

IBM Flex System Fabric CN4093 10 Gb Converged Scalable Switch



Release Notes

for Networking OS 7.8

Note: Before using this information and the product it supports, read the general information in the *Safety information and Environmental Notices* and *User Guide* documents on the *IBM Documentation* CD and the *Warranty Information* document that comes with the product.

Second Edition (June 2014)

© Copyright IBM Corporation 2014

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

P/N: 00CG954

Release Notes

This release supplement provides the latest information regarding IBM Networking OS 7.8 for the CN4093 10 Gb Converged Scalable Switch.

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

- *IBM Networking OS Application Guide for the CN4093 10 Gb Converged Scalable Switch*
- *IBM Networking OS ISCLI Reference for the CN4093 10 Gb Converged Scalable Switch*
- *IBM Networking OS BBI Quick Guide for the CN4093 10 Gb Converged Scalable Switch*
- *CN4093 10 Gb Converged Scalable Switch User's Guide*

The publications listed above are available at the following address:


<http://publib.boulder.ibm.com/infocenter/flexsys/information/index.jsp>

Please keep these release notes with your product manuals.

Hardware Support

N/OS 7.8 software is supported on the CN4093 10 Gb Converged Scalable Switch for the IBM Flex System. The CN4093 10 Gb Converged Scalable Switch (CN4093), shown in Figure 1, is a high performance network switch that features high-capacity Ethernet and Fibre Channel ports, and provides tight integration with IBM Flex System chassis management module.

Figure 1. CN4093 10 Gb Converged Scalable Switch Faceplate



CN4093The CN4093 has the following port capacities:

- Forty-Two 10Gb internal ports (maximum)
- Two 10Gb SFP+ ports
- Two high-capacity QSFP+ ports
- Twelve IBM Omni Ports (SFP+) which can be configured (in pairs) to operate in 10Gb Ethernet mode or 4/8Gb Fibre Channel mode
- One 1Gb RJ-45 external management port
- One 1Gb internal management port
- One mini-USB serial port

The CN4093 SFP+ ports accept any SFP+ Direct Attach Cable (DAC) that complies to the MSA specification for the appropriate protocol and port speeds. However, for the most current list of compatible port transceivers and cables, see the appropriate product part number documentation.

Updating the Switch Software Image

The switch software image is the executable code running on the CN4093. A version of the image comes pre-installed on the device. As new versions of the image are released, you can upgrade the software running on your switch. To get the latest version of software supported for your CN4093, go to the following website:

<http://www.ibm.com/support>

To determine the software version currently used on the switch, use the following switch command:

```
CN4093# show version
```

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

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

For instructions on the typical upgrade process using the CLI, ISCLI, or BBI, see [“Loading New Software to Your Switch” on page 7](#).

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



CAUTION:

N/OS 7.5 is the earliest version of software supported by the CN4093 10 Gb Converged Scalable Switch. Do not use any prior versions with this device.

Special Software Update Issues

When updating to N/OS 7.8, the following special conditions may apply, depending on the version of software currently installed on your switch. These conditions are cumulative: If updating from version 2.0 (for example), follow the recommendations in order, beginning with those for 2.0, and then continue with all that apply, such as for “3.0 and prior,” “4.0 and prior,” and so on.

Updating from IBM Networking OS 7.5

Beginning with N/OS 7.7, the UID 1 default name (USERID) cannot be modified. However, you are allowed to change the UID 1 password, if required. Changes made to the UID 1 name in any switch software version prior to N/OS 7.7 will be lost after an upgrade to N/OS 7.7 or later.

Loading New Software to Your Switch

The CN4093 can store up to two different switch software images (called `image1` and `image2`) as well as special boot software (called `boot`). When you load new software, you must specify where it should be placed: either into `image1`, `image2`, or `boot`.

For example, if your active image is currently loaded into `image1`, you would probably load the new image software into `image2`. This lets you test the new software and reload the original active image (stored in `image1`), if needed.

ATTENTION: When you upgrade the switch software image, always load the new boot image and the new software image before you reset the switch. If you do not load a new boot image, your switch might not boot properly (To recover, see ["Recovering from a Failed Software Upgrade" on page 23](#)).

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

- The image and boot software loaded on an SFTP, FTP, or TFTP server on your network.
Note: Be sure to download both the new boot file and the new image file.
- The hostname or IP address of the SFTP, FTP, or TFTP server
Note: The DNS parameters must be configured if specifying hostnames.
- The name of the new software image or boot file

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

Loading Software via the N/OS CLI

1. Enter the following Boot Options command:

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

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

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

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

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

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

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

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

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

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

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

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

```
Boot Options# image
```

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

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

Specify the image that contains the newly loaded software.

7. Reboot the switch to run the new software:

```
Boot Options# reset
```

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

Loading Software via the ISCLI

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

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

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

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

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

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

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

4. If required by the SFTP, FTP, or TFTP server, enter the appropriate username and password.
5. The switch will prompt you to confirm your request.
Once confirmed, the software will begin loading into the switch.

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

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

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

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


7. Reboot the switch to run the new software:

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

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

Loading Software via BBI

You can use the Browser-Based Interface to load software onto the CN4093. The software image to load can reside in one of the following locations:

- SFTP server
- 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 SFTP/FTP/TFTP server, enter the server's information in the SFTP/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 an SFTP/FTP/TFTP server, enter the file name and click **Get Image**.
 - If you are loading software from your computer, click **Browse**.

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

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

New and Updated Features

N/OS 7.8 for CN4093 10 Gb Converged Scalable Switch (CN4093) has been updated to include several new features, summarized in the following sections. For more detailed information about configuring CN4093 features and capabilities, refer to the complete N/OS 7.8 documentation as listed on [page 3](#).

ACLs

Metering is supported for IPv6 ACLs.

Default Setting - SNMP, Telnet, HTTP

The following default settings are available on the CN4093:

- Telnet, HTTP, and SNMP v1 and v2 are disabled.
- SNMPv3 is enabled in stand-alone and stacking mode.
- No default read or write community strings are configured.
- The default username and password are set to USERID; PASSWORD (with a zero) for accessing the CLI or using a Web browser.
- In switch default boot mode, four default SNMPv3 users are available:
 - User 1 name is `adminmd5` (password `adminmd5`). Authentication used is MD5. Privacy protocol used is DES.
 - User 2 name is `adminsha` (password `adminsha`). Authentication used is SHA. Privacy protocol used is DES.
 - User 3 name is `mmv3_mgr` (password `mmv3_mgr`). Authentication used is MD5. Privacy protocol used is DES. User 3 with the default password is used for EHCM level 1 access. For EHCM level 2 and level 3 access, the CMM generates a random password. EHCM level 2 uses MD5 authentication and DES privacy protocol. EHCM level 3 uses SHA authentication and AES-128 privacy protocol
 - User 4 name is `adminshaaes` (password `Edpq132x!#9Zpx432w`). Authentication used is SHA. Privacy protocol used is AES-128.

In boot strict mode, two default SNMPv3 users are available:

- User 1 name is `mmv3_mgr` (password `mmv3_mgr`). Authentication used is SHA. Privacy protocol used is AES-128.
- User 2 name is `adminshaaes` (password `Edpq132x!#9Zpx432w`). Authentication used is SHA. Privacy protocol used is AES-128.

E_Ports

E_ports (expansion ports) connect two full fabric switches to form an inter-switch link (ISL). Up to four Fibre Channel ISLs can be established between two full fabric switches.

Only Fibre Channel port types can be configured as E-ports. These ports must be members of a Fibre Channel VLAN. Use the following commands to configure E_ports:

```
CN4093(config)# system port <port range> type fc
CN4093(config)# interface fc <port range>
CN4093(config-if)# type e
```

In a stack of switches, ports on the Master or Backup switches can be configured as E-ports.

If different zones were configured on the switches being connected using E_ports, the zones merge to establish a consistent zoning policy across the fabric. The zones on the two switches must belong to any one zoneset.

Note: Zones are added in the merged zoneset only if the zoneset on the individual switch was active.

E-ports cannot be used to form stack trunk links.

Limitations

- IBM Networking OS supports ISL distance up to 3 kms.
- E_ports can be configured only on the IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch. E_ports cannot interoperate with switches from other vendors.
- A maximum of eight switches are supported in a Fibre Channel fabric.
- A maximum of four E_ports can be configured on a switch.

Edge Virtual Bridging (EVB)

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

FCoE Link Aggregation

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

Logging by Severity Level

When you enable logging on the CN4093, log messages are displayed on the console and stored on the switch. To view messages of a particular severity level, use the following command:

```
CN4093(config)# logging console severity <0-7>
```

To store messages of a particular severity level, use the following command:

```
CN4093(config)# logging buffer severity <0-7>
```

Multiple Spanning Tree Protocol (MSTP)

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

Use the following commands to configure MSTP:

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

```
CN4093(config)# spanning-tree mst configuration      (Enter MST configuration mode)
CN4093(config-mst)# name <name>                  (Define the Region name)
CN4093(config-mst)# exit
CN4093(config)# spanning-tree mode mst           (Set mode to Multiple Spanning Trees)
```

3. Map VLANs to MSTP instances:

```
CN4093(config)# spanning-tree mst configuration      (Enter MST configuration mode)
CN4093(config-mst)# instance <instance ID> vlan <vlan number or range>
```

NIST SP 800-131A Compliance

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

The CN4093 10 Gb Converged Scalable Switch can operate in two boot modes:

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

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

By default, HTTP, Telnet, and SNMPv1 and SNMPv2 are disabled on the CN4093. In strict mode, you can enable these protocols if the security policy on the switch is set to "secure," but a message similar to the following will appear:

```
NOTICE mgmt: strict mode: Warning, telnet security does not
meet security strict mode requirements
```

Before enabling strict mode, ensure the following:

- The software version on all connected switches is Networking/ OS 7.8.
- The supported protocol versions and cryptographic cipher suites between clients and servers are compatible. For example: if using SSH to connect to the switch, ensure that the SSH client supports SSHv2 and a strong cipher suite that is compliant with the NIST standard.

- Compliant Web server certificate is installed on the switch, if using BBI.
- A new self-signed certificate is generated for the switch (CN4093(config)# access https generate-certificate). The new certificate is generated using 2048-bit RSA key and SHA-256 digest.
- Protocols that are not NIST SP 800-131A compliant must be disabled or not used.
- Only SSHv2 or higher is used.
- The current configuration, if any, is saved in a location external to the switch. When the switch reboots, both the startup and running configuration are lost.
- Only protocols/algorithms compliant with NIST SP 800-131A specification are used/enabled on the switch. Please see the NIST SP 800-131A publication for details. The following table lists the acceptable protocols and algorithms:

Table 1. Acceptable Protocols and Algorithms

| Protocol/Function | Strict Mode Algorithm | Compatibility Mode Algorithm |
|------------------------|---|--|
| BGP | BGP does not comply with NIST SP 800-131A specification. When in strict mode, BGP is disabled. However, it can be enabled, if required. | Acceptable |
| Certificate Generation | RSA-2048 SHA-256 | RSA 2048 SHA 256 |
| Certificate Acceptance | RSA 2048 or higher SHA 224 or higher | RSA SHA, SHA2 |
| HTTPS | TLS 1.2 only See "Acceptable Cipher Suites" on page 15 ; | TLS 1.0, 1.1, 1.2 See "Acceptable Cipher Suites" on page 15 ; |
| IKE | | |
| Key Exchange | DH Group 24 | DH group 1, 2, 5, 14, 24 |
| Encryption | 3DES, AES-128-CBC | 3DES, AES-128-CBC |
| Integrity | HMAC-SHA1 | HMAC-SHA1, HMAC-MD5 |
| IPSec | | |
| AH | HMAC-SHA1 | HMAC-SHA1, HMAC-MD5 |
| ESP | 3DES, AES-128-CBC, HMAC-SHA1 | 3DES, AES-128-CBC, HMAC-SHA1, HMAC-MD5 |
| LDAP | LDAP does not comply with NIST SP 800-131A specification. When in strict mode, LDAP is disabled. However, it can be enabled, if required. | Acceptable |
| OSPF | OSPF does not comply with NIST SP 800-131A specification. When in strict mode, OSPF is disabled. However, it can be enabled, if required. | Acceptable |

Table 1. Acceptable Protocols and Algorithms

| Protocol/Function | Strict Mode Algorithm | Compatibility Mode Algorithm |
|-------------------------|---|---|
| RADIUS | RADIUS does not comply with NIST SP 800-131A specification. When in strict mode, RADIUS is disabled. However, it can be enabled, if required. | Acceptable |
| Random Number Generator | NIST SP 800-90A AES CTR DRBG | NIST SP 800-90A AES CTR DRBG |
| Secure NTP | Secure NTP does not comply with NIST SP 800-131A specification. When in strict mode, secure NTP is disabled. However, it can be enabled, if required. | Acceptable |
| SLP | SHA-256 or higher RSA/DSA 2048 or higher | |
| SNMP | SNMPv3 only AES-128-CFB-128/SHA1 Note: Following algorithms are acceptable if you choose to support old SNMPv3 factory default users: AES-128-CFB/SHA1 DES/MD5 AES-128-CFB-128/SHA1 | SNMPv1, SNMPv2, SNMPv3 DES/MD5, AES-128-CFB-128/SHA1 |
| SSH/SFTP | | |
| Host Key | SSH-RSA | SSH-RSA |
| Key Exchange | ECDH-SHA2-NISTP521 ECDH-SHA2-NISTP384 ECDH-SHA2-NISTP256 ECDH-SHA2-NISTP224 RSA2048-SHA256 DIFFIE-HELLMAN-GROUP-EXCHANGE-SHA256 DIFFIE-HELLMAN-GROUP-EXCHANGE-SHA1 | ECDH-SHA2-NISTP521 ECDH-SHA2-NISTP384 ECDH-SHA2-NISTP256 ECDH-SHA2-NISTP224 ECDH-SHA2-NISTP192 RSA2048-SHA256 RSA1024-SHA1 DIFFIE-HELLMAN-GROUP-EXCHANGE-SHA256 DIFFIE-HELLMAN-GROUP-EXCHANGE-SHA1 DIFFIE-HELLMAN-GROUP14-SHA1 DIFFIE-HELLMAN-GROUP1-SHA1 |

Table 1. Acceptable Protocols and Algorithms

| Protocol/Function | Strict Mode Algorithm | Compatibility Mode Algorithm |
|-------------------|---|--|
| Encryption | AES128-CTR AES128-CBC 3DES-CBC | AES128-CTR AES128-CBC RIJNDAEL128-CBC BLOWFISH-CBC 3DES-CBC ARCFOUR256 ARCFOUR128 ARCFOUR |
| MAC | HMAC-SHA1 HMAC-SHA1-96 | HMAC-SHA1 HMAC-SHA1-96 HMAC-MD5 HMAC-MD5-96 |
| TACACS+ | TACACS+ does not comply with NIST SP 800-131A specification. When in strict mode, TACACS+ is disabled. However, it can be enabled, if required. | Acceptable |

Acceptable Cipher Suites

The following cipher suites are acceptable (listed in the order of preference) when the CN4093 10 Gb Converged Scalable Switch is in compatibility mode:

Table 2. List of Acceptable Cipher Suites in Compatibility Mode

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

The following cipher suites are acceptable (listed in the order of preference) when the CN4093 10 Gb Converged Scalable Switch is in strict mode:

Table 3. List of Acceptable Cipher Suites in Strict Mode

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

Table 3. List of Acceptable Cipher Suites in Strict Mode

| | | | | | |
|--------|-----|-----|-------------|--------|---------------------------------|
| 0x003C | RSA | RSA | AES_128_CBC | SHA256 | TLS_RSA_WITH_AES_128_CBC_SHA256 |
| 0x000A | RSA | RSA | 3DES | SHA1 | SSL_RSA_WITH_3DES_EDE_CBC_SHA |

Configuring Strict Mode

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

```
CN4093(config)# [no] boot strict enable
```

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

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

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

Please see the *IBM Networking OS 7.8 CN4093 10 Gb Converged Scalable Switch Application Guide* for details on SNMPv3 users.

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

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

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

Limitations

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

- Power ITEs and High-Availability features do not comply with NIST SP 800-131A specification.
- The CN4093 will not discover Platform agents/Common agents that are not in strict mode.
- Web browsers that do not use TLS 1.2 cannot be used.
- Limited functions of the switch managing Windows will be available.

Flexible Port Mapping

Users can change the default licensed port mapping by manually activating or deactivating physical ports within the limitations of the installed licenses' bandwidth. To activate/deactivate a port, use the following command:

```
CN4093(config)# [no] boot port-map <port>
```

Optimized FCoE Traffic Flow

To optimize the FCoE traffic flow, ACL entries are installed by default. Only FCoE to FCoE traffic is optimized. Traffic to and from Fibre Channel nodes is not optimized.

If required, you can disable optimized traffic flow. However, you must first disable FIP snooping. Use the following commands:

```
CN4093(config)# no fcoe fips enable
CN4093(config)# no fcoe optimized-forwarding enable
```

To re-enable optimized traffic flow, use the following command sequence:

```
CN4093(config)# no fcoe fips enable
CN4093(config)# fcoe optimized-forwarding enable
CN4093(config)# fcoe fips enable
```

Optimized traffic flow is enabled for a single switch FCF as well as a stack. In a stack, one or more members may have optimized flow entries installed based on the configuration. You can view the installed optimized ACLs on the respective switches.

Note: Optimized ACL entries cannot be installed on ports 53-64, even when the ports are in FCoE mode.

The following commands provide optimized traffic flow information:

- To check current state:

```
CN4093# show fcoe optimized-forwarding status
```

- To view list of optimized ACLs:

```
CN4093# show fcoe optimized-acls vlan <VLAN ID>
```

OSPFv3 Over IPsec

BBI and SNMP support for OSPFv3 over IPsec has been added.

Private VLANs

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

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

Quality of Service (QoS)

The following commands to view QoS statistics have been added:

- CN4093(config)# show interface port <port number or range> egress-queue-counters {<queue number>|drop}
- CN4093(config)# show interface port <port number or range> egress-queue-rate {<queue number>|drop}

The output of these commands include the following information:

- Number of packets/bytes transmitted per queue
- Rate of packets/bytes transmitted per queue
- Number of packets/bytes dropped per queue
- Rate of packets/bytes dropped per queue

Single IP Management

The master switch in a stack can have a floating management IP address set up on the management interface. In case of master switch failure, the floating management IP will be used by the backup switch taking over management.

SNMP

ACLs

ACLs can be configured to permit or deny SNMP messages from reaching the switch CPU. These restrictions can be applied to:

- In-band and out-of-band SNMP messages
- SNMPv1 and SNMPv2 messages
- IPv4 and IPv6 SNMP messages

The restrictions do not affect other management traffic reaching the CPU.

You can specify the IPv4 and IPv6 subnetworks that you want to permit access to. The access can be read-only or read-write. Use the following commands:

Table 4. SNMP ACL Commands

| ACL | Description |
|---|---|
| <code>CN4093(config)# access management-network <IPv4 address> <subnet mask> snmp-ro</code> | Add an IPv4 SNMP read-only access list. |
| <code>CN4093(config)# access management-network <IPv4 address> <subnet mask> snmp-rw</code> | Add an IPv4 SNMP read-write access list. |
| <code>CN4093(config)# access management-network6 <IPv6 address> <IPv6 prefix> snmp-ro</code> | Add an IPv6 SNMP read-only access list. |
| <code>CN4093(config)# access management-network6 <IPv6 address> <IPv6 prefix> snmp-rw</code> | Add an IPv6 SNMP read-write access list. |
| <code>CN4093(config)# no access management-network <IPv4 address> <subnet mask> snmp-ro</code> | Delete an IPv4 SNMP read-only access list. |
| <code>CN4093(config)# no access management-network <IPv4 address> <subnet mask> snmp-rw</code> | Delete an IPv4 SNMP read-write access list. |
| <code>CN4093(config)# no access management-network6 <IPv6 address> <IPv6 prefix> snmp-ro</code> | Delete an IPv6 SNMP read-only access list. |
| <code>CN4093(config)# no access management-network6 <IPv6 address> <IPv6 prefix> snmp-rw</code> | Delete an IPv6 SNMP read-write access list. |
| <code>CN4093(config)# no access management-network snmp-ro</code> | Clear IPv4 SNMP read-only access lists. |

Table 4. SNMP ACL Commands

| ACL | Description |
|---|--|
| CN4093(config)# no access management-network snmp-rw | Clear IPv4 SNMP read-write access lists. |
| CN4093(config)# no access management-network6 snmp-ro | Clear IPv6 SNMP read-only access lists. |
| CN4093(config)# no access management-network6 snmp-rw | Clear IPv6 SNMP read-write access lists. |

MIBs

Added additional FCoE and Fibre Channel MIBs.

Telnet

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

Unified Fabric Port (UFP)

- UFP can be configured in stacking mode.
- Edge Virtual Bridging (EVB) can be configured with UFP in both stand-alone and stacking mode.

UFP works with other CN4093 features. Please see the *IBM Networking OS CN4093 10 Gb Converged Scalable Switch Application Guide* for details.

Layer 2 Failover

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

Layer 2 failover is not supported on UFP ports in auto mode.

Increased VLAN Limits

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

VMReady

Configuring with UFP and VMReady, the CN4093 can support up to 32 VMGroups with UFP vPorts in auto-mode.

VMReady is supported only on a vPort which is configured in auto-VLAN mode.

802.1Qbg

Configured with Edge Virtual Bridging (EVB), UFP supports up to 256 VLANs on a vPort.

EVB is supported only on a vPort which is configured in auto-VLAN mode.

User Access

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

VRRP - Holdoff Time

Each VRRP router is configured with a priority between 1–254. A bidding process determines which VRRP router is or becomes the master—the VRRP router with the highest priority.

The master periodically sends advertisements to an IPv4 multicast address. As long as the backups receive these advertisements, they remain in the backup state. If a backup does not receive an advertisement for three advertisement intervals, it initiates a bidding process to determine which VRRP router has the highest priority and takes over as master. In addition to the three advertisement intervals, a manually set holdoff time can further delay the backups from assuming the master status.

Resolved Issues

The following known issues have been resolved.

Private VLANs

Traffic with secondary VLAN ID is not forwarded to promiscuous ports. (ID: 70980)

Fibre Channel

- If using BBI for Fibre Channel configuration, you must save the configuration using only the ISCLI command:
`CN4093# copy running-configuration startup-configuration.`
Do not use the save button on the BBI. (ID: XB217551)
- When using Emulex CNA with a fully loaded Flex System chassis, do not provision more than four ports (FCF MAC addresses) in a Fibre Channel VLAN. Provisioning more than four ports may result in FCoE link flaps. (ID: XB203686)
- In NPV mode, do not provision more than 4 uplink ports. Each host (using Emulex CNA) broadcasts FIP control frames to all uplink ports. A large number of such frames may result in session flaps. A maximum of 60 hosts can be provisioned with 4 uplinks. You can provision more hosts using lesser number of uplinks. For example: 80 hosts with 3 uplinks or 120 hosts with 2 uplinks. (ID: XB217235)
- If you modify the pWWN of a Fibre Channel alias, the changed pWWN is saved in the running configuration but is not applied to the FC device until a reboot. (ID: SW216951)
Workaround: To modify the pWWN, first remove the FC alias and then add it back with the changed pWWN.
- If you delete a VLAN that was used to create a zone, then that VLAN parameter cannot be used in the command to delete the zone. (ID: 69390)
Workaround: To delete the zone, use any other VLAN that has a Fibre Channel member port.
- On a CN4093, if there is only one VLAN with a member Fibre Channel port, and if zoning has been configured on the switch, operation to remove the last member Fibre Channel port from the VLAN will be blocked. (ID: 69389)
Workaround: Add another Fibre Channel port from that switch as a member of another VLAN, and then remove the previous port from the VLAN.
- A Spurious Fault Status bit value reported by the FC Module may cause the FC Module to shut down. (ID: XB276665)
- An "8Gb FC removed at port NN" message appears indicating that the FC Transceiver was removed from a port where the transceiver was not removed. When this message appears, FCoE sessions may flap or be totally disconnected. (ID: XB262084, XB270246)
- Excessive FIP Multicast packets may overwhelm the FC Module CPU. This may cause the FC Module to be slow to respond or become completely unresponsive. (ID: XB215762, XB271920, XB272548)

OSPF

OSPF does not choose the ASE with the smallest external metric for type 2 metric ASE. When choosing which ASE to use, OSPF selects the best total cost, and not the smallest external metric. (ID: XB267981)

Miscellaneous

The Switch Type is never displayed at the CLI login prompt or on the BBI login page.
(ID: XB267783)

Supplemental Information

This section provides additional information about configuring and operating the CN4093 and N/OS.

The Boot Management Menu

The Boot Management menu allows you to switch the software image, reset the switch to factory defaults, or to recover from a failed software download.

You can interrupt the boot process and enter the Boot Management menu from the serial console port. When the system displays Memory Test, press **<Shift B>**. The Boot Management menu appears.

```
Resetting the System ...
Memory Test .....

1 - Change booting image
2 - Change configuration block
3 - Boot in recovery mode (tftp and xmodem download of images to
   recover switch)
4 - Xmodem download (for boot image only- use recovery mode for
   application images)
5 - Reboot
6 - Exit

Please choose your menu option: 3
```

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 a software image recovery, press 3 and follow the screen prompts.
- To perform an Xmodem download (boot image only), press 4 and follow the screen prompts.
- To exit the Boot Management menu, press 6. The booting process continues.

Recovering from a Failed Software Upgrade

Use the following procedure to recover from a failed software upgrade.

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports Xmodem download (for example, HyperTerminal, 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 **Boot in recovery mode**. You will see the following display:

```
Entering Rescue Mode.
Please select one of the following options:
    T) Configure networking and tftp download an image
    X) Use xmodem 1K to serial download an image
    R) Reboot
    E) Exit
```

- If you choose option **x** (Xmodem serial download), go to step 5.
 - If you choose option **t** (TFTP download), go to step 6.
5. **Xmodem download**: When you see the following message, change the Serial Port characteristics to 115200 bps:

```
Change the baud rate to 115200 bps and hit the <ENTER> key before
initiating the download.
```

- 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.
- When you see the following message, change the Serial Port characteristics to 9600 bps:

```
Change the baud rate back to 9600 bps, hit the <ESC> key.
```

- When you see the following prompt, enter the image number where you want to install the new software and press **<Enter>**.

```
Install image as image 1 or 2 (hit return to just boot image): 1
```

- The following message is displayed when the image download is complete. Continue to step 7.

```
Installing image as image1...
Image1 updated successfully
Please select one of the following options:
    T) Configure networking and tftp download an image
    X) Use xmodem 1K to serial download an image
    R) Reboot
    E) Exit
```

6. **TFTP download**: The switch prompts you to enter the following information:

```
Performing TFTP rescue. Please answer the following questions (enter
'q' to quit):
IP addr   :
Server addr:
Netmask   :
Gateway   :
Image Filename:
```


- a. Enter the required information and press **<Enter>**.
- b. You will see a display similar to the following:

```

Host IP   : 10.10.98.110
Server IP : 10.10.98.100
Netmask   : 255.255.255.0
Broadcast : 10.10.98.255
Gateway   : 10.10.98.254
Installing image 6.8.3_0S.img from TFTP server 10.10.98.100

```

- c. When you see the following prompt, enter the image number where you want to install the new software and press **<Enter>**.

```

Install image as image 1 or 2 (hit return to just boot image): 1

```

- d. The following message is displayed when the image download is complete. Continue to step 7.

```

Installing image as image1...
Image1 updated successfully
Please select one of the following options:
  T) Configure networking and tftp download an image
  X) Use xmodem 1K to serial download an image
  R) Reboot
  E) Exit

```

7. Image recovery is complete. Perform one of the following steps:
 - Press **r** to reboot the switch.
 - Press **e** to exit the Boot Management menu
 - Press the Escape key (**<Esc>**) to re-display the Boot Management menu.

Recovering a Failed Boot Image

Use the following procedure to recover from a failed boot image upgrade.

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports Xmodem download (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 **4** for **Xmodem download**. You will see the following display:

```

Perform xmodem download

To download an image use 1K Xmodem at 115200 bps.

```

5. When you see the following message, change the Serial Port characteristics to 115200 bps:

```
Change the baud rate to 115200 bps and hit the <ENTER> key before
initiating the download.
```

- a. Press <Enter> to set the system into download accept mode. When the readiness meter displays (a series of "C" characters), start Xmodem on your terminal emulator. You will see a display similar to the following:

```
Extracting images ... Do *NOT* power cycle the switch.
**** RAMDISK ****
Un-Protected 38 sectors
Erasing Flash...
..... done
Erased 38 sectors
Writing to
Flash...9...8...7...6...5...4...3...2...1... done
Protected 38 sectors
**** KERNEL ****
Un-Protected 24 sectors
Erasing Flash...
..... done
Erased 24 sectors
Writing to Flash...9...8...7...6...5...4...3...2...1...
```

- b. When you see the following message, change the Serial Port characteristics to 9600 bps:

```
Change the baud rate back to 9600 bps, hit the <ESC> key.
```

Boot image recovery is complete.

Chassis Management Module

The switch management port IP address can only be configured via the CMM web interface. The switch-based configuration interfaces (such as the menu-based CLI, ISCLI, BBI, etc.) cannot be used for this purpose.

When configuring the IP interface, which is dedicated to the internal management port (IF128, MGT1), you cannot use a subnet that is already configured on any other enabled interface (IF1-127). This results in IF128 being disabled and an IP configuration of all zeros displayed on the CMM user interface. The CMM event log will indicate that a "Duplicate route" was detected.

For example, consider that the interface dedicated to the external management port (EXTM, IF127) is configured or enabled to the following IP address and mask:

```
Interface information:
127: IP4 192.168.71.120 255.255.255.0
```

The switch will reject an attempt made from the CMM CLI to configure the internal management port (MGT1, IF128) to the following IP address and mask:

```
system: switch[1]> ifconfig -i 192.168.71.130 -s 255.255.255.0
```

In this scenario, the switch rejects the attempt by disabling any current configuration on IF128, and responds to the CMM with an IP address, mask, and gateway that contains all zeros.

On the CMM CLI, the resulting condition appears as follows:

```
system: swi tch[1]> i fconfi g
Ethernet ScSE
Enabl ed
-c  static
-i  0.0.0.0
-s  0.0.0.0
-g  0.0.0.0
system: mm[1]> di spl ay l og
1 I IOMod_01 04/03/12 08:02:49 (i omodul e01) Dupli cate route
detected to I/O modul e i omodul e01.
2 I IOMod_01 04/03/12 08:02:49 (i omodul e01) I/O modul e 1 I P
address was changed to 0.0.0.0.
```

VLAGs

For optimal VLAG operation, adhere to the following configuration recommendations:

- Any port-related configuration, such as applied ACLs, should be the same for all ports included in the same VLAG, across both peer switches.
- Configure VLAG health checking as shown in the *Application Guide*.

After configuring VLAG, if you need to change any configuration on the VLAG ports, you must follow the guidelines given below:

- If you want to change the STP mode, first disable VLAG on both the peers. Make the STP mode-related changes and re-enable VLAG on the peers.
- If you have MSTP on, and you need to change the configuration of the VLAG ports, follow the steps below:

On the VLAG Secondary Peer:

1. Shutdown the VLAG ports on which you need to make the change.
2. Disable their VLAG instance using the command:
CN4093 (config)# no vlag adminkey <key> enable (or)
CN4093 (config)# no portchannel <number> enable
3. Change the configuration as needed.

On the VLAG Primary Peer:

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

On the VLAG Secondary Peer:

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

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

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.

Port Mirroring Tags BPDUs

When you perform port mirroring, Spanning Tree BPDUs are VLAN tagged at the monitoring port. This is standard behavior of port mirroring on the CN4093. All mirrored egress traffic is tagged.

Secure Management Network

The following CN4093 attributes are reserved to provide secure management access to and from the chassis management module:

- MGT port (MGT1)
- VLAN 4095
- IP interface 126, 128
- Gateway 4
- STG 128

For more information about remotely managing the CN4093 through the external ports, see “Accessing the Switch” in the *IBM Networking OS 7.8 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 CN4093 attempts to avoid unnecessary key generation. The process generates three server keys:

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

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

Spanning Tree Configuration Tips

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

- Create a separate Spanning Tree Group for each VLAN.
- Manually add all associated VLANs into a single Spanning Tree Group.

When using Layer 2 Trunk Failover, disable Spanning Tree Protocol on external ports.

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 CN4093, 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 CN4093s in the network. Refer to “System Host Log Configuration” in the *Command Reference*.

Trunk Group Configuration Tips

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

- Always configure trunk groups first, on both ends, before you physically connect the links.
- Configure all ports in a trunk group to the same speed (you cannot aggregate 1Gb ports with 10GBASE-SFP+ ports).
- Configure all ports in a trunk group with the same duplex.
- Configure all ports in a trunk group with the same flowcontrol.

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:

```
CN4093(config)# virt vmware scan
```

VRRP Configuration

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

Known Issues

This section describes known issues for N/OS 7.8 on the CN4093 10 Gb Converged Scalable Switch.

BBI

- While accessing BBI pages, the switch may crash. This event is unpredictable and is not related to any particular BBI page or configuration. (ID: 67865)
- If using BBI to add ports or remove ports to/from a Fibre Channel VLAN, first apply the change and then proceed with further configuration. (ID:XB250871)

Boot Configuration Block

In the CLI, the boot configuration command `(CN4093(config)# boot configuration-block)` examines only the initial character of the *block* option. Invalid *block* strings (those other than `active`, `backup`, or `factory`) that use a valid first character (a, b, or f) will be interpreted as the matching valid string. (ID: 42422)

Chassis Management Module (CMM)

- The switch management port IP address cannot currently be configured via the CMM web interface. Use an alternate switch configuration method such as the CLI, ISCLI, BBI, etc. (ID: 64760)
- NTP configuration cannot currently be saved via the CMM web interface. Use an alternate switch configuration method such as the CLI, ISCLI, BBI, etc.

DHCP

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

FCoE

- In N/OS 7.8, the CN4093 supports up to 175 simultaneous FCoE sessions. When this capacity is reached, traffic for additional sessions is dropped, though some host servers and uplink devices may consider all sessions fully established. (ID: 60337, 64842)
- When using FCoE to connect the switch to a Cisco Nexus 5000 (as the external FCF), the DCBX PFC willing flag must be enabled. (ID: 65043)
- Disruption to FCoE connections and FCoE traffic may be expected when changing the LACP mode. It is recommended that the administrator halt FCoE traffic before changing any switch configuration. (ID: 67044)
- In NPV mode, when static LAG is enabled or disabled, multiple MAC addresses are displayed for a single port. You can view this using the `show fcoe fips fcoe` command. However, the MAC addresses get cleared based on the configured timeout value. (ID: XB223966)

Fibre Channel

- Use only the ISCLI or BBI to configure Fibre Channel. IBM N/OS CLI is not supported. After configuring Fibre Channel, save any subsequent configurations only in ISCLI or BBI. If IBM N/OS CLI is used to save any switch configuration, the Fibre Channel configuration will be lost.
- If you need to change the Fibre Channel configuration mode from BBI to ISCLI, first save the configuration using the `save` button on the BBI. Do not use two configuration modes at the same time. (ID: XB205624)
- On two full fabric switches connected using E_ports, if you delete a zoneset and reload the switch(es), the zoneset remains active after the switch comes up. (ID: XB250712).

Workaround: Delete the zoneset after the switch reloads.

- In a topology where local FCFs (in full FC mode or NPV mode) are connected to external FCFs (such as Cisco Nexus 5000 Series Switch) via Ethernet ports, we recommend you configure two different default VLANs: one for the local FCFs and one for the external FCFs. The default VLANs are used to learn the FCoE VLANs supported. (ID: XB251853)

Forwarding Database (FDB)

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

HTTPS

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

IPsec

IPsec does not support virtual links. (ID: 48914)

ISCLI

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

ISCLI Configuration Scripts

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

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

LACP

- If a static trunk on the CN4093 is connected to another CN4093 with LACP configured (but no active LACP trunk), the `CN4093# show portchannel information` command might erroneously report the static trunk as forwarding.
- If you configure LACP (active/passive) on one port, also configure LACP on the partner switch, at the end of the link. If you connect LACP with a static trunk, there will be no connectivity on that link.
- Since LACP trunks use LACPDU packet to maintain trunking with the partner, there is a possibility for those packets to be dropped from an extremely busy trunk. If this happens, some links in the LACP trunk might be removed, then aggregated back to the trunk if an LACPDU is received. To avoid this unstable LACP trunk link, you can add more links to the trunk to increase the bandwidth, or use regular static trunk if there are no more links available.
- Under some conditions, setting the LACP timeout value on partner switches to “short” may cause LACP links to flap in and out of service. If this situation occurs, set the LACP timeout value to “long.” (ID: 63405, 64518)
- Under heavy switch load conditions, LACP links may flap when configured with short timeout mode. To stabilize LACP under heavy load, it is recommended to use the long timeout mode instead. (ID: 66173)

OSPF

- Some changes to OSPF configuration (such as creating a new area or changing an area’s type) may result in OSPF state reconvergence. (ID: 46445, 48483)
- OSPFv3 over IPsec
 - This combination can only be configured only on a per-interface basis. Configuration based on virtual links is not currently supported.
 - The current implementation for OSPFv3 allows the use of only one protocol (AH or ESP) at any given time. AH and ESP cannot be applied together.

Ports and Transceivers

- Under repeated and rapid removal and reinsertion a port transceiver, it is possible that the resulting port state may not be represented accurately within the switch. (ID 32412)
Solution: Once you have removed a transceiver from a switch port, wait five seconds before reinserting any transceiver into the same port. This allows the port to stabilize, and promotes accurate port state information within the switch.
- When the link speed for an external connection is forced (i.e. no Auto-Negotiation) to 100 Mbps and then changed to 10 Mbps, if the external device is changed first, the external device may erroneously report the link as DOWN even after the switch is changed to 10 Mbps.
Solution: At the external device, disconnect and reconnect the cable.

- Interoperability with Older Hubs
The command-line interface might display link up and link down messages continuously for an external port that is connected to certain older hub models configured for 100 Mbps half-duplex. The display might show link up erroneously. This behavior has been observed when connecting the GbESM with the following devices:
 - NETGEAR FE104 100 hub
 - SBS 1000Base-T NIC
 - 3Com Linkbuilder FMS100 Hub 3C250 TX/I
 - 3Com SuperStack II 100TX 3C250C-TX-24/12
 - Nortel Baystack 204 Hub
- If the CN4093 is connected to an application switch which requires a link speed of 100 Mbps half-duplex, then enable auto negotiation on the CN4093 port with `port speed=any, mode=any, fctl=both, and auto=on`.
- Egress packets contribute to statistics on IBM Omni Ports even when link is down or transceivers are not present. (ID: 62639)
- In Ethernet mode (the default), IBM Omni Ports may take longer than dedicated Ethernet ports to reflect changes in port link status. As a result, some traffic loss can be expected while the port transitions to a down state. Also, when using protocols sensitive to link timing (VRRP, OSPF, BGP, LACP, VLAG, IGMP, etc.) it is not recommended to use low- or sub-second timer values. (ID: 64746)

QoS

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

```
CN4093(config)#
show interface port <swunit:port_num> egress-mcast-queue-counters

For example:
CN4093(config)# show interface port 1:24 egress-mcast-queue-counters

Multicast QoS statistics for port 1:24:
QoS Queue 8:
  Tx Packets:                377
  Dropped Packets:           0
  Tx Bytes:                  50883
  Dropped Bytes:             0
```

QSFP+

- The QSFP+ ports do not auto-negotiate. The desired speed must be configured to match on both ends of the connection, and the switch reset for changes to take effect. (ID: 46340)
- After you upgrade switch software and reset the switch, you must configure the QSFP+ port mode. Use the following command (ID: 46858):


```
boot qsfp-40gports <15,19>
```

SLP

When using multi-value attributes that contain a list of comma-separated values, the service reply will match if it contains one or more of the values. It is not required that all values match. (ID: 60086)

SNMP

- During SNMP MIB walks, if you experience timeouts, set the timeout value to 3 seconds or higher in the SNMP application/tool. (IDs: 71913, 71914, 71906)
- If you delete multiple VLANs using SNMP, you may see an error if the SNMP packet size exceeds 1800 bytes. (ID: XB228120)
- When you try to delete an SNMP community, you may see the following error message: (ID: XB222036)

```
Error: do not find correspond additional read community string.
```

Workaround: Reload the switch.
- The image is transferred multiple times when the number of retries is set to be greater than 0. (ID: XB278271)

Stacking

- DHCP has higher priority over static management IP configuration. If you want to configure a static management IP, you must first disable DHCP. (ID: 68589)
- When you issue the `show mac-address-table counters` command, the output is displayed after a delay of approximately 20 seconds. In the mean time, if there are any updates to the table, the output of the command may have incorrect information. (ID: XB218728)

Virtual Link Aggregation Groups

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

VMready

When VMready information commands are issued on a stack member, you may see VM validation messages from groups in the VM table, although the VM is not attached. (ID: XB210686)