IBM Flex System Fabric CN4093 10Gb Converged Scalable Switch

Release Notes

for Networking OS 7.7

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.

First Edition (September 2013)

© Copyright IBM Corporation 2013
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

P/N: 00AY510

Release Notes

This release supplement provides the latest information regarding IBM Networking OS 7.7 for CN4093 10Gb Converged Scalable Switch (referred to as CN4093 throughout this document).

This supplement modifies and extends the following Networking OS documentation for use with N/OS 7.7:

- *IBM Networking OS Application Guide for the CN4093 10Gb Converged Scalable Switch*
- *IBM Networking OS Command Reference for the CN4093 10Gb Converged Scalable Switch*
- *IBM Networking OS ISCLI Reference for the CN4093 10Gb Converged Scalable Switch*
- *IBM Networking OS BBI Quick Guide for the CN4093 10Gb Converged Scalable Switch*
- *CN4093 10Gb 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.

© Copyright IBM Corp. 2013

Hardware Support

N/OS 7.7 software is supported on the CN4093 10Gb Converged Scalable Switch for the IBM Flex System. The CN4093 10Gb 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 10Gb Converged Scalable Switch Faceplate

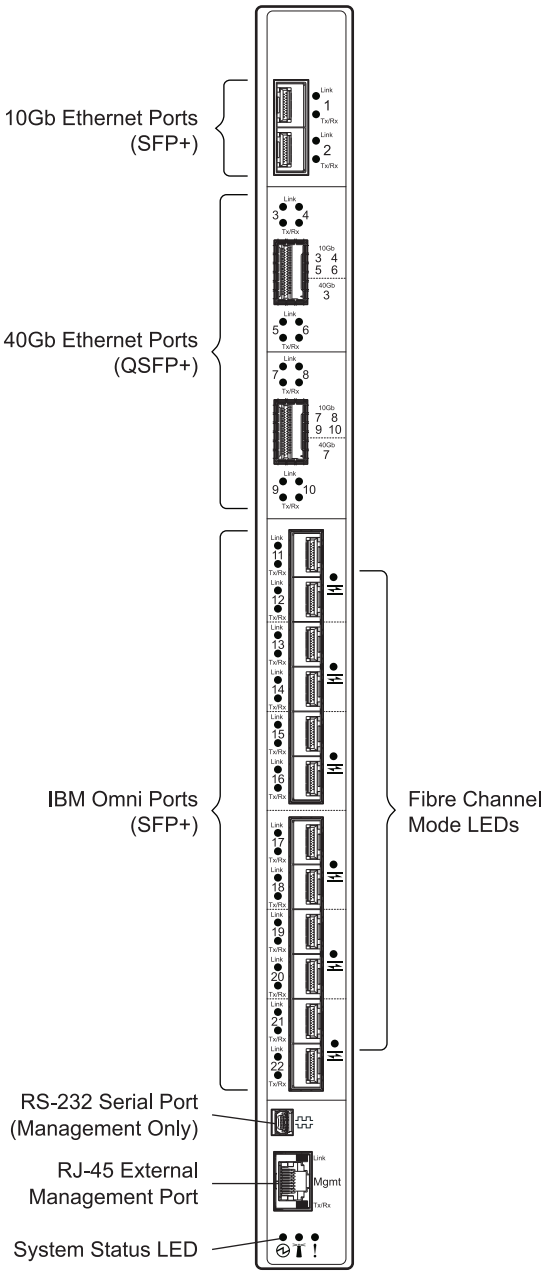

The CN4093 has the following port capacities:

- 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.

© Copyright IBM Corp. 2013

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/systems/support

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

```
>> # /info/sys/gen
```

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 10Gb Converged Scalable Switch. Do not use any prior versions with this device.

Special Software Update Issues

When updating to N/OS 7.7, 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

- In N/OS 7.7, the UID 1 default name is `USERID`, which cannot be modified. However, you are allowed to change the UID password, if required. Changes made to the UID 1 name in N/OS 7.5 will be lost after an upgrade to N/OS 7.7.

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 Software Upgrade" on page 39).

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.

© Copyright IBM Corp. 2013

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 FTP or TFTP directory (for example, `tftpboot`).

4. If required by the FTP or SFTP 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 a 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.

© Copyright IBM Corp. 2013

New and Updated Features

N/OS 7.7 for CN4093 10Gb 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 7.7 documentation as listed on page 3.

DHCP

Host Name Configuration

The CN4093 supports DHCP host name configuration as described in RFC 2132, option 12. DHCP host name configuration is enabled by default.

Host name can be manually configured using the following command:

```
CN 4093(config)# hostname <name>
```

If the host name is manually configured, the switch does not replace it with the host name received from the DHCP server.

After the host name is configured on the switch, if DHCP or DHCP host name configuration is disabled, the switch retains the host name.

The switch prompt displays the host name.

Host name configuration can be enabled/disabled using the following command:

```
CN 4093(config)# [no] system dhcp hostname
```

SYSLOG Server

During switch startup, if the switch fails to get the configuration file, a message can be recorded in the SYSLOG server.

The CN4093 supports requesting of a SYSLOG server IP address from the DHCP server as described in RFC 2132, option 7. DHCP SYSLOG server request option is enabled by default.

Manually configured SYSLOG server takes priority over DHCP SYSLOG server.

Up to two SYSLOG server addresses received from the DHCP server can be used. The SYSLOG server can be learnt over a management port or a data port.

Use the `CN 4093# show logging` command to view the SYSLOG server address.

DHCP SYSLOG server address option can be enabled/disabled using the following command:

```
CN 4093(config)# [no] system dhcp syslog
```

Enhanced Password Security

Switch Login Levels

To enable better switch management and user accountability, three levels or *classes* of user access have been implemented on the CN4093. Levels of access to CLI, Web management functions, and screens increase as needed to perform various switch management tasks. Conceptually, access classes are defined as follows:

- User interaction with the switch is completely passive—nothing can be changed on the CN4093. Users may display information that has no security or privacy implications, such as switch statistics and current operational state information.
- Operators can only effect temporary changes on the CN4093. These changes will be lost when the switch is rebooted/reset. Operators have access to the switch management features used for daily switch operations. Because any changes an operator makes are undone by a reset of the switch, operators cannot severely impact switch operation.
- Administrators are the only ones that may make permanent changes to the switch configuration—changes that are persistent across a reboot/reset of the switch. Administrators can access switch functions to configure and troubleshoot problems on the CN4093. Because administrators can also make temporary (operator-level) changes as well, they must be aware of the interactions between temporary and permanent changes.

Access to switch functions is controlled through the use of unique user names and passwords. Once you are connected to the switch via console, remote Telnet, or SSH, you are prompted to enter a password.

Access to each user level (except admin account) can be disabled by setting the password to an empty value. To disable admin account, use the command: `>>Main# /cfg/sys/access/user/dis`. Admin account can be disabled only if there is at least one user account enabled and configured with administrator privilege.

Strong Passwords

The administrator can require use of Strong Passwords for users to access the CN4093. Strong Passwords enhance security because they make password guessing more difficult.

The following rules apply when Strong Passwords are enabled:

- Minimum length: 8 characters; maximum length: 64 characters
- Must contain at least one uppercase alphabet
- Must contain at least one lowercase alphabet
- Must contain at least one number
- Must contain at least one special character:

  Supported special characters: ! “ # % & ‘ ( ) ; < = >> ? [\] * + , - . / : ^ _ { | } ~
- Cannot be same as the username
- No consecutive four characters can be the same as in the old password

When strong password is enabled, users can still access the switch using the old password but will be advised to change to a strong password while attempting to log in.

© Copyright IBM Corp. 2013

Strong password requirement can be enabled using the following command:

```
>>Main# /cfg/sys/access/user/strongpw ena
```

The administrator can choose the number of days allowed before each password expires. When a strong password expires, the user is allowed to log in one last time (last time) to change the password. A warning provides advance notice for users to change the password.

Locking Accounts

To protect the switch from unauthorized access, the account lockout feature can be enabled. By default, account lockout is disabled. To enable this feature, ensure the strong password feature is enabled. Then use the following command:

```
>> # /cfg/sys/access/user/strongpwd/lockout enable
```

After multiple failed login attempts, the switch locks the user account if lockout has been enabled on the switch.

Re-enabling Locked Accounts

The administrator can re-enable a locked account by reloading the switch or by using the following command:

```
>> # /cfg/sys/access/user/uid <user ID>/clrlock
```

However, the above command cannot be used to re-enable an account disabled by the administrator.

To re-enable all locked accounts, use the following command:

```
>> # /cfg/sys/access/user/strongpwd/clrlock
```

Edge Virtual Bridging (EVB)

The Virtual Station Interface (VSI) database (VSIDB) manager can be configured with an IPv4 or IPv6 address. Use the following command to configure the VSIDB manager IP address:

```
>> Edge Control Protocol Configuration# /cfg/virt/evb/vsidb 1
>> VSI Type DB 1# managrip <IPv4 or IPv6 address>
```

FCoE - Port Trunking

Networking OS 7.7 supports port trunking for FCoE connections. The Link Aggregation (LAG) can be used for separate FCoE traffic, or for Ethernet and FCoE traffic. Ports directly connected to servers cannot be combined in a LAG group.

Uplink ports, connected to the FCF, can be grouped as static or dynamic trunks.

Internal ports cannot be grouped as trunks.

Normal trunk operations such as creating/enabling the trunk, and adding/removing member ports can be performed. When a port is added to a trunk group, FCFs previously detected on the port will be deleted. The deleted FCF may be relearned later. However, this may cause flickering in the network traffic.

Enhanced Transmission Selection (ETS), Priority-based Flow Control (PFC), and Data Center Bridging (DCBX) are configured on a per-port basis. Each port in a trunk must have the same ETS, PFC, and DCBX configuration. When a port ceases to be the trunk group member, its configuration does not change.

Note: FCoE port trunking is not supported in stacking mode.

Fibre Channel

BBI Support

Added BBI support for Fibre Channel configuration.

Note: 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.

SMI-S

Added support for establishing connection with the SMI-S agent via IPv6 management interface using HTTP/HTTPS.

IPv4 Address Conflict Detection

The CN4093 10Gb Converged Scalable Switch uses a simple mechanism to detect if two hosts on the same subnetwork are using the same IPv4 address at the same time. The CN4093 sends a gratuitous ARP request for its own IP address. If it receives an ARP response, it sends a syslog message with the IP address and MAC address of the host that is using its IP address.

The CN4093 sends a gratuitous ARP request in the following situations:

- an IP interface comes up when:
  - the interface is enabled
  - a link comes up
  - a port goes into STP forwarding state
  - a member is added to a VLAN
- the IP address of an IP interface changes

© Copyright IBM Corp. 2013

LACP

Link Aggregation Control Protocol (LACP) is an IEEE 802.3ad standard for grouping several physical ports into one logical port (known as a dynamic trunk group or Link Aggregation group) with any device that supports the standard. Please refer to IEEE 802.3ad-2002 for a full description of the standard.

A port's Link Aggregation Identifier (LAG ID) determines how the port can be aggregated. The Link Aggregation ID (LAG ID) is constructed mainly from the *system ID* and the port's *admin key*, as follows:

- **System ID**: an integer value based on the switch's MAC address and the system priority assigned in the CLI.
- **Admin key:** a port's Admin key is an integer value (1 - 65535) that you can configure in the CLI. Each CN4093 port that participates in the same LACP trunk group must have the same *admin key* value. The Admin key is *local significant*, which means the partner switch does not need to use the same Admin key value.

For example, consider two switches, an Actor (the CN4093) and a Partner (another switch), as shown in Table 1.

Table 1. Actor vs. Partner LACP configuration

Actor Switch	Partner Switch 1
Port 38 (admin key = 100)	Port 1 (admin key = 50)
Port 39 (admin key = 100)	Port 2 (admin key = 50)
Port 40 (admin key = 100)	Port 3 (admin key = 70)

In the configuration shown in Table 1, Actor switch ports 38 and 39 aggregate to form an LACP trunk group with Partner switch ports 1 and 2. Only ports with the same LAG ID are aggregated in the trunk group. Actor switch port 40 is not aggregated in the trunk group because it has a different LAG ID. Switch ports configured with the same admin key on the Actor switch but have a different LAG ID (due to Partner switch admin key configuration or due to partner switch MAC address being different) can be aggregated in another trunk group. i.e. Actor switch port 40 can be aggregated in another trunk group with ports that have the same LAG ID as port 40.

To avoid the Actor switch ports (with the same admin key) from aggregating in another trunk group, you can configure a trunk ID. Ports with the same admin key (although with different LAG IDs) compete to get aggregated in a trunk group. The LAG ID for the trunk group is decided based on the first port that is aggregated in the group. Ports with this LAG ID get aggregated and the other ports are placed in `suspended` mode. As per the configuration shown in Table 1, if port 38 gets aggregated first, then the LAG ID of port 38 would be the LAG ID of the trunk. Port 40 would be placed in suspended mode. When in suspended mode, a port transmits only LACP data units (LACPDUs) and discards all other traffic.

A port may also be placed in suspended mode for the following reasons:

- When LACP is configured on the port but it stops receiving LACPDUs from the partner switch.
- When the port has a different LAG ID because of the partner switch MAC or port LACP key being different. For example: when a switch is connected to two partners.

Trunk ID can be configured using the following command:

```
>># /cfg/l2/trunk <19-36>/adminkey <adminkey of the LAG>
```

Manual Reflective Relay

Reflective Relay (RR) is an operation where the switch forwards a frame back to the port on which it arrived if the destination MAC address is on the same port. When an EVB profile is configured on a port, RR is automatically enabled on the port after capability exchange with the peer, using the IEEE802.1QBG protocol. This is the usual mode of operation.

When the switch interoperates with devices that do not support IEEE 802.1QBG protocols, RR can be manually configured using the following command:

```
>>Main# /cfg/port <port num>/rrforce ena
```

Manual RR and EVB profile cannot be configured on a port at the same time.

Network Time Protocol (NTP)

New commands added to provide the following:

- Detailed information on NTP association:

```
CN 4093(config)# show ntp associations

address            ref clock            st           when(s)
offset(s)
#192.168.13.33      -                    16           -          0
*192.168.13.57      192.168.1.111        3            32         11

* - synced
# - unsynced
```

- Minimize number of syslogs when NTP synchronization fails or system clock is updated:

```
CN 4093(config)# [no] ntp sync-logs        (Enable logs for information on sync
failures)

CN 4093(config)# [no] ntp offset <0-86400>   (Set minimum clock change to trigger
logs)
```

© Copyright IBM Corp. 2013

SNMP

MIBs

- Added MIBs required for accessing MLDv2 information.

Community Strings

Added support for 8 read-only and read-write community strings for SNMP v1 and SNMPv2. If any one of the community strings is matched, then read-only or read-write access will be granted. Use the following commands to add or delete community strings:

```
To add:
>> #/cfg/sys/ssnmp/rcomm-additional <1-32 characters>
(or)
>> #/cfg/sys/ssnmp/wcomm-additional <1-32 characters>

To delete:
>> #/cfg/sys/ssnmp/rcomm-delete <1-32 characters>
(or)
>> #/cfg/sys/ssnmp/wcomm-delete <1-32 characters>
```

Static Multicast ARP

The Microsoft Windows operating system includes the Network Load Balancing (NLB) technology that helps to balance incoming IP traffic among multi-node clusters. In multicast mode, NLB uses a shared multicast MAC address with a unicast IP address. Since the address resolution protocol (ARP) can map an IP address to only one MAC address, port, and VLAN, the packet reaches only one of the servers (the one attached to the port on which the ARP was learnt).

To avoid the ARP resolution, you must create a static ARP entry with multicast MAC address. You must also specify the list of ports through which the multicast packet must be sent out from the gateway or Layer 2/Layer 3 node.

With these configurations, a packet with a unicast IPv4 destination address and multicast MAC address can be sent out as per the multicast MAC address configuration. NLB maps the unicast IP address and multicast MAC address as follows:

Cluster multicast MAC address: 03-BF-W-X-Y-Z; where W.X.Y.Z is the cluster unicast IP address.

You must configure the static multicast ARP entry only at the Layer 2/Layer 3 or Router node, and not at the Layer 2-only node.

IBM Networking OS supports a maximum of 20 static multicast ARP entries.

Note: If you use the ACL profile or IPMC-OPT profile, an ACL entry is consumed for each Static Multicast ARP entry that you configure. Hence, you can configure a maximum of 256 ACL and multicast MAC entries together.The ACL entries have a higher priority. In the default profile, the number of static multicast ARP entries that you configure does not affect the total number of ACL entries.

Configuring Static Multicast ARP

To configure multicast MAC ARP, you must perform the following steps:

- Configure the static multicast forwarding database (FDB) entry: Since there is no port list specified for static multicast ARP, and the associated MAC address is multicast, you must specify a static multicast FDB entry for the cluster MAC address to limit the multicast domain. If there is no static multicast FDB entry defined for the cluster MAC address, traffic will not be forwarded. Use the following command:

```
>> Main# /cfg/l2/fdb/mcast add <cluster MAC address> <port(s)>
```

- Configure the static multicast ARP entry: Multicast ARP static entries should be configured without specifying the list of ports to be used. Use the following command:

```
>> Main# /cfg/l3/arp/static add <destination unicast IP address> <destination
multicast MAC address> <cluster VLAN number>
```

Configuration Example

Consider the following example:

- Cluster unicast IP address: 10.10.10.42
- Cluster multicast MAC address: 03:bf:0A:0A:0A:2A
- Cluster VLAN: 42
- List of individual or port trunks to which traffic should be forwarded: 54 and 56

Following are the steps to configure the static multicast ARP based on the given example:

1. Configure the static multicast FDB entry.

```
>> Main# /cfg/l2/fdb/mcast add 03:bf:0A:0A:0A:2A 42 54 56
```

2. Configure the static multicast ARP entry:

```
>> Main# /cfg/l3/arp/static add 10.10.10.42 03:bf:0A:0A:0A:2A 42
```

You can verify the configuration using the following commands:

- Verify static multicast FDB entry:

```
>> Main# /info/l2/fdb/mcast/find 03:bf:0A:0A:0A:2A

 Multicast Address  VLAN  Port(s)
 -----------------  ----  ---------
 03:bf:0A:0A:0A:2A    42  54 56
```

© Copyright IBM Corp. 2013

- Verify static multicast ARP entry:

```
>> Main# /info/l3/arp/dump

 Current ARP configuration:
  rearp 5
 Current static ARP:
   ip              mac               port  vlan
   --------------- ----------------- ----- ----
   10.10.10.42     03:bf:0A:0A:0A:2A         42
 ------------------------------------------------
 Total number of arp entries : 2
     IP address    Flags    MAC address    VLAN  Age Port
   --------------- ----- ----------------- ----  --- ----
   10.10.10.1       P    fc:cf:62:9d:74:00   42
   10.10.10.42      P    03:bf:0A:0A:0A:2A   42      0
```

Limitations

- You must configure the ARP only in the Layer 2/Layer 3 node or the router node but not in the Layer 2-only node. Networking OS cannot validate if the node is Layer 2-only.
- The packet is always forwarded to all the ports as specified in the Multicast MAC address configuration. If VLAN membership changes for the ports, you must update this static multicast MAC entry. If not, the ports, whose membership has changed, will report discards.
- ACLs take precedence over static multicast ARP. If an ACL is configured to match and permit ingress of unicast traffic, the traffic will be forwarded based on the ACL rule, and the static multicast ARP will be ignored.

Switch Partition

Switch Partition (SPAR) facilitates the creation of multiple partitions within a switch to form a virtual switching context with respect to data plane partition of a switch. Each SPAR defined in a switch represents a switch partition in the data plane hardware. Data plane traffic is not shared between SPARs on the same switch. SPAR provides a simple Ethernet interface connectivity option for connecting Blade server chassis to network infrastructure.

SPAR is implemented as a dedicated VLAN with a set of internal server ports and a single uplink port or link aggregation (LAG). SPAR does not support multiple uplink ports or LAGs, which eliminates the possibility of misconfiguration or loop creation. A port can only be a member of one SPAR.

SPAR operates as a Layer 2 broadcast network. Hosts on the same VLAN, attached to a SPAR can communicate with each other and with the upstream switch. Hosts on the same VLAN, but attached to different SPARs, communicate via the upstream switch.

The default SPAR is SPAR-0.

SPAR operates in two processing modes. The default mode is pass-through domain mode.

- Local Domain: In local domain processing mode, VLAN classification and assignment is based on the user-defined VLAN.

- Pass-through Domain: In pass-through domain processing mode, VLAN classification and assignment is based on the outer tag, which contains the unique domain VLAN ID of the SPAR. The inner tag with the user-defined VLAN remains unchanged.

Note: UFP and SPAR cannot be configured together.

Local Domain Processing

Each SPAR on a switch has a unique VLAN ID, which implicitly provides data separation between SPARs. If multiple networks share the uplink, the upstream switch port must be configured as a 802.1Q trunk port so it can process multiple VLAN traffic from a SPAR. The SPAR domain uses a single uplink port or LAG shared among all the VLANs. For link redundancy or greater bandwidth, the uplinks can be grouped as static or LACP LAG.

If a VLAN is defined on multiple SPARs, the egress port mask is used to prevent communication between the SPARs in the same local domain VLAN. Since port membership of each SPAR is unique, the egress port mask ensures that different SPAR ports in the same local domain VLAN do not communicate with each other.

In local domain processing, all SPAR ports must have the following settings:

- Tagging must be enabled. (`>>Main# /cfg/port/tag ena`)
- TAGIPVID is disabled on all SPAR ports. (`>>Main# /cfg/port/tagipvid dis`)
- PVID is based on any VLAN defined in SPAR. (`>>Main# /cfg/port/pvid <VLAN number>`)

Layer 2 Switching

The CN4093 10Gb Converged Scalable Switch learns MAC+VLAN on a per-port basis and not on a per-SPAR basis. If the switch detects the same MAC+VLAN combination on multiple SPAR ports, it considers it as a station move. Hence, SPAR deployment in local domain topology is restricted within distinct physical networks. Multiple SPAR domains within a physical network must not share the same set of VLANs to avoid the same MAC+VLAN combination appearing on more than one SPAR.

Pass-Through Domain Processing

In this processing mode, each SPAR is identified by its unique VLAN domain ID. Packets are classified based on the SPAR domain ID (outer tag) and not the user-defined VLAN (inner tag). SPAR ports must be configured in tunnel mode.

SPAR provides single or multiple VLAN connectivity through a single uplink port or LAG (static or LACP) without requiring VLAN definition within the SPAR domain.

Pass-through domain operates in Q-In-Q mode. Inside SPAR, different user-defined VLAN traffic is classified into single S-VLAN associated with the SPAR.

Although the uplink can be shared by multiple networks using the pass-through domain, SPAR will not be server-VLAN aware. Hence, multiple VLAN traffic will be mixed together in a single broadcast domain(i.e. broadcast traffic on different VLANs from the upstream network will reach all servers attached to the SPAR pass-through domain). The servers drop the packets if they do not belong to the desired VLAN. The pass-through implementation uses ingress VLAN tagging i.e. TAGIPVID is enabled on all SPAR ports.

© Copyright IBM Corp. 2013

In pass-through domain processing mode, all SPAR ports must have the following settings:

- TAGPVID is disabled. (>>Main# /cfg/port/tagpvid dis)
- TAGIPVID is enabled on all SPAR ports. (>>Main# /cfg/port/tagipvid ena)
- PVID is based on the SPAR DVLAN. (>>Main# /cfg/port/pvid <DVLAN number>)

SPAR VLAN Management

SPAR VLANs use the same 4000 VLAN space available for other applications/features on the switch. The VLAN ID can be in the range of 2 - 4094. VLAN 1 and the management VLAN 4095 are reserved for the default switch context.

A VLAN assigned to a SPAR cannot be used for any other switch application. Similarly, VLAN used by any other switch application cannot be assigned to a SPAR.

SPAR member ports cannot be members of any other VLAN.

Example Configuration

This example includes configuration of SPAR 1 in pass-through mode and SPAR 2 in local domain mode.

1. Create SPAR 1.

```
>>Main# /cfg/spar 1
```

2. Add uplink port to SPAR 1.

```
>>SPAR 1 Configuration# uplink
>>SPAR 1 Uplink Configuration# port Ext 1
>>SPAR 1 Uplink Configuration# ..
```

3. Configure domain mode.

```
>>SPAR 1 Configuration# domain
>>SPAR 1 Domain Configuration# mode passthrough
```

4. Configure SPAR VLAN.

```
>>SPAR 1 Domain Configuration# dvlan
>>SPAR 1 Default VLAN Domain Configuration# sparvid 4081
```

5. Add member ports.

```
>>SPAR 1 Default VLAN Domain Configuration# addsport INTA5-INTA10
>>SPAR 1 Default VLAN Domain Configuration# ..
>>>SPAR 1 Domain Configuration# ..
```

6. Enable SPAR 1.

```
>>SPAR 1 Configuration# ena
>>SPAR 1 Configuration# ..
```

7. Create SPAR 2.

```
>>Configuration# spar 2
```

8. Add uplink port to SPAR 2.

```
>>SPAR 2 Configuration# uplink port Ext 2
```

9. Configure domain mode.

```
>>SPAR 2 Configuration# domain
>>SPAR 2 Domain Configuration# mode local
```

10. Configure SPAR VLAN.

```
>>SPAR 2 Domain Configuration# dvlan
>>SPAR 2 Default VLAN Domain Configuration# sparvid 4082
```

11. Add member ports.

```
>>SPAR 1 Default VLAN Domain Configuration# addsport INTA11-INTA14
>>SPAR 1 Default VLAN Domain Configuration# ..
```

12. Configure local domain 1.

```
>>SPAR 2 Domain Configuration# lvlan 1
>>SPAR 2 Local VLAN Domain 1 Configuration# vid10
>>SPAR 2 Local VLAN Domain 1 Configuration# addsport INTA11-INTA14
>>SPAR 2 Local VLAN Domain 1 Configuration# ena
>>SPAR 2 Local VLAN Domain 1 Configuration# ..
```

13. Configure local domain 2.

```
>>SPAR 2 Domain Configuration# lvlan 2
>>SPAR 2 Local VLAN Domain 2 Configuration# vid20
>>SPAR 2 Local VLAN Domain 2 Configuration# addsport INTA11-INTA14
>>SPAR 2 Local VLAN Domain 2 Configuration# ena
>>SPAR 2 Local VLAN Domain 2 Configuration# ..
```

14. Configure local domain 3.

```
>>SPAR 2 Domain Configuration# lvlan 3
>>SPAR 2 Local VLAN Domain 3 Configuration# vid30
>>SPAR 2 Local VLAN Domain 3 Configuration# addsport INTA11-INTA14
>>SPAR 2 Local VLAN Domain 3 Configuration# ena
>>SPAR 2 Local VLAN Domain 3 Configuration# ..
>>SPAR 2 Domain Configuration# ..
```

© Copyright IBM Corp. 2013

15. Enable SPAR 2.

```
>>SPAR 2 Configuration# ena
```

Unsupported Features

The following features are not supported when SPAR is configured:

- 802.1x
- FCoE
- Hotlinks
- IGMP
- Layer 3 Configuration
- Management VLAN
- Private VLAN
- Protocol VLAN
- Edge Virtual Bridging
- sFlow
- Stacking
- STP, RSTP, MRSTP, PVST
- UFP
- vLAG
- VMAP
- VMready
- VNIC

Limitations

The following restrictions apply:

- SPAR VLANs are automatically added to Spanning Tree Group (STG) 1 regardless of the STP mode selected. STP is turned off on all SPAR ports.
- Trunks (static or LACP) created on the default switch cannot reference any SPAR ports. Trunks must first be defined in the SPAR context before they can be used by SPAR. Use the commands in the following menus to define trunks in the SPAR context:
  `>>Main# /cfg/spar <number>/uplink` menu;
  `>>Main# /cfg/layer 2` menu.
- ACLs defined on the default switch can be used for SPAR ports. However, the following restrictions apply:
  - An ACL cannot be shared across SPAR ports if:
    - An exit port (`>>Main# /cfg/acl/acl <number>/egrport`) is used as a filtering criteria and the exit port does not belong to the same SPAR as the port on which the ACL is applied.
    - A monitor port is used as a filtering criteria, and the monitor port does not belong to the same SPAR as the mirrored port and is not defined on the default switch.
  - The above ACL restrictions apply to all ACLs defined in an ACL group.

- Port mirroring can be configured on SPAR ports with the following restrictions:
  - The monitor port must belong to the same SPAR as the mirrored port, or must be defined on the default switch.
- Layer 2 failover features can be configured on SPAR ports. However, the Layer 2 failover AMON option is not supported. Only the Layer 2 failover MMON option can be used when all ports defined within the trigger belong to the same SPAR.

OSPFv3

Enhancements based on RFC5340.

Persistent Terminal Length

The screen length for the current session can be set using the command: `>>Main# lines` *<0-300>*.

However, when the switch is reloaded, the screen length is set to default.

To set the screen length to be persistent across multiple sessions, use the following commands:

Telnet and SSH:

```
>>Main# /cfg/sys/linevty <0-300>
```

Console:

```
>>Main# /cfg/sys/linecons <0-300>
```

The commands to set a persistent screen length are saved in the startup configuration and will be applied even when the switch is reloaded. If you need to change the screen length for a particular session, you can do so using the command for setting the current session's screen length.

Unified Fabric Port

Virtualization is widely deployed in data centers for isolating traffic and allocating bandwidth. However, virtualization introduces the overhead of managing two network entities: server and network. Unified Fabric Port (UFP) helps reduce this overhead by providing the ability to manage server side network functionality of the Network Interface Card (NIC) by applying network policies defined on the switch.

UFP is an architecture that logically subdivides a high-speed physical link connecting to a server NIC or to a Converged Network Adapter (CNA). UFP provides a switch fabric component to control the NIC. To the server operating system (OS) or to the hypervisor, each channel, appears as an independent physical NIC. Each channel has a unique identity and profile that defines its properties and functionality. The server communicates with the switch over the channel as defined in the channel profile. The channels share the high-speed physical link bandwidth.

For each channel, vNIC on the server side communicates with virtual port on the switch side. Any 10 Gbps internal (server) port can be configured as an UFP port.

© Copyright IBM Corp. 2013

Figure 2. UFP vPorts

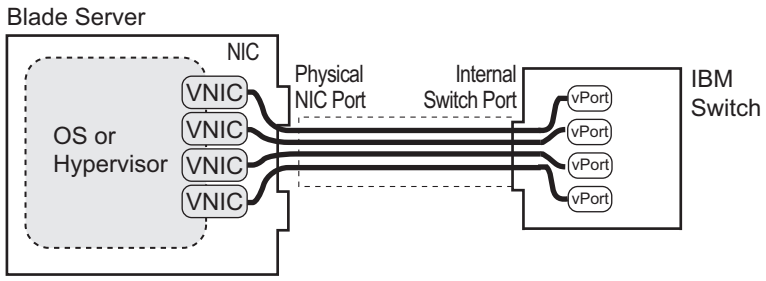

Note: The CN4093 10Gb Converged Scalable Switch does not support UFP and VNIC configuration simultaneously.

UFP Protocol

The UFP protocol is a link-level protocol that runs as a separate instance for each physical communication link established between a server NIC and a switch port. The UFP protocol has three categories of operation:

- Channel Initialization: The server NIC and the switch port negotiate the number of channels and establish channel identifiers. Each UFP channel has a data component and a control component. The two components have the same UFP channel ID.
- Channel Control: For an established channel. the switch can modify channel properties by sending a control message on the UFP channel. While the channel ID is the same for the control and data components, the destination MAC address of the control message frame is a well-known address.
- Channel Data Path: The UFP protocol supports two types of data paths: local domain and pass-through domain. Local domain includes a server with multiple NICs connecting to a single physical switch domain with a single VLAN domain. Pass-though domain includes a server with multiple NICs connecting to multiple physical switching domains, where each domain has its own VLAN.

Limitations

The following restrictions apply when configuring UFP:

- FCoE must be configured only on vPort 2.
- If using Emulex NIC, FCoE can be configured on vPort 0 or vPort 1 of an ITE server NIC.
- UFP port in FCoE mode cannot operate with FIP Auto VLAN feature.
- UFP does not support VMready Local Group configuration.
- UFP cannot be configured in stacking mode.
- VLANs having member vPorts configured in trunk or access modes cannot have member vPorts configured in tunnel mode.
- vPorts on a physical port, if configured in trunk or access mode, must be members of separate VLANs.
- A tunnel mode uplink port can be member of only one VLAN.
- A vPort in trunk mode can be a member of up to 32 VLANs.
- VLANs 4002-4005 are reserved for outer tagging.

- VLAN translation is not applied on packets that egress an UFP port.
- UFP bandwidth is guaranteed lossless only for unicast traffic.

When CEE is on, FCoE vPort must be used for lossless priority traffic. For loss-tolerant priority traffic, a non-FCOE UFP vPort must be used. If lossless and loss-tolerant traffic is mixed, the lossless property of FCoE vPort is not guaranteed.

UFP Control

To enable UFP.

```
>>Main# /cfg/virt/ufp on
```

Channel Initialization

The channel initialization process sequence is as follows:

1. The server NIC sends a UFP channel request to the switch.
2. The switch sends back a UFPCHAN Type-Length-Value (TLV) with the number of UFP channels available. If the switch can satisfy the number channels requested by the server NIC, it sends an acknowledgement and provides channel identifiers for each requested channel.
3. Data and control traffic flowing between the server NIC and the switch port is tagged with the UFP channel identifier to which it belongs.

To initialize channels, UFP uses the Channel Discovery and Configuration Protocol (CDCP) TLV. The fields of the CDCP TLV and exchange sequence of TLVs apply unmodified to UFP channel initialization. After a UFP channel is assigned a channel ID, the switch can change the channel ID at any time.

After the establishment of a UFP channel, the switch and server NIC exchange channel properties in Edge Virtual Bridging (EVB) TLVs. For UFP channels, the bits for the STD and Edge Control Protocol (ECP) capabilities are set in the EVB TLV. The exchanged EVB TLVs provide an indication of ECP readiness at the switch and at the server NIC. Communication between UFP channels on the same physical switch port is not allowed.

Channel Control

After UFP channels are established, the server NIC and switch port exchange UFP messages for control and configuration of each UFP channel. The UFP messages are in a standard TLV format and are of three types:

- Configuration TLV: Used by the switch to send configuration information to the server NIC. One configuration TLV is supported: NIC-Props TLV.
- Operational TLV: Used by the switch and the server NIC to perform runtime operations. Two operational TLVs are supported: Link-Down TLV and Link-Up TLV.
- Information/Statistics TLV: Used for analysis of information and statistics. One TLV is supported: End-of-TLV.

UFP messages are exchanged using ECP frames. The UFP TLV includes the type, length, flags, status, and payload fields.

Channel Data Path

The UFP protocol supports two types of channels:

- Type 1 or local domain: Untagged frames are associated with the channel VLAN ID. If tagging is enabled on the switch port, the frames belong to the VLAN specified in the tag.
- Type 2 or pass-through domain: All tagged and untagged frames are associated with the channel VLAN ID. The tags inserted by the OS are passed through without any modification.

The channel type is set via the EVB TLV sent by the switch to the server NIC. A *type* value of 001b indicates local domain, and a value of 010b indicates pass-through domain.

The switch can change the channel type at any time by sending an updated UFP TLV.

Local Domain

In local domain data path type, a server NIC connects with a switch port that belongs to a single VLAN domain. Switching is based on the inner VLAN tag controlled by the server (or Hypervisor), or on the port PVID. All Layer 2 and Layer 3 features are supported in local domain type.

Pass-through Domain

In pass-through domain data path type, a server NIC may connect with a switch port that belongs to multiple VLAN domains. Each UFP channel is a separate VLAN domain. Switching is based on the outer VLAN tag inserted by the switch at ingress. The outer VLAN tag is based on the UFP channel VLAN ID. The switch strips the outer tag on egress. The inner VLAN tag, controlled by the server (or Hypervisor), is not modified.

To implement pass-through domain, you must configure the virtual port in tunnel mode using the following command:

```
>>Main# /cfg/virt/ufp/port <num>/vport <num>/network mode tunnel
```

Virtual Port Modes

A single physical switch port is configured with virtual ports (vPorts). Each UFP channel connects the server NIC with a switch vPort. Properties, such as native VLAN and bandwidth, defined for a vPort are applied to the traffic that belongs to the vPort.

Note: A maximum of four vPorts can be configured per physical switch port.

vPort-VLAN Mapping

In local domain data path type, the switch and server identify the vPort/vNIC based on the port and VLAN tag in the incoming and outgoing packets. Since no two vPorts carry traffic for the same VLAN, the port+VLAN combination can be uniquely mapped to a vPort. When a vPort is initialized, the switch communicates a list of allowed VLANs, including the native VLAN to which the vPort belongs, to the server NIC. The server NIC uses this information to filter incoming and outgoing traffic based on the VLAN. All packets are single tagged with the vPort's native VLAN.

vPort-S-Tag Mapping

A vPort can also be identified with an S-tag (service tag or outer tag). When a vPort is initialized, the switch communicates the UFP channel ID of the vPort to the server NIC. When the server NIC or switch transmit frames, they add this S-tag to indicate the vPort or vNIC to which the packet is being transmitted. No VLAN mapping is required. Such packets can be single tagged (with native VLAN) or double tagged (with S-tag).

UFP vPort Mode

The UFP mode is configured based on the type of switching domain (single VLAN or multiple VLANs) a vPort is being connected to. For local domain data path types, trunk or access mode is configured. For pass-through domain data path type, tunnel mode is configured.

UFP vPort mode can be configured using the following command:

```
>>Main# /cfg/virt/ufp/port <num>/vport <num>/network mode {access|trunk|tunnel|fcoe}

Dafult mode is 'tunnel'
```

Tunnel Mode

In tunnel mode, a vPort can belong to only one VLAN. An outer tag with the vPort's VLAN ID is inserted in packets that egress the vPort. The inner VLAN tag remains unchanged. The switch processes packets based on the outer tag. When all the ports or vPorts that belong to a particular VLAN are placed in tunnel mode, they belong to one pass-through domain.

Tunnel mode of operation is useful in virtualized environments where it is desired to place all virtual machine (VM) data traffic, which needs to be sent to an upstream switch for Layer 2 or Layer 3 processing, in one domain. In such cases, the UFP port or vPort must be in tunnel mode and the upstream switch port must be in 802.1Q trunk mode.

Note: Two vPorts on a physical port cannot be members of the same VLAN.

Figure 3. Packet pass-through in Tunnel Mode

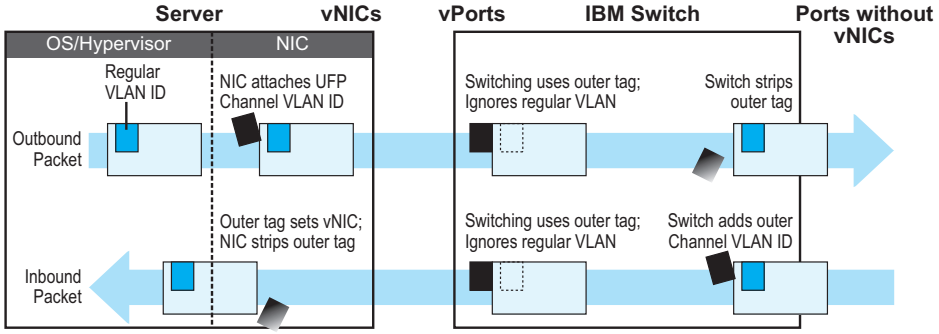

© Copyright IBM Corp. 2013

802.1Q Trunk Mode

In trunk mode, a vPort can carry packets that have inner tag belonging to up to 32 VLANs. This is restricted to a maximum of 4000 VLANs for all UFP vPorts configured on the switch. For each VLAN in the inner tag, a VLAN translation entry is required.

Note: Two vPorts operating in trunk mode on the same physical port cannot carry the same set of VLANs in the inner tag.

Figure 4. Packet passing through in Trunk Mode

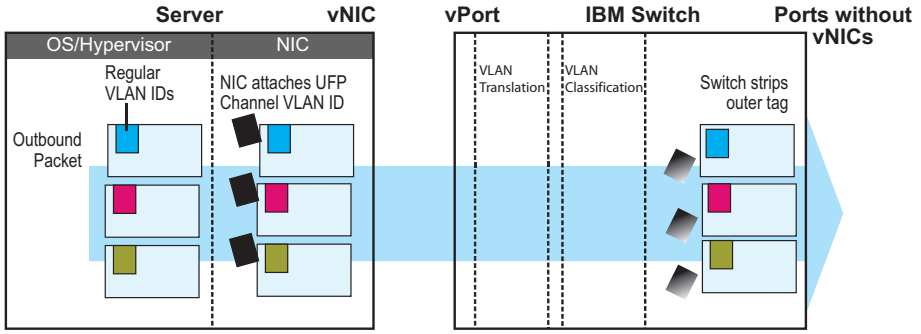

Access Mode

In access mode, a vPort carries packets with inner tags that belong to one VLAN. The vPort is associated with the VLAN defined by using the command:

```
>>Main# /cfg/virt/ufp/port <port number>/vport <port number>/network/defvlan <VLAN Nmber>
```

FCoE Mode

A vPort configured in FCoE mode can only be attached to a Fibre Channel (FC) VLAN. Only one vPort on a physical port can be configured in FCoE mode. A vPort in FCoE mode operates as a local domain data path type with packets being single tagged.

UFP Bandwidth Provisioning

UFP provides one mode of bandwidth provisioning for vPort: Strict Bandwidth Provisioning Mode.

Strict Bandwidth Provisioning Mode

In this mode, the switch and NIC apply bidirectional bandwidth control on the vPort as per the defined configuration. By default, a bandwidth of 2.5 Gbps per vPort is guaranteed. If other vPorts are idle, the bandwidth of a vPort can be up to 10 Gbps. A minimum bandwidth of 1 Gbps is provisioned which can be raised by 100 Mbps increments. The sum of the minimum bandwidth guaranteed for all vPorts together cannot exceed the capacity of the physical link.

A vPort can also be configured with a maximum bandwidth.

This mode works with the port scheduler to avoid unintended packet drop due to policing through EFP metering block. If flow control is enabled, the switch provides a no drop packet forwarding behavior, thereby improving end-to-end Transmission Control Protocol (TCP) throughput performance.

Note: If a vPort is configured with low upper limit, it might lead to head-of-line (HOL) congestion on egress port.

By default, uplink ports have a separate traffic class for storage traffic with guaranteed bandwidth. Rest of the bandwidth is shared equally among other traffic.

Use the following command to configure strict bandwidth provisioning:

```
>>Main# /cfg/virt/ufp/port <num>/vport <num>/qos bw {minbw|maxbw}

minbw - Set minimum guaranteed bandwidth
maxbw - Set maximum allowed bandwidth
```

UFP Configuration Examples

Following is an example configuration of UFP vPorts in access mode.

Example 1: Access Mode

1. Turn on UFP.

```
>>Main# /cfg/virt/ufp on
```

2. Configure internal port as UFP.

```
>>UFP Global Configuration# port INTA1
>>Port INTA1 UFP Configuration# ena
```

3. Configure virtual port.

```
>>Port INTA1 UFP Configuration# vport 1
>>Virtual Port INTA1.1 Configuration# ena
```

4. Configure vPort access mode.

```
>>Virtual Port INTA1.1 Configuration# network
>>Virtual Port INTA1.1 Network Configuration# mode access
```

5. Configure vPort default VLAN.

```
>>Virtual Port INTA1.1 Network Configuration# defvlan 100
```

6. Ensure tagging is disabled on vPort.

```
>>Virtual Port INTA1.1 Network Configuration# deftag d
>>Virtual Port INTA1.1 Network Configuration# ..
```

© Copyright IBM Corp. 2013

7. Specify QoS parameters for the vPort.

```
>>Virtual Port INTA1.1 Configuration# qos
>>Virtual Port INTA1.1 QOS Configuration# bw
>>Virtual Port INTA1.1 Bandwidth Configuration# minbw 25                (in percentage)
>>Virtual Port INTA1.1 Bandwidth Configuration# maxbw 100               (in percentage)
>>Virtual Port INTA1.1 Bandwidth Configuration# /
```

8. Configure PVID of external port 1.

```
>>Main# /cfg/port EXT1
>>Port EXT1# pvid 100
>>Port EXT1# /
```

9. Apply the configuration.

```
>>Main# apply
```

Example 2: Trunk Mode

Following is an example configuration of UFP vPorts in trunk mode.

1. Turn on UFP.

```
>>Main# /cfg/virt/ufp on
```

2. Configure internal port 1 as UFP.

```
>>UFP Global Configuration# port INTA1
>>Port INTA1 UFP Configuration# ena
```

3. Configure virtual port.

```
>>Port INTA1 UFP Configuration# vport 3
>>Virtual Port INTA1.1 Configuration# ena
```

4. Configure vPort trunk mode.

```
>>Virtual Port INTA1.3 Configuration# network
>>Virtual Port INTA1.3 Network Configuration# mode trunk
```

5. Configure vPort default VLAN.

```
>>Virtual Port INTA1.3 Network Configuration# defvlan 100
```

6. Ensure tagging is disabled on vPort.

```
>>Virtual Port INTA1.3 Network Configuration# deftag d
>>Virtual Port INTA1.3 Network Configuration# ..
```

7. Specify QoS parameters for the vPort.

```
>>Virtual Port INTA1.3 Configuration# qos
>>Virtual Port INTA1.3 QOS Configuration# bw
>>Virtual Port INTA1.3 Bandwidth Configuration# minbw 25                  (in percentage)
>>Virtual Port INTA1.3 Bandwidth Configuration# maxbw 100                 (in percentage)
>>Virtual Port INTA1.3 Bandwidth Configuration# ..
>>Virtual Port INTA1.3 Network Configuration# ..
>>Virtual Port INTA1.3 Configuration# ..
>>Port INTA1 UFP Configuration# ..
>>UFP Global Configuration# ..
```

8. Configure internal port 2 as UFP.

```
>>UFP Global Configuration# port INTA2
>>Port INTA2 UFP Configuration# ena
```

9. Configure virtual port.

```
>>Port INTA2 UFP Configuration# vport 3
>>Virtual Port INTA2.3 Configuration# ena
```

10. Configure vPort trunk mode.

```
>>Virtual Port INTA2.3 Configuration# network
>>Virtual Port INTA2.3 Network Configuration# mode trunk
```

11. Configure vPort default VLAN.

```
>>Virtual Port INTA2.3 Network Configuration# defvlan 100
```

12. Ensure tagging is disabled on vPort.

```
>>Virtual Port INTA2.3 Network Configuration# deftag d
>>Virtual Port INTA2.3 Network Configuration# ..
```

13. Specify QoS parameters for the vPort.

```
>>Virtual Port INTA2.3 Configuration# qos
>>Virtual Port INTA2.3 QOS Configuration# bw
>>Virtual Port INTA2.3 Bandwidth Configuration# minbw 25                  (in percentage)
>>Virtual Port INTA2.3 Bandwidth Configuration# maxbw 100                 (in percentage)
>>Virtual Port INTA2.3 Bandwidth Configuration# /
```

14. Enable tagging on external port 1.

```
>>Main# /cfg/port EXT1
>>Port EXT1# tag ena
>>Port EXT1# pvid 100
>>Port EXT1# /
```

© Copyright IBM Corp. 2013

15. Configure VLAN 200 parameters.

```
>>Main# /cfg/l2/vlan 200
>>VLAN 200# ena
>>VLAN 200# def EXT1
>>VLAN 200# addvport INTA1.3
>>VLAN 200# addvport INTA2.3
>>VLAN 200# /
```

16. Configure VLAN 300 parameters.

```
>>Main# /cfg/l2/vlan 300
>>VLAN 300# ena
>>VLAN 300# def EXT1
>>VLAN 300# addvport INTA1.3
>>VLAN 300# addvport INTA2.3
>>VLAN 300# /
```

17. Apply the configuration.

```
>>Main# apply
```

Example 3: Tunnel Mode

Following is an example configuration of UFP vPorts in tunnel mode.

1. Turn on UFP.

```
>>Main# /cfg/virt/ufp on
```

2. Configure internal port as UFP.

```
>>UFP Global Configuration# port INTA1
>>Port INTA1 UFP Configuration# ena
```

3. Configure virtual port.

```
>>Port INTA1 UFP Configuration# vport 1
>>Virtual Port INTA1.1 Configuration# ena
```

4. Configure vPort access mode.

```
>>Virtual Port INTA1.1 Configuration# network
>>Virtual Port INTA1.1 Network Configuration# mode tunnel
```

5. Configure vPort default VLAN.

```
>>Virtual Port INTA1.1 Network Configuration# defvlan 4000
```

6. Ensure tagging is disabled on vPort.

```
>>Virtual Port INTA1.1 Network Configuration# deftag d
>>Virtual Port INTA1.1 Network Configuration# ..
```

7. Specify QoS parameters for the vPort.

```
>>Virtual Port INTA1.1 Configuration# qos
>>Virtual Port INTA1.1 QOS Configuration# bw
>>Virtual Port INTA1.1 Bandwidth Configuration# minbw 25                (in percentage)
>>Virtual Port INTA1.1 Bandwidth Configuration# maxbw 100               (in percentage)
>>Virtual Port INTA1.1 Bandwidth Configuration# ..
>>Virtual Port INTA1.1 Network Configuration# ..
```

8. Configure tagging on external port 1.

```
>>Main# /cfg/port EXT1
>>Port EXT1# tagipvid enable
>>Port EXT1# tagpvid disable
>>Port EXT1# pvid 4000
>>Port EXT1# /
```

9. Apply the configuration.

```
>>Main# apply
```

Example 4: FCoE Mode

Following is an example configuration of UFP vPorts in FCoE mode.

1. Enable CEE.

```
>>Main# /cfg/cee/on
```

2. Enable FIPs.

```
>>Main# /cfg/fcoe/fips/on
```

3. Turn on UFP.

```
>>Main# /cfg/virt/ufp on
```

4. Configure internal port as UFP.

```
>>UFP Global Configuration# port INTA1
>>Port INTA1 UFP Configuration# ena
```

5. Configure virtual port.

```
>>Port INTA1 UFP Configuration# vport 2
>>Virtual Port INTA2.1 Configuration# ena
```

6. Configure vPort access mode.

```
>>Virtual Port INTA2.1 Configuration# network
>>Virtual Port INTA2.1 Network Configuration# mode fcoe
```

© Copyright IBM Corp. 2013

7. Configure vPort default VLAN.

```
>>Virtual Port INTA2.1 Network Configuration# defvlan 1102
```

8. Ensure tagging is disabled on vPort.

```
>>Virtual Port INTA2.1 Network Configuration# deftag d
>>Virtual Port INTA2.1 Network Configuration# ..
```

9. Specify QoS parameters for the vPort.

```
>>Virtual Port INTA2.1 Configuration# qos
>>Virtual Port INTA2.1 QOS Configuration# bw
>>Virtual Port INTA2.1 Bandwidth Configuration# minbw 25                (in percentage)
>>Virtual Port INTA2.1 Bandwidth Configuration# maxbw 100               (in percentage)
>>Virtual Port INTA2.1 Bandwidth Configuration# ..
>>Virtual Port INTA2.1 Network Configuration# ..
```

10. Enable tagging on external port.

```
>>Main# /cfg/port EXT1
>>Port EXT1# tag enable
>>Port EXT1# pvid 1102
>>Port EXT1# /
```

11. Apply the configuration.

```
>>Main# apply
```

VLAN - Ingress VLAN Tagging

Tagging can be enabled on an ingress port. When a packet is received on an ingress port, and if ingress tagging is enabled on the port, a VLAN tag with the port PVID is inserted into the packet as the outer VLAN tag. Depending on the egress port setting (tagged or untagged), the outer tag of the packet is retained or removed when it leaves the egress port.

Ingress VLAN tagging is used to tunnel packets through a public domain without altering the original 802.1Q status.

When ingress tagging is enabled on a port, all packets, whether untagged or tagged, will be tagged again. As shown in Figure 5, when tagging is enabled on the egress port, the outer tag of the packet is retained when it leaves the egress port. If tagging is disabled on the egress port, the outer tag of the packet is removed when it leaves the egress port.

Figure 5. 802.1Q tagging (after ingress tagging assignment)

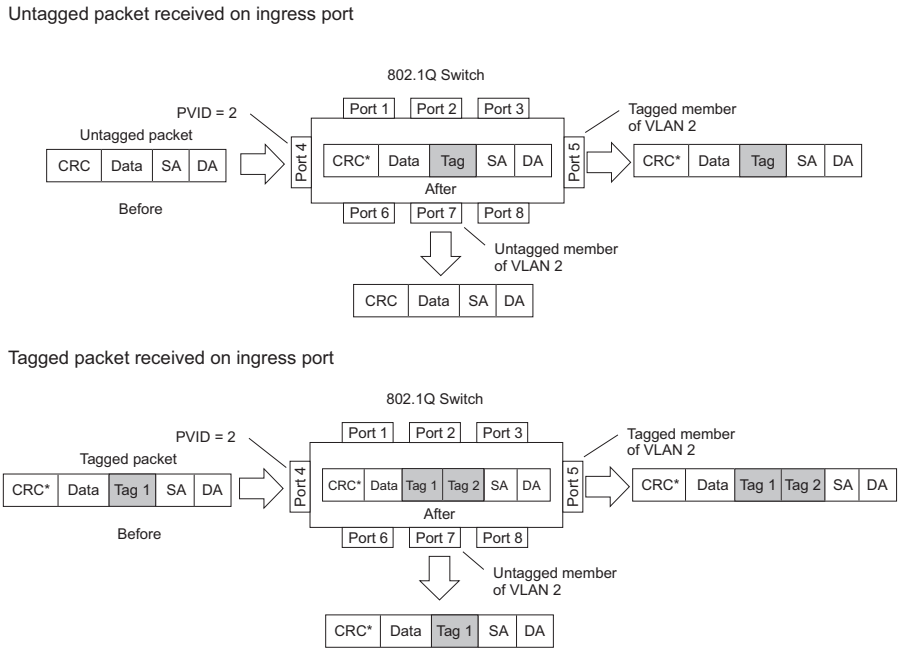

By default, ingress tagging is disabled. To enable ingress tagging on a port, use the following command:

```
>> Main# /cfg/port <number>/tagipvid enable
```

Limitations

Ingress tagging cannot be configured with the following features/configurations:

- VNIC ports
- VMready ports
- UFP ports
- Management ports

VLAG Health Check

Added support for configuring VLAG health check ports with an IPv4 or IPv6 address.

VMready

Up to 2048 VM profiles, 4093 VM groups, 4096 VMs, and 4096 VEs can be configured on the CN4093 10Gb Converged Scalable Switch. Of the total VMs, 2048 can be used in local groups.

© Copyright IBM Corp. 2013

vNIC Groups - Shared Mode

The CN4093 10Gb Converged Scalable Switches support two modes for configuring the vNIC uplinks: dedicated mode and shared mode. The default is the dedicated mode. To enable the shared mode, enter the following command:

```
>> Main# cfg/virt/vnic/ulshare ena
```

In the dedicated mode, only one vNIC group is assigned to an uplink port. This port can be a regular port or a trunk port. The NIC places an outer tag on the vNIC group packets. This outer tag contains the vNIC group VLAN. The uplink NIC strips off the outer tag before sending out the packet.

In the shared mode, multiple vNIC groups can be assigned to an uplink port. This port can be a regular port or a trunk port. The vNIC groups share the uplink. You may assign a few vNIC groups to share an uplink and the other vNIC groups to have a single uplink each. In either case, the switch still operates in shared mode. As in the dedicated mode, the NIC places an outer tag on the vNIC group packets. This outer tag contains the vNIC group VLAN. The uplink NIC does not strip off the outer tag. The vNIC group tag defines the regular VLAN for the packet.This behavior is particularly useful in cases where the downstream server does not set any tag. Effectively, each vNIC group is a VLAN, which you can assign by configuring the VLAN to the vNIC group. You must enable the tag configuration on the uplink port.

The table below compares the configurations of the two modes.

Table 2. Comparison: Dedicated Mode vs. Shared Mode

Configuration Area	Dedicated Mode	Shared Mode
Port	“tagpvid” must be disabled.	“tagpvid” is user configurable.
	“pvid” = vNIC group VLAN.	“pvid” is user configurable.
	“tag” is user configurable.	“tag” must be enabled.
	Port can be added only to the vNIC group VLAN.	Port can be added to multiple VLANs in addition to the vNIC group VLANs that are automatically configured.
	Inserts vNIC group VLAN in the outer tag of ingress packets.	Inserts regular VLAN in the outer tag. VLAN tags are passed to and received from the uplink switch similar to vNIC ports.
		To handle untagged packets, configure the pvid/native VLAN of the uplink port to one of the vNIC group VLANs, and disable “tagpvid”.
VLAN	Add the port to a vNIC group VLAN and delete it from any other VLAN when the vNIC group VLAN is enabled.	Add the port to all vNIC group VLANs that are sharing the port. Do not remove it from any other VLAN.
	Delete the port from the vNIC group VLAN and add it back to the default VLAN 1 when the vNIC group is disabled/deleted or when the vNIC feature is globally disabled.	Remove the port from a vNIC group VLAN when the vNIC group is disabled/deleted. When the vNIC feature is globally disabled or the port is not added in any vNIC group, remove the port from all vNIC group VLANs and add it back to default VLAN 1 if no non-vNIC VLAN exists on the port.
	Do not add a port or trunk to multiple vNIC groups that are enabled.	Can add a port or trunk to multiple vNIC groups that are enabled.
	Do not configure additional VLANs on the uplink ports.	Can configure additional VLANs on the uplink ports.
STP	An uplink port can only be in one STG.	An uplink port can be in multiple STGs.
	When you add a port to a vNIC group, STP is automatically disabled.	When you add a port to a vNIC group, STP is automatically disabled.
	When you remove a port from a vNIC group, STP is automatically reset to factory default.	When you remove a port from a vNIC group, STP is automatically reset to factory default.
Failover	An uplink up/event can trigger the failover state change only of one vNIC group.	An uplink up/event can trigger the failover state change of multiple vNIC groups.

The vNIC group VLAN ID is placed on all vNIC group packets as an “outer” tag. As shown in Figure 6, the outer vNIC group VLAN ID is placed on the packet in addition to any regular VLAN tag assigned by the network, server, or hypervisor.

© Copyright IBM Corp. 2013

Figure 6. Outer and Inner VLAN Tags

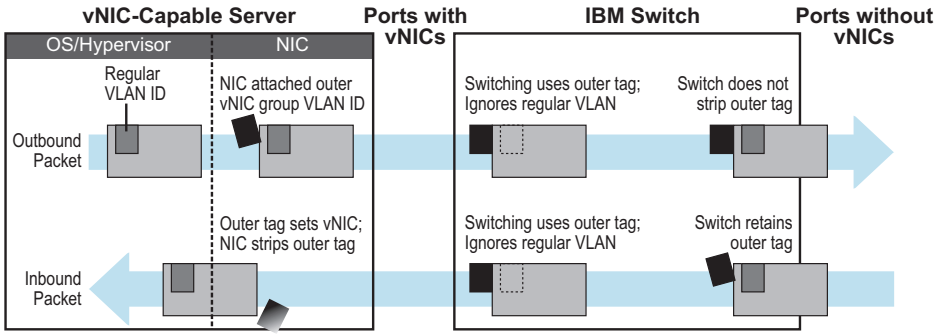

Within the CN4093, all Layer 2 switching for packets within a vNIC group is based on the outer vNIC group VLAN. The CN4093 does not consider the regular, inner VLAN ID (if any) for any VLAN-specific operation.

The outer vNIC group VLAN is not removed by the switch before the packet egresses any internal port or external uplink port. For untagged packets sent by the server, the uplink NIC uses this outer tag to switch the packet to destined VLAN.

The shared mode is useful is cases where the multiple VNIC groups need to share an uplink port. The VNIC group tag defines the user VLAN. Following is an use case:

An ESX server is presented with eight VNICs (four from bay 7 and four from bay 9) used with four virtual switches of the ESX host and with no tagged port groups. A pair of odd/even VNICs is placed within each virtual switch. On the CN4093, four VNIC groups are created and the desired VLAN for each VNIC group is configured. For example, if VNIC group 1 on the CN4093 has four interfaces: 1.1, 2.1, 3.1, 4.1. VNIC group 1 is configured with VLAN 10. Packets coming from any VM connecting with the virtual switch that VMNIC 2 and 3 (VNIC 1.1, 2.1, 3.1, and 4.1 on bay 7 and bay 9) will be assigned with VLAN 10. These packets go out the uplink with VLAN 10 tag. The upstream switch sends these packets to the desired destination on VLAN 10.

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.

© Copyright IBM Corp. 2013

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.
```

   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.
   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.
```

   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
```

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_OS.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: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:

- The ISL should include enough ports to accommodate the peer-to-peer traffic.
- 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*.

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 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 *IBM Networking OS 7.7 Application Guide*.

Note: The external uplink ports (EXTx) cannot be members of management VLANs.

© Copyright IBM Corp. 2013

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:

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

VRRP Configuration

Although the Virtual Router Redundancy Protocol (VRRP) standard permits up to 255 virtual router instances, the N/OS 7.7 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.

© Copyright IBM Corp. 2013

Known Issues

This section describes known issues for N/OS 7.7 on the CN4093 10Gb 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)

Boot Configuration Block

- In the CLI, the boot configuration command (`/boot/conf` <*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 CMM CLI. (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 7.7, 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)

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 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)
- 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.
- 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.
- 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.

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)

© Copyright IBM Corp. 2013

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 a CN4093 is connected to another CN4093 with LACP configured (but no active LACP trunk), the `/info/l2/trunk` command might erroneously report the static trunk as forwarding.
- If you configure LACP (active/passive) on one port, also configure LACP on the partner switch, at the end of the link. If you connect LACP with a static trunk, there will be no connectivity on that link.
- 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)

Private VLANs

- Isolation for secondary VLANs is not honored across stacking interlinks. Traffic between the ports of a secondary VLAN is not isolated when those ports belong to different switches within a stack. Traffic in the secondary VLAN will be properly isolated only for traffic between ports of the same switch. (ID: 68340)
- The sequence in which a private VLAN is configured is not the same as displayed in the output of the `CN 4093(config)# show running-config` command. Hence, if you copy and paste the private VLAN configuration from the output of the above command, the private VLAN configuration will be lost. (ID: 67169)
- Traffic with secondary VLAN ID is not forwarded to promiscuous ports. (ID: 70980)

© Copyright IBM Corp. 2013

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/qsfp40g/add` *<3,7>*

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)

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)

Virtual Link Aggregation Groups

- 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)