IBM System Networking RackSwitch™ G8316

Release Notes

For Networking OS 7.9

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

Second Edition (June 2014)

Release Notes

=Intro==TOR

=Intro==IBM

This release supplement provide the latest information regarding IBM Networking OS 7.9 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* 7.9:

- *IBM Networking OS* 7.9 *Application Guide*
- *IBM Networking OS 7.9 Command Reference*
- *IBM Networking OS* 7.9 *ISCLI Reference*
- *IBM Networking OS 7.9 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

=Intro==TOR All

=Intro==IBM All

> > > > > > > > > Comment: BladeCenter H and HT are for Janice, not Rizzo.

=HW==IBM Janice

=HW==IBM Rizzo

=HW==TOR Piggy

=HW==TOR Piglet

=HW==TOR Scooter

=HW==TOR Gryphon

=HW==TOR Bingo

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

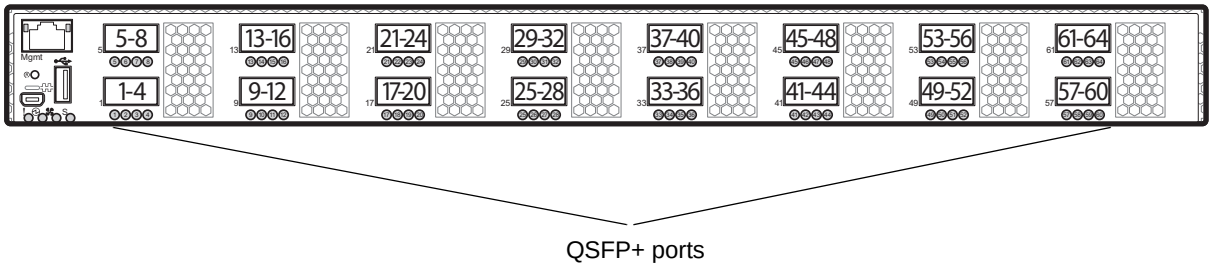

Transceivers and DACs

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

=Intro==All

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/support

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

```
RS G8316# show version
```

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

- Load a new software image and boot image onto an 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.

Updating VLAG Switches with IBM Networking OS 7.x

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

[Old description; I’m not sure whether it is still valid for switches other than Piglet. --Lynn]

1. Shut down all the ports on both VLAG Peers.
2. Upgrade VLAG Peer 1 to 7.x (both OS and Boot Image).
3. Upgrade VLAG Peer 2 to 7.x (both OS and Boot Image).

   Note: Both VLAG peers must be updated with the same version.
4. Save the configuration using the following command:

   ```
   RS G8316(config)# copy running-configuration startup-configuration
   ```
5. Reload VLAG Peer 1.

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

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

```
RS G8316(config)# copy running-configuration startup-configuration
```

2. Use TFTP or FTP to copy the new OS image and boot image onto both vLAG switches.
3. Shut down all ports except the ISL ports and the health check port on the primary switch (Switch 1).

   Note: Do not save this configuration.
4. Reload Switch 1, Switch 2 will assume the vLAG primary role
5. Once Switch 1 has rebooted, Switch 1 will take the secondary role.
6. Shut down all ports except the ISL ports and the health check port on Switch 2.

   Note: Do not save this configuration.
7. Reload Switch 2, Switch 1 will reassume the vLAG primary switch role.
8. Once Switch 2 has reloaded, make sure Switch 1 has transitioned to vLAG primary and Switch 2 has transitioned to secondary.
9. Verify all the vLAG clients have converged using the following command:

```
RS G8316(config)# show vlag information
```

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 14).

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

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

  Note: Be sure to download both the new boot file and the new image file.
- The hostname or IP address of the FTP or TFTP server

  Note: The DNS parameters must be configured if specifying hostnames.
- The name of the new software image or boot file

When the software requirements are met, use one of the following procedures to download the new software to your switch. You can use the N/OS CLI, the ISCLI, or the BBI to download and activate new software.

Loading Software via the ISCLI

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

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

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

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

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

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

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

4. If required by the FTP or TFTP server, enter the appropriate username and password.
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 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 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 an FTP or TFTP server, enter the server's information in the FTP/TFTP Settings section.
4. In the Image Settings section, select the image version you want to replace (Image for Transfer).
   – If you are loading software from an FTP or TFTP server, enter the file name and click **Get Image**.
   – If you are loading software from your computer, click **Browse**.

     In the File Upload Dialog, select the file and click **OK**. Then click **Download via Browser**.

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

New and Updated Features

N/OS 7.9 for RackSwitch 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 7.9 documentation as listed on page 3.

Auto VLAN Tagging on Trunk Ports

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

BGP Community Lite

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

Display BGP Routes

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

Dynamic Policy-Based Routing (PBR)

Dynamic Policy-Based Routing (PBR) allows a switch to isolate traffic to keep tenants from communicating with each other while routing all northbound traffic normally. When a tenant tries to send traffic to a neighboring tenant, the switch will redirect the traffic to an upstream router so that an upstream firewall can decide if that traffic is allowed.

This feature scales better than Virtual Routing and Forwarding (VRF), but it only works in environments that have non-overlapping IP spaces, where each tenant has a unique IP space.

ESN to SNMP

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

IBM N/OS Menu-Based Interface Removal

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

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

IPSec over Virtual Links

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

IPv6 Counter Enhancement

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

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

Layer 3 ARP Table Full

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

Link Aggregation Control Protocol (LACP) Individual Mode

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

OpenFlow 1.3.1 Support

Added features from Openflow Switch Specification Version 1.3.1.

Openflow 1.3.1 Group Support

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

Openflow Support for Static LAG over Edge Ports

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

sFlow Support in Openflow Ports

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

QBG Support

This release Implements the IEEE 802.1Qbg standard, allowing server-network edge virtualization, uniform view of the VMs in the network hypervisors, visibility of VM traffic, and automatic migration of port profiles.

RMON Support (RFC1757, RFC2819)

Remote network (RMON) monitoring devices, often called monitors or probes, are instruments that exist for the purpose of managing a network. This release supports RMON for ethernet statistics, ethernet history and alarm and event groups.

Secure FTP

This release adds support for Secure FTP (sFTP).

Service Location Protocol

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

Spanning Tree Protocol (STP) Range Enhancement

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

SNMP

The following features have been added to SNMP support.

SNMP ACL

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

SNMP Trap Host

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

Use SSH Public Keys for up to 20 Local Switch Users

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

vLAG MSTP Enhancement

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

vLAG PIM with Multicast Sources

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

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

Supplemental Information

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

The Boot Management Menu

=====Everyone except Bingo

=====Bingo Only

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

=====Everyone except Bingo

=====Bingo Only

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:

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

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

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

  On the VLAG Secondary Peer:

  1. Shutdown the VLAG ports on which you need to make the change.
  2. Disable their VLAG instance using the command:
     `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 or when STP is PVRST.

> > > > > > > > **Comment:** Bug 32226—Autonegotiation must be the same on both sides of an external link.

> > > > > > > > **Comment:** Bug 29677 Change to default Internal port autonegotiation setting. Now enabled by default to accommodate Wake-Over-LAN. Expected behavior. No fix planned.

> > > > > > > > **Comment:** Bug 35435. Loss of vCenter connectivity may cause changes to be unsynchronized, even when connectivity is later restored. Forced scan is required.

Resolved Issues

The following known issues have been resolved.

Private VLANs

Traffic with secondary VLAN ID is not forwarded to promiscuous ports. (ID: 70980)

Other Corrections

The following issues have also been resolved:

- Reset due to software exception:
  ID ###

See the software change log included with the software release files for details on these issues, as well as for corrections made in prior releases.

Known Issues

> > > > > > > > **Review Note:** Are any of these issues corrected? Are there new issues to report?

This section describes known issues for N/OS 7.9 on the RackSwitch G8316

ACLs

- Access Control Lists (ACLs) which are configured to match both a destination MAC address and an egress port fail to act when the matching packets are encountered. As a result, ACLs cannot be used to block traffic exiting specific ports for specific static multicast MAC addresses. (ID 32896)

  Solution: Instead of using an ACL to block the traffic, configure a static multicast route that includes all ports other than those you wish to block. Consider an example where you wish to block port EXT1 for DMAC `01:02:03:04:05:FF` on the default VLAN (VLAN 1). In this case, you would add a multicast route that includes all ports except EXT1. For example:

  ```
  # /cfg/l2/fdb/mcast/add <Destination MAC> <VLAN> <Ports list or range to allow>
      -or-
  # /cfg/l2/fdb/mcast/add 01:02:03:04:05:FF 1 INT1-INT14
  EXT2-EXT10
  ```

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

BGP

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

- Maximum number of Autonomous Systems (AS) per path is 20. (ID: 42371) (was Gryphon-FC)

CPU Utilization

- When changing form CLI mode to ISCLI mode and rebooting the switch, the management CPU will inaccurately register as 100% busy for up to five minutes. This is not known to cause any switch functionality issues and can be ignored.

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)
- 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.
- When using vNICs with FCoE, vNIC groups that participate in FCoE cannot include ENodes on both vNIC ports and regular (non-vNIC) ports. Within any given vNIC group, ENodes must be attached only to vNIC ports, or only to regular ports, but not both.

FCOE

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

FIPS

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

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

Forwarding Database (FDB)

From IBM Networking OS 7.9 onwards, MAC address information is no longer learned by control packets such as LACPDUs. This behavior is as expected. (ID: XB253517)

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)

ISCLI

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

LACP

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)

On-Box Scripting

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

  Old Keys:

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

  New Keys:

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

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

Openflow

When you configure a port to use Openflow, spanning tree protocol is automatically disabled on that port. (ID: XB266710)

OSPF

- When a network contains both static and redistributed routes, the static routes are listed as *preferred*. (ID: 45431)

- 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.

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

Precision Time Protocol

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

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)

QoS

When the following command is issued command is issued, "Dropped Packets" and "Dropped Bytes" counters will be displayed as '0' due to hardware limitations: (ID: XB233503)

```
RS G8316(config)#
show interface port <swunit:port_num> egress-mcast-queue-counters

For example:
RS G8316(config)# show interface port 1:24 egress-mcast-queue-counters

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

- If you delete multiple VLANs using SNMP, you may see an error if the SNMP packet size exceeds 1800 bytes. (ID: XB228120)

Spanning Tree

- When using LACP with PVRST, it is not recommended to configure the switch as the STP root bridge. When doing so, traffic can be discarded for up to 30 seconds on affected LACP ports while initial STP path states are being resolved (discarding, learning, forwarding). (ID: 63315)
- After changing from MSTP to a different STP mode, some MSTP commands will still appear in the configuration file. The non-applicable MSTP commands do not affect switch operation and can be ignored. (ID: 64388)

> > > > > > > > **Review Note:** Bug 30172 When changing the TACACS+ secondary password, the next authentication requires the secondary password. Planned to be fixed in 6.1. Shown in Janice docs. Does this currently affect BigBird/Rizzo/Zoe?

UFP

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

Virtual Link Aggregation Groups

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

VLANs

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

VMready

- VMs belonging to different ESX servers cannot ping each other across different VM groups. Because the VM groups belong to different VLANs, this is appropriate and expected behavior. 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)