

BLADEOS™

Release Notes

BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter®

Version 6.1.2

Part Number: BMD00149, December 2009

BLADE
NETWORK TECHNOLOGIES

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

Copyright © 2009 BLADE Network Technologies, Inc., 2350 Mission College Blvd. Suite 600, Santa Clara, California, 95054, USA. All rights reserved. Reference number: BMD00149

This document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this document may be reproduced in any form by any means without prior written authorization of BLADE Network Technologies, Inc. Documentation is provided “as is” without warranty of any kind, either express or implied, including any kind of implied or express warranty of non-infringement or the implied warranties of merchantability or fitness for a particular purpose.

U.S. Government End Users: This document is provided with a “commercial item” as defined by FAR 2.101 (Oct. 1995) and contains “commercial technical data” and “commercial software documentation” as those terms are used in FAR 12.211-12.212 (Oct. 1995). Government End Users are authorized to use this documentation only in accordance with those rights and restrictions set forth herein, consistent with FAR 12.211- 12.212 (Oct. 1995), DFARS 227.7202 (JUN 1995) and DFARS 252.227-7015 (Nov. 1995).

BLADE Network Technologies, Inc. reserves the right to change any products described herein at any time, and without notice. BLADE Network Technologies, Inc. assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by BLADE Network Technologies, Inc. The use and purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of BLADE Network Technologies, Inc.

BLADE OS and BLADE are trademarks of BLADE Network Technologies, Inc. in the United States and certain other countries. VMware is a registered trademark or trademark of VMware, Inc. in the United States and/or other jurisdictions. Any other trademarks appearing in this manual are owned by their respective companies.

Originated in the USA.

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 6.1 for the BNT Virtual Fabric 10Gb Switch Module. This supplement modifies information found in the complete documentation:

- *BLADEOS 6.1 Application Guide* for the BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter
- *BLADEOS 6.1 Command Reference* for the BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter
- *BLADEOS 6.1 ISCLI Reference* for the BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter
- *BLADEOS 6.1 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 6.1 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.

To download 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
 - Boot file: `GbESM-24-10G-6.1.2.0_Boot.img`
 - Image file: `GbESM-24-10G-6.1.2.0_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 above 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>}
```

6. The system prompts you to confirm your request.

Once confirmed, the software will load into the switch.

7. When loading is complete, enter the following command at the Boot Options# prompt:

```
Boot Options# image
```

8. 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.

New and Updated Features

BLADEOS 6.1 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 6.1 documentation as listed on [page 3](#).

Fiber Channel over Ethernet (FCoE)

FCoE is an effort to converge two of the different physical networks in today's data centers. It allows Fibre Channel traffic (such as that commonly used Storage Area Networks, or SANs) to be transported over Ethernet links typically used for high-speed Local Area Networks (LANs). This provides an evolutionary approach toward network consolidation, allowing Fibre Channel equipment and tools to be retained, while leveraging cheap, ubiquitous Ethernet networks for growth.

The BNT Virtual Fabric 10Gb Switch Module (VFSM) with BLADEOS 6.1 software is compliant with the INCITS T11.3, FC-BB-5 FCoE specification.

The following are required for implementing FCoE using the VFSM:

- The VFSM must be connected to the Fibre Channel network via FCF. The QLogic Virtual Fabric Extension Module for IBM BladeCenter is the first FCF currently supported.
- If the FCF is a blade chassis bridge module (such as the supported Qlogic FCF), the connection between the VFSM and the bridge module requires either two or four VFSM external ports.
- For internal VFSM ports participating in FCoE, the connected blade server must use the supported FCoE CNA. The QLogic CNA is currently the first CNA supported for this purpose.
- CEE must be turned on. When CEE is on, DCBX, PFC, ETS, and FIP Snooping are enabled and configured with default FCoE settings. These features may be reconfigured, but must remain enabled in order for FCoE to function.

Converged Enhanced Ethernet (CEE)

CEE refers to a set of IEEE standards designed to allow different physical networks with different data handling requirements to be converged together, simplifying management, increasing efficiency and utilization, and leveraging legacy investments without sacrificing evolutionary growth.

CEE standards were developed primarily to enable Fibre Channel traffic to be carried over Ethernet networks. This required enhancing the existing Ethernet standards to make them lossless on a per-priority traffic basis, and to provide a mechanism to carry converged (LAN/SAN/IPC) traffic on a single physical link. Although CEE standards were designed with FCoE in mind, they are not limited to FCoE installations. CEE features can be utilized in traditional LAN (non-FCoE) networks to provide lossless guarantees on a per-priority basis, and to provide efficient bandwidth allocation based on application needs.

Note – By default, CEE is turned off. Turning CEE on will automatically change some 802.1p QoS and 802.3x standard flow control settings on the VFSM. Read the Application Guide carefully to determine whether you will need to take action to reconfigure expected settings.

FCoE Initialization Protocol (FIP) Snooping

FIP snooping is an FCoE feature. In order to enforce point-to-point links for FCoE traffic outside the regular Fibre Channel topology, Ethernet ports used in FCoE can be automatically and dynamically configured with Access Control Lists (ACLs).

Using FIP snooping, the VFSM examines the FIP frames normally exchanged between the FCF and ENodes to determine information about connected FCoE devices. This information is used to create narrowly tailored ACLs that permit expected FCoE traffic to and from confirmed Fibre Channel nodes, and deny all other undesirable traffic.

For FCoE, the VFSM must be connected to at least one FCF bridge module. The Qlogic Virtual Fabric Extension Module for IBM BladeCenter is the first FCF currently supported.

For internal VFSM ports participating in FCoE, the connected blade server must use the supported FCoE CNA. The Qlogic CNA is currently the first CNA supported for this purpose.

Priority-Based Flow Control (PFC)

Priority-based Flow Control (PFC) is defined in IEEE 802.1Qbb. PFC extends the IEEE 802.3x standard flow control mechanism, allowing the switch to pause some classes of traffic on the port while other traffic on the port continues, based on the 3-bit 802.1p priority value in the 802.1Q VLAN tag.

PFC is vital for FCoE environments, where SAN traffic must remain lossless and should be paused during congestion, while LAN traffic on the same links should be delivered with “best effort” characteristics.

When CEE is turned on, PFC is enabled on priority value 3 by default. Optionally, the administrator can also enable PFC on one other priority value, providing lossless handling for another traffic type, such as for a business-critical LAN application.

Note – For any given port, only one flow control method can be implemented at any given time: either PFC or standard IEEE 802.3x flow control.

Enhanced Transmission Selection (ETS)

ETS is defined in IEEE 802.1Qaz. ETS provides a method for allocating port bandwidth based on 802.1p priority values in the VLAN tag. Using ETS, different amounts of link bandwidth can be specified for different traffic types (such as for LAN, SAN, and management).

ETS is an essential component in a CEE environment that carries different types of traffic, each of which is sensitive to different handling criteria, such as Storage Area Networks (SANs) that are sensitive to packet loss, and LAN applications that may be latency-sensitive. In a single converged link, such as when implementing FCoE, ETS allows SAN and LAN traffic to coexist without imposing contrary handling requirements upon each other.

When CEE is turned on, the default values for ETS configuration as follows:

Figure 2 Default ETS Priority Groups

Typical Traffic Type	802.1p Priority	PGID	Bandwidth Allocation
LAN	0	0	10%
LAN	1		
LAN	2		
SAN	3	1	50%
Latency-Sensitive LAN	4	2	40%
Latency-Sensitive LAN	5		
Latency-Sensitive LAN	6		
Latency-Sensitive LAN	7		

The administrator may reassign 802.1p priority values among up to 8 priority groups (PGIDs), and allocate a percentage of the switch's link bandwidth to each PGID.

PGID 15 (unconfigured by default) is available as a strict priority group. It is typically used for critical traffic, such as network management. After traffic assigned to PGID 15 is served, any remaining link bandwidth is shared among the other groups, according to the configured bandwidth allocation.

Data Center Bridging Exchange Protocol (DCBX)

DCBX is a vital element of CEE. DCBX allows peer CEE devices to exchange information about their advanced capabilities. Using DCBX, neighboring network devices discover their peers, negotiate peer configurations, and detect misconfigurations.

For normal operation of any FCoE implementation on the VFSM, DCBX must remain enabled on all ports participating in FCoE.

DCBX also allows CEE devices to negotiate with each other for the purpose of automatically configuring advanced CEE features such as PFC and ETS. The administrator can determine which CEE feature settings on the switch are communicated to and matched by CEE neighbors, and also which CEE feature settings on the switch may be configured by neighbor requirements.

Virtual NIC Support

Some NICs, such as the Emulex Virtual Fabric Adapter for IBM BladeCenter, can virtualize their resources. This helps resolve issues caused by limited NIC slot availability. By virtualizing a 10Gbps NIC, it can be divided into multiple logical instances known as virtual NICs (vNICs). Each vNIC appears as a regular, independent NIC to the server operating system or hypervisor, with each vNIC using a definable portion of the physical NIC's overall bandwidth.

BLADEOS 6.1 supports the Emulex Virtual Fabric Adapter (VFA) for IBM BladeCenter to provide the following vNIC features:

- The vNIC feature is supported only on switches installed in bay 7 and 9 in the blade chassis.
- Up to four vNICs are supported on each internal switch port.
- vNICs can be grouped together, along with regular internal ports, external ports, and trunk groups, to define vNIC groups for enforcing communication boundaries.
- In the case of a failure on the external uplink ports associated with a vNIC group, the switch can signal affected vNICs for failover while permitting other vNICs connected on the physical port to continue operation.
- Each vNIC can be independently allocated a symmetric percentage of the 10Gbps bandwidth on the link (from NIC to switch, and from switch to NIC).
- The VFSM can be used as the single point of vNIC configuration.

Basic vNIC Configuration

To use vNICs, the following basic configuration is required:

- On the server, be sure the VFA is operating in vNIC mode (the default).
- On the switch, enable the vNIC feature, and set the appropriate ports to vNIC mode, then place vNICs and ports into groups to enforce communication boundaries.

Consider the following example configuration:

Figure 3 A Simple vNIC Group

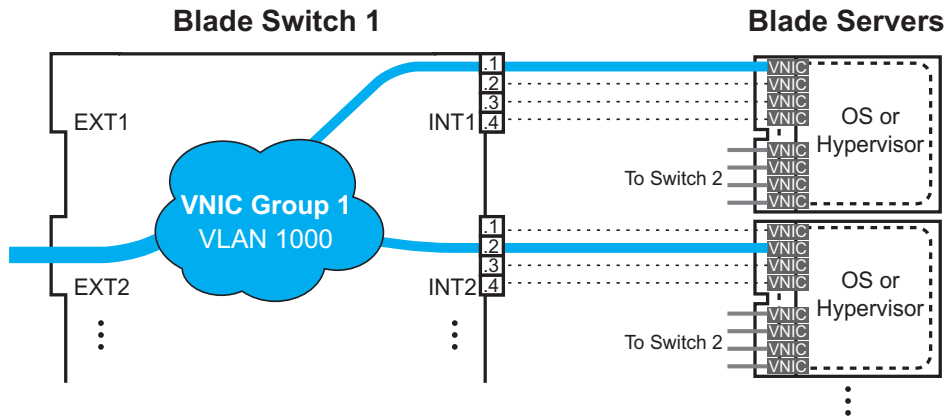


Figure 3 has the following vNIC network characteristics:

- vNIC group 1 is comprised of vNIC pipes INT1.1 and INT2.2, and external uplink port EXT2.
- To enforce vNIC communication boundaries, the VNIC group uses VLAN 1000. The group VLAN is an “outer” tag, which is used in addition to any regular VLAN tag assigned to packets by the network, server, or hypervisor. The outer VLAN is used only between the VFSM and the VFA and is removed by the VFA before packets reach the server OS or hypervisor, or by the switch before packet egress any port that does not contain a vNIC.
- Enabled vNICs are allocated 50% (5 Gbps) of the available bandwidth on their ports (increased from the 25% default value).
- All remaining vNIC pipes on the switch are disabled (by default) and are automatically allocated 0 bandwidth.
- vNIC-aware failover (optional) is enabled for the group.

This simplified example can be configured using the procedure that follows. However, for more detailed instructions on configuring vNIC options, see the *BLADEOS 6.1 Application Guide* and *BLADEOS 6.1 Command Reference* for the BNT Virtual Fabric 10Gb Switch Module for IBM BladeCenter.

The following procedure is used to configure the vNIC example:

1. On the blade server, be sure the VFA is operating in vNIC mode.

A VFA is capable of two modes:

- Virtual NIC (vNIC) mode

This is the default mode on the VFA. In vNIC mode, the blade server operating system or hypervisor configures the VFA as eight separate ethernet devices (four per physical port).

The VFA must operate in this mode to properly use the vNIC features on the switch.

- Physical NIC (pNIC) mode

If vNICs are not required, the VFA may be operated as a dual-port physical NIC.

In this mode, the switch will treat the VFA as a standard (non-virtual) NIC with no vNIC features.

To install the VFA or set operation mode, see the instructions in the *Emulex Virtual Fabric Adapter Installation Guide*.

2. On the VFSM, enable the vNIC feature.

The vNIC feature and options can be configured using the standard switch management options: the BLADEOS CLI, the ISCLI, or the Browser-based Interface (BBI).

In the BLADEOS CLI, the vNIC feature can be enabled with the following command:

```
>> # /cfg/virt/vnic/on
```

3. Configure the appropriate switch ports for vNIC operation:.

By default, the VFSM considers all internal switch ports to be connected to standard (non-virtual) NICs. To enable vNICs on specific ports, and allocate vNIC bandwidth per this example, use the following commands:

```
>> vNIC Global Configuration# port INT1           (Select port INT1)
>> Port INT1 vNICs# vnic 1                       (Select vNIC 1 on the port)
>> vNIC INT1.1# ena                               (Enable the vNIC pipe)
>> vNIC INT1.1# bw 50                             (Allow 60% egress bandwidth)
>> vNIC INT1.2# /cfg/virt/vnic/port INT2         (Select port INT2)
>> Port INT2 vNICs# vnic 2                       (Select vNIC 2 on the port)
>> vNIC INT2.2# ena                               (Enable the vNIC pipe)
>> vNIC INT2.2# bw 50                             (Allow 60% egress bandwidth)
```

Note – vNIC operation may only be enabled on ports connected to peers that support the feature. Enabling the vNIC operation on incompatible ports will display a warning message.

4. To enforce vNIC communication boundaries, add ports, trunks and virtual pipes to the vNIC group.

vNICs can be grouped together, along with internal and external switch ports and trunks, into vNIC groups. Each vNIC group is essentially a separate virtual network within the switch. Elements within a vNIC group have a common logical function and can communicate with each other, while elements in different vNIC groups are separated. The group in this example is defined with the following commands:

```
>> vNIC INT2.2# /cfg/virt/vnic/vnicgrp 1      (Select vNIC group 1)
>> vNIC Group 1# vnicvlan 1000             (Specify the VLAN)
>> vNIC Group 1# addvnic INT1.1            (Add vNIC pipes to the group)
>> vNIC Group 1# addvnic INT2.2
>> vNIC Group 1# addport EXT1
>> vNIC Group 1# failover ena              (Enable vNIC failover for the group)
>> vNIC Group 1# ena                       (Enable the vNIC group)
```

Note – Once VLAN 1000 is configured for vNIC groups, it will not be available for configuration in the regular VLAN menus (/cfg/12/vlan).

5. Apply and save the configuration.

vNIC Interface Names

When in vNIC mode, the VFA presents eight vNICs to the OS or hypervisor (four for each of the two physical NIC ports). Each vNIC is identified in the OS or hypervisor with a different vNIC function number (0-7). vNIC function numbers correlate to vNIC IDs on the switch as follows:

Table 1 vNIC ID Correlation

Virtual Fabric Adapter		Virtual Fabric Switch Module		
PCIe Function	VFA Port	Chassis Bay	vNIC Pipe	vNIC ID
0	0	HSSM 7	1	INT _x .1
2	0	HSSM 7	2	INT _x .2
4	0	HSSM 7	3	INT _x .3
6	0	HSSM 7	4	INT _x .4
1	1	HSSM 9	1	INT _x .1
3	1	HSSM 9	2	INT _x .2
5	1	HSSM 9	3	INT _x .3
7	1	HSSM 9	4	INT _x .4

In [Table 1 on page 15](#), the *x* in the vNIC ID represents the internal switch port to which the VFA port is connected. Each physical VFA port is connected to a different switch bay in the blade chassis.

vNIC Restricted Features

To enforce communication boundaries, vNIC groups are isolated from each other and from other segments on the switch through the use of an outer VLAN tag. As a result, some regular switching and routing functions (especially those that rely on the regular, inner VLAN tag) will not be performed for packets within a vNIC group. The following restrictions apply for vNIC traffic:

- Spanning Tree Protocol and per-VLAN Spanning Tree Protocol is not applicable.
- The inner VLAN tag will not be changed, added, or removed at packet egress.
- Per-VLAN IGMP snooping is performed on the outer vNIC VLAN tag only, and not the inner VLAN tag.
- The inner VLAN tag cannot be used to restrict multicast traffic to a subset of internal ports.
- Because inner VLAN tags are ignored, vNIC group elements assigned with different inner VLAN tags will not be isolated from one another. All ports in the vNIC group will receive multicast traffic for all member's inner VLANs.
- ACL filters will not match against the inner VLAN tag.
- Layer 3 switching and routing protocols between VLANs is not supported to or from vNIC groups.

Note – Although these features are not supported within vNIC groups, they still apply to traffic that is not part of vNIC groups. The VFSM does not prevent configuring these features for regular VLANs, even if those VLANs are also included within vNIC groups.

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.

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. This 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
```

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 6.1 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/12/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]
```

The optional `inport` or `extport` parameter can be specified to apply the action (add or remove the VMAP) for only the internal ports or the external ports within the VLAN. If omitted, the operation will be applied to all ports in the associated VLAN.

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.

Remote Monitoring

BLADEOS 6.1 now supports Remote Monitoring (RMON). RMON allows network devices to exchange network monitoring data, gather cumulative and history statistics for Ethernet interfaces, and create and trigger alarms for user-defined events.

An RMON management application can be used to access RMON MIB information on the VFSM, as described in RFC 1757. The switch supports RMON Group 1 (Statistics), Group 2 (History), Group 3 (Alarms), and Group 9 (Events).

RMON properties are configured globally in the RMON menu, and enabled on a per-port basis in the Port menu:

# / cfg/rmon	<i>(global RMON menu)</i>
-and-	
# / cfg/port <x>/ rmon	<i>(per-port RMON menu)</i>

Link Layer Detection Protocol

BLADEOS 6.1 now supports 802.1AB Link Layer Detection Protocol (LLDP). Using LLDP, the VFSM advertises port and link information to other LLDP-capable devices and accepts their LLDP advertisements for the purpose of discovering pertinent information about remote ports.

Switch port information and any remote device information is stored in a Managed Information Base (MIB). Higher-layer management tools may access the MIB to accumulate and report such information, and even to and discover configuration inconsistencies between systems on the same IEEE 802.3 LAN.

The LLDP configuration menu is accessed using the following CLI command:

/ cfg/12/1ldp

Uni-Directional Link Discovery Protocol

BLADEOS 6.1 now supports the Uni-Directional Link Discovery (UDLD) protocol, compliant with RFC 5171. UDLD operates at Layer 2 in conjunction with existing IEEE 802.3 Layer 1 fault detection mechanisms. It is used between peer devices to detect and disable unidirectional Ethernet links caused, for instance, by mis-wired cable strands, interface malfunctions, or media converter faults.

UDLD is configured on a per-port basis. It is disabled by default. The UDLD configuration menu is available using the following CLI command:

/ cfg/port <x>/ udld

Operation/Administration/Maintenance Protocol

BLADEOS 6.1 now supports IEEE 802.3ah Operation, Administration, and Maintenance (OAM) protocol. OAM allows the switch to detect faults on physical port links. Using OAM, if the Local Information that a port sends does not match the Remote Information received, the link is determined to be in an anomalous condition and is automatically disabled.

OAM is configured on a per-port basis. It is disabled by default. The OAM configuration menu is available using the following CLI command:

```
# /cfg/port <x>/oam
```

sFlow Monitoring

BLADEOS 6.1 now supports sFlow technology for monitoring traffic in data networks. The switch software includes an embedded sFlow agent which can be configured on a per-port basis to sample network traffic and provide continuous statistical report information to a central sFlow analyzer.

sFlow features are disabled by default, but may be configured using the following menu:

```
# /cfg/sys/sflow
```

Internal Loopback Interface

BLADEOS 6.1 now supports up to five loopback interfaces.

A loopback interface is an interface which is assigned an IP address, but is not associated with any particular physical port. The loopback interface is thus always available for higher layer protocols to use and advertise to the general network, regardless of which specific ports are in operation.

Loopback interfaces can be of benefit in a number of protocols, improving access to a switch, as well as increasing its reliability, security, and scalability. In addition, loopback interfaces can add flexibility and simplify management, information gathering, and filtering.

One example of this increased reliability is for OSPF to advertise a loopback interface as an interface route which will be available regardless of the status of individual physical links. This provides a higher probability that the routing traffic will be received and subsequently forwarded.

Further reliability and performance could be provided by configuring parallel BGP paths to a loopback interface on a peer device, which would result in improved load sharing.

Access and security can be improved through filtering. Incoming traffic can be filtered by rules that specify loopback interfaces as the only acceptable destination addresses.

Information gathering and filtering as well as management can potentially be simplified if protocols such as SNMP use loopback interfaces for receiving and sending trap and log type information.

The Loopback Interface configuration menu is accessed using the following CLI command:

```
# /cfg/13/loopif <loopback interface number (1-5)>
```

Rate Limiting

BLADEOS 6.1 now supports traffic rate limits for packets broadcast, multicast, and unknown unicast packets. For each port, the maximum number of packets permitted per second for each packet type can be specified. The following commands have been added to the Port menu (`/cfg/port <x>`) to support rate limiting:

- **brate** <value> | **dis** Broadcast limit, 0 to 2097151 packets per second, or no limit.
- **mrates** <value> | **dis** Multicast limit, 0 to 2097151 packets per second, or no limit.
- **drates** <value> | **dis** Unknown unicast limit, 0 to 2097151 packets per second, or no limit.

Hot Links

BLADEOS 6.1 now supports Hot Links. Hot Links provides basic link redundancy with fast recovery for network topologies that require Spanning Tree to be turned off.

Hot Links allows up to 25 triggers, each of which consists of a pair of layer 2 interfaces that may contain either an individual port or trunk. One interface is the Master, and the other is a Backup. While the Master interface is active and forwarding traffic, the Backup interface is placed in a standby state and blocks traffic. If the Master interface fails, the Backup interface becomes active and forwards traffic. Once the Master interface is restored, it transitions to the standby state and blocks traffic unless the Backup interface fails.

OSPF Enhancements

BLADEOS 6.1 includes multiple enhancements to the VFSM Open Shortest Path First (OSPF) implementation:

■ **Passive Interfaces**

OSPF interfaces may be configured as *passive*. Passive interfaces are included in the router's LSAs, but do not send LSAs, hello packets, or any other OSPF protocol information, and do not consider any received OSPF packets or information. Passive interfaces behave as stub networks, allowing OSPF routing devices to be aware of devices that do otherwise participate in OSPF (either because they do not support it, or because the administrator chooses to restrict OSPF traffic exchange or transit). The following command has been added:

```
# /cfg/13/ospf/if <x>/passive {enable|disable}
```

■ **Point-to-Point Networks**

For LANs that have only two OSPF routing agents (the VFSM and one other device), specifying the interfaces as part of a point-to-point network allows the switch to significantly reduce the amount of routing information it must carry and manage, thereby reducing convergence time and enhancing OSPF efficiency. The following command has been added:

```
# /cfg/13/ospf/if <x>/ptop {enable|disable}
```

■ **Sub-second timers**

To increase OSPF convergence speed, hello and dead timers for OSPF interfaces and virtual interfaces can now specified in milliseconds by adding “ms” to the number. For example:

```
# /cfg/13/ospf/if <x>/hello 200ms (200 milliseconds)
```

■ **Loopback Interface Address**

OSPF can now be configured to use the VFSM internal loopback address in advertising its Router ID.

Layer 2 Failover Enhancements

BLADEOS 6.1 includes multiple enhancements to the Layer 2 Failover feature to support advanced NIC teaming:

- In addition to the automatic monitoring triggers for trunk links, the switch software now supports new manual monitoring triggers. This allows you to define a list of ports and/or static or dynamic trunks to disable when a link failure threshold is reached on set of trigger ports and/or static or dynamic trunks.
- Up to two LACP keys can be used for each failover trigger. Previously, only one per trigger was supported.

Forwarding Database Enhancements

Configuration of the Forwarding Database (FDB) aging feature has been simplified. Because FDB aging required the same value configured in all Spanning Tree Groups (STGs), the per-STG aging parameters have been replaced with a single, global configuration command:

# /cfg/l2/stg <STG number>/brg/aging <value>	(old per-STG command)
-replaced by-	
# /cfg/l2/fdb/aging <value>	(new global command)

ISL Layer 2 Protocol Enhancements

BLADEOS 6.1 now supports additional Layer 2 protocols on Inter-Switch Link (ISL) ports:

- VRRP
- STP
- RSTP/MSTP
- 802.1Q VLAN Tagging
- 802.1p QoS/CoS
- 802.1X Port-Based Access Control (supported in force-auth mode only)
- ACLs

STP Fast Uplink Bridge Priority

With BLADEOS 6.1, the Fast Uplink Convergence bridge priority has been set to 65535.

CLI List and Range Inputs

For CLI commands that allow an individual item to be selected from within a numeric range, lists and ranges of items can now be specified. For example, the `/info/vlan` command permits the following options:

# /info/vlan	(show all VLANs)
# /info/vlan 1	(show only VLAN 1)
# /info/vlan 1,3,4095	(show listed VLANs)
# /info/vlan 1-20	(show range 1 through 20)
# /info/vlan 1-5,90-99,4090-4095	(show multiple ranges)
# /info/vlan 1-5,19,20,4090-4095	(show a mix of lists and ranges)

The numbers in a range must be separated by a dash: `<start of range>-<end of range>`

Multiple ranges or list items are permitted using a comma: `<range or item 1> , <range or item 2>`

Do not use spaces within list and range specifications.

Ranges can also be used to apply the same command option to multiple items. For example, to enable multiple ports with one command:

# /cfg/port 1-4/ena	(Enable ports 1 through 4)
---------------------	----------------------------

Chassis Internal Network

BLADEOS 6.1 now supports BladeCenter Chassis Internal Network (CIN). CIN provides internal connectivity between blade server ports and the internal Advanced Management Module (AMM) port. This allows blade server users to access the AMM via CLI, web-browser, or SNMP session, and allows the AMM to use services on the blades, such as LDAP, SMTP, DNS, and NTP.

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 2 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 6.1 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.

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.

Access Control List Filters

To accommodate VLAN-based ACLs (VMAP) the number of regular QoS ACLs has been reduced to 128.

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 6.1 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 Virtual Router IDs (VRIDs) between 1 and 255, the BLADEOS 6.1 implementation allows the configuration of VRIDs between 1 and 128, corresponding to the number of supported IP interfaces.

Known Issues

This section describes known issues for BLADEOS 6.1 on the BNT Virtual Fabric 10Gb Switch Module.

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/12/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/12/fdb/mcast/add <Destination MAC> <VLAN> <list of ports or ranges to allow>
-or-
# /cfg/12/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/12/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/12/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
```