

Open HPC: an Update Including Intel's work and Interaction with E4S

Dr. Robert W. Wisniewski Chief Software Architect Extreme Scale Computing A21 PI and Lead Architect

September 23, 2019

What is the project's mission and vision?

<u>Mission</u>: 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.

<u>Vision</u>: 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.



Value of the Community



Stable HPC Platform Software that:

- Fuels a vibrant and efficient HPC software ecosystem
- Takes advantage of hardware innovation
 & drives revolutionary technologies
- Eases traditional HPC application development and testing at scale
- Extends to new workloads (ML, analytics, big data)
- Accommodates new environments (i.e., cloud)



OpenHPC: a brief History...

Current Project Members



Member participation interest? Contact <u>Neal Caidin</u> ncaidin@linuxfoundation.org

OpenHPC TSC – Individual Members

- Reese Baird, SpaceX (Maintainer)
- Aaron Blakeman, Intel (Maintainer)

New for 2019-2020

- David Brayford, LRZ (Maintainer)
- Eric Coulter, Indiana University (End-User/Site Representative)
- Chris Downing, Amazon Web Services (Maintainer)
- Craig Gardner, SUSE (Maintainer)
- Oguzhan Herkiloglu, Comodo (Maintainer)
- Michael Karo, Altair (Component Development Representative)
- Alex Lovell-Troy, Cray (Maintainer)
- Takayuki Okamoto, Fujitsu (Maintainer)
- Kevin Pedretti, Sandia National Laboratory (Maintainer)
- Nam Pho, University of Washington (Maintainer)
- Cyrus Proctor, Texas Advanced Computing Center (Maintainer)
- Adrian Reber, Red Hat (Maintainer)
- Karl W. Schulz, UT Austin (Project Lead, Testing Coordinator)
- Derek Simmel, Pittsburgh Supercomputing Center (End-User/Site Representative)
- Chris Simmons, UT Dallas (Maintainer)
- Ashish K Singh, Dell (Maintainer)
- Raja Subramani, Dell (Maintainer)
- Nirmala Sundararajan, Dell (Maintainer)

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





OpenHPC Stack Overview

Operator Interface					Applications						
Fabric Mgmt	Node-level Diags System Diags	Provisioning	User and Software Mgmt	Hardware Monitor And Control	Data Access	Workload Manager	Optimized I/O Libraries	Scalable Debugging And Performance Analysis Tools	High Performance Parallel Libraries	Co frances	Base OS User Space Utilities
					Data Store	Resource Mgmt Runtimes	Pe I/O Services		Compiler and Programming- Model Runtimes	Software Development Toolchain	
			Overlay & Pub-Sub Networks, Identity								
			Base OS Runtime Libraries								
			Node-specific Kernel(s), mOS, Base OS								
Hardware											

Released June 11, 2019

OpenHPC v1.3.8 - S/W components

Functional Areas	Components	components available 89	updates	36%			
Base OS	CentOS 7.6, SLES12 SP4			new with v1.3.8			
Architecture	aarch64, x86_64		I				
Administrative Tools	Conman, Ganglia, Lmod, LosF, Nagio Genders, Shine, Spack, test-suite	os, NHC, pdsh, pdsh-mod-slurm	, prun, EasyE	Build, ClusterShell, mrsh,			
Provisioning	Warewulf, xCAT	•	3 rd Party libraries are built for each compiler/MP				
Resource Mgmt.	SLURM, Munge, PBS Professional, P						
Runtimes	Charliecloud, OpenMP, OCR, Singula						
I/O Services	Lustre client, BeeGFS client*					family	
Numerical/Scientific Libraries	Boost, GSL, FFTW, Hypre, Metis, MFEM, Mumps, OpenBLAS, OpenCoarrays, PETSc, PLASMA, Scalapack, Scotch, SLEPc, SuperLU, SuperLU_Dist, Trilinos					Resulting	
I/O Libraries	HDF5 (pHDF5), NetCDF/pNetCDF (in		repositories currently comprised of ~700 RPMs				
Compiler Families	GNU (gcc, g++, gfortran), Clang/LLVI						
MPI Families	MVAPICH2, OpenMPI, MPICH, Intel						
Development Tools	Autotools, cmake, hwloc, mpi4py, R						
Performance Tools	PAPI, IMB, Likwid, mpiP, pdtoolkit T. Paraver, OSU Microbenchmarks						

v1.3.8.1 released in August with two small updates (MPICH and SLURM)

Target System Architecture (for bare-metal)



Intel driving capability-class open source software compatible with OpenHPC with intent to contribute

- mOS
 - Scalable operating systems
- Unified Control System
 - Unified, Productive (single pane of glass), Reliable
- MPI
 - Scalable, high performance, topology optimized
- GEOPM
 - Global Extensible Open Power Manager
- PMIx
 - Process management with "Instant On"
- DAOS
 - Distribute Asynchronous Object Store

Project Adoption Growth



- Continued access/download growth since initial release at SC'15
- Plots highlight number of unique visitors/month to the OpenHPC build server/repo(s)
- Over 21TB downloaded in 2018

OpenHPC: a building block repository

[Key takeaway]

- OpenHPC provides a collection of pre-built ingredients common in HPC environments; fundamentally it is a software <u>repository</u>
- The repository is published for use with Linux distro package managers:
 MPI stack(s)
 Admin Tools



tack(s) A

dmin Tools



• zypper (SLES)

Scientific Libs

- You can pick relevant bits of interest for your site
 - if you prefer a resource manager that is not included, you can build that locally and still leverage the scientific libraries and development environment
 - similarly, you might prefer to utilize a different provisioning system
 - public package repositories also makes it easy to include desired elements in customized containers (e.g. Docker, Singularity, Charliecloud)

openHPC

How to Request Additional Software?

- We have a simple submission site for new requests:
 - https://github.com/openhpc/submissions 0

0

- Example components added via this mechanism since the v1.2. release (Nov' 16)
 - BeeGFS client o 0
 - xCAT recipe 0
 - hwloc 0
 - 0
 - LLVM/clang 0
 - PLASMA 0

openhpc

- NHC pNetCDF SCOTCH
 - Charliecloud
- MPI4py
 - Likwid
 - MFEM
- SLEPc GeoPM Singularity o PMIx o Dimemas/Extrae,
 - Paraver
 - OpenCoarrays
 - OSU Benchmarks

Components reviewed on a rolling quarterly basis

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	
If other, please describe:	



How E4S and OpenHPC Work Together

• OpenHPC

- Provides installation and base components for installing and running cluster to supercomputer
- Provides generic binaries for all systems (working on targeted builds)

• E4S

- Built from scratch for your system
- Focus on upper layers of system software stack (libraries, runtimes, etc.)
- Targeted for capability-class machines



Thank you for your time

- OpenHPC Home:
- Primary GitHub Site:
- Package Repositories:
- OBS Frontend:
- Component Submission:
- System Registry:
- CI Infrastructure:
- OpenHPC Wiki:
 - includes links to overview paper and past presentations
- Mailing Lists:

http://www.openhpc.community/support/mail-lists/



- http://www.openhpc.community/
- https://github.com/openhpc/ohpc
- http://build.openhpc.community/OpenHPC:/
- https://build.openhpc.community
- https://github.com/openhpc/submissions
- System Registration Form
- http://test.openhpc.community:8080
- https://github.com/openhpc/ohpc/wiki