

ExtremeSwitching™ X8

Simplify, virtualize, and converge enterprise, data center, high-performance computing/big data, and service provider networks.

HIGHLIGHTS

- Enterprise Campus and Data Center aggregation/core, High Performance Computing, Big Data, Cloud Computing, Internet Exchange Point
- High density 1GbE, 10GbE, 40GbE and 100GbE switch for high consolidation
- 384 ports of 100/1000/10000MbE RJ45 or 1GbE RJ45/SFP per switch, 1,152 ports per rack
- 768 ports of 10GbE (SFP+ MSA compliant) per switch, 2,304 ports per rack
- 192 ports of 40GbE QSFP+ per switch, 576 ports per rack
- 32 ports of 100GbE CFP2 per switch, 96 ports per rack
- Low power consumption with 5.6 Watt per 10GbE
- Orthogonal architecture scalable from 1.28Tbps to 20.48Tbps switching capacity
- 2.3 μSec port-to-port latency across fabric, sub-μSec within a blade
- High-scale connectivity with up to 1 Million L2/L3 forwarding entries
- Powered by time-tested, modular ExtremeXOS® operating system with resilient and intelligent features
- Software Defined Networking (SDN) with OpenFlow and OneController™



Today's cloud computing and data center deployments demand highly virtualized, converged and scalable infrastructure. Multi-tenant, content hosting drives the need for a simple yet highly virtualized network fabric, along with the needs for expandability, performance, reliability and manageability. Extreme Networks' open, standards-based Software Defined Architecture (SDA) coupled with advanced virtualization features in its Extensible Operating System (EXOS) and its extensible hardware fabric, creates a uniquely flexible and powerful basis for networks of any size. By taking advantage of best-in-breed hardware and software, customers can build open, orchestrated networks that not only fit any desired configuration but also integrate seamlessly with other standards-based networks and equipment. This results in networks that are simple, fast and smart. Extreme Networks' X8 platform, provides seamless any-to-any connectivity at the hardware layer of SDA through a simple, scalable, and centrally orchestrated network.

Designed for the cloud-scale requirements for Enterprises, Data Centers, Cloud Computing, High Performance Computing (HPC), Big Data Analytics, and Internet Exchange Points (IXP), the ExtremeSwitching X8 provides a low-latency, high performance switch fabric with high-density connectivity solutions ranging from 100MbE to 100GbE, speeds, all in a compact footprint of one-third of a rack. Powered with advanced features the X8 eliminates the use of expensive multi-tier architectures and the associated challenges of inter-device connectivity, up/downlink bandwidth, and latency. The X8 also leverages a low-power design ideal for green operations and high degrees of energy efficiency, resulting in lower Total Cost of Ownership (TCO).

Key Features/Benefits

“CLOUD-SCALE” SWITCHING

The X8 can provide 2.56Tbps to 20.48Tbps total switching capacity (and in between) for traffic growth while supporting 384 1GbE/768 10GbE /192 40GbE, or 32 100GbE wire-speed ports in a single chassis. This translates into 1,152 1GbE/2,304 10GbE/576 40GbE, or 96 100GbE ports in a single rack or any combination of those. The X8 offers 160/320/480/640/960/1280 Gbps unidirectional bandwidth per slot options with current switch fabric modules making it highly extensible platform for all sizes of networks. The X8 design leverages an orthogonal direct mating system between the interface and switch fabric modules, eliminating the performance bottlenecks of backplane or mid-plane designs. This provides the potential to further scale the total switching capacity with next generation fabric modules within the same system, prolonging investment protection

HIGH-DENSITY, HIGH-SCALE SERVER VIRTUALIZATION

The X8, powered by the ExtremeXOS® modular operating system, provides a scalable foundation for virtualized applications. By connecting up to 2 Million Virtual Machines (VM MAC) per system, the X8 enables a highly virtualized network infrastructure.

ExtremeXOS runs on a high-performance control plane and helps simplify network deployment and operation through VM-ready capabilities such as ExtremeXOS Network Virtualization (XNV™) and Extreme Networks Direct Attach™. XNV allows auto-configuration of Virtual Port Profiles (VPP) to automatically detect and provision network policies in a virtualized data center in a hypervisor agnostic manner, providing simplicity in managing VM mobility in data-centers.

Advanced virtualization and overlay features such as Virtual Private LAN System (VPLS) and Multi-Chassis Link Aggregation Group (MLAG), and Virtual Router Redundancy Protocol (VRRP) enable an active/active data center design for the most demanding cloud businesses.

RELIABLE DELIVERY OF CONVERGED STORAGE TRAFFIC

The X8 is an optimal solution for high-throughput storage applications, synchronous replication, and disaster recovery within data center, Big Data and HPC environments. The X8 supports IEEE Data Center Bridging (DCBx) for reliable delivery of IP storage (iSCSI, NFS, GPFS, CIFS) traffic over a single converged network. With IEEE-compliant lossless Ethernet implementation of Priority Flow Control (PFC) and Enhanced Transmission Selection (ETS), storage traffic can be prioritized, queued and bandwidth guaranteed.

For backward compatibility, the X8 also supports transit Fiber Channel over Ethernet (FCoE) for carrying converged fiber channel storage traffic in converged networks, using wire-speed 10GbE or 40GbE IO, providing reliable storage services.

LOW LATENCY FOR FASTER RESPONSE TIME

For low-latency applications such as High Performance Computing (HPC) or Big Data Analytics, the X8 provides only 2.3 microsecond port-to-port latency from any (non-XL) IO module to any IO module with a single-tier, and under 3 microsecond latency with a two-tier fabric using Extreme Networks

The X770 series Top-of-Rack switch with the X8.

Latency is less than a microsecond between the ports on the same IO module.

With the growth of mobile devices, an increasing number of applications are hosted in data centers, resulting in increased east-west and north-south traffic. Though total latency is a result of the compute, storage and transmission latencies, the X8 can deliver the benefits of lower network latency to improve customer experiences and overall application response time, while maximizing transactional revenue per unit time.

ORCHESTRATION THROUGH SOFTWARE DEFINED ARCHITECTURE

The X8 powers the data plane hardware layer of Extreme Networks' Software Defined Architecture (SDA) with highly extensible, scalable and programmable hardware. Coupled with advanced virtualization features in its Extensible Operating System (EXOS), and Software Defined Networking (SDN), flexible network architectures of any size can be defined.

The X8 supports standard-based Open Flow to work with third party and open source SDN controllers and provides a complete SDN solution. The X8 enables SDN by providing a programmable control path and a scalable fabric that can be programmed in an efficient manner. Additionally, the X8 supports the Open Stack plugin for cloud-level integration.

The X8 can be leveraged in hybrid networks for simultaneous SDN and non-SDN based deployments. Off-the-shelf or home-grown SDN based applications can be used for network control, traffic engineering, quality of service, SLA management, monitoring, analytics, virtualization and other advanced implementations.

HIGH SCALE ROUTING

The X8 supports two types of interface modules for advanced L2 and L3/MPLS oriented networks: the “Non-XL” modules for high-density, low-latency moderate-scale edge/aggregation and core applications, and the “XL” modules for the high-scale aggregation/core or border applications. By supporting up to 1 Million Layer 2/ Layer 3 entries on its 40/10GbE and 100/10GbE XL-series modules, and large tables sizes for storing IPv4/IPv6 prefixes and MPLS labels, the X8 eliminates the need for costly traditional routers. Enterprise grade routing protocols including BGP and MPLS are supported for Internet connectivity. XL-series modules can also be deployed for the high-scale edge/aggregation applications where large numbers of host routes or access control lists (ACL) or multicast entries are required, such as managed hosting and cloud.

In enterprise and data centers, the XL and non-XL modules can be deployed together within the same switch without impacting one another - not only collapsing the tiers, but also reducing capex and opex.

RELIABILITY AND SERVICE ASSURANCE

Today’s mission critical data centers and service provider networks cannot afford service blackouts and upset customers. Designed for Tier-3 and 4 data centers and Tier-3 and 2 service providers, the X8 prevents any single point of failure at the hardware level through isolated control and data planes, fully redundant 1+1 management modules, N+1 redundant switching fabric, and N+1 redundant fans. For the power system, the X8 offers N+1 power supply level and N+N power grid level redundancy.

At the software level, modular ExtremeXOS software increases network availability by monitoring independent processes in real time. If any of these processes become unresponsive or stop running they can be automatically restarted without impacting other processes. The X8 supports a set of resiliency and fault-tolerant features, such as patch level In-Service Software Upgrades (ISSU), Multi-Chassis Link Aggregation (M-LAG), Virtual Router Redundancy Protocol (VRRP), and Ethernet Automatic Protection Switching (EAPS)


LOWER OPERATING COSTS, BETTER BUSINESS MARGINS

The X8 is designed for power cooling and space efficiency. Power consumption as low as 5.6 Watts per 10GbE port is enabled through an obstruction-less orthogonal design and variable speed fans that assist in rapid heat dissipation and more efficient use of data center cooling resources. This results in lower utility bills. The high density design of the X8 reduces rack space requirements, providing fast return on investment per RU and per port. Through advanced manageability and automation tools, the X8 reduces the need for staffing and reduces the training learning curve, reducing the total cost of ownership.

Applications


SCALABLE DATA CENTER CORE/SPINE

The X8 is designed for tier 1 through 4 enterprise, multi-tenant and cloud data center core and border deployments. In these applications, the X8 can be used with XL and non-XL series modules to provide a mix of high- and modest- scale connectivity in a spine-leaf architecture. The non-XL series modules provide 10GbE and 40GbE downlink connectivity for leaf switches such as the Extreme Network X670 and X770, while the XL series modules provide a high-scale 10GbE, 40GbE and 100GbE solution for inter-DC, disaster recovery, or cloud/internet connectivity. A pair of X8 switches connected through M-LAG provides a robust and scalable core solution. Virtualization features such as Virtual Routers (VR), and MPLS provide traffic isolation and VM mobility between the Active/Active data centers while the overall network can be orchestrated using Software Defined Architecture components.




EXPANDABLE DATA CENTER EDGE

For high-density, high-throughput environments such as hosted data centers and digital media, the X8 provides wire-speed 1GbE, 10Gb and 40 GbE server and storage connectivity as an end-of-row/middle-of-row aggregation solution. This enables seamless 1GbE to 10GbE server migration using X8 10GbE copper or fiber modules with regular or LAN on Motherboard (LoM) based servers. Edge oversubscription can be controlled through the number of fabric modules populated per system. The X8 is connected to the core using 40GbE or 100GbE uplinks. Features such as VEPA, Port Isolation, and MLAG play key role in building a scalable flat network with convergence features like DCBx and FCoE for storage connectivity while SDN helps to orchestrate the networks end to end.




HIGH PERFORMANCE COMPUTE CLUSTER

In High Performance Computing and Super Computing, where large numbers of compute nodes and storage need to be connected using low-latency 10 and 40GbE, the X8 provides high-density, low-latency access and back-end connectivity. On the access side, X8 or Summit X770 series low-latency compact switches can provide 10GbE or 40GbE connectivity to connect compute nodes within the cluster. On the back-end, the X8 can provide a 2-way or 4-way ECMP high-performance, low-latency HPC core for future cluster expansions, while running multi-100G bandwidth between clusters.



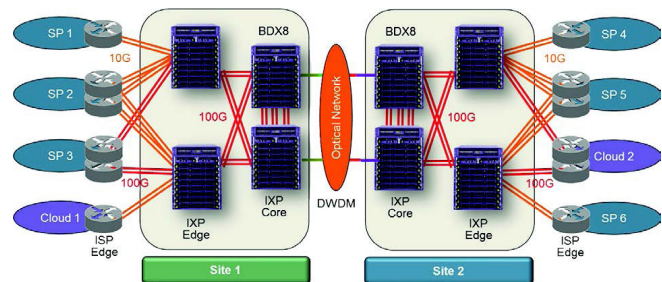
RESILIENT CAMPUS BACKBONE

As mobility, BYOD (bring your own device), video and collaboration applications demand more bandwidth on the campus edge, this puts back pressure on a more scalable and higher performance campus core. For large enterprises with more than one campus, connectivity to other regional offices as well as the enterprise data center requires a very resilient backbone. The X8 can provide a resilient backbone ring based on 10GbE, 40GbE or 100GbE and the time-tested EAPS (Extreme Automatic Protection Service) protocol. Optionally, MPLS could be used on top of the resilient backbone and for the Internet connectivity. Multicasts and QoS features enable collaboration services aggregation through the backbone with BGP/MPLS.




VIRTUALIZED INTERNET EXCHANGE POINT

The X8 offers a great solution for Internet Exchange Points (IXP) due to its density and performance. The IXP can use X8 to provision 10GbE and 100GbE ports on the customer edge and can use the same 10GbE or 100GbE in the core as well. The X8 supports long-range optics on both 10GbE and 100GbE to provide flexible customer connectivity options. An IXP can use the X8 to provision 10GbE and 100GbE services on the customer edge and can consolidate bandwidth with 100GbE links in the core. In addition, features such as MPLS VPLS, LSP load share, Selective VLAN, VLAN Bridging/Isolation, and provide capabilities to virtualize the IXP network for flexible virtual services delivery.
















SERVICE PROVIDER AGGREGATION OFFLOAD

Service Providers face a major challenge. As Internet traffic continues to grow exponentially, the majority of it remains best-effort IP traffic such as consumer video. This does not justify investments put into expensive MPLS and router-based edge and core network tiers, which results in high cost per Megabit. Moreover, conventional routers do not scale cost-effectively when it comes to port density and cost. Advanced and high-performance routing-switches like the X8 can help service providers by building a high-density aggregation layer to terminate more 10GbE ports, which allows them to be consolidated over fewer 40GbE or 100GbE ports toward the PE router. The X8 can serve as a L2, L3 or MPLS extension point. As a result, this offloads the PE routers and prolongs the investments made in the core while lowering the cost/Mb.



Technical Specifications**

COMPONENT	DESCRIPTION	DIMENSIONS			WEIGHT	
		Height	Width Length	Depth	Regular	Loaded
	Base System (48001) <ul style="list-style-type: none"> 14.5 RU Chassis with 2 Management slots, 8 IO slots, 4 Fabric slots and 8 Power Supply slots Includes 5 Fan Trays (for Front to Back airflow) Does not include Management, Fabric, IO, Blank modules or Power Supplies 	25 in 63.5 cm 14.5 RU	18 in 45.7 cm	30 in 76.2 cm	85 Kg 187.4 lb	190.7 Kg 420.6 lb
	Management Module (48021) <ul style="list-style-type: none"> 2 Management Modules for 1+1 control plane redundancy Minimum 1 required 	3 in 7.6 cm	8 in 20.3 cm	18 in 45.7 cm	2.5 Kg 5.5 lb	
	48-port 1GbE Copper Interface Module (48038) <ul style="list-style-type: none"> 8 Interfaces Modules maximum Maximum 384 1GbE wire-speed ports per system Supports 100/1000 MbE copper ports over Cat 5/6/7 cabling Suited for high-density low-latency edge/aggregation applications 	2.3 in 5.8 cm	16.5 in 41.9 cm	18.5 in 46.9 cm	5.1 kg 11.1 lb	
	48-port 1GbE Fiber Interface Module (48039) <ul style="list-style-type: none"> 8 interfaces Modules maximum Maximum 384 1GbE wire speed ports per system Supports 1Gbe SFP with SR/LR/ER/ZR optics and cables Suited for high-density low-latency edge/aggregation applications 	2.3 in 5.8 cm	16.5 in 41.9 cm	18.5 in 46.9 cm	7.3 kg 16.2 lb	
	48-Port 10GbE Copper Interface Module (48040) <ul style="list-style-type: none"> 8 Interface Modules maximum Maximum 384 10GbE wire-speed ports per system Supports 100/1000/10000 MbE (10GBaseT) RJ45 auto-negotiation over Cat6a/7 cabling Suited for high-density edge/aggregation applications 	1.9 in 4.8 cm	17.9 in 45.4 cm	16.5 in 41.9 cm	5 Kg 11.1 lb	
	48-Port 10GbE Fiber Interface Module (48041) <ul style="list-style-type: none"> 8 Interface Modules maximum Maximum 384 10GbE wire-speed ports per system Supports 1/10GbE SFP/SFP+ with SR/LR/ER/ZR optics and cables Suited for high-density, low-latency edge/aggregation/core applications 	3 in 7.6 cm	17 in 43.1 cm	18 in 45.7 cm	5 Kg 11.1 lb	
	12-Port 40GbE Fiber Interface Module (48046) <ul style="list-style-type: none"> 8 Interface Modules maximum Maximum 96 40GbE or 384 10GbE wire-speed ports per system Supports 40GbE QSFP+ with SR4/LR4 optics and cables Suited for high-density, low-latency edge/aggregation/core applications 	3 in 7.6 cm	17 in 43.1 cm	18 in 45.7 cm	5 Kg 11.1 lb	
	24-Port 40GbE Fiber Interface Module (48051) <ul style="list-style-type: none"> 8 Interface Modules maximum Maximum 192 40GbE or 768 10GbE wire-speed ports per system Supports 40GbE QSFP+ with SR4/LR4 optics and cables Suited for high-density, low-latency edge/aggregation/core applications 	3 in 7.6 cm	17 in 43.1 cm	18 in 45.7 cm	5 Kg 11.1 lb	

COMPONENT	DESCRIPTION	DIMENSIONS			WEIGHT	
	4-Port 100GbE Fiber Interface Module (48061) <ul style="list-style-type: none"> 8 Interface Modules maximum Maximum 32 100GbE or 320 10GbE wire-speed ports per system Supports 100GbE CFP2 with SR10/LR4 optics and cables Suited for high-bandwidth aggregation/core applications 	2.3 in 5.8 cm	16.5 in 41.9 cm	18.5 in 46.9 cm	6.6 Kg 14.6 lb	
	12-Port 40GbE "XL" Fiber Interface Module <ul style="list-style-type: none"> 8 Interface Modules maximum Maximum 96 40GbE or 384 10GbE wire-speed ports per system Supports 40GbE QSFP+ with SR4/LR4 optics and cables Large TCAM for up to 1 Million L2/L3/MPLS forwarding entries Suited for high-scale aggregation/core/border applications 	3 in 7.6 cm	17 in 43.1 cm	18.5 in 45.7 cm	7.6 Kg 16.9 lb	
	IO-Blank Panel (48018) <ul style="list-style-type: none"> 8 Blank Panels maximum Required for the IO and Switch Fabric Modules to maintain proper airflow for cooling. If not present, the modules may overheat and shut-down. 	2.3 in 5.8 cm	16.5 in 41.9 cm	17.7 in 44.9 cm	2.1 Kg 4.8 lb	
	2.56 Tbps (10T) Switch Fabric Module (48032) <ul style="list-style-type: none"> Provides 160Gbps per slot (unidirectional) bandwidth per SFM 4 SFM total for 10.24Tbps capacity (must be of same type) Minimum 3 required for wire-speed across all ports 4th SFM required for N+1 redundancy (except for GbE IO) Optionally, 3:1 oversubscription across all IO with 1 SFM and 3:2 oversubscription with 2 SFM Supports all IO modules except 24x40GbE module 	20 in 50.8 cm	3 in 7.6 cm	10 in 25.4 cm	4.1 Kg 9.2 lb	
	5.12 Tbps (20T) Switch Fabric Module (48031) <ul style="list-style-type: none"> Provides 320Gbps per slot (unidirectional) bandwidth per SFM 4 SFM total for 20.48Tbps capacity (must be of same type) Minimum 3 required for wire-speed cross all ports 4th SFM required for N+1 redundancy (except for GbE IO) Optionally, 3:1 oversubscription across all IO with 1 SFM and 3:2 oversubscription with 2 SFM Supports all IO modules 	20 in 50.8 cm	3 in 7.6 cm	10 in 25.4 cm	4.1 Kg 9.2 lb	
	Fan Tray-Spare (48015) <ul style="list-style-type: none"> Front to Back airflow 5 Fan Trays total (with 6 (N+1) Fans per tray) All 5 trays required (and included with the base system) 	24 in 60.9 cm	4 in 10.1 cm	3 in 7.6 cm	2.8 Kg 6.2 lb	
	2500 Watt AC Power Supply (48011) <ul style="list-style-type: none"> 8 Power Supplies total split across 2 bays (for circuit level redundancy) 2500 Watt with 220V, 1250 Watt with 110V N+1 and N+N level power redundancy (depending on the configuration) - per Power Calculator 	1.5 in 3.8 cm	4.2 in 10.6 cm	14.5 in 36.8 cm	2.4 Kg 5.3 lb	

X8 Series

GENERAL SPECIFICATIONS

PERFORMANCE

- 2.56 Tbps to 20.48 Tbps data switching capacity options
- Up to 11.4 billion packets per second (BPPS) forwarding throughput using the I/O modules available today, and up to 15.2 BPPS of fabric throughput
- Cut-Through and Store-and-Forward switching support
- 2.3 micro-second port-to-port latency (64-byte packet) across fabric, sub 1.5 micro-second within (non-XL) IO module

DENSITY

- 100 MbE Port Density:
 - 384 wire-speed 100BASE-T/X RJ45/SFP ports using 48-port 1 or 10 GbE I/O modules
- 1 GbE Port Density:
 - 384 wire-speed 1GBASE-T/X RJ45/SFP ports using 48-port 1 or 10GbE I/O modules
- 10 GbE Port Density:
 - 768 wire-speed 10GBASE-X (SFP+ MSA compliant) ports using 24-port 40 GbE I/O module
 - 384 wire-speed 10GBASE-X SFP+ ports using 48-port 10 GbE or 12-port 40 GbE I/O modules (SFP+ MSA compliant)

- 320 wire-speed 10GBASE-X (SFP+ MSA compliant) ports using 4-port 100 GbE I/O module
- 384 wire-speed 10GBASE-T RJ45 ports using 48-port 10 GbE modules
- 40 GbE Port Density:
 - 192 wire-speed 40GBASE-X QSFP+ ports using 24-port 40 GbE I/O module
 - 96 wire-speed 40GBASE-X QSFP+ ports using 12-port 40 GbE I/O module
- 100 GbE Port Density:
 - 32 wire-speed 100GBASE-X CFP2 ports using 4-port 100 GbE I/O modules

CPU, MEMORY

- Intel i7 Dual Core 2GHz
- 2GB ECC DDR3 SDRAM
- 1GB Compact Flash

LED INDICATORS

- Per port status LED for link/packet activity
- Per port status LED for link/packet activity for breakout (40/100GbE)
- System Status LEDs: management, fan, fabric and interface modules and power supplies

SCALE

Two tiers of scale options using non-XL and XL modules:

PARAMETERS	MODULES		
	10/40 GBE NON-XL	100 GBE NON-XL	40/100 GBE XL
Largest Frame	9216 Bytes		
VLAN	4,094		
MAC Forwarding	128K	384K	1032K
IPv4 Host	16K	130K	268K
IPv4 LPM	16K	16K	1024K
IPv6 Host	6K		64K
IPv6 LPM	8K	8K	512K
IPv4 Multicast	6K	64K	64K
IPv6 Multicast	3K	8K	8K
MPLS Forwarding	8K	16K	64K
ACL	2K	8K	136K
LAG Groups	384		
LAG Members	64		
QoS Queues/Port	8		
Rate Limiting Granularity	8Kbps-1Mbps Ingress and egress bandwidth policing/rate limiting per flow/ACL Egress bandwidth rate shaping per egress queue and per port		

OPERATING SPECIFICATIONS

- Operating Temperature Range: 0° C to 40° C (32° F to 104° F)
- Operating Humidity: 10% to 95% relative humidity, non-condensing
- Operational Altitude: UL/CB Certified to 2000m (6,561ft). Tested to 4000m (13,000 ft) per NEBS GR-64, Issue 3, 04-12 [76].
- Operational Shock: 30 m/s² (3g), 11ms, 60 Shocks
- Operational Random Vibration: 3-500MHz @ 1.5g rms

PACKAGING AND STORAGE SPECIFICATIONS

- Storage/Transportation Temperature: -40° C to 70° C (-40° F to 158° F)
- Storage and Transportation Humidity: 10% to 93% RH
- Packaged Shock (Half Sine):
 - < 50 kg 180 m/s² (10 g), 6 ms, 600 shocks, modules
 - > 50 kg 100 m/s² (6 g), 11 ms, 600 shocks, chassis
- Packaged Sine Vibration: 5-62 Hz @ Velocity 5mm/s, 62-500 Hz @ 0.2G
- Packaged Random Vibration: 5-20Hz @ ASD=1.0 & 20-200Hz @ -3dB/octave
- Packaged drop height
 - @ 39.5" <22 lb (10 kg) modules
 - @ 11.8" <110 lb (50 kg) chassis

REGULATORY/SAFETY

North American Safety

- cULus 60950-1:2007 2nd Ed., Listed Device (U.S.)
- CSA 22.2#60950-1-07 2nd Ed. 2011-12(Canada)
- CSA 22.2#60950-1-03 1st Ed. 2006-07(Canada)
- Complies with FCC 21CFR Chapter1, Subchapter J (U.S. Laser Safety)
- CDRH Letter of Approval (U.S. FDA Approval)

EUROPEAN SAFETY

- CB Scheme, IEC 60950-1:2005+National Deviations
- EN 60825-1:2007 (Lasers Safety)
- 206/95/EC Low Voltage Directive

INTERNATIONAL SAFETY

- GS Mark, EN 60950-1:2006+A11:2009+A1:200+A12:2011
- Taiwan CNS 14336-1(2010) (BSMI)
- AS/NZX 60950-1 (Australia/New Zealand)

EMI/EMC STANDARDS

North American EMC Standards

- FCC CFR 47 part 15 Class A (U.S.)
- ICES-003 Class A (Canada)

European EMC Standards

- EN 55022:2010 Class A
- EN 55024:2010 Class A
- EN 55011:2009+A1:2011
- EN 61000-3-2:2006+A2:2009 (Harmonics)
- EN 61000-3-3:2008 (Flicker)
- EN 61000-6-4:2007+A1:2011 (Emissions for Industrial, Scientific & Medical)
- EN 61000-6-2:2005 (Immunity for Industrial, Scientific & Medical)
- ETSI EN 300 386 v1.6.1 (2012-09) EMC Telecommunications
- 2004/108/EC EMC Directive

International EMC Certifications

- CISPR 22:2010, Class A (International Emissions)
- CISPR 24:2010 Class A (International Immunity)
- IEC 61000-4-2:2008 Electrostatic Discharge, 8kV Contact, 15kV Air, Criteria A
- IEC 61000-4-3:2010 Radiated Immunity 20V/m, Criteria A
- IEC 61000-4-4:2012 Transient Burst, 1kV, Criteria A
- IEC 61000-4-5 2005 Surge, 2kV, 4kV, Criteria A
- IEC 61000-4-6:2008 Conducted Immunity, 0.15-80MHz, 10V/m unmod. RMS, Criteria A
- IEC 61000-4-11:2004 Power Dips & Interruptions, >30%, 25 periods, Criteria C

Country Specific

- Japan Class A (VCCI)
- Australia/New Zealand, RCM
- Taiwan EMC CNS 13438(95) Class A, Safety CNS 14336-1(2010) (BSMI)
 - Brazil Res. 442, Category 3 (ANATEL)
 - China, complex equipment exemption (CCC)
- Mexico (via NRTL Listing)
- South Korea KN22, KN24 (KCC)

TELECOMMUNICATION STANDARDS

- ETSI EN 300 386 v1.6.1 (2012-09) EMC Telecommunications)
- ETSI EN 300 019 (Environmental for Telecommunications)

ENVIRONMENTAL COMPLIANCE

- EU RoHS - 2011/65/EU
- EU WEEE - 2012/19/EU
- China RoHS - SJ/T 11363-2006

IEEE 802.3 MEDIA ACCESS STANDARDS

- IEEE 802.3z 10GBASE-X
- IEEE 802.3ae 10GBASE-X
- IEEE 802.3ba 40GBASE-X
- IEEE 802.3ac VLAN Tag
- IEEE 802.3ad Link Aggregation

OPERATIONAL/TRANSPORTATION STANDARDS

- EN/ETSI 300 019-2-1 v2.1.2 – Class 1.2 Storage
- EN/ETSI 300 019-2-2 v2.2.1 – Class 2.3 Transportation
- EN/ETSI 300 019-2-3 v2.2.2 – Class 3.1e Operational
- EN/ETSI 300 753 v1.2.1 (2009-7) – Acoustic Noise
- ASTM D3580 Random Vibration Unpackaged 1.5G

FAN AND ACOUSTIC NOISE

- Sound pressure for comparison to legacy standards per ISO 7779:2010(E)
 - Low Speed: 60.3 dB(A) (LpA)
 - Medium Speed: 66.0 dB(A) (LpA)
 - High Speed: 82.3 dB(A) (LpA)
- Sound power per ISO 7779:2010(E), ISO 3744:2010(E), ETSI/EN 300 753:2007(E)
 - Low Speed: 72.0 dB(A) (LWAm)
 - Medium Speed: 78.0 dB(A) (LWAm)
 - High Speed: 94.4 dB(A) (LWAm)
- Declared sound power per 300 753 :2011-11) via ISO 9296:2010 &
 - Low Speed: 7.5 bels (LWAd)
 - Medium Speed: 8.1 bels (LWAd)
 - High Speed: 9.7 bels (LWAd)

WARRANTY

- Ltd. 1-Year on Hardware
- 90-Days on Software
- For Warranty Details, Visit:
www.extremenetworks.com/go/warranty

Power Specifications

POWER SUPPLY

- Rated Inputs:
 - Low Range: 100-120VAC, 60/50 Hz, 13 A max each power supply
 - High Range: 200-240VAC, 60/50 Hz, 13 A max each power supply
- Input Ranges:
 - Low Range: 90 -132VAC, 47 - 63 Hz
 - High Range: 185 - 264VAC, 47 - 63 Hz
- Power supply input socket IEC 320 C20
- Power cord input plug IEC 320 C19
- Power cord-sets up to 2m (6.5ft) length require minimum 16 AWG (1.0 mm²) copper stranded wire Power cord-sets greater than 2m (6.5ft) length require minimum 14 AWG (1.25 mm²) copper stranded wire (The power supply cord-set wall plug must be appropriately rated and approved for the country of installation)
- Efficiency 90% typical at full load
- DC voltage output range: 47.5 to 48.5 Vdc
- Nominal DC output:
 - Low Range: 48 Vdc, 25 A maximum each PSU
 - High Range: 48 Vdc, 50 A maximum each PSU
- DC output power:
 - 2500W @ high range for one PSU (See manual for more than one PS)
 - 1250W @ low range for one PSU (See manual for more than one PS)

POWER CONSUMPTION

- Worst case power consumption (including worst case optics) and heat load (With 10Tbps SFM):
 - 8,220 Watt at output of power supplies
 - 9,206 Watt at input of power supplies
 - 31,421 BTU/Hour total heat load
- Worst case power consumption (including worst case optics) and thermal (With 20Tbps SFM):
 - 8,460 Watt at output of power supplies
 - 9,475 Watt at input of power supplies
 - 32,338 BTU/Hour total heat load

EXTREMEXOS SUPPORTED PROTOCOLS

The X8 switch supports ExtremeXOS version 15.1 or later. Supported protocols and features can be reviewed in the latest [ExtremeXOS Data Sheet](#).

EXTREMEXOS SUPPORTED OPTICS

[100 Base SFP Optics](#)

[1000 Base-X SFP](#)

[10/100/1000 Base-T SFP](#)

[10G-Base-X-SF](#)

[40GE QSFP+ Optics](#)

[40GE QSFP+ Passive Copper Cables](#)

[40GE QSFP+ Active Optical Cables](#)

[40GE QSFP+ Fanout Cables](#)

[100G CFP2 Optics](#)

Ordering Information

PART NUMBER	PRODUCT NAME	DESCRIPTION	MAX PER
BASE PRODUCTS			
48001	BDX8-AC	X8 Series chassis with 8 I/O slots. Chassis includes 5 Fan Trays. Power Supplies or Blank Panels are not included	1
BASE MODULE OPTIONS			
48021	BDX-MM1	Management Module 1 for X8 series chassis. 2 modules required for 1+1 redundancy	2
48032	BDXA-FM10T	2.56Tbps Fabric Module for X chassis. Minimum 3 modules required for wirespeed performance, 4 required for N+1 redundancy supporting full 10Tbps	4
48031	BDXA-FM20T	5.12Tbps Fabric Module for BlackDiamond X8 chassis. Minimum 3 modules required for wirespeed performance, 4 required for N+1 redundancy supporting full 20Tbps	4
IO MODULE OPTIONS			
48038	BDXA-G48T	48-Port 1GBASE-T RJ45 module for BlackDiamond X series chassis. Up to 8 modules in the BlackDiamond X8 chassis support up to 384 wirespeed 100/1000MbE (1GbE) copper ports and work with either 2.56 or 5.12Tbps Fabric Modules	8
48039	BDXA-G48X	48-Port 1GBASE-X SFP module for BlackDiamond X series chassis. Up to 8 modules in the BlackDiamond X8 chassis support up to 384 wirespeed 1GbE ports and work with either 2.56 or 5.12Tbps Fabric Modules. Optics and cables are not included	8
48040	BDXA-10G48T	48-Port 10GBASE-T RJ45 module for BlackDiamond X series chassis. Up to 8 modules in the BlackDiamond X8 chassis support up to 384 wirespeed 100/1000/10000MbE (10GbE) copper ports and work with either 2.56 or 5.12Tbps Fabric Modules	8
48041	BDXA-10G48X	48-Port 10GBASE-X SFP+ module for BlackDiamond X series chassis. Up to 8 modules in the BlackDiamond X8 chassis support up to 384 wirespeed 10GbE ports and work with either 2.56 or 5.12Tbps Fabric Modules. Optics and cables are not included	8
48046	BDXA-40G12X	12-Port 40GBASE-X QSFP+ module for BlackDiamond X series chassis. Up to 8 modules in the BDx8 chassis support up to 96 wirespeed 40GbE or 384 wirespeed 10GbE ports and work with either 2.56 or 5.12Tbps Fabric Modules. Optics and cables are not included	8
48047	BDXB-40G12X-XL	12-Port 40GBASE-X QSFP+ (or used as 48-Port 10GBASE-X SFP+) "XL" module for BlackDiamond X series for up to -1 Million TCAM entries scale. Up to 8 modules in the BDx8 chassis support up to 96 wirespeed 40GbE or 384 wirespeed 10GbE ports and work with either 2.56 or 5.12Tbps Fabric Modules. Optics and	8
48051	BDXA-40G24X	24-Port 40GBASE-X QSFP+ module for BlackDiamond X series chassis. Up to 8 modules in the BDx8 chassis support up to 192 wirespeed 40GbE or 768 wirespeed 10GbE ports and only work with 5.12Tbps Fabric Module. Optics and cables are not included	8
48061	BDXB-100G4X	4-Port 100GBASE-X CFP2 module for BlackDiamond X series chassis. Up to 8 modules in the BDx8 chassis support up to 32 wirespeed 100GbE or 320 wirespeed 10GbE ports and work with either 2.56 or 5.12Tbps Fabric Modules. Optics and cables are not included	8
48062	BDXB-100G4X-XL	4-Port 100GBASE-X CFP2 (or used as 40-Port 10GBASE-X SFP+) "XL" module for BlackDiamond X series for up to -1 Million TCAM entries scale. Up to 8 modules in the BDx8 chassis support up to 32 wirespeed 100GbE or 320 wirespeed 10GbE ports and work with either 2.56 or 5.12Tbps Fabric Modules. Optics and cables are not included.	8
ACCESSORIES, POWER SUPPLIES AND FAN (SPARE)			
48015	BDX8-FAN	Fan Tray for BlackDiamond X8 chassis, spare. 5 fan trays required in the system	5
48011	BDX-PSU-AC2500	2500W AC Power Supply for BlackDiamond X series chassis. Up to 8 supported in the BDx8 chassis	8
48018	BDX-IO-BLANK-E	Enhanced Blank Panel for BlackDiamond X series chassis for empty I/O module slot	8
48020	BDX8-MMK	Mid Mount Kit for BlackDiamond X8 chassis.	1
SOFTWARE LICENSES			
48094	BDX-CORE-LIC	Core license for the BlackDiamond X8 chassis for scalable Layer 3 rich applications & OpenFlow Feature Pack.	1
48093	BDX-MPLS-LIC	MPLS Feature Pack license for the BlackDiamond X8 chassis.	1
11011	BDX Direct Attach	Direct Attach Feature Pack for Summit X450a/X460/X480, X650, X670 and BlackDiamond 8800, X Series.	1
POWER CABLES			
10080	Pwr Cord, 16A,NEMA 6-20P,C19	Power Cord, 16A, NEMA 6-20P, IEC320-C19	8
10081	Pwr Cord, 16A,CEE 7/7,C19	Power Cord, 16A, CEE 7/7, IEC320-C19	8
10084	Pwr Cord, 15A,AS/NZZS3112,C19	Power Cord, 15A, AS/NZZS3112, IEC320-C19	8
10087	Pwr Cord, 13A,BS1363,C19	Power Cord, 13A, BS1363, IEC320-C19	8

In support of the Extreme Networks Green initiatives, power cords can be ordered separately but need to be specified at the time order. Please refer to <http://www.extremenetworks.com/product/powercords> for details on power cord availability for this product.

Extreme Supported Optics

The X8 switch supports the optics listed below:

[10/100/1000 BASE-T Optics*](#)

[1000 BASE-X SFP Optics](#)

[10G BASE-X SFP+ Optics](#)

[40G BASE-X QSFP+ Optics](#)

[100G BASE-X CFP2 Optics](#)

**NOTE: Only 1000MbE speed supported on a 10GbE port in a X8 chassis using 10/100/1000 BASE-T Transceiver*

**** EXCEPT FOR THE SECTION TITLED “TECHNICAL SPECIFICATIONS” EXTREME NETWORKS MAKES NO WARRANTY WHATSOEVER WITH RESPECT TO ANY OTHER DATA CONTAINED IN THIS DATA SHEET INCLUDING ANY (A) WARRANTY OF MERCHANTABILITY; OR (B) WARRANTY OF FITNESS FOR A PARTICULAR PURPOSE; WHETHER EXPRESS OR IMPLIED BY LAW, COURSE OF DEALING, COURSE OF PERFORMANCE, USAGE OF TRADE OR OTHERWISE.**