

OFA IWG meeting – 12/18/2012

Attendees

First	Last	Company	11/06/12	11/20/12	12/04/12	12/18/12
David	Duffey	Canonical				
Samantha	Jian-Pielak	Canonical		X		X
Tom	Reu	Chelsio		X		X
Martin	Schlining	DataDirect Networks		X		
Jeff	Kopko	Emulex				
Harry	Cropper	Intel	X	X		X
Mitko	Haralanov	Intel		X		X
Jess	Robel	Intel	X		X	X
Jim	Ryan	Intel				X
Guy	Ergas	Mellanox				X
Rupert	Dance	Software Forge	X	X	X	X
Glenn	Martin	UNH-IOL		X		X
Edward	Mossman	UNH-IOL	X	X	X	X
Bob	Noseworthy	UNH-IOL	X	X		X
Nate	Rubin	UNH-IOL		X	X	

Events

2012 Events

- OFA Interop Logo GA Event 01/07/2013

April 2013 dates

- IBTA PF23 04/01 → 04/12/2013
- OFA April Interop Debug Event 04/08 → 04/12/2013

October 2013

- IBTA PF24 10/07 → 10/18/2013
- OFA Interop Debug Event 10/14 → 10/18/2013

OFA EWG Status

- OFED 3.5 RC3 12/11/2012 – see [release notes](#)
- OFED 3.5 RC4 Release after [Bug 2409](#) is resolved – early January 2013
Unable to build Open MPI 1.6.1/1.6.3 with PSM!
- OFED 3.5 GA One week after RC4
- Finals are over by 12/14/2012 and then there are the Holidays. So GA Logo Testing will start after OFED 3.5 GA is released which hopefully will be the second full week of January 2013

OFA Interop Debug status

- iWARP Report
 - All good
- IB Report
 - There are some issues that UNH-IOL has been working on with affected vendors – confidential!
- RoCE status
 - Testing good – MPI went as expected. RoCE is in early Beta state so it handles a limited test set.
 - In April 2013 Emulex expects to provide full support for OFED 3.x
 - For January 2013, they understand that they will be part of a Beta list
- OFED SRP status
 - Still some issues – Edward is waiting for feedback from Marty.

Review of Interop Test Plan updates

- 1) We reviewed the changes made in the latest OFILG Interoperability Test Plan. These mainly included the changes made by Samantha from Canonical which document the changes required to run the test Plan on a system using Ubuntu 12-04 or 12-10.
 - a) [Interop Test Plan version 1.46](#)
 - b) [Ubuntu - OFA Interop testing notes](#)
- 2) Samantha is still working on Open-MPI test – she will provide update by next meeting

Review of Suggested Interop Program Changes

We reviewed the proposed changes to the OFILP which are necessary to address the budget shortfall of the last few years which the OFA has funded through their general fund. See [Suggested Interop Program Changes for 2013](#)

OFILP expenses: Reduce program costs

- 1) Do no testing for companies that are not OFILG members
 - a) Microsoft
 - b) Oracle (RDS)
 - c) RedHat
 - d) SUSE
 - e) Ubuntu
- 2) Eliminate ULP Testing that is no longer supported or used by OFILG member companies
 - a) iSER
 - b) RDS (Oracle)
 - c) SDP
 - d) MVAPICH
- 3) Eliminate or ask for support from companies that support ULP for marginal protocols
 - a) NFSoRDMA
 - b) uDAPL (IBM GPFS and DB2)
 - c) Vendor SM testing (only test OpenSM)

Comments

Bob Noseworthy

- 1) There is not a linear relation between amount of testing and cost. Even if you are able to reduce the Logo Test time from 4 weeks twice per year to two weeks twice per year, you still need to maintain the cluster, the engineers, the know-how and skills and you have to pay the heating, cooling and electrical bills.
- 2) We have agreed before that \$160K is the floor to operate this program and there comes a point where this just is not feasible for UNH-IOL. I agree that shaving costs is good but the focus must come from the group to achieve our goals and get more participation. We all want program to be successful.

Jim Ryan

- 1) We can continue business as usual or put the whole program at risk. Good middle ground to carefully cut costs and program where we can. Our goal is to grow the program but in the interim we must do both.

Suggested Changes to program from Intel - Concerns coming from members

1) **Qualified Pass**

- a) Allow a product on the Logo List as long as all issues and defects are very clearly outlined. The program is currently black and white. With this change, a product would be allowed on this list with permission and agreement from vendor, and UNH-IOL would provide the details of what passed and what failed.

2) **Arbitration Policy**

- a) Allow additional OFA members to participate in Arbitration process. For example Oracle, IBM, HP and national labs who are members of the OFA.
- b) OFILG members are disqualified from involvement in the arbitration process unless their products are involved.
- c) The OFA IWG will continue to maintain a standing committee (currently Software Forge and UNH-IOL) but we can bring in additional OFA members based on the agreement of the arbitration parties involved.

3) **Components to test**

- a) Allow more feedback from OFILG members regarding qualified products to be included in Interop Logo testing. The concern is to prevent the inclusion of “ringer” products that could cause another OFILG member product to fail.
- b) Potential guidelines for this enhancement: A tested product must have a legitimate or current SKU and be available to the general public within 6 months of the event.

4) **More Transparency – shared objective**

- a) The goal is to make Interop issues more broadly known to the OFA community and the end users whenever possible without breaking NDA guidelines. Be more focused on the needs of the end users.

Comments

- 1) Bob Noseworthy is in broad agreement with OFILP Changes noted above but he is worried about the transparency request because of NDA. He suggests we work on this at Developer Conference.

Opens

- 1) **Bob:** What is OFA doing to include Open Stack and other technologies such as DCB? There is a need for some organization to work with all the competing Ethernet Fabric management solutions like Brocade VCS, SPB, Mellanox UFM, TRILL and others.
 - a) Bob can't do inter-alliance communication but he will facilitate wherever possible.
- 2) **Jim:** wide spread concern that IB verbs are too heavy and can't be part of Exascale. We want OFA to stay relevant and part of Data Center.

Action Requests (ARs)

Rupert Dance

12/04/2012

- 1) Check with all participants of the Interop Debug Event to make sure issues are resolved – [email sent out 12/5/2012](#)
- 2) Log SRP bug with RedHat

11/20/2012

- 1) Check October dates with Mellanox – [sent email directly on 11/20/2012](#)

Older

- 1) Update Logo Program to include OFA Policy regarding OFILG Membership and granting of the Logo
- 2) Check with EWG and XWG about the distribution of OFED Binaries
- 3) Create Logo with version or change Logo Guidelines – [see 4/3/2012 minutes](#)
- 4) Incorporate ICR into Logo Program

UNH-IOL

12/04/2012 - Nate

- 1) The RoCE test plan currently does not reflect the additional tests that we performed (uDAPL and MPI).

11/6/2012 – Edward

- 1) Help Samantha with RDMA Basic Interop tests on Ubuntu
- 2) Work with Mitko and Samantha to see if they want to test Intel HCAs on Ubuntu

10/9/2012

- 1) **Nate:** verify that rping works for RoCE as well as iWARP
- 2) IOL Check the OpenMPI commands for RoCE

7/17/2012

- 1) Bob Noseworthy: provide Michael and OFA IWG a list of DCB vendors who support RoCE

5/1/2012

- 1) Make a separate webpage page to show how Distros have been installed. Include a sanitized kick start file and the update logs.

Open Fabrics Enterprise Distribution (OFED)
Version 3.5-rc3
Release Notes
December 2012

=====
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 subnet or iWARP network.

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)
 - 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 QLogic 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 Mellanox Technologies, Qlogic and IBM as well as iWARP RNICs by Chelsio Communications and Intel.

- 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-2 EN (fw-ConnectX2 Rev 2.9.1200)
 - ConnectX-3 (fw-ConnectX3 Rev 2.10.0700) (FDR and FDR10 Modes are Supported)
 - ConnectX-3 EN (fw-ConnectX3 Rev 2.10.0800)

For official firmware versions please see:

http://www.mellanox.com/content/pages.php?pg=firmware_download

- o Qlogic HCAs:
 - QLogic TrueScale DDR PCIe x8 and x16 HCAs
 - QLogic TrueScale QDR PCIe x8 Gen2 HCAs

- 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 Chelsio RNICs:
 - S310/S320 10GbE Storage Accelerators
 - R310/R320 10GbE iWARP Adapters

- o Intel RNICs:
 - NE020 10Gb iWARP Adapter

1.4 Switches Supported

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

- Qlogic
- Flextronics
- Oracle (Sun)
- Mellanox

1.5 Third Party Packages

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

- MVAPICH2 - 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`

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-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-3.2 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