Highlights

- Delivers agility at all layers of the data center stack
- Provides native 25 GbE server connectivity with flexible 25/40/100 GbE uplink options in a fixed 1U form factor
- Includes a programmable ASIC to accelerate adoption of new protocols and technologies
- Utilizes the SLX Insight Architecture and 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 Extreme Workflow Composer Automation Suites

ExtremeSwitching™ SLX 9140

Next-Generation Programmable Leaf Switch

As data centers and cloud service providers embrace new high-performance servers that support higher workload densities, they increasingly need dense 25/100 GbE switches in leaf-and-spine topologies. The SLX® 9140 switch is designed to help organizations stay ahead of this application- and data-driven networking challenge with a broad choice of interface speeds to seamlessly evolve from existing 1/10/40 GbE to 25/100 GbE capabilities.

The SLX 9140 enables organizations to design networks that accommodate a variety of applications and east-west traffic patterns. And with its high-density, scale-out architecture and low-power design, the SLX 9140 delivers a cost-effective solution that optimizes power, cooling, and data center space. With a rich set of Layer 2 and Layer 3 features and advanced visibility and automation capabilities, the SLX 9140 is built to address dynamic growth in Virtual Machines (VMs), distributed applications, and digital transformation.

SLX 9140 Overview

The SLX 9140 is a fixed 25/100 GbE top-of-rack leaf switch with 24 MB of packet buffer and an overall throughput of 1.8 Tbps in and out/1.2 Bpps non-blocking switching capacity. It offers forty-eight 25 GbE SFP-28 ports and six 100 GbE QSFP-28 ports.

SFP and QSFP ports offer a choice of speeds — including 100, 40, 25, 10, or 1 GbE — along with a wide choice of transceivers and cables. Ports can be mixed, offering flexible design options to cost-effectively support demanding data center and cloud service provider environments. 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.

The SLX 9140 switch offers:

- SFP-28 ports that support 25, 10, and 1 GbE modes
- QSFP-28 ports that support 100/40 GbE or 4x10/25 GbE modes

Together with Extreme Networks IP fabrics, the SLX 9140 and 9240 help transform data center networks by enabling cloud-based architectures that deliver new levels of scale, agility, and operational efficiency. These highly automated, software-driven, and programmable data center fabric design solutions support a wide range of network virtualization options and scale, supporting data center environments ranging from tens to thousands of servers. Moreover, they make it easy for organizations to architect, automate, and integrate current and future data center technologies while they transition to a cloud model on their own timeframe and terms.

The SLX 9140 helps address the increasing agility and analytics needs of digital businesses with innovative network automation and visibility capabilities provided by Extreme Workflow Composer™ and the SLX Insight Architecture.

High Availability and Reliability

The SLX 9140 delivers the high performance and reliability required by modern data centers. It is designed for high availability from both a software and hardware perspective. Key features include:

- A high-availability architecture with a clear separation between the control plane and data plane
- Deep packet buffers and advanced QoS capabilities to streamline execution at high data rates, even for bursty or long-lived traffic flows
- Redundant power supplies and fan modules that minimize single points of failure
- Active/Active Layer 2 multipathing
- 64-way ECMP routing for load balancing and redundancy
- BFD, OSPF3-NSR, and BGP4-GR

Modular, Virtualized Operating System

The SLX 9140 runs Extreme SLX-OS, a fully virtualized Linux-based operating system that delivers process-level resiliency and fault isolation. 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 Extreme 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 9140 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.

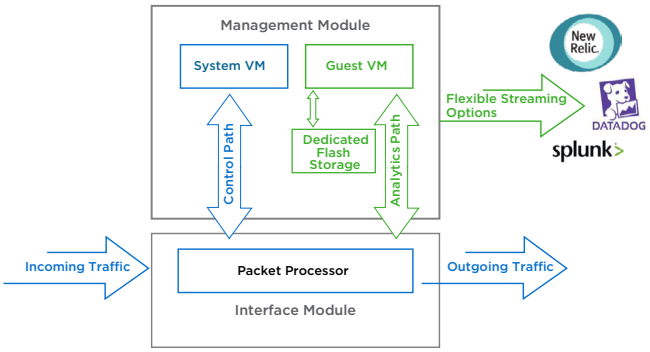

Figure 1: The Extreme SLX Insight Architecture..

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 action on the network. These actions can address end-user 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 Networks 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 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 an open KVM environment that runs third party applications and customized monitoring, troubleshooting, and analytics tools. Enabled by Extreme SLX-OS, this preconfigured guest VM is on the local CPU of each SLX 9140 Switch. It hosts third-party network operations and analytics applications on every device, extending visibility to the entire network.
- **Dedicated Analytics Path:** The SLX Insight Architecture provides an innovative internal analytics path between the packet processor on the SLX 9140 and the SLX Insight Architecture open KVM environment running on the local CPU. This enables applications running in the open KVM environment to extract data without disrupting the forwarding or control plane traffic on the switch.
- **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 9140 provides 128 GB of on-device storage dedicated to visibility applications running in the guest VM, providing real-time data capture for fast and easy access.

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 website 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, Extreme 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).

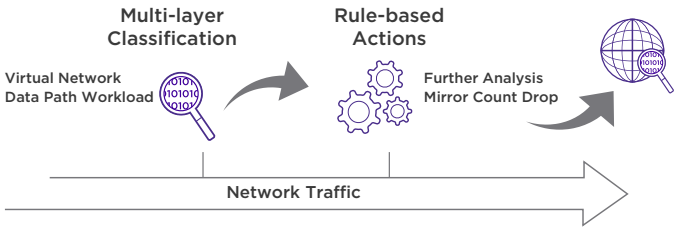

Figure 2: SLX Visibility Services.

VXLAN RIOT-Ready Hardware

The VXLAN Routing Into and Out of Tunnels (RIOT) capability enables intercommunication between data center workloads located across broadcast domains in different subnets. Many switching platforms require two or even three passes through the ASIC for RIOT functionality — either for route and encapsulation; route and decapsulation; or decapsulation, route, and encapsulation. They also tend to unnecessarily waste Ethernet ports for loopback. Ethernet LoopBack LAG (ELBL) is required for RIOT functionality, which reduces the number of available front panel ports on the switch, and each extra pass creates added latency to the RIOT function.

The SLX 9140 hardware supports RIOT, providing a flexible application deployment architecture for new and legacy multitier application workloads. With the SLX 9140, all RIOT functions — including decapsulation, route, and encapsulation — require only one pass through the ASIC. This maintains efficiency of front panel port availability and reduces latency for RIOT.

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.

SLX 9140 leverages 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.

DevOps-Inspired Automation

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

SLX 9140 and Extreme Workflow Composer

The SLX 9140, combined with Extreme Workflow Composer and the Extreme 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 infrastructure; 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 Networks infrastructure

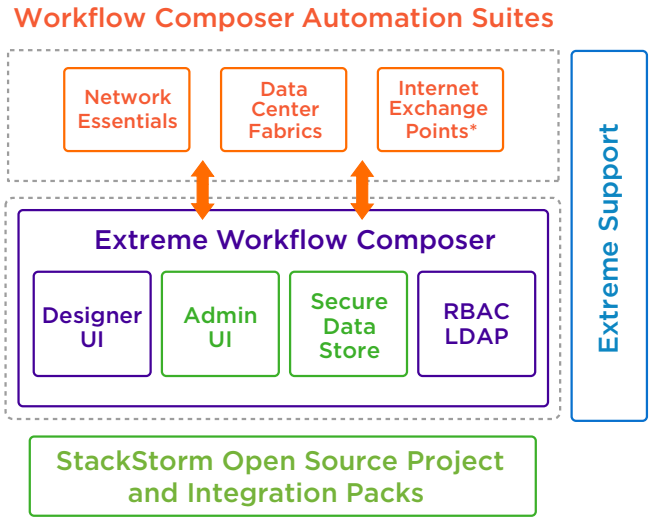

* Availability TBD

Figure 3: The Extreme Workflow Composer Automation Suite Architecture.

Extreme Management Center for Insights, Visibility and Control

The SLX family of switches and routers, including SLX 9140 can be managed by Extreme Management Center (XMC). XMC includes a suite of applications, empowering administrators to deliver a superior quality experience to end users through a single pane of glass 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.

XMC provides a consolidated view of users, devices and applications for wired and wireless networks – from data center 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.

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*.

SLX 9140 Switch Specifications

Switch Specifications	
Form factor	1U
Switching bandwidth (data rate, full duplex)	1.8 Tbps in and 1.8 Tbps out for a sum total of 3.6 Tbps
Forwarding capacity (data rate, full duplex)	(L2) 1.2 Bpps, (L3) 600 Mpps line-rate performance
Dimensions and weight	44.0 cm; 17.3 in. (Width), 44.5 cm; 17.5 in. (Depth), 4.37cm; 1.72 in. (Height) 9.00 kg; 19.8 lb
Port-to-port latency	2.5 usec
Architecture Store and Forward	Supported
25/10/1 GbE ports	48
100/40 GbE ports	6
Power supplies	Two internal, redundant, field-replaceable, load-sharing AC or DC power supplies
Cooling fans	Five field-replaceable fans
Airflow	Rear-to-front or front-to-rear airflow
Dynamically shared packet 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	182 W Two AC PSU, five fan trays, 10% traffic, low fan speed
Maximum power consumption	489 W Two AC PSU, six fan trays, 100% traffic, high fan speed
Power supply rated maximum (AC)	650 W
Switch power consumption	DC PSU 475 W; AC PSU 489 W
Environment	
Temperature	Operating: −5°C to 50°C (front-to-back air flow) −5°C to 45°C (back-to-front air flow) Temporarily up to 55°C (six fan trays) Non-operating and storage: −40°C to 70°C
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, 5-500 Hz
Airflow	134 CFM (estimated with two PSU, six fan trays)
Switch heat dissipation (25°C)	AC PSU 400 W
Acoustics (25°C)	52 dBA
MTBF (25°C)	324,414 hours

SLX 9140 Software Specifications

Software Specifications		
Connector options	• 10/1 GbE SFP+ • 40 GbE QSFP+ • 25 GbE SFP-28 • 100 GbE QSFP-28 • 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 96,000	
Maximum VLANs	4,096	
Maximum routes (in hardware)	Up to 80,000	
Maximum ACLs	2,000	
Maximum members in a standard LAG	36	
Maximum per-port priority pause level	8	
Maximum switches an mLAG can span	2	
Maximum IPv4 unicast routes	48,000	
Maximum IPv6 unicast routes	16,000	
DCB priority flow control classes	8	
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 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 1997 BGP Communities Attribute
- RFC 1908 Coexistence between Version 1 and Version 2 of the Internet-standard Network Management Framework
- 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 2370 OSPF Opaque Link-State Advertisement (LSA)
- RFC 2375 IPv6 Multicast Address Assignments

- RFC 2385 Protection of BGP Sessions with the TCP MD5 Signature Option
- 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 2474 Definition of the Differentiated Services Field in the IPv4 and IPv6 Headers
- RFC 2571 An Architecture for Describing SNMP Management Frameworks
- RFC 2545 Use of BGP-MP Extensions for IPv6
- 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
- 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

Virtualization Support

- VXLAN Routing
- VXLAN Bridging
- VXLAN Tunnel End Point
- VXLAN Multi-VNI

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

Management and Monitoring

- 1588v2 PTP
- 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

SLX 9140 Ordering Information

Part Number	Description
BR-SLX-9140-48V-AC-F	SLX 9140-48V switch AC with front-to-back airflow. 48×25 GbE/10 GbE/1 GbE + 6×100 GbE/40 GbE
BR-SLX-9140-48V-DC-F	SLX 9140-48V switch DC with front-to-back airflow 48×25 GbE/10 GbE/1 GbE + 6×100 GbE/40 GbE
BR-SLX-9140-48V-AC-R	SLX 9140-48V switch AC with back-to-front airflow 48×25 GbE/ 10GbE/ 1GbE + 6×100 GbE/40 GbE
BR-SLX-9140-48V-DC-R	SLX 9140-48V switch DC with back-to-front airflow 48×25 GbE/10 GbE/1 GbE + 6×100 GbE/40 GbE
Upgrade Licenses	
BR-SLX-9140-ADV-LIC	Advanced License for BR-SLX-9140 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.

Extreme™
Customer-Driven Networking

http://www.extremenetworks.com/contact