

## QLogic OFED+ Host Software

QLogic Corporation  
All rights reserved

### Table of Contents

<a href="#">1</a>	<a href="#">Version</a>
<a href="#">2</a>	<a href="#">Changes</a>
<a href="#">2.1</a>	<a href="#">Changes to Hardware Support</a>
<a href="#">2.2</a>	<a href="#">Changes to OS Support</a>
<a href="#">2.3</a>	<a href="#">Changes to Software Components</a>
<a href="#">2.4</a>	<a href="#">Changes to Industry Standards Compliance</a>
<a href="#">3</a>	<a href="#">Bug Fixes</a>
<a href="#">4</a>	<a href="#">Known Issues</a>
<a href="#">5</a>	<a href="#">Additional Information</a>
<a href="#">5.1</a>	<a href="#">Included in this Release</a>
<a href="#">5.2</a>	<a href="#">Operating Systems Supported in this Release</a>
<a href="#">5.3</a>	<a href="#">Lustre and GPFS Versions Verified for this Release</a>
<a href="#">5.4</a>	<a href="#">InfiniBand Host Channel Adapters Supported in this Release</a>
<a href="#">5.5</a>	<a href="#">Performance Tuning</a>
<a href="#">5.6</a>	<a href="#">New Features</a>
<a href="#">5.7</a>	<a href="#">Product Constraints</a>
<a href="#">5.8</a>	<a href="#">Product Limitations</a>
<a href="#">5.9</a>	<a href="#">Other Information</a>
<a href="#">6</a>	<a href="#">Trademarks</a>
<a href="#">7</a>	<a href="#">Notices</a>
<a href="#">8</a>	<a href="#">Contacting Support</a>

## 1 Version

---

These release notes describe the changes, fixes, known issues, and release details that apply to the QLogic OFED+ Host software package for version 6.0.2.1.11.

## 2 Changes

---

The following sections describe the changes that have been made to the QLogic OFED+ Host software package between versions 5.1.0.0.49 and 6.0.2.1.11, including the following releases:

- ◆ 5.1.0.0.49
- ◆ 5.1.0.2.1
- ◆ 5.1.0.2.8

- ◆ 5.1.0.3.14
- ◆ 6.0.0.0.54
- ◆ 6.0.0.1.2
- ◆ 6.0.0.2.7
- ◆ 6.0.1.0.30
- ◆ 6.0.1.1.3
- ◆ 6.0.2.0.28
- ◆ 6.0.2.1.11

For detailed information about any of the previous releases listed, refer to the Release Notes for the specific version.

## 2.1 Changes to Hardware Support

---

Added hardware support for the following releases:

- ◆ Release 5.1.0.0.49
  - QLE7340
  - QLE7342
  - QMH7342
- ◆ Release 5.1.0.2.1
  - None
- ◆ Release 5.1.0.2.8
  - None
- ◆ Release 5.1.0.3.14
  - QME7342
- ◆ Release 6.0.0.1.2
  - None
- ◆ Release 6.0.0.2.7
  - None
- ◆ Release 6.0.1.0.30
  - None
- ◆ Release 6.0.1.1.3
  - None
- ◆ Release 6.0.2.0.28
  - None
- ◆ Release 6.0.2.1.11
  - None

## 2.2 Changes to OS Support

---

Added operating system (OS) support for the following releases:

- ◆ Release 5.1.0.0.49
  - SuSE<sup>®</sup> Linux Enterprise Server (SLES) 11.0 X86\_64 (AMD Opteron<sup>®</sup> and Intel<sup>®</sup> EM64T)
    - (Base) 2.6.27.19-5-default
- ◆ Release 5.1.0.2.1
  - None
- ◆ Release 5.1.0.2.8
  - Added support to allow the 2.6.18-164 kernel from Red Hat<sup>®</sup> Enterprise Linux (RHEL)5 Update 4 to be installed over Red Hat EL5 Update 3, as a kernel security patch.
  - Rocks 5.3

- 2.6.18-164.6.1.el5
- ◆ Release 5.1.0.3.14
  - Red Hat EL5 X86\_64 (AMD Opteron, and Intel EM64T)
    - (Update 4) 2.6.18-164.9.1el5
  - Platform Cluster Manager 1.2b
    - (RHEL 5.4) 2.6.18-164.6.1.el5
- ◆ Release 6.0.0.0.54
  - Red Hat EL4 X86\_64 (AMD Opteron, and Intel EM64T)
    - (Update 7) 2.6.9-78.0.1.ELsmp
    - (Update 8) 2.6.9-89.ELsmp
  - Red Hat EL5 X86\_64 (AMD Opteron, and Intel EM64T)
    - (Update 3) 2.6.18-128.1.1.el5, 2.6.18-128, 1.6.el5, 2.6.18-128.1.10.el5, 2.6.18-128.1.14.el5, 2.6.18-128.7.1.el5
    - (Update 4) 2.6.18-164.6.1.el5, 2.6.18-164.9.1.el5
  - SuSE Linux Enterprise Server (SLES) 10.0 X86\_64 (AMD Opteron and Intel EM64T)
    - (SP2) 2.6.16.60-0.21-default
    - (SP3) 2.6.16.60-0.54.5-smp, 2.6.16.60-0.54.5-default
  - CentOS X86\_64 (AMD Opteron and Intel EM64T)
    - (Update 5.4) 2.6.18-164.el5
  - Scientific Linux X86\_64 (AMD Opteron and Intel EM64T)
    - (Update 5.4) 2.6.18-164.el5
- ◆ Release 6.0.0.1.2
  - None
- ◆ Release 6.0.0.2.7
  - None
- ◆ Release 6.0.1.0.30
  - Red Hat EL5 X86\_64 (AMD Opteron, and Intel EM64T)
    - (Update 5) 2.6.18-194.el5
- ◆ Release 6.0.1.1.3
  - Platform Cluster Manager 2.0.1
    - (RHEL 5.5) 2.6.18-194.el5
- ◆ Release 6.0.2.0.28
  - SuSE Linux Enterprise Server (SLES) 11 X86\_64 (AMD Opteron and Intel EM64T)
    - (Update 1) 2.6.32.12-0.7-default
  - Rocks
    - (Rocks 5.3.2) 2.6.18-194.el5
- ◆ Release 6.0.2.1.11
  - None

See list of supported Operating Systems in [Section 5](#).

## 2.3 Changes to Software Components

---

Changes made to the software components for the following releases:

- ◆ Release 5.1.0.0.49
  - QLogic OFED+ Software
  - QLogic InfiniBand Tools Software
- ◆ Release 5.1.0.2.1
  - QLogic OFED+ Software
  - QLogic InfiniBand Tools Software
- ◆ Release 5.1.0.2.8
  - QLogic OFED+ Software
  - QLogic InfiniBand Tools Software
- ◆ Release 5.1.0.3.14
  - QLogic OFED+ Software
  - QLogic InfiniBand Tools Software
- ◆ Release 6.0.0.0.54
  - QLogic OFED+ Software
  - QLogic InfiniBand Tools Software
- ◆ Release 6.0.0.1.2
  - None
- ◆ Release 6.0.0.2.7
  - QLogic OFED+ Software
  - QLogic InfiniBand Tools Software
- ◆ Release 6.0.1.0.30
  - QLogic OFED+ Software
  - QLogic InfiniBand Tools Software
- ◆ Release 6.0.1.1.3
  - None
- ◆ Release 6.0.2.0.28
  - QLogic OFED+ Software
  - QLogic InfiniBand Tools Software
- ◆ Release 6.0.2.1.11
  - QLogic OFED+ Software

## 2.4 Changes to Industry Standards Compliance

---

- ◆ Supports the OpenFabrics Enterprise Distribution (OFED®) version 1.5.2 software package.
- ◆ Each of the upper layer protocols (ULPs) support their appropriate industry standard. Please see the release notes, found in the docs directory, for a particular ULP compliance level.

## 3 Bug Fixes

---

The following fixes have been made to the QLogic OFED+ Host software package between versions 5.1.0.0.49 and 6.0.2.1.11:

- ◆ Fixed in Release 5.1.0.0.49
  - Test case `mpi_fork` no longer fails when using `openmpi`.
  - When installing QLogic OFED+ Host software or QLogic InfiniBand Fabric Suite software packages on a server configured with Ethernet bonding, the server no longer hangs when IPoIB bonding is being installed.
  - When using `opensm`, `iba_report` and `fabric_info` correctly reports the number of Subnet Managers (SM)s in the fabric.
  - When running with non-QLogic DDR or QDR devices, links now come up at their appropriate speeds.
  - When using OFED's standard `ib_srp` with a QLogic Fabric Virtual I/O Controller (FVIC) Gateway module, data write operations no longer fail (which resulted in a disconnect of the device).
  - When running `mvapich1` over Performance Scaled Messaging (PSM), all commands such as `saquery`, `iba_saquery`, FastFabric tools and others are available, along with non-MPI ULPs and commands.
  
- ◆ Fixed in Release 5.1.0.2.1
  - Added support for systems equipped with more than 16 cores. To run with Open Message Passing Interface (MPI), download and build either Open MPI 1.34 or OpenMPI 1.4.
  
- ◆ Fixed in Release 5.1.0.2.8
  - Added support to allow installation of the 2.6.18-164 kernel on top of RHEL 5 U3.
  - Added support for systems equipped with more than 16 cores. This support includes Open MPI version 1.4, which replaces the older version in previous releases.
  
- ◆ Fixed in Release 5.1.0.3.14
  - None
  
- ◆ Fixed in Release 6.0.0.0.54
  - `saquery -p query` now issues a compliant GetTable (PathRecord) query. This runs correctly when run against InfiniBand Trade Association (IBTA) compliant Fabric Managers (FM)s.
  - An issue in the Mellanox ConnectX firmware, causing it to freeze when doing any of the following, has now been corrected:
    1. When using any command or operation which uses `mstvpd`.
    2. When using `iba_hca_rev`. This utility reports the revision of the firmware on the Host Channel Adapter.
    3. When using QLogic's firmware update tool, the user may have to reboot the system, even if they did not change the firmware.

- Open Subnet Manager (SM) now supports Multicast Forwarding Table Subnet Administration (SA) queries. FastFabric features such as route analysis using `iba_report` or `iba_saquery` are available when using opensm.
  - The OFED command `ibcheckerrors` no longer displays FAILED for xEdge and 12100 switches.
- ◆ Fixed in Release 6.0.0.1.2
    - None
  - ◆ Fixed in Release 6.0.0.2.7
    - None
  - ◆ Fixed in Release 6.0.1.0.30
    - None
  - ◆ Fixed in Release 6.0.1.1.3
    - None
  - ◆ Fixed in Release 6.0.2.0.28
    - When the QLogic OFED+ kit is installed at the time of Platform Cluster Manager (PCM) installation, there is no longer an error message stating that the `qlogic_ofed` kit is incompatible with the OS kit. The text of the error message was:  
Kits incompatible with the OS Kit (rhel-5.5-x86\_64) are detected. Please remove the following kits to continue:  
`qlogic_ofed`
  - ◆ Fixed in Release 6.0.2.1.11
    - The issue of QLogic OFED+ 6.0.2.0.28 is not being compatible with CentOS 5.x operating systems has been resolved with the new QLogic OFED+ 6.0.2.1.11.

## 4 Known Issues

---

The QLogic OFED+ Host software package, version 6.0.2.1.11, has the following known issues:

Known Issue	Workaround
VNIC driver hangs during shutdown when using a gateway. This is a known problem in the Linux kernel in all releases prior to 2.6.27.	None
When a port is down and does not have a LID assigned, <code>clear_p1stats</code> or <code>clear_p2stats</code> will fail against the given port.	None
<code>unregister_netdevice</code> hangs while restarting VNIC service on the host with a TrueScale Host Channel Adapter when using a gateway. This is a known issue in the Linux kernel in all releases prior to 2.6.27.	None

Known Issue	Workaround
<p>When canceling out of a Host Channel Adapter firmware update, the following message is shown:</p> <pre>HCA update failed. Return code: 1 at /usr/bin/qlgc_firmware_tool line 264, STDIN line 2.</pre>	<p>This message can be ignored; it indicates that the Host Channel Adapter firmware was not updated due to the user canceling the operation.</p>
<p>For SLES10 SP2, the <code>--32bit</code> option of <code>INSTALL</code> will not work. For individual Red Hat Package Manager (RPM) installs, the SLES10 SP2 does not distinguish between 32-bit and 64-bit RPMs. Installed RPMs that are 32-bit must be manually uninstalled first and then the appropriate RPM for each package (only 64 bit if available) must be installed.</p>	<p>QLogic recommends using <code>./INSTALL</code>, which automatically performs all necessary uninstalls of old RPMs prior to installing the new 64-bit RPMs.</p>
<p>Issues have been observed when using IPv6 with IPoIB on Red Hat Enterprise Linux 4.</p>	<p>If using IPv6 with IPoIB, QLogic recommends using Red Hat Enterprise Linux 5.</p>
<p>When using <code>opensm</code>, the incorrect Node Description may be reported for hosts after they are rebooted.</p>	<p>Restart <code>opensm</code>.</p> <p>QLogic recommends using the QLogic Fabric Manager to correctly handle changes to host nodenames and host reboots. If using <code>opensm</code> when seeing this problem, <code>opensm</code> can be restarted.</p>
<p>When reinstalling QLogic OFED+, it may try to stop existing instances of <code>opensm</code>. If <code>opensm</code> is not presently running, it will report:</p> <pre>Stopping IB Subnet Manager [FAILED].</pre>	<p>QLogic recommends using the QLogic Fabric Manager.</p> <p>If the QLogic Fabric Manager is installed instead of <code>opensm</code>, this error will not occur.</p>
<p>When using vFabric, the OFED <code>saquery</code> command may use the wrong P-Key and timeout waiting for responses.</p>	<p>QLogic recommends using the <code>iba_saquery</code> tool, which is included with QLogicIB-Basic or QLogicIB-IFS. <code>iba_saquery</code> will work properly when vFabric is configured.</p>
<p>If <code>LD_LIBRARY_PATH</code> is exported inconsistently with the version of <code>openmpi</code> being used, applications may build or run incorrectly. This issue can impact FastFabric tools that use MPI, rebuilding of <code>mpi</code> apps, or rebuilding <code>openmpi</code> itself using the <code>do_build</code> or <code>do_openmpi_build</code> tools.</p>	<p>When using <code>openmpi</code>, make sure <code>PATH</code> and <code>LD_LIBRARY_PATH</code> are not exported specifying a different path than the <code>openmpi</code> path that is being used. The <code>mpi-selector</code> can configure a <code>LD_LIBRARY_PATH</code> for subsequent logins. <code>Openmpi</code> does not require the <code>LD_LIBRARY_PATH</code> to be set.</p>
<p>When using <code>opensm</code>, after bouncing ports on a node, the port may not return to an active state for a period of time. As a result, commands that issue an SA query such as OFED's <code>saquery</code> command, or various FastFabric tools such as <code>iba_report</code> and <code>iba_saquery</code>, may hang waiting for the port to become active and the SA to respond.</p>	<p>Restart <code>opensm</code>.</p> <p>QLogic recommends using the QLogic Fabric Manager, which has much greater resiliency and quicker handling of port state changes.</p>
<p>Test <code>ib_rdma_bw</code> fails if it detects conflicting CPU frequencies.</p>	<p>Use the <code>-F</code> option (do not fail even if <code>cpufreq_ondemand</code> module is loaded) when running this test. This option is not documented, but is required for this environment.</p>
<p>When using vFabric to change an IPoIB application from Networking to Non-Networking, the IPoIB interface may remain in a running state.</p>	<p>After changing the application, restart the network services or bring the interface down/up to force IPoIB to re-query the SM and correct the situation.</p>

Known Issue	Workaround
<p>When installing the QLogic-Basic or QLogic-IFS SW on SLES10SP3, there may be conflicts with software that is already installed on the system. The following message may appear:</p> <pre>error: %preun(ofed-1.4.1-0.14.9.x86_64) scriptlet failed, exit status 1</pre>	<p>Manually uninstall the old version OFED before you install the newer software.</p> <ol style="list-style-type: none"> <li>As root, run the following command:  <pre>rpm -e --noscripts ofed</pre> </li> <li>Re-run the normal installation.</li> </ol>
<p>When installing the QLogic-Basic or QLogic-IFS SW on SLES11SP1, there may be conflicts with software that is already installed on the system. The following message may appear:</p> <pre>error: %preun(ofed-1.4.2-0.9.6.x86_64) scriptlet failed, exit status 1 error: %preun(opensm-3.2.6_20090317-0.1.42.x86_64)scriptlet failed, exit status 1 Unable to uninstall previous OFED RPMs</pre>	<p>Manually uninstall the old version OFED before you install the newer software.</p> <ol style="list-style-type: none"> <li>As root, run the following command:  <pre>rpm -e --noscripts ofed</pre> </li> <li>Re-run the normal installation.</li> </ol>
<p>When uninstalling MVAICH2 (for any verbs or PSM), some files under the <code>/usr/mpi/*/mvapich2*/</code> directory tree that are created at runtime by MVAICH2 may not be removed. One example is <code>mpdlib.pyc</code>.</p>	<p>After uninstalling, remove any undesired files left in the <code>/usr/mpi/</code> directory.</p>
<p>When a long <code>netperf</code> test is run between hosts using QLE7240 and QLE7280 DDR Host Channel Adapters, the PM is unable to get or set port counters for these Host Channel Adapters.</p> <p>For example, the <code>netperf</code> command, <code>netperf -l 360 -H &lt;ib0 IP address for netserver host&gt;</code> (-l 360 means the test will run for 360 seconds; the default is 10 seconds), causes the QLogic IFS tool <code>iba_top</code> to produce warning messages about being unable to get or set port counters.</p>	<p>Do not run long <code>netperf</code> test between hosts using QLE7240 or QLE7280 DDR Host Channel Adapters.</p>
<p>When running the OFED performance benchmark <code>qperf</code> to test the bandwidth of Unreliable Datagram (UD) traffic, the command hangs or times out intermittently when using an InfiniBand MTU of 4096 bytes and the '-t 10' (or larger number of seconds). For example:</p> <pre>qperf hostname -m 4096 -t 10 ud_bw</pre>	<p>Run the command without the <code>-t</code> option. For example:</p> <pre>qperf hostname -m 4096 ud_bw</pre>
<p>When running bidirectional traffic on a QLogic Host Channel Adapter in unreliable connection (UC) mode on certain server types, the Ethernet connection drops or the adapter can exhibit panics.</p>	<p>For an Ethernet connection drop, perform a service network restart.</p> <p>For panic state, power cycle the server to recover.</p>
<p>When performing an IFS installation on SLES 11 SP1, a message similar to the following may appear:</p> <pre>WARNING: -e needs -E or -F</pre>	<p>This message can be ignored.</p>
<p>On SLES 11 or SLES 11 SP1 environments, when running <code>qperf rc_bi_bw</code> or <code>rc_rdma_read_lat</code> tests, the system can have a kernel panic causing it to drop a <code>vmcore</code> file and reboot. This kernel panic also occurs with OFED 1.5.2.</p>	<p>Do not run these <code>qperf</code> tests on SLES 11 environments.</p>

Known Issue	Workaround
Running MPI microbenchmarks over verbs can exhibit low performance using the <code>osu_bw</code> and <code>osu_bibw</code> tests. It is not known how much this affects applications performance.	Use MPIs built for use with PSM, QLogic's high-performance interface for MPI on QLogic Host Channel Adapters. These MPIs are the MPIs that are listed in the <code>mpi-selector --list</code> with " <code>_qlc</code> " as part of the name, or QLogic MPI.
The <code>ib_send_bw</code> microbenchmark that comes with OFED 1.5.2, and with this host software release, had a regression. Consequently, the command hangs when it is run without the <code>-a</code> (all sizes) option. This benchmark is part of the <code>perftest-1.3.0-0.28</code> suite of benchmarks provided with OFED 1.5.2.	Use the <code>-a</code> option.

## 5 Additional Information

---

### 5.1 Included in this Release

---

- ◆ QLogic OFED+ Host software package (6.0.2.1.11) that includes:
  - QLogic OFED+ (1.5.2.1.12)
  - QLogic InfiniBand Tools (6.0.2.0.15)

### 5.2 Operating Systems Supported in this Release

---

The following operating systems (X86\_64) are supported in this release:

- ◆ Red Hat EL4 X86\_64 (AMD Opteron and Intel EM64T):
  - (Update 7) 2.6.9-78.ELsmp, 2.6.9-78.0.1.ELsmp (QLE7000 series, QMH7342 and QME7342 InfiniBand Adapters are not supported)
  - (Update 8) 2.6.9-89.ELsmp
- ◆ Red Hat EL5 X86\_64 (AMD Opteron and Intel EM64T):
  - (Update 3) 2.6.18-128.el5, 2.6.18-128.1.1.el5, 2.6.18-128.1.6.el5, 2.6.18-128.1.10.el5, 2.6.18-128.1.14.el5, 2.6.18-128.7.1.el5
  - (Update 4) 2.6.18-164.el5, 2.6.18-164.6.1.el5, 2.6.18-164.9.1.el5
  - (Update 5) 2.6.18-194.el5
- ◆ SuSE Linux Enterprise Server (SLES) 10.0 X86\_64 (AMD Opteron and Intel EM64T):
  - (SP2) 2.6.16.60-0.21-smp, 2.6.16.60-0.21-default
  - (SP3) 2.6.16.60-0.54.5-smp, 2.6.16.60-0.54.5-default
- ◆ SuSE Linux Enterprise Server (SLES) 11.0 X86\_64 (AMD Opteron and Intel EM64T):
  - (Base) 2.6.27.19-5-default
  - (SP1) 2.6.32.12-0.7-default
- ◆ CentOS X86\_64 (AMD Opteron and Intel EM64T):
  - (Update 5.3) 2.6.18-128.el5
  - (Update 5.4) 2.6.18-164.el5
  - (Update 5.5) 2.6.18-194.el5

- ◆ Scientific Linux X86\_64:
  - (Scientific Linux 5.3) 2.6.18-128.el5
  - (Scientific Linux 5.4) 2.6.18-164.el5
  - (Scientific Linux 5.5) 2.6.18-194.el5
- ◆ Rocks:
  - (Rocks 5.3.2) 2.6.18-194.el5
- ◆ Platform Cluster Manager 2.0.1:
  - (RHEL 5.5) 2.6.18-194.el5

CPU model of Linux kernel can be identified by `uname -m` and `/proc/cpuinfo` as follows:

Model	Uname	/proc/cpuinfo
EM64T	x86_64	Intel CPUs
Opteron®	x86_64	AMD CPUs

**NOTE:** Other combinations (such as i586 uname) are not currently supported.

### 5.3 Lustre and GPFS Versions Verified for this Release

---

The following Lustre versions are confirmed for this release:

- ◆ Lustre 2.0 with QLogic OFED+ running on RHEL 5.4
- ◆ Lustre 1.8.4 with QLogic OFED+ running on RHEL 5.5

The following GPFS versions are confirmed for this release:

- ◆ GPFS 3.3
- ◆ GPFS 3.4.0.3

### 5.4 InfiniBand Host Channel Adapters Supported in this Release

---

The following models of Host Channel Adapters are supported:

- |                       |                      |              |
|-----------------------|----------------------|--------------|
| ◆ QLE7240             | ◆ 7204-HCA-LPX2P-DDR | ◆ MHQH19-XTC |
| ◆ QLE7280             | ◆ MHGA28-XTC         | ◆ 44R8723    |
| ◆ QLE7340             | ◆ MHGH28-XSC         | ◆ 44R8728    |
| ◆ QLE7342             | ◆ MHGH28-XTC         | ◆ 43W4441    |
| ◆ QME7342             | ◆ MHGH29-XSC         | ◆ 46M2220    |
| ◆ QMH7342             | ◆ MHGH29-XTC         | ◆ 46M2199    |
| ◆ 7104-HCA-128LPX-DDR | ◆ MHGS18-XTC         | ◆ 46M2203    |
| ◆ 7104-HCA-LPX1P-DDR  | ◆ MHQH29-XTC         | ◆ 592519-B21 |
| ◆ 7104-HCA-LPX2P-DDR  |                      |              |

## 5.5 Performance Tuning

---

Tuning for verbs performance can be accomplished by making the following changes.

### 5.5.1 Intel CPU Systems (in general)

---

Set the BIOS using the following procedure:

1. Set PCIe parameter `MaxPayload = 256 bytes`
2. Set PCIe parameter `MaxReadRequest = 4096 bytes`
3. Disable all C-States.
4. Disable Intel Hyper-Threading technology

If the BIOS does not support the `MaxPayload` and `MaxReadRequest` settings, add the following parameter to the `/etc/modprobe.conf` file's `options ib_qib` line:

```
pcie_caps=0x51
```

Set the C-State to 0 where there is no BIOS support, by performing the following procedure:

1. Add kernel boot option: `processor.max_cstate=0`
2. Reboot.

Turn off the specified daemons using the following commands:

```
/sbin/chkconfig irqbalance off
```

```
/sbin/chkconfig --level 12345 cpuspeed off    (for RHEL or similar systems)
```

```
/sbin/chkconfig --level 12345 powersaved off  (for SLES systems)
```

```
/sbin/chkconfig --level 12345 haldaemon off
```

```
Set init runlevel = 3
```

### 5.5.2 Intel CPU Systems (additions for Lustre/GPFS over verbs RDMA)

---

One of the following parameters sets up the QLogic driver to accomplish these settings for Lustre/GPFS over verbs RDMA. For single-port Host Channel Adapters, the line would look like one of the following according to the core and CPU type of the node:

- ◆ For 12-core, Intel Westmere, Xeon 56xx 6-core CPU, nodes:

```
options ib_qib singleport=1 pcie_caps=0x51 krcvqs=4 rcvhdrCnt=4096
```

- ◆ For 8-core, Intel Nehalem, Xeon 55xx, nodes:

```
options ib_qib singleport=1 pcie_caps=0x51 krcvqs=8 rcvhdrCnt=4096
```

- ◆ For 8-core, Intel Harpertown, Xeon 54xx, nodes:

```
options ib_qib singleport=1 pcie_caps=0x51 pcie_coalesce=1 krcvqs=8
rcvhdrCnt=4096
```

- ◆ For 16-core, 4x Intel quad-core CPUs, nodes:

```
options ib_qib singleport=1 pcie_caps=0x51 rcvhdrCnt=4096
(Leave all 16 contexts for PSM)
```

### **5.5.3 Intel Nehalem or Westmere CPU Systems (DIMM configuration)**

Compute node memory bandwidth is important for HPC application performance and for storage node performance. On Intel CPUs codenamed Nehalem or Westmere (Xeon 5500 Series or 5600 Series) it is important to have an equal number of DIMMs on each of the three memory channels for each CPU. On the common dual CPU systems, you should use a multiple of six DIMMs for best performance. Details on DIMM configuration for these CPUs at in the following URL:

<http://www.delltechcenter.com/page/04-08-2009+-+Nehalem+and+Memory+Configurations>

(even though only Nehalem is mentioned in the previous information, the same rules apply for Westmere).

### **5.5.4 AMD CPU Systems (in general)**

Turn off the specified daemons using the following commands:

```
/sbin/chkconfig irqbalance off
```

```
/sbin/chkconfig --level 12345 cpuspeed off    (for RHEL or similar systems)
```

```
/sbin/chkconfig --level 12345 powersaved off  (for SLES systems)
```

```
/sbin/chkconfig --level 12345 haldaemon off
```

Set `init runlevel = 3`

## 5.5.5 AMD CPU Systems (additions for Lustre/GPFS over verbs RDMA)

---

One of the following parameters sets up the QLogic driver to accomplish these settings for Lustre/GPFS over verbs RDMA. For single-port Host Channel Adapters, the line would look like one of the following, according to the core and CPU type of the node:

- ◆ For 8-core, AMD Barcelona or Shanghai, nodes:

```
options ib_qib singleport=1 krcvqs=8 rcvhdrCnt=4096
```

- ◆ For 12-core or 24-core, AMD Opteron nodes:

```
options ib_qib singleport=1 krcvqs=4 rcvhdrCnt=4096
```

To modify the previous `options ib_qib` line examples for two-port Host Channel Adapters (for example, QLE7342), divide the `krcvq` parameter by the number of active ports on the Host Channel Adapters. For example, an 8-core system would set the `krcvq` parameter to 4 for a dual-port adapter, and omit the single-port parameter as follows:

- ◆ For 8-core, dual HCA port (e.g. Intel Nehalem/QLE7342) nodes:

```
options ib_qib pcie_caps=0x51 krcvqs=4 rcvhdrCnt=4096
```

## 5.6 New Features

---

QLogic OFED+ Host software package version 6.0.2.1.11 does not have any new features.

QLogic OFED+ Host software package version 6.0.2.0.28 includes the following new features:

- ◆ OFED has been upgraded to version 1.5.2.
- ◆ iSCSI extensions for RDMA (iSER) no longer requires an update to the iSCSI RPMs supplied with the distro. When installing iSER, the iSCSI RPMs will not be updated by the installer.
- ◆ InfiniBand-Bonding is now separately selectable in the install menus. When planning to use Ethernet bonding, but not IPoIB bonding, QLogic recommends using the bonding driver supplied with the distribution.
- ◆ For forward compatibility command line installations, selecting `ipoib` will continue to install IPoIB and InfiniBand bonding. When selecting `ofed_ipoib`, only IPoIB will be installed. To explicitly select InfiniBand bonding, select `ofed_ib_bonding`.

QLogic OFED+ Host software package version 6.0.1.1.3 includes the following new feature:

- ◆ Platform Cluster Manager has been upgraded to version 2.0.1.

QLogic OFED+ Host software package version 6.0.1.0.30 includes the following new feature:

- ◆ OFED has been upgraded to version 1.5.1.

QLogic OFED+ Host software package version 6.0.0.2.7 and 6.0.0.1.2 do not have any new features.

QLogic OFED+ Host software package version 6.0.0.0.54 includes the following new features:

- ◆ OFED has been upgraded to version 1.5.
- ◆ CLI commands `INSTALL` and `iba_config` now have a `--answer keyword=value` option. This option permits command line control over selected questions during both interactive and non-interactive installations. Refer to the *QLogic FastFabric Command Line Interface Reference Guide* for a list of valid keywords.
- ◆ CLI command `INSTALL` now allows `-E` and `-D` options to be specified in conjunction with other installation options (such as `-U`, `-a` or `-i`). As such, a single command can install a component (or many components) and change its autostart to a non-default value. This option can be especially useful as part of the `FF_INSTALL_OPTIONS` or `FF_UPGRADE_OPTIONS` specified in `fastfabric.conf`.
- ◆ CLI command `INSTALL` now has a `--fwupdate` option that can force a firmware downgrade or unconditional firmware rewrite.
- ◆ CLI command `INSTALL` now supports aliases of `psm_mpi` and `verbs_mpi` to limit which MPIS are installed. CLI command `mpi` will install all available MPIS for both verbs and PSM.
- ◆ When Dispersive Routing is enabled, it allows packets sent via an MPI program run over PSM to take any one of several routes through a fabric, thus often increasing performance. The number of routes is determined by the value of 2 to the power of the Lid Mask Control setting (LMC). Because LMC defaults to 0, the default number of routes through the fabric is  $2^0$  or 1. LMC can be set as high as 3, allowing a total number of  $2^3$  or 8 routes through the fabric. Providing these additional routes can reduce fabric congestion, and thus improve performance.
- ◆ Boot over InfiniBand support has been provided using gPXE boot. This feature provides a mechanism by which a boot image is loaded onto a server without requiring the use of Ethernet.
- ◆ CLI commands `iba_config` and `INSTALL` now support a `-v` option. When run with this option, the version number for the software is reported.
- ◆ PSM and all QLogic MPIS now support dispersive routing.
- ◆ PSM and all QLogic MPIS now support the use of the QLogic Distributed SA for obtaining proper Path, vFabric, QOS and Partitioning information during job startup.
- ◆ PSM and all QLogic MPIS now support specification of the SL and PKey for a given job. This can be used when the QLogic Distributed SA is not being used.

- ◆ By default the installation will no longer start `iba_mon`. Use of `iba_mon` is not recommended when using the 6.0 QLogic FM and its new PM capabilities.
- ◆ The QLogic Distributed SA (`qlogic_sa`) is now available. This feature can be used by MPIs using the QLogic optimized PSM API to obtain proper vFabric, LMC and/or Mesh/Torus path information.
- ◆ CLI command `iba_saquery` now supports `vfinfo`, `vfinfocsv` and `vfinfocvs2` output to provide information about virtual fabrics in the QLogic FM.
- ◆ CLI command `iba_saquery` now supports query by pkey (`-k`), ServiceID (`0S`) or SL (`-L`) for `-o path`.
- ◆ CLI command `iba_saquery` now supports query by pkey (`-k`) for `-o mcmember`.
- ◆ Script file `/etc/init.d/iba_portconfig` is a new startup script that can be edited and enabled to force link speed and other link attributes on the Host Channel Adapter each time a server boots.
- ◆ The `do_build`, `do_mvapich_build`, `do_openmpi_build` and `do_mvapich2_build` scripts now support building for QLogic's optimized PSM interface and use the same `-qlc` suffix as the QLogic supplied pre-built MPIs. This interface can also be selected on the command line using the `-Q` option.
- ◆ CLI commands `p1info` and `p2info` now support a `-q` option, that shows the QSFP information for QLogic QDR Host Channel Adapters.
- ◆ CLI commands `iba_smaquery` and `iba_pmaquery` are new low-level tools that issue raw SMA or PMA queries to any node in the fabric.
- ◆ CLI command `iba_paquery` is a new low-level tool to issue queries to the 6.0 QLogic Fabric Manager's PM.

QLogic OFED+ Host software package version 5.1.0.2.1, 5.1.0.2.8, and 5.1.0.3.13 do not have any new features.

QLogic OFED+ Host software package version 5.1.0.0.49 includes the following new features:

- ◆ OFED has been upgraded to version 1.4.2.
- ◆ When installing/uninstalling the user will be prompted for any effects on iSCSI RPMs supplied by vendors other than OFED (indicated by "Voltaire Inc." in rpm VENDOR field). This will prevent QLogic OFED installations from unexpectedly uninstalling iSCSI.
- ◆ Installation prompts for rebuilding OFED RPMs are done prior to building any RPMs, therefore avoiding the need for the user to wait for each one to build before answering the next question. Also, dependency checks for builds are now done and all errors are reported prior to doing any builds. This check provides a concise summary of dependencies. It also allows the build to fail earlier, avoiding a partial install/partial build situation due to missed dependencies.

- ◆ CLI command `depcheck` option `--without` has been added to `INSTALL` to suppress checking of the OS dependencies. This option can allow an install to be forced even if the OS dependencies are missing. This command should be used with care, because using a forced install will typically result in operational problems due to missing OS libraries or tools.
- ◆ MPI Source component checks have been installed for all of the dependencies, which will be needed when building any of the included MPI sources (mvapich, openmpi, and mvapich2) with GNU compilers. Dependency checks for building MPI itself are done only once during the installation. Any subsequent removal of OS RPMs after the install could yield errors when rebuilding the MPIs.
- ◆ Dependency checking has been extended to include additional dependencies when rebuilding RPMs.

## 5.7 Product Constraints

---

- ◆ All installation and uninstallation of QLogic OFED+ Host software package components must be performed using the `INSTALL` or `iba_config` commands. If software is manually installed or uninstalled using other methods (RPM, other scripts, and so on), the installation on the system could become inconsistent and cause unreliable operation, in which case subsequent runs of `INSTALL` or `iba_config` may make incorrect conclusions about the configuration of the system and consequently make incorrect recommendations. If the system becomes inconsistently configured, QLogic recommends running the `./INSTALL` TUI and selecting `ReInstall` on all components. Once the re-installation has started, carefully review all prompts and choices.
- ◆ On SuSE systems, `NETWORKMANAGER` must be set to "no" in `/etc/sysconfig/network/config`. This is the default setting.
- ◆ The `mpirun` command in this release is not compatible with nodes running InfiniPath releases prior to version 2.2.1.
- ◆ OFED SDP has not been qualified for this release. IPoIB is recommended for data transfers.

## 5.8 Product Limitations

---

- ◆ When upgrading from a previous QLogic VNIC driver for QLogic InfiniPath 2.2.1 or earlier, the config file will not be automatically upgraded. Instead, the user must manually copy the `/etc/InfiniBand/qlogic_vnic.cfg` file to `/etc/InfiniBand/qlgc_vnic.cfg`. However, when upgrading from older QuickSilver releases, the `/etc/sysconfig/iba/ics_vnic.cfg` file will be automatically copied to `/etc/InfiniBand/qlgc_vnic.cfg`.
- ◆ QLogic products will auto-negotiate with devices that utilize IBTA-compliant auto-negotiation. When attaching QLogic products to a third-party switch, the bit error rate is optimized if the third-party switch utilizes attenuation-based tuning.

- ◆ At the time of this release, Oracle has not certified InfiniBand adapters for RDS. As a result, QLogic is unable to extend support for RDS at this time.
- ◆ At the time of this release, the PathScale Compiler Suite is not supported for SLES11. As a result, QLogic is unable to extend support for this environment.
- ◆ The 10m Amphenol Cables have exhibited inconsistent reliability and are not recommended for use with TrueScale InfiniBand Adapters.

## 5.9 Other Information

---

- ◆ The recommended tuning for nodes with dual six-core Intel Xeon 5600 Series (Westmere) processors is to configure the Truescale Host Channel Adapters for 14 contexts, for improved message rate and small message performance. The `cfgctxts` driver parameter can be specified in the `/etc/modprobe.conf` file to configure the contexts appropriately. For example, add the following line to `modprobe.conf`:

```
options ib_qib pcie_caps=0x51 cfgctxts=14
```

- ◆ Note that the current release of `mvapich2` may have a memory management issue on certain machines and certain MPI applications. If MPI applications hang when using `MVAPICH2`, try using these settings with your application:

```
MV2_IBA_EAGER_THRESHOLD=16384 MV2_VBUF_TOTAL_SIZE=16384
```

When using QLogic's sample applications, these settings can be passed using the `ofed.mvapich2.params` file (in `/opt/iba/src/mpi_apps` folder). For other applications, they can be passed as part of the `mpirun` command.

- ◆ When running `MVAPICH2`, QLogic recommends turning off RDMA fast path. To turn off RDMA fast path, specify `MV2_USE_RDMA_FAST_PATH=0` in the `mpirun_rsh` command line or set this option in the parameter file for `mvapich2`.
- ◆ Older QLogic InfiniPath 2.2.1 and prior releases will not be recognized in the installation menus and will be treated as if no InfiniBand software is installed.
- ◆ When upgrading from an older QLogic OFED+ release to a new QLogic OFED+, the older QLogic OFED+ release will be uninstalled first.
- ◆ When installing the OFED Debug Info component, make sure that `~/rpmmacros` does not have any lines changing the `%debug_package` macro.

If `%debug_package` is set to `%{nil}`, debuginfo RPMs will not be available for installation.

The current setting for this parameter can be verified using:

```
rpm --eval '%{debug_package}'
```

- ◆ If a non-blank output occurs, then debuginfo is enabled. If a blank line is output, then debuginfo is disabled on the present system.
- ◆ When running QLogic SRP over two ports in round robin mode, performance is degraded.
- ◆ When running verbs-based benchmarks such as `ib_send_bw`, the following message may appear:

```
BUG: soft lockup - CPU#1 stuck for 10s.
```

This message does not impact completion or correctness of the benchmark.

- ◆ The `ib_send_bw` benchmark, when run in UC mode, is written such that it will hang if even one packet is dropped.
- ◆ The QLogic 12000 series switch firmware must be at version 6.0.0.1.2 or later to connect with the TrueScale Host Channel Adapters in this release.
- ◆ QLE7240/QLE7280 performance with Barcelona stepping B2 and the BIOS/kernel workaround for AMD errata 298 can lead to reduced InfiniBand bandwidth, as a result of poor memory bandwidth and latency on the adapters due to the workaround. To fix this problem, upgrade to Barcelona stepping B3 (or later).
- ◆ `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
```

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)
```

- ◆ Additional information on interrupt request (IRQ) affinity settings. The following information will be added to the information in the *QLogic OFED+ Host Software Users Guide*, Appendix F Troubleshooting, section Performance Issues, sub-section Erratic Performance:

To immediately change the processor affinity of an IRQ, execute a command similar to the following, as a root user:

```
echo 01 > /proc/irq/$my_irq/smp_affinity
```

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. Look at the stats in `/proc/interrupts` while the adapter is active to see which CPU is fielding `ib_qib` interrupts.

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

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

This problem is caused by a limitation in Intel MPI/uDAPL's dynamic connection mechanism when MPI processes are not paying sufficient attention to incoming interconnect traffic. To work around this problem, add the following to the mpirun command:

```
genv I_MPI_USE_DYNAMIC_CONNECTIONS 0
```

- ◆ An SRP write performance bug between QLE7240/QLE7280 and LSI (Engenio) storage systems was fixed. To apply this fix, 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 `srp_sg_tablesize` when adding the `ib_srp` module (as root):

```
modprobe ib_srp srp_sg_tablesize=80
```
  - If using `srp_daemon`, modify `max_sect` by adding the following command to `/etc/srp_daemon.conf` file:

```
a id_ext=200500A0B81146A1,ioc_guid=00a0b80200402bef,max_sect=2048
```
- ◆ To ensure proper operation of MPI, the QLogic OFED+ INSTALL sets the memory locking limits in the `/etc/security/limits.conf` file to unlimited, due to the large amounts of memory needed to run the MPIs. Open MPI recommends using unlimited ulimits for lockable memory. For more information, see:

<http://www.open-mpi.org/faq/?category=openfabrics#ib-locked-pages>

An alternative is to use the `--mca mpi_leave_pinned 1` option with mpirun.

- ◆ When users increase the number of processes beyond the number of open files allowed by ulimit, mpirun prints an error message. The ulimit for the number of open files is typically 1024 on both Red Hat and SLES. The message will look similar to the following:

```
MPIRUN.up001: Warning: ulimit for the number of open files is only
1024, but this mpirun request requires at least number of files open
files (sockets). The shell ulimit for open files needs to be increased.
This is due to limit:
descriptors 1024
```

The ulimit can be increased; QLogic recommends an increase of approximately 20 percent over the number of CPUs. For example, in the case of 2048 CPUs, ulimit could be increased to 2500: `ulimit -n 2500`. Increase the ulimit only on the host where mpirun was started, unless the mode of operation allows mpirun from any node.

- ◆ The MPI sample applications installed in `/usr/local/src/mpi_apps` do not support the Intel 11 compiler. When using the Intel 11 compiler, the HPL sample application is expected to fail to build.

QLogic recommends using the GNU compilers, or earlier versions of the Intel compiler when using these sample applications. Alternatively, for HPL runs, HPL can be found on the web along with alternative Math Libraries that may work with newer versions of the Intel compilers.

## 6 Trademarks

---

Accelera, Accelerating Cluster Performance, FastFabric, InfiniCon Systems, InfiniNIC, InfiniPath, InfiniView, Intelligent NIC, Networking University, NetXen, QLogic, the QLogic logo, ReadyPath, SANDoctor, SANSurfer, and SilverStorm are registered trademarks of QLogic Corporation.

To the extent a name or logo does not appear on this list does not constitute a waiver of any and all intellectual property rights that QLogic Corporation or its subsidiaries has established in any of its product, feature, or service names or logos. 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>

E-mail: [support@qlogic.com](mailto:support@qlogic.com)

[Go to Top](#)



© Copyright 2011. All rights reserved worldwide. QLogic, the QLogic logo, and the Powered by QLogic logo are registered trademarks of QLogic Corporation. InfiniBand is a registered trademark and service mark of the InfiniBand Trade Association. OpenFabrics Enterprise Distribution (OFED™) is a trademark of OpenFabrics, Inc. Red Hat is a registered trademark of Red Hat, Inc. SuSE is a registered trademark of Novell, Inc. Lustre is a registered trademark of Cluster File Systems, Inc. Intel is a registered trademark of Intel Corporation. All other brand and product names are trademarks or registered trademarks of their respective owners.