



Troubleshooting Overview

This chapter introduces the basic concepts, methodology, and general troubleshooting guidelines for problems that may occur when configuring and using the Cisco Nexus 5000 Series switch.

This chapter includes the following sections:

- [Troubleshooting Basics](#)
- [Fabric Manager Tools and CLI Commands](#)
- [Failover](#)

Troubleshooting Basics

The following are the basic steps for troubleshooting:

-
- Step 1** Gather information that defines the specific symptoms.
 - Step 2** Identify all potential problems that could be causing the symptoms.
 - Step 3** Systematically eliminate each potential problem (from most likely to least likely) until the symptoms disappear.
-

To identify the possible problems, you need to use a variety of tools and understand the overall configuration. The following chapters in this guide describe many approaches and specific solutions to potential problems.

Troubleshooting a Switch Crash

When a switch crashes, the cause might be from the failure of a process, and results in a reload of the switch.

A crash is usually recorded with a core file on the switch and includes the reason for the crash, such as a failed process. The following can help you determine the cause of the crash:

- Use the **show version** or **show system reset-reason** commands to display the reason for the crash.

```
switch# show system reset-reason
Please look at Note Details
1) At 4054 usecs after Sat Nov 6 15:15:01 2010
   Reason: Reset triggered due to HA policy of Reset
```

REVIEW DRAFT – CISCO CONFIDENTIAL

```
Service: clis hap reset
Version: 4.2(1)N2(1)
```

```
2) At 841383 usecs after Sat Nov 6 14:56:25 2010
Reason: Reset triggered due to HA policy of Reset
Service: clis hap reset
Version: 4.2(1)N2(1)
```

- Use the **show cores** command to determine if a core file was recorded. You also can use the **show process log** command to display the processes and if a core was created.

```
switch#show process log
Process          PID      Normal-exit  Stack  Core  Log-create-time
-----
clis             4023      N           Y      Y     Sat Nov 6 15:14:53 2010
clis             4155      N           Y      N     Sat Nov 6 14:56:18 2010
```

- Use the **show processes log details** command to provide useful information about the reason for the crash:

```
switch# show processes log details
Service: clis
Description: CLI Server

Started at Sat Nov 6 14:59:10 2010 (882984 us)
Stopped at Sat Nov 6 15:14:53 2010 (614588 us)
Uptime: 15 minutes 43 seconds

Start type: SRV_OPTION_RESTART_STATELESS (23)
Death reason: SYSMGR_DEATH_REASON_FAILURE_SIGNAL (2)
Last heartbeat 9.35 secs ago
RLIMIT_AS: 687474867
System image name: n5000-uk9.4.2.1.N2.1.bin
System image version: 4.2(1)N2(1) S0

PID: 4023
Exit code: signal 11 (core dumped)

Threads: 4026 4024 4025
```

- Note the module-number and the PID number in the output of the **show cores** command for the process that crashed. (Usually the module number is 1 for a Nexus 5000 switch.)

```
switch#show cores
Module-num      Instance-num    Process-name    PID      Core-create-time
-----
1               1              clis            4023     Nov 6 15:20
```

- Use the **copy core://module-id/PID ftp:** command to export the file and contact the TAC to obtain an analysis of the file.
- Obtain the timestamp of the crash with the **show version**, **show system reset-reason**, or **show cores** commands. With the **show logging** command, review the events that happened just before the crash.

```
switch# show logging
[snip]
2010 Nov 6 08:00:50 TTPSW-5020SF1 %$ VDC-1 %$ %STP-2-BLOCK_BPDUGUARD: Received BPDU
on port Ethernet103/1/1 with BPDU Guard enabled. Disabling port.
2010 Nov 6 08:00:51 TTPSW-5020SF1 %$ VDC-1 %$ %ETHPORT-2-IF_DOWN_ERROR_DISABLED:
Interface Ethernet103/1/1 is down (Error disabled. Reason:BPDUGuard)
2010 Nov 6 14:56:18 TTPSW-5020SF1 %$ VDC-1 %$ %SYSMGR-2-SERVICE_CRASHED: Service
"clis" (PID 4155) hasn't caught signal 11 (core will be saved).
```

REVIEW DRAFT – CISCO CONFIDENTIAL**Best Practices**

Best practices are the recommended steps you should take to ensure the proper operation of your switch.

- Maintain a consistent Cisco NX-OS release across all your Cisco Nexus 5000 switches.
- Refer to the release notes for your Cisco SAN-OS release for the latest features, limitations, and caveats.
- Enable system message logging.
- Troubleshoot any new configuration changes after implementing the change.
- Use the Device Manager to manage your configuration and detect possible problems before they become critical.

Common Terms

Term	Description
DCBX	Data Center Bridging Exchange
RSTP+	Rapid Spanning Tree Protocol
FCoE	Fibre Channel over Ethernet
FCF	Fibre Channel Forwarder
FIP	FCoE Initialization Protocol
PFC	Priority Flow Control
ETS	Enhanced Transmission Selection
LLDP	Link Layer Discovery Protocol
CEE	Converged Enhanced Ethernet
VNTag	Virtual Network Tag
Lossless Ethernet	No-Drop Ethernet
CNA	Consolidated Network Adapter
HBA	Host Bus Adapter
NPV/NPIV	N Port Virtualizer
VN-Link	Virtual Network Link
FEX	Fabric Extender
PAA	Port Analyzer Adapter
RCF	Reconfigure Fabric
RSCN	Request State Change Notification
Menlo	Cisco FCoE MUX ASIC
FCP	Fibre Channel Protocol
FSPF	Fabric Shortest Path First

REVIEW DRAFT – CISCO CONFIDENTIAL

Fabric Manager Tools and CLI Commands

This section highlights the tools and CLI commands that are commonly used to troubleshoot problems. These tools and commands are a portion of what you may use to troubleshoot your specific problem.

The following chapters in this guide may describe additional tools and commands specific to the symptoms and possible problems covered in that chapter.

NX-OS Tips

Displaying what is required from the configuration

```
switch# show running-config interface
version 4.0(1a)N2(1)

interface vfc29
  no shutdown
  bind interface Ethernet1/29

interface fc2/3
  no shutdown
  switchport speed 1000
  switchport mode SD

interface fc2/4

interface Ethernet1/1
  speed 1000
```

Displaying within Config Mode

With NX-OS, you can display required data from within the configuration mode, so there is no need to back out to the switch prompt.

```
switch(config)# show run
switch(config)# show interface brief
```

Pipe command

```
switch# show logging |
  egrep      Egrep
  grep      Grep
  head      Stream Editor
  last      Display last lines
  less      Stream Editor
  no-more   Turn-off pagination for command output
  sed       Stream Editor
  wc        Count words, lines, characters
  begin     Begin with the line that matches
  count     Count number of lines
  exclude   Exclude lines that match
  include   Include lines that match
```

REVIEW DRAFT – CISCO CONFIDENTIAL**Using the pipe command to only display required keyword**

```
switch# show running-config | include switchport
system default switchport
switchport mode trunk
switchport trunk allowed vlan 1,18
switchport mode fex-fabric
switchport mode fex-fabric
switchport speed 1000
switchport mode SD
no system default switchport shutdown
```

Copy command

```
switch# copy ?
bootflash:      Select source filesystem
core:           Select source filesystem
debug:          Select source filesystem
ftp:            Select source filesystem
licenses        Backup license files
log:            Select source filesystem
modflash:       Select source filesystem
nvram:          Select source filesystem
running-config Copy running configuration to destination
scp:            Select source filesystem
sftp:           Select source filesystem
startup-config Copy startup configuration to destination
system:         Select source filesystem
tftp:           Select source filesystem
volatile:       Select source filesystem
```

Redirecting output

NX-OS allows you to redirect outputs to files and flash areas in the switch.

```
switch# show tech-support aaa > bootflash:ciscolive09

switch# dir
103557265   Apr 01 17:39:22 2009  .tmp-system
      12451   Apr 10 16:36:37 2009  ciscolive09
      49152   Apr 01 17:39:22 2009  lost+found/
20058112   Oct 21 13:10:44 2008  n5000-uk9-kickstart.4.0.0.N1.2.bin
20193280   Apr 01 17:36:37 2009  n5000-uk9-kickstart.4.0.1a.N2.1.bin
76930262   Oct 21 13:11:33 2008  n5000-uk9.4.0.0.N1.2.bin
103557265   Apr 01 17:37:30 2009  n5000-uk9.4.0.1a.N2.1.bin
      4096    Jan 01 00:03:26 2005  routing-sw/
```

Redirecting output of the show tech-support details command

Use the **tac-pac filename** command to redirect the output of the **show tech-support details** command to a file and then gzip the file.

REVIEW DRAFT – CISCO CONFIDENTIAL

The file is stored on bootflash://*filename* provided that there is enough memory available. If you do not specify a filename, NX-OS creates the file as volatile:show_tech_out.gz. Copy the file from the device using the procedure in the copy command section.

```
switch# tac-pac
switch# dir volatile:
374382 Aug 16 17:15:55 2010 show_tech_out.gz
```

From volatile, copy the file to the bootflash, FTP, or TFTP server.

```
switch# copy volatile:show_tech_out.gz ?
bootflash: Select destination filesystem
debug: Select destination filesystem
ftp: Select destination filesystem
log: Select destination filesystem
modflash: Select destination filesystem
nvram: Select destination filesystem
running-config Copy from source to running configuration
scp: Select destination filesystem
sftp: Select destination filesystem
startup-config Copy from source to startup configuration
system: Select destination filesystem
tftp: Select destination filesystem
volatile: Select destination filesystem
```

NX-OS command listing

```
switch# show cli list | include ?
-i    Ignore case difference when comparing strings
-x    Print only lines where the match is a whole line
WORD  Search for the expression

switch# show cli list | include debug | include interface
```

Narrowing scope of keywords

You can use many commands like **grep** and **include** to narrow the scope of a keyword.

```
switch(config-if)# show interface | grep fc
fc2/1 is trunking
fc2/2 is trunking
fc2/3 is up
fc2/4 is down (Administratively down)
vfc29 is up
```

Logging

You can use logging through the CLI or Device Manager. In the following examples, the **logging** command and the Device Manager display severity information:

Viewing Severity Information with the CLI


```
switch(config)# show logging
```

REVIEW DRAFT – CISCO CONFIDENTIAL

```

Logging console:          enabled (Severity: critical)
Logging monitor:         enabled (Severity: notifications)
Logging linecard:        enabled (Severity: notifications)
Logging fex:             enabled (Severity: notifications)
Logging timestamp:       Seconds
Logging server:          enabled
                          {10.91.42.134}
                          server severity:      notifications
                          server facility:      local7
                          server VRF:          management
Logging logflash:        disabled
Logging logfile:         enabled
Name - ciscolive09: Severity - debugging Size - 4194304

```

Viewing Severity Levels in the Device Manager**Ethalyzer and SPAN**

Ethalyzer is a tool that collects frames that are destined to, or originate from, the Nexus 5000 control plane. Node to switch or switch to switch traffic can be seen with this tool.

SPAN is a feature whereby frames that are transient to the switch are copied to a second port for analysis. Node to switch or node to node traffic can be seen via this method.

Ethalyzer

Ethalyzer is a Cisco NX-OS protocol analyzer tool based on the Wireshark open source code. This tool is a command-line version of Wireshark that captures and decodes packets. You can use Ethalyzer to troubleshoot your network and analyze the control-plane traffic.

REVIEW DRAFT – CISCO CONFIDENTIAL

Command	Description
ethanalyzer local sniff-interface	Captures packets sent or received by the supervisor and provides detailed protocol information.
ethanalyzer local sniff-interface brief	Captures packets sent or received by the supervisor and provides a summary of protocol information.
ethanalyzer local sniff-interface limit-captured-frames	Limits the number of frames to capture.
ethanalyzer local sniff-interface limit-frame-size	Limits the length of the frame to capture.
ethanalyzer local sniff-interface capture-filter	Filters the types of packets to capture.
ethanalyzer local sniff-interface display-filter	Filters the types of captured packets to display.
ethanalyzer local sniff-interface decode-internal	Decodes the internal frame header for Cisco NX-OS. Note Do not use this option if you plan to analyze the data using Wireshark instead of NX-OS Ethanalyzer.
ethanalyzer local sniff-interface write	Saves the captured data to a file.
ethanalyzer local sniff-interface read	Opens the captured data file and analyzes it.

Examples

```
switch# ethanalyzer local sniff-interface
No matches in current mode, matching in (exec) mode
  inbound-hi  Inbound(high priority) interface
  inbound-low Inbound(low priority) interface
  mgmt       Management interface
```

```
switch# ethanalyzer local sniff-interface mgmt brief
Capturing on eth0
2008-08-13 01:34:23.776519 10.116.167.244 -> 172.18.217.80 TCP 1106 > telnet [ACK] Seq=0
Ack=0 Win=64040 Len=0
2008-08-13 01:34:23.777752 172.18.217.80 -> 10.116.167.244 TELNET Telnet Data ...
2008-08-13 01:34:23.966262 00:04:dd:2f:75:10 -> 01:80:c2:00:00:00 STP Conf. Root =
32768/00:04:c1:0f:6e:c0 Cost = 57 Port = 0x801d
[snip]
```

The following example is for viewing the Spanning Tree Protocol (STP) and Fibre Channel: Using 0 in the command captures output until you press **Ctrl-C**. The FCID is a well-known name for switch domain controller.

```
switch# ethanalyzer local sniff-interface inbound-hi brief limit-captured-frames 0
```

```
Capturing on eth4

2008-08-13 01:37:16.639896 00:0d:ec:6b:cd:41 -> 01:80:c2:00:00:00 1 0 00:0d:ec:6b:cd:41 ->
01:80:c2:00:00:00 0x0 0x0 STP RST. Root = 32769/00:0d:ec:6b:cd:41 Cost = 0 Port = 0x8093
2008-08-13 01:37:18.639992 00:0d:ec:6b:cd:41 -> 01:80:c2:00:00:00 1 0 00:0d:ec:6b:cd:41 ->
01:80:c2:00:00:00 0x0 0x0 STP RST. Root = 32769/00:0d:ec:6b:cd:41 Cost = 0 Port = 0x8093

[snip]
```


REVIEW DRAFT—CISCO CONFIDENTIAL

```

2008-08-13 01:37:23.220253 00:0d:ec:6b:cd:40 -> fc:fc:fc:ff:ff:fd 4 0 ff.ff.fd ->
ff.ff.fd 0x5384 0xffff SW_ILS ELP
2008-08-13 01:37:23.220615 00:0d:ec:6b:cd:40 -> aa:bb:cc:dd:01:04 4 0 ff.ff.fd ->
ff.ff.fd 0x5384 0x2b3f FC Link Ctl, ACK1
2008-08-13 01:37:23.227202 00:0d:ec:6b:cd:40 -> aa:bb:cc:dd:01:04 4 0 ff.ff.fd ->
ff.ff.fd 0x5384 0x2b3f SW_ILS SW_ACC (ELP)
2008-08-13 01:37:23.229927 00:0d:ec:6b:cd:40 -> fc:fc:fc:ff:ff:fd 4 0 ff.ff.fd ->
ff.ff.fd 0x5384 0x2b3f FC Link Ctl, ACK1

```

Detailed BPD

```

switch# ethanalyzer local sniff-interface inbound-hi limit-captured-frames 0
Capturing on eth4
Frame 1 (57 bytes on wire, 57 bytes captured)
  Arrival Time: Aug 13, 2008 01:41:32.631969000
  [Time delta from previous captured frame: 1218591692.631969000 seconds]
  [Time delta from previous displayed frame: 1218591692.631969000 seconds]
  [Time since reference or first frame: 1218591692.631969000 seconds]
  Frame Number: 1
  Frame Length: 57 bytes
  Capture Length: 57 bytes
  [Frame is marked: False]
  [Protocols in frame: eth:vlan:llc:stp]
[snip]
  DSAP: Spanning Tree BPDU (0x42)
  IG Bit: Individual
  SSAP: Spanning Tree BPDU (0x42)
  CR Bit: Command
  Control field: U, func=UI (0x03)
    00. 00.. = Command: Unnumbered Information (0x00)
    .... ..11 = Frame type: Unnumbered frame (0x03)
[snip]

```

SPAN

The Switched Port Analyzer (SPAN) feature—sometimes called port mirroring or port monitoring—selects network traffic for analysis by a network analyzer. The network analyzer can be a Cisco SwitchProbe, a Fibre Channel Analyzer, or other Remote Monitoring (RMON) probes.

SPAN sources refer to the interfaces from which traffic can be monitored. The Cisco Nexus 5000 Series switch supports Ethernet, virtual Ethernet, Fibre Channel, virtual Fibre Channel, port channels, SAN port channels, VLANs, and VSANs as SPAN sources. With VLANs or VSANs, all supported interfaces in the specified VLAN or VSAN are included as SPAN sources. You can choose the SPAN traffic in the ingress direction, the egress direction, or both directions for Ethernet, virtual Ethernet, Fibre Channel, and virtual Fibre Channel source interfaces:

- Ingress source (Rx)—Traffic entering the switch through this source port is copied to the SPAN destination port.
- Egress source (Tx)—Traffic exiting the switch through this source port is copied to the SPAN destination port.

**Note**

For the Cisco Nexus 5548 Switch, Fibre Channel ports cannot be configured as ingress source ports in a SPAN session.

REVIEW DRAFT – CISCO CONFIDENTIAL**Source Ports**

A source port, also called a monitored port, is a switched interface that you monitor for network traffic analysis. The switch supports any number of ingress source ports (up to the maximum number of available ports on the switch) and any number of source VLANs or VSANs.

A source port has these characteristics:

- Can be of any port type: Ethernet, virtual Ethernet, Fibre Channel, virtual Fibre Channel, port channel, SAN port channel, VLAN, and VSAN.
- Cannot be monitored in multiple SPAN sessions.
- Cannot be a destination port.
- Each source port can be configured with a direction (ingress, egress, or both) to monitor. For VLAN, VSAN, port channel, and SAN port channel sources, the monitored direction can only be ingress and applies to all physical ports in the group. The rx/tx option is not available for VLAN or VSAN SPAN sessions.
- Beginning with Cisco NX-OS Release 5.0(2)N1(1). Port channel and SAN port channel interfaces can be configured as ingress or egress source ports.
- Source ports can be in the same or different VLANs or VSANs.
- For VLAN or VSAN SPAN sources, all active ports in the source VLAN or VSAN are included as source ports.
- The Cisco Nexus 5010 switch supports a maximum of two egress SPAN source ports. This limit does not apply to the Cisco Nexus 5020 Switch and the Cisco Nexus 5548 switch.

SPAN Destinations

SPAN destinations refer to the interfaces that monitors source ports. The Cisco Nexus 5000 Series switch supports Ethernet and Fibre Channel interfaces as SPAN destinations.

Source SPAN	Destination SPAN
Ethernet	Ethernet
Fibre Channel	Fibre Channel
Fibre Channel	Ethernet (FCoE)
Virtual Ethernet	Ethernet
Virtual Fibre Channel	Fibre Channel
Virtual Fibre Channel	Ethernet (FCoE)

Characteristics of Destination Ports

Each local SPAN session must have a destination port (also called a monitoring port) that receives a copy of traffic from the source ports, VLANs, or VSANs. A destination port has these characteristics:

- Can be any physical port, Ethernet, Ethernet (FCoE), or Fibre Channel. Virtual Ethernet and virtual Fibre Channel ports cannot be destination ports.
- Cannot be a source port.
- Cannot be a port channel or SAN port channel group.

REVIEW DRAFT – CISCO CONFIDENTIAL

- Does not participate in spanning tree while the SPAN session is active.
- Is excluded from the source list and is not monitored if it belongs to a source VLAN of any SPAN session.
- Receives copies of sent and received traffic for all monitored source ports. If a destination port is oversubscribed, it can become congested. This congestion can affect traffic forwarding on one or more of the source ports.

Monitor Caveats

Limitations of Nexus 5000 SPAN CoS values are not preserved at the monitor (span) destination.

- Packets coming in on the monitor source with an unknown VLAN tag are spanned out with a 0 VLAN tag (priority tag).
- For Ethernet destination, the monitor session is up only if the destination port is configured as switch port monitor.
- Out of 18 configurable sessions, only two are active (up state). The rest are in down state (hardware resource unavailable).


Configuration limitations: VLAN or port-channel cannot be configured as egress source

- VLAN or port channel cannot be a monitor destination.
- Only two egress sources supported.
- Only one destination port can be configured for a session.

SPAN Configuration

Example:

```
switch(config)# interface fc3/2
switch(config-if)# switchport mode sd
switch(config-if)# switchport speed 1000
switch(config-if)# no shut
switch(config-if)# exit
switch(config)# monitor session 1
switch(config-monitor)# source interface fc3/1 tx
switch(config-monitor)# source interface fc3/1 rx
switch(config-monitor)# destination interface fc3/2
```



Verifying the SPAN Session

Example:

```
switch# show monitor session
SESSION STATE REASON DESCRIPTION
-----
1 up The session is up
```

REVIEW DRAFT – CISCO CONFIDENTIAL

```

switch# show monitor session 1
  session 1
-----
type           : local
state          : up
source intf    :
  rx           : fc3/1
  tx           : fc3/1
  both        : fc3/1
source VLANs   :
  rx           :
source VSANs   :
  rx           :
destination ports : fc3/2

```

Suspending the SPAN Session

Example:

```

switch(config)# monitor session 1 suspend

switch(config)# show monitor session 1
  session 1
-----
type           : local
state          : down (Session suspended)
source intf    :
  rx           : fc3/1
  tx           : fc3/1
  both        : fc3/1
source VLANs   :
  rx           :
source VSANs   :
  rx           :
destination ports : fc3/2

```

Debugging**Command-Line Debugging**

Available debugs depend on features enabled in NX-OS. There are many different options to select when turning on debugs.

Determine the destination of the output:

- Logfile—Data file in switch memory.
- Capture to direct to screen via console, Telnet, or SSH.

You must have administrator privileges to run debugs. Debugs can only be run from the CLI.

Debug Logging

Set the log file as CiscoLive_debugs, using the **debug logfile** command. Then, use the **show debug** command to see name of the debug file.

REVIEW DRAFT – CISCO CONFIDENTIAL

```
switch# debug logfile CiscoLive_debugs
switch# show debug
```

Display debugging to the screen with the following command:

```
switch# show debug logfile CiscoLive_debugs
```

Copy the debug file from MDS to a server with the **copy** command. When you enter the VRF, if none is specified then the default is used.

```
switch# copy log: CiscoLive_debugs tftp:

Enter vrf: management
Enter hostname for the tftp server: 10.91.42.134
Trying to connect to tftp server.....
Connection to Server Established.
|
TFTP put operation was successful
```

To delete the debug logfile, use one of the following commands:

```
switch# clear debug-logfile CiscoLive_debugs

switch# undebug all
```

If you do not use one of these commands, the debug logfile will be cleared and overwritten when the next debug logfile is created. The system only allows one debug logfile to exist.

Debugs to the Direct Telnet Window

- Use a Telnet, SSH, or console application that captures the expected output to buffer or file.
- Undebug all or no debug of a specific debug command is required to turn trace off.
- The debugs are not persistent across reboots
- Most debugs are easy to read and understand, but some are not.

Cisco Discover Protocol

Cisco Discover Protocol (CDP) version 2 is applied to the physical Ethernet interface and only works when enabled at both ends of the link. LLDP standard is derived from CDP.

CDP is used to verify proper connectivity to correct network devices, very useful at switch deployment.

The following example shows the arguments that can be used with the **show CDP** command:

```
switch# show cdp
all          Show interfaces that are CDP enabled
entry       Show CDP entries in database
global      Show CDP global parameters
interface   Show CDP parameters for an interface
neighbors   Show CDP neighbors
traffic     Show CDP traffic statistics

switch# show cdp global
Global CDP information:
  CDP enabled globally
  Sending CDP packets every 60 seconds
  Sending a holdtime value of 180 seconds
  Sending CDPv2 advertisements is enabled
  Sending DeviceID TLV in Default Format
```

REVIEW DRAFT – CISCO CONFIDENTIAL

```

Device ID:TM-6506-1
System Name:
Interface address(es):
    IPv4 Address: 11.1.1.1
Platform: cisco WS-C6506, Capabilities: Router Switch IGMP Filtering
Interface: Ethernet1/4, Port ID (outgoing port): TenGigabitEthernet1/2 ? Verifies proper
port connections
Holdtime: 133 sec

Version:
Cisco Internetwork Operating System Software
IOS (tm) s72033_rp Software (s72033_rp-IPSERVICES_WAN-VM), Version 12.2(18)SXF11, RELEASE
SOFTWARE (fc1)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2007 by cisco Systems, Inc.
Compiled Fri 14-Sep-07 23:09 by kellythw

Advertisement Version: 2
Native VLAN: 1 ? Sent on Native VLAN
Duplex: full

```

Failover

FCoE Traffic

When the Nexus 5000 experiences loss of fabric connectivity, it brings down all the affected vFC interfaces.

The following methods are used to signal the host of loss of connectivity to the FC fabric

- FIP Clear Link Virtual Link to the CNA will be signaled to indicate the 'shut' state of vFC. Throughout the 'shut' period FCF Advertisements indicate 'not available for login'.
- In case the loss of connectivity is over the FCoE network, FIP keep-alives are used by the FCF and the CNA to timeout the login sessions. The keep-alive timers are configurable.

Non-FCoE traffic

Under certain failure scenarios where the access switch has lost all uplink connectivity to the aggregation layer, the CNA needs to be signaled of the loss of LAN connectivity. This helps the CNA failover the host traffic to the standby port. Traditionally, such a failure is signaled by bringing down the host facing link. Bringing down the link achieves two purposes:

- Host is signaled of loss of connectivity.
- The access switch stops forwarding traffic to and from the host-facing link.

However, in the converged network, even though the LAN connectivity is lost at the access switch, the SAN connectivity might still be intact. Bringing down the entire host-facing link is not desirable. Instead, the loss of connectivity is signaled over protocols. Loss of SAN connectivity is signaled using the FIP Clear Virtual Link message. Loss of LAN connectivity is signaled using logical link status TLVs defined in DCBX and VIC protocols.

REVIEW DRAFT – CISCO CONFIDENTIAL**LAN Traffic**

When LAN connectivity is lost for a particular VLAN on the uplinks, the VLAN is also brought down on the host-facing link.

Dedicating a VLAN solely for FCoE traffic helps with shutting down non-FCoE traffic to and from the host-facing link without disrupting FCoE traffic from the same host.

REVIEW DRAFT – CISCO CONFIDENTIAL