IBM Virtual Fabric 10Gb Switch Module for IBM BladeCenter®

Release Notes

For Networking OS 7.8

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 (December 2013)

Release Notes

The IBM Virtual Fabric 10Gb Switch Module (VFSM) is one of up to four switch modules that can be installed in the IBM BladeCenter chassis.

This release supplement provide the latest information regarding IBM Networking OS 7.8 for the Virtual Fabric Switch Module (referred to as VFSM throughout this document).

This supplement modifies and extends the following IBM N/OS documentation for use with *N/OS* 7.8:

- *IBM Networking OS* 7.8 *Application Guide*
- *IBM Networking OS* 7.8 *Command Reference*
- *IBM Networking OS* 7.8 *ISCLI Reference*
- *IBM Networking OS* 7.8 *BBI Quick Guide*
- *IBM Virtual Fabric 10Gb Switch Module Installation Guide*

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

`http://www.ibm.com/support`

Please keep these release notes with your product manuals.

Hardware Support

N/OS 7.8 software is supported only on the IBM Virtual Fabric 10Gb Switch Module (IBM model name 46C7191) for IBM BladeCenter. The Virtual Fabric Switch Module (VFSM) shown in Figure 1. is a high performance Layer 2-3 embedded network switch that features tight integration with IBM BladeCenter H or BladeCenter HT management modules.

Figure 1. Virtual Fabric Switch Module Faceplate

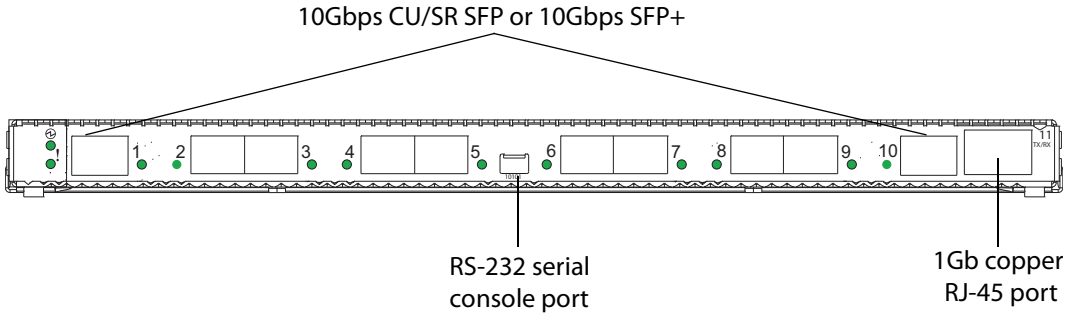

The VFSM has the following port capacities:

- Ten 10Gbps CU/SR SFP or 10Gbps SFP+
- Fourteen 1Gb/10Gb internal ports
- One 10/100/1000Mbps external copper (RJ-45) port
- Two 100Mb internal management ports
- One RS-232 serial port

Updating the Switch Software Image

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

```
VFSM# show version
```

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

- Load a new software image and boot image onto an FTP or TFTP server on your network.
- Transfer the new images to your switch.
- Specify the new software image as the one which will be loaded into switch memory the next time a switch reset occurs.
- Reset the switch.

For instructions on the typical upgrade process using the CLI, ISCLI, or BBI, see "Loading New Software to Your Switch" on page 6.

CAUTION:

Although the typical upgrade process is all that is necessary in most cases, upgrading from (or reverting to) some versions of N/OS or BLADEOS requires special steps prior to or after the software installation process. Please be sure to follow all applicable instructions in the following sections to ensure that your switch continues to operate as expected after installing new software.

Special Software Update Issues

When updating to N/OS 7.8, 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 5.x or Prior

After updating:

- The STG port priority value is different compared to release 5.x and prior. In release 5.x and prior, the priority value could be set to any integer from 0 to 255. In release 6.3 and later, the range is still 0 to 255, but must be specified in increments of 4 (such as 0, 4, 8, 12, and so on).

  If the specified value is not evenly divisible by 4, the value will be automatically rounded down to the nearest valid increment whenever manually changing the priority value, when loading a configuration from prior to release 6.3, and during the software upgrade process. If using STG port priorities, after upgrading to release 6.3 or later, it is recommended that the administrator review the configured values and make any appropriate changes. (ID: 38556)

Updating from BLADEOS 6.1 or Prior

After updating:

- Some time zones are different compared to release 6.1.2 and prior. After upgrading to release 6.3 or later, it is recommended that the administrator review the configured time zone and make any appropriate changes. (ID: 29778)
- If the switch is connected to a bridge module (such as for FCoE), the bridge module BR ports will be disabled when the upgraded software first boots. To restore the proper configuration, re-enable the BR ports and reboot the switch. (ID: 43890)

Updating from BLADEOS 6.5.1 or Prior

After updating:

- The default value for port flow control for external uplink ports is different compared to release 6.5.1 or prior. After upgrading to release 6.5.2 or later, it is recommended that the administrator review the configured flow control settings and make any appropriate changes. (ID: 43781)

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 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 from IBM Networking OS 7.2 or Prior

After updating:

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

Loading New Software to Your Switch

The VFSM 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 17).

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

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

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

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

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

Loading Software via the N/OS CLI

1. Enter the following Boot Options command:

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

2. Enter the name of the switch software to be replaced:

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

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

```
Enter hostname or IP address of FTP/TFTP server: <hostname or IP address>
```

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

```
Enter name of file on FTP/TFTP server: <filename>
```

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

5. Enter your username for the server, if applicable.

```
Enter username for FTP server or hit return for
TFTP server: {<username> /<Enter>}
```

   If entering an FTP server username, you will also be prompted for the password. The system then prompts you to confirm your request. Once confirmed, the software will load into the switch.

6. If software is loaded into a different image than the one most recently booted, the system will prompt you whether you wish to run the new image at next boot. Otherwise, you can enter the following command at the `Boot Options#` prompt:

```
Boot Options# image
```

   The system then informs you of which software image (`image1` or `image2`) is currently set to be loaded at the next reset, and prompts you to enter a new choice:

```
Currently set to use switch software "image1" on next reset.
Specify new image to use on next reset ["image1"/"image2"]:
```

   Specify the image that contains the newly loaded software.

7. Reboot the switch to run the new software:

```
Boot Options# reset
```

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

Loading Software via the ISCLI

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

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

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

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

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

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

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

4. If required by the FTP or TFTP server, enter the appropriate username and password.
5. The switch will prompt you to confirm your request.

   Once confirmed, the software will begin loading into the switch.

6. When loading is complete, use the following commands to enter Global Configuration mode to select which software image (`image1` or `image2`) you want to run in switch memory for the next reboot:

```
Router# configure terminal
Router(config)# boot image {image1|image2}
```

   The system will then verify which image is set to be loaded at the next reset:

```
Next boot will use switch software image1 instead of image2.
```

7. Reboot the switch to run the new software:

```
Router(config)# reload
```

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

Loading Software via BBI

You can use the Browser-Based Interface to load software onto the VFSM. The software image to load can reside in one of the following locations:

- FTP server
- TFTP server
- Local computer

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

1. Click the Configure context tab in the toolbar.
2. In the Navigation Window, select System > Config/Image Control.

   The Switch Image and Configuration Management page appears.

3. If you are loading software from your computer (HTTP client), skip this step and go to the next. Otherwise, if you are loading software from a FTP/TFTP server, enter the server's information in the FTP/TFTP Settings section.
4. In the Image Settings section, select the image version you want to replace (Image for Transfer).
   - If you are loading software from a FTP/TFTP server, enter the file name and click **Get Image**.
   - If you are loading software from your computer, click **Browse**.

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

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

New and Updated Features

N/OS 7.8 for IBM Virtual Fabric 10Gb Switch Module (VFSM) has been updated to include several new features, summarized in the following sections. For more detailed information about configuring VFSM features and capabilities, refer to the complete N/OS 7.8 documentation as listed on page 3.

ACLs

Metering is supported for IPv6 ACLs.

Edge Virtual Bridging (EVB)

VSI Database can be accessed via HTTP or HTTPS. The manager IP can be configured with an IPv4 or IPv6 address.

FCoE Link Aggregation

FCoE LAG can be configured in stacking and stand-alone mode.

Multiple Spanning Tree Protocol (MSTP)

In IBM Networking OS 7.8, VLANs can be mapped to MSTP instances without creating them on the switch. In previous IBM Networking OS releases, the VLANs were created on the switch which often resulted in the switch having multiple unused VLANs.

Use the following commands to configure MSTP:

1. Configure port and VLAN membership on the switch.
2. Configure Multiple Spanning Tree region parameters and set the mode to MSTP.

```
VFSM(config)# spanning-tree mst configuration        (Enter MST configuration mode)
VFSM(config-mst)# name <name>                        (Define the Region name)
VFSM(config-mst)# exit
VFSM(config)# spanning-tree mode mst          (Set mode to Multiple Spanning Trees)
```

3. Map VLANs to MSTP instances:

```
VFSM(config)# spanning-tree mst configuration                (Enter MST configuration mode)
VFSM(config-mst)# instance <instance ID> vlan  <vlan number or range>
```

NIST-800 131a Compliance

The implementations specified in this section are compliant with National Institute of Standards and Technology (NIST) Special Publication (SP) 800-131a. SP 800-131a

The Virtual Fabric Switch Module can operate in two boot modes:

- Compatibility mode (default): This is the default switch boot mode. This mode may use algorithms and key lengths that may not be allowed/acceptable by NIST SP 800-131a specification. This mode is useful in maintaining compatibility with previous releases and in environments that have lesser data security requirements.
- Strict mode: Encryption algorithms, protocols, and key lengths in strict mode are compliant with NIST SP 800-131a specification.

When in boot strict mode, the switch uses Secure Sockets Layer (SSL)/Transport Layer Security (TLS) 1.2 protocols to ensure confidentiality of the data to and from the switch.

Before enabling strict mode, ensure the following:

- The software version on all connected switches is IBM N/OS 7.8.
- The supported protocol versions and cryptographic cipher suites between clients and servers are compatible. For example: if using SSH to connect to the switch, ensure that the SSH client supports SSHv2 and a strong cipher suite that is compliant with the NIST standard.
- Compliant Web server certificate is installed on the switch, if using BBI.
- A new self-signed certificate is generated for the switch (`VFSM(config)# access https generate-certificate`). The new certificate is generated using 2048-bit RSA key and SHA-256 digest.
- Protocols that are not NIST SP 800-131A compliant must be disabled or not used.
- Only SSHv2 or higher is used.
- The current configuration, if any, must be saved in a location external to the switch. When the switch reboots, both the startup and running configuration are lost.
- Only protocols/algorithms compliant with NIST SP 800-131a specification are used/enabled on the switch. Please see the NIST SP 800-131a publication for details. The following table lists the acceptable protocols and algorithms:

Table 1. Acceptable Protocols and Algorithms

Protocol/Function	Strict Mode Algorithm	Compatibility Mode Algorithm
BGP	BGP does not comply with NIST SP 800-131a specification. When in strict mode, BGP is disabled. However, it can be enabled, if required.	Acceptable
Certificate Generation	RSA-2048 SHA-256	RSA 2048 SHA 256
Certificate Acceptance	RSA 2048 or higher SHA 224 or higher	RSA SHA, SHA2
HTTPS	TLS 1.2 only See “Acceptable Cipher Suites” on page 13;	TLS 1.0, 1.1, 1.2 See “Acceptable Cipher Suites” on page 13;

Table 1. Acceptable Protocols and Algorithms

Protocol/Function	Strict Mode Algorithm	Compatibility Mode Algorithm
IKE		
Key Exchange	DH Group 24	DH group 1, 2, 5, 14, 24
Encryption	3DES, AES-128-CBC	3DES, AES-128-CBC
Integrity	HMAC-SHA1	HMAC-SHA1, HMAC-MD5
IPSec		
AH	HMAC-SHA1	HMAC-SHA1, HMAC-MD5
ESP	3DES, AES-128-CBC, HMAC-SHA1	3DES, AES-128-CBC, HMAC-SHA1, HMAC-MD5
LDAP	LDAP does not comply with NIST SP 800-131a specification. When in strict mode, LDAP is disabled. However, it can be enabled, if required.	Acceptable
OSPF	OSPF does not comply with NIST SP 800-131a specification. When in strict mode, OSPF is disabled. However, it can be enabled, if required.	Acceptable
RADIUS	RADIUS does not comply with NIST SP 800-131a specification. When in strict mode, RADIUS is disabled. However, it can be enabled, if required.	Acceptable
Random Number Generator	NIST SP 800-90A AES CTR DRBG	NIST SP 800-90A AES CTR DRBG
Secure NTP	Secure NTP does not comply with NIST SP 800-131a specification. When in strict mode, secure NTP is disabled. However, it can be enabled, if required.	Acceptable
SLP	SHA-256 or higher RSA/DSA 2048 or higher	
SNMP	SNMPv3 only AES-128-CFB-128/SHA1 **Note:** Following algorithms are acceptable if you choose to support old SNMPv3 factory default users: AES-128-CFB/SHA1 DES/MD5 AES-128-CFB-128/SHA1	SNMPv1, SNMPv2, SNMPv3 DES/MD5, AES-128-CFB-128/SHA1

Table 1. Acceptable Protocols and Algorithms

Protocol/Function	Strict Mode Algorithm	Compatibility Mode Algorithm
SSH/SFTP		
Host Key	SSH-RSA	SSH-RSA
Key Exchange	ECDH-SHA2-NISTP521 ECDH-SHA2-NISTP384 ECDH-SHA2-NISTP256 ECDH-SHA2-NISTP224 RSA2048-SHA256 DIFFIE-HELL-MAN-GROUP-EXCHANGE-SHA256 DIFFIE-HELL-MAN-GROUP-EXCHANGE-SHA1	ECDH-SHA2-NISTP521 ECDH-SHA2-NISTP384 ECDH-SHA2-NISTP256 ECDH-SHA2-NISTP224 ECDH-SHA2-NISTP192 RSA2048-SHA256 RSA1024-SHA1 DIFFIE-HELL-MAN-GROUP-EXCHANGE-SHA256 DIFFIE-HELL-MAN-GROUP-EXCHANGE-SHA1 DIFFIE-HELL-MAN-GROUP14-SHA1 DIFFIE-HELL-MAN-GROUP1-SHA1
Encryption	AES128-CTR AES128-CBC 3DES-CBC	AES128-CTR AES128-CBC RIJNDAEL128-CBC BLOWFISH-CBC 3DES-CBC ARCFOUR256 ARCFOUR128 ARCFOUR
MAC	HMAC-SHA1 HMAC-SHA1-96	HMAC-SHA1 HMAC-SHA1-96 HMAC-MD5 HMAC-MD5-96
TACACS+	TACACS+ does not comply with NIST SP 800-131a specification. When in strict mode, TACACS+ is disabled. However, it can be enabled, if required.	Acceptable

Acceptable Cipher Suites

The following cipher suites are acceptable (listed in the order of preference) when the Virtual Fabric Switch Module is in compatibility mode:

Table 2. List of Acceptable Cipher Suites in Compatibility Mode

Cipher ID	Key Exchange	Authentication	Encryption	MAC	Cipher Name
0xC027	ECDHE	RSA	AES_128_CBC	SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
0xC013	ECDHE	RSA	AES_128_CBC	SHA1	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
0xC012	ECDHE	RSA	3DES	SHA1	SSL_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
0xC011	ECDHE	RSA	RC4	SHA1	SSL_ECDHE_RSA_WITH_RC4_128_SHA
0x002F	RSA	RSA	AES_128_CBC	SHA1	TLS_RSA_WITH_AES_128_CBC_SHA
0x003C	RSA	RSA	AES_128_CBC	SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256
0x0005	RSA	RSA	RC4	SHA1	SSL_RSA_WITH_RC4_128_SHA
0x000A	RSA	RSA	3DES	SHA1	SSL_RSA_WITH_3DES_EDE_CBC_SHA
0x0033	DHE	RSA	AES-128_CBC	SHA1	TLS_DHE_RSA_WITH_AES_128_CBC_SHA
0x0067	DHE	RSA	AES_128_CBC	SHA256	TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
0x0016	DHE	RSA	3DES	SHA1	SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA

The following cipher suites are acceptable (listed in the order of preference) when the Virtual Fabric Switch Module is in strict mode:

Table 3. List of Acceptable Cipher Suites in Strict Mode

Cipher ID	Key Exchange	Authentication	Encryption	MAC	Cipher Name
0xC027	ECDHE	RSA	AES_128_CBC	SHA256	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256
0xC013	ECDHE	RSA	AES_128_CBC	SHA1	TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA
0xC012	ECDHE	RSA	3DES	SHA1	SSL_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA
0x0033	DHE	RSA	AES-128_CBC	SHA1	TLS_DHE_RSA_WITH_AES_128_CBC_SHA
0x0067	DHE	RSA	AES_128_CBC	SHA256	TLS_DHE_RSA_WITH_AES_128_CBC_SHA256
0x0016	DHE	RSA	3DES	SHA1	SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA
0x002F	RSA	RSA	AES_128_CBC	SHA1	TLS_RSA_WITH_AES_128_CBC_SHA
0x003C	RSA	RSA	AES_128_CBC	SHA256	TLS_RSA_WITH_AES_128_CBC_SHA256
0x000A	RSA	RSA	3DES	SHA1	SSL_RSA_WITH_3DES_EDE_CBC_SHA

Configuring Strict Mode

To change the switch mode to boot strict mode, use the following command:

```
VFSM(config)# [no] boot strict enable
```

When strict mode is enabled, you will see the following message:

```
Warning, security strict mode limits the cryptographic algorithms used by secure
protocols on this switch. Please see the documentation for full details, and
verify that peer devices support acceptable algorithms before enabling this
mode. The mode change will take effect after reloading the switch and the
configuration will be wiped during the reload. System will enter security strict
mode with default factory configuration at next boot up.

Do you want SNMPV3 support old default users in strict mode (y/n)?
```

When strict mode is disabled, the following message is displayed:

```
Warning, disabling security strict mode. The mode change will take effect after
reloading the switch.
```

You must reboot the switch for the boot strict mode enable/disable to take effect.

Limitations

In IBM N/OS 7.8, consider the following limitation/restrictions if you need to operate the switch in boot strict mode:

- Power ITEs and High-Availability features do not comply with NIST SP 800-131A specification.
- The VFSM will not discover Platform agents/Common agents that are not in strict mode.
- Web browsers that do not use TLS 1.2 cannot be used.
- Limited functions of the switch managing Windows will be available.

Private VLANs

IBM N/OS supports Private VLAN configuration as described in RFC 5517.

Note: Private VLANs feature is not supported in stacking mode.

Telnet

Two attempts are allowed to log in to the switch. After the second unsuccessful attempt, the Telnet client is disconnected via TCP session.

UFP

Unified Fabric Port (UFP) is a cost-effective way to allocate, share and dynamically control network bandwidth between a server and a switch. UFP lets you create multiple virtual connections. The UFP protocol is a link-level protocol that runs a separate instance for each physical communication link established between a server NIC and a switch port. Virtualizing the ports allows you to separate or aggregate port traffic by applying the network policies defined on the switch. Virtualization lessens bottlenecks and provides higher bandwidth while consolidating equipment use.

UFP provides a switch fabric component to control NICs. The server operating system (OS) or hypervisor recognizes each subdivided link (channel) as an independent physical NIC. Each channel has a unique identity and profile that defines its properties and functionality. The server communicates with the switch over the channel as defined in the channel profile. The channels share the high-speed physical link bandwidth.

Please see the *IBM Networking OS IBM Virtual Fabric 10Gb Switch Module Application Guide* for details.

UFP works with other VFSM features.

Layer 2 Failover

UFP can be configured with auto-monitoring or manual monitoring. In auto-monitoring, a vPort is automatically associated with a Failover trigger if it has any VLAN in common with the monitor ports.

Increased VLAN Limits

Configured with UFP and VLANs, a vPort can support maximum 256 VLANs. A UFP port supports 256 VLANs.

VMReady

Configuring with UFP and VMReady, the VFSM can support up to 32 VMGroups with UFP vPorts in auto-mode.

User Access

Up to 20 users can be configured to allow access to the switch. Each user can be configured with a password and access level.

Resolved Issues

The following known issues have been resolved.

Private VLANs

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

Supplemental Information

This section provides additional information about configuring and operating the VFSM 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.

Management Module

- The “Fast POST=Disabled/Enabled” inside the IBM management module Web interface “I/O Module Admin Power/Restart” does not apply to the VFSM.

  Solution: To boot with Fast or Extended POST, go to the “I/O Module Admin/Power/Restart” window. Select the VFSM, and then choose “Restart Module and Run Standard Diagnostics” or “Restart Module and Run Extended Diagnostics.”
- The following table correlates the Firmware Type listed in the IBM management module’s Web interface “Firmware VPD” window to the VFSM software version:

  Table 4. Firmware Type list

Firmware Type	Description
Boot ROM	VFSM Boot code version
Main Application 1	Currently running image
Main Application 2	Backup image

- Within the IBM management module Web interface, the Java applets of “Start Telnet Session” and “Start Web Session” do not support changing of default known ports 23 and 80 respectively.

  Solution: If the Telnet or HTTP port on the VFSM is changed to something other than the default port number, the user must use a separate Telnet client or Web browser that supports specifying a non-default port to start a session to the VFSM user interface.

Management Module/VFSM Connectivity

Currently, the IBM management module is designed to provide one-way control of the VFSM. As a result, the VFSM may lose connectivity to the management module via the management port under the following conditions:

- If new IP attributes are pushed from the management module to the VFSM while the IP Routing table is full, the new attributes will not be applied.

  Solution: Enable “External Management over all ports,” connect to the switch using other interface and then clear the routing table. Then push the IP address from the management module. If this does not work, use Solution 2 below.
- If you execute the /boot/reset CLI command on the VFSM or the VFSM resets itself, the management module might not push the IP attributes to the switch, and connectivity may be lost.

Solution 1: If you should experience any connectivity issues between the switch module and the management module, go to the “I/O Module Configuration” window on the management module’s Web interface. Under the “New Static IP Configuration” section, click **Save** to trigger the management module to push the stored IP attributes to the switch module.

Solution 2: If Solution 1 does not resolve your connectivity issue, then go to the “I/O Module Admin/Power/Restart” window on the management module’s Web interface. Restart the switch module in question.

Solution 3: If this still does not resolve the issue, enable Preserve new IP configuration on all resets setting on the management module and restart the switch module via the “I/O Module Admin/Power/Restart” window on the management module’s Web interface.

Note: As a rule, always use the management module Web interface to change the VFSM management IP attributes (IP address, mask and gateway), and then click Save to push the IP attributes to the switch module. Use of the command-line interface to change the switch module management IP attributes may result in duplicated entries for the management IP Interface in the switch route table and/or loss of connectivity via the management module.

External Port Link Negotiation

Autonegotiation settings for each external switch port should be the same as those of the devices being connected. In a valid configuration, both ends of a port link are set with autonegotiation on, or both ends are set to specific speed and link properties with autonegotiation disabled.

Internal Port Autonegotiation

By default, link autonegotiation is turned on for internal ports. This is in contrast to external ports, where autonegotiation is off by default. Internal ports use autonegotiation in order to support the Wake-Over-LAN (WOL) features of some servers. If an attached server does not support autonegotiation or WOL, turn autonegotiation off for the internal port.

Port Mirroring Tags BPDU Packets

When you perform port mirroring, Spanning Tree BPDU packets are VLAN tagged at the monitoring port. This is standard behavior of port mirroring on the VFSM. All mirrored egress traffic is tagged.

Secure Management Network

The following VFSM attributes are reserved to provide secure management access to and from the IBM management module:

- MGT1 (port 15) and MGT2 (port 16)
- VLAN 4095
- IP interface 128
- Gateway 132

For more information about remotely managing the VFSM through the external ports, see “Accessing the Switch” in the *IBM Networking OS 7.8 Application Guide*.

Note: The external uplink ports (EXTx) cannot be members of management VLANs.

Secure Shell (SSH)

Because SSH key generation is CPU intensive, the VFSM attempts to avoid unnecessary key generation. The process generates three server keys:

1. One key is generated to replace the current server key, if used.
2. A second key is generated as a spare, in case the current server key is used and the specified interval expires.
3. A third key is generated for use at the next reboot.

Therefore, if you never login via SSH, you will only see two key generation events. You may see all three events directly following a reboot. If you want to witness the key generation after the specified interval has expired, then you must login via SSH at least once during each expiration interval.

Spanning Tree Configuration Tips

To ensure proper operation with switches that use Cisco Per VLAN Spanning Tree (PVST+), you must do one of the following:

- Create a separate Spanning Tree Group for each VLAN.
- Ensure that the trunk link between two switches belongs to the same native VLAN.

Syslog Configuration Tip

The *facility* parameter traditionally is used to correlate services (such as IP, CLI, etc.) to messages. This is done to distinguish between the different services that are running in the network/device. However, for the VFSM, there is a single configured facility value (0-7) used on all messages. By configuring a unique facility value for each switch, a single SYSLOG server can distinguish between the various VFSMs in the network. Refer to “System Host Log Configuration” in the *Command Reference.*

Trunk Group Configuration Tips

Please be aware of the following information when you configure trunk groups:

- Always configure trunk groups first, on both ends, before you physically connect the links.
- Configure all ports in a trunk group to the same speed.

vCenter Synchronization

When applying distributed VM group configuration changes, the switch will attempt to synchronize settings with the VMware vCenter for virtualization management. If the vCenter is unavailable, an error message will be displayed on the switch. Be sure to evaluate all error message and take the appropriate actions to ensure the expected changes are properly applied. If corrective actions are not taken, synchronization may remain incomplete when connection with the vCenter is restored.

Solution: When the switch connection with the vCenter is restored, use the following operational command to force synchronization:

```
VFSM(config)# virt vmware scan
```

VRRP Configuration

Although the Virtual Router Redundancy Protocol (VRRP) standard permits up to 255 virtual router instances, the N/OS 7.8 implementation only allows up to 128 virtual router instances (corresponding to the number of supported IP interfaces). Each virtual router instance can be assigned a unique Virtual Router ID (VRID) between 1 and 255.

Known Issues

This section describes known issues for N/OS 7.8 on the IBM Virtual Fabric 10Gb Switch Module

ACLs

- When an Access Control List (ACL) is installed on two different ports, only one statistics counter will be available. The VFSM does not support two different statistics counter for one ACL installed on two different ports.
- The ACL filters for TCP/UDP work properly only on packets that do not have IP options.
- When you assign an ACL (or ACL Group) to one port in a trunk, N/OS does not automatically assign the ACL to other ports in the trunk, and it does not prompt you to assign the ACL to other ports in the trunk. Manually assign each ACL or ACL Group to all ports in the trunk.

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.

Boot Configuration Block

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

Debug

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

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)

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)
- Under some circumstances, ENode ports might not establish FCoE connection. This may occur after a bridge module connection fails and is restored, or when the FCoE VLAN is configured prior to enabling CEE and FIP Snooping. To restore the connection, it may be necessary to disable and re-enable the ENode port. (ID: 50904, 51845)
- The FIP Snooping option for automatic VLAN creation is not recommended for use with the Emulex Virtual Fabric Adapter for IBM BladeCenter. Disable automatic VLAN creation when connecting the switch to an Emulex Virtual Fabric Adapter. (ID: 51529)
- Because the effective bandwidth on stack ports is approximately 94% of line rate, some loss of packets is to be expected during heavy load on stack ports. However, when PFC is implemented for loss-sensitive traffic, loss is confined to regular Ethernet flows and does not affect traffic designated for lossless transit (FCoE). (ID: 61894, 63886)
- 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)
- When using two VFSM switches for FCoE, each switch must be connected to its own FCF. Topologies connecting both switches to a single FCF are not currently supported.
- Is it recommended to use a 40Gbps, 4-port connection from the switch to the FCF bridge module.
- When using FCoE to connect the switch to a Cisco Nexus 5000 (as the FCF), the VFSM must be configured as the STP root bridge. To make the switch become the root bridge, configure the bridge priority value of the switch lower than all other switches and bridges on your network:

  ```
  VFSM(config)# spanning-tree stp <x> bridge priority
  <0-65535>
  ```

Forwarding Database (FDB)

From IBM Networking OS 7.8 onwards, MAC address information is no longer learned by control packets such as LACPDUs. This behavior is as expected. (ID: XB253517)

GMT Displayed While Booting

- While the switch is booting, the system time may be displayed for GMT (time zone 0) in the System Log. However, once the switch has finished booting, the administrator-configured time zone will be used for subsequent log messages.

Hotlinks

- Prior to enabling hotlinks, Spanning-Tree should be globally disabled using the following command (ID: 47917):

```
VFSM(config)# spanning-tree mode dis
```

IGMP

- IPMC filter actions for IGMP with IP options, based on the IGMP `IpmcOptFwd` table state:
  - `Off`: There is no IPMC options forwarding filter for this entry. IPMC packets with options corresponding to the IPMC entry will be dropped.
  - `On`: An IPMC options forwarding filter exists for this entry. If the `FWD` state is `YES`, IPMC packets matching the entry will be forwarded to the same ports as those without the options. If the `FWD` state is `NO`, the packets will be dropped.
- When using IPMC filters for IGMP with IP options, the VFSM supports a limited number of groups. The number of groups with the IGMP `IpmcOptFwd` table state set to `On` is based on the ipmc-opt profile selected during boot (for the new profile to take effect, a switch reboot is required):
  - `boot profile ipmc-opt acls-128`
    Only 1536 groups may be set to `On`. The other groups are set to `Off`.
  - `boot profile ipmc-opt acls-256`
    Only 1280 groups may be set to `On`. The other groups are set to `Off`.
  - `boot profile ipmc-opt acls-384`
    Only 1024 groups may be set to `On`. The other groups are set to `Off`.
  - `boot profile ipmc-opt acls-none`
    Only 1792 groups may be set to `On`. The other groups are set to `Off`.

IKEv2

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

IP Gateways

- Although the switch allows IPv4 gateways numbered 1 through 132 to be configured, the Virtual Fabric Switch Module supports only IPv4 gateways numbered 1 to 4. IPv4 gateways 5 through 132 are not supported and should not be configured. (ID: 42433)

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)

ISCLI Configuration Scripts

- When using the ISCLI, configuration commands are applied to the active switch configuration immediately upon execution. As a result, when using the ISCLI to load a configuration script containing a long list of processor-intensive commands (such as static route definitions), switch response to other management functions (such as Telnet access for additional management sessions) may be slow or even time-out while the switch individually applies each scripted command. (ID 31787)

  Solution: The CLI may be used as an alternative to the ISCLI. Because CLI commands are not fully processed until the CLI `apply` command is given, the equivalent configuration script can be loaded in its entirety and then applied as a whole without undue impact on other management sessions.

Jumbo Frames

- Some ingress jumbo frames (for example, ICMP) are not routed from one VLAN to another VLAN. Jumbo frames are routed across data VLANs.

LACP

- If a static trunk on a VFSM is connected to another VFSM with LACP configured (but no active LACP trunk), the `VFSM# show portchannel information` command might erroneously report the static trunk as forwarding.
- If you configure LACP (active/passive) on one port, also configure LACP on the partner switch, at the end of the link. If you connect LACP with a static trunk, there will be no connectivity on that link.
- 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)

Multicast

- When a static multicast MAC entry is created, multicast traffic is forwarded to all participating ports and also back to the originator. The traffic to the originator is counted toward port statistics, but is dropped and has no other effect on operation. (ID: 64280)

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

- Under repeated and rapid removal and reinsertion a port transceiver, it is possible that the resulting port state may not be represented accurately within the switch. (ID 32412)

  Solution: Once you have removed a transceiver from a switch port, wait five seconds before reinserting any transceiver into the same port. This allows the port to stabilize, and promotes accurate port state information within the switch.
- 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)

QoS Metering

Traffic may exceed the configured maximum burst size of the ACL meter (`/cfg/port <x>/aclqos/meter/mbsize`) by one packet, with that packet remaining In-Profile. Once the ACL meter has been exceeded, additional burst packets fall Out-of-Profile.

SNMP

- Due to backward-compatibility issues, two Routing Information Protocol (RIP) MIBs are available in N/OS: `ripCfg` and `rip2Cfg`. Use the `rip2Cfg` MIB to configure RIPv1 and RIPv2 through SNMP.

  N/OS does not support the standard RIPv2 MIB as described in RFC 1724. Use the `rip2Cfg` MIB to configure RIPv1 and RIPv2 through SNMP.
- Certain SNMP MIB browsers may report an error (such as “OID not increasing”) or may halt processing prior to reaching the end of the MIB. This can occur in cases where duplicate lines appear in the MIB, as when two IP route destinations and masks resolve to the same address. Some MIB browsers may interpret such duplicate lines as an error or as the end of the MIB, while others will ignore the duplicates and continue processing to the end of the file.

  To bypass such browser problems, turn off the OID increment check when processing the MIB. For example, use the `-Cc` option of the snmpwalk function in NetSNMP:

  ```
  snmpwalk -Cc -v 1 -t 60 -c public -m ALL -M $miblocation
  10.13.5.103 $1
  ```
- Port information displayed in MIBs related to port-based VLANs does not distinguish between a regular port or a trunk port. Use the `VFSM(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)

Spanning Tree

- When Spanning Tree Protocol is enabled, STP will be applied only to external switch ports. The internal switch ports will always be in the forwarding (`FWD`) state, regardless of the STP setting. (ID: 56659)
- 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)

Stacking

- If the switch is used in stacking mode, whenever you change the high-gig stacking trunk ports (using the `boot stack higig-trunk` command in the ISCLI) or stack VLAN (using the `boot stack` vlan command in the ISCLI), you must reboot the switch in order for the new high-gig trunk port settings to take effect.
- If you change the stack mode (using the `boot stack mode` command in the ISCLI), you must reboot the switch for the new mode to take effect. (ID: 55756)
- 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 7.8. Upload N/OS 7.6 (not N/OS 7.8) 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 7.8. Upload the N/OS 7.6 switch with N/OS 7.8 prior to joining it to the existing N/OS 7.8 stack. (ID: 69744; ID: 70387)

Statistics

- The counter for empty egress port map discards may increase on ports used for stacking. This is the because all multicast or broadcast traffic gets flooded to all stacking ports. Discard of these packets is normal and expected on the stacking ports. (ID: 45689)

TACACS+

Changing the TACACS+ password for the secondary TACACS+ authentication server causes the authentication to failover from the primary authentication server to the secondary. Subsequent authentication attempts fail when using the primary server password and succeed when using the secondary server password.

Solution: To avoid confusion, set the primary authentication server to use the same password as the secondary server prior to applying the configuration.

Trunking

If a member switch in a stack is no longer bound to the master (such as when the member's `csnum` has been manually changed), the master may be unable to determine the speed settings of ports on the unbound member switch. Because all ports in a trunk must be verified by the master as having the same port speed settings, the unknown settings of the unbound member ports may cause a trunk port mis-match error when applying configuration changes. (ID: 41400)

To correct this, in the master switch configuration, remove the member's ports from the stacking trunk, then bind the member with its new `csnum` to the master and reconfigure the appropriate trunk ports.

Tx Ring Loop

When you create a trunk or link loop between the VFSM and another switch, packets might loop infinitely at line rate within the related links. When this problem occurs, the VFSM continuously displays the following messages at the console:

```
WARNING: packet_sent u: 0, dv_active: tx ring full
packet_sent dcnt=114, public1=110, vcnt=1025
```

Remove the loop to resolve this misconfiguration.

VMready

- After adding a VM to a VM group, `addvmtbl` command entries appear in the configuration dump. The `addvmtbl` commands are added automatically by the switch for the purpose of retaining some VM settings beyond switch reset. These commands may be ignored. They should not be entered manually by the administrator. (ID: 42568)
- 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.
- The total number of user defined ACLs supported is 128. The total number of VMs with ACL metering depends on the availability of user defined ACLs. (ID: 62522)
- When VMready information commands are issued on a stack member, you may see VM validation messages from groups in the VM table, although the VM is not attached. (ID: XB210686)

vNICs

- vNICs Are Enabled When Reverting

  Under some circumstances, using Revert Apply might not revert the vNIC configuration as expected. This can occur if LLDP was in the disabled state in the previously applied configuration, and then enabled in the current configuration either manually (using the `/cfg/l2/lldp/on` command) or automatically as the result of enabling vNICs. Under such circumstances, vNICs that were newly enabled in the current session may not fully return to their prior disabled condition. Affected vNICs will appear to be disabled on the switch, but may remain in the enabled state on the server. This can be resolved in one of two ways:

  – Solution 1: Using LLDP—To allow the switch to send the appropriate vNIC status messages to the servers when Revert Apply is executed, enable LLDP and save the configuration prior to performing commands you may wish to revert:

  ```
  VFSM(config)# lldp enable
  VFSM(config)# <trial commands...>
  VFSM(config)# no lldp enable
  ```

  – Solution 2: If you do not wish to keep LLDP enabled, once Revert Apply is executed, manually enable and disable the vNIC feature to force vNIC synchronization:

  ```
  VFSM(config)# vnic enable
  VFSM(config)# no vnic enable
  ```

- When using vNICs with FCoE, the FIP Snooping option for automatic VLAN creation is not recommended for use with the Emulex Virtual Fabric Adapter for IBM BladeCenter. Disable automatic VLAN creation when connecting the switch to an Emulex Virtual Fabric Adapter. (ID: 51529)
- Bandwidth metering drops excess packets when the configured limits on the vNIC pipe are reached. If using vNICs with an Emulex Virtual Fabric Adapter, any CEE Enhanced Transmission Selection will be ignored unless FCoE is enabled. If vNICs and FCoE are used on the Emulex Virtual Fabric Adapter, bandwidth metering will be used on vNIC pipes 1, 3, and 4, and CEE Enhanced Transmission Selection will be used on vNIC pipe 2 for FCoE traffic.
- When using vNICs for FCoE, the Emulex Virtual Fabric Adapter will establish an FCoE connection to the switch when port is added to an FCoE VLAN. This connection will be formed even if FIP Snooping or CEE are not enabled, and will remain up when FIP Snooping, CEE, or vNICs are disabled. To resolve this, disable and re-enable the affected internal port. (ID: 50904).
- The following vNIC features are not supported on Microsoft Windows servers:
  – vNIC teaming
  – VLAN configuration
- When you change the CEE configuration while vNIC traffic is passing through the switch, the switch may behave in an unpredictable manner, such as receiving IBP/CBP discards. If this happens, reboot the switch to overcome the situation. To avoid this scenario, shut down all the ports before making any CEE-related configuration changes. (ID: 57414)
- The vNIC thread is used only in the stack mode. It is used to communicate the failover state and failover disabled ports from master to back. The thread will only increase when the Master goes down and the Backup takes over. It will not increase when sending and receiving vNIC TLVs through DCBX/LLDP. (ID: 58851)