

Emulex Poll Mode Driver/DPDK Release Notes

Versions: 11.0.242.0

Date: March 1, 2016

Purpose and Contact Information

These release notes describe the new features, current known issues, and technical tips associated with this Broadcom[®] Emulex[®] Poll Mode Driver/DPDK release.

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 tech.support-elx@broadcom.com, 800-854-7112 (US/Canada toll free), +1 714-885-3402 (US/International), or +44 1189-772929 (Europe, Middle East, and Africa).

New Features

- Supports the Red Hat Enterprise Linux RHEL 6.6, RHEL 6.7, RHEL 7.0, and RHEL 7.1 64-bit operating systems.
- Supports DPDK versions 1.8.0 and 2.0.

Known Issue

1. **When using DPDK over SR-IOV on RHEL 7.1 systems, the virtual machine (VM) does not power on after setting the Copy host CPU configuration in the Virtual Machine Manager (VMM).**

Workaround

There are two methods you can use to set up a server for DPDK over SR-IOV.

First method:

1. Make sure you have the latest Broadcom Emulex firmware and PMD driver installed.
2. Enable SR-IOV on the host.
3. Compile DPDK 1.8.0, or 2.0 on the host.
4. Compile the PMD+Surf driver on the host.
5. Using VMM on a RHEL 7.1 or RHEL 7.2 host, create a virtual machine (VM) and install RHEL 7.1 on it.
6. After installation, turn off the VM and enable the Copy host CPU configuration option under processor settings of the guest operating system in the VMM.

Note: If you are using RHEL 7.1, the VM will not boot, and an error stating that the CPU does not support the `invts` feature will appear. Use the command `virsh edit <VM name>` to remove the line `<feature policy='require' name='invts' />` from the XML file, then reboot the VM. (This workaround is not needed for RHEL 7.2.)

Second method:

1. Make sure you have the latest Broadcom Emulex firmware and PMD driver installed.
 2. Enable SR-IOV on the host.
 3. Using VMM on the RHEL 7.1 or RHEL 7.2 host, create a VM and install RHEL 7.1 or RHEL 7.2 on it.
 4. Compile DPDK 1.8.0, or 2.0 on the VM.
 5. Compile the PMD+SURF driver on the VM.
 6. Shut down the VM.
 7. Reboot the VM.
- 2. The *Emulex Poll Mode Driver User Manual* in Section 2.4.3, `max_dpdk_oce_port <int>` must be appended.**

The following text and note must be appended to the existing description in Section 2.4.3, `max_dpdk_oce_port <int>`:

The maximum value is 32.

Note: This parameter must be increased to a number greater than 8 when the total number of be2net NIC functions in a system exceed 8.

The total number of be2net NIC functions present in a system must not exceed 32 to use DPDK.

- 3. The text in the *Emulex Poll Mode Driver User Manual* Section 1.2, Features must be replaced.**

The following text must replace the existing text in Section 1.2, Features.

The Broadcom Emulex PMD includes the following features:

- Concurrent iSCSI support
- Universal multi-channel (UMC) configuration support
 - ◆ PMD installation is available on all four channels per port provided that the PMD is installed on three or four channels.

Note: If you install the Broadcom Emulex PMD on only one or two channels on the port, it must be installed on the third channel, the fourth channel or both. This is necessary because freeing up the resources for the first or second channel does not release the required RSS resources.

- ◆ Four physical functions (PFs) per port
- SR-IOV Configuration support
 - ◆ On Emulex OCe14000B-series adapters, DPDK SRIOV is not supported on the last port of the adapter.
 - With receive side scaling (RSS) enabled, up to nine virtual functions (VFs) are supported with one PF in an SR-IOV configuration. In this configuration, the PMD running on the VFs has access to two RSS rings each.

- With RSS disabled, up to fourteen VFs are supported with one PF in an SR-IOV configuration. For more information on specifying command line parameters on `dpdk_surf.ko` module, See Section 2.4.4 `max_rxq_per_port <int>`.
- In a 4-port adapter, up to nine VFs are supported on each of the first three ports. The fourth port does not support any VFs because of limited resources.

4. The text in the *Emulex Poll Mode Driver User Manual* Section 2.4.4, `max_rxq_per_port <int>` must be replaced.

The following text must replace the existing text in Section 2.4.4, `max_rxq_per_port <int>`.

The `max_rxq_per_port` parameter specifies the maximum number of Rx queues per port provided to the Broadcom Emulex PMD. The default value is 0.

Note: Because of the limited resources on an OCe14000-series adapter, creating a single RSS-enabled Rx queue might fail in SR-IOV with more than nine VFs enabled per port. To support 10 to 14 VFs per port, you must request a single non-RSS Rx queue by setting `max_rxq_per_port` to 1. This forces the Broadcom Emulex PMD to create only one non-RSS Rx queue per OCe14000-series adapter port.

5. The text in the *Emulex Poll Mode Driver User Manual* Section 3.2, SR-IOV Limitations must be replaced.

The following text must replace the existing text in Section 3.2, SR-IOV Limitations.

The Broadcom Emulex PMD has the following limitations when using SR-IOV:

- With RSS enabled, up to nine VFs are supported with one PF in an SR-IOV configuration. In this configuration, the PMD running on the VFs has access to two RSS rings each. RSS is enabled by default and the following command is typically used in the `load_drivers` script: `# insmod $BINARIES/dpdk_surf.ko`. For more information, see Section 2.5.1, Create the `load_drivers` Script.
- With RSS disabled, up to 14 VFs are supported with one PF. To disable RSS, use the following command in the `load_drivers` script: `# insmod $BINARIES/dpdk_surf.ko max_rxq_per_port=1`. For more information, see Section 2.5.1, Create the `load_drivers` Script.
- In a 4-port adapter, up to four VFs are supported on each of the first three ports. The fourth port does not support any VFs because of limited resources.
- Alternative requester ID interpretation (ARI) has not been fully qualified.

For SR-IOV configuration information, refer to the *Emulex Drivers for Linux User Manual*.

Technical Tips

- 1. When using DPDK, the Network Manager can interfere with the underlying be2net host interfaces that are being used as DPDK ports and in some scenarios lead to an Unrecoverable Error in the card.**

Workaround

For each be2net network interface being used by DPDK, disable Network Manager for the interface or disable Network Manager. Also disable any scripts that may access the underlying be2net interface.

```
/etc/sysconfig/network-scripts/ifcfg-eth0
```

```
NM_CONTROLLED="no"
```

Or

```
chkconfig NetworkManager off
```

```
service NetworkManager stop
```

Broadcom, the pulse logo, Connecting everything, the Connecting everything logo, Avago Technologies, and Emulex are among the trademarks of Broadcom Ltd. and/or its affiliates in the United States, certain other countries and/or the EU. Any other trademarks or trade names mentioned are the property of their respective owners. Broadcom Ltd. 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 Ltd. is believed to be accurate and reliable. However, Broadcom Ltd. 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. Copyright © 2015- 2016 Broadcom Ltd. All Rights Reserved.