IBM Flex System Fabric EN4093 and EN4093R 10 Gb 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)

P/N: 00AY517

Release Notes

This release supplement provides the latest information regarding IBM Networking OS 7.7 for EN4093 and EN4093R 10Gb Virtual Fabric Scalable Switch (collectively referred to as EN4093 throughout this document). Where required, differences between the EN4093 and EN4093R are detailed.

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

- *IBM Networking OS Application Guide for the EN4093 and EN4093R 10Gb Virtual Fabric Scalable Switch*
- *IBM Networking OS Command Reference for the EN4093 and EN4093R 10Gb Virtual Fabric Scalable Switch*
- *IBM Networking OS ISCLI Reference for the EN4093 and EN4093R 10Gb Virtual Fabric Scalable Switch*
- *IBM Networking OS BBI Quick Guide for the EN4093 and EN4093R 10Gb Virtual Fabric Scalable Switch*
- *EN4093 and EN4093R 10Gb Virtual Fabric Scalable Switch User's Guide*

The publications listed above are available at the following address:

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

Please keep these release notes with your product manuals.

Hardware Support

N/OS 7.7 software is supported on the EN4093, shown in Figure 1, is a high performance Layer 2-3 embedded network switch that features tight integration with IBM Flex System chassis management module.

Figure 1. EN4093 10Gb Scalable Switch Faceplate

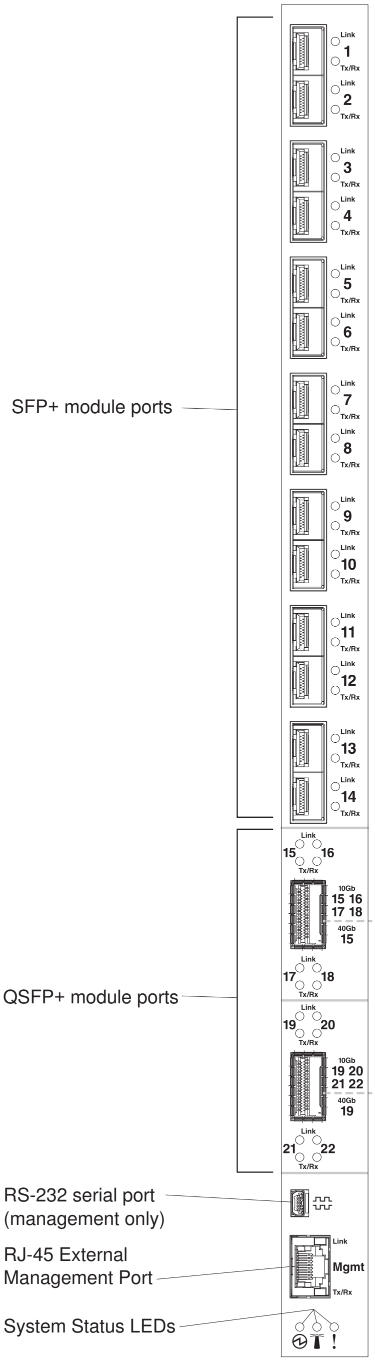

The EN4093 has the following port capacities:

- Fourteen 10Gb SFP+ ports
- Two 40Gb QSFP+ ports
- Fourty-Two 10Gb internal ports (maximum)
- One 1Gb RJ-45 external management port
- One 1Gb internal management port
- One mini-USB serial port

Transceivers

The following transceivers and DACs are available:

Table 1. EN4093 Transceivers and DACs

Description	Part number
Transceivers	
1000Base-SX SFP (MMFiber) transceiver	81Y1622
1000Base-T SFP transceiver 4	81Y1618
1000Base-LX SFP LX transceiver	90Y9424
10GBase-SR SFP+ (MMFiber) transceiver	44W4408
10GBase-SR SFP+ (MMFiber) transceiver	46C3447
IBM BNT SFP+ LR transceiver	90Y9412
IBM BNT QSFP+ 40Gbase-SR transceiver	49Y7884
Direct Attach Cables (DACs)	
1m IBM Passive DAC SFP+	90Y9427
3m IBM Passive DAC SFP+	90Y9430
5m IBM Passive DAC SFP+	90Y9433
1m 40Gb QSFP+ to 4 x 10Gb SFP+ Breakout Cable	49Y7886
3m 40Gb QSFP+ to 4 x 10Gb SFP+ Breakout Cable	49Y7887
5m 40Gb QSFP+ to 4 x 10Gb SFP+ Breakout Cable	49Y7888
10m IBM MTP Fiber Optical Cable (Requires transceiver 49Y7884)	90Y3519
30m IBM MTP Fiber Optical Cable (Requires transceiver 49Y7884)	90Y3521
1m QSFP+ to QSFP+ DAC	49Y7890
3m QSFP+ to QSFP+ DAC	49Y7891

The EN4093 accepts any SFP+ Direct Attach Cable that complies to the MSA specification.

Updating the Switch Software Image

The switch software image is the executable code running on the EN4093. 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 EN4093, 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 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.

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 7.3 or Prior

If VLAG is enabled:

- Both the VLAG peer switches must be upgraded at the same time with the same N/OS version.

Updating from IBM Networking OS 7.5 or Prior

- 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 any of the previous switch software versions will be lost after an upgrade to N/OS 7.7.

Loading New Software to Your Switch

The EN4093 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 20).

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

- The image and boot software loaded on an 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 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 FTP or TFTP server.

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

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

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

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

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

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

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

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

   ```
   Boot Options# image
   ```

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

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

   Specify the image that contains the newly loaded software.

7. Reboot the switch to run the new software:

   ```
   Boot Options# reset
   ```

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

Loading Software via the ISCLI

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

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

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

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

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

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

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

4. If required by the FTP or TFTP server, enter the appropriate username and password.

5. The switch will prompt you to confirm your request.

   Once confirmed, the software will begin loading into the switch.

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

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

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

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

7. Reboot the switch to run the new software:

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

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

Loading Software via BBI

You can use the Browser-Based Interface to load software onto the EN4093. The software image to load can reside in one of the following locations:

- FTP server
- TFTP server
- Local computer

After you log onto the BBI, perform the following steps to load a software image:

1. Click the Configure context tab in the toolbar.
2. In the Navigation Window, select System > Config/Image Control.

   The Switch Image and Configuration Management page appears.
3. If you are loading software from your computer (HTTP client), skip this step and go to the next. Otherwise, if you are loading software from a FTP/TFTP server, enter the server's information in the FTP/TFTP Settings section.
4. In the Image Settings section, select the image version you want to replace (Image for Transfer).
   – If you are loading software from a FTP/TFTP server, enter the file name and click **Get Image**.

   – If you are loading software from your computer, click **Browse**.

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

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

New and Updated Features

N/OS 7.7 for EN4093 and EN4093R 10Gb Virtual Fabric Scalable Switch (EN4093) has been updated to include several new features, summarized in the following sections. For more detailed information about configuring EN4093 features and capabilities, refer to the complete N/OS 7.7 documentation as listed on page 3.

DHCP

Host Name Configuration

The EN4093 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:

```
EN 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:

```
EN 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 EN4093 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 `EN 4093# show logging` command to view the SYSLOG server address.

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

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

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>
```

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 EN4093. 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 EN4093. 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 EN4093. 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 EN4093. 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 EN4093. 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.

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
```

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.

Hot Links

Hot links provides basic link redundancy with fast recovery. Prior to IBM Networking OS 7.7, STP had to be globally disabled for configuring hot links. This restriction is no longer applicable. STP can be globally enabled but must be disabled on the ports used for hot links configuration.

IPv4 Address Conflict Detection

The EN4093 and EN4093R 10Gb Virtual Fabric 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 EN4093 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 EN4093 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

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 EN4093 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 EN4093) and a Partner (another switch), as shown in Table 2.

Table 2. 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 2, 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 2, 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 <65-128>/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:

```
EN 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:

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

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

OpenFlow

OpenFlow architecture consists of a control plane residing outside of the switch (typically on a server) and a data plane residing in the switch. The control plane is called OpenFlow controller. The data plane which resides in the switch consists of a set of flows which determine the forwarding of data packets.

The OpenFlow protocol is described in the OpenFlow Switch Specification 1.0.0

An OpenFlow network consists of simple flow-based switches in the data path, with a remote controller to manage all switches in the OpenFlow network.

OpenFlow maintains a TCP channel for communication of flow management between the controller and the switch. All controller-switch communication takes place over the switch's management network.

Switch Profiles

The EN4093 10Gb Scalable Switch can be used for configuring OpenFlow and legacy switching features simultaneously. However, Layer 2 and Layer 3 switching features can be configured only on the ports that are not OpenFlow ports. Legacy switching ports and OpenFlow ports do not communicate with each other.

Alternately, the switch can be configured as an OpenFlow-only switch if you do not need to configure legacy switching features.

Based on your requirement, select the switch boot profile using the following commands:

- OpenFlow-only: `EN 4093(config)# boot profile openflow`

  The switch will operate only in OpenFlow environment. None of the legacy switching features will be supported.
- OpenFlow and Legacy Switching:
  ```
  EN 4093(config)# boot profile default
  ```
  Legacy switching features can be configured on the non-OpenFlow ports. By default, the switch boots in this profile.

Reload the switch to apply boot profile changes.

For details, see *IBM Flex System Fabric EN4093 and EN4093R 10Gb Virtual Fabric Scalable Switch Application Guide*.

OSPFv3

Enhancements based on RFC5340.

Secure FTP

Added support for configuring SFTP server port other than the standard port 22.

SNMP 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>
```

VLAG

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 EN4093 10Gb Scalable Switch. Of the total VMs, 2048 can be used in local groups.

vNIC Groups - Shared Mode

The EN4093 and EN4093R 10Gb Virtual Fabric 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 3. 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 “tag-pvid”.
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 2, 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.

Figure 2. Outer and Inner VLAN Tags

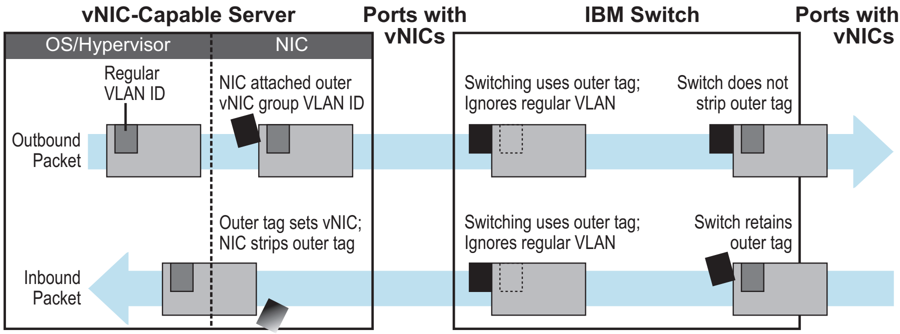

vNIC-Capable Server
Ports with vNICs
IBM Switch
Ports without vNICs
OS/Hypervisor
NIC
Regular VLAN ID
NIC attached outer vNIC group VLAN ID
Switching uses outer tag; Ignores regular VLAN
Switch does not strip outer tag
Outbound Packet
Outer tag sets vNIC; NIC strips outer tag
Switching uses outer tag; Ignores regular VLAN
Switch retains outer tag
Inbound Packet

Within the EN4093, all Layer 2 switching for packets within a vNIC group is based on the outer vNIC group VLAN. The EN4093 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 EN4093, four VNIC groups are created and the desired VLAN for each VNIC group is configured. For example, if VNIC group 1 on the EN4093 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 EN4093 and N/OS.

The Boot Management Menu

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

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

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

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

Please choose your menu option: 3
```

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

- To change the booting image, press 1 and follow the screen prompts.
- To change the configuration block, press 2, and follow the screen prompts.
- To perform a software image recovery, press 3 and follow the screen prompts.
- To perform an Xmodem download (boot image only), press 4 and follow the screen prompts.
- To exit the Boot Management menu, press 6. The booting process continues.

Recovering from a Failed Software Upgrade

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

1. Connect a PC to the serial port of the switch.
2. Open a terminal emulator program that supports Xmodem download (for example, HyperTerminal, CRT, PuTTY) and select the following serial port characteristics:

– Speed:	9600 bps
– Data Bits:	8
– Stop Bits:	1
– Parity:	None
– Flow Control:	None

3. Boot the switch and access the Boot Management menu by pressing <**Shift B**> while the Memory Test is in progress and the dots are being displayed.

4. Select **3** for **Boot in recovery mode**. You will see the following display:

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

   – If you choose option `x` (Xmodem serial download), go to step 5.
   – If you choose option `t` (TFTP download), go to step 6.

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

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

   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

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 EN4093. All mirrored egress traffic is tagged.

Secure Management Network

The following EN4093 attributes are reserved to provide secure management access to and from the chassis management module:

- MGT port (MGT1)
- VLAN 4095
- IP interface 127, 128
- Gateway 4
- STG 128

For more information about remotely managing the EN4093 through the external ports, see “Accessing the Switch” in the *IBM Networking OS 7.7 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 EN4093 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 EN4093, 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 EN4093s 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.

Known Issues

This section describes known issues for N/OS 7.7 on the EN4093 and EN4093R 10Gb Virtual Fabric 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)

- False switch errors might be reported in the CMM event log after the Flex Chassis power is cycled off/on. These events imply that the switch has experienced an electrical or thermal issue, as shown below:

```
Info 08:42 AM           I/O module iomodule02 is within the
recommended
                        temperature range.
Info 08:42 AM Recovery: I/O module iomodule02 current fault.
Err  08:42 AM           I/O module iomodule02 current fault.
Err  08:42 AM           I/O module iomodule02 temperature fault.
Warn 08:42 AM           I/O module iomodule02 is over recommended
                        temperature.
```

  These false switch errors occur and recover within the same minute and only intermittently occur as a result of a chassis power cycle. Ignore these false error messages. (ID: 60687)

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

DHCP

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

FCoE

- In N/OS 7.7, the EN4093 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)
- 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)

HTTPS

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

IPsec

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

ISCLI

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

ISCLI Configuration Scripts

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

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

LACP

- If a static trunk on a EN4093 is connected to another EN4093 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)

Menu-Based CLI

- When dynamic VLAN configuration is enabled for features such as QBG, VMready, or FCoE, whenever automatic changes are made to the switch VLAN configuration, any other unapplied changes pending from menu-based CLI configuration session will be automatically applied. (ID: 65392)

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 EN4093 is connected to an application switch which requires a link speed of 100 Mbps half-duplex, then enable auto negotiation on the EN4093 port with port speed=any, mode=any, fctl=both, and auto=on.

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)
- When upgrading the switch software image to version 7.7, configurations associated with a private VLAN that was disabled using the command: `EN 4093(config)# no private-vlan enable` will be lost. Any related VLAN, IGMP snooping, or IGMP relay configuration will also be lost. (ID: 70912)
- The sequence in which a private VLAN is configured is not the same as displayed in the output of the `EN 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)

QSFP+

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

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

SLP

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

SNMP

During SNMP MIB walks, if you experience timeouts, set the timeout value to 3 seconds or higher in the SNMP application/tool. (IDs: 71913, 71914, 71906)

Stacking

- LLDP status changes are not reflected in the master switch system logs. (ID: 65325)

UFP

- In N/OS 7.7, the EN4093 does not support UFP QoS statistics. (ID: 64021, 64371)