IBM Networking OS™ 7.4 for RackSwitch™ G8316

Release Notes

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 (September 2012)

Release Notes

This release supplement provide the latest information regarding IBM Networking OS 7.4 for the RackSwitch G8316 (referred to as G8316 throughout this document).

This supplement modifies and extends the following IBM N/OS documentation for use with *N/OS* 7.4:

- *IBM Networking OS* 7.4 *Application Guide*
- *IBM Networking OS* 7.4 *Command Reference*
- *IBM Networking OS* 7.4 *ISCLI Reference*
- *IBM Networking OS* 7.4 *BBI Quick Guide*
- RackSwitch G8316 *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 G8316 contains sixteen 40GbE QSFP+ ports. 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 G8316 Front Panel

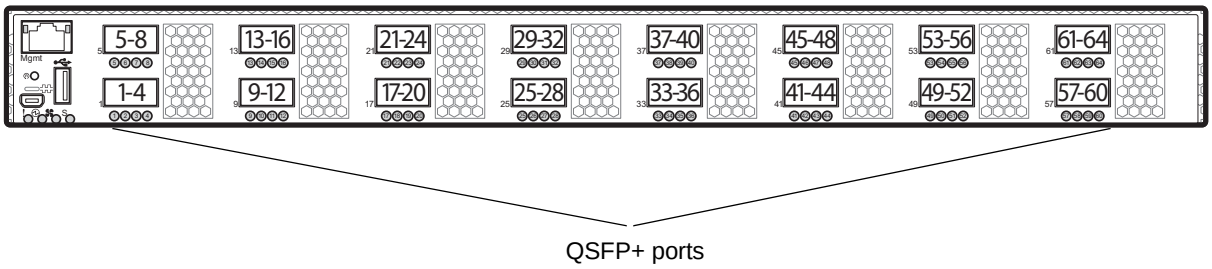

Transceivers and DACs

The following transceivers and Direct Attach Cables (DACs) are available:

Table 1. RackSwitch G8316 Transceivers and DACs

Description	Option part number	Tier 1 CRU part number
QSFP+ 40GBASE-SR4 Optical Fiber Transceiver	49Y7884	49Y7928
QSFP+ 40Gbps 1 meter DAC	49Y7890	49Y7934
QSFP+ 40Gbps 3 meter DAC	49Y7891	49Y7935
QSFP+ to four SFP+ 1 meter breakout DAC	49Y7886	49Y7930
QSFP+ to four SFP+ 3 meter breakout DAC	49Y7887	49Y7931
QSFP+ to four SFP+ 5 meter breakout DAC	49Y7888	49Y7932

The G8316 accepts any QSFP+ 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 G8316. 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 G8316, 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:

```
RS G8316# show boot
```

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 7.4, 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 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 G8316 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 19).

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 G8316. 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 7.4 for RackSwitch G8316 (G8316) has been updated to include several new features, summarized in the following sections. For more detailed information about configuring G8316 features and capabilities, refer to the complete N/OS 7.4 documentation as listed on page 3.

BGP Next Hop Attribute

BGP routing updates sent to a neighbor contain the next hop IP address used to reach a destination. In eBGP, the edge router, by default, sends its own IP address as the next hop address. However, this can sometimes cause routing path failures in Non-Broadcast Multiaccess Networks (NBMA) and when the edge router sends iBGP updates.

To avoid routing failures, you can manually configure the next hop IP address. In case of NBMA networks, you can configure the external BGP speaker to advertise its own IP address as the next hop. In case of iBGP updates, you can configure the edge iBGP router to send its IP address as the next hop.

Next hop can be configured on a BGP peer or a peer group. Use the following commands:

- Next Hop for a BGP Peer

```
RS G8316(config)# router bgp
RS G8316(config-router-bgp)# neighbor <number> next-hop-self
```

- Next Hop for a BGP Peer Group:

```
RS G8316(config)# router bgp
RS G8316(config-router-bgp)# neighbor group <number> next-hop-self
```

Diagnostics Enhancement

The following commands have been added to improve the ability to diagnose system issues:

- `RS G8316(config)#` **`show logging`** [**`messages`**] [**`severity`** *<0-7>*] [**`reverse`**] | [**`head`**|**`last`**] *<line number>*
- `RS G8316(config)#` **`show environment power`**
- `RS G8316(config)#` **`show environment fan`**
- `RS G8316#` **`show version`** [**`brief`**]
- `RS G8316#` **`show tech-support`** [**`l2|l3|link|port`**]
- `RS G8316#` **`system idle`** *<0-60>*
- `RS G8316(config)#` **`logging synchronous`** [**`level`** *<severity-level>***`|all`**]
- `RS G8316#` **`show who`**
- `RS G8316(config)#` **`access user clear`** *<session ID>*
- `RS G8316#` **`show line`**
- `RS G8316#` **`clear line`** *<session ID>*
- `RS G8316# [no]` **`debug lacp packet`**
- `RS G8316# [no]` **`debug spanning-tree bpdu`** [**`receive`**|**`transmit`**]

For detailed description of these commands, see the *IBM Networking OS 7.4 Command Reference* and *IBM Networking OS 7.4 ISCLI Reference Guides*.

Note: Use the debug commands with caution as they can disrupt the operation of the switch under high load conditions.

LLDP

LLDP transmissions can be configured to enable or disable inclusion of the following optional information:

Table 2. LLDP Optional Information Types

Type	Description	Default
portdesc	Port Description	Enabled
sysname	System Name	Enabled
sysdescr	System Description	Enabled
syscap	System Capabilities	Enabled
mgmtaddr	Management Address	Enabled
portvid	IEEE 802.1 Port VLAN ID	Disabled
portprot	IEEE 802.1 Port and Protocol VLAN ID	Disabled
vlanname	IEEE 802.1 VLAN Name	Disabled
protid	IEEE 802.1 Protocol Identity	Disabled
macphy	IEEE 802.3 MAC/PHY Configuration/Status, including the auto-negotiation, duplex, and speed status of the port.	Disabled
powermdi	IEEE 802.3 Power via MDI, indicating the capabilities and status of devices that require or provide power over twisted-pair copper links.	Disabled
linkaggr	IEEE 802.3 Link Aggregation status for the port.	Disabled
framesz	IEEE 802.3 Maximum Frame Size for the port.	Disabled
dcbx	Data Center Bridging Capability Exchange Protocol (DCBX) for the port.	Enabled

Management Interface Default IP Address

To facilitate switch boot up, the in-band and out-of-band management interfaces are configured with factory default IP addresses. These are as follows:

- VLAN 1/ Interface 1: 192.168.49.50/24
- Out-of-band Management Port A: 192.168.50.50/24

If you configure static IP addresses or if DHCP/BOOTP addresses are assigned to these interfaces, the factory default IP addresses will not be applied. By default, DHCP and BOOTP are enabled on the management interfaces.

If you add interface 1 to another VLAN and do not configure any IP address, the factory default IP address will be automatically assigned to the interface.

We recommend that you disable the factory default IP address configuration after the switch boot up and configuration is complete. Use the following command:

```
RS G8316(config)# no system default-ip data
```

OSPFv3 Over IPsec

BBI and SNMP support for OSPFv3 over IPsec has been added.

Persistent Terminal Length

The screen length for the current session can be set using the command: RS G8316# `terminal-length` *<0-300>*.

However, when the switch is reloaded, the screen length is set to default.

To set the screen length to be persistent across multiple sessions, use the following commands:

Telnet and SSH:

```
RS G8316(config)# line vty length <0-300>
```

Console:

```
RS G8316(config)# line console length <0-300>
```

The commands to set a persistent screen length are saved in the startup configuration and will be applied even when the switch is reloaded. If you need to change the screen length for a particular session, you can do so using the command for setting the current session's screen length.

PIM Multicast Routes

In release 7.4, the number of multicast routes (mroutes) available for use with PIM has been increased from 1000 to 2000.

Precision Time Protocol

As defined in the IEEE 1588-2008 standard, Precision Time Protocol (PTP) is a precision clock synchronization protocol for networked measurement and control systems. PTP provides system-wide synchronization accuracy and precision in the sub-microsecond range with minimal network and local clock computing resources. The synchronization is achieved through the exchange of messages: General messages that carry data but need not be time stamped; Event messages that are time stamped and are critical for clock synchronization accuracy.

A PTP network consists of PTP-enabled devices such as switches or routers. These devices consist of the following types of clocks:

- Master clock: In a PTP domain, the clock with the most precise time is considered the master clock. A best master clock algorithm determines the highest quality clock in a network.
- Ordinary clock: An ordinary clock synchronizes its time with the Master clock. The ordinary clock has a bidirectional communication with the master clock. By receiving synchronization/delay response and sending delay request packets, the ordinary clock adjusts its time with the master clock.
- Boundary clock: A boundary clock connects to multiple networks. It synchronizes with the attached master clock and in turn acts as a master clock to all attached ordinary clocks. Boundary clocks help to reduce the effect of jitter in Ethernet-based networks.
- Transparent clock: A transparent clock listens for PTP packets and adjusts the correction field in the PTP event packets as they pass the PTP device.

RackSwitch G8316 supports the configuration of ordinary clock and transparent clock. It cannot be a master clock as the switch does not participate in the master clock selection process.

By default, PTP version 2 is installed on the switch but is globally disabled. Use the following command to globally enable PTP:

```
RS G8316(config)# ptp {ordinary|transparent} enable
```

PTP is configured on switch ports. In case of trunk ports, the PTP configuration must be the same on all ports participating in the same trunk. The switch uses only one port from a trunk (typically the one used by a multicast protocol) to forward PTP packets.

By default, PTP is enabled on all the switch ports. To disable PTP on a port, use the following commands:

```
RS G8316(config)# interface port <port number>
RS G8316(config-if)# no ptp
```

Note: PTP cannot be enabled on management ports.

PTP packets occur on Control Plane Protection (CoPP) queue 36. You can change the CoPP packet rate limit per queue using the following command:

```
RS G8316(config)# qos protocol-packet-control rate-limit-packet-queue
<packet queue number (0-40)>  <packets per second (1-10000)>
```

Ordinary Clock Mode

When the RackSwitch G8316 is configured as an ordinary clock, it synchronizes its clock with the master clock. If the G8316 does not detect a master clock, it will not synchronize its clock with any other device. In this mode, the G8316's clock cannot be modified manually or using another time protocol such as Network Time Protocol (NTP).

As an ordinary clock, the G8316 synchronizes with a single PTP domain. The switch uses a delay-request mechanism to synchronize with the master clock. The switch uses a source IP address for the packets it generates. You can create a loopback interface for this purpose. By default, the switch uses the lowest interface in the VLAN from which the sync messages are received. To assign a loopback interface, use the following command:

```
RS G8316(config)# ip ptp source-interface loopback <interface number>
```

Note: If there are no interfaces on the switch that belong to the VLAN from which the sync messages are received, then the ordinary clock will not function. An error message will be generated. You can view this message using the `RS G8316# show ptp` command.

Transparent Clock Mode

When the G8316 is configured as a transparent clock, its time can be set manually or using any time protocol. You must configure PTPv2 for the transparent clock to function. The switch does not modify PTPv1 packets as they pass through the switch.

As a transparent clock, the G8316 supports syntonization (synchronization of clock frequency but not time) and synchronization with multiple domains.

Event packets received on all ports on the switch that have PTP enabled will be adjusted with the residence time. The switch sends all PTP packets to the multicast group address: 224.0.1.129. You can use Protocol Independent Multicast (PIM), Internet Group Management Protocol (IGMP), or any other multicast protocol to route the PTP packets.

Tracing PTP Packets

PTP packets can be traced on the PTP ports. These packets can be identified by their destination IP address and UDP ports. The following table includes the IEEE standard specification:

Table 3. IEEE Standard PTP Messages

Message	IP Address/UDP Port
PTP-primary: All PTP messages except peer delay mechanism messages	224.0.1.129
PTP-pdelay: Peer delay mechanism messages	224.0.0.107
Event Messages: Sync, delay request, peer delay request, peer delay response	319
General Messages: Announce, follow-up, delay response, peer delay response follow-up, management	320

Viewing PTP Information

The following table includes commands for viewing PTP information:

Table 4. PTP Information Commands

Command	Description
`RS G8316(config)# show ptp`	Displays global PTP information
`RS G8316(config)# show interface port <port number>`	Displays port information including port-specific PTP information
`RS G8316(config)# show ptp counters`	Displays ingress and egress PTP counters

Running Configuration

The following ISCLI command has been added to compare the running configuration with the startup configuration stored in FLASH.

```
RS G8316# show running-config diff
```

SNMP MIBs

- Added MIBs required for accessing LLDP data, as specified in IEEE 802.AB.
- Added MIBs required for accessing MLDv2 information.
- Added entity MIBs, as specified in RFC4133.
- Added MIBs required for managing host resources, as specified in RFC2790.

VLAG

Recovery from Cabling Error or Misconfiguration

In a Virtual Link Aggregation Group (VLAG) topology, if there is any cabling error or misconfiguration, Link Aggregation Control Protocol (LACP) may create loops. In such cases, the switch generates a syslog. The switch also shuts down the VLAG ports based on the following rules:

- If an LACP member port attempts to share the key with another LACP trunk, the port is shutdown.
- If the Primary and Secondary switches detect the misconfiguration or cabling error, then the LACP member port on the Secondary switch is shutdown.

  Note: Ports are shutdown only if LACP is used to configure trunks. This does not apply to static trunks.

Depending on the VLAG topology, you may see any of the following messages:

```
Access switch uses different LACP key to connect to adminkey <LACP key> on
vLAG switches, shutdown related ports

Adminkey <LACP key> isn't allowed to split to two trunks, please check the
configuration on the port which connect to port <port number>

Adminkey <LACP key> is used to connect the vLAG switches to two access
switches, shutdown related ports

Adminkey <LACP key> has different LACP priority on vLAG switches, shutdown
related ports
```

You will need to enable the ports after addressing the issue.

Capacity Enhancement

The maximum number of configurable VLAG instances is as follows:

- **With STP off:** Maximum of 52 VLAG instances
- **With STP on:**
  - **PVRST/MSTP with one VLAG instance per VLAN/STG:** Maximum of 52 VLAG instances
  - **PVRST/MSTP with one VLAG instance belonging to multiple VLANs/STGs:** Maximum of 20 VLAG instances

VMcheck

The G8316 primarily identifies virtual machines by their MAC addresses. An untrusted server or a VM could identify itself by a trusted MAC address leading to MAC spoofing attacks. Sometimes, MAC addresses get transferred to another VM, or they get duplicated.

The VMcheck solution addresses these security concerns by validating the MAC addresses assigned to VMs. The switch periodically sends hello messages on server ports. These messages include the switch identifier and port number. The hypervisor listens to these messages on physical NICs and stores the information, which can be retrieved using the VMware Infrastructure Application Programming Interface (VI API). This information is used to validate VM MAC addresses. Two modes of validation are available: Basic and Advanced.

Use the following command to select the validation mode or to disable validation:

```
RS G8316(config)# [no] virt vmgroup <VM group number> validate {basic|advanced}
```

Basic Validation

This mode provides port-based validation by identifying the port used by a hypervisor. It is suitable for environments in which MAC reassignment or duplication cannot occur.

The switch, using the hello message information, identifies a hypervisor port. If the hypervisor port is found in the hello message information, it is deemed to be a trusted port. Basic validation should be enabled when:

- A VM is added to a VM group, and the MAC address of the VM interface is in the Layer 2 table of the switch.
- A pre-provisioned VM interface that belongs to a VM group connects to the switch.
- A trusted port goes down. Port validation must be performed to ensure that the port does not get connected to an untrusted source when it comes back up.

Use the following command to set the action to be performed if the switch is unable to validate the VM MAC address:

```
RS G8316(config)# virt vmcheck action basic {log|link}

log - generates a log
link - disables the port
```

Advanced Validation

This mode provides VM-based validation by mapping a switch port to a VM MAC address. It is suitable for environments in which spoofing, MAC reassignment, or MAC duplication is possible.

When the switch receives frames from a VM, it first validates the VM interface based on the VM MAC address, VM Universally Unique Identifier (UUID), Switch port, and Switch ID available in the hello message information. Only if all the four parameters are matched, the VM MAC address is considered valid.

In advanced validation mode, if the VM MAC address validation fails, an ACL can be created to drop the traffic received from the VM MAC address on the switch port. Use the following command to specify the number of ACLs to be used for dropping traffic:

```
RS G8316(config)# virt vmcheck acls max <1-256>
```

Use the following command to set the action to be performed if the switch is unable to validate the VM MAC address:

```
RS G8316(config)# virt vmcheck action advanced {log|link|acl}
```

Following are the other VMcheck commands:

Table 5. VMcheck Commands

Command	Description
RS G8316(config)# virt vmware hello {ena|hport <*port number*>|haddr|htimer}	Hello messages setting: enable/add port/advertise this IP address in the hello messages instead of the default management IP address/set the timer to send the hello messages
RS G8316(config)# no virt vmware hello {enable|hport <*port number*>}	Disable hello messages/remove port
RS G8316(config)# [no] virt vmcheck trust <*port number or range*>	Mark a port as trusted; Use the no form of the command to mark port as untrusted
RS G8316# no virt vmcheck acl [mac-address [<*port number*>]|port]	Delete ACL(s): all ACLs/an ACL by MAC address ((optional) and port number) /all ACLs installed on a port

Supplemental Information

This section provides additional information about configuring and operating the G8316 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 ..............................

1 - Change booting image
2 - Change configuration block
3 - Boot in recovery mode (tftp and xmodem download of images to
    recover switch)
4 - Xmodem download (for boot image only - use recovery mode for
    application images)
5 - Reboot
6 - Exit

Please choose your menu option: 3
```

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 a software image recovery, press 3 and follow the screen prompts.
- To perform an Xmodem download (boot image only), press 4 and follow the screen prompts.
- To exit the Boot Management menu, press 6. 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, CRT, 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 **Boot in recovery mode**. You will see the following display:

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

   – If you choose option **x** (Xmodem serial download), go to step 5.
   – If you choose option **t** (TFTP download), go to step 6.

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

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

   a. Press <**Enter**> to set the system into download accept mode. When the readiness meter displays (a series of “C” characters), start Xmodem on your terminal emulator.
   b. When you see the following message, change the Serial Port characteristics to 9600 bps:

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

   c. When you see the following prompt, enter the image number where you want to install the new software and press <**Enter**>.

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

   d. The following message is displayed when the image download is complete. Continue to step 7.

```
Installing image as image1...
Image1 updated successfully
Please select one of the following options:
        T) Configure networking and tftp download an image
        X) Use xmodem 1K to serial download an image
        R) Reboot
        E) Exit
```

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

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

a. Enter the required information and press <**Enter**>.

b. You will see a display similar to the following:

```
        Host IP    : 10.10.98.110
        Server IP  : 10.10.98.100
        Netmask    : 255.255.255.0
        Broadcast  : 10.10.98.255
        Gateway    : 10.10.98.254
Installing image 6.8.3_OS.img from TFTP server 10.10.98.100
```

c. When you see the following prompt, enter the image number where you want to install the new software and press <**Enter**>.

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

d. The following message is displayed when the image download is complete. Continue to step 7.

```
Installing image as image1...
Image1 updated successfully
Please select one of the following options:
        T) Configure networking and tftp download an image
        X) Use xmodem 1K to serial download an image
        R) Reboot
        E) Exit
```

7. Image recovery is complete. Perform one of the following steps:
   – Press **r** to reboot the switch.
   – Press **e** to exit the Boot Management menu
   – Press the Escape key (<**Esc>**) to re-display the Boot Management menu.

Recovering a Failed Boot Image

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

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports Xmodem download (for example, HyperTerminal, CRT, 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 **4** for **Xmodem download**. You will see the following display:

```
Perform xmodem download

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

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

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

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

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

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

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

Boot image recovery is complete.

VLAGs

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

- The ISL should include enough ports to accommodate the peer-to-peer traffic.
- 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 want to add a port to an ISL trunk/VLAN, or if you want to remove any particular port from an ISL trunk/VLAN, first shutdown that port. Make the necessary changes and re-enable the port.
- 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:
     `RS G8316 (config)# no vlag adminkey <key> enable` (or)
     `RS G8316 (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 7.4 on the RackSwitch G8316

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)

FCoE

- 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.

Hotlinks

- Prior to enabling hotlinks, Spanning-Tree should be globally disabled using the following command (ID: 47917):

```
RS G8316(config)# spanning-tree mode dis
```

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)
- Maximum of 1024 IGMP groups are set to 0n for the IP Options feature, regardless of the ipmc-opt profile used to boot (***ipmc-opt acls-128***, ***ipmc-opt acls-256***, ***ipmc-opt acls-384***, ***ipmc-opt acls-none***). (ID: 55931)

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

- IPsec does not support OSPFv3 virtual links. (ID: 48914)
- Packet fragmentation over IPsec is supported in transport mode only. Fragmentation is not available in tunneling mode. (ID: 50291)

Management Port Rate Limiting

Rate-limiter functionality is not available yet for the p4040 processor on the management port. The processor is able to handle the traffic on management port. (ID: 56871)

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. Configuration based on virtual links is not currently supported.
  - 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)

PIM

- When the PIM mroute table is cleared for only one of the intermediate routers in the topology, the maximum number of entries might not be recreated, or it might take 10-15 minutes until the entries are recreated. (ID: 55369)

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` or `interface portchannel` modes. Configuration of this feature is not currently supported in the menu-based CLI. (ID: 62846)

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)

Qlogic NIC

Some link flapping (twice) can occur when you reboot a server with a Qlogic NIC. This happens for the link between Qlogic QLE8152 and the G8316. (ID: 55526)

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)

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. However, ping can be facilitated if IP interfaces with VLAN IDs corresponding to those of the VM groups are configured on the switch.
- 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).
- Bandwidth metering drops excess packets when the configured limits on the vNIC pipe are reached. CEE Enhanced Transmission Selection will be ignored. (ID: 50950)
- 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)