

OFA IWG meeting – 9/11/2012

Attendees

First	Last	Company	07/31/12	08/14/12	08/28/12	09/11/12
David	Duffey	Canonical				X
Tom	Reu	Chelsio	X	X		X
Martin	Schlining	DataDirect Networks			X	X
Michael	Both	GE Intelligent Platforms	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
Jess	Robel	Intel	X	X	X	X
Jim	Ryan	Intel				X
Idan	Kligvasser	Mellanox		X		
Guy	Ergas	Mellanox	X	X		X
Rupert	Dance	Software Forge	X	X	X	X
Glenn	Martin	UNH-IOL	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

- [IBTA PF 22](#) 10/15 → 10/26/2012
- [OFA Interop Debug Event](#) 10/22 → 10/26/2012

David Duffey - Introduction

- 1) David is the directory of technical partnership at Canonical which produces Ubuntu
- 2) He works with component vendors and enterprise server side for example Emulex
- 3) He supports Ubuntu cloud and the work they do with Open Stack and is actively engaged with HP Cloud, T-Mobile, British Telecom and others.
- 4) One of their goals is to get more paying customers at heart of data center. Currently they have some support of IB but it is not directly support yet due to funding issues.
- 5) They want to make a case to join OFILG so they can officially support IB and as a company they are at a tipping point. They are not quite ready because they are focusing on commodity scale out.
- 6) Please get feedback to him if we know of folks using Ubuntu and Open Stack . Rupert has provided his contact information above and he encourages us to contact him and he will respond.

Contact Info

David Duffey
Director of Technical Partnerships
512-850-6776 (work)
512-287-4289 (work fax)
david.duffey@canonical.com

Comments from the group

Bob Noseworthy

- 1) The OFA IWG is here to promote Open Fabric Stack. Our members should promote support for Ubuntu because it will benefit them.
- 2) UNH-IOL will have some Ubuntu deployed and tested in October and may have some onsite support from Ubuntu. We need support from Open Fabrics community.

Jim Ryan

- 1) Rupert asked if there is any problem allowing Canonical to attend the event even though they are not currently and OFILG member. Jim says there is nothing in the By-Laws that would prevent that and so it is not a problem
- 2) Jim says that Canonicals involvement is a huge development and he wants to support their efforts. He would encourage them to continue to attend the work group meetings

User Day event in Zurich

- 1) As of 9/11/2012 attendance was still light weight but IBM is still behind it and wants it supported.
- 2) The [event was postponed](#) until early 2013 as of 9/17/2012

Review of the Interop Test Plan updates submitted by Edward from UNH-IOL

Rob Noseworthy

- 1) The OFA IWG should focus on RoCE and DCB – stay away from fabric testing. Please send comments on demands and interest in RoCE

SRP

Issues in OFED 3.5 RC1

- 1) Nothing works on any target and causes a system failure when you try and load the SRP module. The 6.3 SCSI API does match.
 - a) UNH-IOL tested with both SL 6.2 and SL 6.3 x86_64 using OFED 3.5 RC1
 - b) Hardware was an Intel PCIe Gen 2 system with a Mellanox ConnectX 2 HCA and a DDN 10K SRP target using a Mellanox ConnectX 2 card.
 - c) After issuing the command "modprobe ib_srp", the system becomes unresponsive and then the system hard locks and causes the user to have to reboot the system.
- 2) Rupert has logged this in OFA Bugzilla as [Bug # 2378](#)
- 3) Members want to know why did this not get caught
 - a) Rupert responded that not everyone has targets to test. Also this is part of the OFED release cycle and this was the first release candidate. This problem was identified by both UNH-IOL and DDN. Vendors need to do more testing and get active in testing the RCs. This is Open Source software and we are all responsible for its maintenance and testing.

OSU MPI

Edward suggested removing this section but Rupert reminded the group that this is still an active Open Source MPI and it should remain in the document but moved to a deprecated status. This way vendors and end users can use the test suite if they are still interested in OSU MVAPICH.

iSER

Edward suggested removing this section but Rupert reminded the group that this is still an active ULP and it should remain in the document. We don't currently have any members using iSER but if we did this would be a mandatory test.

NFS RDMA Beta

RDMA over converged to glossary

Inter versus Intra – homo versus hetero

New sections

8.X – OS Installation

8.X.1 – OS Policy

Edward added these sections so we had a framework for more details which will follow regarding the setup for Ubuntu, RHEL, SUSE etc.

Section 9.3

As a requirement for the OFILG Logo, a vendor's DUT must pass all mandatory testing using an agreed upon primary OS and OpenSM. Additional beta testing is performed using secondary, tertiary, and quaternary Operating Systems. This beta testing has no bearing on whether the OFILG Logo is granted to a device, it is purely informative. OpenSM and SL are mandatory

10.6 SRP using OFED

- 1) Marty from DDN has provided some updated information regarding SRP. For example the command `srp_sg_tablesize=255` has been deprecated.
- 2) There will be Part A and Part B for the next event but Part B needs to be implemented. Marty will work with Edward to get this running and then Rupert will add to the test plan.

Interop test Plan review – from 8/28/2012

- 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 dapltest.
- 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 – [Intel will support 1.6.1 – verified 9/11/2012](#)

Mellanox - Guy

Rupert Dance

9/11/2012

- 1) Send out invitation and NDA to David Duffey from Canonical – [done 9/12/2012](#)
- 2) Update Interop test Plan with approved comments and changes from the meeting on 9/11/2012

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 – filed Bug 2378](#)
- 3) Follow up on status of Master report including all Inter OS testing for May 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

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.