

Lenovo RackSwitch G8264

Release Notes

For Enterprise Network Operating System 8.4

LenovoTM

Note: Before using this information and the product it supports, read the general information in the *Safety information and Environmental Notices and User Guide* documents on the *Lenovo Documentation CD* and the *Warranty Information* document that comes with the product.

First Edition (September 2016)

© Copyright Lenovo 2016
Portions © Copyright IBM Corporation 2014.

LIMITED AND RESTRICTED RIGHTS NOTICE: If data or software is delivered pursuant a General Services Administration "GSA" contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925.

Lenovo and the Lenovo logo are trademarks of Lenovo in the United States, other countries or both.

Release Notes

This release supplement provides the latest information regarding Lenovo Enterprise Network Operating System 8.4 for the Lenovo RackSwitch G8264 (referred to as G8264 throughout this document).

This supplement modifies and extends the following Enterprise NOS documentation for use with NOS 8.4:

- *Lenovo RackSwitch G8264 Application Guide for Lenovo Enterprise Network Operating System 8.4*
- *Lenovo RackSwitch G8264 Command Reference for Lenovo Enterprise Network Operating System 8.4*
- *Lenovo RackSwitch G8264 Installation Guide*

The publications listed here are available from the following website:

<http://publib.boulder.ibm.com/infocenter/systemx/documentation/index.jsp>

Please keep these release notes with your product manuals.

Hardware Support


Enterprise NOS 8.4 software is supported on the G8264, a high performance Layer 2-3 network switch

The G8124-E contains 48 10GbE SFP+ ports and four 40GbE QSFP+ ports. The SFP+ ports can be populated with optical or copper transceivers, or Direct Attach Cables (DACs). The QSFP+ ports can be populated with optical QSFP+ transceivers or DACs..

Note: If a DAC is not programmed to meet MSA specifications (including length identifier), the switch disables the port and generates a syslog message indicating that the DAC is not approved.

The G8264 is a 1U rack-mountable aggregation switch. The G8264 uses a wire-speed, non-blocking switching fabric that provides simultaneous wire-speed transport of multiple packets at low latency on all ports.

Figure 1. RackSwitch G8264 front panel



Updating the Switch Software Image

The switch software image is the executable code running on the G8264. A version of the image comes pre-installed on the device. As new versions of the image are released, you can upgrade the software running on your switch. To get the latest version of software supported for your G8264, go to the following website:

<http://www.support.lenovo.com/>

To determine the software version currently used on the switch, use the following switch command:

```
RS8264> show version
```

The typical upgrade process for the software image consists of the following steps:

- Load a new software image and boot image onto an SFTP, FTP, or TFTP server on your network.
- Transfer the new images to your switch.
- Specify the new software image as the one which will be loaded into switch memory the next time a switch reset occurs.
- Reset the switch.

For instructions on the typical upgrade process, see [“Loading New Software to Your Switch”](#) on page 6.

Loading New Software to Your Switch

The G8264 can store up to two different switch software images (called `image1` and `image2`) as well as special boot software (called `boot`). When you load new software, you must specify where it should be placed: either into `image1`, `image2` or `boot`.

For example, if your active image is currently loaded into `image1`, you would probably load the new image software into `image2`. This lets you test the new software and reload the original active image (stored in `image1`), if needed.

Attention: When you upgrade the switch software image, always load the new boot image and the new software image before you reset the switch. If you do not load a new boot image, your switch might not boot properly (To recover, see [“Recovering from a Failed OS Image Upgrade”](#) on page 10).

To load a new software image to your switch, you will need the following:

- The image and boot software loaded on an SFTP, FTP, or TFTP server on your network.

Note: Be sure to download both the new boot file and the new image file.

- The hostname or IP address of the SFTP, FTP, or TFTP server.

Note: The DNS parameters must be configured if specifying hostnames.

- The name of the new software image or boot file.

When the software requirements are met, use the following procedures to download the new software to your switch.

1. In Privileged EXEC mode, enter the following command:

```
RS8264# copy {sftp|tftp|ftp} {image1|image2|boot-image}
```

2. Enter the hostname or IP address of the SFTP, FTP or TFTP server.

```
Address or name of remote host: <name or IP address>
```

3. Enter the name of the new software file on the server.

```
Source file name: <filename>
```

The exact form of the name will vary by server. However, the file location is normally relative to the SFTP, FTP or TFTP directory (for example, `tftpboot`).

4. If required by the SFTP, FTP or TFTP server, enter the appropriate username and password.
5. The switch will prompt you to confirm your request.

Once confirmed, the software will begin loading into the switch.

- When loading is complete, use the following commands to enter Global Configuration mode to select which software image (`image1` or `image2`) you want to run in switch memory for the next reboot:

```
RS8264# configure terminal  
RS8264(config)# boot image {image1|image2}
```

The system will then verify which image is set to be loaded at the next reset:

```
Next boot will use switch software image1 instead of image2.
```

- Reboot the switch to run the new software:

```
RS8264(config)# reload
```

The system prompts you to confirm your request. Once confirmed, the switch will reboot to use the new software.

Note: If you select “No” when asked to confirm the reload, any changes made to the configuration since the last reboot will be lost.

Updating vLAG Switches with Lenovo Network Enterprise OS 8.x

Following are the steps for updating the software and boot images for switches configured with vLAG:

1. Save the configuration on both switches using the following command:

```
RS8264# copy running-config startup-config
```

2. Use FTP, STFP, or TFTP to copy the new Networking OS and boot images onto both vLAG switches. For more details, see [“Loading New Software to Your Switch” on page 6](#).
3. Shutdown all ports except the ISL ports and the health check port on the primary switch (Switch 1).
Note: Do not save this configuration.
4. Reload Switch 1. Switch 2 will assume the vLAG primary role. Once Switch 1 has rebooted, Switch 1 will take the vLAG secondary role.
5. Shutdown all ports except the ISL ports and the health check port on Switch 2.
Note: Do not save this configuration.
6. Reload Switch 2. Switch 1 will assume the vLAG primary role. Once Switch 2 has rebooted, make sure that Switch 1 is now the vLAG primary switch and Switch 2 is now the vLAG secondary switch.
7. Verify the all the vLAG clients have converged using the following command:

```
RS8264> show vlag information
```

Supplemental Information

This section provides additional information about configuring and operating the G8264 and Enterprise NOS.

The Boot Management Menu

The Boot Management menu allows you to switch the software image, reset the switch to factory defaults, or to recover from a failed software download.

You can interrupt the boot process and enter the Boot Management menu from the serial console port. When the system displays Memory Test, press **<Shift + B>**. The Boot Management menu appears.

```
Resetting the System ...
Memory Test .....
.
.
Boot Management Menu
  I - Change booting image
  C - Change configuration block
  R - Boot in recovery mode (tftp and xmodem download of images to
recover switch)
  Q - Reboot
  E - Exit
Please choose your menu option:
```

The Boot Management menu allows you to perform the following actions:

- To change the booting image, press **I** and follow the screen prompts.
- To change the configuration block, press **C**, and follow the screen prompts.
- To perform a TFTP/Xmodem download, press **R** and follow the screen prompts.
- To reboot the switch, press **Q**. The booting process restarts.
- To exit the Boot Management menu, press **E**. The booting process continues.

Recovering from a Failed OS Image Upgrade

The Boot Management menu allows you to perform fundamental device management operations, such as selecting which software image will be loaded, resetting the G8264 to factory defaults or recovering from a failed image download.

Use the following procedure to recover from a failed image upgrade.

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports XModem Download (for example, HyperTerminal, SecureCRT, or PuTTY) and select the following serial port characteristics:
 - Speed: 9,600 bps
 - Data Bits: 8
 - Stop Bits: 1
 - Parity: None
 - Flow Control: None
3. To access the Boot Management menu, you must interrupt the boot process from the Console port. Boot the G8264 and when the system begins displaying Memory Test progress (a series of dots), press **<Shift + B>**.

The Boot Management menu will display:

```
Resetting the System ...
Memory Test .....
.
.
Boot Management Menu
  I - Change booting image
  C - Change configuration block
  R - Boot in recovery mode (tftp and xmodem download of images to
recover switch)
  Q - Reboot
  E - Exit
Please choose your menu option:
```

4. Select **R** for Boot in recovery mode. You will see the following display:

```
Entering Rescue Mode.
Please select one of the following options:
  T) Configure networking and tftp download an image
  X) Use xmodem 1K to serial download an image
  P) Physical presence (low security mode)
  R) Reboot
  E) Exit

Option?:
```

- If you choose option **X** (Xmodem serial download), go to [Step 5](#).
- If you choose option **T** (TFTP download), go to [Step 6](#).

5. **Xmodem download:** When you see the following message, change the Serial Port characteristics to 115,200 bps:

```
Change the baud rate to 115200 bps and hit the <ENTER> key before
initiating the download.
```

- a. Press **<Enter>** to set the system into download accept mode. When the readiness meter displays (a series of "C" characters), start XModem on your terminal emulator. You will see a display similar to the following:

```
... Waiting for the <Enter> key to be hit before the download can
start...
CC
```

- b. When you see the following message, change the Serial Port characteristics to 9,600 bps:

```
Change the baud rate back to 9600 bps, hit the <ESC> key.
```

- c. When you see the following prompt, press **<Enter>** to start installing the image. If the file is a software image, enter the image number:

```
Install image as image 1 or 2 (hit return to just boot image): 1
```

The image install will begin. After the procedure is complete, the Recovery Mode menu will be re-displayed.

```
Please select one of the following options:
  T) Configure networking and tftp download an image
  X) Use xmodem 1K to serial download an image
  P) Physical presence (low security mode)
  R) Reboot
  E) Exit

Option?:
```

Continue to [Step 7](#).

6. **TFTP download:** The switch prompts you to enter the following information:

```
Performing TFTP rescue. Please answer the following questions (enter 'q'
to quit):
IP addr      :
Server addr:
Netmask     :
Gateway     :
Image Filename:
```

a. Enter the required information and press **<Enter>**. You will see a display similar to the following:

```
Host IP      : 10.10.98.110
Server IP    : 10.10.98.100
Netmask     : 255.255.255.0
Broadcast   : 10.10.98.255
Gateway     : 10.10.98.254
Installing image G8264-8.4.1.0_OS.imgs from TFTP server 10.10.98.100
```

b. If the file is a software image, you will be prompted to enter an image number:

```
Install image as image 1 or 2 (hit return to just boot image): 2
```

The following message is displayed when the image download is complete:

```
Image2 updated succeeded
Updating install log. File G8264-8.4.1.0_OS.imgs installed from
10.10.98.100 at 15:29:30 on 12-3-2015
Please select one of the following options:
  T) Configure networking and tftp download an image
  X) Use xmodem 1K to serial download an image
  P) Physical presence (low security mode)
  R) Reboot
  E) Exit

Option?:
```

Continue to [Step 7](#).

7. Image recovery is complete. Perform one of the following steps:

- Press **R** to reboot the switch.
- Press **E** to exit the Boot Management menu.
- Press the Escape key (**<Esc>**) to re-display the Boot Management menu.

Recovering from a Failed Boot Image Upgrade

Use the following procedure to recover from a failed boot image upgrade.

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports Xmodem download (such as HyperTerminal, CRT, or PuTTY) and select the following serial port characteristics:
 - Speed: 9600 bps
 - Data Bits: 8
 - Stop Bits: 1
 - Parity: None
 - Flow Control: None
3. Boot the switch and access the Boot Management menu by pressing **<Shift B>** while the Memory Test is in progress and the dots are being displayed.
4. Select **R** to boot in recovery mode. Then choose option **X** (Xmodem serial download). You will see the following display:

```
Perform xmodem download
To download an image use 1K Xmodem at 115200 bps.
```

5. When you see the following message, change the Serial Port characteristics to 115200 bps:

Change the baud rate to 115200 bps and hit the <ENTER> key before initiating the download.

- a. Press <Enter> to set the system into download accept mode. When the readiness meter displays (a series of "C" characters), start Xmodem on your terminal emulator. You will see a display similar to the following:

```
Extracting images ... Do *NOT* power cycle the switch.
**** RAMDISK ****
UnProtected
38 sectors
Erasing Flash...
..... done
Erased 38 sectors
Writing to Flash... 9... 8... 7... 6... 5... 4... 3... 2... 1... done
Protected 38 sectors
**** KERNEL ****
UnProtected
24 sectors
Erasing Flash...
..... done
Erased 24 sectors
Writing to Flash... 9... 8... 7... 6... 5... 4... 3... 2... 1...
```

- b. When you see the following message, change the Serial Port characteristics to 9600 bps:

Change the baud rate back to 9600 bps, hit the <ESC> key.

Boot image recovery is complete.

VLAGs

For optimal VLAG operation, adhere to the following configuration recommendations:

- Any port-related configuration, such as applied ACLs, must be the same for all ports included in the same VLAG, across both peer switches.
- Configure VLAG health checking as shown in the *Application Guide*.

After configuring VLAG, if you need to change any configuration on the VLAG ports, you must adhere to the following guidelines:

- If you want to change the STP mode, first disable VLAG on both the peers. Make the STP mode-related changes and re-enable VLAG on the peers.
- If you have MSTP on, and you need to change the configuration of the VLAG ports, follow these steps:

On the VLAG Secondary Peer:

1. Shutdown the VLAG ports on which you need to make the change.
2. Disable their VLAG instance using the command:
RS8264 (config)# **no vlag adminkey** <key> **enable**
or
RS8264 (config)# **no portchannel** <number> **enable**
3. Change the configuration as needed.

On the VLAG Primary Peer:

4. Disable the VLAG instance.
5. Change the configuration as needed.
6. Enable the VLAG instance.

On the VLAG Secondary Peer:

7. Enable the VLAG instance.
8. Enable the VLAG ports.

Note: This is not required on non-VLAG ports or when STP is off or when STP is PVRST.

New and Updated Features

Enterprise NOS 8.4 for the G8264 has been updated to include several new features, summarized in the following sections. For more detailed information about configuring G8264 features and capabilities, refer to the complete Enterprise NOS 8.4 documentation as listed on [page 3](#).

UFP Support for Eight VNICs Per Port with Emulex CNAs

Added support for 8 vNICs per port for the UFP feature when working with Emulex Rackserver CNAs.

UFP Support for Four VNICs Per Port with QLogic NICs

Four VNIC per port UFP support is available with QLogic NICs. For details about QLogic NICs that support UFP with Lenovo Rackswitches and the required Firmware and Driver version that needs to be installed on the NICs, see QLogic documentation.

TACACS+ Two Level Authentication

When TACACS+ is used to control switch access and the CLI Enable Mode is configured to require a password, a second authentication for the Enable command will be required.

Support for 256 PVRST Instances

The number of STP groups supported in the PVRST protocol has been increased to 256. For optimal performance, it is recommended that you stay within a maximum of ten interfaces or aggregated links (including vLAGs) with 4K MAC table entries when deploying 256 STP groups.

Security Feature Support

This feature supports Secure I/O Module (SIOM) framework by managing security policies based on the IOM mode (Secure/Legacy). A secure version of LDAP using startTLS and LDAPS is supported. Cryptographic Provisioning is also supported.

UFP+PVLAN with PVLAN Trunk and Promiscuous Port Support

UFP has been enhanced to support Full Private VLAN Trunk and Promiscuous mode on a UFP port.

Certificate Signing Request (CSR)

This feature enhances the certificate management capabilities on the switch by incorporating the ability to generate a Certificate Signed Request which can be submitted to an external Certificate Authority (CA) for obtaining a signed certificate. The capability to support CSR and process the CA signed certificate thereof is made available from multiple user interfaces including BBI, SNMP, and CLI.

Password Encryption

This feature enables all passwords in the switch to be encrypted using industry-standard encryption methods.

ACL Redirect to Trunk Support

This feature enables the switch to steer traffic based on a combination of a physical port and Layer 2-4 protocol header fields.

vLAG Peer Gateway

This feature enables a vLAG switch to act as an active gateway for packets that are addressed to the router MAC address of the vLAG peer.

SNMPv1 Default Community String Removal

This feature removes the default read/write community string for SNMP v1/2c from the factory default configuration.

Default UserID with Default Password That Must Be Reset at First Login

This feature adds a default user "USERID" at UID 1 with default password "PASSWORD" and prompts for a change of the default password at first login.

DHCP Snooping

DHCP Snooping acts like a firewall between untrusted hosts and DHCP servers. It provides security by filtering untrusted DHCP packets and by building/maintaining a DHCP Snooping binding table.

Configurable Syslog Port

This feature adds support for a configurable syslog host server port for both primary and secondary syslog host servers.

STP Debugging Enhancement

The STP display has been enhanced to display the current and previous STP STG root information.

IGMP Reports Over the Current 3K Limit Must Be Forwarded to the Mrouter

IGMP reports over the current 3K limit will be forwarded to an Mrouter. If no Mrouter exists, such IGMP reports will be discarded. IGMPv2 leaves for groups not known by the switch will also be forwarded to the multicast router.

SLP IPv6 Support

SLP has been enhanced with support for IPv6.

MSTP Stacking

MSTP protocol is supported in stacking mode.

Known Issues

This section describes known issues for Enterprise NOS 8.4 on the Lenovo RackSwitch G8264.

Note: Please review the Change History documentation posted with the Switch Firmware to check if any of these issues have been fixed in the latest release.

ACLs

ACL logging does not block traffic sent to the CPU. Use Management ACLs if you need to filter or block inbound traffic. (ID: XB211816)

BGP

Maximum number of route maps that can be added to a BGP peer is 16 (8 route-maps for incoming traffic and 8 for outgoing traffic). (ID: 46448)

BGP Debugging

While enabling or disabling BGP debug for a particular peer/IP address, the logging behavior may not be as expected. Following is a workaround: (ID: 59104)

To enable BGP debug for a particular peer:

1. Enable BGP debug for all the peers.
2. Disable BGP debug for all the peers.
3. Enable BGP debug for a particular peer.

To disable BGP debug for a particular peer:

1. Enable BGP debug for all the peers.
2. Disable BGP debug for all the peers.
3. Enable BGP debug for all the peers except the one for which you want it disabled.

Debug

Enterprise NOS debug commands are for advanced users. Use the debug commands with caution as they can disrupt the operation of the switch under high load conditions. This could be dangerous in mission-critical production environments. Before debugging, check the MP utilization to verify there is sufficient overhead available for the debug functionality. When debug is running under high load conditions, the CLI prompt may appear unresponsive. In most cases, control can be returned by issuing a `no debug <function>` command.

EVB

- When a VM cannot be associated, the console may be flooded with syslog messages stating that the validation has failed. (ID: XB191291)
- Due to a hardware limitation, traffic received by a VM may not conform to the RxRate (receive rate) that you have configured. (ID: 55600)

FCoE

- The FCoE connection between the server and the FCF will be retained even if you disable CEE/FIP/vNIC on the switch. To avoid this scenario, either reboot the switch, or disable and re-enable the ports connected to the sever and the FCF after you disable CEE/FIP/vNIC. (ID:41915)
- By default the "VLAN Name" and "Port and Protocol ID" LLDP TLVs are disabled on a port. These two TLVs are added to the LLDP PDU for each VLAN that is configured in a port. This may cause the length of LLD PDU to exceed the Ethernet packet size if there are nearly 40 or more VLANs configured on a port, or if the VLAN names are too long. There is a possibility that the DCBX TLVs may not be added to the LLDP TLV due to the length. Because of this the FCoE connection will not form on that port. It is recommended to avoid enabling the "VLAN Name" and "Port and Protocol ID" TLV if you have high number of VLANs configured and FCoE is enabled on that port. (ID: 42446)
- The FIP Snooping option for automatic VLAN creation is not recommended for use with the Emulex Virtual Fabric Adapter. Disable automatic VLAN creation when connecting the switch to an Emulex Virtual Fabric Adapter. (ID: 51529)
- When using DCBX to synchronize Priority-base Flow Control (PFC) with a peer (using the PFC TLV option), PFC operation on the switch port will not be disabled as expected when PFC is not available on the peer. To resolve this, manually disable PFC on ports connected to devices that do not participate in PFC. (ID: 62114)
- Disruption to FCoE connections and FCoE traffic may be expected when changing the LACP mode. It is recommended that the administrator halt FCoE traffic before changing any switch configuration. (ID: 67044)
- It is recommended to use the FIP snooping automatic VLAN creation option in FCOE environments, in addition to configuring VLANs manually. The auto-VLAN feature should be disabled only if no additional FCF or ENode ports will be automatically added to the FCOE VLAN. Otherwise, some FCF or ENode ports might not be automatically added to the FCOE VLAN, even if the auto-VLAN feature is later enabled, requiring them to be added manually.
- The FIPS auto-VLAN feature is disabled by default. (ID: XB258382)
- FCoE connections may become temporarily unstable whenever a change occurs to the VLAG virtual port. (ID: XB263734)
- In an event in which connections for multiple ports becomes unstable, FCoE traffic may pause or be dropped as FCoE FDB entries are cleared and reinstated. (ID: XB275415)

FDB

From IBM Networking OS 7.11 onwards, MAC address information is no longer learned by control packets such as LACPDU. This behavior is as expected. (ID: XB253517)

IKEv2

IKEv2 cannot be configured on management ports. Configure IKEv2 only on data ports. (ID: 57427)

IPsec

When configuring IPsec to operate between IBM switches, keys of various lengths are supported. However, when using IPsec to connect with non-IBM devices, the manual policy session keys must be of the following fixed lengths:

- For the AH key:
 - SHA1 = 20 bytes
 - MD5 = 16 bytes
- For the ESP auth key:
 - SHA1 = 20 bytes
 - MD5 = 16 bytes
- For the ESP cipher key:
 - 3DES = 24 bytes
 - AES-cbc = 24 bytes
 - DES = 8 bytes

ISCLI

If a port needs to be a member of more than 500 VLANs, we recommend that you first shut down the port and then add the port as a member of the VLANs. (ID: 70739)

LACP

Under heavy switch load conditions, LACP links may flap when configured with short timeout mode. To stabilize LACP under heavy load, it is recommended to use the long timeout mode instead. (ID: 66173)

On-Box Scripting

- The keys in the returned dictionary from `get_lldpReceive` need to be updated as follows:
 - Old Keys:

```
'index 1': {'Alias': 1,
            'Bad Frame': 'false',
            'DMAC': 'NB',
            'Parameters rxTTL': 0,
            'RCV Frame': 'false',
            'RXI nfo Ageout': 'false',
            'Recei ve State': 'LLDP_WAI T_PORT_OPERATI ONAL',
            'Remote Changed': 'false',
            'SNMP Noti fy': 'false',
            'Too Many Nei ghor': 'false',
            'TooManyNei ghorsTi mer': 0,
```

- New Keys:

```
'index 1': {'Al ias': 1,
            'BadFrame': 'fal se',
            'DMAC': 'NnTB',
            'RCVFrame': 'fal se',
            'RXI nfoAgeout': 'fal se',
            'Recei veState': 'LLDP_WAI T_PORT_OPERATI ONAL',
            'RemoteChanged': 'fal se',
            'SNMPNoti fy': 'fal se',
            'TooManyNei ghor': 'fal se',
            'TooManyNei ghorsTi mer': 0,
            'port': 1,
            'rxTTL': 0},
```

(ID: XB258010)

- The document strings for `ibmpylib.set_var()` and `del_var()` do not automatically update when you add a new function. (ID: XB264941)
- The storage space available for user scripts is 850K. (ID: XB265456)

OpenFlow

- Static FDB flows are stored as ACL flows. (ID: XB262456)
- Using a `port.mod` message to change the `OFPPC_NO_FLOOD` bit does not reprogram the already installed flows. (ID: XB259385)
- When you configure a port to use OpenFlow, spanning tree protocol is automatically disabled on that port. (ID: XB266710)

OSPF

- Some changes to OSPF configuration (such as creating a new area or changing an area's type) may result in OSPF state reconvergence. (ID: 46445, 48483)

- OSPFv3 over IPsec:
 - This combination can only be configured only on a per-interface basis.
 - The current implementation for OSPFv3 allows the use of only one protocol (AH or ESP) at any given time. AH and ESP cannot be applied together.
 - IPsec does not support OSPFv3 virtual links. (ID: 48914)

Ports and Transceivers

In stacking mode, two ports of different link speeds can exist in the same portchannel. This may lead to loss of traffic. (ID: XB278986)

Private VLANs

- The sequence in which a private VLAN is configured is not the same as displayed in the output of the show running-config command. Hence, if you copy and paste the private VLAN configuration from the output of the preceding command, the private VLAN configuration will be lost. (ID: 67169)
- Traffic with secondary VLAN ID is not forwarded to promiscuous ports. (ID: 70980)

QoS

When the following command is issued command is issued, "Dropped Packets" and "Dropped Bytes" counters will be displayed as '0' due to hardware limitations:

```
RS8264(config)# show interface port <swunit:port_num>
egress-mcast-queue-counters
```

For example:

```
RS8264(config)# show interface port 1:24 egress-mcast-queue-counters

Multicast QoS statistics for port 1:24:
QoS Queue 8:
  Tx Packets:                377
  Dropped Packets:           0
  Tx Bytes:                  50883
  Dropped Bytes:             0
```

(ID: XB233503)

QSFP+

- The QSFP+ ports do not auto-negotiate. The desired speed must be configured to match on both ends of the connection, and the switch reset for changes to take effect. (ID: 46340)
- After you upgrade switch software and reset the switch, you must configure the QSFP+ port mode. Use the following command:

```
RS8264(config)# boot qsfp-40gports <1, 5, 9, 13>
```

(ID: 46858)

- When changing a QSFP port from 10G mode to 40G mode, a port error will occur if any previously configured 10G port settings do not apply to the new 40G state, preventing further configuration of the port. The administrator must manually clear the 10G port settings that do not apply to 40G prior to changing modes. (ID: 62576)

Routed Ports

Do not use IBM N/OS CLI, SNMP, or BBI to configure routed ports, or to configure any other feature if a routed port is already configured on the switch.

If a routed port is configured on the switch, the configuration, apply, and save commands are not displayed in IBM N/OS CLI or BBI; in SNMP, you may be able to enter the configuration commands, but you will not be able to save the configuration. (ID: 57983)

sFlow

In some cases, sFlow configured with the minimum polling and sampling rate could cause the switch to get into a hang state with no traffic passing after about 7 days of operations with large volumes of traffic. Please contact Customer Support or the System Engineer before enabling sFlow. (ID: 57045)

SLP

Abbreviated IPv6 addresses are not supported in Service Location Protocol (SLP) strings. All IPv6 addresses used in SLP request strings must be extended. For example, the SLP request:

```
slptool findattrs
service:io-device.Lenovo:management-module://2001::1
```

will not work. Instead, you must use the extended form of the IPv6 address:

```
slptool findattrs
service:io-device.Lenovo:management-module://2001:0000:0000:
0000:0000:0001
```

SNMP

- Port information displayed in MIBs related to port-based VLANs does not distinguish between a regular port or a trunk port. Use the `show mac-address-table static` command to view details on regular ports and trunk ports. (ID: 57194)
- If you delete multiple VLANs using SNMP, you may see an error if the SNMP packet size exceeds 1800 bytes. (ID: XB228120)

Spanning Tree

- When using LACP with PVRST, it is not recommended to configure the switch as the STP root bridge. When doing so, traffic can be discarded for up to 30 seconds on affected LACP ports while initial STP path states are being resolved (discarding, learning, forwarding). (ID: 63315)

- After changing from MSTP to a different STP mode, some MSTP commands will still appear in the configuration file. The non-applicable MSTP commands do not affect switch operation and can be ignored. (ID: 64388)

Stacking

When stacking is enabled, the switch may continue to learn MAC addresses from ports or trunks even though they are in a blocking state. The unexpected MAC addresses represent control packets, not endpoint devices, and do not impact switch performance. (ID:61996)

Virtual Aggregation Link Groups (vLAG)

- The following features are not supported on ports participating in VLAGs:
 - FCoE
 - Hotlinks
 - IGMP relay
 - Private VLANs
 - vNICs
 - UDLD
- In a multi-layer VLAG topology, the VLAG ports may be disabled in CIST if you change the STP mode on the secondary switch to MSTP. (ID: 58696)

VLANs

- When a VLAN appears in the VLAN range for a port in a configuration dump, this does not guarantee that the port is actually a member of that VLAN. The actual port to VLAN mapping can be displayed by using the show vlan command. (ID: XB267491)
- When VLAG ports are removed from a VLAG VLAN, the port list still contains both the VLAG ports just removed and the ISL ports. (ID:XB278681)

VMready

- VMs belonging to different ESX servers cannot ping each other across different VM groups. Because the VM groups belong to different VLANs, this is appropriate and expected behavior.
- On switch ports on which VMs are learned, the switch does not learn the MAC address of the destination host unless the host sends some network traffic. Therefore the switch might not forward packets to the destination host (for instance, when using ping).
 - If you are not using VMready in a VM environment, disable VMready (no virt enable).
 - If you are using VMready, periodically send traffic from the host (for example, ping), so that the host's MAC address is always present in the Forwarding Database (MAC Address Table). (ID: 44946)

- vNIC egress bandwidth control is not strictly enforced on the switch for packets larger than 900 bytes, resulting in greater egress bandwidth from the switch to the server than is configured. However, ingress bandwidth control (from the server to the switch) is strictly enforced. (ID: 50950)

vNICs

- When using vNICs with FCoE, the FIP Snooping option for automatic VLAN creation is not recommended for use with the Emulex Virtual Fabric Adapter. Disable automatic VLAN creation when connecting the switch to an Emulex Virtual Fabric Adapter. (ID: 51529)
- vNIC egress bandwidth control is not strictly enforced on the switch for packets larger than 900 bytes, resulting in greater egress bandwidth from the switch to the server than is configured. However, ingress bandwidth control (from the server to the switch) is strictly enforced. (ID: 50950)