Troubleshoot Link Flap Issue on Nexus 9000

Contents

Introduction

This document describes how to troubleshoot the layer 1 link flap issue on Nexus 9000 switches.

Prerequisites

Requirements

Cisco recommends that you have a familiarity with the Cisco Nexus Operating System (NX-OS) and basic Nexus architecture before you proceed with the information that is described in this document.

Components Used

The information in this document is based on these software and hardware versions:

- N9K-C93180YC-FX
- nxos64-cs.10.2.6.M

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

Understand Link Flap Causes

A link flap is a networking issue where a physical interface on a switch, such as the Nexus 9000, continuously alternates between being up and down. This disruptive behavior can degrade network performance, destabilize the network, and interrupt communication, thereby causing significant inconveniences. Link flaps generally arise from faulty physical layers or protocol synchronization issues.

- Protocol Triggered Link Flap

Protocol-triggered link flaps occur when there is an issue with the protocol synchronization. This could involve protocols such as the Link Aggregation Control Protocol (LACP), Virtual Port-Channel and others. The issue can arise from protocol misconfiguration or packets lost, leading to link instability. Regular monitoring and timely software updates can help prevent this type of link flap.

- Layer 1 Physical Issue

Link flaps can also stem from Layer 1, the physical layer of the network. This often involves physical components such as cables and interfaces. Damaged, loose, or aging cables and malfunctioning interfaces can cause the link to flap. Regular physical inspections and maintenance, including cable checks and interface testing, can help identify and rectify these issues before they lead to link flaps.

This article focuses on layer 1 physical issue troubleshooting.

Identify Link Flap

Link flaps can be easily identified from logs. The example displays a link flap event on port E1/5, where the port goes down and then comes back up later.

<#root>

```
2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-FOP_CHANGED: port-channel100: first operational
2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel100: Ethernet1/5 is down
2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel
2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_BANDWIDTH_CHANGE: Interface port-channel100,bandwidth

2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet1/5 is down (Lin


2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel
2024 Jan 21 05:27:58 N9K-C93180YC-FX %ETHPORT-5-SPEED: Interface Ethernet1/5, operational speed changed
2024 Jan 21 05:27:58 N9K-C93180YC-FX %ETHPORT-5-IF_DUPLEX: Interface Ethernet1/5, operational duplex mod
2024 Jan 21 05:27:58 N9K-C93180YC-FX %ETHPORT-5-IF_RX_FLOW_CONTROL: Interface Ethernet1/5, operational
2024 Jan 21 05:27:58 N9K-C93180YC-FX %ETHPORT-5-IF_TX_FLOW_CONTROL: Interface Ethernet1/5, operational
2024 Jan 21 05:27:58 N9K-C93180YC-FX %ETHPORT-5-SPEED: Interface port-channel100, operational speed chan
2024 Jan 21 05:27:58 N9K-C93180YC-FX %ETHPORT-5-IF_DUPLEX: Interface port-channel100, operational duplex
2024 Jan 21 05:27:58 N9K-C93180YC-FX %ETHPORT-5-IF_RX_FLOW_CONTROL: Interface port-channel100, operation
2024 Jan 21 05:27:58 N9K-C93180YC-FX %ETHPORT-5-IF_TX_FLOW_CONTROL: Interface port-channel100, operation

2024 Jan 21 05:28:02 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-PORT_UP: port-channel100: Ethernet1/5 is up


2024 Jan 21 05:28:02 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-FOP_CHANGED: port-channel100: first operational
2024 Jan 21 05:28:02 N9K-C93180YC-FX %ETHPORT-5-IF_BANDWIDTH_CHANGE: Interface port-channel100,bandwidth
```

```
2024 Jan 21 05:28:02 N9K-C93180YC-FX %ETHPORT-5-IF_UP: Interface Ethernet1/5 is up in mode access
2024 Jan 21 05:28:02 N9K-C93180YC-FX %ETHPORT-5-IF_UP: Interface port-channel100 is up in mode access
```

Identify Layer 1 Link Flap or Protocol-triggered Link Flap

The Ethernet Port Manager (Ethpm) is a process that manages Ethernet interfaces. The Ethpm event history can be utilized to identify the cause of a link flap.

Layer 1 Flap Example

E1/5 experiences a link failure at 05:28:35, with the ethpm transition triggered by **ETH_PORT_FSM_EV_LINK_DOWN**. This indicates a Layer 1 flap.

```
<#root>

2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-chann
2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_BANDWIDTH_CHANGE: Interface port-channel100,bandwidt

2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet1/5 is down (Lin


2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-chann


N9K-C93180YC-FX# show system internal ethpm event-history interface e1/5


[143] 2024-01-21T05:26:02.100255000+00:00 [-] FSM:<Ethernet1/5> Transition:
Previous state: [ETH_PORT_FSM_ST_WAIT_BUNDLE_MEMBER_BRINGUP]
Triggered event: [ETH_PORT_FSM_EV_FIRST_BRINGUP_BUNDLE_MEMBER_DONE]
Next state: [ETH_PORT_FSM_ST_BUNDLE_MEMBER_UP]
[144]

2024-01-21T05:27:35.

783495000+00:00 [-] FSM:<Ethernet1/5> Transition:
Previous state: [ETH_PORT_FSM_ST_BUNDLE_MEMBER_UP]

Triggered event: [ETH_PORT_FSM_EV_LINK_DOWN]


Next state: [FSM_ST_NO_CHANGE]
```

LACP Triggered Flap Example

E1/8 enters an initializing down state at 07:40:07, with the ethpm transition triggered by **ETH_PORT_FSM_EV_EXTERNAL_REINIT_NO_FLAP_REQ**. This indicates a link flap triggered by Link Aggregation Control Protocol (LACP).

```
<#root>

2024 Jan 21 07:37:20 N9K-C93180YC-FX %ETHPORT-5-IF_UP: Interface port-channel200 is up in Layer3
2024 Jan 21 07:40:07 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-chann
2024 Jan 21 07:40:07 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-FOP_CHANGED: port-channel200: first operationa
2024 Jan 21 07:40:07 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel200: Ethernet1/8 is dow
```

```
2024 Jan 21 07:40:07 N9K-C93180YC-FX %ETHPORT-5-IF_BANDWIDTH_CHANGE: Interface port-channel200,bandwidt

2024 Jan 21 07:40:07 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_INITIALIZING: Interface Ethernet1/8 is down (Ini
```

<#root>

```
N9K-C93180YC-FX# show system internal ethpm event-history interface e1/8


[218] 2024-01-21T07:37:20.551880000+00:00 [-] FSM:<Ethernet1/8> Transition:
Previous state: [ETH_PORT_FSM_ST_WAIT_BUNDLE_MEMBER_BRINGUP]
Triggered event: [ETH_PORT_FSM_EV_FIRST_BRINGUP_BUNDLE_MEMBER_DONE]
Next state: [ETH_PORT_FSM_ST_BUNDLE_MEMBER_UP]

[219]

2024-01-21T07:40:07.104339000

+00:00 [-] FSM:<Ethernet1/8> Transition:
Previous state: [ETH_PORT_FSM_ST_BUNDLE_MEMBER_UP]
Triggered event:

[ETH_PORT_FSM_EV_EXTERNAL_REINIT_NO_FLAP_REQ]


Next state: [FSM_ST_NO_CHANGE]
```

Troubleshoot Layer 1 Link Flap

Cisco offers an extensive array of optical modules to accommodate a wide range of speeds, media, and distances. Before connecting a link to the Nexus 9000, ensure that the SFP and cable are compatible with your current software and hardware. You can verify this by:

Cisco Optics-to-Device Compatibility Matrix

Cisco Optics-to-Optics Interoperability Matrix

Layer 1 Issue on NX-OS 10.2.1 and Later Releases

Starting with NX-OS 10.2.1, the Platform Insights Engine (PIE) is supported on all Cloudscale ToR and EoR platforms. PIE is an on-switch real-time root cause analysis application.

Three PIEs can assist you in addressing the Layer 1 link flap issue.

Link Flap PIE

The link flap PIE analyzes link flap events published by user space drivers (USDs) and determines the root cause for a link flap. The PIE publishes the root cause analysis insight to the broker. Link flap events are published by the USDs (PIE client) when a link flaps. The USDs collect all of the relevant data from the ASIC and USD that is required for root cause analysis and publish the data to the broker. The link flap PIE analyzes the data and arrives at the most probable root cause for the flap.

Link Down PIE

The link down PIE finds the root cause for a link not coming up. The USD collects data about an interface when the interface is configured to be up, but the interface operating state is not up. This data is published to the PIE application. The link-down PIE subscribes to these events, receives the data from the broker, and analyzes the data to find the root cause.

Optics PIE

The optics PIE is a continuous monitoring engine that performs a time series analysis of the DOM data collected at regular intervals. By tracking various parameters in the DOM over a period, the PIE arrives at a metric to describe the state of optics for each optical port. The metric is an insight about the trending health of an optical transceiver.

For more information, refer to this PIE document:

Cisco Nexus 9000 Series NX-OS Platform Insights Engine Guide, Release 10.2(x)

PIE Example: Link Flap Caused by Shutting Down and then Re-Enabling the Port on the Peer Side

<#root>

```
2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-FOP_CHANGED: port-channel100: first operational
2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel100: Ethernet1/5 is down
2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel
2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_BANDWIDTH_CHANGE: Interface port-channel100,bandwidth
```

2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet1/5 is down (Lin

```
2024 Jan 21 05:27:35 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel
2024 Jan 21 05:27:58 N9K-C93180YC-FX %ETHPORT-5-SPEED: Interface Ethernet1/5, operational speed changed
<snip>
2024 Jan 21 05:28:02 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-PORT_UP: port-channel100: Ethernet1/5 is up
```

N9K-C93180YC-FX# show pie interface ethernet 1/5 link-flap-rca

2024-01-21 05:27:35

```
 Event Id: 00000068 Ethernet1/5 Source Id: 436209664 RCA Code: 41   >>>PIE event time
```

Reason: Link flapped/down due to Local Fault, check peer

```
 >>>PIE link flap reason
```

N9K-C93180YC-FX# show pie interface ethernet 1/5 transceiver-insights

```
2024-01-21 05:30:12 Event Id: 00000080 Event Class: xcvr DOM DB Event Interface: Ethernet1/5 Health Met
2024-01-21 05:28:12 Event Id: 00000072 Event Class: xcvr DOM DB Event Interface: Ethernet1/5 Health Met
```

PIE Example: Link Down Caused by Shutting Down the Port on the Peer Side

<#root>

```
2024 Jan 21 05:48:38 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-FOP_CHANGED: port-channel100: first operational
2024 Jan 21 05:48:38 N9K-C93180YC-FX %ETH_PORT_CHANNEL-5-PORT_DOWN: port-channel100: Ethernet1/5 is down
2024 Jan 21 05:48:38 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel
2024 Jan 21 05:48:38 N9K-C93180YC-FX %ETHPORT-5-IF_BANDWIDTH_CHANGE: Interface port-channel100,bandwidth
```

2024 Jan 21 05:48:38 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_LINK_FAILURE: Interface Ethernet1/5 is down (Lin

```
2024 Jan 21 05:48:38 N9K-C93180YC-FX %ETHPORT-5-IF_DOWN_PORT_CHANNEL_MEMBERS_DOWN: Interface port-channel
```

N9K-C93180YC-FX# show pie interface ethernet 1/5 link-down-rca

2024-01-21 05:48:48

Event Id: 00000197 Ethernet1/5 Source Id: 436209664 RCA Code: 16 >>>PIE event time

Reason: No PCS alignment detected. Please check Fec, speed, Autoneg configurations with peer

>>>Physical layer failed

N9K-C93180YC-FX# show pie interface ethernet 1/5 transceiver-insights

```
2024-01-21 05:50:12 Event Id: 00000199 Event Class: xcvr DOM DB Event Interface: Ethernet1/5 Health Met
2024-01-21 05:48:12 Event Id: 00000187 Event Class: xcvr DOM DB Event Interface: Ethernet1/5 Health Met
```

Replacing Faulty Parts

Based on the PIE output, it is recommended to replace the potentially faulty component and continue monitoring. If the link flap persists, then a swap test is needed to narrow down the faulty part. A swap test can be conducted by changing one component at a time while keeping everything else unchanged. Ultimately, the link stabilizes after the specific faulty component has been swapped out.

Layer 1 Issue on NX-OS 10.1.2 and Earlier Releases

For NX-OS software releases prior to 10.2(1), PIE support is not available. Several manual steps are required for checking layer 1 link flap.

Verifying Port-Client Event History

This lists all link events on the attached module. Debounce time refers to the duration that an interface waits before notifying the supervisor of a link going down. During this period, the interface waits to see if the link comes back up. This is used to determine whether the link has gone down or is just experiencing a minor flap.

<#root>

N9K-C93180YC-FX# attach module 1

module-1# show system internal port-client link-event

```
*************** Port Client Link Events Log ***************
---- ------ ----- ----- ------
Time PortNo Speed Event Stsinfo
---- ------ ----- ----- ------
```

Jan 21 05:48:38 2024 00122142 Ethernet1/5 ---- DOWN Link down debounce timer stopped and link is down

Jan 21 05:48:37 2024 00993003 Ethernet1/5 ---- DOWN Link down debounce timer started(0x40e50006)

Jan 21 05:45:14 2024 00432606 Ethernet1/5 10G UP SUCCESS(0x0)

Verifying the ASIC Events

These events provide detailed information about each link event.

<#root>

N9K-C93180YC-FX# attach module 1
module-1# show hardware internal tah link-events fp-port 5

324) Jan 21 05:48:37 2024 uSec 992843: Fp 5 : tahusd_isr.c #8469
Port Down with an ASIC interrupt
------------- ASIC MAC/PCS/Serdes REGS (Mac Channel 0) -------------

Link flapped due to Local Fault, check peer

>>>Local Fault means the local

device detected the issue on the receive path.

>>>

Remote Fault means a Local Fault is detected across the link.

Intr Regs 00:0x0000, 01:0x0000, 02:0x0000, 03:0x0010, 07:0x0000, 11:0x0000, 15:0x0000
sts2.bercount : 0x0f00 sts2.erroredblocks : 0x0000
bercounthi : 0x0000 erroredblockhi : 0x0000
counters0.syncloss : 0x0001 counters0.blocklockloss: 0x0001
counters1.highber : 0x0000 counters1.vlderr : 0x0000
counters2.unkerr : 0x0012 counters2.invlderr : 0x0000

Error Code	Explanation

sts2.erroredblocks	Counts errored blocks (higher order bits).
sts2.bercount	Counts bad sync headers (lower order bits).
bercounthi	Counts bad sync headers (higher order bits).
erroredblockhi	Counts errored blocks (higher order bits).
counters0.syncloss	Sync loss
counters0.blocklockloss	Block lock loss
counters1.highber	High BER
counters1.vlderr	Valid Error
counters2.unkerr	Unknown Error
counters2.invlderr	Invalid Error

Checking Digital Optical Monitoring (DOM) Information on Both Sides

There are several pieces of Small Form-factor Pluggable (SFP) information in this output. If any value falls outside the acceptable range in the SFP diagnostic, then the SFP is considered a potentially damaged component and needs to be replaced. In this example, everything is in good order.

<#root>

```
N9K-C93180YC-FX# show interface e1/5 transceiver details


Ethernet1/5
transceiver is present
type is 10Gbase-SR                    >>>SFP type
name is CISCO-OPLINK                  >>>SFP vendor
part number is TPP4XGDS0CCISE2G
revision is 02
serial number is OPMXXXXXXXX          >>>SFP SN
nominal bitrate is 10300 MBit/sec     >>>SFP bitrate
Link length supported for 50/125um OM2 fiber is 82 m
Link length supported for 62.5/125um fiber is 26 m
Link length supported for 50/125um OM3 fiber is 300 m
cisco id is 3
cisco extended id number is 4
cisco part number is 10-2415-03
cisco product id is SFP-10G-SR        >>>SFP PID
cisco version id is V03
```

```
SFP Detail Diagnostics Information (internal calibration)


--------------------------------------------------------------------------------
            Current                Alarms       Warnings
          Measurement           High    Low    High    Low
--------------------------------------------------------------------------------
Temperature

 36.52 C           75.00 C -5.00 C 70.00 C 0.00 C

Voltage

     3.28 V             3.63 V  2.97 V  3.46 V 3.13 V

Current

     6.61 mA          12.00 mA 0.50 mA 11.50 mA 1.00 mA

Tx Power

   -2.70 dBm        1.99 dBm -11.30 dBm -1.00 dBm -7.30 dBm

Rx Power

   -2.40 dBm        1.99 dBm -13.97 dBm -1.00 dBm -9.91 dBm
Transmit Fault Count = 0
--------------------------------------------------------------------------------
Note: ++ high-alarm; + high-warning; -- low-alarm; - low-warning
peer side information is snipped.
```

Swap Test and Replacement of Faulty Parts

If everything appears fine with the previous checks, then a swap test is needed to narrow down the faulty part. A swap test can be conducted by changing one component at a time while keeping everything else unchanged. Eventually, the link stabilizes after the specific faulty component has been swapped out.

Related Information

Nexus 9000 Datasheet

Nexus 9000 Interfaces Configuration Guide

Nexus 9000 Series NX-OS Platform Insights Engine Guide