



Using Dell EqualLogic and Multipath I/O with Citrix XenServer 6.2

Dell Engineering | Donald Williams
November 2013

Revisions

Date	Description
November 2013	Initial release

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

© 2013 Dell Inc. All rights reserved. Reproduction of this material in any manner whatsoever without the express written permission of Dell Inc. is strictly forbidden. For more information, contact Dell.

PRODUCT WARRANTIES APPLICABLE TO THE DELL PRODUCTS DESCRIBED IN THIS DOCUMENT MAY BE FOUND AT: <http://www.dell.com/learn/us/en/19/terms-of-sale-commercial-and-public-sector> Performance of network reference architectures discussed in this document may vary with differing deployment conditions, network loads, and the like. Third party products may be included in reference architectures for the convenience of the reader. Inclusion of such third party products does not necessarily constitute Dell's recommendation of those products. Please consult your Dell representative for additional information.

Trademarks used in this text:

Dell™, the Dell logo, Dell Boomi™, Dell Precision™, OptiPlex™, Latitude™, PowerEdge™, PowerVault™, PowerConnect™, OpenManage™, EqualLogic™, Compellent™, KACE™, FlexAddress™, Force10™ and Vostro™ are trademarks of Dell Inc. Other Dell trademarks may be used in this document. Cisco Nexus®, Cisco MDS®, Cisco NX-OS®, and other Cisco Catalyst® are registered trademarks of Cisco System Inc. EMC VNX®, and EMC Unisphere® are registered trademarks of EMC Corporation. Intel®, Pentium®, Xeon®, Core® and Celeron® are registered trademarks of Intel Corporation in the U.S. and other countries. AMD® is a registered trademark and AMD Opteron™, AMD Phenom™ and AMD Sempron™ are trademarks of Advanced Micro Devices, Inc. Microsoft®, Windows®, Windows Server®, Internet Explorer®, MS-DOS®, Windows Vista® and Active Directory® are either trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Red Hat® and Red Hat® Enterprise Linux® are registered trademarks of Red Hat, Inc. in the United States and/or other countries. Novell® and SUSE® are registered trademarks of Novell Inc. in the United States and other countries. Oracle® is a registered trademark of Oracle Corporation and/or its affiliates. Citrix®, Xen®, XenServer® and XenMotion® are either registered trademarks or trademarks of Citrix Systems, Inc. in the United States and/or other countries. VMware®, Virtual SMP®, vMotion®, vCenter® and vSphere® are registered trademarks or trademarks of VMware, Inc. in the United States or other countries. IBM® is a registered trademark of International Business Machines Corporation. Broadcom® and NetXtreme® are registered trademarks of Broadcom Corporation. Qlogic is a registered trademark of QLogic Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and/or names or their products and are the property of their respective owners. Dell disclaims proprietary interest in the marks and names of others.



Table of contents

Revisions.....	2
1 Preface.....	4
1.1 Audience.....	4
1.2 Related resources	4
2 Introduction.....	5
2.1 Objective.....	5
2.2 Revision information.....	5
3 Using XenServer software iSCSI adapter	6
4 Configuring iSCSI Multipath I/O.....	10
4.1 Check MPIO status	17
A Create a volume on a Dell EqualLogic PS Series Group.....	18
A.1 Create a volume	18
A.2 Create access control records on the volume	18
B XenServer iSCSI performance tuning.....	19
B.1 Tuning iSCSI performance	19
C Technical support and customer service	20



1 Preface

Dell EqualLogic PS Series arrays optimize resources by automating performance and network load balancing. These arrays also offer all-inclusive array management software, host software, and firmware updates at no additional cost.

1.1 Audience

The information in this guide is intended for technology professionals interested in using Dell EqualLogic storage in a Citrix XenServer environment.

1.2 Related resources

For detailed information about PS Series arrays, groups, volumes, array software, and host software, log in to the documentation page of the EqualLogic customer support site at <https://eqsupport.dell.com/secure/login.aspx>.

Additional information about Dell products and services is located at dell.com.

To learn more about EqualLogic products and new releases being planned, visit the EqualLogic page of the Dell TechCenter site at <http://delltechcenter.com/page/EqualLogic>. This site stores articles, demos, online discussions, technical documentation and other assets describing the benefits of EqualLogic products.

For an updated list of products compatible with EqualLogic, reference the following documents:

- The *Dell EqualLogic Configuration Guide* at <http://en.community.dell.com/techcenter/storage/w/wiki/2639.equallogic-configuration-guide.aspx>
- The *Dell EqualLogic Storage Compatibility Matrix* at <http://en.community.dell.com/techcenter/storage/w/wiki/2661.equallogic-compatibility-matrix-07262013.aspx>.



2 Introduction

In today's technology-driven world, data is one of the most valuable assets in any business organization; and it must be available 24 hours, seven days a week. A data outage of any duration and size can result in a tremendous loss of revenue. High availability is commonly achieved by using fault tolerant and redundant components. For example, multiple servers for redundancy, RAID on the disks or even failover clustering on the servers is not enough protection as it relies on a storage area network (SAN) for its data. To ensure high availability, access to shared storage must be readily available and protected against failure. As SANs become populated with more and more data, the potential loss of access to data storage components becomes problematic. This is why many businesses today consider Multipath I/O (MPIO) a must.

2.1 Objective

This document describes how to configure EqualLogic storage, including MPIO, in a Citrix XenServer version 6.2 environment using the software iSCSI adapter. It does not include information on storage connectivity via Citrix StorageLink.

2.2 Revision information

Table 1 Software and firmware used for this document.

Vendor	Model	Software Revision
Dell	PS-6100E	6.0.6, 7.0
Dell	PowerConnect 6248	3.3.3.5
Citrix	XenServer 6.2 (GA)	XenServer release 6.2.0-Build 70446

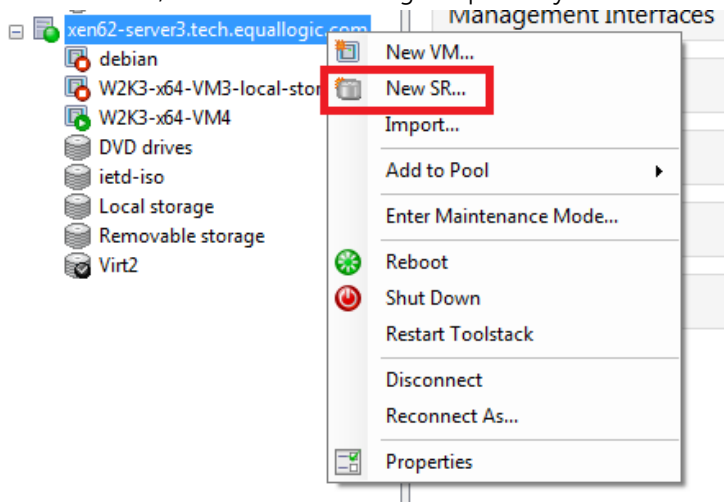
Table 2 Referenced documents

Vendor	Document Title
Dell	<i>PS Series Administrators Guide</i>
Dell	<i>PS Series CLI Guide</i>
Dell	<i>PS Series iSCSI Initiator and Operating System Considerations</i>
Citrix	<i>XenServer 6.2 Administrator Reference Guide</i>
Citrix	<i>XenServer 6.2 Installation Guide</i>

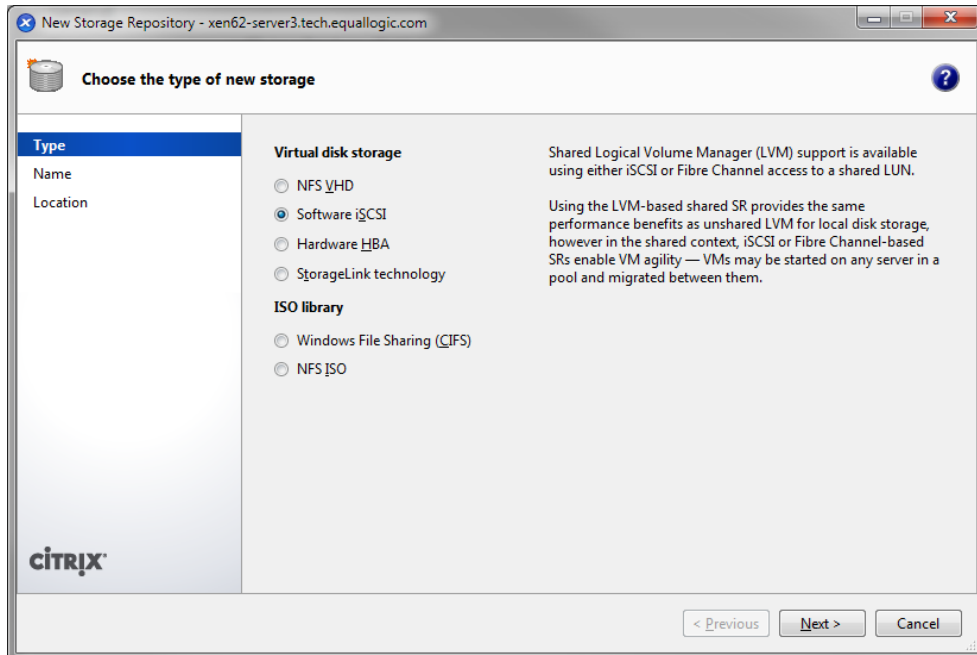


3 Using XenServer software iSCSI adapter

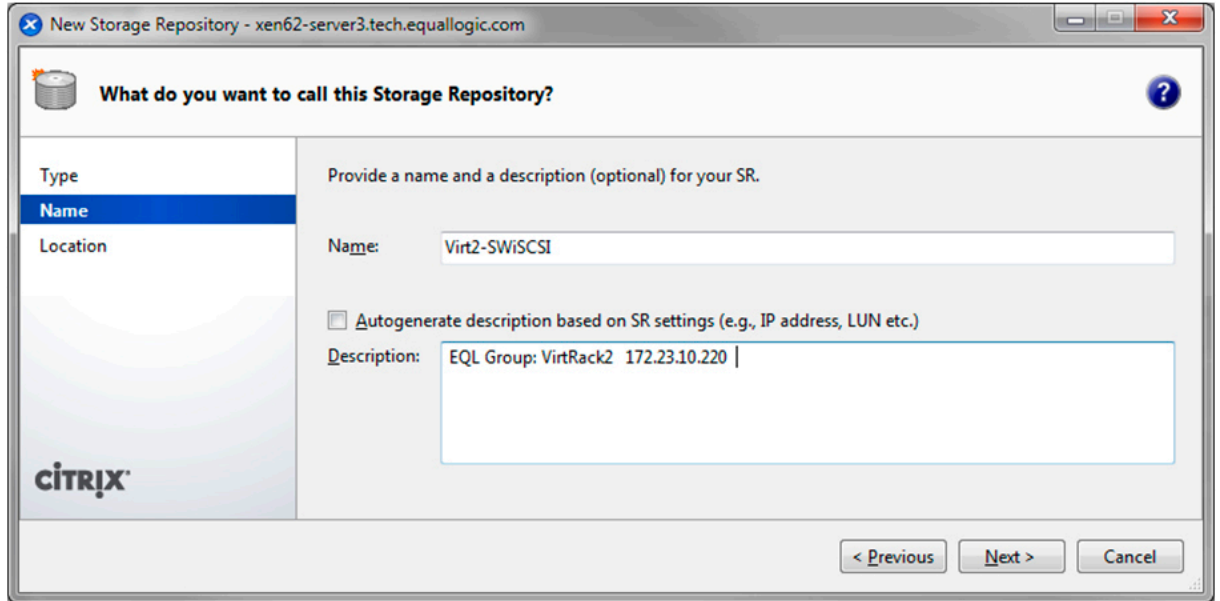
1. Using EqualLogic PS Series management tools, such as Group Manager or the CLI, create an EqualLogic volume. For additional information on creating an EqualLogic volume via the Group Manager CLI, please refer to [Appendix A](#).
 - a. Create access control records on the volume for the Citrix XenServer.
2. In XenCenter, create a new Storage Repository.



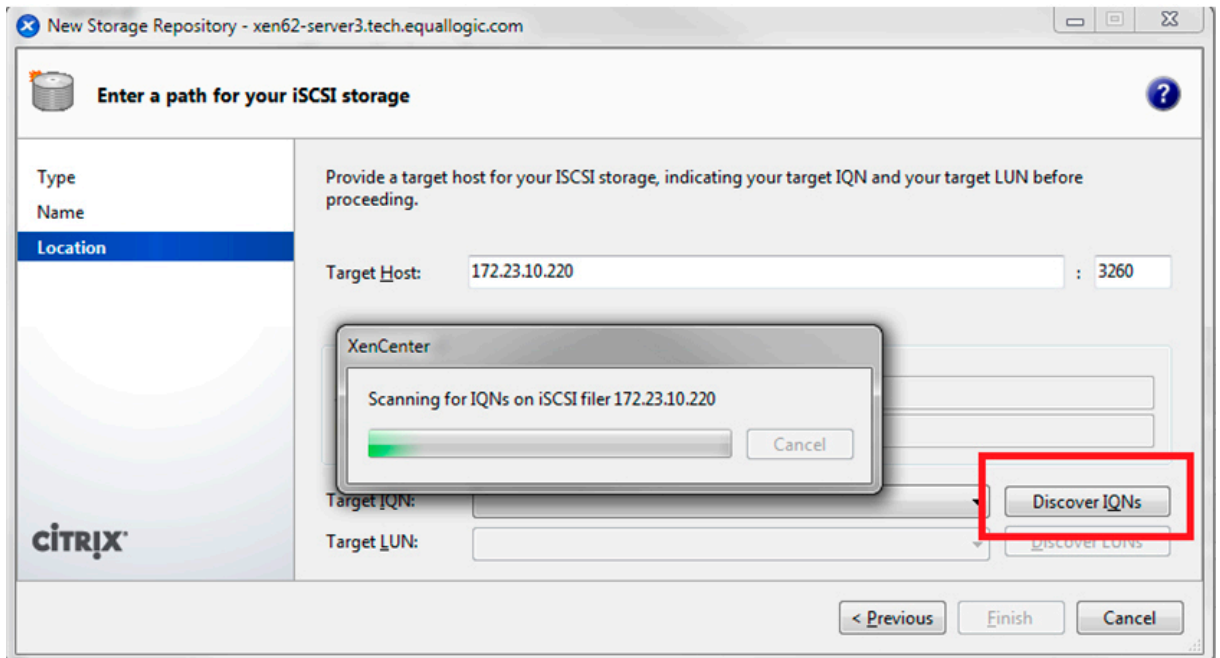
3. Choose the type of new storage.
 - a. Select **Software iSCSI** to use the software iSCSI adapter.
 - b. Click **Next**.




4. Provide a name and description for the Storage Repository.
 - a. To easily identify that the Storage Repository resides on EqualLogic storage, it is recommended to clear the **Autogenerate description based on SR settings** check box and then include the EqualLogic PS Series group information in the description.
 - b. Click **Next**.



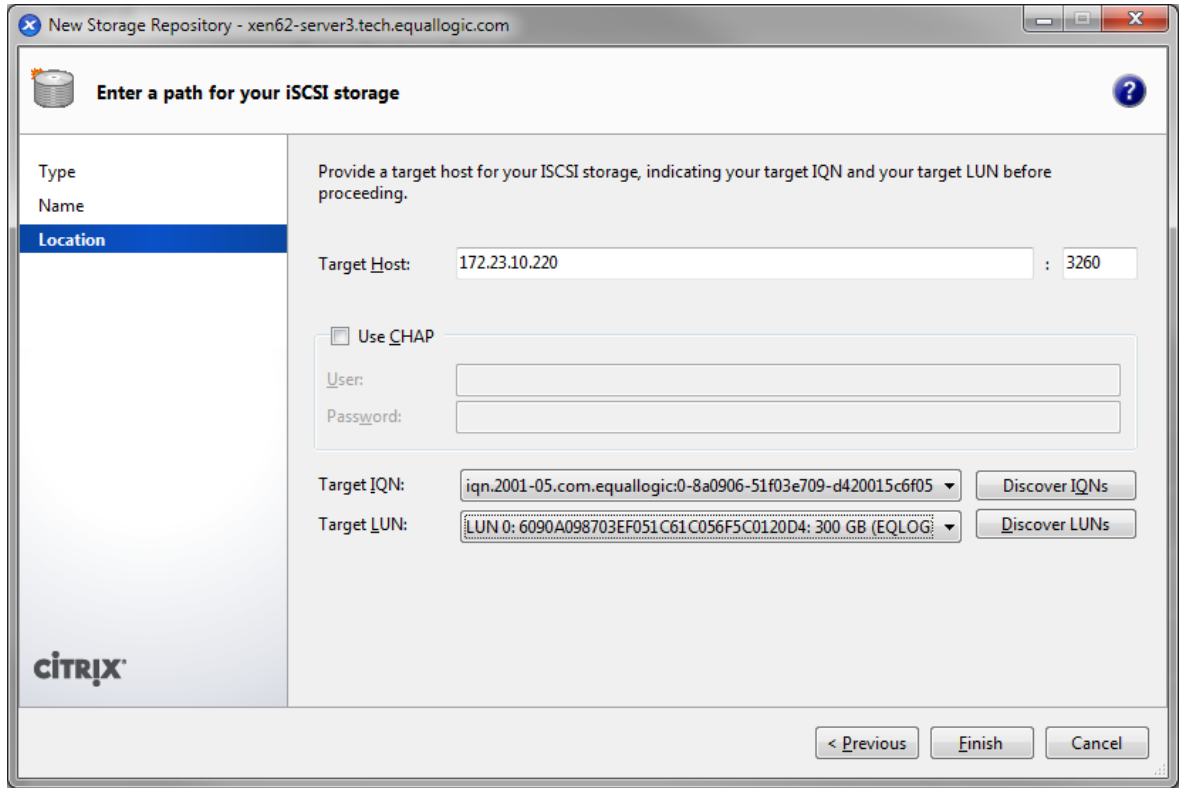
5. Enter a path for the iSCSI storage.
 - a. In the **Target Host** field, enter the EqualLogic Group IP address.
 - b. If CHAP is being used for access control, enter the username and password.
 - c. Click **Discover IQNs**.




6. In the **Target IQN** field, select the EqualLogic volume that was previously created.



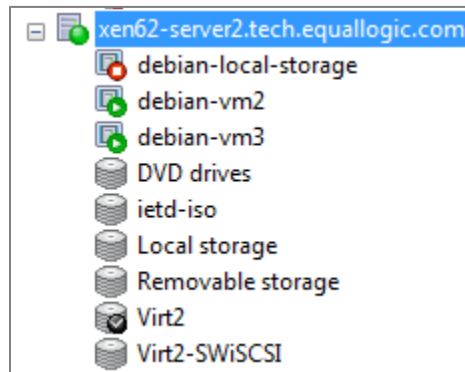
7. Click **Discover LUNs** to connect to the volume.



- Click **Yes** to format the disk and create the Storage Repository.



- The EqualLogic storage repository can be viewed from XenCenter. In addition, you can create, move or copy VMs to the Storage Repository.

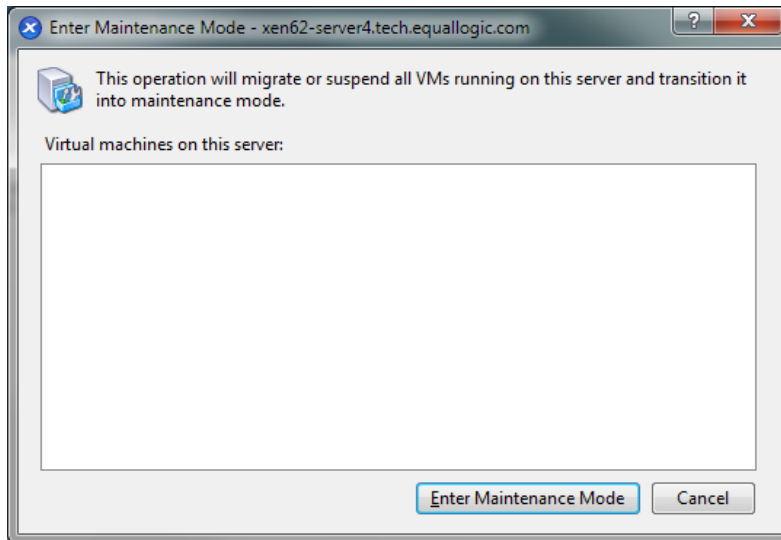


4 Configuring iSCSI Multipath I/O

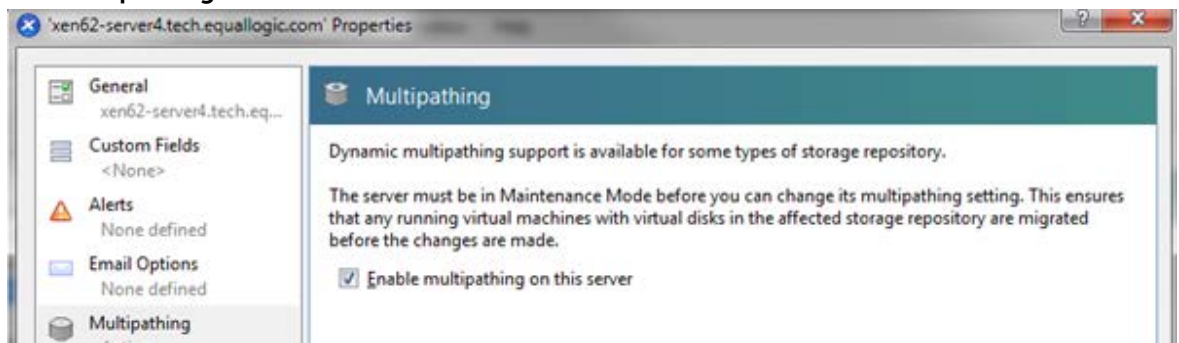
XenServer 6.2 natively supports MPIO with a single discovery address. Patches are available for XenServer 6.0.x and 6.1.x to provide this same feature. Please refer to the following URL for more information: <http://support.citrix.com/article/CTX138429>.

The server must be in maintenance mode before configuring multipathing.

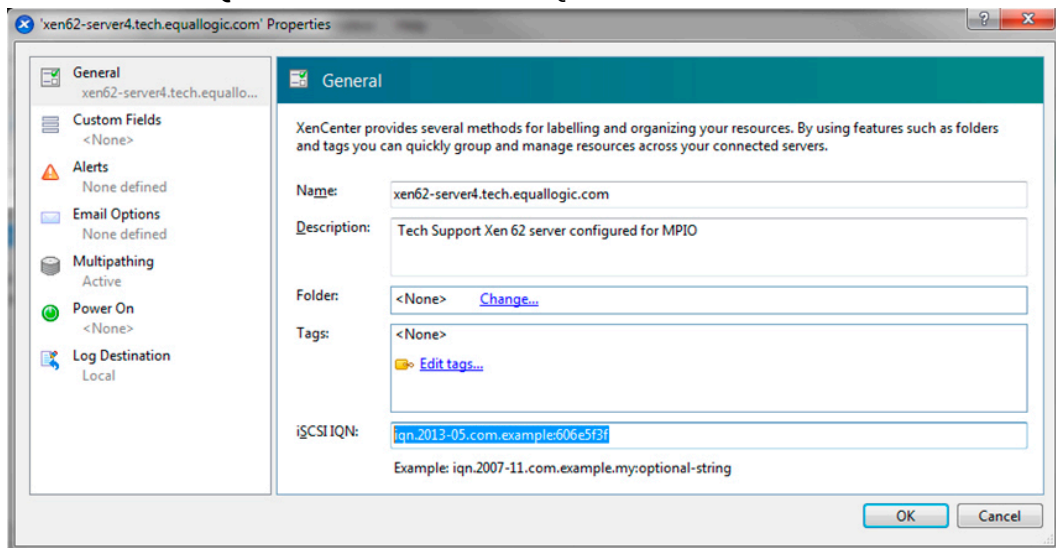
1. In the XenCenter GUI, put the XenServer in maintenance mode.
 - a. Right-click on the server, and select **Enter Maintenance Mode**.
 - b. Click **Enter Maintenance Mode**.



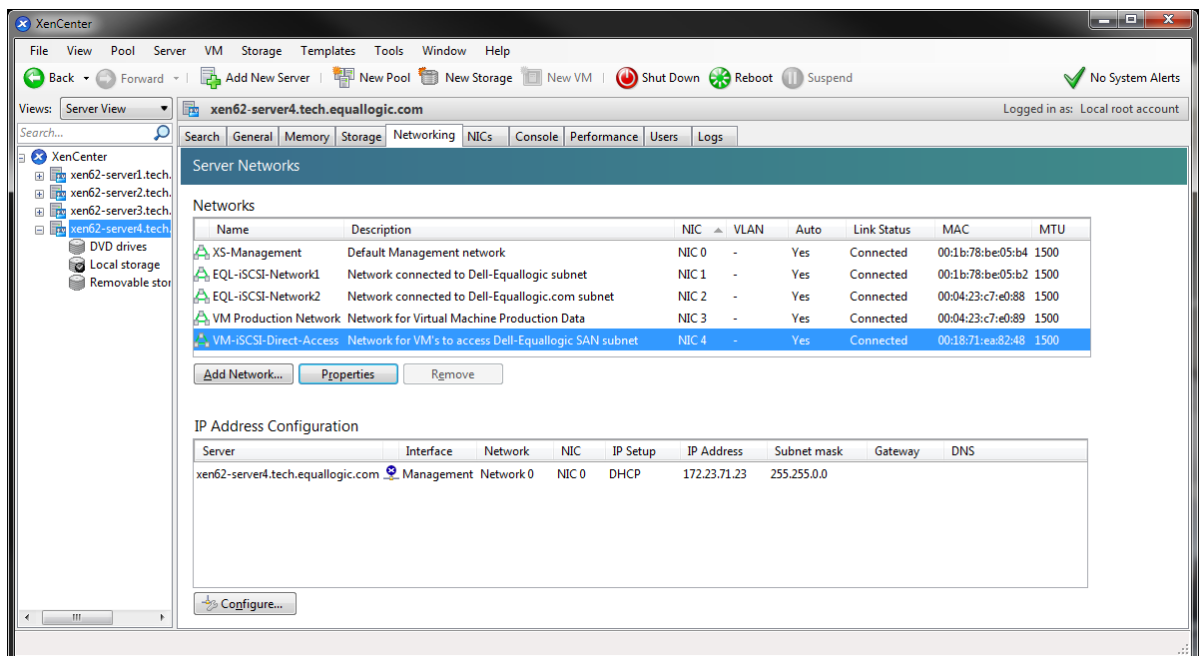
2. Enable multipathing on the XenServer.
 - a. In XenCenter, right-click on the XenServer, select **Properties > Multipathing**, and then **Enable Multipathing on this server**.



3. To easily identify the Citrix XenServer (while monitoring connections, for example), change the XenServer iSCSI IQN name to match the actual hostname. This is not a requirement. To change the XenServer, iSCSI IQN name:
 - a. In XenCenter, right-click on **XenServer**, and select **Properties > General**.
 - b. In the **iSCSI IQN** field, enter the desired IQN.



4. To easily identify the function of the XenServer network interfaces, label the networks accordingly by providing descriptions for each network interface (iSCSI, Management, etc.). This is not a requirement.
 - a. In XenCenter, highlight a XenServer and click on the **Networking** tab to view a list of available networks and network interfaces.



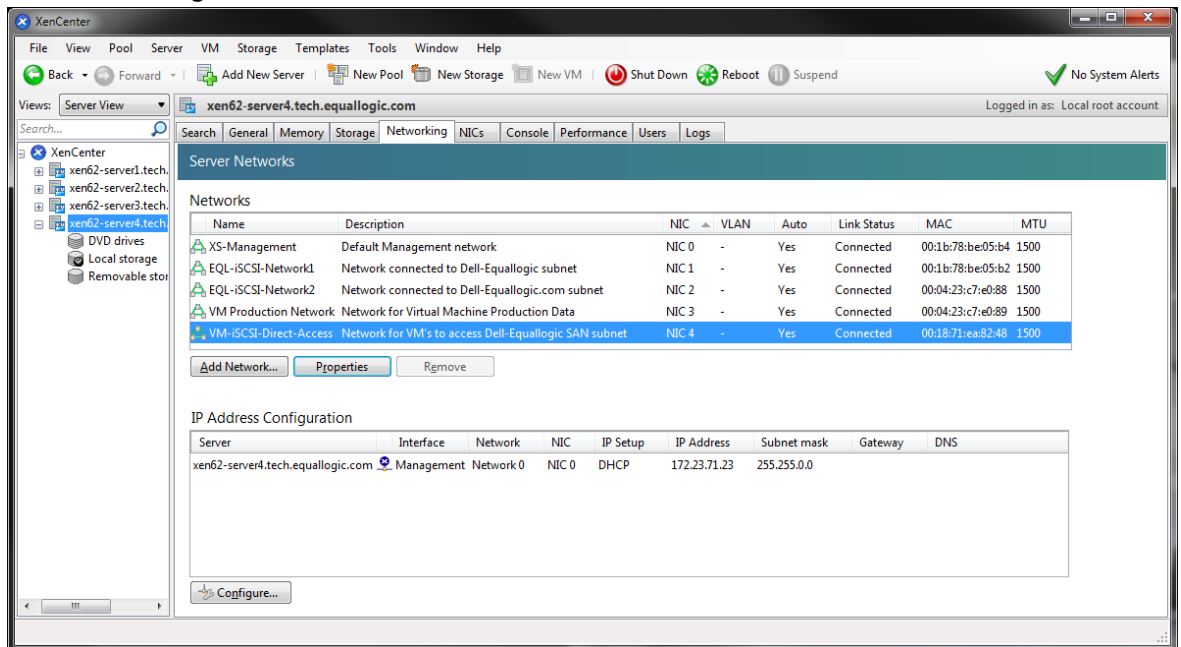
5. Enable Jumbo Frames (optional).

If the network configuration is unknown, try standard frames first. Standard frames have an MTU (Maximum Transmission Unit) size of 1500 bytes. Jumbo Frames have an MTU size of greater than 1500, typically 9000 bytes. Jumbo Frames are not required. However, in a properly configured network, Jumbo Frames can improve network efficiency and lower CPU overhead.

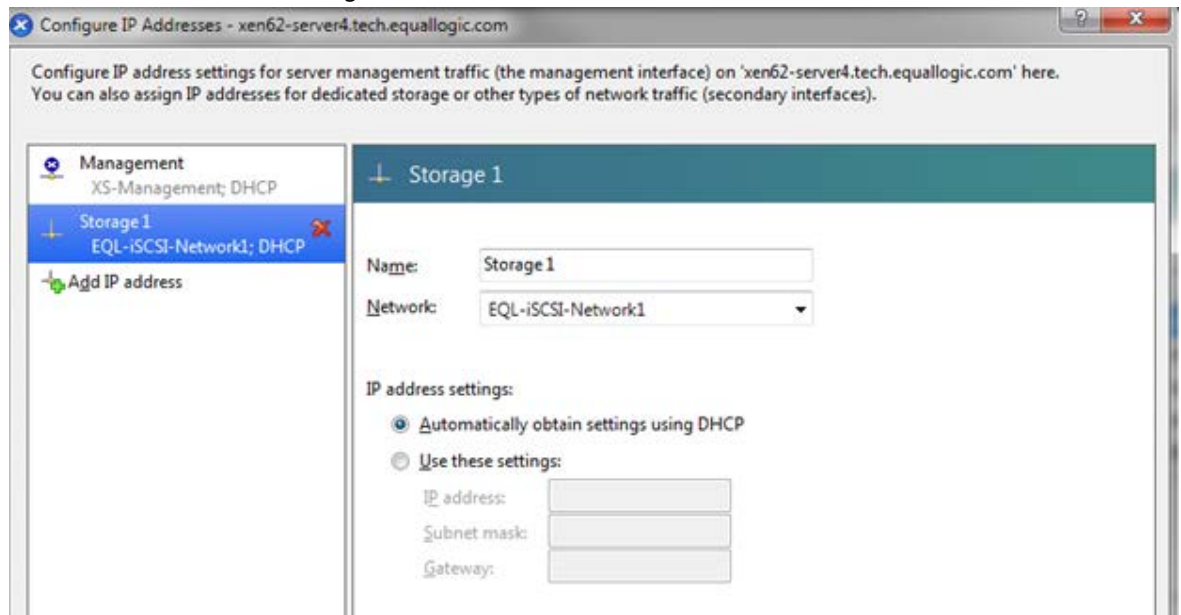
- a. To enable Jumbo Frames in XenServer, modify the properties of all iSCSI networks and set the MTU to 9000. If using Jumbo Frames, enable Jumbo Frames before configuring the network interface IP Addresses.

6. Configure iSCSI Network IP Addresses.

- a. In XenCenter, highlight a XenServer and click on the **Networking** tab
- b. Click **Configure** on the bottom of the screen.



7. Click **Add IP address** to assign an IP address to each iSCSI network interface.



The next steps require access to the console of the XenServer. Access the console from the XenCenter GUI or via SSH.

8. Create the iSCSI interface files for the storage networks.
In the XenServer console, run the following command for each network interface being used for iSCSI:

```
iscsiadm -m iface --op new -I c_iface<n>
```

The value of <n> for Network 0 would be c_iface0.

The value of <n> for Network 1 would be c_iface1.

The value of <n> for Network 2 would be c_iface2.

Example:

```
# iscsiadm -m iface -op new -I c_iface1  
New interface c_iface1 added
```

```
# iscsiadm -m iface -op new -I c_iface2  
New interface c_iface2 added
```

9. Bind the network interfaces to the XenServer Bridged Network.

In the XenServer console, run the following command for each network interface being used for iSCSI:

```
iscsiadm -m iface --op update -I <interface name> -n iface.net_ifacename -v <xen bridged network #>
```

The value of <interface name> for Network 0 would be c_iface0.

The value of <interface name> for Network 1 would be c_iface1.

The value of <interface name> for Network 2 would be c_iface2.

The value of <xen bridged network #> for c_iface0 would be xenbr0.

The value of <xen bridged network #> for c_iface1 would be xenbr1.

The value of <xen bridged network #> for c_iface2 would be xenbr2.

Examples:

```
# iscsiadm -m iface --op update -I c_iface1 -n iface.net_ifacename -v xenbr1
c_iface1 updated.
```

```
# iscsiadm -m iface --op update -I c_iface2 -n iface.net_ifacename -v xenbr2
c_iface2 updated.
```

10. Edit the /etc/multipath.conf file for EqualLogic devices.

Note: If the EqualLogic storage is the only connected platform from the XenServer, then the default multipath.conf file may be renamed. A new one can be created using the sample text below.

To rename the default multipath.conf file, use the following command:

```
#mv /etc/multipath.conf /etc/multipath.conf.default
```



XenServer comes with a simple editor called Nano. Copy this sample text, modify the devices section accordingly and save the configuration file.

```
#nano /etc/multipath.conf
devices {
    device {
        vendor "EQLOGIC"
        product "100E-00"
        path_grouping_policy multibus
        getuid_callout "/sbin/scsi_id -g -u -s /block/%n"
        path_checker tur
        failback immediate
        path_selector "round-robin 0"
        rr_min_io 3
        rr_weight priorities
    }
}
```

The default file will look like this:

```
devices {
#     device {
#         vendor            "APPLE*"
#         product           "Xserve RAID"
#         getuid_callout    "/sbin/scsi_id -g -u -s /block/%n"
#         features          "0"
#         hardware_handler  "0"
#         path_selector     "round-robin 0"
#         path_grouping_policy multibus
#         rr_weight         uniform
#         rr_min_io         1000
#         path_checker      directio
#         prio              const
#     }
}
```

The modified version should look like this:

```
devices {
device {

        vendor "EQLOGIC"
        product "100E-00"
        path_grouping_policy multibus
        getuid_callout "/sbin/scsi_id -g -u -s /block/%n"
        path_checker tur
        failback immediate
        path_selector "round-robin 0"
        rr_min_io 3
        rr_weight priorities
    }
}
```



Note: There is a known issue with Linux multipathing daemon (multipathd) versions built from source prior to May 2011 when using `path_checker readsector0` in the devices section of the `multipath.conf` file. Using this value will cause iSCSI protocol errors when used with Dell EqualLogic PS Series Firmware v7.x. This issue does not exist with versions of Dell EqualLogic PS Series Firmware prior to Version 7.0. Upgrading to Dell EqualLogic PS Series Firmware Version 7.x with `path_checker readsector0` set will result in total connectivity loss.

Using the `tur` (test unit ready) or `directio` algorithm for validating path status is required. See `path_checker tur` in the example above.

11. When done, press **[CTRL] + [X]** to exit, press **[Y]** to answer yes and save the configuration file, and then confirm that the filename is `multipath.conf`.

```
GNU nano 1.3.12 File: /etc/multipath.conf
devices {
    device {
        vendor "EQLOGIC"
        product "100E-00"
        path_grouping_policy multibus
        getuid_callout "/sbin/scsi_id -g -u -s /block/%n"
        path_checker tur
        failback immediate
        path_selector "round-robin 0"
        rr_min_io 3
        rr_weight priorities
    }
}
```

⌘G Get Help ⌘O WriteOut ⌘R Read File ⌘Y Prev Page ⌘K Cut Text ⌘C Cur Pos
⌘X Exit ⌘J Justify ⌘W Where Is ⌘N Next Page ⌘U UnCut Text ⌘T To Spell

12. To apply the changes, it is required to restart the server.



4.1 Check MPIO status

Once a Storage Repository (SR) has been created on the EqualLogic array, verify that multipathing is working correctly. See section 3 titled, "Using XenServer software iSCSI adapter" for instructions on creating an EqualLogic volume and Storage Repository and connecting to it with the software iSCSI adapter.

Verify multipathing is working correctly at the CLI by running the following commands.

```
# multipath -ll
```

Sample output:

```
36090a098703e40f38e1aa5095c01a0ff dm-1 EQLOGIC,100E-00
size=15G features='0' hwhandler='0' wp=rw
|+- policy='round-robin 0' prio=1 status=active
|  `-- 3:0:0:0 sda          8:0   active ready running
`+- policy='round-robin 0' prio=1 status=enabled
   `-- 4:0:0:0 sdb          8:16  active ready running
```

The example above shows that while MPIO is enabled, it is not correctly configured. Two separate policies are each managing a separate device versus one Round Robin policy managing two devices. The sda and sdb devices should be directly underneath a single policy as opposed to each under their own policy. The next example illustrates a correctly configured MPIO policy, where one policy has two iSCSI devices. Both paths are displayed as active and ready.

```
# multipath -ll
36090a098703e9072541a75035c012023 dm-0 EQLOGIC,100E-00
size=15G features='0' hwhandler='0' wp=rw
`+- policy='round-robin 0' prio=1 status=active
   |-- 5:0:0:0 sdb          8:16  active ready running
   `-- 6:0:0:0 sda          8:0   active ready running
```

Note: If more than two interfaces for iSCSI traffic are enabled, there will be additional iSCSI devices (sda, sdb, sdc, etc...) for each iSCSI network defined.



A Create a volume on a Dell EqualLogic PS Series Group

A.1 Create a volume

To create a volume on a Dell EqualLogic PS Series Group using the Dell EqualLogic Group Manager CLI, run the following command:

```
GrpName>volume create <volume name> <size>
```

The following example shows the creation of a 300 GB thin-provisioned volume.

```
GrpName> volume create XS-Server3-Vol0 300G thin-provision
```

Volume creation succeeded.

```
iSCSI target name is iqn.2001-05.com.equallogic:0-8a0906-51f03e709-  
d420015c6f051cc6-xs-server3-vol0
```

A.2 Create access control records on the volume

To enable the volume to be accessed by the XenServer, create an access control record on the volume by running the following command:

```
GrpName>volume select <vol name> access create initiator <initiator name>
```

The following example creates an access control record on the volume using the XenServer's iSCSI IQN.

```
GrpName> volume select XS-Server3-Vol0 access create initiator iqn.2013-  
05.com.equallogic.tech.xen62-server3
```

Access control record created with ID 1.

B XenServer iSCSI performance tuning

B.1 Tuning iSCSI performance.

To optimize XenServer iSCSI performance for use with EqualLogic storage, modify the `/etc/iscsi/iscsid-mpath.conf` file.

It is important to check with all storage vendors if multiple vendor storage devices are connected to the servers.

Modify the following parameters:

```
node.session.cmds_max = 1024 (Default 128)
```

```
node.session.queue_depth = 128 (Default 32)
```

```
iscsi.MaxRecvDataSegmentLength = 262144 (Default 128K)
```

Reboot the XenServer for the changes to take effect.



C Technical support and customer service

Dell support service is available to answer questions about PS Series SAN arrays.

Contacting Dell:

1. If an Express Service Code is available, have it ready.
This code helps the Dell automated support telephone system direct your call more efficiently.
2. If you are a customer in the United States or Canada in need of technical support, call 1-800-945-3355. If not, go to Step 3.
3. Visit support.dell.com/equallogic.
4. Log in, or click **Create Account** to request a new support account.
5. At the top right, click **Contact Us**, and call the phone number or select the link for the type of support needed.

