



NETWORKERS 2004

TROUBLESHOOTING MDS9000 IP STORAGE AREA NETWORKS

SESSION OPT-3052

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

1

Session Objectives

Cisco.com

- **Take existing knowledge of Storage Area Networking Fundamentals and apply them the Data Networking for IP Storage**
- **Expand on the knowledge of the Fibre Channel operations and how iSCSI and FCIP fit in**
- **Introduce the tools of the trade required to troubleshoot and diagnose IP Storage issues with the Cisco MDS and SN5428 Family of Switches**

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

2

OPT-3052 - Agenda

Cisco.com

- **iSCSI MDS and SN5428 Architecture**
- **iSCSI Configuration and Protocol Troubleshooting**
- **Troubleshooting the interaction with the iSCSI host**
- **Overview of FCIP implementation on MDS & SN5428-2**
- **FCIP protocol troubleshooting**

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

3

OPT-3052 Reference Network

Cisco.com

- **All outputs and references for this session will be taken from the same working network built for this Networkers session.**
- **Network is as true to a production SAN as possible**
- **Complete Network Configurations and topology maps are included in the handouts**
- **Use session materials as a workbook for troubleshooting your network**


OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

4

Topology

Cisco.com



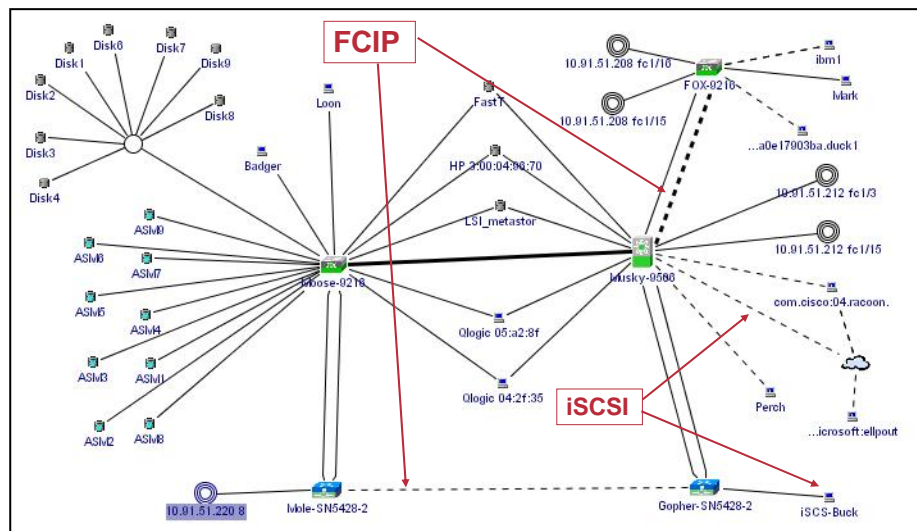
OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

5

Topology as Seen by Fabric Manager

Cisco.com



OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

6

MDS Switch Troubleshooting Levels

Cisco.com

- **Command Line Debugs**
Used by network level 2-3 engineers
- **Analyzers and SPAN**
Used by Level 3 Engineers
- **Fabric Manager & Device Manager**
Used by Network Operators and Level 1 Engineers
- **General Statistics gathering from show commands and GUI Managers**
Network Admin's and Operators Level 1 engineering

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

7

Debug 101: Capture Methods

Cisco.com

- **Direct to screen**
Using Windows hyperterm or any console/telnet/ssh utility
- **To a Log File**
Direct output to file within MDS
- **From Command Line or from within Configuration**
From CLI you can be from admin prompt or dropped into configuration mode to run debugs

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

8

Gathering Traces for Analysis

Cisco.com

All Non Disruptive to Switch Operations and Traffic On The SAN

- Built-in FCAalyzer
- Ethereal on PC
- MDS Port Analyzer Adapter
- Span Use
- Using an external FC Analyzers

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

9

Ethereal: Protocol Analyzer

Cisco.com

- iSCSI, FCIP, SCSI and Fibre Channel Protocols Supported in Ethereal version 0.9.9 and higher. *Use latest Ethereal ,Currently 0.10.3*
- Analyzer can be run in 2 modes depending on what traffic you are troubleshooting.
 - FCAnalyzer is operating in SAN/OS on the supervisor card with output sent direct to telnet/console screen in **Local Mode**, or to ethereal running on a PC on the management IP Network in **Remote Mode**
- These 2 modes are used to capture data going to and from the Fabric Services

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

10

MDS FCAnalyzer (SAN/OS imbedded)

Cisco.com

- Output is displayed to the console in readable sniffer like format
- Is only used to monitor Fibre Channel Traffic to and from supervisor on the MDS9000
 - Traffic like Fabric Login's, FSPF routing, Switch to switch control traffic.
- Output can go direct to your console screen or to a workstation running a color Ethereal program

Note: Span is used for FC port to FC port monitoring and is also used to span out iSCSI ports and FCIP Ports

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

11

FCAnalyzer Detail Trace

Cisco.com

- Command entered at the MDS CLI (while in config mode) to start the FCAnalyzer trace; here we show how to filter for VSAN 1 traffic only
 - MDS-9509(config)# FCAnalyzer local display-filter (mdshdr.vsan==0x01)
- Other MDS built in filters for the imbedded ethereal program are found in the MDS9000 configuration guide
- FCAnalyzer is turned off with "control-C" after output has been displayed to your screen

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

12

Cisco.com

```
Fox-9216(config)# fcanalyzer local display-filter mdshdr.vsan==01
```

The remaining info is Fiber Channel

[illegible]

13

Cisco.com

```
Fox-9216(config)# fcanalyzer local brief display-filter mdshdr.vsan==01
```

- Line by line trace
- Cannot be expanded, use non brief to capture details

ff.75-75.02.03	0x5c0b0fff FCL FS_PLOGI
75.02.03-ff.75	0x5c0b0fff FCL FS_LN RJT (PLOGI)
ff.75-ff.65	0x5d040fff SW_ILS SW_RSCN
ff.75-ff.68	0x5e0bfff SW_ILS SW_RSCN
ff.65-ff.75	0x5d040b8f FCL Link_CD, ACKI
ff.68-ff.75	0x5e0b8bf FCL Link_CD, ACKI
ff.75-ff.65	0x5d040b8f FCL Link_CD, ACKI (SW_RSCN)
ff.75-ff.68	0x5e0b8bf FCL Link_CD, ACKI
ff.65-ff.75	0x5d040b8f FCL Link_CD, ACKI (SW_RSCN)
ff.75-ff.65	0x5d040b8f FCL Link_CD, ACKI
ff.65-ff.75	0x281c0bfff DNS_GE_ID
ff.75-ff.65	0x281c0b8f FCL Link_CD, ACKI
ff.75-ff.68	0x281c0b8f FCL Link_CD, ACKI (GE_ID)
ff.68-ff.75	0xbdb70bfff DNS_GE_ID
ff.75-ff.68	0xbdb70b60 FCL Link_CD, ACKI
ff.75-ff.68	0xbdb70b60 DNS_ACC (GE_ID)
ff.65-ff.75	0x281c0b5f FCL Link_CD, ACKI
ff.68-ff.75	0xbdb70b60 FCL Link_CD, ACKI
75.02.03-68.01.00	0x22080bfff FCP SCSI_Inquiry
75.02.03-68.01.00	0x22080bfff FCP SCSI_Inquiry
75.02.03-68.01.00	0x22080bfff FCP SCSI_Inquiry
75.02.03-68.01.00	0x22080bfff FCP SCSI_Inquiry
75.02.03-68.01.00	0x22080bfff FCP SCSI_Inquiry
75.02.03-68.01.00	0x22080bfff FCP SCSI_Inquiry
68.01.00-75.02.03	0x220c0bfff FCP_FCP_DATA
68.01.00-75.02.03	0x220c0bfff FCP_FCP_RSP_Good
75.02.03-68.01.00	0x220d0bfff FCP SCSI_Inquiry
68.01.00-75.02.03	0x220d0bfff FCP_FCP_DATA
68.01.00-75.02.03	0x220d0bfff FCP_FCP_RSP_Good
75.02.03-68.01.00	0x220f0bfff FCP SCSI_Inquiry
68.01.00-75.02.03	0x220f0bfff FCP_FCP_RSP_Good

OPT-3052
9909 06 2004 X

© 2004 Cisco Systems, Inc. All rights reserved.

MDS FCAnalyzer Remote with Ethereal

Cisco.com

- For capture of internal MDS9000 switch traffic, FC traffic to and from supervisor not FC port to FC port traffic
- Make sure MDS can ping workstation

MDS configuration

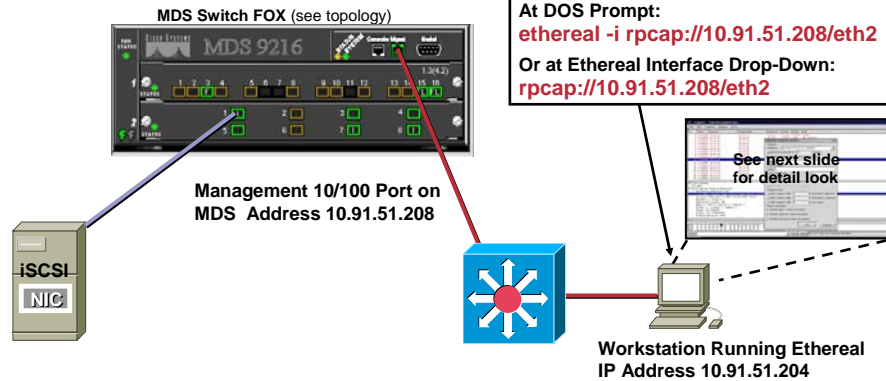
```
IPS-TEST(config)# fcanalyzer remote
10.91.51.204
IPS-TEST# sh fcanalyzer
ActiveClient = 10.91.51.204, DEFAULT
IPS-TEST#clear fcanalyzer turns off analyzer
```

At DOS Prompt:

ethereal -i rpcap://10.91.51.208/eth2

Or at Ethereal Interface Drop-Down:

rpcap://10.91.51.208/eth2



OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

15

Ethereal

File Edit View Capture Analyze Help

No.	Time	Source	Destination	Protocol	Info
66	18.237150	ff.ff.fe	65.05.00	FC ELS	ACC (FLOGI)
67	18.237175	ff.ff.fe	65.05.00	FC ELS	ELS
68	18.237455	ff.fc.65	65.05.00	FC ELS	PLOGI
69	18.237472	ff.fc.65	65.05.00	FC ELS	PLOGI
88	18.252660	ff.ff.fc	65.05.00	FC ELS	ACC (PLOGI)
89	18.252674	ff.ff.fc	65.05.00	FC ELS	ACC (PLOGI)
93	18.253572	ff.ff.fc	65.05.00	dns	ACC (GID_FT)
94	18.253589	ff.ff.fc	65.05.00	dns	ACC (GID_FT)
97	18.254857	ff.ff.fc	65.05.00	FC ELS	ACC (LOGO)
98	18.254871	ff.ff.fc	65.05.00	FC ELS	ACC (LOGO)
101	18.255337	ff.ff.fd	65.05.00	FC ELS	ACC (SCR)

rpcap builds tunnel between this PC and port eth2 on MDS

Filters can be applied before or after capture (only looking at FCID 65.05.00)

Ethereal: Capture Options

Capture

Interface: rpcap://10.91.51.208/eth2

Link-layer header type: [] bytes

Limit each packet to: [68] bytes

Capture packets in promiscuous mode

Filter: []

Capture file(s)

File: []

Use ring buffer Number of files: [0]

Rotate capture file every: [1] second(s)

Display options

Update list of packets in real time

Automatic scrolling in live capture

Capture limits

Stop capture after: [1] packet(s) captured

Stop capture after: [1] kilobyte(s) captured

Stop capture after: [1] second(s)

Name resolution

Enable MAC name resolution

Enable network name resolution

Enable transport name resolution

OK Cancel

Use of Span Feature

Cisco.com

- Used for FC Port to FC Port analyzing
- Same type of tool as used on Cisco Catalyst products
- Ingress and egress ports are sent to a FC port set up as a Span Destination (SD port type)
- No limits to where the ports are located on the MDS switch
- Used to output to third party test equipment or to Cisco Port Analyzer Adapter (PAA)

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

17

Setting up SPAN 1-2-3

Cisco.com

Span drop down selected

One FC interface configured as Span Destination (SD)
This example uses FC 1/11

Destination SD port selected

Source of data selected, in this case iSCSI interface 3/1

Save and close, Span has started

The screenshot shows the Span configuration interface with a table of sessions and sources. The table has columns for Session, Vsan List, and Or Interface (Direction) List. The first row shows Session 1, Vsan List 1, and Or Interface (Direction) List iSCSI3/1 (Rx), iSCSI3/1 (Tx). The table is titled "10.91.51.212 - SPAN".

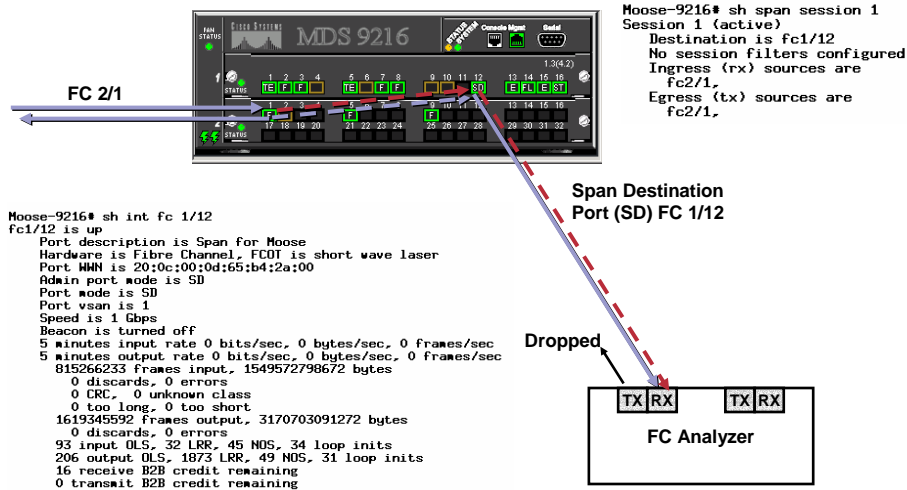
OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

18

Span Configuration with Single FC Port Analyzer Tool

Cisco.com



OPT-3052
9909_06_2004_X

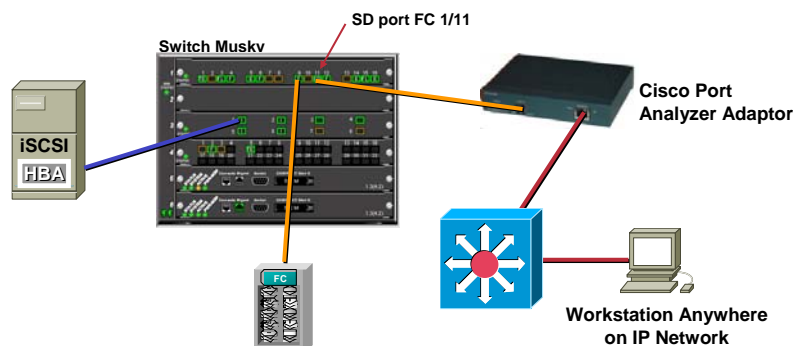
© 2004 Cisco Systems, Inc. All rights reserved.

19

Port Analyzer Connection Using Port Analyzer with Ethereal

Cisco.com

- For capture of device port to device port Fibre Channel traffic to a IP attached workstation running Ethereal
 - Configure Port Analyzer Adapter (dip switch settings, mostly plug and play)
 - Configure SPAN port on MDS9000
 - Set Ethereal to capture on local Ethernet interface



OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

20

MDS9000 Fabric Manager and Device Manager

Cisco.com

- **Fabric Manager & Device Manager are full of great serviceability tools for supporting the SAN Network**
- **Best of both worlds, using JAVA based manager along with CLI to troubleshoot**
- **Ability to run the Graphic manager on any laptop from remote office location anywhere on IP network or from with the Data Center.**
- **We will go over many of these graphic outputs in detail throughout this session**

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

21

OVERVIEW OF MDS AND SN5428-2 IP STORAGE

ISCSI



OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

22

Section Agenda

Cisco.com

- Briefing on iSCSI Protocol
- Understanding the iSCSI Initiator
- Understanding the iSCSI Session
- iSCSI Targets
- Debugging iSCSI Protocol on MDS
- Debugging iSCSI Protocol on SN5428
- Troubleshooting iSCSI Hosts

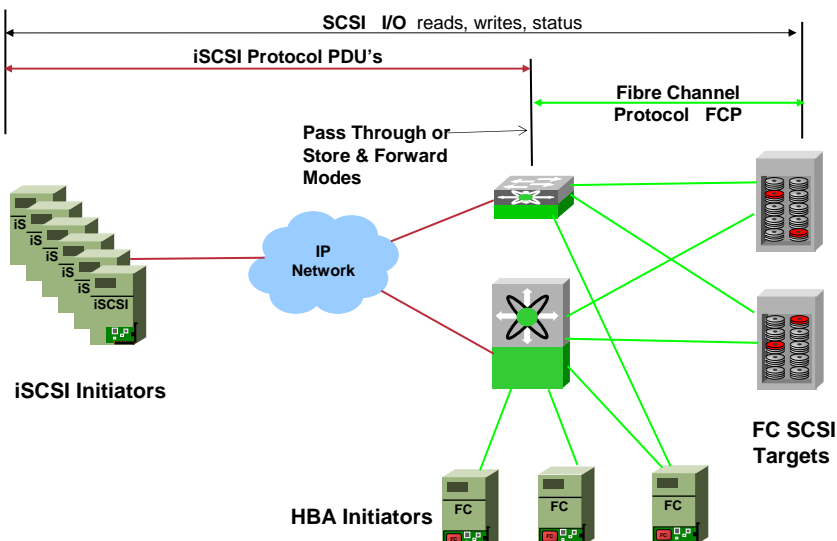
OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

23

Understanding what we are troubleshooting SCSI Data End to End

Cisco.com




OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

24

Understanding what we are troubleshooting iSCSI Initiators iSCSI Targets

Cisco.com



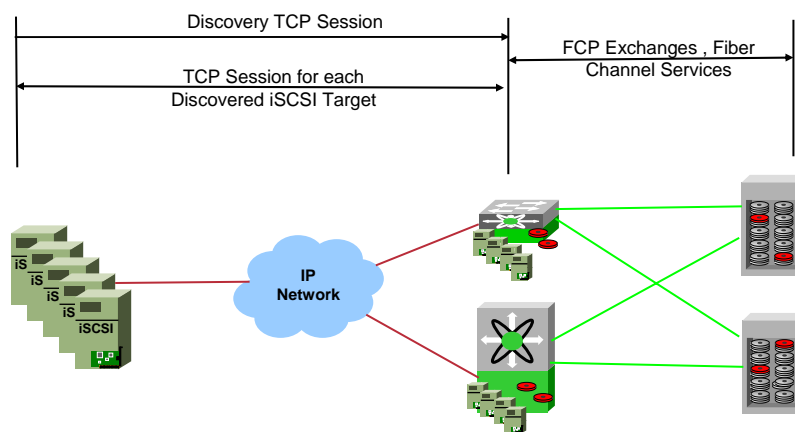
OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

25

Understanding what we are troubleshooting The iSCSI Connection

Cisco.com



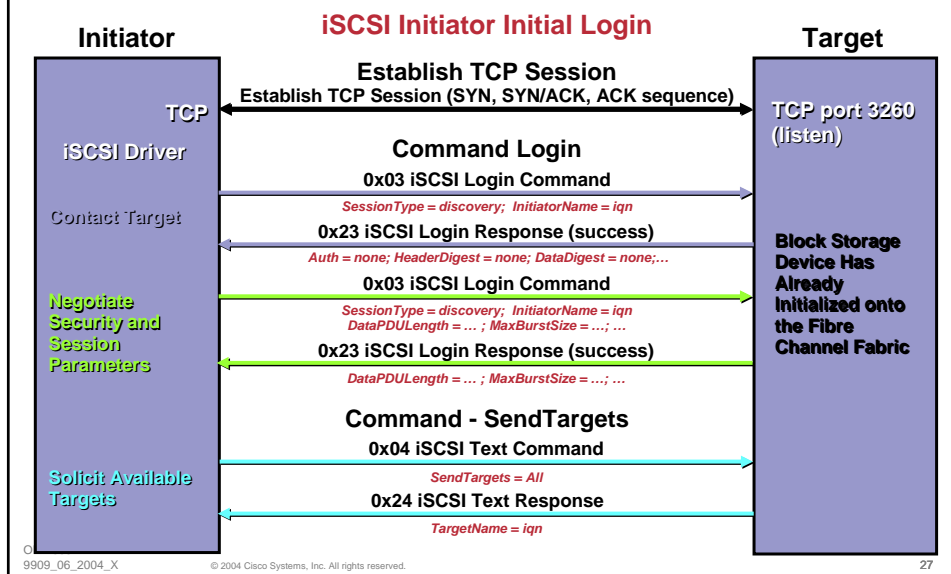
OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

26

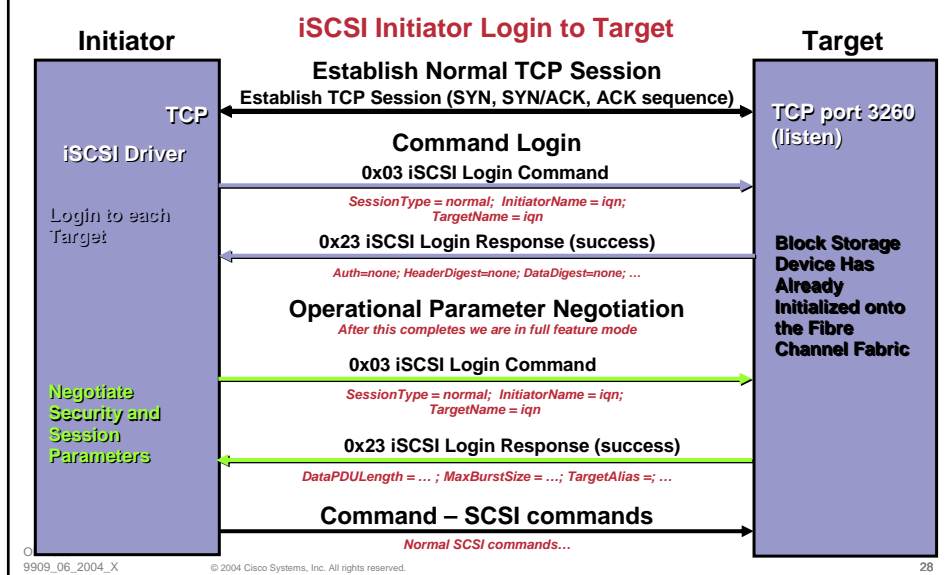
Example: iSCSI Initial Login Sequence for Target Discovery

Cisco.com



Example: iSCSI Login Sequence to Discovered Targets for Normal Operation

Cisco.com



iSCSI on The MDS A Little Visual Help

Fibre Channel Switch to real targets connected in the SAN

Cisco.com

iSCSI interface is here.

Interface needs to be no shut to work and iSCSI process must be enabled

interface iscsi 2/1- 2/8

GiGE IP address goes here. Could be sub-interfaces if using VLANs

GiGE 2/1 - 2/8

Each interface has own IP route tables

GiGE Link

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

29

iSCSI Session Verification

Cisco.com

Initiator 10.1.11.19

Initiator name iqn.1987-05.com.cisco:04.racoon.

Session #1

Discovery session, ISID 00023d000001, Status active

There will be a Session then created for each and every Target that the iSCSI Initiator is configured for. In this case Session #2 called duck-pxe

Session #2

Target duck-pxe

VSAN 602, ISID 00023d000004, Status active, no reservation

Test the iSCSI connection at install

Just point the iSCSI driver at the network portal and you should see a discovery session. This will verify proper driver operation and access into VSAN 1 (default VSAN) This is true for Cisco Driver.

Microsoft Driver does do a Discovery Session but drops it quickly, debugs shall help identify if MS Discovery login occurred

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

30

MDS iSCSI Initiators (A view from Device Manager)

Cisco.com

10.91.51.212 - iSCSI

Initiators

Targets

Session Initiators

iSNS Profiles

Globals

Name or IP Address	VsanMembership	Node Address				Port Address		AuthUser
		Dynamic	Persistent	SystemAssigned	WWN	Persistent	WWN	
10.1.11.19	4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	21:06:00:0d:ec:00:ea:42	true	systemAssigned	
10.1.11.23	4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	21:07:00:0d:ec:00:ea:42	true	Cisco 21:05:00:0d:ec:00:ea:42	
10.1.11.15	4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	21:10:00:0d:ec:00:ea:42	true	Cisco 21:0f:00:0d:ec:00:ea:42	
10.1.100.31	1	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	21:0a:00:0d:ec:00:ea:42	true	Cisco 21:09:00:0d:ec:00:ea:42	
10.1.200.31	4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	21:04:00:0d:ec:00:ea:42	true	Cisco 21:02:00:0d:ec:00:ea:42	
10.1.11.26	4	<input type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	21:0d:00:0d:ec:00:ea:42	true	Cisco 21:0c:00:0d:ec:00:ea:42	

Edit Port WWN...

Create...

Delete

Apply

Refresh

Help

Close

Data retrieved at 17:00:35

IP Address of iSCSI Hosts

pWWN and WWN's can also be assigned by switch or administered

VSAN Membership, can be in multiple VSAN's

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

31

MDS CLI Displays: Show iSCSI Interface

Cisco.com

FOX-9216# show interface iscsi 2/1
iscsi2/1 is up
Hardware is GigabitEthernet
Port WWN is 20:41:00:0b:be:77:72:40
Admin port mode is ISCSI
Port mode is ISCSI
Port vsan is 1j
Speed is 1 Gbps
iSCSI initiator is identified by name
Number of iSCSI session: 4, Number of TCP connection: 4
Configured TCP parameters
Local Port is 3260
PMTU discover is enabled, reset timeout is 3600 sec
Keepalive-timeout is 60 sec
Minimum-retransmit-time is 300 ms
Max-retransmissions 4
Sack is disabled
QOS code point is 0
Forwarding mode: pass-thru
TMF Queueing Mode : disabled
Proxy Initiator Mode : disabled
5 minutes input rate 3752 bits/sec, 469 bytes/sec, 0 frames/sec
5 minutes output rate 192 bits/sec, 24 bytes/sec, 0 frames/sec
iSCSI statistics
Input 596940 packets, 323436852 bytes
Command 259326 pbus, Data-out 286119 pbus, 288606720 bytes
Output 3820532 packets, 4954553084 bytes
Response 259326 pbus (with sense 5), R2T 41140 pbus
Data-in 3468571 pbus, 4771166340 bytes

Identified via name vs. IP address

4 TCP sessions

- 1) Discovery
- 2) Target 1
- 3) Target 2
- 4) Target 3

Pass through Mode


OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

32

MDS CLI Displays: Show Interface GigE

Cisco.com



- Use to see if interface is up
- Used to verify IP address and Mask
- Lookup MAC address to Verify ARP entry in Layer 2 Switch
- Input rates and errors

```
IPS_test# show interface gigabitethernet 2/1
GigabitEthernet2/1 is up
Hardware is GigabitEthernet, address is 0005.3000.2e12
Internet address is 10.1.30.200/24
MTU 1500 bytes, BW 1000000 kbit
Port mode is IPS
Speed is 1 Gbps
Beacon is turned off
5 minutes input rate 45368688 bits/sec, 5671086 bytes/sec, 5341 frames/sec
5 minutes output rate 37979592 bits/sec, 4747449 bytes/sec, 4001 frames/sec
494767128 packets input, 525488052862 bytes
  493 multicast frames, 0 compressed
  0 input errors, 0 frame, 0 overrun 0 fifo
370486468 packets output, 439364990332 bytes, 0 underruns
  0 output errors, 0 collisions, 0 fifo
  0 carrier errors
```

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

33

MDS CLI Displays: Specific iSCSI show commands

Cisco.com

- **global** Show iscsi global config information
- **initiator** Show iscsi initiator information
- **session** Show iscsi session information
- **stats** Show iscsi statistics
- **virtual-target** Show iscsi virtual-targets

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

34

MDS CLI Displays: iSCSI initiator

Cisco.com

IPS_test# show iscsi initiator

detail	Show detailed iscsi node information (Gives you iscsi side and FC side in one output)
fc-session detail	Show detailed FC session details (Information on the Fibre Channel negotiated parameters)
configured	Show iscsi initiator node information (General information much like <sh iscsi session>)
iscsi-session detail	Show detailed iscsi session details (Dig here for TCP details)
summary	Show iscsi target node information (this would be useful if there were remote targets but this feature is in a future release)

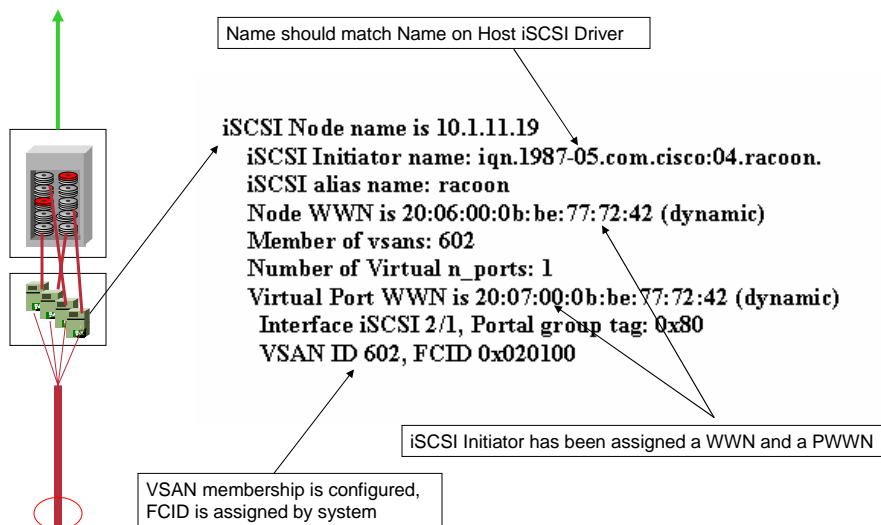
OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

35

MDS CLI Displays: Details on iSCSI initiator in the iSCSI

Cisco.com



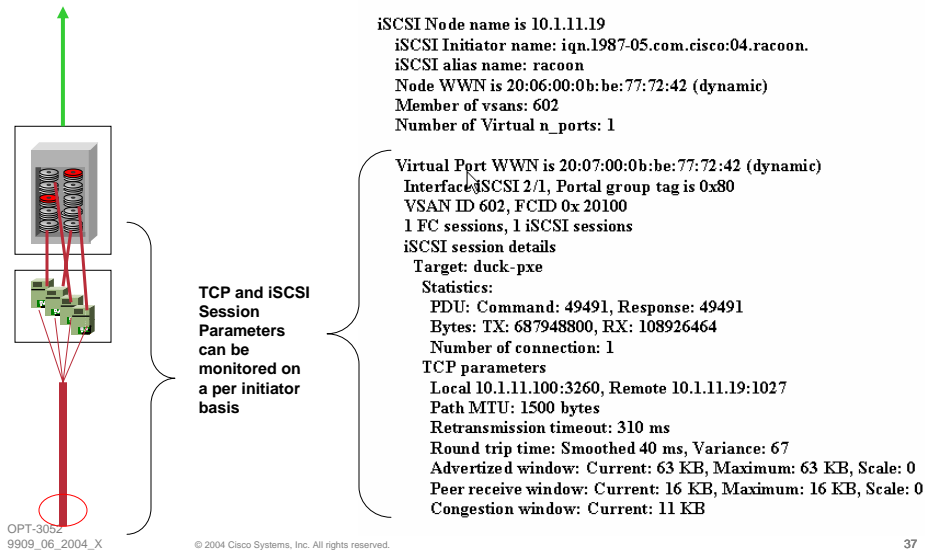
OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

36

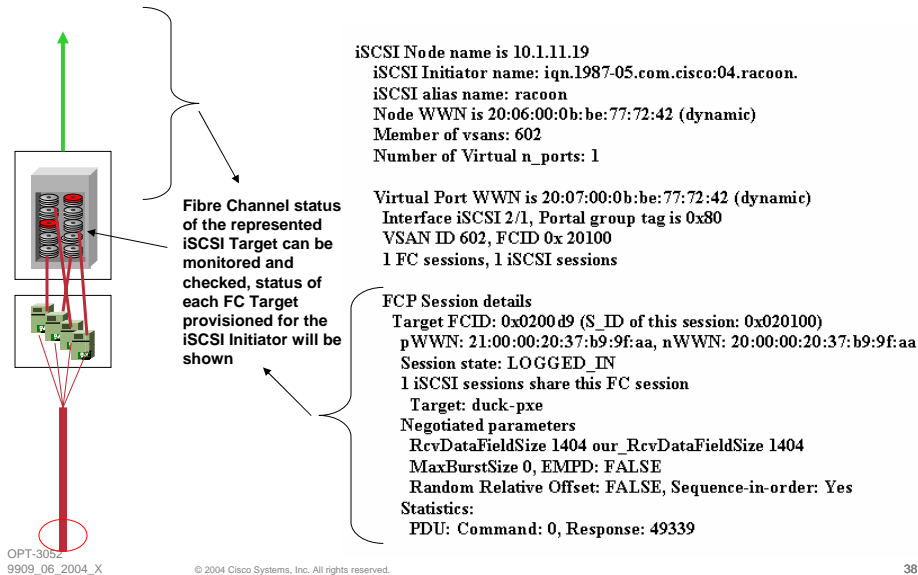
MDS CLI Displays: show iscsi initiator iscsi-session

Cisco.com



MDS CLI Displays: show iscsi initiator fcp-session detail

Cisco.com



MDS iSCSI Targets

Cisco.com

Understanding Virtual-Targets

- A created Virtual-Target is the named Target that is accessed by the iSCSI initiator at Login i.e. Sendtargets response or sent to iSNS server
- Seeing Targets on iSCSI status screen of iSCSI initiator tells you that the initiator has logged in to the configured iSCSI virtual-target on the IPS
- PLOGI to real target pWWN on the FC SAN with good status return is required for Disk Manager of Server OS to see Target LUN's

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

39

Viewing MDS iSCSI Hosts in the Fabric (CLI)

Cisco.com

Show Fibre Channel Name Server

Musky-9506# sh fcns data vsan 4

VSAN 4:

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEAT/RE
0x2d0102	N	21:0c:00:0d:ec:00:ea:42	(Cisco)	scsi-fcp:init isc..w
0x2d0104	N	21:08:00:0d:ec:00:ea:42	(Cisco)	scsi-fcp:init isc..w
0x2d0105	N	21:01:00:0d:ec:00:ea:42	(Cisco)	scsi-fcp:init isc..w

iSCSI Initiators

iSCSI interfaces,
VSAN they are in
and WWN's on
MDS

Show Fabric Login Database

Musky-9506# sh flogi data vsan 4

INTERFACE	VSAN	FCID	PORT NAME	NODE NAME
iscsi3/1	4	0x2d0102	21:0c:00:0d:ec:00:ea:42	21:0d:00:0d:ec:00:ea:42
iscsi3/1	4	0x2d0105	21:01:00:0d:ec:00:ea:42	21:06:00:0d:ec:00:ea:42
iscsi3/2	4	0x2d0104	21:08:00:0d:ec:00:ea:42	21:06:00:0d:ec:00:ea:42

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

40

MDS iSCSI Targets (A view from Device Manager)

Cisco.com

iSCSI Name	Dynamic	Port WWN's		LUN Map (Hex)		Initiator(s) Access List	Advertised Interfaces	Trespas Mode	RevertTo PrimaryPort
		Primary	Secondary	FC Primary	FC Secondary				
ipn racoonboot0	false	22:00:00:20:37:c5:23:56				<input checked="" type="checkbox"/>	gigE3/1	<input type="checkbox"/>	<input type="checkbox"/>
ipn redstonestorageforacorn	false	22:00:00:20:37:c5:26:6d				<input type="checkbox"/>	10.1.11.1902 gigE3/1	<input type="checkbox"/>	<input type="checkbox"/>
ipn storageforperch	false	20:05:00:a0:b0:0c:64:51				<input type="checkbox"/>	10.1.100.31/02 gigE3/5	<input type="checkbox"/>	<input type="checkbox"/>
ipn morestorageforperch	false	22:00:00:20:37:c5:2e:2e				<input type="checkbox"/>	10.1.200.31/02 gigE3/6	<input type="checkbox"/>	<input type="checkbox"/>
ipn storageforloopout	false	22:00:00:04:c7:75:21:a0				<input type="checkbox"/>	10.1.11.26/02 all	<input type="checkbox"/>	<input type="checkbox"/>
ipn storageforloopin	false	22:00:00:04:c7:75:13:30				<input type="checkbox"/>	10.1.11.23/02 all	<input type="checkbox"/>	<input type="checkbox"/>

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

41

MDS iSCSI TCP Values (A view from Device Manager)

Cisco.com

Interface	Port	SACK	KeepAlive (sec)	ReTransmission		SendBuffer (KB)	Bandwidth (KB)		QoS	Enable	PMTU ResetTimeout (sec)
				MinTimeout (ms)	Max		Min	Max			
iscsi3/1	3260	<input type="checkbox"/>	60	300	4	0	15000	1000000	0	<input checked="" type="checkbox"/>	3600
iscsi3/2	3260	<input checked="" type="checkbox"/>	60	300	4	0	15000	1000000	0	<input checked="" type="checkbox"/>	3600
iscsi3/5	3260	<input checked="" type="checkbox"/>	60	300	4	0	15000	1000000	0	<input checked="" type="checkbox"/>	3600
iscsi3/6	3260	<input checked="" type="checkbox"/>	60	300	4	0	15000	1000000	0	<input checked="" type="checkbox"/>	3600
iscsi3/3	3260	<input checked="" type="checkbox"/>	60	300	4	4096	70000	1000000	0	<input checked="" type="checkbox"/>	3600
iscsi3/4	3260	<input checked="" type="checkbox"/>	60	300	4	4096	70000	1000000	0	<input checked="" type="checkbox"/>	3600
iscsi3/7	3260	<input checked="" type="checkbox"/>	60	300	4	4096	70000	1000000	0	<input checked="" type="checkbox"/>	3600
iscsi3/8	3260	<input checked="" type="checkbox"/>	60	300	4	4096	70000	1000000	0	<input checked="" type="checkbox"/>	3600

- Each iSCSI interface has adjustable TCP parameters that apply to all iSCSI sessions using this virtual interface
- The configuration may require tuning based on network design requirements, consult Cisco design guides and configuration guides for proper use
- These configurations can be done via CLI, and by Device Manager.

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

42

Connection Troubleshooting Interface IP Stats

Cisco.com

```
Musky-9506# sh ips stats ip interface gigabitethernet 3/1 detail
Internet Protocol Statistics for port GigabitEthernet 3/1
 0 total received, 0 good, 0 error
 0 reassembly required, 0 reassembled ok, 0 dropped after timeout
 0 Bad checksum, 0 bad header length, 0 source multicast address
 0 packet len < IP header, 0 packets too short, 0 packets too long
 0 fragments dropped for lack of memory, 0 bad fragments
 0 Duplicate fragments dropped, 0 Fragment timed out
 0 cannot forward, 0 bad protocol, 0 bad option, 0 bad version
 0 packets sent, 0 outgoing dropped, 0 dropped no route
 0 fragments created, 0 cannot fragment
```

- Monitor IP Stats for checksum errors and drops
- Clear old counters and capture file or screen

```
clear ips stats ip interface gigabitethernet 3/1
```

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

43

Connection Troubleshooting Interface TCP Stats Page 1 of 2

Cisco.com

```
Musky-9506# sh ips stats tcp interface gigabitethernet 3/1 detail
TCP Statistics for port GigabitEthernet 3/1
```

TCP send stats

```
234 segments, 4704 bytes
98 data, 136 ack only packets
0 control (SYN/FIN/RST), 0 probes, 0 window updates
0 segments retransmitted, 0 bytes
0 retransmitted while on ethernet send queue, 0 packets split
16 delayed acks sent
```

TCP receive stats

```
187 segments, 162 data packets in sequence, 143968 bytes in sequence
0 predicted ack, 88 predicted data
0 bad checksum, 0 multi/broadcast, 0 bad offset
0 no memory drops, 0 short segments
0 duplicate bytes, 0 duplicate packets
0 partial duplicate bytes, 0 partial duplicate packets
0 out-of-order bytes, 0 out-of-order packets
0 packet after window, 0 bytes after window
0 packets after close
93 acks, 4704 ack bytes, 0 ack toomuch, 6 duplicate acks
0 ack packets left of snd_una, 0 non-4 byte aligned packets
0 window updates, 0 window probe
0 pcb hash miss, 0 no port, 0 bad SYN, 0 paws drops
```

Retransmissions & checksum issues for send side can be monitored with this display

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

44

Connection Troubleshooting Interface TCP Stats 2 of 2

Cisco.com

TCP Connection Stats

0 attempts, 0 accepts, 0 established
0 closed, 0 drops, 0 conn drops
0 drop in retransmit timeout, 0 drop in keepalive timeout
0 drop in persist drops, 0 connections drained

TCP Miscellaneous Stats

93 segments timed, 93 rtt updated
0 retransmit timeout, 0 persist timeout
6 keepalive timeout, 6 keepalive probes

TCP SACK Stats

0 recovery episodes, 0 data packets, 0 data bytes
0 data packets retransmitted, 0 data bytes retransmitted
0 connections closed, 0 retransmit timeouts

TCP SYN Cache Stats

0 entries, 0 connections completed, 0 entries timed out
0 dropped due to overflow, 0 dropped due to RST
0 dropped due to ICMP unreachable, 0 dropped due to bucket overflow
0 abort due to no memory, 0 duplicate SYN, 0 no-route SYN drop
0 hash collisions, 0 retransmitted

TCP Active Connections

Local Address	Remote Address	State	Send-Q	Recv-Q
10.1.11.22:3260	10.1.11.26:1375	ESTABLISH	0	0
10.1.11.22:3260	10.1.11.19:1026	ESTABLISH	0	0
10.1.11.22:3260	10.1.11.19:1028	ESTABLISH	0	0
10.1.11.22:3260	10.1.11.19:1029	ESTABLISH	0	0
0.0.0.0:3260	0.0.0.0:0	LISTEN	0	0

Monitor for TCP
Connection issues

Selective ACK's issues

TCP connection protocol
health

clear ips
arp Clear arp table
stats Various ips stats

Active iSCSI TCP
connections, 1 for each
target will be seen

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

45

Level 1 ASIC Troubleshooting GiGE MAC layer stats

Cisco.com

Musky-9506# sh ips stats mac interface gigabitethernet 3/1

Ethernet MAC statistics for port GigabitEthernet3/1

Hardware Transmit Counters

108583222 frame 122456748364 bytes
0 collisions, 0 late collisions, 0 excess collisions
0 bad frames, 0 FCS error, 0 abort, 0 runt, 0 oversize

Hardware Receive Counters

58420066596 bytes, 87353942 frames, 1245637 multicasts,
339288 broadcasts

0 bad, 0 runt, 0 CRC error, 0 length error
0 code error, 0 align error, 0 oversize error

Software Counters

87351646 received frames, 108583222 transmit frames
0 frames soft queued, 0 current queue, 0 max queue
0 dropped, 0 low memory
11855471540430 idle count

SPF or Cable issues,
should not see
collisions on GiGE

- Use to verify health of lower level MAC hardware

clear ips stats mac interface gigabitethernet 3/1

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

46

iSCSI Protocol Login Stats: Part 1 of 2

Cisco.com

Musky-9506# **show iscsi stats detail iscsi3/1**

```
5 minutes input rate 3824 bits/sec, 478 bytes/sec, 0 frames/sec
5 minutes output rate 120 bits/sec, 15 bytes/sec, 0 frames/sec
iSCSI statistics
102878 packets input, 91499160 bytes
  Command 15782 pdus, Data-out 85522 pdus, 86560256 bytes, 0 fragments
  output 58205 packets, 42375944 bytes
  Response 15782 pdus (with sense 2), R2T 11747 pdus
  Data-in 29102 pdus, 39581432 bytes
iSCSI Forward:
  Command: 15782 PDUs (Rcvd: 15782)
  Data-Out (Write): 85522 PDUs (Rcvd 85522), 0 fragments, 86560256 bytes
FCP Forward:
  Xfer_rdy: 11747 (Rcvd: 11747)
  Data-In: 29102 (Rcvd: 29104), 39581432 bytes
  Response: 15782 (Rcvd: 15784), with sense 2
  TMF Resp: 0

iSCSI Stats:
  Login: attempt: 2, succeed: 2, fail: 0, authn fail: 0
  Rcvd: NOP-Out: 1568, Sent: NOP-In: 1568
  NOP-In: 0, Sent: NOP-Out: 0
  TMF-REQ: 0, Sent: TMF-RESP: 0
  Text-REQ: 0, Sent: Text-RESP: 0
  SNACK: 0
  Unrecognized Opcode: 0, Bad header digest: 0
  Command in window but not next: 0, exceed wait queue limit: 0
  Received PDU in wrong phase: 0
  SCSI Busy responses: 0
```

iSCSI Protocol counters, see RFC doc for meanings

Counters on the Fibre Channel side

Login stats and iSCSI ping counters

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

47

iSCSI Protocol Login Stats: Part 2 of 2

Cisco.com

FCP Stats:

```
Total: Sent: 101322
Received: 56653 (Error: 0, Unknown: 0)
Sent: PLOGI: 2, Rcvd: PLOGI_ACC: 2, PLOGI_RJT: 0
PRLI: 2, Rcvd: PRLI_ACC: 2, PRLI_RJT: 0, Error: 0, From initiator: 0
LOGO: 2, Rcvd: LOGO_ACC: 0, LOGO_RJT: 0
PRLO: 2, Rcvd: PRLO_ACC: 0, PRLO_RJT: 0
ABTS: 0, Rcvd: ABTS_ACC: 0
TMF REQ: 0
Self orig command: 2, Rcvd: data: 2, resp: 2
Rcvd: PLOGI: 4, Sent: PLOGI_ACC: 3, PLOGI_RJT: 1
LOGO: 3, Sent: LOGO_ACC: 3, LOGO_RJT: 0
PRLI: 1, Sent: PRLI_ACC: 1, PRLI_RJT: 0
PRLO: 0, Sent: PRLO_ACC: 0, PRLO_RJT: 0
ABTS: 0
```

iSCSI Drop:

```
Command: Target down 0, Task in progress 0, LUN map fail 0
CmdSeqNo not in window 0, No Exchange ID 0, Reject 0
No task: 0
```

Data-Out: 0, Data CRC Error: 0

TMF-Req: 0, No task: 0

FCP Drop:

Xfer_rdy: 0, Data-In: 0, Response: 0

Buffer Stats:

```
Buffer less than header size: 105, Partial: 52962, Split: 75707
Pullup give new buf: 105, Out of contiguous buf: 0, Unaligned m_data: 0
```

iSCSI login and logout counters

CRC errors on network side and data drops on Fibre Channel side

Command to clear these counters
clear counters int iscsi 3/1

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

48

iSCSI Debugging

Cisco.com

- Requirement to attach to IPS module to run debugs.
Can use ? for help

```
IPS-TEST# show module
Mod  Ports  Module-Type          Model          Status
---  -
1     16     1/2 Gbps FC/Supervisor DS-X9216-K9-SUP active *
2     8      IP Storage Module      ok
```

Indication of Attachment
remember keys to exit

Locate IP Blade
slot and show
health of Blade

```
IPS-TEST# attach module 2
Attaching to module 2 ...
To exit type 'exit', to abort type '$.'
module-2#
```

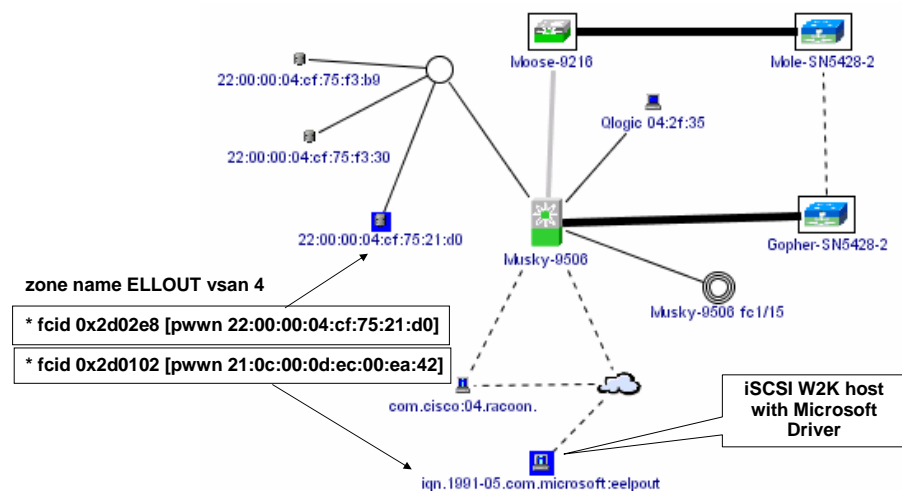
OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

49

Troubleshoot login of iSCSI Host Ellpout

Cisco.com



OPT-3052
9909_06_2004_X



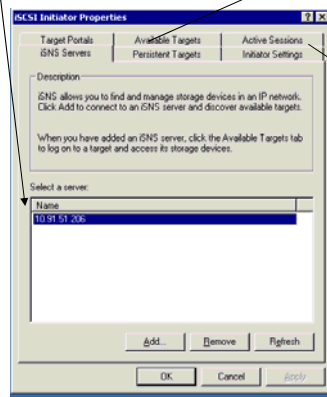
© 2004 Cisco Systems, Inc. All rights reserved.

50


MS iSCSI Driver 101

Cisco.com

iSNS server database is queried for targets available on the network



Driver creates initiator name at time of driver load



OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

51

Re-login: Part 1 of 4 When iSCSI host had acquired target before

Cisco.com

Musky-9506# att mod 3
Attaching to module 3 ...
To exit type 'exit', to abort type '\$.'

Best place to troubleshoot iSCSI host connection is on the IPS blade

module-3# **debug ips iscsi login port 1**

53 Musky-9506 %IPS-5-ISCSI_CONN_DOWN: GigabitEthernet3/1: iSCSI session down from 10.1.11.26 to target iqn.storageforellpout, reason: client send iscsi logout

Console message-----Logout because host sent logout PDU

port1: 1: iscsi_connect_callback: iscsi ds len: 1392, tcp mss: 1460
port1: 2: Rcv InitName: iqn.1991-05.com.microsoft:EELPOUT

Host name should match config on driver

port1: 3: SessionType: Normal

iSCSI host doing Normal login (not discovery, it knows what it wants because it is either manually configured on the host or it received this target from iSNS server)

port1: 4: Rcv TargetName: iqn.storageforellpout

iSCSI host asked for Target iqn.storageforellpout

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

52

Re-login: Part 2 of 4 When iSCSI host had acquired target before

Cisco.com

```
port1: 5:  init NWWN 21:d:0:d:ec:0:ea:42
port1: 6:  Init PWWN 21:0c:00:0d:ec:00:ea:42
port1: 7:  Init NWWN 21:0d:00:0d:ec:00:ea:42
port1: 8:  Target PWWN 22:00:00:04:cf:75:21:d0
port1: 9:  Target NWWN 20:00:00:04:cf:75:21:d0
```

WWN's are re-assigned Dynamically or Statically based on configuration

```
port1: 10: Set our RcvDataFieldSize to TCP MSS 1404 (Pass-thru)
```

Host want's to connect to target with MTU of 1404

```
port1: 11: iscsi_send_login_resp: State: S4.0.1, CSG 0 NSG 1, TBIT 1, transit 0 got_MRPL 0,
sent_MRPL 0, sent_1st_LoginOperNeg 0
port1: 12:  Put TargetPortalGroupTag key: TargetPortalGroupTag=256, new_len 24, out_len 0
port1: 13: Find_common_key: num_proposed 1, num_acceptable 1
port1: 14:  init propose 0
port1: 15:  tgt allows 0
port1: 16: iscsi_send_login_resp auth common index: 0, method: 0
port1: 17:  Add key AuthMethod=None, new_len 15, out_len 25
```

iSCSI authentication in driver config is set to None

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

53

Re-login: Part 3 of 4 When iSCSI host had acquired target before

Cisco.com

```
port1: 18: S4.2: MaxRecvDataSegmentLength found (65536) or T-bit is 1, fc sess 0xd8c80500, fc
state 0, discovery session 0
P+ort1: 19: fcp_init_sess_fsm: got RcvDataFieldSize 1404 from FC target 0x2d02e8
port1: 20: S4.2.1: rcv PLOGI ACC, T-bit 1
```

PLOGI to Fibre Channel Target with MTU agreed to be 1404, good for pass-through mode

```
port1: 21: iscsi_send_login_resp: State: S4.2.1, CSG 1 NSG 3, TBIT 1, transit 1 got_MRPL 1,
sent_MRPL 0, sent_1st_LoginOperNeg 0
port1: 22: Put ImmediateData key: ImmediateData=No, new_len 16, out_len 0
port1: 23: Put MaxRecvDataSegmentLength key: MaxRecvDataSegmentLength=1404, new_len
29, out_len 17
```

MTU of 1404 set in iSCSI PDU negotiation

```
port1: 24: Find_common_key: num_proposed 1, num_acceptable 2
port1: 25:  init propose 0
port1: 26:  tgt allows 1
port1: 27:  tgt allows 0
port1: 28: iscsi_send_login_resp digest common index: 1, method: 0
port1: 29:  Add key HeaderDigest=None, new_len 17, out_len 47
```

No CRC checking on header for this 1 target

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

54

Re-login: Part 4 of 4 When iSCSI host had acquired target before

Cisco.com

```
port1: 30: Find_common_key: num_proposed 1, num_acceptable 2
port1: 31: init propose 0
port1: 32: tgt allows 1
port1: 33: tgt allows 0
port1: 34: iscsi_send_login_resp digest common index: 1, method: 0
port1: 35: Add key DataDigest=None, new_len 15, out_len 65
```

No Data CRC for this Target

```
port1: 36: Add key ErrorRecoveryLevel=0, new_len 20, out_len 81
```

Error Level can be 0, 1, 2 per RFC. Have only seen 0 so far

```
port1: 37: Add key InitialR2T=Yes, new_len 14, out_len 102
```

Target supports Ready to Transfer

```
port1: 38: Send_login_resp: our_MaxRecvDSLen 1404, Rcvd value 65536, FC RcvDataFieldSize 1404
MTU of 1404 is set from iSCSI host to FC storage
```

```
port1: 39: Add key MaxBurstLength=262144, new_len 21, out_len 117
port1: 40: Add key FirstBurstLength=65536, new_len 22, out_len 139
Musky-9506 %IPS-5-ISCISI_CONN_UP: GigabitEthernet3/1: iSCSI session up from 10.1.11.26 10 to
target: iqn.storageforellpout
```

Login all done and SCSI data ready to flow


OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

55

Troubleshoot Login of iSCSI Host Duck1

Cisco.com



OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

56

Cisco iSCSI Driver 101

Cisco.com

Target Portal Discovery IP address 10.1.11.100

Session attributes negotiated during login per iSCSI RFC 3720

IQN name of iSCSI Host

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

57

Cisco iSCSI Driver Discovery & Target Login Debug: Part 1 of 5

module-2# debug ips iscsi login port 1

module-2# Jun 23 10:53:55 2004 port1: 65: iscsi_connect_callback: iscsi ds len: 1392, tcp mss: 1426
port1: 1: Rcv InitName: iqn.1987-05.com.cisco:02.24a0e17903ba.duck1

iSCSI initiator logs in with TCP MSS sizes and name configured on iSCSI Host

port1: 2: Rcv InitAlias: DUCK1
port1: 68: SessionType: Discovery

This iSCSI driver has a Alias in it's driver and it wants to do a discovery session not a normal session

port1: 3: init NWWN 20:9:0:b:be:77:72:42
port1: 4: Init PWWN 20:08:00:0b:be:77:72:42
port1: 5: Init NWWN 20:09:00:0b:be:77:72:42
port1: 6: Target PWWN 00:00:00:00:00:00:00:00
port1: 7: Target NWWN 00:00:00:00:00:00:00:00

iSCSI initiator gets WWN's assigned my MDS

port1: 8: Set our RcvDataFieldSize to TCP MSS 1368 (Pass-thru)

TCP MSS set, again this will be negotiated with PLOGI to target

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

58

Cisco iSCSI Driver Discovery & Target Login Debug: Part 2 of 5

Cisco.com

```
port1: 9: iscsi_send_login_resp: State: S4.0.1, CSG 0 NSG 1, TBIT 1, transit 0 got_MRPL 0,
sent_MRPL 0, sent_1st_LoginOperNeg 0
port1: 10: Put TargetPortalGroupTag key: TargetPortalGroupTag=128, new_len 24, out_len 0
port1: 11: Find_common_key: num_proposed 1, num_acceptable 1
port1: 12: init propose 0
port1: 13: tgt allows 0
port1: 14: iscsi_send_login_resp auth common index: 0, method: 0
port1: 15: Add key AuthMethod=None, new_len 15, out_len 25
```

Everything above is Discovery session parameter exchanges

```
port1: 16: S4.2: MaxRecvDataSegmentLength found (1370) or T-bit is 1, fcscsess 0xd8c80000, fc
state 0, discovery session 1
port1: 17: S4.2: T-bit 1, send login response
port1: 18: iscsi_send_login_resp: State: S4.2, CSG 1 NSG 3, TBIT 1, transit 1 got_MRPL 0,
sent_MRPL 0, sent_1st_LoginOperNeg 0
port1: 19: Put MaxRecvDataSegmentLength key: MaxRecvDataSegmentLength=1392, new_len
29, out_len 0
port1: 20: Add key X-com.cisco.sendAsyncText=Yes, new_len 29, out_len 30
```

Sendtargets for discovery session sent in key exchange

```
port1: 21: iscsi_connect_callback: iscsi ds len: 1392, tcp mss: 1426
```

Another TCP session established for a new session

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

59

Cisco iSCSI Driver Discovery & Target Login Debug: Part 3 of 5

Cisco.com

```
port1: 22: Rcv InitName: iqn.1987-05.com.cisco:02.24a0e17903ba.duck1
port1: 23: Rcv InitAlias: DUCK1
port1: 24: SessionType: Normal
```

New TCP Session from initiator with Alias DUCK1 and that it is a Normal session

```
port1: 25: Rcv TargetName: iqn.storageforduck
```

iSCSI host logs into target he got from sendtargets response

```
port1: 26: Init PWWN 20:08:00:0b:be:77:72:42
port1: 27: Init NWWN 20:09:00:0b:be:77:72:42
port1: 28: Target PWWN 20:05:00:a0:b8:0c:64:51
port1: 29: Target NWWN 20:04:00:a0:b8:0c:64:50
```

WWN's assigned

```
port1: 30: Set our RcvDataFieldSize to TCP MSS 1368 (Pass-thru)
port1: 31: iscsi_send_login_resp: State: S4.0.1, CSG 0 NSG 1, TBIT 1, transit 0 got_MRPL 0,
sent_MRPL 0, sent_1st_LoginOperNeg 0
port1: 32: Put TargetPortalGroupTag key: TargetPortalGroupTag=128, new_len 24, out_len 0
```

MSS parameters set

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

60

Cisco iSCSI Driver Discovery & Target Login Debug: Part 4 of 5

Cisco.com

```
port1: 33: Find_common_key: num_proposed 1, num_acceptable 1
port1: 34: init propose 0
port1: 35: tgt allows 0
port1: 36: iscsi_send_login_resp auth common index: 0, method: 0
port1: 37: Add key AuthMethod=None, new_len 15, out_len 25
port1: 38: S4.2: MaxRecvDataSegmentLength found (524288) or T-bit is 1, fc sess 0xd8c80280, fc state 0, discovery session 0
port1: 39: fcp_init_sess_fsm: got RcvDataFieldSize 1024 from FC target 0x680100
```

Parameters for target session set

```
port1: 40: S4.2.1: rcv PLOGI ACC, T-bit 1
```

PLOGI to SCSI Target on SAN

```
port1: 41: iscsi_send_login_resp: State: S4.2.1, CSG 1 NSG 3, TBIT 1, transit 1 got_MRPL 1, sent_MRPL 0, sent_1st_LoginOperNeg 0
port1: 42: Put ImmediateData key: ImmediateData=No, new_len 16, out_len 0
port1: 43: Put MaxRecvDataSegmentLength key: MaxRecvDataSegmentLength=1024, new_len 29, out_len 17
port1: 44: Add key X-com.cisco.sendAsyncText=No, new_len 28, out_len 47
port1: 45: Add key InitialR2T=Yes, new_len 14, out_len 76
port1: 46: Send_login_resp: our_MaxRecvDSLen 1024, Rcvd value 524288, FC RcvDataFieldSize 1024
```

Login response sent back to iSCSI host with negotiated MSS and other parameters

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

61

Cisco iSCSI Driver Discovery & Target Login Debug: Part 5 of 5

Cisco.com

```
port1: 47: Add key MaxBurstLength=524288, new_len 21, out_len 91
port1: 48: Add key FirstBurstLength=524288, new_len 23, out_len 113
port1: 49: Add key MaxConnections=1, new_len 16, out_len 137
```

PDU's with key values based on iSCSI RFC, these should match values in the displays on Cisco iSCSI driver

```
port1: 50: Find_common_key: num_proposed 2, num_acceptable 2
port1: 51: init propose 0
port1: 52: tgt allows 1
port1: 53: tgt allows 0
port1: 54: iscsi_send_login_resp digest common index: 1, method: 0
port1: 55: Add key HeaderDigest=None, new_len 17, out_len 154
port1: 56: Find_common_key: num_proposed 2, num_acceptable 2
port1: 57: init propose 0
port1: 58: tgt allows 1
port1: 59: tgt allows 0
port1: 60: iscsi_send_login_resp digest common index: 1, method: 0
```

Final iSCSI parameters exchanged and SCSI data allowed to flow

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

62

Other iSCSI Debugs

Cisco.com

debug ips iscsi

all	Turn on all iSCSI debug	fsm	Configure iSCSI FSM debugging
api	Configure iSCSI API debugging	login	Configure iSCSI login debugging
buf	Configure iSCSI buffer debugging	parse	Configure iSCSI parsing debugging
error	Configure iSCSI error debugging	pkt	Configure iSCSI packet debugging
event	Configure iSCSI event debugging	task	Configure iSCSI transaction debugging
fcpevent	Configure iSCSI fcp event debugging	transaction	Configure iSCSI transaction debugging
fcframe	Configure iSCSI fcp frame debugging	warn	Configure iSCSI warning debugging
fcpsm	Configure iSCSI FCP FSM debugging		

Most debugs require a port or session for information you are wanting to debug

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

63

Using The MDS Tools

Cisco.com

- **Use Ethereal on iSCSI Host to debug iSCSI PDU's**
- **Span GigE interface on Catalyst Switch to Ethereal to see iSCSI PDU's in Network**
- **Use FCanalyzer Local or remote to view iSCSI Initiator in the MDS Fabric**
- **Span iSCSI interface on MDS to FC SD port and use PAA and Ethereal to see iSCSI Initiator and Target flows on the MDS Fabric**

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

64

Cisco.com




Cisco.com

- Ability to SPAN off iSCSI interface to SPAN Destination port on MDS
- Use Ethereal to view and troubleshoot SCSI data and exchanges between iSCSI Initiator on the IPS Blade to the Target on the Fabric

OPT-3052
9909 06 2004 X

SN5428-2 iSCSI Debugging Internals

Cisco.com



OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

67

SN5428-2 Debugs

Cisco.com

- Depends on which support mode the SN5428-2 is in
iSCSI router mode —debug scsirouter
 and/or **FCIP mode** —debug FCIP
- Both modes have debug for GigE interface
- Only iSCSI mode has the iSCSI tracing
- Debugging different from typical Cisco IOS
- Debugs are enabled and captured to a circular file by default
- Captured debug information is then displayed via show commands

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

68

SN5428 Debug Outputs

Cisco.com

[Gopher-SN5428-2]# debug scsirouter "name"

iscsitrace	- iSCSI trace facility to look at iSCSI Protocol PDU's
scsitrace	- SCSI trace facility to look at SCSI Protocol CDB's
target	- iSCSI trace of just the target you want

[Gopher-SN5428-2]# show debug

fcswitch	- Show Fibre Channel Switch attributes
fcip	- Show FCIP information (covered in FCIP Section)
interface	- Show interface debug info
iscsitrace	- Show iScsi trace facility global settings
isns	- iSNS debug information (not covered in this session)
mailboxtrace	- Show mailbox trace data (Engineering detail)
rawlunatabase	- Show raw inquiry data from lun(s) discovery
scsirouter	- Scsirouter debug information

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

69

SN5428-2 FC Device Debugging

Cisco.com

This is a complete output of all LUN discovery that is done by the Fibre Channel firmware in the SN5428. The Report LUNs is done to all WWPN's returned to the SN5428 after a Query to the FC Directory Server

show debug rawlunatabase

```
Entry Address = 0x7595b6c
fabricLoginFailureCode=0x0,fabricLoginExtendedCode=0x0,fabricLoginTimeoutCode=0x0
ReportLunsLLDStatus=0x0,ReportLunsLLDStatusModifier=0x0,ReportLunsSCSIStatus=0x0,ReportLunsASCASCQ=0x0,ReportLunsLunCount=1
InquiryLLDStatus=0x0,InquiryLLDStatusModifier=0x0,InquiryLastLunWithLLError=0x0,
InquirySCSIStatus=0x0,InquiryASCASCQ=0x0,InquiryLastLunWithSCSIStatusError=0x0
boolLunsNotSupported=0x0,InquiryLastLunNotSupported=0x0
loopId=0x1,masterState=0x6,slaveState=0x7,loggedIn=0,roles=1,valid=1,portId=0xe6,scanLuns=0x0
numberLuns=0x1,reportAsyncEvent=0x0,node_wwn=0x20000020 0x37a70dd8, port_wwn=0x21000020 0x37a70dd8
lun=0, wwnn=0x20000020 0x37a70dd8, reportAsyncEvent=0x0
stdInquiry data for lun=0x0
bytes0-7=0x00000332 0x8b00500a
vendorId=SEAGATE , product=ST318451FC , revision=0001 device Type=0x0
DeviceIdPage:bytes0-3= 0x0083000c,bytes4-7= 0x01030008,bytes8-11= 0x20000020
:bytes12-15= 0x37a70dd8,bytes16-19=0x00800014,bytes20-23=0x33434330
S/N Page:bytes0-3= 0x00800014 s/n=3CC02AVP00007107EPTS
```

LUN count
returned

WWPN that Inquiry
is being done to

WWN on the LUN
Port

LUN is LUN 0, and this is
the data return from the
report LUN command

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

70

SN5428-2

Show debug scsirouter all scsitrace

Cisco.com

300 lines in the actual trace output, 500-800. These are lines 753-759

Target name configured on SN5428

SCSI CDB command 1a decoded as Mode sense 6

```

753: IntrID 0, TgtID 0, LUN 0, TmStmp 0x00000000
754: IntrID 0, TgtID 0, LUN 0, TmStmp 0x00000000
755: IntrID 1, TgtID 0, LUN 0, TmStmp 0x00001ecb iSCSITotFE LoginReq
      tskTag 0x0, iLmtP 0x0, mbuFP 0x0, scsiTskP 0x0
      scsiTskSt:None, flgs 0x0, destTgtID 0x0, destLun 0x0, dataLen 0x0
      TargetName: lun1
756: IntrID 1, TgtID 0, LUN 0, TmStmp 0x00001ee3 iSCSITotFE CmdReq
      tskTag 0x1, iLmtP 0xffe3e0, mbuFP 0x0, scsiTskP 0x75021d0
      scsiTskSt:Dormant, flgs 0x20, destTgtID 0x1, destLun 0x0, dataLen 0xc0
      CDB: 1a003f00 c0000000 00000000 00000000 (ModeSense_6)
757: IntrID 1, TgtID 0, LUN 0, TmStmp 0x00001ee3 TFEtoIntr InitCDB
      tskTag 0x1, iLmtP 0xffe3e0, mbuFP 0x7548b04, scsiTskP 0x75021d0
      scsiTskSt:Active, direct:Read, attr 0x4, cmdsOutStndng 0x0, dataLength 0xc0
      CDB: 1a003f00 c0000000 00000000 00000000 (ModeSense_6)
758: IntrID 1, TgtID 0, LUN 0, TmStmp 0x00001ee3 IntrTotFE CmdRespIn
      tskTag 0x1, iLmtP 0xffe3e0, mbuFP 0x7548b04, scsiTskP 0x75021d0
      scsiTskSt:Active, direct:Read, scsiTskFlgs:OnActiveQ, destLun 0x0, lldStat 0x0
      scsiStat 0xb02, dataLen 0xc0, residLen 0xc0, cmdsOutStndng 0x1, contrRtn 0x0
759: IntrID 1, TgtID 0, LUN 0, TmStmp 0x00001ee3 Sensedata
      tskTag 0x1, snsLen 0x12
      70000600 0000000a 00000000 29000300 000000c0 000000c0 00000001 00000000
  
```

Direction of command

Line 800 is the last event to occur, so to decode in order you would start from lower line # to higher line #

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

71

SN5428-2

show interface ge1 stats

Cisco.com

```

*[Gopher-SN5428-2]# show interface "ge1" stats
signal: signal detect duplex: full auto-negotiate: complete
flow control: rx pause
  
```

Over all health of GiGE interface can be viewed

errors	receive	transmit
0000000000 rx CRC	0398695391	0254254917 frames
0000000000 rx align	00000631266f235a	00000c0437e5b1c6 octets (hex)
0000000000 rx symbol	0398695391	0254254917 good frames
0000000000 rx phys	00000631266d22fe	00000c0437da6009 good octets (hex)
0000000000 rx fifo dr fr	0000037569	0000004048 BC frames
0000000000 rx sequence	0000041582	0000041589 MC frames
0000000000 rx carrier ext	0000000000	0000000000 XON packets
0000000000 rx length	0000000000	0000000000 XOFF packets
0000000000 rx unsup opcode		packets:
0000000000 rx no buffs	0143204125	4179649503 64 byte
0000000000 rx unsize fr	0218743078	0272001285 65-127 byte
0000000000 rx frag fr	0008130157	0036333525 128-255 byte
0000000000 rx osize fr	0000065126	0036206673 256-511 byte
0000000000 rx jabber evt	0076836458	0232013078 512-1023 byte
0000000000 tx underrun	4246683743	4087985445 1024-1522 byte
0000000000 tx defer evt		

ge1: RX 1.3 MB/s (1.1%) TX 3.1 MB/s (2.6%)

Throughput on GiGE can be monitored on this interface

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

72

SN5428-2 Show int fci1 stats

Cisco.com

```
*[Gopher-SN5428-2]# sh int fci1 stats
```

```
loop: LOOP READY
connection: F Port
Data Rate: 2 Gb/s
port id: 0x10000
ALPA: 0x0
firmware: READY
```

```
SCSI stats:
400254977 status IOCB
0 type 1 IOCB
529 marker IOCB
0 unhandled IOCB
```

```
WWPN: 28 00 00 05 9b a6 93 c0
```

```
SCSI Initiator stats:
```

```
400254977 SCSI Command status IOCBs
Breakdown of Exception status IOCBs:
000000007 LIP Reset Aborts
000000000 DMA Errors
000000000 Port Unavailable
000000023 Port Logged Out
000000000 Port Configuration Changed
000000006 Command Timeout
000000000 Data Overrun
000000000 Write Data Underrun - (No Data)
0000001966 Read Data Underrun - (No Data)
000000000 Queue Full
0 type 1 IOCBs
529 marker IOCBs
0 unhandled IOCBs
```

Continued →

```
0 SCSI Commands terminated by Firmware Restart
0 SCSI Commands terminated by Loop Down
814545 SCSI Commands terminated by Invalid Target
```

```
bytes read bytes written
00000b41bf5fa250 000005ab22defa00 (hex)
```

```
0000000531 Get All Next Requests Sent
0000000531 Get All Next Accepts Received
0000000000 Get All Next Rejects Received
```

```
Fabric Port Login Exceptions:
```

```
Lun discovery information:
portid=0x2d02e8 WWPN = 22 00 00 04 cf 75 21 d0 :
1 luns reported, ALL luns discovered.
portid=0x208da WWPN = 21 00 00 04 cf 67 3e c2 :
1 luns reported, ALL luns discovered.
```

```
SCSI Target stats:
```

```
0 Receive Failures
```

```
WWPN: 28 00 00 05 9b a6 93 c0
```

```
FC stats:
```

ERRORS	GENERAL
0000000000 link fail	0000000000 Loop Initializations (LIP F7)
0000000001 sync loss	0000000000 Loop Failures (LIP F8)
0000000001 sig loss	0000000000 Loop Resets (LIP FF)
0000000000 CRC error	0000000001 Loop Up Events
0000000000 unhand async	0000000000 Loop Down Events
	0000000004 Port Data Base Changed Events
0000000002 loop total	0000000095 Fabric Change Notifications

Complete
Discovery of
LUNs and how
many responded
to report all LUNs

WWN of initiator on
interface fci1 used to
allow access of this port
to the storage Array

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

73

GigE Packet Tracing

Cisco.com

```
*[SN5428-2]# debug interface ge1 trace pktsize 64 pktcnt 100 enable
```

```
*[SN5428-2]# sh debug int ge1 trace stats
Packet trace is enabled: max trace size 64, max packets remaining 96
Traces are being saved to buffer of size 131072 bytes
Packet trace buffer contains 4 packets (0% full)
```

```
*[SN5428-2]# show debug interface "ge1" trace stats
Packet trace is enabled: max trace size 64, max packets remaining 95
Traces are being saved to buffer of size 131072 bytes
Packet trace buffer contains 5 packets (0% full)
```

```
*[SN5428-2]# no debu int ge1 trace enable
```

```
*[SN5428-2]# sh debug interface ge1 trace hex ascii
```

```
first - Number of packets to be displayed from start of trace
last - Number of packets to be displayed from end of trace
<CR>
```

```
*[SN5428-2]# sh debug interface ge1 trace hex ascii
```

```
1: RX, len 370, time 0.000
0000: 01 00 0c cc cc cc 00 05 31 3f ee 83 01 64 aa aa *.....1?..d.*
0010: 03 00 00 0c 20 00 02 b4 e8 e5 00 01 00 0a 53 77 *.....Sw*
0020: 69 74 63 68 00 02 00 08 00 00 00 00 00 03 00 16 *ltch.....*
0030: 47 69 67 61 62 69 74 45 74 68 65 72 6e 65 74 30 *GigabitEthernet0*
```

```
2: RX, len 60, time 1.600
0000: ff ff ff ff ff 00 06 5b ec aa 42 08 06 00 01 *.....[.B...*
0010: 08 00 06 04 00 01 00 06 5b ec aa 42 0a 01 32 0d *.....[.B..2.*
0020: 00 00 00 00 00 00 0a 01 32 15 00 00 00 00 00 00 *.....2.....*
0030: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 *.....*
```

Debugging A Little Different
from IOS Type Outputs

Debug Is Turned On, But Output Is
Directed to a File

Debug Monitoring Is Done with
Showing of Stats

Debug Is Halted With "No"

Show Command Is Used to
Display the Captured Data

The Output From the Captured
File Is Saved and Can Be Shown
Until Another Debug Over Writes
the File

Ascii Output Can Be Selected
for Readable Form

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

74

iSCSI Protocol PDU debugging

Cisco.com

```
[Gopher-SN5428-2]# debug scsirouter BUCK iscsitrace pducnt 500 enable
```

Debug enabled to capture 500 PDU's

```
[Gopher-SN5428-2]# no debug scsirouter BUCK iscsitrace enable
```

Debug turned off with NO

```
[Gopher-SN5428-2]# sh debug scsirouter BUCK iscsitrace short
```

Use show command to view trace on brief output or detail

```
1(0.000): 192.168.1.1:3260 192.168.1.2:1404 CONN OPEN
TCP connection up
2(0.000): 192.168.1.2:1404 192.168.1.1:3260 40000137001c 1 LOGIN/I 329 T CSG=SN NSG=ON
3(0.000): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 1 LOGIN-RSP 39 T CSG=SN NSG=ON ver=0 OK
Login and response (detail output next slide)
4(0.000): 192.168.1.2:1404 192.168.1.1:3260 40000137001c 1 LOGIN/I 263 T CSG=ON NSG=FFP
5(0.000): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 1 LOGIN-RSP 285 T CSG=ON NSG=FFP ver=0 OK
Full Feature Phase-Login complete
```

SCSI data PDU's
Some detail can be seen in detail of trace or use scsi trace utility to see SCSI CDB's

```
6(0.016): 192.168.1.2:1404 192.168.1.1:3260 40000137001c 1 SCSI-CMD 0 F R op=a0 lun=0 ETLen=16
7(0.016): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 1 SCSI-DATA-IN 16 F
8(0.016): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 1 SCSI-RSP 0 F stat=0
9(0.016): 192.168.1.2:1404 192.168.1.1:3260 40000137001c 2 SCSI-CMD 0 F R Inquiry lun=0 ETLen=36
10(0.016): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 2 SCSI-DATA-IN 36 F
11(0.016): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 2 SCSI-RSP 0 F stat=0
12(0.016): 192.168.1.2:1404 192.168.1.1:3260 40000137001c 3 SCSI-CMD 0 F R Inquiry lun=0 ETLen=255
13(0.016): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 3 SCSI-DATA-IN 14 F
14(0.016): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 3 SCSI-RSP 0 F U stat=0
15(0.016): 192.168.1.2:1404 192.168.1.1:3260 40000137001c 4 SCSI-CMD 0 F R Inquiry lun=0 ETLen=255
16(0.016): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 4 SCSI-DATA-IN 16 F
17(0.016): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 4 SCSI-RSP 0 F U stat=0
18(0.016): 192.168.1.2:1404 192.168.1.1:3260 40000137001c 5 SCSI-CMD 0 F R Inquiry lun=0 ETLen=255
19(0.016): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 5 SCSI-DATA-IN 24 F
20(0.016): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 5 SCSI-RSP 0 F U stat=0
21(0.033): 192.168.1.2:1404 192.168.1.1:3260 40000137001c 6 SCSI-CMD 0 F R ReadCap lun=0 ETLen=8
22(0.033): 192.168.1.1:3260 192.168.1.2:1404 40000137001c 6 SCSI-RSP 28 F U stat=2
```

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

75

Output Example of iSCSI trace detail

Cisco.com

```
[Gopher-SN5428-2]# show debug scsirouter "BUCK" iscsitrace first 20
```

TCP Connection up

```
1: 192.168.1.1:3260 <-> 192.168.1.2:1404, len 32, time 0.000
event: CONN OPEN
```

Login PDU OPCODE 43

```
2: 192.168.1.2:1404 -> 192.168.1.1:3260, isid 40000137001c, len 380, time 0.000
```

Login Response

```
opcode: 43 (LOGIN,IMMEDIATE)
transit: 1
continue: 0
current stage: 0 (SECURITY NEGOTIATION)
next stage: 1 (LOGIN OPERATIONAL NEGOTIATION)
max version: 0
min version: 0
dataSegLength: 329
isid: 40000137001c
tsih: 0000
itt: 00000001
cid: 1
cmdSN: 00000001
expStatSN: 00000000
TEXT:
InitiatorName=iqn.1991-05.com.microsoft:g01234567890123456
SessionType=Normal
TargetName=iqn.1987-05.com.cisco:00.82226e828635.buck-lun
AuthMethod=None
```

IQN Name of Host logging in, Normal session

iSCSI Target acquired during login

```
3: 192.168.1.1:3260 -> 192.168.1.2:1404, isid 40000137001c, len 88, time 0.000
opcode: 23 (LOGIN-RSP)
transit: 1
continue: 0
current stage: 0 (SECURITY NEGOTIATION)
next stage: 1 (LOGIN OPERATIONAL NEGOTIATION)
max version: 0
active version: 0
dataSegLength: 39
isid: 40000137001c
tsih: 0000
itt: 00000001
statSN: 00000000
expCmdSN: 00000001
maxCmdSN: 00000008
statusClass: 00 (SUCCESS)
statusDetail: 00 (SUCCESS)
TEXT:
TargetPortalGroupTag=1
AuthMethod=None
```

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

76

SN5428-2

Most debugs can be gathered from GUI

Cisco.com

The screenshot shows the Cisco Storage Network 5428-2 GUI. The top navigation bar includes Monitor, Configuration, Maintenance, Troubleshooting, Support, Home, and Help. The Troubleshooting page is active, displaying a list of links organized into sections: Reset Actions (Clear Counters, Clear Log File, Clear FC Logs, Clear FC Zones), Configuration Data (Boot Config, Running Config, FC Config, FC Link State Database, FC Name Server Database), System Data (Boot, Buffers, Log File Data, Memory, Module, Raw Lun Database, Sessions), Stack, Tasks, SCSI trace status, SCSI trace, SCSI target trace, AAA Debug, AAA Test Request, AAA Test Cancel, AAA Test Query, Debug Files (Current Crash, Saved Crash, Log File, Search Log File, Tail Log File, FC Logs), TacOps, Technical Support, Ping, Config FC Logs, FC Loopback, and FC Beacon. An arrow points from the 'Troubleshooting' navigation tab to the 'Troubleshooting' section header.

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

77

iSCSI Host Troubleshooting

Cisco.com

Make sure that the NIC is configured correctly. Some options must match how the switch is configured.

The screenshot shows the Windows Computer Management console. The 'Device Manager' tree on the left is expanded to show 'Network adapters'. The 'Intel(R) PRO/1000 XT Server Adapter Properties' dialog box is open, showing the 'Advanced' tab. The 'Flow Control' property is set to 'Disabled'. Callout boxes point to the 'Flow Control' property and the 'Jumbo frames larger than 1500' property, which is also set to 'Disabled'. Another callout box points to the 'IP and TCP checksum offload' property, which is set to 'Enabled'. The 'Wake On LAN' property is also visible and set to 'Enabled'.

Flow Control

Jumbo frames larger than 1500

IP and TCP checksum offload

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

78

iSCSI Host Troubleshooting

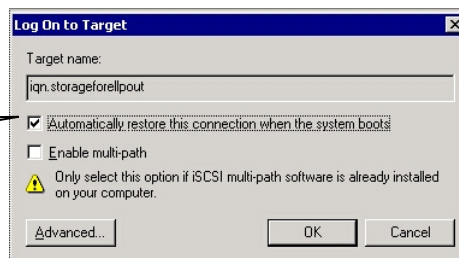
Cisco.com

Take note of the VSAN that the initiator is in. If not using IVR, it must match the storage port VSAN

Fabric Zoning and Array lun masking and security software need to have the PWWN configured

For Microsoft Initiator
Remember to make targets persistent or they will be lost at next reboot

iSCSI Node name is 10.1.11.26
iSCSI Initiator name: iqn.1991-05.com.microsoft:eelpout
iSCSI alias name:
Node WWN is 21:0d:00:0d:ec:00:ea:42 (configured)
Member of vsans: 4
Number of Virtual n_ports: 1
Virtual Port WWN is 21:0c:00:0d:ec:00:ea:42 (configured)
Interface iSCSI 3/1, Portal group tag: 0x100
VSAN ID 4, FCID 0x2d0104




Log On to Target

Target name:
iqn.storageforellpout

☒ Automatically restore this connection when the system boots

☐ Enable multi-path

 Only select this option if iSCSI multi-path software is already installed on your computer.

Advanced... OK Cancel

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

79

iSCSI Host Troubleshooting

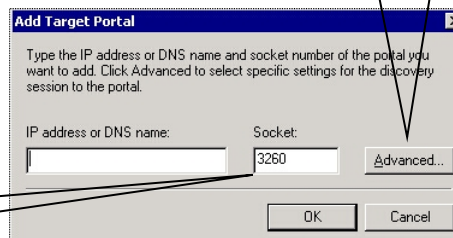
Cisco.com

- Revert to basic configuration with no chap passwords, and default zoning if you are unable to get the iscsi host connected.

- Verify that the scsi target is available from the MDS

Ensure that CHAP passwords are configured correctly on both the discovery panel, and the target panel
This is done under Advanced options

The TCP ports must match on both the driver configuration, and the iSCSI gateway configuration



Add Target Portal

Type the IP address or DNS name and socket number of the portal you want to add. Click Advanced to select specific settings for the discovery session to the portal.

IP address or DNS name: Socket: 3260

Advanced... OK Cancel

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

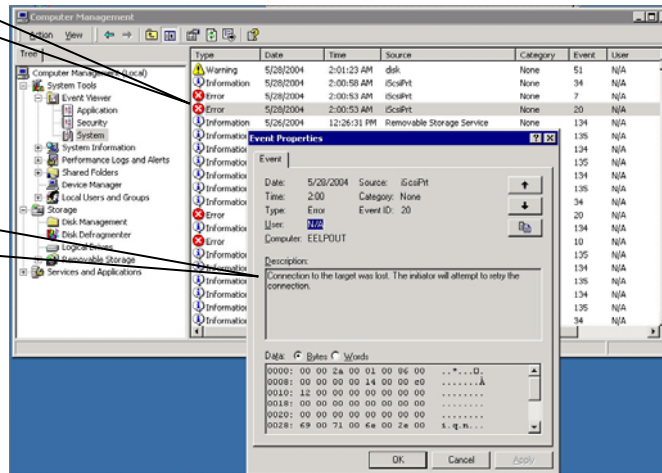
80

iSCSI Host Troubleshooting

Cisco.com

Under Event Viewer – System You can see iScsiPrt error messages

View the message detail. Ensure that the time stamps agree with the event you are investigating.



OPT-3052

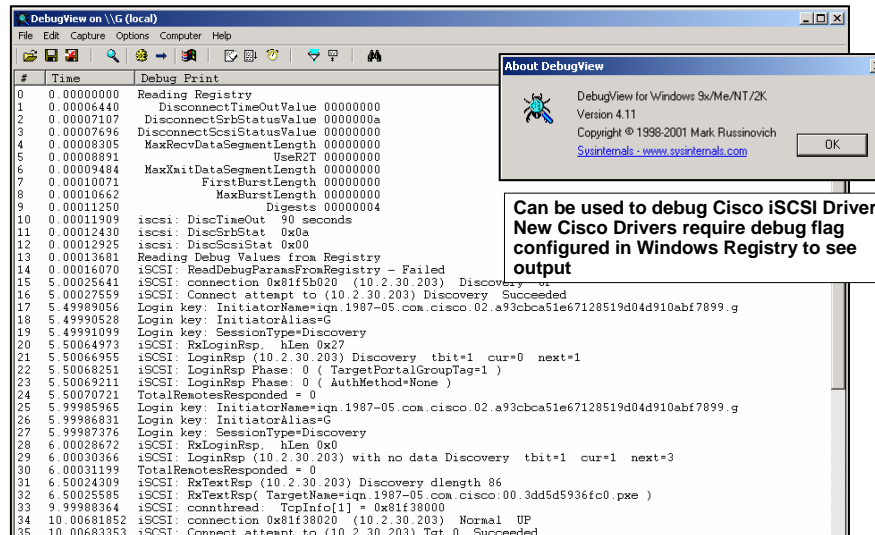
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

81

Windows iSCSI Debug

Cisco.com



Can be used to debug Cisco iSCSI Driver, New Cisco Drivers require debug flag configured in Windows Registry to see output

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

82

Linux Support Commands

Cisco.com

Useful Commands:

- **/var/log/iscsi.log**—This log contains iSCSI diagnostic information if there were problems during the iSCSI kernel module load
- **/usr/bin/iscsi-ls l**—Utility to view current available iSCSI devices on the system (new in 3.2 drivers)
- **/proc/scsi/iscsi**—Identifies the controller number assigned to the iSCSI HBA
- **/etc/rc.d/init.d/iscsi reload**—Rediscover command for new iSCSI devices that have been added to iscsi.conf (does not effect connected targets)
- Read and follow release notes that come with the Linux Driver

To check version of Linux driver loaded:

```
[Linux]# cd/usr/src
```

```
[Linux]# ls -l iscsi
```

```
lrwxrwxrwx 1 root root 17 Feb 28 14:31 iscsi -> linux-iscsi-3.1.1
```

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

83

iSCSI Host Troubleshooting: Check List

Cisco.com

- Is the IP connectivity between host and portal correct?
- Is the Zoning correct for iscsi and storage pwwn?
- Is the Array configured to permit iscsi pwwn to access luns?
- Are the CHAP passwords configured the same in all places?
- Are targets marked 'persistent' in the Microsoft driver?
- Is the NIC card configured correctly for network attachment?
- Ethereal can be run on the host to trace the iSCSI packets.
- Use Disk Manager to confirm that luns are visible to the OS.
- Use the System Event Viewer to check for iScsiPrt errors.
- Read release notes
- Windows requires Hotfixes in some O/S settings

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

84

OVERVIEW OF MDS AND SN5428-2 IP STORAGE - FCIP



OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

85

Section Agenda

Cisco.com

- Understanding the FCIP architecture on MDS
- ISL's and FCIP
- Inter-Vsan Routing Design Troubleshooting
- Troubleshooting FCIP
- Understanding the FCIP architecture on SN5428-2

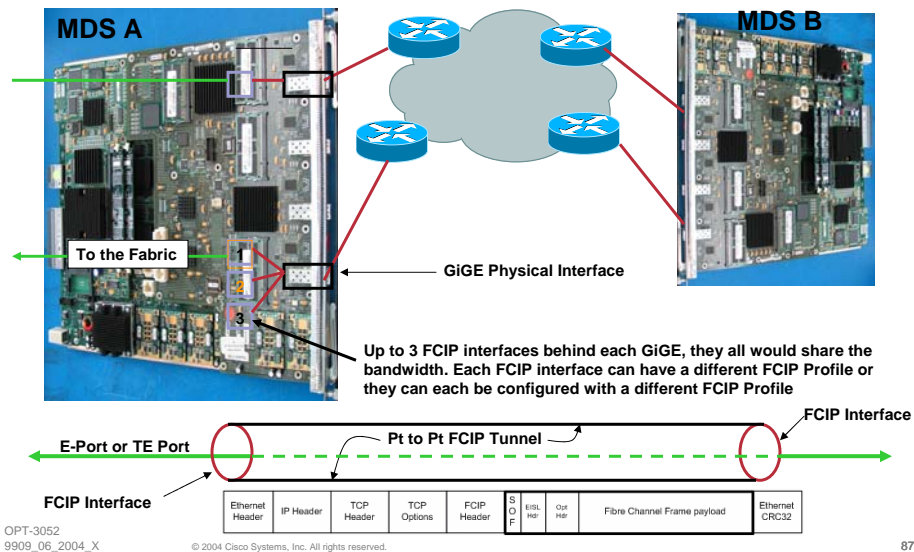
OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

86

FCIP Visual Help


Cisco.com



Network Topology: Datapoints

Cisco.com

- Lab Network has 2 FCIP links between MDS Switch Fox and MDS Switch Musky
- We are using InterVsan Routing Feature (IVR) to limit what FC services and FCID's cross the FCIP Links
- FCIP interfaces are in a Port-channels for High Availability and increased Bandwidth needs
- Features configured to limit what VSAN's can cross the FCIP TE-Port ISL.



OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

88

Profiles and FCIP Interfaces

Cisco.com

Profiles for Tunnels

Profile	Interface	Attached	Enable	IP Port	Remote Profile	Special Frames	Remote Profile ID
1	fcip2	ggf201		50.1.1.2	3225	0x0000000000000000	0
2	fcip3	ggf201		50.1.1.2	3225	0x0000000000000000	0

FCIP Interfaces

Interface	Description	Port/Vlan	Oper	Admin	Oper	Failure Cause	Last Change
fcip2	FCIP Link to Fox	201 TE	up	up	none	none	5/26
fcip3	FCIP Link to Fox	201 TE	up	up	none	none	5/26

Profile Detail Configurations

Listen on	IP Address	Port	SACh	Keep	Retransmission	Send	Bandwidth (Kb/s)	Est Round	Path MTU
1	50.1.1.2	3225	0	60	200	4	0	1000000	15000
2	50.1.1.2	3225	0	60	200	4	0	1000000	15000
3	50.1.1.2	3225	0	60	200	4	0	1000000	15000

GiGE IP interfaces

GiGE IP interfaces

OPT-3052
9909_06_2004_X
© 2004 Cisco Systems, Inc. All rights reserved.

VSAN's and IVR (Inter VSAN Routing)

Cisco.com

Musky-9506# sh ivr vsan-topology active

AFID	SWITCH	WWN	Active	Cfg. VSANS
1	20:00:00:0b:be:77:72:40	yes	yes	200-201
1	20:00:00:0d:ec:00:ea:40	*	yes	50,201

Total: 2 entries in active IVR VSAN-Topology

Current Status: Inter-VSAN topology is ACTIVE

Last activation time: Thu May 27 12:06:06 2004

Switch Moose is not needed to be configured for IVR

zone name IVRZ_IVR_Badger vsan 200

* fcid 0xb10002 [pwwn 10:00:00:00:c9:30:ba:06]

* fcid 0xd100e8 [pwwn 21:00:00:20:37:b9:9f:5f]

zone name IVRZ_IVR_Badger vsan 50

* fcid 0xb10002 [pwwn 10:00:00:00:c9:30:ba:06]

* fcid 0xd100e8 [pwwn 21:00:00:20:37:b9:9f:5f]

OPT-3052
9909_06_2004_X
© 2004 Cisco Systems, Inc. All rights reserved.

Displaying the FCIP Interface

Cisco.com

Verify IP address tunnel end points

TCP parameters and packet flows

Fibre Channel Frame flows across FCIP tunnel, Reassembled frames tell how many FC data frames are being segmented form fit into MTU of IP Network

Compression statistics if using compression

Timestamp errors if syncing with NTP on switches

```

FOX-9216# sh int fcip 6 counters
fcip6
TCP Connection Information
2 Active TCP connections
Control connection: Local 50.1.1.2:3225, Remote 50.1.1.1:65527
Data connection: Local 50.1.1.2:3225, Remote 50.1.1.1:65529
22 Attempts for active connections, 4 close of connections
TCP Parameters
Path MTU 2300 bytes
Current retransmission timeout is 200 ms
Round trip time: Smoothed 2 ms, Variance: 1
Advertized window: Current: 80 KB, Maximum: 103 KB, Scale: 4
Peer receive window: Current: 84 KB, Maximum: 150 KB, Scale: 4
Congestion window: Current: 10 KB, Slow start threshold: 145 KB
5 minutes input rate 57732216 bits/sec, 7216527 bytes/sec, 3530 frames/sec
5 minutes output rate 14705680 bits/sec, 1838210 bytes/sec, 1095 frames/sec
2405074 frames input, 4449922164 bytes
205393 Class F frames input, 19326984 bytes
2199681 Class 2/3 frames input, 4430595180 bytes
0 Reass frames
0 Error frames timestamp error 0
915734 frames output, 1177424204 bytes
205570 Class F frames output, 19355124 bytes
710164 Class 2/3 frames output, 1158069080 bytes
0 Error frames
IP compression statistics
4511410480 input bytes 422682273 compressed input bytes 10.67 ratio
1215167612 output bytes 182800613 compressed output bytes 6.65 ratio
    
```

FOX-9216# clear counters interface fcip 6

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

91

Detailed IP Stats

Cisco.com

FOX-9216# sh ips ?

arp Show/clear arp table

internal Show ips internal information

ip Show route table

stats Various ips stats related commands

status Show status of ips module

eth-trace-logs Enter ethernet interface

event-history Show internal event log information

fcip-trace-log Show ips fcip trace log

info Show internal data structure information

mem-stats Show ips memory information

Detail of IPS information if suspected issue is with IP and TCP

These are great outputs to be captured to troubleshoot connectivity and throughput issues

buffer Show ips buffer stats

dma-bridge Show dma-bridge stats

icmp Show icmp stats

ip Show ip stats

mac Show ethernet mac stats

tcp Show tcp stats

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

92

GiGE Interface Stats

Cisco.com

Musky-9506# show ips stats tcp interface gigabitethernet 3/3 detail
TCP Statistics for port GigabitEthernet3/3

TCP send stats

42259 segments, 72997936 bytes
35693 data, 6420 ack only packets
0 control (SYN/FIN/RST), 0 probes, 146 window updates
0 segments retransmitted, 0 bytes
0 retransmitted while on ethernet send queue, 0 packets split
4049 delayed acks sent

TCP receive stats

30074 segments, 10466 data packets in sequence, 17313376 bytes in sequence
5 predicted ack, 10449 predicted data
0 bad checksum, 0 multi/broadcast, 0 bad offset
0 no memory drops, 0 short segments
0 duplicate bytes, 0 duplicate packets
0 partial duplicate bytes, 0 partial duplicate packets
0 out-of-order bytes, 0 out-of-order packets
0 packet after window, 0 bytes after window
0 packets after close
18550 acks, 72997936 ack bytes, 0 ack toomuch, 0 duplicate acks
0 ack packets left of snd_una, 0 non-4 byte aligned packets
12465 window updates, 0 window probe
0 pcb hash miss, 0 no port, 0 bad SYN, 0 paws drops

Musky-9506# clear ips stats all (or per interface)

TCP Connection Stats

0 attempts, 0 accepts, 0 established
0 closed, 0 drops, 0 conn drops
0 drop in retransmit timeout, 0 drop in keepalive timeout
0 drop in persist drops, 0 connections drained

TCP Miscellaneous Stats

7957 segments timed, 18550 rtt updated
0 retransmit timeout, 0 persist timeout
0 keepalive timeout, 0 keepalive probes

TCP SACK Stats

0 recovery episodes, 0 data packets, 0 data bytes
0 data packets retransmitted, 0 data bytes retransmitted
0 connections closed, 0 retransmit timeouts

TCP SYN Cache Stats

0 entries, 0 connections completed, 0 entries timed out
0 dropped due to overflow, 0 dropped due to RST
0 dropped due to ICMP unreachable, 0 dropped due to bucket overflow
0 abort due to no memory, 0 duplicate SYN, 0 no-route SYN drop
0 hash collisions, 0 retransmitted

TCP Active Connections

Local Address	Remote Address	State	Send-Q	Recv-Q
50.1.2.1:3225	50.1.2.2:65296	ESTABLISH	0	0
50.1.2.1:3225	50.1.2.2:65296	ESTABLISH	0	0
50.1.2.1:3225	0.0.0.0:0	LISTEN	0	0

Excessive retransmits
indicates possible core
drops and/or tcp window
size should be adjusted

1 TCP Connection for
Class F other for Data

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

93

FCIP Considerations

Cisco.com

- Need for proper configuration setup for TCP
- Follow configuration guidelines for Min/Max bandwidth, Round Trip Time settings
- Consider MTU's of 2300 in your designs
- Pay big attention to Retransmissions they will be first indication of network issues and improper network equipment mis-configurations

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

94

Verifying MTU

Is it what you think it is?

Cisco.com

Musky-9506# clear counters interface gigabitethernet 3/3

show interface GiG 3/3
MTU configured under switchport settings.

```
Musky-9506# sh interface gigabitethernet 3/3
GigabitEthernet3/3 is up
Port description is FCIP
Hardware is GigabitEthernet, address is 0005.3000.a408
Internet address is 50.1.2.1/24
MTU 2300 bytes
Port mode is IPS
Speed is 1 Gbps
Beacon is turned off
Auto-Negotiation is turned on
5 minutes input rate 16 bits/sec, 2 bytes/sec, 0 frames/sec
5 minutes output rate 16 bits/sec, 2 bytes/sec, 0 frames/sec
3762808 packets input, 427481218 bytes
0 multicast frames, 0 compressed
0 input errors, 0 frame, 0 overrun 0 fifo
5240056 packets output, 982552407 bytes, 0 underruns
0 output errors, 0 collisions, 0 fifo
0 carrier errors
```

From: show interface FCIP x

```
TCP Parameters
Path MTU 2300 bytes
Current retransmission timeout is 200 ms
Round trip time: Smoothed 2 ms, Variance: 1
Advertized window: Current: 101 KB, Maximum: 39 KB, Scale: 4
Peer receive window: Current: 66 KB, Maximum: 116 KB, Scale: 4
Congestion window: Current: 10 KB, Slow start threshold: 966 KB
5 minutes input rate 1256 bits/sec, 157 bytes/sec, 0 frames/sec
5 minutes output rate 304 bits/sec, 38 bytes/sec, 0 frames/sec
7751 frames input, 14420052 bytes
0 Class F frames input, 0 bytes
7751 Class 2/3 frames input, 14420052 bytes
0 Reass frames
0 Error frames timestamp error 0
8050 frames output, 15809456 bytes
0 Class F frames output, 0 bytes
8050 Class 2/3 frames output, 15809456 bytes
0 Error frames
```

This is controlled via PMTU enabled, If Path MTU discovery times out, this will remain 1500

Check for reassembly of frames, should be zero if MTU set proper to fit FCP frame

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

95

GiGE MTU Path

Cisco.com

The Command Line Interface of the MDS9000 supports extended IP pings to help in determining path availability, MTU and round trip time.

To manually determine the Path MTU, do the following:

- Note the peer IP addresses for each end of the FCIP link under test.
- At the cli, issue the ping command without any operands. Enter the appropriate responses at each of the following prompts, ensuring "y" is answered to the "Extended Commands" prompt.
- Repeat step 2 using different values for "Datagram size" until the largest datagram size is found.

```
9500 Switch# ping
Target IP address: 10.0.2.2
Repeat count: 3
Datagram size: 2194
Timeout in seconds: 2
Extended commands (y/n): y
Source address: 10.0.1.2
Data pattern (hex:16): 1111
PATTERN: 0x1111
PING 10.0.2.2 (10.0.2.2): 2194 data bytes
2202 bytes from 10.0.2.2: icmp seq=0 ttl=253 time=13.5 ms
2202 bytes from 10.0.2.2: icmp seq=1 ttl=253 time=12.9 ms
2202 bytes from 10.0.2.2: icmp seq=2 ttl=253 time=13.0 ms

--- 10.0.2.2 ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 12.9/13.1/13.5 ms
```

PASS

```
9500 Switch# ping
Target IP address: 10.0.2.2
Repeat count: 3
Datagram size: 2198
Timeout in seconds: 2
Extended commands (y/n): y
Source address: 10.0.1.2
Data pattern (hex:16): 1111
PATTERN: 0x1111
PING 10.0.2.2 (10.0.2.2): 2198 data bytes

--- 10.0.2.2 ping statistics ---
3 packets transmitted, 0 packets received, 100% packet loss
```

FAIL

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

96

FCIP Interface Details

Cisco.com

Check:

- VSAN Information
- Proper Configuration settings have been invoked
- Profiles, GiGE interfaces, and Peers are correct

```
Musky-9506# sh int fcip 2
fcip2 is trunking
  Hardware is GigabitEthernet
  Port WWN is 20:8a:00:0d:ec:00:ea:40
  Peer port WWN is 20:4a:00:0b:be:77:72:40
  Description: FCIP Link to FOX
  Admin port mode is E, trunk mode is auto
  Port mode is TE
  vsan is 201
  Belongs to port-channel 3
  Trunk vsans (allowed active) (1,201)
  Trunk vsans (operational) (1,201)
  Trunk vsans (up) (1,201)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
  Using Profile id 1 (interface GigabitEthernet3/3)
  Peer Information
  Peer Internet address is 50.1.2.2 and port is 3225
  Write acceleration mode is off
  IP Compression is enabled and set for higher
  compression ratio
  Special Frame is disabled
  Maximum number of TCP connections is 2
  Time Stamp is disabled
  QOS control code point is 0
  QOS data code point is 0
  B-port mode disabled
```

```
TCP Connection Information
  2 Active TCP connections
  Control connection: Local 50.1.2.1:65516, Remote 50.1.2.2:3225
  Data connection: Local 50.1.2.1:65518, Remote 50.1.2.2:3225
  10 Attempts for active connections, 2 close of connections
TCP Parameters
  Path MTU 2300 bytes
  Current retransmission timeout is 200 ms
  Round trip time: Smoothed 2 ms, Variance: 1
  Advertized window: Current: 101 KB, Maximum: 39 KB, Scale: 4
  Peer receive window: Current: 66 KB, Maximum: 116 KB, Scale: 4
  Congestion window: Current: 10 KB, Slow start threshold: 966 KB
  5 minutes input rate 1256 bits/sec, 157 bytes/sec, 0 frames/sec
  5 minutes output rate 304 bits/sec, 38 bytes/sec, 0 frames/sec
  7751 frames input, 14420052 bytes
  0 Class F frames input, 0 bytes
  7751 Class 2/3 frames input, 14420052 bytes
  0 Reass frames
  0 Error frames timestamp error 0
  8050 frames output, 15809456 bytes
  0 Class F frames output, 0 bytes
  8050 Class 2/3 frames output, 15809456 bytes
  0 Error frames
```

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

97

Steps for FCIP to Come Up

Cisco.com

1. FCIP Protocol Handshake

TCP/IP Path verified by user

FCIP Profiles and the FCIP Interface need to connect from both sides

Debug FSM to view this exchange

2. E-Port FC-SW-2 Standards Handshake

ELP-ESC-EPP-EFP-DIA-MR

Use FCAnalyzer or SPAN to PAA to view exchange

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

98

FCIP Debugs: FCIP Protocol Handshake

Cisco.com

Again, you must attach to the Module you are debugging

MDS9500# debug ips fcip ?

all	Configure FCIP debugging
ctlpkt	Configure FCIP Control Pkt debugging
ctlpkt-detail	Configure verbose FCIP Control Pkt debugging
datapkt	Configure FCIP Data Pkt debugging
datapkt-detail	Configure verbose FCIP Data Pkt debugging
err	Configure FCIP Error debugging
fsm	Configure FCIP FSM debugging (detail look on next set of slides)

port selection debugging for each category

```
<debug ips fcip err port 3>
```

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

99

FCIP MDS Internal Debugging: 1 of 3

Cisco.com

Musky-9506# attach mod 3

Attaching to module 3 ...

To exit type 'exit', to abort type '\$.'

This flow of protocol goes in both directions at same time if default of passive is not set on the profile

module-3# debug ips fcip fsm port 3

module-3# May 28 10:06:46 2004 port3: 1:FCIP: Ethernet link is UP; cancel the debounce timer

Message that GiGE port is up and connected

May 28 10:06:46 2004 port3: 2:FCIP: Eth-if 2102000 UP & RUNNING, try to bring up FCIP tunnels using this IF

May 28 10:06:46 2004 port3: 3:FCIP2: Bring up tunnel because eth-if up

Tunnels will now attempt to connect

May 28 10:06:46 2004 port3: 4:FCIP2: Try to Bring UP the Tunnel

May 28 10:06:46 2004 port3: 5:FCIP2: Start TCP listener with peer: 50.1.2.2:3225

Trying to connect to 51.1.2.2

May 28 10:06:46 2004 port3: 6:FCIP: Create a new listener object for 50.1.2.1:3225

May 28 10:06:46 2004 port3: 7:FCIP: Create FCIP Listener with local info: 50.1.2.1:3225

FCIP2 is listening for connections

May 28 10:06:46 2004 port3: 8:FCIP2: cancel bring up timer

May 28 10:06:46 2004 port3: 9:FCIP2: Start the bringup tunnel timer, timeout: 1120

May 28 10:06:47 2004 port3: 10:FCIP2: Received new TCP connection from peer: 50.1.2.2:65110

First TCP connection coming in from GiGE network (F class data)

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

100

FCIP MDS Internal Debugging: 2 of 3

Cisco.com

May 28 10:06:47 2004 port3: 11:FCIP2: Create a DE 0xd802cb80 for this tunnel
May 28 10:06:47 2004 port3: 12:FCIP2: Bind the DE 0xd802cb80 [1] to tunnel LEP 0x8014a120

De-Encapsulation /Encapsulation and Link End Point created (F-class)

May 28 10:06:47 2004 port3: 13:FCIP2: cancel bring up timer
May 28 10:06:47 2004 port3: 14:FCIP2: Bind DE 1 to TCP-hdl 0xd8c7c000
May 28 10:06:47 2004 port3: 15:FCIP2: Bind DE 1 to eport 0x80148e60
May 28 10:06:47 2004 port3: 16:FCIP2: bind de 1 in eport 0x80148e60, hash = 1 num-conn: 2

Create 1 flow through the FCIP Tunnel

May 28 10:06:47 2004 port3: 17:FCIP2: Received new TCP connection from peer: 50.1.2.2:65108

Second TCP connection coming in from 51.1.2.2 (Class 2/3 Data)

May 28 10:06:47 2004 port3: 18:FCIP2: Create a DE 0xd802ca40 for this tunnel
May 28 10:06:47 2004 port3: 19:FCIP2: Bind the DE 0xd802ca40 [2] to tunnel LEP 0x8014a120

De-Encapsulation /Encapsulation and Link End Point created (Class 2/3 Data)

May 28 10:06:47 2004 port3: 20:FCIP2: cancel bring up timer
May 28 10:06:47 2004 port3: 21:FCIP2: Bind DE 2 to TCP-hdl 0xd8c7c800
May 28 10:06:47 2004 port3: 22:FCIP2: Bind DE 2 to eport 0x80148e60
May 28 10:06:47 2004 port3: 23:FCIP2: bind de 2 in eport 0x80148e60, hash = 2 num-conn: 2

Create 2 flow through the FCIP Tunnel

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

101

FCIP MDS Internal Debugging: 3 of 3

Cisco.com

May 28 10:06:47 2004 port3: 24:FCIP2: Send LINK UP to SUP

Message to Supervisor that FCIP is up

May 28 10:06:47 2004 port3: 25:FCIP2: *** Received eisl frame in E mode

EPP exchange to place E-Port into TE-Port, this means a good ELP has occurred

May 28 10:06:47 2004 port3: 26:FCIP2: SUP-> Set trunk mode: 2
May 28 10:06:47 2004 port3: 27:FCIP2: Change the operational mode to TRUNK

FCIP up as Trunk and operational

FCIP is healthy at this point

- Continue debug of connection if E-Port is:
Segmented, Domain ID Overlap, Zone Merge failure

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

102

FCIP E-Port

Using FCanalyzer in the MDS Fabric: 1 of 2

Cisco.com

From Configuration mode-----do a brief single line only trace of just traffic on VSAN 201

Musky-9506(config)# fcanalyzer local brief display-filter mdshdr.vsan==201 limit-captured-frames 300

ff.ff.fd -> ff.ff.fd 0x28d0 0xffff SW_ILS ELP

Exchange Switch Parameter sent
First thing after sync of FCIP tunnel

ff.ff.fd -> ff.ff.fd 0x28d0 0xa77f FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0x28d0 0xa77f SW_ILS SW_ACC (ELP)
ELP Except from other switch

ff.ff.fd -> ff.ff.fd 0x28d0 0xa77f FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0x28d1 0xffff SW_ILS ESC
Exchange Switch Capabilities sent

ff.ff.fd -> ff.ff.fd 0x28d1 0xa780 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0x28d1 0xa780 SW_ILS SW_ACC (ESC)
Accept from other switch

ff.ff.fd -> ff.ff.fd 0x28d2 0xffff SW_ILS 0x71
ff.ff.fd -> ff.ff.fd 0x28d2 0xa781 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa782 0xffff SW_ILS 0x71
ff.ff.fd -> ff.ff.fd 0xa782 0x28d3 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0x28d2 0xa781 SW_ILS SW_ACC (EPP)
ff.ff.fd -> ff.ff.fd 0x28d2 0xa781 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa782 0x28d3 SW_ILS SW_RJT (EPP)
EPP exchange and ACK

ff.ff.fd -> ff.ff.fd 0xa782 0x28d3 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa783 0xffff SW_ILS 0x71
ff.ff.fd -> ff.ff.fd 0xa783 0x28d4 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa783 0x28d4 SW_ILS SW_ACC (0x71)
ff.ff.fd -> ff.ff.fd 0xa783 0x28d4 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa784 0xffff SW_ILS ELP
ff.ff.fd -> ff.ff.fd 0xa784 0x28d5 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0x28d6 0xffff SW_ILS ELP
ff.ff.fd -> ff.ff.fd 0x28d6 0xa785 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0x28d6 0xa785 SW_ILS SW_ACC (ELP)
ff.ff.fd -> ff.ff.fd 0xa784 0x28d5 SW_ILS SW_RJT (ELP)
ELP from other switch

ff.ff.fd -> ff.ff.fd 0x28d6 0xa785 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa784 0x28d5 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0x28d7 0xffff SW_ILS EFP
ff.ff.fd -> ff.ff.fd 0x28d7 0xa786 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa787 0xffff SW_ILS EFP
ff.ff.fd -> ff.ff.fd 0xa787 0x28d8 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa787 0x28d8 SW_ILS SW_ACC (EFP)
ff.ff.fd -> ff.ff.fd 0x28d7 0xa786 SW_ILS SW_ACC (EFP)
ff.ff.fd -> ff.ff.fd 0x28d7 0xa786 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa787 0x28d8 FC Link Ctl, ACK1
Exchange Fabric Parameters- sent

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

103

FCIP E-Port

Using FCanalyzer in the MDS Fabric: 2 of 2

Cisco.com

ff.ff.fd -> ff.ff.fd 0xa788 0xffff SW_ILS DIA
ff.ff.fd -> ff.ff.fd 0xa788 0x28d9 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa788 0x28d9 SW_ILS SW_ACC (DIA)
ff.ff.fd -> ff.ff.fd 0xa788 0x28d9 FC Link Ctl, ACK1
Domain ID Assignment

ff.ff.fd -> ff.ff.fd 0x28da 0xffff SW_ILS RDI
ff.ff.fd -> ff.ff.fd 0x28da 0xa789 FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0x28da 0xa789 SW_ILS SW_ACC (RDI)
ff.ff.fd -> ff.ff.fd 0x28da 0xa789 FC Link Ctl, ACK1
Request a Specific Domain ID

ff.ff.fd -> ff.ff.fd 0xa78a 0xffff SW_ILS EFP
ff.ff.fd -> ff.ff.fd 0xa78a 0x28db FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa78a 0x28db SW_ILS SW_ACC (EFP)
ff.ff.fd -> ff.ff.fd 0xa78a 0x28db FC Link Ctl, ACK1
Exchange Fabric Parameters from other switch

ff.ff.fd -> ff.ff.fd 0xa78b 0xffff SW_ILS MR
ff.ff.fd -> ff.ff.fd 0xa78b 0x28dc FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0x28dd 0xffff SW_ILS MR
ff.ff.fd -> ff.ff.fd 0x28dd 0xa78c FC Link Ctl, ACK1
ff.ff.fd -> ff.ff.fd 0xa78b 0x28dc SW_ILS SW_ACC (MR)
ff.ff.fd -> ff.ff.fd 0x28dd 0xa78c SW_ILS SW_ACC (MR)
Zone Merge Request and Accept

ff.fc.60 -> ff.fc.69 0xa7a1 0xffff FC-FCS GCAP
ff.fc.69 -> ff.fc.60 0xa7a1 0x28f3 FC Link Ctl, ACK1
ff.fc.60 -> ff.fc.69 0x28f2 0xa7a3 FC Link Ctl, ACK1
Get Management Capabilities

ff.fc.60 -> ff.fc.69 0xa7a2 0xffff FC-FCS GMAL
ff.fc.69 -> ff.fc.60 0xa7a2 0x28f4 FC Link Ctl, ACK1
Get Management Interconnect Elements

ff.fc.60 -> ff.fc.b1 0xa7a5 0xffff dNS GE_PT
ff.fc.b1 -> ff.fc.60 0xa7a5 0x28f5 FC Link Ctl, ACK1
ff.fc.60 -> ff.fc.69 0xa7a4 0xffff dNS GE_PT
ff.fc.69 -> ff.fc.60 0xa7a4 0x28f6 FC Link Ctl, ACK1
Query NS to get Port Type

ff.fc.69 -> ff.fc.60 0xa7a1 0x28f3 FC-FCS (GCAP)
ff.fc.60 -> ff.fc.69 0xa7a1 0x28f3 FC Link Ctl, ACK1
ff.fc.69 -> ff.fc.60 0xa7a2 0x28f4 FC-FCS (GMAL)
ff.fc.60 -> ff.fc.69 0xa7a2 0x28f4 FC Link Ctl, ACK1
ff.fc.69 -> ff.fc.60 0x28f7 0xffff FC-FCS GCAP
ff.fc.69 -> ff.fc.d1 0x28fa 0xffff dNS GE_PT
ff.fc.69 -> ff.fc.d1 0xa7b0 0x2901 FC Link Ctl, ACK1
More management requests from other switch

ff.fc.60 -> b1.00.02 0xa7b2 0xffff FC ELS PLOGI
ff.fc.69 -> d1.00.e8 0x2903 0xffff FC ELS PLOGI
ff.fc.69 -> d1.00.e8 0x2903 0xffff FC Basic Link Svc
ff.fc.69 -> d1.00.e8 0x2904 0xffff FC ELS LOGO
Port Login from MDS switches and then logout

OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

104

Detail Break Down of Frame Capture from FCanalyzer Local

Cisco.com

Fibre Channel
 Exchange Last In: 0
 R_CTL: 0x2(Device_Data/Unsolicited Control)
 Dest Addr: ff.ff.fd
 CS_CTL: 0x00
 Src Addr: ff.ff.fd
 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,
 0... .. = ExgRpd: Exchange Originator
 ..0... .. = SeqRec: Seq Initiator
 ..1... .. = ExgFst: Exchg First
 ...0... .. = ExgLst: NOT exchg last
1... .. = SeqLst: Seq Last
0... .. = Pri: CS_CTL
1... .. = TSI: Transfer Seq Initiative
00... .. = LDF: Last Data Frame - No Info
 (0x000000)
00... .. = A01: no ack required (0x000000)
0... .. = RetSeq: NOT retransmitted sequence
00... .. = AA: ABTS - Cont (0x000000)
0... .. = RelOff: rel offset NOT set
 SEQ_ID: 0x00
 DF_CTL: 0x00
 SEQ_CNT: 0
 OX_ID: 0x2873
 RX_ID: 0xffff
 Parameter: 0x00000000

SW_ILS
 Cmd Code: ELP (0x10)
 Revision: 2
 Flag: 0000
 R_A_TOV: 10000 msecs
 E_D_TOV: 2000 msecs
 Req Eport Name: 20:8a:00:0d:ec:00:ea:40 (00:0d:ec)
 Req Switch Name: 20:00:00:0d:ec:00:ea:40 (00:0d:ec)
 Class F Svc Parameters: (Class F Valid | No X_ID Interlk)
 Max Class F Frame Size: 2112
 Class F Max Concurrent Seq: 1
 Class F E2E Credit: 1
 Class F Max Open Seq: 1
 Class 1 Svc Parameters: (Class 1 Invalid)
 Class 2 Svc Parameters: (Class 2 Valid | Seq Delivery)
 Class 2 Frame Size: 2112
 Class 3 Svc Parameters: (Class 3 Valid | Seq Delivery)
 Class 3 Frame Size: 2112
 ISL Flow Ctrl Mode: R_RDY Flow Ctrl
 Flow Ctrl Param Len: 20
 B2B Credit: 16
 Compatability Param 1: 2112
 Compatability Param 2: 10000
 Compatability Param 3: 2000
 Compatability Param 4: 0

Exchange Link Parameters

Class 2 & # capable

Credits, MTU and Timers

OPT-3052
 9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

105

A Word about IP Routes

Cisco.com

- Depending on circumstances, three possible options:

ip default-gateway

ip default-network

ip route

Directly connected network routes added at configuration of interface

C 10.2.30.0/24 is directly connected, gigabitethernet2/7

ip routes can be statically configure

ip route 192.168.1.0 255.255.255.0 10.1.1.2 interface gigabitethernet 2/1

Default-network places route for 0.0.0.0

ip default-network 10.1.1.0

OPT-3052
 9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

106

debug ips ipstack

Cisco.com

Requires attaching to the IPS-8 Module:

debug ips ipstack

all	Configure ALL IP Stack debugging	in	Configure IP Stack IN debugging
arp	Configure IP Stack ARP debugging	ip	Configure IP Stack IP debugging
arp-detail	Configure verbose IP Stack ARP debugging	ip-detail	Configure verbose IP Stack IP debugging
error	Configure IP Stack Error debugging	route	Configure IP Stack route debugging
eth	Configure IP Stack Ethernet debugging	sdip	Configure IP Stack SDIP debugging
eth-cfg	Configure IP Stack Ethernet Configuration debugging	tcp	Configure IP Stack TCP debugging
eth-chan	Configure IP Stack Ether Channel debugging	udp	Configure IP Stack UDP debugging
eth-detail	Configure verbose debugging of Ethernet layer	udp-detail	Configure verbose IP Stack UDP debugging
icmp	Configure IP Stack ICMP debugging	vrrp	Configure IP Stack VRRP debugging
if	Configure IP Stack interface debugging	vrrp-detail	Configure verbose IP Stack VRRP debugging

Most of these debugs will be used under the guidance of a support technician but are very familiar to Cisco IOS CCIE's


OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

107


Using SPAN And PAA for FCIP

Cisco.com




▲ To PC with Ethereal

Interface FC 1/11 is setup as a SD port so rx & tx traffic will be mirrored to this port



FCIP interfaces are part of a Port-Channel so you will need to SPAN the Port-Channel




OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

108

FCIP and SN5428-2



Cisco.com

SN 5428-2-K9 ge1 ge2 Console Mgmt HA fc1 fc2 fc3 fc4 fc5 fc6 fc7 fc8

Switch Mole Client

FCIP for SAN

Interface	Secondary Interface	IP Address	Netmask
None	None	100.1.1.1	255.255.255.0

Remote Address: 100.1.1.2

Connection Mode:

☐ TCP Server ☒ TCP Client ☐ Raw

Switch Gopher Server

FCIP for SAN

Interface	Secondary Interface	IP Address	Netmask
None	None	100.1.1.2	255.255.255.0

Remote Address: 100.1.1.1

Connection Mode:

☒ TCP Server ☐ TCP Client ☐ Raw

Set Defaults

rxTcpWinSize	1048576	8192-2097152, default 262144
TCP receive window size		
txTcpWinSize	2097152	8192-2097152, default 2097152
TCP transmit window size		
tcpPort	3225	0-65535, default 3225
TCP port to listen on or connect to		
idlePingDelay	15	1-65535, default 15
Idle time (seconds) before sending keep alive packet		
pktTraceMask	65535	0-0xFFFF, 0 sets traces off
Trace mask		
frinHWater	688	1-4294967294, default 688
Maximum frames to FC		
useSelfport	<input checked="" type="checkbox"/>	enable/disable, default is enabled
Use FC B-Port connectivity		
batchTcp	<input checked="" type="checkbox"/>	enable/disable, default is enabled
Batch multiple FC frames in one TCP segment		
compression	<input type="checkbox"/>	enable/disable, default is disabled
Compress data stream		
pacingrate	0	0-999, default is 0
Sends IP packets for remote (mbits/sec)		

TCP Tuning easily configured from GUI

OPT-3052

9909_06_2004_X

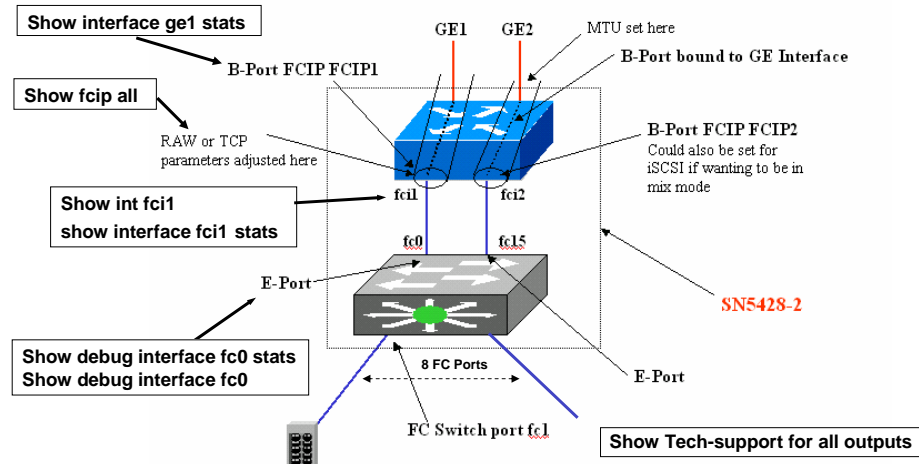
© 2004 Cisco Systems, Inc. All rights reserved.

109

Internal Look at FCIP in SN5428-2

Cisco.com

What the SN5428-2 FCIP configuration looks like on the inside



OPT-3052

9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

110

Information Gathering and Analysis for Troubleshooting FCIP on a SN5428-2

Cisco.com

Display the FCIP B-Port connection that is bound to Interface GE1/2

<show interface fci1>

Stats on the internal Fibre Channel interface that is the B-Port

<show interface fci1 stats> *The same information can be pulled from fci2*

Status on the internal E-Port interface 0 that is the ISL connection through the B-Port

<show debug interface fc0>

Display the internal counters on the E-Ports of FC Switch

<show debug interface fc0 stats> *Same display can be done to show fc15*

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

111

Information Gathering and Analysis for Troubleshooting FCIP on a SN5428-2

Cisco.com

Display GigEthernet counters and usage

<show interface ge1/2 stat>

Configuration parameters available with TCP encapsulation

< fcip fci1/2 destination config>

Condensed view of Fabric Domains

<show fcswitch eport brief>

List of a FC devices logged into this switch, not all switches in fabric

<show fcswitch nameserver>

Show the topology of the Fiber Channel SAN that is connected via FCIP

<show fcswitch eport all>

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

112

SN5428-2 One Stop Look at FCIP

Cisco.com

[Mole-SN5428-2]# show fcip all

Instance Device I/F Network I/F

fcip1 fci1 ge1 100.1.1.1

Instance Description

fcip1 FCIP for SAN

Instance Destination LocalMode IpAddress IsConnected

fcip1 fcip tcpclient 100.1.1.2 TRUE

Instance LinkState

fcip1 UP

fcip1 Trace Status

pkdTracing On, mask 0xffff

mboxTracing On

mboxCmdCount 0

fcip1 Credit Information

Receive Credits 3

Transmit Credits 12

fcip1 Connection Information

idlePingDelay 15

Compression: Off

tcpPort 3225

rxTcpWindowSize 1048576

maxRxTcpWindowSize 1048576

txTcpWindowSize 1048512

txTcpCongestionWindowSize 104380

maxTxTcpWindowSize 1048576

frtn 0

frtnHiWater 688

[Gopher-SN5428-2]# show fcip all

Instance Device I/F Network I/F

fcip2 fci2 ge2 100.1.1.2

Instance Description

fcip2 FCIP for SAN

Instance Destination LocalMode IpAddress IsConnected

fcip2 fcip tcpserver 100.1.1.1 TRUE

Instance LinkState

fcip2 UP

fcip2 Trace Status

pkdTracing On, mask 0xffff

mboxTracing On

mboxCmdCount 0

fcip2 Credit Information

Receive Credits 3

Transmit Credits 12

fcip2 Connection Information

idlePingDelay 15

Compression: Off

tcpPort 3225

rxTcpWindowSize 1048576

maxRxTcpWindowSize 1048576

txTcpWindowSize 1048576

txTcpCongestionWindowSize 1049740

maxTxTcpWindowSize 1048576

frtn 0

frtnHiWater 688

Q and A



Complete Your Online Session Evaluation!

Cisco.com

- WHAT:** Complete an online session evaluation and your name will be entered into a daily drawing
- WHY:** Win fabulous prizes! Give us your feedback!
- WHERE:** Go to the Internet stations located throughout the Convention Center
- HOW:** Winners will be posted on the onsite Networkers Website; four winners per day

OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

115



OPT-3052
9909_06_2004_X

© 2004 Cisco Systems, Inc. All rights reserved.

116