

## RackSwitch G8332

### Product Guide (withdrawn product)

The RackSwitch™ G8332 (shown in Figure 1) provides low latency, lossless performance, and a feature-rich design with key virtualization features, such as Converged Enhanced Ethernet (CEE)/Data Center Bridging (DCB), high availability, and enterprise class Layer 2 and Layer 3 functions. In addition, the RackSwitch G8332 also delivers excellent cost savings as you consider acquisition costs, energy costs, operational expenses, and ease of use and management for a 40 Gbps class switch. The RackSwitch G8332 has 32 QSFP+ ports and is suitable for clients using 10 Gigabit Ethernet or 40 Gigabit Ethernet connectivity (or both).

With latency below 600 nanoseconds, the RackSwitch G8332 is among industry leaders. This fact and the 2.56 Tbps throughput makes the RackSwitch G8332 an ideal offering for latency-sensitive applications, such as high-performance computing, financial applications, hosting environments, and cloud designs in Enterprise, Web 2.0, and VM Centric data centers. In addition, the G8332 supports VMready® with Virtual Vision, which enables the network to be virtual machine (VM) aware, and provides the capability to have a virtualization environment that is simpler and less expensive with exceptional performance.



Figure 1. RackSwitch G8332

### Did you know?

The RackSwitch G8332 supports several types of configurations: 10 Gigabit Ethernet, 40 Gigabit Ethernet, Converged Enhanced Ethernet (CEE/DCB), and iSCSI.

With exceptional port density and flexibility with break-out cables, the RackSwitch G8332 can support up to 96x 10 Gigabit Ethernet connections and still have 8x 40 Gigabit Ethernet uplinks.

The RackSwitch G8332 is SDN ready. The G8332 has support for OpenFlow. With OpenFlow, you can easily create user-controlled virtual networks, optimize performance dynamically, and minimize complexity when used with many industry available OpenFlow controllers.

The RackSwitch G8332 supports network virtualization through overlays as a VXLAN hardware gateway.

The RackSwitch G8332 is Cloud ready. NetBoot support offers automated network provisioning helping reduce operational costs. Support for industry standards like VEPA and VEB 802.1Qbg using the VMready features in Networking OS provide VM-aware networking and automation for multiple hypervisor environments.

## 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-HC3	Feature code for MTM 8036-HC4
Switch			
RackSwitch G8332 (Rear to Front)	8036BRX	A54G	None
RackSwitch G8332 (Front to Rear)	8036BFX	None	A54H
Miscellaneous options			
Console Cable Kit Spare	90Y9462	A2MG	A2MG
Adjustable 19" 4 Post Rail Kit	00D6185	A3KP	A3KP
Recessed 19" 4 Post Rail Kit	00CG089	None	A51M
iDataPlex Rail Kit	90Y3535	None	A1SZ
Air Inlet Duct for 483 mm RackSwitch	00D6060	A3KQ	None
Hot-Swappable, 750W CFF Power Supply Spare (rear to front)	00D5858	A2X7	None
Hot-Swappable, Front-to-Rear 550W CFF Power Supply Spare	00D5961	None	A3FN
Hot-Swappable, Rear-to-Front Fan Assembly Spare	00D6071	A54K	None
Hot-Swappable, Front-to-Rear Fan Assembly Spare	00D6073	None	A54J

The part numbers for the G8332 switches include the following items:

- One RackSwitch G8332 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 G8332 switch supports up to two redundant hot-swap 550 W AC power supplies for front-to-rear air flow and 750 W AC power supply for rear-to-front air flow (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 G8332 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 G8332 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 G8332 (rear-to-front airflow) switch for 4-post rack installations with the Adjustable 4-post Rail Kit (00D6185).

The G8332 (front-to-rear airflow) switch optionally supports recessed 19-inch, 4-post rack kit (00CG089) which is used when the switch is installed in the Intelligent Cluster™ Rack (MT 1410), Enterprise Rack (MT 9363), or PureFlex® System Rack (9363) with NeXtScale™ System. The G8332 (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 G8332 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-HC3 and 8036-HC4
<b>Rack power cables</b>		
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
<b>Line cords</b>		
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
Italy 10A line C13 to CEE 7/7 (2.8M)	39Y7921	6217
Power Cord Taiwan AC plug 10A/250V, 2.8M	00CG265	A53E
Power Cord Taiwan AC plug 15A/125V, 2.8M	00CG267	A53F

## Supported cables and transceivers

With the flexibility of the G8332 switch, clients can use the following types of connectivity:

- For 10 GbE links (supported on ports 2 - 25 only), clients can split out four 10 GbE ports for each 40 GbE port by using QSFP+ DAC breakout cables for distances up to 5 meters. For distances up to 100 meters, the 40GBASE-iSR4 QSFP+ transceivers can be used with OM3 optical MTP-to-LC break-out cables or up to 150 meters with OM4 optical MTP-to-LC break-out cables. For longer distances, the 40GBASE-eSR4 transceivers can be used with OM3 optical break-out cables for distances up to 300 meters or OM4 optical break-out cables for distances up to 400 meters.
- 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 meters, the 40GBASE-SR4/iSR4 QSFP+ transceivers can be used with OM3 multimode fiber or up to 150 meters when OM4 multimode fiber is used. For distances up to 300 meters, the 40GBASE-eSR4 QSFP+ transceivers can be used with OM3 multimode fiber or up to 400 meters when OM4 multimode fiber is used. For distances up to 10 kilometers, the 40GBASE-LR4 QSFP+ transceiver can be used with single mode fiber.

Table 3 lists the supported cables and transceivers.

Table 3. Supported transceivers and direct-attach cables

Description	Part number	Feature code (MTM 8036-HC3 / 8036-HC4)	Maximum quantity supported
<b>QSFP+ transceiver and cables - 40 GbE</b>			
Lenovo 40GBASE-SR4 QSFP+ Transceiver	49Y7884	A1DR	32
Lenovo 40GBASE-iSR4 QSFP+ Transceiver	00D9865	ASTM	32
Lenovo 40GBASE-eSR4 QSFP+ Transceiver	00FE325	A5U9	32
Lenovo 40GBASE-LR4 QSFP+ Transceiver	00D6222	A3NY	32
<b>Optical cables for 40 GbE QSFP+ SR4 transceivers</b>			
Lenovo 10m QSFP+ MTP-MTP OM3 MMF Cable	90Y3519	A1MM	32
Lenovo 30m QSFP+ MTP-MTP OM3 MMF Cable	90Y3521	A1MN	32
Lenovo 10m QSFP+ MTP-MTP OM3 MMF Cable (replaces 90Y3519)	00VX003	AT2U	32
Lenovo 30m QSFP+ MTP-MTP OM3 MMF Cable (replaces 90Y3521)	00VX005	AT2V	32
<b>QSFP+ breakout cables - 40 GbE to 4x10 GbE</b>			
Lenovo 1m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7886	A1DL	24
Lenovo 3m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7887	A1DM	24
Lenovo 5m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7888	A1DN	24
<b>QSFP+ direct-attach cables - 40 GbE</b>			
Lenovo 1m Passive QSFP+ DAC Cable	49Y7890	A1DP	32
Lenovo 3m Passive QSFP+ DAC Cable	49Y7891	A1DQ	32
Lenovo 5m Passive QSFP+ DAC Cable	00D5810	A2X8	32
Lenovo 7m Passive QSFP+ DAC Cable	00D5813	A2X9	32

## Benefits

The RackSwitch G8332 enables customers to easily build 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 RackSwitch G8332 is an ideal tier two switch to use when connecting a number of RackSwitch switches, such as the RackSwitch G8264, RackSwitch G8264T, or RackSwitch G8264CS, or embedded Ethernet offerings for Flex System (SI4093, EN4093R, or CN4093) with 40 Gigabit Ethernet uplink ports at the access layer. Other clients like the RackSwitch G8332 for the investment protection. These clients can use it with their 10 Gigabit Ethernet environments today, but can also use it in the future as they move to 40 Gigabit Ethernet.

The RackSwitch G8332 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 or iSCSI
- 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: 40 Gb Ethernet, Cloud and SDN

The RackSwitch G8332 offers the following benefits:

- **High performance:** This 40 Gigabit Ethernet low latency switch with 2.56 of Tbps throughput provides the best combination of low latency, non-blocking switching, and ease of management. The RackSwitch G8332 is also a single ASIC design, which promises consistent lower port-to-port latency.
- **Lower power and better cooling:** The RackSwitch G8332 uses as little as 270 W of power, which is a fraction of the power consumption of previous 40 Gigabit Ethernet 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 help automatically reducing power consumption.
- **High availability:** The RackSwitch G8332 also comes standard with hot-swap redundant power supplies and fans, making the switch highly reliable 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 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 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 G8332 switch interoperates seamlessly with other vendors' upstream switches.
- **Fault tolerance:** The G8332 switch learns alternative routes automatically and performs faster convergence in the unlikely case of a link, switch, or power failure. The RackSwitch G8332 uses technologies such as L2 trunk failover, advanced VLAN-based failover, VRRP, and Hot Links.
- **OpenFlow enabled:** The RackSwitch G8332 offers the 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:** Multicast supports IGMP Snooping v1, v2, and v3 with 3K IGMP groups, and Protocol Independent Multicast, such as PIM Sparse Mode or PIM Dense Mode.
- **Management:** Network Address Translation (NAT) allows a single device, such as a router, to act as an agent between the Internet (or "public network") and a local (or "private") network. This means that only a single,

unique IP address is required to represent an entire group of local IP addresses, to the external network. In addition to reducing the number of public IP addresses that are required, NAT also provides security by acting as a firewall between internal and external networks.

- **Standards:** The RackSwitch G8332 supports the IEEE 1588 Precision Time Protocol standard, which helps improve synchronization of the network while also improving performance, availability, and flexibility.

## Features and specifications

**Note:** Features and specifications listed in this section are based on Networking OS 7.7.

The RackSwitch G8332 has the following features and specifications:

- Form factor: 1U rack mount switch
  - RackSwitch G8332 Rear-to-Front version for ports located in the rear of the rack matching System x®, ThinkServer®, BladeCenter® and Flex System™ designs
  - RackSwitch G8332 Front-to-Rear version for ports located in the front of the rack matching airflow of iDataPlex and NeXtScale System® designs
- Ports
  - 32 ports for 40 Gb Ethernet QSFP+ transceivers (40GBASE-SR4 or 40GBASE-LR4), QSFP+ to QSFP+ DAC cables (40GBASE-CR4), or QSFP+ to 4x 10 Gb SFP+ break-out cables (ports 2 - 25 only). QSFP+ 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
  - Up to 96 10 Gb Ethernet SFP+ connections (with optional break-out cables)
  - Non-blocking architecture with wire-speed forwarding of traffic and aggregated throughput of 2.56 Tbps
  - Full line rate performance with less than 600 nanosecond 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
  - Port-based and protocol-based VLANs
  - Up to 4095 VLANs supported per switch, with VLAN numbers ranging from 1 to 4095 (VLAN 4095 is used by the management network.)

- 802.1Q VLAN tagging support
  - Ingress VLAN tagging support to tunnel packets through a public domain without altering the original 802.1Q tagging information
  - Private VLANs support
- 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
  - Network Address Translation (NAT)
  - 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
- Distributed Overlay Virtual Ethernet (DOVE) support
  - Transit Switch role tunnels traffic between DOVE nodes
  - DOVE VLAN Gateway role links legacy servers (virtual machines or native OS servers not serviced by a DVS) to the overlay
  - DOVE External Gateway role links the overlay to an external network
- Virtualization
  - 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
  - Native FC Forwarder (FCF) switch operations
  - End-to-end FCoE support (initiator to target)
  - FCoE Initialization Protocol (FIP) support for automatic ACL configuration
  - Supports 2,048 FCoE login sessions per VLAN (up to 4,096 login sessions in total for up to 12 FCF VLANs) with FIP Snooping by using Class ID ACLs
  - Fibre Channel services:
    - Name Server
    - Fabric Controller for State Change Notifications (SCNs)
    - Fabric login services
    - Fabric Configuration Server (FCS)
    - Fabric-Device Management Interface (FDMI)
    - Zoning
- 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
  - 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 802.3ba 40GBASE-SR4 short range fiber optics 40 Gb Ethernet
- IEEE 802.3ba 40GBASE-LR4 long 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
- FC-BB-5 for FCoE

## Connectors and LEDs

Figure 2 shows the front panel of the RackSwitch G8332.




Figure 2. Front panel of the RackSwitch G8332

The front panel of the G8332 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.
- 32x 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+ break out cables are supported in ports 2 - 25 only. Ports 1 and 26 -32 are for 40 Gb connectivity only.
- 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 G8332.




Figure 3. Rear panel of the RackSwitch G8332

The rear panel of the G8332 contains the following components:

- Two redundant, load-sharing hot-swap 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. G8332 network cabling requirements

Transceiver	Standard	Cable	Connector
40 Gb Ethernet			
QSFP+ 40GBASE-SR4 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
40GBASE-iSR4 QSFP+ Transceiver (00D9865)	40GBASE-SR4	10 m or 30 m MTP fiber optics cables (see Table 3); support for up to 100 m with OM3 multimode fiber or up to 150 m with OM4 multimode fiber	MTP
40GBASE-eSR4 QSFP+ Transceiver (00FE325)	40GBASE-SR4	10 m or 30 m MTP fiber optics cables (see Table 3); support for up to 300 m with OM3 multimode fiber or up to 400 m with OM4 multimode fiber	MTP
QSFP+ 40GBASE-LR4 Transceiver (00D6222)	40GBASE-LR4	1310 nm single-mode fiber cable up to 10 km	LC
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+
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

## Warranty

The RackSwitch G8332 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 G8332 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 G8332 switch are as follows:

- Height: 44 mm (1.7 in.)
- Width: 436 mm (17.2 in.)
- Depth: 503 mm (19.8 in.)
- Weight: 11.5 kg (25.0 lb)

## Operating environment

The G8332 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: 270 W

## Agency approvals

The switch conforms to the following regulations:

- Safety certifications
  - UL60950-1
  - CAN/CSA 22.2 No.60950-1
  - TUV GS to EN 60950-1
  - IEC60950-1, all country deviations
  - CNS 14336-1
  - Argentina Smark to IEC60950-1
  - GB4943.1-2011
  - EAC
  - NOM-019
- 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
  - KC Class A
  - CE

## Typical configurations

This section describes some configurations for the RackSwitch G8332.

### 40 GbE aggregation switch for east-west traffic

As more clients seek to flatten their networking to use more of an east-west traffic flow that is required by many of today's applications and especially virtualized environments, the RackSwitch G8332 is an ideal tier-two switch for

connecting multiple access switches or interconnect modules.

One example is when the RackSwitch G8332 is combined with RackSwitch G8264 by using SFP+ or G8264T using 10Gbase-T to the compute nodes, where clients can easily support up to 1536 nodes on 10 Gigabit Ethernet network with as few as a 3:1 oversubscription. Alternatively, clients can also build a smaller network by using the G8264CS while using the Omni Port to connect to their FC/FCoE SAN, as shown in Figure 4 and Figure 5.

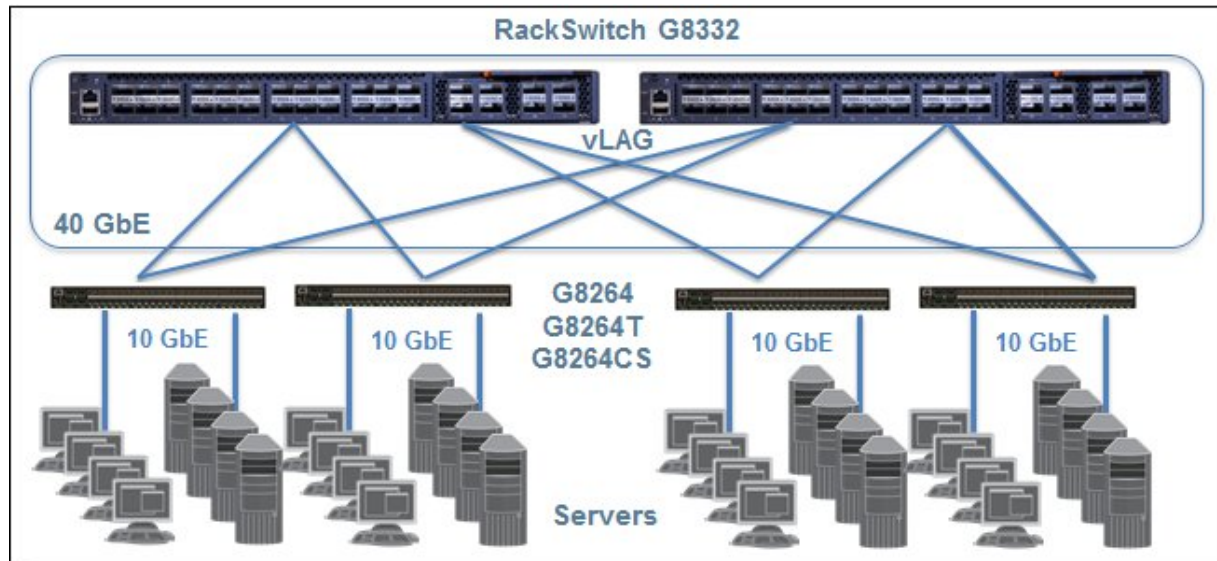


Figure 4. Aggregation switch for east-west traffic

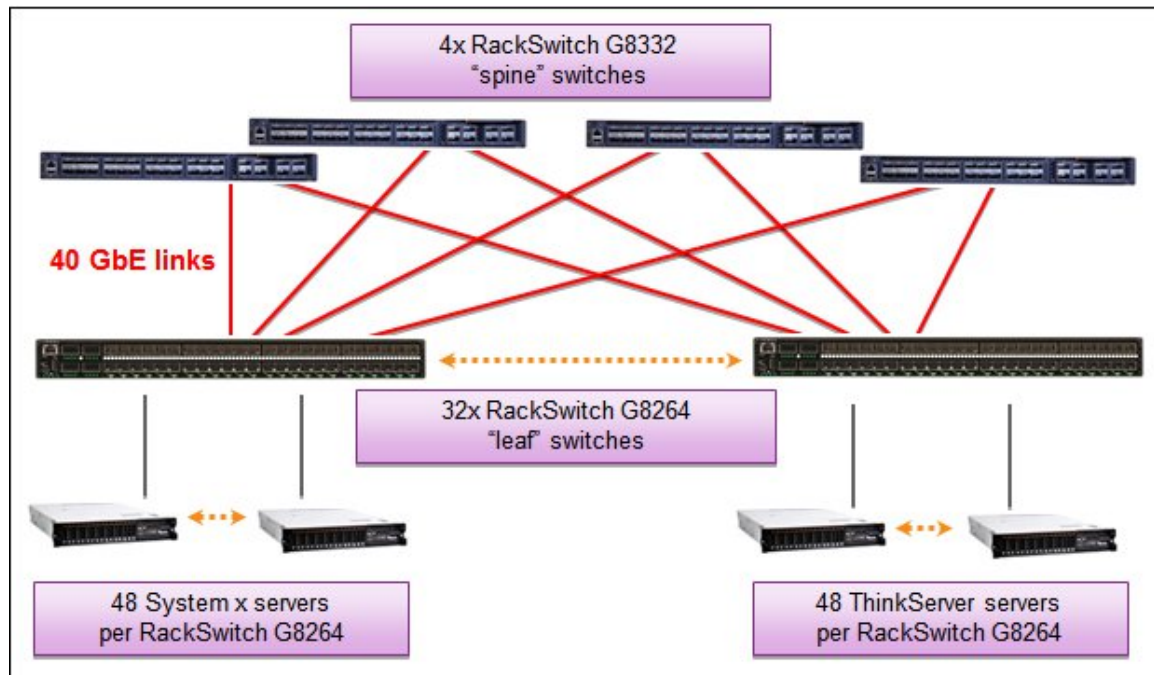


Figure 5. 1536 Node 10 Gigabit Ethernet Network (3:1 Oversubscribed)

Another example is the exceptional flat network implementations that can be built with embedded Ethernet offerings for Flex System (SI4093, EN4093R, or CN4093), which can be used easily to support up to 15 chassis of only a pair of RackSwitch G8332 switches. Figure 6 shows an example with nine chassis.




Figure 6. Flat network design supporting multiple Flex System chassis

### Building a POD

To build even larger networks, the RackSwitch G8332 offering allows for the creation of a simple two-tier design that can support up to 2,816 10 GbE server ports with a low oversubscription ratio of 2.1 to 1. Figure 7 shows an example.

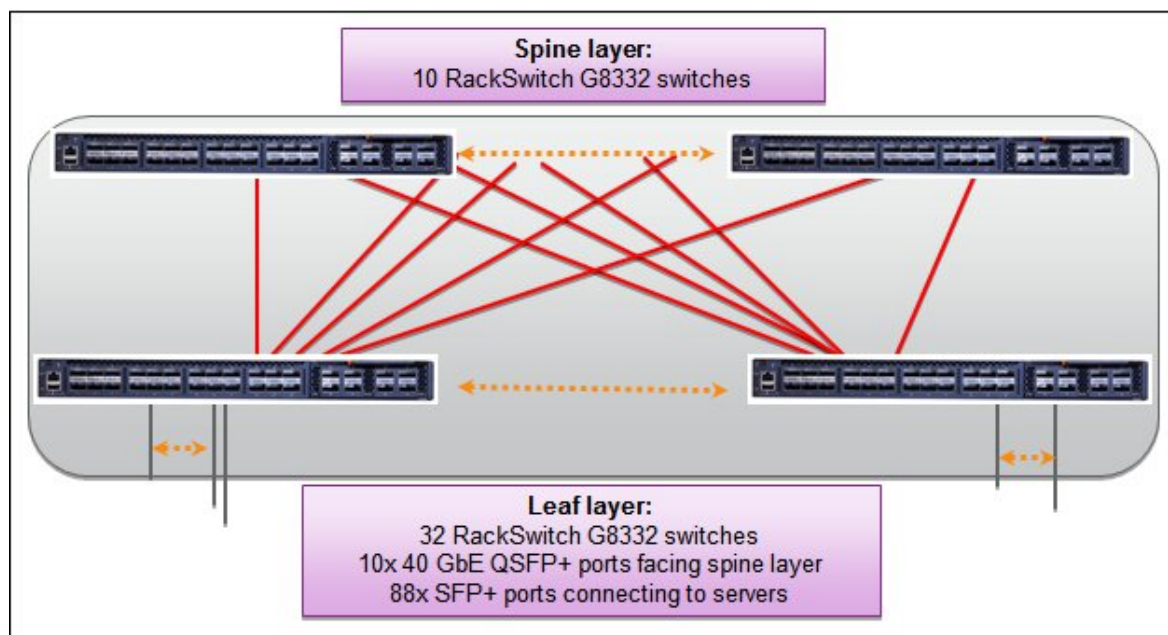


Figure 7. POD design with 2816 10 Gigabit Ethernet ports



## Using the RackSwitch G8332 switch 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 starts to increase in popularity. As more 40 Gigabit Ethernet adapters appear and costs fall, plus applications start looking for the next step beyond 10 Gigabit, the RackSwitch G8332 becomes the ideal solution for connecting multiple servers per rack, as shown in Figure 8.




Figure 8. Edge switch for 40 GbE server connectivity

## Related publications and links

For more information, see the following resources:

- Offering Information page (to search on announcement letters, sales manuals, or both):  
[http://www.ibm.com/common/ssi/index.wss?request\\_locale=en](http://www.ibm.com/common/ssi/index.wss?request_locale=en)

On this page, enter **RackSwitch G8332**, select the information type, and then click **Search**. On the next page, narrow your search results by geography and language.

- RackSwitch G8332 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*
  -

## 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 2024. All rights reserved.

This document, TIPS1139, was created or updated on August 6, 2015.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:  
<https://lenovopress.lenovo.com/TIPS1139>
- Send your comments in an e-mail to:  
[comments@lenovopress.com](mailto:comments@lenovopress.com)

This document is available online at <https://lenovopress.lenovo.com/TIPS1139>.

## 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  
Intelligent Cluster  
NMotion®  
NeXtScale  
NeXtScale System®  
RackSwitch  
System x®  
ThinkServer®  
VMready®  
Virtual Vision  
iDataPlex®

The following terms are trademarks of other companies:

Microsoft® and Hyper-V® 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.