

OFA EWG Meeting – 11/23/2015

Attendees

First	Last	Company	10/12/15	10/26/15	11/09/15	11/23/15
Pradeep	Kankipati	Avago	X	X		X
Chuong	Vuong	Avago		X		
Steve	Wise	Chelsio			X	X
Dwip	Banerjee	IBM		X		
David	Dai	IBM	X	X		X
Paul	Bowden	Intel		X		
Tom	Elken	Intel		X		
Tatyana	Nikolova	Intel	X			
Jim	Ryan	Intel	X			X
Robert	Woodruff	Intel	X	X	X	X
Vladimir	Sokolovsky	Mellanox	X	X	X	X
Jeff	Becker	NASA			X	X
Rupert	Dance	Software Forge	X	X	X	X
John	Jolly	SUSE		X		

Meeting Minutes 11/09/2015:

- https://www.openfabrics.org/downloads/WorkGroups/ewg/EWG_Minutes/2015-11-09-EWG-Meeting-Minutes.pdf

OFED 3.18-1 status

- **RC2 – 10/26/2015**
 - <https://www.openfabrics.org/downloads/OFED/ofed-3.18-1/OFED-3.18-1-rc2.tgz>
- **OFED User Package Status as of 11/23/2015 – updates since 11/09/2015 meeting**
 - https://www.openfabrics.org/downloads/WorkGroups/ewg/OFED_User_Package_Status/OFED-User-Package-Status-2015-11-23.pdf
 - libfabric-1.2.0rc1 – released 11/13/2015

OFED 3.18-1 Bug Status as of 11/23/2015

- https://www.openfabrics.org/downloads/WorkGroups/ewg/OFED_Bugs/OFED-3.xx-Bug-Status-2015-11-23.pdf
- [2492](#) – **Chelsio** - perfest tools enhancement – opened 6/17/2014 – **can we close this?**
- [2553](#) – **Intel** - Unable to run UD traffic using qperf – opened 6/29/2015
 - Assigned to Vinit from Intel who has confirmed issue and working on solution
 - Pradeep Kankipati from Avago is going to ask his team to look at this – 10/26/2015
- [2561](#) - **Mellanox** - IPoIB failure after arp cache flush – opened 7/20/2015
 - Assigned to erezsh@mellanox.com – no action in last month
 - This is also a problem in the current 4.3 kernel and is not easy to fix
- [2562](#) – **Mellanox** - ib_write_bw doesn't support events - enhancement
- [2574](#) – **Intel** - QIB makes SCIF does not work with Xeon-phi – target 3.18-2
- [2580](#) – NFSoRDMA PPC64LE client crashes while trying to mount NFSoRDMA server running on x86 box target 3.18-2

OFED 3.18-x

- OFED 3.18-1 - ready for GA
- OFED 3.18-2 – what do we want to include
 - Distro Versions: RHEL 7.2 and SLES 12 SP1
 - Bug Fixes and new features

Start of the meeting

OFED 3.18-1 Bug Status as of 11/23/2015

- https://www.openfabrics.org/downloads/WorkGroups/ewg/OFED_Bugs/OFED-3.xx-Bug-Status-2015-11-23.pdf
- [2492](#) – **Chelsio** - perfest tools enhancement – opened 6/17/2014 – can we close this? Yes Steve will close it.
- [2553](#) – **Intel** - Unable to run UD traffic using qperf – opened 6/29/2015
 - Assigned to Vinit from Intel who has confirmed issue on SFI hardware but does not have resources to fix it.
 - Pradeep Kankipati from Avago is going to ask his team to look at this – 10/26/2015
 - Pradeep has talked to his team but it will take 2-3 weeks – 11/23/2015
- [2561](#) - **Mellanox** - IPoIB failure after arp cache flush – opened 7/20/2015
 - Assigned to erezsh@mellanox.com – no action in last month
 - This is also a problem in the current 4.3 kernel and is not easy to fix – target 4.5
- [2562](#) – **Mellanox** - ib_write_bw doesn't support events – enhancement – target 3.18-2
- [2574](#) – **Intel** - QIB makes SCIF does not work with Xeon-phi – target 3.18-2
 - **Woody**: he thinks it is fixed in RC2 – he sent email
 - **Woody**: there is a script and a work around.
- [2580](#) – NFSoRDMA PPC64LE client crashes while trying to mount NFSoRDMA server running on x86 box target 3.18-2
 - Jeff spent a lot of time with Chuck Lever and UNH-IOL
 - There is a PCIe error that is going on which they feel has to be firmware or hardware. They feel there is nothing related to NFSRDMA. However this issue does not occur with any other test they run.
 - **David** – in the test plan there is RSocket and other tests that work.
 - Stefan from UNH-IOL is going to try and move the card to another slot or another box.
 - When debug mode is on the mount still fails. Then whittle down to one flag and mount worked and Connectathon failed.
 - **Dave Dai**: they have only one Power 8 Box that supports this. The other is P7. There may be more than one PCIe slot.
 - Power 7 to Power 7 is no problem.
 - Power 7 to x86
 - **Dave and Jeff** are both OK with letting this go until 3.18-2

OFED 3.18-1 - ready for GA

Woody and Vlad are both OK with moving forward on the release.

OFED 3.18-2 – what do we want to include

Distro Versions:

- RHEL 7.2 released on 11/19/2015 – this will require more work required on backports.
- SLES 12 SP1 released on 11/18/2015 – not much work on backports

Bug Fixes and new features

- **Woody**
 - Libfabric 1.2.0 GA is before end of year – enhanced verbs provided for tag messages so MPI over the following providers:
 - MXM
 - Cray GNI
 - OPA but there is no driver in the kernel.

- No DAPL
- **Steve**
 - Does not want to try and pull backports into 3.18-2 – just too much work
- **Pradeep**
 - Why 3.18-2?
- **Woody**
 - We released 3.18-1 so people can use RHEL 7.1 right now. But now people want support for RHEL 7.2 and SLES 12 SP1
 - We are also waiting for RoCE v2 and Omni Path and other stuff. These will all require at least kernel 4.5 and so we need to wait on new kernel.
- **Rupert**
 - Libibverbs – this is still not GA and is in parallel with other stuff like MLNX enhanced drivers. So we must wait and see.
 - Target next big OFED

Opens

Woody

There are two new Co-Chairs for IWG - Paul Bowden from Intel and an engineer from Mellanox.

Super Computer 2105 Report

- Sean went and presented a BoF on OFiWG – fairly well attended.
- New version of MPICH which is the code base of MVAPICH and Intel MPI. The new MPICH 4 architecture is more optimal for Exascale.
- More support for UCX and MPI. UCX is used on the big Coral machines from Mellanox.

ARs

Chelsio – Steve Wise

9/15/2014

- 1) Send patch for Bug [2492](#) upstream to the perftest maintainer – **in process**
 - a) Steve agreed to mark this as Resolved Later until they have resources to allocate to this bug – [11/23/2015](#)

Avago – Pradeep

10/26/2015

- 1) Check with your tem and see if they can help to resolve Bug [2553](#) – Unable to run UD traffic using qperf on RoCE

Intel –IB

7/6/2015

- 1) [2553](#) - Unable to run UD traffic using qperf - **Tom** to follow up
 - a) **Pradeep**: Avago has shared all info with Intel (issue, repro details and root cause) and they have been able to reproduce it. **No action pending from our side - 10/20/2015**
- 2) **Rupert** has given access to RoCE servers at Software Forge - **Assigned to Vinit from Intel on 9/9/2015**

Mellanox – Vlad

10/12/2015

- 1) libibverbs-1.2.0 – Contact Doug and explain the problems

- a) **Woody** says this could require kernel changes for support and may need MLX changes
- b) **Vlad** will check and verify. May need 4.3 - new features will not be available without new kernel.

8/3/2015

- 1) Talk to Eli and others about Bug [2561](#) – IPoIB
- 2) Work with Gill to resolve [2562](#) – ib_write_bw doesn't support events

SFI – Rupert

10/26/2015

- 1) Review and update known issues in the OFED Release notes – [done 11/22/2015](#)

8/4/2015

- 1) We need to add a note in add note in 3.18-1 that ipath and eHCA will be deprecated in OFED 4.3 – [done 11/22/2015](#)