

Lenovo RackSwitch G8264CS

# Release Notes

For Networking OS 8.2

***lenovo***<sup>®</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 (May 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 8.2 for the Lenovo RackSwitch G8264CS (referred to as G8264CS throughout this document).

This supplement modifies and extends the following Lenovo N/OS documentation for use with N/OS 8.2:

- *Lenovo Networking OS 8.2 Application Guide*
- *Lenovo Networking OS 8.2 ISCLI Reference*
- *Lenovo RackSwitch G8264CS 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 8.2 software is supported on the G8264CS, a high performance Layer 2-3 network switch.

The G8264CS is a 1U rack-mountable aggregation switch with unmatched line-rate Layer 2 performance. The G8264CS uses a wire-speed, non-blocking switching fabric that provides simultaneous wire-speed transport of multiple packets at low latency on all ports.

The switch unit contains the following switching ports:

- Thirty-six 10 Gigabit Ethernet (GbE) Small Form Pluggable Plus (SFP+) ports which also support legacy 1 GbE connections
- Four 40 GbE Quad Small Form Pluggable Plus (QSFP+) ports, each of which can optionally be used as four 10 GbE SFP+ ports
- Twelve Omni Ports (port pairs selectable between 10 Gb SFP+ Ethernet mode or 4/8 Gbps Fibre Channel mode).

**Figure 1.** RackSwitch G8264CS front panel



QSFP+ ports can be populated with optical QSFP+ modules or Direct-Attach Cables (DACs), including those that allow the port to be used as four 10 GbE SFP+ ports. SFP+ ports can be populated with optical or copper transceiver modules or DACs.

---

## Updating the Switch Software Image

The switch software image is the executable code running on the G8264CS. 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 G8264CS, 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:

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 5](#).

## Loading New Software to Your Switch

The G8264CS 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 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 the following procedures to download the new software to your switch.

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

```
RS G8264CS(config)# copy {sftp|tftp|ftp} {image1|image2|boot-image}
```

2. Enter the hostname or IP address of the SFTP, 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.

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:

```
RS G8264CS# configure terminal
RS G8264CS(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:

```
RS G8264CS(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 8.2 for Lenovo RackSwitch G8264CS (G8264CS) has been updated to include several new features, summarized in the following sections. For more detailed information about configuring G8264CS features and capabilities, refer to the complete N/OS 8.2 documentation as listed on [page 3](#).

### Backdoor PSIRT Fix

Certain login username/password combinations used for development debugging were hard-coded on the switches and were not changeable. These were a cause of potential security vulnerability. Access to these login username/password combinations has been removed.

### CEE Command Enhancement

The command:

```
RS G8264CS(config)# cee global ets priority-group <priority group>
```

allows configuration of all eight priority groups at the same time.

### BGP Community Lite

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

### Display BGP Routes

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

### Hotlinks + STP

You no longer have to disable STP globally when enabling the Hotlinks feature; you only have to disable STP on Hotlinks trigger interfaces.

### TACACS+ Enhancement

When a TACACS+ server is configured to provide authentication, authorization, and accounting (AAA) services but is unavailable, you will be prompted for a username and password.

### ESN to SNMP

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

## IPSec over Virtual Links

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

## IPv6 Counter Enhancement

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

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

## Link Aggregation Control Protocol (LACP) Individual Mode

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

## RMON Support (RFC1757, RFC2819)

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

## Secure Image Signing

This release adds a signature to the header of the image to ensure that the image loaded onto the switch is a genuine Lenovo official image.

## Service Location Protocol

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

## Spanning Tree Protocol (STP) Range Enhancement

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



## **SNMP**

The following features have been added to SNMP support.

### ***SNMP ACL***

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

### ***SNMP Trap Host***

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

## **Use SSH Public Keys for up to 20 Local Switch Users**

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

## **vLAG MSTP Enhancement**

This enhancement removes STP configuration restrictions, such as changing the MSTP instance and VLAN associations, that were enforced in previous releases when vLAG and MSTP were both enabled. The vLAG interswitch link ports are no

---

## Supplemental Information

This section provides additional information about configuring and operating the G8264CS 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 G8264CS, 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 G8264CS-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:  
RS G8264CS (config)# no vlag adminkey <key> enable (or)  
RS G8264CS (config)# no portchannel <number> enable
3. Change the configuration as needed.

### On the VLAG Primary Peer:

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

### On the VLAG Secondary Peer:

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

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

---

## Known Issues

This section describes known issues for N/OS 8.2 on the Lenovo RackSwitch G8264CS.

**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.

### BBI

- Some versions of Microsoft Internet Explorer version 6.x do not perform HTTP download efficiently. If you have one of these versions, HTTP software download might take much longer than expected (up to several minutes).
- Web-browsers from different vendors may vary in their support of standard features. If you encounter problems using the BBI in a particular browser, a different browser may resolve the issue.

### BGP

Maximum number of route maps that can be added to a BGP peer inbound or outbound is eight. (ID: 46448)

### DHCP

When a static IP address is configured for the management interface, the switch sends a DHCP INFORM packet through the management port, but ignores the returning DHCP ACK packets. (ID: 68071)

### 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

- For FCoE multipathing with UFP and QLogic NetXtreme II NICs in a Flex System Interconnect Fabric solution, the ports along the two FCoE paths (through the two G8264CS switches) should be configured as follows:
  - For one path, the ports should have native vlan 1 and any FCoE vlan
  - For the other path, the ports may have any native vlan id (different from 1) and the FCoE vlan must be 1002.(ID: LV303000)
- Only the ISCLI and BBI are supported for configuring the FCoE feature. After configuring FCoE, save any subsequent configurations only in ISCLI or BBI. If IBM N/OS CLI is used to save any switch configuration, the FCoE configuration will be lost.
- In N/OS 7.8, the G8264CS supports up to 175 simultaneous FCoE sessions. When this capacity is reached, traffic for additional sessions is dropped, though some host servers and uplink devices may consider all sessions fully established. (ID: 60337, 64842)



- When using FCoE to connect the switch to a Cisco Nexus 5000 (as an external FCF), the DCBX PFC willing flag must be enabled. (ID: 65043)
- 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)
- In NPV mode, when static LAG is enabled or disabled, multiple MAC addresses are displayed for a single port. You can view this using the **show fcoe fips fcoe** command. However, the MAC addresses get cleared based on the configured timeout value. (ID: XB223966)

## Fibre Channel

- Use only the ISCLI or BBI to configure Fibre Channel. IBM N/OS CLI is not supported. After configuring Fibre Channel, save any subsequent configurations only in ISCLI or BBI. If IBM N/OS CLI is used to save any switch configuration, the Fibre Channel configuration will be lost.
- If you need to change the Fibre Channel configuration mode from BBI to ISCLI, first save the configuration using the save button on the BBI. Do not use two configuration modes at the same time. (ID: XB205624)
- In a topology where local FCFs (in full FC mode or NPV mode) are connected to external FCFs (such as Cisco Nexus 5000 Series Switch) via Ethernet ports, we recommend you configure two different default VLANs: one for the local FCFs and one for the external FCFs. The default VLANs are used to learn the FCoE VLANs supported. (ID: XB251853)
- A host with an Emulex CNA periodically sends FIP control frames to every FC port in a FCoE VLAN. When there are more than the recommended number of hosts connected to an FCoE VLAN, the FIP control frames being processed by the FC module might exceed its CPU processing power. This will result in FCoE session flaps or loss of connection to the FC module.

To avoid this, consider the following guidelines when configuring the FCoE VLANs to suit the scale of your Data Center Network:

- 1 VLAN comprising of 12 FC Uplink Ports can scale up to 70 servers
- 2 VLANs comprising of 6 FC Uplink Ports each can scale up to 120 servers
- 4 VLANs comprising of 3 FC Uplink Ports each can scale up to 160 servers

## HTTPS

While handling an HTTPS request, the switch may crash if the connection to the client is suddenly terminated during the session. (ID: XB205895)

## IGMP

- The G8264 supports the following IGMP capacities (ID: 45775):
  - IGMP Snooping mode: 3072 IGMP and IPMC groups
  - IGMP Relay mode: 1000 IGMP groups and IPMC groups
- Only 1024 VLANs can be added to IGMP Snooping. Only 8 VLANs can be added to IGMP Relay. (ID: 45781)

## IPsec

- When configuring IPsec to operate between IBM switches, keys of various lengths are supported. However, when using IPsec to connect with non-IBM devices, the manual policy session keys must be of the following fixed lengths:
  - For the AH key:
    - SHA1 = 20 bytes
    - MD5 = 16 bytes
  - For the ESP auth key:
    - SHA1 = 20 bytes
    - MD5 = 16 bytes
  - For the ESP cipher key:
    - 3DES = 24 bytes
    - AES-cbc = 24 bytes
    - DES = 8 bytes
- IPsec does not support OSPFv3 virtual links. (ID: 48914)
- Packet fragmentation over IPsec is supported in transport mode only. Fragmentation is not available in tunneling mode. (ID: 50291)

## ISCLI

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

## LACP

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

- 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

- Egress packets contribute to statistics on IBM Omni Ports even when link is down or transceivers are not present. (ID: 62639)

- In Ethernet mode (the default), IBM Omni Ports may take longer than dedicated Ethernet ports to reflect changes in port link status. As a result, some traffic loss can be expected while the port transitions to a down state. When using protocols sensitive to link failover timing, it is recommended to use regular Ethernet ports first. (ID: 67288)
- There is a time delay of up to 25 seconds when the following command is executed on IBM Omni Ports (ID: 69836):

```
RS8264CS# clear interface port <port number range> counters
```

## QoS

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

```
RS8264CS(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>

## SNMP

- During SNMP MIB walks, if you experience timeouts, set the timeout value to 3 seconds or higher in the SNMP application/tool. (IDs: 71913, 71914, 71906)
- If you delete multiple VLANs using SNMP, you may see an error if the SNMP packet size exceeds 1800 bytes. (ID: XB228120)

## Spanning Tree

- When using LACP with PVRST, it is not recommended to configure the switch as the STP root bridge. When doing so, traffic can be discarded for up to 30 seconds on affected LACP ports while initial STP path states are being resolved (discarding, learning, forwarding). (ID: 63315)

- After changing from MSTP to a different STP mode, some MSTP commands will still appear in the configuration file. The non-applicable MSTP commands do not affect switch operation and can be ignored. (ID: 64388)

## Virtual Link Aggregation groups

In MSTP mode, dynamically changing the Spanning Tree Group of a vLAG-enabled port is not allowed. When Spanning Tree Protocol is enabled, you cannot add vLAG-enabled ports to new VLANs created in the switch without first globally disabling vLAG. (ID: 57336)

## 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).

## vNICs

- Bandwidth metering drops excess packets when the configured limits on the vNIC pipe are reached. Any CEE Enhanced Transmission Selection configuration is ignored. (ID: 50950)
- vNIC egress bandwidth control is not strictly enforced on the switch for packets larger than 900 bytes, resulting in greater egress bandwidth from the switch to the server than is configured. However, ingress bandwidth control (from the server to the switch) is strictly enforced.