

OFA IWG meeting – 11/20/2012

Attendees

First	Last	Company	09/25/12	10/09/12	11/06/12	11/20/12
David	Duffey	Canonical	X			
Samantha	Jian-Pielak	Canonical				X
Tom	Reu	Chelsio	X	X		X
Martin	Schlining	DataDirect Networks	X			X
Jeff	Kopko	Emulex				
Harry	Cropper	Intel	X	X	X	X
Mitko	Haralanov	Intel	X	X		X
Jess	Robel	Intel	X	X	X	
Jim	Ryan	Intel				
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
Bob	Noseworthy	UNH-IOL		X	X	X
Nate	Rubin	UNH-IOL	X	X		X
Bill	Weber	Unisys	X	X		

Events

2012 Events

- OFA December 2012 Interop GA Event **TBD**

April 2013 dates

- IBTA PF 23 04/01 → 04/12/2013
- OFA April Interop Debug Event 04/08 → 04/12/2013

October 2013 IBTA PF24 and OFA Interop Debug Event **Proposed**

- Option 1 October 7-18th
- Option 2 October 14-25th

Causeway

- 1) This is the new OFA collaboration tool and will be used to upload the OFILG documents. To get an account please click the following link: <https://ofa.causewaynow.com/>
- 2) Meeting Minutes from 11-06-2012
 - a) <https://ofa.causewaynow.com/wg/OFLIG/document/73>

SC12 update

- Attendance was approximately 10K which was down a little because the National Labs were not allowed to attend because of the General Services Administration (GSA) scandal.
- From now on SC will rotate in one of the following cities
 - Denver, NO, Portland, SLC
- Report on booth activities
 - We had a series of experts present including Alan Benner from IBTA IWG, David Duffey from Canonical and Josh Simons from VMware and Katherine for Finisar.
- Report on OFA Cluster Demo at SC12

- We accessed the OFA Cluster from SC12 and demoed all of the Interop Test Suite and explained how we did compliance and Interop testing for the IBTA and the OFA.

IBTA and OFA Test Event updates

IBTA Plugfest 22 status

- Very time consuming especially because we developed new Physical layer test suite for FDR devices
- Draft IL list by end of year

OFA Interop Debug status

- iWARP Report - Edward
 - No new news – good news because everything seemed to pass successfully
 - Harry is happy and so is Tom.
- IB Report - Edward
 - Reminder that there will be no reports sent from Interop Debug events but emails will be sent to any vendor with products that are experiencing failures.
 - There were initially some issues with scripts but we got NFSoRDMA working correctly and everything seems to be working well.
 - Fabric Init is 100% OK
 - IPoIB, Open MPI, RDMA, uDAPL all products pass except when using one vendor product. These issues are isolated to this one product. However OPEN MPI run is failing because of this issue.
 - RDS has not been run because Edward recalls that this was given a lower priority for this event.
- SRP – Marty
 - For the October 2012 Interop event, there are now two different tests – the standard one we have been running for many years and a new test which includes the testing of large IO operations (see updated test plan) . Larger SCSI commands can be received at the IO Target. It has been in kernel since 2.6.38 and 2.6.39. However OFED 3.5 RC2 and Daily builds do not match up with RHEL 6.2 and 6.3
 - There have been 4 bugs which caused the system to crash. The first three of the 4 issues have been resolved.
 - [Bug 2374](#): Crash when SRP targets are loaded
 - [Bug 2379](#): Hung task in kernel, also system locks up when primary path to storage is lost
 - [Bug 2396](#): Unloading ib_srp taking longer than usual
 - [Bug 2395](#): Kernel crash using SRP extended sg descriptors, increased max_sectors_kb, when running direct IO.
 - Large IO changes still giving problems. It is in Linux Kernel and not SRP driver.
 - Really does not want OFED 3.5 without this being resolved.
 - If we run both tests with present test script, you won't see it. It only occurs when you increase max_sectors_kb to 2048 and use direct IO you will crash the kernel. SRP caught this but kernel may be susceptible otherwise.
 - SL 6.2, SL6.3, RHEL 6.2 and 6.3 are in the release notes – these have 2.6.32
 - SLES 11 SP2 is a 3.0 and so the problem should not exist.
 - Kernel.org will be fine because it is newer
- RoCE status
 - Nate got MPI working well in a back to back fashion. Also Link Init and uDAPL.

- NFSoRDMA and RDMA Interop and RDMA Stress will be the next things they tackle.

Ubuntu update – Samantha - Testing notes are preliminary

- 1) What is in the kernel is based on OFED 1.4.2
- 2) Narayan Desai is one of the developers who converts and maintains Debian packages
 - a) Samantha says there are 28 or 29 packages needed and he only has 8 or 9 ported
 - b) She updated to the ones he provided and they did not solve any issues she was having
- 3) The group is trying to push OFED into 13.04 Debian which must be completed by February and must be stable. If this happens, then Ubuntu can take advantage of this and will be able to support all the current features.

Review of the Ubuntu [test report](#) prepared by Samantha – Section numbers to be updated

9.2.2.1 Ubuntu - how to enable the drivers.

For Ubuntu 12.04 and 12.10 Server editions, you must run the following commands to enable the IB Interface and then assign the IP address in /etc/network/interfaces

- apt-get install ibutils infiniband-diags srptools mpitests
- modprobe mlx4_ib #Mellanox ConnectX cards
- modprobe rdma_ucm
- modprobe ib_umad
- modprobe ib_ipoib

Note: Most of this is about installing necessary packages to get things to work.

- 1) Most of the commands used here and in the following tests require root-level privileges. Either use 'sudo -i' to simulate a Root login shell or prepend 'sudo' to all the commands.
- 2) The OFED version included in packages and modules available in Ubuntu 12.04 and 12.10 is v.1.4.2.

11.2 IB Fabric Initialization using OFED

11.2.2 Step 10

Note: For Ubuntu, ibdiagnet -t <tp> requires local system name specified. Use ibstat to find a match and do 'ibdiagnet -s <sys name> -t <tp>'.

11.3 IB IPoIB Connect Mode (CM) using OFED

All good

11.4 IB IPoIB Datagram Mode (DM) using OFED

All good

11.6 IB SRP using OFED

11.6.2 Step 7

Note: For Ubuntu, lsscsi is not installed by default. Please do 'apt-get install lsscsi' to install it.

13.2 TI NFS over RDMA using OFED

13.2.1 Step 1).

Note: NFSoRDMA is not installed by default in OFED v.1.4.2.

13.3 TI Reliable Datagram Service (RDS) using OFED

13.3.1 Step 2

Note: Package rds-tools 1.4.1-OFED-1.4.2-1 is required to run rds-info on Ubuntu. Also rds-tcp module needs to be loaded - 'modprobe rds-tcp'

13.5 TI RDMA BASIC Interop

Note: Test cannot be run due to the lack of supported packages for Ubuntu

13.6 TI RDMA Stress Test

Note: Test cannot be run due to the lack of supported packages for Ubuntu

Open MPI testing was prevented because the FDR optical cables were removed in order to complete the IBTA Plugfest testing on these cables. Edward says this is now enabled. But he has one link issue to resolve link issues.

Open Issues

Mitko

- 1) Do the Ubuntu systems have Intel HCAs or Mellanox?
 - a) **Edward:** 2 are ConnectX 2 and 2 are ConnectX 3 HCA.
 - b) **Edward** is willing to swap out. He will connect offline with Mitko.
 - c) **Samantha** says she thinks it is good to test with Intel as well but they have limited resources. They do not plan to provide direct support for RDMA at this point.

Edward

- 1) Wants to thank Jim Ryan and Intel for providing 4 new Intel PCIe Gen 3 servers for use in the OFA Cluster
- 2) Samantha had problems with the RDMA Basic Interop test suite and did not know how to resolve.
AR: Edward take a look at this

Action Requests (ARs)

OFA IWG - New test suite development

- 1) Windows: need to develop test suites for Server 8 and include more tests for WinOFED

Chelsio

Intel iWARP – Harry

Rupert Dance

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

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

11/6/2012 – Edward – validate that the following dates have been reserved

- 1) IBTA PF 23 04/01 → 04/12/2013 – [Edward confirms that they have been reserved 11/20/2012](#)
- 2) OFA April Interop Debug Event 04/08 → 04/12/2013

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.