

# OpenFlow Single-table Implementation for Dell Networking N-Series SDN

A Dell Deployment and Configuration Guide

Dell Networking Solutions Engineering  
October 2016

## Revisions

Date	Description	Authors
October 2016	Initial release	Victor Teeter, Maunish Shah

THIS WHITE PAPER IS FOR INFORMATIONAL PURPOSES ONLY, AND MAY CONTAIN TYPOGRAPHICAL ERRORS AND TECHNICAL INACCURACIES. THE CONTENT IS PROVIDED AS IS, WITHOUT EXPRESS OR IMPLIED WARRANTIES OF ANY KIND.

Copyright © 2016 Dell Inc. or its subsidiaries. All rights reserved. Dell and the Dell EMC logo are trademarks of Dell Inc. in the United States and/or other jurisdictions. All other marks and names mentioned herein may be trademarks of their respective companies.

# Table of contents

1	Introduction .....	5
2	OpenFlow Network Requirements .....	8
2.1	Management Layer.....	9
2.2	Designated ports.....	10
2.2.1	Management ports .....	10
2.2.2	Primary and secondary OpenFlow control ports .....	10
3	OpenFlow software installations.....	11
3.1	NEC PF6800 controller installation and configuration.....	11
3.2	DNOS-OF 1.1 firmware installation .....	12
3.3	Configuring N-Series for OpenFlow.....	13
4	Discovering OpenFlow Switches.....	16
5	Configuring end-to-end flows .....	32
5.1	Creating an L2 flow.....	32
5.1.1	VTN L2 map creation.....	35
5.1.2	vBridge configuration .....	36
5.1.3	vExternal configuration .....	37
5.1.4	Validation .....	43
5.2	Creating an L3 flow.....	48
5.2.1	VTN L3 map creation .....	51
5.2.2	vBridge configuration .....	53
5.2.3	vExternal configuration .....	54
5.2.4	vRouter configuration .....	58
5.2.5	Validation .....	63
6	Failover (Failure Recovery) .....	68
A	Additional Resources .....	73
B	Version information .....	74
C	Glossary of Terms .....	75
D	Support and Feedback .....	77
	About Dell EMC .....	77

## Executive summary

OpenFlow is an innovative technology that separates the control and data planes on a network. This separation allows monitoring and management of all network traffic from a centralized workstation. This in turn allows network administrators to create and apply customized policies to mold traffic paths and patterns to best fit their environments and preferences.

OpenFlow is fast becoming a requirement for certain segments of campus networking products, particularly in education markets. DNOS-OF 1.1 is a campus networking firmware release that runs on existing Dell Networking N-Series hardware. This provides OpenFlow 1.3.4 support as a pure OpenFlow switch, as well as interoperability with OpenFlow cluster controllers.

This document gives network designers an understanding of the basic principles of OpenFlow and software defined networks (SDN). Simplification of design allows for operational efficiencies and enables network architects to deploy an easily managed campus network that addresses the needs of diverse communities.

## Introduction

DNOS-OF 1.1 is a downloadable firmware image available for Dell Networking N-Series hardware that enables OpenFlow 1.3.4 support as a *pure* OpenFlow (OF) switch. This firmware:

- provides basic OpenFlow-mode support on N-Series switches which enables SDN for campus networks
- may co-exist in flash with existing N-Series firmware without affecting image management functionality. This provides the ability to load DNOS-OF 1.1 code from within the network administrator's existing firmware, convert to a pure OF switch, then revert back to former configurations with no impact to existing firmware.
- provides limited, simpler features and functionality to allow for delivery of a quick, low-cost solution that is easy to test in customer lab environments.

It is important to note that DNOS 6.3 for N-Series switches introduced *hybrid* OpenFlow features. See the User Guide for the N-Series switch for more information about the hybrid implementation:

[http://www.dell.com/support/home/us/en/19/Products/ser\\_stor\\_net/networking](http://www.dell.com/support/home/us/en/19/Products/ser_stor_net/networking)

### Data Plane and Control Plane

The control plane has been extracted from the switches and given to the OpenFlow controller, which now manages all control plane responsibilities. The controller keeps track of which links are up or down on the network (across all switches) and manipulates the path of each flow accordingly. This in turn leaves only the data plane responsibilities on the switches to move packets across the OpenFlow network.

OpenFlow creates flows and paths for traffic to pass between OF switches across the network much as an access control list (ACL) creates them in traditional switches. The difference is, when a packet matches an ACL entry, the ACL permits or denies the packet. When a packet matches a defined packet flow entry, the flow forwards the packet to the assigned port. If a packet does not match a defined entry, the flow sends the packet to the OF controller for analysis and further instruction.

Due to this method of separating the data and control planes, most protocols used in traditional networks do not apply to OpenFlow. As an example, administrators may plug in cables between DNOS-OF switches without worrying about network loops.

### 1.1

## Supported switches

DNOS-OF 1.1 provides support for the following families of Dell Networking switches:

- Dell Networking N1500 (N1524, N1524P, N1548, N1548P)
- Dell Networking N2000 (N2024, N2024P, N2048, N2048P)
- Dell Networking N3000 (N3024, N3024P, N3048, N3048P)
- Dell Networking N4000 (N4032, N4064, N4024F, N4064F)

**Note:** DNOS-OF 1.1 does not support the Dell Networking N3024F switch.

## 1.2

## Purpose of this document

Table 1 provides a more detailed description of the purpose of this guide:

Table 1 Document purpose (is / is not)

This document is	This document is not
a supplement to the DNOS-OF User Guide.	a complete OpenFlow reference guide.
for single-table implementation of DNOS-OF only.	for multiple-table implementations of DNOS-OF.
a guide showing examples using NEC controllers.	an NEC OpenFlow controller installation or user guide.
for Dell Networking N-Series running DNOS-OF.	a guide for FTOS OpenFlow switches.
provided to help new users of OpenFlow get a better understanding of the technology.	a guide showing examples using Ryu or other controllers. (The User Guide contains other examples.)

## 1.3

## DNOS-OF firmware download

Locate DNOS-OF v1.1 firmware for download from dell.com by clicking the following link:

<http://www.dell.com/support/home/us/en/19/product-support/product/dell-networking-os-openFlow/drivers>

**Note:** The link also contains a **pfc\_nec.zip** file which will be used later in section 3.1 to enable NEC PFC OF controller support of Dell N-Series switches.

## 1.4

## OpenFlow layered topologies

There are several layers of networking that make up a complete OpenFlow topology. To better explain this, Figure 1 shows the OpenFlow solution broken out into the three basic layers. The bottom layer shows an OpenFlow data network consisting of several switches interconnected to allow for end-to-end flows between any ports. This network is labeled *OpenFlow data network*. These Dell N-Series switches run DNOS-OF 1.1 and contain the *data plane* only. Data on this network normally travels switch-to-switch using the paths in red, but may also use secure channel links to access controllers when no packet flow entry is matched.

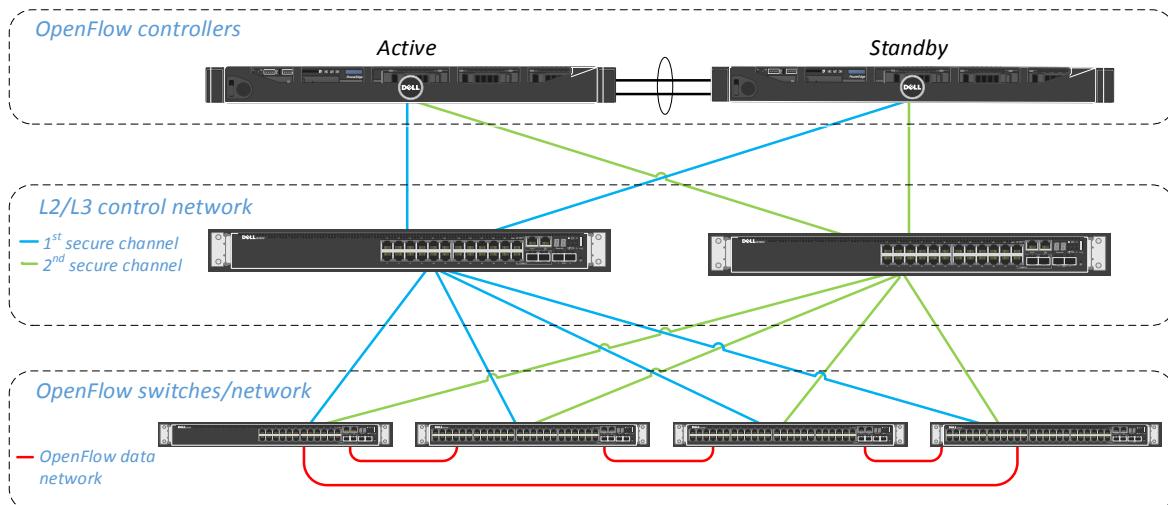


Figure 1 OpenFlow data and control networks

The middle layer in Figure 1 is the L2/L3 control network. This network uses traditional switches and is the liaison between the actual OpenFlow data network (bottom layer) which comprises the data plane, and the OpenFlow controllers (top layer) which comprises the control plane. These switches can be any switch or manufacturer and require no specific configuration for OpenFlow control traffic. Dell recommends using Dell switches running DNOS 6.x, DNOS 9.x or DNOS 10.x.

**Note:** Switches used in the L2/L3 control network are not OpenFlow switches. Traditional switches are required at this layer of the topology.

The top layer shows Dell servers with OpenFlow controller software installed. Network administrators use this software to construct custom-tailored data traffic across their networks focused specifically around their own requirements or preferences. These controllers comprise the *control plane* of the network.

## 2

## OpenFlow Network Requirements

Figure 2 shows the OpenFlow network example that this guide demonstrates using NEC Programmable Flow Control (PFC) controllers. The figure provides example IP addresses and ports to show where the cables connect.

This model uses two Dell R630 Servers to install NEC PF6800 controller software. Obtain the licensed software through NEC's website: [www.necam.com/sdn](http://www.necam.com/sdn).

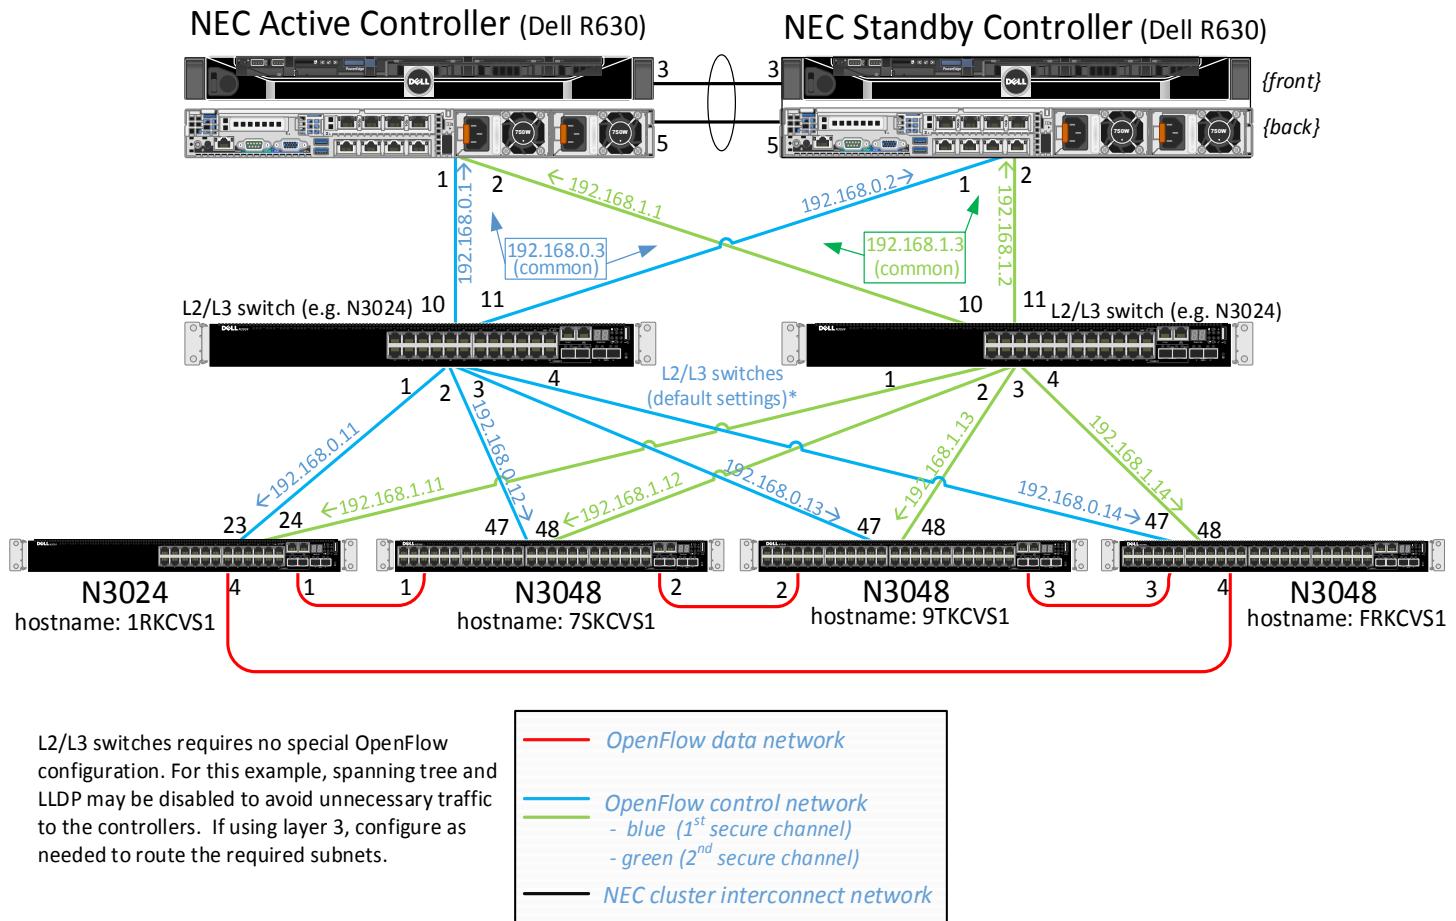


Figure 2 OpenFlow network example

Front and back views of the two Dell R630 servers are displayed as labeled in the figure above.

This example uses Dell Networking N-3000 switches as OpenFlow switches, however any Dell N-Series switch listed under Supported switches [on page 5](#) also works.

**Note:** Hostnames of DNOS-OF switches are predefined and may not be changed.

## 2.1 Management Layer

No network design is complete without a layer for management traffic. This holds especially true for OpenFlow because it allows administrators to create their networks using management stations complete with an OpenFlow controller web application. It also allows administrators to configure the OpenFlow switches to be operated by the NEC controllers. Figure 3 uses dotted lines to show all management network cables coming into a single management switch.

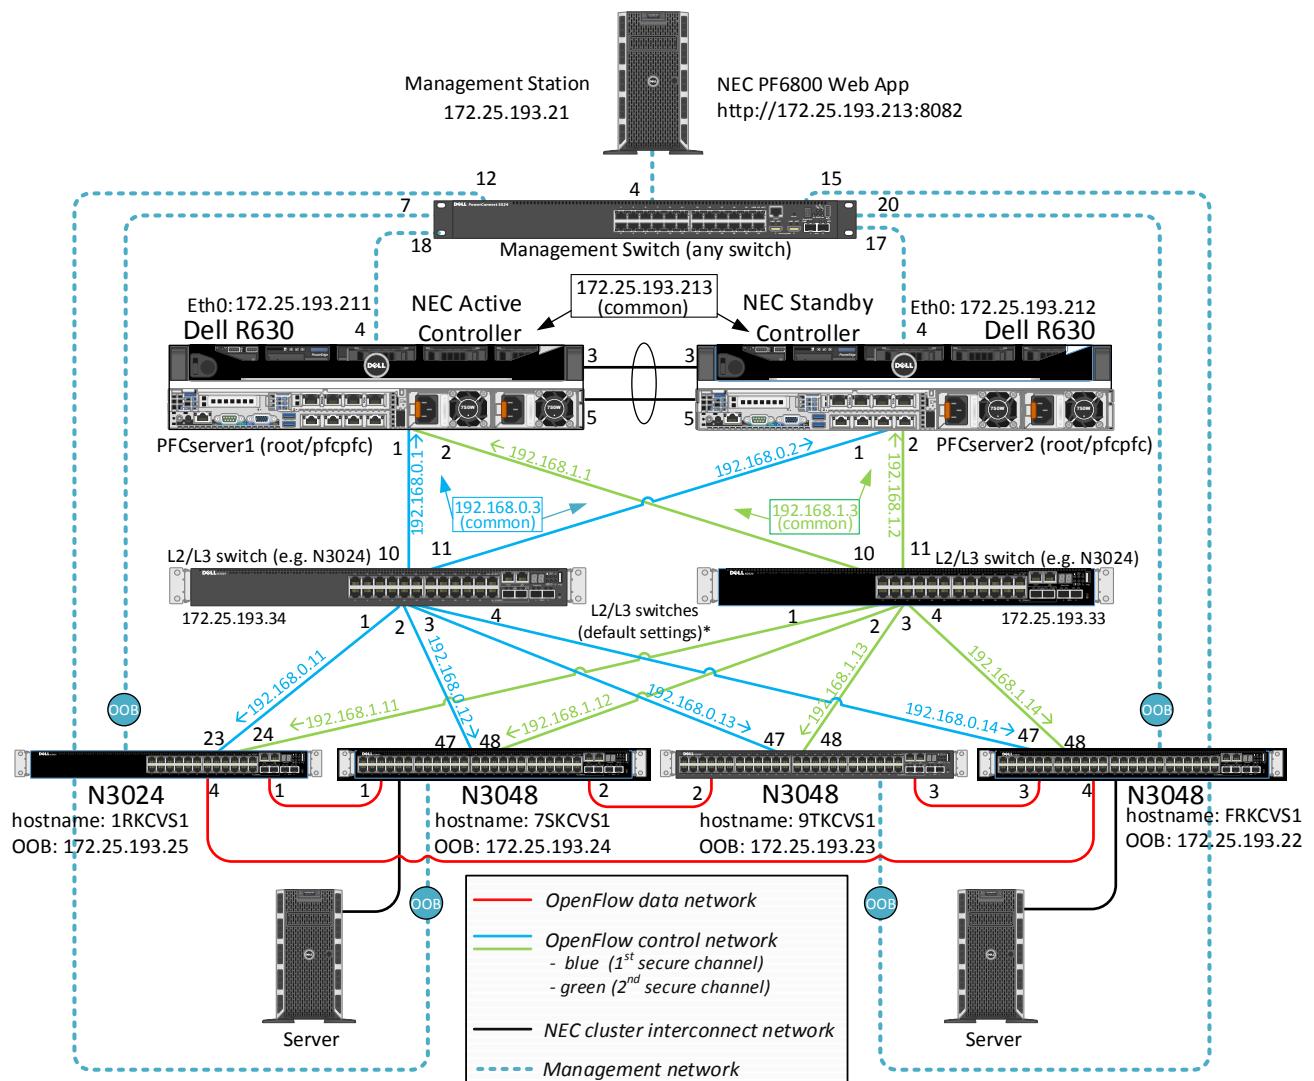


Figure 3 Addition of the management network

The out-of-band (OOB) ports on the N-Series OpenFlow switches are used to allow NEC controllers access. Figure 3 also shows IP addresses for both controllers and the web application client. Each red, green, black,

blue, and dotted blue line represents a cable used to connect two devices on either end. The one and two-digit number beside each cable entering a device (switch or computer) indicates the port used for this particular example. Some ports on the DNOS-OF OpenFlow N-Series switches are specifically designed for a particular use, such as OOB and OF control (Secure Channel) ports. Table 2 below shows the designated ports that must be used for controlling and managing DNOS-OF OpenFlow switches. The example network may be cabled using the information in Table 2 and Figure 3. Additional cables installed between OpenFlow switches will also be utilized by the controller as appropriate, providing both additional bandwidth and failover redundancy.

The IP address labeled “common” between the NEC active and standby controllers is a shared IP address used to connect to both controllers. Using this address allows a seamless network failover recovery to the Standby controller in cases where the Active controller fails.

## 2.2

## Designated ports

DNOS-OF designates ports on the Dell N-Series OpenFlow switches that must only be used for specific purposes. Table 2 below lists each supported switch along with designated ports and their purposes. Use of the ports for purposes other than those defined results in a failed configuration.

Table 2 Designated ports on DNOS-OF switches

Dell Switch	Management Port	Primary OF Control Port	Backup OF Control Port
N1524, N1524P	Port 1	Port 23	Port 24
N1548, N1548P	Port 1	Port 47	Port 48
N2024, N2024P	Port 1	Port 23	Port 24
N2048, N2048P	Port 1	Port 47	Port 48
N3024, N3024P	OOB	Port 23	Port 24
N3048, N3048P	OOB	Port 47	Port 48
N4032, N4032F	OOB	Port 31	Port 32
N4064, N4064F	OOB	Port 63	Port 64

### 2.2.1

### Management ports

Since the Dell N1500 and N2000 series switches do not have OOB ports for management, DNOS-OF designates port 1 as the management port. Port 1 on these two series may not be used for any other purpose. Port 1 on the other N-series switches (i.e. N3000, N4000) may be used as normal. Management ports, whether OOB or port 1, attach to a management switch (see Figure 3) for purposes such as SSH, Telnet, TFTP, and SNMP. This separate network only carries network management traffic.

### 2.2.2

### Primary and secondary OpenFlow control ports

OpenFlow control ports connect to OpenFlow control switches to provide communication between the OpenFlow controller and the DNOS-OF switches. For each switch, the highest numbered port connects to the secondary OF controller. The highest-numbered port minus-one connects to the primary OF controller. For example, the Dell Networking N3048 switch uses ports 47 and 48. These two ports are unavailable for any other purposes.

## OpenFlow software installations

This section guides the reader through installation and configuration of all required software for configuring and enabling the example OpenFlow network shown in Figure 2. This section explains installation of DNOS-OF firmware for converting N-Series switches to pure OpenFlow switches and configuring the NEC controllers to prepare them for action.

### 3.1

## NEC PF6800 controller installation

Obtain the NEC PF6800 software and license through NEC's website ([www.necam.com](http://www.necam.com)). Use the instructions provided by NEC to install the software onto the two Dell R630 servers, as shown in Figure 3 (also shown in top layer of Figure 1). Download and apply the NEC PFC controller patch from Dell after completing installation. The patch enables the NEC controller to support N-series Dell switches. Download the patch (*nec\_pfc.zip*) and follow the instructions provided in the zipped file:

<http://www.dell.com/support/home/us/en/19/product-support/product/dell-networking-os-openflow/drivers>

**Note:** The **pfc\_nec.zip** file includes a file called “Instructions for NEC PFC OF Controller patch installation.” Use these instructions to modify the NEC PFC OpenFlow controller to support Dell N-Series switches.

### 3.1.1

## Configuring the PF6800 controllers

Refer to NEC's *PF6800 Ver. 6.0 Installation Guide* for PF6800 and cluster configuration instructions. The instructions provide complete guidance and direction for PFC setup and license installation, as well as NEC controller cluster configuration and initiation. They also set the cluster to use OpenFlow version 1.3.

As shown in Figure 3, we used the IP address of 172.25.193.213 when configuring the controller. SSH to 172.25.193.213:8082 to access the Active PF6800. Use the commands below to verify the configuration.

### Validating the controller setup

Use the show commands below to validate a successful configuration. The output should show the presence of both active and standby controllers as highlighted. IP addresses shown for each server (pfcserver1 and pfcserver2) are user-defined while configuring the PFC controller. For this example we used 1.1.1.1 and 1.1.1.2.

```
[root@pfcserver1 ~]# pfc_show_cluster_status
Cluster condition is eligible to execute pfc_switch_cluster.

      Date Time          node        IP address   status
2016-08-22 13:04:41  pfcserver1  1.1.1.1      ACT
2016-08-22 13:04:40  pfcserver2  1.1.1.2      SBY

[root@pfcserver1 ~]# pfc_show_cluster_config
[OWN CONFIG]
      Field          Committed data   Registered data
      server name    pfcserver1       -
```

interconnect IP address	1.1.1.1	-
monNP interconnect IP address	-	-
BMC management IP address	-	-
BMC user name	-	-
[OPPOSITE CONFIG]		
Field	Committed data	Registered data
server name	pfcserver2	-
interconnect IP address	1.1.1.2	-
monNP interconnect IP address	-	-
BMC management IP address	-	-
BMC user name	-	-
[COMMON CONFIG]		
Field	Committed data	Registered data
num of redundancy nodes	2	-
floating IP address 1	192.168.0.3	-
floating IP address 2	192.168.1.3	-
floating IP address 3	172.25.193.213	-
monitoring IP address	-	-
time-out value of monIP	30	-
monitoring NIC	-	-
monitoring NP	-	-
heartbeat sensitivity	low	-
watchdog timer	softdog	-
forcibly stopping	disable	-

```
[root@pfcserver1 ~]#
```

## 3.2 DNOS-OF 1.1 firmware installation

Download DNOS-OF firmware and apply it to all N-Series switches participating in the OpenFlow network. Refer to the bottom layer of Figure 1 and the bottom four switches in Figure 3. Locate the current version of DNOS-OF firmware, version 1.1, at the following web address on dell.com:

<http://www.dell.com/support/home/us/en/19/product-support/product/dell-networking-os-openflow/drivers>

The single zip file *dnos-of-1.1.0.x.zip* contains the following three software Images:

- DNOS-OF-N3000-N2000-1.1.x.x.stk (Dell N3000 and N2000 switches)
- DNOS-OF-N4000-1.1.x.x.stk (Dell N4000 switches)
- DNOS-OF-N1500-1.1.x.x.stk (Dell N1500 switches)

See the Dell N-Series User's Manual for OpenFlow firmware installation instructions. The instructions should resemble traditional DNOS firmware installation instructions for these switches, as the process is the same.

Obtain the User's Manual for N-Series hardware from dell.com at the following web address:

<http://www.dell.com/support/home/us/en/19/product-support/product/dell-networking-os-openflow/manuals>

### 3.3

## Configuring N-Series for DNOS OpenFlow

Enter the commands below from the command line to configure each DNOS-OF switch participating in the network. Notice that the hostnames and IP address assignments of each switch can be matched with those represented in Figure 3. DPIDs are user-defined on the command line. In this example the last octet of the control port IP addresses (11-14) are used for each DPID. Immediately after setting a DPID, it is confirmed with a message showing it was set. The DPID is stored in hexadecimal.

After each switch is configured and enabled, connectivity to the OpenFlow controller is automatically confirmed and displays the message “**OFS Connected to Controller**” as shown in each section below.

**Note:** Once configured, the **write** command may be used to save each running-configuration to the startup-configuration.

### 3.3.1 Configuring switch 1 (1RKCVS1)

```
1RKCVS1_console# set openflow mode singletable
```

OpenFlow table processing mode set to singletable

```
1RKCVS1_console# set openflow controller name Test  
Controller Test (index 0) created
```

```
1RKCVS1_console# set openflow controller primary Test 192.168.0.3 6633  
1RKCVS1_console# set openflow controller backup Test 192.168.1.3 6633  
1RKCVS1_console# set openflow primary control port 192.168.0.11 /24  
1RKCVS1_console# set openflow backup control port 192.168.1.11 /24  
1RKCVS1_console# set openflow dpid 11
```

OpenFlow DPID set to 0x0000000000000000b

```
1RKCVS1_console# set openflow controller state Test enabled  
Controller Test enabled (connected)
```

Setting interface Test UP

```
Aug 18 21:17:24.008 ofconnectionmanager: MSG: OFS Connected to Controller  
1RKCVS1_console#
```

### 3.3.2

## Configuring switch 2 (7SKCVS1)

```
7SKCVS1_console# set openflow mode singletable
```

OpenFlow table processing mode set to singletable

```
7SKCVS1_console# set openflow controller name Test  
Controller Test (index 0) created  
  
7SKCVS1_console# set openflow controller primary Test 192.168.0.3 6633  
7SKCVS1_console# set openflow controller backup Test 192.168.1.3 6633  
7SKCVS1_console# set openflow primary control port 192.168.0.12 /24  
7SKCVS1_console# set openflow backup control port 192.168.1.12 /24  
7SKCVS1_console# set openflow dpid 12
```

OpenFlow DPID set to 0x000000000000000c

```
7SKCVS1_console# set openflow controller state Test enabled  
Controller Test enabled (connected)
```

Setting interface Test UP

```
Aug 18 21:21:01.511 ofconnectionmanager: MSG: OFS Connected to Controller  
7SKCVS1_console#
```

### 3.3.3 Configuring switch 3 (9TKCVS1)

```
9TKCVS1_console# set openflow mode singletable
```

```
OpenFlow table processing mode set to singletable  
9TKCVS1_console# set openflow controller name Test  
Controller Test (index 0) created
```

```
9TKCVS1_console# set openflow controller primary Test 192.168.0.3 6633  
9TKCVS1_console# set openflow controller backup Test 192.168.1.3 6633  
9TKCVS1_console# set openflow primary control port 192.168.0.13 /24  
9TKCVS1_console# set openflow backup control port 192.168.1.13 /24  
9TKCVS1_console# set openflow dpid 13
```

OpenFlow DPID set to 0x000000000000000d

```
9TKCVS1_console# set openflow controller state Test enabled  
Controller Test enabled (connected)
```

Setting interface Test UP

```
Aug 18 21:22:13.128 ofconnectionmanager: MSG: OFS Connected to Controller  
9TKCVS1_console#
```

### 3.3.4 Configuring switch 4 (FRKCVS1)

```
FRKCVS1_console# set openflow mode singletable
```

OpenFlow table processing mode set to singletable  
FRKCVS1\_console# **set openflow controller name Test**  
Controller Test (index 0) created

```
FRKCVS1_console# set openflow controller primary Test 192.168.0.3 6633  
FRKCVS1_console# set openflow controller backup Test 192.168.1.3 6633  
FRKCVS1_console# set openflow primary control port 192.168.0.14 /24  
FRKCVS1_console# set openflow backup control port 192.168.1.14 /24  
FRKCVS1_console# set openflow dpid 14
```

OpenFlow DPID set to 0x000000000000000e

```
FRKCVS1_console# set openflow controller state Test enabled  
Controller Test enabled (connected)
```

Setting interface Test UP

```
Aug 18 21:24:24.508 ofconnectionmanager: MSG: OFS Connected to Controller  
FRKCVS1_console#
```

After discovery of the DNOS-OF switches, the OpenFlow controller can begin controlling the traffic between them. The next chapter explains how to discover all switches in the OpenFlow data network.

## Discovering OpenFlow Switches

The steps below are used to demonstrate the Openflow switch discovery process within the NEC controller.

- From each OpenFlow switch command line, ensure the Openflow feature is *disabled* using the following command: **set openflow controller state Test disable**.

**Note:** Though the *Configuring N-Series for DNOS OpenFlow* instructions above on page 13 had enabled the OF protocol as part of the complete configuration, temporarily disabling it on each switch will best demonstrate the discovery process explained in this section.

- If already installed, remove all cables between the Openflow switches. Figure 4 depicts these cables as red lines. Do not remove or disconnect any cables used in the Openflow control network. Figure 4 depicts these cables with blue and green lines.

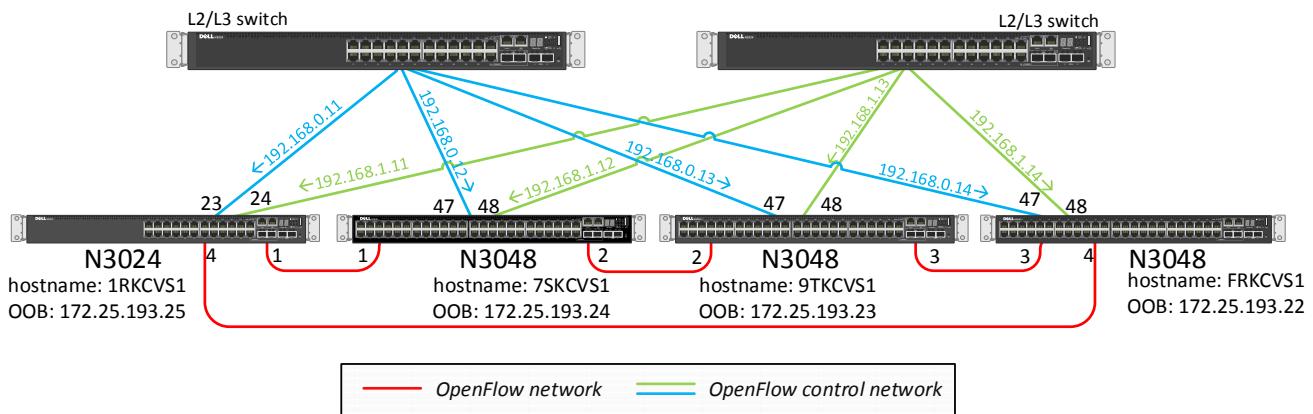


Figure 4 OpenFlow network cabling

The network should now look like Figure 5. Add one OpenFlow network cable at a time for discovery demonstration purposes, as instructed below. The controller topology map and each DNOS-OF switch show this.

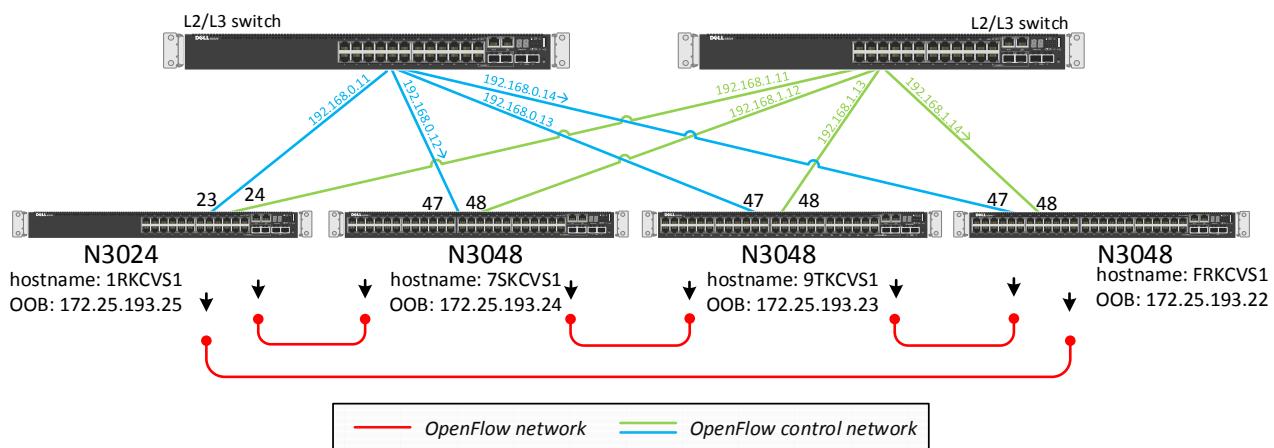


Figure 5 OpenFlow network cables removed

3. Use a web browser to open the controller client interface from a management station. This example uses the NEC PF6800 web app to attach to the NEC controller at <https://172.25.193.213:8082/nwviewer> (the default username/password is Administrator/nwpass). From here, the administrator enables each OpenFlow switch to communicate with the controllers and builds a topology map. Click the PFC icon and view the Map tab to see your available OF switch network (see Figure 6). There are currently no switches available to the controller.

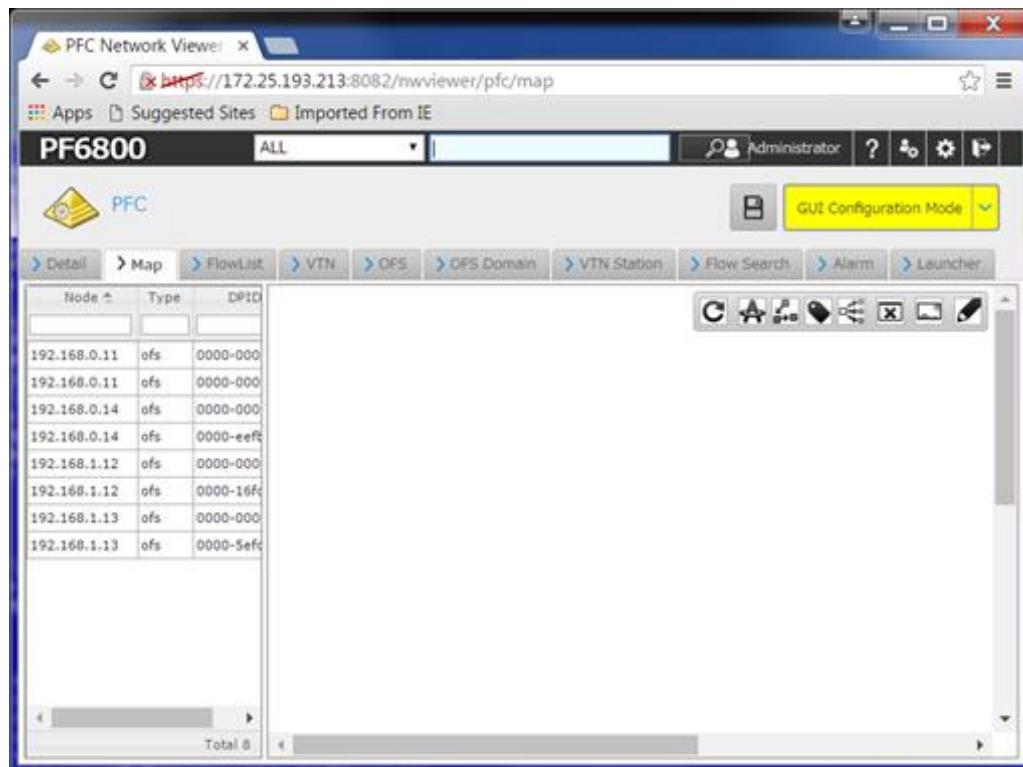


Figure 6 Controller web app

4. Run the following commands from the first OpenFlow switch (1RKCVS1):

The **show openflow flows** command shows the flow entries currently available to the switch.

```
1RKCVS1_console# show openflow flows
```

```
Showing OpenFlow flow entries for all tables
Ingress Port Table (0) Flow Entries
```

```
-----
Number of entries reported = 0
Maximum number of entries for this table = 2000
```

Number of entries actually found = 0

#### VLAN Table (10) Flow Entries

Number of entries reported = 0

Maximum number of entries for this table = 12288

Number of entries actually found = 0

#### Termination MAC Table (20) Flow Entries

Number of entries reported = 0

Maximum number of entries for this table = 512

Number of entries actually found = 0

#### Unicast Routing Table (30) Flow Entries

Number of entries reported = 0

Maximum number of entries for this table = 40960

Number of entries actually found = 0

#### Multicast Routing Table (40) Flow Entries

Number of entries reported = 0

Maximum number of entries for this table = 8191

Number of entries actually found = 0

#### Bridging Table (50) Flow Entries

Number of entries reported = 0

Maximum number of entries for this table = 32767

Number of entries actually found = 0

#### ACL Policy Table (60) Flow Entries

Number of entries reported = 2

Maximum number of entries for this table = 7680

| Flow ID:0xa5c7bb60 Priority:65535 Hard\_time:0 Idle\_time:0 Cookie:0 inPort:mask = 23:0xffffffff  
srcMac:mask = 0000.0000.0000:0000.0000.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0  
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port  
= 0 Destination L4 Port = 0 ICMP Type = 0 pe = 0 ICMP Code = 0 | outPort = -3  
| Flow ID:0xa5c7bb60 Priority:65535 Hard\_time:0 Idle\_time:0 Cookie:0 inPort:mask = 24:0xffffffff  
srcMac:mask = 0000.0000.0000:0000.0000.0000

```
000.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0  
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0  
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port  
= 0 Destination L4 Port = 0 ICMP Ty  
pe = 0 ICMP Code = 0 | outPort = -3
```

Number of entries actually found = 2

-----Multicast Entries-----

```
1RKCVS1_console#
```

**Note:** Notice only the ACL Policy Table has flow entries. Since this is a single table implementation of OpenFlow, all flow entries are listed under the ACL Policy Table. All other tables will remain at 0. Going forward in this document the **show openflow flows 60** command will be used to minimize output and only display the ACL Policy Table entries.

The **show openflow groups** command can also be run to show that no group entries are yet found. These groups will populate once end-to-end traffic flows are established, as seen in later in this document.

```
1RKCVS1_console# show openflow groups  
Showing the OpenFlow groups in the group tables  
No group entries found.
```

```
L2 Interface Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
L2 Rewrite Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
L3 Unicast Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
L2 Multicast Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
L2 Flood Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
L3 Interface Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
L3 Multicast Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
L3 ECMP Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
L2 Overlay Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
MPLS Label Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
MPLS Forward Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
L2 Unfiltered Group: Number of Group Entries:0 Max Group Entries:14336 Max Bucket Entries:1  
OpenFlow Group Statistics
```

-----  
No group entries found.

Run the following command to enable OpenFlow on the first (1RKCVS1) switch:

```
1RKCVS1_console# set openflow controller state Test enabled
```

Setting interface Test UP

```
1RKCVS1_console#
```

5. Press the refresh button on the PFC Network Viewer to show switch 1RKCVS1 (192.168.0.11) on the map.

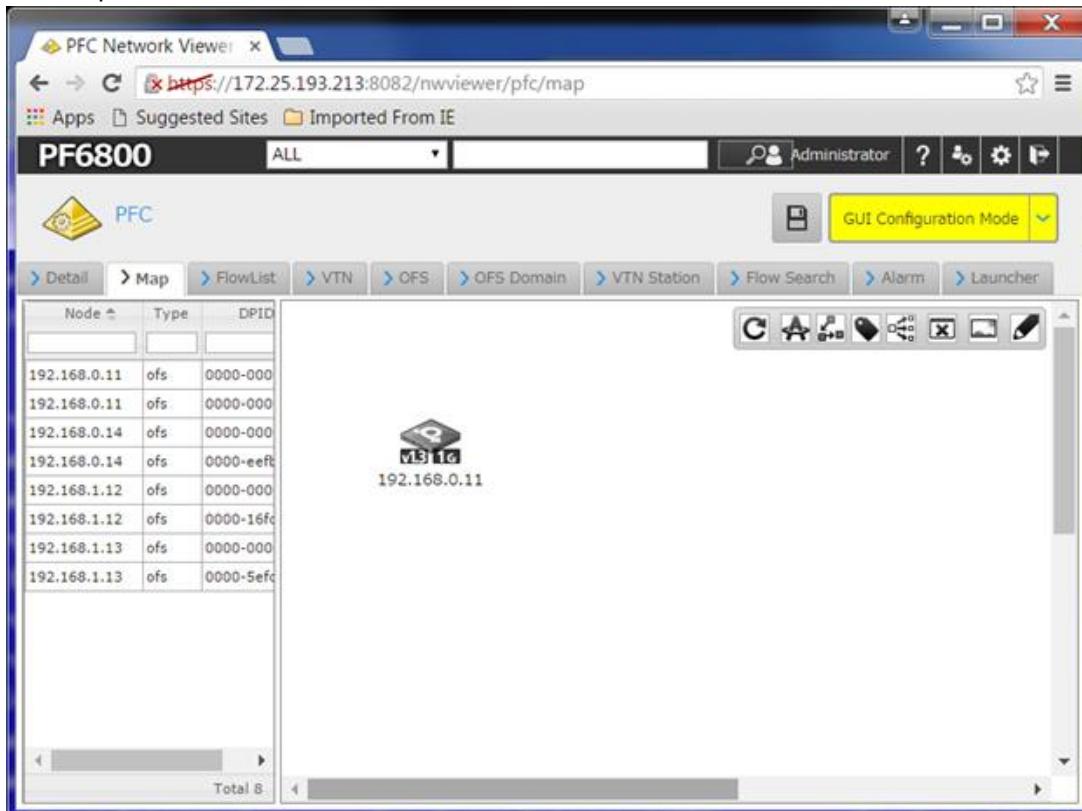


Figure 7 Single OpenFlow switch discovery

**Note:** The NEC controller license must be installed to build network maps

6. Run the **show openflow flows 60** command again from the first OpenFlow switch (1RKCVS1) to see the flow entries now available to the switch.

```
1RKCVS1_console# show openflow flows 60
```

ACL Policy Table (60) Flow Entries

-----  
Number of entries reported = 4

Maximum number of entries for this table = 7680

```

| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:4 inPort:mask = 0:0x0
srcMac:mask = 0000.0000.0000:0000.0000.0000
destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 88cc tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstIp4 = 0.0.0.0/0.0.0.0 srclp6
= ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Type = 0 I
CMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 23:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0
000.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstIp4 = 0.0.0.0/0.0.0.0
srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 24:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0
000.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstIp4 = 0.0.0.0/0.0.0.0
srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
Number of entries actually found = 3

```

1RKCVS1\_console#

- Run the following command from the second OpenFlow switch (7SKCVS1) to enable it as well:

7SKCVS1\_console# **set openflow controller state Test enabled**

Setting interface Test UP

7SKCVS1\_console#

8. Press the refresh button on the PFC Network Viewer to show the addition of the second switch (7SKCVS1) to the map.

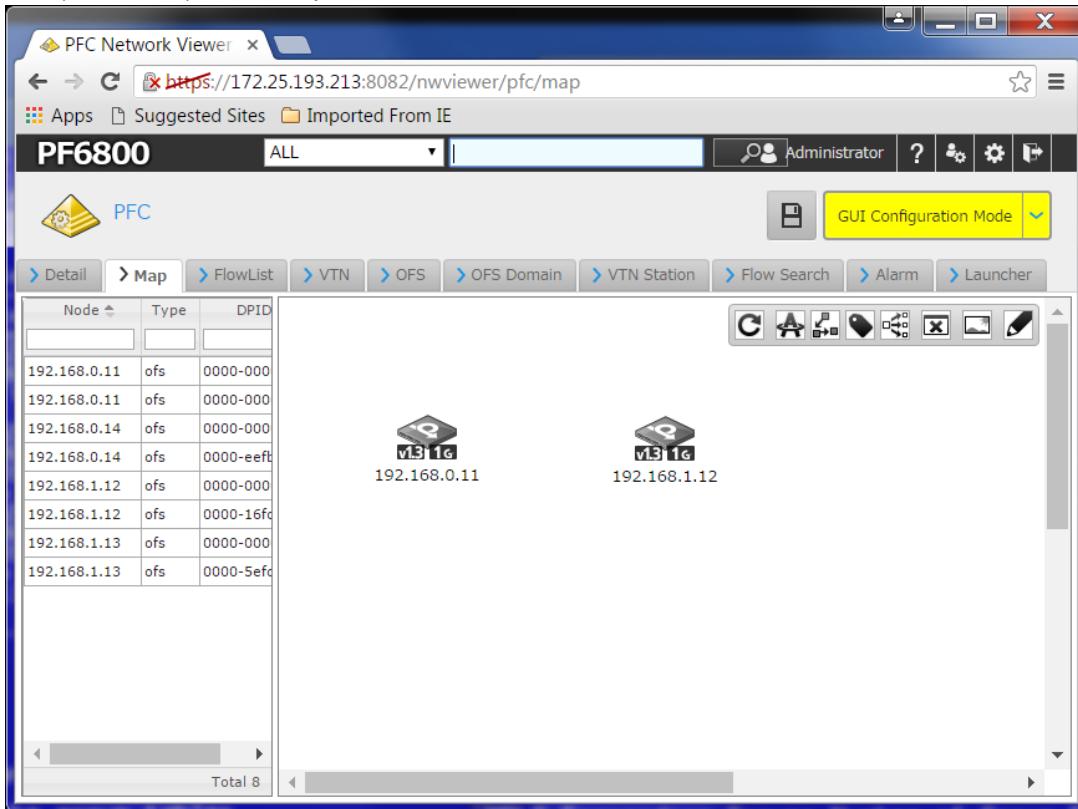


Figure 8 Unconnected OpenFlow switch topology

9. Use the *enable* commands again to enable OpenFlow on the third and fourth switches (9TKCVS1 and FRKCVS1)

```
9TKCVS1_console# set openflow controller state Test enabled
Setting interface Test UP
9TKCVS1_console#
...
FRKCVS1_console# set openflow controller state Test enabled
Setting interface Test UP
FRKCVS1_console#
```

10. Press the refresh button on the PFC Network Viewer to show the remaining switches on the map.

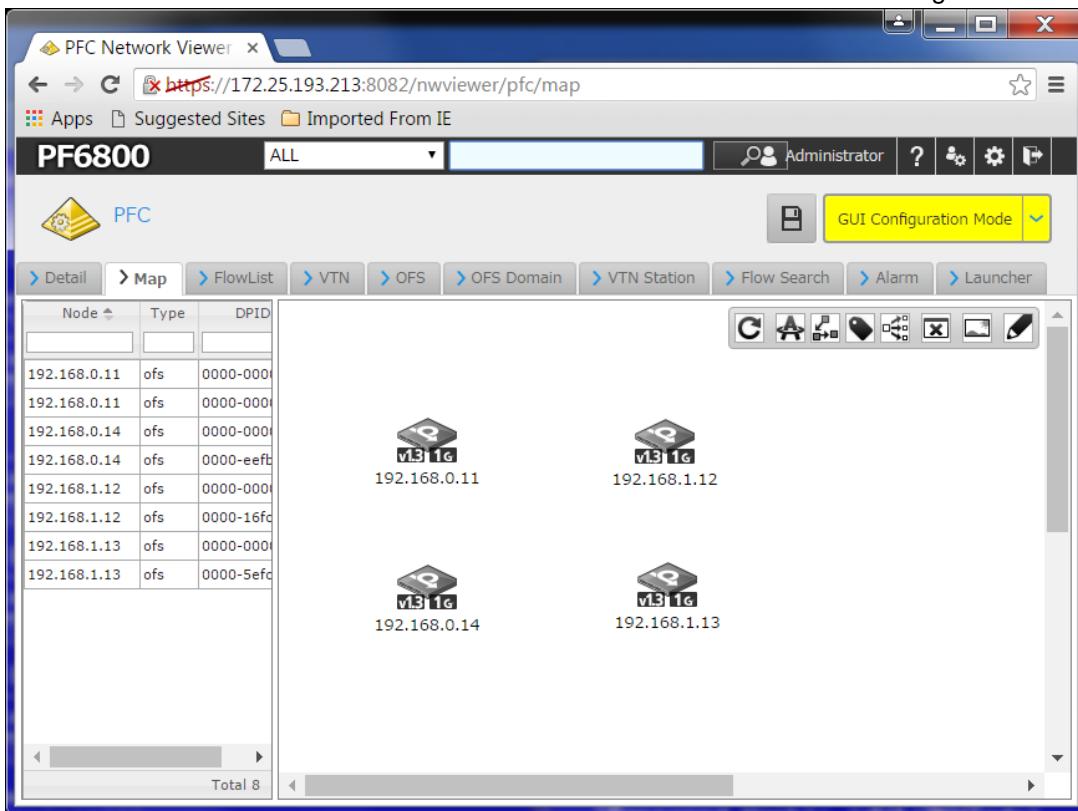


Figure 9 All OpenFlow switches discovered

**Note:** Clicking the pencil icon brings up the edit mode allowing rearrangement of the map.

At this point in the process, there are still no cables connecting the OpenFlow switches. The **show openflow flows** and **show openflow groups** commands continue to provide the same output as seen before.

11. Connect the first two switches (1RKCVS1 and 7SKCVS1) in the OpenFlow network as shown in Figure 10. This demonstration uses port 1 on both switches. Any switch port works except those designated for a particular use, as shown in Table 2 on page 10.

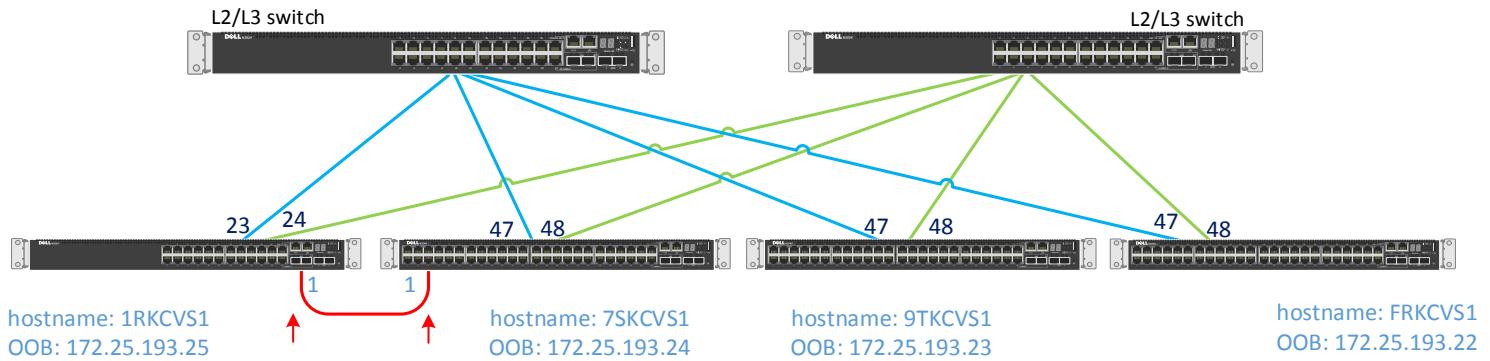


Figure 10 Connecting two OpenFlow switches

12. Refreshing the map on the PFC Network Viewer shows the two switches are now physically connected.

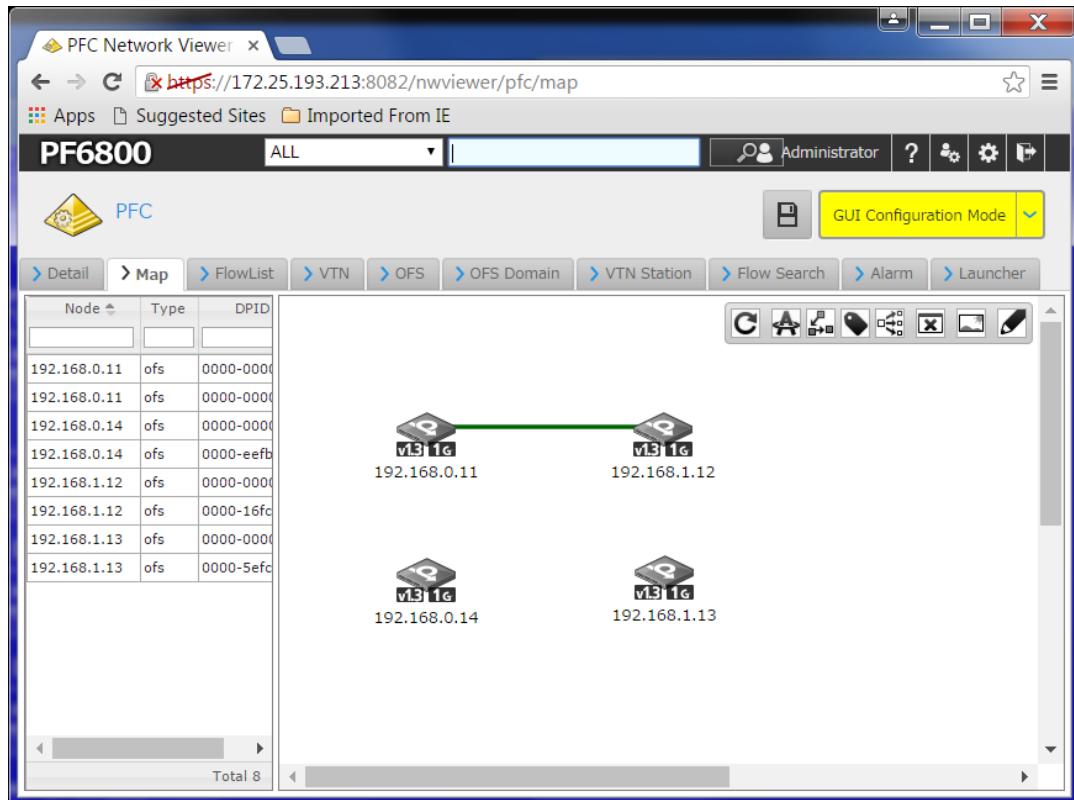


Figure 11 Mapping physical connections between switches

13. Run the command **show openflow flows** on the 1<sup>st</sup> switch to show there are now six reported ACL entries. This number increases by default each time another OpenFlow switch connects directly to the switch.

```
1RKCVS1_console# show openflow flows 60
```

Showing OpenFlow flow entries for all tables  
Ingress Port Table (0) Flow Entries

-----  
Number of entries reported = 0  
Maximum number of entries for this table = 2000  
Number of entries actually found = 0

VLAN Table (10) Flow Entries

-----  
Number of entries reported = 0  
Maximum number of entries for this table = 12288  
Number of entries actually found = 0

Termination MAC Table (20) Flow Entries

-----  
Number of entries reported = 0  
Maximum number of entries for this table = 512  
Number of entries actually found = 0

Unicast Routing Table (30) Flow Entries

-----  
Number of entries reported = 0  
Maximum number of entries for this table = 40960  
Number of entries actually found = 0

Multicast Routing Table (40) Flow Entries

-----  
Number of entries reported = 0  
Maximum number of entries for this table = 8191  
Number of entries actually found = 0

Bridging Table (50) Flow Entries

-----  
Number of entries reported = 0  
Maximum number of entries for this table = 32767  
Number of entries actually found = 0

ACL Policy Table (60) Flow Entries

-----  
Number of entries reported = 6  
Maximum number of entries for this table = 7680  
| Flow ID:0xa5c7bb60 Priority:65535 Hard\_time:0 Idle\_time:0 Cookie:4 inPort:mask = 0:0x0  
srcMac:mask = 0000.0000.0000:0000.0000.0000  
destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 88cc tunnelId = 0 srclp4 =  
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0 srclp6

```

= ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Type = 0 I
CMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:22000 Hard_time:0 Idle_time:0 Cookie:6 inPort:mask = 1:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000 destMac:mask = 0000.0000.0000:0000.0000 etherType = 0000 tunnelId = 0 srclp4
= 0.0.0/0.0.0.0 dstIp4 = 0.0.0/0.0.0
srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Typ
e = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb60 Priority:23000 Hard_time:0 Idle_time:0 Cookie:5 inPort:mask = 1:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =
4094:0xffff srclp4 = 0.0.0/0.0.0 dstIp4 = 0.0.
0.0/0.0.0 srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00
Source L4 Port = 0 Destination L4 Port
= 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 23:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000 destMac:mask = 0000.0000.0000:0000.0000 etherType = 0000 tunnelId = 0
srclp4 = 0.0.0/0.0.0 dstIp4 = 0.0.0/0.0.0
srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 24:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000 destMac:mask = 0000.0000.0000:0000.0000 etherType = 0000 tunnelId = 0
srclp4 = 0.0.0/0.0.0 dstIp4 = 0.0.0/0.0.0
srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3

```

Number of entries actually found = 5

-----Multicast Entries-----

1RKCVS1\_console#

14. Next, connect the 2<sup>nd</sup> and 3<sup>rd</sup> switches (7SKCVS1 and 9TKCVS1) in the OpenFlow network as shown in Figure 12. This demonstration uses port 2 on both switches. Any switch port works except those designated for a particular use as shown in Table 2 on page 10.

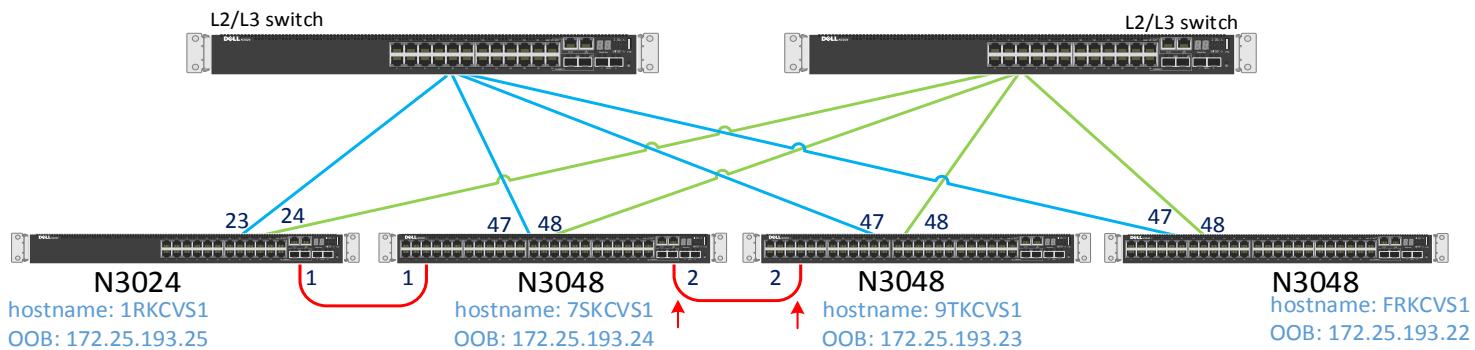


Figure 12 Connecting three OpenFlow switches

15. Refresh the map on the PFC Network Viewer (Figure 13) to show three switches are now physically connected.

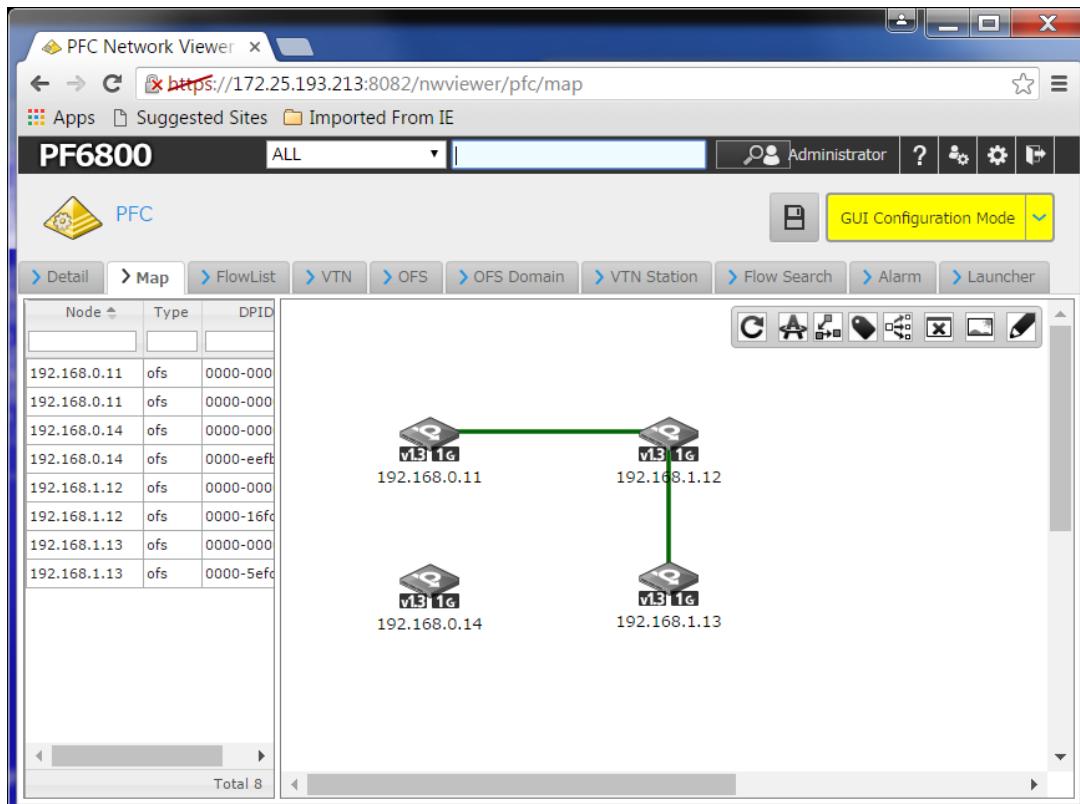


Figure 13 Expanding the OpenFlow topology

16. Run **show openflow flows** on the 1<sup>st</sup> switch again to show that there are still only six reported ACL entries. As previously stated, this number only increases when another OpenFlow switch connects directly to the switch. In this case, the newly added switch does not directly connect to this switch.

```
1RKCVS1_console# show openflow flows 60
```

ACL Policy Table (60) Flow Entries

```
-----
Number of entries reported = 6
Maximum number of entries for this table = 7680
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:4 inPort:mask = 0:0x0
srcMac:mask = 0000.0000.0000:0000.0000.0000
destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 88cc tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0 srclp6
= ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Type = 0
CMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:22000 Hard_time:0 Idle_time:0 Cookie:6 inPort:mask = 1:0xffffffff
srcMac:mask = 0000.0000.0000:0000.00
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Typ
e = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb60 Priority:23000 Hard_time:0 Idle_time:0 Cookie:5 inPort:mask = 1:0xffffffff
srcMac:mask = 0000.0000.0000:0000.00
00.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =
4094:0xffff srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.
0.0/0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00
Source L4 Port = 0 Destination L4 Port
= 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 23:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 24:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
```

Number of entries actually found = 5

```
1RKCVS1_console#
```

17. Connect the fourth switch (FRKCVS1) to both the third switch (9TKCVS1) and the first switch (1RKCVS1) as shown in **Figure 14**. This demonstration uses ports 3 and 4. Use any switch ports except those designated for a particular use as shown in Table 2 on page 10.

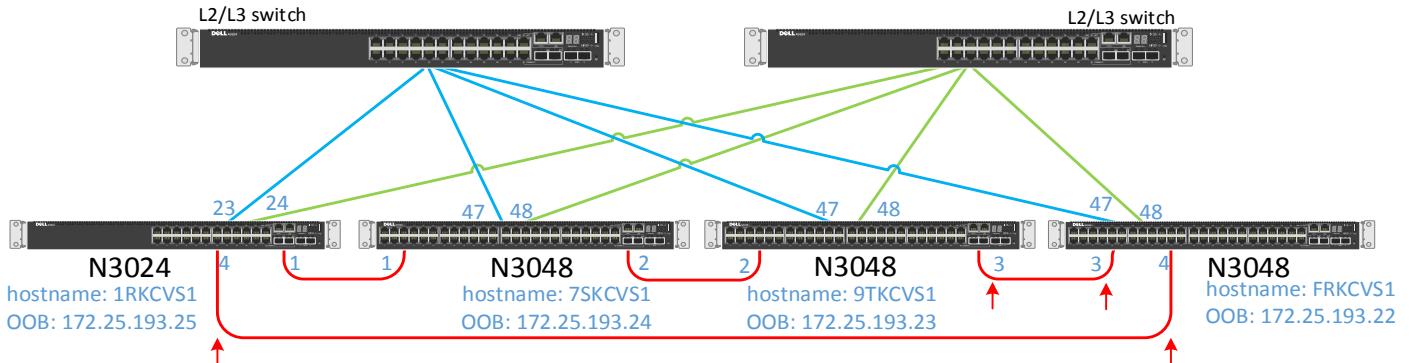


Figure 14 OpenFlow physical network

18. Refresh the map a final time on the PFC Network Viewer to show physical connectivity of all switches in the OpenFlow network.

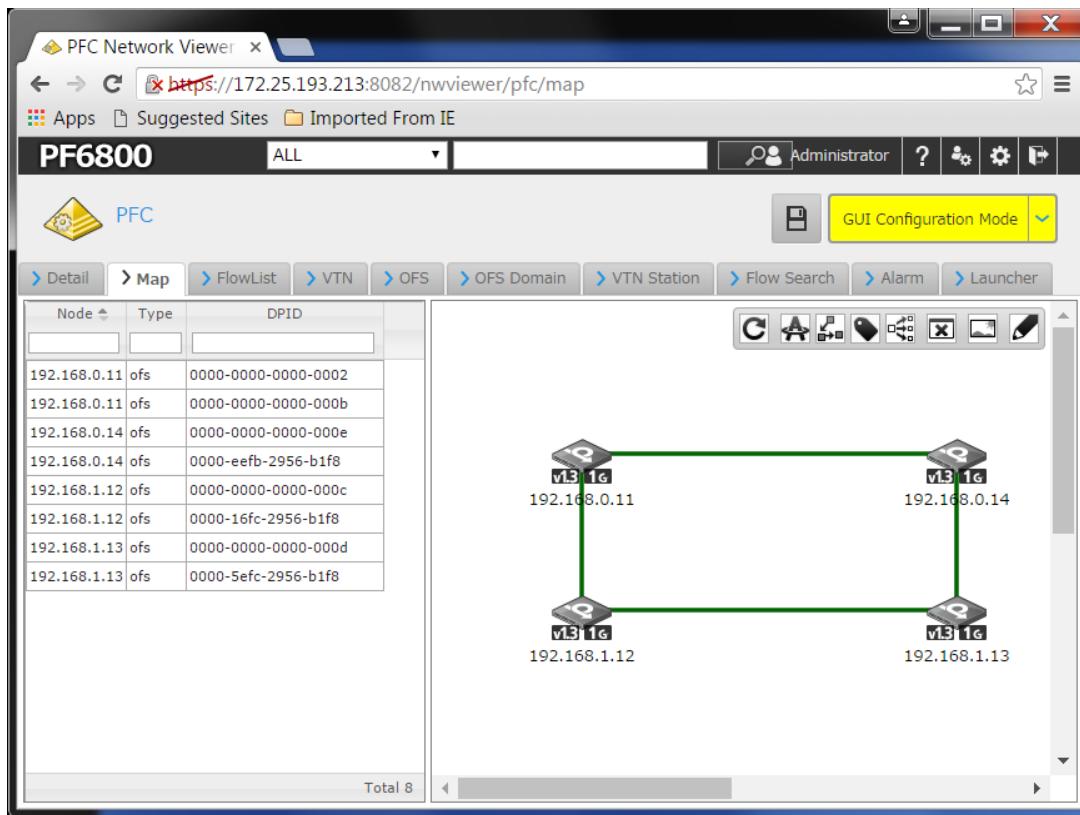


Figure 15 Physical connections between OpenFlow switches

19. Run the **show openflow flows** on the first switch again.

```
1RKCVS1_console# show openflow flows 60
```

#### ACL Policy Table (60) Flow Entries

Number of entries reported = 8

Maximum number of entries for this table = 7680

```
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:12668 inPort:mask = 0:0x0
srcMac:mask = 0000.0000.0000:0000.0000.
0000 destMac:mask = 0000.0000.0000:0000.0000.etherType = 88cc tunneId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0 src
Ip6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Type =
0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:22000 Hard_time:0 Idle_time:0 Cookie:12679 inPort:mask =
1:0xffffffff srcMac:mask = 0000.0000.0000:000
0.0000.0000 destMac:mask = 0000.0000.0000:0000.0000.etherType = 0000 tunneId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.
0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4
Port = 0 Destination L4 Port = 0 ICMP
Type = 0 ICMP Code = 0 | Set output group ID = 0xb0000004 outPort = 0
| Flow ID:0xa5c7bb60 Priority:23000 Hard_time:0 Idle_time:0 Cookie:12678 inPort:mask =
1:0xffffffff srcMac:mask = 0000.0000.0000:000
0.0000.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =
4094:0xffff srclp4 = 0.0.0.0/0.0.0.0 dstlp4 =
0.0.0.0/0.0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00
Source L4 Port = 0 Destination L4 P
ort = 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb60 Priority:22000 Hard_time:0 Idle_time:0 Cookie:12672 inPort:mask =
4:0xffffffff srcMac:mask = 0000.0000.0000:000
0.0000.0000 destMac:mask = 0000.0000.0000:0000.0000.etherType = 0000 tunneId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.
0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4
Port = 0 Destination L4 Port = 0 ICMP
Type = 0 ICMP Code = 0 | Set output group ID = 0xb0000001 outPort = 0
| Flow ID:0xa5c7bb60 Priority:23000 Hard_time:0 Idle_time:0 Cookie:12670 inPort:mask =
4:0xffffffff srcMac:mask = 0000.0000.0000:000
0.0000.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =
4094:0xffff srclp4 = 0.0.0.0/0.0.0.0 dstlp4 =
0.0.0.0/0.0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00
Source L4 Port = 0 Destination L4 P
ort = 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 23:0xffffffff
srcMac:mask = 0000.0000.0000:0000.
000.0000 destMac:mask = 0000.0000.0000:0000.0000.etherType = 0000 tunneId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 24:0xffffffff
srcMac:mask = 0000.0000.0000:0000.
000.0000 destMac:mask = 0000.0000.0000:0000.0000.etherType = 0000 tunneId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.
```

```
srclp6 = :::: dstlp6 = :::: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port  
= 0 Destination L4 Port = 0 ICMP Ty  
pe = 0 ICMP Code = 0 | outPort = -3  
Number of entries actually found = 7
```

```
1RKCVS1_console#
```

The NEC controller now completely controls the OpenFlow network.

## 5

## Configuring end-to-end flows

The procedures in the previous chapters detail DNOS-OF switch configurations that enable any end-to-end flow using a properly cabled OpenFlow controller. To demonstrate, this section uses the NEC PF6800 controller to create both L2 and L3 flows running traffic across each layer to validate the configurations.

## 5.1

### Creating a Layer 2 flow

This example shows how to create an L2 end-to-end flow between two endpoints on the OpenFlow network:

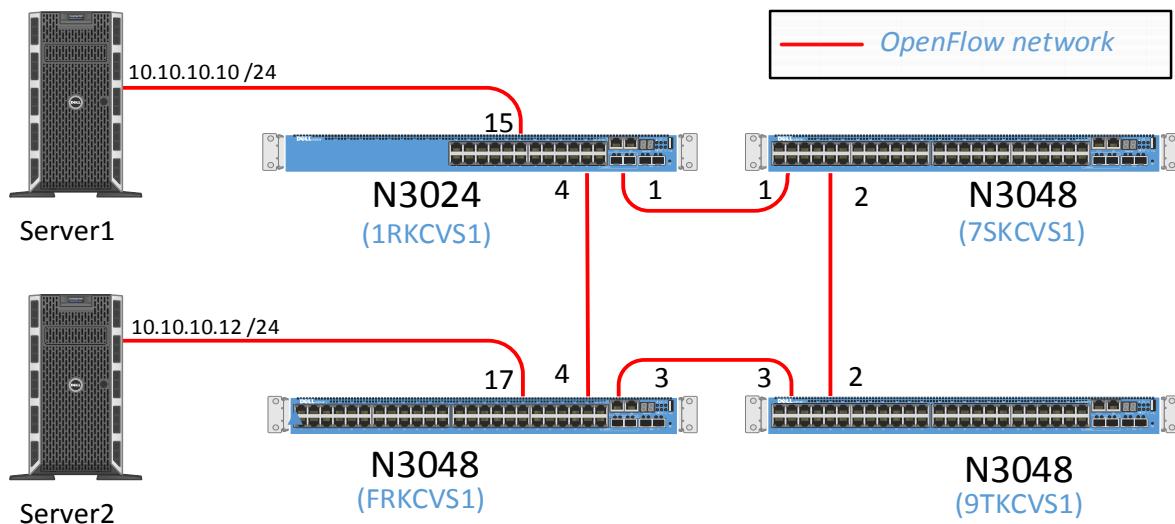


Figure 16 Example topology for a Layer 2 end-to-end flow

This illustration uses two servers running Windows Server 2012 R2. One server connects to port ge1/0/15 on the first switch (1RKCVS1) and the other server connects to port ge1/0/17 on the fourth switch (FRKCVS1). The controller ensures traffic between these two servers takes the shortest path on the OpenFlow network. The shortest path is across the cable that connects port ge1/0/4 on the first switch to port ge1/0/4 on the fourth switch. If the link is broken, the controller ensures proper failover procedures and moves all traffic to an alternative path. In this basic example topology, the only other path between the servers crosses all four switches. It uses ports ge1/0/1 (on two switches), ports ge1/0/2 (on two switches), and ports ge1/0/3 (on two switches).

The output of the **show openflow flows 60** command lists all connected ports on an OF switch. The output below has these ports highlighted to make them easier to locate in this example.

```
1RKCVS1_console# show openflow flows 60
```

#### ACL Policy Table (60) Flow Entries

```
-----  
Number of entries reported = 8  
Maximum number of entries for this table = 7680  
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:4 inPort:mask = 0:0x0  
srcMac:mask = 0000.0000.0000:0000.0000.0000  
destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 88cc tunnelId = 0 srclp4 =  
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0 srclp6  
= ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0  
Destination L4 Port = 0 ICMP Type = 0  
CMP Code = 0 | outPort = -3  
| Flow ID:0xa5c7bb60 Priority:23000 Hard_time:0 Idle_time:0 Cookie:5 inPort:mask = 1:0xffffffff  
srcMac:mask = 0000.0000.0000:0000.00  
0.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =  
4094:0xffff srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.  
0.0/0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00  
Source L4 Port = 0 Destination L4 Port  
= 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0  
| Flow ID:0xa5c7bb60 Priority:22000 Hard_time:0 Idle_time:0 Cookie:8 inPort:mask = 4:0xffffffff  
srcMac:mask = 0000.0000.0000:0000.00  
0.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =  
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0  
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port  
= 0 Destination L4 Port = 0 ICMP Typ  
e = 0 ICMP Code = 0 | Set output group ID = 0xb000000f outPort = 0  
| Flow ID:0xa5c7bb60 Priority:23000 Hard_time:0 Idle_time:0 Cookie:7 inPort:mask = 4:0xffffffff  
srcMac:mask = 0000.0000.0000:0000.00  
0.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =  
4094:0xffff srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.  
0.0/0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00  
Source L4 Port = 0 Destination L4 Port  
= 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0  
| Flow ID:0xa5c7bb60 Priority:2 Hard_time:0 Idle_time:0 Cookie:15 inPort:mask = 15:0xffffffff  
srcMac:mask = 0000.0000.0000:0000.0000  
.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =  
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0  
clp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port =  
0 Destination L4 Port = 0 ICMP Type  
= 0 ICMP Code = 0 | outPort = -3  
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 23:0xffffffff  
srcMac:mask = 0000.0000.0000:0000.0  
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0  
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0  
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port  
= 0 Destination L4 Port = 0 ICMP Ty  
pe = 0 ICMP Code = 0 | outPort = -3  
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 24:0xffffffff  
srcMac:mask = 0000.0000.0000:0000.0  
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0  
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0
```

```
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port  
= 0 Destination L4 Port = 0 ICMP Ty  
pe = 0 ICMP Code = 0 | outPort = -3  
Number of entries actually found = 7
```

-----Multicast Entries-----

```
|cookie:11 priority:22016 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =  
0x0806:0xffff ipProto:mask = 0x00:0x00 inPort  
:15 outports:4, controller  
|cookie:12 priority:22016 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =  
0x0800:0xffff ipProto:mask = 0x70:0xff inPort  
:15 outports:4, controller  
|cookie:13 priority:22016 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =  
0x8035:0xffff ipProto:mask = 0x00:0x00 inPort  
:15 outports:4, controller  
|cookie:14 priority:22000 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =  
0x0000:0x0000 ipProto:mask = 0x00:0x00 inPort  
:15 outports:4,  
1RKCVS1_console#
```

When a server connects to the first switch, the number of entries in the ACL table increases. The number increased to 7 in this example, as shown in the last highlighted area above.

**Note:** OpenFlow flows only increase while a controller is connected to the switch.

Use the following procedure to setup and validate an L2 end-to-end flow on an NEC PF6800 controller. The four major steps are:

1. [VTN L2 map creation](#) maps out a Virtual Tenant Network, a logical layout of the desired network.
2. [vBridge configuration](#) constructs the logical switch settings on the controller.
3. [vExternal \(server\) configuration](#) constructs end-point settings on the controller.
4. [Validation](#) runs tests to verify the setup.

## 5.1.1 VTN L2 map creation

Use the following steps to create a VTN L2 map:

1. From the controller, go to GUI Configuration Mode and the VTN tab to add a new VTN named “L2.” Click **Apply** and **Save**. See Figure 17:

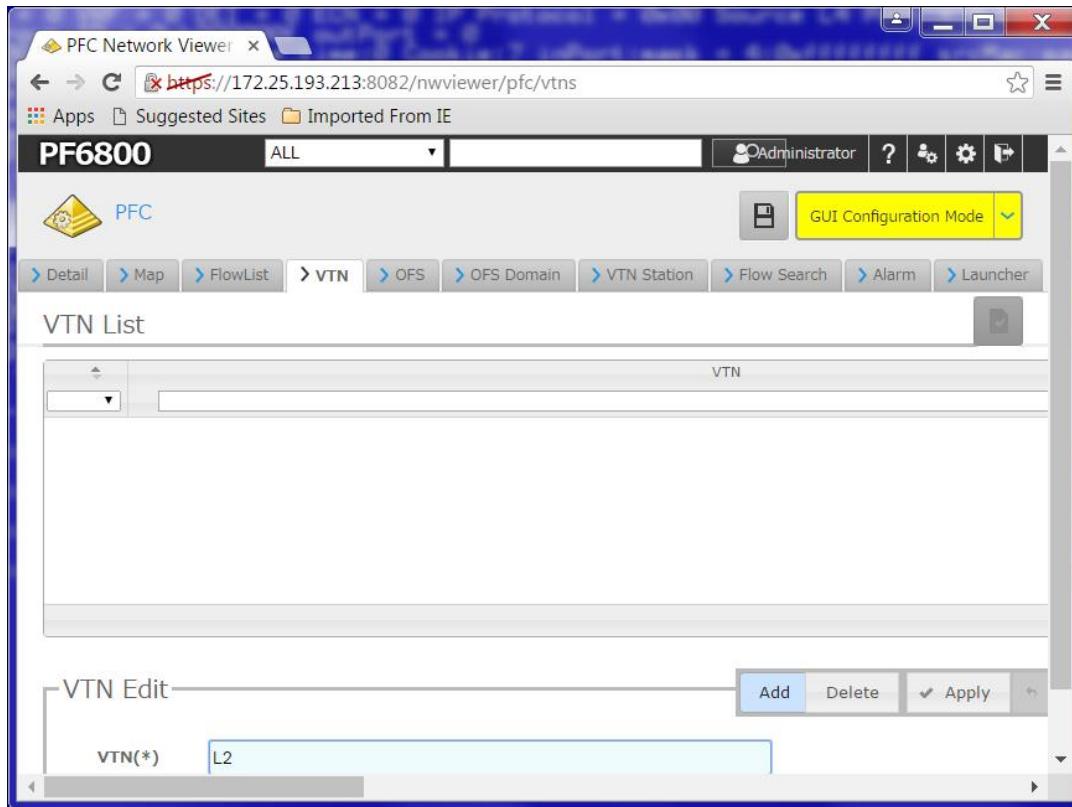


Figure 17 Adding a VTN from the controller

2. Click the **Map** link associated with the newly created L2 (Map) on the VTN tab to open the **VTN Map** tab. The map area displays a blank canvas for creating a virtual network.
3. Click the **Edit** (pencil) icon in the top right corner to open the virtual network building blocks (vRouter, vBridge, and vExternal).
4. Drag a vBridge icon onto the blank canvas and name it **vBR100**. This will be the virtual L2 network.
5. Drag two vExternal icons to the canvas and name them **Server1** and **Server2** to reflect the actual Server1 and Server2 in the example topology.
6. Ctrl-click to select both **Server1** and **vBR100**, then click the “link” button to tell the controller to create a logical link between the two.
7. Use Ctrl-click again to select **Server2** and **vBR100**. Click the “link” button again to tell the controller to create a logical link between them as well.
8. Click **Save**.

The VTN Map should now resemble Figure 18.

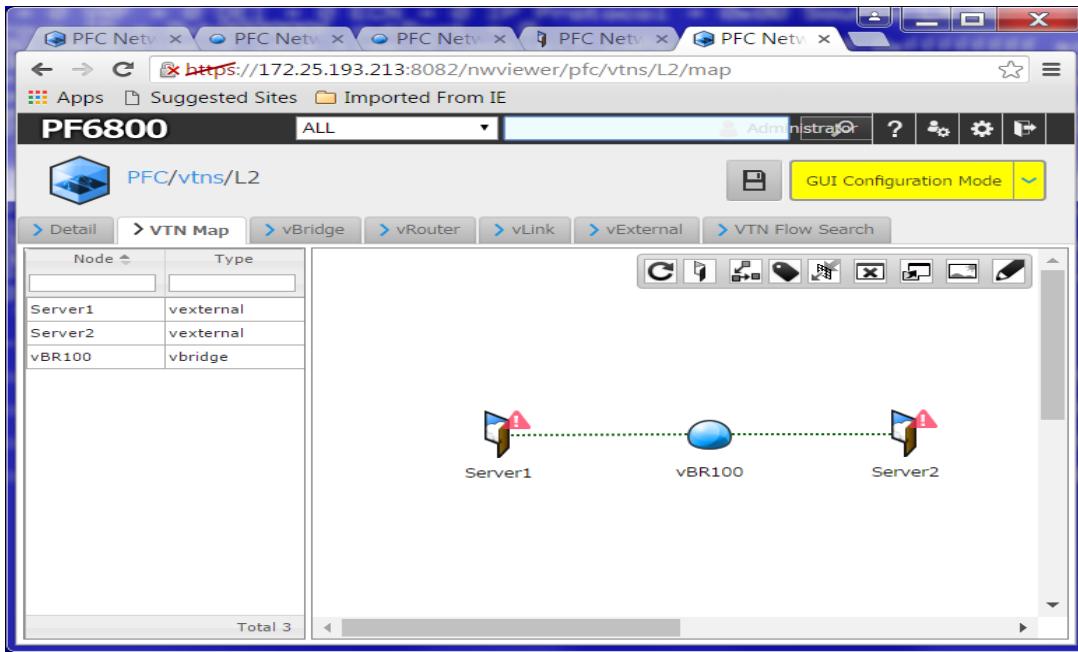


Figure 18 VTN Map of a logical network

### 5.1.2 vBridge configuration

1. Double-click the **vBR100** icon on the map. Click **vBR100** again on the screen presented to open settings.
2. Go to the *VLAN Map* tab for the vBR100, select **tagged** and enter **100** in the *vlan ID* field as shown in Figure 19.

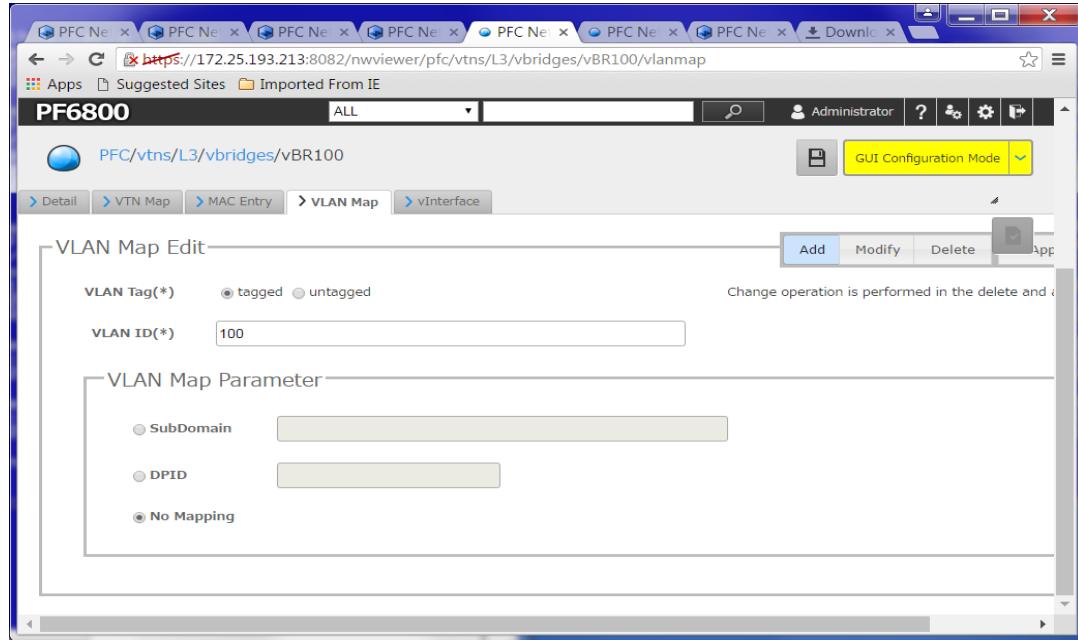


Figure 19 Setting the vBR100 VLAN

3. Click Apply and the Commit (checkmark button).
4. Go back to the VTN map.

### 5.1.3 vExternal (server) configuration

1. Select L2 > vExternal > Server1 and click the **OFS Map** tab.
2. Enter the data path ID (DPID) port and the VLAN ID of 100 as shown in Figure 20.

**Note:** The DPID was defined in section [3.3 - Configuring N-Series for DNOS](#) OpenFlow starting on page 13 and can also be found in the running configuration on each switch.

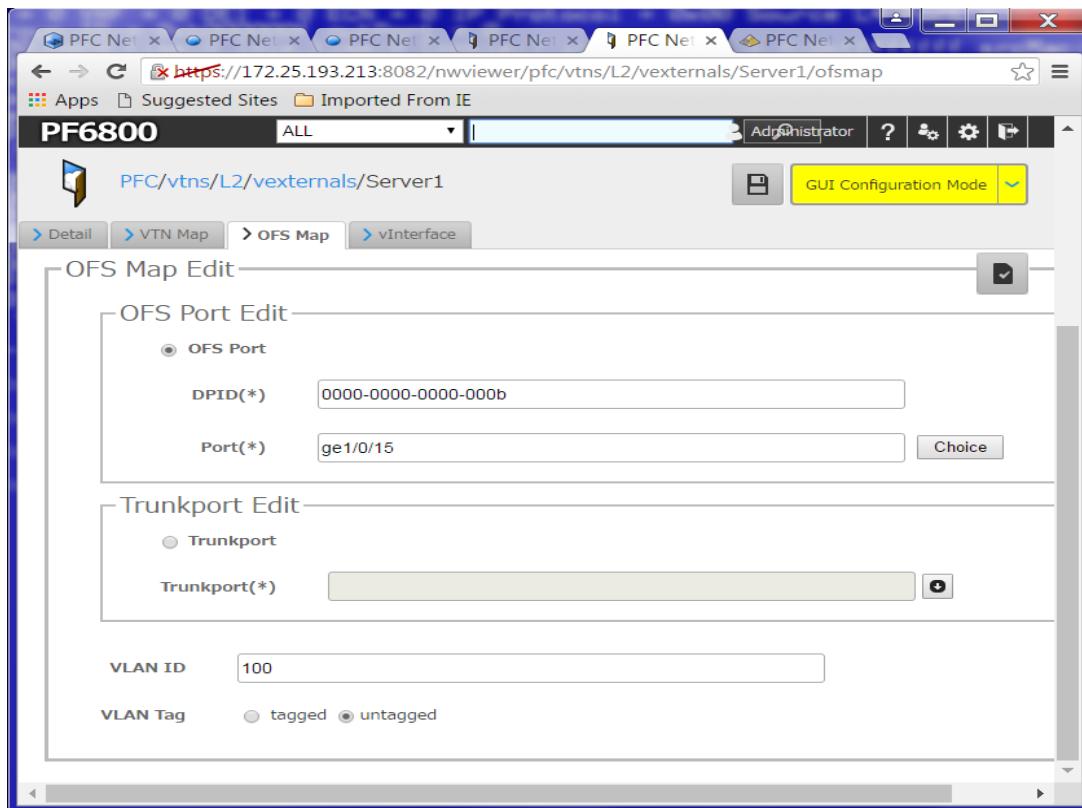


Figure 20     Editing the OFS Map for Server1

3. Click **Commit** (checkmark) on the OFS Map Edit screen to save settings.

Return to the VTN Map to see the physical connection in addition to the logical connection from Server1 to vBR100. See Figure 21:

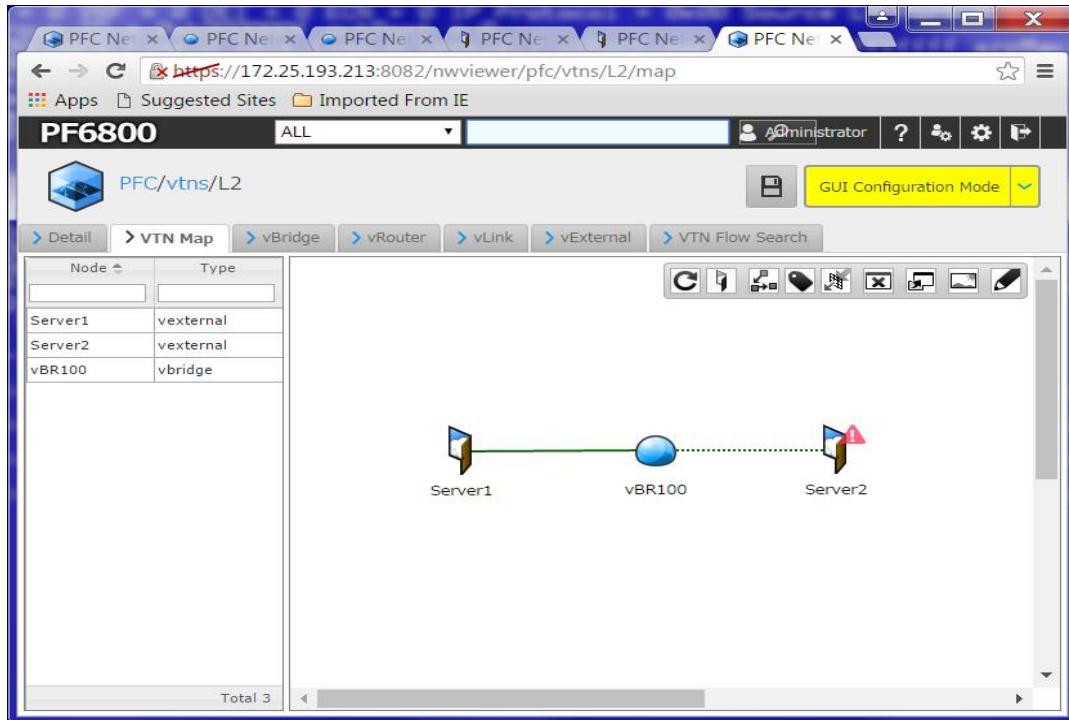


Figure 21      Server1 to vBR100 physical connection

4. Select **L2 > vExternal > Server2** and click the **OFS Map** tab.
5. Enter the data path ID (DPID) port and enter **100** in the VLAN ID field as shown in Figure 22.

**Note:** The DPID was defined in section [3.3 - Configuring N-Series for DNOS](#) OpenFlow starting on page 13 and can also be found in the running configuration on each switch.

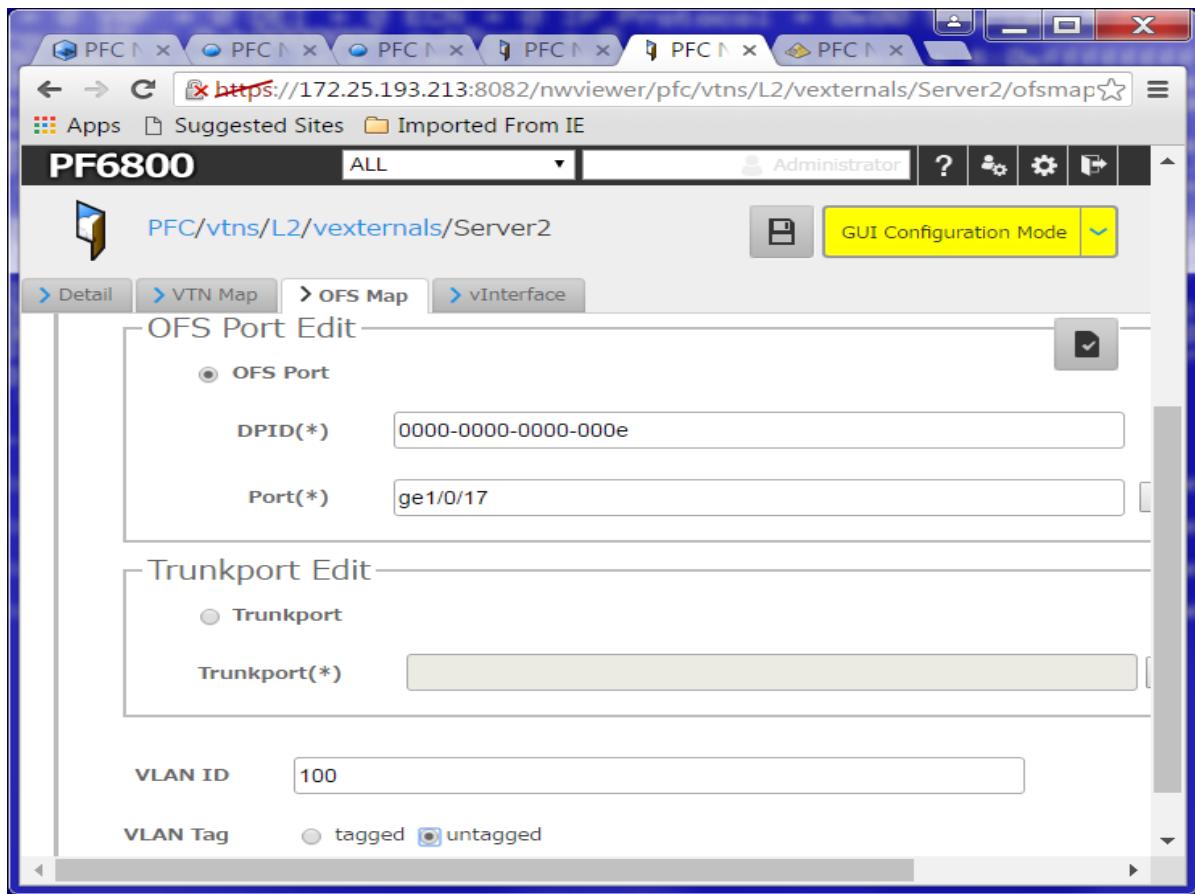


Figure 22 Editing the OFS Map for Server2

6. Click the Commit (checkmark icon) on the OFS Map Edit screen to save settings.

**Hint:** Hovering the mouse pointer over the checkmark icon will display the word “Commit.”

7. Go back to the VTN Map to see the logical connection from Server2 to vBR100 now has a physical connection as well (see Figure 23).

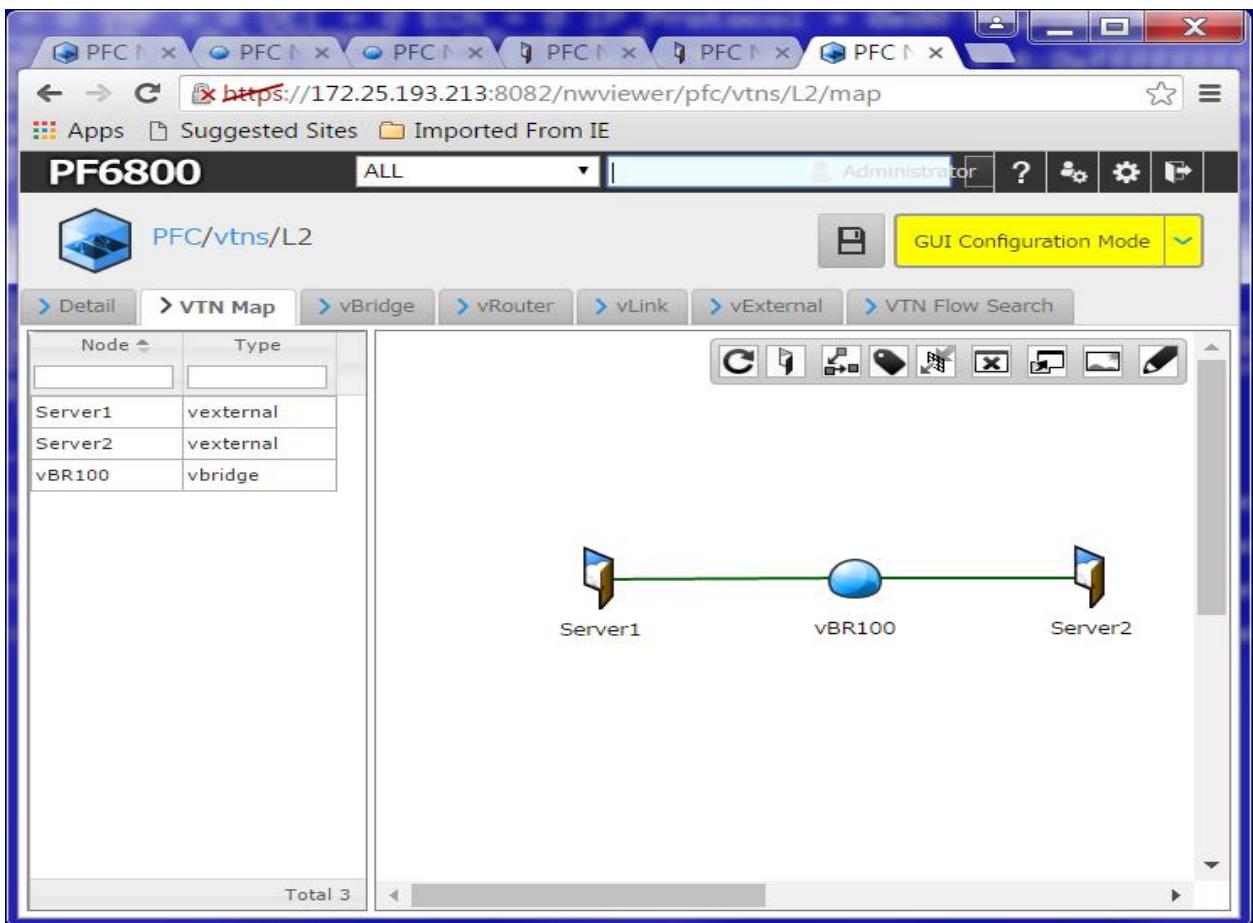


Figure 23 Server1 and Server2 connections to vBR100

- From the first switch (1RKCVS1), run the **show openflow flows 60** and **run openflow groups** commands again to see the highlighted new entries to both flows and groups.

```
1RKCVS1_console# show openflow flows 60
```

ACL Policy Table (60) Flow Entries

```
-----
Number of entries reported = 9
Maximum number of entries for this table = 7680
| Flow ID:0xa5c7bb60 Priority:24576 Hard_time:0 Idle_time:0 Cookie:16 inPort:mask = 0:0x0
srcMac:mask = 0000.0000.0000:0000.0000.000
0 destMac:mask = 0200.0038.0000:ff1f.fff8.0000 etherType = 0000 vlanId:mask = 4094:0xffff
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.
0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4
Port = 0 Destination L4 Port = 0 IC
MP Type = 0 ICMP Code = 0 | Set output group ID = 0x ffe0004 outPort = 0
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:4 inPort:mask = 0:0x0
srcMac:mask = 0000.0000.0000:0000.0000.0000
destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 88cc tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0 srclp6
```

```

= ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Type = 0 I
CMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:23000 Hard_time:0 Idle_time:0 Cookie:5 inPort:mask = 1:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000
00.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =
4094:0xffff srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.
0.0/0.0.0 srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00
Source L4 Port = 0 Destination L4 Port
= 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb60 Priority:22000 Hard_time:0 Idle_time:0 Cookie:8 inPort:mask = 4:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0
srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port =
0 Destination L4 Port = 0 ICMP Typ
e = 0 ICMP Code = 0 | Set output group ID = 0xb000000f outPort = 0
| Flow ID:0xa5c7bb60 Priority:23000 Hard_time:0 Idle_time:0 Cookie:7 inPort:mask = 4:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000
00.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =
4094:0xffff srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.
0.0/0.0.0 srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00
Source L4 Port = 0 Destination L4 Port
= 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb60 Priority:2 Hard_time:0 Idle_time:0 Cookie:15 inPort:mask = 15:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000
.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0 sr
clp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port =
0 Destination L4 Port = 0 ICMP Type
= 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 23:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0
srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port =
0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 24:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0
srclp6 = ::/:: dstIp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port =
0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3

```

Number of entries actually found = 8

-----Multicast Entries-----

```

|cookie:11 priority:22016 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =
0x0806:0xffff ipProto:mask = 0x00:0x00 inPort
:15 outports:4, controller

```

```
|cookie:12 priority:22016 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =  
0x0800:0xffff ipProto:mask = 0x70:0xff inPort  
:15 outports:4, controller  
|cookie:13 priority:22016 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =  
0x8035:0xffff ipProto:mask = 0x00:0x00 inPort  
:15 outports:4, controller  
|cookie:14 priority:22000 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =  
0x0000:0x0000 ipProto:mask = 0x00:0x00 inPort  
:15 outports:4,
```

1RKCVS1\_console# **show openflow groups**

Showing the OpenFlow groups in the group tables

```
Group Type:0, Group VLAN ID:4094 Group Port ID:4  
GroupId=0x0ffe0004 (L2 Interface, VLAN ID = 4094, Port ID = 4): Entry Count=2  
    BucketIndex=0 outputPort=4 popVlanTag=0 vlanPcp=0 vlanCfi=0 dscp=0  
Group Type:11, Group VLAN ID:0 Group Port ID:15  
GroupId=0xb000000f (L2 Unfiltered Interface, VLAN ID = 0, Index = 15): Entry Count=2  
    BucketIndex=0 outputPort = 15 vlanPcp=0 vlanCfi=0 dscp=0
```

```
L2 Interface Group: Number of Group Entries:1 Max Group Entries:1433 Max Bucket Entries:1  
L2 Rewrite Group: Number of Group Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
L3 Unicast Group: Number of Group Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
L2 Multicast Group: Number of Group Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
L2 Flood Group: Number of Group Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
L3 Interface Group: Number of Group Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
L3 Multicast Group: Number of Group Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
L3 ECMP Group: Number of Group Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
L2 Overlay Group: Number of Group Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
MPLS Label Group: Number of Group Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
MPLS Forwarding Group: Number of Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
L2 Unfiltered Group: Number of Entries:1 Max Group Entries:14336 Max Bucket Entries:1  
OpenFlow Group Statistics  
-----
```

```
GroupId = 0x0ffe0004 (L2 Interface, VLAN ID = 4094, Port ID = 4): Group Stats - Duration:177,  
RefCount:1  
GroupId = 0xb000000f (L2 Unfiltered Interface, VLAN ID = 0, Index = 15): Group Stats -  
Duration:2042, RefCount:1  
1RKCVS1_console#
```

## 5.1.4 Validation

To validate the end-to-end connection between Server1 and Server2, simply ping between the two devices. A ping should be successful from Server1 (10.10.10.10) to Server2 (10.10.10.12), and vice versa. After completing a ping, re-run the commands below to see the new entries (highlighted).

```
1RKCVS1_console# show openflow groups
Showing the OpenFlow groups in the group tables
Group Type:0, Group VLAN ID:1 Group Port ID:15
GroupId=0x0001000f (L2 Interface, VLAN ID = 1, Port ID = 15): Entry Count=5
    BucketIndex=0 outputPort=15 popVlanTag=1 vlanPcp=0 vlanCfi=0 dscp=0
Group Type:0, Group VLAN ID:4094 Group Port ID:4
GroupId=0x0ffe0004 (L2 Interface, VLAN ID = 4094, Port ID = 4): Entry Count=5
    BucketIndex=0 outputPort=4 popVlanTag=0 vlanPcp=0 vlanCfi=0 dscp=0
Group Type:1, Group VLAN ID:0 Group Port ID:0
GroupId=0x10000001 (L2 Rewrite, Index = 1): Entry Count=5
    BucketIndex=0 referenceGroupId = 0x0001000f vlanId = 1 srcMac: 00:00:00:00:00:00
dstMac: BC:30:5B:F3:1B:9
Group Type:1, Group VLAN ID:0 Group Port ID:0
GroupId=0x10000002 (L2 Rewrite, Index = 2): Entry Count=5
    BucketIndex=0 referenceGroupId = 0x0ffe0004 vlanId = 4094 srcMac: 00:00:00:00:00:00
dstMac: 02:20:00:38:0
Group Type:11, Group VLAN ID:0 Group Port ID:15
GroupId=0xb000000f (L2 Unfiltered Interface, VLAN ID = 0, Index = 15): Entry Count=5
    BucketIndex=0 outputPort = 15 vlanPcp=0 vlanCfi=0 dscp=0

L2 Interface Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L2 Rewrite Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L3 Unicast Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L2 Multicast Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L2 Flood Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L3 Interface Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L3 Multicast Group: Number of Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L3 ECMP Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L2 Overlay Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
MPLS Label Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
MPLS Forwarding Group: Number of Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L2 Unfiltered Group: Number of Entries:2 Max Group Entries:14336 Max Bucket Entries:1
OpenFlow Group Statistics
-----
GroupId = 0x0001000f (L2 Interface, VLAN ID = 1, Port ID = 15): Group Stats - Duration:19,
RefCount:1
GroupId = 0x0ffe0004 (L2 Interface, VLAN ID = 4094, Port ID = 4): Group Stats - Duration:314,
RefCount:2
GroupId = 0x10000001 (L2 Rewrite, Index = 1): Group Stats - Duration:19, RefCount:1
GroupId = 0x10000002 (L2 Rewrite, Index = 2): Group Stats - Duration:18, RefCount:1
GroupId = 0xb000000f (L2 Unfiltered Interface, VLAN ID = 0, Index = 15): Group Stats -
Duration:2179, RefCount:1
```

```
1RKCVS1_console# show openflow flows 60
```

```
ACL Policy Table (60) Flow Entries
-----
```

Number of entries reported = 11  
 Maximum number of entries for this table = 7680  
 | Flow ID:0xa5c7bb60 Priority:24576 Hard\_time:0 Idle\_time:0 Cookie:16 inPort:mask = 0:0x0  
 srcMac:mask = 0000.0000.0000:0000.0000.000  
 0 destMac:mask = 0200.0038.0000:ff1f.fff8.0000 etherType = 0000 vlanId:mask = 4094:0xffff  
 srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.  
 0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4  
 Port = 0 Destination L4 Port = 0 IC  
 MP Type = 0 ICMP Code = 0 | Set output group ID = 0x ffe0004 outPort = 0  
 | Flow ID:0xa5c7bb60 Priority:24600 Hard\_time:0 Idle\_time:0 Cookie:17 inPort:mask = 0:0x0  
 srcMac:mask = 0000.0000.0000:0000.0000.000  
 0 destMac:mask = 0200.0040.0001:ff03.ffff.ffff etherType = 0000 vlanId:mask = 4094:0xffff srclp4  
 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0/0.  
 0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4  
 Port = 0 Destination L4 Port = 0 IC  
 MP Type = 0 ICMP Code = 0 | Set output group ID = 0x10000001 outPort = 0  
 | Flow ID:0xa5c7bb60 Priority:65535 Hard\_time:0 Idle\_time:0 Cookie:4 inPort:mask = 0:0x0  
 srcMac:mask = 0000.0000.0000:0000.0000.000  
 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 88cc tunnelId = 0 srclp4 =  
 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0 srclp6  
 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0  
 Destination L4 Port = 0 ICMP Type = 0 I  
 CMP Code = 0 | outPort = -3  
 | Flow ID:0xa5c7bb60 Priority:23000 Hard\_time:0 Idle\_time:0 Cookie:5 inPort:mask = 1:0xffffffff  
 srcMac:mask = 0000.0000.0000:0000.00  
 0.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =  
 4094:0xffff srclp4 = 0.0.0.0/0.0.0 dstlp4 = 0.0.  
 0.0/0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00  
 Source L4 Port = 0 Destination L4 Port  
 = 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0  
 | Flow ID:0xa5c7bb60 Priority:22000 Hard\_time:0 Idle\_time:0 Cookie:8 inPort:mask = 4:0xffffffff  
 srcMac:mask = 0000.0000.0000:0000.00  
 0.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4  
 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0  
 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port  
 = 0 Destination L4 Port = 0 ICMP Typ  
 e = 0 ICMP Code = 0 | Set output group ID = 0xb000000f outPort = 0  
 | Flow ID:0xa5c7bb60 Priority:23000 Hard\_time:0 Idle\_time:0 Cookie:7 inPort:mask = 4:0xffffffff  
 srcMac:mask = 0000.0000.0000:0000.00  
 0.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =  
 4094:0xffff srclp4 = 0.0.0.0/0.0.0 dstlp4 = 0.0.  
 0.0/0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00  
 Source L4 Port = 0 Destination L4 Port  
 = 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0  
 | Flow ID:0xa5c7bb60 Priority:2 Hard\_time:0 Idle\_time:0 Cookie:15 inPort:mask = 15:0xffffffff  
 srcMac:mask = 0000.0000.0000:0000.0000  
 .0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =  
 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0 sr  
 clp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port =  
 0 Destination L4 Port = 0 ICMP Type  
 = 0 ICMP Code = 0 | outPort = -3  
 | Flow ID:0xa5c7bb60 Priority:25000 Hard\_time:0 Idle\_time:300 Cookie:18 inPort:mask =  
 15:0xffffffff srcMac:mask = bc30.5bf3.1b9d:fff

```

f.ffff.ffff destMac:mask = 5c26.0a42.4768:ffff.ffff.ffff etherType = 0000 vlanId:mask = 1:0xffff srclp4
= 0.0.0.0/0.0.0.0 dstlp4 = 0.0
.0.0/0.0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00
Source L4 Port = 0 Destination L4 Port
= 0 ICMP Type = 0 ICMP Code = 0 | Set output group ID = 0x10000002 outPort = 0
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 23:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0
000.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0/0.0.0
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb60 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 24:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0
000.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0/0.0.0
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3

```

Number of entries actually found = 10

-----Multicast Entries-----

```

|cookie:11 priority:22016 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =
0x0806:0xffff ipProto:mask = 0x00:0x00 inPort
:15 outports:4, controller
|cookie:12 priority:22016 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =
0x0800:0xffff ipProto:mask = 0x70:0xff inPort
:15 outports:4, controller
|cookie:13 priority:22016 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =
0x8035:0xffff ipProto:mask = 0x00:0x00 inPort
:15 outports:4, controller
|cookie:14 priority:22000 destMac:mask 0100.0000.0000:0100.0000.0000 etherType:mask =
0x0000:0x0000 ipProto:mask = 0x00:0x00 inPort
:15 outports:4,
1RKCVS1_console#

```

Flow directions can also be seen from the NEC controller. From the VTN Map, GUI Monitoring Mode, click the “Flows” button to bring up the list of configured flows. In this case, two flows should be displayed, 65727

(Server1) and 65729 (Server2). Click flow ID 65727 to see the traffic flow direction from Server1 to Server2 (see Figure 24).

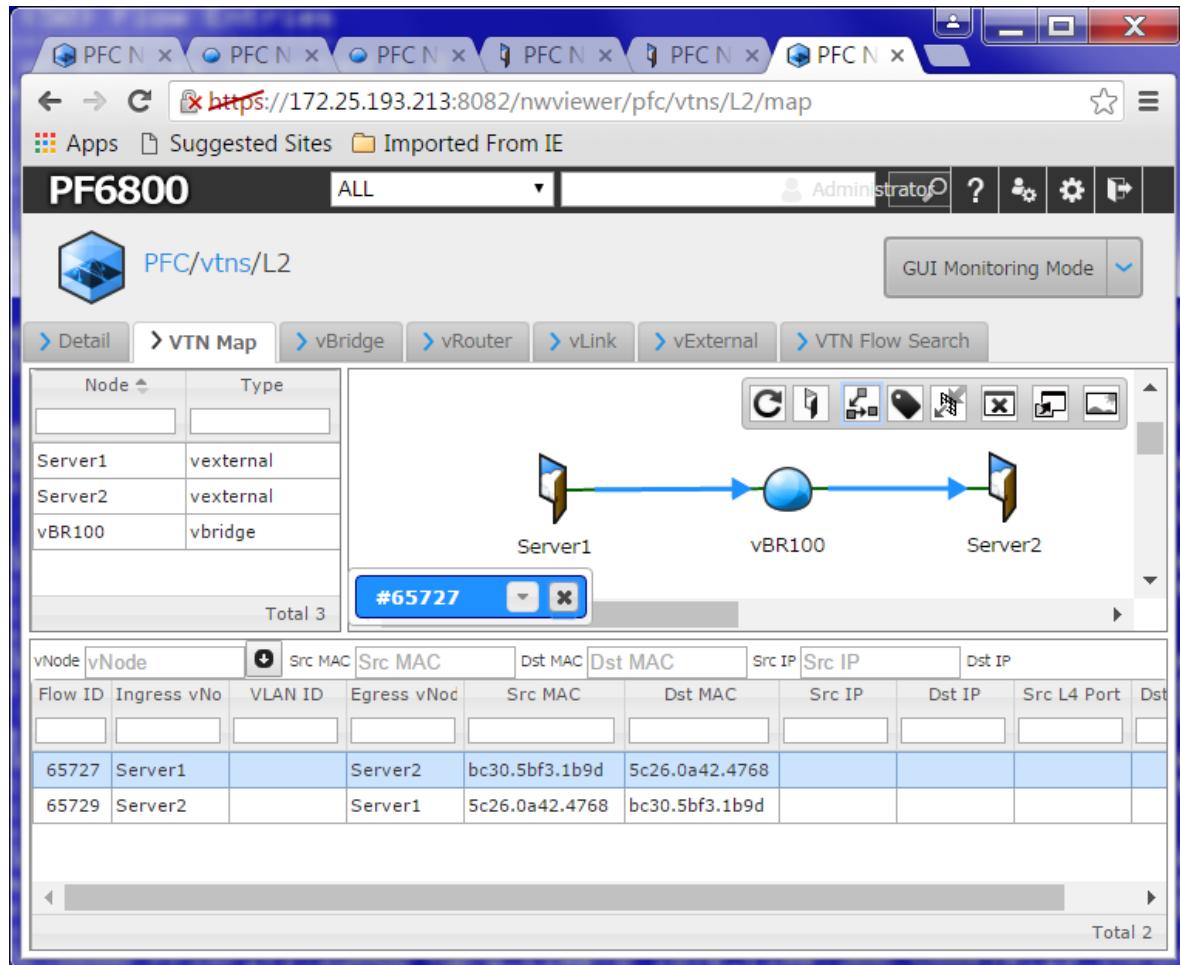


Figure 24 Flow from Server1 to Server2

Click flow ID 65729 (Server2) to see the traffic flow direction from Server2 to Server1 (see Figure 25) .

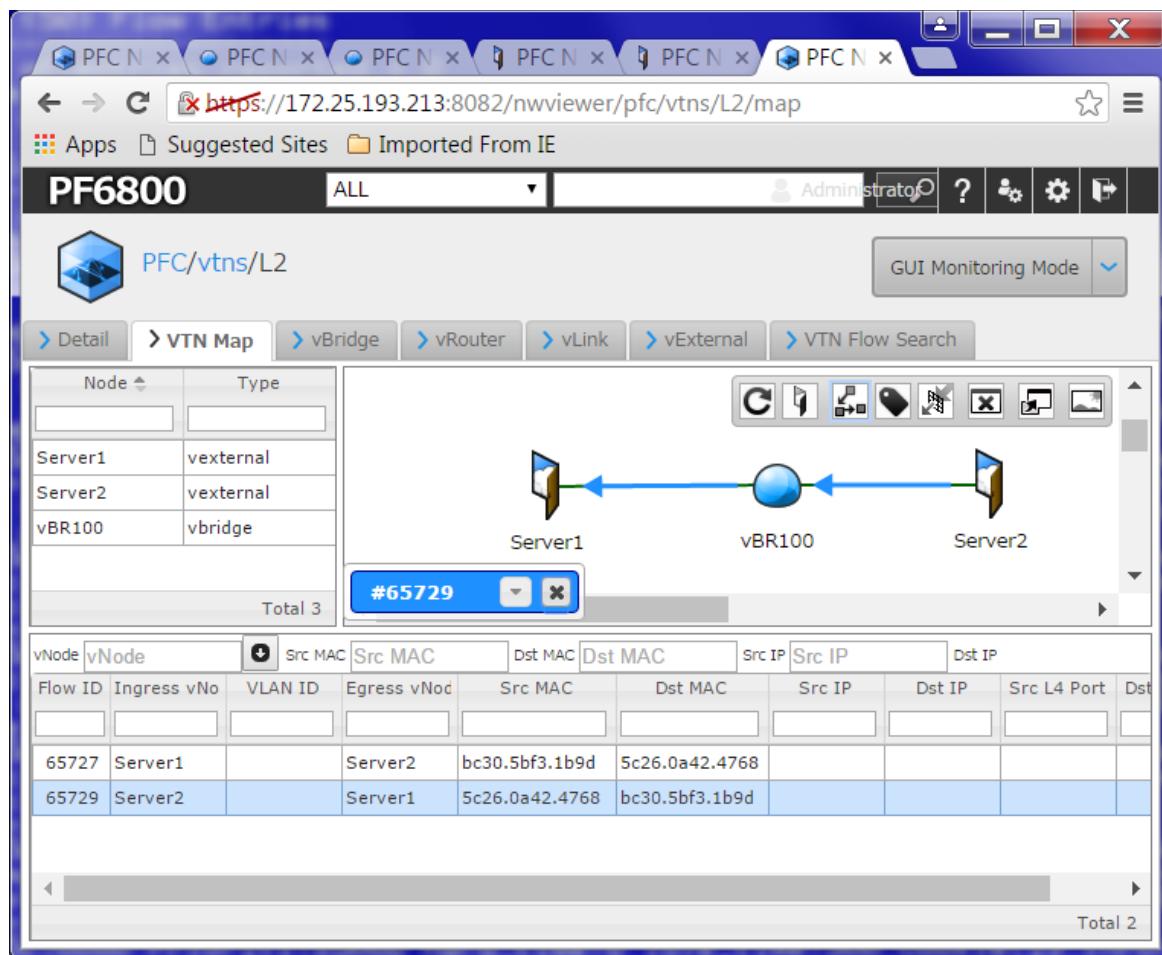


Figure 25 Flow from Server2 to Server1

Network traffic can now pass between the two endpoints, Server1 and Server2.

## 5.2 Creating a Layer 3 flow

This example shows how to create an L3 end-to-end flow between two endpoints on different subnets on the OpenFlow network. See Figure 26.

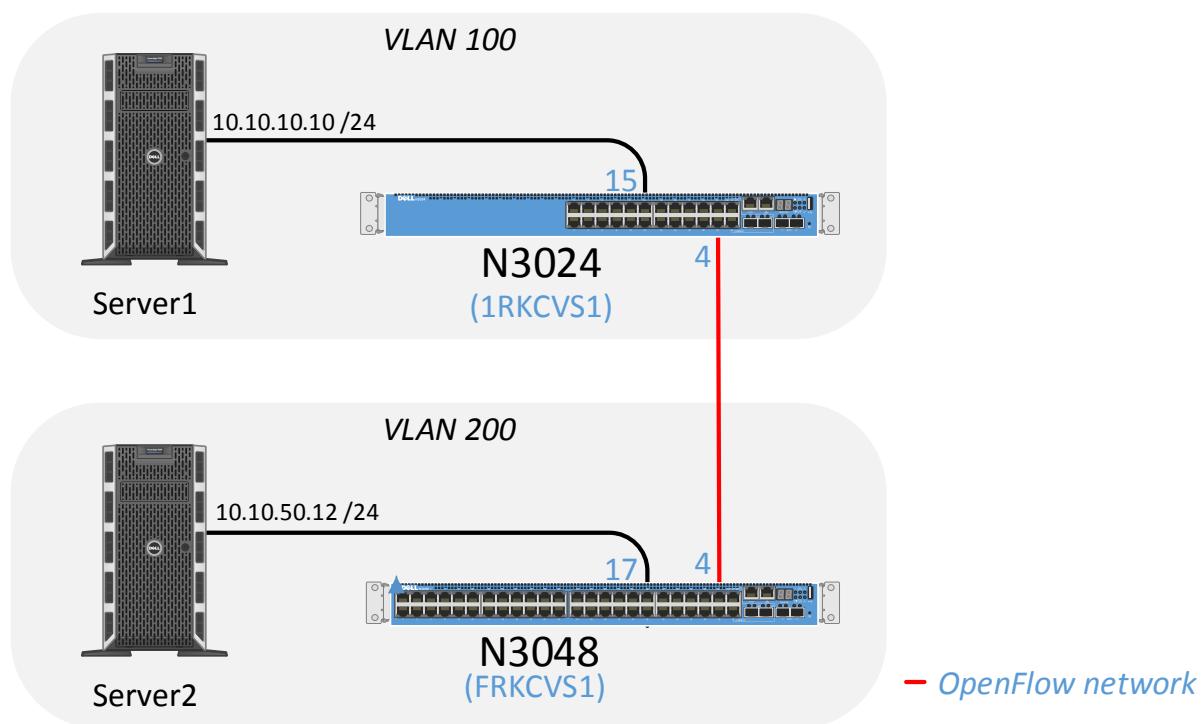


Figure 26 Simple topology for a Layer 3 end-to-end flow

This illustration uses two servers running Windows Server 2012 R2. One connects to port ge1/0/15 on one switch (1RKCVS1). The other connects to port ge1/0/17 on a different switch (FRKCVS1). The controller ensures traffic between these two servers takes the shortest path on the OpenFlow network, and moves traffic to an alternative path if the primary link goes down.

## Redundant Paths

As with any network, best practice is to have alternative routes should the working routes fail. Figure 27 shows an additional path between Server1 and Server2. Notice the figure has the same physical setup as the L2 scenario in Figure 16. The difference is Server2 now occupies a different subnet than Server1 to demonstrate how OpenFlow routes the traffic between VLANs.

See page 68 for a more detailed discussion of Failover.

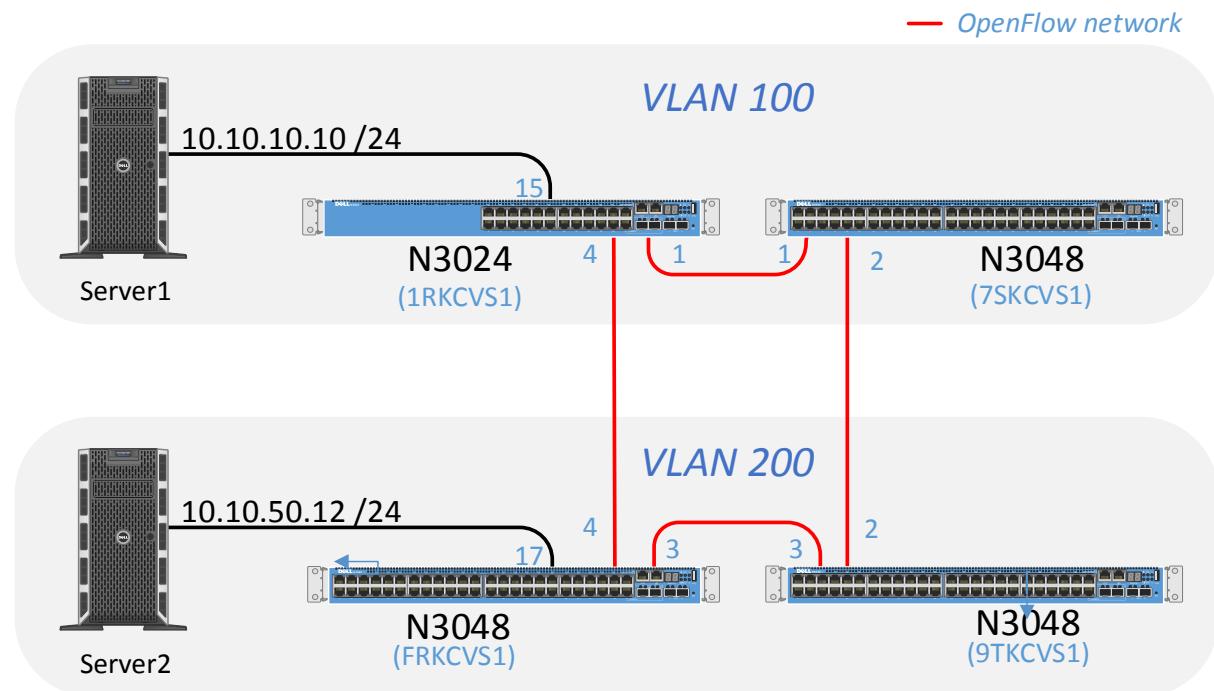


Figure 27 An L3 end-to-end flow topology with alternative route

Use the following procedure to setup and validate an L3 end-to-end flow. The setup is similar to that found in Figure 23, but only Server1 will be connected to vBR100, while Server2 will be connected to another VLAN, vBR200. The two virtual L2 bridges will then be connected together using a virtual L3 router.

Once the steps below are completed, the resulting map will look like Figure 28 .

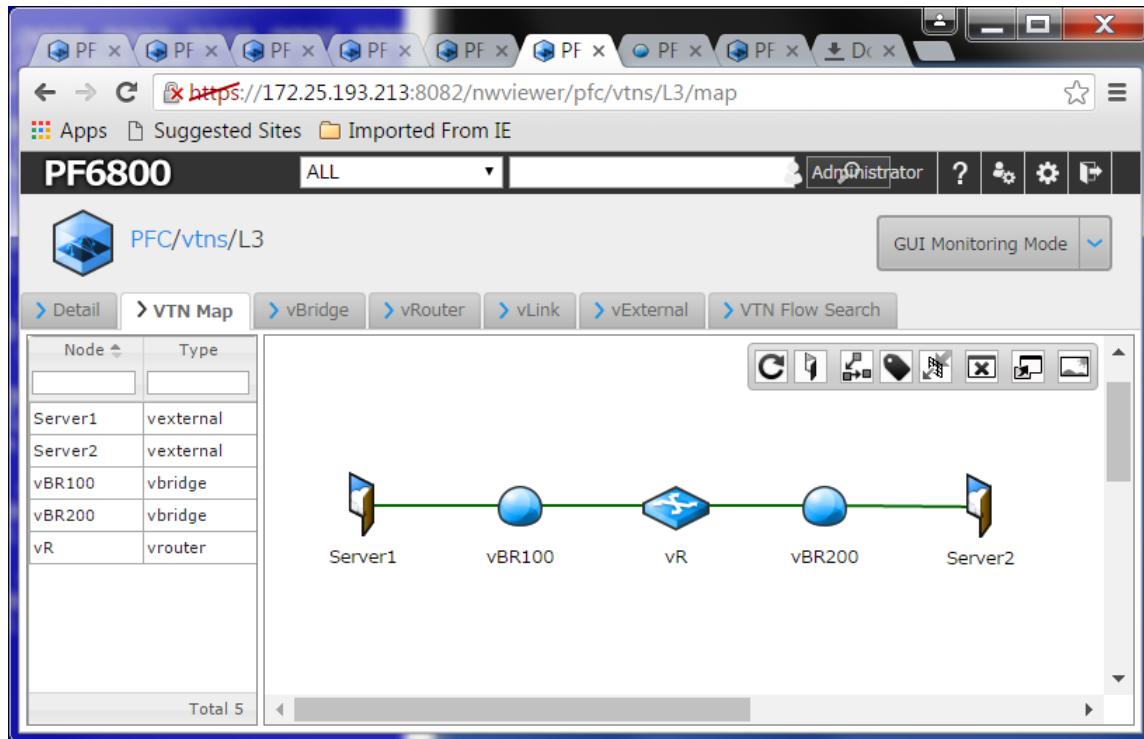


Figure 28 Routed network

Follow the steps below to setup and validate an L3 end-to-end flow on an NEC PF6800 controller. The five major steps are:

1. [VTN L3 map creation](#) – to show the logical layout of the desired network
2. [vBridge configuration \(for two VLANs\)](#) – to construct the logical switch settings on the controller
3. [vExternal \(server\) configuration](#) – to construct end point (server) settings on the controller
4. [vRouter configuration](#) – to construct the logical router settings on the controller
5. [Validation](#)- run tests to verify the setup works properly

**Note:** If using the same Servers from the L2 example, change the IP address of Server2 to 10.10.50.12 /24. Delete the L2 Map created in the L2 example to avoid errors from the controller seen when using a switch to both switch and route between the same two end points.

Set the IP addresses of Server1 to 10.10.10.10 /24 and Server2 to 10.10.50.12 /24.

## 5.2.1 VTN L3 map creation

Use the following steps to create a VTN L3 map:

1. From the controller, go to the GUI Configuration Mode, the VTN tab, and add a new VTN named “L3”. Click Apply and Commit.
2. Click the “Map” link of the newly created “L3 (Map)” on the VTN tab to open the **VTN Map** tab. The map area should display a blank canvas for creating a virtual network.
3. Click the Edit (pencil) icon in the top right corner to open the virtual network building blocks (vRouter, vBridge, and vExternal).
4. Click and drag one vRouter to the canvas and name it vR (default), which will act as the router between two virtual L2 networks.
5. Click and drag two vExternal icons to the canvas and name them Server1 and Server2, which will reflect the actual Server1 and Server2 in the example topology.
6. Click and drag two vBridge icons onto the blank canvas and name them “vBR100” and “vBR200”. These will be the two virtual L2 networks.
7. Use Cntrl-click to select Server1 and vBR100, then click the “link” button to tell the controller to create a logical link between the two. .
8. Use Cntrl-click again to select Server2 and vBR200, clicking the “link” button again to tell the controller to create a logical link between the them as well.
9. Using the same method, link the vR to vBR100, and then link vR to Vbr200.
10. Click Save.

The VTN Map should now look similar to Figure 29, showing two servers linked with two vBR (virtual bridges) and a vR (virtual router).

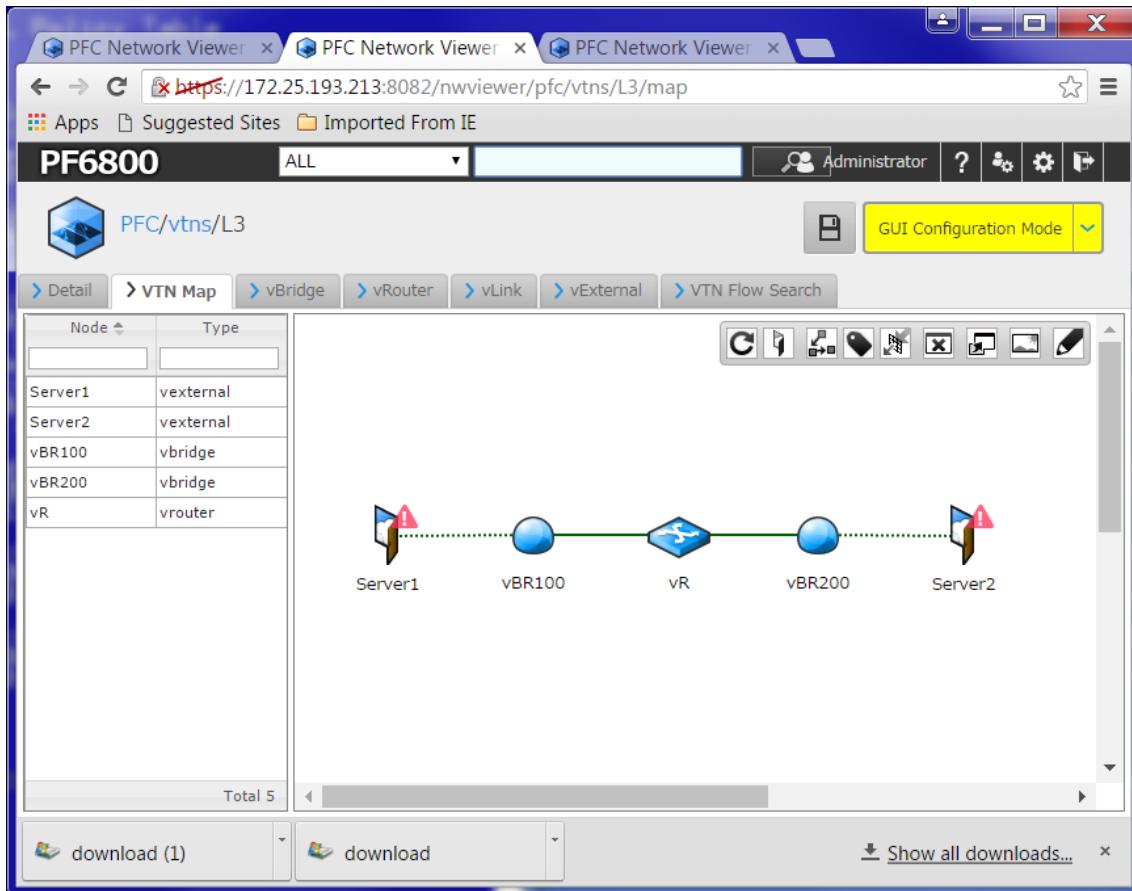


Figure 29      VTN Map

## 5.2.2 vBridge configuration (for two VLANs)

1. Double-click the **vBR100** icon on the map. Click **vBR100** again on the screen presented to open settings.
2. Go to the “VLAN Map” tab (not the VTN Map) for the vBR100 and select **tagged**, and enter **100** in the vlan ID field, as shown in Figure 30:

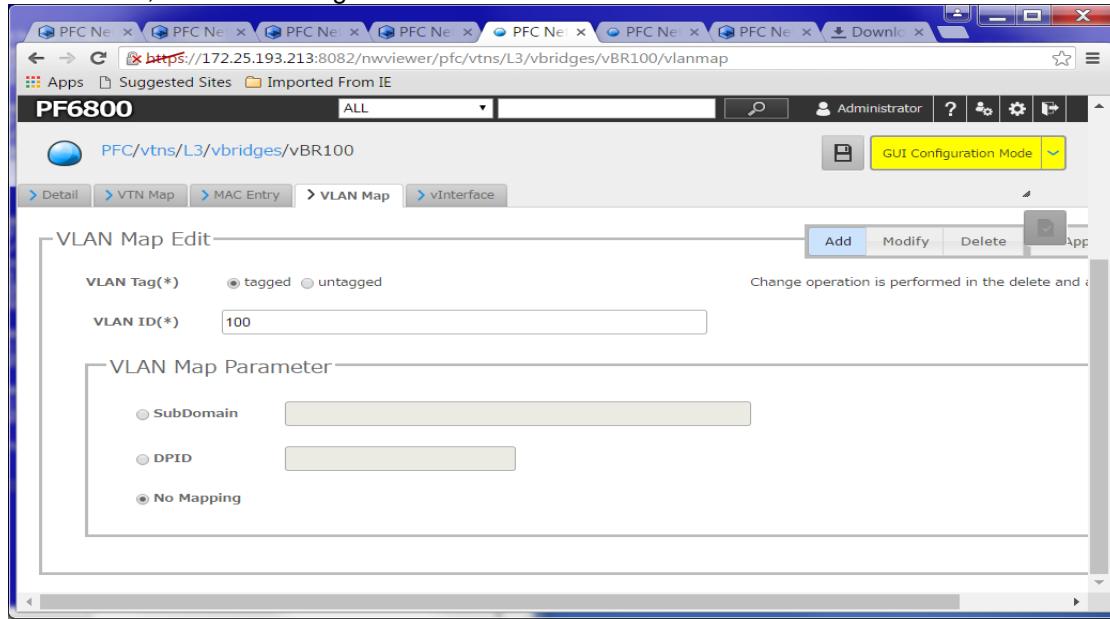


Figure 30 Setting the vBR100 VLAN

3. Set VLAN Map on vBR200 (PFC/vtns/L3/vbridges/vBR200) to VLAN 200 (tagged), as shown in Figure 31:

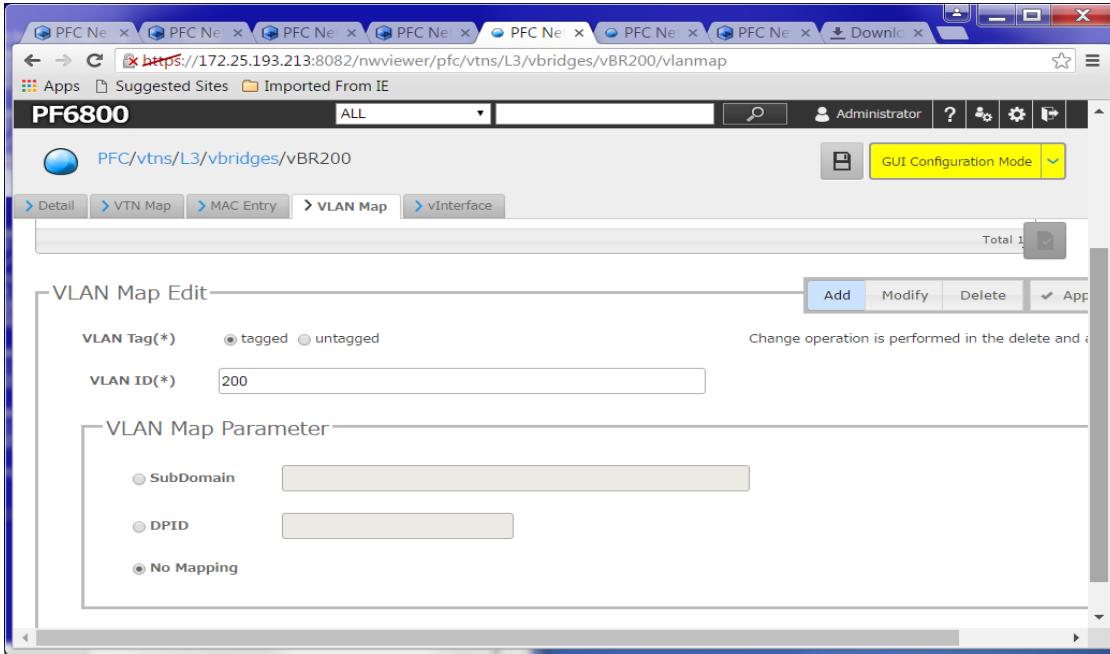


Figure 31 Setting the vBR200 VLAN

### 5.2.3 vExternal (server) configuration

1. Go back to the VTN map.
2. Select L3 > vExternal > Server1, and click on the “OFS Map” tab.
3. Enter the DPID (data path ID = 11, which is entered as hex number 0000-0000-0000-000b), port ge1/0/15, and the VLAN ID of 100 (untagged), as shown in Figure 32.

**Note:** The DPID was defined in section [3.3 - Configuring N-Series for DNOS](#) OpenFlow starting on page 13 and can also be found in the running configuration on each switch.

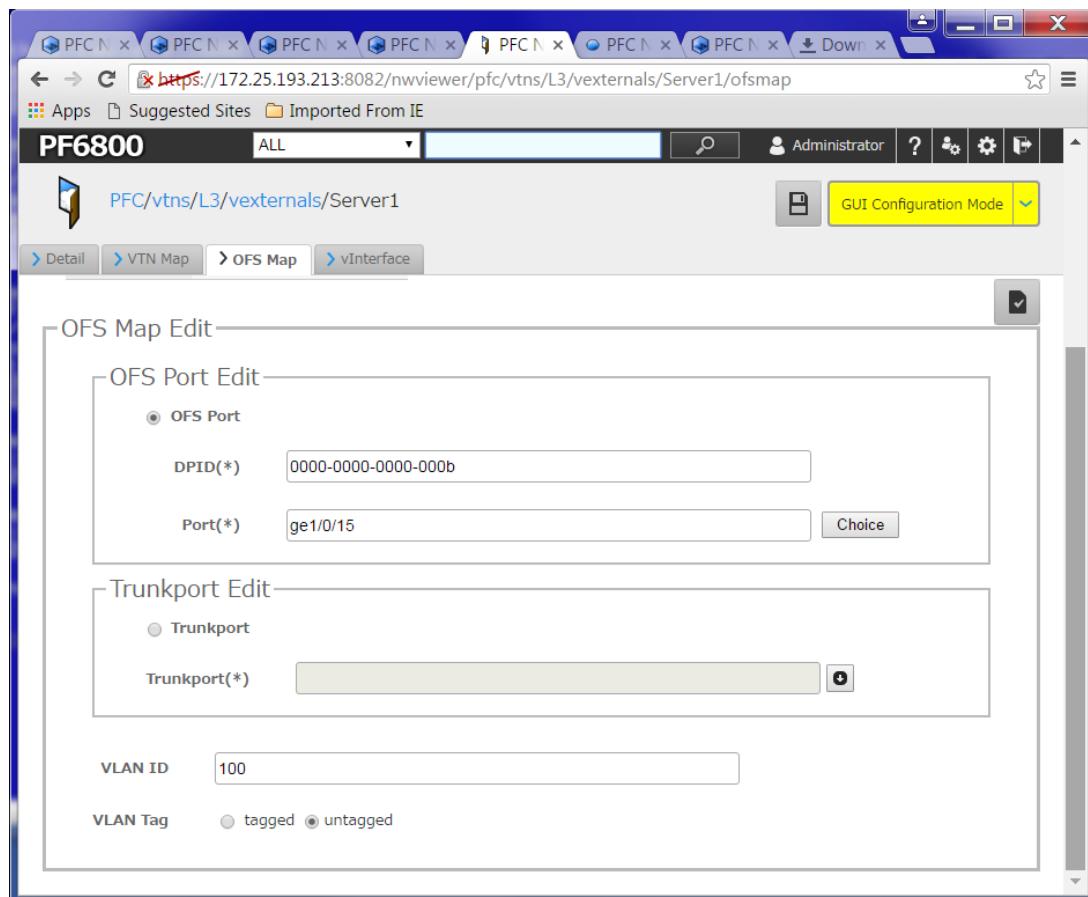


Figure 32 Server1 port settings on the OpenFlow controller

4. Click **Commit** (checkmark) on the OFS Map Edit screen to save settings.

Go to the VTN Map to see that the logical connection from Server1 to vBR100 now has a physical connection as well, represented by the solid line (see Figure 33).

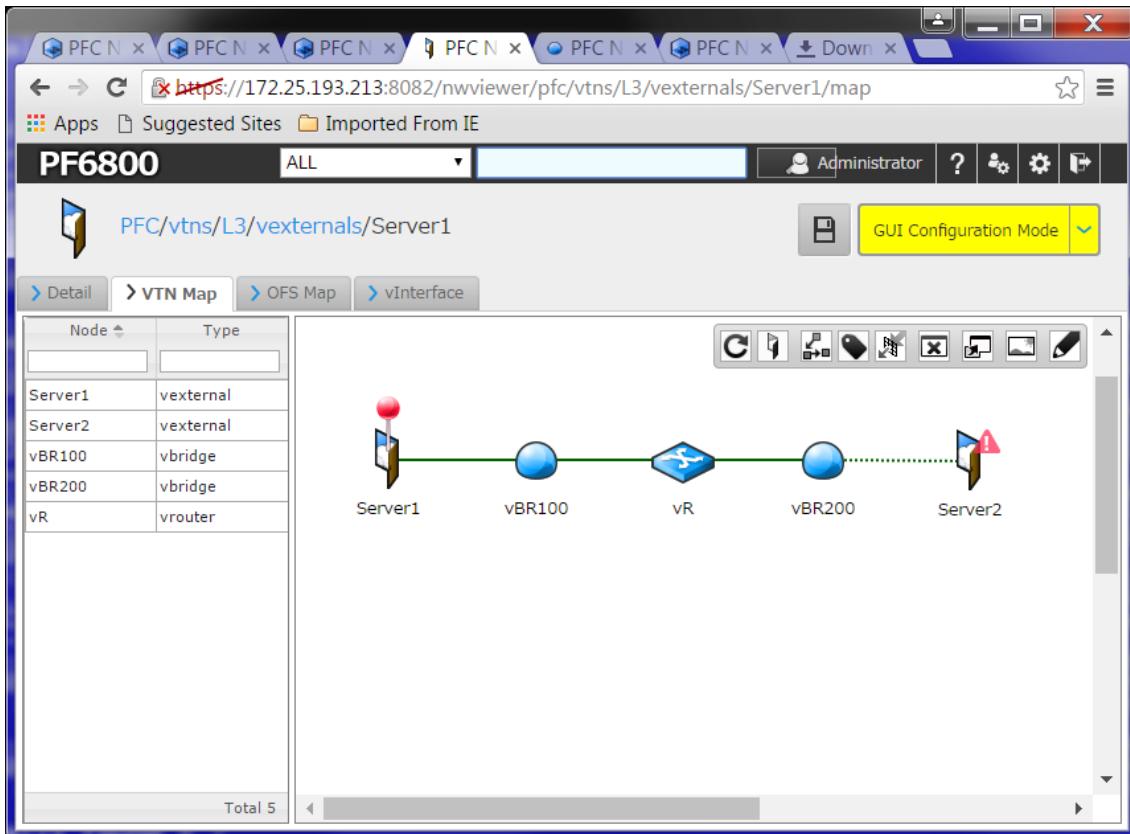


Figure 33 Server1 to vBR100 physical connection

5. Select **L3 > vexternal > Server2** and click the **OFS Map** tab.

6. Enter the data path ID (DPID) of 14 as hex number 0000-0000-0000-000e), port ge1/0/17, and the VLAN ID of 200 (untagged) as shown in Figure 34.

**Note:** The DPID was defined in section [3.3 - Configuring N-Series for DNOS](#) OpenFlow starting on page 13 and can also be found in the running configuration on each switch.

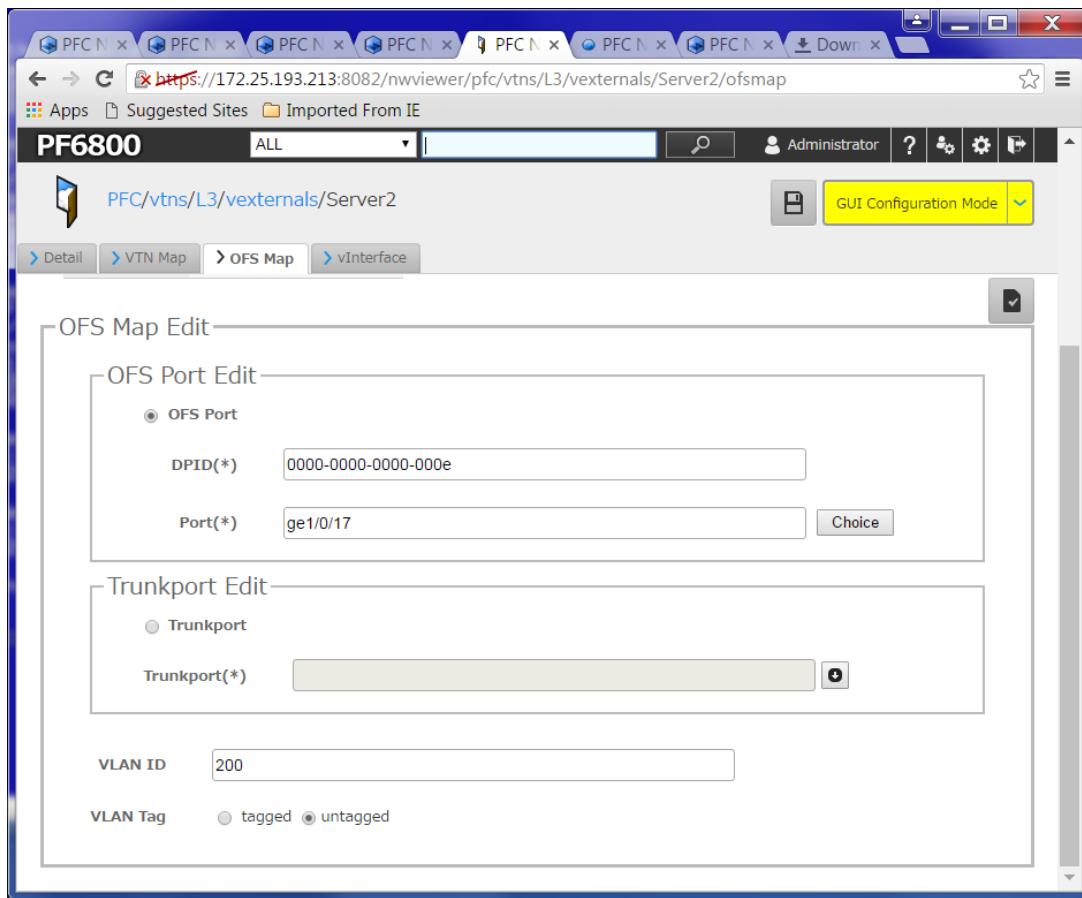


Figure 34 Server2 port settings on the OpenFlow controller

7. Click **Commit** on the OFS Map Edit screen to save settings.

Go to the VTN Map to see that the logical connections from Server1 to Server2 now have physical connections from end to end, represented by solid lines.

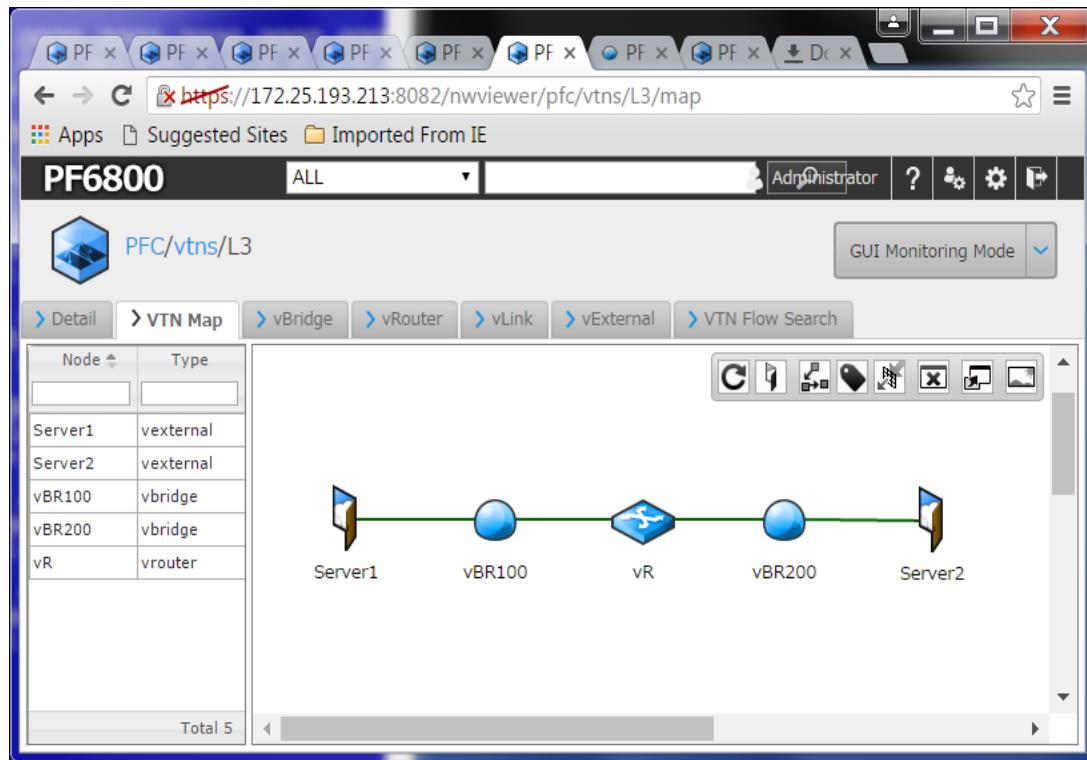


Figure 35 End to End VTN Map (vRouter not configured)

## 5.2.4 vRouter configuration

1. Go to **PFC/vtns/L3/vrouters/vR** to provide each vInterface with an IPv4 address. For Neighbor vNode going to vBR100 (VLAN 100), the IP4 address used for this example is 10.10.10.254 /24, as shown in Figure 36 .

**Note:** You must click **Modify**, then add the IPv4 address, then click **Apply** and the **Commit** checkbox to save.

The screenshot shows the Dell Networking N-Series SDN Controller's web-based management interface. The URL in the browser is <https://172.25.193.213:8082/nwviewer/pfc/vtns/L3/vrouters/vR/interfaces>. The page title is "PF6800". The top navigation bar includes tabs for "Detail", "VTN Map", "Static IP Route", "IP Route", "ARP Entry", "DHCP Relay", and "vInterface". The "vInterface" tab is selected. Below the tabs is a table titled "vInterface List" with the following data:

vInterface	Status	MAC Address	IPv4 Address	Admin	DHCP Relay Interface	Neighbor vNode/vInterface
vRIF000001	up	0225.5cca.64cd	10.10.50.254/24	up	Unregistered	vBR200/vBRIF000002
vRIF000002	up	0225.5cca.66c7	10.10.10.254/24	up	Unregistered	vBR100/vBRIF000002

A message at the bottom right of the table area says "Total 2". Below the table is a "vInterface Edit" form with the following fields:

- Interface(\*): vRIF000002
- MAC Address: 0225.5cca.66c7
- IPv4 Address/Prefix: 10.10.10.254 / 24
- DHCP Relay Interface:  Registered  Unregistered

Buttons for "Add", "Modify", "Delete", and "Commit" are visible above the edit form. A note on the right side of the edit form states: "Change operation is performed in the delete and modify mode".

Figure 36 Setting IP address for the VLAN 100 routing interface

2. For Neighbor vNode going to vBR200 (VLAN 200), the IP4 address used for this example is 10.10.50.254 /24, as shown in Figure 37.

**Note:** You must click **Modify**, then add the IPv4 address, then click **Apply** and the **Commit** checkbox to save.

The screenshot shows the PF6800 GUI interface for managing virtual interfaces. The top navigation bar includes tabs for PFC N, and Down. The URL in the address bar is https://172.25.193.213:8082/nwviewer/pfc/vtns/L3/vrouters/vR/interfaces. The main title is PF6800, and the current tab is vInterface.

**vInterface List:**

vInterface	Status	MAC Address	IPv4 Address	Admin	DHCP Relay Interface	Neighbor vNode/vInterface
vRIF000001	up	0225.5cca.64cd	10.10.50.254/24	up	Unregistered	vBR200/vBRIF0000002
vRIF000002	up	0225.5cca.66c7	10.10.10.254/24	up	Unregistered	vBR100/vBRIF0000002

**vInterface Edit:**

Form fields for editing vRIF000001:

- Interface(\*): vRIF000001
- MAC Address: 0225.5cca.64cd
- IPv4 Address/Prefix: 10.10.50.254 / 24
- DHCP Relay Interface:  Registered  Unregistered

Buttons at the top right of the edit screen: Add, Modify, Delete, Commit.

Figure 37 Setting IP address for the VLAN 200 routing interface

3. Go back to the VTN Map again to see all physical connections are complete between Server1 and Server2 as shown in Figure 38.

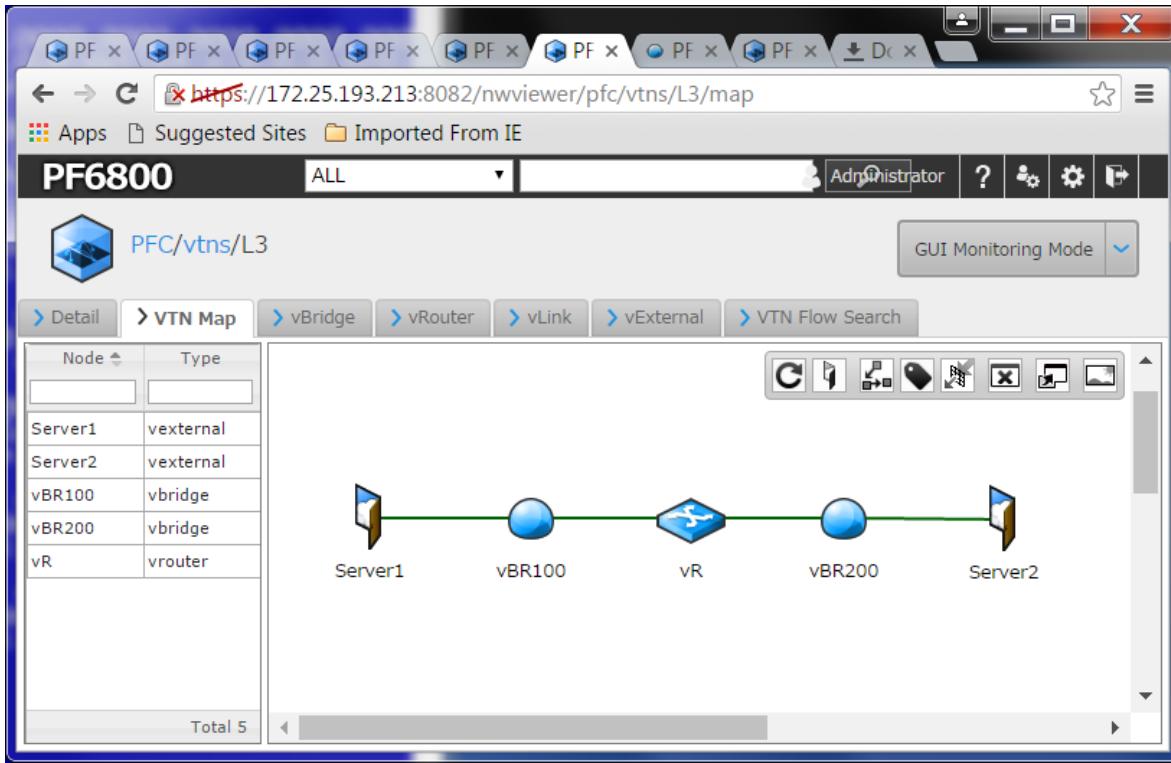


Figure 38 Routed network

- Run **show openflow flows 60** to see the addition of another flow.

```
1RKCVS1_console# show openflow flows 60
```

```
Showing OpenFlow flow entries for table 60, ACL Policy Table
ACL Policy Table (60) Flow Entries
```

```
-----
Number of entries reported = 9
Maximum number of entries for this table = 7680
| Flow ID:0xa5c7bb68 Priority:24576 Hard_time:0 Idle_time:0 Cookie:14 inPort:mask = 0:0x0
srcMac:mask = 0000.0000.0000:0000.0000.000
0 destMac:mask = 0200.0038.0000:ff1f.fff8.0000 etherType = 0000 vlanId:mask = 4094:0xffff
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.
0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4
Port = 0 Destination L4 Port = 0 IC
MP Type = 0 ICMP Code = 0 | Set output group ID = 0x ffe0004 outPort = 0
| Flow ID:0xa5c7bb68 Priority:65535 Hard_time:0 Idle_time:0 Cookie:4 inPort:mask = 0:0x0
srcMac:mask = 0000.0000.0000:0000.0000.0000
destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 88cc tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.0 srclp6
= ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Type = 0 I
CMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb68 Priority:23000 Hard_time:0 Idle_time:0 Cookie:6 inPort:mask = 1:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000.0000
```

```

00.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =
4094:0xffff srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.
0.0/0.0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00
Source L4 Port = 0 Destination L4 Port
= 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb68 Priority:22000 Hard_time:0 Idle_time:0 Cookie:7 inPort:mask = 4:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Typ
e = 0 ICMP Code = 0 | Set output group ID = 0xb000000f outPort = 0
| Flow ID:0xa5c7bb68 Priority:23000 Hard_time:0 Idle_time:0 Cookie:5 inPort:mask = 4:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000
00.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask =
4094:0xffff srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.
0.0/0.0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00
Source L4 Port = 0 Destination L4 Port
= 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb68 Priority:2 Hard_time:0 Idle_time:0 Cookie:13 inPort:mask = 15:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000
.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0 sr
clp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port =
0 Destination L4 Port = 0 ICMP Type
= 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb68 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 23:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb68 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 24:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.
srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port
= 0 Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
Number of entries actually found = 8

```

Run **show openflow groups** to see OpenFlow Groups and statistics.

```
1RKCVS1_console# show openflow groups
```

Showing the OpenFlow groups in the group tables

```

Group Type:0, Group VLAN ID:1 Group Port ID:15
GroupId=0x0001000f (L2 Interface, VLAN ID = 1, Port ID = 15): Entry Count=4
    BucketIndex=0 outputPort=15 popVlanTag=1 vlanPcp=0 vlanCfi=0 dscp=0
Group Type:0, Group VLAN ID:4094 Group Port ID:4
GroupId=0x0ffe0004 (L2 Interface, VLAN ID = 4094, Port ID = 4): Entry Count=4

```

```

BucketIndex=0 outputPort=4 popVlanTag=0 vlanPcp=0 vlanCfi=0 dscp=0
Group Type:1, Group VLAN ID:0 Group Port ID:0
GroupId=0x10000001 (L2 Rewrite, Index = 1): Entry Count=4
    BucketIndex=0 referenceGroupId = 0x0001000f vlanId = 1 srcMac: 00:00:00:00:00:00 dstMac:
BC:30:5B:F3:1B:9D
Group Type:11, Group VLAN ID:0 Group Port ID:15
GroupId=0xb000000f (L2 Unfiltered Interface, VLAN ID = 0, Index = 15): Entry Count=4
    BucketIndex=0 outputPort = 15 vlanPcp=0 vlanCfi=0 dscp=0

L2 Interface Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L2 Rewrite Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L3 Unicast Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L2 Multicast Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L2 Flood Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L3 Interface Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L3 Multicast Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L3 ECMP Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L2 Overlay Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
MPLS Label Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
MPLS Forwarding Group: Number of Entries:2 Max Group Entries:14336 Max Bucket Entries:1
L2 Unfiltered Group: Number of Group Entries:2 Max Group Entries:14336 Max Bucket Entries:1
OpenFlow Group Statistics
-----
GroupId = 0x0001000f (L2 Interface, VLAN ID = 1, Port ID = 15): Group Stats - Duration:67660,
RefCount:1
GroupId = 0x0ffe0004 (L2 Interface, VLAN ID = 4094, Port ID = 4): Group Stats - Duration:73828,
RefCount:1
GroupId = 0x10000001 (L2 Rewrite, Index = 1): Group Stats - Duration:67660, RefCount:1
GroupId = 0xb000000f (L2 Unfiltered Interface, VLAN ID = 0, Index = 15): Group Stats - Duration:135287,
RefCount:1

```

Notice where highlighted, there are no L3 Unicast Group statistics shown under OpenFlow Group Statistics. This will not be seen until after L3 traffic has been processed and routed through the network. The next section will route traffic through the OpenFlow network to ensure it is working correctly.

## 5.2.5 Validation

To validate the end-to-end connection between Server1 and Server2, simply ping between the two devices. A ping should be successful from Server1 (10.10.10.10) to Server2 (10.10.50.12), and vice versa. After completing a ping, re-run the commands below to see new entries (highlighted).

**Note:** Be sure the appropriate default gateways are set on Server1 AND Server2. These should be 10.10.10.254 and 10.10.50.254, respectively.

1. Ping from Server1 to Server2 and vice versa to verify connectivity.
20. Run show openflow flows and show openflow groups again to see the changes.

```
1RKCVS1_console# show openflow flows 60
```

Showing OpenFlow flow entries for table 60, ACL Policy Table  
ACL Policy Table (60) Flow Entries

```
-----  
Number of entries reported = 11  
Maximum number of entries for this table = 7680  
| Flow ID:0xa5c7bb68 Priority:24576 Hard_time:0 Idle_time:0 Cookie:14 inPort:mask = 0:0x0  
srcMac:mask = 0000.0000.0000:0000.0000.000  
0 destMac:mask = 0200.0038.0000:ff1f.fff8.0000 etherType = 0000 vlanId:mask = 4094:0xffff srclp4 =  
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.  
0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port =  
0 Destination L4 Port = 0 IC  
MP Type = 0 ICMP Code = 0 | Set output group ID = 0x ffe0004 outPort = 0  
| Flow ID:0xa5c7bb68 Priority:24600 Hard_time:0 Idle_time:0 Cookie:15 inPort:mask = 0:0x0  
srcMac:mask = 0000.0000.0000:0000.0000.000  
0 destMac:mask = 0200.0040.0006:ff03.ffff.ffff etherType = 0000 vlanId:mask = 4094:0xffff srclp4 =  
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.  
0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port =  
0 Destination L4 Port = 0 IC  
MP Type = 0 ICMP Code = 0 | Set output group ID = 0x10000001 outPort = 0  
| Flow ID:0xa5c7bb68 Priority:65535 Hard_time:0 Idle_time:0 Cookie:4 inPort:mask = 0:0x0  
srcMac:mask = 0000.0000.0000:0000.0000.0000  
destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 88cc tunnelId = 0 srclp4 = 0.0.0.0/0.0.0.0  
dstlp4 = 0.0.0.0/0.0.0.0 srclp6  
= ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0 Destination  
L4 Port = 0 ICMP Type = 0 I  
CMP Code = 0 | outPort = -3  
| Flow ID:0xa5c7bb68 Priority:23000 Hard_time:0 Idle_time:0 Cookie:6 inPort:mask = 1:0xffffffff  
srcMac:mask = 0000.0000.0000:0000.0000.0000  
00.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask = 4094:0xffff  
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.  
0.0.0.0 srclp6 = ::/:: dstlp6 = ::/:: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4  
Port = 0 Destination L4 Port  
= 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0  
| Flow ID:0xa5c7bb68 Priority:22000 Hard_time:0 Idle_time:0 Cookie:7 inPort:mask = 4:0xffffffff  
srcMac:mask = 0000.0000.0000:0000.0000.0000  
00.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunnelId = 0 srclp4 =  
0.0.0.0/0.0.0.0 dstlp4 = 0.0.0.0/0.0.0.
```

```

srclp6 = :::: dstlp6 = :::: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Typ
e = 0 ICMP Code = 0 | Set output group ID = 0xb000000f outPort = 0
| Flow ID:0xa5c7bb68 Priority:23000 Hard_time:0 Idle_time:0 Cookie:5 inPort:mask = 4:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000.0000 destMac:mask = 0200.0040.0000:ff03.fff8.0000 etherType = 0000 vlanId:mask = 4094:0xffff
srclp4 = 0.0.0.0/0.0.0.0 dstlp4 = 0.0.
0.0/0.0.0.0 srclp6 = :::: dstlp6 = :::: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4
Port = 0 Destination L4 Port
= 0 ICMP Type = 0 ICMP Code = 0 | outPort = 0
| Flow ID:0xa5c7bb68 Priority:2 Hard_time:0 Idle_time:0 Cookie:13 inPort:mask = 15:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunneld = 0 srclp4 =
0.0.0/0.0.0.0 dstlp4 = 0.0.0/0.0.0.0 sr
clp6 = :::: dstlp6 = :::: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Type
= 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb68 Priority:25000 Hard_time:0 Idle_time:300 Cookie:17 inPort:mask = 15:0xffffffff
srcMac:mask = bc30.5bf3.1b9d:ffff.ffff.destMac:mask = 0225.5cca.66c7:ffff.ffff.etherType = 0800 vlanId:mask = 1:0xffff srclp4 =
0.0.0/0.0.0.0 dstlp4 = 10.
10.50.12/255.255.255.255 srclp6 = :::: dstlp6 = :::: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol =
0x00 Source L4 Port = 0 Destina
tion L4 Port = 0 ICMP Type = 0 ICMP Code = 0 | Set output group ID = 0x20000001 outPort = 0
| Flow ID:0xa5c7bb68 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 23:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunneld = 0 srclp4 =
0.0.0/0.0.0.0 dstlp4 = 0.0.0/0.0.0
srclp6 = :::: dstlp6 = :::: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
| Flow ID:0xa5c7bb68 Priority:65535 Hard_time:0 Idle_time:0 Cookie:0 inPort:mask = 24:0xffffffff
srcMac:mask = 0000.0000.0000:0000.0000.0000 destMac:mask = 0000.0000.0000:0000.0000.0000 etherType = 0000 tunneld = 0 srclp4 =
0.0.0/0.0.0.0 dstlp4 = 0.0.0/0.0.0
srclp6 = :::: dstlp6 = :::: DSCP = 0 VRF = 0 DEI = 0 ECN = 0 IP Protocol = 0x00 Source L4 Port = 0
Destination L4 Port = 0 ICMP Ty
pe = 0 ICMP Code = 0 | outPort = -3
Number of entries actually found = 10

```

### 1RKCVS1\_console# show openflow groups

Showing the OpenFlow groups in the group tables

```

Group Type:0, Group VLAN ID:1 Group Port ID:15
GroupId=0x0001000f (L2 Interface, VLAN ID = 1, Port ID = 15): Entry Count=5
    BucketIndex=0 outputPort=15 popVlanTag=1 vlanPcp=0 vlanCfi=0 dscp=0
Group Type:0, Group VLAN ID:4094 Group Port ID:4
GroupId=0x0ffe0004 (L2 Interface, VLAN ID = 4094, Port ID = 4): Entry Count=5
    BucketIndex=0 outputPort=4 popVlanTag=0 vlanPcp=0 vlanCfi=0 dscp=0
Group Type:1, Group VLAN ID:0 Group Port ID:0
GroupId=0x10000001 (L2 Rewrite, Index = 1): Entry Count=5

```

```

BucketIndex=0 referenceGroupId = 0x0001000f vlanId = 1 srcMac: 00:00:00:00:00:00 dstMac:
BC:30:5B:F3:1B:9D
Group Type:2, Group VLAN ID:0 Group Port ID:0
GroupId=0x20000001 (L3 Unicast, Index = 1): Entry Count=5
    BucketIndex=0 referenceGroupId = 0x0ffe0004 vlanId = 4094 srcMac: 02:25:5C:CA:64:CD dstMac:
02:20:00:38:00:05
Group Type:11, Group VLAN ID:0 Group Port ID:15
GroupId=0xb000000f (L2 Unfiltered Interface, VLAN ID = 0, Index = 15): Entry Count=5
    BucketIndex=0 outputPort = 15 vlanPcp=0 vlanCfi=0 dscp=0

```

L2 Interface Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
L2 Rewrite Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
L3 Unicast Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
L2 Multicast Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
L2 Flood Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
L3 Interface Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
L3 Multicast Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
L3 ECMP Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
L2 Overlay Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
MPLS Label Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
MPLS Forwarding Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
Entries:1		
L2 Unfiltered Group: Number of Group Entries:2	Max Group Entries:14336	Max Bucket Entries:1
OpenFlow Group Statistics		

---

GroupId = 0x0001000f (L2 Interface, VLAN ID = 1, Port ID = 15): Group Stats - Duration:71265,	
RefCount:1	
GroupId = 0x0ffe0004 (L2 Interface, VLAN ID = 4094, Port ID = 4): Group Stats - Duration:77433,	
RefCount:2	
GroupId = 0x10000001 (L2 Rewrite, Index = 1): Group Stats - Duration:71265, RefCount:1	
<b>GroupId = 0x20000001 (L3 Unicast, Index = 1): Group Stats - Duration:1644, RefCount:1</b>	
GroupId = 0xb000000f (L2 Unfiltered Interface, VLAN ID = 0, Index = 15): Group Stats - Duration:1641,	
RefCount:1	
1RKCVS1_console#	

Notice that an L3 Unicast Group statistic has been added to the OpenFlow Group Statistics list.

You can also see flows from the NEC controller. From the VTN Map, GUI Monitoring Mode, click the “Flows” button to bring up the list of configured flows. For this example, the following two flows should be displayed: 85223 (Server1) and 85220 (Server2).

**Note:** Reactivation of flows may require a ping or other network traffic.

Click on a Flow ID as shown in Figure 39 to see direction of the flow. Clicking on Flow ID 85223 in this example shows the flow from Server1 to Server2.

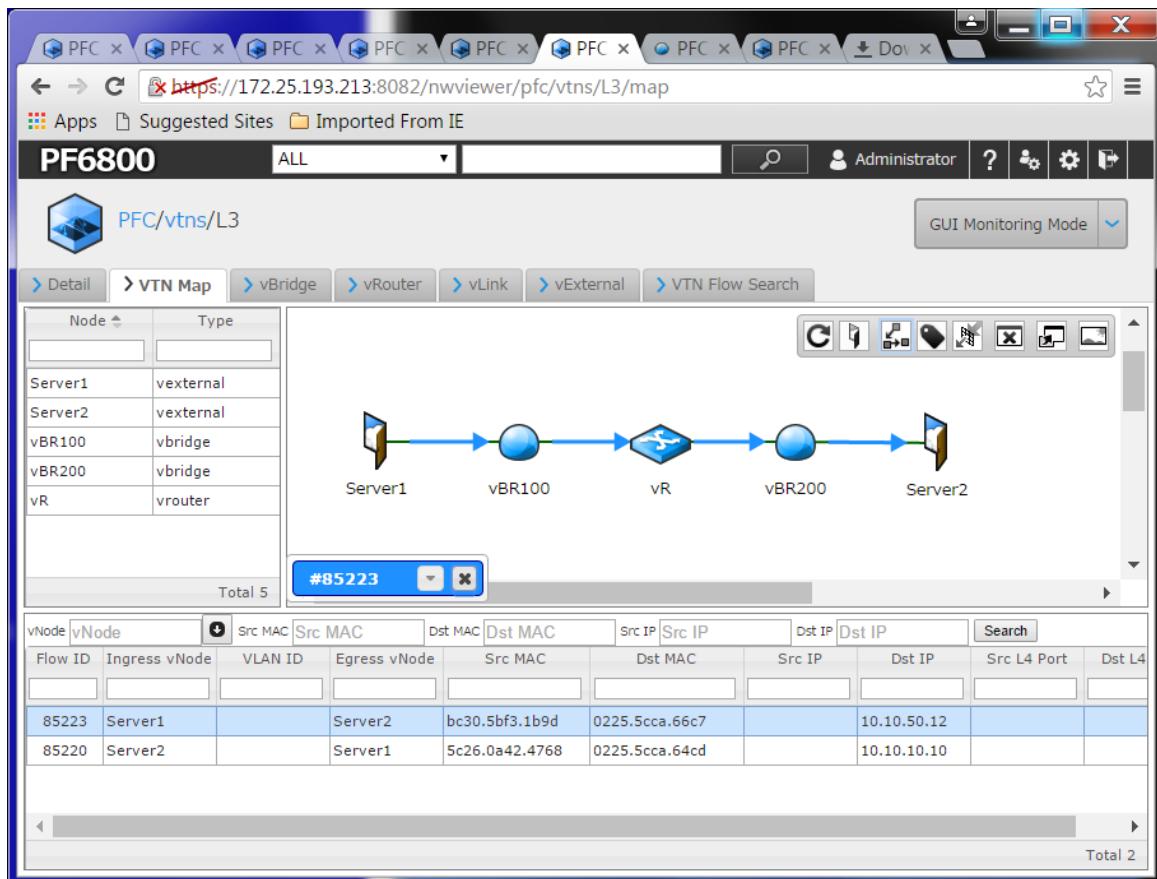


Figure 39 Flow from Server1 to Server2

As Figure 40 shows, clicking on the Flow ID 85220 in this example) shows an illustration of the flow from Server2 to Server1.

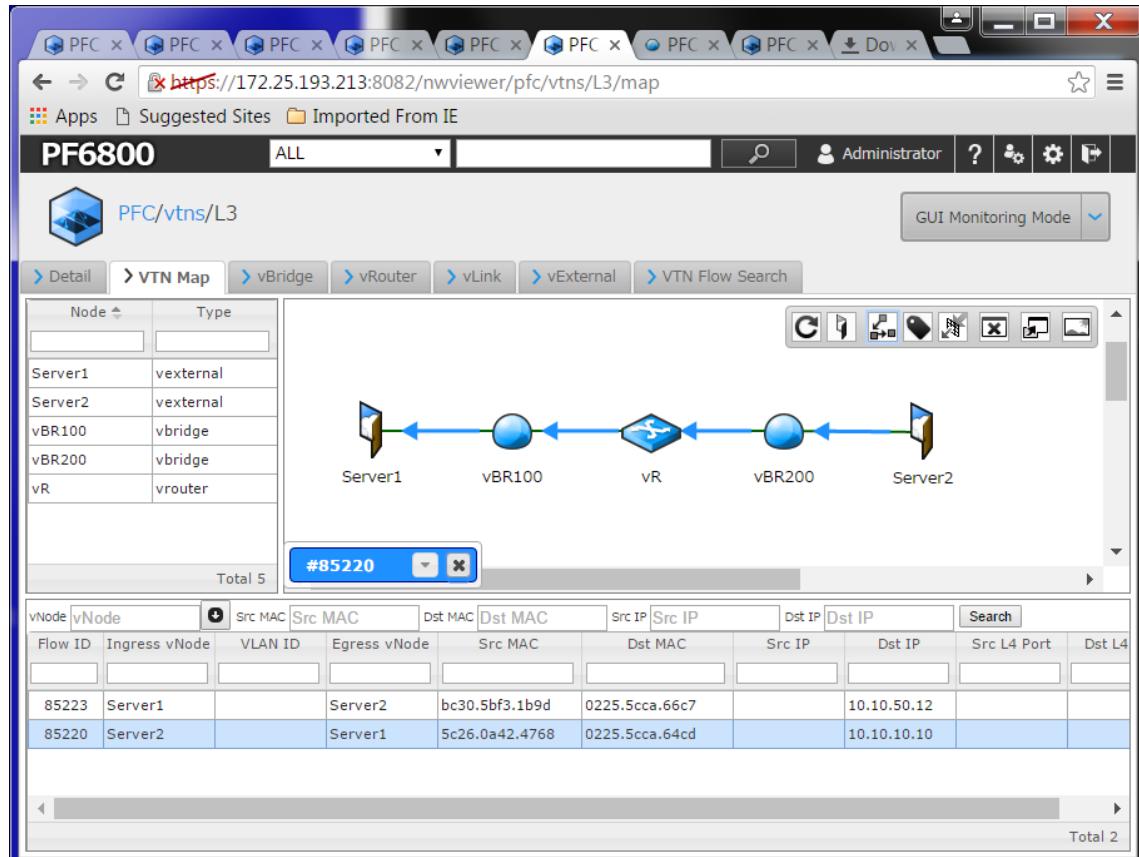


Figure 40 Flow from Server2 to Server1

## Failover (Failure Recovery)

Notice that the L3 example only uses two of the four physical switches to connect the two end points. However, the OpenFlow controller subdivides them into three logical switches including a virtual routing switch. This ability to create virtual devices allows the controller to dynamically use all OpenFlow switches (four in this scenario) in a holistic manner. This achieves proper connectivity between end points and also provides a very efficient failover system.

If equipment fails on the network, the controller moves traffic to the secondary hardware and creates new flows necessary to accomplish the mission. To demonstrate this, remove the cable connecting switches 1RKCVS1 and FRKCVS1, which currently provides the shortest path on the network to connect Server1 and Server2.

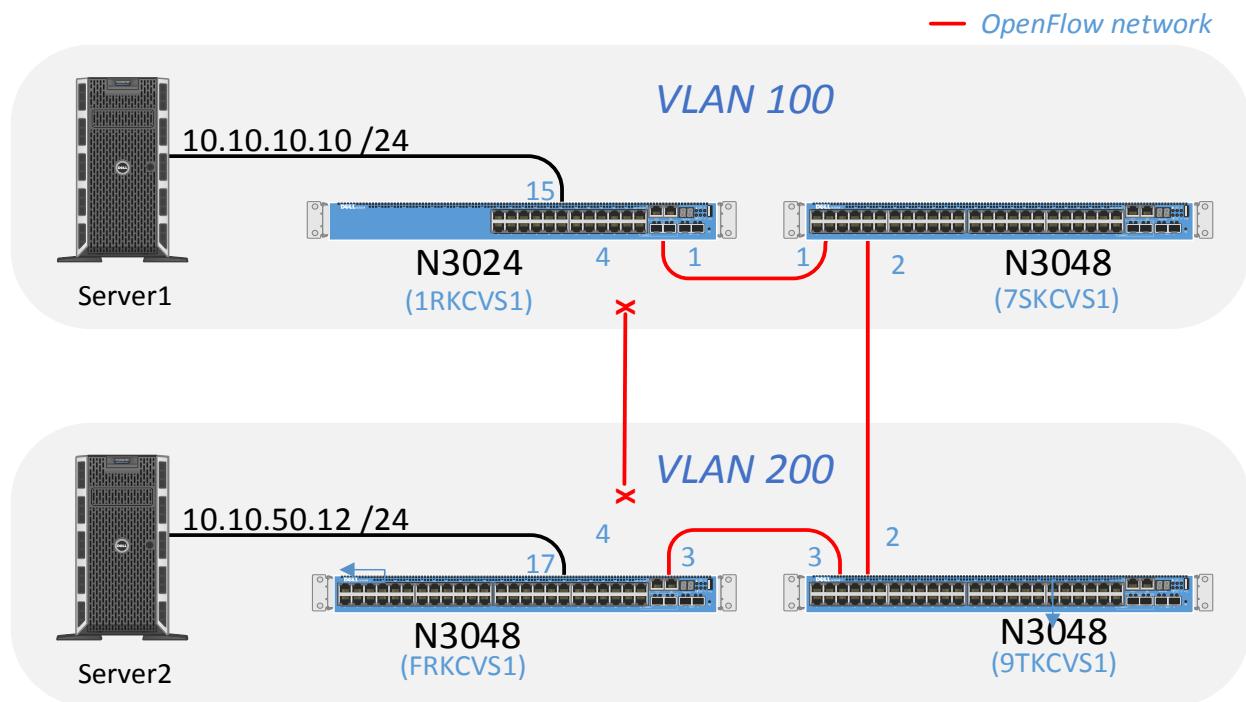


Figure 41 An L3 end-to-end flow topology with alternative route

When a link goes down, the controller immediately uses an alternative path from known physical switch connections discovered in [Chapter 4, Discovering OpenFlow Switches](#) starting on page 16. Refresh the topology map shown in Figure 42 to see that the link is now missing.

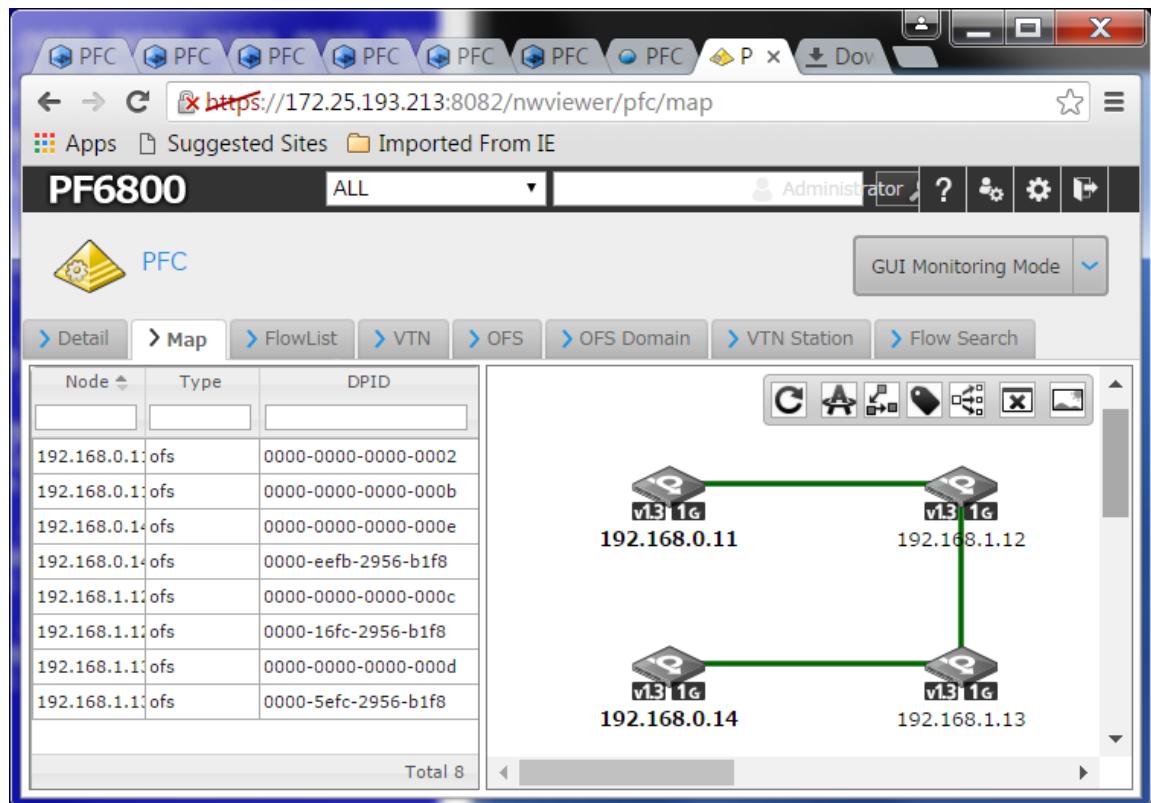


Figure 42 Physical connections between OpenFlow switches

Refreshing the VTN Map however show that the end-to-end flow between Server1 and Server2 is still up and running, as seen in Figure 43 and Figure 44.

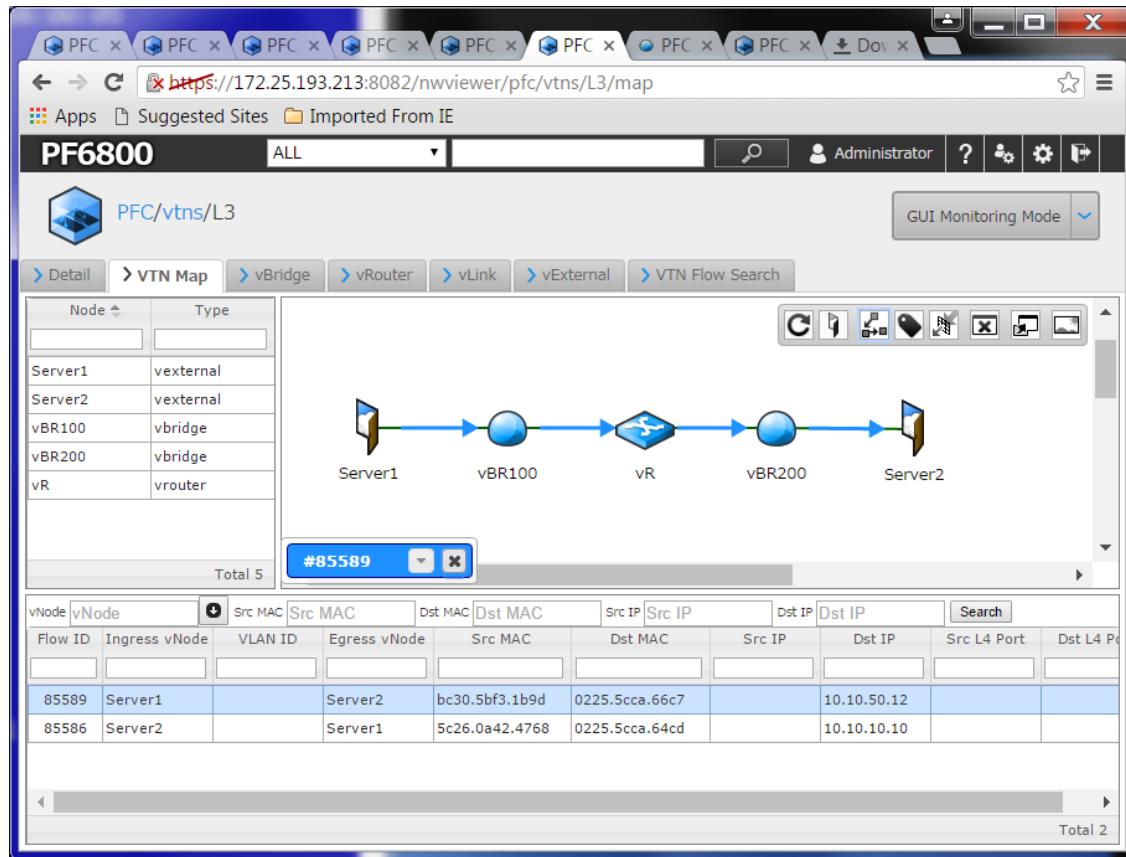


Figure 43 Flow continues after failover (Server1 to Server2)

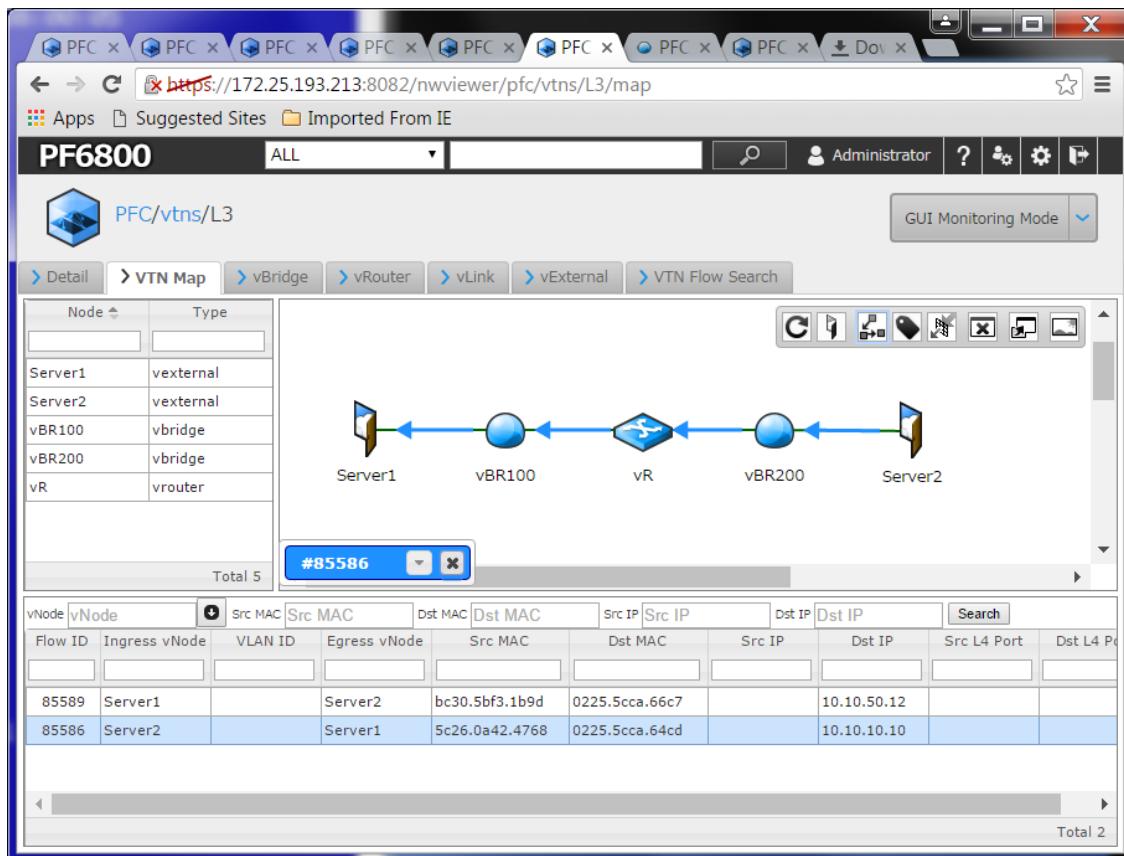


Figure 44 Flow continues after failover (Server2 to Server1)

Recovery time is efficient, missing few if any ping responses when simulating a network failure. In our testing, only one ping request timed out after pulling a cable to force a topology change (Figure 45).

Figure 45 Efficient Failover Recovery

## A

# Additional Resources

[Support.dell.com](http://Support.dell.com) is focused on meeting your needs with proven services and support.

[Dell TechCenter.com](http://Dell TechCenter.com) is an IT Community where you can connect with Dell Customers and Dell employees to share knowledge, best practices, and information about Dell products and installations.

Referenced or recommended Dell publications:

- DNOS-OF User Guide:
  - <http://www.dell.com/support/home/us/en/19/product-support/product/dell-networking-os-openflow/manuals>
- DNOS-OF v1.1 firmware is available for download from dell.com:
  - <http://www.dell.com/support/home/us/en/19/product-support/product/dell-networking-os-openFlow/drivers>
- Open Networking Foundation site
  - <https://www.opennetworking.org/>
- Dell Networking Whitepapers
  - <http://en.community.dell.com/techcenter/networking/p/guides>
- Dell Networking N-Series User Guides
  - <http://en.community.dell.com/techcenter/networking/p/guides#N-series>
- Dell Networking N-Series non-OpenFlow firmware downloads
  - [http://www.dell.com/support/home/us/en/04/Products/ser\\_stor\\_net/networking/net\\_fxd\\_prt\\_swch\\_s](http://www.dell.com/support/home/us/en/04/Products/ser_stor_net/networking/net_fxd_prt_swch_s)

## B

## Version information

This paper was compiled using the following components and versions:

Component	Version
Dell Networking N3000 OpenFlow switches	DNOS-OF firmware version 1.1
Dell Networking N3000 Traditional switches	DNOS firmware version 6.3
Dell Servers	PowerEdge R630
NEC PFC controller	PF6800 version 6.2

# C

## Glossary of Terms

**CLI (Command Line Interface):** Text-based telnet, secure shell (SSH), or serial type interface used for issuing commands to a device.

**Cluster:** Two or more devices connected in a manner that allows one of them to take over the duties of a failed device.

**Control plane:** The part of device architecture concerned with drawing network topology, or routing table information that defines correct handling of incoming information.

**Data plane:** The part of device architecture that determines handling of arriving data.

**DNOS:** Dell Networking Operating System; an operating system running on Dell Networking N-Series switches, such as release 6.3.

**DNOS-OF:** Dell Networking Operating System - OpenFlow; an operating system running on Dell Networking N-Series switches, such as release 1.1.

**DPA:** Data Plane Abstraction, enables OpenFlow performance at scale by utilizing the multiple hardware tables available in physical **switches**.

**DPID:** Data path Identifier. a 64-bit number that uniquely identifies a data path.

**L2 (Layer 2):** an OSI model layer pertaining to switching network packets based on MAC addresses

**L3 (Layer 3):** an OSI model layer that uses IP routing tables to route packets between VLANs

**MAC Address (Media Access Control address):** A MAC address is a hardware-specific address that uniquely identifies each node of a network.

**NEC:** The NEC Corporation ([www.nec.com](http://www.nec.com)) is a provider of information technology services and products.

**OF (OpenFlow):** An instance of an SDN architecture, based on a set of specifications maintained by the Open Networking Forum ([ONF](http://ONF.org)).

**OF-DPA:** Broadcom's OpenFlow Data Plan Abstraction, compliant with the Open Networking Foundation (ONF) OpenFlow specification. It defines and implements a hardware abstraction layer that maps the industry-leading StrataXGS® switch architecture to the OpenFlow™ 1.3.4 logical switch pipeline.

**OpenFlow** - an open standard enabling researchers to run experimental network protocols in campus networks.

**OOB:** Out-of-band (management). A port on a switch that allows management traffic only.

**Secure Channel:** The NEC controller channel that functions as their control plane between the OFS (OpenFlow Switch) and the PFC controller

**SSH:** Secure shell is a network protocol for operating network services securely over an unsecured network.

**SDN:** Software-defined networking allows admins to manage service through lower-level functionality.

**VTN:** Virtual Tenant Network. An isolated virtual network with distinct network policies.

## D Support and Feedback

### Contacting Technical Support

Support Contact Information

Web: <http://Support.Dell.com/>

Telephone: USA: 1-800-945-3355

### Feedback for this document

We encourage readers of this publication to provide feedback on the quality and usefulness of this deployment guide by sending an email to [Dell\\_Networking\\_Solutions@Dell.com](mailto:Dell_Networking_Solutions@Dell.com)

## About Dell EMC

Dell EMC is a worldwide leader in data center and campus solutions, which includes the manufacturing and distribution of servers, network switches, storage devices, personal computers, and related hardware and software. For more information on these and other products, please visit the Dell EMC website at <http://www.dell.com>.