

OFA IWG meeting – 3/6/2012

Attendees

First	Last	Company	01/24/12	02/07/12	02/21/12	03/06/12
Chris	Calandro	Chelsio	X	X		X
Tom	Reu	Chelsio				
Martin	Schlining	DataDirect Networks			X	X
Pradeep	Satyanarayana	IBM	X	X		
Harry	Cropper	Intel	X	X	X	X
Jim	Ryan	Intel	X	X	X	X
Slava	Pekelis	Mellanox		X		
Mitko	Haralanov	Intel	X	X	X	X
Jess	Robel	Intel	X	X	X	X
Rupert	Dance	Software Forge	X	X	X	X
Bill	Boas	System Fabric Works		X		
Chris	Hutchins	UNH-IOL	X	X	X	
Edward	Mossman	UNH-IOL	X			
Bob	Noseworthy	UNH-IOL		X	X	X
Nick	Wood	UNH-IOL	X	X	X	X
Bill	Weber	Unisys	X	X		X

2012 Events

- [IBTA PF 21](#) 04/16→04/20/2012
- [OFA Interop Debug Event](#) 04/23→04/27/2012
- IBTA PF 22 10/15→10/19/2012
- OFA Interop Debug Event 10/22→10/26/2012

OpenFabrics User and Developer Workshop

The three day workshop features leading technical speakers from AMD, Chelsio, Cray, EMC, Fusion I/O, HP, Intel, Lawrence Livermore National Laboratory (LLNL), Mellanox, Microsoft, Myricom, NVIDIA, NASA, Ohio State University, Oak Ridge National Laboratory (ORNL), QLogic, Sandia, Software Forge, OCZ and VMware and includes sessions that cover OFED and RDMA updates and proposed protocols for new OFED and RDMA capabilities for:

- Windows Server 8
- Exascale, including HPC storage file systems
- Bus-based Flash and SSDs and virtualization
- Big data and Hadoop
- Heterogeneous parallel computing
- High-frequency trading and market data

The workshop is open to anyone with the desire to provide collaborative input into the direction and mechanics of OFED; this includes developers, technologists, supporters, end users and business professionals. Non-OFA members welcome.

- 1) We are trying to include speakers from the Distros speakers and get them to discuss where they are going in the future
- 2) Video will be available – should be good quality.

OFED Status

- OFED 1.5.4: GA - released December 5th
- OFED 15.4.1 GA - released Feb 2nd 2012
 - We will most likely use OFED 1.5.4.1 GA for the April Interop event
- OFED 3.2 - Work is underway and will be in synch with Kernel version 3.2. We discussed whether we should consider using an RC for the Interop debug event.
 - Harry - Early RC could be used – Harry will key it up for validation
 - Mitko – no plans yet – source repository not finalized.
 - Bob Noseworthy – quite willing to try
 - Marty – wants to try 3.2 at next Debug – FDR

December-January 2012 - OFA Interop GA Event - Arbitration Status

- There have been a number of arbitration requests submitted and they are all being reviewed by the committee. We are also actively working with the vendors involved to resolve the technical issues.
- All arbitration requests will be resolved by Wednesday March 7 and the Logo Reports will be sent to vendors the following week.

Review of UNH-IOL Suggested updates to the Test Plan.

Bob - goals

- 1) Broader range of relevant testing
- 2) Want to support the Distros to become members
- 3) Need to test the Distros on all the hardware available
- 4) Want to sacrifice where we can.
- 5) One day for each OS for testing – currently this would be 4 days and then one day for mixed OS testing.

Nick – detailed explanation

- 1) Currently we test devices “all to all” and as we expand to include multiple Distros, it takes too long. So we are proposing that we change some of the tests so fewer iterations are run.
- 2) **IPoIB** – this is now all to all and for every SM
 - a) Testing in the past has revealed Boundary issues at 8 KB and other sizes have not caused problems
 - b) **Options four:** 2 below 8KB boundary and 2 above
 - c) **Mitko:** he would rather see more than less – problems can hide everywhere. Four is too little – go with 6 or 8. The middle ranges sometimes cause problems. Chris is supportive and so is Harry.
 - d) **Jess:** feels that 6 maybe the right number
 - e) **Bob** is interested in the scale test as well. Part of scale back is dependent on how many Distros we include. If RHEL joins and SUSE did not then what do we do. We need to make choice that is beneficial to the members.
 - f) **Chris** if we get one it might get the others to join.
 - g) **Rupert:** suggest that we should stick with technical and let OFA decide on membership issues and funding.
- 3) **MPI**
 - a) We currently test MPIs included in OFED – this is Open MPI and MVPICH. OFED 3.2 will be based on the kernel and these will not be included
 - b) Open MPI is the only MPI that covers all platforms such as IB, iWARP. Windows will also be available in future versions. So UNH-IOL has suggested dropping MVAPICH
 - c) **Harry:** leave in test plan but not make it mandatory.
 - d) **Chris:** Chelsio has large customer that uses MVAPICH.
 - e) **Bob:** MVAPICH testing can always be scheduled by members who need to validate support.

- 4) **SM**
 - a) **Open SM** would be the primary SM used to run all of the test suites
 - b) **Vendor SM**: A subset of the tests would be run using each of the vendor SMs submitted. For example IPoIB, MPI and RDMA tests. UNH-IOL philosophy is that historically if SM can properly configure the fabric, then it has no impact on ULP
 - c) **Mitko**: he does not think this will cause problems. **Jess** will look at this later.
- 5) **SRP**
 - a) Program is designed to test ULP and not exhaustive testing of hardware. Therefore the suggestion is that each target only be configured with one volume which will reduce the complexity of testing. Additional volumes and configurations can be scheduled after the event.
 - b) **Marty**: feels that each target should have two volumes. He has had problems that the OS could not read more than one LUN. Inquiries, reads and writes must all be validated.
- 6) **SDP**
 - a) RHEL and SUSE no longer support SDP primarily because of licensing issues. So we suggest changing the Logo Policy to indicate that a ULP will only be tested if supported by the OS. So testing on RHEL and SUSE would be skipped.
 - b) **Harry** is delighted – less is more
- 7) **Events**
 - a) Debug – OS installed and vendors and distributors can apply non-repository based patches. Any fixes would have to be in the repository for the GA.
 - b) GA Event – only the OS and repository can be used. No outside patches.
- 8) Licensing Issues
 - a) We cannot boot MS cluster because licenses have expired.
 - b) There is also a concern with RHEL because valid licenses must be available for each installation in order to access the repositories.

Questions

- 1) Chris – what do you mean by repository?
 - a) We are referring to the Distros libraries which contain all of the additional software available that is not contained in the ISO. RHEL uses YUM and SUSE uses YAST to access these additional packages.
 - b) Chris is supportive
- 1) Mitko – comment of the suggestion that any patches must be in repository before the Logo Event
 - a) This means we will be at the mercy of the Distro schedule
 - b) **Bob**: this is not policy yet and was predicated on whether RHEL and SUSE are members.
 - c) If we adopt this method, OFILG has to dictate that RHEL and SUSE agree to these things.
 - d) **Bob**: what he wants to achieve is getting RHEL, SUSE and MS included.

Action Requests (ARs)

Group 2/21/2012

- 1) Review the Logo Program and prepare updates for the next meeting.
- 2) Review the Test Plan and prepare updates for the next meeting.

Nick

10/11/2011

- 1) Need fail and pass criteria in Link Init procedure
 - a) LED
 - b) Link Width
 - c) Link Speed
 - d) No crashes
 - e) Etc

10/04/2011

- 1) Jess requests that the Link Init and other tests have the pass/fail criteria spelled out in a table – for example what happens if LED only comes up on one device.
- 2) Need to write up an explanation of the table and where the devices must be applied

Rupert Dance

7/5/2011

- 1) Update Logo Program document to include OFA Policy regarding OFILG Membership and granting of the Logo

3/1/2011

- 1) Check with EWG and XWG about the distribution of OFED Binaries – [in process as of 5/2/2011](#)

2/15/2011

- 1) Create Logo with version or change Logo Guidelines

8/31/2010

- 1) Incorporate ICR into Logo Program