BLADEOS™

Release Notes

BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter®

Version 6.3

Part Number: BMD00193, May 2010

BLADE NETWORK TECHNOLOGIES

2350 Mission College Blvd.
Suite 600
Santa Clara, CA 95054
www.bladenetwork.net

Release Notes

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

These release notes provide the latest information regarding BLADEOS 6.3 for the BNT Virtual Fabric 10Gb Switch Module. This supplement modifies information found in the complete documentation:

- *BLADEOS 6.3 Application Guide* for the BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter
- *BLADEOS 6.3 Command Reference* for the BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter
- *BLADEOS 6.3 ISCLI Reference* for the BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter
- *BLADEOS 6.3 BBI Quick Guide* for the BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter
- BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter, *Installation Guide*

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

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

Please keep these release notes with your product manuals.

Hardware Support

BLADEOS 6.3 software is supported only on the BNT Virtual Fabric 10Gb Switch Module (IBM model name 46C7191) for IBM BladeCenter. The Virtual Fabric 10Gb 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 10Gb Switch Module Faceplate

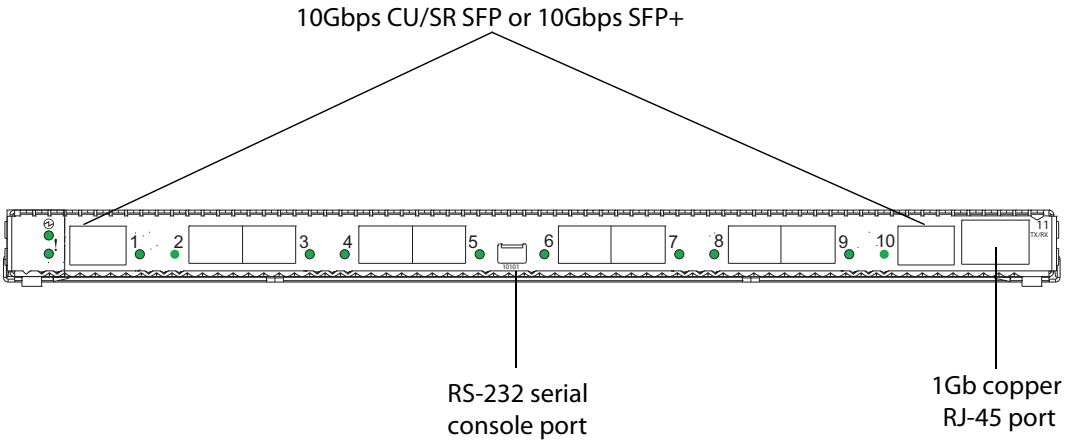

The VFSM has the following port capacities:

- 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 ships with the switch, and 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 available for your VFSM, go to:

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

From the BLADEOS CLI, use the `/boot/cur` command to determine the current software version.

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

- Place the new image onto a FTP or TFTP server on your network, or on a local computer.
- Transfer the new image to your switch.
- Select the new software image to be loaded into switch memory the next time the switch is reset.

Loading New Software to Your Switch

The switch can store up to two different software images, called `image1` and `image2`, as well as 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 9).

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

- The image and boot software loaded on a FTP or TFTP server on your network. For example:
  - Boot file: `GbESM-24-10G-6.3.1.0_Boot.img`
  - Image file: `GbESM-24-10G-6.3.1.0_OS.img`

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
- The name of the new software image or boot file

Note – The DNS parameters must be configured if specifying hostnames.

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

Using the BLADEOS CLI

1. At the `Boot Options#` prompt, enter:

```
Boot Options# 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.

6. The system then prompts you to confirm your request.

   Once confirmed, the software will load into the switch.

7. 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 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.

Using 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 (usually `tftpboot`).

4. Enter your username and password for the server, if applicable.

```
User name: {<username>|<Enter>}
```

5. The system prompts you to confirm your request.

Once confirmed, the software will load into the switch.

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

```
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.
```

Using the 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.

Using 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, CRT, 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.

New and Updated Features

BLADEOS 6.3 for BNT Virtual Fabric 10Gb Switch Module (VFSM) has been updated to include new and enhanced features in support of Virtualization and Fibre Channel over Ethernet.

The list of features below summarizes the updated features. For more detailed information about configuring VFSM features and capabilities, refer to the complete BLADEOS 6.3 documentation as listed on page 3.

Stacking

A *stack* is a group of up to eight Virtual Fabric 10Gb Switch Module switches with BLADEOS that work together as a unified system. A stack has the following properties, regardless of the number of switches included:

- The network views the stack as a single entity.
- The stack can be accessed and managed as a whole using standard switch IP interfaces.
- Once the stacking links have been established (see below), the number of ports available in a stack equals the total number of remaining ports of all the switches that are part of the stack.
- The number of available IP interfaces, VLANs, Trunks, Trunk Links, and other switch attributes are not aggregated among the switches in a stack. The totals for the stack as a whole are the same as for any single switch configured in stand-alone mode.

Stacking Requirements

Before BLADEOS switches can form a stack, they must meet the following requirements:

- All switches must be the same model (Virtual Fabric 10Gb Switch Module).
- Each switch must be installed with BLADEOS, version 6.3 or later. The same release version is not required, as the Master switch will push a firmware image to each differing switch which is part of the stack.
- The recommended stacking topology is a bidirectional ring. To achieve this, two external 10Gb Ethernet ports on each switch must be reserved for stacking.By default, the first two 10Gb Ethernet ports are used.
- The cables used for connecting the switches in a stack carry low-level, inter-switch communications as well as cross-stack data traffic critical to shared switching functions. Always maintain the stability of stack links in order to avoid internal stack reconfiguration.

Stacking Limitations

The VFSM with BLADEOS 6.3 can operate in one of two modes:

- Default mode, which is the regular stand-alone (or non-stacked) mode.
- Stacking mode, in which multiple physical switches aggregate functions as a single switching device.

When in stacking mode, the following stand-alone features are not supported:

- Active Multi-Path Protocol (AMP)
- sFlow port monitoring
- Uni-Directional Link Detection (UDLD)
- Port flood blocking
- BCM rate control
- Link Layer Detection Protocol (LLDP)
- Protocol-based VLANs
- RIP
- OSPF and OSPFv3
- IPv6
- Virtual Router Redundancy Protocol (VRRP)
- Loopback Interfaces
- Router IDs
- Route maps
- Border Gateway Protocol (BGP)
- MAC address notification
- Static MAC address adding
- Static multicast
- MSTP
- IGMP Relay and IGMPv3
- Converge Enhanced Ethernet (CEE)
- Fibre Channel over Ethernet (FCOE)
- Virtual NICs

Note – In stacking mode, switch menus and commands for unsupported features may be unavailable, or may have no effect on switch operation.

VMready

The switch's VMready software makes it *virtualization aware*. Servers that run hypervisor software with multiple instances of one or more operating systems can present each as an independent *virtual machine* (VM) with its own applications. With VMready, the VFSM automatically discovers virtual machines (VMs), virtual switches, and VM NICs (collectively known as virtual entities or VEs), and can distinguish between regular VMs, Service Console Interfaces, and Management Interfaces. BLADEOS 6.3 supports up to 2048 VEs.

VEs may be placed into VM groups on the switch to define communication boundaries: VEs in a given VM group are permitted to communicate with each other, while VEs in different groups are not. VM groups also allow the configuration of group-level settings, such as virtualization policies and ACLs.

The administrator can pre-provision VEs by adding the MAC addresses of potential VEs to a VM group. When a VE with a pre-provisioned MAC address becomes connected to the switch, the switch will automatically apply the appropriate group membership configuration.

The VFSM with VMready detects the migration of VEs across different hypervisors. As VEs move, the VFSM NMotion™ feature automatically moves the appropriate network configuration as well. NMotion gives the switch the ability to maintain assigned group membership and associated policies (such as VLAN Maps and VM policy bandwidth control) when a VE moves to a different port on the switch.

VMready also works with VMware's Virtual Center (vCenter) for advanced VE management. By connecting with the vCenter, the switch can obtain information about distant VEs, push VM configuration profiles to the VEs in distributed VM groups, and enhance VE migration.

VMready is configured from the Virtualization menu, available with the following CLI command:

```
# /cfg/virt
```

Note – The VMready and vNIC features are not supported simultaneously on the same ports.

VLAN Maps

A VLAN map (VMAP) is an Access Control List (ACL) that can be assigned to a VLAN rather than to a switch port as with regular ACLs. In a virtualized environment, VMAPs allow you to create traffic filtering and metering policies that are associated with a VM group VLAN, allowing ACLs to follow VMs as they migrate between hypervisors.

VMAPs are configured from the ACL menu, available with the following CLI command:

```
# /cfg/acl/vmap <1-128>
```

BLADEOS 6.3 supports up to 128 VMAPs. Individual VMAP filters are configured in the same fashion as regular ACLs, except that VLANs cannot be specified as a filtering criteria since the filter is explicitly assigned to a VLAN by nature.

Once a VMAP filter is created, it can be assigned or removed using the following commands:

- For a regular VLAN:

```
/cfg/l2/vlan <VLAN ID>/vmap {add|rem} <VMAP ID> [intports|extports]
```

- For a VM group:

```
/cfg/virt/vmgroup <ID>/vmap {add|rem} <VMAP ID> [intports|extports]
```

When the optional `intports` or `extports` parameter is specified, the action to add or remove the VMAP is applied for only the switch server ports (`intports`) or uplink ports (`extports`). If omitted, the operation will be applied to all ports in the associated VLAN or VM group.

Note – VMAPs have a lower priority than port-based ACLs. If both an ACL and a VMAP match a particular packet, both filter actions will be applied as long as there is no conflict. In the event of a conflict, the port ACL will take priority.

OSPFv3

BLADEOS supports the Open Shortest Path First (OSPF) version 2 and version 3 routing protocols. The OSPFv3implementation conforms to the OSPF version 2 specifications detailed in Internet RFC 1583, and OSPF version 3 specifications in RFC 2328 Appendix G.2 and RFC 2740.

OSPF version 3 is based on OSPF version 2, but has been modified to support IPv6 addressing. In most other ways, OSPFv3 is similar to OSPFv2: They both have the same packet types and interfaces, and both use the same mechanisms for neighbor discovery, adjacency formation, LSA flooding, aging, and so on. The administrator should be familiar with the OSPFv2 concepts covered in the preceding sections of this chapter before implementing the OSPFv3 differences as described in the following sections.

Although OSPFv2 and OSPFv3 are very similar, they represent independent features on the VFSM. They are configured separately, and both can run in parallel on the switch with no relation to one another, serving different IPv6 and IPv4 traffic, respectively.

OSPFv3 command paths are located as follows:

```
>> # /cfg/l3/ospf3                    (OSPFv3 config menu)
>> # /info/l3/ospf3                   (OSPFv3 information menu)
>> # /stats/l3/ospf3                  (OSPFv3 statistics menu)
```

OSPFv3 has numerous improvements that increase the protocol efficiency in addition to supporting IPv6 addressing. These improvements change some of the behaviors in the OSPFv3 network and may affect topology consideration, but have little direct impact on configuration. For example:

- Addressing fields have been removed from Router and Network LSAs.
- Link-local flooding scope has been added, along with a Link LSA. This allows flooding information to relevant local neighbors without forwarded it beyond the local router.
- Flexible treatment of unknown LSA types to make integration of OSPFv3 easier.

BLADEOS 6.3 does not currently support the following OSPFv3 features:

- Multiple instances of OSPFv3 on one IPv6 link.
- Authentication via IPv6 Security (IPsec)

Active MultiPath Protocol

Active MultiPath Protocol (AMP) allows you to connect three switches in a loop topology, and load-balance traffic across all uplinks (no blocking). When an AMP link fails, upstream communication continues over the remaining AMP link. Once the failed AMP link re-establishes connectivity, communication resumes to its original flow pattern.

Each AMP group contains two aggregator switches and one access switch. Aggregator switches support up to 22 AMP groups. Each access switch supports only one AMP group.

The VFSM can be used as an AMP access switch only.

Layer 3 routing protocols are not supported on AMP-configured switches.

Use the following command to access the AMP configuration menu:

```
>> # /cfg/l2/amp
```

Port Trunk Hashing Enhancements

Network traffic is statistically distributed among the ports in a trunk group using an enhanced, RTAG7 hash process. To improve traffic distribution, more bits from more frame attributes are used in the hash. The VFSM now supports the following hashing options, which can used in any combination:

- Layer 2 Source MAC address (enabled by default)
- Layer 2 Destination MAC address (enabled by default)
- Layer 3 Source IP address (enabled by default)
- Layer 3 Destination IP address (enabled by default)
- Ingress port number (enabled by default)
- Layer 4 (TCP, UDP, etc.) port information (enabled by default)

Note – For MPLS packets, Layer 4 port information is excluded from the hash calculation. Other IP fields are used, along with the first two MPLS labels.

Trunk hashing options can be configured using the following command path:

```
>> # /cfg/l2/thash
```

SM IPv4 and IPv4 Configuration Extensions

BLADEOS 6.3 supports EIPAA extensions to the chassis' Advanced Management Module (AMM) that provide more IP address configuration options. The following IP addressing options are now available on the AMM:

- Automatic IPv4 address configuration via DHCPv4
- Automatic IPv6 address configuration via DHCPv6 or Stateless Auto-Configuration (SAC)
- Manual IPv4 and IPv6 static address configuration
- IPv4 and IPv6 address change notification

Management Interface Connection

The switch management interface (for IPv4 or IPv6) will function only if the AMM has an IP address in the same network as the switch.

Also, use caution when disabling the IPV4 management interface from the AMM. This removes the IPv4 routes to the switch, severing communication to the IPv4 management interface on the switch.

DHCPv6 Behavior

The DHCPv6 server provides an IPv6 address only. The server does not provide an IPv6 prefix length or default gateway in the DHCP offer. Instead, this information is expected from an upstream router in a Router Advertisement (RA) packet, and will override the default settings.

The DHCPv6 IP address assignment process is summarized as follows.

1. The DCHP server sends an offer with the base IPv6 address.
2. The switch installs the address using a default prefix length of 64.
3. An upstream router sends an RA that may specify a different prefix length.
4. When this RA is received, the switch will modify the IPv6 address to new prefix length.
5. The switch will also add the source IP address of the RA to the default routers list.

Fibre Channel over Ethernet

BLADEOS 6.3 now supports the following Fibre Channel Forwarders (FCFs):

- QLogic Virtual Fabric Extension Module for IBM BladeCenter
- Cisco Nexus 5000 Series Switches

IGMP Group Capacity

BLADEOS 6.3 supports IGMP groups differently than earlier releases:

The VFSM now supports a maximum of 2048 IGMP entries, on a maximum of 1024 VLANs.

When the switch is in stacking mode, one IGMP entry is allocated for each unique join request, based on the combination of the port, VLAN, and IGMP group address. If multiple ports join the same IGMP group, they require separate IGMP entries, even if using the same VLAN.

In stand-alone (non-stacking) mode, one IGMP entry is allocated for each unique join request, based on the VLAN and IGMP group address only (regardless of the port). If multiple ports join the same IGMP group using the same VLAN, only a single IGMP entry is used.

Clause 73 Backplane Auto-negotiation

Internal ports may be now be configured to use IEEE 802.3 Clause 73 for high-speed backplane Ethernet negotiation. When enabled on a port (and on the connecting device), Clause 73 is used to advertise available link modes and automatically select a mode that takes maximum advantage of shared capabilities. If the connecting device does not support Clause 73, the switch port will next attempt legacy Clause 37 negotiation, and then Clause 37 BAM.

The Clause 73 option can be enabled on internal ports using the following command:

```
# /cfg/port <port>/gig/cl73 enable
```

Other Features

BLADEOS 6.3 now also supports the following features:

- vNIC support has been added to the SNMP MIB.
- BPDU Guard now supports PVSRT/RSTP/MSTP.
- 802.1X RADIUS VLAN assignment for authenticated ports.
- A port ErrDisable option has been added, allowing switch ports to be automatically disabled when abnormal port conditions are detected, and optionally reenabled after a configurable time period.

Supplemental Information

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

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 1 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.

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 *BLADEOS 6.3 Application Guide*.

Note – The external uplink ports (EXT*x*) 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.

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.

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.

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.

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.
- Manually add all associated VLANs into a single Spanning Tree Group.

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 *BLADEOS 6.3 Command Reference*.

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.

VRRP Configuration

Although the Virtual Router Redundancy Protocol (VRRP) standard permits up to 255 virtual router instances, the BLADEOS 6.3 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.

Command Replacements

STP Edge and Link

The following commands have been moved in the CLI:

Old CLI Path	New Path
`/cfg/l2/stg <#>/port <#>/edge`	`/cfg/port <#>/stp/edge`
`/cfg/l2/stg <#>/port <#>/link`	`/cfg/port <#>/stp/link`
`/cfg/l2/mrst/cist/port <#>/edge`	`/cfg/port <#>/stp/edge`
`/cfg/l2/mrst/cist/port <#>/link`	`/cfg/port <#>/stp/link`

The equivalent ISCLI config-if mode commands (under `interface port <#>`) have also been moved:

Old ISCLI Path	New Path
`spanning-tree stp <#> edge`	`spanning-tree edge`
`spanning-tree stp <#> link`	`spanning-tree link-type`
`spanning-tree mstp cist edge`	`spanning-tree edge`
`spanning-tree mstp cist link-type`	`spanning-tree link-type`

When the switch boots with BLADEOS 6.3 software, any prior configuration in the active configuration block will automatically be updated to use the new command paths.

BPDU Guard

BLADEOS 6.3, BPDU Guard is now configured on a per-port basis, rather than globally. The following commands have been moved in the CLI:

Old CLI Path	New Path
`/cfg/l2/bpduguard {ena|dis}`	`/cfg/port <#>/bpduguard {ena|dis}`

The equivalent ISCLI Global Configuration Mode command has moved to the Interface Port Configuration Mode:

Old ISCLI Path	New Path
`(config)# spanning-tree bpdu-guard`	`(config)# interface port <x>` `(config-if)# bpdu-guard`

During upgrade to BLADEOS 6.3, any current BPDU Guard configuration will be automatically converted. If the prior global BPDU setting was enabled, that setting will be removed, and BPDU Guard will be applied to all ports where the port fast-forward feature is enabled. This conversion will also be applied when loading configuration scripts from previous releases.

Note – If you want to apply BPDU Guard to different ports than configured with port fast-forward, after any conversion, review the interface port configuration and make any appropriate changes manually.

Known Issues

This section describes known issues for BLADEOS 6.3 on the BNT Virtual Fabric 10Gb Switch Module.

Software Upgrade Issues

- 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)
- 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)

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.

Access Control Lists

- 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.

Link Aggregation Control Protocol

If a static trunk on a VFSM is connected to another VFSM with LACP configured (but no active LACP trunk), the `/info/l2/trunk` 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.

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.

QoS and Trunking

When you assign an ACL (or ACL Group) to one port in a trunk, BLADEOS 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.

Solution: Manually assign each ACL or ACL Group to all ports in a trunk.

RIP MIBs

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

BLADEOS does not support the standard RIPv2 MIB as described in RFC 1724. Use the `rip2Cfg` MIB to configure RIPv1 and RIPv2 through SNMP.

Trunk and Link 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
```

Solution: Remove the loop to resolve this misconfiguration.

Browser Based Interface

- 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.

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.

Blocking Egress Traffic

Access Control Lists (ACLs) which are configured to match both a destination MAC address and an egress port fail to act when the matching packets are encountered. As a result, ACLs cannot be used to block traffic exiting specific ports for specific static multicast MAC addresses.

Solution: Instead of using an ACL to block the traffic, configure a static multicast route that includes all ports other than those you wish to block. Consider an example where you wish to block port EXT1 for DMAC `01:02:03:04:05:FF` on the default VLAN (VLAN 1). In this case, you would add a multicast route that includes all ports except EXT1. For example:

```
# /cfg/l2/fdb/mcast/add <Destination MAC> <VLAN> <list of ports or ranges to allow>
   -or-
# /cfg/l2/fdb/mcast/add 01:02:03:04:05:FF 1 INT1-INT14 EXT2-EXT10
```

Changing Port 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.

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.

TACACS+ Passwords

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.

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.

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.

vNICs 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:

```
>> # /cfg/l2/lldp/on
>> # apply
>> # save
>> # <trial commands...>
>> # revert apply
```

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:

```
>> # revert apply
>> # /cfg/virt/vnic/on
>> # apply
>> # off
>> # apply
```

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:

```
>> # /oper/virt/vmware/scan
```

SNMP MIB Browser Errors

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
```

Active MultiPath Protocol

For proper AMP operation, all access switches should be configured with a higher priority value (lower precedence) than the aggregators. Otherwise, unexpected AMP keep-alive packets may be forwarded from one aggregator switch to the other, even when its AMP group is disabled.
(ID: 37310)