

June 2019



SLX-OS 18r.1.00c for SLX 9850 and SLX 9540

Release Notes v2.0

Copyright © 2019 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, please see:

www.extremenetworks.com/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:

www.extremenetworks.com/support/policies/software-licensing

Contents

Document history	4
Preface	5
Contacting Extreme Technical Support.....	5
Extreme resources	5
Document feedback.....	6
Overview	7
New SKUs	8
Behavior changes	9
Behavior changes in release 18r.1.00c	9
Behavior changes in release 18r.1.00b	9
Behavior changes in release 18r.1.00a	10
Behavior changes in release 18r.1.00	11
Software Features.....	12
New software features in 18r.1.00c	12
New software features in 18r.1.00b	12
New software features in 18r.1.00a	12
New software features in 18r.1.00	13
CLI commands.....	15
CLI commands introduced in R18r.1.00c	15
CLI commands introduced in R18r.1.00b.....	15
CLI commands introduced in R18r.1.00a	15
CLI commands introduced in R18r.1.00.....	16
RFCs, Standards, and Scalability.....	17
Hardware support.....	18
Supported devices.....	18
Supported power supplies	19
Supported optics.....	19
Software upgrade and downgrade	21
Image file names	21
Upgrade and downgrade considerations.....	21
Limitations and restrictions	30
Defects closed in 18r.1.00c.....	34

Document history

Version	Summary of changes	Publication date
2.0	Updated Defect List	June 2019
1.0	Initial Release	June 2019

Preface

Contacting Extreme Technical Support

As an Extreme customer, you can contact Extreme Technical Support using one of the following methods: 24x7 online or by telephone. OEM customers should contact their OEM/solution provider.

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

- GTAC (Global Technical Assistance Center) for immediate support
- Phone: 1-800-998-2408 (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.
- Email: support@extremenetworks.com. To expedite your message, enter the product name or model number in the subject line.
- GTAC Knowledge - Get on-demand and tested resolutions from the GTAC Knowledgebase or create a help case if you need more guidance.
- The Hub - A forum for Extreme customers to connect with one another, get questions answered, share ideas and feedback, and get problems solved. This community is monitored by Extreme Networks employees but is not intended to replace specific guidance from GTAC.
- Support Portal - Manage cases, downloads, service contracts, product licensing, and training and certifications.

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

- Your Extreme Networks service contract number and/or serial numbers for all involved Extreme Networks products
- A description of the failure
- A description of any action(s) 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

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 www.extremenetworks.com. Product documentation for all supported releases is available to registered users at <https://www.extremenetworks.com/support/documentation/>.

Document feedback

Quality is our first concern at Extreme, and we have made every effort to ensure the accuracy and completeness of this document. However, if you find an error or an omission, or you think that a topic needs further development, we want to hear from you.

You can provide feedback in two ways:

- Use our short online feedback form at <https://www.extremenetworks.com/documentation-feedback/>
- 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.

Overview

SLX-OS 18r.1.00c has multiple customer found defect solutions, along with new software features. Refer to the section, **New software features in 18r.1.00c** for more information.

SLX-OS 18r.1.00a supports the following solutions and features:

- L2 Exchange
 - Egress ACL rate-limiting: This is a key feature for IXP customers to flexibly rate-limit ACL filtered traffic on port/VLAN/BD

- vSLX*
 - IP fabric BGP EVPN VXLAN control plane
 - L2 Exchange control plane

Note:

vSLX is community supported in the Extreme “The Hub” Community pages for Switching and Routing Data Center products.

https://community.extremenetworks.com/extreme/categories/extreme_switchingrouting

Details of support process for vSLX is available in the vSLX guide and release note.

New SKUs

No new SKUs are introduced in this release.

Behavior changes

Behavior changes in release 18r.1.00c

The following system behaviors have changed in this release:

None.

Behavior changes in release 18r.1.00b

The following system behaviors have changed in this release:

The management module for the SLX9850 product family has been upgraded from 16GB of memory to 32GB of memory. The deprecated version of the management module with 16GB of memory is supported up to software release SLX-OS 18r.2.00.

Extreme highly recommends upgrading to the 32GB version of the management module.

If the combination of the 16GB (standby) and 32GB (active) management modules are installed in an SLX 9850, the following RAS log message will appear once the 32GB module becomes active:

M1 | Active | FFDC, WARNING, SLX9850-8, Detected system memory size mismatch on dual MM - active has 32GB and standby has 16GB.

NOTE: A RAS log message will not appear on the console if the 32GB management module is in standby mode.

TCAM profiles

Statistics are supported for all rate-limiting subtypes in the "layer2-ratelimit" TCAM profiles. Also, with this profile, services such as ACLs, BUM traffic, and port rate-limiting share the same TCAM space with L2/L3 ACLs and therefore have a lower priority. Therefore, if traffic matches an ACL entry, rate-limiting is not applied.

VC-Mode Tag

In the previous, "VC-Mode Tag" mode, the ingress PE original packet's PCP is used only to classify packet to the appropriate traffic class (queue) on the ingress PE; egress PE does not use this value. Therefore, with this change in effect, the ingress PE received packet's original PCP is now forced to be stamped on the outgoing AC endpoint at the egress PE. In addition, the PE that is configured with this mode forces all locally switched traffic to *always* honor the original ingress packet's PCP value and uses that value as its outgoing AC endpoint's PCP value.

For example, in a BD setup, there are 2 AC LIFs and a VPLS peer. One AC LIF is configured with a single tagged, and the other AC LIF is configured as a dual tagged AC LIF. If the packet is received from the single tagged LIF and destined to the dual tagged AC LIF, then its PCP is copied to the outgoing packet to both outer and inner VLAN TAGs. If the packet is received from the dual tagged, and destined to the single tagged AC LIF, then the inner tag's PCP will be copied to the outgoing packet's VLAN tag. If the packet is received from the remote peer, then the VLAN tag's PCP that was sent over by the peer will be used towards the destination AC LIF. This case is similar to the single tagged case as mentioned in the

above scenario. With this change in effect, the internal TC will no longer be used for stamping onto the outgoing packet at the egress PE.

IP MTU

In 18r.1.00a, the hardware only supports three MTU values, including the default value of 1500 and two user-defined values. If the limit is reached, the following error message is displayed:

```
%Error: Maximum limit of allowed different IP MTUs reached.
```

Note: This behavior change applies to both global and interface MTU.

In releases prior to 18r.1.00a, the recommendation was to use one the following three values: 1300, 1500 and 9194. If the user configuration did not match these values, the previous lower value is selected, i.e. if 1400 is configured HW will be programmed with 1300, similarly if 9000 is configured HW will be programmed with IP MTU of 1500. However, this created a mismatch between the user-configured value and hardware-programmed value.

Note: The `ipv6 mtu` command is deprecated in this release.

Behavior changes in release 18r.1.00a

The following system behaviors have changed in this release:

- Unknown unicast Storm control feature will bypass VLL (bridge-domain p2p) traffic in L2-optimized profile starting this release.
- If the user tries to bind the policer with configured CIR/EIR value is less than 22000 bps in Fusion/Avalanche the operational CIR/EIR will be zero and the same will be notified to the user via syslog on console.
- When working with different rate-limiting subtypes, the precedence order will be as follows in layer2-rate-limit profile:
ACL RL -> VLAN/BD RL -> Port RL -> BMU storm-control
- The management module for the SLX9850 product family has been upgraded from 16GB of memory to 32GB of memory. The deprecated version of the management module with 16GB of memory is supported up to software release SLX-OS 18r.2.00.

Extreme highly recommends upgrading to the 32GB version of the management module.

If the combination of the 16GB (standby) and 32GB (active) management modules are installed in an SLX 9850, the following RAS log message will appear once the 32GB module becomes active:

```
M1 | Active | FFDC, WARNING, SLX9850-8, Detected system memory size mismatch on dual MM - active has 32GB and standby has 16GB.
```

NOTE: A RAS log message will not appear on the console if the 32GB management module is in standby mode.

IP MTU

In 18r.1.00a, the hardware only supports three MTU values, including the default value of 1500 and two user-defined values. If the limit is reached, the following error message is displayed:

```
%Error: Maximum limit of allowed different IP MTUs reached.
```

Note: This behavior change applies to both global and interface MTU.

In releases prior to 18r.1.00a, the recommendation was to use one of the following three values: 1300, 1500 and 9194. If the user configuration did not match these values, the previous lower value is selected, i.e. if 1400 is configured HW will be programmed with 1300, similarly if 9000 is configured HW will be programmed with IP MTU of 1500. However, this created a mismatch between the user-configured value and hardware-programmed value.

Note: The `ipv6 mtu` command is deprecated in this release.

Behavior changes in release 18r.1.00

The following system behaviors have changed in this release:

- DNS improvement: Now we do not need to configure a domain name using command “ip dns domain-name ...”, along with “ip dns name-server <ip>” command. Previously, we had to configure domain name for DNS resolving to happen.
- Multi-VRF support for NTP client and server: Previously, NTP client tried to reach the server via mgmt-vrf instead of default-vrf. Support for default-vrf was added in this release.
- A defect in the firmware of the SSD used in the SLX 9540 may cause the SSD to stop responding. This is not corrected in the 18r.1.00 release. When this defect happens, the Extreme SLX-OS places the file system into a read-only mode to indicate that the file system is hung. A power-cycle fully recovers the device. An SSD controller firmware update is available, and GTAC can assist you in applying this to your systems. Please refer to Field Notice, FN-2018-422 for more information.

Software Features

New software features in 18r.1.00c

The following software features are new in this release:

- AAA command authorization is supported for TACACS+.
- The egress throughput rate on its breakout ports is increased when a breakout connector is configured on an Ethernet interface.
- MCT upgrade process from SLX-OS 17r.1.01x to SLX-OS 18r.1.00c.
- MCT upgrade process from SLX-OS 18r.1.00 to SLX-OS 18r.1.00c.
- The **show mpls ldp tunnel detail** command now displays the state of the LDP tunnels with counter profiles.
- IP directed broadcast is supported on an interface to all devices within a directly attached network or subnet.

New software features in 18r.1.00b

The following software features are new in this release:

- BUM storm control is now supported at the global and interface-level mode for Ethernet interfaces.
- Extreme VLAN MIB - Supported Enterprise MIB Objects – VLAN extended statistics MIB support
- Q-Bridge - Updated SNMP Q-Bridge MIB details – Q-BRIDGE-MIB is supported for read-only access
- Upgrading to release 18r.1.00b from 2.6 kernel 32-bit systems
- Upgrading to release 18r.1.00b from 2.6 kernel 64-bit systems
- Endpoint tracking enhancements

New software features in 18r.1.00a

The following software features are new in this release:

- Egress ACL-based Rate-limiting
- Source Interface for SNMP Trap/Notification/Inform
- SNMP MIB set support for sysContact, sysName, sysLocation, and ifAdminStatus objects
- vSLX

New software features in 18r.1.00

The following software features are new in this release:

- Endpoint tracking:
 - Minimizes the configuration and management of VLANs on switches in the data center
 - Supports authentication of macs using 802.1x protocol
 - Supports assignment of VLAN using NAC to authenticated mac, dynamic VLAN creation of the same and port vlan assignment
 - Can be used for dynamic VLAN management per host using NAC configuration
- Multiple VLAN Registration Protocol (MVRP):
 - Layer 2 protocol allows the dynamic propagation of VLAN information from device to device; manually configure an MVRP-aware access device with all desired VLANs for the network and all other MVRP-aware devices on the network learn these VLANs
- Layer 2 Loop Detection:
 - In Loose Mode, the LD shutdown action is changed from physical port level shutdown to LIF (Logical Interface) level shutdown. This keeps the physical port up and prevents other VLANs on the same physical port from being impacted.
- Data streaming Enhancements:
 - JSON encoding support
 - TM VoQ statistics
 - Line card CPU and Memory Statistics
 - MPLS traffic statistics data streaming
- Increase scale support for class-maps under the service policy:
 - Number of class maps per policy-map has been changed from 128 to 4K.
- Host and the SLX VM image snapshots:
 - Creates snapshot image of the currently running Host and the SLX VM images
- MPLS Enhancements:
 - Syslog enhancement for LDP session down events
 - Logging LSP down reason in syslog
 - Sort the output of “show mpls ldp tunnel” command by FEC address

BGP EVPN VxLAN based IP Fabric**

BGP EVPN IP Fabric is a controller-less architecture that simplifies data center operations by leveraging open, standards-based protocols to abstract network control plane, data plane, and automation functions from the underlying physical platforms. BGP EVPN Network Virtualization builds upon underlying infrastructure platforms, fabrics, and automation to deliver simplified and secure network operations.

The following features are supported:

- BGP EVPN support: Support for EVPN route types (Inclusive Multicast, MAC/MACIP routes, IPv4/v6 prefix routes, ES routes, AD routes)
- Dynamic tunnel (VxLAN) discovery: Supports Dynamic Tunnel discovery using BGP EVPN
- Bridge Domain Support: BGP-EVPN is supported over basic VLAN and Bridge-Domain
- ARP Suppression: Suppress/reduce the ARP broadcast traffic in an IP fabric.
- Static Anycast Gateway: Static Anycast Gateway allows configuring Static Anycast MAC as gateway for multiple tenant systems in a virtualized data center fabric. Same Gateway address is configured across all TORs for a given Tenant/VLAN combination, thus enabling seamless VM mobility across the leaf switches in an IP Fabric deployment without any need for host gateway configuration changes.
- IP Unnumbered Interfaces: Reduces consumption of IP Address space. Leaf to spine inter-switch point-to-point L3 links are configured as ip unnumbered to conserve IP addresses and optimize hardware resources.
- L2VNI capability: The L2VNI is the MAC/NVE mapping table
- L3VNI/routing capability: Default and non-default VRF routing/L3VNI are supported with BGP-EVPN. Symmetric and Asymmetric IRB are supported
- Logical VTEP: A logical VXLAN tunnel end point (LVTEP) is supported for both Layer 2 and Layer 3 for SLX 9540 only. SLX 9850 is not supported as leaf in VXLAN IP fabric.

REST API Support: All configuration operations supported in CLI are supported via REST. Selected BGP show commands for EVPN are supported with REST.

**Please note that the support for the L3VNI comes with limitations, which will be resolved in the patch release of 18r.1.00. Extreme recommends using 18r.1.00 for evaluation or controlled deployments of L3VNI.

CLI commands

CLI commands introduced in R18r.1.00c

New commands

The following commands are new in this release:

- aaa authorization command
- ip directed-broadcast

Modified commands

The following commands have been modified for this release:

- qos port-speed-up
- seq (rules in MAC extended ACLs)

CLI commands introduced in R18r.1.00b

New commands

The following commands are new in this release:

- storm-control ingress (global)
- endpoint-tracking timeout reauth-period
- qos port-speed-up

Modified commands

The following commands have been modified for this release:

- ip mtu
- show storm-control

Deprecated commands

- ipv6 mtu

CLI commands introduced in R18r.1.00a

New commands

The following command is new in this release:

- service-policy out <policy-name>

Modified commands

The following command have been modified for this release:

- snmp-server host

Deprecated commands

- ipv6 mtu

CLI commands introduced in R18r.1.00

New commands

The following commands are new in this release:

- bypass-lsp (Telemetry)
- clear ip arp suppression-cache
- clear ip arp suppression-statistics
- clear ipv6 nd suppression-cache
- clear ipv6 nd suppression-statistics
- clear mrvp statistics
- endpoint-tracking enable
- fec (telemetry)

Modified commands

The following commands have been modified for this release:

- show arp
- show arp summary
- show ipv6 neighbor
- show ipv6 neighbor summary

Deprecated commands

There are no deprecated commands in this release.

RFCs, Standards, and Scalability

For RFCs, standards, and scale numbers supported in this release, refer to the [Extreme SLX-OS Scale and Standards Matrix for SLX 9850 and SLX 9540](#).

Hardware support

Supported devices

The following devices are supported in this release:

Supported Hardware	Description
BR-SLX9850-4-BND-AC	Extreme SLX 9850 4-slot chassis with 1 management module, 5 switch fabric modules, 2 3000W AC power supplies, 3 fan modules, and accessory kit. Power cord not included.
BR-SLX9850-4-BND-DC	Extreme SLX 9850 4-slot chassis with 1 management module, 5 switch fabric modules, 2 3000W DC power supplies, 3 fan modules, and accessory kit. Power cord not included.
BR-SLX9850-8-BND-AC	Extreme SLX 9850 8-slot chassis with 1 management module, 5 switch fabric modules, 4 3000W AC power supplies, 3 fan modules, and accessory kit. Power cord not included.
BR-SLX9850-8-BND-DC	Extreme SLX 9850 8-slot chassis with 1 management module, 5 switch fabric modules, 4 3000W DC power supplies, and 3 fan modules, and accessory kit. Power cord not included.
BR-SLX9850-10GX72S-M	Extreme SLX 9850 72-port 10 GbE/1 GbE dual-speed (M) interface module with IPv4/IPv6/MPLS hardware support. Requires SFP+ optics for 10 GbE connectivity and SFP optics for 1 GbE connectivity. Supports up to 750,000 MAC. Supports up to 1,500,000 IPv4 routes, 140,000 IPv6 routes with OptiScale™ Internet Routing.
BR-SLX9850-100GX36CQ-M	Extreme SLX 9850 36-port 100 GbE, 60-port 40 GbE, or 240-port 10 GbE flex-speed (M) interface module with IPv4/IPv6/MPLS hardware support. Requires QSFP28 optics for 100 GbE, QSFP+ optics for 40 GbE, and 40 GbE to 10 GbE breakout for 10 GbE connectivity. Supports up to 750,000 MAC. Supports up to 1,500,000 IPv4 routes, 140,000 IPv6 routes with OptiScale™ Internet Routing.
BR-SLX9850-10GX72S-D	Extreme SLX985072-port 10GbE/1GbE (D) interface module with IPv4/IPv6 hardware support. Requires SFP+ optics for 10GbE connectivity and SFP optics for 10GbE connectivity. Supports 750K MAC, 256K IPv4 routes and 64K IPv6 routes with up to 8GB packet buffers
BR-SLX9850-100GX36CQ-D	Extreme SLX 9850 36-port 100GbE, 60-port 40GbE, or 240-port 10GbE flex-speed (D) interface module with IPv4/IPv6 hardware support. Requires QSFP28, QSFP+ optics & 40GbE to 10GbE
BR-SLX9850-100GX12CQ-M	Extreme SLX 9850 12-port 100 GbE, 20-port 40GbE, or 80-port 10GbE flex-speed (M) interface module with IPv4/IPv6/MPLS hardware support. Requires QSFP28, QSFP+ optics & 40GbE to 10GbE breakout (for 10 GbE) connectivity. Supports up to 750,000 MAC. Supports up to 1,500,000 IPv4 routes, 140,000 IPv6 routes with OptiScale™ Internet Routing.
BR-SLX9850-100GX6CQ-M-UPG	6x100G POD SW license to be used with SLX9850-100Gx12CQ-M 100G blade only
XBR-SLX9850-4-S	Extreme SLX9850 Spare 4-slot chassis
XBR-SLX9850-8-S	Extreme SLX9850 Spare 8-slot chassis
BR-SLX9850-MM	Extreme SLX 9850 management module for 4-slot and 8-slot systems, includes 16GB RAM, 2 internal Solid State Drives, 4-Core Intel CPU, 2 USB 3.0 ports, 2 RJ-45 console ports, and 10GbE Services port
BR-SLX9850-4-SFM	Extreme SLX 9850 switch fabric module for 4-slot chassis
BR-SLX9850-8-SFM	Extreme SLX 9850 switch fabric module for 8-slot chassis
XBR-SLX9850-ACPWR-3000	Extreme SLX 9850 AC 3000W power supply for 4- and 8-slot chassis, 90-270V AC input
XBR-SLX9850-DCPWR-3000	Extreme SLX 9850 DC 3000W power supply for 4- and 8-slot chassis
XBR-SLX9850-4-FANM	Extreme SLX 9850 fan module for 4-slot chassis. Fan module has 2 fans
XBR-SLX9850-8-FANM	Extreme SLX 9850 fan module for 8-slot chassis. Fan module has 4 fans
XBR-SLX9850-4-CAB	Extreme SLX 9850 Cable Combo Kit for 4-slot chassis
XBR-SLX9850-8-CAB	Extreme SLX 9850 Cable Combo Kit for 8-slot chassis
XBR-SLX9850-4-SFMPNL	Extreme SLX 9850 switch fabric module blank panel for 4-slot chassis
XBR-SLX9850-8-SFMPNL	Extreme SLX 9850 switch fabric module blank panel for 8-slot chassis
XBR-SLX9850-PWRPNL	Extreme SLX 9850 power supply blank panel for 4-slot and 8-slot chassis
XBR-SLX9850-IMPNL	Extreme SLX 9850 interface module blank panel for 4-slot and 8-slot chassis
XBR-SLX9850-MMPNL	Extreme SLX 9850 management module blank panel for 4-slot and 8-slot chassis
XBR-SLX9850-4-4PRM-KIT	Extreme SLX 9850 four-post rack mounting kit for 4-slot chassis. Include 27-31" flush and recessed mounting
XBR-SLX9850-4-2PRM-KIT	Extreme SLX 9850 two-post rack mounting kit for 4-slot chassis. Include telco flush and midplane mounting
XBR-SLX9850-8-4PRM-KIT	Extreme SLX 9850 four-post rack mounting kit for 8-slot chassis. Include flush and recessed mounting
XBR-SLX9850-8-2PRM-KIT	Extreme SLX 9850 two-post rack mounting kit for 8-slot chassis. Include telco flush and midplane mounting

Supported Hardware	Description
BR-SLX-9540-24S-AC-F	Extreme SLX 9540-24S Switch AC with Front to Back airflow. Supports 24x10GE/1GE + 24x1GE ports
BR-SLX-9540-24S-DC-F	Extreme SLX 9540-48S Switch DC with Front to Back airflow. Supports 48x10GE/1GE + 6x100GE/40GE
BR-SLX-9540-24S-AC-R	Extreme SLX 9540-24S Switch AC with Back to Front airflow. Supports 24x10GE/1GE + 24x1GE ports
BR-SLX-9540-24S-DC-R	Extreme SLX 9540-24S Switch DC with Back to Front airflow. Supports 24x10GE/1GE + 24x1GE ports
BR-SLX-9540-48S-AC-F	Extreme SLX 9540-48S Switch AC with Front to Back airflow. Supports 48x10GE/1GE + 6x100GE/40GE
BR-SLX-9540-48S-DC-F	Extreme SLX 9540-48S Switch DC with Front to Back airflow. Supports 48x10GE/1GE + 6x100GE/40GE
BR-SLX-9540-48S-AC-R	Extreme SLX 9540-48S Switch AC with Back to Front airflow. Supports 48x10GE/1GE + 6x100GE/40GE
BR-SLX-9540-48S-DC-R	Extreme SLX 9540-48S Switch DC with Back to Front airflow. Supports 48x10GE/1GE + 6x100GE/40GE
BR-SLX-9540-24S-COD	Upgrade 24x1GE to 24x10GE/1GE
BR-SLX-9540-2C-POD	Ports on Demand for 2x100GE/40GE Uplinks
BR-SLX-9540-ADV-LIC-P	Advanced Feature License for MPLS, BGP-EVPN, CE2.0, NSX, OptiScale™ Internet Routing (for Extreme SLX 9540-24S and 9540-48S)

Supported power supplies

- Extreme SLX 9850 AC 3000W power supply for 4- and 8-slot chassis, 90-270V AC input
- Extreme SLX 9850 DC 3000W power supply for 4- and 8-slot chassis, 48V DC input

Supported optics

Part Number	Description
1G-SFP-TX	MODULE, MINI-GBIC, TX, 1000BASE, RJ45
1G-SFP-SX-OM	1000BASE-SX SFP OPTIC, MMF LC
1G-SFP-SX-OM-8	1000BASE-SX SFP OPTIC, MMF LC 8
1G-SFP-LX-OM	1000BASE-LX SFP OPTIC, SMF LC
1G-SFP-LX-OM-8	1000BASE-LX SFP OPTIC, SMF LC 8
1G-SFP-LHA-OM	1000BASE-LHA SFP OPTIC, SMF, LC CONN
1G-SFP-BXD	1000BASE-BXD SFP OPTIC SMF
1G-SFP-BXU	1000BASE-BXU SFP OPTIC SMF
10G-SFP-USR	10G USR SFP+ TRANS 100M OVER MMF
10G-SFP-SR	10G SR SFP+ TRANS 300M OVER MMF
10G-SFP-SR-8	10G SR-8 SFP+ TRANS 300M OVER MMF 8
10G-SFP-LR	10G LR SFP+ TRANS 10KM OVER SMF
10G-SFP-LR-8	10G LR SFP+ TRANS 10KM OVER SMF 8
10G-SFP-ER	10G ER SFP+ TRANS 40KM OVER SMF
10G-SFP-ZR	10GBASE-ZR SFP+ optic (LC), for up to 80km over SMF
10GE-SFP-AOC-0701	10GE SFP+ Direct Attach Cables 7m - Active Optical cables
10GE-SFP-AOC-1001	10GE SFP+ Direct Attach Cables 10m - Active Optical cables
10G-SFP-TWX-0101	10 GbE SFP+ optics Twinax Active Copper cable: 1m
10G-SFP-TWX-0301	10 GbE SFP+ optics Twinax Active Copper cable: 3m
10G-SFP-TWX-0501	10 GbE SFP+ optics Twinax Active Copper cable: 5m
40G-QSFP-SR4	40G QSFP+ SR4 TRANS 100M OVER MMF
40G-QSFP-SR4-INT	40G QSFP+ 100M OVER MMF 10G BREAKOUT
40G-QSFP-ESR4-INT	40G QSFP+ 300M OVER MMF 10G BREAKOUT
40G-QSFP-LR4	40G QSFP+ LR4 TRANS 10KM OVER SMF
40G-QSFP-QSFP-C-0101	40G QSFP+ TO QSFP+ ACTIVE COPPER 1M

Part Number	Description
40G-QSFP-QSFP-C-0301	40G QSFP+ TO QSFP+ ACTIVE COPPER 3M
40G-QSFP-QSFP-C-0501	40G QSFP+ TO QSFP+ ACTIVE COPPER 5M
40G-QSFP-QSFP-AOC-1001	40G QSFP+ to QSFP+ ACTIVE OPTICAL CABLE 10M
40G-QSFP-4SFP-C-0101	4X10GE QSFP+TO4SFP+ COPPER BREAKOUT 1M
40G-QSFP-4SFP-C-0301	4X10GE QSFP+TO4SFP+ COPPER BREAKOUT 3M
40G-QSFP-4SFP-C-0501	4X10GE QSFP+TO4SFP+ COPPER BREAKOUT 5M
40G-QSFP-4SFP-AOC-1001	4X10GE QSFP+TO4SFP+ Fiber BREAKOUT 10M
100G-QSFP28-CWDM4-2KM	100GBASE CWDM4 QSFP TRANS LC 2KM OVER SM
100G-QSFP28-SR4	100G QSFP28 SR4 TRANS 100M OVER MMF
100G-QSFP28-LR4L-2KM	100G QSFP28 LR4 LITE TRANS 2KM OVER SMF
100G-QSFP28-LR4-10KM	100G QSFP28 LR4 TRANS 10KM OVER SMF
100G-QSFP28-LR4-LP-10KM	100G QSFP28 LR4 LOWPOWER 2KM OVER SMF
100G-QSFP-QSFP-P-0101	100G QSFP Passive Direct Attach Copper Cable, 1M
100G-QSFP-QSFP-P-0301	100G QSFP Passive Direct Attach Copper Cable, 3M
100G-QSFP-QSFP-P-0501	100G QSFP Passive Direct Attach Copper Cable, 5M
100G-QSFP-QSFP-AOC-1001	100G QSFP Direct Attach Active Optical Cable, 10M
10G-SFPP-USR-E	10GE USR SFP+,HIGH RX SENSITIVITY
10G-SFPP-USR-8-E	10GE USR SFP+,HIGH RX SENSITIVITY (8-pack)
10G-SFP-USR-SA	10GE USR SFP+ OPTIC (LC),RANGE 100M MMF, TAA
10G-SFP-SR-S	10GBASE-SR, SFP+OPTIC(LC), 300M MMF, 70C
10G-SFP-LR-SA	10GBASE-LR, SFP+ OPTIC (LC),10KM OVERSMF, TAA, 70C
10G-SFP-BXU-S	10GE LR SFP+ OPTIC (LC) BIDIRECTIONAL UP
10G-SFP-BXD-S	10GE LR SFP+ OPTIC (LC) BIDIRECTIONAL DO
*Methode SP7051	Methode SP7051-BRCD SFP+ 10G-Base-T (10G speed only)
*Inphi IN-Q2AY2-XX	Inphi 100G QSFP-28 ColorZ DWDM (80km)

*Optics reference qualified and should be purchased from the respective vendors. Extreme doesn't sell these.

Software upgrade and downgrade

Image file names

Download the following images from www.extremenetworks.com.

Image file name	Description
slxos18r.1.00c.tar.gz	SLX-OS 18r.1.00c software
slxos18r.1.00c_all_mibs.tar.gz	SLX-OS 18r.1.00c MIBS
slxos18r.1.00c.md5	SLX-OS 18r.1.00c md5 checksum

Upgrade/downgrade considerations using firmware download CLI through fullinstall

The fullinstall CLI option is supported through the firmware download when upgrading from release SLX-OS 17r.1.01a to SLX-OS 17r.2.01. The fullinstall CLI option is NOT supported with USB.

Upgrade and downgrade considerations

- Upgrade from a 32-bit to 32-bit SLX-OS is performed using 'coldboot' option
- Upgrade from a 32-bit to 64-bit SLX-OS is a two-step sequential process as shown below:
 - 1) Upgrade using 'coldboot' to 17r.1.01a
 - 2) Upgrade using 'fullinstall' to 64-bit SLX OS
- Upgrade/Downgrade using 'fullinstall' takes up to 60 minutes for completion as compared to 25 minutes for 'coldboot'
- Upgrade from a 64-bit to 64-bit SLX-OS is performed using 'coldboot' option
- It is recommended to use 7zip or WinRAR to Un-compress the SLXOS tar file
- When firmware upgrade or downgrade is performed, following matrix can be used as a reference.

To	16r.1.0 (32-bit)	17r.1.0 (32-bit)	17r.1.01 (32-bit)	17r.1.01a (32-bit)	17r.2.0 (64-bit)	18r.1.0 (64-bit)
From						
16r.1.0 (32-bit)	N/A	N/A	N/A	N/A	Two Step Process: 1. Upgrade to 17r.1.01a 2. Upgrade to 17r.2.0	Two Step Process: 1. Upgrade to 17r.1.01a 2. Upgrade to 18r.1.0
17r.1.0 (32-bit)	N/A	N/A	N/A	N/A	Two Step Process: 1. Upgrade to 17r.1.01a 2. Upgrade to 17r.2.0	Two Step Process: 1. Upgrade to 17r.1.01a 2. Upgrade to 18r.1.0
17r.1.01 (32-bit)	N/A	N/A	N/A	N/A	Two Step Process: 1.Upgrade to 17r.1.01a 2.Upgrade to 17r.2.0	Two Step Process: 1. Upgrade to 17r.1.01a 2. Upgrade to 18r.1.0
17r.1.01a (32-bit)	N/A	N/A	N/A	N/A	fullinstall	fullinstall
17r.2.0 (64-bit)	Two Step Process: 1. Downgrade to 17r.1.01a 2. Coldboot to 16r.1.0	Two Step Process: 1. Downgrade to 17r.1.01a 2. Coldboot to 17r.1.0	Two Step Process: 1.Downgrade to 17r.1.01a 2. Coldboot to 16r.1.0	fullinstall	coldboot	coldboot
18r.1.0 (64-bit)	Two Step Process: 1. Downgrade to 17r.1.01a 2. Coldboot to 16r.1.0	Two Step Process: 1. Downgrade to 17r.1.01a 2. Coldboot to 17r.1.0	Two Step Process: 1.Downgrade to 17r.1.01a 2.Coldboot to 17r.1.01	fullinstall	coldboot	coldboot

Upgrade Steps from 32-bit to 64-bit SLX-OS

1. Make sure the device is running SLXOS 17r.1.01a or later, if not, please see the 17r.1.01 documentation on how to upgrade to that release.
2. Upgrade to SLX-OS 18r.1.00 using fullinstall
3. Save Configuration

To save the config, run

```
copy running-config startup-config
```

4. Firmware download with “fullinstall” option from source directory

```
device# firmware download fullinstall ftp user releaseuser password releaseuser file  
release.plist directory <path> host <host_ip>
```

Notes:

Firmware download with the “fullinstall” option will retain the startup configuration file, and upon auto reboot of the device, it will replay the startup configuration file automatically.

Firmware [download] upgrade support from SLXOS 18r.1.00b [Linux Kernel 2.6] to SLXOS 19.1.0 [Linux Kernel 4.14] is available from SLXOS 18r.1.00b onwards using "fullinstall" additional keyword.

Upgrade/downgrade using firmware download CLI through USB:

- Upgrade from SLX-SLX 17r.1.01a to SLX-OS 17r.2.01 is supported via firmware download CLI with “fullinstall” option.
- Upgrade from SLX 17r.1.01b to SLX-OS 17r.2.01a or later is supported via firmware download CLI with “fullinstall” option.
- USB based FWD upgrade from SLX-OS 17r.1.01a (32-bit) to SLX-OS 17r.2.01 (64-bit) or later is supported with “fullinstall” option.
- USB3.0 used for firmware download can be in VFAT or EXT4 format.

Instruction to check and upgrade FPGAs/CPLDs:

Refer to the *SLX-OS Upgrade Guide* for all variations on upgrading SLX-OS.

FPGA/CPLD versions:

SLX-9850	Release Date
MM sys FPGA	08/25/2016
LC sys FPGA	08/30/2016
SFM sys FPGA	08/04/2016
SLX-9540	Release Date
Sys FPGA	02/09/2017
CPLD 0	02/09/2017
CPLD 1	02/09/2017


MCT Upgrade Process

This section describes the process to upgrade MCT cluster nodes with minimum traffic loss disruption.

The MCT upgrade process is divided into the following sections:

1. MCT upgrade process from SLX-OS 17r.1.01x to SLX-OS 18r.1.00c (32-bit OS to 64-bit OS)
2. MCT upgrade process from SLX-OS 18r.1.00 to SLX-OS 18r.1.00c (64-bit OS to 64-bit OS)

The steps in the MCT upgrade process use the following nomenclature for MCT nodes: Node A and Node B.



MCT upgrade process from SLX-OS 17r.1.01x to SLX-OS 18r.1.00c:

This section describes the procedure to upgrade MCT cluster nodes from SLX-OS 17r.1.01x to SLX-OS 18r.1.00c and later releases with minimal traffic loss disruption.

This is a **32-bit OS to 64-bit OS upgrade** and hence uses the **firmware download** command with **fullinstall** option in order to perform the upgrade.

1. Configure client isolation mode under the cluster to be loose on Node A and on Node B respectively using the client-isolation loose command. For example:

```
Device(config)# cluster <Name of the cluster> <cluster-id>  
Device(config-cluster-1)# client-isolation loose
```

2. Isolate Node A from the network using the following steps:
 - a. Disable the MCT client-interfaces on Node A using client-interfaces-shutdown command under cluster configuration section.

```
Device-A(config-cluster-1)# client-interfaces-shutdown
```

b. Disable the link connected to MCT peer node and uplink to the core network.

This would result in all CCEP traffic to switch to Node B within 30 seconds depending on scale and other parameters.

3. Copy running-configuration to startup-configuration on node A.
4. Upgrade Node A to the 18r.1.00c release using the **firmware download fullinstall** command. While the upgrade on node A is in progress, the traffic would continue to pass through node B.
5. Verify that once the node comes UP, the member-vlan configuration under the cluster section is removed.
6. Create an evpn template and add to the existing configuration on Node A. For example:

```
Device(config)# evpn <evpn-instance-name>
route-target both auto ignore-as
rd auto
vlan add <NUMBER: 1-4090>                (If VLAN config is present)
bridge-domain add <NUMBER: 1-4090>      (If L2VPN config is present)
```

7. Isolate Node B from the network using the following steps. Please note that there is complete traffic loss at this step.
 - a. Disable the MCT clients from the Node B using **client-interfaces-shutdown** command under cluster configuration section.

```
Device-B(config-cluster-1)# client-interfaces-shutdown
```

- b. Disable the link connected to MCT peer node and uplink to the core.

Note: This step is suggested at this stage to avoid traffic duplication if L2VPN configuration is present. If L2VPN config is not present, enter the **no client-interfaces-shutdown** command on Node A before isolating Node B to minimize traffic loss. (Swap Step 7 and 9)

8. Copy running-configuration to startup-configuration on Node B.
9. Enable the interface towards the peer MCT node (ICL interface) and the uplink to the core network on Node A. (The ICL link would still be down since Node-B is isolated before this step. This is performed so that after Node B gets upgraded, the ICL link will come up once no shut is performed on the ICL link on Node-B.)

10. Bring Node A back to the network by entering the **no client-interfaces-shutdown** command under cluster configuration.

```
Device-A(config-cluster-1)# no client-interfaces-shutdown
```

This would result in all CCEP traffic to switch to Node A within 30 seconds depending on scale and other parameters.

11. Upgrade Node B to the 18r.1.00c release using the **firmware download fullinstall** command. While the upgrade on node B is in progress, the traffic would continue to pass through node A.
12. Verify that once the Node B comes UP, the member-vlan configuration under the cluster section is removed.
13. Create an evpn template and add to the existing configuration on Node B. For example:

```
Device-B(config)# evpn <evpn-instance-name>  
route-target both auto ignore-as  
rd auto  
vlan add <NUMBER: 1-4090> (If VLAN config is present)  
bridge-domain add <NUMBER: 1-4090> (If L2VPN config is present)
```

14. Enable the interface towards the peer MCT node (ICL) and the uplink to the core network on Node B.
15. Verify if the BGP session between the MCT peers is established and the cluster is up.
16. Bring Node B back to the network by entering the **no client-interfaces-shutdown command** under cluster configuration.

```
Device-B(config-cluster-1)# no client-interfaces-shutdown
```

17. Copy running-config to startup-config on both the nodes.

Additional upgrade considerations for upgrading SLX9850 from 17r.1.01a or 17r.1.01b to 18r.1.00c

When upgrading a SLX9850 from 17r.1.01a or 17r.1.01b to 18r.1.00c, if TPVM is installed in the system, you must un-install it by running the “tpvm uninstall” command before starting firmware download. Otherwise, it will cause system initialization issue. After the system is upgraded, you can install the TPVM image from 18r.1.00 by running the “tpvm install” command.

MCT upgrade process from SLX-OS 18r.1.00 to SLX-OS 18.1.00c:

This section describes the procedure to upgrade MCT cluster nodes from SLX-OS 18r.1.00 or 18r.1.00ax patch or 18r.1.00b patch to SLX-OS 18r.1.00c patch and later releases with minimal traffic loss disruption.

This is a **64-bit OS to 64-bit OS upgrade** and hence uses the **firmware download** command with **coldboot** option to perform the upgrade.

1. Configure client isolation mode under the cluster to be loose on Node A and Node B respectively using the client-isolation loose command. For example:

```
Device(config)# cluster <Name of the cluster> <cluster-id>
```

```
Device(config-cluster-1)# client-isolation loose
```

2. Isolate Node A from the network using the following steps:
 - a. Disable the MCT client-interfaces on Node A using **client-interfaces-shutdown** command under cluster configuration section.

```
Device-A(config-cluster-1)# client-interfaces-shutdown
```

- b. Interface connected to MCT peer node (ICL interface) must be left in **no shut** state.
- c. Disable uplink to the core network.

This causes all CCEP traffic to switch to Node B within 30 seconds depending on the scale and other parameters.

3. Copy running-config to startup-config on node A.
4. Upgrade Node A using **firmware download** command with **coldboot** option to the 18r.1.00c image. While the upgrade on node A is in progress, the traffic would continue to pass through node B.
5. Verify if Node A is back online after the upgrade and has completed initialization.
6. Isolate Node B from the network using the following steps.

Note: There is complete traffic loss at this step.

- a. Disable the MCT client-interfaces on Node B using **client-interfaces-shutdown** command under cluster configuration section.

```
Device-B(config-cluster-1)# client-interfaces-shutdown
```

- b. Interface connected to MCT peer node (ICL interface) must be left in **no shut** state.
- c. Disable uplink to the core network.

Note: This step is suggested at this stage in order to avoid traffic duplication if L2VPN configuration is present. If L2VPN configuration is not present, perform **no client-interfaces-shutdown** on Node A before isolating Node B in order to minimize traffic loss. (Swap Step-6 and Step-9)

- 7. Copy running-configuration to startup-configuration on Node B.
- 8. Enable the uplink to the core network on Node A. (The ICL interface would be up by now since we did not shut it prior to upgrade.)
- 9. Bring Node A back to the network by configuring the **no client-interfaces-shutdown** command under cluster configuration on Node A. This would result in all CCEP traffic to switch to Node A within 30 seconds depending on the scale and other parameters.

Device-A(config-cluster-1)# **no client-interfaces-shutdown**

- 10. Upgrade Node B to 18r.1.00c release using the **firmware download** command with **coldboot** option. While the upgrade on node B is in progress, the traffic would continue to pass through node A.
- 11. Verify that once the Node B comes UP, the uplink to the CORE network on Node B is configured to come up.
- 12. Verify if BGP session between MCT peers is established and the cluster is up.
- 13. Bring Node B back to the network by bringing the client-interfaces UP using the following command under cluster configuration.

Device-B(config-cluster-1)# **no client-interfaces-shutdown**

- 14. Copy running-config to startup-config on both the nodes.

Limitations and restrictions

Egress ACL-based Rate Limiting:

- Support in “layer2-ratelimit” TCAM profile only.
- Support CE ports only (that is, not support for MPLS uplinks).
- Broadcast, multicast and unknown unicast packets not supported.
- Port channel is not supported.
- Rate limit counters (conform/violate) not supported.

Additional Limitations

- Egress RL is designed to support the packet receiving at one physical port but transmitting on the different physical port. If the packets are received and transported on the same physical port, ingress rate-limit should be deployed.
- If multiple VLANs on the same ingress port belong to the same BD, and the egress ACL rate limiting is configured to rate limit one of the VLANs, all VLAN traffic is rate limited. A workaround is to add matching source or destination MAC address along with the VLAN in the ACL.
- Ingress ACL RL and egress ACL RL do not work together on the same flow of traffic.

CLI configuration design considerations for Rate limiting:

SLX 9850 (4 slots)

- 156 class maps per tower if bind to interface (VOQ limitation).
- 2k class maps per system (Supported hardware entry).
- 128 class maps for port-channel in default LAG profile, 128 LAG total (LAG hardware entry).
- 256 class map for port-channel in profile-1 LAG profile, 256 LAG total (LAG hardware entry).
- 1k policy maps per system (Software scaling).
- 32K class-maps per system (Software scaling).

Note: For SLX 9850 (8 slots), the VOQ limitation per tower is cut in half.

SLX 9540

- 64 class maps for port-channel in default LAG profile, 128 LAG total. (Support hardware entry).
- 128 class maps for port-channel in LAG profile-1, 128 LAG total. (Support hardware entry).
- 2k class maps per tower or system (Supported hardware entry).
- 1k policy maps per system (Software scaling).
- 32K class-maps per system (Software scaling).

QOS resource can be running out in following cases:

1. Out of resource when user bind the policy to a port or channel, in this case SW shall fail the command with error message.
2. Out of resource when user add new class to existing policy. SW shall fail the command to add new class with error message.
3. Out of resource when user add a new port to a LAG. In this case SW can't fail the command due to design limitation. Instead, it will send raslog to inform the user (the user need to check raslog and remove the config if resource running out).

sFlow: sFlow packet samples that are collected are inaccurate based on probability and low rate-limit.

L2 ACL: Unintentional traffic leaking can occur in a short period time (within 10 ms) during the adding of an L2 and L3 ACL.

Cos to TC mapping

- "qos map cos-traffic-class cosTC" command has known issue in this release and not taking effect for port channel.

VPLS VC

- In certain situations, VC peer flaps can happen in the VPLS network due to excessive amount of multicast traffic. To protect the control plane protocols, the following configuration is recommended on all ingress interfaces.
 - Apply BUM rate limit per interface

```
storm-control ingress broadcast limit-bps <rate in bps>
storm-control ingress multicast limit-bps <rate in bps>
storm-control ingress unknown-unicast limit-bps <rate in bps>
```

Rate limit values should be calculated based on amount of multicast traffic expected on the interface. Unknown-unicast should be as low as possible.

- Apply MCAST rate limit per forwarding ASIC

```
qos rx-queue multicast best-effort-rate <rate in kbps>
```

Command must be configured on one interface per ASIC. Actual rate depends of amount of expected MCAST traffic per forwarding ASIC.

BFD:

- Sessions with less than 300ms timer may flap in scale conditions
- Known issues with BFD when BFD is configured over multi-slot LAG, or multi-hop session over ECMP paths

L3VPN: Known issues with Peer-group, RR-group and Prefix-list ORF

FRR facility backup

- VPLS/VLL Bypass traffic will not work when router/untagged VE interfaces configured as MPLS uplink ports

MCT L3 cases are not supported when ICL interface is configured as router/untagged VE

it is required for all MPLS uplinks to be tagged interfaces to use FRR bypass for VLL/VPLS/L3VPN applications

Routing over VPLS

- pw-profile must be configured with tagged mode only under the bridge-domain instance for routing with VPLS.
- It is required for all MPLS uplinks to be tagged VE interfaces to support VEoVPLS.

Internet Routes Scaling

- It is recommended that the internet routes scaling features be enabled with internet peering configurations, as qualified by Extreme
- Feature is supported with default VRF only; default VRF and non-default VRF should not be co-existing when default VRF is configured with Internet routes scaling feature

L3VPN jumbo limitation

- The IPMTU value configured in CLI is applicable, if outgoing routing interface is an undelay IP interface (VE or L3 port); the IPMTU value configured in CLI is not applicable if the outgoing interface is uplink for IPoMPLS, L3VPN traffic, or ICL for MCT peers. Jumbo frames over MPLS/L3VPN tunnels can be accepted based the port L2MTU values.

EVPN IP Fabric

- IPv6 Static Anycast Gateway is not supported.

Storm-control

- Counters for Broadcast and Multicast storm-control are not supported in layer2-optimized-1 profile.

Increase scale support for class-maps under the service policy

- The ACL/VLAN/BD Rate Limiting scale numbers are dependent on tcam profile configured. Basically, based on the tcam entries reserved for the feature, user can scale number of policers/stats for appropriate application.

Consider below example with tcam profile "layer2-optimised-1".

- Create 2K Vlan/BD based class-maps and 2K ACL based class-maps associate those with policy-map pmap1.

- Configure 1k distinct policer attributes (cir/cbs/eir/ebs) for all the policy-map/class-map combination and bind the policy-map pmap1 to any interface.
- Now overall there will be 4K policers active for that interface with 4k distinct class-maps (match criteria).
- Note: The 4K policers (class-maps) scale will not be applicable to port-channel. There are only 1,215 policers are reserved for port-channels.
- Based on the requirement user must set the tcam profile and must reboot the box for activating the same.

Defects closed in 18r.1.00c

This section lists software defects with Critical, High, and Medium Technical Severity closed with and without a code change as of **6/19/2019** in 18r.1.00c.

Note: Parent Defect ID is the customer found Defect ID. The Issue ID is the tracking number uniquely used to check in the fix for each major release.

Parent Defect ID:	SLXOS-22514	Issue ID:	SLXOS-30535
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Traffic Management
Reported in Release:	SLXOS 17r.1.01a	Technology:	Rate Limiting and Shaping
Symptom:	100G interfaces on SLX 9850 may not achieve line rate egress throughput.		
Condition:	On a L2VPN network 100G interfaces on SLX 9850 may not achieve line rate of egress through put.		
Workaround:	Augment performance with additional interfaces as required.		

Parent Defect ID:	SLXOS-37463	Issue ID:	SLXOS-37463
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	SLXOS 18r.1.00a	Technology:	BGP4 - IPv4 Border Gateway Protocol
Symptom:	A warning is seen on console as follows. "No. of prefix received from BGP Peer 2000:31:1:8::153: exceeds warning limit 0"		
Condition:	When the maximum prefix config is at the IPv6 neighbor level and the ipv6 address-family activate cmd at the peer-group level and the device is reloaded with that saved config.		
Workaround:	Remove maximum prefix config at the neighbor and re-config.		

Parent Defect ID:	SLXOS-38394	Issue ID:	SLXOS-38394
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Security
Reported in Release:	SLXOS 18r.1.00b	Technology:	HTTP/HTTPS
Symptom:	Enabling or disabling HTTP(S) service on Virtual Routing and Forwarding(VRF) name other than management VRF("mgmt-vrf") may not succeed.		
Condition:	Enabling or disabling HTTP(S) service fails on chassis based devices		
Workaround:	None		

Parent Defect ID:	SLXOS-38406	Issue ID:	SLXOS-38406
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Traffic Management
Reported in Release:	SLXOS 18r.1.00b	Technology:	Rate Limiting and Shaping
Symptom:	When egress Rate-limiter is applied on port-channel, and system is rebooted, then egress Rate-Limiter was not working.		
Condition:	When system was rebooted with Egress RL applied on port-channel		
Workaround:	After reboot, reapply egress RL.		

Parent Defect ID:	SLXOS-38447	Issue ID:	SLXOS-38447
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	SLXOS 18r.1.00b	Technology:	BGP4 - IPv4 Border Gateway Protocol
Symptom:	L3VPN traffic may not be forwarded.		
Condition:	The VRFs are configured first in an order and then L3VPN config (route-target, route-distinguisher etc.) is done in a different order to those VRFs, so that the label allocation to VRFs does not happen in the order of creation of VRFs. Now if HA fail over is done, L3 VPN traffic may not be forwarded.		

Parent Defect ID:	SLXOS-27981	Issue ID:	SLXOS-38496
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 17r.1.01ah	Technology:	CLI - Command Line Interface
Symptom:	If user changes the startup config file to let management IP and default gateway in a different subnet from the management IP and default gateway that currently configured on SLX, and perform "copy tftp startup-config", then reload, SLX will keep the previously configured management IP and gateway after reload		
Condition:	User changes the startup config file to have management IP and gateway address in a different subset from the management IP and gateway that currently configured on SLX		
Workaround:	Before reload the system, remove the management IP and default gateway from system using CLI		

Parent Defect ID:	SLXOS-20017	Issue ID:	SLXOS-38877
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 17r.1.00a	Technology:	SNMP - Simple Network Management Protocol
Symptom:	Unexpected behavior with SLX		
Condition:	While trying to write port alias using SNMP application.		
Workaround:	None		

Parent Defect ID:	SLXOS-38980	Issue ID:	SLXOS-38980
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18r.1.00b	Technology:	SNMP - Simple Network Management Protocol
Symptom:	Snmp daemon terminates and restarts on HA failover.		
Condition:	Occurs only when there is an SNMP host configured with source-interface as management mm-ip.		
Workaround:	Keep source-interface configuration as default		

Parent Defect ID:	SLXOS-38299	Issue ID:	SLXOS-39007
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Other
Reported in Release:	SLXOS 18x.1.00a	Technology:	Other
Symptom:	Sometimes, a panic dump may be seen while rebooting the setup.		
Condition:	This is a rare condition which may be seen while device is rebooting or when sending high rate traffic to CPU.		
Workaround:	N/A		

Parent Defect ID:	SLXOS-39185	Issue ID:	SLXOS-39185
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Traffic Management
Reported in Release:	SLXOS 18r.1.00b	Technology:	Rate Limiting and Shaping
Symptom:	rate limiting clear command is not working		
Condition:	execute show command and clear command, data still shows even after clear command		

Parent Defect ID:	SLXOS-39214	Issue ID:	SLXOS-39214
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	MPLS
Reported in Release:	SLXOS 18r.1.00b	Technology:	LDP - Label Distribution Protocol
Symptom:	LDP sessions stay down despite ping functioning between the peers.		
Condition:	socket supporting LDP session is terminated by HA failover or route change. Note that the problem very rarely happens.		

Parent Defect ID:	SLXOS-39319	Issue ID:	SLXOS-39319
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Other
Reported in Release:	SLXOS 18r.1.00c	Technology:	Other
Symptom:	On the serial console, user observes output overwritten and becoming unreadable.		
Condition:	On the serial console, user changes window size and observe output overwritten and becoming unreadable.		

Parent Defect ID:	SLXOS-39220	Issue ID:	SLXOS-39349
Severity:	S1 - Critical		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18s.1.01a	Technology:	SNMP - Simple Network Management Protocol
Symptom:	LLDP-MIB::lldpLocPortId value is not correct (appears corrupted) when queried via SNMP GET operation.		
Condition:	Issue occurs only for SNMP GET operation (on LLDP-MIB::lldpLocPortId). SNMP GET-NEXT and snmpwalk returns correct values.		
Workaround:	<ol style="list-style-type: none"> 1. Use SNMP GET-NEXT or snmpwalk instead of SNMP GET when querying LLDP-MIB::lldpLocPortId via SNMP. 2. Use CLI to query (LLDP-MIB::lldpLocPortId) instead of SNMP, if it's feasible. 		

Parent Defect ID:	SLXOS-38901	Issue ID:	SLXOS-39427
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Network Automation and Orchestration
Reported in Release:	SLXOS 18r.1.00aa	Technology:	NETCONF - Network Configuration Protocol
Symptom:	Seeing error while using rpc to get port channel config.		
Condition:	If Insight enable mmlid 1 command is enabled.		

Parent Defect ID:	SLXOS-39445	Issue ID:	SLXOS-39445
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Security
Reported in Release:	SLXOS 18r.1.00c	Technology:	TACACS & TACACS+
Symptom:	DCM daemon termination will be observed while executing any CLI		
Condition:	When AAA Command Authorization is enabled and the configured tacacs+ server are not reachable.		
Workaround:	Make the configured Tacacs+ server reachable.		

Parent Defect ID:	SLXOS-38336	Issue ID:	SLXOS-39626
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18x.1.00a	Technology:	CLI - Command Line Interface
Symptom:	Overlay-gateway configuration doesn't show up in running-config.		
Condition:	Overlay-gateway configuration doesn't show up in running-config after firmware upgrade with ZTP (Zero touch provisioning),		
Workaround:	none		

Parent Defect ID:	SLXOS-38493	Issue ID:	SLXOS-39702
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Other
Reported in Release:	SLXOS 18r.1.00aa	Technology:	Other
Symptom:	file transfer may be affected if destination port is udp 646 and pkt has fragment offset.		
Condition:	During file transfer if destination port is udp 646 between source and destination without "mpls ldp" being enabled on the box. UDP packet with destination port 646 is trapped to cpu even without mpls being enabled on the box.		

Parent Defect ID:	SLXOS-34817	Issue ID:	SLXOS-39748
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLXOS 18r.2.00	Technology:	VRP - VLAN Registration Protocol
Symptom:	User will observe that the REST API response for mvrp interface related operational GET command fetches the value from the global bucket for non-MVRP interfaces.		
Condition:	Issue was seen only when trying to fetch MVRP information using REST API infrastructure for interfaces where MVRP was not configured.		
Workaround:	No		

Parent Defect ID:	SLXOS-39783	Issue ID:	SLXOS-39784
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLXOS 18r.1.00a	Technology:	OAM - Operations, Admin & Maintenance
Symptom:	nf_contrack feature in Linux kernel track all IP packets coming to CPU. It can cause nf_contrack table full issue & fragmented packet drop issue.		
Condition:	There is no specific condition trigger this, by default the service is up and running.		

Parent Defect ID:	SLXOS-29369	Issue ID:	SLXOS-39838
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLXOS 18r.2.00	Technology:	Port Mirroring
Symptom:	MTU of a destination mirror port may be a non-default MTU.		
Condition:	1. When Global MTU is configured on the device, 2. A port is configured as a destination-mirror port		

Parent Defect ID:	SLXOS-25731	Issue ID:	SLXOS-39974
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLXOS 17s.1.02b	Technology:	MCT - Multi-Chassis Trunking
Symptom:	MCT daemon termination followed by switch reload		
Condition:	MCT daemon terminates when client server sends the LACP oper key as 0.		
Workaround:	Remove 'esi auto lacp' config		

Parent Defect ID:	SLXOS-39618	Issue ID:	SLXOS-40059
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	MPLS
Reported in Release:	SLXOS 18r.1.00aa	Technology:	MPLS VPLS - Virtual Private LAN Services
Symptom:	Peers MPLS interface VE MAC learned as remote VPLS mac.		
Condition:	Issue seen in egress PE node, when receiving VPLS packet has inner payload DA MAC as 0100.5e00.xxxx		

Parent Defect ID:	SLXOS-40076	Issue ID:	SLXOS-40076
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	SDN
Reported in Release:	SLXOS 18r.1.00c	Technology:	OpenFlow
Symptom:	Openflow flows are not pushed in Openflow profile 3 (Tcam profile)		
Condition:	Issue will be seen in openflow-profile-3 tcam while pushing openflow flows		
Workaround:			

Parent Defect ID:	SLXOS-40143	Issue ID:	SLXOS-40143
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Security
Reported in Release:	SLXOS 18r.1.00c	Technology:	RADIUS
Symptom:	On console following error message will be seen - "Dcmd[3617]: pam_radius_auth: Could not open configuration file /etc/raddb/server: No such file or directory"		
Condition:	When REST/RESTCONF query is given		

Parent Defect ID:	SLXOS-40087	Issue ID:	SLXOS-40367
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18r.2.00	Technology:	High Availability
Symptom:	hasmd hang which was killed by SWD and switch reloaded in external login attach.		
Condition:	the issue may happen in brutal force login attack.		
Workaround:	None		

Parent Defect ID:	SLXOS-24384	Issue ID:	SLXOS-40383
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 17r.2.01	Technology:	CLI - Command Line Interface
Symptom:	ha chassisreboot command from ha inline help is not needed.		
Condition:	ha chassisreboot should have been obsoleted.		
Workaround:	Please use reload system command.		

Parent Defect ID:	SLXOS-24114	Issue ID:	SLXOS-40435
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLXOS 17r.2.00	Technology:	OAM - Operations, Admin & Maintenance
Symptom:	After the devices boots up, the user will see the dcmd.sh, ccmd.sh, and netstat defunct processes.		
Condition:	The defunct processes will show up when the user runs the "ps aux" command.		
Workaround:	None. They are not harmful and so the user can just ignore them.		

Parent Defect ID:	SLXOS-39058	Issue ID:	SLXOS-40466
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Other
Reported in Release:	SLXOS 18r.1.00a	Technology:	Other
Symptom:	Switch reloaded with panic dump, impacting the data traffic forwarding.		
Condition:	High rate of software assisted layer 3 forwarding of traffic, causing connection tracking table to fill up.		
Workaround:	NA		

Parent Defect ID:	SLXOS-40476	Issue ID:	SLXOS-40477
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18r.2.00	Technology:	High Availability
Symptom:	During DOS attacks, flood of disable pam_unix log messages are seen on console		
Condition:	DOS attacks on system		
Workaround:	Configure syslog server to redirect these messages		

Parent Defect ID:	SLXOS-38229	Issue ID:	SLXOS-40484
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18r.1.00ab	Technology:	LLDP - Link Layer Discovery Protocol
Symptom:	NOS CLI "show lldp neighbors" command failed to fetch the neighbor details.		
Condition:	LLDP must be configured on the SLX device.		
Workaround:	CMSH "show lldp neighbors" command can be used to fetch the LLDP neighbor details		

Parent Defect ID:	SLXOS-38762	Issue ID:	SLXOS-40532
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Other
Reported in Release:	SLXOS 18r.1.00aa	Technology:	Other
Symptom:	JSON output of REST bridge-domain config has duplicated URN part way through the output.		
Condition:	For vlans configured more than 100, REST bridge-domain config has duplicated URN part way through the output.		
Workaround:	No		

Parent Defect ID:	SLXOS-40884	Issue ID:	SLXOS-40884
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18r.1.00c	Technology:	Software Installation & Upgrade
Symptom:	Firmwaredownload may fail with dpkg confd error messages on console.		
Condition:	While upgrade/downgrade using normal firmwaredownload /fullinstall.		

Parent Defect ID:	SLXOS-39462	Issue ID:	SLXOS-40929
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLXOS 18r.1.00ac	Technology:	MCT - Multi-Chassis Trunking
Symptom:	MAC is not being updated with the new ifindex on the remote LC.		
Condition:	MAC move from one client interface to another client interface when same mac changes from CCR to CCL at the same time on remote LC.		

Parent Defect ID:	SLXOS-40994	Issue ID:	SLXOS-40994
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Other
Reported in Release:	SLXOS 18r.1.00c	Technology:	Other
Symptom:	/var/run directory not present at bootup so /var/run/racoon2 was not created.		
Condition:	/var/run and /var/run/racoon2 should be present at bootup time		
Workaround:	add mkdir /var/run and -p option for create /var/run/racoon2 in sysinit script		

Parent Defect ID:	SLXOS-39963	Issue ID:	SLXOS-41019
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLXOS 18r.1.00b	Technology:	sFlow
Symptom:	SFLOW controller failed to capture few flows on bi-directional traffic.		
Condition:	SFLOW configuration enabled on interface.		

Parent Defect ID:	SLXOS-39538	Issue ID:	SLXOS-41117
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Security
Reported in Release:	SLXOS 18r.1.00aa	Technology:	AAA - Authentication, Authorization, and Accounting
Symptom:	Unexpected reload		
Condition:	When TACACS authorization fails on re-try		
Workaround:	Make sure we have the proper network connectivity to avoid TACACS authorization fails at first attempt.		

Parent Defect ID:	SLXOS-41166	Issue ID:	SLXOS-41168
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18r.1.00b	Technology:	CLI - Command Line Interface
Symptom:	Unexpected reload of the device.		
Condition:	Protocol lldp has dot1-tlv/dot3-tlv config and when "show lldp neighbors detail" command is issued.		
Workaround:	None		

Parent Defect ID:	SLXOS-40759	Issue ID:	SLXOS-41327
Severity:	S1 - Critical		
Product:	SLX-OS	Technology Group:	MPLS
Reported in Release:	SLXOS 18r.1.00aa	Technology:	MPLS VPLS - Virtual Private LAN Services
Symptom:	Not able to program MPLS tunnel		
Condition:	Power-off/on line card on PE router		

Parent Defect ID:	SLXOS-40826	Issue ID:	SLXOS-41816
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Other
Reported in Release:	SLXOS 18r.1.00aa	Technology:	Other
Symptom:	SLX device experience unexpected sudden reload.		
Condition:	FWD daemon termination cause the sudden reload.		