



OpenHPC Community BoF

Karl W. Schulz, David Brayford, Derek Simmel, Thomas Sterling,
Scott Suchyta, Nirmala Sundararajan

OpenHPC Technical Steering Committee (TSC) Members


November 15, 2017 • Denver, CO

Outline

Part I: Presentation (~30 min)

- What's New / Changes since last year
- Software in latest release
- Future work

Part II: Open Forum (~30 min)



Current Project Members



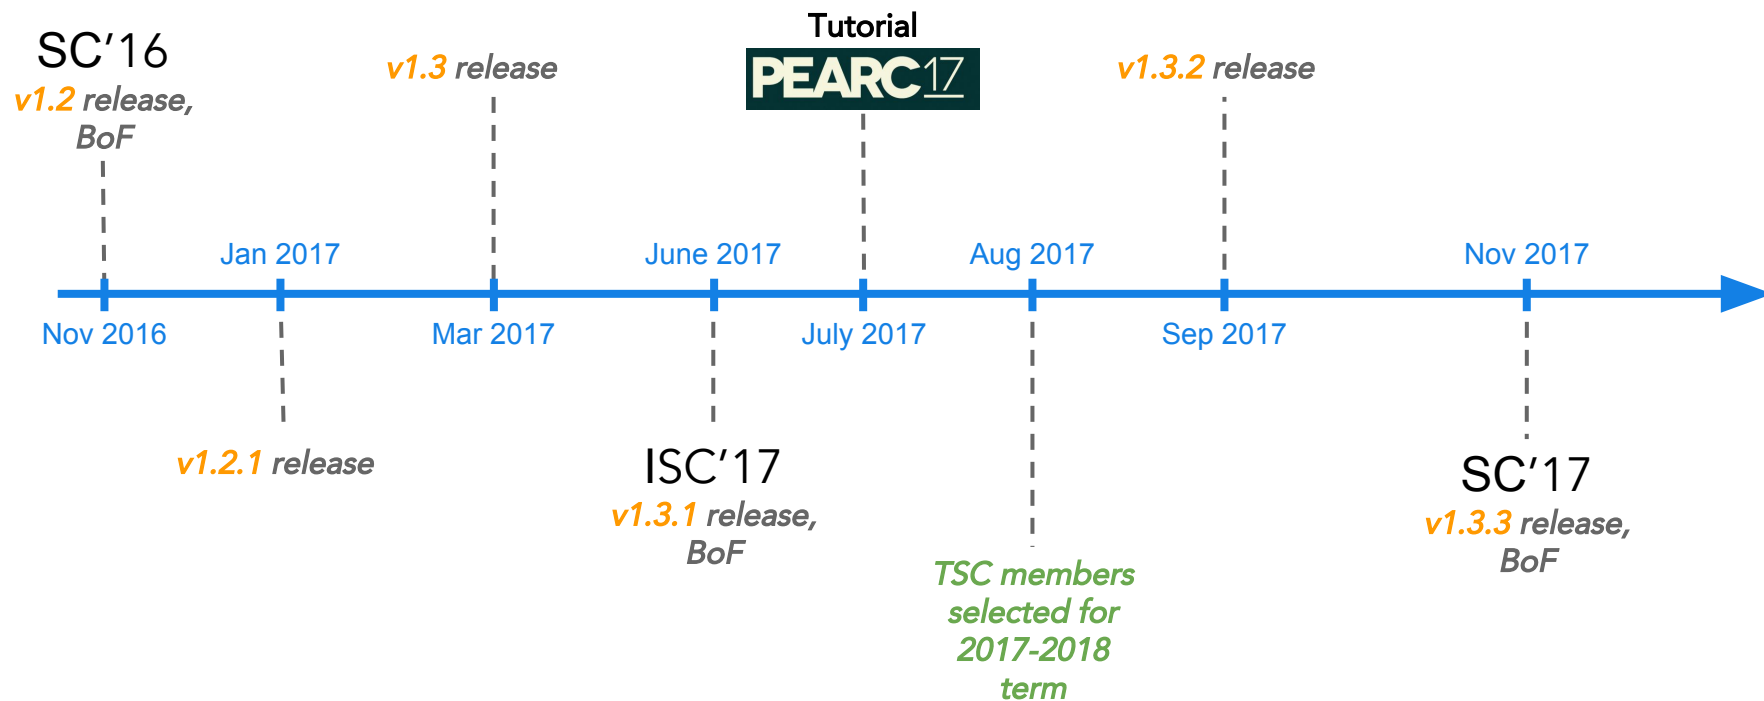
Indiana
University



University of
Cambridge


Mixture of academics, labs, and industry

High-Level Events/Releases Since Last Year



OpenHPC Technical Steering Committee (TSC)

Role Overview



Note: we completed election of TSC members for the 2017-2018 term in August:

- terms are for 1 year

OpenHPC TSC – Individual Members

- Reese Baird, Intel (Maintainer)
- David Brayford, LRZ (Maintainer)
- [Eric Coulter, Indiana University \(End-User/Site Representative\)](#)
- [Leonordo Fialho, ATOS \(Maintainer\)](#)
- Todd Gamblin, LLNL (Component Development Representative)
- Craig Gardner, SUSE (Maintainer)
- [Renato Golin, Linaro \(Testing Coordinator\)](#)
- Jennifer Green, Los Alamos National Laboratory (Maintainer)
- Douglas Jacobsen, NERSC (End-User/Site Representative)
- Chulho Kim, Lenovo (Maintainer)
- [Janet Lebens, Cray \(Maintainer\)](#)
- Thomas Moschny, ParTec (Maintainer)
- [Nam Pho, NYU \(Maintainer\)](#)
- [Cyrus Proctor, Texas Advanced Computing Center \(Maintainer\)](#)
- [Adrian Reber, Red Hat \(Maintainer\)](#)
- [Joseph Stanfield, Dell EMC \(Maintainer\)](#)
- Karl W. Schulz, Intel (Project Lead, Testing Coordinator)
- [Jeff Schutkoske, Cray \(Component Development Representative\)](#)
- Derek Simmel, Pittsburgh Supercomputing Center (End-User/Site Representative)
- Scott Suchyta, Altair (Maintainer)
- Nirmala Sundararajan, Dell EMC (Maintainer)

New members for 2017-2018

Interested in participating next year?

- *expect call for nominations next summer (June/July)*

Updates / New Items

will next highlight some new items/changes since last year's BoF

Switched from ISOs -> distribution tarballs

- If your system does not have external network access, or you just prefer to mirror things locally, we provide a mechanism to download entire repo
- Previously we provided ISOs, now these are tarball archives available at:

<http://build.openhpc.community/dist>

- A “`make_repo.sh`” script is provided that will set up a locally hosted OpenHPC repository using the contents from downloaded tarball

Index of /dist/1.3.1

Name	Last modified	Size
 Parent Directory		-
 OpenHPC-1.3.1.CentOS_7.aarch64.tar	2017-06-16 20:54	1.1G
 OpenHPC-1.3.1.CentOS_7.src.tar	2017-06-16 20:57	9.2G
 OpenHPC-1.3.1.CentOS_7.x86_64.tar	2017-06-16 20:55	2.8G
 OpenHPC-1.3.1.SLE_12.aarch64.tar	2017-06-16 20:50	839M
 OpenHPC-1.3.1.SLE_12.src.tar	2017-06-16 20:52	8.8G
 OpenHPC-1.3.1.SLE_12.x86_64.tar	2017-06-16 20:50	2.3G
 OpenHPC-1.3.1.md5s	2017-06-16 21:03	1.6K

```
# tar xf OpenHPC-1.3.1.CentOS_7.x86_64.tar
# ./make_repo.sh
--> Creating OpenHPC.local.repo file in /etc/yum.repos.d
--> Local repodata stored in /root/repo
```

```
# yum repolist | grep OpenHPC
OpenHPC-local          OpenHPC-1.3 - Base
OpenHPC-local-updates  OpenHPC-1.3.1 - Updates
```


More Generic Repo Paths

- Starting with the v1.3 release, we adopted more generic paths for underlying distros

```
[OpenHPC]
name=OpenHPC-1.3 - Base
baseurl=http://build.openhpc.community/OpenHPC:/1.3/CentOS_7
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-OpenHPC-1

[OpenHPC-updates]
name=OpenHPC-1.3 - Updates
baseurl=http://build.openhpc.community/OpenHPC:/1.3/updates/CentOS_7
gpgcheck=1
gpgkey=file:///etc/pki/rpm-gpg/RPM-GPG-KEY-OpenHPC-1
```

- Similar approach taken for SLES12 repo config -> SLE_12

Release Roadmap

Release	Target Release Date	Expectations
1.3.4	Feb/March 2018	New component additions and version upgrades.

- We have had some requests for a roadmap for future releases
- A high-level roadmap is now maintained on GitHub [wiki](#)

Release	Date
1.3.3	November 08, 2017
1.3.2	September 07, 2017
1.3.1	June 16, 2017
1.3	March 31, 2017
1.2.1	January 24, 2017
1.2	November 12, 2016
1.1.1	June 21, 2016
1.1	April 18, 2016
1.0.1	February 05, 2016
1.0	November 12, 2015

Meta packages

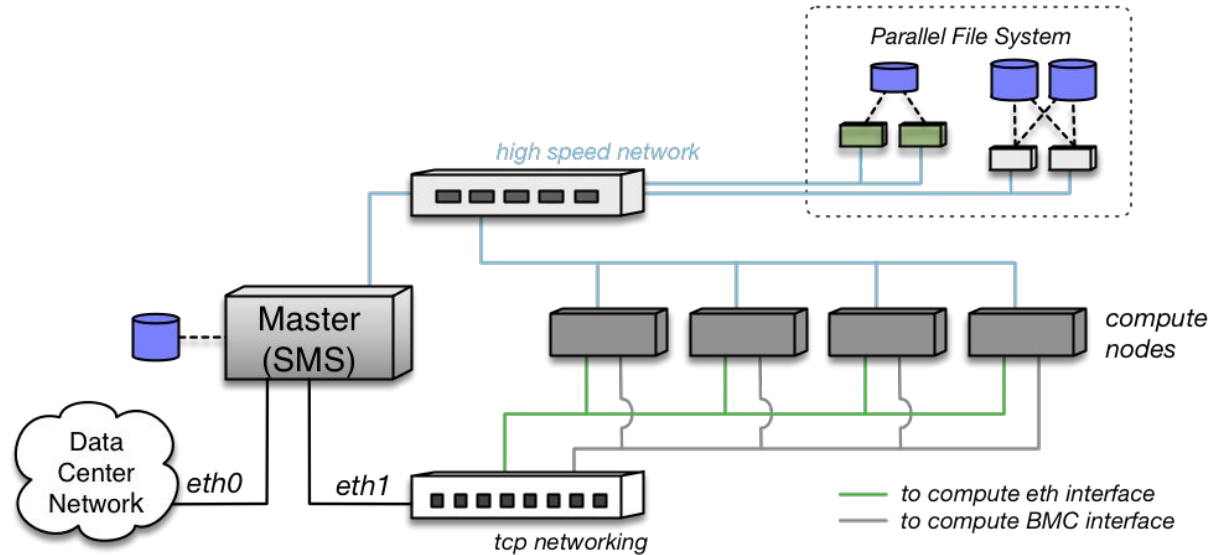
- Meta RPM packages were introduced with v1.3.1 release
 - these replace previous use of groups/patterns
 - general naming convention remains the same
 - names that begin with “ohpc-*” are typically metapackages
 - intended to group related collections of RPMs by functionality
- There is no requirement to use, but we leverage them in the recipes for convenience
- Available package list shown in Appendix E

Table 2: Available OpenHPC Meta-packages

Group Name	Description
ohpc-autotools	Collection of GNU autotools packages.
ohpc-base	Collection of base packages.
ohpc-base-compute	Collection of compute node base packages.
ohpc-ganglia	Collection of Ganglia monitoring and metrics packages.
ohpc-gnu7-io-libs	Collection of IO library builds for use with GNU compiler toolchain.
ohpc-gnu7-mpich-parallel-libs	Collection of parallel library builds for use with GNU compiler toolchain and the MPICH runtime.
ohpc-gnu7-mvapich2-parallel-libs	Collection of parallel library builds for use with GNU compiler toolchain and the MVAPICH2 runtime.
ohpc-gnu7-openmpi-parallel-libs	Collection of parallel library builds for use with GNU compiler toolchain and the OpenMPI runtime.
ohpc-gnu7-parallel-libs	Collection of parallel library builds for use with GNU compiler toolchain.
ohpc-gnu7-perf-tools	Collection of performance tool builds for use with GNU compiler toolchain.
ohpc-gnu7-python-libs	Collection of python related library builds for use with GNU compiler toolchain.
ohpc-gnu7-runtimes	Collection of runtimes for use with GNU compiler toolchain.
ohpc-gnu7-serial-libs	Collection of serial library builds for use with GNU compiler toolchain.
ohpc-intel-impi-parallel-libs	Collection of parallel library builds for use with Intel(R) Parallel Studio XE toolchain and the Intel(R) MPI Library.
ohpc-intel-io-libs	Collection of IO library builds for use with Intel(R) Parallel Studio XE software suite.
ohpc-intel-mpich-parallel-libs	Collection of parallel library builds for use with Intel(R) Parallel Studio XE toolchain and the MPICH runtime.
ohpc-intel-mvapich2-parallel-libs	Collection of parallel library builds for use with Intel(R) Parallel Studio XE toolchain and the MVAPICH2 runtime.
ohpc-intel-openmpi-parallel-libs	Collection of parallel library builds for use with Intel(R) Parallel Studio XE toolchain and the OpenMPI runtime.
ohpc-intel-perf-tools	Collection of performance tool builds for use with Intel(R) Parallel Studio XE toolchain.
ohpc-intel-python-libs	Collection of python related library builds for use with Intel(R) Parallel Studio XE toolchain.
ohpc-intel-runtimes	Collection of runtimes for use with Intel(R) Parallel Studio XE toolchain.
ohpc-intel-serial-libs	Collection of serial library builds for use with Intel(R) Parallel Studio XE toolchain.
ohpc-nagios	Collection of Nagios monitoring and metrics packages.
ohpc-slurm-client	Collection of client packages for SLURM.
ohpc-slurm-server	Collection of server packages for SLURM.
ohpc-warewulf	Collection of base packages for Warewulf provisioning.

Recipes updated for Ethernet only fabrics

- Some users reported confusion over what configuration or installation items were optional in an ethernet-only cluster
- Recipes updated to hopefully improve this



[Tech Preview] designation removed from ARM recipes

```
[train01@sms001 ~]$ module avail
```

```
x86_64
```

```
----- /opt/ohpc/pub/moduledeps/gnu7-mpich -----
adios/1.11.0  mpiP/3.4.1      petsc/3.7.6      scorep/3.0
boost/1.63.0  mumps/5.1.1    phdf5/1.10.0    sionlib/1.7.1
fftw/3.3.6   netcdf-cxx/4.3.0  scalapack/2.0.2  superlu_dist/4.2
hypr/2.11.1  netcdf-fortran/4.4.4  scalasca/2.3.1  tau/2.26.1
imb/4.1      netcdf/4.4.1.1  scipy/0.19.0    trilinos/12.10.1

----- /opt/ohpc/pub/moduledeps/gnu7 -----
R/3.3.3      metis/5.1.0    numpy/1.12.1    openmpi/1.10.7
gs1/2.3      mpich/3.2 (L)  ocr/1.0.1      pdtoolkit/3.23
hdf5/1.10.0  mvapich2/2.2  openblas/0.2.19  superlu/5.2.1

-----
EasyBuild/3.2.1
autotools (L)
```

```
train01@cavium1:~> module avail
```

```
aarch64
```

```
----- /opt/ohpc/pub/moduledeps/gnu7-mpich -----
adios/1.11.0  mpiP/3.4.1      petsc/3.7.6      scorep/3.0
boost/1.63.0  mumps/5.1.1    phdf5/1.10.0    sionlib/1.7.1
fftw/3.3.6   netcdf-cxx/4.3.0  scalapack/2.0.2  superlu_dist/4.2
hypr/2.11.1  netcdf-fortran/4.4.4  scalasca/2.3.1  tau/2.26.1
imb/4.1      netcdf/4.4.1.1  scipy/0.19.0    trilinos/12.10.1

----- /opt/ohpc/pub/moduledeps/gnu7 -----
R/3.3.3      metis/5.1.0    ocr/1.0.1      pdtoolkit/3.23
gs1/2.3      mpich/3.2 (L)  openblas/0.2.19  superlu/5.2.1
hdf5/1.10.0  numpy/1.12.1  openmpi/1.10.7

----- /opt/ohpc/pub/modulefiles -----
EasyBuild/3.2.1  hwloc/1.11.6  singularity/2.3
autotools (L)  ohpc (L)  valgrind/3.12.0
```

Latest v1.3.3 release includes newer Warewulf that leverages iPXE for multiple-architecture provisioning support

Bare-metal recipes for aarch64 now tested in OpenHPC CI infrastructure

OpenHPC providing consistent development environment to the end user across multiple architectures.


PEARC'17

- We presented our first tutorial in July at the PEARC'17 conference: *Getting Started with OpenHPC*
 - Phase I:
 - Target system architecture
 - Stack overview
 - Packaging conventions
 - Module conventions
 - Development environment
 - Integration testing
 - Upgrade paths
 - Phase II:
 - Hands-on system installations
 - Leveraged CI infrastructure kindly hosted at TACC
- Tutorial material available at: <https://goo.gl/NyiDmr>



Addition of new compiler/MPI variants

- We have started to introduce new compiler/MPI variants over time:
 - **gnu7** compiler variant introduced in v1.3.1 release
 - **llvm5** compiler variant introduced in v1.3.3
 - **openmpi3** MPI variant introduced in v1.3.3
- In the case of a fresh install, recipes default to installing the new variant along with matching runtimes and libraries
- If upgrading a previously installed system, administrators can opt-in to enable the new variant
- meta-packages provide a convenience mechanism to add on:



```
[sms]# yum -y install ohpc-gnu7-perf-tools \  
ohpc-gnu7-serial-libs \  
ohpc-gnu7-io-libs \  
ohpc-gnu7-python-libs \  
ohpc-gnu7-runtimes \  
ohpc-gnu7-mpich-parallel-libs \  
ohpc-gnu7-openmpi-parallel-libs \  
ohpc-gnu7-mvapich2-parallel-libs
```

An arrow points from the text 'meta-packages provide a convenience mechanism to add on:' to the terminal screenshot.

Coexistence of multiple compiler/MPI variants

Example: admin upgraded from v1.3 to v1.3.1 and opted-in to install gnu7 variant

```
$ module list
Currently Loaded Modules:
  1) autotools  2) prun/1.1  3) gnu7/7.1.0  4) mvapich2/2.2  5) ohpc

$ module avail

----- /opt/ohpc/pub/moduledeps/gnu7-mvapich2 -----
adios/1.11.0  imb/4.1  netcdf-cxx/4.3.0  scalapack/2.0.2  sionlib/1.7.1
boost/1.63.0  mpiP/3.4.1  netcdf-fortran/4.4.4  scalasca/2.3.1  superlu_dist/4.2
fftw/3.3.6  mumps/5.1.1  petsc/3.7.6  scipy/0.19.0  tau/2.26.1
hypre/2.11.1  netcdf/4.4.1.1  phdf5/1.10.0  scorep/3.0  trilinos/12.10.1

----- /opt/ohpc/pub/moduledeps/gnu7 -----
R_base/3.3.3  metis/5.1.0  numpy/1.12.1  openmpi/1.10.7
gsl/2.3  mpich/3.2  ocr/1.0.1  pdtoolkit/3.23
hdf5/1.10.0  mvapich2/2.2 (L)  openblas/0.2.19  superlu/5.2.1

----- /opt/ohpc/pub/modulefiles -----
EasyBuild/3.2.1  gnu/5.4.0  ohpc (L)  singularity/2.3
autotools (L)  gnu7/7.1.0 (L)  papi/5.5.1  valgrind/3.12.0
clustershell/1.7.3  hwloc/1.11.6  prun/1.1 (L)
```

previously
installed from
1.3 release

Variety of installation recipes now available

Latest v1.3.3 release provides recipes with multiple resource managers, OSes, provisioners, and architectures:

- [Install_guide-CentOS7-Warewulf-PBSPPro-1.3.3-x86_64.pdf](#)
- [Install_guide-CentOS7-Warewulf-SLURM-1.3.3-aarch64.pdf](#)
- [Install_guide-CentOS7-Warewulf-SLURM-1.3.3-x86_64.pdf](#)
- [Install_guide-CentOS7-xCAT-SLURM-1.3.3-x86_64.pdf](#)
- [Install_guide-SLE_12-Warewulf-PBSPPro-1.3.3-x86_64.pdf](#)
- [Install_guide-SLE_12-Warewulf-SLURM-1.3.3-aarch64.pdf](#)
- [Install_guide-SLE_12-Warewulf-SLURM-1.3.3-x86_64.pdf](#)

- Additional provisioner (xCAT) added in v1.3.1
- iPXE based Warewulf added in v1.3.3



Initially, we started off with only a single recipe

Template scripts

Note that simple template bash scripts are provided that encapsulate commands presented in the guides. Also provide input file for local settings:

```
# yum/zypper install docs-ohpc

# ls /opt/ohpc/pub/doc/recipes/**/*/recipe.sh
/opt/ohpc/pub/doc/recipes/centos7/aarch64/warewulf/slurm/recipe.sh
/opt/ohpc/pub/doc/recipes/centos7/x86_64/warewulf/pbspro/recipe.sh
/opt/ohpc/pub/doc/recipes/centos7/x86_64/warewulf/slurm/recipe.sh
/opt/ohpc/pub/doc/recipes/centos7/x86_64/xcat/slurm/recipe.sh
/opt/ohpc/pub/doc/recipes/sles12/aarch64/warewulf/slurm/recipe.sh
/opt/ohpc/pub/doc/recipes/sles12/x86_64/warewulf/pbspro/recipe.sh
/opt/ohpc/pub/doc/recipes/sles12/x86_64/warewulf/slurm/recipe.sh

# ls /opt/ohpc/pub/doc/recipes/*/input.local
/opt/ohpc/pub/doc/recipes/centos7/input.local
/opt/ohpc/pub/doc/recipes/sles12/input.local
```

Local System Settings can be supplied to input file:

```
# compute hostnames
c_name[0]=c1
c_name[1]=c2
...
# compute node MAC addresses
c_mac[0]=00:1a:2b:3c:4f:56
c_mac[1]=00:1a:2b:3c:4f:56
...
```

input.local + recipe.sh == installed system

Test Suite

Based on request at last year's SC'16 BoF, the OpenHPC test suite is now available as an installable RPM (introduced with v1.3 release)

```
# yum/zypper install test-suite-ohpc
```

- creates/relies on “ohpc-test” user to perform user testing (with accessibility to run jobs through resource manager)
- related discussion added to recipes in Appendix C

```
[sms]# su - ohpc-test
[test@sms ~]$ cd tests
[test@sms ~]$ ./configure --disable-all --enable-fftw
checking for a BSD-compatible install... /bin/install -c
checking whether build environment is sane... yes
...
----- SUMMARY -----
Package version..... : test-suite-1.3.0
Build user.....       : ohpc-test
Build host.....       : sms001
Configure date.....   : 2017-03-24 15:41
Build architecture... : aarch64
Host operating system..... : CentOS
Compiler Families..... : gnu7
MPI Families.....     : mpich mvapich2 openmpi3
Resource manager ..... : SLURM
Provisioner .....    : Warewulf
Test suite configuration..... : long
```

Requesting Additional Software

- Recall that we have a simple submission site for new requests:
 - <https://github.com/openhpc/submissions>
 - TSC completed 3 review cycles this year
- Items added via this mechanism since v1.2. release (Nov' 16)
 - BeeGFS client
 - xCAT installation recipe
 - hwloc
 - Singularity
 - LLVM/clang
 - PLASMA
 - pNetCDF
 - SCOTCH
 - SLEPc
 - PMIx
 - MPI4py

Next Submission Deadline: December 8th, 2017

Software Name
Public URL
Technical Overview
Latest stable version number
Open-source license type
Relationship to component? <input type="checkbox"/> contributing developer <input type="checkbox"/> user <input type="checkbox"/> other
If other, please describe:
Build system <input type="checkbox"/> autotools-based <input type="checkbox"/> CMake <input type="checkbox"/> other
If other, please describe:

OpenHPC v1.3.3 Release

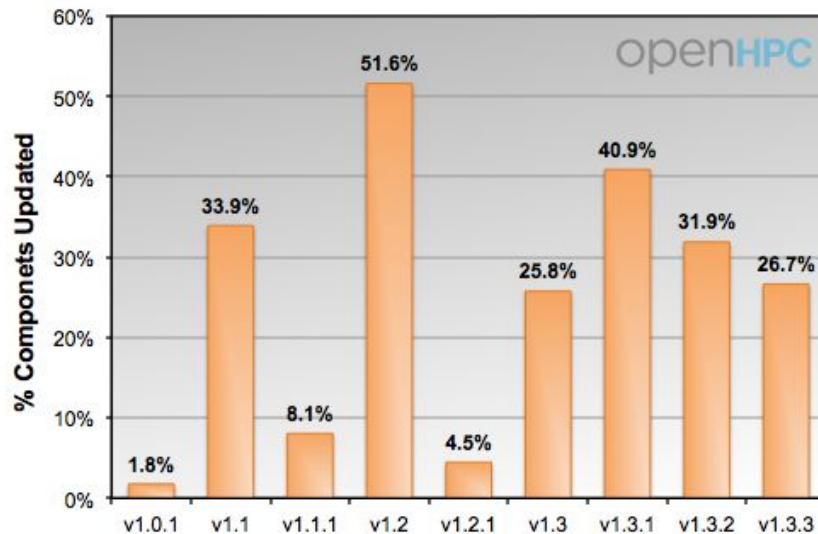
November 2017

OpenHPC v1.3.3 - Current S/W components

Functional Areas	Components
Base OS	CentOS 7.4, SLES12 SP3 new with v1.3.3
Architecture	aarch64, x86_64
Administrative Tools	Conman, Ganglia, Lmod, LosF, Nagios, pdsh, pdsh-mod-slurm, prun, EasyBuild, ClusterShell, mrsh, Genders, Shine, Spack, test-suite
Provisioning	Warewulf, xCAT
Resource Mgmt.	SLURM, Munge, PBS Professional, PMIx
Runtimes	OpenMP, OCR, Singularity
I/O Services	Lustre client, BeeGFS client
Numerical/Scientific Libraries	Boost, GSL, FFTW, Hypre, Metis, Mumps, OpenBLAS, 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), Clang/LLVM
MPI Families	MVAPICH2, OpenMPI, MPICH
Development Tools	Autotools, cmake, hwloc, mpi4py , R, SciPy/NumPy, Valgrind
Performance Tools	PAPI, IMB, mpiP, pdtoolkit TAU, Scalasca, ScoreP, SIONLib

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

Component growth/update history



of pre-packaged RPMs available (v1.3.3)

Base OS	x86_64	aarch64	noarch
CentOS 7	517	309	34
SLES 12	521	317	34

OpenHPC and Containers



- The v1.3.2 release of OpenHPC (June 2017) added Singularity
 - <https://github.com/singularityware/singularity>
 - Singularity containers can be used to package entire scientific workflows, software, libraries, and even data
 - Workflow utilizes container images:
 - as you enter and work within the Singularity container, you are physically located inside of this image
 - image grows and shrinks in real time as you install or delete files
 - if you want to copy a container, you copy the image.
- An OpenHPC v1.3.3 container containing the base development environment is now published on Singularity Hub (a registry for sharing images)
 - <https://www.singularity-hub.org/collections/243>

OpenHPC and Containers



- Quick example: build/run MPI executable

```
$ module load singularity
$ singularity exec shub://openhpc/ohpc:1.3.3.e17 /bin/bash

Singularity> module swap openmpi3 mpich
Singularity> module list

Currently Loaded Modules:
  1) prun/1.2   2) gnu7/7.2.0  3) ohpc   4) mpich/3.2

Singularity> mpicc tests/mpi/hello.c
Singularity> mpirun -np 4 ./a.out

Hello, world (4 procs total)
--> Process #   0 of   4 is alive. -> sms017
--> Process #   1 of   4 is alive. -> sms017
--> Process #   2 of   4 is alive. -> sms017
--> Process #   3 of   4 is alive. -> sms017
```

Future Work

- Stateful installation recipe with xCAT
- Inclusion of LIKWID performance analysis tool
- Compatibility builds with ARM HPC compiler
- Continued s/w updates and releases....

Open Discussion

Open Discussion - Potential Topics

- release cadence
 - have been on a roughly quarterly cadence
 - is this useful, or too much?
- upgrades
 - majority of repo access now on 1.3 branch
 - people able to upgrade ok?
 - comments on opt-in approach for new compiler/mpi variants
 - do folks use the recipe.sh scripts to aid in installation?
 - would an “upgrade” option to these scripts be helpful?
- interest in more tutorials?
- what would you like to see from OpenHPC going forward?
 - general comments/questions
 - are folks using ohpc collateral in cloud/virtual machines?
 - interest in using machine learning software in HPC environment?

Pointers to Online Resources

- OpenHPC Home: <http://www.openhpc.community/>
- Primary GitHub Site: <https://github.com/openhpc/ohpc>
- Package Repositories: <http://build.openhpc.community/OpenHPC:/>
- OBS Frontend: <https://build.openhpc.community>
- Component Submission: <https://github.com/openhpc/submissions>
- System Registry: [System Registration Form](#)
- CI Infrastructure: <http://test.openhpc.community:8080>
- OpenHPC Wiki: <https://github.com/openhpc/ohpc/wiki>
 - includes links to overview paper and past presentations

- Mailing Lists: <http://www.openhpc.community/support/mail-lists/>
 - [openhpc-announce](#)
 - [openhpc-users](#)
 - [openhpc-devel](#)