



## InfiniPath 2.2.1 Drivers and Software for QLogic QHT7xxx and QLE7xxx HCAs

This software license applies only to QLogic customers.  
QLogic Corporation.  
All rights reserved.

### Table of Contents

1. [Version](#)
2. [Changes](#)
3. [Bug Fixes](#)
4. [Known Issues](#)
  - 4.1 [Upgrade All Compute Nodes](#)
  - 4.2 [Conflicts With RHEL4 U5, U6 and RHEL5 OpenFabrics Packages](#)
  - 4.3 [Restart Driver after Installing infinipath-kernel RPM](#)
  - 4.4 [BIOS Settings](#)
  - 4.5 [Restart System for ulimit Changes to take Effect](#)
  - 4.6 [mpi-frontend Versions 2.2/2.2.1 Not Compatible with Jobs Running Older InfiniPath Versions](#)
  - 4.7 [MPD Requires mpi-frontend-2.2 and python](#)
  - 4.8 [Compatibility Between Other MPIs and Fortran Compiled Binaries](#)
  - 4.9 [Performance Issues with Barcelona Stepping B2 and BIOS/kernel Workarounds](#)
  - 4.10 [mpirun -debug Option Fails to Open xterm Windows](#)
  - 4.11 [Additional Information on Interrupt ReQuest \(IRQ\) Affinity Settings](#)
  - 4.12 [Explicitly Remove the ipoibtools RPM if Upgrading from the 2.1 Release to the 2.2.1 Release](#)
  - 4.13 [Workaround in Case of IntelMPI/uDAPL Error "unexpected DAPL event 4008"](#)
  - 4.14 [Problem Sending Packets on ipath\\_ether on RHEL5 with the 2.6.18-92 Kernel](#)
  - 4.15 [Open MPI Queries MTU from .ini File Instead of Switch](#)
  - 4.16 [Set Parameters to Enable Improved SRP Write Performance](#)
  - 4.17 [Error when Compiling Open MPI Programs with gcc](#)
5. [Additional Information](#)
6. [Trademarks](#)
7. [Notices](#)
8. [Contacting Support](#)

### 1. Version

---

These release notes describe the changes, fixes, known issues, and release details that apply to InfiniPath drivers and software, version 2.2.1.

Note that the OpenFabrics\*, OpenSM\*, OtherHCAs\*, and OtherMPIs\* RPMs with this release remain at version 2.2. They are identical to the RPMs shipped with InfiniPath 2.2, and were not rebuilt.

## 2. Changes

---

The following changes have been made to the InfiniPath software between versions 2.2 and 2.2.1.

### Hardware Support

- No additional hardware support.

### OS Support

- Added kernel support (2.6.18-92) for Red Hat Enterprise Linux 5.1 (RHEL 5.1), Scientific Linux 5.1, and CentOS 5.1 (Rocks 5.1).
- Removed support for Red Hat Fedora 6.

### Software Components

- QLogic Host-based Subnet Manager works with InfiniPath.
- Software also includes the OpenFabrics 1.3-based library and utility RPMs, which are identical to those shipped with the 2.2 release.
- Software works with the Lustre cluster file system Version 1.6.5.1. For more information about downloading, installing, and configuring Lustre, go to <http://www.lustre.org>.

### Industry Standards Compliance

- IBTA and PCIe compliant
- OFED-certified

## 3. Bug Fixes

---

The following fixes have been made to the InfiniPath software between versions 2.2 and 2.2.1.

- Fixed an issue that prevented CentOS 4.6 from installing properly.
- Intel MPI over uDAPL performance and reliability now comparable to other MPIs over IB Verbs.
- Fixed incorrect check for max physical address in TID; affects systems with 64GB or greater memory.
- Fixed SRP write performance bug between QLE7240/QLE7280 and LSI (Engenio) storage systems (Bug 14457). For details, see [4.16 Set Parameters to Enable Improved SRP Write Performance](#).

## 4. Known Issues

---

The InfiniPath software, version 2.2.1, has the known issues described in the following sections:

- [4.1 Upgrade All Compute Nodes](#)
- [4.2 Conflicts with RHEL4 U5, U6 and RHEL5 OpenFabrics Packages](#)
- [4.3 Restart Driver after Installing infinipath-kernel RPM](#)
- [4.4 BIOS Settings](#)
- [4.5 Restart System for ulimit Changes to take Effect](#)
- [4.6 mpi-frontend Versions 2.2/2.2.1 Not Compatible with Jobs Running Older InfiniPath Versions](#)
- [4.7 MPD Requires mpi-frontend-2.2 and python](#)
- [4.8 Compatibility Between Other MPIs and Fortran Compiled Binaries](#)
- [4.9 Performance Issues with Barcelona Stepping B2 and BIOS/kernel Workaround](#)
- [4.10 mpirun -debug Option Fails to Open xterm Windows](#)
- [4.11 Additional Information on Interrupt ReQuest \(IRQ\) Affinity Settings](#)
- [4.12 Explicitly Remove the ipoibtools RPM if Upgrading from the 2.1 Release to the 2.2.1 Release](#)
- [4.13 Workaround in Case of IntelMPI/uDAPL Error "unexpected DAPL event 4008"](#)
- [4.14 Problem Sending Packets on ipath\\_ether on RHEL5 with the 2.6.18-92 Kernel](#)
- [4.15 Open MPI Queries MTU from .ini File Instead of Switch](#)
- [4.16 Set Parameters to Enable Improved SRP Write Performance](#)
- [4.17 Error when Compiling Open MPI Programs with gcc](#)

### 4.1 Upgrade All Compute Nodes

---

Upgrade all compute nodes to the InfiniPath 2.2.1 RPMs at the same time. Jobs may not be mixed between compute nodes running InfiniPath 2.2.1 and older InfiniPath releases. Upgrade all installed InfiniPath RPMs at the same time. The mpirun command in this release is not compatible with nodes running earlier InfiniPath releases.

## 4.2 Conflicts with RHEL4 U5, U6 and RHEL5 OpenFabrics Packages

---

OpenFabrics on the RHEL4 and RHEL5 distributions has the following two issues:

- Some of the OpenFabrics RPMs (most notably `ibutils`) have library dependencies that are distribution-specific.
- The OpenFabrics packages contained in this release are newer than the ones that are distributed with RHEL4 U5 and RHEL5, and can cause conflicts during installation. The following RPM has been added to the RHEL tarball to work around the conflicts problem:  
`rhel4-ofed-fixup-2.2-xxx.rhel4_psc.noarch.rpm`

where

'xxx' is the build number

You can find this file in the OpenFabrics directory. You must install it while installing the other OpenFabrics RPMs. Installing this RPM will not affect Fedora 6 users.

## 4.3 Restart Driver after Installing `infinipath-kernel` RPM

---

After installing the `infinipath-kernel` RPM, you must restart the InfiniPath driver before the changes will take effect. To do this, enter the command:

```
/etc/init.d/infinipath restart
```

## 4.4 BIOS Settings

---

The BIOS settings have the following two issues:

- You must enable Advanced Configuration and Power Interface (ACPI).
- Set the Memory Type Range Registers (MTRR) mapping to "Discrete".

The InfiniPath driver uses MTRR to enable write by combining to the on-chip transmit buffers. This option improves write bandwidth to the QLogic chip by writing multiple words in a single bus transaction (typically 64). This option applies only to `x86_64` systems.

However, some BIOS do not have the MTRR mapping option. It may be referred to in a different way, depending on the chipset, vendor, BIOS, or other factors. For example, it is sometimes referred to as a 32-bit memory hole, which must be enabled.

A script, `ipath_mtrr`, sets MTRR registers to enable maximum performance from the InfiniPath driver. This script is available as part of the InfiniPath package.

You can also check and adjust these BIOS settings using the BIOS Setup utility. For specific instructions, follow the hardware documentation that came with your system. Also see the *QLogic HCA and InfiniPath Software Install Guide* and *QLogic HCA and InfiniPath Software Users Guide*.

## 4.5 Restart System for `ulimit` Changes to take Effect

---

Changes have been made to set '`ulimit -l 131072`' in `/etc/initscript`. This allows locking of up to 128MB of memory per process. You must reboot your system for these changes to take effect. This is not an issue when upgrading from a current (2.2x) release.

## 4.6 `mpi-frontend` Versions 2.2/2.2.1 Not Compatible with Jobs Running Older InfiniPath Versions

---

Neither the 2.2 nor 2.2.1 version of the `mpi-frontend` RPM can be used for jobs running on nodes with InfiniPath versions older than 2.2.

However, the 2.2.1 `mpi-frontend` RPM is compatible with jobs running on nodes with 2.2 software.

Similarly, the 2.2 `mpi-frontend` is compatible with jobs running on nodes with 2.2.1 software.

You must install the `mpi-frontend-*` RPM on all nodes that will use `mpirun`. Alternatively, you can use the `mpirun` option `-distributed=off`.

## 4.7 MPD Requires `mpi-frontend-2.2.1` and `python`

---

To use MPD, you must install the software packages `mpi-frontend-2.2.1*.rpm` and `python` (available with your distribution) on all nodes.

#### **4.8 Compatibility Between Other MPIs and Fortran Compiled Binaries**

---

The MPIs listed under "OtherMPIs RPMs" have been compiled for PSM to support the following versions of the compilers on both RHEL4U4 and SLES 10:

- (GNU) gcc 4.1.0
- (PathScale) pathcc 3.0
- (PGI) pgcc 7.1-5
- (Intel) icc 10.1.012 (icc 10.1, Build 20080112 Package ID: l\_cc\_p\_10.1.012)

Users should expect binary compatibility between these MPIs and Fortran compiled binaries according to the following rules:

- The gfortran version must be 4.1.x (g77 is not supported).
- The pathf77 pathf90 versions can be any 2.x or 3.x compiler.
- The pgf77 and pgf90 version must be 7.x.
- The ifort version must be 10.x.

#### **4.9 Performance Issues with Barcelona Stepping B2 and BIOS/kernel Workaround**

---

QLE7240/QLE7280 performance with Barcelona stepping B2 and the BIOS/kernel workaround for AMD errata 298 will lead to reduced InfiniBand bandwidth in some circumstances. This is due to poor memory bandwidth and memory latency on those parts because of the workaround.

To fix this problem, upgrade to Barcelona stepping B3 (or later).

#### **4.10 mpirun -debug Option Fails to Open xterm Windows**

---

A race condition may occur when there are multiple processes trying to lock the \$HOME/.Xauthority file concurrently. An error similar to this may occur:

```
/usr/X11R6/bin/xauth: error in locking authority file /home/<user>/.Xauthority
```

To work around this problem, disable X11 forwarding by setting the environment variable \$MPI\_SHELL\_X='ssh -x'. Then set the DISPLAY environment variable with the mpirun -display option. For example:

```
env MPI_SHELL_X='ssh -x' mpirun -debug -display \  
<X11_display_hostname:display_number.screen_number> [...] (continued on same line)
```

#### **4.11 Additional Information on Interrupt ReQuest (IRQ) Affinity Settings**

---

These points add to the information in Appendix C.6.4 "Erratic Performance" in the *QLogic HCA and InfiniPath Software Users Guide*.

- You can use a command similar to this to immediately change the processor affinity of an IRQ:  
# echo 01 > /proc/irq/\$my\_irq/smp\_affinity  
However, the contents of the smp\_affinity file may not reflect the expected values, even though the affinity change has taken place.
- If the driver is reloaded, the affinity assignment will revert to the default, so you will need to reset it to the desired value.

You can look at the stats in /proc/interrupts while the adapter is active to observe which CPU is fielding ib\_ipath interrupts.

#### **4.12 Explicitly Remove the ipoitools RPM if Upgrading from the 2.1 Release to the 2.2.1 Release**

---

The ipoitool RPM is now obsolete in OFED 1.3. It will be left behind if upgrading from the 2.1 release to the 2.2.1 release. You can explicitly remove ipoitools with following command:

```
# rpm -e ipoitools
```

However, leaving the RPM will not cause problems.

#### 4.13 Workaround in Case of IntelMPI/uDAPL Error "unexpected DAPL event 4008"

---

The following error may occur on rare occasions with IntelMPI/uDAPL:

```
"unexpected DAPL event 4008 from ..."
```

To work around this, add the following to your mpirun command:

```
-genv I_MPI_USE_DYNAMIC_CONNECTIONS 0
```

This problem is caused by a limitation in Intel MPI/uDAPL's dynamic connection mechanism when MPI processes are not sufficiently attentive to incoming interconnect traffic.

#### 4.14 Problem Sending Packets on ipath\_ether on RHEL5 with the 2.6.18-92 Kernel

---

In some conditions (seen only on RHEL5 2.6.18-92 kernels in testing) the `ipath_ether` module stops sending packets. In some cases it recovers by itself, in other cases the `ipath_ether` interface must be configured down and back up to recover. (Bug 14877) In this case, use these commands (as root):

```
# ifconfig eth2 down
# ifconfig eth2 up
```

#### 4.15 Open MPI Queries MTU from .ini File Instead of Switch

---

Open MPI sets the QP MTU for the QLE7240/QLE7280 to 4K, even if connected to a 2K switch. This can result in MPI errors which can hang the MPI job. (Bug 14938)

To workaround this problem, change the MTU value in the `$prefix/share/openmpi/mca-btl-openib-hca-params.ini` file from "4096" to "2048". For example:

```
[QLogic InfiniPath PCIe]
vendor_id = 0x1fc1,0x1077
vendor_part_id = 16,29216
use_eager_rdma = 1
mtu = 2048
```

**NOTE:** You must change this in the `.ini` file on each node in the cluster.

#### 4.16 Set Parameters to Enable Improved SRP Write Performance

---

An SRP write performance bug between QLE7240/QLE7280 and LSI (Engenio) storage systems was fixed. To apply this fix, you need to modify the `srp_sg_tablesize` and `max_sect` parameters.

To modify `srp_sg_tablesize`, do one of the following:

- Add the following line to `/etc/modprobe.conf`:  
`options ib_srp srp_sg_tablesize=80`
- Define it when you add the `ib_srp` module (as root):  
`# modprobe ib_srp srp_sg_tablesize=80`

If you are using `srp_daemon`, modify `max_sect` by adding the following command to `/etc/srp_daemon.conf`:  
`a id_ext=200500A0B81146A1,ioc_guid=00a0b80200402bef,max_sect=2048`

#### 4.17 Error when Compiling Open MPI Programs with gcc

---

On RHEL 5, 5.1 and 5.2 systems, when compiling Open MPI programs with `gcc`, you will see an error. For example:

```
/usr/mpi/gcc/openmpi-2.2/bin/mpicc -I/usr/mpi/gcc/openmpi-2.2/include -o bw bw.c;
```

```
-----
The Open MPI wrapper compiler was unable to find the specified compiler gcc4 in your PATH.
Note that this compiler was either specified at configure time or in one of several
possible environment variables.
-----
```

```
make: *** [BW] Error 1
```

Workaround: Edit this file:

```
/usr/mpi/gcc/openmpi-2.2/share/openmpi/mpicc-wrapper-data.txt
```

and change "gcc4" to "gcc".

**NOTE:** This will be fixed in the next release.

## 5. Additional Information

---

Please refer to the "Troubleshooting Appendix" in the *QLogic HCA and InfiniPath Software Install Guide* for more information on installation issues.

In an upcoming release, `ipath_ether` will be deprecated. QLogic recommends using the standardized protocol IPoIB in Connected Mode to provide the same functionality.

## 6. Trademarks

---

Linux is a registered trademark of Linus Torvalds.

QLA, QLogic, the QLogic logo, and InfiniPath are registered trademarks of QLogic Corporation.

Red Hat and all Red Hat-based trademarks are trademarks or registered trademarks of Red Hat, Inc. SuSE is a registered trademark of SuSE Linux AG.

All other brand and product names are trademarks or registered trademarks of their respective owners.

## 7. Notices

---

Information furnished in this document is believed to be accurate and reliable. However, QLogic Corporation assumes no responsibility for its use, nor for any infringements of patents or other rights of third parties which may result from its use. QLogic Corporation reserves the right to change product specifications at any time without notice. Applications described in this document for any of these products are only for illustrative purposes. QLogic Corporation makes no representation nor warranty that such applications are suitable for the specified use without further testing or modification. QLogic Corporation assumes no responsibility for any errors that may appear in this document.

## 8. Contacting Support

---

Please feel free to contact your QLogic approved reseller or QLogic Technical Support at any phase of integration for assistance. QLogic Technical Support can be reached by the following methods:

Web: <http://support.qlogic.com>

North America Contact Information

Email: [support@qlogic.com](mailto:support@qlogic.com)

Phone: (952) 932-4040

Support contact information for other regions of the world is available at the QLogic website:

<http://support.qlogic.com>

[Go to Top](#)



© Copyright 2008. All rights reserved worldwide. QLogic, the QLogic logo, and the Powered by QLogic logo are registered trademarks of QLogic Corporation. All other brand and product names are trademarks or registered trademarks of their respective owners.