

OFA IWG meeting – 7/31/2012

Attendees

First	Last	Company	05/29/12	07/03/12	07/17/12	07/31/12
Tom	Reu	Chelsio	X	X	X	X
Martin	Schlining	DataDirect Networks	X	X		
Michael	Both	GE Intelligent Platforms			X	X
David	Tetley	GE Intelligent Platforms				X
Harry	Cropper	Intel		X	X	X
Mitko	Haralanov	Intel		X	X	X
Tatyana	Nikolova	Intel	X			
Jess	Robel	Intel	X	X	X	X
Jim	Ryan	Intel	X	X		
Idan	Kligvasser	Mellanox	X		X	
Guy	Ergas	Mellanox				X
Slava	Pekelis	Mellanox	X	X		
Rupert	Dance	Software Forge	X	X	X	X
Glenn	Martin	UNH-IOL		X	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	X

Introduction to Guy Ergas from Mellanox

- 1) Guy is taking over the responsibilities for the Mellanox in the OFILG. Idan will no longer be participating.
- 2) He is new in IB but has been performing a similar role in Ethernet for some time.

2012 Events

- [OFA User Day](#) 08/06 → 08/07/2012
- [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
 - 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

OFA Interop GA Event Status Updates

- 1) Arbitration resolution status
 - a) There were two issues and one of them has been resolved. The second issue involves a problem with uDAPL and we are working with the new release in OFED 3.5 to see if the problem is resolved. We suspect it is a software issue with uDAPL.
- 2) Logo List is ready to be published. He will send out link to the page for review. The page won't include any products that are under arbitration
- 3) Preview will include all products.
- 4) Master report – lots of effort to determine how participants get benefits.

Test Plan Updates for October 2012 Debug Interop event

TI RDS	Not supported by iWARP
TI SDP	Not supported by iWARP

TI MPI - Intel	Deprecated
TI MPI - Open MPI	Mandatory
TI MPI - OSU - MVAPICH 2	Deprecated

- 1) **Harry** reminds us that we agreed to deprecate both the Intel MPI and the OSA MPI testing because we are no longer performing these tests as part of the OFA Logo Validation.
 - a) Rupert committed to having the test plan updated for the next meeting.
- 2) **David Tetley** (GE) Why no MVAPICH
 - a) Rupert: we feel that the Open MPI test suite covers the needed MPI testing and there is a lot of overlap. We have added Distro testing of RHEL, SUSE, Scientific Linux and Windows so we need to consolidate the effort where we can without sacrificing test suites which are important to the OFILG members.
 - b) Intel MPI is not being executed but Intel tests are being run in Open MPI
- 3) Michael Both
 - a) He would like us to add RoCE testing to the OFILP program.
 - b) He is working to get GE to become a member.
 - c) Customers are the end users. They are running a full set of tests.
- 4) Guy: he has just started looking into RoCE
 - a) **AR:** Review with Mellanox management and see if they are able to support RoCE testing for October 2012.

GE Status

- 1) Small cluster of 5 nodes connected to 10 GbE via RoCE using Mellanox HCA
- 2) Seeing a test lock up using the Intel MPI doing a gather. AXIS software did a gather meaning all nodes send to target node at one time. Data would stop progressing and seeing packet drop and not forward progress.

- 3) They were surprised that this would happen on standard 10 GbE switch. They have been told that switches need support for DCB to relieve congestions. All nodes backing off at same time. Just received switch SX1036 with DCB support and they will try same tests.
- 4) Their switch can support DCB by enabling in SW

Bob Noseworthy

There are a number of problems we need to address in order to support RoCE Interop testing

- 1) RDMA over Converged Ethernet – compliance testing. Ideally the IBTA would add compliance testing for RoCE just as they do IB transport compliance testing in the IBTA.
- 2) Interop of RoCE – we need more end point vendors and switch vendors to get involved in order to really test interoperability.
- 3) DCB is effectively the QOS of switch. It is what ensures a Lossless Fabric when the transport is Ethernet.
- 4) Fabric interconnect – there is more to supporting RoCE than just DCB. You have to consider what fabric interconnect is being used. There are many competing standards out there and the industry has not settled on one standard. Some of the options are listed below and Spanning Tree is not really a viable option for this technology. So if we add RoCE Interoperability testing we would need to decide which fabric interconnects we would support. See <http://www.enterprisenetworkingplanet.com/datacenter/will-trill-or-shortest-path-bridging-win-out.html> for some details
 - a) Spanning tree
 - b) Trill
 - c) Shortest Path Bridging

SDP support in OFED 3.X

- 1) Bob thinks the decision is based on the end users and the membership of OFILG. UNH-IOL would also like to get their EXS protocol into the OFED but not necessarily the OFILP program.
- 2) Mitko – not sure what the long term support will be. They issue SDP that comes with OFED. He will try for more feedback for next meeting.
- 3) Jess – she will ask around as well.
- 4) Guy – he will check.

Action Requests (ARs)

OFA IWG - New test suite development

- 1) RoCE test suite.
- 2) Windows
 - a) Need to develop test suites for Server 8
 - b) Include more tests for WinOFED

UNH-IOL

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.

Rupert Dance

7/31/2012

- 1) Update Interop Test plan for the Mandatory Tests

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

Guy

7/31/2012

- Review with Mellanox management and see if they are able to support RoCE testing for October 2012.