



Meeting of the Technical Steering Committee (TSC) Board

Wednesday, May 20th, 2020
11:00am ET

Meeting Logistics

- <https://zoom.us/j/556149142>
- United States : +1 (646) 558-8656
 - Meeting ID: 556 149 142

Antitrust Policy Notice

- Linux Foundation meetings involve participation by industry competitors, and it is the intention of the Linux Foundation to conduct all of its activities in accordance with applicable antitrust and competition laws. It is therefore extremely important that attendees adhere to meeting agendas, and be aware of, and not participate in, any activities that are prohibited under applicable US state, federal or foreign antitrust and competition laws.
- Examples of types of actions that are prohibited at Linux Foundation meetings and in connection with Linux Foundation activities are described in the Linux Foundation Antitrust Policy available at <http://www.linuxfoundation.org/antitrust-policy>. If you have questions about these matters, please contact your company counsel, or if you are a member of the Linux Foundation, feel free to contact Andrew Updegrave of the firm of Gesmer Updegrave LLP, which provides legal counsel to the Linux Foundation.

Agenda/Updates

- Announcements:
 - **Reminder: TSC call for nominations has been published**
 - nominations due by Friday, June 12 (tsc-nominations@OpenHPC.groups.io)
 - Upcoming deadlines:
 - **SC'20**
 - Tutorials: **Extended to May 08** – Thanks David B. for submitting!
 - BoFs: Due July 31, 2020
-

- PEARC'20 tutorial/cloud working group updates (csim)
- Package conflict update
- UCX/libfabric (continued from last time)
- Build/CI results
- Proposed RC recipes
- Draft high-level summary of 2.0 RC1
- Reminder on intent to use new repo

OHPC Cloud Working Group Updates

- Virtual PEARC confirmed; asked us to fill out another survey
 - Asked if willing to prerecord
 - Available for office hours outside of time slot
 - Available for Monday or Thursday
- Looking for ideas on how best to present self-directed educational content
 - CentOS 7 and OHPC 1.3.9 working
 - Simmons starts CentOS 8 and OHPC 2.0 next week
 - Students start documentation this week

Package conflicts update

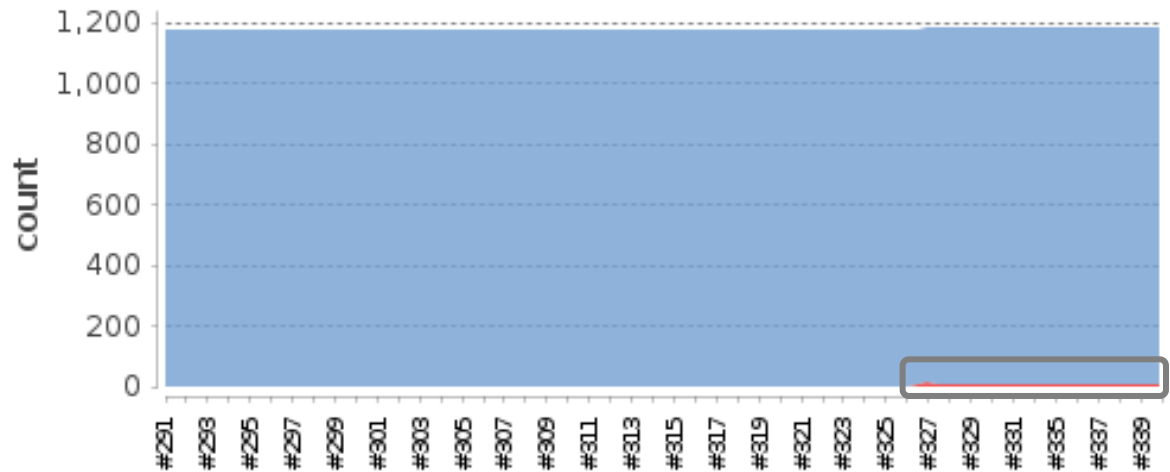
- Recall that we were getting package install conflicts for /usr/lib/.build-id files on rhel systems
- Added the following for global disable of these links:

```
# Disable generation of .build-id links  
%global _build_id_links none
```


- gcc+intel (or gcc+arm) packages can now co-exist peacefully
- much better CI results with gcc+intel:
 - went from 1,209 passing CI tests with gcc+intel to 2,294 tests presently

mpich: ucx/libfabric (continued discussion)

- From last time:
 - MPICH ch4 build using UCX introduced some repeatable regression failures in several packages running on tcp network



- Since then, rebuilt mpich against our build of libfabric
 - regressions have been resolved (on CI systems at TACC)
 - 1 or 2 intermittent regressions using tcp on Linaro aarch64



mpich: ucx/libfabric (continued discussion)

- I mentioned last time that I do not believe we can have a monolithic build that is capable of using both ucx and libfabric
 - reached out to mpich contact at Argonne to hopefully confirm/deny
 - here is response from Ken:

*"You *can* have a monolithic build, but it is not recommended. There is no long-term preference between the two, but note that Mellanox Infiniband support is much better with ch4:ucx. Both builds should support TCP networks for maximum compatibility."*

- Given this, I don't think it makes sense to try and maintain a monolithic build for mpich
 - also, since we see more consistent results with libfabric, propose we go with that for default build (but, recall that IB large message performance is better out of the box with UCX)
 - we can provide a mpich-ucx-gnu9 build for folks who want to use UCX (drop in replacement)
- For openmpi4, we can continue to have a monolithic build and our latest has been updated to use ohpc builds of ucx and libfabric
- Updated MPI table in docs:

Table 1: Available MPI variants

	Ethernet (TCP)	InfiniBand	Intel® Omni-Path
MPICH (libfabric)	✓	✓	✓
MPICH (ucx)	✓	✓	✓
MVAPICH2		✓	
MVAPICH2 (psm2)			✓
OpenMPI (libfabric/ucx)	✓	✓	✓

- need to revisit packaging slightly


2.0 Build Status (cont)

- Current package counts:
 - ~~800 RPMs as of 3/11/2020~~
 - ~~905 RPMs as of 3/25/2020~~
 - ~~982 RPMs as of 4/8/2020~~
 - ~~1003 RPMs as of 4/22/2020~~
 - ~~1021 RPMs as of 5/6/2020~~
 - 1064 RPMs as of 5/20/2020


Base OS	aarch64	x86_64	noarch
CentOS 8	174	334	29
Leap 15	172	326	29

2.0 (cont.) – CI updates

centos8.1/x86/WW/slurm (3/10/2020)



centos8.1/x86/WW/slurm (05/20/2020)



- As of 3/10:
 - ~~610 user level tests passing~~
- As of 3/25:
 - ~~910 user level tests passing~~
- As of (4/08):
 - ~~982 user level tests passing~~
- As of (4/22):
 - ~~1,132 user level tests passing~~
- As of (5/06):
 - ~~1,140 user level tests passing~~
- As of this morning (5/20):
 - 1,173 user-level tests passing

2.0 (cont.) – CI updates


2.0RC1	All	+	S	Categorized - Job	Last Success	Last Failure	Last Duration	Test Result
			[-]	! .. » [aarch64]	5 days 22 hr - #64	9 hr 49 min - #22	1 hr 42 min	N/A
				✓ (2.0) - (centos8.1,aarch64) (warewulf+slurm) (fabric=eth)	5 days 22 hr - #64	6 days 9 hr - #63	1 hr 42 min	0 of 1,102 failed (-1)
				! (2.0) - (leap15.1,aarch64) (warewulf+pbspro) (fabric=eth)	N/A	19 hr - #6	2 hr 0 min	5 of 1,080 failed (+1)
				! (2.0) - (leap15.1,aarch64) (warewulf+slurm) (fabric=eth)	1 mo 16 days - #3	9 hr 49 min - #22	7 min 57 sec	1 of 1,088 failed (-1)
			[-]	🔄 .. » [x86_64] - CentOS 8	1 hr 44 min - #457	5 hr 5 min - #97	1 hr 3 min	N/A
				✓ (2.0) - (centos8.1,x86_64) (warewulf+slurm) (fabric=eth) - UEFI	1 hr 44 min - #457	2 days 19 hr - #435	1 hr 3 min	0 of 1,173 failed (±0)
				🔄 (2.0) - (centos8.1,x86_64) (warewulf+slurm) (fabric=ib) + psxe	4 days 16 hr - #68	5 hr 5 min - #97	6 min 1 sec	0 of 2,294 failed (±0)
			[-]	! .. » [x86_64] - Leap15	1 hr 54 min - #434	2 hr 28 min - #20	1 hr 9 min	N/A
				🔄 (2.0) - (leap15.1,x86_64) (warewulf+pbspro) (fabric=eth)	3 hr 52 min - #47	6 days 15 hr - #7	58 min	0 of 1,151 failed (±0)
				✓ (2.0) - (leap15.1,x86_64) (warewulf+slurm) (fabric=eth)	1 hr 54 min - #434	9 hr 54 min - #432	1 hr 9 min	0 of 1,159 failed (±0)
				! (2.0) - (leap15.1,x86_64) (warewulf+slurm) (fabric=ib)	1 mo 26 days - #9	2 hr 28 min - #20	3 min 37 sec	5 of 2,370 failed (+5)

← intermittent

← much improved!

Proposed RC1 Recipes

- Subset of the 10 recipes we have in v1.3.9 (6 total)
 - CentOS 8.1
 - Warewulf+SLURM (x86_64)
 - Warewulf+SLURM (aarch64)
 - Leap 15.1
 - Warewulf+SLURM (x86_64)
 - Warewulf+SLURM (aarch64)
 - Warewulf+PBS Pro (x86_64)
 - Warewulf+PBS Pro (aarch64)



OpenHPC (v2.0.0 RC1)
Cluster Building Recipes

CentOS 8.1 Base OS
Warewulf/SLURM Edition for Linux* (x86_64)

Release
Candidate

VERSION
2.X

Document Last Update: 2020-05-20
Document Revision: 8e9f8f45b

2.0 RC1 – Summary of Major Changes

- Introduces support for **CentOS8.1** and **Leap 15.1**
 - not backwards compatible with 1.3.x – intended for fresh installs with new OS versions
 - packaging modifications to most development packages to facilitate easier custom package rebuilds locally (a la OHPC_CFLAGS, etc)
- MPICH updated to use new ch4 interface (with libfabric) -> now supports IB fabrics
- OpenMPI variant updated to openmpi4 (with libfabric/UCX support)
- gcc compiler variant updated to gcc9
 - note: gcc10 just came out on May 7th ☹
- SLURM 20.x series (using configless option in recipes)
- introduces subset of ARM HPC compiler builds (along with compatibility package)
- Deprecated ohpc builds:
 - nagios (rely on distro version instead)
 - munge (rely on distro version instead)
 - ganglia
 - mpiP
 - ocr
 - mrsh
- New components:
 - libfabric-ohpc
 - ucx-ohpc

DRAFT items to highlight
with RC1 release

2.0 RC1 – Known Issues

DRAFT items to highlight with RC1 release

- xCAT based recipes not yet available
- Lustre client and BeeGFS testing not yet available
- Packages not available:
 - dimmemas
 - Boost with Intel compilers
 - Trilinos with Intel compilers (due to Boost dependency)
 - geopm with Intel compilers
 - openmpi4-pmix and companion slurm-pmix
- Highlight subset of packages available for arm-hpc
 - mpich/openmpi4
 - fftw
 - gsl
 - hdf5
 - metis
 - plasma
 - ptscotch
 - superlu,superu_dist
 - etc....

} removed from relevant meta-packages

New Repository reminder

- Recall that we have discussed previously divorcing the OBS build system and user-facing package repository
- Looked at hosting packages directly in s3 on EC2, or a separate EC2 instance
 - in the end, I like the standalone EC2 instance as we can take advantage of the log monitoring facility we have in place for apache (to track repository usage) and make it easy to browse the repo
 - <https://repos.openhpc.community> is fired up again
- Last steps prior to RC release:
 - tidy up everything we have
 - finalize the 6 recipes (need to include arm-hpc compiler discussion in aarch64)
 - highlight updated build from src rpm option
 - create Jenkins release mechanism to push to new repository
 - add Jenkins CI job to test directly from repos.openhpc.community
 - finalizing mpich packaging so that both ucx and libfabric can be installed at the same time
 - other thoughts??

