



Extreme OS ONE Switching and Routing v22.2.3.0 Troubleshooting Reference Guide

Comprehensive Solutions and Management

9041124-00 Rev AA
August 2026



Copyright © 2026 Extreme Networks, Inc. All rights reserved.

Legal Notice

Extreme Networks, Inc. reserves the right to make changes in specifications and other information contained in this document and its website without prior notice. The reader should in all cases consult representatives of Extreme Networks to determine whether any such changes have been made.

The hardware, firmware, software or any specifications described or referred to in this document are subject to change without notice.

Trademarks

Extreme Networks and the Extreme Networks logo are trademarks or registered trademarks of Extreme Networks, Inc. in the United States and/or other countries.

All other names (including any product names) mentioned in this document are the property of their respective owners and may be trademarks or registered trademarks of their respective companies/owners.

For additional information on Extreme Networks trademarks, see: <https://www.extremenetworks.com/about-extreme-networks/company/legal/trademarks>

Open Source Declarations

Some software files have been licensed under certain open source or third-party licenses.

End-user license agreements and open source declarations can be found at: <https://www.extremenetworks.com/support/policies/open-source-declaration/>





Table of Contents


Abstract.....	v
Preface.....	vi
Text Conventions.....	vi
Documentation and Training.....	vii
Open Source Declarations.....	viii
Training.....	viii
Help and Support.....	viii
Subscribe to Product Announcements.....	ix
Send Feedback.....	ix
What's New in this Document.....	10
Certificate Management Issues.....	11
Event Management Issues.....	12
Link Aggregation Control Protocol (LACP).....	12
Troubleshoot Port-Channel Related Issues.....	12
Troubleshoot Link Layer Discovery Protocol (LLDP).....	16
Troubleshoot Multi-Chassis Link Aggregation (MLAG).....	17
Simple Network Management Protocol (SNMP).....	18
Installation and Deployment Issues.....	20
Zero-Touch Provisioning (ZTP).....	20
Network Infrastructure Components Issues.....	21
Address Resolution Protocol (ARP) and Neighbor Discovery (ND)	21
ARP Data Structure.....	21
MLAG Interfaces in ARP ND.....	22
Media Access Control (MAC)-DB.....	22
Nexthop Table.....	22
Entries Advertised Towards Unified Forwarding Table (UFTm) or MLAG.....	22
Enabling Detailed Debugs in ARP ND.....	23
Show Command Changed Output.....	23
Service Access Gateway (SAG) IP or MAC Behavior.....	24
Reconstruct Configuration Database (CDB) from Tech Support.....	26
Reconstruct System Database (SDB) from Tech Support.....	26
Troubleshoot Bidirectional Forwarding Detection (BFD).....	27
Troubleshoot BFD Flap Issues.....	28
Troubleshoot Kernel Crash.....	29
L3 interface.....	30
Logs and Configs.....	30
ECMP and RH Debugging.....	30
Limitations and Restrictions.....	31
Bridge Domain	31

Limitations and Restrictions.....	31
Troubleshoot Application-Specific Integrated Circuit (ASIC).....	31
Internal Message.....	32
Security Management Issues.....	33
Secure Shell (SSH) Services.....	33
Device-Level Commands.....	33
Log Files to Check.....	34
SSH Server Configuration.....	34
Core Files.....	34
Tech Support Files.....	34
Telnet Services.....	34
Device-Level Commands.....	34
Log Files to Check.....	35
Tech Support Files.....	35
Token Management.....	35
Troubleshoot Terminal Access Controller Access Control System Plus (TACACS+)	
Management.....	36
Troubleshoot Radius Management.....	36
Network Time Protocol (NTP) Management.....	37
Device Commands.....	37
Log Files.....	37
Additional Troubleshooting Steps.....	38
Troubleshoot Lightweight Directory Access Protocol (LDAP) Management.....	38
User Management.....	39
Device-Level Commands.....	39
Log Files.....	39
Tech Support Files.....	39
CLI Shell.....	39
Troubleshooting Hardware Ternary Content-Addressable Memory (TCAM) Infra and	
System Access Control Lists (ACL).....	39
Debug Commands.....	40
Key Concepts.....	40
Understanding Output.....	40
Troubleshooting User Access Control Lists (ACL).....	41
User ACL Features and Classifier Microservices.....	41
Configuration Flow.....	41
Classifier Microservice.....	41
Debug Commands.....	42
Verification.....	42
RBID.....	42
User ACL End-To-End Debug Flow Chart.....	43
Troubleshooting gRPC Network Management Interface (gNMI) and gRPC Network	
Operations Interface (gNOI).....	44
Device Logs and Debugging.....	44
Tech Support Checklist.....	44
Restrictions and Limitations.....	45
YANG Paths.....	45
Troubleshoot Service Access Gateway (SAG).....	45



Abstract

The *Extreme OS ONE Switching and Routing Troubleshooting Reference Guide* version 22.2.3.0 delivers comprehensive technical solutions for diagnosing and resolving issues across supported platforms, including Extreme 8520, 8720, 8730, and 8820. Designed for advanced IT professionals, it details procedures for backup and restore operations, certificate and device management, installation, deployment, licensing, and upgrades.



Preface

Read the following topics to learn about:

- The meanings of text formats used in this document.
- Where you can find additional information and help.
- How to reach us with questions and comments.

Text Conventions

Unless otherwise noted, information in this document applies to all supported environments for the products in question. Exceptions, like command keywords associated with a specific software version, are identified in the text.

When a feature, function, or operation pertains to a specific hardware product, the product name is used. When features, functions, and operations are the same across an entire product family, such as Extreme Networks switches, the product is referred to as *the switch*.

Table 1: Notes and warnings

Icon	Notice type	Alerts you to...
	Tip	Helpful tips and notices for using the product
	Note	Useful information or instructions
	Important	Important features or instructions
	Caution	Risk of personal injury, system damage, or loss of data
	Warning	Risk of severe personal injury

Table 2: Text

Convention	Description
screen displays	This typeface indicates command syntax, or represents information as it is displayed on the screen.
The words <i>enter</i> and <i>type</i>	When you see the word <i>enter</i> in this guide, you must type something, and then press the Return or Enter key. Do not press the Return or Enter key when an instruction simply says <i>type</i> .
Key names	Key names are written in boldface, for example Ctrl or Esc . If you must press two or more keys simultaneously, the key names are linked with a plus sign (+). Example: Press Ctrl+Alt+Del
<i>Words in italicized type</i>	Italics emphasize a point or denote new terms at the place where they are defined in the text. Italics are also used when referring to publication titles.
NEW!	New information. In a PDF, this is searchable text.

Table 3: Command syntax

Convention	Description
bold text	Bold text indicates command names, keywords, and command options.
<i>italic</i> text	Italic text indicates variable content.
[]	Syntax components displayed within square brackets are optional. Default responses to system prompts are enclosed in square brackets.
{ x y z }	A choice of required parameters is enclosed in curly brackets separated by vertical bars. You must select one of the options.
x y	A vertical bar separates mutually exclusive elements.
< >	Nonprinting characters, such as passwords, are enclosed in angle brackets.
...	Repeat the previous element, for example, <i>member[member...]</i> .
\	In command examples, the backslash indicates a “soft” line break. When a backslash separates two lines of a command input, enter the entire command at the prompt without the backslash.

Documentation and Training

Find Extreme Networks product information at the following locations:

[Current Product Documentation](#)

[Release Notes](#)

[Hardware and Software Compatibility](#) for Extreme Networks products

[Extreme Optics Compatibility](#)

[Other Resources](#) such as articles, white papers, and case studies

Open Source Declarations

Some software files have been licensed under certain open source licenses. Information is available on the [Open Source Declaration](#) page.

Training

Extreme Networks offers product training courses, both online and in person, as well as specialized certifications. For details, visit the [Extreme Networks Training](#) page.

Help and Support

If you require assistance, contact Extreme Networks using one of the following methods:

[Extreme Portal](#)

Search the GTAC (Global Technical Assistance Center) knowledge base; manage support cases and service contracts; download software; and obtain product licensing, training, and certifications.

[The Hub](#)

A forum for Extreme Networks customers to connect with one another, answer questions, and share ideas and feedback. This community is monitored by Extreme Networks employees, but is not intended to replace specific guidance from GTAC.

[Call GTAC](#)

For immediate support: (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2800. For the support phone number in your country, visit www.extremenetworks.com/support/contact.

Before contacting Extreme Networks for technical support, have the following information ready:

- Your Extreme Networks service contract number, or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any actions already taken to resolve the problem
- A description of your network environment (such as layout, cable type, other relevant environmental information)
- Network load at the time of trouble (if known)
- The device history (for example, if you have returned the device before, or if this is a recurring problem)
- Any related RMA (Return Material Authorization) numbers

Subscribe to Product Announcements

You can subscribe to email notifications for product and software release announcements, Field Notices, and Vulnerability Notices.

1. Go to [The Hub](#).
2. In the list of categories, expand the **Product Announcements** list.
3. Select a product for which you would like to receive notifications.
4. Select **Subscribe**.
5. To select additional products, return to the **Product Announcements** list and repeat steps 3 and 4.

You can modify your product selections or unsubscribe at any time.


Send Feedback

The User Enablement team at Extreme Networks has made every effort to ensure that this document is accurate, complete, and easy to use. We strive to improve our documentation to help you in your work, so we want to hear from you. We welcome all feedback, but we especially want to know about:

- Content errors, or confusing or conflicting information.
- Improvements that would help you find relevant information.
- Broken links or usability issues.

To send feedback, email us at Product-Documentation@extremenetworks.com.

Provide as much detail as possible including the publication title, topic heading, and page number (if applicable), along with your comments and suggestions for improvement.



What's New in this Document

There are no content changes for this guide for the Extreme OS ONE Switching and Routing 22.2.3.0 software release.

Extreme OS ONE Switching and Routing 22.2.3.0 and later releases support Extreme 8520, Extreme 8720, Extreme 8730, and Extreme 8820 hardware platforms.



Note

Although many software and hardware configurations are tested and supported for this release, all possible configurations and scenarios are beyond this document's scope.


For additional information, see the *Extreme OS ONE Switching and Routing Release Notes* for this version.



Certificate Management Issues

Use this topic to learn about the log and certificate storage locations for troubleshooting certificate-related issues. The following table summarizes where logs and certificates are stored and the naming conventions for certificate files to assist in troubleshooting related issues:

Category	Location/File Name
Log Files	• tmp/tierra/trace/security/security-RASTrace.log • /tmp/tierra/trace/security/security-RASTrace-DError.log • TS-sanity-<timestamp>/tmp/tierra/trace/security/security-RASTrace.log • TS-sanity-<timestamp>/tmp/tierra/trace/security/security-RASTrace-DError.log
Certificate Storage	Imported or generated certificates are stored in the directory: /var/data/cert-mgmt/appcert/<ssl-profile-id>/
File Formats	• CA Certificate: ca_<ssl-profile-id>.pem • App Certificate: <ssl-profile-id>.pem • Key File: <ssl-profile-id>.key



Event Management Issues

[Link Aggregation Control Protocol \(LACP\) on page 12](#)

[Troubleshoot Link Layer Discovery Protocol \(LLDP\) on page 16](#)

[Troubleshoot Multi-Chassis Link Aggregation \(MLAG\) on page 17](#)

[Simple Network Management Protocol \(SNMP\) on page 18](#)

Link Aggregation Control Protocol (LACP)

Use this topic for LACP (Link Aggregation Control Protocol) troubleshooting.

Troubleshoot Port-Channel Related Issues

You can troubleshoot port-channel issues.

About This Task

Follow this procedure to troubleshoot port-channel issues.

For troubleshooting port-channel issues, focus on LACP (Link Aggregation Control Protocol) since it manages dynamic port-channel creation, updates, and deletion. As a best practice, collect LACP-related logs, configs, and database contents, as described in the following table:

Data	Path/Command
RASTrace	/tmp/tierra/trace/lacp/RASTrace.log
Journal Control Logs	journalctl-b greplacp
Kubernetes Log	kubectllogs-ntierra<LACPPODNAME>
Service lacp docker container log	/var/log/containers/lacp-*
Config DB dump for interface metadata	show system internal cdb path/interfaces show system internal cdb path/lacp
State DB dump for interface metadata	show system internals db path/interfaces show system internals db path/lacp
Debug commands for LACP	system internal service lacp command <debug_command>
Debug Functions' dump in tech-support	/var/ms-commands/lacp.t

Procedure

1. Check whether LACP is enabled globally.

```
show running-config lacp
protocol lacp
```

2. Check the port-channel-related information.

```
show interface port-channel 10
port-channel 10 is up
  MTU 9216 Bytes
  IfIndex 0x400000a
  Mac address is 00:04:96:d6:55:47
  Port mode is Full Duplex, 100G
  MinLinks is 1
  LagType is LACP
  Active Members in this channel: Eth 0/23:1, Eth 0/23:2, Eth 0/23:3, Eth 0/23:4
  Members in this channel: Eth 0/23:1, Eth 0/23:2, Eth 0/23:3, Eth 0/23:4
```

```
system internal service lacp command dumpAllPortChannelData
```

```
Dumping All Port-channel Data Structures:
```

```
*****
```

```
-----
---
PoName       : port-channel 10
ID           : 10
If index     : 0x400000 a
Description  : Port-Channel Interface
Type        : LACP
MinLinks     : 1
AdminState   : UP
OperState    : UP
LagSpeed     : 7
Active List  : Eth 0/23:1, Eth 0/23:2, Eth 0/23:3, Eth 0/23:4
Member List  : Eth 0/23:1, Eth 0/23:2, Eth 0/23:3, Eth 0/23:4
```

```
Actor Aggregator Info
```

```
-----
Admin Key    : 10
System Id    : 00:04:96:d6:55:1c
System Priority : 32768
Oper Key     : 10
```

```
Partner Aggregator Info
```

```
-----
Oper Key     : 10
System Id    : 000496d683e0
System Priority : 32768
```

3. Check if member ports are OPER-UP and MinLinks condition is satisfied.

```
show interface brief | include 23:
Eth 0/23:1 9216   UP   UP   25G   0x10002e1   Ethernet 0/23:1
Eth 0/23:2 9216   UP   UP   25G   0x10002e2   Ethernet 0/23:2
Eth 0/23:3 9216   UP   UP   25G   0x10002e3   Ethernet 0/23:3
Eth 0/23:4 9216   UP   UP   25G   0x10002e4   Ethernet 0/23:4
```

- Confirm that lacp-mode is active on at least one side (from "show running-config lacp").

```
Interface ethernet 0/23:1
  Channel-group 10 active
  no shutdown
!
```

- Check whether sync, collecting, and distributing states are set.

```
show lacp interface ethernet 0/23:1
interface Eth 0/23:1 is up
Channel group is 10 port channel is Po10
  PDUs sent      : 15677
  PDUs rcvd      : 15726
  LACP Rx errors  : 0
  LACP TX errors  : 0
  LACP unknown errors : 0
  LACP errors     : 0
  Local Port      : Eth 0/23:1 MAC Address = 00:04:96: d6:55:1c
  System Identifier = 80:00:00:04:96:d6:55:1c
  Port Identifier  = 0x8000, 0x117
  Operational key  = 10
  LACP_Activity    = active
  LACP_Timeout     = Long Timeout (30s)
  Synchronization = IN_SYNC
  Collecting       = true
  Distributing     = true

system internal service lacp command dumpAllPortMemberData
Dumping All available Port-Member Data
*****
-----
  IntfName
  PortNum
  PortPriority
  AggId
  Agg AdminStatus
  ethernet 0/23:1
  PortEnabled
  LinkOperStatus

  Actor Port Info
  -----
  Admin Key      : 10
  Admin Port Pri : 32768
  Admin State    : ACTIVE, DEFAULT
  Oper Key       : 10
  Oper Port Pri  : 32768
  Oper State     : ACTIVE, AGGREGATING, SYNC, COLLECTING, DISTRIBUTING

  Partner Port Info
  -----
  Oper Key       : 10
  Oper Port Pri  : 32768
  Oper State     : ACTIVE, AGGREGATING, SYNC, COLLECTING, DISTRIBUTING
```

- Verify whether TX and RX counters are incrementing.

```
show counters lacp
Port          in-pkts  out-pkts  TxErr  RxErr  unknownErr  LacpErr
-----
Channel group: 10
ethernet 0/23:1  15720   15671     0      0      0           0
ethernet 0/23:2  15720   15671     0      0      0           0
```

ethernet 0/23:3	15720	15671	0	0	0	0
ethernet 0/23:4	15720	15671	0	0	0	0

Note the following about the packet counters:

- in-pkts increments every 1 second
- out-pkts increments every 10 seconds

7. Verify whether RX PDUs are actually received (in the file /var/log/containers/lacp-xxxxxx_lacp-xxxxxx.log).

```
16777956(0x10002e4) - RX:
00000000 01 80 c2 00 00 02 00 04    96 d6 83 ff 88 09 01 01
|-----|
00000010 01 14 80 00 00 04 96 d6    83 e0 00 0a 80 00 04 17
|-----|
00000020 3d 00 00 00 02 14 80 00    00 04 96 d6 55 1c 00 0a    |=-----U---|
00000030 80 00 04 17 3d 00 00 00    03 10 00 00 00 00 00 00    |----
=====|
00000040 00 00 00 00 00 00 00 00    00 00 00 00 00 00 00 00
|-----|
00000050 00 00 00 00 00 00 00 00    00 00 00 00 00 00 00 00
|-----|
00000060 00 00 00 00 00 00 00 00    00 00 00 00 00 00 00 00
|-----|
00000070 00 00 00 00 00 00 00 00    00 00 00 00                |-----|
```

8. Verify whether the TX offload information is correct for all member ports.

```
system internal service lacp command dumpAllTxOffloadInfo

Dumping TX Offload Information
*****
IntfIndex           : 16777953
Enabled              : true
Interval             : 30s
reconciliationInProgress : false
PDU Start
00000000 01 80 c2 00 00 02 00 04    96 d6 55 47 88 09 01 01    |-----
UG-----|
00000010 01 14 80 00 00 04 96 d6    55 1c 00 0a 80 00 01 17    |-----
U-----|
00000020 3d 00 00 00 02 14 80 00    00 04 96 d6 83 e0 00 0a    |
=====|
00000030 80 00 01 17 3d 00 00 00    03 10 00 00 00 00 00 00    |----
=====|
00000040 00 00 00 00 00 00 00 00    00 00 00 00 00 00 00 00
|-----|
00000050 00 00 00 00 00 00 00 00    00 00 00 00 00 00 00 00
|-----|
00000060 00 00 00 00 00 00 00 00    00 00 00 00 00 00 00 00
|-----|
00000070 00 00 00 00 00 00 00 00    00 00 00 00
|-----|
PDU End
```

9. Check whether PSDB data is proper and is in sync with the current state (required to support the restart of LACP-MS).

```
system internal service lacp command dumpAllPsDBData

Dumping All PSDB Data
*****
***** Dumping PSDB Data 1 *****
*****
Interface Number: ethernet 0/23:1
```

```

IfIndex: 16777953 (0x10002e1)      PortId: -2147483369      Key: 10
PortSpeed: 7      Port Number: 0x 117      Port Priority: 32768
AggId: 10
AggAttached:
AggId: 10 hwAggId: 67108874 aggDescription: "Port-Channel Interface" aggName:
"port-channel 10" aggType:1 aggMinLinks:1 aggMaxBundle:64 aggSpeed: 12 config:
{interval: 30000000000 mode:1 sys temIdMac:"00:04:96:d6:55:1c" systemPriority:32768}
actor Adminkey: 10 actorOperKey: 10 portAlgorithm: "\x00\x80\xc2\x01"
partnerSystemId:"\x00\x04\x964\xe0" partnerSystemPriority: 32768 adminState:true
ready:true portNumList:279 portN umList:535 portNumList:791 portNumList: 1047 Lag
Hash:1 is MemSpeedSame: trueAggSelected: 1 Port Enabled: true LACP Enabled: true
Begin: true ReadyN: true
Link Oper Status: true Log Enable: true MAC Properties: mac:"\x00\x04\x96\xd6UG"
speed:7 duplex:1 mtu:9216
Actor Admin: lacpSystem: {actor_System_priority: 32768 actor System:
"\x00\x04\x96\xd6U\x1c"} key: 10 port_pri: 32768 port: 279 state:65 Actor Oper:
lacpSystem: {actor_System_priority: 32768 actor System: "\x00\x04\x96\xd6U\x1c"} key:
10 port_pri: 32768 port:279 state:61 Partner Admin: lacpSystem: {actor_System:
"\x00\x00\x00\x00\x00\x00"} state:69
Partner Oper: lacpSystem: {actor_System_priority: 32768 actor System:
"\x00\x04\x964\xe0"} key: 10 port_pri: 32768 port:279 state:61
rxmPrevState: 1 ptxmPrevState: 1 txmPrevState: 4 actorCdmPrevState: 1
partnerCdmPrevState: 1 markerResponder PrevState: 1
rxmCurrState: 7 ptxmCurrState: 4 txmCurrState: 4 actorCdmCurrState: 2
partnerCdmCurrState: 6 arkerResponderCurrState: 2
AggPortDebug: AggPortDebugInformation ID: 279 AggPort DebugRxState:7
AggPortDebugLastRxTime:93257190 AggPortDebugMuxState:3 AggPort DebugMuxReason: "from
Port Config Event AggSelected equa ls Selected" AggPortDebugActorChurnState:2 Agg
PortDebugPartnerChurnState:6 AggPortDebugActorSync TransitionCount:1

```

Troubleshoot Link Layer Discovery Protocol (LLDP)

You can troubleshoot LLDP (Link Layer Discovery Protocol) issues.

About This Task

Follow this procedure to troubleshoot LLDP (Link Layer Discovery Protocol) issues.

The following are a few logs/configs/DB contents useful for debugging.

Data	Path/Command
RASTrace	/tmp/tierra/trace/lldp/RASTrace.log
Journal Control Logs	journalctl -b grep lldp
Kubernetes Log	kubect1 logs -n tierra <LLDP POD NAME>
Service lldp docker container log	/var/log/containers/lldp-*
Config DB dump for interface metadata	show system internal cdb / interfacesDevice# show system internal cdb /lldp
State DB dump for interface metadata	show system internal sdb / interfacesDevice# show system internal sdb /lldp
Debug commands for LLDP	system internal service lldp command <debug_command>
Debug Functions' dump in tech-support	/var/ms-commands/lldp.txt

Procedure

1. Verify LLDP Global Status.
 - a. Check if LLDP is enabled globally.
 - b. If disabled, LLDP will not send or receive neighbor information.
2. Verify LLDP on the Interface.
 - a. LLDP can be disabled on specific ports.
 - b. Check per-interface config.
 - c. Ensure RX & TX are enabled.
3. Check Neighbor Discovery. See if LLDP neighbors are being learned.

If the neighbor is missing, it could be:

 - Disabled LLDP on the neighbor
 - Connectivity issue
4. Verify LLDP timers.

Long Timers may delay neighbor visibility.

 - a. Check hello-time & hold time.
 - b. Ensure it aligns with the peer device.
5. Check if LLDP TX & RX Counters are incrementing at intervals.
6. Check Physical & Link Layer Health.
 - a. Verify if the links are up.
 - b. Check for packet/frame drops, errors in interface counters.
7. Check Microservice Health (8730 / OS ONE).

On OS ONE, LLDP runs as a microservice. Verify it's up. Check pod status.
8. Validate these steps on peer devices too.
9. Collect the Necessary tech support & debug Logs.

Troubleshoot Multi-Chassis Link Aggregation (MLAG)

You can troubleshoot MLAG (Multi-Chassis Link Aggregation) issues.

About This Task

Follow this procedure to troubleshoot MLAG (Multi-Chassis Link Aggregation) issues.

The following are a few logs or configs that are useful for debugging this agent.

Data	Path/Command to get the details
RASTrace	/tmp/tierra/trace/mlag/mlag-RASTrace.log /var/log/tierra/mlag
Service mlag docker container log	/var/log/containers/mlag-*
Config DB dump for interface metadata	show system internal cdb path /mlag

Data	Path/Command to get the details
State DB dump for interface metadata	show system internal sdb path /mlag
Persistent State DB dump for interface agent config	NA

Procedure

1. Verify the peer keepalive creation.
 - a. Ensure the Primary Keepalive's Source Interface is bound to the default VRF. Verify with the following command::


```
# show vrf default-vrf
```
 - b. Confirm the IfIndex is assigned to the Source Interface. Verify with the following command:


```
# show interface brief
```
 - c. Use the **netstat** command to verify the TCP connection
2. Verify the MAC sync issues between MLAG peers.
 - a. Verify MAC addresses are present in the mac-address-table.


```
show mac-address-table bridge-domain all
```
 - b. Check the internal gRPC connection between UFTM MS and MLAG MS.
 - c. Ensure interested Bridge Domains (BDs) are present on both MLAG peers.
 - d. Verify MLAG interfaces and ISL tunnels are present in UFTM.
3. Verify the following if the back node is not showing the primary node's MAC.
 - a. Check if System MAC, BDs, and Management IP are synced between MLAG peer nodes.
 - b. Use the curl command to debug and show received MAC, Management IP, Role, and BDs of the peer.


```
curl 0:9004/show-peer --data detail=1
```

Simple Network Management Protocol (SNMP)

For SNMP (Simple Network Management Protocol) troubleshooting, check the log file locations and review the limitations as described in the following table:

Category	Details
Device Commands	show running-config system snmp-server show running-state system snmp-server show system internal cdb /system/snmp-servers show system internal sdb /system/snmp-servers
Log Files	- tmp/tierra/trace/snmp-agent/snmp-agent-RASTrace.log /tmp/tierra/trace/snmp-agent/snmp-agent-RASTrace-DError.log /var/log/containers/snmp-agent-<containerID>

Category	Details
Core Files	/mnt/on1/usrdata/coredumps/
Net-SNMP Troubleshooting	• Check /etc/snmp/snmpd.conf configuration file. • View /var/log/snmpd.log for logs.
Limitations	• SNMP read-only operations are supported; set operations are not supported. • SNMP traps may be delayed due to retry mechanisms during device boot-up.




Installation and Deployment Issues

[Zero-Touch Provisioning \(ZTP\)](#) on page 20

Zero-Touch Provisioning (ZTP)

When troubleshooting ZTP (Zero-Touch Provisioning) issues, refer to log file locations and review the limitations as described in the following table:

Category	Description
Log Files	• MS-related logs: /tmp/tierra/trace/ztp/ztp-RASTrace.log • Python script logs: /var/data/ztp/ztpdhcp/log/pythonZtp.log • Switch startup config file: /var/data/ztp/ztpdhcp/log/configReplay.log • CLI logs: /var/data/ztp/ztpdhcp/log/ztp.log (accessible via show ztp logs)
Additional Info	• All log files are included in techSupport. • Abort ZTP with ztp dhcp cancel if it's stuck or misconfigured. • Debug ZTP config files from HTTPS servers if device errors point to parameter issues. • Ensure firmware version matches supported_nos in ZTP config file
Limitations	• No IPv6 support (DHCPv6 lacks options 66 and 67). • No inband interface support. • No app upgrade support via ZTP



Network Infrastructure Components Issues

[Address Resolution Protocol \(ARP\) and Neighbor Discovery \(ND\) on page 21](#)
[Reconstruct Configuration Database \(CDB\) from Tech Support on page 26](#)
[Reconstruct System Database \(SDB\) from Tech Support on page 26](#)
[Troubleshoot Bidirectional Forwarding Detection \(BFD\) on page 27](#)
[Troubleshoot Kernel Crash on page 29](#)
[L3 interface on page 30](#)
[Bridge Domain on page 31](#)
[Troubleshoot Application-Specific Integrated Circuit \(ASIC\) on page 31](#)
[Internal Message on page 32](#)

Address Resolution Protocol (ARP) and Neighbor Discovery (ND)

Use this topic to learn about troubleshooting address resolution protocol (ARP) and neighbor discovery (ND).

ARP Data Structure

Verify the following key data structures in ARP (Address Resolution Protocol):

- Interface Neighbor Cache (show-ifnbr)
 - This shows the neighbor cache entries/information associated with interfaces, which includes IP to MAC mappings in ARP.
 - Example: `curl 0:9003/show-ifnbr` or `curl --data "ifName=ve 100&af=4" 0:9003/show-ifnbr`
- MAC Database (show-macdb)
 - Displays the MAC database, which includes MAC-SRN requests for L2 interface resolution towards UFTM.
 - Example: `curl --data "ifName=ve 100" 0:9003/show-macdb`
- Nexthop Table (show-nhop-table)
 - Display the nexthop information for neighbors, important for routing decisions that use ARP to resolve next-hop MAC addresses.

- Example: `curl --data "ifName=ve 100" 0:9003/show-nhop-table`
- Additionally:
 - Interface Details (`show-if`) provides interface information but is not directly related to ARP data structures. For example, `curl --data "ifName=ve 100" 0:9003/show-if`
 - Clear Interface Neighbour Cache (`clear-if`) can be used to flush ARP entries. For example, `c`

MLAG Interfaces in ARP ND

To check MLAG (Multi-Chassis Link Aggregation) interfaces in ARP ND (Address Resolution Protocol Neighbor Discovery), use the following commands:

- **show-mlag-bds:** Displays the MLAG bridge domain database
- **show-mlag-intfs:** Displays the MLAG interface database
- **show-mlag-tunnels:** Displays the MLAG tunnel database

Media Access Control (MAC)-DB

MAC (Media Access Control) addresses are associated with L3 interface entries for ARP learned on VE interfaces. For ARP entries learned on VE interfaces, the MAC addresses requiring Layer 2 interface resolution are added to the MAC database.

After the Layer 2 interface resolution from UFTM, the MAC database includes both the Layer 2 interface and the MAC address. Until the MAC database is fully resolved, the ARP entries on VE interfaces are not forwarded to UFTM, meaning they do not appear in the **show ip route** output.

Nexthop Table

The global table for L3 IfEntry and MAC addresses includes a list of IP addresses awaiting L2 interface resolution for each entry. These IP addresses will be resolved to MAC addresses through L2 interface resolution in the MAC database. The table includes IP addresses from both VE and non-VE interfaces.

Entries Advertised Towards Unified Forwarding Table (UFTm) or MLAG

The following commands can be used to display information about advertised entries towards UFTm (Unified Forwarding Table) or MLAG (Multi-Chassis Link Aggregation). UFTm (Unified Forwarding Table – multicast) represents the multicast forwarding resources in hardware.:

- **show-ifnbr:** This command displays the list of advertised IP neighbors for each neighbor, including the source and destination clients.

- **show-ifnbr:** This command displays the list of advertised IP neighbors for each neighbor, including the source and destination clients.

Enabling Detailed Debugs in ARP ND

To enable detailed debugs in ARP ND (Address Resolution Protocol Neighbor Discovery), run the following curl command:

```
curl --data "feature=DevDebug&status=true" 0:9003/set-feature
```

This command enables detailed debugging with information like Goid, modname, and lineno. The output will display the current feature status, including the newly enabled DevDebug feature.

```
l3ProxyARP: false
l3LearnFRequest: true
ArpSuppress: false
DevDebug: true
ArpAgeTime : 1200.
```

Show Command Changed Output

The following show commands display ARP (Address Resolution Protocol) entries for specific interfaces or IP addresses.

Show IP ARP Interface

The **show ip arp interface** command displays ARP entries for a specified interface.

```
show ip arp interface ve 100
VRF:default-vrf
-----
Total number of IPv4 arp entries : 4
Ip Address      Mac Address      Type      Interface      L2 Interface      Age
=====
77.1.1.2        00:10:94:00:00:01 Local    ve 100         unresolved        4h49m28s
99.1.1.1        00:16:3e:7a:b3:03 Local    ve 100         unresolved        4h49m28s
10.10.10.2      00:10:94:00:00:03 Local    ve 100         unresolved        4h49m28s
67.1.1.111     00:10:94:00:00:11 Static   ve 100         unresolved        4h49m28s
```

This command displays the ARP (Address Resolution Protocol) entries for the specified interface (ve 100) on the default VRF (Virtual Routing and Forwarding) of the device. The output includes the following information:

- **Total number of IPv4 arp entries:** the total number of ARP entries for the specified interface.
- **IP Address:** the IP address for the ARP entry.
- **Mac Address:** the MAC address associated with the IP address.
- **Type:** the type of ARP entry, which can be Local , Static, Dynamic, MLAG-Static, MLAG, EVPN-Static, EVPN.
- **Interface:** the interface associated with the ARP entry.
- **L2 Interface:** the L2 (Layer 2) interface associated with the ARP entry.
- **Age:** the age of the ARP entry in hours, minutes, and seconds.

Show IP ARP VRF

The **show ip arp vrf** command displays ARP entries for a specified IP address in a specific VRF.

```
show ip arp vrf default-vrf 67.1.1.111
Type      : Static
-----
Mac-address      : 00:10:94:00:00:11      Sequence-no :
L3-interfane Name : ve 100                Flags       : 0
L2-interface Name : unresolved           Age         : 4h49m53s
vm2#
vm2# show ip arp vrf default-vrf 67.1.1.111
```

This command displays the ARP entry for the specified IP address (67.1.1.111) in the default VRF (default-vrf) of the device. The output includes the following information:

- Type: the type of ARP entry, which can be Local , Static, Dynamic, MLAG-Static, MLAG, EVPN-Static, EVPN.
- Mac-address: the MAC address associated with the IP address.
- Sequence-no: the sequence number associated with the ARP entry.
- L3-interface Name: the Layer 3 interface associated with the ARP entry.
- Flags: any flags associated with the ARP entry.
- L2-interface Name: the Layer 2 interface associated with the ARP entry.
- Age: the age of the ARP entry in hours, minutes, and seconds.

Service Access Gateway (SAG) IP or MAC Behavior

Use this topic to learn about the SAG (Service Access Gateway) IP or MAC (Media Access Control) behavior for different scenarios:

Sending ARP Responses

- When an ARP request is received for the **SAG IP**, the ARP daemon (ARP ND) replies using the **SAG MAC** as both the source MAC in the Ethernet frame and in the ARP response.
- If the ARP request is for the **Physical IP**, ARP ND responds with the **Physical MAC** as the source MAC in both the Ethernet frame and the ARP reply.

Sending ARP Refresh/Request

- The **source MAC** and **sender MAC** in ARP refreshes or requests are always set to the **Physical MAC** (this is currently a known defect and should be fixed to use the SAG MAC).
- The **sender IP** is set to the **SAG IP** if available; otherwise, it uses the **Physical IP**.

Retries Behavior

- For retries up to half the maximum allowed, the **destination MAC (DMAC)** is set to a unicast MAC address.
- After half the maximum retries, the DMAC changes to a broadcast MAC address.

- For IPv6, the broadcast MAC is replaced with a multicast MAC address; unicast MAC remains unchanged.

Table 4: Summary Table of IP/MAC Behavior for ARP/Neighbor Discovery

Type	With/without SAG	SRC MAC (L2)	DEST MAC L2	Sender MAC	Sender IP	Remarks
Ipv4	only physical IP	PHY MAC	Broadcast	PHY MAC	PHY IP	New ARP request
Ipv4	only physical IP	PHY MAC	unicast	PHY MAC	PHY IP	ARP response
Ipv4	only physical IP	PHY MAC	Unicast/Broadcast	PHY MAC	PHY IP	ARP refresh
Ipv6	only physical IP	PHY MAC	Multicast	PHY MAC	PHY IP	New ARP request
Ipv6	only physical IP	PHY MAC	Unicast	PHY MAC	PHY IP	ARP response
Ipv6	only physical IP	PHY MAC	Unicast/Multicast	PHY MAC	PHY IP	ARP refresh
Ipv4	SAG+physical IP	SAG MAC	Broadcast	SAG MAC	SAG IP	New ARP request
Ipv4	SAG+physical IP	SAG MAC	unicast	SAG MAC	SAG IP	ARP response (SAG IP request)
Ipv4	SAG+physical IP	PHY MAC	unicast	PHY MAC	PHY IP	ARP response (Phy IP request)
Ipv4	SAG+physical IP	SAG MAC	Unicast/Broadcast	SAG MAC	SAG IP	ARP refresh
Ipv6	SAG+physical IP	SAG MAC	Multicast	SAG MAC	SAG IP	New ARP request
Ipv6	SAG+physical IP	SAG MAC	Unicast	SAG MAC	SAG IP	ARP response (SAG IP request)
Ipv6	SAG+physical IP	PHY MAC	Unicast	PHY MAC	PHY IP	ARP response (Phy IP request)
Ipv6	SAG+physical IP	SAG MAC	Unicast/Multicast	SAG MAC	SAG IP	ARP refresh
Ipv4	only SAG IP	SAG MAC	Broadcast	SAG MAC	SAG IP	New ARP request
Ipv4	only SAG IP	SAG MAC	unicast	SAG MAC	SAG IP	ARP response
Ipv4	only SAG IP	SAG MAC	Unicast/Broadcast	SAG MAC	SAG IP	ARP refresh

Table 4: Summary Table of IP/MAC Behavior for ARP/Neighbor Discovery (continued)

Type	With/without SAG	SRC MAC (L2)	DEST MAC L2	Sender MAC	Sender IP	Remarks
Ipv6	only SAG IP	SAG MAC	Multicast	SAG MAC	SAG IP	New ARP request
Ipv6	only SAG IP	SAG MAC	Unicast	SAG MAC	SAG IP	ARP response
Ipv6	only SAG IP	SAG MAC	Unicast/Multicast	SAG MAC	SAG IP	ARP refresh

Reconstruct Configuration Database (CDB) from Tech Support

You can reconstruct CDB (Configuration Database) from tech support.

About This Task

Follow this procedure to reconstruct CDB from tech support using Valkey CDB.

Procedure

Start the Valkey server inside a container from the Tech Support files.

- a. Navigate to the directory containing Valkey files.

```
cd /home/<username>/TOS-25831/TS-DUT3-250205T1257/mnt/onl/config/db/cdb
```

- b. List the available 'aof' files.

```
ls
# cdb.aof.31.base.aof  cdb.aof.manifest  cdb.aof.31.incr
```

- c. Run a Valkey container, mounting the directory with the aof files.

```
docker run --rm -v /home/<username>/TOS-25831/TS-DUT3-250205T1257/mnt/onl/config/db/
cdb:/data --name construct-cdb engartifacts1.extremenetworks.com:8099/tierra-os/
valkey/valkey:8.0.1-alpine3.20 valkey-server --appendonly yes --appendfilename
cdb.aof.31.incr.aof
```

- d. Connect to Valkey to verify the data.

```
docker exec -it construct-cdb valkey-cli
127.0.0.1:6379> keys *
```

Reconstruct System Database (SDB) from Tech Support

You can reconstruct SDB (System Database) from tech support.

About This Task

Follow this procedure to reconstruct SDB from tech support using Valkey CDB.

Procedure

Start the Valkey server inside a container from the Tech Support files.

- a. Navigate to the directory containing Valkey files.

```
cd /home/<username>/mctissue/tech-support/
```

- b. List the available aof files.

```
find . -name sdb_dump.rdb
./mnt/onl/config/db/sdb_dump.rdb
```

- c. Run a Valkey container, mounting the directory with the rdb files.

```
docker run --rm -v /home/<username>/mctissue/tech-support/mnt/onl/config/db:/
data --name construct-sdb engartifacts1.extremenetworks.com:8099/tierra-os/valkey/
valkey:8.0.1-alpine3.20 valkey-server --dbfilename sdb_dump.rdb
```

- d. Connect to Valkey to verify the data.

```
docker exec -it construct-sdb valkey-cli
127.0.0.1:6379> keys *
```

Troubleshoot Bidirectional Forwarding Detection (BFD)

You can troubleshoot issues with BFD (Bidirectional Forwarding Detection).

About This Task

Follow this procedure to successfully establish a BFD session on hardware.

BFD sessions and clients are managed by the BFD microservice. To debug BFD-related issues, collect or debug BFD-related data.

The following are the useful logs and configs:

Data	Path/Command to get details
RASTrace	/tmp/tierra/trace/bfd/bfd-RASTrace*.log /tmp/tierra/trace/bfd/bfd-RASTrace-Excess*.log /tmp/tierra/trace/bfd/bfd-RASTrace-DError.log
BFDLogstoconsole	/var/log/tierra/bfd/bfd.log
ConfigDBdump	show system internal cdb path /bfd
StateDBdump	show system internal sdb path /bfd
PersistentStateDBdump	Not Applicable

Procedure

1. **Check Enable-BFD Config:** Verify if enable-bfd config is present for BGP neighbor, peer-group, or static-route.
2. **Check BFD Profile:** Ensure a specific profile is mentioned in the above config with appropriate intervals and multipliers.
3. **ARP/ND Resolution:** Verify if ARP or ND is resolved for the BFD Destination IP using the CLI.
4. **Ping Test:** Confirm reachability by verifying if ping to the destination IP is successful.

5. **Rx Packets:** Check if Rx packets are observed from session counters (show bfd neigh details).
6. **Hardware Neighbor Objects:** Verify if hardware neighbor objects for this dest-ip have proper MAC, L2/L3 interface details.
7. **Hardware Egress Objects:** Verify if hardware egress objects for the destination IP have the correct MAC address and Layer 2/Layer 3 interface details.

Troubleshoot BFD Flap Issues

You can troubleshoot BFD (Bidirectional Forwarding Detection) flap issues.

About This Task

Follow this procedure to identify and isolate the module responsible for BFD flap issues in most common cases.

Key Points for Troubleshooting and Isolating Modules

- Check for issues in BFD-ms when:
 - BFD create requests are not sent across the HAL interface
 - Profile information is not filled correctly
 - There is a mismatch in state info, diag info, or other info between CLI show output and curl command output
 - There is a mismatch in session states between FWD HAL and BFD-ms
- For issues related to BFD flap or session state INIT/DOWN, correlate with other system events, such as:
 - Interface flaps
 - ARP entry flush/relearn
 - NHID, L2Intf, and route info correctness
 - ARP/route entries in FWD-HAL DB

Procedure

1. Find the Session ID for Curl Commands.

To get the session ID for a given destination IP, use the show bfd neighbor command with the details option.

```
48y# sh bfd nei dest-ip 10.1.1.2 vrf all details
```

The output includes a Session ID, which can be used to retrieve other internal BFD states

The following are the Deriving Session ID:

- For TD3/TD4/VT
 - Session ID = LD value - 1
- For Jericho2 platforms
 - For IPv4 sessions: Session ID = (LD-24000)/4
 - For IPv6 sessions: Session ID = LD/4

2. Check the StateDB Contents.

Use the **show-bfd-session** curl command to dump internal states and information received from FWD-HAL.

```
curl 10.52.1.52:9005/show-bfd-session --data sessionID=1
```

The output includes:

- Local and remote session states
- Profile information
- Interface and client details
- Session ID and HW session ID

On 8820 platform, the software initiates the session and hands it over to hardware, capturing Tx/Rx packet information and FSM transitions.

3. Check the Interface Info.

Use the **show-common-session-info** curl command to display interface information as received from Interface-Mgr's state DB updates.

```
curl 0:9005/show-common-session-info --data sessionID=1
```

The output includes:

- SIP, DIP, AFI, VRF ID, interface index, and interface name
- Session ID and profiles

4. Check the Client Communication.

Use the **show-client** curl command to display client communication information.

```
curl 0:9005/show-client --data clientName=STATIC_ROUTE
```

The output includes:

- Create, update, and delete requests and responses
- Session status reports sent to the client

Troubleshoot Kernel Crash

You can troubleshoot a kernel crash issue.

About This Task

Follow this procedure to troubleshoot the kernel crash issue.

By analysing the backtrace and crash dump, you can determine the root cause of the kernel crash and take the necessary steps to fix the issue.

When a kernel crash occurs, a vmcore is generated and stored in the `/mnt/on1/usrdata/coredumps` directory.

Procedure

1. Copy the vmcore: Transfer the vmcore to a build machine for analysis.
2. Download the kernel file: Obtain the kernel file (vmlinux-6.6.54-yocto-standard) from the specified location and save it to the build machine.

3. Install the crash tool: Ensure the crash tool (version 8 or higher) is installed on the build machine.
4. Run the crash tool.
The example output from running crash shows kernel and system details, including CPU, uptime, panic message, and backtrace (bt) of the crashing task.

Example

- **Crash Analysis**

To analyze the crash, use the **crash vmlinux-6.6.54-yocto-standard vmcore-23_24_18** command. This will initiate the crash analysis, providing detailed information about the kernel crash, including the backtrace.

- **Backtrace Analysis**

The backtrace provides a step-by-step view of the kernel crash, including the functions and modules involved. Analyze the backtrace to identify the root cause of the issue.

```
crash> bt
PID: 2796 TASK: ffff952093775580 CPU: 16 COMMAND: "IQPKT, TX"
#0 [ffffb0a8c0933a48] machine_kexec at ffffffff6f04151
#1 [ffffb0a8c0933aa0] __crash_kexec at ffffffff705019e
#2 [ffffb0a8c0933b28] __dev_queue_xmit at ffffffff7f4a290
#3 [ffffb0a8c0933c08] exc_page_fault at ffffffff848ad1c
#4 [ffffb0a8c0933c30] asm_exc_page_fault at ffffffff8601316
```

L3 interface

Use this topic to learn about troubleshooting issues related to L3 interface.

For L3 interface-related issues, focus on the following microservices:

- Interface-mgr
- Fwd-hal

Logs and Configs

- **RASTrace:** /tmp/tierra/trace/<service-name>/*-RASTrace*.log
- **Kubernetes Log:** kubectl logs -n tierra <POD NAME>
- **Container logs:** /var/log/containers/<service-name>

ECMP and RH Debugging

- Check ECMP group mode (RH or Normal)
- Monitor ECMP table usage (limit: 16K for RH mode)
- Look for warning messages in logs: "ECMP resource utilization exceeded RH threshold limit."

Limitations and Restrictions

- Mode changes (Normal to RH or vice versa) are disruptive.
- Exceeding 50% ECMP shared table capacity with RH mode groups changes requests to Normal mode.
- No restrictions on Normal mode ECMP group resource usage.

Bridge Domain

Use this topic to learn about troubleshooting issues related to bridge domain.

- **Services Involved:** Service interface-mgr, mftm, and fwd-hal handle bridge domain creation, updates, and deletion.
- **Debugging Data:** Collect logs and configs from these services, including:
 - RASTrace logs: `/tmp/tierra/trace/<service-name>/*-RASTrace.log`
 - Kubernetes logs: `kubect1 logs -n tierra <classifiers POD NAME>`
 - Container logs: `/var/log/containers/<service-name>`
- **Tech Support:** Generate `tech-support tgz` for further debugging.

Limitations and Restrictions

TD4 Limitations

- Double-tagged traffic on single-tagged LIF: Inner VLAN remains unmodified, only outer VLAN is updated.
- Double-tagged traffic on single-tagged LIF to untagged LIF: Inner VLAN remains unmodified, only outer VLAN is removed

J2 Limitations

- Mac-aging-time programmed in hardware may differ from user-configured time due to meta-cycle-time and number-of-cycles adjustments.
- CDB stores user-configured time, while SDB stores hardware-programmed time

Troubleshoot Application-Specific Integrated Circuit (ASIC)

You can troubleshoot issues related to ASIC (Application-Specific Integrated Circuit).

About This Task

Follow this procedure to troubleshoot issues with ASIC.

Procedure

1. Verify ASIC Initialisation.
 - a. Look for log messages in RASLogs: `InitAdapter: BCM adapter initialization failed.`
 - b. Verify HW device detection using `lspci | grep Broadcom` command.

```
15:00.0 Ethernet controller: Broadcom Inc. and subsidiaries Device b891 (rev 01)
```

2. Verify ASIC Init RAS Logs.
 - a. Check for debug logs from fwd-hal service, including:
 - Platform type and chassis name
 - InitAdapter success or failure messages

The following is an example of log entries:


```
{ "Level": "debug", "Service": "fwd-hal", "LogID": 0, "Topic": 4, "Calling entrypoint with ": "x86_64-extremenetworks-8730-32d-r0", "Time" 2025-08-05 11:12:35.098 UTC +0000", "Msg": "as ChassisName" }

{ "Level": "debug", "Service": "fwd-hal", "LogID": 0, "Topic": 4, "Time": "2025-08-05 11:12:35.098 UTC +0000", "Msg": "Platform Type 3" } { "Level": "debug", "Service": "fwd-hal", "LogID": 0, "Topic": 4, "entrypoint: return code ": 0, "Time": "2025-08-05 11:12:40.462 UTC +0000", "Msg": "InitAdapter: BCM adapter initialized successfully" }
```

Internal Message

Use the **show system internal msg-svc** command to display information about the consumers, topics, and connected client applications for the internal NATS messaging service.

For detailed information on command syntax, parameters, and usage guidelines, see the *Extreme OS ONE Switching and Routing v22.2.2.0 Command Reference*.



Security Management Issues

- [Secure Shell \(SSH\) Services](#) on page 33
- [Telnet Services](#) on page 34
- [Token Management](#) on page 35
- [Troubleshoot Terminal Access Controller Access Control System Plus \(TACACS+\) Management](#) on page 36
- [Troubleshoot Radius Management](#) on page 36
- [Network Time Protocol \(NTP\) Management](#) on page 37
- [Troubleshoot Lightweight Directory Access Protocol \(LDAP\) Management](#) on page 38
- [User Management](#) on page 39
- [CLI Shell](#) on page 39
- [Troubleshooting Hardware Ternary Content-Addressable Memory \(TCAM\) Infra and System Access Control Lists \(ACL\)](#) on page 39
- [Troubleshooting User Access Control Lists \(ACL\)](#) on page 41
- [Troubleshooting gRPC Network Management Interface \(gNMI\) and gRPC Network Operations Interface \(gNOI\)](#) on page 44
- [Troubleshoot Service Access Gateway \(SAG\)](#) on page 45

Secure Shell (SSH) Services

You can troubleshoot issues related to SSH (Secure Shell).

Use this topic to learn about device-level commands, log files, and tech support files for troubleshooting SSH services.

Device-Level Commands

- **show running-config system ssh-server**
- **show running-state system ssh-server**
- **show system internal cdb path /system/ssh-servers**
- **show system internal sdb path /system/ssh-servers**
- **systemctl status sshd**

Log Files to Check

- /tmp/tierra/trace/security/security-RASTrace.log
- /tmp/tierra/trace/security/security-RASTrace-DError.log
- /tmp/tierra/trace/base-svc/base-svc-RASTrace.log
- /tmp/tierra/trace/base-svc/base-svc-RASTrace-DError.log
- /var/log/containers/security-<containerID>
- /var/log/daemon.log
- /var/log/auth.log

SSH Server Configuration

```
/etc/ssh/sshd_config
```

Core Files

```
/mnt/onl/usrdata/coredumps/
```

Tech Support Files

Located in TS-sanity-<timestamp> directory:

- ShowOutput.txt
- security/security-RASTrace.log
- security/security-RASTrace-DError.log
- base-svc/base-svc-RASTrace.log
- base-svc/base-svc-RASTrace-DError.log
- /etc/ssh/sshd_config
- /var/log/daemon.log
- /var/log/auth.log
- /var/log/containers/security-<containerID>
- /mnt/onl/usrdata/coredumps/

Telnet Services

You can troubleshoot issues related to tenet services.

Use this topic to learn about device-level commands, log files, and tech support files for troubleshooting telnet services.

Device-Level Commands

- **show running-config system telnet-server**
- **show running-state system telnet-server**

- `show system internal cdb path /system/telnet-servers`
- `show system internal sdb path /system/telnet-servers`
- `systemctl status xinetd*`

Log Files to Check

- `/tmp/tierra/trace/security/security-RASTrace.log`
- `/tmp/tierra/trace/security/security-RASTrace-DError.log`
- `/tmp/tierra/trace/base-svc/base-svc-RASTrace.log`
- `/tmp/tierra/trace/base-svc/base-svc-RASTrace-DError.log`

Tech Support Files

Located in TS-sanity-<timestamp> directory:

- `ShowOutput.txt`
- `security/security-RASTrace.log`
- `security/security-RASTrace-DError.log`
- `base-svc/base-svc-RASTrace.log`
- `base-svc/base-svc-RASTrace-DError.log`
- `/var/ms-commands/base-svc`

Token Management

You can troubleshoot issues related to token management.

Use this topic to learn about common errors and possible causes that occur during token management.

Investigate and resolve token-related issues by examining syslog messages and verifying token validity.

- **Token Expiration Error**
 - Syslog message: "Token validation failed. Reason: token expired"
 - Possible causes: Indicates that the token has exceeded its valid period.
- **Token Validation Error**
 - Syslog message: "Token validation failed. Reason: Token used in this session is not valid."
 - Possible causes: Invalid token, signature verification failure (for example, "crypto/rsa: verification error").

Troubleshoot Terminal Access Controller Access Control System Plus (TACACS+) Management

You can troubleshoot issues related to TACACS+ (Terminal Access Controller Access Control System Plus).

About This Task

Follow this procedure to troubleshoot issues related to TACACS+.

Procedure

1. Verify connectivity.

Run the following command to check the connection to the TACACS+ server:

```
ping <server_IP> vrf <vrf_name>
```

For example,

```
ping 10.24.15.200 vrf mgmt-vrf
```

2. Check server configuration.

- a. Ensure the server address and credentials are correct.
- b. Verify the shared secret matches the Remote TACACS+ server.

3. Analyze packet capture.

Check for errors like "Invalid AUTHEN/START packet (check keys)" indicating a mismatch in the shared secret.

4. Review authentication logs.

- a. Verify the SSH authentication failures.

```
/var/log/auth.log
```

- b. Verify the gNMI (gRPC Network Management Interface) authentication failures.

```
/tmp/tierra/trace/security/security-RASTrace.log
```

5. Review device configuration and logs.

- **show running-config system aaa server-group tacacs+**
- **show running-state system aaa server-group tacacs+**
- `/etc/tacacs.conf`
- `/var/log/auth.log`
- `security-RASTrace.log` and `base-svc-RASTrace.log` for additional error details

6. Review Tech Support file, such as `TS-sanity-<timestamp>/var/data/disk/techsupport/ShowOutput.txt` and similar logs for further troubleshooting.

Troubleshoot Radius Management

You can troubleshoot issues related to the RADIUS (Remote Authentication Dial-In User Service) server.

About This Task

Follow this procedure to troubleshoot issues related to the RADIUS server.

Procedure

1. Check Connectivity.

- a. Ping the RADIUS server.

```
ping 10.24.15.200 vrf mgmt-vrf
```

- b. Verify server address and credentials.

2. Verify Configuration.

- a. Check CA certificates for secured RADIUS.
- b. Ensure user role is properly configured on the server.

3. Debug and Logs.

- a. Run FreeRADIUS in foreground for debugging.

```
udo /usr/sbin/freeradius -fxx -l stdout
```

- b. Check device logs and tech support output.

```
show running-config system aaa server-group radius
show running-state system aaa server-group radius
show system internal cdb path /system/aaa/server-groups/server-
group[name=radius]
```

Check the other log files, such as `/var/log/auth.log` and `/tmp/tierra/trace/security/security- RASTrace.log`).

Note that the GNMI command accounting is not supported.

Results

Network Time Protocol (NTP) Management

Use this topic to learn about troubleshooting issues with NTP (Network Time Protocol).

Device Commands

1. Check NTP configuration.

```
show running-config system ntp
show running-state system ntp
```

2. Verify NTP status.

```
show ntp status
show ntp association
```

3. Check system logs.

```
show system internal cdb path /system/ntp
show system internal sdb path /system/ntp
```

Log Files

Check various log files for errors:

- `/tmp/tierra/trace/security/security-RASTrace.log`
- `/tmp/tierra/trace/security/security-RASTrace-DError.log`

- /tmp/tierra/trace/base-svc/base-svc-RASTrace.log

Additional Troubleshooting Steps

- Check NTP service status:

```
systemctl status ntpd
```

- Test NTP connectivity:

```
ping <address>  
ntpdate -q <address>
```

- Use `vrfcmd.sh` to test NTP on specific VRF:

```
vrfcmd.sh <vrf-id> ping <address>  
vrfcmd.sh <vrf-id> ntpdate -q <address>
```

Troubleshoot Lightweight Directory Access Protocol (LDAP) Management

You can troubleshoot issues with LDAP (Lightweight Directory Access Protocol).

About This Task

Follow this procedure to troubleshoot issues with LDAP.

Procedure

1. Check Connectivity: Check connectivity with the remote LDAP server over the configured VRF.

```
ping 10.24.15.200 vrf mgmt-vrf
```
2. Verify Server Details: Verify the server address, credentials, and CA certificates (for secure LDAP) of the configured server.
3. Verify Role Mapping: Ensure user group roles match supported device roles.
4. Verify device logs and configs.
 - show running-config system aaa server-group ldap
 - show running-state system aaa server-group ldap

Check the following internal logs and config files:

- var/log/auth.log
 - /etc/ldap.conf
 - security-RASTrace.log
 - base-svc-RASTrace.log
5. Review tech support files and trace logs for errors.
 - TS-sanity-<timestamp>/var/data/disk/techsupport/ShowOutput.txt
 - TS-sanity-<timestamp>/tmp/tierra/trace/security/security-RASTrace.log
 - TS-sanity-<timestamp>/tmp/tierra/trace/base-svc/base-svc-RASTrace.log

User Management

Use this topic to learn about device-level commands, log files, and tech support files for troubleshooting user management.

Device-Level Commands

- show running-config system aaa authentication
- show running-state system aaa authentication
- show system internal cdb path /system/aaa/authentication
- show system internal sdb path /system/aaa/authentication

Log Files

- /tmp/tierra/trace/security/security-RASTrace.log
- /tmp/tierra/trace/security/security-RASTrace-DError.log
- /tmp/tierra/trace/base-svc/base-svc-RASTrace.log
- /tmp/tierra/trace/base-svc/base-svc-RASTrace-DError.log
- /etc/shadow

Tech Support Files

Located in TS-sanity-<timestamp> directory:

- ShowOutput.txt - security/security-RASTrace.log
- security/security-RASTrace-DError.log
- base-svc/base-svc-RASTrace.log
- base-svc/base-svc-RASTrace-DError.log
- /etc/shadow copy

CLI Shell

Check the following logs for CLI shell troubleshooting:

```
var/log/tierra/cli/*
```

```
TS-sanity-<timestamp>/var/log/tierra/cli/*
```

Troubleshooting Hardware Ternary Content-Addressable Memory (TCAM) Infra and System Access Control Lists (ACL)

Use this topic to learn about troubleshooting hardware TCAM (Ternary Content-Addressable Memory) infra and system ACLs (Access Control List).

Debug Commands

Switch Commands

- **system internal service fwd-hal command xecute-adapter-cmd args func=acl-show-groups**
- **system internal service fwd-hal command xecute-adapter-cmd args func=acl-show-entries [param0=<group_id>] [argList=0xRBID]**
- **system internal service fwd-hal command xecute-adapter-cmd args func=acl-clear-counters [param0=<group_id>] [argList=0xRBID]**

Shell Commands

- fp commands:
 - **show [stage|group|entry|presel|hint] [id] [brief]**
 - **list actions/qualifiers <stage> <presel>**
 - **list dropcodes <stage> <qual> [<dropcode>]**
 - **list cpureasons <stage> [<cpureason>]**
- flexctr commands:
 - **stat show <statPoolID> <statID>**
 - **stat clear <statPoolID> <statID>**

Key Concepts

- Logical Groups: Created at device init, with software index used by HAL
- Stats Action ID: Td4 uses flex counter, requiring stats pool and stat ID allocation
- DFP: Hardwired group, not logical, without h/w ID
- Group Priority: For action resolution
- Preselector: Splits logical group to multiple groups using preselect TCAM

Understanding Output

- TCAM Logical Groups and Stats Pools
 - Initialization: Created at device init with software index for HAL data structures
 - Stats Action ID: Td4 uses flex counters, allocating stats pool and stat IDs (number of stat IDs = number of possible entries)
- Key Terms
 - DFP: Hardwired group without hardware ID
 - Qset and Actions: Qualifiers and actions in a group
 - Mode and Width: Single, double, triple, or intraslice-double
 - Stage: IFP (Ingress), VFP (Lookup), EFP (Egress)
 - Hint: Identifies stats pool internally
 - Entries: Total and free entries in a group
 - Preselector: Splits logical groups into subgroups (not used in this case)

- Group Details
 - Group Priority: For action resolution
 - Physical Slices: 2 slices (slice-0 and slice-1) for double-wide groups
 - Allotted Entries: 1023 (1 less than total 1024 for make and break)
 - Counters: Allotted and free
- System ACL and Control PDU
 - GroupSwIndex: SW index of the group
 - Total Entries: Total entries in the group
 - Packet Count: Incrementing when packets hit entries
- Show Entries
 - Param0 and Param1: Specify group SW index and RBID (Router Bridge ID)
 - Packet Count: Value of packets hitting entries, incrementing with each hit

Troubleshooting User Access Control Lists (ACL)

Use this topic to learn about troubleshooting user ACLs (Access Control List).

User ACL Features and Classifier Microservices

Classifier functionality is achieved through microservices running in the Gen4OS. These microservices are the building blocks of classifier features.

Configuration Flow

1. User configures device via CLI or gNMI.
2. CLI sends gNMI message to API-GW.
3. API-GW processes and validates config, writing it to CDB.
4. API-GW publishes config to message bus and responds to client.
5. Classifier MS collects message and processes it.
6. Classifier MS sends message to forward HAL to program hardware.
7. API-GW collects state and internal data from message bus.
8. API-GW publishes data to SDB and PSDB.
9. Hardware filters packets and applies actions.
10. Telegraf queries forward HAL for statistics and updates SDB.

Classifier Microservice

- Collection of features (ACL, QoS, Route Policy, and Mirror) for classifying and treating packets. Packed as feature packages and running as separate threads within the classifier microservice.
- Common packages (IntfMgmt, BdMgmt, Logging, Hal Msg, MsgPublish) provide base building blocks for features and act as bridge between classifier features and other microservices.

Debug Commands

- `system internal service classifiers command help`
- `system internal service classifiers command show-classifier`
- `system internal service classifiers command show-acl`
- `system internal service classifiers command show-acl-attachment`
- `system internal service classifiers command show-acl-bd-db`
- `system internal service classifiers command show-acl-binding`
- `system internal service classifiers command show-acl-ifindex-db`


Verification

1. Check classifier service status.
2. Verify classifier's common infra packages have proper data.
3. Attach ACL in CLI and corresponding classifiers output, and then verify classifier output.
4. Verify the ACL Binding Data and get Rule Block ID (RBID).
 - RBID Generation: The classifier's binding logic combines data to create a unique 64-bit RBID for each ACL.
 - RBID Purpose: This RBID is used to identify the ACL across the system and is sent to HAL for programming.
 - HAL Programming: The RBID is a programmable entity in hardware, reflected in HAL DBs and BCM shell, serving as a key reference point.

RBID

- 64-bit number uniquely identifying an ACL across the system.
- Created by binding logic in classifier.
- Sent to HAL for programming and reflected in HAL DBs and BCM shell.

User ACL End-To-End Debug Flow Chart



Troubleshooting gRPC Network Management Interface (gNMI) and gRPC Network Operations Interface (gNOI)

When troubleshooting gNMI (gRPC Network Management Interface) and gNOI (gRPC Network Operations Interface) issues, refer to log file locations and consider the listed limitations for effective debugging.

- Verify gRPC Server Status: Run the **show running-state system grpc-server** command to ensure at least one instance is enabled.
- View gNMI Statistics: Use the **show grpc-server gnmi statistics** command to view external stream subscription details.
- Test with Open-Source gNMI Client: Utilize `gnmic` with the `--debug` flag for connection troubleshooting and error response analysis.
- Analyze gRPC Error Codes: Inspect error codes and response messages in proprietary clients for root cause analysis.
- Timeouts and Connection Issues: If a client starts a streaming RPC but doesn't consume responses, a 10s timeout is applied. Subscription is canceled. Connection timeouts depend on client-specified context during connection setup.
- Device Logs
 - Service logs: `/tmp/tierra/trace` (API Gateway, Security, Ingress Gateway)
 - Show trace command can be used to verify `/tmp/tierra/trace` logs per microservice
 - Audit Logs: show logging audit config

Device Logs and Debugging

- Service Logs: Check `/tmp/tierra/trace` for API Gateway, Security, and Ingress Gateway logs. Use **show trace** command to verify logs per microservice.
- Audit Logs: Run **show logging audit config** for Syslog Messages and Event Logs.
 - Key Event IDs:
 - 6001: gRPC request completed successfully
 - 6009: Failed to initialize microservice
 - 12002: gRPC server connection status
 - 12003: gRPC server termination status

Tech Support Checklist

- Device-Specific Information: Check `var/data/disk/techsupport/ShowOutput.txt`.
- Microservice-Related Data: Refer `var/ms-commands/base-svc.txt` for pod statuses, deployment configuration, and memory usage.
- Resource Consumption: Run **kubectl top pods -A** to view resource consumption.
- Pod Status: Run **kubectl get pods -A** to list pod statuses.

- Container-Specific Logs: Check `/var/log/containers` for detailed backtrace information.
- Core Dumps: Saved at `mnt/on1/usrdata/coredumps/`. Use `dlv` tool to debug core files.

Restrictions and Limitations

- gRPC Server Configuration:
 - Unique port number and VRF name combination required.
 - Modifying settings not allowed when enabled.
- Subscription Limitations:
 - Duplicate subscription paths not supported.
 - Wildcard (*) keypaths not supported in Get and Set RPCs.
- Scalability:
 - Supports up to 32 concurrent sessions.
 - Handles a maximum of 256 subscription paths.
- gNMI RPC Limitations:
 - Not fully compliant with the latest gNMI specification.
 - Poll-based subscriptions not supported.
 - Certain YANG paths not stored in SDB and require explicit querying.

YANG Paths

```
/network-instances/network-instance/protocols/protocol/bgp/rib
/network-instances/network-instance/tables-network/ipv4-unicast
/network-instances/network-instance/tables-network/ipv6-unicast
/network-instances/network-instance/tables-network/ethernet
/network-instances/network-instance/tables-network/next-hop-groups
/network-instances/network-instance/tables-network/next-hops
/network-instances/network-instance/tables-network/ipv4-neighbor
/network-instances/network-instance/tables-network/ipv6-neighbor
```

Troubleshoot Service Access Gateway (SAG)

You can troubleshoot SAG (Service Access Gateway) issues.

About This Task

Follow this procedure to troubleshoot SAG issues.

Procedure

1. Verify VE Interface Status: Ensure the VE interface is up and attached to a VRF.
2. Check SAG MAC Termination.
 - a. For XGS-based devices: Use the `I2 show` command from the BCM shell to verify that the SAG MAC address is correctly programmed in hardware.
 - b. For DNX-based devices: Use the `system internal service fwd-hal` command `adapter-function-exec args func=I2-show-bd` command to check the SAG MAC configuration.

3. Verify Tunnel Status. Check if the tunnel is up. If not:
 - a. Verify configuration correctness.
 - b. Check for tunnel endpoint reachability issues.
4. Troubleshoot Traffic Drop Issues.
 - a. Check if the tunnel is up.
 - b. Verify tunnel records in hardware by dumping the tunnel DB using the system internal service fwd-hal command adapter-function-exec args func=tunnel-show-dbs command.
 - c. Ensure overlay reachability records (for example, route records) are fine and point to the tunnel's gport.
5. **Check Log Files:** Review log files for any errors encountered during SAG MAC programming.
6. **Generate Tech-Support Tgz:** Collect the tech-support tgz file for further SAG or LVTEP debugging.