

Lenovo RackSwitch G8264

# Release Notes

for Networking OS 8.3

**Lenovo**<sup>TM</sup>

**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 *Lenovo Documentation CD* and the *Warranty Information* document that comes with the product.

First Edition (August 2015)

© Copyright Lenovo 2015  
Portions © Copyright IBM Corporation 2014.

**LIMITED AND RESTRICTED RIGHTS NOTICE:** If data or software is delivered pursuant a General Services Administration "GSA" contract, use, reproduction, or disclosure is subject to restrictions set forth in Contract No. GS-35F-05925.

Lenovo and the Lenovo logo are trademarks of Lenovo in the United States, other countries, or both.

---

## Release Notes

This release supplement provides the latest information regarding Lenovo Networking OS 8.3 for the Lenovo RackSwitch G8264 (referred to as G8264 throughout this document).

This supplement modifies and extends the following Lenovo N/OS documentation for use with N/OS 8.3:

- *Lenovo RackSwitch G8264 Application Guide for Networking OS 8.3*
- *Lenovo RackSwitch G8264 ISCLI Command Reference for Networking OS 8.3*
- *Lenovo RackSwitch G8264 Installation Guide*

The publications listed above are available from the following website:

<http://publib.boulder.ibm.com/infocenter/systemx/documentation/index.jsp>

Please keep these release notes with your product manuals.


# Hardware Support

N/OS 8.3 software is supported on the G8264, a high performance Layer 2-3 network switch.

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

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

**Figure 1.** RackSwitch G8264 front panel:



---

## 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.support.lenovo.com/>

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 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 6](#).

**Note:** Due to Lenovo firmware security guidelines, it is not possible at this time to upgrade or migrate a switch from IBM Networking OS to Lenovo Networking OS.

## Updating VLAG Switches with Lenovo Networking OS 8.x

Following are the steps for updating the software image and boot image for 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 OS image and boot image onto both vLAG switches.
3. Shut down all ports except the ISL ports and the health check port on the primary switch (Switch 1).

**Note:** Do not save this configuration.

4. Reload Switch 1, Switch 2 will assume the vLAG primary role
5. Once Switch 1 has rebooted, Switch 1 will take the secondary role.
6. Shut down all ports except the ISL ports and the health check port on Switch 2.

**Note:** Do not save this configuration.

7. Reload Switch 2, Switch 1 will reassume the vLAG primary switch role.

8. Once Switch 2 has reloaded, make sure Switch 1 has transitioned to vLAG primary and Switch 2 has transitioned to secondary.
9. Verify all the vLAG clients have converged using the following command:

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

## 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.

**Attention:** 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 Software Upgrade” on page 10](#)).

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 the following procedures to download the new software to your switch.

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

```
RS8264(config)# copy {sftp|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:

```
RS8264# configure terminal  
RS8264(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:

```
RS8264(config)# reload
```

The system prompts you to confirm your request. Once confirmed, the switch will reboot to use the new software.

**Note:** If you select “No” when asked to confirm the reload, any changes made to the configuration since the last reboot will be lost.

---

## New and Updated Features

N/OS 8.3 for RackSwitch 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 8.3 documentation as listed on [page 3](#).

### BGP Pre Append AS path

This enhancement facilitates the process of routing maps by influencing the BGP route selection so that one path is preferred over another. This is especially useful when customers are dual home to two different Internet Service Providers (ISPs) and they want to have a primary path to one ISP and use the other ISP as a backup path.

### Dynamic ARP Inspection

This enhancement enables the device to intercept and examine all ARP requests and response packets in a subnet. The feature discards those packets with invalid IP to MAC address bindings. This capability protects the network from some man-in-the-middle attacks.

### EasyConnect

This feature is designed to simplify switch configuration by applying predefined configuration modes. Once launched, the user is requested a minimal set of input information (e.g. hostname, netmask, server and uplink ports or VLAG) and the tool is automatically customizing the switch accordingly.

### Increase MAX ARPs to 16K

This is an enhancement to support 16K ARP entries on the platforms using Trident+ and T2 chipsets.

### OpenSSL RC4 Vulnerability

This release disables the usage of the RC4 cipher in the VMWare client to avoid security vulnerabilities.

### Python Off-Box API

NOS python API (NPAPI) enables administrators to manage device with Off-Box Python program.

MSSC Plug-in support only.

### Python On-Box Scripting

The release supports the monitoring of additional VLAG/BGP/VRRP/IGMP events and L3 APIs including BGP, IGMP and VRRP.



## **SNSC**

This release adds the option to disable terminal prompting in SNSC.

## **Spanning Tree Protocol (STP) Range Enhancement**

Existing Spanning Tree Protocol (STP) commands now support configuration of a range of STP groups at once. STP range groups configurations are also displayed in a configuration dump.

## **Service Location Protocol**

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

## **VRRP**

This enhancement allows VRRP to work in two ways under vLAG topology:

- Full Active-Active: both vLAGs perform L3 traffic routing for the related VRRP domain.
- Half Active-Active: one vLAG performs L3 traffic routing while the second one manages L2 forwarding for the related VRRP domain.

---

## 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 at 10:55:49 Sun Apr 19, 2015...
Memory Test .....
Boot Management Menu
I - Change booting image
C - Change configuration block
R - Boot in recovery mode (tftp and xmodem download of images to recover
switch)
Q - Reboot
E - Exit
Please choose your menu option: I
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 **I** and follow the screen prompts.
- To change the configuration block, press **C**, and follow the screen prompts.
- To perform a TFTP/Xmodem download, press **R** and follow the screen prompts.
- To reboot the switch, press **Q**. The booting process restarts
- To exit the Boot Management menu, press **E**. The booting process continues.

### *Recovering from a Failed Software 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. To access the Boot Management menu, you must interrupt the boot process from the Console port. Boot the G8264, and when the system begins displaying Memory Test progress (a series of dots), press **<Shift B>**. The Boot Management menu will display:

```
Resetting at 10:55:49 Sun Apr 19, 2015...
Memory Test .....
Boot Management Menu
I - Change booting image
C - Change configuration block
R - Boot in recovery mode (tftp and xmodem download of images to recover
switch)
Q - Reboot
E - Exit
Please choose your menu option:
```

4. Select **R** for **Boot in recovery mode**. You will see the following display:

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

Option?:
```

- If you choose option **X** (Xmodem serial download), go to step 5.
- If you choose option **T** (TFTP download), go to step 6.

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

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

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

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

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

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

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

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

Option?:
```

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

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

- a. Enter the required information and press **<Enter>**.
- b. You will see a display similar to the following:

```
Host IP    : 10.10.98.110
Server IP  : 10.10.98.100
Netmask    : 255.255.255.0
Broadcast  : 10.10.98.255
Gateway    : 10.10.98.254
Installing image G8264-8.2.1.0_OS.img from TFTP server 10.10.98.100
```

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

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

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

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

Option?:
```

7. Image recovery is complete. Perform one of the following steps:

- Press **r** to reboot the switch.
- Press **e** to exit the Boot Management menu
- Press the Escape key (**<Esc>**) to re-display the Boot Management menu.

## Recovering from a Failed Boot Image Upgrade

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

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports Xmodem download (such as HyperTerminal, CRT, or PuTTY) and select the following serial port characteristics:
  - o Speed: 9600 bps
  - o Data Bits: 8
  - o Stop Bits: 1
  - o Parity: None
  - o 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 **X** for **Xmodem download**. You will see the following display:

```
Perform xmodem download

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

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

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

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

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

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

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

Boot image recovery is complete.

## VLAGs

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

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

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

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

### On the VLAG Secondary Peer:

1. Shutdown the VLAG ports on which you need to make the change.
2. Disable their VLAG instance using the command:  
RS8264 (config)# no vlag adminkey <key> enable (or)  
RS8264 (config)# no vlag portchannel <number> enable
3. Change the configuration as needed.

### On the VLAG Primary Peer:

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

### On the VLAG Secondary Peer:

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

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

---

## Known Issues

This section describes known issues for N/OS 8.3 on the RackSwitch G8264.

**Note:** Please review the Change History documentation posted with the Switch Firmware to check if any of these issues have been fixed in the latest release.

### 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 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

- Lenovo 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

- 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 Lenovo 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)

## 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.

## 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 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)

```
G8264(config)#
show interface port <swunit:port_num> egress-mcast-queue-counters

For example:
RS8264CS(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>**

## Routed Ports

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 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 **RS G8052(config)# 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)

## 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)

## Virtual Link Aggregation groups

- The following features are not supported on ports participating in VLAGs:
  - Hotlinks
  - IGMP relay
  - Private VLANs
- 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

- vNIC egress bandwidth control is not strictly enforced on the switch for packets larger than 900 bytes, resulting in greater egress bandwidth from the switch to the server than is configured. However, ingress bandwidth control (from the server to the switch) is strictly enforced. (ID: 50950)
- When 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)

