



OcNOS®

Open Compute Network Operating System for Data Centers Version 7.0.0

Troubleshooting Guide

February 2026

© 2026 IP Infusion Inc. All Rights Reserved.

This documentation is subject to change without notice. The software described in this document and this documentation are furnished under a license agreement or nondisclosure agreement. The software and documentation may be used or copied only in accordance with the terms of the applicable agreement. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or any means electronic or mechanical, including photocopying and recording for any purpose other than the purchaser's internal use without the written permission of IP Infusion Inc.

IP Infusion Inc.
3979 Freedom Circle
Suite 900
Santa Clara, California 95054
<http://www.ipinfusion.com/>

For support, questions, or comments via E-mail, contact:
support@ipinfusion.com

Trademarks:

IP Infusion and OcNOS are trademarks or registered trademarks of IP Infusion. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Use of certain software included in this equipment is subject to the IP Infusion, Inc. End User License Agreement at <http://www.ipinfusion.com/license>. By using the equipment, you accept the terms of the End User License Agreement.

Contents

Preface	6
Audience	6
Conventions	6
IP Infusion Product Release Version	6
Related Documentation	7
Feature Availability	7
Migration Guide	7
IP Maestro Support	7
Technical Support	7
Technical Sales	7
Technical Documentation	7
.....	9
CHAPTER 1 Debugging and Logging	10
About Debugging	10
Start Debugging Output	10
Log to Standard Output	10
Log to a File	11
Stop Debugging	11
Debug Modes	11
CHAPTER 2 Debugging Kernel Crash	12
Kernel Dump Extraction and Analysis	12
With Technical Support Archive	12
Without Technical Support Archive	12
Basic Commands in Crash	14
CHAPTER 3 Layer 2 Switching	16
Spanning Tree Protocol	16
BPDU Guard	18
BPDU Filter	19
Root Guard	19
VLANs	19
Interfaces	20
Link Light for the Port does not come on	20
Other	21
Link Aggregation	21
Multi-Chassis Link Aggregation	22
LLDP	22
802.1x	23
VLAN Cross-Connect	25
IGMP Snooping	25
PVLAN	26

CHAPTER 4	Layer 3 Routing	27
	Missing Route	27
	RIP	27
	No RIP Adjacency	27
	BGP	28
	OSPFv2	32
	Neighborship Formation	32
	Virtual Link Neighborship Formation	33
	Multi-area Adjacency Formation	34
	Graceful Restart (using Link Local Signaling) Neighborship Formation	34
	OSPFv3	35
	Neighborship Formation	35
	VRRP	37
	BFD	38
	IS-IS	38
	Adjacency Problems	38
	Routing Update Problems	39
	FAQ	41
CHAPTER 5	Multicast	44
	IGMP	44
	IGMP not Active	44
	Groups not added after Receiving Dynamic Reports	45
	IGMP proxy is not active	45
	MLD	46
	MLD not Active	46
	Groups not Added on Receiving Dynamic Reports	47
	MLD Proxy not Active	47
	PIM	48
	RPF Neighbor not Reachable	48
	RPF Interface not What is Expected	48
	RP not Reachable	49
	Mroute not Created for PIM-SSM	49
	MSDP	50
	MSDP Peering not Established	50
CHAPTER 6	HQoS	51
CHAPTER 7	Data Center and Virtualization	53
	VXLAN	53
	VXLAN-EVPN	54
CHAPTER 8	Security	55
	DHCP Snooping	55
	DHCP Snooping IP Source Guard	55
	DHCP Snooping over MLAG	56
	DHCPv6-Prefix Delegation	56

CHAPTER 9	SNMP	58
CHAPTER 10	DHCP Client and Relay	59
CHAPTER 11	Remote Logging	60
CHAPTER 12	System Management	61
	License Troubleshooting	62
	Zero Touch Provisioning	65
CHAPTER 13	TACACS+ and AAA	66
	AAA console authentication via Management VRF	67
CHAPTER 14	Configuration Management	69
	Replacing the Start-up Configuration File	69
	OcNOS Datastores	70

Preface

This guide describes how to configure OcNOS.

Audience

This guide is intended for network administrators and other engineering professionals who configure OcNOS.

Conventions


[Table 1](#) on page 6 shows the conventions used in this guide.

Table 1: Conventions

Convention	Description
Italics	Emphasized terms or titles of books
Note:	Special instructions, suggestions, or warnings
<code>monospaced type</code>	Code elements such as commands, parameters, files, and directories

IP Infusion Product Release Version

Each integer in release number indicates Major, Minor, and Maintenance release versions. Build numbers that follow the release numbers are for internal tracking and verification of the software build process and are visible to customers as part of the software version number.



Product Name: IP Infusion Product Family

Major Version: New customer-facing functionality that represents a significant change to the code base; including, a significant marketing change or direction in the product.

Minor Version: Enhancements or extensions to existing features, changes to address external needs, or internal improvements might be motivated by improvements to satisfy new sales regions or marketing initiatives.

Maintenance Version: A collection of product bugs or hotfixes usually scheduled every 30 or 60 days, based on the number of hotfixes.

Related Documentation

For information about installing OcNOS, see the *Installation Guide* for your platform.

Feature Availability

Each OcNOS SKU contains a set of supported features. For a list of available features based on the SKU that you purchased. Refer to the *Feature Matrix*.

Migration Guide

Check the *Migration Guide* for necessary configuration changes before migrating from one version of OcNOS to another.

IP Maestro Support

Monitor devices running OcNOS Release 6.3.4-70 and above using IP Maestro software.

Technical Support

IP Infusion maintains an online technical support site that provides a variety of technical support programs for licensed OcNOS customers at the [Technical Assistance Center](#).

Customers and partners enjoy full access to the support website. The site allows customers and partners to open technical support calls, update open calls with new information, and review the status of open or closed calls. The password-protected site includes technical documentation, Release Notes, and descriptions of service offerings.

Technical Sales

Contact the IP Infusion sales representative for more information about the OcNOS solution.

Technical Documentation

For core commands and configuration procedures, visit: [Product Documentation](#).

For training videos, visit: [OcNOS Free Training Videos](#).

For a list of supported platforms and SKUs of OcNOS features, refer to the [OcNOS Feature Matrix](#).

Disclaimer

The global documentation site is evolving to provide an enhanced website user experience for select topics included in this release. Some guides are now available outside the existing documentation library and can be accessed directly from custom documentation landing pages. These guides offer robust in-built search functionality.

For the latest documentation, visit the product-specific documentation landing page and select the relevant guide.

Comments

If you have comments, or need to report a problem with the content, contact techpubs@ipinfusion.com.

CHAPTER 1 Debugging and Logging

OcNOS has a comprehensive debugging and logging facility in various protocols and components. This chapter describes how to start and stop debugging and logging. The protocol debug commands are in the corresponding command reference sections.

About Debugging

In OcNOS, every protocol has debug commands that log parameter-specific information. For example, using the `debug ospf nsm` command results in OcNOS writing all messages exchanged between OSPF and NSM such as interface, bandwidth, and address updates.

Note: By default, Syslog is enabled.

You can direct the output from the debug command to:

- Standard output (stdout)
- A file

OcNOS generates debug output until the `no` form of the `debug` command is given.

Start Debugging Output

To start debugging output, turn on the debug options by giving the relevant `debug` command and enable logging (logging level daemon-name <0-7>). For example:

```
> enable
# configure terminal
(config)# log file <filename>
(config)# debug <protocol> (parameter)
(config)# logging level <protocol> <0-7>
(config)# exit
```

Log to Standard Output

To direct debugging output to stdout, give the `terminal monitor` command.

```
# terminal monitor
```

This is a sample output of the `debug cml events` command displayed on the terminal:

```
7001-PEER#show debugging cml
CML terminal debugging status:
CML event debugging is on
CML smi debugging is on
7001-PEER#
```

Log to a File

To send debugging output to a file:

1. Use the `log file` command and specify the path and file name where the information is to be logged.

When logging to a file, you can simultaneously log to stdout by using the `terminal monitor` command.

2. Use the `no` form of the command to turn off logging to a file:

```
(config)# no log file (filename)
```

Stop Debugging

To turn off debugging, use the `no debug` command. When a protocol is specified with the `no debug` command, debugging is stopped for the specified protocol. To stop all debugging, use the `all` parameter with these commands.

```
(config)# no debug bgp events
```

Debug Modes

Debug commands act differently depending on the mode within which they are entered.

- In configure mode, debug session commands persist across switch reboots (if configuration is saved), and can be viewed in the running configuration using the `show running-config` command.
- In privileged exec mode, debug session commands do not appear in the running configuration, and can be viewed only by entering the `show debug` command. Additionally, exec debug commands do not persist across switch reboots.

CHAPTER 2 Debugging Kernel Crash

The kernel dump tools (kdump-tools) package facilitates the configuration and management of kernel crash dumps in Linux systems. It automates the setup of kdump, a mechanism that captures the system memory (vmcore) when the kernel encounters a critical failure. This captured vmcore is invaluable for post-mortem analysis and debugging.

Kernel Dump Extraction and Analysis

To analyze a kernel memory dump, extract the relevant files and run the crash analysis utility.

Note: This utility does not function if the system has less than 7 GB of memory.

With Technical Support Archive

1. Create a technical support archive using CLI.

```
#cmlsh
>en
#show cores
#show techsupport all
#exit
The system saves the archive at:
/var/log/OcNOS_tech_support_all_<DATE>_<TIME>.tar.gz
```

2. Extract the technical support archive and core dump files.

```
#tar -xf /var/log/OcNOS_tech_support_all_<date>_<time>.tar.gz
#tar -xf core_kdump_<timestamp>.tar
This extracts the vmcore at:
./<timestamp>/dump.<timestamp>
```

3. Extract debug-enabled kernel image (vmlinux tar file).

```
cd /lib/debug/
tar -xf vmlinux-<kernel_version>.tar.gz
```

4. Analyze the dump using crash.

```
crash /lib/debug/vmlinux /path/to/vmcore
```

Without Technical Support Archive

1. Extract the kdump archive (tar files).

```
cd /var/log/crash/cores/
tar -xf core_kdump_<timestamp>.tar
```

2. Extract debug-enabled kernel image (vmlinux tar file).

```
cd /lib/debug/
tar -xf vmlinux-<kernel_version>.tar.gz
```

3. To analyze the dump, run the crash command:

```
crash /lib/debug/vmlinux /path/to/vmcore
#crash /lib/debug/vmlinux /var/log/crash/cores/<timestamp>/dump.<timestamp>
```

crash 8.0.6

Copyright (C) 2002-2025 Red Hat, Inc.

Copyright (C) 2004, 2005, 2006, 2010 IBM Corporation

Copyright (C) 1999-2006 Hewlett-Packard Co

Copyright (C) 2005, 2006, 2011, 2012 Fujitsu Limited

Copyright (C) 2006, 2007 VA Linux Systems Japan K.K.

Copyright (C) 2005, 2011, 2020-2024 NEC Corporation

Copyright (C) 1999, 2002, 2007 Silicon Graphics, Inc.

Copyright (C) 1999, 2000, 2001, 2002 Mission Critical Linux, Inc.

Copyright (C) 2015, 2021 VMware, Inc.

This program is free software, covered by the GNU General Public License, and you are welcome to change it and/or distribute copies of it under certain conditions. Enter "help copying" to see the conditions. This program has absolutely no warranty. Enter "help warranty" for details.

GNU gdb (GDB) 10.2

Copyright (C) 2021 Free Software Foundation, Inc.

License GPLv3+: GNU GPL version 3 or later <<http://gnu.org/licenses/gpl.html>>

This is free software: you are free to change and redistribute it.

There is NO WARRANTY, to the extent permitted by law.

Type "show copying" and "show warranty" for details.

This GDB was configured as "x86_64-pc-linux-gnu".

Type "show configuration" for configuration details.

Find the GDB manual and other documentation resources online at:

<<http://www.gnu.org/software/gdb/documentation/>>.

For help, type "help".

Type "apropos word" to search for commands related to "word".

KERNEL: /lib/debug/vmlinux [TAINTED]

DUMPFIL: /var/log/crash/cores/202506160433/dump.202506160433 [PARTIAL DUMP]

CPUS: 4

DATE: Mon Jun 16 04:33:22 UTC 2025

UPTIME: 09:59:00

LOAD AVERAGE: 0.18, 0.29, 0.40

TASKS: 229

NODENAME: DUT1

RELEASE: 6.1.76-g8720581cb

VERSION: #1 SMP PREEMPT_DYNAMIC Wed Mar 5 11:17:59 UTC 2025

MACHINE: x86_64 (2400 Mhz)

MEMORY: 16 GB

PANIC: "Kernel panic - not syncing: sysrq triggered crash"

PID: 7826

COMMAND: "bash"

TASK: ffff888106216e00 [THREAD_INFO: ffff888106216e00]

CPU: 0

STATE: TASK_RUNNING (PANIC)

crash>

crash>

crash> **bt**

PID: 4534 TASK: ffff888105de8000 CPU: 0 COMMAND: "bash"

```
#0 [ffffc900005afc78] machine_kexec at ffffffff81056fe1
#1 [ffffc900005afcc8] __crash_kexec at ffffffff8113bf92
#2 [ffffc900005afd88] panic at ffffffff81e5878a
#3 [ffffc900005afe08] sysrq_handle_crash at ffffffff81733781
#4 [ffffc900005afe10] __handle_sysrq_cold at ffffffff81e81771
#5 [ffffc900005afe40] write_sysrq_trigger at ffffffff8173422f
#6 [ffffc900005afe50] proc_reg_write at ffffffff8135a050
#7 [ffffc900005afe68] vfs_write at ffffffff812d5c52
#8 [ffffc900005aff00] ksys_write at ffffffff812d6246
#9 [ffffc900005aff38] do_syscall_64 at ffffffff81ebf182
#10 [ffffc900005aff50] entry_SYSCALL_64_after_hwframe at ffffffff820000dc
RIP: 00007f59a8dee2c0 RSP: 00007fca5d70058 RFLAGS: 00000202
RAX: ffffffffda RBX: 0000000000000002 RCX: 00007f59a8dee2c0
RDX: 0000000000000002 RSI: 0000555df62882c0 RDI: 0000000000000001
RBP: 0000555df62882c0 R8: 0000000000000007 R9: 0000000000000073
R10: 0000000000000000 R11: 0000000000000202 R12: 0000000000000002
R13: 00007f59a8ec9760 R14: 0000000000000002 R15: 00007f59a8ec49e0
ORIG_RAX: 0000000000000001 CS: 0033 SS: 002b
```

crash>

Basic Commands in Crash

After launching the crash shell (indicated by the crash> prompt), use the following commands to inspect the system state at the time of the crash:

bt: Displays the stack trace for all tasks. This helps identify where each process was executing when the crash occurred.

ps: Lists all processes and their statuses, providing insight into the system's process table at the moment of the crash.

vm: Shows virtual memory information, useful for diagnosing memory-related issues.

`files`: Displays open files for a specific process, aiding in understanding resource utilization.

`help`: Lists all available commands within the crash utility for further exploration.

CHAPTER 3 Layer 2 Switching

This chapter contains steps to resolve Layer 2 switching issues.

Spanning Tree Protocol

This section shows how to resolve Spanning Tree Protocol (STP) issues.

Symptom/Cause	Solution
STP convergence Interface might be shutdown. Incorrect command syntax BPDU is getting dropped	Use these commands: <code>show spanning-tree</code> Check the Root Id and Bridge ID. Non Root bridge should have the bridge id of the peer bridge <code>show spanning-tree statistics bridge 1</code> Check the bpdv transmitting and receiving. If receiving is zero means packet is not receiving. In this issue check the port state. Whether its up or down. <code>show interface description</code> Check the interface status. <code>show running-config</code> Verify the whether the command configured successfully.
Port flapping Might be receiving the superior and inferior BPDU simultaneously Handling of the BPDU is improper	Use these commands: <code>debug mstp all</code> <code>debug mstp cli</code> <code>debug mstp packet rx</code> <code>debug mstp packet tx</code> <code>debug mstp protocol detail</code> <code>debug mstp timer detail</code>
BPDU sending and receiving Interface is admin shutdown No associated physical interfaces All associated physical interfaces are down	Use these commands: <code>show running-config</code> Verify if the interface is not shut down by the user. <code>show spanning-tree statistics bridge 1</code> Check the bpdv transmitting and receiving. If receiving is zero, it means packet is not receiving. <code>show spanning-tree statistics interface p8p1 bridge 1</code> Check the packets send and receive per interface wise.

Symptom/Cause	Solution
<p>RSTP convergence</p> <p>Interface is admin shutdown No associated physical interfaces All associated physical interfaces are down</p>	<p>Use these commands:</p> <pre>show running-config</pre> <p>Verify if the interface is not shut down by the user.</p> <pre>show spanning-tree statistics bridge 1</pre> <p>Check the bpdv transmitting and receiving. If receiving is zero means packet is not receiving.</p> <pre>show spanning-tree</pre> <p>In the output part please check the Root Id and Bridge ID. A non root bridge should have the bridge id of the peer bridge.</p> <pre>show spanning-tree statistics interface p8p1 bridge 1</pre> <p>Check the packets send and receive per interface wise.</p>
<p>MSTP convergence</p> <p>Interface is admin shutdown No associated physical interfaces All associated physical interfaces are down</p>	<p>Use these commands:</p> <pre>show running-config</pre> <p>Verify if the interface is not shut down by the user.</p> <pre>show spanning-tree statistics bridge 1</pre> <p>Check the bpdv transmitting and receiving. If receiving is zero means packet is not receiving.</p> <pre>show spanning-tree mst detail</pre> <p>Check the Root Id and Bridge ID. A non Root bridge should have the bridge id of the peer bridge.</p> <pre>show spanning-tree statistics interface p8p1 bridge 1</pre> <p>Check the packets send and receive per interface wise.</p>
<p>RPVST convergence</p> <p>Interface is admin shutdown No associated physical interfaces All associated physical interfaces are down</p>	<p>Use these commands:</p> <pre>show running-config</pre> <p>Verify if the interface is not shut down by the user.</p> <pre>show spanning-tree statistics bridge 1</pre> <p>Check the bpdv transmitting and receiving. If receiving is zero means packet is not receiving.</p> <pre>show spanning-tree rpvst detail</pre> <p>Check the Root Id and Bridge ID. A non Root bridge should have the bridge id of the peer bridge.</p> <pre>show spanning-tree statistics interface p8p1 bridge 1</pre> <p>Check the packets send and receive per interface wise.</p>

BPDU Guard

Symptom/Cause	Solution
BPDU is not handled properly Interface is admin shutdown No associated physical interfaces All associated physical interfaces are down	Use these commands: <code>show spanning-tree</code> This command have bpdu-guard configured a field. If its not showing configured (on) then its not get configured. <code>show running-config</code> Check the configuration of the root-guard. If it showing off means the bpdu guard is not configured. If it configured successfully then upon receiving the superior BPDU the port state will down. Once the port admin goes down then need to bring it up manually.

BPDU Filter

Symptom/Cause	Solution
BPDU is not handled properly Interface is admin shutdown No associated physical interfaces All associated physical interfaces are down	Use these commands: <pre>show spanning-tree</pre> This command have bpdud-filter configured a field. If its not showing configured (on) then its not get configured. <pre>show running-config</pre> Check the configuration of the root-guard. If configured successfully then upon receiving the superior BPDU the bpdud-filter configured port will loose its bpdud-filter and it will become normal port.

Root Guard

Symptom/Cause	Solution
BPDU is not handled properly Interface is admin shutdown No associated physical interfaces All associated physical interfaces are down	Use these commands: <pre>show spanning-tree</pre> This command have root guard configured a field. If its not showing configured then its not get configured. <pre>show running-config</pre> Check the configuration of the root-guard. On receiving superior BPDU then root guard port should move to the root-inconsistent. If its not behaving like this then handling of the packet is improper in the control plan.

VLANs

Symptom/Cause	Solution
Unable to create/delete VLAN with a bridge Bridge is not VLAN aware Incorrect command syntax Maximum number of vlans created already VLAN created or deleted already	Use these commands: <pre>show vlan brief</pre> Confirm if the VLAN is created. <pre>show running-config</pre> Verify the type of bridge configured and the number of VLANS already configured on the bridge.

Symptom/Cause	Solution
VLAN tagged/untagged packets are not egressed/ingressed properly Incorrect configuration on interface Interface may be in down state Interface may be in spanning-tree disabled state	Use these commands: <pre>show vlan brief</pre> Confirm if the VLAN is created. <pre>show running-config</pre> Verify the type of bridge configured and the number of VLANs already configured on the bridge. <pre>show interface IF_NAME</pre> View the input/output/dropped packet counters. <pre>clear counters IF_NAME</pre> Use this command to clear counter statistics for the interface.
VLAN interface is not operational Interface is admin shutdown No associated physical interfaces All associated physical interfaces are down	Use these commands: <pre>show running-config</pre> Verify if the interface is not shut down by the user. <pre>show vlan brief</pre> Check if the VLAN is bound to physical interface <pre>show interface</pre> Verify if the member interface is UP and RUNNING

Interfaces

Link Light for the Port does not come on

Symptom/Cause	Solution
No cable connected.	Connect cable from switch to a known good device.
Wrong Port	Make sure that both ends of the cable are plugged into the correct ports.
Device has no power	Ensure that both devices have power.
Wrong cable type	Verify the cable selection.
Bad cable	Swap suspect cable with known good cable. Look for broken or missing pins on connectors.
Loose connection	Check for loose connections. Sometimes a cable appears to be seated in the jack, but is not. Unplug the cable and reinsert it.
Patch Panels	Bypass the patch panel if possible to rule it out.
Media Convertors	Bypass the media convertor (e.g fiber-to-copper) if possible to rule it out.

Other

Symptom/Cause	Solution
Bad Port or Module Port or Interface or Module not enabled	Move the cable to a known good port to troubleshoot a suspect port or module.
Interface is down	Use this command: <pre>show interface brief</pre> Look for "Status" and "Reason". If the "Status" is down and "Reason" is AD (Administratively Down), it means that the administrator has disabled the interface. Bring the interface up: <pre>no shutdown</pre>
VLAN interface is down	Use these commands: <pre>show interface brief</pre> Check if the corresponding VLAN interface "Status" is up or down. Also check the value in the "Reason" column. <pre>show vlan brief</pre> Check if the "State" field to see if it is in "SUSPEND". Enable the van: <pre>vlan <vlan-name> bridge <bridge-id> state enable"</pre> Possible cause might be no <code>switchport</code> given on the interface attached to VLAN due to which van is detached from the interface. Enable <code>switchport</code> on that particular interface and attach the van and bridge again. <pre>show running-config interface <interface-name></pre> View the configuration.

Link Aggregation

Symptom/Cause	Solution
LACP packets not transmitted properly Interface is down Channel group is configured as passive at both the nodes	Use these commands: <pre>show lacp-counter</pre> Verify if there are any packets are sent out of interface. <pre>show interface IF_NAME description</pre> Check if the interface is up <pre>show running config</pre> To list the LACP running configuration
LACP packets not received properly Interface is down Channel group is configured as passive at both the nodes	Use these commands: <pre>show lacp-counter</pre> Verify if there are any packets are sent out of interface. <pre>show interface IF_NAME description</pre> Check if the interface is up

Symptom/Cause	Solution
LACP sync is not up Interface is down Channel group is configured as passive at both the nodes Link bandwidth mismatch MTU mismatch VLAN mismatch DUPLEX mismatch	Use these commands: <pre>show etherchannel detail</pre> Check the Actor system ID, partner system ID, interface coming under the aggregation and sync bit <pre>show etherchannel summary</pre> Check the sync bit. <pre>show interface IF_NAME description</pre> Check if the interface is up, Bandwidth, duplex and MTU. Check if Active mode is configured at one

Multi-Chassis Link Aggregation

Symptom/Cause	Solution
MLAG Domain adjacency between TORs to determine active / standby during switchover	Use these commands: <pre>show mlag domain summary</pre> Verify domain and MLAG sync configuration. <pre>show lagd mlag 1</pre> To list the properties / Neighbor partner info <pre>show lagd stp-config</pre> STP MLAG Configurations Present in LAGD
To list the MCCPDU received on the system/ mcec debugs	Use these commands: <pre>show mcec statistics</pre> <pre>clear mcec statistics</pre> Command to display /clear the mccpdu statistics <pre>show mlag stp-synchronization status</pre> To check STP MLAG Configuration / Sync status <pre>debug mcec (timer event hello info cli mac-sync stp all)</pre> Command to enable debug messages for MCEC module

LLDP

Symptom/Cause	Solution
LLDP packets not transmitted properly Interface is down LLDP agent not enabled LLDP tx not enabled	Use these commands: <pre>show lldp interface IF_NAME</pre> Verify if there are any packets discarded. <pre>show interface IF_NAME description</pre> Check if the interface is up

Symptom/Cause	Solution
LLDP packets not received properly Interface is down LLDP agent not enabled LLDP rx not enabled	Use these commands: <pre>show lldp interface IF_NAME</pre> Verify if there are any packets received in error. <pre>show interface IF_NAME description :</pre> Check if the interface is up
LLDP system name, system description not seen in the peer Interface is down LLDP agent not enabled LLDP rx/tx not enabled Basic management capabilities are not enabled.	Use these commands: <pre>show lldp interface IF_NAME neighbor</pre> Check the neighbor information <pre>show running config</pre> Check if system name/description are configured at the sending side Note: This applies to many other lldp capabilities like ieee-8021-org-specific, ieee-8023-org-specific, and med capabilities. The needed capabilities must be enabled on the device using lldp tlv-select.
All VLANs configured not displayed in the remote VLAN details of the peer device VLAN not configured for the interface	This is a special case where there are more than 82 vlans and the TLV size is increasing the MTU. As vlan details are encoded after encoding all the other TLVs. So whenever the buffer becomes full, the end of TLV is sent automatically and the rest of the vlan information is not displayed. This is the expected behavior.
LLDP med capabilities are not displayed in the neighbor Interface is not up LLDP agent not enabled Tx/rx not enabled LLDP tlv-type med capabilities not enabled LLDP med dev-type has to be end point class3 at least on one side. It doesn't sent its med capabilities if net-connect is enabled until it receives med info from the neighbor	Use these commands: <pre>show lldp interface IF_NAME</pre> Verify if tx/rx, med is enabled. <pre>show interface IF_NAME descriptio</pre> Check if the interface is up <pre>show lldp interface IF_NAME neighbor</pre> Verify the med related information of neighbor

802.1x

Symptom/Cause	Solution
Packets does not reach to device Loose connection between device and XSUPPLICANT	Following XSUPPLICANT scripts need to be modified based upon the connections (on which interface we need to allow EAPOL packets and on which interfaces we have to deny) and the IP address assigned. <pre>/usr/local/etc/1x/ md5-example.conf</pre> <pre>/usr/local/etc/1x/startup.sh</pre> <pre>/usr/local/etc/1x/startup2.sh</pre> The very first packet sent to device should be EAPOL packet. Data packets will be dropped if the first packet is not EAPOL.

Symptom/Cause	Solution
<p>Packets do not reach to that interface of device which is connected to radius server.</p> <p>Packets must be getting dropped at kernel level due to some malformed fields.</p> <p>The kernel must not be lifting those packets to the protocol level.</p> <p>The decoding of the packet could have some issue.</p> <p>The other interface is down.</p>	
<p>Packets do not reach to Radius Server</p> <p>Loose connection between device and RADIUS SERVER.</p> <p>IP Configured at device for RADIUS SERVER must not match the IP configured at the RADIUS SERVER's interface connected with device.</p>	
<p>Radius Server does not reply back to device</p> <p>The RADIUS SERVER must not be having the XSUPPLICANT details with it</p>	<p>Check following files:</p> <pre>/usr/local/etc/raddb/users</pre> <p>The entry corresponding to the mac address of the xsupplicant interface which is connected to device should be updated.</p> <pre>00:02:A5:4E:FF:83 Auth-Type := eap, User-Password == "00:02:A5:4E:FF:83" Tunnel-type:0 = 13, Tunnel-Medium-Type:0 = 6, Tunnel-Private-Group-ID:0 = 201, Reply-Message = "Hello, %u"</pre> <pre>/usr/local/etc/raddb/clients.conf</pre> <p>Ip should be updated</p> <pre>client 10.10.10.40/24 { secret = authd shortname = device }</pre>
<p>Radius Server replies back to device with "access challenge"</p> <p>Authentication issues at RADIUS SERVER</p>	<p>Credentials must be verified properly.</p> <p>XSUPPLICANT should retrigger the EAPOL packets</p>
<p>FDB does not get updated at device accordingly</p> <p>Problem with authentication of the particular MAC or port.</p>	<p>Retrigger the EAPOL packet and check above mentioned scenarios one by one to find whether packets are getting dropped somewhere or there is some problem with the credential matching.</p>

VLAN Cross-Connect

Symptom/Cause	Solution
Cross-connect is not up Check cross-connect status Interface is down	Use these commands: <pre>show cross-connect</pre> Verify cross-connect is up. <pre>show interface IF_NAME</pre> Check if the interface is up
Cross-connect packets not received properly Interface is down Check cross-connect status Check interface stats	Use these commands: <pre>show interface IF_NAME</pre> Check if the interface is up <pre>show cross-connect</pre> Verify cross-connect is up. <pre>show interface IF_NAME counters rate mbps</pre> Verify interface stats are incrementing properly

IGMP Snooping

Symptom/Cause	Solution
IGMP Report not learned IGMP Snooping disabled globally. IGMP Snooping disabled on vlan interface. Port on which report is sent in discarding state (xSTP). Destination MAC of the frame does not correspond to Destination IP multicast address. TTL of the IP packet not 1. For IGMPv3 report destination address of IP packet not 224.0.0.22.	Use these commands: <pre>show igmp snooping interface</pre> Check if IGMP snooping is enabled globally and on interface. If underlying L2 interface is in forwarding state in spanning tree then it will be listed below vlan interface. <pre>show running-config</pre> Confirm if igmp snooping is disabled globally or on interface.
IGMP reports forwarded by switch Snooping switch sending reports only when it receives an IGMP query	Use these commands: <pre>show igmp snooping mrouter VLAN-IFNAME</pre> Confirm mrouter interface learned
IGMP Queries sent with Source IP as zero Switch sending queries with SIP as 0 when it is configured as querier, otherwise queries received from router will be forwarded as is to all interfaces bound to that VLAN.	Use these commands: <pre>show igmp snooping interface VLAN-IFNAME</pre> Check if querier is enabled on the interface <pre>show running-config</pre> Confirm the configuration of querier.

Symptom/Cause	Solution
<p>Traffic not forwarded from the source specified in the IGMP membership Report.</p> <p>The resulting state of the source might be in exclude list OR The group might have been expired</p>	<p>Use these commands:</p> <pre>show igmp snooping groups detail</pre> <p>Confirm if the source is in exclude list or include list of the Group and to check expiry timer, if it is dynamic group then it will not be shown in this command after expiry.</p>

PVLAN

Symptom/Cause	Solution
<p>Private VLAN cannot be created</p> <p>VLAN is not created. Incorrect command Syntax Given Bridge Group is incorrect.</p>	<p>Use these commands:</p> <pre>show vlan brief</pre> <p>Confirm if the VLAN is created.</p> <pre>show running-config</pre> <p>Confirm the bridge type created. Private Vlan cannot be created for provider bridges</p>
<p>With the issue of “no shutdown” command, vlan interface is not activated</p> <p>Due to the nature of Private VLANs, you cannot activate the VLAN interface for isolated or community VLANs. You can only activate the VLAN interface that belongs to primary VLAN.</p>	<p>Use this command:</p> <pre>show vlan private-vlan bridge BRIDGE_GROUP</pre> <p>Confirm the type of Private vlans configured</p>
<p>Not able to associate a Secondary PVLAN to the Primary PVLAN</p> <p>Secondary Private VLAN type is Isolated and one isolated PVLAN is already associated with the particular Primary PVLAN. Secondary PVLAN is already associated with some other Primary PVLAN.</p>	<p>Use this command:</p> <pre>show vlan private-vlan bridge BRIDGE_GROUP</pre> <p>Confirm the types of PVLANs configured and associations of Secondary PVLANs with Primary PVLANs.</p>
<p>Not able to configure an interface as private-vlan host-port</p> <p>Interface is not access-port</p>	<p>Use this command:</p> <pre>show interface IFNAME</pre> <p>Verify the port mode. Only Access ports can be configured as host-ports.</p>

CHAPTER 4 Layer 3 Routing

This chapter contains steps to resolve Layer 3 routing issues.

Missing Route

When end-to-end connectivity is a problem:

1. Check if you can ping your own interface and other devices on directly connected networks.
2. Check if the route is installed in the system:

```
show ip route
```

3. Check the configuration:

```
show running-config
```

4. Check if the interface is in the correct network:

```
show ip interface brief
```

RIP

This section contains steps to resolve RIP issues.

No RIP Adjacency

Symptom/Cause	Solution
Interface administratively shut down	<pre>show ip interface brief</pre> <p>Make sure that the interface is not administratively shut down.</p> <p>Bring up the interface using the <code>no shutdown</code> command, if needed. Use the <code>show interface</code> command to verify that the interface is up.</p>
RIP not enabled on the interface	<pre>show ip rip interface</pre> <p>Verify that RIP is enabled for the interface.</p> <pre>network</pre> <p>Enable RIP on the interface</p>
Interface configured as a passive interface	<p>Make sure that the interface is not configured as a passive interface using the <code>show run</code> command</p> <p>If the interface is configured as passive, give the <code>no passive interface</code> command</p>
RIP advertisements not sent and received on the interface	<p>Use a packet sniffer (such as Ethereal or TCP dump) or log messages to verify the RIP advertisements. To turn on logging:</p> <pre>debug rip event or debug rip packet</pre>

Symptom/Cause	Solution
One router configured as RIPv1 and the other router as RIPv2	Configure the router running RIPv2: <pre>ip rip send version 1-compatible ip rip receive version 1 2</pre>

BGP

This section contains steps to resolve BGP issues.

A common mistake is incomplete meshing of IBGP. IBGP must be fully meshed if Route Reflectors or confederations are not used. There is no way of learning automatically from the network.

Symptom/Cause	Solution
IBGP neighbors not coming up BGP sessions not formed IBGP neighbors stuck in active/ idle state	<p>Check for IP reachability between the BGP neighbors using ping command. If ping fails, configure IGP in the network for IBGP reachability or configure static route.</p> <p>If IP reachability is there, check for the BGP configuration on both sides. Ensure that the AS numbers on both sides (peers) and the local AS / remote AS are same for IBGP. Use “show run bgp” to verify the configuration.</p> <p>Verify that the bgp router ids are different on both sides using “show run bgp” or show ip bgp Summary” command.</p> <p>If the IBGP peering is over any address family other than IPV4/unicast, explicit neighbor activation need to be configured for that address-family. Eg:- neighbor < peer address> activate (for address-families other than default/ipv4 unicast)</p> <p>If configurations are correct, check for tcp connection status using “netstat -tln” command at shell prompt. If tcp connections are not in CONNECTED State, check whether the BGP neighborhood is configured over loopback interface or normal/physical interface.</p> <p>If the neighborhood/peering is over loopback interface on both sides, then explicit “neighbor <peer loopback address> update source loopback” need to be configured on both peers. If loopback interface is used only on one side (peer), the above configuration need to be applied on the other side (peer).</p> <p>If bgp md5 authentication is enabled, ensure that the passwords match on both sides. Verify that “debug ip bgp” does not throw any message like “No md5 digest *****” or “Invalid md5 digest”.</p> <p>IBGP session should come up and verify the IBGP neighbors are in Established state using the command “show ip bgp summary” or “show ip bgp neighbors”.</p>

Symptom/Cause	Solution
<p>EBGP neighbors not coming up EBGP sessions not formed EBGP neighbors stuck in Active/idle state</p>	<p>Check for IP reachability between the BGP neighbors using ping command. If ping fails, verify IP address config and configure static route if the ebgp peers are not directly connected.</p> <p>Verify the EBGP config is correct using the “show run bgp” command. For EBGP the AS numbers are different on both side and the local AS/remote AS should be reverse on both peers.</p> <p>Verify that the bgp router ids are different on both sides, using “show run bgp” or show ip bgp Summary” command.</p> <p>If the EBGP peering is over any address family other than IPV4/unicast, explicit neighbor activation need to be configured for that address-family. Eg:- neighbor < peer address> activate (for address-families other than default/ipv4 unicast).</p> <p>If configurations are proper, check for tcp connection status using “netstat -tln” command at shell prompt. If tcp connections are not in CONNECTED State, check whether the BGP neighborship is configured over loopback interface or normal/physical interface.</p> <p>If the neighborship/peering is over loopback interface on both sides then explicit “neighbor <peer loopback address> update source loopback” need to be configured on both peers. If loopback interface is used only on one side (peer), the above config need to be applied on the other side (peer).</p> <p>If EBGP neighbors are not directly connected, configure the bgp multihop config with appropriate ttl value. The “ttl” value should be greater than or equal to number of hops, the bgp peer is away. Eg: neighbor < peer address> ebgp-multihop <ttl value>.</p> <p>If bgp md5 authentication is enabled, ensure that the passwords match on both sides. Verify that “debug ip bgp” does not throw any message like “No md5 digest *****” or “Invalid md5 digest “.</p> <p>EBGP session should come up and verify the EBGP neighbors are in Established state using the command “Show ip bgp summary” or “show ip bgp neighbors”.</p>

Symptom/Cause	Solution
BGP peering flaps BGP session flaps BGP establishes and drops	<p>Check for the BGP peer flap cause/reason from the notification field in "show ip bgp neighbor < peer address>".</p> <p>If the reason is "Hold timer expiry" then keepalives are not received/ processed.</p> <p>If keepalives are not received, check the peer router "show ip bgp summary" and verify that the sent msg field is getting incremented.</p> <p>If Keepalives are received but not processed then keepalives are getting stuck behind the Update messages. In other words the update messages are taking longer time to process or the system is configured beyond the claimed BGP limits. The keepalives getting queued can be verified from the "show ip bgp summary" output In Q/OutQ field. If the system is configured with in the scalability limits, please call up support team for help.</p> <p>If the reason is "Administer configured" then the administrator has modified the fields which could trigger BGP session reset. Eg:- BGP Graceful restart configuration. If the flapping is continuous, check for interface flap, bad connectors and bgp process unstable at other end and traffic and rate limiting parameters. Check for interoperability issues too.</p> <p>If the reason is "Malformed packet" then BGP has received a invalid packet from the peer. This is mostly caused by oversized BGP packet as BGP TLVs are variable length. Please call up the support team for help.</p> <p>If none of these, then the issue could be PATH MTU related. Try pinging the peer with various size packet like above 1500 bytes etc. Normal ping succeeds but extended ping fails. Trouble shooting the path MTU issues is beyond the scope of this document. Please call up the support team.</p>

Symptom/Cause	Solution
<p>Missing routes Routes not seen in BGP table Irregular network connectivity</p>	<p>For IBGP peers, the topology should be fully meshed if no Route Reflectors are configured. If Route Reflectors are used, ensure IBGP is established between all iBGP peer and the RRs and between the RRs. If confederations are configured to avoid IBGP mesh, ensure that the confederation configuration is proper.</p> <p>Route origination issues - If BGP is not originating the route and If the route is injected via "network prefix" command, Verify that the exact match (not best match) for the route exists in the RIB (IP RIB) using the "show ip route include prefix ". If exact match does not exist, add a static route for the exact prefix and verify the static route is active in the RIB. BGP should originate the route now.</p> <p>Aggregate route origination issue - If route aggregation is configured and aggregate route is missing in BGP table ("Show ip bgp") but available in IP RIB ("show ip route"), then the configuration is not complete. For route aggregation, route needs to be there in BGP table. Add the aggregated route in BGP via the "network <aggregate prefix/mask>". Now the aggregated route can be seen in the "Show ip bgp" command also.</p> <p>Originator/cluster id collisions - If Router reflectors are used and topology is proper, still routes not populated. The problem could be clashing cluster ids and router/originator ids. To debug this issue, check for the BGP table from the originator of the route using "show ip bgp " Check whether the route is advertised to the peer using the "show ip bgp neighbor <peer address> advertised-routes" command. Check whether the route is received at the peer using the "show ip bgp neighbor <peer address>" at the peer router. If route is not received, check for the access-list configured on the peer router. If access list are also permitting the route, enable bgp debugs and clear the bgp session using "clear ip bgp <peer-address> out "and figure out the reason for the denial of the route. If the reason for route ignore/drop is "originator id same as router id" Change the router ids in one of the router. If the reason for route ignore/drop is "reflected from same cluster" then the All RRs are not peered with each other. Ensure ibgp peering between all RRs.</p> <p>Filtering issues - If Routes are not seen with proper configuration, look for the filters/route-maps configured at the inbound and outbound of each peer. Verify that there are no typos in the filters or in the expressions used for filtering. Verify that access-list configured are for exact match.</p> <p>Community config/mismatch issues – verify the configuration has "neighbor <peer address> send community" to advertise the community strings. Ensure that the community match strings configured via route-map are matching/correct/proper.</p> <p>IBGP meshing issues – Routes are marked as valid and internal but not marked as the best in the "show ip bgp" output. The nexthop for the route (mostly ebgp route) is not reachable and hence not advertised to peer. Configure static route or IGP for nexthop reachability. Another reason for missing route in an IBGP network is that not fully meshed. Ensure that IBGP is fully meshed or use RR + peer group config for ease of maintenance</p>
<p>High CPU utilization</p>	<p>In case of recursive BGP routes, ie where BGP nexthop is also learned via BGP, recursive look up may end up in invalid/illegal. IT is advisable to learn the BGP Nexthop via IGP also. Route flapping could be one cause for High CPU utilization. Routing loops caused by lack of IBGP meshing. Use RR or ensure IBGP full meshing.</p>

OSPFv2

This section contains steps to resolve OSPFv2 issues.

Neighborship Formation

Symptom/Cause	Solution
OSPF not enabled on the interface or the corresponding interface is down.	Use network command to enable ospf on a particular network address. Do “no shutdown” on the corresponding interface. Use these commands: <pre>show ip ospf interface show ip ospf show running-config ospf show ip ospf neighbor show ip ospf neighbor detail</pre>
OSPF Interface network address mismatch or subnet mask mismatch	Use these commands to find the output/find the mismatch: <pre>show ip ospf interface show ip interface brief show running-config ospf show ip ospf neighbor show ip ospf neighbor detail.</pre>
Hello packet attribute mismatch I.e. interface hello interval or dead interval mismatch	Use these commands to find the output/mismatch: <pre>show running-config ospf show ip ospf interface debug ospf nsm debug ospf packet hello detail</pre>
OSPF interface configured as passive interface. Passive interfaces will not send hello packets.	Use these commands: <pre>show ip ospf interface show ip ospf show running ospf</pre>
OSPF option field mismatch i.e. Ospf External Routing Capability mismatch due to stub area mismatch.	Configure same stubby-area. Use these commands: <pre>debug ospf packet hello show ip ospf show running ospf</pre>
OSPF gets stuck in two-way state when there is no DR/BDR elected i.e. both routers configured with priority 0	One router should be configured with a priority other than 0. Use these commands: <pre>show ip ospf neighbor show running ospf</pre>

Symptom/Cause	Solution
OSPF can get stuck in Exstart/ exchange state when the maximum transmission unit (MTU) do not match	Use these commands: show ip ospf interface show running ospf show ip ospf neighbor
OSPF network type mismatch can cause timers (hello-interval) mismatch	Use these commands: show ip ospf interface debug ospf packet hello show running ospf

Virtual Link Neighborship Formation

Symptom/Cause	Solution
OSPF same router-id and loopback address	Use these commands: show ip ospf show running-config ospf show ip ospf interface
OSPF area mismatch for virtual-link	Use these commands: show running-config ospf show ip ospf virtual-link
OSPF virtual-link router-id mismatch causing virtual-link down	Use these commands: show running-config ospf show ip ospf virtual-link show ip ospf interface show ip ospf neighbor
OSPF not enabled on the interface or the corresponding interface is down causing vlink down	Use these commands: show ip ospf show running-config ospf show ip ospf virtual-link show ip ospf interface

Multi-area Adjacency Formation

Symptom/Cause	Solution
OSPF Multi-Area Adjacency happens only if the adjacency over the primary link is up.	Use these commands: <pre>show ip ospf neighbor show ip ospf interface</pre>
Multi-area configuration mismatch, such as area-id or process-id mismatch.	Use these commands: <pre>show running-config show ip ospf multi-area-adjacencies show ip ospf neighbor</pre>
Multi-Area link remains down, if the Router is not an "ABR"	Use these commands: <pre>show ip ospf show ip ospf multi-area-adjacency.</pre>
Multi-Area link will be removed if the primary link is down	Use these commands: <pre>show running-config show ip ospf interface show ip ospf multi-area-adjacency</pre>
For multi access networks, it is required to specify the neighbor's interface address. Multi-area configuration mismatch such as when the neighbor's address does not match any of the OSPF networks or the neighbor's address is incorrect.	Use these commands: <pre>show running-config ospf show ip ospf multi-area-adjacency</pre>
Multi-area adjacency not applicable for the area on which virtual-link is enabled and vice-versa.	Use this command: <pre>show running-config ospf</pre>

Graceful Restart (using Link Local Signaling) Neighborhood Formation

Symptom/Cause	Solution
Daemon restarted after maximum router dead interval of all OSPF interfaces on restarting node	Use this command on restarting node: <pre>show ip ospf interface</pre>

Symptom/Cause	Solution
Capability mismatch between restarting and helper node	<p>Enable <code>capability lls</code> in router node.</p> <p>Use these commands on both nodes:</p> <pre>show ip ospf show running-config ospf</pre>
Configurations not saved before restart	<p>Use any one of these commands and check the saved file in <code>/usr/local/etc</code>:</p> <pre>write <file_path> write file write memory write terminal</pre>
Two routers on same segment get restarted at same time	<p>Use these commands on restarting node:</p> <pre>debug ospf nfsm debug ospf packet hello recv detail show ip ospf neighbor</pre>
Helper router got its LSDB changed while helping	<p>Use these commands on restarting node:</p> <pre>debug ospf nfsm debug ospf packet hello recv detail show ip ospf neighbor</pre>
"Resync timeout timer" is not sufficient to sync the whole database	<p>Use these commands on restarting node:</p> <pre>debug ospf nfsm show ip ospf neighbor</pre>

OSPFv3

This section contains steps to resolve OSPFv3 issues.

Neighborship Formation

Symptom/Cause	Solution
OSPF not enabled on the interface or the corresponding interface is down	<p>Use network command to enable OSPFv3 on a particular network address. Do no shutdown on the corresponding interface. Use these commands:</p> <pre>show ipv6 ospf interface show ipv6 ospf show running-config ospfv3 show ipv6 ospf neighbor show ipv6 ospf neighbor detail</pre>

Symptom/Cause	Solution
Hello packet attribute mismatch i.e. interface hello interval or dead interval mismatch	Use these commands to find the output/mismatch: <pre>show running-config ospfv3 show ipv6 ospf interface debug ipv6 ospf n fsm debug ipv6 ospf packet hello detail</pre>
OSPF option field mismatch i.e. Ospf External Routing Capability mismatch due to stub area mismatch	Configure the same stubby-area. Use these commands: <pre>debug ipv6 ospf packet hello show ipv6 ospf show running-config ospfv3</pre>
OSPF gets stuck in two-way state when there is no DR/BDR elected i.e. both routers configuration with priority 0	One router should be configured with a priority other than 0. Use these commands: <pre>show ipv6 ospf neighbor show running-config ospfv3</pre>
OSPF can get stuck in Exstart/Exchange state when the maximum transmission unit (MTU) does not match	Use these commands: <pre>show ipv6 ospf interface show running-config ospfv3 show ipv6 ospf neighbor</pre>
OSPF network type mismatch can cause timers (hello-interval) mismatch	Use these commands: <pre>show ipv6 ospf interface debug ipv6 ospf packet hello show running-config ospfv3</pre>
OSPF same router-id and loopback address	Use these commands: <pre>show ipv6 ospf show running-config ospfv3 show ipv6 ospf interface</pre>
Ospf area mismatch for virtual-link	Use these commands: <pre>show running-config ospfv3 show ipv6 ospf virtual-link</pre>
OSPF virtual-link router-id mismatch cause virtual-link down.	Use these commands: <pre>show running-config ospfv3 show ipv6 ospf virtual-link show ipv6 ospf interface show ipv6 ospf neighbor</pre>

Symptom/Cause	Solution
OSPF vlink Global remote address is NULL causing virtual-link adjacency down.	Use these commands: <pre>show ipv6 ospf interface show ipv6 ospf virtual-link</pre>
OSPF not enabled on the interface or the corresponding interface is down causing vlink down.	Use these commands: <pre>show ipv6 ospf show running-config ospfv3 show ipv6 ospf virtual-link show ipv6 ospf interface</pre>

VRRP

This section contains steps to resolve VRRP issues.

Symptom/Cause	Solution
Incorrect VRRP States	<p>Make sure the interfaces are up and running by using the <code>show interface</code> command. If the interface is down, give the <code>no shutdown</code> command in <code>interface</code> mode to bring up the interface.</p> <p>Or</p> <p>Use the <code>ifconfig IFNAME up</code> command to bring up the interface.</p> <p>Make sure the interface has an IP address in the that matches the subnet of the VRRP session's virtual-ip. If not exist, add the ip address in <code>interface</code> mode using the command "<code>ip address A.B.C.D/M</code>"</p> <p>The configured IP address is displayed in the output of "<code>show running-config interface <IFNAME></code>"</p> <p>Make sure that both VRRP routers can reach each other by pinging. If both routers cannot reach each other, check the network connections for the default Master and default Backup routers.</p> <p>Check the advertisement interval on Master and Backup routers. The advertisement interval must be the same on both. The default advertisement interval = 1 second.</p> <p>Use the <code>advertisement-interval</code> command in router mode to configure the advertisement interval.</p>

BFD

This section contains steps to resolve BFD issues.

Symptom/Cause	Solution
BFD session not created	<p>Use the <code>show run</code> command to make sure that the interface is not administratively shutdown. If <code>shutdown</code> is configured, remove this configuration with the <code>no shutdown</code> command.</p> <p>Use the <code>show interface</code> command to make sure that the interface is up.</p> <p>Check if minimum BFD configuration exists for the respective client. Client list includes static route, RIP, OSPF, BGP, ISIS. For configuration refer BFD command reference section.</p> <p>Check if BFD is administratively down. If BFD is admin-down, session will show state as Admin-Down for 1 minute and then session information is deleted. No session will be seen after that.</p> <p>The next hop for which BFD is enabled should be reachable. For example in OSPF, OSPF neighbor adjacency should be formed for BFD session to be created to monitor the neighbor next hop. This applies to all clients.</p> <p>If configuration is as expected, then enable BFD debug logs and check if packets are received, transmitted for that session. To turn on logging, enter the <code>debug bfd all</code> command</p>
BFD session down	<p>Check the diagnostic reason which can be Neighbor Signaled Down (Remote admin Down) or Control Detect Expiry (non-receipt of hello packets from neighbor for detect multiplier time).</p> <p><code>show bfd session detail</code></p>

IS-IS

Adjacency Problems

Normally, the causes for IS-IS adjacency related problems are because of link failures and configuration mistakes. The link failures can be easily detected by checking `show isis interface` command.

The first step to troubleshoot IS-IS adjacency is the `show clns neighbors` command which displays the basic configuration. The `show clns is-neighbors` command can also be used but it lists only neighboring routers, while the former command lists all types of adjacencies both for IS-IS and for ES-IS.

Symptom/Cause	Solution
<p>Expected adjacencies not formed or not seen in the output of <code>show clns neighbors</code></p> <p>The interfaces might have been administratively shut down</p> <p>The IS-IS configurations are incorrect.</p> <p>The levels configured on the interfaces might be mismatching</p> <p>The IP subnet configured is incorrect.</p> <p>Duplicate system ID is configured in IS-IS area.</p>	<p>Check for link failures. We can verify whether all the interfaces on which IS-IS adjacency is enabled are in UP state. This can be confirmed by checking the output of <code>sh ip interface brief</code> command. Once the state of the interface is verified to be in UP state, do a ping to the other end and check the connectivity is proper. If the ping fails, that means the physical connectivity problem exists, which should be resolved before going further.</p> <p>If the link is fine, check the IS-IS configurations. IS-IS is enabled in two steps, first is creating isis process as shown below:</p> <pre>router isis net 49.0010.0000.0001.00</pre> <p>A misconfigured NET command could also lead to IS-IS adjacency problem. Next step is to attach this process to the appropriate interfaces with command <code>ip router isis</code>.</p> <p>By default, IS-IS router processes have Level 1-2 capability. We can configure IS-IS routers to be level-1 only or level-2 only by using the <code>is-type</code> command. If the levels are mismatching between two directly connected routers adjacency is not formed. By correcting the levels the adjacency issue can be resolved.</p> <p>If the directly connected routers are not present in the same IP subnet, adjacencies are not formed. In this case, the hellos received will be rejected as the interface addresses are not in same subnet. Hence make sure that the interfaces are configured in the same IP subnet.</p> <p>The system ID configured is unique throughout the network. If same system ID is configured for any two IS-IS routers adjacency will not be established. By enabling debug logs we can find the interface which is sending hellos with duplicate system ID. Hence make sure that system IDs are different across the network.</p>
<p>The adjacency state of some or all neighbors is stuck in INIT state in <code>show clns neighbors</code></p> <p>Authentication problems.</p> <p>The MTU (Maximum Transmission Unit) mismatch has occurred.</p>	<p>If IS-IS authentication is enabled, first resolve this issue. If IS-IS authentication is not enabled, the problem might be with mismatched MTUs. First verify the authentication configurations on the routers. Make sure that proper authentication method is configured and the passwords are consistent. Sometimes, authentication may be enabled on only one side in which case the other side will not be able to process the hellos and hence the neighbor state will be stuck in INIT state.</p> <p>If the authentication problem is resolved and still the neighbor state is stuck in INIT, the possible cause will be because of mismatched MTU. The <code>show isis interface</code> command displays the MTU size of the link. You can also enable debug logs and check the MTU of the packets being sent on the interfaces.</p>

Routing Update Problems

This section covers IS-IS routing update problems based on the assumption that there are no adjacency problems.

As a first step, check the contents of the LSPs. The `show isis database detail` command gives the detail of the specific LSP contents as shown below:

```
sh isis database level-1 0010.0000.0001.00-00 detail
Area 1:
IS-IS Level-1 Link State Database:
LSPID          LSP Seq Num  LSP Checksum  LSP Holdtime  ATT/P/OL
0010.0000.0001.00-00  0x00000002  0x7746        1172          0/0/0
Area Address: 49
```

```

NLPID:      0xCC
IP Address:  11.11.11.1
Metric:     10      IS 0010.0000.0001.01
Metric:     10      IP 11.11.11.0 255.255.255.0

```

Another important command is `show isis topology` which displays all known routers.

```
sh isis topology
```

Area 1:

IS-IS paths to level-1 routers

System Id	Metric	Next-Hop	Interface	SNPA
0010.0000.0001	10	0010.0000.0001	p7p1	0800.275a.1f66
0010.0000.0002	--			

IS-IS paths to level-2 routers

System Id	Metric	Next-Hop	Interface	SNPA
0010.0000.0001	10	0010.0000.0001	p7p1	0800.275a.1f66
0010.0000.0002	--			

Symptom/Cause	Solution
Route advertisement	<p>Most route advertisement problems are caused because of incorrect configurations. If some routes are missing, use the show commands mentioned above to find more information about topology and LSPs. Using the knowledge about topology you can narrow down the problem to a single router.</p> <p>As IS-IS is a link-state protocol, the routers depend on LSPs to learn the topology and routing information. If a route is missing, it might be because the routers did not receive the original LSP or the LSP was received corrupted and hence was purged. In such cases, enabling debug logs for LSPs will help to determine the problem.</p>
Route flaps	<p>Route flaps occur because of unstable links between the routers. As the links flap between UP and DOWN state, they induce SPF calculation resulting in high CPU utilization which might lead to crashes</p> <pre>show clns neighbors detail</pre> <p>Route flaps can be identified by checking the Uptime for the neighbor:</p> <p>In these cases physical connectivity should be stabilized to resolve the problem.</p> <p>In more complicated cases, the flapping might happen because of corrupted LSP storms or a routing loop. In such cases enabling debug logs for spf calculation helps to know which LSPs are changing more frequently and triggering SPF calculations. Care should be taken while enabling logs for these cases as CPU will be already overloaded.</p>

Symptom/Cause	Solution
Route redistribution	<p>IS-IS allows external routes like static, connected, or routes from other protocols to be distributed into IS-IS levels. The main reasons for external routes missing in routing table could be because of the wrong configuration. The source of the route may not be active in which case the external route will not be installed in the routing table.</p> <pre>show ip route</pre> <p>Verify that the source of the route is active.</p> <p>In case of static routes, the nexthop might not be reachable as the interface is down and hence it is not present in the routing table.</p> <pre>show ip interface brief</pre> <p>Check the interface state.</p> <p>External routes might be filtered in the route-map or distribute-list at the source, in which case it is the correct behavior to not install the route in the routing table.</p>

FAQ

Q: What happens to the RIP learned routes whose network address is same as one of the directly connected IPs, when the prefix length is greater than, less than, or equal to, that of the directly connected interface?

All routes with greater prefix length and less prefix length will be displayed in the RIP routing table. The equal prefix-length entry will be discarded.

For example, if the 1.1.0.0/24, 1.1.0.0/25 and 1.1.0.0/23 routes are redistributed into RIP, and a neighbor has an interface with IP 1.1.0.2/24, the RIP route entry of the 1.1.0.0/24 network will not be available. But, the 1.1.0.0/23 and 1.1.0.0/25 entries will be available in the RIP routing table. The 1.1.0.0/24 entry is available as a directly connected entry in the routing table. If this interface is down, the RIP route will become active, until the interface comes up.

Q: When redistributing other routes to RIP, why does the redistributed route always override the routes learned by RIP?

RIP learned routes have lower priority than redistributed routes (for example, connected/static/ISIS). Therefore, the redistributed routes always override the routes learned by RIP.

Q: Why is the BGP session reset after any BGP capability is configured?

In BGP, capabilities are advertised in the OPEN message during session initialization. If a capability is enabled or disabled after the session is established, the BGP session needs to be reset, and a new capability is included in the OPEN message.

Q: How do I set up “neighbor send-community” in BGP?

The “neighbor send-community” is set up by default. By default, on receiving the communities attribute, the router re-announces them to the neighbor. This command does not appear in the list of available commands in the Router mode. It is visible only when the user has used the `no neighbor send-community` command. Please refer to the *BGP Command Reference* for details on how to use this command.

Use the `show ip bgp neighbor` command to confirm that the neighbor send-community is set up.

Q: What is Route Reflection used for?


Route Reflection is used in IBGP to resolve the IBGP full mesh problem. Configuring one or more routers as Route Reflectors reduces the number of connections between BGP peers within an Autonomous System (AS).

The Route Reflecting BGP peer has to be configured with the peer addresses of all its route reflection clients. The Route Reflector is responsible for:

- Sending updates from a client peer to other client peers, as well as non-client peers.
- Sending updates from non-client peers to client peers.

Q: Can you explain more about the BGP “neighbor next-hop-self” command?

The neighbor next-hop-self command is only effective in the case of IBGP. For example:



On executing the `neighbor <a.b.c.d> next-hop-self` command on R1, the R1 router advertises the routes (if any) with the next-hop attribute that equals the R1 IP address.

This command is useful for advertising the routes learned by R1 from other routers, and are not reachable by R2.

Q: Does the “no bgp graceful restart” command turn off graceful restart? Or does it just set the parameters back to the default?

The `no bgp graceful-restart` command turns off the graceful restart functionality.

The `no bgp graceful-restart restart-time` and `no bgp graceful-restart stalepath-time` commands reset the timer values to the default values.

Q: Which BGP draft version is used to test conformance?

IP Infusion Inc. uses the IXIA ANVL test suite for testing conformance. The latest version of ANVL is based on draft-ietf-idr-bgp4-26.txt.

Q: Can I change the Area ID of an existing OSPF network configuration?

No, you cannot change the Area ID without deleting the existing configuration. You need to remove the network A.B.C.D/X area Y before changing the area ID of this network.

For example, entering “network 10.73.0.0/16 area 0.0.0.5” when “network 10.73.0.0/16 area 0.0.0.1” already exists, will display a warning message.

Q: How do I display information about max-age? I used the “show ip ospf database max-age” command: max-age was not displayed.

The `show ip ospf database max-age` command maintains a list of all the LSAs in the database that have reached the max-age (3600 seconds).

If the LSAs have not reached the max-age, it is not displayed. To test the functionality of max-age, follow these steps:

1. Connect two routers, R1 and R2 both of them running the OSPF daemon.
2. After a few minutes, kill the OSPF daemon on one of the routers (say on R2).
3. Wait for one hour to get a display of the list of LSAs that have reached the max-age on R1.

The following is a sample output of the `show ip ospf database max-age` command on R1:

```
# show ip ospf database max-age
```

```
OSPF Router process 100 with ID (3.3.3.4)
MaxAge Link States:
Link type: 7
Link State ID: 37.37.37.0
Advertising Router: 3.3.3.1
LSA lock count: 6
Link type: 7
Link State ID: 10.0.0.0
Advertising Router: 3.3.3.1
LSA lock count: 6
Link type: 7
Link State ID: 20.255.37.37
Advertising Router: 3.3.3.1
LSA lock count: 6
```

Q: How do I create a secondary loopback address? This address has to be advertised by LSAs to make it reachable from other routers and hosts.

Configure a secondary loopback address as follows:

```
(config)# interface lo
(config-if)# ip address A.B.C.D/32
```

For this loopback address to be advertised by LSAs, enable OSPF on this interface by configuring the routing process, and specifying the Process ID. The Process ID should be a unique positive integer identifying the routing process. Then define the interface and associate the area ID with the interface.

```
(config)# router ospf [process id]
(config-router)# network A.B.C.D/32 area 0
```

CHAPTER 5 Multicast

This chapter contains steps to resolve multicast issues.

IGMP

IGMP not Active

Symptom/Cause	Solution
Multicast routing is not configured	Use these commands: <code>show ip mvif</code> Check whether vif is created for the interface. <code>show ip igmp interface</code> Check the status of igmp <code>ip multicast-routing</code> Configure multicast routing.
Interface is down	Use these commands: <code>show ip interface brief</code> Check whether interface has an ip address or the status is down. <code>ip address A.B.C.D/subnet-mask></code> If the address field in the show ip interface brief output shows “unassigned”, configure IP address on the interface with mask. Alternatively, configured ip address can also be checked using “show running-config interface <ifname>” For ex: <code>ip address 2.2.2.2/24</code> <code>no shutdown</code> Bring up admin status of the interface. <code>show ip mvif</code> Check whether vif is created for the interface.
No IGMP configuration exists.	Use these commands: <code>ip igmp</code> Configure igmp on the interface <code>show ip igmp interface</code> Check whether igmp is enabled. <code>show running-config interface <ifname></code> Verify if igmp is configured. Note: IGMP can also be enabled by pim sparse/dense mode configuration.

Groups not added after Receiving Dynamic Reports

Symptom/Cause	Solution
IGMP is not active	Follow the steps above
IGMPv3 report should have one or more sources when the group-record type is INCLUDE	Use this command: <code>show ip igmp groups</code> The output displays the added group
Interface is configured as igmp proxy	Use these commands: <code>show ip igmp proxy</code> Check whether igmp proxy is enabled. <code>show running-config interface IFNAME</code> Check whether "ip igmp proxy service" configuration exists.

IGMP proxy is not active

Verify that IGMP proxy is enabled and active:

```
show ip igmp proxy
```

Symptom/Cause	Solution
Host interface not configured	Use these commands: <code>show ip igmp proxy</code> Check the upstream interface. <code>show running-config interface IFNAME</code> Check if the configuration exists. <code>ip igmp proxy-service</code> Configure interface as host interface (enable proxy-service).
Mroute-proxy interface not configured	Use these commands: <code>show ip igmp proxy</code> Check whether mroute is configured. If there is no output, mroute-proxy is not configured. <code>ip igmp mroute-proxy upstream-IFNAME</code> Configure mroute proxy on an interface.
Mismatch between proxy service interface and host interface	Use these commands: <code>show ip igmp proxy</code> Check the upstream interface. <code>show running-config</code> Check which interface has the proxy-service configured. <code>ip igmp proxy-service</code> Configure proxy service on the upstream interface shown by <code>show ip igmp proxy</code> .

Symptom/Cause	Solution
IP Multicast routing not configured	<p>Use these commands:</p> <pre>show running-config multicast</pre> <p>Check whether ip multicast routing is configured.</p> <pre>ip multicast-routing</pre> <p>if the configuration is not show in the running-config, use this command to configure ip multicast routing.</p>
Interface is down	Follow the steps explained above

MLD

MLD not Active

Check whether vif is created for the interface:

```
show ipv6 mif
```

Symptom/Cause	Solution
IPv6 Multicast routing is not configured	<p>Use these commands:</p> <pre>show ipv6 mld interface</pre> <p>Check the status of igmp</p> <pre>ipv6 multicast-routing</pre> <p>Use this command to configure multicast routing.</p>
Interface is down	<p>Use these commands:</p> <pre>show ipv6 interface brief</pre> <p>Check whether interface has an ipv6 address or if the status is down. If the address field is “unassigned”, configure IPv6 address on the interface with mask. For ex:</p> <pre>ipv6 address 3ffe:506::1/48</pre> <pre>no shutdown</pre> <p>Use this command to bring up admin status of the interface.</p>
No MLD configuration exists	<p>Use these commands:</p> <pre>ipv6 mld</pre> <p>Use this command to configure mld on the interface</p> <pre>show ipv6 mld interface</pre> <p>Check whether MLD is enabled.</p> <pre>show running-config interface IFNAME</pre> <p>Verify if mld is configured.</p> <p>Note: MLD can also be enabled by ipv6 pim sparse/dense mode configuration.</p>

Groups not Added on Receiving Dynamic Reports

Symptom/Cause	Solution
MLD is not active	<p>Follow the steps above</p> <p>MLDv2 report should have one or more sources when the group-record type is INCLUDE.</p> <pre>show ipv6 mld groups</pre> <p>The output displays the added group</p>
Interface is configured as mld proxy	<p>Use these commands:</p> <pre>show ipv6 mld proxy</pre> <p>Check whether mld proxy is enabled.</p> <pre>show running-config interface IFNAME</pre> <p>Check whether “ipv6 mld proxy service” configuration exists.</p>

MLD Proxy not Active

Verify mld proxy is enabled and active:

```
show ipv6 mld proxy
```

Symptom/Cause	Solution
Host interface is not configured	<p>Use these commands:</p> <pre>show ipv6 mld proxy</pre> <p>Check the upstream interface.</p> <pre>show running-config interface IFNAME</pre> <p>Check if the configuration exists.</p> <pre>ipv6 mld proxy-service</pre> <p>Configure interface as host interface (enable proxy-service).</p>
Mroute-proxy interface not configured	<p>Use these commands:</p> <pre>show ipv6 mld proxy</pre> <p>Check whether mroute is configured. If there is no output, mroute-proxy is not configured.</p> <pre>ipv6 mld mroute-proxy upstream-IFNAME</pre> <p>Configure mroute proxy on an interface.</p>
Mismatch between proxy service interface and host interface.	<p>Use these commands:</p> <pre>show ipv6 mld proxy</pre> <p>Check the upstream interface.</p> <pre>show running-config</pre> <p>Check which interface has the proxy-service configured.</p> <pre>ipv6 mld proxy-service</pre> <p>Configure proxy service on the upstream interface shown in the “show ip igmp proxy”.</p>

Symptom/Cause	Solution
IPv6 Multicast routing not configured	Use these commands: <pre>show running-config multicast</pre> Check whether ipv6 multicast routing is configured. <pre>ipv6 multicast routing</pre> If the configuration is not shown in the running-config, use this command to configure ipv6 multicast routing.
Interface is down	Follow the steps explained above

PIM

RPF Neighbor not Reachable

Symptom/Cause	Solution
PIM sparse mode is not configured on the interfaces	Use these commands: <pre>show ip pim neighbor</pre> Check the pim neighbour as the downstream router. <pre>ip pim sparse-mode</pre> Use this command to configure sparse mode on all the concerned interfaces
Multicast routing is not configured	Use this command: <pre>ip multicast-routing</pre> Configure multicast routing on the interface.
Interface is not active	Use these commands: <pre>show ip interface brief</pre> If the address field is "unassigned", configure IP address on the interface with a mask. Alternatively, configured ip address can also be checked using <code>show running-config interface IFNAME</code> . For ex: <pre>ip address 2.2.2.2/24</pre> <pre>no shutdown</pre> Bring up the admin status of the interface if down.

RPF Interface not What is Expected

Symptom/Cause	Solution
Some other route exists.	Use this command: <pre>show ip rpf <source-address></pre> Check whether RPF interface is as expected; If not some other route exists to the source which is causing the multicast packet to be dropped due to RPF check failure.

RP not Reachable

Check whether an RP is reachable from all the routers in the PIM domain:

```
show ip pim rp-mapping
```

Symptom/Cause	Solution
RP not configured on all the routers for static RP	Use this command: <pre>ip pim rp-address <RP-address></pre> Configure rp on all the routers in the pim domain.
VIF for RP is not created	Use these commands: <pre>show ip mvif</pre> Verify whether VIF is created for the RP. <pre>ip pim sparse-mode</pre> If RP is configured on the loopback interface of the router then pim sparse mode configuration must be present on the loopback interface. <pre>ip multicast-routing</pre> Multicast routing must be configured on each router.

Mroute not Created for PIM-SSM

Note: RP configuration is not required for PIM-SSM.

Symptom/Cause	Solution
Static group configuration without source	Use these commands: <pre>show ip igmp groups detail</pre> Check whether group recode is include and source list has one or more sources. <pre>no ip igmp static-group <group-addr></pre> <pre>no ip igmp join-group <group-addr></pre> Remove the static/join group without source, on the interface.
IGMPv2 report present for the group	Use these commands: <pre>show ip igmp groups detail</pre> Group mode should not be exclude. Check whether group mode is exclude; remove the v2 group-record using command below. <pre>clear ip igmp groups <group-addr></pre> Use this command to remove the igmp v2 group.

MSDP

MSDP Peering not Established

Symptom/Cause	Solution
MSDP not configured	Use these commands: <code>show ip msdp peer</code> Check whether msdp peering is established. <code>ip msdp peer <peer-address></code> Configure msdp peering on the RPs in different PIM domains.
Password mismatch	Use these commands: <code>show running-config</code> Check the configured password for msdp peers. <code>no ip msdp <wrong-password> peer <peer IP address></code> Unconfigure wrong password on the MSDP peer. <code>ip msdp <password> peer <peer IP address></code> Configure the same password as that of its corresponding MSDP peer.

CHAPTER 6 HQoS

This chapter contains steps to resolve HQoS issues.

Symptom/Cause	Solution
3-Level queuing hierarchy: <ul style="list-style-type: none">Queueing hierarchy at an interface is not knownDefault configurations of queues not known	<pre>show policy-map interface IFNAME</pre> <p>If an egress (queuing) policy-map is attached to interface, check its hierarchy and queue statistics (counters).</p> <pre>show policy-map type queuing <name></pre> <p>Check the configurations of a queueing policy-map.</p>
L2 and L3 QoS/Trust IEEE 802.1p/ (DSCP): <ul style="list-style-type: none">Wrong queue selection	<p>Queue selection for L2 and L3 ports are based on incoming packet's cos and dscp value respectively.</p> <pre>trust cos/dscp</pre> <p>Change the queue selection based on cos or dscp value.</p> <pre>show running config interface <name></pre> <p>Show if any user configured trust.</p>
Rate Limiting: <ul style="list-style-type: none">Packet drops ingress interfaceRate limiting not working	<p>If a particular flow of packets are sent at rate lesser than received rate and packets are dropped at ingress interface, there is a possibility of that flow is being policed (rate limiting) by policer. Policers are part of ingress policy-map which are attached to ingress interface to limit the rate of "particular flows".</p> <pre>show running config interface <if-name></pre> <p>Show any ingress policy-maps that are attached to interface</p> <pre>show policy-map type qos <policy-map name></pre> <p>Show the matching flow and the policer details</p>
Shaping and Bandwidth Reservation Per Queue: <ul style="list-style-type: none">Packet drop at egress queuePackets sent at lesser rate than received	<p>If there is no congestion at egress interface, but there is packet drop at queue and outgoing traffic rate is lesser than incoming rate, shaping would have been configured at that queue.</p> <pre>show policy-map type queueing <egress-policy-map></pre> <p>Check if any shaping is configured on the queues of the interface.</p> <pre>show policy-map type queueing <egress-policy-map></pre> <p>Check if there is congestion and queues are strict-priority; packets will flow according to the configured minimum bandwidth.</p>
WRR/WFQ/SP Scheduling Per Queue: <ul style="list-style-type: none">Different traffic flows are transmitted/processed at different rate	<p>If there is congestion at interface, if queue configured for SP scheduling, the packets are processed according to minimum bandwidth of queue. If queue configured for WRR scheduling, packets are processed according to the WRR weight of the queue.</p> <pre>show policy-map type queueing <egress-policy-map></pre> <p>Check the minimum bandwidth/weight.</p>

Symptom/Cause	Solution
Weighted RED <ul style="list-style-type: none">Packets are dropped only at some or all queues.	WRED/taildrop is used decrease the queue size/limit resulting in loss of packet. <code>show policy-map type queueing <egress-policy-map></code> Verify WRED/taildrop configuration.
802.1p remarking: <ul style="list-style-type: none">cos / dscp value is getting remarkedPackets are going to different queues than the one selected by incoming cos/dscp value	If packets of particular flow are going to different queue than the queue for their cos/dscp value, a remarking could have happened. <code>show policy-map type queueing <egress-policy-map></code> Check if remarking is done using an ingress policy-map. <code>qos map cos <in-cos> <out-cos></code> Configure remarking.

CHAPTER 7 Data Center and Virtualization

This chapter contains steps to resolve data center and virtualization issues.

VXLAN

Symptom/Cause	Solution
"map access port" for a physical port is not successful	Ensure that the access port that you are trying to map to a vxlan vniid is part of a bridge.
"map access port-vlan" for an interface is not successful	Ensure that the access port that you are trying to map to a vxlan vniid is part of a bridge. Ensure that the vlan which you are trying to map is configured on the access port.
"map network tunnel" is not successful for unicast tunnel	Verify that the physical interface used for tunnel is up and running. Ensure that the ip route is present on VTEP.
Can we have tunnel source ip as loopback interface ip?	Yes, you can use loopback IP as tunnel source IP.
Known unicast traffic not flowing from one end to another end after configuring static entries	Ensure that all the interfaces are up and running. Ensure that source mac address is learnt on source VTEP on each VTEP using "show bridge". Ensure that you can ping tunnel destination IP, ie. you have a route to reach the destination VTEP. Ensure that the mac address configured in static entry is that of correct destination host.
BUM traffic is not flowing	Ensure that multicast tunnel is configured on the VTEPs. Unicast tunnel does not support BUM traffic, only known unicast traffic.
Having different VLAN's on each access side isnt working.	For input VTEP, traffic is allowed based on the vlan configured on the access port but on destination VTEP after decapsulation, vlan check is not happening.
Multicast tunnel installation is unsuccessful.	Verify that the physical interface used for tunnel is up and running. Ensure that "tunnel add interface <if-name>" is configured in the multicast tunnel.
Multicast traffic not flowing from one end to another end after configuring VTEPs successfully.	Ensure that all the interfaces are up and running. Ensure that the intermediate routers are supporting multicast. Ensure that the VTEPs have joined the multicast group correctly, ie the tunnel destination ip(multicast ip) is correct on the participating VTEPs. Ensure that source mac address is learnt on source VTEP on each VTEP using "show bridge".

VXLAN-EVPN

Symptom/Cause	Solution
<code>vxlان host-reachability-protocol evpn-bgp vrfblue</code> is not successful	Make sure that the vrf which you are trying to map is configured.
Can we have vtep-ip-global as loopback ip	Yes.This ip is the source ip of the tunnel.
<code>sh bgp l2vpn evpn</code> does not show the multicast entries	Make sure that your route is reachable using any static/dynamic (isis/ospf...) protocol. Also make sure that bgp neighborhood is established.
Ping is not working for a particular vnid	Ensure that all the interfaces are up and running Ensure that source mac address is learnt on source VTEP on each VTEP using <code>show bridge</code> . Ensure that you can ping tunnel destination IP, ie. you have a route to reach the destination VTEP. Ensure that bgp neighborhood is established, have proper router-id's configured Ensure that you have imported/exported the route correctly using route-target.Ensure unique RD in all VRFs Also ensure that on all the VTEPs, access port is mapped to respective/relevant VNID

CHAPTER 8 Security

This chapter contains steps to resolve security issues.

DHCP Snooping

Symptom/Cause	Solution
DHCP packets not received DHCP snooping not enabled on bridge. DHCP snooping for that vlan not enabled on bridge	Use this command: <code>show ip dhcp snooping bridge BRIDGEID</code> Make sure DHCP snooping is enabled on the bridge
DHCP snooping entries not visible The interface which the ip address assigned might be a trust port	Use this command: <code>show ip dhcp snooping bridge BRIDGEID</code> Make sure the interface connected to host should be untrusted. If it is showing trust for that interface, untrust the interface to see the entry in the table.

DHCP Snooping IP Source Guard

Symptom/Cause	Solution
Not able to enable ip source guard on interface	Use this command: <code>show ip dhcp snooping bridge BRIDGEID</code> Make sure DHCP snooping is enabled on the bridge
Unable to execute the ip source guard mode merge command	Use the above command to make sure ipsg is enabled on that interface and then only merge will be accepted.
How to see the policies used as part of IP source guard on interface	Use this command: <code>show ip verify source interface IFNAME</code> Shows the entries learned as part this interface and the same is pushed as policies.

DHCP Snooping over MLAG

Symptom/Cause	Solution
DHCP packets not received DHCP snooping not enabled on bridge. DHCP snooping for that vlan not enabled on bridge	Use this command: <code>show ip dhcp snooping bridge BRIDGEID</code> Make sure DHCP snooping is enabled on the bridge.
DHCP snooping entries not visible The interface which the ip address assigned might be a trust port	Use this command: <code>show ip dhcp snooping bridge BRIDGEID</code> Make sure the interface connected to host should be untrusted. If it is showing trust for that interface, untrust the interface to see the entry in the table.
DHCP packets are not synced between MLAG active-active/active-standby nodes.	Use this command: <code>show mcec statistics</code> Make sure MLAG domain adjacency is up and neighbor is in-sync.
DHCP packets are dropped.	Use this command: <code>show ip dhcp snooping bridge BRIDGEID</code> Make sure that the MLAG interface facing towards the server is trusted.

DHCPv6-Prefix Delegation

Symptom/Cause	Solution
Prefix are not delegated	Use this config command: <code>no ipv6 nd suppress-ra</code> Make sure this command is enabled on the requesting router host connected interface.

Symptom/Cause	Solution
Prefixes are not delegated with varying prefix length	Use this config command: <code>ipv6 address PREFIX_FROM_SERVER ::1:0:0:0:1/64</code> Suffix should start with "::" and mask should be 64.
Prefixes are not learnt on Requesting router	Use this command: <code>show ipv6 dhcp interface</code> Make sure that prefix delegation is enabled on that interface.

CHAPTER 9 SNMP

Follow these steps to resolve SNMP issues.

1. The SNMP daemon (`snmpd`) must be running which enables the protocol. This can be checked at the Linux prompt.
2. Any required configuration such as community strings and so on must be set up in `snmpd.conf`. This can be checked at the Linux prompt.
3. If traps are enabled, the `trapd` daemon must be running. This can be checked at the Linux prompt.
4. Enable logging (logging level *daemon-name* <0-7>) and traces (`debug snmp-server`) to find more about failure of a command.
5. `Snmpget` or `snmpwalk` may timeout, in case `show techsupport` command is in progress. In such case, provide the timeout (option -t) while invoking `snmpget` or `snmpwalk`.

CHAPTER 10 DHCP Client and Relay

This chapter contains steps to resolve DHCP client and relay issues.

Symptom/Cause	Solution
DHCP client does not get an IP address after a request to server	The <code>dhclient</code> process may not be running for the input interface, look for the <code>dhclient</code> process “ <code>ps -aef</code> ” at the <code>:linux</code> prompt. Look for the error messages in <code>/var/log/message</code> .
DHCP packet loss	Capture packets on interested interfaces on the client/server side. Enable <code>dhclient</code> and <code>dhcp</code> (server) in debug mode to get messages (error or pkt info)

CHAPTER 11 Remote Logging

This chapter contains steps to resolve remote logging issues.

Symptom/Cause	Solution
No logs are seen in <code>/var/log/messages</code>	If syslog messages are not being written, check whether <code>/etc/rsyslog.conf</code> file contains <code>*.* /var/log/messages</code> . This entry is added by <code>hostpd</code> during start up, irrespective of any logging server is configured. Also, verify the <code>rsyslogd</code> process is running (<code>pgrep rsyslog</code>).
No logs are seen	The logging to a logfile or syslog depends on the configured logging level. If the logging level is debug then enable "debug <module>".

CHAPTER 12 System Management

This chapter contains steps to resolve system management issues.

Symptom/Cause	Solution																																																												
Non availability of telnet/ssh service	<p>When the node is booting up, we disable all remote access. Upon the start of <code>hostpd</code>, the service <code>xinetd</code> starts.</p> <p>Make sure <code>hostpd</code> is running or started during init sequence of board initialization, and <code>xinted</code> service is running. Execute at the Linux prompt and verify listening socket:</p> <pre>ip netns exec zebosfib1 netstat -tln</pre> <p>Active Internet connections (only servers)</p> <table><thead><tr><th>Proto</th><th>Recv-Q</th><th>Send-Q</th><th>Local Address</th><th>Foreign Address</th></tr><tr><th>State</th><th>PID/Program name</th><th></th><th></th><th></th></tr></thead><tbody><tr><td>tcp</td><td>0</td><td>0</td><td>127.0.0.1:705</td><td>0.0.0.0:*</td></tr><tr><td>LISTEN</td><td>30044/snmpd</td><td></td><td></td><td></td></tr><tr><td>tcp</td><td>0</td><td>0</td><td>0.0.0.0:199</td><td>0.0.0.0:*</td></tr><tr><td>LISTEN</td><td>30044/snmpd</td><td></td><td></td><td></td></tr><tr><td>tcp6</td><td>0</td><td>0</td><td>:::22</td><td>:::*</td></tr><tr><td>LISTEN</td><td>29997/xinetd</td><td></td><td></td><td></td></tr><tr><td>tcp6</td><td>0</td><td>0</td><td>:::23</td><td>:::*</td></tr><tr><td>LISTEN</td><td>29997/xinetd</td><td></td><td></td><td></td></tr><tr><td>tcp6</td><td>0</td><td>0</td><td>:::830</td><td>:::*</td></tr><tr><td>LISTEN</td><td>29997/xinetd</td><td></td><td></td><td></td></tr></tbody></table>	Proto	Recv-Q	Send-Q	Local Address	Foreign Address	State	PID/Program name				tcp	0	0	127.0.0.1:705	0.0.0.0:*	LISTEN	30044/snmpd				tcp	0	0	0.0.0.0:199	0.0.0.0:*	LISTEN	30044/snmpd				tcp6	0	0	:::22	:::*	LISTEN	29997/xinetd				tcp6	0	0	:::23	:::*	LISTEN	29997/xinetd				tcp6	0	0	:::830	:::*	LISTEN	29997/xinetd			
Proto	Recv-Q	Send-Q	Local Address	Foreign Address																																																									
State	PID/Program name																																																												
tcp	0	0	127.0.0.1:705	0.0.0.0:*																																																									
LISTEN	30044/snmpd																																																												
tcp	0	0	0.0.0.0:199	0.0.0.0:*																																																									
LISTEN	30044/snmpd																																																												
tcp6	0	0	:::22	:::*																																																									
LISTEN	29997/xinetd																																																												
tcp6	0	0	:::23	:::*																																																									
LISTEN	29997/xinetd																																																												
tcp6	0	0	:::830	:::*																																																									
LISTEN	29997/xinetd																																																												
Failure to authenticate a user	<p>If the basic files for Linux authentication of a user are missing/corrupted, the login to the node is denied. Using console root user, make sure the <code>/etc/passwd</code> file has an entry for the user trying to login. Look for authentication errors are in <code>/var/log/messages</code>, for more about such failures.</p>																																																												
Remote access to the node via telnet/ssh hangs	<p>The shell <code>imish/cmlsh</code> is configured for all OcNOS users, except for user <code>root</code>, which is accessible via console only. If the module <code>imi</code> or <code>cml</code>d is not responding, then there will be no <code>imish/cmlsh</code> prompt after successful login.</p> <p>The system monitoring module (<code>pservd</code>) restarts such hung modules, recovering hang states of one of more modules. Look for the core directory (<code>/var/log/crash/cores</code>) and syslog messages in <code>/var/log/messages</code> to find the actions from system monitoring module.</p>																																																												

Symptom/Cause	Solution
Continuous restart of any module	<p>If any module is restarting continuously, disable monitoring such module via:</p> <pre>no software-watchdog <module name></pre> <p>If the NSM/HSL module crashes or hangs, the system reboots.</p> <p>The system does not reboot automatically when the earlier two reboots were due to HSL or NSM crashes during the initial few minutes of board boot up. This is to stop continuous reboots of the system due to NSM/HSL crashes.</p> <p>There is no mechanism to disable this except for disabling pservd service. Stop the service pservd to disable it.</p> <p>If module pservd is hung, it will be restarted in 5 mins.</p>
Deleting ZebOS.conf loses management IP address	<p>During ONIE installation, if you do not configure a static IP address, OcNOS boots and gets an IP address for eth0 (management port) through DHCP and updates the <code>/etc/network/interfaces</code> file. Once you configure a static IP address from the OcNOS command line and save the configuration, OcNOS updates <code>/etc/network/interfaces</code> and changes the method used to configure eth0 from dhcp to static.</p> <p>In this scenario, if you delete <code>ZebOS.conf</code>, then the management IP address is lost and you can only recover management access by assigning an IP address via the console.</p>
sys-update install <installer> failure	<ul style="list-style-type: none"> • No free space left on system. Minimum 1 GB space is needed: remove some files to make available space > 1GB on device. • Binaries not compatible with the board: use proper installer file for the respective board. • Installer not downloaded properly, try again: downloaded installer file is not complete. • Source Interface not found. • OcNOS version you are trying to upgrade is already Installed: no need to upgrade again, you have the same version already installed. • File not found on board: installer file is not present on board for given path, provide valid path for installer file. • File not found on server: installer file is not present on the server provided in the link, provide valid link for installer file. • Server connection timed out: waited 60 seconds for server to respond. • Unsupported protocol: the ftp, http, tftp, and file protocols are supported. • Invalid installer: installer file is not valid. • % Source interface is not up : Ensure source interface is UP <p>Note: When the sys-update operation stops without any error, check whether the IP reachability is there to download the installer file.</p>

License Troubleshooting

Note: If you install OcNOS version 1.3.8 (or later) for the first time on a device and then perform license activation, the activated license is deactivated if you install any version before 1.3.8. To recover, the license has to be activated/installed again.

Symptom/Cause	Error	Solution
Failure: license get <url> / license refresh	License file (IPI-DEVICEID.bin) Not Found	license is not present on system, use "license get <url>" to install the license
	License installation failed due to incorrect Device ID in the License file, please use the relevant device specific license file	Downloaded license is not for the current device and it is removed. Use "license get <url>" to install the correct license.
	The allowed time to process response has expired	License file lifetime has expired, but this is not an actual license expiration error. Also this lifetime value is not visible to the user. So download the license from FNO portal using "license get <url>" again. Note: The Lifetime field is the lifetime of the capability response, in seconds, after which the response is considered "stale" and cannot be processed by the client or server. IPInfusion has the lifetime set to 3628800 seconds or 42 days. If a capability response is created and held without installing for more than the specified period (42 days), it turns stale and the target device would not be able to process this.
	Response is out of order with previous responses, also show license is not reflecting the new license features.	User have already installed a license which is downloaded more recently than the current license. But once you land in this error case, re-installation of either of these two licenses will not be helpful anymore. So download the license from portal freshly and install "license get <url>" command.
	Failed to create trusted storage	Remove the contents of /cfg/license/ then install the license using "license get <url>" CLI.
	Invalid license file	License file might be corrupt, so download and install the license from FNO license portal using "license get <url>". If it still fails, validate the checksum of the license file in /cfg/license/bin/ with the one downloaded from the FNO portal.
	Start date for the license is in the future	Correct the system clock and issue the "license refresh" command to install the license
	Empty license file	Download the license from FNO license portal, and install using "license get <url>".
	Failed to process capability response / Failed to process the license file	Remove the contents of /cfg/license/ then install the license using "license get <url>" CLI.
	Command "license get" is not installing the given license file, but processes old license and fails.	Correct the system clock and issue "license get <url>" to install the license again.

Symptom/Cause	Error	Solution
	License is not matching with device software	License file SKU is not compatible with device software, please map the right SKU, then generate and install the license.
	Empty license response received: license is not mapped with SKU or the license server exhausted its limit	Select a SKU while generating license from FNO license portal or increase the license pool on the license server to accommodate more devices

Zero Touch Provisioning

Symptom/Cause	Solution
ONIE Image/IP address not fetched from DHCP server	Ensure DHCP server is up and reachable from device.
ONIE Image/IP address not fetched from DHCP server	Ensure DHCP server config file has proper info for this device. Things to check in dhcpd.conf file: <ul style="list-style-type: none"> • Device MAC address is proper, if MAC address based config is used. • Device VCI is proper in DHCP config file. Use onie-sysinfo command to check the same. • Syntax and value provided for DHCP options to be used by this device is proper.
Error: license & config already exist on device. Skipping ZTP provided data check.	As device already had old license and config, ZTP provided info are discarded
Error: Lease info from DHCP server not found. Skipping	Ensure that DHCP server is up and reachable. Install license and then config manually once device is UP
Error: Unable to download the startup config mentioned at ZTP/DHCP server!	Ensure that ZTP provided config file path is reachable and having download permission. Install license and then config manually once device is UP.
Error: Unable to download the license provisioned through ZTP/DHCP server!	Ensure that ZTP provided license file path is reachable and having download permission. If license path was provided in DHCP server config, ensure license file with this Device ID exists. Install license and then config manually once device is UP.
Error: eth0 is not configured using dhcp	DHCP server didn't install IP address as dynamic. Please install license/config manually
Error: ZTP provided config didn't get applied successfully	Ensure valid license was installed or provided by ZTP server.

CHAPTER 13 TACACS+ and AAA

This chapter contains steps to resolve TACACS+ and AAA issues.

Symptom/Cause	Solution
Server Not Reachable case verification via <code>/var/log/messages</code> or system logs	Make sure the TACACS+ server is running. Try login via ssh / telnet, if login fails check the following via console: <code>show system log include PAM-tacplus</code> PAM-tacplus : Connection failed srv 0: Transport endpoint is not connected. The above message confirms that the TACACS+ server is not reachable or not running.
Server Not Reachable case verification Via Enabling the AAA error-enable functionality	Execute the command <code>aaa authentication login error-enable vrf management</code> Try login via telnet /ssh. It will display <code>Remote TACACS servers unreachable</code> and will fail to login.
In-correct TACACS+ Username or Password.	Try login via ssh / telnet. if login fails check the following via console: <code>show system log include pam status</code> <code>pam_sm_authenticate: exit with pam status: 7</code> The above messages confirms that the authentication is failed.
In-correct TACACS+ key	Try login via ssh / telnet. if login fails check the following via console <code>show system log include tac_authen_read</code> <code>tac_authen_read: inconsistent reply body, incorrect key?</code> The above messages confirms that incorrect key.
User account locked	<p>By default, a user account is locked when a user gives an incorrect password 4 times. Once a user account is locked, by default the lock is cleared after 1200 secs (20 minutes). The Alert Operlog below appears when a user is locked.</p> <pre>"OcNOS : HOSTP : ALERT : [USER_MGMT_ACCOUNT_LOCKED_1]: Threshold for unsuccessful authentication attempts exceeded by user 'test'. User account will be unlocked after '1200' seconds."</pre> <p>You can configure the maximum fail attempts and unlock timeout using these commands:</p> <pre>aaa local authentication attempts max-fail <1-25> (The default maximum fail authentication attempts is 4) aaa local authentication unlock-timeout <1-3600> (The default unlock timeout for a locked user is 1200 seconds)</pre> <p>To manually clear the lock of a user, give this command:</p> <pre>clear aaa local user lockout username USERNAME</pre> <p>Possible causes of a user getting locked:</p> <ul style="list-style-type: none">• Incorrect password given more than the <code>max-fail</code> attempts.• When copying <code>show running-config</code> output manually and pasting to a file, make sure the <code>username</code> command is in a single line. If there is an embedded newline character in the password, the login fails.

Note: Configure the `aaa authentication login default vrf management group tacacs+ local` or `aaa authentication login default fallback error local vrf management` to fall back to local user authentication if a TACACS+ server is not reachable.

The same user can be present locally and in TACACS+ server, but the password can be different.

If a TACACS+ server is not reachable then, use the locally configured password to login.

AAA console authentication via Management VRF

Symptom/Cause	Solution
Enabling console authentication via the TACACS server and having the TACACS server reachable only through the Management (MGMT) VRF leads to login failures. This occurs because console login operates in the default VRF, while the TACACS server resides in the MGMT VRF. As a result, the TACACS client in the default VRF cannot establish communication with the server in the MGMT VRF.	<p>To establish reachability of the TACACS Server from the Default VRF, follow these steps:</p> <p>Utilize the default loopback interface as the source interface for TACACS. You can also select any other interface within the Default VRF.</p> <p>Introduce a static route in the Default VRF to reach the TACACS server. This step facilitates route leaking.</p> <p>Implement a static route in the MGMT VRF to access the loopback interface in the Default VRF. This is another instance of route leaking.</p> <p>Configure a static route in the TACACS server, enabling it to reach the loopback interface within the Default VRF.</p>

Example:

Upon implementing the configurations provided below, console authentication now takes place via the TACACS server (10.12.159.141), which is accessible through the MGMT VRF.

```
OcNOS#sh running-config aaa
aaa authentication login console group tacacs+

OcNOS#sh running-config tacacs+
feature tacacs+
tacacs-server login host 10.12.159.141 seq-num 1 key 7
0x67efdb4ad9d771c3ed8312
b2bc74cedb

OcNOS#sh running-config interface lo
!
interface lo
 ip address 127.0.0.1/8
 ip address 1.1.1.1/24 secondary
 ipv6 address ::1/128
!

OcNOS#sh running-config ip route
!
ip route 10.12.159.141/32 eth0
ip route vrf management 1.1.1.1/32 lo global
!

OcNOS#show ip interface eth0 brief
```

'*' - address is assigned by dhcp client

Interface	IP-Address	Admin-Status	Link-Status
eth0	*10.12.159.117	up	up

Route to be Added in TACACS Server:

```
ip route add 1.1.1.1/32 via 10.12.159.117
```

CHAPTER 14 Configuration Management

This chapter contains steps to resolve device configuration and management issues.

Symptom/Cause	Solution
CLI running configuration output does not match the output of NetConf 'get-config' output, and this conflict is leading to an inability to alter device configuration. This situation can arise when the device running-config or start-up configuration is not synced with the device configuration store.	You can use this command to recover the device from this state: <ul style="list-style-type: none"><code>reload flush-db</code>: This command removes the current start-up database. The board reboots after you give this command.

Replacing the Start-up Configuration File

To replace a start-up configuration file, you must update the configuration database as well. To do this, follow these steps:

1. Copy the configuration file to the startup configuration (`copy file` command)
2. Reboot and apply the configuration to the device and the configuration database (`reload flush-db` command)

If the start-up configuration file and the configuration database are not synchronized, then the system does not allow configurations to be done and will result in an unstable state. Therefore, you must keep the configuration file and the configuration database synchronized.

This example shows replacing a start-up configuration file and then synchronizing it to the configuration database:

```
#copy file /home/ZebOS-TEST.conf startup-config
Copy Success
#
#reload flush-db
The system has unsaved changes.
Would you like to save them now? (y/n): n

Configuration Not Saved!
Are you sure you would like to reset the system? (y/n): y
```

For both of these prompts, you must specify whether to save or discard the changes. Abnormal termination of the session without these inputs can impact the system behavior.

For the `unsaved changes` prompt:

Would you like to save them now?

Always say “no” to this prompt because otherwise the command takes the current *running configuration* and applies it to the current start-up configuration.

OcNOS Datastores

Table 14-1 describes the datastores used by OcNOS.

Table 14-1: OcNOS datastores

Name	Location	Description
Candidate configuration datastore	volatile memory	A configuration datastore that can be manipulated without impacting the device's current configuration and that can be committed to the running configuration datastore and stored in volatile memory.
Running configuration datastore	CML_RD.db	A configuration datastore holding the complete configuration currently active on the device. The running configuration data store always exists.
Startup configuration datastore	CML_DB.db	The configuration datastore holding the configuration loaded by the device when it boots. This separates the startup configuration datastore from the running configuration datastore.
ZebOS.conf	/cfg/usr/local/etc/ ZebOS.conf	The <code>show running-config</code> command queries protocol module data structures to get the running configuration and display it. A <code>write</code> given at the command line writes the running configuration to the <code>ZebOS.conf</code> file. The <code>ZebOS.conf</code> file contains device startup configuration in human-readable format (same as <code>show running-config</code> output).

Note: Do not alter these datastores manually. Always use OcNOS/NetConf commands to read and update these files.