

SLX-OS 18x.1.00b for SLX 9030-48S and SLX 9030-48T

Release Notes v1.0

August 2019

9035843-00 Rev AA

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
Overview	6
Software Features.....	7
Important Notes.....	17
Hardware Support.....	19
Software Upgrade and Downgrade	22
Limitations and Restrictions.....	23
Defects	26

Document history

Version	Summary of changes	Publication date
1.0	Initial Release	August 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.

Overview

Supported Platforms

SLX-OS 18x.1.00b release supports the following SLX platforms:

- SLX 9030-48S: 48 x 10/1G + 4 x 100/40G
- SLX 9030-48T: 48 x 10/1G/100M + 4 x 100/40G

Supported Modes

Standalone bare-metal mode

SLX 9030-48S and SLX 9030-48T are fixed 1U switches that are used as the datacenter leaf switches. SLX 9030-48S provides 48x10/1G fiber ports for server connectivity along with 4x100/40G uplink ports. SLX 9030-48T provides 48x10/1G/100M copper ports for server connectivity along with 4x100/40G uplink ports. These switches are released as a part of SLX-OS 18x.1.00b.

- High density 40G/100G spine-leaf connection
- Native 10/G/1G/100M server connectivity at the leaf
- High performance VXLAN routing
- Port-to-port Latency: 1 microsec

Software Features

New software features in 18x.1.00b

The following software features are supported in 18x.1.00a:

- Endpoint tracking
- IP Directed Broadcast

New software features in 18x.1.00

The following software features are supported in 18x.1.00:

- Layer 2/Layer 3 Forwarding, VLAN, LLDP, ARP/NDP, STP/MSTP, OSPF, BGP, VRRP, and Default VRF only
- Layer 2 and Layer 3 ACL
- LAG
- SPAN/SFlow
- Zero Touch Provisioning (ZTP) (No reboot required after ZTP)
- SNMP/REST
- Telemetry Streaming for CPU/Memory/Interface Statistics
- Shaper/Scheduler, COS/DSCP Remark – Port and Flow based
- QoS – Policers
- ACL on LAG
- Thermal Policy
- Layer 3 Scale (Host Routes only)
- Platform
 - All basic platform features
 - Breakout
 - Interface LED (new)
- VxLAN
 - Layer 2 Gateway
 - Router port and untagged VLAN VW
- User defined VRF
- Static Anycast Gateway (SAG)/Unnumbered Interface

- MCT
- Operational Diag
- ARP Suppression
- Conversational ARP
- ICMP Rate-limiting (IP Stack)
- Layer 3 MCT
- Layer 3 VxLAN
- DHCP Relay
- Multicast IGMPv1v2v3 Snooping Control plane and Data Path support
- QoS – WRED
- LVTEP
- Source Suppression with LVTEP and ARP Suppression

Unsupported features in the 18x.1.00a release:

- BFD
- MPLS
- VPLD (VLL)
- MVRP
- ISIS
- GuestVM
- UDLD

The following solution is supported in SLX-OS 18x.1.00a:

- EVPN VxLAN based Network Virtualization Overlay

BGP-EVPN (VxLAN) – EVPN VxLAN based Network Virtualization Overlay

BGP eVPN Network Virtualization is a controller-less architecture that simplifies data center operations by leveraging open, standards-based protocols to abstract network control plane, data plane, and automation functions from the underlying physical platforms. As an integral part of the Extreme open data center design stack elements, BGP eVPN Network Virtualization builds upon underlying infrastructure platforms, fabrics, and automation to deliver simplified and secure network operations.

The following table lists the set of new features coming in SLX-OS 18x.1.00

Feature Name	Feature Description
BGP eVPN	Standards based, Controller-less Network Virtualization Overlays with VxLAN encapsulation. Provides automatic VxLAN tunnel end point discovery, end host MAC and MAC-IP learning over the control plane.
ARP Suppression	Suppress/reduce the ARP broadcast traffic in an IP fabric.
Static Anycast Gateway	Static Anycast Gateway allows configuring Static Anycast MAC as gateway for multiple tenant systems in a virtualized data center fabric. Same Gateway address is configured across all TORs for a given Tenant/VLAN combination, thus enabling seamless VM mobility across the leaf switches in an IP Fabric deployment without any need for host gateway configuration changes.
Conversational ARP	ARP entries for active conversations only (helps optimize ARP table size)
IP Unnumbered Interfaces	Reduces consumption of IP Address space. Leaf to spine inter-switch point-to-point L3 links are configured as ip unnumbered (/31 subnets) to conserve IP addresses and optimize hardware resources.
L2 VNI capability	The L2VNI is the MAC/NVE mapping table
L3 VNI	The L3VNI is IP prefix/NVE mapping table
Dynamic tunnel (VxLAN) discovery	Supports Dynamic Tunnel discovery using BGP EVPN.
Cluster Management	Configuration management between MCT nodes for logical VTEP is supported.
Manageability, Monitoring, Debugging	NetConf, RESTful API provisioning, VRF support for Telnet/SNMP/SSH, VxLAN tunnel traffic statistics, Show/debug commands

Endpoint tracking

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.

Endpoint Tracking

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 nonauthenticated MAC addresses supported for each endpoint-tracking-enabled port is 8000 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.

RADIUS VSA support

This feature supports the following VSA format.

VSA format

Type (26)	Length	Vendor ID (Extreme Vendor ID)	
Vendor ID (cont'd)		Sub-attr (Sub-type)	Sub-length
Value			

The VSA fields are as follows.

VSA fields

Field	Description
Type	8-bit field. Always 26 as it represents a VSA attribute.
Length	8-bit field. Length of entire attribute, including type and length fields.
Vendor ID	4 octets, encoding the Extreme Vendor ID.
Sub-attr (Sub-type)	8-bit field, indicating class of command.
Sub-length	8-bit field. Length of the Sub-attr (attribute), including sub-type and sub-length fields.
Value	A string, integer, or IP address based on Sub-type

As a RADIUS Access response, only one VSA attribute is provided, for Egress-vlan. The Egress-vlan VSA format has sub-type 216 and a value field of type integer. The format is shown below.

Egress-vlan VSA format

Type (26)	Length	Vendor ID (Extreme Vendor ID)	
Vendor ID (cont'd)		Sub-type (216)	Sub-length
egress VLAN_ID			

For MAC reauthentication, RADIUS sends a Change of Authorization (CoA) of Code 43 with the VSA. This has a sub-type of 1 and the Value field is the string "subscriber:command=reauthenticate". The format is shown below.

MAC reauthentication format

Type (26)	Length	Vendor ID (Extreme Vendor ID)	
Vendor ID (cont'd)		Sub-type (1)	Sub-length
subscriber:command=reauthenticate			

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
```

By default, reauthentication for each session on the port is disabled. However, you can optionally set a timer value for reauthentication, by means of the endpoint-tracking timeout reauth-period command, as in the following example.

```
device(conf-if-eth-0/1)# endpoint-tracking timeout reauth-period 86400
```

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 IVID.
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 SLXX

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 detail	Displays detailed VLAN information.

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.

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.

Consolidated features in SLX-OS 18x.1.00a

Following table lists the features present in SLX-OS 18x.1.00a.

Layer 2 Switching	
<ul style="list-style-type: none">• Layer 2 Access Control Lists (ACLs)• Address Resolution Protocol (ARP) RFC 826• IGMP v1/v2/v3 Snooping• MAC Learning and Aging• Link Aggregation Control Protocol (LACP) IEEE 802.3ad/802.1AX• Virtual Local Area Networks (VLANs)• VLAN Encapsulation 802.1Q• BD Support• Per-VLAN Spanning Tree (PVST+/PVRST+)• Rapid Spanning Tree Protocol (RSTP) 802.1w• Multiple Spanning Tree Protocol (MSTP) 802.1s	<ul style="list-style-type: none">• Pause Frames 802.3x• Static MAC Configuration• Multi-Chassis Trunking (MCT)• VXLAN extension tunnels• Overlay services: overlay gateway instances, overlay transit instances (on spine nodes)• Endpoint tracking (EPT)
Layer 3 Routing	
<ul style="list-style-type: none">• Border Gateway Protocol (BGP4+)• DHCP Helper• OSPF v2/v3• Static routes• IPv4/v6 ACL• Route Policies• 32-Way ECMP• VRF Lite• VRF-aware OSPF, BGP, VRRP, static routes• VRRP v2 and v3• Anycast Gateway over VxLAN	<ul style="list-style-type: none">• IPv4/IPv6 dual stack• ICMPv6 Route-Advertisement Guard IPv6 ACL packet filtering• BGP-Allow AS• IPv6 routing• Multi-VRF• L3 over Bridge Domains (BD)
Automation and Programmability	
<ul style="list-style-type: none">• gRPC Streaming protocol and API• REST API with YANG data model• Python	<ul style="list-style-type: none">• DHCP automatic provisioning• NETCONF API
Quality of Service	

<ul style="list-style-type: none"> • ACL-based QoS • Two Lossless priority levels for QoS • Class of Service (CoS) IEEE 802.1p • DSCP Trust • DSCP to Traffic Class Mutation • DSCP to CoS Mutation • DSCP to DSCP Mutation • CoPP (Control Plane Policing) 	<ul style="list-style-type: none"> • Random Early Discard • Per-port QoS configuration • ACL-based Rate Limit • Dual-rate, three-color token bucket • ACL-based remarking of CoS/DSCP/Precedence • ACL-based sFlow • Scheduling: Strict Priority (SP), Deficit Weighted Round-Robin (DWRR)
Management and Monitoring	
<ul style="list-style-type: none"> • Zero-Touch Provisioning (ZTP) • IPv4/IPv6 management • Industry-standard Command Line Interface (CLI) • NETCONF API • REST API with YANG data model • SSH/SSHv2 • Link Layer Discovery Protocol (LLDP) IEEE 802.1AB • MIB II RFC 1213 MIB • Syslog (RASlog, AuditLog) • Management VRF • Switched Port Analyzer (SPAN) • Telnet 	<ul style="list-style-type: none"> • SNMP v1, v2C, v3 • sFlow version 5 • Out-of-band management • RMON-1, RMON-2 • NTP • Management Access Control Lists (ACLs) • Role-Based Access Control (RBAC) • Range CLI support • Python • DHCP Relay • SLX-OS and Linux Shell Interoperability
Security	
<ul style="list-style-type: none"> • RADIUS – Authentication and Authorization • AAA • TACACS+ • Secure Shell (SSHv2) • TLS 1.1, 1.2 • HTTP/HTTPS 	<ul style="list-style-type: none"> • BPDU Drop • Secure Copy Protocol • Control Plane Protection • SFTP
IP Fabric	
<ul style="list-style-type: none"> • Controllerless Network Virtualization (BGP-EVPN) • ARP suppression • Conversational ARP • Static Anycast Gateway 	<ul style="list-style-type: none"> • Logical VTEP (Static and EVPN) • IP Un-numbered interface • RIOT (Routing In and Out of Tunnel) (v4 and v6)
Platform	
<ul style="list-style-type: none"> • 1G/10G/40G/100G Auto speed detection • Multi Speed(1G/10G) Optic Support 	40G 4x10G /100G 4x25 Breakout Support

Important Notes

Zero Touch Provisioning (ZTP)

- ZTP is enabled by default on SLX switches from factory or by “write erase”. Upon switch power-on or reboot by “write erase”, it will automatically connect to DHCP server through both management interface and inband ports with connection for firmware to download and configuring the switch based on the DHCP configuration.
- If the switch does not have a DHCP server connected or the DHCP server is not configured for ZTP, the switch will keep searching the DHCP server for ZTP.

The serial console of the switch will display ZTP message as following:

```
ZTP, Sat Nov 17 07:55:37 2018, ===== ZTP start =====  
ZTP, Sat Nov 17 07:55:37 2018, disable raslog  
ZTP, Sat Nov 17 07:55:37 2018, CLI is ready  
ZTP, Sat Nov 17 07:55:49 2018, inband ports are enabled  
ZTP, Sat Nov 17 07:55:49 2018, serial number = 1818N-41522  
ZTP, Sat Nov 17 07:55:49 2018, model name = EN-SLX-9030-48S  
ZTP, Sat Nov 17 07:55:49 2018, use both management interface and inband interfaces  
ZTP, Sat Nov 17 07:55:49 2018, checking inband interfaces link status  
ZTP, Sat Nov 17 07:56:43 2018, find link up on interfaces: eth0 Eth0.1 Eth0.9 Eth0.10 Eth0.11  
ZTP, Sat Nov 17 07:56:43 2018, start dhcp process on interfaces: eth0 Eth0.1 Eth0.9 Eth0.10  
Eth0.11  
ZTP, Sat Nov 17 07:56:53 2018, get no dhcp response from all interfaces  
ZTP, Sat Nov 17 07:56:53 2018, retry in 10 seconds  
ZTP, Sat Nov 17 07:57:03 2018, inband ports are enabled  
ZTP, Sat Nov 17 07:57:03 2018, serial number = 1818N-41522  
ZTP, Sat Nov 17 07:57:03 2018, model name = EN-SLX-9030-48S  
ZTP, Sat Nov 17 07:57:03 2018, use both management interface and inband interfaces  
ZTP, Sat Nov 17 07:57:03 2018, checking inband interfaces link status  
ZTP, Sat Nov 17 07:57:04 2018, find link up on interfaces: eth0 Eth0.1 Eth0.6 Eth0.9 Eth0.10  
Eth0.11  
ZTP, Sat Nov 17 07:57:04 2018, start dhcp process on interfaces: eth0 Eth0.1 Eth0.6 Eth0.9  
Eth0.10 Eth0.11
```

ZTP, Sat Nov 17 07:57:14 2018, get no dhcp response from all interfaces

ZTP, Sat Nov 17 07:57:14 2018, retry in 10 seconds

You need to login onto the serial console, wait for the above message to show up to confirm ZTP has been triggered, and then run “dhcp ztp cancel” and “reload system” to cancel the ZTP operation.

```
SLX# dhcp ztp cancel
```

```
ZTP is canceled.
```

```
SLX#
```

```
SLX#
```

```
SLX# dhcp ztp cancel
```

```
ZTP is not enabled.
```

```
SLX#
```

After dhcp ztp cancel, reload is not required.

```
SLX# SLX# reload system
```

```
Warning: This operation will cause the chassis to reboot and  
requires all existing telnet, secure telnet and SSH sessions to be  
restarted.
```

```
Unsaved configuration will be lost. Please run `copy running-config startup-config` to save the  
current configuration if not done already.
```

```
Are you sure you want to reboot the chassis [y/n]? y
```

```
[ 940.360081] VBLADE: vblade_control: FEPORTS_DISABLE
```

```
xpDma::quiesce:307 devId=0
```

```
xpDriverWrapper::quiesce:146 devId=0
```

```
FABOS_BLADE_MSG_BL_DISABLE received in HSLUA for chip 0
```

```
2017/03/27-21:14:13, [RAS-1007], 567,, INFO, SLX9030, System is about to reload.
```

Hardware Support

SLX 9030 Hardware and License SKUs

SKU	Description
EN-SLX-9030-48S-4C-AC-F	Extreme SLX 9030-48S Switch AC with Front to Back Airflow, Supports 48x10GE/1GE + 4x100GE/40GE
EN-SLX-9030-48S-4C-AC-R	Extreme SLX 9030-48S Switch AC with Back to Front Airflow, Supports 48x10GE/1GE + 4x100GE/40GE
EN-SLX-9030-48S-4C	Extreme SLX 9030-48S Switch with No Power supplies, no fans, Supports 48x10GE/1GE + 4x100GE/40GE
EN-SLX-9030-48T-4C-AC-F	Extreme SLX 9030-48T 10GBaseT Switch AC with Front to Back Airflow, Supports 48x10GE/1GE + 4x100GE/40GE
EN-SLX-9030-48T-4C-AC-R	Extreme SLX 9030-48T 10GBaseT Switch AC with Back to Front Airflow, Supports 48x10GE/1GE + 4x100GE/40GE
EN-SLX-9030-48T-4C	Extreme SLX 9030-48T 10GBaseT Switch with No Power supplies, no fans, Supports 48x10GE/1GE + 4x100GE/40GE
EN-SLX-9030-ADV-LIC-P	SLX 9030 Advanced Feature License for BGP-EVPN, gRPC

Supported power supplies

The following table lists the power supplies that are available for the devices supported in this release:

SKU	Description
17115	Fan module, Front to Back airflow
17116	Fan module, Back to Front airflow
10960	770W AC power supply, Front -to-Back airflow
10961	770W AC power supply, Back-to-Front airflow
10962	1100W DC power supply, Front -to-Back airflow
10963	1100W DC power supply, Back-to-Front airflow

Supported optics

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

Description	SKU	MFG Part Number
1000Base-SX SFP optic, MMF, LC connector, Optical Monitoring Capable	1G-SFP-SX-OM	33210-100
1000Base-LX SFP optic, SMF, LC connector, Optical Monitoring Capable	1G-SFP-LX-OM	33211-100
1000BASE-TX SFP Copper, RJ-45 Connector	1G-SFP-TX	33002-100
1GE COPPER SFP,1-PK, ROHS	1G-SFP-000190	57-1000042-02
10GBASE-USR, SFP+ optic (LC), target range 100m over MMF,	10G-SFP-USR	57-1000130-01
10GBASE-USR, SFP+ optic (LC), target range 100m over MMF 70C TAA COMPLIANT	10G-SFP- USR-SA	57-1000343-01
10GBASE-SR, SFP+ optic (LC), target range 300m over MMF , 80C	10G-SFP-SR	57-0000075-01
10GBASE-SR,SFP+ MMF LC CONNECTOR , 70C	10G-SFP-SR-S	57-1000340-01
10GBASE-SR, SFP+ optic (LC), target range 300m over MMF 70C-TAA COMPLIANT	10G-SFP-SR-SA	57-1000344-01
10GBASE-LR, SFP+ optic (LC), for up to 10km over SMF	10G-SFP-LR	57-0000076-01
10GBASE-LR,SFP+ SMF LC CONNECTOR (No TAA), 70C	10G-SFP-LR-S	57-1000341-01
10GBASE-LR, SFP+ optic (LC), for up to 10km over SMF 70C - TAA COMPLIANT	10G-SFP-LR-SA	57-1000345-01
10GBASE-ER SFP+ optic (LC), for up to 40km over SMF	10G-SFP-ER	57-0000058-01
10GBASE-ZR SFP+ optic (LC), for up to 80km over SMF	10G-SFP-ZR	57-1000180-01
10GBASE-ZRD SFP+ optic (LC), for up to 80km over SMF	10G-SFP-ZRD-T	57-1000266-01
ACTIVE DIRECT ATTACHED SFP+ COPPER, 1MTR, 1-PK	10G-SFP-TWX-0101	58-1000026-01
ACTIVE DIRECT ATTACHED SFP+ COPPER, 3MTR,1-PK	10G-SFP-TWX-0301	58-1000027-01
ACTIVE DIRECT ATTACHED SFP+ COPPER, 5MTR,1-PK	10G-SFP-TWX-0501	58-1000023-01
PASSIVE DIRECT ATTACHED SFP+ COPPER, 1MTR, 1-PK	10G-SFP-TWX-P-0101	58-1000024-01
PASSIVE DIRECT ATTACHED SFP+ COPPER, 3MTR,1-PK	10G-SFP-TWX-P-0301	58-1000025-01
PASSIVE DIRECT ATTACHED SFP+ COPPER, 5MTR,1-PK	10G-SFP-TWX-P-0501	58-1000019-01
10GE SFP+ Direct Attached Active Optical Cable, 7m, 1-pack	10G-SFP-AOC-0701	57-1000273-01
10GE SFP+ Direct Attached Active Optical Cable, 10m, 1-pack	10G-SFP-AOC-1001	57-1000274-01
10GBASE-T SFP+	10338	908711-10
40GBASE-SR4 QSFP+ optic (MTP 1x8 or 1x12), 100m over MMF, 1-pack	40G-QSFP-SR4-1	57-1000128-01
40GBASE-SR4 QSFP+ optic (MTP 1x8 or 1x12), 100m over MMF, compatible with 10GBASE-SR, 10G breakout-capable, 1- pack	40G-QSFP-SR4-INT	57-1000129-01
40GBase-LR4 QSFP+ optic (LC), for up to 10km over SMF, 1- pack	40G-QSFP-LR4-1	57-1000263-01

40GBase-LR4 QSFP+ to 4 SFP+ optic (LC), for up to 10km over SMF, 1-pack	40G-QSFP-LR4-INT	57-1000477-01
40GBASE-LM4 QSFP+, 1310nm, 160m over duplex LC OM4 MMF, 2km over duplex LC SMF	40G-QSFP-LM4	57-1000325-01
40GBase-ER4 QSFP+ optic (LC), for up to 40km over SMF	40G-QSFP-ER4-1	57-1000327-01
40GE SR QSFP+ optic (LC), Bidirectional, 100m over OM3 MMF	40G-QSFP-SR-BIDI	57-1000339-01
4x10GE QSFP+ to 4 SFP+ Active Copper Cable 1 M	40G-QSFP-4SFP-C-0101	58-0000051-01
4x10GE QSFP+ to 4 SFP+ Active Copper Cable 3 M	40G-QSFP-4SFP-C-0301	58-0000052-01
4x10GE QSFP+ to 4 SFP+ Active Copper Cable 3 M	40G-QSFP-4SFP-C-0501	58-0000053-01
40GE Direct Attached QSFP+ to QSFP+ Active Copper cable, 1m, 1-pack	40G-QSFP-QSFP-C-0101	58-0000041-01
40GE Direct Attached QSFP+ to QSFP+ Active Copper cable, 3m, 1-pack	40G-QSFP-QSFP-C-0301	58-0000042-01
40GE Direct Attached QSFP+ to QSFP+ Active Copper cable, 5m, 1-pack	40G-QSFP-QSFP-C-0501	58-0000043-01
40GE QSFP Direct Attached Passive Copper Cable, 1m, 1-pack	40G-QSFP-QSFP-P-0101	58-0000033-01
40GE QSFP Direct Attached Passive Copper Cable, 5m, 1-pack	40G-QSFP-QSFP-P-0501	58-0000035-01
40GE Direct Attached QSFP+ to QSFP+ Active Optical Cable, 10m, 1-pack	40G-QSFP-QSFP-AOC-1001	57-1000306-01
100 GbE QSFP28 optic (LC), LR4 low power, for distances up to 10 km over SMF	100G-QSFP28-LR4-LP-10KM	57-1000338-01
100 GbE QSFP28 optic (LC), LR4-lite, for distances up to 2 km over SMF	100G-QSFP28-LR4L-2KM	57-1000329-01
100 GbE QSFP28 optic (MTP 1x12), SR4, for distances up to 100 m over MMF	100G-QSFP28-SR4	57-1000326-01
100GBASE-ESR4 QSFP+ optic (MTP 1x8 or 1x12), eSR4, for distances up to 300 m over MMF	100G-QSFP-ESR4	57-1000352-01
100 GbE QSFP28 optic (LC), CWDM4, for distances up to 2 km over SMF	100G-QSFP28-CWDM4-2KM	57-1000336-01
100GE Direct Attached QSFP+ to QSFP+ Active Optical Cable, 10m, 1-pack	100G-QSFP-QSFP-AOC-1001-10m	57-1000347-01
100GE Direct Attached QSFP-28 to QSFP-28 Passive Copper cable, 1m	100G-QSFP-QSFP-P-0101	58-0000044-01
100GE Direct Attached QSFP-28 to QSFP-28 Passive Copper cable, 3m	100G-QSFP-QSFP-P-0301	58-0000045-01

Software Upgrade and Downgrade

Image file names

Download the following images from www.extremenetworks.com.

Image file name	Description
slxos18x.1.00b.tar.gz	SLX-OS 18x.1.00b software
slxos18x.1.00b_all_mibs.tar.gz	SLX-OS 18x.1.00b MIBS
slxos18x.1.00b.md5	SLX-OS 18x.1.00b md5 checksum

Migration path

Recommended upgrade/downgrade migration path.

To \ From	SLX 18x.1.00_CR	SLX 18x.1.00	SLX18x.1.00a
SLX 18x.1.00_CR	NA	default-config	default-config
SLX 18x.1.00	default-config	NA	coldboot
SLX 18x.1.00a	default-config	coldboot	NA

NOTE: For MCT it is recommended to upgrade one node at a time.

Limitations and Restrictions

IGMPv2 snooping

- Unknown multicast packets on IGMP snooping enabled vlans are dropped
- In IGMPv3, only Include mode is supported

IP Fabric

- A few IP Anycast gateways are not active after save and reload of the system.
- Sometimes moving L3VNI from Vlan to BD impacts traffic forwarding
- SAG scale is supported for 256 VE interfaces.
- When static route next hop is pointing to a Tunnel nexthop, it is not installing the route correctly and results in traffic going to wrong tunnel after trigger like reboot
- OSPF neighborhood is not establishing over VXLAN tunnel
- Arp suppression is only supported on vlans 1-512. Vlan 513-4096 and BD 1-1024 will not be able to support arp-suppression.

Layer 2

- **System supports maximum scale of 2k mac-addresses on the system.**
- Few MAC addresses are not learnt when traffic is sent at scale (64K mac addresses) due to hash collisions.
- Packets get transiently duplicated when member interface in port channel is toggled.
- Traffic not load balanced on LAG when destination IP gets incremented and hashing is set only for dst-ip.
- Bridge-domain command does not accept range.
- "show system internal vxlan" command is not working.
- show mac-address count interface port-channel <>" is showing wrong "remote address count".
- Point to Point (P2P) Bridge Domain is not supported.

Layer 3

- **VRRP**
 - Unable to configure virtual IP with VRRP v3 if the ipv4 prefix length is /31
- **IP**
 - "show interface loopback <>" displays incorrect mtu value.
 - "Duplicate IP address with null mac messages detected for static arp configured while enabling STP"
- **BGP:**
 - Under vrf AF, "no export map" is accepted even though there is no <cr>.
 - Under vrf AF, "no export map", the 'no' operation of 'export map' command does not work

- When the max-route limit for v4/v6 prefixes in a VRF is removed, it is not dynamically updating the route limit, nor it is asking user to clear the route table for the config to take effect
 - **OSPF:**
 - Connected, OSPF routes exported to BGP EVPN are not working with export-map

MCT

- Traffic flooding after cluster deploy/no-deploy on high scale.
- After clear mac dynamic - Mac learning is very slow and also causes MAC still present in HW.
- When traffic received from MCT client switches from MCT node 1 to node 2, the source MAC address on MCT node 2 does not change from remote(CCR) to local (Dynamic-CCL). This will remain until the Dynamic-CCL MAC learnt on the remote MCT node ages out. After that, the CCR MAC will be updated to Dynamic-CCL as expected. During this transition, there could be a brief traffic impact at high scale. This happens when traffic first hashes to MCT node 1, then switches to Node 2. Since CCR also points to the local client, the brief traffic impact is only seen if the MAC on MCT node 1 ages out.

NetConf

- When "clear ip bgp nei all" executed multiple times may cause process instability.

BFD

- BFD is not supported

Platform

- "Error:soc_reg_above_64_field32_read(IQM_OCCUPIED_BD_COUNTERr) failed for unit 0 core 0 " is seeing on console when collecting SS.
- Custom RPC call "show-ntp" is providing empty output instead of the active NTP server
- Account log is showing SSH instead of unique tty identifier while login through SSH; after login, it's behaving correct.
- Login message says " admin connected using console" instead of "telnet" when accessing the switch via telnet
- NTP is using wrong ip as source ip to send ntp packet out of OOB mgmt-interface after sh/no shut of mgmt.
- Switches are taking default option as 'y', instead of 'no input' for 'y/n' prompt for 'reload system', 'firmware" etc commands.
- Issues with special-characters in password.
 - Dollar sign (\$), double-quote sign ("), and single-quote (') are not supported by the firmware download command.
 - Double-quote (") is not supported the copy support command.
 - Single-quote (') is not supported by the copy config command.
- The "copy <file> running-config" will always return success if you only specify the configuration file path but not the file name.

Port Mirroring (SPAN)

- "sflow-profile" command under policy-map is not supported.

QOS

- "show qos flow-control interface all" shows port-channel member interfaces two times in the output.
- L2 QOS egress maps not supported.
- Flow Control not supported.
- Priority flow control not supported.
- Cee-map cli not supported.

REST API

- Account log is showing wrong hostname while login through REST (HTTP).

Endpoint Tracking

- With Endpoint Tracking enabled on switchports, do not enable STP.
- Issues are seen with higher scales.
- ACLs are not supported on Endpoint Tracking enabled ports.

Telemetry Streaming

- Unreachable GRPC Client causes in stability in the system.

Defects

SBs—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 for SLX-OS 18x.1.00a

This section lists software defects with Critical, High, and Medium Technical Severity closed with a code change as of August 2019 in SLX-OS 18x.1.00a.

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

Parent Defect ID:	SLXOS-38336	Issue ID:	SLXOS-38336
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-39126	Issue ID:	SLXOS-39126
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLXOS 18x.1.00a	Technology:	Other

Symptom:	When mac moves from EPT enabled CEP port to remote EPT enabled CCEP member port , then on local node port vlan membership of CEP port and vlan is not removed .
Condition:	MCT topology with EPT enabled ports and mac move
Workaround:	None

Parent Defect ID:	SLXOS-39222	Issue ID:	SLXOS-39222
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLXOS 18x.1.00a	Technology:	Other
Symptom:	In EPT MCT case when switch is reloaded with traffic on , sometimes mac's are not synced across MCT.		
Condition:	MCT with EPT enabled ports		
Workaround:	None		

Parent Defect ID:	SLXOS-39233	Issue ID:	SLXOS-39233
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLXOS 18x.1.00a	Technology:	Other
Symptom:	In some cases , mac remain in pending state in hardware .		
Condition:	MCT with EPT enabled ports		
Workaround:	None		

Parent Defect ID:	SLXOS-39238	Issue ID:	SLXOS-39238
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Layer 2 Switching
Reported in Release:	SLXOS 18x.1.00a	Technology:	Other
Symptom:	In some cases , pending mac is not deleted from hardware , which does not allow relearning of the mac		
Condition:	MCT with EPT ports.		
Workaround:	None		

Parent Defect ID:	SLXOS-39220	Issue ID:	SLXOS-39312
Severity:	S1 - Critical		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18s.1.01a	Technology:	SNMP - Simple Network Management Protocol
Symptom:	LLDP-MIB::lldpLocPortId value is not correct (appears corrupted) when queried via SNMP GET operation.		
Condition:	Issue occurs only for SNMP GET operation (on LLDP-MIB::lldpLocPortId). SNMP GET-NEXT and snmpwalk returns correct values.		
Workaround:	<ol style="list-style-type: none"> 1. Use SNMP GET-NEXT or snmpwalk instead of SNMP GET when querying LLDP-MIB::lldpLocPortId via SNMP. 2. Use CLI to query (LLDP-MIB::lldpLocPortId) instead of SNMP, if it's feasible. 		

Parent Defect ID:	SLXOS-25731	Issue ID:	SLXOS-39972
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-40476	Issue ID:	SLXOS-40548
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Management
Reported in Release:	SLXOS 18r.2.00	Technology:	High Availability
Symptom:	During DOS attacks, flood of disable pam_unix log messages are seen on console		
Condition:	DOS attacks on system		
Workaround:	Configure syslog server to redirect these messages		

Parent Defect ID:	SLXOS-40907	Issue ID:	SLXOS-40908
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		

Parent Defect ID:	SLXOS-42679	Issue ID:	SLXOS-42679
Severity:	S2 - High		
Product:	SLX-OS	Technology Group:	Other
Reported in Release:	SLXOS 18x.1.00a	Technology:	Other
Symptom:	Unexpected reload of device		
Condition:	Management cluster is in broken state.		
Workaround:			