

InForm OS 2.2.3/2.2.4 IBM AIX 6.1 Implementation Guide

Topics in this Guide

1.0 Notices	2
2.0 Advisories	2
3.0 Introduction	3
4.0 Supported Configurations	3
5.0 InForm OS Upgrade Considerations	3
6.0 Port Configuration Quick Reference	4
7.0 Host Server Setup	4
8.0 InServ Storage Server Setup	6
9.0 Fabric Setup	7
10.0 InServ Storage Server Usage	15
11.0 Host Usage	15

1.0 Notices

For Printed Material

Copyright © 2008 3PAR Inc. All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, without the prior written consent of 3PAR Inc., 4209 Technology Drive, Fremont, CA 94538.

For Electronic Material

The user may print one copy for personal use. Production of multiple copies or copies for sale is strictly prohibited.

Trademarks

3PAR, InServ, InForm, InSpire and Serving Information are registered trademarks of 3PAR Inc. All other trademarks and registered trademarks are the property of their respective owners.

2.0 Advisories

The following alerts appear throughout this guide:



NOTE: Notes are reminders, tips, or suggestions that supplement the information 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.

3.0 Introduction

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, this information must be used 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.

4.0 Supported Configurations

Supported configurations are listed in the InForm OS 2.2.4 Configuration Matrix. Supported configurations must follow the installation and configuration steps described in the relevant Implementation Guide(s).

5.0 InForm OS Upgrade Considerations

Refer to the InForm OS Upgrade Pre-Planning Guide (PN 320-1307) for information and planning of an online InForm OS upgrade.

6.0 Port Configuration Quick Reference

Table 1-1. Port Configuration

InServ Storage Server Ports: Configure InServ Storage Server ports to port personality 14 for Fabric configurations.

```
# controlport persona 14 <X:X:X>
```

Configure InServ Storage Server ports to port personality 13 for Direct Connect configurations.

```
# controlport persona 13 <X:X:X>
# controlport vcn disable -f <X:X:X>
```

Disabling VCN prevents port drop when a new LUN is exported or removed; optional.

6.1 Host Ports

No configuration is needed.



NOTE: The host server ports will not be discovered on the InServ Storage Server until the `cfgmgr` command is run on the host server.

7.0 Host Server Setup

This document covers configurations using 3PAR MPIO for AIX for path failure management.

For configurations using 3PAR MPIO for AIX, complete the HBA and Driver Installation and Configuration and the 3PAR MPIO for AIX Installation and Setup.

7.1 HBA and Driver Installation and Configuration

For HBA installation instructions, driver support and usage guidelines, refer to the IBM Installation and Usage Guide for each product type. The required drivers are located on Volume 1 of the IBM Base Installation CDs, and are supplied with the hardware kit from IBM.

After the installation of the host HBAs, power up the AIX host server.

- 1 To display the driver versions and ROS levels of each host HBA type on the host use the `lslpp -l <filesetname>` command to display the AIX OS driver set filenames for the FC HBAs.

IBM FC 5716 filesetname	devices.pci.df1000fa.com
IBM FC 575X filesetname	devices.pci.df1000fd.com

- 2 To display the microcode levels on the FC HBAs, use the following command:

```
# lscfg -vp -l fcs*
DEVICE      LOCATION      DESCRIPTION
fcs0        U7879.001.DQD0XD8-P1-C3-T1  FC Adapter
.
.
.
ROSROS Level and ID.....02881914
```

7.2 3PAR MPIO for AIX Installation and Setup

This procedure applies to either a new installation or an existing installation where InServ Storage Servers Virtual Volumes (VVs) already exist on an AIX 6.1 system. This installation must be performed by a user logged into the AIX system as the superuser. Installation of the 3PAR MPIO requires a system reboot to become effective. The 3PAR MPIO software will create an Active/Active fail-over environment to the 3PAR Storage Devices. SMIT MPIO can be used to configure or manage the MPIO environment if settings other than the default active/active environment are desired. The 3PAR MPIO package is not to be used in conjunction with VERITAS Volume Manager.

- 1 Load the distribution CD containing the 3PAR MPIO for IBM AIX into the CD drive.
- 2 Use `smit update_all` to install the 3PAR MPIO for IBM AIX from the distribution CD.
- 3 Restart the AIX Server.

Please refer to the document for 3PAR MPIO v2.2 for IBM AIX for upgrade or removal instructions or for additional information.



CAUTION: When removing a host port from an InServ Storage Server and removing a host definition using the InForm CLI `removehost -rvl` command, remove the MPIO host definition from the AIX host first by using the Smit MPIO remove path before issuing the `removehost -rvl` command. Failure to follow this order may result in I/O hang conditions.

8.0 InServ Storage Server Setup



REQUIRED: This step must be completed before connecting the InServ Storage Server port to a device.

8.1 Direct Connect

- 1 Set each storage server port that will be directly connected to the InServ Storage Server to port persona 13.

```
# controlport persona 13 <X:X:X>
```

To prevent port drop from occurring on LUN export or removal, use the following InForm CLI command:

```
controlport vcn disable -f <X:X:X>
```

where `<X:X:X>` is the port location, expressed as `node:slot:port`.



NOTE: Port drop only occurs if no file systems have been created on InServ Storage Server LUNs, or LUNs are being exported for the first time to an AIX host. If this value is not set, simply issuing `cfgmgr` from the host will cause the host ports to show up again on InServ Storage Server. This setting is optional.

- 2 Verify that each port has the appropriate personality defined:

```
# showport -par
N:S:P ConnType CfgRate Class2 VCN -----Persona----- IntCoal
4:0:1 loop auto disable enable (13) g_ven, g_hba, aix, 0, DC enabled
```

- 3 Connect the InServ Storage Server ports to the FC HBAs for a direct connect configuration.

- 4 Run the `cfgmgr` command on the host server.
- 5 Verify the cabling for correctness using the InForm CLI command `showhost`. Each host HBA port's WWN should appear associated with the storage server port to which it is connected.

```
# showhost
2   sqa-03-EMX           10000000C939867D   4:0:1
10000000C939867E   5:0:1
```

8.2 Fabric Connect

- 1 Set each storage server port that will connect to a Brocade, McData, QLogic, or Cisco fabric to port personality 14.

```
# controlport persona 14 <X:X:X>
```

Where `<X:X:X>` is the port location, expressed as node:slot:port.

- 2 Verify each port has the appropriate personality defined:

```
# showport -par
N:S:P ConnType CfgRate Class2 VCN -----Persona----- IntCoal
4:0:2 point auto disable enable (14) g_ven, g_hba, aix, 0, FA enabled
```

9.0 Fabric Setup

Brocade switch ports that will connect to a host server HBA port or to an InServ Storage Server port should be set in their default mode. On Brocade 3xxx switches running Brocade firmware 3.0.2 or later, verify that each switch port is in the correct mode using 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 should be in their default modes as type "GX-Port", with a speed setting of "Negotiate".
- Cisco switch ports that connect to InServ Storage Server ports or host HBA ports should be set to AdminMode = FX and AdminSpeed = auto port, with the speed set to auto negotiate.



NOTE: The AIX Host setting to enable dynamic tracking is recommended for all Fabric.

- 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".

9.1 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 (RSCNs) that are irrelevant to 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 reconnection 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.

9.1.1 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 1-1).

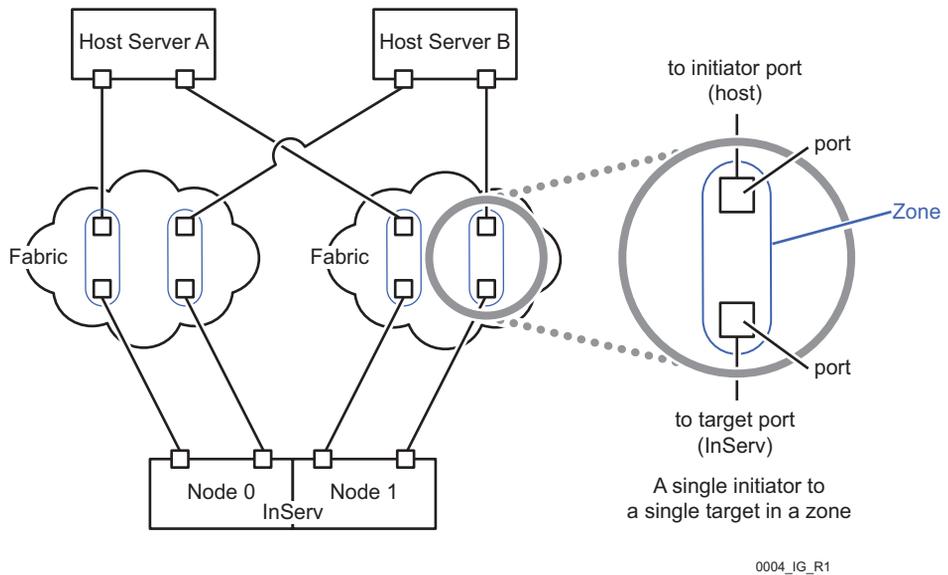
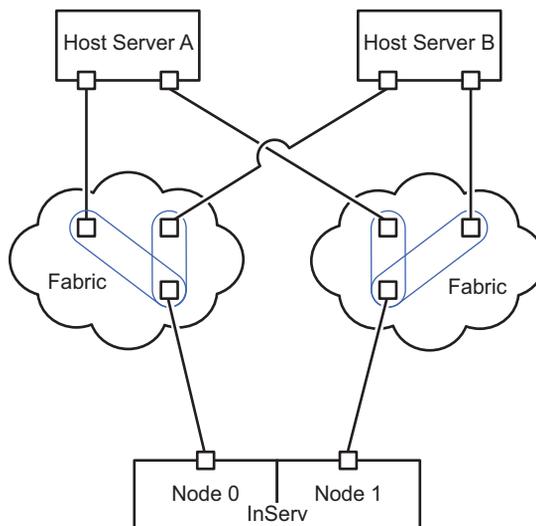


Figure 1-1. Single Initiator to Single Target Zoning No Fan-In/No Fan-Out

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

Fan-out refers to an InServ Storage Server port that is connected to more than one host port, as shown in [Figure 1-2](#).



0003_IG_R1

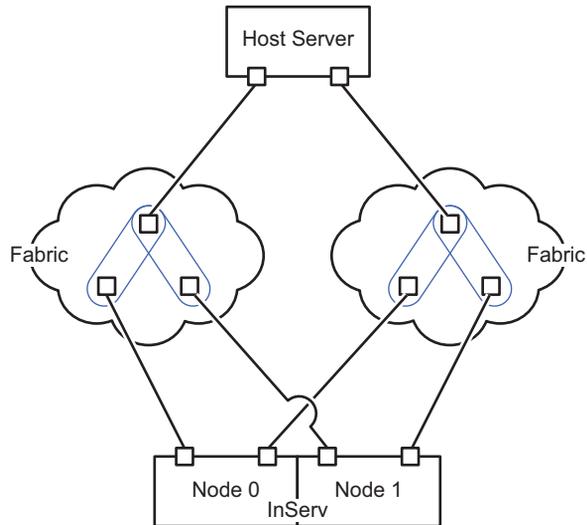
Figure 1-2. Single Initiator to Single Target Zoning with Fan-Out



NOTE: A maximum of 64 host server ports can fan-out from a single InServ Storage Server port.

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

Fan-in refers to a host server port connected to many InServ Storage Server ports. This is shown in [Figure 1-3](#).



0002_IG_R1

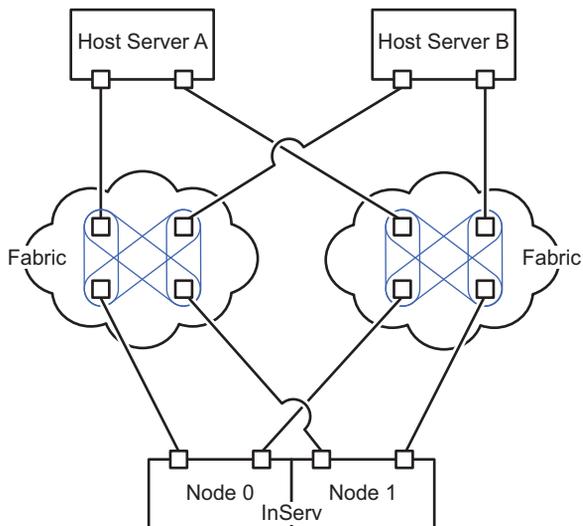
Figure 1-3. Single Initiator to Single Host Target Zoning with Fan-In



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

9.1.4 Single Initiator to Single Target Zoning with Mixed Fan-In and Fan-Out Configurations

The following figure (Figure 1-4) shows a single initiator to a single target zoning with fan-in and fan-out from one InServ Storage Server to multiple host servers.



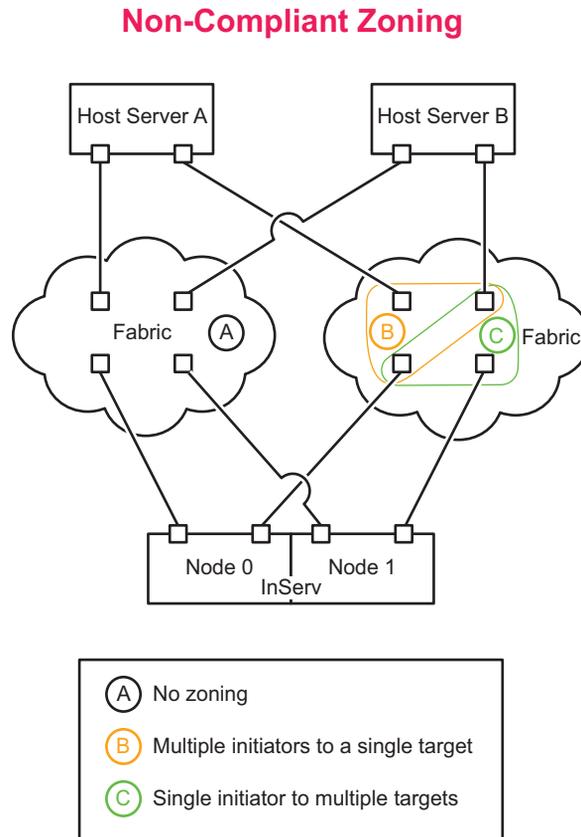
0005_IG_R1

Figure 1-4. Single Initiator to Single Target Zoning with Fan-In and Fan-Out

9.1.5 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 1-5](#).



0001_IG_R1

Figure 1-5. Non-Compliant Zoning

9.2 General Note on Fabrics

When host server ports can access multiple targets on fabric zones, the assigned target number (which is assigned by the host driver) for each discovered target can change when the host server is booted and some targets are not present in the zone. This may change the device node access point for devices during a host server reboot. This issue can occur with any fabric-connected storage, and is not specific to the InServ Storage Server.

9.2.1 Target Port Limitations

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

- Maximum of 64 host server ports per InServ Storage Server port, with a maximum total of 1,024 host server 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 amongst 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.

10.0 InServ Storage Server Usage

Follow the instructions for creating Virtual Volumes (VVs) and Virtual LUNs (VLUNs) in the *InForm OS CLI Administrator's Manual* while adhering to these cautions and guidelines:

- This configuration supports sparse LUNs (meaning that LUNs may be skipped). LUNs may also be exported in non-ascending order (e.g. 0, 5, 7, 3).
- Even though the InServ Storage Server supports the exportation of VLUNs with LUNs in the range from 0 to 65535, only create VLUNs with LUNs in the range from 0 to 511. Testing at 3PAR was limited to 512 LUNs per storage server port (target) using LUN numbers 0 to 511.
- Exported VLUNs will not be registered on the host until `cfgmgr` is run on the host.
- The maximum LUN size that is supported by the AIX host is 2TB.

11.0 Host Usage

When VLUNs are exported, they will not automatically appear in the host's host device list. After a new VLUN is exported from a storage server port it will not appear on the host until a rescan is issued from the AIX system using the host command `cfgmgr`.

To check for devices after a rescan is issued, use `lsdev -Cc disk`:

```
# lsdev -Cc disk

hdisk1 Available 06-08-02      3PAR InServ Virtual Volume
hdisk2 Available 06-08-02      3PAR InServ Virtual Volume
hdisk3 Available 06-08-02      3PAR InServ Virtual Volume
hdisk4 Available 06-08-02      3PAR InServ Virtual Volume
hdisk5 Available 06-08-02      3PAR InServ Virtual Volume
```

When removing InServ Storage Server VVs from the AIX host, it is good practice to first remove the `hdisk` definition from the AIX host using the `rmdev -l hdiskN -d` command, then remove the VV from the host on server. Performing a clean removal in this fashion insures the `hdisk` entry is removed from the AIX device data base so that if another LUN is exported in the future with the same LUN and characteristics, a device mismatch does not occur on the AIX host.

SMIT MPIO can be used to configure or manage the MPIO environment if settings other than the default active/active environment are desired.

11.1 Dynamic Tracking and Fast_Fail

When dynamic tracking of FC Devices is enabled, the FC adapter driver can detect when the Fiber Channel N_Port ID of a device changes and re-route traffic destined for that device to the new address while the devices are still online. Examples of events that can cause an N_Port ID to change are moving a cable between a switch and storage device from one switch port to another, connecting two separate switches via an Inter-Switch Link (ISL), and possibly rebooting a switch.

- Dynamic tracking set to enabled is recommended for all Fabrics. All testing at 3PAR was performed with Dynamic Tracking enabled for all fabrics types.
- Dynamic tracking can be enabled using 'SMIT Devices'.
- FC Adapter.
- FC SCSI I/O Controller Protocol Device.
- Change / Show Characteristics of a FC SCSI Protocol Device.
- Dynamic Tracking of FC Devices 'yes'.

It is also recommended that the value for 'fast_fail' be enabled rather than the default value of 'delayed_fail'. This speeds up recovery time in the event of a path failure. This value can be changed at the same location that dynamic tracking is set.

11.2 High Availability Cluster Multi-Processing for AIX (HACMP)

3PAR supports HACMP v5.4.1 when using 3PAR MPIO 2.2 for IBM AIX when using enhanced concurrent volume groups in multi-host environments. The use of non-concurrent volume groups is not supported. There are no other special considerations for using HACMP. Refer to the *3PAR Multipath I/O 2.2 for IBM AIX User's Guide* for additional information. Refer to IBM HACMP documentation for HACMP planning, setup and usage. These IBM HACMP documents are of particular importance:

- Concepts and Facilities Guide
- Planning and Installation Guide
- Administration Guide