



June 2020

SLX-OS 18s.1.03a for SLX 9140 and SLX 9240 Release Notes

Copyright Statement and Legal Notices

Copyright © 2020 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

Copyright Statement and Legal Notices	2
Document History.....	4
Preface	5
New Software Features.....	7
Supported Optics	17
Software Upgrade and Downgrade	20
Limitations and Restrictions.....	27
Defects	34

Document History

Version	Summary of changes	Publication date
1.0	Initial Release	June 2020
1.1	Updated Upgrade Matrix	June 2020

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 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 <http://www.extremenetworks.com/documentation-feedback-pdf/>
- Email us at internalinfodev@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.

New Software Features

This section addresses features introduced in the current release as well as those introduced in the previous release.

SLX-OS 18s.1.03a New features summary

- Support for **VXLAN Transit mode**
- Support for **VXLAN Native IP mode**
- Support for a custom service attribute for TACACS+
- New CLI commands
 - `attribute-list service-attribute`
 - `mode native-ip`
 - `mode vxlan-transit`
- Modified commands
 - `ipv6 access-list`
 - `uda-key profile`
 - `uda-profile-apply`

SLX-OS 18s.1.03a New features detail

VXLAN Transit Mode

An interface configured in Transit mode filters based on VXLAN inner headers, without stripping the outer headers.

By default, tunneled flows are filtered based on inner headers only when you configure the stripping of outer headers on the ingress interface. Otherwise, filtering is based on the outer headers. In VXLAN Transit mode, Layer 2 and Layer 3 policies are always applied on inner headers, irrespective of whether you have enabled outer header stripping.

If the VXLAN traffic that ingresses on a Transit interface has an inner IPv6 header, the Layer 3 and UDA policies can filter and match on the full 128-bit Source and Destination IP addresses. On regular interfaces, filtering is limited based on lower or upper 64-bits of SIP and DIP.

Non-VXLAN traffic that ingresses on an interface in VXLAN Transit mode is treated as an Ethernet frame with an unknown Layer 3 Protocol. If necessary, you can filter these frames using Layer 2 or UDA policies.

You can use the **profile npb-parser** command to select the Default or VXLAN Transit profile.

You can use the **mode vxlan-transit** command to enable VXLAN Transit mode on an interface.

VXLAN traffic on a Transit interface is parsed differently from VXLAN traffic on a regular interface. The outer headers are not copied into the Token Header Buffer and are not available for policy lookups, including the UDA. The UDA profile on a VXLAN Transit interface does not include any outer headers and is defined to match the inner and payload headers.

For example, here is the UDA profile definition on a regular interface:

```
device(config)# uda-key profile UDA_PROFILE1
device(config-uda-key)# flow header0 ETHERNET header1 IPV4 header2 UDP header3 VXLAN header
TUN_ETHERNET header5 IPV4 header6 TCP header7 PAYLOAD4
```

On a Transit interface, the UDA profile is defined as follows:

```
device(config)# uda-key profile UDA_PROFILE2
device(config-uda-key)# flow header0 ETHERNET header1 IPV4 header2 TCP header3 PAYLOAD32
```

VXLAN Native IP Mode

An interface configured in Native (non-tunnel) IP mode parses only the native headers (outer Layer 2, Layer 3, and Layer 4). Parsing stops after the outer Layer 4 header and tunneled headers (if any) are not parsed.

For IPv6 flows, Native IP mode also supports matching and filtering based on full 128-bit Source and Destination addresses. Because tunneled headers are not parsed in Native IP mode, the Token Header Buffer always has sufficient space to copy the full IPv6 SIP and DIP.

You can use the **mode native-ip** command to enable Native mode on an interface.

UDA behaves the same way in Native IP mode as it does on a regular interface, except when you are trying to match an IPv6 flow. Use the header type **IPV6_128** to match an IPv6 flow. For example:

```
device(config)# uda-key profile UDA_PROFILE1
device(config-uda-key)# flow header0 ETHERNET header1 IPV6_128 header2 TCP header3 PAYLOAD32
```

Custom Service Attribute for TACACS+

With the **service-attribute** option for the **tacacs-server host attribute-list** command, administrators can configure a service attribute through the CLI. If a service-attribute mismatch occurs between a device and a server, or if the brcd-role is not configured on the TACACS+server, then the following actions are taken based on the mode:

- NPB mode: Users are denied access to the device. An audit log or RASlog will indicate that authorization failed due to a service-attribute mismatch or the lack of configuration of the brcd-role. Users can access the device only in a User role.
- DC mode: Users are given access to the device, but only in a User role.

IPv6 ACL Configuration

The **ipv6 access-list** command is enhanced with new options to support the matching and masking of any set of bits in the SIP or DIP.

Use the **prefix_len** option to define the number of high-order, contiguous bits of an address that comprise the IPv6 prefix to match.

Use the **mask** option to specify the bits to mask, not match. Specify a 128 bitmap in the format XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX:XXXX.

SLX-OS 18s.1.03 New features summary

SLX-OS 18s.1.0.3 is the fifth release in a series for SLX Switching platforms. SLX 9140 and SLX 9240 are the target platforms for this release and are mainly focused on the Network Packet Broker (NPB) and Datacenter solutions. No new hardware platform is added in this release, and only software features are added.

NOTE. This document includes information that is supported in previous release.

The key features for SLX-OS 18s.1.03 are focused on NPB and Datacenter Solution features enhancing manageability, user experience on SLX.

The new features are as follows:

Data Center features:

TACACS+ AAA Command Authorization:

With the introduction of AAA Command Authorization feature, authorization request will now be sent to configured TACACS+ server when TACACS+ command authorization is configured. Execution of a particular command for a particular user will be allowed or denied based on the accept and deny rule configured on the TACACS+ server for that user.

DHCP Relay Source Interface Configuration:

This feature mainly used in SAG environment where SAG will act as DHCP relay agent address. The path from SAG to DHCP server works without any problem. Also, path from DHCP server to the SAG exists but the response may arrive at a different switch than the original one as the SAG is inherently distributed across many switches. In this environment, the unique loopback address will act as Gateway IP Address to forward response back to the client. The link selection sub-option will decide which subnet is correlated to the DHCP request.

Extreme Cloud Connect:

ZTP+ is a mechanism through which the device interacts with Extreme Management Center (XMC) by installing a cloud connector (CC) plugin on the device. The CC will reach out to XMC to notify that it is now on the network and will go through several states validating the version of software, configuration, configuration changes and configuration and status updates.

App telemetry on SLX 9140:

This feature helps to extract network analytics, for example, application name, flow pathing, bandwidth, and latency from Extreme Networks SLX switch platforms. It uses the sFlow and the Encapsulated Remote Switched Port Analyzer (ERSPAN) protocols to extract and transport specific raw packets and sampled flows from the SLX-OS switches to Extreme Analytics processing engines for further analysis.

NPB features:

Load Balancing Support for Logical NPB Grid:

This feature allows load balancing among a set of destinations. We introduce a new construct called the load balance destination group which has a set of destinations. Depending on the frame hash, the frame is sprayed to one of the members of such a load balance group. Each load balance group

has one or more destinations. A policy result is a PBF destination group which can comprise of a one or more destination load balance groups and/or destinations.

Egress Packet Truncation:

The egress packet-truncation feature enables you to truncate source packets to a certain number of octets and forward it to a particular destination. Truncation is performed at the 'flow' level. The truncation functionality is achieved by defining a truncation profile, which contains a user-specified interface and truncation size and assigning it to a route-map. The truncation profile determines the final length of the egress frames on the selected flows.

SLX-OS 18s.1.03 New features details

TACACS+ AAA Command Authorization

Prior to this release, authorization was enforced by the Extreme's role-based access control (RBAC) protocol at the device level. With the introduction of AAA Command Authorization feature, authorization request will now be sent to configured TACACS+ server when TACACS+ command authorization is configured by running *'aaa authorization command tacacs+'*.

A few key points regarding this feature are:

- Authorization will be done for all the users including – 'admin'
- If TACACS + server doesn't find the user configuration for the user executing the CLI (in its configuration file), then it will reject all the commands for that particular user (even 'admin').
- To avoid the above situation, user with name 'DEFAULT' can be created on the TACACS+ server.
- Execution of every command will pass through authorization process, when AAA authorization is enabled.
- TACACS+ authorization is disabled by default.
- Along with TACACS+ config, user can also configure 'local' option so that when configured TACACS + servers are not reachable, execution of a command can happen based on local RBAC rules.
- If the 'local' option (after *tacacs+*) is not selected while configuring *'aaa authorization'* and if all the configured TACACS+ servers are not reachable then, execution of all commands will fail.

In order to recover from this situation, a fallback approach has been implemented. In this case, only 'admin' user will be allowed to execute only *'aaa authorization'* command so that the 'admin' user can unconfigure *'aaa authorization'* by running – *'aaa authorization command none'* or can select the 'local' option by running – *'aaa authorization command tacacs+ local'*.

DHCP Relay Source Interface Configuration

In modern DCs SAGs are of great value helping to reduce the amount of IP addresses needed. DHCP helps to reduce the time to deliver resources and makes the environment more flexible. In this context using the SAG as an DHCP relay agent should be possible. In the current implementation in SLX-OS the direction from the SAG to the DHCP server works without problems. And also the path back from the DHCP server to the SAG exists. But by its nature the SAG is distributed over some or many TOR switches and the response from the DHCP server might arrive at a different switch than the original sender.

In this scenario using unique loopback address as GIADDR along with Relay Agent Information Option (Option82) can be useful. Client connected leaf node can use unique loopback IP as GIADDR and can add Option82 with circuit-ID, remote-ID and link-selection sub-options. Link-selection Option82 sub-option can have the client connected network address (the SAG IP network). If DHCP server supports Option-82 link selection sub-option, then the server can be configured to allocate IP address based on link selection sub-option. If not, then server can be configured to allocate based on circuit-ID/remote-ID sub-options. DHCP server reply with Gateway IP Address as DIP which is unique loopback address of the relay agent.

Few key points regarding this feature are:

- The link selection sub-option takes on the normal role of the Gateway IP Address in relaying to the DHCP server which subnet is correlated to the DHCP request.
- When using this sub-option, the Gateway IP Address continues to be present but only relays the IP address that is to be used by the DHCP server to communicate with.
- Global Option-82 should be enabled along with gateway address configuration in relay agent.
- DHCP server should be configured to support GIADDR and Option-82 link-selection sub-option.
- Consumes unique IP address per node.
- RFC 3527 supports only DHCPv4 relay, so this feature will not be supported for IPv6.
- On downgrade feature won't be supported. [Extreme Cloud Connect](#)

ZTP+ or Enhanced ZTP is a mechanism through which the device interacts with Extreme Management Center (XMC) by installing a cloud connector (CC) plugin on the device. The CC will reach out to XMC to notify that it is now on the network and will go through several states validating the version of software, configuration, configuration changes and configuration and status updates. By having the CC initiate communication with XMC, it will also support XMC in the cloud and allow for access behind a company's firewall. ZTP+ uses HTTPS to provide secure communications between XMC and the device.

Few key points regarding this feature are:

- To support ZTP+, SLX devices should launch Cloud Connector process in ZTP mode.
- Cloud Connector, will be run one-time completion mode i.e. the CC will not run in persistent mode, rather it will bail out on completion of ZTP+ state machine.
- Since SLX OS has native ZTP support, both native ZTP and ZTP+ functionalities are available. It is up to the user choice to configure appropriate.
- Current ZTP+ support is only with XMC and not cloud based management center.
- Only out-of-band management is supported in this release.
- Configuration supported via ZTP+
 - Image upgrade (This requires XMC configuration, prior to bring up of devices) via SCP, FTP
 - Static Management IP
 - Gateway
 - DNS
 - Host name
 - SNMP v3 configuration
 - NTP server (time zone is not supported due to mismatch in XMC and Switch configuration)

App telemetry on 9140

The primary purpose of Application Telemetry is to extract network analytics, for example, application name, flow pathing, bandwidth, and latency from Extreme Networks SLX switch platforms.

This feature is supported only on the SLX 9140. You can use either sFlow or Application telemetry or both at the same time, as they can co-exist on a switch. sFlow must be enabled to use Application Telemetry.

Application Telemetry uses the sFlow and the Encapsulated Remote Switched Port Analyzer (ERSPAN) protocols to extract and transport specific raw packets and sampled flows from the SLX-OS switches to Extreme Analytics processing engines for further analysis. When a switch is added as a telemetry source, an Extreme Management Center (XMC) server runs a Tcl script that configures the switch automatically. Manual configuration is also supported.

The first set of raw traffic information is produced by highly specific ingress ACLs (processed within the SLX-OS hardware); these ACLs are applied at the system level (on all interfaces) of an SLX-OS switch to match specific packet types (for example, TCP SYN or DNS packets) with the purpose of mirroring these packets to the Extreme Analytics engine for further analysis. The ACL-filtered traffic is encapsulated and transported by means of the ERSPAN protocol towards Analytics Engines where, as with the ACLs, processing is handled by the SLX-OS hardware.

The second set of raw traffic is generated by the standard sFlow protocol and is enabled on all SLX-OS interfaces. By its very nature, sFlow is a sampled packet technology that is processed within the SLX-OS CPU. The sampled traffic is then transported over UDP to the analytics engine for further analysis.

The analytics engine processes the ERSPAN flows to extract application details, network flows, network response time for TCP-based flows, application response time for HTTP, HTTPS (SSL), DNS, DHCP, and so on. The sFlow information is used to deduce the bandwidth calculations of the individual flows and applications.

The result is that the application name, network response time, and application response time extracted from the ERSPAN mirrored traffic provides the basic Application Telemetry flow. When a sampled sFlow is matched to a basic flow, an enhanced Application Telemetry flow is produced that contains packet and byte counters, along with the details of a network to a switch interface.

A few key points regarding this feature are:

- The feature is enabled and disabled at the global level.
- A telemetry policy file is copied to the switch, by means of an XMC server script, over TFTP.
- SLX-OS saves 133 Application Telemetry filters in a telemetry.pol file.
- The IP address of the sFlow agent (the management IP address of the switch) and the IP address of the first collector in the default VRF acts as the source and destination for Generic Routing Encapsulation (GRE) configuration.
- Only one sFlow collector is supported for this feature. If multiple collectors are configured, the first collector configured with the default VRF is selected. This feature supports only the first sFlow IPv4 collector with the default VRF.
- If no sFlow collector is configured with the default VRF, an error is returned when the feature is enabled.
- The feature uses three new TCAM profiles, app-tele-l2-l3-iacl, app-tele-l3-iqos-l2-iacl and app-tele-l3-iqos-l3-iacl to optimize hardware resources.

NPB feature - Egress Packet Truncation

The egress packet-truncation feature enables you to truncate source packets to a certain number of octets and forward it to a particular destination. Truncation is performed at the 'flow' level. The truncation functionality is achieved by defining a truncation profile, which contains a user-specified interface and truncation size, and assigning it to a route-map. The truncation profile determines the final length of the egress frames on the selected flows.

To reference an interface in a truncation profile, it must first be placed into loopback mode. Only four truncation profiles can be created and referenced in route maps. A route-map or route-map stanza can reference a single truncation profile any number of times.

NOTE

On supported platforms, extra 10 bytes are appended to the frame after truncation. These 10 bytes includes the 4 byte FCS. For example, if you set the truncation size as 256 bytes, the frame size on the wire will be 266.

A few key points regarding this feature are:

- New CLI for creating a Truncation Profile.
- The interface configured under the profile will be configured as a loopback interface (along with the required settings to achieve truncation).
- An interface cannot be associated with more than one Truncation Profile.
- The interface cannot be used for any other purpose and it cannot have any config (route-map, strip-*, ...).
- Upto 4 Truncation Profiles are supported in this release.

NPB feature - Load Balancing Support for Logical NPB Grid

This feature allows load balancing among a set of destinations. We introduce a new construct called the load balance destination group which has a set of destinations. Depending on the frame hash, the frame is sprayed to one of the members of such a load balance group. Each load balance group has one or more destinations. The following are the changes.

- New CLI introduced to create load balance group, add, modify and remove members.
- New option to the destination group CLI. The destination group used to have a set of destinations. Now, we allow a destination group to have a set of destinations and/or destination-load balance groups.
- Show cli for destination load balance groups.
- A maximum of 31 members are supported for a given destination load balance group.

Consolidated Features in SLX-OS 18s.1.03

The following table lists the features introduced since SLX-OS 18s.1.00.

NPB Mode Features

<i>Header Stripping</i>	
<ul style="list-style-type: none"> • 802.1BR • VN-Tag • MPLS Label (EoMPLS & IPoMPLS) • GTP -U-v1 • VXLAN Encap • ERSPAN-II • NVGRE Encap 	<ul style="list-style-type: none"> • Per port support of header stripping, enabled or disabled via CLI • Tag stripping: 802.1BR or VN-ag (either one is supported) • Tunnel encapsulations stripping VXLAN, NVGRE, ERSPAN-II/GTP-U-v1/MPLS • Filter traffic using policy engine, based on values of fields in the tags/encapsulations in addition to standard L2/L3/L4 fields (outer and/or inner) • Multiple stripping configurations per port.
<i>Transparent VLAN</i>	
<ul style="list-style-type: none"> • Aggregation • Replication • VLAN filtering • VLAN tag add • VLAN tag delete • Combination of VLAN delete and VLAN add with header stripping. • Max TVF domains 	<ul style="list-style-type: none"> • Aggregate flows from multiple taps to a single egress interface. • Replicate flows from a single tap to multiple egress interfaces • Filter flows from tap to forward or drop based on route map policies • Outermost VLAN tag in the forwarded frame will be deleted • New VLAN tag will be added in standard canonical format • Route maps to be applied on ports or port-channels. • Maximum supported TVF domains is 4096

<i>Flex ACLs</i>	
<ul style="list-style-type: none"> • Super ACL capability • Limited deep packet inspection (DPI) 	<ul style="list-style-type: none"> • Deep packet inspection of tunneled traffic to filter specific flows, especially traffic that cannot be filtered using standard or extended MAC/IP ACLs. • Uses Flex ACLs (new for SLX 9240 and SLX 9140 platforms). Dictionary format CLI • Super ACL capability for traffic (tunneled or not) to match packet fields spanning across well-known layers.
<i>Scale Improvements</i>	
<ul style="list-style-type: none"> • L3 • L2 • Flex • Per Core (2 core per switch) 	<ul style="list-style-type: none"> • IP policy-based forwarding entries (IPACL): 2048 (IPV4+IPV6) • MAC policy-based forwarding entries (L2ACL): 4000 • Flex policy-based forwarding entries (Flex): 1024 • Ports per LAG: 64 • TVF domains: 4096
<i>VLAN</i>	<ul style="list-style-type: none"> • 400 VLANs

<i>NPB enhancements</i>	
<ul style="list-style-type: none"> • Support for LLDP in NPB mode • Reporting packet drop counts along packet forwarding path • Support for Telemetry streaming profiles, including LLDP link status and neighbor info • Logical NPB Grid for Forte (NSH Tagging for node identification in fabric) 	<ul style="list-style-type: none"> • NPB grid load balancing(new) • Egress packet truncations(new)
<ul style="list-style-type: none"> • Onboard packet capture • Internal loopback support 	<ul style="list-style-type: none"> • Onboard packet capture - capture ingress/egress data frames in PCAP format for a given port in NPB mode only, one port at a time. Auto stop after capturing designated number of frames • Internal loopback - service chaining in NPB operations. Deep packet header inspection.

<i>New Optics Qualified</i>	
<ul style="list-style-type: none"> • 25G SFP28 LR • 40G Bi Di media • 10GBASE-T SFP+ • 100G DAC cable support for 5 m reach 	<ul style="list-style-type: none"> • 10504 • 40G Bidirectional media support • 10338 • 100G-QSFP-QSFP-P-0501

<i>Misc Features</i>	
Port Breakout Support	Support for 4x25G
Dynamic Breakout Support	Eliminates the need to reload the system when breakout or non-breakout on ports.

Supported Optics

For a list of supported fiber-optic transceivers that are available from Extreme, refer to the latest version of the Extreme Optics Family Data Sheet available online at www.extremenetworks.com.

Description	Orderable PN	P/N
1000Base-SX	1G-SFP-SX-OM	33210-100
1000Base-LX	1G-SFP-LX-OM	33211-100
1GE Copper SFP (Pseudo-Branded)	1G-SFP-TX	33002-100
1GE Copper SFP (BR-Branded)	1G-SFP-000190	57-1000042-02
10GE USR SFP+	10G-SFP-USR	57-1000130-01
10GE USR SFP+, 70C TAA	10G-SFP-USR-SA	57-1000343-01
10GE SR SFP+, 85C	10G-SFP-SR	57-0000075-01
10GE SR SFP+, 70C	10G-SFP-SR-S	57-1000340-01
10GE SR SFP+, 70C TAA	10G-SFP-SR-SA	57-1000344-01
10GE LR SFP+, 85C	10G-SFP-LR	57-0000076-01
10GE LR SFP+, 70C	10G-SFP-LR-S	57-1000341-01
10GE LR SFP+, 70C TAA	10G-SFP-LR-SA	57-1000345-01
10GE AOC 7M	10GE-SFP-AOC-0701	57-1000273-01
10GE AOC 10M	10GE-SFP-AOC-1001	57-1000274-01
10GE Direct Attach 5M Active	10G-SFP-TWX-0501	58-1000023-01
10GE Direct Attach 1M Active	10G-SFP-TWX-0101	58-1000026-01
10GE Direct Attach 3M Passive	10G-SFP-TWX-P-0301	58-1000025-01
Description	Orderable PN	P/N
10GE Direct Attach 5M Passive	10G-SFP-TWX-P-0501	58-1000019-01
25G SR	25G-SFP28-SR	57-1000342-01
25GE Direct Attach 01M Passive	25G-SFP28-TWX-P-0101	58-0000064-01

25GE Direct Attach 03M Passive	25G-SFP28-TWX-P-0301	58-0000065-01
40GE QSFP+ SR4	40G-QSFP-SR4-1	57-1000128-01
4x10GE QSFP+ LR4, 10km,	40G-QSFP-LR4-INT	57-1000477-01
40GE BiDi QSFP+	40G-QSFP-SR-BIDI	57-1000339-01
40GE QSFP+ LR4, 10KM, 70C	40G-QSFP-LR4-1	57-1000263-01
40GE QSFP+ SR4 to 10G-SR SFP+	40G-QSFP-SR4-INT	57-1000129-01
40GE QSFP to QSFP 1M Cable(Passive)	40G-QSFP-C-0101	58-0000033-01
40GE QSFP to QSFP 3M Cable(Passive)	40G-QSFP-C-0301	58-0000034-01
40GE QSFP to QSFP 5M Cable(Passive)	40G-QSFP-C-0501	58-0000035-01
4x10GE QSFP+ to 4 SFP+ Active copper cable - 1m	40G-QSFP-4SFP-C-0101	58-0000051-01
4x10GE QSFP+ to 4 SFP+ Active copper cable - 3m	40G-QSFP-4SFP-C-0301	58-0000052-01
4x10GE QSFP+ to 4 SFP+ Active copper cable - 5m	40G-QSFP-4SFP-C-0501	58-0000053-01
40GE QSFP to QSFP cable - 10m AOC	40G-QSFP-QSFP-AOC-1001	57-1000306-01
100GE QSFP28 SR4	100G-QSFP28-SR4	57-1000326-01
100GE QSFP28 LR4 (3.5W)	100G-QSFP28-LR4-LP-10KM	57-1000338-01
100GE QSFP28 CWDM	100G-QSFP28-CWDM4-2KM	57-1000336-01
100G QSFP28 Active Optical (10m)	100G-QSFP-QSFP-AOC-1001	57-1000347-01
Description	Orderable PN	P/N
100GE QSFP28 LRL 2km	100G-QSFP28-LR4L-2KM	57-1000329-01

Note: 10GE LR SFP+, 85C multi speed optic can operate on 10G as well as 1G.

New optics supported starting with SLX18s.1.01

Description	Orderable PN	P/N
25G SFP28 LR (10km), Single Mode, LC-connector, 70degC	25G-SFP28-LR	10504

Supported Mellanox 10G optics:

- 10G USR SFP+
- 10G SR SFP+
- 10G LR SFP+

DAC cables:

- 40G-QSFP-QSFP-P-0X01: passive 40G direct attached copper cables (X = 1, 3, 5m reach)
- 40G-QSFP-QSFP-C-0X01: active 40G direct attached copper cables (X = 1, 3, 5m reach)
- 40G-QSFP-4SFP-C-0X01: active 40G direct attached breakout copper cables (X = 1, 3, 5m reach)
- 100G-QSFP-QSFP-P-0101: 100GE Direct Attached QSFP-28 to QSFP-28 Passive Copper cable, 1m
- 100G-QSFP-QSFP-P-0301: 100GE Direct Attached QSFP-28 to QSFP-28 Passive Copper cable, 3m

Software Upgrade and Downgrade

This section includes information that supports both the current and previous release.

SLX-OS 18s.1.03a

Image file names

Download the following images from www.extremenetworks.com.

Image file name	Description
slxos18s.1.03a.tar.gz	SLX-OS 18s.1.03a_ software
slxos18s.1.03a_all_mibs.tar.gz	SLX-OS 18s.1.03a_ MIBS
slxos18s.1.03a.md5	SLX-OS md5 checksum
tpvm2.1.0.tar.gz	TPVM image

To Install SLX-OS 18s.1.03 from the network:

Run command: **firmware download scp host** *<ip-address>* *<directory>*

Where: *<directory>* is where the image is downloaded.

To Install SLX-OS 18s.1.03 from a USB device, follow the steps below:

Step 1: Copy unzipped SLX-OS firmware to the USB device under the firmware directory.

- For upgrade from releases through 18s.1.01, the directory structure is */brocade/firmware/<build>*.
- For upgrade from 18s.1.02, the directory structure is */slxos/firmware/<build>*.

Step 2: Plug the USB device into the switch on which you want to download the firmware.

Step 3: Execute the **usb on** command from the CLI prompt.

Step 4: Execute the following: **firmware download usb** *<full path of the firmware>*

TPVM

This section addresses upgrading and downgrading TPVM across releases in this series.

Upgrade and downgrade procedures have changed. Refer to “TPVM package upgrade and downgrade between 18s1.03 and 18s1.01” below.

Package support matrix

Release Name	TPVM package name	TPVM package location
slxos17s.1.xx	vm-swbd2900-1.0.0-1.i386.deb	<releaseserver>/slxoss/slxos17s.1.xx/dist/SWBD2900/vm-swbd2900-1.0.0-1.i386.deb
slxos18s.1.01	vm-swbd2900-1.0.0-1.i386.deb	<releaseserver>/slxoss/slxos18s.1.01/dist/SWBD2900/vm-swbd2900-1.0.0-1.i386.deb
slxos18s.1.02 slxos18s.1.03	tpvm-2.1.0-1.i386.deb	<releaseserver>/slxoss/slxos18s.1.xx/dist/SWBD2900/tpvm-2.1.0-1.i386.deb

TPVM package between 18s1.02 and 18s1.03

TPVM package between 18s1.02 and 18s1.03 is compatible, if TPVM package was installed before upgrade or downgrade, then after upgrade or downgrade, user can use “tpvm start” to start TPVM.

TPVM package upgrade and downgrade between 18s1.01 and 18s1.03

TPVM Package upgrade from slxos18s.1.01 to slxos18s.1.03

1. First, uninstall the existing TPVM package using following SLX-OS CLI
tpvm uninstall
2. Upgrade device with slxos18s.1.03 release using firmware download command.
3. Remove existing TPVM package located at following path in device
SLX-OS VM using linux shell login prompt
rm -rf /tftpboot/SWBD2900/vm-swbd2900-*.deb
rm -rf /mnt/tftpboot/SWBD2900/vm-swbd2900-*.deb
4. scp/ftp following TPVM package from release/build server
<TPVM release url>/SWBD2900/tpvm-2.1.0-1.i386.deb to device following directory on device
/tftpboot/SWBD2900/
5. Install new TPVM package using following SLX-OS CLI
tpvm install
6. Use following SLX-OS CLI to check TPVM install status and start TPVM
show tpvm status
tpvm start

TPVM Package downgrade from slxos18s.1.03 to slxos18s.1.01

1. Uninstall the existing TPVM package using following SLX-OS CLI
tpvm uninstall
2. Remove existing TPVM package located at following path in device
SLX-OS VM using linux shell login prompt
rm -rf /tftpboot/SWBD2900/tpvm-*.deb
3. Upgrade device with slxos18s.1.01 release using firmware download command.
4. Download the SLX-OS firmware
5. Install new TPVM package using following SLX-OS CLI
tpvm install
6. Use following SLX-OS CLI to check TPVM install status and start TPVM
show tpvm status
tpvm start

Migration path

Default Mode: Recommended upgrade/downgrade migration paths

To	SLX 17s.1.00a	SLX 17s.1.01	SLX 17s.1.02	SLX 17s.1.02x	SLX 18s.1.00	SLX 18s.1.01	SLX 18s.1.01x	SLX 18s.1.03	SLX 18s.1.03a
From									
SLX 17s.1.00a	NA	FWDL coldboot	FWDL coldboot	FWDL coldboot	*	*	*	*	*
SLX 17s.1.01	Default – config	NA	FWDL coldboot	Default-config	*	*	*	*	*
SLX 17s.1.02	Default – config	FWDL coldboot	NA	FWDL coldboot	FWDL coldboot	FWDL coldboot	FWDL coldboot	*	*
SLX 17s.1.02x	Default – config	Default-config	FWDL coldboot	NA	FWDL coldboot	FWDL coldboot	FWDL coldboot	*	*
SLX 18s.1.00	*	*	Default - config	Default - config	NA	FWDL coldboot	FWDL coldboot	*	*
SLX 18s.1.01	*	*	Default - config	Default - config	FWDL coldboot	NA	FWDL coldboot	FWDL coldboot	FWDL coldboot
SLX 18s.1.01x	*	*	Default - config	Default - config	FWDL coldboot	FWDL coldboot	NA	FWDL coldboot	FWDL coldboot
SLX 18s.1.03	*	*	*	*	*	FWDL coldboot	FWDL coldboot	NA	FWDL coldboot
SLX 18s.1.03a	*	*	*	*	*	FWDL Coldboot	FWDL coldboot	FWDL coldboot	NA

***NOTE:** For SLX 17s.1.00/a/1, the recommended path is first to install the SLX17s.1.02x release, and then the SLX 18s.1.01 release. For an MCT cluster, it recommended that only one node be upgraded at a time. Wait for the first node to come up completely before upgrading the second node.

NPB Mode: Recommended upgrade/downgrade migration paths

To From	17s.1.00	17s.1.00a	17s.1.01	17s.1.02	17s.1.02x	18s.1.00	18s.1.01	18s.1.02	18s.1.03	18s.1.03a
17s.1.00	NA	FWDL coldboot	FWDL coldboot	FWDL coldboot	FWDL coldboot	*	*	*	*	*
17s.1.00a	FWDL-coldboot	NA	FWDL coldboot	FWDL coldboot	FWDL coldboot	*	*	*	*	*
17s.1.01	Default – config	Default – config	NA	FWDL coldboot	Default-config	*	*	*	*	*
17s.1.02	Default – config	Default – config	FWDL coldboot	NA	FWDL coldboot	FWDL coldboot	FWDL coldboot	FWDL coldboot	*	*
17s.1.02x	Default – config	Default – config	Default-config	FWDL coldboot	NA	FWDL coldboot	FWDL coldboot	FWDL coldboot	FWDL coldboot	*
18s.1.00	*	*	*	Default – config	Default - config	NA	FWDL coldboot	FWDL coldboot	FWDL coldboot	FWDL coldboot
18s.1.01	*	*	*	Default – config	Default - config	FWDL coldboot	NA	FWDL coldboot	FWDL coldboot	FWDL coldboot
18s.1.02	*	*	*	Default – config	Default - config	FWDL coldboot	FWDL coldboot	NA	FWDL coldboot	FWDL coldboot
18s.1.03	*	*	*	*	Default - config	FWDL coldboot	FWDL coldboot	FWDL coldboot	NA	FWDL coldboot
18s.1.03a	*	*	*	*	*	*	FWDL Coldboot	FWDL coldboot	FWDL coldboot	NA

***NOTE:** For SLX 17s.1.00/a/1, the recommended path is first to install the SLX17s.1.02x release, and then the SLX 18s.1.02 release.

Recommendations

Upgrading EFA over TPVM from 17s.1.0x to 18s.1.03

Direct upgrade from 17s.1.0x is not supported because of linux version upgrade and TPVM decoupling. Follow the steps below:

1. Perform upgrade from 17s.1.0x to 18s.1.01x
2. Perform upgrade from 18s.1.01x to 18s.1.03

Upgrading EFA over TPVM from 18s.1.01a/b/c to 18s.1.03 with Ubuntu version 16.04

Do the following to upgrade the EFA over TPVM application.

1. Log in to TPVM with the TPVM IP address.
\$ ssh -l root <TPVM_IP>
2. Copy the EFA over TPVM database and logs backup to an external server.
\$ service efa-server stop
\$ scp /var/efa/efa.db <server_DB_Location>
\$ scp /var/log/efa/efa.log <server_log_Location>
3. Stop and uninstall TPVM.
\$ tpvm stop
\$ tpvm uninstall
4. Upgrade the device to 18s.1.03
<Switch># start-shell
Entering Linux shell for the user: admin
[admin@<Switch>]#
[admin@<Switch>]# su
Password: <password>
[root@<Switch>]#
5. Remove existing TPVM package located at following path in device
SLX-OS VM using linux shell login prompt
rm -rf /tftpboot/SWBD2900/vm-swbd2900-*.deb
rm -rf /mnt/tftpboot/SWBD2900/vm-swbd2900-*.deb
6. scp/ftp following TPVM package from release/build server
<TPVM release url>/ SWBD2900/ tpvm-2.1.0-1.i386.deb to device following directory on device
/tftpboot/SWBD2900/

7. With the device upgraded to 18s.1.03, execute the **efa deploy** command.

During the execution of **efa deploy** command, the below warnings will be seen. These can be ignored, as they are deemed harmless.

/dev/mapper/nbd0p1 not set up by udev: Falling back to direct node creation.

/dev/mapper/nbd0p2 not set up by udev: Falling back to direct node creation.

/dev/mapper/nbd0p5 not set up by udev: Falling back to direct node creation.

```
$ efa deploy
```

8. Log in to TPVM and verify the TPVM version.

```
$ ssh -l admin <TPVM_IP>
```

9. Restore the EFA over TPVM database and logs.

```
$ sudo su <-- Provide TPVM root password
```

```
$ systemctl stop efa-server
```

```
$ mv /var/efa/efa.db /tmp/
```

```
$ mv /var/log/efa/efa.log /tmp/
```

```
$ scp <server_EFA_DB_Location> /var/efa/
```

```
$ scp <server_EFA_Log_Location> /var/log/efa/
```

```
$ systemctl start efa-server
```

10. Verify that EFA over TPVM is running and verify the version.

```
$ ps -ef | grep -i efa
```

The EFA over TPVM version should be 2.0.1.

11. Before executing the **efa deconfigure** command, execute the **efa configure** command at least once following the upgrade.

Limitations and Restrictions

NPB limitations and restrictions

- When switching from NPB to default mode, the user should de-configure the following items and reload the system:
 - TVF domains, NPB policy route-map, and route-map set next-hop-tvf-domain
- When switching from default to NPB mode, the user should revert the system to default-configuration and reload the system.
- To achieve the maximum L2/L3 ACL rules, the ACLs must be applied equally among the following two port groups:
 - 9140
 - Port Group 0: eth0/1-36
 - Port Group 1: eth0/37-54
 - 9240
 - Port group0: eth 0/1-0/16
 - Port group1: eth 0/17-0/32
- With 4K TVF/route-maps scale, the system takes longer to load on config replay.
- IPv6 GTP packets are not supported for NPB L3 ACL filtering or GTP HTTPS filtering.

NPB Header stripping

- 802.1BR and VN tag are mutually exclusive on an interface.
 - Allowed only in the outer ETH.
- MPLS labels can number up to maximum of 4.
- ERSPAN stripping – Type 2 is supported. Type 1 is obsolete.
- Parser block can parse only up to 128 bytes of ingress frame.
- When both 802.1BR/VN-tag and GTP stripping are enabled, only 802.1BR/VN-tag is stripped
- When both 802.1BR/VN-tag and MPLS label stripping are enabled, only MPLS labels are stripped
- IPv6 SIP and DIP are only 64 bits each (upper or lower).
- Needs appropriate profile
- VLAN Delete will always remove the first tag
 - C in C-tag frames
 - C1 in C1+C2 tag frames
 - S in S+C tag frames
- VLAN add can only add C-VLAN tag.
- VLAN add/delete is ignored when GTP strip is enabled.

NPB Flex ACLs

- Up to 8 headers in layer stack can be accessed.
- Each flex word can be up to 4 bytes (with mask).
- Payload bytes (if available) can be 4/8/16/32 bytes.

Onboard packet capture

- Captured frames are rate limited to 256 PPS from hardware.
- Frames are truncated to 256 bytes.
- Auto stop occurs after capturing designated number of frames.
- The PCAP file is deleted automatically upon reboot.
- PCAP is supported only on one port at a time – ingress or egress and not both.

Internal loopback

- No frames will go out of the service port, even if it is connected to an external device. Hence it is suggested that the user configure only unused ports as loopback ports.
- A shut/no shut is required on a member port to bring it up, both while attaching it to a port-channel and detaching it from a port-channel.
- It is suggested, not to have sfp present in ports, configured in loopback mode. In case, sfp is present, the sequence to configure port in loopback mode, is to shut it first, configure loopback phy, change speed if required, and then do a no shut.

NPB Grid and load balancing

- NPB Grid relies on LLDP for neighbor detection hence all SLX nodes would send these frames before EVM comes down and disables them.
- 8191 PBF destinations are supported.
- 16383 PBF destination groups are supported. Each destination can have maximum of 64 members.
- 8192 PBF destination load balance groups are supported, with each group having a maximum of 31 members.
- NPB grid encapsulates frames in NSH header for forwarding, hence extra bytes are added at aggregator and removed at the last hop.

Egress packet truncation

- An excess of 10 bytes would be added in addition to requested truncation size. This includes the 4 bytes FCS.
- Maximum of 4 truncation profiles are supported.
- Truncation uses loopback ports internally as specified in a truncation profile, a truncation interface has to be put in loopback mode before using it for a truncation profile.
- Incomplete truncation profiles would cause frames to get dropped and results in no-forwarding. Frame size and truncation interface has to be set for a profile to make it complete.

Datacenter feature limitations and restrictions

TACACS+ Command Authorization

- REST/NETCONF support for TACACS+ authorization is not present
- TACACS+ Command authorization is not supported during config replay
- Exit and Quit command is not supported for authorization

DHCP Relay Source Interface Configuration

Following are the limitations for this feature:

- Consumes unique IP address per node.
- RFC 3527 supports only DHCPv4 relay, so this feature will not be supported for IPv6.

App telemetry

The following points summarize the limitations of the Application Telemetry feature:

- Flex ACLs are used internally to support this feature and are not user configurable.
- ERSPAN encapsulation internally uses one hardware SPAN session out of four available sessions. If all hardware SPAN sessions are already exhausted and the user tries to enable this feature by means of app-telemetry enable command, an error message appears.
- When content of the app-telemetry.pol file is changed, user must remove and reapply file by using rules under the provided configuration command.
- When the switch is reloaded, the app-telemetry.pol file is read and the telemetry ACLs are installed if the configuration has been saved. If the correct telemetry profile is not loaded on the switch and the feature is enabled, an error message is issued.
- Telemetry rules and ACL statistics are not persistent following a system reload.
- Application telemetry feature can support up to 1024 TCAM entries.

MAC rACLs

- MAC rACLs are not supported, as previously documented in the section “Guidelines for rACLs” in the Extreme SLX-OS Security Configuration Guide, 17s.1.02.

ACL

- Egress ACLs, Flow-Based QOS not supported on Ports and Port-Channel/MCT interfaces on SLX 9140, SLX 9240

ARAS

- Host data Collection, Ceclone backup and restore through ipv6 address is not supported.

IGMPv2 snooping

- When upgrade from 17s.100a/17s.100 to 17s.1.02, default startup query interval of 31 seconds is changed to 100sec in the running config for IGMPv2 snooping.

IP Fabric

- ACLs names are case-sensitive on Management interface.
- In rare scenarios, Ping to BGP EVPN installed prefix route host may fail, though the route is present in control plane and in hardware.
- With Scale, traffic convergence takes long time in IP Fabric for symmetric and asymmetric scenarios.
- IPv6 symmetric or asymmetric routing is not supported on SLX9240 platforms when used as Leaf nodes.
- Principal election is not pre-emptive in node join scenario.
- nsh encapsulation not supported over IPv6 neighbor.
- BFD session does not break if there is alternate path.
- Host route feature is not supported for IPv6 traffic (/128).
- MCT cluster formation takes some time in forming the cluster in scale scenarios.
- Might notice unnecessary GARP(for host address) packets seen in the network.
- Range command is not supported for BD
- BFD sessions may flap when the BFD interval is configured less than 300 msec.
- Customer tagged frames cannot be passed over VXLAN tunnel.
- Multiple flapping of CCEP ports from each nodes one after other might result in no DF elected for some VLANs. Reboot the node to recover
- Under certain circumstances when Layer 3 protocols like OSPF are run over MCT, the session might get stuck. Workaround is to reboot the switches or clear the arp suppression cache
- DF may not be elected on one of the MCT peers after upgrade. Work around is to do MCTno deploy/deploy

Layer 2

- In RSTP, when native vlan is shut, it affects convergence of vlan traffic when interop with cisco devices.

Layer3

- VRRP
 - “show vrrp summary” and “show ipv6 vrrp summary” will display all sessions in default vrf.
- BGP
 - Extended community filters support is not available.

Multicast

- Frame corruption might occur while performing high rate of replication with traffic flowing at line rate

NetConf

- Netconf configuration for startup-config datastore is not supported
- Configuring multiple commands in a single request is supported for configuration/deletion of vlan, switch port, trunk port, VE and rules under IP ACL only.
- Range is not supported.
- Maximum 16 sessions supported.

Overlay Transit Service

- Configuration download to startup and reload with 256 overlay class map each having 1024 rules takes 1 hours 40 minutes approximately

Platform

- DIAG:
 - Diag related commands work only under /offline_diag directory.
 - Diag portloopbacktest with external loopback plug is not supported on SLX9240 platform.

Port Mirroring (SPAN)

- Only Flow based SPAN supported for port channel. Member ports of port channel can be enabled with port SPAN.
- Deny rules in service ACL is pass through in Flow based QoS. Only permit rules with SPAN action will result in Flow based mirroring
- In class map if SPAN action coexists with QOS action (e.g. DSCP marking which results in frame editing), original packet will be mirrored and not reflect the frame editing done as per the QOS action.

Port-Security:

- OUI Mac Addresses are not supported.

PTP

- Rest API operational-state GET will not correctly display the output of the following PTP "show" commands:
 - show ptp clock foreign-masters record
 - show ptp corrections
- No REST API URL for "show ptp port-interface Ethernet|port-channel"

QoS

- FB QoS - Cos Marking, DSCP Marking, Sflow, SPAN
 - SPAN with L2 ACL in egress direction (SLX 9240)
 - Flow-based QoS is not supported in egress direction
- QoS – WRED

- Byte counter is not available as part of show qos red statistics CLI for port-channel
- QoS – Pause/PFC/Buffer Management
 - PFC and Flow-control statistics are not supported due to hardware limitation
 - Max allowed tx buffer in SLX9140 is 3000 and not 8000.

REST API

- REST configuration for startup-config datastore is not supported.
- Only one command can be configured with one REST request. Configuring multiple commands in a single request is not supported.
- Pagination and Range is not supported.
- Maximum 30 sessions are supported.

REST API/NetConf Operational-state calls

- HTTP Status throw message “501 Not Implemented” while trying to get operational state for top resource (rest/operational-state) using REST API, User can query operational-state at feature level.
- Operational-state calls not supported for Overlay GW and Visibility Services features.
- Yang files for Unsupported features like MPLS, ISIS are available and operational-state call returns empty value or “404 not found”
- Operational-state calls for supported feature mat not be accurate and may return “404 not found” or empty value, not advisable to use it

Security

- Login authentication service (aaa authentication login cli)
 - With “local” option specified as secondary authentication service, local authentication will be tried only when the primary authentication service (TACACS+/RADIUS/LDAP) is either unreachable or not available.
 - When login authentication configuration is modified, the user sessions are not logged out. All connected user sessions can be explicitly logged out using “clear sessions” CLI.
- ACLs are not supported for egress traffic flows on management interfaces.
- Configuring TACACS+ or RADIUS without a key is not supported. If no key is configured, the switch uses a default key of “sharedsecret”. If the specific vrf is not mentioned, mgmt.-vrf will be taken as default.
- There is a possibility that locked user accounts will get unlocked after a reboot if the running-config (before reboot) is different from startup-config of user accounts.
- Encrypted text (taken from running-config of any user account password with encryption turned on) should not be used as input for clear-text password for the same user. This may result in login failure of the user subsequently.

sFlow

- If Port based and flow based sflow is enabled on an interface, Port based sflow takes effect
- Flow-based Sflow is not supported on port-channel and its member ports
- Port-based Sflow not supported on port-channel but supported on member ports
- There will be no counter samples when only flow based sampling is enabled.

- When multiple sampling rates are applied on an interface through multiple class-maps, the lowest sample-rate will take the effect.

SNMP

- Warning messages while loading MIBs
- Certain MIB browsers may show warning messages while loading MIBs when dependent MIB is already not loaded. For example, in RFC 3289 MIB, DIFFSERV-MIB module has dependency on INTEGRATED-SERVICES-MIB module which is defined in the same RFC. However, DIFFSERV-MIB occurs first in the file and hence may throw a warning since INTEGRATED-SERVICES-MIB is not loaded yet. It should not be an issue as long as the MIB objects show up in the MIB browser. To avoid the warning, place the dependent MIB module file in the same folder with name as <MIB MODULE>.mib or <MIB MODULE>.my (ex: INTEGRATED-SERVICES-MIB.mib) ...”

Telemetry Streaming

- Running gRPC server on non-default port not supported.

Traffic

- On the SLX 9140 and SLX 9240 switches, traffic destined to 128.0.0.0/16 block is dropped.
- Hash collisions may be observed with higher scale in Route, ARP/Mac and/or Tunnel tables resulting in entries not getting programmed.

TPVM

- Upgrade and downgrade procedures have changed. Refer to “TPVM” in the “Software Upgrade and Downgrade” section.
- The **tpvm password** command is not supported. Unexpected behavior can result.

Defects

Closed with or without code changes in SLX-OS 18s.1.03a

Parent Defect ID:	SLXOS-45136	Issue ID:	SLXOS-45139
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 17s.1.02	Technology:	Other
Symptom:	Switch reported FFDC with RAS-1004 Software verify error detected warning message.		
Condition:	This is very rare scenario to hit this issue.		

Parent Defect ID:	SLXOS-46001	Issue ID:	SLXOS-46002
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLXOS 18s.1.03	Technology:	Telemetry
Symptom:	Below parameters are missing from Telemetry profile interface default_interface_statistics: out-link-utilization and in-link-utilization parameters		
Condition:	When upgrade from lower versions to 18s.1.03a with telemetry profile configured.		

Parent Defect ID:	SLXOS-46603	Issue ID:	SLXOS-46752
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	SLXOS 18s.1.03	Technology:	BGP4 - IPv4 Border Gateway Protocol
Symptom:	EVPN MACs not getting learnt.		
Condition:	When the MAC moves from one multi-chassis trunk cluster to another.		

Parent Defect ID:	SLXOS-47400	Issue ID:	SLXOS-47401
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	SLXOS 18s.1.03	Technology:	BGP4 - IPv4 Border Gateway Protocol
Symptom:	NLRI-entries were NOT cleared before clearing peer safi entry. can cause loop and unexpected reload		
Condition:	Upon execution of cmd "no neighbor <ip-address> activate"		
Workaround:	Validates nlri-entries and then clean up the peer safi entry. Peer safi entry is cleaned as part of timer callback.		

Parent Defect ID:	SLXOS-47502	Issue ID:	SLXOS-47503
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18s.1.03	Technology:	CLI - Command Line Interface
Symptom:	TPVM display under CLI "show interface status".		
Condition:	When TPVM is not installed and executed "show interface status"		

Parent Defect ID:	SLXOS-47956	Issue ID:	SLXOS-47957
Severity:	S1 - Critical		
Product:	SLX-OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	SLXOS 18s.1.03	Technology:	IP Addressing
Symptom:	Traffic loss for few hosts		
Condition:	1.Deleting the VLAN+VE interface 2.Any route updates on the device(like addition/deletion of route)		

Parent Defect ID:	SLXOS-49149	Issue ID:	SLXOS-49840
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Security
Reported in Release:	SLXOS 20.1.2	Technology:	User Accounts & Passwords
Symptom:	Admin user can get the root privileges		
Condition:	when user try to use start-shell, python, OSCMD from admin login		

Parent Defect ID:	SLXOS-50265	Issue ID:	SLXOS-50266
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLXOS 18s.1.03	Technology:	Telemetry
Symptom:	Unexpected reload		
Condition:	When the telemetry config is enabled for PBR profile		

Parent Defect ID:	SLXOS-50542	Issue ID:	SLXOS-50750
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18s.1.03	Technology:	CLI - Command Line Interface
Symptom:	Unexpected reload		
Condition:	When "no breakout" is being configured and breakout port has "deny inner-gtp-https" configuration already.		
Workaround:	Remove "deny inner-gtp-https" configuration on breakout port before doing "no breakout"		

Parent Defect ID:	SLXOS-50879	Issue ID:	SLXOS-51108
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	SLXOS 18s.1.03	Technology:	Static Routing (IPv4)
Symptom:	Unexpected reload		
Condition:	When we configure the "set interface null0" under the route-map configuration.		

Closed with or without code changes in SLX-OS 18s.1.03

Parent Defect ID:	SLX-OS-18587	Issue ID:	SLX-OS-18587
Severity:	S4 - Low		
Product:	SLX-OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	SLX-OS 17s.1.02	Technology:	OSPF - IPv4 Open Shortest Path First
Symptom:	ASBR status is printed No, when a switch is ABR in NSSR area.		
Condition:	When switch is configured ABR in NSSA area , ASBR status still shows No.		
Workaround:	This is show command print issue and no impact on functionality. As per functionality the switch does act as ASBR		

Parent Defect ID:	SLX-OS-18629	Issue ID:	SLX-OS-18629
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLX-OS 17s.1.02	Technology:	Configuration Fundamentals
Symptom:	When configuring timestamp on a range of interfaces, only the first interface in the range is being configured.		
Condition:	Configure timestamp configuration on a range of interfaces		
Workaround:	Configure timestamp on an individual port basis.		

Parent Defect ID:	SLX-OS-19337	Issue ID:	SLX-OS-19337
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLX-OS 17s.1.00a	Technology:	Hardware Monitoring

Symptom:	The status light on the switch blinks from Amber to Green even though one FAN is missing.		
Condition:	A missing fan causes the switch status LED to blink.		

Parent Defect ID:	SLX-OS-21044	Issue ID:	SLX-OS-21044
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLX-OS 17s.1.02	Technology:	Hardware Monitoring
Symptom:	10GE SFP+ optics used with Mellanox QSA Adapter may not link up.		
Condition:	When 10GE SFP+ optic is used with Mellanox QSA Adapter the port may not link up and "Unqualified SFP transceiver" logs would be reported on the console.		

Parent Defect ID:	SLX-OS-21059	Issue ID:	SLX-OS-21059
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLX-OS 17s.1.02	Technology:	Telemetry
Symptom:	Random link down issues noticed with 100G optics		
Condition:	Even after configuring both Tx and Rx link fault signaling to off, random link (remaining) down issues noticed with 100G optics.		
Workaround:	Configure link fault signaling "rx off tx on" as a workaround		

Parent Defect ID:	SLX-OS-24366	Issue ID:	SLX-OS-24366
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLX-OS 17s.1.02	Technology:	MCT - Multi-Chassis Trunking
Symptom:	"Show mac-address-table interface" doesn't display learnt Mac-Address on Logical interface.		
Condition:	Execute ?clear mac-address-table dynamic logical-interface <name>? command instead of "clear mac-address-table dynamic".		
Workaround:	Execute `clear mac-address-table dynamic bridge-domain <id>?`		

Parent Defect ID:	SLX-OS-24793	Issue ID:	SLX-OS-24793
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	SLX-OS 17s.1.02	Technology:	OSPFv3 - IPv6 Open Shortest Path First
Symptom:	256 OSPFv3 interfaces are supported on default VRF and in non-default 1000+ OSPFv3 interfaces can be configured.		
Condition:	OSPFv3 interface scale above 256 interfaces in default VRF.		

Parent Defect ID:	SLX-OS-26258	Issue ID:	SLX-OS-26258
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLX-OS 18s.1.00	Technology:	Software Installation & Upgrade
Symptom:	Default high threshold value for "Current mA" field shown in the output of CLI command "show default threshold" is incorrect.		
Condition:	For SFP type 40GSRINT, the CLI output of "show defaults threshold sfp type 40GSRINT" displays incorrect high threshold for the "Current mA" field.		

Parent Defect ID:	SLX-OS-26341	Issue ID:	SLX-OS-26341
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLX-OS 18s.1.01	Technology:	Software Installation & Upgrade
Symptom:	After upgrade from 17s build to 18s build, sometimes, "show tpvm status" displays runtime error, and indicates use "tpvm install force" to clear the error. The message to use "tpvm install force" is not correct, this command is not supported, it needs to be removed from message.		
Condition:	Upgrade from 17s build to 18s build		
Workaround:	A workaround is to uninstall TPVM before upgrade.		

Parent Defect ID:	SLX-OS-26996	Issue ID:	SLX-OS-26996
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLX-OS 18s.1.01	Technology:	VLAN - Virtual LAN
Symptom:	This issue is seen only when CLIs "link-error-disable" and "link-fault-signaling" are being configured for an interface first time		
Condition:	when CLI "loopback phy" is not configured.		
Workaround:	No		

Parent Defect ID:	SLX-OS-26743	Issue ID:	SLX-OS-31039
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	SLX-OS 18s.1.01	Technology:	BGP4 - IPv4 Border Gateway Protocol
Symptom:	User might observe that the REST API for BGP EVPN IP Fabric is giving some discrepancies for operational data.		
Condition:	User is using REST to query BGP EVPN IP Fabric Operational DB		

Parent Defect ID:	SLX-OS-28744	Issue ID:	SLX-OS-31448
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLX-OS 17r.1.01af	Technology:	Hardware Monitoring
Symptom:	"show system monitor" is not supported, however there is no any functional impact and only display issue.		
Condition:	"show system monitor" is not supported, however there is no any functional impact and only display issue.		

Parent Defect ID:	SLX-OS-36052	Issue ID:	SLX-OS-37628
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLX-OS 17s.1.03	Technology:	CLI - Command Line Interface
Symptom:	[PI-RESTAPI] Device is getting "application communication failure" after shutdown http server with user-defined vrf		
Condition:	Shutdown http server with user-defined vrf		
Workaround:	None		

Parent Defect ID:	SLX-OS-38108	Issue ID:	SLX-OS-38415
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLX-OS 18r.1.00a	Technology:	Licensing
Symptom:	LICD termination while upgrading the code from 18r.1.0.0a to 18r.1.0.0aa.		
Condition:	LICD termination while upgrading the code from 18r.1.0.0a to 18r.1.0.0aa.		

Parent Defect ID:	SLX-OS-38942	Issue ID:	SLX-OS-38942
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLX-OS 18s.1.01	Technology:	LAG - Link Aggregation Group
Symptom:	Traffic ingressing on a Tunnel and egressing out of interface 0/9 and 0/9:1 on SLX 9240s can get dropped.		
Condition:	The issue happens when all these conditions are true, 1. It's SLX 9240 2. Ingress is a VXLAN tunnel or the ICL interface (NSH tunnel) 3. Egress interface is 0/9 or 0/9:1		

Parent Defect ID:	SLX-OS-28689	Issue ID:	SLX-OS-39224
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLX-OS 17r.2.01	Technology:	Configuration Fundamentals
Symptom:	Symptom 1) Unable to configure large number of VLANs under an EVPN instance Symptom 2) Unable to restore a startup-config file containing a large number of VLANs under an EVPN instance.		
Condition:	Condition 1) When a vlan configuration, more than 253 characters long, is attempted when configuring an EVPN instance a length validation error is noticed. Condition 2) A vlan configuration, more than 253 characters long, when split and configured under a EVPN instance is saved in the startup-config file and restored the system may not accept the configuration.		
Workaround:	Workaround 1) Split the VLAN configuration into multiple lines. Workaround 2) Reconfigure the VLANs		

Parent Defect ID:	SLX-OS-38336	Issue ID:	SLX-OS-39629
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLX-OS 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:	SLX-OS-40106	Issue ID:	SLX-OS-40106
Severity:	S4 - Low		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLX-OS 18s.1.02	Technology:	Other
Symptom:	Output of 'show running-config' contains "advertise bgp-auto-nbr-tlv" even though the config is not supported in NPB mode.		
Condition:	Issue is seen only when switch is running in NPB mode.		

Parent Defect ID:	SLX-OS-40846	Issue ID:	SLX-OS-40846
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 3 Routing/Network Layer
Reported in Release:	SLX-OS 17s.1.02b	Technology:	ARP - Address Resolution Protocol
Symptom:	Traffic to/from DHCP host is not routed when the DHCP IP is assigned to a new host. The ARP for such host does not age out when age out timer expires.		
Condition:	DHCP Server is sending ACK packets to relay agent even when the client address is known. Mostly seen with Windows DHCP server.		
Workaround:	Use command : clear arp ip <IP address>		

Parent Defect ID:	SLX-OS-41166	Issue ID:	SLX-OS-41745
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLX-OS 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:	SLX-OS-38299	Issue ID:	SLX-OS-42609
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Other
Reported in Release:	SLX-OS 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:	SLX-OS-42673	Issue ID:	SLX-OS-42675
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Other
Reported in Release:	SLX-OS 18x.1.00a	Technology:	Other
Symptom:	Unexpected reload		
Condition:	When the management cluster is in broken state.		

Parent Defect ID:	SLX-OS-37521	Issue ID:	SLX-OS-37521
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Monitoring
Reported in Release:	SLX-OS 18s.1.02	Technology:	Telemetry
Symptom:	Traffic Stream with offset value 0x01000200 matching radius1 UDA ACL instead of ZOOM-UDP1 UDA ACL		
Condition:	<p>Two app-telemetry rules are configured and traffic for second rule will also match the first configured rules.</p> <p>For example:</p> <p style="padding-left: 40px;">offset / mask for radius1 is 0x01000000 / 0xff000000 and for ZOOM-UDP1 is 0x01000200 / 0xfffff00</p> <p>Any traffic which matches ZOOM-UDP1 will always match radius1 as well. As radius1 tcam entry is programmed before ZOOM entry, so traffic is hitting radius1 entry (only one tcam entry can take hit in same region).</p>		
Workaround:	As ZOOM-UDP1 entry is more specific than radius1 entry, it needs to be programmed before radius1 entry.		