3PAR VMware ESX 3.X - 4.X Implementation Guide

3PAR, Inc.
4209 Technology Drive
Fremont, CA 94538 USA

Part No. 320-200195 Rev C
September 2010

Revision Notice
This is the third release of this manual. A complete revision history is provided at the end of this document.

Changes
The material in this document is for information only and is subject to change without notice. While reasonable efforts have been made in the preparation of this document to assure its accuracy, 3PAR Inc. assumes no liability resulting from errors or omissions in this document or from the use of the information contained herein.
3PAR reserves the right to make changes in the product design without reservation and without notification to its users.

Updates to the Documentation at 3PAR Central
Any updates to this document, or to other 3PAR technical documents, can be found by logging in to 3PAR Central's Document Control System from 3PAR's Support page at http://support.3PAR.com.

3PAR Technical Support and Services
Contact your local service provider for technical support and services at: http://www.3PAR.com/services.html.

Sales and Ordering Information
For sales and ordering information, contact:
3PAR Inc.
4209 Technology Drive
Fremont, CA 94538 USA
Telephone: 510-413-5999
Fax: 510-413-5699
Email: salesinfo@3PAR.com

Reader Comments and Suggestions
Please email your comments and suggestions about this document to ReaderComments@3PAR.com.

Copyrights

Trademarks

Table of Contents

1 Introduction

In this chapter

This implementation guide provides information for establishing communications between an InServ® Storage Server and a VMware ESX host. General information is also provided on the basic steps required to allocate storage on the InServ Storage Server that can then be accessed by the ESX host.

The information contained in this implementation guide is the outcome of careful testing of the 3PAR InServ Storage Server with as many representative hardware and software configurations as possible.

REQUIRED: For predictable performance and results with your 3PAR InServ Storage Server, you must use the information in this guide in concert with the documentation set provided by 3PAR for the InServ Storage Server and the documentation provided by the vendor for their respective products.

REQUIRED: All installation steps should be performed in the order described in this implementation guide.

1.1 Supported Configurations

For complete details on supported host configurations, consult the latest 3PAR InForm Configuration Matrix that is available on DCS.

1.2 InForm OS Upgrade Considerations

Refer to the InForm OS Upgrade Pre-Planning Guide (PN 320-200158) for information and planning of an online InForm OS upgrade.

1.3 Audience

This implementation guide is intended for system and storage administrators who perform and manage the system configurations and resource allocation for 3PAR InServ Storage Servers. The tasks described in this guide assume that the administrator is familiar with VMware ESX and the 3PAR Inform OS. Although this guide attempts to provide the basic information that is required to establish communications between the InServ Storage Server and the VMware ESX host, and to allocate the required storage for a given configuration, the appropriate 3PAR documentation must be consulted in conjunction with the VMware ESX host and HBA vendor documentation for specific details and procedures.

This implementation guide does NOT intend to reproduce any third-party product documentation. For details about host servers, HBAs, fabric and Ethernet switches, and non-3PAR software management tools, consult the appropriate third-party documentation.

1.4 Related Documentation

The following documents also provide information related to InServ Storage Servers and the InForm Operating System:

For information about…	Read the…
Specific platforms supported	3PAR InForm Configuration Matrix
CLI commands and their usage	*InForm OS Command Line Interface Reference*
Using the InForm Management Console to configure and administer InServ Storage Servers	*3PAR InForm OS Management Console Online Help*
InServ Storage Server concepts and terminology	*3PAR InForm OS Concepts Guide*
Determining InServ Storage Server hardware specifications, installation considerations, power requirements, networking options, and cabling	*3PAR InServ S-Class/T-Class Storage Server Physical Planning Manual* or the *3PAR InServ E-Class/F-Class Storage Server Physical Planning Manual*
Identifying storage server components and detailed alert information	*3PAR InForm OS Messages and Operator's Guide*
Using 3PAR Remote Copy	*3PAR Remote Copy User's Guide*
Using 3PAR CIM	*3PAR CIM API Programming Reference*

1.5 Organization

This guide is organized as follows:

- Chapter 1, *Introduction* (this chapter), provides an overview of this guide, including information on audience, related documentation, and typographical conventions.
- Chapter 2, *Configuring the InServ Storage Server for Fibre Channel*, explains how to establish a fibre channel connection between the InServ Storage Server and the host depending on the operating environment.
- Chapter 3, *Configuring the InServ Storage Server for iSCSI*, explains how to establish an iSCSI connection between the InServ Storage Server and the host.
- Chapter 4, *Configuring the Host for a Fibre Channel Connection*, describes the steps required to set up the host to communication over a fibre channel connection.
- Chapter 5, *Configuring the Host for FC over Ethernet (FCoE) Connection*, describes the process for establishing a Fibre Channel over Ethernet (FCoE) connection between the InServ Storage Server and the host.
- Chapter 6, *Configuring the Host for an iSCSI Connection*, describes the steps that are necessary set up the host to communicate over an iSCSI connection.
- Chapter 7, *Allocating Storage for Access by the ESX Host*, provides information that is useful when creating and managing storage for a given operating environment and connection configuration.
- Chapter 8, *Booting from the InServ Storage Server*, describes the process that is required to boot the ESX Server from the SAN.

This guide also contains a revision history for your reference.

1.6 Typographical Conventions

This guide uses the following typographical conventions:

Typeface	Meaning	Example
ABCDabcd	Used for dialog elements such as titles, button labels, and other screen elements.	When prompted, click **Finish** to complete the installation.
ABCDabcd	Used for paths, filenames, and screen output.	Open the file `/kernel/drv/lpfc.conf`
ABCDabcd	Used to differentiate user input from screen output.	`# cd /opt/3par/console`
<ABCDabcd>	Used for variables in filenames, paths, and screen output.	Modify the content string by adding the `-P<x>` option after `-jar inform.jar`
<ABCDabcd>	Used for variables in user input.	`#.\java -jar inform.jar -P<x>`

1.7 Advisories

To avoid injury to people or damage to data and equipment, be sure to observe the cautions and warnings in this guide. ***Always be careful when handling any electrical equipment.***

NOTE: Notes are reminders, tips, or suggestions that supplement the procedures included in this guide.

CAUTION: Cautions alert you to actions that can cause damage to equipment, software, or data.

REQUIRED: Requirements signify procedures that must be followed as directed in order to achieve a functional and supported implementation based on testing at 3PAR.

WARNING: Warnings alert you to actions that can cause injury to people or irreversible damage to data or the operating system.

2 Configuring the InServ Storage Server for Fibre Channel

In this chapter

This chapter explains how to establish a Fibre Channel connection between the InServ Storage Server and a VMware ESX host and covers InForm OS 2.2.X and 2.3.X versions. For information on setting up the physical connection for a particular InServ Storage Server, see the appropriate 3PAR installation manual.

2.1 Configuring the InServ Storage Server Running InForm O.S. 2.3.X

This section describes how to connect the Inserv Storage Server to an ESX Host over a Fibre Channel network when running InForm OS 2.3.X.

For information on setting up a connection for OS 2.2.X and earlier, see *2.2 Configuring the InServ Storage Server Running InForm O.S. 2.2.X* on page 2.5.

NOTE: By default, the QLogic and Emulex Drivers for the VMware ESX Server support failover. For failover support using the QLogic or Emulex driver, virtual volumes (VVs) should be simultaneously exported down multiple paths to the host server. To do this, create a host definition on the InServ Storage Server that includes the WWNs of multiple HBA ports on the host server and then export the VLUNs to that host definition.

REQUIRED: The following setup must be completed before connecting the InServ Storage Server port to a device.

2.1.1 Setting Up the Ports

Before connecting the InServ Storage Server to a host, the connection type and mode must be specified. To set up the InServ Storage Server ports for a direct or fabric connection, complete the following steps for each port.

1 To determine if a port has already been configured for a host port, issue the CLI `showport -par` command.

```
%showport -par
N:S:P Connmode ConnType CfgRate MaxRate Class2   UniqNodeWwn VCN      IntCoal
2:0:1 disk     loop     auto    4Gbps   disabled disabled    disabled enabled
2:0:2 disk     loop     auto    4Gbps   disabled disabled    disabled enabled
2:4:1 host     point    auto    4Gbps   disabled disabled    disabled enabled
2:4:2 host     point    auto    4Gbps   disabled disabled    disabled enabled
3:0:1 disk     loop     auto    4Gbps   disabled disabled    disabled enabled
3:0:2 disk     loop     auto    4Gbps   disabled disabled    disabled enabled
3:4:1 host     point    auto    4Gbps   disabled disabled    disabled enabled
3:4:2 host     point    auto    4Gbps   disabled disabled    disabled enabled
```

2 If the port has not been configured, take the port offline before configuring it for the ESX host by issuing the CLI `controlport offline [node:slot:port]`command.

CAUTION: Before taking a port offline in preparation for a direct or fabric connection, you should verify that the port has not been previously defined and that it is not already connected to a host as this would interrupt the existing host connection. If an InServ port is already configured for a direct or fabric connection, you can ignore step 2 as you do not have to take the port offline.

3 To configure the port for the host server, issue the `controlport config host -ct loop/point [node:slot:port]`command with the appropriate option for the `-ct` parameter.

For a direct connection:

Use the `-ct loop` parameter to specify a direct connection.

For a fabric connection:

Use the `-ct point` parameter to specify a fabric connection.

4 Issue the `controlport rst` command to reset and register the new port definitions.

The following example shows how to set up a fabric connected port.

```
% controlport offline 1:5:1
% controlport config host -ct point 1:5:1
% controlport rst 1:5:1
```

2.1.2 Creating the Host Definitions (Host Persona)

Before connecting the ESX host to the InServ Storage Server, you need to create a host definition that specifies a valid host persona for each InServ Storage Server that is to be connected to a host HBA port through a fabric or direct connection. ESX uses the generic legacy host personality of 6 for both the QLogic and Emulex HBAs. The following steps show how to create the host definition.

1 To create host definitions, issue the `createhost` command with the `-persona` option to specify the persona and the host name. For example:

```
% createhost -persona 6 ESXserver1 10000000C9724AB2 10000000C97244FE
```

2 To verify that the host has been created, issue the `showhost` command.

```
Id Name       Persona         -WWN/iSCSI_Name- Port
 0 ESXserver1 Generic-legacy 10000000C9724AB2 ---
                             10000000C97244FE ---
```

NOTE: See the 3PAR InForm OS Command Line Reference or the IMC help for complete details on using the `controlport`, `createhost` and `showhost` commands.

2.2 Configuring the InServ Storage Server Running InForm O.S. 2.2.X

This section describes the steps that are required to connect the Inserv Storage Server to an ESX host over a fibre channel network and to create the host definitions when running InForm OS 2.2.X.

NOTE: For configurations that are intended to have more than one host type (for example: an ESX host and a Windows host) connected to a shared InServ Storage Server HBA port via fabric connections, refer to the *Heterogeneous Host Support Guide* in DCS for the required InServ Storage Server port settings and interoperability considerations.

NOTE: By default, the QLogic and Emulex Drivers for the VMware ESX Server support failover. For failover support using the QLogic or Emulex driver, virtual volumes (VVs) should be simultaneously exported down multiple paths to the host server. To do this, create a host definition on the InServ Storage Server that includes the WWNs of multiple HBA ports on the host server and then export the VLUNs to that host definition.

REQUIRED: The following setup must be completed before connecting the InServ Storage Server port to a device.

2.2.1 Setting Up the Ports

Before connecting the InServ Storage Server to a host, the InServ Storage Server port personality must be specified. To set up the InServ Storage Server ports for a direct or fabric connection, issue the appropriate set of InForm CLI `controlport` commands for each port.

For Direct Connections use Personality 1 with VCN disabled

```
# controlport persona 1 <node:slot:port>
# controlport vcn disable -f <node:slot:port>
```

Verify port personality 1, connection type loop, using the CLI `showport -par` command.

```
N:S:P ConnType CfgRate MaxRate  Class2      VCN -----------Persona------------ IntCoal
(...)
1:4:2     loop    auto   2Gbps disable disabled *(1) g_ven, g_hba, g_os, 0, DC enabled
(...)
```

For Fabric Connection use Personality 7 with VCN disabled

```
# controlport persona 7 <node:slot:port>
# controlport vcn disable -f <node:slot:port>
```

Verify port personality 7, connection type point, using the CLI `showport -par` command.

```
N:S:P ConnType CfgRate MaxRate  Class2      VCN -----------Persona------------ IntCoal
(...)
1:4:2    point    auto   2Gbps disable disabled *(7) g_ven, g_hba, g_os, 0, FA enabled
(...)
```

2.2.2 Creating the Host Definitions

Before connecting the ESX host to the InServ Storage Server, you need to create a host definition for each InServ Storage Server that is to be connected to a host HBA port through a fabric or direct connection.

1. To create host definitions, issue the InForm CLI `createhost` command with the appropriate host name. For example:

```
# createhost ESXserver1 10000000C9724AB2 10000000C97244FE
```

2. To verify that the host has been created, issue the InForm CLI `showhost` command.

```
Id Name        -WWN/iSCSI_Name- Port
 0 ESXserver1 10000000C9724AB2 ---
              10000000C97244FE ---
```

NOTE: See the 3PAR InForm OS Command Line Reference or the IMC help for complete details on using the `controlport`, `showport`, `createhost` and `showhost` commands.

2.3 Setting Up the Fabric

This section provides information for establishing a fabric connection.

NOTE: For Direct Connect configuration, skip this Fabric Setup section.

- Brocade switch ports used in this configuration should be set to the default mode.
- On Brocade 3xxx switches running Brocade firmware 3.2.0 or later, use the Brocade telnet interface and the `portcfgshow` command as follows:

```
brocade2_1:admin> portcfgshow
Ports              0  1  2  3    4  5  6  7
-----------------+--+--+--+--+----+--+--+--
Speed             AN AN AN AN   AN AN AN AN
Trunk Port        ON ON ON ON   ON ON ON ON
Locked L_Port     .. .. .. ..   .. .. .. ..
Locked G_Port     .. .. .. ..   .. .. .. ..
Disabled E_Port   .. .. .. ..   .. .. .. ..
                  where AN:AutoNegotiate, ..:OFF, ??:INVALID.
```

- McData switch or director ports are set to the default modes as G-ports, with a speed setting of **Negotiate**. Point to point settings required on **McData Fabric Data Rate** should be set to **Auto-negotiate**.
- Set Cisco switch ports that connect to InServ Storage Server ports or host HBA ports to **AdminMode = FX** and **AdminSpeed = auto** with the speed setting **auto negotiate**.
- QLogic switch ports should be set to port type **GL-port** and port speed **auto-detect**. QLogic switch ports that connect to InServ should be set to I/O Stream Guard **disable** or **auto** but never **enable**.

2.3.1 Fabric Zoning

REQUIRED: When establishing zoning with the InServ Storage Server, there must only be a single initiator zoned with a single target. If a customer experiences an issue using another zoning approach, 3PAR may require the customer to implement this zoning approach as part of troubleshooting and/or corrective action.

Fabric zoning controls which fabric connected-devices have access to each other on the fabric. The required use of single initiator to single target zoning isolates the host server and InServ Storage Server ports from Registered State Change Notifications (RSCN) that are irrelevant for these ports.

Zoning is achieved by associating the devices World Wide Names (WWNs) or ports to be connected through the fabric. While both the WWN and the port zoning methods can be used with the InServ Storage Server, the WWN zoning method is recommended since the zone survives the changes of ports resulting from cable re-connection for a given device on a fabric. In the following explanations an initiator port (initiator for short) refers to a host server HBA port and a target port (target for short) refers to an InServ Storage Server HBA port.

Fabric zoning should be employed, using the methods provided by the switch vendor, to create relationships between host server HBA ports and storage server ports before connecting the host server HBA ports or InServ Storage Server ports to the fabric(s).

After connecting each host server HBA port and InServ Storage Server port to the fabric(s), verify the switch and zone configurations using the InForm CLI `showhost` command, to ensure that each initiator is zoned with the correct target.

NOTE: In the examples in the following sections, a fabric can be one or more Fibre Channel switches or directors.

2.3.2 Single Initiator to Single Target Zoning No Fan-In No Fan-Out

In a single initiator to single target zoning, no fan-in, no fan-out configuration, each HBA port is connected to only one InServ Storage Server port (Figure 2-1).

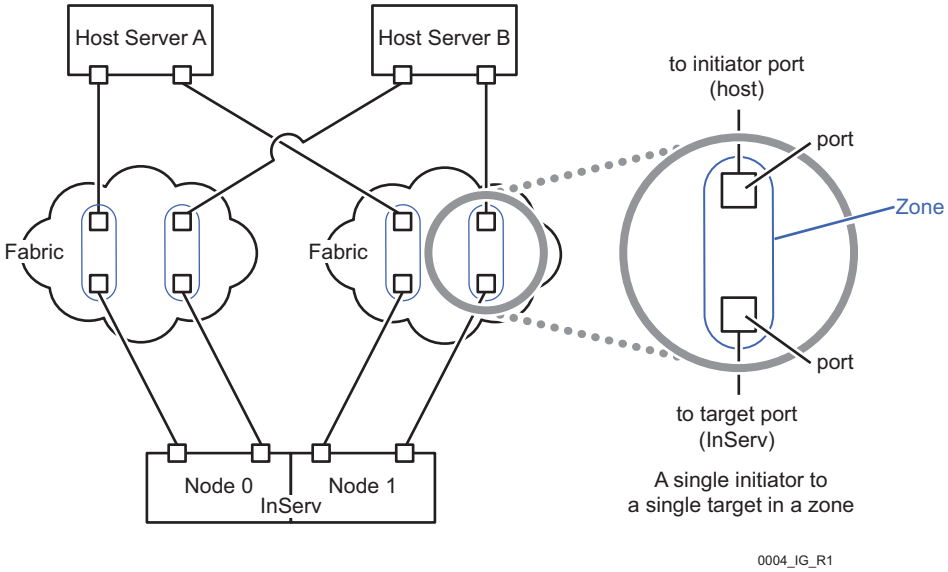

Figure 2-1. Single Initiator to Single Target Zoning No Fan-In/No Fan-Out

2.3.3 Single Initiator to Single Target Zoning with Fan-Out from One InServ Port to Multiple Host Ports

Fan-out refers to an InServ Storage Server port that is connected to more that one host port, as shown in Figure 2-2.

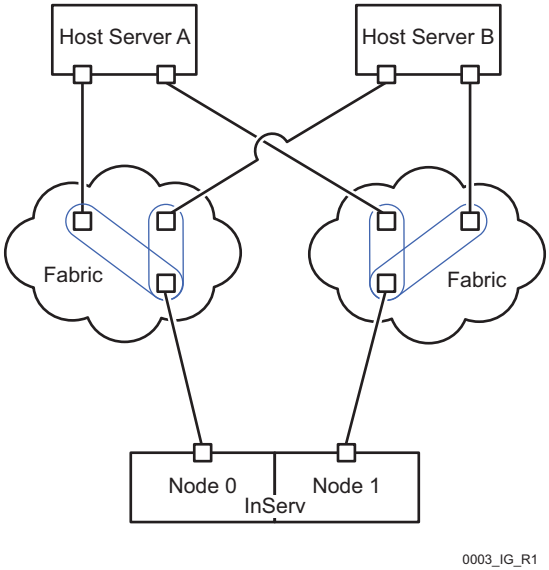

Figure 2-2. Single Initiator to Single Target Zoning with Fan-Out

NOTE: Observe the following limits in a VMware ESX / 3PAR Storage Server configuration:

- A maximum of 16 host server ports can fan-out from a single 2Gbps InServ Storage Server port.
- A maximum of 32 host server ports can fan-out from a single 4Gbps InServ Storage Server port.

2.3.4 Single Initiator to Single Target Zoning with Fan-In from Multiple InServ Ports to One Host Port

Fan-in refers to a host server port connected to many InServ Storage Server ports. This is shown in Figure 2-3.

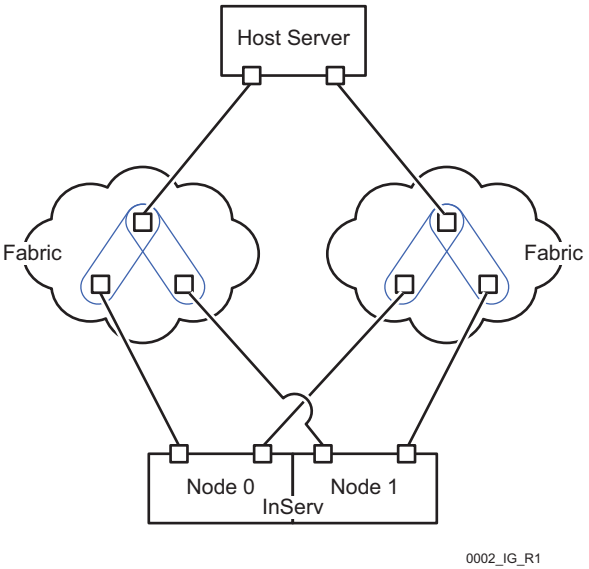

Figure 2-3. Single Initiator to Single Host Target Zoning with Fan-In

NOTE: A maximum of four InServ Storage Server ports can fan-in to a single host server port.

2.3.5 Single Initiator to Single Target Zoning with Mixed Fan-In/Out Configurations

The following figure (Figure 2-4) shows a single initiator to a single target zoning with fan-in and fan-out from one InServ Storage Server to multiple host servers.

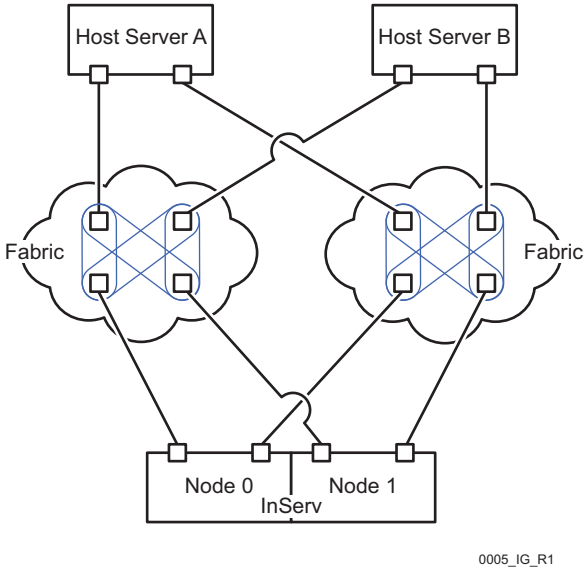

Figure 2-4. Single Initiator to Single Target Zoning with Fan-In and Fan-Out

2.3.6 Non-Compliant Zoning Examples

In the following examples, the zoning rule of one initiator zoned to one target is not respected.

Non-compliant zoning is shown in Figure 2-5.

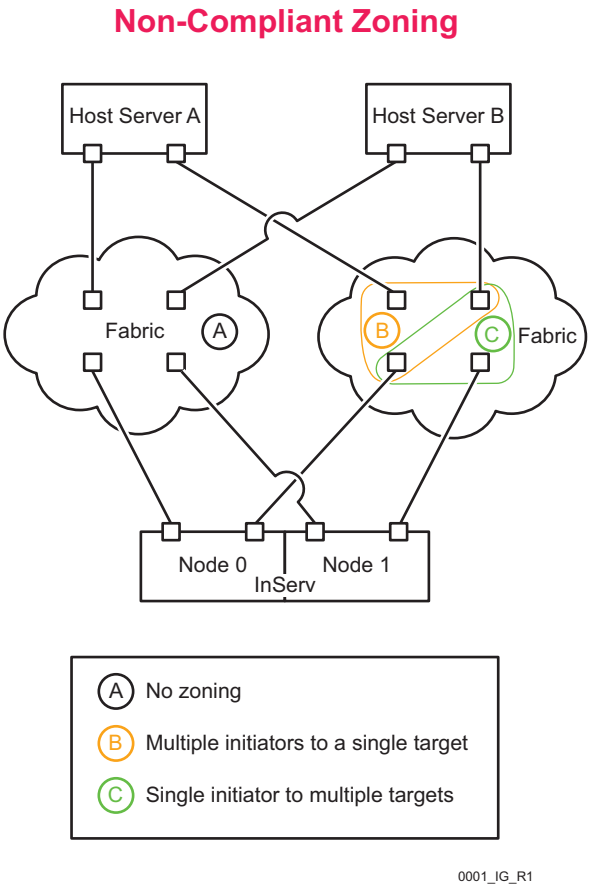

Figure 2-5. Non-Compliant Zoning

2.3.7 General Note on Fabrics

The following section describes general limitations for target ports.

2.3.7.1 Target Port Limitations

In order to avoid the overwhelming of a target port and to ensure continuous I/O operations, observe the following limitations on a target port:

- Maximum host server initiator ports per InServ Storage Server target port in an ESX server configuration:
  - 16 host server ports per 2Gbps InServ Storage Server port with VMware Adaptive Queue Depth or Target Throttling enabled on ESX servers (refer to *4.4 Performance Considerations for Multiple Host Configurations* on page 4.11 for details).
  - 32 host server ports per 4Gbps InServ Storage Server port with VMware Adaptive Queue Depth or Target Throttling enabled on ESX servers (refer to *4.4 Performance Considerations for Multiple Host Configurations* on page 4.11 for details).
- Maximum total of 1,024 host server initiator ports per InServ Storage Server.
- I/O queue depth on each InServ Storage Server HBA model as follows:
  - QLogic 2G: 497
  - LSI 2G: 510
  - Emulex 4G: 959
  - 3PAR HBA 4G: 1638
- The I/O queues are shared among the connected host server HBA ports on a first-come, first-served basis.
- When all queues are in use and a host HBA port tries to initiate I/O, it will receive a target queue full response from the InServ Storage Server port. This can result in erratic I/O performance on each host server. If this condition occurs, each host server should be throttled so that they cannot overrun the InServ Storage Server port's queues when all host servers are delivering their maximum number of I/O requests. Please consult *4.4 Performance Considerations for Multiple Host Configurations* on page 4.11 for details.

3 Configuring the InServ Storage Server for iSCSI

In this chapter

This chapter explains how to establish an iSCSI connection between the InServ Storage Server and the VMware ESX host. For information on setting up the physical connection, see the appropriate 3PAR installation manual.

3.1 Setting Up the Ports for an iSCSI Connection

To establish an iSCSI connection between the InServ Storage Server and the ESX host, you need to set up each InServ Storage Server iSCSI target port that will be connected to an iSCSI Initiator as described in the following steps.

1 Issue the InForm CLI `showport` command to check the current settings of the iSCSI ports:

```
# showport -iscsi
N:S:P   State  IPAddr Netmask Gateway TPGT  MTU Rate DHCP iSNS_Prim iSNS_Sec iSNS_Port
0:1:1 offline 0.0.0.0 0.0.0.0 0.0.0.0   11 1500  n/a    0   0.0.0.0  0.0.0.0      3205
0:1:2 offline 0.0.0.0 0.0.0.0 0.0.0.0   12 1500  n/a    0   0.0.0.0  0.0.0.0      3205
1:1:1 offline 0.0.0.0 0.0.0.0 0.0.0.0  111 1500  n/a    0   0.0.0.0  0.0.0.0      3205
1:1:2 offline 0.0.0.0 0.0.0.0 0.0.0.0  112 1500  n/a    0   0.0.0.0  0.0.0.0      3205
```

2 Issue the InForm CLI `controliscsiport` command to set up the IP address and netmask address of the iSCSI target ports.

```
# controliscsiport addr 10.1.1.100 255.255.255.0 -f 0:1:1
# controliscsiport addr 10.1.1.102 255.255.255.0 -f 1:1:1
# controliscsiport addr 10.1.1.101 255.255.255.0 -f 0:1:2
# controliscsiport addr 10.1.1.103 255.255.255.0 -f 1:1:2
```

NOTE: Make sure the IP switch ports, where the InServ iSCSI target port(s) and iSCSI Initiator host to which it is connected, are able to communicate with each other, by using the ping command on the host. (The VMware ESX Server iSCSI Initiator must be configured to perform this operation per Chapter 6, *Configuring the Host for an iSCSI Connection*.

NOTE: There is a limit of 64 host server iSCSI initiator ports connected to any one InServ Storage Server Target port.

3.2 Creating the iSCSI Host Definition on an InServ Running InForm OS 2.3.x

Create a host definition that ties all of the connections from a single host to a host name. Prior to creating a host definition using the following steps, the InServ iSCSI target ports must have been set up and an iSCSI connection/session must be established. The iSCSI connection/session is established by following the steps in *3.1 Setting Up the Ports for an iSCSI Connection* on page 3.2 and the steps in Chapter 6, *Configuring the Host for an iSCSI Connection* through *6.5 Configuring the VMware iSCSI Initiator* on page 6.12 (ESX host set-up).

The following example of host definition creation depicts a VMware iSCSI Initiator "iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56" on an ESX server (the only iSCSI initiator for this server in this case) connecting through a vlan to a pair of InServ iSCSI ports. The host definition is given the name "ESX1" and the host persona is set to 6 (Generic-legacy).

1 Issue the InForm CLI `showhost` command to verify that the host iSCSI Initiators are connected to the InServ iSCSI target ports.

```
# showhost
Id Name     Persona         ----------------WWN/iSCSI_Name---------------- Port
   --       --              iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56  0:1:2
                            iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56  1:1:2
```

2 Issue the InForm CLI `createhost` command to create the appropriate host definition entry.

```
# createhost -iscsi -persona 6 ESX1 iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56
```

3 Issue the InForm CLI `showhost` command to verify that the host entry has been created.

```
# showhost
Id Name     Persona         ----------------WWN/iSCSI_Name---------------- Port
 0 ESX1     Generic-legacy  iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56  0:1:2
                            iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56  1:1:2
```

4 To test the connection, create some temporary virtual volumes and export the VLUNS to the host.

NOTE: Consult Chapter 7, *Allocating Storage for Access by the ESX Host* for complete details on creating, exporting and discovering storage for an iSCSI connection.

5 On the ESX iSCSI Initiator host, perform a rescan and verify that the disks have been discovered.

3.3 Creating the iSCSI Host Definition on an InServ Running Inform OS 2.2.x

Create a host definition that ties all of the connections from a single host to a host name. Prior to creating a host definition using the following steps, the InServ iSCSI target ports must have been set up and an iSCSI connection/session must be established. The iSCSI connection/session is established by following the steps in *3.1 Setting Up the Ports for an iSCSI Connection* on page 3.2 and the steps in Chapter 6, *Configuring the Host for an iSCSI Connection* through section *6.5 Configuring the VMware iSCSI Initiator* on page 6.12 (ESX host set-up).

The following example of host definition creation depicts a VMware iSCSI Initiator "iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56" on an ESX server (the only iSCSI initiator for this server in this case) connecting through a vlan to a pair of InServ iSCSI ports. The host definition is given the name "ESX1".

1 Issue the InForm CLI `showhost` command to verify that the host iSCSI Initiators are connected to the InServ iSCSI target ports.

```
# showhost
Id Name     ----------------WWN/iSCSI_Name---------------- Port
   --       iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56  0:1:2
            iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56  1:1:2
```

2 Issue the InForm CLI `createhost` command to create the appropriate host entry.

```
# createhost -iscsi ESX1 iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56
```

3 Issue the InForm CLI `showhost` command to verify that the host entry has been created.

```
# showhost
Id Name     ----------------WWN/iSCSI_Name---------------- Port
 0 ESX1     iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56  0:1:2
            iqn.1998-01.com.vmware:sqahpbc02icm5-40e25c56  1:1:2
```

4 To test the connection, create some temporary virtual volumes and export the VLUNS to the host.

NOTE: Consult Chapter 7, *Allocating Storage for Access by the ESX Host* for complete details on creating, exporting and discovering storage for an iSCSI connection.

5 On the ESX iSCSI Initiator host, perform a rescan and verify that the disks have been discovered.

3.4 Setting Up and Configuring CHAP Authentication

Enabling Host CHAP is an option that can be set up at the ESX system administrator's discretion. The following example outlines the procedures for host (initiator) CHAP which is available in ESX 3.x. As of ESX 4.0, mutual (bidirectional, initiator-target) CHAP is also available.

1 Issue the InForm CLI `showhost` command to verify that a host definition has been created on InServ Storage Server for the ESX Server that will have CHAP enabled.

```
# showhost
Id Name         ----------------WWN/iSCSI_Name---------------- Port
 0 ESX1         iqn.1998-01.com.vmware:hpdl380-01-11a38a59     0:1:2
                iqn.1998-01.com.vmware:hpdl380-01-11a38a59     1:1:2
```

The following example uses the CHAP secret (CHAP password) `host_secret3` for the ESX host. Be aware that the CHAP secret must be at least 12 characters long.

2 On the ESX server's VI / vSphere client, open the **Storage Adapter** Tab, then select the **iSCSI Software Adapter**, and select the **Properties** link. Select the **CHAP Authentication** tab.

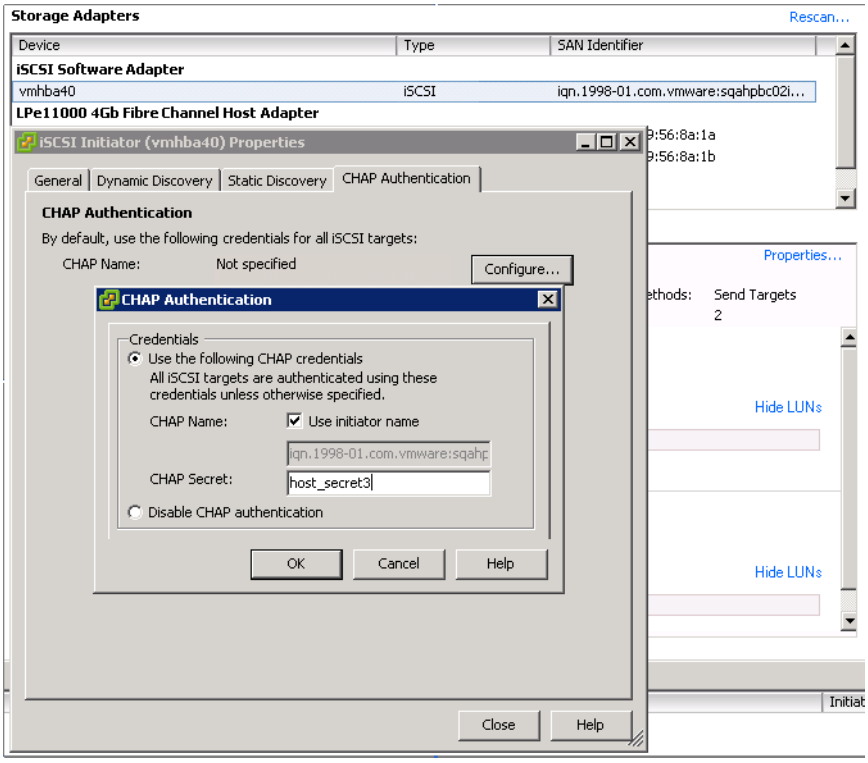

3 Select the **Use the following CHAP credentials.** radio button.

4 Check the Box for **Use Initiator Name**.

5 Enter the CHAP Secret (must be at least 12 characters long).

6 Click **OK** when you are done. A warning screen appears indicating that a reboot of the ESX server is required.

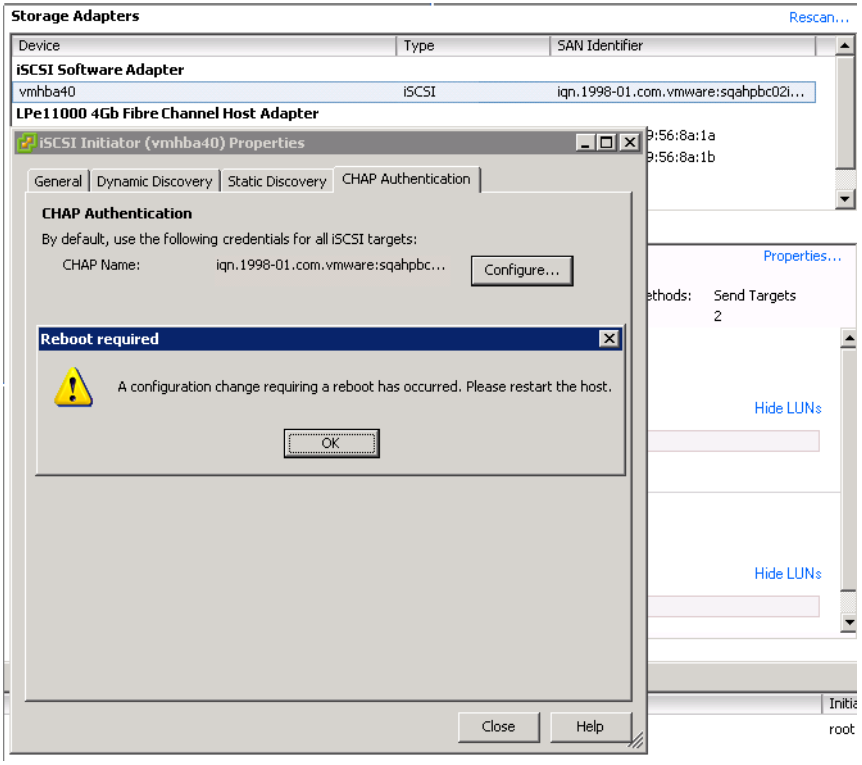

7 Click **OK** again to confirm.

8 On the 3PAR InServ Storage Server, issue the InForm CLI `sethost` command with the `initchap` parameter to set the CHAP secret for the ESX host.

```
# sethost initchap -f host_secret3 ESX1
```

NOTE: If mutual CHAP on ESX is being configured, then target CHAP will need to be configured on the InServ Storage Server as well as initiator CHAP. Set target CHAP secret using the InForm CLI `sethost` command with the `targetchap` parameter.

```
# sethost targetchap -f host_secret3 ESX1
```

Issue the InForm CLI `showhost` command with the `-chap` parameter to verify that the specified CHAP secret has been set for the host definition.

```
# showhost -chap
Id Name -Initiator_CHAP_Name- -Target_CHAP_Name-
 0 ESX1 ESX1 --
```

9 Reboot the ESX host.

4 Configuring the Host for a Fibre Channel Connection

In this chapter

This chapter describes the procedures and considerations that are required to set up an ESX host to communication with an InServ Storage Server over a fibre channel connection.

4.1 Installing the HBA and Drivers

Before setting up the ESX host, make sure the host adapters are installed and operating properly. If necessary, consult the documentation provided by the HBA vendor for instructions.

Drivers for VMware supported HBAs are included as part of the ESX OS installation package supplied by VMware. Updates and/or patches for the HBA drivers can be acquired through VMware support.

4.2 Installing Virtual Machine Guest Operating System (GOS)

The VMware ESX Server documentation lists recommended Virtual Machine guest operating systems (GOS) and their installation and setup as Virtual Machines. Refer to the VMware ESX Server documentation for information on setting up your Virtual Machine configuration.

NOTE: VMware and 3PAR recommend the LSI logic adapter emulation for Windows 2003 Servers. The LSI Logic adapter is also the default option for Windows 2003 when creating a new Virtual Machine. 3PAR testing has noted a high incidence of Windows 2003 Virtual Machine failures during an ESX multipath failover/failback event when the BUS Logic adapter is used with Windows 2003 VMs.

NOTE: 3PAR testing indicates that the SCSI timeout value for Virtual Machine guest operating systems should be 60 seconds in order to successfully ride out path failovers at the ESX layer. Most guest operating systems supported by VMware have a default SCSI timeout value of 60 seconds, but this value should be checked and verified for each GOS installation. In particular, Red Hat 4.x guest operating systems should have their SCSI timeout value changed from their default value of 30 seconds to 60 seconds.

This command line can be used to set the SCSI timeout on all SCSI devices presented to a Red Hat 4.x Virtual Machine to 60 seconds:

```
find /sys -name timeout | grep "host.*target.*timeout" | xargs -n 1 echo "echo 60 >"|sh
```

This must be added as a line in /etc/rc.local of the Red Hat 4.x guest OS in order for the timeout change to be maintained with a Virtual Machine reboot.

Example of modified /etc/rc.local file:

```
# cat /etc/rc.local
#!/bin/sh
#
# This script will be executed *after* all the other init scripts.
# You can put your own initialization stuff in here if you don't
# want to do the full Sys V style init stuff.

find /sys -name timeout | grep "host.*target.*timeout" | xargs -n 1 echo
"echo 60 >"|sh
touch /var/lock/subsys/local
```

CAUTION: VMware KB Article 51306: VMware has identified a problem with RHEL5 (GA), RHEL4 U4, RHEL4 U3, SLES10 (GA), and SLES9 SP3 guest operating systems. Their file systems may become read-only in the event of busy I/O retry or path failover of the ESX Server's SAN or iSCSI storage.

Because of this known issue, 3PAR does not recommend, and does not support the usage of RHEL5 (GA), RHEL4 U4, RHEL4 U3, SLES10 (GA), and SLES9 SP3 as guest operating systems for Virtual Machines on VMware ESX Servers attached to 3PAR InServ Storage Servers.

CAUTION: The use of the N-Port ID Virtualization (NPIV) feature introduced with VMware ESX 3.5 - 4.0, allowing Virtual Ports/WWNs to be assigned to individual Virtual Machines, is not recommended and not supported with 3PAR InForm OS.

4.3 Fibre Channel Multipath Failover Considerations and I/O Load Balancing

The VMware ESX Server 3.0 - 3.5 includes failover with multi-path support to maintain a constant connection between the ESX host and the InServ Storage Server array. VMware terms this multipath support as "active/active" with a choice of two "path policies" called "FIXED" or "MRU". As of ESX 4.0, a third path policy choice of "round robin" is available. The path policies can be modified on a per InServ volume (LUN) basis by right clicking on the device listing and selecting the "properties" function from the VI / vSphere Client menu. A pop-up window allows you to 'manage paths' whereby the properties of the paths to the volume that was previously selected can be modified. Using this control, you can select the path policy, and specify which path is the active preferred path to a volume on the storage array, or which path is the standby path within the FIXED path policy scheme. Additionally, paths can be disabled to prevent any traffic over a specific path to a volume on the storage array.

The VI / vSphere client GUI interface only allows for settings to be changed on a volume by volume (LUN by LUN) basis. The GUI interface is appropriate and preferred for use in managing I/O paths within the FIXED path policy scheme. Refer to section 4.3.1 for procedures on implementing and configuring the round-robin path policy on ESX 4.0 with 3PAR InServ.

- "Active/active, with a path policy of "round-robin" is the preferred multipath implementation for ESX 4.0. Refer to section 4.3.1 for procedures on implementing and configuring the round-robin path policy on ESX 4.0 with 3PAR InServ.
- "Active/active, with a path policy of "fixed" and the preferred/active paths manually set to balance I/O load evenly across all paths is the preferred multipath implementation for ESX 3.0 - 3.5.
  - In the event the active path is detected as having failed or has been disabled either at the fabric switch, or on the storage array, all ESX Server I/O to the storage array continues by failing over to a 'standby' path. When the ESX server detects that the preferred path has been recovered or is enabled, I/O from the ESX server then resumes on the preferred path -- assuming a preferred path policy had previously been set to that path.

- I/O from the ESX server should be manually distributed or balanced when 2 or more paths exist to more than one InServ volume on the storage array. Manually balancing the loads across available paths may improve I/O performance. This path load balancing to the storage array is dependant on the number of I/O's that are targeted for specific volumes on the storage array. Tuning I/O's to specific volumes on specific paths to the storage array varies from configuration to configuration and is totally dependant on the workload from the ESX Server and the virtual machines to the devices on the storage array.

The following vSphere client screen shot depicts a LUN with five I/O paths in a FIXED I/O policy scheme. The path marked Active (I/O) with the '*' in the Preferred column is the path chosen as preferred, and is the path to which all I/O is currently assigned for the given LUN. The other paths listed are active, but in 'standby' mode. The paths in active 'standby' will not be used for I/O traffic for this LUN unless the preferred path fails.

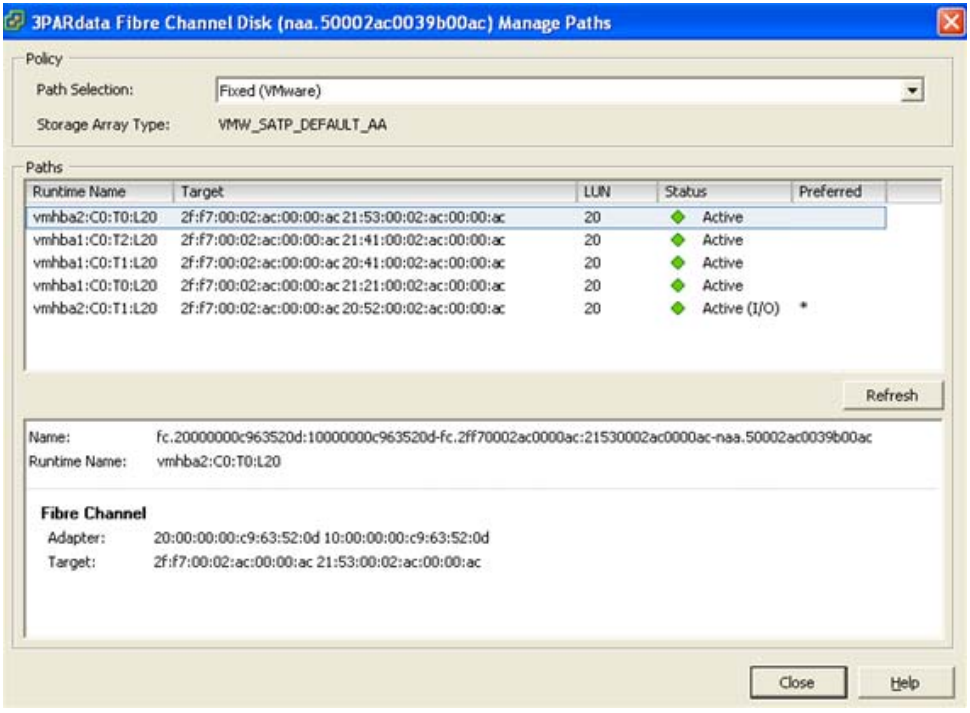

- Active/active, with a path policy of MRU (Most Recent Used) does not maintain or re-instate balancing of I/O load after a failover/failback multipath event. This could leave I/O in an unplanned for and unbalanced state which may yield significant I/O performance issues. Implementation of a MRU (Most Recent Used) path policy is not recommended by 3PAR.

NOTE: If I/O is active to a LUN and an attempt is made to modify the path policy, a failure can occur indicating:

```
"error during the configuration of the host:
sysinfoException; Status=Busy: Message=Unable to Set".
```

If this problem occurs while attempting to change the path policy, reduce the I/O's to that LUN and then try making the desired changes.

For additional information on this topic, please refer to the chapter on "Multipathing" contained in the VMware SAN Configuration Guide.

4.3.1 Configuring Round Robin Multipathing on ESX 4.0 for Fibre Channel

With Version 4.0 onwards, VMware supports a round robin I/O path policy for Active/Active storage arrays such as 3PAR InServ Storage Sever. An Active/Active round robin I/O path policy is the preferred configuration for ESX 4.x, however, the policy is not enabled by default. The default I/O path policy for ESX 4.0 is VMware's FIXED path policy scheme.

Managing a round robin I/O path policy scheme though the VI / vSphere client GUI for a large network can be cumbersome and a challenge to maintain because the policy must be specified for each LUN individually and updated whenever new devices are added.

VMware has provided a mechanism that allows the default storage array SATP plug-in rules to be changed to specify the use of a round robin I/O policy globally for any newly discovered LUN's via esxcli commands on the ESX server.

CAUTION: The procedure for changing the SATP plug-in rules to use the Round Robin I/O multipathing policy is only applicable for VMware hosts using 3PARdata InServ Storage Server LUN's. If the host is sharing storage from other vendors, then you should consider what affect changing the rules will have on the storage environment as a whole before making any I/O policy changes.

The VMware native multipathing has two important plug-ins: a Storage Array Type plug-in (SATP) which handles path failover and monitor's path health, and a Path Selection plug-in (PSP) which chooses the best path and routes I/O requests for a specific logical device. The

default SATP VMW_SATP_DEFAULT_AA is mapped to VMW_PSP_FIXED which is VMware's "Fixed", preferred path, I/O multipathing policy (an Active/Passive multipathing scheme).

To show the current default mapping of the Path Selection Policy, issue the following excli command:

```
# esxcli nmp satp list | grep "VMW_SATP_DEFAULT_AA"
Name Default PSP Description
VMW_SATP_DEFAULT_AA VMW_PSP_FIXED Supports non-specific active/active arrays
```

To show the current default SATP rules for FC and iSCSI devices, issue the following excli command:

```
# esxcli nmp satp listrules | egrep "VMW_SATP_DEFAULT_AA"
Name      Vendor Model Driver Transport Options Claim Options
Description
VMW_SATP_DEFAULT_AA    fc     Default for Fibre Channel

VMW_SATP_DEFAULT_AA    iscsi     Default for iSCSI
```

Use the following procedures to change the SATP plug-in rules. These changes require the use of esxcli commands on each ESX server.

1. Make a copy of the device configuration file esx.conf before changing the SATP configuration rule.

```
# cp /etc/vmware/esx.conf /etc/vmware/esx.conf-orig
```

   The contents of the esx.conf file are enabled in the vmkernel at bootup time and get applied to the devices that are discovered.

NOTE: VMware specifically warns not to directly edit the esx.conf file.

2. 3PAR devices are applied with default FC/iSCSI SATP to PSP rule VMW_SATP_DEFAULT_AA / VMW_PSP_FIXED preferred path multipathing. Change the default mapping rule to SATP VMW_SATP_DEFAULT_AA / VMW_PSP_RR resulting in Active/Active round robin multipathing by issuing the following excli command.

```
# esxcli nmp satp setdefaultpsp --satp VMW_SATP_DEFAULT_AA --psp VMW_PSP_RR
```

The default PSP for VMW_SATP_DEFAULT_AA is now VMW_PSP_RR. The above steps can be performed with I/O being served and does not cause any disruption. The change does not require a reboot. It will only be effective for newly discovered LUN's.

3 Verify the changes using the following excli command:

```
# esxcli nmp satp list | grep "VMW_SATP_DEFAULT_AA"
Name Default        PSP        Description
VMW_SATP_DEFAULT_AA VMW_PSP_RR Supports non-specific active/active
arrays
```

CAUTION: The change of the default PSP to VMW_PSP_RR affects all storage devices (FC, iSCSI) that use the SATP VMW_SATP_DEFAULT_AA. If a host server is sharing multiple storage vendors along with 3PAR, and if the other connected storage does not support Active/Active Round Robin multipathing using the same SATP VMW_SATP_DEFAULT_AA, then it's multipathing will be adversely effected. If the other storage uses a different SATP of its own, then the change of SATP VMW_SATP_DEFAULT_AA mapping to VMW_PSP_RR should be done to take advantage of round robin multipathing. You can check a given device SATP-PSP relationship using the `esxcli nmp device list` or `esxcli nmp device list -d <device id>` command.

Example: If 3PAR and storage X are connected to the same host using VMW_SATP_DEFAULT_AA and since storage X does not have it's own SATP then it might cause an issue if storage X does not support Round Robin multipathing. If 3PAR and storage Y are sharing the same host since storage Y has it's own SATP (VMW_SATP_Y) and 3PAR will use (VMW_SATP_DEFAULT_AA) there will be no conflict and the change could be done. You can check the available SATP rules using the `esxcli nmp satp list` command.

The following sample output from a Fibre Channel configuration shows a newly exported LUN that has been automatically set to Round-Robin. Note each path status shown as "Active (I/O)". The path status for an iSCSI configuration would be the same.

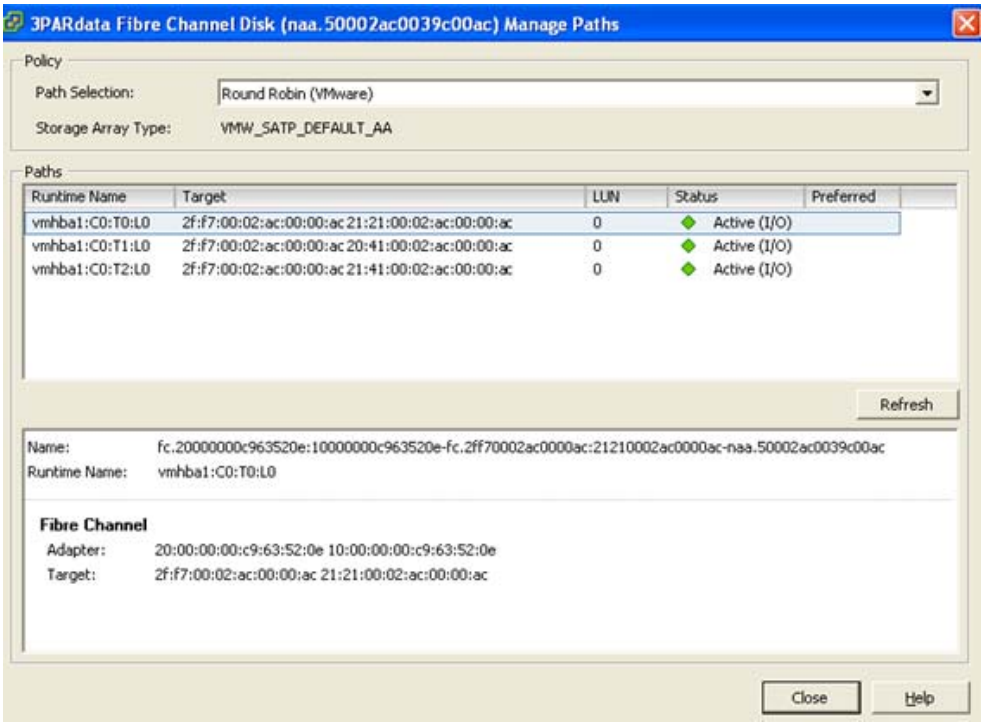

Or

```
# esxcli nmp device list -d naa.50002ac0005800ac
naa.50002ac0005800ac
Device Display Name: 3PARdata Fibre Channel Disk (naa.50002ac0005800ac)
Storage Array Type: VMW_SATP_DEFAULT_AA
Storage Array Type Device Config:
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0;lastPathIndex=none}
Working Paths: vmhba1:C0:T2:L33, vmhba1:C0:T1:L33, vmhba1:C0:T0:L33
```

The changes will be effective for new devices but not for existing devices. In other words, for a newly exported LUN, the SATP/PSP uses the VMW_SATP_DEFAULT_AA / VMW_PSP_RR A/A multipathing, but existing devices still use VMW_SATP_DEFAULT_AA / VMW_PSP_FIXED. The policy is set on a device is saved in esx.conf file and persists even after a reboot, however, VMware specifically warns not to directly edit the esx.conf file.

Multipath I/O path policy changes on existing LUNs must be done on a LUN by LUN basis. The VI / vSphere client GUI may be used for changing the path policy on previously exported LUNs

on an individual basis, or scripting of esxcli commands may be used to create a batch process to make the changes.

For the VI / vSphere method; navigate to the "Manage Paths" screen for each Disk (LUN) and use the Path Selection drop-down menu to select round robin path policy. The path policy change will take place immediately.

For the esxcli method; put together a script that uses the following commands:

List all the 3PAR devices present on the Host

```
# esxcli nmp device list | grep -i naa.50002ac | grep -v Device
naa.50002ac0005800ac
naa.50002ac003b800ac
naa.50002ac0039300ac
```

Change the I/O path policy to round robin for each device identified in the previous output

```
# esxcli nmp device setpolicy --device naa.50002ac0005800ac --psp VMW_PSP_RR
```

Verify that the change has been made

```
# esxcli nmp device list -d naa.50002ac0005800ac
naa.50002ac0005800ac
    Device Display Name: 3PARdata Fibre Channel Disk (naa.50002ac0005800ac)
    Storage Array Type: VMW_SATP_DEFAULT_AA
    Storage Array Type Device Config:
    Path Selection Policy: VMW_PSP_RR
    Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0;lastPathIndex=none}
    Working Paths: vmhba1:C0:T2:L33, vmhba1:C0:T1:L33, vmhba1:C0:T0:L33
```

NOTE: If I/O is active to a LUN and an attempt is made to modify the path policy, a failure can occur indicating:

`error during the configuration of the host: sysinfoException; Status=Busy: Message=Unable to Set"`

If this problem occurs while attempting to change the path policy, reduce the I/O's to that LUN and then try making the desired changes.

4.4 Performance Considerations for Multiple Host Configurations

The information in this section should be considered when using multiple ESX hosts, (or other hosts in conjunction with ESX hosts), that are connected in a fan-in configuration to a pair of InServ Storage Server ports.

NOTE: 3PAR recommends changing the ESX Scsi.ConflictRetries from its default value of 80 to a value of 200 when connected to an InServ Storage Server running InForm OS version 2.2.4 or prior. This change lengthens the time allowed for I/O retries due to ESX SCSI-2 reservation conflicts on VMFS LUNs caused by delayed processing of SCSI-2 reservation commands on the InServ Storage Server, thereby helping to avoid VM I/O failure.

Changing this value can be achieved through use of the VMware ESX VI / vSphere Client: **ESX server > Configuration tab > Software Advanced Settings > Scsi >**, scroll to **Scsi.ConflictRetries** and change the value in the field. Click **OK** to complete the change. A reboot is not required for the change to take effect.

- For all VMware ESX server releases up through ESX 3.5 update 3

The default behavior of an ESX 3.5 update 3 and older servers to “Queue Full” SCSI messages from the InServ Storage Server is to treat them as valid commands and to continue sending data. When continued outstanding commands are being sent to an InServ Storage Server port, the port cannot recover and stops responding for attached hosts.

This type of action is critical where QLogic HBAs are used on the InServ Storage Server because, when the storage port stops responding, the QLogic driver on the InServ Storage Server has to issue a reset of the affected port.

The time in which the InServ Storage Server port is at full capacity and the reset time of the port does not trigger a failover in the ESX server since the ESX server never detects the port going away. This results in a Virtual Machine crash.

- As of VMware ESX server release 3.5 update 4 (and including ESX 4.0 GA)

An algorithm has been added that allows ESX to respond to "Queue Full" and "Busy" SCSI messages from the storage array. The "Queue Full" and "Busy" response by ESX is to back off of I/O for a period of time thus helping to prevent overdriving of the InServ Storage Server

ports. This feature should be enabled as part of an ESX - 3PAR InServ Storage Server deployment.

The "Queue Full"and "Busy" LUN throttling algorithm is disabled by default. To enable the algorithm, complete the following steps:

4 From the VI / vSphere client, select the ESX host **Configuration tab --> Advanced Settings --> Disk.**

5 Scroll to find and adjust the following 3PAR recommended settings:

QFullSampleSize = 32

QFullThreshold = 4

Enabling this algorithm puts ESX "Queue Full" / "Busy" response in line with other 3PAR supported operating systems. With the algorithm enabled, no additional I/O throttling scheme on ESX 3.5 update 4 and newer ESX servers is necessary. Please consult the additional information regarding the ESX "Queue Full" / "Busy" response algorithm found in VMware KB 1008113.

4.4.1 Follow-up Actions/Recommendations for ESX 3.5 update 3 and older:

The behavior of the InServ Storage Server port being over run is limited only to ESX 3.5 update 3 and older servers.

Other 3PAR supported operating systems handle Queue Full notifications by backing off on I/O for a period of time. 3PAR currently does not have any InServ firmware to address this issue but there are recommendations that can help prevent this problem from occurring as detailed in the following sections.

Considering that the ESX 3.5 update 3 and older server default behavior with regard to “Queue Full” messages does not initiate a pause or throttling of I/O, it is important to plan your ESX environments accordingly. This involves tuning the work load of the total number of ESX servers so they do not overrun the Queue Depth of the target InServ HBA Ports.

4.4.1.1 Recommendations for ESX Hosts Attached to a Storage Port on the InServ

For performance and stability, 3PAR recommends that no more than 16 hosts be attached to any one 2Gbps InServ Storage Server port, and no more than 32 hosts attached to any one 4Gbps InServ Storage Server port. This is due to the 3PAR recommendation to throttle the total aggregate of the outstanding commands from all hosts to be less then the queue depth of the InServ Storage Server port and throughput.

NOTE: These recommendations are guidelines for best practice. Adding more than the recommended ESX hosts should only be attempted when the total expected workload is calculated and shown not to overrun either the queue depth or throughput of the storage port.

4.4.1.2 Modifying the Tuneable Parameters for Queue Depth Throttling in ESX 3.X

The default settings for target port queue depth on the ESX Server can be modified to ensure that the total workload of all servers will not overrun the total queue depth of the target InServ storage port. The method endorsed by 3PAR is to limit the queue depth on a per-target basis. This recommendation comes from the simplicity in limiting the number of outstanding commands on a target (InServ port) per ESX server.

The following values can be set on the instances of an HBA in an ESX operating system. These values limit the total number of outstanding commands that the operating system routes to one target.

1) ESX Emulex HBA Target Throttle = tgt_queue_depth
2) ESX Qlocic HBA Target Throttle = ql2xmaxdepth

The formula is as follows:

```
(3PAR port queue depth) / (total number of ESX severs attached) = recommended
target port queue depth.
```

The 3PAR port queue depth limitations used for the calculations are from the listings in *2.3.7.1 Target Port Limitations* on page 2.15.

Example 1 (set up as follows:)

- Qlogic 2G FC HBAs installed on the InServ.
- 16 ESX hosts attached to a QLogic port on the InServ.

Formula:

```
497 / 16 = 31.xxx (recommended max target port queue depth = 31)
```

Example 2 (set up as follows:)

- LSI 2G FC HBA installed on the InServ.

- 12 ESX hosts attached to a LSI port on the InServ.

Formula:

```
510 / 12 = 42.xxx (recommended max target port queue depth = 42)
```

Setting tgt_queue_depth for Emulex in ESX 3.X (Example)

To set the `tgt_queue_depth` for an Emulex FC HBA in ESX 3.X to something other than the default requires a multi step process.

1. Shutdown all of the virtual machines.
2. Log into the ESX service console as root.
3. Make a backup copy of `/etc/vmware/esx.conf`.

```
cp /etc/vmware/esx.conf /etc/vmware/esx.conf.orig
```

4. Identify the Emulex HBA module that is currently loaded.

```
vmkload_mod –l | grep –i lpfc
```

Depending on the module of the HBA, the module can be one of the following:

- lpfcdd_7xx
- lpfcdd_732
- lpfc_740

5. The Target Queue depth can now be modified via the command line using VMware supplied binaries.

The example shows the lpfc_740 module. Use the appropriate module based on the outcome of step 4.

```
esxcfg-module –s “lpfc_tgt_queue_depth=31” lpfc_740
esxcfg-boot –b
```

6. You can check to see that the change has been implemented as follows:

```
esxcfg-module -q
```

7 Reboot the ESX server for the changes to take affect.

Upon boot up, the ESX server will now be throttled to a maximum of 31 (as per example) outstanding commands to the target (InServ port).

Setting ql2xmaxdepth for QLogic in ESX 3.X (Example)

To set the `ql2xmaxdepth` for an QLogic FC HBA in ESX 3.X to something other than the default requires a multi step process.

1 Shutdown all of the virtual machines.

2 Log into the ESX service console as root.

3 Make a backup copy of `/etc/vmware/esx.conf`.

```
cp /etc/vmware/esx.conf /etc/vmware/esx.conf.orig
```

4 Identify the QLogic HBA module that is currently loaded.

```
vmkload_mod -l | grep qla2300
```

Depending on the model of the HBA, the module can be one of the following:

qla2300_707 (ESX Server 3.0.x)

qla2300_707_vmw (ESX 3.5)

5 The target port queue depth can now be modified via the service console command line using VMware supplied binaries.

The example shows the qla2300_707 module. Use the appropriate module based on the outcome of step 4.

```
esxcfg-module -s ql2xmaxqdepth=42 qla2300_707
esxcfg-boot –b
```

The server will now need to be rebooted for the changes to take affect.

6 Reboot the ESX server.

Upon boot up, the ESX server will now be throttled to a max of 42 (as per example) outstanding commands to the target (InServ port).

NOTE: For additional information on changing the queue depth for HBAs, refer to the *VMware Fibre Channel SAN Configuration Guide*.

4.5 ESX / ESXi 4.1 Additional Feature Considerations

ESX / ESXi 4.1 introduces new features related to storage I/O control and integration with storage arrays. 3PAR recommends the usage of features Storage I/O Control (SIOC) and vStorage APIs for Array Integration (VAAI) with ESX / ESXi 4.1 - 3PAR InServ system configurations.

4.5.1 Storage I/O Control (SIOC)

The SIOC feature allows for a new level of monitoring and control of I/O from individual Virtual Machines to a 3PAR InServ Storage Array at the datastore level and across ESX / ESXi servers in a VMware Cluster.

Further information regarding the Storage I/O Control (SIOC) feature and considerations for its deployment may be found in the VMware technical white paper "Storage I/O Control Technical Overview and Considerations for Deployment"

http://www.vmware.com/files/pdf/techpaper/VMW-vSphere41-SIOC.pdf

4.5.2 vStorage APIs for Array Integration (VAAI)

In partnership with VMware, 3PAR has developed an ESX / ESXi 4.1 plug-in that enables a new set of SCSI commands to be used by ESX / ESXi 4.1 in conjunction with 3PAR InServ. VMware refers to this newly incorporated set of SCSI commands as the "primitives".

ESX extensions that make use of these primitives are collectively referred to as vStorage APIs for Array Integration (VAAI). The VMware primitives enable an ESX/ESXi host to convey virtual machine operations to storage hardware at a meta level instead of at the traditional data level. This reduces operational latency and traffic on the FC fabric / iSCSI network. Some of these primitives enable the storage hardware to participate in block allocation and deallocation functions for virtual machines. These primitives are also known as hardware offloads.

A Brief description of the "primitives":

- Full Copy (XCOPY), enables the storage array to make full copies of data within the array without having to have the ESX Server read and write the data. This offloads some data copy processes to the storage array.
- Block Zeroing (WRITE-SAME), enables the storage array to zero-out a large number of blocks within the array without having to have the ESX Server write the zeros as data and helps expedite the provisioning of VMs. This offloads some of the file space zeroing functions to the storage array.
- Hardware Assisted Locking (ATS), provides an alternative to SCSI reservations as a means to protect the metadata for VMFS cluster file systems and helps improve the scalability of large ESX server farms sharing a datastore.

Support for VMware VAAI functionality is available via installation of the 3PAR VAAI Plug-in 1.1.0 on ESX / ESXi 4.1 in combination with 3PAR InForm OS version 2.3.1 MU2 (minimum).

Refer to the site http://www.3par.com/services/download.html for further information on VMware VAAI, the 3PAR VAAI Plug-in for ESX / ESXi 4.1 installation package, and the 3PAR VAAI Plug-in 1.1.0 for VMware vSphere 4.1 User's Guide.

5 Configuring the Host for FC over Ethernet (FCoE) Connection

In this chapter

This chapter describes the procedures for setting up an ESX software Fibre Channel over Ethernet (FCoE) configuration with an InServ Storage Server. The instructions in this chapter should be used in conjunction with the VMware ESX Server Configuration Guide for FCoE Software Initiator Setup.

5.1 Configuring an InServ Port for a FCoE Host Connection

There is nothing unique that needs to be configured on the InServ ports to be able to connect a host with a Converged Network Adapter (CNA) card configured with FCoE. As far as the InServ system is concerned, the initiator coming from the FCoE card thru the FCoE Forwarder switch is treated as just another Fibre Channel device connecting to the InServ ports.

The same guidelines described in Chapter 2, *Configuring the InServ Storage Server for Fibre Channel* and Chapter 4, *Configuring the Host for a Fibre Channel Connection* must be followed when a server with a CNA card configured with FCoE is connected to InServ ports.

NOTE: For specific configurations that support FCoE CNAs and forwarder switches, please refer to the 3PAR 2.3.x Configuration Matrix.

5.2 Configuring a Converged Network Adapter (CNA) and FCoE Forwarder Switch

The following summary describes the general steps you should follow to configure a Converged Network Adapter (CNA) and FCoE Forwarder Switch.

NOTE: For complete and detailed instructions for configuring a server with a given Converged Network Adapter, please refer to the CNA manufacturer documentation. To configure the FCoE Forwarder switch, please refer to the switch vendor documentation.

1. Install the CNA card in the server just like any other PCIe card - refer to the server vendor documentation.
2. Install the CNA card driver following the CNA card installation instructions (it assumes the server is already running a supported Operating System).
3. Physically connect the server CNA card ports to the FCoE Forwarder switch and configure the FCoE Forwarder switch ports - refer to the switch vendor documentation.
4. Configure the InServ ports per the guidelines in section *2.1 Configuring the InServ Storage Server Running InForm O.S. 2.3.X* on page 2.2 and connect the InServ port either to the FCoE Forwarder FC switch ports or the Fibre Channel fabric connected to the FCoE Forwarder.
5. Create FC zones for the host initiators ports and the InServ target port. Once the initiators have logged in to the InServ target ports, create a host definition and provision storage to the host.

NOTE: It is not possible to connect a server with a CNA directly to the InServ system. A FCoE Forwarder switch must be used.

6 Configuring the Host for an iSCSI Connection

In this chapter

This chapter describes the procedures for setting up an ESX software iSCSI configuration with an InServ Storage Server. The instructions in this chapter should be used in conjunction with the VMware ESX Server Configuration Guide for iSCSI Software Initiator Setup.

6.1 InStalling iSCSI on the VMware ESX Server

Software iSCSI drivers for VMware supported NICs are included as part of the ESX OS installation package supplied by VMware. Updates and/or patches for the software iSCSI drivers can be acquired through VMware support.

The following illustration shows an example of an ESX iSCSI Software Initiator configuration with 2 Servers:

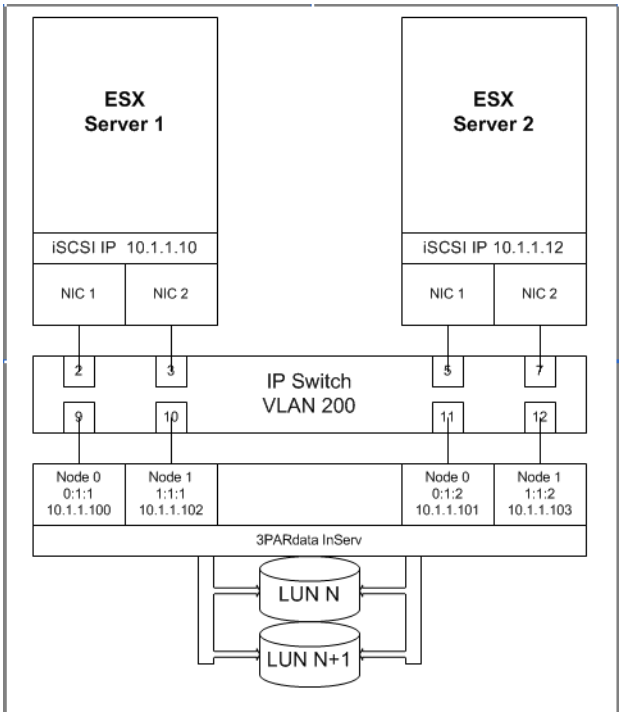

Figure 6-1. ESX iSCSI Software Initiator Configuration

NOTE: When multiple teamed NICs are configured, all InServ iSCSI ports and ESX Server iSCSI NICs must be in the same VLAN of the IP switch.

6.2 Installing Virtual Machine Guest Operating System (GOS)

The VMware ESX Server documentation lists recommended Virtual Machine guest operating systems (GOS) and their installation and setup as Virtual Machines. Refer to the VMware ESX Server documentation for information on setting up your Virtual Machine configuration.

NOTE: VMware and 3PAR recommend the LSI logic adapter emulation for Windows 2003 Servers. The LSI Logic adapter is also the default option for Windows 2003 when creating a new Virtual Machine. 3PAR testing has noted a high incidence of Windows 2003 Virtual Machine failures during an ESX multipath failover/failback event when the BUS Logic adapter is used with Windows 2003 VMs.

NOTE: 3PAR testing indicates that the SCSI timeout value for Virtual Machine guest operating systems should be 60 seconds in order to successfully ride out path failovers at the ESX layer. Most guest operating systems supported by VMware have a default SCSI timeout value of 60 seconds, but this value should be checked and verified for each GOS installation. In particular, Red Hat 4.x guest operating systems should have their SCSI timeout value changed from their default value of 30 seconds to 60 seconds.

This command line can be used to set the SCSI timeout on all SCSI devices presented to a Red Hat 4.x Virtual Machine to 60 seconds:

```
find /sys -name timeout | grep "host.*target.*timeout" | xargs -n 1 echo "echo 60 >"|sh
```

This must be added as a line in /etc/rc.local of the Red Hat 4.x guest OS in order for the timeout change to be maintained with a Virtual Machine reboot.

Example of modified /etc/rc.local file:

```
# cat /etc/rc.local
#!/bin/sh
#
# This script will be executed *after* all the other init scripts.
# You can put your own initialization stuff in here if you don't
# want to do the full Sys V style init stuff.

find /sys -name timeout | grep "host.*target.*timeout" | xargs -n 1 echo
"echo 60 >"|sh
touch /var/lock/subsys/local
```

CAUTION: VMware KB Article 51306: VMware has identified a problem with RHEL5 (GA), RHEL4 U4, RHEL4 U3, SLES10 (GA), and SLES9 SP3 guest operating systems. Their file systems may become read-only in the event of busy I/O retry or path failover of the ESX Server's SAN or iSCSI storage.

Because of this known issue, 3PAR does not recommend, and does not support the usage of RHEL5 (GA), RHEL4 U4, RHEL4 U3, SLES10 (GA), and SLES9 SP3 as guest operating systems for Virtual Machines on VMware ESX Servers attached to 3PAR InServ Storage Servers.

6.3 Creating a VMkernel Port

The following steps describe how to set up a VMKernel port.

1. Log into the VI / Vsphere Client and select the server from the inventory panel. The hardware configuration page for this server appears.
2. Click the **Configuration** tab and click **Networking**.
3. Click the **Add Networking** Link. The Add Network wizard appears.

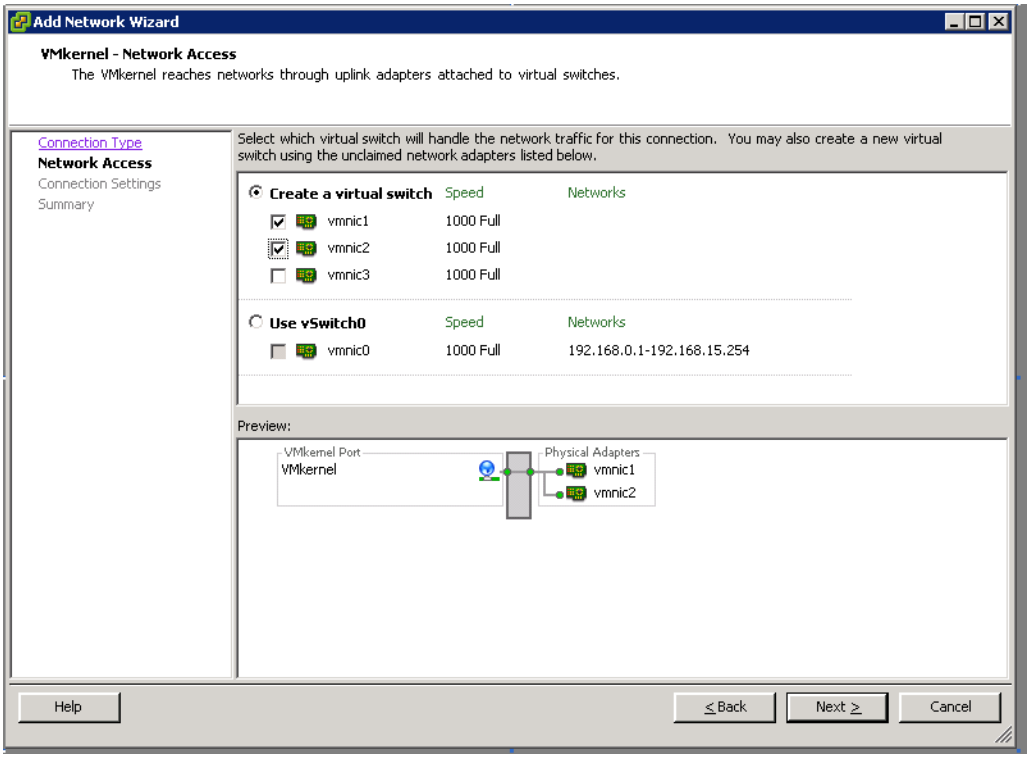

4. Select **VMkernel** and click **Next**. This lets you connect to the VMkernel which runs services for iSCSI storage to the physical network. The Network Access page appears.
5. Select the **Create a Virtual Switch** and select the NICs that will be used for iSCSI. (In this example, 2 NICs are selected to configure an active/active teamed NICs that will connect to the 3PAR storage array.)

6 Configure Active/Active NIC teaming by bringing up all of the NIC adapters being used as "Active Adapters" in the vSwitch Properties. For each ESX server, using the **VI / vSphere client > Configuration tab > Networking > Properties >** click on the "**Edit**" radio button > highlight and use the "**Move Up**" radio button to bring each of the NIC adapters being used for NIC teaming from the "Standby Adapters" or "Unused Adapters" section to the "Active Adapters" section.

The screen below identifies that this has been completed for NIC adapters vmnic1 and vmnic2.

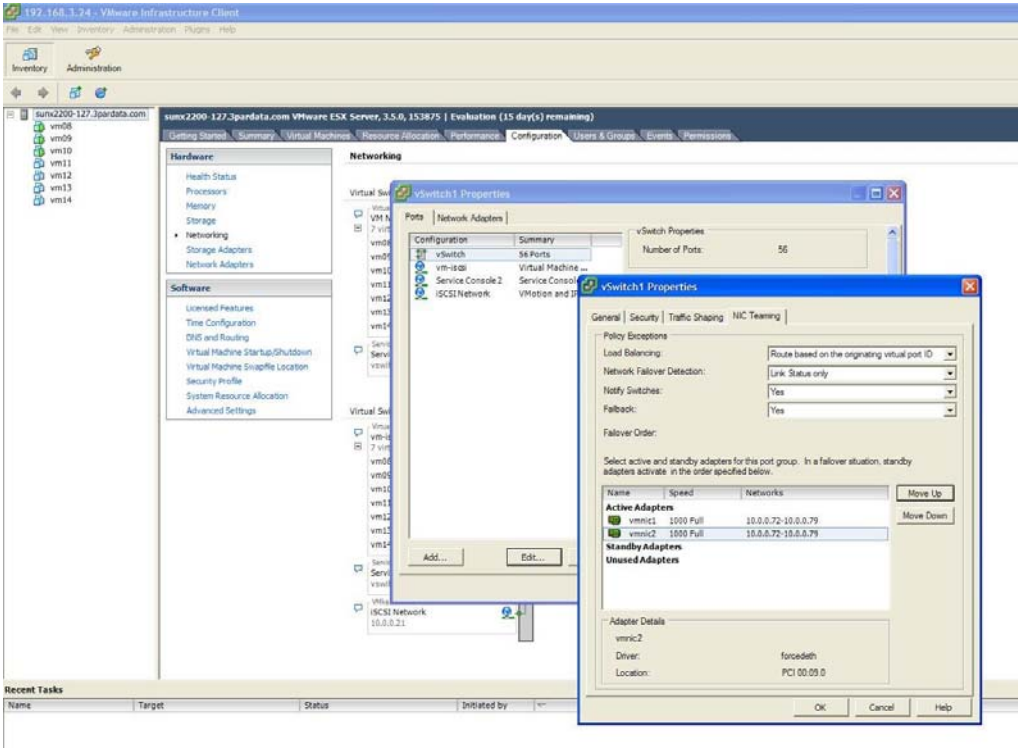

7 Click **Ok** to complete.

NOTE: 3PAR recommends an Active/Active NIC Teaming configuration for best failover performance.

8 Click **Next**.

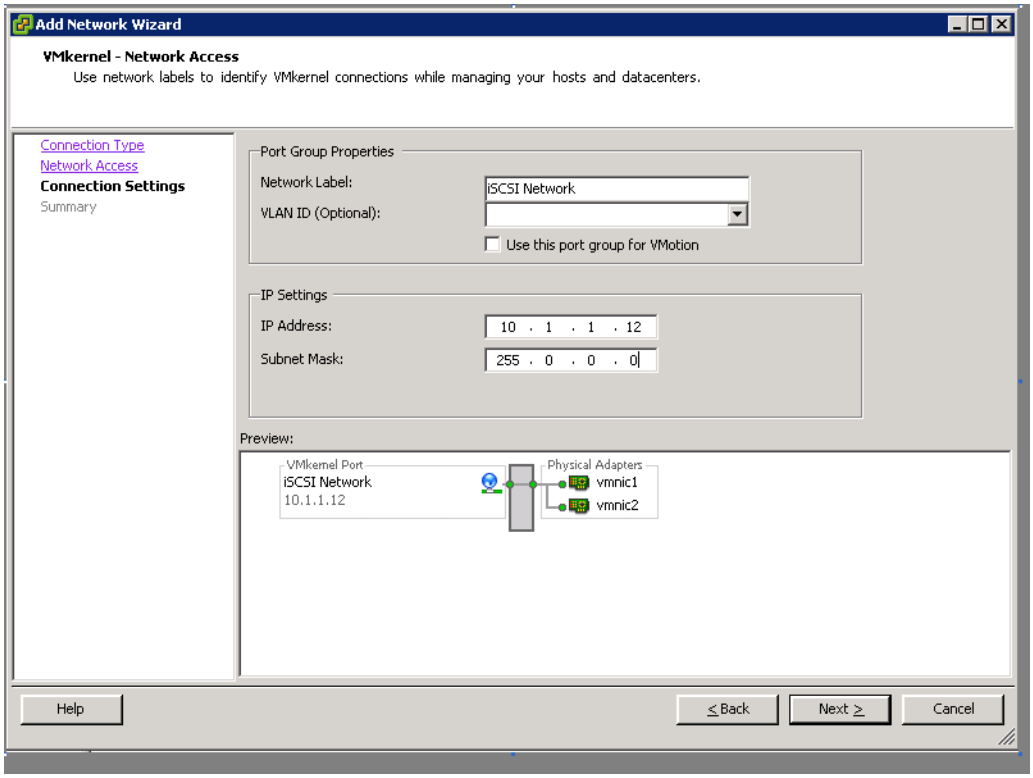

9 Select a network label and the IP address that will be used for the iSCSI network.

10 Click **Next**.

A window shows your settings about to be completed.

11 Click **Close**.

A window will appear stating that no DNS setting and gateways have been set.

12 Add a VMkernel default gateway that is in the same subnet as the iSCSI network. Refer to the VMware ESX Server Configuration Guide for detailed instructions regarding these settings.

13 Click **Finish** when you have completed all the necessary changes.

6.4 Configuring a Service Console Connection for the iSCSI Storage

The following steps describe how to configure the Service Console connection for the iSCSI storage.

1. From the **Configuration** tab, and the **Networking** tab, click on the Properties for the vSwitch associated with the VMkernel port you just created for the iSCSI network.
2. Click **Add**.

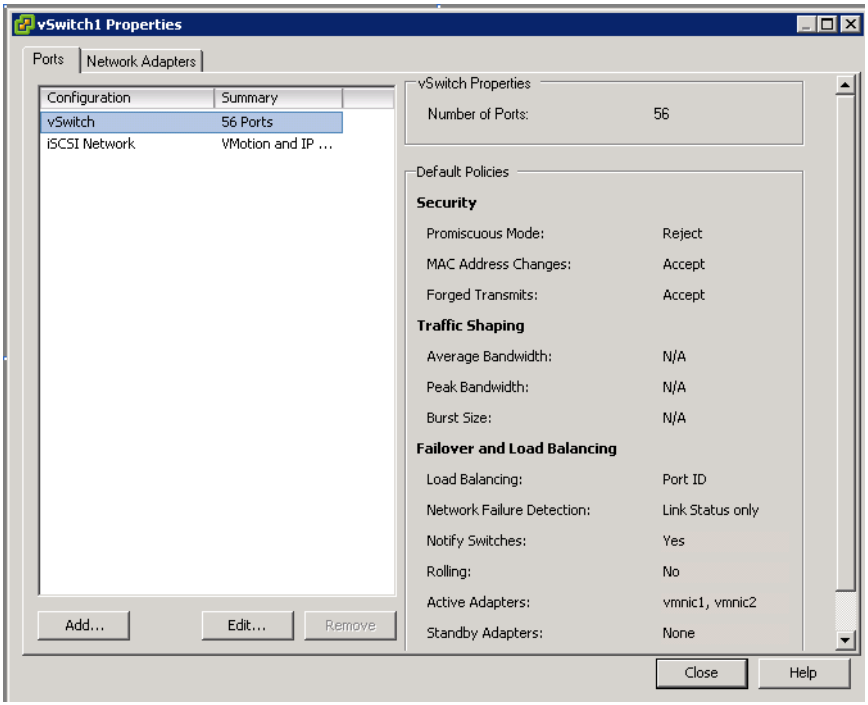

3. Select the radio button for **Service Console** to add support for host management traffic.

4 Click **Next**.

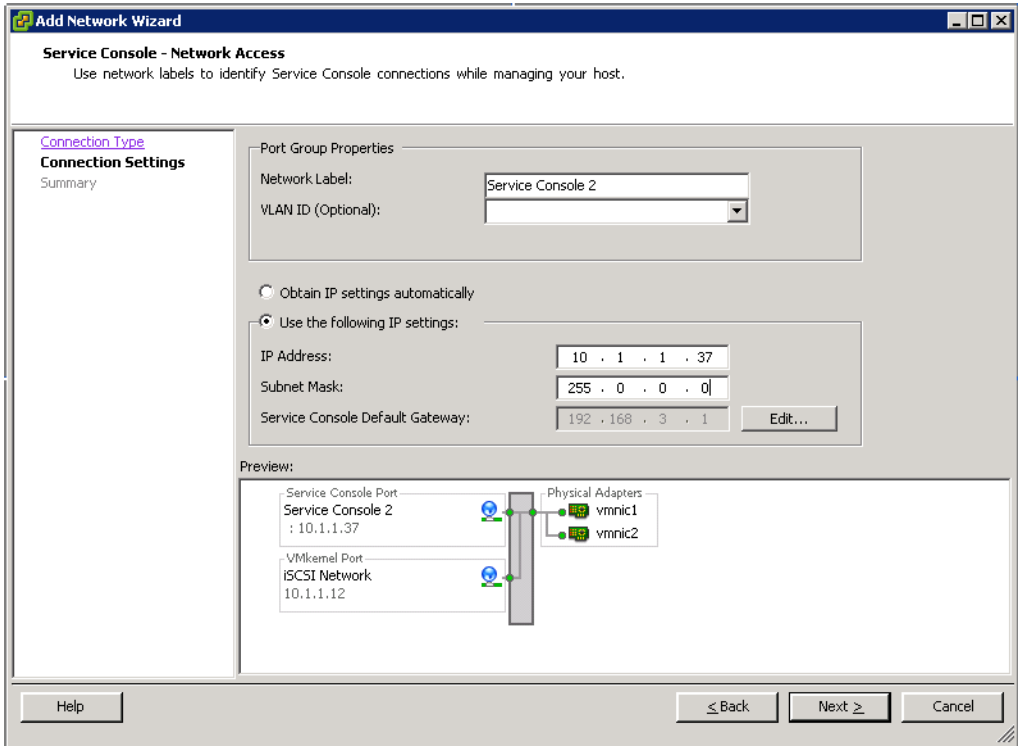

5 Enter the IP address for the Service console used to communicate with the iSCSI software initiator. The IP address must be in the same subnet as the iSCSI.

6 Click **Next**. A window appears showing the changes / additions that have been made.

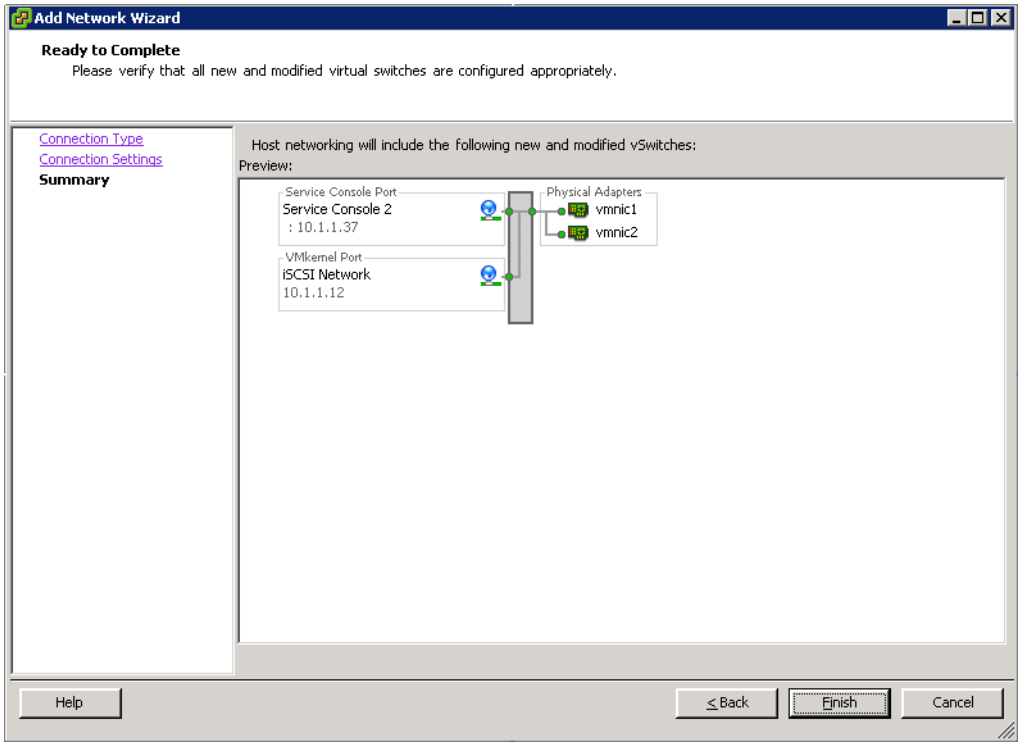

7 Click **Finish**.

8 Close all windows associated with the network configuration.

9 Check the Configuration Display.

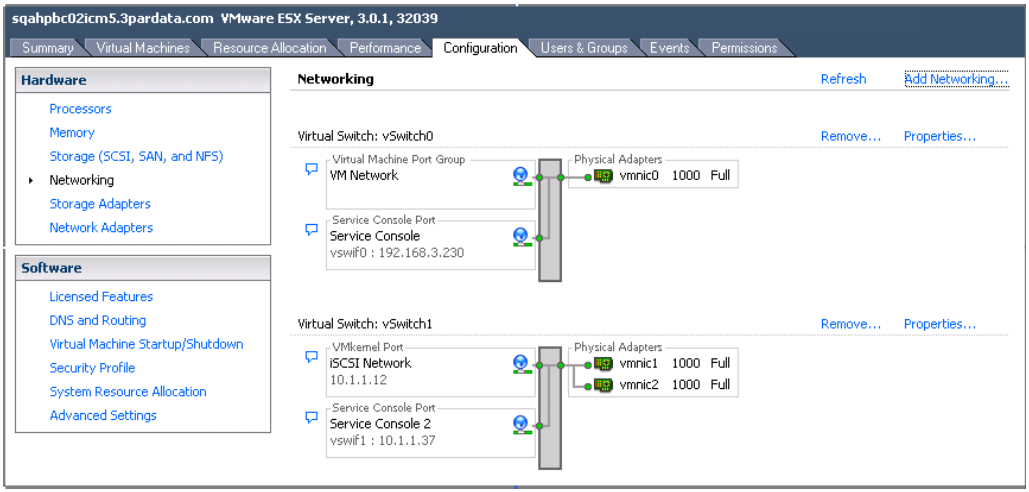

The new network configuration is displayed with the addition of the iSCSI network.

You should now be able to ping the InServ Storage Server ports that were previously defined from the COS.

6.5 Configuring the VMware iSCSI Initiator

The following steps show how to configure the VMware iSCSI Initiator.

1 With the **Configuration** tab selected, click the **Security Profile** option from the Software menu box.

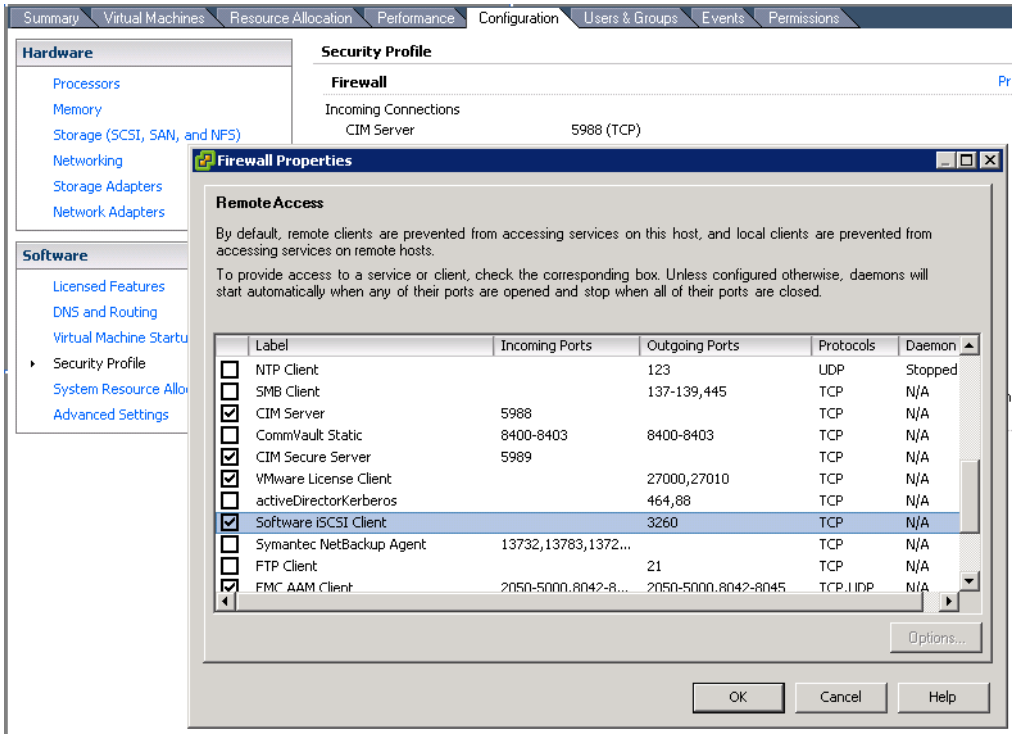

2 Open up the ports that will be used for the iSCSI connection, then click **OK**.

3 The iSCSI software initiator needs to be enabled before the ESX server can use it. Click on the **Storage Adapters** option in the Hardware menu box.

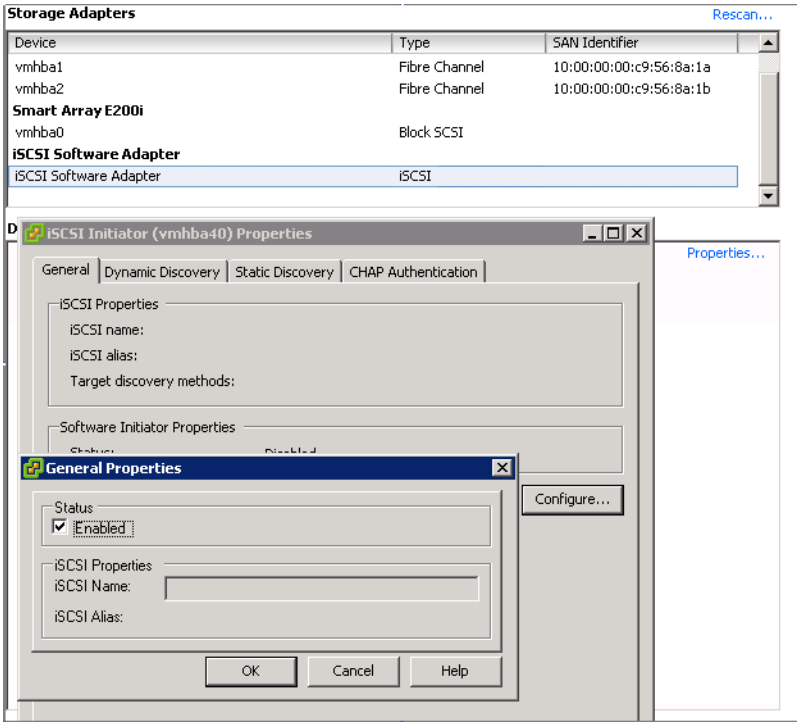

4 From the ESX Server Configuration tab, select the **iSCSI Software Adapter**.

5 Click on the **Properties** tab.

6 Select the **General** tab.

7 Click **Configure...**.

8 Select the **Enabled** check box for the status.

9 click **OK**.

10 Click the **Dynamic Discovery** tab.

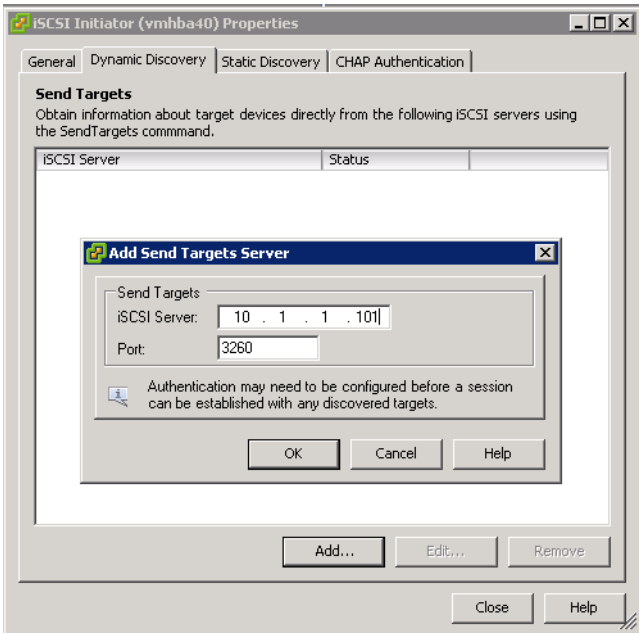

11 Click **Add...**

12 Enter the IP address of one of the previously defined InServ iSCSI ports.

13 Click **OK**.

14 Add additional InServ iSCSI ports if they exist and have been defined on the InServ.

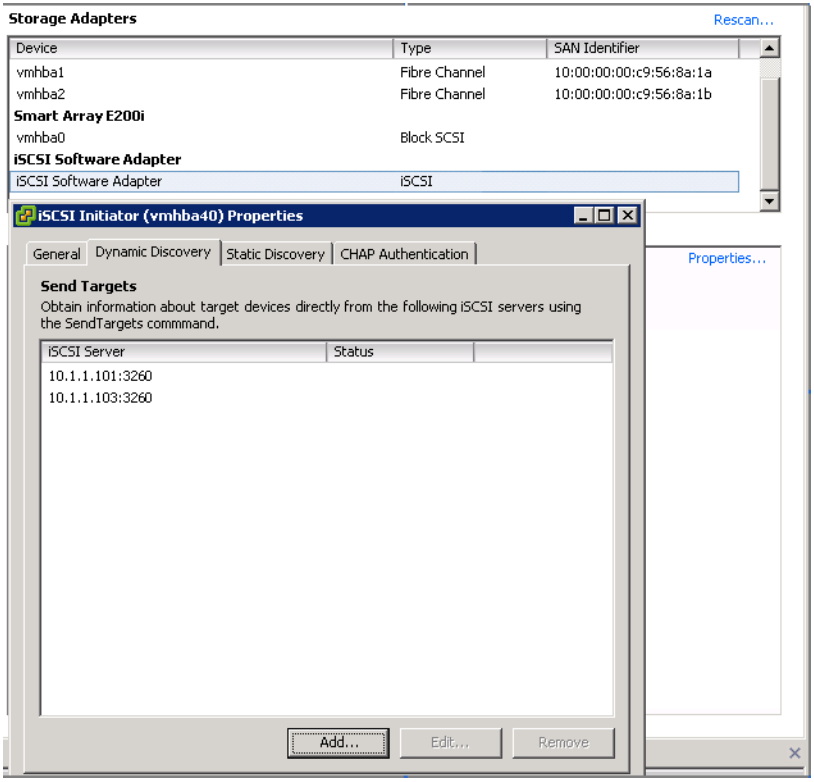

15 When all of the desired InServ iSCSI ports have been added to the **Dynamic Discovery** tab, close this window.

16 Reboot the ESX Server**.** If Virtual Machines are active, shut them down.

The ESX Server and InServ should now be configured for use. Using the showhost command on the InServ, The new iSCSI connections should now show as present.

```
# showhost
Id Name         ----------------WWN/iSCSI_Name---------------- Port
   --           iqn.1998-01.com.vmware:hpdl380-01-11a38a59     0:1:2
                iqn.1998-01.com.vmware:hpdl380-01-11a38a59     1:1:2
```

As new LUNs are exported to the ESX Server iSCSI host, a rescan must be performed on the iSCSI software adapter. This is accomplished in the same manner that Fiber Channel LUNs would be rescanned.

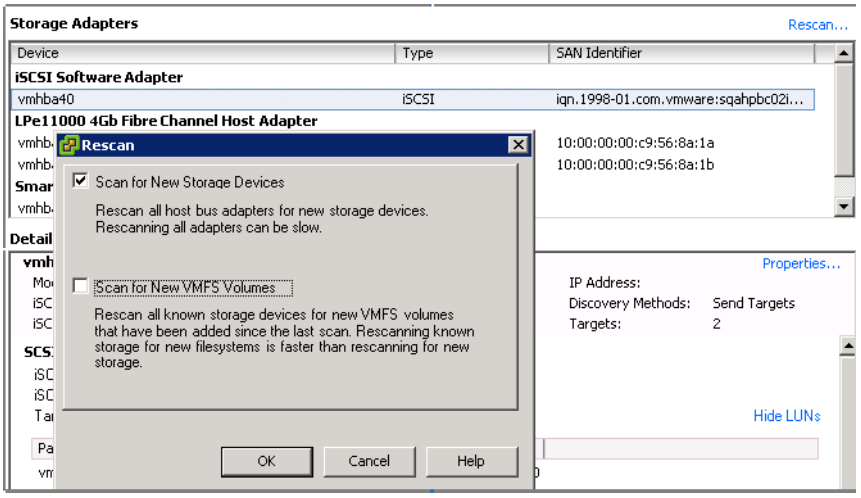

Click **Rescan**, select the **Scan for New Storage Devices** check box, then click **OK** to rescan for new LUNs exported to the ESX iSCSI Server.

6.6 iSCSI Failover Considerations and Multipath Load Balancing

The VMware ESX Server 3.0 - 3.5 includes failover with multi-path support to maintain a constant connection between the ESX host and the InServ Storage Server array. VMware terms this multipath support as "active/active" with a choice of two "path policies" called "FIXED" or "MRU". As of ESX 4.0, a third path policy choice of "round robin" is available. The path policies can be modified on a per InServ volume (LUN) basis by right clicking on the device listing and selecting the "properties" function from the VI / vSphere Client menu. A pop-up window allows you to 'manage paths' whereby the properties of the paths to the volume that was previously selected can be modified. Using this control, you can select the path policy, and specify which path is the active preferred path to a volume on the storage array, or which path is the standby path within the FIXED path policy scheme. Additionally, paths can be disabled to prevent any traffic over a specific path to a volume on the storage array.

The VI / vSphere client GUI interface only allows for settings to be changed on a volume by volume (LUN by LUN) basis. The GUI interface is appropriate and preferred for use in managing I/O paths within the FIXED path policy scheme. Refer to section 4.3.1 for procedures on implementing and configuring the round-robin path policy on ESX 4.0 with 3PAR InServ.

- "Active/active, with a path policy of "round-robin" is the preferred multipath implementation for ESX 4.0. Refer to section 4.3.1 for procedures on implementing and configuring the round-robin path policy on ESX 4.0 with 3PAR InServ.
- "Active/active, with a path policy of "fixed" and the preferred/active paths manually set to balance I/O load evenly across all paths is the preferred multipath implementation for ESX 3.0 - 3.5.
  - In the event the active path is detected as having failed or has been disabled either at the fabric switch, or on the storage array, all ESX Server I/O to the storage array continues by failing over to a 'standby' path. When the ESX server detects that the preferred path has been recovered or is enabled, I/O from the ESX server then resumes on the preferred path -- assuming a preferred path policy had previously been set to that path.
  - I/O from the ESX server should be manually distributed or balanced when 2 or more paths exist to more than one InServ volume on the storage array. Manually balancing the loads across available paths may improve I/O performance. This path load balancing to the storage array is dependant on the number of I/O's that are targeted for specific volumes on the storage array. Tuning I/O's to specific volumes on specific paths to the storage array varies from configuration to configuration and is totally dependant on the workload from the ESX Server and the virtual machines to the devices on the storage array.

The following vSphere client screen shot depicts a LUN with five I/O paths in a FIXED I/O policy scheme. The path marked Active (I/O) with the '*' in the Preferred column is the path chosen as preferred, and is the path to which all I/O is currently assigned for the given LUN. The other paths listed are active, but in 'standby' mode. The paths in active 'standby' will not be used for I/O traffic for this LUN unless the preferred path fails.

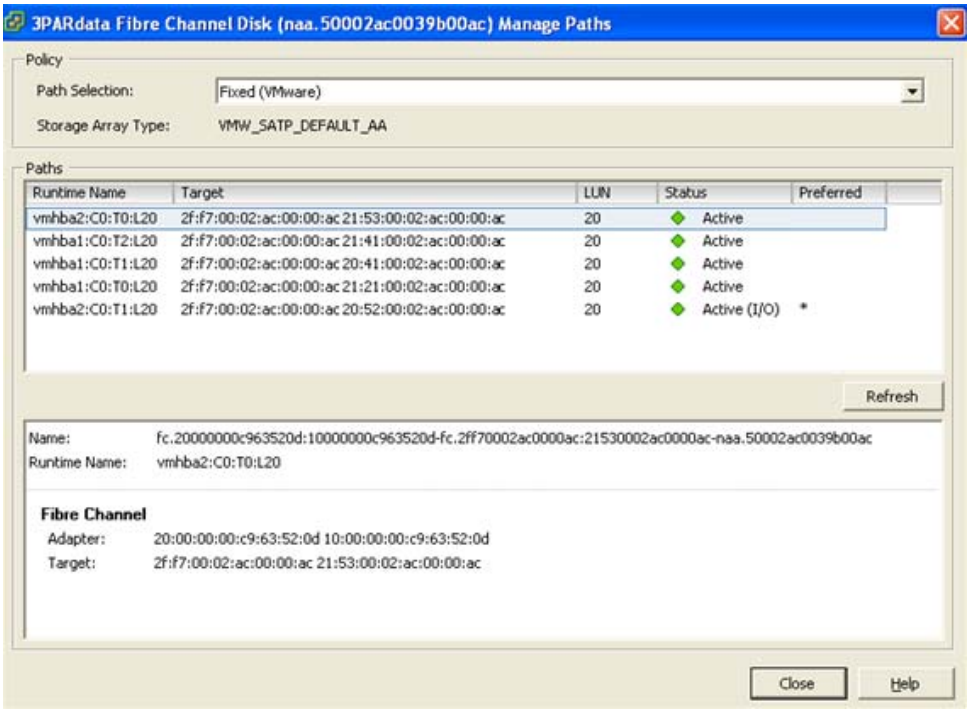

- Active/active, with a path policy of MRU (Most Recent Used) does not maintain or re-instate balancing of I/O load after a failover/failback multipath event. This could leave I/O in an unplanned for and unbalanced state which may yield significant I/O performance issues. Implementation of a MRU (Most Recent Used) path policy is not recommended by 3PAR.

NOTE: If I/O is active to a LUN and an attempt is made to modify the path policy, a failure can occur indicating:

```
"error during the configuration of the host:
sysinfoException; Status=Busy: Message=Unable to Set".
```

If this problem occurs while attempting to change the path policy, reduce the I/O's to that LUN and then try making the desired changes.

For additional information on this topic, please refer to the chapter on "Multipathing" contained in the VMware SAN Configuration Guide.

6.6.1 Configuring Round Robin Multipathing on ESX 4.0

With Version 4.0 onwards, VMware supports a round robin I/O path policy for Active/Active storage arrays such as 3PAR InServ Storage Sever. An Active/Active round robin I/O path policy is the preferred configuration for ESX 4.x, however, the policy is not enabled by default. The default I/O path policy for ESX 4.0 is VMware's FIXED path policy scheme.

Managing a round robin I/O path policy scheme though the VI / vSphere client GUI for a large network can be cumbersome and a challenge to maintain because the policy must be specified for each LUN individually and updated whenever new devices are added.

VMware has provided a mechanism that allows the default storage array SATP plug-in rules to be changed to specify the use of a round robin I/O policy globally for any newly discovered LUN's via esxcli commands on the ESX server.

CAUTION: The procedure for changing the SATP plug-in rules to use the Round Robin I/O multipathing policy is only applicable for VMware hosts using 3PARdata InServ Storage Server LUN's. If the host is sharing storage from other vendors, then you should consider what affect changing the rules will have on the storage environment as a whole before making any I/O policy changes.

The VMware native multipathing has two important plug-ins: a Storage Array Type plug-in (SATP) which handles path failover and monitor's path health, and a Path Selection plug-in (PSP) which chooses the best path and routes I/O requests for a specific logical device. The default SATP VMW_SATP_DEFAULT_AA is mapped to VMW_PSP_FIXED which is VMware's "Fixed", preferred path, I/O multipathing policy (an Active/Passive multipathing scheme).

To show the current default mapping of the Path Selection Policy, issue the following excli command:

```
# esxcli nmp satp list | grep "VMW_SATP_DEFAULT_AA"
Name Default PSP Description
VMW_SATP_DEFAULT_AA VMW_PSP_FIXED Supports non-specific active/active arrays
```

To show the current default SATP rules for FC and iSCSI devices, issue the following excli command:

```
# esxcli nmp satp listrules | egrep "VMW_SATP_DEFAULT_AA"
Name      Vendor Model Driver Transport Options Claim Options
Description
VMW_SATP_DEFAULT_AA    fc     Default for Fibre Channel

VMW_SATP_DEFAULT_AA    iscsi     Default for iSCSI
```

Use the following procedures to change the SATP plug-in rules. These changes require the use of esxcli commands on each ESX server.

1 Make a copy of the device configuration file esx.conf before changing the SATP configuration rule.

```
# cp /etc/vmware/esx.conf /etc/vmware/esx.conf-orig
```

The contents of the esx.conf file are enabled in the vmkernel at bootup time and get applied to the devices that are discovered.

NOTE: VMware specifically warns not to directly edit the esx.conf file.

2 3PAR devices are applied with default FC/iSCSI SATP to PSP rule VMW_SATP_DEFAULT_AA / VMW_PSP_FIXED preferred path multipathing. Change the default mapping rule to SATP VMW_SATP_DEFAULT_AA / VMW_PSP_RR resulting in Active/Active round robin multipathing by issuing the following excli command.

```
# esxcli nmp satp setdefaultpsp --satp VMW_SATP_DEFAULT_AA --psp VMW_PSP_RR
```

The default PSP for VMW_SATP_DEFAULT_AA is now VMW_PSP_RR. The above steps can be performed with I/O being served and does not cause any disruption. The change does not require a reboot. It will only be effective for newly discovered LUN's.

3 Verify the changes using the following excli command:

```
# esxcli nmp satp list | grep "VMW_SATP_DEFAULT_AA"
Name Default        PSP        Description
VMW_SATP_DEFAULT_AA VMW_PSP_RR Supports non-specific active/active arrays
```

CAUTION: The change of the default PSP to VMW_PSP_RR affects all storage devices (FC, iSCSI) that use the SATP VMW_SATP_DEFAULT_AA. If a host server is sharing multiple storage vendors along with 3PAR, and if the other connected storage does not support Active/Active Round Robin multipathing using the same SATP VMW_SATP_DEFAULT_AA, then it's multipathing will be adversely effected. If the other storage uses a different SATP of its own, then the change of SATP VMW_SATP_DEFAULT_AA mapping to VMW_PSP_RR should be done to take advantage of round robin multipathing. You can check a given device SATP-PSP relationship using the `esxcli nmp device list` or `esxcli nmp device list -d <device id>` command.

Example: If 3PAR and storage X are connected to the same host using VMW_SATP_DEFAULT_AA and since storage X does not have it's own SATP then it might cause an issue if storage X does not support Round Robin multipathing. If 3PAR and storage Y are sharing the same host since storage Y has it's own SATP (VMW_SATP_Y) and 3PAR will use (VMW_SATP_DEFAULT_AA) there will be no conflict and the change could be done. You can check the available SATP rules using the `esxcli nmp satp list` command.

The following sample output from a Fibre Channel configuration shows a newly exported LUN that has been automatically set to Round-Robin. Note each path status shown as "Active (I/O)". The path status for an iSCSI configuration would be the same.

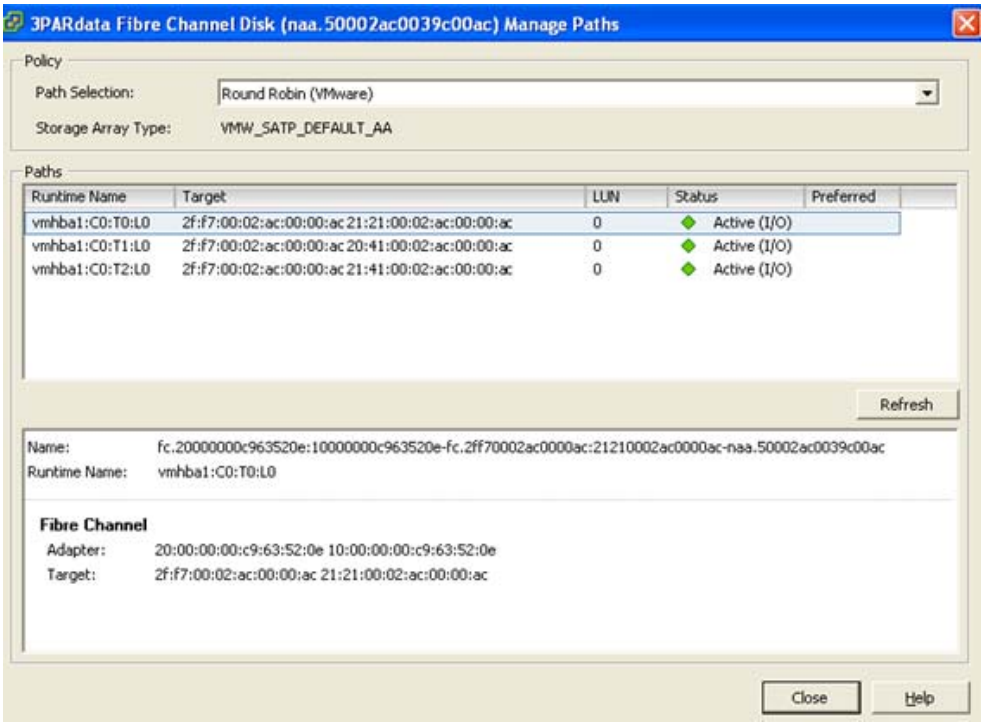

Or

```
# esxcli nmp device list -d naa.50002ac0005800ac
naa.50002ac0005800ac
Device Display Name: 3PARdata Fibre Channel Disk (naa.50002ac0005800ac)
Storage Array Type: VMW_SATP_DEFAULT_AA
Storage Array Type Device Config:
Path Selection Policy: VMW_PSP_RR
Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0;lastPathIndex=none}
Working Paths: vmhba1:C0:T2:L33, vmhba1:C0:T1:L33, vmhba1:C0:T0:L33
```

The changes will be effective for new devices but not for existing devices. In other words, for a newly exported LUN, the SATP/PSP uses the VMW_SATP_DEFAULT_AA / VMW_PSP_RR A/A multipathing, but existing devices still use VMW_SATP_DEFAULT_AA / VMW_PSP_FIXED. The policy is set on a device is saved in esx.conf file and persists even after a reboot, however, VMware specifically warns not to directly edit the esx.conf file.

Multipath I/O path policy changes on existing LUNs must be done on a LUN by LUN basis. The VI / vSphere client GUI may be used for changing the path policy on previously exported LUNs

on an individual basis, or scripting of esxcli commands may be used to create a batch process to make the changes.

For the VI / vSphere method; navigate to the "Manage Paths" screen for each Disk (LUN) and use the Path Selection drop-down menu to select round robin path policy. The path policy change will take place immediately.

For the esxcli method; put together a script that uses the following commands:

List all the 3PAR devices present on the Host

```
# esxcli nmp device list | grep -i naa.50002ac | grep -v Device
naa.50002ac0005800ac
naa.50002ac003b800ac
naa.50002ac0039300ac
```

Change the I/O path policy to round robin for each device identified in the previous output

```
# esxcli nmp device setpolicy --device naa.50002ac0005800ac --psp VMW_PSP_RR
```

Verify that the change has been made

```
# esxcli nmp device list -d naa.50002ac0005800ac
naa.50002ac0005800ac
    Device Display Name: 3PARdata Fibre Channel Disk (naa.50002ac0005800ac)
    Storage Array Type: VMW_SATP_DEFAULT_AA
    Storage Array Type Device Config:
    Path Selection Policy: VMW_PSP_RR
    Path Selection Policy Device Config:
{policy=rr,iops=1000,bytes=10485760,useANO=0;lastPathIndex=none}
    Working Paths: vmhba1:C0:T2:L33, vmhba1:C0:T1:L33, vmhba1:C0:T0:L33
```

NOTE: If I/O is active to a LUN and an attempt is made to modify the path policy, a failure can occur indicating:

```
error during the configuration of the host: sysinfoException;
Status=Busy: Message=Unable to Set"
```

If this problem occurs while attempting to change the path policy, reduce the I/O's to that LUN and then try making the desired changes.

6.7 Performance Considerations for Multiple Host Configurations

The information in this section should be considered when using multiple ESX hosts, (or other hosts in conjunction with ESX hosts), that are connected in a fan-in configuration to a pair of InServ Storage Server ports.

NOTE: 3PAR recommends changing the ESX Scsi.ConflictRetries from its default value of 80 to a value of 200 when connected to an InServ Storage Server running InForm OS version 2.2.4 or prior. This change lengthens the time allowed for I/O retries due to ESX SCSI-2 reservation conflicts on VMFS LUNs caused by delayed processing of SCSI-2 reservation commands on the InServ Storage Server, thereby helping to avoid VM I/O failure.

Changing this value can be achieved through use of the VMware ESX VI / vSphere Client as follows:

ESX server > Configuration tab > Software Advanced Settings > Scsi >

Scroll to **Scsi.ConflictRetries** and change the value in the field. Click **OK** to complete the change. A reboot is not required for the change to take effect.

6.8 ESX / ESXi 4.1 Additional Feature Considerations

ESX / ESXi 4.1 introduces new features related to storage I/O control and integration with storage arrays. 3PAR recommends using Storage I/O Control (SIOC) and vStorage APIs for Array Integration (VAAI) with ESX / ESXi 4.1 for 3PAR InServ system configurations.

6.8.1 Storage I/O Control (SIOC)

The SIOC feature allows for a new level of monitoring and control of I/O from individual Virtual Machines to a 3PAR InServ Storage Array at the datastore level and across ESX / ESXi servers in a VMware Cluster.

Further information regarding the Storage I/O Control (SIOC) feature and considerations for its deployment may be found in the VMware technical white paper "Storage I/O Control Technical Overview and Considerations for Deployment"

http://www.vmware.com/files/pdf/techpaper/VMW-vSphere41-SIOC.pdf

6.8.2 vStorage APIs for Array Integration (VAAI)

In partnership with VMware, 3PAR has developed an ESX / ESXi 4.1 plug-in that enables a new set of SCSI commands to be used by ESX / ESXi 4.1 in conjunction with 3PAR InServ Servers. VMware refers to this newly incorporated set of SCSI commands as the "primitives".

ESX extensions that make use of these primitives are collectively referred to as vStorage APIs for Array Integration (VAAI). The VMware primitives enable an ESX/ESXi host to convey virtual machine operations to storage hardware at a meta level instead of at the traditional data level. This reduces operational latency and traffic on the FC fabric / iSCSI network. Some of these primitives enable the storage hardware to participate in block allocation and deallocation functions for virtual machines. These primitives are also known as hardware offloads.

A Brief description of the "primitives":

- Full Copy (XCOPY), enables the storage array to make full copies of data within the array without having to have the ESX Server read and write the data. This offloads some data copy processes to the storage array.
- Block Zeroing (WRITE-SAME), enables the storage array to zero-out a large number of blocks within the array without having to have the ESX Server write the zeros as data and thus speeds up provisioning of VMs. This offloads some of the file space zeroing functions to the storage array.
- Hardware Assisted Locking (ATS), provides an alternative to SCSI reservations as a means to protect the metadata for VMFS cluster file systems to improve the scalability of large ESX server farms sharing a common datastore.

Support for VMware VAAI functionality is available via installation of the 3PAR VAAI Plug-in 1.1.0 on ESX / ESXi 4.1 in combination with 3PAR InForm OS version 2.3.1 MU2 (minimum).

Refer to the site http://www.3par.com/services/download.html for further information on VMware VAAI, The 3PAR VAAI Plug-in for ESX / ESXi 4.1 installation package, and the 3PAR VAAI Plug-in 1.1.0 for VMware vSphere 4.1 User's Guide.

7 Allocating Storage for Access by the ESX Host

In this chapter

This chapter describes the basic procedures that are required to create and export virtual volumes so they can be utilized by the VMware ESX host. For complete details on creating and managing storage on the InServ Storage Server, consult the appropriate 3PAR documentation.

7.1 Creating Storage On the InServ Storage Server

This section describes the general recommendations and restrictions for exporting 3PAR InServ storage to an ESX host.

Consult the InForm Management Console help and the 3PAR Command Line Reference for complete details on creating virtual volumes for a given InForm OS. In addition, for a detailed discussion on using Thin Provisioned volumes and strategies for creating VLUNs, consult *VMware on 3PAR Utility Storage* located in DCS.

NOTE: To create thinly provisioned Virtual Volumes, a 3PAR Thin Provisioning license is required.

7.1.1 Creating Virtual Volumes for InForm OS 2.2.4 and Beyond

After devising a plan for allocating space for the ESX host, you need to create the required virtual volumes on the InServ Storage Server. Volumes can be fully provisioned from a CPG or can be thinly provisioned.

Using the InForm Management Console:

1. From the Menu bar, select:

   Actions > Provisioning>VV>Create Virtual Volume

2. Use the **Create Virtual Volume** wizard to create a base volume.
3. Select one of the following options from the **Provisioning** list:

- **Fully Provisioned from PDs**
- **Fully Provisioned from CPG**
- **Thinly Provisioned**

Using the CLI:

To create a fully provisioned or thinly provisioned virtual volume, issue the following CLI command:

```
createvv [options] <usr_CPG> <VV_name> [.<index>] <size>[g|G|t|T]
```

Here is an example:

```
# createvv -cnt 5 TESTLUNs 5G
```

7.1.2 Creating Virtual Volumes for InForm OS 2.2.3 and Earlier

When running InForm OS 2.2.3 and earlier, the `createaldvv` command can be used to create Virtual Volumes on the InServ Storage Server that can then be exported and discovered by the ESX host. Here is the general form of the command:

```
createaldvv [options] <vvname>[.<index>] <size>[g|G|t|T]
```

Here is an example:

```
# createaldvv -cnt 5 TESTLUNs 5G
```

This will create 5 virtual volumes of 5 Gig each in size, and fully provisioned from PDs.

NOTE: To create a fully provisioned or thinly provisioned virtual volume, consult the InForm CLI Reference.

7.2 Exporting LUNs to an ESX Host

This section explains how to export Virtual Volumes created on the InServ Storage Server as VLUNs for the ESX host.

When exporting VLUNS to the VMware ESX host, you should observer the following guidelines:

- New VLUNs that are exported while the host server is running will not be registered until a Bus Rescan is initiated. This may be performed from the ESX Server Management Interface (VI client / vSphere client).
- Disks can be added to a Virtual Machine with the Virtual Machine powered up. However, to remove a disk, the virtual machine must be powered off. This is a limitation of the ESX Server.
- The maximum number of LUNs on a single ESX HBA port is 256, and 256 total LUNs on the ESX Server. Internal devices, such as local hard drives and CD drives, are counted as a LUN in the ESX Server LUN count.
- VLUNs can be created with any LUN number in the range from 0 to 255 (VMware ESX Server limitation).
- iSCSI LUNs and FC LUNs are treated as any other LUN on the ESX Server. No special requirements or procedures are needed to use iSCSI LUNs.
- The ESX limitation for the largest LUN that can be utilized by the ESX server is 2047 GB.
- Sparse LUN numbering, (meaning LUN numbers can be skipped), is supported by VMware ESX Server. A LUN 0 is not required.

For failover support using the QLogic or Emulex drivers, virtual volumes (VVs) should be exported down multiple paths to the host server simultaneously. To facilitate this task, create a host definition on the InServ Storage Server that includes the WWNs of multiple HBA ports on the host server and export the VLUNs to that host definition.

It has been noted by 3PAR that provisioning several Virtual Machines (VMs) to a smaller number of large LUNs, versus a single VM per single LUN provides better overall results. Further examination and explanation of this recommendation is outlined in the document "VMware on 3PAR Utility Storage" published in 3PAR DCS.

Concerning TPVVs; ESX Server VMFS-3 does not write data to the entire volume at initialization and it can be used with Thin Provisioned Virtual Volumes without any configuration changes

to VMFS. A further examination of this subject, recommendations and limitations are explored in the 3PAR document VMware on 3PAR Utility Storage.

7.2.1 Creating a Virtual Logical Unit Number (VLUN) for Export

Creation of a Virtual Logical Unit Number (VLUN) template enables export of a Virtual Volume (VV) as a VLUN to one or more ESX hosts. There are four types of VLUN templates:

- port presents - created when only the `node:slot:port` are specified. The VLUN is visible to any initiator on the specified port.
- host set - created when a host set is specified. The VLUN is visible to the initiators of any host that is a member of the set.
- host sees - created when the `hostname` is specified. The VLUN is visible to the initiators with any of the host's World Wide Names (WWNs).
- matched set - created when both `hostname` and `node:slot:port` are specified. The VLUN is visible to initiators with the host's WWNs only on the specified port.

You have the option of exporting the LUNs through the Inform Management Console or the CLI.

Using the Inform Management Console:

1. From the Menu bar, select **Actions>Provisioning>VLUN>Create VLUN**.
2. Use the Export Virtual Volume dialog box to create a VLUN template.

Using the CLI:

To create a port presents VLUN template, issue the following command:

```
createvlun [options] <VV_name | VV_set> <LUN> <node:slot:port
```

To create a host sees or host set VLUN template, issue the following command:

```
createvlun [options] <VV_name | VV_set> <LUN> <host_name/set>
```

To create a matched set VLUN template, issue the following command:

```
createvlun [options] <VV_name | VV_set> <LUN> <node:slot:port>/
<host_name>
```

Here is an example:

```
createvlun -cnt 5 TESTLUNs.0 0 hostname/hostdefinition
```

Consult the InForm Management Console help and the 3PAR Command Line Reference for complete details on exporting volumes and available options for the InForm OS version that is being used on the InServ Storage Server. Please note that the commands and options available for creating a virtual volume may vary for earlier versions of the InForm OS.

7.3 Discovering LUNs on VMware ESX Hosts

This section provides tips for discovering LUNs that are being utilized by the ESX host.

Once the Server is built, the preferred method for configuring and managing the use of the ESX Server is through a VI / vSphere Client Management Interface and Virtual Center / vSphere Server.

New VLUNs that are exported while the host server is running will not be registered until a Bus Rescan is initiated. This may be performed from the ESX Server Management Interface (VI client / vSphere client). If failover support is utilized (recommended), view all LUNs and their respective paths using the menu from the ESX Server Configuration storage Interface tabs.

Disks can be added to a Virtual Machine with the Virtual Machine powered up. However, to remove a disk, the virtual machine must be powered off. This is a limitation of the ESX Server.

The maximum number of LUNs on a single ESX HBA port is 256, and 256 total LUNs on the ESX Server. Internal devices, such as local hard drives and CD drives, are counted as a LUN in the ESX Server LUN count.

7.4 Removing Volumes

After removing a vLUN exported from the InServ, perform a ESX host bus adapter Rescan. ESX will update the disk inventory upon rescan. This applies to both Fibre Channel and iSCSI.

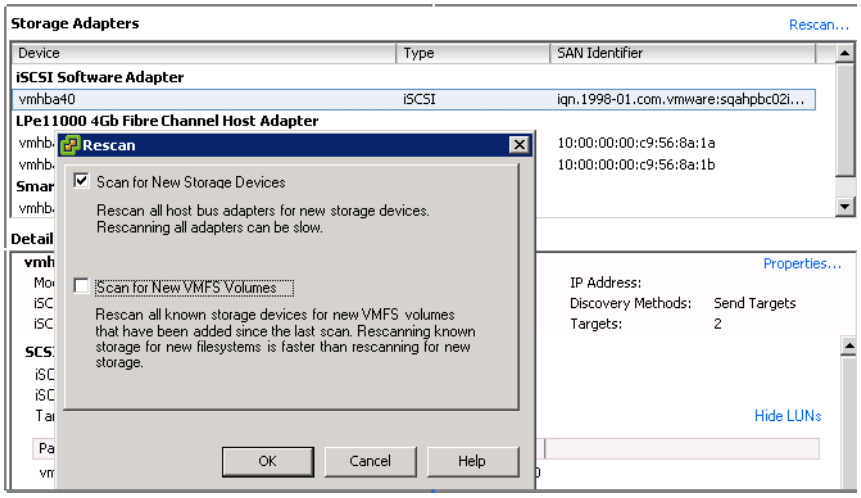

It is advisable to remove the disk/LUN from the Virtual Machine (VM) inventory before removing it from InServ and ESX inventory.

8 Booting from the InServ Storage Server

In this chapter

This chapter provides a general overview of the procedures that are required to boot the VMware ESX Operating System from the SAN.

8.1 Booting the VMware ESX Server from the InServ Storage Server

In a boot-from-SAN environment, each ESX server's operating system is installed on a the InServ Storage Server, rather than on the host's internal disk. In this situation, you should create a separate virtual volume for each ESX server to be used for the boot image.

Here are the general steps in this process:

- Boot the HBA BIOS
- Enable the HBA port to boot
- Perform the required zoning
- Create a virtual volume and export it as a VLUN to the ESX host
- Discover the LUN and designate it as bootable through the HBA BIOS

For detailed information, consult the VMware *Fibre Channel SAN Configuration Guide.*

The VMware ESX Server has an option that allows the VMware Base OS to be installed and booted from a SAN or InServ Storage Server virtual storage device. You can choose this option during the initial installation phase of the VMware Server Installation. Please refer to the VMware documentation for further information regarding 'SANboot'.

3PAR makes the following general recommendations for preparing the host HBAs for a SAN boot deployment:

NOTE: The NVRAM settings on HBAs can be changed by any server in which they are installed. These settings will persist in the HBA even after it is removed from a server. To obtain the correct settings for this configuration, you must return all NVRAM settings to their default values.

1 After installation of the HBAs, reset all of the HBA NVRAM settings to their default values.

NOTE: Each HBA port is reported as a host bus adapter and the HBA settings should be set to default.

2 Enter the HBA program during Server boot by pressing **ALT+Q** or **CTRL+Q** for QLogic HBAs or **ALT+E** for Emulex HBAs. Change and save all HBA settings to their default settings.

NOTE: When using a McData Fabric, set the HBA topology to 'point to point.

3 Reboot the host computer.

Revision History

Release level	Revision summary
320-200195 Rev A December 2009	First release of this Implementation Guide with content restructuring for improved access and coverage of VMware ESX 3.X - 4.X connected to InServ Storage Servers running all currently supported InForm OS versions.
320-200195 Rev B December 2009	Updated screen shot.
320-200195 Rev C September 2010	Updated Chapter 3 with setup information for InForm OS 2.2.x and 2.3.x. Added new Chapter 5 for FCoE implementations. Added information for VAAI support.