Lenovo Networking Best Practices: Transceivers, Cables and Media

Includes details on optical and copper media standards

Describes supported options on Lenovo switches

Includes discussions of connection speeds up to 100Gb

Discusses recommended media for various distances between devices

Scott Lorditch

LENOVO PRESS

Abstract

This paper describes the available options for connecting Lenovo® switches to each other and within a network. The recommended options depend on the desired speed of the connection and the distance between devices. This paper is intended for network designers and administrators who are involved in choosing the media to use in their network.

At Lenovo Press, we bring together experts to produce technical publications around topics of importance to you, providing information and best practices for using Lenovo products and solutions to solve IT challenges.

See a list of our most recent publications at the Lenovo Press web site:

http://lenovopress.com

Do you have the latest version? We update our papers from time to time, so check whether you have the latest version of this document by clicking the **Check for Updates** button on the front page of the PDF. Pressing this button will take you to a web page that will tell you if you are reading the latest version of the document and give you a link to the latest if needed. While you're there, you can also sign up to get notified via email whenever we make an update.

Contents

Introduction

This chapter provides information about the physical connections that are available on Lenovo Networking switches. It includes information about cables, transceivers, and low-level port settings for these connections, and specific considerations and guidelines.

This chapter includes the following topics:

Considerations for cabling and transceivers

This section describes the physical cabling requirements and other considerations for the various ports on Lenovo Networking switches.

Topics in this section:

Copper cabling standards – Cat 3 through Cat 6 and beyond

In this section, we describe copper cabling and connector standards. Most copper connections use what is referred to as an "RJ-45" connector. This designation originated from the world of telephony, and is more properly called an "8P8C" (eight position, eight conductor) connector. Actual RJ-45 connectors for multi-line phones do not interoperate with those used for Ethernet.

The copper cabling used for Ethernet was originally referred to as "UTP" – unshielded twisted-pair. These cables mostly do have shielding today, to enable higher bandwidth and range.

Copper cable standards are designated as categories, and come from the TIA/EIA 568 and ISO/IEC 11801 standards, as follows:

- Cat-3: voice telephony, 56Kb modems, 10BaseT (10Mb) Ethernet
- Cat-4: 16Mb Token Ring, no longer used
- Cat-5 and Cat-5e: 100BaseTX (2-pair) and 1000BaseT (4-pair). Most existing category 5 installations meet the 5e standard. This was very common in data centers until recently.
- Cat-6 and Cat-6a: 10GBaseT. Category 6 has a 55m range and Cat-6a has a 100m range. There is no Cat-6e.
- Cat-7/a: used for pre-standard 10GBaseT
- Cat-8.1 and 8.2: correspond to IEC11801 class I and II; short range (36m) 10Gb cabling. These could displace "DAC" cables (see below) in the future.

© Copyright Lenovo 2019. All rights reserved.

Fiber standards – OC-x

Lenovo supports two primary types of fiber connections today:

- Multi-mode fiber (MMF): ranges typically 100-150m; 50 micron diameter
- Single-mode fiber (SMF): much higher ranges, up to 10Km or more. 9 micron diameter. This is not typically used in a data center.

Multi-mode fiber quality is measured by “OM” designations, as follows:

- Older fiber types
  - FDDI quality: 62.5 micron MMF, with a maximum range of 220m with 1 GbE
  - OM1: 62.5 micron, maximum range of 275m with 1 GbE
  - OM2: 50 micron, maximum range of 550m with 1 GbE
- Current types
  - OM3:
    - 10Gb: maximum range of 300m
    - 40 Gb: maximum range of 100m (SR4)
    - 100Gb: maximum range of 100m (SR10)
  - OM4 – recommended for new builds.
    - 10Gb: maximum range of 400m
    - 40Gb: maximum range of 150m
    - 100Gb: maximum range of 150m

Fiber Connectors

Most connections use the “LC” (“little connector”). 40Gb and 100Gb multi-mode fiber runs require four pair of fiber strands and typically use the MPO/MTP connectors which support six pair of strands.

Lenovo also supports 40Gb “BiDi” (Bi-Directional) transceivers, designed to operate on one pair of multi-mode fibers by using multiple wavelengths. 100Gb BiDi is planned.

Optical break-out patch cords are available; these connect pairs of fibers from an MPO/MTP connector to four LC connectors.

RJ45 ports: 10Mb, 100MB and 1Gb ports

Ports in the 10/100/1000 Mb - 1 Gb speed range are available in two types: copper RJ45 (both built-in and SFP-based) and SFP-based fiber. This section describes the characteristics of these ports and modules.

Dedicated RJ45 Management ports

For switches that have a dedicated RJ45 Ethernet management port, this port supports 10/100/1000 Mb speeds and is for switch management only. It has no connection to the data ports; that is, packets that are coming in and out of the dedicated Ethernet management port can never be switched into or out of any of the data ports, and vice versa.

The dedicated management ports do not support any form of tagging, send only untagged packets, and expect to receive untagged packets. Under ENOS, when you are looking at the VLANs that are allowed on ports on a switch (for example, by using the `show int info` or `show int trunk` commands), dedicated management ports show that they are associated

with VLAN 4095 (under CNOS they are part of `vrf management` and have no VLAN membership). For all externally visible dedicated management ports, consider this VLAN to be a place holder to show that the port is not connected to the fabric; that is, it is not really sending and receiving packets on VLAN 4095. The relevant standard identifies VLAN 4095 as “reserved”; valid VLAN numbers are from 1-4094.

Built-in RJ45 data ports

RJ45 data ports adhere to standard IEEE 802.3 10/100/1000 Mb rules and default to auto negotiate. These ports are normal user data ports that can be configured and used by end nodes to pass packets between ports. Switches with built-in RJ45 data ports include the RackSwitch™ G8052, G7028, and G7052. The NE1032T and NE1072T switches have 10Gb RJ-45 ports, usually labeled as 10GBaseT.

SFP RJ45 and optical modules

All 1 Gb SFP ports that use approved SFP modules (RJ45 or fiber) support 1 Gbps speed only. No support for 10/100Mb is available in any SFP/SFP+ ports on Lenovo switches.

Lenovo switches support only Lenovo approved SFP modules, and usually do not support other vendors’ SFP modules. Attempting to use non-approved SFP modules will cause that port to become disabled until the non-supported SFP module is removed and replaced with an approved SFP.

All Lenovo SFP fiber modules use an LC type of connector. The type of fiber cabling that is supported by each SFP module varies by the module type. For more information about types of fibers and distances that are supported for various modules, see Table 1 on page 9.

10 Gb connections

For almost all 10 Gb switches, the 10 Gb ports use SFP+ modules, except for the following switches:

- NE1032T and NE1072T; these switches use built-in 10 GbaseT RJ45 connectors
- Internal ports on the NE2552E embedded switch module. These ports can run at either 10Gb or 25Gb depending on the configuration of the embedded server. These connections are hard-wired between the switch module and the enclosure-based server via the mid-plane of the enclosure.

The 10 Gb switches that support SFP+ based ports support a range of Lenovo transceivers, in addition to Direct Attach Cable (DAC), which is also known as TwinAx. This section describes these transceivers and cables with some considerations on usage.

Optical 10 Gb

The 10 Gb switches support a range of 10 Gb optical modules, and support might vary between switch models. As with 1 Gb SFP modules, Lenovo switches support only Lenovo approved SFP+ modules, and attempting to use a non Lenovo approved SFP+ module usually renders that port inoperable until an approved SFP+ module is installed. All Lenovo SFP+ modules use an LC type of connector. The type of fiber cabling that is used by each module varies by the module type. Lenovo also supports a 1/10Gb SFP+ transceiver, which is capable of either speed.

For more information about the types of fibers and distances that are supported, see Table 1 on page 9.

DAC

Unlike SFP+ modules, Lenovo does not lock out other vendors DACs, but does have specific limitations on lengths of DACs which are supported. Typically, 5 meter or shorter passive DACs work on all Lenovo switches that have SFP+ ports. Some Lenovo switches also support longer DAC cables, in addition to Active DACs. For more information about what your switch supports, see the installation guide and other documentation for that product.

In most instances, any MSA-compliant DAC works in Lenovo switches.

Because DACs are a single cable that potentially connects two different model or vendor switches, any DAC used must be compatible with both sides of the connection. For example, if one side supports a 10 meter DAC and the other side supports only a maximum of a 5 meter DAC, you cannot use a 10 meter DAC (you must use a 5 meter or shorter DAC). The same condition is true of active and passive DAC. If one side supports active and passive and the other side supports only passive, you can use only a passive DAC. In the rare instance where the two sides do not have a common supported DAC (for example, one side supports only active DAC, and the other side supports only passive DAC), DAC is not suitable for this environment. Instead, consider an optical SFP+ module.

AOC

Active Optical Cables (AOC) are similar in appearance to DAC cables, which use twinaxial copper wiring. They include transceivers on each end and so are not subject to vendor validation. Greater length is available with AOCs than with DACs because they use optical signaling.

10GBaseT

Only the NE1032T and NE1072T switches support 10GBaseT ports. These RJ45 10GBaseT ports can be 1 Gb or 10 Gb and have a default setting for auto negotiation. The 1 Gb connections can use Cat-5 or better cabling for up to 100 meter distances. Use Cat-6 or better cable for 10 Gb connections. Cat-6 cable supports up to 50 m at 10 Gb speeds, and Cat-6A cable supports up to 100 m at 10 Gb.

Considerations for 10 Gb ports

Consider the following points regarding 10 Gb ports:

- 10GBaseT ports on the NE1032T and NE1072T only support 1Gb and 10Gb.
- Similarly, there 10GBaseT SFP+ transceivers, which are supported in many Lenovo switches which have SFP+ capable ports (G70xx, G8264, G8272, G8296). These transceivers support only 10Gb, and do not support 1Gb or lower speeds.
- 10/100 Mb speeds are not supported for SFP+ ports on Lenovo switches.
- All Lenovo SFP+ 10 Gb optical modules use LC connectors. The other side of the link can be any type of connector (for example, LC, SC) if the optical type is the same (for example SR signaling on both sides, or LR signaling on both sides). Attempting to attach unlike modules (that is, connecting an SR transceiver to an LR transceiver) results in a poor connection or no connection at all.
- Also offered and supported is a 1/10Gb SFP+ SR transceiver.

40 Gb connections

Lenovo uses a QSFP+ format for all of its 40 Gb connections and supports DAC and optical modules for these ports. QSFP+ transceivers are physically larger than SFP/SFP+/SFP28 transceivers and require ports of that size on switches and other devices.

Optical 40 Gb

As with 1 Gb and 10 Gb transceiver modules, Lenovo switches support only Lenovo approved QSFP+ modules. Attempting to use a non Lenovo approved QSFP+ module renders that port inoperable until an approved QSFP+ module is installed.

The physical connection on xSR4-based (SR4, iSR4, and eSR4) optical 40 Gb modules is MPO (which is also referred to by the trademarked name of MTP). The LR4 optical module uses an LC connector. Lenovo also supports a 40Gb SR “BiDi” (bi-directional) transceiver. This transceiver has a QSFP+ form factor, but an LC connector, and uses a single pair of fibers.

Two different wavelengths of light are used, to allow 40Gb of traffic to be carried. The use case for this transceiver is to enable an existing single-pair fiber plant, typically one which had been used for 10Gb, to be re-purposed for 40Gb thus avoiding the need to install a four-pair fiber plant. Support for 100Gb SR BiDi is planned.

DAC and AOC cables

The following rules for 40 Gb DAC (which is also known as TwinAx) and AOC cables are similar to 10 Gb DAC:

- Lenovo does not lock out other vendors’ QSFP+ based cables.
- Lenovo does have restrictions on type (Active/Passive) and length.
- Both devices on each side of a QSFP+ DAC/AOC cable must agree on supported type (Active/Passive) and maximum length.

Most Lenovo QSFP+ 40 Gb ports also support being changed to a 4 x 10 Gb mode of operation. In that mode, they support a 4 x 10 Gb break out DAC (1 x QSFP+ male connector on one end, and 4 x SFP+ male connectors on the other end), or an MPO to 4xLC fiber connection.

Considerations for 40 Gb ports

Consider the following points regarding 40 Gb ports:

- The following Lenovo switches have ports which can run at 40Gb:
  - G8332, G8272, NE2572 and NE10032 have 100Gb ports which can be set to run at 40Gb using the `hardware profile portmode` command.
  - NE1032T and NE1072T
  - External 100Gb uplink ports on the NE2552E can be configured to run at 40Gb
- Most Lenovo 40 Gb ports support being operated as 1x 40 Gb port, or as 4x 10 Gb ports. To change between 1x 40 Gb and 4x10 Gb mode, use the `boot qsfp-40Gports X` or `no boot qsfp-40Gports X` commands (`X` = the port number), on switches running ENOS. The CNOS equivalent is `hardware profile portmode`, which has several options depending on the switch model.

  How the 40 Gb port is configured for speed can be seen with the `show boot qsfp` (CNOS: `show hardware profile portmode`) command, and any time the mode is changed, a reboot is necessary for the change between modes can take effect.
- Use the command `show boot qsfp` (`show hardware profile portmode`) to see the current mode for each QSFP+ port and the mode it assumes after the next reload
- One exception to the ability of a Lenovo QSFP+ port to be set to a 4x 10 Gb port is the G8332 switch. For this switch, the first port and last 7 ports (port 1 and ports 26 - 32) support 40 Gb only and cannot be set to 4x 10 Gb mode. Inversely, ports 2 - 25 can be set for 4x 10 Gb mode on the G8332 if desired (but default to 40 Gb).

Under CNOS, physical ports are always numbered consecutively, but ports in 4x10Gb mode are numbered as `ethernet 1/<port number>/<1-4>`.

- For the G8332 switch (with ENOS), the port number formatting of the 40 Gb ports is different from the other 40 Gb ports on other Lenovo switches and has the following considerations:
  - By default, all 40 Gb ports are numbered 1 - 32.
  - When ports are placed into 4x 10 Gb mode and the switch is reloaded for the change to take effect, the numbering for these four converted 10 Gb ports is X/1 - X/4 (where X is the base port number, and the /1 through /4 are for physically broken out 10 Gb ports). For example, if port 2 were converted to 4x 10 Gb ports, the numbering becomes 2/1, 2/2, 2/3, and 2/4 after the reload to take effect.
  - Only ports 2 - 25 can be broken out in to 4x 10 Gb modes. Port 1 and ports 26 - 32 cannot be broken out to 4x 10 Gb; the ports can operate in 40 Gb mode only.
- When a 40 Gb port is used as 4x 10 Gb ports, it is common to use a break out fiber or DAC, with a single QSFP_+ connection on one and four connections (eg four LC or four SFP+) on the other.
- The 40 Gb optical modules are available in various types, such as SR4, eSR4, and iSR4. Attaching a like-to-like module can be done (that is, SR4 to SR4); however, mixing different module types on each end might lead to issues.

  If the range is within 100 m, attaching the different variations of SR4 works.
- When you are using an optical breakout cable and configuring a 40 Gb port for 4x 10 Gb ports, the use of one of the SR4 variations (SR4, eSR4, or iSR4) on the 40 Gb side and attaching to a 10 Gb SR optical module has the following limitations:
  - Attaching to any SFP+ based SR module works, but attaching to some older non-SFP+ 10 Gb SR format modules (for example, X2 or Xenpak) might not work.
  - When you are using an SR4 or an iSR4 and breaking this configuration out to 10 Gb SFP+ SR modules, it might be necessary to have some minimum length cable to attenuate the signal from the SR4 module to not over drive the 10 Gb SR module.
  - When you are using an eSR4 for this type of breakout, you often can attach to SR modules of any type (SFP+, X2, or Xenpak)

25 Gb and 100 Gb connections

Consider the following points regarding 25 Gb and 100 Gb ports:

- 25Gb and 100Gb speeds are relatively new. They are displacing 10Gb and 40Gb on servers (NICs) as well as on switches.
- 25Gb transceivers are designated as SFP28, but are the same size as SFP+. 100Gb transceivers are designated as QSFP28 but are the same size as QSFP transceivers.
- The following Lenovo switches have ports which can be configured to operate at 25Gb and/or 100Gb:
  - NE2572 and NE2552E have 25Gb SFP28 ports and 100Gb QFP28 ports which can run at 4x25Gb with appropriate transceivers/cables.
    - The 100Gb port will take a QSFP28 optical transceiver which will use a MPO fiber connector; the fiber patch cord will have four LC connectors on the other end. SR and LR optics are supported.
    - Using a DAC breakout, the cable will have a QSFP28 connector on one end and four SFP28 connectors on the other end.

  - 25Gb ports are supported using SFP28 optical transceivers, or DAC (or AOC) cables with SFP28 transceiver heads on each end. The optical transceivers use LC connectors. At present, SFP28 transceivers are available only with SR optics; LR and ER are not available.

- The NE2572, NE2552E and NE10032 switches can configure ports to carry varying bandwidths but there are restrictions.
  - 25Gb ports can be set to operate at 10Gb. For the NE2572, groups of four contiguous ports are configured, not individual ports.
  - 100Gb ports can be set to operate at 40Gb, or to carry 4x10Gb, 4x25Gb or 2x50Gb.
  - There are restrictions on the speeds which the ASIC used in these switches can support; these restrictions are being loosened with each succeeding release of firmware. Please check the latest edition of the Application Guide for your switch for details.

Transceiver considerations

This section describes transceiver considerations.

Supported transceiver modules and limitations

For more information about part numbers for 1 Gb, 10 Gb and 40 Gb transceivers for Lenovo Networking switches, see the documentation for that specific product. Lenovo locks out transceivers other than those from Lenovo, or some legacy transceivers from IBM to ensure only fully tested and approved modules are used.

For more information about determining supported modules for a switch, see the *Lenovo Networking Product Catalog* available from:

http://bit.ly/2oe4iaW

Supported light levels on optical modules

When optical transceivers are used, it is important to understand what is acceptable for send and receive light levels. If a module is transmitting or receiving unexpected light levels, it can cause errors and in the case of excessive transmit light levels, premature failure of a module.

One of the more common issues with fiber connections is low light level, but how low is too low? Table 1 lists expected minimum and maximum light levels. If you receive signaling below and above these values, investigate and resolve the cause (most frequently dirty or defective optical cabling is the cause of low light, but it can also be caused by an incorrect cabling type or by bad SFP/SFP+/QSFP+ optical modules). Although these levels are part of the specifications, not all optics are equal and some might have issues when they are getting close to these levels. A good practice is to avoid running too close to the highs or lows that these numbers represent.

Table 1 General light levels and distances for common optical modules

Speed	Module P/N	TX Output Level	RX Input Level	Max Distances
1 Gb	SX SFP	112 µW (minimum) 501 µW (maximum)	20 µW (minimum) 1000 µW (maximum)	275 m OM1 MMF 550 m OM2 MMF
1 Gb	LX SFP	112 µW (minimum) 501 µW (maximum)	10 µW (minimum) 501 µW (maximum)	10k SMF

Speed	Module P/N	TX Output Level	RX Input Level	Max Distances
10 Gb	SR SFP+	186 µW (minimum) 794 µW (maximum)	102 µW (minimum) 794 µW (maximum)	26 m FDDI grade MMF 33 m OM1 MMF 82 m OM2 MMF 300 m OM3 MMF 400 m OM4 MMF
10 Gb	LR SFP+	151 µW (minimum) 1122 µW (maximum)	38 µW (minimum) 1112 µW (maximum)	10k SMF
10 Gb	ER SFP+	339 µW (minimum) 2512 µW (maximum)	26 µW (minimum) 794 µW (maximum)	40k SMF
40 Gb	SR4 QSFP+	174 µW (minimum) 1738 µW (maximum)	112 µW (minimum) 1738 µW (maximum)	100 m OM3 MMF 150 m OM4 MMF
40 Gb	iSR4 QSFP+	158 µW (minimum) 1738 µW (maximum)	1122 µW (minimum) 1738 µW (maximum)	100 m OM3 MMF 150 m OM4 MMF
40 Gb	eSR4 QSFP+	178 µW (minimum) 794 µW (maximum)	102 µW (minimum) 794 µW (maximum)	300 m OM3 MMF 400 m OM4 MMF
40 Gb	LR4 QSFP+	200 µW (minimum) 1698 µW (maximum)	43 µW (minimum) 1698 µW (maximum)	10k SMF
25 Gb	SR SFP28	630 uW (minimum) 1737 uW (maximum)	251 uW (minimum) 2187 uW (maximum)	70 m OM3 MMF 100 m OM4 MMF
25 Gb	LR SFP28	199 uW (minimum) 1584 uW (maximum)	74 uW (minimum) 2238 uW (maximum)	10k SMF
100 Gb	SR4 QSFP28	630 uW (minimum) 1737 uW (maximum)	251 uW (minimum) 2187 uW (maximum)	70 m OM3 MMF 100 m OM4 MMF
100 Gb	LR4 QSFP28	371 uW (minimum) 2818 uW (maximum)	138 uW (minimum) 3548 uW(maximum)	10k SMF

Important: Not all Lenovo switches support every module that is listed in Table 1, and some switches have limitations on how many of a certain type of module can be installed at one time (for example, the G8264 supports up to a maximum of six ER modules to be installed at one time).

Also, some optics might exceed these ratings for minimums, maximums, or both. For more information about limitations and exact minimums and maximums for your modules and switches, see the appropriate documentation.

On Lenovo switches, current light levels for a specific optical module can be seen via a GUI and command-line interface (CLI). The following commands are used to determine the light level of an optical module that uses ENOS syntax:

- The ENOS command to show all SFP/SFP+/QSFP+ ports and any associated transceivers light levels that are rated in µW (micro watts) is as follows:

  ```
  show interface transceiver
  ```

Considerations for low-level interface configurations

This section describes the low-level commands that can be applied to interfaces that affect basic operation on the physical medium.

Topics in this section:

Speed, duplex, and auto negotiation settings

All built-in (data and management) RJ45 ports support different speeds and duplex. This support includes negotiating 10 Mb, 100 Mb or 1000 Mb (1 Gb), and full or half duplex, in addition to independently enabling or disabling auto negotiation of these features.

One way in which Lenovo switches under ENOS vary from Cisco switches is what happens to auto negotiation when a port is hardcoded for speed and duplex. When you hardcode speed and duplex on a Cisco switch, it also automatically disables auto negotiation. With Lenovo switches, hardcoding speed and duplex does not disable auto negotiation (it means that it attempts to auto negotiate with the other side only for the speed/duplex that is hard set and still continue to negotiate other elements that are part of auto negotiation, such as auto-MDI/MDI-X and Flow Control). To fully disable auto negotiation on a Lenovo port, the command `no auto` must be run on that port. Before you set a port to `no auto`, the speed and duplex must first be set on that interface.

Under CNOS, the `speed`, `duplex`, and `flowcontrol` commands are used to control these settings. The `[no] auto` command under CNOS controls auto policy provisioning when using integration with VMware; this is a significant change.

Ports that are running 1 Gb and faster should auto negotiate. The reason is that auto negotiation negotiates speed, duplex, and flow control support (pause frames) and auto MDI/MDI-X (auto straight-through or cross over cable detection). Therefore, if auto negotiation is disabled, the type of cable (straight-through or cross over) becomes a possible issue.

With auto negotiation enabled, you can use either cable type and it automatically adapts to that cable. With auto negotiation disabled, if the switch port is connecting to another switch port, it must be a cross over. If the switch port is connecting to an end device (for example, a host NIC or router), it must be a straight through cable.

If 10/100 operation is desired, you must decide whether to hardcode speed and duplex. This decision is mostly based on what the customer standard is for these 10 Mb or 100 Mb connections.

For Lenovo 1 Gb ports that are part of an installable module (SFP), these ports are only 1 Gb and cannot be set to 10 Mb or 100 Mb.

Flow control

Flow control is also a feature that can be hardcoded or negotiated. By default, most 10 Gb and 40 Gb ports on Lenovo switches have flow control that is disabled (some Lenovo embedded switches default their 10 Gb server-facing ports to flow control that is enabled, but

the uplink 10 Gb and 40 Gb ports for these switches have default settings to disable flow control).

Flow control can be enabled on these interfaces; however, it is not a good idea. In particular, avoid flow control on most links between networking devices that might carry many conversations from many end devices.

This practice (flow control that is disabled between network devices) is true for most vendors, not only Lenovo. The reason is that a pause frame (used in flow control) that is received on a port tells that port to stop allowing packets for some time (defined in the pause frame), and it affects all incoming packets (except other pause frames). Therefore, if an upstream switch receives a pause frame that is generated by a downstream device, it discards all traffic from all devices sending in to this link, not only a single distressed device that generated the pause frames downstream.

The CEE (Converged Enhanced Ethernet) add Priority Flow Control (PFC), which enhances the original flow control with the ability to pause only traffic in a certain priority category, rather than all traffic on a port.

PFC is turned on with the `cee pfc enable` command, either globally or on a per-port basis. It is necessary to configure the various priority levels with the `cee pfc priority` command unless the default settings are acceptable.

Jumbo Frame considerations

All Lenovo Networking switches support jumbo frames up to a minimum of 9 KB (9216 bytes) packet sizes, which is often referred to as maximum transmission unit (MTU). Under ENOS, jumbo support to the maximum capability of the switch is enabled automatically and can not be disabled. Under some earlier releases of CNOS, jumbo support must be explicitly enabled on each port with the `mtu` command. Table 2, list which Lenovo switches support which size jumbo frame packets.

Important: Although jumbo frames are supported by default on all data interfaces on Lenovo switches, jumbo frames are not supported on any management interface on these switches.

Default MTU for CNOS switches on logical interfaces such as `interface VLAN <x>` is still set at 1560; this is to avoid interoperability issues with BGP. The MTU can be configured to support jumbo frames if desired.

Table 2 lists the maximum frame sizes for each RackSwitch top-of-rack switch model.

Table 2 MTU support for all Lenovo RackSwitch switches

Switch model	Max Jumbo/MTU frame size	Comments
RackSwitch G7028	12,288	ENOS only
RackSwitch G7052	12,288	ENOS only
RackSwitch G8052	9,216	ENOS only
RackSwitch G8264	9,216	ENOS only
RackSwitch G8332	9,216	ENOS or CNOS for physical interfaces

Switch model	Max Jumbo/MTU frame size	Comments
Switches running CNOS	9,216	Default for physical interfaces, but not for L3-interfaces, where the default is 1560.

To check whether a port is sending or receiving jumbo frames, the maintenance counters can be inspected on most switches; for example, by using the `show int port X maintenance-counter` command on ENOS, as shown in Example 1.

Example 1 Edited output showing frame size counters (ENOS)

```
G8264CS#show int port 17 maint
----------------------------------------------------------------
Maintenance statistics for port 17:
<<<Snip>>>
Receive 64 Byte Frame Counter
HW: MAC_GRx64 : 481000
Receive 65 to 127 Byte Frame Counter
HW: MAC_GRx127 : 5761
Receive 128 to 255 Byte Frame Counter
HW: MAC_GRx255 : 17954
Receive 256 to 511 Byte Frame Counter
HW: MAC_GRx511 : 9224
Receive 512 to 1023 Byte Frame Counter
HW: MAC_GRx1023 : 0
Receive 1024 to 1518 Byte Frame Counter
HW: MAC_GRx1518 : 0
Receive 1519 to 1522 Byte Good VLAN Frame Counter
HW: MAC_GRx1522 : 0
Receive 1519 to 2047 Byte Frame Counter
HW: MAC_GRx2047 : 0
Receive 2048 to 4095 Byte Frame Counter
HW: MAC_GRx4095 : 0
Receive 4096 to 9216 Byte Frame Counter
HW: MAC_GRx9216 : 0
Receive 9217 to 16838 Byte Frame Counter
HW: MAC_GRx16383 : 0
<<<Snip>>>
Transmit 64 Byte Frame Counter
HW: MAC_GTx64 : 0
Transmit 65 to 127 Byte Frame Counter
HW: MAC_GTx127 : 0
Transmit 128 to 255 Byte Frame Counter
HW: MAC_GTx255 : 17605
Transmit 256 to 511 Byte Frame Counter
HW: MAC_GTx511 : 0
Transmit 512 to 1023 Byte Frame Counter
HW: MAC_GTx1023 : 0
Transmit 1024 to 1518 Byte Frame Counter
HW: MAC_GTx1518 : 0
Transmit 1519 to 1522 Byte Good VLAN Frame Counter
HW: MAC_GTx1522 : 0
Transmit 1519 to 2047 Byte Frame Counter
HW: MAC_GTx2047 : 0
Transmit 2048 to 4095 Byte Frame Counter
HW: MAC_GTx4095 : 0
Transmit 4096 to 9216 Byte Frame Counter
HW: MAC_GTx9216 : 0
Transmit 9217 to 16383 Byte Frame Counter
```

```
HW: MAC_GTx16383 : 0
<<<Snip>>>
```

The equivalent command for CNOS is shown in Example 2.

Example 2 Jumbo frame counters under CNOS

```
G8272#sho int eth 1/1 counters detailed all
Ethernet1/1
  64 bit counters:
   <many counters snipped>
  All Port Counters:
   0.                          InPackets = 2242707
   1.                           InOctets = 167862042
   2.                         InUcastPkts = 0
   3.                         InMcastPkts = 2242707
   4.                         InBcastPkts = 0
   5.                         InJumboPkts = 0
   6.                  StormSuppressPkts = 0
   7.                         OutPackets = 836268
   8.                          OutOctets = 70896423
   9.                        OutUcastPkts = 1
  10.                        OutMcastPkts = 790150
  11.                        OutBcastPkts = 46117
  12.                        OutJumboPkts = 0
```

Author

Scott Lorditch is a Consulting System Engineer for Lenovo. He performs network architecture assessments and develops designs and proposals for solutions that involve Lenovo Networking products. He also developed several training and lab sessions for technical and sales personnel. Scott joined IBM as part of the acquisition of Blade Network Technologies® and joined Lenovo as part of the System x® acquisition from IBM. Scott spent almost 20 years working on networking in various industries, as a senior network architect, a product manager for managed hosting services, and manager of electronic securities transfer projects. Scott holds a BS degree in Operations Research with a specialization in computer science from Cornell University.

Thanks to the following people for their contributions to this project:

- Philippe Cauche
- David Watts

This paper is based on a chapter in the Lenovo Press book, *Lenovo Networking Best Practices for Configuration and Installation*. Thanks to the authors:

- Scott Irwin
- Scott Lorditch
- Ted McDaniel
- William Nelson
- Matt Slavin
- Megan Gilge

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.
1009 Think Place - Building One
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 2019. All rights reserved.
Note to U.S. Government Users Restricted Rights -- Use, duplication or disclosure restricted by Global Services Administration (GSA) ADP Schedule Contract

This document was created or updated on February 4, 2019.

Send us your comments via the **Rate & Provide Feedback** form found at
http://lenovopress.com/lp1071

Trademarks

Lenovo, the Lenovo logo, and For Those Who Do are trademarks or registered trademarks of Lenovo in the United States, other countries, or both. These and other Lenovo trademarked terms are marked on their first occurrence in this information with the appropriate symbol (® or ™), indicating US registered or common law trademarks owned by Lenovo at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of Lenovo trademarks is available on the Web at http://www.lenovo.com/legal/copytrade.html.

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

Blade Network Technologies®	RackSwitch™	System x®
Lenovo®	Lenovo(logo)®	

The following terms are trademarks of other companies:

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