

OFA IWG meeting – 8/28/2012

Attendees

First	Last	Company	07/17/12	07/31/12	08/14/12	08/28/12
Tom	Reu	Chelsio	X	X	X	
Martin	Schlining	DataDirect Networks				X
Michael	Both	GE Intelligent Platforms	X	X		
David	Tetley	GE Intelligent Platforms		X	X	
Pradeep	Satyanarayana	IBM			X	
Harry	Cropper	Intel	X	X	X	X
Mitko	Haralanov	Intel	X	X	X	X
Tatyana	Nikolova	Intel				
Jess	Robel	Intel	X	X	X	X
Jim	Ryan	Intel				
Idan	Kligvasser	Mellanox	X		X	
Guy	Ergas	Mellanox		X	X	
Slava	Pekelis	Mellanox				
Rupert	Dance	Software Forge	X	X	X	X
Glenn	Martin	UNH-IOL	X	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
Bill	Weber	Unisys	X	X		X

2012 Events

- [Zurich Workshop](#) 09/24 → 9/25/2012
- [IBTA PF 22](#) 10/15 → 10/26/2012
- [OFA Interop Debug Event](#) 10/22 → 10/26/2012

OFA EWG Status

- 1) OFED 3.5 RC1 is available
 - a) <http://www.openfabrics.org/downloads/OFED/ofed-3.5/OFED-3.5-rc1.tgz>
- 2) OFED-3.5-rc1 Main Changes from OFED 3.2-rc1
 - a) Based on linux-3.5 kernel
 - b) Updated packages:
 - i) ibacm-1.0.7
 - ii) librdmacm-1.0.16
 - iii) dapl-2.0.36
 - iv) libibumad-1.3.8
 - v) infiniband-diags-1.6.1
 - vi) opensm-3.3.15
 - c) Removed -SDP
 - i) libsdp, sdpnetstat
 - d) Linux Operating Systems:
 - i) RedHat EL6.2 2.6.32-220.el6
 - ii) RedHat EL6.3 2.6.32-279.el6
 - iii) SLES11 SP2 3.0.13-0.27-default
 - iv) kernel.org 3.5 *
- 3) MPI to be tested with OFED 3.5
 - a) MVAPICH2 1.8

- b) Open MPI version 1.6.1
- c) Intel MPI: the released version is 4.0.3
- OFED 3.5 needs to have a technology-preview branch

Membership updates

Bob Noseworthy

- 1) Canonical (Ubuntu) is looking into involvement with RDMA. They visited UNH-IOL last week. They were representing the Open Compute Network.
- 2) They would like to see the OFA software be fully supported in Ubuntu but this may not happen in time for the October 2012 Interop event.

Discussion regarding the support of RoCE

- 1) Rupert discussed RoCE with the OFA Board of Directors and asked for help with getting RoCE vendors more engaged.
- 2) Bob commented that we need more support from Channel Adapter manufacturers before we can go to the DCB switch vendors. Emulex and QLogic and both good ones to ask for support
- 3) Bob suggests that we get involved with the Ethernet Alliance (EA) and see if we can have a joint event where we test RoCE. There is a liaison with EA and there will be an event right around the time we do Interop. April, 8th is the first opportunity
 - a) **First Step:** EA OFA event in April with RoCE – no membership requirement. Will OFA require membership
 - b) **Second Step:** Convince Switch vendors of Logo Value.
- 4) Harry from Intel wants to know more about the EA and this proposed event.
 - a) There is a [Terra Fabric](#) event this fall and another event in April 2012.
 - b) Bob explains that they want RDMA engine available on 40 Gb/s – this would be very beneficial to make this work with EA. He asks Harry to let us know if Intel has 40 Gb/s support.

Interop Debug Event for October 2012

Operating Systems

- 1) SL 6.3 - mandatory
- 2) RHEL 6.3
- 3) CentOS 6.3
- 4) SLES11 SP2

OFED

- 1) We will use the latest RC available from the OFED 3.X series
- 2) Right now OFED 3.5 RC2 is pending release

Vendor responses to these choices

Intel – no issues with these OS

Marty – No issues

Jess will follow up.

Interop test Plan review

- 1) Ethernet Fabric tests are not implemented. They were written with Spanning Tree in mind and would need to be updated and generalized. For example we could generalize to say if using some Ethernet fabric it could be tested
 - a) SPB
 - b) Trill
- 2) Bob feels we should remove the following
 - a) Ethernet Fabric Initialize and others but leave Init. Remove re-convergence and Failover.
 - b) Bob feel not appropriate to do spanning tree
 - c) Harry does not care about these tests – at least Spanning Tree
- 3) UNH-IOL currently only has one switch available for iWARP testing and that is the Fujitsu which has both SFP+ and MicroGiGaCN
- 4) Remove iWARP Connectivity because this testy suite relies on of Java scripts which depend on an old utility in OFED which is no longer supported and does not work with OFED 3.5
- 5) iSER
 - a) Bob says he has a company has an iSER target so leave it in the test plan. The test may need updating.
- 6) NFS-RDMA
 - a) Do we need to do Connecathon? Can we pare this test down so it is not so time consuming?
 - b) We need to valid base functionality.
- 7) SDP – remove this test suite because it is no longer supported in the Distros or OFED because of IP and Licensing issues.
- 8) uDAPL
 - a) 3.5 does not have any known problems but 1.5.4.1 did. We think it was daplttest.
- 9) RDMA Stress Test
 - a) This requires a switch fabric and until we get the DCB community involved we cannot test it. Make sense for IB but not for iWARP right now.
 - b) Will be important for RoCE when DCB handles BW
- 10) Remove Intel MPI – this is because it is not Open Source and it has not been updated in some time.
- 11) Open MPI
 - a) We need to settle on which version to use. We need some vendor feedback. Harry says that they use 1.6 but Rupert points out that the current version is 1.6.1 and he would rather see us use the latest GA version. Harry will review with his engineers.
 - b) Open MPI 1.6.1
- 12) MVAPICH – this is deprecated and no longer mandatory.
 - a) Goal is to go through all mandatory and as much as Beta as possible.
 - b) We will keep in the document.

Event Requirements

- 1) WinOFED
- 2) We will not test for October because Microsoft is not ready with server 2012. UNH-IOL will not be preparing environment for October.
- 3) Marty
 - a) Cannot keep WinOFED running for more than 24 hours. No changes since WinOFED 2.0 for SRP.

Inter versus Intra OS testing

- 1) Bob would prefer testing a larger group of homogeneous testing within a given Operating system. For example we would have an entire cluster with RHEL, then SLES and then CentOS.
- 2) This is good for vendors to test all their hardware in each OS.
- 3) Passed by acclamation. No longer test between Operating systems but test all devices within each OS.
- 4) Cannot guarantee full coverage especially for October 2012

Summary for Test Plan

- 1) Test to eliminate
 - a) Link Init
 - b) OpenSM
 - c) NFSoRDMA
 - d) RDMA Stress

- 2) Extra OS Testing – this become the first tests that are run.
 - a) Fabric Init
 - b) Open MPI
 - c) IPoIB

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

Intel iWARP – Harry

8/28/2012

- 1) Find out if Intel has support for 40 Gb/s that would work in the EA Plugfest.
- 2) Does Intel have any problems using Open MPI 1.6.1 – current GA

Mellanox - Guy

7/31/2012

- Review with Mellanox management and see if they are able to support RoCE testing for October 2012.
 - [Response from Mellanox on 8/14/2012](#): At this time point we have the feeling that it will be very hard to perform RoCE testing without interfering with FDR program that we have in focus for the upcoming event.

Rupert Dance

8/28/2012

- 1) Ask the OFA Marketing group to announce the release of the new OFA Interoperability Logo List – [done 9/11/2012](#)
- 2) 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](#)
- 3) Follow up on status of Master report including all Inter OS testing for May 2012

7/31/2012

- 1) Update Interop Test plan for the Mandatory Tests – [done 8/14/2012 and 8/20/2012](#)

Older

- 1) Update Logo Program document 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

8/14/2012

- 1) Update website page with invitation for October Interop event – [done 8/20/2012](#)

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.