IBM System Networking RackSwitch™ G8124/G8124-E

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

The RackSwitch G8124/G8124-E is an all 10Gb Ethernet rackable aggregation switch with unmatched line-rate Layer 2/3 performance. It uses a wire-speed, non-blocking switching fabric that provides simultaneous wire-speed transport of multiple packets at low latency on all ports.

This release supplement provide the latest information regarding IBM Networking OS 7.7 for the RackSwitch G8124/G8124-E (collectively referred to as G8124 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 G8124 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

N/OS 7.7 software is supported on the G8124, a high performance Layer 2-3 network switch.

The G8124 contains 24 ten Gigabit Small Form-factor, Pluggable (SFP+) slots and two 1Gb management ports. The 10Gb SFP+ slots can accept 1Gb copper transceivers, 10Gb optical transceivers, or Direct Attach Cables (DAC).

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 G8124 Front Panel

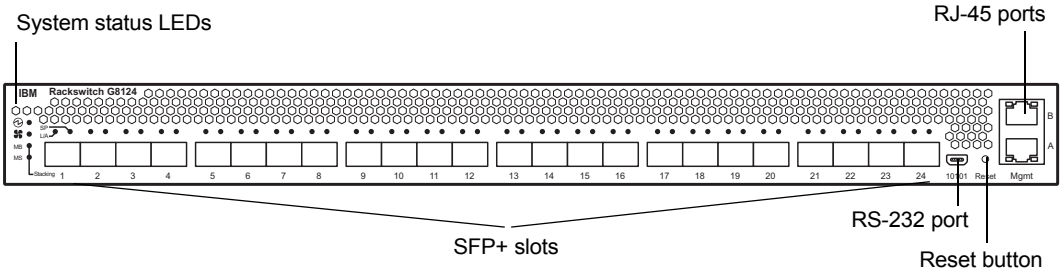

Transceivers

The G8124 accepts any of the following transceivers available from BLADE Network Technologies:

Table 1. Recommended SFP+ Transceiver

Part number	Description
BN-CKM-S-T	SFP Transceiver, 1000Base-T Copper
BN-CKM-S-SX	SFP Transceiver, 1000Base-SX Short Range Fiber
BN-CKM-S-LX	SFP Transceiver, 1000Base-LX Long Range Fiber
BN-CKM-SP-SR	SFP+ Transceiver, 10GBase-SR Short Range
BN-CKM-SP-LR	SFP+ Transceiver, 10GBase-LR Long Range

The G8124 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 G8124. 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 G8124, 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 G8124# 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 8.

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 1.x

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

```
RS G8124(config)# boot configuration-block factory
```

After updating:

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

  If the NTP interval value is lower than 5, then after software upgrade the NTP interval is set to the minimum value of 5. (ID: 36500)
- The default values and range values for IGMP report timeout parameter are different for release 5.x and later as compared to release 1.1:
  - On release 1.1 the range for IGMP report timeout is <130-1225> seconds with a default of 260 seconds.
  - On release 5.x and later, the range is <1-255> minutes with a default of 10 minutes.

  During upgrade, the value of IGMP report timeout is set to the new default value (10). The value does not appear in the running configuration output. (ID: 35578)
- On release 1.1, the default setting for Hotlinks BPDU is `enabled`, and on release 5.x and later the default setting is `disabled`. During upgrade, the Hotlinks BPDU command is set to `disabled`. (ID: 36385)
- The default value for the `access https` command is different compared to release 1.x. On release 1.1 the default setting is `enabled`, and on release 5.x and later, the setting is `disabled`. During upgrade, `access https` is set to `disabled`. (ID: 36834)

Reverting to BLADEOS 1.x or Prior

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

Updating from BLADEOS 5.x or Prior

After updating:

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

Updating from BLADEOS 5.1 or Prior

When you upgrade the G8124 from release 5.1 or prior, the configuration block is converted to match the new software.

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

Updating from BLADEOS 6.5.1 or Prior

After updating:

- The default value for port flow control is different compared to release 6.5.1 or prior. After upgrading to release 6.5.2 or later, it is recommended that the administrator review the configured flow control settings and make any appropriate changes. (ID: 43781)
- Previously configured static MAC addresses must be reconfigured after the upgrade (ID: 35659)
- The administrator may no longer configure the number of IGMP queries sent when a Leave message is received. The count is set to 2 after the upgrade. (ID: 36638)
- TACACS+ secure back door is disabled during the upgrade. If you use TACACS+ secure back door, you must re-enable it after resetting the switch. (ID: 34707)
- During software upgrade, the system time zone setting is lost. Re-configure the system time zone. (ID: 36493)

Updating from BLADEOS 6.6 or Prior

After updating:

- The default value for port autonegotiation is different compared to release 6.6 or prior. Starting with BLADEOS 6.7, autonegotiation is turned on by default. Autonegotiation cannot be configured for 10-Gig links. After upgrading, it is recommended that the administrator review the port settings and make any appropriate changes.
- 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 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 G8124 can store up to two different switch software images (called `image1` and `image2`) as well as special boot software (called `boot`). When you load new software, you must specify where it should be placed: either into `image1`, `image2`, or `boot`.

For example, if your active image is currently loaded into `image1`, you would probably load the new image software into `image2`. This lets you test the new software and reload the original active image (stored in `image1`), if needed.

CAUTION:
When you upgrade the switch software image, always load the new boot image and the new software image before you reset the switch. If you do not load a new boot image, your switch might not boot properly (To recover, see "Recovering from a Failed Upgrade" on page 16).

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

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

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

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

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

Loading Software via the 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 G8124. 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 G8124 (G8124) has been updated to include several new features, summarized in the following sections. For more detailed information about configuring G8124 features and capabilities, refer to the complete N/OS 7.7 documentation as listed on page 3.

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:

```
RS G8124(config-router-bgp)# bestpath as-path multipath-relax
```

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 G8124. 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 G8124. 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 G8124. 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 G8124. 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: `RS G8124(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 G8124. 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:

```
RS G8124(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:

```
RS G8124(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:

```
RS G8124(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:

```
RS G8124(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.

Data Center Bridging (DCBX) is 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 G8124 uses a simple mechanism to detect if two hosts on the same subnetwork are using the same IPv4 address at the same time. The G8124 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 G8124 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 G8124 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 G8124) 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:

```
RS G8124(config)# portchannel <53-104> lacp key <adminkey of the LAG> suspend-individual
```

VLAG

Protocol Independent Multicast (PIM)

Note: This section is applicable only to G8124-E.

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).

Supplemental Information

This section provides additional information about configuring and operating the G8124 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:
     `RS G8124 (config)# no vlag adminkey <key> enable` (or)
     `RS G8124 (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 G8124

BGP Debug

While enabling or disabling BGP debug for a particular peer/IP address, the logging behavior may not be as expected. Following is a workaround: (ID: 59104)

To enable BGP debug for a particular peer:

1. Enable BGP debug for all the peers.
2. Disable BGP debug for all the peers.
3. Enable BGP debug for a particular peer.

To disable BGP debug for a particular peer:

1. Enable BGP debug for all the peers.
2. Disable BGP debug for all the peers.
3. Enable BGP debug for all the peers except the one for which you want it disabled.

Boot Configuration Block

- In the CLI, the boot configuration command (`/boot/conf` <*block*>) examines only the initial character of the *block* option. Invalid *block* strings (those other than `active`, `backup`, or `factory`) that use a valid first character (a, b, or f) will be interpreted as the matching valid string. (ID: 42422)

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.

Deployment Profiles

- When changing from a different deployment profile, if a resource in the new profile has less capacity than is in use in the prior profile, an error message may be logged when the mode is changed. (ID: 64009)

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)
- 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)

IKEv2

- IKEv2 cannot be configured on management ports. Configure IKEv2 only on data ports. (ID: 57427)

IPsec

- When configuring IPsec to operate between IBM switches, keys of various lengths are supported. However, when using IPsec to connect with non-IBM devices, the manual policy session keys must be of the following fixed lengths:
  - For the AH key:
    - SHA1 = 20 bytes
    - MD5 = 16 bytes
  - For the ESP auth key:
    - SHA1 = 20 bytes
    - MD5 = 16 bytes
  - For the ESP cipher key:
    - 3DES = 24 bytes
    - AES-cbc = 24 bytes
    - DES = 8 bytes

ISCLI

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

LACP

- Even if STP is disabled, when changing the LACP mode to off (from active or passive mode), the port is placed in the DISC state to prevent network loops. (ID: 42768)
- Under heavy switch load conditions, LACP links may flap when configured with short timeout mode. To stabilize LACP under heavy load, it is recommended to use the long timeout mode instead. (ID: 66173)

OSPF

- When reverting from BLADEOS 6.4 (or later) to BLADEOS 6.3 (or prior), OSPF areas 3 through 5 (if configured) are consolidated under OSPF area 0 instead of being removed. If this is not the desired behavior, delete OSPF areas 3, 4, and 5 (if configured) prior to reverting to BLADEOS 6.3 or prior, or verify expected OSPF area 0 configuration after reverting. (ID: 43327)
- 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)

PIM

- PIM may be configured via the ISCLI only. PIM configuration and information is not available using the CLI menu system, the BBI, or via SNMP. (ID: 38443, 39279, 39445, 39849, 40046)
- PIM supports standard multicast frame sizes. However, uncommon use of jumbo frames for multicasts has not been confirmed for PIM operation.
- PIM Source-Specific Multicast (PIM-SSM) is not supported.
- Anycast RP is not supported.
- PIM RP filters are not supported.
- PIM is not supported simultaneously with vNICs or FCoE.
- When using the `clear ip pim mroute` command to clear a large list of PIM neighbor entries, the PIM state on the switch can lose synchronization with its PIM neighbors. If this occurs, synchronize PIM by globally disabling and then re-enabling PIM on the switch.

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)

Private VLANs

- The sequence in which a private VLAN is configured is not the same as displayed in the output of the `RS G8124(config)# show running-config` command. Hence, if you copy and paste the private VLAN configuration from the output of the above command, the private VLAN configuration will be lost. (ID: 67169)
- Traffic with secondary VLAN ID is not forwarded to promiscuous ports. (ID: 70980)

Rate Limiting

- The operational precision of rate limits set for bcast and mcast is statistical. Rate limit accuracy increases when rate limits are set above 128 Mbps. (ID: 47506)

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)

Routing Profile

- The G8124 does not support the VMready or IGMP snooping features while the Routing deployment profile is used.

SNMP

- Port information displayed in MIBs related to port-based VLANs does not distinguish between a regular port or a trunk port. Use the `RS G8124(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)

Statistics

- Due to a hardware limitation, traffic that has no route to destination will be discarded by the switch, but this information will not be displayed in any statistics command. (ID:58975)

VLAG

- The following features are not supported on ports participating in VLAGs:
  - FCoE
  - Hotlinks
  - IGMP relay
  - Private VLANs
  - vNICs
- This known issue is applicable only to G8124-E.

  In a VLAG with PIM Dense Mode topology, if IGMP snooping is enabled on the Layer 2 VLAG switches, the Mrouter on the Layer 2 VLAG switches will experience continuous flapping. To avoid this issue, we recommend that you configure the aggregation layer VLAG switch as the IGMP Querier. (ID: 68717)

VMready

- The G8124 does not support the VMready feature while the Routing deployment profile is used.
- 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.

vNICs

- When using vNICs for iSCSI, the operation to clone a VM on an iSCSI disk may time out, leaving the VM uncopied.