Highlights

- Delivers agility at all layers of the data center stack
- Provides high-density 100/40 GbE spine and leaf connectivity in a 1U fixed form factor
- Includes a programmable ASIC to accelerate adoption of new protocols and technologies
- Utilizes the Extreme SLX Insight Architecture and Extreme SLX Visibility Services for flexible, real-time monitoring of virtualized, dynamic workloads to streamline troubleshooting
- Provides payload timestamping to more accurately set and measure performance SLAs
- Incorporates turnkey and customizable cross-domain workflow automation for the entire network lifecycle through Extreme Workflow Composer and network automation suites

ExtremeSwitching™ SLX 9240

Programmable, Flexible, High-Density Switch

As data centers and cloud service providers embrace new high-performance servers and distributed applications, they increasingly need dense 100/40 GbE switches for leaf and spine configurations. Traditionally, infrastructure has been slow to evolve, and it can be a barrier to innovation. With flexibility at all layers of the data center stack, IT teams can drive agility. The ExtremeSwitching SLX 9240 Switch is designed to help organizations stay ahead of this application- and data-driven network transformation without compromising performance.

Programmable Switching Platform

The SLX 9240 delivers the high-density 100 GbE connectivity required by high-end enterprise and cloud data centers. The underlying hardware is programmable, enabling a faster transition to emerging protocols and new technologies. Workload visibility combined with end-to-end network visibility helps infrastructure teams continue to improve SLAs as they increase network virtualization. In addition, the SLX 9240, coupled with Extreme Workflow Composer™ and Extreme Workflow Composer Automation Suites, simplifies end-to-end network management, including turnkey provisioning, validation, and troubleshooting of workflows. And with Extreme Management Center (XMC), users gain centralized management, automation, analytics, visibility and control from the wireless edge to the datacenter.

SLX 9240 Overview

The SLX 9240 is a fixed 100/40 GbE leaf and spine switch in a 1U form factor that supports 24 MB of dynamically shared packet buffer and an overall throughput of 3.2 Tbps/1.3 Bpps. It can be configured as a leaf switch or connect to 40 or 100 GbE uplinks from leaf switches, such as the ExtremeSwitching SLX 9140.

By leveraging this high-density switch, data center networks can dramatically improve power, space, and cooling efficiencies, even at scale. A programmable ASIC enables the adoption of new protocols and technologies through an OS, rather than a forklift upgrade. Payload timestamping improves the accuracy of performance SLA setting and measurement.

Modular, Virtualized Operating System

The SLX 9240 runs Extreme SLX-OS, a fully virtualized Linux-based operating system that delivers process level resiliency and fault isolation. The SLX-OS supports advanced switching features and is highly programmable with support for REST API with the YANG data model, Python, and NETCONF—enabling full lifecycle automation with Workflow Composer. It is based on Ubuntu Linux, which offers all the advantages of open source and access to commonly used Linux tools.

SLX-OS runs in a virtualized environment over a KVM hypervisor, with the operating system compartmentalized and abstracted from the underlying hardware. The core operating system functions for the SLX 9240 are hosted in the system VM.

This approach provides clean failure domain isolation for the switch operating system while leveraging the x86 ecosystem—thereby removing single-vendor lock-in for system tools development and delivery. In addition, it supports a guest VM, which is an open KVM environment for running third-party and customized monitoring, troubleshooting, and analytics applications.

Embedded Network Visibility

The SLX Insight Architecture and SLX Visibility Services deliver a new approach to network monitoring and troubleshooting that makes it faster, easier, and more cost-effective to obtain the comprehensive, real-time visibility needed for network operations and automation. This innovative approach provides comprehensive visibility from the network to the workload, and triggers network actions. These actions can address enduser application or service needs, and provide context-rich data for additional analysis, automation, and reporting. For details, read Visibility in the Modern Data Center with Extreme Switches and Routers.

SLX Insight Architecture

The SLX Insight Architecture leverages an innovative combination of SLX-OS software and SLX hardware features to provide pervasive visibility into the network without impacting normal network operation or performance. This flexible and open solution enables organizations to deploy their choice of third-party or customized monitoring and troubleshooting tools directly in the network—providing real-time visibility to meet specific business and operational needs across the network. This enables organizations to improve service and application assurance, as well as dramatically reduce operational impact and cost.

As shown in Figure 1, key components of the SLX Insight Architecture include:

- **Guest VM** - The SLX Insight Architecture provides an open KVM environment that runs third-party applications and customized monitoring, troubleshooting, and analytics tools. Enabled by SLX-OS, this preconfigured guest VM is on each SLX 9240 Switch. It hosts third-party network operations and analytics applications on every device, extending visibility to the entire network.
- **Flexible Streaming** - The SLX Insight Architecture provides API streaming, enabling captured data to be delivered to analytics applications off the platform for additional analysis, visualization and reporting, or logging and archiving.
- **Dedicated Analytics Storage** - The SLX 9240 provides 128 GB of on-device storage dedicated to visibility applications running in the guest VM, providing real-time data capture for easy and fast access.

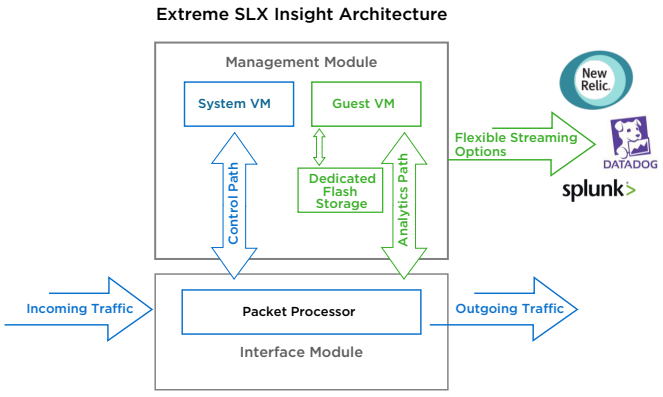

Figure 1: The SLX Insight Architecture

SLX Visibility Services

As network complexity increases, isolated data points at the physical or virtual network layer provide little insight into the criticality of an issue. For example, bursty storage backup traffic slowing down an internal Web site is a lower priority than a slowdown for a revenue-generating application. Network administrators need workload context across the network to ensure the appropriate action is taken in each case.

SLX Visibility Services help simplify network operations with embedded visibility from the physical network to application workloads. By combining physical and virtual network traffic data with overlay and workload information across multiple network layers, this solution enables diverse, rule-based actions to maintain performance and mitigate risk. Other key functions include:

- Pervasive visibility at scale across the network for seamless support of highly distributed multitier application workloads
- Rich multilayer classification (such as IP and MAC addresses, port numbers, VNIs) and workload matching with network-wide scale
- Automated application of rule-based actions (such as count, drop, mirror, sFlow) to incoming network traffic
- Further actions outside the switch, including pushing context-rich data to the SLX Insight Architecture, Workflow Composer, and third-party analytics and monitoring applications

SLX Visibility Services are embedded into SLX switches, reducing the operational complexity of managing network visibility at scale (see Figure 2).

Extreme Visibility Services Figure

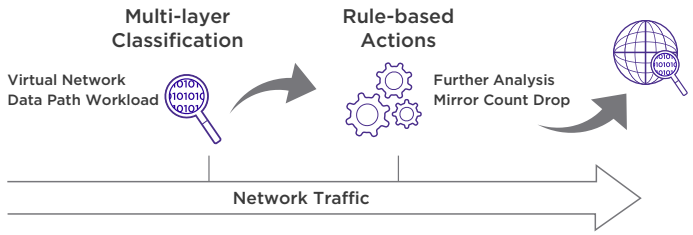

Figure 2: Extreme SLX Visibility Services

SLX Packet Broker for Carrier-Grade Visibility

Network traffic growth in carrier networks is driving scalability and management challenges that service providers must urgently address. Wire-speed visibility into these networks is critical to enable effective control of this traffic as it scales. Part of a broad portfolio of network visibility solutions, the SLX 9240 can be deployed as a Network Packet Broker. From the aggregation layer of the network, SLX 9240 curates traffic and steers replicated flows to relevant analytics platforms to provide (for instance) DDoS mitigation, congestion management, differentiated treatment of voice and video, and billing. The SLX 9240 offers 100 GbE (32), 40 GbE (32), 25 GbE (128) and 10 GbE (128) interface capacities.

For more information on the full portfolio of Extreme Network Packet Broker solutions, visit extremenetworks.com/networkpacketbroker.

Ansible

Ansible Network modules deliver the benefits of simple, powerful, agentless automation to network administrators. Ansible SLX network modules can be used to configure, test and validate existing network state on the SLX family of devices including 9240.

Extreme Management Center for Unified Insight, Visibility and Control

High levels of virtualization, containerization and cloud environments, combined with the enormous traffic, limit visibility in the modern data center. To address that dynamic challenge, SLX switches and routers, can be managed by Extreme Management Center (XMC). XMC includes a suite of applications, empowering administrators to deliver a superior quality experience to users through a single consolidated view and a common set of tools to provision, manage and troubleshoot the network. It works across wired and wireless networks, from the edge to the data center and private cloud. It gives you the granular visibility and real-time analytics, to make data-based business decisions.

XMC provides a consolidated view of users, devices and applications for wired and wireless networks – from datacenter to edge. Zero touch provisioning lets one quickly bring new infrastructure online. A granular view of users, devices and applications with an easy to understand dashboard, enables efficient inventory and network topology management. XMC also provides ecosystem integration, includes off the box integrations with major enterprise data center virtual environments such as VMWare, OpenStack and Nutanix to provide VM visibility and enforce security settings. Get more information on Extreme Management Center.

Cross-Domain Automation for IT Operations

To unleash new levels of business innovation and competitive advantage, many organizations are embracing digital transformation. Their success depends on building an agile business, and, in the digital era, IT agility is achievable only with centralized, cross-domain automation.

Extreme SLX 9240 leverages Extreme Workflow Composer, powered by StackStorm. With its nearly 2,000 pre-built points of integration, this DevOps-inspired, event-driven automation platform enables cross-domain workflows and straightforward integration with disparate IT technologies, platforms, and policies to provide split-second, reliable execution of service provisioning and remediation. Extreme Workflow Composer Automation Suites are specifically designed to speed up time-to-value by providing complete network lifecycle automation. For more details, read the Extreme Workflow Composer Automation Platform At A Glance.

Speed Up Time to Value with Turnkey Automation Suites

As organizations address the primary barrier to IT agility—the network—they need automation that is easy to deploy by operators with limited skills, that delivers value immediately, and that provides more than Day 0 provisioning. Extreme Workflow Composer Automation Suites (Figure 3) provide turnkey, customizable network automation for out-of-box functionality that delivers immediate value to the business, while the workflows provide automation for the entire lifecycle: provisioning, validation, troubleshooting, and remediation. As a result, IT organizations can adopt automation at their own pace, deploy services, resolve issues faster, and eliminate a barrier to IT agility. For details, read the Extreme Workflow Composer Automation Suites At A Glance.

DevOps-Inspired Automation

Streamline end-to-end IT operations and increase IT agility with event-driven, cross-domain automation.

SLX 9240 and Workflow Composer

The SLX 9240, combined with Workflow Composer and the Workflow Composer Automation Suites, delivers automation for provisioning, validation, troubleshooting, and remediation of network services:

- Unleash IT agility by eliminating cross-domain service provisioning, troubleshooting, and remediation delays
- Accelerate time-to-value and time-to-resolution with automation suites designed, built, and tested for Extreme networks; easily customized as skills and requirements change
- Leverage the power of DevOps methodologies and popular open source technologies that embrace industry best practices, as well as a thriving technical community for peer collaboration and innovation
- Increase agility beyond Day 0 by automating the entire network lifecycle—provisioning, validation, troubleshooting, and remediation of Extreme network infrastructure

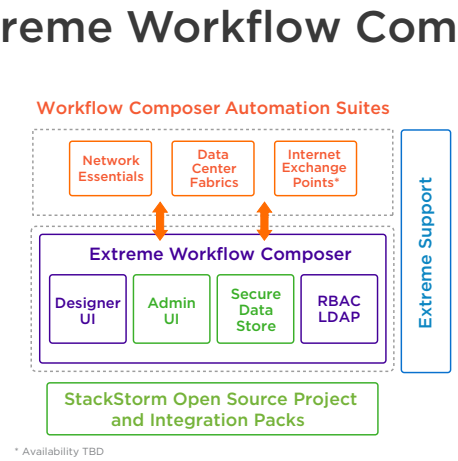

Figure 3: The Extreme Workflow Composer Automation Suite Architecture

SLX 9240 Switch Specifications	
Form Factor	1U
Switching Bandwidth (data rate, full duplex)	3.2 Tbp in and 3.2 Tbps out for a sum total of 6.4 Tbps
Forwarding Capacity (data rate, full duplex)	(L2) 2.4 Bpps, (L3) 600 Mpps line-rate performance
Dimensions and Weight	440 mm; 17.32 in. (Width), 444.7 mm; 17.5 in. (Depth) 43.7mm; 1.72 in. (Height) 9.07 kg; 20 lb
Port-to-port latency	2.5 usec
Architecture Store and Forward	Supported
100/40 GbE Ports	32
Power Supplies	Two internal, redundant, field-replaceable, load-sharing AC or DC power supplies
Cooling Fans	Five field-replaceable fans
Dynamically Shared Packer Buffer	24 MB
Power	
Power Inlet (AC)	C13
Input Voltage	90 V to 264 V or 40.8 V to 60 V DC
Input Line Frequency	47 Hz to 63 Hz
Inrush Current	25 A peak
Maximum Current	12 A/AC, 14 A/DC
Typical Power Consumption	84 W | Two AC PSU, five fan trays, 10% traffic, low fan speed
Maximum Power Consumption	581 W | Two AC PSU, six fan trays, 100% traffic, high fan speed
Power Supply Rated Maximum (AC)	650 W
Switch Power Consumption	DC PSU 563 W; AC PSU 581 W
Environment	
Humidity	5% to 95% at 50°C
Altitude	Up to 3,000 m safety; 60 m to 4,000 m operational
Shock (operational)	20 G, 11 ms, half-sine wave
Vibration (operational)	1 G sine, 0.4 gms random, 5 Hz to 500 Hz
Airflow	134 CFM (estimated with two PSU, six fan trays)
Acoustics (25°C)	52 dBA
MTBF (25°C)	327,539 hours
Software Specifications	
Connector Options	10 GbE SFP+ (via splitter cable) 100 GbE QSFP-28 40 GbE QSFP+ Out-of-band Ethernet management: 10/100/1000 Mbps RJ-45 Console management: RJ45 serial port and USB type-C port with serial communication device class support Storage: USB port, standard-A plug
Maximum MAC Addresses	Up to 48,000
Maximum VLANs	4,096
Maximum Routes (in hardware)	Up to 40,000
Maximum ACLs	512
Maximum Members in a Standard LAG	16
Maximum Per-Port Priority Pause Level	8
Maximum Switches an mLAG Can Span	2
Maximum IPv4 Unicast Routes	48,000
Maximum IPv6 Unicast Routes	8,000
DCB Priority Flow Control Classes	4
Maximum Jumbo Frame Size	10,000 bytes
QoS Priority Queues (per port)	8

IEEE Compliance

- **Ethernet**
  - IEEE 802.1D Spanning Tree Protocol
  - IEEE 802.1s Multiple Spanning Tree
  - IEEE 802.1w Rapid Reconfiguration of Spanning Tree Protocol
  - IEEE 802.3 Ethernet
  - IEEE 802.3ad Link Aggregation with LACP
  - IEEE 802.3ae 10G Ethernet
  - IEEE 802.1Q VLAN Tagging
  - IEEE 802.1p Class of Service Prioritization and Tagging
  - IEEE 802.1v VLAN Classification by Protocol and Port
  - IEEE 802.1AB Link Layer Discovery Protocol (LLDP)
  - IEEE 802.3x Flow Control (Pause Frames)
  - IEEE 802.3ab 1000BASE-T
  - IEEE 802.3z 1000BASE-X

RFC Compliance

- **General Protocols**
  - RFC 768 User Datagram Protocol (UDP)
  - RFC 783 TFTP Protocol (revision 2)
  - RFC 791 Internet Protocol (IP)
  - RFC 792 Internet Control Message Protocol (ICMP)
  - RFC 793 Transmission Control Protocol (TCP)
  - RFC 826 ARP
  - RFC 854 Telnet Protocol Specification
  - RFC 894 A Standard for the Transmission of IP Datagram over Ethernet Networks
  - RFC 959 FTP
  - RFC 1027 Using ARP to Implement Transparent Subnet Gateways (Proxy ARP)
  - RFC 1112 IGMP v1
  - RFC 1157 Simple Network Management Protocol (SNMP) SNMP SNMP v1 and v2c
  - RFC 1305 Network Time Protocol (NTP) Version 3
  - RFC 1492 TACACS+
  - RFC 1519 Classless Inter-Domain Routing (CIDR)
  - RFC 1584 Multicast Extensions to OSPF
  - RFC 1765 OSPF Database Overflow
  - RFC 1812 Requirements for IP Version 4 Routers
  - RFC 1908 Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework
  - RFC 1908 Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework
  - RFC 1997 BGP Communities Attribute
  - RFC 2068 HTTP Server
  - RFC 2131 Dynamic Host Configuration Protocol (DHCP)
  - RFC 2154 OSPF with Digital Signatures (Password, MD-5)
  - RFC 2236 IGMP v2
  - RFC 2267 Network Ingress Filtering Option—Partial Support
  - RFC 2328 OSPF v2 RFC 2385 Protection of BGP Sessions with the TCP MD5 Signature Option
  - RFC 2370 OSPF Opaque Link-State Advertisement (LSA)
  - RFC 2375 IPv6 Multicast Address Assignments
  - RFC 2439 BGP Route Flap Damping
  - RFC 2460 Internet Protocol, Version 6 (v6) Specification (on management interface)
  - RFC 2462 IPv6 Stateless Address Auto-Configuration
  - RFC 2464 Transmission of IPv6 Packets over Ethernet Networks (on management interface)
  - RFC 2545 Use of BGP-MP Extensions for IPv6
  - RFC 2474 Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers
  - RFC 2571 An Architecture for Describing SNMP Management Frameworks
  - RFC 2578 Structure of Management Information Version 2
  - RFC 2579 Textual Conventions for SMIv2
  - RFC 2580 Conformance Statements for SMIv2
  - RFC 2710 Multicast Listener Discovery (MLD) for IPv6
  - RFC 2711 IPv6 Router Alert Option
  - RFC 2740 OSPFv3 for IPv6
- **General Protocols (cont.)**
  - RFC 2865 Remote Authentication Dial-In User Service (RADIUS)
  - RFC 3101 The OSPF Not-So-Stubby Area (NSSA) Option
  - RFC 3137 OSPF Stub Router Advertisement
  - RFC 3176 sFlow
  - RFC 3392 Capabilities Advertisement with BGPv4
  - RFC 3410 Introduction and Applicability Statements for Internet Standard Management Framework
  - RFC 3411 An Architecture for Describing SNMP Frameworks
  - RFC 3412 Message Processing and Dispatching for the SNMP
  - RFC 3413 Simple Network Management Protocol (SNMP) Applications
  - RFC 3414 User-based Security Model
  - RFC 3415 View-based Access Control Model
  - RFC 3416 Version 2 of SNMP Protocol Operations
  - RFC 3417 Transport Mappings
  - RFC 3418 Management Information Base (MIB) for the SNMP
  - RFC 3584 Coexistence between Version 1, Version 2, and Version 3 of the Internet-standard Network
  - RFC 3587 IPv6 Global Unicast Address Format RFC 4291 IPv6 Addressing Architecture
  - RFC 3623 Graceful OSPF Restart—IETF Tools
  - RFC 3768 VRRP
  - RFC 3826 The Advanced Encryption Standard (AES) Cipher Algorithm in the SNMP User-based Security Model
  - RFC 4271 BGPv4
  - RFC 4443 ICMPv6 (replaces 2463)
  - RFC 4456 BGP Route Reflection
  - RFC 4510 Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map
  - RFC 4724 Graceful Restart Mechanism for BGP
  - RFC4750 OSPFv2.MIB
  - RFC 4861 IPv6 Neighbor Discovery
  - RFC 4893 BGP Support for Four-Octet AS Number Space
  - RFC 5082 Generalized TTL Security Mechanism (GTSM)
  - RFC 5880 Bidirectional Forwarding Detection (BFD)
  - RFC 5881 Bidirectional Forwarding Detection (BFD) for IPv4 and IPv6 (Single Hop)RFC 5882 Generic Application of Bidirectional Forwarding Detection (BFD)
  - RFC 5883 Bidirectional Forwarding Detection (BFD) for Multihop Paths
  - RFC 5942 IPv6 Neighbor Discovery
  - RFC 7348 Virtual eXtensible Local Area Network (VxLAN)
  - RFC 7432 BGP-EVPN—Network Virtualization Using VXLAN Data Plane
- **SSH/SCP/SFTP**
  - RFC 4250 Secure Shell (SSH) Protocol Assigned Numbers
  - RFC 4251 Secure Shell (SSH) Protocol Architecture
  - RFC 4252 Secure Shell (SSH) Authentication Protocol
  - RFC 4253 Secure Shell (SSH) Transport Layer Protocol
  - RFC 4254 Secure Shell (SSH) Connection Protocol
  - RFC 4344 SSH Transport Layer Encryption Modes
  - RFC 4419 Diffie-Hellman Group Exchange for the Secure Shell (SSH) Transport Layer Protocol

- **MIBS**
  - RFC 2674 Bridge MIB
  - RFC 2819 RMON Groups 1, 2, 3, 9
  - RFC 2863 The Interfaces Group MIB
  - RFC 3826 SNMP-USM-AES-MIB
  - RFC 4022 TCP MIB
  - RFC 4113 UDP.MIB
  - RFC 4133 Entity MIB (Version 3); rmon.mib, rmon2.mib, sflow_v5.mib, bridge.mib, pbridge.mib, qbridge.mib, rstp. mib, lag.mib, lldp.mib, lldp_ext_dot1.mib, lldp_ext_dot3.mib
  - RFC 4273 BGP-4 MIB
  - RFC 4292 IP Forwarding MIB
  - RFC 4293 Management Information Base for the Internet Protocol (IP)
  - RFC 4750 OSPFv2.MIB
  - RFC 7331 BFD MIB

Layer 2 Switching

- Conversational MAC Learning
- Virtual Link Aggregation Group (vLAG) spanning
- Layer 2 Access Control Lists (ACLs)
- Address Resolution Protocol (ARP) RFC 826
- Layer 2 Loop prevention in an overlay environment
- MLD Snooping
- IGMP v1/v2 Snooping
- MAC Learning and Aging
- Link Aggregation Control Protocol (LACP) IEEE 802.3ad/802.1AX
- Virtual Local Area Networks (VLANs)
- VLAN Encapsulation 802.1Q
- Per-VLAN Spanning Tree (PVST+/PVRST+)
- Rapid Spanning Tree Protocol (RSTP) 802.1w
- Multiple Spanning Tree Protocol (MSTP) 802.1s
- STP PortFast, BPDU Guard, BPDU Filter
- STP Root Guard
- Pause Frames 802.3x
- Static MAC Configuration
- Multi-Chassis Trunking (MCT)

Layer 3 Routing

- Border Gateway Protocol (BGP4+)
- DHCP Helper
- Layer 3 ACLs
- IGMPv2
- OSPF v2/v3
- Static routes
- IPv4/v6 ACL
- Bidirectional Forwarding Detection (BFD)
- 64-Way ECMP
- VRF Lite
- VRF-aware OSPF, BGP, VRRP, static routes
- VRRP v2 and v3
- IPv4/IPv6 dual stack
- ICMPv6 Route-Advertisement Guard
- Route Policies
- IPv6 ACL packet filtering
- BGP Additional-Path
- BGP-Allow AS
- BGP Generalized TTL Security Mechanism (GTSM)
- BGP Peer Auto Shutdown
- IPv6 routing
- OSPF Type-3 LSA Filter
- Wire-speed routing for IPv4 and IPv6 using any routing protocol
- BGP-EVPN Control Plane Signaling RFC 7432
- BGP-EVPN VXLAN Standard-based Overlay
- Multi-VRF
- IP Unnumbered Interface
- VRRP-E

Automation and Programmability

- gRPC Streaming protocol and API
- REST API with YANG data model
- Python
- PyNOS libraries
- DHCP automatic provisioning
- NETCONF API

High Availability

- BFD

Quality of Service

- 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
- Random Early Discard
- Per-port QoS configuration
- ACL-based Rate Limit
- Dual-rate, three-color token bucket

Quality of Service (cont.)

- ACL-based remarking of CoS/DSCP/Precedence
- ACL-based sFlow
- Scheduling: Strict Priority (SP), Deficit Weighted Round-Robin (DWRR)

Management and Monitoring

- 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
- 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 Option 82 Insertion
- DHCP Relay
- Timestamping

Security

- Port-based Network Access Control 802.1X
- RADIUS
- AAA
- TACACS+
- Secure Shell (SSHv2)
- TLS 1.1, 1.2
- HTTP/HTTPS
- BPDU Drop
- Lightweight Directory Access Protocol (LDAP)
- Secure Copy Protocol
- Control Plane Policing (CPP)
- LDAP/AD
- SFTP
- Port Security

Ordering Information

Part Number	Description
BR-SLX-9240-32C-AC-F	SLX 9240-32C Switch AC with front-to-back airflow. 32×100 GbE/40 GbE
BR-SLX-9240-32C-DC-F	SLX 9240-32C Switch DC with front-to-back airflow. 32×100 GbE/40 GbE
BR-SLX-9240-32C-AC-R	SLX 9240-32C Switch AC with back-to-front airflow. 32×100 GbE/40 GbE
BR-SLX-9240-32C-DC-R	SLX 9240-32C Switch DC with back-to-front airflow. 32×100 GbE/40 GbE
Upgrade Licenses	
BR-SLX-9240-ADV-LIC	**Advanced License for BR-SLX-9240** License includes OVSDB integration, BGP EVPN, Guest VM, gRPC, 1588 BC, Timestamping, TPVM and NPB feature. The NPB feature set includes the following features: Traffic aggregation, Traffic replication (Transparent VLAN Flooding), L2 and L3 ACL, Route-map, Hash based load-balancing, and Timestamping.