

# Open Fabrics Enterprise Distribution (OFED)

Version 3.5-rc5

Release Notes

January 2013

## Table of Contents

1. Overview, which includes:
  - OFED Distribution Rev 3.5 Contents
  - Supported Platforms and Operating Systems
  - Supported HCA and RNIC Adapter Cards and Firmware Versions
  - Tested Switch Platforms
  - Third party Test Packages
  - OFED sources
2. Change log
3. Known Issues

## 1. Overview

These are the release notes of OpenFabrics Enterprise Distribution (OFED) release 3.5. The OFED software package is composed of several software modules, and is intended for use on a **computer cluster constructed as an InfiniBand Fabric, an iWARP Network or a RoCE Fabric.**

Note: If you plan to upgrade the OFED package on your cluster, please upgrade all of its nodes to this new version.

### 1.1 OFED 3.5 Contents

The OFED package contains the following components:

- OpenFabrics core and ULPs:
  - IB HCA drivers (mthca, mlx4, qib, ehca)
  - iWARP RNIC driver (cxgb3, cxgb4, nes)
  - **ROCE and Ethernet support**
  - core
  - Upper Layer Protocols: IPoIB, SRP Initiator and target, iSER Initiator and target, RDS, uDAPL, qlgc\_vnic and NFS-RDMA\*.
- OpenFabrics utilities:
  - OpenSM (OSM): InfiniBand Subnet Manager
  - Diagnostic tools
  - Performance tests
- Extra packages:
  - infinipath-psm: Performance-Scaled Messaging API, an accelerated interface to Intel(R) HCAs
- Sources of all software modules (under conditions mentioned in the modules' LICENSE files)
- Documentation

## 1.2 Supported Platforms and Operating Systems

---

- o CPU architectures:
  - x86\_64
  - x86
  - ppc64
  - ia64
  
- o Linux Operating Systems:
  - RedHat EL6.2                    2.6.32-220.el6
  - RedHat EL6.3                    2.6.32-279.el6
  - SLES11 SP2                    3.0.13-0.27-default
  - kernel.org                    3.5        \*

\* Minimal QA for these versions.

## 1.3 HCAs and RNICs Supported

---

This release supports IB HCAs by **IBM, Intel** and Mellanox Technologies, **iWARP RNICs** by **Chelsio Communications** and **Intel** and **RoCE adapters** by **Emulex, IBM** and **Mellanox**.

### InfiniBand Adapters

- o IBM HCAs:
  - GX Dual-port SDR 4x IB HCA
  - GX Dual-port SDR 12x IB HCA
  - GX Dual-port DDR 4x IB HCA
  - GX Dual-port DDR 12x IB HCA
  
- o Intel (formerly QLogic) HCAs:
  - Intel(R) True Scale DDR PCIe x8 and x16 HCAs
  - Intel(R) True Scale QDR PCIe x8 Gen2 HCAs
  
- o Mellanox Technologies HCAs (SDR, DDR and QDR Modes are Supported):
  - ConnectX(R) and ConnectX EN (fw-25408 Rev 2.9.1000)
  - ConnectX-2 (fw-ConnectX2 Rev 2.9.1000)
  - ConnectX-3 (fw-ConnectX3 Rev 2.10.0700) (FDR and FDR10 Modes are Supported)

For official firmware versions please see:

[http://www.mellanox.com/content/pages.php?pg=firmware\\_download](http://www.mellanox.com/content/pages.php?pg=firmware_download)

### iWARP Adapters

- o Chelsio RNICs:
  - S310/S320 10GbE Storage Accelerators
  - R310/R320 10GbE iWARP Adapters
  
- o Intel RNICs:
  - NE020 10Gb iWARP Adapter

## RoCE Adapters

- o Emulex
  - Emulex® OneConnect Universal CAN for RoCE
- o IBM
  - IBM Flex System EN4132 2-port 10 GbE RoCE
- o Mellanox
  - ConnectX-2 EN (fw-ConnectX2 Rev 2.9.1200)
  - ConnectX-3 EN (fw-ConnectX3 Rev 2.10.0800)

## 1.4 Switches Supported

-----  
This release was tested with switches and gateways provided by the following companies:

- Flextronics
- Intel (formerly QLogic)
- Mellanox
- Oracle (Sun)

## 1.5 Third Party Packages

-----  
The following third party packages have been tested with **OFED 3.2**:

- MVAICH2 - 1.8
- Intel MPI - 4.0.3
- Open MPI - 1.6/1.6.1

## 1.6 OFED Sources

-----  
All sources are located under `git://git.openfabrics.org/`

Linux:

-----  
URL: `git://git.openfabrics.org/compat-rdma/linux-3.5.git`  
Branch: `for-ofed-3.5`

- Linux kernel subtree that includes files relevant for the OFED project only.  
Based on v3.5. Used to shorten git clone time.  
Note: the regular Linux git tree can be used as well.

compat:

-----  
URL: `git://git.openfabrics.org/compat-rdma/compat.git`  
Branch: `ofed`

- Based on compat project (<https://github.com/mcgrof/compat>). The compat module provides functionality introduced in newer kernels to older kernels through a set of header files and exported symbols.  
See <https://github.com/mcgrof/compat/wiki> for details.
- Used to replace `kernel_addons` in the previous OFED kernel tree.

compat-rdma:

-----

URL: [git://git.openfabrics.org/compat-rdma/compat-rdma.git](http://git.openfabrics.org/compat-rdma/compat-rdma.git)

Branch: master

User level Sources are downloaded from <http://www.openfabrics.org/downloads/>  
as written in the BUILD\_ID

The kernel sources are based on Linux 3.5 mainline kernel. Its patches  
are included in the OFED sources directory.  
For details see HOWTO.build\_ofed.

The list of maintainers is available under:  
<http://www.openfabrics.org/downloads/MAINTAINERS>

=====  
2. Change log  
=====

OFED-3.5-rc5 Main Changes from OFED 3.5-rc4  
-----

perfctest-2.0-0.25.g971259d

OFED-3.5-rc4 Main Changes from OFED 3.5-rc3  
-----

RDMA/nes: Fixes for OFED-3.5 RC3  
RDMA/cxgb4: Keep the maximum number of stag limited to T4\_MAX\_NUM\_STAG  
Updated packages:  
    infinipath-psm-3.1-364.1140\_open

OFED-3.5-rc3 Main Changes from OFED 3.5-rc2  
-----

RDMA/nes: Fix for terminate timer crash  
RDMA/nes: Fix for crash when adding timer and timer is pending  
RDMA/nes: Fix for incorrect multicast address  
RDMA/nes: Fix for crash when adding timer  
RDMA/nes: Locking fix to prevent race conditions  
RDMA/nes: Changing nes driver version  
RDMA/nes: Fixes for PAU handling  
cxgb4: Initialize data structures before using.  
cxgb4: Fix initialization of SGE\_CONTROL register  
cxgb4: Add embedded firmware and its configuration file support  
cxgb4/iw\_cxgb4: Add LE hash collision bug fix for active and passive offloaded connections  
cxgb4: Fix unable to get UP event from the LLD  
cxgb4: Remove unnecessary #ifdef condition  
cxgb4: Fix build warnings introduced after T4 configuration file support  
cxgb4: Add support for T4 configuration file  
iw\_cxgb4: only log rx\_data warnings if cpl status is non zero.  
RDMA/cxgb4: fix mpa version mismatch  
RDMA/cxgb4: abort connections that receive unexpected streaming mode data  
RDMA/cxgb4: Insert hwtid in pass\_accept\_req instead in pass\_establish  
RDMA/cxgb4: Remove the limitation on maximum number of STAGs  
linux-next-cherry-picks: Fix build failure for kernel < 3.2  
IB/qib: remove EBUSY patch  
ib\_srp: Avoid that module removal can trigger a deadlock  
ib\_srp: Unbreak build on SLES 11 SP2

ib\_srp: Add post-3.5 upstream patches

ib\_srp: Avoid use-after-free during module unload on pre-2.6.36 kernels

Updated packages:

mstflint-1.4-1.23.g5c2032a

perftest-2.0-0.8.gaeb562e

qperf-0.4.9

OFED-3.5-rc2 Main Changes from OFED 3.5-rc1

---

compat-rdma: Add SRP backport

compat-rdma: /etc/init.d/openibd: Fix LSB header

compat-rdma: IB/qib: linux-3.6 patches backported

compat-rdma: iw\_cxgb4: Fix bug 2369 in OFED bugzilla

compat-rdma: IB/qib: fix compliance regression in 3.5

compat-rdma: RDMA/nes: Add linux-next-pending patches

compat-rdma: RDMA/nes: Updated backports

compat-rdma: NFSRDMA RHEL6.3 backport

compat-rdma: NFSRDMA SLES11SP2 backport

compat-rdma: linux-next-cherry-picks: RDMA/ucma.c: Different fix for ucma context uid=0, causing iWarp RDMA applications to fail in connection establishment

Updated packages:

infinipath-psm-3.0.1-115.1015\_open

perftest-1.4.0-0.80.gd1763bd

qperf-0.4.7-0.2.gf3f7001

OFED-3.5-rc1 Main Changes from OFED 3.2-rc1

---

1. Based on linux-3.5 kernel

2. Updated packages:

ibacm-1.0.7

librdmacm-1.0.16

dapl-2.0.36

Management:

libibumad-1.3.8

infiniband-diags-1.6.1

opensm-3.3.15

3. Removed libsdp, sdpnetstat

OFED-3.5-rc1 Main Changes from OFED 1.5.4.1

---

1. Based on linux-3.5 kernel

2. kernel-ib RPM renamed to compat-rdma

3. Removed MPI packages

4. Updated packages:

libibverbs-1.1.6

libmlx4-1.0.4

dapl-2.0.35

ibacm-1.0.6

perftest-1.4.0-0.66.g9074687

Management:

libibmad-1.3.9  
opensm-3.3.14  
infiniband-diags-1.6.0

OFED-1.5.4 Main Changes from OFED 1.5.3.2

---

1. Add kernel 3.0 support
2. Add RHEL6.2 support
3. Add NFSRDMA to the list of packages to be installed by default on supported platforms
4. Add new InfiniBand link speeds: EDR/FDR/FDR10
5. OFA Kernel: See ofa\_kernel.changelog.txt
6. Updated packages:
  - libibverbs-1.1.4-1.24.gb89d4d7
  - libmlx4-1.0.1-1.20.g6771d22
  - libnes-1.1.3
  - libcxgb3-1.3.1
  - libcxgb4-1.2.0
  - dapl-2.0.34
  - ibacm-1.0.5
  - librdmacm-1.0.15
  - infinipath-psm-2.9-926.1005\_open
  - ibutils-1.5.7-0.1.g05a9d1a
  - perftest-1.3.0-0.56.gd6f6093
  - mstflint-1.4-1.18.g1adcfbf

Management:

libibmad-1.3.8  
infiniband-diags-1.5.12  
opensm-3.3.12

MPI:

mvapich2-1.7-5.2011svn5031

OFED-1.5.3.2 Main Changes from OFED 1.5.3.1

---

1. Add RHEL5.7 support
2. Add RHEL6.1 support
3. Add kernel.org 2.6.37 - 2.6.39 support
4. ib-bonding-0.9.0-43 (SLES10 SP4 support)

OFED-1.5.3.1 Main Changes from OFED 1.5.3

---

1. Fix QIB support on SLES10 SP3

OFED-1.5.3 Main Changes from OFED 1.5.2

---

1. Updated packages:
  - dapl-2.0.32

librdmacm-1.0.14.1  
ibacm-1.0.4  
libibmad-1.3.7  
libibumad-1.3.7  
infiniband-diags-1.5.8  
opensm-3.3.9  
ibutils-1.5.7  
libibverbs-1.1.4-1.22.g7257cd3.tar.gz  
libmlx4-1.0.1-1.18.gb810a27.tar.gz  
libnes-1.1.1  
infinipath-psm-1.14  
libsdp-1.1.108-0.15.gd7fdb72  
mstflint-1.4-0.9.gc683022  
sdpnetstat-1.60-0.3.gb6e7425  
perftest-1.3.0-0.42.gf350d3d

- MPI

openmpi-1.4.3-1  
mvapich2-1.6-2

2. Added RHEL6 support
3. Added RHEL5.6 support
4. Improved RAW Ethernet QP support (requires ConnectX/ConnectX2 FW >= 2.8.0600)
5. NFSoRDMA was removed from the list of packages in the install script to prevent its installation by default.

To force NFSoRDMA installation, add the following lines to the ofed.conf:

```
nfsrdma=y  
rnfs-utils=y
```

To get ofed.conf, run

```
# ./install.pl -p
```

After adding the lines above to ofed.conf, to install OFED run

```
# ./install.pl -c ofed.conf
```

#### OFED-1.5.2 Main Changes from OFED 1.5.1

- 
1. Added RHEL5.5, SLES11 SP1 and kernel 2.6.33 support
  2. /etc/modprobe.conf deprecated  
ib\_ipoib use /etc/modprobe.d/ib\_ipoib.conf  
ib\_sdp use /etc/modprobe.d/ib\_sdp.conf  
mlx4\_en use /etc/modprobe.d/mlx4\_en.conf
  3. Added RAW Ethernet QP support in libibverbs
  4. Added multicast support in perftest
  5. Added new package: ibacm. See ibacm\_release\_notes.txt
  6. Added new package infinipath-psm.
  7. Updated packages:
    - libibverbs to rev 1.1.4-0.14.gb6c138b
    - libehca to rev 1.2.2-0.1.g69e1a88
    - libmlx4 to rev 1.0-0.13.g4e5c43f
    - libnes to rev 1.1.0
    - librdmacm to rev 1.0.13
    - dapl to rev 2.0.30

compat-depl to rev 1.2.19  
libsdp to rev 1.1.103  
mstflint to rev 1.4-0.6.g63179f4  
perftest to rev 1.3.0-0.27  
rds-tools to rev 2.0.4  
Management:  
  opensm to rev 3.3.7  
  libibmad to rev 1.3.6  
  libibumad to rev 1.3.6  
  infiniband-diags to rev 1.5.7  
MPI:  
  openmpi to rev 1.4.2-1  
  mvapich2 to rev 1.5.1-2.p1  
  mpitests to rev 3.2-923

#### OFED-1.5.1 Main Changes from OFED 1.5

---

1. Added RoCEE support - see RoCEE\_README.txt
2. Added enhanced atomic operations to ConnectX (kernel only).  
  See mlx4\_release\_notes.txt.
3. Updated Open MPI to rev 1.4.1-2ofed
4. Updated MVAPICH2 to rev 1.4.1
5. Updated DAPL to rev 2.0.27
6. Updated libnes to rev 1.0.1
7. Updated librdmacm to rev 1.0.11
8. Removed tvflash RPM
9. NFS-RDMA is not supported on SLES10 SP3
10. Fixed IPv6 support and IPv4 routing corner cases for RDMA CM
11. Bug fixes

#### OFED-1.5 Main Changes from OFED 1.4.2

---

Note: For details regarding the various changes, please see the release notes for each package in the docs directory.

##### 2.1 General changes

- o Kernel code based on 2.6.30
- o libraries location - all userspace libraries can be downloaded from <http://www.openfabrics.org/downloads/>  
  See BUILD\_ID for exact locations
- o Qlogic moved the low level driver from ipath to qib.

##### 2.2 SDP

- o Performance improvements
- o Zero copy in beta level

##### 2.3 uDAPL

- o New UCM provider (ofa-v2-mlx4\_0-1u) with IB UD-based CM per process.  
  It is more scalable than rdma\_cm (cma) or socket cm (scm).
- o Common code base with WinOF 2.1
- o Bug fixes



## 2.4 perfctest

### o Renamed tests:

- ib\_rdma\_bw -> rdma\_bw
- ib\_rdma\_lat -> rdma\_lat

## 2.5 Management

### o OpenSM

- Support for Mesh Analysis for LASH routing algorithm
- Reloadable OpenSM configuration (preliminary implementation)
- Routing paths sorted balancing (for UpDown and MinHops)
- Weighted LID matrices calculation (for UpDown, MinHop and DOR)
- I/O nodes connectivity (for FatTree)

## 2.6 MPI:

- OSU MVAPICH 1.2.0
- Open MPI 1.4
- OSU MVAPICH2 1.4
- MPI tests 3.2

## 2.7 iSER:

### o Available on:

- kernel.org: 2.6.30 and higher
- RHEL5.4

## 2.8 NFS-RDMA

- o Added support for RHEL5.4, SLES10 SP3, kernel.org 2.6.25 and 2.6.30. Kernels 2.6.26 and 2.6.27 are not supported
- o NFS-RDMA is in Beta level

---

## 3. Known Issues

---

The following is a list of general limitations and known issues of the various components of the OFED 3.5-rc2 release.

1. When upgrading from an earlier OFED version, the installation script does not stop the earlier OFED version prior to uninstalling it.  
Workaround: Stop the old OFED stack (/etc/init.d/openibd stop) before upgrading to OFED 3.5 or reboot the server after OFED installation.
2. Memory registration by the user is limited according to administrator setting. See "Pinning (Locking) User Memory Pages" in OFED\_tips.txt for system configuration.
3. Fork support from kernel 2.6.12 and above is available provided that applications do not use threads. fork() is supported as long as the parent process does not run before the child exits or calls exec(). The former can be achieved by calling wait(childpid), and the latter can be achieved by application specific means. The Posix system() call is supported.
4. The qib driver is supported only on 64-bit platforms.
5. IPoIB: brctl utilities do not work on IPoIB interfaces. The reason for that is that these utilities support devices of type Ethernet only.
6. In case uninstall is failing, check the error log and remove the remaining RPMS manually using 'rpm -e <rpms list>'.

7. On SLES11.x, set `allow_unsupported_modules` parameter to 1 in file:  
/etc/modprobe.d/unsupported-modules. Without this the modules will not  
load.

Note: See the release notes of each component for additional issues.