

RackSwitch G8264

Product Guide (withdrawn product)

The RackSwitch™ G8264 that leverages 10Gb SFP+ and 40Gb QSFP+ Ethernet technology is specifically designed for the data center. It is ideal for today's big data, cloud, and optimized workload solutions. It is an enterprise class Layer 2 and Layer 3 full featured switch that delivers line-rate, high-bandwidth switching, filtering, and traffic queuing without delaying data. Large data-center grade buffers help keep traffic moving, while the redundant power and fans along with numerous high availability features help provide high availability for business sensitive traffic.

The RackSwitch G8264, shown in Figure 1, is ideal for latency sensitive applications such as high performance computing clusters and financial applications. In addition to the 10 Gb Ethernet (GbE) and 40 GbE connections, the G8264 also has the capability for traditional 1 GbE connections. The G8264 supports the newest protocols including Data Center Bridging/Converged Enhanced Ethernet (DCB/CEE) for Fibre Channel over Ethernet (FCoE) in addition to iSCSI and network-attached storage (NAS).



Figure 1. RackSwitch G8264

Did you know?

The G8264 supports Virtual Fabric, which helps clients significantly reduce cost and complexity related to I/O requirements of many virtualization deployments. Virtual Fabric helps reduce the number of multiple I/O adapters to a single dual-port 10 G adapter and reduces the number of cables and required upstream switch ports. By using Virtual Fabric, you can carve a dual-port 10 G server adapter into 8 virtual network interface cards (vNICs) and create dedicated virtual pipes between the adapter and switch for optimal performance, higher availability, and improved security. With Virtual Fabric, you can make dynamic changes and allocate bandwidth per vNIC so that you can adjust it over time without downtime.

The RackSwitch G8264 is the first 10 GbE switch to offer 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. With OpenFlow, clients can easily create user-controlled virtual networks, optimize performance dynamically, and minimize complexity when used with an industry compliant OpenFlow controller.

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 7309-HC3	Feature code for MTM 7309-HC4
Switch			
RackSwitch G8264 (Rear to Front)	7309G64	A1AC	None
RackSwitch G8264 (Front to Rear)	730964F	None	A1AD
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, 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	00D6073	None	A54J
Hot-Swappable, Rear-to-Front Fan Assembly Spare	00D6071	A54K	None

The part numbers for the G8264 switches include the following items:

- One RackSwitch G8264 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 G8264 switch supports up to two hot-swap power supplies (two power supplies come standard with the switch) and up to four 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 G8264 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 G8264 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 G8264 (rear-to-front airflow) switch for 4-post rack installations.

The G8264 (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 1410 Intelligent Cluster Rack, 9363 Enterprise Rack, or PureFlex System Rack with NeXtScale System. The G8264 (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 G8264 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 7309-HC3 and 7309-HC4
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 G8264 switch, clients can take advantage of the technologies that they require for multiple environments:

- For 1 GbE links, clients can use RJ-45 SFP transceivers with UTP cables up to 100 m. Clients that need longer distances can use a 1000BASE-SX transceiver, which can drive distances up to 220 meters by using 62.5 μ multi-mode fiber and up to 550 meters with 50 μ multi-mode fiber, or the 1000BASE-LX transceivers that support distances up to 10 kilometers using single-mode fiber (1310 nm).
- For 10 GbE server or upstream network connectivity, clients can use direct-attached copper (DAC) SFP+ cables for in-rack cabling and distances up to 7 m. These DAC cables have SFP+ connectors on each end,

and they do not need separate transceivers. For longer distances, the 10GBASE-SR transceiver can support distances up to 300 meters over OM3 multimode fiber or up to 400 meters over OM4 multimode fiber. The 10GBASE-LR transceivers can support distances up to 10 kilometers on single mode fiber. For extended distances, the 10GBASE-ER transceivers can support distances up to 40 kilometers on single mode fiber. 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. For distances up to 10 km, the 40GBASE-LR QSFP+ transceiver can be used with single mode fiber with LC connectors.

Table 3 lists the supported cables and transceivers.

Table 3. Supported transceivers and direct-attach cables

Description	Part number	Feature code (MTM 7309-HC3 / 7309-HC4)	Maximum quantity supported
SFP transceivers - 1 GbE			
Lenovo 1000BASE-T SFP Transceiver (does not support 10/100 Mbps)	00FE333	A5DL	48
Lenovo 1000BASE-SX SFP Transceiver	81Y1622	3269	48
Lenovo 1000BASE-LX SFP Transceiver	90Y9424	A1PN	48
SFP+ transceivers - 10 GbE			
Lenovo 10GBASE-SR SFP+ Transceiver	46C3447	5053	48
Lenovo 10GBASE-LR SFP+ Transceiver	90Y9412	A1PM	48
Lenovo 10GBASE-ER SFP+ Transceiver	90Y9415	A1PP	6
Optical cables for 1 GbE SFP SX and 10 GbE SFP+ SR transceivers			
Lenovo 0.5m LC-LC OM3 MMF Cable	00MN499	ASR5	48
Lenovo 1m LC-LC OM3 MMF Cable	00MN502	ASR6	48
Lenovo 3m LC-LC OM3 MMF Cable	00MN505	ASR7	48
Lenovo 5m LC-LC OM3 MMF Cable	00MN508	ASR8	48
Lenovo 10m LC-LC OM3 MMF Cable	00MN511	ASR9	48
Lenovo 15m LC-LC OM3 MMF Cable	00MN514	ASRA	48
Lenovo 25m LC-LC OM3 MMF Cable	00MN517	ASRB	48
Lenovo 30m LC-LC OM3 MMF Cable	00MN520	ASRC	48
SFP+ passive direct-attach cables - 10 GbE			
Lenovo 0.5m Passive SFP+ DAC Cable	00D6288	A3RG	48
Lenovo 1m Passive SFP+ DAC Cable	90Y9427	A1PH	48
Lenovo 1.5m Passive SFP+ DAC Cable	00AY764	A51N	48
Lenovo 2m Passive SFP+ DAC Cable	00AY765	A51P	48
Lenovo 3m Passive SFP+ DAC Cable	90Y9430	A1PJ	48
Lenovo 5m Passive SFP+ DAC Cable	90Y9433	A1PK	48
Lenovo 7m Passive SFP+ DAC Cable	00D6151	A3RH	48

SFP+ active direct-attach cables - 10 GbE			
Lenovo 1m Active SFP+ DAC Cable	95Y0323	A25A	48
Lenovo 3m Active SFP+ DAC Cable	95Y0326	A25B	48
Lenovo 5m Active SFP+ DAC Cable	95Y0329	A25C	48
Lenovo 1m Active DAC SFP+ Cable (replaces 95Y0323)	00VX111	AT2R	48
Lenovo 3m Active DAC SFP+ Cable (replaces 95Y0326)	00VX114	AT2S	48
Lenovo 5m Active DAC SFP+ Cable (replaces 95Y0329)	00VX117	AT2T	48
QSFP+ transceiver and cables - 40 GbE			
Lenovo 40GBASE-SR4 QSFP+ Transceiver	49Y7884	A1DR	4
Lenovo 40GBASE-LR4 QSFP+ Transceiver	00D6222	A3NY	4
Optical cables for 40 GbE QSFP+ SR4 transceivers			
Lenovo 10m QSFP+ MTP-MTP OM3 MMF Cable	90Y3519	A1MM	4
Lenovo 30m QSFP+ MTP-MTP OM3 MMF Cable	90Y3521	A1MN	4
Lenovo 10m QSFP+ MTP-MTP OM3 MMF Cable (replaces 90Y3519)	00VX003	AT2U	4
Lenovo 30m QSFP+ MTP-MTP OM3 MMF Cable (replaces 90Y3521)	00VX005	AT2V	4
QSFP+ breakout cables - 40 GbE to 4x10 GbE			
Lenovo 1m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7886	A1DL	4
Lenovo 3m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7887	A1DM	4
Lenovo 5m Passive QSFP+ to SFP+ Breakout DAC Cable	49Y7888	A1DN	4
QSFP+ direct-attach cables - 40 GbE			
Lenovo 1m Passive QSFP+ DAC Cable	49Y7890	A1DP	4
Lenovo 3m Passive QSFP+ DAC Cable	49Y7891	A1DQ	4
Lenovo 5m Passive QSFP+ DAC Cable	00D5810	A2X8	4
Lenovo 7m Passive QSFP+ DAC Cable	00D5813	A2X9	4

Benefits

The RackSwitch G8264 is considered particularly suited for these environments:

- The G8264 supports several types of configurations: 1 GbE, 10 GbE, 40 GbE, vNIC, Converged Enhanced Ethernet (CEE/DCB), NAS, and iSCSI.
- The G8264 with SFP+ technology is the ideal solution for clients who are looking for low latency and low power consumption.
- The G8264 supports stacking for up to eight switches using a single switch image and configuration file that shares one IP address and one management interface for simplified management.
- The G8264 supports Data Center Bridging (DCB), the group of protocols from IEEE that provide lossless Ethernet and that allow clients to reduce the costs of implementing either:
 - NAS or iSCSI convergence and priority-based flow control
 - FCoE by aggregating switch ports before connecting to more costly upstream Fibre Channel gateway devices.
- The G8264 can be configured in easy connect mode to allow for transparent and simple connectivity to the upstream network, enabling easy connectivity to upstream Cisco, Juniper or other networks without having to change those networks.
- Lenovo is a leader in network virtualization, offering the on-switch VMready® software that reduces the complexity of managing virtual machines (VMs) in the network. For more information, see VMready at: <http://www.ibm.com/systems/networking/software/vmready>
- The G8264 is SDN ready with its OpenFlow support. With OpenFlow, you can easily create user-controlled virtual networks, optimize performance dynamically, and minimize complexity when used with an OpenFlow controller.

- The G8264 plays a vital role with Virtual Fabric, which reduces costs and complexity in environments where four or more NICs are needed per server. An example is virtualization, where clients often need up to eight NICs per server. For more information, see Virtual Fabric at: <http://www.ibm.com/systems/networking/software/virtualfabric.html>

The RackSwitch G8264 offers the following benefits:

- **High performance**
The 10 Gb/40 Gb switch provides the best combination of low latency, non-blocking line-rate switching, and ease of management. It has a throughput of 1.28 Tbps.
- **Lower power and better cooling**
The RackSwitch G8264 uses as little as 330 W of power, which is a fraction of the power consumption of most competitive offerings. Unlike side-cooled switches, which can cause heat recirculation and reliability concerns, the front-to-rear or rear-to-front cooling design of the G8264 reduces data center air conditioning costs by having airflow match the servers in the rack. In addition, variable speed fans help to automatically reduce power consumption.
- **Stacking**
With the G8264, a single switch image and configuration file can be used for up to eight switches, sharing only one IP address and one management interface.
- **Virtual Fabric**
The G8264 can help customers address I/O requirements for multiple NICs while reducing cost and complexity. Virtual Fabric allows you to carve a physical NIC into multiple virtual NICs (between 2 - 8 vNIC) and to create a virtual pipe between the adapter and the switch for improved performance, availability, and security.
- **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, and Microsoft Hyper-V.
- **Layer 3 functionality**
The G8264 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**
RackSwitch switches perform seamlessly with other vendors' upstream switches.
- **Fault tolerance**
RackSwitch switches learn alternate routes automatically and perform 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 HotLink.
- **Multicast**
The G8264 supports IGMP Snooping v1, v2, and v3 with 2 K IGMP groups. It also supports Protocol Independent Multicast (PIM), such as PIM Sparse Mode or PIM Dense Mode.
- **Converged fabric**
The G8264 switch supports CEE and connectivity to FCoE gateways. CEE helps enable clients to combine storage, messaging traffic, VoIP, video, and other data on a common data center Ethernet infrastructure. FCoE helps enable 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.
- **OpenFlow enabled**

The RackSwitch G8264 is the first 10 GbE switch to offer benefits of OpenFlow. OpenFlow is the new open 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. The G8264 is also the ideal switch to use with industry compliant OpenFlow controllers.

- **Transparent networking capability**
With a simple configuration change to Easy Connect Mode, the RackSwitch G8264 becomes a transparent network device, invisible to the core, eliminating network administration concerns of Spanning Tree Protocol configuration/interoperability, VLAN assignments and avoids any possible loops.

By emulating a host NIC to the data center core, it accelerates the provisioning of VMs by eliminating the need to configure the typical access switch parameters.

Table 4 compares the features of the RackSwitch G8264 with the RackSwitch G8264T.

Table 4. Comparison of the RackSwitch G8264 and RackSwitch G8264T features

	RackSwitch G8264	RackSwitch G8264T
Port types	Based on SFP+	Based on 10GBase-T
Latency	Lower latency - 880 nanosecond	3.2 microseconds
Stacking	Yes	No
Typical power	Lower power - 330 watts	Competitive versus other 10GBase-T - 385 W
Copper cable length	Up to 7 m	Up to 100 m
Total distance	Up to 40 km with Transceiver	Limited to 100 m
TCO	Less than other SFP+	Lower TCO – for 100 m and less (based on adapter, cables and switch)

Features and specifications

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

The RackSwitch G8264 has the following features and specifications:

- **Form factor:** 1U rack mount switch
 - RackSwitch G8264 Rear-to-Front version for ports located in the rear of the rack matching System x®, BladeCenter® and Flex System® designs
 - RackSwitch G8264 Front-to-Rear version for ports located in the front of the rack matching airflow of iDataPlex® and NeXtScale™ System designs
- **Ports**
 - 48 ports for 1 Gb or 10 Gb Ethernet SFP/SFP+ transceivers (support for 1000BASE-SX, 1000BASE-LX, 1000BASE-T, 10GBASE-SR, 10GBASE-LR, or 10GBASE-ER) or SFP+ direct-attach copper cables. SFP+ modules and DAC cables are not included and must be purchased separately.
 - 4 ports for 40 Gb Ethernet QSFP+ transceivers, QSFP+ to QSFP+ DAC cables, or QSFP+ to 4x 10 Gb SFP+ break-out cables. QSFP+ modules and DAC cables are not included and must be purchased separately.
 - 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**
 - 1 Gb, 10 Gb, and 40 Gb Ethernet ports for bandwidth optimization and performance
 - Up to 64 10 Gb Ethernet SFP+ connections (with optional break-out cables)
 - Non-blocking architecture with wire-speed forwarding of traffic and aggregated throughput of 1.28 Tbps
 - Up to 960 Million packets per second (Mpps) with 880 nanoseconds 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
- 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
 - Ingress VLAN tagging support to tunnel packets through a public domain without altering the original 802.1Q tagging information
 - 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
 - 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
 - Up to 128 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 NICs (vNICs) with Ethernet, iSCSI, or FCoE traffic on vNICs
 - Unified Fabric Port (UFP)
 - Ethernet or FCoE traffic is supported on UFP virtual ports (vPorts)
 - Supports up to 1024 VLAN for the virtual ports
 - Supports either VMready or 802.1Qbg Edge Virtual Bridging (not both) on the same physical port
 - Integration with L2 Failover
 - 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
- 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
 - FCoE Link Aggregation Group (LAG) support
 - Supports 2,048 FCoE sessions with FIP Snooping by using Class ID ACLs
- Stacking
 - Up to eight switches in a stack - single IP management
 - 802.1Qbg support
 - vNIC or UFP support
 - Support for UFP with 802.1Qbg
- Manageability
 - Simple Network Management Protocol (SNMP V1, V2 and V3)
 - HTTP browser GUI
 - Telnet interface for CLI
 - Secure Shell (SSH)
 - Secure FTP (sFTP)
 - Service Location Protocol (SLP)
 - Link Layer Discovery Protocol (LLDP)
 - Serial interface for CLI
 - Scriptable CLI
 - Firmware image update (TFTP and FTP)
 - Network Time Protocol (NTP) and Precision Time Protocol (PTP) for switch clock synchronization
 - 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
- Stacking
- VMware Virtual Center (vCenter) for VMready
- Routing Information Protocol (RIP)
- Border Gateway Protocol (BGP)
- Protocol Independent Multicast (PIM)
- Virtual Router Redundancy Protocol (VRRP)
- sFLOW

The following features are not supported with Stacking (for a full list of features refer to the Networking OS Application Guide):

- FCoE
- IGMP Relay, IGMP Querier, and IGMPv3
- IPv6
- Policy-based routing
- Routing protocols (RIP, OSPF, BGP)
- sFLOW
- Virtual Router Redundancy Protocol (VRRP)

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 (frame tagging on all ports when VLANs are enabled)
- 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
- IEEE 802.3ab 1000BASE-T copper twisted pair Gigabit Ethernet
- 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 802.3ae 10GBASE-ER extended 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
- IEEE 802.3x Full-duplex Flow Control
- IEEE 802.3z 1000BASE-SX short range fiber optics Gigabit Ethernet
- IEEE 802.3z 1000BASE-LX long range fiber optics Gigabit Ethernet

Connectors and LEDs

Figure 2 shows the front panel of the RackSwitch G8264.

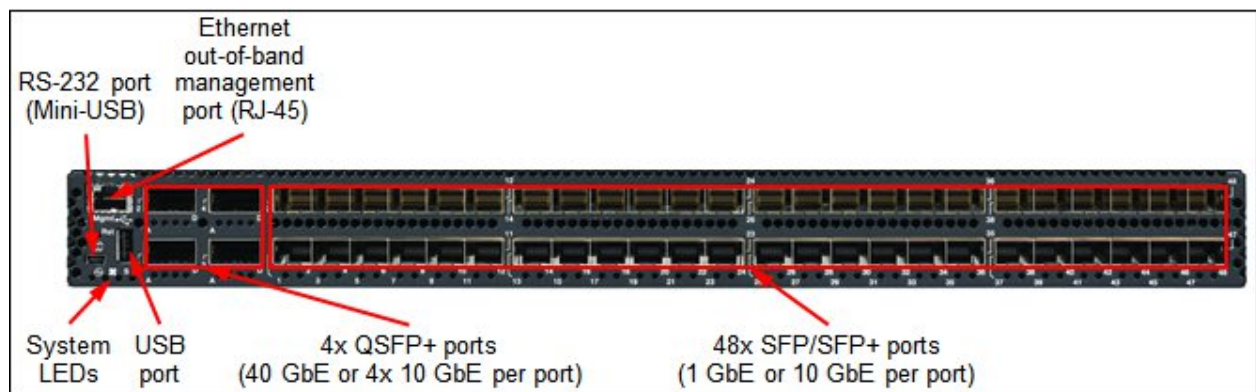


Figure 2. Front panel of the RackSwitch G8264

The front panel of the G8264 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.
- 48x SFP/SFP+ port connectors to attach SFP/SFP+ transceivers for 1 Gb or 10 Gb Ethernet connections or DAC cables for 10 Gb Ethernet connections.
- 4x QSFP+ port connectors to attach QSFP+ transceivers for 40 Gb Ethernet connections or DAC cables for 40 Gb or 4x 10 Gb Ethernet connections.
- 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 G8264.




Figure 3. Rear panel of the RackSwitch G8264

The rear panel of the G8264 contains the following components:

- Two Hot-swappable AC power supplies (IEC 320-C14 power connector)
- Four hot-swappable fan assemblies

Network cabling requirements

The network cables that can be used with the switch are listed in Table 5.

Table 5. G8264 network cabling requirements

Transceiver	Standard	Cable	Connector
40 Gb Ethernet			
40GBASE-SR4 QSFP+ Transceiver (49Y7884)	40GBASE-SR4	10 m or 30 m MTP fiber optics cables supplied by Lenovo (see Table 3); support for up to 100 m with OM3 multimode fiber or up to 150 m with OM4 multimode fiber	MTP
40GBASE-LR4 QSFP+ 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; 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			
10GBASE-SR SFP+ 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
10GBASE-LR SFP+ Transceiver (90Y9412)	10GBASE-LR	1310 nm single-mode fiber cable up to 10 km	LC
10GBASE-ER SFP+ Transceiver (90Y9415)	10GBASE-ER	1310 nm single-mode fiber cable up to 40 km	LC
Direct attach cable	10GSFP+Cu	SFP+ DAC cables up to 7 m (see Table 3)	SFP+
1 Gb Ethernet			
1000BASE-T SFP Transceiver (00FE333)	1000BASE-T	UTP Category 5, 5E, and 6 up to 100 meters	RJ-45
1000BASE-SX SFP Transceiver (81Y1622)	1000BASE-SX	Up to 30 m with fiber optic cables supplied by Lenovo (see Table 3); 850 nm multimode fiber cable up to 550 m (50 μ) or up to 220 m (62.5 μ)	LC
1000BASE-LX SFP Transceiver (90Y9424)	1000BASE-LX	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

Warranty

The RackSwitch G8264 comes with a standard 3-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 G8264 switch through Lenovo Services:

- Warranty service upgrades (3, 4, 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
- 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

Lenovo 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 Lenovo service upgrade 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 G8264 switch are as follows:

- Height: 44 mm (1.7 in.)
- Width: 440 mm (17.3 in.)
- Depth: 381 mm (15.0 in.)
- Weight: 6.4 kg (14.1 lb)

Operating environment

The G8264 switch is supported in the following operating environment:

- Temperature: 0 to 45 °C (32 to 113 °F).
- Relative humidity: Non-condensing, 10 - 85%
- Altitude: up to 3,049 m (10,000 feet)
- Acoustic noise: Less than 65 dB
- Airflow: Front-to-rear or rear-to-front cooling with redundant variable speed fans for reduced power draw
- 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
 - UL-UL60950-1 (First Edition)
 - C-UL to CAN/CSA 22.2 No.60950-1 (First Edition)
 - TUV/GS to EN 60950-1, Amendment A1-A4, A11
 - CB-IEC60950-1, all country deviations
- 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
- NEBS
 - GR-63-Core: NEBS, Physical Protection
 - GR-1089-Core: EMC and Electrical Safety for Network Telecommunications Equipment
 - Non PoE models: 24 port and 48 port
- Environmental
 - Reduction of Hazardous Substances (ROHS) 6

Typical configurations

The following configurations are the most typical for using the RackSwitch G8264:

- Rack-optimized server aggregation: 10 GbE attached rack servers
- Rack-optimized server aggregation: 1 GbE attached rack servers
- 2560-node 1 GbE network: Using next-generation 64-port 1U switches
- Virtual Fabric configuration with System x servers using the Emulex Virtual Fabric Adapter
- IP storage over 10 GbE: Changing the economics of storage

Rack-optimized server aggregation: 10 GbE attached rack servers

The RackSwitch G8264 is an ideal networking device for port aggregation, allowing the consolidation of one to three racks full of servers before connecting to an upstream switch.

- High-concentration of rack-optimized servers, for example:
 - System x 1U servers with a 10 Gb adapter installed
 - Easily supports 40 servers per rack.
 - Easily supports up to 48 servers or more, depending on uplink bandwidth required across two racks, including storage or UPS.
 - System x 2U servers with a 10 Gb adapter installed
 - Most clients will have up to twenty 2U servers per rack.
 - With the G8264 switch being able to support up to 64 SFP+ ports, a client can easily use a switch across three racks of servers, and more likely two for redundancy.
 - BladeCenter using any of the following modules in the chassis:
 - BladeCenter 1/10 Gb Uplink Ethernet Switch Module
With only two 10 Gb SFP+ ports per chassis, clients with larger data centers can easily support 24 chassis or more per RackSwitch G8264.
 - BladeCenter Virtual Fabric 10 Gb Switch Module

With support of up to ten 10 Gb SFP+ ports per switch, clients can easily support 4 - 6 chassis per RackSwitch, and if they are not concerned about oversubscription, they can support up to 24 or more per RackSwitch G8264.

- 10 Gb Ethernet Pass-Thru Module for BladeCenter
 For many clients looking to use a pass-through module, upstream costs can be expensive, especially if you are connecting to a core switch and are looking at up to fourteen 10 Gb ports per chassis. The G8264 provides a much more cost-effective alternative and can support up to four fully populated BladeCenter H chassis per RackSwitch. This alternative is also much better for larger BladeCenter H environments compared to a smaller 24-port switch.

- Low-profile, high-performance, 48/64-port 10 GbE switch needed for aggregation function per rack.

Figure 4 shows 10 GbE port aggregation, and Table 6 lists the features and benefits.




Figure 4. 10 GbE port aggregation

Table 6. Features and benefits

Features	Benefits
Line-rate, non-blocking, all 64-ports	Supports massive compute and virtualization workloads
Deterministically low latency	Faster application response times
Support for IGMP Snooping and L3 forwarding	Enables high-bandwidth, low-latency multicast applications
Standards-based Layer 2/3 protocols; industry standard CLI	Interoperates with the existing network; no learning curve

Rack-optimized server aggregation: 1 GbE attached rack servers

The design goal is to have the RackSwitch G8264 (10 GbE) switch at the distribution layer and RackSwitch G8000/G8052 (1 GbE) at the Edge/Access layer.

The G8264 is also ideal at the distribution layer when you are looking to aggregate connections for access layer switches like the G8000 or G8052 which support forty-eight 1 GbE ports to the servers and four 10 GbE uplinks. In either scenario, the G8264 provides exceptional value by providing the following support:

- Up to 12 G8000/G8052 switches (576 servers) when using the four 10 Gb uplinks per G8000/G8052, while still having the four 40 Gb uplinks from the G8264 to the core
- Up to 16 G8000/G8052 switches (768 servers) when using the four 10 Gb uplinks per G8000/G8052 and using 10 Gb SFP+ break-out cables in the four 40 Gb ports with no uplinks to a core

Keep in mind the following additional details:

- Logical configuration: Configure G8000/G8052 for Layer 2 and apply static routes for L3 forwarding
- Full Layer 2/3 Feature Set: STP, MSTP, RSTP, PVRST+; RIP v1/2, static routes, OSPF
- Security: 802.1X; RADIUS/TACACS+; Wire Speed ACLs, SSH v1, v2; HTTPS Secure BBI
- QoS: Up to eight queues per port, IEEE 802.1p and DiffServ prioritization

Figure 5 shows the aggregation design for 1 GbE attached rack servers.




Figure 5. 1 GbE port aggregation

2560-node 1 GbE network: Using next-generation 64 port 1U switches

The following high-level configuration is an example of how a configuration might look if the G8052 is used as leaf switches and the G8264 is used as a spine switch.

- All active links
 - Traffic load balanced across all links using Layer 3 ECMP
- 64 RackSwitch G8052 Leaf Switches for 2560 non-blocking 1 GbE ports
 - 48x 10/100/1000 server-facing ports each
 - 4x 10GBASE-SR fiber uplinks each
- Four RackSwitch G8264 *spine* switches for non-blocking core

Figure 6 shows the 2560-node logical design.




Figure 6. The 2560-node logical design

Virtual Fabric configuration with System x servers using the Emulex Virtual Fabric Adapters

The following example describes an Virtual Fabric implementation:

- Virtual Fabric for System x is built using industry standards and is the ideal solution for customers that require more than two NICs per server. Over the past couple of years many clients have seen their I/O requirements per server increase anywhere from 4 to 6 to 8 or even more NICs per server, especially with the adoption of virtualization. What Virtual Fabric does is provide clients with the ability to leverage a virtual NIC approach leveraging 10 Gb Ethernet technology to help reduce cost and complexity while also achieving better performance and more flexibility. It provides the following benefits:
 - Reduced costs
Acquisition and operating by using fewer adapters, cables, and upstream switch ports.
 - Reduced complexity
Fewer items to manage, based standards, making management easier with the addition of high availability and better security.
 - Better performance
A significant amount more I/O bandwidth per server and lower latency than traditional 1 Gb Ethernet.
 - Flexibility
The ability to carve up a dual-port 10 Gb adapter into 4 - 8 virtual NICs and create virtual pipes between the adapter and switch for higher availability and security. Not only can customers dynamically allocate I/O bandwidth to each virtual NIC, Virtual Fabric also provides the ability to change those allocations as needed. For example, a customer might have one configuration for workload during the day and another during the evenings. Virtual Fabric has the ability to make those changes on the fly without downtime.
- Virtual Fabric is designed in partnership with Emulex. This solution requires the following components:
 - Emulex 10 Gb Virtual Fabric Adapters for System x
For more details, see *Emulex 10Gb Virtual Fabric Adapter II and III Family for System x* , TIPS0844 at: <http://www.redbooks.ibm.com/abstracts/tips0844.html>
 - RackSwitch G8264

Figure 7 shows the Virtual Fabric logical design.




Figure 7. The Virtual Fabric logical design

IP storage over 10 GbE: Changing the economics of storage

G8264 benefits for IP storage applications are as follows:

- Lossless Ethernet Fabric
- Low latency ideal for HPC clusters, cloud, and financial applications
- Line-rate, high-bandwidth performance
- Low-power consumption with fewer components
- Low-cost, pay-as-you-grow 10 Gb/40 Gb storage network

Figure 8 shows IP storage connectivity.

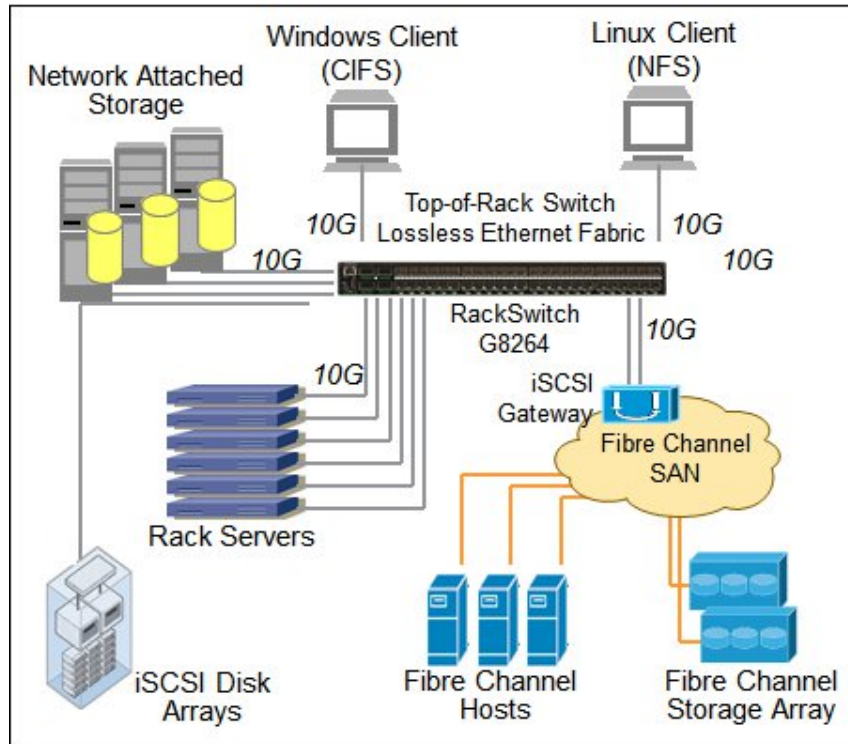


Figure 8. IP storage connectivity

Related publications

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 G8264 for System x and iDataPlex, select the information type, and then click **Search**. On the next page, narrow your search results by geography and language.
- RackSwitch G8264 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>
- Virtual Fabric
<http://www.ibm.com/systems/networking/software/virtualfabric.html>

Related product families

Product families related to this document are the following:

- [10 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 2025. All rights reserved.

This document, TIPS0815, was created or updated on June 9, 2015.

Send us your comments in one of the following ways:

- Use the online Contact us review form found at:
<https://lenovopress.lenovo.com/TIPS0815>
- Send your comments in an e-mail to:
comments@lenovopress.com

This document is available online at <https://lenovopress.lenovo.com/TIPS0815>.

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®

NeXtScale System®

System x®

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.

ibm.com® is a trademark of IBM in the United States, other countries, or both.

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