



Meeting of the Technical Steering Committee (TSC) Board

Wednesday, April 22nd, 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

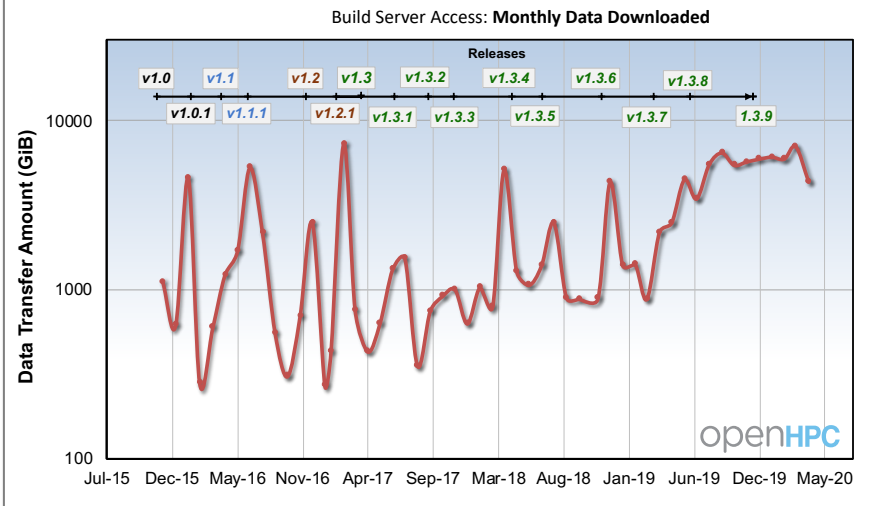

- Linux Foundation meetings involve participation by industry competitors, and it is the intention of the Linux Foundation to conduct all of its activities in accordance with applicable antitrust and competition laws. It is therefore extremely important that attendees adhere to meeting agendas, and be 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:
 - ISC in-person event **cancelled** – will be a digital event instead
 - still have not heard anything yet regarding our accepted BoF
- Upcoming deadlines:
 - SC'20
 - Tutorials?: Due ~~April 16, 2020~~, Extended to April 30 – *Adrian leading a submission)
 - BoFs: Due July 31, 2020
- ARM compiler build work back on track

-
- Quarterly access stats
 - PEARC'20 tutorial/cloud working group updates (csim)
 - [Netdata](#) eval (csim)
 - Additional component deprecation
 - Component Review #9 Results
 - 2.0 build status


Updated Usage/Access Statistics (thru Q1 2020)



- Stats for build/repo server (tracking # of unique visitors per month and amount of data downloaded):

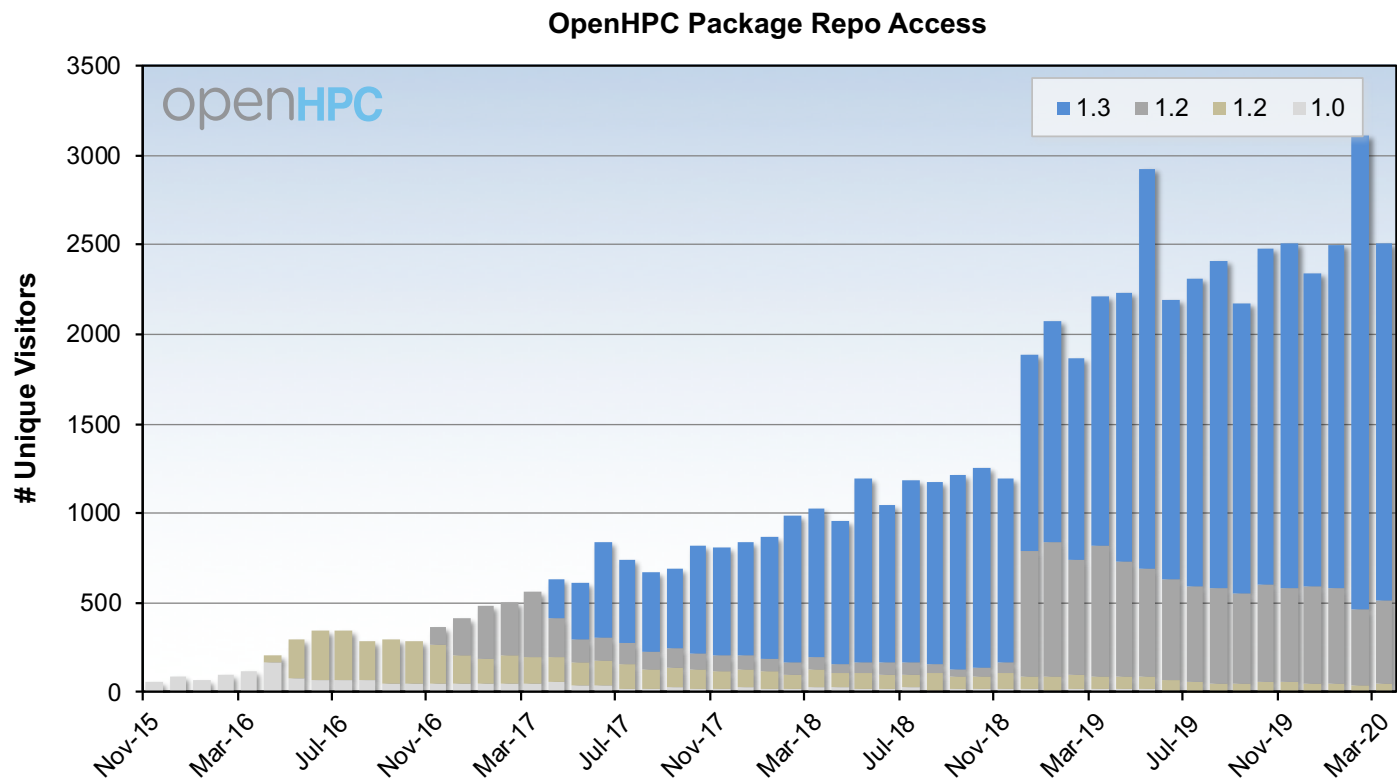
<http://build.openhpc.community>

- Averaging ~5.7 TB/month download in 2020



Updated Usage/Access Stats (thru Q1 2020)

- These stats monitor access specifically to package repository metadata (typically expected to be via yum/zypper)
- Repo access binned by minor version




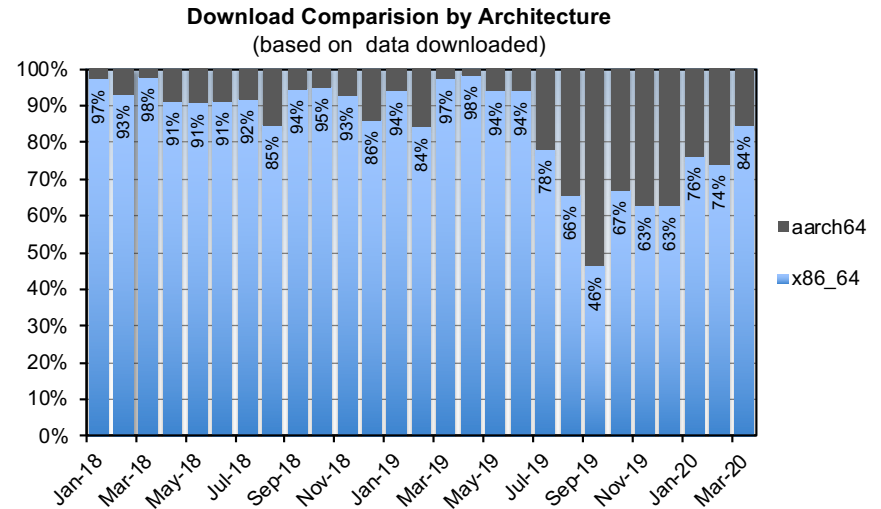
Updated Usage/Access Stats (thru Q1 2020)

Architecture specific metrics:

- To provide some characterization, we scrape the access logs to analyze two architecture specific file types:

`(aarch64 | x86_64) .rpm`
`(aarch64 | x86_64) .tar`

- Plots compare percentages for the amount of data xfer'ed and the # of unique visitors accessing the (aarch64|x86_64) files



OHPC Cloud Working Group Updates (csim)


- PEARC is now confirmed to be a virtual [event](#)
- Was no cloud working group meeting this week
- Private GitHub repo for project was moved to the OHPC organizational GitHub account
- Let us know if you want access
- Have successfully deployed the “full cluster” version
- Beginning to work more on documentation and getting the OHPC version finished up

Netdata Eval/Overview (csim)

- A GitHub request for the inclusion of Netdata as a monitoring solution for OHPC was suggested
- Ganglia's last release was 2014?
- Netdata is:
 - Real-time health monitoring and performance troubleshooting solution
 - Not a drop in replacement for Ganglia
 - More distributed by design but can be used in a client / server mode with configuration
 - ... AMAZING

Netdata cont.

- Very lightweight
 - 3-5% CPU of a single core when sampling every second
 - Half that when collecting only (more on that in a bit)
 - 40-80 MB of RAM used depending on settings
- Amazingly intuitive interface



Netdata basic concepts

- Each node can either have a database of results locally or not
- If database is set to none, the node is a collector only; You can then set it up in a client/server configuration.
- If a database is configured, a web portal is also run on that node to display the results. Default is port 19999
- Database types:
 - RAM – monitoring data is stored in RAM only; ~100 MB will give you about 6 hours of backlogs depending on the number of collectors that get configured
 - Dbengine – RAM + disk; most recent data is in RAM and older data gets written to disk; Web UI has access to all data just a little slower accessing older data

What netdata doesn't do today

- At this time, it doesn't provide "cluster views"
- It has no concept of a larger machine
- It is not easy to compare metrics across multiple nodes
- Netdata Cloud (releasing next week) might provide some solutions
- However, Netdata Cloud will require using their cloud service initially

BigTex Grid (1 sources) (tree view)

CPU Total:	13648
Hosts up:	850
Hosts down:	0

Longer Term Storage and Visualization

- Any netdata node can also save data to a backend
- Many backends supported
 - Graphite, InfluxDB, KairosDB, Blueflood, ElasticSearch
 - OpenTSDB, InfluxDB, KairosDB, MongoDB
 - json document DBs, Prometheus, TimescaleDB
- Only one backend per node at a time but you can chain together multiple netdata nodes with different backends
- Netdata can down sample the data to reduce storage at the cost of fidelity
- At UTD, we're moving forward with netdata + graphite + grafana

Netdata + Graphite and Grafana

- Netdata runs on every cluster node and forwards those results to a master netdata server
- No local database is used on compute nodes
- A netdata master server will down sample data to 10 second intervals
- Store all cluster node netdata data in a Graphite database
- Graphite web front end can then be used to make multi-host “charts on demand” to assist with debugging
- Grafana will read from the Graphite database (and other data stores)
- Grafana dashboards can be written to integrate Netdata and Queue system logs/accounting to create new capabilities
- Goal is to have a historical job dashboard where clicking on a ‘jobid’ gives you access to all of the netdata metrics during the lifetime of your job and a platform to analyze the data

Netdata worth considering for OHPC?

Thoughts? Questions?

Additional component deprecation

- We discussed deprecation of mpiP/ocr last time
- Would like to discuss doing the same for one additional package for which execution issues have cropped up in CI:
 - mrsh
 - provides functionality like ssh/scp, but uses **munge** for underlying authentication
 - not part of critical path in current installation recipes (optional)
 - last release was April 2016
 - relies on tcp_wrappers library which was deprecated in rhel7 and not available in rhel8 (<https://access.redhat.com/solutions/3906701>)

Component Review #9 Results

Component Name	# of Reviewers	# of accepts	# of rejects	Avg. Priority
Libfabric	6	4	2	7.0
OpenFlightHPC User Suite	7	1	5	2.2
LLNL Magpie	5	5	0	6.3

- Discussion on libfabric:
 - those who voted to reject were doing so in favor of using the distro provided version (1.8.x available in both Leap and CentOS8)
 - i'm also inclined to do that but update MPI builds accordingly and document/test installation in recipes

2.0 Build Status

- Have enabled a good number of packages in 2.0 Factory
 - introduction of arm1 compiler variant means 10 compiler/MPI permutations for MPI-based builds

Administrative Tools

Package	Built?	Notes
conman	✓	
docs	✓	
examples	✓	
ganglia		
genders	✓	
lmod-defaults	✓	added (3/10/20)
lofs	✓	
mrsh	✓	fixed (3/8/20)
nagios		
nagios-plugins		
ndoutils		
nhc	✓	added (3/10/20)
npe		
pdsh	✓	
prun	✓	
test-suite	✓	

Compiler Families

Package	Built?	Notes
gcc9	✓	
intel-compatibility	✓	
arm-compatibility	✓	
llvm		

Development Tools

Package	Built?	Notes
easybuild	✓	
autoconf	✓	
automake	✓	
cmake	✓	
hwloc	✓	
libtool	✓	
python-mpi4py	✘	3/24/20 – arm1 failures
python-numpy	✘	3/24/20 – arm1 failures
python-scipy	✘	3/19/20 – arm1 failures
spack	✓	[still targets root usage though]
valgrind	✓	

MPI Families

Package	Built?	Notes
impi-compatibility	✓	no longer requires install of intel compat. package
mpich	✓	
mvapich2	✓	
openmpi4	✓	

Serial Libs

Package	Built?	Notes
R	✓	
GSL	✓	
metis	✓	
openblas	✓	
plasma	✓	arm1 enabled (4/21/20)
superlu	✓	arm1 fixed (3/9/20)

✘ = partial builds success

✓ = all builds complete

X = all builds fail

2.0 Build Status (cont.)

Parallel Libraries

Package	Built?	Notes
boost	✘	intel error
fftw	✓	build fixed
hypr	✘	intel/arm1 blas issue
mfem	✘	needs deps
mumps	✘	arm1/impi issue
opencoarrays	✘	leap/openmpi4 failures
petsc	✘	intel/arm1 issue
scotch	✓	arm1 fixed (3/9/20)
scalapack	✓	arm1 fixed (3/9/20)
slepc	✘	waiting on petsc deps
superlu_dist	✘	2 Leap 15.1 failures
trilinos	✘	Intel/ARM failures (4/8/20)

Resource Management

Package	Built?	Notes
PBS Pro	✘	need libical-devel on aarch/centos
pmix	✓	
slurm	✓	

Performance Tools

Package	Built?	Notes
dimemas	✘	need some boost builds, suse failures
extrae	✘	arm1 failures
geopm	✘	intel failures (4/18/20)
imb	✘	updated to v2019.6, all built except for arm1 failures (4/16/20)
likwid	✓	(4/7/20)
msr-safe	✓	(4/20/20)
mpiP		
omb	✓	3/4/20
papi	✓	
paraver	✓	added 3/5/20
pdtoolkit	✘	intel/leap failure (3/5/20)
scalasca	✘	need scorep deps (3/5/20)
scorep	✘	variety of issues (3/5/20)
tau	✘	need pdtoolkit on intel/leap and arm1 (3/24/20)

Runtimes

Package	Built?	Notes
charliecloud	✓	3/23/20
eer		
singularity	✓	updated to use newer go version (4/6/20)

✘ = partial builds success

✓ = all builds complete

X = all builds fail

2.0 Build Status (cont.)

I/O Libraries

Package	Built?	Notes
adios	✘	intel failures (and deps)
hdf5	✘	Leap/arm1 failures
netcdf-cxx	✘	Leap/arm1 failures
netcdf-fortran	✘	Leap/arm1 failures
phdf5	✘	arm1 failures
pnetcdf	✓	
sionlib	✘	arm1 failures

Provisioning

Package	Built?	Notes
warewulf	✓	starting to run in CI (3/10/20)

✘ = partial builds success

✓ = all builds complete

✘ = all builds fail


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

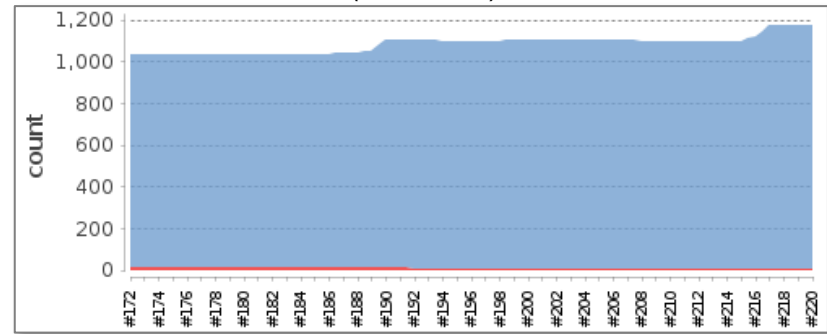
Base OS	aarch64	x86_64	noarch
CentOS 8	165	316	29
Leap 15	162	302	29

2.0 (cont.) – CI updates

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



centos8.1/x86/WW/slurm (4/22/2020)



- From last time (3/10):
 - ~~610 user level tests passing~~
- As of 3/25:
 - ~~910 user level tests passing~~
- As of this morning (4/08):
 - ~~982 user level tests passing~~
- As of this morning (4/22):
 - 1,132 user-level tests passing

2.0 (cont.) – CI updates

add description

S	Categorized - Job	Last Success	Last Failure	Last Duration	Test Result
2.0	All +				
–	.. » [aarch64]	18 days - #3	4 hr 9 min - #1	7 min 57 sec	N/A
🚫	(2.0) - (centos8.1,aarch64) (warewulf+slurm) (fabric=eth)	19 days - #47	13 hr - #54	29 min	5 of 1,103 failed (-9)
🚫	(2.0) - (leap15.1,aarch64) (warewulf+pbspro) (fabric=eth)	N/A	4 hr 9 min - #1	1 hr 1 min	14 of 800 failed (+14)
🔄	(2.0) - (leap15.1,aarch64) (warewulf+slurm) (fabric=eth)	18 days - #3	5 days 20 hr - #10	7 min 57 sec	13 of 808 failed (-13)
–	.. » [x86_64] - CentOS 8	28 days - #3	1 hr 31 min - #222	3 min 14 sec	N/A
🚫	(2.0) - (centos8.1,x86_64) (warewulf+slurm) (fabric=eth)	1 mo 13 days - #27	1 hr 31 min - #222	3 min 17 sec	3 of 1,174 failed (-4)
🚫	(2.0) - (centos8.1,x86_64) (warewulf+slurm) (fabric=ib) + psxe	28 days - #3	3 hr 49 min - #50	3 min 14 sec	5 of 1,210 failed (-14)
–	.. » [x86_64] - Leap15	28 days - #9	59 min - #295	3 min 37 sec	N/A
🚫	(2.0) - (leap15.1,x86_64) (warewulf+pbspro) (fabric=eth)	N/A	3 hr 39 min - #5	38 min	12 of 841 failed (±0)
🚫	(2.0) - (leap15.1,x86_64) (warewulf+slurm) (fabric=eth)	1 mo 6 days - #6	59 min - #295	3 min 48 sec	12 of 1,014 failed (±0)
🚫	(2.0) - (leap15.1,x86_64) (warewulf+slurm) (fabric=ib)	28 days - #9	28 days - #15	3 min 37 sec	35 of 823 failed (±0)

v1.3.9 totals

1,166 tests
1,156 tests
1,166 tests

1,237 tests
2,510 tests

1,203 tests
1,549 tests

- Enabled some pbspro jobs on Leap
- We are unfortunately down 1 of our Jenkins aarch64 nodes in the UK
 - host throwing h/w error during bios post

Moving forward with a preview release

- Reasonably content with ethernet only CI tests
- Propose to:
 - finalize MPI builds with libfabric (and UCX where applicable)
 - agree/enable on default build flags
 - achieve similar/reasonable test coverage on InfiniBand runs
 - work thru an updated release and test process as the public packages will be pushed to repos.openhpc.community
 - push out a subset of recipes (but coverage for both Leap and CentOS)
 - Warewulf only
 - No arm1 compiler coverage initially
 - document known issues and packages currently unavailable