

IBM System Networking RackSwitch™ G8264



Release Notes

For Networking OS 7.8

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.

First Edition (December 2013)

© Copyright IBM Corporation 2013

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 7.8 for the RackSwitch G8264 (referred to as G8264 throughout this document).

This supplement modifies and extends the following IBM N/OS documentation for use with N/OS 7.8:

- *IBM Networking OS 7.8 Application Guide*
- *IBM Networking OS 7.8 Command Reference*
- *IBM Networking OS 7.8 ISCLI Reference*
- *IBM Networking OS 7.8 BBI Quick Guide*
- *RackSwitch G8264 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 contains forty-eight 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.

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.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 6](#).



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 7.8, 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 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/l2/update <x>	/cfg/l2/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. Shut down all the ports on both VLAG Peers.
2. Upgrade VLAG Peer 1 to 7.x (both OS and Boot Image).
3. Upgrade VLAG Peer 2 to 7.x (both OS and Boot Image).
Note: Both VLAG peers must be updated with the same version.
4. Save the configuration using the following command:

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

5. Reload VLAG Peer 1.
6. Reload VLAG Peer 2.
7. Turn on the required ports.

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.



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 17](#)).

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 N/OS CLI

1. Enter the following Boot Options command:

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

2. Enter the name of the switch software to be replaced:

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

3. 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>
```

4. 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).

5. 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 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.
5. The switch will prompt you to confirm your request.
Once confirmed, the software will begin loading into the switch.

6. 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:

```
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.
```

7. 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 BBI

You can use the Browser-Based Interface to load software onto the G8264. 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 a FTP/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 a FTP/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 7.8 for RackSwitch G8264 (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 N/OS 7.8 documentation as listed on [page 3](#).

ACLs

Metering is supported for IPv6 ACLs.

Edge Virtual Bridging (EVB)

VSI Database can be accessed via HTTP or HTTPS. The manager IP can be configured with an IPv4 or IPv6 address.

Multiple Spanning Tree Protocol (MSTP)

In IBM Networking OS 7.8, VLANs can be mapped to MSTP instances without creating them on the switch. In previous IBM Networking OS releases, the VLANs were created on the switch which often resulted in the switch having multiple unused VLANs.

Use the following commands to configure MSTP:

1. Configure port and VLAN membership on the switch.
2. Configure Multiple Spanning Tree region parameters and set the mode to MSTP.

```
RS8264(config)# spanning-tree mst configuration      (Enter MST configuration mode)
RS8264(config-mst)# name <name>                    (Define the Region name)
RS8264(config-mst)# exit
RS8264(config)# spanning-tree mode mst             (Set mode to Multiple Spanning Trees)
```

3. Map VLANs to MSTP instances:

```
RS8264(config)# spanning-tree mst configuration      (Enter MST configuration mode)
RS8264(config-mst)# instance <instance ID> vlan <vlan number or range>
```

NIST SP 800-131A Compliance

The implementations specified in this section are compliant with National Institute of Standards and Technology (NIST) Special Publication (SP) 800-131A.

The RackSwitch G8264 can operate in two boot modes:

- Compatibility mode (default): This is the default switch boot mode. This mode may use algorithms and key lengths that may not be allowed/acceptable by NIST SP 800-131A specification. This mode is useful in maintaining compatibility with previous releases and in environments that have lesser data security requirements.
- Strict mode: Encryption algorithms, protocols, and key lengths in strict mode are compliant with NIST SP 800-131A specification.

When in boot strict mode, the switch uses Secure Sockets Layer (SSL)/Transport Layer Security (TLS) 1.2 protocols to ensure confidentiality of the data to and from the switch.

Before enabling strict mode, ensure the following:

- The software version on all connected switches is IBM N/OS 7.8.
- The supported protocol versions and cryptographic cipher suites between clients and servers are compatible. For example: if using SSH to connect to the switch, ensure that the SSH client supports SSHv2 and a strong cipher suite that is compliant with the NIST standard.
- Compliant Web server certificate is installed on the switch, if using BBI.
- A new self-signed certificate is generated for the switch (RS8264(config)# access https generate-certificate). The new certificate is generated using 2048-bit RSA key and SHA-256 digest.
- Protocols that are not NIST SP 800-131A compliant must be disabled or not used.
- Only SSHv2 or higher is used.
- The current configuration, if any, must be saved in a location external to the switch. When the switch reboots, both the startup and running configuration are lost.
- Only protocols/algorithms compliant with NIST SP 800-131A specification are used/enabled on the switch. Please see the NIST SP 800-131A publication for details. The following table lists the acceptable protocols and algorithms:

Table 1. Acceptable Protocols and Algorithms

Protocol/Function	Strict Mode Algorithm	Compatibility Mode Algorithm
BGP	BGP does not comply with NIST SP 800-131A specification. When in strict mode, BGP is disabled. However, it can be enabled, if required.	Acceptable
Certificate Generation	RSA-2048 SHA-256	RSA 2048 SHA 256
Certificate Acceptance	RSA 2048 or higher SHA 224 or higher	RSA SHA, SHA2
HTTPS	TLS 1.2 only See “Acceptable Cipher Suites” on page 13 ;	TLS 1.0, 1.1, 1.2 See “Acceptable Cipher Suites” on page 13 ;
IKE		
Key Exchange	DH Group 24	DH group 1, 2, 5, 14, 24
Encryption	3DES, AES-128-CBC	3DES, AES-128-CBC
Integrity	HMAC-SHA1	HMAC-SHA1, HMAC-MD5
IPSec		
AH	HMAC-SHA1	HMAC-SHA1, HMAC-MD5
ESP	3DES, AES-128-CBC, HMAC-SHA1	3DES, AES-128-CBC, HMAC-SHA1, HMAC-MD5
LDAP	LDAP does not comply with NIST SP 800-131A specification. When in strict mode, LDAP is disabled. However, it can be enabled, if required.	Acceptable

Table 1. Acceptable Protocols and Algorithms

Protocol/Function	Strict Mode Algorithm	Compatibility Mode Algorithm
OSPF	OSPF does not comply with NIST SP 800-131A specification. When in strict mode, OSPF is disabled. However, it can be enabled, if required.	Acceptable
RADIUS	RADIUS does not comply with NIST SP 800-131A specification. When in strict mode, RADIUS is disabled. However, it can be enabled, if required.	Acceptable
Random Number Generator	NIST SP 800-90A AES CTR DRBG	NIST SP 800-90A AES CTR DRBG
Secure NTP	Secure NTP does not comply with NIST SP 800-131A specification. When in strict mode, secure NTP is disabled. However, it can be enabled, if required.	Acceptable
SLP	SHA-256 or higher RSA/DSA 2048 or higher	
SNMP	SNMPv3 only AES-128-CFB-128/SHA1 Note: Following algorithms are acceptable if you choose to support old SNMPv3 factory default users: AES-128-CFB/SHA1 DES/MD5 AES-128-CFB-128/SHA1	SNMPv1, SNMPv2, SNMPv3 DES/MD5, AES-128-CFB-128/SHA1
SSH/SFTP		
Host Key	SSH-RSA	SSH-RSA
Key Exchange	ECDH-SHA2-NISTP521 ECDH-SHA2-NISTP384 ECDH-SHA2-NISTP256 ECDH-SHA2-NISTP224 RSA2048-SHA256 DIFFIE-HELL- MAN-GROUP-EXCHANGE-SHA256 DIFFIE-HELL- MAN-GROUP-EXCHANGE-SHA1	ECDH-SHA2-NISTP521 ECDH-SHA2-NISTP384 ECDH-SHA2-NISTP256 ECDH-SHA2-NISTP224 ECDH-SHA2-NISTP192 RSA2048-SHA256 RSA1024-SHA1 DIFFIE-HELL- MAN-GROUP-EXCHANGE-SHA 256 DIFFIE-HELL- MAN-GROUP-EXCHANGE-SHA 1 DIFFIE-HELL- MAN-GROUP14-SHA1 DIFFIE-HELL- MAN-GROUP1-SHA1
Encryption	AES128-CTR AES128-CBC 3DES-CBC	AES128-CTR AES128-CBC RIJNDAEL128-CBC BLOWFISH-CBC 3DES-CBC ARCFOUR256 ARCFOUR128 ARCFOUR

Table 1. Acceptable Protocols and Algorithms

Protocol/Function	Strict Mode Algorithm	Compatibility Mode Algorithm
MAC	HMAC-SHA1 HMAC-SHA1-96	HMAC-SHA1 HMAC-SHA1-96 HMAC-MD5 HMAC-MD5-96
TACACS+	TACACS+ does not comply with NIST SP 800-131A specification. When in strict mode, TACACS+ is disabled. However, it can be enabled, if required.	Acceptable

Acceptable Cipher Suites

The following cipher suites are acceptable (listed in the order of preference) when the RackSwitch G8264 is in compatibility mode:

Table 2. List of Acceptable Cipher Suites in Compatibility Mode

Cipher ID	Key Exchange	Authentication	Encryption	MAC	Cipher Name
0xC027	ECDHE	RSA	AES_128_CBC	SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
0xC013	ECDHE	RSA	AES_128_CBC	SHA1	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
0xC012	ECDHE	RSA	3DES	SHA1	SSL_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
0xC011	ECDHE	RSA	RC4	SHA1	SSL_ECDHE_RSA_WITH_RC4_128_SHA
0x002F	RSA	RSA	AES_128_CBC	SHA1	TLS_RSA_WITH_AES_128_CBC_SHA
0x003C	RSA	RSA	AES_128_CBC	SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256
0x0005	RSA	RSA	RC4	SHA1	SSL_RSA_WITH_RC4_128_SHA
0x000A	RSA	RSA	3DES	SHA1	SSL_RSA_WITH_3DES_EDE_CBC_SHA
0x0033	DHE	RSA	AES_128_CBC	SHA1	TLS_DHE_RSA_WITH_AES_128_CBC_SHA
0x0067	DHE	RSA	AES_128_CBC	SHA256	TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
0x0016	DHE	RSA	3DES	SHA1	SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA

The following cipher suites are acceptable (listed in the order of preference) when the RackSwitch G8264 is in strict mode:

Table 3. List of Acceptable Cipher Suites in Strict Mode

Cipher ID	Key Exchange	Authentication	Encryption	MAC	Cipher Name
0xC027	ECDHE	RSA	AES_128_CBC	SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
0xC013	ECDHE	RSA	AES_128_CBC	SHA1	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
0xC012	ECDHE	RSA	3DES	SHA1	SSL_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
0x0033	DHE	RSA	AES_128_CBC	SHA1	TLS_DHE_RSA_WITH_AES_128_CBC_SHA
0x0067	DHE	RSA	AES_128_CBC	SHA256	TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
0x0016	DHE	RSA	3DES	SHA1	SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA
0x002F	RSA	RSA	AES_128_CBC	SHA1	TLS_RSA_WITH_AES_128_CBC_SHA
0x003C	RSA	RSA	AES_128_CBC	SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256
0x000A	RSA	RSA	3DES	SHA1	SSL_RSA_WITH_3DES_EDE_CBC_SHA

Configuring Strict Mode

To change the switch mode to boot strict mode, use the following command:

```
RS8264(config)# [no] boot strict enable
```

When strict mode is enabled, you will see the following message:

```
Warning, security strict mode limits the cryptographic algorithms used by secure protocols on this switch. Please see the documentation for full details, and verify that peer devices support acceptable algorithms before enabling this mode. The mode change will take effect after reloading the switch and the configuration will be wiped during the reload. System will enter security strict mode with default factory configuration at next boot up.
```

```
Do you want SNMPV3 support old default users in strict mode (y/n)?
```

Please see the *IBM Networking OS 7.8 RackSwitch G8264 Application Guide* for details on SNMPv3 users.

When strict mode is disabled, the following message is displayed:

```
Warning, disabling security strict mode. The mode change will take effect after reloading the switch.
```

You must reboot the switch for the boot strict mode enable/disable to take effect.

Limitations

In IBM N/OS 7.8, consider the following limitation/restrictions if you need to operate the switch in boot strict mode:

- Power ITEs and High-Availability features do not comply with NIST SP 800-131A specification.
- The G8264 will not discover Platform agents/Common agents that are not in strict mode.
- Web browsers that do not use TLS 1.2 cannot be used.
- Limited functions of the switch managing Windows will be available.

OpenFlow

IBM Networking OS supports OpenFlow protocol as described in the OpenFlow Switch Specifications 1.0.0 and 1.3.1.

The RackSwitch G8264 can be configured to run either version 1.0 or version 1.3 of the OpenFlow standard. By default, the switch uses OpenFlow 1.0.

For details, please see the *IBM Networking OS 7.8 RackSwitch G8264 Application Guide*.

Private VLANs

IBM N/OS supports Private VLAN configuration as described in RFC 5517.

Note: Private VLANs feature is not supported in stacking mode.

Scripting

You can create and execute local Python script on switches to create small programs that allow the switch to automatically provision itself, perform fault monitoring, upgrade the image files, or auto-generate configuration file. You can use local scripts as a key part of your auto-provisioning solutions. You can also manage scripts on the RackSwitch G8264.

You can implement version control systems, automatically generate alerts, create custom logging tools, and automate management of network devices. Using Python scripts, you can perform many functions that can be performed through the CLI. In addition to configuration, you can notify users by sending email or updating the SYSLOG.

See the *IBM Networking OS Python Programming Guide for the RackSwitch G8264* for details on how to create and execute Python scripts.

Telnet

Two attempts are allowed to log in to the switch. After the second unsuccessful attempt, the Telnet client is disconnected via TCP session closure.

User Access

Up to 20 users can be configured to allow access to the switch. Each user can be configured with a password and access level.

Resolved Issues

The following known issues have been resolved.

Private VLANs

- The sequence in which a private VLAN is configured is not the same as displayed in the output of the `RS8264(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)

Supplemental Information

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

**** VMLINUX ****

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, should 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 follow the guidelines given below:

- 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 the steps below:

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.

Known Issues

This section describes known issues for N/OS 7.8 on the RackSwitch G8264

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)

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-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.

Forwarding Database (FDB)

From IBM Networking OS 7.8 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

- 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)

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

- The port speed setting is not configurable for Finisar SFPs. Updating from BLADEOS 6.5 (or prior) to N/OS 6.8 (or later) will result in port speed configuration settings being reset to default values for ports using Finisar SFPs. (ID: 55063)

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)

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(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
```

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)

Routed Ports

- IBM N/OS CLI, SNMP, or BBI should not be used 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)

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(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)

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)
- Stacks using N/OS 7.6 will not accept new members that use any other version N/OS, including N/OS 7.8. Upload N/OS 7.6 (not N/OS 7.8) on the individual switch prior to joining it to a stack that uses N/OS 7.6. (ID: 69744)
- Switches using N/OS 7.6 will not automatically join an existing stack that is loaded with any other version of N/OS, including N/OS 7.8. Upload the N/OS 7.6 switch with N/OS 7.8 prior to joining it to the existing N/OS 7.8 stack. (ID: 69744; ID: 70387)
- During failover/failback operation of the stack Master, there may be some traffic loss for up to 40 seconds. (ID: 70568)
- DHCP has higher priority over static management IP configuration. If you want to configure a static management IP, you must first disable DHCP using the command: `RS8264(config)# no system dhcp mac <MAC address>`. (ID: 71181)
- When you issue the `show mac-address-table counters` command, the output is displayed after a delay of approximately 20 seconds. In the mean time, if there are any updates to the table, the output of the command may have incorrect information. (ID: XB218728)

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)

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).
- When VMready information commands are issued on a stack member, you may see VM validation messages from groups in the VM table, although the VM is not attached. (ID: XB210686)

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)
- When you change the CEE configuration while vNIC traffic is passing through the switch, the switch may behave in an unpredictable manner, such as receiving IBP/CBP discards. If this happens, reboot the switch to overcome the situation. To avoid this scenario, shut down all the ports before making any CEE-related configuration changes. (ID: 57414)