

OFA IWG meeting – 7/17/2012

Attendees

First	Last	Company	05/15/12	05/29/12	07/03/12	07/17/12
Tom	Reu	Chelsio		X	X	X
Martin	Schlining	DataDirect Networks	X	X	X	
Michael	Both	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	X	
Idan	Kligvasser	Mellanox		X		X
Slava	Pekelis	Mellanox		X	X	
Rupert	Dance	Software Forge	X	X	X	X
Glenn	Martin	UNH-IOL			X	X
Edward	Mossman	UNH-IOL			X	X
Bob	Noseworthy	UNH-IOL	X		X	X
Nate	Rubin	UNH-IOL			X	X
Bill	Weber	Unisys	X	X		X

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

- OFED 3.2
 - There will be no 3.2 GA Release. The first GA release will be based on the 3.5 Kernel
 - OFED 3.2 RC1 has been released but does not support RHEL 6.3 but OFED 1.5.4.1 does.
- OFED 3.5 work is expected to start the week of July 1st
- MPI to be tested with OFED 3.5
 - MVAPICH2 1.8
 - Open MPI version 1.6.1
 - Intel MPI: the released version is 4.0.3
- OFED 3.5 needs to have a technology-preview branch

Introduction to Michael Both from GE Intelligent Platforms

- 1) They have developed two commercial applications called [AXISview](#) and [AXISFlow](#)
- 2) They provide an API like MPI but are designed to appease customers who want something lighter weight.
- 3) The applications are portable across OS
- 4) AXISFlow shows data running live.
 - a) Portable with many interconnects that can work together. They have released IB and RoCE implementations.
 - b) It is currently being used in both military and medical applications.
 - c) It can handle large number of nodes and cores per processor. Testing currently deals with testing up to 5 node and 4-8 cores per processor.
- 5) **InfiniBand** showed no signs of issues and everything worked with **sockets** as well using 1 GbE or 10 GbE.

- 6) They have also used the application on **RoCE** – same verbs as IB and yet the test fails oddly. They tried different MTU sizes but it made no difference. Nodes and switch are all consistent.
 - a) They are trying to characterize and understand the problem. They also noted that Open MPI seems to have the same issues.
 - b) Can run it in any system with RoCE on Ethernet.
 - c) Not sure it is hardware or software.
 - d) They are using a commercial switch and Mellanox chips.
- 7) They use RoCE because they do not have custom IB switches. They have custom Ethernet 1 GbE switches. They hope customers will easily be able to switch.
- 8) **Bob Noseworthy**:
 - a) They need to look at RoCE with DCB vendors involved. Any end uses should be using DCB vendors' products.
 - b) Michael has a couple of commercial switches. These are 10 G switches. Both they also test with 1 GE and 10 GbE on sockets.
 - c) **Bob**: there is currently no certification for DCB – Metro is the closed. UNH-IOL has [Data Center Bridging](#) (DCB) effort. OFA should pursue and engage the rest of the community. A DCB switch is critical for RoCE just as a compliant IB switch is critical for proper operation of InfiniBand.
 - i) He is not suggesting compliance as a pre-requisite. He will get list of RoCE switch vendors. He notes that even our iWARP testing uses a Fujitsu switch donated to the group.
 - ii) We need help from the OFA community to get DCB vendors involved.
- 9) **Mike**: Using SL 6.1. He will give all the hardware details including hardware and Open MPI.
- 10) **Idan**: he will test once we get details.

OFA Interop Event Status Updates

- 1) Edward and his team are currently working on Inter OS test reports. This is the scenario in which a cluster is built with a mixture of different OS installations and all the functionality is tested.
- 2) All Logo Reports were sent out by UNH-IOL on July 3rd and therefore technically the arbitration submission deadline should be July 17th. Rupert has extended the arbitration claim deadline to the end of the week which is July 20th 2012
- 3) **Bob Noseworthy** – there is at least one official arbitration claim issue for certain.
 - a) Note there are no Gen 3 Platforms.
 - b) Latency and performance can be affected.
- 4) **Bob**: wants a logo list in two weeks.
 - a) What about the master list for all the Inter OS testing? He wants to publish publically.
 - i) **Rupert** suggested that we ask the OFA XWG to review before we make a final decision.
- 5) **Harry** - when is next event
 - a) Rupert: The OFED 3.X Interop debug event schedule for July has been cancelled because OFED was not ready The next Interop Debug event will be on 10/22 → 10/26/2012 and it will most likely be based on the 3.5 kernel.
- 6) Mandatory Tests
 - a) Bob reminded the group that we need OFA IWG members to participate in the development of RoCE test plans and to extend the Windows testing

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

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

Slava

3/20/2012

Check with Mellanox about participating in RoCE testing