BROADCOM®

Emulex® Boot for the Fibre Channel and NVMe over FC Protocols

User Guide
Release 14.2

Table of Contents

Chapter 1: Introduction

This guide describes installing, updating, enabling, and configuring the Emulex® FC boot code. This guide also describes the Emulex boot from SAN (BFS) implementation and its operation with distinct hardware and operating system requirements. BFS is the booting of a server directly from an operating system image located on a SAN by way of an Emulex adapter using the Emulex FC and FC-NVMe boot code.

When booting from SAN, the storage device is typically identified by its WWPN and a LUN for FC and by an NQN and namespace for NVMe over FC. By extending the server system BIOS, BFS functionality is provided by the boot code contained on an Emulex adapter in the server. When properly configured, the adapter then permanently directs the server to boot from a logical unit (disk) on the SAN as if it were a local disk.

NOTE: Emulex drivers support multipath boot configurations. Refer to your storage vendor's documentation for details on how to configure multipath booting.

NOTE: Screen captures in this document are for illustrative purposes only. Your system information might vary.

1.1 Emulex Boot Code Files

Emulex provides the following types of boot code:
- The x86 Boot code works with the existing system BIOS on x86 systems.
- The UEFI Boot code provides system boot capability through the use of the UEFI shell. It also functions on UEFI 2.*x*-based x64 platforms and Arm aarch64 platforms through the HII interface.

Emulex is unique in providing Universal Boot images of boot code. Universal Boot images contain x86 Boot code and UEFI Boot code. These images provide multiplatform support for BFS.

Universal Boot transparently determines your system platform type and automatically executes the proper boot code image in the adapter. These boot code images reside in the adapter flash memory, which allows for easier adapter portability and configuration between servers.

Boot code is configured using the following methods:
- x86 Boot code uses the Emulex FC BIOS utility, which is included in the x86 Boot code.
- UEFI Boot code is configured using the UEFI HII.

1.2 Abbreviations

The following table lists the acronyms and abbreviations used in this document.

Table 1: Acronyms and Abbreviations

Acronym/Abbreviation	Description
ANA	Asymmetric Namespace Access
BBS	BIOS Boot Specification
BFS	boot from SAN
D_ID	destination identifier
DUD	driver update disk
EDD	enhanced disk device
EUID	effective unique identifier
FA-PWWN	Fabric Assigned Port Word Wide Name
FCA	Fibre Channel adapter
FCP	Fibre Channel Protocol
FEC	forward error correction
FL_Port	fabric loop port
GPT	GUID partition table
GUID	globally unique identifier
HPE	Hewlett Packard Enterprise
IDE	Integrated Drive Electronics
I/O	input/output
MBR	master boot record
MPIO	multipath I/O
MSA	Modular Smart Array
NGUID	Namespace globally unique identifier
NIS/NIS+	Network Information Service/Network Information Service Plus
NQN	NVMe qualified name
PLOGI	Port Login
SFS	SAN foundation software
UUID	universally unique identifier

Chapter 2: Booting from SAN

This chapter provides instructions for installing and using all types of boot code provided by Emulex.

When booting from SAN, the storage device is typically identified by its WWPN and a LUN for an FC boot disk and by an NQN and namespace for an NVMe disk. By extending the server system BIOS, BFS capability is provided by the boot code contained on an Emulex adapter in the server. If properly configured, the adapter then permanently directs the server to boot from a logical unit (disk) on the SAN as if it were a local disk.

NOTE: Not all procedures are required. Emulex HBAs usually ship from the factory with the latest version of boot code installed and enabled, so you do not need to install or enable boot code in those cases.

This chapter describes setting up a system to boot from SAN. The specific procedure to follow is determined by the system architecture and the operating system.

Table 2: Boot from SAN Procedures

Architecture	Operating System	Procedure
x86_64	Windows Server	Section 2.1, Windows Server
	Linux	Section 2.2.1, Configuring BFS on Linux or VMware
	VMware	
	Solaris (SFS driver)	Section 2.3.1, Configuring BFS on Solaris (x86 and UEFI)
Arm aarch64	Linux (RHEL, OpenEuler, Kylin, UniontechOS, Ubuntu)	Section 2.2.1, Configuring BFS on Linux or VMware

2.1 Windows Server

This section describes configuring BFS for Windows systems and installing Windows Server on an FC boot disk.

2.1.1 Configuring Boot from SAN on Windows Server (x86)

To configure BFS, perform these steps:

1. If necessary, enable the boot code on the adapter (see Chapter 3, Updating and Enabling the Boot Code).

2. Enable the adapter to BFS (see Section 5.3, Enabling an Adapter to Boot From SAN).

3. If you want to use a topology other than the default, you must change the topology setting before configuring boot devices (see Section 5.7.1, Changing the Topology).

4. Configure boot devices (see Section 5.5, Configuring Boot Devices).

5. If desired, configure the boot options on the adapter (see Section 5.6, Configuring Advanced Adapter Parameters).

6. Install the operating system on an FC boot disk (see Section 2.1.3, Installing Windows Server on an FC Boot Disk).

2.1.2 Configuring Boot from SAN on Windows Server (UEFI)

To configure boot from SAN, perform these steps:

1. If necessary, enable the boot code on the adapter (see Chapter 3, Updating and Enabling Boot Code).

3. Enable the adapter to boot from SAN (see Section 4.4, Setting Boot from SAN).

4. Add boot devices (see Section 4.6, Adding Boot Devices).

6. If desired, change the order in which targets are discovered by the boot driver (see Section 4.8, Changing the Boot Device Order).

7. If desired, configure the boot options on the adapter (see Section 4.9, Configuring HBA and Boot Parameters).

8. Install the operating system on an FC boot disk (see Section 2.1.3, Installing Windows Server on an FC Boot Disk).

2.1.3 Installing Windows Server on an FC Boot Disk

This procedure installs a supported Windows Server operating system from a local media (CD/DVD/USB) onto an unformatted FC disk drive and configures the system to boot from the SAN disk drive. The LUNs must be created in the SAN storage device and zoned appropriately to the host adapter WWN before starting the installation.

NOTE: The computer system BIOS might be pre configured to boot from another controller to take precedence over the Emulex adapter during boot. If this occurs, you must disconnect or disable the other adapter. This allows you to configure and build the operating system on the drive connected to the Emulex adapter.

NOTE: Restrict the number of paths (from the HBA to the boot disk) to one during installation. Additional paths can be added after the operating system is installed and the multipath or MPIO software is configured.

1. From the server system setup, ensure that the local boot media is the first device in the boot order list.

2. Enable the adapter BIOS setting to allow SAN boot in the Emulex x86 Boot utility or the UEFI HII.

3. Configure the boot target and LUN in the Emulex x86 Boot utility or the UEFI HII to point to the desired target (for configuring x86 Boot code, see Section 5.5, Configuring Boot Devices, and for configuring UEFI Boot code, see Section 4.6, Adding Boot Devices).

4. Boot the host server with the Windows Server operating system media inserted. Follow the on-screen prompts to install the appropriate version of Windows Server.

   The Windows installation exposes all available and visible LUNs as disks and partitions numbered 0 to *N*, where *N* is the highest number available. These numbers typically are the LUN numbers assigned by the array.

5. Select the disk on which you want to install the operating system.

NOTE: If the FC disk is not visible at this point, download the driver kit from the Broadcom website to your system. Extract and copy the driver files to a formatted disk or USB drive.
Click **Load Driver**. Browse to the disk or USB device where the driver is located to load the Storport Miniport driver for the appropriate operating system. The driver location and driver are displayed in the **Select driver to be installed** window. After the driver is loaded, the FC disks should be visible.

6. Select the disk and follow the system prompts in the Windows installation.

7. After the installation is complete, a Boot Option variable called Windows Boot Manager is populated with a Media Device path pointing to the Windows bootloader utility.

   The Windows Boot Manager option is inserted as the first boot option in the boot order list. The CD/DVD boot is the second device in the boot order list.

   Upon reboot, the system boots from the LUN set up on the SAN.

   If multipath to the HBA is required, install or enable multipathing software (or Windows native MPIO feature) and reboot the server. After rebooting, verify the multipath configuration.

2.1.3.1 Directing a UEFI-Based Server to a Windows Server Operating System Image (Installed as UEFI-Aware) Already Installed on the SAN

This installation procedure assumes that a LUN exists in the SAN storage device, that it is zoned appropriately to the host adapter WWN, and that a UEFI-aware operating system resides on the target LUN.

To direct a UEFI-based server to a Windows Server operating system image, perform these steps:

1. Enable BFS in the UEFI HII.

2. Configure the boot target and LUN in the UEFI HII to point to the desired target (see Section 4.6, Adding Boot Devices).

3. Select **Boot Manager** from the System UEFI configuration manager.

4. Select **Add Boot Option**.

5. Identify the desired target in the list, and continue down the explorer path until you locate the `bootmgfw.efi` file. This file is the bootloader utility for your Windows Server UEFI-aware operating system installation.

6. Enter a boot device description (such as win2K16_UEFI_SAN) and optional data (if desired) for this device and select **Commit Changes.**

7. From the Boot Manager, select **Change Boot Order.**

8. Move the previous input description name (win2K16_UEFI_SAN) to the desired position in the boot order.

9. Select **Commit Changes.**

   The **Start Options** list now reflects the boot order changes. Upon reboot, the server is able to boot from this target LUN on the SAN.

2.1.3.2 GUID Partition Table

The GPT provides a more flexible mechanism for partitioning disks than the older MBR partitioning scheme that has been common to PCs. MBR supports four primary partitions per hard drive and a maximum partition size of 2 TB. If the disk is larger than 2 TB (the maximum partition size in a legacy MBR), the size of this partition is marked as 2 TB and the rest of the disk is ignored.

The GPT disk itself can support a volume up to 2^{64} blocks in length (for 512-byte blocks, this is 9.44 ZB). The GPT disk can also theoretically support unlimited partitions.

NOTE: By default, Microsoft Windows Server operating systems install with a GPT-formatted disk on a UEFI-aware server.

For more information on the GUID partition table, refer to the Microsoft website and search for the terms *Windows* and *GPT FAQ.*

2.1.4 Configuring Boot from SAN for NVMe over FC on a Windows Server

BFS for NVMe over FC in Windows is supported only on the following operating systems:

- Windows Server 2016
- Windows Server 2019
- Windows Server 2022

To configure BFS for NVMe over FC on a Windows Server, perform the following steps:

1. Follow the instructions in Section 4.15, NVMe over FC Boot Settings, to configure NVMe over FC boot and to add an NVMe boot device using the UEFI HII.

NOTE: Before starting the installation, zone the target WWN appropriately to the initiator WWNs. Create the namespace of appropriate size in the NVMe target and map it to the initiator NQNs (for instructions see the note in Section 4.15.1, Enabling NVMe over FC BFS).

2. Install the operating system on an NVMe over FC boot disk (see Section 2.1.3, Installing Windows Server on an FC Boot Disk).

If the disk number of a GPT disk is changed during loading of the out-of-box driver, the disk will be offline and you will be unable to clear the read-only attribute. To work around this, you must first change the SAN policy to `online all` before loading the out-of-box driver and set the disk as a raw disk.

To set the SAN policy, perform the following steps:

1. Go to the command prompt by typing `Shift+F10`.
2. Enter the `diskpart` utility by typing `diskpart`.
3. Change the SAN policy in `diskpart` by typing `san policy=onlineall`.
4. Exit the `diskpart` utility by typing `exit`.

To set the disk as a raw disk, perform the following steps:

1. Go to the command prompt by typing `Shift+F10`.
2. Enter the `diskpart` utility by typing `diskpart`.
3. Select the disk by typing `sel disk <disk_number>`.
4. Set the selected disk as a raw disk by typing `clean`.
5. Exit the `diskpart` utility by typing `exit`.

2.2 Linux and VMware

This section describes configuring BFS on supported Linux and VMware operating systems.

2.2.1 Configuring BFS on Linux or VMware

To configure BFS on Linux or VMware, perform these steps:

1. If necessary, install or update the boot code on the adapter. See Chapter 3, Updating and Enabling the Boot Code.
2. If necessary, enable the boot code on the adapter. See Chapter 3, Updating and Enabling the Boot Code.

3. Enable the adapter to boot from SAN. For x86 Boot code, see Section 5.3, Enabling an Adapter to Boot From SAN, and for UEFI Boot code, see Section 4.4, Setting Boot from SAN.

4. If you want to use a topology other than the default, you must change the topology setting before configuring boot devices. For x86 Boot code, see Section 5.7.1, Changing the Topology, and for UEFI Boot code, see Section 4.9.1, Changing the Topology.

5. Configure boot devices. For x86 Boot code, see Section 5.5, Configuring Boot Devices, and for UEFI Boot code, see Section 4.6, Adding Boot Devices.

6. If desired, configure the boot options on the adapter. For x86 Boot code, see Section 5.6, Configuring Advanced Adapter Parameters, and for UEFI Boot code, see Section 4.9, Configuring HBA and Boot Parameters.

7. Use the driver on the operating system distribution disk to boot the system. If necessary, you can then update the driver to the desired version.

2.2.2 Installing Linux or VMware on an FC Boot Disk

The LUNs should be created in the SAN storage device and zoned appropriately to the host adapter's WWN before you start the installation.

NOTE: The computer system BIOS might be pre configured to boot from another controller to take precedence over the Emulex adapter during boot. If this occurs, you must disconnect or disable the other adapter. This allows you to configure and build the operating system on the drive connected to the Emulex adapter.

Ensure that all paths (from the HBA to the boot disk) are enabled during installation.

1. From the server system setup, ensure that CD/DVD is the first device in the boot order list.

2. If necessary, install or update the boot code on the adapter. See Chapter 3, Updating and Enabling Boot Code.

3. If necessary, enable the boot code on the adapter. See Chapter 3, Updating and Enabling Boot Code.

4. Configure the boot target and LUN in the Emulex x86 Boot utility (see Section 5.5, Configuring Boot Devices) or in the UEFI HII (see Section 4.6, Adding Boot Devices) to point to the desired target.

5. Boot the host server with the operating system DVD inserted. Follow the on-screen prompts to install the operating system.

   The installation exposes all available and visible LUNs as disks and partitions.

6. Select the disk on which you want to install the operating system.

NOTE: On Linux distributions, if the FC disk is not visible at this point, contact Broadcom Technical Support for the DUD. The DUD ISO image contains the drivers.
Restart the installation process. On the installation splash screen, press **Tab** and in the boot command line, add the `linux dd` option to boot the command line. This action prompts you for driver disks, at which point you must provide the DUD ISO image.
Select the disk and follow the system prompts in the Linux or VMware installation.

7. After the installation is complete, a Boot Option variable specific to the operating system installed is populated with a Media Device path pointing to the bootloader utility.

   Upon reboot, the system boots from the LUN set up on the SAN.

8. After rebooting, verify the multipath configuration.

2.2.3 Configuring BFS for NVMe over FC on Linux

BFS for NVMe over FC in Linux is supported only on the following operating systems:

- SLES 12 SP5
- SLES 15 SP3 and later
- RHEL 8.5 and later
- RHEL 9.0 and later

To configure BFS for NVMe over FC on Linux, perform the following steps:

1. Follow the instructions in Section 4.15, NVMe over FC Boot Settings, to configure NVMe over FC boot and to add an NVMe boot device using the UEFI HII.

NOTE: Before starting the installation, zone the target WWN appropriately to the initiator WWNs. Create the namespace of appropriate size in the NVMe target and map it to the initiator NQNs (for instructions see the note in Section 4.15.1, Enabling NVMe over FC BFS).

NOTE: Step 2 to Step 13 are not required on RHEL 9.1 and SLES 15 SP4 and later. Step 2 to Step 13 might not be required on SLES 15 SP3.

2. Obtain the `elx-nvmefc-dm-<build>-n.tar.gz` pre-installation kit from www.broadcom.com. The pre installation kit includes Emulex scripts that run during the operating system installation process and that set up BFS on the boot device.

3. To prepare a pre-installation USB storage device, continue to Step 4. Otherwise, to prepare a pre-installation CD or DVD, perform the following steps:

    a. Create a folder from which an ISO image will be generated by typing the following command:
    ```
    # mkdir <directory_name>
    ```
    b. Copy the `elx-nvmefc-dm-<build>-n.tar.gz` file that you obtained in Step 2 to the directory you created in Step a by typing the following command:
    ```
    # cp elx-nvmefc-dm-<build>-n.tar.gz ISOROOT
    ```
    c. Generate the ISO image by typing the following command:
    ```
    # genisoimage -o sample.iso ISOROOT
    ```
    Messages similar to the following are displayed:
    ```
    I: -input-charset not specified, using utf-8 (detected in locale settings)
    Total translation table size: 0
    Total rockridge attributes bytes: 0
    Total directory bytes: 0
    Path table size(bytes): 10
    Max brk space used 0
    4800 extents written (9 MB)
    ```
    d. Insert a blank CD or DVD into the computer's disc writer, and burn the ISO image onto the CD by typing the following command:
    ```
    # wodim -v speed=2 dev=/dev/cdrom sample.iso
    ```
    e. Continue to Step 5.

4. To prepare a pre-installation USB storage device, perform the following steps:

   a. Locate the USB storage device by typing the following command:

      ```
      # fdisk -l
      ```

      Messages similar to the following are displayed:

      ```
      Disk /dev/sda: 15.9 GB, 15938355200 bytes, 31129600 sectors
      Units = sectors of 1 * 512 = 512 bytes
      Sector size (logical/physical): 512 bytes / 512 bytes
      I/O size (minimum/optimal): 512 bytes / 512 bytes
      Disk label type: dos
      Disk identifier: 0x0d7584c8
      ```

   b. Mount a partition on the USB device by typing the following command:

      ```
      # mount /dev/sda1 /mnt
      ```

   c. Copy the `elx-nvmefc-dm-<build>-n.tar.gz` file you obtained in Step 2 to the mounted partition by typing the following command:

      ```
      # cp elx-nvmefc-dm-<build>-n.tar.gz /mnt
      ```

5. Attach the operating system installation media to the server, and reboot or power on the server using the UEFI Boot Menu.

6. When the installation GUI appears, open a command-line shell by pressing **Ctrl**+**Alt**+**F5**.

7. Access the pre-installation device you created in Step 3 or Step 4, and copy its contents to a temporary directory by performing the following steps:

   a. Create a temporary directory on which to mount the pre-installation medium by typing the following command:

      ```
      # mkdir /<directory_name>
      ```

   b. Either insert the CD or DVD you created in Step 3 or attach the USB storage device you created in Step 4, and mount the device to the temporary directory by typing one of the following commands:

      For the CD or DVD, where `/dev/cdrom` is the device detected as the CD or DVD device:

      ```
      # mount /dev/cdrom /<directory_name>
      ```

      For the USB storage device, where /dev/sda1 is the device detected as the USB storage device:

      ```
      # mount /dev/sda1 /<directory_name>
      ```

   c. Copy the contents of the mounted device to a temporary subdirectory of the directory you created in Step a by typing the following command:

      ```
      # cp -rp /<directory_name> /<subdirectory_name>
      ```

   d. Unmount the device you mounted in Step b by typing the following command:

      ```
      # unmount /<directory_name>
      ```

   e. Remove the CD, DVD, or USB storage device.

8. Change to the temporary directory to which the pre-installation kit was copied by typing the following command:

   ```
   # cd /<subdirectory_name>
   ```

9. Untar the pre-installation media kit by typing the following command:

   ```
   # tar zxf elx-nvmefc-dm-<build>-n.tar.gz
   ```

10. Change directory to the extracted kit directory by typing the following command:

    ```
    # cd elx-nvmefc-dm-<build>-n/
    ```

11. Extract the kit that corresponds to the operating system distribution that you are installing by typing the following command:

    ```
    # tar zxf nvmefc-dm-<build>.n.<operating_system_distro>.tar.gz
    ```

12. Change the directory to the extracted directory by typing the following command:

    ```
    # cd nvmefc-dm-<build>.n.<operating_system_distro>
    ```

13. Run the `nvme_configure_bfs.sh` script by typing the following command and following the prompts that appear:

    ```
    # ./nvme_configure_bfs.sh
    ```

For SLES installations, proceed with the instructions in Section 2.2.3.1, Continuing a SLES Installation. For RHEL installations, proceed with the instructions in Section 2.2.3.2, Continuing a RHEL Installation

2.2.3.1 Continuing a SLES Installation

This section describes the remaining steps required to complete the SLES operating system installation.

1. Return to the operating system installation GUI by pressing **Ctrl+Alt+F7**.

2. Select the product to install and click **Next**.

3. Proceed with the operating system installation. When the **Suggested Partitioning** window appears, select the device name that corresponds to the subsystem NQN and namespace ID that you configured for booting in the UEFI HII.

4. Partition the device as desired (refer to the operating system documentation for details).

5. Proceed with the operating system installation. When the **Installation Settings** window appears, follow these steps:

    a. Click **Software**. The **Software Selection and System Tasks** window appears.
    b. Click **Details**.
    c. Select the **Search** tab, type **nvme-cli** in the **Search** field, and click **Search**.
    d. Select the **nvme-cli** option, and click **Accept**.
    e. Click **Continue** to install any other packages that you require, and continue the installation process.

After the installation is complete, the system reboots using the newly installed media.

2.2.3.2 Continuing a RHEL Installation

This section describes the remaining steps required to complete the RHEL operating system installation.

1. Return to the operating system installation GUI by pressing **Ctrl+Alt+F6**.

2. Proceed with the operating system installation. When the **Installation Summary** window appears, select **Installation Destination** in the **System** section.

3. Click the **Refresh** link. The **Rescan Disks** popup appears.

4. Click **Rescan Disks**.

5. After **Disk rescan complete** appears, click **OK**.

6. Select the device name that corresponds to the subsystem NQN and namespace ID you configured for booting in the UEFI HII.

7. Partition the device as desired.

8. Proceed with the operating system installation.

9. When you are prompted to reboot the system, do not reboot. Instead, reopen the command-line shell by pressing **Ctrl+Alt+F6**.

10. Run the post_install_config.sh script file, which is one of the Emulex script files that was copied to the system in Section 2.2.3, Configuring BFS for NVMe over FC on Linux, by typing the following command:

    `# ./post_install_config.sh`

NOTE: Step 10 is not required on RHEL 9.1.

11. Return to the operating system installation GUI by pressing **Ctrl+Alt+F6**.

12. Click **Reboot** to complete the operating system installation.

After the installation is complete, the system reboots using the newly installed media.

NOTE: The operating system installer has a known issue in which the installer fails to set a UEFI boot entry. To work around this issue, perform the following steps:

1. Press **F11** to enter the UEFI Boot Menu.

2. Select the UEFI boot entry that is mapped to the adapter port that is configured for BFS.

3. Press **Enter**. The UEFI utility automatically adds a UEFI boot path for the adapter, and the operating system boots.

2.2.4 Configuring BFS for NVMe over FC on VMware

BFS for NVMe over FC in VMware is supported only on ESXi 7.0 U1 and subsequent operating systems.

To configure BFS for NVMe over FC on ESXi 7.0 U1 and subsequent operating systems, perform the following steps:

1. Follow the instructions in Section 4.15, NVMe over FC Boot Settings, to configure NVMe over FC boot and to add an NVMe boot device using the UEFI utility.

NOTE: Before starting the installation, zone the target WWN appropriately to the initiator WWNs. Create the namespace of appropriate size in the NVMe target and map it to the initiator NQNs (for instructions, see the note in Section 4.15.1, Enabling NVMe over FC BFS).

2. Attach the operating system installation media to the server, and reboot or power on the server using the UEFI Boot Menu.

3. When the installation GUI appears, wait for the **Welcome to the VMware ESXi (*x.x*) Installation** screen.

4. Press **Alt+F1** to switch to the ESXi console window.

5. Enter `root` as the login credential and press **Enter**. Leave the password blank and press **Enter** to continue.

6. Since `hostd` is not available use `localcli` instead of `esxcli` to perform operations in this shell.

7. At the command prompt, run the `localcli nvme info get` command to obtain the ESXi host NQN.

8. Make a note of this NQN string and provide it to the storage administrator to configure the boot namespace in the NVMe target.

9. Issue the `localcli storage core adapter list` command to obtain a list of available FC adapters.

10. After configuring the host NQN on the target, trigger a discovery at the host to discover the newly configured namespace for boot. Use the `localcli storage core adapter list` command to obtain a list of available FC adapters. Use one of the following methods to trigger a discovery:

    a. Perform a LIP reset to a chosen FC adapter:
    ```
    #localcli storage core adapter list localcli storage san fc reset -A <vmhba#>
    ```

    b. Use the VMware native nvme fabrics command to discover an NVMe controller on a specific target port through a specific NVMe adapter:
    ```
    #localcli nvme fabrics discover -a<vmhba> -W=<Target WWNN> -w=<Target WWPN>
    ```
    For example:
    ```
    #localcli nvme fabrics discover -a vmhba66 -W=20:2c:00:a0:98:c1:fe:52 - w=20:6d:00:a0:98:c1:fe:52
    ```

11. Type the following commands to verify if the new subsystem and namespace are discovered:
    ```
    #localcli nvme controller list
    #localcli nvme namespace list
    ```

12. Press **ALT+F2** to return to the installation GUI and continue the installation process.

13. The newly assigned namespace appears in the **Select a Disk to Install or Upgrade** window under the list of remote storage devices.

14. Select the displayed namespace and continue the operating system installation following the on-screen prompts.

15. After successful operating system installation, ensure that the Emulex adapter is configured as the first boot device in the BIOS and proceed with boot.

2.3 Solaris

This section explains how to configure BFS on the Solaris operating system.

2.3.1 Configuring BFS on Solaris (x86 and UEFI)

To configure BFS, perform these steps:

1. If necessary, enable the boot code on the adapter (see Chapter 3, Updating and Enabling Boot Code).

2. Enable the adapter to BFS.For x86 Boot code, see Section 5.3, Enabling an Adapter to Boot From SAN, and for UEFI Boot code, see Section 4.4, Setting Boot from SAN.

3. If you want to use a topology other than the default, you must change the topology setting before configuring boot devices. For x86, see Section 5.7.1, Changing the Topology, and for UEFI, see Section 4.9.1, Changing the Topology.

4. Configure boot devices. For x86, see Section 5.5, Configuring Boot Devices, and for UEFI, see Section 4.6, Adding Boot Devices.

5. If desired, configure the boot options on the adapter. For x86, see Section 5.6, Configuring Advanced Adapter Parameters. For UEFI, see Section 4.9, Configuring HBA and Boot Parameters.

6. Boot the Solaris installation CD/DVD and follow the prompts.

NOTE: If you need help determining the LUNs to select for BFS, see Section 2.3.1.1, Determining LUNs to Select for BFS.

2.3.1.1 Determining LUNs to Select for BFS

To determine which LUNs to select, perform these steps:

1. Open a terminal window and leave it open.

2. In the terminal window, select the LUN you will use as the SAN boot disk (not the local drive) using the `luxadm probe` command. This command shows all the available LUNs. Record this LUN information, which is used throughout this procedure. LUN 0 is used in the example:

```
luxadm probe
Found Fibre Channel device(s):
Node WWN:50060e8003823800 Device Type:Disk device
Logical Path:/dev/rdsk/c5t226000C0FF9833AFd6s2
Node WWN:50060e8003823800 Device Type:Disk device
Logical Path:/dev/rdsk/c5t226000C0FF9833AFd6s2
Node WWN:50060e8003823800 Device Type:Disk device
```

3. Copy the <*/dev/rdsk/nnn*> part of the path statement for a drive.

4. In the terminal window, use the `luxadm display` command to show the WWPN or the LUN for which you selected the path in the prior step:

```
luxadm display </dev/rdsk/nnn>
```

5. Record this LUN or WWPN information for use in the procedure.

2.3.2 Installing Solaris from a Network Image

The system must have a DVD drive and must be part of the site's network and naming service. If you use a naming service, the system must already be in a service, such as NIS, NIS+, DNS, or LDAP. If you do not use a naming service, you must distribute information about this system by following your site's policies.

NOTE: This procedure assumes that the system is running the Volume Manager. If you are not using the Volume Manager to manage media, refer to the Oracle system administration documentation that covers devices and file systems.

To install from a network image, perform these steps:

1. Log on as a superuser or equivalent.

2. Insert the Solaris DVD in the system's drive.

3. Create a directory to contain the DVD image.
   ```
   # mkdir -p <install_dir_path>
   ```
   The <*install_dir_path*> specifies the directory where the DVD image is to be copied.

4. Change to the `Tools` directory on the mounted disk.
   ```
   # cd /cdrom/cdrom0/Solaris_11/Tools
   ```

5. Remove the `SUNWemlxu` and `SUNWemlxs` from the /<*install_dir_path*>/Solaris/Tools/Boot directory.

6. Unzip the `lpfc` driver to a temporary directory:

7. `pkgadd -R/<install_dir_path>/Solaris/Tools/Boot -d/tmp`

8. Modify the `elxfc.conf` file to use persistent binding.

9. Copy the DVD image in the drive to the install server's hard disk.

   ```
   # ./setup_install_server install_dir_path
   ```

   The `<install_dir_path>` specifies the directory where the DVD image is to be copied.

   The `setup_install_server` command indicates whether you have enough disk space available for the Solaris software disk images. To determine available disk space, use the `df -kl` command.

10. Decide whether you need to make the install server available for mounting:

    – If the install server is on the same subnet as the system to be installed or you are using DHCP, you do not need to create a boot server. Proceed to Step 11.

    – If the install server is not on the same subnet as the system to be installed and you are not using DHCP, complete the following steps:

    a. Verify that the path to the install server's image is shared appropriately.

       ```
       # share | grep <install_dir_path>
       ```

       The `<install_dir_path>` specifies the path to the installation image where the DVD image was copied:

       • If the path to the install server's directory is displayed and **anon=0** is displayed in the options, proceed to Step 11.

       • If the path to the install server's directory is not displayed or **anon=0** is not in the options, continue and make the install server available to the boot server. Using the `share` command, add this entry to the `/etc/dfs/dfstab` (all on one line).

       ```
       share -F nfs -o ro,anon=0 -d "install server directory" <install_dir_path>
       ```

    b. Verify that the `nfsd` daemon is running, or start the `nfsd` daemon.

       • If the install server is running the current Solaris release or a compatible version, type the following command.

       ```
       # svcs -l svc:/network/nfs/server:default
       ```

       • If the `nfsd` daemon is online, continue to Step c. If the `nfsd` daemon is not online, start it.

       ```
       # svcadm enable svc:/network/nfs/server
       ```

    c. Share the install server.

       ```
       # shareall
       ```

11. Change directories to root (/).

    ```
    # cd /
    ```

12. Eject the Solaris DVD.

13. (Optional) Patch the files that are located in the miniroot on the net install image that was created by `setup_install_server`. Patching a file might be necessary if a boot image has problems. For more information, refer to the Solaris operating system documentation.

2.3.2.1 Installing Solaris by Migrating an Image from a Local SCSI Disk

To install Solaris by migrating an image from a local SCSI disk, perform these steps:

1. Type the following at the OBP prompt:

   ```
   show-devs
   ```

   The ID information for each found adapter is displayed, for example:

   ```
   /pci@5d,700000/lpfc@1 select-dev
   ```

2. Select the Emulex adapter on which you want to enable BFS by entering the path to the adapter, for example:

   ```
   " /pci@5d,700000/lpfc@1" select-dev
   ```

3. To view the current boot device ID, type the following command:

   ```
   show-devs
   " /pci@5d,700000/lpfc@1" select-dev /* to select lpfc@1 (for example) */
   .boot-id
   ```

   Make a note of the WWPN, D_ID, or AL_PA returned from the probe and write down the corresponding boot entry.

4. To enable boot from the SAN, set the boot device ID to the SAN device from which you want to boot, for example:

```
" /pci@5d,700000/lpfc@1" select-dev
wwpn|did|alpa lun target_id set-boot-id
unselect-dev
```

   where:

   – `wwpn|did|alpa` is the device WWPN, D_ID, or AL_PA of the storage device.
   – `lun` is the LUN number in hexadecimal. To enter it in decimal, enter `d# [lun]`.
   – `target_id` is the target ID in hexadecimal. To enter it in decimal, enter `d# [target_id]`.

NOTE: Use the WWPN in most cases. The D_ID and AL_PA might change between boots, causing the SAN boot to fail, unless the D_ID and AL_PA are specifically configured to not change between boots.

   Example 1: alpa = e1, lun = 100 (decimal), and target id = 10 (decimal):

```
alpa e1 d# 100 d# 10 set-boot-id
```

   Example 2: wwpn = 50000034987AFE, lun = af (hexadecimal), and target id = 10 (decimal):

```
wwpn 50000034987AFE af d# 10 set-boot-id
```

   Example 3: did = 6312200, lun = 25 (hexadecimal), and target id = f (hexadecimal):

```
did 6312200 25 f set-boot-id
```

5. Boot to the original local disk to set up the newly defined FC disk. Type the following command:

```
boot <local_disk>
```

   where `<local_disk>` is the complete path or the alias of the original boot disk.

6. Run the format utility by typing the following command:

```
format
```

7. Select the target disk to become the new boot disk (for example, `c1t1d0`).

8. Select the partition option and partition the disk as desired.

9. Select the label option and write a volume label to the target disk.

   For help with the format utility, refer to the man page, `man format`.

10. Install the boot on partition 0 of the target disk. (Type this command on one line.)

```
installboot /usr/platform/ `uname -i`/lib/fs/ufs/bootblk /dev/rdsk/c1t1d0s0
```

11. Create a file system for each partition that contains a mounted file system.

```
newfs -v /dev/rdsk/c1t1d0s0 (becomes root)
newfs -v /dev/rdsk/c1t1d0s6 (becomes usr)
newfs -v /dev/rdsk/c1t1d0s7 (becomes export/home)
```

12. Create temporary mount points for the new partitions.

```
mkdir root2
mkdir usr2
mkdir export2
```

13. Mount, copy, and then unmount the `usr2` file system.

```
mount /dev/dsk/c1t1d0s6 /usr2
c0t0d0s6 ufsdump 0f - /dev/rdsk/c0t0d0s6 | (cd /usr2; ufsrestore rf -)
umount /usr2
```

14. Copy the `export/home` file system.

```
mount /dev/dsk/c1t1d0s7 /export2
ufsdump 0f - /dev/rdsk/c0t0d0s7 | (cd /export2; ufsrestore rf -)
umount /export2
```

15. Perform a copy.
```
mount /dev/dsk/clt1d0s0 /root2
ufsdump 0f - /dev/rdsk/c0t0d0s0 | (cd /root2; ufsrestore rf -)
```

16. Edit `/root2/etc/vfstab`, changing the controller number, target number, and LUN number to point to the new FC boot disk. For example, if the FC boot disk is c1t1d0, replace all local disk entries of c0t0d0 with c1t1d0.

    Currently the file shows:

    `/dev/dsk/c0t0d0s1` (swap)

    `/dev/dsk/c0t0d0s0 and /dev/rdsk/c0t0d0s0` (root)

    `/dev/dsk/c0t0d0s6 and /dev/rdsk/c0t0d0s6` (usr)

    `/dev/dsk/c0t0d0s7 and /dev/rdsk/c0t0d0s7` (export)

    Edit the file to show:

    `/dev/dsk/c1t1d1s1` (swap)

    `/dev/dsk/clt1d0s0 and /dev/rdsk/c1t1d0s1` (root)

    `/dev/dsk/clt1d0s6 and /dev/rdsk/c1t1d0s6` (usr)

    `/dev/dsk/clt1d0s7 and /dev/rdsk/c1t1d0s7` (export)

17. Reboot the system.
```
sync
sync
halt
reset-all
```

18. Boot to disk.
```
boot disk
```
    The system should boot to the FC disk.

19. View the current dump device setting.
```
dumpadm
```

20. Change the dump device to the swap area of the FC drive.
```
dumpadm -d /dev/dsk/c1t1d0s1
```
    where `/dev/dsk/c1t1d0s1` is a sample path to the swap area of the FC drive.

Chapter 3: Updating and Enabling the Boot Code

Use the Emulex utilities to update the firmware and boot code to a newer version and enable boot code. The utility that you use depends on the operating system and, in some cases, the driver type or system architecture. Table 3 indicates the utilities that you can use to install and update the boot code.

Table 3: Utilities that Update and Enable the Boot Code

Operating System	Emulex HBA Manager GUI	Emulex HBA Manager CLI	UEFI HII	elxmgmt	FCA emlxadm	Management Utilities
Windows	X	X	X	—	—	X
Linux	X	X	X	—	—	X
Solaris	X	X	X	—	X	—
VMware	—	—	X	X	—	X

After you decide which utility to use, refer to the appropriate procedure:

- Emulex HBA Manager GUI:
  - For Windows, Linux, and Solaris, refer to the *Emulex HBA Manager Application User Guide*.
- Emulex HBA Manager CLI:
  - For Windows, Linux, and Solaris, refer to the *Emulex HBA Manager Application Command Line Interface User Guide*.
- UEFI HII: See Chapter 4, Configuring Boot Using the UEFI HII.
- elxmgmt: VMware only; refer to the *Emulex Drivers for VMware ESXi User Guide*.
- Management utility: Refer to the *Emulex Elxflash and LpCfg Utilities User Guide.*

Chapter 4: Configuring Boot Using the UEFI HII

You can configure boot functions using the UEFI HII. If you have several adapters in your system, the UEFI system firmware or boot code uses the highest version present on your adapters.

NOTE: UEFI boot supports booting from a 4K native disk.

NOTE: UEFI NVMe support is not available in the Arm aarch64 UEFI driver.

4.1 Prerequisites

Before you configure any UEFI parameters on your system, ensure that all of the adapters in the system have the same and latest UEFI version to prevent incompatibilities with the firmware or boot code support.

NOTE: Adapters with older versions of UEFI boot code are managed by the more recent version.

4.2 Starting the UEFI HII

Depending on the server UEFI configuration, the UEFI HII might appear under different setup menus in the server system firmware or BIOS.

Navigation through the menus and configuration windows in the UEFI HII depends on the system BIOS. Refer to the server documentation for navigation instructions.

To configure an Emulex HBA using the UEFI HII, perform these steps:
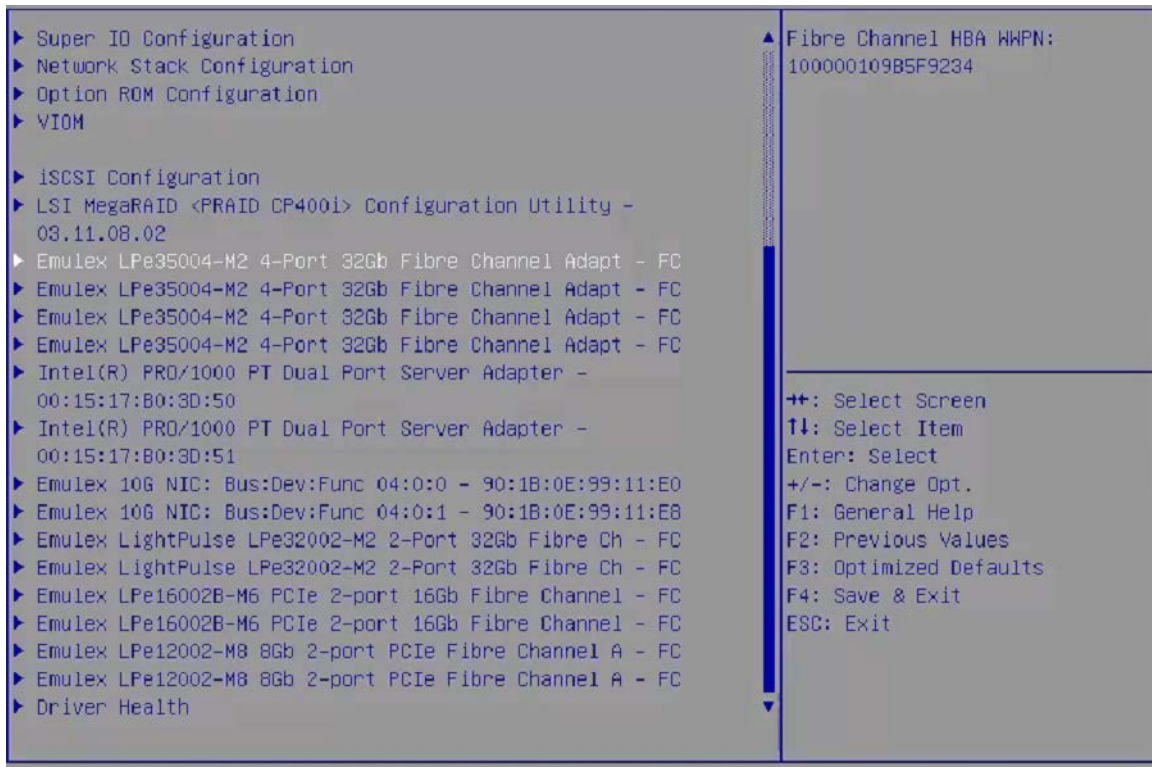
1. During system boot, follow the instructions that appear on the screen to access the UEFI HII, such as pressing the server specific hotkey.

NOTE: After you enter the UEFI HII, any dynamic changes made to the SAN are not reflected unless the system is reset or the adapter is reconnected.

2. Select the menu item that allows you to configure devices.

   The window that appears displays a list of all the adapters in the system (Figure 1). Your list might vary depending on the installed adapters.

Figure 1: Adapter List

3. Select the adapter that you want to configure and press **Enter**.

The main configuration menu is displayed (Figure 2).

Figure 2: Main Configuration Menu


```
              Emulex LightPulse LPe32002-M2 2-Port 32Gb Fibre Ch - FC

  001: LPe32002-M2      PCIe8.0GT/s , x8              Set to Enable to SCAN Fabric.
  LPe32002-M2 Port Name : 10000090FA942A0E
  Seg#: 00 Bus#: 02 Dev#: 00 Func#: 00

  Set Boot from SAN              <Enable>
▶ Scan for Fibre Devices
▶ Add Boot Device
▶ Delete Boot Device
▶ Change Boot Device Order
▶ Configure HBA and Boot Parameters
▶ Set Emulex Adapter to Default Settings
▶ Display Adapter Info
▶ Legacy Only Configuration Settings
▶ Request RESET or RECONNECT to Make Changes Active
▶ Emulex Firmware Update Utility
▶ Emulex NVMe over FC Boot Settings



                      F9=Reset to Defaults        F10=Save
  ↑↓=Move Highlight    <Enter>=Select Entry        Esc=Exit
```

NOTE: UEFI NVMe support is not available on the Arm aarch64 UEFI driver. The **Emulex NVMe over FC Boot Settings** option is unavailable and is grayed-out.

4.3 Configuring Boot in the UEFI HII

The UEF HII allows you to perform the following tasks to configure UEFI boot code:

- Set BFS (see Section 4.4, Setting Boot from SAN).
- Scan for Fibre devices (see Section 4.5, Scanning for Fibre Devices).
- Add and delete boot devices (see Section 4.6, Adding Boot Devices and Section 4.7, Deleting Boot Devices).
- Change boot device order (see Section 4.8, Changing the Boot Device Order).
- Configure HBA and boot parameters (see Section 4.9, Configuring HBA and Boot Parameters).
- Set adapters to their default settings (see Section 4.10, Resetting Emulex Adapters to Their Default Values).
- Display adapter information (see Section 4.11, Displaying Adapter Information).
- Update legacy-only configuration settings (see Section 4.12, Legacy-Only Configuration Settings).
- Request reset or reconnect to make changes active (see Section 4.13, Requesting a Reset or Reconnect).
- Update the Emulex firmware (see Section 4.14, Emulex Firmware Update Utility).
- Set Emulex NVMe over FC boot settings (see Section 4.15, NVMe over FC Boot Settings).

NOTE: Changes made to parameters common to UEFI and x86 drivers are changed in both driver configuration utilities.
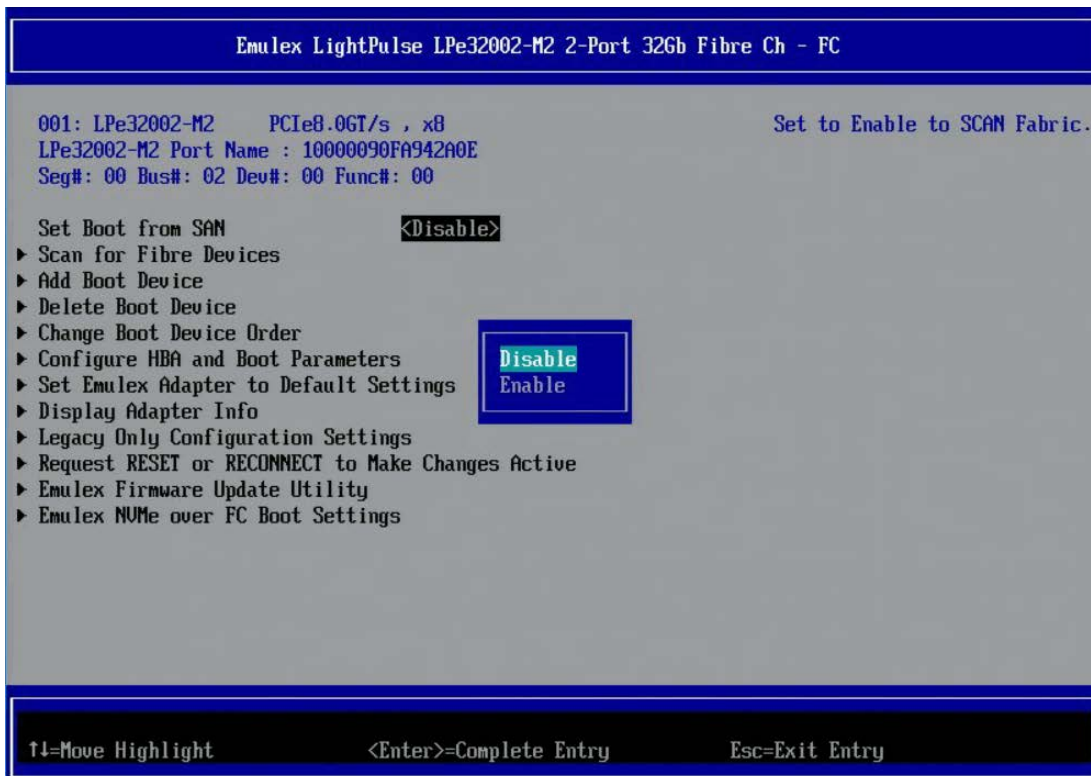
4.4 Setting Boot from SAN

NOTE: If FC-SP-2 authentication (DH-CHAP) is enabled in the HBA Manager application, BFS is not supported.

BFS must be enabled if you want to configure the adapter for FC or NVME over FC boot.

To set BFS, perform these steps:

1. From the adapter list (Figure 1), select the adapter whose BFS setting you want to change and press **Enter**.

2. From the main configuration menu (Figure 2), select **Set Boot from SAN** and press **Enter**.

   A popup (Figure 3) allows you to enable or disable BFS.

Figure 3: Boot from SAN Popup

3. Make your selection and press **Enter**.

   The new boot from SAN setting is displayed (Figure 4).

Figure 4: New Boot from SAN Setting
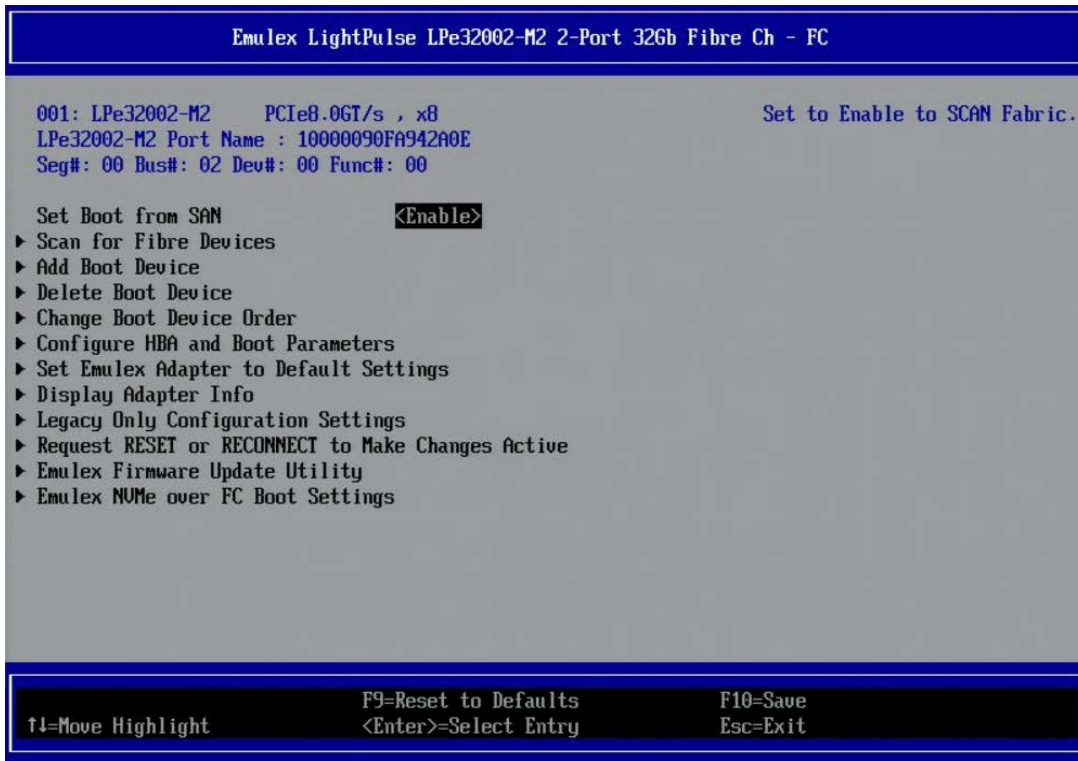
4. Save changes as instructed by the UEFI HII on your system, or press **Esc** to exit and save changes when prompted.

4.5 Scanning for Fibre Devices

NOTE: The following information pertains to scanning for Fibre Channel (Fibre) devices:

- After you enter the UEFI HII, any dynamic changes made to the SAN are not reflected unless the system is reset or the adapter is reconnected.
- For trunking SAN changes, the system must be reset. Re-enter the utility and perform a scan.
- The adapter link is brought up either by performing a scan for Fibre devices or by adding boot devices.

To scan for Fibre devices, perform these steps:

1. From the adapter list screen (Figure 1), select the adapter on which you want to scan for Fibre devices and press **Enter**.

2. From the main configuration menu (Figure 2), select **Scan for Fibre Devices** and press **Enter**.

   The **SAN Discovery Target List** window (Figure 5). This is only a list of discovered target devices to determine SAN connectivity.

Figure 5: SAN Discovery Target List Window
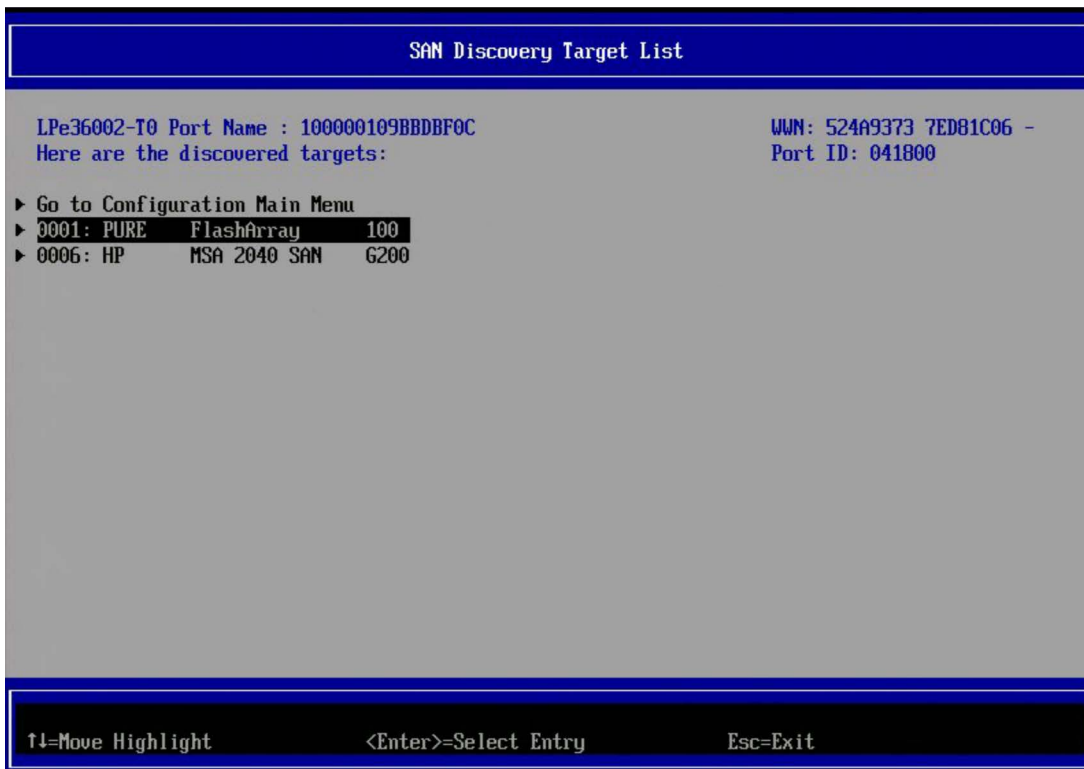
4.6 Adding Boot Devices

NOTE: After you enter the UEFI HII, any dynamic changes made to the SAN are not reflected unless the system is reset or the adapter is reconnected.

NOTE: The adapter link is brought up by either adding a boot device or by performing a scan for Fibre devices.
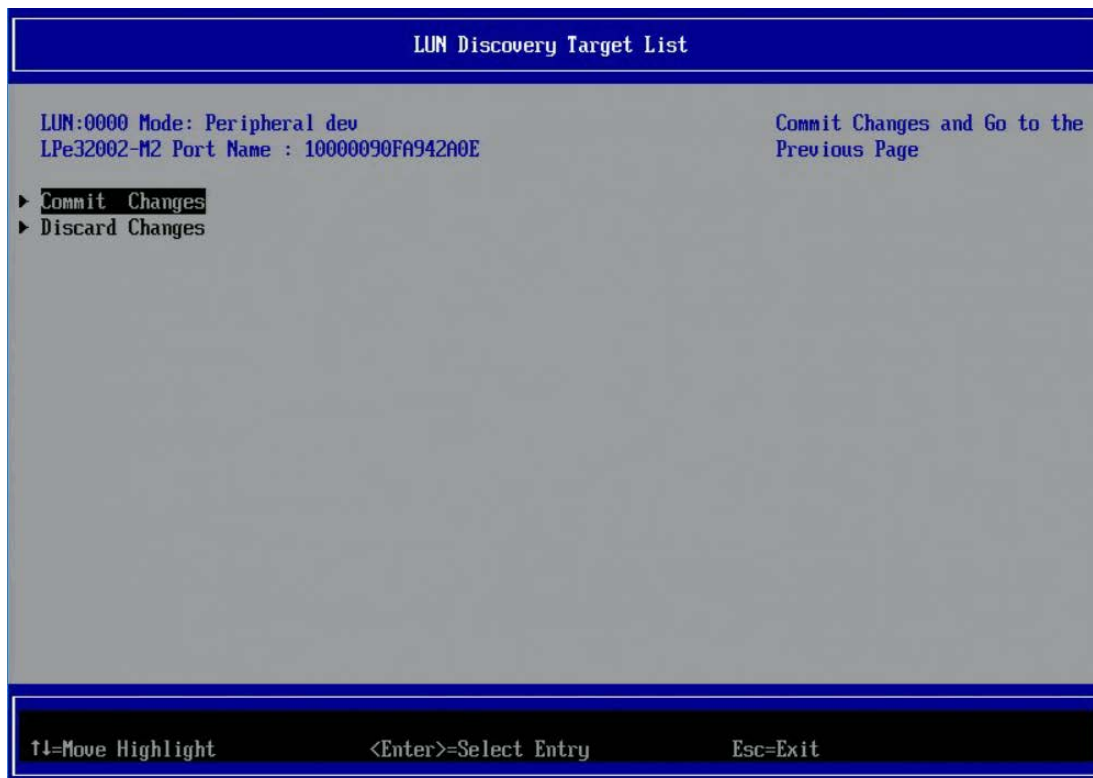
To add a boot device, perform these steps:

1. From the adapter list (Figure 1), select the adapter to which you want to add a boot device and press **Enter**.

2. From the main configuration menu (Figure 2), select **Add Boot Device** and press **Enter**.
    The **SAN Discovery Target List** window is displayed. (Figure 5).

3. Select the target you want and press **Enter**.
    A list of bootable LUNs is displayed (Figure 6)

Figure 6: Bootable LUNs

```
         LPe32002-M2 Port Name : 10000090FA942A0E

 WWN: 207000C0 FF1E206D                     HP       MSA 2040 SAN
 Here are the discovered LUNs:              G200

 ▶ LUN:0000 Mode: Peripheral dev
 ▶ LUN:0001 Mode: Peripheral dev
 ▶ LUN:0002 Mode: Peripheral dev
 ▶ LUN:0004 Mode: Peripheral dev




 ↑↓=Move Highlight          <Enter>=Select Entry        Esc=Exit
```

4. Select the boot device you want to add and press **Enter**.
    You are prompted to commit or discard your changes (Figure 7).

Figure 7: Commit and Discard Changes Prompt

5. Select **Commit Changes** and press **Enter**.

4.7 Deleting Boot Devices

To delete boot devices, perform these steps:

1. From the adapter list (Figure 1), select the adapter from which you want to delete a boot device and press **Enter**.

2. From the main configuration menu (Figure 2), select **Delete Boot Device** and press **Enter**.
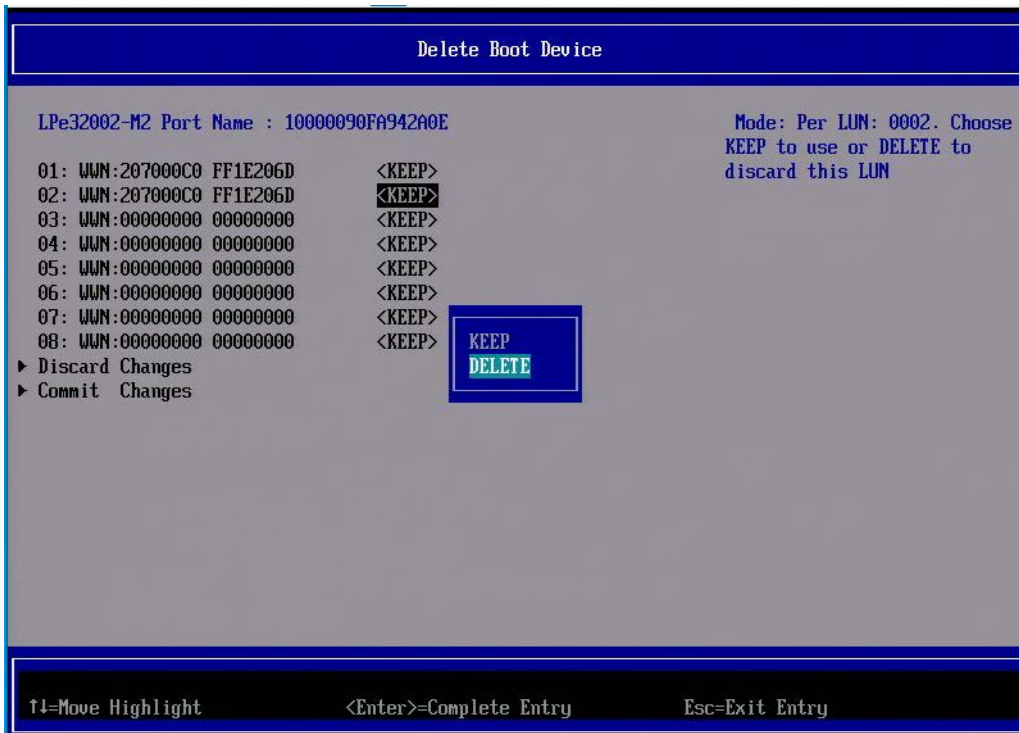   The **Delete Boot Device** window is displayed (Figure 8).

Figure 8: Delete Boot Device Window
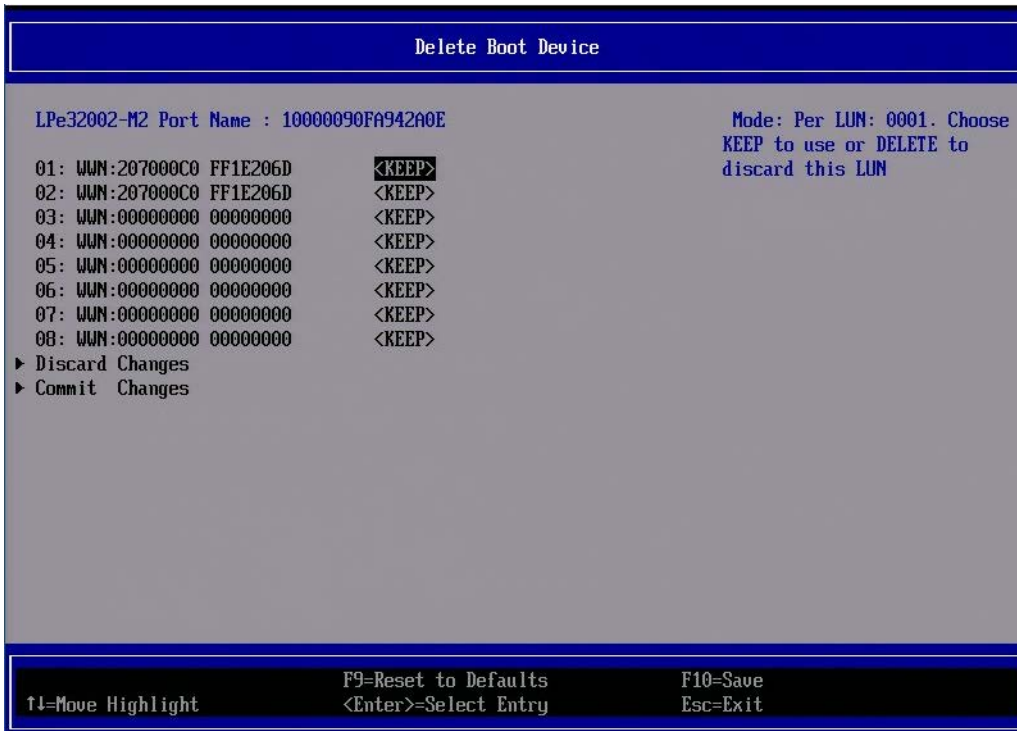
3. Select the boot device that you want to delete and press **Enter**.

Figure 9: Selected Boot Device: Highlighted

4. A popup with **<KEEP>** and **<DELETE>** is displayed as shown in Figure 9. Select **<DELETE>** and press **Enter**.

     NOTE: You can delete multiple boot devices.

5. Select **Commit Changes** and press **Enter**. The boot device is removed and the main configuration menu (Figure 2) is displayed.

4.8 Changing the Boot Device Order

NOTE: Changing the boot device order only changes the order in which the targets are discovered by the boot driver. The system BIOS controls the order that prioritizes the available boot devices in a system.

To change the boot device order, perform these steps:

1. From the adapter list screen (Figure 1), select the adapter whose boot device order you want to change and press **Enter**.

2. From the main configuration menu (Figure 2), select **Change Boot Device Order** and press **Enter**.
   The **Change Boot Device Order** window is displayed (Figure 10).

Figure 10: Change Boot Device Order Window
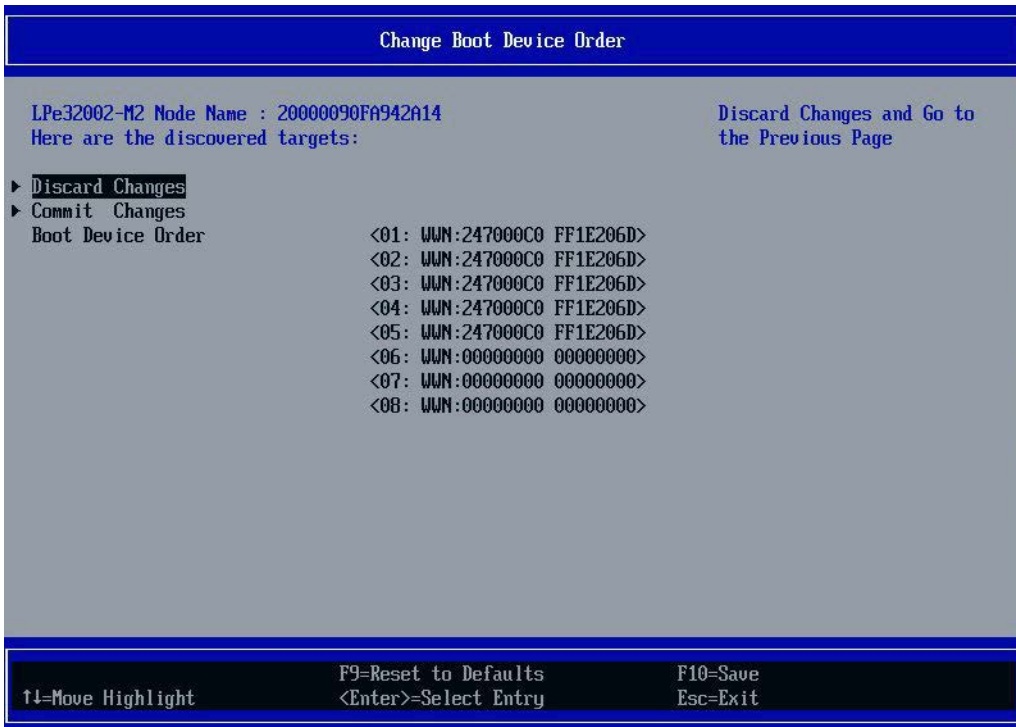
3. Select **Boot Device Order**.

    The current boot device order is displayed (Figure 11).

Figure 11: Boot Device Order

4. Press **Enter**.

   The cursor moves to the list of boot devices (Figure 12).

Figure 12: Boot Device Order Selected
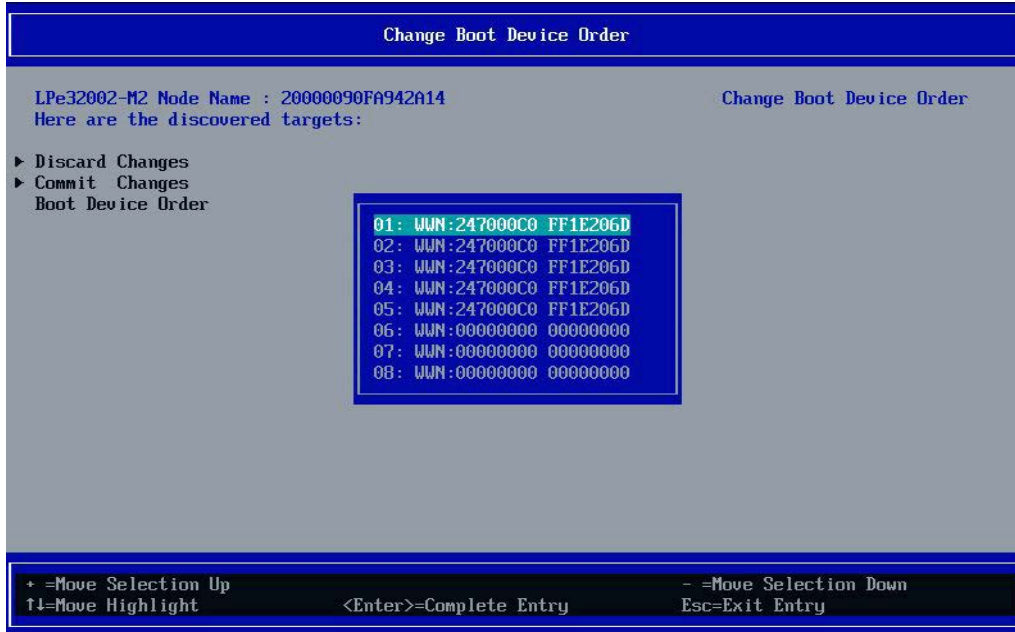
5. Select the device whose boot order you want to change. Use the **+** or **-** keys to change the order of the selected device and press **Enter**.

   The new boot device order is displayed.

6. Select **Commit Changes** and press **Enter** to confirm the changes.

4.9 Configuring HBA and Boot Parameters

The UEFI HII enables you to configure the following HBA parameters:

Parameter	Note
Topology	—
PLOGI retry timer	—
Force link speed	—
Maximum LUNs per target	—
Boot target scan method	—
Enable or disable Brocade FA-PWWN	FA-PWWN is no longer supported; the feature will be disabled in a future release.
Enable or disable Brocade Boot LUN	Brocade Boot LUN is no longer supported, the feature will be disabled in a future release.
Enable or disable 16G Forward Error Correction	—
Trunking mode	Trunking is supported only on LPe35000-series and LPe36000-series adapters. Before you configure trunking on the Emulex adapter, follow the instructions from Brocade for configuring trunking on the switch.

Figure 13: Configure HBA and Boot Parameters Window (LPe31000-Series and LPe32000-Series HBAs)
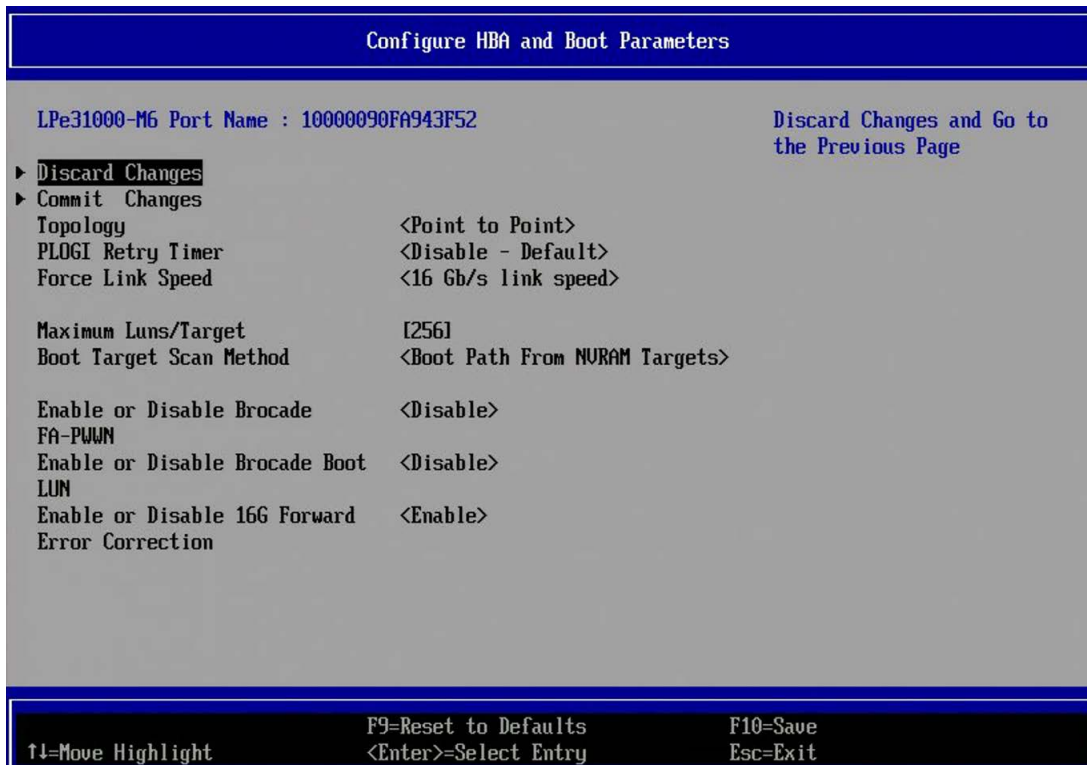
Figure 14: Configure HBA and Boot Parameters Window (LPe35000-Series and LPe36000-Series HBAs)
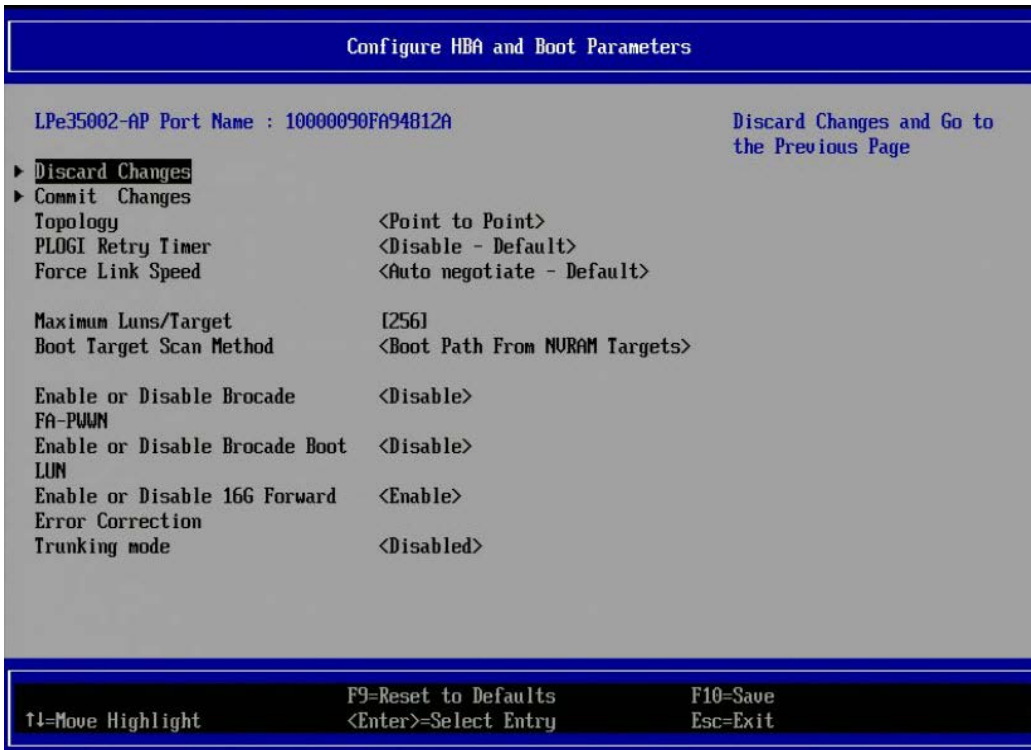
4.9.1 Changing the Topology

The available topology options are shown in Figure 15.

NOTE: FC-AL is private loop only (direct-attach).

- Point-to-point
- FC-AL (Not supported in 64G LPe36000-series adapters)

To select the adapter topology, perform these steps:

1. From the adapter list (Figure 1), select the adapter whose topology you want to change and press **Enter**.

2. From the main configuration menu (Figure 2), select **Configure HBA and Boot Parameters** and press **Enter**.

3. In the **Configure HBA and Boot Parameters** window, navigate to **Topology Selection** and press **Enter**.
   Information similar to Figure 15 is displayed.

Figure 15: Topology Popup
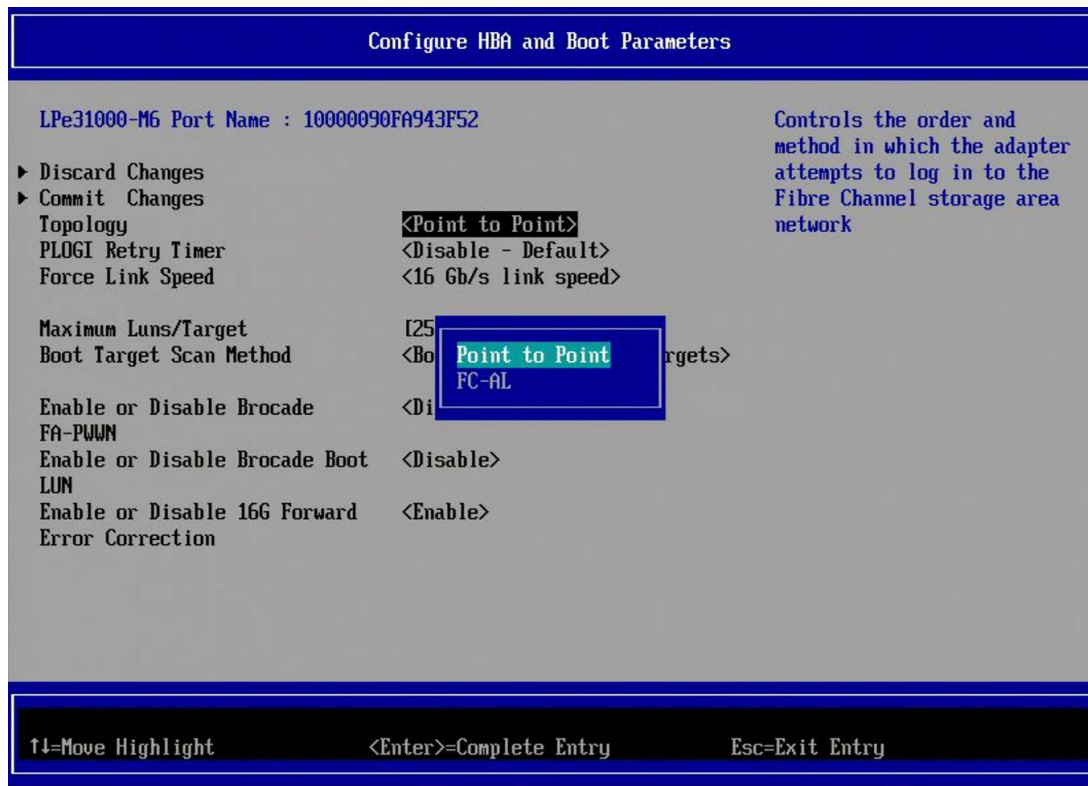
4. Select a topology and press **Enter**.
   The window is refreshed with the new value.

NOTE: The presence of a fabric is detected automatically in the case of point-to-point topology.

5. Select **Commit Changes** and press **Enter**.

4.9.2 Changing the PLOGI Retry Timer

This option allows you to set the interval for the Port Login retry timer. This option is especially useful for Tachyon-based RAID arrays. Under rare occasions, a Tachyon-based RAID array resets itself, and the port goes offline temporarily in the loop. When the port comes to life, the Port Login retry interval scans the loop to discover this device.

You can choose one of these options:

- Disable – default
- 50 ms
- 100 ms
- 200 ms

To change the timer value, perform these steps:

1. From the adapter list (Figure 1), select the adapter whose Port Login retry timer information you want to change and press **Enter**.

2. From the main configuration menu (Figure 2), select **Configure HBA and Boot Parameters** and press **Enter**.

3. In the **Configure HBA and Boot Parameters** window, navigate to **PLOGI Retry Timer** and press **Enter**.
   The **PLOGI Retry Timer** popup appears (Figure 16).

Figure 16: PLOGI Retry Timer Popup
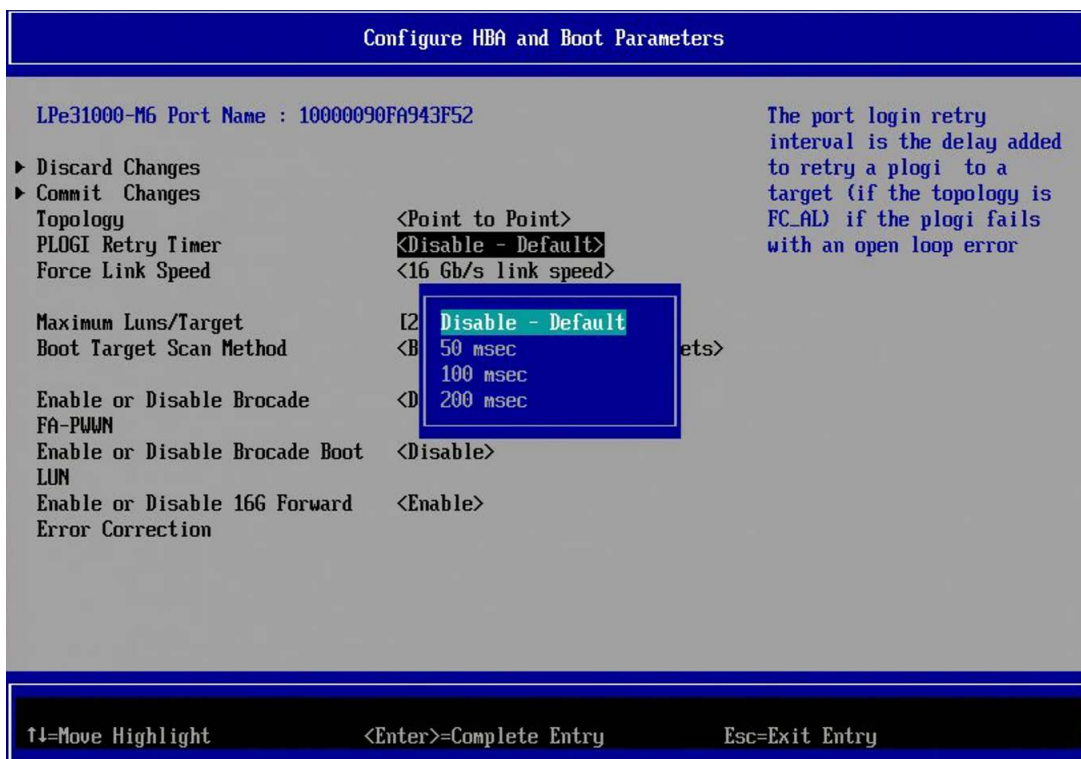
4. Select a retry timer value and press **Enter**.
   The window is refreshed with the new value.

5. Select **Commit Changes** and press **Enter**.

4.9.3 Changing the Link Speed

Use this feature to change, or force, the link speed between ports instead of autonegotiating it. Supported link speeds are specific to the adapter model and the optics installed. The menu displays only the options that are valid for the selected adapter.

Possible link speed choices include the following (depending upon your adapter model):
- Autonegotiate – default
- 8Gb/s
- 16Gb/s
- 32Gb/s
- 64Gb/s

To change the link speed, perform these steps:

1. From the adapter list (Figure 1), select the adapter whose link speed you want to change and press **Enter**.

2. From the main configuration menu (Figure 2), select **Configure HBA and Boot Parameters** and press **Enter**.

3. In the **Configure HBA and Boot Parameters** window, navigate to **Force Link Speed** and press **Enter**.

   The **Force Link Speed** popup appears (Figure 17).

Figure 17: Force Link Speed Popup
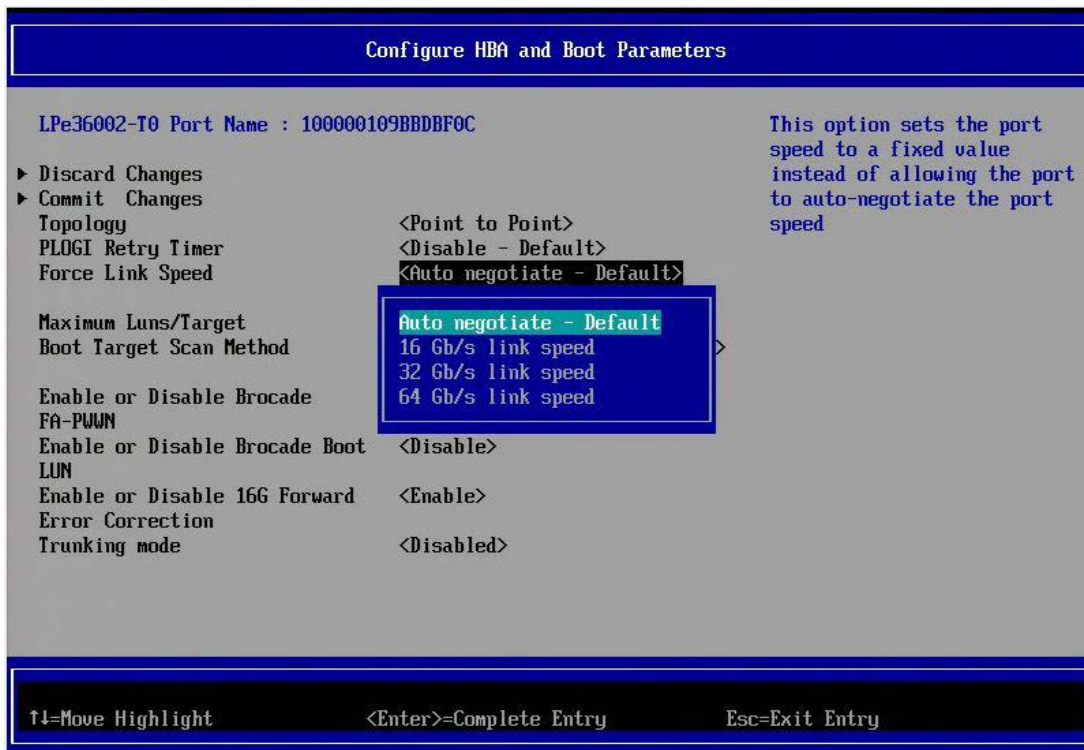
4. Select a link speed and press **Enter**.

   The screen is refreshed with the new value.

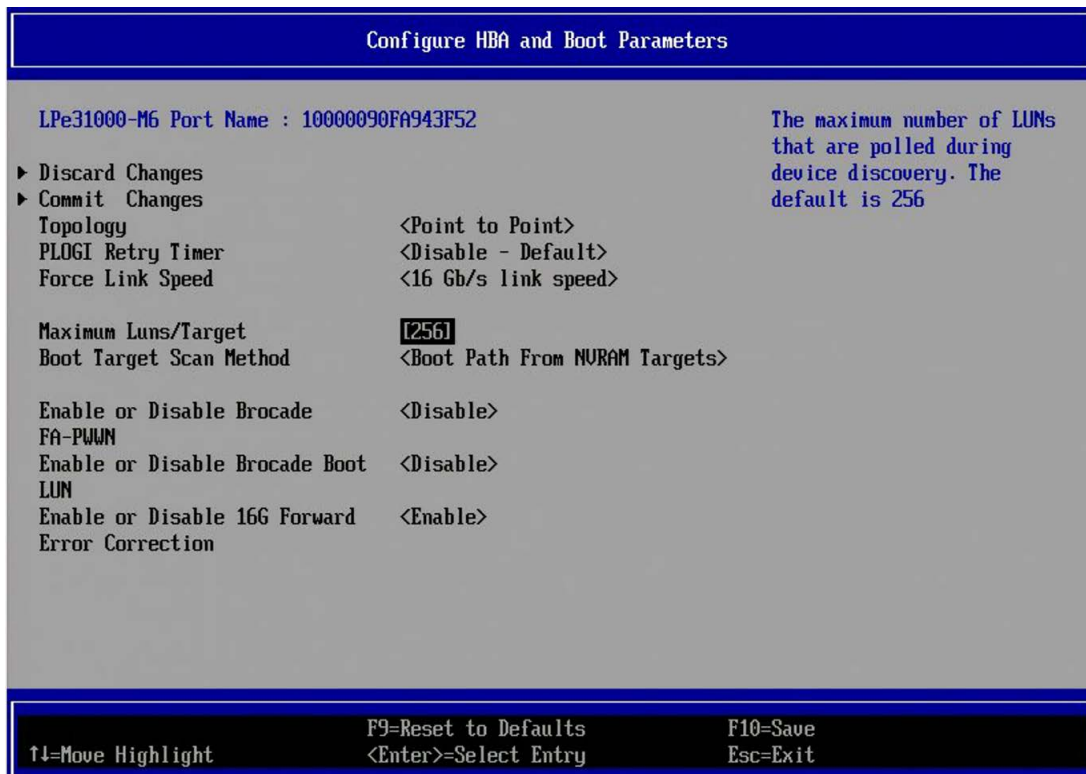5. Select **Commit Changes** and press **Enter**.

4.9.4 Changing the Maximum Number of LUNs per Target

The maximum number of LUNs represents the maximum number of LUNs that are polled during device discovery. You can see the LUNs listed under the **Add Boot Device** menu. The minimum value is 1, and the maximum value is 4096. The default is 256.

To change the maximum number of LUNs, perform these steps:

1. From the adapter list (Figure 1), select the adapter whose maximum LUNs per target information you want to change and press **Enter**.

2. From the main configuration menu (Figure 2), select **Configure HBA and Boot Parameters** and press **Enter**.

3. In the **Configure HBA and Boot Parameters** window, navigate to **Maximum LUNs/Target** and press **Enter**.

4. Type the maximum LUN value (between 1 and 4096) and press **Enter** (Figure 18).

   The screen is refreshed with the new value.

Figure 18: Maximum LUNs/Target Entry Window

NOTE: The default value is 256, which is typically the maximum number of LUNs in a target device. A larger number of LUNs causes discovery to take more time.

5. Select **Commit Changes** and press **Enter**.

4.9.5 Changing the Boot Target Scan Method

You can configure up to eight boot entries for fabric point-to-point or private loop configurations.

You can select one of the following boot scan options:

- **Boot Path from NVRAM Targets** – Boot scan discovers only LUNs that are saved to the adapter's NVRAM. Select up to eight attached devices to use as potential boot devices. Limiting discovery to a set of eight selected targets can greatly reduce the time it takes for the UEFI driver to complete discovery.
- **Boot Path from Discovered Targets** – Boot scan discovers all devices that are attached to the FC port. Discovery can take a long time on large SANs if this option is selected.
- **Do Not Create a Boot Path**.
- Boot Scan from **EFI FC Scan Level** – Allows third-party software to toggle between boot path from NVRAM and boot path from Discovered Targets by manipulating an UEFI system NVRAM variable. After the scan is set to EFIFCScanLevel, the scan method can be changed without entering the UEFI boot configuration utility. **EFI FC Scan Level: NVRAM Targets** or **Discovered Targets** – Sets the boot path from either NVRAM targets or discovered targets, and subsequently allows third-party software to toggle between NVRAM targets and discovered targets by manipulating the value of the `EFIFCScanLevel` variable maintained by the UEFI system firmware or boot code.
  Thus, the scan method can be changed without entering the UEFI HII.

  If `EFIFCScanLevel` is selected, the scan is determined by the value of the `EFIFCScanLevel` variable maintained by the UEFI system firmware or boot code. The value of this variable can be changed either by using the menu in the UEFI boot configuration utility or by using third-party software.

To change the boot target scan method, perform these steps:

1. From the adapter list (Figure 1), select the adapter whose boot target scan method you want to change and press **Enter**.

2. From the main configuration menu (Figure 2), select **Configure HBA and Boot Parameters** and press **Enter**.

3. In the **Configure HBA and Boot Parameters** window, navigate to **Boot Target Scan Method** and press **Enter**.

   The **Boot Target Scan Method** popup appears (Figure 19).

Figure 19: Boot Target Scan Method Menu Popup
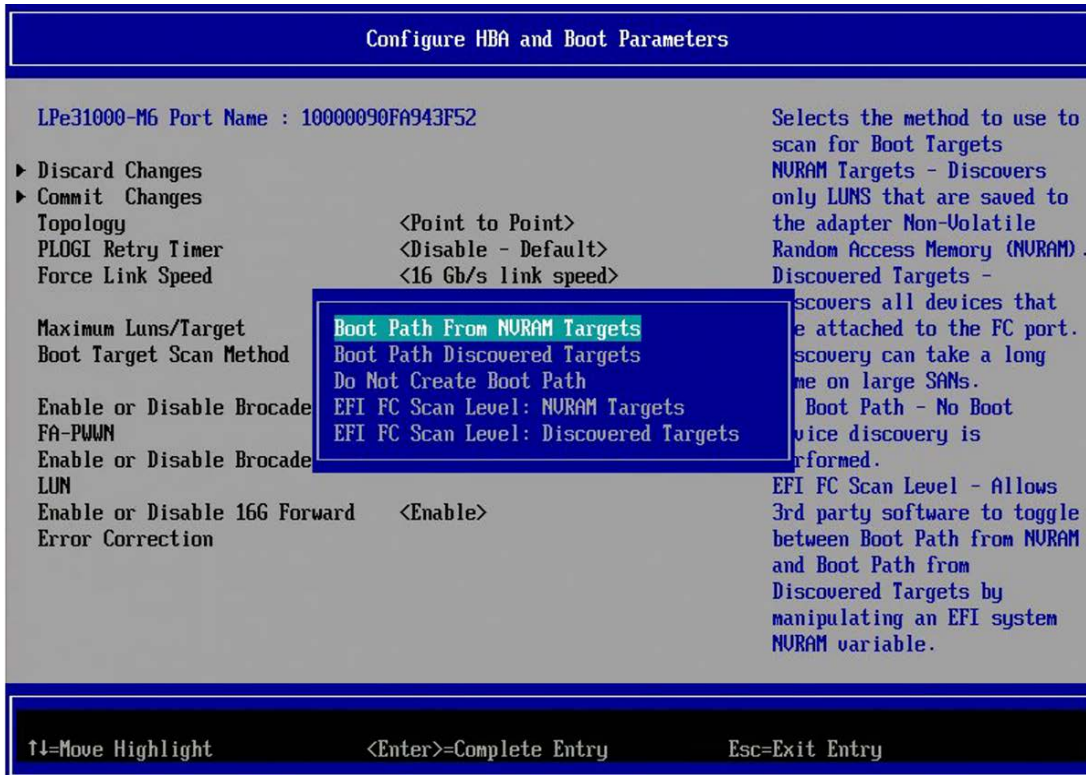
4. Select a boot scan setting and press **Enter**.

    The screen is refreshed with the new value. If you have a large SAN and set the boot path to **Boot Path Discovered Targets**, discovery takes a long time.

5. Select **Commit Changes** and press **Enter**.

4.9.6 Configuring Brocade FA-PWWN

NOTE: FA-PWWN is no longer supported; the feature will be disabled in a future release.

4.9.7 Configuring Brocade Boot LUN

NOTE: Brocade boot LUN is no longer supported, the feature will be disabled in a future release.

4.9.8 Configuring 16G Forward Error Correction

FEC improves the reliability of FC links by recovering from bit errors that occur in high-speed networks. FEC helps minimize or avoid data-stream errors that can lead to application performance degradation or outages.

This option enables or disables FEC for 16Gb/s link speed. For link speeds 32Gb/s and higher, FEC is always enabled. The default setting is enabled.

To enable or disable the FEC, perform these steps:

1. From the adapter list (Figure 1), select the adapter that you want to enable or disable and press **Enter**.

2. From the main configuration menu (Figure 2), select **Configure HBA and Boot Parameters** and press **Enter**.

3. In the **Configure HBA and Boot Parameters** window (Figure 13 or Figure 14), navigate to **Enable or Disable 16G Forward Error Correction**. Press **Enter**.

    The FEC popup appears (Figure 20).

Figure 20: FEC Popup
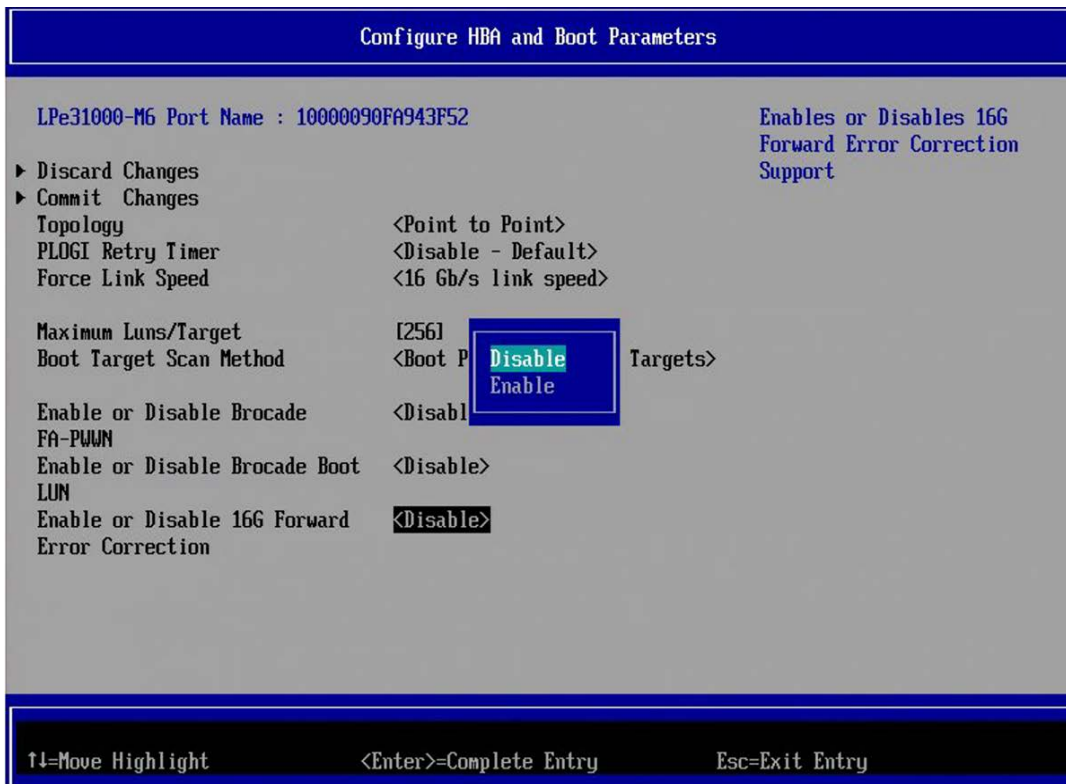
4. Make your selection and press **Enter**.

    The configuration is changed, and the new setting is displayed.

5. Select **Commit Changes** and press **Enter**.

6. Reboot the system.

4.9.9 Selecting Trunking

Trunking enables you to combine multiple physical FC links to form a single logical link (aggregated port). The aggregated port's maximum link speed is the sum of the maximum link speeds of the individual physical links comprising the aggregated port. The actual link speed of the aggregated port depends on the states (active or non-active) of the individual physical links comprising the aggregated port.

The physical links comprising an aggregated port are referred to as lanes. Both 2-lane and 4-lane aggregated ports are supported.

The LPe35002 and LPe36002 adapters support 2-lane port aggregation. If 2-lane port aggregation is configured on a dual-port adapter, the two physical links are combined to form a single 2-lane aggregated port whose aggregate speed is two times the maximum speed supported, assuming both physical links are active.

For example:
- On an LPe35002-series adapter, an aggregated port comprised of two physical links running at 32Gb/s each has a potential logical (aggregate) link speed of 64Gb/s.
- On an LPe36002-series adapter, an aggregated port comprised of two physical links running at 64Gb/s each has a potential logical (aggregate) link speed of 128Gb/s.

An LPe35004 adapter supports both 2-lane port aggregation and 4-lane port aggregation as follows:
- If 2-lane port aggregation is configured on an LPe35004 adapter, the four physical links on the adapter are divided among two separate aggregated ports. The two lowest-numbered physical links form the first aggregated port, and the two highest-numbered physical links form the second aggregated port.
- If 4-lane port aggregation is configured on an LPe35004 adapter, all four physical links are combined to form a single 4-lane trunk whose aggregate speed is potentially 128Gb/s, assuming all four links are active, each at 32Gb/s link speed.

NOTE: The following information pertains to trunking:

- Trunking (also referred to as FC port aggregation) is supported only on LPe35000-series and LPe36000-series adapters.

- Trunking is not supported at 8Gb/s speeds, and the link will not come up at this speed.

- Dynamic D_Port cannot coexist with the trunking feature. If trunking is enabled, the firmware automatically disables Dynamic D_Port.

If you are using an adapter that does not support trunking, the Select Trunking option is not available.

This option enables or disables trunking. The default setting is disabled.

NOTE: Changing the trunking mode resets the adapter to factory default settings. You must configure the HBA and boot parameters after changing the trunking mode for BFS to work.

To enable or disable trunking, perform these steps:

NOTE: Before you enable or disable trunking on the Emulex adapter, follow the instructions from Brocade to enable or disable trunking on the corresponding ports in the switch.

1. From the adapter list (Figure 1), select the adapter for which you want to enable or disable trunking and press **Enter**.

2. From the main configuration menu (Figure 2), select **Configure HBA and Boot Parameters** and press **Enter**.

3. In the **Configure HBA and Boot Parameters** window, navigate to **Trunking mode**. Press **Enter**.

4. The **Trunking** popup appears.

   In LPe35002 adapters, the options available for selection are as follows (Figure 21).
   – Disabled
   – 2-Lane

   In LPe35004 adapters, the options available for selection are as follows:
   – Disabled
   – 2-Lane
   – 4-Lane

Figure 21: Trunking Selection Menu in LPe35002 Adapters
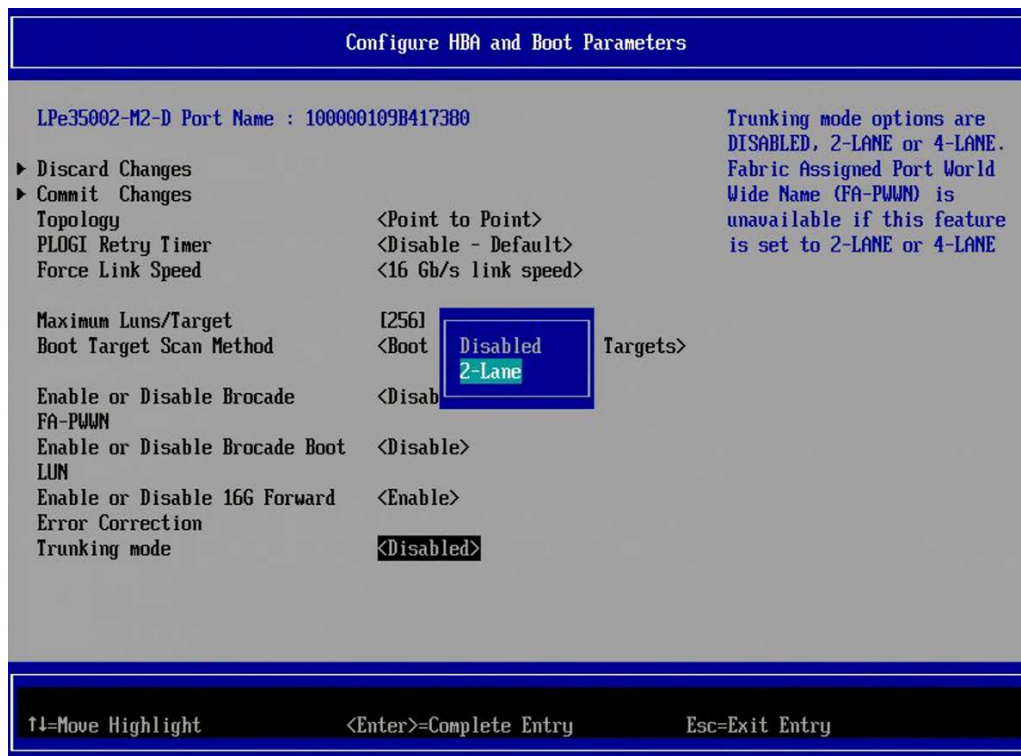
5. Select the trunking mode from the available options.

   Disable mode disables trunking for the adapter.

6. Make your selection and press **Enter**.

   The new setting is displayed.

7. Select **Commit Changes** and press **Enter**.

8. Reboot the system. If trunking has been enabled, only the trunked ports are visible in the main configuration menu. If trunking has been disabled, all the physical ports are visible in the main configuration menu.

4.9.9.1 Viewing the Trunking Configuration

NOTE: When trunking is enabled, BFS must be enabled and the system must be rebooted for the physical port status to be up and for the link speed to be shown.

Select the main configuration menu (Figure 2). The link speed and physical link status are shown under **Trunking mode** (Figure 22).

Figure 22: Trunking Enabled for LPe35002 Adapters
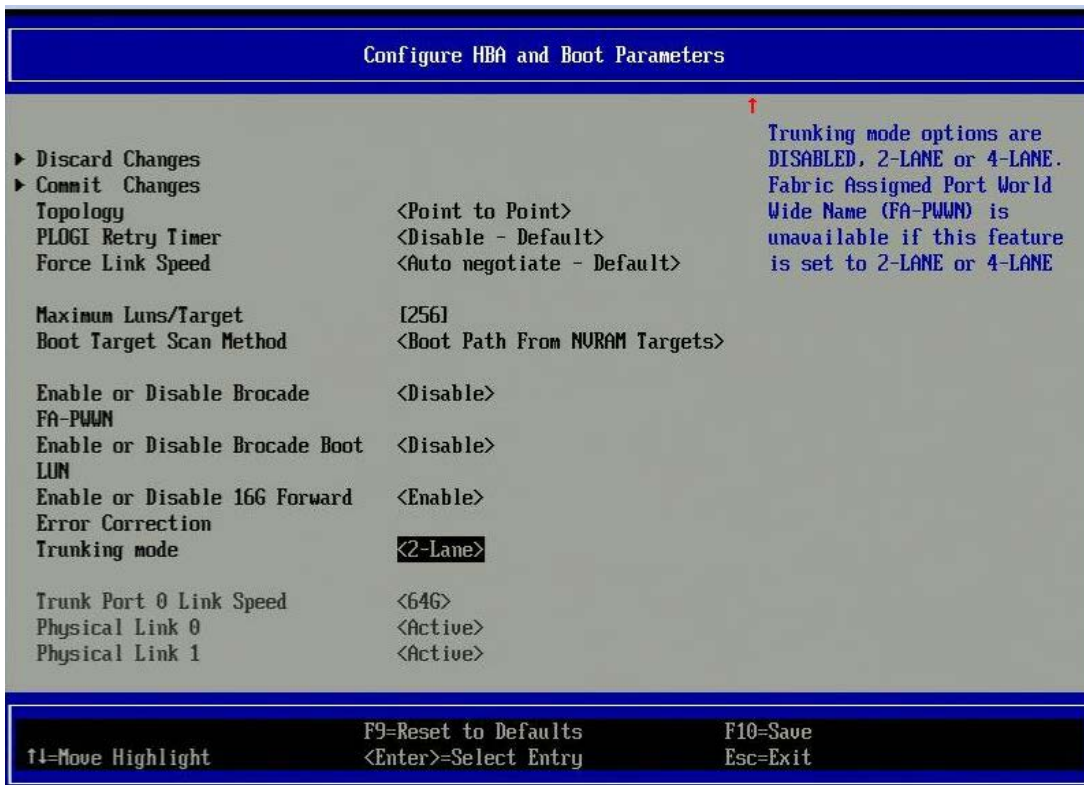
4.10 Resetting Emulex Adapters to Their Default Values

NOTE: The following information pertains to resetting default values:

- The reconnect feature is not applicable while selecting trunking, as the procedure requires a system reboot.
- Resetting the adapter to its default settings clears all entries made while configuring the boot device, as detailed in Section 4.5, Configuring Boot Devices.
- A system reboot is required after resetting an adapter to its default values.

Resetting the system defaults also clears the NVRAM target list and sets all FC and NVMe boot devices to 0.

Default settings are listed in Table 4.

Table 4: Boot Code Defaults

Parameter	Default	Valid Values
Boot from SAN	Disabled	■ Enabled ■ Disabled
Topology	Point-to-point (for LPe31000-series, LPe32000-series, LPe35000-series, and LPe36000-series adapters)	■ FC-AL (supported only on 8Gb/s adapters) ■ Point-to-Point
PLOGI Retry Timer	Disabled	■ Disabled ■ 50 ms ■ 100 ms ■ 200 ms
Link Speed	(Auto negotiate)	■ Auto negotiate (the adapter speed is selected automatically based on its model). ■ 8Gb/s ■ 16Gb/s ■ 32Gb/s ■ 64Gb/s **NOTE:** Adapters with a 32G SFP support link speeds 32Gb/s,16Gb/s, and 8Gb/s. Adapters with a 64G SFP support link speeds 64Gb/s, 32Gb/s, and 16Gb/s.
Max LUNs Setting	256	0 to 4096
Boot Target Scan Method	Boot path from NVRAM targets	■ Boot Path from NVRAM Targets ■ Boot Path from Discovered Targets ■ Do Not Create a Boot Path ■ EFI FC Scan Level NVRAM Targets ■ EFI FC Scan Level Discovered Targets
Brocade FA-PWWN	**NOTE:** FA-PWWN is no longer supported; the feature will be disabled in a future release.	—
Brocade Boot LUN	**NOTE:** Brocade boot LUN is no longer supported, the feature will be disabled in a future release.	—
16G Forward Error Correction	Enabled	■ Enabled ■ Disabled

Table 4: Boot Code Defaults (Continued)

Parameter	Default	Valid Values
Emulex NVMe Boot Settings	Disabled	■ Enable or Disable NVMe over FC boot - Enabled - Disabled ■ Generate and Use Port Based Host NQN - Enabled - Disabled
Trunking	Disabled	■ 4-Lane Trunk ■ 2-Lane Trunk ■ Disabled

To set Emulex adapters to their default settings, perform these steps:

NOTE: The Emulex UEFI configuration utility does not immediately display the adapter's default values after they have been set. When you reset or reconnect the adapter. The Emulex UEFI configuration utility displays the new values.

1. From the adapter list (Figure 1), select the adapter whose default settings you want to change and press **Enter**.

2. From the main configuration menu (Figure 2), select **Set Emulex Adapters to Default Settings** and press **Enter.**
   The **Set Emulex Adapter to Default Settings** window appears (Figure 23).

3. Select **Set Adapter Defaults** and press **Enter**.

Figure 23: Setting Adapter Defaults Window

4. Perform a system reboot after resetting the adapter to its default values.

4.11 Displaying Adapter Information

The **Adapter Information** window displays the following information about the selected adapter:

■ PCI information for the HBA

■ HBA revision ID

■ Link status

■ BFS status

■ Link speed

■ Topology

■ Firmware version

■ UEFI Boot version

■ Trunking mode (LPe35000-series and LPe36000-series adapters only)

To display adapter information, perform these steps:

1. From the adapter list (Figure 1), select the adapter whose information you want to view and press **Enter**.

2. From the main configuration menu (Figure 2), select **Display Adapter Info** and press **Enter**.

   The Controller Information window displays information about the selected adapter (Figure 24).

NOTE: Only Dell adapters show the firmware family version.

Figure 24: Controller Information Window
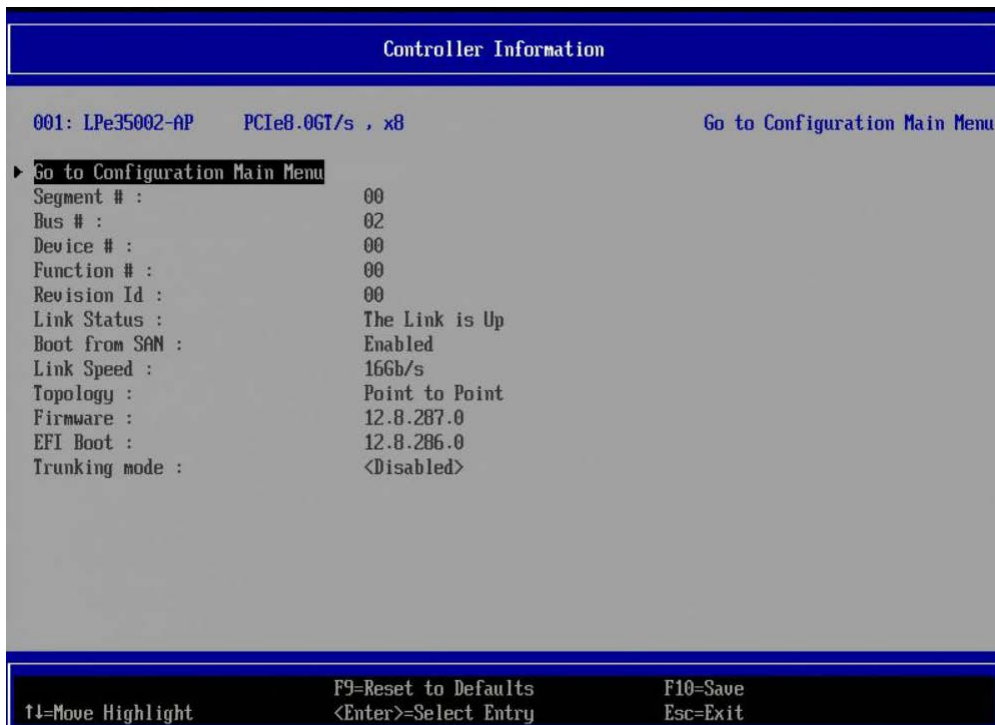
NOTE: When trunking is enabled, the maximum aggregate link speed is always displayed regardless of the state of each physical link.

4.12 Legacy-Only Configuration Settings

Using the UEFI HII, you can set parameters that are specific to the x86 boot driver, which is also referred to as the Legacy boot driver.

To change legacy-only configuration settings, perform these steps:

1. From the adapter list (Figure 1), select the adapter whose legacy settings you want to change and press **Enter**.

2. From the main configuration menu (Figure 2), select **Legacy Only Configuration Settings** and press **Enter**.

   The **Legacy Only Configuration Settings** window displays settings that apply only to the Emulex legacy boot driver (Figure 25).

Figure 25: Legacy Only Configuration Settings Window
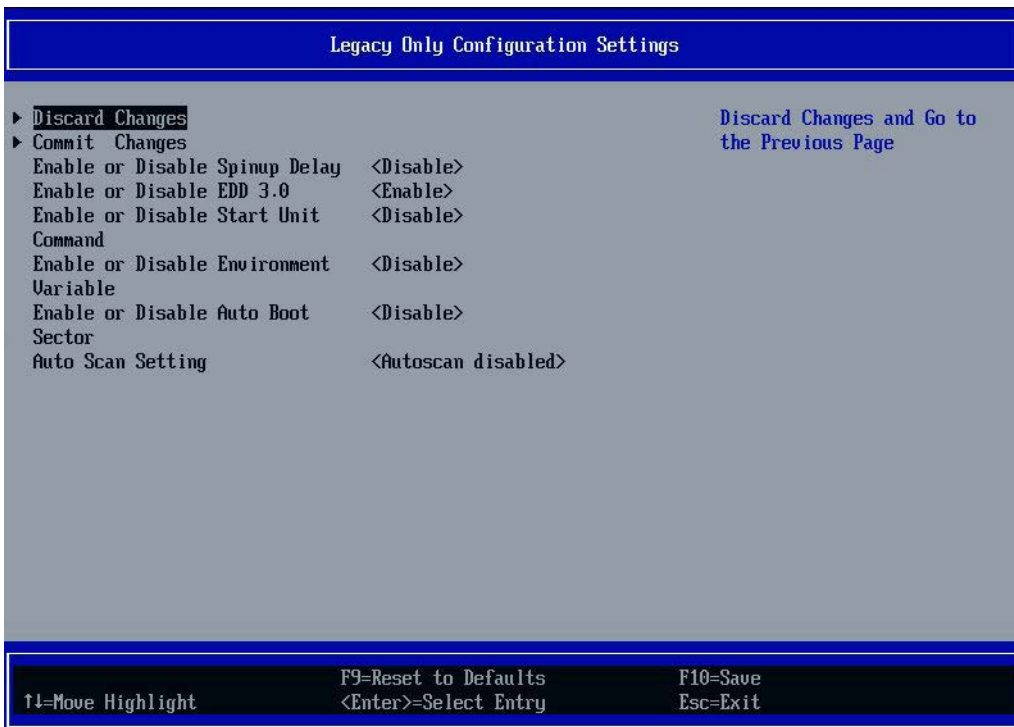
3. Make the desired changes to any of these settings and select **Commit Changes**.

See Chapter 4, Emulex FC BIOS Utility, for a description of the legacy-only settings displayed in Figure 25.

4.13 Requesting a Reset or Reconnect

After you have made changes to the UEFI HII, you might need to reset or reconnect the adapter for the changes to take effect. Enabling the **Request RESET or RECONNECT to Make Change Active** option causes a reset or reconnect to occur automatically after you exit the UEFI HII. This option is disabled by default, and you must enable it if you want the adapter to automatically reset or reconnect. Otherwise, the changes do not take effect until you manually reboot the server.

Reset forces the system to reboot after you exit the UEFI HII.

Reconnect unloads and reloads the driver. This allows any changes to take effect immediately without requiring a reboot. The Reconnect feature also requires a UEFI 2.5-compliant BIOS to enable this functionality. If the server BIOS is not UEFI 2.5-compliant, this option is unavailable.

NOTE: The Reconnect feature is not applicable while configuring 16G FEC and selecting trunking, as all these procedures require a system reboot.

To request a reset or reconnect, perform these steps:

1. From the adapter list (Figure 1), select the adapter that you want to reset or reconnect and press **Enter**.

2. From the main configuration menu (Figure 2), select **Request RESET or RECONNECT to Make Changes Active**.

   The **Request RESET or RECONNECT to Make Changes Active** window appears (Figure 26).

Figure 26: Request RESET or RECONNECT to Make Changes Active Window – Reset

Figure 27: Request RESET or RECONNECT to Make Changes Active Window – Reconnect
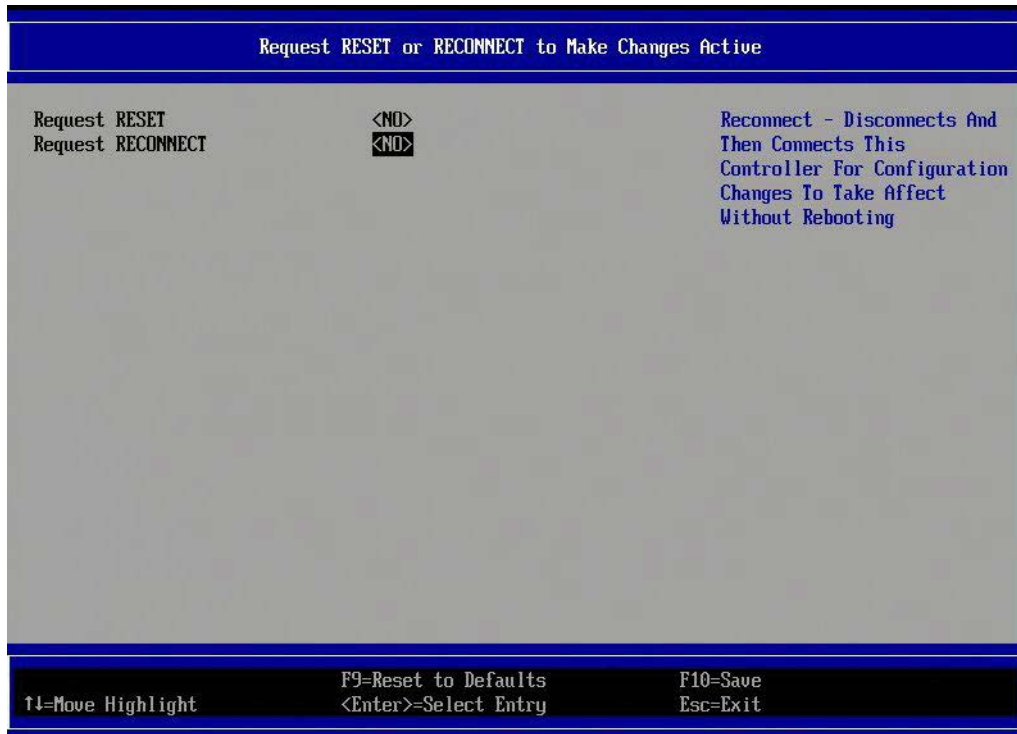
3. Do one of the following:

    To request a reset:

    a. Set **Request RESET** to **YES** to notify the system BIOS that you want to perform a system reset.

    b. Exit the UEFI HII and return to the system BIOS set up screens.

    A popup window prompts you to reset the system. Any changed Emulex adapter settings become active during the subsequent reboot.

    To request a reconnect:

    Set **Request RECONNECT** to **YES** to notify the system BIOS that you want to perform a disconnect and reconnect on the adapter port.

    The changed Emulex adapter settings are activated without a system reset.

4.14 Emulex Firmware Update Utility

The Emulex UEFI boot driver supports firmware updates by implementing the UEFI Firmware Management Protocol.

NOTE: If a secure version of firmware (version 11.0 or later) is installed on an LPe31000-series or LPe32000-series adapter and you want to update to an earlier unsecure version of firmware, you must remove the secure firmware jumper block before performing the update. Refer to the *Emulex Fibre Channel Host Bus Adapters Installation Guide* for more information.

To update the firmware on an adapter, perform the following steps:

1. From the adapter list (Figure 1), select the adapter whose firmware you want to update and press **Enter**.

2. From the main configuration menu (Figure 2), select **Emulex Firmware Update Utility**. Press **Enter**.
   The **Select Device** window appears (Figure 28).

Figure 28: Select Device Window

3. A list of devices appears. Select the storage device that contains the firmware file. The device can be a USB flash drive or a USB image file with the firmware file.

4. After you select the storage device, press **Enter**.
   The **Select File or Folder** window appears.

Figure 29: Select File or Folder Window

5. Navigate to the folder that contains the firmware image.

6. Select the firmware file and press **Enter**.

    The **Start Firmware Update** window appears (Figure 30).

Figure 30: Start Firmware Update Window
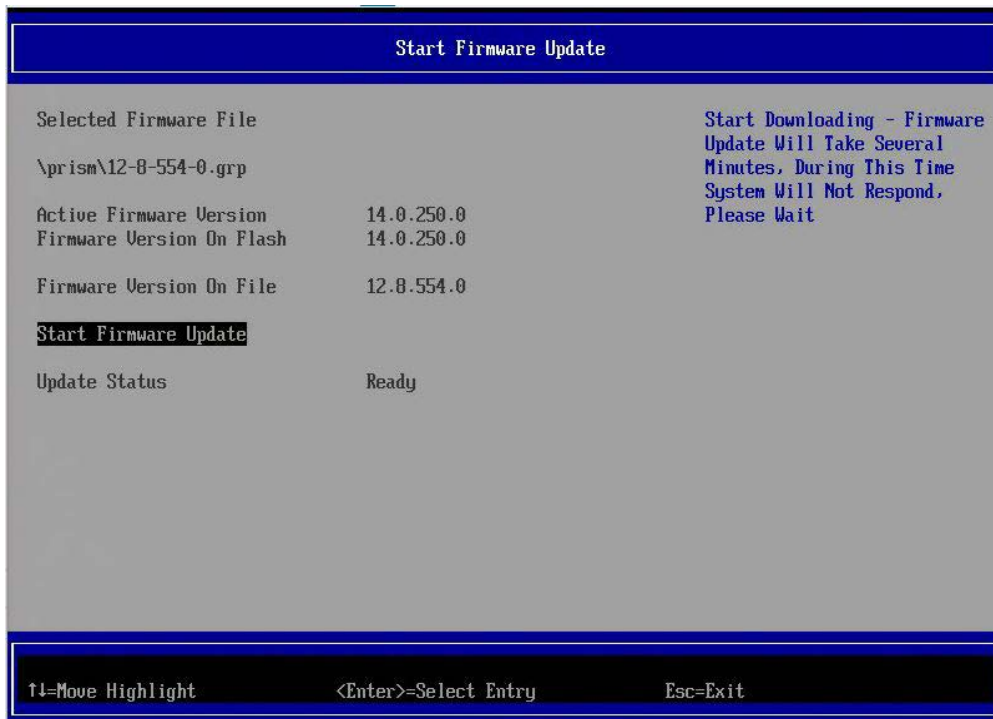
7. Select **Start Firmware Update** and press **Enter**. On systems that support UEFI version 2.7 or later, a popup with a message similar to the following appears:

   The Firmware update may take a while to complete, During this time the system will not respond.

   Click **OK** or press **Enter** to acknowledge the popup, the firmware update will begin. When the operation is complete, the **Update Status** indicates the completion status.

NOTE: In case the firmware image is not compatible, an error message similar to the following is displayed:
   Firmware image incompatible with flash part is displayed.

4.15 NVMe over FC Boot Settings

Enabling NVMe over FC BFS allows the driver to scan for both NVMe and FCP devices. The limit for the total number of devices scanned on the Fabric is 256 for FC and 256 for NVMe. When NVMe over FC is enabled, the UEFI driver scans for NVMe devices in addition to FCP devices. The UEFI driver runs in either FCP mode or FCP and NVMe mode. Running NVMe mode only is not supported.

NOTE: NVMe support is not available on the Arm aarch64 UEFI driver.

4.15.1 Enabling NVMe over FC BFS

NOTE: The following information pertains to NVMe over FC BFS:

- Fabric and Direct-connect point-to-point topology is supported for NVMe over FC BFS.
- All namespaces require a unique identifier (like a UUID, GUID, or EUI64) configured on the namespace for the UEFI HII to be able to discover and configure the namespace.
- When configuring the NVMe target system, you can use the initiator NQN that is displayed in the **Emulex NVMe over FC Boot Settings** window (Figure 31).
- BFS must be enabled and saved, and a reboot or reconnect is required to make the changes active before making other changes.

The default for NVMe over FC is disabled.

To configure NVMe over FC BFS, perform the following steps:

1. From the adapter list (Figure 1), select the adapter port for which you want to configure NVMe over FC BFS and press **Enter**.

2. From the main configuration menu (Figure 2), select **Emulex NVMe over FC Boot Settings**. Press **Enter**.

   The **Emulex NVMe over FC Boot Settings** window appears (Figure 31).

Figure 31: Emulex NVMe over FC Boot Settings Window
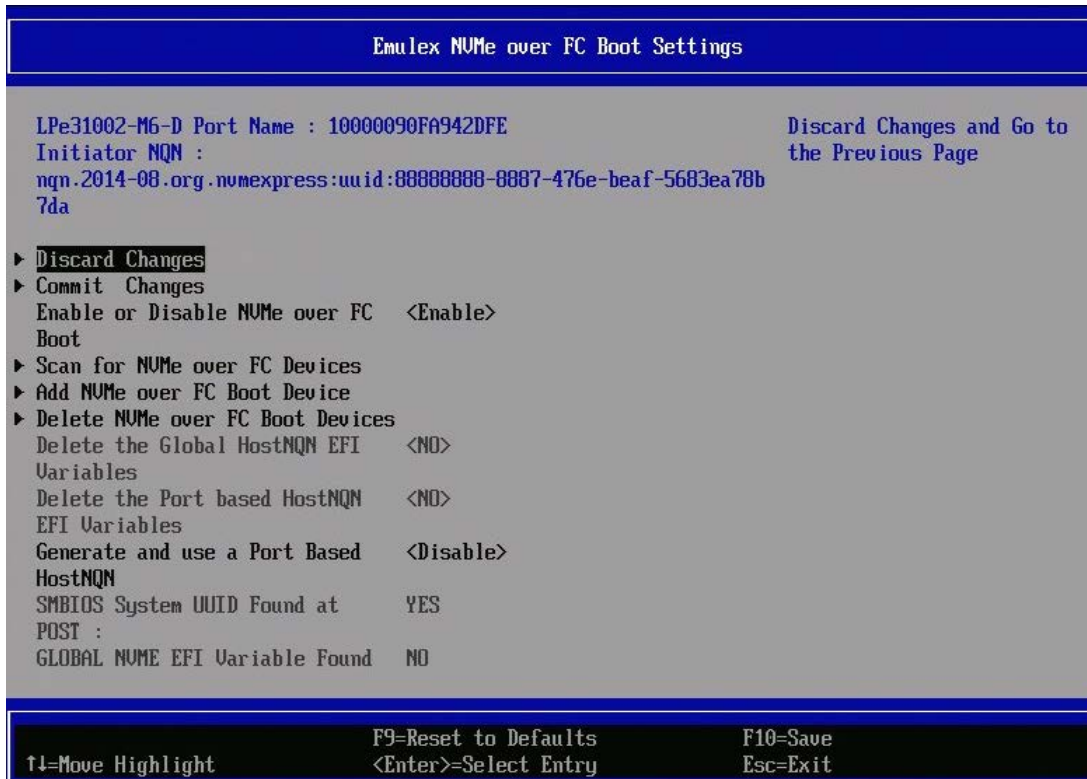
3. Select **Enable or Disable NVMe over FC Boot** and press **Enter**.

4. The **Emulex NVMe over FC Boot Settings** popup appears (Figure 32).

Figure 32: Emulex NVMe over FC Boot Settings Popup
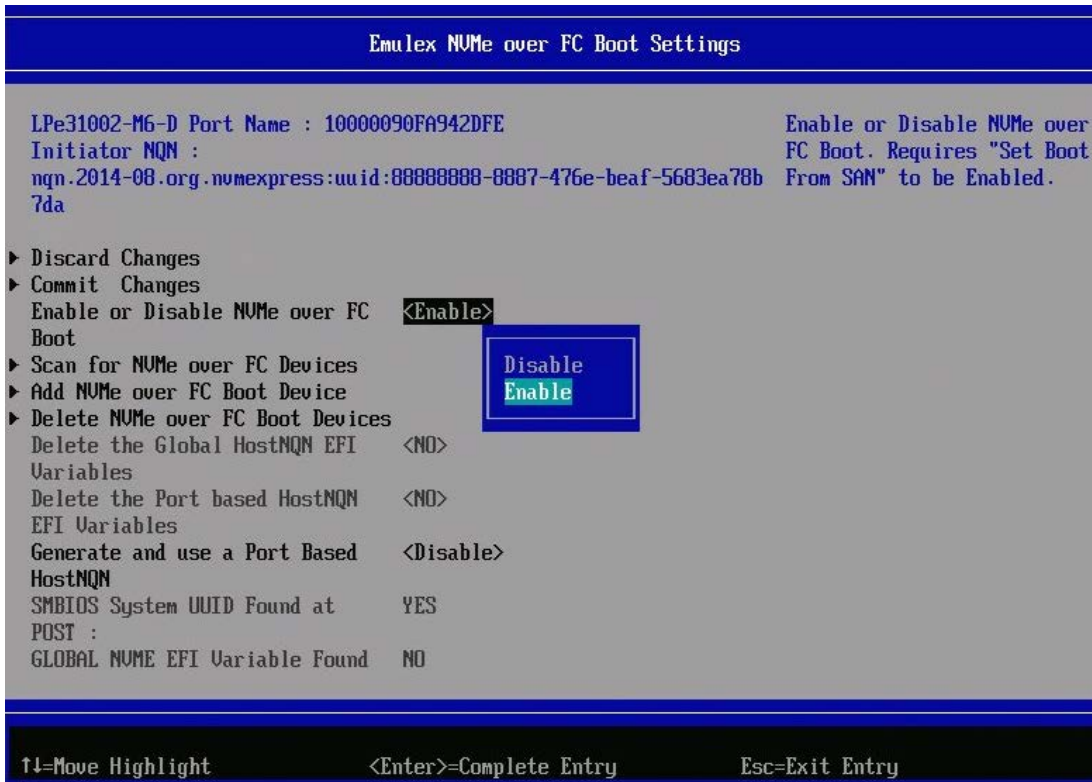
5. Make your selection, and press **Enter**.

The utility displays the new NVMe over FC boot setting (Figure 31).

6. Save the changes with **Commit Changes**.

4.15.2 Adding an NVMe over FC Boot Device

NOTE: Use only lowercase letters when adding an initiator NQN to attach an Emulex adapter to an NVMe array.

To add an NVMe over FC boot device, perform the following steps:

1. From the adapter list (Figure 1), select the adapter to which you want to add an NVMe over FC boot device and press **Enter**.

2. From the main configuration menu (Figure 2), select **Emulex NVMe over FC Boot Settings**. Press **Enter**.

    The **Emulex NVMe over FC Boot Settings** window appears (Figure 31).

3. Select **Add NVMe over FC Boot Device** and press **Enter**.

    The **Discovered NVMe over FC Controller List** window is displayed.

Figure 33: Discovered NVMe over FC Controller List Window

The window displays the discovered controllers and the firmware version of each controller.

4. Select a controller and press **Enter**. The **Discovered NVMe over FC Subsystem NQN List** window is displayed (Figure 34).

Figure 34: Discovered NVMe over FC Subsystem NQN List Window
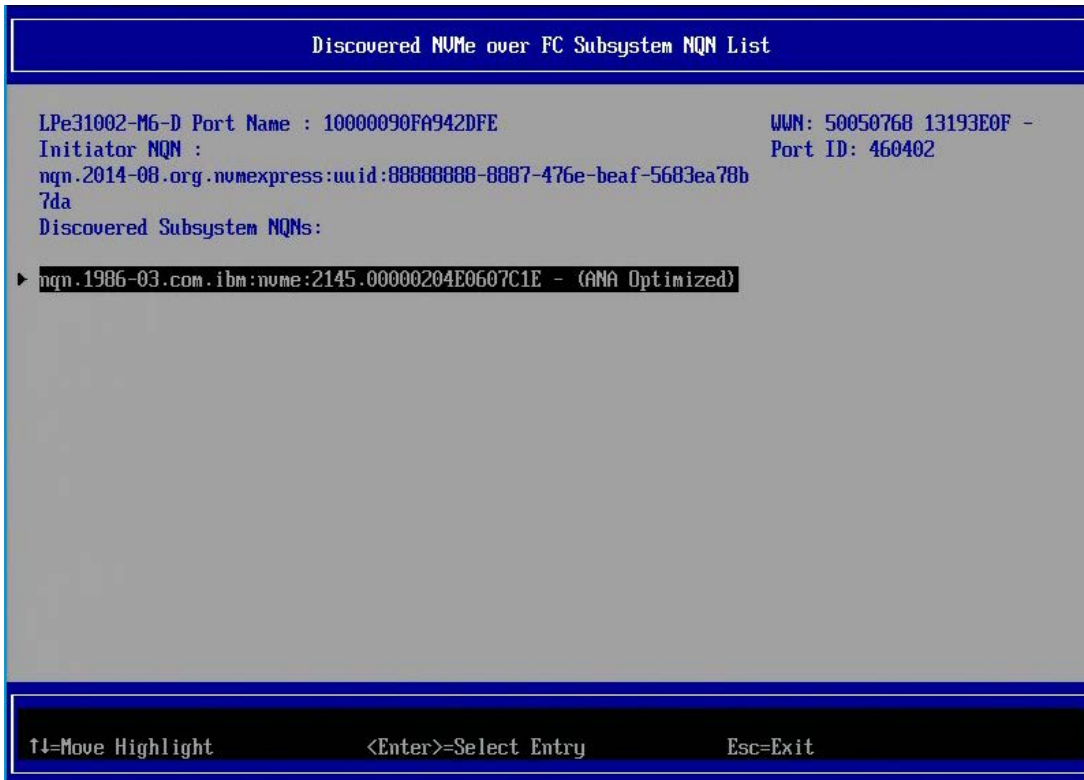
5. Make a selection and press **Enter**.

NOTE: Some arrays support Asymmetric Namespace Access (ANA). If the array supports ANA, the ANA state of the subsystem/namespace is listed.

Select and map only the NVMe over FC namespaces that are tagged with (ANA optimized) if they are listed.

The **Discovered NVMe over FC Namespace List** window is displayed (Figure 35).

Figure 35: Discovered NVMe over FC Namespace List Window
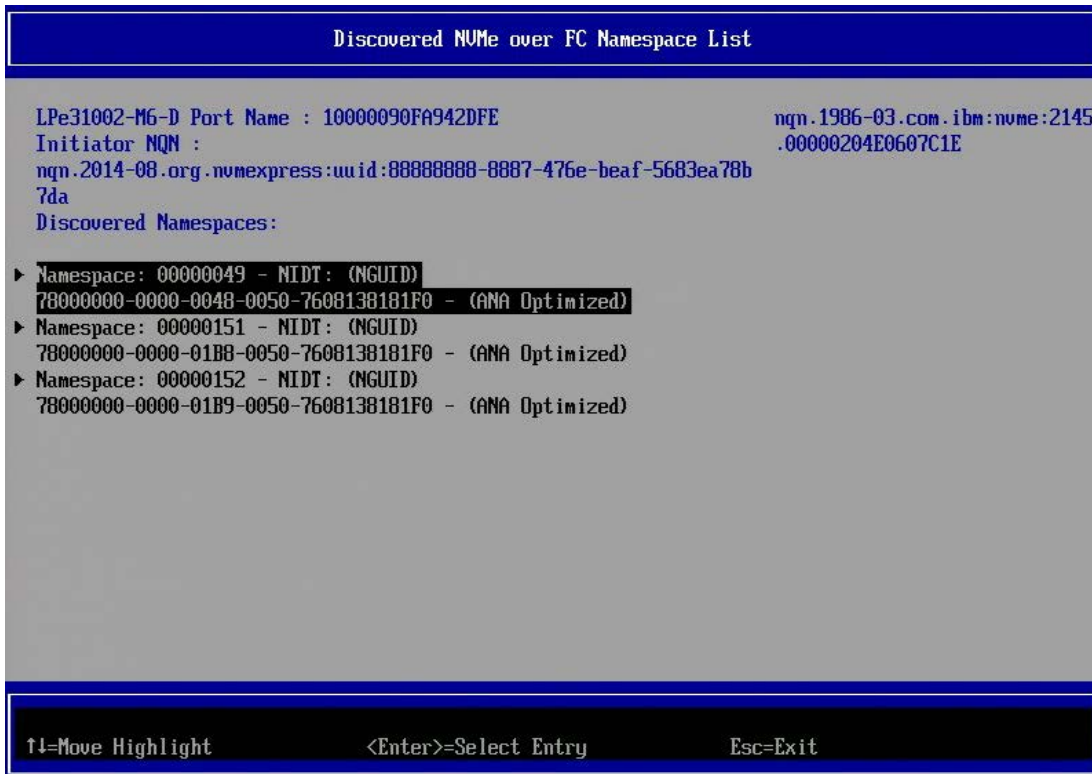
The EUID64, NGUID, or UUID is displayed next to the namespace.

6. Select the namespace that you wish to use as your boot device and press **Enter**.

NOTE: Some arrays support ANA. If the array supports ANA, the ANA state of the subsystem/namespace is listed.
Select and map only the NVMe over FC namespaces that are tagged with (ANA optimized) if they are listed.

The **NVMe over FC Namespace Boot Device** window is displayed (Figure 36).

Figure 36: NVMe over FC Namespace Boot Device Window
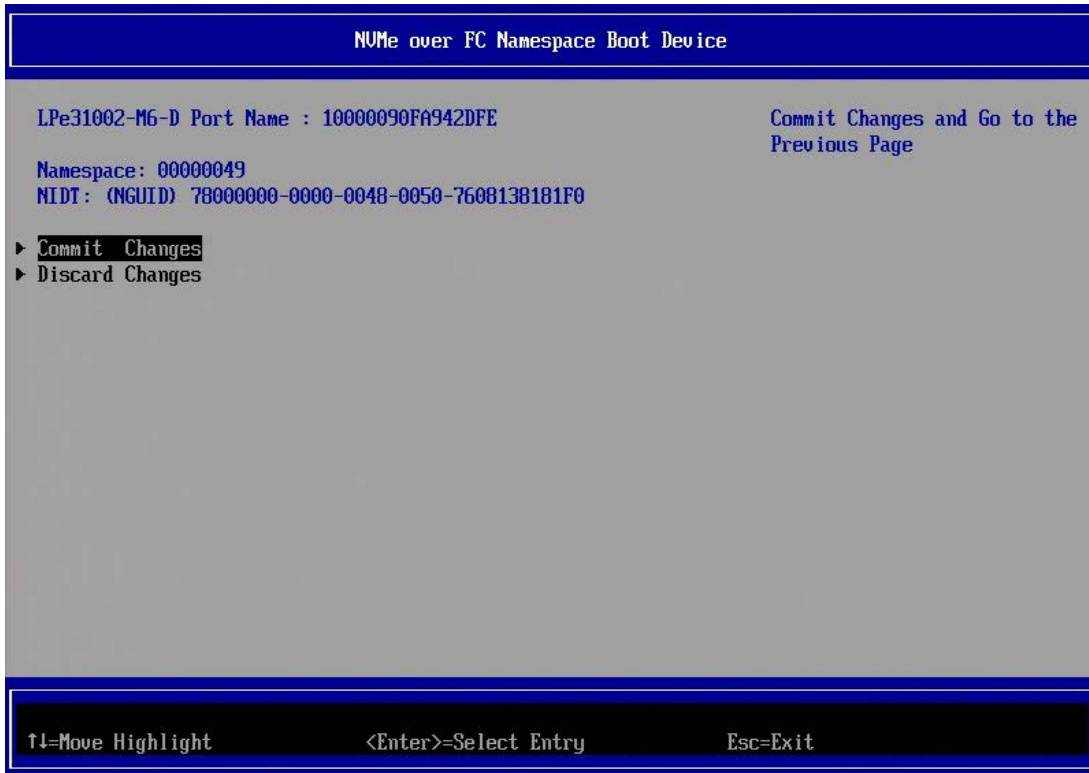
7.   Select **Commit Changes** to save the changes.

4.15.3 Deleting an NVMe over FC Boot Device

To delete a previously added NVMe over FC boot device, perform the following steps:

1. From the adapter list (Figure 1), select the adapter whose NVMe over FC boot device you want to delete and press **Enter**.

2. From the main configuration menu (Figure 2), select **Emulex NVMe over FC Boot Settings**. Press **Enter**.

   The **Emulex NVMe over FC Boot Settings** window appears (Figure 31).

3. Select **Delete NVMe over FC Boot Devices** and press **Enter**.

   The **Delete NVMe over FC Boot Devices** window is displayed.

Figure 37: Delete NVMe over FC Boot Devices Window
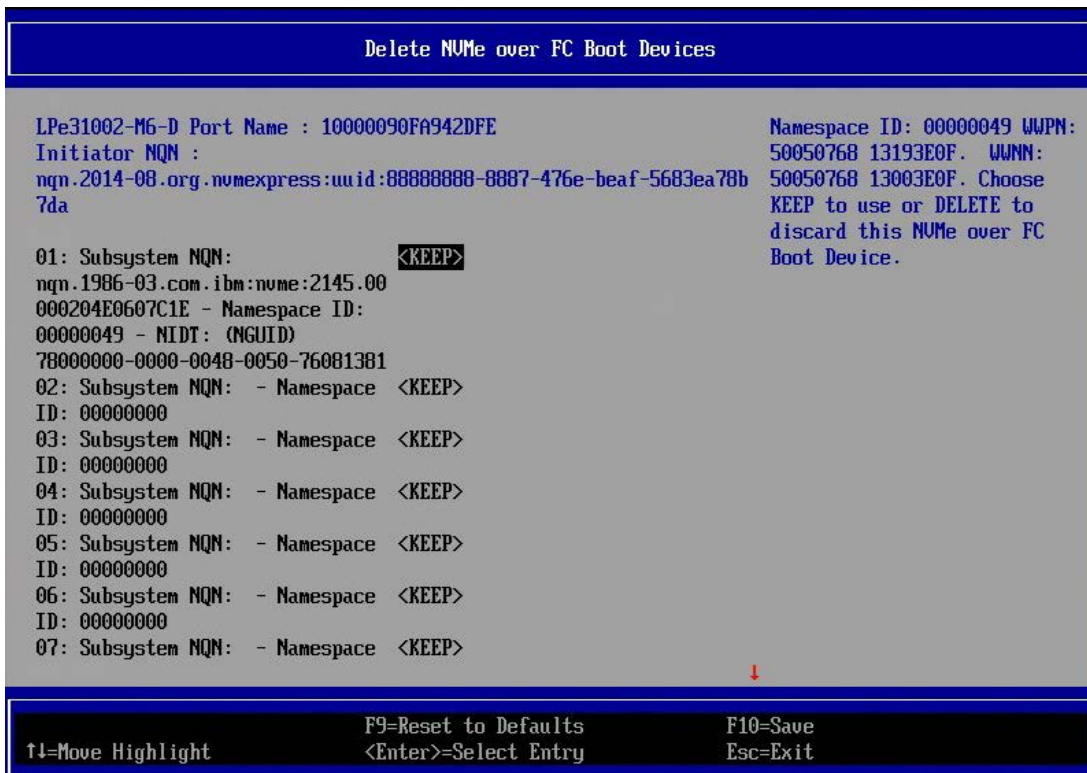
The scrollable list displays the subsystem NQN and namespace ID of each saved boot entry. You can choose to keep or delete a previously configured NVMe boot device.

4. Press **Enter** after highlighting **<KEEP>** next to the subsystem NQN and namespace ID of the entry whose FC boot devices you want to delete. The corresponding help text displays the WWPN and WWNN of the selected entry.

   The popup in the following figure appears.

Figure 38: Delete NVMe over FC Boot Devices Popup
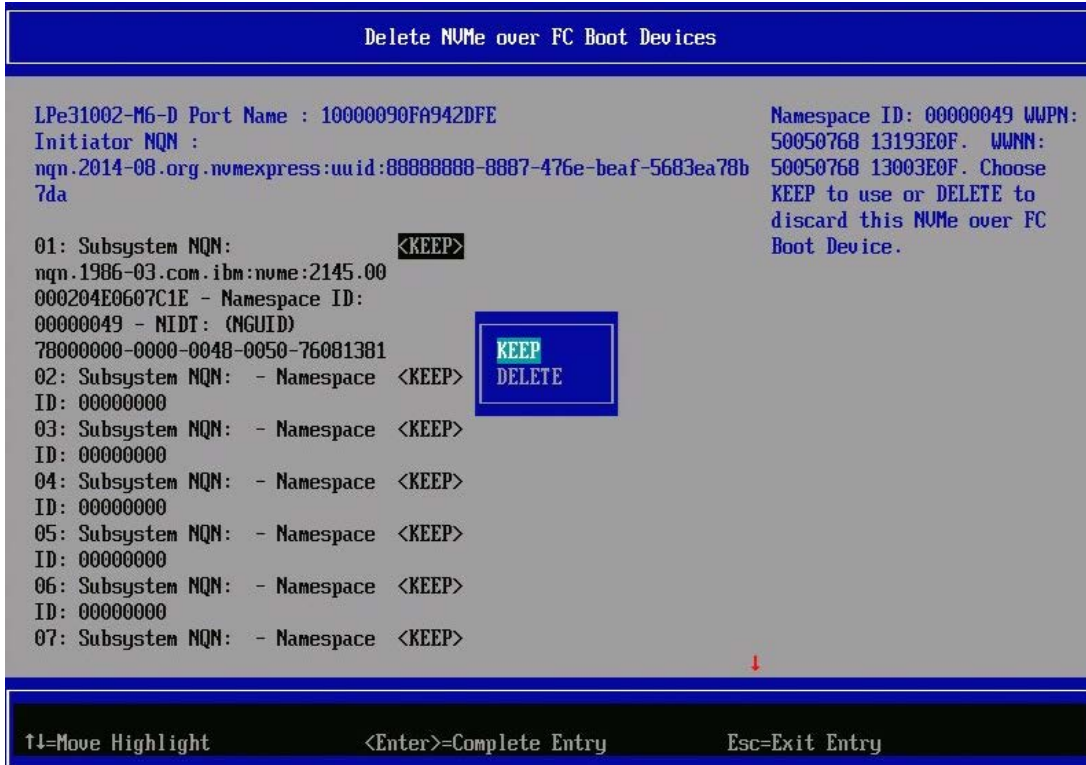
5. Select **DELETE** and press **Enter**.

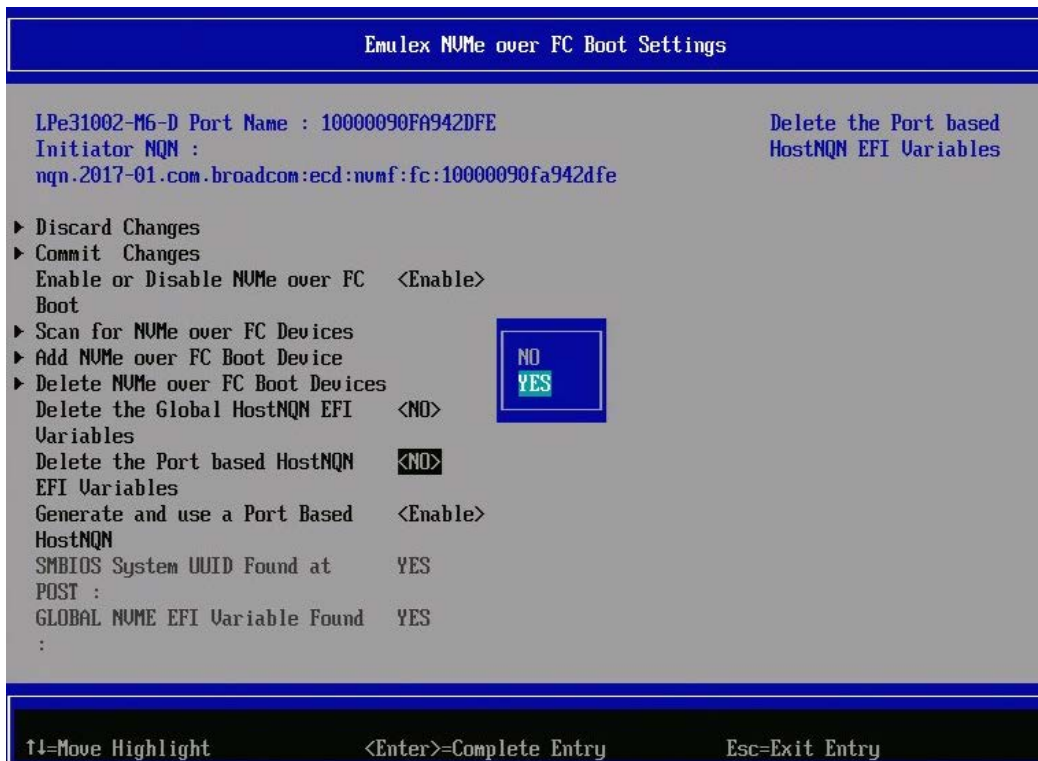6. Select **Commit Changes** and press **Enter**. The NVMe FC boot device is removed, and the **Emulex NVMe over FC Boot Settings** window is displayed (Figure 31).

4.15.4 Deleting the Global HostNQN EFI Variables

NOTE: Use this option in cases where there is an issue with NVMe EFI variables.

The **Delete the Global HostNQN EFI Variables** option is available only if NVMe EFI variables exist. Otherwise, this option is grayed out. For NVMe EFI variables to exist, you must first enable the **Generate and use a Port Based HostNQN** option and perform a reset/reconnect. NVMe EFI variables can also be created by operating system installers in certain circumstances. For instructions on enabling the **Generate and use a Port Based HostNQN** option, see Section 4.15.6, Generating and Using a Port Based HostNQN.

To delete Global HostNQN EFI variables generated by the Emulex UEFI boot, perform the following steps:

1. From the adapter list (Figure 1), select the adapter for which you want to delete the Global HostNQN EFI variables and press **Enter**.

2. From the main configuration menu (Figure 2), select **Emulex NVMe over FC Boot Settings**. Press **Enter**.

   The **Emulex NVMe over FC Boot Settings** window appears (Figure 31).

3. Select **Delete the Global HostNQN EFI Variables** and press **Enter**.

   The **Delete the Global HostNQN EFI Variables** window is displayed.

Figure 39: Delete the Global HostNQN EFI Variables Window
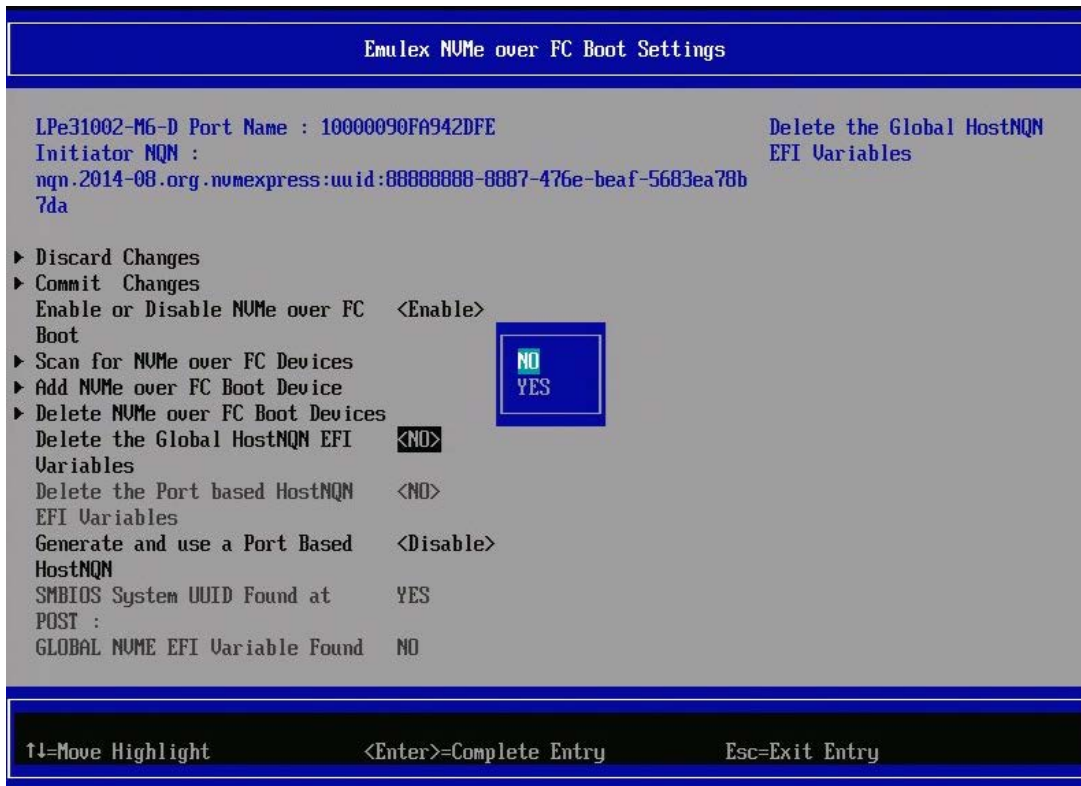
4. Select **Yes** and press **Enter**.

   The global NVMe EFI variables are deleted.

   The **Delete the Global HostNQN EFI Variables** option is grayed out.

5. Select **Commit Changes** and press **Enter**. The Global HostNQN EFI variables are removed and the **Emulex NVMe over FC Boot Settings** window is displayed (Figure 31).

4.15.5 Deleting the Port Based HostNQN EFI Variables

To delete a port based HostNQN EFI variable generated by the Emulex UEFI boot, perform the following steps:

1. From the adapter list (Figure 1), select the adapter for which you want to delete the port based HostNQN EFI variable and press **Enter**.

2. From the main configuration menu (Figure 2), select **Emulex NVMe over FC Boot Settings**. Press **Enter**.

    The **Emulex NVMe over FC Boot Settings** window appears (Figure 31).

3. Select **Delete the Port based HostNQN EFI Variables** and press **Enter**.

    The **Delete the Port based HostNQN EFI Variables** window is displayed.

Figure 40: Delete the Port based HostNQN EFI Variables Window

4. Select **Yes** and press **Enter**.

    The port based NVMe EFI variables are deleted. The **Delete the Port based HostNQN EFI Variables** option is grayed out.

5. Select **Commit Changes** and press **Enter**. The Port based HostNQN EFI variables are removed, and the **Emulex NVMe over FC Boot Settings** window is displayed (Figure 31).

4.15.6 Generating and Using a Port Based HostNQN

The default setting for the **Generate and use a Port Based HostNQN** option is disabled. During UEFI driver initialization, the Emulex UEFI boot driver will attempt to access the SMBIOS table and acquire the system UUID. If a valid UUID is obtained, the status of SMBIOS System UUID is shown as **YES**, and the Emulex UEFI boot driver generates and uses a HostNQN based on the system UUID for initiating the NVMe discovery process. This allows all ports to use the same initiator NQN value and boot via multiple HBA ports on a system.

NOTE: On some systems during UEFI driver initialization, the UUID may not be available from the SMBIOS table. In such cases, the Emulex UEFI boot driver automatically enables the **Generate and use a Port Based HostNQN** option and the status of **SMBIOS System UUID Found at POST** is shown as **NO**.

When the **Generate and use a Port Based HostNQN** option is enabled either by the Emulex UEFI boot driver or manually, the UEFI boot creates and uses a WWPN-based HostNQN and HostID value for initiating the NVMe discovery process. This allows booting only via one FC port.

After you enable the **Generate and use a Port Based HostNQN** option and perform a reset/reconnect, the NVMe HostNQN EFI variables exist and the **Delete the Port Based HostNQN EFI Variable** option is available for selection.

NOTE: For Linux, Broadcom recommends using the default setting of disabled for the **Generate and use a Port Based HostNQN** option.

For Windows and ESXi operating systems, you must enable the **Generate and use a Port Based HostNQN** option.

To generate and use a port-based HostNQN, perform the following steps:

1. From the adapter list (Figure 1), select the adapter for which you want to generate and use a port-based HostNQN. Press **Enter**.

2. From the main configuration menu (Figure 2), select **Emulex NVMe over FC Boot Settings**. Press **Enter**.
   The **Emulex NVMe over FC Boot Settings** window appears (Figure 31).

3. Select **Generate and use a Port Based HostNQN** and press **Enter**.
   The **Generate and Use a Port Based HostNQN** window is displayed.

Figure 41: Generate and use a Port Based HostNQN Window
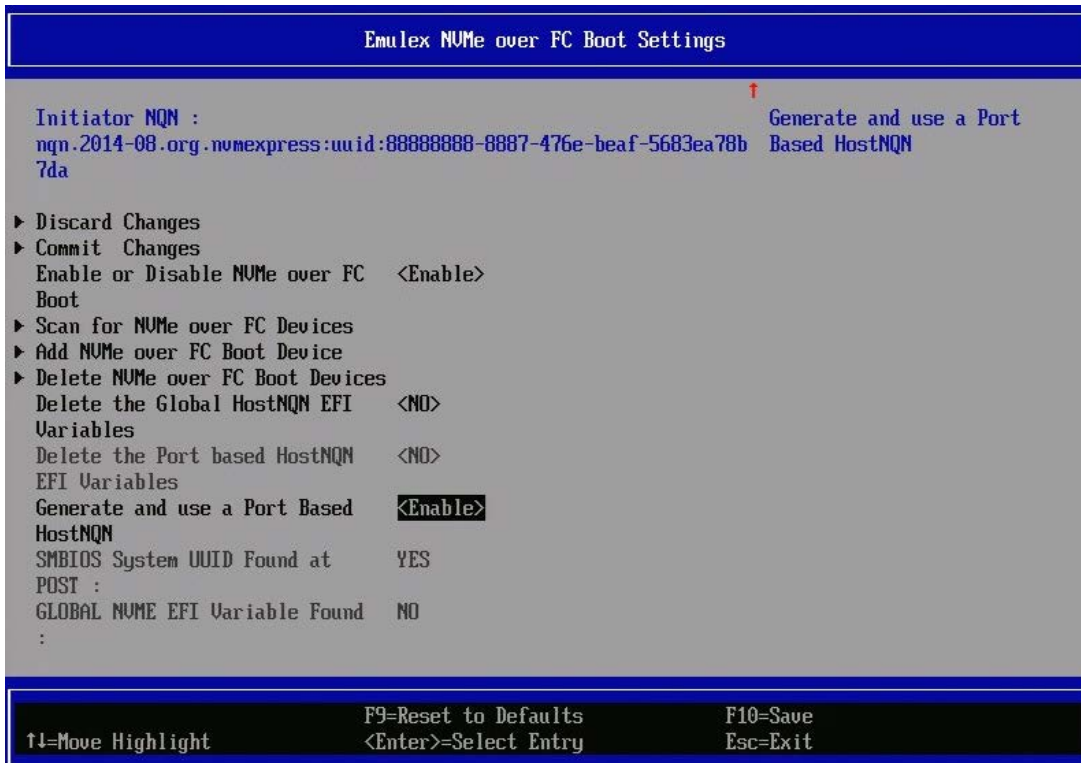
4. Select **Enable** and press **Enter**.

5. Select **Commit Changes** and press **Enter**. The adapter is enabled to generate and use a HostNQN based on the port WWPN, and the **Emulex NVMe over FC Boot Settings** window is displayed (Figure 31).

6. Reset or reconnect the adapter for the changes to take effect (see Section 4.13, Requesting a Reset or Reconnect).

   The initiator NQN string changes, and the WWN-based NQN is displayed.

4.15.7 Scanning for NVMe over FC Devices

NOTE: Use only lowercase letters when adding an initiator NQN to attach an Emulex adapter to an NVMe array.

To scan for NVMe over FC devices, perform the following steps:

1. From the adapter list (Figure 1), select the adapter for which you want to scan NVMe over FC devices. Press **Enter**.

2. From the main configuration menu (Figure 2), select **Emulex NVMe over FC Boot Settings**. Press **Enter**.

   The **Emulex NVMe over FC Boot Settings** window appears (Figure 31).

3. Select **Scan for NVMe over FC Devices** and press **Enter**.

   A list of discovered NVMe over FC controllers is displayed (Figure 33).

4.16 Enabling or Disabling the HPE Shared Memory Feature (HPE Systems Only)

During a POST, this setting enables or disables the HPE Shared Memory feature, as defined in the Option Card Sensor Data Reporting specification.

NOTE: This menu is for HPE systems only and does not appear on other systems.

NOTE: The HPE Shared Memory feature is available only on Gen9 systems.

To enable or disable the HPE Shared Memory feature, perform these steps:

1. Access the HPE **System Configuration** UEFI menu (Figure 42).

Figure 42: HPE System Configuration Menu
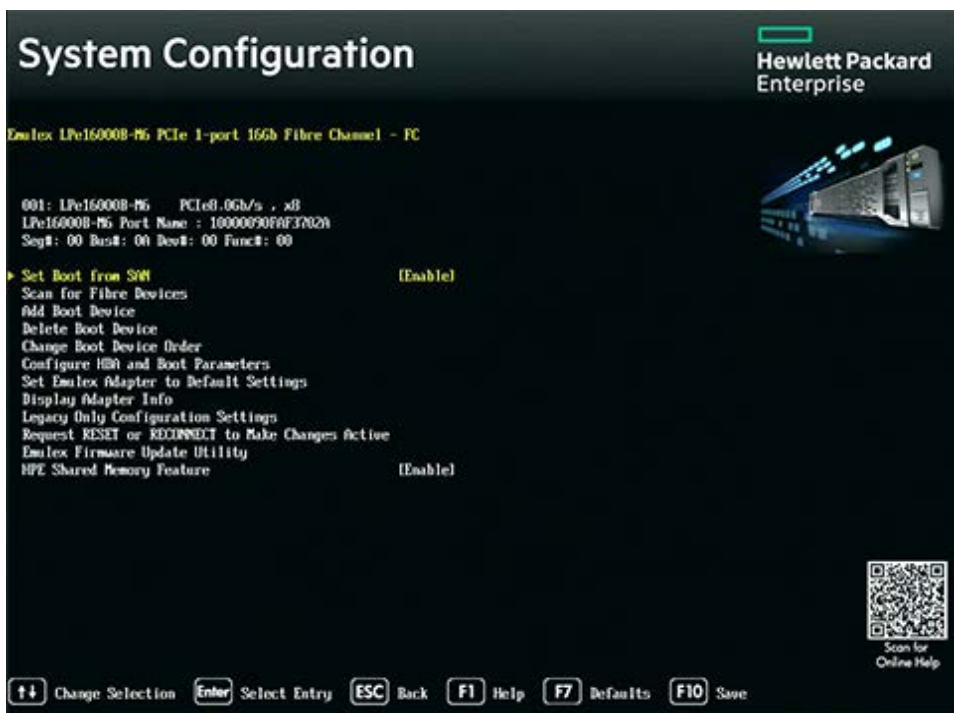
2. Navigate to **HPE Shared Memory Feature** and press **Enter**.

   The **HPE Shared Memory Feature** popup appears (Figure 43).

Figure 43: HPE Shared Memory Feature Popup
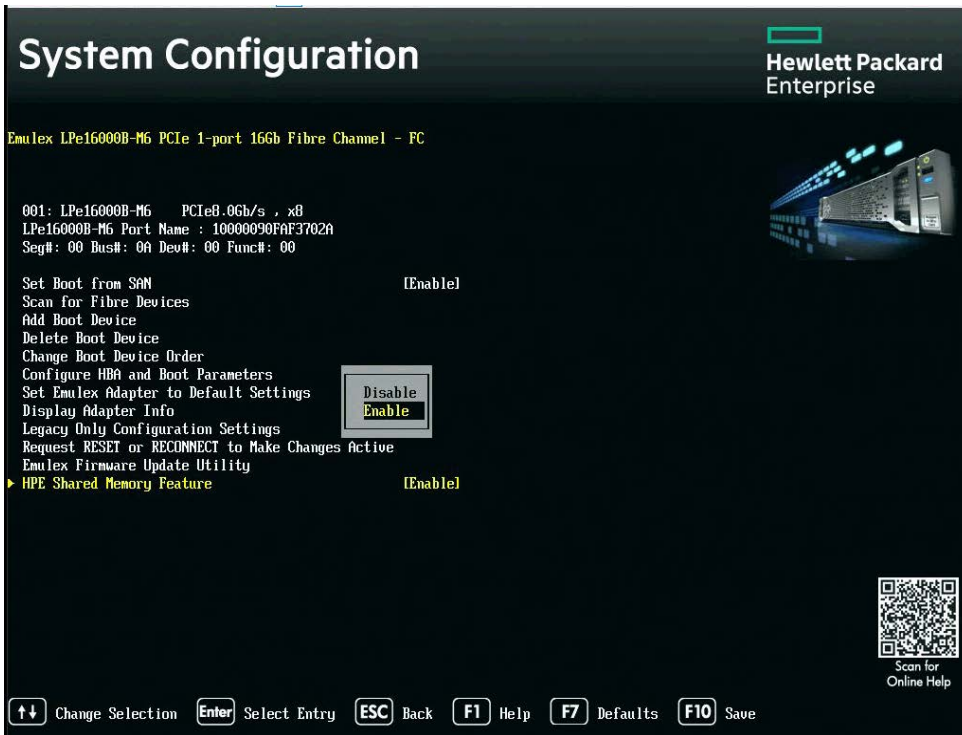
3. Make your selection, and press **F10** to save it.

Chapter 5: Emulex FC BIOS Utility

Before using the Emulex FC BIOS utility, ensure that the boot code is loaded and enabled on the adapter as described in Chapter 3, Updating and Enabling the Boot Code.

NOTE: The following information pertains to the Emulex FC BIOS utility:
- This chapter reflects the most recent release of the Emulex FC BIOS utility. Some selections might not be available if you are using an older version of the utility.
- Changes made to parameters common to UEFI and x86 drivers are changed in both configuration utilities.
- After exiting the Emulex FC BIOS utility, the system will reboot regardless of whether any changes were made.
- After you log in to the switch, connected adapter ports appear as G_Ports until you log in to the fabric.

During POST, the Emulex BIOS is executed if the legacy Option ROM is enabled in the system BIOS.

If the Emulex x86 Boot code is enabled on an adapter, the boot code performs the following tasks:
- Initializes the HBAs in the system.
- If BFS is enabled, attempts to bring up links on the ports.
- Checks for configured boot devices that are present or discoverable on the FC SAN.

A total of up to eight boot devices are displayed on the POST window, depending on how many boot devices are configured in the Emulex FC BIOS utility on the HBA. The boot devices are displayed only if they are discovered on the SAN.

The discovered order depends on the BIOS of the HBA or port that is executed first. Only the first eight targets discovered are displayed. If boot devices are discovered on the SAN, the Emulex x86 Boot code prints out a message identifying the adapter port that the target is mapped from and the disk targets that are discovered or mapped from that port.

```
--Adapter 01 LPe32002-M2:    S_ID:032D00  PCI Bus, Device, Function (25,00,00)
DID:011E00  WWPN:21010002AC019040  LUN:108
DID:011E00  WWPN:21010002AC019040  LUN:110
```

Each series of Option ROMs can manage up to eight boot devices.

NOTE: Since there is a limit of eight boot devices, and since the XROM HBA execution order depends on the adapter BIOS, you can configure no more than eight boot devices per adapter family.

5.1 Navigating the Emulex FC BIOS Utility

Use the following methods to navigate the Emulex FC BIOS utility:
- Press the up and down arrows on your keyboard to move through and select menu options or configuration fields. If multiple adapters are listed, use the up and down arrows to scroll to the additional adapters.
- Press **PageUp** to scroll to the previous page.
- Press **PageDn** to scroll to the next page.
- Press **Enter** to select a menu option, to select a changed value, to select a row in a window, or to change a configuration default.
- Press **Esc** to go back to the previous menu.

5.2 Starting the Emulex FC BIOS Utility

To start the Emulex FC BIOS Utility, perform these steps:

1. Turn on the computer and press and hold down **Alt** or **Ctrl** and press **E** immediately (within 4 seconds of the Emulex bootup message).

    An adapter listing window is displayed (Figure 44).

NOTE: The following information pertains to the Emulex FC BIOS utility:

- ● Links should be connected and established before you enter the utility; otherwise, you might receive an error message.
- ● After you enter the utility, any dynamic changes made to the SAN are not reflected unless the system is reset.
- ● If the bootup message does not appear, you must enable the x86 Boot code. See Section 5.3, Enabling an Adapter to Boot From SAN, for more information.
- ● After you exit the utility, the system will reboot regardless of whether any changes were made.
- ● The utility support is server specific. Some servers might not support the Emulex FC BIOS utility. In such cases, all configuration must be done from the UEFI HII.

Figure 44: Adapter Listing Window


```
              Emulex FC BIOS Utility, KE 14.0.198.6031

        This utility displays and saves changes when selected.
        You will be prompted to reboot for all changes to take effect.

                  Emulex FC Ports in the System:

    1. LPe35004-M2:     Bus:AF Dev:00 Func:00    WWPN: 100000109B65DC1A
    2. LPe35004-M2:     Bus:AF Dev:00 Func:01    WWPN: 100000109B65DC1B
    3. LPe35004-M2:     Bus:AF Dev:00 Func:02    WWPN: 100000109B65DC1C
    4. LPe35004-M2:     Bus:AF Dev:00 Func:03    WWPN: 100000109B65DC1D




            Enter <Esc> to exit  <PageDn> to Next Page
                <↑/↓> to Highlight, <Enter> to Select
```

2. Select the adapter to configure and press **Enter**.

    The main configuration menu is displayed (Figure 45).

Figure 45: Main Configuration Menu

The Link Status indicates the state of the physical link. For instructions on how to bring the link up, see Section 5.4, Scanning for Target Devices.

Under normal circumstances, you would first configure boot devices using the Emulex FC BIOS utility (see Section 5.5, Configuring Boot Devices). However, in the following two situations, you must perform the indicated procedure first:

■ The adapter is not enabled to BFS: You must enable the adapter's BIOS to BFS (see Section 5.3, Enabling an Adapter to Boot From SAN).

■ If you want to use a topology other than the default, you must change the topology setting before configuring boot devices (see Section 5.7.1, Changing the Topology).

5.3 Enabling an Adapter to Boot From SAN

To enable an adapter to boot from SAN, perform these steps:

1. From the main configuration menu (Figure 45), select **Enable/Disable Boot from SAN** and press **Enter**.

NOTE: Adapters are disabled by default.

BFS must be enabled on one adapter to use remote boot functionality. After you enable an adapter, the status of the Boot BIOS changes to enabled as shown in Figure 46.

NOTE: x86 Boot code supports a maximum of 16 ports.

Figure 46: Boot BIOS Status Window

5.4 Scanning for Target Devices

NOTE: The adapters bring up the link while scanning for target devices (see Section 5.4, Scanning for Target Devices) or while configuring boot devices (see Section 5.5, Configuring Boot Devices).

To scan for target devices, perform these steps:

1. From the adapter listing window (Figure 44), select the adapter that you want to scan for target devices and press **Enter**.

2. From the main configuration menu (Figure 45), select **Scan for Target Devices** and press **Enter**. A list of the discovered targets is displayed (Figure 47). This is only a list of discovered target devices to determine SAN connectivity.

Figure 47: Devices Present on the Adapter

5.5 Configuring Boot Devices

This option supports FC_AL private loop and point-to-point. When operating in loop (FC_AL) topology, the system automatically determines whether you are configured for a private loop. For the **Configure Boot Devices** option, the eight boot entries are zero by default (**D** key).

Use the **Boot Devices** window (Figure 48) to configure up to eight boot entries for fabric point-to-point or private loop configurations. The first adapter is usually in the lowest PCI slot in the system. This device is the only boot device, and it is the only device exported to the multi-boot menu.

NOTE: The following information pertains to topology:

- Loop topology (FC-AL) is supported only at speeds equal to 8Gb/s or slower.
- If you want to use a topology other than the default, you must change the topology setting before configuring boot devices (see Section 5.7.1, Changing the Topology). For FC-AL, each adapter has a default AL_PA of 01 (hexadecimal).

To configure boot devices, perform these steps:

1. On the main configuration menu (Figure 45), select **Configure Boot Devices.**

   A list of eight boot devices is shown (Figure 48). The primary boot device is the first entry shown, and it is the first bootable device.

   If the first boot entry fails due to a hardware error, the system can boot from the second bootable entry. If the second boot entry fails, the system boots from the third bootable entry and so on, if it is configured to do so in the system BIOS on a BBS system.

NOTE: The Emulex FC BIOS only presents the boot devices to the system BIOS. The system BIOS must enumerate and attempt to boot from the drive as the primary hard drive from which to boot. See Section 5.9, Using Multipath BFS.

Figure 48: List of Saved Boot Devices Window

2. Select a boot entry and press **Enter**.

   A window similar to Figure 49 is displayed.

Figure 49: Device Selection List Example (Array) Window

NOTE: To minimize the amount of time needed to locate the boot device, select the drive with the lowest AL_PA as the boot device when connected to devices on an FC-AL.

3. Select **00** and press **Enter** to clear the selected boot entry, or select the WWPN or D_ID of a device to configure. A window similar to Figure 50 is displayed.

Figure 50: Starting LUN Offset Window

4. If you select a device, you are prompted for the starting LUN. Select the starting LUN by using the up and down arrow keys (Figure 50). The starting LUN can be any number from 0 to 255.

Figure 51: LUN Listing Window

You can define 256 LUNs per target. As shown in Figure 51, only eight LUNs per page are shown, from the starting LUN offset, up to 256 consecutive LUNs.

5. Use the arrow key up and down to select the LUN you want to set up as a boot device.

    The **Boot Devices** window is displayed (Figure 52 and Figure 53). You can choose to boot the device using WWPN or D_ID.

Figure 52: Boot Devices Window Using D_ID

Figure 53: Boot Devices Window Using WWPN

6. Use the up arrow and down arrow keys (Figure 53) to select the boot method you want. If you select to boot the device by WWPN, the WWPN of the selected entry is used to configure the boot device (Figure 54). If the D_ID is selected, then the D_ID is used to configure the boot device (Figure 55).

Figure 54: Primary Boot Device Setup Window Listed by WWPN

Figure 55: Primary Boot Device Setup Window Listed by D_ID

7. Press **Enter** to select the change.

8. Press **Esc** until you exit the Emulex FC BIOS utility.
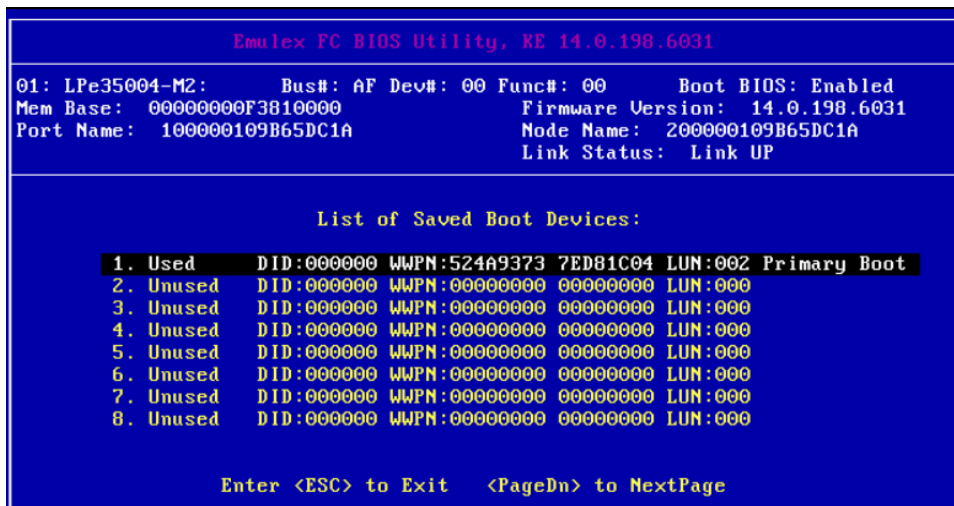
9. Reboot the system for the new boot path to take effect.

5.6 Configuring Advanced Adapter Parameters

The Emulex FC BIOS utility has numerous options that can be modified to provide for different behavior. Use the `BIOS` utility to perform the following tasks:

- Change the PLOGI retry timer.
- Enable or disable spinup delay.
- Set autoscan.
- Enable or disable EDD 3.0.
- Enable or disable the start unit command.
- Enable or disable the environment variable.
- Enable or disable the auto boot sector.

To access the advanced adapter parameters configuration menu, from the main configuration menu (Figure 45) select **Configure Advanced Adapter Parameters** and press **Enter.**

The advanced adapter parameters configuration menu is displayed (Figure 56).

Figure 56: Advanced Adapter Parameters Configuration Menu

Default settings are acceptable for most installations.

To reset all values to their defaults, from the main configuration menu (Figure 45) select **Reset Adapter Defaults** and press **Enter**.

5.6.1 Changing the PLOGI Retry Timer

This option is useful for Tachyon-based RAID arrays. In rare situations, a Tachyon-based RAID array resets itself and the port goes offline temporarily. When the port returns to operation, the Port Login scans the loop to discover this device. The Port Login retry interval is the time it takes for one PLOGI to scan the whole loop (if 126 AL_PAs are on the loop). You can choose the following options:

- No PLOGI retry: 0 ms – default
- 50 ms takes 5 to 6 seconds per device
- 100 ms takes 12 seconds per device
- 200 ms takes 22 seconds per device

To set the interval for the PLOGI retry timer, perform these steps:

1. On the main configuration menu (Figure 45), select **Configure Advanced Adapter Parameters** and press **Enter**.

   The adapter configuration menu is displayed (Figure 56).

2. Select **Change PLOGI Retry Timer** and press **Enter.** Information similar to Figure 57 is displayed.

Figure 57: Change the PLOGI Retry Timer Window

3. Select the retry timer interval.

4. Press **Enter** to accept the new interval.

5. Press **Esc** until you exit the Emulex FC BIOS utility.

6. Reboot the system.

5.6.2 Enabling or Disabling the Spinup Delay

This option allows you to enable or disable the disk spinup delay. The factory default setting is disabled.

If at least one boot device has been defined, and the spinup delay is enabled, the BIOS searches for the first available boot device.

- If a boot device is present, the BIOS boots from it immediately.
- If a boot device is not ready, the BIOS waits for the spinup delay and, for up to 3 additional minutes, continues the boot scanning algorithm to find another multi-boot device.

If boot devices have not been defined and autoscan is enabled, the BIOS waits for 5 minutes before scanning for devices.

- In a private loop, the BIOS attempts to boot from the lowest target AL_PA it finds.
- In an attached fabric, the BIOS attempts to boot from the first target found in the NameServer data.

To enable or disable the spinup delay, perform these steps:

1. On the main configuration menu (Figure 45), select **Configure Advanced Adapter Parameters** and press **Enter**.

   The adapter configuration menu is displayed (Figure 56).

2. Select **Enable or Disable Spinup Delay** and press **Enter**.

   The enable or disable spinup delay window (Figure 58) is displayed.

Figure 58: Enable or Disable Spinup Delay Window
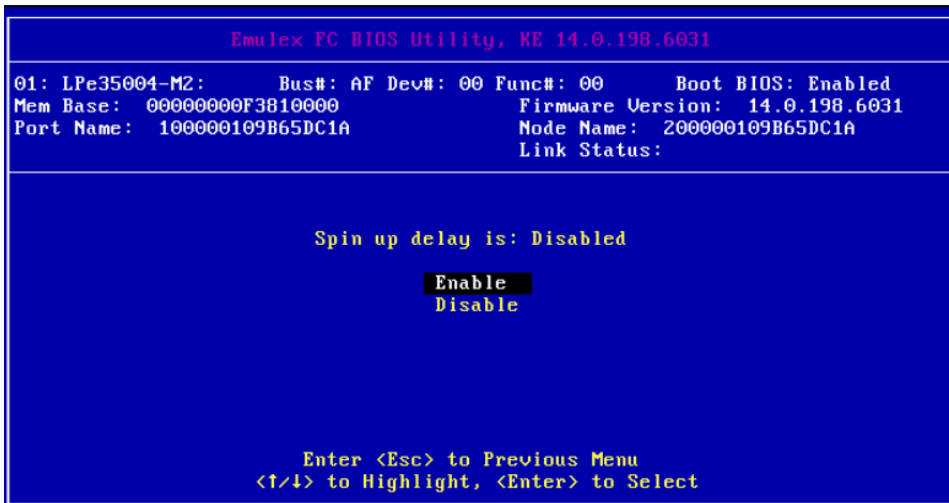
3. Select whether to enable or disable the spinup delay.

4. Press **Enter** to accept the new value.

5. Press **Esc** until you exit the Emulex FC BIOS utility.

6. Reboot the system.

5.6.3 Setting Autoscan

This option allows you to attempt to boot from the first found disk reported by the switch on the SAN. The factory default is disabled. If more than one adapter with the same PCI bus number exists in the system, and each has a boot drive attached, the first PCI-scanned adapter is the boot adapter.

The following autoscan options are available:

- Autoscan disabled – default.
- Any first device – The first adapter issues a name server inquiry to the switch and the first reported FC disk becomes the boot device. The adapter attempts to log in to a private loop. The first successfully scanned device becomes the boot device. Only this device is exported to the multi-boot menu.
- First LUN 0 device.
- First NOT LUN 0 device (a device other than LUN 0).

To set autoscan, perform these steps:

1. From the main configuration menu (Figure 45), select **Configure Advanced Adapter Parameters** and press **Enter**.

   The adapter configuration menu is displayed (Figure 56).

2. Select **Autoscan setting** and press **Enter**.

   Figure 59 is displayed.

Figure 59: Autoscan Setting Menu

3. Select the desired autoscan setting.

4. Press **Enter** to accept the new option.

5. Press **Esc** until you exit the Emulex FC BIOS utility.

6. Reboot the system.

5.6.4 Enabling or Disabling EDD 3.0

EDD 3.0 provides additional data to the operating system bootloader during the INT-13h function 48h (get device parameters) call. This information includes the path to the boot device and disk size. The default setting for EDD 3.0 is enabled.

NOTE: If EDD 3.0 is disabled, EDD 2.1 is enabled.

To enable or disable EDD 3.0, perform these steps:

1. From the main configuration menu (Figure 45), select **Configure Advanced Adapter Parameters** and press **Enter**.
    The adapter configuration menu is displayed (Figure 56).

2. Select **Enable or Disable EDD 3.0**.
    The **EDD 3.0** window (Figure 60) is displayed.

Figure 60: EDD 3.0 Window

3. Select the desired EDD 3.0 setting.

4. Press **Enter** to accept the new setting.

5. Press **Esc** until you exit the Emulex FC BIOS utility.

6. Reboot the system.

5.6.5 Enabling or Disabling the Start Unit Command

This option allows a start unit command to be sent to a particular disk. This option might be needed if the disk you want to boot from is not yet spun up or started. You must know the specific LUN to issue the SCSI start unit command. The default setting is disabled.

To enable or disable the start unit command, perform these steps:

1. From the main configuration menu (Figure 45), select **Configure Advanced Adapter Parameters** and press **Enter**.

   The adapter configuration menu is displayed (Figure 56).

2. Select **Enable or Disable Start Unit Command** and press **Enter**.

   The **Start Unit Command** window (Figure 61) is displayed.

Figure 61: Start Unit Command Window

3. Select the desired start unit command setting.

4. Press **Enter** to accept the new setting.

5. Press **Esc** until you exit the Emulex FC BIOS utility.

6. Reboot the system.

5.6.6 Enabling or Disabling the Environment Variable

This option sets the boot controller order if the system supports the environment variable. The default setting is disabled.

To enable or disable the environment variable, perform these steps:

1. From the main configuration menu (Figure 45), select **Configure Advanced Adapter Parameters** and press **Enter**.
   The adapter configuration menu is displayed (Figure 56).

2. Select **Enable or Disable Environment Variable** and press **Enter**.
   The **Environment Variable** window (Figure 62) is displayed.

Figure 62: Environment Variable Window

3. Select the desired setting.

4. Press **Enter** to accept the new setting.

5. Press **Esc** until you exit the Emulex FC BIOS utility.
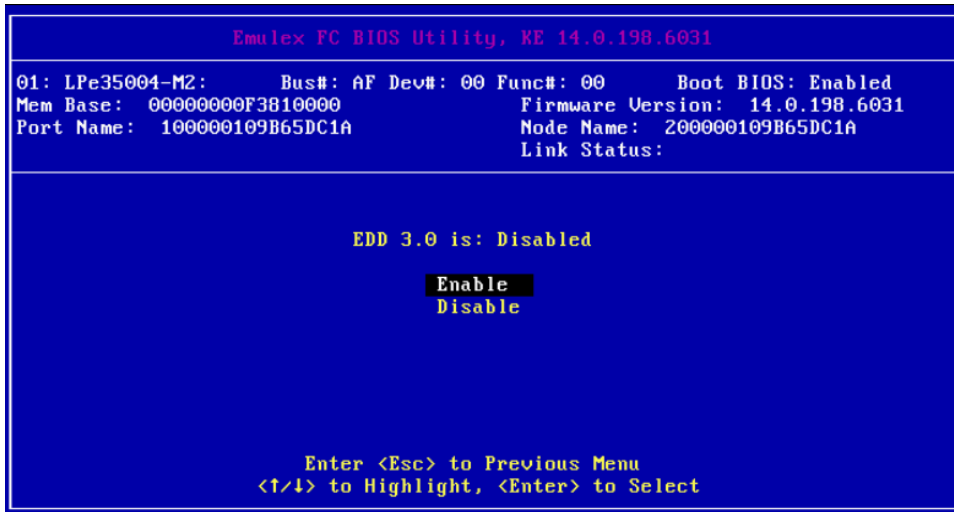
6. Reboot the system.

5.6.7 Enabling or Disabling Auto Boot Sector

This option automatically defines the boot sector of the target disk for the migration boot process, which applies only to HP MSA1000 arrays. If no partition is on the target, the default boot sector format is 63 sectors. The default setting is disabled.

To enable or disable the auto sector format, perform these steps:

1. From the main configuration menu (Figure 45), select **Configure Advanced Adapter Parameters** and press **Enter**.
   The adapter configuration menu is displayed (Figure 56).

2. Select **Enable or Disable Auto Boot Sector** and press **Enter**.
   The **Auto Boot Sector** window (Figure 63) is displayed.

Figure 63: Auto Boot Sector Window

3. Select the desired setting.

4. Press **Enter** to accept the new setting.

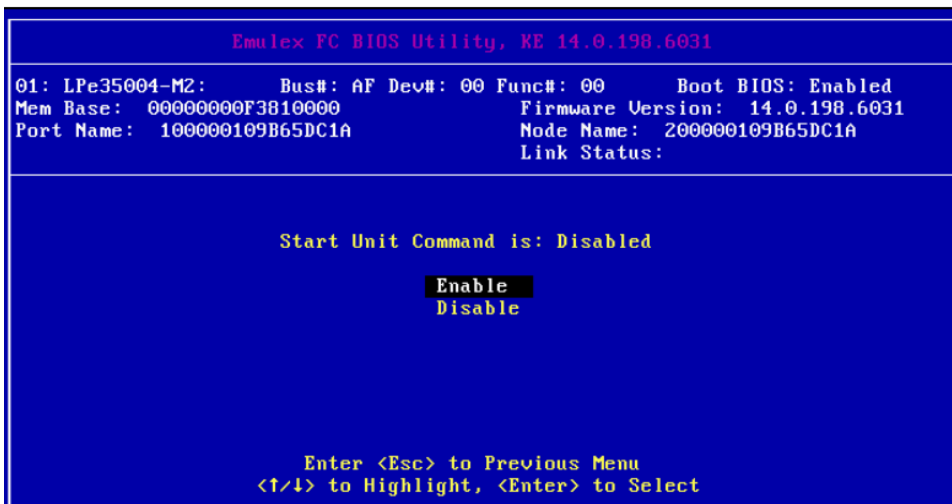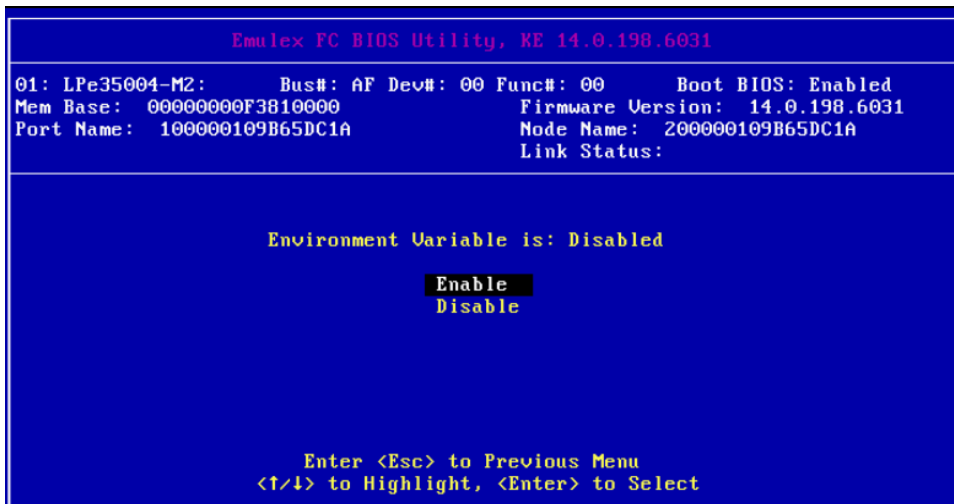5. Press **Esc** until you exit the Emulex FC BIOS utility.

6. Reboot the system.

5.7 Configuring Adapter Firmware Parameters

The Emulex FC BIOS utility has numerous options that can be modified to provide for different behavior. Use the Emulex FC BIOS utility to perform the following tasks:

■ Select a topology.
■ Enable or disable Brocade® FA-PWWN.

NOTE: FA-PWWN is no longer supported; the feature will be disabled in a future release.

■ Select trunking.
■ Select a link speed.
■ Enable or disable 16G FEC.

NOTE: If you want to use a topology other than the default, you must change the topology setting before configuring boot devices (see Section 5.7.1, Changing the Topology).

To access the adapter firmware parameters configuration menu, from the main configuration menu (Figure 45) select **Configure Adapter Firmware Parameters** and press **Enter.**

The adapter firmware parameters configuration menu is displayed (Figure 64).

Figure 64: Adapter Firmware Parameters Configuration Menu

Default settings are acceptable for most installations.

To reset all values to their defaults, from the main configuration menu (Figure 45), select **Reset Adapter Defaults** and press **Enter**.

5.7.1 Changing the Topology

You can configure the following topologies:

- Point-to-Point
- FC-AL

NOTE: For FC-AL, only direct-attach is supported.

To select the adapter topology, perform these steps:

1. On the main configuration menu (Figure 45), select **Configure Adapter Firmware Parameters** and press **Enter**.

   The adapter firmware parameters configuration menu is displayed (Figure 64).

2. Select **Topology Selection** and press **Enter**.

   Information similar to Figure 65 is displayed.

The topology options available are shown in Figure 65.

Figure 65: Topology Menu

5.7.2 Enabling or Disabling the Brocade FA-PWWN

NOTE: FA-PWWN is no longer supported; the feature will be disabled in a future release.

5.7.3 Selecting Trunking

Trunking enables you to combine multiple physical FC links to form a single logical link (aggregated port). The aggregated port's maximum link speed is the sum of the maximum link speeds of the individual physical links comprising the aggregated port. The actual link speed of the aggregated port depends on the states (active or non-active) of the individual physical links comprising the aggregated port.

The physical links comprising an aggregated port are referred to as lanes. Both 2-lane and 4-lane aggregated ports are supported.

The LPe35002 and LPe36002 adapters support 2-lane port aggregation. If 2-lane port aggregation is configured on a dual-port adapter, the two physical links are combined to form a single 2-lane aggregated port whose aggregate speed is two times the maximum speed supported, assuming both physical links are active.

For example:
- On an LPe35002-series adapter, an aggregated port comprised of two physical links running at 32Gb/s each has a potential logical (aggregate) link speed of 64Gb/s.
- On an LPe36002-series adapter, an aggregated port comprised of two physical links running at 64Gb/s each has a potential logical (aggregate) link speed of 128Gb/s.

An LPe35004 adapter supports both 2-lane port aggregation and 4-lane port aggregation as follows:
- If 2-lane port aggregation is configured on an LPe35004 adapter, the four physical links on the adapter are divided among two separate aggregated ports. The two lowest-numbered physical links form the first aggregated port, and the two highest-numbered physical links form the second aggregated port.
- If 4-lane port aggregation is configured on an LPe35004 adapter, all four physical links are combined to form a single 4-lane trunk whose aggregate speed is potentially 128Gb/s, assuming all four links are active, each at 32Gb/s link speed.

NOTE: The following information pertains to trunking:
- Trunking (also referred to as FC port aggregation) is supported only on LPe35000-series and LPe36000-series adapters.
- Trunking is not supported at 8Gb/s speeds, and the link will not come up at this speed.
- Dynamic D_Port cannot coexist with the trunking feature. If trunking is enabled, the firmware automatically disables Dynamic D_Port.

If you are using an adapter that does not support trunking, the Select Trunking option is not be available.

NOTE: Changing the trunking mode resets the adapter to factory default settings. You must configure the HBA and boot parameters after changing the trunking mode for BFS to work.

To enable or disable trunking, perform these steps:

NOTE: Before you enable or disable trunking on an Emulex adapter, follow the instructions from Brocade to enable or disable trunking on the switch.

1. On the main configuration menu (Figure 45), select **Configure Adapter Firmware Parameters** and press **Enter**.

   The adapter firmware parameters configuration menu is displayed (Figure 64).

2. Select **Trunking Selection** and press **Enter**.

Depending on the adapter you are using, information similar to Figure 66 is displayed.

Figure 66: Trunking Menu for LPe35004 Adapters

3. Select the trunking mode from the available options.

Disable Trunking disables trunking for the adapter.

4. Press **Enter** to accept the trunking selection.

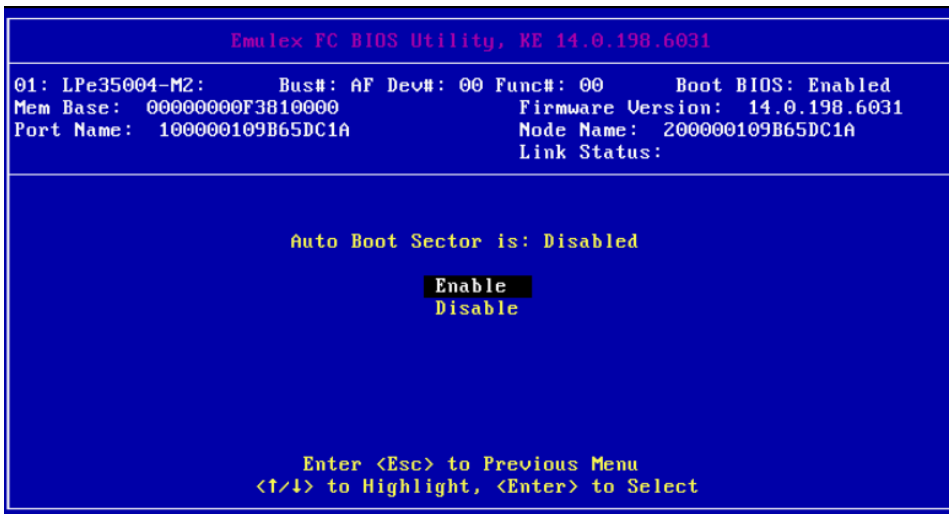5. Press **Esc** until you exit the Emulex FC BIOS utility.

6. Reboot the system.

If trunking has been enabled, only the trunked ports are visible in the main configuration menu. If trunking has been disabled, all the physical ports are visible in the main configuration menu.

NOTE: When trunking is enabled, the maximum aggregate link speed is always displayed regardless of the state of each physical link.

5.7.3.1 Viewing the Trunking Configuration

NOTE: When trunking is enabled, boot from SAN must be enabled and the system must be rebooted for the physical port status to be up and for the link speed to be shown.

1. From the adapter listing window (Figure 44), select the adapter for which you want to view trunking information and press **Enter**.

2. From the main configuration menu (Figure 45), select **Scan for Target Devices** to bring the links up for the trunk.

3. Press **Esc** to return to the main configuration menu.

4. On the main configuration menu (Figure 45), select **Configure Advanced Firmware Parameters** and press **Enter**.

   The adapter firmware configuration menu is displayed (Figure 64).

5. Select **Trunking Selection** and press **Enter**.

   Information similar to Figure 67 or Figure 68 is displayed.

   Figure 67 shows the trunking menu for LPe35004 (4-port) adapters with a 2-lane trunk configured.

   Figure 68 shows the trunking menu for LPe35004 (4-port) adapters with a 4-lane trunk configured.

   The physical link status, trunk link status, and trunk link speed are displayed.

Figure 67: Trunking Menu for LPe35004 Adapter (2-Lane)

Figure 68: Trunking Menu for LPe35004 Adapter (4-Lane)

```
         Emulex FC BIOS Utility, KE 14.0.198.6031
┌────────────────────────────────────────────────────────────┐
│ 01: LPe35004-M2:     Bus#: AF Dev#: 00 Func#: 00    Boot BIOS: Disabled │
│ Mem Base:   00000000F3810000          Firmware Version:  14.0.198.6031 │
│ Port Name:   100000109B65DC1A         Node Name:  200000109B65DC1A │
│                                       Link Status:  Link UP │
└────────────────────────────────────────────────────────────┘

    Trunking:  4-Lane Trunk        Trunk P0:
                                   ----------
    ┌─────────────────┐
    │ Disable Trunking│           Link Speed:  128 Gb
      2-Lane Trunk
      4-Lane Trunk               Phys. Link 0:Active
                                 Phys. Link 1:Active
                                 Phys. Link 2:Active
                                 Phys. Link 3:Active


             Enter <Esc> to Previous Menu
           <↑/↓> to Highlight, <Enter> to Select
```

5.7.4 Changing the Link Speed

NOTE: Supported link speeds are specific to the adapter model and the optics installed. Examples of LPe35000-series and LPe36000-series adapters are shown in Figure 69.

The default link speed is Auto Select, which automatically selects the link speed based on the adapter model.

NOTE: The following information pertains to the link speed:

- When trunking is enabled, 8Gb/s link speed is not available because it is not supported. When the link speed is changed on a trunk port, all physical links of the trunk port are forced. The new maximum aggregate link speed is shown in the trunking selection after reboot.
- The maximum aggregate link speed is always shown even if some ports are inactive.

Possible link speeds (depending upon your adapter model):
- Auto Select – default
- 8Gb/s
- 16Gb/s
- 32Gb/s
- 64Gb/s

NOTE: Some older devices do not support auto-link speed detection; therefore, the link speed must be specified to match the speed of those devices.

To change the adapter's link speed, perform these steps:

1. From the main configuration menu (Figure 45), select **Configure Advanced Firmware Parameters** and press **Enter**.

   The adapter firmware parameters configuration menu is displayed (Figure 64).

2. Select **Link Speed Selection** and press **Enter**.

   The **Link Speed** menu (Figure 69) is displayed.

Figure 69: Link Speed Menu for LPe35000-Series and LPe36000-Series Adapters

3. Use the up arrow and down arrow keys to select the desired link speed you would like to use.

4. Press **Enter** to accept the new link speed.

5. Press **Esc** until you exit the Emulex FC BIOS utility.

6. Reboot the system.

5.7.5 Enabling or Disabling 16G Forward Error Correction

FEC improves the reliability of FC links through the recovery from bit errors that occur in high-speed networks. FEC helps minimize or avoid data-stream errors that can lead to application performance degradation or outages.

This option enables or disables FEC for 16Gb/s link speed. For link speeds 32Gb/s and higher, FEC is always enabled. The default setting is enabled.

To enable or disable FEC, perform these steps:

1. From the main configuration menu (Figure 45), select **Configure Adapter Firmware Parameters** and press **Enter**.

    The adapter firmware parameters configuration menu is displayed (Figure 64).

2. Select **16G Forward Error Correction** and press **Enter**.

    The **16G Forward Error Correction** window (Figure 70) is displayed.

Figure 70: 16G Forward Error Correction Window

3. Select the desired setting.

4. Press **Enter** to accept the new setting.

5. Press **Esc** until you exit the Emulex FC BIOS utility.

6. Reboot the system.

5.8 Resetting the Adapter to Default Values

NOTE: Resetting the adapter to default values clears all entries made while configuring the boot device, as detailed in Section 5.5, Configuring Boot Devices.

The Emulex FC BIOS utility enables you to reset boot code parameters to their factory default settings. These defaults are listed in Table 5.

Table 5: Boot Code Defaults

Parameter	Default	Valid Values
`Boot from SAN`	Disabled	■ Enabled ■ Disabled
`EDD 3.0`	EDD 3.0	■ Disabled (EDD 2.1) ■ Enabled (EDD 3.0)
`PLOGI Retry Timer`	Disabled	■ Disabled ■ 50 ms ■ 100 ms ■ 200 ms
`Spinup Delay`	Disabled	■ Enabled ■ Disabled
`Auto Scan`	Disabled	■ Enabled ■ Disabled
`Start Unit`	Disabled	■ Enabled ■ Disabled
`Environmental Variable`	Disabled	■ Enabled ■ Disabled
`Auto Boot Sector`	Disabled	■ Enabled ■ Disabled
`Brocade FA-PWWN`	**NOTE:** FA-PWWN is no longer supported; the feature will be disabled in a future release.	—
`Trunking`	Disabled	■ 4-Lane Trunk ■ 2-Lane Trunk ■ Disabled
`Topology`	Point-to-point	■ FC-AL ■ Point-to-Point
`Link Speed`	AutoSelect (highest successfully negotiated speed)	■ 8Gb/s ■ 16Gb/s ■ 32Gb/s ■ 64Gb/s **NOTE:** Adapters with a 32G SFP support link speeds 32Gb/s, 16Gb/s, and 8Gb/s. Adapters with a 64G SFP support link speeds 64Gb/s, 32Gb/s, and 16Gb/s.
`16G Forward Error Correction`	Enabled	■ Enabled ■ Disabled

To reset parameters to their factory default values, perform these steps:

1. On the main configuration menu (Figure 45), select **Reset Adapter Defaults** and press **Enter**.

   A window similar to Figure 71 asks if you want to reset to the default settings.

Figure 71: Reset Adapter Defaults Window
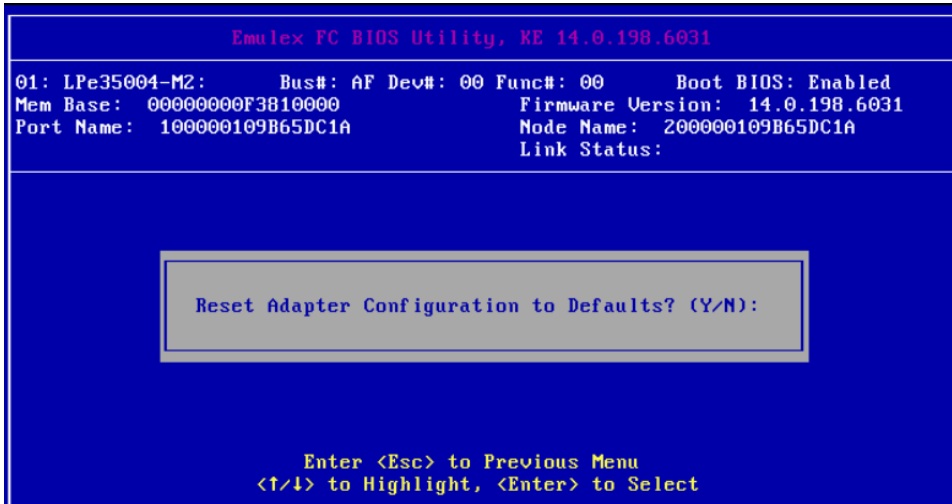
2. Press **Y**.

   All settings revert to their factory default values.

5.9 Using Multipath BFS

Multi-boot BIOS is in compliance with the BBS. The system must have a multi-boot system BIOS to use this feature. A multi-boot system BIOS allows you to select any boot disk in the system BIOS setup menu. The boot disk can be an FC drive, a SCSI drive, an IDE drive, a USB device, or a drive. The Emulex BIOS supplies up to eight drives to the system BIOS menu. The multi-boot system BIOS can override the FC drive that is selected in the Emulex FC BIOS utility.

If your system supports multi-boot BBS, the local boot disk (drive C) is the first multi-boot entry on the system BIOS setup menu. The list of entries is determined by the list of configured boot entries in the Emulex FC BIOS utility. For example:

```
Adapter 1: boot_entry0, boot_entry1
Adapter 2: boot_entry2, boot_entry3
```

The order of boot entries exported to multi-boot (BBS) is as follows:

```
boot_entry0, boot_entry1, boot_entry2, and boot_entry3.
```

However, multi-boot allows changing the boot order in the server BIOS, which allows any disk to become the C drive.

Chapter 6: Troubleshooting

Circumstances exist in which your system might operate in an unexpected manner. This chapter describes several of these circumstances and offers one or more workarounds for each situation.

6.1 x86 Boot

The following table lists x86 Boot issues and their resolutions.

Table 6: x86 Boot Troubleshooting

Issue	Situation	Resolution
The bootup message does not appear as the system boots	You want to access the BIOS utility, but the bootup message does not appear.	Make sure that x86 Boot code has been loaded and enabled.
`Retry This Adapter` message	The message `Retry This Adapter` appears during BIOS scanning.	Check the hardware configuration, or reconfigure the adapter BIOS using the Emulex FC BIOS utility.
`Cannot Mount Root File System` message (Solaris SFS driver)	The message `Cannot Mount Root File System` appears during bootup.	Make sure the correct storage device is identified in the `scsi_vhci.conf` file.
`Cannot Find UNIX Kernel` message (Solaris SFS driver)	The message `Cannot Find UNIX Kernel` appears during bootup.	Set up the correct LUN to boot in the Emulex FC BIOS utility. The correct LUN can be seen at the end of the Device Address line when you issue a `luxadm display <device>` command. Refer to the luxadm documentation from Oracle for more information.
`No Such Partition` message (Solaris SFS driver)	The message `No Such Partition` appears during bootup.	Make sure the correct boot device is selected at the GRUB menu. Refer to the GRUB documentation from Oracle and the `/boot/grub/menu.lst` for more information.
`Link is down due to Unsupported Optic Installed` or `Unsupported Optic`	These messages appear while scanning for targets or configuring boot devices in the Emulex FC BootBIOS or during a POST. `Link is down due to Unsupported Optic Installed` is displayed, or in the **Link Speed Selection** menu, `Unsupported Optic Installed` is displayed, or in the **Link Status** field, `Unsupported Optic` is displayed.	Replace any unsupported optics with supported optics.

Table 6: x86 Boot Troubleshooting

Issue	Situation	Resolution
Multipath BFS installation fails.	On OpenEuler 20.03 SP1/SP3 LTS, KylinV10 SP1/SP2, UniontechOS V20 1050E, and FusionOS 22 installing the operating system via FC BFS with multipath enabled, fails.	Perform the following steps: 1. Install the operating system in single path. 2. Enable multipath from the operating system 3. Add the boot device to the multipath and rebuild `initramfs`.

6.2 UEFI Boot

The following table lists UEFI Boot issues and their resolutions.

Table 7: UEFI Boot Troubleshooting

Issue	Situation	Resolution

`Link is down due to Unsupported Optic Installed`
or
`Unsupported Optic`
	These messages appear while scanning for targets or configuring boot devices in the UEFIHII or during a POST.
	`Link is down due to Unsupported Optic Installed` is displayed,
	or
	in the **Link Speed Selection** menu,
	`Unsupported Optic Installed`
	is displayed,
	or
	in the **Link Status** field,
	`Unsupported Optic`
	is displayed.
		Replace any unsupported optics with supported optics.

You are not able to perform a UEFI secure boot	UEFI secure boot fails	For the UEFI secure boot to function, the following requirements must be met: ■ Server hardware support must be enabled. ■ The operating system must support UEFI secure boot. ■ The UEFI driver must support UEFI secure boot (UEFI signed). ■ The operating system device driver must support UEFI secure boot (operating system-specific).
NVMe storage is not found, and the UEFI boot menu does not boot the operating system.	If the operating system was installed using NVMe over FC BFS, and you are updating the firmware from version 12.8.xxx.xx to version 14.0.xxx.xx, NVMe storage might not be found, and the UEFI boot menu might not boot the operating system.	Enable the **Generate and use a Port Based HostNQN** option.

Table 7: UEFI Boot Troubleshooting

Issue	Situation	Resolution
NVMe storage is not found, and the UEFI boot menu does not boot the operating system.	If the operating system was installed using NVMe over FC BFS, and you are updating the firmware from version 14.0.xxx.xx to version 12.8.xxx.xx, NVMe storage might not be found, and the UEFI boot menu might not boot the operating system.	Update the SAN storage configuration and add an adapter WWPN-based HostNQN to the storage array.
Multipath BFS installation fails.	On OpenEuler 20.03 SP1/SP3 LTS, KylinV10 SP1/SP2, UniontechOS V20 1050E, and FusionOS 22 installing the operating system via FC BFS with multipath enabled, fails.	Perform the following steps: 1. Install the operating system in single path. 2. Enable multipath from the operating system 3. Add the boot device to the multipath and rebuild `initramfs`.

Appendix A: License Notices

A.1 Secure Hash Algorithm (SHA-1) Notice

```
/*
 * Written by Aaron D. Gifford <me@aarongifford.com>
 *
 * Copyright 1998, 2000 Aaron D. Gifford. All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 * 1. Redistributions of source code must retain the above copyright
 *    notice, this list of conditions and the following disclaimer.
 * 2. Redistributions in binary form must reproduce the above copyright
 *    notice, this list of conditions and the following disclaimer in the
 *    documentation and/or other materials provided with the distribution.
 * 3. Neither the name of the copyright holder nor the names of contributors
 *    may be used to endorse or promote products derived from this software
 *    without specific prior written permission.
 *
 * THIS SOFTWARE IS PROVIDED BY THE AUTHOR(S) AND CONTRIBUTORS ``AS IS'' AND
 * ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
 * ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR(S) OR CONTRIBUTORS BE LIABLE
 * FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
 * DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
 * OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
 * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
 * LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
 * OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
 * SUCH DAMAGE.
 */
```

A.2 Secure Hash Algorithm (SHA-2) Notice

```
/*
 * AUTHOR:Aaron D. Gifford - http://www.aarongifford.com/
 *
 * Copyright (c) 2000-2001, Aaron D. Gifford
 * All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 * 1. Redistributions of source code must retain the above copyright
 *    notice, this list of conditions and the following disclaimer.
 * 2. Redistributions in binary form must reproduce the above copyright
 *    notice, this list of conditions and the following disclaimer in the
 *    documentation and/or other materials provided with the distribution.
 * 3. Neither the name of the copyright holder nor the names of contributors
 *    may be used to endorse or promote products derived from this software
 *    without specific prior written permission.
 *
 * THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTOR(S) ``AS IS'' AND
 * ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
 * ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTOR(S) BE LIABLE
 * FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
 * DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
 * OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
 * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
 * LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
 * OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
 * SUCH DAMAGE.
 *
 */
```

A.3 OpenSSL Notice

This is a copy of the current LICENSE file inside the CVS repository.

```
LICENSE ISSUES
==============
The OpenSSL toolkit stays under a dual license, i.e. both the conditions of the OpenSSL License and the original SSLeay license
apply to the toolkit. See below for the actual license texts. Actually both licenses are BSD-style Open Source licenses. In case
of any license issues related to OpenSSL please contact openssl-core@openssl.org.
OpenSSL License
---------------
/* ====================================================================
 * Copyright (c) 1998-2011 The OpenSSL Project. All rights reserved.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 *
 * 1. Redistributions of source code must retain the above copyright
 *    notice, this list of conditions and the following disclaimer.
 *
 * 2. Redistributions in binary form must reproduce the above copyright
 *    notice, this list of conditions and the following disclaimer in
 *    the documentation and/or other materials provided with the
 *    distribution.
 *
 * 3. All advertising materials mentioning features or use of this
 *    software must display the following acknowledgment:
 *    "This product includes software developed by the OpenSSL Project
 *    for use in the OpenSSL Toolkit. (http://www.openssl.org/)"
 *
 * 4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to
 *    endorse or promote products derived from this software without
 *    prior written permission. For written permission, please contact
 *    openssl-core@openssl.org.
 *
 * 5. Products derived from this software may not be called "OpenSSL"
 *    nor may "OpenSSL" appear in their names without prior written
 *    permission of the OpenSSL Project.
 *
 * 6. Redistributions of any form whatsoever must retain the following
 *    acknowledgment:
 *    "This product includes software developed by the OpenSSL Project
 *    for use in the OpenSSL Toolkit (http://www.openssl.org/)"
 *
 * THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT ``AS IS'' AND ANY
 * EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR
 * PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR
 * ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL,
 * SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT
 * NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES;
 * LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
 * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT,
 * STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE)
 * ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED
 * OF THE POSSIBILITY OF SUCH DAMAGE.
 * ====================================================================
 *
 * This product includes cryptographic software written by Eric Young
 * (eay@cryptsoft.com). This product includes software written by Tim
 * Hudson (tjh@cryptsoft.com).
 *
 */
Original SSLeay License
-----------------------
/* Copyright (C) 1995-1998 Eric Young (eay@cryptsoft.com)
 * All rights reserved.
 *
 * This package is an SSL implementation written
 * by Eric Young (eay@cryptsoft.com).
```

```
 * The implementation was written so as to conform with Netscapes SSL.
 *
 * This library is free for commercial and non-commercial use as long as
 * the following conditions are adhered to. The following conditions
 * apply to all code found in this distribution, be it the RC4, RSA,
 * lhash, DES, etc., code; not just the SSL code. The SSL documentation
 * included with this distribution is covered by the same copyright terms
 * except that the holder is Tim Hudson (tjh@cryptsoft.com).
 *
 * Copyright remains Eric Young's, and as such any Copyright notices in
 * the code are not to be removed.
 * If this package is used in a product, Eric Young should be given attribution
 * as the author of the parts of the library used.
 * This can be in the form of a textual message at program startup or
 * in documentation (online or textual) provided with the package.
 *
 * Redistribution and use in source and binary forms, with or without
 * modification, are permitted provided that the following conditions
 * are met:
 * 1. Redistributions of source code must retain the copyright
 *    notice, this list of conditions and the following disclaimer.
 * 2. Redistributions in binary form must reproduce the above copyright
 *    notice, this list of conditions and the following disclaimer in the
 *    documentation and/or other materials provided with the distribution.
 * 3. All advertising materials mentioning features or use of this software
 *    must display the following acknowledgment:
 *    "This product includes cryptographic software written by
 *    Eric Young (eay@cryptsoft.com)"
 *    The word 'cryptographic' can be left out if the routines from the library
 *    being used are not cryptographic related :-).
 * 4. If you include any Windows specific code (or a derivative thereof) from
 *    the apps directory (application code) you must include an acknowledgment:
 *    "This product includes software written by Tim Hudson (tjh@cryptsoft.com)"
 *
 * THIS SOFTWARE IS PROVIDED BY ERIC YOUNG ``AS IS'' AND
 * ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE
 * IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE
 * ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR CONTRIBUTORS BE LIABLE
 * FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL
 * DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS
 * OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION)
 * HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT
 * LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY
 * OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF
 * SUCH DAMAGE.
 *
 * The licence and distribution terms for any publically available version or
 * derivative of this code cannot be changed. i.e. this code cannot simply be
 * copied and put under another distribution licence
 * [including the GNU Public Licence.]
```

A.4 mbed TLS

```
Apache License
Version 2.0, January 2004
http://www.apache.org/licenses/
TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION
1. Definitions.
   "License" shall mean the terms and conditions for use, reproduction,
   and distribution as defined by Sections 1 through 9 of this document.
   "Licensor" shall mean the copyright owner or entity authorized by
   the copyright owner that is granting the License.
   "Legal Entity" shall mean the union of the acting entity and all
   other entities that control, are controlled by, or are under common
   control with that entity. For the purposes of this definition,
   "control" means (i) the power, direct or indirect, to cause the
   direction or management of such entity, whether by contract or
   otherwise, or (ii) ownership of fifty percent (50%) or more of the
   outstanding shares, or (iii) beneficial ownership of such entity.
   "You" (or "Your") shall mean an individual or Legal Entity
```

```
    exercising permissions granted by this License.
    "Source" form shall mean the preferred form for making modifications,
    including but not limited to software source code, documentation
    source, and configuration files.
    "Object" form shall mean any form resulting from mechanical
    transformation or translation of a Source form, including but
    not limited to compiled object code, generated documentation,
    and conversions to other media types.
    "Work" shall mean the work of authorship, whether in Source or
    Object form, made available under the License, as indicated by a
    copyright notice that is included in or attached to the work
    (an example is provided in the Appendix below).
    "Derivative Works" shall mean any work, whether in Source or Object
    form, that is based on (or derived from) the Work and for which the
    editorial revisions, annotations, elaborations, or other modifications
    represent, as a whole, an original work of authorship. For the purposes
    of this License, Derivative Works shall not include works that remain
    separable from, or merely link (or bind by name) to the interfaces of,
    the Work and Derivative Works thereof.
    "Contribution" shall mean any work of authorship, including
    the original version of the Work and any modifications or additions
    to that Work or Derivative Works thereof, that is intentionally
    submitted to Licensor for inclusion in the Work by the copyright owner
    or by an individual or Legal Entity authorized to submit on behalf of
    the copyright owner. For the purposes of this definition, "submitted"
    means any form of electronic, verbal, or written communication sent
    to the Licensor or its representatives, including but not limited to
    communication on electronic mailing lists, source code control systems,
    and issue tracking systems that are managed by, or on behalf of, the
    Licensor for the purpose of discussing and improving the Work, but
    excluding communication that is conspicuously marked or otherwise
    designated in writing by the copyright owner as "Not a Contribution."
    "Contributor" shall mean Licensor and any individual or Legal Entity
    on behalf of whom a Contribution has been received by Licensor and
    subsequently incorporated within the Work.
2.  Grant of Copyright License. Subject to the terms and conditions of
    this License, each Contributor hereby grants to You a perpetual,
    worldwide, non-exclusive, no-charge, royalty-free, irrevocable
    copyright license to reproduce, prepare Derivative Works of,
    publicly display, publicly perform, sublicense, and distribute the
    Work and such Derivative Works in Source or Object form.
3.  Grant of Patent License. Subject to the terms and conditions of
    this License, each Contributor hereby grants to You a perpetual,
    worldwide, non-exclusive, no-charge, royalty-free, irrevocable
    (except as stated in this section) patent license to make, have made,
    use, offer to sell, sell, import, and otherwise transfer the Work,
    where such license applies only to those patent claims licensable
    by such Contributor that are necessarily infringed by their
    Contribution(s) alone or by combination of their Contribution(s)
    with the Work to which such Contribution(s) was submitted. If You
    institute patent litigation against any entity (including a
    cross-claim or counterclaim in a lawsuit) alleging that the Work
    or a Contribution incorporated within the Work constitutes direct
    or contributory patent infringement, then any patent licenses
    granted to You under this License for that Work shall terminate
    as of the date such litigation is filed.
4.  Redistribution. You may reproduce and distribute copies of the
    Work or Derivative Works thereof in any medium, with or without
    modifications, and in Source or Object form, provided that You
    meet the following conditions:
    (a) You must give any other recipients of the Work or
    Derivative Works a copy of this License; and
    (b) You must cause any modified files to carry prominent notices
    stating that You changed the files; and
    (c) You must retain, in the Source form of any Derivative Works
    that You distribute, all copyright, patent, trademark, and
    attribution notices from the Source form of the Work,
    excluding those notices that do not pertain to any part of
    the Derivative Works; and
    (d) If the Work includes a "NOTICE" text file as part of its
    distribution, then any Derivative Works that You distribute must
```

include a readable copy of the attribution notices contained
within such NOTICE file, excluding those notices that do not
pertain to any part of the Derivative Works, in at least one
of the following places: within a NOTICE text file distributed
as part of the Derivative Works; within the Source form or
documentation, if provided along with the Derivative Works; or,
within a display generated by the Derivative Works, if and
wherever such third-party notices normally appear. The contents
of the NOTICE file are for informational purposes only and
do not modify the License. You may add Your own attribution
notices within Derivative Works that You distribute, alongside
or as an addendum to the NOTICE text from the Work, provided
that such additional attribution notices cannot be construed
as modifying the License.

You may add Your own copyright statement to Your modifications and
may provide additional or different license terms and conditions
for use, reproduction, or distribution of Your modifications, or
for any such Derivative Works as a whole, provided Your use,
reproduction, and distribution of the Work otherwise complies with
the conditions stated in this License.

5. Submission of Contributions. Unless You explicitly state otherwise,
   any Contribution intentionally submitted for inclusion in the Work
   by You to the Licensor shall be under the terms and conditions of
   this License, without any additional terms or conditions.
   Notwithstanding the above, nothing herein shall supersede or modify
   the terms of any separate license agreement you may have executed
   with Licensor regarding such Contributions.

6. Trademarks. This License does not grant permission to use the trade
   names, trademarks, service marks, or product names of the Licensor,
   except as required for reasonable and customary use in describing the
   origin of the Work and reproducing the content of the NOTICE file.

7. Disclaimer of Warranty. Unless required by applicable law or
   agreed to in writing, Licensor provides the Work (and each
   Contributor provides its Contributions) on an "AS IS" BASIS,
   WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or
   implied, including, without limitation, any warranties or conditions
   of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A
   PARTICULAR PURPOSE. You are solely responsible for determining the
   appropriateness of using or redistributing the Work and assume any
   risks associated with Your exercise of permissions under this License.

8. Limitation of Liability. In no event and under no legal theory,
   whether in tort (including negligence), contract, or otherwise,
   unless required by applicable law (such as deliberate and grossly
   negligent acts) or agreed to in writing, shall any Contributor be
   liable to You for damages, including any direct, indirect, special,
   incidental, or consequential damages of any character arising as a
   result of this License or out of the use or inability to use the
   Work (including but not limited to damages for loss of goodwill,
   work stoppage, computer failure or malfunction, or any and all
   other commercial damages or losses), even if such Contributor
   has been advised of the possibility of such damages.

9. Accepting Warranty or Additional Liability. While redistributing
   the Work or Derivative Works thereof, You may choose to offer,
   and charge a fee for, acceptance of support, warranty, indemnity,
   or other liability obligations and/or rights consistent with this
   License. However, in accepting such obligations, You may act only
   on Your own behalf and on Your sole responsibility, not on behalf
   of any other Contributor, and only if You agree to indemnify,
   defend, and hold each Contributor harmless for any liability
   incurred by, or claims asserted against, such Contributor by reason
   of your accepting any such warranty or additional liability.

END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

   To apply the Apache License to your work, attach the following
   boilerplate notice, with the fields enclosed by brackets "[]"
   replaced with your own identifying information. (Don't include
   the brackets!) The text should be enclosed in the appropriate
   comment syntax for the file format. We also recommend that a
   file or class name and description of purpose be included on the
   same "printed page" as the copyright notice for easier
   identification within third-party archives.

```
Copyright [yyyy] [name of copyright owner]
Licensed under the Apache License, Version 2.0 (the "License");
you may not use this file except in compliance with the License.
You may obtain a copy of the License at
    http://www.apache.org/licenses/LICENSE-2.0
Unless required by applicable law or agreed to in writing, software
distributed under the License is distributed on an "AS IS" BASIS,
WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.
See the License for the specific language governing permissions and
limitations under the License.
```