IBM Networking OS 7.5

Release Notes

for the EN2092 1Gb Ethernet Scalable Switch,
Second Edition (replaces 88Y7950)

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

Second Edition (November 2012)

P/N: 00D2340

Release Notes

This release supplement provides the latest information regarding IBM Networking OS 7.5 for the EN2092 1Gb Ethernet Scalable Switch (referred to as EN2092 throughout this document).

This supplement modifies and extends the following IBM Networking OS documentation for use with Networking OS 7.5:

- *IBM Networking OS Application Guide for the EN2092 1Gb Ethernet Scalable Switch*
- *IBM Networking OS Command Reference for the EN2092 1Gb Ethernet Scalable Switch*
- *IBM Networking OS ISCLI Reference for the EN2092 1Gb Ethernet Scalable Switch*
- *IBM Networking OS BBI Quick Guide for the EN2092 1Gb Ethernet Scalable Switch*
- *EN2092 1Gb Ethernet 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

Networking OS 7.5 software is supported only on the EN2092 1Gb Ethernet Scalable Switch for the IBM Flex System. The EN2092 1Gb Ethernet Scalable Switch (EN2092), 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. EN2092 1Gb Ethernet Scalable Switch Faceplate

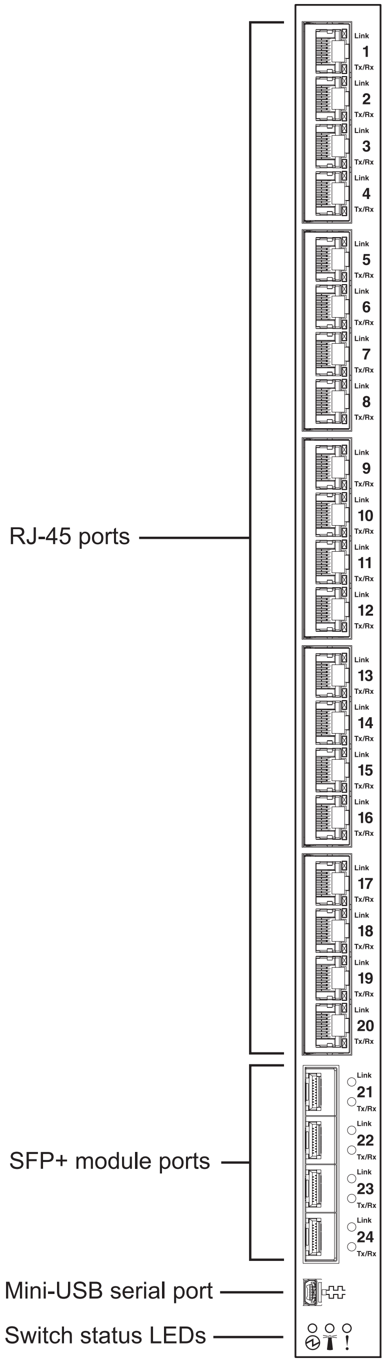

The EN2092 has the following port capacities:

- Twenty 1Gb RJ-45 ports
- Four 10Gb SFP+ ports
- Twenty-Eight 1Gb internal ports (maximum)
- One 1Gb internal management port
- One mini-USB serial port

Transceivers

The following transceivers and DACs are available:

Table 1. EN2092 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
Direct Attach Cables (DACs)	
1m IBM Passive DAC SFP+	90Y9427
3m IBM Passive DAC SFP+	90Y9430
5m IBM Passive DAC SFP+	90Y9433

Updating the Switch Software Image

The switch software image is the executable code running on the EN2092. 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 EN2092, 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.

  Note: Software loading options also include SFTP. Some screen prompts may appear slightly different than depicted in this document.
- 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 6.

ATTENTION: Although the typical upgrade process is all that is necessary in most cases, upgrading from (or reverting to) some versions of Networking 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.

Loading New Software to Your Switch

The EN2092 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 menu-based CLI, the ISCLI, or the BBI to download and activate new software.

Loading Software via the Menu-Based 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 EN2092. 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

Networking OS 7.5 for EN2092 1Gb Ethernet Scalable Switch (EN2092) has been updated to include several new features, summarized in the following sections. For more detailed information about configuring EN2092 features and capabilities, refer to the complete Networking OS 7.5 documentation as listed on page 3.

Diagnostics Enhancement

The following commands have been added to improve the ability to diagnose system issues:

- `Router(config)#` **`show logging`** [**`messages`**] [**`severity`** *<0-7>*] [**`reverse`**] **|** [**`head`**|**`last`**] *<line number>*
- `Router(config)#` **`show environment power`**
- `Router(config)#` **`show environment fan`**
- `Router#` **`show version`** [**`brief`**]
- `Router#` **`show tech-support`** [**`l2|l3|link|port`**]
- `Router#` **`system idle`** *<0-60>*
- `Router(config)#` **`logging synchronous`** [**`level`** *<severity-level>***`|all`**]
- `Router#` **`show who`**
- `Router(config)#` **`access user clear`** *<session ID>*
- `Router#` **`show line`**
- `Router#` **`clear line`** *<session ID>*
- `Router# [no]` **`debug lacp packet`**
- `Router# [no]` **`debug spanning-tree bpdu`** [**`receive`**|**`transmit`**]

For detailed description of these commands, see the *IBM Networking OS 7.5 Command Reference* and *IBM Networking OS 7.5 ISCLI Reference Guides*.

Note: Use the debug commands with caution as they can disrupt the operation of the switch under high load conditions.

LLDP

LLDP transmissions can be configured to enable or disable inclusion of the following optional information:

Table 2. LLDP Optional Information Types

Type	Description	Default
`portdesc`	Port Description	Enabled
`sysname`	System Name	Enabled
`sysdescr`	System Description	Enabled
`syscap`	System Capabilities	Enabled
`mgmtaddr`	Management Address	Enabled
`portvid`	IEEE 802.1 Port VLAN ID	Disabled
`portprot`	IEEE 802.1 Port and Protocol VLAN ID	Disabled
`vlanname`	IEEE 802.1 VLAN Name	Disabled
`protid`	IEEE 802.1 Protocol Identity	Disabled
`macphy`	IEEE 802.3 MAC/PHY Configuration/Status, including the auto-negotiation, duplex, and speed status of the port.	Disabled
`powermdi`	IEEE 802.3 Power via MDI, indicating the capabilities and status of devices that require or provide power over twisted-pair copper links.	Disabled
`linkaggr`	IEEE 802.3 Link Aggregation status for the port.	Disabled
`framesz`	IEEE 802.3 Maximum Frame Size for the port.	Disabled
`dcbx`	Data Center Bridging Capability Exchange Protocol (DCBX) for the port.	Enabled
`dcbx`	Data Center Bridging Capability Exchange Protocol (DCBX) for the port.	Enabled

OSPFv3 Over IPsec

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

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.

Running Configuration

The following ISCLI command has been added to compare the running configuration with the startup configuration stored in FLASH.

```
Router# show running-config diff
```

Secure FTP

IBM Networking OS supports Secure FTP (SFTP) to the switch. SFTP uses Secure Shell (SSH) to transfer files. SFTP encrypts both commands and data, and prevents passwords and sensitive information from being transmitted openly over the network.

All file transfer commands include SFTP support along with FTP and TFTP support. SFTP is available through the menu-based CLI, ISCLI, BBI, and SNMP.

The following examples illustrate SFTP support for menu-based CLI commands:

Download software image file:

```
# /boot/gtimg
Enter name of switch software image to be replaced
["image1"|"image2"|"boot"]:  image2
Enter hostname or IP address of SFTP/FTP/TFTP server: 10.10.10.1
Enter name of file on SFTP/FTP/TFTP server: filename
Enter username for SFTP/FTP server or hit return for TFTP server: name
Enter password for username on SFTP/FTP server:
Enter the port to use for downloading the image
["data"|"extm"|"mgt"]:  data
Enter method of transporting [sftp | ftp]: sftp

image2 currently contains Software Version 7.3.0
 that was downloaded at 18:30:39 Thu Jun 7, 2012.
New download will replace image2 with file "filename"
 from SFTP/FTP/TFTP server 10.10.10.1.
Connecting via DATA port.
Confirm download operation [y/n]:
```

Download HTTPS certificate:

```
# /cfg/sys/access/https/gtca
Enter hostname or IP address of SFTP/TFTP server: 10.10.10.1
Enter name of file on SFTP/TFTP server: filename
Enter username for SFTP server or hit return for TFTP server: name
Enter password for username on SFTP server:
Enter the port to use for downloading the file
["data"|"extm"|"mgt"]:  data
Confirm download operation [y/n]:
```

The following examples illustrate SFTP support for ISCLI commands:

```
Router# copy sftp {image1|image2|boot-image} [mgt-port|data-port]
         (Copy software image from SFTP server to the switch)

Router# copy sftp {ca-cert|host-cert|host-key} [mgt-port|data-port]
         (Copy HTTPS certificate or host key from SFTP server to the switch)
```

Service Location Protocol

Service Location Protocol (SLP) allows the switch to provide dynamic directory services that helps users find servers by attributes rather than by name or address. SLP eliminates the need for a user to know the name of a network host supporting a service. SLP allows the user to bind a service description to the network address of the service.

Service Location Protocol is described in RFC 2608.

Note: SLP is not supported on the internal management port (MGT).

SLP defines specialized components called agents that perform tasks and support services as follows:

- User Agent (UA) supports service query functions. It requests service information for user applications. The User Agent retrieves service information from the Service Agent or Directory Agents. A Host On-Demand client is an example of a User Agent.
- Service Agent (SA) provides service registration and service advertisement. **Note**: In this release, SA supports UA/DA on Linux with SLPv2 support.
- Directory Agent (DA) collects service information from Service Agents to provide a repository of service information in order to centralize it for efficient access by User Agents. There can only be one Directory Agent present per given host.

The Directory Agent acts as an intermediate tier in the SLP architecture, placed between the User Agents and the Service Agents, so they communicate only with the Directory Agent instead of with each other. This eliminates a large portion of the multicast request or reply traffic on the network, and it protects the Service Agents from being overwhelmed by too many service requests.

Services are described by the configuration of attributes associated with a type of service. A User Agent can select an appropriate service by specifying the attributes that it needs in a service request message. When service replies are returned, they contain a Uniform Resource Locator (URL) pointing to the service desired, and other information, such as server load, needed by the User Agent.

Active DA Discovery

When a Service Agent or User Agent initializes, it can perform Active Directory Agent Discovery using a multicast service request and specifies the special, reserved service type (`service:directory-agent`). Active DA Discovery is achieved through the same mechanism as any other discovery using SLP.

The Directory Agent replies with unicast service replies, which provides the URLs and attributes of the requested service.

SLP Configuration

Use the following ISCLI commands to configure SLP on the switch:

Table 3. SLP ISCLI Commands

Command Syntax and Usage
`[no] ip slp enable` Enables or disables SLP on the switch. **Command mode:** Global configuration
`[no] ip slp active-da-discovery enable` Enables or disables Active DA Discovery. **Command mode:** Global configuration
`ip slp active-da-discovery start-wait-time` *<1-10>* Configures the wait time before starting Active DA Discovery, in seconds. The default value is 3 seconds. **Command mode:** Global configuration
`clear ip slp directory-agents` Clears all Directory Agents learned by the switch. **Command mode:** Global configuration
`show ip slp information` Displays SLP information. **Command mode:** All
`show ip slp directory-agents` Displays Directory Agents learned by the switch. **Command mode:** All
`show ip slp user-agents` Displays User Agents information. **Command mode:** All
`show ip slp counters` Displays SLP statistics. **Command mode:** All
`clear ip slp counters` Clears all Directory Agents learned by the switch. **Command mode:** Global configuration

SNMP MIBs

- Added MIBs required for accessing LLDP data, as specified in RFC2737 and IEEE 802.AB.
- Added MIBs required for accessing MLDv2 information.
- Added entity MIBs, as specified in RFC2737 and RFC4133.
- Added MIBs required for managing host resources, as specified in RFC1514 and RFC2790.

VMcheck

The EN2092 primarily identifies virtual machines by their MAC addresses. An untrusted server or a VM could identify itself by a trusted MAC address leading to MAC spoofing attacks. Sometimes, MAC addresses get transferred to another VM, or they get duplicated.

The VMcheck solution addresses these security concerns by validating the MAC addresses assigned to VMs. The switch periodically sends hello messages on server ports. These messages include the switch identifier and port number. The hypervisor listens to these messages on physical NICs and stores the information, which can be retrieved using the VMware Infrastructure Application Programming Interface (VI API). This information is used to validate VM MAC addresses. Two modes of validation are available: Basic and Advanced.

Use the following command to select the validation mode or to disable validation:

```
>>Main# /cfg/virt/vmgroup <VM group number>/validate {basic|advanced|disable}
```

Basic Validation

This mode provides port-based validation by identifying the port used by a hypervisor. It is suitable for environments in which MAC reassignment or duplication cannot occur.

The switch, using the hello message information, identifies a hypervisor port. If the hypervisor port is found in the hello message information, it is deemed to be a trusted port. Basic validation should be enabled when:

- A VM is added to a VM group, and the MAC address of the VM interface is in the Layer 2 table of the switch.
- A pre-provisioned VM interface that belongs to a VM group connects to the switch.
- A trusted port goes down. Port validation must be performed to ensure that the port does not get connected to an untrusted source when it comes back up.

Use the following command to set the action to be performed if the switch is unable to validate the VM MAC address:

```
>>Main# /cfg/virt/vmcheck/action/basic {log|link}

log - generates a log
link - disables the port
```

Advanced Validation

This mode provides VM-based validation by mapping a switch port to a VM MAC address. It is suitable for environments in which spoofing, MAC reassignment, or MAC duplication is possible.

When the switch receives frames from a VM, it first validates the VM interface based on the VM MAC address, VM Universally Unique Identifier (UUID), Switch port, and Switch ID available in the hello message information. Only if all the four parameters are matched, the VM MAC address is considered valid.

In advanced validation mode, if the VM MAC address validation fails, an ACL can be created to drop the traffic received from the VM MAC address on the switch port. Use the following command to specify the number of ACLs to be used for dropping traffic:

```
>>Main# /cfg/virt/vmcheck/acls <1-256>
```

Use the following command to set the action to be performed if the switch is unable to validate the VM MAC address:

```
>>Main# /cfg/virt/vmcheck/action/advanced {log|link|acl}
```

Following are the other VMcheck commands:

Table 4. VMcheck Commands

Command	Description
`>>Main# /cfg/virt/vmware/hello {ena|dis|addport <port number>| rmport <port number>|haddr|htimer|cur}`	Hello messages setting: enable/disable/add port/remove port/advertise this IP address in the hello messages instead of the default management IP address/set the timer to send the hello messages/view current hello message settings
`>>Main# /cfg/virt/vmcheck/{trust|notrust} <port number>`	Mark a port as trusted/untrusted
`>>Main# /cfg/virt/vmcheck/cur`	View current VMcheck settings
`>>Main# /oper/virt/vmcheck/acl/{remall|remmac [<port number>]|remport <port number>}`	Delete ACL(s): all ACLs/an ACL by MAC address ((optional) and port number) /all ACLs installed on a port

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
```

- 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. IBM 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.

Supplemental Information

This section provides additional information about configuring and operating the EN2092 and Networking 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

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 (IF1-126). This results in port MGT1 not accepting the configuration and an IP configuration of all zeros displayed on the CMM user interface.

For example, consider that an external interface (IF1) is configured or enabled to the following IP address and mask:

```
Interface information:
1: 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
```

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 EN2092. All mirrored egress traffic is tagged.

Secure Management Network

The following EN2092 attributes are reserved to provide secure management access to and from the chassis management module:

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

For more information about remotely managing the EN2092 through the external ports, see “Accessing the Switch” in the *IBM Networking OS 7.5 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 EN2092 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 EN2092, 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 EN2092s 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).

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 Networking OS 7.5 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 Networking OS 7.5 on the EN2092 1Gb Ethernet Scalable Switch.

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)

- NTP configuration can only be saved via the CMM web interface. NTP configuration using the switch-based configuration interfaces (such as the menu-based CLI, ISCLI, BBI, etc.) will be overridden by the CMM whenever the switch or CMM are restarted. (ID: 60460)

Hotlinks

- Prior to enabling hotlinks, Spanning-Tree should be globally disabled using the following command (ID: 47917):

```
Router(config)# spanning-tree mode dis
```

IPsec

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

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.

Jumbo Frames

- Some ingress jumbo frames (for example, ICMP) are not routed from one VLAN to another VLAN. Jumbo frames are routed across data VLANs.

LACP

- If a static trunk on a EN2092 is connected to another EN2092 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)

Menu-Based Command Line Interface

- 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 switch 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 EN2092 is connected to an application switch which requires a link speed of 100 Mbps half-duplex, then enable auto negotiation on the EN2092 port with port speed=any, mode=any, fctl=both, and auto=on.

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)