

# Troubleshoot Nexus 5000 Port-channel Loadbalancing

## Contents

[Introduction](#)

[Prerequisites](#)

[Requirments](#)

[Components Used](#)

[Problem](#)

[Network Diagram](#)

[Troubleshoot](#)

[Scenario 1: Multicast traffic forwarding when FP enabled on vPC peer-link](#)

[Scenario 2: Multicast loadbalancing within port-channel](#)

[Solution](#)

[Usefeul Commands](#)

[Related Information](#)

[Known defects](#)

## Introduction

This document describes how to troubleshoot port-channel load balancing in Nexus 5000 switches with multicast traffic.

Contributed by Sivakumar Sukumar, Cisco TAC Engineer.

## Prerequisites

### Requirments

Cisco recommends that you have knowledge of these topics:

- Cisco Nexus 5672UP and router e.g ASR supports multicast
- Basic understanding of Virtual Port-channel (vPC) , fapric path(FP) and Multicat(MC) technology

### Components Used


This document is not restricted to specific software and hardware versions.

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, make sure that you understand the potential impact of any command.

# Problem

Multicast Traffic is not equally distributed between Port-Channels as well as links within Port-Channel.

## Network Diagram




Created by Paint X

## Troubleshoot

### Scenario 1: Multicast traffic forwarding when FP enabled on vPC peer-link

When fabricpath running only between vPC link, the multicast traffic from host traverses over Peer-link to upstream router

If fabricpath disabled (on vPC PL), then MC Traffic is distributed over the Port-channels to the L3 GW's (ASR) and does not transverse the vPC PL.




Created by Paint X

## Scenario 2: Multicast loadbalancing within port-channel

Traffic not evenly load-balanced and always uses one link inside a port-channel.

Port-channel 15 bundles eth 1/1 and eth 1/8



Created by Paat K

```
394(config-if)# sh int port-c 15 | i pps
input rate 248 bps, 0 pps; output rate 301.67 Mbps, 377.54 Kpps
input rate 248 bps, 0 pps; output rate 301.67 Mbps, 377.54 Kpps
394(config-if)# sh int eth 1/8 | i pps
input rate 168 bps, 0 pps; output rate 280.01 Mbps, 145.79 Kpps
394(config-if)# sh int eth 1/1 | i pps
input rate 80 bps, 0 pps; output rate 10.08 Mbps, 231.76 Kpps
```

## Solution

When MC traffic hitting SVI on Nexus is on a FP enabled VLAN, switch forwards via one of the MC Forwarding Tag (FTag). Refer the link below to know more about FTag. Since there is only one FP interface on the switch i.e the vPC peer-link, the FTag tree prefers peer-link interface as soon as enters the switch. The traffic cannot go to upstream interfaces (to the router) because they are not FP enabled interfaces

```
show fabricpath isis topology summary
FabricPath IS-IS Topology Summary
Fabricpath IS-IS domain: default
MT-0
Configured interfaces: port-channel99
Max number of trees: 2 Number of trees supported: 2
Tree id: 1, ftag: 1, root system: 002a.6ab9.20c1, 3941
Tree id: 2, ftag: 2 [transit-traffic-only], root system: 002a.6ab6.9ac1, 3940
Ftag Proxy Root: 002a.6ab9.20c1 show fabricpath switch-id
Total Switch-ids: 4
=====
SWITCH-ID  SYSTEM-ID  FLAGS  STATE  STATIC EMULATED/ANYCAST
-----+-----+-----+-----+-----+-----
[E] 394    002a.6ab6.9ac1 Primary Confirmed No    Yes
```

```

394          002a.6ab9.20c1 Primary Confirmed No    Yes
* 3940       002a.6ab6.9ac1 Primary Confirmed Yes   No
3941        002a.6ab9.20c1 Primary Confirmed Yes   No show fabricpath isis database detail | egrep

```

**"Hostname|Affinity|Numgraphs"**

```
Hostname : 394 Length : 14
```

```
Affinity :
```

```
Nickname : 394 Numgraphs: 1 Graph-id: 1
```

```
Hostname : 395 Length : 14
```

```
Affinity :
```

```
Nickname : 394 Numgraphs: 1 Graph-id: 2
```

Verify if the MC traffic comes with different src/dest ip/mac/port to make an optimal hash using loadbalancing algorithm configured on the switch. Run commands above to check if there is any problem with port-channel load balancing.

```
394(config-if)# show mac address-table
```

Legend:

\* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC

age - seconds since last seen,+ - primary entry using vPC Peer-Link

VLAN	MAC Address	Type	age	Secure	NTFY	Ports/SWID.SSID.LID
+ 925	0000.0000.0a01	dynamic	0	F	F	3339.0.0
+ 925	0000.0000.0a4f	dynamic	0	F	F	3339.0.0
+ 925	0000.0000.0b11	dynamic	0	F	F	3339.0.0
+ 925	0000.0037.4e8d	dynamic	0	F	F	3339.0.0
* 925	002a.6a31.5f41	static	0	F	F	3339.0.0

```
394(config-if)# show int port-c 15 | i pps
```


```
input rate 248 bps, 0 pps; output rate 301.67 Mbps, 377.54 Kpps 394(config-if)# show int eth
```

```
1/8 | i pps
```

```
input rate 168 bps, 0 pps; output rate 175.60 Mbps, 145.79 Kpps
```

```
394(config-if)# sh int eth 1/1 | i pps
```

```
input rate 80 bps, 0 pps; output rate 126.08 Mbps, 231.76 Kpps
```



Created by Paint X

```

394(config-if)# show port-channel load-balance forwarding-path interface port-channel 15 vlan
925 src-ip 10.1.1.1 dst-ip 231.1.1.1 dst-mac 0100.5e01.0101 src-mac 0000.0037.4e8d
Missing params will be substituted by 0's.
Load-balance Algorithm on switch: source-dest-ip
crc_hash: 231 Polynomial: CRC10b      Outgoing port id  Ethernet1/8
Param(s) used to calculate load-balance:
  seed: 0xe
  vlan: 0x39d
  dst-ip: 231.1.1.1
  src-ip: 10.1.1.1
  dst-mac: 0100.5e01.0101
394(config-if)# show port-channel load-balance forwarding-path
interface port-channel 15 vlan 925 src-ip 10.1.1.2 dst-ip 231.1.1.2 dst-mac 0100.5e01.0102 src-
mac 0000.0000.0a01
Missing params will be substituted by 0's.
Load-balance Algorithm on switch: source-dest-ip
crc_hash: 250 Polynomial: CRC10b Outgoing port id  Ethernet1/1
Param(s) used to calculate load-balance:
  seed: 0xe
  vlan: 0x39d
  dst-ip: 231.1.1.2
  src-ip: 10.1.1.2
  dst-mac: 0100.5e01.0102
  src-mac: 0000.0000.0a01

```

## Usefeul Commands

- **show port-channel load-balance forwarding-path interface port-channel <num> vlan src-ip dst-ip dst-mac src-mac**
- **show fabricpath isis database detail | egrep "Hostname|Affinity|Numgraphs"**
- **show system internal rtm sdb ftag multicast**

- **show fabricpath isis trees multideestination 1**
- **show fabricpath route switchid**
- **show fabricpath isis topology summary**

## Related Information

- <http://www.cisco.com/c/en/us/support/docs/switches/nexus-5000-series-switches/116303-technote-nexus-00.html>
- <http://www.cisco.com/c/en/us/support/docs/switches/nexus-7000-series-switches/117297-technote-rpf-00.html>
- [http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/6\\_x/nx-os/fabricpath/configuration/guide/b-Cisco-Nexus-7000-Series-NX-OS-FP-Configuration-Guide-6x/b-Cisco-Nexus-7000-Series-NX-OS-FP-Configuration-Guide-6x\\_chapter\\_0100.html#concept\\_1ADF06ED94EE493AB8C5906B65029F80](http://www.cisco.com/c/en/us/td/docs/switches/datacenter/sw/6_x/nx-os/fabricpath/configuration/guide/b-Cisco-Nexus-7000-Series-NX-OS-FP-Configuration-Guide-6x/b-Cisco-Nexus-7000-Series-NX-OS-FP-Configuration-Guide-6x_chapter_0100.html#concept_1ADF06ED94EE493AB8C5906B65029F80)

## Known defects

- Cisco Bug ID [CSCvb13924](#) vPC+ multicast flooded on peer-link irrespective of affinity
- Cisco Bug ID [CSCts77757](#) L3 PO load-balance displays incorrect interface