



OpenHPC Community BOF

Karl W. Schulz, Intel

David Brayford, Leibniz Supercomputing Centre

November 16 ♦ Salt Lake City, Utah

Today's Agenda

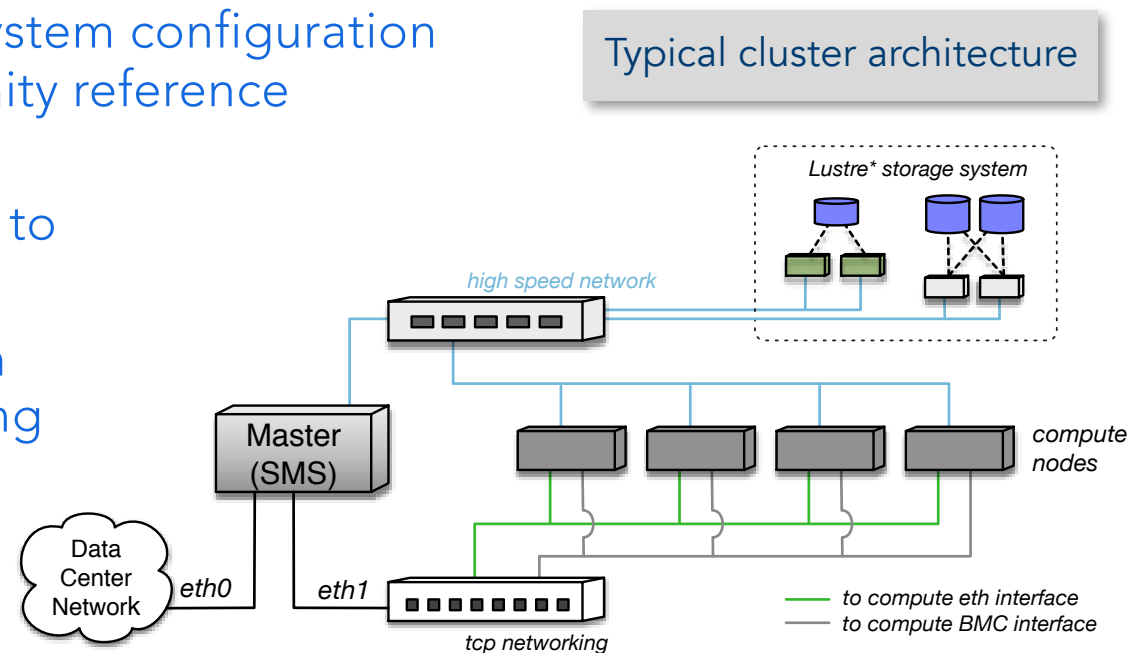
- Motivation/Intro
- What's New/Accomplishments
 - mission/vision, members
 - governance overview
 - component submission process
 - system registry form
 - infrastructure updates
 - v1.2 release - new goodies
- Future work
- Open Forum

Background Motivation for Community Effort

- Many sites spend considerable effort aggregating a large suite of open-source projects to provide a capable HPC environment for their users:
 - necessary to build/deploy HPC focused packages that are either absent or do not keep pace from distro providers
 - local packaging or customization frequently tries to give software versioning access to users (e.g. via modules or similar equivalent)
 - hierarchal packaging necessary for multiple compiler/mpi families
- On the developer front, many successful projects must engage in continual triage and debugging regarding configuration and installation issues on HPC systems

What is OpenHPC?

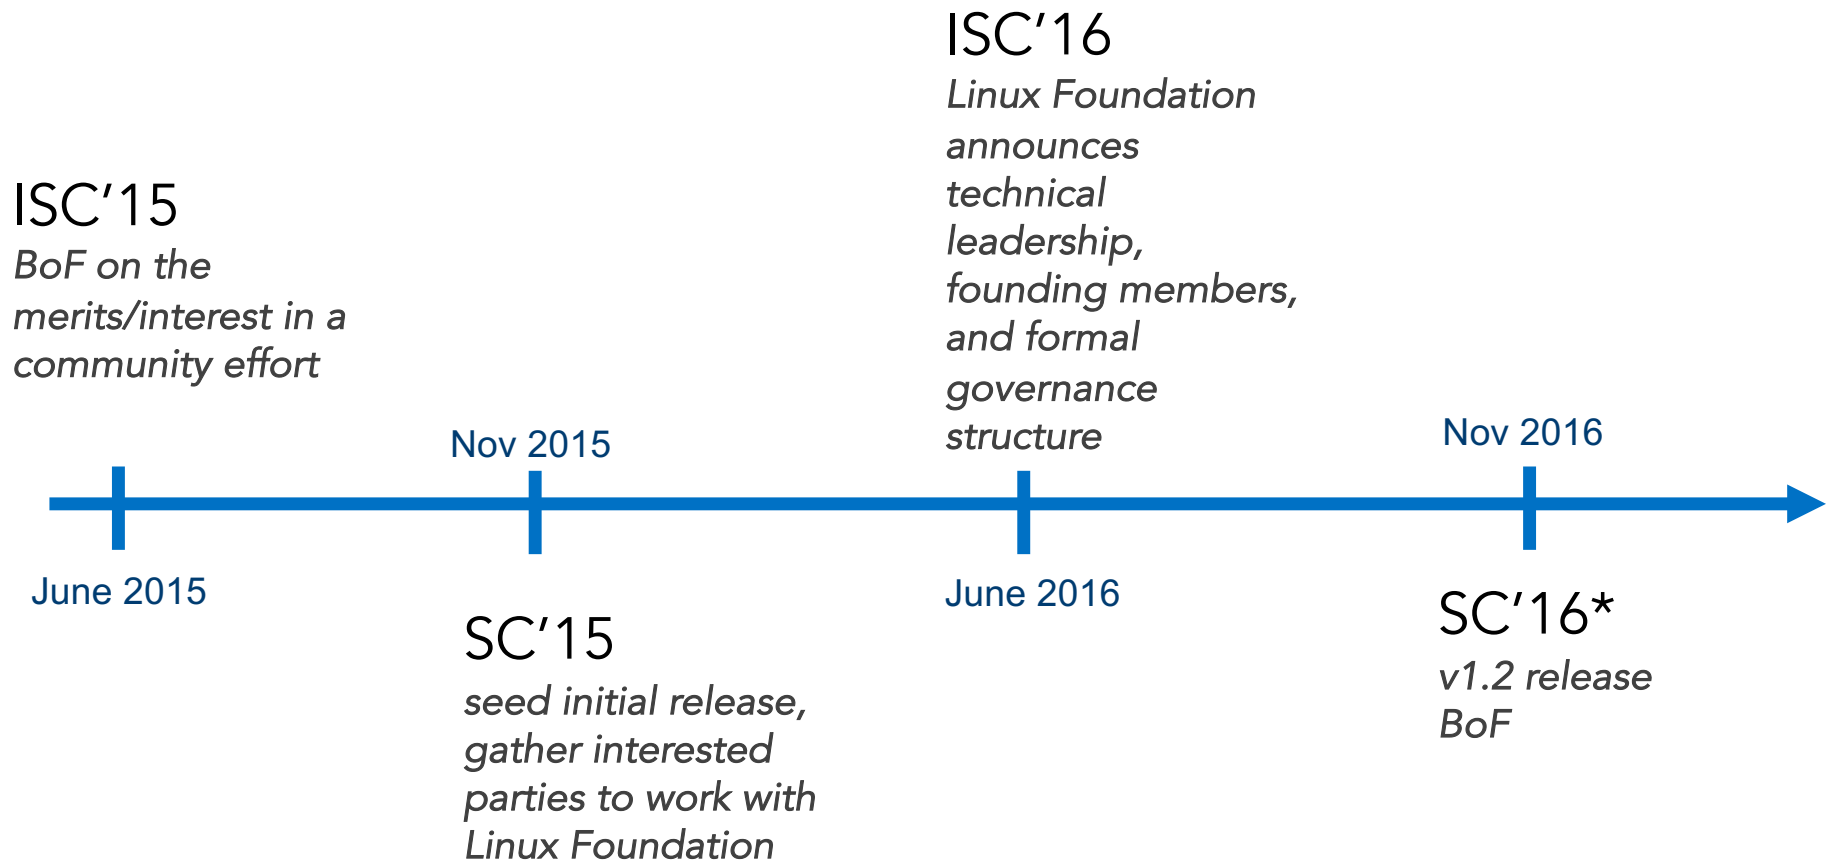
- OpenHPC is a community effort endeavoring to:
 - provide collection(s) of pre-packaged components that can be used to help install and manage flexible HPC systems throughout their lifecycle
 - leverage standard Linux delivery model to retain admin familiarity (ie. package repos)
 - allow and promote multiple system configuration recipes that leverage community reference designs and best practices
 - implement integration testing to gain validation confidence
 - provide additional distribution mechanism for groups releasing open-source software
 - provide a stable platform for new R&D initiatives



OpenHPC: Mission and Vision

- **Mission**: *to provide a reference collection of open-source HPC software components and best practices, lowering barriers to deployment, advancement, and use of modern HPC methods and tools.*
- **Vision**: OpenHPC components and best practices will enable and accelerate innovation and discoveries by broadening access to state-of-the-art, open-source HPC methods and tools in a consistent environment, supported by a collaborative, worldwide community of HPC users, developers, researchers, administrators, and vendors.

OpenHPC: a brief History...



OpenHPC: Project Members



Argonne
National
Laboratory



Altair

ARM Atos



Barcelona
Supercomputing
Center
Centro Nacional de Supercomputación



CEA



CRAY



FUJITSU



Hewlett Packard
Enterprise



Lawrence Livermore
National Laboratory



Lenovo



sgi



TACC

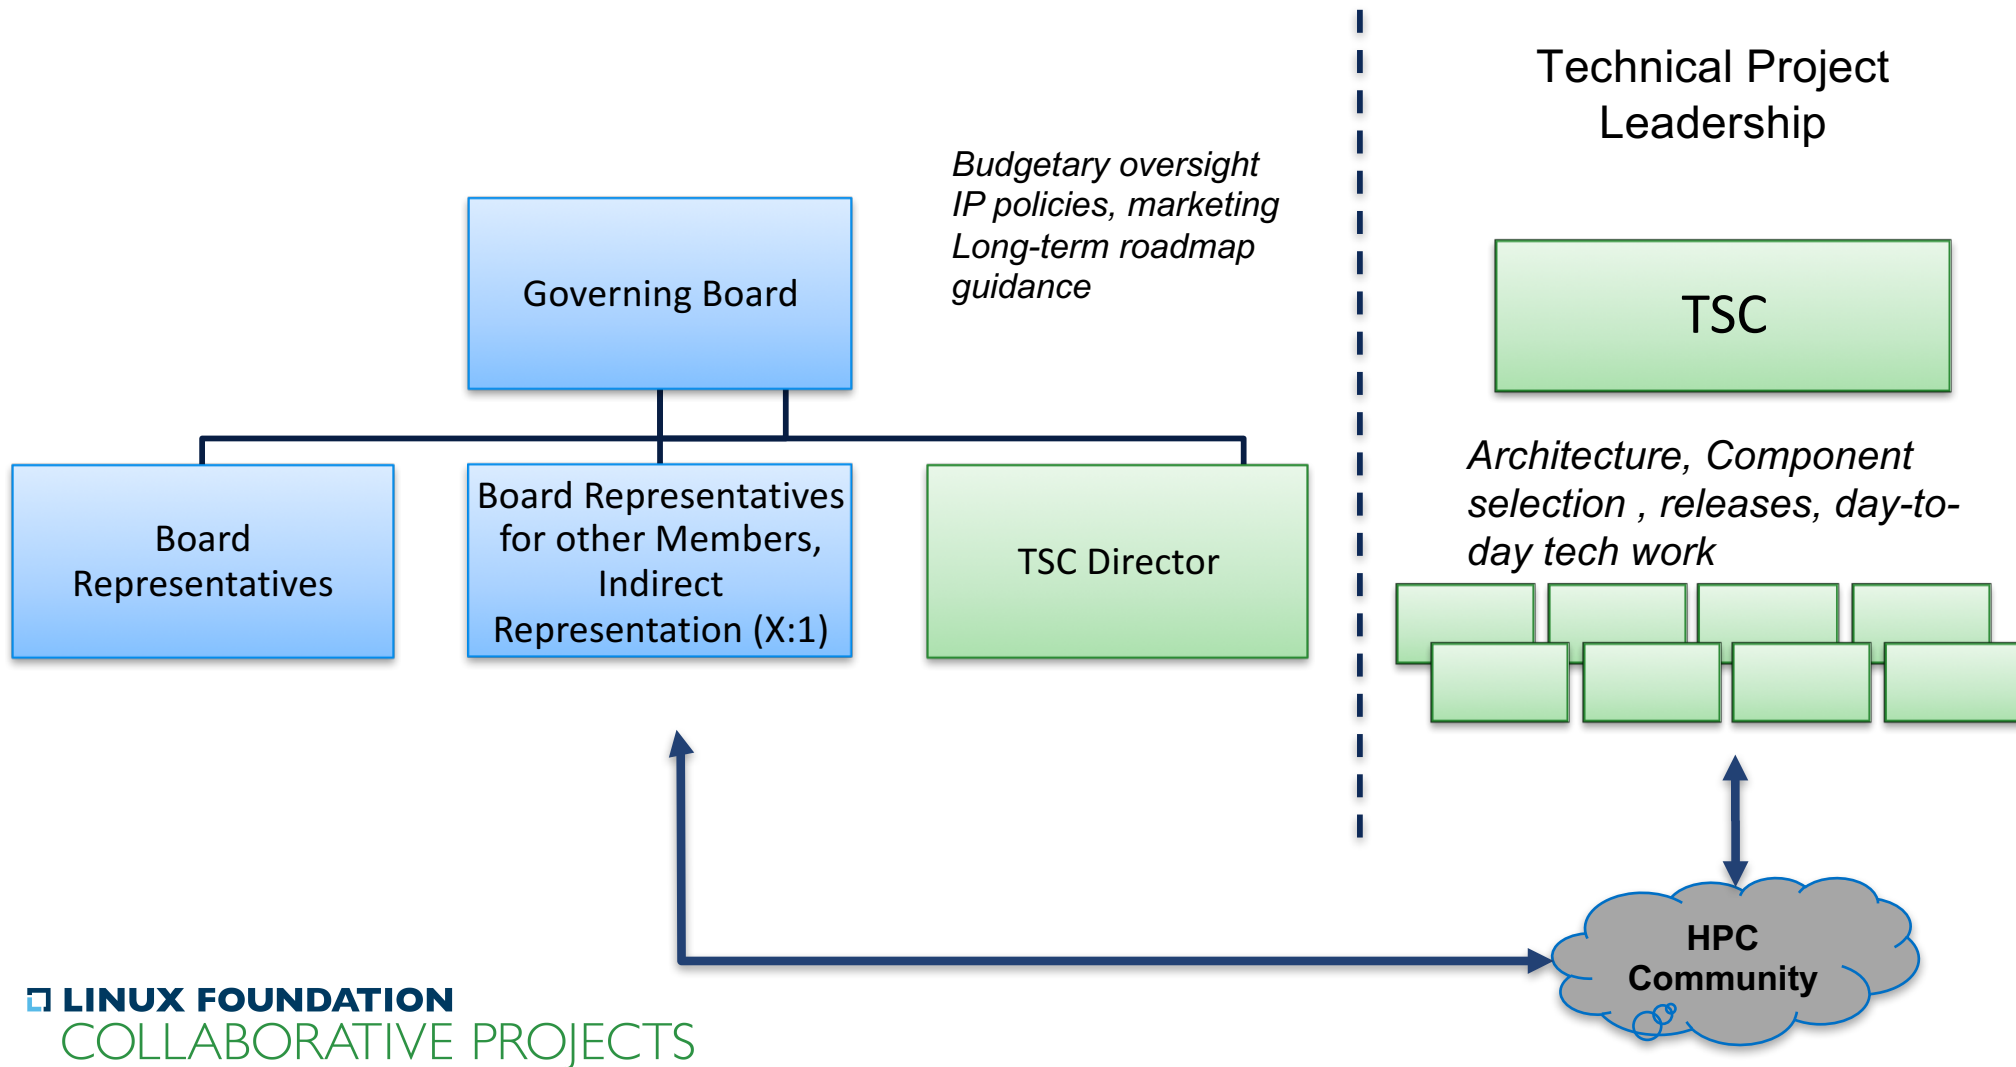
UNIVA

*Mixture of Academics, Labs,
OEMs, and ISVs/OSVs*

Project member participation interest? Please contact
Kevlin Husser or Jeff ErnstFriedman
jernstfriedman@linuxfoundation.org

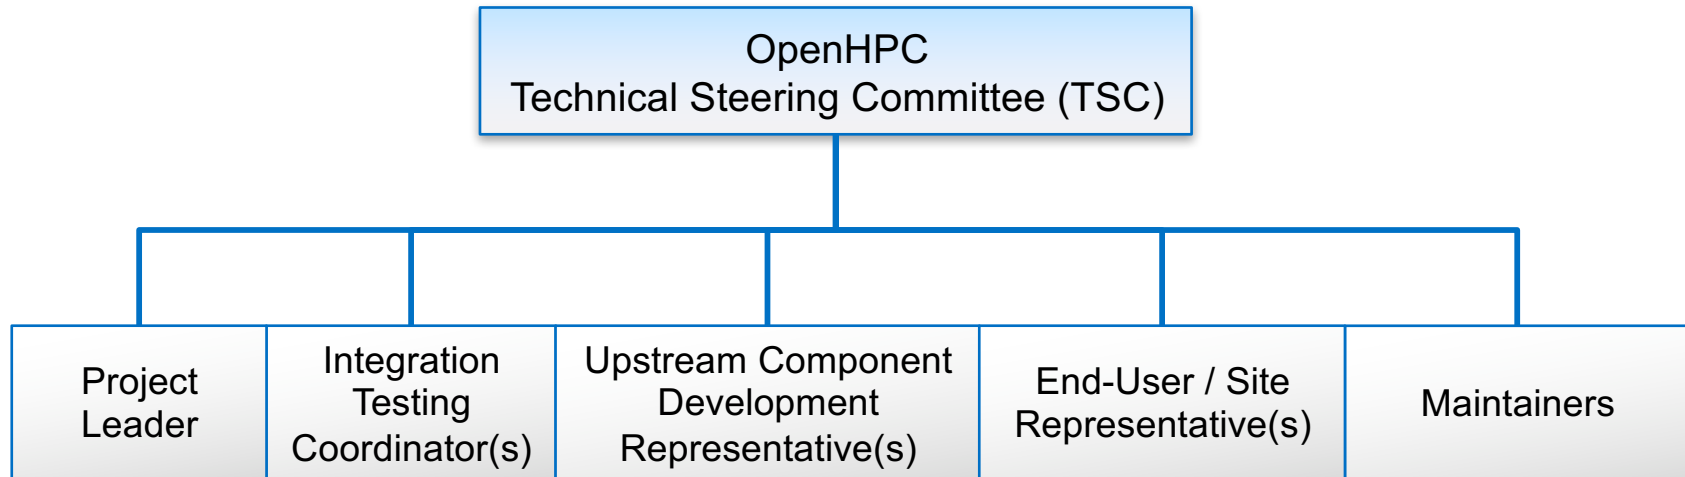
Community Governance Overview

Governing Board + Technical Steering Committee



OpenHPC Technical Steering Committee (TSC)

Role Overview



OpenHPC TSC – Individual Members

- Reese Baird, Intel (Maintainer)
- Pavan Balaji, Argonne National Laboratory (Maintainer)
- David Brayford, LRZ (Maintainer)
- Todd Gamblin, Lawrence Livermore National Labs (Maintainer)
- Craig Gardner, SUSE (Maintainer)
- Yiannis Georgiou, ATOS (Maintainer)
- Balazs Gerofi, RIKEN (Component Development Representative)
- Jennifer Green, Los Alamos National Laboratory (Maintainer)
- Eric Van Hensbergen, ARM (Maintainer, Testing Coordinator)
- Douglas Jacobsen, NERSC (End-User/Site Representative)
- Chulho Kim, Lenovo (Maintainer)
- Greg Kurtzer, Lawrence Berkeley National Labs (Component Development Representative)
- Thomas Moschny, ParTec (Maintainer)
- Karl W. Schulz, Intel (Project Lead, Testing Coordinator)
- Derek Simmel, Pittsburgh Supercomputing Center (End-User/Site Representative)
- Thomas Sterling, Indiana University (Component Development Representative)
- Craig Stewart, Indiana University (End-User/Site Representative)
- Scott Suchyta, Altair (Maintainer)
- Nirmala Sundararajan, Dell (Maintainer)

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

Component Additions?

- A common question posed to the project has been how to request new software components? In response, the TSC has endeavored to formalize a simple submission/review process
- Submission site went live in September:
<https://github.com/openhpc/submissions>
- Expecting to do reviews quarterly
 - completed first iteration for v1.2 release
 - next submission deadline: **January 27th, 2017**

Subset of information requested during submission process

Software Name

Public URL

Technical Overview

Latest stable version number

Open-source license type

Relationship to component?

- contributing developer
- user
- other

If other, please describe:

Build system

- autotools-based
- CMake
- other

Opt-in System Registry Now Available

System Registry

- Interested users can now register their usage on a public system registry
- Helpful for us to have an idea as to who is potentially benefitting from this community effort
- Accessible from top-level GitHub page

OpenHPC System Registry

This opt-in form can be used to register your system to let us (and the community) know that you are using elements of OpenHPC.

* Required

Name of Site/Organization *

Your answer

What OS distribution are you using? *

CentOS/RHEL

SLES

Other: _____

Site or System URL

Your answer

Infrastructure

Project CI infrastructure

- TACC is kindly hosting some CI infrastructure for the project (Austin, TX)
- Using for build servers and continuous integration (CI) testbed. Housing point for maintainers and interested parties to examine test results

<http://test.openhpc.community:8080>

Many thanks to TACC and vendors for hardware donations!!: Intel, Cavium, Dell



Community Test System (CI) - Jenkins

<http://test.openhpc.community:8080>

OpenHPC CI Infrastructure

Thanks to the Texas Advanced Computing Center (TACC) for hosting support and to Intel, Cavium, and Dell for hardware donations.

 add description

1.1.1 **1.2** All Interactive admin +

S	W	Name ↓	Last Success	Last Failure	Last Duration	
✓	🔗	(1.2) - (centos7.2,x86_64) - (warewulf+pbspro) - long cycle	1 day 8 hr - #39	N/A	58 min	▶
✓	🔗	(1.2) - (centos7.2,x86_64) - (warewulf+pbspro) - short cycle	1 day 4 hr - #155	4 days 10 hr - #109	14 min	▶
✓	🔗	(1.2) - (centos7.2,x86_64) - (warewulf+slurm) - long cycle	2 days 4 hr - #244	4 days 4 hr - #219	1 hr 0 min	▶
✓	🔗	(1.2) - (centos7.2,x86_64) - (warewulf+slurm) - short cycle	2 hr 46 min - #554	8 days 15 hr - #349	14 min	▶
✓	🔗	(1.2) - (centos7.2,x86_64) - (warewulf+slurm+PXSE) - long cycle	1 day 6 hr - #39	4 days 10 hr - #20	2 hr 29 min	▶
✓	🔗	(1.2) - (sles12sp1,x86_64) - (warewulf+pbspro) - short cycle	1 day 3 hr - #166	4 days 10 hr - #86	12 min	▶
✓	🔗	(1.2) - (sles12sp1,x86_64) - (warewulf+slurm) - short cycle	1 day 2 hr - #259	8 days 20 hr - #72	14 min	▶
✓	🔗	(1.2) - (sles12sp1,x86_64) - (warewulf,slurm) - long test cycle	1 day 5 hr - #97	6 days 19 hr - #41	54 min	▶
⚠	🔗	(1.2) - aarch64 - (centos7.2) - (warewulf+slurm)	2 days 21 hr - #3	N/A	0.41 sec	▶
⚠	🔗	(1.2) - aarch64 - (sles12sp1) - (warewulf+slurm)	1 day 8 hr - #45	2 days 21 hr - #41	2 hr 13 min	▶

These tests periodically install bare-metal clusters from scratch using OpenHPC recipes and then run a variety of integration tests.

OpenHPC v1.2 Release

November 12, 2016

OpenHPC v1.2 - Current S/W components

Functional Areas	Components
Base OS	CentOS 7.2, SLES12 SP1
Architecture	x86_64, aarch64 (Tech Preview)
Administrative Tools	Conman, Ganglia, Lmod, LosF, Nagios, pdsh, prun, EasyBuild, ClusterShell, mrsh, Genders, Shine, Spack
Provisioning	Warewulf
Resource Mgmt.	SLURM, Munge, PBS Professional
Runtimes	OpenMP, OCR
I/O Services	Lustre client (community version)
Numerical/Scientific Libraries	Boost, GSL, FFTW, Metis, PETSc, Trilinos, Hypre, SuperLU, SuperLU_Dist, Mumps, OpenBLAS, Scalapack
I/O Libraries	HDF5 (pHDF5), NetCDF (including C++ and Fortran interfaces), Adios
Compiler Families	GNU (gcc, g++, gfortran)
MPI Families	MVAPICH2, OpenMPI, MPICH
Development Tools	Autotools (autoconf, automake, libtool), Valgrind, R, SciPy/NumPy
Performance Tools	PAPI, IMB, mpiP, pdtoolkit TAU, Scalasca , ScoreP , SIONLib

new with v1.2

Notes:

- Additional dependencies that are not provided by the BaseOS or community repos (e.g. EPEL) are also included
- 3rd Party libraries are built for each compiler/MPI family (8 combinations typically)
- Resulting repositories currently comprised of ~300 RPMs


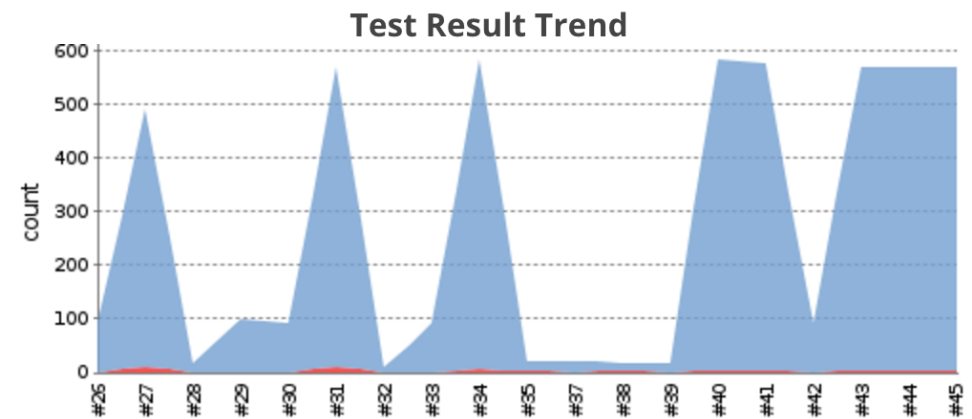
ARM Tech Preview

- Latest v1.2 release includes builds for **aarch64**
- Recipes/packages being made available as a Tech Preview
 - additional work required for provisioning
 - significant majority of development packages testing OK, but there are some caveats
 - please see <https://github.com/openhpc/ohpc/wiki/ARM-Tech-Preview> for latest info
- Eric Van Hensbergen will highlight test systems and observations in porting to aarch64

OpenHPC 1.2 Testing Dashboard (aarch64)

- Failing Modules
 - GSL
 - PAPI
 - TAU
- No IB testbed yet, so no MVAPICH2 tests
- Other Missing/Disabled:
 - Nagios and Ganglia (on SLES 12 SP 1)
 - *-intel-*
 - *-*-impi

Base OS	x86_64	aarch64	noarch
CentOS 7.2	310	198	43
SLES 12 SP1	312	205	43



OpenHPC 1.2 Tested Systems (aarch64)

- Gigabyte Cavium ThunderX
 - 2x48 core w/128GB of Memory
 - Gigabit Ethernet
 - Tested with SLES 12 SP 1
- SoftIron AMD Seattle
 - 1x8 core w/16GB of Memory
 - Gigabit Ethernet
 - Tested with SLES 12 SP 1
- Xgene Mustang
 - 1x8 core w/16GB of Memory
 - Gigabit Ethernet
 - Tested with CentOS 7.2



Future steps for ARM (aarch64)

- Work with OpenHPC to roll forward to SLES 12 SP 2 & CentOS 7.3 to get better base-level ARM support
- InfiniBand testing
- Fix provisioning workarounds via modifications of Warewulf and/or GRUB2

New Architecture Builds: Tech Preview for ARM

```
karl@sms001:~> module avail
```

x86_64

```
----- /opt/ohpc/pub/moduledeps/gnu-mpich -----
adios/1.10.0  mpiP/3.4.1      petsc/3.7.0      scorep/3.0
boost/1.61.0  mumps/5.0.2    phdf5/1.8.17    sionlib/1.7.0
fftw/3.3.4    netcdf/4.4.1    scalapack/2.0.2  superlu_dist/4.2
hypre/2.10.1  netcdf-cxx/4.2.1  scalasca/2.3.1  tau/2.25.2
imb/4.1       netcdf-fortran/4.4.4  scipy/0.18.0    trilinos/12.6.4

----- /opt/ohpc/pub/moduledeps/gnu -----
R_base/3.3.1  metis/5.1.0    numpy/1.11.1    openmpi/1.10.4
gsl/2.2.1     mpich/3.2 (L)  ocr/1.0.1       pdtoolkit/3.22
hdf5/1.8.17   mvapich2/2.2   openblas/0.2.19  superlu/5.2.1

----- /opt/ohpc/pub/modulefiles -----
EasyBuild/2.9.0  clustershell/1.7.2  ohpc (L)  prun/1.1 (L)
autotools (L)  gnu/5.4.0 (L)  papi/5.4.3  valgrind/3.11.0
```

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

aarch64

```
----- /opt/ohpc/pub/moduledeps/gnu-mpich -----
adios/1.10.0  mpiP/3.4.1      petsc/3.7.0      scorep/3.0
boost/1.61.0  mumps/5.0.2    phdf5/1.8.17    sionlib/1.7.0
fftw/3.3.4    netcdf/4.4.1    scalapack/2.0.2  superlu_dist/4.2
hypre/2.10.1  netcdf-cxx/4.2.1  scalasca/2.3.1  tau/2.25.2
imb/4.1       netcdf-fortran/4.4.4  scipy/0.18.0    trilinos/12.6.4

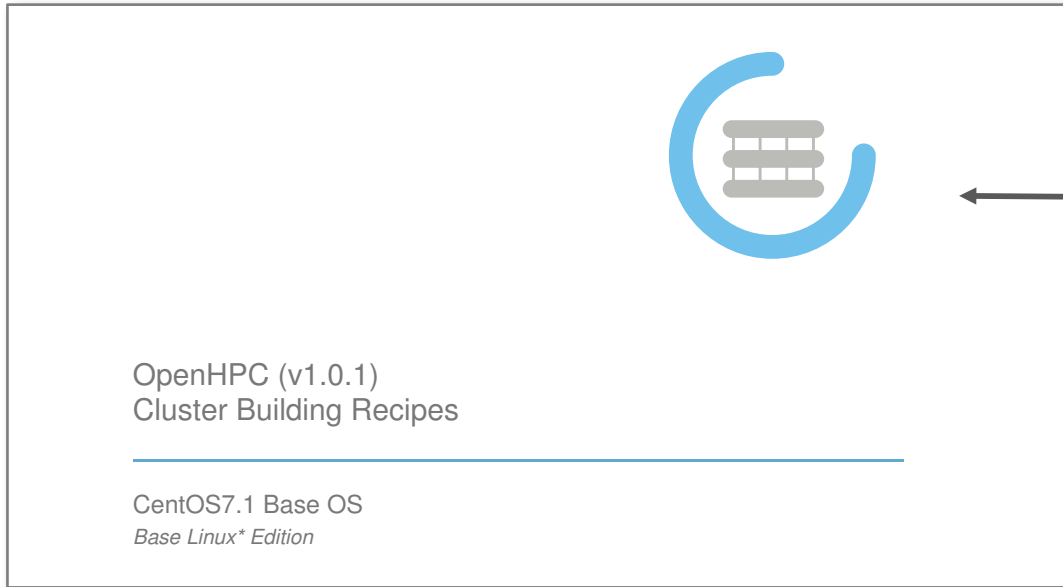
----- /opt/ohpc/pub/moduledeps/gnu -----
R_base/3.3.1  metis/5.1.0    numpy/1.11.1    openmpi/1.10.4
gsl/2.2.1     mpich/3.2 (L)  ocr/1.0.1       pdtoolkit/3.22
hdf5/1.8.17   mvapich2/2.2   openblas/0.2.19  superlu/5.2.1

----- /opt/ohpc/pub/modulefiles -----
EasyBuild/2.9.0  clustershell/1.7.2  ohpc (L)  prun/1.1 (L)
autotools (L)  gnu/5.4.0 (L)  papi/5.4.3  valgrind/3.11.0
```

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

Expanded Recipes

Choose your own adventure...



Initially, we started off with a single recipe with the intent to expand

Latest v1.2 release continues to expand with multiple resource manager options:

- [Install_guide-CentOS7.2-PBSPPro-1.2-x86_64.pdf](#)
- [Install_guide-CentOS7.2-SLURM-1.2-aarch64.pdf](#)
- [Install_guide-CentOS7.2-SLURM-1.2-x86_64.pdf](#)
- [Install_guide-SLE_12_SP1-PBSPPro-1.2-x86_64.pdf](#)
- [Install_guide-SLE_12_SP1-SLURM-1.2-aarch64.pdf](#)
- [Install_guide-SLE_12_SP1-SLURM-1.2-x86_64.pdf](#)

Expanded Recipes

Choose your own adventure

Multiple options now extend to the template recipe scripts that encapsulate commands presented in the guides:

```
# yum/zypper install docs-ohpc
```

```
# ls /opt/ohpc/pub/doc/recipes/*/*/*/*/*/recipe.sh
/opt/ohpc/pub/doc/recipes/centos7.2/x86_64/warewulf/pbspro/recipe.sh
/opt/ohpc/pub/doc/recipes/centos7.2/x86_64/warewulf/slurm/recipe.sh
/opt/ohpc/pub/doc/recipes/sles12sp1/x86_64/warewulf/pbspro/recipe.sh
/opt/ohpc/pub/doc/recipes/sles12sp1/x86_64/warewulf/slurm/recipe.sh
```

```
# ls /opt/ohpc/pub/doc/recipes/*/input.local
/opt/ohpc/pub/doc/recipes/centos7.2/input.local
/opt/ohpc/pub/doc/recipes/sles12sp1/input.local
```

```
# 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

Summary

- Community formalized as Linux Foundation collaborative project in May, 2016
- Technical Steering Committee (TSC) has been working together since the beginning of the summer
 - established a starting component selection process
 - latest release (Nov. 2016) incorporated additions based on this process
 - e.g. MPICH, PBS Pro, Scalasca/ScoreP
- Endeavoring to lower the barrier to entry for HPC and provide flexible building blocks for use by seasoned veterans

Future Work

- Continued component additions
 - interested parties are encouraged to utilize the component submission process for s/w addition requests
- Recipe(s) with different provisioning system
 - xCAT based recipe
- Policies/procedure for TSC/maintainer appointments

Open Forum

Now, let's open up for general discussion. Potential topics:

- what are your experiences so far?
- what's good/bad?
- images/containers for usage in cloud environments - interest in publishing?
- what would you like to see from OpenHPC going forward?
- general feedback, discussion...

Information/Places to Interact

<http://openhpc.community> (general info)

<https://github.com/openhpc/ohpc> (GitHub site)

<https://github.com/openhpc/submissions> (new submissions)

<https://build.openhpc.community> (build system/repos)

<http://www.openhpc.community/support/mail-lists/> (email lists)

- openhpc-announce
- openhpc-users
- openhpc-devel



Thanks for your Time