

RackSwitch G8316

Product Guide (withdrawn product)

The RackSwitch™ G8316 (shown in Figure 1) provides low latency, lossless performance, and a feature-rich design with key virtualization features, such as Converged Enhance Ethernet (CEE)/Data Center Bridging (DCB), high availability, and enterprise class Layer 2 and Layer 3 functions. In addition, the G8316 also delivers excellent cost savings as you consider acquisition costs, energy costs, operational expense, and ease of use and management for a 40 Gb class switch. The RackSwitch G8316 is suitable for clients using 10 Gb or 40 Gb connectivity (or both).

The G8316 Ethernet aggregation switch enables customers to own an end-to-end flat 2-tier network that is based on industry standards using RackSwitch switches and System x and ThinkServer servers. For example, the G8316 is an ideal tier two switch to use when connecting a number of RackSwitch G8264 40 Gb uplink ports at the access layer. Other clients like the G8316 as an investment protection switch, as they can use it with their 10 Gb Ethernet environments today, but can also leverage it in the future as they move to 40 Gb Ethernet.

With latency below 1 microsecond, the G8316 is an industry leader. This and the 1.28 Tbps throughput makes the G8316 an ideal offering for latency-sensitive applications, such as high-performance computing, financial applications, hosting environments, and cloud designs. In addition, VMready® with Virtual Vision enables the network to be Virtual Machine (VM) aware, and provides the capability to have a virtualization environment that is simpler, less expensive, and provides exceptional performance.



Figure 1. RackSwitch G8316

Did you know?

The RackSwitch G8316 supports several types of configurations: 10 Gb, and 40 Gb, Virtual NIC, Converged Enhanced Ethernet (CEE/DCB), and iSCSI.

The G8316 supports Data Center Bridging (DCB), which is the IEEE's group of protocols that provide lossless Ethernet and allows for clients to reduce the costs of implementing either NAS or iSCSI convergence and priority-based flow control.

Uses VMready on-switch software that helps reduce the complexity of managing virtual machines (VMs) in the network (the license for this capability comes standard with the purchase of the switch and is part of the Networking Operating System).

The RackSwitch G8316 is OpenFlow enabled. With OpenFlow, you can easily create user-controlled virtual networks, optimize performance dynamically, and minimize complexity.



Part number information

The part numbers to order the switch and additional options are shown in Table 1.

Table 1. Part numbers and feature codes for ordering

Description	Part number	Feature code for MTM 8036-HC2	Feature code for MTM 8036-HC1
Switch			
RackSwitch G8316 (Rear to Front)	8036ARX	A2MZ	None
RackSwitch G8316 (Front to Rear)	8036AFX	None	A2N0
Miscellaneous options			
Console Cable Kit Spare	90Y9462	A2MG	A2MG
Adjustable 19" 4 Post Rail Kit	00D6185	A3KP	A3KP
iDataPlex Rail Kit	90Y3535	None	A1SZ
Air Inlet Duct for 483 mm RackSwitch	00D6060	A3KQ	None
Hot-Swappable, Front-to-Rear Power Supply Spare	49Y7937	None	A2MJ
Hot-Swappable, Rear-to-Front Power Supply Spare	49Y7938	A2MH	None
Hot-Swappable, Front-to-Rear Fan Assembly Spare	49Y7939	None	A2MF
Hot-Swappable, Rear-to-Front Fan Assembly Spare	88Y6026	A2ME	None

The part numbers for the G8316 switches include the following items:

- One RackSwitch G8316 with two power supplies and four fan assemblies (rear-to-front airflow or front-to-rear airflow)
- Generic Rack Mount Kit (2-post)
- Console Cable Kit that includes:
 - RJ-45 (plug) to RJ-45 (plug) serial cable (1 m)
 - Mini-USB to RJ-45 (jack) adapter cable (0.2 m) with retention clip
 - DB-9 to RJ-45 (jack) adapter
- Warranty Flyer
- Important Notices Flyer
- Documentation CD-ROM

Note: Power cables are not included and must be ordered separately (see Table 2 for details).

The G8316 switch supports up to two redundant hot-swap 450 W AC power supplies (two power supplies come standard with the switch) and up to four redundant hot-swap fan assemblies (four fan assemblies come standard with the switch). Spare power supplies and fan assemblies can be ordered, if required. Each Power Supply Spare option contains one hot-swap power supply (rear-to-front or front-to-rear), and each Fan Assembly Spare option contains one hot-swap fan assembly (rear-to front or front-to-rear).

The G8316 switch also comes standard with the Console Cable Kit for management through a serial interface. Spare serial management cables can be ordered, if required. The Console Cable Kit Spare option contains the following items:

- RJ-45 (plug) to RJ-45 (plug) serial cable (1 m)
- Mini-USB to RJ-45 (jack) adapter cable (0.2 m) with retention clip
- DB-9 to RJ-45 (jack) adapter



The G8316 switch supports optional adjustable 19-inch, 4-post rack installation kit, part number 00D6185. Optionally, Air Inlet Duct, part number 00D6060, can be ordered with the G8316 (rear-to-front airflow) switch for 4-post rack installations with the Adjustable 4-post Rail Kit (00D6185).

The G8316 (front-to-rear airflow) switch also supports 4-post iDataPlex® rack kit (90Y3535) which is used when the switch is installed in the iDataPlex Rack.

The G8316 switch ships standard without any AC power cables. Table 2 lists the part numbers and feature codes to order the power cables (two power cables are required per switch).

Table 2. Power cables

Description	Part number	Feature code for MTM 8036-HC2 and 8036-HC1
Rack power cables		
1.5m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	39Y7937	6201
2.8m, 10A/100-250V, C13 to IEC 320-C20 Rack Power Cable	39Y7938	6204
4.3m, 10A/100-250V, C13 to IEC 320-C14 Rack Power Cable	39Y7932	6263
Line cords		
European 10A line C13 to CEE 7/7 (2.8M)	39Y7917	6212
Denmark 10A line C13 to DK2-5A (2.8M)	39Y7918	6213
Switzerland 10A line C13 to SEV 1011 (2.8M)	39Y7919	6216
Israel 10A line C13 to SI 32 (2.8M)	39Y7920	6218
South Africa 10A line C13 to SABS 164/1 (2.8M)	39Y7922	6214
United Kingdom 10A line C13 to BS 1363 (2.8M)	39Y7923	6215
Australia/NZ 10A line C13 to SAA-AS C112 (2.8M)	39Y7924	6211
Korea 7A line C13 to KETI 15A/250V (2.8M)	39Y7925	6219
India 6A line C13 to Fig 68 (2.8M)	39Y7927	6269
China 6A line C13 to GB 2099.1 (2.8M)	39Y7928	6210
Brazil 10A line C13 to NBR 6147 (2.8M)	39Y7929	6223
Argentina 10A line C13 to IRAM 2063 (2.8M)	39Y7930	6222
10A/250V C13 to NEMA 6-15P 2.8m power cord	46M2592	A1RF
Japan 10A/100V C13 to JIS C-8303 2.8m power cord	46M2593	A1RE

Supported cables and transceivers

With the flexibility of the G8316 switch, clients can take advantage of the technologies that they require for multiple environments:

- For 10 GbE links, clients can use the 10GBASE-SR transceivers for distances up to 300 meters over OM3 multimode fiber or up to 400 meters over OM4 multimode fiber with LC connectors. For longer distances, the 10GBASE-LR transceivers can support distances up to 10 kilometers on single mode fiber with LC connectors. The use of an SFP+ transceiver in a QSFP+ port requires the optional QSFP-to-SFP+ adapter (00D9676). To increase the number of available 10 GbE ports, clients can split out four 10 GbE ports for each 40 GbE port using QSFP+ DAC Breakout Cables for distances up to 5 meters. For distances up to 100 m, optical MTP-to-LC break-out cables can be used with the 40GBASE-SR4 transceiver, but Lenovo does not supply these optical breakout cables.
- For 40 GbE to 40 GbE connectivity, clients can use the affordable QSFP+ to QSFP+ DAC cables for distances up to 7 meters. For distances up to 100 m, the 40GBASE-SR4 QSFP+ transceiver can be used with OM3



multimode fiber with MTP connectors or up to 150 m when using OM4 multimode fiber with MTP connectors.

Table 3 lists the supported cables and transceivers.

Table 3. Supported transceivers and direct-attach cables

Description	Part number	Feature code (MTM 8036-HC1 / 8036-HC2)	Maximum quantity supported
QSFP+ transceiver and cables - 40 GbE			
Lenovo 40GBASE-SR4 QSFP+ Transceiver	49Y7884	A1DR	16
Optical cables for 40 GbE QSFP+ SR4 transceivers			
Lenovo 10m QSFP+ MTP-MTP OM3 MMF Cable	90Y3519	A1MM	16
Lenovo 30m QSFP+ MTP-MTP OM3 MMF Cable	90Y3521	A1MN	16
Lenovo 10m QSFP+ MTP-MTP OM3 MMF Cable (replaces 90Y3519)	00VX003	AT2U	16
Lenovo 30m QSFP+ MTP-MTP OM3 MMF Cable (replaces 90Y3521)	00VX005	AT2V	16
QSFP+ breakout cables - 40 GbE to 4x10 GbE			
Lenovo 1m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7886	A1DL	16
Lenovo 3m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7887	A1DM	16
Lenovo 5m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7888	A1DN	16
QSFP+ direct-attach cables - 40 GbE			
Lenovo 1m Passive QSFP+ DAC Cable	49Y7890	A1DP	16
Lenovo 3m Passive QSFP+ DAC Cable	49Y7891	A1DQ	16
Lenovo 5m Passive QSFP+ DAC Cable	00D5810	A2X8	16
Lenovo 7m Passive QSFP+ DAC Cable	00D5813	A2X9	16
QSFP to SFP+ adapter			
Mellanox QSFP to SFP+ Adapter	00D9676	ARZH	16
SFP+ transceivers - 10 GbE*			
SFP+ SR Transceiver	46C3447	5053	16
SFP+ LR Transceiver	90Y9412	A1PM	16
Optical cables for 10 GbE SFP+ SR transceivers			
Lenovo 0.5m Passive SFP+ DAC Cable	00D6288	A3RG	16
Lenovo 1m Passive SFP+ DAC Cable	90Y9427	A1PH	16
Lenovo 1.5m Passive SFP+ DAC Cable	00AY764	A51N	16
Lenovo 2m Passive SFP+ DAC Cable	00AY765	A51P	16
Lenovo 3m Passive SFP+ DAC Cable	90Y9430	A1PJ	16
Lenovo 5m Passive SFP+ DAC Cable	90Y9433	A1PK	16
Lenovo 7m Passive SFP+ DAC Cable	00D6151	A3RH	16

* Require QSFP to SFP+ Adapter (00D9676).

Benefits

The RackSwitch G8316 is considered particularly suited for these environments:

- Clients who are deploying 10 GbE on servers or blade chassis and require 40 GbE upstream aggregation to build a POD or cluster
 - Flex System embedded switches – SI4093, EN4093R, CN4093
 - Also ideal for 2nd tier networking when using with RackSwitch G8264/T/CS



- Client who are deploying 40 GbE server connectivity
- Clients looking to converge their SAN and LAN on to one network via NAS, iSCSI, or FCoE
- Looking for ways to reduce I/O cost (CAPEX) – adapters, cables, transceivers & upstream network
- Looking to reduce complexity (OPEX) – less to manage and lower energy cost
- Applications demanding better performance and lower latency
- Clients looking for investment protection: Cloud and SDN

The RackSwitch G8316 offers the following benefits:

- **High performance:** This 10/40Gb low latency (880 nanoseconds) switch with 1.28 Tbps throughput provides the best combination of low latency, non-blocking line-rate switching, and ease of management. The G8316 is also a single ASIC design, which promises consistent lower port-to-port latency compared with other vendors with multiple chip designs, which causes port-to-port latency to be inconsistent and unpredictable.
- **Lower power and better cooling:** The RackSwitch G8316 uses as little as 330 W of power, which is a fraction of the power consumption of many competitive offerings. The front-to-rear or rear-to-front cooling design reduces data center air conditioning costs by having airflow match the servers in the rack. In addition, variable speed fans assist in automatically reducing power consumption.
- **High Availability:** The G8316 also comes standard with hot-swap redundant power supplies and fans, making the switch highly reliable, ready to use, and easy to service in the unlikely event of a failure.
- **VM-aware networking:** VMready software on the switch simplifies configuration and improves security in virtualized environments. VMready automatically detects virtual machine movement between physical servers and instantly reconfigures each VM's network policies across VLANs to keep the network up and running without interrupting traffic or impacting performance. VMready works with all leading VM providers, such as VMware, Citrix, Xen, Microsoft Hyper-V, Red Hat KVM, and IBM PowerVM.
- **Layer 3 functionality:** The G8316 switch includes Layer 3 functionality, which provides security and performance benefits, as inter-VLAN traffic stays within the switch. This switch also provides the full range of Layer 3 protocols from static routes for technologies, such as Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) for enterprise customers.
- **Seamless interoperability:** The G8316 switch interoperate seamlessly with other vendors' upstream switches.
- **Fault tolerance:** The G8316 switch learns alternate routes automatically and performs faster convergence in the unlikely case of a link, switch, or power failure. The switch uses proven technologies like L2 trunk failover, advanced VLAN-based failover, VRRP, and Hot Links.
- **OpenFlow enabled:** The RackSwitch G8316 offers benefits of OpenFlow. OpenFlow is the new open application programming interface (API) that enables the network administrator to easily configure and manage virtual networks that control traffic on a "per-flow" basis. It creates multiple independent virtual networks and related policies without dealing with the complexities of the underlying physical network and protocols.
- **Multicast:** This supports IGMP Snooping v1, v2, and v3 with 2K IGMP groups, as well as Protocol Independent Multicast, such as PIM Sparse Mode or PIM Dense Mode.
- **Converged fabric:** The G8316 switch supports CEE and Data Center Bridging, which is ideal for storage connectivity (NAS and iSCSI). CEE enables clients to combine storage, messaging traffic, VoIP, video, and other data on a common data center Ethernet infrastructure. Data Center Bridging helps with iSCSI and FCoE with features such as Priority-based Flow Control, Enhanced Transmission Selection, and Congestion Notifications. FCoE enables highly efficient block storage over Ethernet for consolidating server network connectivity. As a result, clients can deploy a single server interface for multiple data types, which can simplify both deployment and management of server network connectivity, while maintaining the high availability and robustness required for storage transactions.

Features and specifications

Note: Features and specifications listed in this section are based on Networking OS 7.9.

The RackSwitch G8316 has the following features and specifications:



- Form factor: 1U rack mount switch
 - RackSwitch G8316 Rear-to-Front version for ports located in the rear of the rack matching System x®, ThinkServer®, BladeCenter® and Flex System® designs
 - RackSwitch G8316 Front-to-Rear version for ports located in the front of the rack matching airflow of the iDataPlex design
- Ports
 - 16 ports for 40 Gb Ethernet QSFP+ transceivers (40GBASE-SR4), SFP+ transceivers (10GBASE-SR or 10GBASE-LR with the optional QSFP-to-SFP+ adapter, 00D9676), QSFP+ to QSFP+ DAC cables (40GBASE-CR4), or QSFP+ to 4x 10 Gb SFP+ break-out cables. QSFP+ and SFP+ modules and DAC cables are not included and must be purchased separately (see Table 3).
 - One 10/100/1000 Ethernet port (RJ-45 connector) for out of band (OOB) management
 - One RS-232 serial port (mini-USB connector) that provides an additional means to configure the switch
 - One USB port for mass storage devices
- Scalability and performance
 - 40 Gb Ethernet ports with optional 10 GbE support for bandwidth optimization and performance
 - Non-blocking architecture with wire-speed forwarding of traffic and aggregated throughput of 1.28 Tbps
 - Full line rate performance with less than 1 microsecond switching latency
 - Media access control (MAC) address learning: automatic update, support for up to 128,000 MAC addresses
 - Up to 126 IP interfaces per switch (IP interface 128 is reserved for out-of-band management)
 - Static and LACP (IEEE 802.3ad) link aggregation, up to 64 trunk groups with up to 32 ports per trunk group
 - Support for jumbo frames (up to 9,216 bytes)
 - Broadcast/multicast storm control
 - IGMP snooping to limit flooding of IP multicast traffic
 - IGMP filtering to control multicast traffic for hosts participating in multicast groups
 - Configurable traffic distribution schemes over trunk links based on source/destination IP or MAC addresses, or both
 - Fast port forwarding and fast uplink convergence for rapid STP convergence
- Availability and redundancy
 - Virtual Router Redundancy Protocol (VRRP) for Layer 3 router redundancy
 - IEEE 802.1D STP for providing L2 redundancy
 - IEEE 802.1s Multiple STP (MSTP) for topology optimization, up to 32 STP instances are supported by a single switch
 - IEEE 802.1w Rapid STP (RSTP) provides rapid STP convergence for critical delay-sensitive traffic like voice or video
 - Per-VLAN Rapid STP (PVRST) enhancements
 - Layer 2 Trunk Failover to support active/standby configurations of network adapter teaming on compute nodes
 - Hot Links provides basic link redundancy with fast recovery for network topologies that require Spanning Tree to be turned off
- VLAN support
 - Up to 4095 VLANs supported per switch, with VLAN numbers ranging from 1 to 4095 (VLAN 4095 is used by the management network.)
 - Port-based and protocol-based VLANs
 - 802.1Q VLAN tagging support
 - Private VLANs support as defined in RFC 5517
- Security
 - VLAN-based, MAC-based, and IP-based access control lists (ACLs)
 - 802.1x port-based authentication
 - Multiple user IDs and passwords
 - User access control
 - Radius, TACACS+ and LDAP authentication and authorization
 - NIST 800-131A Encryption
 - Selectable encryption protocol
- Quality of Service (QoS)
 - Support for IEEE 802.1p, IP ToS/DSCP, and ACL-based (MAC/IP source and destination addresses, VLANs) traffic classification and processing



- Traffic shaping and re-marking based on defined policies
 - Eight output Class of Service (COS) queues per port for processing qualified traffic
 - Weighted Random Early Detection (WRED) with Explicit Congestion Notification (ECN) to help avoid congestion
 - Control plane protection (CoPP)
 - IPv4/IPv6 ACL metering
- IP v4 Layer 3 functions
 - Host management
 - IP forwarding
 - IP filtering with ACLs, up to 256 IPv4 ACLs supported
 - VRRP for router redundancy
 - Support for up to 128 static routes
 - Routing protocol support (RIP v1, RIP v2, OSPF v2, BGP)
 - Support for policy-based routing (PBR)
 - Support for DHCP Relay
 - Support for IGMP snooping and IGMP relay
 - Support for Protocol Independent Multicast (PIM) in Sparse Mode (PIM-SM) and Dense Mode (PIM-DM).
- IPv6 Layer 3 functions
 - IPv6 host management
 - IPv6 forwarding
 - Support for static routes
 - Support for OSPF v3 routing protocol
 - IPv6 filtering with ACLs, up to 128 IPv6 ACLs supported
- OpenFlow 1.0 and 1.3.1 support
- Virtualization
 - Virtual Fabric vNIC (virtual NIC) support with Ethernet, iSCSI, or FCoE traffic on vNICs
 - Virtual link aggregation groups (vLAGs)
 - 802.1Qbg Edge Virtual Bridging (EVB) is an emerging IEEE standard for allowing networks to become virtual machine (VM)-aware.
 - Virtual Ethernet Bridging (VEB) and Virtual Ethernet Port Aggregator (VEPA) are mechanisms for switching between VMs on the same hypervisor.
 - Edge Control Protocol (ECP) is a transport protocol that operates between two peers over an IEEE 802 LAN providing reliable, in-order delivery of upper layer protocol data units.
 - Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) allows centralized configuration of network policies that will persist with the VM, independent of its location.
 - EVB Type-Length-Value (TLV) is used to discover and configure VEPA, ECP, and VDP.
 - VMready support
 - Up to 4,096 virtual entities (VEs)
 - Automatic VE discovery
 - Up to 4,093 local or distributed VM groups for VEs
 - NMotion™ feature for automatic network configuration migration
- Converged Enhanced Ethernet
 - Priority-Based Flow Control (PFC) (IEEE 802.1Qbb) extends 802.3x standard flow control to allow the switch to pause traffic based on the 802.1p priority value in each packet's VLAN tag.
 - Enhanced Transmission Selection (ETS) (IEEE 802.1Qaz) provides a method for allocating link bandwidth based on the 802.1p priority value in each packet's VLAN tag.
 - Data Center Bridging Capability Exchange Protocol (DCBX) (IEEE 802.1AB) allows neighboring network devices to exchange information about their capabilities.
- Fibre Channel over Ethernet (FCoE)
 - FC-BB5 FCoE specification compliant
 - FCoE transit switch operations
 - FCoE Initialization Protocol (FIP) support for automatic ACL configuration
 - Supports 2,048 FCoE sessions with FIP Snooping by using Class ID ACLs
- Manageability
 - Industry-standard command line interface (isCLI)
 - Simple Network Management Protocol (SNMP V1, V2 and V3)
 - HTTP/HTTPS browser GUI
 - Telnet interface for CLI



- Secure Shell (SSH) v1 and v2 for CLI
- Secure Copy (SCP) for uploading and downloading the switch configuration via secure channels
- Service Location Protocol (SLP)
- Link Layer Discovery Protocol (LLDP) for discovering network devices
- Serial interface for CLI
- Scriptable CLI
- Dual software images
- Firmware image update via TFTP, FTP, or Secure FTP (sFTP)
- Network Time Protocol (NTP) and Precision Time Protocol (PTP) for switch clock synchronization
- Netconf (XML)
- Switch Center management application
- Monitoring
 - Switch LEDs for port status and switch status indication
 - Remote Monitoring (RMON) agent to collect statistics and proactively monitor switch performance
 - Port mirroring for analyzing network traffic passing through switch
 - Change tracking and remote logging with syslog feature
 - Support for sFLOW agent for monitoring traffic in data networks (separate sFLOW analyzer required elsewhere)

The following features are not supported with IPv6:

- Bootstrap Protocol (BOOTP) and DHCP
- RADIUS, TACACS+ and LDAP
- VMware Virtual Center (vCenter) for VMready
- Routing Information Protocol (RIP)
- Border Gateway Protocol (BGP)
- Protocol Independent Multicast (PIM)
- Virtual Router Redundancy Protocol (VRRP)
- sFLOW



Standards supported

The switch supports the following standards:

- IEEE 802.1AB Data Center Bridging Capability Exchange Protocol (DCBX)
- IEEE 802.1D Spanning Tree Protocol (STP)
- IEEE 802.1p Class of Service (CoS) prioritization
- IEEE 802.1s Multiple STP (MSTP)
- IEEE 802.1Q Tagged VLAN
- IEEE 802.1Qbg Edge Virtual Bridging
- IEEE 802.1Qbb Priority-Based Flow Control (PFC)
- IEEE 802.1Qaz Enhanced Transmission Selection (ETS)
- IEEE 802.1x port-based authentication
- IEEE 802.1w Rapid STP (RSTP)
- IEEE 802.3 10BASE-T Ethernet (management port only)
- IEEE 802.3ab 1000BASE-T copper twisted pair Gigabit Ethernet (management port only)
- IEEE 802.3ad Link Aggregation Control Protocol
- IEEE 802.3ae 10GBASE-SR short range fiber optics 10 Gb Ethernet
- IEEE 802.3ae 10GBASE-LR long range fiber optics 10 Gb Ethernet
- IEEE 802.3ba 40GBASE-SR4 short range fiber optics 40 Gb Ethernet
- IEEE 802.3ba 40GBASE-CR4 copper 40 Gb Ethernet
- IEEE 802.3u 100BASE-TX Fast Ethernet (management port only)
- IEEE 802.3x Full-duplex Flow Control

Connectors and LEDs

Figure 2 shows the front panel of the RackSwitch G8316.




Figure 2. Front panel of the RackSwitch G8316

The front panel of the G8316 contains the following components:

- LEDs that display the status of the switch and the network.
- One Mini-USB RS-232 console port that provides an additional means to configure the switch.
- One USB port for mass storage devices.



- 16x QSFP+ port connectors to attach QSFP+ transceivers for 40 Gb Ethernet connections or DAC cables for 40 Gb or 4x 10 Gb Ethernet connections. QSFP+ ports support SFP+ modules with an optional QSFP-to-SFP+ adapter (00D9676)
- One RJ-45 10/100/1000 Mb Ethernet port for out-of-band management
- An Ethernet link OK LED and an Ethernet Tx/Rx LED for each Ethernet port on the switch.

Figure 3 shows the rear panel of the RackSwitch G8316.




Figure 3. Rear panel of the RackSwitch G8316

The rear panel of the G8316 contains the following components:

- Two redundant, load-sharing hot-swap 450 W AC power supplies (IEC 320-C14 power connector)
- Four redundant hot-swap fan assemblies



Network cabling requirements

The network cables that can be used with the switch are listed in Table 4.

Table 4. G8316 network cabling requirements

Transceiver	Standard	Cable	Connector
40 Gb Ethernet			
QSFP+ 40GBASE-SR Transceiver (49Y7884)	40GBASE-SR4	10 m or 30 m MTP fiber optics cables (see Table 3); support for up to 100/150 m with OM3/OM4 multimode fiber	MTP
Direct attach cable	40GBASE-CR4	QSFP+ to QSFP+ DAC cables up to 7 m (see Table 3); QSFP+ to 4x SFP+ DAC break-out cables up to 5 m for 4x 10 GbE SFP+ connections out of a 40 GbE port (see Table 3)	QSFP+
10 Gb Ethernet			
SFP+ SR Transceiver (46C3447)*	10GBASE-SR	Up to 30 m with fiber optic cables supplied by Lenovo (see Table 3); 850 nm OM3 multimode fiber cable up to 300 m or up to 400 m with OM4 multimode fiber	LC
SFP+ LR Transceiver (90Y9412)*	10GBASE-LR	1310 nm single-mode fiber cable up to 10 km	LC
Management ports			
1 GbE management port	1000BASE-T	UTP Category 5, 5E, and 6 up to 100 meters	RJ-45
RS-232 management port	RS-232	DB-9-to-mini-USB or RJ-45-to-mini-USB console cable (comes standard with the switch)	Mini-USB

* Requires Mellanox QSFP to SFP+ Adapter (00D9676).

Warranty

The RackSwitch G8316 comes with a standard 1-year hardware warranty with Next Business Day (NBD), 9x5, Customer Replaceable Unit (CRU) warranty service from Lenovo. Software Upgrade Entitlement is based on the switch's warranty or post warranty extension and service contracts. Optional warranty and maintenance upgrades are available for the G8316 switch through Lenovo:

- Warranty service upgrades (3 or 5 years)
 - 24x7 onsite repair with 2-hour target response time
 - 24x7 onsite repair with 4-hour target response time
 - 9x5 onsite repair with 4-hour target response time
 - 9x5 onsite repair with next business day target response time
- Maintenance (post-warranty) service offerings (1 or 2 years)
 - 24x7 onsite repair with 2-hour target response time
 - 24x7 onsite repair with 4-hour target response time
 - 9x5 onsite repair with 4-hour target response time
 - 9x5 onsite repair with next business day target response time

Warranty service upgrade offerings are region-specific, that is, each region might have its own service types, service levels, response times, and terms and conditions. Not all covered types of warranty service offerings might be available in a particular region.

For more information about the Lenovo warranty service upgrade offerings that are available in your region, visit the Product Selector at the following website:

<https://www-304.ibm.com/sales/gss/download/spst/servicepac>



Physical specifications

The approximate dimensions and weight of the G8316 switch are as follows:

- Height: 44 mm (1.7 in.)
- Width: 439 mm (17.3 in.)
- Depth: 445 mm (19.0 in.)
- Weight: 10.0 kg (22.0 lb)

Operating environment

The G8316 switch is supported in the following operating environment:

- Temperature: 0 to 40 °C (32 to 104 °F).
- Relative humidity: Non-condensing, 10 - 90%
- Altitude: up to 1,800 m (6,000 feet)
- Acoustic noise: Less than 65 dB
- Airflow: Front-to-rear or rear-to-front cooling
- Electrical input: 50-60 Hz, 100-240 V AC auto-switching
- Typical power: 330 W

Agency approvals

The switch conforms to the following regulations:

- Safety certifications
 - UL60950-1
 - CAN/CSA 22.2 No.60950-1
 - EN 60950-1
 - IEC60950-1
 - NOM NYCE 019
 - GOST R MEK 60950-1
 - GB4943-2001
- Electromagnetic compatibility certifications
 - FCC 47CFR Part 15 Class A
 - EN 55022 Class A
 - ICES-003 Class A
 - VCCI Class A
 - AS/NZS CISPR 22 Class A
 - CISPR 22 Class A
 - EN 55024
 - EN 300386
 - CE
- Environmental
 - Reduction of Hazardous Substances (ROHS) 6

Typical configurations

This section describes two popular configurations.

40 GbE aggregation switch for east-west traffic

As more clients look to flatten their networking to exploit more of an east-west traffic flow that is required by a number



of today's applications and especially virtualized environments, the RackSwitch G8316 is an ideal tier-two switch for connecting multiple RackSwitch G8264 switches, which are likely being used at the edge to connect to multiple 10 GbE ports on servers, as shown in Figure 4.




Figure 4. Aggregation switch for east-west traffic

Using the RackSwitch G8316 as an edge switch for servers using 40 Gb Ethernet adapters

For environments such as HPC and others where I/O performance is important, 40 Gb is increasing in popularity. As more 40 Gb Ethernet adapters appear and costs fall in line, and applications start looking for the next step beyond 10 Gb, the RackSwitch G8316 is the ideal solution for connecting multiple servers per rack, as shown in Figure 5.




Figure 5. 40 GbE edge switch server connectivity

Related publications and links

For more information, see the following references:

- Offering Information page (to search on announcement letters, sales manuals, or both):
http://www.ibm.com/common/ssi/index.wss?request_locale=en
 On this page, enter **G8316**, select the information type, and then click **Search**. On the next page, narrow your search results by geography and language.
- RackSwitch G8316 product publications:
<http://ibm.com/support/entry/portal/documentation>
 - *Application Guide*
 - *Industry-Standard CLI Reference*
 - *Browser-Based Interface (BBI) Quick Guide*
 - *Menu-based CLI Command Reference*
- VMready:
<http://www.ibm.com/systems/networking/software/vmready>

Related product families

Product families related to this document are the following:

- [40 Gb Ethernet Connectivity](#)
- [Top-of-Rack Switches](#)

Notices

Lenovo may not offer the products, services, or features discussed in this document in all countries. Consult your local Lenovo representative for information on the products and services currently available in your area. Any reference to a Lenovo product, program, or service is not intended to state or imply that only that Lenovo product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any Lenovo intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any other product, program, or service. Lenovo may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

Lenovo (United States), Inc.
8001 Development Drive
Morrisville, NC 27560
U.S.A.
Attention: Lenovo Director of Licensing

LENOVO PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. Lenovo may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

The products described in this document are not intended for use in implantation or other life support applications where malfunction may result in injury or death to persons. The information contained in this document does not affect or change Lenovo product specifications or warranties. Nothing in this document shall operate as an express or implied license or indemnity under the intellectual property rights of Lenovo or third parties. All information contained in this document was obtained in specific environments and is presented as an illustration. The result obtained in other operating environments may vary. Lenovo may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Any references in this publication to non-Lenovo Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this Lenovo product, and use of those Web sites is at your own risk. Any performance data contained herein was determined in a controlled environment. Therefore, the result obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

© Copyright Lenovo 2022. All rights reserved.

This document, TIPS0842, was created or updated on June 8, 2015.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/TIPS0842>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/TIPS0842>.



Trademarks

Lenovo and the Lenovo logo are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. A current list of Lenovo trademarks is available on the Web at <https://www.lenovo.com/us/en/legal/copytrade/>.

The following terms are trademarks of Lenovo in the United States, other countries, or both:

Lenovo®
BladeCenter®
Flex System
NMotion®
RackSwitch
System x®
ThinkServer®
VMready®
Virtual Vision
iDataPlex®

The following terms are trademarks of other companies:

Hyper-V® and Microsoft® are trademarks of Microsoft Corporation in the United States, other countries, or both.

Other company, product, or service names may be trademarks or service marks of others.

