BLADEOS™

Release Notes

RackSwitch™ G8124

Version 6.3

Part Number: BMD00188-A, April 2010

BLADE NETWORK TECHNOLOGIES

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

Release Notes

The RackSwitch G8124 (G8124) is an all 10Gb Ethernet rackable aggregation switch with unmatched line-rate Layer 2/3 performance. The G8124 uses a wire-speed, non-blocking switching fabric that provides simultaneous wire-speed transport of multiple packets at low latency on all ports.

These release notes provide the latest information regarding BLADEOS 6.3 for the RackSwitch G8124.

This supplement modifies information found in the complete documentation:

- *BLADEOS 6.3 Application Guide* for the RackSwitch G8124
- *BLADEOS 6.3 Command Reference* for the RackSwitch G8124
- *BLADEOS 6.3 ISCLI Reference* for the RackSwitch G8124
- *BLADEOS 6.3 Browser-Based Interface Quick Guide* for the RackSwitch G8124
- *RackSwitch G8124 Installation Guide*

The publications listed above are available from the support website:

http://www.bladenetwork.net/support_services_rackswitch.html

Please keep these release notes with your product manuals.

Hardware Support

This BLADEOS 6.3 software is supported only on the RackSwitch G8124. The G8124 is a high performance Layer 2-3 network switch.

The G8124 contains 24 ten Gigabit Small Form-factor, Pluggable (SFP+) slots and two 1Gb management ports. The 10Gb SFP+ slots can accept 1Gb copper transceivers, 10Gb optical transceivers, or Direct Attach Cables (DAC).

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

Figure 1 RackSwitch G8124 Front Panel

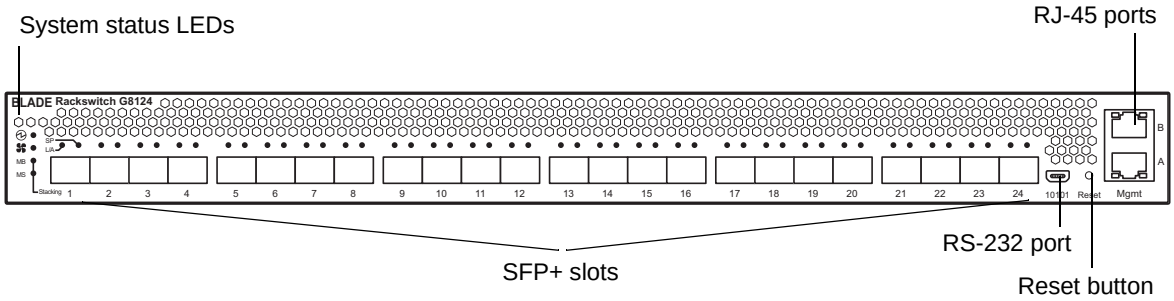

Transceivers

The G8124 accepts any of the following transceivers available from BLADE Network Technologies:

Table 1 Recommended SFP+ Transceiver

Part number	Description
BN-CKM-S-T	SFP Transceiver, 1000Base-T Copper
BN-CKM-S-SX	SFP Transceiver, 1000Base-SX Short Range Fiber
BN-CKM-S-LX	SFP Transceiver, 1000Base-LX Long Range Fiber
BN-CKM-SP-SR	SFP+ Transceiver, 10GBase-SR Short Range
BN-CKM-SP-LR	SFP+ Transceiver, 10GBase-LR Long Range

The G8124 accepts any SFP+ Direct Attach Cable that complies to the MSA specification.

Updating the Switch Software Image

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

http://www.bladenetwork.net/support_services_rackswitch.html

Click on software updates. Use the following command to determine the current software version:

```
RS 8124# show boot
```

To upgrade the software image on your switch, perform the following tasks:

- Load the new software image and boot image onto a TFTP server on your network.
- Transfer the new software image and boot image from the TFTP server to your switch.
- Select the new software image to be loaded into switch memory the next time the switch is reset.

Caution—Before you upgrade from software version 1.x, it is recommended that you save the previous configuration block on a TFTP server, and set the configuration block to factory default, as follows:

```
RS 8124# boot configuration-block factory
```

Configuration Upgrade Notes

When you upgrade the G8124 from release 5.1 or prior, the configuration block is converted to match the new software.

Most configuration data is automatically converted to equivalent commands and ranges. However, some older configuration data has no equivalent on release 5.2 or later, and is not converted. For example, ACL commands are different prior to release 5.2 and are not converted. Log messages list commands that were not converted during the upgrade. You must manually configure those features that were not converted during the upgrade.

If you revert from software image 5.x or later to software image 1.x, the configuration file is cleared and reset to the factory default.

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-image`.

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 7).

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

- The image and boot software loaded on a TFTP server on your network
- The hostname or IP address of the 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 the following procedure to download the new software to your switch.

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

```
RS 8124# copy tftp {image1|image2|boot-image}
```

2. Enter the hostname or IP address of the 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 normally is relative to the TFTP directory (usually `tftpboot`).

The system prompts you to confirm your request.

4. Select the image to run upon the next software reload:

```
RS 8124# boot image {image1|image2}
```

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 Features

This section summarizes the functionality of new features available for the RackSwitch G8124 in BLADEOS 6.3.

For more detailed information about configuring G8124 features and capabilities, please refer to the complete BLADEOS 6.3 documentation as listed on page 3.

Deployment Profiles

The BLADEOS software for the RackSwitch G8124 can be configured to operate in different modes for different deployment scenarios. Each deployment profile sets different capacity levels for basic switch resources, such as the number of IP routes and ARP entries, in order to optimize the switch for different types of networks:

- Default Profile—This profile is recommended for general network usage. Switch resources are balanced to provide moderate capacity for IP routes, ARP entries, ACLs, and VMAPs.
- Routing Profile—This is a special deployment profile. It is recommended when more IP routes are required on the switch. In order to provide the additional IP routes, the number of ARP entries is reduced, and the ACL and VMAP features are unsupported.

The properties of each mode are compared in the following table.

Table 2 Deployment Mode Comparison

Switch Feature	Capacity in Default Profile	Capacity in Routing Profile
IP routes	3,072	9,388
ARP entries	2,048	1,000
OSPF ECMP routes	2684	9,000
ACLs	127	*unsupported*
VMAPs	127	*unsupported*
VM Policy Bandwidth Control	available	*unsupported*

The following ISCLI command is used to change the deployment profile:

```
RS 8124(config)# boot profile {default|routing}   (Select deployment profile)
```

Note – The switch must be rebooted in order for the new mode to take effect.

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 G8124 with BLADEOS 6.3 software is compliant with the INCITS T11.3, FC-BB-5 FCoE specification.

The following are required for implementing FCoE using the G8124:

- The G8124 must be connected to the Fibre Channel network via FCF (such as a Cisco Nexus 5000 Series Switch).
- For internal G8124 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 are enabled and configured with default FCoE settings. These features may be reconfigured, but must remain enabled in order for FCoE to function.
- FIP snooping must be turned on. When FIP snooping is turned on, the feature is enabled on all ports by default. The administrator can disable FIP snooping on individual ports that do not require FCoE, but FIP snooping must remain enabled on all FCoE ports in order for FCoE to function.
- FCoE and vNIC features are mutually exclusive. FCoE is supported only when the G8124 vNIC feature is disabled.
- The FCoE and vNIC features are not supported simultaneously on the same G8124.

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 G8124. 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 G8124 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 G8124 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 G8124 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.

By default, FIP Snooping is turned off. The CEE feature is required to be turned on prior to turning FIP snooping on.

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 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

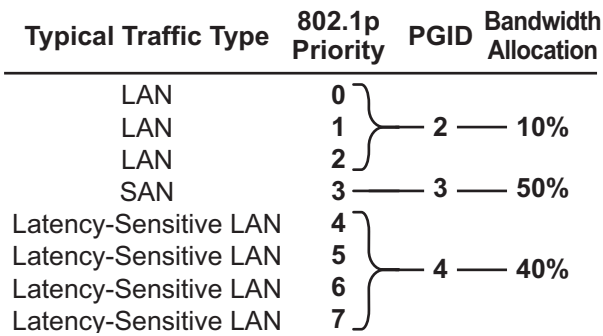

The administrator may reassign 802.1p priority values among up to 8 priority groups (PGIDs), and allocate a percentage of the switch's link bandwidth to each PGID.

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.

Note – Actual bandwidth used by any specific PGID may vary from configured values by up to 10% of the available bandwidth in accordance with 802.1Qaz ETS standard. For example, a setting of 10% may be served anywhere from 0% to 20% of the available bandwidth at any given time.

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 G8124, 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, ETS, and (for some CNAs) FIP. 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.

Note – The DCBX and vNIC features are not supported simultaneously on the same G8124.

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 G8124 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. BLADEOS 6.3 supports up to 2048 VEs.

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 G8124 with VMready also detects the migration of VEs across different hypervisors. As VEs move, the G8124 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 using the following command paths:

```
RS 8124(config)# virt vmgroup ?
RS 8124(config)# virt vmware ?
RS 8124(config)# virt vmprofile ?
RS 8124(config)# virt vmpolicy ?
```

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 using the following command path:

```
RS 8124(config)# access-control vmap <VMAP ID> ?
```

BLADEOS 6.3 supports up to 127 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 regular VLANs, use config-vlan mode:

```
RS 8124(config)# vlan <VLAN ID>
RS 8124(config-vlan)# [no] vmap <VMAP ID> [serverports|
      non-serverports]
```

- For a VM group, use the global configuration mode:

```
RS 8124(config)# [no] virt vmgroup <ID> vmap <VMAP ID>
      [serverports|non-serverports]
```

When the optional `serverports` or `non-serverports` parameter is specified, the action to add or remove the VMAP is applied for either the switch server ports (`serverports`) or uplink ports (`non-serverports`). 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.

To support VMAPs, the number of regular ACLs has been reduced to 128.

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 G8124. 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:

```
RS 8124(config)# ipv6 router ospf                    (OSPFv3 router config mode)
RS 8124(config-router-ospf3)# ?

RS 8124(config)# interface ip <Interface number>     (Configure OSPFv3)
RS 8124(config-ip-if)# ipv6 ospf ?                   (OSPFv3 interface config)

RS 8124# show ipv6 ospf ?                            (Show OSPFv3 information)
```

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)

IGMP Group Capacity

BLADEOS 6.3 supports IGMP groups differently than earlier releases:

The G8124 now supports a maximum of 1024 IGMP entries, on a maximum of 1024 VLANs.

One IGMP entry is allocated for each unique join request, based on the VLAN and IGMP group address. If multiple ports join the same IGMP group using the same VLAN, only a single IGMP entry is used.

Resolved Issues

Switch Reset with No Flash Dump

(ID: 34688) Fixed a condition where an unexpected reset occured without a flash dump being stored. When the show command (**show ?**) is executed continuously, the switch no longer runs out of memory and resets.

IGMP

(ID: 35226) In release 5.1, IGMP Snooping was supported only on 8 VLANs. This is corrected in release 5.2 and later, where IGMP Snooping is supported on all enabled VLANs.

SNMP

vNIC support has been added to the SNMP MIB.

Known Issues

This section describes known issues for the G8124 and BLADEOS 6.3.

Software Upgrade Issues

- Previously configured static MAC addresses must be reconfigured after the upgrade (ID: 35659)
- The administrator may no longer configure the number of IGMP queries sent when a Leave message is received. The count is set to 2 after the upgrade. (ID: 36638)
- TACACS+ secure backdoor is disabled during the upgrade. If you use TACACS+ secure backdoor, you must re-enable it after resetting the switch. (ID: 34707)
- The range value for NTP interval is different compared to 1.x software. On release 1.1 the range is <1-10080> and on release 5.x and later the range is <5-44640>.

  If the NTP interval value is lower than 5, then after software upgrade the NTP interval is set to the minimum value of 5. (ID: 36500)
- During software upgrade, the system time zone setting is lost. Re-configure the system time zone. (ID: 36493)
- The default values and range values for IGMP report timeout parameter are different for release 5.x and later as compared to release 1.1:
  - On release 1.1 the range for IGMP report timeout is <130-1225> seconds with a default of 260 seconds.
  - On release 5.x and later, the range is <1-255> minutes with a default of 10 minutes.

  During upgrade, the value of IGMP report timeout is set to the new default value (10). The value does not appear in the running configuration output. (ID: 35578)
- On release 1.1, the default setting for Hotlinks BPDU is `enabled`, and on release 5.x and later the default setting is `disabled`. During upgrade, the Hotlinks BPDU command is set to `disabled`. (ID: 36385)
- The default value for the **`access https`** command is different compared to release 1.x. On release 1.1 the default setting is `enabled`, and on release 5.x and later, the setting is `disabled`. During upgrade, **`access https`** is set to `disabled`. (ID: 36834)
- The default for the Layer-3 hash is different compared to release 5.x and prior. In release 5.x, the source IP address (SIP) was the default used to generate the Layer-3 hash. In release 6.3, source and destination IP addresses (SIP-DIP) are used as the default. (ID: 39733)
- Some time zones are different compared to release 5.x and prior. After upgrading to release 6.3, it is recommended that the administrator review the configured time zone and make any appropriate changes. (ID:29778)

SNMP

- SNMP read and write functions are enabled by default. For best security practices, if these functions are not needed for your network, it is recommended that you disable these functions prior to connecting the switch to your network. (ID: 40056)

1Gb Fiber Transceivers

- The following limitation applies when using 1Gb Fiber transceivers: The uni-directional link failure detection is not fully functional and the switch may not bring down the link if the Tx fiber cable is severed. This limitation only applies to 1Gb fiber transceivers. (ID: 29759)

LACP

- An LACP trunk with standby ports cannot be changed from the interface portchannel (trunk) menu commands. Interface portchannel commands are applied only to the active ports in an LACP portchannel, so if one port in the LACP portchannel is down, any configuration applied on the portchannel will not be applied to the down port, and the validation routine flags this as an inconsistent configuration. (ID: 36144)

OSPF

- When OSPFv3 is enabled, the OSPF backbone area (0.0.0.0) is created by default and is always active. (ID: 37932)

Port Mirroring

- G8124 port 1 cannot be used as a port monitor for ACL mirroring. (ID: 40416)
- 127.

VMready

- VM entries are not automatically removed when the corresponding port is changed from a server port to a regular port. Be sure that all VMs associated with the target port are removed prior to redefining the port state. (ID: 39044)

vNICs

- When a port is removed from a vNIC group, flow control for the port is automatically disabled. (ID: 38880)

Statistics

- Switch statistics are not updated for ICMP functions (such as `ping`) made using the management ports. (ID: 38997)
- The “all events” counter for OSPFv3 includes the total number of changes associated with any OSPFv3 interface, including changes to internal states. (ID: 38783)

ECMP and Trunk Hashing

- Layer 3 hashing options (source IP address and destination IP address) enabled under ECMP route hashing (`ip route ecmphash`) or port trunk hashing (`portchannel hash`) apply to both ECMP and trunk features (the enabled settings are cumulative). If unexpected ECMP route hashing occurs, disable the unwanted `sip` or `dip` option set in trunk hashing. Likewise, if unexpected trunk hashing occurs, disable any unwanted options set in ECMP route hashing. (ID: 38580)
- Source port (`sport`) and/or destination port (`dport`) options for the ECMP route hash (`ip route ecmphash`) are supported only when Layer 4 `tcpl4` and/or `udpl4` options are enabled. Conversely, when `tcp14` and/or `udpl4` are enabled, hashing options for `sport` and/or `dport` must also be enabled. (ID: 39586)

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)