CHAPTER 1

Troubleshooting Overview

This chapter introduces the basic concepts, methodology, and tools to use for troubleshooting problems that may occur when configuring and using a Cisco MDS 9000 Family switch. The two most common symptoms of problems occurring in a storage network are:

- Ahost not accessing its allocated storage
- An application not responding after attempting to access the allocated storage

To identify the possible problems, you need to use a variety of tools and understand the overall storage environment. For this reason, this chapter describes a number of general troubleshooting tools and procedures in addition to those that are specific to the Cisco MDS 9000 family. This chapter also provides a plan of attack for investigating storage issues. Refer to the other chapters in this book for detailed explanations of specific issues.

This chapter includes the following sections:

Introduction

Some basic questions should be answered before you go into too much detail about specific problems and solutions. A process of elimination can determine which network components have been subject to change and therefore may be the cause of your problems. The main steps you should follow to identify a problem in a SAN environment include:

1. Verify physical connectivity and registration to the fabric
2. Verify storage subsystem and server configuration
3. Verify end-to-end connectivity and fabric configuration

This section provides a series of questions that may be useful when troubleshooting a problem with a Cisco MDS 9000 family switch or connected devices. Use the answers to these questions to plan a course of action and to determine the scope of the problem. For example, if a host can only see some of the LUNs, (but not all of them) on an existing subsystem, then fabric-specific issues (FSPF, ISLs, FCNS) do not need to be investigated, as they are currently working correctly. The fabric components can therefore be eliminated from the problem.

By answering the questions in the following subsections, you can determine the paths you need to follow and the components that you should investigate further. These questions are independent of host, switch or subsystem vendor.

Basics

- Is this a newly installed system or an existing installation? (It could be a new SAN, host or subsystem, or new LUNs exported to an existing host.)
- Has the host ever been able to see its storage?
- Does the host recognize any LUNs in the subsystem?
- Are you trying to solve an existing application problem (too slow, too high latency, excessively long response time) or did the problem show up recently?
- What changed in the configuration or in the overall infrastructure immediately before the applications started to have problems?

Basic Connectivity

- Are you using the correct fiber (SM or MM)?
- Did you check for a broken fiber?
- Is the LED on the connected module Fibre Channel port green, and do the LEDs on any HBA/Storage Subsystem ports indicate normal functionality?
- Is there a LUN masking policy applied on the storage subsystem? If yes, is the server allowed to see the LUNs exported by the storage array?
- Is there any LUN masking policy configured on the host? Did you enable the server to see all the LUNs it can access?
- If LUN masking software is used, is the host's PWWN listed in the LUN masking database?
- Is the subsystem configured for NPort?

Examine the FLOGI database on the two switches that are directly connected to the host HBA and subsystem ports. Also, verify that both ports (attached port on MDS-A and MDS-B) are members of the same VSAN. If both devices are listed in the FCNS database then ISLs are not an issue.

Fibre Channel End-to-End Connectivity

Answering the following questions will help to determine if end-to-end fibre channel connectivity exists from a host or subsystem perspective:

- Does the host list the subsystem's WWPN or FCID in its logs?
- Does the subsystem list the host's WWPN or FCID in its logs or LUN masking database?
- Can the host complete a port login (PLOGI) to the storage subsystem?
- Is there any SCSI exchange that takes place between the server and the disk array?
- Is the HBA configured for NPort?

You can use the HBA configuration utilities or the host system logs to determine if the subsystem PWWN or FCID is listed as a device. This can validate that FSPF is working correctly.

Fabric Issues

- Are both the host bus adapter (HBA) and the subsystem port successfully registered with the fabric name server?
- Does the correct PWWN for the Server HBA and the storage subsystem port show up on the correct port in the FLOGI database? In other words, is the device plugged into the correct port?
- Does any single zone contain both devices? The zone members can be WWNs, FCIDs.
- Is the zone correctly configured and part of the active configuration or zoneset within the same VSAN?
- Do the ISLs show any VSAN isolation?
- Do the host and storage belong to the same VSAN?
- Are any parameters, such as FSPF, Static Domain Assignment, VSAN or Zoning, mismatched in the configuration of the different switches in the fabric?

Using Host Diagnostic Tools

Most host systems provide utilities or other tools that you can use for troubleshooting the connection to the allocated storage. For example, on a Windows system, you can use the Diskmon or Disk Management tool to verify accessibility of the storage and to perform some basic monitoring and administrative tasks on the visible volumes.

Alternatively, you can use Iometer, an I/O subsystem measurement and characterization tool, to generate a simulated load and measure performance. Iometer is a public domain software utility for Windows, originally written by Intel, that provides correlation functionality to assist with performance analysis.

Iometer measures the end-to-end performance of a SAN without cache hits. This can be an important measurement because if write or read requests go to the cache on the controller (a cache hit) rather than to the disk subsystems, performance metrics will be artificially high. You can obtain Iometer from SourceForge.net at the following URL:

```
http://sourceforge.net/projects/iometer/
```

Iometer is not the only I/O generator you can use to simulate traffic through the SAN fabric. Other popular I/O generators and benchmark tools used for SAN testing include Iozone and Postmark. Iozone is a file system benchmark tool that generates and measures a variety of file operations. It has been ported to many systems and is useful for performing a broad range of file system tests and analysis.

Postmark was designed to create a large pool of continually changing files, which simulates the transaction rates of a large Internet mail server.

PostMark generates an initial pool of random text files in a configurable range of sizes. Creation of the pool produces statistics on continuous small file creation performance. Once the pool is created, PostMark generates a specified number of transactions, each of which consists of a pair of smaller transactions:

- Create file or Delete file
- Read file or Append file

Benchmark is available from Network Appliance, Inc. at the following URL:

http://www.netapp.com/tech_library/3022.html

Benchmarking tools offer a variety of capabilities and you should select the one that provides the best I/O characteristics of your application environment.

Utilities provided by the Sun Solaris operating system let you determine if the remote storage has been recognized and exported to you in form of a raw device or mounted file system, and to issue some basic queries and tests to the storage. You can measure performance and generate loads using the **iostat** utility, the **perfmeter** GUI utility, the **dd** utility, or a third-party utility like Extreme SCSI.

Every UNIX version provides similar utilities, but this guide only provides examples for Solaris. Refer to the documentation for your specific operating system for details.

Using Cisco MDS 9000 Family Tools

If the server does not see its storage and you cannot use the information available on the host side to determine the root cause of the problem, you can obtain additional information from a different viewpoint using the troubleshooting tools provided with the Cisco MDS 9000 family of switches. This section introduces these tools and describes the kinds of problems for which you can use each tool. It includes the following topics:

Command-Line-Interface (CLI)

The Cisco MDS 9000 Family CLI lets you configure and monitor a Cisco MDS 9000 Family switch using a local console or remotely using a Telnet or SSH session. The CLI provides a command structure similar to Cisco IOS® software, with context-sensitive help, show commands, multi-user support, and roles-based access control.

CLI Debug

The Cisco MDS 9000 Family of switches includes an extensive debugging feature set for actively troubleshooting a storage network. Using the CLI, you can enable debugging modes for each switch feature and view a real-time updated activity log of the control protocol exchanges. Each log entry is time-stamped and listed in chronological order. Access to the debug feature can be limited through the CLI roles mechanism and can be partitioned on a per-role basis. While debug commands show realtime information, the **show** commands can be used to list historical information as well as realtime.

Note You can log debug messages to a special log file, which is more secure and easier to process than sending the debug output to the console.

By using the '?' option, you can see the options that are available for any switch feature, such as FSPF. A log entry is created for each entered command in addition to the actual debug output. The debug output shows a time-stamped account of activity occurring between the local switch and other adjacent switches.

You can use the debug facility to keep track of events, internal messages and protocol errors. However, you should be careful with using the debug utility in a production environment, because some options may prevent access to the switch by generating too many messages to the console or if very CPU-intensive may seriously affect switch performance.

Note It is a good idea to open a second Telnet or SSH session before entering any debug commands. That way, if the debug output comes too fast to stop it in the output window, you can use the second session to enter the **undebug all** command to stop the debug message output.

The following is an example of the output from the **debug flogi event** command

```
switch# debug flogi event interface fc1/1
Dec 10 23:40:26 flogi:     current state [FLOGI_ST_FLOGI_RECEIVED]
Dec 10 23:40:26 flogi:     current event [FLOGI_EV_VALID_FLOGI]
Dec 10 23:40:26 flogi:     next state    [FLOGI_ST_GET_FCID]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:     current state [FLOGI_ST_GET_FCID]
Dec 10 23:40:26 flogi:     current event [FLOGI_EV_VALID_FCID]
Dec 10 23:40:26 flogi:     next state    [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:     current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:     current event [FLOGI_EV_CONFIG_DONE_PENDING]
Dec 10 23:40:26 flogi:     next state    [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:     current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:     current event [FLOGI_EV_RIB_RESPOSE]
Dec 10 23:40:26 flogi:     next state    [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:     current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:     current event [FLOGI_EV_NAME_SERVER_REG_RESPONSE]
Dec 10 23:40:26 flogi:     next state    [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:     current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:     current event [FLOGI_EV_ACL_CFG_RESPONSE]
Dec 10 23:40:26 flogi:     next state    [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:     current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:     current event [FLOGI_EV_ZS_CFG_RESPONSE]
Dec 10 23:40:26 flogi:     next state    [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute_all: done processing event FLOGI_EV_ZS_CFG_RESPONSE
Dec 10 23:40:26 flogi:     current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:     current event [FLOGI_EV_LCP_RESPONSE]
Dec 10 23:40:26 flogi:     next state    [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi: fu_fsm_execute: ([1]21:00:00:e0:8b:08:96:22)
Dec 10 23:40:26 flogi:     current state [FLOGI_ST_PERFORM_CONFIG]
Dec 10 23:40:26 flogi:     current event [FLOGI_EV_CONFIG_DONE_COMPLETE]
Dec 10 23:40:26 flogi:     next state    [FLOGI_ST_FLOGI_DONE]
```

The following is a summary of the available debug commands:

Table 1-1 Debug commands

Debug command	Purpose
acl	Enable acl debugging
all	Enable all debugging

Table 1-1 Debug commands (continued)

Debug command	Purpose
ascii-cfg	Configure ascii-cfg debugging
bootvar	Enable bootvar debugging
callhome	Enable debugging for Callhome
fc2	Configure FC2 debugging
fcc	Enable FCC debugging
fcdomain	Enable fcdomain debugging
fcfwd	Enable fcfwd debugging
fcns	Debug name server
fcs	Configure Fabric Configuration Server Debugging
flogi	Configure flogi debug
fspf	Configure FSPF debugging
hardware	Debug hardware, kernel loadable module parameters
ipconf	Enable IP configuration debugging
ipfc	Enable IPFC debugging
klm	Debug kernel loadable module parameters
logfile	Direct debug output to logfile
mcast	Enable mcast debugging
mip	Debug multiple IP kernel driver
module	Configure LC Manager debugging
ntp	Debug NTP module
platform	Configure Platform Manager debugging
port	Configure port debugging
port-channel	Enable port-channel debug
qos	Configure QOS Manager Debugging
rdl	Configure RDL debugging
rib	Configure rib debugging
rscn	Configure RSCN debugging
scsi-target	Configure scsi target daemon debugging
security	Configure debugging for security/accounting
sensor	Enable Sensor Mgr debugging
span	Configure SPAN debug
system	Enable System debugging
tlport	Configure TL Port debugging
vni	Enable virtual network interface debugging
vrrp	Enable vrrp debugging
vsan	Enable VSAN manager debugging

Table 1-1 Debug commands (continued)

Debug command	Purpose
vsh	Enable vsh debugging
vshd	Configure vshd debugging
wwn	Configure WWN Manager Debugging
xbar	Enable xbar debugging
xbc	Enable Xbar Client debugging
zone	Zone server debug commands

FC Ping and FC Traceroute

Note FC Ping and FC traceroute are used to troubleshoot problems with connectivity and path choices. They are not designed for use in identifying or resolving performance issues.

Ping and Traceroute are two of the most useful tools for troubleshooting TCP/IP networking problems. The *Ping* utility generates a series of *echo* packets to a destination across a TCP/IP internetwork. When the *echo* packets arrive at the destination, they are re-routed and sent back to the source. Using *Ping*, you can verify connectivity and latency to a particular destination across an IP routed network. *Traceroute* operates in a similar fashion, but can also determine the specific path that a frame takes to its destination on a hop-by-hop basis. These tools have been migrated to Fibre Channel for use with the Cisco MDS 9000 Family switches and are called *FC Ping* and *FC Traceroute*. You can use *FC Ping* and *FC Traceroute* from the CLI or from the Cisco Fabric Manager.

FC Ping allows you to ping a Fibre Channel *N_Port* or end device. By specifying the FC_ID or Fibre Channel address, you can send a series of frames to a target *N_Port*. Once these frames reach the outgoing *F_Port*, they are looped back to the source and a time-stamp is taken. *FC Ping* helps you to verify the connectivity and latency to an end *N_Port*. *FC Ping* uses the PRLI Extended Link Service, and verifies the presence of a FC entity in case of positive or negative answers.

FC Traceroute is slightly different than the IP equivalent because both the outbound and return paths are recorded as these may differ in a switched Fibre Channel network. The *FC Traceroute* command identifies the path taken on a hop-by-hop basis and includes a timestamp at each hop in both directions.

FC Ping and *FC Traceroute* are useful tools to check for network connectivity problems or verify the path taken toward a specific destination. You can use *FC Traceroute* to test the connectivity of TE ports along the path between the generating switch and the switch closest to the destination.

Note FC Trace will only work across EISL links.

The following is an example of output from the **fcping** command:

```
switch# fcping fcid 0xef02c9 vsan 1
28 bytes from  0xef02c9  time = 1408 usec
28 bytes from  0xef02c9  time = 379 usec
28 bytes from  0xef02c9  time = 347 usec
28 bytes from  0xef02c9  time = 361 usec
28 bytes from  0xef02c9  time = 363 usec

5 frames sent, 5 frames received, 0 timeouts
Round-trip min/avg/max = 347/571/1408 usec
```

The following is an example of output from the **fctrace** command:

```
switch# fctrace fcid 0xef0000 vsan 1
Route present for :  0xef0000
20:00:00:05:30:00:59:de(0xfffcee)
Latency: 0 msec
20:00:00:05:30:00:58:1e(0xfffc6c)
Timestamp Invalid.
20:00:00:05:30:00:59:1e(0xfffcef)
Latency: 0 msec
20:00:00:05:30:00:59:1e(0xfffcef)
Latency: 174860 msec
20:00:00:05:30:00:58:1e(0xfffc6c)
```

Cisco Fabric Manager

Cisco Fabric Manager provides fabric-wide management capabilities including discovery, multiple switch configuration, network monitoring, and troubleshooting. It provides the troubleshooting features described in the following topics:

- Fabric Manager and Device Manager, page 1-8
- Analyzing Switch Device Health, page 1-9
- Analyzing End-to-End Connectivity, page 1-10
- Analyzing Switch Fabric Configuration, page 1-10
- Analyzing the Results of Merging Zones, page 1-11
- Alerts and Alarms, page 1-12

Note For detailed information about using Cisco Fabric Manager, refer to the *Cisco MDS 9000 Family Fabric Manager User Guide*.

Fabric Manager and Device Manager

Fabric Manager provides a map of the discovered fabric and includes tables that display statistical information about the switches in the fabric. You can also select troubleshooting tools from the Fabric Manager Troubleshooting menu.

Note When you click on a zone or VSAN in Fabric Manager, the members of the zone or VSAN are highlighted on the Fabric Manager Map pane.

Device Manager provides a graphic display of a specific switch and shows the status of each port on the switch. From Device Manager, you can drill down to get detailed statistics about a specific switch or port.

Figure 1-1 shows the Fabric Manager Summary View window.

Figure 1-1 Cisco Fabric Manager Summary View

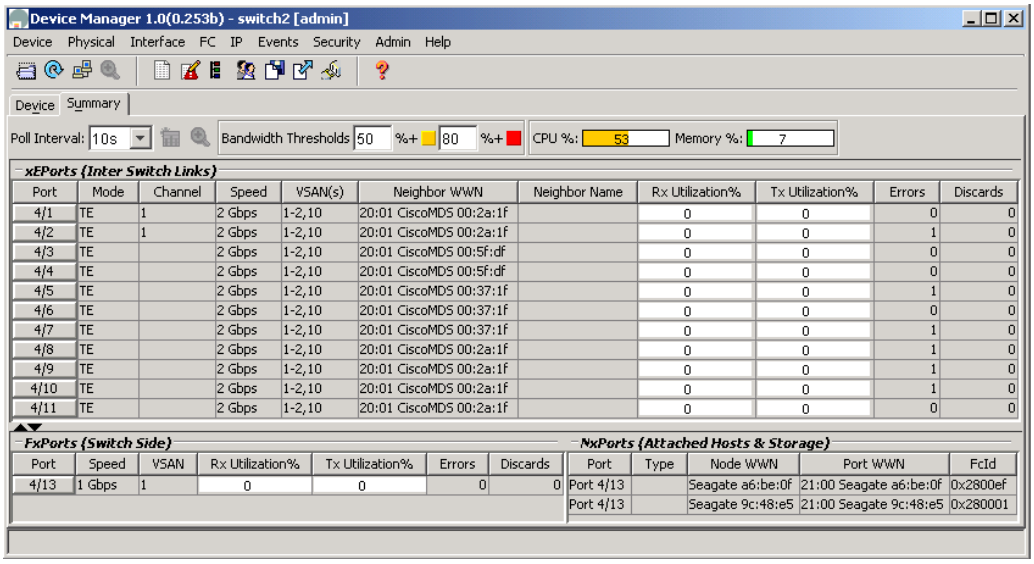

The Summary View window lets you analyze switch performance issues, diagnose problems, and change parameters to resolve problems or inconsistencies. This view shows aggregated statistics for the active Supervisor Module and all switch ports. Information is presented in tabular or graphical formats, with bar, line, area, and pie chart options. You can also use the Summary View to capture the current state of information for export to a file or output to a printer.

Analyzing Switch Device Health

Choose the Switch Health option from the Fabric Manager Troubleshooting menu to determine the status of the components of a specific switch.

Figure 1-2 Switch Health Analysis Window

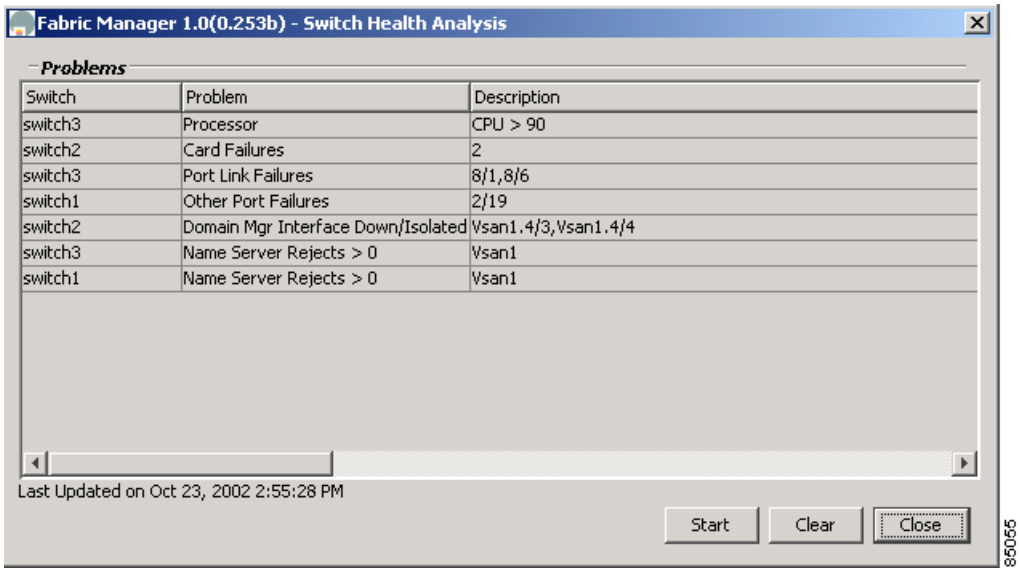

The Switch Health Analysis window displays any problems affecting the selected switches.

Analyzing End-to-End Connectivity

Select the End to End Connectivity option from the Fabric Manager Troubleshooting menu to determine connectivity and routes among devices with the switch fabric. The connectivity tool checks to see that every pair of end devices in an active zone can talk to each other, using a Ping test and by determining if they are in the same VSAN. This option uses versions of the **ping** and **traceroute** commands modified for Fibre Channel networks.

The End to End Connectivity Analysis window displays the selected end points with the switch to which each is attached, and the source and target ports used to connect it.

The output shows all the requests which have failed. The possible descriptions are:

- Ignoring empty zone—No requests are issued for this zone.
- Ignoring zone with single member—No requests are issued for this zone.
- Source/Target are unknown—No nameserver entries exist for the ports or we have not discovered the port during discovery.
- Both devices are on the same switch—The devices are not redundantly connected.
- No paths exist.
- Only one unique path exists.
- VSAN does not have an active zone set.
- Average time... micro secs—The latency value was more than the threshold supplied.

Analyzing Switch Fabric Configuration

Select the Fabric Configuration option from the Fabric Manager Troubleshooting menu to analyze the configuration of a switch by comparing the current configuration to a specific switch or to a policy file. You can save a switch configuration to a file and then compare all switches against the configuration in the file.

Figure 1-3 Fabric Configuration Analysis Window

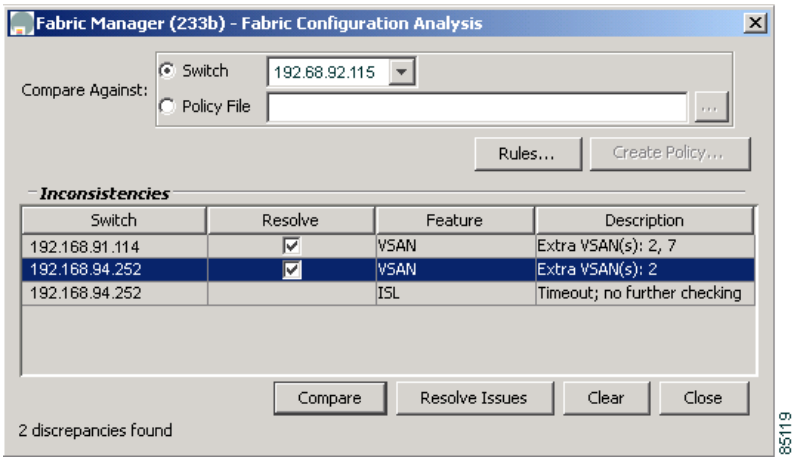

You use a policy file to define the rules to be applied when running the Fabric Checker. When you create a policy file, the system saves the rules selected for the selected switch.

Analyzing the Results of Merging Zones

Cisco Fabric Manager provides a very useful tool for troubleshooting problems that occur when merging zones configured on different switches.

Select the **Zone Merge** option on the Fabric Manager Troubleshooting menu to determine if two connected switches have compatible zone configurations.

Figure 1-4 Zone Merge Analysis Window

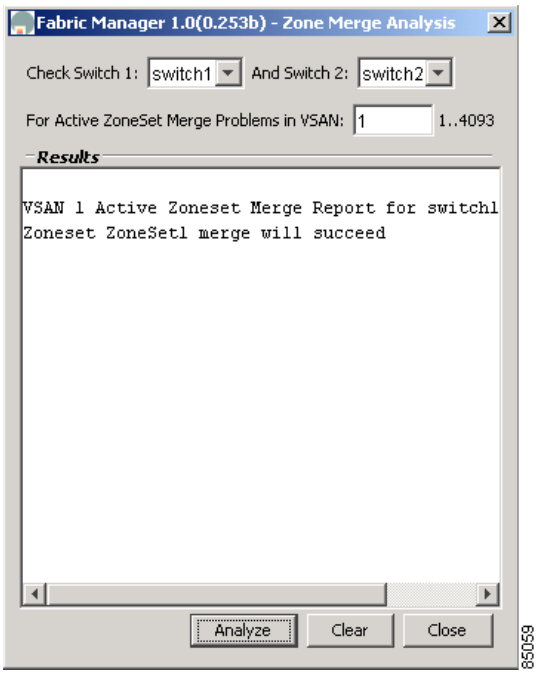

The Zone Merge Analysis window displays any inconsistencies between the zone configuration of the two selected switches.

You can use the following options on the Fabric Manager Troubleshooting menu to verify connectivity to a selected object or to open other management tools:

- Traceroute—Verify connectivity between two end devices that are currently selected on the Map pane.
- Device Manager— Launch Device Manager for the switch selected on the Map pane.
- Command Line Interface—Open a Telnet or SSH session for the switch selected on the Map pane.

Figure 1-5 Traceroute Window

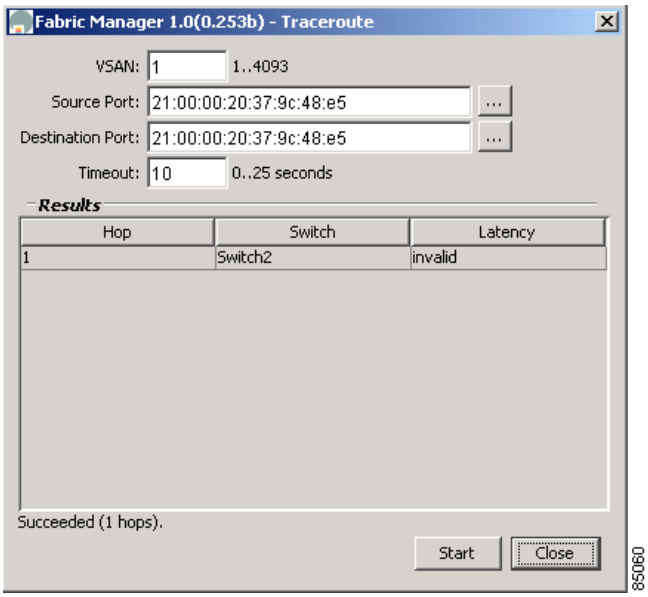

Alerts and Alarms

You can configure and monitor SNMP, RMON, Syslog, and Call Home alarms and notifications using the different options on the Device Manager Events menu. SNMP provides a set of preconfigured traps and informs that are automatically generated and sent to the destinations (trap receivers) that you identify. The RMON Threshold Manager lets you configure thresholds for specific events that trigger log entries or notifications. You can use either Fabric Manager or Device Manager to identify Syslog servers that will record different events or to configure Call Home, which can alert you through e-mail messages or paging when specific events occur.

SCSI Target Discovery

For more information about SCSI target discovery, refer to the *Cisco MDS 9000 Family Configuration Guide*.

The Fibre Channel name service is a distributed service in which all connected devices participate. As new SCSI target devices attach to the fabric, they register themselves with the name service, which is then distributed among all participating fabric switches. This information can then be used to help determine the identity and topology of nodes connected to the fabric.

For the Cisco MDS 9000 Family of switches, the SCSI Target Discovery feature has been added to provide added insight into connected SCSI targets. This feature allows the switch to briefly log into connected SCSI target devices and issue a series of SCSI inquiry commands to help discover additional information. The additional information that is queried includes logical unit number (LUN) details including the number of LUNs, the LUN IDs, and the sizes of the LUNs.

This information is then compiled and made available to through CLI commands, through the Cisco Fabric Manager, and also via an embedded SNMP MIB which allows the information to be easily retrieved by an upstream management application. Using the *SCSI Target Discovery* feature, you can have a much more detailed view of the fabric and its connected SCSI devices.

The following is an example of output from the discover scsi-target command:

```
switch# discover scsi-target local remote
discovery started
switch# show scsi-target lun vsan 1
- ST318203FC from SEAGATE (Rev 0004)
  FCID is 0xef02b5 in VSAN 1, PWWN is 21:00:00:20:37:46:78:97
  --------------------------------------------------------------------------------
  LUN    Capacity  Status  Serial Number    Device-Id
         (MB)
  --------------------------------------------------------------------------------
  0x0    18210     Online  LRA2510000007027 C:1 A:0 T:3 20:00:00:20:37:46:78:97
- ST318203FC from SEAGATE (Rev 0004)
  FCID is 0xef02b6 in VSAN 1, PWWN is 21:00:00:20:37:5b:cf:b9
  --------------------------------------------------------------------------------
  LUN    Capacity  Status  Serial Number    Device-Id
         (MB)
  --------------------------------------------------------------------------------
  0x0    18210     Online  LR94873000007029 C:1 A:0 T:3 20:00:00:20:37:5b:cf:b9
- ST318203FC from SEAGATE (Rev 0004)
  FCID is 0xef02b9 in VSAN 1, PWWN is 21:00:00:20:37:18:6f:90
  --------------------------------------------------------------------------------
  LUN    Capacity  Status  Serial Number    Device-Id
         (MB)
--------------------------------------------------------------------------------
0x0    18210     Online  LR18591800001004 C:1 A:0 T:3 20:00:00:20:37:18:6f:90
```

Note This tool can be effective to find out the number of LUNs exported by a storage subsystem, but it may be ineffective when LUN Zoning/LUN Security tools are used.

SNMP and RMON Support

The Cisco MDS 9000 Family of switches provides extensive SNMPv1, v2, and v3support, including Management Information Bases (MIBs) and notifications (traps and informs).

The applications provided by Cisco that use SNMP include Fabric Manager, Cisco SSE and Cisco RME. Also, the SNMP standard allows any third-party applications that support the different MIBs to manage and monitor Cisco MDS 9000 Family switches.

SNMP v3 provides extended security. Each switch can be selectively enabled or disabled for SNMP service. In addition, each switch can be configured with a method of handling SNMPv1 and v2 requests.

Note During initial configuration of your switch, the system prompts you to define SNMP v1 or V2 community strings and to create a SNMP v3 username and password.

Cisco MDS 9000 Family switches support over 50 different MIBs, which can be divided into the following six categories:

- IETF Standards-based Entity MIBs (for example, RFC273□ENTITY-MIB)□These MIBs are used to report information on the physical devices themselves in terms of physical attributes etc.
- Cisco-Proprietary Entity MIBs (for example, CISCO-ENTITY-FRU-CONTROL-MIB)□These MIBs are used to report additional physical device information about Cisco-only devices such as their configuration.

- IETF IP Transport-oriented MIBs (for example, RFC2013□UDP-MIB)□These MIBs are used to report transport-oriented statistics on such protocols as IP, TCP, and UDP. These transports are used in the management of the Cisco MDS 9000 Family through the OOB Ethernet interface on the Supervisor module.
- Cisco-Proprietary Storage and Storage Network MIBs (for example, NAME-SERVER-MIB)□These MIBs were written by Cisco to help expose information that is discovered within a fabric to management applications not connected to the fabric itself. In addition to exposing configuration details for features like zoning and Virtual SANs (VSANs) via MIBs, discovered information from sources like the FC-GS-3 Name Server can be pulled via a MIB. Additionally, MIBs are provided to configure/enable features within the Cisco MDS 9000 Family. There are over 20 new MIBs provided by Cisco for this information and configuration capability.
- IETF IP Storage Working Group MIBs (for example, ISCSI-MIB)□While many of these MIBs are still work-in-progress, Cisco is helping to draft such MIBs for protocols such as iSCSI and Fibre Channel-over-IP (FCIP) to be standardized within the IETF.
- Miscellaneous MIBs (for example, SNMP-FRAMEWORK-MIB)□There are several other MIBs provided in the Cisco MDS 9000 Family switches for tasks such as defining the SNMP framework or creating SNMP partitioned views.

You can use SNMPv3 to assign different SNMP capabilities to specific roles.

Cisco MDS 9000 Family switches also support Remote Monitoring (RMON) for Fibre Channel. RMON provides a standard method to monitor the basic operations of network protocols providing connectivity between SNMP management stations and monitoring agents. RMON also provides a powerful alarm and event mechanism for setting thresholds and sending notifications based on changes in network behavior.

The RMON groups that have been adapted for use with Fibre Channel include the *AlarmGroup* and *EventGroup*. The *AlarmGroup* provides services to set alarms. Alarms can be set on one or multiple parameters within a device. For example, you can set an RMON alarm for a specific level of CPU utilization or crossbar utilization on a switch. The *EventGroup* lets you configure events that are actions to be taken based on an alarm condition. The types of events that are supported include *logging*, *SNMP traps*, and *log-and-trap*.

Note To configure events within an RMON group, use the **Events** > **Threshold Manager** option from Device Manager.

Using RADIUS

RADIUS is fully supported for the Cisco MDS 9000 Family switches through the Fabric Manager and the CLI. RADIUS is a protocol used for the exchange of attributes or credentials between a head-end RADIUS server and a client device. These attributes relate to three classes of services:

- Authentication
- Authorization
- Accounting

Authentication refers to the authentication of users for access to a specific device. You can use RADIUS to manage user accounts for access to Cisco MDS 9000 Family switches. When you try to log into a switch, the switch validates you with information from a central RADIUS server.

Authorization refers to the scope of access that you have once you have been authenticated. With Cisco MDS 9000 Family switches, assigned roles for users can be stored in a RADIUS server along with a list of actual devices that the user should have access to. Once the user has been authenticated, then switch can then refer to the RADIUS server to determine the extent of access the user will have within the switch network.

Accounting refers to the log information that is kept for each management session in a switch. This information may be used to generate reports for troubleshooting purposes and user accountability. Accounting can be implemented locally or remotely (using RADIUS).

The following is an example of an accounting log entries.

```
switch# show accounting log
Sun Dec 15 04:02:27 2002:start:/dev/pts/0_1039924947:admin
Sun Dec 15 04:02:28 2002:stop:/dev/pts/0_1039924947:admin:vsh exited normally
Sun Dec 15 04:02:33 2002:start:/dev/pts/0_1039924953:admin
Sun Dec 15 04:02:34 2002:stop:/dev/pts/0_1039924953:admin:vsh exited normally
Sun Dec 15 05:02:08 2002:start:snmp_1039928528_172.22.95.167:public
Sun Dec 15 05:02:08 2002:update:snmp_1039928528_172.22.95.167:public:Switchname
set to Switch
```

Note The accounting log only shows the beginning and ending (start and stop) for each session.

Using Syslog

Syslog lets you store a chronological log of system messages locally or sent to a central Syslog server. Syslog messages can also be sent to the console for immediate use. These messages can vary in detail depending on the configuration that you choose. Syslog messages are categorized into 7 severity levels from *debug to critical* events. You can limit the severity levels that are reported for specific services within the switch. For example, you may wish only to report *debug* events for the *FSPF* service but record all severity level events for the *Zoning* service.

A unique feature within the Cisco MDS 9000 Family switches is the ability to send RADIUS accounting records to the Syslog service. The advantage of this feature is that you can consolidate both types of messages for easier correlation. For example, when you log into a switch and change an FSPF parameter, Syslog and RADIUS provide complimentary information that will help you formulate a complete picture of the event.

Using Fibre Channel SPAN

For more information about configuring SPAN, refer to the *Cisco MDS 9000 Family Configuration Guide*.

You can use the Switched Port Analyzer (SPAN) utility to perform detailed troubleshooting or to take a sample of traffic from a particular application host for proactive monitoring and analysis. This utility is most helpful when you have a Fibre Channel protocol analyzer available and you are monitoring user traffic between two FC IDs.

When you have a problem in your storage network that you cannot solve by fixing the device configuration, you typically need to take a look at the protocol level. You can use debug commands to look at the control traffic between an end node and a switch. However, when you need to focus on all the traffic originating from or destined to a particular end node such as a host or a disk, you can use a protocol analyzer to capture protocol traces.

To use a protocol analyzer, you must insert the analyzer in-line with the device under analysis, which disrupts input and output (I/O) to and from the device. This problem is worse when the point of analysis is on an Inter-Switch Link (ISL) link between two switches. In this case, the disruption may be significant depending on what devices are downstream from the severed ISL link.

In Ethernet networks, this problem can be solved using the SPAN utility, which is provided with the Cisco Catalyst Family of Ethernet switches. SPAN has also been implemented with the Cisco MDS 9000 Family of switches for use in Fibre Channel networks. SPAN lets you take a *copy* of all traffic and direct it to another port within the switch. The process is non-disruptive to any connected devices and is facilitated in hardware, which prevents any unnecessary CPU load. Using Fibre Channel SPAN, you can connect a Fibre Channel analyzer, such as a Finisar analyzer, to an unused port on the switch and then SPAN a copy of the traffic from a port under analysis to the analyzer in a non-disruptive fashion.

SPAN allows you to create up to 16 independent *SPAN* sessions within the switch. Each session can have up to four unique sources and one destination port. In addition, you can apply a filter to capture only the traffic received or the traffic transmitted. With Fibre Channel SPAN, you can even capture traffic from a particular Virtual SAN (VSAN).

To start the SPAN utility use the CLI command **span session** *session_num*, where *session_num* identifies a specific SPAN session. When you enter this command, the system displays a submenu, which lets you configure the destination interface and the source VSAN or interfaces.

```
switch2# config terminal
switch2(config)# span session 1
<<Create a span session>>

switch2(config-span)# source interface fc1/8
<<specify the port to be spanned>>

switch2(config-span)# destination interface fc1/3
<<specify the span destination (SD) port>>

switch2(config-span)# end
switch2# show span session 1
Session 1 (active)
   Destination is fc1/1
   No session filters configured
   Ingress (rx) sources are
     fc1/8,
   Egress (tx) sources are
     fc1/8,
```

Using Cisco Network Management Products

This section describes network management tools that are available from Cisco and are useful for troubleshooting problems with Cisco MDS 9000 Family switches and connected devices. It includes the following topics:

Cisco MDS 9000 Port Analyzer Adapter

The Cisco MDS 9000 Port Analyzer Adapter is a stand-alone adapter card that converts Fibre Channel (FC) frames to Ethernet frames by encapsulating each Fibre Channel frame into an Ethernet frame. This product is meant to be used for analyzing SPAN traffic from a Fibre channel port on a Cisco MDS 9000 Family switch.

The Cisco Port Analyzer Adapter provides two physical interfaces:

- An FC interface that connects to the SPAN port of a Cisco MDS 9000 Family switch
- A 100/1000 Mb/s Ethernet port that forwards the encapsulated Fibre Channel traffic with a broadcast destination MAC Address

Note The Cisco Port Analyzer Adapter does not support half-duplex mode and for this reason, it will not work when connected to a hub.

The Cisco Port Analyzer Adapter provides the following features:

- Encapsulates FC frames into Ethernet frames
- Sustains 32 max size FC frames burst (in 100 Mb/s mode)
- Line rate at 1Gb/s (for FC frames bigger than 91bytes)
- 64KBytes of onboard frame buffer
- Configurable option for Truncating FC frames to 256 Bytes - for greater burst
- Configurable option for Deep Truncating FC frames to 64 Bytes - best frames burst
- Configurable option for Ethernet Truncating FC frames to 1496 Bytes - max size E-net frames
- Configurable option for No Truncate Mode - sends jumbo frames on E-net side.
- Packet Counter (Indicates number of previous packet drops)
- SOF/EOF type information embedded
- 100/1000 Mb/s Ethernet interface - option on board
- Auto Configuration on power up
- Fibre Channel and Ethernet Link up indicator - LEDs.
- Checks FC frame CRC

When used in conjunction with the open source protocol analyzer, Ethereal (http//www.ethereal.com), the Cisco Port Analyzer Adapter provides a cost-effective and powerful troubleshooting tool. It allows any PC with a Ethernet card to provide the functionality of a flexible Fibre Channel analyzer. For more information on using the Cisco Port Analyzer Adapter see the *Cisco MDS 9000 Family Port Analyzer Adapter Installation and Configuration Guide*.

Cisco Fabric Analyzer

For more information about using the Cisco Fabric Analyzer, refer to the *Cisco MDS 9000 Family Configuration Guide*.

The ultimate tool for troubleshooting network protocol problems is the protocol analyzer. Protocol analyzers promiscuously capture network traffic and completely decode the captured frames down to the protocol level. Using a protocol analyzer, you can conduct a detailed analysis by taking a sample of a

storage network transaction and by mapping the transaction on a frame-by-frame basis, complete with timestamps. This kind of information lets you pinpoint a problem with a high degree of accuracy and arrive at a solution more quickly. However, dedicated protocol analyzers are expensive and they must be placed locally at the point of analysis within the network.

With the Cisco Fabric Analyzer, Cisco has brought Fibre Channel protocol analysis within a storage network to a new level of capability. Using Cisco Fabric Analyzer, you can capture Fibre Channel control traffic from a switch and decode it without having to disrupt any connectivity, and without having to be present locally at the point of analysis.

The Cisco Fabric Analyzer consists of three main components:

- An agent embedded in the Cisco MDS 9000 Family switches. This agent can be selectively enabled to promiscuously capture designated control traffic.
- A text-based interface to the control and decoded output of the analyzer.
- GUI-based client application that you can install on any workstation to provide a full-function interface to the decoded data.

The text-based interface to the Cisco Fabric Analyzer is a CLI-based program for controlling the analyzer and providing output of the decoded results. Using the CLI-based interface, you can remotely access an Cisco MDS 9000 Family switch, using Telnet or a secure method such as Secure Shell (SSH). You can then capture and decode Fibre Channel control traffic, which offers a convenient method for conducting detailed, remote troubleshooting. In addition, because this tool is CLI-based, you can use roles-based policies to limit access to this tool as required.

The GUI-based implementation (Ethereal) can be installed on any Windows or Linux workstation. This application provides an easier-to-use interface that is more easily customizable. The GUI interface lets you easily sort, filter, crop, and save traces to your local workstation.

The Ethereal application allows remote access to Fibre Channel control traffic and does not require a Fibre Channel connection on the remote workstation.

The Cisco Fabric Analyzer lets you capture and decode Fibre Channel traffic remotely over Ethernet. It captures Fibre Channel traffic, encapsulates it in TCP/IP, and transports it over an Ethernet network to the remote client. The remote client then deencapsulates and fully decodes the Fibre Channel frames. This capability provides flexibility for troubleshooting problems in remote locations.

The Cisco Fabric Analyzer captures and analyzes control traffic coming to the Supervisor Card. This tool is much more effective than the debug facility for packet trace and traffic analysis, because it is not very CPU intensive and it provides a graphic interface for easy analysis and decoding of the captured traffic.

```
switch# config terminal
switch(config)# fcanalyzer local brief
Capturing on eth2
  0.000000   ff.ff.fd -> ff.ff.fd   SW_ILS 1    0x59b7 0xffff  0x7 ->  0xf HLO
  0.000089   ff.ff.fd -> ff.ff.fd   FC     1    0x59b7 0x59c9 0xff ->  0x0 Link Ctl, ACK1
  1.991615   ff.ff.fd -> ff.ff.fd   SW_ILS 1    0x59ca 0xffff 0xff ->  0x0 HLO
  1.992024   ff.ff.fd -> ff.ff.fd   FC     1    0x59ca 0x59b8  0x7 ->  0xf Link Ctl, ACK1

fcanalyer example of fully decoded frame.
switch2(config)# fcanalyzer local
Capturing on eth2
Frame 1 (96 bytes on wire, 96 bytes captured)
    Arrival Time Jan 13, 2003 135038.787671000
    Time delta from previous packet 0.000000000 seconds
    Time relative to first packet 0.000000000 seconds
    Frame Number 1
    Packet Length 96 bytes
    Capture Length 96 bytes
Ethernet II, Src 00000000000a, Dst 00000000ee00
    Destination 00000000ee00 (00000000ee00)
```

```
    Source 00000000000a (00000000000a)
    Type Vegas FC Frame Transport (0xfcfc)
MDS Header(SOFf/EOFn)
    MDS Header
        Packet Len 66
        .... 0000 0001 11.. = Dst Index 0x0007
        .... ..00 1111 1111 = Src Index 0x00ff
        .... 0000 0000 0001 = VSAN 1
    MDS Trailer
        EOF EOFn (3)
Fibre Channel
    R_CTL 0x02
    Dest Addr ff.fc.7e
    CS_CTL 0x00
    Src Addr ff.fc.7f
    Type SW_ILS (0x22)
    F_CTL 0x290000 (Exchange Originator, Seq Initiator, Exchg First, Seq Last,
CS_CTL, Transfer Seq Initiative, Last Data Frame - No Info, ABTS - Abort/MS, )
    SEQ_ID 0x11
    DF_CTL 0x00
    SEQ_CNT 0
    OX_ID 0x5a06
    RX_ID 0x0000
    Parameter 0x00000000
SW_ILS
    Cmd Code SW_RSCN (0x1b)
    0010 .... = Event Type Port is offline (2)
    .... 0000 = Address Format Port Addr Format (0)
    Affected Port ID 7f.00.01
    Detection Function Fabric Detected (0x00000001)
    Num Entries 1
    Device Entry 0
    Port State 0x20
    Port Id 7f.00.01
    Port WWN 1000000530005f1f (000530)
    Node WWN 1000000530005f1f (000530)
```

However, the Cisco Fabric Analyzer is not the right tool for troubleshooting end-to-end problems because it cannot access any traffic between the server and storage subsystems. That traffic is switched locally on the linecards, and does not reach the Supervisor card. In order to debug issues related to the communication between server and storage subsystems, you need to use Fibre Channel SPAN with an external protocol analyzer.

There are two ways you can start the Cisco Fabric Analyzer from the CLI.

- **fcanalyzer local**—Launches the text-based version on the analyzer directly on the console screen or on a file local to the system.
- **fcanalyzer remote** *ip address*—Activates the remote capture agent on the switch, where *ip address* is the address of the management station running Ethereal.

CiscoWorks RME

CiscoWorks Resource Manager Essentials (RME) is a set of CiscoWorks applications that provide comprehensive resource management capabilities. With the introduction of the Cisco MDS 9000 Family of switches, CiscoWorks RME has been extended to provide resource management services to a Cisco MDS 9000 Family storage network.

CiscoWorks RME comprises a set of resource management services. The following list outlines the services provided by CiscoWorks RME for the Cisco MDS 9000 Family of switches.

- Inventory Manager—Provides a facility to gather and audit a detailed hardware and software inventory of all Cisco MDS 9000 Family devices deployed in the storage network. A reporting facility is included to generate inventory reports
- Configuration Manager—Maintains an active repository of device configuration files for devices that are managed. It provides facility to upload and download configuration files to/from devices and a facility to log a record in the Change Audit log database when a new version of the configuration file is archived. Standard reports can be generated for configuration management inventory and activity.
- Configuration Editor—Provides a powerful web-based editor that allows multiple configuration files to be checked out of the configuration archive, be updated or changed, and then either saved locally or downloaded to the device.
- Net Show—Provides a simplified web-based show command interface, allowing show commands to be run against multiple switches or routers to enhance and simplify network troubleshooting.
- Software Image Manager—Simplifies version management and routine deployment of software updates to Cisco devices through wizard-assisted planning, scheduling, downloading, and monitoring of software updates
- Syslog Analyzer—Filters Syslog messages logged by Cisco devices and displays explanations of probable causes and recommended actions. This tool also helps facilitate manual parsing of Syslog files for reporting purposes.

CiscoWorks RME provides a system that can manage hardware, software, and configuration inventory across multiple infrastructures including storage networks, LANs, MANs, and WANs.

Using Other Troubleshooting Products

This section describes products from other vendors that you might find useful when troubleshooting problems with your storage network and connected devices. It includes the following topics:

Fibre Channel Testers

Fibre Channel testers are generally used to troubleshoot low-level protocol functions (such as Link Initialization). Usually these devices operate at 1- or 2-Gbps and provide the capability to create customized low-level Fibre Channel primitive sequences.

Fibre Channel testers are primarily used to ensure physical connectivity and low-level protocol compatibility, such as with different operative modes like Point-to-Point or Loop mode.

Fibre Channel testers and more generalized optical testers may used to spot broken cables, speed mismatch, link initialization problems and transmission errors. These devices sometimes incorporate higher-level protocol analysis tools and may bundled with generic protocol analyzers.

Fibre Channel Protocol Analyzers

An external protocol analyzer (for example from Finisar), is capable of capturing and decoding link level issues and the fibre channel ordered sets which comprise the fibre channel frame. The Cisco Port Analyzer Adapter, does not capture and decode at the ordered set level.

A Fibre Channel protocol analyzer captures transmitted information from the physical layer of the Fibre Channel network. Because these devices are physically located on the network instead of at a software re-assembly layer like most Ethernet analyzers, Fibre Channel protocol analyzers can monitor data from the 8b/10b level all the way to the embedded upper-layer protocols.

Fibre Channel network devices (HBAs, switches, and storage subsystems) are not able to monitor many SAN behavior patterns. Also, management tools that gather data from these devices are not necessarily aware of problems occurring at the Fibre Channel physical, framing, or SCSI upper layers for a number of reasons.

Fibre Channel devices are specialized for handling and distributing incoming and outgoing data streams. When devices are under maximum loads, which is when problems often occur, the device resources available for error reporting are typically at a minimum and are frequently inadequate for accurate error tracking. Also, Fibre Channel host bus adapters (HBAs) do not provide the ability to capture raw network data.

For these reasons, a protocol analyzer may be more important in troubleshooting a storage network than in a typical Ethernet network. There are a number of common SAN problems that occur in deployed systems and test environments that are visible only with a Fibre Channel analyzer. These include the following:

- Credit starvation
- missing, malformed, or non-standard-compliant frames or primitives
- protocol errors