

IBM System Networking RackSwitch™ G8264-T



# Release Notes

For Networking OS 7.9

**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 *IBM Documentation* CD and the *Warranty Information* document that comes with the product.

**Second Edition (June 2014)**

**© Copyright IBM Corporation 2014**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

## Release Notes

This release supplement provide the latest information regarding IBM Networking OS 7.9 for the RackSwitch G8264-T (referred to as G8264-T throughout this document).

This supplement modifies and extends the following IBM N/OS documentation for use with N/OS 7.9:

- *IBM Networking OS 7.9 Application Guide*
- *IBM Networking OS 7.9 ISCLI Reference*
- *RackSwitch G8264-T Installation Guide*

The publications listed above are available from the IBM support website:

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

Please keep these release notes with your product manuals.


---

## Hardware Support

The G8264-T contains forty-eight 10G/1GbaseT copper ports and four 40GbE QSFP+ ports. The four 40GbE QSFP+ ports can also work as sixteen 10GbE 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.

Figure 1. RackSwitch G8264T Front Panel



## Transceivers

The following transceivers and DACs are available:

*Table 1. RackSwitch G8264-T Ordering Information*

<b>Part number</b>	<b>Description</b>
<b>Transceivers</b>	
BN-CKM-QS-SR	QSFP+ 40GBASE-SR4 Optical Fiber Transceiver
<b>Direct Attach Cables (DACs)</b>	
BN-SP-CBL-1M	SFP+ 10Gbps 1 meter DAC
BN-SP-CBL-3M	SFP+ 10Gbps 3 meter DAC
BN-SP-CBL-5M	SFP+ 10Gbps 5 meter DAC
BN-QS-QS-CBL-1M	QSFP+ 40Gbps 1 meter DAC
BN-QS-QS-CBL-3M	QSFP+ 40Gbps 3 meter DAC
BN-QS-QS-CBL-5M	QSFP+ 40Gbps 5 meter DAC
BN-QS-SP-CBL-1M	QSFP+ to four SFP+ 1 meter breakout DAC
BN-QS-SP-CBL-3M	QSFP+ to four SFP+ 3 meter breakout DAC
BN-QS-SP-CBL-5M	QSFP+ to four SFP+ 5 meter breakout DAC

The G8264-T accepts any SFP+ Direct Attach Cable that complies to the MSA specification.

---

## Updating the Switch Software Image

The switch software image is the executable code running on the G8264-T. 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-T, go to the following website:

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

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

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

- Load a new software image and boot image onto an 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 using the CLI, ISCLI, or BBI, see [“Loading New Software to Your Switch” on page 10](#).



### CAUTION:

**Although the typical upgrade process is all that is necessary in most cases, upgrading from (or reverting to) some versions of N/OS or BLADEOS requires special steps prior to or after the software installation process. Please be sure to follow all applicable instructions in the following sections to ensure that your switch continues to operate as expected after installing new software.**

## Special Software Update Issues

When updating to N/OS 7.9, the following special conditions may apply, depending on the version of software currently installed on your switch. These conditions are cumulative: If updating from version 2.0 (for example), follow the recommendations in order, beginning with those for 2.0, and then continue with all that apply, such as for “3.0 and prior,” “4.0 and prior,” and so on.

### Updating from BLADEOS 1.x

Before you upgrade from software version 1.x, it is recommended that you save the previous configuration block on a TFTP server, and set the configuration block to factory default, as follows:

```
RS8264-T(confi g)# boot confi gurati on-bl ock factory
```

After updating:

- The range value for dest-lookup-threshold, broadcast-threshold and multicast-threshold parameters are different compared to release 1.x. On release 1.1 the range is <1-33554431> and in release 5.x and later, the range is <0-2097151>.

During upgrade, the dest-lookup-threshold, broadcast-threshold and multicast-threshold parameters are unchanged if the values are within the range available on release 7.9. But any parameters that fall outside the range available on release 7.9 are set to the maximum value of 2097151. (ID: 35936)

- The range value for NTP interval is different compared to 1.x software. On release 1.1 the range is <1-10080> and on release 5.x and later the range is <5-44640>.

If the NTP interval value is lower than 5, then after software upgrade the NTP interval is set to the minimum value of 5. (ID: 36099 | 36500)

- The default values and range values for IGMP report timeout parameter are different for release 5.x and later as compared to release 1.1:
  - On release 1.1 the range for IGMP report timeout is <130-1225> seconds with a default of 260 seconds.
  - On release 5.x and later, the range is <1-255> minutes with a default of 10 minutes.

During upgrade, the value of IGMP report timeout is set to the new default value (10). The value does not appear in the running configuration output. (ID: 36131 | 35578)

- On release 1.1, the default setting for IP routing is `di sabl ed`, and on release 5.x and later the default setting for IP routing is `enabl ed`. During software upgrade, IP routing is set to the new default (`enabl ed`). (ID: 36217)
- During software upgrade from release 1.1 to release 5.x or later, Uplink Failure Detection (UFD) is converted to Layer 2 Failover. The UFD LtM and LtD options are converted to Failover trigger 1 MMON monitor and MMON control. If the LtD is configured using a combination of ports, trunks, and LACP *admin keys*, then the UFD configuration is cleared during conversion upgrade, and log message is displayed. (ID: 36220)
- Release 5.x and later uses a different command to set the SSH port to its default value. After upgrade, use the following command (ID: 36382): **default t ssh port**
- On release 1.1, the default setting for Hotlinks BPDU is `enabl ed`, and on release 5.x and later the default setting is `di sabl ed`. During upgrade, the Hotlinks BPDU command is set to `di sabl ed`. (ID: 36385 | 36385)
- On release 1.1 the maximum number of characters allowed for Hotlinks trigger name is 33, and on release 5.x and later the maximum is 32 characters. During upgrade, only the first 32 characters of Hotlinks Trigger names are saved. A warning message is displayed. (ID: 36471)
- On release 5.x and later the `hal f dupl ex` option is not available for static trunks or LACP trunks. During upgrade from release 1.1 or prior, the `dupl ex hal f` setting is cleared from the trunk configuration. All other parameters are maintained. (ID: 36486)
- On release 5.x and later, BPDU Guard can be configured for ports only (not trunks). During upgrade from release 1.x, the BPDU Guard setting is applied to each member port in a trunk. (ID: 36512)

- DHCP settings are different compared to release 1.x:  
On release 1.1 DHCP can be configured for any IP interface. For interface 1, DHCP is enabled by default. For the rest of the IP interfaces, DHCP is disabled by default.  
On release 5.x and later, DHCP is enabled/disabled globally, but only IP interface 1 can get an IP address dynamically via DHCP. The default value for DHCP in 5.x and later is enabled.  
In some cases, the entire 1.x configuration can be lost during the upgrade (for example, if some IP interfaces have DHCP enabled, but IP interface 1 has DHCP disabled). (ID: 36536)
- The commands saved in release 1.1 are displayed in logging messages after upgrade, but some commands might not be displayed under the correct configuration menu. (ID: 36705)
- The running configuration display for the following command changes on release 5.x and later:  
`no snmp-server link-trap port x enable`  
On release 1.1 ports in the range 49-52 are displayed using port numbers.  
On release 5.x and later, ports in the range 49-52 are displayed using port alias. (ID: 36833)
- The default value for the `access https` command is different compared to release 1.x. On release 1.1 the default setting is enabled, and on release 5.x and later, the setting is disabled. During upgrade, `access https` is set to disabled. (ID: 36834)
- The running configuration output for the following command is different compared to release 1.x: **`interface port dot1x mode force-unauthorized`**  
On release 5.x and later, the display commands are as follows (ID 36980):  
`interface port 11  
dot1x mode force-unauth`
- On release 5.x and later, there is no command to disable/enable FDB learning for trunks, only for switch ports. (ID: 37102)

## Reverting to BLADEOS 1.x or Prior

If you revert from software image 5.x or later to software image 1.x, the configuration file is cleared and reset to the factory default.

- The default for the Layer-3 hash is different compared to release 5.x and prior. In release 5.x, the source IP address (SIP) was the default used to generate the Layer-3 hash. In release 6.3 and above, source and destination IP addresses (SIP-DIP) are used as the default. (ID: 39733)
- Some time zones are different compared to release 5.x and prior. After upgrading to release 6.3 or above, it is recommended that the administrator review the configured time zone and make any appropriate changes. (ID:29778)

## Updating from BLADEOS 5.1 or Prior

When you upgrade the G8264-T from release 5.1 or prior, the configuration block is converted to match the new software.

Most configuration data is automatically converted to equivalent commands and ranges. However, some older configuration data has no equivalent on release 5.2 or later, and is not converted. For example, ACL commands are different prior to release 5.2 and are not converted. Log messages list commands that were not converted during the upgrade. You must manually configure those features that were not converted during the upgrade.

## Updating from BLADEOS 6.4 or Prior

After updating:

- The default for STP/PVST Protection mode is different compared to release 6.4 and prior. In release 6.6, STP/PVST Protection is disabled by default. After upgrading, review the STP settings and make any appropriate changes.
- The default for static route health check is different compared to release 6.4 and prior. In release 6.6, static route health check is disabled by default. After upgrading, review the static route health check settings and make any appropriate changes.
- The legacy FDB update rate has been deprecated in favor of independent hotlinks FDB updates in all switch configuration interfaces.

Interface	Old Commands	New Commands
Menu CLI	/cfg/l 2/update <x>	/cfg/l 2/hotlink/sndrate <x>
ISCLI	spanning-tree uplinkfast max-update-rate <x>	hotlinks fdb-update-rate <x>
BBI	Configure   Layer 2   Uplink Fast   Update Rate  Dashboard   Layer 2   Uplink Fast   STP Uplink Fast Rate	Configure   Layer 2   Hot Links   FDB update rate  Dashboard   Layer 2   Hot Links   FDB update rate

These changes are also reflected in the SNMP MIB.

After upgrading, review the hotlinks FDB settings and make any appropriate changes

- The CLI BGPTOECMP option has been deprecated.

## Updating from BLADEOS 6.6 or Prior

After updating:

- The default mode for Spanning Tree is different compared to prior releases. The default mode is now PVRST. After upgrading, it is recommended that the administrator review the STP settings and make any appropriate changes.



## Updating from IBM Networking 6.8 or Prior



### CAUTION:

If the current software version on your switch is 6.8 or prior, first upgrade the switch software image to 6.9 and reset the switch. Then load the 7.2 boot image and software image.

## Updating from IBM Networking OS 6.9 or Prior



### CAUTION:

When you upgrade the switch software image, you must load the new boot image and the new software image before you reset the switch.

After updating:

- The default settings of SNMP community strings have changed. Check the new settings and reconfigure as appropriate.

## Updating from IBM Networking OS 7.2 or Prior

After updating:

- The default time zone setting is different compared to release 7.2 and prior. In the prior releases, a default setting of US Pacific Time was used. In release 7.4 and above, no default is assumed. For switches that use the default US Pacific Time setting, after upgrading to release 7.4 or above it is recommended that the administrator review the configured time zone and make any appropriate changes. (ID: 60469)

## Updating VLAG Switches with IBM Networking OS 7.x

Following are the steps for updating the software image and boot image for switches configured with VLAG:

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

```
RS8264-T(config)# copy running-configuration startup-configuration
```

2. Use TFTP or FTP to copy the new OS image and boot image onto both vLAG switches.
3. Shut down 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
5. Once Switch 1 has rebooted, Switch 1 will take the secondary role.
6. Shut down all ports except the ISL ports and the health check port on Switch 2.  
**Note:** Do not save this configuration.
7. Reload Switch 2, Switch 1 will reassume the vLAG primary switch role.

8. Once Switch 2 has reloaded, make sure Switch 1 has transitioned to vLAG primary and Switch 2 has transitioned to secondary.
9. Verify all the vLAG clients have converged using the following command:

```
RS8264-T(config)# show vlag information
```

## Loading New Software to Your Switch

The G8264-T 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.



### CAUTION:

**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 Upgrade” on page 16](#)).**

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

- The image and boot software loaded on an 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 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 one of the following procedures to download the new software to your switch. You can use the N/OS CLI, the ISCLI, or the BBI to download and activate new software.

## Loading Software via the ISCLI

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

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

2. Enter the hostname or IP address of the 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 FTP or TFTP directory (for example, `tftpboot`).

4. If required by the FTP or TFTP server, enter the appropriate username and password.

- 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 (i mage1 or i mage2) you want to run in switch memory for the next reboot:

```
Router# configure terminal
Router(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:

```
Router(config)# reload
```

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

## Loading Software via the N/OS CLI

- Enter the following Boot Options command:

```
>> # /boot/gtimg
```

- Enter the name of the switch software to be replaced:

```
Enter name of switch software image to be replaced
["image1"/"image2"/"boot"]: <image>
```

- Enter the hostname or IP address of the FTP or TFTP server.

```
Enter hostname or IP address of FTP/TFTP server: <hostname or IP address>
```

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

```
Enter name of file on FTP/TFTP server: <filename>
```

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

- Enter your username for the server, if applicable.

```
Enter username for FTP server or hit return for
TFTP server: {<username>/<Enter>}
```

If entering an FTP server username, you will also be prompted for the password. The system then prompts you to confirm your request. Once confirmed, the software will load into the switch.

6. If software is loaded into a different image than the one most recently booted, the system will prompt you whether you wish to run the new image at next boot. Otherwise, you can enter the following command at the `Boot Options#` prompt:

```
Boot Options# image
```

The system then informs you of which software image (`image1` or `image2`) is currently set to be loaded at the next reset, and prompts you to enter a new choice:

```
Currently set to use switch software "image1" on next reset.  
Specify new image to use on next reset ["image1"/"image2"]:
```

Specify the image that contains the newly loaded software.

7. Reboot the switch to run the new software:

```
Boot Options# reset
```

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

## Loading Software via BBI

You can use the Browser-Based Interface to load software onto the G8264-T. The software image to load can reside in one of the following locations:

- FTP server
- TFTP server
- Local computer

After you log onto the BBI, perform the following steps to load a software image:

1. Click the Configure context tab in the toolbar.
2. In the Navigation Window, select System > Config/Image Control.  
The Switch Image and Configuration Management page appears.
3. If you are loading software from your computer (HTTP client), skip this step and go to the next. Otherwise, if you are loading software from an FTP or TFTP server, enter the server's information in the FTP/TFTP Settings section.
4. In the Image Settings section, select the image version you want to replace (Image for Transfer).
  - If you are loading software from an FTP or TFTP server, enter the file name and click **Get Image**.
  - If you are loading software from your computer, click **Browse**.  
In the File Upload Dialog, select the file and click **OK**. Then click **Download via Browser**.

Once the image has loaded, the page refreshes to show the new software.

---

## New and Updated Features

N/OS 7.9 for RackSwitch G8264-T (G8264-T) has been updated to include several new features, summarized in the following sections. For more detailed information about configuring G8264-T features and capabilities, refer to the complete N/OS 7.9 documentation as listed on [page 3](#).

### Auto VLAN Tagging on Trunk Ports

This enhancement facilitates the process of adding trunk ports in VLANs by automatically adding them to all VLANs from their allowed ranges. By default, all VLANs are allowed on each port. When a port is configured as trunk port, it is automatically added to all VLANs from its allowed range. Also, when a new VLAN is created, all trunk ports which have that VLAN in their allowed ranges are automatically added to it.

### BGP Community Lite

BGP community strings can be advertised in updates to neighbors. You can configure a switch to attach a community string to the route updates it sends to peers, and the switch will not make any routing changes or alterations to the community string when receiving updates with a community string attached.

### Display BGP Routes

There is an option to display BGP advertised routes that have been advertised to a specific neighbor.

### ESN to SNMP

This feature enables SNMP access to the Electronic Serial Number of the switch.

### Full Private VLAN

This release supports Private VLAN configurations as described in RFC 5517.

### IBM N/OS Menu-Based Interface Removal

The IBM N/OS menu-based CLI is not supported as of this release.

All switches will boot up with the Industry-Standard CLI (ISCLI). The existing NOS CLI configuration can still be recognized and correctly converted to provide smooth migration for customers who have NOS CLI configuration.

### IPSec over Virtual Links

OSPFv3 over IPSec on Virtual Links is needed to complete NIST IPSec certification for OSPFv3 traffic. IPSec is needed to secure IPv6 traffic. The feature will use IPv6 Authentication Header (AH) to provide authentication and IPv6 Encapsulating Security Payload (ESP) to provide authentication and confidentiality to virtual link packets.

## IPv6 Counter Enhancement

This release adds CLI and corresponding SNMP MIB objects for IPv6 counters. The feature provides support for the IPv6 neighbor cache table and statistics, such as:

- current number of installed entries
- maximum number of entries supported by the router
- high water of the IPv6 neighbor cache table
- clearing statistics

## Layer 3 ARP Table Full

When the Layer 3 ARP table is full, the switch will generate a new trap message in addition to the existing syslog message.

## Link Aggregation Control Protocol (LACP) Individual Mode

When this feature is enabled on an LACP portchannel, if a member port of the portchannel does not receive any LACPDU over a period of time, it will be treated as a normal port that may forward data traffic according to its STP state.

## OpenFlow 1.3.1 Support

Added features from Openflow Switch Specification Version 1.3.1.

### Openflow 1.3.1 Group Support

Support for Openflow groups, in accordance with Openflow 1.3.1, has been added. Actions associated with flow entries can direct packets to a group.

### Openflow Support for Static LAG over Edge Ports

This feature will allow user to configure Openflow static LAG port as edge port on the switch. You can configure multiple Openflow LAG ports and physical ports as edge ports as required.

### sFlow Support in Openflow Ports

This release adds sampling support for packets received on Openflow ports configured for this feature. An sFlow server should be configured to be reachable via non-Openflow data port or management port for this functionality to work.

## Password Fix-Up Mode

Password fix-up mode enables admin user account recovery if administrator access is lost. To use this mode, you must contact Support to obtain access to this feature. You can also disable password fix-up functionality to let the administrator of the switch decide whether to enable password fix-up mode to cover security concerns.

## Service Location Protocol

Service location Protocol (SLP) provides a dynamic configuration mechanism for applications in local area networks. Applications are modeled as clients that need to find servers attached to any of the available networks within an enterprise.

## Spanning Tree Protocol (STP) Range Enhancement

Existing Spanning Tree Protocol (STP) commands now support configuration of a range of STP groups.

## SNMP

The following features have been added to SNMP support.

### SNMP ACL

This feature is an enhancement to add access control for SNMP requests.

### SNMP Trap Host

This feature implements the SNMP interface for getting and setting SNMP host configuration for traps.

## Use SSH Public Keys for up to 20 Local Switch Users

The feature allows users to login to a switch via SSH using public key authentication instead of password authentication. When SSH is enabled the switch supports both password and public key authentication. The switch now supports up to 20 SSH public key users.

## vLAG MSTP Enhancement

This enhancement removes STP configuration restrictions, such as changing the MSTP instance and VLAN associations, that were enforced in previous releases when vLAG and MSTP were both enabled. The vLAG interswitch link ports are no longer error-disabled when there is an MSTP region mismatch between the vLAG switches. Instead, a recurring warning message appears during the duration of the configuration mismatch.

## vLAG PIM with Multicast Sources

Protocol Independent Multicast (PIM) is a routing protocol that routes multicast traffic from multicast sources to receivers. This enhancement enables support for vLAG PIM with multicast sources connected in the Layer 2 domain behind the vLAG ports by defining the vLAG PIM protocol behavior and traffic routing across different multicast source and receivers.

Multicast sources and receivers can be connected anywhere in the vLAG PIM environment. You can now use vLAG PIM in a multi-tier tenant environment with Layer 2 vLAG on the bottom tier and Layer 3 vLAG on the top tier.

---

## Supplemental Information

This section provides additional information about configuring and operating the G8264-T and N/OS.

### 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
1 - Change booting image
2 - Change configuration block
3 - Xmodem download
4 - Exit

Please choose your menu option: 1
Current boot image is 1. Enter image to boot: 1 or 2: 2
Booting from image 2
```

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

- To change the booting image, press 1 and follow the screen prompts.
- To change the configuration block, press 2, and follow the screen prompts.
- To perform an Xmodem download, press 3 and follow the screen prompts.
- To exit the Boot Management menu, press 4. The booting process continues.

### Recovering from a Failed Upgrade

Use the following procedure to recover from a failed software 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, 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 **3** for **Xmodem download**. When you see the following message, change the Serial Port characteristics to 115200 bps:

```
## Switch baudrate to 115200 bps and press ENTER ...
```



5. 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.
6. Select the Boot Image to download. The XModem initiates the file transfer. When the download is complete, a message similar to the following is displayed:

```

yzModem - CRC mode, 62494(SOH)/0(STX)/0(CAN) packets, 6 retries
Extracting images ... Do *NOT* power cycle the switch.
**** VMLI NUX ****
Un-Protected 10 sectors
Erasing Flash..... done
Writing to Flash..... done
Protected 10 sectors
**** RAMDISK ****
Un-Protected 44 sectors
Erasing Flash..... done
Writing to Flash..... done
Protected 44 sectors
**** BOOT CODE ****
Un-Protected 8 sectors
Erasing Flash..... done
Writing to Flash..... done
Protected 8 sectors

```

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

```
## Switch baudrate to 9600 bps and press ESC ...
```

8. Press the Escape key (**<Esc>**) to re-display the Boot Management menu.
9. Select **3** to start a new **XModem Download**. When you see the following message, change the Serial Port characteristics to 115200 bps:

```
## Switch baudrate to 115200 bps and press ENTER ...
```

10. Press **<Enter>** to continue the download.

11. Select the OS Image to download. The XModem initiates the file transfer. When the download is complete, a message similar to the following is displayed:

```
yzModem - CRC mode, 27186(SOH)/0(STX)/0(CAN) packets, 6 retries  
  
Extracting images ... Do *NOT* power cycle the switch.  
  
**** Switch OS ****  
  
Please choose the Switch OS Image to upgrade [1|2|n] :
```

12. Select the image number to load the new image (1 or 2). It is recommended that you select 1. A message similar to the following is displayed:

```
Switch OS Image 1 ...  
  
Un-Protected 27 sectors  
  
Erasing Flash..... done  
  
Writing to Flash..... done  
  
Protected 27 sectors
```

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

```
## Switch baudrate to 9600 bps and press ESC ...
```

14. Press the Escape key (<Esc>) to re-display the Boot Management menu.
15. Select **4** to exit and boot the new image.

## 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-T (config)# no vlag adminkey <key> enable (or)  
RS8264-T (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.

---

## Known Issues

This section describes known issues for N/OS 7.9 on the RackSwitch G8264-T

### 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. (ID: 46448)

### BGP Debug

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

- IBM N/OS 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-based 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.

## FCOE

FCoE connections flap whenever a change occurs to the vLAG virtual port. (ID: XB263734)

## FIPS

The FIPS auto-VLAN feature is "Disable" by default. (ID: XB258382)

In an event in which multiple ports on a switch are flapped, FCoE traffic may drop or pause due to FCoE FDB entries being flushed and reinstalled. (ID: XB275415)

## Forwarding Database (FDB)

From IBM Networking OS 7.9 onwards, MAC address information is no longer learned by control packets such as LACPDUs. 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 shutdown the port and then add the port as a member of the VLANs. (ID: 70739)

## LACP

### On-Box Scripting

- You need to update the keys in the returned dictionary from `get_LLDPReceive` as follows: (ID: XB258010)

Old Keys:

```
'index 1': {'Alias': 1,
            'Bad Frame': 'false',
            'DMAC': 'NB',
            'Parameters rxTTL': 0,
            'RCV Frame': 'false',
            'RXInfo Ageout': 'false',
            'Receive State': 'LLDP_WAIT_PORT_OPERATIONAL',
            'Remote Changed': 'false',
            'SNMP Notify': 'false',
            'Too Many Neighbor': 'false',
            'TooManyNeighborsTimer': 0,
```

New Keys:

```
'index 1': {'Alias': 1,
            'BadFrame': 'false',
            'DMAC': 'NnTB',
            'RCVFrame': 'false',
            'RXInfoAgeout': 'false',
            'ReceiveState': 'LLDP_WAIT_PORT_OPERATIONAL',
            'RemoteChanged': 'false',
            'SNMPNotify': 'false',
            'TooManyNeighbor': 'false',
            'TooManyNeighborsTimer': 0,
            'port': 1,
            'rxTTL': 0},
```

- The document string 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

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

## Precision Time Protocol

When using the PTP Transparent Clock on the switch, there may be variations in the residence time for PTP packets traversing the switch. The corrections stored in the Follow-Up/Delay-Response packets will correctly take into account the residence time. However, other PTP devices that receive event packets that pass through the switch (thus obtaining a residence time correction from the switch) must be configured to be resilient to residence time variations. For example, some PTP devices provide stiffness filters which help the device compute an average of the path delay. (ID: 61657)

## Private VLANs

- The sequence in which a private VLAN is configured is not the same as displayed in the output of the `RS8264-T(config)# show running-config` command. Hence, if you copy and paste the private VLAN configuration from the output of the above 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, "Dropped Packets" and "Dropped Bytes" counters will be displayed as '0' due to hardware limitations: (ID: XB233503)

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

## 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 (ID: 46858):  
**boot qsfp-40gports <1, 5, 9, 13>**
- 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)



## 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)

## SNMP

- Port information displayed in MIBs related to port-based VLANs does not distinguish between a regular port or a trunk port. Use the `RS8264-T(config)# 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)

## UFP

The command `show ufp information port <x>` does not show disabled vPorts. (ID: XB267210)

## Virtual Link Aggregation Groups

- 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 vl an` 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). (ID: 44946)
  - 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).

## 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)