IBM System Networking RackSwitch™ G8264

Release Notes

For Networking OS 7.7

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 (June 2013)

Release Notes

This release supplement provide the latest information regarding IBM Networking OS 7.7 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* 7.7:

- *IBM Networking OS* 7.7 *Application Guide*
- *IBM Networking OS* 7.7 *Command Reference*
- *IBM Networking OS* 7.7 *ISCLI Reference*
- *IBM Networking OS* 7.7 *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 One10GbE 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

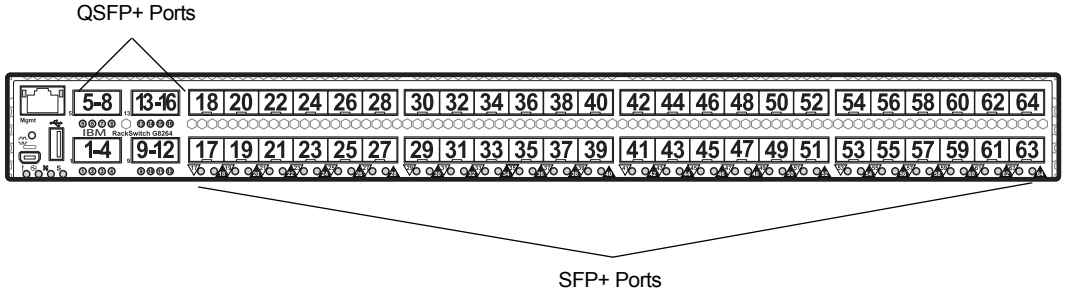

Transceivers

The following transceivers and DACs are available:

Table 1. RackSwitch G8264 Ordering Information

Part number	Description
	Transceivers
BN-CKM-S-T	SFP 1000BASE-T Copper Transceiver
BN-CKM-S-SX	SFP 1000BASE-SX Short Range Fiber Transceiver
BN-CKM-S-LX	SFP 1000BASE-LX Long Range Fiber Transceiver
BN-CKM-S-ZX	SFP 1000BASE-ZX Extra Long Range Fiber Transceiver
BN-CKM-SP-SR	SFP+ 10GBASE-SR Short Range Optical Fiber Transceiver
BN-CKM-SP-LR	SFP+ 10GBASE-LR Long Range Optical Fiber Transceiver
BN-CKM-SP-ER	SFP+ 10GBASE-ER Extended Range Optical Fiber Transceiver
BN-CKM-QS-SR	QSFP+ 40GBASE-SR4 Optical Fiber Transceiver
	Direct Attach Cables (DACs)
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 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. 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/systems/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 7.

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 7.7, 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 FDP 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 has 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)

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 21).

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 7.7 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 7.7 documentation as listed on page 3.

ARP - Local Proxy

An IP interface or a routed port that has local proxy ARP enabled allows the RackSwitch G8264 to respond to all ARP requests for IP addresses within the subnetwork, and to forward all traffic between hosts in the subnetwork. This feature is useful when hosts in a subnetwork are separated at Layer 2 with features such as Private VLAN. After responding to an ARP request, the G8264 sends an ARP request to the destination host for creating an ARP entry, if such an entry does not already exist in the ARP cache. If VRRP is enabled, the G8264 uses the virtual router MAC address of the master in the ARP response. If VRRP is not enabled, the G8264 uses the switch base MAC address.

When local proxy ARP is enabled on an interface, ICMP redirects must be disabled globally.

This feature can be enabled using the following commands:

On an IP interface:

```
RS8264(config-ip-if)# ip local-proxy-arp
```

On a routed port:

```
RS8264(config-if)# ip local-proxy-arp
```

Border Gateway Protocol

Multipath Relax

BGP multipath relax functionality allows load balancing across different autonomous system paths that have equal AS path length. This functionality can be enabled using the command:

```
RS8264(config-router-bgp)# bestpath as-path multipath-relax
```

DHCP

Host Name Configuration

The G8264 supports DHCP host name configuration as described in RFC 2132, option 12. DHCP host name configuration is enabled by default.

Host name can be manually configured using the following command:

```
RS8264(config)# hostname <name>
```

If the host name is manually configured, the switch does not replace it with the host name received from the DHCP server.

After the host name is configured on the switch, if DHCP or DHCP host name configuration is disabled, the switch retains the host name.

The switch prompt displays the host name.

Host name configuration can be enabled/disabled using the following command:

```
RS8264(config)# [no] system dhcp hostname
```

SYSLOG Server

During switch startup, if the switch fails to get the configuration file, a message can be recorded in the SYSLOG server.

The G8264 supports requesting of a SYSLOG server IP address from the DHCP server as described in RFC 2132, option 7. DHCP SYSLOG server request option is enabled by default.

Manually configured SYSLOG server takes priority over DHCP SYSLOG server.

Up to two SYSLOG server addresses received from the DHCP server can be used. The SYSLOG server can be learnt over a management port or a data port.

Use the `RS8264# show logging` command to view the SYSLOG server address.

DHCP SYSLOG server address option can be enabled/disabled using the following command:

```
RS8264(config)# [no] system dhcp syslog
```

Edge Virtual Bridging (EVB)

IPv6 Support

The Virtual Station Interface (VSI) database (VSIDB) manager can be configured with an IPv4 or IPv6 address. Use the following command to configure the VSIDB manager IP address:

```
RS8264(config)# virt evb vsidb 1
RS8264(conf-vsidb)# host <IPv4 or IPv6 address>          (Set VSI database Manager IP)
```

Configuring EVB in Stacking Mode

A *stack* is a group of up to [eight] RackSwitch G8264 switches with IBM Networking OS that work together as a unified system. The switches in a stack are interconnected by a stack trunk in a local ring topology.

An operational stack must contain one Master and one or more Members, as follows:

- **Master**

  One switch controls the operation of the stack and is called the Master. The Master provides a single point to manage the stack. A stack must have one and only one Master. The firmware image, configuration information, and run-time data are maintained by the Master and pushed to each switch in the stack as necessary.

- **Member**

  Member switches provide additional port capacity to the stack. Members receive configuration changes, run-time information, and software updates from the Master.

- **Backup**

  One member switch can be designated as a Backup to the Master. The Backup takes over control of the stack if the Master fails. Configuration information and run-time data are synchronized with the Master.

EVB can be configured on any port in the stack. Use the Master to configure EVB on a port in the stack. The port numbers in a stack use the following format:

<switch number>:<port number>

The Master process the EVB-related information for all the switch ports in a stack. The Master performs the VSIDB synchronization. The Master synchronizes all EVB changes with the Backup.

If the Master fails, the Backup takes over control of the stack. The VSI associations on the Master ports are lost. All other VSI associations remain unchanged.

Enhanced Password Security

Switch Login Levels

To enable better switch management and user accountability, three levels or *classes* of user access have been implemented on the G8264. Levels of access to CLI, Web management functions, and screens increase as needed to perform various switch management tasks. Conceptually, access classes are defined as follows:

- User interaction with the switch is completely passive—nothing can be changed on the G8264. Users may display information that has no security or privacy implications, such as switch statistics and current operational state information.
- Operators can only effect temporary changes on the G8264. These changes will be lost when the switch is rebooted/reset. Operators have access to the switch management features used for daily switch operations. Because any changes an operator makes are undone by a reset of the switch, operators cannot severely impact switch operation.

- Administrators are the only ones that may make permanent changes to the switch configuration—changes that are persistent across a reboot/reset of the switch. Administrators can access switch functions to configure and troubleshoot problems on the G8264. Because administrators can also make temporary (operator-level) changes as well, they must be aware of the interactions between temporary and permanent changes.

Access to switch functions is controlled through the use of unique user names and passwords. Once you are connected to the switch via console, remote Telnet, or SSH, you are prompted to enter a password.

Access to each user level (except admin account) can be disabled by setting the password to an empty value. To disable admin account, use the command: `RS8264(config)# no access user administrator-enable`.
Admin account can be disabled only if there is at least one user account enabled and configured with administrator privilege.

Strong Passwords

The administrator can require use of Strong Passwords for users to access the G8264. Strong Passwords enhance security because they make password guessing more difficult.

The following rules apply when Strong Passwords are enabled:

- Minimum length: 8 characters; maximum length: 64 characters
- Must contain at least one uppercase alphabet
- Must contain at least one lowercase alphabet
- Must contain at least one number
- Must contain at least one special character:

  Supported special characters: ! “ # % & ‘ ( ) ; < = >> ? [\] * + , - . / : ^ _ { | } ~
- Cannot be same as the username
- No consecutive four characters can be the same as in the old password

When strong password is enabled, users can still access the switch using the old password but will be advised to change to a strong password while attempting to log in.

Strong password requirement can be enabled using the following command:

```
RS8264(config)# access user strong-password enable
```

The administrator can choose the number of days allowed before each password expires. When a strong password expires, the user is allowed to log in one last time (last time) to change the password. A warning provides advance notice for users to change the password.

Locking Accounts

To protect the switch from unauthorized access, the account lockout feature can be enabled. By default, account lockout is disabled. To enable this feature, ensure the strong password feature is enabled. Then use the following command:

```
RS8264(config)# access user strong-password lockout
```

After multiple failed login attempts, the switch locks the user account if lockout has been enabled on the switch.

Re-enabling Locked Accounts

The administrator can re-enable a locked account by reloading the switch or by using the following command:

```
RS8264(config)# access user strong-password clear local user lockout username
                  <user name>
```

However, the above command cannot be used to re-enable an account disabled by the administrator.

To re-enable all locked accounts, use the following command:

```
RS8264(config)# access user strong-password clear local user lockout all
```

FCoE - Port Trunking

IBM N/OS 7.7 supports port trunking for FCoE connections. The Link Aggregation (LAG) can be used for separate FCoE traffic, or for Ethernet and FCoE traffic. Ports directly connected to servers cannot be combined in a LAG group.

Uplink ports, connected to the FCF, can be grouped as static or dynamic trunks.

Normal trunk operations such as creating/enabling the trunk, and adding/removing member ports can be performed. When a port is added to a trunk group, FCFs previously detected on the port will be deleted. The deleted FCF may be relearned later. However, this may cause flickering in the network traffic. We recommend that you make trunk group changes, if any, prior to live FCoE traffic.

Priority-based Flow Control (PFC) and Data Center Bridging (DCBX) are configured on a per-port basis. Each port in a trunk must have the same ETS, PFC, and DCBX configuration. When a port ceases to be the trunk group member, its configuration does not change.

Hot Links

Hot links provides basic link redundancy with fast recovery. Prior to IBM Networking OS 7.7, STP had to be globally disabled for configuring hot links. This restriction is no longer applicable. STP can be globally enabled but must be disabled on the ports used for hot links configuration.

IPv4 Address Conflict Detection

The RackSwitch G8264 uses a simple mechanism to detect if two hosts on the same subnetwork are using the same IPv4 address at the same time. The G8264 sends a gratuitous ARP request for its own IP address. If it receives an ARP response, it sends a syslog message with the IP address and MAC address of the host that is using its IP address.

The G8264 sends a gratuitous ARP request in the following situations:

- an IP interface comes up when:
  - the interface is enabled
  - a link comes up
  - a port goes into STP forwarding state
  - a member is added to a VLAN
- the IP address of an IP interface changes

LACP

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standard for grouping several physical ports into one logical port (known as a dynamic trunk group or Link Aggregation group) with any device that supports the standard. Please refer to IEEE 802.3ad-2002 for a full description of the standard.

A port's Link Aggregation Identifier (LAG ID) determines how the port can be aggregated. The Link Aggregation ID (LAG ID) is constructed mainly from the *system ID* and the port's *admin key*, as follows:

- **System ID**: an integer value based on the switch's MAC address and the system priority assigned in the CLI.
- **Admin key:** a port's Admin key is an integer value (1 - 65535) that you can configure in the CLI. Each G8264 port that participates in the same LACP trunk group must have the same *admin key* value. The Admin key is *local significant*, which means the partner switch does not need to use the same Admin key value.

For example, consider two switches, an Actor (the G8264) and a Partner (another switch), as shown in Table 2.

Table 2. Actor vs. Partner LACP configuration

Actor Switch	Partner Switch 1
Port 7 (admin key = 100)	Port 1 (admin key = 50)
Port 8 (admin key = 100)	Port 2 (admin key = 50)
Port 9 (admin key = 100)	Port 3 (admin key = 70)

In the configuration shown in Table 2, Actor switch ports 7 and 8 aggregate to form an LACP trunk group with Partner switch ports 1 and 2. Only ports with the same LAG ID are aggregated in the trunk group. Actor switch port 9 is not aggregated in the trunk group because it has a different LAG ID. Switch ports configured with the same admin key on the Actor switch but have a different LAG ID (due to Partner switch admin key configuration or due to partner switch MAC address being different) can be aggregated in another trunk group. i.e. Actor switch port 9 can be aggregated in another trunk group with ports that have the same LAG ID as port 9.

To avoid the Actor switch ports (with the same admin key) from aggregating in another trunk group, you can configure a trunk ID. Ports with the same admin key (although with different LAG IDs) compete to get aggregated in a trunk group. The LAG ID for the trunk group is decided based on the first port that is aggregated in the group. Ports with this LAG ID get aggregated and the other ports are placed in `suspended` mode. As per the configuration shown in Table 2, if port 7 gets aggregated first, then the LAG ID of port 7 would be the LAG ID of the trunk. Port 9 would be placed in suspended mode. When in suspended mode, a port transmits only LACP data units (LACPDUs) and discards all other traffic.

A port may also be placed in suspended mode for the following reasons:

- When LACP is configured on the port but it stops receiving LACPDUs from the partner switch.
- When the port has a different LAG ID because of the partner switch MAC or port LACP key being different. For example: when a switch is connected to two partners.

Trunk ID can be configured using the following command:

```
RS8264(config)# portchannel <65-128> lacp key <adminkey of the LAG> suspend-individual
```

LLDP - Stacking Mode

LLDP can be configured in stacking mode. LLDP can be configured only on the ports that are not used to create the stack. The LLDP configuration menus on the stacking ports are disabled.

When configuring LLDP on a port, use the correct port syntax.

Only the Master and Backup process the LLDP packets. The member switches do not transmit any LLDP information. The stack Master sends LLDP PDUs for all ports in the stack.

Microburst Detection

Microburst detection helps to identify peaks in data traffic at millisecond intervals. A maximum threshold is configured at the ingress port. Following commands can be used to check the status of an ingress port:

```
RS8264# show microburst microburst-status    (View microburst state of ingress ports)
RS8264# show microburst port-log             (View buffer utilization on ingress ports)
```

Microburst detection is disabled by default. To enable, use the following command:

```
RS8264# microburst enable
```

Reflective Relay

Reflective Relay (RR) is an operation where the switch forwards a frame back to the port on which it arrived if the destination MAC address is on the same port. When an EVB profile is configured on a port, RR is automatically enabled on the port after capability exchange with the peer, using the IEEE802.1QBG protocol. This is the usual mode of operation.

When the switch interoperates with devices that do not support IEEE 802.1QBG protocols, RR can be manually configured using the following command:

```
RS8264(config-if)# reflective-relay force enable
```

Manual RR and EVB profile cannot be configured on a port at the same time.

VLAG

Protocol Independent Multicast (PIM)

Added support for PIM configuration in a VLAG topology.

In a VLAG topology, IBM Networking OS supports PIM in Sparse Mode (PIM-SM) and Dense Mode (PIM-DM).

Health Check

Added support for configuring VLAG health check ports with an IPv4 or IPv6 address.

VMready

Up to 2048 VM profiles, 4093 VM groups, 4096 VMs, and 4096 VEs can be configured on the RackSwitch G8264. Of the total VMs, 2048 can be used in local groups.

vNIC Enhancements

vNIC Uplink Modes

The switch supports two modes for configuring the vNIC uplinks: dedicated mode and shared mode. The default is the dedicated mode. To enable the shared mode, enter the following command:

```
RS8264(config)# vnic uplink-share
```

In the dedicated mode, only one vNIC group is assigned to an uplink port. This port can be a regular port or a trunk port. The NIC places an outer tag on the vNIC group packets. This outer tag contains the vNIC group VLAN. The uplink NIC strips off the outer tag before sending out the packet.

In the shared mode, multiple vNIC groups can be assigned to an uplink port. This port can be a regular port or a trunk port. The vNIC groups share the uplink. You may assign a few vNIC groups to share an uplink and the other vNIC groups to have a single uplink each. In either case, the switch still operates in shared mode. As in the dedicated mode, the NIC places an outer tag on the vNIC group packets. This outer tag contains the vNIC group VLAN. The uplink NIC does not strip off the outer tag. The vNIC group tag defines the regular VLAN for the packet.This behavior is particularly useful in cases where the downstream server does not set any tag. Effectively, each vNIC group is a VLAN, which you can assign by configuring the VLAN to the vNIC group. You must enable the tag configuration on the uplink port.

The table below compares the configurations of the two modes.

Table 3. Comparison: Dedicated Mode vs. Shared Mode

Configuration Area	Dedicated Mode	Shared Mode
Port	"tagpvid" must be disabled.	"tagpvid" is user configurable.
	"pvid" = vNIC group VLAN.	"pvid" is user configurable.
	"tag" is user configurable.	"tag" must be enabled.
	Port can be added only to the vNIC group VLAN.	Port can be added to multiple VLANs in addition to the vNIC group VLANs that are automatically configured.
	Inserts vNIC group VLAN in the outer tag of ingress packets.	Inserts regular VLAN in the outer tag. VLAN tags are passed to and received from the uplink switch similar to vNIC ports.
		To handle untagged packets, configure the pvid/native VLAN of the uplink port to one of the vNIC group VLANs, and disable "tagpvid".
VLAN	Add the port to a vNIC group VLAN and delete it from any other VLAN when the vNIC group VLAN is enabled.	Add the port to all vNIC group VLANs that are sharing the port. Do not remove it from any other VLAN.
	Delete the port from the vNIC group VLAN and add it back to the default VLAN 1 when the vNIC group is disabled/deleted or when the vNIC feature is globally disabled.	Remove the port from a vNIC group VLAN when the vNIC group is disabled/deleted. When the vNIC feature is globally disabled or the port is not added in any vNIC group, remove the port from all vNIC group VLANs and add it back to default VLAN 1 if no non-vNIC VLAN exists on the port.
	Do not add a port or trunk to multiple vNIC groups that are enabled.	Can add a port or trunk to multiple vNIC groups that are enabled.
	Do not configure additional VLANs on the uplink ports.	Can configure additional VLANs on the uplink ports.
STP	An uplink port can only be in one STG.	An uplink port can be in multiple STGs.
	When you add a port to a vNIC group, STP is automatically disabled.	When you add a port to a vNIC group, STP is automatically disabled.
	When you remove a port from a vNIC group, STP is automatically reset to factory default.	When you remove a port from a vNIC group, STP is automatically reset to factory default.
Failover	An uplink up/event can trigger the failover state change only of one vNIC group.	An uplink up/event can trigger the failover state change of multiple vNIC groups.

LACP Trunks

The uplink used by vNIC groups can be a regular port, static trunk, or a dynamic trunk. If you are using a dynamic trunk for the uplink, you must configure the same LACP admin key for the vNIC uplink ports you want to group as a trunk on the upstream switch. You cannot connect these vNIC uplink ports to multiple upstream switches unless these switches are a single logical switch from LAG perspective.

Use the command below to add the LACP trunk as the uplink for a vNIC group.

```
RS8264(config)# vnic vnicgroup <vNIC group number>
RS8264(vnic-group-config)# key <LACP admin key>
```

VRRP - Next Hop Tracking

VRRP can be configured to track next hops. A health check mechanism, using ICMP ping or ARP requests, is used to track the next hop. If the health check succeeds or fails, the priority of the virtual router, for which the next hop tracking was configured, changes based on the configured tracking-priority-increment value.

If a VRRP group is enabled, its priority is calculated as the sum of all priorities for all active next hops. If an active next hop belongs to two virtual routers, then that next-hop's priority is added twice.

VRRP next hop tracking can be configured using the following commands:

Enable VRRP tracking on next hop:

```
RS8264(config)# virtual-router <x> track next-hops
```

Configure next hop IP address:

```
RS8264(config-vrrp)# virtual-router <x> next-hop <IP address> [ICMP|ARP]
[<interval>] [<retries>]

Default Values:
Health check protocol: ICMP
Interval: 2 seconds
Retries: 3
```

Note: A maximum of four unique next-hop IP addresses can be configured.

Configure priority for each next hop.

The value is added to the virtual router's priority when the next hop is active, and subtracted, when the next hop goes down:

```
RS8264(config-vrrp)# tracking-priority-increment next-hop <0-254>
```

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 7.7 on the RackSwitch G8264

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

- If you have configured the maximum number of ACLs supported on the G8264, any subsequent VSI ASSOCIATE request that requires an ACL may cause an invalid entry to be displayed in the output of the `>> Main# /info/virt/evb/vdp/vms` command. Typically, such invalid entries display a large value (e.g. 49093) in the TxACL column. (ID: 55254)
- 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.

IGMP

- The G8264 supports the following IGMP capacities (ID: 45775):
  - IGMP Snooping mode: 3072 IGMP and IPMC groups
  - IGMP Relay mode: 1000 IGMP groups and IPMC groups
- Only 1024 VLANs can be added to IGMP Snooping. Only 8 VLANs can be added to IGMP Relay. (ID: 45781)

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)
- 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)

OpenFlow

- When Openflow is enabled, any configuration for STP and other Layer 2 or Layer 3 features will be ignored. Configuration settings, logs, and files will continue to list Layer 2 and Layer 3 features as previously configured, rather than explicitly deconfigured or disabled. (ID: 61677)
- When FDB is used, multicast addresses can be added or deleted, but not modified. (ID: 66696)

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)

Port Routing

- Using the ISCLI, routing can be enabled or disabled on a per-port basis using the `[no] switchport` command in the `interface port` mode. Configuration of this feature is not currently supported in the menu-based CLI. (ID: 62846)

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)

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)
- Isolation for secondary VLANs is not honored across stacking interlinks. Traffic between the ports of a secondary VLAN is not isolated when those ports belong to different switches within a stack. Traffic in the secondary VLAN will be properly isolated only for traffic between ports of the same switch. (ID: 68340)
- Traffic with secondary VLAN ID is not forwarded to promiscuous ports. (ID: 70980)

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)

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 7.7. Upload N/OS 7.6 (not N/OS 7.7) 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 7.7. Upload the N/OS 7.6 switch with N/OS 7.7 prior to joining it to the existing N/OS 7.7 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)

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).

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)