

CCF HPC China 2023

OpenHPC

Dr. Christopher S. Simmons
csim@mit.edu

OpenHPC Project Lead

Aug 23, 2023



Current Project Members

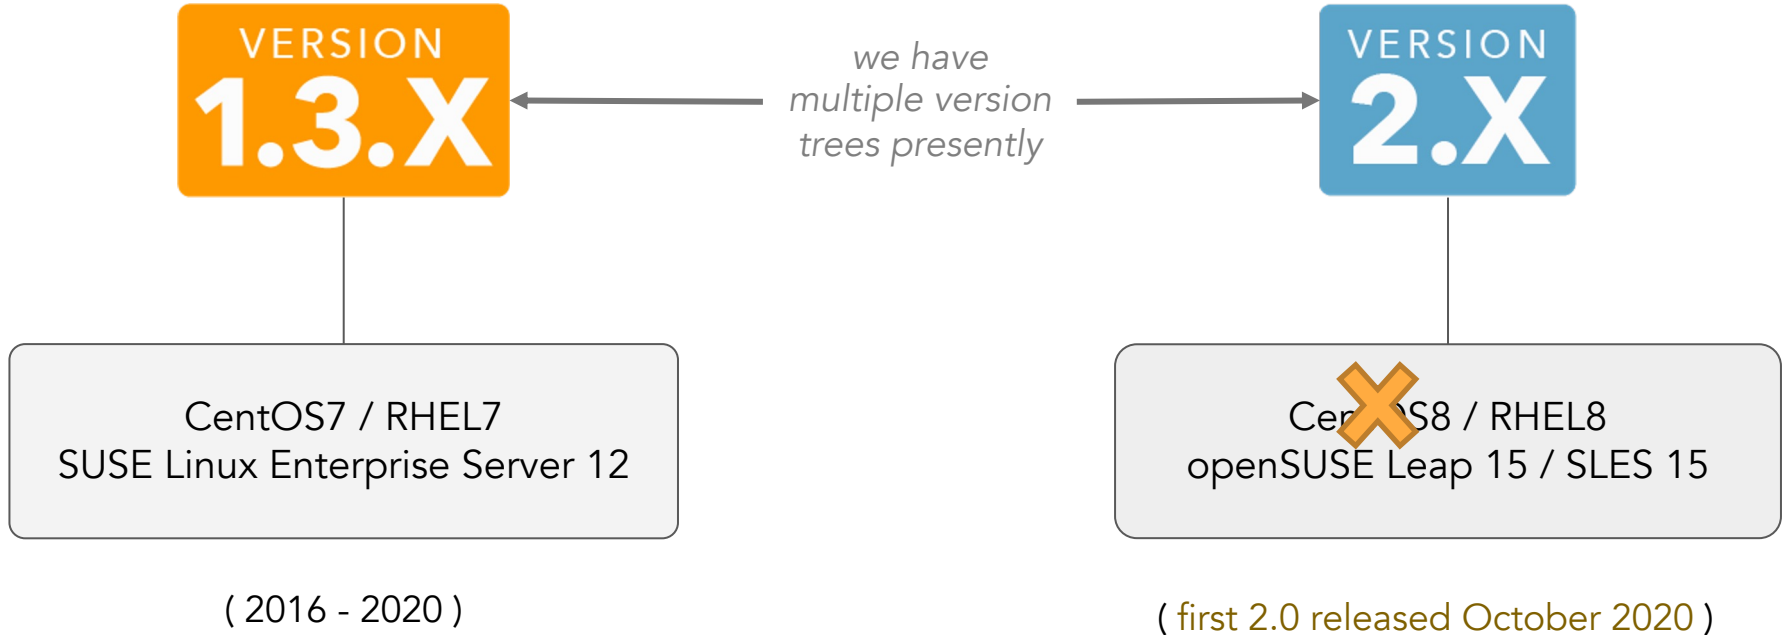


OpenHPC TSC – Individual Members

- Reese Baird, SpaceX (Maintainer)
- David Brayford, HPE (Maintainer)
- Alfred Egger, University of Salzburg (End-User/Site Representative)
- [Martin Grigorov, Huawei \(Maintainer\)](#)
- Sol Jerome, UMass Amherst (End-User/Site Representative)
- Michael Karo, Altair (Upstream Component Development Representative)
- [Forrest Ling, easyHPC \(End-User/Site Representative\)](#)
- Adrian Reber, Red Hat (Maintainer)
- Phil Regier, SUSE (Maintainer)
- Karl W. Schulz, AMD Research (Testing Coordinator)
- Jeremey Siadal, Intel (Maintainer)
- Derek Simmel, PSC (End-User/Site Representative)
- [Christopher Simmons, MGHPCC \(Maintainer, Project Lead\)](#)
- Caeser Stoica, Lenovo (Maintainer)
- Jason Wells, Harvard (End-User/Site Representative)

Interested in participating?
TSC nominations done in June yearly


OpenHPC: multi distro support




Community Growth Snapshots

Project Adoption Growth

Average # of visitors/month



Average download/month (TB)



- Summary of access/download growth since initial release at SC'15
- Plots of unique visitors/month and TB/month to the OpenHPC build server/repo(s)
- Significant uptake with 2.x releases and over 150K this year
- 110+ TB per month and growing in 2024; Up from <50 TB per month in 2022

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:

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

PEARC21

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

Type: Multiple Answer

| RHEL 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

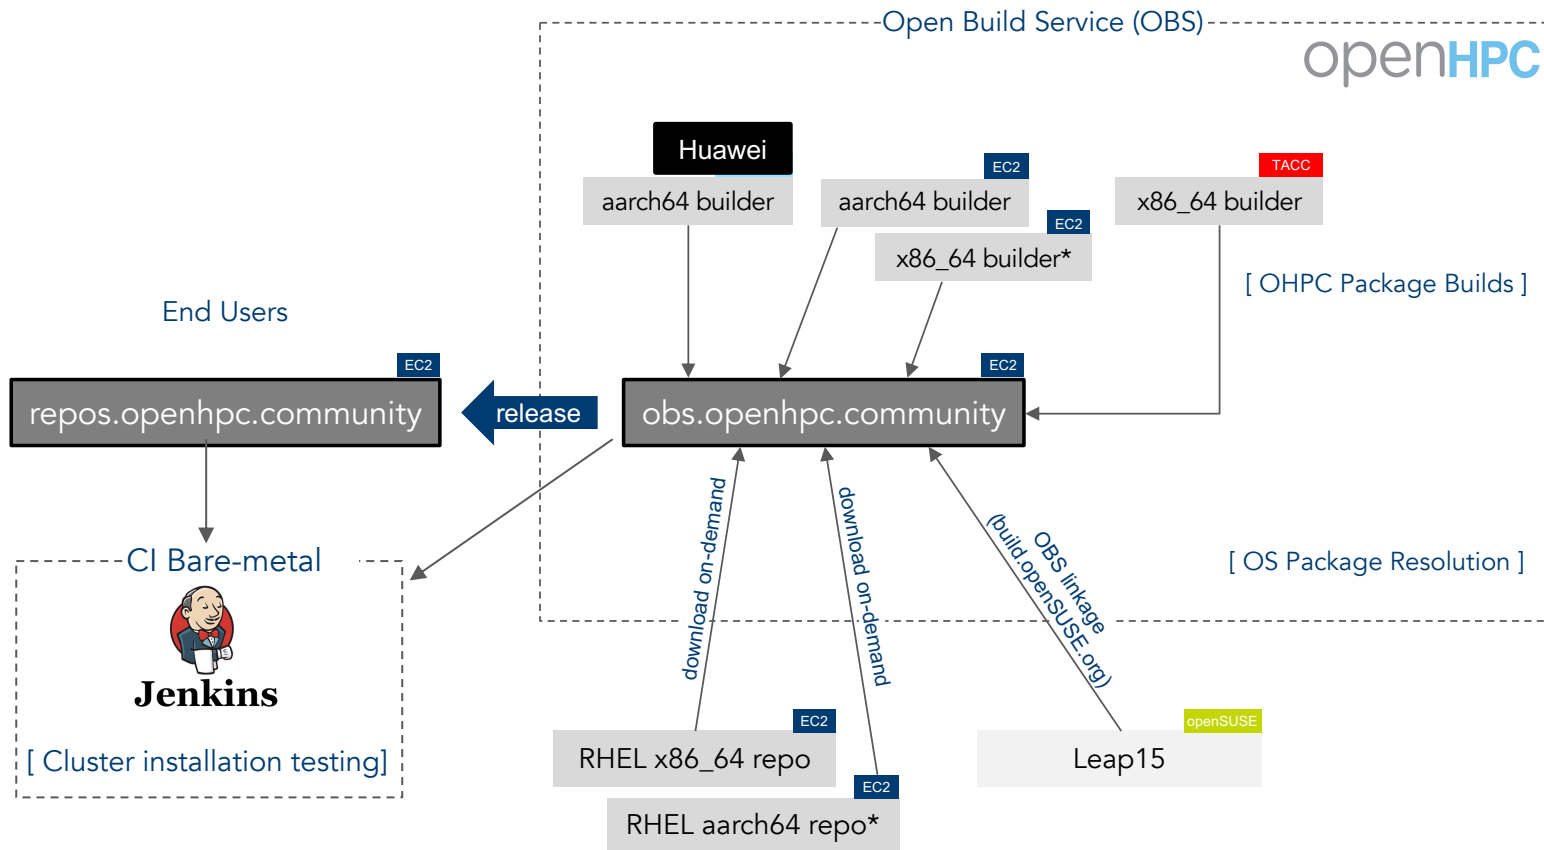
- **Test:**

- 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  **Rocky Linux™**
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 v2.6 - S/W components

| Functional Areas | Components |
|--------------------------------|---|
| | components available 79 |
| Base OS | RHEL 8.6 , 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, Warewulf4 |
| Resource Mgmt. | SLURM, Munge, OpenPBS, 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

v2.6: 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
- **2.6** 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-2.6-x86_64.pdf](#)
- [Install_guide-Rocky8-Warewulf-SLURM-2.6-x86_64.pdf](#)
- [Install_guide-Leap_15-Warewulf-OpenPBS-2.6-x86_64.pdf](#)
- [Install_guide-Leap_15-Warewulf-SLURM-2.6-x86_64.pdf](#)

aarch64:

- [Install_guide-Rocky8-Warewulf-PBSPro-2.6-aarch64.pdf](#)
- [Install_guide-Rocky8-Warewulf-SLURM-2.6-aarch64.pdf](#)
- [Install_guide-Leap_15-Warewulf-PBSPro-2.6-aarch64.pdf](#)
- [Install_guide-Leap_15-Warewulf-SLURM-2.6-aarch64.pdf](#)


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

v2.6: Update Highlights

- Updated target distro support to EL 8.6 and Leap 15.3
- New compiler variant (gnu12) introduced with this release
- SLURM build updated to include optional REST API interface package
- Updated SLURM recipes to include cgroups.conf
- Fix for Leap 15.3 provisioning with Warewulf (<https://github.com/openhpc/ohpc/issues/1602>)
- Tech preview build of Warewulf 4.x
- Introduced in 2.6.1, support for upstream singularity or apptainer in OHPC provided meta-packages
- v2.6: Known Issues
 - Several package builds with the latest Intel classic compiler exhibited problems with the OpenHPC test suite.
 - Affected packages are plasma, TAU, and superlu_dist
 - Users are advised to stick with the gcc (gnu12) compiler variant if using these packages

2.6 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`
- Usage is similar to previous releases, but made easier now by the fact that the compiler can be installed directly from an online repository
 - convenience package 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
```


2.6 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
 - 2.6 variants 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
```

2.6 updates (cont.)

<https://developer.arm.com/downloads/-/arm-compiler-for-linux>

- 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 compilers first locally

```
# 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) arm22/22.1              10) libfabric/1.13.0
 3) ohpc               7) arm1/compat
 4) binutils/10.2.0   8) hwloc/2.5.0
```

armDeveloper

IP Explorer Documentation Downloads Community Support

Developing on Arm Architecture and Processors Tools and Software

Home / Downloads / Arm Compiler for Linux

Download Arm Compiler for Linux

Download the latest version

Download the Arm Compiler for Linux (22.1) package for your OS:

Packages contain Arm C/C++/Fortran Compiler and Arm Performance Libraries.

| | | |
|--|--------------|---------|
| arm-compiler-for-linux_22.1_RHEL-7_aarch64.tar | RHEL 7 | 1.62 GB |
| arm-compiler-for-linux_22.1_RHEL-8_aarch64.tar | RHEL 8 | 1.62 GB |
| arm-compiler-for-linux_22.1_SLES-15_aarch64.tar | SLES 15 | 1.62 GB |
| arm-compiler-for-linux_22.1_Ubuntu-18.04_aarch64.tar | Ubuntu 18.04 | 1.61 GB |
| arm-compiler-for-linux_22.1_Ubuntu-20.04_aarch64.tar | Ubuntu 20.04 | 1.61 GB |

Standalone ArmPL for Ubuntu 20.04

| | |
|--|-----------|
| arm-performance-libraries_22.1_Ubuntu-20.04_gcc-10.2.tar | 325.53 MB |
| arm-performance-libraries_22.1_Ubuntu-20.04_gcc-11.2.tar | 324.95 MB |
| arm-performance-libraries_22.1_Ubuntu-20.04_gcc-9.3.tar | 253.89 MB |

Additional Future Items

- 2.7 items:
 - Updated to support RHEL 8.8
 - job launch support with PMIx reintroduced for OpenMPI
 - component packaging for use with Arm Compiler
 - Updated Intel-compiled OHPC packages for Intel oneAPI 2023 using classic compiler
 - ???? <your input here>
- 3.0 items:
 - Support for OpenEuler and testing against Kunpeng (Huawei's ARM architecture)
 - Support for RHEL 9.2 and SUSE Leap 15.5
 - New compiler variant (gnu13) to be introduced with this release
 - Component packaging for use with Arm Compiler
 - Support for Warewulf4, and PMIx for OpenMPI on SLURM

Additional Future Items for 3.0 cont.

- Default optimization for RHEL packages on 9.0 has been increased from x86-64 to x86-64-v2
- -march=x86-64-v2 added for GNU and Intel packages for OHP3.0 for EL 9.2
- Plans to make it easier for sites to redeploy with v3 and v4 optimization to get better performance on “newer” architectures

In the summer of 2020, AMD, Intel, Red Hat, and SUSE [collaborated](#) to define three x86-64 microarchitecture levels on top of the x86-64 baseline. The three microarchitectures group together CPU features roughly based on hardware release dates:

- **x86-64-v2** brings support (among other things) for vector instructions up to Streaming SIMD Extensions 4.2 (SSE4.2) and Supplemental Streaming SIMD Extensions 3 (SSSE3), the POPCNT instruction (useful for data analysis and bit-fiddling in some data structures), and CMPXCHG16B (a two-word compare-and-swap instruction useful for concurrent algorithms).
- **x86-64-v3** adds vector instructions up to AVX2, MOVBE (for big-endian data access), and additional bit-manipulation instructions.
- **x86-64-v4** includes vector instructions from some of the AVX-512 variants.

3.0 Status

- Testing has completed for:
 - All 3 distros
 - on both architectures
 - with both schedulers
 - with warewulf3
 - on ethernet and infiniband for x86
 - and ethernet for ARM
- OBS builds have been moved to the repo staging area

3.x

OpenHPC CI Infrastructure

Thanks to the Texas Advanced Computing Center (TACC) and Linaro for hosting support. Thanks also to Intel, Marvell, Cavium, Dell and Works on Arm for hardware support.

3.0

All

| S | W | Name ↓ | Last Success | Last Failure | Last Duration |
|---|---|---|------------------|------------------|---------------|
| ⊗ | ☁ | (3.0) - (leap15.4,x86_64) (warewulf+slurm) (fabric=eth) | N/A | 1 mo 24 days #5 | 4 min 1 sec |
| ✔ | ☁ | (3.0) - (leap15.5,aarch64) (warewulf+openpbs) (fabric=eth) | 2 hr 39 min #6 | 7 hr 57 min #4 | 2 hr 2 min |
| ✔ | ☁ | (3.0) - (leap15.5,aarch64) (warewulf+slurm) (fabric=eth) | 1 day 16 hr #26 | 1 day 19 hr #25 | 1 hr 46 min |
| ✔ | ☁ | (3.0) - (openeuler22.03,aarch64) (warewulf+openpbs) (fabric=eth) | 11 days #18 | 12 days #17 | 2 hr 16 min |
| ✔ | ☁ | (3.0) - (openeuler22.03,aarch64) (warewulf+slurm) (fabric=eth) | 11 days #82 | 18 days #81 | 2 hr 4 min |
| ✔ | ☁ | (3.0) - (rocky9.2,aarch64) (warewulf+openpbs) (fabric=eth) | 6 hr 15 min #8 | 8 days 21 hr #6 | 48 min |
| ✔ | ☁ | (3.0) - (rocky9.2,aarch64) (warewulf+slurm) (fabric=eth) | 9 hr 57 min #18 | 9 days 18 hr #16 | 1 hr 41 min |
| ✔ | ☁ | (3.0) - (rocky9.2,aarch64) (warewulf+slurm) (fabric=eth) + arm hpc compiler | 3 days 19 hr #5 | 3 days 23 hr #4 | 2 hr 27 min |
| ✔ | ☀ | (3.0) - (rocky9.2,x86_64) (warewulf+openpbs) (fabric=ib) - UEFI | 2 hr 21 min #660 | 1 day 2 hr #652 | 1 hr 2 min |
| ✔ | ☁ | (3.0) - (rocky9.2,x86_64) (warewulf+slurm) (fabric=ib) | 3 hr 45 min #789 | 12 hr #786 | 1 hr 7 min |

Icon: S M L

Icon legend

📡 Atom feed for all

📡 Atom feed for failures

📡 Atom feed for just latest builds

Testing 3.0 from the staging area

- Install the OHPC release from <http://repos.openhpc.community/.staging/OpenHPC/3/>
- Edit the repo file to point to the staging area instead of the release area
- Release expected by end of this month

```
[OpenHPC]
name=OpenHPC-3 - Base
baseurl=http://repos.openhpc.community/.staging/OpenHPC/3/EL_9
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-OpenHPC-3
```

```
[OpenHPC-updates]
name=OpenHPC-3 - Updates
baseurl=http://repos.openhpc.community/.staging/OpenHPC/3/updates/EL_9
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-OpenHPC-3
```

Thank you!