

IBM System Networking RackSwitch™ G8264



Release Notes

For IBM Networking OS™ 7.11

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 (November 2014)

© Copyright IBM Corporation 2014

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Release Notes

This release supplement provide the latest information regarding IBM Networking OS 7.11 for the RackSwitch G8264 (referred to as G8264 throughout this document).

This supplement modifies and extends the following IBM N/OS documentation for use with N/OS 7.11:

- *IBM Networking OS 7.11 Application Guide*
- *IBM Networking OS 7.11 ISCLI Reference*
- *IBM RackSwitch G8264 Installation Guide*

The publications listed above are available from the IBM support website:

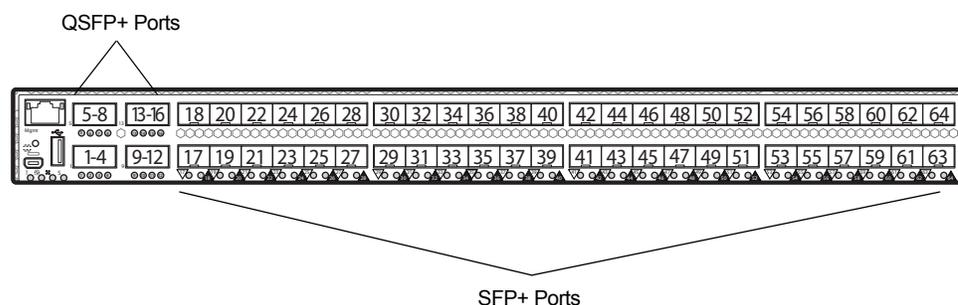
<http://www.ibm.com/support>

Hardware Support

The G8264 contains forty-eight 10GbE SFP+ ports and four 40GbE QSFP+ ports. The SFP+ ports can be populated with optical transceivers, copper transceivers, or Direct Attach Cables (DACs). The QSFP+ ports can be populated with optical QSFP+ transceivers or DACs.

Note: If a DAC is not programmed to meet MSA specifications (including length identifier), the switch disables the port and generates a syslog message indicating that the DAC is not approved.

Figure 1. RackSwitch G8264 Front Panel



Updating the Switch Software Image

The switch software image is the executable code running on the G8264. A version of the image comes pre-installed on the device. As new versions of the image are released, you can upgrade the software running on your switch. To get the latest version of software supported for your G8264, go to the following website:

<http://www.ibm.com/support>

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

```
RS8264# 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 SFTP, 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, see “[Loading New Software to Your Switch](#)” on page 16.



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 7.11, the following special conditions may apply, depending on the version of software currently installed on your switch. These conditions are cumulative: If updating from version 2.0 (for example), follow the recommendations in order, beginning with those for 2.0, and then continue with all that apply, such as for “3.0 and prior,” “4.0 and prior,” and so on.

Updating from BLADEOS 6.4 or Prior

After updating:

- The default for STP/PVST Protection mode is different compared to release 6.4 and prior. In release 6.6, STP/PVST Protection is disabled by default. After upgrading, review the STP settings and make any appropriate changes.
- The default for static route health check is different compared to release 6.4 and prior. In release 6.6, static route health check is disabled by default. After upgrading, review the static route health check settings and make any appropriate changes.

- The legacy FDP update rate has been deprecated in favor of independent hotlinks FDB updates in all switch configuration interfaces.

Interface	Old Commands	New Commands
Menu CLI	/cfg/l2/update <x>	/cfg/l2/hotlink/sndrate <x>
ISCLI	spanning-tree uplinkfast max-update-rate <x>	hotlinks fdb-update-rate <x>
BBI	Configure Layer 2 Uplink Fast Update Rate Dashboard Layer 2 Uplink Fast STP Uplink Fast Rate	Configure Layer 2 Hot Links FDB update rate Dashboard Layer 2 Hot Links FDB update rate

These changes are also reflected in the SNMP MIB.

After upgrading, review the hotlinks FDB settings and make any appropriate changes

- The CLI `BGPTOECMP` option has been deprecated.

Updating from BLADEOS 6.6 or Prior

After updating:

- The default mode for Spanning Tree is different compared to prior releases. The default mode is now PVRST. After upgrading, it is recommended that the administrator review the STP settings and make any appropriate changes.

Updating from IBM Networking 6.8 or Prior



CAUTION:

If the current software version on your switch is 6.8 or prior, first upgrade the switch software image to 6.9 and reset the switch. Then load the 7.2 boot image and software image.

Updating from IBM Networking OS 6.9 or Prior



CAUTION:

When you upgrade the switch software image, you must load the new boot image and the new software image before you reset the switch.

After updating:

- The default settings of SNMP community strings have changed. Check the new settings and reconfigure as appropriate.

Updating to IBM Networking OS 7.x or Later in a VLAG Environment

Use the following process for updating the software image and boot image for all G8264 switches configured with VLAG:

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

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

2. Use TFTP or FTP to copy the new software 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.
Once Switch 1 has rebooted, Switch 1 will take the secondary role.
5. Shut down all ports except the ISL ports and the health check port on Switch 2.
Note: Do not save this configuration.
6. Reload Switch 2. Switch 1 will reassume the VLAG primary switch role.
7. Once Switch 2 has reloaded, make sure Switch 1 has transitioned to VLAG primary and Switch 2 has transitioned to secondary.
8. Verify all the VLAG clients have converged using the following command:

```
RS8264(config)# show vlag information
```

Updating from IBM Networking OS 7.2 or Prior

After updating:

- The default time zone setting is different compared to release 7.2 and prior. In the prior releases, a default setting of US Pacific Time was used. In release 7.4 and above, no default is assumed. For switches that use the default US Pacific Time setting, after upgrading to release 7.4 or above it is recommended that the administrator review the configured time zone and make any appropriate changes. (ID: 60469)

Updating from IBM Networking OS 7.8 or Prior

VLAN implementation was changed in release 7.9.

In order to improve the overall configuration of VLANs, major functional changes have been made regarding configuration of VLAN mapping for tagged/trunk ports when creating VLANs, enabling tagging or trunk mode on a port, and adding or removing VLANs from a tagged/trunk port.

For those upgrading from a release prior to 7.9, examine the following material to determine how your network may be affected, and to make any appropriate configurations changes after the upgrade.

Managing Tagged Ports in the ISCLI

Table 1 describes the functional differences between VLAN-related ISCLI commands before and after N/OS 7.9.

Table 1. ISCLI VLAN-Related Commands Before and After N/OS 7.9

ISCLI Command	Prior to N/OS 7.9	Starting With N/OS 7.9
<code>vlan <VLAN ID range></code>	Creates regular VLANs with no member ports.	Creates regular VLANs and automatically adds all tagged/trunk ports that have the VLANs in their allowed VLAN ranges as members.
<code>switchport mode trunk</code>	The port inherits the access PVID/Native VLAN as the trunk PVID/Native VLAN, and the port is a member of this VLAN only. The port must be added manually to other configured VLANs.	Adds the port to all configured VLANs that are allowed for the port (default: all regular VLANs) and inherits the access PVID/Native VLAN as the trunk PVID/Native VLAN. Any VLAN created afterward is associated automatically with the port if the VLAN is in the port's allowed range.
<code>switchport trunk native vlan <VLAN ID></code>	Sets the port's PVID/Native VLAN to the given VLAN ID. The PVID is created if it doesn't exist, and the port is added as a member of its PVID. If the PVID is not present in the port's current allowed VLAN range, the PVID is automatically added to the range.	Same as earlier releases, except that the new PVID/Native VLAN must be present in the port's current allowed VLAN range. Otherwise, an error message is generated.
<code>switchport trunk allowed vlan [optional parameters]</code>	The port must be configured in trunk mode (<code>switchport mode trunk</code>) before issuing this command.	You can issue this command on either access or trunk mode, but it takes effect only when the port is in trunk mode.
<code>switchport trunk allowed vlan all</code>	Adds the port to all regular VLANs configured when the command is issued. The port is not added automatically to VLANs created afterward.	Adds the port to all regular VLANs, present and future.

Table 1. ISCLI VLAN-Related Commands Before and After N/OS 7.9 (continued)

ISCLI Command	Prior to N/OS 7.9	Starting With N/OS 7.9
<code>switchport trunk allowed vlan <VLAN ID range></code>	Sets a new allowed VLAN range for the port. Removes the port from the VLANs that are no longer present in the new range, and adds the port to the configured VLANs present in the new range. Non-existing VLANs in the range are created automatically and the port is added to them. An error message is generated if the new range contains reserved or internal VLANs.	Sets a new allowed VLAN range for the port. Removes the port from the VLANs that are no longer present in the new range, and adds the port to the configured VLANs present in the new range (skipping any reserved / internal VLANs). No new VLANs are created automatically unless the new range doesn't have any configured VLANs. In this case, the lowest-numbered VLAN from the range is created. The port is also added automatically to any VLAN created afterward.
<code>switchport trunk allowed vlan none</code>	Removes the port from all VLANs, and moves it to the default VLAN or to a private VLAN if the port is in private VLAN mode and mapped to a private VLAN. This command appears in the configuration output <code>display/file</code> only when the port is in private VLAN mode and mapped to a Private VLAN.	Removes the port from all VLANs, and moves it to the default VLAN. This command doesn't appear in the configuration output <code>display/file</code> .
<code>switchport trunk allowed vlan add <VLAN ID range></code>	Adds the VLAN ID range to the port's current allowed VLAN list, and adds the port to the configured VLANs in the range. Creates the non-existing VLANs in the range, and adds the port to them. An error is generated if the range contains reserved or internal VLANs.	Adds the VLAN ID range to the port's current allowed VLAN list, and adds the port to the configured VLANs in the range. You can specify non-existing VLANs in the range, but they are not created automatically. The range may also include reserved or internal VLANs, but the port is not added to them.

Table 1. ISCLI VLAN-Related Commands Before and After N/OS 7.9 (continued)

ISCLI Command	Prior to N/OS 7.9	Starting With N/OS 7.9
<code>switchport trunk allowed vlan remove <VLAN ID range></code>	Removes the VLAN ID range from the port's allowed VLAN list, and disassociates the port from those VLANs. If the port's PVID/ Native VLAN is included in the range, the lowest-numbered VLAN in the remaining allowed range becomes the new PVID/Native VLAN. If all VLANs in the current allowed range are removed, the port is placed in the default VLAN.	Removes the VLAN ID range from the port's allowed VLAN list, and disassociates the port from those VLANs. If the port's PVID/ Native VLAN is included in the range, the lowest-numbered VLAN in the remaining allowed range becomes the new PVID/Native VLAN. If there are no configured VLANs in the remaining allowed range, the lowest-numbered regular VLAN is created automatically. If all VLANs in the current allowed range are removed, an error message is generated.
<code>no switchport trunk allowed vlan</code>	Removes the port from all VLANs, and moves it to the default VLAN.	Sets the default allowed VLAN range for the port (i.e. the port is added to all regular VLANs).

Managing Tagged Ports in the BBI and SNMP

In releases prior to N/OS 7.9, enabling trunk mode or tagging on a port in BBI/SNMP will add the port to the PVID/ Native VLAN only. You have to manually add the tagged port to the other VLANs.

Starting with N/OS 7.9, enabling trunk mode or tagging on a port will add the port to all configured VLANs and other VLANs created afterward. You will have to remove any undesired VLANs from the port, because there is no equivalent BBI/SNMP operation for the `switchport trunk allowed vlan` command in ISCLI.

Tagged Ports in Configuration Outputs

In releases prior to N/OS 7.9, the configuration dump/file will show only the VLANs in which a tagged/trunk port is a member.

Starting with N/OS 7.9, the configuration dump/file will show the VLANs configured for the port to be associated with, and that may include configured VLANs, non-existing VLANs, internal or reserved VLANs. The actual and operational VLAN and port associations are shown by the `show vlan` and `show interface information` commands in ISCLI (or its equivalent in BBI or SNMP.)

Tagged Ports in QBG VLANs

In releases prior to N/OS 7.9, the `switchport trunk allowed <range>` command converts a dynamic QBG VLAN to static type

Starting with N/OS 7.9, the conversion is done when the range contains QBG VLANs only.

Stacking Ports

When issuing tagged/trunk port-related commands on stacking ports in N/OS 7.9 or later the configuration is either ignored or an error message is generated, because stacking ports should always be a member of all VLANs.

Tagged Ports Configuration Scenario

[Table 2](#) illustrates the differences between N/OS 7.9 and previous releases when configuring VLANs and associating VLANs with tagged/trunk ports. Some command outputs in the table were edited for brevity. The port numbers may not accurately reflect the actual port numbering in some switches.

Note: The same ISCLI command such as `switchport mode trunk` applied on two switches in a network, with one switch running NOS 7.9 or later and the other running an older release, may result in mismatched VLAN configurations between the ports connecting the two switches. This may lead to problems such as loss of traffic and connectivity.

Table 2. VLAN Tagging Configuration Scenario Before and After N/OS 7.9

Prior to N/OS 7.9	Starting with N/OS 7.9
<i>Initial factory configuration:</i>	<i>Initial factory configuration:</i>
<pre>RS8264(config)#show running-config Current configuration: ! version "7.8" ... ! end</pre>	<pre>RS8264(config)#show running-config Current configuration: ! version "7.9" ... ! end</pre>
<pre>RS8264(config)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 1 ...</pre>	<pre>RS8264(config)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 1 ...</pre>
<pre>RS8264(config)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 1 ... VLANs:1</pre>	<pre>RS8264(config)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 1 ... VLANs: 1</pre>

Table 2. VLAN Tagging Configuration Scenario Before and After N/OS 7.9 (continued)

Prior to N/OS 7.9	Starting with N/OS 7.9
<p>Create VLANs 100 and 200, and enable trunk mode on port 1:</p>	<p>Create VLANs 100 and 200, and enable trunk mode on port 1:</p>
<pre>RS8264(config)#vlan 100 Warning: VLAN 100 was assigned to STG 100. VLAN 100 is created.</pre>	<pre>RS8264(config)#vlan 100 Warning: VLAN 100 was assigned to STG 100. VLAN 100 is created.</pre>
<pre>RS8264(config-vlan)#vlan 200 Warning: VLAN 200 was assigned to STG 73. VLAN 200 is created.</pre>	<pre>RS8264(config-vlan)#vlan 200 Warning: VLAN 200 was assigned to STG 73. VLAN 200 is created.</pre>
<pre>RS8264(config-vlan)#interface port 1 RS8264(config-if)#switchport mode trunk</pre>	<pre>RS8264(config-vlan)#interface port 1 RS8264(config-if)#switchport mode trunk</pre>
<pre>RS8264(config-if)#show running-config Current configuration: ! ... interface port 1 switchport mode trunk switchport trunk allowed vlan 1 exit ! vlan 100 name "VLAN 100" ! vlan 200 name "VLAN 200" ... ! end</pre>	<pre>RS8264(config-if)#show running-config Current configuration: ! ... interface port 1 switchport mode trunk exit ! vlan 100 name "VLAN 100" ! vlan 200 name "VLAN 200" ... ! end</pre>
<pre>RS8264(config-if)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 1 ... 100 VLAN 100 ena empty 200 VLAN 200 ena empty</pre>	<pre>RS8264(config-if)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 1 ... 100 VLAN 100 ena 1 200 VLAN 200 ena 1</pre>
<pre>RS8264(config-if)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 1, Tagging/Trunk-mode ... VLANs:1</pre>	<pre>RS8264(config-if)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 1, Tagging/Trunk-mode ... VLANs: 1,100,200</pre>

Table 2. VLAN Tagging Configuration Scenario Before and After N/OS 7.9 (continued)

Prior to N/OS 7.9	Starting with N/OS 7.9
<p>Create VLAN 300:</p>	<p>Create VLAN 300:</p>
<pre>RS8264(config)#vlan 300 Warning: VLAN 300 was assigned to STG 46. VLAN 300 is created.</pre>	<pre>RS8264(config)#vlan 300 Warning: VLAN 300 was assigned to STG 46. VLAN 300 is created.</pre>
<pre>RS8264(config-vlan)#show running-config Current configuration: ! ... interface port 1 switchport mode trunk switchport trunk allowed vlan 1 exit ! vlan 100 name "VLAN 100" ! vlan 200 name "VLAN 200" ! vlan 300 name "VLAN 300" ... ! end</pre>	<pre>RS8264(config-vlan)#show running-config Current configuration: ! ... interface port 1 switchport mode trunk exit ! vlan 100 name "VLAN 100" ! vlan 200 name "VLAN 200" ! vlan 300 name "VLAN 300" ... ! end</pre>
<pre>RS8264(config-vlan)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 1 ... 100 VLAN 100 ena empty 200 VLAN 200 ena empty 300 VLAN 300 ena empty</pre>	<pre>RS8264(config-vlan)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 1 ... 100 VLAN 100 ena 1 200 VLAN 200 ena 1 300 VLAN 300 ena 1</pre>
<pre>RS8264(config-vlan)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 1, Tagging/Trunk-mode ... VLANs:1</pre>	<pre>RS8264(config-vlan)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 1, Tagging/Trunk-mode ... VLANs: 1,100,200,300</pre>

Table 2. VLAN Tagging Configuration Scenario Before and After N/OS 7.9 (continued)

Prior to N/OS 7.9	Starting with N/OS 7.9
<i>Add VLAN 300 to port 1:</i>	<i>Add VLAN 300 to port 1:</i>
<pre>RS8264(config)#interface port 1 RS8264(config-if)#switchport trunk allowed vlan add 300</pre>	<pre>RS8264(config)#interface port 1 RS8264(config-if)#switchport trunk allowed vlan add 300</pre>
<pre>RS8264(config-if)#show running-config Current configuration: ! ... interface port 1 switchport mode trunk switchport trunk allowed vlan 1,300 exit ! vlan 100 name "VLAN 100" ! vlan 200 name "VLAN 200" ! vlan 300 name "VLAN 300" ... ! end</pre>	<pre>RS8264(config-if)#show running-config Current configuration: ! ... interface port 1 switchport mode trunk exit ! vlan 100 name "VLAN 100" ! vlan 200 name "VLAN 200" ! vlan 300 name "VLAN 300" ... ! end</pre>
<pre>RS8264(config-if)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 1 ... 100 VLAN 100 ena empty 200 VLAN 200 ena empty 300 VLAN 300 ena 1</pre>	<pre>RS8264(config-if)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 1 ... 100 VLAN 100 ena 1 200 VLAN 200 ena 1 300 VLAN 300 ena 1</pre>
<pre>RS8264(config-if)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 1, Tagging/Trunk-mode ... VLANs:1 300</pre>	<pre>RS8264(config-if)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 1, Tagging/Trunk-mode ... VLANs: 1,100,200,300</pre>

Table 2. VLAN Tagging Configuration Scenario Before and After N/OS 7.9 (continued)

Prior to N/OS 7.9	Starting with N/OS 7.9
<i>Remove VLAN 1 from port 1:</i>	
<pre>RS8264(config)#interface port 1 RS8264(config-if)#switchport trunk allowed vlan remove 1 RS8264(config-if)#show running-config Current configuration: ! ... interface port 1 switchport mode trunk switchport trunk allowed vlan 300 switchport trunk native vlan 300 exit ! vlan 100 name "VLAN 100" ! vlan 200 name "VLAN 200" ! vlan 300 name "VLAN 300" ... ! end RS8264(config-if)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 2 ... 100 VLAN 100 ena empty 200 VLAN 200 ena empty 300 VLAN 300 ena 1 RS8264(config-if)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 300, Tagging/Trunk-mode ... VLANs:300</pre>	<pre>RS8264(config)#interface port 1 RS8264(config-if)#switchport trunk allowed vlan remove 1 RS8264(config-if)#show running-config Current configuration: ! ... interface port 1 switchport mode trunk switchport trunk allowed vlan 2-4095 switchport trunk native vlan 100 exit ! vlan 100 name "VLAN 100" ! vlan 200 name "VLAN 200" ! vlan 300 name "VLAN 300" ... ! end RS8264(config-if)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 2 ... 100 VLAN 100 ena 1 200 VLAN 200 ena 1 300 VLAN 300 ena 1 RS8264(config-if)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 100, Tagging/Trunk-mode ... VLANs: 100,200,300</pre>

Table 2. VLAN Tagging Configuration Scenario Before and After N/OS 7.9 (continued)

Prior to N/OS 7.9	Starting with N/OS 7.9
<i>Restore default VLAN membership for port 1:</i>	
<pre>RS8264(config)#interface port 1 RS8264(config-if)#no switchport trunk allowed vlan Every port has to be a member of at least one VLAN, ports will be added to default VLAN.</pre>	<pre>RS8264(config)#interface port 1 RS8264(config-if)#no switchport trunk allowed vlan</pre>
<pre>RS8264(config-if)#show running-config Current configuration: ! ... interface port 1 switchport mode trunk switchport trunk allowed vlan 1 exit ! vlan 100 name "VLAN 100" ! vlan 200 name "VLAN 200" ! vlan 300 name "VLAN 300" ... ! end</pre>	<pre>RS8264(config-if)#show running-config Current configuration: ! ... interface port 1 switchport mode trunk switchport trunk native vlan 100 exit ! vlan 100 name "VLAN 100" ! vlan 200 name "VLAN 200" ! vlan 300 name "VLAN 300" ! ... end</pre>
<pre>(config-if)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 1 ... 100 VLAN 100 ena empty 200 VLAN 200 ena empty 300 VLAN 300 ena empty</pre>	<pre>(config-if)#show vlan VLAN Name Status Ports ----- 1 Default VLAN ena 1 ... 100 VLAN 100 ena 1 200 VLAN 200 ena 1 300 VLAN 300 ena 1</pre>
<pre>RS8264(config-if)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 1, Tagging/Trunk-mode ... VLANs:1</pre>	<pre>RS8264(config-if)#show interface port 1 Current port 1 configuration: enabled, PVID/Native-VLAN 100, Tagging/Trunk-mode ... VLANs: 1,100,200,300</pre>

Loading New Software to Your Switch

The G8264 can store up to two different switch software images (called `image1` and `image2`) as well as special boot software (called `boot`). When you load new software, you must specify where it should be placed: either into `image1`, `image2`, or `boot`.

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



CAUTION:

When you upgrade the switch software image, always load the new boot image and the new software image before you reset the switch. If you do not load a new boot image, your switch might not boot properly (To recover, see [“Recovering from a Failed Upgrade” on page 22](#)).

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

- The image and boot software loaded on an SFTP, 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 SFTP, 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, 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 SFTP, FTP or TFTP directory (for example, `tftpboot`).

4. If required by the SFTP, 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.

- 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.
```

- 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

- Enter the following Boot Options command:

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

- Enter the name of the switch software to be replaced:

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

- 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>
```

- 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).

- 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 G8264. The software image to load can reside in one of the following locations:

- FTP server
- TFTP server
- Local computer

After you log onto the BBI, perform the following steps to load a software image:

1. Click the Configure context tab in the toolbar.
2. In the Navigation Window, select System > Config/Image Control.
The Switch Image and Configuration Management page appears.
3. If you are loading software from your computer (HTTP client), skip this step and go to the next. Otherwise, if you are loading software from an SFTP, 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 SFTP, 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 7.11 for RackSwitch G8264 (G8264) has been updated to include several new features, summarized in the following sections. For more detailed information about configuring G8264 features and capabilities, refer to the complete N/OS 7.11 documentation as listed on [page 3](#).

ARP Information Includes VLAN 4095

Information for the management VLAN (VLAN 4095) is now included in the show ip arp command output:

```
RS8264(config)# show ip arp
-----
Mgmt ARP entries:

Total number of Mgmt arp entries : 2
  IP address  Flags  MAC address  VLAN  Age Port
-----
  10.100.121.132  P    08:17:f4:62:64:fe  4095  MGT
  10.100.121.1    00:22:00:ad:45:00  4095  1 MGT

Data ARP entries:

Current ARP configuration:
  rearp 5

Current static ARP:
  IP address  MAC address  Port  VLAN
-----
  10.10.10.42  03:bf:0a:0a:0a:2a  42

Total number of arp entries : 2
  IP address  Flags  MAC address  VLAN  Age Port
-----
  10.10.10.1  P    08:17:f4:62:64:00  42
  10.10.10.42  P    03:bf:0a:0a:0a:2a  42
```

BGP Supports Prepending up to 32 AS Numbers

The BGP feature has been extended in N/OS 7.11. You can now specify up to 32 AS numbers (increased from 3) that you want to prepend to a matched route.

```
RS8264(config-route-map)# as-path-preference <AS number list (up to 32 entries)>
```

To specify the AS number list, place one space between each of the entries. For example:

```
RS8264(config-route-map)# as-path-preference 2 3 4 5 20 22
```

OpenFlow Supports Static LACP

When the G8264 is configured to use version 1.3 of the OpenFlow standard, port trunk groups can be added to OpenFlow instances. A trunk aggregates its member ports to form a logical port with increased bandwidth. You can add an existing static trunk group (portchannel) or static LACP trunk group to an OpenFlow instance using the following commands:

```
RS8264(config)# openflow instance <instance ID>  
RS8264(config-openflow-instance)# member portchannel <trunk ID>
```

where the *trunk ID* (the logical port ID) is derived from the original trunk configuration, based on the trunk type:

- Static trunk group (portchannel)

```
RS8264(config)# portchannel <trunk ID> port <port list>
```

- Static LACP trunk group

```
RS8264(config)# portchannel <trunk ID> lacp key <LACP admin key>
```

Once added to the instance, the trunk ports inherit the OpenFlow data properties such as MAC learning turned off, flood blocking turned on, and STP disabled.

The trunk link remains active as long as at least one member port is up. The trunk link speed is an aggregation of the speed of the individual member ports. If any port in the trunk goes down, the overall trunk link speed is decreased accordingly.

To add a static trunk group or static LACP trunk group to the edge ports list, use the following command:

```
RS8264(config-openflow-instance)# edgeport portchannel <trunk ID>
```

Stacking Master Access to the Member Switch CLI

For easier administration, you can use the Master switch to access the CLI of any Member switch. Use this feature to display specific Member switch information such as log messages, system information, and statistics.

Member CLI access is available regardless of whether the Member switch is operating in the attached state (prior to joining the stack), or in the active state (after successfully joining the stack). Member CLI access is available even without an assigned IP address or a physical connection to the serial console on the Member.

To use this feature, log in to the Master switch (via its local serial console or through a remote Telnet or SSH connection) and issue the following command:

```
RS8264# connect <asnum>
```

where *asnum* is the attached switch number of the target Member switch.

For example, to access the Member switch identified by attached switch number 3, enter the following command:

```
RS8264#connect 3
Connecting via DATA port.
Connected to 127.117.210.5.
Escape character is '^]'.

IBM Networking Operating System RackSwitch G8264, Stack.

Enter login username:
```

Log in to the Member switch as usual. If the Member switch is in the active state (after successfully joining the stack), the prompt will include the Member's configured switch number (`csnum`), which may be different from the assigned switch number (`asnum`):

```
RS8264-3>
```

To return to the Master switch prompt, log out of the Member switch.

Supplemental Information

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

The Boot Management Menu

The Boot Management menu allows you to switch the software image, reset the switch to factory defaults, or to recover from a failed software download.

You can interrupt the boot process and enter the Boot Management menu from the serial console port. When the system displays Memory Test, press **<Shift B>**. The Boot Management menu appears.

```
Resetting the System ...
Memory Test .....

Boot Management Menu
1 - Change booting image
2 - Change configuration block
3 - Xmodem download
4 - Exit

Please choose your menu option: 1
Current boot image is 1. Enter image to boot: 1 or 2: 2
Booting from image 2
```

The Boot Management menu allows you to perform the following actions:

- To change the booting image, press 1 and follow the screen prompts.
- To change the configuration block, press 2, and follow the screen prompts.
- To perform an Xmodem download, press 3 and follow the screen prompts.
- To exit the Boot Management menu, press 4. The booting process continues.

Recovering from a Failed Upgrade

Use the following procedure to recover from a failed software upgrade.

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports XModem Download (for example, HyperTerminal, SecureCRT, PuTTY) and select the following serial port characteristics:
 - Speed: 9600 bps
 - Data Bits: 8
 - Stop Bits: 1
 - Parity: None
 - Flow Control: None
3. Boot the switch and access the Boot Management menu by pressing **<Shift B>** while the Memory Test is in progress and the dots are being displayed.
4. Select **3** for **Xmodem download**. When you see the following message, change the Serial Port characteristics to 115200 bps:

```
## Switch baudrate to 115200 bps and press ENTER ...
```

5. Press **<Enter>** to set the system into download accept mode. When the readiness meter displays (a series of “C” characters), start XModem on your terminal emulator.
6. Select the Boot Image to download. The XModem initiates the file transfer. When the download is complete, a message similar to the following is displayed:

```

yzModem - CRC mode, 62494(SOH)/0(STX)/0(CAN) packets, 6 retries
Extracting images ... Do *NOT* power cycle the switch.

**** VMLINUX ****

Un-Protected 10 sectors
Erasing Flash..... done
Writing to Flash.....done
Protected 10 sectors

**** RAMDISK ****

Un-Protected 44 sectors
Erasing Flash..... done
Writing to Flash.....done
Protected 44 sectors

**** BOOT CODE ****

Un-Protected 8 sectors
Erasing Flash..... done
writing to Flash.....done
Protected 8 sectors

```

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

```
## Switch baudrate to 9600 bps and press ESC ...
```

8. Press the Escape key (**<Esc>**) to re-display the Boot Management menu.
9. Select **3** to start a new **XModem Download**. When you see the following message, change the Serial Port characteristics to 115200 bps:

```
## Switch baudrate to 115200 bps and press ENTER ...
```

10. Press **<Enter>** to continue the download.

11. Select the OS Image to download. The XModem initiates the file transfer. When the download is complete, a message similar to the following is displayed:

```
yzModem - CRC mode, 27186(SOH)/0(STX)/0(CAN) packets, 6 retries  
  
Extracting images ... Do *NOT* power cycle the switch.  
  
**** Switch OS ****  
  
Please choose the Switch OS Image to upgrade [1|2|n] :
```

12. Select the image number to load the new image (1 or 2). It is recommended that you select 1. A message similar to the following is displayed:

```
Switch OS Image 1 ...  
  
Un-Protected 27 sectors  
  
Erasing Flash..... done  
  
Writing to Flash.....done  
  
Protected 27 sectors
```

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

```
## Switch baudrate to 9600 bps and press ESC ...
```

14. Press the Escape key (<Esc>) to re-display the Boot Management menu.
15. Select **4** to exit and boot the new image.

VLAGs

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

- Any port-related configuration, such as applied ACLs, 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:
RS8264 (config)# **no vlag adminkey** <key> **enable** (or)
RS8264 (config)# **no portchannel** <number> **enable**
3. Change the configuration as needed.

On the VLAG Primary Peer:

4. Disable the VLAG instance.
5. Change the configuration as needed.
6. Enable the VLAG instance.

On the VLAG Secondary Peer:

7. Enable the VLAG instance.
8. Enable the VLAG ports.

Note: This is not required on non-VLAG ports or when STP is off or set for PVRST mode.

Known Issues

This section describes known issues for N/OS 7.11 on the G8264.

ACLs

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

BGP

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

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.

Debug

- IBM N/OS debug commands are for advanced users. Use the debug commands with caution as they can disrupt the operation of the switch under high load conditions. This could be dangerous in mission-critical production environments. Before debugging, check the MP utilization to verify there is sufficient overhead available for the debug functionality. When debug is running under high load conditions, the CLI prompt may appear unresponsive. In most cases, control can be returned by issuing a `no debug <function>` command.

EVB

- When a VM cannot be associated, the console may be flooded with syslog messages stating that the validation has failed. (ID: XB191291)
- Due to a hardware limitation, traffic received by a VM may not conform to the RxRate (receive rate) that you have configured. (ID: 55600)

FCoE

- The FCoE connection between the server and the FCF will be retained even if you disable CEE/FIP/vNIC on the switch. To avoid this scenario, either reboot the switch, or disable and re-enable the ports connected to the sever and the FCF after you disable CEE/FIP/vNIC. (ID:41915)
- By default the "VLAN Name" and "Port and Protocol ID" LLDP TLVs are disabled on a port. These two TLVs are added to the LLDP PDU for each VLAN that is configured in a port. This may cause the length of LLD PDU to exceed the Ethernet packet size if there are nearly 40 or more VLANs configured on a port, or if the VLAN names are too long. There is a possibility that the DCBX TLVs may not be added to the LLDP TLV due to the length. Because of this the FCoE connection will not form on that port. It is recommended to avoid enabling the "VLAN Name" and "Port and Protocol ID" TLV if you have high number of VLANs configured and FCoE is enabled on that port. (ID: 42446)
- The FIP Snooping option for automatic VLAN creation is not recommended for use with the Emulex Virtual Fabric Adapter. Disable automatic VLAN creation when connecting the switch to an Emulex Virtual Fabric Adapter. (ID: 51529)
- When using DCBX to synchronize Priority-base Flow Control (PFC) with a peer (using the PFC TLV option), PFC operation on the switch port will not be disabled as expected when PFC is not available on the peer. To resolve this, manually disable PFC on ports connected to devices that do not participate in PFC. (ID: 62114)
- Disruption to FCoE connections and FCoE traffic may be expected when changing the LACP mode. It is recommended that the administrator halt FCoE traffic before changing any switch configuration. (ID: 67044)
- It is recommended to use the FIP snooping automatic VLAN creation option in FCOE environments, in addition to configuring VLANs manually. The auto-VLAN feature should be disabled only if no additional FCF or ENode ports will be automatically added to the FCOE VLAN. Otherwise, some FCF or ENode ports might not be automatically added to the FCOE VLAN, even if the auto-VLAN feature is later enabled, requiring them to be added manually.
- The FIPS auto-VLAN feature is disabled by default. (ID: XB258382)
- FCoE connections may become temporarily unstable whenever a change occurs to the VLAG virtual port. (ID: XB263734)
- In an event in which connections for multiple ports becomes unstable, FCoE traffic may pause or be dropped as FCoE FDB entries are cleared and reinstated. (ID: XB275415)

FDB

- From IBM Networking OS 7.11 onwards, MAC address information is no longer learned by control packets such as LACPDUs. This behavior is as expected. (ID: XB253517)

IKEv2

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

IPsec

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

ISCLI

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

LACP

- Under heavy switch load conditions, LACP links may flap when configured with short timeout mode. To stabilize LACP under heavy load, it is recommended to use the long timeout mode instead. (ID: 66173)

On-Box Scripting

- 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 `ibmppylib.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

- Some changes to OSPF configuration (such as creating a new area or changing an area's type) may result in OSPF state reconvergence. (ID: 46445, 48483)
- OSPFv3 over IPsec:
 - This combination can only be configured only on a per-interface basis.
 - The current implementation for OSPFv3 allows the use of only one protocol (AH or ESP) at any given time. AH and ESP cannot be applied together.
 - IPsec does not support OSPFv3 virtual links. (ID: 48914)

Ports and Transceivers

- The port speed setting is not configurable for Finisar SFPs. Updating from BLADEOS 6.5 (or prior) to N/OS 6.8 (or later) will result in port speed configuration settings being reset to default values for ports using Finisar SFPs. (ID: 55063)
- In stacking mode, two ports of different link speeds can exist in the same portchannel. This may lead to loss of traffic. (ID: XB278986)

Private VLANs

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

QoS

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

```
RS8264(config)# show interface port <swunit:port_num>
egress-mcast-queue-counters
```

For example:

```
RS8264(config)# show interface port 1:24 egress-mcast-queue-counters

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

QSFP+

- The QSFP+ ports do not auto-negotiate. The desired speed must be configured to match on both ends of the connection, and the switch reset for changes to take effect. (ID: 46340)
- After you upgrade switch software and reset the switch, you must configure the QSFP+ port mode. Use the following command (ID: 46858):
boot qsfp-40gports <1, 5, 9, 13>
- When changing a QSFP port from 10G mode to 40G mode, a port error will occur if any previously configured 10G port settings do not apply to the new 40G state, preventing further configuration of the port. The administrator must manually clear the 10G port settings that do not apply to 40G prior to changing modes. (ID: 62576)

Routed Ports

- IBM N/OS CLI, SNMP, or BBI should not be used to configure routed ports, or to configure any other feature if a routed port is already configured on the switch. If a routed port is configured on the switch, the configuration, apply, and save commands are not displayed in IBM N/OS CLI or BBI; in SNMP, you may be able to enter the configuration commands, but you will not be able to save the configuration. (ID: 57983)

sFlow

- In some cases, sFlow configured with the minimum polling and sampling rate could cause the switch to get into a hang state with no traffic passing after about 7 days of operations with large volumes of traffic. Please contact Customer Support or the System Engineer before enabling sFlow. (ID: 57045)

SNMP

- Port information displayed in MIBs related to port-based VLANs does not distinguish between a regular port or a trunk port. Use the `show mac-address-table static` command to view details on regular ports and trunk ports. (ID: 57194)
- If you delete multiple VLANs using SNMP, you may see an error if the SNMP packet size exceeds 1800 bytes. (ID: XB228120)

Stacking

- When stacking is enabled, the switch may continue to learn MAC addresses from ports or trunks even though they are in a blocking state. The unexpected MAC addresses represent control packets, not endpoint devices, and do not impact switch performance. (ID:61996)
- Stacks using N/OS 7.6 will not accept new members that use any other version N/OS, including N/OS 7.11. Upload N/OS 7.6 (not N/OS 7.11) on the individual switch prior to joining it to a stack that uses N/OS 7.6. (ID: 69744)
- Switches using N/OS 7.6 will not automatically join an existing stack that is loaded with any other version of N/OS, including N/OS 7.11. Upload the N/OS 7.6 switch with N/OS 7.11 prior to joining it to the existing N/OS 7.11 stack. (ID: 69744; ID: 70387)
- During failover/failback operation of the stack Master, there may be some traffic loss for up to 40 seconds. (ID: 70568)
- DHCP has higher priority over static management IP configuration. If you want to configure a static management IP, you must first disable DHCP using the command: `RS8264(config)# no system dhcp mac <MAC address>`. (ID: 71181)
- When you issue the `show mac-address-table counters` command, the output is displayed after a delay of approximately 20 seconds. In the mean time, if there are any updates to the table, the output of the command may have incorrect information. (ID: XB218728)

STP

- 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)

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)

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.
- 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).
- 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)

vNICs

- When using vNICs with FCoE, the FIP Snooping option for automatic VLAN creation is not recommended for use with the Emulex Virtual Fabric Adapter. Disable automatic VLAN creation when connecting the switch to an Emulex Virtual Fabric Adapter. (ID: 51529)
- vNIC egress bandwidth control is not strictly enforced on the switch for packets larger than 900 bytes, resulting in greater egress bandwidth from the switch to the server than is configured. However, ingress bandwidth control (from the server to the switch) is strictly enforced. (ID: 50950)

