

# vPC Layer 3 Backup Routing with F1 and Peer-Gateway



Document ID: 116740

Contributed by Andy Gossett, Cisco TAC Engineer.  
Dec 16, 2013

## Contents

### Introduction

#### Prerequisites

- Requirements
- Components Used

#### Configure

- Network Diagram
- Peer-Gateway Overview
- vPC L3 Backup Routing with F1 and Peer-Gateway
- Peer-Gateway Exclude VLAN

#### Verify

#### Troubleshoot

## Introduction

This document describes Layer 3 (L3) backup routing in a virtual port channel (vPC) setup. Cisco recommends that you use the *peer-gateway exclude-vlan* command when you use F1 modules on the peer-link.

**Note:** If the vPC peer link is configured on a Cisco Nexus 32-port 1/10 Gigabit Ethernet (F1-Series) module (N7K-F132XP-15), you must include the L3 backup routing VLAN in the VLAN list specified by the *peer-gateway exclude-vlan* command.

See Cisco Nexus 7000 Series NX-OS Release Notes, Release 5.1: New Software Features: Layer 3 Backup Routing VLAN for details on the new *peer-gateway exclude-vlan* command.

## Prerequisites

### Requirements

There are no specific requirements for this document.

### Components Used

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

- Cisco Nexus 7000 Series Switch, Release 5.1(3) and later
- Mixed chassis with M1 and F1 line cards

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.

## Configure

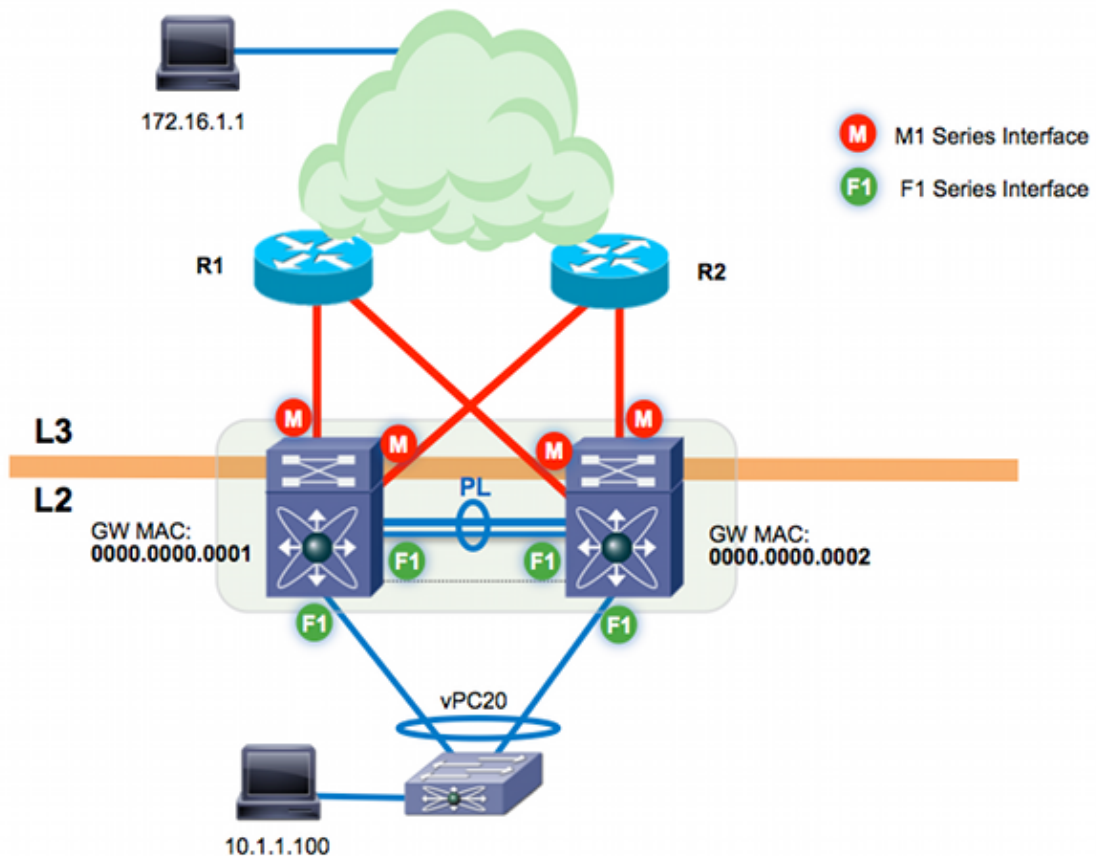
### Notes:

Use the Command Lookup Tool (registered customers only) in order to obtain more information on the commands used in this section.

The Output Interpreter Tool (registered customers only) supports certain *show* commands. Use the Output Interpreter Tool in order to view an analysis of *show* command output.

## Network Diagram

The topology used in this document is:

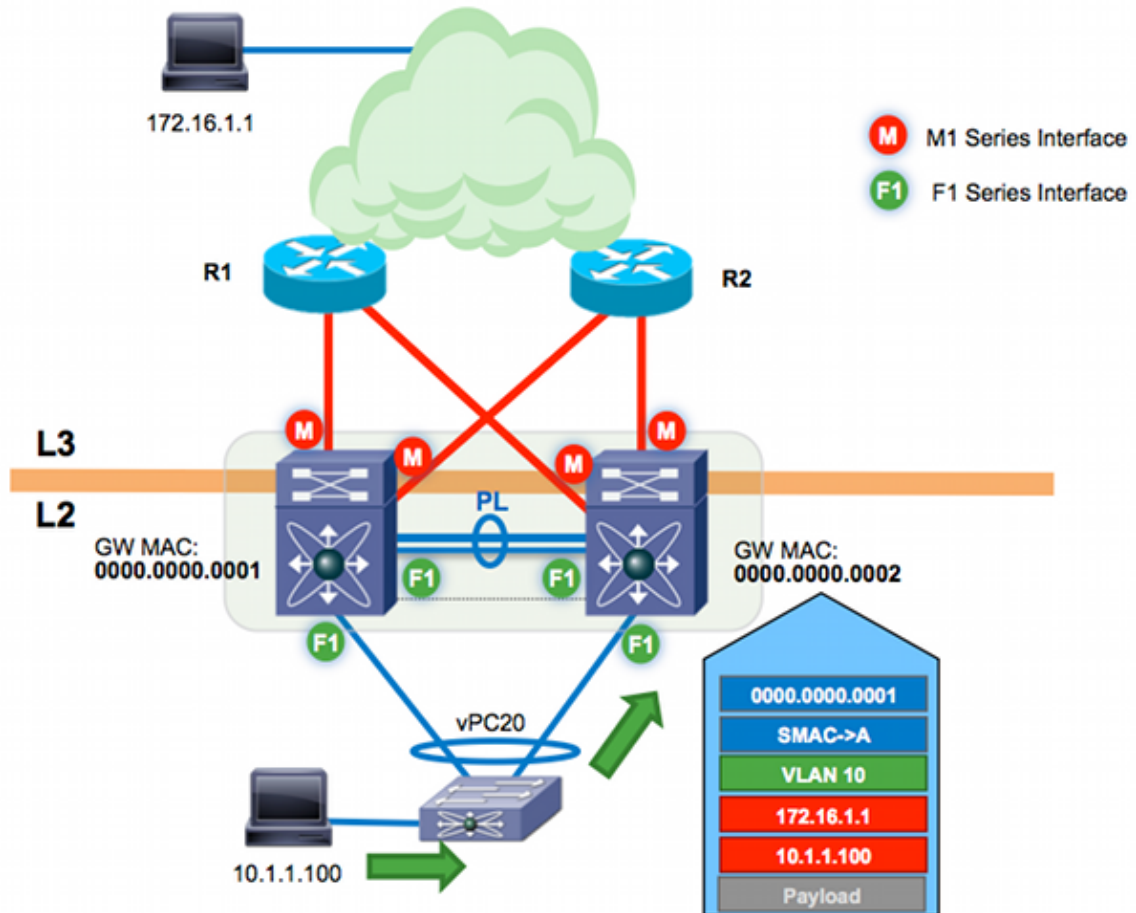


The vPC peer-link is built on F1 modules. M1 modules are allocated to the VDC for proxy-routing functionality; the M1 modules terminate the L3 uplinks into the core layer. There are two Cisco Nexus 7000 switches:

- n7k-agg1 (MAC 0000.0000.00001)
- n7k-agg2 (MAC 0000.0000.00002)

## Peer-Gateway Overview

Peer-gateway is a vPC feature that allows vPC peer devices to act as a gateway for traffic destined to the MAC address of their peers. In this example, a host in VLAN 10 (10.1.1.100) sends a frame northbound to the host 172.16.1.1. The gateway for the host in VLAN10 is n7k-agg1 (MAC 0000.0000.00001).



The destination MAC address for the frame is toward the n7k-agg1 MAC (0000.0000.0001). The Layer 2 (L2) switch connects to the Cisco Nexus 7000 switches through a vPC. As a result, this frame can hash toward n7k-agg1 or n7k-agg2. In this example, the port-channel load balancing algorithm hashes the frame on the link connected to n7k-agg2.

n7k-agg1 is configured in the same vPC domain as n7k-agg2, and peer-gateway is enabled. As a result, n7k-agg2 programs the MAC address for n7k-agg1 with the Gateway (G) flag in the MAC table for all switch virtual interfaces (SVIs) allowed across the peer-link – and vice versa.

```
n7k-agg2# show mac address-table vlan 10 address 0000.0000.0001
```

Legend:

```
* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
age - seconds since last seen,+ - primary entry using vPC Peer-Link,
(T) - True, (F) - False
```

VLAN	MAC Address	Type	age	Secure	NTFY	Ports/SWID.SSID.LID
G 10	0000.0000.0001	static	-	F	F	sup-eth1(R)

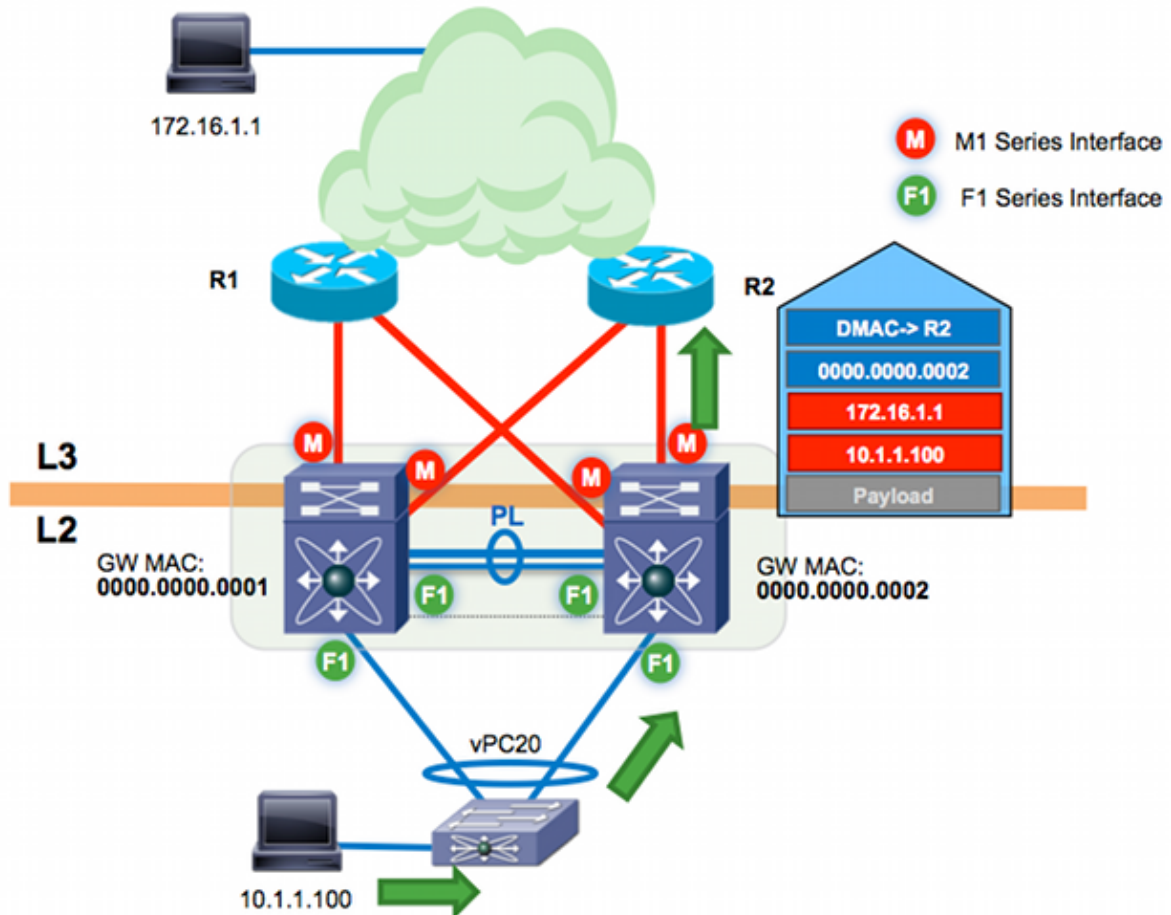
Since the Gateway flag is set for MAC 0000.0000.0001, n7k-agg2 performs a L3 lookup and routes this frame on behalf of n7k-agg1.

```

n7k-agg2# show ip route 172.16.1.1
IP Route Table for VRF "default"
'*' denotes best ucast next-hop
'**' denotes best mcast next-hop
'[x/y]' denotes [preference/metric]
'%<string>' in via output denotes VRF <string>

172.16.1.0/30, ubest/mbest: 2/0
  *via 10.0.0.5, Eth3/4, [110/8], 00:20:40, ospf-1, intra
  *via 10.0.0.13, Eth3/3, [110/8], 00:20:40, ospf-1, intra

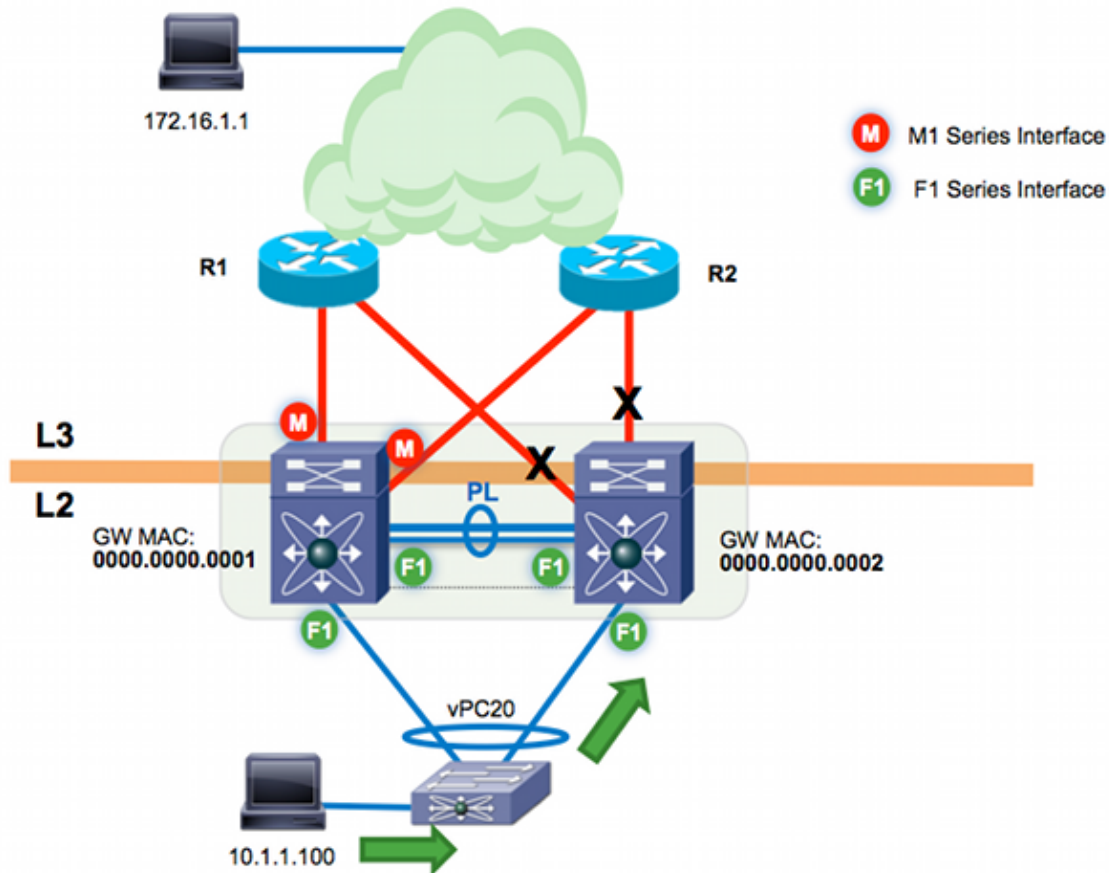
```



See Cisco Nexus 7000 Series NX-OS Interfaces Configuration Guide, Release 6.x: Configuring vPCs: vPC Peer-Gateway for more details.

## vPC L3 Backup Routing with F1 and Peer-Gateway

vPC L3 backup routing refers to traffic routed between vPC peers over the peer-link. Assume the two L3 uplinks on n7k-agg2 (from the previous example) are now down. If there is a routing protocol such as Open Shortest Path First (OSPF) or Enhanced Interior Gateway Routing Protocol (EIGRP) that is running between the two Cisco Nexus 7000 switches on one of the vPC VLANs, n7k-agg2 has an alternate route across the peer-link.



```
n7k-agg2# show ip route 172.16.1.1
(some output omitted)
172.16.1.0/30, ubest/mbest: 1/0
  *via 10.99.1.1, Vlan99, [110/48], 00:00:04, ospf-1, intra
```

```
n7k-agg2# show ip arp 10.99.1.1
IP ARP Table
Total number of entries: 1
Address      Age      MAC Address      Interface
10.99.1.1    00:13:02  0000.0000.0001  Vlan99
```

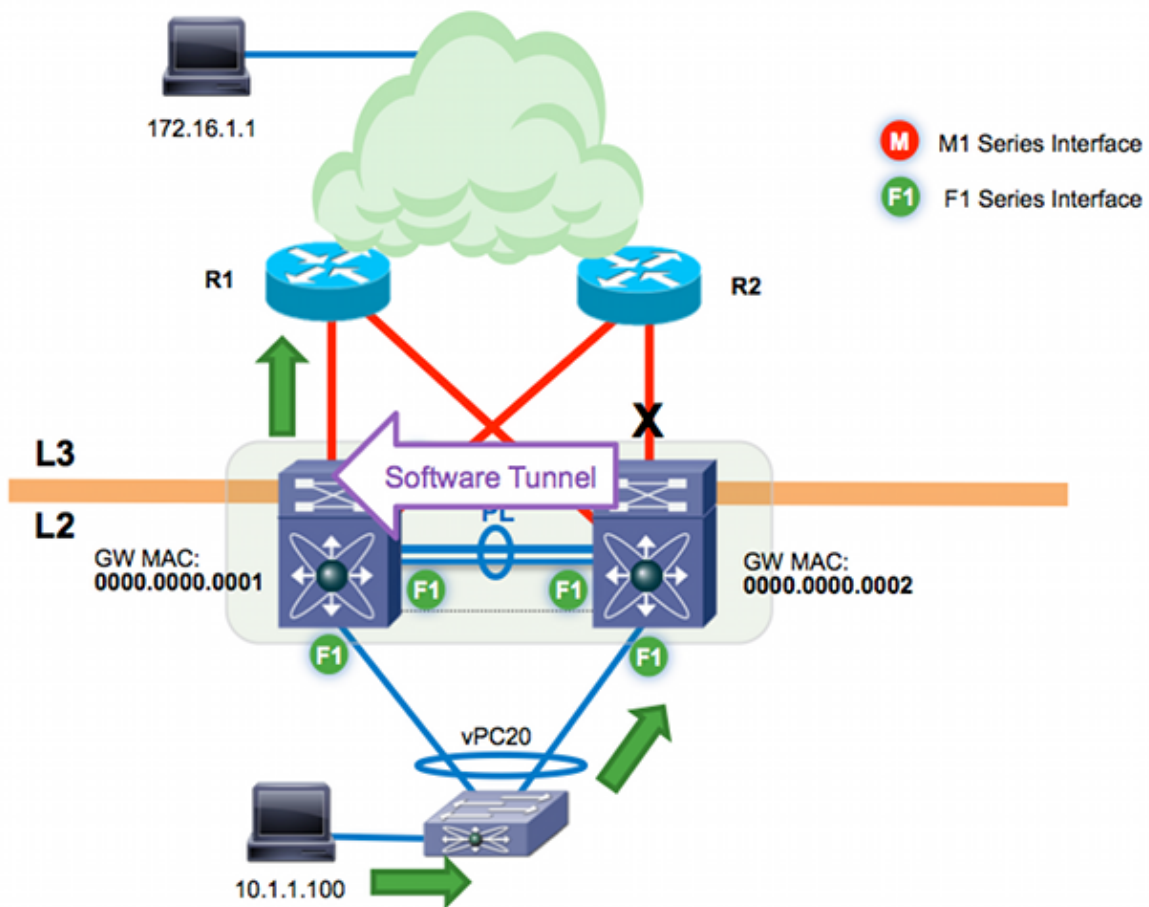
```
n7k-agg2# show mac address-table vlan 99 address 0000.0000.0001
```

Legend:

\* - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC  
age - seconds since last seen, + - primary entry using vPC Peer-Link,  
(T) - True, (F) - False

VLAN	MAC Address	Type	age	Secure	NTFY	Ports/SWID.SSID.LID
G 99	0000.0000.0001	static	-	F	F	sup-eth1(R)

The L3 next hop to the 172.16.1.1 destination is n7k-agg1 MAC 0000.0000.0001 on VLAN 99. Any VLAN allowed on the vPC peer-link is, by definition, a vPC-enabled VLAN. VLAN 99 is a vPC-enabled VLAN. Because peer-gateway is enabled, VLAN 99 is programmed with the Gateway flag. This traffic flow is tunneled in software between the two Cisco Nexus 7000 switches when F1 modules are used for the peer-link.



Use ethanalyzer in order to see this flow on the inband. Because ethanalyzer captures only traffic sent to the CPU for software processing, you do not see traffic that is successfully forwarded in hardware.

```
n7k-agg2# ethanalyzer local interface inband capture-filter "host 10.1.1.100
and host 172.16.1.1"
Capturing on inband
2013-10-29 17:30:00.638106 10.1.1.100 -> 172.16.1.1 ICMP Echo (ping) request
2013-10-29 17:30:00.647949 10.1.1.100 -> 172.16.1.1 ICMP Echo (ping) request
2013-10-29 17:30:00.657941 10.1.1.100 -> 172.16.1.1 ICMP Echo (ping) request
2013-10-29 17:30:00.667943 10.1.1.100 -> 172.16.1.1 ICMP Echo (ping) request
2013-10-29 17:30:00.678179 10.1.1.100 -> 172.16.1.1 ICMP Echo (ping) request
2013-10-29 17:30:00.687948 10.1.1.100 -> 172.16.1.1 ICMP Echo (ping) request
2013-10-29 17:30:00.697948 10.1.1.100 -> 172.16.1.1 ICMP Echo (ping) request
2013-10-29 17:30:00.707944 10.1.1.100 -> 172.16.1.1 ICMP Echo (ping) request
2013-10-29 17:30:00.717947 10.1.1.100 -> 172.16.1.1 ICMP Echo (ping) request
2013-10-29 17:30:00.728246 10.1.1.100 -> 172.16.1.1 ICMP Echo (ping) request
10 packets captured
```

Traffic switched in software can experience delay and extreme packet loss due to control plane policing (CoPP) and hardware rate-limiters. Overall performance is slower for software forwarding than hardware forwarding.

In summary, because of the hardware implementation of proxy-forwarding on F1, traffic that meets these requirements will be tunneled in the software:

1. The L3 next hop for a vPC device is its vPC peer on a vPC-enabled VLAN.
2. The Gateway flag is set for the next hop MAC address.
3. F1 interfaces are used on the peer-link.

## Peer-Gateway Exclude VLAN

Use the *peer-gateway exclude-vlan vlan-number* command in order to allow L3 backup routing to be performed in hardware with F1 modules on the peer-link. In this example, the two Cisco Nexus 7000 switches are running an OSPF Interior Gateway Protocol (IGP) on VLAN 99. Therefore, you should exclude peer-gateway only on VLAN 99 in order to allow this traffic flow to be forwarded in hardware.

```
n7k-agg2(config)# vpc domain 102
n7k-agg2(config-vpc-domain)# peer-gateway exclude-vlan 99
Warning:
!! Previous peer-gateway config has been overwritten!!
```

You can verify the exclusion with the *show vpc* command:

```
n7k-agg2# show vpc | grep -i gateway
Peer Gateway                : Enabled
Peer gateway excluded VLANs : 99
```

n7k-agg2 no longer has the Gateway flag set for the n7k-agg1 MAC (0000.0000.0001) on VLAN 99:

```
n7k-agg2# show ip route 172.16.1.1
(some output omitted)
172.16.1.0/30, ubest/mbest: 1/0
  *via 10.99.1.1, Vlan99, [110/48], 00:00:04, ospf-1, intra
```

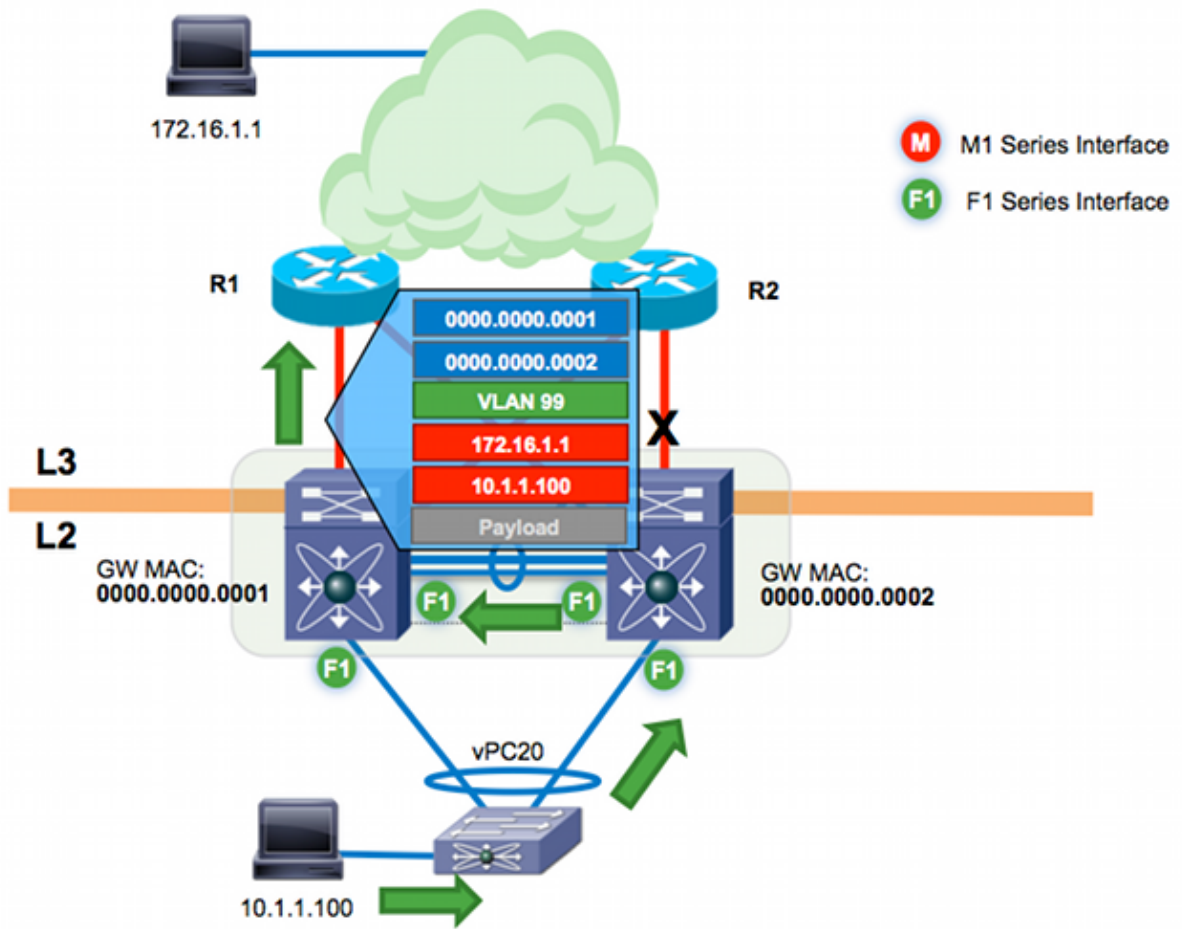
```
n7k-agg2# show ip arp 10.99.1.1
IP ARP Table
Total number of entries: 1
Address      Age      MAC Address      Interface
10.99.1.1    00:13:02 0000.0000.0001 Vlan99
```

```
n7k-agg2# show mac address-table vlan 99 address 0000.0000.0001
```

Legend:

```
  * - primary entry, G - Gateway MAC, (R) - Routed MAC, O - Overlay MAC
  age - seconds since last seen,+ - primary entry using vPC Peer-Link,
  (T) - True, (F) - False
  VLAN      MAC Address      Type      age      Secure NTFY Ports/SWID.SSID.LID
-----+-----+-----+-----+-----+-----+-----+-----
* 99        0000.0000.0001      static      -          F      F vPC Peer-Link
```

As a result, traffic routed into VLAN 99 with a destination MAC of 0000.0000.0001 can be L2-switched across the peer-link in hardware.



## Verify

Verification procedures are included within the configuration steps.

## Troubleshoot

There is currently no specific troubleshooting information available for this configuration.