

OFA IWG meeting – 3/26/2013

Attendees

First	Last	Company	02/12/13	02/26/13	03/12/13	03/26/13
Samantha	Jian-Pielak	Canonical	X		X	
Tom	Reu	Chelsio	X	X	X	
Martin	Schlining	DataDirect Networks	X	X	X	
Jeff	Kopko	Emulex				X
Alex	Nicholson	Emulex	X	X	X	
Brad	Benton	IBM	X	X	X	X
Pradeep	Satyanarayana	IBM	X	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		
Yevgeny	Petrilin	Mellanox				
Rupert	Dance	Software Forge	X	X	X	X
Glenn	Martin	UNH-IOL			X	
Edward	Mossman	UNH-IOL	X	X	X	X
Bob	Noseworthy	UNH-IOL	X			
Nate	Rubin	UNH-IOL	X	X		X

2013 Events

April 2013 dates

- [IBTA PF23](#) 04/01 → 04/12/2013
- [OFA April Interop Debug Event](#) 04/08 → 04/12/2013

October 2013

- IBTA PF24 10/07 → 10/18/2013
- OFA Interop Debug Event 10/14 → 10/18/2013

OFA Conferences – Monterey California

- [OFA User Conference](#) 04/18 → 04/19/2013
- [OFA Developer Workshop](#) 04/21 → 04/23/2013
- **Early Bird Deadline** **Extended:** 03/29/2013

OFA Developer Workshop Sponsorship – Jim Ryan

- 1) There are three levels for OFA Members and a special for the IBTA cable vendors
- 2) The deadline is April 5th

Bronze \$1000

- A sponsor logo on all workshop promotion, materials and OFA website

Silver \$1500

- A sponsor logo on all workshop promotion, materials and OFA website
- Ability to have branded giveaways given to attendees at check-in

Gold \$3500

- A table in the Marriott lobby

- Sponsor logo on all workshop promotion, materials and OFA website
- Permission to display products
- Permission to distribute documentation and giveaways.
- For sponsors who already have at least one registrant for the workshop, there's a \$250 discount

IBTA cables \$1750

- A table in an exhibit room with the InfiniBand cluster that will be on display
- Sponsor logo on all workshop promotion, materials and OFA website
- Permission to display products
- Permission to distribute documentation and giveaways.
- One registration pass for the workshop

OFA EWG Status

- OFED 3.5 GA Released 02/14/2013
- OFED 3.5.1 Planned Support for RHEL 6.4
- The EWG is proposing to do one or more Technology Previews based on OFED 3.5.1.
 - Intel Xeon Phi
 - Inclusion of Emulex and IBM drivers
- CentOS Current Distribution 6.4
- RHEL Current Distribution 6.4
- Scientific Linux Current Distribution 6.3

UNH-IOL - OS

- 1) **Edward** wants to stick with Scientific Linux 6.3 for Interop Debug Event
- 2) SL is on RC2 and may be able to release SL 6.4 in time for GA
- 3) We considered using SL 6.3 for the Interop Debug and SL 6.4 for GA but the OFA Logo Policy states that we must use the same version for both events. This is in order to give vendors sufficient time to test at the Debug event and be sure they are ready for the Interop GA event.
- 4) **Jeff** wants to know the cutoff date.
 - a) They have been working on SL 6.3 but would consider 6.4 for Logo GA
- 5) **Mitko**:
 - a) He does not feel it makes sense to use 6.3 for Interop for Debug and 6.4 for GA Event. For example what happens if we get a failure in 6.4
- 6) Conclusion
 - a) We will stick with SL 6.3 for both the OFA Interop Debug and GA Events.

OFED

OFED 3.5.1 is essentially the same as OFED 3.5 but with support for RHEL 6.4

- 1) **Mitko**: he would be willing to try using OFED 3.5.1. We could target 3.5.1 with 3.5 as a backup.
- 2) **Pradeep**
 - a) Why not use OFED 3.5 since IBM will not use RHEL 6.4 and we are going to use SL 6.3 so what the point of using OFED 3.5.1
- 3) Brad
 - a) They are not using RDMA built into RHEL but rather the OFED 3.5 distribution
- 4) Jeff
 - a) He supports OFED 3.5
- 5) **Conclusion**
 - a) OFED 3.5

March Interop GA Event Status

- 1) Edward and UNH-IOL reporting on IB
 - a) 90% done with Mandatory testing
 - b) NFSoRDMA is last test to accomplish
 - c) They were considering doing Inter OS testing between RHEL and SUSE as we have done in previous event. Rupert suggest that we should eliminate it since we will not be doing it going forward and the Interop Debug event is so close.
- 2) RoCE
 - a) All good results so far
- 3) iWARP
 - a) There have a few more tests to run but they have not seen any changes since the Debug event.
- 4) Reports
 - a) They will try to notify vendors of any failures before the start Interop Debug.
- 5) Hardware and Firmware is good so far
 - a) What is the status of the IBM RCA?
 - i) **Pradeep**: Available early next week
 - b) **Jeff**: they don't have the vendor switch they want to bring to the event but expect it on 3/27-3/28/2013

April Interop Debug Event – questions from Jeff

- 1) OFED stack 3.5 or 3.5.1? – 3.5
 - a) We decided on 3.5
- 2) Scientific Linux 6.3 or 6.4? – 6.3
 - a) We decided on 6.3
- 3) Date for next Interop GA Event
- 4) The IL Policy indicates that this should occur within 3-4 weeks of the conclusion of the Interop Debug event. It is sometimes delayed if we find Interoperability issues that are hard to resolve or if there is a delay in the release of the OFED GA as we experienced with OFED 3.5
- 5) Any expected OFED stack or Linux OS changes expected
 - a) No
- 6) OFA Test Plan changes
 - a) Any changes expected for the RPING, DAPL and MPI command lines?
 - i) **Nate**: There have been no test plan updates but there may be changes once we have a RoCE fabric. So far they have only been able to do parts 1-2 of uDAPL because they were only doing peer-peer testing.
 - ii) Emulex is asking because they want to be prepared and run all tests in advance.
 - b) Min/Max number of QP's?
 - i) **Nate**: No changes
 - c) Anything new like perfest or any other benchmark testing?
 - i) Rupert: no changes to my knowledge. Also as a matter of policy both UNH-IOL and the OFILG try to stay away from performance testing.
- 7) Beta Test Plans – we intend to develop these tests during the April 2013 Interop Debug Event.
 - a) RoCE Fabric Init

- b) RoCE IPoCE
 - c) RoCE Ethernet Gateway
 - d) RoCE Fibre Channel Gateway
- 8) Enhancements – these were requests from IBM and we will entertain these as time permits.
- a) RDMA_CM.
 - b) RSocketS over RoCE
 - c) IPv6
 - d) Bonding over RoCE Ethernet interfaces
- 9) RoCE switch vendors
- a) **Rupert:** to my knowledge there are no additional switch vendors for this event.
- 10) Fabric configurations
- a) Size, vendors etc
 - b) We will have an expanded RoCE fabric because we will have both IBM and Emulex. In addition Emulex hopes to bring a pre-production switch from their partner Arista.

RSocketS – Sean Hefty

- 1) RSocketS is a user space implementation, so there's nothing that needs to be pushed upstream. An initial release of RSocketS was included in OFED 3.5 as part of the librdmacm 1.0.16 release. The 1.0.17 release extended RSocketS capabilities with more features.
- 2) There should be an rsocket.7 man page installed as part of the librdmacm, along with a sample application (rstream) and a socket preload library (librspreload.so). The preload library is installed in %lib%/rdma by default to avoid potential compatibility issues.
- 3) RSocketS should run over InfiniBand or RoCE. **iWARP support is not yet available.**
 - a) Pradeep says that iWARP does not support write immediate and that is why it is not supported.
- 4) Rupert created a first draft of a test plan which was based on the old SDP test. This will need refinement after more experimentation is done with rstream and the RSocketS protocol. The draft test is available here:
 - a) Test: <https://ofa.causewaynow.com/wg/OFLIG/documentRevision/download/136>
 - b) Man Pages: <https://ofa.causewaynow.com/wg/OFLIG/documentRevision/download/134>

Interop Test Spec Updates – Version 1.47

- 1) Most of the IBM suggestions have been incorporated into the latest version of the test spec which is 1.47 <https://ofa.causewaynow.com/wg/OFLIG/documentRevision/download/135>
- 2) Other changes
 - a) Updated the requirements for the OFILP for April 2013
 - b) RoCE Updates suggested by IBM
 - c) Updated Section 9 regarding the OS and OFED
 - d) Added RoCE Stress Test to Table 16 and updated Section 13.6 RDMA Stress test
 - e) Added RSocketS to the list of tests for future adoption
 - f) Added IPv6 to the list of tests for future adoption
 - g) Added Bonding for RoCE Ethernet interfaces to the list of tests for future adoption
- 3) All changes are noted by the Blue Bar on the right of the document.

Action Requests (ARs)

Rupert Dance

2/26/2013

- 1) Secure 50-60 new cables to get the OFA Cluster up and ready for GA testing – done 2/28/2013
- 2) OpenFabrics Interoperability Logo Program Changes – in process 3/24/2013
- 3) Contact Sean Hefty about Rsockets and SDP – done 3/25/2013

12/04/2012

- 1) Log SRP bug with RedHat

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

3/26/2013

Edward: find out from Bob the status of the DCB switch he was going to try and loan to the event

2/14/2013

Bob Noseworthy: do you have contacts for IBM within DCB. Bob will check with Blade Networks as well.

10/9/2012

- 1) **Nate:** verify that rping works for RoCE as well as iWARP
- 2) IOL Check the OpenMPI commands for RoCE