

FabricPath Reverse Path Forwarding Check for a VPC+ Emulated Switch-ID



Document ID: 117297

Contributed by Al Bryant, Cisco TAC Engineer.
Jan 17, 2014

Contents

Introduction

Prerequisites

- Requirements
- Components Used

Topology

Background Information

Procedure

- Step 1. Check the RPF State
 - Software Check
 - Hardware Check
- Step 2. Check the vPC+ Affinity

Affinity Selection

Conclusion

Introduction

This document describes how to determine which interface to use for the Reverse Path Forwarding (RPF) check with an Emulated Switch-ID (ESID).

Prerequisites

Requirements

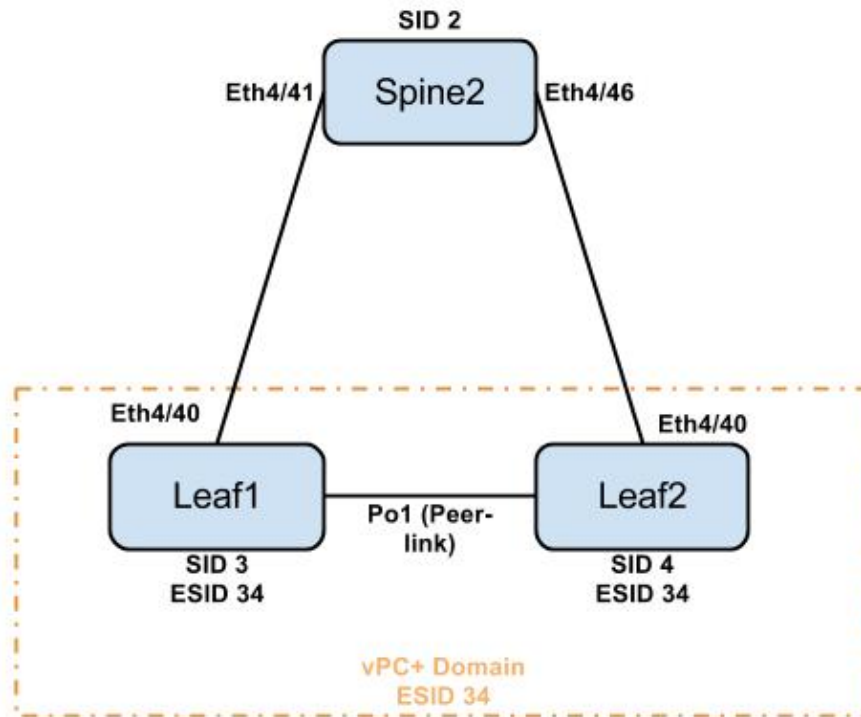
Cisco recommends that you have knowledge of basic FabricPath terminology and concepts. Specifically, users should be familiar with Forwarding Tags (FTAGs) and how they are used. For more information on FTAGs, see Mapping Out the Multidestination Tree for a FTAG.

Note: The terms FTAG, Trees, and FTAG Trees are used interchangeably in this document. At the time this document was written, only FTAG1 and FTAG2 have been implemented. There might be additional FTAGs in the future.

Components Used

The information in this document is based on the Nexus 7000. Some commands might vary on other Nexus platforms, but the concepts will remain the same.

Topology



Note: This is only a partial topology. There is another Spine1 connected to Spine2 and both Leaf1 and Leaf2, but it is excluded from this document for simplicity. Spine1 is the root for FTAG1 and Spine2 is the root for FTAG2.

Background Information

In a FabricPath domain, each switch is identified by a switch-ID (SID). In the case of a Virtual Port Channel+ (vPC+) setup, each peer switch has two SIDs. One for the physical switch and one for the logical switch that is formed by vPC+. The logical SID is shared on both peers and is known as an ESID.

When a multidestination packet is received on a vPC+ interface and forwarded into the FabricPath domain, the packet is encapsulated with a FabricPath header and the source SID is set to the ESID. Once this packet is received on a FabricPath core port, an RPF check is performed in order to ensure the packet was received on the correct interface. In short, the RPF check cross references the source SID and FTAG of the packet to the interface it was received on. If it was not received on the RPF interface, the packet fails the RPF check and is dropped.

RPF checks only come into play for multidestination packets which include unknown unicast, multicast, and broadcast packets. RPF checks are not performed on unicast packets because they can come in multiple interfaces with Equal-Cost Multi-Pathing. Multidestination packets should only come in one interface per source SID based on the FTAG tree.

In this document, the RPF check for an ESID is investigated. An ESID represents a vPC+ domain. This means it is one logical switch that consists of two physical switches. If you consider the topology diagram, you will see that there are two physical links to one ESID which is typically the case. Remember, the RPF check only allows multidestination packets to be accepted in one interface based on the FTAG tree.

Procedure

Step 1. Check the RPF State

In this example, Spine2 is the root for FTAG2. First, check Spine2 for the current RPF interface for the ESID 34 and FTAG2. There are two places to check the RPF status, which are in the software and in the hardware. Both outputs show that that Eth4/41 is the RPF interface for FTAG2, ESID 34.

Software Check

```
Spine2# show 12 multicast trees
```

```
<snip>
```

```
ftag/2, topo/0, Switch-id 34), uptime: 04:55:59, isis
  Outgoing interface list: (count: 1, '*' is the preferred interface)
* Interface Ethernet4/41, [RPF] [admin distance/115] uptime: 02:09:32, isis
```

Hardware Check

First, determine the VDC number and the module the packet ingresses.

```
Spine2# show vdc
```

```
Switchwide mode is m1 f1 m1x1 f2 m2x1
```

vdc_id	vdc_name	state	mac	type	lc
6	Spine2	active	84:78:ac:0b:60:46	Ethernet	f2

Then, attach to the module of the receiving interface in order to check the hardware state.

```
Spine2# attach module 4
```

```
module-4# show fabricpath unicast routes vdc 6 ftag 2 switchid 34
```

```
Route in VDC 6
```

```
-----
FTAG | SwitchID | SubSwitchID | Loc/Rem | RPF | RPF Intf | Num Paths | Merge V
-----
0002 | 0034 | 0000 | Remote | Yes | Eth4/41 | 0 | 1
-----
```

```
PD Information for Prefix:
```

FE num	ADDR TYPE	HTBL ADDR	TCAM ADDR	SWSI
10	HASH TABLE	00001440	000000ff	0000124c
11	HASH TABLE	00001440	000000ff	0000124c

Step 2. Check the vPC+ Affinity

In order to understand why Spine2 chose Eth4/41 to Leaf1 instead of Eth4/45 to Leaf2 when both are part of the ESID, you must first understand the concept of an affinity.

In a non-vPC+ setup, multidestination packets are primarily forwarded on the first tree in the topology, which is Tree1. In a vPC+ environment, each tree (FTAG1 or FTAG2) has an affinity for one or the other peer

switch. In this situation, broadcast, unknown unicast, and non-IP multicast frames traverse the tree for which the particular peer switch has an affinity.

You need to check which peer has the affinity for FTAG2. In order to do this, log in to one of the leaf switches in the vPC+ domain. The check can be done from either peer switch because both should know their peer's affinity as well as their own.

```
Leaf1# show fabricpath isis database detail / i Affinity/Hostname/Nickname
```

<snip>

```
Hostname      : Leaf2                Length : 5
Affinity      :
  Nickname: 34 Numgraphs: 1 Graph-id: 1      Nickname      :
  Priority: 0 Nickname: 4 BcastPriority: 64
  Priority: 0 Nickname: 34 BcastPriority: 0
  Nickname Migration :

Hostname      : Leaf1                Length : 5
Affinity      :
  Nickname: 34 Numgraphs: 1 Graph-id: 2
  Nickname Migration :
  Nickname      :
  Priority: 0 Nickname: 3 BcastPriority: 64
  Priority: 0 Nickname: 34 BcastPriority: 0
```

With this information you can see that if a multdestination frame is received on Leaf1, the switch forwards it along the FTAG2 tree. If the frame is received on Leaf2, it is forwarded along the FTAG1 tree.

For more information on FTAG trees, see Mapping Out the Multidestination Tree for a FTAG.

This information is used on Spine2 in order to build the RPF state for the ESID 34. This is Eth4/41 is used as the RPF interface for FTAG2 ESID 34.

Affinity Selection

The FTAG affinity is selected with this method:

- There is a ranking system based off of the SID. The SID is derived from the pool of MAC addresses allocated to the chassis in question.
- The lowest system ID has the highest rank.
- The vPC+ peer with the lowest system ID, and therefore highest rank, has the affinity for the FTAG1.
- The second lowest system ID, and therefore second highest rank, has the affinity for the FTAG2.

Conclusion

The RPF state is built in accordance with the vPC+ affinity state on each peer. The RPF interface to an ESID for a given FTAG is the interface that connects to the peer switch with the affinity.