

Citrix XenServer with Dell SC Series Storage Configuration and Deployment

Dell Storage Engineering
January 2017

Revisions

Date	Description
January 2016	Initial XenServer 6.5 release
January 2017	Updated for XenServer 7.0 and SCOS 7.1

Acknowledgements

Author: Chuck Armstrong

The information in this publication is provided “as is.” Dell Inc. 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 software described in this publication requires an applicable software license.

Copyright © 2016 - 2017 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA [1/23/2017] [Configuration and Deployment] [3132-CD-V]

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

Table of contents

Revisions.....	2
Acknowledgements.....	2
1 Introduction.....	5
1.1 Scope.....	5
1.2 Audience.....	5
1.3 Document conventions.....	5
1.4 Terminology.....	6
2 Citrix XenServer and Dell Storage product overview.....	8
2.1 Citrix XenServer.....	8
2.2 SC Series storage.....	8
3 XenServer storage overview.....	9
3.1 Shared iSCSI storage.....	9
3.1.1 NIC bonding compared to iSCSI MPIO.....	10
3.2 Shared Fibre Channel storage.....	11
3.3 SR-to-VM mapping.....	11
3.4 Multipathing.....	11
3.4.1 Enabling multipathing in XenCenter.....	12
4 XenServer storage using Open-iSCSI initiator (software).....	13
4.1 Open-iSCSI initiator setup with SC Series arrays.....	15
4.2 Multipath with dual subnets.....	15
4.2.1 XenServer Open-iSCSI initiator configuration.....	16
4.2.2 Assign NIC functions using the XenCenter management GUI.....	17
4.2.3 Assign NIC functions using the XE CLI.....	19
4.2.4 Configure Server Objects in Dell Storage Manager.....	21
4.3 Multipath with a single subnet.....	24
4.3.1 Configure bonded storage NICs.....	25
4.3.2 Assign NIC functions to the bond using the XenCenter management GUI.....	29
4.3.3 Assign NIC functions to the bond using the XE CLI.....	30
4.3.4 Configure Server Objects in Dell Storage Manager.....	31
5 XenServer Storage using iSCSI HBA.....	34
5.1 Configure iSCSI HBA.....	34
5.2 Connect to SC Series iSCSI control ports.....	38

- 5.3 Configure Server Objects in Dell Storage Manager39
- 6 XenServer Storage using Fibre Channel HBA42
 - 6.1 Configure Server Objects in Dell Storage Manager42
- 7 Create new Storage Repository (SR).....46
 - 7.1 Software iSCSI SRs46
 - 7.1.1 Identify SC Series storage iSCSI targets46
 - 7.1.2 Create new Software iSCSI SR.....47
 - 7.1.3 Verify multipath status50
 - 7.2 Create SR with hardware HBA (iSCSI and FC)51
- A Technical support and resources54
 - A.1 Related documentation.....54

1 Introduction

This document provides examples, tips, recommended settings, and other storage guidelines a system administrator can follow while configuring a Citrix® XenServer® environment to connect to Dell EMC™ SC Series storage. Frequently asked questions regarding various SC Series storage features are also addressed.

For additional installation and configuration information, Dell EMC recommends reviewing related [XenServer](#) and [XenCenter](#) documentation, which is publicly available on the [Citrix Product Documentation](#) website.

1.1 Scope

This paper covers the steps required to configure a Citrix XenServer environment to use SC Series storage and includes best practices for both Fibre Channel and iSCSI environments. This document is focused on XenServer 7.0 and its related features, however, most of the information provided also applies to XenServer 6.x.

1.2 Audience

This paper is intended for storage administrators, network administrators, SAN system designers, storage consultants, or anyone tasked with configuring a SAN infrastructure for Dell SC Series storage when used to support a Citrix XenServer environment.

It is assumed that readers have received formal training or have advanced working knowledge of:

- Installation and configuration of Citrix XenServer
- Configuration and operation of SC Series storage
- Guest operating systems in use (such as Microsoft® Windows Server® or Linux®)
- Citrix XenServer 7.0 Administrator's Guide

Note: The specific information contained within this document is based on the test environment built for the creation of this document. Actual configuration details may vary in any other environment.

1.3 Document conventions

Table 1 lists the formatting conventions used in this document.

Table 1 Document conventions

Format	Description	Example
Command-line text	User command-line input	<code>iscsiadm -m node --login</code>
<i>Italic command-line text</i>	Placeholder or variable	<code>new_initiator_iqn</code>

1.4 Terminology

The following terms are used throughout this document:

Note: Definitions identified with an asterisk (*) are provided by the *Citrix XenServer 7.0 Administrator's Guide*, which is available on the [Citrix Product Documentation](#) website.

Fault domain (FD): A set of hardware components that share a single point of failure. For controller-level redundancy, fault domains are created for SC Series storage to maintain connectivity in the event of a controller failure. In a dual-switch topology, each switch acts as a fault domain with a separate subnet and VLAN. Failure of any component in an FD will not impact the other FD.

iSCSI offload engine (iSOE): Technology that can free processor cores and memory resources to increase I/O operations per second (IOPS) and reduce processor utilization.

iSCSI Qualified Names (IQNs): Unique iSCSI initiator (host server) or iSCSI target (storage) addresses are referred to as iSCSI Qualified Names (IQNs). IQNs are the identifiers used for iSCSI connectivity between host servers and iSCSI storage platforms.

Link aggregation group (LAG): A group of Ethernet switch ports configured to act as a single high-bandwidth connection to another switch. Unlike a stack, each individual switch must still be administered separately and function independently.

Local area network (LAN): A network carrying traditional IP-based client communications.

Logical unit (LUN): A number identifying a logical device, usually a volume that is presented by an iSCSI or Fibre Channel storage controller.

Multipath I/O (MPIO): A host-based software layer that manages multiple paths for load balancing and redundancy in a storage environment.

Native VLAN and default VLAN: The default VLAN for a packet that is not tagged with a specific VLAN or has a VLAN ID of 0 or 1. When a VLAN is not specifically configured, the switch default VLAN will be utilized as the native VLAN.

Network interface card (NIC): A network interface card or network interface controller is an expansion board inserted into the computer/server so that the computer/server can connect to a network. Most NICs are designed for a particular type of network (typically Ethernet) protocol (typically TCP/IP) and media.

Physical Block Devices (PBDs)*: Physical Block Devices represent the interface between a physical server and an attached Storage Repository (SR). PBDs are connector objects that allow a given SR to be mapped to a XenServer host. PBDs store the device configuration fields that are used to connect to and interact with a given storage target.

Storage area network (SAN): A Fibre Channel, Ethernet, or other specialized network infrastructure specifically designed to carry block-based traffic between one or more servers to one or more storage and storage inter-process communications systems.

Storage Repositories (SRs)*: A Storage Repository is a particular storage target, in which virtual machine (VM) Virtual Disk Images (VDIs) are stored.

Virtual Block Devices (VBDs)*: Virtual Block Devices are connector objects (similar to the PBD described previously) that allows mappings between VDIs and VMs. In addition to providing a mechanism for attaching (also called plugging) a VDI into a VM, VBDs allow for the fine-tuning of parameters regarding Quality of Service (QoS), statistics, and the bootability of a given VDI.

Virtual Disk Images (VDIs)*: A Virtual Disk Image is a storage abstraction which represents a virtual hard disk drive in a VM. VDIs are the fundamental unit of virtualized storage in XenServer. VDIs are persistent, on-disk objects that exist independently of XenServer hosts.

Virtual LAN (VLAN): A method of virtualizing a LAN to make it appear as an isolated physical network. VLANs can reduce the size of and isolate broadcast domains. VLANs still share resources from the same physical switch and do not provide any additional QoS services such as minimum bandwidth, quality of a transmission, or guaranteed delivery.

Virtual machine (VM)*: A virtual machine (VM) is a computer composed entirely of software that can run its own operating system and applications as if it were a physical computer. A VM behaves exactly like a physical computer and contains its own virtual (software-based) CPU, RAM, hard disk, and NIC.

2 Citrix XenServer and Dell Storage product overview

This section provides an overview of Citrix XenServer and the Dell SC Series storage technologies presented in this paper.

2.1 Citrix XenServer

Citrix XenServer is a leading server virtualization and hypervisor management platform that enables reductions in total cost of ownership for desktop, cloud, and server virtualization infrastructures. The ability to consolidate and contain workloads on XenServer provides a means for any organization to address the challenges present in today's IT data center by transforming their compute infrastructure.

2.2 SC Series storage

SC Series storage is the Dell EMC enterprise storage solution featuring multi-protocol support and self-optimizing, tiering capabilities. SC Series storage can be configured with all flash, as a hybrid system, or with only traditional spinning disks and features automatic migration of data to the most cost-effective storage tier. Efficient thin provisioning and storage virtualization enable disk capacity usage only when data is actually written, enabling a pay-as-you-grow architecture. This self-optimizing system can reduce overhead cost and free up the administrator for other important tasks.

3 XenServer storage overview

XenServer environments utilize shared storage as part of the virtualization platform. Shared storage can be connected through iSCSI, Fibre Channel (FC), or NFS. The information contained in this document is focused on iSCSI and Fibre Channel block storage connectivity because the SC Series storage platform is a block storage platform. The addition of FluidFS storage (a file-based NAS platform) to the SC Series platform can provide NFS storage, but is not covered as part of this document.

Note: Additional information regarding FluidFS storage platforms can be found on the page, [FluidFS technical content](#).

Block storage devices in the form of iSCSI or FC LUNs are presented to the XenServer hosts, from which SRs are created. SRs, which are connected to all XenServer hosts participating in the XenServer pool, are the entity on which VDIs reside. VDIs are recognized by VMs as physical disks, and are persistent on the SRs.

3.1 Shared iSCSI storage

XenServer using SC Series storage provides support for shared SRs on iSCSI-attached LUNs. iSCSI LUNs can use the Open-iSCSI software initiator or a supported iSCSI host bus adapter (HBA). Figure 1 and Figure 2 illustrate how both the Open-iSCSI software initiator and iSCSI HBA environments look conceptually.




Figure 1 Shared iSCSI storage using the Open-iSCSI software initiator




Figure 2 Shared iSCSI storage using iSCSI HBAs

3.1.1 NIC bonding compared to iSCSI MPIO

NIC bonding improves XenServer host resiliency by using two physical NICs as if they were one. If one NIC within the bond fails, the host network traffic will automatically be rerouted through the second NIC. NIC bonding supports active/active mode, but only supports load balancing of VM traffic across the physical NICs. Each virtual network interface will only use one of the links in the bond at a time. Load balancing for non-VM traffic does not occur with NIC bonding.

Multi-Path Input/Output (MPIO) also provides host resiliency by using two physical NICs. MPIO uses round robin to balance the storage traffic between all available storage targets on a Dell SC Series storage array. By spreading the load between multiple SC Series storage targets, bottlenecks can be avoided while providing NIC, subnet, and switch redundancy.

If all Front End iSCSI ports on the SC Series storage are on the same subnet, NIC bonding is the recommended option because XenServer iSCSI MPIO requires at least two separate subnets. In this configuration, all iSCSI connections will use the same physical NIC because bonding does not support active/active connections for anything but VM traffic. For this reason, it is recommended that Front End iSCSI ports be configured for two subnets. This allows load balancing across all NICs and failover with MPIO.

3.2 Shared Fibre Channel storage

XenServer using SC Series storage provides support for shared SRs on FC-attached LUNs using FC HBAs. Figure 3 illustrates how an environment using FC HBAs looks conceptually.




Figure 3 Shared Fibre Channel storage

3.3 SR-to-VM mapping

XenServer is capable of deploying a many-to-one, VM-to-SR (volume) deployment. The best number of VMs per SR largely depends on the workload and IOPS requirement of the VMs being deployed. When multiple VDIs share an SR, they also share the disk queue for that SR on the host. For this reason, care should be taken to prevent bottleneck conditions on the SR. Additionally, replication and DR become a factor when hosting multiple VMs on an SR. This is due to replication and recovery taking place on a per-SR (volume) basis.

3.4 Multipathing

Multipathing allows for failures in storage adapters (Open-iSCSI and HBA), switch ports, SAN switches, and storage I/O ports. Use of multipathing is recommended to increase availability and redundancy for production deployments of XenServer when hosting critical VM workloads.

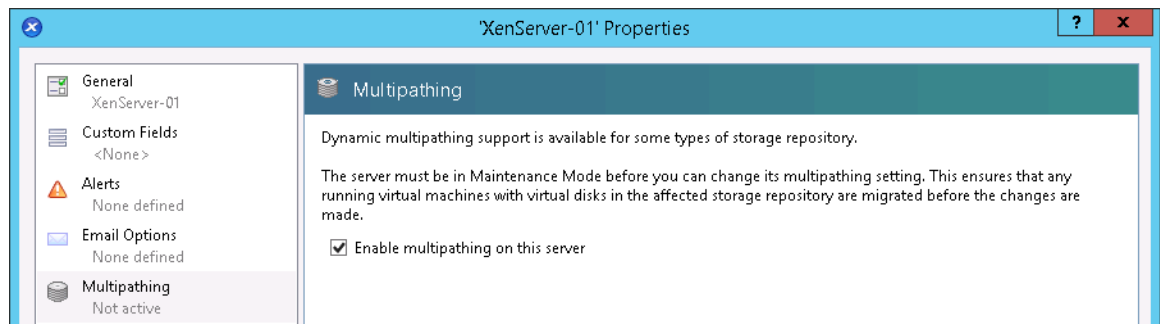
XenServer supports Active/Active multipathing for iSCSI and FC protocols for I/O data paths. Dynamic multipathing uses a round-robin mode load balancing algorithm resulting in active traffic on all paths during normal operation. Multipathing can be enabled using XenCenter or the command line interface. Enabling

multipathing requires a restart of the XenServer host and should be enabled before storage is added to the server. Only use multipathing when there are multiple paths to the storage.

Note: Additional information on multipathing with XenServer hosts can be found in the *Citrix XenServer 7.0 Administrator's Guide*, which is available on the [Citrix Product Documentation](#) website.

3.4.1 Enabling multipathing in XenCenter

1. Right-click the server in XenCenter and select **Enter Maintenance Mode**.
2. Right-click the server and select **Properties**.
3. In the **Properties** window, select **Multipathing**.
4. Check the **Enable multipathing on this server** box and click **OK**.
5. The server will need to be restarted for multipathing to take effect.



4 XenServer storage using Open-iSCSI initiator (software)

XenServer iSCSI storage repositories are supported with SC Series storage through the use of the Open-iSCSI initiator.

Shared iSCSI SRs using the software-based host initiator are capable of supporting VM agility using XenMotion — VMs can be started on any XenServer host in a resource pool and migrated between them with no noticeable interruption.

iSCSI SRs utilize the entire LUN specified at creation time and may not span more than one LUN. Support for the Challenge-Handshake Authentication Protocol (CHAP) is provided for client authentication, during both the data-path-initialization and the LUN-discovery phases.

Note: Use dedicated network adapters for iSCSI traffic. The default connection can be used, but the best practice is to separate iSCSI and network traffic.

XenServer hosts support a single iSCSI initiator, which is automatically created and configured with a random iSCSI Qualified Name (IQN) during host installation. iSCSI targets commonly provide access control through iSCSI initiator IQN lists, so all iSCSI targets/LUNs to be accessed by a XenServer host must be configured to allow access by the host initiator IQN. Similarly, targets/LUNs to be used as shared iSCSI SRs must be configured to allow access by all host IQNs in the resource pool.

Changing the default XenServer IQN to one that is consistent with a naming schema in the iSCSI environment is recommended. The XenServer host IQN value can be modified using the XenCenter GUI, or the XE CLI.

To set the host IQN using XenCenter:

Right-click the host, select **Properties**, enter the desired **iSCSI IQN**, and click **OK**.




Figure 4 XenCenter: Set host IQN

To set the host IQN using the XE CLI:

From the XE CLI, execute the following command to modify the host IQN.

```
xe host-param-set uuid=host_uuid other-config:iscsi_iqn=new_initiator_iqn
```




Figure 5 XE CLI: Set host IQN

Caution: Do not change the XenServer host IQN with iSCSI SRs attached. Doing so can result in failures connecting to new targets or existing SRs.

Caution: When changing the host (Initiator) IQN, it is imperative that every iSCSI target and initiator have a unique IQN. If a non-unique IQN identifier is used, data corruption and/or denial of LUN access can occur.

4.1 Open-iSCSI initiator setup with SC Series arrays

When planning an iSCSI network, isolate the iSCSI traffic from management traffic through the use of separate switches and subnets. Failure to follow this best practice may result in compromised reliability or performance.

When implementing multipathing with iSCSI storage, be certain none of the redundant iSCSI paths are configured within the same subnet or on the same physical network as the management interface. If this occurs, the iSCSI initiator will not be able to successfully establish a session over any iSCSI path on the management network or subnet.

There are two commonly used ways to implement multipathing with the XenServer Open-iSCSI initiator to connect to SC Series storage:

Multipath with dual subnets: In this configuration, the Front End iSCSI control ports on the SC Series storage are on two separate subnets. This option uses MPIO for multipathing. This is the recommended option when high availability (HA) is required.

Multipath with single subnet: In this configuration, the Front End iSCSI control ports on the SC Series storage are on the same subnet. This option uses NIC bonding for path failover. This method is an optional solution when the servers have a single iSCSI storage NIC and HA is not required.

4.2 Multipath with dual subnets

Using XenServer Open-iSCSI multipathing with dual subnets to properly connect to SC Series storage requires the following:

- XenServer 6.5 or later
- iSCSI using two unique, dedicated storage NICs and subnets. The two subnets should be different from the XenServer management network to comply with Citrix best practices
- Multipathing enabled on all XenServer hosts in the pool
- iSCSI target IP addresses for the SC Series storage Front End control ports. For the example included in this document, the iSCSI FE control ports on SC Series storage are assigned IP addresses 10.10.10.100/24 and 10.20.10.100/24

In this configuration, the SC Series storage is configured with the iSCSI Front End ports on two separate subnets, different from the management interface. The SC Series storage is configured with two control ports, one on each subnet. Multipathing is controlled through MPIO.

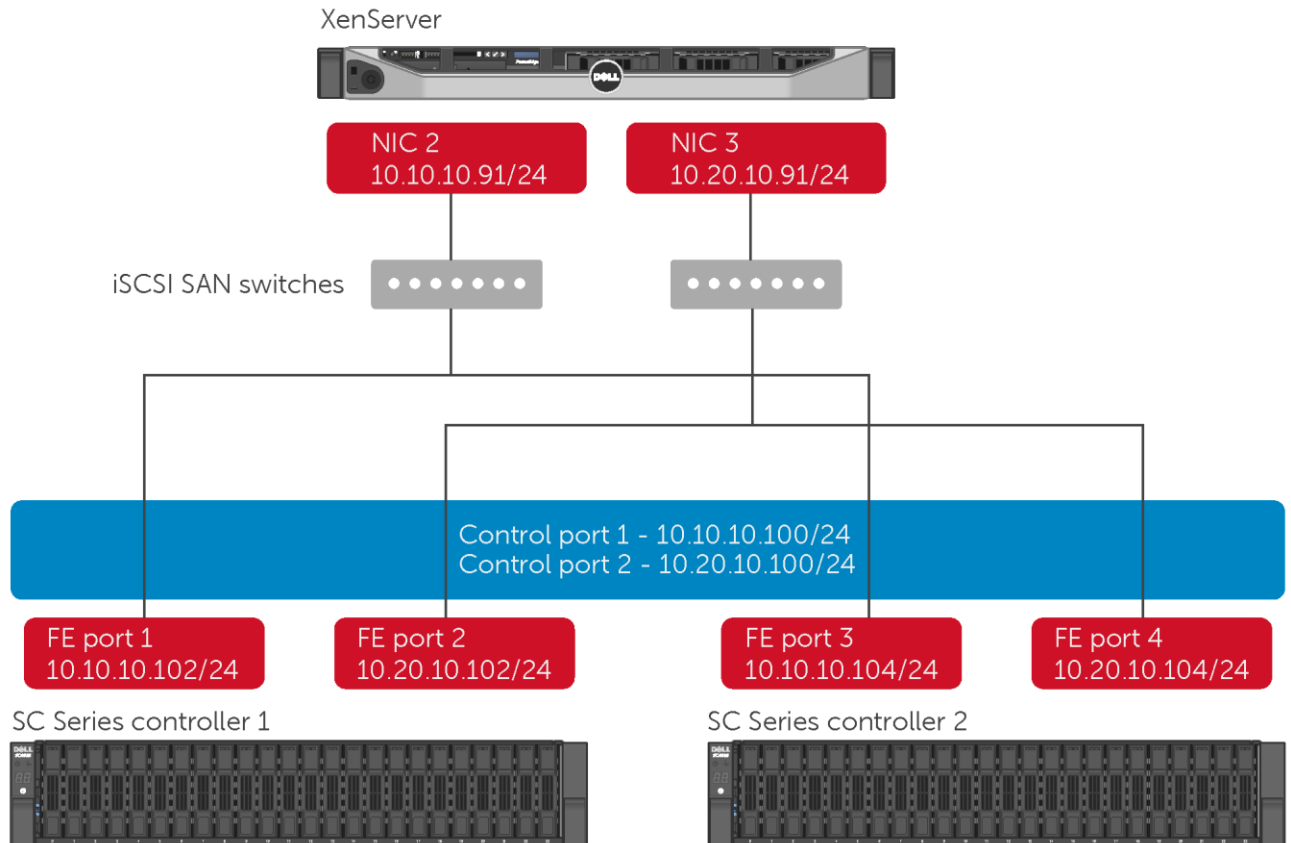


Figure 6 Open-iSCSI: dual subnets and MPIO

4.2.1 XenServer Open-iSCSI initiator configuration

The XenCenter management GUI or the XE CLI can be used to configure dedicated NICs for iSCSI storage traffic use. Assigning a NIC for iSCSI use will prevent the use of the NIC for other functions such as host management. However, appropriate network configuration is also required to ensure the NIC is used for the desired traffic. For example, to dedicate a NIC to iSCSI storage traffic, the NIC, storage target, switch, and VLAN (if a VLAN is used) must be configured so the iSCSI storage target is only accessible over the assigned NIC.

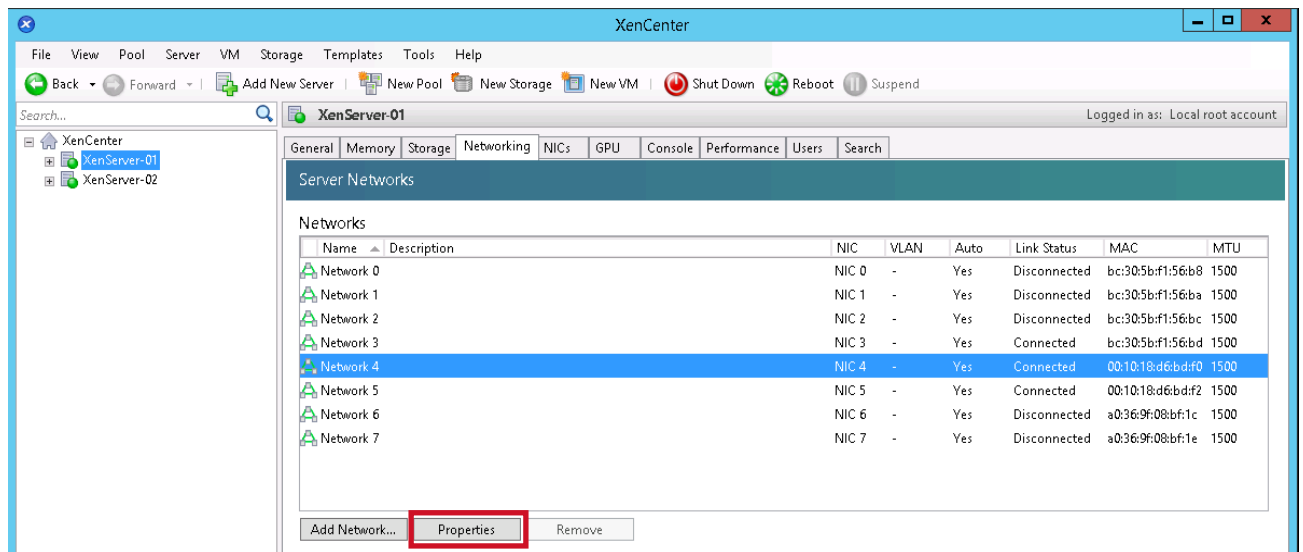
Ensure that the dedicated NICs used for iSCSI storage use a separate IP subnet that is not routable from the XenServer management interface. Enforcing this ensures storage traffic will not be directed over the management interface after a host reboot, which would otherwise be possible due to the initialization of the network interfaces.

4.2.2 Assign NIC functions using the XenCenter management GUI

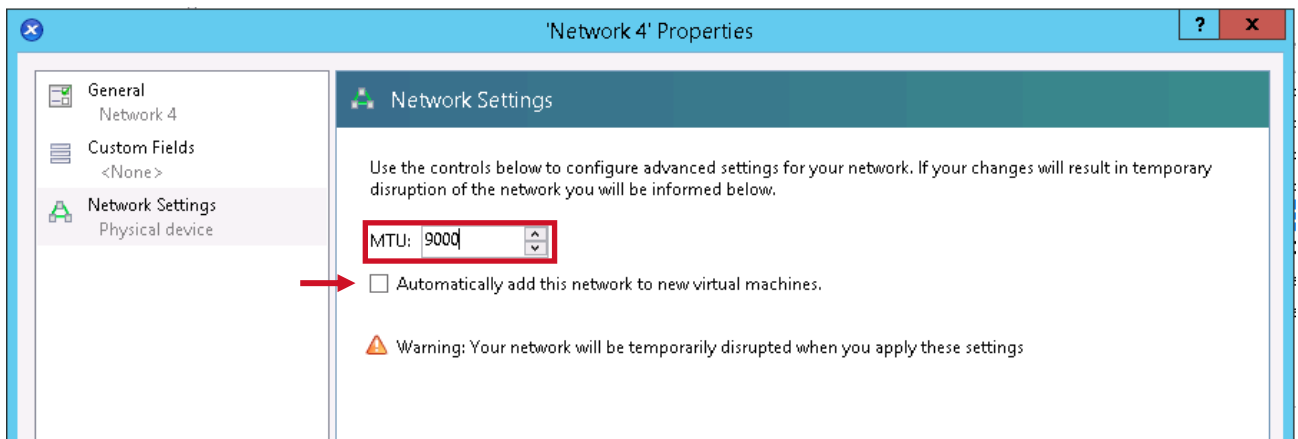
To perform these steps using the CLI rather than the XenCenter GUI, see section [4.2.3](#). Be sure Multipathing has been enabled as shown in section [3.4.1](#).

4.2.2.1 Optional steps: implementing Jumbo Frames

1. In the XenCenter management GUI, navigate to the **Infrastructure** view, navigate through the objects, select the desired XenServer host, select the **Networking** tab, select the desired **Network** from the list, and click **Properties**.



2. Select **Network Settings**, input the **MTU** of 9000 (default value is 1500), and deselect the checkbox so this network will not be added to new VMs.



3. Repeat steps 1–2 for each additional network dedicated for iSCSI storage.

4.2.2.2 Required steps

Note: If Jumbo Frames are to be used, the steps in section 4.2.2.1 must be completed prior to executing the following steps.

1. In the XenCenter management GUI, navigate to the **Infrastructure** view, navigate through the objects, select the desired XenServer host, select the **Networking** tab, and click **Configure**.

The screenshot displays the XenCenter management console for XenServer-01. The 'Networking' tab is active, showing a list of server networks and their IP configurations. The 'Server Networks' table is as follows:

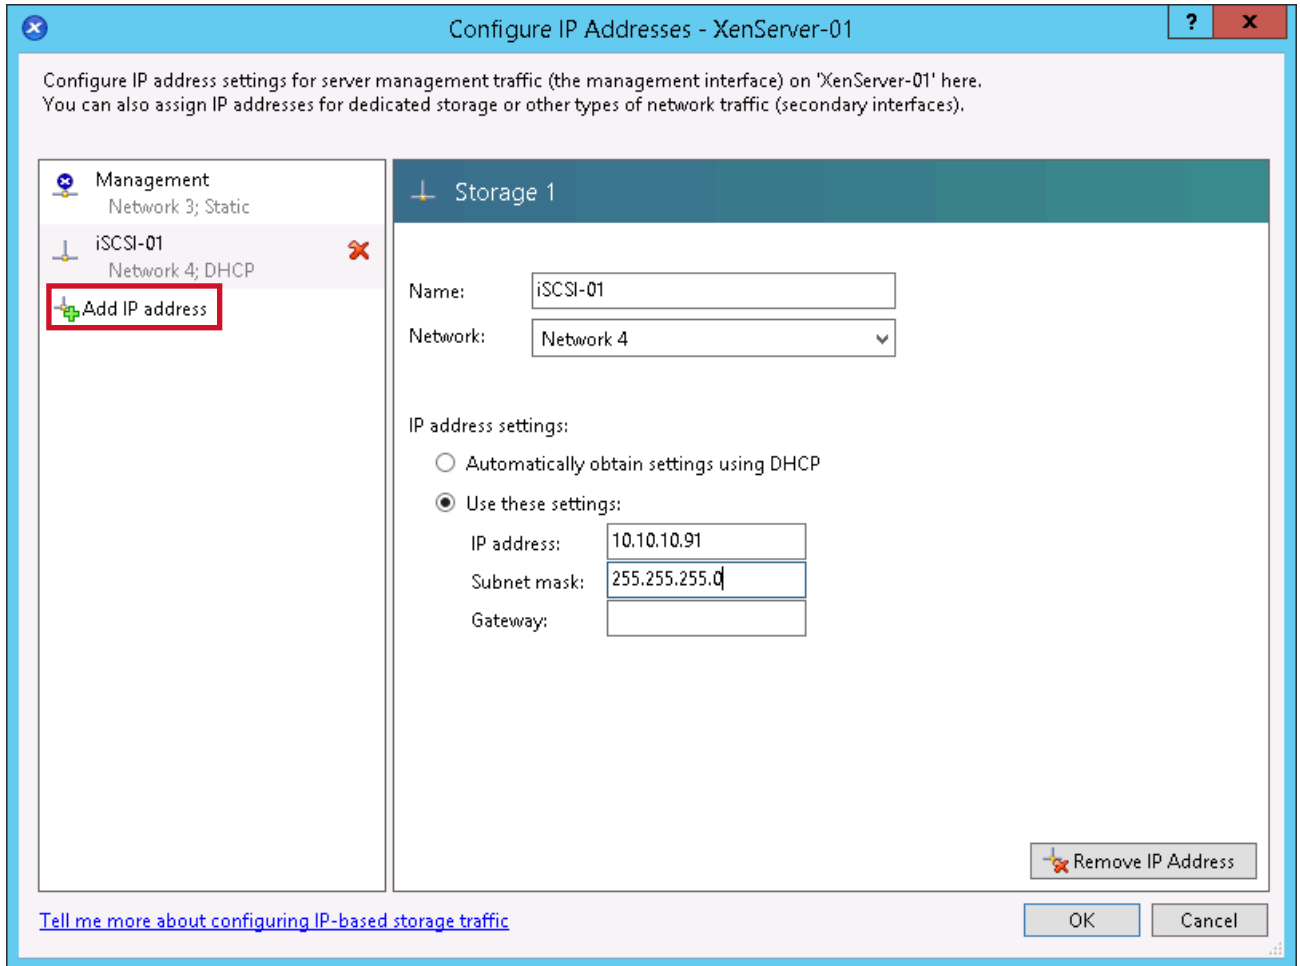
Name	Description	NIC	VLAN	Auto	Link Status	MAC	MTU
Network 0		NIC 0	-	Yes	Disconnected	bc:30:5b:f1:56:b8	1500
Network 1		NIC 1	-	Yes	Disconnected	bc:30:5b:f1:56:ba	1500
Network 2		NIC 2	-	Yes	Disconnected	bc:30:5b:f1:56:bc	1500
Network 3		NIC 3	-	Yes	Connected	bc:30:5b:f1:56:bd	1500
Network 4		NIC 4	-	No	Connected	00:10:18:d6:bd:f0	9000
Network 5		NIC 5	-	Yes	Connected	00:10:18:d6:bd:f2	1500
Network 6		NIC 6	-	Yes	Disconnected	a0:36:9f:08:bf:1c	1500
Network 7		NIC 7	-	Yes	Disconnected	a0:36:9f:08:bf:1e	1500

Below the networks table, the 'IP Address Configuration' table is shown:

Server	Interface	Network	NIC	IP Setup	IP Address	Subnet mask	Gateway	DNS
XenServer-01	Management	Network 3	NIC 3	Static	10.211.17.91	255.255.255.192	10.211.17.126	10.211.18.4,10.211.18.5

A red box highlights the 'Configure...' button at the bottom left of the IP Address Configuration table.

2. Click **Add IP address**, enter the desired name, select the **Network** from the drop-down box, and enter the **IP address settings** for the dedicated iSCSI storage NIC. Click **OK**.



3. Repeat steps 1-2 for each additional NIC dedicated for iSCSI storage.

4.2.3 Assign NIC functions using the XE CLI

If NIC functions were assigned using the XenCenter GUI in section 4.2.2, please skip to section 4.2.4 to configure the server objects in Dell Storage Manager. Be sure Multipathing has been enabled as shown in section 3.4.1.

4.2.3.1 Optional steps: implementing Jumbo Frames

1. Get the PIF UUID for the interface:
 - For a standalone XenServer host: Execute `xe pif-list` to list the PIFs on the server.
 - If the XenServer host is part of a pool:
 - a. Execute `xe host-list` to retrieve a list of the hosts and UUIDs.
 - b. Execute `xe pif-list host-uuid=host-uuid` to list the PIFs on the selected host.
2. Set the MTU parameter to 9000 (default value is 1500):

```
xe pif-param-set other-config:mtu=9000 uuid=Pif-UUID
```

3. Repeat this process for each eth interface dedicated for iSCSI storage traffic on each XenServer host connecting to the SC Series storage.

4.2.3.2 Required steps

Note: If Jumbo Frames are to be used, the steps in section 4.2.3.1 must be completed prior to executing the following steps.

1. Ensure that the physical interface (PIF) is on a separate subnet or that routing is configured to suit your network topology, forcing the desired traffic over the selected PIF.
2. Get the PIF UUID for the interface:
 - For a standalone XenServer host: Execute `xe pif-list` to list the PIFs on the server.
 - If the XenServer host is part of a pool:
 - i. Execute `xe host-list` to retrieve a list of the hosts and UUIDs.
 - ii. Execute `xe pif-list host-uuid=host-uuid` to list the PIFs on the selected host.
3. Set up an IP configuration for the PIF, adding appropriate values for the mode parameter, and if using static IP addressing, add values for the IP, netmask, gateway (if required), and DNS parameters:

```
xe pif-reconfigure-ip mode=DHCP|static uuid=pif-uuid
```

Example: `xe pif-reconfigure-ip mode=static ip=10.10.10.91 netmask=255.255.255.0 gateway=10.10.10.1 uuid=Pif-UUID`

Note: When setting IP information for iSCSI connection, the gateway parameter is only required if iSCSI traffic must route to another IP subnet and has access to the appropriate router.

4. Set the PIFs disallow-unplug parameter to true:

```
xe pif-param-set disallow-unplug=true uuid=Pif-UUID
```

5. Set the management purpose of the interface:

```
xe pif-param-set other-config:management_purpose="iSCSI-01" uuid=Pif-UUID
```

6. Disable automatic assignment of network to new VMs:

```
xe pif-param-set other-config:automatic="false" uuid=Pif-UUID
```

7. Repeat steps 3–6 for each eth interface dedicated for iSCSI storage traffic on each XenServer host connecting to the SC Series storage. For iSCSI MPIO configurations, a minimum of two eth interfaces on each XenServer host, on separate subnets is recommended.


Note: XE CLI uses `xe pif` to identify ethx devices where *x* identifies which device. The XenCenter management GUI identifies the same devices with the NIC *x* designation, where *x* identifies which device. While the designation differs (eth vs NIC), the *x* identifier is consistent between utilities.

Note: For more information on this topic see the *Citrix XenServer 7.0 Administrator's Guide*, which is available on the [Citrix Product Documentation](#) website.


4.2.4 Configure Server Objects in Dell Storage Manager

Use the following steps to configure the server object for access to the SC Series storage:


1. In Dell Storage Manager Client, go to the **Storage** tab, drill down through the array to highlight the **Servers** object, click **Create Server**.




2. Enter the name to identify the XenServer host, select the operating system from the drop-down list, and click **Manually Add HBA**.




3. Select **iSCSI** from the **HBA Port Type** drop-down list, enter the previously defined iSCSI IQN in the **WWN or iSCSI Name** field, and click **OK**.



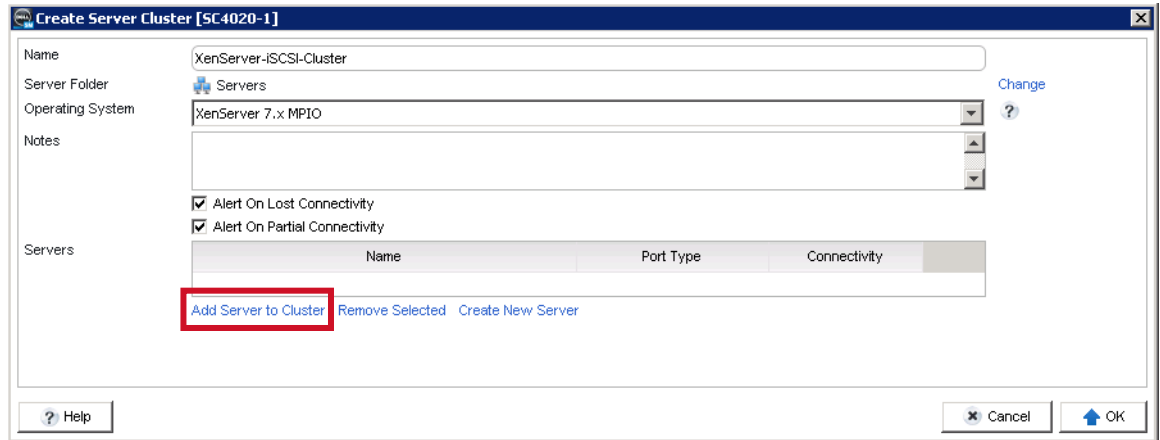
4. In the completed Create Server window, click **OK**.



5. Repeat steps 1–3 for each XenServer to be added to the pool.
6. With **Servers** still highlighted, click **Create Server Cluster**.



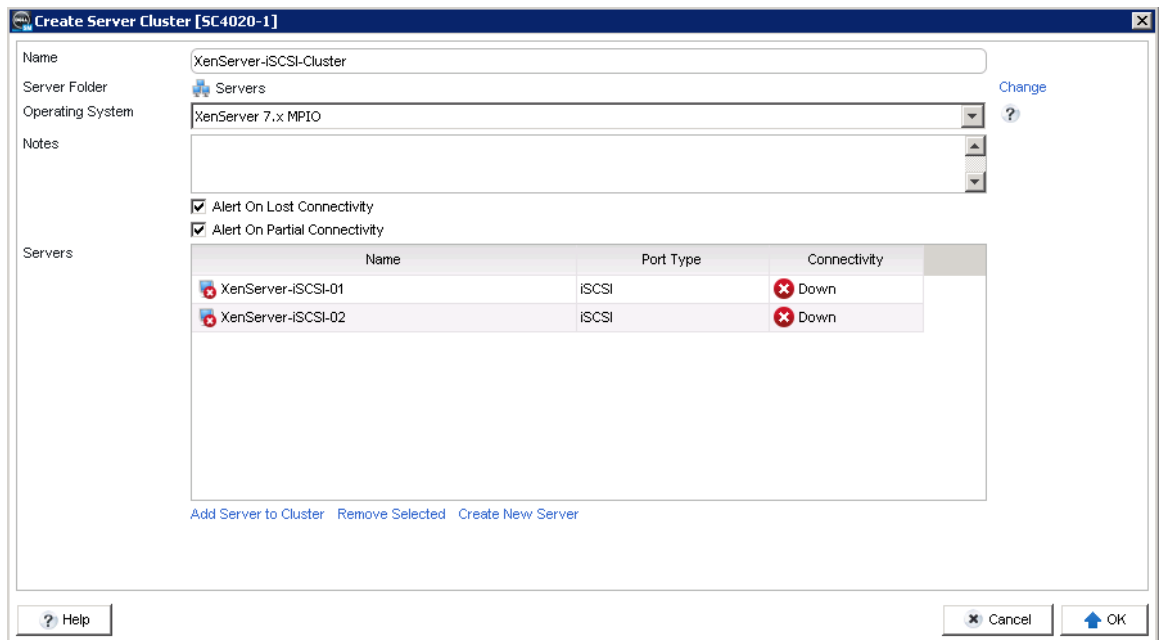
7. Enter the **Name** to identify the XenServer Cluster, select the **Operating System** from the drop-down list, and click **Add Server to Cluster**.



8. Select the server from list and click **OK**.



9. Repeat steps 6-7 for each XenServer host that will be part of the XenServer pool.
10. Click **OK** to create the server cluster.



After creating the server and server cluster objects, volumes can be created and mapped to the servers. For XenServers in a pool, map LUNs to the server cluster object to ensure all servers use the same LUN number.

Continue to [section 7](#) to create Storage Resources from these volumes on the XenServer or XenServer pool.

Note: See the [Create a Cluster Object in Enterprise Manager](#) and [Creating and mapping a volume in Enterprise Manager](#) videos for additional information.

4.3 Multipath with a single subnet

Assigning a network to storage will prevent the use of the network for other functions such as host management, but requires that the appropriate configuration be in place in order to ensure the network is used for the desired traffic. For example, to dedicate a network to storage traffic, the NIC, storage target, switch, and VLAN (if a VLAN is used) must be configured such that the target is only accessible over the assigned NIC. This allows use of standard IP routing to control how traffic is routed between multiple NICs within a XenServer.

Before dedicating a network interface as a storage interface for use with iSCSI SRs, ensure that the dedicated interface uses a separate IP subnet which is not routable from the main management interface. If this is not enforced, then storage traffic may be directed over the main management interface after a host reboot, due to the order in which network interfaces are initialized.

The process for configuring multipathing in a single subnet environment is similar to that of a dual subnet environment. The key difference is that redundancy is handled by the bonded network adapters. The requirements for Open-iSCSI multipathing with SC Series storage in a single subnet are as follows:

- XenServer 6.5 or later
- iSCSI using two unique, dedicated storage NICs bonded together to act as a single interface
- One network storage interface using the bonded interface on a dedicated subnet that is different from the XenServer management network to comply with Citrix best practices
- Multipathing enabled on all XenServer hosts in the pool
- iSCSI target IP addresses for the SC Series storage Front End control ports (in this example, the IP addresses for the control port are 10.10.10.100 and 10.10.10.110)




Figure 7 Open-iSCSI: single subnet and MPIIO

4.3.1 Configure bonded storage NICs


In this configuration, redundancy to the network is provided by two bonded NICs. Bonding the two NICs will create a new bonded interface that network interfaces will be associated with. This will create multiple paths with one storage IP address on the server.

The creation of the bonded interface differs based on membership in a pool. If the XenServer host is part of a XenCenter pool, configuration is performed through the XenCenter management GUI. However, if the XenServer host is a standalone (not part of a pool), configuration is performed with the XE CLI as shown in [section 4.3.1.3](#). Be sure Multipathing has been enabled as shown in [section 3.4.1](#).

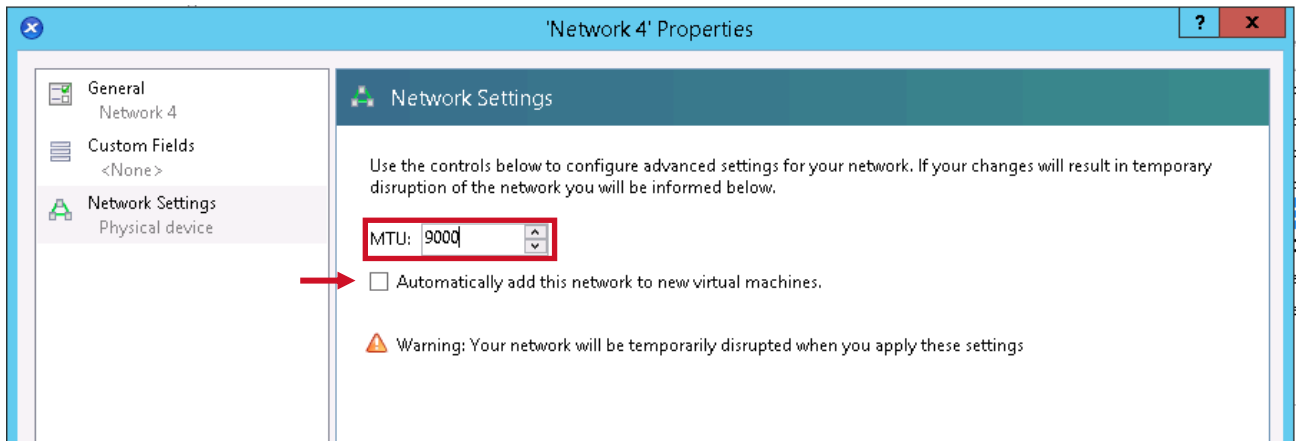
Note: Create NIC bonds as part of the initial resource pool creation, prior to joining additional hosts to the pool. This will allow the bond configuration to be replicated to new hosts as they join the pool.

4.3.1.1 Optional steps: implementing Jumbo Frames on a host participating in a pool

1. In the XenCenter management GUI, navigate to the **Infrastructure** view, drill down through the objects, select the desired XenServer host, select the **Networking** tab, select the desired network from the list, and click **Properties**.



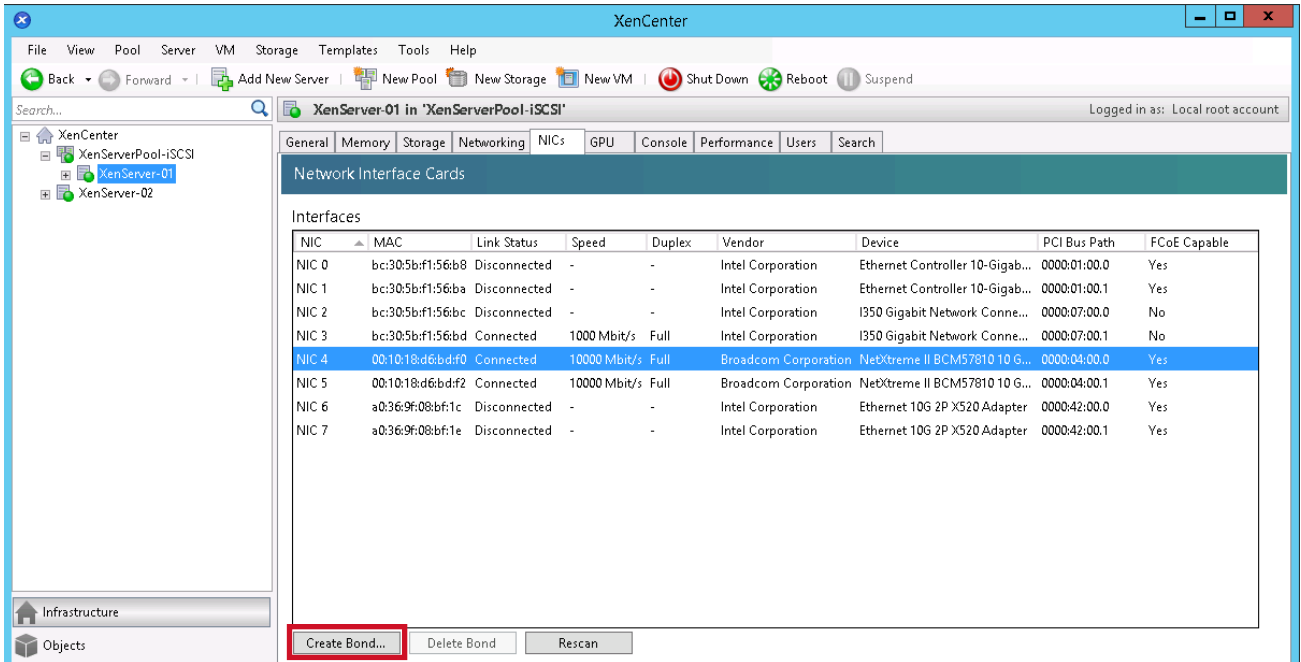
2. Select **Network Settings**, input the **MTU** of 9000 (default value is 1500), and deselect the checkbox so this network will not be added to new VMs.



3. Repeat steps 1–2 for each additional network dedicated for iSCSI storage.

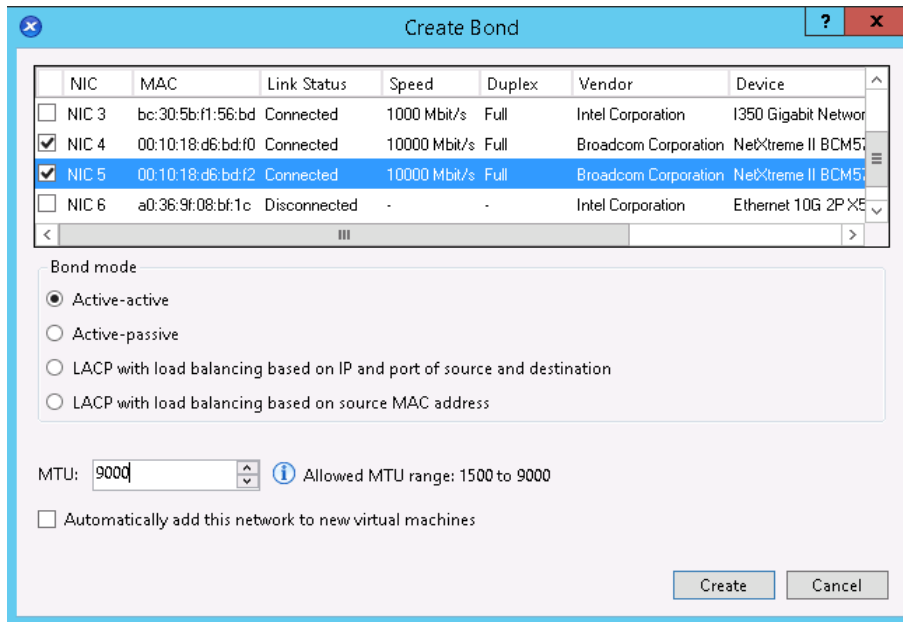
4.3.1.2 Creating a NIC bond in a XenServer host participating in a pool

1. From the XenCenter GUI, select the server, click the **NIC** tab, and click **Create Bond**.



2. Select the NICs you would like to participate in the bond, select the **Bond** mode (default is Active-active), modify the MTU (default is 1500), and click **Create**.

Note: If Jumbo Frames are to be used, the MTU must be set (as shown in section 4.3.1.1) prior to creating the bonded interface.



- Once complete, there will be a new bonded NIC displayed in the list of NICs.

Interfaces				
NIC	MAC	Link Status	Speed	Duplex
Bond 4+5	00:10:18:d6:bd:f0	Connected	10000 Mbit/s	Full

4.3.1.3 Optional steps: implementing Jumbo Frames on a standalone host

- Get the PIF UUID for the interface:
 - For a standalone XenServer host: Execute `xe pif-list` to list the PIFs on the server.
 - If the XenServer host is part of a pool:
 - Execute `xe host-list` to retrieve a list of the hosts and UUIDs.
 - Execute `xe pif-list host-uuid=host-uuid` to list the PIFs on the selected host.
- Set the MTU parameter to 9000 (default value is 1500):


```
xe pif-param-set other-config:mtu=9000 uuid=Pif-UUID
```
- Repeat this process for each eth interface dedicated for iSCSI storage traffic on each XenServer host connecting to the SC Series storage.

4.3.1.4 Creating a NIC bond in a standalone XenServer host using the XE CLI

Note: If Jumbo Frames are to be used, the MTU must set as shown in section 4.3.1.3 prior to creating the bonded interface.

- Create the network to which the bonded interface will connect:


```
xe network-create name-label=desired_name
```
- Set the MTU parameter to 9000 (default value is 1500):


```
xe network-param-set MTU=9000 uuid=Network_UUID
```
- Get the PIF UUIDs for the NICs to be included in the bond:



```
xe pif-list
```
- Create the bonded interface:


```
xe bond-create network-uuid=Network_UUID pif-uuids=Pif-UUID_1,Pif-UUID_2,<mode=<balance-slb | active-backup | lacp>> (default mode is Active-active)
```


4.3.2 Assign NIC functions to the bond using the XenCenter management GUI

To perform these steps using the CLI rather than the XenCenter GUI, see section [4.3.3](#). Be sure Multipathing has been enable a shown in section [3.4.1](#).

1. In the XenCenter management GUI, navigate to the Infrastructure view, drill down through the objects, select the desired XenServer host, select the **Networking** tab, and click **Configure**.



2. Click **Add IP address**, enter the desired name, select the Network from the drop-down box, enter the IP address settings for the dedicated iSCSI storage NIC, and click **OK**.



4.3.3 Assign NIC functions to the bond using the XE CLI

If NIC functions were assigned using the XenCenter GUI in section [4.3.2](#), please skip to section [4.3.4](#) to configure the server objects in Dell Storage Manager. Be sure Multipathing has been enable a shown in section [3.4.1](#).

1. Get the PIF UUID for the bond interface:
 - If on a standalone server, use `xe pif-list` to list the PIFs on the server.
 - If on a host in a resource pool:
 - a. Type `xe host-list` to retrieve a list of the hosts and UUIDs.
 - b. Use the command `xe pif-list host-uuid=host-uuid` to list the host PIFs.
2. Set up an IP configuration for the PIF, adding appropriate values for the mode parameter, and if using static IP addressing the IP, add values for the netmask, gateway (if required) and DNS parameters:

```
xe pif-reconfigure-ip mode=DHCP|static uuid=pif-uuid
```

Example: `xe pif-reconfigure-ip mode=static ip=10.10.10.91
netmask=255.255.255.0 gateway=10.10.10.1 uuid=Pif-UUID`

Note: When setting IP information for iSCSI connection, the gateway parameter is only required if iSCSI traffic must route to another IP subnet and has access to the appropriate router.

3. Set the PIF disallow-unplug parameter to true:

```
xe pif-param-set disallow-unplug=true uuid=Pif-UUID
```


Note: Using `xe pif` uses `ethx` devices where `x` is the same as NIC `x` in the XenCenter management GUI.

Note: For more information on this topic, see the *Citrix XenServer 7.0 Administrator's Guide*, which is available on the [Citrix Product Documentation](#) website.

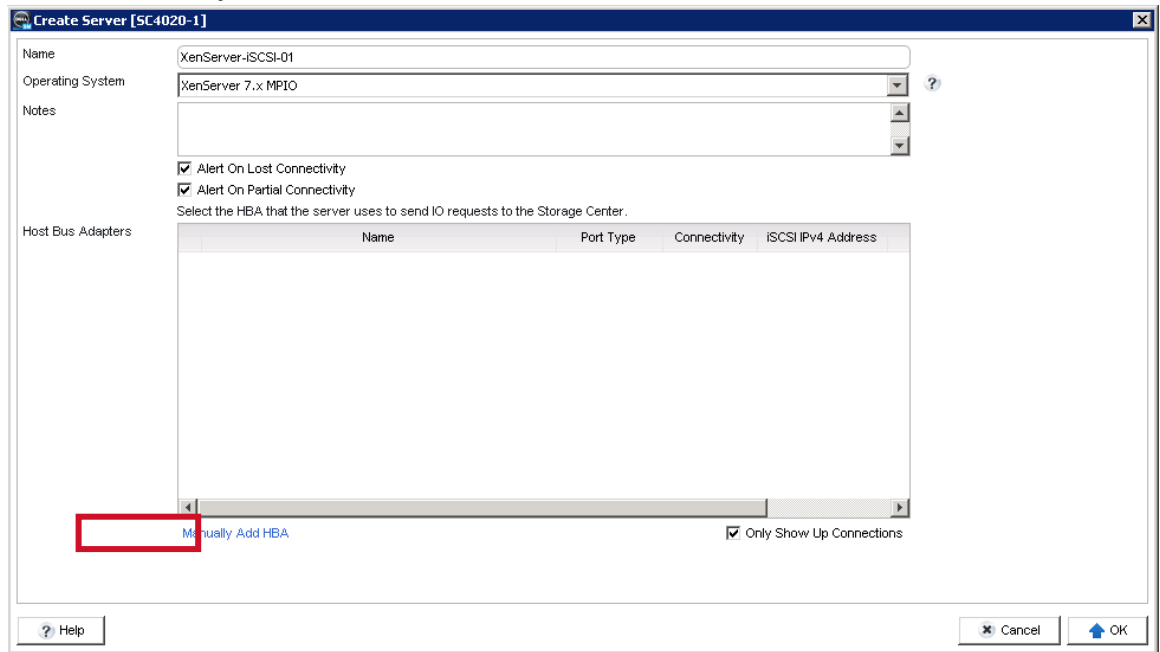
4.3.4 Configure Server Objects in Dell Storage Manager

Use the following steps to configure the server object for access to the SC Series storage:

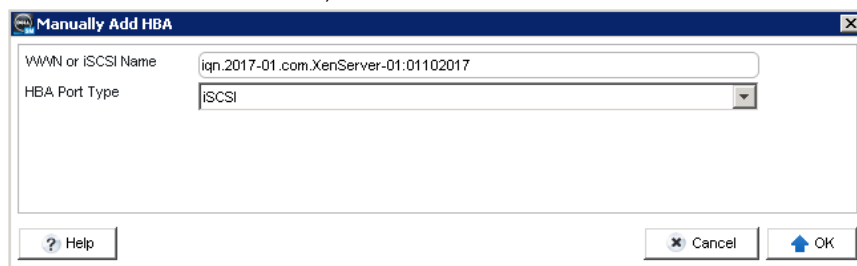
1. In Dell Storage Manager Client, go to the **Storage** tab, navigate through the array to highlight the **Servers** object, click **Create Server**.



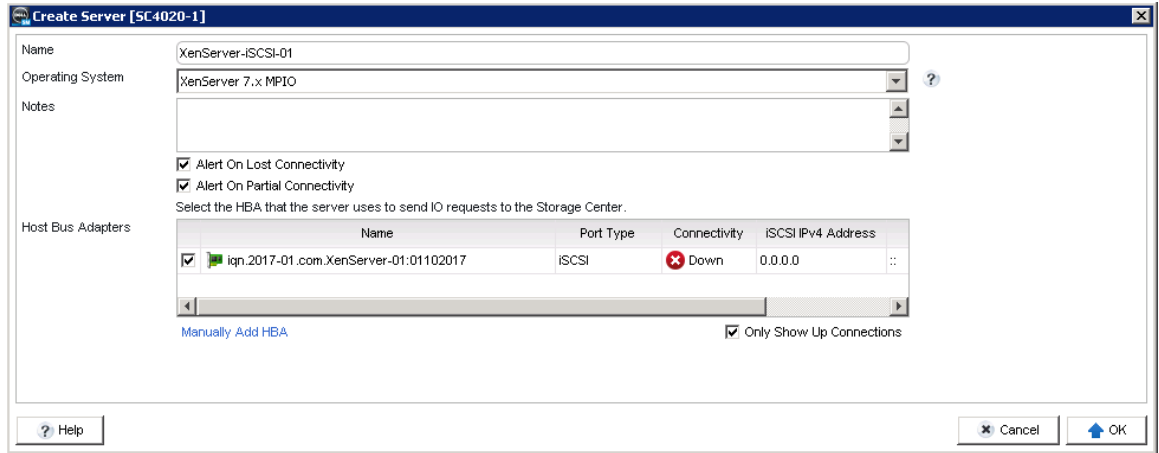
2. Enter the name to identify the XenServer host, select the operating system from the drop-down list, and click **Manually Add HBA**.




3. Select **iSCSI** from the **HBA Port Type** drop-down list, enter the previously defined iSCSI IQN in the **WWN or iSCSI Name** field, and click **OK**.



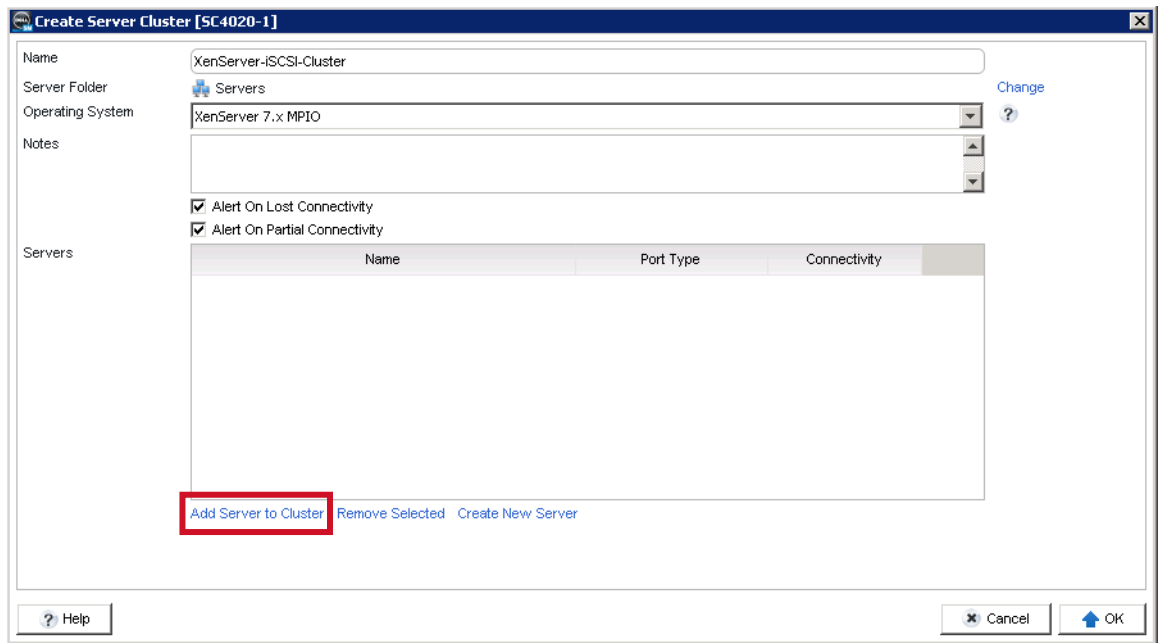
- In the completed Create Server window, click **OK**.



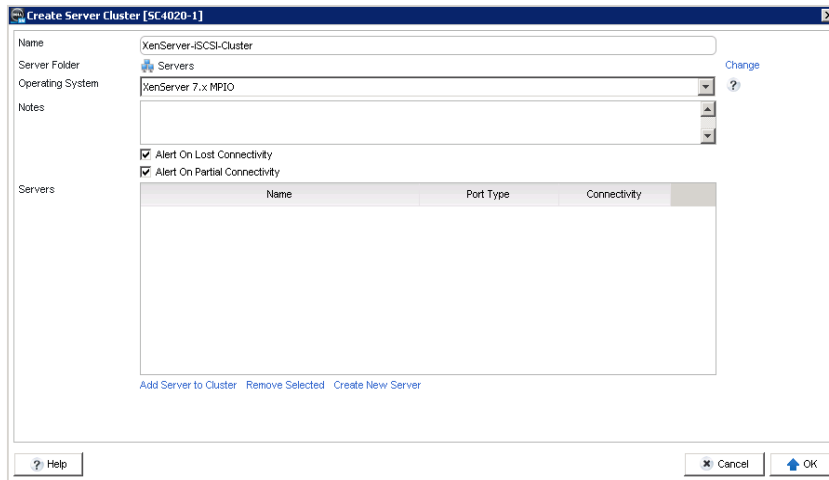
- Repeat steps 2–3 for each XenServer to be added to the pool.
- With **Servers** still highlighted, click **Create Server Cluster**.




- Enter the **Name** to identify the XenServer Cluster, select the **Operating System** from the drop-down list, and click **Add Server to Cluster**.



8. Select the server from list and click **OK**.



9. Repeat steps 6-7 for each XenServer host that will be part of the XenServer pool.
10. Click **OK** to create the server cluster.



After creating the server and server cluster objects, volumes can be created and mapped to the servers. For XenServers in a pool, map LUNs to the server cluster object to ensure all servers use the same LUN number.

Continue to [section 7](#) to create Storage Resources from these volumes on the XenServer or XenServer pool.

Note: See the [Create a Cluster Object in Enterprise Manager](#) and [Creating and mapping a volume in Enterprise Manager](#) videos for additional information.

5 XenServer Storage using iSCSI HBA

If using an iSCSI HBA to create an iSCSI SR, use the CLI from the XenServer host or the BIOS-level management interface to update target information. Depending on which HBA is being used, the initiator IQN for the HBA needs to be configured. Given the type of HBA used, the documentation for that HBA should be consulted to configure the IQN. Once the IQN has been configured for the HBA, use Dell Storage Manager Client to create a new LUN. However, instead of using the XenServer IQN, specify the IQN of the various ports of the HBA. Do this for every XenServer host in the pool. The QLogic® HBA CLI is included in the XenServer host and located at the following absolute path:

```
/opt/QLogic_Corporation/QConvergeConsoleCLI/qauccli
```

If using Emulex® iSCSI HBAs, consult the Emulex documentation for instructions on installing and configuring the HBA.

5.1 Configure iSCSI HBA

Enter the following command:

```
[root@XenServer-01 ~]# /opt/QLogic_Corporation/QConvergeConsoleCLI/qauccli
```

If using Emulex iSCSI HBAs, consult the Emulex documentation for instructions on installing and configuring the HBA.

For the purposes of an example, this guide illustrates how the QLogic iSCSI HBA CLI quaccli can be used to configure an IP addresses on a dual-port QLE4062C iSCSI HBA adapter, add the iSCSI server to the SC Series storage, and configure a LUN for the server. This setup will also utilize multipathing since there are two iSCSI HBA ports.

1. From the XenServer console, launch a command prompt and execute the following:

```
/opt/QLogic_Corporation/QConvergeConsoleCLI/qauccli
```

Note: This configuration can also be performed during the server boot by pressing [Ctrl] + [Q] when prompted.

```

-----
Program Version: 1.1.00.13 Driver Version: 5.01.00.08 IC: 2
FW Version: 3.0.1.33 Type: Copper
Current HBA/Port Information: HBA Alias:
HBA: 0 Port: 0 HBA Port Index: 1 HBA Model: QLE4062C
IP Address: 0.0.0.0 Link: Up
IPv6 Protocol is currently disabled.
Port iSCSI Name: iqn.2000-04.com.qlogic:qle4062c.gs40717a35854.1
Port iSCSI Alias:
-----

1. Display Program Version Information
2. Host Level Info & Operations
3. HBA Level Info & Operations
4. Port Level Info & Operations
5. List All QLogic iSCSI HBA Ports detected
6. Help
7. Select HBA Port
8. Refresh
9. Exit
enter selection: █

```

2. Configure the IP address for the iSCSI HBA:
 - a. To set the IP address for the HBA, choose option 4 (Port Level Info & Operations), then option 2 (Port Network Settings Menu).
 - b. Enter option 4 (Select HBA Port) to select the appropriate HBA port, then select option 2 (Configure IP Settings).

```

-----
Program Version: 1.1.00.13 Driver Version: 5.01.00.08 IC: 2
FW Version: 3.0.1.33 Type: Copper
Current HBA/Port Information: HBA Alias:
HBA: 0 Port: 0 HBA Port Index: 1 HBA Model: QLE4062C
IP Address: 0.0.0.0 Link: Up
IPv6 Protocol is currently disabled.
Port iSCSI Name: iqn.2000-04.com.qlogic:qle4062c.gs40717a35854.1
Port iSCSI Alias:
-----

1. Display Program Version Information
2. Host Level Info & Operations
3. HBA Level Info & Operations
4. Port Level Info & Operations
5. List All QLogic iSCSI HBA Ports detected
6. Help
7. Select HBA Port
8. Refresh
9. EXIT
enter selection: 4
PORT Level Info & Operations Menu
-----

Program Version: 1.1.00.13 Driver Version: 5.01.00.08 IC: 2
FW Version: 3.0.1.33 Type: Copper
Current HBA/Port Information: HBA Alias:
HBA: 0 Port: 0 HBA Port Index: 1 HBA Model: QLE4062C
IP Address: 0.0.0.0 Link: Up
IPv6 Protocol is currently disabled.
Port iSCSI Name: iqn.2000-04.com.qlogic:qle4062c.gs40717a35854.1
Port iSCSI Alias:
-----

1. Port Information
2. Port Network Settings Menu
3. Edit Configured Port Settings Menu
4. Port Restore Factory Defaults
5. Port Diagnostic Menu
6. Bootcode Settings Menu
7. --> Target Level Info & Operations
8. Select HBA Port
9. Refresh
10. EXIT
enter selection: 2
Port Network Settings Menu

```

- c. Enter the appropriate IP settings for the HBA adapter port. When finished, exit and save or select another HBA port to configure. In this example another HBA port will be configured.

```

-----
Program Version: 1.1.00.13 Driver Version: 5.01.00.08 IC: 2
FW Version: 3.0.1.33 Type: Copper
Current HBA/Port Information: HBA Alias:
HBA: 0 Port: 0 HBA Port Index: 1 HBA Model: QLE4062C
IP Address: 0.0.0.0 Link: Up
IPv6 Protocol is currently disabled.
Port iSCSI Name: iqn.2000-04.com.qlogic:qle4062c.gs40717a35854.1
Port iSCSI Alias:
-----

1. Display Network Settings
2. Configure IP Settings
3. iSNS Settings
4. Select HBA Port
5. Save changes and reset HBA (if necessary)
6. Refresh
7. Exit
enter selection: 2
Enable IPv4 [on] :
DHCP to obtain IPv4 Network Information: [off] :
IP_Address [0.0.0.0] :10.10.36.121
IP_Subnet_Mask [0.0.0.0] :255.255.0.0
IP_Gateway [0.0.0.0] :10.10.10.1
Enable IPv6 [off] :

```

- d. From the Port Network Settings menu, select option 4 to select an additional HBA port to configure. Enter 2 to select the second HBA port.
- e. Once the second HBA port is selected, choose option 2 (Configure IP Settings) from the Port Network Settings menu to input the appropriate IP settings for the second HBA port.

```

1. Display Network Settings
2. Configure IP Settings
3. iSNS Settings
4. Select HBA Port
5. Save changes and reset HBA (if necessary)
6. Refresh
7. Exit
enter selection: 4
1. HBA: 0 Port: 0 HBA Port Index: 1 HBA Model: QLE4062C
  HBA Serial Number: (G340717A35854) FW Version: 3.0.1.33 Type: Copper
  IP Address: 10.10.36.121
  Alias:
  iSCSI Name: iqn.2000-04.com.qlogic:qle4062c.gs40717a35854.1
2. HBA: 0 Port: 1 HBA Port Index: 2 HBA Model: QLE4062C
  HBA Serial Number: (G340717A35854) FW Version: 3.0.1.33 Type: Copper
  IP Address: 0.0.0.0
  Alias:
  iSCSI Name: iqn.2000-04.local.techsol:tundra1
Select an HBA Port Index : 2
Port Network Settings Menu
-----
Program Version: 1.1.00.13 Driver Version: 5.01.00.08 IC: 2
FW Version: 3.0.1.33 Type: Copper
Current HBA/Port Information: HBA Alias:
HBA: 0 Port: 1 HBA Port Index: 2 HBA Model: QLE4062C
IP Address: 0.0.0.0 Link: Up
IPv6 Protocol is currently disabled.
Port iSCSI Name: iqn.2000-04.local.techsol:tundra1
Port iSCSI Alias:
-----

1. Display Network Settings
2. Configure IP Settings
3. iSNS Settings
4. Select HBA Port
5. Save changes and reset HBA (if necessary)
6. Refresh
7. Exit
enter selection: 2
Enable IPv4 [on] :
DHCP to obtain IPv4 Network Information: [off] :
IP_Address [0.0.0.0] :10.10.36.122
IP_Subnet_Mask [0.0.0.0] :255.255.0.0
IP_Gateway [0.0.0.0] :10.10.10.1
Enable IPv6 [off] :

```

- f. Choose option 5 (Save changes and reset HBA (if necessary)), then select **Exit** until the main menu is displayed.

The iSCSI name or IQN can also be changed using the qaucli utility:

1. Select option 4 (Port Level Info & Operations) from the main menu.
2. Select option 3 (Edit Configured Port Settings).
3. Select option 3 (Port Firmware Settings).
4. Select option 7 (Configure Advanced Settings).
5. Press <Enter> until reaching iSCSI_Name.
6. Enter a unique IQN name for the adapter.

5.2 Connect to SC Series iSCSI control ports

The next step is to establish a target from XenServer so that it registers with the SC Series storage.

1. From the main interactive qaucli menu, select option 4 (Port Level Info & Operations).
2. From the Port Level Info & Operations menu, select option 7 (Target Level Info & Operations).
3. On the HBA target menu screen, select option 6 (Add a Target).
 - a. Press [Enter] until reaching the TGT_TargetIPAddress option.
 - b. Enter the target IP address of the SC Series iSCSI control ports, and repeat for each target. In this example, 10.10.10.100 and 10.20.10.100 are used, which are the iSCSI control ports on the SC Series storage.

```
HBA Target Menu
-----
Program Version: 1.1.00.13 Driver Version: 5.01.00.08 IC: 2
FW Version: 3.0.1.33 Type: Copper
Current HBA/Port Information: HBA Alias:
HBA: 0 Port: 1 HBA Port Index: 2 HBA Model: QLE4062C
IP Address: 10.10.36.122 Link: Up
IPv6 Protocol is currently disabled.
Port iSCSI Name: iqn.2000-04.local.techsol:tundra.iscsi.1
Port iSCSI Alias:
-----

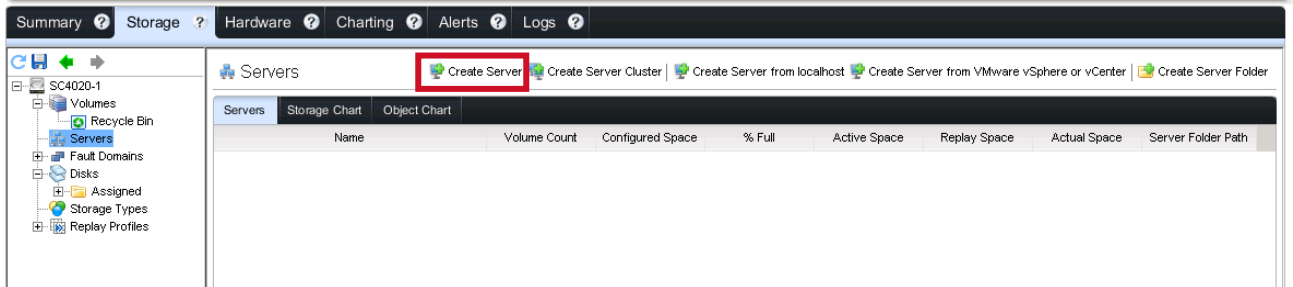
 1. List Targets
 2. Display Target Information
 3. Bind Target
 4. Unbind Target (persistent targets only)
 5. Configure Target Parameters
 6. Add A Target
 7. Configure Target Authentication Menu
 8. List LUN information
 9. Save Target/CHAP changes
10. Select HBA Port
11. Refresh
12. Exit
enter selection: 6
IPv6 Target? [off]:
TGT_iSCSI_Name [] :
TGT_Port [3260] :
TGT_TargetIPAddress [0.0.0.0] :10.10.64.1
```

- c. Once all targets are entered for HBA 0, select option 9 to the save the port information.
 - d. Select option 10 to select the second HBA port.
4. Enter option 12 to exit. Enter YES to save the changes.
 5. Exit out of the qaucli utility.

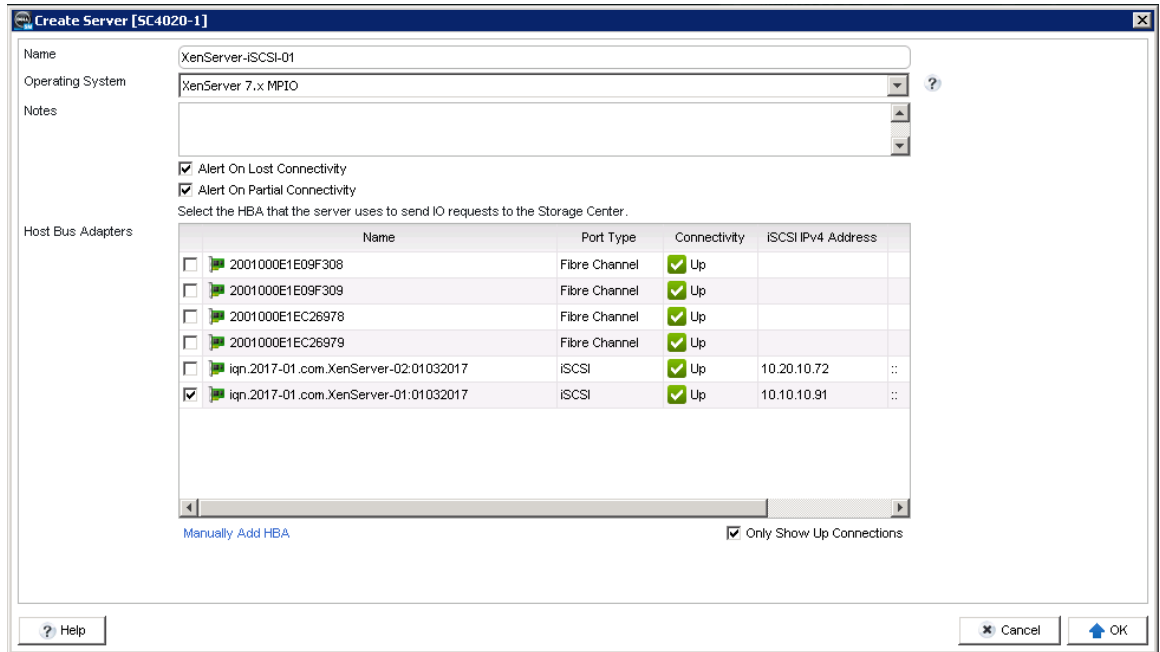
5.3 Configure Server Objects in Dell Storage Manager

Use the following steps to configure the server object for access to the SC Series storage:

1. In Dell Storage Manager Client, go to the **Storage** tab, drill down through the array to highlight the **Servers** object, and click **Create Server**.

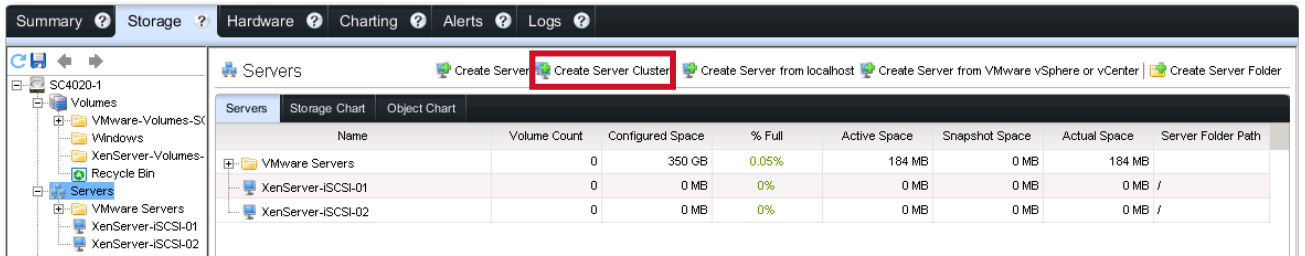


2. Enter the Name to identify the XenServer host, select the Operating System from the drop-down list, select the appropriate IQN listed in the Host Bus Adapters section, and click **OK**.

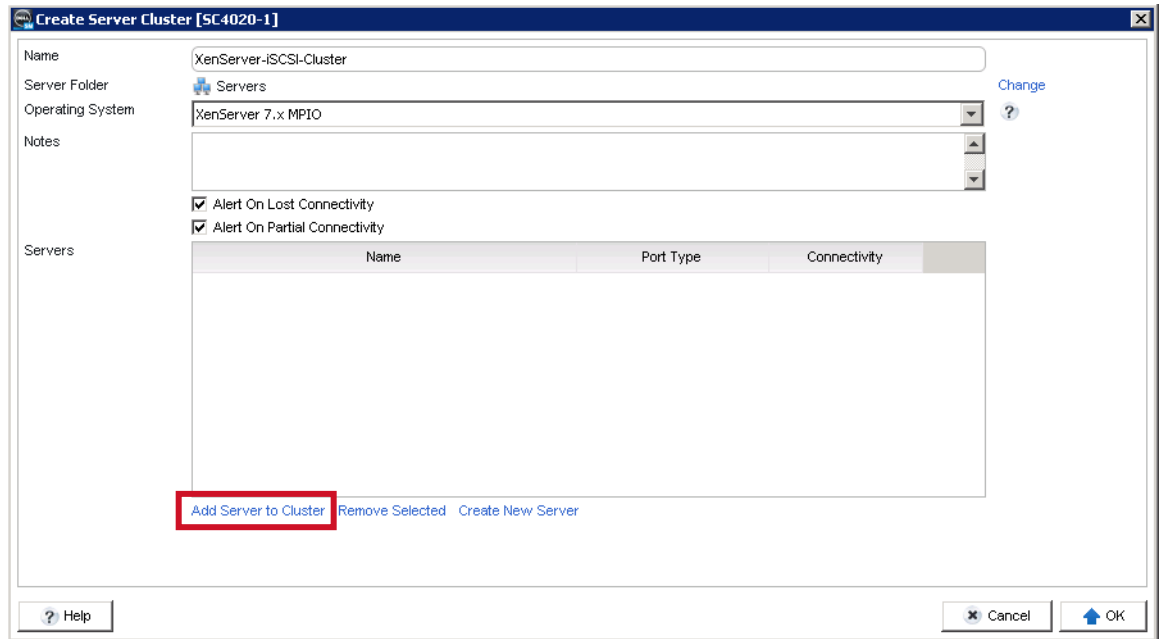


3. Repeat steps 1–2 for each XenServer to be added to the pool.

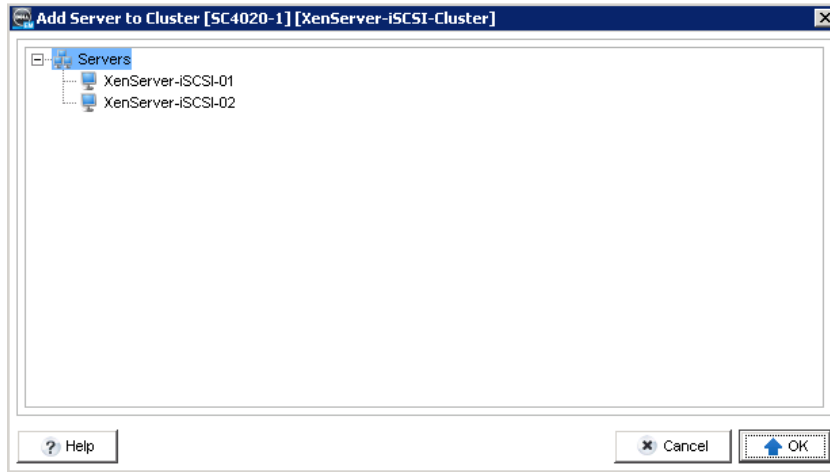
4. With **Servers** still highlighted, click **Create Server Cluster**.



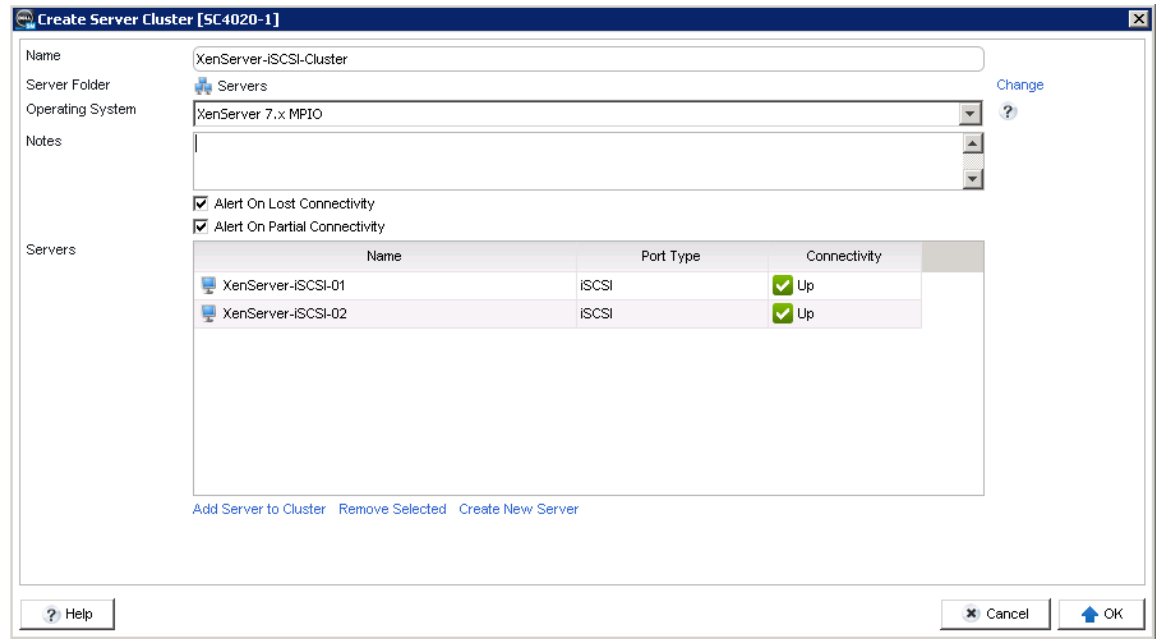
5. Enter the **Name** to identify the XenServer Cluster, select the **Operating System** from the drop-down list, and click **Add Server to Cluster**.



6. Select the server from the list and click **OK**.



7. Repeat steps 5–6 for each XenServer host that will be part of the XenServer pool.
8. Click OK to create the server cluster.



After creating the server and server cluster objects, volumes can be created and mapped to the servers. For XenServers in a pool, map LUNs to the server cluster object to ensure all servers use the same LUN number.

Continue to [section 7](#) to create Storage Resources from these volumes on the XenServer or XenServer pool.

Note: See the [Create a Cluster Object in Enterprise Manager](#) and [Creating and mapping a volume in Enterprise Manager](#) videos for additional information.

6 XenServer Storage using Fibre Channel HBA

XenServer Fibre Channel (FC) Storage Repositories are supported with SC Series storage through the use of FC HBAs. Supported HBA brands include QLogic and Emulex.

Shared FC SRs using a FC HBA are capable of supporting VM agility using XenMotion: VMs can be started on any XenServer host in a resource pool and migrated between them with no noticeable interruption.

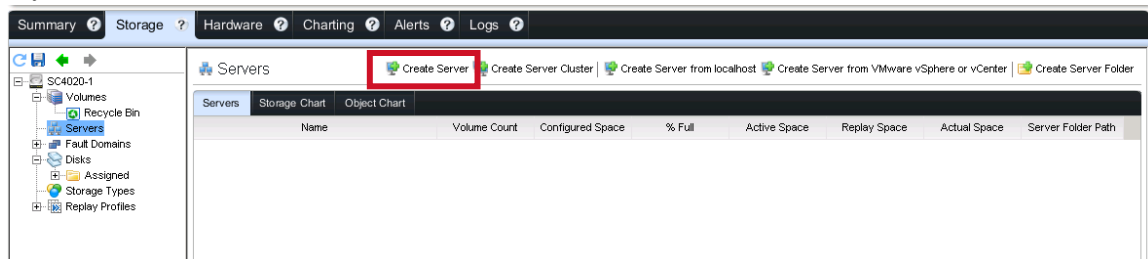
The following sections detail the steps involved in adding a new Fibre-Channel-connected volume to a XenServer pool. Be sure Multipathing has been enable a shown in section [3.4.1](#).

This section assumes all Fibre Channel connections have been properly configured and zoned following best practices found in the Dell Enterprise Solutions document, [How to zone a Fibre Channel Compellent array](#).


6.1 Configure Server Objects in Dell Storage Manager

Use the following steps to configure the server object for access to the SC Series storage:


1. In Dell Storage Manager, go to the **Storage** tab, navigate through the array to highlight the servers object, and click **Create Server**.



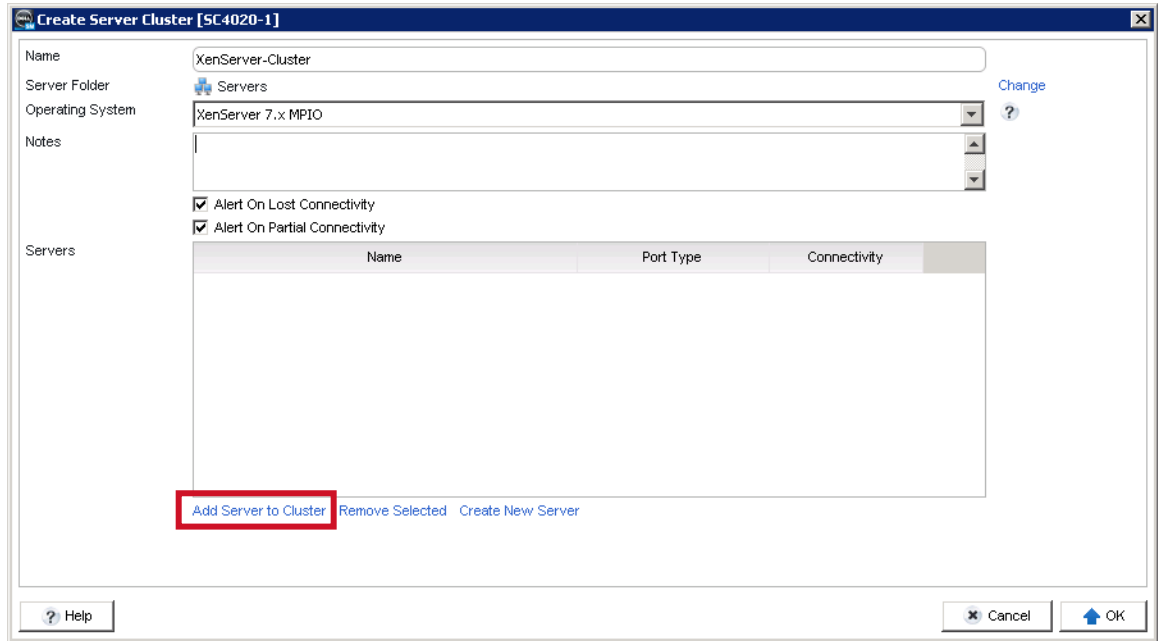
2. Enter the name to identify the XenServer host, select the operating system from the drop-down list, select the appropriate FC WWN listed in the Host Bus Adapters section, and click **OK**.




3. Repeat steps 1–2 for each XenServer to be added to the pool.
4. With **Servers** still highlighted, click **Create Server Cluster**.



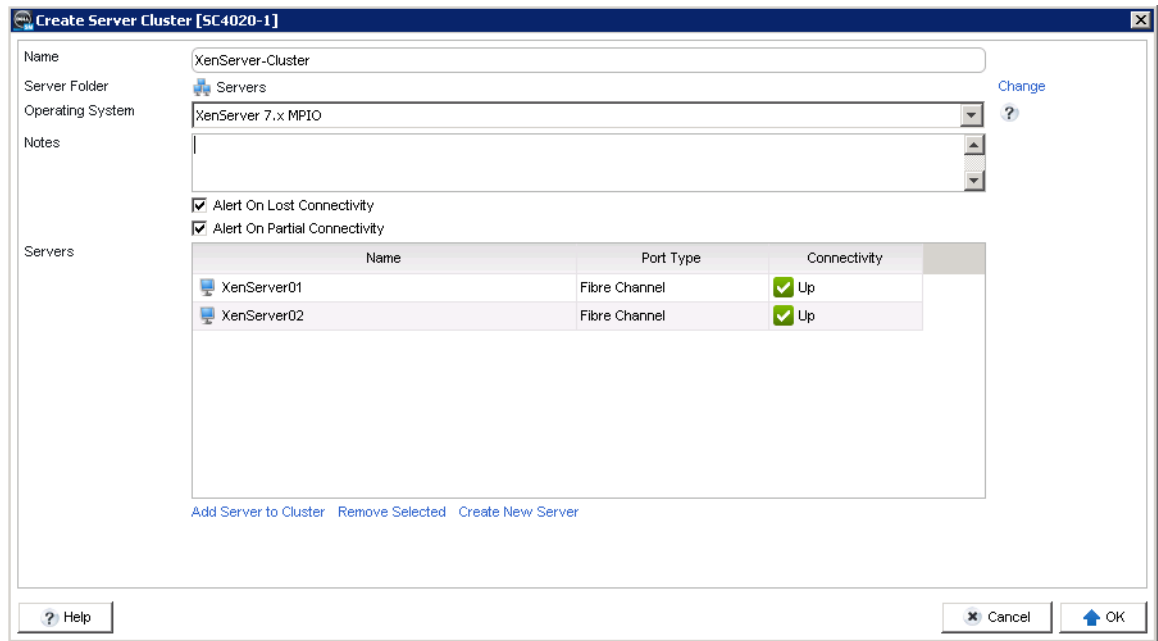
5. Enter the name to identify the XenServer cluster, select the operating system from the drop-down list, and click **Add Server to Cluster**.



6. Select the server from the list and click **OK**.



7. Repeat steps 5–6 for each XenServer host that will be part of the XenServer pool.
8. Click **OK** to create the server cluster.



After creating the server and server cluster objects, volumes can be created and mapped to the servers. For XenServers in a pool, map LUNs to the server cluster object to ensure all servers use the same LUN number.

Continue to [section 7](#) to create Storage Resources from these volumes on the XenServer or XenServer pool.

Note: See the [Create a Cluster Object in Enterprise Manager](#) and [Creating and mapping a volume in Enterprise Manager](#) videos for additional information.

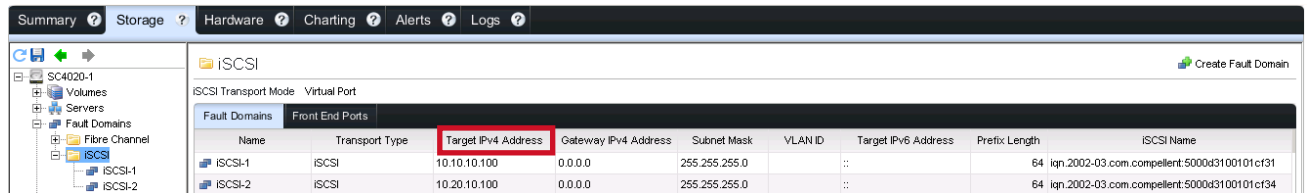
7 Create new Storage Repository (SR)

Once the volumes are mapped to the servers, they can be added to the XenServer using XenCenter or the CLI. The following steps detail adding storage using XenCenter. The steps for adding storage through the CLI can be found in the *Citrix XenServer 7.0 Administrator's Guide* available on the [Citrix Product Documentation](#) website.

7.1 Software iSCSI SRs

7.1.1 Identify SC Series storage iSCSI targets

To gather SC Series storage iSCSI target information: Within Dell Storage Manager, go to the **Storage** tab, navigate to **Fault Domains**, and highlight **iSCSI**. This should display the target IPv4 addresses of the control ports; one for each fault domain. The target addresses should be on the same IP subnet as the server's storage NICs. Figure 8 shows a multipath-dual subnet example, where Figure 9 shows a single subnet bonded example.

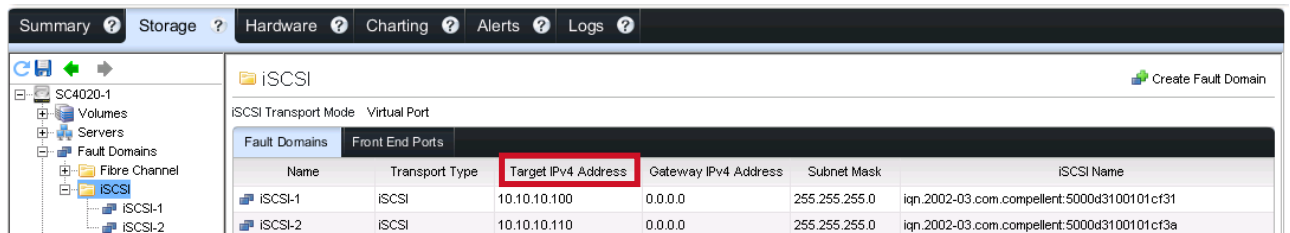


Name	Transport Type	Target IPv4 Address	Gateway IPv4 Address	Subnet Mask	VLAN ID	Target IPv6 Address	Prefix Length	iSCSI Name
iSCSI-1	iSCSI	10.10.10.100	0.0.0.0	255.255.255.0		::	64	iqn.2002-03.com.compellent:5000d3100101cf31
iSCSI-2	iSCSI	10.20.10.100	0.0.0.0	255.255.255.0		::	64	iqn.2002-03.com.compellent:5000d3100101cf34

Figure 8 Multipath – Dual Subnets: Identifying control port IP addresses

In the multipath – Dual subnet example, the IP addresses are:

- 10.10.10.100/24
- 10.20.10.100/24



Name	Transport Type	Target IPv4 Address	Gateway IPv4 Address	Subnet Mask	iSCSI Name
iSCSI-1	iSCSI	10.10.10.100	0.0.0.0	255.255.255.0	iqn.2002-03.com.compellent:5000d3100101cf31
iSCSI-2	iSCSI	10.10.10.110	0.0.0.0	255.255.255.0	iqn.2002-03.com.compellent:5000d3100101cf3a

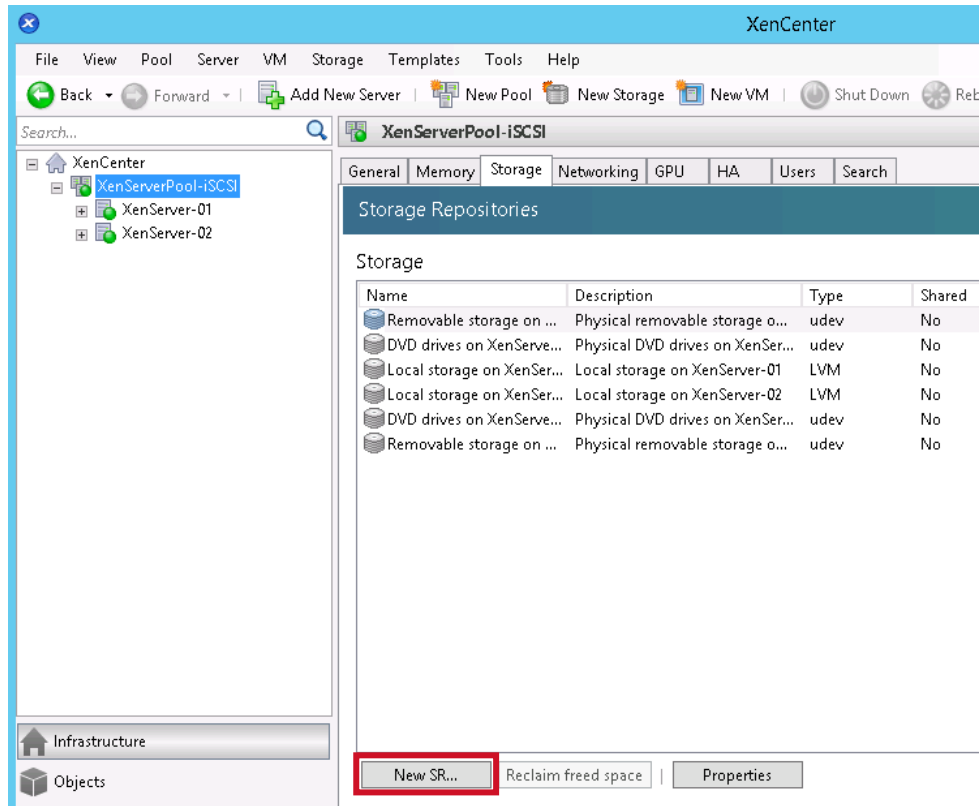
Figure 9 Multipath – Single Subnet (Bonded NIC): Identifying control port IP address

In the multipath – Single subnet using Bonded NICs example, the IP address is:

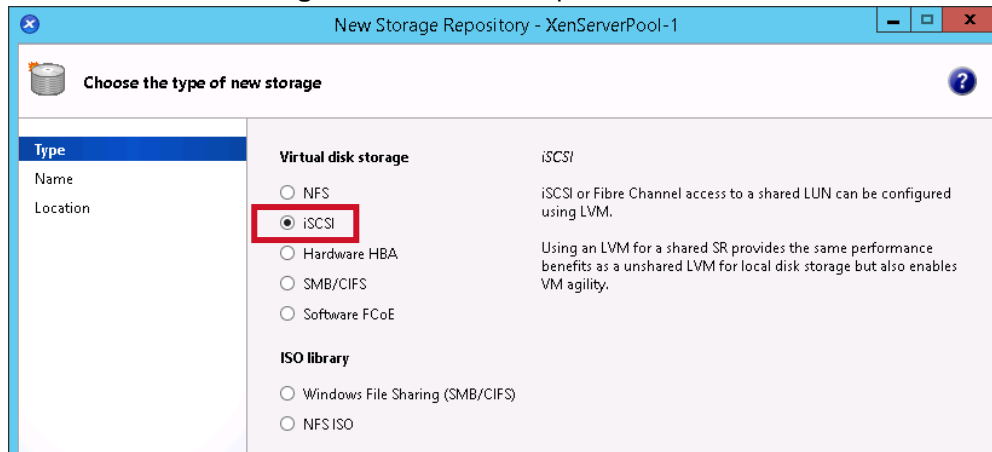
- 10.10.10.100/24
- 10.10.10.110/24

7.1.2 Create new Software iSCSI SR

1. In XenCenter, navigate to the XenServer Pool, select the **Storage** tab, and click **New SR**.

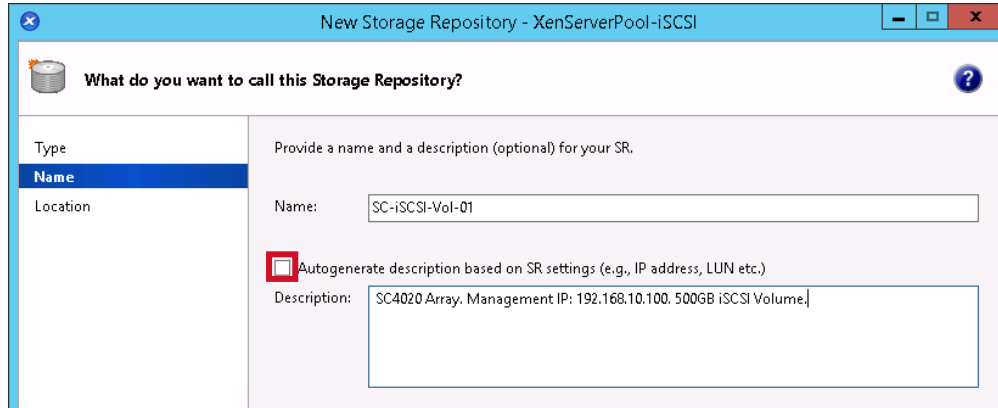


2. Under **Virtual disk storage**, select the **iSCSI** option, and click **Next**.

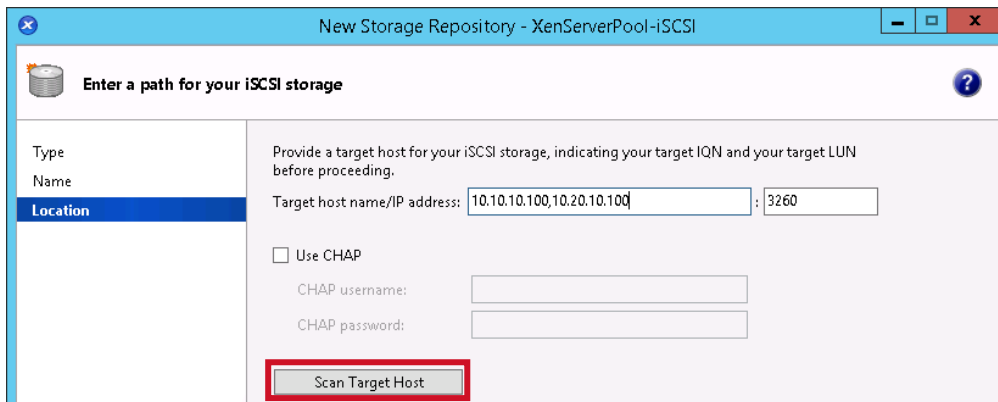


3. Give the new storage repository a name and click **Next**.

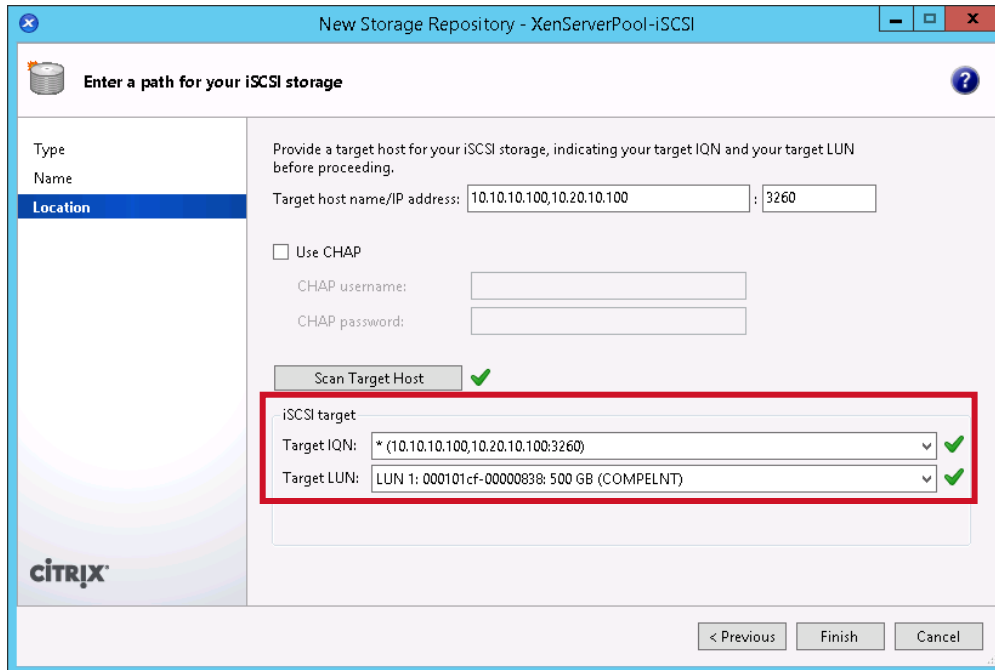
Note: Uncheck the box, **Autogenerate description based on SR settings**, to enter additional description information such as SC Series array IP address, array model information, capacity, or administrative contact information.



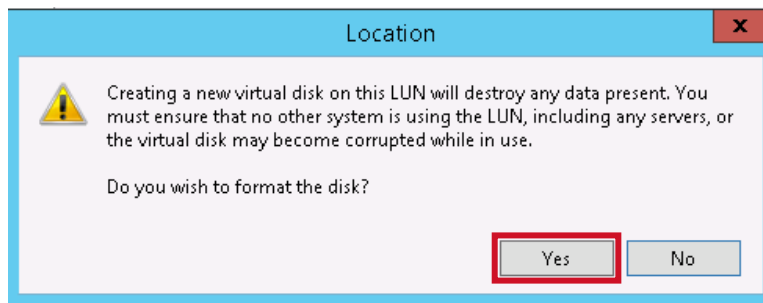
4. In the Target hostname/IP address field, enter both SC Series storage iSCSI control port IP addresses, comma-separated, and click **Scan Target Host**.



5. In the **Target IQN** drop-down list, select the **(Control Port IP Address)*.
6. In the **Target LUN** drop-down list, select the LUN on which the new SR will reside.
7. Click **Finish**.



8. Click **Yes** to format the disk.



The new SR is now available to the server or pool. Repeat these steps for mapping and adding storage for any additional SRs.

7.1.3 Verify multipath status

To view the status of the multipath from the XenCenter GUI, select the new SR from the list of object and expand the Multipathing dropdown section as shown in Figure 10.

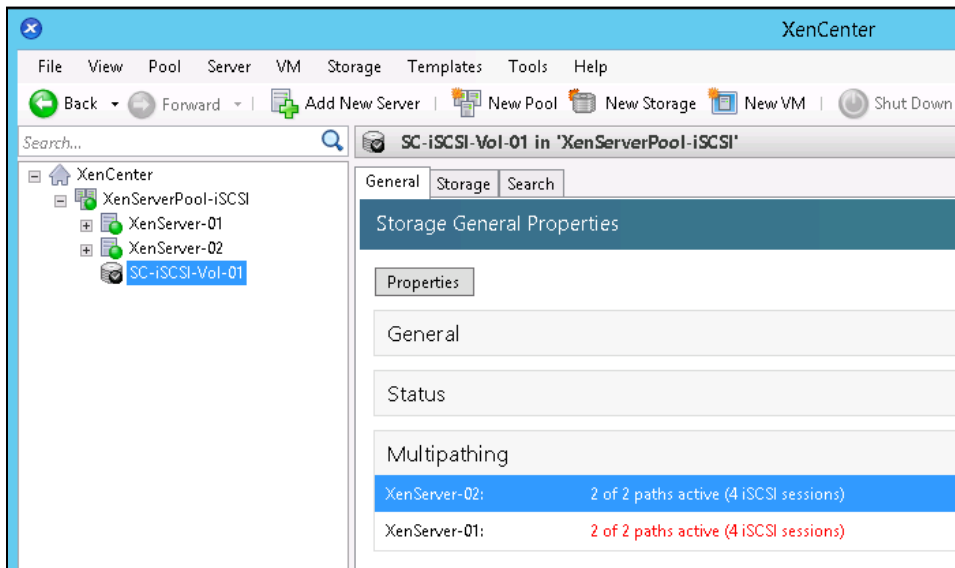


Figure 10 Display multipath status from GUI

To view the status of the multipath from the CLI, use the following command:

```
mpathutil status
```

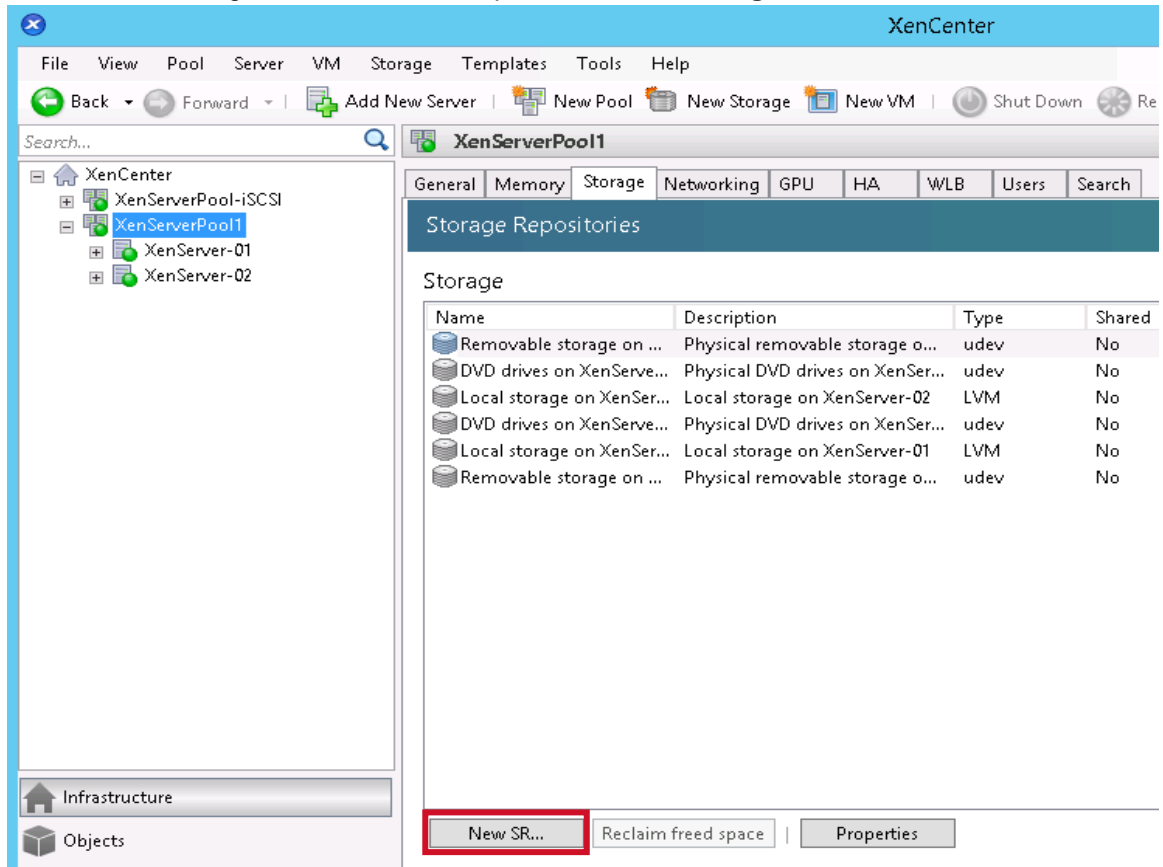
```
[root@XenServer-01 ~]# mpathutil status
show topology
36000d3100101cf000000000000000838 dm-1 COMPELNT,Compellent Vol
size=500G features='1 queue_if_no_path' hwhandler='0' wp=rw
`-+- policy='service-time 0' prio=1 status=active
  |- 64:0:0:1 sdc 8:32 active ready running
  `-- 67:0:0:1 sdd 8:48 active ready running
[root@XenServer-01 ~]#
```

Figure 11 Display multipath status from CLI

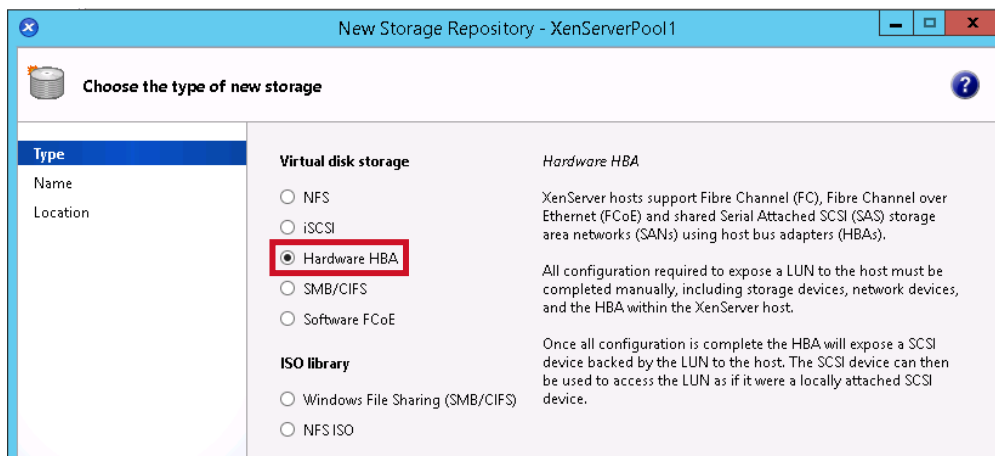
Note: If mpathutil status does not show multiple paths as expected, confirm multipathing was enabled as shown in [section 3.4.1](#).

7.2 Create SR with hardware HBA (iSCSI and FC)

1. In XenCenter, navigate to the XenServer pool, select the **Storage** tab, and click **New SR**.



2. Under **Virtual disk storage**, select the **Hardware HBA** option, and click **Next**.



3. Give the new storage repository a name and click **Next**.

Note: Uncheck the box, **Autogenerate description based on SR settings**, to enter additional description information such as SC Series array IP address, array model information, capacity, or administrative contact information.

What do you want to call this Storage Repository?

Type

Name

Location

Provide a name and a description (optional) for your SR.

Name: SC-FC-Vol-01

Autogenerate description based on SR settings (e.g., IP address, LUN etc.)

Description: SC4020 Array. Management IP: 192.168.10.100. 500GB FC Volume.

4. Select the desired LUN from the list and click **Next**.

Select the LUN to reattach or create a new SR on

Type

Name

Location

Summary

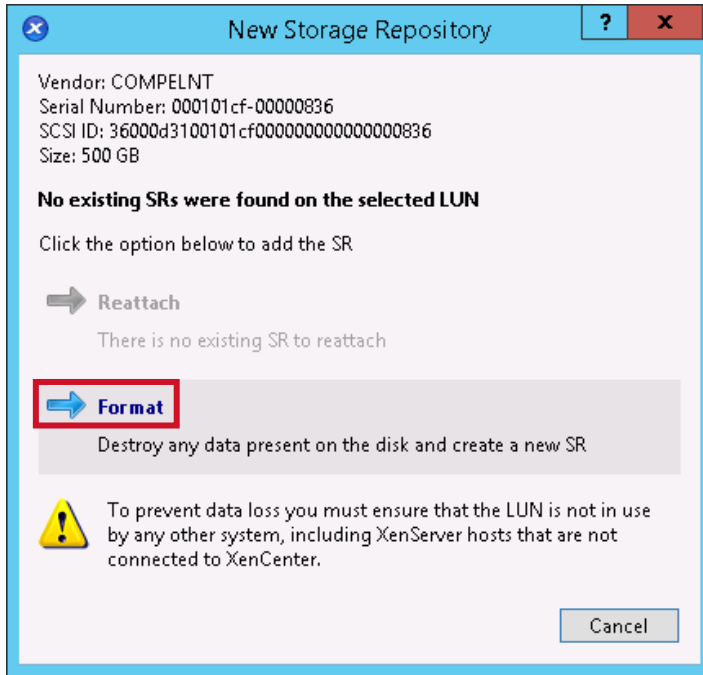
Choose LUNs to create a new SR on:

Size	Serial	ID	Details
COMPELNT			
500 GB	000101cf-00000836	36000d3100101cf00000...	8:0:1:1

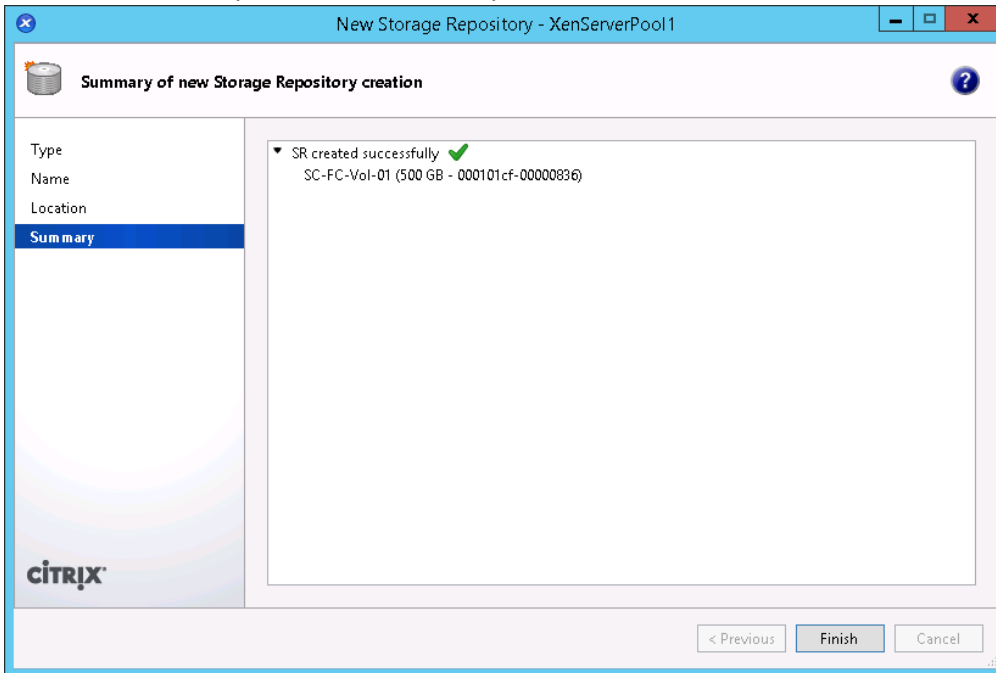
Select All Clear All

< Previous Next > Cancel

5. If this is a new SR (has not been created or used before), click **Format** to prepare the SR for use.



6. Click **Finish** to complete the SR creation process and make the SR available to the server or pool



A Technical support and resources

Dell.com/support is focused on meeting customer needs with proven services and support.

[Dell TechCenter](#) is an online technical community where IT professionals have access to numerous resources for Dell EMC software, hardware and services.

[Storage Solutions Technical Documents](#) on Dell TechCenter provide expertise that helps to ensure customer success on Dell EMC Storage platforms.

A.1 Related documentation

See the following referenced or recommended Dell and Citrix XenServer publications:

- [Citrix XenServer 7.0 Installation Guide](#)
- [Citrix XenServer 7.0 Administrator's Guide](#)
- [Citrix XenCenter Managing Storage Repositories \(SRs\) Documents](#)
- Dell Enterprise Manager video: [Create a Cluster Object in Enterprise Manager](#)
- Dell Enterprise Manager video: [Creating and mapping a volume in Enterprise Manager](#)
- Dell Enterprise Solutions: [How to zone a Fibre Channel Compellent array](#)
- [Dell Storage Compatibility Matrix](#)
- [Switch Configuration Guides for Dell PS Series or SC Series storage](#)