

OFA IWG meeting – 10/9/2012

Attendees

First	Last	Company	08/28/12	09/11/12	09/25/12	10/09/12
David	Duffey	Canonical		X	X	
Tom	Reu	Chelsio		X	X	X
Martin	Schlining	DataDirect Networks	X	X	X	
Harry	Cropper	Intel	X	X	X	X
Mitko	Haralanov	Intel	X	X	X	X
Jess	Robel	Intel	X	X	X	X
Jim	Ryan	Intel		X		
Guy	Ergas	Mellanox		X		X
Rupert	Dance	Software Forge	X	X	X	X
Glenn	Martin	UNH-IOL			X	
Edward	Mossman	UNH-IOL	X	X	X	
Bob	Noseworthy	UNH-IOL		X		X
Nate	Rubin	UNH-IOL	X	X	X	X
Bill	Weber	Unisys	X	X	X	X

2012 Events

- [IBTA PF 22](#) 10/15 → 10/26/2012
- [OFA Interop Debug Event](#) 10/22 → 10/26/2012

OFED 3.5-rc2 is available.

The release is available at: <http://www.openfabrics.org/downloads/OFED/ofed-3.5/OFED-3.5-rc2.tgz>

Updated packages:

- infinipath-psm-3.0.1-115.1015_open
- perftest-1.4.0-0.80.gd1763bd
- qperf-0.4.7-0.2.gf3f7001

Supported Platforms and Operating Systems

- CPU architectures:
 - x86_64
 - x86
 - ppc64
 - ia64

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

Testing done on OFED 3.5 RC2

- 1) UNH-IOL – no known problems. SRP issue seems to be resolved
- 2) Software Forge – tested with SL 6.3 and ran the CIWG Test Suite – no issues found
- 3) Intel iWARP: Harry – has no data
- 4) Intel IB: Mitko – no data

Results of OFILG Motion

Motion 1: Allow David Duffy from Canonical (Ubuntu) to attend OFA Interop Debug Event in October 2012

Motion 2: Allow Christoph Lameter to attend OFA Interop Debug Event in October 2012

Motion 1 Discussion: During a regular OFA IWG meeting , the OFA IWG discussed David Duffey's request to attend the October 2012 Interop Debug Event . Support was given by acclamation. This vote is written record of the group's approval

Motion 2 Discussion: Christoph Lameter from Jump Trading is a Kernel developer associated with the Linux Foundation and he has worked closely with David Duffey to support Ubuntu. Chrisoph was also a valuable participant in the OFA User Day event in August 2012. Christoph is requesting that the group grant permission for him to attend even though he is not a member of the OFILG.

Final Note: Both parties have signed the NDA required for OFILG members.

Company	First	Last	Motion 1	Motion 2	Date
Chelsio	Tom	Reu	Yes	Yes	10/03/12
DataDirect Networks	Martin	Schlining	Yes	Yes	10/05/12
Intel	Jess	Calciano	Yes	Yes	10/04/12
Intel	Harry	Cropper	Yes	Yes	10/03/12
Intel	Mitko	Haralanov	Yes	Yes	10/04/12
Intel	Jim	Ryan	Yes	Yes	10/03/12
Mellanox	Guy	Ergas	Yes	Yes	10/02/12
NetApp	Abdel	Sadek			
Software Forge	Rupert	Dance	Yes	Yes	10/02/12
UNH-IOL	Bob	Noseworthy			

Yes Votes	8	8
No Votes	0	0
Abstain Votes	0	0

Review of OFA-IWG Interoperability Test Plan-v1.45-v1

- 1) Change the following RoCE tests to TBD – done 10/9/2012 version 1.45-v2
 - a) RoCE IPoCE
 - b) RoCE InfiniBand Gateway
 - c) RoCE Fibre Channel Gateway
- 2) RoCE Ethernet Gateway → RoCE IB Gateway – done 10/9/2012 version 1.45-v2
- 3) Add Fabric Path to the list of Ethernet Fabric technologies – done 10/9/2012 version 1.45-v2
- 4) **Bob Noseworthy**
 - a) We cannot solve all problems with RoCE right away. iWARP runs on TCP/IP and we expect it to work whereas RoCE uses DCB and Ethernet Fabrics which are still under development.
 - b) FCIA and Ethernet Alliance are running flat topologies and they are not even running Spanning Tree.
 - c) The OFILG could decide to never look at fabric and just run flat as well. But in the long term end users will be using fabric in Data Center and validation will be required.
 - d) QFabric and MLAG are other Ethernet Fabrics available today
 - e) Validation of the Ethernet Fabric could be another variable to verify. It is like changing IB SM. We would go from using a Cisco stack of switches to Brocade set of switches
 - f) DCB provides QoS and can run on flat fabric or managed Ethernet Fabric.
 - g) We need to understand how DCBX Application TLV for RoCE
 - h) Even if we ignore fabric problems we need to set ETS for each topology and ensure traffic from the end station.
- 5) Change section 5.3 from QSFP → SFP – done 10/9/2012 version 1.45-v2
- 6) Rupert asked Tom and Harry to review Section 5.4.6 Vendor Specific Notes and see if the Note is still accurate
 - a) **Note:** Currently there is no interoperability between cxgb3 and nes if peer2peer is enabled. Both nes and cxgb3 have their own proprietary ways of doing "client must send the first pdu". The Chelsio parameter file /sys/module/iw_cxgb3/parameters/peer2peer should be modified on all hosts to contain the appropriate value for each test. For example: the value must be set to '0' for the **iWARP Connectivity test** and set to '1' for the uDAPL test.
- 7) Page 64: fix the cross reference for 13.2.2.2 – done 10/9/2012 version 1.45-v2
- 8) Page 65 remove iWARP since it is not supported – done 10/9/2012 version 1.45-v2
 - 4) [For iWARP hosts with Chelsio RNICs] Set the peer2peer bit to one with the following command:
 - a) For T3s: echo -n 1 > /sys/module/iw_cxgb3/parameters/peer2peer
 - a) For T4s: echo -n 1 > /sys/module/iw_cxgb4/parameters/peer2peer
- 9) **IOL AR:** check the OpenMPI commands for RoCE
 - a) For InfiniBand & RoCE:

```
$MPIHOME/bin/mpirun -mca btl openib,self,sm -mca pml ob1 -mca btl_openib_flags \ 306 -np $NP -hostfile $HOSTFILE $IMBHOME/IMB-MPI1
```
 - a) For iWarp:

```
$MPIHOME/bin/mpirun -mca btl openib,self,sm -mca pml ob1 -mca \ btl_openib_flags 306 -mca btl_openib_receive_queues P,65536,256,192,128 -np \ $NP -hostfile $HOSTFILE $IMBHOME/IMB-MPI1
```
- 10) Edward: it looks like you specified duplicate SRP parameters – done 10/9/2012 version 1.45-v2

```
Load srp module with cmd_sg_entries=255 allow_ext_sg=1 allow_ext_sg=1 {
a) Example: modprobe ib_srp cmd_sg_entries=255 allow_ext_sg=1
allow_ext_sg=1
b) Let it initialize
```

Opens

Bob Noseworthy

- 1) He is seeing increased interest in a larger event. DCB, FCoEthernet, iWARP 40 Gigabit. So can we have larger forum. We may want to break off and have a separate RoCE event. One event versus three disparate events.
- 2) EA and OFA should work more closely – Rupert suggests that perhaps the OFA should try and include Ethernet Fabrics
- 3) Need to implement DCBX Application TLV for RoCE
- 4) Goal is to make RoCE mandatory for April. We added the Beta so we can be ready. This way there will be Mandatory tests and an official Logo.
- 5) He is assuming there is no 40 GbE iWARP but can handle it if notified soon.

Rupert

We will be using the new OFA Causeway collaboration tool which will allow us to store our documents in a central location and use version control. It will also provide a place for meeting minutes and will track all of our email threads. I will be sending out invites to join.

Next OFA IWG meeting: Depends on the number of participants but we may not have a call or if we do it will be as a meeting from the IOL

Action Requests (ARs)

OFA IWG - New test suite development

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

Chelsio

10/9/2012

- 1) Rupert gave Tom a reminder to Chelsio to send in registration – done 10/9/2012

Intel iWARP – Harry

8/28/2012

- 1) Find out if Intel has support for 40 Gb/s that would work in the EA Plugfest.

Rupert Dance

9/25/2012

- 1) Update Interop test Plan with RoCE – done 10/9/2012

8/28/2012

- 1) Rupert follow up with DDN about SRP problems with WinOFED 2.0 and with OFED 3.5 RC1 – in process as of 9/7/2012 – filed [Bug 2378](#) – resolved in OFED 3.5 RC2
- 2) Follow up on status of Master report including all Inter OS testing for May 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

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.