

# BFD for BGP in AdminDown State on Nexus7000

## Contents

[Introduction](#)

[Background Information](#)

[Problem: BFD doesn't come up for BGP over an SVI on Nexus 7000](#)

[Analysis](#)

[Solution](#)

## Introduction

This document describes the problem with Bidirectional Forwarding Detection (BFD), which doesn't come up for Border Gateway Protocol (BGP) over a Switch Virtual Interface (SVI) on Nexus 7000.

## Background Information

BFD for BGP is supported on Nexus starting 5.0(2) release but there are certain limitations to it. BFD over SVI across Virtual Port-Channel (VPC) is not supported on 6.2(12) release. BFD over SVI over FabricPath is supported from 7.2(0)D1(1) release. However, BFD over SVI should be supported on regular Link Aggregation Control Protocol (LACP) port-channel interface.

## Problem: BFD doesn't come up for BGP over an SVI on Nexus 7000

BFD remains in admin down state but BGP session comes up fine when the BFD session is over an SVI interface across a regular port-channel. This is the configuration of the SVI interface:

```
interface Vlan1012
description Connected-to-N7k-2
no shutdown
mtu 9202
mac-address 0022.0022.0022
bfd interval 50 min_rx 50 multiplier 3
bfd echo-rx-interval 50
bfd ipv4 interval 50 min_rx 50 multiplier 3
bfd ipv6 interval 50 min_rx 50 multiplier 3
bfd ipv4 echo-rx-interval 50
bfd ipv6 echo-rx-interval 50
vrf member ROUTING-TRANSIT
ip flow monitor Monitor-x input sampler Sampler-x
ipv6 flow monitor Monitor-x-IPv6 input sampler Sampler-x
no ip redirects
ip address 10.1.12.0/31
```

When the BGP is configured to use BFD, the BFD session remains in **AdminDown** state but the BGP session comes up. The Tx count for the BFD session increments or is 0 but the Rx count

remains 0 all the time.

```
N7k-1#show bfd nei vrf all details
```

OurAddr	NeighAddr	LD/RD	RH/RS	Holdown(mult)	State
Int	Vrf				
10.0.12.0	10.0.12.1	1107296261/0	Down	N/A (3)	Down
Vlan1012	ROUTING-TRANSIT				

Session state is Down and not using echo function

Local Diag: 0, Demand mode: 0, Poll bit: 0, Authentication: None

MinTxInt: 2000000 us, MinRxInt: 2000000 us, Multiplier: 3

Received MinRxInt: 0 us, Received Multiplier: 3

Holdown (hits): 0 ms (0), Hello (hits): 2000 ms (0)

Rx Count: 0, Rx Interval (ms) min/max/avg: 0/0/1 last: 50999 ms ago

Tx Count: 0, Tx Interval (ms) min/max/avg: 0/0/0 last: 0 ms ago

Registered protocols: bgp

Downtime: 0 days 0 hrs 1 mins 3 secs

Last packet: Version: 0 - Diagnostic: 0

**State bit: AdminDown** - Demand bit: 0 Poll bit: 0 - Final bit: 0 Multiplier: 3 -

Length: 24 My Discr.: 0 - Your Discr.: 0 Min tx interval: 0 - Min rx interval: 0 Min Echo interval: 0 - Authentication bit: 0 Hosting LC: 4, Down reason: No Diagnostic, Reason not-

hosted: None N7k-1#show ip bgp vrf all summary

BGP summary information for VRF ROUTING-TRANSIT, address family IPv4 Unicast

BGP router identifier 10.1.12.0, local AS number 65535

BGP table version is 13, IPv4 Unicast config peers 1, capable peers 1

0 network entries and 0 paths using 0 bytes of memory

BGP attribute entries [0/0], BGP AS path entries [0/0]

BGP community entries [0/0], BGP clusterlist entries [0/0]

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.1.12.1	4	65535	5406	5407	13	0	0	00:01:10	0

The problem persists even if the VDC is deleted and recreated and the whole configuration is made from the start.

## Analysis

From the logs it is noticed that there is a TCAM Error message that seems to cause this behavior. When Netflow is configured on the SVI, this is the error message seen on the Nexus.

```
N7k %$ VDC-1 %$ %ACLQOS-SLOT4-2-ACLQOS_FAILED: ACLQOS failure: feature combination not supported on VDC-2 VLAN 1012 for : Netflow Sampler (SVI), Netflow Sampler (SVI), BFD
```

This error message is seen when there are combinations of features used. When feature combinations that cannot coexist are configured on the same Ternary Content Addressable Memory (TCAM) banks, TCAM failures are seen and some of those features do not function properly.

## Solution

If Access Control List (ACL) TCAM bank mapping is enabled, it can help to overcome this issue. The ACL TCAM bank mapping allows TCAM banks to accommodate more feature combinations in a more predictable manner. Features are pre-classified into feature groups, which are further predefined into feature classes according to which features are allowed to coexist in a TCAM bank. ACL TCAM bank mapping allows you to configure a set of features at the same time and reduces multiple results that can accumulate when feature combinations that cannot coexist are configured on the same TCAM banks. The TCAM bank mapping can be configured using the

command **hardware access-list resource feature bank-mapping in default VDC**. To view how the features are mapped to TCAM banks, use the command **show system internal access-list feature bank-chain map vlan-vlan ingress module *slot-number***.

After TCAM bank mapping is configured, the BFD session comes out of AdminDown state and the %ACLQOS error log is not noticed again.