

Contents

- [1 Troubleshooting IP Storage Services](#)
- [2 Overview](#)
 - ◆ [2.1 iSCSI Restrictions](#)
 - ◆ [2.2 iSLB Restrictions](#)
- [3 Initial Troubleshooting Checklist](#)
 - ◆ [3.1 Common Troubleshooting Tools in Fabric Manager](#)
 - ◆ [3.2 Common Troubleshooting Commands in the CLI](#)
- [4 IP Issues](#)
 - ◆ [4.1 Verifying Basic Connectivity](#)
 - ◇ [4.1.1 Verifying Basic Connectivity Using Device Manager](#)
 - ◇ [4.1.2 Verifying Basic Connectivity Using the CLI](#)
 - ◆ [4.2 Verification of Switch Connectivity](#)
 - ◇ [4.2.1 Verifying Switch Connectivity Using Fabric Manager](#)
 - ◇ [4.2.2 Verifying Switch Connectivity Using the CLI](#)
 - ◆ [4.3 Verification of Static IP Routing](#)
 - ◇ [4.3.1 Verifying Static IP Routing Using Device Manager](#)
 - ◇ [4.3.2 Verifying Static IP Routing Using the CLI](#)
 - ◆ [4.4 Cannot Assign IP Address to an Interface](#)
- [5 FCIP Issues](#)
 - ◆ [5.1 One-to-One FCIP Tunnel Creation and Monitoring](#)
 - ◆ [5.2 Configuring the First Switch with the CLI](#)
 - ◆ [5.3 Displaying the Default Values Using the CLI](#)
 - ◆ [5.4 Setting the Static Route for FCIP Tunnels Using the CLI](#)
 - ◆ [5.5 Debugging the Configuration of the Second Switch Using the CLI](#)
 - ◆ [5.6 Displaying the Debug Output from FCIP Tunnel Supervisor Using the CLI](#)
 - ◆ [5.7 Displaying the Debug Output from the FCIP Tunnel IPS Module Using the CLI](#)
 - ◆ [5.8 Verifying the Configuration of the Profiles Using the CLI](#)
 - ◆ [5.9 Verifying the Establishment of the FCIP Tunnel Using the CLI](#)
 - ◆ [5.10 Verifying the Establishment of Default TCP Connections for Each Configured FCIP Tunnel Using the CLI](#)
 - ◆ [5.11 Verifying the Statistics of the ASIC Chip on Each Gigabit Ethernet Port Using the CLI](#)
 - ◆ [5.12 Ethereal Screen Captures of the TCP Connection and FCIP Tunnels](#)
 - ◆ [5.13 One-to-Three FCIP Tunnel Creation and Monitoring](#)
 - ◆ [5.14 Displaying the Configuration of the First Switch Using the CLI](#)
 - ◆ [5.15 Creating the FCIP Interface for the Second Tunnel Using the CLI](#)
 - ◆ [5.16 FCIP Profile Misconfiguration Examples](#)
 - ◆ [5.17 Displaying Incorrect or Nonexistent IP Address for an FCIP Profile Using the CLI](#)
 - ◆ [5.18 Displaying Configuration Errors When Bringing Up a Tunnel on a Selected Port Using the CLI](#)
 - ◆ [5.19 FCIP Interface Misconfiguration Examples](#)
 - ◆ [5.20 Displaying FCIP Misconfiguration Examples Using the CLI](#)
 - ◆ [5.21 Displaying the FCIP Interface as Administratively Shut Down Using the CLI](#)
 - ◆ [5.22 Displaying the Debug Output from the Second Switch Using the CLI](#)
 - ◆ [5.23 Displaying Passive Mode Set on Both Sides of the FCIP Tunnel with the CLI](#)
 - ◆ [5.24 Displaying a Time Stamp Acceptable Difference Failure Using the CLI](#)
 - ◆ [5.25 FCIP Special Frame Tunnel Creation and Monitoring](#)
 - ◆ [5.26 Configuring and Displaying an FCIP Tunnel with Special Frame Using the CLI](#)
 - ◆ [5.27 Special Frame Misconfiguration Example](#)
 - ◆ [5.28 Troubleshooting FCIP Link Flaps](#)
 - ◆ [5.29 Troubleshooting FCIP ISL Link Failures](#)

- ◆ [5.30 Troubleshooting FCIP and Compression](#)
- [6 iSCSI Issues](#)
 - ◆ [6.1 Troubleshooting iSCSI Authentication](#)
 - ◆ [6.2 Displaying iSCSI Authentication Using Fabric Manager](#)
 - ◆ [6.3 Displaying iSCSI Authentication Using the CLI](#)
 - ◆ [6.4 Troubleshooting User Name and Password Configuration](#)
 - ◆ [6.5 Verifying iSCSI User Account Configuration Using Fabric Manager](#)
 - ◆ [6.6 Verifying iSCSI User Account Configuration Using the CLI](#)
 - ◆ [6.7 RADIUS Configuration Troubleshooting](#)
 - ◆ [6.8 Verifying RADIUS Key and Port for Authentication and Accounting](#)
 - ◆ [6.9 Troubleshooting RADIUS Routing Configuration](#)
 - ◆ [6.10 Displaying the Debug Output for RADIUS Authentication Request Routing Using the CLI](#)
 - ◆ [6.11 Troubleshooting Dynamic iSCSI Configuration](#)
 - ◆ [6.12 Checking the Configuration](#)
 - ◆ [6.13 Performing Basic Dynamic iSCSI Troubleshooting](#)
 - ◆ [6.14 Useful Show Commands to Debug Dynamic iSCSI Configuration](#)
 - ◆ [6.15 show iscsi session detail Command Output](#)
 - ◆ [6.16 show iscsi remote-node initiator Command Output](#)
 - ◆ [6.17 show iscsi local-node Command Output](#)
 - ◆ [6.18 show fens data vsan 1 Command Output](#)
 - ◆ [6.19 show flogi database vsan 1 Command Output](#)
 - ◆ [6.20 Virtual Target Access Control](#)
 - ◆ [6.21 Useful Show Commands to Debug Static iSCSI Configuration](#)
 - ◆ [6.22 show iscsi session detail Command Output](#)
 - ◆ [6.23 show iscsi stats Command Output](#)
 - ◆ [6.24 show iscsi stats detail Command Output](#)
 - ◆ [6.25 show fens database Command Output](#)
 - ◆ [6.26 show flogi database Command Output](#)
 - ◆ [6.27 show iscsi remote-node iscsi-session-detail tcp-parameters Command Output](#)
- [7 iSCSI TCP Performance Issues](#)
 - ◆ [7.1 CLI Commands Used to Access Performance Data](#)
 - ◆ [7.2 Understanding TCP Parameters for iSCSI](#)
 - ◆ [7.3 Lab Setup](#)
 - ◆ [7.4 Configuring from the Bottom Switch Using the CLI](#)
 - ◆ [7.5 Verifying Connectivity Between Client and IPS iSCSI Service](#)
 - ◆ [7.6 TCP Parameter Changes](#)
 - ◆ [7.7 Displaying the Gigabit Ethernet Interface](#)
 - ◆ [7.8 Verifying that the Host Is Configured for High MTU or MSS with the CLI](#)
- [8 iSLB Issues](#)
 - ◆ [8.1 iSLB Configuration Not Distributed to All Switches in the Fabric](#)
 - ◆ [8.2 iSCSI Initiator and Virtual Target Configuration Not Distributed](#)
 - ◆ [8.3 iSLB Configuration, Commit, or Merge Failed--"VSAN ID is Not Yet Configured"](#)
 - ◆ [8.4 iSLB Configuration, Commit, or Merge Failed--"Failed to Allocate WWN"](#)
 - ◆ [8.5 iSLB Configuration, Commit, or Merge Failed--"Duplicate WWN Found as..."](#)
 - ◆ [8.6 iSLB Configuration, Commit, or Merge Failed--"Duplicate Node Name"](#)
 - ◆ [8.7 iSLB Configuration Failed--"Pending iSLB CFS Config Has Reached Its Limit..."](#)
 - ◆ [8.8 iSCSI Disable Failed--"Cannot Disable Iscsi - Large Iscsi Config Present..."](#)
 - ◆ [8.9 iSLB Commit Timeout](#)
 - ◆ [8.10 Session Down--"pWWN in Use At Remote Switch"](#)
 - ◆ [8.11 Redirected Session Does Not Come Up](#)
 - ◆ [8.12 iSLB Zones Not Present in Active Zone Set](#)
 - ◆ [8.13 Traffic Description After iSLB Commit or Activation of Zone Set](#)

- ◆ [8.14 VRRP Master Overutilized](#)
- ◆ [8.15 iSLB Zone Set Activation Failed](#)
- ◆ [8.16 iSLB CFS Commit Fails](#)
- ◆ [8.17 Resolving an iSLB Merge Failure](#)

Troubleshooting IP Storage Services

This section describes how to identify and resolve IP storage services problems that might occur in the Cisco MDS 9000 Family products. It includes the following sections:

- Overview
- Initial Troubleshooting Checklist
- IP Issues
- FCIP Issues
- iSCSI Issues
- iSCSI TCP Performance Issues
- iSLB Issues


Overview

Using open-standard, IP-based technology, the Cisco MDS 9000 Family IP Storage (IPS) module enables you to extend the reach of Fibre Channel SANs. The switch can connect separated SAN islands through IP networks using FCIP, and allow IP hosts to access Fibre Channel storage using the iSCSI protocol.

The IPS module allows you to use FCIP and iSCSI features. It supports the full range of features available on other switching modules, including VSANs, security, and traffic management. The IPS module can be used in any Cisco MDS 9000 Family switch and has eight Gigabit Ethernet ports. Each port can run the FCIP and iSCSI protocols simultaneously.

FCIP transports Fibre Channel frames transparently over an IP network between two Cisco MDS 9000 Family switches or other FCIP standards-compliant devices (see Figure 20-1). Using the iSCSI protocol, the IPS module provides IP hosts access to Fibre Channel storage devices. IP host-initiated iSCSI commands are encapsulated in IP, and sent to an MDS 9000 IPS port. There, the commands are routed from the IP network into a Fibre Channel network, and forwarded to the intended target.

Figure 20-1 Connecting MDS 9000 Family Switches Over IP



 **Note:**

- The iSCSI feature is specific to the IPS module and is available in Cisco MDS 9200 Switches or Cisco MDS 9500 Directors.
- The Cisco MDS 9216i switch and the 14/2 Multiprotocol Services (MPS-14/2) module also allow you to use Fibre Channel, FCIP, and iSCSI features. The MPS-14/2 module is available for use in any switch in the Cisco MDS 9200 Series or Cisco MDS 9500 Series.
- The iSCSI feature is not supported on the Cisco Fabric Switch for HP c-Class Bladesystem and Cisco Fabric Switch for IBM BladeCenter.

iSCSI Restrictions

iSCSI has the following limits in Cisco SAN-OS Release 3.0(1) and later:

- Maximum iSCSI sessions on a switch--5000
- Maximum iSCSI sessions per IPS por (not proxy initiator mode)--500
- Maximum iSCSI sessions per IPS port (proxy initiator mode)--500
- Maximum concurrent iSCSI session creations per port

If more iSCSI sessions try to come up simultaneously on a port, the initiator gets a temporary error and then the initiator retries.

- If iSLB CFS is enabled, you must use Device Manager to commit any iSCSI global configuration changes made through Fabric Manager.

iSLB Restrictions

iSLB has the following restrictions in Cisco SAN-OS Release 3.0(1) and later:

- Maximum iSLB initiators in a physical fabric--2000.
- Maximum number of iSCSI sessions per IPS port in either transparent or proxy initiator mode--500.
- Maximum number of switches in a fabric that can have iSLB with CFS enabled--4.

- iSLB can be used without CFS distribution but if iSLB auto-zone feature is used, traffic disruption occurs when any zone set is activated.
- Maximum number of initiators in the pending configuration--200. Before adding more initiators, you must commit the configuration first.
- If there are more than 200 initiators in the running configuration, you must lower the number of initiators to below 200 before disabling iSCSI.
- If IVR and iSLB features are enabled in the same fabric, there should be at least one switch in the fabric that has both of these features enabled. That switch must do any zoning related configuration and activation (for normal zones, IVR zones, or iSLB zones) or there may be traffic disruption in the fabric.
- iSLB VRRP load balancing is based on the number of initiators and not on the number of sessions. If you configure an initiator to see more targets than other initiators (resulting in more sessions on this initiator), you should configure this initiator with a higher load metric.
- iSLB should not be configured with Fabric Manager. Use Device Manager, which supports iSLB with CFS distribution.

Initial Troubleshooting Checklist

Begin troubleshooting IP storage services issues by checking the following issues:

Checklist	Check off
Verify licensing requirements. See the <i>Cisco MDS 9000 Family Fabric Manager Configuration Guide</i> .	
Verify that you are not configuring IPsec with IPv6.	
Verify that auto-zone and CFS distribution are enabled for iSLB.	
If an Ethernet PortChannel is configured between the IPS module and an Ethernet switch, verify that any VRRP load balancing policy on the Ethernet switch is based on source/destination IP address.	
Verify that iSCSI is enabled on the required switches in the fabric, and for the required Gigabit Ethernet interfaces. See the <i>Cisco MDS 9000 Family Fabric Manager Configuration Guide</i> for details.	

Common Troubleshooting Tools in Fabric Manager

Use the following Fabric Manager procedures to access IP interfaces, FCIP, and iSCSI:

- Choose **Switches > Interfaces** to access IP interfaces.
- Choose **ISLs > FCIP** to access FCIP.
- Choose **End Devices > iSCSI** to access iSCSI.

Common Troubleshooting Commands in the CLI

Use the following commands to troubleshoot IP interface, FCIP, and iSCSI issues:


- **show ip**
- **show ips arp**
- **show ipv6 traffic**

- **show ips ipv6**
- **show fcip**
- **show iscsi**

Use the following commands to troubleshoot iSLB issues:

- **show islb initiator [configured]** - Displays all iSLB initiators that have logged into the switch. Use the configured keyword to see all iSLB initiators that have been configured.
- **show islb session** - Verifies that all expected iSCSI sessions are up.
- **show islb merge status** - Displays the status of iSLB merge in the fabric. If the merge is in progress, it shows the identity of the two switches in the fabric. If the merge failed, it shows the reason for the merge failure.
- **show islb status** - Displays whether iSLB CFS distribution is enabled in the fabric and if a CFS session is active.
- **show islb cfs-session status** - Displays the result of the last CFS operation applied from the local switch. If the operation failed, it shows the reason for the failure.
- **show islb vrrp [assignment | interface | summary | vr]** - Shows the iSLB load balancing information with details on the load on each interface and the mapping of initiator to iSCSI interface for every initiator in the fabric.
- **show logging log** - Displays the log file that captures system messages from all modules.

Use the following commands as directed by your customer support representative to further troubleshoot iSLB issues:

 **Note:** To issue commands with the **internal** keyword for troubleshooting purposes, you must have a user account that contains the network-admin role.

- **show ips internal event-history errors** - Displays the errors encountered by the IPS manager.
- **show ips internal event-history msgs** - Displays the message transaction events history for the IPS manager.

Use the following commands to troubleshoot the iSLB initiator and initiator target configuration:

- **show ips internal info islb initiator *node-name*** - Displays the internal data structure for the iSLB initiator.
- **show ips internal event-history iscsi initiator *name*** - Displays the initiator state machine transitions.
- **show ips internal info islb zoneset** - Displays the internal data structure for iSLB zone sets.
- **show ips internal info islb [fc-addr-list | fc-port | fc-port-wwn-tree | hashtable | initiator-mapping | mib-index | nv-pss | scsi-lu-ext]** - Displays the internal data structures for iSLB objects.

Use the following commands to troubleshoot iSLB CFS:

- **show ips internal info islb cfs** - Displays internal data structures for iSLB CFS.
- **show ips internal event-history islb** - Displays the iSLB CFS state machine transitions.

Use the following command to troubleshoot load balancing:

- **show ips internal info islb vrrp [assignment | interface | metric | session]** - Displays the internal data structures for the iSLB load balancing feature.


Use the following **debug** commands to gather more information for iSLB:

- **debug ips error**
- **debug ips islb [config | config-detail | flow | flow-detail]**
- **debug ips islb cfs error**
- **debug ips islb cfs [config | config-detail | flow | flow-detail]**
- **debug ips islb vrrp error**
- **debug ips islb vrrp [flow | flow-detail]**

IP Issues

If you suspect that all or part of your IP connection has failed, you can verify that by performing one or more of the procedures in this section. Using these procedures, you can verify connectivity for IEEE 802.1Q, EtherChannel, and VRRP for iSCSI. This section includes the following topics:

- Verifying Basic Connectivity
- Verification of Switch Connectivity
- Verification of Static IP Routing
- Cannot Assign IP Address to an Interface

 **Note:** If you configure secondary VRRP IPv6 addresses on an IPFC VSAN interface, before a downgrading to a release prior to Cisco Release 3.0(1), you must remove the secondary VRRP IPv6 addresses. This is required only when you configure IPv6 addresses.

Verifying Basic Connectivity

Use the procedures in this section to verify that you have IP connectivity.

Verifying Basic Connectivity Using Device Manager

To verify basic connectivity using Device Manager, follow these steps:

1. Choose **IP > Routes** to verify the static route to the remote device.
 2. Choose **Interface > Ethernet and iSCSI** to verify that the Gigabit Ethernet interface is up.
-

Verifying Basic Connectivity Using the CLI

To verify basic connectivity using the CLI, follow these steps:

1. Use the **ping** or the **ping ipv6** command to perform a basic check of host reachability and network connectivity.

```
switch# ping 11.18.185.121
PING 11.18.185.121 (172.18.185.121): 56 data bytes
64 bytes from 11.18.185.121: icmp_seq=0 ttl=128 time=0.3 ms
64 bytes from 11.18.185.121: icmp_seq=1 ttl=128 time=0.1 ms
64 bytes from 11.18.185.121: icmp_seq=2 ttl=128 time=0.2 ms
64 bytes from 11.18.185.121: icmp_seq=3 ttl=128 time=0.2 ms
64 bytes from 11.18.185.121: icmp_seq=4 ttl=128 time=0.1 ms
64 bytes from 11.18.185.121: icmp_seq=5 ttl=128 time=0.1 ms

--- 11.18.185.121 ping statistics ---
6 packets transmitted, 6 packets received, 0% packet loss
round-trip min/avg/max = 0.1/0.1/0.3 ms
```

2. If the ping fails, use the **traceroute** or the **traceroute ipv6** command to determine where connectivity is failing.

```
switch# traceroute 11.18.185.121
traceroute to 11.18.185.121 (11.18.185.121), 30 hops max, 38 byte packets
 1  11.18.189.129 (11.18.189.129)  0.413 ms  0.257 ms  0.249 ms
 2  11.18.0.33 (11.18.0.33)  0.296 ms  0.260 ms  0.258 ms
 3  11.81.254.69 (11.81.254.69)  0.300 ms  0.273 ms  0.277 ms
 4  * * *
 5  * * *
```

3. Use **show ip route** or the **show ipv6 route** command to verify the static route to the remote device.

```
switch # show ip route

Codes: C - connected, S - static

Default gateway is 11.18.185.97

C 11.18.185.96/27 is directly connected, mgmt0

C 11.18.189.128/26 is directly connected, gigabitethernet4/7
```

4. Use the **clear ips arp** or **clear ipv6 neighbor** command to clear the Address Resolution Protocol (ARP) or neighbor cache to verify that the activity you are viewing is the most current.

```
switch# clear ips arp interface gigabitethernet 4/7

arp clear successful
```

5. Use the **show ips arp** or the **show ips ipv6 neighbors** command to verify the hardware address for the remote device.


```
switch# show ips arp interface gigabitethernet 4/7
Protocol      Address      Age (min)   Hardware Addr  Type   Interface
Internet     172.18.185.97    0    00:d0:01:3b:38:0a  ARPA  GigabitEthernet4/7
Internet     172.18.189.156    0    00:08:02:b3:45:1b  ARPA  GigabitEthernet4/7
```


6. Use the **show interface** command to verify that the Gigabit Ethernet interface is up.

```
GigabitEthernet4/7 is up
  Hardware is GigabitEthernet, address is 0005.3000.9f58
  Internet address is 172.18.189.137/26
  MTU 1500 bytes, BW 1000000 Kbit
  Port mode is IPS
  Speed is 1 Gbps
  Beacon is turned off
  5 minutes input rate 688 bits/sec, 86 bytes/sec, 0 frames/sec
  5 minutes output rate 312 bits/sec, 39 bytes/sec, 0 frames/sec
  156643 packets input, 16859832 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  144401 packets output, 7805631 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors
```

Verification of Switch Connectivity

Use the procedures in this section to verify connectivity to a destination switch.

 **Note:** The FC ID variable used in these procedures is the domain controller address; it is not a duplication of the domain ID.

Verifying Switch Connectivity Using Fabric Manager

To verify connectivity to a destination switch using Fabric Manager, follow these steps:

1. Choose **Fabricxx > VSANxx > Domain Manager** to display the domain ID for the destination switch.
2. Concatenate the domain ID with FFFC to obtain the domain controller address. For example, if the domain ID is 0xda(218), the concatenated ID is 0xffcda.
3. Choose **Tools > Ping...** to verify reachability to the destination switch.

Verifying Switch Connectivity Using the CLI

To verify connectivity to a destination switch using the CLI, follow these steps:

1. Use the **show fcdomain domain-list vsan** command to display the domain ID for the destination switch.

```
switch# show fcdomain domain-list vsan 200

Number of domains: 7

Domain ID WWN
-----
0x01(1) 20:c8:00:05:30:00:59:df [Principal]
0x02(2) 20:c8:00:0b:5f:d5:9f:c1
```

```
0x6f(111) 20:c8:00:05:30:00:60:df
0xda(218) 20:c8:00:05:30:00:87:9f [Local]
0x06(6) 20:c8:00:0b:46:79:f2:41
0x04(4) 20:c8:00:05:30:00:86:5f
0x6a(106) 20:c8:00:05:30:00:f8:e3
```

2. Concatenate the domain ID with FFFC to obtain the domain controller address. For example, if the domain ID is 0xda(218), the concatenated ID is 0xffcda.

3. Use the **fcping** command to verify reachability to the destination switch.

```
switch# '''fcping fcid 0xFFFFCDA vsan 200 '''
28 bytes from 0xFFFFCDA time = 298 usec
28 bytes from 0xFFFFCDA time = 260 usec
28 bytes from 0xFFFFCDA time = 298 usec
28 bytes from 0xFFFFCDA time = 294 usec
28 bytes from 0xFFFFCDA time = 292 usec

5 frames sent, 5 frames received, 0 timeouts
Round-trip min/avg/max = 260/288/298 usec
```

Verification of Static IP Routing

Use the procedures in this section to verify static IP routing.

Verifying Static IP Routing Using Device Manager

1. Choose **IP > Routes** in Device Manager to verify the static IP routes.

Verifying Static IP Routing Using the CLI

To verify static IP routes using the CLI, follow these steps:

1. Use the **show ip route config** or the **show ipv6 route** command to verify the routes configured.

```
switch# show ip route config
Destination      Gateway           Mask Metric      Interface
  default        172.17.8.1       0.0.0.0      0          mgmt0
  11.2.36.0       11.3.36.1       255.255.252.0 0
  11.2.56.0       11.3.56.1       255.255.252.0 0
  11.3.36.0       0.0.0.0         255.255.252.0 0 GigabitEthernet8/7
  11.3.56.0       0.0.0.0         255.255.252.0 0 GigabitEthernet8/8
  172.17.8.0      0.0.0.0         255.255.255.0 0          mgmt0
```

2. Use the **show ip route** or the **show ipv6 route** command to verify that the IP routes are still present.

```
switch# show ip route

Codes: C - connected, S - static

Default gateway is 172.17.8.1

C 172.17.8.0/24 is directly connected, mgmt0
S 11.2.36.0/22 via 11.3.36.1, gigabitethernet8/7
C 11.3.36.0/22 is directly connected, gigabitethernet8/7
C 11.3.56.0/22 is directly connected, gigabitethernet8/8
S 11.2.56.0/22 via 11.3.56.1, gigabitethernet8/8
```

Cannot Assign IP Address to an Interface

You may encounter a problem when assigning an IP address to an interface. If that IP address is already in use by another interface (for example, a remote VRRP interface), you may see the following message:


Invalid configuration: this IP address overlaps with another interface in network

You can recover from this problem by using the **shutdown** CLI command on that VRRP interface, programming the IP address, and then using the **no shutdown** CLI command on that VRRP interface.

FCIP Issues

This section contains information on troubleshooting FCIP tunnels with and without special frames and includes the following topics:

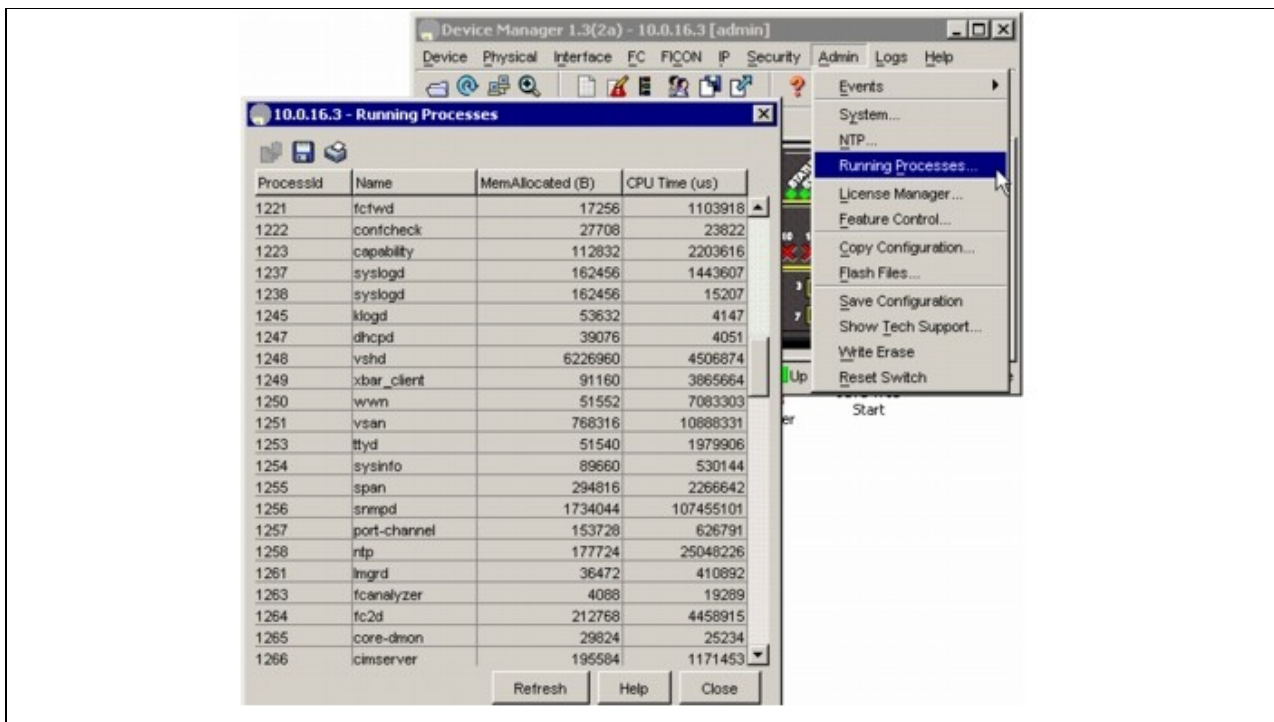
- One-to-One FCIP Tunnel Creation and Monitoring
- One-to-Three FCIP Tunnel Creation and Monitoring
- FCIP Profile Misconfiguration Examples
- FCIP Interface Misconfiguration Examples
- FCIP Special Frame Tunnel Creation and Monitoring
- Special Frame Misconfiguration Example
- Troubleshooting FCIP Link Flaps
- Troubleshooting FCIP ISL Link Failures
- Troubleshooting FCIP and Compression

 **Note:** FCIP tape acceleration does not work if the FCIP port is part of a PortChannel or if there are multiple paths between the initiator and the target port. Such a configuration might cause SCSI discovery failure or broken write or read operations.

One-to-One FCIP Tunnel Creation and Monitoring

This section describes the configuration for one-to-one FCIP tunnel with FCIP debug activated (MDS2) and without debug activated (MDS1). Figure 20-2 shows the one-to-one topology used for configuration.

Figure 20-2 One-to-One Topology



Configuring the First Switch with the CLI

To configure the first switch using the CLI, follow these steps:

1. Enter configuration mode.
2. Set the interface.

```
MDS1(config)# interface gigabitethernet 2/8
```

Set the IP address.

```
MDS1(config-if)# ip address 10.10.10.2 255.255.255.0
```

4. Enter **no shutdown**.

```
MDS1(config-if)# no shutdown
```

5. Enter the profile number and profile mode.

```
MDS1(config)# fcip profile 28
```

The profile number can be any number between 1 to 255

6. Enter the IP address of the local Gigabit Ethernet port that will be the endpoint of the FCIP tunnel.

```
MDS1(config-profile)# ip address 10.10.10.2
```

7. Exit profile mode.

```
MDS1(config-profile)# exit
```

8. Set the FCIP interface and enter interface mode.

```
MDS1(config)# interface fcip 28
```

The interface FCIP can be any number between 1 to 255 and does not need to be the same as the profile number. In this example the same number is used for simplicity.

9. Specify a profile to use.

```
MDS1(config-if)# use-profile 28
```

The FCIP interface will use the local FCIP profile. The FCIP profile binds the FCIP interface to the physical Gigabit Ethernet port and configures the TCP settings used by the FCIP interface.

```
MDS1(config-if)# peer-info ipaddr 10.10.11.2
```

The IP address in this example indicates the remote endpoint IP address of the FCIP tunnel.

```
MDS1(config-if)# no shutdown
MDS1(config-if)# end
```

Displaying the Default Values Using the CLI

The following example displays the default values from the **show running-config** command:

```
MDS1# show running-config

Building Configuration ...
  fcip profile 28
ip address 10.10.10.2
port 3225
tcp keepalive-timeout 60
tcp max-retransmissions 4
tcp pmtu-enable reset-timeout 3600
tcp initial-retransmit-time 100
tcp window-size 64
vsan database
vsan 2 name grumpy_02

interface fcip28
no shutdown
use-profile 28
peer-info ipaddr 10.10.11.2

ip route 10.10.11.0 255.255.255.0 10.10.10.1
```

Setting the Static Route for FCIP Tunnels Using the CLI

The static route must be set for FCIP tunnels. This route could also be `ip route 10.10.11.0 255.255.255.0 interface gigabitethernet 2/8`.

```
ips heartbeat
ips hapreset
ips boot
  interface GigabitEthernet2/8
ip address 10.10.10.2 255.255.255.0
```

(This is the IP address used by the FCIP profile.)

```
no shutdown
```

Debugging the Configuration of the Second Switch Using the CLI

The following example shows the configuration of a switch (MDS2) with debug mode activated. To activate debug mode for this situation, run the **debug ips flow fcip** command on a separate terminal.

```
MDS2(config)# fcip profile 28
Mar 10 21:41:04 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id 32222)
Mar 10 21:41:04 ips: Create Entity 28
Mar 10 21:41:04 ips: entity28: add to config pss

MDS2(config-profile)# ip address 10.10.11.2
Mar 10 21:41:15 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id 32258)
Mar 10 21:41:15 ips: entity28: IP address changed to 10.10.11.2
Mar 10 21:41:15 ips: entity28: IP 10.10.11.2 configured for interface GigabitEthernet2/8
Mar 10 21:41:15 ips: entity28: Apply the entity config and save to config pss
Mar 10 21:41:15 ips: entity28: add to config pss

MDS2(config-profile)# exit

MDS2(config)# interface fcip 28
Mar 10 21:41:46 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id 32358)
Mar 10 21:41:46 ips: Verified FCIP28 Create:0
Mar 10 21:41:46 ips: FCIP28: Verified Create:0
Mar 10 21:41:46 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id 32360)
Mar 10 21:41:46 ips: FCIP28: Creating FCIP tunnel
Mar 10 21:41:46 ips: FCIP28: add to admin pss
Mar 10 21:41:46 ips: FCIP28: add to run-time pss
Mar 10 21:41:46 ips: FCIP28: log: 0 phy: 0 state: 0 syslog: 0

MDS2(config-if)# use-profile 28
Mar 10 21:42:23 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id 32480)
Mar 10 21:42:23 ips: FCIP28: Process tunnel configuration event
Mar 10 21:42:23 ips: FCIP28: Change Entity-id from 0 to 28
Mar 10 21:42:23 ips: FCIP: Optimal IF lookup for GigabitEthernet2/8 is GigabitEthernet2/8
Mar 10 21:42:23 ips: FCIP28: bind with GigabitEthernet2/8 (phy GigabitEthernet2/8)
Mar 10 21:42:23 ips: FCIP28: Queueing bind tunnel to src if event to tunnel FSM resource: 0
Mar 10 21:42:23 ips: Locked fcip_if_fsm for MTS_OPC_IPS_FCIP_CMI_REQUEST(msg id 32480)
Mar 10 21:42:23 ips: FCIP28: Send bind for GigabitEthernet2/8 to PM (phy GigabitEthernet2/8)
Mar 10 21:42:23 ips: FCIP28: add to run-time pss
Mar 10 21:42:23 ips: FCIP28: log: 2087000 phy: 2087000 state: 0 syslog: 0
Mar 10 21:42:23 ips: Dequeued mts msg MTS_OPC_IPS_CFG_FCIP_IF(mts opc 1905, msg id 7304)
Mar 10 21:42:23 ips: Hndlr MTS_OPC_IPS_CFG_FCIP_IF (mts_opc 1905 msg_id 7304)
Mar 10 21:42:23 ips: FCIP28: Got a tunnel param pull request from LC
Mar 10 21:42:23 ips: Added to pending queue event-id [29] event-cat [2]
Mar 10 21:42:23 ips: FCIP28: Queueing Process a Pull Request event to Pending queue resource: 0
Mar 10 21:42:23 ips: Dequeued mts msg MTS_OPC_PM_FCIP_BIND(mts opc 335, msg id 32495)
Mar 10 21:42:23 ips: Hndlr MTS_OPC_PM_FCIP_BIND (mts_opc 335 msg_id 32495)
Mar 10 21:42:23 ips: FCIP28: Success received from PM for bind to GigabitEthernet2/8 (phy
```

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
GigabitEthernet2/8)
Mar 10 21:42:23 ips: FCIP28: Bind-resp event processing bind...
Mar 10 21:42:23 ips: FCIP28: add to run-time pss
Mar 10 21:42:23 ips: FCIP28: log: 2087000 phy: 2087000 state: 1 syslog: 0
Mar 10 21:42:23 ips: FCIP28: Last reference....
Mar 10 21:42:23 ips: FCIP28: Update the tunnel param and save to PSS
Mar 10 21:42:23 ips: FCIP28: add to admin pss
Mar 10 21:42:23 ips: FCIP28: add to run-time pss
Mar 10 21:42:23 ips: FCIP28: log: 2087000 phy: 2087000 state: 1 syslog: 0
Mar 10 21:42:23 ips: Unlocked fcip_if_fsm for MTS_OPC_IPS_FCIP_CMI_REQUEST(msg id 32480)
Mar 10 21:42:23 ips: Dequeued pending queue msg event_id [29] cat [2]
Mar 10 21:42:23 ips: (ips_demux) Mts Opcode is 1905, id is 7304
Mar 10 21:42:23 ips: FCIP28: Processing Pull Config Request
Mar 10 21:42:23 ips: FCIP28: Bound to entity 28 port: 3225 ip: 10.10.11.2

MDS2(config-if)# peer-info ipaddr 10.10.10.2
Mar 10 21:43:01 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_CMI_REQUEST(mts opc 3321, msg id
32616)
Mar 10 21:43:01 ips: FCIP28: Process tunnel configuration event
Mar 10 21:43:01 ips: FCIP28: Change Peer IP from 0.0.0.0 to 10.10.10.2 and port from 3225
to 3225
Mar 10 21:43:01 ips: FCIP28: Queueing Set tunnel param event to tunnel FSM resource: 0
Mar 10 21:43:01 ips: Locked fcip_if_fsm for MTS_OPC_IPS_FCIP_CMI_REQUEST(msg id 32616)
Mar 10 21:43:01 ips: FCIP28: Send tunnel params to LC to DPP: 7
Mar 10 21:43:01 ips: Dequeued mts msg MTS_OPC_IPS_FCIP_SET_LC_TUNNEL_PARAM(mts opc 1897,
msg id 7358)
Mar 10 21:43:01 ips: Hndlr MTS_OPC_IPS_FCIP_SET_LC_TUNNEL_PARAM (mts_opc 1897 msg_id 7358)
Mar 10 21:43:01 ips: In handler : Received resp code: 0
Mar 10 21:43:01 ips: FCIP28: Received the tunnel params from LC
Mar 10 21:43:01 ips: FCIP28: Update the tunnel param and save to PSS
Mar 10 21:43:01 ips: FCIP28: add to admin pss
Mar 10 21:43:01 ips: FCIP28: add to run-time pss
Mar 10 21:43:01 ips: FCIP28: log: 2087000 phy: 2087000 state: 1 syslog: 0
Mar 10 21:43:01 ips: Unlocked fcip_if_fsm for MTS_OPC_IPS_FCIP_CMI_REQUEST(msg id 32616)

MDS2(config-if)#

MDS2(config-if)# no shutdown

MDS2(config-if)# Mar 10 21:43:32 ips: Dequeued mts msg
MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc 3114, msg id 32737)
Mar 10 21:43:32 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
32737)
Mar 10 21:43:32 ips: Dequeued mts msg MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc
3114, msg id 32778)
Mar 10 21:43:32 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
32778)
Mar 10 21:43:32 ips: Dequeued mts msg MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc
3114, msg id 32783)
Mar 10 21:43:32 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
32783)
```

Displaying the Debug Output from FCIP Tunnel Supervisor Using the CLI

The following example shows the debug output from the supervisor of the FCIP tunnel:

```
MDS2(config)# interface fcip 28
MDS2(config-if)# no shutdown
MDS2(config-if)# Mar 10 22:59:46 ips: fu_priority_select: - setting fd[3] for select call
```

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
- found data in FU_PSEL_Q_CAT_MTS queue, fd(3), usr_q_info(1)
Mar 10 22:59:46 ips: fu_priority_select_select_queue: round credit(0)
Mar 10 22:59:46 ips: curr_q - FU_PSEL_Q_CAT_CQ, usr_q_info(3), priority(4), credit(0),
empty
Mar 10 22:59:46 ips: Starting a new round
Mar 10 22:59:46 ips: fu_priority_select: returning FU_PSEL_Q_CAT_MTS queue, fd(3),
usr_q_info(1)
Mar 10 22:59:46 ips: Dequeued mts msg MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc
3114, msg id 47540)
Mar 10 22:59:46 ips: ips_mts_hdlr_pm_logical_port_state_change_range:
Mar 10 22:59:46 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
47540)
Mar 10 22:59:46 ips: fu_fsm_execute_all: match_msg_id(0), log_already_open(0)
Mar 10 22:59:46 ips: fu_fsm_execute_all: null fsm_event_list
Mar 10 22:59:46 ips: fu_fsm_engine: mts msg
MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(msg_id 47540) dropped
Mar 10 22:59:46 ips: fu_priority_select: - setting fd[3] for select call - found data in
FU_PSEL_Q_CAT_MTS queue, fd(3), usr_q_info(1)
Mar 10 22:59:46 ips: fu_priority_select_select_queue: round credit(6)
Mar 10 22:59:46 ips: curr_q - FU_PSEL_Q_CAT_CQ, usr_q_info(3), priority(4), credit(3),
empty
Mar 10 22:59:46 ips: fu_priority_select: returning FU_PSEL_Q_CAT_MTS queue, fd(3),
usr_q_info(1)
Mar 10 22:59:46 ips: Dequeued mts msg MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc
3114, msg id 47589)
Mar 10 22:59:46 ips: ips_mts_hdlr_pm_logical_port_state_change_range:
Mar 10 22:59:46 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
47589)
Mar 10 22:59:46 ips: fu_fsm_execute_all: match_msg_id(0), log_already_open(0)
Mar 10 22:59:46 ips: fu_fsm_execute_all: null fsm_event_list
Mar 10 22:59:46 ips: fu_fsm_engine: mts msg
MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(msg_id 47589) dropped
Mar 10 22:59:46 ips: fu_priority_select: - setting fd[3] for select call - found data in
FU_PSEL_Q_CAT_MTS queue, fd(3), usr_q_info(1)
Mar 10 22:59:46 ips: fu_priority_select_select_queue: round credit(4)
Mar 10 22:59:46 ips: curr_q - FU_PSEL_Q_CAT_CQ, usr_q_info(3), priority(4), credit(2),
empty
Mar 10 22:59:46 ips: fu_priority_select: returning FU_PSEL_Q_CAT_MTS queue, fd(3),
usr_q_info(1)
Mar 10 22:59:46 ips: Dequeued mts msg MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(mts opc
3114, msg id 47602)
Mar 10 22:59:46 ips: ips_mts_hdlr_pm_logical_port_state_change_range:
Mar 10 22:59:46 ips: Hndlr MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE (mts_opc 3114 msg_id
47602)
Mar 10 22:59:46 ips: fu_fsm_execute_all: match_msg_id(0), log_already_open(0)
Mar 10 22:59:46 ips: fu_fsm_execute_all: null fsm_event_list
Mar 10 22:59:46 ips: fu_fsm_engine: mts msg
MTS_OPC_PM_LOGICAL_PORT_STATE_CHANGE_RANGE(msg_id 47602) dropped
```

Displaying the Debug Output from the FCIP Tunnel IPS Module Using the CLI

The following example shows the debug output from the IPS module of the FCIP tunnel:

```
MDS2# attach module 2
module-2# debug ips fcip fsm port 8
(This is the Gigabit Ethernet port 2/8.)

Mar 13 19:18:19 port8: 2700:FCIP28: Received new TCP connection from peer:
10.10.10.2:65455
Mar 13 19:18:19 port8: 2701:FCIP: (fcip_de_create): DE = 0xdc02ca40
Mar 13 19:18:19 port8: 2702:FCIP28: Create a DE 0xdc02ca40 for this tunnel
Mar 13 19:18:19 port8: 2703:FCIP28: Bind the DE 0xdc02ca40 [1] to tunnel LEP 0x801ebac0
Mar 13 19:18:19 port8: 2704:FCIP28: Bind DE 1 to TCP-hdl 0xdc489800
```


Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
Mar 13 19:18:19 port8: 2705:FCIP28: Bind DE 1 to eport 0x801eaaa0
Mar 13 19:18:19 port8: 2706:FCIP28: bind de 1 in eport 0x801eaaa0, hash = 1 num-conn: 2
Mar 13 19:18:19 port8: 2707:FCIP28: Received new TCP connection from peer:
10.10.10.2:65453
Mar 13 19:18:19 port8: 2708:FCIP: (fcip_de_create): DE = 0xdc02cb40
Mar 13 19:18:19 port8: 2709:FCIP28: Create a DE 0xdc02cb40 for this tunnel
Mar 13 19:18:19 port8: 2710:FCIP28: Bind the DE 0xdc02cb40 [2] to tunnel LEP 0x801ebac0
Mar 13 19:18:19 port8: 2711:FCIP28: Bind DE 2 to TCP-hdl 0xdc488800
Mar 13 19:18:19 port8: 2712:FCIP28: Bind DE 2 to eport 0x801eaaa0
Mar 13 19:18:19 port8: 2713:FCIP28: bind de 2 in eport 0x801eaaa0, hash = 2 num-conn: 2
Mar 13 19:18:19 port8: 2714:FCIP28: Send LINK UP to SUP
Mar 13 19:18:20 port8: 2715:FCIP28: *** Received eisl frame in E mode
Mar 13 19:18:20 port8: 2716:FCIP28: SUP-> Set trunk mode: 2
Mar 13 19:18:20 port8: 2717:FCIP28: Change the operational mode to TRUNK
Mar 13 19:18:20 port8: 2718:FCIP28: Tunnel bringup debounce timer callback, try to bring
up tunnel
Mar 13 19:18:20 port8: 2719:FCIP28: Tunnel is already in oper UP state, don't try to
bring up again...
```

Verifying the Configuration of the Profiles Using the CLI

Use the **show fcip profile** command to verify that the configuration of the profiles are correct. The IP address and TCP port are the ports to listen on, and both are adjustable in the FCIP profile. The following example displays all the default values that are adjustable while configuring the FCIP profile.

```
MDS1# show fcip profile
```

```
-----
ProfileId      Ipaddr        TcpPort
-----
28             10.10.10.2    3225
```

```
MDS1# show fcip profile 28
```

```
FCIP Profile 28
```

```
    Listen Port is 3225
TCP parameters
    SACK is disabled
    PMTU discover is enabled, reset timeout is 3600 sec
    Keep alive is 60 sec
    Minimum retransmission timeout is 100 ms
    Maximum number of re-transmissions is 4
    Advertised window size is 64 KB
```

Verifying the Establishment of the FCIP Tunnel Using the CLI

Use the **show interface fcip** command to verify that the FCIP tunnel is established and that traffic is passing through.

```
MDS1# show interface fcip 28
```

```
FCIP28 is trunking
```

```
    Hardware is GigabitEthernet
    Port WWN is 20:5e:00:05:30:00:59:de
    Peer port WWN is 20:5e:00:0b:5f:d5:9f:c0
    Admin port mode is auto, trunk mode is on
    Port mode is TE
```

(The FCIP tunnel will be either E (ISL or TE (EISL) passing through multiple VSANs.)

```
vsan is 1
```

```
    Trunk vsans (allowed active) (1-2)
    Trunk vsans (operational) (1-2)
```

Cisco_MDS_SanOS_Troubleshooting_Guide_-_Troubleshooting_IP_Storage_Services

```
Trunk vsans (up) (1-2)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
Using Profile id 28 (interface GigabitEthernet2/8)
```

(This is the FCIP profile and the Gigabit Ethernet being used by the FCIP tunnel.)

Peer Information

```
Peer Internet address is 10.10.11.2 and port is 3225
```

(This is the remote endpoint's IP address and listening port.)

Special Frame is disabled

(The special frame for verification of a remote MDS is not being used.)

Maximum number of TCPconnections is 2

(The default is 2 TCP connections being used, one for class F and the other for class 2 and 3.)

Time Stamp is disabled

(The time stamp can be activated under the FCIP interface.)

B-port mode disabled

TCP Connection Information

```
2 Active TCP connections
```

```
Control connection: Local 10.10.10.2:3225, Remote 10.10.11.2:65519
```

(The above is class F traffic.)

```
Data connection: Local 10.10.10.2:3225, Remote 10.10.11.2:65521
```

(This is class 2,3 traffic.)

```
6 Attempts for active connections, 3 close of connections
```

TCP Parameters

```
Path MTU 1500 bytes
```

```
Current retransmission timeout is 100 ms <<< Default, adjusted under
```

```
Round trip time: Smoothed 10 ms, Variance: 5
```

(This is the calculated round trip time of the FCIP tunnel. Large round trip times will require in

```
Advertized window: Current: 64 KB, Maximum: 64 KB, Scale: 1
```

(This is the local advertised TCP window size, and the default is 64 KB.)

```
Peer receive window: Current: 64 KB, Maximum: 64 KB, Scale: 1
```

(This is the remote endpoint advertised TCP window size.)

```
Congestion window: Current: 2 KB
```

(This is the minimum window size used during congestion, and it is not configurable.)

```
5 minutes input rate 136 bits/sec, 17 bytes/sec, 0 frames/sec
```

```
5 minutes output rate 136 bits/sec, 17 bytes/sec, 0 frames/sec
```

```
2288 frames input, 211504 bytes
```

```
2288 Class F frames input, 211504 bytes
```

```
0 Class 2/3 frames input, 0 bytes
```

```
0 Error frames
```

```
2288 frames output, 211520 bytes
```

```
2288 Class F frames output, 211520 bytes
```

```
0 Class 2/3 frames output, 0 bytes
```

```
0 Error frames 0 reassign frames
```

```
MDS1# show interface fcip 28 brief
```

```
-----
Interface  Vsan   Admin  Admin  Status      Oper  Profile  Port-channel
           Mode   Mode   Trunk                                     Mode
           Mode
-----
```

```
fcip28      1      auto   on      trunking    TE    28      --
```

```
MDS1# show interface fcip 28 counters brief
```

```
-----
Interface                Input (rate is 5 min avg)      Output (rate is 5 min avg)
-----
                          Rate      Total          Rate      Total
                          Mbits/s  Frames        Mbits/s  Frames
-----
fcip28                    18         0              18         0
-----
```

(These are the frames that averaged over 5 minutes and the total count of frames since the last **clear counters** command was issued, or since the last tunnel up.)

Verifying the Establishment of Default TCP Connections for Each Configured FCIP Tunnel Using the CLI

Verify that two default TCP connections are established for each FCIP tunnel configured, one for control traffic and one for data traffic.

```
MDS1# show ips stats tcp interface gigabitethernet 2/8
```

```
TCP Statistics for port GigabitEthernet2/8
```

```
Connection Stats
```

```
6 active openings, 8 accepts
```

```
6 failed attempts, 0 reset received, 8 established
```

```
Segment stats
```

```
295930 received, 1131824 sent, 109 retransmitted
```

(Excessive retransmits indicate possible core drops and/or that the TCP window size should be adjusted.)

```
0 bad segments received, 0 reset sent
```

```
TCP Active Connections
```

```
Local Address      Remote Address      State      Send-Q  Recv-Q
10.10.10.2:3225    10.10.11.2:65519   ESTABLISH  0       0
```

(This is used for F control traffic only.)

```
10.10.10.2:3225    10.10.11.2:65521   ESTABLISH  87568   0
```

(Send-Q increasing during read-only test.)

```
10.10.10.2:3225    0.0.0.0:0          LISTEN     0       0
```

(The TCP listen port is ready for new TCP connections.)

You can use the following command to verify that traffic is incrementing on the Gigabit Ethernet port of the FCIP tunnel:

```
MDS1# show ips stats mac interface gigabitethernet 2/8
```

```
Ethernet MAC statistics for port GigabitEthernet2/8
```

```
Hardware Transmit Counters
```

```
1074898 frame 1095772436 bytes
```

```
0 collisions, 0 late collisions, 0 excess collisions
```

```
0 bad frames, 0 FCS error, 0 abort, 0 runt, 0 oversize
```

```
Hardware Receive Counters
```

```
33488196 bytes, 298392 frames, 277 multicasts, 16423 broadcasts
```

```
0 bad, 0 runt, 0 CRC error, 0 length error
```

```
0 code error, 0 align error, 0 oversize error
```

```
Software Counters
```

```
298392 received frames, 1074898 transmit frames
```

```
0 frames soft queued, 0 current queue, 0 max queue
```

```
0 dropped, 0 low memory
```

Verifying the Establishment of Default TCP Connections for Each Configured FCIP Tunnel Using the CLI

Verifying the Statistics of the ASIC Chip on Each Gigabit Ethernet Port Using the CLI

Traffic statistics can be verified on the internal ASIC chip on each Gigabit Ethernet port:

```
MDS1# show ips stats flamingo interface gigabitethernet 2/8

Flamingo ASIC Statistics for port GigabitEthernet2/8
  Hardware Egress Counters
    2312 Good, 0 bad protocol, 0 bad header cksum, 0 bad FC CRC
  (Good frames and CRC error frames can be monitored.)

  Hardware Ingress Counters
  (Verify good increments on the active tunnel.)

    2312 Good, 0 protocol error, 0 header checksum error
    0 FC CRC error, 0 iSCSI CRC error, 0 parity error
  Software Egress Counters
    2312 good frames, 0 bad header cksum, 0 bad FIFO SOP
    0 parity error, 0 FC CRC error, 0 timestamp expired error
    0 unregistered port index, 0 unknown internal type
    0 RDL, 0 RDL too big RDL, 0 TDL ttl_1
    3957292257 idle poll count, 0 loopback, 0 FCC PQ, 0 FCC EQ
    Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
  Software Ingress Counters
    2312 Good frames, 0 header cksum error, 0 FC CRC error
    0 iSCSI CRC error, 0 descriptor SOP error, 0 parity error
    0 frames soft queued, 0 current Q, 0 max Q, 0 low memory
    0 out of memory drop, 0 queue full drop
    0 RDL, 0 too big RDL drop
    Flow Control: 0 [0], 0 [1], 0 [2], 0 [3]
```

Ethereal Screen Captures of the TCP Connection and FCIP Tunnels

Figure 20-3, Figure 20-4, and Figure 20-5 are screen captures taken with Ethereal of TCP connections being established, and of FCIP tunnels. Note that FCIP tunnel activation is the same as an FC EISL becoming active (such as ELP, ESC, and EFP). The following traces were captured after configuration on both MDS 9000 Family switches, and the last **no shutdown** was entered on switch MDS1. All settings are default (for example, SACK is disabled, and the TCP window is set to 64 K).

Figure 20-3 First Capture of TCP Connection




Figure 20-4 shows more of the trace, with frame 13 being the first FCIP frame. This frame carries the FC Standard ELP.

Figure 20-4 Second Capture of TCP Connection

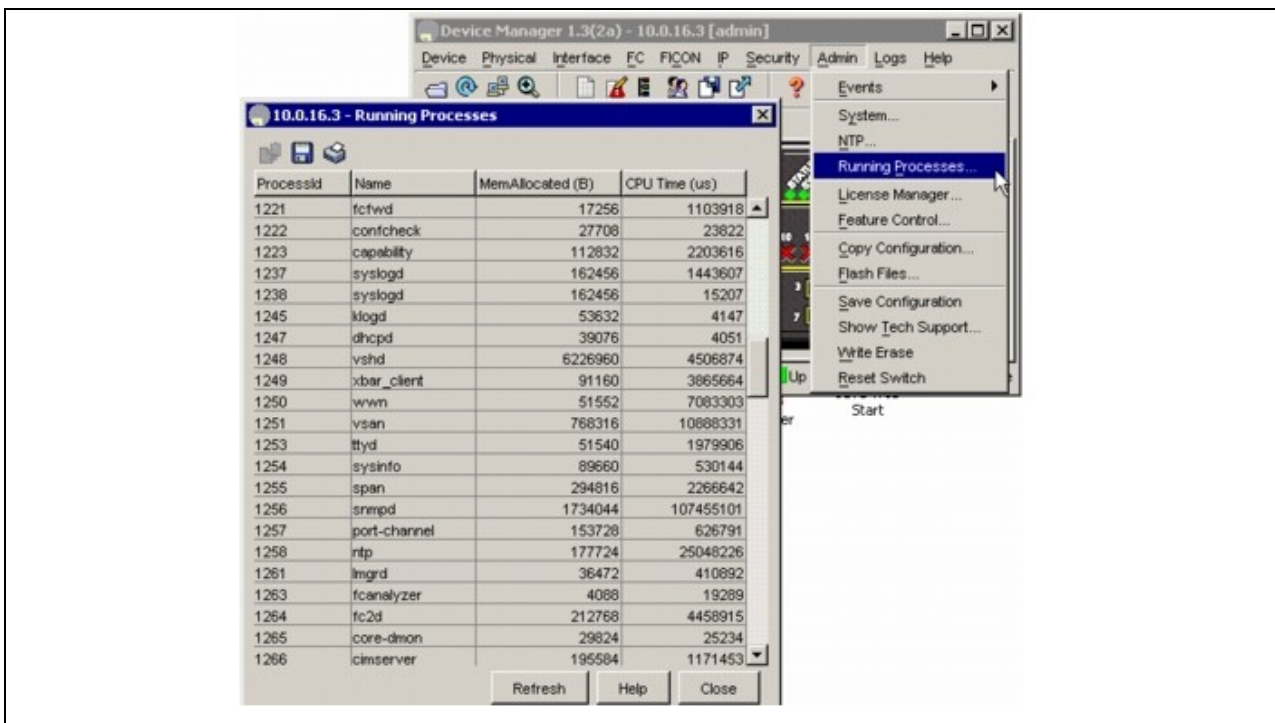



Figure 20-5 shows the FC portion of the EISL initialization over the FCIP tunnel.


Figure 20-5 Third Capture of TCP Connection



One-to-Three FCIP Tunnel Creation and Monitoring

Figure 20-6 shows the configuration of switch MDS1 for three tunnels from one Gigabit Ethernet port.

Figure 20-6 MDS1 Configured for Three FCIP Tunnels



Displaying the Configuration of the First Switch Using the CLI

The following example shows the configuration of switch MDS1 for three tunnels from one Gigabit Ethernet port:

```
MDS1(config)# fcip profile 21
MDS1(config-profile)# ip address 10.10.10.2
MDS1(config-profile)# exit
MDS1(config)# interface fcip 21
MDS1(config-if)# use-profile 21
MDS1(config-if)# peer-info ipaddr 10.10.11.2
MDS1(config-if)# no shutdown
MDS1(config-if)# exit
MDS1(config)# ip route 10.10.11.0 255.255.255.0 10.10.10.1
MDS1(config)# ip route 10.10.11.0 255.255.255.0 interface gigabitethernet 2/1
```

Creating the FCIP Interface for the Second Tunnel Using the CLI

Now the interface FCIP is created for the second tunnel. The same FCIP profile is used for this example. A separate FCIP profile can be used for each FCIP interface if desired.

```
MDS1(config-if)#
MDS1(config-if)# interface fcip 23
MDS1(config-if)# use-profile 21
MDS1(config-if)# peer-info ipaddr 10.10.8.2
MDS1(config-if)# no shutdown
MDS1(config-if)# exit
MDS1(config)#
```

Now the FCIP interface is created for the third tunnel.

```
MDS1(config)# interface fcip 28
MDS1(config-if)# use-profile 21
MDS1(config-if)# peer-info ipaddr 10.10.7.2
MDS1(config-if)# no shut
MDS1(config-if)# end
MDS1(config)#
```

FCIP Profile Misconfiguration Examples

The examples in this section show FCIP profile misconfigurations.

Displaying Incorrect or Nonexistent IP Address for an FCIP Profile Using the CLI

```
MDS22(config)# fcip profile 21
MDS22(config-profile)# ip addr 1.1.1.1
MDS22(config-profile)# ip addr 34.34.34.34
MDS22(config-profile)# exit
MDS22(config)# exit
MDS22# show fcip profile 21
FCIP Profile 21
```

```
Internet Address is 34.34.34.34
```

(In this line, the interface Gigabit Ethernet port is not shown. This means the IP address is not

```
Listen Port is 3225
```

```
TCP parameters
```

```
SACK is disabled
```

```
PMTU discover is enabled, reset timeout is 3600 sec
```

```
Keep alive is 60 sec
```

```
Minimum retransmission timeout is 300 ms
```

```
Maximum number of re-transmissions is 4
```

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
Advertised window size is 64 KB
MDS22# config t
Enter configuration commands, one per line. End with CNTL/Z.
MDS22(config)# interface gigabitethernet 2/5
MDS22(config-if)# ip addr 34.34.34.34 255.255.255.0
MDS22(config-if)# no shutdown
MDS22(config-if)# end
MDS22# show fcip profile 34
error: fcip profile not found
MDS22# show fcip profile 21
FCIP Profile 21
    Internet Address is 34.34.34.34 (interface GigabitEthernet2/5)
(In this line, the Gigabit Ethernet port is now shown and the FCIP profile is bound to a physical

Listen Port is 3225
    TCP parameters
        SACK is disabled
        PMTU discover is enabled, reset timeout is 3600 sec
        Keep alive is 60 sec
        Minimum retransmission timeout is 300 ms
        Maximum number of re-transmissions is 4
        Advertised window size is 64 KBThe following example shows a configuration error
when using multiple FCIP profiles on one physical Gigabit Ethernet port.
MDS2(config)# fcip profile 21
MDS2(config-profile)# ip address 10.10.11.2
error: fcip another profile exists with same port & ip
(Multiple FCIP profiles can be used on one physical Gigabit Ethernet port, but each profile must h

MDS2(config-profile)# port 32
(Change the TCP listen port on the profile. The default is 3225.)

MDS2(config-profile)# ip address 10.10.11.2
(The IP address for the Gigabit Ethernet port 2/1 is now accepted, and two FCIP profiles are using

MDS2(config-profile)# end
MDS2# show fcip profile 21
FCIP Profile 21
    Internet Address is 10.10.11.2 (interface GigabitEthernet2/1)
    Listen Port is 32
(This is a new TCP listen port.)

TCP parameters
    SACK is disabled
    PMTU discover is enabled, reset timeout is 3600 sec
    Keep alive is 60 sec
    Minimum retransmission timeout is 300 ms
    Maximum number of re-transmissions is 4
    Advertised window size is 64 KB
MDS2# show fcip profile 28
FCIP Profile 28
    Internet Address is 10.10.11.2 (interface GigabitEthernet2/1)
    Listen Port is 3225
(This is the default listen port.)

TCP parameters
    SACK is disabled
    PMTU discover is enabled, reset timeout is 3600 sec
    Keep alive is 60 sec
    Minimum retransmission timeout is 300 ms
    Maximum number of re-transmissions is 4
    Advertised window size is 64 KB
```


Displaying Configuration Errors When Bringing Up a Tunnel on a Selected Port Using the CLI

The following example shows a configuration FCIP error when bringing a tunnel up on the selected port. This could be either an FCIP profile issue or an FCIP interface issue. Both sides must be configured correctly.

```
MDS2(config)# fcip profile 21
MDS2(config-profile)# port 13
(Change the TCP listen port on switch MDS2.)

MDS2(config-profile)# end
MDS2(config)# interface fcip 21
MDS2(config-if)# passive-mode
(Put interface fcip 21 in passive mode to guarantee MDS1 initiates a TCP connection.)

module-2# debug ips fcip fsm port 1
module-2# Mar 14 23:08:02 port1: 863:FCIP21: SUP-> Set Port mode 1
Mar 14 23:08:02 port1: 864:FCIP21: SUP-> Port VSAN (1) already set to same value
Mar 14 23:08:02 port1: 865:FCIP21: SUP-> Trunk mode (1) already set to same value
Mar 14 23:08:02 port1: 866:FCIP21: SUP-> Enable tunnel ADMIN UP
Mar 14 23:08:02 port1: 867:FCIP21: Try to Bring UP the Tunnel
Mar 14 23:08:02 port1: 868:FCIP21: Start TCP listener with peer: 10.10.10.2:13
(This debug output from switch MDS2 shows that the FCIP tunnel will not come up because switch MDS

Mar 14 23:08:02 port1: 869:FCIP: Create a new listener object for 10.10.11.2:13
Mar 14 23:08:02 port1: 870:FCIP: Create FCIP Listener with local info: 10.10.11.2:13
MDS1(config)# interface fcip 21
MDS1(config-if)# peer-info ip 10.10.11.2 port 13
(The remote end FCIP interface must be configured to establish a TCP connection on a port that is

MDS1(config-if)# end
MDS1# show interface fcip 21
fcip21 is trunking
(The FCIP tunnel is now up.)

Hardware is GigabitEthernet
  Port WWN is 20:42:00:05:30:00:59:de
  Peer port WWN is 20:42:00:0b:5f:d5:9f:c0
  Admin port mode is auto, trunk mode is on
  Port mode is TE
  vsan is 1
  Trunk vsans (allowed active) (1-2)
  Trunk vsans (operational) (1-2)
  Trunk vsans (up) ()
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) (1-2)
  Using Profile id 21 (interface GigabitEthernet2/1)
  Peer Information
    Peer Internet address is 10.10.11.2 and port is 13
    Special Frame is disabled
  Maximum number of TCP connections is 2
  Time Stamp is disabled
  B-port mode disabled
  TCP Connection Information
    2 Active TCP connections
    Control connection: Local 10.10.10.2:65188, Remote 10.10.11.2:13
(The port is 13 as configured.)

Data connection: Local 10.10.10.2:65190, Remote 10.10.11.2:13
  174 Attempts for active connections, 5 close of connections
MDS2# show ips stats tcp interface gigabitethernet 2/1
TCP Statistics for port GigabitEthernet2/1
```

Connection Stats

44 active openings, 2 accepts
26 failed attempts, 0 reset received, 20 established

Segment stats

2515 received, 2342 sent, 0 retransmitted
0 bad segments received, 0 reset sent

TCP Active Connections

Local Address	Remote Address	State	Send-Q	Recv-Q
10.10.11.2:13	10.10.10.2:65188	ESTABLISH	0	0
(The port is 13 as configured.)				
10.10.11.2:13	10.10.10.2:65190	ESTABLISH	0	0
(The port is 13 as configured.)				
10.10.11.2:13	0.0.0.0:0	LISTEN	0	0
0.0.0.0:3260	0.0.0.0:0	LISTEN	0	0

FCIP Interface Misconfiguration Examples

The examples in this section show FCIP interface misconfigurations.

Displaying FCIP Misconfiguration Examples Using the CLI

The following example shows that the peer-info IP address of the remote endpoint is missing. The debug output is from the IPS module.

```
Module-2# debug ips fcip fsm port 1
module-2# Mar 14 21:37:05 port1: 38:FCIP21: SUP-> Set Port mode 1
Mar 14 21:37:05 port1: 39:FCIP21: SUP-> Port VSAN (1) already set to same value
Mar 14 21:37:05 port1: 40:FCIP21: SUP-> Trunk mode (1) already set to same value
Mar 14 21:37:05 port1: 41:FCIP21: SUP-> Enable tunnel ADMIN UP
Mar 14 21:37:05 port1: 42:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:37:05 port1: 43:FCIP21: Bring up tunnel Failed, peer-ip not set
(The peer IP address is not set.)
```

```
MDS2# show interface fcip 21
fcip21 is down (Link failure or not-connected)
Hardware is GigabitEthernet
Port WWN is 20:42:00:0b:5f:d5:9f:c0
Admin port mode is auto, trunk mode is on
vsan is 1
Using Profile id 21 (interface GigabitEthernet2/1)
Peer Information
(This line shows the Peer Information as empty. The line should read "Peer Internet address is 10.
```

```
Special Frame is disabled
Maximum number of TCP connections is 2
Time Stamp is disabled
B-port mode disabled
TCP Connection Information
0 Attempts for active connections, 0 close of connections
5 minutes input rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
5 minutes output rate 0 bits/sec, 0 bytes/sec, 0 frames/sec
0 frames input, 0 bytes
0 Class F frames input, 0 bytes
0 Class 2/3 frames input, 0 bytes
0 Error frames
0 frames output, 0 bytes
0 Class F frames output, 0 bytes
0 Class 2/3 frames output, 0 bytes
```

```
0 Error frames 0 reassign frames
```

Displaying the FCIP Interface as Administratively Shut Down Using the CLI

The following example shows that the FCIP interface is administratively shut down. The debug output is from the IPS module.

```
Module-2# debug ips fcip fsm port 1
```

```
Module-2# Mar 14 21:32:27 port1: 1:FCIP21: Create tunnel with ifindex: a00014
Mar 14 21:32:27 port1: 2:FCIP21: Get the peer info from the SUP-IPS-MGR
Mar 14 21:32:27 port1: 3:FCIP21: SUP-> Disable tunnel: already in disable state
Mar 14 21:32:27 port1: 4:FCIP21: SUP-> Set Port mode 1
Mar 14 21:32:27 port1: 5:FCIP21: SUP-> Set port index: 21
Mar 14 21:32:27 port1: 6:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:32:27 port1: 7:FCIP21: Tunnel in admin down state
(The tunnel needs no shut down on the FCIP interface.)

Mar 14 21:32:27 port1: 8:FCIP21: SUP-> Set port VSAN: 1
Mar 14 21:32:27 port1: 9:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:32:27 port1: 10:FCIP21: Tunnel in admin down state
Mar 14 21:32:27 port1: 11:FCIP21: SUP-> Set port WWN: 0x2042000b5fd59fc0
Mar 14 21:32:27 port1: 12:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:32:27 port1: 13:FCIP21: Tunnel in admin down state
(The tunnel needs no shut down on the FCIP interface.)

Mar 14 21:32:27 port1: 14:FCIP21: SUP-> Set trunk mode: 1
Mar 14 21:32:27 port1: 15:FCIP21: SUP-> Set source IF: 2080000
Mar 14 21:32:27 port1: 16:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:32:27 port1: 17:FCIP21: Tunnel in admin down state
Mar 14 21:32:27 port1: 18:FCIP21: SUP-> Switch WWN: 0x2000000b5fd59fc0
Mar 14 21:32:27 port1: 19:FCIP21: Try to Bring UP the Tunnel
Mar 14 21:32:27 port1: 20:FCIP21: Tunnel in admin down state
Mar 14 21:32:27 port1: 21:FCIP21: SUP-> Response to SB's pull all tunnel info
Mar 14 21:32:27 port1: 22:FCIP21: SUP-> Set peer port: 3225 current port: 3225
Mar 14 21:32:27 port1: 23:FCIP21: peer port has same value, do nothing
Mar 14 21:32:27 port1: 24:FCIP21: Set number of tcp connection 2
Mar 14 21:32:27 port1: 25:FCIP21: SUP-> Set Local listen IP: 10.10.11.2 current ip
0.0.0.0
Mar 14 21:32:27 port1: 26:FCIP21: SUP-> Set Local listen Port: 3225 current port 3225
Mar 14 21:32:27 port1: 27:FCIP21: SUP-> Enable PMTU Discovery, timeout 3600
Mar 14 21:32:27 port1: 28:FCIP21: SUP-> Set round-trip time to 300 ms. Current value 100
ms
Mar 14 21:32:27 port1: 29:FCIP21: SUP-> Set keep-alive time to 60 sec. current value 60
sec
```

```
MDS2# show interface fcip 21
```

```
fcip21 is down (Administratively down)
  Hardware is GigabitEthernet
  Port WWN is 20:42:00:0b:5f:d5:9f:c0
  Admin port mode is auto, trunk mode is on
  vsan is 1
  Using Profile id 21 (interface GigabitEthernet2/1)
  Peer Information
    Peer Internet address is 10.10.10.2 and port is 3225
    Special Frame is disabled
    Maximum number of TCP connections is 2
  Local MDS trying to connect to remote end point on port 13 and remote end point set to
  default listen port 3225
```

```
MDS2# show interface fcip 21
fcip21 is down (Link failure or not-connected)
  Hardware is GigabitEthernet
  Port WWN is 20:42:00:0b:5f:d5:9f:c0
  Admin port mode is auto, trunk mode is on
  vsan is 1
  Using Profile id 21 (interface GigabitEthernet2/1)
  Peer Information
    Peer Internet address is 10.10.10.2 and port is 13
```

```
MDS1# show fcip profile 21
FCIP Profile 21
  Internet Address is 10.10.10.2 (interface GigabitEthernet2/1)
  Listen Port is 3225
  TCP parameters
    SACK is disabled
    PMTU discover is enabled, reset timeout is 3600 sec
    Keep alive is 60 sec
    Minimum retransmission timeout is 300 ms
    Maximum number of re-transmissions is 4
    Advertised window size is 64 KB
```

Displaying the Debug Output from the Second Switch Using the CLI

The following debug output is from switch MDS2:

```
Mar 14 23:26:07 port1: 1340:FCIP21: Start TCP listener with peer: 10.10.10.2:3225
Mar 14 23:26:07 port1: 1341:FCIP: Create a new listener object for 10.10.11.2:3225
Mar 14 23:26:07 port1: 1342:FCIP: Create FCIP Listener with local info: 10.10.11.2:3225
Mar 14 23:26:07 port1: 1343:FCIP21: Create a DE 0xd802d140 for this tunnel
Mar 14 23:26:07 port1: 1344:FCIP21: Bind the DE 0xd802d140 [1] to tunnel LEP 0x80111570
Mar 14 23:26:07 port1: 1345:FCIP21: Start the active connection [1] to 10.10.10.2:13
Mar 14 23:26:07 port1: 1346:FCIP21: Create a DE 0xd802cdc0 for this tunnel
Mar 14 23:26:07 port1: 1347:FCIP21: Bind the DE 0xd802cdc0 [2] to tunnel LEP 0x80111570
Mar 14 23:26:07 port1: 1348:FCIP21: Start the active connection [2] to 10.10.10.2:13
(The switch is attempting to create a TCP connection on port 13. The creation port must match the

Mar 14 23:26:07 port1: 1349:FCIP21: Active Connect creation FAILED [1]
Mar 14 23:26:07 port1: 1350:FCIP21: Delete the DE [1]0xd802d140
Mar 14 23:26:07 port1: 1351:FCIP21: Delete the DE object [1] 0xd802d140
Mar 14 23:26:07 port1: 1352:FCIP21: Try 7 to bring up the tunnel
Mar 14 23:26:07 port1: 1353:FCIP21: Start the bringup tunnel timer, timeout: 64000
Mar 14 23:26:07 port1: 1354:FCIP21: Active Connect creation FAILED [2]
Mar 14 23:26:07 port1: 1355:FCIP21: Delete the DE [2]0xd802cdc0
Mar 14 23:26:07 port1: 1356:FCIP21: Set lep operation state to DOWN
Mar 14 23:26:07 port1: 1357:FCIP21: Delete the DE object [2] 0xd802cdc0
Mar 14 23:26:07 port1: 1358:FCIP21: Try 8 to bring up the tunnel
Mar 14 23:26:07 port1: 1359:FCIP21: Start the bringup tunnel timer, timeout: 128000
MDS2(config-if)# peer-info ipaddr 10.10.10.2 port 3225
(This changes the start active connection port to match the default port 3225.)
```

Or you can use this command:

```
MDS2(config-if)# no peer-info ipaddr 10.10.10.2 port 13
(Removing port 13 will also set it to the default of 3225.)
```

```
MDS2# show interface fcip 21
fcip21 is trunking
```

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
Hardware is GigabitEthernet
Port WWN is 20:42:00:0b:5f:d5:9f:c0
Peer port WWN is 20:42:00:05:30:00:59:de
Admin port mode is auto, trunk mode is on
Port mode is TE
vsan is 1
Trunk vsans (allowed active) (1-2)
Trunk vsans (operational) (1-2)
Trunk vsans (up) (1-2)
Trunk vsans (isolated) ()
Trunk vsans (initializing) ()
Using Profile id 21 (interface GigabitEthernet2/1)
Peer Information
  Peer Internet address is 10.10.10.2 and port is 3225
  Special Frame is disabled
Maximum number of TCP connections is 2
Time Stamp is disabled
B-port mode disabled
TCP Connection Information
  2 Active TCP connections
  Control connection: Local 10.10.11.2:65330, Remote 10.10.10.2:3225
  Data connection: Local 10.10.11.2:65332, Remote 10.10.10.2:3225
```

Displaying Passive Mode Set on Both Sides of the FCIP Tunnel with the CLI

In the following example, passive mode is set on both sides of the FCIP tunnel:

```
module-2# Mar 14 23:49:06 port1: 1870:FCIP21: SUP-> Set Port mode 1
Mar 14 23:49:06 port1: 1871:FCIP21: SUP-> Port VSAN (1) already set to same value
Mar 14 23:49:06 port1: 1872:FCIP21: SUP-> Trunk mode (1) already set to same value
Mar 14 23:49:06 port1: 1873:FCIP21: SUP-> Enable tunnel ADMIN UP
Mar 14 23:49:06 port1: 1874:FCIP21: Try to Bring UP the Tunnel
Mar 14 23:49:06 port1: 1875:FCIP21: Start TCP listener with peer: 10.10.10.2:3225
Mar 14 23:49:06 port1: 1876:FCIP: Create a new listener object for 10.10.11.2:3225
Mar 14 23:49:06 port1: 1877:FCIP: Create FCIP Listener with local info: 10.10.11.2:3225
Mar 14 23:49:06 port1: 1878:FCIP21: Passive mode set, don't initiate TCP connection
(A TCP connection will not be established when passive mode is set.The Gigabit Ethernet port will
```

```
MDS2# show interface fcip 21
fcip21 is down (Link failure or not-connected)
  Hardware is GigabitEthernet
  Port WWN is 20:42:00:0b:5f:d5:9f:c0
  Admin port mode is auto, trunk mode is on
  vsan is 1
  Using Profile id 21 (interface GigabitEthernet2/1)
  Peer Information
    Peer Internet address is 10.10.10.2 and port is 3225
    Passive mode is enabled
(Passive mode is set, so a TCP connection will not be established.)
```

```
Special Frame is disabled
MDS1# show interface fcip 21
fcip21 is down (Link failure or not-connected)
  Hardware is GigabitEthernet
  Port WWN is 20:42:00:05:30:00:59:de
  Admin port mode is auto, trunk mode is on
  vsan is 1
  Using Profile id 21 (interface GigabitEthernet2/1)
  Peer Information
    Peer Internet address is 10.10.11.2 and port is 3225
    Passive mode is enabled
(Both sides are set to passive mode. You must change one or both sides to no passive-mode under th
```

```
Special Frame is disabled
MDS2(config)# interface fcip 21
MDS2(config-if)# no passive-mode
(Change one or both sides to no passive-mode.)
```

```
MDS2# show interface fcip 21
fcip21 is trunking
```

Displaying a Time Stamp Acceptable Difference Failure Using the CLI

The following example shows a time stamp acceptable difference failure, or no NTP server connected to synchronize clocks. When using time stamps, the MDS switch must be a synchronized clock. NTP is configurable on the MDS 9000 switch.

```
MDS2(config)# interface fcip 21
MDS2(config-if)# time-stamp
module-2# debug ips fcip fsm port 1
Mar 15 00:01:35 port1: 3248:FCIP21: IPS-> Enable timestamp acceptable difference 1000
(Time stamp is enabled under the FCIP interface. The default acceptable difference is 1000.)

Mar 15 00:01:35 port1: 3249:FCIP21: IPS-> acc diff in sec: 0x1 frac: 0x0
Mar 15 00:01:35 port1: 3250:FCIP21: Sending response code: 0
Mar 15 00:01:48 port1: 3251:FCIP21: Time stamp tolerance check failed local time:
0x3e726d6c2db994b7 tolerance: 0x100000000 recv time: 0x3e7251ace20db73a
(The timestamp difference failed the acceptable difference.)

Mar 15 00:01:48 port1: 3252:FCIP21: Time stamp tolerance check failed local time:
0x3e726d6c2db994b7 tolerance: 0x100000000 recv time: 0x3e7251ace20db73a
Mar 15 00:01:48 port1: 3253:FCIP21: Time stamp tolerance check failed local time:
0x3e726d6c2db994b7 tolerance: 0x100000000 recv time: 0x3e7251ace20db73a
<<< cut >>>
Mar 15 00:01:48 port1: 3290:FCIP21: Time stamp tolerance check failed local time:
0x3e726d6c2db994b7 tolerance: 0x100000000 recv time: 0x3e7251ace20db73a
Mar 15 00:01:48 port1: 3291:FCIP21: (fcip_de_rcv): Previous partial packet -
Concatenating
Mar 15 00:01:48 port1: 3292:FCIP21: Time stamp tolerance check failed local time:
0x3e726d6c2db994b7 tolerance: 0x100000000 recv time: 0x3e7251ace20db73a
Mar 15 00:01:48 port1: 3293:FCIP21: FCIP frame len 0x300 is not within correct range <<<
?? >>>
Mar 15 00:01:48 port1: 3294:FCIP21: Delete the DE [2]0xd802d680
Mar 15 00:01:48 port1: 3295:FCIP21: replace the eport entry at index: 1
Mar 15 00:01:48 port1: 3296:FCIP21: DE [-670902656] 0x00000002 terminate tcp connection
0xd8072800
(The TCP connection is disconnected because the time stamp difference is too large.)

Mar 15 00:01:48 port1: 3297:FCIP21: Delete the DE object [2] 0xd802d680
Mar 15 00:01:48 port1: 3298:FCIP21: Delete the DE [1]0xd802cf00
Mar 15 00:01:48 port1: 3299:FCIP21: Unregister from flamingo port_index: 0x21
Mar 15 00:01:48 port1: 3300:FCIP21: Send Link down to SUP
Mar 15 00:01:48 port1: 3301:FCIP21: Start the bringup tunnel timer, timeout: 18470
Mar 15 00:01:48 port1: 3302:FCIP21: replace the eport entry at index: 0
Mar 15 00:01:48 port1: 3303:FCIP21: Set lep operation state to DOWN
Mar 15 00:01:48 port1: 3304:FCIP21: DE [-670904576] 0x00000001 terminate tcp connection
0xd8072c00
Mar 15 00:01:48 port1: 3305:FCIP21: Delete the DE object [1] 0xd802cf00
Mar 15 00:01:50 port1: 3306:FCIP21: Received new TCP connection from peer:
10.10.10.2:65066
(The TCP connection begins trying to reestablish the connection.)

Mar 15 00:01:50 port1: 3307:FCIP21: Tunnel is not ADMIN UP state, reject new TCP
connection from 10.10.10.2:65066
```

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
Mar 15 00:01:50 port1: 3308:FCIP21: Received new TCP connection from peer:
10.10.10.2:65064
Mar 15 00:01:50 port1: 3309:FCIP21: Tunnel is not ADMIN UP state, reject new TCP
connection from 10.10.10.2:65064
Mar 15 00:01:56 port1: 3310:FCIP21: SUP-> Set Port mode 1
Mar 15 00:01:56 port1: 3311:FCIP21: SUP-> Port VSAN (1) already set to same value
Mar 15 00:01:56 port1: 3312:FCIP21: SUP-> Set trunk mode: 1
Mar 15 00:01:56 port1: 3313:FCIP21: SUP-> Enable tunnel ADMIN UP
Mar 15 00:01:56 port1: 3314:FCIP21: Try to Bring UP the Tunnel
Mar 15 00:01:56 port1: 3315:FCIP21: tunnel bring-up debounce timer set, wait for timer to
pop
(Connect the NTP server or synchronized clocks, or increase the acceptable difference.)
```

```
module-2# debug ips fcip fsm port 1
module-2#
Jan 14 14:22:08 port1: 854886:FCIP21: IPS-> Enable timestamp acceptable difference 2000
Jan 14 14:22:08 port1: 854887:FCIP21: IPS-> acc diff in sec: 0x2 frac: 0x0
(The time stamp acceptable difference passes and the tunnel continues to be brought up.)
```

```
module-2#
module-2# Jan 14 14:22:39 port1: 854932:FCIP21: Received new TCP connection from peer:
10.10.10.2:64172
Jan 14 14:22:39 port1: 854933:FCIP21: Create a DE 0xd802d5c0 for this tunnel
Jan 14 14:22:39 port1: 854934:FCIP21: Bind the DE 0xd802d5c0 [1] to tunnel LEP 0x80111570
Jan 14 14:22:39 port1: 854935:FCIP21: Bind DE 1 to TCP-hdl 0xd8071000
Jan 14 14:22:39 port1: 854936:FCIP21: Bind DE 1 to eport 0x80110550
Jan 14 14:22:39 port1: 854937:FCIP21: bind de 1 in eport 0x80110550, hash = 1 num-conn: 2
Jan 14 14:22:39 port1: 854938:FCIP21: Received new TCP connection from peer: 10
.10.10.2:64170
Jan 14 14:22:39 port1: 854939:FCIP21: Create a DE 0xd802c900 for this tunnel
Jan 14 14:22:39 port1: 854940:FCIP21: Bind the DE 0xd802c900 [2] to tunnel LEP
0x80111570
Jan 14 14:22:39 port1: 854941:FCIP21: Bind DE 2 to TCP-hdl 0xd8070000
Jan 14 14:22:39 port1: 854942:FCIP21: Bind DE 2 to eport 0x80110550
Jan 14 14:22:39 port1: 854943:FCIP21: bind de 2 in eport 0x80110550, hash = 2 n
um-conn: 2
Jan 14 14:22:39 port1: 854944:FCIP21: Send LINK UP to SUP
Jan 14 14:22:39 port1: 854945:FCIP21: *** Received eisl frame in E mode
Jan 14 14:22:39 port1: 854946:FCIP21: SUP-> Set trunk mode: 2
Jan 14 14:22:39 port1: 854947:FCIP21: Change the operational mode to TRUNK
MDS2# show interface fcip 21
fcip21 is trunking
  Hardware is GigabitEthernet
  Port WWN is 20:42:00:0b:5f:d5:9f:c0
  Peer port WWN is 20:42:00:05:30:00:59:de
  Admin port mode is auto, trunk mode is on
  Port mode is TE
  vsan is 1
  Trunk vsans (allowed active) (1-2)
  Trunk vsans (operational) (1-2)
  Trunk vsans (up) (1-2)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
  Using Profile id 21 (interface GigabitEthernet2/1)
  Peer Information
    Peer Internet address is 10.10.10.2 and port is 3225
    Special Frame is disabled
  Maximum number of TCP connections is 2
  Time Stamp is enabled, acceptable time difference 2000 ms
  B-port mode disabled
  TCP Connection Information
```

Figure 20-7 shows a trace of time stamp difference failure.

Figure 20-7 Trace of Time-stamp Difference Failure


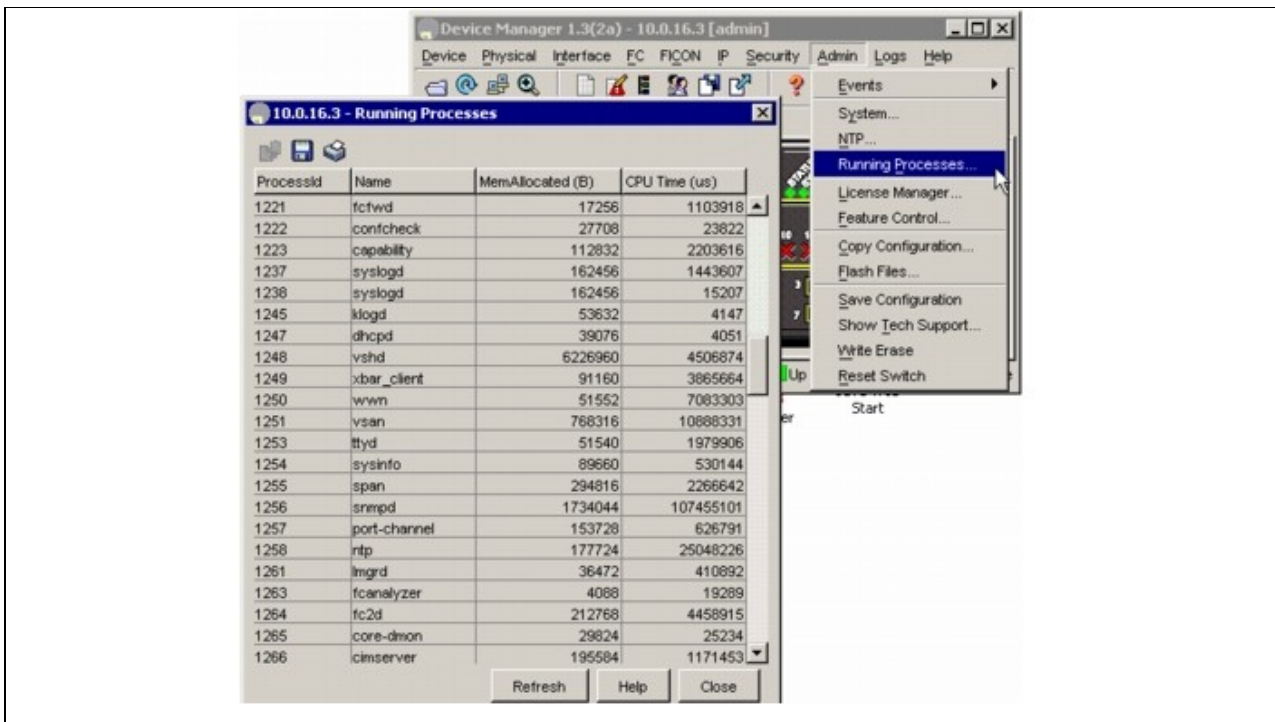


Figure 20-8 shows a trace of time stamp difference accepted.

Figure 20-8 Trace of Time-Stamp Difference Accepted



FCIP Special Frame Tunnel Creation and Monitoring

The FCIP tunnel configuration (see the "One-to-One FCIP Tunnel Creation and Monitoring" section and the "One-to-Three FCIP Tunnel Creation and Monitoring" section) must be completed before adding the FCIP special frame configuration. This section describes how to correctly configure and show an FCIP tunnel with a special frame.

Configuring and Displaying an FCIP Tunnel with Special Frame Using the CLI

To configure and display an FCIP tunnel with a special frame using the CLI, follow these steps:

1. Use the **show wwn switch** command to verify the WWN of each MDS 9000 Family switch end point.

```
MDS2# ''show wwn switch''
Switch WWN is 20:00:00:0b:5f:d5:9f:c0
```

2. Use the **special-frame peer-wwn** command to enable the FCIP special frame that is used in creating the FCIP tunnel.

```
MDS2(config)# interface fcip 21
MDS2(config-if)# special-frame peer-wwn 20:00:00:05:30:00:59:de profile-id 1
module-2#
Jan 14 15:25:38 port1: 857314:FCIP21: SUP-> Set Port mode 1
Jan 14 15:25:38 port1: 857315:FCIP21: SUP-> Port VSAN (1) already set to same value
Jan 14 15:25:38 port1: 857316:FCIP21: SUP-> Trunk mode (1) already set to same value
Jan 14 15:25:38 port1: 857317:FCIP21: SUP-> Enable tunnel ADMIN UP
Jan 14 15:25:38 port1: 857318:FCIP21: Try to Bring UP the Tunnel
Jan 14 15:25:38 port1: 857319:FCIP21: Start TCP listener with peer: 10.10.10.2:3225
Jan 14 15:25:38 port1: 857320:FCIP: Create a new listener object for 10.10.11.2:3225
Jan 14 15:25:38 port1: 857321:FCIP: Create FCIP Listener with local info: 10.10.11.2:3225
Jan 14 15:25:38 port1: 857322:FCIP21: Create a DE 0xd802cd00 for this tunnel
Jan 14 15:25:38 port1: 857323:FCIP21: Bind the DE 0xd802cd00 [1] to tunnel LEP 0x80111570
Jan 14 15:25:38 port1: 857324:FCIP21: Start the active connection [1] to 10.10.10.2:3225
Jan 14 15:25:38 port1: 857325:FCIP21: Create a DE 0xd802db40 for this tunnel
Jan 14 15:25:38 port1: 857326:FCIP21: Bind the DE 0xd802db40 [2] to tunnel LEP 0x80111570
Jan 14 15:25:38 port1: 857327:FCIP21: Start the active connection [2] to 10.10.10.2:3225
Jan 14 15:25:38 port1: 857328:FCIP21: Active Connect creation SUCCEEDED [1]
Jan 14 15:25:38 port1: 857329:FCIP21: Bind DE 1 to TCP-hdl 0xd8072c00
Jan 14 15:25:38 port1: 857330:FCIP21: Setup for Special Frame handling: I'm Originator
(This begins the Special Frame setup of the Originator.)
Jan 14 15:25:38 port1: 857331:FCIP21: Send the SF as Originator & wait for response
(The Special Frame is sent.)
Jan 14 15:25:38 port1: 857332:FCIP21: Setup timer to wait for SF
Jan 14 15:25:38 port1: 857333:FCIP21: Active Connect creation SUCCEEDED [2]
(The Special Frame is correctly configured with the WWN of the remote MDS 9000 switch.)
Jan 14 15:25:38 port1: 857334:FCIP21: Bind DE 2 to TCP-hdl 0xd8072000
Jan 14 15:25:38 port1: 857335:FCIP21: Setup for Special Frame handling: I'm Originator
Jan 14 15:25:38 port1: 857336:FCIP21: Send the SF as Originator & wait for response
Jan 14 15:25:38 port1: 857337:FCIP21: Setup timer to wait for SF
Jan 14 15:25:38 port1: 857338:FCIP21: processing SF frame, I'm Originator
Jan 14 15:25:38 port1: 857339:FCIP21: Bind DE 1 to eport 0x80110550
Jan 14 15:25:38 port1: 857340:FCIP21: bind de 1 in eport 0x80110550, hash = 1 num-conn: 2
Jan 14 15:25:38 port1: 857341:FCIP21: processing SF frame, I'm Originator
Jan 14 15:25:38 port1: 857342:FCIP21: Bind DE 2 to eport 0x80110550
Jan 14 15:25:38 port1: 857343:FCIP21: bind de 2 in eport 0x80110550, hash = 2 num-conn: 2
Jan 14 15:25:38 port1: 857344:FCIP21: Send LINK UP to SUP
Jan 14 15:25:39 port1: 857345:FCIP21: SUP-> Set trunk mode: 2
Jan 14 15:25:39 port1: 857346:FCIP21: Change the operational mode to TRUNK
Jan 14 15:25:39 port1: 857347:FCIP21: *** Received non-eisl frame in TE mode 64 64
```


3. Use the **show interface fcip** command to verify that a special frame is enabled.

```
MDS2# show interface fcip 21
fcip21 is trunking
  Hardware is GigabitEthernet
  Port WWN is 20:42:00:0b:5f:d5:9f:c0
  Peer port WWN is 20:42:00:05:30:00:59:de
  Admin port mode is auto, trunk mode is on
  Port mode is TE
  vsan is 1
  Trunk vsans (allowed active) (1-2)
  Trunk vsans (operational) (1-2)
  Trunk vsans (up) (1-2)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
  Using Profile id 21 (interface GigabitEthernet2/1)
  Peer Information
    Peer Internet address is 10.10.10.2 and port is 3225
    Special Frame is enabled
    Peer switch WWN is 20:00:00:05:30:00:59:de
  Maximum number of TCP connections is 2
  Time Stamp is enabled, acceptable time difference 3000 ms
  B-port mode disabled
  TCP Connection Information
    2 Active TCP connections
    Control connection: Local 10.10.11.2:64792, Remote 10.10.10.2:3225
    Data connection: Local 10.10.11.2:64794, Remote 10.10.10.2:3225
    372 Attempts for active connections, 345 close of connections
  TCP Parameters
    Path MTU 1500 bytes
    Current retransmission timeout is 300 ms
    Round trip time: Smoothed 10 ms, Variance: 5
    Advertized window: Current: 64 KB, Maximum: 64 KB, Scale: 1
    Peer receive window: Current: 64 KB, Maximum: 64 KB, Scale: 1
    Congestion window: Current: 2 KB, Slow start threshold: 1048560 KB
```

4. Use the **show wwn switch** command on the remote switch to verify the peer switch WWN.

Figure 20-9 shows a trace of an FCIP tunnel with a special frame.

Figure 20-9 Trace of FCIP Tunnel with a Special Frame



Special Frame Misconfiguration Example

The following example shows an incorrect peer WVN when using a special frame.

Example 20-1 Annotated Example of Incorrect Peer WVN with Special Frame Enabled

```

module-2# Jan 14 15:14:30 port1: 855278:FCIP21: SUP-> Set Port mode 1
Jan 14 15:14:30 port1: 855279:FCIP21: SUP-> Port VSAN (1) already set to same value
Jan 14 15:14:30 port1: 855280:FCIP21: SUP-> Trunk mode (1) already set to same
Jan 14 15:14:30 port1: 855281:FCIP21: SUP-> Enable tunnel ADMIN UP
Jan 14 15:14:30 port1: 855282:FCIP21: Try to Bring UP the Tunnel
Jan 14 15:14:30 port1: 855283:FCIP21: Start TCP listener with peer: 10.10.10.2:3225
Jan 14 15:14:30 port1: 855284:FCIP: Create a new listener object for 10.10.11.2:3225
Jan 14 15:14:30 port1: 855285:FCIP: Create FCIP Listener with local info: 10.10.11.2:3225
Jan 14 15:14:30 port1: 855286:FCIP21: Create a DE 0xd802d240 for this tunnel
Jan 14 15:14:30 port1: 855287:FCIP21: Bind the DE 0xd802d240 [1] to tunnel LEP 0x80111570
Jan 14 15:14:30 port1: 855288:FCIP21: Start the active connection [1] to 10.10.10.2:3225
Jan 14 15:14:30 port1: 855289:FCIP21: Create a DE 0xd802d200 for this tunnel
Jan 14 15:14:30 port1: 855290:FCIP21: Bind the DE 0xd802d200 [2] to tunnel LEP 0x80111570
Jan 14 15:14:30 port1: 855291:FCIP21: Start the active connection [2] to 10.10.10.2:3225
Jan 14 15:14:30 port1: 855292:FCIP21: Active Connect creation SUCCEEDED [1]
Jan 14 15:14:30 port1: 855293:FCIP21: Bind DE 1 to TCP-hdl 0xd8072c00
Jan 14 15:14:30 port1: 855294:FCIP21: Setup for Special Frame handling: I'm Originator
Jan 14 15:14:30 port1: 855295:FCIP21: Send the SF as Originator & wait for response
Jan 14 15:14:30 port1: 855296:FCIP21: Setup timer to wait for SF
Jan 14 15:14:30 port1: 855297:FCIP21: Active Connect creation SUCCEEDED [2]
Jan 14 15:14:30 port1: 855298:FCIP21: Bind DE 2 to TCP-hdl 0xd8072000
Jan 14 15:14:30 port1: 855299:FCIP21: Setup for Special Frame handling: I'm Originator
Jan 14 15:14:30 port1: 855300:FCIP21: Send the SF as Originator & wait for response
Jan 14 15:14:30 port1: 855301:FCIP21: Setup timer to wait for SF
Jan 14 15:14:30 port1: 855302:FCIP21: TCP Received a close connection [1] reason 1
Jan 14 15:14:30 port1: 855303:FCIP21: Delete the DE [1]0xd802d240
Jan 14 15:14:30 port1: 855304:FCIP21: DE [-670903744] 0x00000001 terminate tcp connection
0xd8072c00
Jan 14 15:14:30 port1: 855305:FCIP21: Delete the DE object [1] 0xd802d240
Jan 14 15:14:30 port1: 855306:FCIP21: lep not bound, close only de [1]
Jan 14 15:14:30 port1: 855307:FCIP21: TCP Received a close connection [2] reason 1

```

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
Jan 14 15:14:30 port1: 855308:FCIP21: Delete the DE [2]0xd802d200
Jan 14 15:14:30 port1: 855309:FCIP21: Set lep operation state to DOWN
Jan 14 15:14:30 port1: 855310:FCIP21: Start the bringup tunnel timer, timeout: 38740
Jan 14 15:14:30 port1: 855311:FCIP21: DE [-670903808] 0x00000002 terminate tcp connection
0xd8072000
Jan 14 15:14:30 port1: 855312:FCIP21: Delete the DE object [2] 0xd802d200
Jan 14 15:14:30 port1: 855313:FCIP21: lep not bound, close only de [2]
Jan 14 15:14:31 port1: 855314:FCIP21: Received new TCP connection from peer:
10.10.10.2:64050
Jan 14 15:14:31 port1: 855315:FCIP21: Create a DE 0xd802d080 for this tunnel
Jan 14 15:14:31 port1: 855316:FCIP21: Bind the DE 0xd802d080 [1] to tunnel LEP 0x80111570
Jan 14 15:14:31 port1: 855317:FCIP21: Bind DE 1 to TCP-hdl 0xd8072000
Jan 14 15:14:31 port1: 855318:FCIP21: Setup for Special Frame handling: I'm Responder
Jan 14 15:14:31 port1: 855319:FCIP21: Setup timer to wait for SF
Jan 14 15:14:31 port1: 855320:FCIP21: processing SF frame, I'm Responder
Jan 14 15:14:31 port1: 855321:FCIP21: Source FC fabric name in SF (0x20000005300059de)
does not match LEP's peer fabric WWN (0x20010005300059df)
Jan 14 15:14:31 port1: 855322:FCIP21: Delete the DE [1]0xd802d080
Jan 14 15:14:31 port1: 855323:FCIP21: Set lep operation state to DOWN
Jan 14 15:14:31 port1: 855324:FCIP21: DE [-670904192] 0x00000001 terminate tcp connection
0xd8072000
Jan 14 15:14:31 port1: 855325:FCIP21: Delete the DE object [1] 0xd802d080
Jan 14 15:14:31 port1: 855326:FCIP21: Received new TCP connection from peer:
10.10.10.2:64048
Jan 14 15:14:31 port1: 855327:FCIP21: Create a DE 0xd802d200 for this tunnel
Jan 14 15:14:31 port1: 855328:FCIP21: Bind the DE 0xd802d200 [1] to tunnel LEP 0x80111570
Jan 14 15:14:31 port1: 855329:FCIP21: Bind DE 1 to TCP-hdl 0xd8072c00
Jan 14 15:14:31 port1: 855330:FCIP21: Setup for Special Frame handling: I'm Responder
Jan 14 15:14:31 port1: 855331:FCIP21: Setup timer to wait for SF
Jan 14 15:14:31 port1: 855332:FCIP21: processing SF frame, I'm Responder
Jan 14 15:14:31 port1: 855333:FCIP21: Source FC fabric name in SF (0x20000005300059de)
does not match LEP's peer fabric WWN (0x20010005300059df)
Jan 14 15:14:31 port1: 855334:FCIP21: Delete the DE [1]0xd802d200
Jan 14 15:14:31 port1: 855335:FCIP21: Set lep operation state to DOWN
Jan 14 15:14:31 port1: 855336:FCIP21: DE [-670903808] 0x00000001 terminate tcp connection
0xd8072c00
Jan 14 15:14:31 port1: 855337:FCIP21: Delete the DE object [1] 0xd802d200
Jan 14 15:14:37 port1: 855338:FCIP21: Received new TCP connection from peer:
10.10.10.2:64046
Jan 14 15:14:37 port1: 855339:FCIP21: Create a DE 0xd802d5c0 for this tunnel
Jan 14 15:14:37 port1: 855340:FCIP21: Bind the DE 0xd802d5c0 [1] to tunnel LEP 0x80111570
Jan 14 15:14:37 port1: 855341:FCIP21: Bind DE 1 to TCP-hdl 0xd8071000
Jan 14 15:14:37 port1: 855342:FCIP21: Setup for Special Frame handling: I'm Responder
Jan 14 15:14:37 port1: 855343:FCIP21: Setup timer to wait for SF
Jan 14 15:14:37 port1: 855344:FCIP21: processing SF frame, I'm Responder
Jan 14 15:14:37 port1: 855345:FCIP21: Source FC fabric name in SF (0x20000005300059de)
does not match LEP's peer fabric WWN (0x20010005300059df)
Jan 14 15:14:37 port1: 855346:FCIP21: Delete the DE [1]0xd802d5c0
Jan 14 15:14:37 port1: 855347:FCIP21: Set lep operation state to DOWN
Jan 14 15:14:37 port1: 855348:FCIP21: DE [-670902848] 0x00000001 terminate tcp connection
0xd8071000
Jan 14 15:14:37 port1: 855349:FCIP21: Delete the DE object [1] 0xd802d5c0
Jan 14 15:14:37 port1: 855350:FCIP21: Received new TCP connection from peer:
10.10.10.2:64044
Jan 14 15:14:37 port1: 855351:FCIP21: Create a DE 0xd802cac0 for this tunnel
Jan 14 15:14:37 port1: 855352:FCIP21: Bind the DE 0xd802cac0 [1] to tunnel LEP 0x80111570
Jan 14 15:14:37 port1: 855353:FCIP21: Bind DE 1 to TCP-hdl 0xd8071400
Jan 14 15:14:37 port1: 855354:FCIP21: Setup for Special Frame handling: I'm Responder
Jan 14 15:14:37 port1: 855355:FCIP21: Setup timer to wait for SF
Jan 14 15:14:37 port1: 855356:FCIP21: processing SF frame, I'm Responder
Jan 14 15:14:37 port1: 855357:FCIP21: Source FC fabric name in SF (0x20000005300059de)
does not match LEP's peer fabric WWN (0x20010005300059df)
Jan 14 15:14:37 port1: 855358:FCIP21: Delete the DE [1]0xd802cac0
Jan 14 15:14:37 port1: 855359:FCIP21: Set lep operation state to DOWN
```

```
Jan 14 15:14:37 port1: 855360:FCIP21: DE [-670905664] 0x00000001 terminate tcp connection
0xd8071400
Jan 14 15:14:37 port1: 855361:FCIP21: Delete the DE object [1] 0xd802cac0
```

Troubleshooting FCIP Link Flaps

If you have an FCIP link that flaps, adjust the TCP keepalive and max retransmission values. In Fabric Manager, choose **ISLs > FCIP**, select the **Profiles** tab and set the Keepalive field. In the CLI, use the **tcp keepalive-timeout** and **tcp max-retransmissions** commands in FCIP profile submode.

Troubleshooting FCIP ISL Link Failures

If you have an FCIP ISL link that experiences unexplained failures, the FCIP profile in use may not be configured to match the underlying physical link characteristics and the shared traffic constraints.

For example, if you have a 200 Mbps physical link and you have configured two FCIP tunnels across that link, each with 155 Mbps traffic, link failures will occur because the physical connection cannot handle the traffic.

To resolve the problem, tune the TCP parameters in the FCIP profile to match the characteristic and constraints of the physical link and any other traffic shared on that link.

See "Verifying the Configuration of the Profiles Using the CLI" section for details on the adjustable parameters in the FCIP profile.

Troubleshooting FCIP and Compression

If you have an FCIP tunnel between an IPS module and a Cisco MDS 14/2 module, Cisco MDS 9221i switch, or MDS 18/4 module, use the same compression mode on both sides of the FCIP tunnel. In this specific configuration, avoid compression mode 1, because the Cisco MDS module could send compressed traffic faster than the IPS module could process.

iSCSI Issues


This section contains information on troubleshooting iSCSI and includes the following topics:

- Troubleshooting iSCSI Authentication
- Displaying iSCSI Authentication Using Fabric Manager
- Displaying iSCSI Authentication Using the CLI
- Troubleshooting User Name and Password Configuration
- RADIUS Configuration Troubleshooting
- Troubleshooting RADIUS Routing Configuration
- Troubleshooting Dynamic iSCSI Configuration

Troubleshooting iSCSI Authentication

iSCSI user login authentication is required with the Cisco MDS 9000 Family switch. There are two ways to authenticate iSCSI users: either locally in the switch's configuration file or using the RADIUS server database.

Figure 20-10 shows a successful iSCSI login for the Windows 2000 driver.

Figure 20-10 Successful iSCSI Login Status Window

On Solaris systems, a successful login is found in the /var/adm/messages directory and should look similar to the following example:

```
Mar 14 12:53:23 ca-sun1 iscsid[12745]: [ID 702911 daemon.notice] discovery process for
172.22.91.223 finished, exiting
Mar 14 12:58:45 ca-sun1 iscsid[12802]: [ID 448557 daemon.notice] logged into
DiscoveryAddress 172.22.91.223:3260 isid 023d0040
Mar 14 12:58:45 ca-sun1 iscsid[12802]: [ID 702911 daemon.notice] iSCSI target 2 =
iqn.com.domainname.vrrp-11.gw.21000020375aff77 at0
Mar 14 12:58:45 ca-sun1 iscsid[12809]: [ID 529321 daemon.notice] logged into target
iqn.com.domainname.vrrp-11.gw.21000020375aff77 7
Mar 14 12:58:45 ca-sun1 iscsid[12802]: [ID 702911 daemon.notice] iSCSI target 3 =
iqn.com.domainname.vrrp-11.gw.21000020374baf02 at0
Mar 14 12:58:45 ca-sun1 iscsid[12810]: [ID 529321 daemon.notice] logged into target
iqn.com.domainname.vrrp-11.gw.21000020374baf02 7
```

Figure 20-11 shows a failed iSCSI login for the Windows 2000 driver.

Figure 20-11 Failed iSCSI Login Status Window

Cisco_MDS_SanOS_Troubleshooting_Guide_-_Troubleshooting_IP_Storage_Services

```
switch# show radius-server
retransmission count:3
timeout value:5
```


following RADIUS servers are configured:

```
171.71.49.197:
    available for authentication on port:1812
    available for accounting on port:1813
    RADIUS shared secret:radius
```

Adjust the RADIUS timeout and retransmission accordingly, as they have a default value of 1 second and 1 time.

Figure 20-12 shows a Windows-based RADIUS server configuration.

Figure 20-12 Windows-Based RADIUS Server Configuration Dialog Box



If the items in Figure 20-12 match your switch's configuration, then verify that the client user name and password also match those in the RADIUS server.

The following example shows the output of the **debug security radius** command, if the iSCSI client logs in successfully.

```
switch#
switch# Mar  4 23:16:20 securityd: received CHAP authentication request for user002
Mar  4 23:16:20 securityd: RADIUS is enabled, hence it will be tried first for CHAP
authentication
Mar  4 23:16:20 securityd: reading RADIUS configuration
Mar  4 23:16:20 securityd: opening radius configuration for group:default
Mar  4 23:16:20 securityd: opened the configuration successfully
Mar  4 23:16:20 securityd: GET request for RADIUS global config
Mar  4 23:16:20 securityd: got back the return value of global radius configuration
operation:success
```

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
Mar  4 23:16:20 securityd: closing RADIUS pss configuration
Mar  4 23:16:20 securityd: opening radius configuration for group:default
Mar  4 23:16:20 securityd: opened the configuration successfully
Mar  4 23:16:20 securityd: GETNEXT request for radius index:0 addr:
Mar  4 23:16:20 securityd: got some reply from 171.71.49.197
Mar  4 23:16:20 securityd: verified the response from:171.71.49.197
Mar  4 23:16:20 securityd: RADIUS server sent accept for authentication request for
user002
Mar  4 23:16:25 securityd: received CHAP authentication request for user002
Mar  4 23:16:25 securityd: RADIUS is enabled, hence it will be tried first for CHAP
authentication
Mar  4 23:16:25 securityd: reading RADIUS configuration
Mar  4 23:16:25 securityd: opening radius configuration for group:default
Mar  4 23:16:25 securityd: opened the configuration successfully
Mar  4 23:16:25 securityd: GET request for RADIUS global config
Mar  4 23:16:25 securityd: got back the return value of global radius configuration
operation:success
Mar  4 23:16:25 securityd: closing RADIUS pss configuration
Mar  4 23:16:25 securityd: opening radius configuration for group:default
Mar  4 23:16:25 securityd: opened the configuration successfully
Mar  4 23:16:25 securityd: GETNEXT request for radius index:0 addr:
Mar  4 23:16:25 securityd: got some reply from 171.71.49.197
Mar  4 23:16:25 securityd: verified the response from:171.71.49.197
Mar  4 23:16:25 securityd: RADIUS server sent accept for authentication request for
user002
Mar  4 23:16:25 securityd: got some reply from 171.71.49.197
Mar  4 23:16:25 securityd: verified the response from:171.71.49.197
Mar  4 23:16:25 securityd: RADIUS server sent accept for authentication request for
user002
```

The previous example shows that the iSCSI client has been authenticated three times, first for the switch login, and the second and third times for the iSCSI driver login. The switch sends RADIUS attributes 1, 3, 4, 5, 6, 60 and 61 to the RADIUS server. The RADIUS server only needs to respond with **request accept** or **request reject**.

The following example shows a RADIUS authentication.

```
639 2003y3m14d 15h12m48s -----
640 2003y3m14d 15h12m48s Message Type=Access_Request
641 2003y3m14d 15h12m48s ID=243, Length=90
642 2003y3m14d 15h12m48s User name=user002
643 2003y3m14d 15h12m48s NAS IP address=2887147911
644 2003y3m14d 15h12m48s CHAP password=?j:<žWøøë-K-èÛ<]
645 2003y3m14d 15h12m48s CHAP challenge=n8NÝgø$"__ó4}ôx
646 2003y3m14d 15h12m48s NAS port=1426
647 2003y3m14d 15h12m48s NAS port type=5
648 2003y3m14d 15h12m48s Service type=8
649 2003y3m14d 15h12m48s User (user002) authenticate OK.
650 2003y3m14d 15h12m54s -----
651 2003y3m14d 15h12m54s Message Type=Access_Request
652 2003y3m14d 15h12m54s ID=60, Length=90
653 2003y3m14d 15h12m54s User name=user002
654 2003y3m14d 15h12m54s NAS IP address=2887147911
655 2003y3m14d 15h12m54s CHAP password=_¿Èò_à!_AèC0__`
656 2003y3m14d 15h12m54s CHAP challenge=_/ôæ?×!âßÈ 4_žZH
657 2003y3m14d 15h12m54s NAS port=1426
658 2003y3m14d 15h12m54s NAS port type=5
659 2003y3m14d 15h12m54s Service type=8
660 2003y3m14d 15h12m54s User (user002) authenticate OK.
661 2003y3m14d 15h12m54s -----
662 2003y3m14d 15h12m54s Message Type=Access_Request
663 2003y3m14d 15h12m54s ID=179, Length=90
```

```

664 2003y3m14d 15h12m54s User name=user002
665 2003y3m14d 15h12m54s NAS IP address=2887147911
666 2003y3m14d 15h12m54s CHAP password=
667 2003y3m14d 15h12m54s CHAP challenge=#ùÊÝü{"__"Ž_Ux
668 2003y3m14d 15h12m54s NAS port=1426
669 2003y3m14d 15h12m54s NAS port type=5
670 2003y3m14d 15h12m54s Service type=8
671 2003y3m14d 15h12m54s User (user002) authenticate OK.

```

--5Àurfàxh

Troubleshooting RADIUS Routing Configuration

The switch sends the RADIUS authentication request from the mgmt0 interface, so the correct route to the RADIUS server must be defined. If no correct route is defined, the switch may send the RADIUS request from the Gigabit Ethernet port. In that case, the RADIUS server returns the accept to the Gigabit Ethernet port and the switch does not get the response.

Displaying the Debug Output for RADIUS Authentication Request Routing Using the CLI

The following example shows the output from the **debug security radius** command.

```

switch# Mar  5 00:51:13 securityd: received CHAP authentication request for user002
Mar  5 00:51:13 securityd: RADIUS is enabled, hence it will be tried first for CHAP
authentication
Mar  5 00:51:13 securityd: reading RADIUS configuration
Mar  5 00:51:13 securityd: opening radius configuration for group:default
Mar  5 00:51:13 securityd: opened the configuration successfully
Mar  5 00:51:13 securityd: GET request for RADIUS global config
Mar  5 00:51:13 securityd: got back the return value of global radius configuration
operation:success
Mar  5 00:51:13 securityd: closing RADIUS pss configuration
Mar  5 00:51:13 securityd: opening radius configuration for group:default
Mar  5 00:51:13 securityd: opened the configuration successfully
Mar  5 00:51:13 securityd: GETNEXT request for radius index:0 addr:
Mar  5 00:51:18 securityd: sending data to 171.71.49.197
Mar  5 00:51:18 securityd: waiting for response from 171.71.49.197
Mar  5 00:51:23 securityd: sending data to 171.71.49.197
Mar  5 00:51:23 securityd: waiting for response from 171.71.49.197
Mar  5 00:51:28 securityd: sending data to 171.71.49.197
Mar  5 00:51:28 securityd: waiting for response from 171.71.49.197
Mar  5 00:51:33 securityd: trying out next server
Mar  5 00:51:33 securityd: no response from RADIUS server for authentication user002
Mar  5 00:51:33 securityd: doing local chap authentication for user002
Mar  5 00:51:33 securityd: local chap authentication result for user002:user not present

```

Troubleshooting Dynamic iSCSI Configuration

A physical Fibre Channel target (target pWWN) presented as an iSCSI target, makes the physical target accessible to an iSCSI initiator. The IPS module presents physical Fibre Channel targets as iSCSI targets to iSCSI initiators in one of two ways: dynamic mapping or static mapping.

By default, the IPS module does not automatically import Fibre Channel targets. Either dynamic or static mapping must be configured before the IPS module makes Fibre Channel targets available to iSCSI initiators. When both are configured, statically mapped Fibre Channel targets have the configured name. Targets that are not mapped will be advertised with the name created by the conventions explained in this section.

Checking the Configuration

Use the following guidelines to verify the configuration of the Gigabit Ethernet interface:

- Ensure that you are configuring the proper slot or port.
- Ensure that the Gigabit Ethernet interfaces are not shut down. Each Gigabit Ethernet interface is "partnered" with a virtual iSCSI interface. For iSCSI to operate on a particular Gigabit Ethernet, the virtual iSCSI interface for that port must be in a no shutdown state:
 - ◆ Choose **Switches > Interfaces > Gigabit Ethernet** in Fabric Manager.
 - ◆ Use the **interface** CLI command:

```
interface Gigabit Ethernet 3/1
no shutdown
.
.
.
interface iscsi 3/1
no shutdown
```

- Verify that the IP parameters are correct.
- Verify that the authentication on the Gigabit Ethernet interface (none or chap) matches the authentication configured on the iSCSI initiator.

 **Note:** Configuring authentication at the interface level overrides the global authentication setting.

- Verify that the Gigabit Ethernet switchport parameters are correct (MTU, mode, and so on.).

Performing Basic Dynamic iSCSI Troubleshooting

Use the following guidelines to perform basic dynamic iSCSI troubleshooting:

- Enable dynamic mapping of Fibre Channel targets:
 - ◇ Choose **End Devices > iSCSI**, click the **Initiators** tab and check the **Dynamic** check box in Fabric Manager to allow iSCSI targets to be discovered by the logged-in iSCSI initiators.
 - ◇ Use the **iscsi import target fc** CLI command to allow iSCSI targets to be discovered by the logged-in iSCSI initiators.
- Dynamic iSCSI configuration places all iSCSI initiators logging into the MDS 9000 switch into VSAN 1 by default.
- Any zoning in effect on the default VSAN (VSAN1) will also be applied to iSCSI-connected devices.

Useful Show Commands to Debug Dynamic iSCSI Configuration

The **show** commands in this section are used to debug dynamic iSCSI configuration. The following command output indicates correctly established iSCSI sessions. Run these commands on your switch and compare the output with these samples to help identify possible issues.

- **show iscsi session detail**
- **show iscsi remote-node initiator**
- **show iscsi stats**
- **show iscsi stats detail**
- **show iscsi local-node**
- **show fens data vsan 1**
- **show flogi database vsan 1**

show iscsi session detail Command Output

```
switch# show iscsi session detail
Initiator iqn.1987-05.com.cisco.02.F984BCA7E08C307E2D87A099B2D452F3.FULLMOON (FULLMOON)
  Session #1 (index 2)
    Target iqn.com.domainname.IPS-TEST.02-07.gw.202300a0b80b14da
    VSAN 1, ISID 000000000000, TSID 134, Status active, no reservation
    Type Normal, ExpCmdSN 44, MaxCmdSN 53, Barrier 0
    MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
    DataSeqInOrder No, InitialR2T Yes, ImmediateData No
    Registered LUN 0, Mapped LUN 0
    Stats:
      PDU: Command: 42, Response: 36
      Bytes: TX: 4960, RX: 0
    Number of connection: 1
    Connection #1
      Local IP address: 0xa021ec8, Peer IP address: 0xa021eca
      CID 0, State: LOGGED_IN
      StatSN 43, ExpStatSN 0
      MaxRecvDSLength 524288, our_MaxRecvDSLength 1024
      CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
      AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
      Version Min: 0, Max: 0
      FC target: Up, Reorder PDU: No, Marker send: No (int 0)
      Received MaxRecvDSLen key: Yes
```

show iscsi remote-node initiator Command Output

```
switch# show iscsi remote-node initiator
iSCSI Node name is iqn.1987-05.com.cisco.02.F984BCA7E08C307E2D87A099B2D452F3.FULLMOON
  iSCSI alias name: FULLMOON
  Node WWN is 20:0c:00:0b:be:77:72:42 (dynamic)
  Member of vsans: 1
  Number of Virtual n_ports: 1
  Virtual Port WWN is 20:0d:00:0b:be:77:72:42 (dynamic)
  Interface iSCSI 2/7, Portal group tag: 0x86
  VSAN ID 1, FCID 0x750105
```

show iscsi local-node Command Output

```
switch# show iscsi local-node
target: iqn.com.domainname.IPS-TEST.02-07.gw.202300a0b80b14da
  Port WWN 20:23:00:a0:b8:0b:14:da , VSAN 1
  Auto-created node
```

show fcns data vsan 1 Command Output

```
switch# show fcns data vsan 1
VSAN 1:
```

FCID	TYPE	PWWN	(VENDOR)	FC4-TYPE:FEATURE
0x750000	N	20:23:00:a0:b8:0b:14:da	(SymBios)	scsi-fcp:target
0x750102	N	10:00:00:00:c9:30:ba:06	(Emulex)	scsi-fcp:init
0x750105	N	20:0d:00:0b:be:77:72:42		scsi-fcp:init isc..w
0x750201	N	50:08:05:f3:00:04:96:71		scsi-fcp
0x750301	N	50:08:05:f3:00:04:96:79		scsi-fcp
0x750400	N	20:00:00:02:3d:07:05:c0	(NuSpeed)	scsi-fcp:init

show flogi database vsan 1 Command Output

```
switch# show flogi database vsan 1
```

INTERFACE	VSAN	FCID	PORT NAME	NODE NAME
fc1/1	1	0x750400	20:00:00:02:3d:07:05:c0	10:00:00:02:3d:07:05:c0
fc1/6	1	0x750000	20:23:00:a0:b8:0b:14:da	20:22:00:a0:b8:0b:14:d9
fc1/8	1	0x750102	10:00:00:00:c9:30:ba:06	20:00:00:00:c9:30:ba:06
fc1/9	1	0x750201	50:08:05:f3:00:04:96:71	50:08:05:f3:00:04:96:70
fc1/10	1	0x750301	50:08:05:f3:00:04:96:79	50:08:05:f3:00:04:96:70
iscsi2/7	1	0x750105	20:0d:00:0b:be:77:72:42	20:0c:00:0b:be:77:72:42

Virtual Target Access Control

Use the following guidelines when creating a virtual target:

- Did you specify the correct pWWN?
- If you are creating a virtual target from a subset of LUN(s) of a physical device, did you specify the correct Fibre Channel (physical) LUN(s) and iSCSI (virtual) LUN(s)?
- If using an access list to control access to the virtual target, did you specify the correct initiator(s)? If you are not using an access list to restrict access, did you choose **End Devices > iSCSI**, select the **Targets** tab and check the **Initiator Access All** check box in Fabric Manager or use the **all-initiator-permit** CLI option to insure all initiators have access?
- If restricting access to a particular interface(s), did you specify the correct Gigabit Ethernet interface(s)?

Useful Show Commands to Debug Static iSCSI Configuration

The **show** commands in this section are used to debug static iSCSI configuration. The following command output indicates correctly established iSCSI sessions. Run these commands on your switch and compare the output with these samples to help identify possible issues.

- **show iscsi session detail**
- **show iscsi stats**
- **show iscsi stats detail**
- **show fcns data vsan 5**
- **show flogi data vsan 5**

- show iscsi remote-node iscsi-session-detail tcp-parameters

show iscsi session detail Command Output

```
switch# show iscsi session detail
Initiator iqn.1987-05.com.cisco.02.8cb3c18879bf356ce18e09679103235f.my-kayak (MY-KAYAK)

Session #1 (index 84)
  Target iqn.com.domainname.IPS-TEST.02-08.gw.2200002037c52d6d
  VSAN 5, ISID 00023d000054, TSID 135, Status active, no reservation
  Type Normal, ExpCmdSN 1356, MaxCmdSN 1366, Barrier 0
  MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
  DataSeqInOrder No, InitialR2T Yes, ImmediateData No
  Registered LUN 0, Mapped LUN 0
  Stats:
    PDU: Command: 13, Response: 13
    Bytes: TX: 1344, RX: 0
  Number of connection: 1
  Connection #1
    Local IP address: 0xa011d64, Peer IP address: 0xa011d65
    CID 0, State: LOGGED_IN
    StatSN 1356, ExpStatSN 0
    MaxRecvDSLength 524288, our_MaxRecvDSLength 1392
    CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
    AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
    Version Min: 0, Max: 0
    FC target: Up, Reorder PDU: No, Marker send: No (int 0)
    Received MaxRecvDSLen key: Yes

Session #2 (index 85)
  Target iqn.com.domainname.IPS-TEST.02-08.gw.2200002037c52e2e
  VSAN 5, ISID 00023d000055, TSID 135, Status active, no reservation
  Type Normal, ExpCmdSN 1356, MaxCmdSN 1366, Barrier 0
  MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
  DataSeqInOrder No, InitialR2T Yes, ImmediateData No
  Registered LUN 0, Mapped LUN 0
  Stats:
    PDU: Command: 13, Response: 13
    Bytes: TX: 1344, RX: 0
  Number of connection: 1
  Connection #1
    Local IP address: 0xa011d64, Peer IP address: 0xa011d65
    CID 0, State: LOGGED_IN
    StatSN 1356, ExpStatSN 0
    MaxRecvDSLength 524288, our_MaxRecvDSLength 1392
    CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
    AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
    Version Min: 0, Max: 0
    FC target: Up, Reorder PDU: No, Marker send: No (int 0)
    Received MaxRecvDSLen key: Yes

Session #3 (index 86)
  Target iqn.com.domainname.IPS-TEST.02-08.gw.2200002037c52356
  VSAN 5, ISID 00023d000056, TSID 135, Status active, no reservation
  Type Normal, ExpCmdSN 1356, MaxCmdSN 1366, Barrier 0
  MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
  DataSeqInOrder No, InitialR2T Yes, ImmediateData No
  Registered LUN 0, Mapped LUN 0
  Stats:
    PDU: Command: 13, Response: 13
    Bytes: TX: 1344, RX: 0
  Number of connection: 1
  Connection #1
```

Cisco_MDS_SanOS_Troubleshooting_Guide_-_Troubleshooting_IP_Storage_Services

```
Local IP address: 0xa011d64, Peer IP address: 0xa011d65
CID 0, State: LOGGED_IN
StatsN 1356, ExpStatSN 0
MaxRecvDSLength 524288, our_MaxRecvDSLength 1392
CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
Version Min: 0, Max: 0
FC target: Up, Reorder PDU: No, Marker send: No (int 0)
Received MaxRecvDSLen key: Yes
```

Session #4 (index 87)

```
Target iqn.com.domainname.IPS-TEST.02-08.gw.2200002037c5260a
VSAN 5, ISID 00023d000057, TSID 135, Status active, no reservation
Type Normal, ExpCmdSN 1356, MaxCmdSN 1366, Barrier 0
MaxBurstSize 0, MaxConn 0, DataPDUInOrder No
DataSeqInOrder No, InitialR2T Yes, ImmediateData No
Registered LUN 0, Mapped LUN 0
```

Stats:

```
PDU: Command: 13, Response: 13
Bytes: TX: 1344, RX: 0
```

Number of connection: 1

Connection #1

```
Local IP address: 0xa011d64, Peer IP address: 0xa011d65
CID 0, State: LOGGED_IN
StatsN 1356, ExpStatSN 0
MaxRecvDSLength 524288, our_MaxRecvDSLength 1392
CSG 3, NSG 3, min_pdu_size 48 (w/ data 48)
AuthMethod none, HeaderDigest None (len 0), DataDigest None (len 0)
Version Min: 0, Max: 0
FC target: Up, Reorder PDU: No, Marker send: No (int 0)
Received MaxRecvDSLen key: Yes
```

show iscsi stats Command Output

```
switch# show iscsi stats iscsi2/7
iscsi2/7
  5 minutes input rate 3336 bits/sec, 417 bytes/sec, 0 frames/sec
  5 minutes output rate 120 bits/sec, 15 bytes/sec, 0 frames/sec
  iSCSI statistics
    4112871 packets input, 4022464380 bytes
    303100 Command pdus, 3740086 Data-out pdus, 3815901300 Data-out bytes, 0
  fragments
    1283306 packets output, 778111088 bytes
    303069 Response pdus (with sense 3163), 195108 R2T pdus
    715480 Data-in pdus, 715214528 Data-in bytes
```

show iscsi stats detail Command Output

```
switch# show iscsi stats detail
iscsi2/7
  5 minutes input rate 3336 bits/sec, 417 bytes/sec, 0 frames/sec
  5 minutes output rate 120 bits/sec, 15 bytes/sec, 0 frames/sec
  iSCSI statistics
    4113028 packets input, 4022586092 bytes
    303140 Command pdus, 3740200 Data-out pdus, 3816015476 Data-out bytes, 0
  fragments
    1283382 packets output, 778114736 bytes
    303109 Response pdus (with sense 3163), 195141 R2T pdus
    715480 Data-in pdus, 715214528 Data-in bytes
  iSCSI Forward:
    Command: 303140 PDUs (Received: 303140)
    Data-Out (Write): 3740200 PDUs (Received 3740200), 0 fragments, 3816015476 b
```

show iscsi session detail Command Output


```

ytes
  TMF Request: 0 (Received 28)
  FCP Forward:
    Xfer_rdy: 195141 (Received: 195141)
    Data-In: 715480 (Received: 715622), 715214528 bytes
    Response: 303109 (Received: 303322), with sense 3163
    TMF Resp: 0
  iSCSI Stats:
    Login: attempt: 16726, succeed: 114, fail: 16606, authen fail: 0
    Rcvd: NOP-Out: 36164, Sent: NOP-In: 36160
      NOP-In: 0, Sent: NOP-Out: 0
      TMF-REQ: 28, Sent: TMF-RESP: 0
      Text-REQ: 39, 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
  FCP Stats:
    Total: Sent: 4110679
      Received: 1281518 (Error: 0, Unknown: 0)
    Sent: PLOGI: 66367, Rcvd: PLOGI_ACC: 71, PLOGI_RJT: 66296
      PRLI: 71, Rcvd: PRLI_ACC: 71, PRLI_RJT: 0, Error resp: 0
      LOGO: 0, Rcvd: LOGO_ACC: 0, LOGO_RJT: 0
      ABTS: 87, Rcvd: ABTS_ACC: 0
      TMF REQ: 0
      Self orig command: 213, Rcvd: data: 142, resp: 213
    Rcvd: PLOGI: 614, Sent: PLOGI_ACC: 490
      LOGO: 197, Sent: LOGO_ACC: 111
      PRLI: 0, Sent: PRLI_ACC: 0
      ABTS: 183
  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
      Persistent Resv 0   Data-Out: 0, TMF-Req: 0
  FCP Drop:
    Xfer_rdy: 0, Data-In: 0, Response: 0
  Buffer Stats:
    Buffer less than header size: 48475, Partial: 2524437, Split: 3550971
    Pullup give new buf: 48475, Out of contiguous buf: 0, Unaligned m_data: 0
show fcns database Command Output

```

```
switch# show fcns data vsan 5
```

```
VSAN 5:
```

```

-----
FCID          TYPE  PWWN                                (VENDOR)          FC4-TYPE:FEATURE
-----
0x610002      N     20:0b:00:0b:be:77:72:42              scsi-fcp:init isc..w
0x6101e1      NL    22:00:00:20:37:c5:2d:6d (Seagate)         scsi-fcp:target
0x6101e2      NL    22:00:00:20:37:c5:2e:2e (Seagate)         scsi-fcp:target
0x6101e4      NL    22:00:00:20:37:c5:23:56 (Seagate)         scsi-fcp:target
0x6101e8      NL    22:00:00:20:37:c5:26:0a (Seagate)         scsi-fcp:target
Total number of entries = 5

```

show fcns database Command Output

```
switch# show fcns data vsan 5
```

```
VSAN 5:
```

```

-----
FCID          TYPE  PWWN                                (VENDOR)          FC4-TYPE:FEATURE
-----
0x610002      N     20:0b:00:0b:be:77:72:42              scsi-fcp:init isc..w
0x6101e1      NL    22:00:00:20:37:c5:2d:6d (Seagate)         scsi-fcp:target

```

show iscsi stats detail Command Output

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
0x6101e2    NL    22:00:00:20:37:c5:2e:2e (Seagate)    scsi-fcp:target
0x6101e4    NL    22:00:00:20:37:c5:23:56 (Seagate)    scsi-fcp:target
0x6101e8    NL    22:00:00:20:37:c5:26:0a (Seagate)    scsi-fcp:target
Total number of entries = 5
```

show flogi database Command Output

```
switch# show flogi data vsan 5
```

```
-----
INTERFACE  VSAN    FCID          PORT NAME          NODE NAME
-----
fc1/12     5       0x6101e8     22:00:00:20:37:c5:26:0a  20:00:00:20:37:c5:26:0a
fc1/12     5       0x6101e4     22:00:00:20:37:c5:23:56  20:00:00:20:37:c5:23:56
fc1/12     5       0x6101e2     22:00:00:20:37:c5:2e:2e  20:00:00:20:37:c5:2e:2e
fc1/12     5       0x6101e1     22:00:00:20:37:c5:2d:6d  20:00:00:20:37:c5:2d:6d
iscsi2/8   5       0x610002     20:0b:00:0b:be:77:72:42  20:0a:00:0b:be:77:72:42
```

```
Total number of flogi = 5.
```

show iscsi remote-node iscsi-session-detail tcp-parameters Command Output

```
switch# show iscsi remote-node iscsi-session-detail tcp-parameters
iSCSI Node name is iqn.1987-05.com.cisco.02.8cb3c18879bf356ce18e09679103235f.my-kayak
iSCSI alias name: MY-KAYAK
Node WWN is 20:0a:00:0b:be:77:72:42 (dynamic)
Member of vsans: 5
Number of Virtual n_ports: 1
Virtual Port WWN is 20:0a:00:0b:be:77:72:42 (dynamic)
Interface iSCSI 2/8, Portal group tag is 0x87
  VSAN ID 0, FCID 0x0
  No. of FC sessions: 1
  No. of iSCSI sessions: 1
  iSCSI session details
    Target node:
    Statistics:
      PDU: Command: 0, Response: 0
      Bytes: TX: 0, RX: 0
      Number of connection: 1
    TCP parameters
      Connection Local 10.1.29.100:3260, Remote 10.1.29.101:1026
      Path MTU 1500 bytes
      Current retransmission timeout is 310 ms
      Round trip time: Smoothed 179 ms, Variance: 33
      Advertized window: Current: 62 KB, Maximum: 62 KB, Scale: 0
      Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
      Congestion window: Current: 63 KB
  VSAN ID 5, FCID 0x610002
  No. of FC sessions: 4
  No. of iSCSI sessions: 4
  iSCSI session details
    Target node: iqn.com.domainname.IPS-TEST.02-08.gw.2200002037c5260a
    Statistics:
      PDU: Command: 13, Response: 13
      Bytes: TX: 1344, RX: 0
      Number of connection: 1
    TCP parameters
      Connection Local 10.1.29.100:3260, Remote 10.1.29.101:1048
      Path MTU 1500 bytes
      Current retransmission timeout is 300 ms
```

show fcns database Command Output

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
Round trip time: Smoothed 165 ms, Variance: 35
Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
Congestion window: Current: 63 KB

Target node: iqn.com.domainname.IPS-TEST.02-08.gw.2200002037c5260a
Statistics:
  PDU: Command: 13, Response: 13
  Bytes: TX: 1344, RX: 0
  Number of connection: 1
TCP parameters
  Connection Local 10.1.29.100:3260, Remote 10.1.29.101:1048
  Path MTU 1500 bytes
  Current retransmission timeout is 300 ms
  Round trip time: Smoothed 165 ms, Variance: 35
  Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
  Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
  Congestion window: Current: 63 KB

Target node: iqn.com.domainname.IPS-TEST.02-08.gw.2200002037c5260a
Statistics:
  PDU: Command: 13, Response: 13
  Bytes: TX: 1344, RX: 0
  Number of connection: 1
TCP parameters
  Connection Local 10.1.29.100:3260, Remote 10.1.29.101:1048
  Path MTU 1500 bytes
  Current retransmission timeout is 300 ms
  Round trip time: Smoothed 165 ms, Variance: 35
  Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
  Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
  Congestion window: Current: 63 KB

Target node: iqn.com.domainname.IPS-TEST.02-08.gw.2200002037c5260a
Statistics:
  PDU: Command: 13, Response: 13
  Bytes: TX: 1344, RX: 0
  Number of connection: 1
TCP parameters
  Connection Local 10.1.29.100:3260, Remote 10.1.29.101:1048
  Path MTU 1500 bytes
  Current retransmission timeout is 300 ms
  Round trip time: Smoothed 165 ms, Variance: 35
  Advertized window: Current: 61 KB, Maximum: 62 KB, Scale: 0
  Peer receive window: Current: 63 KB, Maximum: 63 KB, Scale: 0
  Congestion window: Current: 63 KB
```

iSCSI TCP Performance Issues

Generally there are two segments that affect the iSCSI performance. First is the Fibre Channel side flow control mechanism, buffer-to-buffer credits (BB_credits), and Fibre Channel maximum frame size. Second is the TCP/IP side segment.

As in all TCP/IP-related throughput issues, the most important criteria are the Receive/Send Window Sizes on both TCP endpoints, RTT (round-trip time), actual available bandwidth between the TCP peers, the MSS (maximum segment size), and the support for higher MTUs between the peers.

CLI Commands Used to Access Performance Data

Use the following CLI commands to access performance data:


- **show iscsi remote-node iscsi-session-detail tcp-parameters**
- **show ips stats tcp interface gigabitethernet slot/port detail**
- **show interface iscsi slot/port**
- **show interface gigabitethernet slot/port**
- **show interface fc slot/port**
- **show iscsi remote-node fcp-session-detail**

Understanding TCP Parameters for iSCSI

The default MTU size of an Ethernet network is 1500, while the Fibre Channel networks generally support maximum frame sizes of 2148 bytes. This means that an iSCSI gateway must divide the Fibre Channel frames into two TCP segments or IP fragments while transferring from the Fibre Channel side to the IP side depending on how this division is implemented within the device.

This section refers to the scenario in Figure 20-13.

Figure 20-13 IPS Window Scaling



The IPS module adjusts the Receive Data Field Size that it advertises to its Fibre Channel partner, according to the MTU that is configured on the corresponding Gigabit Ethernet port of an iSCSI client.

If left to the default MTU size, the Fibre Channel frame size from the target device is decreased to match the maximum Ethernet frame size, so the switching of the packet through the switch is faster. One point of performance tuning is to increase the MTU of the IP network between the peers.

Jumbo support was enabled for the IPS ports, and the MTU for the VLAN corresponding to these ports was increased.

The second point of performance tuning is to increase the TCP window size of the iSCSI endpoints. Depending on the latency between the iSCSI client and the IPS, this will need fine tuning. The switch's iSCSI configuration defines the TCP window size in kilobytes.

Any value starting with 64 K (> 65535 = 0xFFFF bytes) will automatically trigger TCP window scaling according to [RFC 1323](#). The IPS TCP window scaling begins only when the remote peer (iSCSI client in this case) requests it. This means that you need to configure the TCP stack of your client to trigger this functionality (see Figure 20-13).

For the Fibre Channel side, depending on the direction of the traffic, the BB_credit of the ports corresponding to the input interfaces (sending/receiving traffic to/from the iSCSI side) could be increased, especially in the case of local Gigabit Ethernet attached iSCSI clients.

Lab Setup

This is the lab setup that was used in collecting the performance-related information.

- The server was an IBM Pentium III server: Dual CPU @ 1.13 GHz.
- The TCP window size at both ends was set to 1 MB (1024 K).
- The IBM ESS Shark had a hardcoded BB_credit value of 64 (not configurable).
- The **fcrxbbcredit** on the corresponding switch port (fc1/3) was set to the same value.
- The C4 and C8 represented the corresponding port WWNs (pWWN) for the IBM Shark storage subsystem. The full pWWN is as follows:
 - ◆ C4 50:05:07:63:00:c4:94:4c (in VSAN 778)
 - ◆ C8 50:05:07:63:00:c8:94:4c (in VSAN 777)

Configuring from the Bottom Switch Using the CLI

The following example is the configuration for the 9216 switch shown in Figure 20-13:

```
iscsi initiator name iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas
pWWN 20:05:00:0c:30:6c:24:42
  vsan 777
  vsan 778
iscsi virtual-target name shark_nas
pWWN 50:05:07:63:00:c8:94:4c fc-lun 0000 iscsi-lun 0000 secondary-pwwn
50:05:07:63:00:c4:94:4c
pWWN 50:05:07:63:00:c8:94:4c fc-lun 0001 iscsi-lun 0001 secondary-pwwn
50:05:07:63:00:c4:94:4c
initiator iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas permit
interface GigabitEthernet2/1
ip address 10.48.69.251 255.255.255.192
iscsi authentication none
no shutdown
vrrp 1
priority 110
address 10.48.69.250
(This is the iSCSI target IP address for the Windows iSCSI client.)

no shutdown
interface iscsi2/1
tcp pmtu-enable
tcp window-size 1024
(To increase the receive window size of the IPS module (in kilobytes).)

tcp sack-enable
no shutdown
```

Verifying Connectivity Between Client and IPS iSCSI Service

The following example verifies the connectivity between the client and the IPS iSCSI service:

```
MDS_BOTTOM# show ips stats tcp interface gigabitethernet 2/1
TCP Statistics for port GigabitEthernet2/1
  Connection Stats
    0 active openings, 24 accepts
    0 failed attempts, 0 reset received, 24 established
  Segment stats
    7047380 received, 56080130 sent, 0 retransmitted
    0 bad segments received, 0 reset sent
  TCP Active Connections
    Local Address      Remote Address      State      Send-Q  Recv-Q
    10.48.69.250:3260  10.48.69.233:1026  ESTABLISH  0       0
    10.48.69.250:3260  10.48.69.233:1057  ESTABLISH  34560   0
    0.0.0.0:3260       0.0.0.0:0          LISTEN     0       0
MDS_BOTTOM# show flogi database vsan 777
-----
INTERFACE  VSAN    FCID          PORT NAME          NODE NAME
-----
fc1/3      777     0x610000      50:05:07:63:00:c8:94:4c  50:05:07:63:00:c0:94:4c
iscsi2/1   777     0x610001      20:05:00:0c:30:6c:24:42  20:00:00:0c:30:57:5e:c2
Total number of flogi = 2.
MDS_BOTTOM# show fcns database vsan 777
VSAN 777:
-----
FCID        TYPE  PWWN          (VENDOR)          FC4-TYPE:FEATURE
-----
0x610000    N     50:05:07:63:00:c8:94:4c (IBM)             scsi-fcp:target fc..
0x610001    N     20:05:00:0c:30:6c:24:42             scsi-fcp:init isc..w
Total number of entries = 2
MDS_BOTTOM#
MDS_BOTTOM# 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      DS-X9308-SMIP     ok
Mod  Sw      Hw      World-Wide-Name (s) (WWN)
---  ---
1    1.1(0.133c)  1.0    20:01:00:0c:30:57:5e:c0 to 20:10:00:0c:30:57:5e:c0
2    1.1(0.133c)  0.2    20:41:00:0c:30:57:5e:c0 to 20:48:00:0c:30:57:5e:c0
Mod  MAC-Address(es)          Serial-Num
---  ---
1    00-0b-be-f8-7f-00 to 00-0b-be-f8-7f-04  JAB070804Q3
2    00-05-30-00-a8-56 to 00-05-30-00-a8-62  JAB070205am
* this terminal session
MDS_BOTTOM# show iscsi remote
iSCSI Node name is iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas
iSCSI alias name: SHARK-NAS
Node WWN is 20:00:00:0c:30:57:5e:c2 (dynamic)
Member of vsans: 777, 778
Number of Virtual n_ports: 1
Virtual Port WWN is 20:05:00:0c:30:6c:24:42 (configured)
  Interface iSCSI 2/1, Portal group tag: 0x1001
    VSAN ID 778, FCID 0x7c0000
    VSAN ID 777, FCID 0x610001
MDS_BOTTOM# show iscsi local
target: shark_nas
  Port WWN 50:05:07:63:00:c8:94:4c
(This is the port of the Shark connected to MDS 9216_Bottom.)
```

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

Secondary PWWN 50:05:07:63:00:c4:94:4c
(This is the port of the Shark connected to MDS 9216_Top.)

Configured node

```
No. of LU mapping: 2
  iscsi LUN: 0000, FC LUN: 0000
  iscsi LUN: 0001, FC LUN: 0001
No. of initiators permitted: 1
  initiator iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas is permitted
  all initiator permit is disabled
```

MDS_BOTTOM#

MDS_BOTTOM# show interface iscsi 2/1

```
iscsi2/1 is up
  Hardware is GigabitEthernet
  Port WWN is 20:41:00:0c:30:57:5e:c0
  Admin port mode is ISCSI
  Port mode is ISCSI
  Speed is 1 Gbps
  Number of iSCSI session: 2, Number of TCP connection: 2
  Configured TCP parameters
  Local Port is 3260
  PMTU discover is enabled (default)
```

(This is especially required if there are devices without jumbo support in the path. The initial

```
  Keepalive-timeout 60
  Initial-retransmit-time 300
```

(If there is high delay between the peers, this parameter that can be adjusted. There's no real fo

```
  Max-retransmissions 8
  Window-size 1024000
  Sack is enabled
Forwarding mode: pass-thru
5 minutes input rate 410824 bits/sec, 51353 bytes/sec, 1069 frames/sec
5 minutes output rate 581291520 bits/sec, 72661440 bytes/sec, 53302 frames/sec
iSCSI statistics
  1072393 packets input, 51482588 bytes
  1072305 Command pdus, 0 Data-out pdus, 0 Data-out bytes, 0 fragments
  53430805 packets output, 72837086312 bytes
  1072273 Response pdus (with sense 9), 0 R2T pdus
  52358444 Data-in pdus, 70272402880 Data-in bytes
```

MDS_BOTTOM# show iscsi remote initiator iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas

iscsi tcp

iSCSI Node name is iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas

```
iSCSI alias name: SHARK-NAS
Node WWN is 20:00:00:0c:30:57:5e:c2 (dynamic)
Member of vsans: 777, 778
Number of Virtual n_ports: 1
Virtual Port WWN is 20:00:00:0c:30:57:5e:c2 (configured)
  Interface iSCSI 2/1, Portal group tag is 0x1001
  VSAN ID 0, FCID 0x 0
  No. of FC sessions: 1
  No. of iSCSI sessions: 1
  iSCSI session details
  Target node:
```

```
  Statistics:
```

```
    PDU: Command: 0, Response: 0
    Bytes: TX: 0, RX: 0
    Number of connection: 1
```

```
  TCP parameters
```

```
    Local 10.48.69.250:3260, Remote 10.48.69.233:1026
    Path MTU: 1500 bytes
    Retransmission timeout: 300 ms
    Round trip time: Smoothed 150 ms, Variance: 31
```

Cisco_MDS_SanOS_Troubleshooting_Guide_-_Troubleshooting_IP_Storage_Services

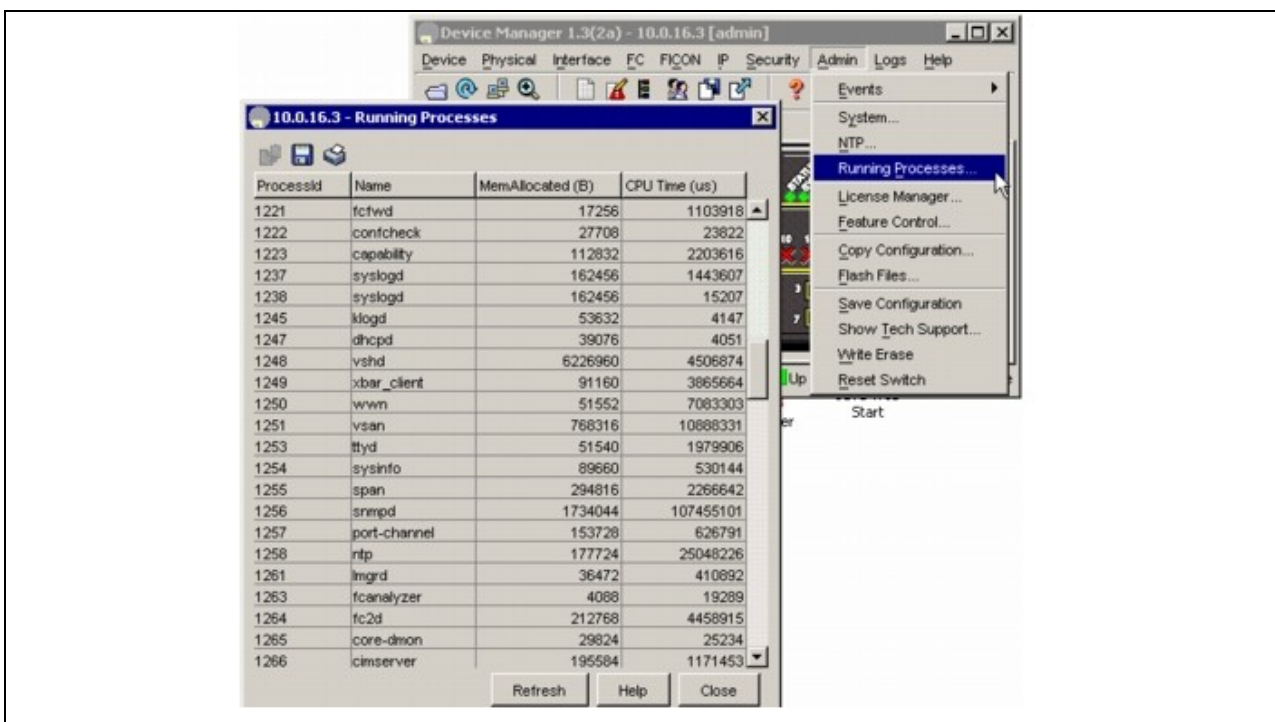
```
Advertized window: Current: 998 KB, Maximum: 1000 KB, Scale: 4
Peer receive window: Current: 1000 KB, Maximum: 1000 KB, Scale: 4
Congestion window: Current: 12 KB
VSAN ID 777, FCID 0x610001
No. of FC sessions: 1
No. of iSCSI sessions: 1
iSCSI session details
  Target node: shark_nas
  Statistics:
    PDU: Command: 392051, Response: 392042
    Bytes: TX: 25692593152, RX: 0
    Number of connection: 1
  TCP parameters
    Local 10.48.69.250:3260, Remote 10.48.69.233:1057
    Path MTU: 1500 bytes
    Retransmission timeout: 300 ms
    Round trip time: Smoothed 2 ms, Variance: 1
```

(Watch out for these numbers. The output is for a TCP session that goes only through one Gigabit E

Advertized window: Current: 1000 KB, Maximum: 1000 KB, Scale: 4
(This is the window size set on the Windows client. See Figure 20-14.)

Peer receive window: Current: 1000 KB, Maximum: 1000 KB, Scale: 4
(This is the window size set on the IPS iSCSI interface. See Figure 20-14.)


Figure 20-14 Congestion Window: Current: 24 kB



TCP Parameter Changes

To change TCP parameters in the Windows registry, use the registry parameters shown in Figure 20-14 as an example.

Setting the Tcp1323Opts (circled in green) to 3, sets two bits on, one for window scaling and the other for the time-stamp option. We are only interested in the window scaling here.

 **Caution:** Editing the registry is a very high risk operation, it can render the system unusable, requiring a reinstallation of the entire operating system. Only advanced users should perform this operation.

Displaying the Gigabit Ethernet Interface

Choose **Switches > Interfaces > Gigabit Ethernet** using Fabric Manager to view the Gigabit Ethernet status.

Also you can use the **show interface** CLI command to view the Gigabit Ethernet status (see Example 20-2).

Example 20-2 Annotated Output of show interface gigabitethernet CLI Command

```
MDS_BOTTOM# show interface gigabitethernet 2/1
GigabitEthernet2/1 is up
  Hardware is GigabitEthernet, address is 0005.3000.a85a
  Internet address is 10.48.69.251/26
  MTU 1500 bytes, BW 1000000 Kbit
Port mode is IPS
  Speed is 1 Gbps
  Beacon is turned off
  5 minutes input rate 3957384 bits/sec, 494673 bytes/sec, 6716 frames/sec
  5 minutes output rate 609420144 bits/sec, 76177518 bytes/sec, 53267 frames/sec
  6979248 packets input, 514206826 bytes
    0 multicast frames, 0 compressed
    0 input errors, 0 frame, 0 overrun 0 fifo
  55551272 packets output, 79456286344 bytes, 0 underruns
    0 output errors, 0 collisions, 0 fifo
    0 carrier errors
```

Better throughput can be achieved if the MTU of both the client NIC and the IPS Gigabit Ethernet interface is changed to a higher MTU, provided the network in the middle supports jumbo frames.

Use the **show ips stats tcp** CLI command to view TCP statistics (see Example 20-3). If the retransmitted value in the segment continues to increase, it shows that either the IP network in the middle has issues or the TCP peer has problems acknowledging the data that the IPS sends to it. If the MTU of this interface is higher than the MSS of the iSCSI client, then the split packets value increases. For example, the client MTU default is 1500, which equates to an MSS value of 1460, but the IPS Gigabit Ethernet MTU changed to 2500.

Example 20-3 show ips stats tcp Command Output

```
MDS_BOTTOM# show ips stats tcp interface gigabit 2/1 detail

TCP Statistics for port GigabitEthernet2/1

  TCP send stats
    56252632 segments, 76746280484 bytes
    56100434 data, 152173 ack only packets
    1 control (SYN/FIN/RST), 0 probes, 24 window updates
    0 segments retransmitted, 0 bytes
  0 retransmitted while on ethernet send queue, 0 packets split
  3 delayed acks sent

  TCP receive stats
    7068115 segments, 1061853 data packets in sequence, 54245464 bytes in sequence
    0 predicted ack, 187 predicted data
    0 bad checksum, 0 multi/broadcast, 0 bad offset
    0 no memory drops, 0 short segments
```

Cisco_MDS_SanOS_Troubleshooting_Guide_--_Troubleshooting_IP_Storage_Services

```
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
7067879 acks, 76746255713 ack bytes, 0 ack toomuch, 21 duplicate acks
0 ack packets left of snd_una, 0 non-4 byte aligned packets
5980106 window updates, 0 window probe
50 pcb hash miss, 0 no port, 0 bad SYN, 0 paws drops

TCP Connection Stats
0 attempts, 24 accepts, 24 established
22 closed, 2 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
7054414 segments timed, 7067879 rtt updated
0 retransmit timeout, 0 persist timeout
19 keepalive timeout, 19 keepalive probes

TCP SACK Stats
0 recovery episodes, 54218621 data packets, 77791012992 data bytes
0 data packets retransmitted, 0 data bytes retransmitted
1 connections closed, 0 retransmit timeouts

TCP SYN Cache Stats
24 entries, 24 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, 2 no-route SYN drop
0 hash collisions, 0 retransmitted

TCP Active Connections
Local Address      Remote Address      State      Send-Q   Recv-Q
10.48.69.250:3260  10.48.69.233:1026  ESTABLISH  0        0
10.48.69.250:3260  10.48.69.233:1057  ESTABLISH  29296    0
0.0.0.0:3260       0.0.0.0:0          LISTEN     0        0
```

Use the **show iscsi remote-node fcp-session-detail** CLI command to view details of the session status. (See Example 20-4.) The RcvDataFieldSize will be set to the maximum 2048 if the MTU is increased on the Gigabit Ethernet interface that corresponds to this iSCSI remote node. Use the Target FCID field to verify that the local port, rather than a remote port that is reached through an ISL link, is used for the storage target to avoid suboptimal access to storage.

Example 20-4 show iscsi remote-node fcp-session-detail Command Output

```
MDS_BOTTOM# show iscsi remote-node fcp-session-detail
iSCSI Node name is iqn.1987-05.com.cisco:02.75af2f95624c.shark-nas
iSCSI alias name: SHARK-NAS
Node WWN is 20:00:00:0c:30:6c:24:42 (dynamic)
Member of vsans: 777, 778
Number of Virtual n_ports: 1

Virtual Port WWN is 20:00:00:0c:30:6c:24:42 (configured)
Interface iSCSI 2/1, Portal group tag is 0x1001
VSAN ID 0, FCID 0x610001
No. of FC sessions: 1
No. of iSCSI sessions: 1
```

FCP Session details

```

Target FCID: 0x000000 (S_ID of this session: 0x000000)
    pWWN: 00:00:00:00:00:00:00:00
    nWWN: 00:00:00:00:00:00:00:00
Session state: INIT
1 iSCSI sessions share this FC session
  Target:
Negotiated parameters
  RcvDataFieldSize 2048 our_RcvDataFieldSize 1392
  MaxBurstSize 0, EMPD: FALSE
  Random Relative Offset: FALSE, Sequence-in-order: Yes
Statistics:
  PDU: Command: 0, Response: 0
VSAN ID 777, FCID 0x610001
No. of FC sessions: 1
No. of iSCSI sessions: 1
FCP Session details

```

```

Target FCID: 0x610000 (S_ID of this session: 0x610001)
    pWWN: 50:05:07:63:00:c8:94:4c
    nWWN: 50:05:07:63:00:c8:94:4c
Session state: LOGGED_IN
1 iSCSI sessions share this FC session
  Target: shark_nas
Negotiated parameters
  RcvDataFieldSize 2048 our_RcvDataFieldSize 1392
  MaxBurstSize 0, EMPD: FALSE
  Random Relative Offset: FALSE, Sequence-in-order: Yes
Statistics:
  PDU: Command: 0, Response: 1612007

```

Verifying that the Host Is Configured for High MTU or MSS with the CLI

To get the real benefit of an increased MTU and higher Fibre Channel frame size, the path between the iSCSI client and the IPS iSCSI interface (as well as the host NIC) has to be capable of supporting this high MTU.

If you do not have access to the host, one way to see if the host is also configured for high MTU/MSS (as well as the path in the middle) is to check the split packets field in the **show ips stats tcp** display.

However this is a generic display for all TCP sessions. That is, if you have some hosts with high MTU-capable NICs, and some others without, it may be difficult to assess which is which. (See Example 20-5.)

Example 20-5 Sample Output for Low Packet Split Count

```

MDS_Top# show ips stats tcp interface gigabitethernet 2/1 detail (truncated output)
TCP Statistics for port GigabitEthernet2/1
  TCP send stats
    10 segments, 240 bytes
    5 data, 5 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
  ...

  TCP Active Connections
    Local Address      Remote Address      State      Send-Q  Recv-Q
    10.48.69.250:3260  10.48.69.233:1026  ESTABLISH  0       0
    10.48.69.250:3260  10.48.69.233:1040  ESTABLISH  0       0
    0.0.0.0:3260      0.0.0.0:0         LISTEN    0       0

```

Afterward, traffic starts flowing from the FC storage towards the server that is connected via iSCSI to the IPS.

Example 20-6 Sample Output for Large Packet Split Count

```
MDS_Top# show ips stats tcp interface gigabitethernet 2/1 detail
TCP Statistics for port GigabitEthernet2/1
  TCP send stats
    715535 segments, 943511612 bytes
    712704 data, 2831 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, 345477 packets split
...

```

iSLB Issues

This section describes common troubleshooting issues for iSLB and includes the following topics:

- iSLB Configuration Not Distributed to All Switches in the Fabric
- iSCSI Initiator and Virtual Target Configuration Not Distributed
- iSLB Configuration, Commit, or Merge Failed?"VSAN ID is Not Yet Configured"
- iSLB Configuration, Commit, or Merge Failed?"Failed to Allocate WWN"
- iSLB Configuration, Commit, or Merge Failed?"Duplicate WWN Found as..."
- iSLB Configuration, Commit, or Merge Failed?"Duplicate Node Name"
- iSLB Configuration Failed?"Pending iSLB CFS Config Has Reached Its Limit..."
- iSCSI Disable Failed?"Cannot Disable Iscsi - Large Iscsi Config Present..."
- iSLB Commit Timeout
- Session Down?"pWWN in Use At Remote Switch"
- Redirected Session Does Not Come Up
- iSLB Zones Not Present in Active Zone Set
- Traffic Description After iSLB Commit or Activation of Zone Set
- VRRP Master Overutilized
- iSLB Zone Set Activation Failed
- iSLB CFS Commit Fails
- Resolving an iSLB Merge Failure

iSLB Configuration Not Distributed to All Switches in the Fabric

Symptom iSLB configuration is not distributed to all switches in the fabric.

Table 20-1 iSLB Configuration Not Distributed to All Switches in the Fabric

Symptom	Possible Cause	Solution
iSLB configuration not distributed to all switches in the fabric.	Not all switches are running Cisco SAN-OS Release 3.0(1) or later.	Update switches to Cisco SAN-OS Release 3.0(1) or later.
	CFS distribution is not enabled for	Enable CFS distribution for iSLB. Use the show cfs application name islb CLI

	iSLB.	<p>command to determine if CFS distribution is enabled. Or use the show cfs peers name islb CLI command and check to see if any switches are missing from the output.</p> <p>Use the islb distribute CLI command to enable CFS distribution.</p>
--	-------	--

iSCSI Initiator and Virtual Target Configuration Not Distributed

Symptom iSCSI initiator and virtual target configuration is not distributed to the fabric.

Table 20-2 iSCSI Initiator and Virtual Target Configuration Not Distributed

Symptom	Possible Cause	Solution
iSCSI initiator and virtual target configuration is not distributed to the fabric.	Normal operation.	<p>Only the following iSCSI and iSLB configuration is distributed:</p> <ul style="list-style-type: none"> • iSLB initiator and iSLB initiator targets • iSLB VRRP load-balancing configuration • iSCSI global authentication parameters (authentication algorithm and CHAP user name or password) • iSCSI dynamic initiator mode (iSCSI, iSLB, or deny)

iSLB Configuration, Commit, or Merge Failed--"VSAN ID is Not Yet Configured"

Symptom iSLB configuration, commit, or merge failed with error VSAN ID is not yet configured.

Table 20-3 iSLB Configuration, Commit, or Merge Failed--"VSAN ID is Not Yet Configured"

Symptom	Possible Cause	Solution
iSLB configuration, commit, or merge failed with error VSAN ID is not yet configured.	The VSAN ID for one of the initiators is not configured on all switches in the fabric.	<p>Check the output of the show islb cfs-session status, show islb merge status, and show ips internal event-history error CLI command for details on which initiator VSAN ID is not configured on a switch.</p> <p>Use the vsan database vsan vsan-id CLI command to add the VSAN ID, or remove the VSAN ID from the</p>

		initiator configuration.
--	--	--------------------------

iSLB Configuration, Commit, or Merge Failed--"Failed to Allocate WWN"

Symptom iSLB configuration, commit, or merge failed with error Failed to allocate WWN.

Table 20-4 iSLB Configuration, Commit, or Merge Failed--"Failed to Allocate WWN"

Symptom	Possible Cause	Solution
iSLB configuration, commit, or merge failed with error Failed to allocate WWN.	The pWWN or nWWN for one for the initiators could not be reserved from the WWN manager. This implies that the particular WWN is already in use.	<p>Check the output of the show islb cfs-session status, show islb merge status and show ips internal event-history error CLI commands for details on the specific WWN and initiator in error.</p> <p>To fix the problem, use another WWN or allow the system to assign one for the initiator using the static [nWWN pWWN] system-assign command.</p>

iSLB Configuration, Commit, or Merge Failed--"Duplicate WWN Found as..."

Symptom iSLB configuration, commit, or merge failed with error Duplicate WWN found as ..

Table 20-5 iSLB Configuration, Commit, or Merge Failed--"Duplicate WWN Found as ..."

Symptom	Possible Cause	Solution
iSLB configuration, commit, or merge failed with error Duplicate WWN found as ..	The pWWN or nWWN for one for the initiators is already in use by another initiator.	<p>Check the output of the show islb cfs-session status, show islb merge status and show ips internal event-history error CLI commands for details on the specific WWN and initiator in error.</p> <p>To fix the problem, use another WWN or allow the system to assign one for the initiator using the static [nWWN pWWN] system-assign command.</p>

iSLB Configuration, Commit, or Merge Failed--"Duplicate Node Name"

Symptom iSLB configuration, commit, or merge failed with error Duplicate node name.

- align="left" valign="bottom"

Symptom	Possible Cause	Solution
iSLB configuration, commit, or merge failed with error Duplicate node name.	Node name of one of the iSLB initiators is the same as an existing iSCSI initiator.	Check the output of the show islb cfs-session status , show islb merge status and show ips internal event-history error CLI commands for details on the specific initiator in error. To fix the problem, use a different node name.

iSLB Configuration Failed--"Pending iSLB CFS Config Has Reached Its Limit..."

Symptom iSLB configuration failed with error Pending iSLB CFS config has reached its limit..

Table 20-7 iSLB Configuration Failed--"Pending iSLB CFS Config Has Reached Its Limit..."

Symptom	Possible Cause	Solution
iSLB configuration failed with error Pending iSLB CFS config has reached its limit.	The limit of 200 initiators in the pending database has been reached, so no more configuration is allowed	Use the islb commit CLI command to commit the outstanding changes.

iSCSI Disable Failed--"Cannot Disable Iscsi - Large Iscsi Config Present..."

Symptom iSCSI disable failed with error Cannot disable iSCSI - large iSCSI config present.

Table 20-8 iSCSI Disable Failed--"Cannot Disable Iscsi - Large Iscsi Config Present..."

Symptom	Possible Cause	Solution
iSCSI disable failed with error Cannot disable iSCSI - large iSCSI config present..	There are more than 200 initiators in the running config, so iSCSI disable is not allowed.	Delete initiators from the configuration until you have less than 200 initiators. Then use the no iscsi enable CLI command to disable iSCSI.

iSLB Commit Timeout

Symptom iSLB commit timeout.

Table 20-9 iSLB Commit Timeout

iSLB Configuration, Commit, or Merge Failed--"Duplicate Node Name"

Symptom	Possible Cause	Solution
iSLB commit timeout.	When a large configuration is present, it is possible for the iSLB commit to take a long time.	Check the output of the show islb cfs-session status CLI command to get the status of the commit.

Session Down--"pWWN in Use At Remote Switch"

Symptom Session down with error pWWN in use at remote switch.

Table 20-10 Session Down--"pWWN in Use At Remote Switch"

Symptom	Possible Cause	Solution
Session down with error pWWN in use at remote switch.	An initiator pWWN can be used only once in the fabric.	<p>If the same initiator tries to log in to two iSCSI ports at the same time, both ports will initially allow the sessions to come up and then try to reserve the pWWN in the fabric. If it is detected that this pWWN is already in use, the session will be destroyed.</p> <p>To fix the problem, use another WWN or allow the system to assign one for the initiator using the static [nWWN pWWN] system-assign CLI command.</p>

Redirected Session Does Not Come Up

Symptom Redirected session does not come up.

Table 20-11 Redirected Session Does Not Come Up

Symptom	Possible Cause	Solution
Redirected session does not come up.	Connection may be down, or initiator to interface mapping may be missing.	<p>Use the ping CLI command to verify that the connection between the redirected port and initiator is up.</p> <p>Use the show logging logfile CLI command to check the system messages to determine what the session creation failure reason is if any.</p> <p>Use the show interface brief CLI command to verify that iSCSI and Gigabit Ethernet interfaces are up.</p> <p>Bring down and then bring up the initiator and try again to see if the error is persistent.</p>

		<p>There are times initiators do not attempt to make connections to the redirected interface.</p> <p>Use the debug ips islb vrrp flow CLI command to check if the redirection is performing correctly.</p> <p>Use the show islb vrrp summary CLI command to see if the initiator to the interface mapping is set up.</p>
--	--	--

iSLB Zones Not Present in Active Zone Set

Symptom iSLB zones not present in active zone set.

Table 20-12 iSLB Zones Not Present in Active Zone Set

Symptom	Possible Cause	Solution
iSLB zones not present in active zone set.	Active zone set not configured.	Check if an active zone set is configured. If this is not the case then create and activate a new zone set for the VSAN in question. Then use the islb zoneset activate CLI command to trigger iSLB zoning.
	Zone set activation failed.	If an active zone set is configured, then check for activation failures. See the "Traffic Description After iSLB Commit or Activation of Zone Set" section.

Traffic Description After iSLB Commit or Activation of Zone Set

Symptom Traffic description after iSLB commit or activation of zone set (normal, IVR, or iSLB).

Table 20-13 Traffic Disruption After iSLB Commit or Activation of Zone Set

Symptom	Possible Cause	Solution
Traffic description after iSLB commit or activation of zone set (normal, IVR, or iSLB).	An iSLB commit must be done from the switch that has IVR configured.	Commit the iSLB configuration from a switch that has both IVR and iSLB enabled. Use the islb commit CLI command.
	Any zone set activation (normal, iSLB, or IVR) must be done from the switch that has IVR configured.	Activate the zone set from a switch that has both IVR and iSLB enabled.

Use the islb zoneset activate CLI command.

VRRP Master Overutilized

Symptom VRRP master is overutilized.

Table 20-14 VRRP Master Overutilized

Symptom	Possible Cause	Solution
VRRP master is overutilized.	iSCSI interface parameters do not match the rest of the interfaces in the VRRP group.	Verify the interface parameters for all interfaces in the VRRP group. Or use the show vrrp CLI command to view which interfaces are in the VRRP group, then use the show interface iscsi CLI command.
	Load metric needs to be adjusted.	Raise the load metric. Use the metric CLI command in iSLB configuration mode. The default value is 1000.

iSLB Zone Set Activation Failed

Symptom iSLB zone set activation failed.

Table 20-15 iSLB Zone Set Activation Failed

Symptom	Possible Cause	Solution
iSLB zone set activation failed.	iSLB auto-zone is enabled but CFS distribution is not enabled.	Enable CFS distribution for iSLB to share load across multiple switches. Use the islb distribute CLI command on each switch in the fabric.
	Zone set activation is not from switch with IVR and iSLB enabled.	Activate the zone set from a switch that has IVR and iSLB enabled. Use the islb zoneset activate CLI command.
	Another zone set activation is in progress.	Only one zoning related action can occur at the same time (zone, IVR zone, or iSLB zone configuration or activation). Wait until the zone set activation completes and then retry the iSLB zone set activation.
	Zoning database is locked because a configuration is	Only one zoning related action can occur at the same time (zone, IVR zone, or iSLB zone configuration or activation). Commit the

	pending.	<p>existing configuration change or discard the changes.</p> <p>Use the show islb status, islb commit, or islb abort CLI command to view the status, to commit the changes or to discard the changes, respectively.</p> <p>Also, verify that no zone or IVR zone changes are pending.</p>
--	----------	--

iSLB CFS Commit Fails

Symptom iSLB CFS commit fails.

Table 20-16 iSLB CFS Commit Fails


Symptom	Possible Cause	Solution
iSLB CFS commit fails.	Zone set activation is not from the switch with IVR and iSLB enabled.	<p>Activate the zone set from a switch that has IVR and iSLB enabled.</p> <p>Use the islb commit CLI command.</p>
	Another zone set activation is in progress.	Wait until the zone set activation completes and then retry the iSLB zone set activation.
	Zoning database is locked because a configuration is pending.	<p>Commit the existing configuration change or discard the changes.</p> <p>Use the show islb status, islb commit, or islb abort CLI command to view the status, to commit the changes or use islb abort to discard the changes.</p> <p>Also, verify that no zone or IVR zone changes are pending.</p>

Resolving an iSLB Merge Failure

To resolve an iSLB merge failure using the CLI, follow these steps:

1. Determine the cause of merge failure using the output of the **show islb merge status** and the **show ips internal event-history error** commands.
2. If the reason for the merge failure is the VSAN configuration, configure the VSAN on all the switches.

3. Log in to the switch in the fabric whose running configuration you want to keep and issue the **islb commit** command.

 **Note:** The iSLB configuration on other switches will be overwritten. A commit after a merge failure synchronizes the fabric configuration to the running- config of the switch where the commit was performed.

Back to Main Page: [Cisco MDS SAN-OS Troubleshooting Guide](#)
