

PEARC22

OpenHPC Community BoF

Christopher Simmons, Derek Simmel, Jason Wells and Eric Coulter

OpenHPC Technical Steering Committee (TSC) Members

July 13, 2022 • Boston, Massachusetts

Live Question Tool
<https://utd.link/ohpc>



openHPC


<https://openhpc.community>



Outline

- Part I: Presentation (~30 min)
 - Community members and growth snapshots
 - CentOS8 EOL announcement and resulting activities
 - Previous, latest, and upcoming releases

- Part II: Open Forum (~30 min)



Current Project Members



OpenHPC TSC – Individual Members

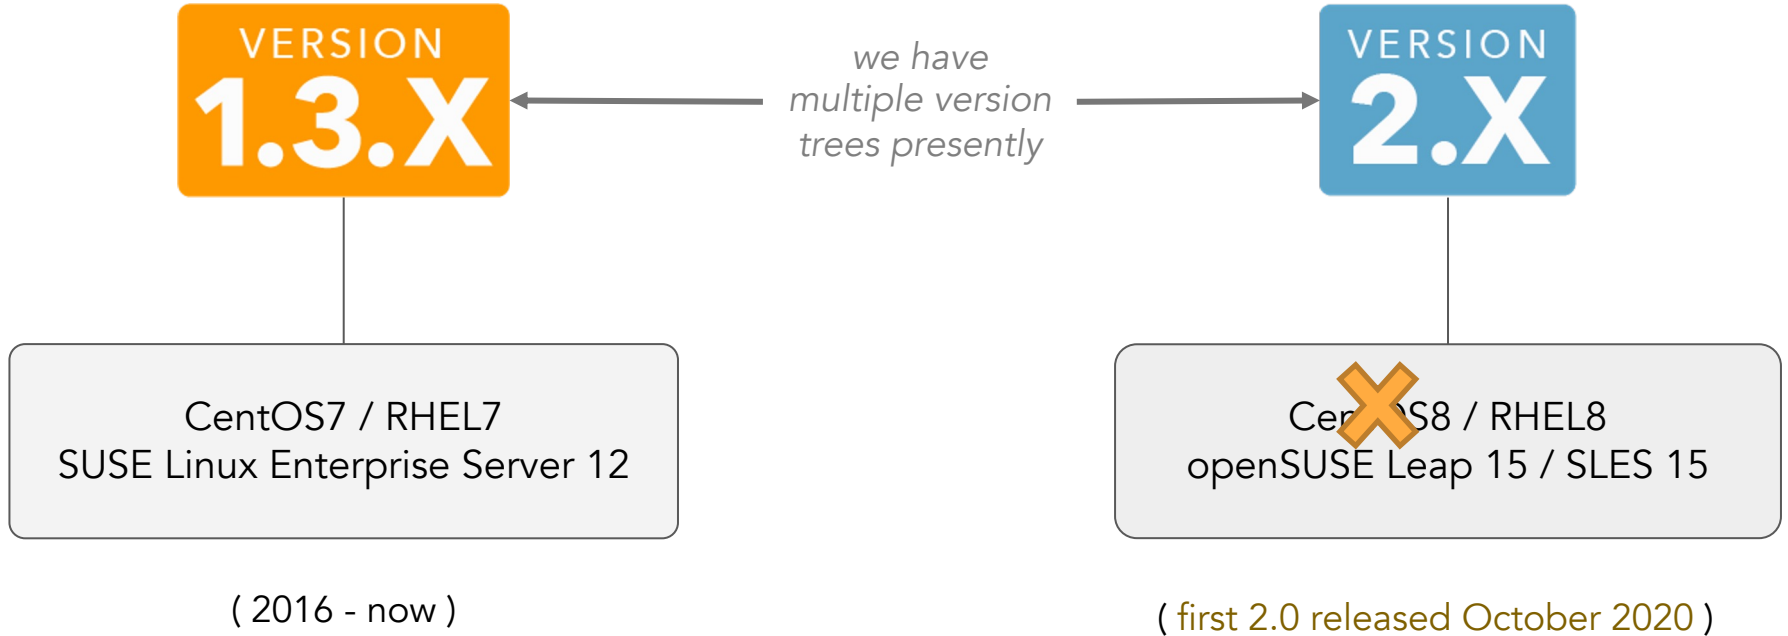
- Reese Baird, SpaceX (Maintainer)
- David Brayford, Intel (Maintainer)
- Eric Coulter, Georgia Institute of Technology (End-User/Site Representative)
- Chris Downing, Amazon Web Services (Maintainer)
- [Richard Henwood, Arm \(Maintainer\)](#)
- Caetano Melone, Stanford (Maintainer)
- [Amarpal Kapoor, Intel \(Maintainer\)](#)
- Michael Karo, Altair (Component Development Representative)
- Cyrus Proctor, Spectral Quantum Tech. (Maintainer)
- Adrian Reber, Red Hat (Maintainer)
- Karl W. Schulz, UT Austin (Testing Coordinator)
- Jeremy Siadal, Intel (Maintainer)
- Derek Simmel, Pittsburgh Supercomputing Center (End-User/Site Representative)
- Christopher Simmons, UT Dallas (Maintainer, Interim Project Lead)
- [Caesar Stoica, Lenovo \(Maintainer\)](#)
- Nirmala Sundararajan, Dell (Maintainer)
- [Jason Wells, Harvard \(End-User/Site Representative\)](#)

New members for 2021-2022

Interested in participating?
TSC nominations done in June yearly

<https://github.com/openhpc/ohpc/wiki/Governance-Overview>

OpenHPC: multi distro support



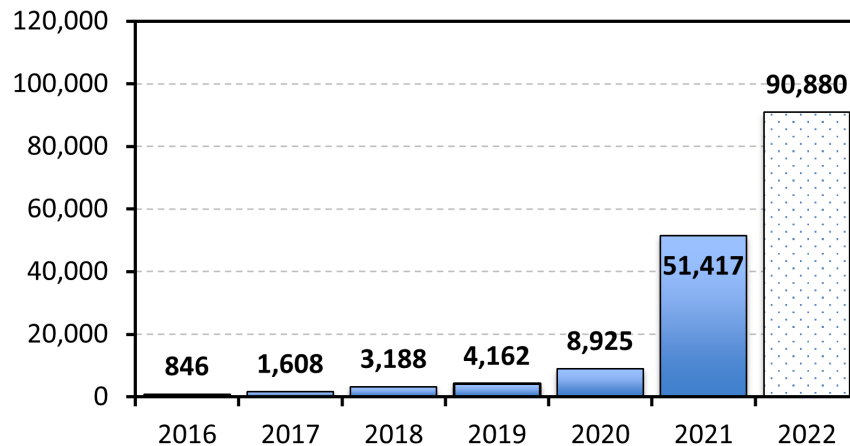
Community Growth Snapshots

Project Adoption Growth

Repo Server(s) Access: Unique Monthly Visitors



Average # of visitors/month



- Summary of access/download growth since initial release at SC'15
- Plots highlight number of unique visitors/month to the OpenHPC build server/repo(s)
- Significant uptake with 2.x releases and just short of 104K in June '22!
- 45 TB per month and growing in 2022; Up from 26 TB per month in 2021

CentOS 8 Business

CentOS8 Announcement - Dec 2020

CentOS Project shifts focus to CentOS Stream

Tuesday, 8, December 2020 Rich Bowen Uncategorized 709 Comments

The future of the CentOS Project is CentOS Stream, and over the next year we'll be shifting focus from CentOS Linux, the rebuild of Red Hat Enterprise Linux (RHEL), to CentOS Stream, which tracks just *ahead* of a current RHEL release. **CentOS Linux 8, as a rebuild of RHEL 8, will end at the end of 2021** CentOS Stream continues after that date, serving as the upstream (development) branch of Red Hat Enterprise Linux.

Meanwhile, we understand many of you are deeply invested in CentOS Linux 7, and we'll continue to produce that version through the remainder of the [RHEL 7 life cycle](#).

CentOS Stream will also be the centerpiece of a major shift in collaboration among the CentOS Special Interest Groups (SIGs). This ensures SIGs are developing and testing against what becomes the next version of RHEL. This also provides SIGs a clear single goal, rather than having to build and test for two releases. It gives the CentOS contributor community a great deal of influence in the future of RHEL. And it removes confusion around what "CentOS" means in the Linux distribution ecosystem.

When CentOS Linux 8 (the rebuild of RHEL8) ends, your best option will be to migrate to CentOS Stream 8, which is a small delta from CentOS Linux 8, and has regular updates like traditional CentOS Linux releases. If you are using CentOS Linux 8 in a production environment, and are concerned that CentOS Stream will not meet your needs, we encourage you to contact Red Hat about options.

- Like most folks, we were caught off guard by this announcement to discontinue CentOS8 on Dec. 31, 2021
- Through 2021, CentOS has been the preferred distro in use by OpenHPC users
- Initially considered multiple alternative options:
 - CentOS8 Stream
 - RHEL8 proper
 - binary-compatible RHEL8 clones
 - solicited community feedback to help guide our path...

Some Community Polling Results Regarding CentOS8



OpenHPC Community BoF

Chat Questions **Polls**

If you are currently using CentOS 8, what do you plan to use as a replacement going forward?

2 days ago

- 9% CentOS 8 Stream
- 11% RHEL 8
- 13% AlmaLinux
- 7% openSUSE Leap
- 4% Oracle Linux
- 50% Rocky Linux**
- 7% Other

If you are currently using CentOS 8, what do you plan to use as a replacement going forward?

Vote submitted. [Change poll answer](#)

Poll Results:

Percentage	OS
14%	CentOS 8 Stream
10%	RHEL 8
0%	AlmaLinux
7%	OpenSUSE Leap
3%	Oracle Linux
55%	Rocky Linux
10%	Other

PEARC21

If using a RHEL variant, which one(s):

Type: Multiple Answer

Variant	Percentage
RHEL	44%
Rocky	33%
AlmaLinux	0%
Oracle Linux	11%
Other	55%

Community Plans for supporting RHEL8 Variants

- Based on community feedback and additional infrastructure testing, we have pivoted as follows for RHEL-based usage (starting with the v2.4 release):

- **Build:**

- ohpc packages are built directly against RHEL proper (using community entitlements)
- OBS-based build system infrastructure updated to support this change

- **Test:**

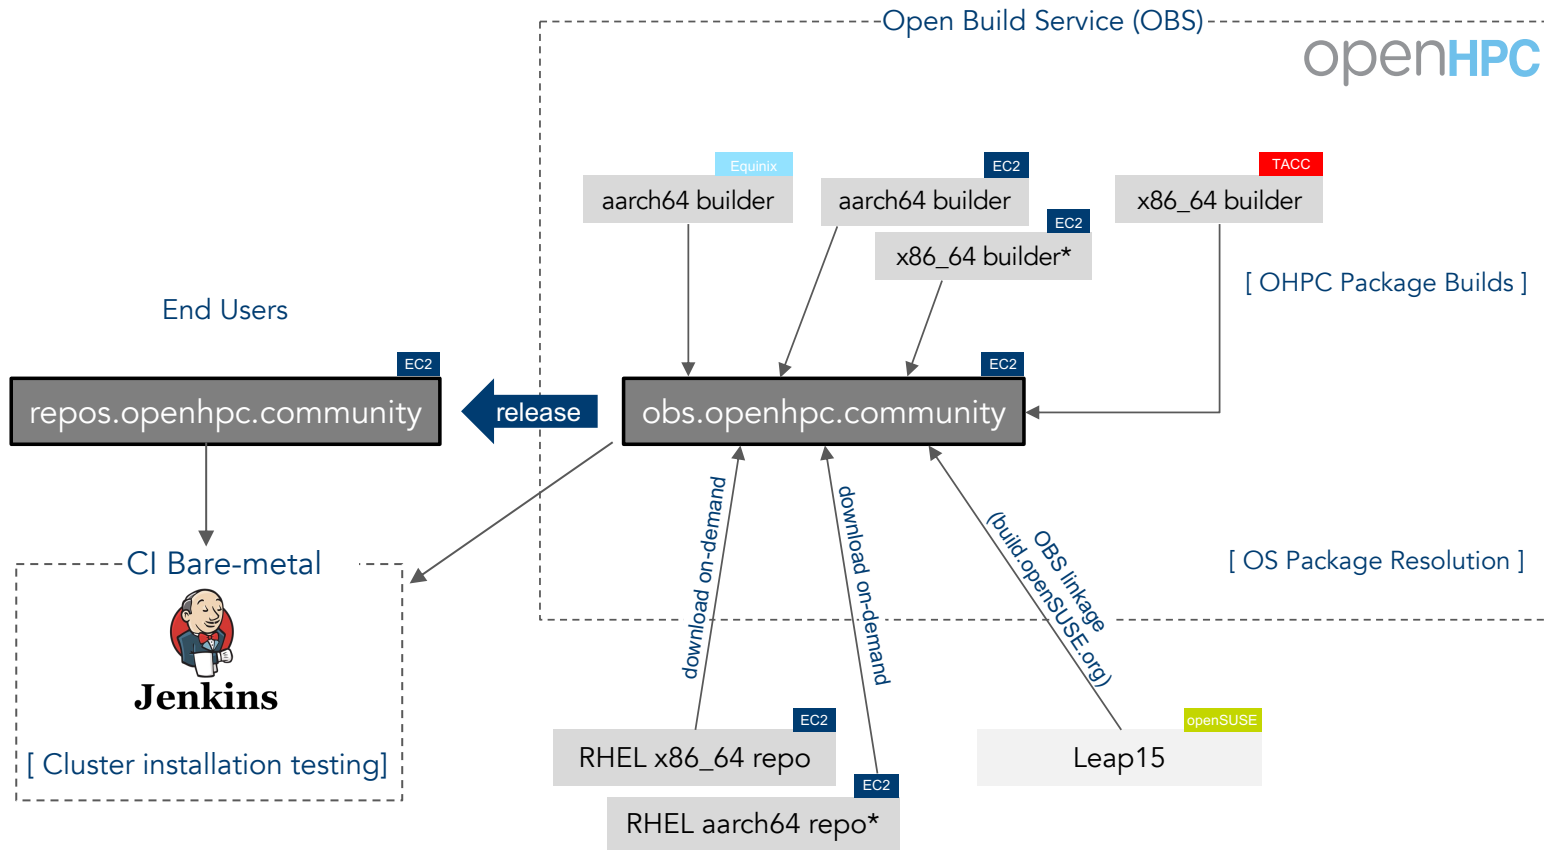
- example installation recipes for RHEL updated to use a binary compatible clone
- continuous integration (CI) infrastructure also updated to leverage alternative RHEL clone
- based on initial community feedback, we chose **Rocky8** as the basis for example recipes
 - *Other RHEL8 binary clones should also be compatible*



- Note: we continue to support OpenSUSE Leap 15.x as well



Updated OpenHPC Build/Delivery Architecture (2.4+)



Release Updates

OpenHPC v2.4 - S/W components

Functional Areas	Components
	components available 79
Base OS	RHEL 8.5 , OpenSUSE Leap 15.3
Architecture	x86_64, aarch64
Administrative Tools	Conman, Lmod, LosF, Nagios, NHC, pdsh, pdsh-mod-slurm, prun, EasyBuild, ClusterShell, Genders, Shine, Spack, test-suite
Provisioning	Warewulf3
Resource Mgmt.	SLURM, Munge, OpenPBS, PMIx, Magpie
Runtimes	Charliecloud, Singularity
I/O Services	Lustre client (community version), BeeGFS client
Numerical/Scientific Libraries	Boost, GSL, FFTW, Hypre, Metis, MFEM, Mumps, OpenBLAS, OpenCoarrays, PETSc, PLASMA, Scalapack, Scotch, SLEPc, SuperLU, SuperLU_Dist, Trilinos
I/O Libraries	HDF5 (pHDF5), NetCDF/pNetCDF (including C++ and Fortran interfaces), Adios
Compiler Families	GNU (gcc, g++, gfortran), Intel oneAPI Toolkit* , ARM Alinea Studio*
Transport Layers	Libfabric, UCX
MPI Families	MVAPICH2, OpenMPI, MPICH, Intel oneAPI HPC Toolkit*
Development Tools	Autotools, cmake, hwloc, mpi4py, R, SciPy/NumPy, Valgrind
Performance Tools	Dimemas, Extrae, GeoPM, IMB, Likwid, msr-safe, OSU Micro-Benchmarks, PAPI, Paraver, pdtoolkit, Scalasca, ScoreP, SIONLib, TAU

Additional dependencies not provided by BaseOS or community repos are also included

v2.4: Installation recipes available

[Key takeaway]

- In addition to being a package repository, OpenHPC provides validated recipes for bare-metal system installs
- Recipes organized by OS, architecture, and key administrative components
- **2.4** release includes **8** different recipes:
 - CentOS8 -> Rocky8
- the **docs-ohpc** RPM installs these recipes (along with shell scripts encapsulating all commands)

Recipes: /opt/ohpc/pub/doc/recipes

Rocky8

Leap 15

Warewulf



PBS

SLURM

x86_64:

- [Install_guide-Rocky8-Warewulf-OpenPBS-2.4-x86_64.pdf](#)
- [Install_guide-Rocky8-Warewulf-SLURM-2.4-x86_64.pdf](#)
- [Install_guide-Leap_15-Warewulf-OpenPBS-2.4-x86_64.pdf](#)
- [Install_guide-Leap_15-Warewulf-SLURM-2.4-x86_64.pdf](#)

aarch64:

- [Install_guide-Rocky8-Warewulf-PBSPro-2.3-aarch64.pdf](#)
- [Install_guide-Rocky8-Warewulf-SLURM-2.3-aarch64.pdf](#)
- [Install_guide-Leap_15-Warewulf-PBSPro-2.3-aarch64.pdf](#)
- [Install_guide-Leap_15-Warewulf-SLURM-2.3-aarch64.pdf](#)

can use these guides as starting point for bare-metal installs

v2.4: Additional Update Highlights

- Refactored our hwloc packaging:
 - now installs into fixed path
 - OpenPBS and SLURM builds updated to use ohpc variant of hwloc
- Compiler update: gcc 9.3 -> gcc 9.4
- RMS updates (both openPBS and SLURM)
- MPI and transport updates:
 - openmpi: 4.0.5 -> 4.1.1
 - mpich: 3.3.2 -> 3.4.2
 - mvapich2: 2.3.4 -> 2.3.6
 - libfabric: 1.11.2 -> 1.13.0
 - ucx 1.9.0 -> 1.11.2
- Updated compatibility package for use with ARM and Intel vendor compilers
- Variety of other component version updates (44% of packages from v2.3)
 - see release notes for specifics
- Caveats:
 - parallel file-system clients not yet available for BeeGFS or Lustre

v2.4 and v2.5: Parallel File System Client Status

- Lustre
 - Upstream client would not build on kernels at 2.4 release time
 - At 2.5 release time, upstream did build however kernel panics in our CI environment
- BeeGFS
 - Initially broken due to symbol changes in kernel related to InfiniBand
 - As of this morning, works with RDMA over IB but some known issues related to RDMA over RoCE
 - https://doc.beegfs.io/7.2.6/release_notes.html

```
# Add BeeGFS client software and dependencies to master host
[sms]# wget -P /etc/yum.repos.d https://www.beegfs.io/release/beegfs-7.2.1/dists/beegfs-rhel8.repo
[sms]# yum -y install kernel-devel gcc elfutils-libelf-devel
[sms]# yum -y install beegfs-client beegfs-helperd beegfs-utils

# Enable OFED support in client
[sms]# perl -pi -e "s/~buildArgs=-j8/buildArgs=-j8 BEEGFS_OPENTK_IBVERBS=1/" \
    /etc/beegfs/beegfs-client-autobuild.conf


# Define client's management host
[sms]# /opt/beegfs/sbin/beegfs-setup-client -m ${sysmgmt_host}

[sms]# systemctl start beegfs-helperd
# Build kernel and mount file system
[sms]# systemctl start beegfs-client
```

Change to: **beegfs_7.2.6**

2.4 updates (cont.)

- Intel repackaged the previous PSXE compiler variants within the oneAPI Toolkit (also introduced new clang-based variants)
- Have introduced updated compatibility packages that enables usage with oneAPI classic compiler variants: **icc**, **icpc**, **ifort** (we tested with v2021.4.0)
- Usage is similar to previous releases, but made easier now by the fact that the compiler can be installed directly from an online repository
 - **new convenience package introduced which will setup the oneAPI repository locally: intel-oneapi-toolkit-release-ohpc**



```
# Enable Intel oneAPI and install OpenHPC compatibility packages
[sms]# yum -y install intel-oneapi-toolkit-release-ohpc
[sms]# yum -y install intel-compilers-devel-ohpc
[sms]# yum -y install intel-mpi-devel-ohpc
```

```
[sms]# rpm -ql intel-oneapi-toolkit-release-ohpc
/etc/yum.repos.d/oneAPI.repo
```

2.4 updates (cont.)

- Note: the newer compatibility package relies on a utility shipped with oneAPI packages to generate modulefiles for locally installed versions
 - e.g. `/opt/intel/oneapi/modulefiles-setup.sh`
- You will thus see more module dependencies that get loaded
- Additional note: if installing oneAPI compilers via package managers, these will land in `/opt/intel`
 - need to make this path available on computes in example installation recipes
 - 2.4 variants updated to call out sharing over NFS directly

```
$ module swap gnu9 intel
Loading compiler version 2021.4.0
Loading tbb version 2021.4.0
Loading compiler-rt version 2021.4.0
Loading debugger version 10.2.4
Loading mkl version 2021.4.0
```

2.4 updates (cont.)

- Compatibility package for Arm compiler has also been updated to work with newer release
- In this case, need to download install package separately and install the

```
# Install OpenHPC compatibility packages
[sms]# zypper install arm1-compilers-devel-ohpc
```

```
$ module swap gnu9 arm1
$ which armclang
/opt/ohpc/pub/arm/arm-linux-compiler/bin/armclang
```

```
$ module list

Currently Loaded Modules:
 1) autotools          5) clang-autocomplete/21.1   9) ucx/1.11.2
 2) prun/2.2           6) arm21/21.1                10) libfabric/1.13.0
 3) ohpc               7) arm1/compat
 4) binutils/10.2.0   8) hwloc/2.5.0
```

<https://developer.arm.com/tools-and-software/server-and-hpc/downloads/arm-allinea-studio/openhpc>

arm Developer

Home | Tools and Software | Server and HPC | Downloads | Arm Allinea Studio | OpenHPC

OpenHPC

Arm Compiler for Linux for OpenHPC

Here you can download the variant of Arm Compiler for Linux suitable for OpenHPC users.

If you do not use a system based on OpenHPC, return to the Arm Allinea Studio downloads page to download the version suitable for your system:

[Download for other systems](#)

Version 2.5 released (29 May 2022)

Version 2.5 (29 May 2022)

[General Updates]

- * updated SLURM version to address CVE-2022-29500

[Component Version Changes]

- * docs-ohpc (2.4.0 -> 2.5.0)
- * slurm-contribs-ohpc (20.11.8 -> 20.11.9)
- * slurm-devel-ohpc (20.11.8 -> 20.11.9)
- * slurm-example-configs-ohpc (20.11.8 -> 20.11.9)
- * slurm-libpmi-ohpc (20.11.8 -> 20.11.9)
- * slurm-ohpc (20.11.8 -> 20.11.9)
- * slurm-openlava-ohpc (20.11.8 -> 20.11.9)
- * slurm-pam_slurm-ohpc (20.11.8 -> 20.11.9)
- * slurm-perlapi-ohpc (20.11.8 -> 20.11.9)
- * slurm-slurmctld-ohpc (20.11.8 -> 20.11.9)
- * slurm-slurmdbd-ohpc (20.11.8 -> 20.11.9)
- * slurm-slurmd-ohpc (20.11.8 -> 20.11.9)
- * slurm-sview-ohpc (20.11.8 -> 20.11.9)
- * slurm-torque-ohpc (20.11.8 -> 20.11.9)

Mentorship Program

- Mentorship program for 2022
 - 4 students selected out of ~25 applicants
 - Project 1: Replacement for Ganglia / new monitoring stack
 - Project 2: Better integration of OHPC software with EasyBuild
 - Project 3: Adding Elastic Fabric Adaptor support to OHPC
 - Project 4: TBD
- Mentorships for 2023 will open in May 2023
 - Apply via the LFX Mentorship portal
 - Join the OHPC Users email list to be notified when open

Additional Future Items

- one final 1.3.x release as 1.3.10 in Q3 with support for RHEL/CentOS 7.9
 - updating resource managers
 - test against newer minor distro versions
- 2.6 items:
 - addition of new component: hpc-workspace
 - Warewulf4
 - job launch support with PMIx
 - component packaging for use with Arm Compiler
 - ???? <your input here>
- Initial “cloud” recipes in Q4 with support for AWS and OpenStack; more to come

Open Discussion

