


Highlights

- Transforms networks to deliver cloud scale, agility, and operational efficiency with Extreme data center fabrics
- Supports 1, 10, and 40 GbE options for optimal flexibility and scale
- Meets today's application demands with high performance and low latency
- Delivers line-rate throughput for all ports and packet sizes
- Fits into any data center design by leveraging 10 GbE/40 GbE uplinks, Ports on Demand (PoD), and Capacity on Demand (CoD)
- Maximizes network availability with efficiency and resiliency
- Supports storage environments with advanced flexibility
- Automates infrastructure provisioning, validation, troubleshooting, and remediation workflow



ExtremeSwitching[™] VDX 6940

Advanced Features Enable Data Center Transformation

The data center is evolving, driving requirements for infrastructure that can support dynamic growth in Virtual Machines (VMs), distributed applications, and Big Data, as well as the transition to cloud-based computing—without compromising performance. The VDX[®] 6940 and the Extreme VDX family of switches deliver the performance, flexibility, scale, and efficiency essential to modern data centers, including cloud and highly virtualized environments.

Optimizes Scale-Out Architecture

VDX switches are designed to help organizations stay ahead of application-driven network change by enabling agile growth through a scale-out architecture. This approach offers three key benefits. First, it enables network expansion as a business grows over time. Data center administrators can horizontally scale spine switches as the number of leaf switches increases. Second, scale-out architecture enables the creation of resilient network fabrics, eliminating a single point of failure and potential downtime. Third, and most important, scale-out network architecture delivers a compelling economic benefit. Unlike scale-up, a scale-out model reduces upfront investment. And by using high-density fixed switches, it lowers the Total Cost of Ownership (TCO), reducing power, cooling, and data center space.

VDX 6940 Switches

The VDX 6940-36Q is a fixed 40 Gigabit Ethernet (GbE)-optimized switch in a 1U form factor. It offers 36 40 GbE QSFP+ ports and can be deployed as a spine or leaf switch (see Figure 1). Each 40 GbE port can be broken out into four independent 10 GbE SFP+ ports, providing a total of 144 10 GbE SFP+ ports. Deployed as a spine, it provides options to connect either 40 GbE or 10 GbE uplinks from leaf switches. By deploying this high-density, compact switch, data center administrators can reduce their TCO through savings on power, space, and cooling.

In a leaf deployment, 10 GbE and 40 GbE ports can be mixed, offering flexible design options to cost-effectively support demanding data center and service provider environments. As in other VDX platforms, the VDX 6940-36Q offers a Ports on Demand (PoD) licensing model. VDX 6940-36Q is available with 24 ports or 36 ports. The 24-port model offers a lower entry point for organizations that want to start small and grow their networks over time. By installing a software license, organizations can upgrade their 24-port switch to the maximum 36-port switch.



Figure 1: The VDX 6940-36Q Switch provides 36 40 GbE ports.

The VDX 6940-144S is a 10 GbE-optimized switch with 40 GbE or 100 GbE uplinks in a 2U form factor. It offers 96 native 1/10 GbE SFP/SFP+ ports and 12 40 GbE QSFP+ ports, or 4 100 GbE QSFP28 ports. In addition, the VDX 6940-144S provides flexibility for uplinks, 40 GbE or 100 GbE from leaf switches, and delivers greater cross-sectional bandwidth for east-west traffic. Purpose-built for scale-out data centers and service providers, it enables aggregation of multiple server racks with a single switch.

With its high density, the VDX 6940-144S is an ideal fit for middle-of-row or end-of-row data center deployments (see Figure 2). By connecting multiple racks to a single cut-through switch, this architecture offers greater cross-sectional bandwidth for applications. Servers from multiple racks can be just one hop away from one another. The compact 2U form factor saves space, also reducing power and cooling, resulting in lower TCO.

VDX 6940 switches provide the advanced feature set that data centers require while delivering the high performance and low latency virtualized environments demand. Together with Extreme data center fabrics, these switches 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 breadth of network virtualization options and scale for 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 that addresses their needs, on their own timetable and on their terms.



Figure 2: The VDX 6940-144S Switch provides 96 1/10 GbE ports and 12 40 GbE or 4 100 GbE ports.

Transforms Networks to Deliver New Levels of Scale, Agility, and Operational Efficiency

VDX switches enable organizations to evolve their data center networks at their own pace, with full investment protection. As the foundation for several data center architectures, VDX switches support Extreme Networks IP fabrics, Extreme VCS® fabrics, as well as network virtualization, including controller-based network virtualization architectures, such as VMware NSX, and standards-based (BGP-EVPN) controller-less architectures with Extreme BGP-EVPN Network Virtualization for architectural flexibility (see Figure 3).

For organizations seeking automated provisioning capabilities to improve IT agility, VDX switches, together with Extreme VCS Fabric technology, accelerate time to value through automated provisioning of network devices and network virtualization. Automated service and resource upgrades further reduce ongoing maintenance time and costs. High availability is achieved through non-disruptive In-Service Software Upgrade (ISSU) and self-healing fabrics.

Extreme Data Center Fabrics and Network Virtualization Options

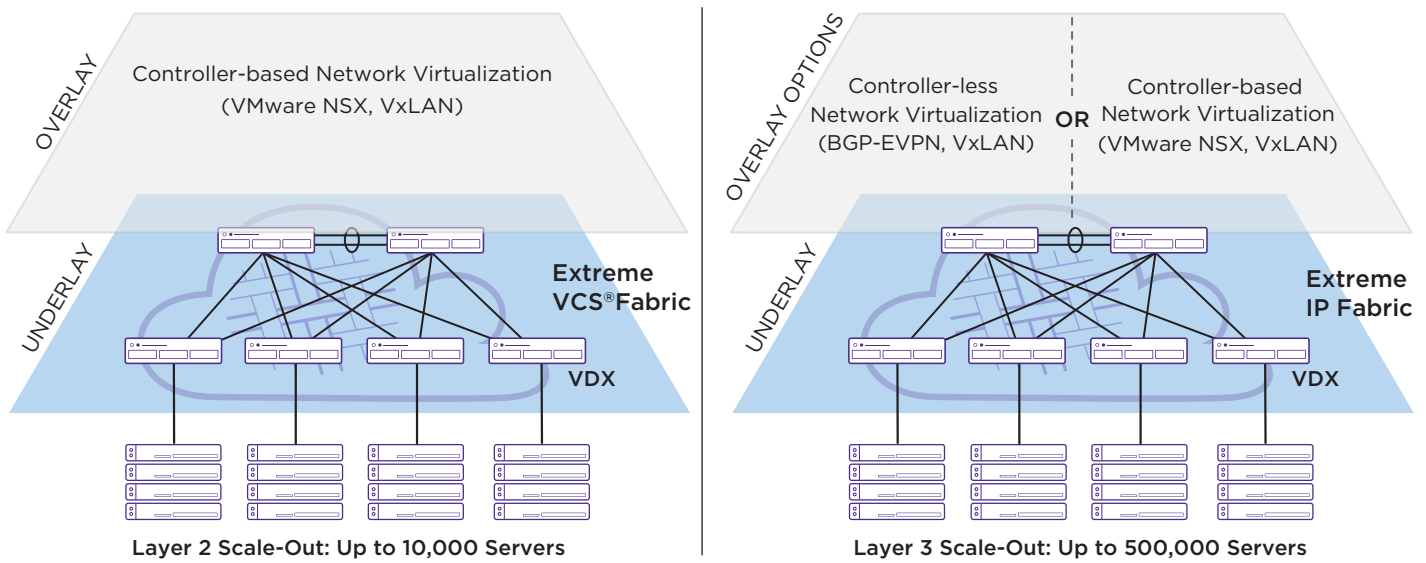


Figure 3: Multiple network architectures offer the flexibility that can help organizations rapidly adapt to changing business conditions and traffic patterns.

Optionally, for DevOps-centric organizations, Extreme VDX switches can be provisioned using Extreme Workflow Composer™ and Extreme Workflow Composer Automation Suites.

[Read more about Extreme data center fabrics.](#)

Turnkey and Customizable Lifecycle Automation

Organizations that aim to automate the entire network lifecycle but lack sufficient engineering resources can leverage Workflow Composer, a server-based, DevOps-inspired network automation platform powered by StackStorm. The Workflow Composer platform automates the entire infrastructure lifecycle — from provisioning and validation to troubleshooting and remediation. It also integrates across IT domains for end-to-end event-driven workflow automation. For more information, see the [Extreme Workflow Composer At-A-Glance](#).

Designed to run with the Workflow Composer platform, Workflow Composer Automation Suites are ideal for IT organizations that seek to embrace automation yet possess limited automation training or time. The suites provide out-of-the-box network lifecycle automation for commonly performed tasks, and are packaged to address major use cases.

The automation suites include:

- **Network Essentials:** Basic building blocks to help organizations with limited resources get up and running quickly, including workflows that automate steps common to most networks.
- **Data Center Fabrics:** A collection of workflows specific to provisioning, troubleshooting, and remediating data center fabrics, including Extreme IP fabric deployments.
- **Internet Exchange Points:** Workflows to automate steps specifically associated with Layer 2 Internet exchange connectivity, such as tenant provisioning and maintenance.

Each automation suite includes documentation and a collection of turnkey yet customizable workflows, services, sensors, actions, and rules. Organizations can use Extreme Automation Suites as-is or as starter kits for building or customizing workflows specific to their data center requirements to reduce time-to-value. For more information, see the [Extreme Workflow Composer Automation Suites At-A-Glance](#).

Additionally, VDX switches offer programmability and interoperability options through a PyNOS Library and YANG model-based REST and Netconf APIs. Cloud orchestration and control through OpenStack and OpenDaylight-based SDN controller support enable full network integration with compute and storage resource provisioning and management.

Extreme Management Center for Insights, Visibility and Control

The VDX family of switches, including VDX 6940 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.

Meets Today's Application Demands with High Performance and Low Latency

As data centers virtualize more of their servers and VM density per server increases, organizations will require higher bandwidth connectivity to support the explosion of data and application processing. With 1/10 GbE and 40 GbE options, VDX 6940 switches deliver the high-performance computing needed to keep up with the demands of a virtualized data center, allowing organizations to reduce network congestion, improve application performance, and meet the capacity required by 10 GbE servers. The 40 GbE and 100 GbE uplinks can easily aggregate high-bandwidth traffic and reduce bottlenecks that occur when aggregating multiple 10 GbE or 40 GbE connections, keeping data center networks working at peak performance.

In an Extreme VCS fabric, VDX 6940 switches also help maximize network utilization with hardware-based Extreme Inter-Switch Link (ISL) Trunking. Organizations can create a 120 GbE Extreme ISL trunk by utilizing three 40 GbE ports or 12 10 GbE ports between VDX 6940 switches. The Extreme ISL trunk is automatically formed between two Extreme VDX switches when they are linked together,

allowing traffic to be equally distributed among all ports. This increases link efficiency and limits traffic disruptions, especially during high-traffic times. Also, 40 GbE and 10 GbE trunking is supported between VDX 6940, 6740, and 8770 switches. Refer to the [Extreme Networks OS Management Configuration Guide](#) for more information.

Extreme Metro VCS technology provides an innovative solution to interconnect data centers and their traffic flows over distance, guaranteeing supported traffic characteristics. Metro VCS technology configured for regular Ethernet traffic supports 10 GbE ISLs up to 80 km, 40 GbE ISLs up to 40 km, and 100 GbE ISLs up to 40 km. To configure Metro VCS technology for lossless traffic applications (DCB/FCoE), refer to the Metro VCS Pre-deployment Guide for details.

VDX 6940 switches offer a unique balance between two conflicting attributes—buffer and latency. Niche products with very high buffer often suffer from high latency, and products with ultra-low latency are not a good fit for data centers with bursty traffic. VDX 6940 switches, with a purpose-built data center chip, excel in optimizing buffer and latency to deliver better application performance. These switches deliver 800 ns any-port-to-any-port latency. In addition, they offer an industry-leading 24 MB deep buffer. This provides the buffering capacity to handle increases in traffic, especially during peak times when ports are congested, allowing traffic to be distributed across the ports. VDX 6940 switches also feature a single ASIC design, instead of multiple ASIC designs commonly found in other switches. This improves performance since all ports communicate via the one ASIC.

Cloud and Big Data Environments

VDX 6940 switches deliver optimized buffer and latency and high performance to enable greater cross-sectional bandwidth for east-west traffic—exactly what cloud workloads demand. These switches offer the flexibility needed to scale out networks, deliver intelligence to more effectively manage VM mobility, as well as provide an SDN-enabled and programmable infrastructure. In addition, VDX 6940 switches offer an advanced feature set that non-virtual and Big Data environments require. With 10, 40, and 100 GbE options for designing oversubscribed or non-oversubscribed networks, high throughput, and optimized buffer and latency, the VDX 6940 is an ideal switch for Big Data applications. Together with Extreme data center fabrics, VDX 6940 switches can simplify network design and operations for both cloud and Big Data network fabrics.

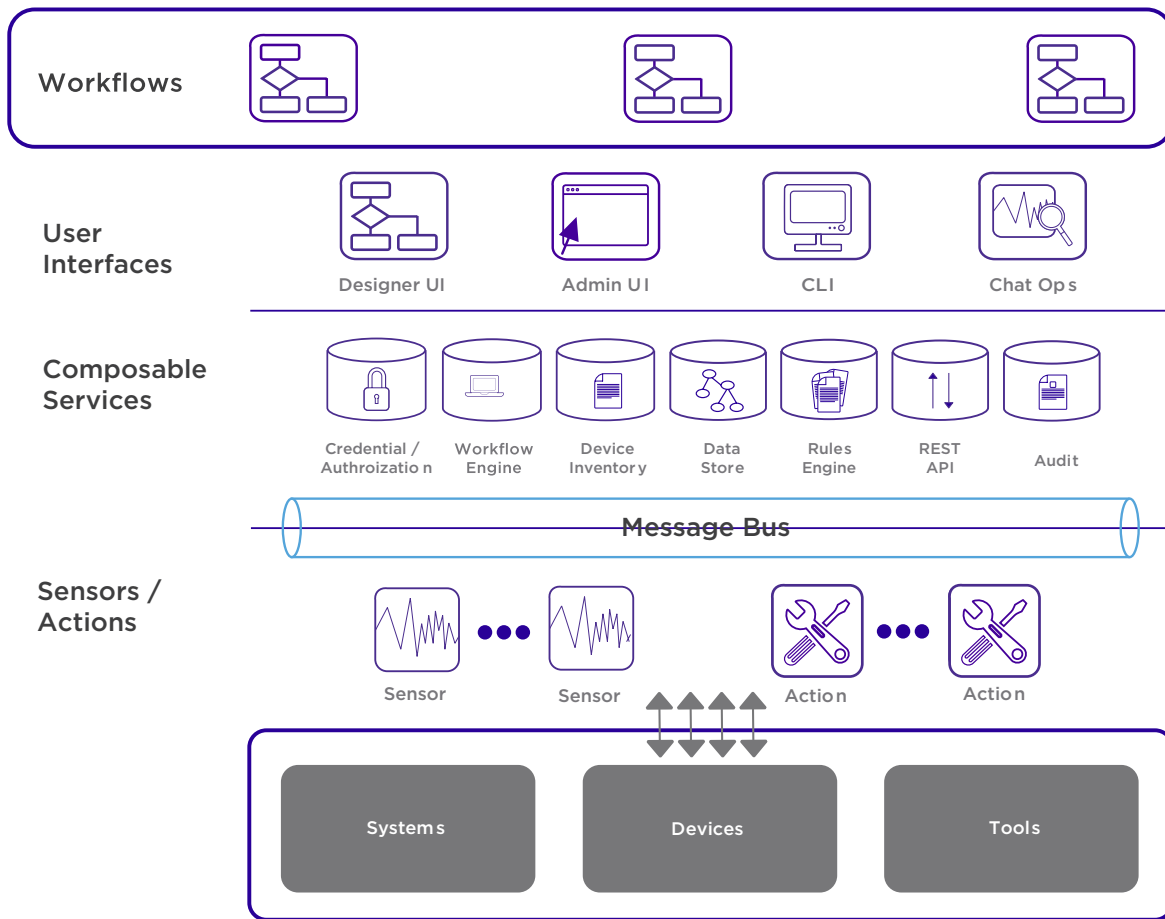


Figure 4: The Workflow Composer architecture brings workflow-centric, cross-domain network automation to IT operations.

Fits into Any Data Center Design

VDX 6940-36Q and 6940-144S front access ports are positioned to enable easy server or switch connectivity and to simplify cabling. With a choice of front-to-back or back-to-front airflow, these switches are ideal for deployments connecting servers, storage, and other switches, as well as for providing compatibility for either hot aisle or cold aisle data center designs. For the 40 GbE ports that also support 10 GbE connections with breakout cables, the switches offer the flexibility needed to support a mixed environment as data centers or service providers transition to higher bandwidth.

VDX 6940 switches are designed to connect data centers in multiple ways to meet individual design requirements. This flexible design provides investment protection, giving organizations a single switch that can support varying data center requirements. The following features help organizations meet their evolving needs:

- **10/40/100 GbE uplinks:** The 40 GbE QSFP+ ports offer the flexibility to expand and interconnect the network infrastructure intelligently and efficiently while reducing bottlenecks. VDX 6940 switches offer the option to separate each 40 GbE uplink into four 10 GbE uplinks via breakout cables. As capacity and need increase, organizations can revert to 40 GbE when ready.
- **Ports on Demand:** Ports on Demand (PoD) enables organizations to activate additional ports. They can purchase the number of ports that they currently need and seamlessly scale up later by simply applying a software license. This flexible and cost-efficient “pay as you grow” licensing model solves scalability challenges by allocating IT resources as needed.

Maximizes Network Availability with Efficiency and Resiliency

Extreme data center fabrics create a more efficient and resilient network, and deliver the high performance and high reliability required by modern data centers, including cloud and highly virtualized environments.

Optimizing East-West Traffic

Traditional data centers are architected with a rigid, three-tier tree topology optimized for the north-south traffic flow of client-server computing environments, compromising performance, increasing latency, and creating bottlenecks. With the increased prevalence of virtualization and distributed applications, data center network traffic is now predominantly east-west, or server-server. Extreme data center fabrics were designed and optimized to address these traffic patterns by moving traffic through any of the active paths and avoiding the multiple hops required in other tiered topologies.

In-Service Software Upgrade

The VDX 6940 delivers a highly efficient In-Service Software Upgrade (ISSU) by leveraging a software model that uses a dual-OS infrastructure on a multi-core CPU. This enables data center administrators to deliver enterprise-class business continuity on fixed switches during a software upgrade/downgrade process. This software change process is non-disruptive to Layer 2, Layer 3, Fibre Channel, and Fibre Channel over Ethernet (FCoE) traffic. Moreover, the ISSU implementation is hardware-optimized, thus reducing the time it takes to complete the upgrade/downgrade process.

Advanced Storage Support

VDX 6940 switches provide advanced storage support with multiple storage connectivity options, including FCoE, iSCSI, and NAS. They also feature Data Center Bridging (DCB), which enables the reliable exchange of storage traffic over the LAN, eliminating packet loss when network congestion occurs and allocating bandwidth as needed to keep the network running efficiently. Moreover, VDX 6940 switches offer NAS Auto QoS intelligence to prioritize delay-sensitive IP storage traffic within the fabric and to help ensure consistent performance while decreasing latency.

Maximizing Investments

To help optimize technology investments, Extreme Networks and its partners offer complete solutions that include professional services, technical support, and education. For more information, contact a Extreme Networks sales partner or visit www.extremenetworks.com.

VDX 6940-36Q and 6940-144S Feature Overview

Overview	VDX 6940-36Q	VDX 6940-144S
Form factor	1U	2U
Switching bandwidth (data rate, full duplex)	2.88 Tbps	2.88 Tbps
Switch performance	2.16 Bpps	2.16 Bpps
Port-to-port latency	800 ns	800 ns
Dimensions and weight	Width: 44 cm (17.32 in.) Height: 4.4 cm (1.73 in.) Depth: 43.8 cm (17.22 in.) Weight: 8.9 kg (19.6 lb) without media; 10.1 kg (22.3 lb) with media	Width: 44 cm (17.32 in.) Height: 8.7 cm (3.41 in.) Depth: 48.5 cm (19.1 in.) Weight: 15.46 kg (34.1 lb) empty; 17.32 kg (38.2 lb) with media
10 GbE SFP+ ports	144 10 GbE ports using breakout cables	96 fixed 10 GbE ports and additional 48 10 GbE ports with breakout cables
1 GbE SFP	0	96
1/10 GBASE-T	0	0
40 GbE QSFP+	36	12
100 GbE QSFP28	0	4
40 GbE Ports on Demand (PoD)	24, 36	6, 12
10 GbE Ports on Demand (PoD)	N/A	64, 80, 96
Power supplies	Two internal, redundant, field-replaceable, load-sharing AC/DC power supplies	Two internal, redundant, field-replaceable, load-sharing AC/DC power supplies
Cooling fans	5 field-replaceable fans (FRUs); each fan FRU has 2 fans	4 field-replaceable fans (FRUs); each fan FRU has 1 fan
Airflow	Rear-to-front or front-to-rear airflow	Rear-to-front or front-to-rear airflow

VDX 6940-36Q and 6940-144S Specifications

Specifications		
Scalability Information¹		
Connector options	Out-of-band Ethernet management: RJ-45 (fixed) Console management: RJ45 to RS-232 (fixed) Firmware and diagnostic: USB	
Maximum VLANs	4,096	
Maximum MAC addresses	112,000 (L2 MACs); fabric-wide 256,000 (CML)	
Maximum port profiles (AMPP)	512	
Maximum members in a standard LAG	16	
Maximum per-port priority pause level	8	
Maximum switches that a vLAG can span	8	
Maximum members in a vLAG	64	
Maximum jumbo frame size	9,216 bytes	
Queues per port	8	
DCB Priority Flow Control (PFC) classes	8	
Maximum ACLs	10,000	
Maximum ARP entries	84,000	
Maximum IPv4 unicast routes	12,000	
Maximum IPv6 unicast routes	3,000	
DCB Priority Flow Control (PFC) classes	8	
HA/ISSU	ISSU fully supported	
Mechanical		
Enclosure	Front-to-rear, rear-to-front airflow; 1U (6940-36Q) and 2U (6940-144S), 19-inch EIA-compliant; power from non-port side	
Environmental		
	VDX 6940-36Q	VDX 6940-144S
Temperature	Operating: 0°C to 40°C (32°F to 104°F) Non-operating and storage: -25°C to 70°C (-13°F to 158°F)	Operating: 0°C to 40°C (32°F to 104°F) Non-operating and storage: -25°C to 70°C (-13°F to 158°F)
Humidity	Operating: 5% to 93% at 50°C (122°F) Non-operating and storage: 10% to 95% at 60°C (140°F)	Operating: 5% to 93% at 50°C (122°F) Non-operating and storage: 10% to 95% at 60°C (140°F)
Altitude	Operating: 0 to 3,000 m (9,842 ft.) Non-operating and storage: 0 to 12,000 m (39,370 ft.)	Operating: 0 to 3,000 m (9,842 ft.) Non-operating and storage: 0 to 12,000 m (39,370 ft.)
Shock	Operating: 15 G, 11 ms, half-sine wave Non-operating and storage: 33 G, 11 ms, half-sine wave, 5/ea axis, and 33 G 6 ms, trapezoidal, 1/ea axis	Operating: 15 G, 11 ms, half-sine wave Non-operating and storage: 33 G, 11 ms, half-sine wave, 5/ea axis, and 33 G 6 ms, trapezoidal, 1/ea axis
Vibration	Operating: 1 G sine, 0.5 gms random, 5 to 500 Hz Non-operating and storage: 2.4 G sine, 1.1 gms random, 5 to 500 Hz	Operating: 1 G sine, 0.5 gms random, 5 to 500 Hz Non-operating and storage: 2.4 G sine, 1.1 gms random, 5 to 500 Hz
Airflow	Maximum: 159 CMH (94 CFM) Minimum: 65 CMH (38 CFM)	Maximum: 370 CMH (218 CFM) Minimum: 95 CMH (56 CFM)

Specifications		
Power	VDX 6940-36Q	VDX 6940-144S
Power supplies	Two internal, redundant, field-replaceable, load-sharing power supplies	Two internal, redundant, field-replaceable, load-sharing power supplies
Power inlet	C13	C13
Input voltage	100 to 240 V or 48 V DC	100 to 240 V or 48 V DC
Input line frequency	50 to 60 Hz	50 to 60 Hz
Inrush current	30 A peak at cold start and 50 A peak at warm start for <10 ms, 10 A peak for cycles 10 ms to 150 ms, <7 A peak for >150 ms	40 A peak at either cold or warm start <10 ms, 10 A peak for cycles 10 ms to 150 ms, <7 A peak for >150 ms
Maximum current	6 A at 100 VAC 7 A at 90 VAC	6 A at 100 VAC 7 A at 85 VAC
Typical power consumption	At 100 VAC 2.18 A, 215 W, 733.81 BTU/hr At 200 VAC 1.12 A, 212 W, 723.57 BTU/hr Input current calculated with 1 PSU; watts and BTU/hour with 2 PSUs	At 100 VAC 3.288 A, 327 W, 1,116 BTU/hr At 200 VAC 1.639 A, 321 W, 1,095 BTU/hr Input current calculated with 1 PSU; watts and BTU/hour with 2 PSUs
Maximum power consumption	At 100 VAC 2.84 A, 282 W, 962.5 BTU/hr At 200 VAC 1.44 A, 276 W, 942 BTU/hr Input current calculated with 1 PSU; watts and BTU/hour with 2 PSUs	At 100 VAC 2.706 A, 512 W, 1,747 BTU/hr At 200 VAC 1.375 A, 500 W, 1,706 BTU/hr Input current calculated with 1 PSU; watts and BTU/hour with 2 PSUs

Safety Compliance

- CAN/CSA C22.2 No. 60950-1-07 including A1 / UL
- 60950-1-07, Ed. 2 including A1
- CAN/CSA-C22.2 No. 60950-1 Second Edition
- EN 60950-1 Second Edition +A1/A12
- IEC 60950-1 Second Edition +A1
- GB 4943.1-2011 and GB9254-2008
- CNS14336-1(99)
- CSA/NRTL Certification

EMC

- FCC Class A
- ICES-003 Class A
- VCCI-Class A
- CE
- C-Tick
- BSMI
- KCC Class A

Immunity

- EN55024

Environmental Regulatory Compliance

- RoHS-6 (with lead exemption) Directive 2002/95/EC
- NEBS-Compliant

Standards Compliance

Extreme VDX 6940 products conform to the following Ethernet standards:

- 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

The following draft versions of the Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) standards are also supported on the:

- Extreme VDX 6940:
- IEEE 802.1Qbb Priority-based Flow Control
- IEEE 802.1Qaz Enhanced Transmission Selection
- IEEE 802.1 DCB Capability Exchange Protocol (Proposed under the DCB Task Group of IEEE 802.1 Working Group)
- FC-BB-5 FCoE (Rev 2.0)

RFC Support

- 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 IGMPv1
- RFC 1157 Simple Network Management Protocol (SNMP) v1 and v2
- 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 2068 HTTP Server
- RFC 2131 Dynamic Host Configuration Protocol (DHCP)
- RFC 2154 OSPF with Digital Signatures (Password, MD-5)
- RFC 2236 IGMPv2
- RFC 2267 Network Ingress Filtering
- RFC 2328 OSPF v2
- RFC 2370 OSPF Opaque Link-State Advertisement (LSA) Option—Partial Support
- RFC 2375 IPv6 Multicast Address Assignments
- RFC 2385 Protection of BGP Sessions with the TCP MD5 Signature Option
- RFC 2439 BGP Route Flap Damping
- RFC2460 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 2711 IPv6 Router Alert Option
- 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 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 3587 IPv6 Global Unicast Address Format
- RFC 3623 Graceful OSPF Restart—IETF Tools
- RFC 3768 VRRP
- RFC 4271 BGPv4
- RFC 4291 IPv6 Addressing Architecture
- RFC 4292 IP Forwarding MIB
- RFC 4293 Management Information Base for the Internet Protocol (IP)
- RFC 4443 ICMPv6 (replaces 2463)
- RFC 4456 BGP Route Reflection
- RFC 4510 Lightweight Directory Access Protocol (LDAP): Technical Specification Road Map
- RFC 4601 Protocol Independent Multicast—Sparse Mode (PIM-SM): Protocol Specification (Revised)
- RFC 4724 Graceful Restart Mechanism for BGP
- RFC 4861 IPv6 Neighbor Discovery
- RFC 4893 BGP Support for Four-Octet AS Number Space
- RFC 5082 Generalized TTL Security Mechanism (GTSM)
- RFC 5280 TLS client authenticating the server certificate
- RFC 5709 OSPFv2 HMAC-SHA Cryptographic Authentication
- 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 6187 SSH authentication using X.509v3 digital certificates and validating that against Pragma Fortress SSH client
- RFC 6960 TLS client authentication doing X.509v3 certificate revocation check dynamically using Online Certificate Status Protocol (OCSP)
- RFC 7166 Supporting Authentication Trailer for OSPFv3 instead of IPsec RFC 7432 BGP-EVPN—Network Virtualization Using VXLAN Data Plane

IPv6 Routing

- RFC 2545 Use of BGP-MP Extensions for IPv6
- RFC 2740 OSPFv3 for IPv6

IPv6 Multicast

- RFC 2710 Multicast Listener Discovery (MLD) for IPv6

Extreme Networks Network OS Software Capabilities

	VCS Fabrics	IP Fabrics
Software Scalability		
Maximum switches in a fabric	48	Unlimited
Maximum ECMP paths in a fabric	32	32
Maximum LAGs in a fabric	2,000	Unlimited
Layer 2 Switching		
Service Node Load Balancing BFD/ARP Optimizations	X	X
Conversational MAC Learning	X	X
Virtual Link Aggregation Group (vLAG) spanning	X	X
Layer 2 Access Control Lists (ACLs)	X	X
Supports 2K ingress and egress ACLs	X	X
Edge Loop Detection (ELD)	X	X
Address Resolution Protocol (ARP) RFC 826	X	X
Private VLANs	X	
Maintenance Mode/Graceful Traffic Diversion	X	
Distributed VXLAN Gateway	X	
Diagnostic Ports	X	
IP Maps Support	X	
L2 Loop prevention in an overlay environment		X
High availability/In-Service Software Upgrade—hardware-enabled	X	X
IGMP snooping support for multicast flooding	X	X
IGMPv1/v2 Snooping	X	X
IGMPv3	X	X
MAC Learning and Aging	X	X
Link Aggregation Control Protocol (LACP) IEEE 802.3ad/802.1AX	X	X
Virtual Local Area Networks (VLANs)	X	X
VLAN Encapsulation 802.1Q	X	X
Per-VLAN Spanning Tree (PVST+ /PVRST+)	X	X
Rapid Spanning Tree Protocol (RSTP) 802.1w	X	X
Multiple Spanning Tree Protocol (MSTP) 802.1s	X	X
STP PortFast, BPDU Guard, BPDU Filter	X	X
STP Root Guard	X	X
Pause Frames 802.3x	X	X
Static MAC Configuration	X	X
Uni-Directional Link Detection (UDLD)	X	X
Uplink switch for VDX switches, Extreme Networks VCS fabrics, and the Extreme Networks VCS Virtual Fabric feature	X	
Transparent LAN Services	X	
L2 Traceroute for VXLAN	X	X
BUM Storm Control	X	X
Layer 3 Switching		
Border Gateway Protocol (BGP4+)	X	X
DHCP Helper	X	X
Layer 3 ACLs	X	X
Multicast: PIM-SM, IGMPv2	X	X
OSPF v2/v3	X	X

	VCS Fabrics	IP Fabrics
Static routes	X	X
IPv4/v6 ACL	X	X
Policy-Based Routing (PBR)	X	X
Bidirectional Forwarding Detection (BFD)	X	X
32-Way ECMP	X	X
VRF Lite	X	X
VRF-aware OSPF, BGP, VRRP, static routes	X	X
VRRP v2 and v3	X	X
uRPF for IPv4 and IPv6	X	
IPv4/IPv6 dual stack	X	X
IPv6 ACL packet filtering	X	X
BGP automatic neighbor discovery for IP fabric		X
BGP Additional-Path	X	X
BGP-Allow AS	X	X
BGP Generalized TTL Security Mechanism (GTSM)	X	X
BGP graceful shutdown for maintenance mode		X
BGP Peer Auto Shutdown	X	X
Multicast Refactoring	X	X
IPv6 routing	X	X
OSPF Type-3 LSA Filter	X	X
Wire-speed routing for IPv4 and IPv6 using any routing protocol	X	X
BGP-EVPN Control Plane Signaling RFC 7432		X
BGP-EVPN VXLAN Standard-based Overlay		X
Multi-VRF	X	X
IP Unnumbered Interface		X
Intersubnet Routing (Symmetric and Asymmetric)		X
IP over Port Channel		X
VRRP-E	X	X
Fabric Virtual Gateway	X	X
Static Anycast Gateway		X
ARP Suppression		X
Automation and Programmability		
OpenFlow 1.3	X	X
REST API with YANG data model	X	X
Puppet	X	X
Python	X	X
PyNOS libraries	X	X
VMware vRealize plugins	X	X
DHCP automatic fabric provisioning	X	X
Netconf API	X	X
Multitenancy and Virtualization		
TRILL FGL-based VCS Virtual Fabric feature	X	
Virtual fabric extension	X	
VM-Aware Network Automation	X	
BFD for virtual fabric extension	X	
Automatic Migration of Port Profiles (AMPP)	X	X

	VCS Fabrics	IP Fabrics
DCB		
Priority-based Flow Control (PFC) 802.1Qbb	X	
Enhanced Transmission Selection (ETS) 802.1Qaz	X	
Manual configuration of lossless queues for protocols other than FCoE and iSCSI	X	
Data Center Bridging Exchange (DCBX)	X	
DCBX Application Type-Length-Value (TLV) for FCoE and iSCSI	X	
IP storage		
Inter-Switch Link (ISL)	X	
Deep on-chip packet buffer	X	X
Auto QoS for NAS	X	X
VCS fabric auto forming/auto healing	X	X
Fibre Channel/FCoE		
Multi-hop Fibre Channel over Ethernet (FCoE); requires Extreme Networks VCS Fabric technology	X	
FC-BB5 compliant Fibre Channel Forwarder (FCF)	X	
Native FCoE forwarding	X	
FCoE to Fibre Channel Bridging	X	
FCoE on VDX 8770	X	
FCoE on QSFP+ port	X	
Multi-hop Access Gateway Support	X	
End-to-end FCoE (initiator to target)	X	
FCoE Initialization Protocol (FIP) v1 support for FCoE device login and initialization	X	
Name Server-based zoning	X	
Supports connectivity to FIP Snooping Bridge (FSB) device	X	
FCoE traffic over standard LAG	X	
Interface Binding	X	
Dual Personality Ports	X	
Logical SANs	X	
High-Availability		
ISSU L2 and L3	X	X
BFD	X	X
OSPF3-NSR	X	X
BGP4-GR	X	X
Management Module Failover	X	X
Quality of Service (QoS)		
ACL-based QoS	X	X
Eight priority levels for QoS	X	X
Class of Service (CoS) IEEE 802.1p	X	X
DSCP Trust	X	X
DSCP to Traffic Class Mutation	X	X
DSCP to CoS Mutation	X	X
DSCP to DSCP Mutation	X	X
Random Early Discard	X	X
Per-port QoS configuration	X	X
ACL-based Rate Limit	X	X
Dual-rate, three-color token bucket	X	X
ACL-based remarking of CoS/DSCP/Precedence	X	X

	VCS Fabrics	IP Fabrics
ACL-based sFlow	X	X
Scheduling: Strict Priority (SP), Deficit Weighted Round-Robin (DWRR), Hybrid	X	X
Scheduling (Hybrid)		
Queue-based Shaping	X	X
Flow-based QoS	X	X
Management and Monitoring		
Logical chassis management	X	
IPv4/IPv6 management	X	X
Industry-standard Command Line Interface (CLI)	X	X
Netconf API	X	X
REST API with YANG data model	X	X
VDX Plugin for OpenStack	X	X
Link Layer Discovery Protocol (LLDP) IEEE 802.1AB	X	X
MIB II RFC 1213 MIB	X	X
Switch Beaconing	X	X
Management VRF	X	X
Switched Port Analyzer (SPAN)	X	X
Telnet	X	X
SNMP v1, v2C, v3	X	X
sFlow RFC 3176	X	X
Out-of-band management	X	X
Remote SPAN (RSPAN)	X	X
RMON-1, RMON-2	X	X
NTP	X	X
Management Access Control Lists (ACLs)	X	X
Role-Based Access Control (RBAC)	X	X
Range CLI support	X	X
UDLD	X	X
OpenStack Neutron ML2 plugin	X	X
Python	X	X
Puppet	X	X
Distributed Configuration Management	X	
Maps switch health monitoring	X	
Security		
Port-based Network Access Control 802.1X	X	X
RADIUS (AAA)	X	X
TACACS+	X	X
Secure Shell (SSHv2)	X	X
BPDU Drop	X	X
Lightweight Directory Access Protocol (LDAP)	X	X
Secure Copy Protocol	X	X
Port Security	X	X

VDX 6940-36Q and 6940-144S Software License Ordering Information

Software SKU	Description
BR-VDX6940-FCOE	FCoE software license for VDX 6940
BR-VDX6940-144S-16-10GPOD	16x10 GbE Ports on Demand (PoD) license for VDX 6940-144S
BR-VDX6940-144S-6X40G-POD	6x40 GbE or 2x100 GbE Ports on Demand (PoD) license for VDX 6940-144S
BR-VDX6940-36Q-12X40G-POD	12x40 GbE Ports on Demand (PoD) license for VDX 6940-36Q

VDX 6940-36Q and 6940-144S Hardware Ordering Information

See the VDX Transceiver Support Matrix for optics and cable ordering details.

Hardware SKU	Description
BR-VDX6940-36Q-AC-F	VDX 6940-36Q base system with 36 40 GbE QSFP+ ports, AC power supply, non-port-side exhaust airflow
BR-VDX6940-36Q-AC-R	VDX 6940-36Q base system with 36 40 GbE QSFP+ ports, AC power supply, port-side exhaust airflow
BR-VDX6940-24Q-DC-F	VDX 6940-36Q base system with 24 40 GbE QSFP+ ports, DC power supply, non-port-side exhaust airflow
BR-VDX6940-24Q-DC-R	VDX 6940-36Q base system with 24 40 GbE QSFP+ ports, DC power supply, port-side exhaust airflow
BR-VDX6940-24Q-AC-F	VDX 6940-36Q base system with 24 40 GbE QSFP+ ports, AC power supply, non-port-side exhaust airflow
BR-VDX6940-24Q-AC-R	VDX 6940-36Q base system with 24 40 GbE QSFP+ ports, AC power supply, port-side exhaust airflow
BR-VDX6940-64S-AC-R	VDX 6940-144S base system with 64 10 GbE SFP+ ports, AC power supply, port-side exhaust airflow
BR-VDX6940-64S-AC-F	VDX 6940-144S base system with 64 10 GbE SFP+ ports, AC power supply, non-port-side exhaust airflow
BR-VDX6940-96S-AC-R	VDX 6940-144S base system with 96 10 GbE SFP+ ports, AC power supply, port-side exhaust airflow
BR-VDX6940-96S-AC-F	VDX 6940-144S base system with 96 10 GbE SFP+ ports, AC power supply, non-port-side exhaust airflow
BR-VDX6940-144S-AC-R	VDX 6940-144S base system with 96 10 GbE SFP+ ports and up to 12 40 GbE QSFP+ ports or up to 4 100 GbE QSFP28 ports, AC power supply, port-side exhaust airflow
BR-VDX6940-144S-AC-F	VDX 6940-144S base system with 96 10 GbE SFP+ ports and up to 12 40 GbE QSFP+ ports or up to 4 100 GbE QSFP28 ports, AC power supply, non-port-side exhaust airflow
BR-VDX6940-64S-DC-F	VDX 6940-144S base system with 64 10 GbE SFP+ ports, DC power supply, non-port-side exhaust airflow
BR-VDX6940-64S-DC-R	VDX 6940-144S base system with 64 10 GbE SFP+ ports, DC power supply, port-side exhaust airflow



<http://www.extremenetworks.com/contact>

©2019 Extreme Networks, Inc. All rights reserved. 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 are the property of their respective owners. For additional information on Extreme Networks Trademarks please see <http://www.extremenetworks.com/company/legal/trademarks>. Specifications and product availability are subject to change without notice. 12150-0119-18