

SLX-OS 18s.1.01b Release for the ExtremeSwitching SLX 9140 and SLX 9240 Platforms, Release Notes v1.0

July 2019

Copyright Statement and Legal Notices

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

Copyright Statement and Legal Notices.....	2
Document History	4
Preface	4
Software Features.....	6
Supported Optics	18
Documentation Supporting SLX-OS	21
Software Upgrade and Downgrade	22
Limitations and Restrictions.....	28
Defects	30

Document History

Version	Summary of changes	Publication date
1.0	Initial Release	July 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 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.

Software Features

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

SLX-OS 18s.1.01b

SLX-OS 18s.1.0.1b release is the fifth release in a series for SLX platforms – the SLX 9140 is the focus of the new features in this release, while the SLX 9240 remains supported as in the previous release. 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.1.01b are focused on enhancing manageability and user experience on SLX.

The new features are as follows:

- **Endpoint tracking:** MAC authentication using RADIUS protocol on SLX 9140
- **Embedded Fabric Automation (EFA-Lite):** Application support on the SLX 9140 and SLX 9240, interoperating with the SLX 9030 as a leaf and the SLX 9540 as a leaf and border leaf

This release supports the EFA-Lite bring-up of **Clos topology** networks with SLX 9030 and SLX 9540 as leaf nodes, in addition to the bring-up of the SLX 9140 as a leaf as supported in the previous release. This release introduces support for the EFA-Lite bring-up of **non-Clos topology** networks with the SLX 9140 as a rack platform.

NOTE: It is recommended that EFA-Lite run on one of the rack SLX 9140 platforms, as well as on one of the spine SLX 9240 platforms.

The SLX 9140/SLX 9240 are fixed 1-RU switching platforms based on programmable ASICs that enable the adoption of new protocols and technologies. These switches were released as a part of SLX-OS 17s.1.00, SLX-OS 18s.1.00, and SLX-OS 18s.1.01 to support enhanced Network Packet Broker (NPB) pipeline and SLX switching and routing feature. The features introduced in 18s.1.01 are as follows:

- High-density 40G/100G spine-leaf connection
- SLX 9140 native 1G/10G/25G server connectivity at the leaf
- High performance VXLAN routing
- Payload timestamping to enable accurate measurement of performance SLAs
- Port-to-port latency: ~2.5 microseconds
- Architecture: store and forward
- Enhanced NPB Pipeline
- Support for 4x10G and 4x25G with dynamic breakout

The major focus of this SLX-OS release is on the feature enhancements supporting End Point Tracking (EPT) and Data Center, with Embedded Fabric Automation (EFA-Lite) for DC Fabric use cases.

Endpoint tracking – MAC authentication using RADIUS protocol

The endpoint tracking feature minimizes the configuration and management of VLANs on switches in the data center, by means of MAC authentication using RADIUS protocol.

Overview

- Managing VLANs on top-of-rack (TOR) switches is tedious for the administrator.
- Having VLANs provisioned ahead of time increases the size of the active topology for control protocols such as STP and RSTP, increasing convergence times.
- Flood (unknown unicast, broadcast) traffic can unnecessarily eat up bandwidth on the TOR-to-EOR (end of row) links.
- For the case when the virtual machine (VM) sends or receives tagged traffic, flood traffic can consume CPU cycles on every server that is connected to the network.
- The dynamic VLAN feature allows SLX-OS to create, prune, and open VLANs on the switch dynamically as they are needed by the VMs. This enables the VLAN to follow the VM as it migrates between servers in the data center.

This feature provides the following to remedy the above drawbacks:

- Association of MAC addresses to specific VLANs.
- Once a MAC address is authorized and the VLAN is not already provisioned on the switch, the dynamic (1) creation of the VLAN to which this MAC is associated, (2) configuration of the tag, and (3) addition of the feature-enabled port on which the MAC was detected.
- Once the last MAC address using the VLAN is deleted or aged out, deletion of the VLAN and associated resources on the switch.

The endpoint tracking feature also authorizes the VM. When a VM (and MAC address) is authorized, SLX-OS dynamically creates the VLAN that is required for the VM to send traffic. If a VM shuts down or is moved, its VLAN is pruned to preserve bandwidth. In this way the network responds to changes in the VM network.

Additional details and considerations

Note the following:

- The user enables or disables this feature on a Layer 2 port (switchport), by means of the endpoint-tracking enable command. This feature is also supported on LAG and MCT ports.
- SLX-OS communicates with a RADIUS server or XMC-NAC (Extreme Management Center Network Access Control) for MAC authentication information, which can map a MAC address to a VLAN. RADIUS VLAN assignment triggers the creation of the VLAN and the port VLAN membership. SLX-OS expects MAC-to-VLAN binding information from RADIUS. VNI (VXLAN Network Identifier) assignment is not supported.
- The maximum number of authenticated or non-authenticated MAC addresses supported for each endpoint-tracking-enabled is 2000 per system.

- Port VLAN membership is local to SLX-OS. The removal of port VLAN membership or the deletion of the VLAN is not propagated to the RADIUS server.
- On an access port, if RADIUS does not provide MAC-to-VLAN mapping, SLX-OS assigns a default VLAN. Only the first RADIUS VLAN assignment is honored, and subsequent MAC addresses cannot override the existing assignment.
- For VMs sending tagged traffic, if the RADIUS authentication response does not have a VLAN assignment for the MAC, SLX-OS creates the VLAN (assuming the VLAN does not already exist) with the received packet's tag and adds the port to the VLAN as tagged. Dynamically created VLANs are not saved across reboots, and therefore are not part of the running configuration.
- When a port becomes part of a VLAN dynamically, this is not reflected in the running configuration of the interface.
- In reauthentication, RADIUS sends two attributes as part of a Change of Authorization (CoA) request (a CoA is an unsolicited message from RADIUS to the switch to trigger an authentication-related action): (1) Calling-station-id, a value field that holds the MAC being authenticated, and VSA, a vendor-specific attribute value field that holds commands and parameters for commands. RADIUS matches the outstanding requests response by means of a CoA identifier.

Enabling and disabling endpoint tracking on a port

This feature is enabled on a per-port basis, on a port in Layer 2 mode (switchport). This section provides additional details and illustrates the configuration.

Endpoint tracking on an access port

When endpoint tracking is enabled on an access port, initial MAC learning occurs on the default VLAN and after authentication if there is RADIUS VLAN assignment. Note the following:

- If RADIUS assigns the VLAN, the VLAN is created if it does not exist. The Layer 2 forwarding entry is added with the RADIUS-assigned VLAN.
- If RADIUS does not assign the VLAN, the Layer 2 forwarding entry is added with the default VLAN.

The following illustrates the use of the endpoint-tracking enable command in switchport access mode (the default).

```
device# configure terminal
device(config)# interface ethernet 0/1
device(conf-if-eth-0/1)# switchport
device(conf-if-eth-0/1)# endpoint-tracking enable
```

Endpoint tracking on a trunk port

On a trunk port, packets are always tagged. To allow any tagged traffic on the port, ingress VLAN filtering is disabled in the forwarding plane when endpoint tracking is enabled. As a result, all the tagged packets are trapped to the CPU as a Layer 2 learning event, which is later authenticated by RADIUS. A MAC learning event that is generated has wire tag information as part of the VLAN information. After authentication, a RADIUS-assigned dynamic VLAN is added for egress forwarding on the port.

The following illustrates the use of the **endpoint-tracking enable** command in switchport trunk mode.


```

device# configure terminal
device(config)# interface ethernet 0/1
device(conf-if-eth-0/1)# switchport
device(conf-if-eth-0/1)# switchport mode trunk
device(conf-if-eth-0/1)# endpoint-tracking enable

```

Once endpoint tracking is enabled on a trunk port, the following occurs:

1. The source MAC lookup fails to find a matching entry in hardware and the packet is sent to the CPU for Layer 2 learning with a tag in place of the VID.
2. When received by the CPU, the wire tag information is stored as a VLAN, and the MAC address is sent for authentication.
3. If authentication is successful and the RADIUS VLAN assignment matches the wire tag, or if there is no VLAN assignment from RADIUS, VLAN creation is triggered and the egress VIF is set for the port to allow forwarding and flooding.

Verifying configurations

A variety of show commands are available to verify the configuration of endpoint tracking, as described in the following table.

Endpoint tracking show commands for SLXS

Command	Description
show mac-address-table endpoint-tracking authenticated	Displays authenticated MAC addresses that are learned on all ports that are enabled for endpoint tracking.
show mac-address-table endpoint-tracking authentication-failed	Displays nonauthenticated MAC addresses that are learned on all ports that are enabled for endpoint tracking.
show mac-address-table endpoint-tracking authenticated interface	Displays authenticated MAC addresses that are learned on a specific port that is enabled for endpoint tracking.
show mac-address-table endpoint-tracking authentication-failed interface	Displays nonauthenticated MAC addresses that are learned on a specific port that is enabled for endpoint tracking.
show vlan brief	Displays all dynamically created VLANs and port/VLAN membership.

VM MAC aging and flush

When a VMs MAC address is deleted because of aging or a flush operation, if it is the last MAC on the port, SLX-OS removes the port from the VLAN and checks to see if the VLAN is associated with other ports. If there are no other associations, the VLAN is also deleted.

MAC reauthentication

MAC reauthentication lets the RADIUS server send unsolicited messages to SLX-OS, to relearn the MAC address of the VM. Note the following:

- SLX-OS stores the CoA request identifier and uses the same identifier in the response (ACK/NAK).
- A RADIUS reauthentication request without calling-station-id is returned with a NAK.
- A RADIUS reauthentication request with a calling-station-id that is not present in the switch is returned with a NAK.
- A RADIUS request with a different vendor-id is silently ignored by the switch.
- Duplicate CoA messages are rejected at the switch.
- Re-authentication can be achieved by means of either a CoA or a Disconnect message. In the case of a CoA message, the VM is not removed and traffic is not disrupted during re-authentication. In the case of a Disconnect message, the VM is removed from the switch. When new traffic from the VM is relearned, authentication is triggered, and traffic is disrupted.

Additional considerations for EPT

Please note the following:

- Logical Interfaces (LIF) are not supported on EPT ports
- Spanning-tree protocol (STP) should be disabled on a port before the port is configured for EPT

MCT support

A VLAN that is created dynamically on one MCT peer node is communicated to other peer node. Similarly, if a cluster client edge port (CCEP) port becomes part of a VLAN dynamically on one MCT peer node, this is communicated to the other peer node. The receiving MCT peer node, depending on the message, creates a dynamic VLAN or dynamic port VLAN membership.

The dynamic deletion of a VLAN or port VLAN membership is triggered by the last local MAC deletion. In addition, similar to the addition case, this information is also communicated to the peer.

Static VLAN dependency

A statically configured VLAN has a higher precedence over a dynamically created VLAN. Note the following:

- If a dynamically created VLAN is configured as static as well, removing the static VLAN also removes the dynamic VLAN information from the system.
- If a dynamically created port/VLAN membership is configured as static as well, removing the static port/VLAN membership also removes the dynamic port/VLAN membership from the system.

IP directed broadcast on an interface

A directed broadcast is an IP broadcast to all devices within a directly attached network or subnet. You can enable IP directed broadcast on a Layer 3 interface. The Layer 3 interface can be a physical Ethernet interface or a VE interface. When the device receives a packet with a destination IP address as broadcast IP and IP directed broadcast is enabled on an interface through which the destination network is reachable, the interface floods the packet to all hosts of this network. IP directed broadcast is supported on both default and user-defined VRFs.

Enabling IP directed broadcast on an interface

By default, IP directed broadcast is disabled on the interfaces of the device.

Perform the following steps to enable IP directed broadcast on an Ethernet interface.

1. From privileged EXEC mode, access global configuration mode.
device# configure terminal
2. Specify the interface.
device(config)# interface ethernet 0/2
3. Enable IP directed broadcast on the interface.
device(config-if-eth-0/2)# ip directed-broadcast.

EFA 2.0.0

With this release Extreme Fabric Automation support is added. Extreme Fabric Automation (EFA) is also known as Data Center Automation Application (DCA) is a Go-based, scalable Golang-based application that orchestrates the following installations:

- 3-stage IP Fabric
- 5-stage IP Fabric
- Tenant Aware Networks

For details refer to:

SLX-OS 18s.1.01b Release for the ExtremeSwitching SLX 9140 and SLX 9240 Platforms, Release Notes v1
9036079-01 Rev AA

- Extreme Fabric Automation 2.0.0 Administration Guide v1.0
- Extreme Fabric Automation 2.0.0 Release Notes v1.0

SLX-OS 18s.1.01

The following are the major features added in SLX-OS 18s.1.01 release.

Feature Name	Use case
IGMPv2 snooping with MCT	DC
SNMP Enhancements	DC
Management ACL to block ICMP timestamp in response packet	DC
Internal loopback support	NPB
On-board packet capture	NPB
EFA with additional support for SLX 9240 as a leaf	DC
New Optics Qualification 25G-LR Media	ALL
Qualify SLX 9240 as a high-density leaf	DC

SLX 9140 and SLX 9240 as Network Packet Broker

SLX HW can be used as standard switching/routing or in NPB-only mode. NPB features are enabled only in NPB mode with the following enhance header stripping and Flex ACL features with advance NPB scale. The following table summarizes the NPB features introduced with SLX-OS 18s.1.00.

Feature Name	Feature Description
NPB Parser	The ability to parse new set of protocol on existing hardware. VXLAN, NVGRE, ERSPAN, IP-GTP-IP, IP-GRE-IP, IP-IP, EoMPLS, IPoMPLS, IPv4/IPv6/ARP . Offset agnostic parsing up to inner L4/payload parsing. Payload (4/8/16/32) bytes follow the last possible parsed header.
Header Stripping	The ability to strip the header, for example, tunnel encapsulation and BR/VN tags for customer tools to analyze traffic that may or may not be able handle some of the tags or packet encapsulations during traffic analysis.
Flex/UDA Match ACL	The ability to filter based on deep packet inspection (DPI) or combination of MAC, IP fields using user-defined Flex ACLs (new for SLX 9240 and SLX 9140 platforms). Parses relatively deep into the packet. In MLX and SLX 95430 the UDA is based on offset/pattern match.

Consolidated Features in SLX-OS 18s.1.01

The following table lists the features introduced in SLX-OS 18s.1.01.

NPB Mode Features

Header Stripping	
<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.
Transparent VLAN	
<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
Flex ACLs	
<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.

	<ul style="list-style-type: none"> • 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.
Scale Improvements	
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
VLAN	<ul style="list-style-type: none"> • 400 VLANs
NPB enhancements	
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.
New optics Qualified	
25G SFP28 LR	<ul style="list-style-type: none"> • 10504

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

Data Center Solutions

IGMPv2 Snooping with MCT	
	<ul style="list-style-type: none"> • IGMPv2 snooping support on MCT. • BGP EVPN based IGMP state syncing. • IGMP VLAN and Group Based DF election for traffic forwarding to MCT clients. <p>Note:- Bridge Domain (BD) case is not supported.</p>
SNMP Enhancements	
	<ul style="list-style-type: none"> • Support to disable SNMP traps for interface link status change events for Ethernet, Loopback, VE, and port-channel interfaces. • Allow SNMP server to be disabled/enabled for specific VRFs or for all the VRFs.
Mgmt ACL to block ICMP timestamp in response packet	
	<ul style="list-style-type: none"> • Allows ICMP timestamp requests and responses to be dropped, by default, as a security policy. • If an ACL with a “<i>permit icmp any any</i>” rule is applied to the management interface, such a rule permits ICMP timestamp requests but ICMP timestamp response is blocked.
Qualify SLX 9240 as a high-density leaf	
	<ul style="list-style-type: none"> • Allows to deploy SLX 9240 as a high-density leaf node in an IP fabric environment, similar to SLX 9140. • Routing in and out of Tunnel (RIOT) is not supported with IPv6. • Reduced scale numbers compared to 9140.

Embedded Fabric Automation Lite

EFA-Lite is an application that can be installed on the TPVM (Third Party Virtual Machine) on the SLX 9240 (spine). The application is bundled as part of the SLX-OS firmware and can be used to configure an IP Fabric on the SLX 9240 and SLX 9140. EFA-Lite is documented in the “Embedded Fabric Automation” chapter of the *Extreme SLX-OS IP Fabrics Configuration Guide*. The following platform roles are supported: SLX 9140 as a leaf, and SLX 9240 as a spine and leaf.

Embedded Fabric Automation (EFA-Lite)	
<p>EFA-Lite application is bundled as part of the SLX firmware and can be used to configure an IP Fabric on the SLX 9240 and SLX 9140.</p> <p>NOTE: See “Upgrading to EFA-Lite 2.0.0 with TPVM Ubuntu version 16.04” later in this document</p>	<ul style="list-style-type: none">• The EFA-Lite application runs on the SLX TPVM to manage the IP Fabric.• EFA-Lite provides simplified mechanism to configure IP Fabric, once the IP addresses of the devices are provided.• Support for SLX 9240 as a leaf is added.

Additional Considerations. Please note the following:

- The default ICMP rate limiting behavior is changed. For the **ip icmp rate-limiting** and **ipv6 icmpv6 rate-limiting**, the default is changed from 1000 to 0.

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
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
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

25G SFP28 LR (10km), Single Mode, LC-connector, 70degC	25G-SFP28-LR	10504
--	--------------	-------

Mellanox supports the following 10G optics:

- 10G USR SFP+
- 10G SR SFP+
- 10G LR SFP+ in RC2

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

Documentation Supporting SLX-OS

SLX-OS 18s.1.01a

For documents supporting the most recent previous release, see the following:

<https://www.extremenetworks.com/support/documentation/version-18s-1-01a/>

For additional documentation support, see the following:

SLX-OS 18s.1.01

<https://www.extremenetworks.com/support/documentation/slx-s-series-software-18s-1-01/>

Software Upgrade and Downgrade

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

SLX-OS 18s.1.01b

Image file names

Download the following images from www.extremenetworks.com.

Image file name	Description
slxos18s.1.01b.tar.gz	SLX-OS 18s.1.01b_ software
slxos18s.1.01b_all_mibs.tar.gz	SLX-OS 18s.1.01b_ MIBS
slxos18s.1.01b.md5	SLX-OS md5 checksum

To Install SLX-OS 18s.1.01b 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.1.01b from a USB device, follow the steps below:

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

Migration path

Recommended upgrade/downgrade migration paths in default mode:

NOTE: Only if upgrading to **SLX 18s.1.01** or **SLX 18s.1.01x** with “default-config” option, then “app-tele-l2-l3-iacl” HW profile is activated.

To	SLX17s.1.00a	SLX17s.1.01	SLX17s.1.02	SLX17s.1.02x	SLX 18s.1.00	SLX 18s.1.01	SLX 18s.1.01x
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	FWD coldboot	FWD coldboot	FWD coldboot
SLX 17s.1.02x	Default – config	Default-config	FWDL coldboot	NA	FWD coldboot	FWD coldboot	FWD coldboot
18s.1.00	*	*	Default - config	Default - config	NA	FWD coldboot	FWD coldboot
18s.1.01	*	*	Default - config	Default - config	FWD coldboot	NA	FWD coldboot
SLX 18s.1.01x	*	*	Default - config	Default - config	FWD coldboot	FWD 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.

Recommendations

Upgrading EFA-Lite from 18s.1.01a to 18s.1.01b with Ubuntu version 16.04

Do the following to upgrade the EFA-Lite application.

1. Log in to TPVM with the TPVM IP address.

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

2. Copy the EFA-Lite database and logs backup to an external server.

```
$ service efa-server stop
```

```
$ scp /var/efa/efa.db <server_EFA-Lite_DB_Location>
```

```
$ scp /var/log/efa/efa.log <server_EFA-Lite_Log_Location>
```

3. Stop and uninstall TPVM.

```
$ tpvm stop
```

```
$ tpvm uninstall
```

4. With the device upgraded to 18s.1.01b, execute the **efa deploy** command.

```
$ efa deploy
```

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

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

6. Restore the EFA-Lite 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
```

7. Verify that EFA-Lite is running and verify the version.

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

The EFA-Lite version should be 2.0.0.

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

SLX-OS 18s.1.01

Migration path

Recommended upgrade/downgrade migration paths in NPB mode:

NOTE: Starting with SLX 17s.1.01 NPB feature is supported only with “Advanced feature” license.

- **SLX 17s.1.02/x/ab/18s.1.00 to SLX 18s.1.01 Upgrade**
 - Save running-config to startup-config
 - Make sure SLX 9240 / 9140 is in NPB mode
 - Make sure ADVANCED FEATURES license is present
 - Firmware Upgrade to SLX 18s.1.01

- **SLX 18s.1.01 to SLX 17s.1.02/x/ab/18s.1.00 Downgrade**
 - Save running-config to startup-config
 - Take back-up of the config
 - Firmware downgrade with default-config option
 - Restore the config

To	SLX 17s.1.00	SLX 17s.1.00a	SLX17s.1.01	SLX17s.1.02	SLX17s.1.02x	SLX 18s.1.00	SLX 18s.1.01
From							
SLX 17s.1.00	NA	Default - config	Default - config	Default - config	Default - config	*	*
SLX 17s.1.00a	Default - config	NA	Default - config	Default - config	Default - config	*	*
SLX 17s.1.01	Default-config	Default - Config	NA	FWD coldboot	Default - config	*	*
SLX 17s.1.02	Default-config	Default - Config	FWDL coldboot	NA	FWD coldboot	FWD coldboot	FWD coldboot
SLX17s.1.02x		Default - Config	Default - Config	FWD coldboot	NA	FWD coldboot	FWD coldboot

	Default-config						
SLX 18s.1.00	*	*	*	Default - config	Default - config	NA	FWD coldboot
SLX 18s.1.01	*	*	*	Default - config	Default - config	FWD 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.

Upgrading to EFA-Lite 1.1.0 with TPVM Ubuntu version 16.04

Do the following to upgrade the EFA-Lite application.

In the previous version, SLX-OS 17s.1.02b used TPVM Ubuntu version 14.04 on an SLX-9240. This task upgrades the version to 16.04.

1. Log in to TPVM with the TPVM IP address.

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

2. Copy the EFA-Lite database and logs backup to an external server.

```
$ service efa-server stop
```

```
$ scp /var/efa/efa.db <server_EFA-Lite_DB_Location>
```

```
$ scp /var/log/efa/efa.log <server_EFA-Lite_Log_Location>
```

3. Stop and uninstall TPVM.

```
$ tpvm stop
```

```
$ tpvm uninstall
```

4. With the device upgraded to 18s.1.01, execute the **efa deploy** command.

```
$ efa deploy
```

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

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

6. Restore the EFA-Lite 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-Lite_DB_Location> /var/efa/
```

```
$ scp <server_EFA-Lite_Log_Location> /var/log/efa/
```

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

7. Verify that EFA-Lite is running and verify the version.

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

The EFA-Lite version should be 1.1.0.

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

Limitations and Restrictions

This section applies to 18s.1.01.

NPB limitations and restrictions

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

NPB Header stripping

- a. 802.1BR and VN tag are mutually exclusive on an interface.
 - Allowed only in the outer ETH.
- b. MPLS labels can number up to maximum of 4.
- c. ERSPAN stripping – Type 2 is supported. Type 1 is obsolete.
- d. Parser block can parse only up to 128 bytes of ingress frame.
- e. When both 802.1BR/VN-tag and GTP stripping are enabled, only 802.1BR/VN-tag is stripped
- f. When both 802.1BR/VN-tag and MPLS label stripping are enabled, only MPLS labels are stripped
- g. IPv6 SIP and DIP are only 64 bits each (upper or lower).
 - Needs appropriate profile
- h. 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
- i. VLAN add can only add C-VLAN tag.
- j. VLAN add/delete is ignored when GTP strip is enabled.

NPB Flex ACLs

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

Onboard packet capture

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

Internal loopback

- a. 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.
- b. 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.
- c. 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.

Non-NPB limitations and restrictions

This section applies to 18s.1.01a.

End Point Tracking:

It is not suggested to enable spanning tree protocol on the EPT enabled port.

ACL is not supported on the EPT enabled port.

Defects

Defects 18s.1.01b

TSBs—Critical issues to consider prior to installing this release

Technical Support Bulletins (TSBs) provide detailed information about high priority defects or issues present in a release. The following sections specify all current TSBs that have been identified as being a risk to or resolved with this specific release. Please review carefully and refer to the complete TSB for relevant issues prior to migrating to this version of code. Refer to “Contacting Extreme Technical Support” at the beginning of this document.”

Closed with or without code changes in SLX-OS 18s.1.01a

This section lists software defects with Critical, High, and Medium Technical Severity closed with or without a code change as of June 21, 2019.

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-25731	Issue ID:	SLXOS-39971
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-38336	Issue ID:	SLXOS-39975
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:	Use ZTP for both the firmware upgrade as well as replaying the config file.		

Parent Defect ID:	SLXOS-26345	Issue ID:	SLXOS-39980
Severity:	S3 - Medium		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18s.1.01	Technology:	Software Installation & Upgrade

Symptom:	After upgrade from 17s build to 18s build, sometimes "show tpvm status" displays run time environment error while TPVM itself is working good.
Condition:	Upgrade from 17s build to 18s build
Workaround:	A workaround is to uninstall TPVM before upgrade.

Parent Defect ID:	SLXOS-40749	Issue ID:	SLXOS-40782
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLXOS 18s.1.01b	Technology:	VLAN - Virtual LAN
Symptom:	switch may undergo unexpected reload under rare conditions		
Condition:	EPT is enabled with reauthentication timer configured		
Workaround:	Do not configure re-authentication timer with EPT		
Solution:	This is not seen without reauth timer configured. Please validate		

Parent Defect ID:	SLXOS-40907	Issue ID:	SLXOS-40907
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLXOS 18s.1.01b	Technology:	VLAN - Virtual LAN
Symptom:	Endpoint tracking configuration persists even after disabling		
Condition:	Endpoint tracking configuration is disabled when Port channel is down		
Workaround:	Enable and disable Endpoint tracking with Port channel up		