September 2025

Extreme SLX-OS 20.7.1b

Release Notes

Supporting ExtremeRouting and ExtremeSwitching SLX 9740, SLX 9640, SLX 9540, SLX 9250, SLX 9150, Extreme 8720, Extreme 8520, and Extreme 8820

Part Number 9039215-02 Rev AA

Contents

Document History

Version	Summary of changes	Publication date
AA	Initial version for 20.7.1b	September 2025

Preface

Getting Help

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, call (800) 998 2408 (toll-free in U.S. and Canada) or 1 (408) 579 2826. 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 Service Notifications

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

1. Go to www.extremenetworks.com/support/service-notification-form.
2. Complete the form. All fields are required.
3. Select the products for which you want to receive notifications.
   Note: You can change your product selections or unsubscribe at any time.
4. Select **Submit**.

Extreme Resources

Visit the Extreme website to locate related documentation for your product and additional Extreme resources.

White papers, data sheets, and the most recent versions of Extreme software and hardware manuals are available at https://www.extremenetworks.com/. Product documentation for all supported releases is available to registered users at https://www.extremenetworks.com/support/documentation/.

Document Feedback

The Information Development team at Extreme Networks has made every effort to ensure the accuracy and completeness of this document. We are always striving to improve our documentation and help you work better, 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 in the document
- Broken links or usability issues

You can provide feedback in the following ways:

- In a web browser, select the feedback icon and complete the online feedback form.
- Access the feedback form at http://www.extremenetworks.com/documentation-feedback-pdf/.
- Email us at documentation@extremenetworks.com.

Provide the publication title, part number, and as much detail as possible, including the topic heading and page number if applicable, as well as your suggestions for improvement.

Release Overview

Release SLX-OS 20.7.1b provides the following features:

- None

Release SLX-OS 20.7.1a provides the following features:

- MCT Cluster Resiliency improvements for I2C failure trigger case

Release SLX-OS 20.7.1 provides the following features:

- MPLS L3VPN Inter-AS option A support
- QoS support using MPLS EXP bits for MPLS L3VPN
- Maximum SAG IP address per VE scale increased
- SNMP trap for Link toggle events for a PO interface enhancement
- Remove private 32-bit ASN using `remove-private-as` command

Behavior Changes

The following are the behavioral changes for SLX-OS 20.7.1b

- No behavioral changes were introduced in this release

The following are the behavioral changes for SLX-OS 20.7.1a

- No behavioral changes were introduced in this release

The following are the behavioral changes for SLX-OS 20.7.1

- Due to resource constraints on SLX 9740 and Extreme 8820, FEC counter retrieval for 100G ports on these platforms have been disabled.
- Zero Touch Provisioning (ZTP) is not supported on in-band ports on SLX 9540, SLX 9640, SLX 9740, and Extreme 8820, since the packet trap entries were causing a CPU congestion (from SLX-OS 20.6.1 release onwards).

Software Features

The following features are added in SLX-OS 20.7.1b:

- None

The following features are added in SLX-OS 20.7.1a:

Feature Name	Supported in Platforms	Description
Improving MCT Cluster Resiliency for I2C failures	Extreme 8720, SLX 9250, SLX 9740 and Extreme 8820	Ensure MCT Cluster, under I2C failure trigger condition on a MCT primary switch, is robust and allows traffic switchover to the peer in a reasonable convergence time.

The following features are added in SLX-OS 20.7.1:

Feature Name	Supported in Platforms	Description
MPLS L3VPN Inter-AS option A support	SLX 9740 and Extreme 8820	Qualified MPLS L3VPN Inter-AS option A for a Border Leaf device. This is validated with a scale of 64 VRFs
QoS using MPLS EXP for L3VPN	SLX 9740 and Extreme 8820	QoS support using MPLS EXP bits provided for L3VPN solution
Enhancement to SNMP Trap for Link toggle events	All Platforms	SNMP Traps are generated for LAG member ports, while avoiding duplicate Traps when there is no Admin/Oper state change
Max. SAG IP address per VE scale increased	All Platforms	Static Anycast Gateway (SAG) IP addresses per VE interface can be configured up to max scale of 4000
Remove private 32-bit ASN using remove-private-as command	All Platforms	BGP 32-bit private ASNs can now be removed using the CLI command

CLI Commands

The following commands were added, modified, or deprecated for the 20.7.1b release:

New commands for 20.7.1b

- No new commands were added in this release

Modified commands for 20.7.1b

- No commands were modified in this release

Deprecated commands for 20.7.1b

- No commands were deprecated in this release

The following commands were added, modified, or deprecated for the 20.7.1a release:

New commands for 20.7.1a

- No new commands were added in this release

Modified commands for 20.7.1a

- No commands were modified in this release

Deprecated commands for 20.7.1a

- No commands were deprecated in this release

The following commands were added, modified, or deprecated for the 20.7.1 release:

New commands for 20.7.1

- No commands were added

Modified commands for 20.7.1

- `dscp (QOS Mode)`
- `exp`
- `qos-mpls map dscp-exp`
- `qos-mpls map exp-dscp`
- `qos-mpls map exp-traffic-class`
- `qos-mpls map-apply dscp-exp`
- `qos-mpls map-apply exp-dscp`
- `qos-mpls map-apply exp-traffic-class`
- `neighbor (BGP Router Mode)`
- `show ip bgp`
- `show ipv6 bgp`
- `show qos-mpls maps dscp-exp`
- `show qos-mpls maps exp-dscp`
- `show qos-mpls maps exp-traffic-class`

Deprecated commands for 20.7.1

- `qos-mpls map traffic-class-exp`
- `qos-mpls map-apply traffic-class-exp`
- `show qos-mpls maps traffic-class-exp`

Changes in Future Releases

- TLS version 1.1 is deprecated from SLX-OS 20.7.1 and will be removed completely in SLX-OS 20.7.2. Users of TLS version 1.1 should start using TLS version 1.2 or higher.

Hardware Support

Supported devices and software licenses

Supported devices	Description
SLX9740-40C	Extreme SLX 9740-40C Router. Base unit with 40x100GE/40GE capable QSFP28 ports, 2 unpopulated power supply slots, 6 unpopulated fan slots
SLX9740-40C-AC-F	Extreme SLX 9740-40C-AC-F Router. Base unit with 40x100GE/40GE capable QSFP28 ports, 2 AC power supplies, 6 fan modules
SLX9740-80C	Extreme SLX 9740-80C Router. Base unit with 80x100GE/40GE capable QSFP28 ports, 4 unpopulated power supply slots, 4 unpopulated fan slots
SLX9740-80C-AC-F	Extreme SLX 9740-80C-AC-F Router. Base unit with 80x100GE/40GE capable QSFP28 ports, 4AC power supplies, 4 fan modules
SLX9740-ADV-LIC-P	Advanced Feature License for MPLS, BGP-EVPN and Integrated Application Hosting for Extreme SLX 9740
SLX9150-48Y-8C	Extreme SLX 9150-48Y Switch with two empty power supply slots, six empty fan slots. Supports 48x25GE/10GE/1GE + 8x100GE/40GE.
SLX9150-48Y-8C-AC-F	Extreme SLX 9150-48Y Switch AC with Front to Back Airflow. Supports 48x25GE/10GE/1GE + 8x100GE/40GE with dual power supplies, six fans.

Supported devices	Description
SLX9150-48Y-8C-AC-R	Extreme SLX 9150-48Y Switch AC with Back to Front Airflow. Supports 48x25GE/10GE/1GE + 8x100GE/40GE with dual power supplies, six fans.
SLX9150-48XT-6C	Extreme SLX 9150-48XT 10GBaseT Switch with two empty power supply slots, six empty fan slots, Supports 48x10GE/1GE + 6x100GE/40GE.
SLX9150-48XT-6C-AC-F	Extreme SLX 9150-48XT 10GBaseT Switch AC with Front to Back Airflow, Supports 48x10GE/1GE + 6x100GE/40GE with dual power supplies, six fans.
SLX9150-48XT-6C-AC-R	Extreme SLX 9150-48XT 10GBaseT Switch AC with Back to Front Airflow, Supports 48x10GE/1GE + 6x100GE/40GE with dual power supplies, six fans.
SLX9150-ADV-LIC-P	SLX 9150 Advanced Feature License for GuestVM, Analytics Path, PTP, BGP-EVPN.
SLX9250-32C	SLX 9250-32C Switch with two empty power supply slots, six empty fan slots. Supports 32x100/40GE.
SLX9250-32C-AC-F	SLX 9250-32C Switch AC with Front to Back Airflow. Supports 32x100GE/40GE with dual power supplies, six fans.
SLX9250-32C-AC-R	SLX 9250-32C Switch AC with Back to Front Airflow. Supports 32x100GE/40GE with dual power supplies, six fans.
SLX9250-ADV-LIC-P	SLX 9250 Advanced Feature License for GuestVM, Analytics Path, BGP-EVPN.
BR-SLX-9540-48S-AC-R	SLX 9540-48S Switch AC with Back to Front airflow (Non-port Side to port side airflow). Supports 48x10GE/1GE + 6x100GE/40GE. (1+1) redundant power supplies and (4+1) redundant fans included.
BR-SLX-9540-48S-AC-F	SLX 9540-48S Switch AC with Front to Back airflow (Port-side to non-port side airflow). Supports 48x10GE/1GE + 6x100GE/40GE. (1+1) redundant power supplies and (4+1) redundant fans included.
BR-SLX-9540-24S-DC-R	SLX 9540-24S Switch DC with Back to Front airflow (Non-port Side to port side airflow). Supports 24x10GE/1GE + 24x1GE ports.
BR-SLX-9540-24S-DC-F	SLX 9540-24S Switch DC with Front to Back airflow (Port-side to non-port side airflow). Supports 24x10GE/1GE + 24x1GE ports.
BR-SLX-9540-24S-AC-R	SLX 9540-24S Switch AC with Back to Front airflow (Non-port Side to port side airflow). Supports 24x10GE/1GE + 24x1GE ports.
BR-SLX-9540-24S-AC-F	SLX 9540-24S Switch AC with Front to Back airflow (Port-side to non-port side airflow). Supports 24x10GE/1GE + 24x1GE ports.
BR-SLX-9540-48S-DC-R	SLX 9540-48S Switch DC with Back to Front airflow (Non-port Side to port side airflow). Supports 48x10GE/1GE + 6x100GE/40GE. (1+1) redundant power supplies and (4+1) redundant fans included.
BR-SLX-9540-48S-DC-F	SLX 9540-48S Switch DC with Front to Back airflow (Port-side to non-port side airflow). Supports 48x10GE/1GE + 6x100GE/40GE. (1+1) redundant power supplies and (4+1) redundant fans included.
BR-SLX-9540-24S-COD-P	Upgrade 24x1GE to 24x10GE/1GE for SLX 9540
BR-SLX-9540-ADV-LIC-P	Advanced Feature License for SLX 9540
EN-SLX-9640-24S	Extreme SLX 9640-24S Router. Supports 24x10GE/1GE + 4x100GE/40GE. (24S+4C sku no Power supplies or Fans)
EN-SLX-9640-24S-12C	Extreme SLX 9640-24S Router. Supports 24x10GE/1GE + 12x100GE/40GE. (All ports 24S+12C sku with no Power supplies or Fans)

Supported devices	Description
EN-SLX-9640-24S-AC-F	Extreme SLX 9640-24S Router AC with Front to Back airflow. Supports 24x10GE/1GE + 4x100GE/40GE.(1 Power supply 6 Fans)
EN-SLX-9640-24S-12C-AC-F	Extreme SLX 9640-24S Router AC with Front to Back airflow. Supports 24x10GE/1GE + 12x100GE/40GE.(1 Power supply 6 Fans)
EN-SLX-9640-4C-POD-P	Extreme SLX 9640 Ports on Demand License for 4 ports of 100GE/40GE Uplinks
EN-SLX-9640-ADV-LIC-P	Extreme SLX 9640 Advanced Feature License
8720-32C	Extreme 8720-32C Switch with two empty power supply slots, six empty fan slots and a 4-post rack mount kit, Supports 32x100/40GE
8720-32C-AC-F	Extreme 8720-32C Switch with front to back airflow, Supports 32x100/40G with two AC power supplies, six fans and a 4-post rack mount kit
8720-32C-AC-R	Extreme 8720-32C Switch with back to front airflow, Supports 32x100/40G with dual AC power supplies, six fans and a 4-post rack mount kit
8720-32C-DC-F	Extreme 8720-32C Switch with front to back airflow, Supports 32x100/40G with dual DC power supplies, six fans and a 4-post rack mount kit
8720-32C-DC-R	Extreme 8720-32C Switch with back to front airflow, Supports 32x100/40G with dual DC power supplies, six fans and a 4-post rack mount kit
8520-48Y-8C	Extreme 8520-48Y Switch with two empty power supply slots, six empty fan slots; Ships with one 4-post rack mount kit; Supports 48x25/10/1G and 8x100/40G ports
8520-48Y-8C-AC-F	Extreme 8520-48Y Switch with front-back airflow; Ships with two AC power supplies, six fans, one 4-post rack mount kit; Supports 48x25/10/1G and 8x100/40G ports
8520-48Y-8C-AC-R	Extreme 8520-48Y Switch with back-front airflow; Ships with two AC power supplies, six fans, one 4-post rack mount kit; Supports 48x25/10/1G and 8x100/40G ports
8520-48Y-8C-DC-F	Extreme 8520-48Y Switch with front-back airflow; Ships with two DC power supplies, six fans, one 4-post rack mount kit; Supports 48x25/10/1G and 8x100/40G ports
8520-48Y-8C-DC-R	Extreme 8520-48Y Switch with back-front airflow; Ships with two DC power supplies, six fans, one 4-post rack mount kit; Supports 48x25/10/1G and 8x100/40G ports
8520-48XT-6C	Extreme 8520-48XT Switch with two empty power supply slots, six empty fan slots; Ships with one 4-post rack mount kit; Supports 48x10/1G copper ports and 6x100/40G fiber ports
8520-48XT-6C-AC-F	Extreme 8520-48XT Switch with front-back airflow; Ships with two AC power supplies, six fans, one 4-post rack mount kit; Supports 48x10/1G copper ports and 6x100/40G fiber ports
8520-48XT-6C-AC-R	Extreme 8520-48XT Switch with back-front airflow; Ships with two AC power supplies, six fans, one 4-post rack mount kit; Supports 48x10/1G copper ports and 6x100/40G fiber ports
8520-48XT-6C-DC-F	Extreme 8520-48XT Switch with front-back airflow; Ships with two DC power supplies, six fans, one 4-post rack mount kit; Supports 48x10/1G copper ports and 6x100/40G fiber ports

Supported devices	Description
8520-48XT-6C-DC-R	Extreme 8520-48XT Switch with back-front airflow; Ships with two DC power supplies, six fans, one 4-post rack mount kit; Supports 48x10/1G copper ports and 6x100/40G fiber ports
8000-PRMR-LIC-P	Extreme 8000 Premier Feature License (includes Integrated Application Hosting)
8820-40C	Extreme 8820-40C base unit with 40x100GE/40GE QSFP28 ports with 2 unpopulated power supply slots, 6 unpopulated fan slots and a 4-post rack mount kit
8820-40C-AC-F	Extreme 8820-40C with Front-Back airflow. Base unit with 40x100GE/40GE QSFP28 ports with 2 AC power supplies, 6 fan modules and a 4-post rack mount kit
8820-40C-AC-R	Extreme 8820-40C with Back-Front airflow. Base unit with 40x100GE/40GE QSFP28 ports with 2 AC power supplies, 6 fan modules and a 4-post rack mount kit
8820-40C-DC-F	Extreme 8820-40C with Front-Back airflow. Base unit with 40x100GE/40GE QSFP28 ports with 2 DC power supplies, 6 fan modules and a 4-post rack mount kit
8820-40C-DC-R	Extreme 8820-40C with Back-Front airflow. Base unit with 40x100GE/40GE QSFP28 ports with 2 DC power supplies, 6 fan modules and a 4-post rack mount kit
8820-80C	Extreme 8820-80C. Base unit with 80x100GE/40GE QSFP28 ports with 4 unpopulated power supply slots, 4 unpopulated fan slots and a 4-post rack mount kit
8820-80C-AC-F	Extreme 8820-80C with Front-Back airflow. Base unit with 80x100GE/40GE QSFP28 ports with 4 AC power supplies, 4 fan modules and a 4-post rack mount kit
8820-80C-AC-R	Extreme 8820-80C with Back-Front airflow. Base unit with 80x100GE/40GE QSFP28 ports with 4 AC power supplies, 4 fan modules and a 4-post rack mount kit
8820-80C-DC-F	Extreme 8820-80C with Front-Back airflow. Base unit with 80x100GE/40GE QSFP28 ports with 4 DC power supplies, 4 fan modules and a 4-post rack mount kit
8820-80C-DC-R	Extreme 8820-80C with Back-Front airflow. Base unit with 80x100GE/40GE QSFP28 ports with 4 DC power supplies, 4 fan modules and a 4-post rack mount kit

Supported power supplies, fans, and rack mount kits

XN-ACPWR-1600W-F	SLX 9740 Fixed AC 1600W Power Supply Front to Back. Power cords not included Extreme 8820 Fixed AC 1600W Power Supply Front to Back. Power cords not included
XN-ACPWR-1600W-R	SLX 9740 Fixed AC 1600W Power Supply Back to Front. Power cords not included. Extreme 8820 Fixed AC 1600W Power Supply Back to Front. Power cords not included
XN-DCPWR-1600W-F	SLX 9740 Fixed DC 1600W Power Supply Front to Back. Power cords not included Extreme 8820 Fixed DC 1600W Power Supply Front to Back. Power cords not included
XN-DCPWR-1600W-R	Extreme 8820 Fixed DC 1600W Power Supply Back to Front. Power cords not included.
XN-FAN-003-F	SLX 9740 FAN Front to Back airflow for SLX9740-40C Extreme 8820 FAN Front to Back airflow for 8820-40C
XN-FAN-003-R	SLX 9740 FAN Back to Front airflow for SLX9740-40C Extreme 8820 FAN Back to Front airflow for 8820-40C
XN-FAN-004-F	SLX 9740 FAN Front to Back airflow for SLX9740-80C Extreme 8820 FAN Front to Back airflow for 8820-80C
XN-FAN-004-R	SLX 9740 FAN Back to Front airflow for SLX9740-80C Extreme 8820 FAN Back to Front airflow for 8820-80C
XN-4P-RKMT299	2-Post Rail Kit for SLX 9740-40C
XN-2P-RKMT300	2-Post Rail Kit for SLX 9740-80C
XN-4P-RKMT301	4-Post Rail Kit for SLX 9740-80C
XN-4P-RKMT302	4-Post Rail Kit for SLX 9740-40C
XN-ACPWR-750W-F	AC 750W PSU, Front to Back Airflow supported on VSP 7400, SLX 9150, SLX 9250, X695, Extreme 8720, Extreme 8520
XN-ACPWR-750W-R	AC 750W PSU, Back to Front Airflow supported on VSP 7400, SLX 9150, SLX 9250, X695, Extreme 8720, Extreme 8520
XN-DCPWR-750W-F	DC 750W PSU, Front to Back Airflow supported on VSP 7400, SLX 9150, SLX 9250, X695, Extreme 8720, Extreme 8520
XN-DCPWR-750W-R	DC 750W PSU, Back to Front Airflow supported on VSP 7400, SLX 9150, SLX 9250, X695, Extreme 8720, Extreme 8520
XN-FAN-001-F	Front to back Fan for use in VSP 7400, SLX 9150, SLX 9250, X695, Extreme 8720, Extreme 8520
XN-FAN-001-R	Back to Front Fan for use in VSP 7400, SLX 9150, SLX 9250, X695, Extreme 8720, Extreme 8520
XN-4P-RKMT298	Four post rack mount rail kit supported on VSP 7400, SLX 9150, SLX 9250, X695, Extreme 8720, Extreme 8520
XN-2P-RKMT299	Two post rack mount rail kit supported on VSP 7400, SLX 9150, SLX 9250, X695, Extreme 8720, Extreme 8520, Extreme 8820
XN-2P-RKMT300	2-Post Rail Kit for Extreme 8820-80C
XN-4P-RKMT301	4-Post Rail Kit for Extreme 8820-80C

XN-4P-RKMT302	4-Post Rail Kit for Extreme 8820-40C

Supported Optics and Cables

For a complete list of all supported optics, see **Extreme Optics** at https://optics.extremenetworks.com/.

Supported FEC modes

SLX 9250 and Extreme 8720

Port Type	Media Type	Default FEC Mode	Supported FEC Modes
100G	Passive DAC	RS-FEC	RS-FEC Disabled
100G	SR4	RS-FEC	RS-FEC Disabled
100G	LR4/PSM4	Disabled	RS-FEC Disabled
25G	Breakout DAC SR	Auto-Neg	RS-FEC FC-FEC Auto-Neg Disabled
25G	Breakout SR4	FC-FEC	RS-FEC FC-FEC Disabled
25G	Breakout LR	RS-FEC	RS-FEC FC-FEC Disabled

SLX 9740 and Extreme 8820

Port Type	Media Type	Default FEC Mode	Supported FEC Modes
100G	Passive DAC	RS-FEC	RS-FEC Disabled
100G	SR4	RS-FEC	RS-FEC Disabled
100G	LR4/PSM4	Disabled	RS-FEC Disabled
25G	Breakout DAC SR	FC-FEC	FC-FEC RS-FEC Disabled
25G	Breakout SR4	FC-FEC	FC-FEC RS-FEC Disabled
25G	Breakout LR	RS-FEC	RS-FEC FC-FEC Disabled

SLX 9150 and Extreme 8520

Port Type	Media Type	Default FEC Mode	Supported FEC Modes
100G	Passive DAC	RS-FEC	RS-FEC Disabled
100G	SR4	RS-FEC	RS-FEC Disabled
100G	LR4/PSM4	Disabled	RS-FEC Disabled
25G(Native)	DAC	Auto-Neg	RS-FEC FC-FEC Auto-Neg Disabled
25G(Native)	SFP	FC-FEC	RS-FEC FC-FEC Disabled
25G(Native)	LR	RS-FEC	RS-FEC FC-FEC Disabled

SLX 9540 and SLX 9640

Port Type	Media Type	Default FEC Mode	Supported FEC Modes
100G	Passive DAC	RS-FEC	RS-FEC Disabled
100G	SR4	RS-FEC	RS-FEC Disabled
100G	LR4/PSM4	Disabled	RS-FEC Disabled
25G	Breakout LR	RS-FEC	RS-FEC FC-FEC Disabled

Software Download and Upgrade

For more information about the various methods of upgrading to SLX-OS 20.7.1b see the *Extreme SLX-OS Software Upgrade Guide.*

Image files

Download the following images from www.extremenetworks.com.

Image file name	**Description**
SLX-OS_20.7.1b.tar.gz	SLX-OS 20.7.1b software
SLX-OS_20.7.1b_mibs.tar.gz	SLX-OS 20.7.1b MIBS
SLX-OS_20.7.1b.md5	SLX-OS 20.7.1b md5 checksum
SLX-OS_20.7.1b-digests.tar.gz	SLX-OS 20.7.1b sha checksum
SLX-OS_20.7.1b-releasenotes.pdf	Release Notes

Baseboard Management Controller (BMC) firmware upgrade

- With SLX-OS 20.6.1 onwards, BMC firmware update will be performed along with SLX-OS update on BMC supported platforms. This upgrade will happen only if the installed BMC firmware version is older than the version bundled along with the SLX-OS firmware. Supported SLX platforms are Extreme 8520, Extreme 8720, Extreme 8820 and SLX 9740.
- No new SLX-OS CLI was introduced for BMC firmware upgrade, as this being an implicit BMC firmware update.
- With this new feature, BMC firmware image is bundled as part of SLX-OS image. When the user updates the OS, and, if BMC firmware version on the device is found to be older than the BMC image bundled with SLX-OS image, the BMC image bundled with SLX shall be updated on BMC along with SLX-OS update.
- By design, only BMC firmware upgrade is supported – downgrade is not supported.
- BMC firmware upgrade will occur with all supported SLX-OS upgrade methods – incremental, full install and net install
- In case the BMC upgrade fails, “firmware download” of SLX-OS will continue without any disruption.
- During BMC upgrade, IPMI/BMC connectivity will be impacted. Hence intermittent RASLOGS (e.g. FW-1404 and EM-1050, HIL-1404 etc) from environmental monitoring daemon may be observed. These intermittent RASLOG messages will disappear only after the device is reloaded.
- Existing BMC configuration will be preserved even after the BMC is updated.
- Limitations -
    - There is a small increase in SLX-OS installation time (around 4 to 7 minutes), if BMC firmware is also upgraded.
    - Intermittent RASLOGS or FFDC messages are generated due to interruption at BMC/IPMI channel.

Extreme 8820

To From	20.4.3/a/b	20.5.1/a	20.5.2a	20.6.1/a/b	20.6.2/a	20.6.3a/b	20.7.1/a/b
20.4.3 (Factory Image)	For upgrade: normal firmware download / coldboot						
20.5.1/a							
20.5.2a							
20.6.1/a/b							
20.6.2/a							
20.6.3a/b							
20.7.1/a/b							

Extreme 8720

To From	20.4.3/a/b	20.5.1/a	20.5.2a	20.6.1/a/b	20.6.2/a	20.6.3a/b	20.7.1/a/b
20.4.3/a/b	For upgrade and downgrade: normal firmware download / coldboot						
20.5.1/a							
20.5.2a							
20.6.1/a/b							
20.6.2/a							
20.6.3a/b							
20.7.1/a/b							

Extreme 8520

To From	20.4.3/a/b	20.5.1/a	20.5.2a	20.6.1/a/b	20.6.2/a	20.6.3a/b	20.7.1/a/b
20.4.3/a/b	For upgrade and downgrade: normal firmware download / coldboot						
20.5.1/a							
20.5.2a							
20.6.1/a/b							
20.6.2/a							
20.6.3a/b							
20.7.1/a/b							

SLX 9740

To From	20.4.3/a/b	20.5.1/a	20.5.2a	20.6.1/a/b	20.6.2/a	20.6.3a/b	20.7.1/a/b
20.4.3/a/b	For upgrade and downgrade: normal firmware download / coldboot						
20.5.1/a							

To From	20.4.3/a/b	20.5.1/a	20.5.2a	20.6.1/a/b	20.6.2/a	20.6.3a/b	20.7.1/a/b
20.5.2a							
20.6.1/a/b							
20.6.2/a							
20.6.3a/b							
20.7.1/a/b							

SLX 9540 and SLX 9640

To From	20.4.3/a/b	20.5.1/a	20.5.2a	20.6.1/a/b	20.6.2/a	20.6.3a/b	20.7.1/a/b
20.4.3/a/b	For upgrade and downgrade: normal firmware download / coldboot						
20.5.1/a							
20.5.2a							
20.6.1/a/b							
20.6.2/a							
20.6.3a/b							
20.7.1/a/b							

SLX 9150 and SLX 9250

To From	20.4.3/a/b	20.5.1/a	20.5.2a	20.6.1/a/b	20.6.2/a	20.6.3a/b	20.7.1/a/b
20.4.3/a/b	For upgrade and downgrade: normal firmware download / coldboot						
20.5.1/a							
20.5.2a							
20.6.1/a/b							
20.6.2/a							
20.6.3a/b							
20.7.1/a/b							

Upgrade and Downgrade considerations for Threshold Monitor configuration:

Downgrade Considerations:

1. If the configured value for CPU "limit" exceeds valid range in older release [0-80] then downgrade will be blocked with error. User can reconfigure CPU “limit” in the range [0-80] and downgrade.

2. If the configured value for Memory "high-limit" exceeds valid range in older release [0-80] or if it is less than the default value of "limit" in older release [60], then downgrade will be blocked with error. User can reconfigure Memory “high-limit” in the range [60-80] and downgrade.

3. If the startup file has “actions” configured as “snmp” or “all”, then config replay process triggered in firmware full-install downgrade, will lead all the corresponding threshold-monitor CLI parameters, such as poll, retry, to reset to respective default values.

Upgrade Considerations:

1. If the startup file has “Memory limit and /or low-limit” configured, then config replay process triggered in firmware full-install downgrade, will lead all the corresponding threshold-monitor CLI parameters, such as poll, retry, to reset to respective default values.

SLX TPVM Support Information

SLX-OS 20.7.1x supports TPVM 4.6.1 and later and TPVM 4.7.0 and later, on all platforms.

Upgrading the TPVM without configuration persistence (Legacy upgrade method)

Upgrading TPVM from 4.5.x to 4.6.x or 4.7.x

Consider the following when upgrading TPVM for 4.6.x or 4.7.x

- To perform full upgrade from TPVM 4.5.x to latest 4.6.x, do the following:
  - Upgrade to SLX-OS 20.6.x while the existing TPVM 4.5.x installation continues to run
  - Copy the new *tpvm-4.6.X-X.amd64.deb* to */tftpboot/SWBD2900* directory on the SLX device.
  - Install latest TPVM 4.6.x using **tpvm upgrade** command
- To perform full upgrade from TPVM 4.5.x or 4.6.x to latest 4.7.x, do the following:
  - Upgrade to SLX-OS 20.7.x while the existing TPVM 4.5.x or 4.6.x installation continues to run
  - Copy the new *tpvm-4.7.X-X.amd64.deb* to */tftpboot/SWBD2900* directory on the SLX device.
  - Install latest TPVM 4.7.x using **tpvm upgrade** command

Notes:

- Security updates are added to the TPVM image and to the separate Debian file used for incremental TPVM update. Main TPVM image size is ~2.0 GB and the TPVM incremental update Debian file size is ~0.5 GB. You must have at least 1GB of free space on the switch before proceeding with the `tpvm upgrade incremental` command.
- Ubuntu Linux distribution on TPVM is upgraded to 20.04 LTS from TPVM version 4.6.0 onwards. As Ubuntu Linux distribution on TPVM is upgraded to 20.04 LTS incremental upgrade is not supported, upgrading TPVM from 4.5.x to 4.6.x needs a full upgrade. Please refer to the respective TPVM 4.6.x Release notes for more information.

- Ubuntu Linux distribution on TPVM is upgraded to 22.04 LTS from TPVM version 4.7.0 onwards. As Ubuntu Linux distribution on TPVM is upgraded to 22.04 LTS incremental upgrade is not supported, upgrading TPVM from 4.5.x or 4.6.x to 4.7.x needs a full upgrade. Please refer to the respective TPVM 4.7.x Release notes for more information.
- The latest version in the TPVM 4.6.x branch, TPVM 4.6.25, has security updates till July 1, 2025. Main TPVM image size is ~2.1 GB and the TPVM incremental update Debian file size is ~0.8 GB.
- The latest version in the TPVM 4.7.x branch, TPVM 4.7.8, has security updates till August 4, 2025. Main TPVM image size is ~2.2 GB and the TPVM incremental update Debian file size is ~0.7 GB.
- Updates within the same series of TPVM releases, for example, between a version of 4.6.x and another version of 4.6.x, incremental upgrades are supported. Use the **tpvm upgrade incremental** command to do the upgrade.

Limitations and Restrictions

Copy flash to startup and reload with TPVM

setNTPServer and setLDAPServer statuses are reported as failed in the output of the `show tpvm status-history`. After reload, TPVM is expected to be running when the above configurations are re-applied. When the TPVM is not running and the NTP and LDAP configurations are applied, these errors are seen. This is a limitation as reapplying NTP and LDAP configurations are not supported.

You need to have minimum 1GB free space on TPVM when you try to perform the security patch upgrade using the command `tpvm upgrade incremental …`

TPVM upgrade incremental command and file support is available only from 4.5 if we try to perform the incremental upgrade from 4.4.0 to latest, the upgrade fails and ask to perform the tpvm upgrade.

TPVM upgrade incremental command will not be supported when you try TPVM deploy in config mode and TPVM upgrade incremental command will not support with snapshot option.

Do not use the **tpvm upgrade incremental** command to upgrade the patches with *tpvm-4.X.X-X.amd64.deb*. Use the *tpvm_inc_upg-4.X.X-X.amd64.deb* image file to perform incremental upgrades.

Similarly, do not use the *tpvm_inc_upg-4.X.X-X.amd64.deb* image file to perform full upgrade. Do not use this file to perform **tpvm deploy** in *config mode* and *option*.

Port macro restrictions on breakout port configuration on SLX 9740 and Extreme 8820

A port macro (PM) is a port group. Each PM has 4 ports, which are contiguous. PM0 has ports 0/1-0/4, PM1 has ports 0/5-0/8, PM2 has ports 0/9-0/12, and so on.

Only the odd ports can be split to 4x10G or 4x25G using the breakout cables: 0/1, 0/3, 0/9, 0/11, 0/13, 0/15, 0/17, 0/19, 0/21, 0/23, 0/25, 0/27, 0/29, 0/31, 0/33, 0/35, 0/37, 0/39, 0/41, 0/43, 0/49, 0/51, 0/53, 0/55, 0/57, 0/59, 0/61, 0/63, 0/65, 0/67, 0/69, 0/71, 0/73, 0/75, 0/77, and 0/79.

Breaking out these ports using the breakout cables results in 72 interfaces for the SLX 9740-40/Extreme 8820-40C and 144 interfaces for the SLX 9740-80C/Extreme 8820-80C.

- Ports 5-8 and 45-48 cannot be broken up and are supported only in 100G.
- For any PM, 40G and 10G ports cannot coexist with 25G ports. The following configurations are not supported:

PM Configuration	Examples
If any port is configured as 40G or 4x10G breakout, no 4x25G breakout is allowed unless the 40G ports will be removed as part of the breakout operation.	• If 0/3 or 0/4 is 40G, you cannot configure 0/1 as 4x25G breakout. • If 0/1 is 4x10G breakout, you cannot configure 0/3 as 4x25G breakout. • If 0/3 is 4x10G breakout, you cannot configure 0/1 as 4x25G breakout. • If 0/1 or 0/2 is 40G, you can configure 0/1 as 4x25G breakout because 0/1 and 0/2 will be removed. • If 0/3 or 0/4 is 40G, you can configure 0/3 as 4x25G breakout because 0/3 and 0/4 will be removed.
If 4x25G breakout is configured, no 40G or 4x10G.	• If 0/1 is configured as 4x25G breakout, you cannot configure 0/3 or 0/4 as 40G. • If 0/1 is configured as 4x25G breakout, you cannot configure 0/3 as 4x10G breakout. • If 0/3 is configured as 4x25G breakout, you cannot configure 0/1 or 0/2 as 40G. • If 0/3 is configured as 4x25G breakout, you cannot configure 0/1 as 4x10G breakout.

QoS

- PCP remarking is not supported for SLX 9740 and Extreme 8820.
- Egress rate limiting in a Bridge Domain configuration is not supported for SLX 9740 and Extreme 8820.
- DSCP-COS map is not supported for SLX 9740 and Extreme 8820.
- On SLX 9640 platform, L3 QoS is not supported for VxLAN L3 gateway.
- On SLX 9540 and SLX 9640, if Trust-DSCP feature is enabled, then non-IP packets will take only the default traffic class value. *For more details, refer the QoS section of SLX-OS 20.6.2 Traffic Management guide.*
- QoS support using MPLS EXP is supported only in SLX 9740 and Extreme 8820 (for L3VPN Uniform mode). DSCP-EXP, EXP-TrafficClass and EXP-DSCP maps are supported.
- DSCP Mutation and EXP-DSCP are mutually exclusive.

Others

- sflow sampling does not work for VLL when BUM rate limiting is applied on interface in SLX 9740 and Extreme 8820.
- sflow sample traffic to CPU is rate limited. You can use the **qos cpu slot** command to change the rate.

- When Resilient Hashing CLI is enabled or disabled, or the *max-path* value is changed, it may cause **BFD sessions** in **related VRFs** to go down. However, **BFD sessions in unrelated VRFs will not be affected.**
- Resilient Hashing feature is supported only on SLX 9150, SLX 9250, SLX 9740, Extreme 8720, Extreme 8520 and Extreme 8820. Other platforms are not supported.
- Resilient Hashing supports 32K flowset entries for Extreme 8720 and Extreme 8520.

Open Config Telemetry Support

- User authentication not supported.
- gNMI calls through inband interfaces not supported.
- Usage of wild cards is not supported.
- gNMI SET is not supported.
- gNMI ON CHANGE subscription is not supported.

SNMP

- Not all counters related to UDP, and TCP MIBs are supported.
- Configuring an in-band port into a Management VRF requires SNMP agent reload.

Maximum Logical Interfaces or LIFs scale

Maximum Logical Interface (LIF) (Port-VLAN/Port-Bridge Domain (BD)) associations supported on SLX 9150, SLX 9250, Extreme 8520, Extreme 8720 is 14200. Since VLAN and BD resources share the same hardware table memory space, the max scale of one has a trade-off with the scale of the other. That is, for example, the maximum Port-BD associations cannot be scaled to 14200 when the combined scale of VLAN and BDs exceeds 8096.

IPv6 Manageability support on TPVM

- The TPVM management interface can be configured with a single IPv6 address. You can configure an IPv4 address in addition to the IPv6 address. Configuring IPv4 address is optional.
- tpvm stop and tpvm start commands must be issued to configure the TPVM management interface's IPv4 and IPv6 address.

Removal of DF towards IP Fabric (Local Bias support for LVTEP)

- Single-homed LVTEP client (spine uplink DOWN in one of the MCT nodes) is not supported.
- Need to have backup routing over ICL to reach the spines in case of uplink failure.

ICMP and ICMPv6 redirect

Enable/disable ICMP and ICMPv6 redirect are only available on SLX 9540 and SLX 9640. On these platforms, these are only supported on physical ports.

Transporting IPv6 traffic over GRE IPv4 Tunnel

- If GRE feature is enabled, IPv6 ACL filters to drop OSPFv3 packets will not work for SLX 9740 and Extreme 8820 platforms.
- Multicast traffic is not supported over IPv6 GRE overlay. Multicast packets will be dropped.
- IPv6 ACL is not supported on GRE tunnel.
- IPv4 and IPv6 control packets over the GRE Tunnel are not accounted for in the GRE tunnel statistics.

- DSCP value from the inner IPv6 packet is not copied to outer GRE header on SLX 9540 and SLX 9640 platforms.

Flow Based Mirroring

(Applicable to SLX 9150, SLX 9250, Extreme 8720 and Extreme 8520 platforms)

- Flow based ingress mirroring does not support port-channel port as a mirroring source port.
- Flow based ingress mirroring supports VLAN as a mirroring source port, but VLAN range is not supported.

MPLS over GRE

(Applicable to SLX 9150, SLX 9250, Extreme 8720 and Extreme 8520 platforms)

Transit MPLSoGRE and dual-tag BD LIF are mutually exclusive on the same interface (Ethernet or Port-channel) - both features cannot co-exist on the same interface.

- MPLSoGRE traffic will be impacted on an interface where dual-tagged BD LIF is configured.
- Other interfaces, without a dual-tagged BD LIF, are not impacted.

Unsupported characters in SLX-OS and TPVM passwords

The following characters are not supported in the SLX-OS and TPVM passwords.

- & (ampersand)
- \ (backslash)
- ‘ (single quote)

Open Defects

Open defects in SLX-OS 20.7.1b

The following software defect is open in SLX-OS 20.7.1b as of September 2025.

Parent Defect ID:	SLXOS-78542	**Issue ID:**	SLXOS-78542
Severity:	S2 - Major		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.7.1b
Symptom:	Slow increase in memory usage of DCMD		
Condition:	When show logging raslog cmd executed periodically		
Workaround:	Avoid frequent execution of high-output commands. The show logging raslog command can generate large outputs, which may not release memory efficiently when executed periodically.		

Open defects in SLX-OS 20.7.1a

No software defects are open in SLX-OS 20.7.1a as of April 2025.

Open defects in SLX-OS 20.7.1

The following software defects are open in SLX-OS 20.7.1 as of April 2025:

Parent Defect ID:	SLXOS-76092	**Issue ID:**	SLXOS-76092
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.4.3c
Symptom:	Self originated External LSA corresponds to the static route is removed from LSDB causing traffic loss.		
Condition:	Happens randomly on any of the core routers in the topology.		
Workaround:	Unconfigure and configure static route corresponding to the external LSA which is missing from LSDB.		
Recovery:	Unconfigure and configure static route corresponding to the external LSA which is missing from LSDB.		

Parent Defect ID:	SLXOS-76527	**Issue ID:**	SLXOS-76527
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3
Symptom:	BGP crashed when BGP Flowspec enabled.		
Condition:	When BGP Flowspec enabled.		
Workaround:	NA		
Recovery:	NA		

Parent Defect ID:	SLXOS-77649	**Issue ID:**	SLXOS-77649
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.7.1
Symptom:	Intermittently observed BGPd crashing on L3VPN MPLS test bed with scale of 310 VRFs while withdrawing routes on BL/PE node		
Condition:	This is due to BGP RIB-OUT module out race condition.		
Workaround:	The crash is not seen with lower scale of 64 VRFs		

Parent Defect ID:	SLXOS-77773	**Issue ID:**	SLXOS-77773
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.7.1
Symptom:	BGPd crash observed on L3VPN MPLS test bed with a scale of 310 VRFs While doing vpnv4/vpnv6 neighbor deactivate/activate and removing/adding 'bgp next-hop loopback <>' config on all 310 VRFs under AF4/AF6.		
Condition:	This is due to BGP RIB-OUT module out race condition.		
Workaround:	This crash is not seen if you only run activate/deactivate of all 310 VRF's without adding bgp next hop as loop back		

Parent Defect ID:	SLXOS-77952	**Issue ID:**	SLXOS-77952
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.7.1
Symptom:	Unexpected Error message is observed in console during the firmware upgrade.		
Condition:	When the image is upgraded from 20.6.3x to 20.7.1. An error message is displayed in the console *init* logs.		
Workaround:	None		
Recovery	This error message is cosmetic and does not have any impact during installation. The installation proceeds normally and a fully functional system is achieved at the end of the installation process. No intervention is required.		

Defects Closed with Code Changes

Defects closed with code changes in SLX-OS 20.7.1b

The following software defects were closed with code changes in SLX-OS 20.7.1b as of September 2025:

Parent Defect ID:	SLXOS-77904	**Issue ID:**	SLXOS-78122
Severity:	S2 - Major		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3a
Symptom:	No ARP response for few VEs, when multiple IPv4 Addresses are configured on a Ve interface.		
Condition:	Issue can be observed when multiple IPv4 Addresses are configured on a Ve interface with user configured VRF present on the Ve interface and when the system is reloaded.		
Workaround:	Issue will not be seen when either user configured VRF is not present on the VE interface and reloaded OR the IPv4 Addresses are configured manually during run time.		

Parent Defect ID:	SLXOS-78131	**Issue ID:**	SLXOS-78192
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3a
Symptom:	Traffic originating from an L2 leaf fails to reach the Border Leaf (BL), resulting in dropped packets. This impacts routable traffic dependent on the BL's VE-MAC.		
Condition:	During the maintenance (update/reboot) of the Kubernetes servers, an unexpected packet loop caused the Border Leaf VE-MAC to be incorrectly learned as a dynamic MAC address, resulting in routable traffic being forwarded to the wrong destination.		
Workaround:	Remove and Re-add the impacted VLANs from the EVPN instance, so that VE MAC will be resynced by the BGP.		

Parent Defect ID:	SLXOS-77789	**Issue ID:**	SLXOS-78198
Severity:	S2 - Major		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.1bb
Symptom:	When the underlay network path experiences a brief disruption, VxLAN tunnels may go down and subsequently come back up. In some cases, after this tunnel flap, traffic may not be forwarded through the correct tunnel. Instead, it may be misdirected to an incorrect VxLAN tunnel. This issue is more likely to occur in topologies with a large number of VxLAN tunnels.		
Condition:	In environments with a high number of VxLAN tunnels, if multiple tunnels flap simultaneously due to an underlay path failure, they may be re-established in a different order. Under certain timing conditions, this can result in VxLAN hardware resources being allocated differently. As a result, traffic may be incorrectly forwarded to the wrong tunnel endpoint.		

Workaround:	To restore proper forwarding behavior, manually clear and re-establish the BGP EVPN sessions using the following command: "clear bgp evpn neighbor all" This forces a tunnel re-creation and typically resolves the forwarding issue.

Parent Defect ID:	SLXOS-78029	**Issue ID:**	SLXOS-78278
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.7.1a
Symptom:	10G SFP optic FTLX8574D3BCL-EX was getting detected as un-qualified on these platforms - 9640, 9540, 9150-F, 8520-F		
Condition:	When device come up or these specific optics are inserted, observed message as unqualified. After validation, these optics are now added to the qualified list.		

Parent Defect ID:	SLXOS-77960	**Issue ID:**	SLXOS-78279
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.2a
Symptom:	sFlow sample packet is not showing the interface speed properly.		
Condition:	sFlow is configured on Port-channel interface in an SLX 9540 or SLX 9640 device.		

Parent Defect ID:	SLXOS-77873	**Issue ID:**	SLXOS-78280
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3ab
Symptom:	User may observe "Message Generic Error" on console as syslog-ng configuration fails and stops running.		
Condition:	while configuring secure port 6514 for syslog-ng		

Parent Defect ID:	SLXOS-77869	**Issue ID:**	SLXOS-78281
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3a
Symptom:	SLXOS software is not detecting the VPN service label update.		
Condition:	Per-Next-Hop Label Allocation: labels are assigned based on the next-hop IP address or exit point, rather than the VRF itself. Analysis: At T1, SLX receives a VPN route from a vendor device with label X1 (assuming per-next-hop label allocation at the vendor node). At T2, SLX receives the same VPN route update with label X2 due to a path or next-hop change at the vendor node. In this scenario, this issue can happen.		

Parent Defect ID:	SLXOS-77864	**Issue ID:**	SLXOS-78282
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3a
Symptom:	Device may crash when 'breakout' and 'no breakout' is executed.		
Condition:	When 'breakout' and then 'no breakout' is executed on the 100G interface while SNMP IFMIB query is issued parallelly.		
Workaround:	.		

Parent Defect ID:	SLXOS-77837	**Issue ID:**	SLXOS-78283
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3a
Symptom:	Frequent EVPN dampening warning message on the console		
Condition:	After a Port-channel flap on the EVPN Multi-homing node		

Parent Defect ID:	SLXOS-77807	**Issue ID:**	SLXOS-78284
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3ab
Symptom:	The EVPN routing table selects an incorrect best path.		
Condition:	This happens when a VRF receives the same route from multiple sources with different local preference values and imports them into the EVPN table.		
Workaround:	None		
Recovery:	None		

Parent Defect ID:	SLXOS-77787	**Issue ID:**	SLXOS-78285
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.2b
Symptom:	UDP packet with destination port 8888 is trapped to control-plane		
Condition:	When UDP packet with destination port 8888 is received through the VPLS tunnel		

Parent Defect ID:	SLXOS-77770	**Issue ID:**	SLXOS-78286
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.1bb
Symptom:	Learning ARP from a host in another subnet (non-connected)		
Condition:	GARP is received from a non-connected host		

Parent Defect ID:	SLXOS-77754	**Issue ID:**	SLXOS-78287
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.1b
Symptom:	Self originated Type5 LSA corresponds to redistribution source will be flushed when a router is becoming non-translator for an NSSA area. Also the LSAs will have checksum errors and leads to the previous		

	instance of same LSA exceeding MAX_AGE and stays on the neighboring routers forever leading to continuous SPF runs.
Condition:	Issue will be seen when an NSSA ABR will become non-translator for one of its attached NSSA areas.
Workaround:	None
Recovery:	None

Parent Defect ID:	SLXOS-78140	**Issue ID:**	SLXOS-78288
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.1bd
Symptom:	The output of "show ip route next-hop ref-routes" does not display all learned routes for a next-hop; it is limited to 500 entries.		
Condition:	Occurs when more than 500 routes are learned with the same next-hop.		
Workaround:	None		
Recovery:	None		

Parent Defect ID:	SLXOS-77373	**Issue ID:**	SLXOS-78289
Severity:	S2 - Major		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3
Symptom:	BGP, LLDP flaps and ICMP packet loss between leaf and spine node is observed.		
Condition:	When unknown multicast IGMP packets are received on a leaf nodes and more than one Vxlan tunnel is present between the leaf nodes.		

Parent Defect ID:	SLXOS-77831	**Issue ID:**	SLXOS-78290
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3
Symptom:	Able to configure the 'crypto cert expiry-level info' value irrespective of minor, major, critical values		
Condition:	Whenever configure the 'crypto cert expiry-level info' less then minor or major or critical values.		

Parent Defect ID:	SLXOS-78258	**Issue ID:**	SLXOS-78332
Severity:	S2 - Major		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3a
Symptom:	Trusted Peer command will be failing on SLX		
Condition:	Customer has installed TPVM and trying to configure trusted-peer in a multi node EFA deployment.		
Workaround:	Include the line "services: files" in /etc/nsswitch.conf		

Parent Defect ID:	SLXOS-78319	**Issue ID:**	SLXOS-78342
Severity:	S3 - Moderate		

Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3ab
Symptom:	sFlow samples are not being forwarded to the sFlow collector.		
Condition:	sFlow is configured on Port-channel member port in an Extreme 8820 device.		

Parent Defect ID:	SLXOS-78261	**Issue ID:**	SLXOS-78350
Severity:	S2 - Major		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3
Symptom:	Intermittent or partial traffic drops may occur when VE-MACs are inconsistently learned as Dynamic on an MCT node following network loops. Traffic is dropped when it reaches the MCT node where the VE-MAC is marked as Dynamic, while it is forwarded correctly on the node where the VE-MAC remains classified as EVPN-Static.		
Condition:	The issue occurs when a traffic loop is present on the connected devices of the L2 leaf switches, causing the VE-MAC to be incorrectly learned as a Dynamic MAC on the edge ports.		
Workaround:	Manually identify and clear VE MACs that are learned as Dynamic MAC addresses on the affected VLANs to trigger a re-sync from BGP.		

Parent Defect ID:	SLXOS-78220	**Issue ID:**	SLXOS-78374
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3a
Symptom:	One of the multi-path next-hop for NSSA route is showing in reverse order.		
Condition:	Router receives equal cost NSSA LSAs for a prefix with forwarding address set.		
Workaround:	None		
Recovery:	None		

Parent Defect ID:	SLXOS-78197	**Issue ID:**	SLXOS-78476
Severity:	S2 - Major		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.7.1a
Symptom:	Unable to login to the device. "SLX-OS is not ready. Please login later" message is seen while trying to login.		
Condition:	When device was upgraded from 20.6.3ac to 20.7.1a or vice-versa		
Workaround:	None		

Parent Defect ID:	SLXOS-77462	**Issue ID:**	SLXOS-78497
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3
Symptom:	Gradual increase in DCMD memory utilization		
Condition:	When XCO triggers deep device discovery		

Parent Defect ID:	SLXOS-78069	**Issue ID:**	SLXOS-78498
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3b
Symptom:	Slow increase in memory usage of DCMD		
Condition:	When show commands, "show process cpu", "show bfd neighbor", "show ssh server status" are run periodically		

Note for SLXOS-78069: This fix addresses memory issues caused by periodic execution of specific show commands (specifically, `show process cpu`, `show bfd neighbor`, and `show ssh server status`). However, memory issues are still being observed when `show logging raslog` is executed periodically. Due to the release timeline for SLX-OS version 20.7.1b, this issue was not addressed in this release. A separate defect **SLXOS-78542** (*Memory Spike in DCMD Process During Automated 'show' Command Execution on EFA-Managed SLX Switch*) has been created to track and resolve this issue in the subsequent release.

Defects closed with code changes in SLX-OS 20.7.1a

The following software defects were closed with code changes in SLX-OS 20.7.1a as of April 2025:

Parent Defect ID:	SLXOS-77710	**Issue ID:**	SLXOS-77710
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3a
Technology Group:	Management	**Technology:**	Other
Symptom:	Unable to login SLX with throwing error ""SLX-OS is not ready. Please login later".		
Condition:	SLX made to run with multiple CLI/NETCONF session.		

Parent Defect ID:	SLXOS-77952	**Issue ID:**	SLXOS-77952
Severity:	S3 - Moderate		
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.7.1
Technology Group:	Management	**Technology:**	Software Installation & Upgrade
Symptom:	Unexpected Error message is observed in console during the firmware upgrade.		
Condition:	When the image is upgrade from 20.6.3x to 20.7.1. An error message is displayed in the console init logs.		
Workaround:	None		
Recovery:	This error message is cosmetic and does not have any impact during installation. The installation proceeds normally and a fully functional system is achieved at the end of the installation process. No intervention is required.		

Defects closed with code changes in SLX-OS 20.7.1

The following software defects were closed with code changes in SLX-OS 20.7.1 as of April 2025:

Parent Defect ID:	SLXOS-76158	**Issue ID:**	SLXOS-76158
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.1

Symptom:	Receive-ACL configuration is accepting a duplicate rule
Condition:	After loading a new startup or running configuration file
Workaround:	Remove and re-add the ACL configuration

Parent Defect ID:	SLXOS-76501	**Issue ID:**	SLXOS-76501
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.1
Symptom:	Traffic will always egress on the Traffic Class 0		
Condition:	Configure and unconfigure dscp-traffic-class map on MCT ICL link		

Parent Defect ID:	SLXOS-76556	**Issue ID:**	SLXOS-76556
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.4.3c
Symptom:	Transient traffic loops on a Cluster Client Endpoint (CCEP) can cause unexpected MAC movements and MAC reachability issues in the MCT IP fabric.		
Condition:	Traffic on a Cluster Client Endpoint (CCEP) may transiently loop back for a few milliseconds due to MCT MAC sync delays, particularly during MAC aging or learning.		
Recovery:	MAC movements recover quickly if the destination MAC is active; otherwise, recovery takes longer until it's relearned on the correct interface.		

Parent Defect ID:	SLXOS-76971	**Issue ID:**	SLXOS-76971
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.4.3d
Symptom:	Packets are forwarded with an incorrect MPLS label, resulting in a traffic outage		
Condition:	Peers establish LDP session for label distribution		

Parent Defect ID:	SLXOS-77206	**Issue ID:**	SLXOS-77206
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3b
Symptom:	System reload is observed		
Condition:	Adding and removing bgp rpki server configs multiple times triggers the crash		

Parent Defect ID:	SLXOS-77281	**Issue ID:**	SLXOS-77281
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.1a
Symptom:	Transit traffic redirected over the MPLS ECMP nexthop will get dropped		
Condition:	RSVP/LDP tunnels load sharing at transit routers		
Workaround:	Consider removing the load sharing at transit routers.		

Parent Defect ID:	SLXOS-77282	**Issue ID:**	SLXOS-77282
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.2a

Symptom:	on SLX-9250, When upgrading to SLX20.6.1 or higher firmware on devices with DC PSU. The DC PSU presence looks toggling sporadically based on the RASLogs.		
Condition:	Issue is seen only on SLX20.6.1 or later firmware images		
Workaround:	These RASLOGS can be Suppressed, as it is not a functional issue.		
Recovery:	Fix is done to take care of the DC PSU and recovery can be done by moving to firmware image with the fix 20.6.3a.		

Parent Defect ID:	SLXOS-77307	**Issue ID:**	SLXOS-77307
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.1a
Symptom:	"ecn" option is not available under "qos red-profile" CLI command		
Condition:	ECN support is not enabled in Extreme 8720 platform		

Parent Defect ID:	SLXOS-77310	**Issue ID:**	SLXOS-77310
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.1a
Symptom:	RED with ECN configuration is rejected when PFC is enabled		
Condition:	Configuring PFC and RED with ECN together for the same Traffic Class		

Parent Defect ID:	SLXOS-77320	**Issue ID:**	SLXOS-77320
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3c
Symptom:	IGP cost is not considered for bgp l3vpn route best path calculation, which results in non-optimal path selection for routes.		
Condition:	When 'nexthop-mpls follow-igp-metric' is configured under bgp address family.		
Workaround:	NA		
Recovery:	NA		

Parent Defect ID:	SLXOS-77327	**Issue ID:**	SLXOS-77327
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.4.3ac
Symptom:	HSLagt daemon is terminated		
Condition:	1. Monitor session is configured with source interface as VLAN/VE 2. Executing “debug sflow show span session all” CLI command		

Parent Defect ID:	SLXOS-77343	**Issue ID:**	SLXOS-77343
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3c
Symptom:	Broadcast DHCP packets received on a CCEP interface were incorrectly looping between MCT peers over the ICL, leading to excessive DHCP packet punting to the CPU. Since DHCP and VRRP packets share the same CPU queue, some VRRP control packets were dropped, resulting in VRRP failures.		
Condition:	When a DHCP relay agent is configured on any VE, the handling of broadcast DHCP packet flooding is managed by the software. Upon receiving a broadcast DHCP packet on a CCEP/CEP port, the MCT		

	forwards the packet to the software, which then floods it across the L2 domain. This process includes duplicating the packet to local client ports and transmitting it to the peer MCT node over the ICL. Upon receiving this packet over the ICL, the peer MCT node was expected to flood it only to local CEP/CCEP ports. However, due to the local client port-channel being down, the packet was incorrectly flooded back to the ICL with the intention of reaching the client via the peer MCT node. This caused the same packet to return to the originator MCT node, triggering another round of flooding. The problem becomes worse if the client port-channel is down on both MCT nodes, as the packet continuously loops between the MCT nodes over the ICL.
Workaround:	The loop can be resolved by globally disabling DHCP relay flooding using the CLI command "ip dhcp relay disable-flooding". However, there is no per-VE CLI option to restrict flooding for specific VLANs or BDs. As a result, this global setting may inadvertently block some clients from reaching DHCP servers.

Parent Defect ID:	SLXOS-77374	**Issue ID:**	SLXOS-77374
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.2a
Symptom:	When BGP session goes down SNMP trap generated has localPort and RemotePort are always zero.		
Condition:	When BGP session goes down.		
Workaround:	NA		
Recovery:	NA		

Parent Defect ID:	SLXOS-77420	**Issue ID:**	SLXOS-77420
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.2
Symptom:	when doing SNMP walk mib object the TX Power value is not displaying properly.		
Condition:	when executing the below SNMP command snmpwalk -v2c -c public 10.38.135.176 1.3.6.1.4.1.1588.3.1.8.1.2.1		

Parent Defect ID:	SLXOS-77505	**Issue ID:**	SLXOS-77505
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3
Symptom:	"Identification data corruption detected" is thrown on the console.		
Condition:	During optic removal and insertion or during reload on some breakout ports randomly.		
Workaround:	No workaround.		

Parent Defect ID:	SLXOS-77564	**Issue ID:**	SLXOS-77564
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.2a
Symptom:	SLX firmware commit shows "command failed" even if the commit is successful.		

Condition:	During SLX firmware upgrade or commit.

Parent Defect ID:	SLXOS-77569	**Issue ID:**	SLXOS-77569
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3c
Symptom:	BFD sessions are flaps with VSP routers.		
Condition:	When peer is VSP routers.		
Workaround:	NA		
Recovery:	NA		

Parent Defect ID:	SLXOS-77618	**Issue ID:**	SLXOS-77618
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3c_CVR
Symptom:	L3 routing table pointing to the previous best path		
Condition:	With BGP PIC enabled, during path change, the route table may continue to point to the old best path		

Parent Defect ID:	SLXOS-77663	**Issue ID:**	SLXOS-77663
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3a
Symptom:	Random link flaps observed in the platforms supporting external PHY.		
Condition:	Default config, seen when there are more connections present.		

Parent Defect ID:	SLXOS-77690	**Issue ID:**	SLXOS-77690
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3c
Symptom:	HSLagt daemon is terminated leading to a reboot of the node		
Condition:	When collecting the Support Save on SLX 9540 and SLX 9640		

Parent Defect ID:	SLXOS-77696	**Issue ID:**	SLXOS-77696
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.5.3
Symptom:	SLX node is reporting a BGP Keepalive TTL exceeded Timeout error		
Condition:	It is applicable for Extreme 8520, Extreme 8720, SLX 9150 and SLX 9250 platforms only, when the node is acting as a transit Layer 2 device between BGP peers		

Parent Defect ID:	SLXOS-77746	**Issue ID:**	SLXOS-77746
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.3b
Symptom:	Unable to configure password-policy for lower character		
Condition:	password-policy for lower character is not reflected on TPVM when it's configured from TPVM config mode on SLX		

Parent Defect ID:	SLXOS-77906	**Issue ID:**	SLXOS-77906
Product:	SLX-OS	**Reported in Release:**	SLXOS 20.6.2ac
Symptom:	Unexpected reboot on the SLX device.		
Condition:	Default config, seen when there are more connections present.		

Defects Closed without Code Changes

Defects closed without code changes in SLX-OS 20.7.1b

No software defects were closed in SLX-OS 20.7.1b without code changes as of September 2025.

Defects closed without code changes in SLX-OS 20.7.1a

No software defects were closed in SLX-OS 20.7.1a without code changes as of April 2025.

Defects closed without code changes in SLX-OS 20.7.1

No software defects were closed in SLX-OS 20.7.1 without code changes as of April 2025.