Pluribus NETWORKS

Netvisor® ONE

Data Collection and Troubleshooting
Practical Guide
Printed in the U.S.A.
10.10.2017

Document Feedback
Feedback about this document may be sent to doc-feedback@pluribusnetworks.com

Table of Contents

Introduction

In many cases problems occur that do not directly point to a particular failure or misconfiguration. In those cases many pieces of data must be collected to begin the troubleshooting process. First, obvious problems must be ruled out. Then, more detailed troubleshooting data can be gathered to start honing-in on the specific problem.

It is critical to understand and document the physical and logical topology of the environment - the physical connections (cables and ports) between switches and the switches purpose in the overall design will contribute greatly to the troubleshooting process.

This guide will focus on many commands that provide status and configuration data of the switch.

Switch Shell vs CLI Access

Netvisor provides the following types of Administrator access for Switch Management:

- **Netvisor CLI**
  - Provides Configuration and Troubleshooting capabilities
  - To Login via CLI, use the network-admin account and customer can create other user accounts with CLI Login Privileges

```
network-admin@10.9.20.201
Password:
Last login: Tue Mar 15 15:10:16 2016
Netvisor OS Command Line Interface 2.3
Connected to Switch spine1; nvOS Identifier:0xc0001c3; Ver: 2.3.203038860
CLI (network-admin@munich-spine1) >
```

- **Netvisor Shell**
  - Shell based Switch access is by defaulted restricted with Challenge/Response to avoid any undesired changes that can affect netvisor's critical internal files/directories.
  - Please contact Pluribus Support for Challenge-Response details.
  - Access to Shell with Challenge-Response method.

```
MacBook-Pro:~ $ ssh admin@10.9.20.201
Password:
Challenge: 6B9X-E2TY-CVFG-2
Please contact pnsupport@pluribusnetworks.com to procure response against the challenge. The
response key times out in 3 minutes and needs to be timely entered.
Enter response:
admin@spine1:~$
```

CLI Scope Setting

Most configuration commands require the "scope" setting. There are two options for this setting: local and fabric. A command executed with the local scope does not propagate to the other members of the fabric while a command executed with the fabric scope does.

Common Symptoms

The symptoms of a network problem may vary depending on the volume and type of network traffic using the network. However, the following are common symptoms that generally accompany network problems.

- Slow or unresponsive applications (servers)
- Loss of connectivity to the affected network
- Route flapping because of VM/MAC moves
- High link utilization or high packet drop activities

Information Required for Pluribus TAC Case

To open a support case, the following information's are required:

- Customer/Account Name
- Appliance Serial Number
- Appliance Software version
- Problem/Issue Description in detail with any error message or screenshots.
- Is this issue impacting your production Network - Yes/No
- Appliance Logs
    - System, Event & Audit log files - /nvOS/log directory
    - nvOSd.log* - /var/nvOS/log/ directory
    - Core files (If any) - /var/nvOS/log/cores directory
    - vport tables - /var/nvOS/etc/vport/ directory

Procedure to collect the appliance logs:

1. Login to the appliance shell
   admin@<applianceIP>

2. Change the current directory to /var/tmp/
   cd /var/tmp/

3. Run this command to create a log snapshot (without the core files)
   tar -czcf logsnapshot.log.tar.gz /nvOS/log /var/nvOS/ /var/log/ --exclude='/var/nvOS/log/cores/*'

4. Copy the created Log Snapshot file <logsnapshot.log.tar.gz> from the appliance using scp command or scp application like winscp.. To connect the appliance using SCP, we need to enable SFTP via CLI cmd: **admin-sftp-modify enable**

Troubleshooting Process Overview

This process can of course be adapted for a given situation but this template can serve as a roadmap for troubleshooting many different types of problems. This process starts with the underlying hardware and builds logically from there to identify problems or eliminate potential problems as the troubleshooting proceeds.

Hardware Troubleshooting

Determine if the switch hardware is having any problems. Gather status data to determine the overall health of the underlying hardware. A systemic problem in the platform may present itself as a network communication or configuration error.

L1 Troubleshooting

Verify that the ports, cables, and transceivers are functioning properly. Similar to the hardware troubleshooting, a problem with the physical layer may present itself as a network communication or configuration error.

Foundational Object Validation

Verify that the underlying foundational objects (fabric, trunk, cluster, and vlag) are functioning properly. These objects provide basic communication through the fabric and redundancy mechanisms for data path failures. Failures in these objects should be easy to troubleshoot but ensure these objects are stable and functioning before continuing with the troubleshooting process.

L2 Troubleshooting

Assuming the prior steps have not solved the problem, begin L2 troubleshooting. This should include pre-requisite information about the problem such as; any user reported symptoms, IP addresses of the systems involved, switch ports of the systems involved, and VLAN membership of the systems involved.

L3 Troubleshooting

Assuming the prior steps have not solved the problem, begin L3 troubleshooting. This should include pre-requisite information about the problem such as; any user reported symptoms, network topology including routed segments, IP addresses of the systems involved, and network segments of the systems involved.

Other Tools and Configuration

There are several other commands that provide access to other tools and configuration information useful in troubleshooting.

Hardware Troubleshooting

Frist confirm the overall health of the switch itself. The following commands validate basic communications and overall health of the device.

Verify Management Port Settings

The management interface must be configured correctly for CLI access. If the management interface is not functional check its configuration using the serial port or VGA/USB ports to confirm the settings shown below are properly configured for your management network.

Confirm the management interface is configured properly with the command:
`switch-setup-show`

```
CLI (network-admin@pnswitch1) > switch-setup-show
switch-name:         pnswitch1
mgmt-ip:             10.9.7.30/16
mgmt-ip6:            fe80::f68e:38ff:fe06:8416/64
mgmt-link-state:     up
mgmt-link-speed:     1g
in-band-ip:          192.168.1.50/24
gateway-ip:          10.9.9.1
gateway-ip6:         10.9.9.1
dns-ip:              10.20.41.1
dns-secondary-ip:    8.8.8.8
domain-name:         pluribusnetworks.com
ntp-server:          0.us.pool.ntp.org
timezone:            America/New_York
date:                2016-11-16,15:23:16
phone-home:          yes
hostid:              184549751
analytics-store:     default
enable-host-ports:   yes
device-id:           D2SYX42
```

Verify Management Services

Access to services through the management interface can be enabled/disabled as needed for your environment. However, if a service is misconfigured it may appear that something has failed.

Confirm the management interface services are configured properly with the command: `admin-service-show`

```
CLI (network-admin@pnswitch1) > admin-service-show
switch       if   ssh nfs web web-ssl web-ssl-port web-port snmp net-api icmp
------------ ---- --- --- --- ------- ------------ -------- ---- ------- ----
pnswitch1 mgmt on  on  on  off     443          80       off  off     on
pnswitch1 data on  on  off off     443          80       off  off     on
```

Validate the Hardware

It is usually a good idea to eliminate hardware problems first. Confirm the hardware is in good health using the following commands.

Confirm components are functioning properly with the command: `switch-info-show`

```
CLI (network-admin@pnswitch1) > switch-info-show
switch:          pnswitch1
model:           S4048-ON
chassis-serial:  1626PN8500046
system-mem:      3.84G
switch-device:   OK
```

```
fan1-status:     OK
fan2-status:     OK
fan3-status:     OK
fan4-status:     OK
fan5-status:     OK
fan6-status:     OK
ps1-status:      OK
ps2-status:      FAULT
```

Confirm power and cooling components are functioning properly with the command: `switch-status-show`

```
CLI (network-admin@pnswitch1) > switch-status-show
switch       name        value units     state
------------ ----------- ----- --------- -----
pnswitch1 FAN1        9980  rpm       ok
pnswitch1 FAN2        10147 rpm       ok
pnswitch1 FAN3        9942  rpm       ok
pnswitch1 FAN4        9929  rpm       ok
pnswitch1 FAN5        10082 rpm       ok
pnswitch1 FAN6        10134 rpm       ok
pnswitch1 CPU Temp    39    degrees-C ok
pnswitch1 Switch Temp 42    degrees-C ok
pnswitch1 Sfpp Temp   37    degrees-C ok
pnswitch1 Qsfp Temp   34    degrees-C ok
```

NOTE: The state column should report “ok” for all objects.

Validate the Software

Confirm the switch is running the latest version of the Netvisor operating environment. The `phone-home` setting in the `switch-setup-show` command must be set to “yes” for this command to report as shown in the sample output.

Confirm the software release with the command: `software-show`

```
CLI (network-admin@pnswitch1) > software-show
version:           2.6.0-2011941
track:             2.6-release
use-proxy:         no
```

To upgrade your software, use the following syntax:

```
CLI (network-admin@Spine1) >software-upgrade package nvOS-X.X.X-XXX.platform.pkg
```

Customers can download the ONVL Upgrade package from Pluribus Cloud

L1 Troubleshooting

Assuming the minimal configuration settings are correct and the major components of the switch are functioning properly, then proceed to the next steps in troubleshooting. Verify the underlying physical ports, transceivers, and cables are functioning properly before validating other configuration settings. Nothing works if the underlying hardware is not working.

Port Status

Confirm the ports are functioning properly with the command: `port-show`

```
CLI (network-admin@pnswitch1) > port-show format switch,port,status

switch       port status

------------ ---- ----------------------------

pnswitch1 0    up,PN-internal,stp-edge-port

pnswitch1 0    up,PN-internal

pnswitch1 0    up,PN-internal

pnswitch1 1    up,vlan-up

pnswitch1 5    up,vlan-up
```

Bezel-Port Mapping

To display the port numbering on ONVL Platforms, use the bezel-portmap-show command:

```
CLI (network-admin@ pnswitch1) > bezel-portmap-show

switch       port bezel-intf

------------ ---- ----------

tac-dell-sw1 1    1

tac-dell-sw1 2    2

tac-dell-sw1 3    3

tac-dell-sw1 4    4

tac-dell-sw1 5    5

….

tac-dell-sw1 50   49.2

tac-dell-sw1 51   49.3

tac-dell-sw1 52   49.4

tac-dell-sw1 53   50

tac-dell-sw1 54   50.2

tac-dell-sw1 55   50.3

switch       port bezel-intf

------------ ---- ----------

tac-dell-sw1 56   50.4
```

```
tac-dell-sw1 57   51
tac-dell-sw1 58   51.2
tac-dell-sw1 59   51.3
tac-dell-sw1 60   51.4
tac-dell-sw1 61   52
tac-dell-sw1 62   52.2
tac-dell-sw1 63   52.3
tac-dell-sw1 64   52.4
tac-dell-sw1 65   53
tac-dell-sw1 66   53.2
tac-dell-sw1 67   53.3
tac-dell-sw1 68   53.4
tac-dell-sw1 69   54
tac-dell-sw1 70   54.2
tac-dell-sw1 71   54.3
tac-dell-sw1 72   54.4
```

Transceiver Status

Validate the transceivers are supported with the command: port-xcvr-show

```
CLI (network-admin@pnswitch1) > port-xcvr-show
switch       port vendor-name      part-number      serial-number
------------ ---- ---------------- ---------------- -------------
pnswitch1 1    3M               1410-P17-00-2.00 Y20B210407
pnswitch1 5    3M               1410-P17-00-2.00 Y20B210834
CLI (network-admin@pnswitch1) >
```

Transceiver Link Quality

Validate the link quality and port speed with the command: **port-phy-show**

```
CLI (network-admin@ pnswitch1) > port-phy-show

switch      port state speed eth-mode   max-frame link-quality  learning def-vlan

----------- ---- ----- ----- ---------- --------- ------------- -------- --------

pnswitch1 14   up    10000 10Gbase-cr 12280     good (51/39)  on       4092

pnswitch1 16   up    10000 10Gbase-cr 1540      great (57/42) on       4091

pnswitch1 41   up    10000 10Gbase-cr 1540      good (32/43)  on       33

pnswitch1 51   up    10000 10Gbase-cr 12280     great (58/40) off      1

pnswitch1 52   up    10000 10Gbase-cr 12280     good (58/39)  off      1

pnswitch1 65   up    10000 10Gbase-cr 12280     good (35/44)  on       1

pnswitch1 66   up    10000 10Gbase-cr 12280     good (39/43)  on       1

pnswitch1 69   down  10000 10Gbase-cr 12280     great (64/52) off      1

pnswitch1 70   down  10000 10Gbase-cr 12280     great (64/49) off      1

pnswitch1 71   down  10000 10Gbase-cr 12280     great (64/52) off      1

pnswitch1 72   down  10000 10Gbase-cr 12280     great (64/53) off      1

tac-f64-sw4 11   down  10000 10Gbase-cr 1540      none (-1/-1)  on       1

tac-f64-sw4 12   up    10000 10Gbase-cr 1540      good (36/40)  on       4090

tac-f64-sw4 14   down  10000 10Gbase-cr 1540      none (-1/-1)  on       1
```

Note: Anything other than good or great link quality should be reason to investigate the cable/transceivers.

Foundational Object Validation

Pluribus Fabric

The fabric is an essential object for switch operations. when you add switches to the fabric, all switches are under a single management domain which is highly available through multiple link aggregation and load balancing between network resources. The fabric performs a classic database 3-phase commit for configuration changes. All members of the fabric must accept the configuration changes before the change is made in the fabric.

We highly recommend to do in-band fabric is because of redundancy when compared to Mgmt.
Important things to consider for a Healthy Fabric:

- All fabric nodes must be able to reach each other.
- Pluribus fabric traffic consumes some bandwidth and in-band based setup, the fabric traffic is prioritized over other switch traffic.
- We have no control over traffic prioritization on MGMT Network (customer provided),

Fabric Network can be modified using the command: fabric-local-modify

```
CLI (network-admin@pnswitch1) > fabric-local-modify
vlan                          VLAN assigned to fabric
fabric-network                fabric administration network
control-network               control plane network
fabric-advertisement-network network to send fabric advertisements on
```

Fabric Status

Confirm the fabric is in good health before continuing to other configured objects.

Verify basic fabric information with the command: `fabric-show, fabric-info`

```
CLI (network-admin@pnswitch1) > fabric-show
name              id                vlan fabric-network control-network tid  fabric-advertisement-network
----------------- ----------------- ---- -------------- --------------- ---- ----------------------------
Dell1             b000177:581bb720 0    in-band        in-band         7    inband-mgmt

CLI (network-admin@pnswitch1) > fabric-info
name:                         Dell1
id:                           b000177:581bb720
vlan:                         0
fabric-network:               in-band
control-network:              in-band
tid:                          7
fabric-advertisement-network: inband-mgmt
```

Verify fabric node state with the command: `fabric-node-show`

```
CLI (network-admin@pnswitch1) > fabric-node-show format name,fab-name,fab-tid,state,device-state,
name         fab-name fab-tid state  device-state
------------ -------- ------- ------ ------------
pnswitch1 Dell1    7       online ok
pnswitch2 Dell1    7       online ok
```

The state represents communication status between members of the fabric and the device status represents the overall health of the switch.

Trunk Status

A trunk (link aggregation group or LAG) can be configured automatically or defined manually. They are used for inter-switch communication (auto-LAG) or general network connectivity (manual LAG). If configured, they provide a critical communication path.

Verify trunk (LAG) status with the command: `trunk-show`

```
CLI (network-admin@pnswitch1) > trunk-show format switch,name,ports,speed,lacp-mode,status
switch    name                ports speed   lacp-mode status
--------- ------------------- ----- ------- --------- ------
pnswitch1 ports1-4            1-4   10g     off
pnswitch1 ports5-8            5-8   10g     off
pnswitch1 ports9-12           9-12  10g     off
pnswitch1 ports13-16          13-16 10g     off
```

Trunks can be configured with or without LACP. The following example shows the Trunk based LACP options.

```
CLI (network-admin@pnswitch1) > trunk-create name port1-4 ports 1,4 lacp-mode
       off          LACP is off
       passive      LACP passive mode
       active       LACP active mode
```

Cluster Status

The cluster and VLAG objects provide the underlying redundancy structure for network communications. If the network design calls for redundancy check that the cluster and VLAG objects are functioning properly.

Verify cluster status with the command: `cluster-show`

```
CLI (network-admin@pnswitch1) > cluster-show
name              state  cluster-node-1 cluster-node-2 tid ports    remote-ports
----------------- ------ -------------- -------------- --- -------- ------------
pnclusterodd      online pnswitch1      pnswitch3      15  4,36,128 4,36,129
pnclustereven     online pnswitch2      pnswitch4      0   4,8,128  4,8,129
CLI (network-admin@pnswitch1) >
```

Cluster communications is dependent on a direct physical link(s) between two switches. For the cluster to function properly that physical link must be functioning.

VLAG Status

Verify VLAG status with the command: `vlag-show`

```
CLI (network-admin@pnswitch1) > vlag-show
name                    cluster       mode          switch    port          peer-switch peer-port     statu
local-state lacp-mode
----------------------- ------------  ------------- --------- ------------- ----------- ------------- -----
--------- ---------
pnvlag1                 pnclusterodd  active-active pnswitch1 trunk-to-plus pnswitch3   trunk-to-plus norma
enabled,up  off
pnvlag2                 pnclustereven active-active pnswitch2 49            pnswitch4   18            norma
enabled,up  active

CLI (network-admin@pnswitch1) >
```

The VLAG relies on the underlying cluster. Confirm the VLAG status is normal and the state is “enabled,up”. If there are problems with a VLAG, work back through the objects it depends on – the cluster, and ultimately physical ports and cables.

L2 Troubleshooting

VLAN configuration provides network isolation beyond the physical ports. By default, when VLANs are created they are assigned to all physical ports. This default behavior can be changed when creating a VLAN (it can be assigned to no ports or a subset of ports). Also, all the physical ports of the switch are in “trunk” mode, meaning they are expecting packets to be tagged (802.1Q) and they can process traffic from any VLAN (tagged packet).

If traffic is not getting through to its destination confirm the VLAN configuration.

Spanning Tree

To build a loop-free topology, switches (“bridges”) have to determine the root bridge and compute the port roles, root, designated, or blocked. The use of RSTP is recommended for ad hoc networks that interoperate in a heterogeneous, multi-vendor switch environment.

RSTP is enabled on the switch by default on Pluribus Switches.

```
CLI (network-admin@pnswitch1*) > stp-show
enable:              no
stp-mode:            rstp
bpdus-bridge-ports: yes
bridge-id:           64:0e:94:28:0a:08
bridge-priority:     32768
hello-time:          2
forwarding-delay:    15
max-age:             20
mst-max-hops:        20
mst-config-name:     Pluribus
mst-config-digest:
cluster-mode:        slave
```

To display the STP state, use the following command: stp-state-show

```
CLI (network-admin@pnswitch1*) > stp-state-show
vlan:             1-4093,4095
ports:            none
root-port(peer):  0
disabled:         1-13,15,17-31,33-50,53-64,67-72
learning:         none
forwarding:       14,16,32,65-66,128
discarding:       none
edge:             65-72
vlan:             4094
ports:            none
root-port(peer):  0
disabled:         67-72
learning:         none
forwarding:       51-52,65-66,128
discarding:       none
edge:             65-72
```

To display the information about STP on ports, use the following command: stp-port-show

```
CLI (network-admin@pnswitch1*) > stp-port-show
port block filter edge bpdu-guard root-guard priority cost
---- ----- ------ ---- ---------- ---------- -------- -----
32   off   off    no   no         no         128      20000
128  off   off    no   no         no         128      1000
```

VLAN Status

Verify VLAN status with the command: `vlan-show`

```
CLI (network-admin@pnswitch1) > vlan-show id 34
id scope  name    active stats ports                        untagged-ports active-edge-ports
-- ------ ------- ------ ----- ---------------------------- -------------- -----------------
34 fabric vlan-34 yes    yes   49-51,54,56-57,61-72,130,255 none           none
CLI (network-admin@pnswitch1) >
```

This command will show the relationship between the VLAN and the ports that allow that VLANs traffic. If traffic is not getting through to its destination confirm the VLAN configuration.

Port and VLAN Relationship

Verify port and VLAN relationship with the command: `port-show vlan ID`

```
CLI (network-admin@pnswitch1) > port-show vlan 34
port ip          mac               vlan vxlan hostname  status         config
---- ----------- ----------------- ---- ----- --------- -------------- ------
65   172.16.34.5 66:0e:94:04:7f:91 34   0     pnswitch1 up,PN-internal
65   172.16.34.1 00:00:5e:00:01:c9 34   0     pnswitch1 up,PN-internal
```

Similar to the VLAN status command, this command will show the relationship between the VLAN and the ports that allow that VLAN traffic.

VLAN Assignment per Port

Verify VLAN assignment per port with the command: `port-vlan-show ports ##`

```
CLI (network-admin@pnswitch1) > port-vlan-show ports 43
port vlans untagged-vlan description       active-vlans
---- ----- ------------- ----------------- ------------
43   4091  4091          to C4500-1 ten1/4 none
```

Similar to the VLAN status command, this command will show the relationship between the VLAN and the ports that allow that VLAN traffic.

Monitoring and tracking a specific MAC address with the L2 table within the fabric is a powerful troubleshooting tool. MAC address and port relationship is also an important troubleshooting tool. Beyond the L2 table the Pluribus switch also creates and tracks vPorts. vPorts provide a fabric wide view on MAC addresses and their relationship to VLANs and other critical information.

L2 Table Status

View the MAC addresses status for a switch using: `l2-table-show`

```
CLI (network-admin@pnswitch1) > l2-table-show
mac               vlan ip              ports       state         hostname  status      migrate
----------------- ---- --------------- ----------- ------------- --------- ----------- -------
00:00:46:43:f3:bd 855  10.80.160.76    61          active                              10
00:00:47:7c:f6:b0 528  10.55.28.122    61          active                              10
00:00:47:ac:51:73 518  10.55.19.103    61          active                              10
00:00:47:ac:c2:d8 524  10.55.24.170    61          active                              11
66:0e:94:0c:c8:cb 412  10.80.112.2     65          active,static pnswitch1 PN-internal 153
00:00:47:a1:eb:35 526  10.55.27.0      61          active                              10
00:00:46:4b:3d:ae 845  10.80.97.145    61          active                              10
00:00:4c:06:92:a2 514  10.81.114.106   61          active                              9
00:00:46:4b:3d:e2 845  10.80.97.171    61          active                              10
00:00:47:ac:51:11 518  10.55.19.54     61          active                              11
CLI (network-admin@pnswitch1) >
```

vPort Status

View the vPort information using: `vport-show`

```
CLI (network-admin@pnswitch1) > vport-show mac 00:00:4c:06:91:f8
mac               vlan ip            ports state  migrate
----------------- ---- ------------ ----- ------ -------
00:00:4c:06:91:f8 514  10.81.114.21 61    active 9
CLI (network-admin@pnswitch1) >
```

vPort History Status

View the vPort history (for a specific MAC address in the following example) using: `vport-history-show`

```
CLI (network-admin@pnswitch1) > vport-history-show mac 00:00:4c:06:91:f8
time           log-type  mac               vlan ip            ports state  local-ports local-state migrate
-------------- --------- ----------------- ---- ------------ ----- ------ ----------- ----------- -------
11-17,21:52:30 save      00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
11-17,21:57:07 l2-modify 00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
11-17,22:02:07 l2-modify 00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
11-17,22:07:07 l2-modify 00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
11-17,22:12:07 l2-modify 00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
11-17,22:17:06 l2-modify 00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
11-17,22:22:06 l2-modify 00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
11-17,22:27:06 l2-modify 00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
11-17,22:32:06 l2-modify 00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
11-17,22:37:06 l2-modify 00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
11-17,22:42:06 l2-modify 00:00:4c:06:91:f8 514  10.81.114.21 61    active 61          active      9
CLI (network-admin@pnswitch1) >
```

L3 Troubleshooting

Layer 3 communications is governed by the vRouter. If traffic is not flowing between Layer 2 domains then confirm the vRouters and vRouter interfaces are configured properly. In addition, OSPF and BGP configuration commands are reviewed as well.

vRouter Status

Verify vrouter status with the command: `vrouter-show`

```
CLI (network-admin@pnswitch1) > vrouter-show
name                  type    scope  vnet     vnet-service state   router-type hw-router-mac     hw-vrid hw-vrrp-id router-id   proto-mu
ospf-redistribute
--------------------- ------- ------ -------- ------------ ------- ----------- ----------------- ------- ---------- ----------- --------
---------------
pnswitch2-test        vrouter fabric test     dedicated    enabled hardware    66:0e:94:17:e7:fa 0       210        10.60.7.115 pim-spar
pnswitch1-test        vrouter fabric test     dedicated    enabled hardware    66:0e:94:0c:07:d0 0       210        10.60.7.114 pim-spar
pnswitch2-test2       vrouter fabric test2    dedicated    enabled hardware    66:0e:94:0c:c8:cb 1       204        10.60.7.110 pim-spar
static
pnswitch1-test2       vrouter fabric test2    dedicated    enabled hardware    66:0e:94:17:cf:c6 1       204        10.60.7.111 pim-spar
static
pnswitch2-test3       vrouter fabric test3    dedicated    enabled hardware    66:0e:94:04:7f:91 2       201        10.60.7.104 pim-spar
pnswitch1-test3       vrouter fabric test3    dedicated    enabled hardware    66:0e:94:0b:cd:85 2       201        10.60.7.105 pim-spar
pnswitch2-test4       vrouter fabric test4    dedicated    enabled hardware    66:0e:94:04:5c:6d 3       202        10.60.7.106 pim-spar
pnswitch1-test4       vrouter fabric test4    dedicated    enabled hardware    66:0e:94:0b:a5:1b 3       202        10.60.7.107 pim-spar
pnswitch2-test5       vrouter fabric test5    dedicated    enabled hardware    66:0e:94:04:f8:2b 4       205        10.60.7.112 pim-spar
pnswitch1-test5       vrouter fabric test5    dedicated    enabled hardware    66:0e:94:0b:a1:f1 4       205        10.60.7.113 pim-spar
CLI (network-admin@pnswitch1) >
```

vRouter Interface Status

Verify vrouter interface status with the command: `vrouter-interface-show`

```
CLI (network-admin@pnswitch1) > vrouter-interface-show vlan 512
vrouter-name          nic      ip             assignment mac               vlan vxlan if   exclusive nic-config nic-state vrrp-id vrrp-pr
vrrp-priority vrrp-state
--------------------- -------- -------------- ---------- ----------------- ---- ----- ---- --------- ----------- ---------- -------- -------
------------- ----------
pnswitch2-test2 eth0.512 10.81.112.2/24 static     66:0e:94:0c:c8:cb 512  0     data no        enable     up
pnswitch2-test2 eth2.512 10.81.112.1/24 static     00:00:5e:00:01:cc 512  0     data no        enable     up        204     eth0.512
master
pnswitch1-test2 eth1.512 10.81.112.3/24 static     66:0e:94:17:cf:c6 512  0     data no        enable     up
pnswitch1-test2 eth3.512 10.81.112.1/24 static     00:00:5e:00:01:cc 512  0     data no        enable     down      204     eth1.512
slave
CLI (network-admin@pnswitch1) >
```

VRRP

Verify the vrrp configured vRouter interfaces with the command: vrouter-interface-show

```
CLI network-admin@pnswitch1 >vrouter-interface-show format all layout vertical
vrouter-name: vrrp-router1
nic: eth0.100
ip: 192.168.11.3/24
assignment: static
mac: 66:0e:94:dd:18:c4
vlan: 100
vxlan: 0
if: data
alias-on:
exclusive: no
nic-config: enable
nic-state: up
vrouter-name: vrrp-router1
nic: eth1.100
ip: 192.168.11.2/24
assignment: static
mac: 00:00:5e:00:01:0a
vlan: 100
vxlan: 0
if: data
alias-on:
exclusive: no
nic-config: enable
nic-state: up
vrrp-id: 10
vrrp-primary: eth1.100
vrrp-priority: 100
vrrp-state: master
vrouter-name: vrrp-router2
nic: eth3.100
ip: 192.168.11.4/24
assignment: static
mac: 66:0e:94:21:54:07
vlan: 100
vxlan: 0
if: data
alias-on:
exclusive: no
nic-config: enable
nic-state: up
vrouter-name: vrrp-router2
nic: eth3.100
ip: 192.168.11.2/24
assignment: static
mac: 00:00:5e:00:01:0a
vlan: 100
vxlan: 0
if: data
alias-on:
exclusive: no
nic-config: enable
nic-state: down
Pluribus Networks Configuration Guide
pluribusnetworks.com 87
vrrp-id: 10
```

```
vrrp-primary: eth3.100
vrrp-priority: 50
vrrp-state: slave
```

vRouter Static Routes

Verify vrouter static routes with the command: `vrouter-static-route-show`

```
CLI (network-admin@pnswitch1) > vrouter-static-route-show network 172.16.0.0/16
vrouter-name          network         gateway-ip     distance
--------------------- --------------- -------------- --------
pnswitch2-test3       172.16.14.0/24  198.105.67.124 200
pnswitch2-test3       172.16.51.0/24  172.16.6.6     200
pnswitch1-test3       172.16.14.0/24  198.105.67.124 200
pnswitch1-test3       172.16.51.0/24  172.16.6.6     200
pnswitch2-test4       172.16.0.0/12   10.95.1.2      200
pnswitch2-test4       172.16.5.0/24   10.37.15.10    200
pnswitch2-test4       172.16.8.0/24   10.37.15.10    200
pnswitch2-test4       172.16.14.0/24  10.37.0.18     200
pnswitch2-test4       172.16.20.0/24  10.37.0.14     200
pnswitch2-test4       172.16.50.0/24  10.37.15.10    200
pnswitch2-test4       172.16.51.0/24  10.37.15.10    200
pnswitch2-test4       172.16.200.0/21 10.37.0.18     200
pnswitch2-test4       172.16.208.0/21 10.37.0.18     200
pnswitch1-test4       172.16.0.0/12   10.95.1.2      200
CLI (network-admin@pnswitch1) >
```

vRouter OSPF Configuration

Verify vrouter OSPF configuration with the command: `vrouter-ospf-show`

```
CLI (network-admin@pnswitch1) > vrouter-ospf-show vrouter-name pnswitch2-test4
vrouter-name          network        ospf-area
--------------------- -------------- ---------
pnswitch2-test4       10.60.26.33/32 202
pnswitch2-test4       10.60.26.45/32 202
pnswitch2-test4       10.60.7.106/32 202
pnswitch2-test4       10.60.26.32/30 202
pnswitch2-test4       10.60.26.44/30 202
pnswitch2-test4       10.95.0.0/18   202
pnswitch2-test4       10.37.0.0/20   202
pnswitch2-test4       10.37.63.0/24  202
CLI (network-admin@pnswitch1) >
```

Verify vrouter OSPF area configuration with the command: `vrouter-ospf-area-show`

Verify vrouter OSPF neighbor configuration with the command: `vrouter-ospf-neighbor-show`

vRouter BGP Configuration

Verify vrouter BGP configuration with the command: `vrouter-bgp-show`

Verify vrouter BGP neighbor configuration with the command: `vrouter-bgp-neighbor-show`

vRouter Packet Relay Configuration

Verify vrouter packet relay configuration with the command: `vrouter-packet-relay-show`

```
CLI (network-admin@pnswitch1) > vrouter-packet-relay-show vrouter-name pnswitch2-test4
vrouter-name          forward-proto forward-ip    nic
--------------------- ------------- ------------- --------
pnswitch2-test4 dhcp          10.55.200.10  eth3.295
pnswitch2-test4 dhcp          10.55.200.200 eth3.295
pnswitch2-test4 dhcp          10.55.200.10  eth3.500
pnswitch2-test4 dhcp          10.55.200.200 eth3.500
pnswitch2-test4 dhcp          10.55.200.10  eth3.605
pnswitch2-test4 dhcp          10.55.200.200 eth3.605
CLI (network-admin@pnswitch1) >
```

VxLAN

VxLAN and VLE configuration has three steps.

1. Associate VxLAN to VLAN. For VLE, VxLAN mode will be transparent.

```
vlan-create id 3015 vxlan 5003015 vxlan-mode transparent scope local description Voztelecom-Espanix no-stats
ports 5 untagged-ports 5
```

```
CLI (network-admin@SW-BCN-S4048-1*) > vlan-show id 3015

id   type   vxlan   vxlan-type vxlan-mode  scope description        active stats ports untagged-ports
active-edge-ports

---- ------ ------- ---------- ----------- ----- ------------------ ------ ----- ----- -------------- ------
-----------

3015 public 5003015 user       transparent local Voztelecom-Espanix yes    no    5     5              none

CLI (network-admin@SW-BCN-S4048-1*) >
```

2. Create a Tunnel

```
tunnel-create scope local name SW-BCN-S4048-1-TO-SW-MAD2-S4048-1 vrouter-name SW-BCN-S4048-1 local-ip
10.40.31.1 remote-ip 10.40.21.1
```

```
CLI (network-admin@SW-BCN-S4048-1*) > tunnel-show name SW-BCN-S4048-1-TO-SW-MAD2-S4048-1

scope name                            type  vrouter-name   local-ip   remote-ip  router-if next-hop  next-
hop-mac      nexthop-vlan remote-switch active state error route-info    ports

----- --------------------------------- ----- -------------- ---------- ---------- --------- --------- -----
------------ ------------ ------------- ------ ----- ----- ------------- -----

local SW-BCN-S4048-1-TO-SW-MAD2-S4048-1 vxlan SW-BCN-S4048-1 10.40.31.1 10.40.21.1 eth0.4011 10.0.13.2
66:0e:94:ea:72:7a 4092         0             yes    ok          10.40.21.0/30 128

CLI (network-admin@SW-BCN-S4048-1*) >
```

3. Associate VxLAN to tunnel

```
tunnel-vxlan-add name SW-BCN-S4048-1-TO-SW-MAD2-S4048-1 vxlan 5003015


CLI (network-admin@SW-BCN-S4048-1*) > tunnel-vxlan-show name SW-BCN-S4048-1-TO-SW-MAD2-S4048-1

name                              vxlan

--------------------------------- -------

SW-BCN-S4048-1-TO-SW-MAD2-S4048-1 5003015
```

4. Vrouter interface status

```
CLI (network-admin@SW-BCN-S4048-1*) > vrouter-interface-show nic eth0.4011

vrouter-name   nic       ip            assignment mac               vlan vlan-type if   vm-nic-type
exclusive nic-config nic-state mtu  sriov-vf mirror-traffic

-------------- --------- ------------- ---------- ----------------- ---- --------- ---- ----------- --------
- ---------- --------- ---- -------- --------------

SW-BCN-S4048-1 eth0.4011 10.40.31.1/30 static     66:0e:94:ca:22:39 4011 public    data none        no
enable     up        9216 false    false

CLI (network-admin@SW-BCN-S4048-1*) >
```

5. Route to Tunnel next-hop

```
CLI (network-admin@SW-BCN-S4048-1*) > vrouter-routes-show network 10.40.21.1

vrouter-name   network       type interface next-hop  distance metric

-------------- ------------- ---- --------- --------- -------- ------

SW-BCN-S4048-1 10.40.21.0/30 bgp  eth0.4092 10.0.13.2 20       0

CLI (network-admin@SW-BCN-S4048-1*) >
```

Multicast

L3 Table Status

Verify the L3 table status with the command: `l3-table-show`

```
CLI (network-admin@pnswitch1) > l3-table-show ip 10.81.114.21
mac               ip           vlan vxlan rt-if    l3-flags
----------------- ------------ ---- ----- -------- --------
00:00:4c:06:91:f8 10.81.114.21 514  0     eth0.514
```

Forwarding Information Base (FIB) Status

Verify the FIB status with the command: `vrouter-fib-routes-show`

```
CLI (network-admin@pnswitch1) > vrouter-fib-routes-show ip 10.81.114.21
vrid ip           prelen nexthop      if-ip       vlan flags ecmp-group state
---- ------------ ------ ------------ ----------- ---- ----- ---------- -----
1    10.81.114.21 32     10.81.114.21 10.81.114.2 514        -1         up
```

FIB ARP Status

Verify the FIB ARP status with the command: `vrouter-fib-arps-show`

```
CLI (network-admin@pnswitch1) > vrouter-fib-arps-show mac 00:00:4c:06:91:f8
switch    ip           if-id vlan mac               flags
--------- ------------ ----- ---- ----------------- -----
pnswitch1 10.81.114.21 7     514  00:00:4c:06:91:f8
pnswitch2 10.81.114.21 7     514  00:00:4c:06:91:f8
```

Quagga Logs

This command displays the Quagga based routing Information from Netvisor CLI:

```
CLI (network-admin@Spine1) > help vrouter-vtysh-cmd
vrouter-vtysh-cmd display output of a Quagga show command
name name-string name of service config
cmd cmd-string any Quagga show/debug/undebug/no debug command

Show Output Examples
CLI (network-admin@Spine1) > vrouter-vtysh-cmd name vr1 cmd "show ip route"

Codes: K - kernel route, C - connected, S - static, R - RIP,
O - OSPF, I - IS-IS, B - BGP, P - PIM, A - Babel, > - selected route, * - FIB route

C>* 100.1.1.0/24 is directly connected, eth1.100 C>* 127.0.0.0/8 is directly connected, lo0 CLI (network-
admin@Spine1) > vrouter-vtysh-cmd name vr1 cmd "show running-config" Building configuration...
Current configuration:
!
hostname vr1
log file zebra.log
!
password zebra
enable password zebra
!
interface eth1.100
 ipv6 nd suppress-ra
 multicast
 no link-detect
!
interface lo0
 no link-detect
!
ip forwarding
ipv6 forwarding
!
line vty
!
end
```

Quagga logs files, such as Zebra, OSPF, OSPF6, BGP, BFD and RIP can also be viewed directly from the console using the Netvisor OS CLI.

```
CLI network-admin@pnswitch1 > vrouter-log-show vrouter-name vrouter-name protocol
       vrouter-log-show Displays vrouter protocol logs one or more of the following options:
       vrouter-name vrouter-name protocol   Specify name of the vRouter.
       zebra|ospf|ospf6|bgp|bfd|rip         Specify the name of the protocol files to view.
```

Quagga log files accumulate on the switch, so we highly recommend clearing these logs after the troubleshooting session

```
       vrouter-log-clear Clears vrouter protocol logs from a protocol log file
       one or more of the following options:

       vrouter-name vrouter name protocol     Specify the name of the vRouter service.
       zebra|ospf|ospf6|bgp|bfd|rip          Specify the name of the log file to clear.
```

Connection Analytics

Connection-show

This command shows every connection and details like protocol type, connection state, src/dst IP's, latency, In/Out bytes usage, etc. Using this command, Administrators can analyze network performance and even isolate the network/host as the cause of the problem. Connection-Show can be used to review the following (not limited to) scenarios:

- Connections that happened between 8.00AM and 10.00AM today
  - *Connection-show start-time 2016-09-21T08:00 end-time 2016-09-21T10:00c*
- Connections specific to a Host
  - *Connection-show src-ip < Ip address >*
- To see which IP is receiving the most traffic
  - *Connection-show within-last 5m sort-desc total-bytes, sum-by dst-ip*
- To see all recent mac-moves
  - *l2-history-show sort-desc migrate,*

Ports-Stats-show

This command provides better insight on a physical port and the connected hosts on that port.

Some scenarios covered by port-stats-show:

- Display port statistics for a particual port.
- Identify the busiest physical port on the switch
  - *port-stats-show sort-desc ibytes*
- Identify and show the recent traffic on ports
  - *port-stats-show show-diff-interval 5s sort-desc ibytes,*

```
CLI (network-admin@pnswitch1) > port-stats-show port 64

time     port ibytes iUpkts iBpkts iMpkts iCongDrops ierrs obytes oUpkts oBpkts oMpkts oCongDrops oerrs mtu-errs

-------- ---- ------ ------ ------ ------ ---------- ----- ------ ------ ------ ------ ---------- ----- --------

09:40:38 64   11.6G  31.2M  7.21K  109M   0          0     3.30G  21.0M  8.75K  904K   0          0     0

CLI (network-admin@pnswitch) > port-stats-show sort-desc ibytes
```

time	port	ibytes	iUpkts	iBpkts	iMpkts	iCongDrops	ierrs	obytes	oUpkts	oBpkts	oMpkts	oCongDrops	oerrs	mtu-errs
09:42:36	62	870G	10.8G	1.15K	3.32M	0	0	866G	10.7G	513K	3.08M	0	0	0
09:42:36	61	865G	10.7G	1.17K	66.6K	0	0	874G	10.8G	516K	11.0M	0	0	0
09:42:36	60	553G	5.95G	427	22.6K	0	0	510G	5.46G	474	12.7K	0	0	0
09:42:36	59	551G	5.93G	454	36.6K	0	0	510G	5.46G	384	16.1K	0	0	0
09:42:36	55	469G	4.99G	396	28.6K	0	0	512G	5.48G	2.36K	18.0K	0	0	0

Other Tools and Configuration

There are other tools that can assist in troubleshooting. In addition, there are several configuration commands that should be reviewed.

vFlow Introduction

The vFlow feature allows an administrator to monitor, capture, or manipulate network flows. vFlows are created by specifying matching criteria to isolate particular network traffic, then prescribing an action to take once the matching traffic is isolated. vFlows are very powerful network tools and provide troubleshooting information on specific network traffic.

vFlow Status

View the configured vFlows with the `vflow-show` command. Note that the action of this command is to drop the matching traffic.

View vFlow configuration with the command: `vflow-show`

```
CLI (network-admin@pnswitch1) > vflow-show name DMZ-isolate-vlan-214-in-permit-20
name                 scope  type  ether-type src-ip                 precedence action
-------------------------------- ------ ----- ---------- --------------------------
vflow-test           fabric vflow ipv4    172.16.214.0/255.255.255.0   13       drop
```

vFlow Statistics

The vflow facility also captures statistics about the defined vflows. This is useful to measure activity for a given vflow.

View vflow statistics with the command: `vflow-stats-show`

```
CLI (network-admin@pnswitch1) > vflow-stats-show
switch    name                                pkts  bytes cpu-pkts cpu-bytes drops drop-bytes
--------- ----------------------------------- ----- ----- -------- --------- ----- ----------
pnswitch1 Fabric-Keepalive                    1.67M 1.12G 1.67M    1.12G     9     6.03K
pnswitch1 iSCSI-ACL-vlan-872-permit-1         0     0     0        0         0     0
```

Other vFlow Commands

The vflow facility is extensive and provides many other features and functions. Following are some of the other vflow commands, refer to the product technical documentation for a complete description of the vflow commands.

```
vflow-class-show           display virtual flow class information
vflow-create               create a virtual flow definition for L2 or L3 IP
vflow-show                 display virtual flow information
vflow-snoop                display the packet headers of flows directed to the server-switch CPU
vflow-stats-show           display packet statistics or logs for the vflow
```

vflow-snoop

Capturing and Analyzing traffic is an extremely valuable tool for troubleshooting. Snooping only works if you use the parameters, copy-to-cpu or to-cpu. The copy-to-cpu parameter ensures that the data plane forwards the packets and sends a copy to the CPU. Use this parameter if you want traffic to flow through the switch.

Use the `vflow-snoop` , we can capture and analyze interesting traffic. Example below:

```
CLI network-admin@pnswitch1 > vflow-create name snoop_ssh scope local action copy-to-cpu src-port 22 proto
vflow-add-filter name snoop_ssh
```

```
CLI network-admin@pnswitch1 > vflow-snoop name snoop_ssh

switch: pleiades24, flow: snoop_ssh, port: 41, size: 230, time: 10:56:57.05785917 src-mac: 00:15:17:ea:f8:7
dst-mac: f4:6d:04:0e:77:60, etype: ip src-ip: 10.9.11.18, dst-ip: 10.9.10.65, proto: tcp src-port: 22, dst-
port: 62356
 switch: pleiades24, flow: snoop_ssh, port: 41, size: 118, time: 10:56:57.05922560 src-mac: 00:15:17:ea:f8:
dst-mac: f4:6d:04:0e:77:60, etype: ip src-ip: 10.9.11.18, dst-ip: 10.9.10.65, proto: tcp src-port: 22, dst-
port: 62356
```

Flowtrace

Flowtrace is Netvisor shell based capture utility, which can be used to capture interesting traffic for troubleshooting purposes.

Example of Flowtrace command to capture icmp packets

flowtrace --proto icmp --src-ip <ip> --dst-ip <ip> -e13 -r

```
Tracing icmp from x.x.x.x -> y.y.y.y
Tracing icmp from y.y.y.y -> x.x.x.x
Jan23.09:44:21 icmp x.x.x.x -> y.y.y.y echo request ident 17656 seq 1
1.0                    -> Switch1 port 20 x.x.x.x:05:56:a0 -> x.x.x.x:1a:8f vlan 200  ttl 64
1.1   Switch1 port 69 ->
1.0                    -> Switch1 port 45 x.x.x.x:f2:9d:ff -> x.x.x.x:d4:9c:b6 vlan 4092 ttl 62
```

Flowtrace Shell command options:

```
MODE (one of):
  --live                   : trace from live packets (default)
  --conn                   : trace from connection-show data
  --vflows <list>          : trace from copy-to-cpu or to-cpu vflows (comma separated list)
  --from-pcap <file>       : trace from pcap-ng file written by flowtrace
  --fabconn                : trace from fabric-connection-show data
LIVE REQUIRED:
  --proto [icmp|tcp|udp]   : protocol type
  --client-ip                : client ip
  --server-ip                : server ip
LIVE OPTIONS:
  --precedence|-e <num>    : set vflow precedence
  --trace-replies|-r       : also swap src/dst ip and trace
  --no-exit-port|-n        : don't look up exit port (faster)
  --collect-time|-c <time> : set time to collect data
  --trace-memory|-m <val>  : table memory size
  --vflow-class            : set vflow-class
CONN/FABCONN REQUIRED:
  --client-ip                : client ip
  --server-ip                : server ip
CONN/FABCONN OPTIONS:
  --client-port <num>      : client tcp port
  --server-port <num>      : server tcp port
  --time                   : connection data at time
  --start-time             : start time for data collection
  --end-time               : end time for data collection
  --duration               : duration of time for data collection
  --since-start            : start-time set to last nvOSd start
  --older-than             : data older than time
  --within-last            : data within last time
  --show-bytes             : show bytes transferred
  --unscaled               : do not scale large values
  --trace-memory|-m <val>  : table memory size
TO-PCAP-FILE OPTIONS:
```

```
  --to-pcap <file>         : write packets to pcap-ng file
  --no-trace               : don't trace, just write to file
FROM-PCAP-FILE OPTIONS:
  --ip                     : filter by client or server ip
  --l4-port                : filter by client or server L4 port
FABCONN OPTIONS:
  --tid <num>              : transaction id (only valid at endpoints)
COMMON OPTIONS:
  --user|-u name           : user name for authentication
  --pass|-p pass           : password for authentication
  --host|-h ip[:port]      : connect to remote switch
  --add-fabric|-f          : additional fabric [user@]ip[:port] (may be specified multiple times)
  --debug|-d               : print debug info
```

The following configuration settings should only be changed under the direction of Pluribus Networks support personnel.

Port Storm Control Configuration

View the port storm control settings with the command: `port-storm-control-show`

```
CLI (network-admin@pnswitch1) > port-storm-control-show
'switch       port speed unknown-ucast-level unknown-mcast-level broadcast-level
------------ ---- ----- -------------------- -------------------- ---------------
pnswitch1 1    10g   30%                  30%                  30%
pnswitch1 2    10g   30%                  30%                  30%
pnswitch1 3    10g   30%                  30%                  30%
pnswitch1 4    10g   30%                  30%                  30%
pnswitch1 5    10g   30%                  30%                  30%
pnswitch1 6    10g   30%                  30%                  30%
pnswitch1 7    10g   30%                  30%                  30%
pnswitch1 8    10g   30%                  30%                  30%
pnswitch1 9    10g   30%                  30%                  30%
pnswitch1 10   10g   30%                  30%                  30%
```

L2 Setting Configuration

View the L2 control settings with the command: `l2-setting-show`

```
CLI (network-admin@pnswitch1*) > l2-setting-show
aging-time(s):       300
software-aging:      on
l2-max-count:        1200000
l2-cur-count:        82
l2-active-count:     19
l2-max-mem:          2.01G
l2-cur-mem:          144K
l2-checker:          disabled
l2-checker-interval: 10m
l3-arp-max-count:    1200000
l3-arp-cur-count:    42
l3-arp-max-mem:      632M
l3-arp-cur-mem:      22.6K
```

System Control Settings

View the system control settings with the command: sys-flow-setting-show

```
CLI (network-admin@pnswitch1*) > system-settings-show
optimize-arps:               on
lldp:                        on
optimize-nd:                 on
reactivate-mac:              on
reactivate-vxlan-tunnel-mac: on
manage-unknown-unicast:      on
manage-broadcast:            on
auto-trunk:                  on
auto-host-bundle:            off
routing-over-vlags:          off
```

Logging and Monitoring

Log messages can provide useful troubleshooting information. However, the variety of log messages precludes covering each possible log message. Instead, log messages can be captured for analysis by Pluribus Networks support personnel.

System Log Status

View the system log messages with the command: `log-system-show`

```
CLI (network-admin@pnswitch1) > log-system-show

category time                            name      code  level message
-------- ------------------------------- --------- ----- ----- ----------------------------------
-------------------------------------------------
system   2015-11-18,10:11:22.601638-07:00 mac_moved 11017 note  MOVING MAC=00:00:5e:00:01:c9, vlan
203 from port 130 () to port 65 :: flow-cb
system   2015-11-18,10:11:22.601967-07:00 mac_moved 11017 note  MOVING MAC=00:00:5e:00:01:c9, vlan
207 from port 130 () to port 65 :: flow-cb
system   2015-11-18,10:11:22.621108-07:00 mac_moved 11017 note  MOVING MAC=00:00:5e:00:01:c9, vlan
202 from port 130 () to port 65 :: flow-cb
system   2015-11-18,10:11:22.641608-07:00 mac_moved 11017 note  MOVING MAC=00:00:5e:00:01:c9, vlan
208 from port 130 () to port 65 :: flow-cb
```

Event Log Status

View the event log messages with the command: `log-events-show`

```
CLI (network-admin@pnswitch1) > log-event-show
category time                            name                code  event-type port message
-------- ------------------------------- ------------------- ----- ---------- ---- --------------
-------------------------------------------------
event    2015-11-08,02:34:25.605809-07:00 stp_port_state      11026 port            STP Port State
Change: port=57 vlan=4094: Disabled -> Forwarding
event    2015-11-08,02:34:25.622182-07:00 adj_trunk_create    11106 port            ADJ Create
Auto Trunk: port1=58 port2=55
event    2015-11-08,02:34:25.652698-07:00 port_down           11003 port       58   down
event    2015-11-08,02:34:25.710632-07:00 port_up             11002 port       55   up
event    2015-11-08,02:34:25.710981-07:00 adj_trunk_port_add  11107 port            ADJ Trunk Port
Add: trunk=auto-128 port=63
```

Audit Log

View the Audit Log messages with the command : log-audit-show

```
CLI (network-admin@pnswitch1) > log-audit-show
category time                            name   code  user          client-addr message
-------- ------------------------------- ------ ----- ------------- ----------- -------
audit    2017-09-20,22:20:39.550756-07:00 logout 11100 network-admin             logout
audit    2017-09-20,22:21:09.560493-07:00 login  11099 network-admin             login
audit    2017-09-20,22:21:10.306290-07:00 logout 11100 network-admin             logout
audit    2017-09-20,22:21:40.316108-07:00 login  11099 network-admin             login
audit    2017-09-20,22:21:41.054306-07:00 logout 11100 network-admin             logout
```

SNMP

The SNMP daemon runs as a service and is launched by using the following command:

```
CLI network-admin@ pnswitch1 > admin-service-modify if mgmt snmp
```

To view the SNMP Service status, use the following command: admin-service-show

```
CLI (network-admin@pnswitch1) > admin-service-show
switch      if   ssh nfs web web-ssl web-ssl-port web-port snmp net-api icmp
----------- ---- --- --- --- ------- ------------ -------- ---- ------- ----
pnswitch1 mgmt on  off on  off     443          80       on   on      on
pnswitch1 data on  off off off     443          80       off  on      on
```

SNMP – Support MIBs: -

- IfTable
- IfXTable
- EntPhySensorTable

Additional SNMP Commands:

```
snmp-show                                                          display SNMP information
        user-name snmp-user user-name                              username
        name snmp-oid name                                         SNMP OID name
        show-type walk|get|get-next                                type of community
snmp-community-show                                                display SNMP communities
for SNMPv1
        [ community-string community-string-string ]               community name
        [ community-type read-only|read-write ]                    community type
snmp-vacm-show                                                     display View Access
Control Models (VACM)
        [ user-type rouser|rwuser ]                                SNMP user type
        [ user-name snmp-user user-name ]                          SNMP administrator name
        [ oid-restrict oid-restrict-string ]                       restrict OID
        [ view view-string ]                                       view type
        [ auth|no-auth ]                                           authentication required
        [ priv|no-priv ]                                           privileges
snmp-user-show                                                     display SNMPv3 users
        [ user-name user-name-string ]                             SNMP user name
        [ auth|no-auth ]                                           authentication required
        [ auth-hash md5|sha ]                                      Hashing algorithm for
authentication
        [ priv|no-priv ]                                           privileges
snmp-trap-enable-modify                                            modify SNMP notifications
about link conditions
    one or more of the following options:
        link-up-down|no-link-up-down                               link is up or down
        default-monitors|no-default-monitors                       default monitoring
        physical-sensors|no-physical-sensors                       Temperature, fan speed,
etc.
        low-disk-space|no-low-disk-space                           Low disk space
        low-disk-space-threshold low-disk-space-threshold-string   Threshold value of low-
disk-space in %
        system-usage|no-system-usage                               Memory and CPU usage
        high-system-usage-threshold high-system-usage-threshold-string  Threshold value of
system-usage in %
        login-failure|no-login-failure                             Incorrect passwords on
login
        lacp-status|no-lacp-status                                 LACP status
```

```
        vport-modified|no-vport-modified                                vPort modification
        stp-port-modified|no-stp-port-modified                          STP port modified
        mirror-to-cpu|no-mirror-to-cpu                                  Mirror to CPU configured
        stp-port-state-failed|no-stp-port-state-failed                  STP Port State Failed
        link-congestion-detected|no-link-congestion-detected            Congestion detected at
port
        fabric-node-state-changed|no-fabric-node-state-changed          Fabric Node State Changed
snmp-trap-enable-show                                                   display information about
SNMP traps
```

Syslog

To display the current syslog configuration, use command: admin-syslog-show

```
CLI network-admin@switch > admin-syslog-show
name    scope   host           port   message-format
log-all fabric  172.16.21.67   514    legacy
```

To specify sending the syslog messages in structured format, per RFC5424, add the message-format option to the configuration.

```
CLI network-admin@switch > admin-syslog-modify name log-all message-format structured
```

About
Pluribus Networks

Pluribus Networks is simplifying the Software-Defined Data Center with its simple, dynamic and secure Adaptive Cloud Fabric architecture, enabling organizations to build scalable private and public clouds that improve service velocity, performance, and reliability. The company's innovative Netvisor software virtualizes open networking hardware to build a holistic, distributed network that is more intelligent, automated and resilient. The company's Insight Analytics platform leverages embedded telemetry and other data sources to enable pervasive visibility across the network to reveal network and application performance that speeds troubleshooting and improves operational and security intelligence. Pluribus Networks has received venture funding from Temasek Holdings, NEA, Menlo Ventures, and AME Cloud Ventures. Pluribus Networks is headquartered in San Jose, California, with development and support centers in Bangalore, India; Phoenix, AZ; Dallas, TX; and Hong Kong, PRC. For additional information contact Pluribus Networks at info@pluribusnetworks.com, or visit www.pluribusnetworks.com. Follow us on Twitter @pluribusnet.

Pluribus Networks, Inc.
6001 America Center Drive Suite 450
San Jose, CA 95002
1 650 289 4717
1 855 438 8638

Pluribus NETWORKS