



Troubleshooting SAN Switching Issues

A storage area network (SAN) is a network of storage devices that provide data storage for servers.

This chapter describes how to identify and resolve problems that can occur with a SAN and the Cisco Nexus 5000 Series switch.

This chapter includes the following sections:

- [Overview](#)
- [NPV](#)
- [Zoning](#)
- [SAN Port Channels](#)
- [FC Services](#)
- [Cisco Fabric Services](#)
- [VSANs](#)
- [Registers and Counters](#)

Overview

The two most common symptoms of problems in a storage network are:

- A host cannot access its allocated storage.
- An application does not respond after attempting to access the allocated storage.

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

Consider the following questions to determine the status of your installation:

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

REVIEW DRAFT – CISCO CONFIDENTIAL**General SAN troubleshooting steps**

-
- Step 1** Obtain information on problems in your fabric.
 - Step 2** Verify physical connectivity between your switches and end devices.
 - Step 3** Verify registration to your fabric for all SAN elements.
 - Step 4** Verify the configuration for your end devices (storage subsystems and servers).
 - Step 5** Verify end-to-end connectivity and fabric configuration.
-

NPV**NP Uplink ports on NPV edge switch are stuck in initializing state**

NP uplink ports connected to the core NPIV switch do not come online and are stuck in an initializing state.

Possible Cause

The core switch might not have been enabled for NPIV.

Example:

```
switch(config-if)# sh int fc2/2
fc2/2 is down (Initializing)
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Port WWN is 20:42:00:0d:ec:a4:3b:80
Admin port mode is NP, trunk mode is on
```

Solution

- Check the status of the NPV external interfaces.
Check that the NPIV is enabled on the core switch.

Example:

```
switch(config-if)# sh npv status
npiv is disabled
disruptive load balancing is disabled
External Interfaces:
=====
Interface: fc2/1, State: Failed(NPIV is not enabled in upstream switch)
Interface: fc2/2, State: Failed(NPIV is not enabled in upstream switch)
Interface: san-port-channel 200, State: Down
```

- If NPIV is disabled, then enable NPIV on the core switch.

Example:

```
switch(config)# feature npiv
```

REVIEW DRAFT – CISCO CONFIDENTIAL**Server interface does not come up and “NPV upstream port not available” message appears**

A server port connected to the NPV edge switch does not come online, and the **show interface** command indicates a status of NPV upstream port not available.

Possible Cause

The upstream NP_Port(s) and the downstream server F_Port(s) on the NPV edge switch may not be in the same VSAN.

Example:

```
switch# sh int fc2/7
fc2/7 is down (NPV upstream port not available)
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Port WWN is 20:47:00:0d:ec:a4:3b:80
Admin port mode is F, trunk mode is off
snmp link state traps are enabled
Port vsan is 99
Receive data field Size is 2112
```

Solution

- Check the VSAN membership of the upstream port and the server port.

Example:

```
switch# show vsan membership
vsan 1 interfaces:
fc2/1 fc2/2 fc2/3 fc2/4
fc2/5 fc2/6 san-port-channel 200
vsan 99 interfaces:
fc2/7 fc2/8
```

- In the example above, notice that the upstream ports (fc2/1-2) are in VSAN 1, and the server ports (fc2/7-8) are in VSAN 99.

Move the NP ports on the NPV edge, and the F ports on the NPIV core into the same VSAN as the server ports.

Example:

```
switch(config)# vsan database
switch(config-vsan-db)# vsan 99 interface fc2/1-2
switch(config-if)# vsan database
switch(config-vsan-db)# vsan 99 interface fc1/17-18
Traffic on fc1/17 may be impacted. Do you want to continue? (y/n) y
Traffic on fc1/18 may be impacted. Do you want to continue? (y/n) y
```

**Note**

Alternatively, if the NPIV core and NPV edge switch are F_Port Trunking capable switches, then that would be the recommended configuration.

Uneven load balancing on the NPV NP ports

An examination of NP upstream ports that are members in the same VSAN reveals that uneven load balancing is occurring.

Possible Cause

REVIEW DRAFT – CISCO CONFIDENTIAL

This may be normal and a direct result of the default SID/DID load balancing that is done before the Nexus 5000 4.2(1)N1 release.

Solution

If the upstream switch is an MDS switch that is running 4.1(3) code or above, and it is a NPV F_Port Trunking capable switch, the preferred configuration would be to run the F_Port Trunking Port Channeling feature.

Example (NPIV core):

```
pod3-9222i(config)# feature npiv
pod3-9222i(config)# feature fport-channel-trunk

pod3-9222i(config)# interface port-channel 1
pod3-9222i(config-if)# switchport mode f
pod3-9222i(config-if)# switchport trunk mode on
pod3-9222i(config-if)# channel mode active
pod3-9222i(config-if)# interface fc2/13, fc2/19
pod3-9222i(config-if)# switchport mode f
pod3-9222i(config-if)# switchport rate-mode dedicated
pod3-9222i(config-if)# switchport trunk mode on
pod3-9222i(config-if)# channel-group 1 force
```

In this example, fc2/13 and fc2/19 are added to port channel 100 and are disabled. Do the same operation on the switch at the other end of the port channel, then do no shutdown at both ends to bring them up.

Example:

```
pod3-9222i(config-if)# no shut
```

Example (NPV Edge):

```
pod7-5020-51(config)# interface san-port-channel 1
pod7-5020-51(config-if)# switchport mode np
pod7-5020-51(config-if)# switchport trunk mode on
pod7-5020-51(config-if)# interface fc2/1-2
pod7-5020-51(config-if)# switchport mode np
pod7-5020-51(config-if)# switchport trunk mode on
pod7-5020-51(config-if)# channel-group 1
```

In this example, fc2/1 and fc2/2 are added to port channel 1 and are disabled. Do the same operation on the switch at the other end of the port channel, then do no shutdown at both ends to bring them up

Example:

```
pod7-5020-51(config-if)# no shut
```

Server on downstream NPV edge switch does not login to the fabric

The server connected to the downstream NPV edge switch does not log in to the fabric.

Possible Cause

The server on the downstream NPV edge switch does not log in to the fabric, and/or you see a “waiting for FLOGI” message.

Example:

```
switch# show npv status
npiv is enabled
```

REVIEW DRAFT – CISCO CONFIDENTIAL

Server Interfaces:

=====

Interface: fc1/6, VSAN: 1, NPIV: No, State: Waiting for FLOGI

Solution

- Verify the configuration of both the NPV edge and core switches. If you are not running the F_Port trunking feature, then verify that there are no VSAN mismatches and that the server ports, NPV NP ports, NPIV Core F_Ports, and storage ports are all in the same VSAN and all are online.
- If the configuration is correct and you can determine where the problem might be, you can collect an Ethalyzer trace and verify that the Fabric Login (FLOGI) frame is being received and sent to the NPIV core as a Fabric Discovery (FDISC) command.

Example Ethalyzer trace:

```
switch# ethanalyzer local sniff-interface inbound-hi display-filter "!!lc && !stp"
limit-captured-frames 0 write bootflash:npv-trace
Capturing on eth4
```

- Recreate the problem by flapping the NPV-attached server port. The trace will be written to bootflash and can be copied off the switch by using the following:

copy bootflash: ftp:

- After the trace has been copied, you can now open and verify the flow using Wireshark.

Example normal NPV login flow:

```
Server -----> FLOGI -----> NPV Edge Switch Fabric Login frame =
FLOGI

NPV Edge Switch -----> FDISC -----> NPIV Core Switch Fabric DISCOVERY
frame maps parameters
from Server FLOGI

NPV Core Switch -----> Accept -----> NPV Edge Switch NPIV Core assigns an
FCID with the Accept
to the FDISC from NPV
Edge Switch

NPV Edge Switch -----> Accept -----> Server Accept to original
Server FLOGI with FCID
assigned from NPIV
Core Switch
```

Locating exact port that server is physically attached to

NPIV switches lose visibility into the physical port that a downstream NPV-connected server is attached to. The following process can be used to identify that physical port.

Possible Cause

When you have an NPIV core switch that has several downstream NPV edge switches attached, you might want to locate the exact port that a server is physically attached to.

REVIEW DRAFT – CISCO CONFIDENTIAL**Solution**

- Identify the PWWN of the server and the corresponding switch that it is attached to.

Example:

```
NPV-Core(config-if)# show flogi database
```

```
fc1/16 100 0xee00e4 21:00:00:04:cf:17:66:b7 20:00:00:04:cf:17:66:b7
fc1/16 100 0xee00e8 21:00:00:04:cf:17:66:0e 20:00:00:04:cf:17:66:0e
fc1/25 100 0xee0100 20:41:00:0d:ec:a3:da:40 20:64:00:0d:ec:a3:da:41
fc1/26 100 0xee0200 20:42:00:0d:ec:a3:da:40 20:64:00:0d:ec:a3:da:41
fc1/26 100 0xee0201 21:00:00:c0:dd:12:04:f3 20:00:00:c0:dd:12:04:f3
```

In the example, the server is identified by this address:

```
fc1/26 100 0xee0201 21:00:00:c0:dd:12:04:f3 20:00:00:c0:dd:12:04:f3
```

and the switch is identified by this address

```
fc1/26 100 0xee0200 20:42:00:0d:ec:a3:da:40 20:64:00:0d:ec:a3:da:41
```

- Identify the IP address of the NPV edge switch.

Example:

```
NPV-Core(config-if)# sh fcns database npv
VSAN 100:
```

```
20:64:00:0d:ec:a3:da:41 172.18.217.51 fc2/1 20:00:00:0d:ec:51:0c:00 fc1/25
20:64:00:0d:ec:a3:da:41 172.18.217.51 fc2/2 20:00:00:0d:ec:51:0c:00 fc1/26
```

- Telnet to the NPV edge switch.

Example:

```
NPV-Core(config-if)# telnet 172.18.217.51
```

- Identify the PWWN of the server.

Example:

```
switch-NPV-Edge# show npv flogi-table
```

```
vfc3 100 0xee0201 21:00:00:c0:dd:12:04:f3 20:00:00:c0:dd:12:04:f3 fc2/2
```

- If the interface is a FCoE (VFC) interface as shown in the previous example, use the show interface vfc3 command to see which port that the VFC is physically bound to.

VSANs stuck in initializing state after configuring the 4.2(1)N1 F_Port trunking feature

Using the **show interface** command of the F_Port trunking port channel or trunking member of the port channel indicates that certain VSANs are in an initializing state and do not come online.

Possible Cause

After configuring the 4.2(1)N1 F_Port Trunking feature, VSANs on the trunk ports appear to be stuck in an initializing state.

REVIEW DRAFT – CISCO CONFIDENTIAL

Example:

```
switch(config-if)# sh int fc2/1
fc2/1 is trunking
Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
Port WWN is 20:41:00:0d:ec:a4:3b:80
Admin port mode is NP, trunk mode is on
snmp link state traps are enabled
Port mode is TNP
Port vsan is 1
Speed is 4 Gbps
Transmit B2B Credit is 16
Receive B2B Credit is 16
Receive data field Size is 2112
Beacon is turned off
Belongs to san-port-channel 200
Trunk vsans (admin allowed and active) (1,99,200)
Trunk vsans (up) (1,99)
Trunk vsans (isolated) ()
Trunk vsans (initializing) (200)
```

Under the Trunk Failures tab of Fabric Manager, you might also see the trunk VSAN listed. However, this may be normal. If no downstream devices are logged in for a particular VSAN, that VSAN stays in initializing state.

Solution

For the VSANs that you are working with, verify by using the following command:

Example:

```
switch# show npv flogi-table

fc2/7 99 0xba0002 10:00:00:00:00:02:00:00 10:00:00:00:00:00:02:00 Spo200
fc2/8 99 0xba0003 10:00:00:00:00:01:00:00 10:00:00:00:00:00:01:00 Spo200
Total number of flogi = 2.
```

In this example, no devices are logged into VSAN 200.

Zoning

Cannot activate zoneset and cannot configure zoning in enhanced zoning mode

The zone set cannot be activated and zoning cannot be configured in enhanced zoning mode. The error message “Zoning database update in progress, command rejected” might be received.

Possible Cause

Another user on the same switch or on a different switch is holding the enhanced zoning configuration lock.

Solution

Release the zoning lock with the following:

-
- Step 1** Determine which switch (domain/ip address) has the lock.
 - Step 2** Determine which user has the lock on that switch.

REVIEW DRAFT – CISCO CONFIDENTIAL

Step 3 Clear the lock for that user on that switch.

- On the same switch, enter the **show zone status vsan** *<vsan-id>* command to determine which user holds the lock.

Example:

```
switch1# show zone status vsan 200
VSAN: 200 default-zone: deny distribute: active only Interop: default
mode: enhanced merge-control: allow
session: remote [dom: 121][ip: 171.165.98.20] <<==
```

In this example the remote switch with the IP address of 171.165.98.20 has the lock.

- Connect to the remote switch and enter the **show zone status vsan** command.

Example:

```
switch2# show zone status vsan 200

VSAN: 200 default-zone: deny distribute: active only Interop: default
mode: enhanced merge-control: allow
session: cli [remi] <<==
```

In the example, user Remi is holding the enhanced zoning lock.

- On the remote switch (N5K2 in the example), release the lock with the **no zone commit vsan** *<vsan-id>* command.
- To confirm that the lock had been cleared, enter the **show zone status vsan** *<vsan-id>* command. At this point, the session parameter should appear as none.
- If the lock still persists, remove the lock from the switch that holds the lock with the **clear zone lock** command.
- If the lock continues to persist, use the following commands to collect information to aid further analysis:

```
show zone internal vsan <vsan-id>
show zone status vsan <vsan-id>
show fcdomain domain-list vsan <vsan-id>
show users
show tech-support zone
show tech-support device-alias
show logging
```

Host cannot communicate with storage

In initial SAN deployments or after topology changes in the SAN, some hosts might not be able to communicate with storage. The initiator cannot access the LUNs that were allocated for them in the storage array.

Possible Cause

If the host and storage are connected to two different switches, the ISL link, (the xE port connecting both switches) might be isolated.

The xE port might be isolated in a specific VSAN for possible reasons:

- Misconfigured fabric timers

REVIEW DRAFT – CISCO CONFIDENTIAL

- Misconfigured port parameters
- Mismatched zoning

Solution

To resolve the VSAN isolation on the TE port:

- Use the **show interface fc <slot/port>** command on the TE port to determine the VSAN number. The isolated VSAN number must match the VSAN number where the host and the storage are connected to. In the display output, you see the Trunk vsans (isolated) (**Vsan <vsan-id>**).
- Use the **show port internal info interface fc <slot/port>** command to determine the root cause of the VSAN isolation.

Possible Cause

Host and storage are not in the same VSAN.

Solution

- Use the **show vsan membership** command to verify that both the host and the storage are in the same VSAN.
- If the host and the storage are in different VSANs, in the configuration mode use the commands **vsan database** and **vsan <vsan-id> interface fc <slot/port>** to move the interface connected to the host and storage devices into the same VSAN.

Possible Cause

The host and storage are not in the same zone. The zone is not in the active zone set. There is no active zone set and default zone policy is set to deny.

Solution

- Use the command **show zone status <vsan-id>** to determine if the default zone policy is set to deny.

Example:

```
switch# show zone status
VSAN: 1 default-zone: deny distribute: active only Interop: default
mode: basic merge-control: allow
session: none
```

The state default zone policy permit means all nodes can see all other nodes. Deny means all nodes are isolated when not explicitly placed in a zone.

If you are not using zoning, you can change the default zone policy with **zone default-zone permit**, but this is not a best practice.

- Use the **show zone member** command for host and storage to verify that they are both in the same zone. If they are not in the same zone, use the **zone name <zonename> <vsan-id>** command to create a zone in that VSAN.

Example:

```
switch(config)# zone name testzone vsan 100
switch(config-zone)# member pwn 21:00:00:20:37:9e:02:3e
switch(config-zone)# member pwn 21:00:00:c0:dd:12:04:ce
```

Use the **show zone vsan <vsan-id>** command to verify that host and storage are now in the same zone.

- Use the **show zoneset active vsan <vsan-id>** command to verify the name of the active zone set.

REVIEW DRAFT – CISCO CONFIDENTIAL

If the zone that has the host and storage is not in the active zoneset, use the **zoneset name** command from the configuration mode to enter the zoneset sub-mode and use the **member** command to add the zone to the active zone set.

Example:

```
switch(config # zoneset name testzoneset vsan 100
switch(config-zoneset)# member testzone
```

- Use the **zoneset activate** command to activate the zone set.

Example:

```
switch(config)# zoneset activate testzoneset vsan 100
```

Zone merge failure when two switches connect using E or TE port

A zone merge failure can occur when two switches connect using the E or TE port.

Possible log messages that can be seen in the **show logging** log are shown in the example.

Example:

```
%ZONE-2-ZS_MERGE_FAILED: %$VSAN 1%$ Zone merge failure, isolating interface fc2/1 error:
Received rjt from adjacent switch:[reason:0]
%ZONE-2-ZS_MERGE_FAILED: %$VSAN 1%$ Zone merge failure, isolating interface fc1/2 error:
Member mismatch
%ZONE-2-ZS_MERGE_ADJ_NO_RESPONSE: Adjacent switch not responding,isolating interface
%ZONE-2-ZS_MERGE_FULL_DATABASE_MISMATCH: Zone merge full database mismatch on interface
```

Possible Cause

Two switches may have the same zone set name and the same zone names, but different zone members.

When merging switch fabrics, you must ensure that the zones in both active zone sets have unique names, or that any zones with the same name have exactly the same members. If either of these conditions are not met, then the E port connecting the two fabrics will appear in an isolated state.

The process to merge switch fabrics is as follows:

- The software compares the protocol versions. If the protocol versions differ, then the ISL is isolated.
- If the protocol versions are the same, then the zone policies are compared. If the zone policies differ, then the ISL is isolated.
- If the zone merge options are the same, then the comparison is implemented based on the merge control setting.
 - If the setting is restrict, the active zone set and the full zone set should be identical. Otherwise the link is isolated.
 - If the setting is allow, then the merge rules are used to perform the merge. The host and storage are not in the same zone. The zone is not in the active zoneset. There is no active zoneset and default zone policy is set to deny.

Solution

If there is a zone merge failure, the issue can be resolved by using one of the following methods:

- Modify the zone members in both zone sets to match and eliminate the conflict.
 - Use the **show zoneset active vsan <vsan-id>** command on both switches to compare the zones and their respective members.
 - Change the membership of one of the zones to match the other zone of the same name.

REVIEW DRAFT – CISCO CONFIDENTIAL

- Deactivate the zone set on one of the switches and restart the zone merge process.
 - Use the **no zoneset activate name** <zonesetname> <vsan-id> command to deactivate the zone set configuration from one of the switch.
 - Use the **show zoneset active** command to confirm that the zone set has been removed.
 - Use the **shutdown** command to shut down the connection to the zone to be merged, and use the **no shutdown** command to reactivate the connection to the zone to be merged.
 - Use the **show zoneset active** <vsan-id> to verify that all the members are correct and use the **show interface fc** <slot/port> to verify that the VSAN is not isolated.
- Explicitly import or export a zone set between the switches to synchronize them.
 - Use the **zoneset import interface** <interface-number> **vsan** <vsan-id> command or the **zoneset export interface** <interface-number> **vsan** <vsan-id> command to overwrite the active zone set on one of the switches.
 - Use the **show interface fc** <slot/port> to verify that the VSAN is not isolated after this disruptive operation.

Zone set activation failure

When a zone set activation failure occurs, the possible log messages that can be seen in the **show logging** log are shown in the example.

Example:

```
ZONE-2-ZS_CHANGE_ACTIVATION_FAILED: Activation failed.
ZONE-2-ZS_CHANGE_ACTIVATION_FAILED_RESN: Activation failed : reason
```

Possible Cause

Zone set activation can fail if a new switch joins the fabric when the size of the zone database is larger than 2048 KB.

Solution

- Use the **show zone analysis active vsan** <vsan-id> command to analyze the active zone set database. Verify that the formatted size does not exceed 2048 KB.

If the 2048 KB limit is exceeded, then some zones or devices within a zone must be removed.

Example:

```
switch# show zone analysis active vsan 100
Zoning database analysis vsan 100
Active zoneset: vsm_vem_v100_zs [-]
Activated at: 13:13:44 UTC May 27 2010
Activated by: Merge [ Interface san-port-channel 100 ]
Default zone policy: Deny
Number of devices zoned in vsan: 1/9 (Unzoned: 8)
Number of zone members resolved: 1/3 (Unresolved: 2)
Num zones: 1
Number of IVR zones: 0
Number of IPS zones: 0
Formatted size: 92 bytes / 2048 Kb
```

- Use the **show zone internal change event-history vsan** <vsan-id> command to determine the zone set activation problem.
- To further troubleshoot this issue, capture the output from the **show tech-support zone** command and the **show logging log** command.

REVIEW DRAFT – CISCO CONFIDENTIAL**Full zone database synchronization failure across two switches**

A full zone database synchronization failure may occur when two switches connect using the E or TE port and have different zone set distribution policies. As a result of a fabric isolation/merge, one fabric might not have the full zone set database in the running configuration.

Possible Cause

The zone set distribution takes effect while sending merge requests to the adjacent switch or while activating a zone set.

The zone distribute policy can be set differently on two switches and that could cause synchronization failure.

Solution

Use the **show zone status** command to verify the distribution policy on both switches.

Example:

```
VSAN: 100 default-zone: deny distribute: active only Interop: default
mode: basic merge-control: allow
```

When the distribute policy is set to active only the active zone set is distributed. Also verify that the distribute policy is set to full.

To enable the full zone set and active zone set distribution to all switches on a per-VSAN basis in the configuration mode, use the **zoneset distribute full vsan <vsan-id>** command.

Mismatched default zone policy in switches in VSAN causes unexpected results when accessing storage

A mismatched default zone policy in all switches in the VSAN in the basic zone mode might cause unexpected results for any hosts accessing storage.

Possible Cause

If the default zone policy is set to permit and if there is no active zone set for VSAN, then all the members of the VSAN can see all the other nodes.

Solution

One approach is to migrate the zone operation mode from basic to enhanced. Enhanced zoning synchronizes the zone configuration across all switches in the VSAN. This eliminates the possibility of mismatched default zone policies.

- Use the **show zone status** command to display the status of the zone.

```
VSAN: 300 default-zone: deny distribute: active only Interop: default
mode: basic merge-control: allow
```

- Use the **zone default-zone** command to set the default zone policy and use the **zone mode enhanced <vsan-id>** command to set the operation to enhanced zoning mode.

Another approach is the foollowing:

- Use the **show zone status** command in all switches in the VSAN to verify the operation mode and the default-zone policy.
- Use the **zone mode basic** command to change any switches that are not in basic mode.

REVIEW DRAFT – CISCO CONFIDENTIAL

- Use the **zone default-zone** command on each switch in the VSAN to set the same default zone policy.

SAN Port Channels

Fibre channel port is down when trying to connect switches via SAN Port Channel

When trying to connect switches using SAN port channel, the Fibre Channel port is down.

The **show interface brief** command produces

```
fc slot/port is down (Error disabled - Possible port channel misconfiguration)
```

Possible Cause

One of the SAN port channel compatibility parameters is misconfigured in the configuration.

A compatibility check ensures that the same parameter settings are used in all physical ports in the channel. Otherwise, they cannot become part of a port channel. The compatibility check is performed before a port is added to the port channel.

The check ensures that the following parameters and settings match at both ends of a port channel:

- Capability parameters:
type of interface, Gigabit Ethernet at both ends, or Fibre Channel at both ends.
- Administrative compatibility parameters:
speed, mode, rate mode, port VSAN, allowed VSAN list, and port security..
- Operational parameters:
remote switch WWN and trunking mode.

Solution

- Use **show san-port-channel compatibility-parameters** to verify which parameters need to be checked in the configuration.

Generally, if the configuration is fixed and the FC port is shut or no shut, the port recovers normally.

- If the issue persists with a different error message, debug further by running one or more of the following commands:

```
show port internal info interface fc <slot/port>
show port internal event-history interface fc <slot/port>
show san-port-channel internal event-history errors
show logging log | grep fc <slot/port>
show san-port-channel internal event-history all
show tech-support detail > bootflash:showtechdet
```

Newly added Fibre Channel interface does not come online in a SAN Port Channel

When a new Fibre Channel interface is added, it does not come online in a SAN port channel.

REVIEW DRAFT – CISCO CONFIDENTIAL

The following error message during the configuration operation may appear.

```
Command failed: port not compatible [reason]
```

Possible Cause

Port channel mode is configured as on.

If you use the default ON mode to avoid inconsistent states across switches and to maintain consistency across switches, then the ports shut down.

Solution

Explicitly enable the ports again using the **no shutdown** command.

Possible Cause

Interface parameters are not compatible with the existing SAN port channel.

Solution

Use the force option to force the physical interface to take on the parameters of the SAN port channel. In the interface sub-configuration mode, use the **channel-group <channel-group number> force** command.

Cannot configure trunking

Trunking cannot be configured under the interface configuration mode.

The following error message may appear in the CLI output:

```
error:invalid switchport config
```

Possible Cause

Trunking protocol is disabled.

Solution

Enable trunking by using the **trunk protocol enable CLI** command.

VSAN traffic does not traverse trunk

The VSAN traffic is not able to traverse the trunk.

A host cannot gain access to a target that is on the same VSAN and connected to two different switches using TE ports. The VSAN traffic is not able to traverse the trunk. Depending on the path from host to target, you may observe a performance degradation or you may not be able to access any disks.

Possible Cause

VSAN is not listed in the allowed-active VSAN list.

Solution

Add VSAN to the allowed-active list by using the **switchport trunk allowed vsan** command.

REVIEW DRAFT – CISCO CONFIDENTIAL**xE port is isolated in a specific VSAN under interface of SAN Port Channel**

The xE port is isolated in a specific VSAN that is under an interface of a SAN port channel.

The following error message may appear in the logging log:

```
"%$VSAN <VSAN#>%$ Interface port-channel <channel #>, vsan <vsan #> is down (isolation due to [cause])".
```

Possible Cause

The xE port can be isolated in a specific VSAN for many reasons:

- Fabric timers might be misconfigured.
- Port parameters might be misconfigured.
- Zoning mismatch.

Solution

To resolve the VSAN isolation on the TE port, use the **show interface fc** *<slot/port>* command on the TE port to determine the VSAN number. The isolated VSAN number must match the VSAN number where the host and the storage are connected to.

In the output of the command, look for information such as Trunk vsans (isolated) (Vsan <number>).

Use the **show port internal info interface san-port-channel** *<number>* command to determine the cause of the VSAN isolation.

Cannot create a san-port-channel interface

A SAN port channel interface cannot be created.

The following error message may appear while in configuration mode:

```
failed to create port-channel channel-id:
```

Possible Cause

The user receives the following message:

```
failed to create port-channel channel-id: all port-channels have been created [max channel number reached]
```

**Note**

You can create a maximum of four SAN port channels (including Release NX-OS 4.2(1)N1(1)). This is a software limitation.

Solution

If you need to create a SAN port channel with a specific number, but four SAN port channels were already configured, then you have to delete one of the SAN port channels that is not actively used. Use the **no interface san-port-channel** *<x>* command to delete one of the SAN port channels.

Possible Cause

You receive the following message:

```
Channel group X is already an Ethernet port channel
```

Solution

REVIEW DRAFT – CISCO CONFIDENTIAL

You need to choose another number between 1 to 256 to configure the SAN port channel.

Use the **show port-channel usage** command to determine the numbers that were used for the existing port channels.

Example:

```
show port-channel usage
Total 3 port-channel numbers used
=====
Used : 198 - 199 , 500
Unused: 1 - 197 , 200 - 499 , 501 - 4096
(some numbers may be in use by SAN port channels)
```

FC Services

This section includes an overview of troubleshooting Cisco Fibre Channel Services followed by a description of common problems and their solutions.

Overview

Fibre Channel fabrics provide a set of services for its clients, which are the Fibre Channel nodes. These Fibre Channel services (FC services) allow the nodes to interact with the storage network to exchange information, such as connection state, connection parameters, configuration, topology changes, and so on.

The FC services can be accessed through login into ports that hold a well known address (WKA). WKAs are port FC IDs that are reserved for internal use of the fabric, usually fabric services.

The following table describes the well-known addresses and the service associated with each: (Source: www.t11.org)

Well Known Address	Description
x'FF FC 01' to x'FF FC FE'	Reserved for Domain Controllers
x'FF FF F0'	Reserved for N_Port Controller
x'FF FF F1' to x'FF FF F3'	Reserved
x'FF FF F4'	Event Service (FC-GS-5)
x'FF FF F5'	Multicast Server (FC-PH3)
x'FF FF F6'	Clock Synchronization Server (FC-PH3)
x'FF FF F7'	Security Key Distribution Service (FC-PH3)
x'FF FF F8'	Alias Server (FC-PH2)
x'FF FF F9'	Quality of Service Facilitator-Class4 (FC-PH2)
x'FF FF FA'	Management Service (FC-GS-5)
x'FF FF FB'	Time Service (FC-GS-5)
x'FF FF FC'	Directory Service (FC-GS-5)

REVIEW DRAFT – CISCO CONFIDENTIAL

Well Known Address	Description
x'FF FF FD'	Fabric Controller
x'FF FF FE'	F_Port Controller
x'FF FF FF'	Broadcast Address/Server

Fibre channel port remains in initializing state

A fibre channel F type port does not come online and is stuck in an initializing state.

The **show interface fc** <slot/port> command displays the following message.

```
fc slot/port is down (Initializing)
```

A Fibre Channel port goes into the initialization state after a successful completion of link-level initialization. For F type ports, the next step is to complete the FLOGI (fabric login) process. The port remains in the initialization state until the FLOGI process completes.

Possible Cause

The port is up because the link partner has put itself into a bypass mode.

Solution

Use the **show hardware internal fc-mac** <slot-number> **port** <port-number> **statistics** command to check whether the Class-3 input counter is increasing after the successful completion of link initialization.

Example:

```
switch# show hardware internal fc-mac 2 port 1 statistics
ADDRESS      STAT                                          COUNT
-----
0x0000003c   FCP_CNTR_MAC_RX_LOSS_OF_SYNC                0x1
0x0000003d   FCP_CNTR_MAC_RX_BAD_WORDS_FROM_DECODER      0x50
0x00000042   FCP_CNTR_MAC_CREDIT_IG_XG_MUX_SEND_RRDY_REQ 0x152
0x00000043   FCP_CNTR_MAC_CREDIT_EG_DEC_RRDY            0x7c
0x00000061   FCP_CNTR_MAC_DATA_RX_CLASS3_FRAMES          0x130
0x00000062   FCP_CNTR_MAC_DATA_RX_CLASSF_FRAMES          0x22
0x00000069   FCP_CNTR_MAC_DATA_RX_CLASS3_WORDS           0x61c98
0x0000006a   FCP_CNTR_MAC_DATA_RX_CLASSF_WORDS           0xff0
0x00000065   FCP_CNTR_MAC_DATA_TX_CLASS3_FRAMES          0x52
0x00000066   FCP_CNTR_MAC_DATA_TX_CLASSF_FRAMES          0x2a
0x0000006d   FCP_CNTR_MAC_DATA_TX_CLASS3_WORDS           0x944c
0x0000006e   FCP_CNTR_MAC_DATA_TX_CLASSF_WORDS           0xec4
0xffffffff   FCP_CNTR_LINK_RESET_IN                      0x1
0xffffffff   FCP_CNTR_OLS_IN                             0x1
0xffffffff   FCP_CNTR_NOS_IN                             0x1
0xffffffff   FCP_CNTR_LRR_IN                             0x2
0xffffffff   FCP_CNTR_LINK_RESET_OUT                     0x1
0xffffffff   FCP_CNTR_OLS_OUT                            0xa
0xffffffff   FCP_CNTR_NOS_OUT                            0x2
0xffffffff   FCP_CNTR_LRR_OUT                            0xb
0xffffffff   FCP_CNTR_LINK_FAILURE                       0x2
```

Possible Cause

The FLOGI packet was dropped somewhere in the data path, starting from FC-MAC to FLOGI server.

Solution

Consider the following solutions:

REVIEW DRAFT – CISCO CONFIDENTIAL

- Use the **show hardware internal fc-mac <slot-number> port <port-number> statistics** command to check for Class-3 packet counters.
- Analyze the output of the **show flogi internal all interface fc <slot/port>** command for a possible drop of FLOGI packets somewhere in the path.
- Check the Fport server fault-injection table for any Invalid, Drop FLOGI packets.
- Use the shut CLI command followed by the no shut command to disable and enable the FC slot/port.
- If this does not clear the problem, try moving the connection to a different port on the same or another FC module.
- If the problem continues to persist, use the following commands to collect information to aid in further analysis:

```
show tech-support flogi
show logging log | grep fc <slot/port>
show port internal info interface fc <slot/port>
show port internal event-history interface fc <slot/port>
show tech-support detail > bootflash:showtechdet
show platform fwm info pif fc <slot/port> {find the gatos instance for the port}
show platform fwm info gatos-errors 13 {check for the non-zero counters for drops}
```

Capture debug Flogi with the following:

```
switch# debug logfile flogi_debug
switch# debug flogi all
switch(config)# int fc <slot/port>
switch(config-if)# shut
switch(config-if)# no shut
switch(config-if)# undebug all
switch# dir log: {check if you have the file in log: directory}
      31      Aug 03 13:45:13 2010  dmesg
  34941      Aug 06 07:21:15 2010  flogi_debug

switch# copy log:flogi_debug ftp://x.y.z.w {or use tftp/scp/sftp}
```

Specific VSAN traffic is not being routed through SAN fabric

Each configured VSAN needs to support a separate set of fabric services. One such service is the FSPF routing protocol, which can be independently configured per VSAN. You may see that specific VSAN traffic is not being routed if inappropriate traffic engineering capabilities are used.

Possible Cause

The FSPF hello interval is misconfigured.

The following example shows possible log messages from the **show logging** command log.

Example:

```
FSPF-3-HELLO_MISMATCH: %$VSAN <vsan-id>%$ Mismatch in Hello timer in the Hello packet on
interface san-port-channel <channel-id>
%FSPF-3-FC2_PROC_ERR: %$VSAN <vsan-id>%$ Error in processing HELLO packet on interface
san-port-channel <channel-id>, Error = Bad packet received
```

Solution

To resolve a wrong hello interval on an ISL using the NX-OS CLI, perform the following steps.

REVIEW DRAFT – CISCO CONFIDENTIAL

- Step 1** Either use the **debug fspf all** command and look for wrong hello interval messages or check the last messages in the **show logging** command log for an error message.

The debug output generates the following messages:

```
fspf: Wrong hello interval for packet on interface 40000c7 in VSAN 200
fspf: Error in processing hello packet , error = Bad packet received
```

- Step 2** Use the **undebug all** command to turn off debugging.

**Tip**

Open a second Telnet or SSH session before entering any debug commands. If the debug output overwhelms the current session, you can use the second session to enter the **undebug all** command to stop the debug message output.

- Step 3** Use the **show fspf vsan <vsan-id> interface** command to view the FSPF configuration on both switches.

Example:

```
switch# show fspf vsan 200 interface port-channel 200
FSPF interface port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 40 s, Dead 80 s, Retransmit 5 s
FSPF State is INIT
Statistics counters :
  Number of packets received : LSU 3 LSA 3 Hello 136 Error packets 3
  Number of packets transmitted : LSU 3 LSA 3 Hello 182 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 0
```

```
switch# show fspf vsan 200 interface san-port-channel 200
FSPF interface san-port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is INIT

Statistics counters :
  Number of packets received : LSU 3 LSA 3 Hello 185 Error packets 169
  Number of packets transmitted : LSU 3 LSA 3 Hello 139 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 24
```

**Note**

In the Example:

- The hello timer is not set to the default (20 seconds) on the first switch. Check the neighboring switch (Nexus 5000) configuration to make sure it matches.
- FSPF is not in FULL state. This indicates a problem.

- Step 4** In the interface configuration mode, change the fspf hello-interval value to match the same values on both switches.

Example:

```
switch(config)# interface san-port-channel 200
switch(config-if)# fspf hello-interval 40 vsan 200
```

REVIEW DRAFT – CISCO CONFIDENTIAL

Step 5 Verify that the FSPF is in FULL state after the change.

```
switch(config-if)# show fspf vsan 200 interface san-port-channel 200
FSPF interface san-port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 40 s, Dead 80 s, Retransmit 5 s
FSPF State is FULL
Neighbor Domain Id is 0x18(24)
Neighbor Interface is san-port-channel 200 (0x000400c7)

Statistics counters :
  Number of packets received : LSU 7 LSA 7 Hello 238 Error packets 218
  Number of packets transmitted : LSU 7 LSA 7 Hello 180 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 32
```

Possible Cause

The FSPF dead interval is misconfigured.

The following example shows possible log messages from the **show logging** command:

Example:

```
%FSPF-3-HELLO_MISMATCH: %$VSAN <vsan-id>%$ Mismatch in Dead timer in the Hello packet on
interface san-port-channel <channel-id>
N5K-2 %FSPF-3-FC2_PROC_ERR: %$VSAN <vsan-id>%$ Error in processing HELLO packet on
interface san-port-channel <channel-id>, Error = Bad packet received
```

Solution

To identify a mismatch of dead intervals on an ISL using the NX-OS CLI, perform the following steps:

Step 1 Either use the **debug fspf all** command and look for wrong dead interval messages or check the last messages in the **show logging** command log for an error message.

The debug output generates the following messages:

```
fspf: Wrong hello interval for packet on interface 40000c7 in VSAN 200
fspf: Error in processing hello packet , error = Bad packet received
```

Step 2 Use the **undebug all** command to turn off debugging.

**Tip**

Open a second Telnet or SSH session before entering any debug commands. If the debug output overwhelms the current session, you can use the second session to enter the **undebug all** command to stop the debug message output.

Step 3 Use the **show fspf vsan <vsan-id> interface** command to view the FSPF configuration on both switches.

Example:

```
switch# show fspf vsan 200 interface san-port-channel 200
FSPF interface san-port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 20 s, Dead 120 s, Retransmit 5 s
FSPF State is INIT

Statistics counters :
  Number of packets received : LSU 4 LSA 4 Hello 27 Error packets 4
```

REVIEW DRAFT – CISCO CONFIDENTIAL

```
Number of packets transmitted : LSU 4 LSA 4 Hello 38 Retransmitted LSU 0
Number of times inactivity timer expired for the interface = 0
```

```
switch# show fspf vsan 200 interface port-channel 200
FSPF interface port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is INIT
```

```
Statistics counters :
  Number of packets received : LSU 4 LSA 4 Hello 41 Error packets 35
  Number of packets transmitted : LSU 4 LSA 4 Hello 29 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 4
```

**Note**

In the example:

- The dead timer is not set to the default (80 seconds) on the first switch. Check the neighboring switch (MDS) configuration to make sure it matches.
- FSPF is not in FULL state. This indicates a problem.

Step 4 In the interface configuration mode, change the fspf dead-interval value so that the same values match on both switches.

```
switch(config)# interface san-port-channel 200
switch(config-if)# fspf dead-interval 80 vsan 200
```

Step 5 Verify that the FSPF is in FULL state after the change. Ensure that there is a route for VSAN traffic with the **show fspf internal route vsan <vsan-id>** command.

Example:

```
switch# show fspf internal route vsan 200
```

```
FSPF Unicast Routes
```

```
-----
VSAN Number  Dest Domain  Route Cost  Next hops
-----
          200      0x18(24)      125 san-port-channel 200
```

```
switch# show fspf vsan 200 interface san-port-channel 200
FSPF interface san-port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is FULL
Neighbor Domain Id is 0x18(24)
Neighbor Interface is san-port-channel 200 (0x000400c7)
```

```
Statistics counters :
  Number of packets received : LSU 8 LSA 8 Hello 47 Error packets 4
  Number of packets transmitted : LSU 8 LSA 8 Hello 70 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 0
```

Possible Cause

There is a region mismatch on the switch.

REVIEW DRAFT – CISCO CONFIDENTIAL

The following example shows possible log messages from the **show logging** command log:

Example:

```
%FSPF-3-BAD_FC2_PKT: %$VSAN 200%$ Received bad FC2 packet on interface san-port-channel
<channel-id> : Packet received for non existant region in VSAN
```

Solution

To identify a region mismatch problem on a switch using the NX-OS CLI, perform the following.

Step 1 Use the **show fspf vsan <vsan-id>** command to display the currently configured region in a VSAN.

Example (region value is 2; default region value is 0):

```
switch# show fspf vsan 200
FSPF routing for VSAN 200
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 2
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 2000 msec
Local Domain is 0x22(34)
Number of LSRs = 1, Total Checksum = 0x00000c10

Protocol constants :
  LS_REFRESH_TIME = 30 minutes (1800 sec)
  MAX_AGE          = 60 minutes (3600 sec)

Statistics counters :
  Number of LSR that reached MaxAge = 0
  Number of SPF computations         = 0
  Number of Checksum Errors          = 0
  Number of Transmitted packets :   LSU 0 LSA 0 Hello 19 Retranmsitted LSU 0
  Number of received packets :     LSU 0 LSA 0 Hello 0 Error packets 18
```

Step 2 Use the **debug fspf all** command and look for nonexistent region messages.

Example:

```
fspf: Hello timer reached for interface san-port-channel 200 in VSAN 200
fspf: FC2 packet received for non existant region 0 in VSAN 200
fspf: FC2 packet received for non existant region 0 in VSAN 200
```

The neighboring switch-advertising region is 0. FSPF is in the init state for each ISL.

Example:

```
switch# show fspf vsan 200 interface san-port-channel 200
FSPF interface san-port-channel 200 in VSAN 200
FSPF routing administrative state is active
Interface cost is 125
Timer intervals configured, Hello 20 s, Dead 80 s, Retransmit 5 s
FSPF State is INIT

Statistics counters :
  Number of packets received : LSU 0 LSA 0 Hello 0 Error packets 0
  Number of packets transmitted : LSU 0 LSA 0 Hello 49 Retransmitted LSU 0
  Number of times inactivity timer expired for the interface = 9
```

Step 3 Use the **undebug all** command to turn off debugging.

Step 4 Use the **show fspf vsan <vsan-id>** command to show FSPF configuration and check the autonomous region.

REVIEW DRAFT – CISCO CONFIDENTIAL

Example:

```
switch# show fspf vsan 200
FSPF routing for VSAN 200
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 2
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 2000 msec
Local Domain is 0x22(34)
Number of LSRs = 1, Total Checksum = 0x00000c10
```

```
switch# show fspf vsan 200
FSPF routing for VSAN 200
FSPF routing administration status is enabled
FSPF routing operational status is UP
It is an intra-domain router
Autonomous region is 0
SPF hold time is 0 msec
MinLsArrival = 1000 msec , MinLsInterval = 2000 msec
Local Domain is 0x18(24)
Number of LSRs = 2, Total Checksum = 0x00014f9f
```

- Step 5** Use the `fspf config vsan` command to enter the FSPF configuration mode and use the `region` command to change the region. The region must match on all switches in the VSAN.

Example:

```
switch(config)# fspf config vsan 200
switch(config-(fspf-config))# region 0
```

Fibre channel port is suspended due to too many invalid FLOGIs

A Fibre Channel node that is connected to an NPV feature-enabled Cisco Nexus 5000 switch or a Cisco Nexus 5000 switch that is running in fabric mode cannot log into the SAN fabric due to a FLOGI rejection.

The following example shows possible log messages from the `show logging` command log.

Example:

```
%FLOGI-1-MSG_FLOGI_REJECT_FCID_ERROR: %$VSAN <vsan-id>%$ [VSAN <vsan-id>, Interface
fcslot/port/: mode[F]] FLOGI rejected - FCID allocation failed.
PORT-5-IF_DOWN_TOO_MANY_INVALID_FLOGIS: %$VSAN <vsan-id>%$ Interface fc slot/port is down
(Suspended due to too many invalid flogis
```

The status of the interface shows `invalidFlogis`.

```
show interface fc slot/port brief
fc slot/port      <vsan-id>      F      --      invalidFlogis
```

Possible Cause

The FC ID persistency table for that VSAN might be full. If the Nexus 5000 Series switch is configured as an NPV edge switch, the FC ID persistency table of the NPV core switch might be full.

FC IDs:

When an N port logs into a Cisco Nexus 5000 Series switch, it is assigned an FC ID. By default, the persistent FC ID feature is enabled. If this feature is disabled, the following situations can occur:

REVIEW DRAFT – CISCO CONFIDENTIAL

- An N port logs into a Cisco Nexus 5000 Series switch. The WWN of the requesting N port and the assigned FC ID are retained and stored in a volatile cache. The contents of this volatile cache are not saved across reboots.
- The switch is designed to preserve the binding FC ID to the WWN on a best-effort basis. For example, if one N port disconnects from the switch and its FC ID is requested by another device, this request is granted and the WWN with the initial FC ID association is released.
- The volatile cache stores up to 4000 entries of WWN to FC ID binding. If this cache is full, a new (more recent) entry overwrites the oldest entry in the cache. In this case, the corresponding WWN to FC ID association for the oldest entry is lost.
- N ports receive the same FC IDs if disconnected and reconnected to any port within the same switch (as long as it belongs to the same VSAN).

Persistent FC IDs can be purged selectively. Static entries and FC IDs currently in use cannot be deleted.

Solution

Check for FLOGI error messages with the **show flogi internal** command.

Example:

```
show flogi internal event-history debugs

222) Event:E_FLOGI_DEBUG, length:309, at 989582 usecs after Thu Jun 17
09:03:01 2010
fs_print_port_stats(10049): Port Stats for fc2/1, after cleanup:
  timestamp: Wed Jun 17 07:03:01 2010
  MSG_FLOGI: 52
  MSG_FC2_LS_RJT_OUT: 51
  EXCEPTION_CANNOT_ALLOCATE_FCID: 51
  EXCEPTION_TIMEOUT: 1
  EXCEPTION_FC2_INVALID_XCHG: 1
  tot_internal_exceptions: 51, since: Thu Dec 31 17:00:00 1969
```

```
show flogi internal errors

52) Event:E_DEBUG, length:119, at 977471 usecs after Thu Jun 17 09:03:01
2010
  [102] Interface fc2/1, nwwn 20:01:00:1b:32:af:d6:8c, pwwn
21:01:00:1b:32:af:d6:8c: flogi is valid; exchange is INVALID.
```

Use the **show fcdomain address-allocation** command to check the FC domain address allocation table for any free FC IDs: (If NPV is enabled, enter the command on the NPV core switch.)

Example:

```
show fcdomain address-allocation

VSAN 1
Free FCIDs: 0xe73f4f to 0xe73fff
           0xe7ff00 to 0xe7fffe

Assigned FCIDs: 0xe70000 to 0xe73f4e
               0xe74000 to 0xe7feff
               0xe7ffff

Reserved FCIDs: 0xe7ffff

Number free FCIDs: 432
Number assigned FCIDs: 65104
Number reserved FCIDs: 1
```


REVIEW DRAFT – CISCO CONFIDENTIAL

To find the auto area-list and the persistent FCIDs, use the **show flogi auto-area-list** command and the **show fcdomain fcid persistent** command.

Example:

```
show flogi auto-area-list
Fcid area allocation company id info:
<...>
  00:14:5E
  00:1B:32
  00:50:2E
  00:E0:69
  00:E0:8B

show fcdomain fcid persistent {entire AREA reserved for OUI 00:E0:8B}

102    21:01:00:1b:32:2f:7f:63    0x020003    SINGLE FCID    YES    DYNAMIC
102    21:00:00:1b:32:0f:7f:63    0x020004    SINGLE FCID    YES    DYNAMIC

102    21:00:00:e0:8b:89:a7:07    0x021c00    ENTIRE AREA    YES    DYNAMIC
102    21:00:00:e0:8b:88:e9:22    0x024300    ENTIRE AREA    YES    DYNAMIC
```

If there are not enough FCIDs, you can purge dynamic and unused FC IDs in the specified VSAN with the **purge fcdomain** command.

Example:

```
switch# purge fcdomain fcid vsan <vsan-id>
```

The ports will soon come up.

It is also possible that HBAs are trying to login with S_ID != 0x0.

If this is the situation and there is nothing in the persistency table for the WWN of the HBA, try to assign the S_ID used by HBA to the HBA itself.

If the S_ID is already in use or is in the wrong domain, the request is rejected by fcdomain. After a number of retries, the port is suspended.

When HBAs get into this mode, they try to log in with every FCID in the FCID space, from 0x00.00.01 up to all the 0xDD.AA.PP numbers.

This behavior can be seen in the **show flogi internal event-history msgs** command output {HBA is trying to login with different FCIDs}

Example:

```
841) Event:E_FLOGI_LRX, length:20, at 56079 usecs after Tue Jun 22 15:40:59 2010
      WWN: 21:01:00:1b:32:af:d6:8c VSAN: 1 ifindex: fc2/1 FCID: 0x000032

886) Event:E_FLOGI_RX, length:20, at 897472 usecs after Tue Jun 22 15:40:58 2010
      WWN: 21:01:00:1b:32:af:d6:8c VSAN: 1 ifindex: fc2/1 FCID: 0x000030

888) Event:E_FLOGI_FAIL, length:20, at 884758 usecs after Tue Jun 22 15:40:58 2010
      WWN: 21:01:00:1b:32:af:d6:8c VSAN: 1 ifindex: fc2/1 ev_id: 21
      rjt reason: 7 OPC: MTS_OPC_DM_GET_FCIDS(275)

903) Event:E_FLOGI_RX, length:20, at 835015 usecs after Tue Jun 22 15:40:58 2010
      WWN: 21:01:00:1b:32:af:d6:8c VSAN: 1 ifindex: fc2/1 FCID: 0x00002f
```

In this case, the solution is to manually configure an entry in the persistency table for the WWN of the HBA as shown in the following example. An alternative is to power-cycle the device. This usually makes the HBA start with a normal FLOGI with S_ID=0x0.

Example:

REVIEW DRAFT – CISCO CONFIDENTIAL

```
switch# conf t
switch(config)# fcdomain fcid database
switch(config-fcid-db)# vsan <vsan-id> wwn 50:05:08:b2:00:71:c8:c2 fcid 0x6fee00 area
```

If the problem continues to persist, use the following commands to collect information to aid further analysis.

```
Show tech-support flogi
Show tech-support fcdomain
Show logging log
show port internal info interface fc <slot/port>
show port internal event-history interface fc <slot/port>
show tech-support detail > bootflash:showtechdet
```

Capture debug flogi and debug fcdomain via following below steps:

```
switch# debug logfile flogi_fcdomain
switch# debug flogi all
switch# debug fcdomain all

switch(config)# int fc <slot/port>
switch(config-if)# shut
switch(config-if)# no shut
switch(config-if)# undebug all
switch# dir log: {check if you have the file in log: directory}
      31      Aug 03 13:45:13 2010  dmesg
55941      Aug 05 07:21:15 2010  flogi_fcdomain

switch# copy log:flogi_fcdomain ftp://x.y.z.w {or use tftp/scp/sftp}
```

Having stale FCNS entries for Fibre Channel nodes

The Fibre Channel nodes are able to be logged (FLOGI) in to the SAN fabric, but the FCNS entries for those nodes are incomplete. Servers cannot reach their targets.

As a result, **fc4-types:fc4_features** will be empty in FCNS database.

Possible Cause

The Fibre Channel nodes may not be registering their FC4 types and FC4 features in the FCNS database in a topology where Nexus 5000 Series switches are configured as NPV core (feature NPIV) and connected to legacy gateway switches. The **fc4-types:fc4_features** can be verified by the **show fcns database detail** command as shown in the following example:

Example:

```
switch# show fcns da fcid 0x621400 detail vsan 2
-----
VSAN:2      FCID:0x621400
-----
port-wwn (vendor)      :21:01:00:1b:32:a3:d7:2c
                        [z7095ib-1_T]
node-wwn                :20:01:00:1b:32:a3:d7:2c
class                   :3
node-ip-addr            :0.0.0.0
ipa                     :ff ff ff ff ff ff ff ff
fc4-types:fc4_features  :
symbolic-port-name     :
symbolic-node-name     :
port-type               :N
port-ip-addr            :0.0.0.0
fabric-port-wwn        :20:d9:00:0d:ec:e0:0e:80
```

REVIEW DRAFT – CISCO CONFIDENTIAL

```

hard-addr                :0x000000
permanent-port-wnn (vendor) :20:11:00:05:1e:06:da:ea
Connected Interface      :fc2/2
Switch Name (IP address) :N5K (10.200.220.13)

```

Some legacy gateway switches might require that the area part of the FCID be the same for the switch and for all the blades logged in through that port.

However, because of an old issue with Qlogic HBAs, the Cisco Nexus 5000 domain server assigns a separate area for each Qlogic HBA that matches a certain OUI by default. Therefore, a conflict between legacy gateway requirements and the Cisco domain allocation scheme exists. Cisco still implements this set up to support old existing Qlogic HBAs in the field.

Solution

Configure **no fcid-allocation area company** <oui> for all used Qlogic OUIs (ensuring flat FCID allocation in the future), force all affected blades to log out of the fabric, delete the already created persistent FCID entry from the Nexus 5000 switch configuration, and allow the blade to log in again.

In the following **show flogi database** command output, all devices obtain a unique area id (x01, x08, x0c):

Example:

```

Fc2/1  2      0x620104  20:10:00:05:1e:5e:6a:85  10:00:00:05:1e:5e:6a:85
Fc2/1  2      0x620800  21:01:00:1b:32:a3:c0:2e  20:01:00:1b:32:a3:c0:2e
Fc2/1  2      0x620c00  21:01:00:1b:32:33:8b:8e  20:01:00:1b:32:33:8b:8e

```

Because of the specific area ID requirement of the legacy switch, the last two blades must also have area x01. To force the Qlogic adapters to log in again and obtain FCID in 0x6201xx range, do the following steps:

Step 1 Configure (force) the future FCID allocation scheme to be flat for all WWNs matching the OUIs that are in this situation.

```
switch(configure)# no fcid-allocation area company 0x001B32
```

Step 2 Force the FCID under reconfiguration to log out of the fabric.

**Note**

If you shut down the Nexus 5000 interface that serves as the primary uplink for that server, it only will log in through another one. The appropriate method is to shut down the affected blade and ensure that the FLOGI for the WWN is gone.

Step 3 Delete the automatically created configuration entry for persistent FCID allocation as shown in the following example:

Example:

```
switch(config)# fcdomain fcid database
switch(config-fcid-db)# no vsan 2 wwn 21:01:00:1b:32:a3:c0:2e fcid 0x620800 area dynamic
```

Step 4 Bring up the blade and ensure that it gets a proper fcid.

Example:

```

Fc2/1  2      0x620104  20:10:00:05:1e:5e:6a:85  10:00:00:05:1e:5e:6a:85
Fc2/1  2      0x620123  21:01:00:1b:32:a3:c0:2e  20:01:00:1b:32:a3:c0:2e

```

REVIEW DRAFT – CISCO CONFIDENTIAL

Interface is isolated because of FC Domain ID overlap

The Fibre Channel or SAN port channel interface of a Cisco Nexus 5000 switch (in fabric mode) that is connected to an FC switch with the xE port type is isolated because of a domain overlap. The following example shows possible logging messages in the **show logging** command log:

Example:

```
PORT-5-IF_DOWN_DOMAIN_OVERLAP_ISOLATION: Interface fc <slot/port> is down (Isolation due to domain overlap).
%FCDOMAIN-2-EPORT_ISOLATED: %$VSAN <vsan-id>%$ Isolation of interface san-port-channel <channel-id> (reason: domain ID assignment failure)
%FCDOMAIN-2-EPORT_ISOLATED: %$VSAN <vsan-id>%$ Isolation of interface san-port-channel <channel-id> (reason: other side Eport indicates isolation)
```

Possible Cause

Two switch fabrics might not merge. If two fabrics with two or more switches are connected, have at least one assigned domain ID in common, and the auto-reconfigure option is disabled (this option is disabled by default), then the E ports that are used to connect the two fabrics will be isolated due to domain ID overlap.

In a Fibre Channel network, the principal switch assigns domain IDs when a new switch is added to an existing fabric. However, when two fabrics merge, the principal switch selection process determines which one of the preexisting switches becomes the principal switch for the merged fabric.

The election of the new principal switch is characterized by the following rules:

- A switch with a populated domain ID list takes priority over a switch that has an empty domain ID list. The principal switch becomes the one in the fabric with the populated domain ID list.
- If both fabrics have a domain ID list, the priority between the two principal switches is determined by the configured switch priority. This is a parameter that can be set by the user. The lower the value of the parameter, the higher the priority.
- If the principal switch cannot be determined by the two previous criteria, the principal switch is then determined by the WWNs of the two switches. The lower the value of the WWN, the higher the switch priority.

Solution

To resolve an FC domain ID overlap, you can change the overlapping static domain ID by manually configuring a new static domain ID for the isolated switch, or disable the static domain assignment and allow the switch to request a new domain ID after a fabric reconfiguration.

To assign a static domain ID using the NX-OS CLI

All devices attached to the switch in the VSAN get a new FC ID when a new domain ID is assigned. Some hosts or storage devices may not function as expected if the FC ID of the host or storage device changes.

To verify FC domain ID overlap and reassign a new domain ID using the CLI, perform the following steps:

Step 1 Enter the **show interface fc <slot/port>** command to view the isolation error message for the E port.

Example:

```
switch(config)# show int fc 2/2
fc2/2 is down (Isolation due to domain other side eport isolated)
  Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
  Port WWN is 20:42:00:0d:ec:d5:fe:00
```

REVIEW DRAFT – CISCO CONFIDENTIAL

```
Admin port mode is E, trunk mode is off
snmp link state traps are enabled
Port vsan is 3
```

Enter the **show interface san-port-channel <channel-id>** command to view the isolation error for the specific VSAN.

Example:

```
switch(config)# show interface san-port-channel 200
san-port-channel 200 is trunking (Not all VSANs UP on the trunk)
  Hardware is Fibre Channel
  Port WWN is 24:c8:00:0d:ec:d5:a3:80
  Admin port mode is auto, trunk mode is on
  snmp link state traps are enabled
  Port mode is TE
  Port vsan is 1
  Speed is 8 Gbps
  Trunk vsans (admin allowed and active) (1,200)
  Trunk vsans (up) (1)
  Trunk vsans (isolated) (200)
  Trunk vsans (initializing) ()
```

Step 2 Use the **show fcdomain domain-list vsan <vsan-id>** command to view which domains are currently in your fabric.

Example (switch is isolated because of a domain ID 44 overlap):

```
switch(config)# show fcdomain domain-list vsan 3

Number of domains: 1
Domain ID          WWN
-----
0x2c(44)    20:03:00:0d:ec:3f:a5:81 [Local] [Principal]

switch(config)# show fcdomain domain-list vsan 3

Number of domains: 1
Domain ID          WWN
-----
0x2c(44)    20:03:00:0d:ec:d5:fe:01 [Local] [Principal]
```

If the isolation occurred for a specific VSAN under a SAN port channel interface, you can view the error with the **show port internal info interface san-port-channel <channel-id> vsan <vsan-id>** as shown in the following example:

Example:

```
switch(config)# show port internal info interface san-port-channel 200 vsan 200

san-port-channel 200, Vsan 200 - state(down), state reason(Isolation due to domain other
side eport isolated), fcid(0x000000)
port init flag(0x10000), num_active_ports (2),
Lock Info: resource [san-port-channel 200, vsan 200]
  type[0] p_gwrap[(nil)]
    FREE @ 159645 usecs after Thu Aug  5 13:35:00 2010
  type[1] p_gwrap[(nil)]
    FREE @ 159964 usecs after Thu Aug  5 13:35:00 2010
  type[2] p_gwrap[(nil)]
    FREE @ 450507 usecs after Tue Aug  3 14:14:08 2010
0x50c8efc7
current state [TE_FSM_ST_ISOLATED_DM_ZS]
RNID info not found.
first time elp: 0
Peer ELP Revision: 3
```

REVIEW DRAFT – CISCO CONFIDENTIAL

- Step 3** Use the **fcdomain domain** *<domain-id>* [**static** | **preferred**] **vsan** *<vsan-id>* command to change the domain ID for one of the overlapping domain IDs.
- The **static** option tells the switch to request that particular domain ID. If it does not obtain that particular address, it will isolate itself from the fabric.
 - The **preferred** option has the switch request a specified domain ID. If that ID is unavailable, it will accept another ID.
- Step 4** Use the **fcdomain restart vsan** command to restart Domain Manager.
- While the static option can be applied to runtime after a disruptive or nondisruptive restart, the preferred option is applied to runtime only after a disruptive restart.

**Note**

A domain ID restart is disruptive. The Fibre Channel nodes that are logged into that domain will be logged out and logged back in. A disruptive reconfiguration might affect data traffic.

To assign a dynamic domain ID after a fabric reconfiguration

You can use fabric reconfiguration to reassign domain IDs and resolve any overlapping domain IDs. If you enable the auto-reconfigure option on both switches before connecting the fabric, a disruptive reconfiguration (RCF) occurs. The RCF functionality automatically forces a new principal switch selection and causes new domain IDs to be assigned to the different switches.

To use fabric reconfiguration to reassign domain IDs for a particular VSAN using the NX-OS CLI, perform the following these steps:

- Step 1** Use the **show fcdomain domain-list** command to determine if you have statically assigned domain IDs on the switches.
- Step 2** If you have statically assigned domain IDs, use the **no fcdomain domain** command to remove the static assignments.
- Step 3** Use the **show fcdomain vsan** *<vsan-id>* command to determine if you have the RCF reject option enabled.

Example:

```
switch# show fcdomain vsan 3
The local switch is the Principal Switch.
```

```
Local switch run time information:
```

```
State: Stable
Local switch WWN:    20:03:00:0d:ec:d5:fe:01
Running fabric name: 20:03:00:0d:ec:d5:fe:01
Running priority: 128
Current domain ID: 0x2c(44)
```

```
Local switch configuration information:
```

```
State: Enabled
FCID persistence: Enabled
Auto-reconfiguration: Disabled
Contiguous-allocation: Disabled
Configured fabric name: 20:01:00:05:30:00:28:df
Optimize Mode: Disabled
Configured priority: 128
Configured domain ID: 0x2c(44) (preferred)
```

REVIEW DRAFT – CISCO CONFIDENTIAL

Principal switch run time information:

```
Running priority: 128
Interface          Role          RCF-reject
-----
fc2/2              Isolated     Enabled
-----
```

- Step 4** If you have the rcf-reject option enabled, use the **interface** command and then the **no fcdomain rcf-reject vsan <vsan-id>** command in interface mode.

Example:

```
switch(config)# interface fc 2/2
switch(config-if)# no fcdomain rcf-reject vsan 3
switch(config-if)#
```

- Step 5** Use the **fcdomain auto-reconfigure vsan <vsan-id>** command in the EXEC mode on both switches to enable auto-reconfiguration after a Domain Manager restart.

- Step 6** Use the **fcdomain restart vsan <vsan-id>** command to restart Domain Manager.

This is a disruptive operation and disruptive reconfiguration and can affect data traffic.

Cisco Fabric Services

This section includes an overview of troubleshooting Cisco Fabric Services (CFS) followed by a description of common problems and their solutions.

Overview

Begin troubleshooting CFS issues by checking the following:

- Verify that CFS is enabled for the same applications on all affected switches.
- Verify that CFS distribution is enabled for the same applications on all affected switches.
If the CFS Regions feature is in use, verify that the application is in the same region on all the affected switches.
- Verify that there are no pending changes for an application and that a CFS commit was issued for any configuration changes in a CFS-enabled application.
- Verify that there are no unexpected CFS locked sessions.
Clear any unexpected locked sessions.

Verifying CFS using CLI

To verify CFS using the CLI, follow these steps:

- Step 1** By default, CFS distribution is enabled. Applications can distribute data and configuration information to all CFS-capable switches in the fabric where the applications exist. This is the normal mode of operation. To determine the state of CFS distribution on a switch, enter the **show cfs status** command.

Example:

REVIEW DRAFT – CISCO CONFIDENTIAL

```
switch(config)# show cfs status
Distribution : Enabled
Distribution over IP : Enabled - mode IPv4
IPv4 multicast address : 239.255.70.83
IPv6 multicast address : ff15::efff:4653
Distribution over Ethernet : Disabled

switch(config)# show cfs merge status name rscn
```

- Step 2** To verify that an application is listed and enabled, issue the **show cfs application** command to all switches.

Example:

```
switch# show cfs application

-----
Application      Enabled   Scope
-----
fwm               Yes      Physical-eth
ntp               No       Physical-fc-ip
stp               Yes      Physical-eth
fscm              Yes      Physical-fc
role              No       Physical-fc-ip
rscn              No       Logical
radius            No       Physical-fc-ip
fctimer           No       Physical-fc
syslogd           No       Physical-fc-ip
callhome          No       Physical-fc-ip
fcdomain          No       Logical
device-alias      Yes      Physical-fc

Total number of entries = 12
```

**Note**

The Physical scope means that CFS applies the configuration for that application to the entire switch. The Logical scope means that CFS applies the configuration for that application to a specific VSAN.

- Step 3** Verify the set of switches in which an application is registered with CFS, using the **show cfs peers name application-name** command for physical scope applications, and the **show cfs peers name <application-name> vsan <vsan-id>** command for logical scope applications.

Example:

```
switch# show cf peers name device-alias

Scope      : Physical-fc
-----
Switch WWN          IP Address
-----
20:00:00:0d:ec:da:6e:00 172.25.183.124          [Local]
20:00:00:0d:ec:24:5b:c0 172.25.183.123
20:00:00:0d:ec:50:09:00 172.25.183.42

Total number of entries = 3
```

**Note**

The **show cfs peers name <application-name>** command displays the peers for all VSANs when applied to a logical application.

Example:

REVIEW DRAFT – CISCO CONFIDENTIAL

```

switch(config)# show cfs peers name rscn

Scope      : Logical [VSAN 1]
-----
Domain Switch WWN          IP Address
-----
106  20:00:00:0d:ec:da:6e:00 172.25.183.124          [Local]
98   20:00:00:0d:ec:24:5b:c0 172.25.183.123
238  20:00:00:0d:ec:50:09:00 172.25.183.42

Total number of entries = 3

Scope      : Logical [VSAN 10]
-----
Domain Switch WWN          IP Address
-----
82   20:00:00:0d:ec:da:6e:00 172.25.183.124          [Local]
5    20:00:00:0d:ec:50:09:00 172.25.183.42
83   20:00:00:0d:ec:24:5b:c0 172.25.183.123

Total number of entries = 3

Scope      : Logical [VSAN 50]
-----
Domain Switch WWN          IP Address
-----
66   20:00:00:0d:ec:da:6e:00 172.25.183.124          [Local]
28   20:00:00:0d:ec:24:5b:c0 172.25.183.123
235  20:00:00:0d:ec:50:09:00 172.25.183.42

Total number of entries = 3

Scope      : Logical [VSAN 100]
-----
Domain Switch WWN          IP Address
-----
90   20:00:00:0d:ec:da:6e:00 172.25.183.124          [Local]
100  20:00:00:0d:ec:24:5b:c0 172.25.183.123
111  20:00:00:0d:ec:50:09:00 172.25.183.42

Total number of entries = 3

```

- Step 4** To determine if all the switches in the fabric constitute one CFS fabric, or a multitude of partitioned CFS fabrics, enter the **show cfs merge status name <application-name>** command and the **show cfs peers name <application-name>** command and compare the outputs. If the two outputs contain the same list of switches, the entire set of switches constitutes one CFS fabric. When this is the case, the merge status should always show success at all switches.

Example:

```

switch(config)# show cfs merge status name rscn

Logical [VSAN 1] Merge Status: Success [ Thu Aug  5 11:33:50 2010 ]
Local Fabric
-----
Domain Switch WWN          IP Address
-----
98   20:00:00:0d:ec:24:5b:c0 172.25.183.123          [Merge Master]
238  20:00:00:0d:ec:50:09:00 172.25.183.42
106  20:00:00:0d:ec:da:6e:00 172.25.183.124

```

REVIEW DRAFT – CISCO CONFIDENTIAL

switch

Total number of switches = 3

Logical [VSAN 10] Merge Status: Success [Thu Aug 5 11:36:43 2010]
Local Fabric

```
-----
Domain Switch WWN          IP Address
-----
83      20:00:00:0d:ec:24:5b:c0 172.25.183.123      [Merge Master]
5       20:00:00:0d:ec:50:09:00 172.25.183.42
82      20:00:00:0d:ec:da:6e:00 172.25.183.124
switch
```

Total number of switches = 3

Logical [VSAN 50] Merge Status: Success [Thu Aug 5 11:36:23 2010]
Local Fabric

```
-----
Domain Switch WWN          IP Address
-----
28      20:00:00:0d:ec:24:5b:c0 172.25.183.123      [Merge Master]
235     20:00:00:0d:ec:50:09:00 172.25.183.42
66      20:00:00:0d:ec:da:6e:00 172.25.183.124
switch
```

Total number of switches = 3

Logical [VSAN 100] Merge Status: Success [Thu Aug 5 11:33:50 2010]
Local Fabric

```
-----
Domain Switch WWN          IP Address
-----
100     20:00:00:0d:ec:24:5b:c0 172.25.183.123      [Merge Master]
111     20:00:00:0d:ec:50:09:00 172.25.183.42
90      20:00:00:0d:ec:da:6e:00 172.25.183.124
switch
```

Total number of switches = 3

If the list of switches in the **show cfs merge status name** command output is shorter than that of the **show cfs peers name** command output, then the fabric is partitioned into multiple CFS fabrics and the merge status may show that the merge has failed, is pending, or is waiting.

Merge failure troubleshooting

During a merge, the merge managers in the merging fabrics exchange their configuration databases with each other. The application on one of the fabrics merges the information, decides if the merge is successful, and informs all switches in the combined fabric of the status of the merge.

When a merge is successful, the merged database is distributed to all switches in the combined fabric and the entire new fabric remains in a consistent state. A merge failure indicates that the merged fabrics contain inconsistent data that could not be merged.

REVIEW DRAFT – CISCO CONFIDENTIAL

If a new switch is added to the fabric and the merge status for any application shows In Progress for a prolonged period of time, then there may be an active session for that application in a switch. Check the lock status for that application on all the switches by using the **show cfs lock** command. If any locks exist, the merge does not proceed. Commit the changes or clear the session lock so that the merge proceeds.

**Note**

Merge failures must be analyzed correctly. Exercise caution when choosing a switch for blank commit. Small configurations may wipe out the large configurations.

Recovering from a Merge Failure with the CLI

To recover from a merge failure using the CLI, perform the following steps:

- Step 1** To identify a switch that shows a merge failure, enter the **show cfs merge status name <application-name>** command.

Example:

```
switch(config)# show cfs merge status name ntp

Physical-fc-ip Merge Status: Success [ Thu Aug  5 11:47:58 2010 ]
Local Fabric
-----
Switch WWN                IP Address
-----
20:00:00:0d:ec:da:6e:00  172.25.183.124          [Merge Master]
                           switch
```

Total number of switches = 1

```
switch(config)# show cfs merge status name ntp

Physical-fc-ip Merge Status: Success [ Thu Aug  5 11:43:39 2010 ]
Local Fabric
-----
Switch WWN                IP Address
-----
20:00:00:0d:ec:50:09:00  172.25.183.42          [Merge Master]
                           MDS-9134
20:00:00:0d:ec:da:6e:00  172.25.183.124
```

Total number of switches = 2

- Step 2** For a more detailed description of the merge failure, enter the **show cfs internal session-history name <application name> detail** command.

Example:

```
switch(config)# show cfs internal session-history name ntp
-----
Time Stamp                Source WWN                Event
User Name                 Session ID
-----
Thu Aug  5 11:45:19 2010  20:00:00:0d:ec:da:6e:00  LOCK_ACQUIRED
admin                               34684
Thu Aug  5 11:45:19 2010  20:00:00:0d:ec:da:6e:00  COMMIT[2]
admin                               34689
Thu Aug  5 11:45:20 2010  20:00:00:0d:ec:da:6e:00  LOCK_RELEASED
```

REVIEW DRAFT – CISCO CONFIDENTIAL

```
admin                               34684
```

- Step 3** Enter configuration mode and enter the commit command for the application to restore all peers in the fabric to the same configuration database.

Example:

```
switch(config)# ntp commit
switch(config)# show cfs merge status name ntp
```

```
Physical-fc-ip Merge Status: Success [ Thu Aug  5 11:51:02 2010 ]
Local Fabric
```

Switch WWN	IP Address	
20:00:00:0d:ec:50:09:00	172.25.183.42	[Merge Master]
20:00:00:0d:ec:da:6e:00	172.25.183.124	

switch

Total number of switches = 2

```
switch(config)# show cfs merge status name ntp
```

```
Physical-fc-ip Merge Status: Success [ Thu Aug  5 11:51:02 2010 ]
Local Fabric
```

Switch WWN	IP Address	
20:00:00:0d:ec:50:09:00	172.25.183.42	[Merge Master]
	MDS-9134	
20:00:00:0d:ec:da:6e:00	172.25.183.124	

Total number of switches = 2

Lock failure troubleshooting

In order to distribute a configuration in the fabric, a lock must first be acquired on all switches in the fabric. Once accomplished, a commit can be issued which distributes the data to all switches in the fabric before releasing the lock.

When a lock has been acquired by another application peer, you cannot commit new configuration changes. This is a normal situation and you should postpone any changes to an application until the lock is released. Use the troubleshooting steps in this section only if you believe the lock has not been properly released.

A lock occurs when an administrator configures a change for a CFS-enabled application. If two administrators on the same switch attempt to configure the same application, only one administrator is given the lock. The other administrator is prevented from making changes to that application until the first administrator commits a change or discards any changes. Use the `show cfs lock name` command to determine the name of the administrator who holds the lock for an application. You should check with that administrator before clearing the lock.

A CFS lock can also be held by another switch in your fabric. Use the `show cfs peers name` command to determine all the switches that participate in the CFS distribution for the application. Then use the `show cfs lock name` command on each switch to determine who owns the CFS lock for that application. You should check with that administrator before clearing the lock. Use the CFS abort option to release the lock without distributing the data to the fabric.

REVIEW DRAFT – CISCO CONFIDENTIAL**Resolving lock failure issues using the CLI**

To resolve a lock failure using the CLI, perform the following steps:

- Step 1** Enter the **show cfs lock name <name>** command to determine the lock holder.

Example:

```
switch(config)# show cfs lock name ntp
```

```
Scope      : Physical-fc-ip
```

Switch WWN	IP Address	User Name	User Type
20:00:00:0d:ec:50:09:00	172.25.183.42	admin	CLI/SNMP v3

```
Total number of entries = 1
```

- Step 2** For a detailed description of the lock failure, enter the **show cfs internal session-history name application <name> detail** command.

Example:

```
switch(config)# show cfs internal session-history name ntp detail
```

Time Stamp	Source WWN	Event
User Name	Session ID	
Thu Aug 5 11:51:02 2010	20:00:00:0d:ec:da:6e:00	LOCK_REQUEST
admin	35035	
Thu Aug 5 11:51:02 2010	20:00:00:0d:ec:da:6e:00	LOCK_ACQUIRED
admin	35035	
Thu Aug 5 11:51:03 2010	20:00:00:0d:ec:da:6e:00	COMMIT[2]
admin	35040	
Thu Aug 5 11:51:03 2010	20:00:00:0d:ec:da:6e:00	LOCK_RELEASE_REQUEST
admin	35035	
Thu Aug 5 11:51:03 2010	20:00:00:0d:ec:da:6e:00	LOCK_RELEASED
admin	35035	
Thu Aug 5 12:03:18 2010	20:00:00:0d:ec:50:09:00	REMOTE_LOCK_REQUEST
admin	284072	
Thu Aug 5 12:03:18 2010	20:00:00:0d:ec:50:09:00	LOCK_OBTAINED
admin	284072	

- Step 3** If the lock is being held by a remote peer, enter the application-name commit command or an application-name abort command at that switch.

Example:

An example of the **<application-name> commit** command follows:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ntp commit
switch(config)#
```

Example:

An example of the **<application-name> abort** command follows:

```
switch# configure terminal
Enter configuration commands, one per line. End with CNTL/Z.
switch(config)# ntp abort
switch(config)#
```

REVIEW DRAFT – CISCO CONFIDENTIAL

System state inconsistent and locks being held

An inconsistent system state occurs for one of the following situations:

- When locks are not held on all of the switches in the fabric.
- When locks are held on all switches in the fabric, but a session does not exist with the lock holding the switch.

In either case, it is necessary to use the clear option to release the locks.

Clearing locks using the CLI

When a lock is being held on a remote peer and entering the `<application-name> commit` command or the `<application-name> abort` command does not clear the lock, issue the `clear <application-name> session` command to clear all locks in the fabric. After all locks are cleared, a new distribution must be started to restore all the switches in the fabric to the same state.

Example:

```
switch# clear ntp session
switch# config terminal
switch(config)# ntp commit
switch(config)#
```

Distribution status verification

After configuring an application and committing the changes, you may want to verify that CFS is distributing the configuration change throughout the fabric or VSAN.

Verifying distribution using the CLI

Use the `show cfs lock name <application-name>` command to determine if a distribution is in progress on the fabric. If the application does not show in the output, the distribution has completed.

Example:

```
switch(config)# show cfs lock name ntp

Scope      : Physical-fc-ip
-----
Switch WWN          IP Address          User Name    User Type
-----
20:00:00:0d:ec:50:09:00 172.25.183.42      admin       CLI/SNMP v3

Total number of entries = 1
```

CFS regions troubleshooting

The following rules apply to CFS Regions:

- When using CFS Regions, an application on a given switch can only belong to one region at a time.
- CFS Regions are only applicable to applications within the physical scope. You cannot create a CFS Region in the logical scope of an application.

REVIEW DRAFT – CISCO CONFIDENTIAL

- Assigning a region to an application takes precedence in distribution over its initial physical scope.
- CFS Regions configuration is not supported for deregistered applications (conditional services) or a physical scope application that is currently locked.
- Regions 1 through 200 are available for user configuration. Regions 201 through 255 are reserved regions and are not available for user configuration.

Distribution failure

To resolve a configuration distribution failure to all switches for a CFS Region, perform the following steps:

Step 1 Verify that application distribution is enabled. For more information, see “[Overview](#)” section on page 1-31.

Step 2 Verify that the application is in the same region on all switches.

Using the CLI from each switch, enter the **show cfs <application> name <application-name>** command.

Example (for device-alias application):

```
switch(config)# show cfs lock name ntp
```

```
Scope      : Physical-fc-ip
```

Switch WWN	IP Address	User Name	User Type
20:00:00:0d:ec:50:09:00	172.25.183.42	admin	CLI/SNMP v3

```
Total number of entries = 1
```

Example (application is capable of being merged; application is in Region Default):

```
switch(config)# show cfs application name device-alias
```

```
Enabled      : Yes
Timeout      : 20s
Merge Capable : Yes
Scope        : Physical-fc
Region       : Default
```

Example (application is capable of being merged; application is in Region 1):

```
switch# show cfs application name device-alias
```

```
Enabled : Yes
Timeout : 20s
Merge Capable : Yes
Scope : Physical-fc
Region : 1
```

REVIEW DRAFT – CISCO CONFIDENTIAL**Regions for conditional service**

When a conditional service goes down (deregisters with CFS), it loses its region configuration. When the conditional service is restarted, it is automatically placed into the default region. To avoid this situation, reconfigure the appropriate region information for the conditional service before starting it again.

Changing regions

If you move an application from one region to another, you might encounter a database mismatch when attempting a merge. To identify and resolve the conflicts, see [“Merge failure troubleshooting” section on page 1-34](#).

**Note**

When an application is moved from one region to another (including the default region), it loses all history.

VSANs

This section includes an overview of troubleshooting VSANs followed by a description of common problems and their solutions.

Overview

Most VSAN problems can be avoided by following the best practices for VSAN implementation.

However, if necessary, you can use the fabric analysis tool in Fabric Manager to verify different categories of problems such as VSANs, zoning, FC domain, admin issues, or switch-specific or fabric-specific issues.

Fabric Manager provides the configuration consistency check tool.

To use the Fabric Configuration option to analyze the configuration of a switch, follow these steps:

-
- Step 1** From the Fabric Manager tools menu, choose **Health > Fabric Configuration**.
The Fabric Configuration Analysis dialog box appears.
- Step 2** Determine whether you want to compare the selected switch to another switch or to a policy file.
- To compare the selected switch to another switch, select **Policy Switch** and then select a switch from the drop-down list of switches.
 - To make a policy file comparison, select **Policy File** and then click the button on the right to browse your file system and select a policy file (*.XML).
- Step 3** Click **Rules** to set the rules to apply when running the Fabric Configuration Analysis tool.
The Rules window appears.
- Step 4** Change the existing rules as appropriate and click **OK**.
- Step 5** Click **Compare** to have the system to compare the configuration.
The results of the analysis are displayed.

REVIEW DRAFT – CISCO CONFIDENTIAL

- Step 6** In the Resolve column, select the issues that you want to resolve.
- Step 7** Click **Resolve Issues** to resolve the identified issues.
- Step 8** Click **Clear** to remove the contents of the window.
- Step 9** Click **Close** to complete the operation and close the window.

For more information about the configuration consistency check tool, see the *Cisco DCNM Fundamentals Guide, Release 5.x*.

**Note**

When suspending or deleting VSANs, make sure that you suspend and unsuspend one VSAN at a time. You should wait a minimum of 60 seconds after you enter the **vsan suspend** command before you enter any other configuration command. If you fail to wait, some Fibre Channel interfaces or member ports in a port channel might become suspended or error-disabled.

Troubleshooting a SAN problem involves gathering information about the configuration and connectivity of individual devices as well as the status of the entire SAN fabric.

VSAN Troubleshooting Activities**Common troubleshooting tools in Fabric Manager**

Verify the VSAN with the following Fabric Manager procedures:

- To view the VSAN configuration in the Information pane, select **Fabricxx > VSANxx**.
- To view the VSAN members, select **Fabricxx > VSANxx**, then click the **Host** or **Storage** tab in the Information pane.
- To view the FC domain configuration in the Information pane, select **Fabricxx > VSANxx > Domain Manager**.

Common troubleshooting CLI commands

Use the following CLI commands to display VSAN, FC domain, and FSPF information:

```
show vsan
show vsan <vsan-id>
show vsan membership
show interface fc <slot/port> trunk <vsan-id>
show <vsan-id> membership
show vsan membership interface fc <slot/port>
```

Checklist

Check for the following:

- Verify the domain parameters for switches in the VSAN.
- Verify the physical connectivity for any problem ports or VSANs.
- Verify that both devices are in the name server.
- Verify that both end devices are in the same VSAN.
- Verify that both end devices are in the same zone.

REVIEW DRAFT – CISCO CONFIDENTIAL**Nexus 5000 trunk port does not connect to upstream SAN switch**

The Nexus 5000 trunk port does not connect to the upstream SAN switch because:

- Status of the trunk port connected to the upstream switch is isolated.
- The switch port trunk mode is enabled on both sides.
- Physical cabling has been checked and verified.
- Ports are up on both switches.

By examining the interface state and querying the interface, the issue is displayed as shown in the following example.

Example:

```
switch(config-if)# show interface brief
```

```
-----
Interface  Vsan  Admin  Admin  Status          SFP  Oper  Oper  Port
          Mode  Trunk  Mode
          Mode
fc2/3      1      E      on     isolated        swl  --   --   --
```

```
switch(config-if)# show interface fc 2/3
fc2/3 is down (Isolation due to no common vsans with peer on trunk)
  Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
  Port WWN is 20:43:00:0d:ec:da:6e:00
  Admin port mode is E, trunk mode is on
```

Possible Cause

The VSAN allow list for both interfaces is not the same. Specifically, there is no common VSAN allowed on both interfaces.

This situation might be caused by the following:

- No common VSANs on both switches.
- The trunk allowed VSAN members that do not contain common members.

In the example, the trunk VSAN allow list on the Nexus 5000 and MDS FC interfaces do not match.

Solution

Determine the connected ports and resolve the allowed VSANs on the trunk for both FC interfaces.

Example:

```
switch(config-if)# show run interface fc 2/3

!Command: show running-config interface fc2/3
!Time: Wed Aug  4 16:06:04 2010

version 4.2(1)N1(1)

interface fc2/3
  switchport mode E
  switchport trunk allowed vsan 1
  no shutdown

switch(config-if)# show run interface fc 1/1

!Command: show running-config interface fc1/1
!Time: Wed Aug  4 16:20:07 2010
```

REVIEW DRAFT – CISCO CONFIDENTIAL

```

version 5.0(1a)

interface fc1/1
  switchport rate-mode dedicated
  switchport mode E
  switchport trunk allowed vsan 100
  no shutdown

switch(config-if)# interface fc 2/3
switch(config-if)# switchport trunk allowed vsan
add all
switch(config-if)# switchport trunk allowed vsan add 100
switch(config-if)# show run interface fc 2/3

!Command: show running-config interface fc2/3
!Time: Wed Aug 4 16:07:25 2010

version 4.2(1)N1(1)

interface fc2/3
  switchport mode E
  switchport trunk allowed vsan 1
  switchport trunk allowed vsan add 100
  no shutdown

switch(config-if)# switchport trunk allowed vsan add 1
switch(config-if)# show run interface fc 1/1

!Command: show running-config interface fc1/1
!Time: Wed Aug 4 16:20:54 2010

version 5.0(1a)

interface fc1/1
  switchport rate-mode dedicated
  switchport mode E
  switchport trunk allowed vsan 1
  switchport trunk allowed vsan add 100
  no shutdown

fc2/3 is trunking
  Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
  Port WWN is 20:43:00:0d:ec:da:6e:00
  Peer port WWN is 20:01:00:0d:ec:24:5b:c0
  Admin port mode is E, trunk mode is on
  snmp link state traps are enabled
  Port mode is TE
  Port vsan is 1
  Speed is 4 Gbps
  Transmit B2B Credit is 250
  Receive B2B Credit is 16
  Receive data field Size is 2112
  Beacon is turned off
  Trunk vsans (admin allowed and active) (1,100)
  Trunk vsans (up) (1,100)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()

switch(config-if)# show interface brief

-----
Interface Vsan Admin Admin Status SFP Oper Oper Port
          Mode Trunk          Mode Speed Channel

```

REVIEW DRAFT – CISCO CONFIDENTIAL

			Mode				(Gbps)	
fc2/3	1	E	on	trunking	swl	TE	4	--

Nexus 5000 E port (non-trunking) does not connect to upstream SAN switch

The Nexus 5000 E port does not connect to the upstream SAN switch because:

- The status of the interconnected non-trunking E ports shows that the status is up. However, all Fibre Channel services are not working between the switches.
- Devices in the same VSAN do not appear in the FCNS database for both switches.
- The **show topology** command does not list peer switch information.
- Zones show members are not logged in.

Example:

```
switch(config-vsan-db)# show interface brief
```

```
-----
Interface  Vsan  Admin  Admin  Status          SFP  Oper  Oper  Port
          Mode  Trunk  Mode
          Mode
-----
fc2/4      50    E      off    up              swl  E     2    --
```

```
switch(config-if)# show interface brief
```

```
-----
Interface  Vsan  Admin  Admin  Status          SFP  Oper  Oper  Port
          Mode  Trunk  Mode
          Mode
-----
fc1fc1/2   100   E      off    up              swl  E     2    --
```

The FC topology does not show a valid peer interface.

Example:

```
switch(config-if)# show topo
```

```
FC Topology for VSAN 100 :
```

```
-----
Interface  Peer Domain Peer Interface  Peer IP Address
-----
fc1/2      0x42(66)   Port 65795    ::
```

The zoneset shows one member is not active

```
switch(config-vsan-db)# show zoneset active vsan 100
zoneset name ZoneSet_Host_Storage vsan 100
zone name Zone_Host_Storage vsan 100
pwwn 20:00:00:25:b5:00:20:0e [Host]
* fcid 0x5a0000 [pwwn 50:0a:09:81:86:78:39:66] [Storage]
```

```
switch(config-if)# show zoneset active vsan 100
zoneset name ZoneSet_Host_Storage vsan 100
zone name Zone_Host_Storage vsan 100
* fcid 0x640114 [pwwn 20:00:00:25:b5:00:20:0e] [Host]
pwwn 50:0a:09:81:86:78:39:66 [Storage]
```

REVIEW DRAFT – CISCO CONFIDENTIAL

The storage and hosts are in the correct VSAN.

Example:

```
switch(config-vsan-db)# show flogi database vsan 100
```

```
-----
INTERFACE          VSAN    FCID          PORT NAME          NODE NAME
-----
fc2/2              100    0x5a0000    50:0a:09:81:86:78:39:66 50:0a:09:80:86:78:39:66
                    [Storage]
```

```
switch(config-if)# show flogi database vsan 100
```

```
-----
INTERFACE          VSAN    FCID          PORT NAME          NODE NAME
-----
fc4/2              100    0x640114    20:00:00:25:b5:00:20:0e 20:00:00:25:b5:02:02:09
                    [Host]
```

Possible Cause

The error is displayed by the **show interface brief** command and the **show vsan membership** command. They show that the E port on one switch is in the wrong VSAN.

The non-trunking E port on one switch is in the wrong VSAN. (VSAN 100 is the correct VSAN.)

Example:

```
switch(config-if)# show interface brief
```

```
-----
Interface  Vsan   Admin  Admin  Status      SFP   Oper  Oper  Port
          Mode  Trunk  Mode
          Mode
-----
fc1fc1/2  100   E      off    up           sw1   E     2    --
```

Solution

Move the non-trunking E-port into VSAN 100.

Example:

```
switch(config-vsan-db)# vsan 100 interface fc 2/4
Traffic on fc2/4 may be impacted. Do you want to continue? (y/n) [n] y
```

The zone set is now active and the FC topology is correct.

Example:

```
switch(config-if)# show zoneset active vsan 100
zoneset name ZoneSet_Host_Storage vsan 100
  zone name Zone_Host_Storage vsan 100
    * fcid 0x640114 [pwwn 20:00:00:25:b5:00:20:0e] [Host]
    * fcid 0x5a0000 [pwwn 50:0a:09:81:86:78:39:66] [Storage]
switch(config-if)# show topology
```

FC Topology for VSAN 100 :

```
-----
Interface  Peer Domain Peer Interface  Peer IP Address
-----
fc1/2     0x5a(90)          fc2/4  172.25.183.124
```

REVIEW DRAFT – CISCO CONFIDENTIAL

Communication problem between host and storage devices

The communication problem between host and storage devices is because:

- Zones are active.
- Both host and storage are logged into the SAN.
- The storage port is not logged into the active zone set.

Example:

```
zoneset name ZoneSet_Host_Storage vsan 100
zone name Zone_Host_Storage vsan 100
* fcid 0x640114 [pwwn 20:00:00:25:b5:00:20:0e] [Host]
  pwwn 50:0a:09:81:86:78:39:66 [Storage]
```

Possible Cause

The host or storage port are in the wrong VSAN.

Example:

```
switch(config)# show fcns database

VSAN 50:
-----
FCID          TYPE  PWWN                               (VENDOR)          FC4-TYPE:FEATURE
-----
0x420000      N     50:0a:09:81:86:78:39:66 (NetApp)          scsi-fcp:target
                                   [Storage]
```

Solution

Move the storage port to the correct VSAN. (VSAN 100 is the correct VSAN in the example.).

Example:

```
switch(config)# show flogi database vsan 50
-----
INTERFACE      VSAN  FCID          PORT NAME          NODE NAME
-----
fc2/2          50    0x420000      50:0a:09:81:86:78:39:66 50:0a:09:80:86:78:39:66
                                   [Storage]
```

Total number of flogi = 1.

```
switch(config)# vsan database
switch(config-vsan-db)# vsan 100 interface fc 2/2
Traffic on fc2/2 may be impacted. Do you want to continue? (y/n) [n] y
switch(config-vsan-db)# show zoneset active vsan 100
zoneset name ZoneSet_Host_Storage vsan 100
zone name Zone_Host_Storage vsan 100
* fcid 0x640114 [pwwn 20:00:00:25:b5:00:20:0e] [Host]
* fcid 0x5a0000 [pwwn 50:0a:09:81:86:78:39:66] [Storage]
```

VSAN is down between switches

The VSAN is down between switches because:

- VSAN is configured on both switches.
- Trunk allow list allows the VSAN.

REVIEW DRAFT – CISCO CONFIDENTIAL

- VSAN reported to be down (Initializing state).
- Zones are active.
- Both host and storage are logged into the SAN.

In this failure, the storage port is not logged into the active zone set.

After examining the interface, the error can be seen as in the following example.

Example:

```
switch(config-if)# show interface fc 2/4 trunk vsan 10
fc2/4 is trunking
  Vsan 10 is down (Isolation due to domain id assignment failure)

switch(config-if)# show port internal info interface fc 2/4 | grep Isolation
  fc2/4, Vsan 10 - state(down), state reason(Isolation due to domain id assignment
failure), fcid(0x000000)
  fc2/4, Vsan 50 - state(down), state reason(Isolation due to vsan not configured on
peer), fcid(0x000000)
```

Possible Cause

The VSANs might have the same static Domain ID configured.

Example:

```
switch(config-if)# show fcdomain domain-list vsan 10

Number of domains: 1
Domain ID          WWN
-----          -
0x53(83)          20:0a:00:0d:ec:da:6e:01 [Local] [Principal]

switch(config)# show fcdomain domain-list vsan 10

Number of domains: 1
Domain ID          WWN
-----          -
0x53(83)          20:0a:00:0d:ec:24:5b:c1 [Local] [Principal]
```

Solution

Change the Domain ID on one of the VSANs.

Example:

```
switch(config)# vsan database
switch(config-vsan-db)# vsan 10 suspend
switch(config-vsan-db)# no vsan 10 suspend
switch(config-vsan-db)# show interface fc 2/4

Number of domains: 1
Domain ID          WWN
-----          -
0x52(82)          20:0a:00:0d:ec:da:6e:01 [Local] [Principal]

switch(config-vsan-db)# sho interface fc 2/4 | begin Trunk
  Trunk vsans (admin allowed and active) (1,10,50,100)
  Trunk vsans (up) (1,10,50,100)
  Trunk vsans (isolated) ()
  Trunk vsans (initializing) ()
```

REVIEW DRAFT – CISCO CONFIDENTIAL

Registers and Counters

Identifying physical layer issues

To troubleshoot physical layer issues with Fibre Channel SFP optics, use the following command:

```
switch# show interface fc x/y transceiver details
```

In the following example, you can see that the results of the command contains useful information such as supported speed, nominal bit rate, and link lengths supported for the SFP.

Example:

```
switch# show interface fc 3/1 transceiver details
fc3/1 sfp is present
  name is CISCO-FINISAR
  part number is FTLF8524P2BNL-C2
  revision is 0000
  serial number is FNS0928K161
  fc-transmitter type is short wave laser w/o OFC (SN)
  fc-transmitter supports intermediate distance link length
  media type is multi-mode, 62.5m (M6)
  Supported speed is 400 MBytes/sec
  Nominal bit rate is 4300 MBits/sec
  Link length supported for 50/125mm fiber is 150 m(s)
  Link length supported for 62.5/125mm fiber is 70 m(s)
  cisco extended id is unknown (0x0)

no tx fault, no rx loss, in sync state, Diag mon type 104
```

The command also provides detailed SFP diagnostic information and warnings and alarms, if any.

Example:

```
SFP Detail Diagnostics Information
-----
                Alarms                Warnings
                High                 Low                 High                 Low
-----
Temperature 41.50 C      95.00 C      -25.00 C      90.00 C      -20.00 C
Voltage      3.45 V       3.90 V       2.70 V       3.70 V       2.90 V
Current      7.18 mA      17.00 mA     1.00 mA     14.00 mA     2.00 mA
Tx Power     -4.41 dBm     -2.00 dBm   -11.74 dBm  -2.00 dBm   -11.02 dBm
Rx Power     -4.40 dBm     1.00 dBm   -20.00 dBm  -1.00 dBm   -18.24 dBm
Transmit Fault Count = 0
-----
Note: ++ high-alarm; + high-warning; -- low-alarm; - low-warning
```

Two example outputs from the command follow. The first shows a low alarm for Rx Power. The second shows low alarms for Tx, Rx, and Current. The interface for the second example was in an Error Disabled state due to the bit error rate being too high.

Low Alarm for RxPower

```
-----
                Alarms                Warnings
                High                 Low                 High                 Low
-----
Temperature 35.02 C      70.00 C      0.00 C      70.00 C      0.00 C
```


REVIEW DRAFT – CISCO CONFIDENTIAL

```

Voltage      0.00 V          0.00 V          0.00 V          0.00 V          0.00 V
Current     7.22 mA           16.00 mA        2.00 mA        14.00 mA        2.40 mA
Tx Power    -0.57 dBm         1.00 dBm       -8.21 dBm      0.00 dBm       -7.21 dBm
Rx Power    -18.86 dBm --     1.00 dBm      -16.58 dBm     0.00 dBm      -14.44 dBm

```

 Note: ++ high-alarm; + high-warning; -- low-alarm; - low-warning

Low Alarms for Current, Tx Power and R x Power

```

-----
Alarms
High      Low      Warnings
High      Low      High      Low
-----
Temperature 32.75 C      70.00 C      0.00 C      70.00 C      0.00 C
Voltage      0.00 V      0.00 V      0.00 V      0.00 V      0.00 V
Current     0.00 mA --   16.00 mA     2.00 mA     14.00 mA     2.40 mA
Tx Power    N/A --       1.00 dBm    -8.21 dBm    0.00 dBm     -7.21 dBm
Rx Power    -22.22 dBm -- 1.00 dBm   -16.58 dBm   0.00 dBm    -14.44 dBm

```

 Note: ++ high-alarm; + high-warning; -- low-alarm; - low-warning

In the following example that the command does not provide detailed transceiver information for Twinax (copper).

```

switch# sh interface ethernet 1/19 transceiver details
Ethernet1/19
  sfp is present
  name is Molex Inc.
  part number is 74752-1301
  revision is E
  serial number is 733010037
  nominal bitrate is 0 Mbits/sec
  Link length supported for 50/125mm fiber is 0 m(s)
  Link length supported for 62.5/125mm fiber is 0 m(s)
  cisco id is --
  cisco extended id number is 4

  Invalid calibration

```

Displaying FcoE bound Ethernet interface counters

The **show interface ethernet** command, has two versions, brief and detailed. Examples of each follow.

Brief Version

Example:

**Note**

Ensure that the jumbo frames are incrementing, as well as RX or TX pause frames counters, if any. Tx might indicate a congestion problem.

```

switch# show interface ethernet 1/4
Ethernet1/4 is up
  Hardware: 1000/10000 Ethernet, address: 000d.ecd5.a38b (bia 000d.ecd5.a38b)
  MTU 1500 bytes, BW 10000000 Kbit, DLY 10 usec,

```

REVIEW DRAFT – CISCO CONFIDENTIAL

```

reliability 255/255, txload 1/255, rxload 1/255
Encapsulation ARPA
Port mode is trunk
full-duplex, 10 Gb/s, media type is 10g

[snip]

RX
 9507 unicast packets  918874 multicast packets  3473 broadcast packets
931854 input packets  76225281 bytes
7121 jumbo packets  0 storm suppression packets
0 runts  0 giants  0 CRC  0 no buffer
0 input error  0 short frame  0 overrun  0 underrun  0 ignored
0 watchdog  0 bad etype drop  0 bad proto drop  0 if down drop
0 input with dribble  0 input discard
0 Rx pause

TX
3986 unicast packets  294583 multicast packets  36307 broadcast packets
334876 output packets  46873259 bytes
1227 jumbo packets
0 output errors  0 collision  0 deferred  0 late collision
0 lost carrier  0 no carrier  0 babble
2266 Tx pause
24 interface resets

```

Detailed Version

The output of the detailed version of the **show interface ethernet** command is shown in different parts in the following example. It includes both normal traffic counters, as well as counters for physical layer and protocol errors. These counters should be monitored anytime there is a connectivity or performance issue.

Example:

```
switch# sh interface ethernet 1/4 counters detailed all Ethernet1/4
```

```

64 bit counters:
0.          rxHCTotalPkts = 931881
1.          txHCTotalPkts = 335522
2.          rxHCUnicastPkts = 9507
3.          txHCUnicastPkts = 3986
4.          rxHCMulticastPkts = 918901
5.          txHCMulticastPkts = 295229
6.          rxHCBroadcastPkts = 3473
7.          txHCBroadcastPkts = 36307
8.          rxHCOctets = 76228116
9.          txHCOctets = 46926647
10.         rxTxHCPkts64Octets = 1065359
11.         rxTxHCpkts65to127Octets = 105246
12.         rxTxHCpkts128to255Octets = 43798
13.         rxTxHCpkts256to511Octets = 13822
14.         rxTxHCpkts512to1023Octets = 30742
15.         rxTxHCpkts1024to1518Octets = 88
16.         rxTxHCpkts1519to1548Octets = 0
17.         rxHCTrunkFrames = 895722
18.         txHCTrunkFrames = 69387
19.         rxHCDropEvents = 0

```

All Port Counters:

```

0.          InPackets = 931881
1.          InOctets = 76228116
2.          InUcastPkts = 9507
3.          InMcastPkts = 918901

```

REVIEW DRAFT – CISCO CONFIDENTIAL

```
4. InBcastPkts = 3473
5. InJumboPkts = 7121
6. StormSuppressPkts = 0
7. OutPackets = 335522
8. OutOctets = 46926647
9. OutUcastPkts = 3986
10. OutMcastPkts = 295229
11. OutBcastPkts = 36307
12. OutJumboPkts = 1227
13. rxHCPkts64Octets = 889975
14. rxHCPkts65to127Octets = 26702
15. rxHCPkts128to255Octets = 6072
16. rxHCPkts256to511Octets = 1913
17. rxHCpkts512to1023Octets = 11
18. rxHCpkts1024to1518Octets = 87
19. rxHCpkts1519to1548Octets = 0
20. txHCPkts64Octets = 175384
21. txHCPkts65to127Octets = 78544
22. txHCPkts128to255Octets = 37726
23. txHCPkts256to511Octets = 11909
24. txHCpkts512to1023Octets = 30731
25. txHCpkts1024to1518Octets = 1
26. txHCpkts1519to1548Octets = 0
27. ShortFrames = 0
28. Collisions = 0
29. SingleCol = 0
30. MultiCol = 0
31. LateCol = 0
32. ExcessiveCol = 0
33. LostCarrier = 0
34. NoCarrier = 0
35. Runts = 0
36. Giants = 0
37. InErrors = 0
38. OutErrors = 0
39. InputDiscards = 0
40. BadEtypeDrops = 0
41. IfDownDrops = 0
42. InUnknownProtos = 0
43. txErrors = 0
44. rxCRC = 0
45. Symbol = 0
46. txDropped = 0
47. TrunkFramesTx = 69387
48. TrunkFramesRx = 895722
49. WrongEncap = 0
50. Babbles = 0
51. Watchdogs = 0
52. ECC = 0
53. Overruns = 0
54. Underruns = 0
55. Dribbles = 0
56. Deferred = 0
57. Jabbers = 0
58. NoBuffer = 0
59. Ignored = 0
60. bpduOutLost = 0
61. cos0OutLost = 0
62. cos1OutLost = 0
63. cos2OutLost = 0
64. cos3OutLost = 0
65. cos4OutLost = 0
66. cos5OutLost = 0
67. cos6OutLost = 0
```

REVIEW DRAFT – CISCO CONFIDENTIAL

```

68.                cos7OutLost = 0
69.                RxPause = 0
70.                TxPause = 2266
71.                Resets = 0
72.                SQETest = 0
73.                InLayer3Routed = 0
74.                InLayer3RoutedOctets = 0
75.                OutLayer3Routed = 0
76.                OutLayer3RoutedOctets = 0
77.                OutLayer3Unicast = 0
78.                OutLayer3UnicastOctets = 0
79.                OutLayer3Multicast = 0
80.                OutLayer3MulticastOctets = 0
81.                InLayer3Unicast = 0
82.                InLayer3UnicastOctets = 0
83.                InLayer3Multicast = 0
84.                InLayer3MulticastOctets = 0
85.                InLayer3AverageOctets = 0
86.                InLayer3AveragePackets = 0
87.                OutLayer3AverageOctets = 0
88.                OutLayer3AveragePackets = 0

```

Understanding Fibre Channel interface counters

The **show interface** command is very useful when troubleshooting physical layer or performance issues with a Fibre Channel interface.

In the output of the command, observe the input/output counters and any input/output discards or errors.

When input discards increment, the FC packet does not have a valid route in the Forwarding Information Base (FIB). All packets not having a route are considered discards and sent to the supervisor. These packets are NOT dropped; however, they are policed before being sent to the supervisor. You should also check the MAC ASIC for errors.

When output discards increment, packets are timing out in egress because the egress is too slow. Check the attached device because it may be a slow draining receiver that is not responding, or not replenishing buffer credits. This causes back pressure to occur on the Nexus 5000 FC interface.

Example:

```

switch# show interface fc2/1
fc2/1 is trunking
  Hardware is Fibre Channel, SFP is short wave laser w/o OFC (SN)
  Port WWN is 20:41:00:0d:ec:a4:02:80

[snip]

  1 minute input rate 5048 bits/sec, 631 bytes/sec, 9 frames/sec
  1 minute output rate 6752 bits/sec, 844 bytes/sec, 9 frames/sec
  36398816 frames input, 2422447564 bytes
    0 discards, 0 errors
    0 CRC, 0 unknown class
    0 too long, 0 too short
  36368010 frames output, 3213593392 bytes
    0 discards, 0 errors
  1 input OLS, 1 LRR, 0 NOS, 0 loop inits
  1 output OLS, 2 LRR, 0 NOS, 0 loop inits
  16 receive B2B credit remaining
  250 transmit B2B credit remaining
  0 low priority transmit B2B credit remaining
Interface last changed at Thu Jan 28 18:26:30 2010

```

REVIEW DRAFT – CISCO CONFIDENTIAL**Troubleshooting Fibre Channel MAC issues**

The **show hardware** command is very useful when troubleshooting FC physical layer issues.

Show hardware internal fc-mac <x> port <y> statistics

The output of the command contains the following useful information:

- Physical layer information

FCP_CNTR_MAC_RX_LOSS_OF_SYNC - Loss of Sync received counter

- Performance information

FCP_CNTR_MAC_CREDIT_IG_XG_MUX_SEND_RRDY_REQ - Receiver Ready's Sent

FCP_CNTR_MAC_CREDIT_EG_DEC_RRDY - Receiver Ready's Received

- Class 3 normal traffic counters

FCP_CNTR_MAC_DATA_RX_CLASS3_FRAMES

FCP_CNTR_MAC_DATA_RX_CLASSF_FRAMES

FCP_CNTR_MAC_DATA_RX_CLASS3_WORDS

FCP_CNTR_MAC_DATA_RX_CLASSF_WORDS

FCP_CNTR_MAC_DATA_TX_CLASS3_FRAMES

FCP_CNTR_MAC_DATA_TX_CLASSF_FRAMES

FCP_CNTR_MAC_DATA_TX_CLASS3_WORDS

FCP_CNTR_MAC_DATA_TX_CLASSF_WORDS

- Fibre Channel primitive sequences

FCP_CNTR_LINK_RESET_IN - Link Resets Received

FCP_CNTR_OLS_OUT- Offline Sequences Sent

FCP_CNTR_NOS_OUT - Not Operational Sequence Sent

FCP_CNTR_LRR_OUT - Link Reset Responses Sent

FCP_CNTR_LINK_FAILURE

Example:

```
switch# show hardware internal fc-mac 2 port 1 statistics
```

ADDRESS	STAT	COUNT
0x0000003c	FCP_CNTR_MAC_RX_LOSS_OF_SYNC	0x5
0x0000003d	FCP_CNTR_MAC_RX_BAD_WORDS_FROM_DECODER	0xec
0x00000042	FCP_CNTR_MAC_CREDIT_IG_XG_MUX_SEND_RRDY_REQ	0x5ec
0x00000043	FCP_CNTR_MAC_CREDIT_EG_DEC_RRDY	0xc41
0x00000061	FCP_CNTR_MAC_DATA_RX_CLASS3_FRAMES	0x5d2
0x00000062	FCP_CNTR_MAC_DATA_RX_CLASSF_FRAMES	0x1a
0x00000069	FCP_CNTR_MAC_DATA_RX_CLASS3_WORDS	0x140b14
0x0000006a	FCP_CNTR_MAC_DATA_RX_CLASSF_WORDS	0xdcc
0x00000065	FCP_CNTR_MAC_DATA_TX_CLASS3_FRAMES	0xc24
0x00000066	FCP_CNTR_MAC_DATA_TX_CLASSF_FRAMES	0x1d
0x0000006d	FCP_CNTR_MAC_DATA_TX_CLASS3_WORDS	0x4b9538
0x0000006e	FCP_CNTR_MAC_DATA_TX_CLASSF_WORDS	0xabc
0xffffffff	FCP_CNTR_LINK_RESET_IN	0x2
0xffffffff	FCP_CNTR_OLS_OUT	0x5
0xffffffff	FCP_CNTR_NOS_OUT	0x2
0xffffffff	FCP_CNTR_LRR_OUT	0x7
0xffffffff	FCP_CNTR_LINK_FAILURE	0x2

REVIEW DRAFT – CISCO CONFIDENTIAL

When troubleshooting FC performance problems, review the R_RDY, Link Reset, and Link Reset Response counters. These help determine buffer to buffer credit problems that could lead to performance issues.

Troubleshooting Fibre Channel forwarding issues

To troubleshoot Fibre Channel forwarding issues, it is important to know that GATOS is the MAC/Forwarding ASIC on the Nexus 5000 switch. This section describes commands specific to this ASIC.

Each Fibre Channel interface is assigned a GATOS number. To understand forwarding issues, you must find the GATOS number for the specific FC interface.

Consider the following example:

```
switch# sh platform fwm info pif fc2/1
dump pif info: ifindex 0x1080000 dump_all 0 verbose 1
fc2/1: slot 1 port 0 state 0x0 pi_if 0x88bbb74 fwimpd ctx 0x889d4ec
fc2/1: oper_mode 0x1 rcvd_rbind: No
fc2/1: iftype 0x1 encap 0x5 bound_if? N #lifs 1 fwimpd ctx 0x88bd74c
fc2/1: lif_blk(pi) 0x8523da4 vif_id_alloc_bmp 0x887360c
fc2/1: cfg_lif_blk_size 0 lif_blk_base(pi) 1922 lif_blk_size(pi) 1
fc2/1: cfg_lif_blk_size(pi) 0
fc2/1: if_flags 0x0 num_sub_lif_tbls 0 Num HIFs pinned 0
fc2/1 pd: lif_entries 1 if_map_idx 49 if_lid 33 if_fcoe_lid 34
fc2/1 pd: reverse ifmap lookup 'same' ifmap_idx 49
fc2/1: SAT_HIF Port?: No
```

In the following part of the example, notice that `gatos_num 13` is the GATOS instance for Fibre Channel interface 2/1:

```
fc2/1 pd: slot 1 logical port num 4 gatos_num 13 fwm_inst 0 fc 0
fc2/1 pd: pif_type 'data fc'(2) hw_present 1 port map idx 49
fc2/1 pd: fabric a info: voq 0-1 port_id 29 connected 1 up 1
fc2/1 pd: fabric b info: voq 0-1 port_id 29 connected 1 up 1
fc2/1 pd: subported 1 primary 1 atherton 0
fc2/1 pd: sup_src_dst_if 17 lif_blk 0-0
fc2/1 pd: policer info: uc (sel 2) mc (sel 1) bc (sel 0)
fc2/1 pd: mac-addr 000d.eca4.02b4
```

In the following part of the example, notice that the command also provides forwarding drop and discard informatio:

```
fc2/1 pd: tx stats: bytes 4958178736 frames 36360131 discard 0 drop 0
fc2/1 pd: rx stats: bytes 2421909296 frames 36390924 discard 0 drop 0
```

You can also display GATOS errors for the GATOS instance that corresponds to the FC interface. In the following part of the example, notice that the command only shows non-zero counters.

```
switch# show platform fwm info gatos-errors 13
Printing non zero Gatos error registers:
DROP_FCF_SW_VSAN_IDX_MISS: res0 = 60 res1 = 0
DROP_FCF_SW_DOMAIN_IDX_MISS: res0 = 489036 res1 = 0
DROP_FCF_SW_TBL_MISS: res0 = 489036 res1 = 0
DROP_NO_FABRIC_SELECTED: res0 = 489036 res1 = 0
DROP_VLAN_MASK_TO_NULL: res0 = 489036 res1 = 0
```

In the first of the two parts, drops and discards were described. Notice that the drop and discard counters are separate for vEthernet and VFC interfaces. Review the output in the second example to help correlate the reason for the drops.

REVIEW DRAFT – CISCO CONFIDENTIAL

```

switch# show platform fwm info pif ethernet 1/4
dump pif info: ifindex 0x1a003000 dump_all 0 verbose 1
Eth1/4: slot 0 port 3 state 0x0 pi_if 0x876acb4 fwimpd ctx 0x876171c
Eth1/4: oper_mode 0x100000 rcvd_rbind: No
Eth1/4: iftype 0x1 encap 0x1 bound_if? Y #lifs 1 fwimpd ctx 0x879f70c
Eth1/4: lif_blk(pi) 0x87cc9a4 foo vif_id_alloc_bmp 0x88313f4
Eth1/4: 0
Eth1/4: cfg_lif_blk_size 0 lif_blk_base(pi) 512 lif_blk_size(pi) 128
Eth1/4: cfg_lif_blk_size(pi) 0
Eth1/4: if_flags 0x0 num_sub_lif_tbls 0 Num HIFs pinned 0
Eth1/4: max_hifpc_mbrs 0, max_hif_ports 0
Eth1/4 pd: lif_entries 1 if_map_idx 8 if_lid 35 if_fcoe_lid 36
Eth1/4 pd: reverse ifmap lookup 'same' ifmap_idx 8
Eth1/4: SAT_HIF Port?: No
Eth1/4 pd: slot 0 logical port num 3 gatos_num 0 fwm_inst 0 fc 0
Eth1/4 pd: pif_type 'data eth'(1) hw_present 1 port map idx 8
Eth1/4 pd: fabric a info: voq 0-7 port_id 55 connected 1 up 1
Eth1/4 pd: fabric b info: voq 0-7 port_id 55 connected 1 up 1
Eth1/4 pd: supported 0 primary 1 atherton 0
Eth1/4 pd: sup_src_dst_if 6 lif_blk 384-511
Eth1/4 pd: policer info: uc (sel 2) mc (sel 1) bc (sel 0)
Eth1/4 pd: mac-addr 000d.ecd5.a38b

```

In the following part of the example, notice the drops in the second line:

```

Eth1/4 pd: tx stats: bytes 50256531 frames 336488 discard 0 drop 0
Eth1/4 pd: rx stats: bytes 6718252 frames 77220 discard 0 drop 845482

```

In the following part of the example, notice that the Fcoe counters are separate:

```

Eth1/4 pd fcoe: tx stats: bytes 2927716 frames 3919 discard 0 drop 0
Eth1/4 pd fcoe: rx stats: bytes 15307492 frames 9470 discard 0 drop 0

```

In the following part of the example, notice that the command helps find the cause of the drops:

```

switch# show platform fwm info gatos-errors 0
Printing non zero Gatos error registers:
DROP_INGRESS_FW_PARSING_ERROR: res0 = 93 res1 = 0
DROP_SRC_VLAN_MBR: res0 = 2567226 res1 = 0
DROP_FCF_SW_DOMAIN_IDX_MISS: res0 = 2445 res1 = 0
DROP_FCF_SW_TBL_MISS: res0 = 2445 res1 = 0
DROP_NO_FABRIC_SELECTED: res0 = 2556 res1 = 0
DROP_VLAN_MASK_TO_NULL: res0 = 2556 res1 = 0
DROP_SRC_MASK_TO_NULL: res0 = 522 res1 = 0

```

REVIEW DRAFT – CISCO CONFIDENTIAL