BROADCOM®

Release Notes

Emulex® Drivers for Linux

Release 12.4.243.20

Purpose and Contact Information

These release notes describe the new features, resolved issues, FC and NVMe driver known issues, and FC and NVMe technical tips associated with this release of the Emulex® drivers for Linux.

For the latest product documentation, go to www.broadcom.com. If you have questions or require additional information, contact an authorized Broadcom® Technical Support representative at ecd-tech.support@broadcom.com.

New Features (Version 12.4.243.20)

1. Adds support for RHEL 8.1.
2. Adds support for AMD Epyc 7002-series processors (PCIe 4.0).

New Features (Version 12.4.243.9)

3. Adds support for sequence-level error recovery (SLER) for NVMe over FC. (LPe31000-series, LPe32000-series, and LPe35000-series adapters only)
4. Adds support for RHEL 7.7.
5. Adds internal and external loopback test support on trunked ports. (LPe35000-series adapters only)
6. Adds driver support for performance scaling across sockets.
7. Emulex-signed RPM files are provided for SUSE drivers beginning with this release.
8. Discontinues support for the following operating systems:
   - RHEL 6.9
   - RHEL 6.10
   - RHEL 7.4
   - Citrix XenServer 6.x
   - Citrix XenServer 7.0

Resolved Issues (Version 12.4.243.20)

This release resolves an issue on LPe35000-series adapters with out-of-box Linux drivers that could cause a kernel panic.

Resolved Issues (Version 12.4.243.9)

1. On LPe35000-series adapters with trunking enabled, if I/O is running on NVMe and FCP drivers and a LIP (Loop Initialization Protocol) is issued, a call trace is not generated and the port is not reset.
2. On certain SLES 15 servers, after completion of an internal loopback test on LPe31000-series and LPe32000-series adapters, there is no longer any delay in the link-up event.
3. On RHEL 8.0, call traces are no longer observed while running I/O on 4K block size namespaces.
4. On SLES 15 SP1, the out-of-box `lpfc` driver correctly defaults to setting the `lpfc_enable_fc4_type` driver parameter to 3 to enable NVMe support.
5. The issue of the SLES 12 SP4 driver with MultiQueue enabled not resuming I/Os after a link toggle or link reset is resolved.
6. The issue of invalid link fault code messages being generated, followed by mailbox command failures, is resolved.
7. The issue of NVMe I/O stopping if the initiator server administrator continually resets the NVMe controllers for more than one hour while the system is under an I/O load, on Linux operating systems, is fixed.
8. On RHEL 8.0, the issue of I/O not resuming on NVMe namespaces after a storage node failover, is resolved.
9. On SLES 15 SP1, `nvme-cli` auto discovery no longer fails.

FC Driver Known Issues (Version 12.4.243.20 and Version 12.4.243.9)

1. Known issue related to software migration.

   Beginning with software release 11.2, FC HBAs and OneConnect® adapters have independent software kits. Before updating earlier drivers and applications to the software in release 12.4.2, refer to the *Emulex Software Kit Migration User Guide* for special instructions and considerations for using the 11.2 and later software kits for FC HBAs and OneConnect adapters.

2. PCI Hot Plug might cause applications, such as the Emulex OneCommand® Manager application or third-party applications that use the Emulex libraries (for example, an HBA API), to malfunction.

   Workaround

   a. Stop all applications that are accessing the FC HBA API interface (Emulex OneCommand Manager application or third-party applications) before performing PCI Hot Plug of an FC adapter.

   b. Use the following command to stop the Emulex OneCommand Manager application:
   `#/usr/sbin/ocmanager/stop_ocmanager`

   c. After performing PCI Hot Plug of the adapter, restart the applications.

3. SCSI errors might occur on deletion of vPorts or PCI Hot Unplug.

   On occasion, the kernel might report SCSI errors when deleting vPorts through the `sysfs` interface or performing a PCI Hot Unplug of an Emulex adapter:

   ```
   kernel: Synchronizing SCSI cache for disk
   kernel: FAILED
   ```

   Or:

   ```
   SCSI error: return code = 0x00010000
   ```

   Workaround

   None. Ignore these messages; they do not indicate a functional failure.

4. An issue exists while deleting vPorts when devices are in use.

   Emulex provides management utilities that allow you to delete vPorts. However, no mechanism exists for the FC driver to detect whether devices accessed through that vPort are in use. This situation means that you can delete a vPort when devices accessible through the vPort are mounted or when I/O is outstanding to the device. When file systems are mounted on vPorts and vPorts are deleted, the file systems still appear to be mounted; however, they are inaccessible.

   Workaround

   Before deleting vPorts, you must prepare the system affected by the vPort deletion accordingly, by unmounting all the devices accessible through the vPorts and ensuring that there is no outstanding I/O.

5. A devloss timeout occurs after swapping ports.

   The driver might not finish discovery when two initiator ports are swapped. This situation causes all devices accessible through one or both of these initiator ports to time out and all I/O to fail.

   Workaround

   Do one of the following:

   – When swapping cables, replace each cable, one at a time, and allow discovery to finish before replacing the next cable. To determine if discovery is finished, read the `state sysfs` parameter.
   – When swapping cables, allow the devloss timeout to occur before replacing the cables (this action fails all outstanding I/O).

6. The Linux Loader (LILO) boot loader is not supported on i386 and x86_64 architectures.

   The LILO boot loader on i386 and x86_64 architectures is not supported for this driver. If the LILO boot loader is used, after the FC driver package is installed and upon reboot, an incorrect initial ramdisk is used, and the system might not boot correctly.

   Workaround

   The boot loader that is supported with this driver is Grand Unified Bootloader (GRUB), which is the default boot loader for most Linux distributions. LILO is an older boot loader used on i386 and x86_64 architectures only. GRUB works correctly with the driver package installation script.

7. An error message `Failed to issue SLI_CONFIG ext-buffer` might be displayed when multiple queue operations are performed.

   Workaround

   During firmware update operations, do not perform queue operations, such as resetting the adapter, the bus, the target, or the host.

8. Revision A of the FOIT AFCT-57F3TMZ-ELX (16GFC longwave optic transceiver) does not support D_Port for MDS Diagnostic for Cisco switches.

   Workaround

   None.

9. Neither Revision A nor Revision B of the FOIT AFCT-57F3TMZ-ELX (16GFC longwave optic transceiver) supports D_Port for Brocade® switches.

   Workaround

   None.

10. The `lpfc` driver fails to install on the SLES 12 SP3 GA kernel.

    Workaround

    The minimum kernel on which the `lpfc` driver can be installed is maintenance kernel 718 (4.4.126-94.22.1) from SUSE, dated April 23, 2018.

11. Boot from SAN is not supported if FC-SP-2 authentication (DHCHAP) is enabled.

    Workaround

    None.

12. Due to limitations in the FC-LS-3 RDP ELS definition, FC switches do not issue RDP commands on trunked links.

    Workaround

    None.

13. The remote switched diagnostic test will fail with a Latency Err-Drop error, if you run diagnostics on multiple HBAs simultaneously.

    Workaround

    Run diagnostics on only one port at a time.

14. On SLES 12 and SLES 15, the driver might not unload due to BTRFS.

    If a discovered SCSI device has a BTRFS partition on it, the upper layers maintain a reference count on the driver which will cause it to not be unloaded. After the SCSI layer discovers an FC device and you see the following BTRFS message:

    ```
    [ 5.502273] sd 0:0:0:1: [sdd] Attached SCSI disk
    [ 5.552774] BTRFS: device fsid 69ca50b4-c061-4f66-9ada-0287d76269a0 devid 1 transid 274 /dev/sdd2
    ```

    The BTRFS subsystem takes a reference count on the `lpfc` driver when this happens. This extra reference count might stop the administrator from unloading the `lpfc` driver. In order to unload the driver, the reference count for it must be 0:

    ```
    # cat /sys/module/lpfc/refcnt
    0
    #
    ```

15. Brocade switches using Fabric OS® version 8.2.1B and earlier might encounter the following issues with DHCHAP authentication:
    - You might not be able to configure the secret pair between the switch and the HBA.
    - When authentication is enabled on the switch, and authentication is disabled on the HBA, the switch disables the port, but it does not issue the expected status messages.
    - The Brocade switch authenticates the HBA port when authentication is disabled on the switch or when a frame is dropped.

    Workaround

    None.

16. Firmware dump is unavailable for LPe35000-series adapters using standard Emulex applications, such as the OneCommand Manager application and the OneCapture™ utility, on the following Linux inbox drivers:
    - SLES 15
    - RHEL 7.6 (except the RHEL 7.6 errata kernel-3.10.0-990.el7)
    - Ubuntu 18.04 (except the Ubuntu 18.04.3 HWE release)

    To perform a firmware dump on these inbox drivers, see FC Driver Technical Tips, item 5, in these release notes.

17. The SLES 12 SP4 driver with MultiQueue enabled does not resume I/Os after a link toggle or link reset.

    Workaround

    No workaround is available for the currently released driver version. This issue is to be fixed by SUSE in the operating system.

18. Loading a Broadcom ECD-signed driver on a system using legacy BIOS might result in an error message similar to the following:

    ```
    Request for unknown module key 'Broadcom Inc.: Emulex Connectivity Division:
    d17ecabc92cd490989959b37f05f0eda48c53895' err -11
    ```

    or

    ```
    PKCS#7 signature not signed with a trusted key
    ```

    Workaround

    None. This is a benign error message, and it can be ignored.

19. On RHEL 8.0 inbox systems, the GUI displays incorrect trunking information.

    Workaround

    Use the RHEL 8.0 out-of-box driver.

20. On some inbox Linux distributions, the `lpfc_enable_mds_diags` driver parameter cannot be enabled dynamically.

    Workaround

    Enable the `lpfc_enable_mds_diags` parameter temporarily by issuing the following commands:

    ```
    rmmod lpfc
    modprobe lpfc lpfc_enable_mds_diags=1
    ```

    After the diagnostics are complete, reload the driver without the `lpfc_enable_mds_diags` parameter.

21. Do not leave the `enable-mds-diags` driver parameter set to 1. It can cause switch diagnostic failures and OCM D_Port failures, when connected to a Brocade switch.

    Workaround

    Set the `enable-mds-diags` driver parameter to 1, only when running diagnostic tests while connected to supported Cisco switches. Disable the `enable-mds-diags` driver parameter when not running diagnostic tests.

22. In the latest inbox RHEL and SLES operating system releases, SCSI multi-queue might be enabled by default, which could dramatically increase the amount of preallocated I/O buffers. Depending on your system configuration, out-of-memory errors might occur on boot.

    Workaround

    If the system cannot boot because of memory issues, perform the following steps:

    a. Temporarily change the FC driver parameters by adding the following command to the kernel boot command line:
       `lpfc.lpfc_hdw_queue=1 lpfc.lpfc_fcp_io_sched=0`

       Refer to the operating system documentation for details.

    b. Install the latest release 12.4 out-of-box FC driver.

23. On certain Linux distributions, the inbox driver might fail to discover LUNs, and the link might remain down. This issue occurs if the firmware version used is 12.4.243.8 or later.

    Workaround

    Use the out-of-box driver version 12.4.64.0 or later for proper discovery of LUNs.

24. While installing the driver kit, if an instance of `elx-lpfc-vector-map` RPM exists, the driver kit `elx-lpfc-extras` RPM fails to install and might result in an RPM conflict error.

    Workaround

    Uninstall the `elx-lpfc-vector-map` RPM using the following commands:

    ```
    # rpm -e --allmatches elx-lpfc-vector-map --noscripts
    # rm -f /usr/lib/dracut/modules.d/89lpfc-vectormap
    # dracut -f
    ```

    Install the driver kit, after you have successfully uninstalled the `elx-lpfc-vector-map` RPM.

NVMe Driver Known Issues (Version 12.4.243.20 and Version 12.4.243.9)

1. In the latest SLES 12, SLES 15, and RHEL 7.x operating system variants, the `nvme list` command might not display all namespaces when the maximum number of namespaces (255) are created on the target.

   Workaround

   Use the `lsblk` command to display all namespaces.

2. Running the `nvmetcli clear` command on an NVMe target when the NVMe initiator port is down might cause the target system to stop responding. Make sure that the NVMe initiator port is UP and that all NVMe devices are disconnected from the initiator before running the `nvmetcli clear` command on the target.

   Workaround

   None.

3. The deliberate faulting of NVMe discovery commands (also called jamming) is not supported. If a discovery command is faulted during initial linkup or during the LIP linkup recovery time, NVMe discovery fails on the affected controller.

   Workaround

   Perform a LIP on the initiator link, and remove the condition that is faulting discovery.

4. If you are adding a subsystem dynamically on a target, you must issue an LIP on the initiator host port that is bound to the target. If you are adding new namespaces dynamically to existing subsystems, you must perform a manual scan on the initiator. See NVMe Driver Technical Tips for more information.

5. Configuring initiator ports as NVMe over FC is not supported on ports that are configured as Fabric assigned Port World Wide Names (FA-PWWNs).

6. Target port reset using `hbacmd` or the Emulex OneCommand Manager application is not supported in this release.

7. Dynamic NVMe target management (including reconfiguration) and dynamic initiator rescan are not supported. Any reconfiguration requires a target server reboot.

   Workaround

   After the target server reboots and the configuration is activated, connected initiators might require a link bounce to rescan the target.

8. Unloading the FCP driver using the `modprobe -r` command might cause issues on the initiator before NVMe devices are disconnected.

   Workaround

   Unload the driver with `rmmod lpfc`, or if `modprobe –r lpfc` is required, wait for the device loss period of 60 seconds before unloading the driver.

9. On systems with version 1.7 of the `nvme-cli` installed, connection attempts might fail, and you might not see NVMe-FC storage. Console log messages and Systemd log messages indicate the failure to connect.

   Workaround

   Do one of the following:

   – Contact the operating system vendor to obtain a newer version of the `nvme-cli` utility.
   – Download and use the latest `nvme-cli` utility at https://github.com/linux-nvme/nvme-cli.

10. An I/O error can occur while resetting an NVMe controller. Error messages similar to the following might be shown, and the I/O request could fail.

```
nvme nvme0: NVME-FC{0}: controller reset complete
print_req_error: I/O error, dev nvme0c2n3, sector 11715.
```

Workaround

Upgrade to the latest operating system kernel version.

The minimum supported kernel versions are:

- RHEL7.6 3.10.0-930.el7 and later
- SLES12 SP3 4.4.155-94.50.1 and later
- SLES12 SP4 GA kernel and later
- SLES15 maintenance update kernel 4.12.14-25.16.1 and later

FC Driver Technical Tips

1. Locked optics are supported on Emulex LPe31000-series and LPe32000-series adapters.

   The adapters perform the following operations:

   - Detect and enable both Broadcom or Emulex certified SFP optics.
   - For firmware revision 11.x and later, unqualified optics are disabled, the link is down, and an error message is written to the log file.
   - The `lpfc` out-of-box driver revision 11.x and later show this message, and the link will not come up.
     `"3176 Port Name [wwpn] Unqualified optics - Replace with Avago optics for Warranty and Technical support"`

   When a 32 Gb/s optic is installed in an Emulex LPe31000-series, LPe32000-series, or LPe35000-series adapter, the link supports 32 Gb/s, 16 Gb/s, and 8 Gb/s speeds.

   When a 16 Gb/s optic is installed in an Emulex LPe31000-series or LPe32000-series adapter, the link supports 16 Gb/s, 8 Gb/s, and 4 Gb/s speeds.

2. For secure boot of the RHEL operating system, download the public key file for the RHEL operating system from www.broadcom.com and place it in the system's key ring before installing the signed `lpfc` driver for the RHEL operating system. Refer to the Red Hat documentation for instructions on adding a kernel module to the system.

3. D_Port and FA-PWWN cannot be enabled simultaneously.

   Workaround

   If D_Port is enabled and you want to enable FA-PWWN, you must first disable D_Port. If FA-PWWN is enabled and you want to enable D_Port, you must first disable FA-PWWN.

4. D_Port and DHCHAP cannot be enabled simultaneously.

   Workaround

   If D_Port is enabled and you want to enable DHCHAP, you must first disable D_Port. If DHCHAP is enabled and you want to enable D_Port, you must first disable DHCHAP.

5. To perform a firmware dump on the inbox drivers listed in FC Driver Known Issues (Version 12.4.243.20 and Version 12.4.243.9), item 16, perform the following steps:

   a. Ensure that the latest Emulex OneCommand Manager core application kit for Linux is installed on the host.

   b. Download the LPe35000-series HBA FW Dump Tool from the Broadcom website, at www.broadcom.com.

   c. Type the following commands to extract and install the tool:

   ```
   tar xzf elx-lpe-35000-fw-dump-tool-<version>-ds-1.tar.gz
   cd elx-lpe-35000-fw-dump-tool-<version>-ds-1/
   rpm -ivh elx-lpe-35000-fw-dump-tool-<version>.<platform>.rpm
   ```

A window similar to the following is displayed.

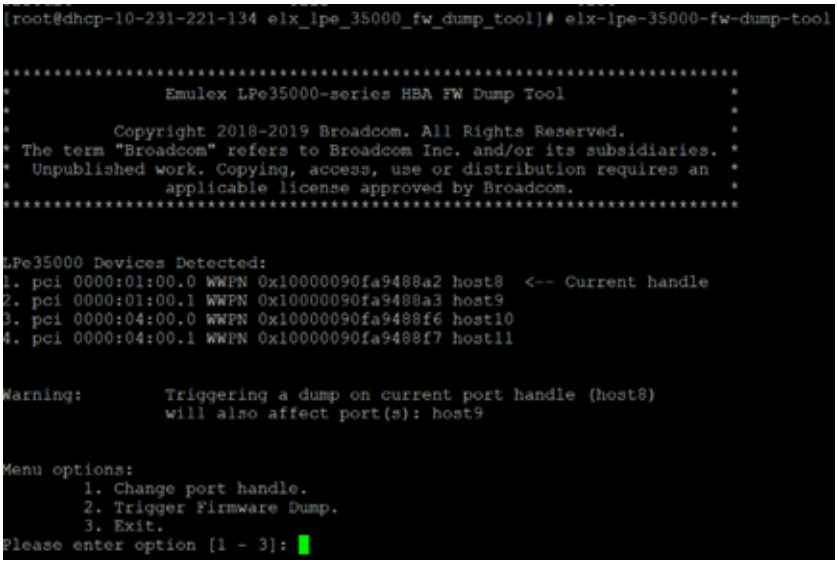

d. Type **1** and press **Enter** to select the port for which you want to perform the firmware dump.

A window similar to the following is displayed.

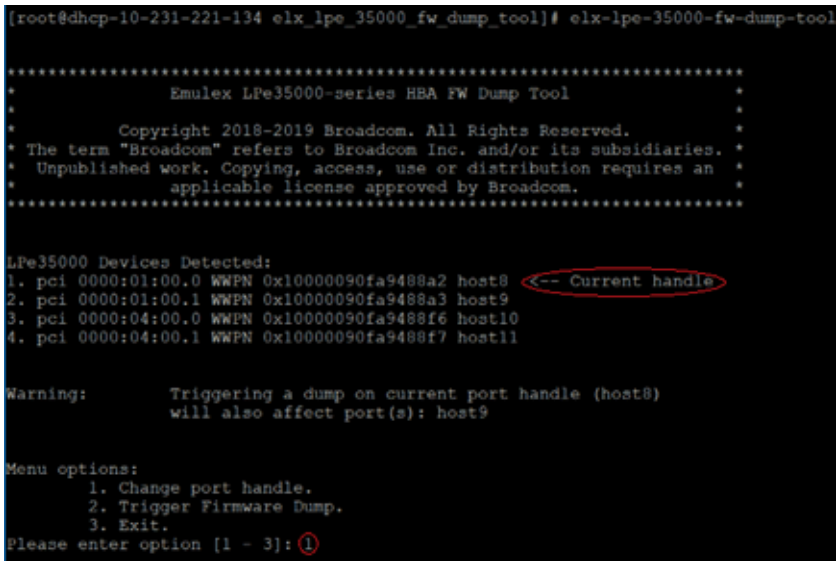

NOTE: If initiating a firmware dump on a specific port will also affect another port on the same adapter ASIC, a warning message notifies you of this fact.

e. Type the number of the port for which you want to perform the firmware dump, and press **Enter**.

A window similar to the following is displayed.

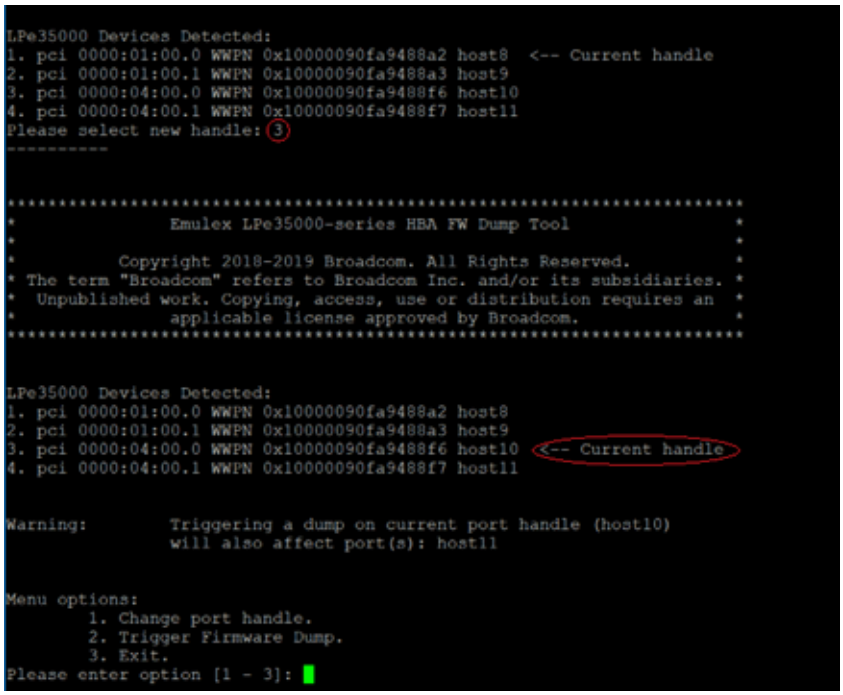

f. When the desired port is identified as the current port, type **2** and press **Enter** to initiate the firmware dump on the current port.

g. If the firmware dump is successful, the firmware dump tool closes automatically.

h. Restart the `elxhbamgrd` service by typing the following command:

```
service elxhbamgrd restart
```

This allows applications such as the Emulex OneCommand Manager application to be refreshed.

By default, the firmware dump file is available on the host in the `/var/opt/Emulex/ocmanager/Dump/` directory after the `elxhbamgrd` service has been restarted.

On SLES 12 SP4 or SLES 15, initiating a firmware dump might result in an error message similar to the following:

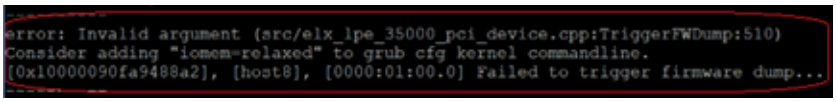

If this message appears, do the following:

a. Add the following kernel command line to the host grub configuration file:

```
iomem=relaxed
```

b. Reboot the host.

c. Repeat the process to initiate a firmware dump.

NOTE: For the Ubuntu x86_64 architecture operating system, the OneCommand Manager application is unavailable. You must use the Emulex OneCapture utility to read the firmware dump from flash. You can download and use the latest Emulex OneCapture utility for Ubuntu from the Broadcom website at www.broadcom.com.

For instructions on reading the firmware dump file, refer to the “Output File” section under “Running OneCapture on Linux, Citrix, and Solaris” in the *Emulex OneCapture Utility User Guide*.

NVMe Driver Technical Tips

1. Creation of N_Port ID Virtualization (NPIV) connections on initiator ports that are configured for NVMe over FC is not supported. However, initiator ports can connect to Fibre Channel Protocol (FCP) and NVMe targets simultaneously.
2. NVMe disks might not reconnect after a device timeout greater than 60 seconds has occurred.

   Workaround

   You must perform a manual scan or connect via the `nvme connect-all` CLI command.
3. To manually scan for targets or dynamically added subsystems, type the following command (all on one line):
   `nvme connect-all --transport=fc --host-traddr=nn-<initiator_WWNN>:pn-<initiator_WWPN> --traddr=nn-<target_WWNN>:pn-<target_WWPN>`

   where:
   - <*initiator_WWNN*> is the WWNN of the initiator, in hexadecimal.
   - <*initiator_WWPN*> is the WWPN of the initiator, in hexadecimal.
   - <*target_WWNN*> is the WWNN of the target, in hexadecimal.
   - <*target_WWPN*> is the WWPN of the target, in hexadecimal.

   For example:
   ```
   # nvme connect-all --transport=fc --host-traddr=nn-0x20000090fa942779:pn-0x10000090fa942779 --traddr=nn-0x20000090fae39706:pn-0x10000090fae39706
   ```
4. To enable autoconnect on the SLES 15 SP1 operating system, install `nvme-cli` from the operating system distribution media.
5. This release supports FC-NVMe specification version 1.19.
6. If the following files are present on the system after the operating system is installed, the operating system has already installed NVMe over FC autoconnect facilities. Do not install the Emulex autoconnect script file for inbox NVMe over FC drivers.
   - `/usr/lib/systemd/system/nvmefc-boot-connections.service`
   - `/usr/lib/systemd/system/nvmefc-connect@.service`
   - `/usr/lib/udev/rules.d/70-nvmefc-autoconnect.rules`

Broadcom, the pulse logo, Connecting everything, Avago Technologies, Avago, the A logo, Brocade, Emulex, Fabric OS, OneCapture, OneCommand, and OneConnect are among the trademarks of Broadcom and/or its affiliates in the United States, certain other countries, and/or the EU.

Copyright © 2012–2019 Broadcom. All Rights Reserved.

The term “Broadcom” refers to Broadcom Inc. and/or its subsidiaries. For more information, please visit www.broadcom.com.

Broadcom reserves the right to make changes without further notice to any products or data herein to improve reliability, function, or design. Information furnished by Broadcom is believed to be accurate and reliable. However, Broadcom does not assume any liability arising out of the application or use of this information, nor the application or use of any product or circuit described herein, neither does it convey any license under its patent rights nor the rights of others.

BROADCOM®