

OFA EWG Meeting – 7/6/2015

Attendees

First	Last	Company	05/11/15	06/08/15	06/22/15	07/06/15
Eddie	Wai	Broadcom	X	X		X
Steve	Wise	Chelsio	X	X		X
Pradeep	Kankipati	Avago	X	X	X	X
Pradeep	Satyanarayana	IBM		X		
Arlin	Davis	Intel		X		X
Mike	Marciniszyn	Intel			X	X
Tatyana	Nikolova	Intel	X	X	X	X
Jim	Ryan	Intel	X	X	X	
Robert	Woodruff	Intel	X	X	X	
Vladimir	Sokolovsky	Mellanox	X	X	X	X
Jeff	Becker	NASA	X		X	
Rupert	Dance	Software Forge	X	X	X	X
Keith	Hoodlet	UNH-IOL				X
Dave	Wyman	UNH-IOL	X	X		

Agenda

Meeting Minutes 6/22/2015:

- https://www.openfabrics.org/downloads/WorkGroups/ewg/EWG_Minutes/2015-06-22-EWG-Meeting-Minutes.pdf

OFED 3.18 status

- **OFED 3.18 RC3**
 - <https://www.openfabrics.org/downloads/OFED/ofed-3.18/OFED-3.18-rc3.tgz>
- **OFED 3.18**
 - First Daily Build 10/28/2014
 - RC1 – 3/11/2015
 - RC2 – 4/13/2015
 - RC3 – 6/22/2015
 - GA – target is 7/9/2015
- **OFED User Package Status as of 7/5/2015**
 - https://www.openfabrics.org/downloads/WorkGroups/ewg/OFED_User_Package_Status/OFED-User-Package-Status-2015-07-05.pdf
 - libiwpm
 - mstflint
 - perftest
- **OFED 3.18 Bug Status as of 7/5/2015**
 - https://www.openfabrics.org/downloads/WorkGroups/ewg/OFED_Bugs/OFED-3.xx-Bug-Status-2015-07-05.pdf
- **OFED 3.18 open issues**
 - [2553](#) - Unable to run UD traffic using qperf
 - [2554](#) - ibscif driver fails to bring up scif device on kernels >= 3.11
- **OFED 3.18-1** – issues to be resolved in the **next OFED**
 - [2544](#) – infinipath-psm won't install with libfabric
 - [2545](#) - ipath compilation fails on SLES11 SP4 and RHEL6.7
 - [2547](#) – libibscif and xeon-phi issues
 - [2551](#) – libfabric compile issue on ppc64
 - [2555](#) - ib_send_bw –run_infinately fails to terminate on the server after client disconnects

- **Release notes status for OFED 3.18 GA**

- https://www.openfabrics.org/downloads/WorkGroups/ewg/OFED_User_Package_Status/OFED-Release%20Notes-Status-2015-07-05.pdf
- Release notes updates either missing or needed

Package	Package Update	Notes Update	Needs Update	Maintainer
diags_release_notes.txt	02/10/15	05/24/14	Yes	Ira Weiny
ibpd	08/19/14		Yes	Jianxin Xiong
ibsim	01/30/14		Yes	Hal Rosenstock
ibutils_release_notes.txt	03/20/14	08/19/13	Yes	Oren Kladnitsky
ipath_release_notes.txt	07/24/14	08/19/13	Yes	Mike Marciniszyn
ipoib_release_notes.txt		05/19/14	Yes	Vlad Sokolovsky
libibcm	10/26/09		Yes	Sean Hefty
libibmad	02/10/15		Yes	Ira Weiny
libibscif	08/19/14		Yes	Jerrie Coffman
libibumad	03/03/15		Yes	Hal Rosenstock
libibverbs	05/05/14		Yes	Roland Drier
libipathverbs	07/24/14		Yes	Mike Marciniszyn
libmlx4	05/07/14		Yes	Roland Drier
libmlx5	01/12/15		Yes	Eli Cohen
libmthca	07/06/11		Yes	Vlad Sokolovsky
libnes	02/18/14		Yes	Tatyana Nikolova
mpss			Yes	Phil Clayton
mstflint_release_notes.txt	06/17/15	05/18/14	Yes	Oren Kladnitsky
opensm_release_notes.txt	12/17/14	05/18/14	Yes	Hal Rosenstock
perftest	07/04/15		Yes	Gil Rockah
psm_release_notes.txt	10/10/14	05/25/14	Yes	Mike Marciniszyn
qib_release_notes.txt		05/25/14	Yes	Mike Marciniszyn
srptools_release_notes.txt	02/11/15	05/25/14	Yes	Bart Van Assche

Next Meeting: July 20th, 2015

Bug Status

OFED 3.18 open issues

- [2553](#) - Unable to run UD traffic using qperf - Changed from enhancement to normal. Mike says the solution looks reasonable. qperf is used internally by Intel for True Scale products but they have not used it with UD. There is no need to hold up OFED 3.18 GA and so we have targeted OFED 3.18-1.
- [2554](#) - ibscif driver fails to bring up scif device on kernels >= 3.11 – Arlin says that they already have a fix for this but they had not provided it to Vlad. Since we are making changes to OFED 3.18, he would like to include this.

OFED 3.18-1 – issues to be resolved in the next OFED – these are all issues that have little impact on the core functions of OFED and so we decided to wait for OFED 3.18-1

- [2544](#) – infinipath-psm won't install with libfabric
- [2545](#) - ipath compilation fails on SLES11 SP4 and RHEL6.7 – Distros not supported in OFED 3.18
- [2547](#) – libibscif and xeon-phi issues
- [2551](#) – libfabric compile issue on ppc64 – disabled by default on ppc64 in OFED 3.18
- [2555](#) - ib_send_bw –run_infinately fails to terminate on the server after client disconnects – new perftest bug that uses a somewhat esoteric flag. Non critical

Bug [2557](#) – Critical - loading of various rdma modules fail in RHEL7.1

Steve:

- 1) The problem is the result of the difference between inbox and OFED modules.
- 2) He did a typical OFED install using the command: *install -all*
 - a) After he rebooted the problems showed up
 - b) Sometimes rdma drivers from core get loaded and not the OFED versions
- 3) He did not change openibd.conf and he does not have mlx cards.

Vlad and Rupert

- 1) From [Wikipedia](#): initramfs, abbreviated from "initial RAM file system", is the successor of initrd. It is a cpio archive of the initial file system that gets loaded into memory during the Linux startup process. The Linux kernel mounts it as root file system and starts the init process from it. This will complete certain tasks before the real root file system is loaded; thus, initramfs **needs to contain all of the device drivers** and tools needed to mount the real root file system. Additional notes available from [Ubuntu](#).
- 2) So initramfs is causing problems because it is loading all the “inbox” drivers which conflict with the OFED device drivers. Initramfs has all the core modules.
- 3) Vlad says there is already logic in openibd to unload everything first and then start new modules. However it will only work for Mellanox modules and devices. If a module is loaded from initramfs, then openibd will unload modules loaded by initramfs and reload from OFED. This maybe a problem with RHEL 7.1 because we have not seen it before. The modules are in /boot/ initramfs.
- 4) There are several solutions to this problem.
 - a) Modify openibd so that it unloads **all** “device drivers” loaded by initramfs and then load OFED device drivers.
 - b) Create a blacklist for initramfs to prevent certain modules from loading: see <http://ubuntuforums.org/showthread.php?t=166624>
 - c) Unload the modules manually and then load the correct OFED Modules. Not all modules have to be unloaded – for example you don't need to do IPoIB or ib_core components. It is only low level drivers.
 - d) Add it to release notes as a known issue.
- 5) Vlad will create a daily build with a fix in openibd and ask Steve to test it.

Tatyana

- 1) You can list the content from initramfs and then see the modules and see what is loaded at boot time.
- 2) The command is [lsinitrd](#)
- 3) In her experience this is a new problem with RHEL 7. One solution is to remove the modules loaded by initramfs.

Discussion regarding whether we should test this with a daily build or an RC

Vlad: things we need to do in the new build

- 1) Tatyana committed some changes that need to be updated. The fix is a kernel patch
 - 2) Create a fix for Bug [2557](#) and the initramfs problem
-
- 1) **Arlin** - He feels that since it is a Kernel patch then we should do an RC. If we do another RC then Intel will submit changes as well.
 - 2) **Rupert** - I feel we should do an RC. The problem described in Bug [2557](#) seems fairly serious and is listed as critical. We need to make sure we get this fixed correctly and test it well.
 - 3) **Steve Wise** – he leans toward a new RC. We should innumerate the issues and send an email so we can confirm and validate and then go to GA.
 - 4) **Tatyana** – she can go with daily or RC
 - 5) **Pradeep** – if we can fix this problem and insure that it is well contained, then he is OK with Daily Build
 - 6) **Mike** – he feels it is ok to do it with a daily build
 - 7) **Eddie** – they have no code in OFED but they really want the GA version to be released. If we can fix this and therefore the issue is not critical, he would vote to do a daily build and then release the GA.

Conclusion:

Arlin:

- 1) If we understand what is in daily build and what needs to be test and when it will get done, then he is OK with a daily build.
- 2) Here are the changes needed:
 - a) **Tatyana:** updates to drivers/infiniband/hw/nes/nes_cm.c
 - b) Bug [2554](#): Driver issues with ibscif and SLES 12 – It is fixed but not sent to Vlad. He would like to add because it is very isolated. **Vlad** is OK including this patch because it is not a default package that is compiled automatically.
 - c) Bug [2557](#) – see discussion above.
- 3) **Daily Build** will be submitted as soon as the above changes have been submitted to Vlad.
- 4) **Deadline** for review and approval of Daily Build 7/13/2015
- 5) **GA Release:** 7/14/2015

Bug [2553](#) - Unable to run UD traffic using qperf

- 1) Intel uses qperf internally by for True Scale products but they have not used it with UD. Mike says the proposed solution looks reasonable.
- 2) Rupert changed the Bug from enhancement to normal.
- 3) Mike feels there is no need to hold up OFED 3.18 GA and so we have targeted OFED 3.18-1.

Release Notes

- 1) We reviewed the list of release notes that are either missing or need to be updated. The details are available in the following PDF:
 - a) https://www.openfabrics.org/downloads/WorkGroups/ewg/OFED_User_Package_Status/OFED-Release%20Notes-Status-2015-07-05.pdf
- 2) Tatyana will do release notes for libnes
- 3) Arlin will follow up on MPSS and libibscif
- 4) Mike has already committed to review and update qib, ipath and others that are part of the Intel packages.
- 5) Vlad will follow up with Mellanox staff and get their release notes committed.

Rupert

- 1) There are two user packages that show Roland as the maintainer and we discussed whether these should be reassigned now that Doug Ledford is the RDMA maintainer.
- 2) Vlad felt that we should ask Doug Ledford to take over libibverbs
- 3) Vlad says that he has another Mellanox developer for libmlx4 – he will change it or send Rupert the info.

Opens

Arlin – GA is firm so no other opens.

Meeting adjourned at ~12:45

ARs

All package maintainers

5/11/2015

- 1) Make sure we get release notes updated and each vendor should get updates in – [partial 7/5/2015](#)

Chelsio – Steve Wise

9/15/2014

- 1) Send patch for Bug [2492](#) upstream to the perftest maintainer and ask them to check out and accept the patch – [in process](#)

Emulex

Intel –IB

7/6/2015

- 1) Arlin will follow up on MPSS and libibscif release notes
- 2) Mike - review and update qib, ipath and others release notes that are part of the Intel packages.

Intel iWARP – Tatyana

7/6/2015

- 1) Update kernel patch submitted to Vlad
- 2) Create release notes for libnes.

Mellanox - Vlad

7/6/2015

- 1) Vlad will follow up with Mellanox staff and get their release notes committed.

NFSRDMA – Jeff

SFI – Rupert

6/22/2015

- 1) Ask Phil Cayton – Patrick McCormick to do a release notes for MPSS – [sent request 7/5/2015](#)

7/6/2015

- 1) Contact Doug Ledford about libibverbs – [done 7/6/2015](#)
- 2) Send out summary regarding the daily build required for the release of OFED 3.18 GA – [done 7/6/2015](#)

UNH-IOL