

OFA IWG meeting – 9/10/2013

Attendees

First	Last	Company	07/30/13	08/13/13	08/27/13	09/10/13
Tom	Reu	Chelsio	X			X
Martin	Schlining	DataDirect Networks		X	X	
Jeff	Kopko	Emulex		X	X	X
Pradeep	Satyanarayana	IBM	X	X		X
Harry	Cropper	Intel		X	X	X
Mitko	Haralanov	Intel		X		
Jess	Robel	Intel	X	X	X	X
Yevgeny	Petrilin	Mellanox				
Rupert	Dance	Software Forge	X		X	X
James	English	UNH-IOL	X	X	X	X
Nate	Landolt	UNH-IOL	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	X	X

2013-2014 Events

Interop and Plugfest Events

- **October 2013**
 - [IBTA PF24](#) 10/07 → 10/18/2013
 - [OFA Interop Debug Event](#) 10/14 → 10/18/2013
- **April 2014**
 - IBTA PF25 03/31 → 04/11/2014
 - OFA Interop Debug Event 04/07 → 04/11/2014

OFA Training

- January 13 → 16 2014

OFA EWG Status

- OFED 3.5-1
 - [OFED 3.5-1 RC1](#)
 - OFED 3.5-1 GA – goal is to release on 9/16/2013
- OFED 3.5-2
 - [OFED 3.5-2 Daily Build](#) – updated on 8/27/2013
 - OFED 3.5-2 RC1 – goal is to release on 9/16/2013
 - Glenn is testing SL 6.4
- OFED 3.11 or 3.12
 - RC in time for the October 2013 Interop Debug Event.
 - GA sometime in November

October 2013 Interop Debug Event

Rupert asked all vendors to submit OFA Interop Debug registration as soon as possible because it is important for planning and setting up the cluster.

IBM

- 1) 10/40 GbE Switches – We discussed the possibility of IBM providing the lab with some 10/40 GbE switches and Pradeep will look into it.
- 2) 40 RCA – Pradeep will also look into the status of providing 40 GbE RCA.

Emulex

- 1) They were the only one with 40 GbE adapters – Pradeep does not know whether they will bring 40 GbE adapters but will check into it.
- 2) **AR:** Mellanox can you join in.
- 3) Emulex intends to have all products listed for the October 2013 GA List.

Bob Noseworthy

- 1) If and when DCB testing is ready to be tested, he would like to see this happen. Ethernet Fabric testing should also be pushed forward.
- 2) Does Emulex or IBM want to test a specific Ethernet Fabric?
- 3) We have gaps – should the program address this or do we leave to market.

Chelsio

- 1) He will check for support 3.5-1 and his products and he will check and see if iWARP testing should include different switches

Intel

- 1) He is OK with the testing being done with NE020
- 2) UNH-IOL says that this is the switch configuration they are using: Fujitsu
 - a) 20 CX4 ports
 - b) 4 QSFP port with QSFP → SFP+
 - c) NE020 are CX4 based so we need a hybrid switch
 - d) Could use Fujitsu to connect

Test Plan review

NFSoRDMA

- 1) The comment about IB only applies to the line mentioning OpenSM
- 2) James did some testing on RoCE. They had a snag mounting server shares.
- 3) **AR:** Nate and James will share problems with Jeff and Pradeep.

Rockets Test Procedure

- 1) We reviewed the draft test procedure available here:
 - a) <https://ofa.causewaynow.com/wg/Interop-WG/document/158>
- 2) Synchronous, Asynchronous, blocking and verify calls are possible.
- 3) Capital S All runs through all iterations.
- 4) There are 5 separate tests and they do all size transfers between all clients and servers

- 5) So far he has only got the sockets calls working but will work with Jeff to get others resolved

Jeff

- 1) I will provide new firmware for Emulex adapters which will likely resolve the issues with RSocketes.
- 2) This draft test suite is RSocketes Basic. Once we have that completed, we need to implement a similar test to the RDMA Stress test. We could use MPI over RSocketes for stress.

RoCE over long distance - Bob Noseworthy

- 1) Dr. Russell has pursued this to aid his RDMA research work.
- 2) If interested we may be able to pursue this by possibly linking to other RDMA enabled sites, or possibly "looping-back" to provide a latent link.
- 3) Regarding long distance RoCE, this can be achieved at 10Gbps with an ANUI WAN simulator (uptown 250ms latency, plus controlled, thus repeatable, network impairments) For faster speeds, we have SM fiber spools of varying lengths (up to 25km on a single spool).
- 4) Personally I would still like to see DCB validated to ensure no loss of packets due to congestion. This would/could also draw in DCB switch vendors to the program.
- 5) Similarly, RoCE is moot without an Ethernet fabric solution. Would this be TRILL? SPB? QFabric, FabricPath, etc etc.
- 6) Will the OFA be addressing these challenges? Some arguably are best for the market to decide, but it does impair the efficacy of the logo program (in my opinion at least)
- 7) External Networks such as DYNES is harder but could be used
- 8) ANUI is a 10 GbE device and is modular but will only be one link at 10 GbE. So we could do this between Emulex and IBM. This could be an Alpha test and could be developed for future events. Bob can bring this over for part of the event. Latency is easy –pattern loss can also be generated.
- 9) Jeff wants to know the model of ANUI because they have one in the lab.
 - a) Bob says it is a GEM 3500 model and Bob will get it.

iSER

- 1) Jeff has a team in Roswell that does embedded IO and they are working on that

OpenFramework Working Group (OFWG)

The OFA has formed a new work group whose stated mission is as follows:

- 1) Develop, test, and distribute:
 - a) An extensible, open source framework that provides access to high-performance fabric interfaces and services.
 - b) Extensible, open source interfaces aligned with ULP and application needs for high-performance fabric services.
- 2) OFWG will not create specifications, but will work with standards bodies to create interoperability as needed
- 3) You can join the group here: <http://lists.openfabrics.org/cgi-bin/mailman/listinfo>

Opens

- 1) Pradeep – will we use OFED 3.11
 - a) Rupert: I don't think 3.11 or 3.12 will be ready in time for the Interop Debug Event. So we will use the latest RC of OFED 3.5-2.

Action Requests (ARs)

Emulex – Jeff

9/10/2013

- 1) Provide FW updates to UNH-IOL for NFSoRDMA testing
- 2) Check on status of iSER

8/27/2013

- 1) Review the changes to NFSoRDMA and see how they apply to RoCE
- 2) Are there changes needed for the iSER test plan so it will support RoCE

IBM - Pradeep

9/10/2013

- 1) Check on the availability of 40 GbE RCA and Switches for the Interop Debug Event.

Software Forge - Rupert Dance

9/10/2013

- 1) Contact Mellanox and ask them to participate in RoCE Interop testing – done 9/18/2013
- 2) The comment about IB only applies to the line mentioning OpenSM – fix test plan – done 9/24/2013

8/27/2013

- 1) Contact IBM for 40 GB switch – requested that Pradeep look into this – done 9/10/2013
- 2) Jeff has contacted their vendor and they will supply a 40 GB switch – done 8/30/2013

7/30/2013

- 1) The group needs to make a determination about what version of OFED to use for testing in October.
 - a) OFED 3.5-2 – the group decided that we should focus on this version – 8/27/2013
 - b) OFED 3.5-11- this is a possibility but it is not certain when the first RC will be available.
- 2) We need to discuss what's new and what other features need to be tested. For example what to do with RoCE and what about higher speed RoCE and IB.

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

UNH-IOL

9/10/2013

- 1) Nate and James will work with Jeff and Pradeep to resolve issues they have found with NFSoRDMA over RoCE

8/27/2013

- 1) **Edward:** follow up with Glenn on SLES 11 – done 9/6/2013
- 2) **Edward:** follow up with Dr Russell regarding RoCE over long distance – done 8/30/2013
- 3) **Nate Rubin:** create an RSocket test plan for review at the next meeting IWG meeting on 9/10/2013 - – done 9/10/2013
- 4) **Glenn** needs to test the NFS over RDMA changes on RoCE – in process 9/10/2013

7/30/2013

- 1) Develop RSocket tests for IB and RoCE – [done 9/6/2013](#)
- 2) Execute the full uDAPL test Suite for RoCE and make sure there are no issues