STORAGETEK®

PART NUMBER
96107

EDITION NUMBER
4

FLEXLINE™

ONLINE STORAGE SYSTEMS

INSTALLATION INSTRUCTIONS FOR AIX

Open Systems Host Attach

PRODUCT TYPE
SOLUTIONS

FlexLine™ 200/300 Series

Installation Instructions for AIX
open systems host attach

Information Control

Should this publication be found, please return it to StorageTek, One StorageTek Drive, Louisville, CO 80028-5214, USA. Postage is guaranteed.

Information in this file is subject to change. In the event of changes, the information will be revised. Comments concerning its contents should be directed to:

FlexLine Document Manager
Storage Technology Corporation
One StorageTek Drive
Louisville, CO 80028-2129

or send email to glsfs@stortek.com

Export Destination Control Statement

These commodities, technology or software were exported from the United States in accordance with the Export Administration Regulations. Diversion contrary to U.S. law is prohibited.

Proprietary Information Statement

Copyright Statement

Summary of Changes

The following is the history and summary of changes for this publication.

Date	Edition	Description
May 2003	First	Initial Release
Nov 2003	Second	Updates to commands and driver information. Also added a CRC webpage to obtain a permanent license for failover software.
Jun 2004	Third	Updates to failover instructions and commands
Oct 2004	Fourth	Updates to reflect the new FlexLine terminology

This page intentionally left blank.

AIX - Installation Instructions

1

The following pages help ensure that you order and configure a viable solution for your RS6000 operating environment running AIX 4.3.3, 5.1, or 5.2. Customers and StorageTek personnel should also be familiar with the following documents for your particular FlexLine storage solution.

- Storage system Site Preparation Guide
- SANtricity Storage Manager Concepts Guide
- Storage system Planning Checklist
- Preconfiguration and Maintenance Setup Instructions

Introduction

This manual is organized to help plan your RS6000 AIX configuration, provide order information, and help to install and attach your storage system to the RS6000 host server running AIX. The following should be used in conjunction with the Siebel configurator.

Plan your site configuration

Prior to placing an order for StorageTek storage equipment, one should make sure of the system needs. First measure the locations of each current and future host, storage system, and other network components. Fill in this form with the requested information.

✔	Parameter	Outcome
	AIX 4.3.3, 5.1, or 5.2 operating level	
	Distance between host and storage system	
	Distance between other devices	
	Number of HBAs needed (two are recommended)	
	Type and number of storage systems being installed	
	Number of FC cables needed	

Select Host Adapter (s)

Pick the type and quantity of the fibre channel PCI Host Bus Adapters (HBAs) needed for your host systems.

Recommended:

Quantity	Description	Rating
	Cambex PC2000 (Failover and Non-failover)	200 MB/sec
	IBM 6228 and 6239 (Failover and Non-failover)	200 MB/sec

May also be used:

Quantity	Description	Rating
	Cambex PC1000 (Failover and Non-failover)	100 MB/sec
	IBM 6227 (Failover and Non-failover)	100 MB/sec

Select the Fibre Channel Cables

Select the length of fiber optic cable and the quantity for connection to the host systems. The following are examples of optical cables for the FlexLine storage systems that may be ordered through StorageTek. Write in the quantities needed.

Quantity	2 Gb FC Optical Cable Description	Part #
	10 meters (32.8 feet) LC to LC Duplex	10800221
	50 meters (164 feet) LC to LC Duplex	10800222
	100 meters (328 feet) LC to LC Duplex	10800223
Note: Other lengths are available. Consult the Siebel Configurator.		

There are also LC to SC cables that would allow one to use a 2 Gb HBA with lower speed switches. The following cables are available in this event.

Quantity	FC Optical Cable Description	Part #
	10 meters (32.8 feet) LC to SC	10800231
	50 meters (164 feet) LC to SC	10800232
	100 meters (328 feet) LC to SC	10800233

If installing through a switch, you may be installing a set of the following.

Quantity	1 Gb FC Optical Cable Description	Part #
	10 meters (32.8 feet) SC to SC Duplex	10800134
	50 meters (164 feet) SC to SC Duplex	10800137
	100 meters (328 feet) SC to SC Duplex	10800138
Note: Other lengths are available. Consult the Siebel Configurator.		

Check your quantities against your earlier distance measurements to ensure your order includes all cable attachments. Fibre Channel cables attach to the HBA, and the other end, typically, mates with the control module's host interface.

Note: Unless there is an existing management station, cable connections to the management station should be via 10BaseT Ethernet.

Fibre Channel short wave optical cables may also be referred to as FC optical (LC to LC Duplex) cables going from the host (or switch) to the host minihub.

Place your Order

Make sure that you have the following.

✔	Description
	IBM RS/6000 Server running AIX 4.3.3, 5.1, or 5.2
	HBA(s)
	Fibre Channel cables to configure your storage system

When you make a new order, a storage management software/documentation CD-ROM along with a CD insert, commonly called a Quick Start Guide will be in your host attach kit that matches your storage system and OS type. The Quick Start Guide will walk you through your installation—first referring you to the Preconfiguration and Maintenance Setup Instructions, P/N 95966, and then to any special instructions written to support your particular operating system installation. In addition depending on your architecture, you may chose to order the Cambex HBA driver CD as well as the Cambex DPF failover driver CD. If the Cambex DPF failover driver is used, a license will have to be purchased for each host that it is installed on.

Note: If you need information about installing HBAs, drivers, or cables; refer to the Client Server Disk installation supplement, P/N 96015, or to your host installation manuals.

Installation

Once you have ordered your hardware and software, read the following information on how to attach your FlexLine storage systems solution to an IBM RS/6000 server running AIX 4.3.3, 5.1, or 5.2

Before You Begin

The materials needed for the installation are listed below. Check to see that you have each of these items. The following pages provide more information about these materials (see headings) with specific steps and details on what to do.

In most cases, actual commands have been provided to help you with your installation, or help you configure or modify settings within the host or on your particular FlexLine™ 200/300 Series storage system solution (e.g. FLX210, FLX240, or FLX280).

Before you begin, make sure you have these items:

1. An IBM RS/6000 Server running AIX 4.3.3, 5.1, or 5.2
2. One or more HBA(s)
3. SFP Transceivers for the storage system
4. Fibre Channel cable(s) that will connect the host to the storage system
5. A fully configured storage system.

Note: In the following sections we will refer to fcs0 and fcs1 to identify the HBAs. The name and numbers on your host depend on the type and how many HBAs you have in your host. For Cambex HBAs they are identified as fscsi(n). For IBM HBAs they are identified as fcs(n). For the purposes of examples we will use fcs(n). If you are using Cambex HBAs, replace fcs with fscsi in all commands. If you have more than two HBAs in your host, use the lower numbered HBA/fcs# for controller A and the higher one for controller B.

Storage Systems and Storage Manager

StorageTek's FlexLine™ 200/300 Series storage systems provide access to high-end storage using Fibre Channel and RAID technology. A smorgasbord of hosts and applications are available to help ensure the very best in data access and acquisition - no matter how complex your network configuration.

The storage management documentation introduces you to the latest in StorageTek's ongoing enhancements in FlexLine technology, as well as the enhancements that your storage system can provide to your RS6000 AIX operating system.

In addition to this manual, both hardware and software information is available on the CRC and (as backup) on the following engineering web sites.

http://gandalf.stortek.com/ctp/index.html

http://gandalf.stortek.com/ctp/manager/index.html

http://gandalf.stortek.com/ctp/installation/index.html

http://gandalf.stortek.com/ctp/cabinet/index.html

The CRC will also have NVSRAM, controller and ESM firmware, drive code, and various other files that may be downloaded for upgrade purposes.

Note: Documents are also available on *Documents on CD*.

Make sure that your minihub speed is set to correspond with your HBA and switch speed. Also make sure you run your fibre channel cables to separate controllers on the storage system. You may run FC cables through a switch or directly to the controllers. The following photograph shows the typical FlexLine fibre channel host inputs.

Note: SFP Transceivers are required to make all host interface connections.

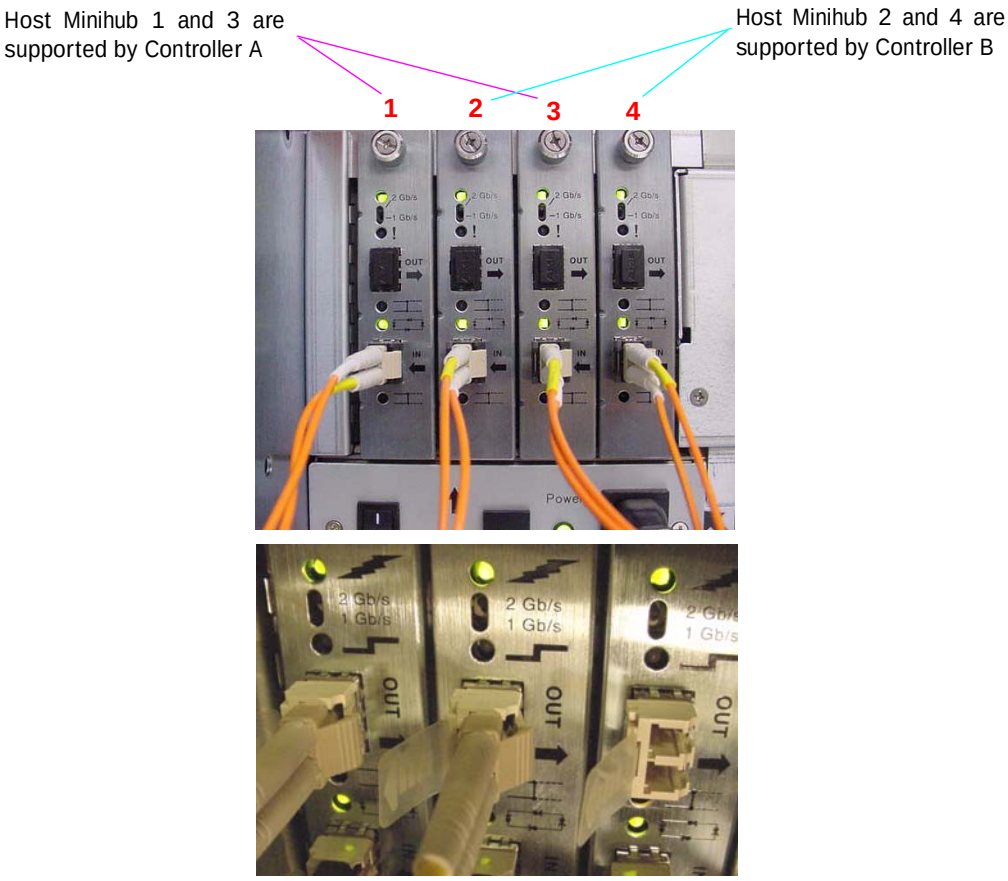

Note: The drive interface minihubs for high-end FlexLine are connected to both controller A and controller B. There is no need to ensure back-end fibre channel loops originate from any one host minihub.

The HBAs in the AIX host are referred to as fcs0 and fcs1. Once the HBAs are installed on the host, run the `lsdev` command:

```
lsdev -Ccadapter
```

The above command returns a list of adapters, some of which will be fcs(n) for the IBM HBAs, or fscsi(n) for Cambex HBAs. For each of the HBA entries (with the appropriate value of (n)), run the following commands:

```
lscfg -vl fcs(x)

lscfg -vl fcs(y)
```

Record the FCS number and the Network Address field (wwn) of the HBA from the output of the commands. This network address field is the value to be used for the Host Port Identifier field in the storage management software administration when administering the mapping administration for the volumes that this host will see. Check the Z9 field on the IBM HBAs for the correct version of microcode. Refer to “HBA WWN and other parameters” on page 26 for the appropriate versions of microcode.

Determine which HBA is fcs0, and label the cable and HBA with fcs0 / controller A. Determine which HBA is fcs1, and label the cable and HBA with fcs1 / controller B. You want to connect the cable from fcs0 to controller A or the switch zone for fcs0/controller A. You want to connect the cable from fcs1 to controller B or the switch zone for fcs1/controller B. The determination of which HBA is fcs0 or fcs1 may require cable pulls and rescan for devices later on in the process.

Prior to the install/upgrade on the host, storage management operations will need to be done. If this is a new install, one of the hosts in the network needs to have the storage management software installed. If you already have the software installed on an existing configured host, make sure it is up-to-date with the latest version of software. If no host has been configured, you can attach your laptop running SMclient 8.4x or 9.x (or later) to do the storage management procedures. This will enable you to assign drives, LUNs, and volumes on the array through a TCP/IP connection. You may also chose to install the storage management software on the AIX host and configure it to use a TCP/IP connection to configure the array before doing the AIX host installation/upgrade.

If the management software is going to be installed/upgraded on the AIX host, use the provided storage management software/document CD to:

- get the storage manager installation guide (inside the CD’s Docs folder)
- perform the install/upgrade of the storage management software on the AIX host.

The above activities need to be done prior to going to the next section.

Storage System Configuration

Refer to (for example) chapter one of the *BladeStore Storage Manager Installation Guide* P/N 96098 for doing the various storage management steps. Upgrade information will be added to the CRC or engineering web pages to provide assistance.

The following is preconfiguration information for a new installation. Information for an existing storage system is provided on page 9.

Note: For all installations/upgrades where software components are being changed, the following procedures must be done prior to any software modifications. This is to ensure a clean removal/install of software components, which may fail if devices are being accessed or viewed.

- stop access to storage devices
- removal of devices on the AIX host (not the storage array)
- disconnection of cables, as appropriate

Preconfiguring storage systems on new installations

The user needs to determine whether or not they are going to use failover. There are separate NVSRAM host selections for enabled and disabled failover. The differences are that in the non-failover NVS host setting, the unowned LUNs are only presented from the owning controller. The failover setting presents all LUNs down both paths, and depends on the HBA failover driver to determine the path presented to the O/S.

The standard method of preconfiguring the storage system should be followed prior to attaching to the host. Make sure that the FC cables are NOT connected to the HBAs. This is to ensure that no disks are visible until the entire software configuration is completed. When the software configuration procedure is completed, then the cables will be connected and devices scanned for with the `cfgmgr` command.

1. Run the following commands.
   ```
   rmdev -Rdl fcs0
   rmdev -Rdl fcs1
   ```
   Make sure that all devices were removed and no error messages exist about devices being busy or accessed. If there are errors, disconnect cables and reboot.

2. Disconnect the FC cables from the HBAs.

3. Using SMclient 8.4x on the AIX host or any attached management station with SM 8.4x or higher attached to the ethernet ports, download the correct type and latest version of firmware and NVSRAM (if needed) to both controllers. Download the firmware first, followed by the NVSRAM.

Note: SANtricity Storage Manager version 9.1 is available, but it does not provide a failover option for AIX. It will be avaialble once testing is complete.

Controller Firmware version (example):

BC84	FW_05460700_05400400.dlp or higher

Controller NVSRAM version (example):

BC84	N4884-5460855-002.dlp or higher

4. Build the RAID volume groups, volumes (LUNs), mappings for host groups, hosts, host ports, and volume mappings into host groups and LUN numbers. When performing these procedures, the following issues apply.

   - When creating the volumes, make a table of what controller (preferred path) is assigned to each volume (LUN number) for later reference when you set up the mappings for a volume to a host port. It is **CRITICAL** that the administration matches the physical connectivity between the fcs# and the controller that it is connected to, otherwise the host may not see the LUNs.
   - When setting the host type field:
     - Select host type AIX for non-failover version (presents LUNs down owning controller path only).
     - Select host type AIX_FO for failover version (presents LUNs down both paths).
     - Select the default host type if no partitioning installed.

5. An architecture decision concerning the access volume/UTM LUN (LUN 31 by default) for each host group accessed by an AIX host needs to be made.
   The issues are:

   - Is there an IBM storage array connected in the architecture? If there is, the access volume must not be mapped, thereby not visible to the host. If it is visible, it will corrupt the access to the IBM array as well as the StorageTek array.
   - If SMagent is not going to be used to do inband storage management (i.e. will only use TCP/IP connection for storage management operations), then the access volume must not be mapped, thereby not visible to the host. If it is visible, it may corrupt the access to the StorageTek array depending on the next bullet item.
   - Be aware that the access volume/UTM LUN cannot be used if the AIX packages `devices.fcp.disk.array.*` are installed to support IBM storage arrays. If these software packages are installed (by default they are installed and may be reinstalled by software patches)

and the access volume is visible, it will corrupt access to the StorageTek array. Check if the packages are installed by running the command:

```
lslpp -L devices.fcp.disk.array.*
```

If there is no IBM storage array connected in the architecture, we recommend removing the AIX packages as follows.

```
installp -u devices.fcp.disk.array.*
```

This will allow you to map and use the access LUN for inband storage management operations if you chose. If AIX software patches are applied, check if the packages were reinstalled and remove them if they were.

Upgrading an Existing Configuration

If upgrading an existing storage system that is installed and configured on AIX, or upgrading software on the AIX host; perform the following actions to remove existing devices on the AIX host prior to doing any software modifications. A complete backup of data should be performed prior to this procedure.

If any of the following steps do not complete and indicate device busy, do not proceed until the application that is using the device is stopped. For example you may need to do an SMagent stop command to stop usage of the access LUNs, stop file system access by applications, or cd out of a mounted file system. All devices must be successfully unmounted, varied off, and removed prior to continuing with software upgrades, otherwise major host corruption will occur.

1. Unmount any mounted file systems using the `umount` command or `smit`
2. Vary offline any volume groups using `varyoffvg` or `smit`
3. If using inband storage management, run the following command.

   ```
   SMagent stop
   ```

4. If using Cambex DPF failover driver, run the following command.

   ```
   stopsrc -s cbxdpfar
   ```

5. Determine the HBAs being used by using the `lsdev –Cc adapter` command.
6. Remove the adapters and scsi devices using the `rmdev –Rdl fcs(n)` `for IBM HBAs, or rmdev –Rdl fscsi(n) for Cambex HBAs`

   Example for an IBM HBA

```
# rmdev –Rdl fcs4
hdisk1  deleted
hdisk2  deleted
hdisk3  deleted
fscsi4  deleted
fcs4  deleted
```

Example for a Cambex HBA

```
# rmdev –Rdl fscsi5
hdisk4  deleted
hdisk5  deleted
fscsi5  deleted
```

7. Disconnect FC cables from HBAs
8. If using inband storage management, run the following command.

   ```
   SMagent start
   ```

9. If using Cambex DPF failover driver, run the following command.

   ```
   startsrc -s cbxdpfar
   ```

10. An architecture decision concerning the access volume/UTM LUN (LUN 31 by default) for each host group on the storage system accessed by the AIX host needs to be made. The mapping of the access volume on the storage system or the removal of host software that conflicts with it are discussed in the following issues.

    - Is there an IBM storage array connected in the architecture? If there is, the access volume must not be mapped, thereby not visible to the host. If it is visible, it will corrupt the access to the IBM array as well as the StorageTek array.
    - If SMagent is not going to be used to do inband storage management (i.e. will only use TCP/IP connection for storage management operations), then the access volume must not be mapped, thereby not visible to the host. If it is visible, it may corrupt the access to the StorageTek array depending on the next bullet item.
    - Be aware that the access volume/UTM LUN cannot be used if the AIX packages `devices.fcp.disk.array.*` are installed to support IBM storage arrays. If these software packages are installed (by default they are installed and may be reinstalled by software patches) and the access volume is visible, it will corrupt access to the StorageTek array. Check if the packages are installed by running the

command:
`lslpp -L devices.fcp.disk.array.*`

If there is no IBM storage array connected in the architecture, we recommend removing the AIX packages at the appropriate time in the upgrade process as follows.

```
installp -u devices.fcp.disk.array.*
```

Note: Do not do the command at this time because it needs to be done in the proper sequence of the software removal process.

This will allow you to map and use the access LUN for inband storage management operations if you chose. If AIX software patches are applied, check if the packages were reinstalled and remove them if they were.

AIX Host Installation Preparation

The following sections provide the necessary configuration, hardware and software requirements, and the things you should know—to prepare and install FlexLine™ 200 Series and 300 Series storage systems. Please read this information to minimize installation problems.

Hardware

Check that you have the appropriate HBA for your planned configuration:

- Cambex PC2000 (or PC1000) or IBM 6227/6228/6239 (for both failover and non-failover configurations).

Software

Check to see that you have the appropriate host software levels, as indicated below.

For all AIX versions you should be at the latest maintenance level and latest emergency patches because various file system (jfs2), HACMP, and HBA device driver issues have been recently (and are still being) corrected.

Cambex HBA drivers – version 2.5.10.5 or higher (standard) for AIX 5.1 and 5.2, Cambex 2.0.10.5 for AIX 4.3.3.

Cambex DPF Failover drivers – version 1.5.3.33 package or higher for AIX 5.1 and 5.2; Cambex 1.0.3.33 package or higher for AIX 4.3.3

Note: The Cambex failover driver must be individually purchased for **each** host. Distribution without purchase is illegal.

SN4116/4108 Switch firmware – 2.4.1c or later

AIX Host Installation/Upgrade Procedures

Use the following procedures as appropriate after performing all preliminary and preparation procedures as previously discussed.

Note: If you are using HACMP, you must also be using the Cambex DPF failover driver since it is the DPF driver that detects the “all paths bad” condition and triggers a node shut down to force resource groups to another node in an attempt to correct for the failure condition.

Cambex HBA Driver Remove/Install/Upgrade

If you’re using the Cambex HBAs and your driver is out-of-date, or if you’re installing for the first time,

- check the current installed driver version with the following command

  `lslpp -L PC1000.driver.obj`

- deinstall the existing Cambex HBA driver (if out-of-date) using `installp –u PC1000.driver.obj`

- ftp from the CRC to the target AIX system and uncompress the HBA driver file if necessary, or mount the Cambex HBA Driver CD and cd to the directory that contains the HBA driver file.

- install the latest PC1000 Cambex HBA driver (that supports both the PC1000 and PC2000 HBAs) with the command:
  `installp –aXd filename.image all`

  where filename is the filename of the latest Cambex HBA driver for your version of AIX, for example:
  PC1000.fcp.2_0_10_5.image (AIX 4.3 version)
  PC1000.fcp.2_5_10_5.image (AIX 5.1/5.2 version)

Adding odm data base device definitions for a non-failover solution

If using the Cambex DPF failover solution, skip this section.

If you are **not** using the Cambex DPF failover package, this step adds entries to the odm device data base to support the storage system in a non-failover architecture.

1. `mkdir /usr/lpp/stkfc`
2. `cd /usr/lpp/stkfc`
3. ftp the file `stkODM_R4.2.tar (or higher)` from the CRC site into the current directory
4. untar the file:

```
tar xvf stkODM_R4.2.tar (or higher)
```

5. install the odm entries using:

   ```
   ./install_stkODM
   ```

6. If the earlier architecture decision concerning the access LUN mapping versus removal of software packages `devices.fcp.disk.array.*` was to remove the software packages, then remove the software packages now with:

   ```
   installp -u devices.fcp.disk.array.*
   ```

If this is a new install, a reboot is necessary at this time to pick up the ODM changes. Install/upgrade HBAs while the system is down for this reboot.

Remove/Install/Upgrading the Cambex DPF Failover Driver

If you are using the Cambex DPF failover solution, perform the following steps.

1. Uninstall the old Cambex DPF failover driver (if installed) with the following command.

   ```
   installp –u cbxdpf.driver.obj
   ```

2. If the earlier architecture decision concerning the access LUN mapping versus removal of software packages `devices.fcp.disk.array.*` was to remove the software packages, then remove the software packages now with:

   ```
   installp -u devices.fcp.disk.array.*
   ```

3. FTP from the CRC to the target AIX system and uncompress the driver file if necessary, or mount the Cambex Failover Driver CD and cd to the directory that contains the driver file,

   and install it using:

   ```
   installp –acd filename.image all
   ```

   where filename is the filename of the latest Cambex DPF failover driver for your version of AIX, for example:
   cbxdpf_1_0_3_33.image (AIX 4.3 version)
   cbxdpf_1_5_3_33.image (AIX 5.1/5.2 version)

4. The DPF license will need to be added at this time.

   Apply the license key.

   If you have the license key, run the following command.

   `smitty devices` >> Cambex Dynamic Path Management >> Manage Licensing >> Add License from the Keyboard.

If no license is available at install time, use the 30-day trial key:

```
cd /usr/lpp/cbxdpf

mv license.30day license

/etc/methods/cfgcbxdpf -K check
```

or `smitty devices` >> Cambex Dynamic Path Management >> Manage Licensing >> Check License Expiration.

There are two different licenses available for the DPF failover driver. One for Cambex HBAs and another for IBM HBAs.

To request a permanent license key.

For Cambex HBAs - E-Mail the output of `/usr/lpp/cbxdpf/dpfinfo` to support@cambex.com

For IBM HBAs - You will need to provide the output of `uname -a` and fill out the form located at:

http://www.support.storagetek.com/GlobalNavigation/Support/ToolsAndServices/Tools/swkeys/SoftwareKeys.htm

If you miss this license step or your 30-day license expires, you will lose access to all of your LUNs when you reboot or if you run the `cfgmgr` command to rescan for LUNs due to a configuration change. The error output from the cfgmgr command will be something like:

method error (/etc/methods/cfgcbxdpf -l cbx?)

for each LUN. The ? will be the number that the Cambex driver associates with the LUN.

5. Modify the Cambex Dynamic Path Management parameters as follows. This makes the autorecovery software detect failures and restore to normal operations quicker.

   smitty devices >> Cambex Dynamic Path Management >> Manage AutoRecovery Subsystem >> Change/Show AutoRecovery Subsystem Parameters. Change the parameters to the following:

   - Time interval to check paths - 10
   - Number of successful inquiries for unfail - 2
   - Refresh from ODM - 300

Installing ODM Entries for Access LUNs

This section creates the ODM entries that allow the SANtricity software to be able to use the access LUNs for inband storage management.

If the earlier architecture decision concerning the access LUN mapping versus removal of software packages `devices.fcp.disk.array.*` was to not remove the software packages and to remove the mapping of the access LUNs, thereby not being able to see the access LUNs, skip this section.

This section is run regardless of whether you're using a failover or non-failover solution.

1. `mkdir /usr/lpp/stkfc`
2. `cd /usr/lpp/stkfc`
3. ftp the file access.tar from the CRC site into the current directory.
4. untar the file.

   ```
   tar xvf access.tar
   ```

5. install the odm entries using:

   ```
   ./install_access
   ```

Note: If the Cambex failover driver is removed or upgraded, this procedure must be rerun because the Cambex removal process removes the above ODM entries.

Installing Compatibility Software

There are three installation packages (stk.svadisk, stk.renumber, stk.hacmp) available on the Cambex DPF Failover CD or on the CRC site in the file compatibility.tar . The next steps will transfer the packages onto the host.

1. `mkdir /usr/lpp/stkfc`
2. `cd /usr/lpp/stkfc`
3. ftp the file compatibility.tar from the CRC site or copy off of the CD into the current directory.
4. untar the file.

   ```
   tar xvf compatibility.tar
   ```

The stk.svadisk package changes the set of attributes of disks on SVAs (V2X/V2X2, V960, or 9500) to optimize the disk I/O performance. This package is required for the installation of the stk.renumber package.

The stk.renumber package resolves the ghost disk problem that occurs when multiple AIX servers see the same set of disks.

The stk.hacmp package helps notify the user when the Cambex DPF failover software detects FC error conditions and triggers HACMP node failover when the DPF driver detects the "all paths bad" to a volume(s).

The software packages need to be installed in the following order.

1. stk.svadisk (V1.1.1.0 or higher)
2. stk.renumber (V1.2.2.0 or higher)
3. stk.hacmp (V1.2.2.0 or higher) (if HACMP installed)

Warning: At this time (04/26/04) there is a bug in HACMP 5.1 that will corrupt/delete the ODM table entries for HACMP event node_up and node_down if no cluster is currently defined before installing/removing the stk.hacmp package. IY56470 fixes this problem. This fix is being posted on StorageTek's CRC site as well as on the Cambex DPF failover driver CD. You must install this fix prior to installing/removing the stk.hacmp package. If you do not do this, the odm entries may be deleted and you will have to remove and reinstall all of the HACMP software and patches. Install the fix by running the following command in the directory where the fix file is mounted/download to.

```
emgr -e iy56470.050404.epkg.Z
```

For an HACMP environment, you must install all three stk.* packages. For a non-HACMP environment, you must install stk.svadisk and stk.renumber packages.

Check for current installed versions by running the following commands.

```
lslpp -L stk.svadisk

lslpp -L stk.renumber

lslpp -L stk.hacmp
```

Install the appropriate set of packages as follows (remove old versions if needed).

```
installp -aXd stk.svadisk all

installp -aXd stk.renumber all

installp -aXd stk.hacmp all
```

Verifying Cable Connections and Scanning for LUNs

1. When scanning for LUNs, there must not be any hosts that have the volume/LUN varied on or mounted. Otherwise the LUN will not be able to be accessed to read the physical volume identification (PVID) off of it. When scanning for LUNs after running the `rmdev` command, make sure that no other hosts have the volumes varied on.
2. If the HBA to fcs# pairing is known, reconnect your FC cables to the correct HBAs.
3. If the HBA to fcs# pairing is not known, perform the following procedure.
   - Connect the FC cable labeled fcs0/controller A to one of the HBAs

- Run `cfgmgr`
- Run `lsdev - Cc disk`
- If no hdisks are seen, move the one connected cable to the other HBA, and rerun the previous two commands.
- Once hdisks are detected and using a failover architecture, run the following command
  `/usr/lpp/cbxdpf/dpfutil listall`

  the output will look similar to the following.

```
Device    Active                          Standby
hdisk1    cbx0 (fscsi1 0,0)
hdisk2    cbx1 (fscsi1 0,1)
hdisk3    cbx6 (fscsi1 0,2)
hdisk4    cbx7 (fscsi1 0,3)
```

  In the above example fscsi1 indicates that we are connected to the fcs1 HBA. In this case we are connected to the wrong HBA. We want to be connected to fcs0 HBA. If this is the case, we will fix the cabling later in the procedure.

- Once hdisks are detected and not using a failover architecture, run the following command
  - `rmdev -Rdl fcs1`
  - if the command removes hdisks, this indicates that we are connected to the fcs1 HBA. In this case we are connected to the wrong HBA. We want to be connected to fcs0 HBA. If this is the case, we will fix the cabling later in the procedure.
- If the cabling for fcs0/controller A is correct, connect the fcs1/controller B cable to the empty HBA, and run `cfgmgr`. Label the HBAs with the correct fcs numbers.
- If the cabling for fcs0/controller A is incorrect, then perform the following steps.
  - `rmdev -Rdl fcs1`
  - move the FC cable for fcs0/controller A to the other HBA
  - connect the FC cable for fcs1/controller B to the empty HBA
  - label the HBAs with the correct fcs numbers
  - run `cfgmgr`.

Verifying Volumes/LUNs/hdisks

The new hdisks should show up in smit or using the `lsdev` command.

The following example is for a failover solution:

```
# lsdev -Cc disk

hdisk0 Available 10-60-00-10,0   16 Bit LVD SCSI Disk Drive

hdisk1 Available 10-68-05        STK Universal XPort
```

```
hdisk2 Available 20-58-05          STK Universal XPort
hdisk3 Available 10-68-01-01  Storagetek BladeStore with Path Failover
hdisk4 Available 10-68-02-01  Storagetek BladeStore with Path Failover
hdisk5 Available 10-68-03-01  Storagetek BladeStore with Path Failover
hdisk6 Available 10-68-04-01  Storagetek BladeStore with Path Failover
```

The STK Universal XPort entries are for the access volumes/LUNs.

IF YOU ARE NOT USING THE FAILOVER SOLUTION, THE INSTALLATION IS COMPLETE. Proceed to the "Performance considerations" section on page 22. The following sections cover the failover solution.

If using the failover solution, run the command:
`/usr/lpp/cbxdpf/dpfutil listall`

This will display all paths with preferred paths under Active:

```
Device    Active                          Standby
hdisk1    cbx0 (fscsi0 0,0)               cbx4 (fscsi1 0,0)
hdisk2    cbx1 (fscsi0 0,1)               cbx5 (fscsi1 0,1)
hdisk3    cbx6 (fscsi1 0,2)               cbx2 (fscsi0 0,2)
hdisk4    cbx7 (fscsi1 0,3)               cbx3 (fscsi0 0,3)
```

Default Primary Path and Advanced Configurations

There are other DPF commands (via DPFutil / SMIT) that are not documented here. We strongly suggest using only the commands specified within this document. Modifying other settings incorrectly within the DPF failover driver may prevent the automatic recovery process from doing its operations correctly.

These other commands, that are available but not documented here, deal with administrating preferred/primary paths. The default primary path is the one specified on the array when creating the volume/LUN (the owning controller/ preferred path). This default primary path is not displayed by these undocumented commands. When you run the DPFutil listall command, you will see that the active path matches this default primary path that is administered on the array. This default adminstration should be sufficient for most installations.

If implementing an advanced configuration where multiple paths to the same controller exist (more than two paths to LUN configuration), then these undocumented administration commands can be used to manually administer/ distribute the primary paths across all of the available paths to the same controller, or bypass certain paths (like switch interlink connections) as a primary path. If you are an advanced user, you should limit your changes to ONLY administrating the primary path. You should not attempt to force path movement since this will have bad interactions with the background

autorecovery process. It is best to experiment with the advanced administration and display capabilities to understand how it all interacts. Once you understand how all of the commands work and interact, create the appropriate shell script that contains the DPFutil commands that administer/display the primary paths for this custom installation.

Warning: These administered changes for primary paths will be lost on a upgrade of the Cambex DPF Failover driver, so they should be scripted (via DPFutil commands, not SMIT) for your installation and saved to be used in the event of a Cambex DPF Failover driver upgrade.

Testing Failover

The following sections describe testing the various path failover scenerios.

In these test procedure examples, a path failure was caused by removing the fibre cable from controller A while doing I/O to all four LUNs. I/O should resume after a few seconds. It is important to note that a LUN will not failover until I/O is attempted to it, thus triggering a failure of the current path. The failure of the standby path is detected by the background DPF autorecovery process. This will detect the failure of the backup path within a few minutes.

Display the path status using the `dpfutil listall` command or `Smit Devices >> Cambex Fibre Channel Adapter >> Change Cambex Dynamic Path Management / Show Device Status`.

Normal Condition

Example outputs before cable was pulled:

```
      /usr/lpp/cbxdpf/dpfutil listall

# Device    Active                          Standby
  hdisk1    cbx0 (fscsi0 0,0)               cbx4 (fscsi1 0,0)
  hdisk2    cbx1 (fscsi0 0,1)               cbx5 (fscsi1 0,1)
  hdisk3    cbx6 (fscsi1 0,2)               cbx2 (fscsi0 0,2)
  hdisk4    cbx7 (fscsi1 0,3)               cbx3 (fscsi0 0,3)
```

Single Path FAILOVER Condition

Disconnect the cable from controller A with active I/O to all LUNs. After a few minutes the `dpfutil listall` command displays the following.

```
# Device   Active                    Standby
  hdisk1   cbx4 (fscsi1 0,0)       Fcbx0 (fscsi0 0,0)
  hdisk2   cbx5 (fscsi1 0,1)       Fcbx1 (fscsi0 0,1)
  hdisk3   cbx6 (fscsi1 0,2)       Fcbx2 (fscsi0 0,2)
  hdisk4   cbx7 (fscsi1 0,3)       Fcbx3 (fscsi0 0,3)
```

When the “single path bad” condition has been detected for a LUN, a single error message is generated in the error log as follows.

Identifier:5960EC86 Label:DPF_FAILOVER Text:”MAIN PATH WRAPPED TO BACK-UP PATH” and the active I/O is switched to the good path.

Repair the failing path condition, e.g. cable, SFP transceiver, controller, or HBA

Verify that the correct paths are restored using `dpfutil listall`

```
# Device    Active                          Standby
  hdisk1    cbx0 (fscsi0 0,0)               cbx4 (fscsi1 0,0)
  hdisk2    cbx1 (fscsi0 0,1)               cbx5 (fscsi1 0,1)
  hdisk3    cbx6 (fscsi1 0,2)               cbx2 (fscsi0 0,2)
  hdisk4    cbx7 (fscsi1 0,3)               cbx3 (fscsi0 0,3)
```

This indicates that the daemon has recovered the path. The restore daemon will redistribute the LUNs back to their preferred paths.

Single Path FAILOVER Condition while REBOOTING

If a path is failed while the host is rebooted, the following will be displayed by the `dpfutil listall` command.

```
# Device    Active                          Standby
  hdisk1    cbx4 (fscsi1 0,0)
  hdisk2    cbx5 (fscsi1 0,1)
  hdisk3    cbx6 (fscsi1 0,2)
  hdisk4    cbx7 (fscsi1 0,3)
```

Repair the failure condition, and then run the `cfgmgr` command.

All Paths Bad Condition without HACMP

Disconnect all the cables from controller A and B with active I/O to all LUNs. After a few minutes the `dpfutil listall` command displays the following.

```
# Device   Active                    Standby
  hdisk1   cbx4 (fscsi1 0,0)       Fcbx0 (fscsi0 0,0)
  hdisk2   cbx5 (fscsi1 0,1)       Fcbx1 (fscsi0 0,1)
  hdisk3   cbx6 (fscsi1 0,2)       Fcbx2 (fscsi0 0,2)
  hdisk4   cbx7 (fscsi1 0,3)       Fcbx3 (fscsi0 0,3)
```

Note: The failover driver will always keep the active path as nonfailed as part of the recovery process.

When the “all paths bad” condition has been detected for a LUN, the following error messages are generated in the error log. You will get multiple instances of the following message as the failover software attempts to find a good path. Identifier:5960EC86 Label:DPF_FAILOVER Text:”MAIN PATH WRAPPED TO BACK-UP PATH”

After a retry limit has been reached for finding a good path, the following error message is generated.

Identifier:FDB66493 Label:DPF_UNRECOV Text:”THRESHOLD HAS BEEN REACHED”

The pattern of errors repeats for each LUN until the problem is corrected.

All Paths Bad Condition with HACMP

Disconnect all the cables from controller A and B with active I/O to all LUNs. After a few minutes the `dpfutil listall` command displays the following.

```
# Device   Active                    Standby
  hdisk1   cbx4 (fscsi1 0,0)       Fcbx0 (fscsi0 0,0)
  hdisk2   cbx5 (fscsi1 0,1)       Fcbx1 (fscsi0 0,1)
  hdisk3   cbx6 (fscsi1 0,2)       Fcbx2 (fscsi0 0,2)
  hdisk4   cbx7 (fscsi1 0,3)       Fcbx3 (fscsi0 0,3)
```

Note: The failover driver will always keep the active path as nonfailed as part of the recovery process.

When the “all paths bad” condition has been detected for a LUN, the following error messages are generated in the error log. You will get multiple instances of the following message as the failover software attempts to find a good path.
Identifier:5960EC86 Label:DPF_FAILOVER Text:”MAIN PATH WRAPPED TO BACK-UP PATH”

After a retry limit has been reached for finding a good path, the following error message is generated.

Identifier:FDB66493 Label:DPF_UNRECOV Text:”THRESHOLD HAS BEEN REACHED”

Mail is also sent to root with the following message.
“HACMP NOTICE: cbxdpf unrecoverable failure on disk *x*“

Then a *halt -q* is automatically performed to shut down the node if other HACMP nodes are still up. This is done to force another HACMP node to bring the resources online. If no other HACMP nodes are available, it will leave the resources in an “all paths bad” condition, with continuing error messages for each LUN, until the problem is rectified.

Performance considerations

There are several parameters of the disk and HBA driver that affect performance.

These are addressed in the odm database parameters. The ODM settings for the storage system have already been set. See default settings in the example below.

1. The osdisk (the default disk device type) ODM settings have default queue type set to NONE and queue depth set to 1. The updated ODM settings are type = SIMPLE and depth=255. These may be changed using `lsattr` command to view and `chdev` command to modify these attributes. NOTE: The relative effect of this parameter is limited by the controller, which is independently capable of 512 concurrent I/Os. The combined pair with mirrored cache, however, presents an array limitation of 768 concurrent I/Os. You can display the current parameters for an hdisk with the `lsattr -El hdisk(n)` command. You can display the current parameters for an HBA with the `lsattr -El fcs(n)` command.
2. The default osdisk ODM settings set the cache mode pages in such a way as to disable any read ahead caching and turn off write cache enable. The updated ODM settings remove the cache mode changes.
3. The updated osdisk ODM adds the max_coalesce and max_transfer attributes. The max_coalesce allows the scsi driver to coalesce large sequential I/O into larger single I/O's which helps performance on those types of I/O patterns. The max_transfer allows tuning of the maximum size of a transfer packet, with the default setting of 0x40000.

AIX - Installation Support

2

The following pages are provided to support your attachment to the RS6000 AIX operating system. This chapter will:

- help you resolve problems
- let you know what HBAs are supported
- provide information about switch configuration.

Introduction

This document was created as a transfer of knowledge from engineering to our support personnel. It describes the various methods of gathering information, troubleshooting, fixing, and tuning the FlexLine™ 200/300 Series storage systems - AIX solution with either Cambex PC1000, Emulex LP8000, or newer HBAs (as documented in chapter one). All the examples are taken from an actual install, but should be the same for all future storage system solutions with only the inquiry strings being different.

Overview

Problems should only be escalated if the troubleshooting methods in this document are exhausted. The information required to analyze the problem should be gathered prior to escalation.

Problem escalation

Escalation, based on the problem type, requires the following information.

Initial install

detailed description of problem
storage system Profile
lsdev –Cc disk output
lsdev –Cc adapter output
cfgmgr –v > /tmp/cfgmgr.log file
lsattr –El hdisk(n) (from hdisk with problem)
ps –ef > /tmp/process.log
lslpp –l > /tmp/version.log

Errors after initial install success

detailed description of problem
storage system Profile
storage system event log (at least last 500 entries)
lsdev –Cc disk output
lsdev –Cc adapter output
errpt –a output
lsattr –El hdisk(n) (from hdisk with problem)
ps –ef > /tmp/process.log
lslpp –l > /tmp/version.log

Supported HBAs and Operating System

The following sections provide the necessary configuration, hardware, and software requirements for the supported HBAs. Currently Cambex PC1000, PC2000, IBM 6227, IBM 6228, and IBM 6239 cards are supported.

Operating system: AIX 4.3.3, 5.1, and 5.2 are supported at the latest maintenance level.

If having problems loading HBA drivers or detecting HBAs, refer to the HBA vendor web page for troubleshooting hints.

Cambex PC1000

The Cambex PC1000 HBA is a 1 Gigabit fibre HBA based on the Qlogic 2200 chip. The host attach kit provides both the Cambex DPF failover driver and Cambex HBA driver for AIX. The failover driver is not free, but a licensed product.

Cambex PC2000

The Cambex PC2000 HBA is a 2 Gigabit fibre HBA based on the Qlogic 2300 chip.

Cambex HBA Driver installation

Install the driver by mounting the cdrom

(something like: `mount –rv cdrfs /dev/cd0`)

Change to the directory containing the driver and run the following command.

```
installp –acd PC1000.image all
```

The output text should show that the driver installed successfully.

Verifying HBAs are found

Run cfgmgr after driver and HBA installation to detect and configure the HBAs.

The command lsdev –Cc adapter will display a list of installed adapters. It will look like this:

AIXPCI3/> lsdev -Cc adapter

sa0 Available 01-S1 Standard I/O Serial Port

sa1 Available 01-S2 Standard I/O Serial Port

sa2 Available 01-S3 Standard I/O Serial Port

siokma0 Available 01-K1 Keyboard/Mouse Adapter

fda0 Available 01-D1 Standard I/O Diskette Adapter

scsi0 Available 10-60 Wide SCSI I/O Controller

scsi1 Available 10-68 Wide/Fast-20 SCSI I/O Controller

scsi2 Available 10-70 Wide/Fast-20 SCSI I/O Controller

scsi3 Available 30-58 Wide SCSI I/O Controller

mga0 Available 30-68 GXT120P Graphics Adapter

sioka0 Available 01-K1-00 Keyboard Adapter

ppa0 Available 01-R1 Standard I/O Parallel Port Adapter

ent0 Available 10-78 IBM 10/100 Mbps Ethernet PCI Adapter (23100020)

ent1 Available 10-80 IBM PCI Ethernet Adapter (22100020)

sioma0 Available 01-K1-01 Mouse Adapter

fscsi5 Available 30-60 Cambex Fibre Channel I/O Controller

fcs0 Available 20-58 FC Adapter

The HBAs (fcs(n) for IBM HBAs, or fscsi(n) for Cambex HBAs) should show as “Available”. If not for Cambex HBAs, check for correct driver installation by `using lslpp –l PC1000.driver.image`. This will show if installed and version.

AIXPCI3/> lslpp -l PC1000.driver.image

Fileset	Level	State	Description

Path: /usr/lib/objrepos

PC1000.driver.obj	2.5.10.5	COMMITTED	PC1000 Fibre Channel Adapter Driver (AIX 5L)

If this is good, check that HBA(s) are correctly seated, defective, etc.

HBA WWN and other parameters

For IBM HBAs use `lscfg –vl fcs(n)` (where fcs(n) is from lsdev –Cc adapter) to get more information about the HBA. For Cambex HBAs use `lscfg –vl fscsi(n)` (where fscsi(n) is from lsdev –Cc adapter) to get more information about the HBA. This provides the Network Address field (WWN) needed for partitioning and the microcode version (bios) field (Z9). At this time the current versions of microcode for IBM HBAs is as follows. For the IBM 6239 HBA, the Z9 field should be `HS1.00X5`. For the IBM 6228 HBA, the Z9 field should be `CS3.91A1`. For the IBM 6227 HBA, the Z9 field should be `SS3.30X1`. The following example is for the IBM 6239 HBA.

For the latest version of microcode for the IBM HBA, go to the IBM web site for information on the latest version of released microcode and instructions on how to upgrade the microcode.

Warning: Do not upgrade microcode from a non-IBM source (e.g. Emulex web site) since this may disable the card/host.

```
fcs1             U0.1-P1-I1/Q1  FC Adapter

      Part Number.................00P4295
      EC Level....................A
      Serial Number...............1E32508B9F
      Manufacturer................001E
      Feature Code/Marketing ID...5704
      FRU Number..................     00P4297
      Device Specific.(ZM)........3
      Network Address.............10000000C934A595
      ROS Level and ID............02E01035
      Device Specific.(Z0)........2003806D
      Device Specific.(Z1)........00000000
      Device Specific.(Z2)........00000000
      Device Specific.(Z3)........03000909
      Device Specific.(Z4)........FF601032
      Device Specific.(Z5)........02E01035
      Device Specific.(Z6)........06631035
      Device Specific.(Z7)........07631035
      Device Specific.(Z8)........20000000C934A595
      Device Specific.(Z9)........HS1.00X5
      Device Specific.(ZA)........H1D1.00X5
      Device Specific.(ZB)........H2D1.00X5
      Device Specific.(YL)........U0.1-P1-I1/Q1
```

The network address is the WWN of the HBA.

Running a trace

If there are I/O errors that are causing problems, a trace may be taken by using the following procedure. Start the trace and perform the failing operation. Stop the trace as soon as possible after the failing condition in order to minimize the chances of wrapping the trace. The function starts the AIX trace facility, so the bos.sysmgmt.trace fileset needs to be installed.

Start trace using:

/usr/lpp/cbxfc/runtrace

Stop trace using:

trcstop

Save trace in memory to file using:

trcrpt > /tmp/trc.log

Device Detection

The storage system should be connected to the HBA(s) and configured correctly. This includes upgrading the storage system firmware/nvsram to the latest versions.

What you should see using the failover driver

If all is working correctly, the lsdev –Cc disk command (for BladeStore) should show the devices:

```
# lsdev -Cc disk
```

```
hdisk0 Available 10-60-00-10,0 16 Bit LVD SCSI Disk Drive
hdisk1 Available 10-68-05      STK Universal XPort
hdisk2 Available 20-58-05      STK Universal XPort
hdisk3 Available 10-68-01-01   Storagetek BladeStore with Path Failover
hdisk4 Available 10-68-02-01   Storagetek BladeStore with Path Failover
hdisk5 Available 10-68-03-01   Storagetek BladeStore with Path Failover
hdisk6 Available 10-68-04-01   Storagetek BladeStore with Path Failover
```

Common mistakes (using a BladeStore example)

Symptom: Failover does not work and I see some BladeStore volumes on each path.

Possible problem:
a) May not have Cambex failover driver installed.
b) May not have Cambex DPF license installed.
c) May have selected AIX for host type instead of AIX_FO.

Symptom: I see the UTM LUNs in lsdev but the Client cannot find the BladeStore.

Possible Problem: The agent software is started at boot time or install time and scans for UTM LUNs. If they are not found (because of problems like

not mapping UTM LUNs), agent needs to be cycled to rediscover the UTM LUNs once they show up in lsdev. This is done by:

Type: SMagent stop
SMagent stopped
Type: SMagent start
After a few seconds, messages like the following will appear:
Checking device /dev/rhdisk2 : Skipping
Checking device /dev/rhdisk5 : Skipping
Checking device /dev/rhdisk6 : Skipping
Checking device /dev/rhdisk7 : Skipping
Checking device /dev/rhdisk8 : Skipping
Checking device /dev/rhdisk9 : Skipping
Checking device /dev/rhdisk10 : Skipping
Checking device /dev/rhdisk11 : Skipping
Checking device /dev/rhdisk12 : Skipping
Checking device /dev/rhdisk13 : Skipping
Checking device /dev/rhdisk14 : Skipping
Checking device /dev/rhdisk15 : Skipping
Checking device /dev/rhdisk16 : Activating
Checking device /dev/rhdisk3 : Skipping
Checking device /dev/rhdisk4 : Skipping
Checking device /dev/rhdisk17 : Activating

Running...
There should be 2 activated UTM LUNs. If they do not show up in this scan list but DO show up in lsdev, it is possible that the UTM LUNs were removed using rmdev without first stopping agent. This can leave the UTM LUNs reserved. Cycle the controllers power to release these.

Other troubleshooting methods

If the adapters are active with the lsdev –Cc adapter command but the disks don't show up, run cfgmgr with the –v (verbose) parameter and send it to a file. (cfgmgr –v > /tmp/cfg1.log).

Edit this file and search for the fcs(n) (for IBM HBAs) or fscsi(n) (for Cambex HBAs) devices that are the HBAs. There will usually be some indication of whether any targets (controllers) were discovered.

The following procedures cover special tasks required.

Removing existing configuration on the host

If upgrading an existing storage system installed and configured on AIX, perform the following actions to remove existing devices.

- Unmount any mounted file systems using the umount command or smit
- Vary offline any volume groups using varyoffvg or smitty

- If using inband storage management, run the following command.

  ```
  SMagent stop
  ```

- If using Cambex DPF failover driver, run the following command.

  ```
  stopsrc -s cbxdpfar
  ```

- Determine the Cambex adapter name (fscsi(n)) (or IBM name fcs(n)) by using the lsdev –Cc adapter command. Locate the name by finding the Camber or IBM adapter entries. In the following examples, fscsi4 and fscsi5 are the Cambex adapters and fcs0 is an IBM adapter.

  fscsi4 Available 20-58 Cambex Fibre Channel I/O Controller

  fscsi5 Available 20-60 Cambex Fibre Channel I/O Controller

  fcs0 Available10-68 Fibre Channel Adapter

- Remove the adapters and scsi devices using the rmdev –Rdl fscsi(n)

  Or fcs(n) for IBM.

Examples:

rmdev –Rdl fscsi4

hdisk1 deleted

hdisk2 deleted

hdisk3 deleted

fscsi4 deleted

rmdev –Rdl fscsi5

hdisk4 deleted

hdisk5 deleted

fscsi5 deleted

- If using inband storage management, run the following command.

  ```
  SMagent start
  ```

- If using Cambex DPF failover driver, run the following command.

  ```
  startsrc -s cbxdpfar
  ```

Switch Configurations

Testing was done with 4116 and 3800 switches on both IBM and Cambex HBAs. The same zoning rules apply for failover as all other platforms.

For those that need a refresher, a zoning web page is available at:

```
http://gandalf.stortek.com/ctp/zoning/
```

NEED MORE INFORMATION?

www.storagetek.com

ABOUT STORAGETEK

Storage Technology Corporation (NYSE: STK) is a $2 billion global company that enables businesses, through its information lifecycle management strategy, to align the cost of storage with the value of information. The company's innovative storage solutions manage the complexity and growth of information, lower costs, improve efficiency and protect investments. For more information, visit www.storagetek.com, or call 1.800.275.4785 or 01.303.673.2800.

WORLD HEADQUARTERS

Storage Technology Corporation
One StorageTek Drive
Louisville, Colorado 80028 USA
1.800.525.0369