tier:** Sets the preferred tier for initial data placement and subsequent data relocation (if applicable) to the highest-performing disk drives with available space.
- **Lowest Available Tier:** Sets the preferred tier for initial data placement and subsequent data relocation (if applicable) to the most cost-effective disk drives with available space.

Note: If the datastore was provisioned on a storage system that was added by importing an access profile from the EMC Unified Storage Access Control Utility, the tiering policy cannot be changed if the storage administrator did not grant access to the storage pool where the datastore resides.

21. Click **Finish**.

After you click **Finish**, the Unified Storage Management feature:

- Creates a LUN in the selected storage volume.
- Assigns the LUN to the designated volume.
- Adds the newly bound LUN to the storage group associated with the selected ESX/ESXi hosts, and provisions the LUN to the hosts over FC or iSCSI.
- Creates the VMFS datastore on the selected ESX/ESXi hosts if VMFS is chosen.

22. Click **Configuration** > **Storage** to see the newly provisioned storage.

Provisioning storage for a VMFS datastore or RDM volume on VPLEX

1. Right-click the object and select **EMC** > **Storage Manager** > **Provision Storage**.

Note: The object can be a host, cluster, folder, or datacenter. If you choose a cluster, folder, or datacenter, then all ESX/ESXi hosts within the object will be attached to the newly provisioned storage.

2. In the **Select Storage Type** window, select **Disk/LUN** and click **Next**.

3. In the **Choose Storage Array** window, select a storage system from the list.
If the VPLEX system is not listed, click **Add**. The **Add Credentials** wizard appears. [“Add VPLEX” on page 35](#) provides instructions for adding a VPLEX system.

4. In the **Choose VMFS Version** window, select **VMFS-5** or **VMFS-3**.

Note: Do not select VMFS-5 if the datastore is going to be accessed by ESXi hosts with an ESX version earlier than 5.0.

5. For VMFS-3 datastores, in the **Maximum File Size** list box, select a maximum file size.
6. Click **Next**.
7. In the **Choose Storage Details** window, make the following selections:
 - a. At **Volume Type**, select **VMFS Datastore** and provide the **Datastore Name**, or select **RDM Volume**.

Note: RDM volumes are bound to a single virtual machine and cannot be shared across multiple virtual machines, unlike VMFS datastores. EMC recommends using VMFS datastores unless a one-to-one mapping between physical and virtual storage is required.

- b. At **VPLEX Virtual Volume**, provide the following information:
 - **Preserve Thin Provisioning:** Select the checkbox to enable thin provisioning of the storage volumes when they are claimed.

Note: If the storage volumes are already claimed, enabling thin provisioning will not take effect.
 - **Use Storage At:**
 - For VPLEX Local model: Select the cluster checkbox, which then enables the **Create Local Mirror** checkbox. Select **Create Local Mirror** to create mirroring on the cluster.
 - For VPLEX Metro model: Select one of the two cluster checkboxes if you want to create a local volume on the cluster. Select both cluster checkboxes if you want to create a distributed volume. Select **Create Local Mirror** for either or both clusters to create mirroring on the specified cluster.

8. Click **Next**.

Note: **Next** is unavailable if the **Datastore Name** under **Volume Type** is null and if you have not selected at least one cluster under **Use Storage At**.

9. In the **Choose Storage Volumes** window, make the following selections:
 - a. Select a storage system from the **Storage Array** list box.

Note: The list box displays only storage systems that have available storage volumes.

- b. Select a storage volume for the first leg, first mirror, second leg, and second mirror.

For back-end volumes, you can choose from **1** (local volume, no mirrors) to **4** (distributed volume, local mirrors on both sides).

If you select more than one back-end volume, the following options are displayed for choosing subsequent volumes:

- **Same Size** (default)—Displays all available volumes equal to the smallest previously selected volume
- **Larger**—Displays all available volumes larger than the smallest previously selected volume
- **All volumes**—Displays all available volumes

Selection example:

- i. First volume: You select 50 GB.
- ii. Second volume: You select **All Volumes** and then select 10 GB for the first local mirror.

You are warned that selection of this volume will cause wasted space on other legs.
- iii. Third volume: You select **Same Size**.

You are asked to select from a list of all available 10 GB volumes.
- iv. Fourth volume: You select **Larger Volumes**.

You are asked to select from a list of all available volumes larger than 10 GB. You select a 20 GB volume and are subsequently warned that 10 GB will be wasted.

When the REST APIs are issued, an extent of 10 GB is created on each volume and a 10 GB device is created from those extents. The devices are then combined to form the 10 GB virtual volume.

10. Click **Next**.

11. Select **Storage Views** and choose from the default storage views that are listed.

The default storage views are those with a list of hosts that are equivalent to those that you selected to do the provisioning.

If storage views are not listed, click **Details** and select storage views from the table.

If the selected storage views do not contain all hosts that you selected to do the provisioning, you receive a warning message and are asked if you want to continue.

12. Click **Next**.

13. Choose a consistency group and click **Next**.

If you are creating a local volume, a consistency group is optional. However, if you are creating a distributed volume, a consistency group is required and the **Next** button is disabled if you do not make a selection.

After you click **Next**, a summary of your selections is displayed.

14. Click **Finish**.

Provisioning storage for a VMFS datastore or RDM volume on XtremIO

1. Right-click the object and select **EMC > Storage Manager > Provision Storage**.

Note: The object can be a host, cluster, folder, or datacenter. If you choose a cluster, folder, or datacenter, then all ESX/ESXi hosts within the object will be attached to the newly provisioned storage.

2. In the **Select Storage Type** window, select **Disk/LUN** and click **Next**.
3. In the **Choose Storage Array** window, select a storage system from the list.

If the XtremIO system is not listed, click **Add**. The **Add Credentials** wizard appears. Follow the instructions at [“Add XtremIO” on page 36](#) to add an XtremIO system.

4. In the **Choose VMFS Version** window, select **VMFS-5** or **VMFS-3**.
5. In the **Choose Storage Details** window, follow these steps:

- a. In the **Volume Type** panel, select the volume type.
- b. For a VMFS datastore, type a **Datastore Name**.

Note: The XtremIO datastore name must be no longer than 26 characters.

- c. In the **LUN Properties** panel, if you are provisioning a datastore on a single host, select a value for **LUN ID** or use the default value.

The **Type** is **Normal (512 LBs)**.

Note: If you are provisioning a datastore on a datacenter, folder, or cluster, USM assigns the LUN ID automatically.

- d. Designate a value for **Capacity** of the LUN.

The maximum capacity is the maximum logical size of the X-brick.

- e. Click **Finish**.

Extending storage

NFS and VMFS datastores can be extended when they start to run out of free space.

All NFS datastores can be extended, but the following restrictions apply when extending VMFS datastores:


- ◆ The datastore to be extended must be on a thick or thin LUN on an EMC CLARiiON, EMC VNX, or EMC VMAX, or on a thin LUN on EMC VNXe.
- ◆ The datastore must not be provisioned from a RAID group.
- ◆ The datastore must not be an RDM volume.
- ◆ The datastore must not be on a metaLUN.
- ◆ The datastore must not span multiple extents.

- ◆ The datastore must not be on a thick LUN on a VNXe iSCSI virtual disk.
- ◆ The datastore must not be on a striped metaLUN.

Note: If the datastore was provisioned on a storage system that was added by importing an access profile from the EMC Unified Storage Access Control Utility, the extend operation will fail if the storage administrator did not grant access to the storage pool where the datastore resides.

To extend a datastore, follow these steps:

1. From the vSphere Client home screen, click **Hosts and Clusters**.
2. Select a virtual machine that is connected to the vCenter Server.
3. Right-click the datastore to be extended, and select **EMC > Storage Manager > Extend Storage**.



4. In the **Extend Capacity by** field, type the additional capacity to add to the datastore and then select a unit of measure from the list box to the right.
5. For VNX thin NFS datastores: In the **Extend Max Capacity by** field, type the additional capacity to add to the datastore maximum capacity and then select a unit of measure.
6. Click **OK**.

Compressing and decompressing file storage system objects

This section presents information about compressing and decompressing virtual machines, hosts, clusters, datastores, folders, and datacenters on a file storage system.

Compressing file storage system objects

Compression can be performed on the following objects in vSphere:

- ◆ Virtual machine: The virtual machine will be compressed.
- ◆ Host: All virtual machines on the host will be compressed.
- ◆ Cluster: All virtual machines in the cluster will be compressed.
- ◆ Datastore: All virtual machines in the datastore will be compressed.
- ◆ Folder: All virtual machines in the folder will be compressed.
- ◆ Datacenter: All virtual machines in the datacenter will be compressed.

If the Celerra Data Deduplication or VNX File Deduplication and Compression feature is turned off for the datastore on which a compress operation is executed, the feature will automatically turn on deduplication on the underlying Celerra or VNX file system. This feature has to be enabled to allow compression to occur on the file system.

- ◆ Compression ratios observed in preliminary tests are generally between 30 and 50 percent, depending on the data that is being compressed.
- ◆ The overall performance of compressed virtual machines is within 10 percent of uncompressed virtual machines.
- ◆ Creating a Full Clone of a compressed virtual machine takes about 25 percent longer than cloning a non-compressed virtual machine.
- ◆ Enable caching on the file system that you are performing compress operations on, as a best practice.

A virtual disk that is added to a compressed virtual machine will not be compressed because it was not present when the compression operation was performed. To compress the new virtual disk, repeat the compression operation on the virtual machine. It is not necessary to decompress the virtual machine first. Only the new, uncompressed virtual disks will be affected by the compression. Virtual disks that are already compressed will be skipped.

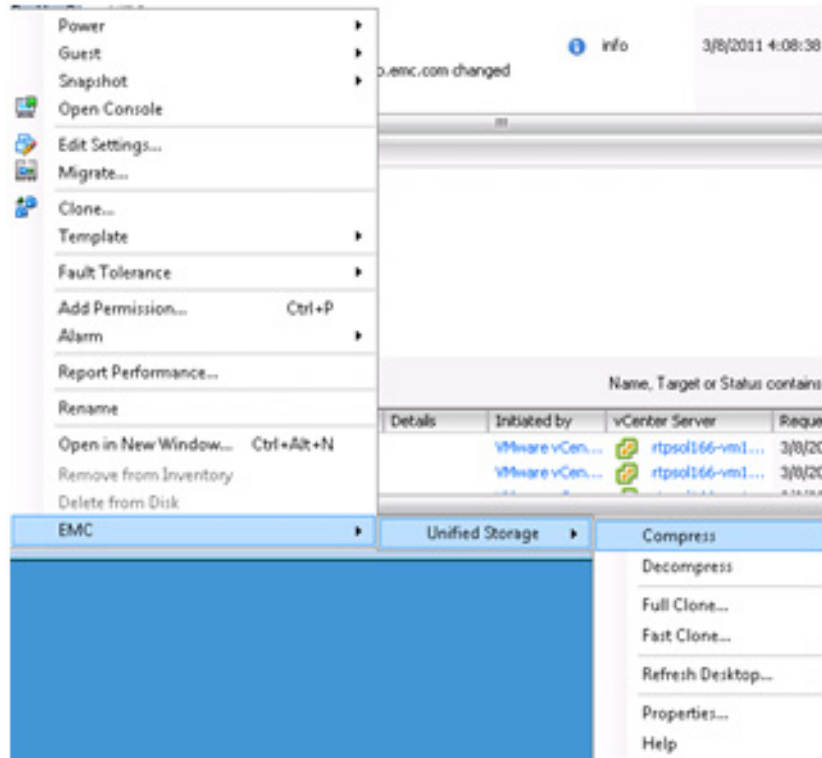
Virtual machine compression can be used when archiving gold image virtual machines. The feature does not allow compression of a virtual machine that has or has had working Fast Clones. The solution is to create a Full Clone of the master virtual machine that has Fast Clones to create a copy of the virtual machine, which can then be compressed and archived.

Compressing a virtual machine decreases the size on disk beyond its nominal size. If you choose to compress a host, cluster, folder, datastore, or datacenter, then the VMDK files associated with all virtual machines within the selected object will be compressed by the NAS platform. The feature does not compress VSWP files.

Note: The compression operation will fail if the storage administrator did not provide DHSM credentials. Compression on VNXe requires ASA access.


Compress a virtual machine

1. Right-click the virtual machine to be compressed, and select **EMC > Storage Manager > Compress**.



Note: If Celerra is not found, the message “Celerra not found for <virtual machine name>” appears. Navigate to **Home > Solutions and Applications > EMC Celerra** to add the Celerra and try again.

You can follow the compress task in the **Recent Tasks** pane at the bottom of the page.



2. Right-click **EMC > Storage Manager** and select **Properties** to show space usage before and after compression.

Compress a cluster, host, folder, datastore, or datacenter

1. Right-click the object to be compressed and select **EMC > Storage Manager > Compress**.

The following warning is displayed: **All VMs will be compressed. Compression of VMs residing on the VNX block datastore will cause compression of the VNX block datastore. Do you want to continue?**

2. Click **Yes** to compress all the virtual machines associated with the object.

The vSphere task panel displays the status and details of the Unified Storage Management log file.

Known issues for compression

- ◆ The compress operation cannot be done on a Version file (a base file that has or has had Fast Clones).
- ◆ The compress operation cannot be done on a Branch File (a Fast Clone).
- ◆ If compression is run on a file that resides on a file system without AutoExtend enabled and with less than 1 MB of available space, the compress operation fails with a “No Space” error. You are given the option to extend the file system for a successful compression.

Decompressing file storage system objects

Decompressing a virtual machine restores the size on disk to its nominal size before it was compressed. If you decompress a host, cluster, folder, datacenter, or datastore, then the files associated with all virtual machines within the selected object will be decompressed by the NAS platform.

Note: The decompression operation fails if the storage administrator did not provide DHSM credentials. Decompression on VNXe requires ASA access.

Decompress a cluster, host, folder, datastore, or datacenter

Note: When you are decompressing one or more virtual machines, the NFS datastore must have adequate storage capacity. If sufficient space is not available, you are given the option to extend the file system for a successful decompression.

1. Right-click the compressed object.
2. Select **EMC > Storage Manager**.
3. Select **Decompress**.

Note: You can follow the decompression task in the **Recent Tasks** pane at the bottom of the page.

Compressing and decompressing VNX block storage objects

Compression on the block side differs from file-side compression. File compression compresses the virtual machine disk (VMDK), while block compression compresses the whole datastore or LUN. With VSI Unified Storage Management, you can use vSphere to compress virtual machines and datastores that reside on VNX block storage systems.

Note: For this release, VSI Unified Storage Management supports block compression only of thin LUNs on VNX block storage systems.

Compressing and decompressing datastores

Controls for compressing or decompressing datastores on VNX block storage systems are available in the datastore property page. If **Turn on compression** is selected, the datastore is compressed.

Compress a datastore

To compress a datastore, from the **Block VNX Properties** dialog box, select **Turn on compression** and then click **Apply** or **OK**.

Note: Clicking **Apply** applies the setting and leaves the dialog box open. Clicking **OK** applies the setting and closes the dialog box.

If compression is not supported on the datastore, an error message is displayed.

The task panel of vSphere displays the compression status and error messages.

After compressing the datastore, the consumed capacity of the datastore decreases.

Decompress a datastore

To decompress the datastore, from the datastore **Properties** dialog box, clear **Turn on compression** and then click **Apply** or **OK**.

The decompression status is displayed in the datastore **Properties** dialog box.

Note: You can also compress or decompress objects on the block storage system from the **EMC > Storage Manager** menu.

Known limitations for datastore compression

VSI 5.6.3 only supports compression of block datastores that are hosting virtual machines with OS VMDK. Other disks will not be compressed.

Compressing and decompressing virtual machines

Controls for compressing or decompressing virtual machines on a VNX block storage system are available in the virtual machine **Properties** dialog box. If **Turn on compression** is selected, the virtual machine is compressed.

To compress or decompress a virtual machine, follow these steps:

1. From the **Virtual machine properties** dialog box, select or clear **Turn on compression**, and then click **OK** or **Apply**.
2. In the dialog box that appears, click **Yes** to compress or decompress the virtual machine on the backend storage.

Notes:

- ◆ Compressing a virtual machine on the VNX block also compresses the entire datastore.
- ◆ You can also compress or decompress the virtual machine from the **EMC > Storage Manager** menu.

- ◆ When the compression operation is committed successfully, the progress is displayed in the vSphere task panel. If the compression operation fails, status and error messages are displayed in the vSphere task panel.
- ◆ When a virtual machine is being compressed or decompressed, **Turn on compression** and **Apply** in the virtual machine property page are disabled and the status is displayed.

Block compression and decompression troubleshooting

Table 1 lists errors that can occur when compressing or decompressing objects on VNX block storage systems with possible solutions:

Table 1 Block compression and decompression troubleshooting

Problem	Solution
System internal error message is displayed.	VSI Unified Storage Management interacts with other modules and libraries, which could be unstable during development. To avoid excessive error messages, all unexpected exceptions from external logic are caught and displayed in the GUI as System internal error , while the detailed exception information is recorded in the log file, EMC.VSI.VSphere4.Features.USD.UnifiedStorageManagement.txt , for debugging.
The solution to the problem does not appear in the plug-in front end.	To check for errors that are not found in the plug-in front end, use the Unisphere client.
Compression of an object fails.	Only thin LUNs on the VNX block can be compressed.
Compression of the block datastore fails.	VSI 5.6.3 only supports compression of block datastores that are hosting virtual machines with OS on VMDK. Other disks will not be compressed.
Decompression fails.	Decompression pauses when pool consumption has reached the system-defined threshold.
The datastore is already compressed or The datastore is not compressed error message is displayed.	Trying to compress a datastore that is already compressed or trying to decompress a datastore that is not compressed causes an error message.

Enabling and disabling block deduplication on VNX systems

The Unified Storage Management feature of VSI supports block deduplication on next-generation VNX systems. You can enable block deduplication at the feature, pool, or LUN level.

The following restrictions apply:

- ◆ Deduplication is available only for pool LUNs on a system that has the deduplication feature enabled.
- ◆ You can enable either deduplication or compression, but not both.

- ◆ If you enable or disable deduplication on a pool LUN that has VNX Snapshots associated with it, all snapshots are deleted when the operation is completed.
- ◆ You cannot set tiering policies for a deduplicated LUN.
- ◆ You can report savings from deduplication at the pool level but not at the LUN level.

Prerequisites

Before you can enable deduplication, prepare your environment as follows:

1. Install Deduplication Enabler.
2. For thin provisioning, install Virtual Provisioning Enabler.
3. Create a VNX block storage array.
4. Create a datastore by provisioning storage on a host:
 - Storage Type: LUN
 - Storage Array: <array created in the previous step>
 - Storage Pool: Pool
 - File System Version: VMFS-3 or VMFS-5
 - Datastore Type: VMFS
5. Create the virtual machine by selecting the datastore you created in the previous step.

Enabling deduplication

You can enable deduplication as follows:

- ◆ On a datastore or virtual machine via the **Properties** panel
- ◆ On various objects, including a datastore, virtual machine, host, host cluster, datacenter, and folder, via the context menu

On a datastore or virtual machine via the Properties panel

1. In the **Properties** menu of the datastore or virtual machine object, select **Turn on deduplication** and then click **Apply**.

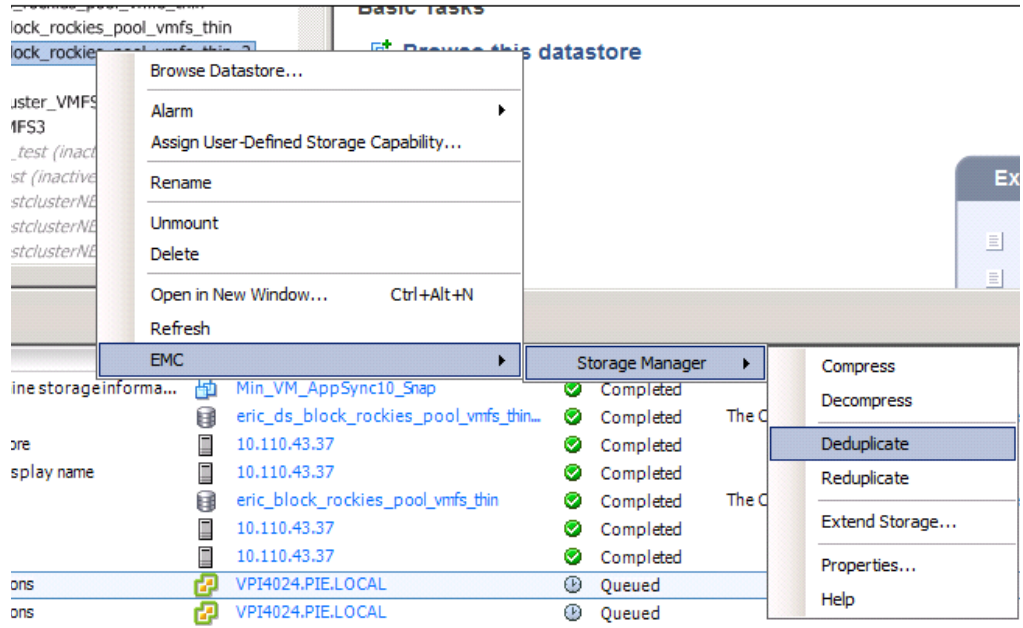
Messages are displayed to warn you that you are implementing deduplication on the LUN and that any VNX Snapshots will be deleted.

2. To proceed, click **Yes** in response to each warning.

On various objects via the context menu

Right-click the object and select **EMC > Storage Manager > Deduplicate**.

This action retrieves all virtual machines residing on the selected object and applies deduplication on each of them, one by one.



Disabling deduplication

You can disable deduplication as follows:

- ◆ On a datastore or virtual machine via the **Properties** panel: Clear **Turn on deduplication** and then click **Apply**.
- ◆ On various objects, including a datastore, virtual machine, host, host cluster, datacenter, and folder: Right-click the object and select **EMC > Storage Manager > Reduplicate**.

This action retrieves all virtual machines residing on the selected object and applies reduplication on each of them, one by one.

Cloning virtual machines

This section describes the Fast Clone, Full Clone, and Native Clone operations for virtual machines.

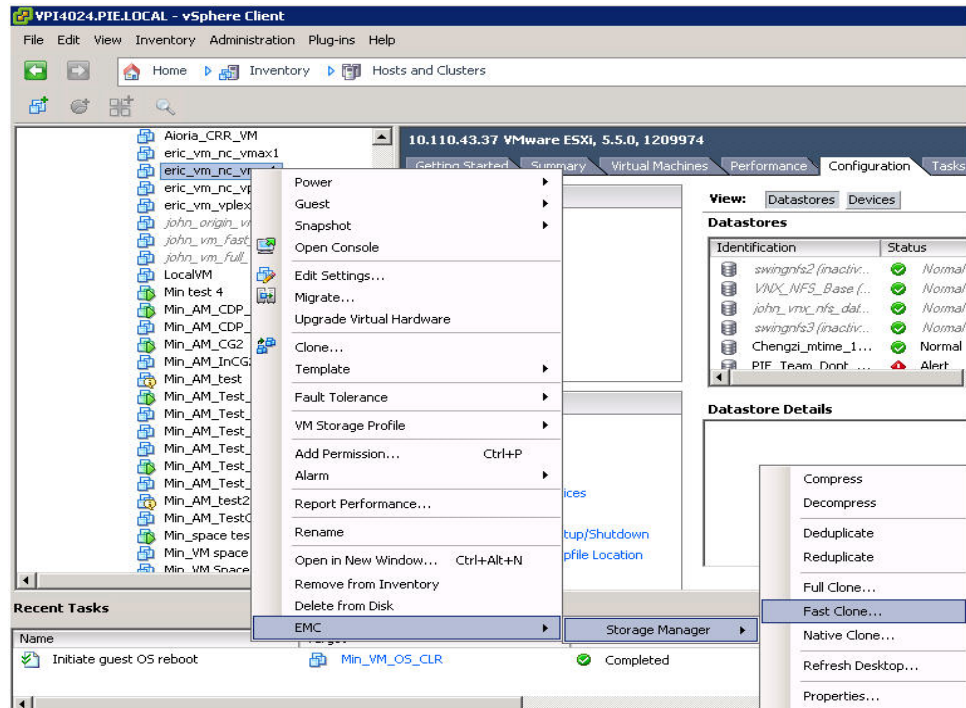
Creating Fast Clones

Fast Clone is a NAS feature that creates a file-based snapshot of a virtual machine that maintains a relationship with its parent virtual machine.

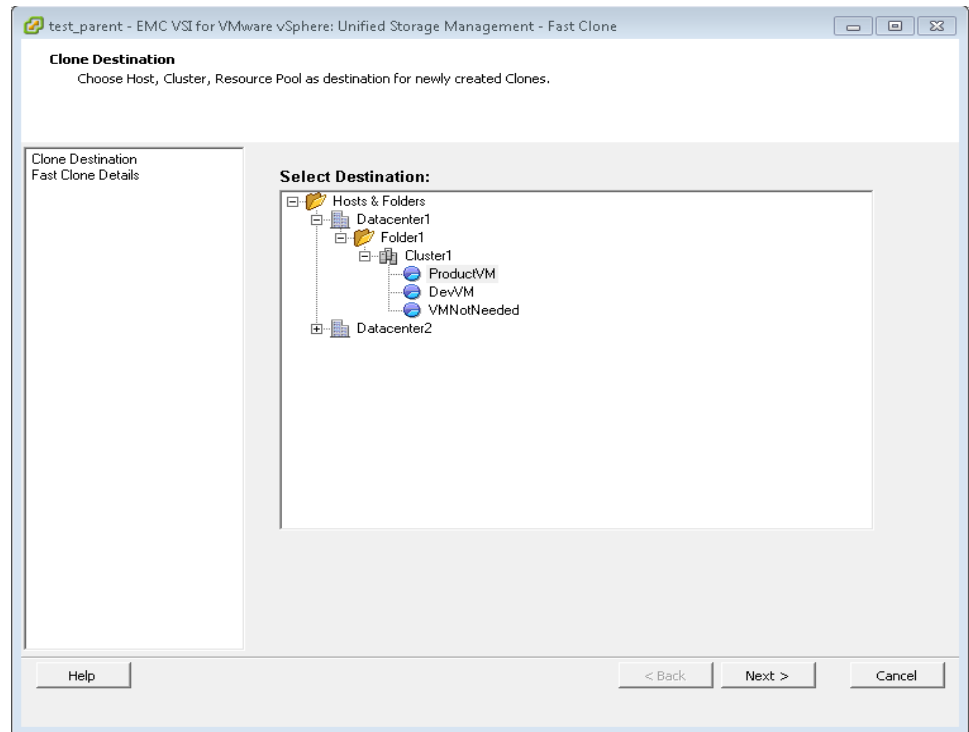
Notes:

- ◆ Once a virtual machine has Fast Clones, it can no longer be compressed. Fast Clones are created in the same NFS datastore as the parent virtual machine.
- ◆ You can only create Fast Clones from uncompressed virtual machines.
- ◆ The Fast Clone operation will fail if the storage administrator does not provide DHSM credentials. Creating Fast Clones on VNXe requires ASA access.

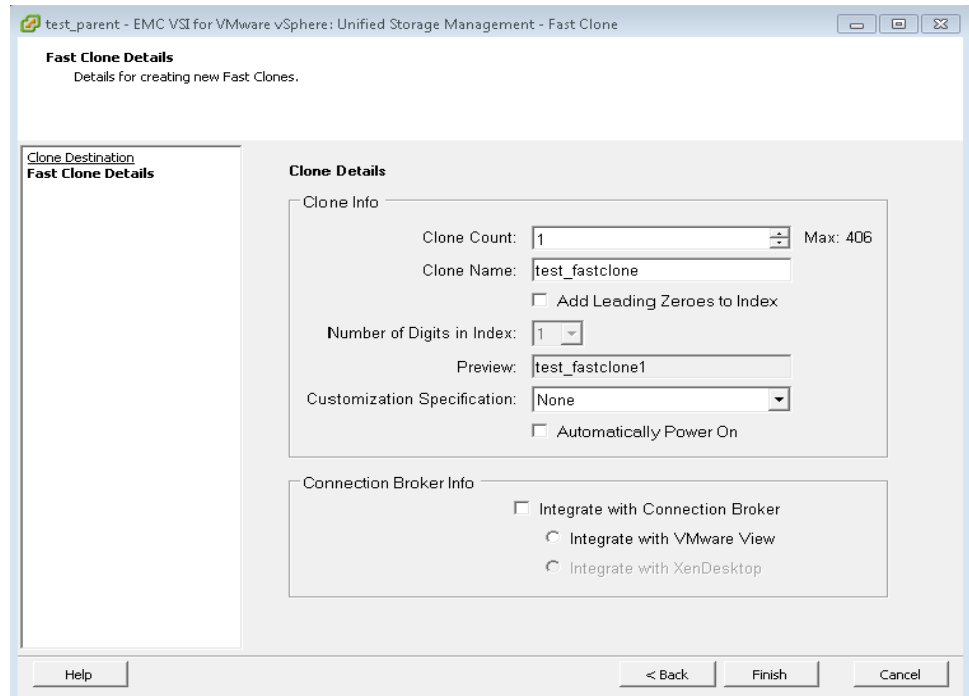
1. Right-click the virtual machine and select **EMC > Storage Manager > Fast Clone**.



2. Select the destination for the clone, and then click **Next**.




3. In the **Clone Details** window, complete the following fields:
 - **Clone Count** is the number of clones to create.
If one clone is created, the name will be the Clone Name you specify. If multiple clones are created, they will be numbered. For example, virtual machine X would have clones called X00001, X00002, X00003, and so on.
 - **Max** (to the right of the **Clone Count** field) is the maximum number of clones that can be created as determined by the number of CPU cores on the selected destination. By default, the maximum number of virtual machines per core is 12.
 - **Clone Name** is the name you specify for a clone.
 - **Add Leading Zeroes to Index** check box allows users to decide whether the clones will have leading zeroes in the file names.
 - **Number of Digits in Index** field specifies the total number of digits that will be appended to the end of the clone name.
 - **Preview** field shows the names of the clones to be created.
 - **Customization Specification** provides a list of all customized definitions from the customization specifications manager.
 - **Automatically Power On** check box automatically powers on the virtual machine clones.
 - **Integrate with Connection Broker** check box allows users to integrate desktops with VMware View Manager or Citrix XenDesktop.



4. Click **Finish**.

Note: You cannot compress a Fast Cloned virtual machine or its parent.

You can follow the progress of the Fast Clone process in the **Recent Tasks** pane at the bottom of the page.



Known issues for creating Fast Clones

The following are known issues for creating Fast Clones of virtual machines:

- ◆ You cannot create a Fast Clone of a virtual machine that is a Fast Clone itself.
- ◆ You can create Fast Clones of a gold image virtual machine only on the same file system. Fast cloning across file systems is not supported.
- ◆ A gold image virtual machine that has Fast Clones cannot be deleted. Using the **Delete from Disk** option in vCenter on a virtual machine of this nature simply removes the virtual machine from inventory; the VMDK file associated with the virtual machine is retained and all other files (for example SWAP, VMX, and so on) are deleted.
- ◆ A master virtual machine that has Fast Clones cannot be compressed.

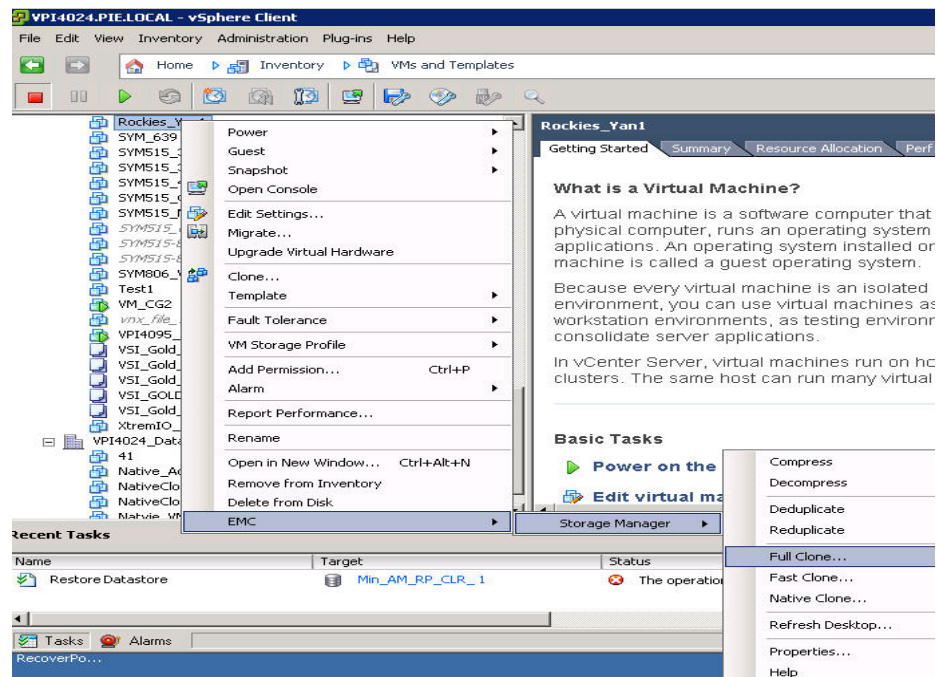
Creating Full Clones

The Full Clone feature creates a complete and independent copy of a virtual machine.

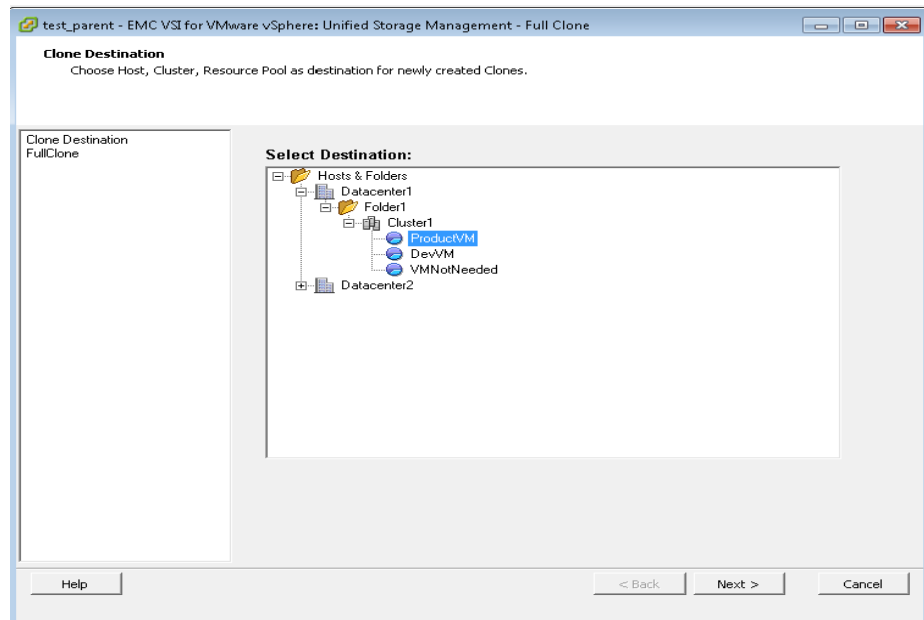
Notes:

- ◆ You can only create Full Clones on file systems on the same Data Mover.
- ◆ File systems with auto extend enabled are automatically extended during the Full Clone task if space runs out.
- ◆ The Full Clone operation fails if the storage administrator does not provide DHSM credentials.
- ◆ Creating Full Clones on VNXe requires ASA access.

1. Right-click the virtual machine to be cloned, and select **EMC > Storage Manager > Full Clone**.



2. Select the destination for the Full Clone and then click **Next**.



3. In the **Full Clone Details** window, complete the following fields:

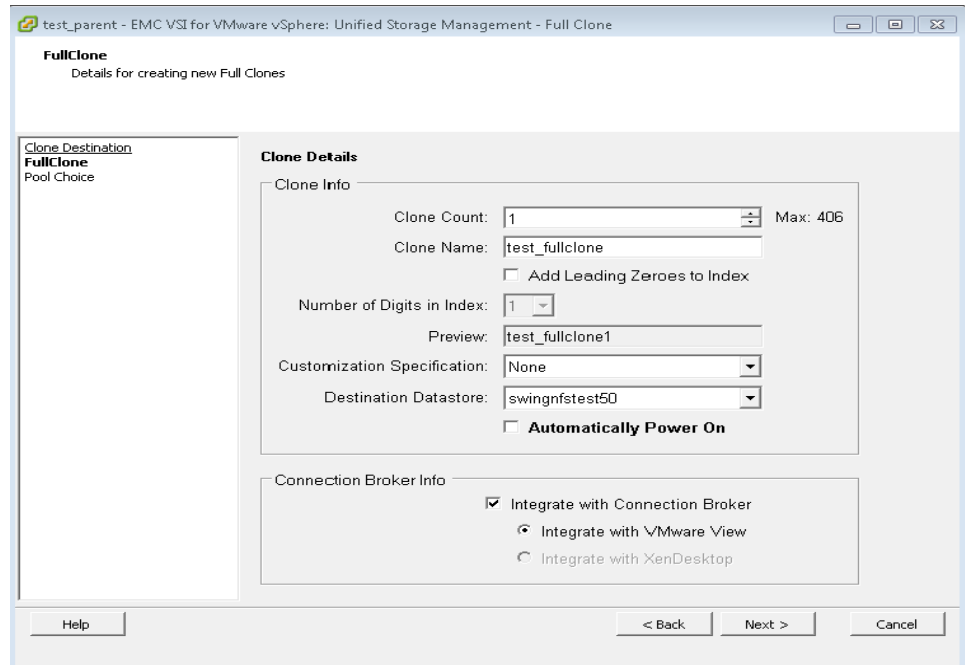
- **Clone Count** is the number of clones to create.

If one clone is created, the name will be the Clone Name you specify. If multiple clones are created, they will be numbered. For example, virtual machine X would have clones called X00001, X00002, X00003, and so on.

Note: **Max** (to the right of the **Clone Count** field) is the maximum number of clones that can be created as determined by the number of CPU cores on the selected destination. By default, the maximum number of virtual machines per core is 12.

- **Clone Name** is the name you choose for a clone.
- **Add Leading Zeroes to Index** allows users to decide whether the clones will have leading zeroes in the file names.
- **Number of Digits in Index** specifies the total number of digits that will be appended to the end of the clone name.
- **Preview** shows the names of the clones that will be created.
- **Customization Specification** provides a list of all customized definitions from the customization specifications manager.
- **Destination Datastore** is the target to store the clones.
- **Automatically Power On** automatically powers on the virtual machine clone.

- **Integrate with Connection Broker** allows users to integrate desktops with VMware View Manager or Citrix XenDesktop.



You can follow the progress of the Full Clone in the **Recent Tasks** pane at the bottom of the page.

Recent Tasks			
Name	Target	Status	Details
Full Clone	DauidsTest	In Progress	[DevTest] /full_clones/clone0001

4. Click **Finish**.

After each Full Clone is completed, the virtual machine is reconfigured.

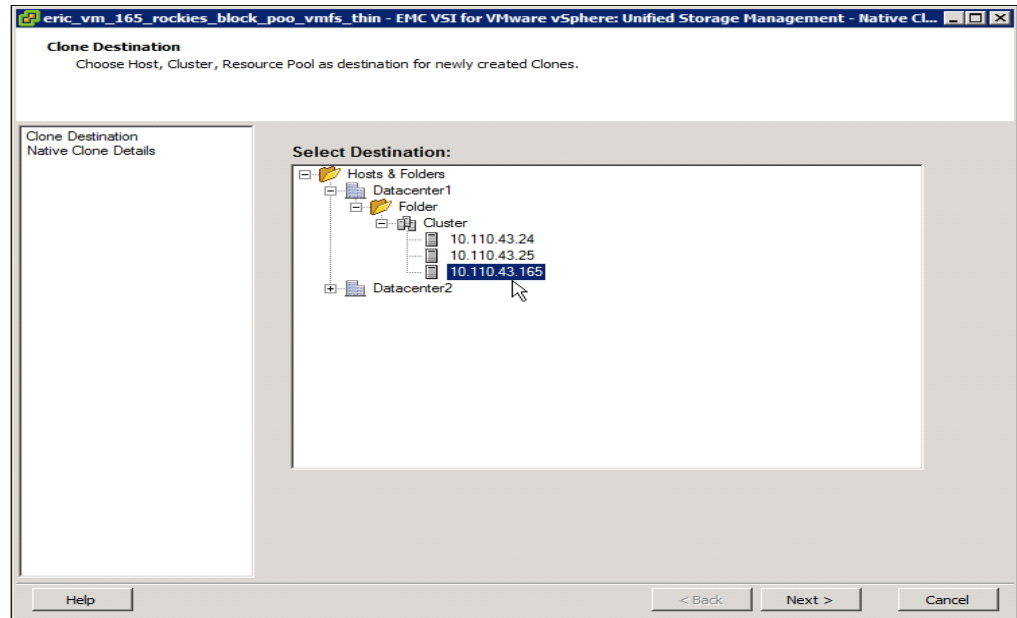
Creating Native Clones

The Native Clone feature uses the VMware Native Clone API to create a clone of a virtual machine in a VMFS datastore.

Note: Datastores with auto extend enabled are automatically extended during the Native Clone task if space runs out. However, VSI does not currently support automatic extensions for the following: Striped VMAX datastores, striped VPLEX datastores, striped XtremIO datastores, and VNX datastores residing on a VNX RAID group.

1. Right-click the virtual machine to be cloned and select **EMC > Storage Manager > Native Clone**.

2. Select a host, cluster, or resource pool as the destination for the new clone, and click **Next**.



3. In the **Native Clone Details** window, complete the following fields:

- **Clone Count** is the number of clones to create.

If one clone is created, the name will be the Clone Name you specify. If multiple clones are created, they will be numbered. For example, virtual machine X would have clones called X00001, X00002, X00003, and so on.

Max (to the right of the **Clone Count** field) is the maximum number of clones that can be created as determined by the number of CPU cores on the selected destination. By default, the maximum number of virtual machines per core is 12.

- **Clone Name** is the name you choose for a clone.
- **Add Leading Zeroes to Index** allows you to decide whether the clones will have leading zeroes in the file names.
- **Number of Digits in Index** specifies the total number of digits that will be appended to the end of the clone name.
- **Preview** shows the name of the clones that will be created.
- **Customization Specification** provides a list of all customized definitions from the customization specifications manager.
- **Destination Datastore** is the target to store the clones.

The available datastores are the VMFS datastores residing on the destination you chose in the previous step.

- **Filtered** allows you to limit the available datastores to those that are in same registered storage system in which the source virtual machine is residing, providing the clone with the benefit of accelerated performance.

If you select **Filtered**, a warning message is displayed, advising you that the filtering process will take several minutes.

- **Automatically Power On** automatically powers on the virtual machine clone.
- **Integrate with Connection Broker** allows users to integrate desktops with VMware Horizon View or Citrix XenDesktop.

4. Click **Next**.

Integrating clones with VMware Horizon View

Clone integration with VMware Horizon View allows you to add desktops to Horizon View pools.

1. Create the clone as follows:
 - **Fast Clones:** Follow the steps in [“Creating Fast Clones” on page 56](#) until you reach the **Fast Clone Details** dialog box.
 - **Full Clones:** Follow the steps in [“Creating Full Clones” on page 59](#) until you reach the **Full Clone Details** dialog box.
 - **Native Clones:** Follow the steps in [“Creating Native Clones” on page 61](#) until you reach the **Native Clone Details** dialog box.
2. Select **Integrate with Connection Broker** and then **Integrate with VMware View** to add desktops to Horizon View pools.
3. In the **Clone Name** field, type a name for the clone.

4. Optionally, select **Add Leading Zeroes to Index**.
5. Click **Next**.
6. In the **VMware View Server** list box, select a VMware View Server.
7. Select **Add VMs to a new pool** or **Add VMs to Existing Pool**.
 - The **Add VMs to a New Pool** selection allows you to add the virtual machines to a new pool that you create. If you choose this option, you must create a new pool.
 - The **Add VMs to Existing Pool** selection allows you to add the virtual machines to a pool that has already been created.
8. Click **Next**.
9. For a new pool, in the **Unique ID** field, type a unique pool name.
10. Optionally, in the **Display Name** field, type a display name.
11. In the **Desktop Persistence** list box, select **Persistent** or **Non-Persistent**.
12. Click **Next**.
13. Complete the new pool settings.
 - The **When the virtual machine is not in use** list box provides the following options:
 - Do nothing (virtual machine remains powered on)
 - Always on (ensures virtual machine is always powered on)
 - Suspend VM
 - Power on VM
 - The **Automatic logoff after disconnect** list box provides the following options:
 - Immediately
 - Never
 - After... x minutes after disconnect
 - The **Allow users to reset their desktop** checkbox allows users to reset their own desktops.
 - The **Default display protocol** field allows users to select either **Microsoft RDP** or **Teradici PC over IP**.
 - The **Adobe Flash quality** list box allows you to select **Low**, **Medium**, or **High** quality.
 - The **Adobe Flash throttling** list box allows you to select **Disabled**, **Conservative**, **Moderate**, or **Aggressive** levels of Adobe Flash throttling.
14. Click **Finish**.

Integrating clones with Citrix XenDesktop

Clone integration with XenDesktop allows you to add desktops to XenDesktop machine catalogs.

1. Create the clone as follows:
 - Fast Clones: Follow the steps in “[Creating Fast Clones](#)” on page 56 until you reach the **Fast Clone Details** dialog box.
 - Full Clones: Follow the steps in “[Creating Full Clones](#)” on page 59 until you reach the **Full Clone Details** dialog box.
 - Native Clones: Follow the steps in “[Creating Native Clones](#)” on page 61 until you reach the **Native Clone Details** dialog box.
2. Select **Integrate with Connection Broker**.
3. Select **Integrate with XenDesktop**.
4. Click **Next**.
5. In the **XenDesktop Controller** list box, select the IP address of a XenDesktop Controller.
6. Select **Add VMs to a New Machine Catalog** or **Add VMs to an Existing Machine Catalog**.
7. To use an existing machine catalog, select the machine catalog from the list box.
8. Select **Add VMs to a New Desktop Group** or **Add VMs to an Existing Desktop Group**.
9. To use an existing desktop group, select the desktop group from the list box.
10. Click **Next** or **Finish**.
11. To create a new machine catalog, in the **Machine Catalog Name** field type a name for the new machine catalog.

Note: The Unified Storage Management feature will reject a machine catalog name that is the same as a machine catalog that already exists, is longer than 64 characters, or uses the following prohibited characters: \/:#.*?=<>[](){}.

12. In the **Description** field, optionally type a description for the new machine catalog.
13. Click **Next** or **Finish**.
14. To create a new desktop group, in the **Desktop Group Name** field type a name for the new desktop group.

Note: The Unified Storage Management feature will reject a desktop group name that is the same as a desktop group that already exists, is longer than 64 characters, or uses the following prohibited characters: \/:#.*?=<>[](){}.

15. In the **Display Name** field, optionally type a display name for the new desktop group.

Note: The Unified Storage Management feature will reject a desktop group display name that uses the following prohibited characters: \/:#.*?=<>[](){}.

16. In the **Description** field, optionally type a description for the new group.
17. Click **Finish**.

Refreshing desktops

The refresh operation can be performed on the following objects in vSphere:

- ◆ Virtual machine—Fast Cloned virtual machine will be refreshed.
- ◆ Host—All Fast Cloned virtual machines on the host will be refreshed.
- ◆ Cluster—All Fast Cloned virtual machines in the cluster will be refreshed.
- ◆ Folder—All Fast Cloned virtual machines in the folder will be refreshed.
- ◆ Datacenter—All Fast Cloned virtual machines in the datacenter will be refreshed.

To refresh an object:

1. Right-click the object to be refreshed.
2. Select **EMC > Storage Manager > Refresh Desktop**.

Uninstalling Unified Storage Management

To uninstall the Unified Storage Management feature from the VSI client, complete the following steps.

Note: Uninstalling the Unified Storage Management does not uninstall the VSI framework.

1. From the vSphere Client **Home** screen, click **EMC**.
The VSI home window appears.
2. Right-click **Unified Storage Management** and click **Uninstall**.
3. Close vSphere Client.

Uninstalling the EMC Unified Storage Access Control Utility

To uninstall the Access Control Utility, follow these steps:

1. Navigate to the **Control Panel**.
2. Double-click **Add or Remove Programs**.
3. Click **Remove** on the entry for **EMC Unified Storage Access Control for VMware**.
4. Click **Yes**.

CHAPTER 5

Troubleshooting

This chapter presents these topics:

- ◆ Known problems and limitations 68
- ◆ Logs 68
- ◆ Technical notes 68
- ◆ EMC Sales and Customer Service contacts 69

Known problems and limitations

Refer to the *VSI for VMware vSphere: Unified Storage Management Release Notes* for the current list of known problems and limitations.

Logs

The log file for EMC VSI for VMware vSphere: Unified Storage Management can be viewed in the VSI log viewer.

1. Navigate to **Home** > **Solutions and Applications** > **EMC**.
2. Click **Logging** in the **Settings** pane.
3. Select **EMC.VSI.VSphere4.Features.USD.UnifiedStorageManagement.txt** from the **Log Files** list box.

Technical notes

- ◆ The EMC VSI for VMware vSphere: Unified Storage Management feature adheres to the limits defined by VMware in the document titled *Configuration Maximums VMware vSphere 5.0 (or Configuration Maximums VMware vSphere 5.1)*, which is available on the VMware website:
<http://www.vmware.com>
- ◆ If you want to extend the size of your gold image virtual machine, you must extend the gold image datastore with the desired size and increase the hard disk size of the virtual machine in the virtual machine settings.
- ◆ For maximum performance, run vSphere Client on the same network as the vCenter Server or storage system.
- ◆ The feature does not support Celerra Nested Mount File Systems (NMFS).
- ◆ Compression is not supported for virtual machines protected by VMware Fault Tolerance.
- ◆ Fast clones are not supported for virtual machines protected by VMware Fault Tolerance.
- ◆ Celerra or VNX file system names and mount paths with special characters are not supported by the feature.
- ◆ Only thin LUNs on VNX block storage systems can be compressed.
- ◆ VSI 5.6.3 only supports compression of block datastores that are hosting virtual machines with operating systems on VMDK. Other disks will not be compressed.
- ◆ Microsoft Windows 2012 does not support Unisphere CLI; therefore, the Unified Storage Management feature does not support the VNXe platform on Windows 2012.
- ◆ Failure of Unified Storage Management operations related to VPLEX, such as credentials management and provisioning, might be caused indirectly by low performance on the VPLEX array. Low performance on VPLEX could lead to behavior changes at the VPLEX API level, which is not yet supported by the Unified Storage Management feature.

EMC Sales and Customer Service contacts

For EMC VSI for VMware vSphere: Unified Storage Management support, review the application-specific Help files or Product Guide. If the documentation does not address your question, and if you have a valid EMC service contract, contact EMC Customer Support at:

- ◆ United States: 1-800-782-4362 (SVC-4EMC)
- ◆ Canada: 1-800-543-4782 (543-4SVC)
- ◆ Worldwide: + 1-508-497-7901

Include the following information in your support request:

- ◆ Type of hardware, including any network hardware, if applicable
- ◆ Operating system
- ◆ Exact wording or screenshots of any messages that appeared on the screen
- ◆ Complete description of the issue, including action taken prior to the incident
- ◆ Log files from the following path (non-64-bit Windows operating systems): \Program Files\VMware\Infrastructure\Virtual Infrastructure Client\Plugins\EMC Unified Storage\log
- ◆ Log files from the following path (64-bit Windows versions): \Program Files (x86)\VMware\Infrastructure\Virtual Infrastructure Client\Plugins\EMC Unified Storage\log
- ◆ Troubleshooting steps

EMC product and licensing information can be obtained as follows:

Product information: For documentation, release notes, and software updates, or for information about EMC products, licensing, and service, go to <https://support.emc.com>.

APPENDIX A

Configuration File Parameters

This appendix presents these topics:

- ◆ [Configuration file name and location.....](#) 72
- ◆ [Parameters within the configuration file](#) 72

Configuration file name and location

EMC VSI for VMware vSphere: Unified Storage Management stores the configuration file at the following location:

```
\AppData\Roaming\EMC\Unified Storage\Config\unified_plugin.ini
```

Parameters within the configuration file

HostNfsMaxVolumes_ESX3

- Default = 32
- The maximum number of NFS datastores that can be mounted on an ESX server with version 3.x. This value is specified by VMware. If an ESX server has a lesser value set and reaches the maximum number of NFS mounts, then the feature automatically increases the NFS.MaxVolumes parameter on the ESX server to this value.

HostNfsMaxVolumes_ESX4

- Default = 64
- The maximum number of NFS datastores that can be mounted on an ESX server with version 4.x. This value is specified by VMware. If an ESX server has a lesser value set and reaches the maximum number of NFS mounts, then the feature automatically increases the NFS.MaxVolumes parameter on the ESX server to this value.

HostNfsMaxVolumes_ESX5

- Default = 256
- The maximum number of NFS datastores that can be mounted on an ESX server with version 5.x. This value is specified by VMware. If an ESX server has a lesser value set and reaches the maximum number of NFS mounts, then the feature automatically increases the NFS.MaxVolumes parameter on the ESX server to this value.

HostVmfsMaxVolumes_ESX

- Default = 256
- The maximum number of VMFS datastores that can be created on an ESX host.

HostTcpIpHeapSize_ESX3

- Default = 30
- When the feature provisions a datastore, it checks the Net.TcpipHeapSize setting. If the setting does not equal this value, then the feature displays a message indicating that you should change the setting and reboot the ESX server.

HostTcpIpHeapSize_ESX4

- Default = 32
- When the feature provisions a datastore, it checks the Net.TcpIpHeapSize setting. If the setting does not equal this value, then the feature displays a message indicating that you should change the setting and reboot the ESX server.

HostTcpIpHeapSize_ESX5

- Default = 32
- When the feature provisions a datastore, it checks the Net.TcpIpHeapSize setting. If the setting does not equal this value, then the feature displays a message indicating that you should change the setting and reboot the ESX server.

HostTcpIpHeapMax_ESX3

- Default = 120
- When the feature provisions a datastore, it checks the Net.TcpIpHeapMax setting. If the setting does not equal this value, then the feature displays a message indicating that you should change the setting and reboot the ESX server.

HostTcpIpHeapMax_ESX4

- Default = 128
- When the feature provisions a datastore, it checks the Net.TcpIpHeapMax setting. If the setting does not equal this value, then the feature displays a message indicating that you should change the setting and reboot the ESX server.

HostTcpIpHeapMax_ESX5

- Default = 128
- When the feature provisions a datastore, it checks the Net.TcpIpHeapMax setting. If the setting does not equal this value, then the feature displays a message indicating that you should change the setting and reboot the ESX server.

DHSMTimeout

- Default = 86400000
- The amount of time in milliseconds that the feature waits for a DHSM request before timing out. DHSM requests are sent to the Data Mover for cloning, compression, and properties features.

HttpTimeout

- Default = 0 (infinite)
- The amount of time in milliseconds that the feature waits for an HTTP request before timing out. HTTP requests are sent to the Celerra or VNX Control Station for all features of the plug-in.

PingTimeout

- Default = 5000
- When you add a Celerra or VNX for file to the feature, the feature first pings the network address you entered to verify that it exists. This value represents the ping timeout in milliseconds.

HostHeartbeatFrequency

- Default = 12
- When you provision a datastore with the feature and leave **Set Timeout Settings** selected on the **Advanced** dialog box, the feature sets this parameter on the ESX server to this value (if it is not already set).

HostHeartBeatMaxFailures

- Default = 10
- When you provision a datastore with the feature and leave **Set Timeout Settings** selected on the **Advanced** dialog box, the feature sets this parameter on the ESX server to this value (if it is not already set).

HostHeartBeatDelta

- Default = 5
- When you provision a datastore with the feature and leave **Set Timeout Settings** selected on the **Advanced** dialog box, the feature sets this parameter on the ESX server to this value (if it is not already set).

HostHeartBeatTimeout

- Default = 5
- When you provision a datastore with the feature and leave **Set Timeout Settings** selected on the **Advanced** dialog box, the feature sets this parameter on the ESX server to this value (if it is not already set).

DatastoreNameMaxLength

- Default = 42
- The maximum length allowed for the datastore name when you provision a datastore with the feature.

MaxNumberOfVmsPerCore

- Default = 12
- The value used to calculate the maximum number of clones that can be assigned to a host or cluster. If you are fast or full cloning a virtual machine to a destination host or cluster, the feature determines how many cores are in the host or cluster and multiplies that by this value to determine the maximum number of clones for that destination.

HttpRetries

- Default = 10
- The number of times the feature auto-retries an HTTP request if a request fails. HTTP requests are sent to the Celerra or VNX Control Station for all features of the plug-in.

MaxFileSystemSizeMB

- Default = 16769024
- The maximum Celerra or VNX file system size in megabytes. This value is specified by EMC.

MinFileSystemSizeMB

- Default = 3
- The minimum size that the feature displays in the list of storage pools when creating new NFS datastores for VNX and Celerra. Storage pools are not displayed if their available capacity is smaller than this size.

MinSharedFolderSizeMB

- Default = 10240
- The minimum size of an NFS shared folder server on VNXe. The value impacts the minimum size of an NFS datastore that can be created on VNXe. The feature checks this value when you enter the NFS datastore size in the feature's provisioning wizard for VNXe. If you enter a smaller value, the wizard cannot proceed with subsequent steps.

HighWaterMarkMin

- Default = 50
- The minimum file system high watermark percentage that is allowed by Celerra or VNX. This value is specified by EMC.

HighWaterMarkMax

- Default = 99
- The maximum file system high watermark percentage that is allowed by Celerra or VNX. This value is specified by EMC.

DefaultHighWaterMark

- Default = 90
- The default high watermark percentage for Celerra or VNX file systems with Virtual Provisioning enabled on the **Advanced** dialog box when you are provisioning a datastore in the feature.

CreateFileSystemTimeout

- Default = 240000
- The timeout for storage provisioning and extension tasks in milliseconds.

LogLevel

- Default = Info
- Specifies the type of messages that get logged. Options are:
 1. Error
 2. Info
 3. Debug

NeoCIMXMLPort

- Default = 5958
- This is a deprecated setting, which is no longer valid for the current version of the feature. However, it is referenced in the storage profile for VNXe.

Do not change the value of this field. Storage profiles breakage is expected with any value other than the default value.

CustomizeScsiLunDisplayName

- Default = True
- Specifies whether the feature displays VMFS datastores and RDM volumes with friendly device names. Options are:
 1. True
 2. False