

OFA IWG meeting – 2/12/2013

Attendees

First	Last	Company	01/15/13	01/29/13	02/12/13
Samantha	Jian-Pielak	Canonical	X	X	X
Tom	Reu	Chelsio			X
Martin	Schlining	DataDirect Networks	X	X	X
Alex	Nicholson	Emulex			X
Brad	Benton	IBM			X
Pradeep	Satyanarayana	IBM			X
Harry	Cropper	Intel	X	X	X
Mitko	Haralanov	Intel	X	X	X
Jess	Robel	Intel	X	X	X
Jim	Ryan	Intel			
Guy	Ergas	Mellanox			
Rupert	Dance	Software Forge		X	X
Glenn	Martin	UNH-IOL	X	X	
Edward	Mossman	UNH-IOL	X	X	X
Bob	Noseworthy	UNH-IOL		X	X
Nate	Rubin	UNH-IOL	X	X	X

2013 Events

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 RC5 Released 01/17/2013
- OFED 3.5 RC6 Released 01/29/2013
- OFED 3.5 GA Released 02/14/2013

Brad Benton IBM

- 1) He has been a representative for OFA for some years
- 2) He has been trying internally to get IBM involved in Interop for a number of years.
- 3) They will provide their PowerPC platform and want to Interop with Intel platforms.
- 4) RoCE is newest of the lot.
- 5) They are using the standard Mellanox hardware and they are interested in Interoperability with Emulex as well.

Interop Debug event versus Interop GA

- 1) IBM will target April Interop Debug event
- 2) Time line is OK

Alex Nicolson – Emulex

- 1) Engineering side of the house. In charge RDMA initiative

Logo GA Testing - Edward

- 1) He is concerned about getting the Logo GA testing completed so we can focus on the Interop Debug event – he feels the time is tight.

Brad Benton – notes on the operating system support

- 1) CentOS and SL does not have support for PowerPC. He will need us to us RHEL
- 2) Edward we have all the RHEL Distro on hand.
- 3) Brad notes that Fedora also supports PowerPC

Discussion regarding next OFED - OFED 3.5.1 OFED 3.8 RC6

The EWG is trying to decide whether to make the next OFED release based on the latest Linux kernel or do a point release of OFED 3.5 for example OFED 3.5.1. If we create 3.5.1 we can make it an industry preview and include new drivers and updates such as the Emulex drivers, Mellanox ConnectIB etc. If we go with the 3.8 kernel, it is already frozen and we will not be able to put anything upstream.

User Comments

- 1) Harry – no issues if we go either way.
- 2) Intel – Mitko - no issues if we go either way.
- 3) DDN – Marty - no issues if we go either way.
- 4) IBM – Brad – they are beholden to SW provided by existing silicon providers. Only things might be what occur during debug event.
 - a) AIX will not be in April.
- 5) Emulex - Alex Nicolson
 - a) Will research
- 6) UNH-IOL – Bob
 - a) He wants to enable all participants of the OFILG. Current Logo testing was really delayed and do not want to see that happen again. Might be better to allow debug to update OFED. He supports true heterogeneous RoCE Interop. RoCE is still in Beta.
 - b) Logo List is supposed to be every six months and not wait for the next GA.

RoCE Updates

- 1) Test
 - a) Link Init
 - b) MPI
 - c) uDAPL
- 2) Brad
 - a) What was the hardware configuration?
 - i) This was adapter to adapter and no switch. 10 GbE
 - b) He is not sure April will be able to provide 40 GbE. April will be 10 GbE. Does not have to be DCB but need Global pause.
 - c) IBM will hope to bring 40 GBE for October
 - d) **AR:** do you have contacts for IBM within DCB. Bob will check with Blade Networks as well.
- 3) Bob – questions for RoCE Vendors

- a) What to add in bridges
 - b) What hardware will be provided?
 - c) Need to know speed if you are brining
 - d) They have Fulcrum 10 GbE
 - e) He has some 40 GbE but port-type. Support DCB on one port.
 - f) Bob wants to promote switches. Not everyone makes Converged Ethernet product.
 - g) He wants PFC protocol and others to work as well.
- 4) Alex
- a) He hopes to do 40 GbE
 - b) He will bring switch if possible.

Ballot - OpenFabrics Interoperability Logo Program Changes

- 1) Qualified Pass
 - a) Allow a product on the Logo List as long as all issues and defects are very clearly outlined.
- 2) Arbitration Committee
 - a) Allow additional OFA members to participate in Arbitration process.
- 3) Determine components to test
 - a) Allow more feedback from OFILG members regarding qualified products to be included in Interop Logo testing.
- 4) More Transparency and Openness
 - 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.

Voting – Unanimous approval

- Tom Reu - Chelsio Communications
- Brad Benton - IBM
- Harry Cropper - Intel
- Rupert Dance - Software Forge
- Edward Mossman - University of New Hampshire InterOperability Laboratory
- Nate Rubin - University of New Hampshire InterOperability Laboratory

Ballot: Change to OFILP testing priorities

- 1) Enhance RoCE testing
 - a) Review and enhance the RoCE Test Suite to support our new members: Emulex and IBM
- 2) Testing for OFILG members only
 - a) In order to provide maximum benefits to our OFILG members and to meet the 2013 budget constraints, we will suspend testing for companies that are not members of the OFILG. This includes Microsoft , Oracle (RDS), RedHat, SUSE, Ubuntu
- 3) Suspend Testing for ULPs that are no longer supported by OFED or used by OFILG member companies
 - a) Suspend testing of iSER, RDS, SDP, MVAPICH
- 4) Suspend ULP testing for non-critical protocols or vendor specific products

- a) NFSoRDMA and Vendor SM testing (only test OpenSM)

Voting – Unanimous approval

- Martin Schlining - Data Direct Networks
- Brad Benton - IBM
- Jim Ryan - Intel
 - Comment: I'm being supportive, but other votes from Intel count
- Harry Cropper - Intel
- Rupert Dance - Software Forge
 - Comment: I vote to approve the changes to the program as reviewed and approved by the OFA BOD
- Edward Mossman - University of New Hampshire InterOperability Laboratory
- Nate Rubin - University of New Hampshire InterOperability Laboratory

Brad Benton – IBM – other tests he would like to see implemented

- 1) RDMA CM across heterogeneous – x86 and PowerPC
 - a) PowerPC – 4 weeks
- 2) RoCE devices nature – wants stress test with traffic concurrent with RDMA. Tests to see if drivers handle multiple traffic of different types
 - a) Bob – using ETS has not been implemented. He thinks we could use IOMeter.
 - i) Loading 40 GbE link could be tough. 10 GbE is easy
 - ii) Need to validate more than just the end station. He underrated we can do a lot with switch vendors.
 - b) Marty – no one using RoCE in SRP as far as he knows
- 3) Bob – he is more interested in QoS at this point than Fabric Management - Trill SFB
- 4) Alex – will study

Action Requests (ARs)

Rupert Dance

1/29/2013

- 1) Check with Mellanox regarding the state of their firmware updates and the readiness to do GA testing
- 2) Discuss with the EWG and determine what are next OFED release
- 3) Get Brad Benton of IBM and Emulex to join next meeting – [done 2/12/2013](#)

12/04/2012

- 1) Log SRP bug with RedHat

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

2/14/2013

Bob Noseworthy: do you have contacts for IBM within DCB. Bob will check with Blade Networks as well.

12/04/2012 - Nate

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

10/9/2012

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