

# *A5x00 Troubleshooting Guide*



Joy Silva  
John Schell  
Sun Network Storage  
April 1999

Copyright 1998 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, SunSoft, SunDocs, SunExpress, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

**RESTRICTED RIGHTS:** Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

---

Copyright 1998 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, SunSoft, SunDocs, SunExpress, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Please  
Recycle



Adobe PostScript

# A5x00 Troubleshooting Guide

## Table of Contents



<i>Chapter</i>	<i>Page</i>
<b>1. Product Summary</b>	<b>2</b>
Hardware Overview	
Software Overview	
Drivers and boot messages	
A5x00 Diagnosis Information	
<b>2. A5x00 Fibre Channel Topology</b>	<b>12</b>
Loop node display	
FC disk architecture	
<b>3. Introduction to Troubleshooting</b>	<b>18</b>
Methodology	
Process of Elimination (POE)	
<b>4. Troubleshooting Steps</b>	<b>24</b>
Solid Failures and hung loops - bottom up	
Intermittant Errors and software problems - top down	
<b>5. Advanced Troubleshooting Techniques</b>	<b>36</b>
CRC and scsi parity errors	
rdls information	
<b>6. A5x00 Related Messages Interpretation</b>	<b>44</b>
<b>7. Service Issues</b>	<b>57</b>
FRU replacemant	
Booting issues	
<b>A. Diagnostics - STORtools, lbf, SunVTS, Symon</b>	<b>63</b>
<b>B. Revision Maintenance Procedures</b>	<b>75</b>
<b>C. Additional Resources</b>	<b>80</b>
<b>Glossary</b>	<b>82</b>

This guide is primarily intended for use by experienced system support engineers who already have a good understanding of the product. Although chapter one does contain a brief overview of the hardware and software components, it is more a quick reference than a detailed explanation and is not meant to be a training guide. Please utilize the additional resources as listed in appendix C for comprehensive product information.



# Chapter 1 - Product Summary

## Hardware Overview

FRU Block Diagram

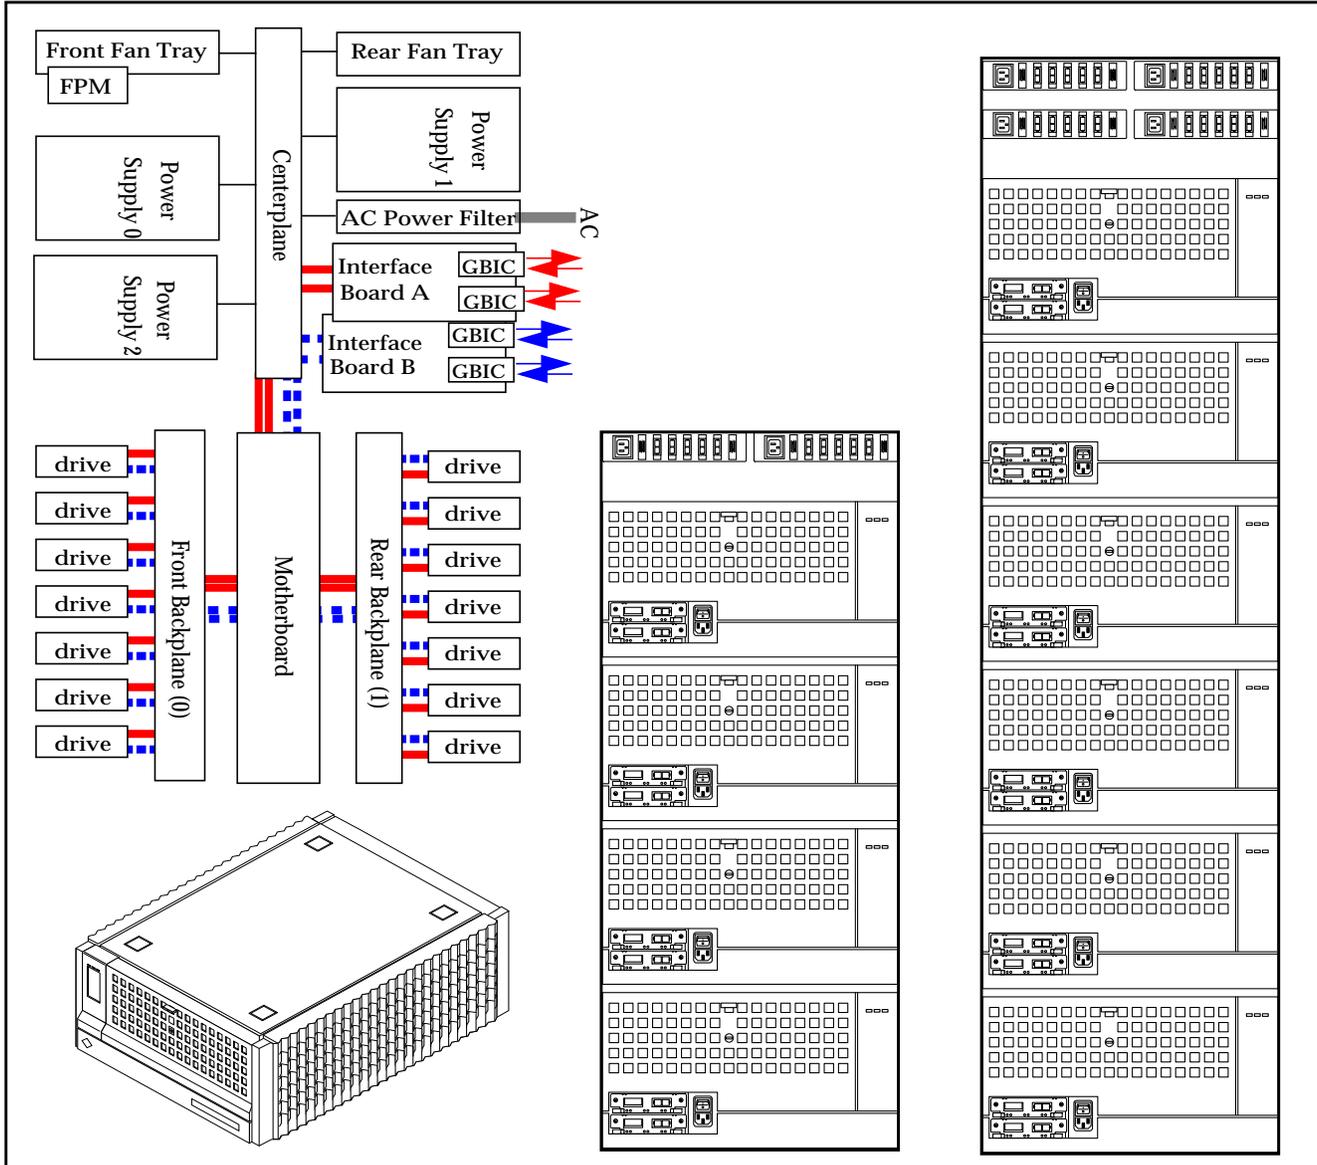
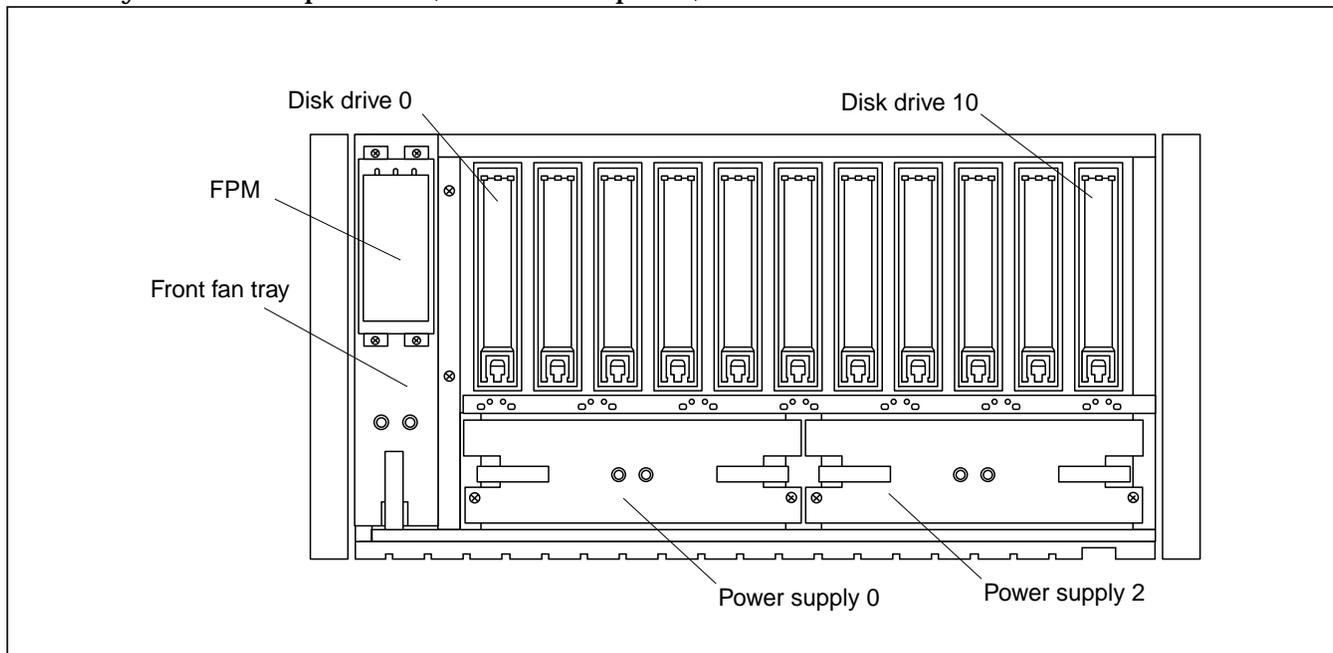


Table top array

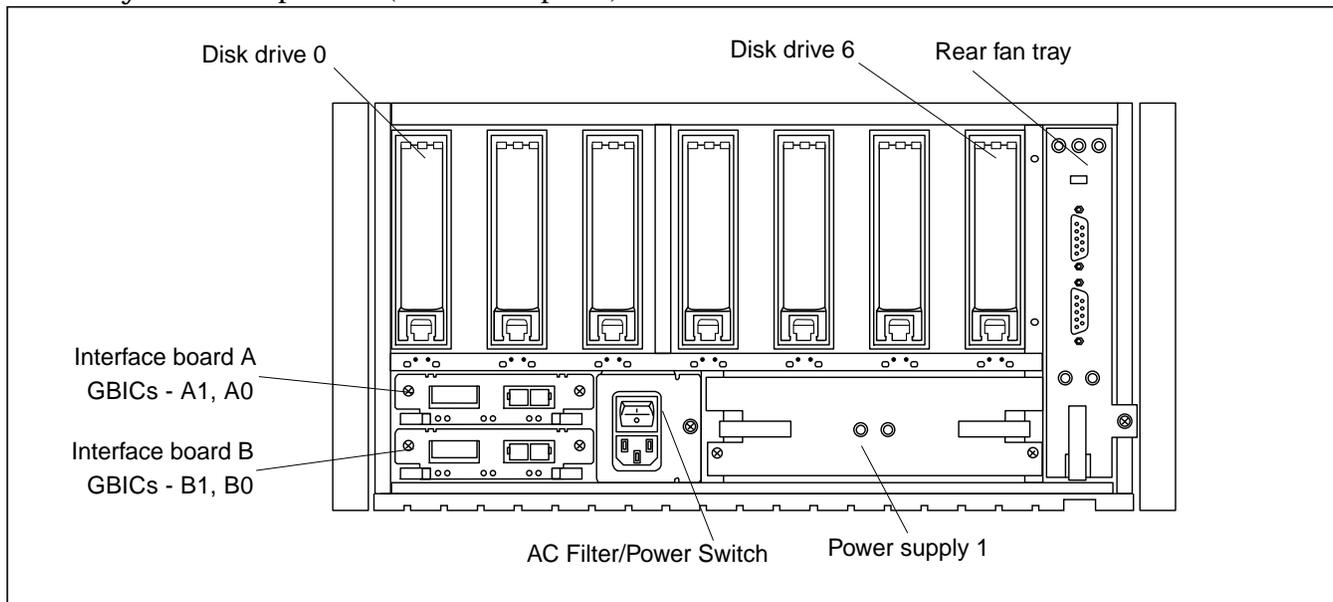
56 inch rack for 4 arrays  
and 2 hubs

72 inch rack for 6 arrays  
and 4 hubs

## Array Front Components (11 disk backplane)

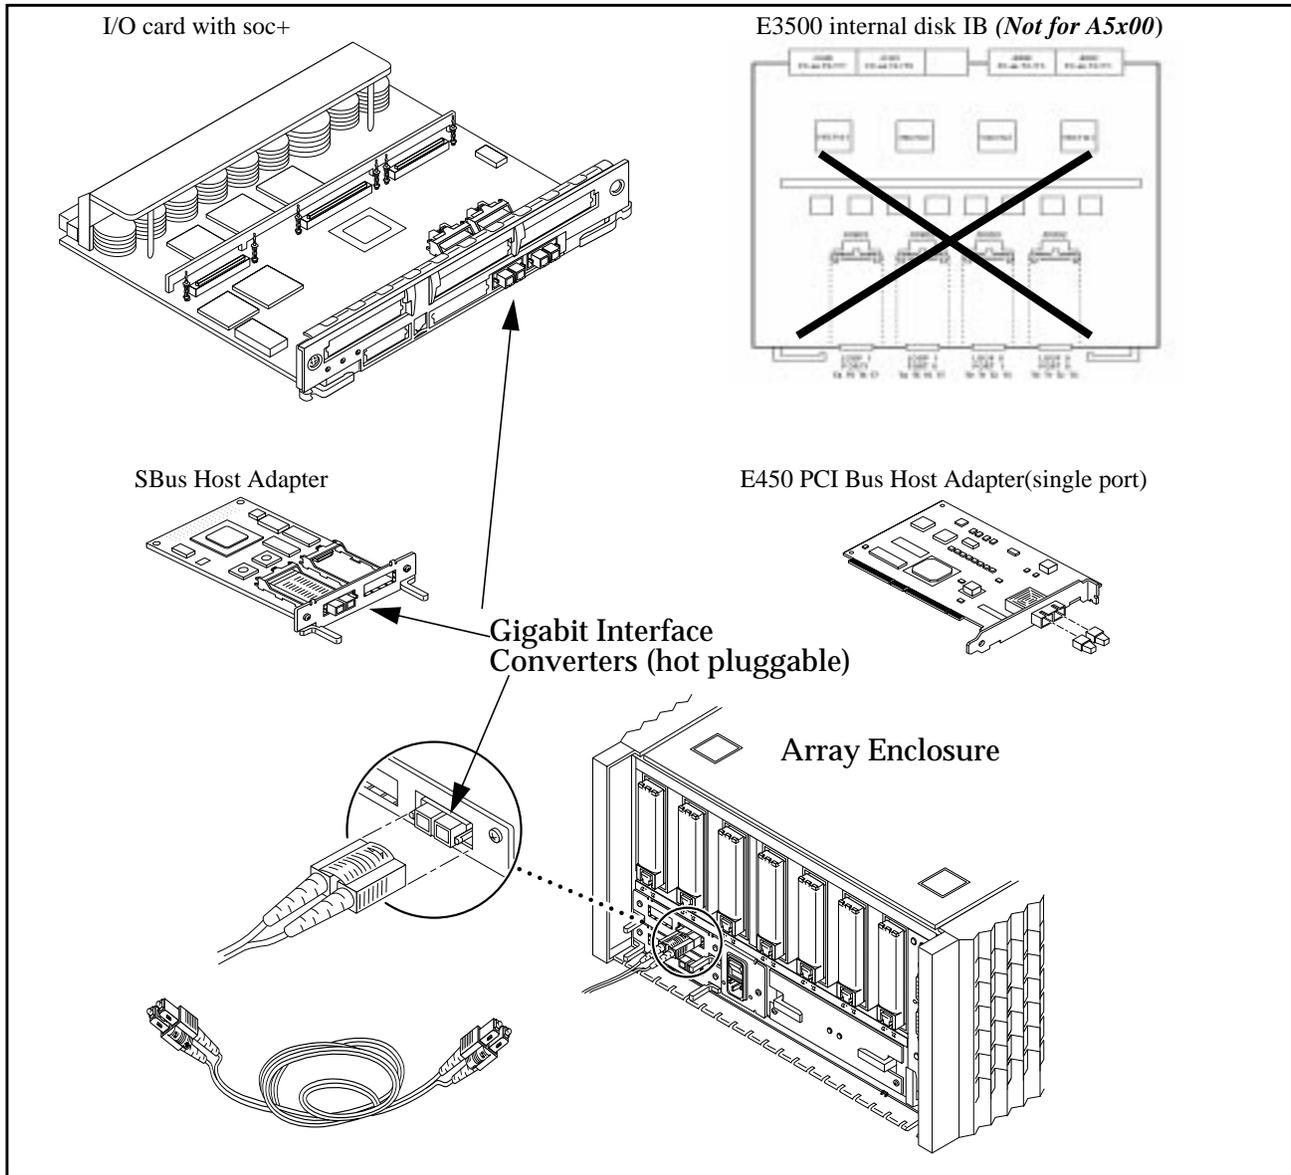


## Array Rear Components (7 disk backplane)

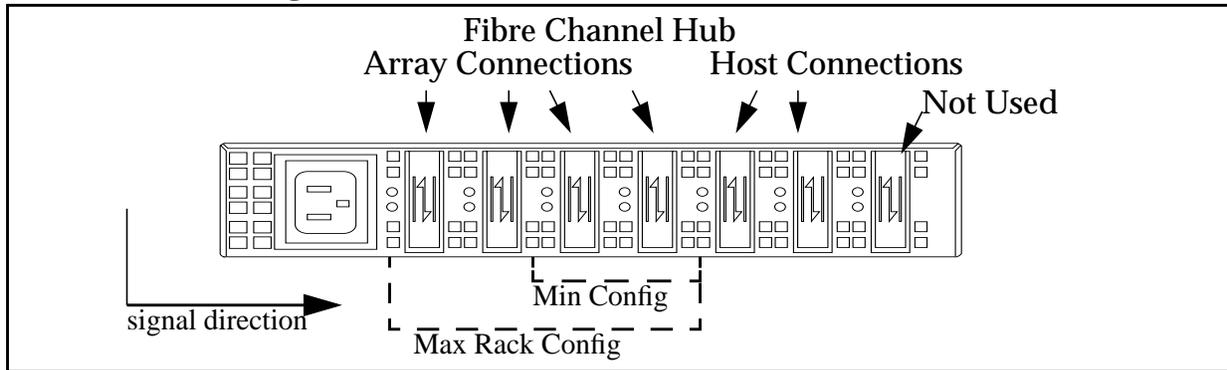




## Host Connections

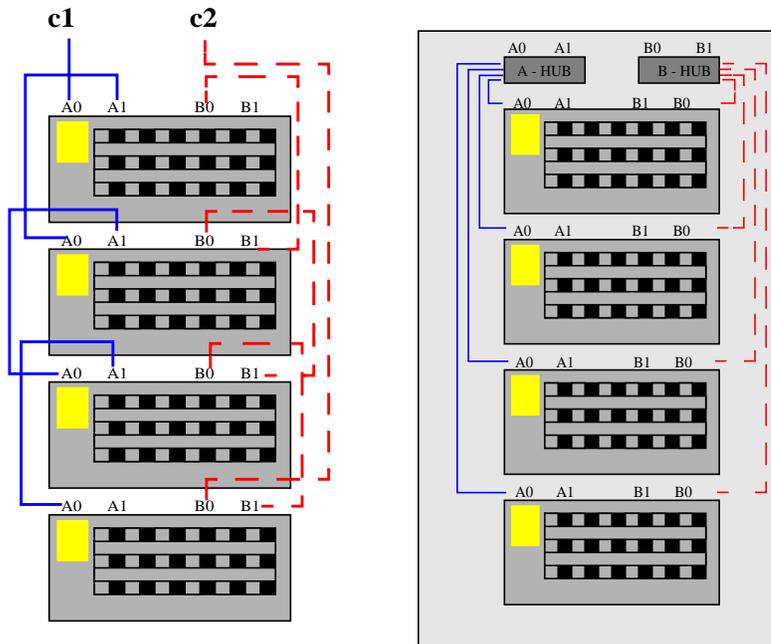


## Hub connection configuration



The factory installed hub cable configurations increases signal margins between the arrays and the hosts. The signal direction is left to right and its best to minimize the loop distance for data sent from the array. Hub slot 6 is not used. Slots 4 and 5 are for host connections starting with 4. Slots 0 through 3 are for arrays starting with 3.

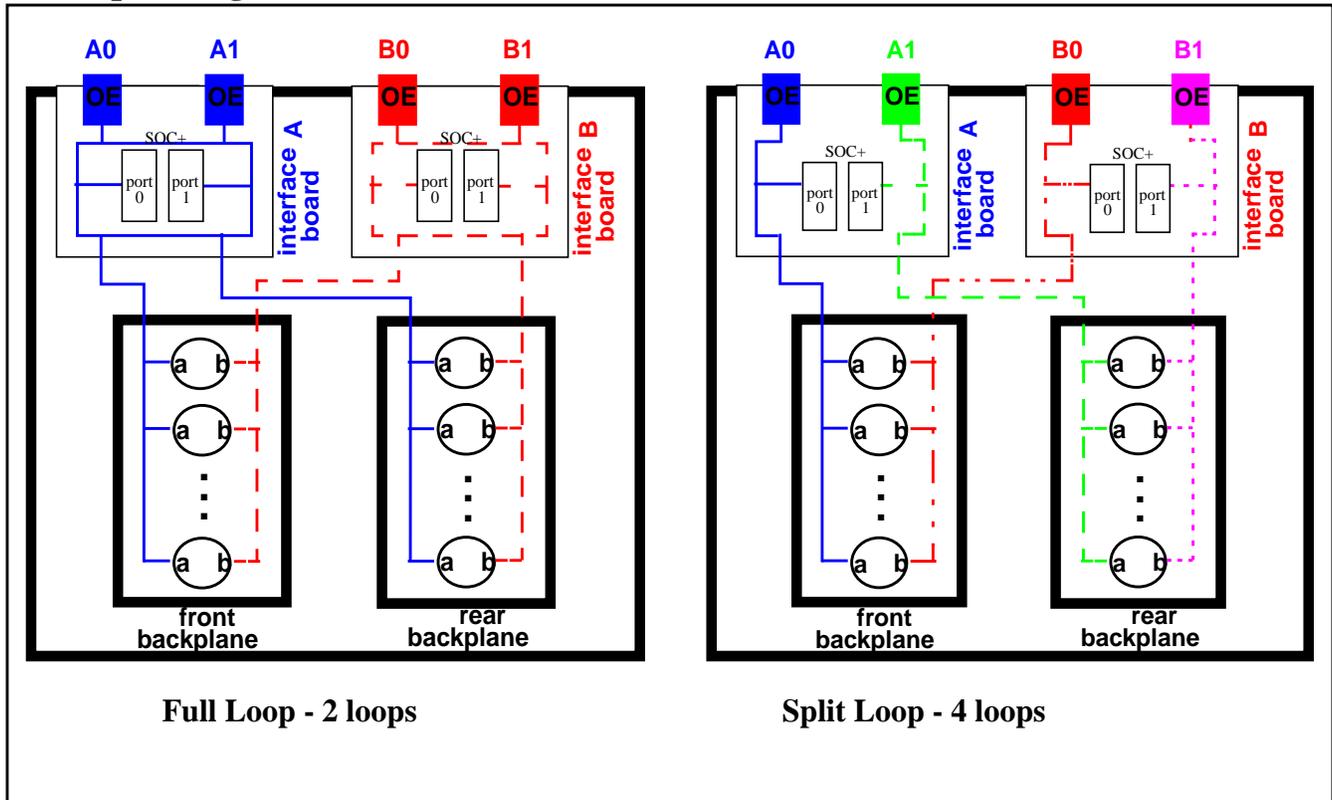
## .Daisy chains and Hubs



Multiple arrays can be connected to the same host connect/s and fc-al loop/s using daisy chained cables or hubs. The functionality and performance is identical but hubs are generally preferred for the cabling simplicity and higher fault tolerance.



## Loop Configurations



The A5x00 arrays can be configured to run as either 2 or 4 fc-al loops.

In the default full loop or 2 loop configuration all disks on both backplanes can be accessed using any of the 4 interface board (IB) ports.

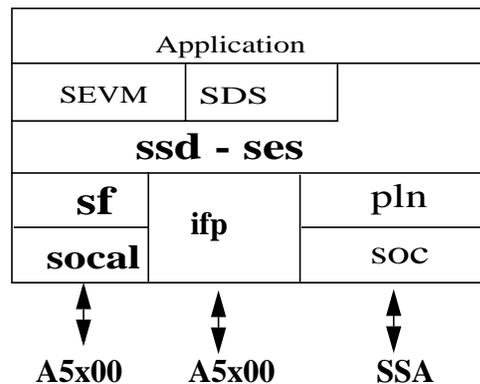
In the “split loop” or 4 loop configuration the front and rear backplanes are completely separate and must be accessed using both the 0 and 1 Interface Board (IB) connections. Connections A0 and B0 access the disks on the front backplane, connections A1 and B1 access the rear backplane.

Note the soc+ ports. These are the ses (scsi enclosure services) nodes used to monitor and control the enclosure. In full loop mode there will be two ses nodes on the loop, in split loop mode there will be one.

## Software Overview

### Solaris drivers:

- **socal** - Serial optical controller for arbitrated loop driver for the SOC+ host bus adapter(HBA)
- **sf** - SCSA compliant driver for FC-AL private loops (SSA pln equivalent). Interfaces with the **ssd** **socal** and **ses** drivers. Provides command mapping and throttling services.
- **ifp** - PCI bus equivalent to the **sbus** **socal** and **sf**.
- **ssd** - Disk driver for fiber channel arbitrated loop(FC-AL) disks and SSA disks.
- **ses** - SCSI enclosure service device driver. Communicates with the arrays interface board (IB) to monitor the physical conditions within the array and access the status and configuration services.



### Driver informational message interpretation:

There are many informational messages generated by the A5x00 drivers, particularly during the boot sequence. These messages can be used to aid in determining the configuration and the various names given to devices that get attached. This section describes common informational messages and how to interpret them.



## Fibre Channel Loop is ONLINE

```
Oct 10 14:28:12 MYSERVER unix: ID[SUNWssa.socal.link.6010] socal5: port 1: Fibre Channel Loop is ONLINE
```

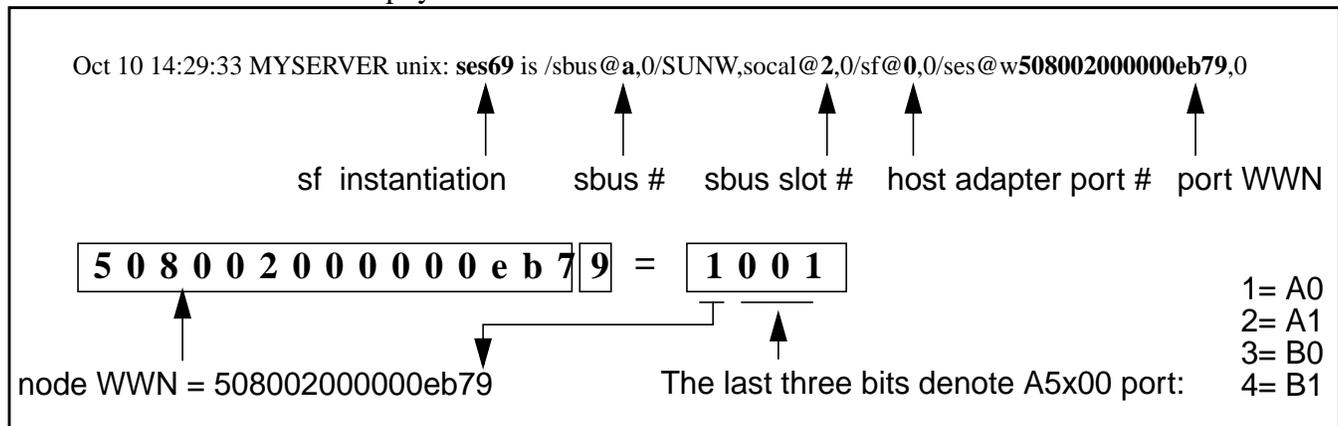
This message is displayed whenever a fibre channel loop becomes ONLINE. You would also see the online lightning bolt displayed on the array FPM (front panel module) at this time which usually indicates a functional connection.

## ses device attach

```
Oct 10 14:29:33 MYSERVER unix: ses69 at sf10: target 6d lun 0
```

```
Oct 10 14:29:33 MYSERVER unix: ses69 is /sbus@a,0/SUNW,socal@2,0/sf@0,0/ses@w508002000000eb71,0
```

This message is displayed whenever a new ses (SCSI Enclosure Services) device is attached. Each A5x00 has two ses devices per loop when configured in single loop mode, one when in split loop mode. The ses device files are used for inclosure information and control. This message ties the ses driver instantiation to the physical slot information:



**Note** - The array path uses a unique world wide numbers (WWN) to identify the array. The WWN is stored on the interconnect assembly. You must use the `luxadm remove_device` and `insert_device` commands when replacing the interconnect asseby to rebuild the ses device paths.





---

## PCI Bus ifp driver examples

```
Mar 25 16:13:00 tm6 unix: PCI-device: SUNW,ifp@1, ifp #1
Mar 25 16:13:00 tm6 unix: pci1077,21001 is /pci@4,2000/SUNW,ifp@1
Mar 25 16:13:00 tm6 unix: ifp1:Firmware Version: v1.11

Mar 25 16:13:00 tm6 unix: ssd116 at pci1077,21005: name w21000020370704be,0, bus address 1b
Mar 25 16:13:00 tm6 unix: ssd116 is /pci@1f,2000/SUNW,ifp@1/ssd@w21000020370704be,0
Mar 25 16:13:00 tm6 unix: <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>

Mar 25 16:13:00 tm6 unix: ses2 at pci1077,21001: name w5080020000006e3b,0, bus address 29
Mar 25 16:13:00 tm6 unix: ses2 is /pci@4,2000/SUNW,ifp@1/ses@w5080020000006e3b,0
```

## A5x00 Diagnosis Information

A5x00 diagnosis information can be obtained from a variety of sources. The hardware itself can be directly diagnosed using the front panel module (FPM) and LED indicators. When the hardware is not easily accessible or there is an intermittent problem, system message files and utilities can be used. In addition to the standard Solaris commands and utilities there is the STORtools package which can be installed for array management and diagnostic purposes. The following is a summary of all the available information sources.

Solaris generic:

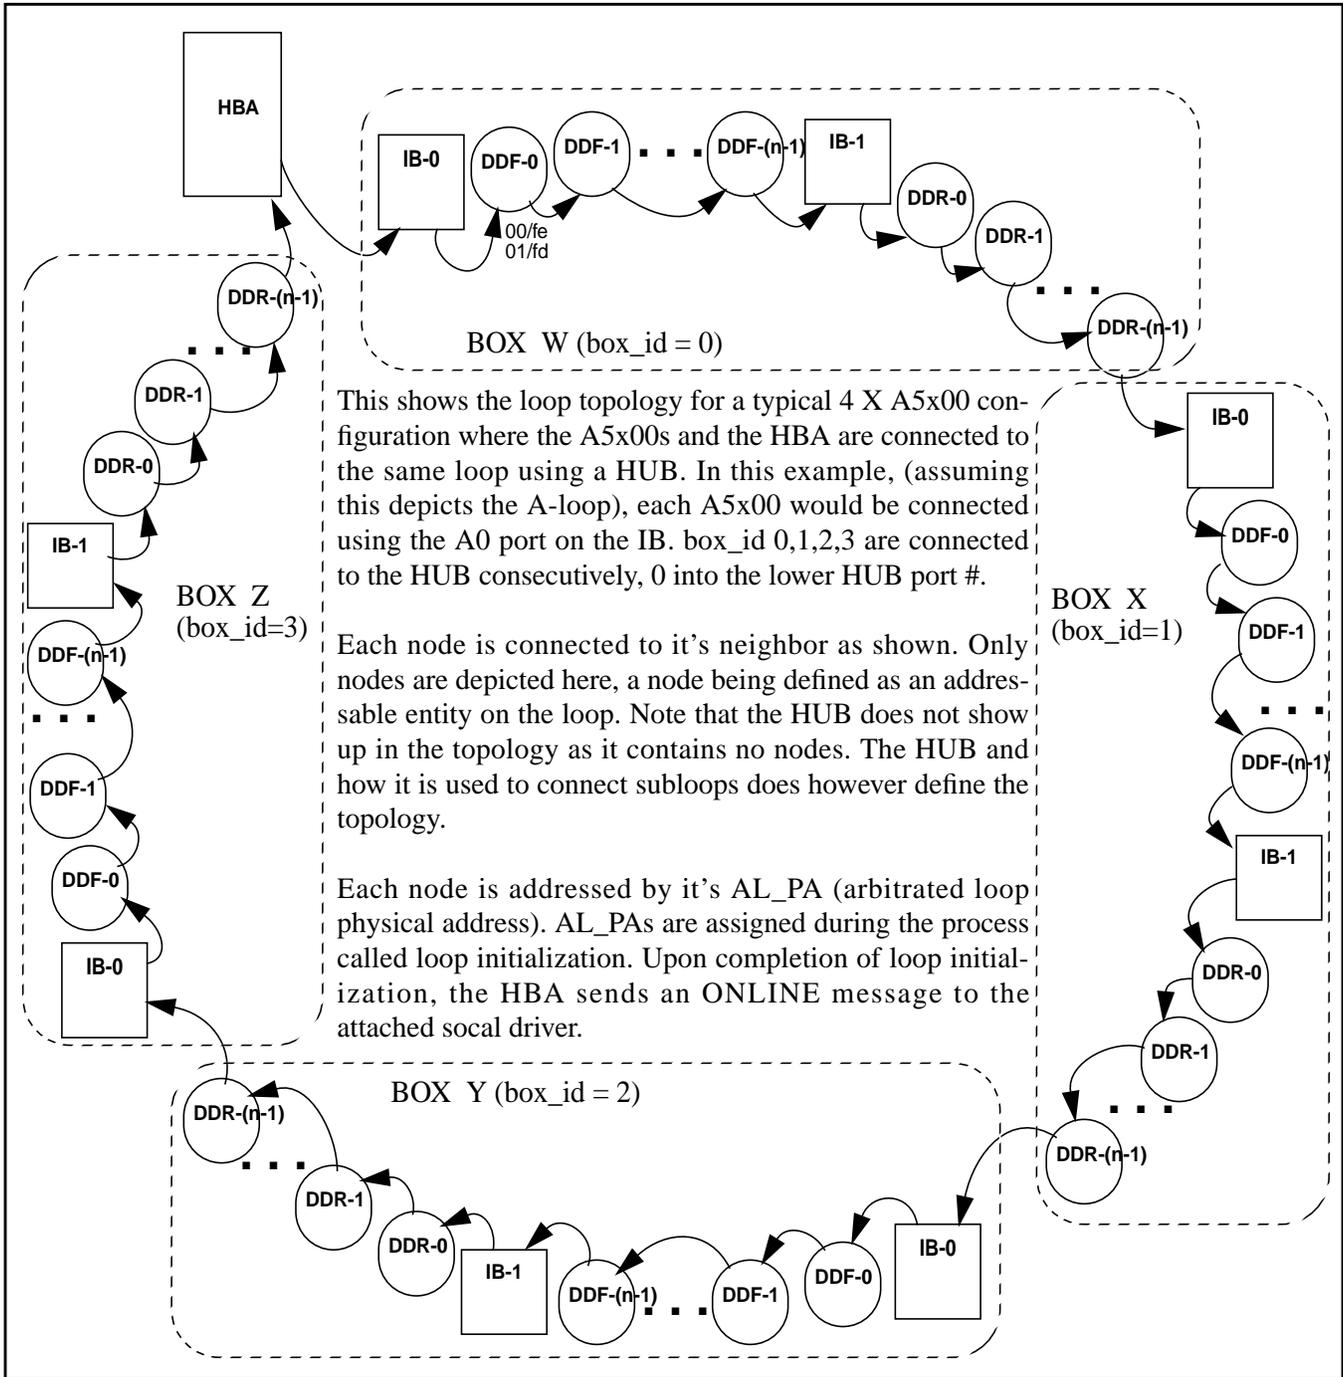
- **/var/adm/messages** - Still the main source for device event and error messages.
- **format** - The standard solaris format utility can be used for disk information without ses.
- **SunVTS** - Contains diagnosis and fault isolation tests for the array enclosure and host adapter using ses and firmware implemented soc+ diagnostics.
- **SyMon** - Can be used to get array status and error information with event notification.

A5x00 specific:

- **FPM and LEDs** - Whenever there is a solid hardware problem on the array the easiest method of locating the failing FRU will be with the FPM and LED indicators.
- **luxadm** - The luxadm command uses the ses features of the array to display status, configuration and revision information for the array and disk drives.
- **STORtools** - Comprehensive package for A5x00 revision checking, configuration management, monitoring with event notification and fc-al loop diagnostics.



# Chapter 2 - A5x00 Fibre Channel Topology





Loop topology can be viewed with the ‘luxadm -e dump\_map’ command. The output shows the position on the loop of each node with the AL\_PA, ID, world wide number (WWN) and type.

Example:

```
# luxadm -e dump_map a|more
Pos AL_PA ID Hard_Addr Port WWN Node WWN Type
0 2 7c 2 20020800207e147d 200c0800207e147d 0xff (Unknown Type)
1 d2 d d2 508002000002bbb9 508002000002bbb8 0xd (SES device)
2 ef 0 ef 210000203713121a 200000203713121a 0x0 (Disk device)
3 e8 1 e8 2100002037131529 2000002037131529 0x0 (Disk device)
4 e4 2 e4 2100002037131098 2000002037131098 0x0 (Disk device)
5 e2 3 e2 210000203713126e 200000203713126e 0x0 (Disk device)
6 e1 4 e1 21000020370ed14a 20000020370ed14a 0x0 (Disk device)
7 e0 5 e0 21000020371313e6 20000020371313e6 0x0 (Disk device)
8 dc 6 dc 21000020371311d4 20000020371311d4 0x0 (Disk device)
9 b5 1d b5 508002000002bbba 508002000002bbb8 0xd (SES device)
10 cd 10 cd 21000020370ee53b 20000020370ee53b 0x0 (Disk device)
11 cc 11 cc 21000020371315c6 20000020371315c6 0x0 (Disk device)
12 cb 12 cb 210000203713176a 200000203713176a 0x0 (Disk device)
13 ca 13 ca 210000203713172b 200000203713172b 0x0 (Disk device)
14 c9 14 c9 2100002037131543 2000002037131543 0x0 (Disk device)
15 c6 16 c6 21000020370e6c7f 20000020370e6c7f 0x0 (Disk device)
16 9e 2d 9e 5080020000025ac1 5080020000025ac0 0xd (SES device)
17 b2 20 b2 21000020370ee2bb 20000020370ee2bb 0x0 (Disk device)
18 b1 21 b1 21000020370ee060 20000020370ee060 0x0 (Disk device)
19 ae 22 ae 21000020370ee100 20000020370ee100 0x0 (Disk device)
20 ad 23 ad 21000020370ee07a 20000020370ee07a 0x0 (Disk device)
21 ac 24 ac 21000020370edce9 20000020370edce9 0x0 (Disk device)
22 ab 25 ab 21000020370ee114 20000020370ee114 0x0 (Disk device)
23 aa 26 aa 21000020370edf7b 20000020370edf7b 0x0 (Disk device)
24 75 3d 75 5080020000025ac2 5080020000025ac0 0xd (SES device)
25 98 30 98 21000020370edf89 20000020370edf89 0x0 (Disk device)
26 97 31 97 21000020370edea5 20000020370edea5 0x0 (Disk device)
27 90 32 90 21000020370ee0fb 20000020370ee0fb 0x0 (Disk device)
28 8f 33 8f 21000020370ee099 20000020370ee099 0x0 (Disk device)
29 88 34 88 21000020370eded3 20000020370eded3 0x0 (Disk device)
30 84 35 84 21000020370ee043 20000020370ee043 0x0 (Disk device)
31 82 36 82 21000020370edfc8 20000020370edfc8 0x0 (Disk device)
32 5a 4d 5a 508002000001e619 508002000001e618 0xd (SES device)
33 72 40 72 2100002037130ff2 2000002037130ff2 0x0 (Disk device)
34 71 41 71 210000203713136d 200000203713136d 0x0 (Disk device)
35 6e 42 6e 2100002037131622 2000002037131622 0x0 (Disk device)
```

First A5x00  
Front backplane

Rear backplane

Second A5x00  
Front backplane

Rear backplane

~ Example cut ~



- Pos - Is the position of the node in the array.
- AL\_PA and Hard Addr - The AL\_PA is the arbitrated loop physical address or hard address, the mappings for this can be found in the fc-al specification.
- ID - The ID, also known as the target number, is a hexadecimal byte which breaks down as follows:

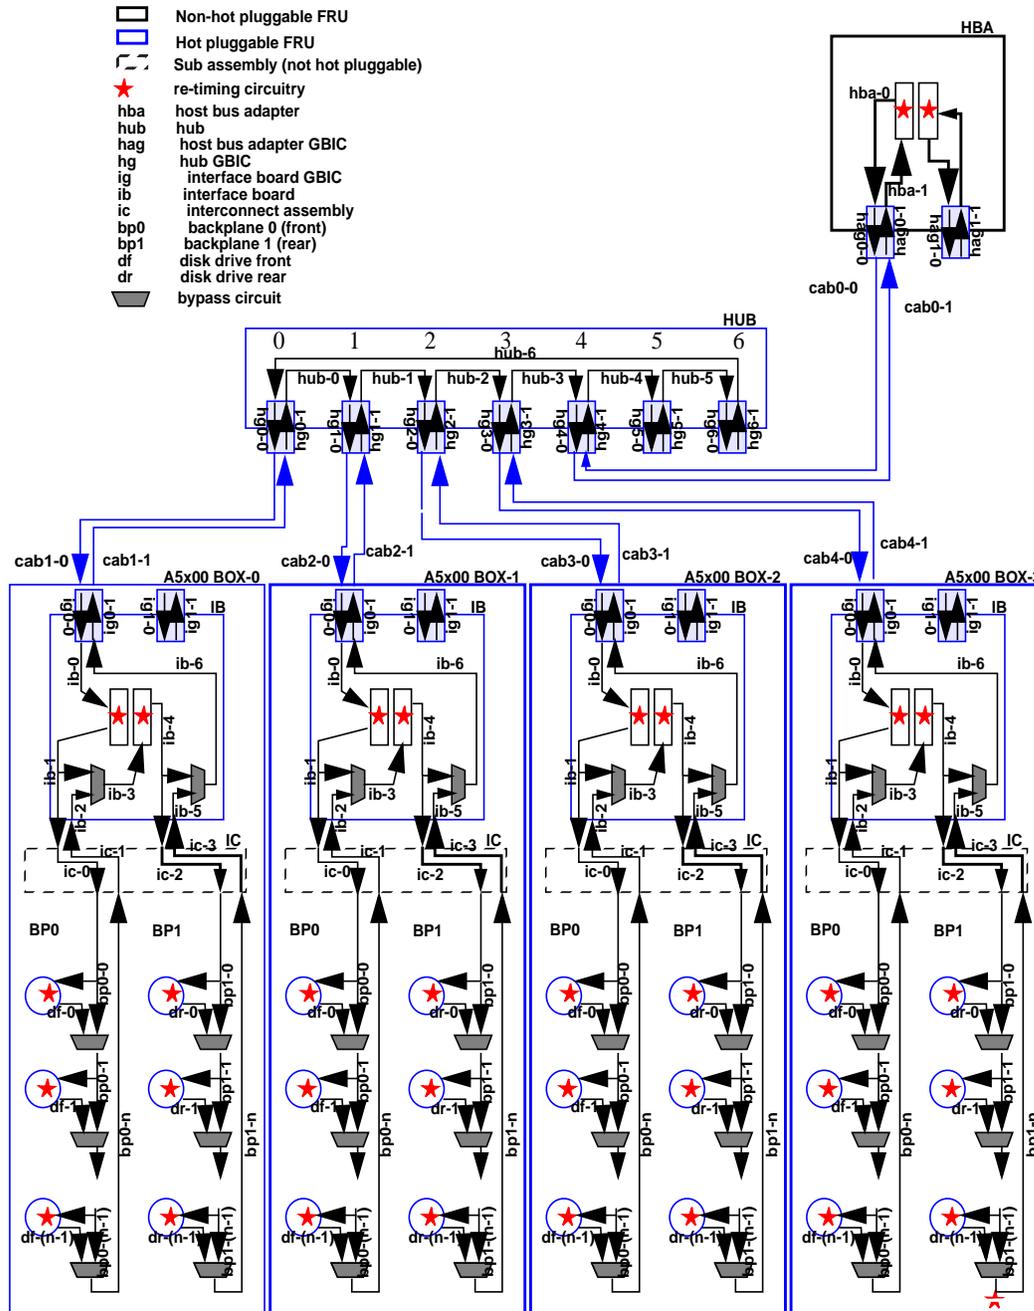
0	Box <sub>i</sub> ID	Back Plane	Slot <sub>i</sub> Number	
---	---------------------	---------------	--------------------------	--

- Box ID - 0 to 3. Each array on the loop must have a unique box id number.
- Back Plane - 0=Front, 1=Back
- Slot Number - The back plane slot number of a disk, 0 to 6 for the 14 disk array or 0 to 10 for 22 disk array.
- Port WWN - World wide number of the drive port (second digit is 1=port A, 2=port B)
- Node WWN - World wide number of the device.
- Type - Type of device on the loop.

**Node ID 32 in hexadecimal would breakdown to 0011,0010 in binary. That would be box Id 1, backplane 1, slot 2. Also 32 hex is 50 in decimal which is where the t50 number comes from (c5t50d0s0).**

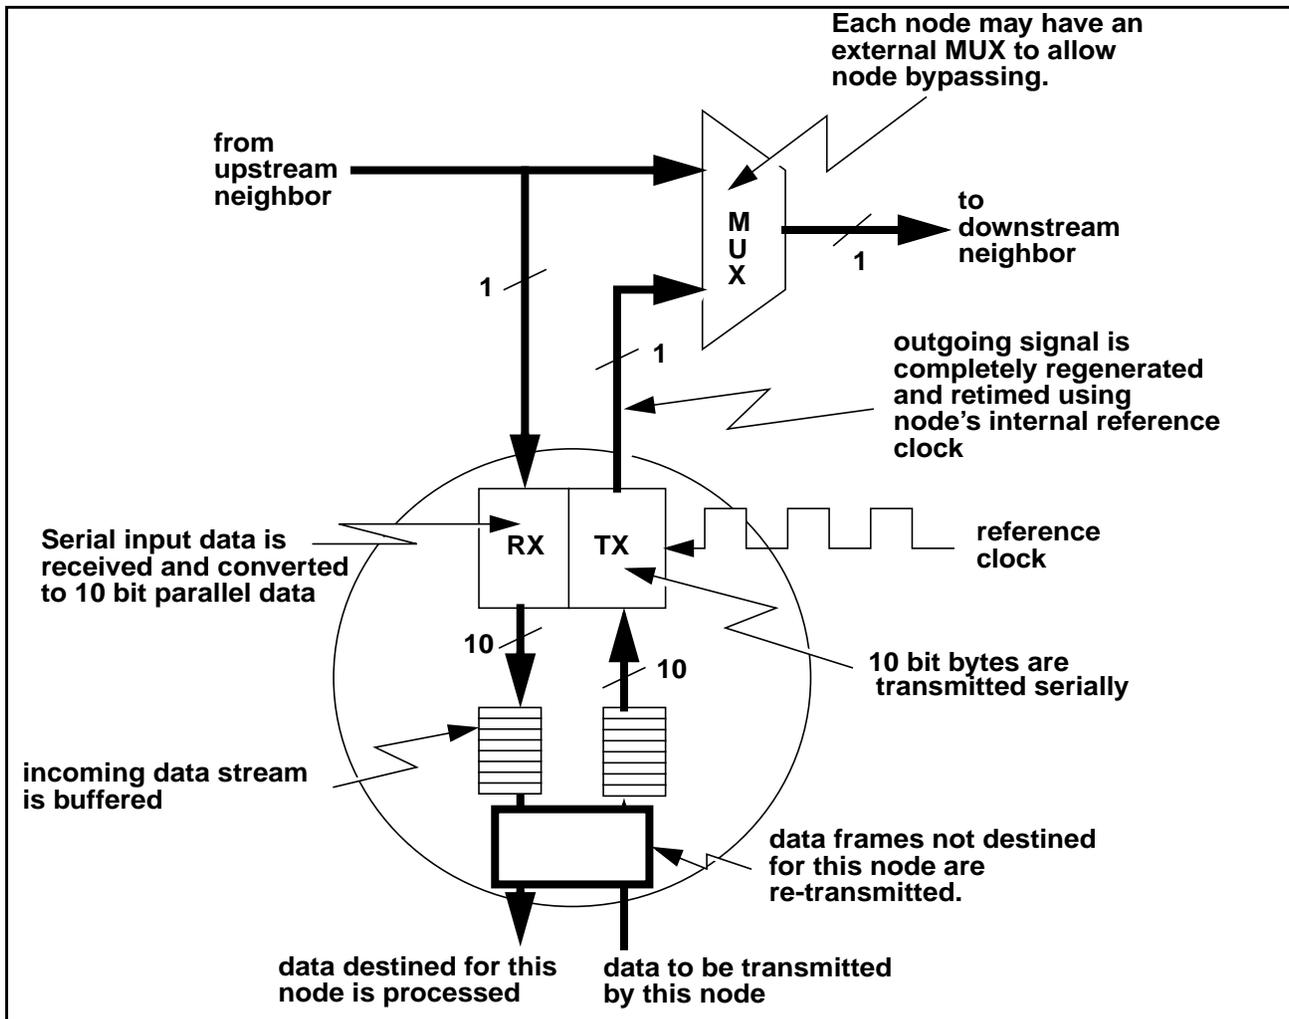
The following diagram further illustrates the fc-al loop topology showing the data flow direction through the hub and arrays.

# A5x00 hardware diagram showing direction of data flow



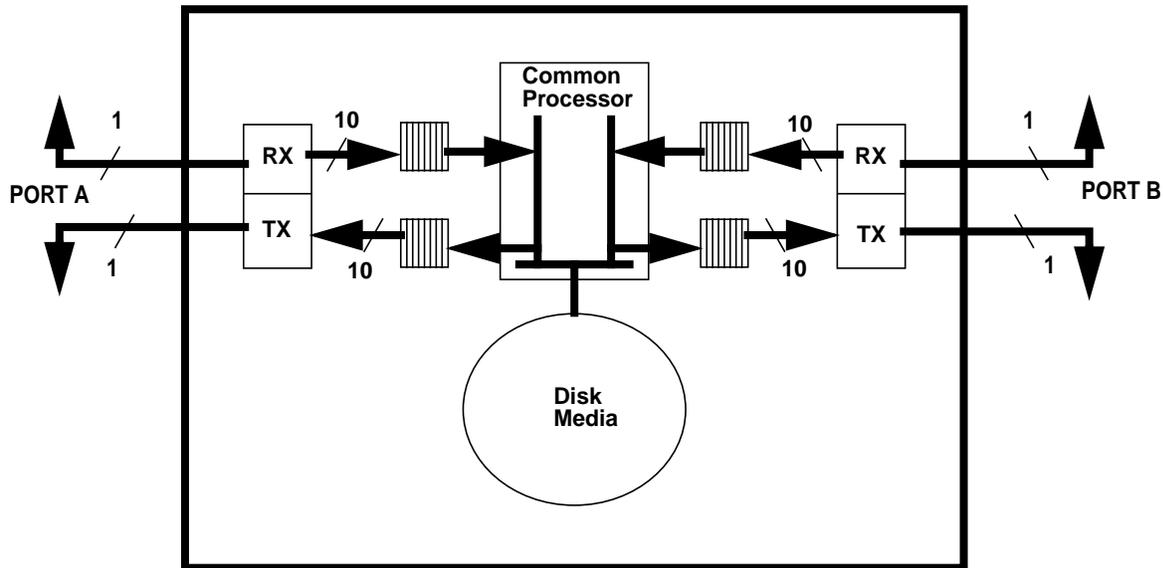


## Node fibre channel connect



Each node on the loop has an incoming and outgoing 1 Gbit/sec serial interface. The incoming signal from the upstream neighbor is received and converted to parallel 10 bit data and buffered. This incoming data is then either processed or is retransmitted depending on the destination. Each node will incur a latency penalty to allow incoming bytes to be analyzed and to allow smoothing to deal with incoming/outgoing clock skew. Note the 10 bit interfaces to the transceiver. All 8 bit bytes are encoded prior to transmission into 10 bit bytes that ensure a frequent and even number of 1 to 0 and 0 to 1 transitions. This ensures proper operation of PLL's (Phase Locked Loops). Note that the retransmitted serial data data stream is completely regenerated. Any signal degradation/jitter on the incoming signal is cleaned up by this regeneration. This has the effect of masking a signal integrity problem from the upstream neighbor.

## Basic drive architecture for fibre channel drives



### Dual port loop interactions:

- Fibre channel domains for loopA and B are independent
- Common processors services both A and B loop
- There exist some loop events that can effect the other port
- Some broken loop conditions can generate soft errors on the other loop:
  - driver recovery may require issuing a drive reset. This can cause both loops to be reinitialized.
  - excessive time servicing loop recovery on one port may lead to starvation/performance impact on other loop.
- Some catastrophic drive failures could cause failures on both loops
- Internal drive firmware resets will result in both loops resetting

### How to avoid loop interactions and resulting problems:

- Check loop integrity at installation
- Disconnect/repair broken loops as soon as possible
- Mirror or RAID. Don't make disk a single point of failure.
- Place mirror halves and RAID sub disks on separate loop pairs whenever possible
- Do not configure A and B loops to the same host adapter.

The A5x00 is a complex storage subsystem made up of many interrelated hardware, software and firmware components. The first step in troubleshooting, which could also be considered preventative maintenance, is to make sure all software, drivers and hardware firmware is up to the current revisions. Many problems may already have been detected and corrected with software patches and newer hardware firmware. The A5x00 and SSA products share a “Software/Firmware Configuration Matrix” which contains the most recent versions of patches and firmware. Always reference the latest matrix when maintaining an A5x00 array subsystem. The latest matrix or patch info can be found at the following web sites:

<http://storageweb.eng/tm/photon/main/> (Sun internal web site)

<http://sunsolve.sun.com/> - SSA/A5x00 Configuration Matrix and general patch information

The STORtools array management and diagnostic package contains an easy to use Check Revisions function which automates the matrix and storage checking.

### Troubleshooting Methodologies

Technical troubleshooting of any electronic equipment with interrelated software and hardware can be categorized into 2 main methodologies; “Bottom Up” and “Top Down”. Employment decisions for these methods are discussed below.

### Bottom Up Troubleshooting

Bottom up troubleshooting refers to the technique of checking the lowest levels of error indicators first, then using higher level loop integrity tests and FRU isolation techniques to identify faulty hardware. Progressing to disk exercisers and operating system utilities or software applications as final functional verification. Bottom Up Troubleshooting is best used when diagnosing the following types of situations:

- New hardware verification
- Solid hardware faults
- Hung fc-al loops

The following lists standard bottom up checks and tests:

- Hardware error LED indicators
- FPM error notices and information
- Manual Process of Elimination (POE) for FRU isolation
- luxadm commands, ls -l commands and the format utility and to verify correct solaris device files
- STORtools Loop Integrity tests and FRU Isolation Tool

---

## Top Down Troubleshooting

Top Down troubleshooting refers to the use of higher level software messages and log files for problem determination as a first step. Then using progressively lower level disk loop integrity tests, disk exercisers and FRU isolation until a problem is found. The main objective being to prevent or minimise any system down time. Top Down troubleshooting is best used when diagnosing the following types of problems:

- Intermittent or transient errors
- Software driver problems
- Device file and configuration problems

The tools used for initial top down troubleshooting are employed while the system is up and running. The following lists various standard methods for the A5x00 and fc-al loops:

- Look for storage related messages in /var/adm/messages and identify any suspect A5x00s and fibre loops. (STORtools message summary)
- Check revisions of software packages, patches and hardware. (STORtools Check Revisions)
- Verify correct device file paths
- Check any related software, configuration or startup files for recent changes. (STORtools Check Current Status can detect software and firmware changes)
- Analyze any core files or panics.
- Search sunsolve for any known related bugs and problems.
- Run STORtools Loop Integrity Test diagnostics to try and reproduce the problem.
- If above loop tests fail perform FRU Isolation.

When a problem is not reproducible, hardware may be suspect and replaced. Always make only one change at a time and carefully monitor results. When possible it is best to restore the original hardware before replacing another part, this eliminates the introduction of additional unknown problem sources.

The problem can generally be consider solved after the problem has not reoccurred for twice the intermittences period (if system panicked once weekly let run for two weeks).

***Use of the STORtools package is highly recommended although not required for troubleshooting problems in A5x00 and fc-al loops. Reference the STORtools User Guide with its troubleshooting flow charts for details.***

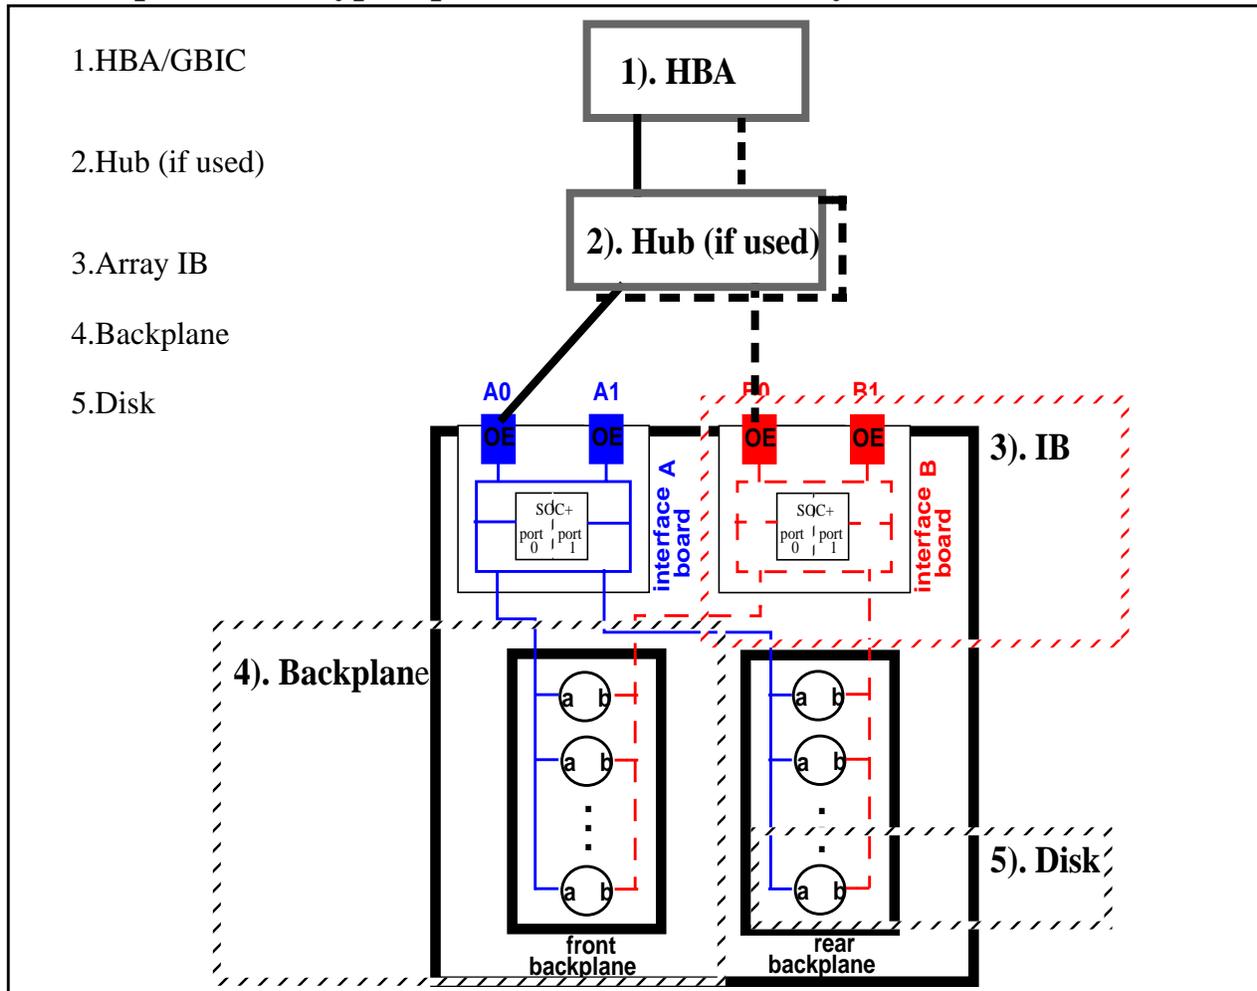


## FRU Isolation with Process Of Elimination Troubleshooting

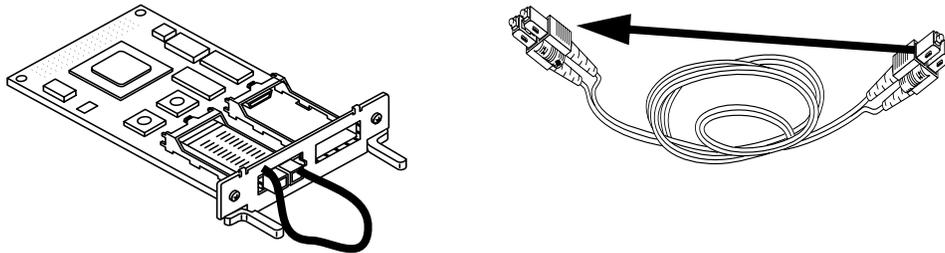
Both the bottom up and top down methodologies may employ a FRU Isolation technique using a process of elimination. Process of Elimination (POE) troubleshooting refers to the method of finding what is not causing the problem (what works) in order to locate a bad component. You start out testing the smallest possible configuration, then continue adding components until the failure is reproduced. The last component added is typically the source of trouble.

The configuration components for an A5x00 fc-al loop subsystem would be:

### 5 POE loopback and Bypass points for the A5x00 subsystem



**HBA Loop Back** - You can create a loop back cable by dismantling the ends of a cable and looping one-cable back to itself. Keep each side with one grey and one yellow connector end.



Or loop back connector plugs can be purchased from fibre cable suppliers. Below is one vendor source:

<http://www.cablexpress.com/>

Part Number	Description	Cost
12833	SM Male/Male Fiber Wrap Plug	\$16.47

**Hub Loopback** - To perform a loop back from the HBA to a hub, remove all GBICs except the one for the suspect host. This will cause the hub to send all frames back to the HBA.

**Array IB loopback** - Bypass both backplane ports for the loop under test. The other loop can remain fully operational if the array is using dual loops.

**Disk Bypass** - You can select an individual disk and bypass the disk from one loop. The disk can remain functional on the other loop.

**Array Bypasses** - There are five areas of component bypass in the array.

**GBIC** - isolates array, will break loop in a daisy-chain config

**Backplane** - removes backplane and all disks from both loops

**Backplane Port** - removes backplane and disks from one loop

**Disk** - Removes disk from both loops

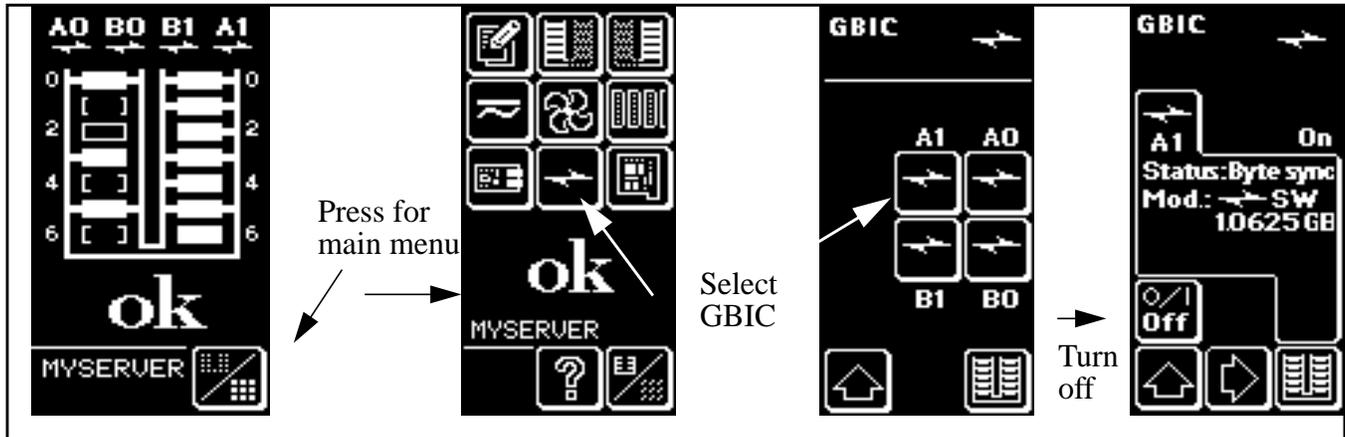
**Disk Port** - Removes disk from one loop

All bypasses are done using the array FPM, see the following examples:



To internally bypass the GBIC select the lightning bolt from the FPM main menu, select the GBIC and then press the 1/0 OFF button (see illustration below).

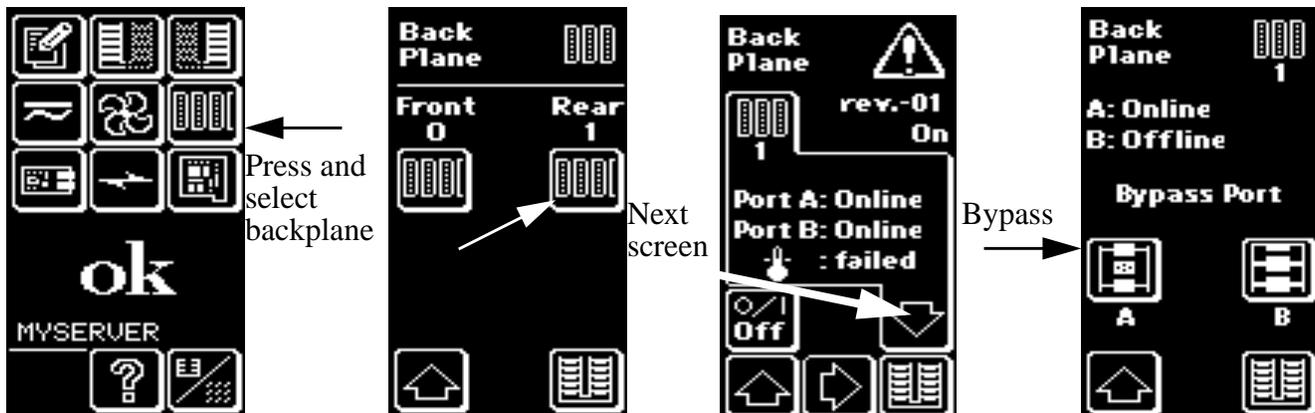
### GBIC bypass example



To loop back through the IB, bypass the IB/loop port for both backplanes.

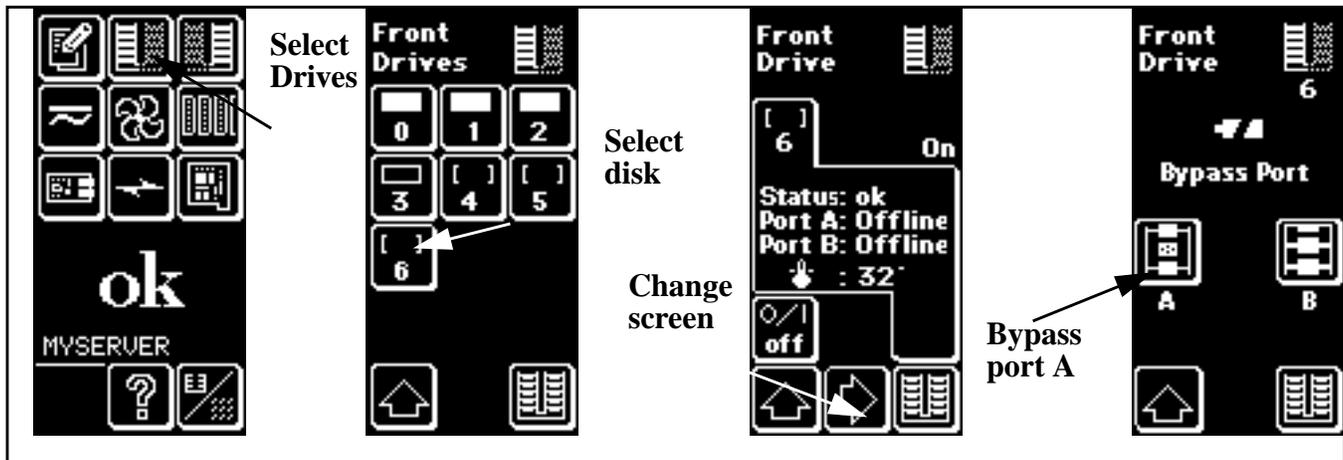
To internally bypass the backplane port select the backplane, next screen, bypass the port.

### Backplane port bypass example



To bypass a disk, select drives, select disk, change screen, bypass port.

*Disk port bypass example*



**Note** - Component bypasses are not preserved after a reset or power on. Default state after power on is all components enabled.

## Solid Failure or hung loop - bottom up troubleshooting

When the A5x00 array or fc-al loop cannot be seen or accessed from the host you should first use the hardware indicators for information.

Summary of steps:

- 1). Check FPMs (errors, node count)
- 2). Isolate array
- 3). Isolate FRU

### 1). Check FPMs

Check the array FPMs for any error conditions and resolve any reported hardware problems.

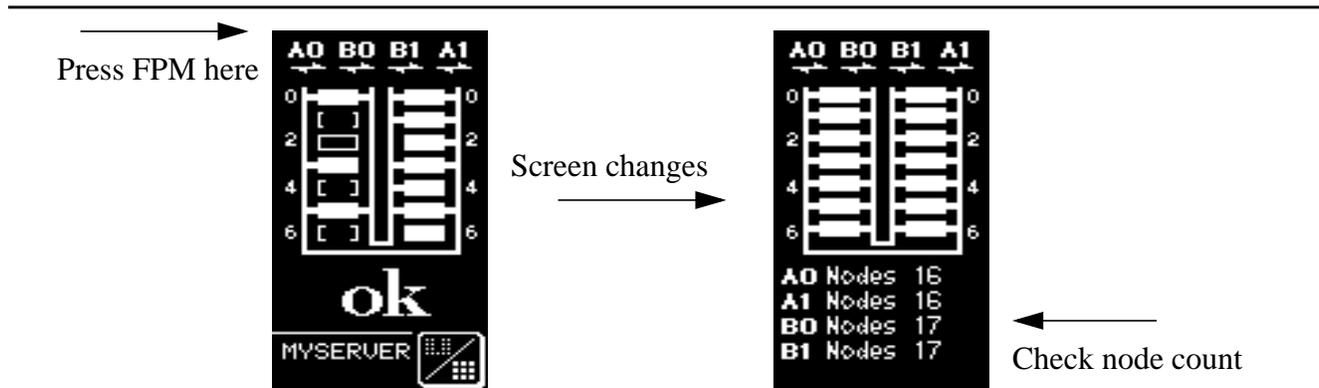
Check node count. This can be seen by pressing the top of the FPM, the node count will appear at the bottom (see illustration below).

In a healthy loop you will see a node count for both the A and B loops that reflects the actual number of nodes on the loop. All disks count as one node, ses nodes (soc+ chip ports) count as 2 (1 for split loop) and each host connection counts as 1. For example a single A5100 array with no hosts and 14 disks will have a count of 16. A 4 array cabinet with hubs will have a node count of  $16 * 4 = 64$  with no hosts attached.

If all nodes (disk and IBs) are seen the loop is probably healthy and the problem is likely in the host use the Top Down Troubleshooting steps to verify the software device files and basic loop integrity.

*A 0 node count indicates a hung loop, proceed to step 2.*

#### FPM Node check example



---

## 2). Isolate arrays

If this is a single direct attached array proceed to step 3 (your battle is half over).

For either daisy chained or hub connected arrays you need to isolate to an array.

For hubs you can use the hub as a central connection point and unplug all GBICs from the hub.

For daisy chains you will need to manually uncable the arrays from the daisy chain or use the FPM to internally bypass the GBIC, just be aware that this will break a daisy chain loop.

When all arrays are isolated (either by pulling cables/GBICs or using the FPM) recheck the node counts using the FPMs of the arrays (remember a single full array should have a count of 16 or 24). If one of the arrays has a bad count that's probably where the problem is located.

Enable the other arrays, or plug in the cables/GBICs, and recheck the node count (should now be something like  $16 * 3 = 48$ ). Plug in the host cable/GBIC and the array should come ONLINE. If the array had not previously been configured and is lacking device files you may need to perform a "luxadm insert" command on the host to create the device files.

Proceed to step 3 to diagnose the faulty array.

## 3). Isolate FRU

Now that the problem has been narrowed down to a single array we can use the array's own bypass FPM options to isolate the bad component.

Remember you can bypass, or loopback, at the GBIC, backplane and disk. We've already looped back at the IB by bypassing the GBIC with the FPM or by unplugging the cable/GBIC so let's start with the backplane.

Check the node count, if the count is now 9 (7 disks, 2 IB) then the bypassed backplane has the problem.

Turn the backplane back on and use the FPM to select then bypass each disk one at a time to see if a bad disk is causing the problem.

Check the node count each time until the bad disk is found.

If the back plane still hangs the loop with all disks bypassed then the back plane itself is bad.

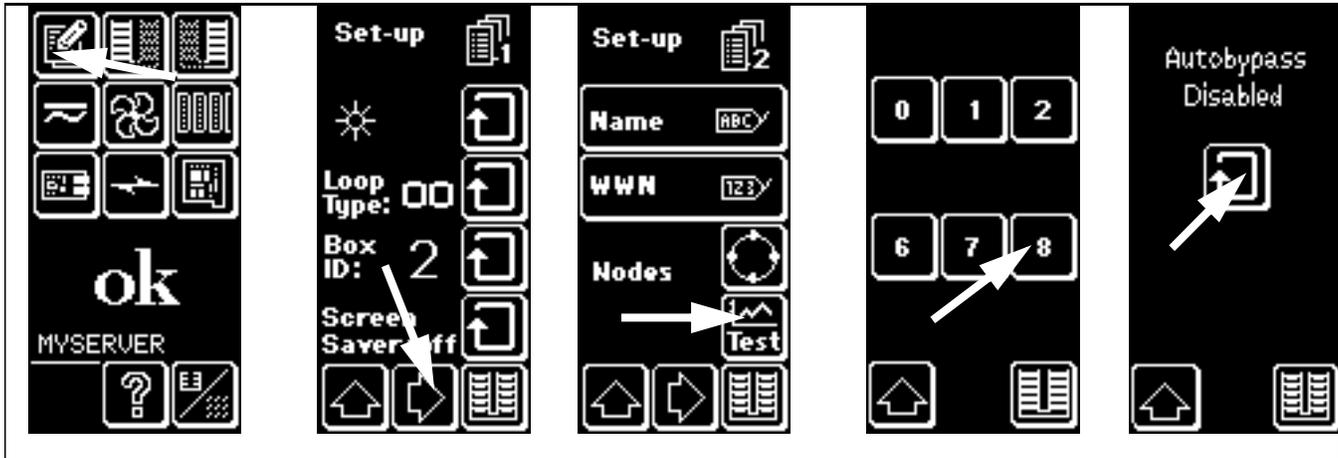
With both backplanes and all GBICs bypassed the node count should be 2 for the IB, if not then the IB itself is bad.



There is also an autobypass feature that can be enabled with the FPM. *This must be left disabled for normal operation*, but can be enabled to assist with troubleshooting.

To enable the autobypass feature, select the stsetup menu, next screen, test, 8.

### *Enable AutoBypass example*



**Note** - Array power and cooling errors are monitored by the ses nodes and will be reported with the fault LEDs on the box. These failures are not reported to the system or seen in the /var/adm/messages file. To get software notification of array power and cooling errors the system should run STORtools, symon, or customized scripts which use luxadm to obtain status.

---

# Intermittant Errors or software problems - top down troubleshooting

Summary of steps:

- 1). Identify suspect fc-al loop
- 2). Check revisions
- 3). Check device files
- 4). Loop Integrity Test
- 5). FRU Isolation

## 1). Identify suspect loop

Failing arrays and loops can be proactively detected by monitoring the messages file for excessive errors, this can be automated using the storage monitor cron setup function in STORtools. An example email message from the storage monitor appears as:

```
Date: Mon, 4 Jan 1999 16:01:02 -0800
From: STORtools@host2
To: joy.silva@Ebay, root@host2
Subject: STORtools ERROR message
```

The following log entries indicate STORedge system problems.  
Please take appropriate action.

```
ERROR: socal0: port 0 has received 93 socal OFFLINE messages
```

This example shows that “socal0: port 0” is receiving loop offlines or LIPs (loop initialization process). These offline conditions may be normal as the loop will perform a LIP for configuration and recovery purposes. However 10 or more of these in a 24 hour period warrants further investigation in the form of running the Loop Integrity Test to ensure that the loop is healthy.

The actual messages as seen in /var/adm/messages appear as:

```
Dec 1 16:44:17 host2 unix: ID[SUNWssa.socal.link.5010] socal0: port 0: Fibre Channel is OFFLINE
Dec 1 16:44:20 host2 unix: ID[SUNWssa.socal.link.6010] socal0: port 0: Fibre Channel Loop is
ONLINE
Dec 1 16:44:23 host2 unix: ID[SUNWssa.socal.link.5010] socal0: port 0: Fibre Channel is OFFLINE
Dec 1 16:44:23 host2 unix: ID[SUNWssa.socal.link.6010] socal0: port 0: Fibre Channel Loop is
ONLINE
```

To run the Loop Integrity Test you should first identify the controller number (“c” number) of the suspect loop. This is best done using the STORtools Display Configuration, Configuration Summary option.



### Example:

```
socal0      /devices/sbus@1f,0/SUNW,socal@3,0:0
sf0        Port 0 Controller c1  18 Drive(s)
           Enclosure box0 Loop A
           Enclosure box1 Loop A

socal0      /devices/sbus@1f,0/SUNW,socal@3,0:1
sf1        Port 1 Controller c2  18 Drive(s)
           Enclosure box0 Loop B
           Enclosure box1 Loop B
```

In this example “socal0: port 0” is sf0 or c1.

If the STORtools package is not available you can use the following manual procedure to identify failing loops:

- 1). Locate the sbus location of the host adapter by searching through the messages file for “socal0”.
- 2). Find the loop and array/s associated with the host adapter and port using a luxadm probe command.

### Example:

```
# grep "socal0 is" /var/adm/messages
~
Apr  2 15:27:48 host2 unix: SUNW,socal0 is /sbus@1f,0/SUNW,socal@3,0
~
# luxadm probe -p
Found
SENA          Name:box1    Node WWN:5080020000007c10
Logical Path:/dev/es/ses0
Physical Path:/devices/sbus@1f,0/SUNW,socal@3,0/sf@0,0/ses@w5080020000007c12,0:0
Logical Path:/dev/es/ses3
Physical Path:/devices/sbus@1f,0/SUNW,socal@3,0/sf@0,0/ses@w5080020000007c11,0:0
Logical Path:/dev/es/ses4
Physical Path:/devices/sbus@1f,0/SUNW,socal@3,0/sf@1,0/ses@w5080020000007c13,0:0
Logical Path:/dev/es/ses5
Physical Path:/devices/sbus@1f,0/SUNW,socal@3,0/sf@1,0/ses@w5080020000007c14,0:0
SENA          Name:box0    Node WWN:5080020000045768
Logical Path:/dev/es/ses1
Physical Path:/devices/sbus@1f,0/SUNW,socal@3,0/sf@0,0/ses@w5080020000045769,0:0
Logical Path:/dev/es/ses2
Physical Path:/devices/sbus@1f,0/SUNW,socal@3,0/sf@0,0/ses@w508002000004576a,0:0
Logical Path:/dev/es/ses6
Physical Path:/devices/sbus@1f,0/SUNW,socal@3,0/sf@1,0/ses@w508002000004576b,0:0
Logical Path:/dev/es/ses7
Physical Path:/devices/sbus@1f,0/SUNW,socal@3,0/sf@1,0/ses@w508002000004576c,0:0
```

The loop causing the OFFLINE messages is connected to the HBA at “sbus@1f,0” port 0 or “sf@0”. The effected arrays are named “box0” and “box1”. This example shows only two arrays configured, but there could have been up to four on the effected loop.

## 2). Check Revisions

After identifying the suspect loop the next step is to check for correct revisions of all hardware and software. Again STORtools is the easiest method but revisions can be verified manually as well.

Appendix B contains the procedures for manual revision procedures.

## 3). Check Device Files

The A5x00 enclosure will not respond to software commands such as luxadm if the devices files for the ses nodes are not correct. These device files can be incorrect if arrays have been removed or reconfigured without using the ‘luxadm remove’ and luxadm insert’ commands or if the system was reconfigured with a loop in error.

Symptoms of incorrect or missing ses files will depend on the problem. Worse case the luxadm commands will appear to hang while waiting for a long timeout on a bad device file, or the command may return “No Network Array enclosures found” although the array disks can be accessed and viewed in the format utility so its known that the loops are functional. An additional symptom would be ses files which refer to the wrong array as shown in the following example.

```
# luxadm probe
Found
SENA          Name:b      Node WWN:5080020000035168
  Logical Path:/dev/es/ses0
  Logical Path:/dev/es/ses1
  Logical Path:/dev/es/ses7
  Logical Path:/dev/es/ses12
  Logical Path:/dev/es/ses13
SENA          Name:c      Node WWN:508002000001e618
  Logical Path:/dev/es/ses2
  Logical Path:/dev/es/ses3
  Logical Path:/dev/es/ses14
  Logical Path:/dev/es/ses15
SENA          Name:a      Node WWN:508002000002bbb8
  Logical Path:/dev/es/ses4
  Logical Path:/dev/es/ses8
  Logical Path:/dev/es/ses9
SENA          Name:d      Node WWN:5080020000032510
  Logical Path:/dev/es/ses5
  Logical Path:/dev/es/ses6
```

**Extra ses file**

**Missing ses file**



```
Logical Path:/dev/es/ses10
Logical Path:/dev/es/ses11
```

Use of the `luxadm insert` command (or a boot -r) to resolve has no effect because the files already exist although they are incorrect. Example:

```
# luxadm insert
Please hit <enter> when you have finished adding photon(s):

Waiting for Loop Initialization to complete...
New Logical Nodes under /dev/es:
New Logical Nodes under /dev/dsk and /dev/rdisk :
```

```
#
```

**Files can be correctly recreated by removing, then creating with luxadm insert. Example:**

```
# rm /dev/es/*
# luxadm insert
Please hit <enter> when you have finished adding photon(s):
```

```
Waiting for Loop Initialization to complete...
New Logical Nodes under /dev/es:
    ses0
    ses1
    ses2
    ses3
    ses4
    ses5
    ses6
    ses7
    ses8
    ses9
    ses10
    ses11
    ses12
    ses13
    ses14
    ses15
New Logical Nodes under /dev/dsk and /dev/rdisk :
```

```
# luxadm probe
```

```
Found
```

```
SENA          Name:a      Node WWN:508002000002bbb8
  Logical Path:/dev/es/ses0
  Logical Path:/dev/es/ses1
  Logical Path:/dev/es/ses8
  Logical Path:/dev/es/ses9
```

```
SENA          Name:d      Node WWN:5080020000032510
  Logical Path:/dev/es/ses2
  Logical Path:/dev/es/ses3
```

← **Correct 4 ses files**

```
Logical Path:/dev/es/ses10
Logical Path:/dev/es/ses11
SENA      Name:b      Node WWN:5080020000035168
Logical Path:/dev/es/ses4
Logical Path:/dev/es/ses5
Logical Path:/dev/es/ses12
Logical Path:/dev/es/ses13
SENA      Name:c      Node WWN:508002000001e618
Logical Path:/dev/es/ses6
Logical Path:/dev/es/ses7
Logical Path:/dev/es/ses14
Logical Path:/dev/es/ses15
#
```

← **Correct 4 ses files**

If instance numbers need correcting the `/etc/path_to_inst` file may require modification. The system will require a reboot when the `path_to_inst` file is edited. Use great caution when modifying this file and create a backup first. Seek support if unfamiliar with `path_to_inst`. Corruption of this file can prevent a system from booting.

#### 4). Loop Integrity Test

The best way to perform loop integrity testing is with the STORtools diagnostics, but the lbf (loop back frame) diagnostic and/or disk loading can be run standalone. Currently the lbf diagnostic can only be used to test fc-al loops connected to SBus based soc+ host adapters. PCI bus based or other vendor adapters will required alternate diagnostics. Reference the STORtools User Guide for details on running loop integrity tests with STORtools. Reference Appendix A for information on running standalone lbf or alternate diagnostics.

---

**Caution** – Loop Integrity testing will effect I/O performance the degree of which depends on method used for testing. If performance impact is not acceptable or unknown the loop and possibly the arrays will need to be taken offline.

---

If the loop integrity test fails you will need to perform FRU Isolation to identify the failing hardware component.

#### 5). FRU Isolation

FRU isolation will always require the process of elimination method of troubleshooting (reference chapter 3) although this method will vary depending or whether the failure is reproduced using disk access or a loop back test (lbf). The following procedures can assist with this process (reference the flowcharts in the STORtools User Guide if using STORtools diagnostics).



---

**Caution** – FRU Isolation will always require that the loop be taken offline and may further require that the arrays be taken offline. Techniques to offline the loop and arrays will vary by environment.

---

### **FRU Isolation Steps using loop back testing:**

- 1). Attached loopback cable to Host Adapter GBIC. Test. PASS - goto step 2.  
FAIL - replace the GBIC. Test. PASS - goto step 8.  
FAIL - replace Host Adapter. Test. PASS - goto step 8.  
FAIL - possible bad replacement part, Loop back cable or I/O card.
- 2) If a hub is used remove all array GBICs. Test.  
PASS (or hub not used) - goto step 3.  
FAIL - replace hub GBIC. Test. PASS - goto step 8.  
FAIL - replace cable. Test. PASS - goto step 8.  
FAIL - replace hub. Test. PASS - goto step 8.  
FAIL - possible bad replacement part or host adapter GBIC.
- 3). For a single array proceed to step 4, else locate problem array by connecting one array to the loop at a time and retesting. When test fails on an array proceed with step 4 for that array. If multiple arrays fail, repeat step 4 for each array.
- 4). If a hub is used attached a loop back cable to the array hub GBIC. Test.  
PASS (or hub not used) - goto step 5.  
FAIL - replace GBIC. Test. PASS - goto step 8.  
FAIL - possible bad hub port or other GBIC.
- 5). Reconnect array to the loop. Bypass both backplanes. Test.  
PASS - goto step 6.  
FAIL - replace array GBIC. Test. PASS - goto step 8.  
FAIL - replace array IB. Test. PASS - goto step 8.  
FAIL - replace cable. Test. PASS - goto step 8.  
FAIL - possible bad replacement part or other GBIC.
- 6). Enable the front backplane. Test. PASS - goto step 7.  
FAIL - for each disk on the backplane do the following:  
Bypass disk. Test. FAIL - enable disk, continue to next disk.  
PASS - enable disk, replace disk. Test. PASS - goto step 8.  
FAIL - possible bad backplane, replacement disk, IB or GBIC.

---

If test fails for all or multiple disks, replace backplane. Test.  
PASS - go to step 8,  
FAIL - possible bad IB, GBIC or Interconnect Assembly.

- 7). Enable the rear backplane. Test.  
FAIL - repeat the following for each disk on the backplane:  
Bypass disk. Test. FAIL - enable disk, continue to next disk.  
PASS - enable disk, replace disk. Test. PASS - goto step 8.  
FAIL - possible bad backplane, replacement disk, IB or GBIC.

If test fails for all disks, replace backplane. Test.  
PASS - goto step 8.  
FAIL - possible bad IB, GBIC or Interconnect Assembly.

If problem can only be reproduced when both backplanes are enabled. Enable both backplanes but enable only the minimum disk configuration (14 slot min config is f3, f6, r0, r3, r6. 22 slot min config is f0, f5, f10, r0, r3, r6, r10). Test.

PASS - enable each disk until test fails then replace failing disk.  
FAIL - possible problem cause is one of the min config disks, or any other loop hardware (backplanes, IB, GBICs, Interconnect Assembly).

- 8). Reconnect and enable all hardware. Test.  
FAIL - goto step 1  
PASS - exit.

When the loop and arrays are passing the basic Loop Integrity Test. It is recommended that you further verify the loop integrity using the Disk Read Loop Test as some failures can only be produced under heavy load. Reference the User Guide for more information.

## **FRU Isolation steps with disk access testing:**

### **Multiple Array Configuration**

- 1). For a single array proceed to step 3, else locate problem array by connecting one array to the loop at a time and retesting. When test fails on an array proceed with step 3 for that array. If multiple arrays fail, repeat step 3 for each array.

If all arrays fail goto step 2.



- 2). When all arrays fail independently the problem cause is probably in a GBIC, host adapter, cable or hub (if used). Leave one array connected and replace the following hardware until the test passes.
  - Host adapter to hub(if used) GBICs.
  - Hub(if used) to array GBICs.
  - Hub and or Cable/s.
  - Array IB.
  - Host adapter.

Test PASS - goto step 7.

Test FAIL - possible bad replacement parts, or array failure. Repeat step 2 with alternate array.

### **Single Array Configuration**

- 3). Bypass the front backplane. Test. PASS goto step 4.  
FAIL - enable front backplane, bypass rear backplane. Test.  
PASS - goto step 6.  
FAIL - both backplanes failing indicates possible problem  
cause is not in the backplanes and disks, goto step 5.
- 4). Enable the front backplane and disable the rear backplane. Test.  
PASS - goto step 5.  
FAIL - goto step 6.
- 5). When both backplanes either pass or fail independently the problem cause may not be in the backplanes and disks or there may be multiple problems.

For a single array configuration replace the following hardware until the test passes.

- Host Adapter GBIC.
- Array GBIC.
- Cable.
- Array IB.
- Host adapter.
- Array Interconnect Assembly.

For a multiple array configuration replace the following until the test passes (note that the host adapter/GBIC and hub/GBIC have already been tested).

- Array to hub GBICs.
- Cable.

- Array IB.
- Hub (could be a bad port on the hub).
- Array Interconnect Assembly.

If test continues to fail there may be multiple problems in the backplanes and or disks. Try testing with only the minimum configuration of disks enabled (14 slot min config is f3, f6, r0, r3, r6. 22 slot min config is f0, f5, f10, r0, r3, r6, r10).

If test continues to fail possible problem cause is one of the min config disks, or any other loop hardware (backplanes, IB, GBICs, Interconnect Assembly).

- 6). When the Test passes with one backplane bypassed, that backplane contains the problem. To further diagnose enable the suspect backplane and do the following for each disk:
  - Bypass disk. Test. FAIL - enable disk, continue to next disk.
  - PASS - enable disk and replace. Test. PASS - goto step 7.
  - FAIL - possible bad backplane, replacement disk, IB or GBIC.

If test fails for all disks, replace backplane. Test.  
PASS - goto step 7.  
FAIL - possible bad IB, GBIC or Interconnect Assembly.
- 7). Reconnect and enable all hardware elements. Test. PASS - exit.  
FAIL - goto step 1.



### CRC Error Analysis

Single bit errors inside fc-al frames are detected by CRC (cyclic redundancy check). The associated error message logged will depend on the type of frame:

#### Read data - CRC error. Data traveling from the disk to the host.

```
Feb 16 12:30:32 area51-29 unix: WARNING: /sbus@2,0/SUNW,socal@1,0/sf@0,0 (sf0):  
Feb 16 12:30:32 area51-29 unix:      Transport error on target=0x65: Fibre Channel CRC Error on frames
```

#### Write data - scsi parity error . Data traveling from the host to the disk.

```
Feb 16 23:49:13 area51-29 unix: WARNING: /sbus@6,0/SUNW,socal@d,10000/sf@0,0/  
ssd@w21000020370ee043,0 (ssd162):  
Feb 16 23:49:13 area51-29 unix:      Error for Command: write(10)          Error Level: Retryable  
Feb 16 23:49:13 area51-29 unix:      Requested Block: 433410          Error Block: 433410  
Feb 16 23:49:13 area51-29 unix:      Vendor: SEAGATE                Serial Number: 9824V60670  
Feb 16 23:49:13 area51-29 unix:      Sense Key: Aborted Command  
Feb 16 23:49:13 area51-29 unix:      ASC: 0x47 (scsi parity error), ASCQ: 0x0, FRU: 0x3
```

#### Command - timeout . Data traveling from the host to the disk.

```
Feb 16 12:30:27 area51-29 unix: WARNING: /sbus@2,0/SUNW,socal@1,0/sf@0,0/  
ssd@w21000020370ee100,0 (ssd12):  
Feb 16 12:30:27 area51-29 unix:      SCSI transport failed: reason `timeout`: retrying command
```

All of these messages generally indicate a loop integrity problem if excessive (more than 10 a day).

---

**Note** - The direction and originator of the data is important to understand in order to pin point the error source.

---



To locate the error source first check to see if all the errors are logged by a single loop node, if this is the case that device is most likely the problem.

If the errors come from all the nodes in an array or array backplane the problem might be caused by a bad IB, backplane, GBIC or disk. Careful mapping of the errors with an understanding of where the data originates can help to pin point the source of the errors. Use STORtools message summary to map the error locations. Or manually examine the /var/adm/messages file and note all nodes logging CRC and scsi parity errors.

### STORtools Message Summary Report

This example shows excessive CRC errors logged by all disks in the array “a” and front backplane disks in array “b”.

```

Controller
-----
/devices/sbus@6,0/SUNW,socal@d,10000:0 (c3)      Socal          sf  OFFLINE  CRC
/devices/sbus@a,0/SUNW,socal@d,10000:0 (c4)      socal1: port 0 sf2
                                                socal2: port 0 sf4

```

```

Disks      Name          ssd      offline  WARNING  CRC WARN
-----
c3t0d0     a,f0          ssd76    0        3        2
c3t100d0   d,f4          ssd80    0        0        0
c3t101d0   d,f5          ssd74    0        0        0
c3t102d0   d,f6          ssd58    0        0        0
c3t112d0   d,r0          ssd29    0        0        0
c3t113d0   d,r1          ssd47    0        0        0
c3t114d0   d,r2          ssd56    0        0        0
c3t115d0   d,r3          ssd82    0        0        0
c3t116d0   d,r4          ssd37    0        0        0
c3t117d0   d,r5          ssd72    0        0        0
c3t118d0   d,r6          ssd70    0        0        0
c3t16d0     a,r0          ssd41    0        1        1
c3t17d0     a,r1          ssd40    0        0        0
c3t18d0     a,r2          ssd52    0        3        2
c3t19d0     a,r3          ssd55    0        1        1
c3t1d0      a,f1          ssd46    0        0        0
c3t20d0     a,r4          ssd33    0        0        0
c3t22d0     a,r6          ssd64    0        1        1
c3t2d0      a,f2          ssd69    0        3        3
c3t32d0     b,f0          ssd159   0        0        0
c3t33d0     b,f1          ssd160   0        1        1
c3t34d0     b,f2          ssd161   0        1        1
c3t35d0     b,f3          ssd157   0        1        1
c3t36d0     b,f4          ssd164   0        3        3
c3t37d0     b,f5          ssd168   0        1        1

```



---

```
c3t38d0    b,f6          ssd156    0          1          1
c3t3d0     a,f3          ssd49     0          2          2
c3t48d0    b,r0         ssd166    0          0          0
c3t49d0    b,r1         ssd163    0          0          0
c3t4d0     a,f4         ssd28     0          0          0
c3t50d0    b,r2         ssd158    0          0          0
c3t51d0    b,r3         ssd169    0          0          0
c3t52d0    b,r4         ssd167    0          0          0
c3t53d0    b,r5         ssd162    0          0          0
c3t54d0    b,r6         ssd165    0          0          0
c3t5d0     a,f5         ssd35     0          0          0
c3t64d0    c,f0         ssd50     0          0          0
c3t65d0    c,f1         ssd48     0          0          0
c3t66d0    c,f2         ssd31     0          0          0
c3t67d0    c,f3         ssd45     0          0          0
c3t68d0    c,f4         ssd73     0          0          0
c3t69d0    c,f5         ssd42     0          0          0
c3t6d0     a,f6          ssd60     0          1          1
c3t70d0    c,f6         ssd78     0          0          0
c3t80d0    c,r0         ssd44     0          0          0
c3t81d0    c,r1         ssd62     0          0          0
c3t82d0    c,r2         ssd77     0          0          0
c3t83d0    c,r3         ssd65     0          0          0
c3t84d0    c,r4         ssd57     0          0          0
c3t85d0    c,r5         ssd59     0          0          0
c3t86d0    c,r6         ssd54     0          0          0
c3t96d0    d,f0         ssd61     0          0          0
c3t97d0    d,f1         ssd53     0          0          0
~ example cut ~
```

Also check the messages file for any scsi parity errors and command timeouts and map those locations. Example:

```
# grep parity /var/adm/messages
Feb 16 23:49:13 area51-29 unix:          ASC: 0x47 (scsi parity error), ASCQ: 0x0, FRU: 0x3

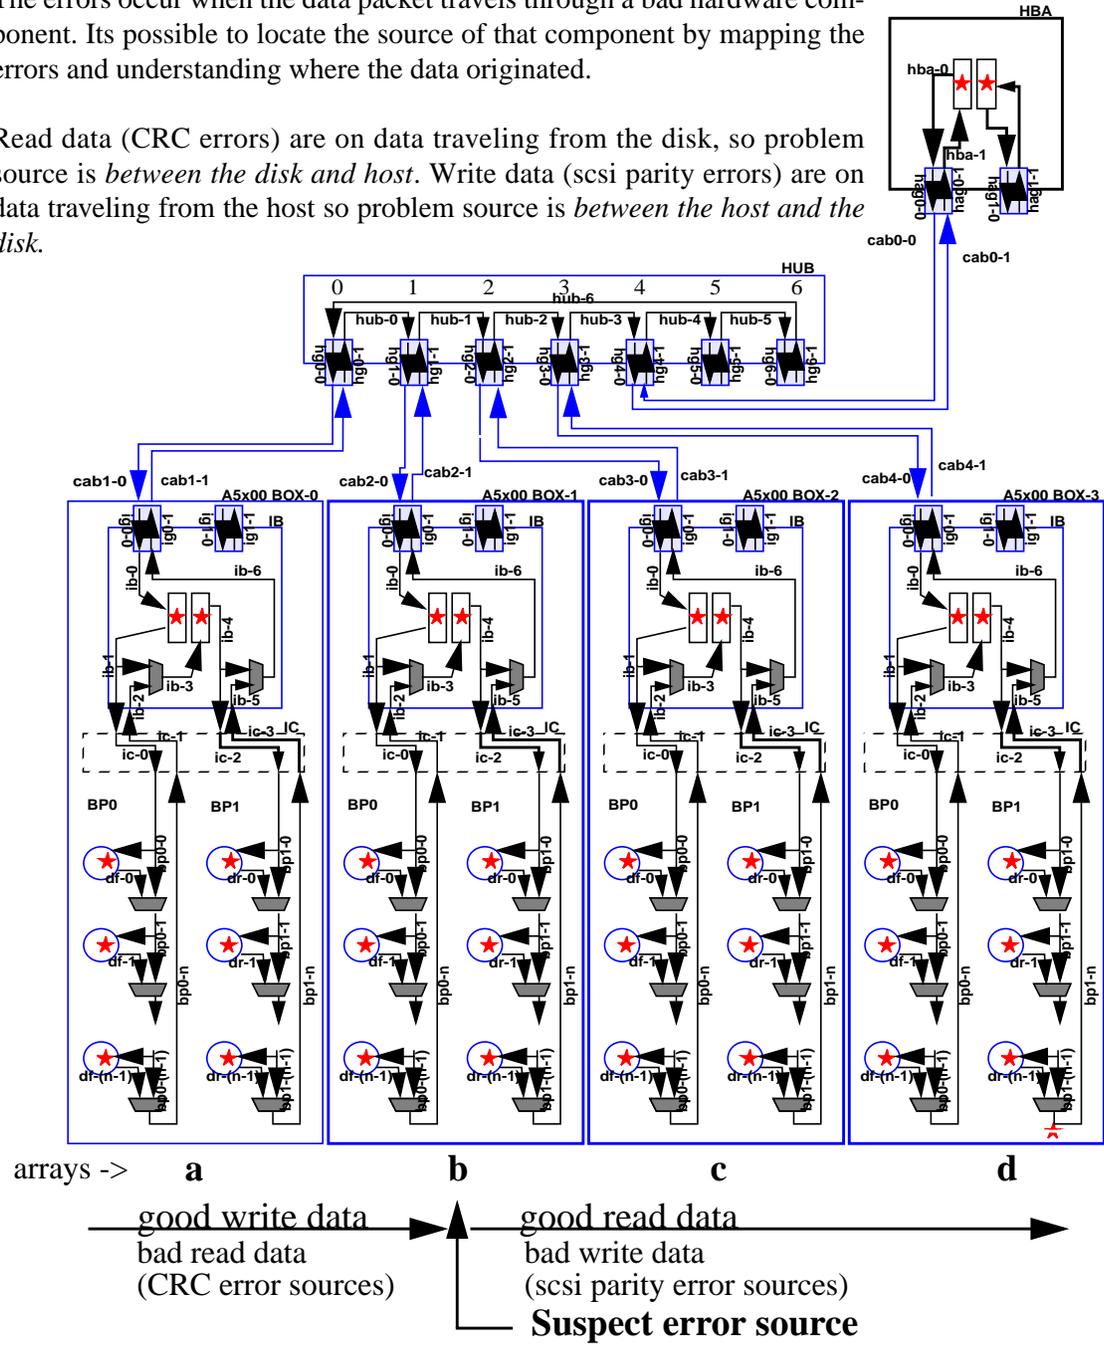
# grep timeout /var/adm/messages
Feb 17 12:41:05 area51-29 unix:          SCSI transport failed: reason 'timeout': retrying command
Feb 17 12:52:12 area51-29 unix:          SCSI transport failed: reason 'timeout': retrying command
```

You will need to examine the full message to locate the error source for these errors. Future STOR-tools message summaries may include checks for these errors as well and will sort by array and slot to help make this process easier. Note that all the nodes getting errors (underlined in the example) are located in the array name “a” both backplanes and the array named “b” front backplane only. Only one scsi parity error occurred and it was at b,r1 (array b rear slot 1).

The next step is to map these errors into the suspect fc-al and array/s configuration. The following example uses the diagram show earlier of a 4 array fc-al loop.

The errors occur when the data packet travels through a bad hardware component. Its possible to locate the source of that component by mapping the errors and understanding where the data originated.

Read data (CRC errors) are on data traveling from the disk, so problem source is *between the disk and host*. Write data (scsi parity errors) are on data traveling from the host so problem source is *between the host and the disk*.





The problem source for this example was a bad front backplane in array named “b”. Although the problem could have been a disk, the IB or interconnect assembly. Also this problem was resolved using loop integrity testing and POE troubleshooting.

## Read Data Link Status (RDLS)

The luxadm expert mode command “rdls” will read the loop nodes Link Error Status Block. Use of this information for troubleshooting purposes has not been fully investigated but may be used to assist in identifying a suspect array on a problem loop. Just be aware that the error source is usually upstream from the sync loss symptom. Higher sync loss counts are normal for disks installed in the first backplane slot, probably due to signal noise while crossing the interconnect assembly.

Reference page 271 of the Fibre Channel Physical and Signaling Interface(FC-PH) Specification for more information.

**The following example was generated on a 4 array loop with a bad GBIC in the second array.**

```
# luxadm -e rdls a | more
```

Link Error Status information for loop: /devices/sbus@6.0/SUNW,socal@d,10000:0

al_pa	lnk fail	sync loss	signal loss	sequence err	invalid word	CRC
2	123456	7866121	7867562	4750704	8675309	123456
d2	0	0	0	0	0	0
ef	0	101	0	0	0	0
e8	0	0	0	0	0	0
e4	0	0	0	0	0	0
e2	0	0	0	0	0	0
e1	0	0	0	0	0	0
e0	0	0	0	0	0	0
dc	0	0	0	0	0	0
b5	0	0	0	0	0	0
cd	0	3088	0	0	0	0
cc	0	44	0	0	0	0
cb	0	0	0	0	0	0
ca	0	0	0	0	0	0
c9	0	0	0	0	0	0
c6	0	0	0	0	0	0
9e	0	0	0	0	0	0
<b>b2</b>	<b>4</b>	<b>2092701</b>	<b>0</b>	<b>0</b>	<b>0</b>	<b>0</b>
b1	3	11609	0	0	0	0
ae	3	21651	0	0	0	0
ad	3	22652	0	0	0	0
ac	3	4479	0	0	0	0
ab	3	1528	0	0	0	0
aa	3	4675	0	0	0	0



75	0	0	0	0	0	0	0
98	4	139450265	0	0	0	0	0
97	3	316171	0	0	0	0	0
90	3	2073	0	0	0	0	0
8f	0	0	0	0	0	0	0
88	0	0	0	0	0	0	0
84	3	1280	0	0	0	0	1
82	0	0	0	0	0	0	0
5a	0	0	0	0	0	0	0
72	0	29	0	0	0	0	0
71	0	0	0	0	0	0	0
6e	0	0	0	0	0	0	0
6d	0	0	0	0	0	0	0
6c	0	0	0	0	0	0	0
6b	0	0	0	0	0	0	0
6a	0	0	0	0	0	0	0
45	0	0	0	0	0	0	0
55	0	32	0	0	0	0	0
54	0	0	0	0	0	0	0
53	0	0	0	0	0	0	0
52	0	0	0	0	0	0	0
51	0	0	0	0	0	0	0
4e	0	0	0	0	0	0	0
4d	0	0	0	0	0	0	0
29	0	0	0	0	0	0	0
3a	0	32	0	0	0	0	0
39	0	0	0	0	0	0	0
36	0	0	0	0	0	0	0
35	0	0	0	0	0	0	0
34	0	0	0	0	0	0	0
33	0	0	0	0	0	0	0
32	0	0	0	0	0	0	0
1	0	0	0	0	0	0	0
25	0	323	0	0	0	0	0
23	0	0	0	0	0	0	0
1f	0	0	0	0	0	0	0
1e	0	0	0	0	0	0	0
1d	0	0	0	0	0	0	0
1b	0	0	0	0	0	0	0
18	0	0	0	0	0	0	0

array "c"

array "d"

NOTE: These LESB counts are not cleared by a reset, only power cycles. These counts must be compared to previously read counts.

\*This example was generated on a system running older drivers that had false numbers for the host adapter, this is fixed in the current driver release (6). Seagate 9GB7200rpm barracuda disks do not report invalid word counts. The new 10Krpm cheetah disks do, although no troubleshooting information has been investigated using the invalid word counts.



In the above example the problem was caused by a bad GBIC receiver on the array name “d”. But the problem could have been caused by a bad GBIC transmitter on the upstream array “a”. Bad IBs, fibre cables, hub GBICs or disks could also have caused similar symptoms. The main thing to remember is the problem source of sync loss is upstream of the symptom.

Use of the ‘luxadm -e dump\_map’ command or the STORtools display configuration, rawdump, can help map the al\_pa (arbitrated loop physical address) to the array. Don’t assume the arrays are in order as shown here. Array cables can be easily switched at the hub changing the loop order.

### Example luxadm -e dump\_map display:

```
# luxadm -e dump_map a|more
Pos AL_PA ID Hard_Addr Port WWN Node WWN Type
0 2 7c 2 20020800207e147d 200c0800207e147d 0xff (Unknown Type)
1 d2 d d2 508002000002bbb9 508002000002bbb8 0xd (SES device)
2 ef 0 ef 210000203713121a 200000203713121a 0x0 (Disk device)
3 e8 1 e8 2100002037131529 2000002037131529 0x0 (Disk device)
4 e4 2 e4 2100002037131098 2000002037131098 0x0 (Disk device)
5 e2 3 e2 210000203713126e 200000203713126e 0x0 (Disk device)
6 e1 4 e1 21000020370ed14a 20000020370ed14a 0x0 (Disk device)
7 e0 5 e0 21000020371313e6 20000020371313e6 0x0 (Disk device)
8 dc 6 dc 21000020371311d4 20000020371311d4 0x0 (Disk device)
9 b5 1d b5 508002000002bbba 508002000002bbb8 0xd (SES device)
10 cd 10 cd 21000020370ee53b 20000020370ee53b 0x0 (Disk device)
11 cc 11 cc 21000020371315c6 20000020371315c6 0x0 (Disk device)
12 cb 12 cb 210000203713176a 200000203713176a 0x0 (Disk device)
13 ca 13 ca 210000203713172b 200000203713172b 0x0 (Disk device)
14 c9 14 c9 2100002037131543 2000002037131543 0x0 (Disk device)
15 c6 16 c6 21000020370e6c7f 20000020370e6c7f 0x0 (Disk device)
16 9e 2d 9e 5080020000025ac1 5080020000025ac0 0xd (SES device)
17 b2 20 b2 21000020370ee2bb 20000020370ee2bb 0x0 (Disk device)
18 b1 21 b1 21000020370ee060 20000020370ee060 0x0 (Disk device)
19 ae 22 ae 21000020370ee100 20000020370ee100 0x0 (Disk device)
20 ad 23 ad 21000020370ee07a 20000020370ee07a 0x0 (Disk device)
21 ac 24 ac 21000020370edce9 20000020370edce9 0x0 (Disk device)
22 ab 25 ab 21000020370ee114 20000020370ee114 0x0 (Disk device)
23 aa 26 aa 21000020370edf7b 20000020370edf7b 0x0 (Disk device)
24 75 3d 75 5080020000025ac2 5080020000025ac0 0xd (SES device)
25 98 30 98 21000020370edf89 20000020370edf89 0x0 (Disk device)
26 97 31 97 21000020370edea5 20000020370edea5 0x0 (Disk device)
27 90 32 90 21000020370ee0fb 20000020370ee0fb 0x0 (Disk device)
28 8f 33 8f 21000020370ee099 20000020370ee099 0x0 (Disk device)
29 88 34 88 21000020370eded3 20000020370eded3 0x0 (Disk device)
30 84 35 84 21000020370ee043 20000020370ee043 0x0 (Disk device)
31 82 36 82 21000020370edfc8 20000020370edfc8 0x0 (Disk device)
32 5a 4d 5a 508002000001e619 508002000001e618 0xd (SES device)
~ Example cut ~
```



## Example of STORtools raw dump display:

Name	Enclosure	WWN	ALPA	ID	ses	Physical Path
a,A0	ses0	508002000002bbb8	d2	d	ses68	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ses@w508002000002bbb9,0:0						
a,A1	ses1	508002000002bbb8	b5	1d	ses69	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ses@w508002000002bbba,0:0						
b,A0	ses6	5080020000025ac0	9e	2d	ses86	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ses@w5080020000025ac1,0:0						
b,A1	ses7	5080020000025ac0	75	3d	ses87	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ses@w5080020000025ac2,0:0						
c,A0	ses4	508002000001e618	5a	4d	ses74	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ses@w508002000001e619,0:0						
c,A1	ses5	508002000001e618	45	5d	ses75	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ses@w508002000001e61a,0:0						
d,A0	ses2	5080020000032510	29	6d	ses70	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ses@w5080020000032511,0:0						
d,A1	ses3	5080020000032510	1	7d	ses71	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ses@w5080020000032512,0:0						

Name	Disks	WWN	ALPA	ID	ssd	Physical Path
a,f0	c3t0d0	200000203713121a	ef	0	ssd76	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w210000203713121a,0:a,raw						
a,f1	c3t1d0	2000002037131529	e8	1	ssd46	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w2100002037131529,0:a,raw						
a,f2	c3t2d0	2000002037131098	e4	2	ssd69	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w2100002037131098,0:a,raw						
a,f3	c3t3d0	200000203713126e	e2	3	ssd49	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w210000203713126e,0:a,raw						
a,f4	c3t4d0	20000020370ed14a	e1	4	ssd28	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w21000020370ed14a,0:a,raw						
a,f5	c3t5d0	20000020371313e6	e0	5	ssd35	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w21000020371313e6,0:a,raw						
a,f6	c3t6d0	20000020371311d4	dc	6	ssd60	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w21000020371311d4,0:a,raw						
a,r0	c3t16d0	20000020370ee53b	cd	10	ssd41	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w21000020370ee53b,0:a,raw						
a,r1	c3t17d0	20000020371315c6	cc	11	ssd40	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w21000020371315c6,0:a,raw						
a,r2	c3t18d0	200000203713176a	cb	12	ssd52	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w210000203713176a,0:a,raw						
a,r3	c3t19d0	200000203713172b	ca	13	ssd55	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w210000203713172b,0:a,raw						
a,r4	c3t20d0	2000002037131543	c9	14	ssd33	/devices/sbus@6,0/SUNW,socal@d,10000/sf@0,0
/ssd@w2100002037131543,0:a,raw						

~ Example cut ~

## ERROR/WARNING Messages Interpretation

All drivers generate information, notifications and warnings that get logged into the file `/var/adm/messages`. These are informational and can be used to determine the state of one's system. Specific drivers and even applications determine the content and format of what gets logged. These messages also can be instrumental in aiding diagnosis of failures/events and create a historical record for the system. In order to be useful for diagnosis, many of the messages may appear extraneous and/or cryptic. An attempt is made here to describe those WARNINGS, NOTIFICATIONS and MESSAGES that are generated that pertain to A5x00 subsystems, how to interpret and how to react to them.

The fibre channel protocol and the host drivers are designed to be fairly robust. Occasionally warnings and/or messages are generated to the console that do not indicate failures but do tend to cause alarm for users. Nearly all peripherals perform internal retries often without generating any output. Disk drive firmware has fairly complex retry algorithms that quietly retry failures, only reporting an actual failure when retry counts are exhausted. Our driver philosophy has been to generate these warning messages so that diagnosis of real problems may be facilitated.

The bottom line is that WARNING messages while often indicating a failure condition are not always cause for alarm.

**The following are some common categories of messages with explanations:**

### MESSAGES

Messages are informational only and do not imply a failure condition. Messages are sent to the console without any preface (i.e. WARNING, FATAL ERROR, etc.)

For example:

```
Feb 23 13:19:52 MYSERVER unix: <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
Feb 23 13:19:52 MYSERVER unix: ssd9 at sf1: name w21000020370b9195,0, bus address c6
Feb 23 13:19:52 MYSERVER unix: ssd9 is /sbus@6,0/SUNW,socal@1,0/sf@1,0/ssd@w21000020370b9195,0
```

### NOTICES

NOTICES usually indicate some change in condition, usually not a good change.

For example:

```
Dec 29 20:03:13 C-BFIS21 unix: NOTICE: vxdmp: Path failure on 118/211
```

This NOTICE indicates that one of the paths to vxdmp device 118/211 is no longer available. NOTICES however may be informational only:

```
Dec 23 21:35:56 C-BFIS21 unix: NOTICE: lane0: will attempt to start ELAN in 13 sec
Dec 23 21:35:56 C-BFIS21 unix: q93b on ba0: Data Link Up
Dec 23 21:35:57 C-BFIS21 unix: NOTICE: lane0: joined ELAN
```

An example of an informational notice indicating driver progress.

## WARNINGS

WARNINGS may or may not indicate a failure. Typically retry logic takes care of the problem while some WARNINGS do indicate some sort of failure event. Warning messages are prefaced at the console with the keyword “WARNING”. Examples of “fatal” WARNINGS:

```
Nov 21 21:20:18 MYSERVER unix: WARNING: /sbus@6,0/SUNW,socal@2,0/sf@1,0 (sf7):
Nov 21 21:20:18 MYSERVER unix:   Offline Timeout
```

This warning indicates that the link (sf7) went offline and failed to return to an online state. Examples of “non-fatal” WARNINGS:

```
Feb 23 14:07:21 momoney unix: WARNING: /sbus@6,0/SUNW,socal@1,0/sf@1,0/ssd@w21000020370bb75e,0
(ssd80):
Feb 23 14:07:21 momoney unix:   SCSI transport failed: reason 'timeout': retrying command
```

This particular warning indicates a command timeout that will be retried.

## OFFLINE/ONLINE sequences

The fibre channel loops may from time to time get re-initialized causing service to the loop to momentarily suspended during this initialization. All outstanding commands on this particular loop will be automatically retried as soon as the loop initialization is complete and normal operation will resume.

### Common causes of OFFLINE/ONLINE (loop re-initialization)

- Soft or hard addition/removal of a device on the loop.
- Power cycle of device on the loop.
- Forced loop-init by driver recovery algorithms.
- A box reset following a download.
- Temporary loss of sync on the loop.



## Fibre Channel Messages

Fibre channel medium has many similarities to ethernet although for the A5x00 it is employed for the purpose of an I/O channel. Fibre Channel loops are specified to have a BER (bit error rate) less than 10E-12. Actual BER is better than 10E-13 and may be as clean as 10E-15. It is normal to occasionally experience a bit error that results in a corrupted frame. One difference between ethernet protocol and fibre channel protocol for SCSI (FCP\_SCSI) is that in FCP-SCSI, when lost frames are discarded, the entire command request is retried rather than retransmit individual frames at the lower levels. A difference between the A5x00 fibre channel messaging and ethernet messaging is that the A5x00 will generate a message for every dropped frame. The end result is a potentially noisy message file whenever links are generating higher than normal CRC errors. This unfortunately means that messages resulting from CRC errors must be quantified before determining what if any maintenance action is required. STORtools diagnostics do support message summary capabilities to aid in this quantification.

It is important to note that the reporting of these soft errors not necessarily be interpreted as system problems. Commands affected by dropped frames due to CRC errors are retried by the drivers. Statistically data frames are usually lost and it is rare to experience CRC errors on the retried command.

Running at 100Mb/s it is conceivably possible to generate a CRC error every 30 minutes and still meet the fibre channel minimum specification of 10E-12! Since however the normal operating BER is measured to be much higher, more than a few a day indicates a potential margin problem and should be diagnosed and repaired in order to decrease the likelihood that the loop will fail due to further degradation in the future. The urgency to repair loops exhibiting CRC errors is dependent on the frequency of the errors, whether or not dual porting is employed and whether or not some form of RAID or mirroring is employed. Typically, even if a loop completely fails, there still must be a second path failure AND a second subdisk failure before data can actually be lost.

The following selected messages are fibre channel loop related and are treated separately from the normal disk errors.

### **WARNING: ASC: 0x47 (scsi parity error)**

```
Sep 29 06:54:57 MSERVER unix: WARNING: /sbus@a,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037078057,0
(ssd709):
Sep 29 06:54:57 MSERVER unix:      Error for Command: write(10)      Error Level: Retryable
Sep 29 06:54:57 MSERVER unix:      Requested Block: 93671      Error Block: 93671
Sep 29 06:54:57 MSERVER unix:      Vendor: SEAGATE              Serial Number: 9735J27064
Sep 29 06:54:57 MSERVER unix:      Sense Key: Aborted Command
Sep 29 06:54:57 MSERVER unix:      ASC: 0x47 (scsi parity error), ASCQ: 0x0, FRU: 0x3
```

**Interpretation:** Indicates a drive discarded a data frame it received for a write command due to a

detected fibre channel frame CRC error. Although this appears to be a drive message, it usually indicates marginal signal quality for this particular fibre channel loop.

- **Probable Cause:** Fibre channel loop margin problem (GBIC, IB, BP, HUB, HA, IC, cables)
- **Error Level:** Retryable, Fatal after 3 unsuccessful retries (very unlikely to ever happen).
- **Action:**
  - <11 (retryable) per day per disk - Monitor only. Run Loop Integrity test when convenient.
  - >10 (retryable) per day per disk OR > 5 per day per loop for 3 consecutive days - Schedule loop integrity testing/FRU isolation. May be advisable to disconnect loop to avoid performance degradation or interference on other loop.
  - > 0 (fatal) - Schedule FRU isolation.

## **WARNING: Fibre Channel CRC Error**

```
Oct 4 23:16:45 MSERVER unix: WARNING: /sbus@e,0/SUNW,socal@2,0/sf@0,0 (sf14):
Oct 4 23:16:45 MSERVER unix: Transport error on target=0x51: Fibre Channel CRC Error on
frames
Oct 4 23:16:45 MSERVER unix: WARNING: /sbus@e,0/SUNW,socal@2,0/sf@0,0/ssd@w22000020370bef82,0
(ssd780):
```

- **Interpretation:** Indicates the sbus host adapter discarded a data frame it received for a read command due to a detected fibre channel frame CRC error. Although this appears to be a drive message, it usually indicates marginal signal quality for this particular fibre channel loop.
- **Probable Cause:** Fibre channel loop margin problem (GBIC, IB, BP, HUB, HA, IC, cables).
- **Error Level:** Retryable, Fatal after 3 unsuccessful retries (very unlikely to ever happen).
- **Action:**
  - <11 (retryable) per day per loop - Monitor only. Run Loop Integrity test when convenient.
  - >10 (retryable) per day per loop OR > 5 per day per loop for 3 consecutive days - Schedule loop integrity testing/FRU isolation. May be advisable to disconnect loop to avoid performance degradation or interference on other loop.
  - > 0 (fatal) - Schedule FRU isolation.

## **WARNING: transport failed: reason 'timeout'**

```
Oct 10 15:32:10 dwh02-A unix: WARNING: /sbus@16,0/SUNW,socal@2,0/sf@0,0/
ssd@w22000020370b4f6f,0 (ssd1223):
Oct 10 15:32:10 dwh02-A unix: SCSI transport failed: reason 'timeout': retrying command
```

- **Interpretation:** After ~70 seconds, ssd driver still had not received response for command. This is a rather generic error message that could be caused by several things including loop integrity problem



(CRC error on command frame), drive hung/non-responsive, driver/firmware lost command, etc. In the absence of other loop integrity indicators (CRC/SCSI parity errors) look for drive problems.

- **Probable Cause:** Fibre channel loop margin problem (GBIC, IB, BP, HUB, HA, IC, cables) Invalid drive state
- **Error Level:** Retryable, Fatal after 3 unsuccessful retries
- **Action:** Determine if this is a probable fibre channel integrity problem (look for CRC/SCSI parity errors).

If this looks like a fibre channel integrity issue:

- <11 (retryable) per day per loop - Monitor only. Run Loop Integrity test when convenient.
- >10 (retryable) per day per loop OR > 5 per day per loop for 3 consecutive days. - Schedule loop integrity testing/FRU isolation. May be advisable to disconnect to avoid performance degradation or interference on other loop.
- > 0 (fatal) - Schedule FRU isolation.

If this is not a fibre channel integrity issue:

- < 5 per day per drive - Monitor only
- > 5 per day per drive - Try replacing drive if this appears isolated to a single drive.

## WARNING: OFFLINE/ONLINE

```
Oct 10 04:12:15 MYSERVER unix: ID[SUNWssa.socal.link.5010] socal7: port 0: Fibre Channel is OFFLINE
```

```
Oct 10 04:06:15 MYSERVER unix: ID[SUNWssa.socal.link.6010] socal7: port 0: Fibre Channel Loop is ONLINE
```

- **Interpretation:** Indicates the fibre channel loop corresponding to socal7 port 0 has been re-initialized.  
Possible causes:
  - Hot plug of device.
  - Some node on loop initiated loop init after loss of sync on incoming fibre channel stream or detected some other error.
- **Probable Cause:** Fibre channel loop margin problem (GBIC, IB, BP, HUB, HA, IC, cables) or Drive error/reset.
- **Error Level:** Retryable, Loop reinitializes and is immediately usable again.
- **Action:**
  - <11 (retryable) per day per loop - Monitor only. Run Loop Integrity test when convenient.
  - >10 (retryable) per day per loop OR > 5 per day per loop for 3 consecutive days - Schedule loop integrity testing/FRU isolation. May be advisable to disconnect to avoid performance degradation or interference on other loop.

---

## WARNING: write cnt xx al\_pa 0x\*

```
Oct 12 12:13:21 MYSERVER unix: sf7: write cnt 28 al_pa 0xae
```

- **Interpretation:** Internal host adapter DMA failure
- **Probable Cause:** none.
- **Error Level:** Retryable.
- **Action:** none.

## WARNING: soc lilp map failed

```
Oct 11 14:32:06 dwh02-A unix: WARNING: /sbus@e,0/SUNW,socal@2,0/sf@1,0 (sf15):  
Oct 11 14:32:06 dwh02-A unix: soc lilp map failed status=0x5
```

- **Interpretation:** Fibre channel went OFFLINE before socal driver could obtain the loop map.
- **Probable Cause:** Fibre channel loop margin problem (GBIC, IB, BP, HUB, HA, IC, cables).
- **Error Level:** Retryable, Loop reinitializes and is immediately usable again.
- **Action:**
  - <11 (retryable) per day per loop - Monitor only. Run Loop Integrity test when convenient.
  - >10 (retryable) per day per loop OR > 5 per day per loop for 3 consecutive days. - Schedule loop integrity testing/FRU isolation. May be advisable to disconnect to avoid performance degradation or interference on other loop.

## WARNING: OFFLINE / Offline Timeout / target offlined

```
Oct 11 14:31:02 MSERVER unix: ID[SUNWssa.socal.link.5010] socal7: port 1: Fibre Channel is  
OFFLINE  
Oct 11 14:31:48 MSERVER unix: WARNING: /sbus@e,0/SUNW,socal@2,0/sf@1,0 (sf15):  
Oct 11 14:31:48 MSERVER unix: Offline Timeout  
Oct 11 14:31:48 MSERVER unix: sf15: target 0x72 al_pa 0x1f offlined
```

- **Interpretation:** Indicates the fibre channel loop corresponding to socal7 port 1 has gone offline and failed to become ONLINE for over 45 seconds. Any targets on this loop will also be “offlined” until this loop returns to an online state.  
Possible causes:
  - Catastrophic loop hardware failure (GBIC, cable, HUB, etc.)
  - Removal of cable
  - Power loss on other side of cable.
  - Hung drive or ses node preventing loop initialization fibre channel stream or detected some other error.
- **Probable Cause:** Fibre channel loop margin problem (GBIC, IB, BP, HUB, HA, IC, cables).



- **Error Level:** Fatal, although loop may return at a later time. With vxdmp, this will cause a path failure if I/Os are being issued to that path.
- **Action:** Determine root cause and correct immediately. If link returns but repeats this process, disconnect loop temporarily until loop integrity testing can be performed.

### **WARNING: ONLINE/ONLINE; host adapter fw date code**

```
Oct 11 12:27:58 MSERVER unix: WARNING: /sbus@e,0/SUNW,socal@2,0/sf@0,0/ssd@w22000020370b875c,0
(ssd518):
Oct 11 12:27:58 MSERVER unix:   SCSI transport failed: reason 'reset': retrying command
Oct 11 12:28:02 MSERVER unix: WARNING: /sbus@e,0/SUNW,socal@2,0/sf@0,0/ssd@w22000020370b7e44,0
(ssd490):
Oct 11 12:28:02 MSERVER unix:   SCSI transport failed: reason 'reset': retrying command
Oct 11 12:28:02 MSERVER unix: ID[SUNWssa.socal.driver.3010] socal7: host adapter fw date code:
<not available>
Oct 11 12:28:03 MSERVER unix: ID[SUNWssa.socal.link.6010] socal7: port 0: Fibre Channel Loop
is ONLINE
Oct 11 12:28:03 MSERVER unix: ID[SUNWssa.socal.link.6010] socal7: port 1: Fibre Channel Loop
is ONLINE
```

- **Interpretation:** Indicates that the driver has just reset the HBA. Note that both ports become ONLINE without going OFFLINE. This is due to the fact that HBA could not send OFFLINEs since it was reset. The string “host adapter fw date code” is printed whenever the HBA is reset. This resetting after the initial driver attach during boot is usually due to error recovery invoked when the driver receives no responses to loop re-initialization commands issued to the HBA (after 2 seconds). There is a bug filed under the A5x00 patch release 5 that causes this to occur unnaturally during OFFLINE conditions on one loop while both loops are under heavy load. This is fixed in A5x00 release 6, but may still occur for other error recovery reasons.
- **Probable Cause:** Driver/HA firmware bug or broken HA.
- **Error Level:** Retryable.
- **Action:** Can ignore since it’s retryable. Using separate HBAs for each loop can help.

### **WARNING: transport failed: reason ‘tran\_err’**

```
Oct 10 04:06:15 MYSERVER unix: WARNING: /sbus@e,0/SUNW,socal@2,0/sf@0,0/
ssd@w22000020370b87d8,0 (ssd432):
Oct 10 04:06:15 MYSERVER unix:   SCSI transport failed: reason 'tran_err': retrying command
```

- **Interpretation:** Indicates a command is being retried due to the fibre loop going OFFLINE. Any outstanding commands to devices on a loop that goes OFFLINE will generate this message. When the link returns to an ONLINE state, all commands will be retried. These are secondary affects of the OFFLINE.

- **Error Level:** Retryable.
- **Action:** Ignore.

### **WARNING: transport failed: reason ‘reset’**

```
Oct 11 12:27:58 MSERVER unix: WARNING: /sbus@e,0/SUNW,socal@2,0/sf@0,0/ssd@w22000020370b875c,0
(ssd518):
Oct 11 12:27:58 MSERVER unix:   SCSI transport failed: reason ‘reset’: retrying command
```

- **Interpretation:** Indicates a command is being retried due to second level error recovery for another command. This is a side affect of a previous error.
- **Error Level:** Retryable.
- **Action:** Ignore.

### **WARNING: Target Reset Failed/sf\_reset\_all**

```
Oct 10 04:12:13 MYSERVER unix: sf14:   Target 0x0 Reset Failed
Oct 10 04:12:15 MYSERVER unix: sf14:   sf_reset all
Oct 10 04:12:15 MYSERVER unix: ID[SUNWssa.socal.link.5010] socal7: port 0: Fibre Channel is
OFFLINE
```

- **Interpretation:** Informational only. Indicates second level error recovery was invoked which force the fibre channel loop to re-initialize. Note the OFFLINE that results.
- **Probable Cause:** A previous command timed out and target reset was unsuccessful.
- **Error Level:** None.
- **Action:** Address the error that preceded the second level recovery.

### **WARNING: unknown LS\_Command**

```
Oct 11 14:32:20 MYSERVER unix: WARNING: ID[SUNWssa.socal.link.3010] socal7: unknown LS_Command,
5000000
```

- **Interpretation:** Indicates a Fibre channel “link services” command has been received that is not supported by the driver. This usually indicates an attempted login from another host adapter on the loop.
- **Error Level:** None. Informational only.
- **Action:** Ignore.

### **ELS 0x\* to target 0x\* retrying**

```
Oct  3 12:31:06 MSERVER unix: ID[SUNWssa.socal.link.6010] socal4: port 0: Fibre Channel Loop
is ONLINE
Oct  3 12:31:07 MSERVER unix: sf8:   ELS 0x0 to target 0x7d retrying
Oct  3 12:31:07 MSERVER unix: sf8:   ELS 0x0 to target 0x5d retrying
```



---

```
Oct  3 12:31:07 MSERVER unix: sf8:      ELS 0x0 to target 0x3d retrying
Oct  3 12:31:08 MSERVER unix: sf8:      ELS 0x0 to target 0xd retrying
```

- **Interpretation:** Following a loop ONLINE, the socal driver issues several ELS (Fibre Channel Extended Link Service) commands to each device on the loop. Some of these devices may be busy due to small command queues (especially the SES devices) and will force the driver to retry the command. This will naturally occur more frequently with more than one host adapter on the same loop. These retries are normal.
- **Error Level:** Retryable.
- **Action:** Ignore.

### ELS 0x\* to target 0x\* timed out

```
Dec  6 13:05:10 MYSERVER unix: sf1:      ELS 0x5 to target 0x76 timed out
```

- **Interpretation:** Following a loop ONLINE, the socal driver issues several ELS (Fibre Channel Extended Link Service) commands to each device on the loop. A timeout indicates retries were exhausted attempting to communicate with this node. A broken node (ssd or ses device) or one in a confused state
- **Probable Cause:** Confused drive/IB state, broken drive, IB or too many consecutive busies when attempting to send commands to device. Also could be an indicator of a severe loop integrity problem.
- **Error Level:** Retryable after secondary recovery (loop init, target reset, etc.)
- **Action:** First determine if there is a potential loop integrity problem. Look for other loop integrity signs (CRC/scsi parity errors, etc.).

### WARNING: INCOMPLETE DMA XFER on sbus

```
Jul 15 10:22:12 MYSERVER unix: WARNING: /sbus@7,0/SUNW,socal@0,0/sf@0,0 (sf6):
Jul 15 10:22:12 MYSERVER unix:      Transport error on target=0x35:  Fiber Channel:INCOMPLETE DMA XFER on sbus
Jul 15 10:22:12 MYSERVER unix: WARNING:/sbus@7,0/SUNW,socal@0,0/sf@0,0/ssd@w22000020370cf656,0 (ssd33):
Jul 15 10:22:12 MYSERVER unix:      SCSI transport failed: reason 'incomplete': retrying
```

- **Interpretation:** This is the symptom of a parity error detected during an SBUS DMA from the socal host adapter to the host memory during a read operationA5x00. This parity error is detected approximately 97% of the times it occurs leaving a 3% chance that it will not be detected, resulting in potential data corruption. **THIS REQUIRES IMMEDIATE ATTENTION!** The root cause is a bad SOC+ chip factory escape. These should no longer be shipped out of the factory and should also be eliminated during installation when the host adapter test is executed.

- **Probable Cause:** Bad Host Adapter
- **Error Level:** Retryable. This will be fatal in future releases.
- **Action:** First, force any volumes on this loop to be failed in vxvm. This is to remove any potentially corrupted data disks on this loop. Second, replace this host adapter immediately.

## vxvm/vxio/vxdmp Messages

### WARNING: vxvm:vxio: Plex <> detached

### WARNING: vxvm:vxio: write error on Plex

```
Nov 22 00:44:30 MYSERVER unix: WARNING: vxvm:vxio: write error on Plex vredo02-02 of volume vredo02 offset 2411085 length 256
Nov 22 00:44:30 MYSERVER unix: WARNING: vxvm:vxio: Plex vredo02-02 detached from volume vredo02
Nov 22 00:44:30 MYSERVER unix: WARNING: vxvm:vxio: c4t6d0-01 Subdisk failed in plex vredo02-02 in vol vredo02
```

- **Interpretation:** A write error has bubbled up to the volume manager and the I/O has failed. This implies that retries were exhausted, multiple paths didn't help. The affected volume will run in degraded mode following the Plex detach.
- **Probable Cause:** Many. Analyze proceeding messages
- **Error Level:** Fatal. Application may be unaffected if this is the only Plex failure in a mirrored or RAIDed volume.
- **Action:** Root cause failures leading up to the vxio error.

## Disk Messages

### WARNING: ASC:0x15 (random positioning error):

```
Dec 29 13:54:58 C-BFIS21 unix: WARNING: /sbus@45,0/SUNW,socal@0,0/sf@0,0/
ssd@w22000020370ba6ad,0 (ssd72):
Dec 29 13:54:58 C-BFIS21 unix: Error for Command: read(10) Error Level: Retryable
Dec 29 13:54:58 C-BFIS21 unix: Requested Block: 3687 Error Block: 3687
Dec 29 13:54:58 C-BFIS21 unix: Vendor: SEAGATE Serial Number: 9744M44374
Dec 29 13:54:58 C-BFIS21 unix: Sense Key: Hardware Error
Dec 29 13:54:58 C-BFIS21 unix: ASC: 0x15 (random positioning error), ASCQ: 0x1, FRU: 0x1
```

- **Interpretation:** Drive head could not position correctly. Could be due to bad drive, drive movement,



drive vibration, improper/broken drive guides, abnormal chassis vibration (look for multiple drives exhibiting same errors in same chassis).

- **Error Level:** Retryable, Fatal after 3 unsuccessful retries.
- **Action:**
  - <11 (retryable) per day per disk - Monitor only.
  - >10 (retryable) per day per disk - Perform corrective action at convenience.
  - > 0 (fatal) - Perform corrective action immediately.

### **WARNING: ASC:0x9 (track following error):**

```
Nov 30 23:07:18 MYSERVER unix: WARNING: /sbus@59,0/SUNW,socal@1,0/sf@0,0/
ssd@w22000020370d6857,0 (ssd46):
Nov 30 23:07:18 MYSERVER unix:      Error for Command: read(10)  Error Level: Retryable
Nov 30 23:07:18 MYSERVER unix:      Requested Block: 45312   Error Block: 45409
Nov 30 23:07:18 MYSERVER unix:      Vendor: SEAGATE   Serial Number: 9812S73160
Nov 30 23:07:18 MYSERVER unix:      Sense Key: Soft Error
Nov 30 23:07:18 MYSERVER unix:      ASC: 0x9 (track following error), ASCQ: 0x0, FRU: 0x10
```

- **Interpretation:** Drive head could not position correctly. Could be due to bad drive, drive movement, drive vibration, improper/broken drive guides, abnormal chassis vibration (look for multiple drives exhibiting same errors in same chassis).
- **Error Level:** Retryable, Fatal after 3 unsuccessful retries.
- **Action:**
  - <11 (retryable) per day per disk - Monitor only.
  - >10 (retryable) per day per disk - Perform corrective action at convenience
  - > 0 (fatal) - Perform corrective action immediately

### **WARNING: Unhandled Sense Key 'Write Protected':**

```
Dec  6 10:22:41 dwh02-A unix: WARNING: /sbus@17,0/SUNW,socal@0,0/sf@1,0/ssd@w22000020370b49dd,0
(ssd498):
Dec  6 10:22:41 dwh02-A unix:      Unhandled Sense Key 'Write Protected
```

- **Interpretation:**
- **Error Level:** Retryable, Fatal after 3 unsuccessful retries.
- **Action:**
  - <11 (retryable) per day per disk - Monitor only.
  - >10 (retryable) per day per disk - Perform corrective action at convenience
  - > 0 (fatal) - Perform corrective action immediately

## Fibre Channel Messages - Symptom Action summary tables:

Message / symptom	Frequency	Action Required
OFFLINE messages	1-10 per day (per loop)	None. Monitor only. Loop integrity testing can be performed at next time of convenience.
	>10 per day OR > 5 per day for 3 consecutive days (per loop)	Schedule Loop Integrity testing. Physically disconnecting this loop from the host can minimize impact of retries (dual ported configuration only).
CRC / SCSI parity error messages	1-10 per day (per loop)	None. Monitor only. Loop integrity testing can be performed at next time of convenience.
	>10 per day OR > 5 per day for 3 consecutive days (per loop)	Schedule Loop Integrity testing. Physically disconnecting this loop from the host can minimize impact of retries (dual ported configuration only).
Loop Integrity Test FAILs AND 1-10 errors per day	N/A	Schedule FRU isolation when convenient. Mandatory downtime should be avoided
Loop Integrity Test FAILs AND > 10 errors per day	N/A	Schedule FRU isolation for this loop. To avoid immediate downtime requirements consider disconnecting this loop from the system temporarily until FRU isolation can be scheduled (dual ported configs only).
Loop Integrity Test PASSes AND 1-10 errors per day	N/A	None. Monitor Only. FRU isolation may not be successful due to intermittence of failure.
Loop Integrity Test PASSes AND > 10 errors per day	N/A	Escalate this to engineering. This is either a Loop Integrity Test problem or some other type of failure/issue.
Offline Timeout (vxdmp path failure)	> 0	Schedule FRU isolation for this loop. To avoid immediate downtime requirements consider disconnecting this loop from the system temporarily until FRU isolation can be scheduled (dual ported configs only).
Unknown LS_Command	N/A	Ignore.
ELS ... to target ... retrying	N/A	Ignore.
transport failed: reason: 'reset'	N/A	Ignore.
sf_reset_all	N/A	Ignore.
transport failed: reason: 'tran_err'	N/A	Ignore.



Message / symptom	Frequency	Action Required
transport failed: reason: 'timeout'	1-10 per day	None. Monitor only. Loop integrity testing can be performed at next time of convenience.
	>10 per day	Schedule Loop Integrity testing. Physically disconnecting this loop from the host can minimize impact of retries (dual ported configuration only).
soc liip map failed	1-10 per day	None. Monitor only. Loop integrity testing can be performed at next time of convenience.
	>10 per day	Schedule Loop Integrity testing. Physically disconnecting this loop from the host can minimize impact of retries (dual ported configuration only).
Target Reset Failed/ sf_reset_all	N/A	Ignore.
host adapter reset ("host adapter fw date code") followed by both ports ONLINE without any OFFLINES	N/A	Ignore.
INCOMPLETE DMA XFER on sbus	> 0	Manually fault all plexes on this loop to prevent possible corrupted data from being saved. Replace HA immediately..
ELS to target timed out	1-10 per day (per loop)	None. Monitor only. Loop integrity testing can be performed at next time of convenience.
	>10 per day OR > 5 per day for 3 consecutive days (per loop)	Schedule Loop Integrity testing. Physically disconnecting this loop from the host can minimize impact of retries (dual ported configuration only).
random positioning error track following error	1-10 per day (per drive)	None. Monitor only.
	>10 per day OR > 5 per day for 3 consecutive days (per drive)	Reseat disk. If problems persist, replace disk.



Reference the Installation and Service Manual for complete hardware installation and repair procedures. Be aware that fc-al technology, FRU hot plugability and disk world wide numbers (WWNs) necessitate more involved procedures.

## This chapter will introduce the following service issues:

- FRU replacement
  - Disks
  - Interface Board (IB) and Interconnect Assembly (IA)
- Array Booting

## FRU Replacement

### Disks

Hot swap replacement of disk drives require additional steps if the drives are under software raid control. The SEVM steps are documented in the service manual but we'll summarize here. SDS steps will follow.

#### 6 Disk replacement steps for SEVM:

- 1. Remove the disk from VM (vxdiskadm - 4)
- 2. Offline the disk from VM (vxdiskadm - 11)
- 2. luxadm remove the disk from solaris and the array (this will delete the /devices and /dev entries). The -F option may be required. Repeat for all hosts if multi-hosted.
- 3. luxadm insert the disk into the array and solaris (this will create the /devices and /dev entries with the new disk's wwn). Repeat for all hosts
- 4. vxdctl enable to notify VM of new disk
- 5. Replace disk in VM or SDS (vxdiskadm - 5)

#### SEVM Example:

```
# vxdiskadm
~
4 Remove a disk for replacement
11 Disable (offline) a disk device
~
# luxadm remove /dev/dsk/c2t32d0s2
```

```
WARNING!!! Please ensure that no filesystems are mounted on these device(s).
All data on these devices should have been backed up.
```

```
The list of devices which will be removed is:
1: Box Name "array1" front slot 0
```



---

Please enter 'q' to Quit or <Return> to Continue: <CR>

stopping: Drive in "array1" front slot 0...Done

offlining: Drive in "array1" front slot 0...Done

Hit <Return> after removing the device(s).

Drive in Box Name "array1" front slot 0

Removing Logical Nodes:

Removing c2t32d0s0

Removing c2t32d0s1

Removing c2t32d0s2

Removing c2t32d0s3

Removing c2t32d0s4

Removing c2t32d0s5

Removing c2t32d0s6

Removing c2t32d0s7

#

# **luxadm insert array1,f0**

The list of devices which will be inserted is:

1: Box Name "array1" front slot 0

Please enter 'q' to Quit or <Return> to Continue: <CR>

Hit <Return> after inserting the device(s).

Drive in Box Name "array1" front slot 0

Logical Nodes under /dev/dsk and /dev/rdisk :

c2t32d0s0

c2t32d0s1

c2t32d0s2

c2t32d0s3

c2t32d0s4

c2t32d0s5

c2t32d0s6

c2t32d0s7

#

# **vxdctl enable**

# **vxdiskadm**

~

5 Replace a failed or removed disk

~

#

## 7 possible Disk Replacement Steps for SDS

- 1). Preserve disk label if needed
- 2). Metareplace the disk slices if it has not been hotspared (optional)
- 3). `luxadm remove -F` the disk. The `-F` is required as SDS doesn't offline disks. Repeat for all hosts if multi-hosted
- 4). `luxadm insert` the new disk. Repeat for all hosts
- 5). Label the disk if needed
- 6). Replace `metadb` if needed
- 7). Enable new disk slices with `metareplace -e`

### SDS Example:

```
# prtvtoc /dev/rdisk/clt3d0s2 > /tmp/clt3d0.vtoc
# luxadm remove -F /dev/rdisk/clt3d0s2
~
# luxadm insert a,f3
~
# cat /var/tmp/clt3d0.vtoc | fmthard -s - /dev/rdisk/clt3d0s2
fmthard: New volume table of contents now in place.
# metadb -d clt3d0s0; metadb -a clt3d0s0
# metareplace -e d0 clt3d0s0
d0: device clt3d0s0 is enabled
#
```



---

## IB and Interconnect Assembly Replacement

The A5x00 IBs contains the arrays functional firmware. When replacing IBs or IAs, its important to understand how the firmware downloads function in the array.

### There are 3 sets of firmware in the array

- Downloadable functional code in the IB
- Static or “golden” firmware code in the IB
  - \* Resides on EEPROM and only loaded at the factory. This is a “known good” code which can be used if download firmware is corrupt, but the static firmware can be down rev.
- Downloadable code on the motherboard (part of the IA)

### There are 3 ways to modify firmware in the array IBs

- luxadm download command
  - \* A temporary fw download to single IB only. Will revert back to original version on reset or power cycle (*not recommended usage*)
  - \* Use the “-s” or save option for permanent downloads to both IBs and the mother board
- On power up or reset the IBs compare code with the mother board and on mismatch will download firmware from mother board, even if the IB’s code is a higher revision.
- New Interconnect Assemblies (with mother board) with newer or older firmware revisions will download that fw to the IBs.
  - \* Note that replacing the IA will change the array WWN and reset the boxID and name.

The IBs have the working downloadable code, plus a static stored version. The motherboard has no working firmware only a downloadable stored version, but this is considered the version master and will download any IBs with a different version, either **upgrading or downgrading the IB working code**.

---

## Booting Issues

**There are 3 main considerations when systems boot from the A5x00**

- 1. You cannot upgrade firmware while there is any I/O going to the array.
- 2. Each host should boot to a separate loop.
- 3. Backup tapes cannot restore a crashed/replaced disk with a different WWN.

### Firmware upgrades when booting from A5x00

#### Two choices for firmware upgrades

- Boot from an alternate device (preferred)\*
- Mirror to an alternate device and detach the array mirror (but you have to reboot anyway to load new drivers, and the loss of redundancy and performance impact of a long resync may be worse than the minimal downtime of the alternate boot device method)

#### \* Boot cdrom firmware upgrade procedure:

- Install the needed A5x00 patches (see matrix)
- Boot cdrom or alternate device
- Copy the needed firmware files to the /tmp directory.
- Download firmware
- reboot system

Example:

```
ok boot cdrom -s
~
# mount -o ro /dev/dsk/clt0d0s0 /a
# cp /a/usr/lib/locale/C/LC_MESSAGES/ibfirmware /tmp
# cp /a/usr/lib/firmware/fc_s/fcal_s_fcode /tmp
# cp -r /patchdir/106129-02 /tmp
# umount /a
# luxadm download -s -f /tmp/ibfirmware ARRAYNAME
# luxadm fcal_s_download -f /tmp/fcal_s_fcode
# cd /tmp/106129-02
# ./download
```



---

**There are 2 main things to remember:**

- You want to run all the new drivers with the new firmware. This is why you update everything before rebooting.
- You don't want to download any firmware while there is any I/O going on and the only way to guarantee that is to have the array completely idle.

## **Multi-Host booting to an A5x00**

### **Each host should boot to a separate loop**

When multiple hosts are booting to the same array fc-al loop its possible that one host's loop init (the loop must LIP to add the host as a new loop node) will cause another host to not see the loop "device not found" and require another boot command. This rare occurrence only happens during simultaneous boots and before the OS is able to handle the loop init offline condition with I/O retries.

The work around for this condition is to simple have each hosts primary boot path on alternate loops. HostA boots to A loop with B as alternate, and HostB boots to the B loop with A as alternate.

If four hosts are booting to the array you may want to consider running in split loop mode so you have four loops.

## **Boot disk recovery on an A5x00**

### **Backup tapes cannot restore a crashed/replaced disk with a different WWN.**

Because the disk has a unique WWN the device files of the original disk will not match the replacement disk. The software can be restored from a backup tape, but the system will not boot as it tries to access a disk with a different WWN.

The work around for this problem is to always mirror the boot disk when booting from an A5x00. This will additionally prevent system down time on boot disk failure.

Reference Symptom and Resolution Database (SRDB) document number 17658 for a recovery procedure by manually editing the /devices entries to replace the WWN, this has some risk and its best to maintain a good boot disk mirror.



## STORtools

The STORtools package was created to simplify the monitoring and troubleshooting of StorEdge A5x00 disk array installations. The tools provide an easy to use menu driven front end program with task explanations and usage help. Command line utilities are provided for advanced customized use and have standard man pages for online documentation. The STORtools User Guide contains complete information including detailed troubleshooting flow charts.

STORtools was designed to address the following 4 main areas of A5x00 maintenance:

### **Revision checking**

- patches
- firmware
- minimum disk configuration

### **Configuration Management**

- Display/map the A5x00 fc-al loop configurations
- Configuration status checking for hardware failure and software or firmware change

### **Monitoring and Notification**

- storage messages summary for predictive failure analysis
- crontab setup for configuration status and messages checking with email notification of problems

### **Diagnostics**

- Loop integrity tests
- Host adapter tests
- Disk exercisers
- Installation verification load tests

Reference the STORtools User Guide for more information.



---

## Example of stormenu:

# stormenu

NOTICE: Use of this software and associated packages is authorized ONLY by personnel trained in it's use and trained on the Products for which it was intended to be used on.

The execution of this toolkit assumes the user accepts the licensing terms and agreements as contained in the file NOTICE.txt.

---- Offline port(s) (HBA tests only) ----  
---- Offline ports: /devices/sbus@1f,0/SUNW,socal@0,0:1 ----

---- FCAL HBA(s) within the STORtools domain ----  
---- Online: /devices/sbus@1f,0/SUNW,socal@0,0:0 ----  
      (c1) (sf0) (socal0:Port 0)  
.....(7 drive(s))

PASS

Press <return> to continue ...

A5000 Storage Management                   Version 3.1  
SBUS Menu: Main

-----

- [1] Check Revisions
- [2] Create Configuration Snapshot File
- [3] Display Configuration
- [4] Check Current Status
- [5] Storage Monitor
- [6] Run Diagnostics
- [7] Installation Verification

- [?] Display Help
- [q] Quit

Current Date:                   Thu Apr  1 15:36:48 PST 1999  
Date of Gold Snapshot:         Mon Mar 22 16:30:00 PST 1999

Filename of Gold Snapshot:  
      /var/opt/STORtools/logs/Golden\_Snapshot\_990322\_1629

Select an operation to perform:

## LBF

The lbf (loop back frame) diagnostic is an unbundled test which is used to test the integrity of an FC-AL subsystem. This diagnostic is the core of the STORtools Loop Integrity Test and its recommended that the stormenu diagnostics submenu be used for loop testing, but lbf can be run from the command line. The diagnostic can be obtain with the STORtools package. Information on STORtools can be found at the following Sun internal web site, or from a Sun representative.

<http://storageweb.eng/tm/STORtools/>

To run the diagnostic you will need to pass it the physical address of the loop to test. Use the output of luxadm probe to get the /devices path, then add a ':0' or ':1' to the end to specify GBICs 0 and 1. An example of obtaining the address and running a simple test is below:

```
# luxadm probe -p
Found
SENA          Name:array1 Node WWN:50800200000009e8
  Logical Path:/dev/es/ses0
  Physical Path:/devices/sbus@3,0/SUNW,socal@0,0/sf@0,0/ses@w50800200000009e9,0:0
  Logical Path:/dev/es/ses1
  Physical Path:/devices/sbus@3,0/SUNW,socal@0,0/sf@0,0/ses@w50800200000009ea,0:0
  Logical Path:/dev/es/ses4
  Physical Path:/devices/sbus@7,0/SUNW,socal@0,0/sf@0,0/ses@w50800200000009eb,0:0
  Logical Path:/dev/es/ses5
  Physical Path:/devices/sbus@7,0/SUNW,socal@0,0/sf@0,0/ses@w50800200000009ec,0:0
```

Or if the array is not being seen you can get the address of the soc+HBA with:

```
# luxadm fcal
Found Path to 2 FC100/S Cards
  Device: /devices/sbus@3,0/SUNW,socal@0,0
  Detected FC100/S Version: @(#) FCode 1.10 97/09/15
  Device: /devices/sbus@7,0/SUNW,socal@0,0
  Detected FC100/S Version: @(#) FCode 1.10 97/09/15
Complete
```

```
# lbf
Rev 1.3
Usage: lbf [-x <xfer size, bytes>|-k <xfer size, kbytes>] [-n <number of passes>] [-t <test
pattern>] [-c (compare on)] [-v (verbose mode) | -q (quiet mode) | -o (throughput mode)]<device>

# lbf -t 7e7e7e7e -k 64k -n 6000 /devices/sbus@f,0/SUNW,socal@0,0:0
PASS    Pat=0x7e7e7e7e Cnt=6000 /devices/sbus@f,0/SUNW,socal@0,0:0
```

Example of lbf failure:



```
# lbf -t 7e7e7e7e -k 64k -n 6000 /devices/sbus@f,0/SUNW,socal@0,0:0
FAIL=41 Pat=0x7e7e7e7e Cnt=0 /devices/sbus@f,0/SUNW,socal@0,0:0
#
```

For failures that are difficult to reproduce you can run the test with varying data patterns and large block sizes. lbf generates a stream of frames that are transmitted and received by the specified socal port. The data received is compared against the data transmitted and error messages are produced whenever there is a failure to receive what has been transmitted. The following script transmits a variety of data patterns which generate known to be worst case bit patterns that can affect the circuitry. Be advised that some failures can only be reproduced using heavy disk I/O while running lbf.

```
#!/bin/sh
# name this lbf.sh
set -x
for i in 31 7e bf de df ee fc fd
do
lbf -t $i$i$i$i -n 10000 -k 64 $i
sleep 1
done
```

For failures that can only be reproduced under heavy load. The STORtools Disk Read Loop Test can be used, or you can run any disk loader program to produce load then run lbf to check for errors. Below is an example of running a simple dd command as a disk loader:

```
#!/bin/csh -xf
#
# Usage dd.csh <controller#>
#
foreach i (`ls /dev/rdisk/$1*s2`)
/bin/dd if=$i of=/dev/null bs=128k &
end
```

The following script can be used to kill all the dd processes (use with great caution):

```
#!/bin/csh
#
# Kill all processes with string in them
#
# Usage: kill_all <string>
# (use with extreme caution)
#
ps -lef | grep $1 | awk '{print "kill -9", $4}'
```

---

**Note** – Some highly intermittent failures may require concurrent disk loading and lbf testing to both loops of an array subsystem.

---

Once a loop has been determined to fail, a process of elimination (POE) technique is used to find the failing FRU. This is a process of systematically checking loop components by dividing the loop into pieces and retesting. As this isolation can be performed to an online system, it's important to be careful not to disturb other active loops! Reference the troubleshooting information in chapters 3 and 4 for details.

*For loop testing of hardware that doesn't support lbf, such as the current PCI HBA, you can use disk loading to generate heavy loop traffic and monitor the /var/adm/messages file for error indicators. FRU Isolation can be accomplished using the “FRU Isolation steps with disk access testing” as described in chapter 4.*



## OBP

There are several commands which can be executed at the Open Boot Prom(OBP) on a halted system to assist with A5x00 diagnosis.

To verify that the social host adapter is recognized by the system type the “show-devs” command and look for the social device.

```
ok show-devs
~
/sbus@3,0/SUNW,socal@0,0
/sbus@3,0/SUNW,fas@3,8800000
/sbus@3,0/SUNW,hme@3,8c00000
/sbus@3,0/SUNW,socal@0,0/sf@1,0
/sbus@3,0/SUNW,socal@0,0/sf@0,0
/sbus@3,0/SUNW,socal@0,0/sf@1,0/ssd
/sbus@3,0/SUNW,socal@0,0/sf@0,0/ssd
~
```

To test the social adapter and verify node (disk/ses) connectivity.

```
ok cd /sbus@3,0/SUNW,socal@0,0
ok open
ok words
social-diag-all          social-diag-extlb
social-diag-srds1b      social-diag-intlb
social-diag-hcb         social-diag-xram          reset
selftest      close      open          power-on
social-diag   set-fc-port  fcps-go      show-wwn
.wwn&alpa    target?      myalpa-len   fcloop-on
no-online-msg timeout?     set-timeout  timeout-msecs
rsp-wait-msecs      ha>wwn?      wwn-invalid? re-login
fc-port            pln-port-wwns pln-wwns     map-out      map-in
dma-sync           dma-map-out   dma-map-in   dma-free
dma-alloc          decode-unit
ok social-diag-all
XRAM test passed.
HCB test passed.
Internal Loopback passed (port 0 ).
Internal Loopback passed (port 1 ).
SERDES Loopback passed (port 0 ).
SERDES Loopback passed (port 1 ).
ok selftest
ok open
ok show-wwn
WWN 21000020370791b6 Loopid dc
WWN 50800200      7672 Loopid b5
WWN 21000020370b1547 Loopid cd
```



---

WWN 2100002037080c19 Loopid cc  
WWN 21000020370b14e3 Loopid cb  
WWN 21000020370b14ac Loopid ca  
WWN 21000020370792c0 Loopid c9  
WWN 21000020370b1615 Loopid c7  
WWN 21000020370b1569 Loopid c6  
WWN 50200200 2afa Loopid 1  
WWN 50800200 9eb Loopid 9e  
WWN 2200002037070769 Loopid b2  
WWN 2200002037070ad7 Loopid b1  
WWN 22000020370707b7 Loopid ae  
WWN 2200002037070d28 Loopid ad  
WWN 2200002037070c71 Loopid ac  
WWN 2200002037070bf4 Loopid ab  
WWN 2200002037070bb8 Loopid aa  
WWN 50800200 9ec Loopid 75  
WWN 2200002037070c53 Loopid 98  
WWN 2200002037070be5 Loopid 97  
WWN 2200002037070dd7 Loopid 90  
WWN 22000020370708f8 Loopid 8f  
WWN 220000203707078f Loopid 88  
WWN 220000203707063a Loopid 84  
WWN 2200002037070c6f Loopid 82  
WWN 50800200 7671 Loopid d2  
WWN 21000020370b14a7 Loopid ef  
WWN 21000020370b15f5 Loopid e8  
WWN 21000020370b1549 Loopid e4  
WWN 21000020370b1994 Loopid e2  
WWN 2100002037080aa9 Loopid e1  
WWN 21000020370b1588 Loopid e0  
ok



---

## SunVTS

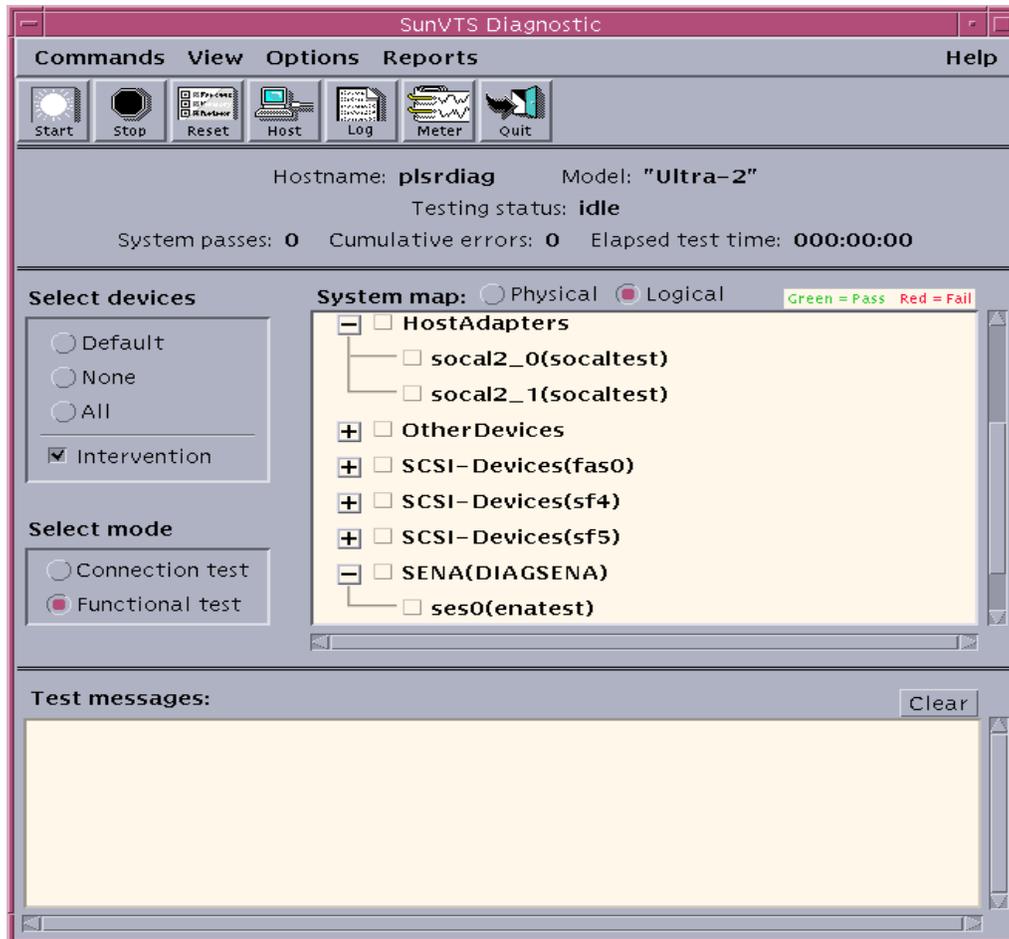
SunVTS is Sun's Solaris based diagnostic tool with an easy to use user interface. Photon support is available in SunVTS 2.1. lbf tests are in SunVTS 2.1.2 and later. SunVTS can be used with photon for validation, exercising, and fault isolation.

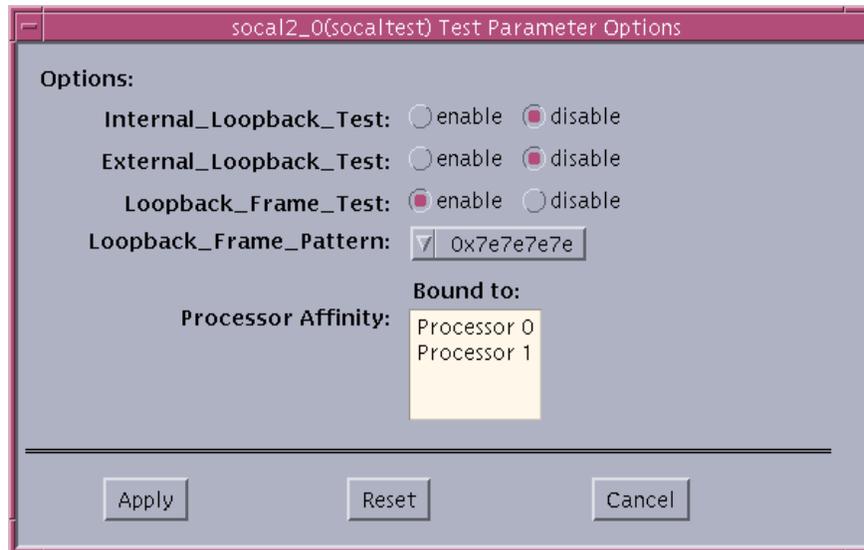
New modules have been added to SunVTS for A5x00 support.

- 1). enatest: This covers the Photon enclosure as well as the connections between the host and the enclosure.
- 2). socaltest: This targets the SOC+ based host adapter card.
- 3). lbf Test: The Loopback Frame Test sends out a buffer initialized with the selected pattern and compares it with the looped back frame. It passes if the two compare and fails if there are any bad compares.
- 4). lbf Pattern: The list of selectable patterns for the Loopback Frame Test

The tests are selectable from the SunVTS GUI. The enatest is grouped under "SENA". The ascii name of the enclosure is displayed along with the SENA group node for easy identification. The socaltest is grouped under "HostAdapters". The lbf test is available via socaltest.

Examples below:





In addition to these new tests, the disktest has been enhanced to be safe in a multi-pathed environment. This is accomplished by disallowing concurrent read/write testing on a disk through multiple paths to avoid potential data corruption. However, concurrent execution of the test in read-only mode on a disk is allowed even if there is a test instance running in the read/write mode through an alternate path.

For additional information on SunVTS reference the SunVTS documentation.

## SYMON

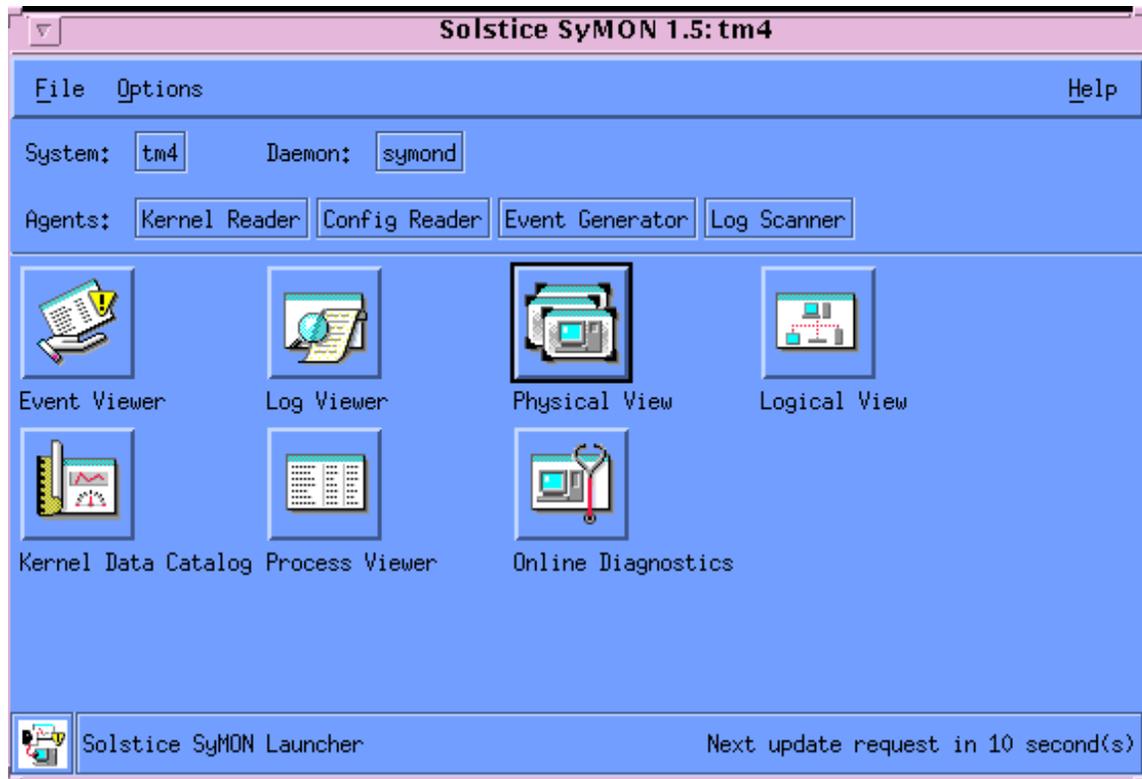
Solstice SyMON, the comprehensive system monitoring tool for Ultra Enterprise 2, 150, 450, 3000-6000 and SS1000E/SC2000E servers, dramatically increases the system administrator's ability to pinpoint a particular server's status. System administrators can use Solstice SyMON to monitor server performance and prevent certain failures from occurring.

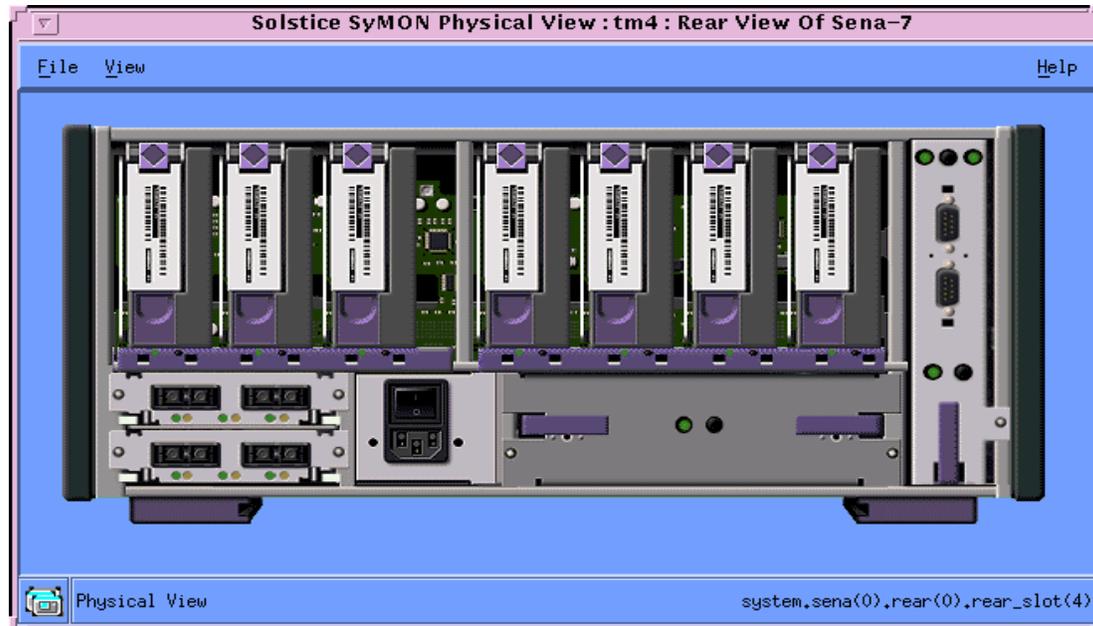
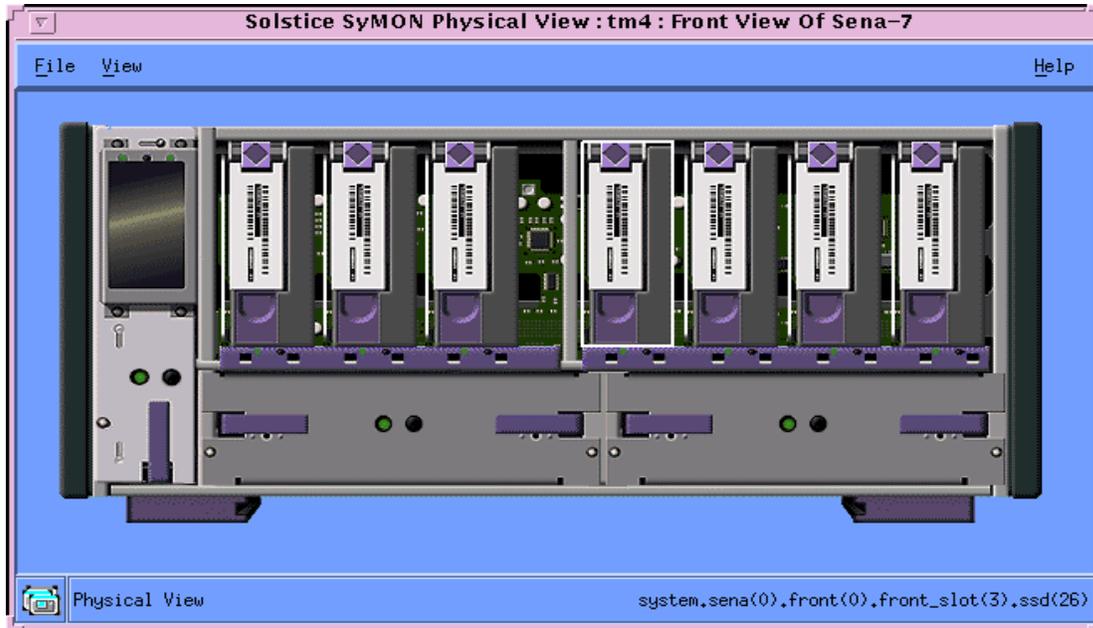
Solstice SyMON 1.5 is the follow-up release of Solstice SyMON 1.4, introduced August, 1997. This release includes support for the Sun Enterprise Network Array A5x00.

Information on downloading Symon can be obtained at the following web sites:

SunSolve internal site <http://sunsolve.corp.sunsolve.com/symon.html>  
SunSolve external site <http://sunsolve2.sun.com/sunsolve/symon/>

Below are examples of the Symon GUI and A5x00 physical views as provided by Symon:







This section contains procedures and examples of the commands used to maintain the various software and firmware revisions which control the A5x00 storage array. Always reference the latest Configuration Matrix for current patch and firmware compatibility and version information. The latest patch information or matrix can be obtained from a SunService representative or found at the following web sites:

<http://storageweb.eng/tm/photon/main/> (Sun internal use only)  
<http://docs.sun.com/> (under Storage, Sun Enterprise Network Array)

The following are the four main areas to be aware of:

- 1. Solaris release and installed patches
- 2. Host Bus Adapter FCODE
- 3. Array Interface Board Firmware
- 4. Disk Firmware

## 1. Solaris Release and Installed Patches

To check the Solaris release:

```
# uname -a
SunOS tm2 5.6 Generic sun4u sparc SUNW,Ultra-Enterprise
# cat /etc/release
          Solaris 2.6 s297s_smccDesktop_37cshwp SPARC
          Copyright 1996 Sun Microsystems, Inc.  All Rights Reserved.
          Manufactured in the USA 18 July 1997
#
```

To check for an installed patch:

```
# showrev -p | grep 105375
Patch: 105375-04 Obsoletes: Requires: Incompatibles: Packages: SUNWluxal, SUNWluxdv,
SUNWluxop
#
```

For instructions on patch installation read the patch README text file.

## 2. Host Bus Adapter Fcode

To check the version of sbus adapter Fcode use the prtconf command:



```
# prtconf -vp|grep FCode
    version: '@(#) FCode 1.11 97/12/07'
    version: '@(#) FCode 1.11 97/12/07'
    version: '@(#) FCode 1.11 97/12/07'
```

#  
To check the Fcode version of a PCI adapter look for the boot message in the /var/adm/messages file.

```
# more /var/adm/messages
~
Mar 12 16:35:30 mississippi unix: PCI-device: scsi@1, ifp #0
Mar 12 16:35:30 mississippi unix: pci1077,21000 is /pci@1f,2000/scsi@1
Mar 12 16:35:30 mississippi unix: ifp0: Firmware Version: v1.11
~
```

To upgrade the Fcode follow the instructions in the related patch README file. Downloads must be done on idle arrays and the system rebooted after, below is an example of an Sbus upgrade instruction:

```
# luxadm fcal_s_download -f /usr/lib/firmware/fc_s/fcal_s_fcode
Found Path to 1 FC100/S Cards
  Device: /devices/io-unit@f,e0200000/sbi@0,0/SUNW,socal@2,0
WARNING!! This program will update the Fcode in this FC100/S Sbus Card.
This may take a few (5) minutes. Please be patient.
Do you wish to continue ? (y/n) y
Loading FCode: /usr/lib/firmware/fc_s/fcal_s_fcode
Loading 0x40000 bytes from /usr/lib/firmware/fc_s/fcal_s_fcode at offset 0x0
Filling with zeroes...
selecting bank 0
selecting bank 1
selecting bank 2
selecting bank 3
Erasing...
FEPROM erased after 118 attempts
Programming...
selecting bank 0
selecting bank 1
selecting bank 2
selecting bank 3
Programming done
  Complete
```

### 3. Array Interface Board (IB) Firmware

This firmware is referred to as the “IB firmware” because that is where the functional code resides. However the Interconnect Board Assembly of the array contains a master copy which is used to download any IB module which is installed. Be aware that this automatic download occurs even if the IB module has a newer revision of firmware. This protects against the array ever running dif-



ferent versions in the 2 IB modules.

When ever replacing either an IB module or an Interconnect Assembly you should always use the array front panel module to check the firmware revision of the IB modules. You can also check the version with the luxadm command. Example:

```
# luxadm disp box1
                               (luxadm version: 1.36 98/02/26)
                               SENA
                               DISK STATUS
SLOT  FRONT DISKS              (Node WWN)          REAR DISKS          (Node WWN)
0     On (O.K.)                2000002037070769   On (O.K.)           2000002037070c53
1     On (O.K.)                2000002037070ad7   On (O.K.)           2000002037070be5
2     On (O.K.)                20000020370707b7   On (O.K.)           2000002037070dd7
3     On (O.K.)                2000002037070d28   On (O.K.)           20000020370708f8
4     On (O.K.)                2000002037070c71   On (O.K.)           200000203707078f
5     On (O.K.)                2000002037070bf4   On (O.K.)           200000203707063a
6     On (O.K.)                2000002037070bb8   On (O.K.)           2000002037070c6f
                               SUBSYSTEM STATUS
FW Revision:1.01 Box ID:1   Node WWN:5080020000009e8   Enclosure Name:box1
Power Supplies (0,2 in front, 1 in rear)
  0 O.K.(rev.-02) 1 O.K.(rev.-02) 2 O.K.(rev.-02)
Fans (0 in front, 1 in rear)
  0 O.K.(rev.-05) 1 O.K.(rev.-00)
ESI Interface board(IB) (A top, B bottom)
  A: O.K.(rev.-04)
    GBIC module (1 on left, 0 on right in IB)
    0 O.K.(mod.-05) 1 Not Installed
  B: O.K.(rev.-04)
    GBIC module (1 on left, 0 on right in IB)
    0 O.K.(mod.-05) 1 Not Installed
Disk backplane (0 in front, 1 in rear)
  Front Backplane: O.K.(rev.-05)
    Temperature sensors (on front backplane)
    0:39°C 1:39°C 2:39°C 3:37°C 4:39°C 5:40°C
    6:42°C (All temperatures are NORMAL.)
  Rear Backplane: O.K.(rev.-05)
    Temperature sensors (on rear backplane)
    0:40°C 1:40°C 2:40°C 3:37°C 4:39°C 5:40°C
    6:42°C (All temperatures are NORMAL.)
Interconnect assembly
  O.K.(rev.-03)
Loop configuration
  Loop A is configured as a single loop.
  Loop B is configured as a single loop.
Language      USA English
#
```



To download new firmware be sure to use the “-s” or save option to update the Interconnect Assembly as well as the IBs, and make sure the array is idle. Example:

```
# luxadm download -s box1
Using the default file /usr/lib/locale/C/LC_MESSAGES/ibfirmware.
#
```

---

**Note** – You cannot upgrade from revision 1.01 to 1.05 or higher directly. If you have an array running 1.01 you will need to get a copy of the 1.03 firmware. Upgrade to 1.03, then to 1.05. Example below:

---

```
# luxadm download -s -f /tmp/ibfirmware.1.03 box1
# luxadm download -s box1
```

This is an isolated unusual restriction which should not reoccur once the array is above revision 1.01, so future upgrades will be single step. The 1.03 firmware can be obtained from the older array patches (105310-04 for Solaris2.5.1, and 105375-04 for Solaris 2.6) or from a Sun representative.

## 4. Disk Firmware

The array disks also contain firmware which may need periodic updates. You can use the luxadm command to check disk firmware revisions. Example:

```
# luxadm disp /dev/rdisk/c2*s2 | grep Revision
Revision:      1178
Revision:      1178
Revision:      1178
Revision:      1178
Revision:      117E
Revision:      1178
Revision:      1178
Revision:      1178
Revision:      1178
Revision:      117E
#
```

The first two characters are the revision the last are a vendor hardware(HW) code. ie:

11	7E
└─	└─
Revision	HW code



To upgrade disk firmware follow the instructions in the appropriate patch (106129-xx) README.

**Example:**

```
# ./download
Firmware Download Utility, V3.11

***** WARNING *****
NO OTHER ACTIVITY IS ALLOWED DURING FIRMWARE UPGRADE!!!
No other programs including any volume manager (e.g. Veritas,
SDS, or Vold) should be running. Other host systems sharing
any I/O bus with this host must either be offline or
disconnected. Any interruption (e.g. power loss) during
upgrade can result in damage to devices being upgraded.

Any disk to be upgraded should first have it's data backed up.
*****

Searching for devices...

DISK DEVICES:
Device          Rev   Product
c1t32d0-c1t54d0: 1378  SEAGATE  ST19171FC -- SUN9.0G
Total Devices:  11

Enter command: prog
Upgrading disks...

c1t53d0: Successful download
c1t52d0: Successful download
c1t38d0: Successful download
c1t33d0: Successful download
c1t48d0: Successful download
c1t32d0: Successful download
c1t51d0: Successful download
c1t35d0: Successful download
c1t54d0: Successful download
c1t49d0: Successful download
c1t50d0: Successful download
c1: recovery delay, 67 sec.

DISK DEVICES:
No devices found to upgrade!

Enter command: quit

#
```



## Appendix C - Additional Resources

---

### Product Documentation

#### Hardware

Configuration Guide	805-0264
Product Note	805-1018
Installation and Service Manual	802-7573
Installation Tasks and Documentation Guide	805-1903
GBIC Installation card	802-7571
SBus FC-100 Host Adapter Installation and Service Manual	802-7572
PCI FC-100 Host Adapter Installation and Service Manual	805-3682
FC-100 Hub Installation and Service Manual	805-0315
Caster Base Installation Guide	802-5034
X-Option Installation Card	805-0071

#### Rackmount documentation

Rackmount Placement Matrix	802-6945
FC-100 Hub Rackmounting Installation Card	805-1990
Rackmounting Card	805-0070

#### Software

Solaris 2.5.1 Hardware: 8/97 Platform Notes: Using luxadm Software	805-1637
Solaris 2.6 Platform Notes: Using luxadm Software	805-2040
Sun Enterprise Volume Manager 2.5 User Guide	805-1603
Sun Enterprise Volume Manager 2.5 System Administrators Guide	805-1607
Sun Enterprise Volume Manager 2.5 Installation Guide	805-1604
Sun Enterprise Volume Manager 2.5 Software Release Notes	805-1605
Sun Enterprise Volume Manager 2.6 System Administrators Guide	805-5706
Sun Enterprise Volume Manager 2.6 Software Installation Guide	805-5707
Sun Enterprise Volume Manager 2.6 Software Release Notes	805-5708
Solstice DiskSuite 4.2 Installation and Product Notes	805-5960
Solstice DiskSuite 4.2 User Guide	805-5961

### Training Courses

#### Lecture Lab

SunU: Photon Disk Array - 2 day (FT957W)
SunEd: TBD



---

## **Web Based**

Sun Enterprise Network Array - 4 hours self paced (FS16W.Z)

## **Web Sites**

<http://www.sun.com/storage/a5000/> - general information

<http://docs.sun.com/> - product note with patch information

### **Sun Internal Web Sites**

<http://thedance.ebay/hardware/photon/> - documentation

<http://storageweb.eng/tm/photon/main/> - technical marketing site

<http://storageweb.eng/tm/STORtools/> - STORtools information

## **White papers**

Sun A5000 and SEVM Configuration Guide (SunWIN token #73710)

A5000 JTF Performance Report

Sun StorEdge A5000 Architecture

StorEdge A5000 Installation Check List and Planning Guide

STORtools User Guide



## Glossary

---

**AL\_PA:** Arbitrated Loop Physical Address. An address given to a device on a loop.

**automatic datapath failover:** The ability of the Sun Enterprise Network Array to immediately select an alternate path on a loop when an existing path fails. Automatic datapath failover is made possible by a combination of the multi-path hardware and the Sun Enterprise Volume Manager(SEVM) DMP capabilities. DMP is available with SEVM version 2.5 and later.

**backplane:** The hardware component that contains data and power connections and status LEDs for the disk drives.

**CLI:** Command-line interface. A character-based user interface, in contrast to a graphical user interface (GUI). A command line is a string of characters beginning with a command, often followed by arguments, including options, filenames, and other expressions, and terminated by the end-of-line character.

**dynamic load balancing:** A method of increasing the bandwidth between two points on a loop by adding another line and making the traffic load approximately equal between the lines.

**DMP:** Dynamic multi-pathing. The ability for an I/O request to be automatically routed via various physical pathways to the same disk device.

**enclosure:** The chassis, the core mechanical structure that contains the mechanical support and shielding necessary to mount the enclosed components.

**fabric:** A combination of switching elements joining a number of Fibre Channel nodes.

**FC-AL:** Fibre Channel Arbitrated Loop. A standard for a shared access loop, in which a number of devices are connected (as opposed to point-to-point transmission). FC-AL also refers to a loop that meets FC-AL requirements.

**FCP:** Fibre Channel Protocol for SCSI. The rules that govern the transmission of SCSI commands and data using Fibre Channel technology. FCP is also the name of the standard that defines those rules.

**FEPRM:** Flash-erasable programmable read-only memory. A type of non-volatile memory that can be overwritten with new information. This type of read-only memory is used on the SOC+ chip on the interface boards for storing firmware.

**firmware:** Operating system, boot instructions, or other software permanently encoded in ROM on a chip. Some kinds of firmware -- such as that used for the Sun Enterprise Network Array, can be overwrit-

---

ten when stored in an EPROM or FEPROM chip.

**FPM:** Front panel module. The device attached on the front fan tray of the Sun Enterprise Network Array that allows you to view information about the disk array and enter commands to change the configuration and other settings.

**FRU:** Field-replaceable unit. A hardware assembly designed for relatively quick replacement at a customer site when it fails.

**GBIC:** Gigabit Interface Converter. A hot-swappable module that converts optical to electrical and electrical to optical signals.

**GUI:** Graphical user interface. A method of interacting with a computer that involves the manipulation of objects on the computer screen that control the computer's activities, usually with a mouse or other selection device. Contrast with command-line interface.

**hot-swappable:** A hardware feature that allows the user to insert or remove a device while a computing device such as a disk array is operating. Most FRUs in the Sun Enterprise Network Array are hot-swappable.

**HBA:** Host Bus Adapter. The module which adapts system bus signals to those used by the A5x00.

**interconnect assembly:** The motherboard and centerplane, which are fastened together as a unit and contain data and power connections for the interface boards, backplanes, and power supplies.

**LRC:** Loop redundancy circuit. A switch located between a node and the rest of a loop that can automatically bypass the failed portion of the loop.

**node:** An entity on a loop recognizable by a WWN or addressed by the loop as a device; an addressable point on a loop.

**Node WWN:** Node World Wide Number. A unique number that specifies a device.

**Open Storage Networks:** Sun's new model for high-performance enterprise storage, under continuing development. The Sun Enterprise Network Array is the first product to reflect the Open Storage Networks approach.

**Port WWN:** Port World Wide Number. A unique number that specifies the port used to access a device.



---

**POE:** Process Of Elimination. A process used for troubleshooting where you break a system down into progressively smaller components while running tests to locate a failure. You determine or eliminate the functional components to isolate the bad.

**RAID:** Redundant array of inexpensive (or independent) disks. An approach to data storage that combines a number of disks to serve as a single virtual disk. The benefits are high levels of performance, availability, and -- depending on the RAID level used -- redundant storage of data for greater security.

**RAS:** Reliability, availability, serviceability. A set of standards used to evaluate mass storage subsystems in corporate environments. The Sun Enterprise Network Array ranks higher in RAS than any other Sun mass storage product.

**SES:** SCSI-3 Enclosure Services command set. The SCSI standard interface for reporting environmental data from a device to the host. SES firmware performs diagnostics on components in the disk array and controls the front panel module (FPM).

**SOC+:** Serial Optical Controller Plus. A Sun-proprietary ASIC that implements FC-AL for the host adapter and the Sun Enterprise Network Array.

**WWN:** Word Wide Number (also known as World Wide Name). See Node WWN and Port WWN.





**THE NETWORK IS THE COMPUTER™**

Sun Microsystems, Inc.  
901 San Antonio Road  
Palo Alto, CA 94303  
415 960-1300