Lenovo Flex System Fabric CN4093 10 Gb Converged Scalable Switch

Release Notes

for Networking OS 8.3

Lenovo

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 Lenovo *Documentation* CD and the *Warranty Information* document that comes with the product.

Fourth Edition (August 2016)

Release Notes

This release supplement provides the latest information regarding Lenovo Networking OS 8.3 for the CN4093 10 Gb Converged Scalable Switch.

This supplement modifies and extends the following Lenovo N/OS documentation for use with N/OS 8.3:

- *CN4093 10 Gb Converged Scalable Switch Application Guide for Lenovo Networking OS 8.3*
- *CN4093 10 Gb Converged Scalable Switch Command Reference for Lenovo Networking OS 8.3*
- *CN4093 10 Gb Converged Scalable Switch Installation Guide*

The publications listed 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

Lenovo Networking OS 8.3 software is supported on the CN4093 10 Gb Converged Scalable Switch for the Lenovo 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 that can change between Ethernet and Fibre Channel modes, and provides tight integration with the Lenovo Flex System chassis management module.

Figure 1. CN4093 10 Gb Converged Scalable Switch Faceplate

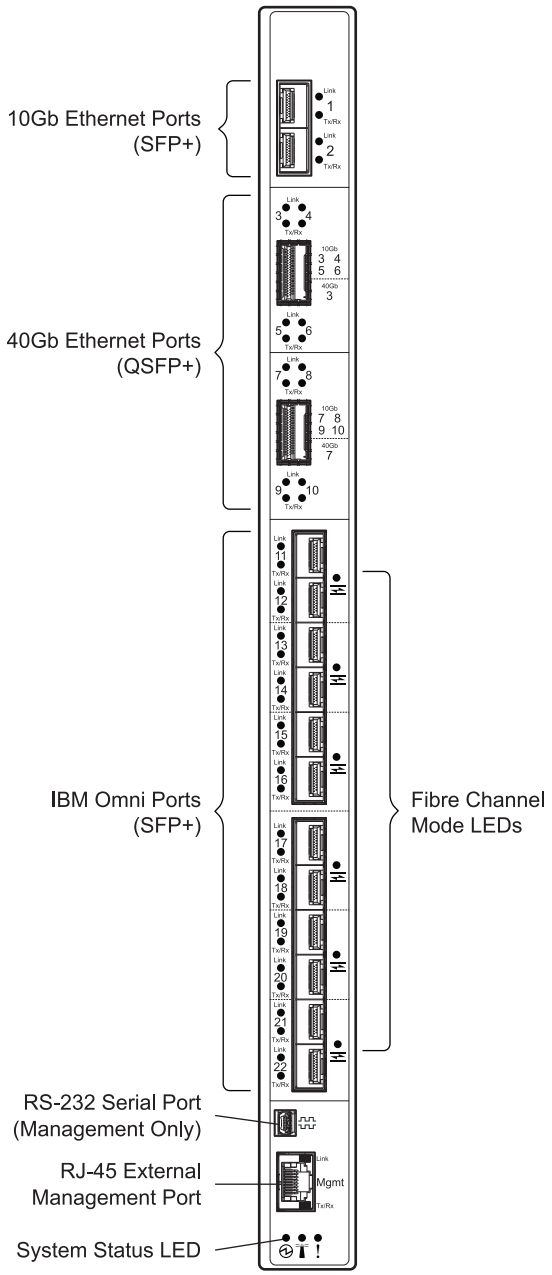

The CN4093 has the following port capacities:

- Forty-Two internal ports (maximum)
- Two 10Gb SFP+ ports
- Two high-capacity QSFP+ ports
- Twelve 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

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 a 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.
- Reload the switch.

For instructions on the typical upgrade process, see “Loading New Software to Your Switch” on page 7.

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 Image Upgrade" on page 12).

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

- The image and boot software loaded on a 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 the following procedure to download the new software to your switch.

Loading Software via the Command Line Interface

Follow these steps to load software onto your switch:

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

```
CN4093# 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:

```
CN4093# configure terminal
CN4093(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:

```
CN4093(config)# reload
```

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

   Note: If you select "No" when asked to confirm the reload, any changes made to the configuration since the last reboot will be lost.

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.

Updating vLAG Switches with Lenovo Networking OS 8.x

Below are the steps for updating the software and boot images for switches configured with vLAG:

1. Save the configuration on both switches using the following command:

```
CN4093# copy running-config startup-config
```

2. Use FTP, STFP or TFTP to copy the new Networking OS and boot images onto both vLAG switches. For more details, see “Loading Software via the Command Line Interface” on page 8.

3. Shutdown all ports except the ISL ports and the health check port on the primary switch (Switch 1).

   Note: Do not save this configuration.

4. Reload Switch 1. Switch 2 will assume the vLAG primary role. Once Switch 1 has rebooted, Switch 1 will take the vLAG secondary role.

5. Shutdown all ports except the ISL ports and the health check port on Switch 2.

   Note: Do not save this configuration.

6. Reload Switch 2. Switch 1 will assume the vLAG primary role. Once Switch 2 has rebooted, make sure that Switch 1 is now the vLAG primary switch and Switch 2 is now the vLAG secondary switch.

7. Verify the all the vLAG clients have converged using the following command:

```
CN4093> show vlag information
```

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 ..............................

Boot Management Menu
        I - Change booting image
        C - Change configuration block
        R - Boot in recovery mode (tftp and xmodem download of images to
recover switch)
        Q - Reboot
        E - Exit
Please choose your menu option:
```

The Boot Management menu allows you to perform the following actions:

- To change the booting image, press **I** and follow the screen prompts.
- To change the configuration block, press **C** and follow the screen prompts.
- To perform a software image recovery, press **R** and follow the screen prompts.
- To reboot the switch press **Q** and follow the screen prompts.
- To exit the Boot Management menu, press **E**. The booting process continues.

Recovering from a Failed Image Upgrade

The Boot Management menu allows you to perform fundamental device management operations, such as selecting which software image will be loaded, resetting the CN4093 to factory defaults or recovering from a failed image download.

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

1. Connect a PC to the serial Console port of the CN4093.
2. Open a terminal emulator program that supports XModem Download (for example, HyperTerminal, SecureCRT or PuTTY) and select the following port characteristics:
   - Speed: 9,600 bps
   - Data Bits: 8
   - Stop Bits: 1
   - Parity: None
   - Flow Control: None
3. To access the Boot Management menu, you must interrupt the boot process from the Console port. Boot the CN4093, and when the system begins displaying Memory Test progress (a series of dots), press **<Shift + B>**.

   The Boot Management menu will appear:

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

Boot Management Menu
        I - Change booting image
        C - Change configuration block
        R - Boot in recovery mode (tftp and xmodem download of images to
recover switch)
        Q - Reboot
        E - Exit
Please choose your menu option:
```

4. Select **R** to boot in recovery mode. The following menu will appear:

```
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
        P) Physical presence (low security mode)
        R) Reboot
        E) Exit

Option? :
```

   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 115,200 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:

```
... Waiting for the <Enter> key to be hit before the download can
start...
CC
```

   b. When you see the following message, change the serial port characteristics to 9,600 bps:

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

   c. Press **<Enter>** to start installing the image. If the file is a software image, enter the image number:

```
install software image 1 or 2 (hit return to just boot image):
```

   The image install will begin. After the procedure is complete, the Recovery Mode menu will be re-displayed.

```
Extracting images ... Do *NOT* power cycle the switch.
Installing Root Filesystem:
Image signature verified.  100%
Installing Kernel:
Image signature verified.  100%
Installing Device Tree:
Image signature verified.  100%
Installing Boot Loader:  100%
Updating install log. File image installed from xmodem at 18:06:02 on
13-3-2015
Please select one of the following options:
        T) Configure networking and tftp download an image
        X) Use xmodem 1K to serial download an image
        P) Physical presence (low security mode)
        R) Reboot
        E) Exit

Option? :
```

   Continue to Step 7.

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>**. 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 8.3.0_OS.img from TFTP server 10.10.98.100
```

   b. If the file is a software image, you will be prompted to enter an image number:

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

   The following message is displayed when the image download is complete:

```
Image2 updated succeeded
Updating install log. File 8.3.0_OS.img installed from 10.10.98.100 at
15:29:30 on 12-3-2015
Please select one of the following options:
        T) Configure networking and tftp download an image
        X) Use xmodem 1K to serial download an image
        P) Physical presence (low security mode)
        R) Reboot
        E) Exit

Option? :
```

   Continue to Step 7.

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.

Chassis Management Module

The switch management port IP address can only be configured via the CMM web interface. The switch-based configuration interfaces (the ISCLI and BBI) 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:switch[1]> ifconfig
Ethernet ScSE
Enabled
-c  static
-i   0.0.0.0
-s   0.0.0.0
-g   0.0.0.0
system:mm[1]> displaylog
1  I  IOMod_01  04/03/12  08:02:49  (iomodule01) Duplicate route detected
to I/O module iomodule01.
2  I  IOMod_01  04/03/12  08:02:49  (iomodule01) I/O module 1 IP 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 these guidelines:

- 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 these steps:

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 vlag 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 must 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 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 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 *Lenovo Networking OS 8.3 Application Guide*.

Note: The external uplink ports (EXT*x*) cannot be members of management VLANs.

Secure Shell (SSH)

Because SSH key generation is CPU intensive, the 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 flow control.

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 8.3 implementation only allows up to 128 virtual router instances (corresponding to the number of supported IP interfaces). Each virtual router instance can be assigned a unique Virtual Router ID (VRID) between 1 and 255.

New and Updated Features

Lenovo Networking OS 8.3 for 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 8.3 documentation as listed on page 3.

4xVRRP w/VLAG

VRRP can work as Full Active-Active or Half Active-Active under a two tier vLAG topology. Full Active-Active means both two tier vLAGs can route L3 traffic for the related VRRP domain. Half Active-Active means vLAGs will do L2/L3 forwarding for the related VRRP domain based on the local and peer VRRP role.

Allow any port Native VLAN ID to be used for the FIP VLAN Discovery Protocol

This feature removes restriction of having native VLAN 1 on FCoE ports. With this feature, user will be able to configure FCoE ports with any native VLAN (other than FCoE VLAN).

BGP preappend AS Path

Prepending AS path allows the switch to add its own AS number multiple times to the outbound routes to influence the BGP route selection of routers in other AS. With this feature, the switch can influence the BGP route selection so that one path is preferred over another.

This is especially useful when customers are dual home to two different Internet Service Providers (ISPs) and they want to have a primary path to one ISP and use the other ISP as a backup path.

Certificate Signing Request (CSR)

This feature enhances the certificate management capabilities on the switch by incorporating the ability to generate a Certificate Signed Request which can be submitted to an external Certificate Authority (CA) for obtaining a signed certificate. The capability to support CSR and process the CA signed certificate thereof is made available from multiple user interfaces including BBI, SNMP and CLI.

Display ARP entries for VLAN 4095

This is an enhancement to display the ARP table for management VLAN 4095.

Dual Speed 1/10G SFP+ Transceiver

SFP+ ports are supporting now dual speed transceivers (1G/10G). The user is able to configure the speed which will be used.

EHCM 1.02i specification update

Updates EHCM MIB to align with 1.02i specifications.

FlexPort 2.0

Flexible Port mapping now has stacking and System Interconnect Fabric support. Changing the current port-map no longer requires a reboot. Installing and uninstalling FoD keys no longer requires a reboot. When a Trial FoD key expires, the switch no longer remove all configuration files and then reboot. It will just revert to the default port-map of the current license level.

LDAP configurable user name

This feature enables the switch's Lightweight Directory Access Protocol (LDAP) Remote Authentication feature to function with the Microsoft Server LDAP Authentication. This support is implemented by allowing a user to override the default Username Attribute passed in with the LDAP authentication request with a configured string.

No terminal prompting

This feature implements a new CLI command "**`[no] terminal dont-ask`**" to turn off prompting for all CLI commands that would ask for user confirmation to proceed. This command will disable CLI confirmation prompts for the current session only.

NPV - Support for automatic disruptive load balancing

This feature triggers NPV load balancing automatically when FC uplinks are recovered from failure or new FC uplinks are introduced.

Stacking

The following features have been added to stacking:

Improved master failover

On older stack implementation the master to backup failover time used to depend on the number of units in stack and also on the size and complexity of the configuration. With this change the dependence of configuration will be reduced. For example in ring-stack, if the stack was big and the config used to have 1K VLANs, the whole hand-shake during master failover would take about 120-140 seconds. With this change the master failover will take less than 10 seconds.

Until the hand-shake is done between the master and the members, not all the internal structures of the new master are properly updated, so protocol packets loss during this time is to be expected. With this change the stack should become more stable during this transition.

This change can also impact a lot of features during master failover since the timing is also changing a lot and the configuration sent from the master to the members that used to previously be in the stack is suppressed (during master failover, the configuration doesn't change, so there is no reason to send it over once more).

Local preference

This feature tries to enhance the distributed trunk (portchannel) hashing for known unicast traffic. When the feature is enabled, unicast packets that need to exit the stack through a certain portchannel are hashed using only the portchannel member ports of the ingress switch. Otherwise, if it's disabled or if the ingress switch doesn't have any member ports, the packet is hashed using all portchannel member ports across the whole stack. This option reduces the traffic bandwidth over the stacking links.

Master access to Members console

Master access to Members console is a feature that provides the user with the ability to access the Member switches from Master switch CLI by using the attached switch number.

RMON checkpoint support for master/backup synchronization

RMON feature on stacking did not have the checkpoints to synchronize the master with backup in case of a master failover. This support was added so some stats/history collected by the master switch will be kept even after the failover.

Rolling upgrades

This feature will enable the user to upgrade the software on a stack without losing access to the servers during the process. If a rolling upgrade (staggered upgrade) is started the image will be downloaded and programmed on the whole stack and later the units will get rebooted one by one.

SNMP trap relay from members to the master

The SNMP traps generated on member will also be sent to the master switch to be visible there.

Stack Bind command

When used in a ring stack, the command will have the following effect: all the units which are already attached to the stack master, but not yet configured, will get a configurated unit number and all these changes will apply in one step. Before this command to be introduced the user used to bind the units one by one and this used to take more time. If this global bind command is used, the user will not control exactly the configured unit number of the stack members and will need to identify them after the bind is over and do the rest of the configuration accordingly.

Staggered reload

This feature will enable the user to reboot a stack without losing access to the servers during the process. This should not be used as a part of an image upgrade/downgrade process. For image upgrade we have staggered upgrade. It can be used for all other moments when the user feels the need to reboot the whole stack - and the images on the stack members are the same. The units reboot (one by one) logic is similar with the logic from staggered upgrade reloading part.

Syslog save to member or master flash

New user interface commands were introduced to retrieve syslogs from master switch for both master and member switches in the stack – for those members which are attached to the stack. By default any syslog will be saved in the local flash – the flash of the switch where was generated. Commands to enable/disable this flash saving were also added – since syslog server can be used instead.

System information and statistics

The following commands have been added in ISCLI to display system information and statistics (cpu/memory/threads stats) on the master for a specified members:

show processes

show processes cpu [swn <*csnum*>**]**

show processes memory [swn <*csnum*>**]**

show processes thread [swn <*csnum*>**]**

show sys-info [swn <*csnum*>**]**

Also supported in SNMP: mpCpuStatsTable, mpCpuStatsTable, hwInfoTable, cpuUtilProcessStatsTable, chassisInfoTable.

TFTP block size (RFC 2348)

When downloading an image/configuration over the TFTP from a server, a maximum size of pacekts of 512 bytes (TFTP standard) was used.

User-defined description for each switch number of the stack

This provides the ability in the stacking MIB to allow for a user-provided description of each member of the stack for identification purposes. The following objects were added:

- in the ISCLI:

  stack switch_number <*csnum*> **decription**

- in the SNMP:

  .1.3.6.1.4.1.20301.2.5.12.1.1.1.1.4 - SwitchDescrCurCfg

  .1.3.6.1.4.1.20301.2.5.12.1.1.2.1.6 – SwitchDescrNewCfg

Zero-config on member for the stack VLAN

For this feature the user can still change the stack VLAN on each switch in the stack, but when the switches will join a stack only the master stack VLAN value will count because all other switches will overwrite the operational VLAN stack with the stack VLAN pushed by the master in the discovery packets. If the master of the stack will go down, the operational value of the stack VLAN will not change even if the backup switch, now the new operational member of the stack has another value for it. The ISCLI will not actually change, but the functionality will do. Until now if the members have a different stack VLAN than the master, the switch will not be able to join the stack and some errors can be seen on the console, while with the new improvement the stack will form anyway and all the switches in the stack will use the master stack VLAN value.

STP range enhancement

Existing STP commands are enhanced to support configuration of a range of STP groups at a time.

Tech support enhancement

Adds CLI command line text before each corresponding output section in the **`show tech-support`** dump.

UFP support for 8 VNICs per port

Added support for 8 vNICs per port for the UFP feature.

UFP support for PVLAN trunk & promiscuous port

Allow UFP vNIC port to be assigned as Private VLAN trunk and promiscuous port.

Known Issues

This section describes known issues for N/OS 8.3 on the CN4093 10 Gb Converged Scalable Switch.

Note: Please review the Change History documentation posted with the Switch Firmware to check if any of these issues have been fixed in the latest release.

FCoE

VMware command causes FCoE host to loose connectivity with storage. (ID: 7400)

- FCOE connections would be lost when the /sbin/llddpnetmap script is run on the server resulting in detection of multiple LLDP peers on the port.
- The connections would not be restored until the port is shutdown/no shutdown.

Omni Ports

- For timing-sensitive Ethernet protocols and latency-sensitive Ethernet data traffic, we recommend you use the native Ethernet ports instead of the Omni Ports in Ethernet mode.

  Omni Ports are optimized for Fibre Channel (FC) traffic. When in Ethernet mode, the Omni Ports exhibit the following sub-optimal characteristics:

  - Higher latency
  - Slower link event detection

    During heavy activity in the module handling Omni Port processing, additional delays in link event detection times could potentially impact link timing sensitive protocols (including but not limited to VRRP, OSPF, BGP, LACP, VLAG, and IGMP). Extended periods of heavy Omni Port activity is not a normal scenario, but when it happens in conjunction with link timing sensitive protocol usage, unpredictable performance results may be observed. (ID: 64746)

- Egress packets contribute to statistics on Omni Ports even when link is down or transceivers are not present. (ID: 62639)

Stacking

Port mirroring doesn't capture all packets. (ID: 37782)

- In stacking environment port mirroring is not working for ingress/egress traffic with more than one 802.1p priority in scenarios where mirror end monitor ports are configured on different units. As a workaround we recommend that monitor and mirror ports to be configured on the same stack unit.

Statistics

The unicast traffic counter is not incremented for member unit port when sending tagged traffic. (ID: 37800)

- Packets that have invalid length in EtherType/Length field from Ethernet header are counted as multicast packets even though they are unicast packets (invalid means that the length specified is not the same as packet length).

STP

A topology change incorrectly appears when port goes down after disabling STP Portfast. (ID: 38072)

- If STP mode is MSTP, shutting down a port which state is operational spanning tree portfast and not administrative portfast may cause spanning tree topolgy change and messages will be prompted on console. This topology change will not influence the traffic.

Not all 16 STGs converge simultaneously after a root bridge change. (ID: 7478)

- On a setup with 16 active STGs, 3-5 of them converge after 30 seconds instead of 2 sec as normal behavior.

UFP

Traffic behavior when having ETS and Bandwidth mode enabled on different ports. (ID: 5706)

- The bandwidth ratios of PGs for ETS can't be guaranteed when running traffic between ETS UFP-enabled port and Bandwidth mode UFP-enabled port. Based on the current implementation of UFP VLAN, for ingress traffic into UFP-enabled port, it won't copy the inner priority to outer priority. For UFP QoS Bandwidth mode port, all the traffic sent from the server is assigned to PG0, regardless of the priority set in the inner TAG. As a result, for PFC priority and ETS it's only guaranteed minimum bandwidth for PGs.
- Also, when a vPort is expected to be lossless, it should be configured with FCoE network mode.

Virtual Link Aggregation Groups (vLAG)

Mrouter/IGMP group synchronization fails after primary or secondary DUT reload in a vLAG scenario. (ID: 40177)

- When topology change occurs on the STP root which is a VLAG switch and it is rebooted, IGMP groups are removed.
- IGMP groups are removed when STP Root bridge, which is also a vLAG switch, recovers after a reboot.