E4S: Extreme-Scale Scientific Software Stack

https://e4s.io

Sameer Shende
Performance Research Laboratory, OACISS, U. Oregon

sameer@cs.uoregon.edu

Tutorial, NCSA Webinar.

Extreme-scale Scientific Software Stack (E4S)
https://e4s.io

- E4S is a community effort to provide open source software packages for developing, deploying, and running scientific applications on HPC platforms.
- E4S provides both source builds and containers of a broad collection of HPC software packages.
- E4S exists to accelerate the development, deployment and use of HPC software, lowering the barriers for HPC users.
- E4S provides containers and turn-key, from-source builds of 80+ popular HPC software packages:
  - MPI: MPICH and OpenMPI
  - Development tools: TAU, HPCToolkit, and PAPI
  - Math libraries: PETSc and Trilinos
  - Data and Viz tools: Adios, HDF5, and Paraview

Extreme-scale Scientific Software Stack (E4S) https://e4s.io

- Spack [http://spack.io] is the primary means for software delivery
- SDKs: collection of related ECP ST products where coordination across package teams will improve usability and practicies, and foster community growth among teams that develop similar and complimentary capabilities. An SDK involves several products.
- Containers of pre-built binaries of ECP ST products.
- Container runtimes supported
  - Docker: Dockerhub: exascaleproject/sdk:AHM19
  - Charliecloud
  - Shifter
  - Singularity
  - Inception at NCAR
- VirtualBox Open Virtualization Appliance (OVA) image that contains these runtimes
- MPI replacement strategies to use native network interconnect

What are containers

- A lightweight collection of executable software that encapsulates everything needed to run a single specific task
  - Minus the OS kernel
  - Based on Linux only
- Processes and all user-level software is isolated
- Creates a portable* software ecosystem
- Think `chroot` on steroids
- Docker most common tool today
  - Available on all major platforms
  - Widely used in industry
  - Integrated container registry via Dockerhub

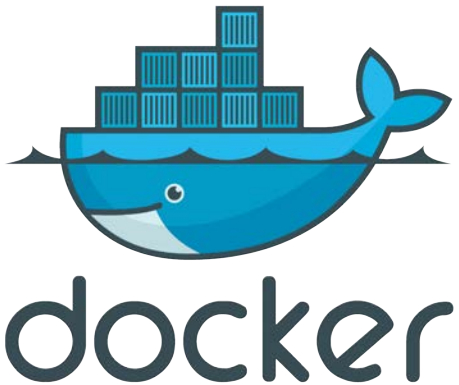

* Container slides from: Andrew Younge, Sandia, “Getting Started with Containers on HPC”, ISC-HPC 2019 tutorial

Hypervisors and Containers

- Type 1 hypervisors insert layer below host OS
- Type 2 hypervisors work as or within the host OS
- Containers do not abstract hardware, instead provide “enhanced chroot” to create isolated environment
- Location of abstraction can have impact on performance
- All enable custom software stacks on existing hardware

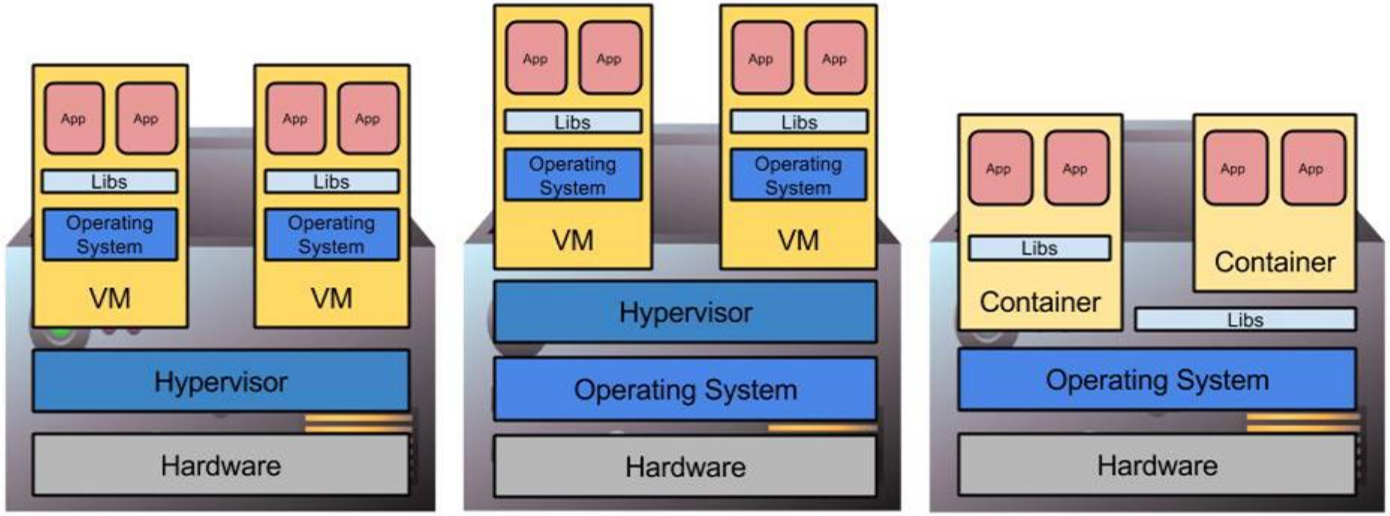

Type 1 Hypervisor Type 2 Hypervisor Containers

Background

- Abstracting hardware and software resources has had profound impact on computing
- Virtual Machines to Cloud computing in the past decade
  - Early implementations limited by performance
  - HPC on clouds: FutureGrid, Magellan, Chameleon Cloud, Hobbes, etc
  - Some initial successes, but not always straightforward
- OS-level virtualization a bit different
  - User level code packaged in container, can then be transported
  - Single OS kernel shared across containers and provides isolation
  - Cgroups traditionally multiplexes hardware resources
  - Performance is good, but OS flexibility is limited

Containers in Cloud Industry

- Containers are used to create large-scale loosely coupled services
- Each container runs just 1 user process – “micro-services”
  - 3 httpd containers, 2 DBs, 1 logger, etc
- Scaling achieved through load balancers and service provisioning
- Jam many containers on hosts for increased system utilization
- Helps with dev-ops issues
  - Same software environment for developing and deploying
  - Only images changes are pushed to production, not whole new image (CoW).
  - Develop on laptop, push to production servers
  - Interact with github similar to developer code bases
  - Upload images to ”hub” or “repository” whereby they can just be pulled and provisioned

Containers

- Containers are gaining popularity for software management of distributed systems
- Enable way for developers to specify software ecosystem
- US DOE High Performance Computing (HPC) resources need to support emerging software stacks
  - Applicable to DevOps problems seen with large HPC codes today
  - Support new frameworks & cloud platform services
- But HPC systems are very dissimilar from cloud infrastructure
  - MPI-based bulk synchronous parallel workloads are common
  - Scale-out to thousands of nodes
  - Performance is paramount

Container features in HPC

- **BYOE** - Bring-Your-Own-Environment
  - Developers define the operating environment and system libraries in which their application runs.
- **Composability**
  - Developers explicitly define how their software environment is composed of modular components as container images,
  - Enable reproducible environments that can potentially span different architectures.
- **Portability**
  - Containers can be rebuilt, layered, or shared across multiple different computing systems
  - Potentially from laptops to clouds to advanced supercomputing resources.
- **Version Control Integration**
  - Containers integrate with revision control systems like Git
  - Include not only build manifests but also with complete container images using container registries like Docker Hub.

Container features not wanted in HPC

- **Overhead**
  - HPC applications cannot incur significant overhead from containers
- **Micro-Services**
  - Micro-services container methodology does not apply to HPC workloads
  - 1 application per node with multiple processes or threads per container
- **On-node Partitioning**
  - On-node partitioning with cgroups is not necessary (yet?)
- **Root Operation**
  - Containers allow root-level access control to users
  - In supercomputers this is unnecessary and a significant security risk for facilities
- **Commodity Networking**
  - Containers and their network control mechanisms are built around commodity networking (TCP/IP)
  - Supercomputers utilize custom interconnects w/ OS kernel bypass operations

HPC Containers

- Docker not good fit for running HPC workloads
  - Security issues
    - Can't allow root on shared resources
  - Lack of HPC architecture support
    - No batch integration
    - Assumes local resources
    - Assumes commodity TCP/IP
- Many different container options in HPC

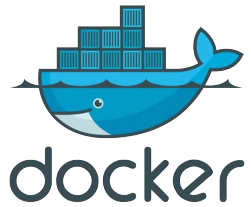

Shifter Singularity Charliecloud …

Spack

- E4S uses the Spack package manager for software delivery
- Spack provides the ability to specify versions of software packages that are and are not interoperable.
- Spack is a build layer for not only E4S software, but also a large collection of software tools and libraries outside of ECP ST.
- Spack supports achieving and maintaining interoperability between ST software packages.
- Acknowledgement: The remaining Spack slides in this presentation are from a talk given by the Spack PI, Todd Gamblin, CASC, LLNL.

- Next: Motivation for Spack!

Scientific software is becoming extremely complex

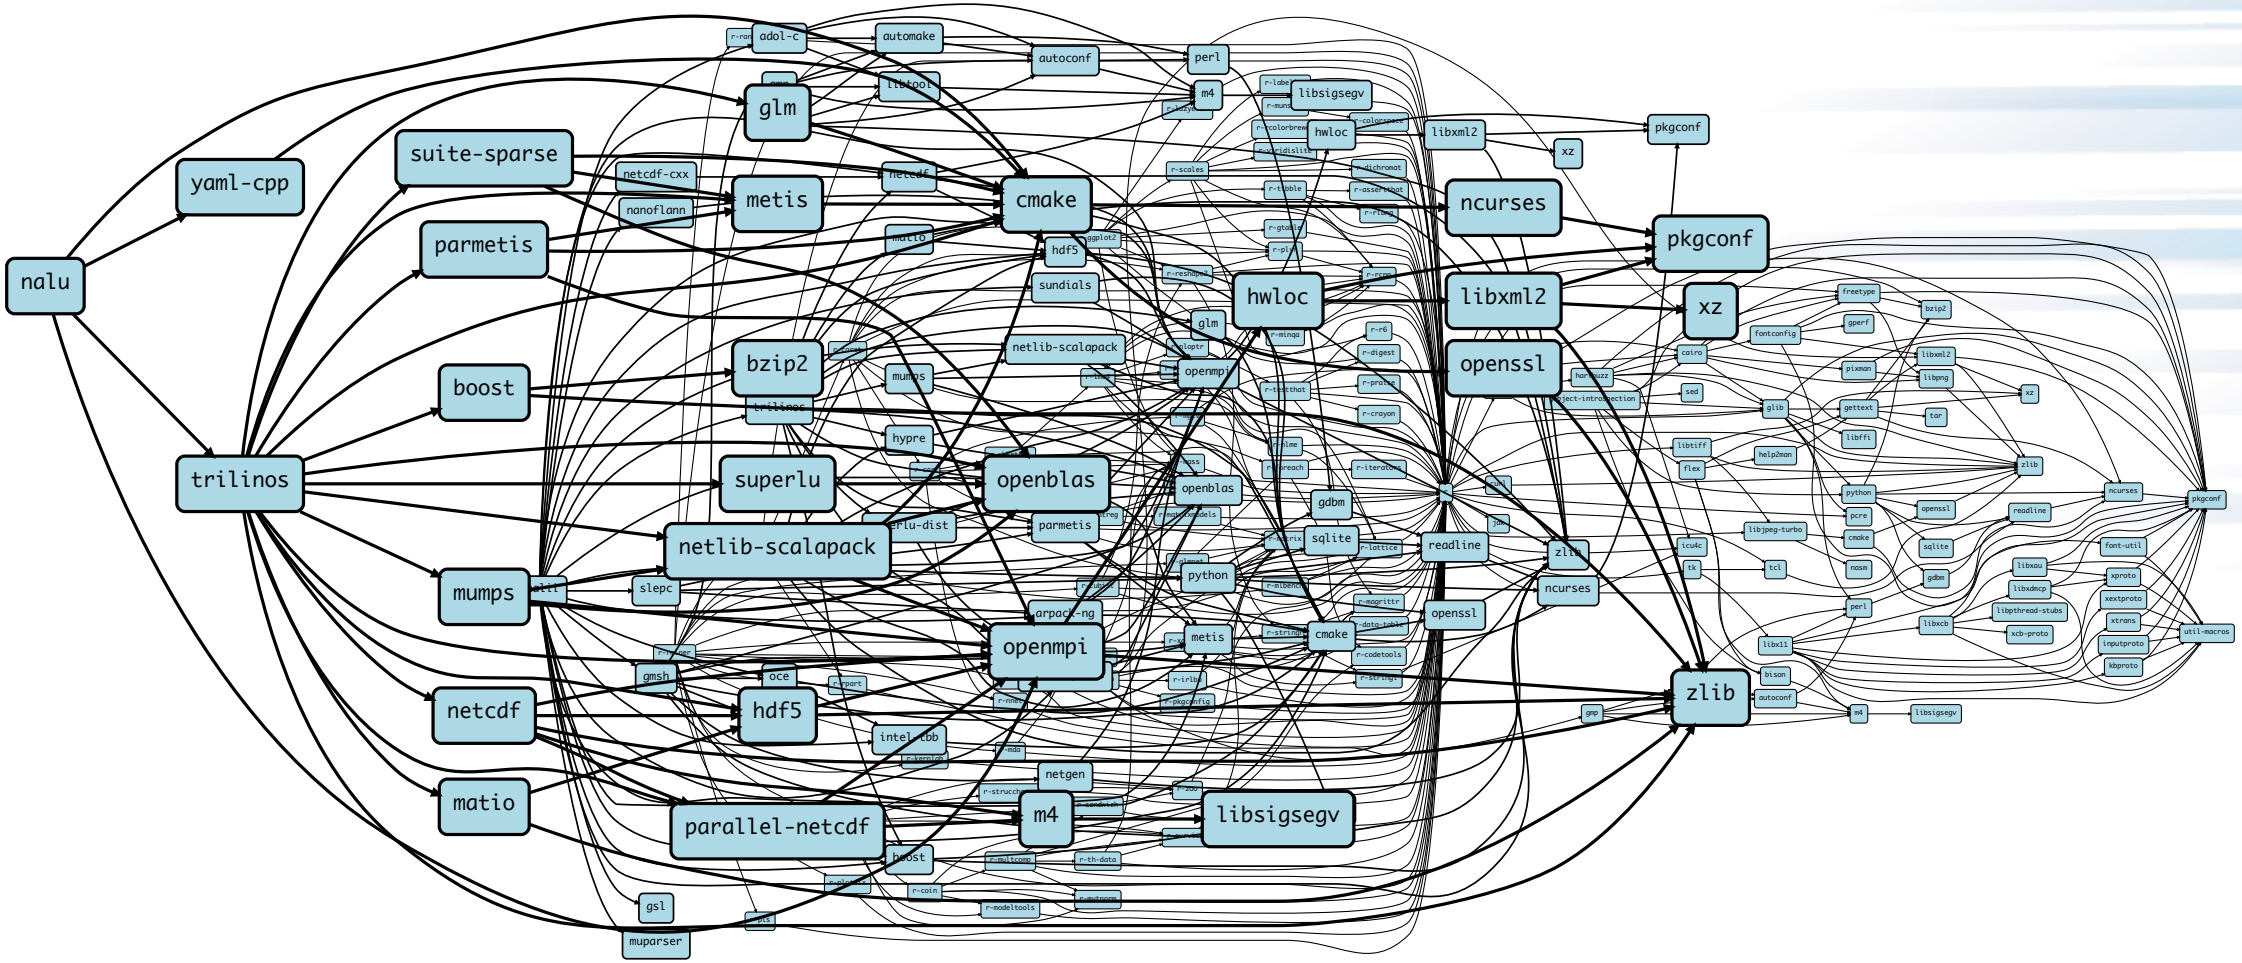

Nalu: Generalized Unstructured Mesh Navier-Stokes Wind Farm Modeling on Hybrid CPU/GPU Platforms

deal.II: C++ Finite Element Library

MarkyFlow: Data-Driven Path Mining

Even proprietary codes are based on many open source libraries

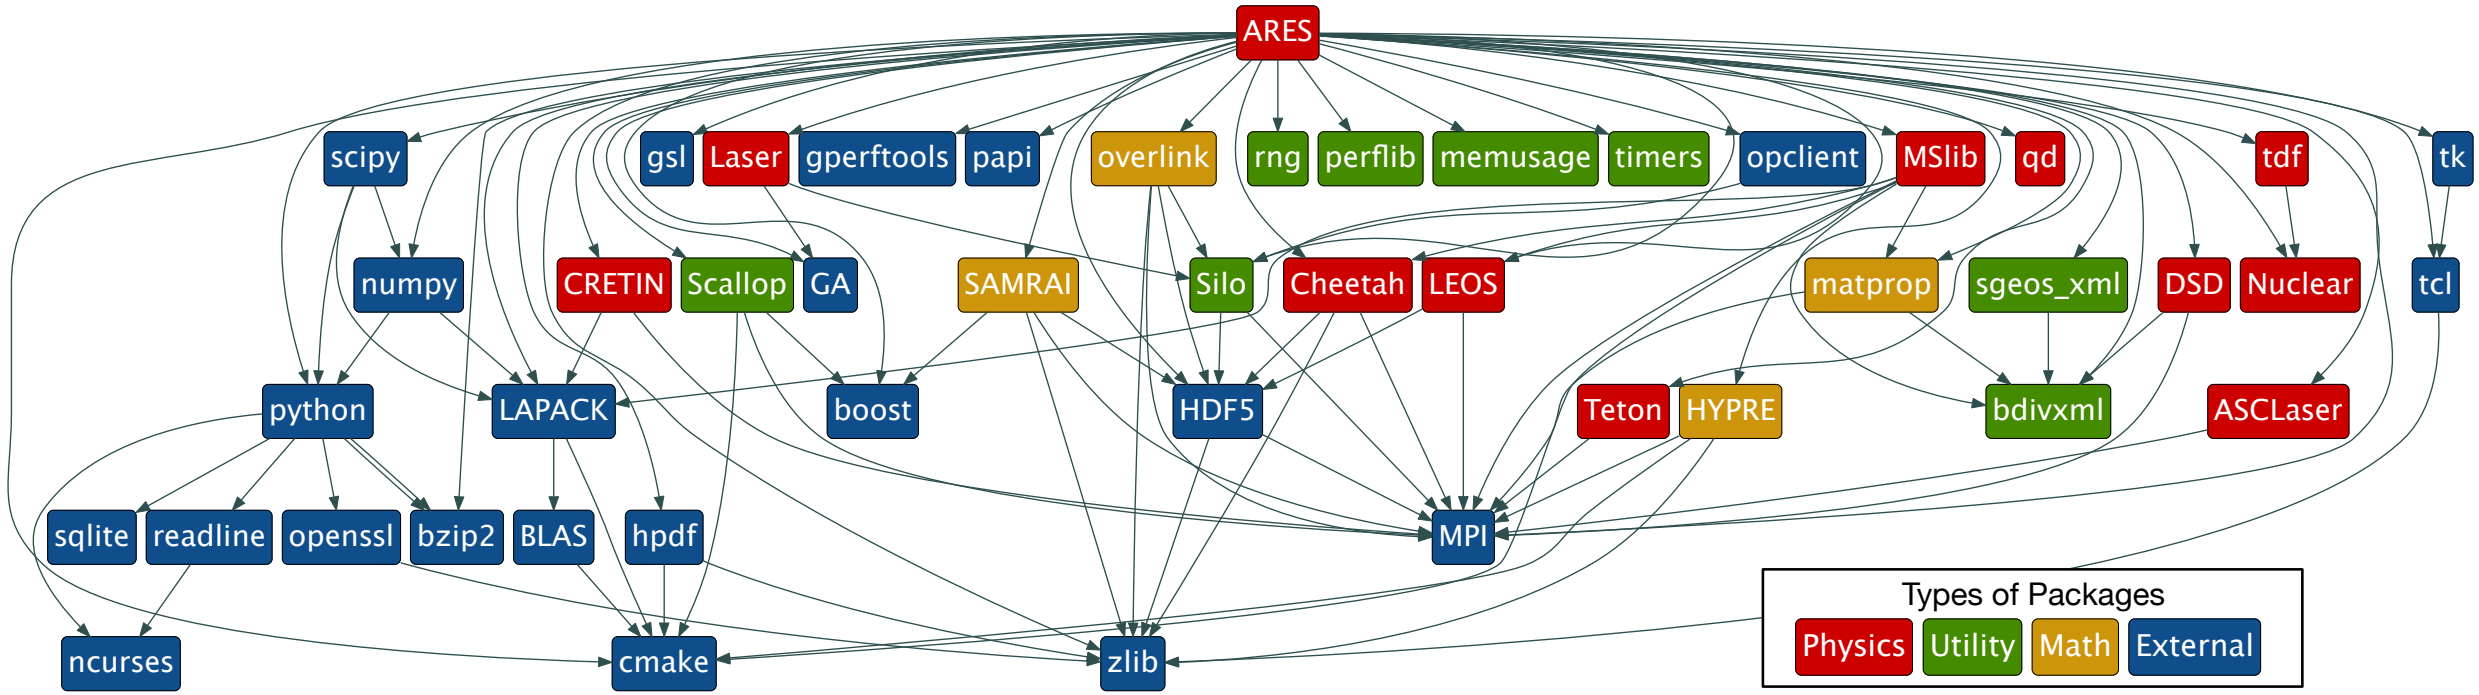

- Half of this DAG is external (blue); *more* than half of it is open source
- Nearly *all* of it needs to be built specially for HPC to get the best performance

The Exascale Computing Project is building an entire *ecosystem*

15+ applications x **80+ software packages** x **5+ target architectures/platforms**
Xeon Power KNL
NVIDIA ARM Laptops?

x **Up to 7 compilers**
Intel GCC Clang XL
PGI Cray NAG

x **10+ Programming Models**
OpenMPI MPICH MVAPICH OpenMP CUDA
OpenACC Dharma Legion RAJA Kokkos

x **2-3 versions of each package + external dependencies**

= up to **1,260,000** combinations!

- Every application has its own stack of dependencies.
- Developers, users, and facilities dedicate (many) FTEs to building & porting.
- Often trade reuse and usability for performance.

We must make it easier to rely on others' software!

How to install software on a Mac laptop, circa 2013

How to install software on a supercomputer

1. Download all 16 tarballs you need
2. Start building!

3. Run code
4. **Segfault!?**
5. Start over…

What about modules?

- Most supercomputers deploy some form of *environment modules*
  - TCL modules (dates back to 1995) and Lmod (from TACC) are the most popular

```
$ gcc
-bash: gcc: command not found

$ module load gcc/7.0.1
$ gcc –dumpversion
7.0.1
```

- Modules don’t handle installation!
  - They only modify your environment (things like PATH, LD_LIBRARY_PATH, etc.)
- Someone (likely a team of people) has already installed gcc for you!
  - Also, you can *only* `module load` the things they’ve installed

What about containers?

- **Containers provide a great way to reproduce and distribute an already-built software stack**
- **Someone needs to build the container!**
  - This isn't trivial
  - Containerized applications still have hundreds of dependencies
- **Using the OS package manager inside a container is insufficient**
  - Most binaries are built unoptimized
  - Generic binaries, not optimized for specific architectures
- **Developing with an OS software stack can be painful**
  - Little freedom to choose versions

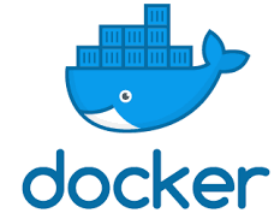

We need something more flexible to **build** the containers

Spack is a flexible package manager for HPC

- How to install Spack (works out of the box):

```
$ git clone https://github.com/spack/spack
$ . spack/share/spack/setup-env.sh
```

- How to install a package:

```
$ spack install hdf5
```

- HDF5 and its dependencies are installed within the Spack directory.
- Unlike typical package managers, Spack can also install many variants of the same build.
  - Different compilers
  - Different MPI implementations
  - Different build options

Visit spack.io

github.com/spack/spack

@spackpm

EXASCALE COMPUTING PROJECT

Spack provides the *spec* syntax to describe custom configurations

```
$ spack install mpileaks                          unconstrained
$ spack install mpileaks@3.3                      @ custom version
$ spack install mpileaks@3.3 %gcc@4.7.3           % custom compiler
$ spack install mpileaks@3.3 %gcc@4.7.3 +threads  +/- build option
$ spack install mpileaks@3.3 cxxflags="-O3 -g3"   setting compiler flags
$ spack install mpileaks@3.3 os=cnl10 target=haswell  setting target for X-compile
$ spack install mpileaks@3.3 ^mpich@3.2 %gcc@4.9.3    ^ dependency information
```

- Each expression is a ***spec*** for a particular configuration
  - Each clause adds a constraint to the spec
  - Constraints are optional – specify only what you need.
  - Customize install on the command line!
- Spec syntax is recursive
  - Full control over the combinatorial build space

`spack list` shows what packages are available

```
$ spack list
==> 3041 packages.
abinit                  glew                    nalu                    py-fastaindex           r-cairo                 r-viridislite
abyss                   glfmultiples            nalu-wind               py-fasteners            r-callr                 r-visnetwork
accfft                  glib                    namd                    py-faststructure        r-car                   r-vsn
ack                     glibmm                  nano                    py-filelock             r-caret                 r-webshot
activeharmony           glimmer                 nanoflann               py-fiona                r-category              r-whisker
adept-utils             glm                     nanopb                  py-fiscalyear           r-catools               r-withr
adios                   global                  nasm                    py-flake8               r-cdcfluview            r-xde
adios2                  globalarrays            nauty                   py-flake8-polyfill      r-cellranger            r-xgboost
adlbx                   globus-toolkit          ncbi-magicblast         py-flask                r-checkmate             r-xlconnect
adol-c                  glog                    ncbi-rmblastn           py-flask-compress       r-checkpoint            r-xlconnectjars
aegean                  gloo                    ncbi-toolkit            py-flask-socketio       r-chemometrics          r-xlsx
aida                    glpk                    nccl                    py-flexx                r-chron                 r-xlsxjars
albany                  glproto                 nccmp                   py-fn                   r-circlize              r-xmapbridge
albert                  glvis                   ncdu                    py-fparser              r-class                 r-xml
alglib                  gmake                   ncftp                   py-funcsigs             r-classint              r-xml2
allinea-forge           gmap-gsnap              ncl                     py-functools32          r-cli                   r-xnomial
allinea-reports         gmime                   nco                     py-future               r-clipr                 r-xtable
allpaths-lg             gmodel                  ncurses                 py-futures              r-cluster               r-xts
alquimia                gmp                     ncview                  py-fypp                 r-clustergeneration     r-xvector
alsa-lib                gmsh                    ndiff                   py-gdbgui               r-clusterprofiler       r-yaml
aluminum                gmt                     nek5000                 py-genders              r-cner                  r-yapsa
amg                     gnat                    nekbone                 py-genshi               r-coda                  r-yaqcaffy
amg2013                 gnu-prolog              nekcem                  py-geopandas            r-codetools             r-yarn
amp                     gnupg                   nektar                  py-gevent               r-coin                  r-zlibbioc
ampliconnoise           gnuplot                 neovim                  py-git-review           r-colorspace            r-zoo
amrex                   gnutls                  nest                    py-git2                 r-combinat              r3d
amrvis                  go                      netcdf                  py-gnuplot              r-complexheatmap        racon
andi                    go-bootstrap            netcdf-cxx              py-goatools             r-compositions          raft
angsd                   gobject-introspection   netcdf-cxx4             py-gpaw                 r-convevol              ragel
ant                     googletest              netcdf-fortran          py-greenlet             r-corhmm                raja
antlr                   gotcha                  netgauge                py-griddataformats      r-corpcor               randfold
ants                    gource                  netgen                  py-guidata              r-corrplot              random123
ape                     gperf                   netlib-lapack           py-guiqwt               r-covr                  randrproto
. . .
```

- Spack has over 3,000 builtin package recipes.

`spack find` shows what is installed

```
$ spack find
==> 103 installed packages.
-- linux-rhel7-x86_64 / gcc@4.4.7 ------------------------------
ImageMagick@6.8.9-10  glib@2.42.1        libtiff@4.0.3     pango@1.36.8          qt@4.8.6
SAMRAI@3.9.1          graphlib@2.0.0     libtool@2.4.2     parmetis@4.0.3        qt@5.4.0
adept-utils@1.0       gtkplus@2.24.25    libxcb@1.11       pixman@0.32.6         ravel@1.0.0
atk@2.14.0            harfbuzz@0.9.37    libxml2@2.9.2     py-dateutil@2.4.0     readline@6.3
boost@1.55.0          hdf5@1.8.13        llvm@3.0          py-ipython@2.3.1      scotch@6.0.3
cairo@1.14.0          icu@54.1           metis@5.1.0       py-nose@1.3.4         starpu@1.1.4
callpath@1.0.2        jpeg@9a            mpich@3.0.4       py-numpy@1.9.1        stat@2.1.0
dyninst@8.1.2         libdwarf@20130729  ncurses@5.9       py-pytz@2014.10       xz@5.2.0
dyninst@8.1.2         libelf@0.8.13      ocr@2015-02-16    py-setuptools@11.3.1  zlib@1.2.8
fontconfig@2.11.1     libffi@3.1         openssl@1.0.1h    py-six@1.9.0
freetype@2.5.3        libmng@2.0.2       otf@1.12.5salmon  python@2.7.8
gdk-pixbuf@2.31.2     libpng@1.6.16      otf2@1.4          qhull@1.0

-- linux-rhel7-x86_64 / gcc@4.8.2 ------------------------------
adept-utils@1.0.1  boost@1.55.0  cmake@5.6-special  libdwarf@20130729  mpich@3.0.4
adept-utils@1.0.1  cmake@5.6     dyninst@8.1.2      libelf@0.8.13      openmpi@1.8.2

-- linux-rhel7-x86_64 / intel@14.0.2 ---------------------------
hwloc@1.9  mpich@3.0.4  starpu@1.1.4

-- linux-rhel7-x86_64 / intel@15.0.0 ---------------------------
adept-utils@1.0.1  boost@1.55.0  libdwarf@20130729  libelf@0.8.13  mpich@3.0.4

-- linux-rhel7-x86_64 / intel@15.0.1 ---------------------------
adept-utils@1.0.1  callpath@1.0.2  libdwarf@20130729  mpich@3.0.4
boost@1.55.0       hwloc@1.9       libelf@0.8.13      starpu@1.1.4
```

- All the versions coexist!
  - Multiple versions of same package are ok.
- Packages are installed to automatically find correct dependencies.
- Binaries work *regardless of user's environment*.
- Spack also generates module files.
  - Don't *have* to use them.

The Spack community is growing rapidly

- **Spack simplifies HPC software for:**
  - Users
  - Developers
  - Cluster installations
  - The largest HPC facilities
- **Spack is central to ECP's software strategy**
  - Enable software reuse for developers and users
  - Allow the facilities to consume the entire ECP stack
- **The roadmap is packed with new features:**
  - Building the ECP software distribution
  - Better workflows for building containers
  - Stacks for facilities
  - Chains for rapid dev workflow
  - Optimized binaries
  - Better dependency resolution

Visit spack.io

github.com/spack/spack

@spackpm

EXASCALE COMPUTING PROJECT

Exascale Platform Preparation

- SDK Exascale platform preparation is focused on interoperable delivery.
- ST products from SDKs are released in the Extreme-scale Scientific Software Stack (E4S) [https://e4s.io].
  - E4S: a community effort to provide open source software packages for developing, deploying, and running scientific applications on HPC platforms
- E4S containers and Spack based builds currently support the following pre-exascale systems:
  - Theta at ALCF (Cray XC).
  - Cori at NERSC (Cray XC).
  - Summit, Sierra, Butte, RZAnsel (IBM Power 9 AC922).
  - Linux x86_64 systems at LANL (Grizzly), Sandia (Voltrino), LLNL (Quartz).
  - Other NSF platforms including Frontera (TACC).
- E4S preparation for future Exascale systems includes testing on AMD and Intel systems.

Integration and Interoperability: E4S

- E4S is released twice a year. Two versions have been released to date and we are planning for a release at SC19. The E4S 0.2 release supports:
  - Containers and turn-key, from-source builds of 80+ popular HPC software packages
  - 37 full release ECP ST products including:
    - MPI: MPICH and OpenMPI
    - Development tools: TAU, HPCToolkit, and PAPI
    - Math libraries: PETSc and Trilinos
    - Data and Viztools: Adios, HDF5, and Paraview
  - Limited access to 10 additional ECP ST products
  - Docker
  - Singularity
  - Shifter
  - Charliecloud
  - Inception
  - Open Virtualization Appliance (OVA) for VirtualBox features Spack, E4S containers, and support for container environments

Integration and Interoperability: E4S on AWS

- E4S AWS public image ami-063e830287b86155c (US-West-2 Oregon) has following container runtimes:
  - Docker
  - Shifter
  - Singularity
  - Charliecloud
- Spack with base PMR components
- E4S full featured Singularity image
  - (exascaleproject/sdk:AHM19)
- Used in ISC-HPC 2019 tutorials
- Used as base image for NASA GEOS-Chem E4S public image
- Resources provided by AWS AI/ML team

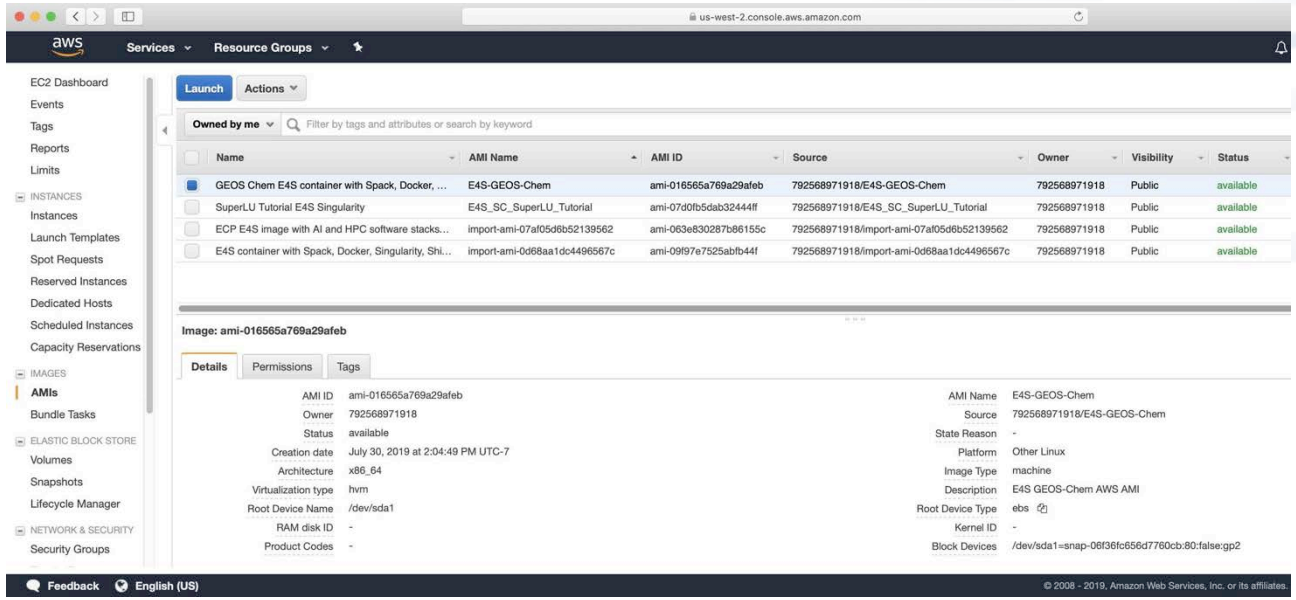

Reproducible, Customizable Container Builds & Spack Mirrors

- E4S provides base images and recipes for building Docker containers based on SDKs
  - Git: https://github.com/UO-OACISS/e4s
  - Base images released (September 2019):
    - UBI 7.6 (RHEL Universal Binary Image for container builds) for x86_64
    - Centos 7.6 for x86_64
    - Ubuntu 18.04 for x86_64
    - UBI 7.6 (RHEL) for ppc64le
- E4S provides build caches for Spack for native bare-metal as well as container builds based installation of ST products
  - Build caches: https://oaciss.uoregon.edu/e4s
    - The build cache model can be extended to target platforms, and can be managed by facilities staff when appropriate.

E4S Build Cache Binaries

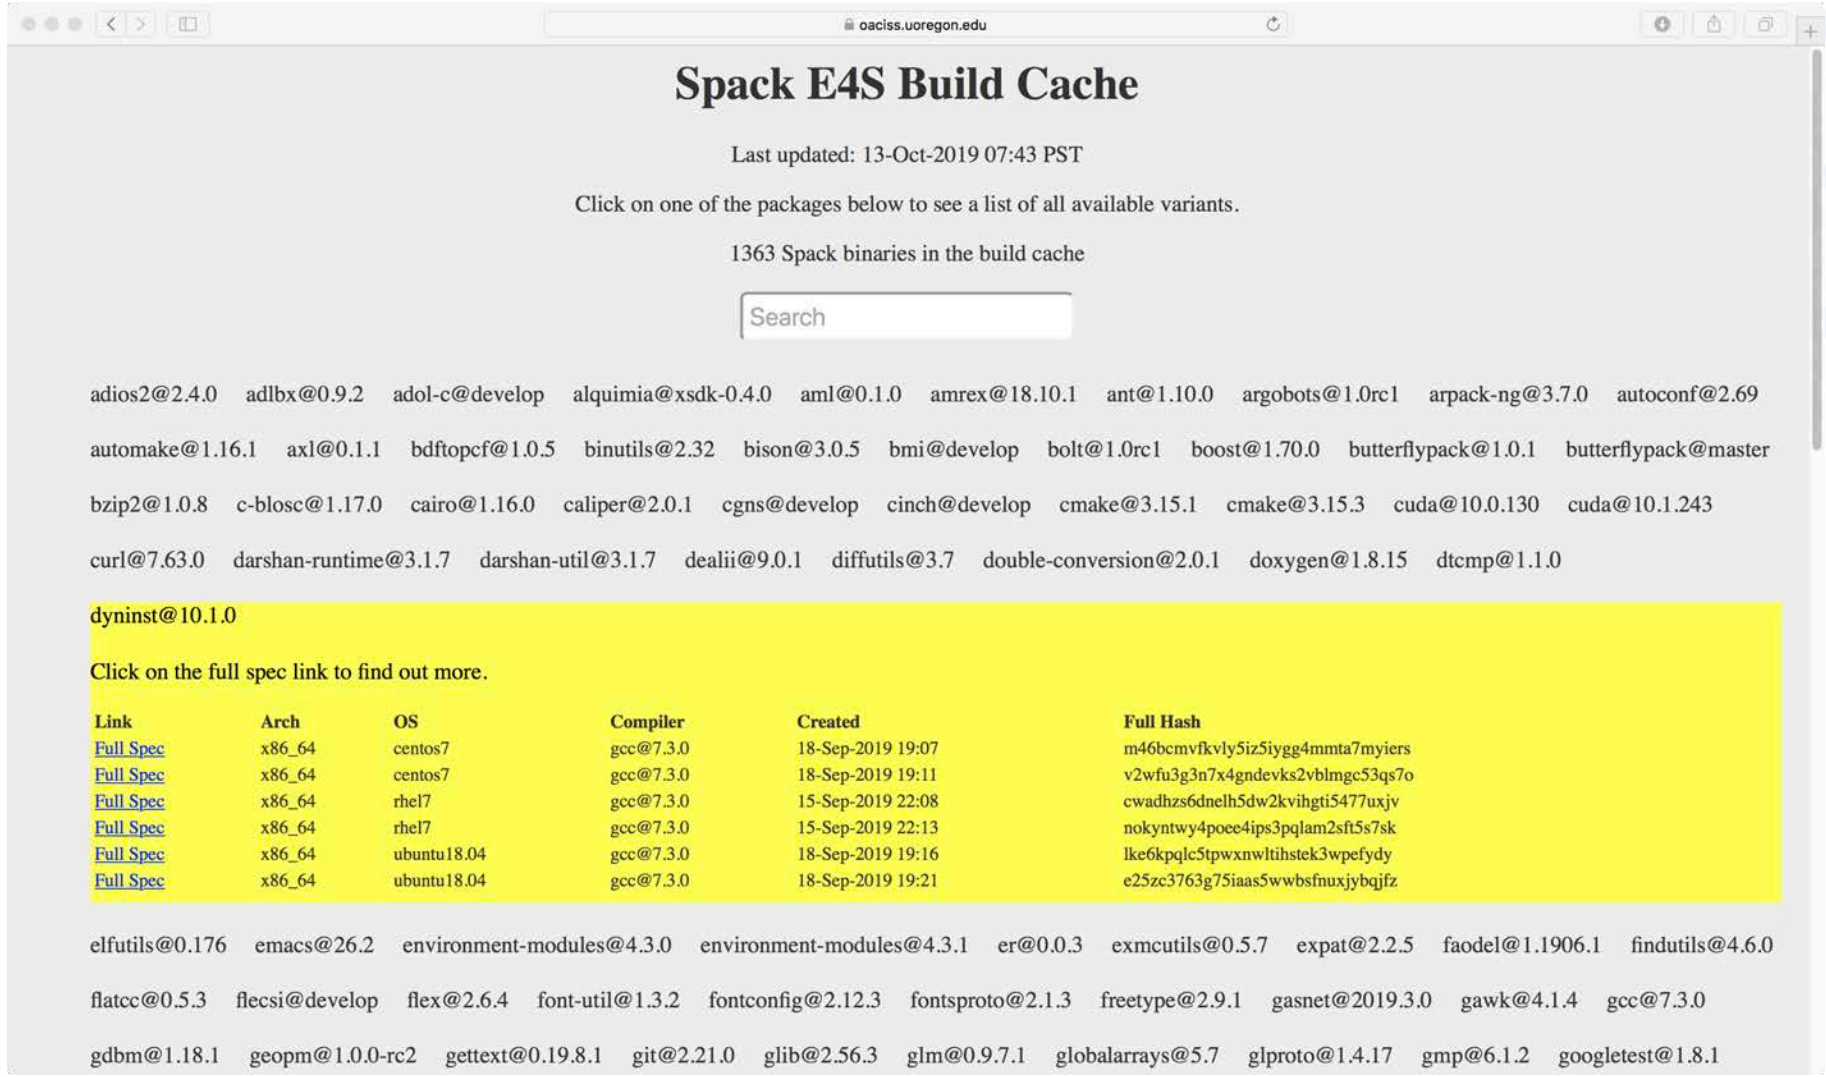

https://oaciss.uoregon.edu/e4s

Reproducible Container Builds using E4S Base Images

UO-OACISS / e4s

Unwatch 3 | Star 0 | Fork 0

Code | Issues 0 | Pull requests 0 | Projects 0 | Wiki | Security | Insights | Settings

Branch: master — e4s / docker / ubi7 / x86_64 / custom / superlu / spack.yaml — Find file | Copy path

sameershende Adding a custom image for SC19 SuperLU tutorial. a0b948d 10 days ago

1 contributor

35 lines (34 sloc) | 579 Bytes — Raw | Blame | History

```
spack:
  packages:
    all:
      compiler: [gcc@7.3.0]
      variants: +mpi
      providers:
        mpi: [mpich]
      buildable: true
      version: []
      paths: {}
      modules: {}
    mpich:
      version: [3.2.1]
      variants: ~wrapperrpath
      buildable: true
      providers: {}
      paths: {}
      modules: {}
      compiler: []
    gcc:
      version: [7.3.0]
      buildable: true
      providers: {}
      paths: {}
      modules: {}
      compiler: []
  specs:
  - superlu-dist
  - petsc
  - mfem
  - strumpack
  - butterflypack
  - openblas
  view: false
```

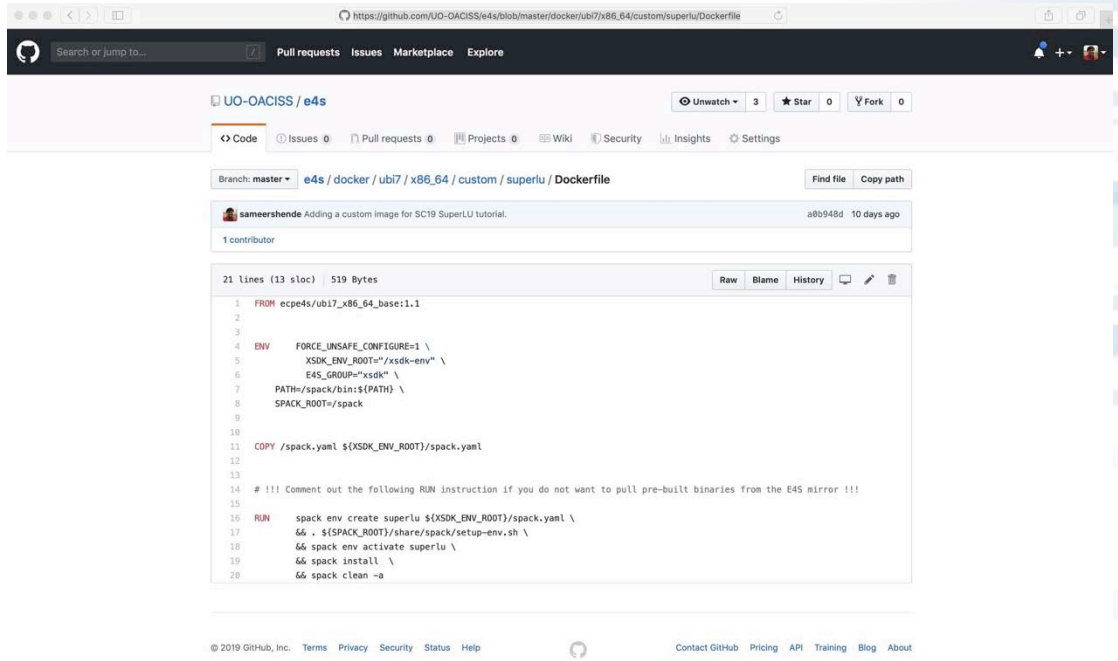

- PMR SDK base image (UBI 7.6) has Spack build cache mirror and GPG key installed.
- Base image has GCC and MPICH configured for MPICH ABI level replacement (with system MPI).
- Customized container build using binaries from E4S Spack build cache for fast deployment.
- No need to rebuild packages from the source code.
- Same recipe for container and native bare-metal builds with Spack!

Reproducible Base Images on Dockerhub

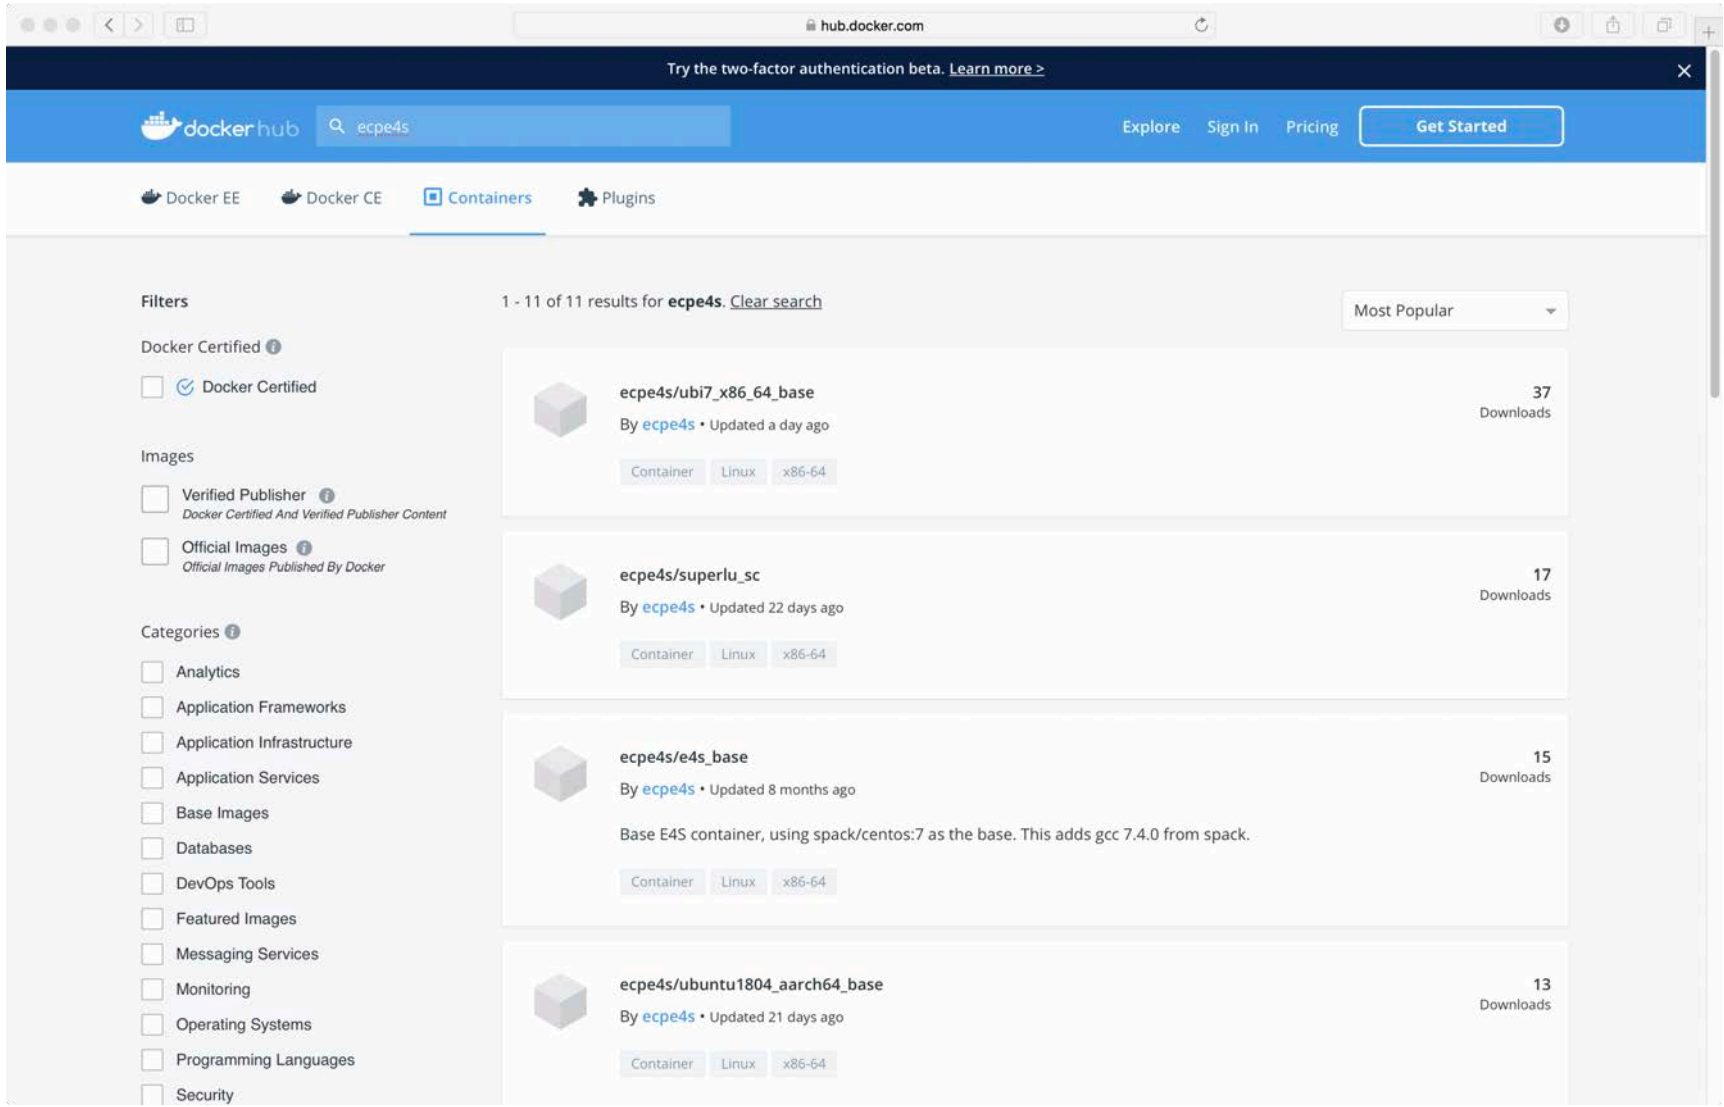

- ecpe4s

- x86_64
- ppc64le
- aarch64

- Centos 7.6
- Ubuntu 18.04
- RHEL/UBI 7.6

Docker Recipes on GitHub

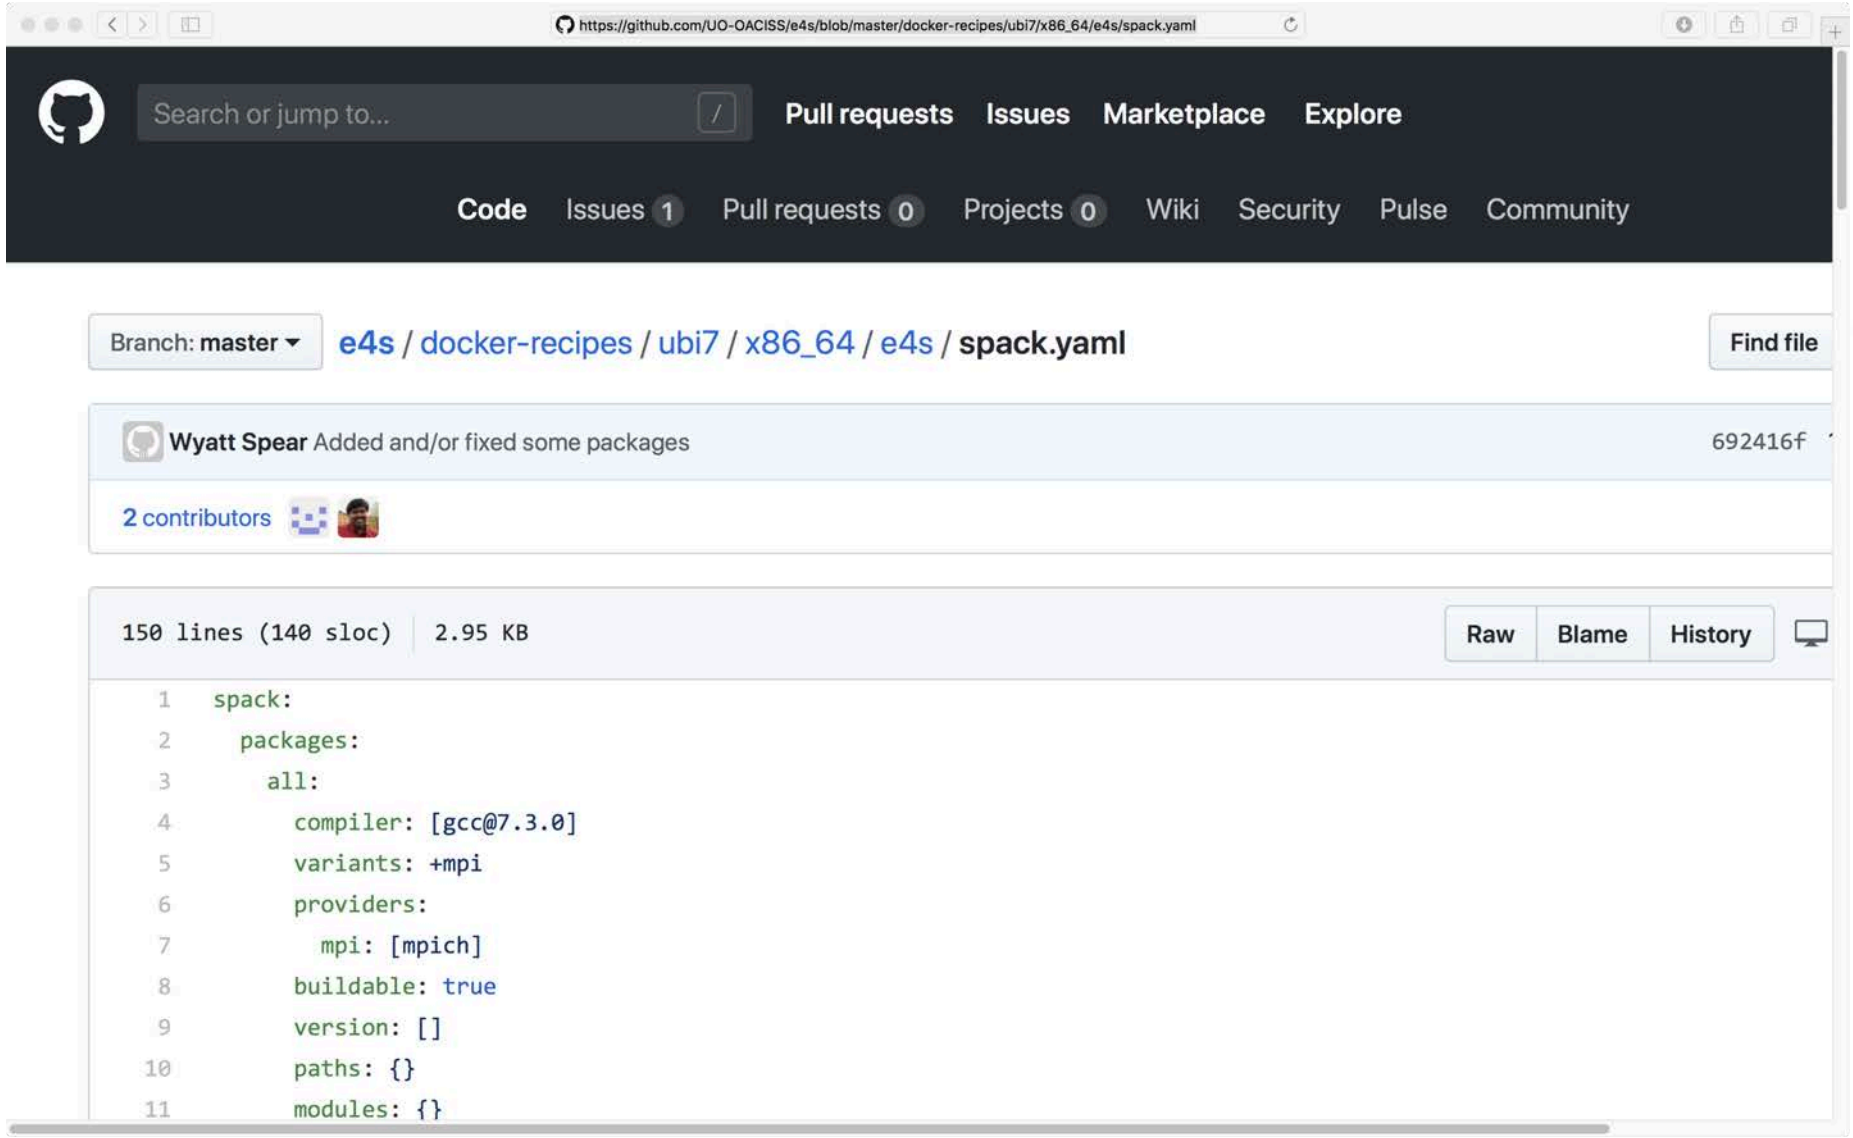

- Base images
- SDKs
- E4S

https://github.com/UO-OACISS/e4s

Spack Build Caches from E4S Base Images

https://oaciss.uoregon.edu/e4s/x86_64/build_cache/linux-rhel7-x86_64/gcc-7.3.0/

Index of /e4s/x86_64/build_cache/linux-rhel7-x86_64/gcc-7.3.0

Name	Last modified	Size	Description
Parent Directory		-	
adios2-2.4.0/	19-Sep-2019 06:25	-	
arpack-ng-3.7.0/	12-Sep-2019 12:35	-	
autoconf-2.69/	17-Sep-2019 17:24	-	
automake-1.16.1/	17-Sep-2019 17:31	-	
axl-0.1.1/	19-Sep-2019 06:28	-	
binutils-2.32/	15-Sep-2019 21:57	-	
bison-3.0.5/	28-Aug-2019 08:17	-	
boost-1.70.0/	19-Sep-2019 06:16	-	
butterflypack-1.0.1/	12-Sep-2019 12:35	-	
bzip2-1.0.8/	17-Sep-2019 17:47	-	
c-blosc-1.17.0/	27-Aug-2019 09:04	-	
caliper-2.0.1/	06-Sep-2019 08:41	-	
cinch-develop/	28-Aug-2019 18:57	-	
cmake-3.15.1/	17-Sep-2019 17:57	-	
cmake-3.15.3/	18-Sep-2019 13:17	-	
cuda-10.0.130/	05-Sep-2019 12:59	-	
cuda-10.1.243/	11-Sep-2019 13:42	-	
curl-7.63.0/	17-Sep-2019 18:03	-	
darshan-runtime-3.1.7/	19-Sep-2019 06:22	-	
darshan-util-3.1.7/	27-Aug-2019 09:02	-	
diffutils-3.7/	18-Sep-2019 14:23	-	
dtcmp-1.1.0/	19-Sep-2019 06:28	-	
dyninst-10.1.0/	15-Sep-2019 22:13	-	
elfutils-0.176/	15-Sep-2019 21:58	-	
environment-modules-4.3.0/	17-Sep-2019 17:59	-	
er-0.0.3/	19-Sep-2019 06:29	-	
expat-2.2.5/	17-Sep-2019 17:24	-	
faodel-1.1906.1/	19-Sep-2019 06:27	-	
findutils-4.6.0/	17-Sep-2019 17:24	-	
flecsi-develop/	18-Sep-2019 14:44	-	
flex-2.6.4/	28-Aug-2019 08:17	-	
freetype-2.9.1/	27-Aug-2019 11:48	-	
gasnet-2019.3.0/	27-Aug-2019 10:18	-	
gdbm-1.18.1/	17-Sep-2019 17:47	-	
gettext-0.19.8.1/	17-Sep-2019 17:26	-	
git-2.21.0/	17-Sep-2019 18:03	-	
glm-0.9.7.1/	29-Aug-2019 12:43	-	
glproto-1.4.17/	27-Aug-2019 11:43	-	
googletest-1.8.1/	27-Aug-2019 09:09	-	
gotcha-1.0.2/	10-Sep-2019 13:01	-	
hdf5-1.10.5/	19-Sep-2019 06:21	-	
help2man-1.47.8/	28-Aug-2019 08:17	-	
hpctoolkit-2019.08.14/	15-Sep-2019 22:17	-	
hwloc-1.11.11/	15-Sep-2019 18:27	-	
hypre-2.16.0/	17-Aug-2019 09:08	-	
hypre-2.17.0/	29-Aug-2019 12:55	-	
inputproto-2.3.2/	27-Aug-2019 11:48	-	
intel-mkl-2019.3.199/	05-Sep-2019 13:09	-	
intel-tbb-2019.4/	10-Sep-2019 12:49	-	
intel-xed-2019.03.01/	10-Sep-2019 12:57	-	
kbproto-1.0.7/	28-Aug-2019 08:27	-	
kokkos-2.8.00/	10-Aug-2019 07:35	-	
kokkos-2.9.00/	27-Aug-2019 10:16	-	
kvtree-1.0.2/	19-Sep-2019 06:25	-	
legion-19.06.0/	27-Aug-2019 10:19	-	
libbsd-0.9.1/	17-Sep-2019 17:30	-	
libdwarf-20180129/	10-Sep-2019 12:57	-	
libedit-3.1-20170329/	27-Aug-2019 11:48	-	
libelf-0.8.13/	10-Sep-2019 12:48	-	
libfabric-1.8.0/	27-Aug-2019 09:04	-	
libffi-3.2.1/	05-Sep-2019 13:15	-	
libiberty-2.31.1/	10-Sep-2019 12:52	-	
libice-1.0.9/	27-Aug-2019 11:48	-	
libiconv-1.15/	17-Sep-2019 17:59	-	
libjpeg-turbo-2.0.2/	28-Aug-2019 08:28	-	
libmonitor-2018.07.18/	10-Sep-2019 12:48	-	
libpciaccess-0.13.5/	17-Sep-2019 17:30	-	
libpfm4-4.10.1/	10-Sep-2019 13:02	-	
libpng-1.6.34/	27-Aug-2019 09:03	-	
libpthread-stubs-0.4/	27-Aug-2019 09:03	-	
libquo-1.3/	27-Aug-2019 10:16	-	
libsigsegv-2.11/	17-Sep-2019 17:54	-	
libsm-1.2.2/	27-Aug-2019 11:43	-	
libsodium-1.0.17/	27-Aug-2019 09:02	-	

- x86_64 build cache
  - 40 GB on disk

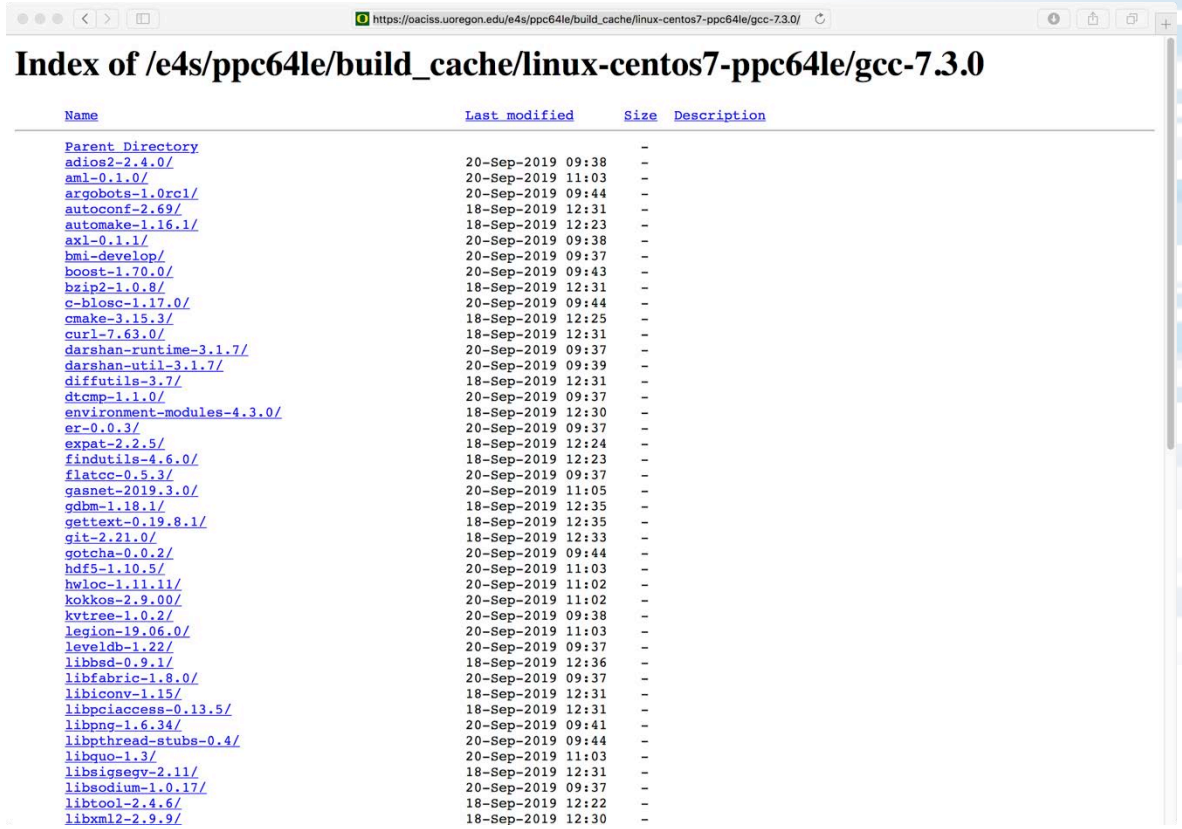

https://oaciss.uoregon.edu/e4s/ppc64le/build_cache/linux-centos7-ppc64le/gcc-7.3.0/

Index of /e4s/ppc64le/build_cache/linux-centos7-ppc64le/gcc-7.3.0

Name	Last modified	Size	Description
Parent Directory		-	
adios2-2.4.0/	20-Sep-2019 09:38	-	
aml-0.1.0/	20-Sep-2019 11:03	-	
argobots-1.0rc1/	20-Sep-2019 09:44	-	
autoconf-2.69/	18-Sep-2019 12:31	-	
automake-1.16.1/	18-Sep-2019 12:23	-	
axl-0.1.1/	20-Sep-2019 09:38	-	
bmi-develop/	20-Sep-2019 09:37	-	
boost-1.70.0/	20-Sep-2019 09:43	-	
bzip2-1.0.8/	18-Sep-2019 12:31	-	
c-blosc-1.17.0/	20-Sep-2019 09:44	-	
cmake-3.15.3/	18-Sep-2019 12:25	-	
curl-7.63.0/	18-Sep-2019 12:31	-	
darshan-runtime-3.1.7/	20-Sep-2019 09:37	-	
darshan-util-3.1.7/	20-Sep-2019 09:39	-	
diffutils-3.7/	18-Sep-2019 12:31	-	
dtcmp-1.1.0/	20-Sep-2019 09:37	-	
environment-modules-4.3.0/	18-Sep-2019 12:30	-	
er-0.0.3/	20-Sep-2019 09:37	-	
expat-2.2.5/	18-Sep-2019 12:24	-	
findutils-4.6.0/	18-Sep-2019 12:23	-	
flatcc-0.5.3/	20-Sep-2019 09:37	-	
gasnet-2019.3.0/	20-Sep-2019 11:05	-	
gdbm-1.18.1/	18-Sep-2019 12:35	-	
gettext-0.19.8.1/	18-Sep-2019 12:35	-	
git-2.21.0/	18-Sep-2019 12:33	-	
gotcha-0.0.2/	20-Sep-2019 09:44	-	
hdf5-1.10.5/	20-Sep-2019 11:03	-	
hwloc-1.11.11/	20-Sep-2019 11:02	-	
kokkos-2.9.00/	20-Sep-2019 11:02	-	
kvtree-1.0.2/	20-Sep-2019 09:38	-	
legion-19.06.0/	20-Sep-2019 11:03	-	
leveldb-1.22/	20-Sep-2019 09:37	-	
libbsd-0.9.1/	18-Sep-2019 12:36	-	
libfabric-1.8.0/	20-Sep-2019 09:37	-	
libiconv-1.15/	18-Sep-2019 12:31	-	
libpciaccess-0.13.5/	18-Sep-2019 12:31	-	
libpng-1.6.34/	20-Sep-2019 09:41	-	
libpthread-stubs-0.4/	20-Sep-2019 09:44	-	
libquo-1.3/	20-Sep-2019 11:03	-	
libsigsegv-2.11/	18-Sep-2019 12:31	-	
libsodium-1.0.17/	20-Sep-2019 09:37	-	
libtool-2.4.6/	18-Sep-2019 12:22	-	
libxml2-2.9.9/	18-Sep-2019 12:30	-	

- IBM Power 9 (ppc64le) build cache
  - 2.6 GB on disk
  - early stages of effort
  - Initial ARM 64 build cache is underway

Docker container of E4S

% docker pull exascaleproject/sdk:AHM19

- Using USB stick or images from https://e4s.io:
- % gunzip −c ecp.tgz | docker load
  % docker images
- Mount home directory:

% docker −i −v $HOME:$HOME −t exascaleproject/sdk:AHM19 /bin/bash

% which spack

% cp −r /usr/local/packages/ecp/demo . ; cd demo; cat README

Using Shifter at NCSA BlueWaters

Load shifter module and E4S image on the compute node

- Allocate a node
  - % qsub -I -l nodes=1:ppn=32 -l walltime=01:15:00 -l gres=shifter16
- Load the shifter module
  - % module load shifter
- Pull the image (once)
  - % shifterimg pull exascaleproject/sdk:AHM19
- Launch the image
  - % shifter --image=exascaleproject/sdk:AHM19 -- /bin/bash
  - % unset CRAYPE_VERSION; . /etc/bashrc
  - % spack find

Extreme-scale Scientific Software Stack (E4S) https://e4s.io

- Containers for HPC that include ECP ST products.

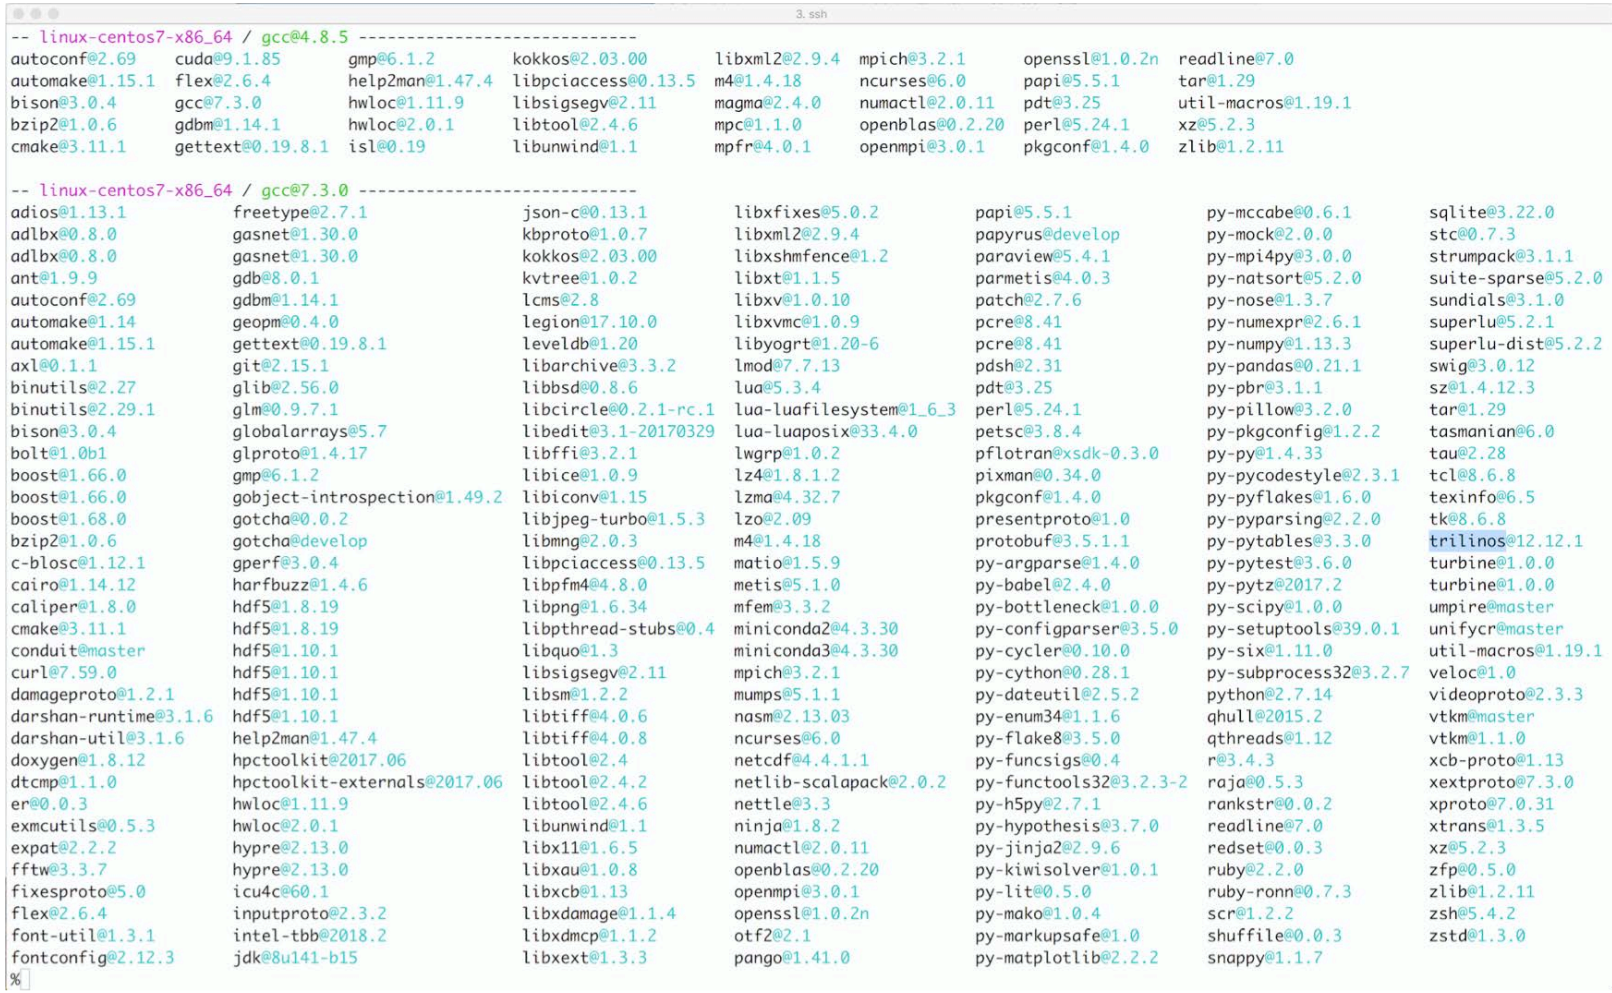

```
3. ssh
-- linux-centos7-x86_64 / gcc@4.8.5 ------------------------------
autoconf@2.69    cuda@9.1.85      gmp@6.1.2        kokkos@2.03.00      libxml2@2.9.4  mpich@3.2.1      openssl@1.0.2n  readline@7.0
automake@1.15.1  flex@2.6.4       help2man@1.47.4  libpciaccess@0.13.5 m4@1.4.18      ncurses@6.0      papi@5.5.1      tar@1.29
bison@3.0.4      gcc@7.3.0        hwloc@1.11.9     libsigsegv@2.11     magma@2.4.0    numactl@2.0.11   pdt@3.25        util-macros@1.19.1
bzip2@1.0.6      gdbm@1.14.1      hwloc@2.0.1      libtool@2.4.6       mpc@1.1.0      openblas@0.2.20  perl@5.24.1     xz@5.2.3
cmake@3.11.1     gettext@0.19.8.1 isl@0.19         libunwind@1.1       mpfr@4.0.1     openmpi@3.0.1    pkgconf@1.4.0   zlib@1.2.11

-- linux-centos7-x86_64 / gcc@7.3.0 ------------------------------
adios@1.13.1           freetype@2.7.1               json-c@0.13.1          libxfixes@5.0.2          papi@5.5.1             py-mccabe@0.6.1        sqlite@3.22.0
adlbx@0.8.0            gasnet@1.30.0                kbproto@1.0.7          libxml2@2.9.4            papyrus@develop        py-mock@2.0.0          stc@0.7.3
adlbx@0.8.0            gasnet@1.30.0                kokkos@2.03.00         libxshmfence@1.2         paraview@5.4.1         py-mpi4py@3.0.0        strumpack@3.1.1
ant@1.9.9              gdb@8.0.1                    kvtree@1.0.2           libxt@1.1.5              parmetis@4.0.3         py-natsort@5.2.0       suite-sparse@5.2.0
autoconf@2.69          gdbm@1.14.1                  lcms@2.8               libxv@1.0.10             patch@2.7.6            py-nose@1.3.7          sundials@3.1.0
automake@1.14          geopm@0.4.0                  legion@17.10.0         libxvmc@1.0.9            pcre@8.41              py-numexpr@2.6.1       superlu@5.2.1
automake@1.15.1        gettext@0.19.8.1             leveldb@1.20           libyogrt@1.20-6          pcre@8.41              py-numpy@1.13.3        superlu-dist@5.2.2
axl@0.1.1              git@2.15.1                   libarchive@3.3.2       lmod@7.7.13              pdsh@2.31              py-pandas@0.21.1       swig@3.0.12
binutils@2.27          glib@2.56.0                  libbsd@0.8.6           lua@5.3.4                pdt@3.25               py-pbr@3.1.1           sz@1.4.12.3
binutils@2.29.1        glm@0.9.7.1                  libcircle@0.2.1-rc.1   lua-luafilesystem@1_6_3  perl@5.24.1            py-pillow@3.2.0        tar@1.29
bison@3.0.4            globalarrays@5.7             libedit@3.1-20170329   lua-luaposix@33.4.0      petsc@3.8.4            py-pkgconfig@1.2.2     tasmanian@6.0
bolt@1.0b1             glproto@1.4.17               libffi@3.2.1           lwgrp@1.0.2              pflotran@xsdk-0.3.0    py-py@1.4.33           tau@2.28
boost@1.66.0           gmp@6.1.2                    libice@1.0.9           lz4@1.8.1.2              pixman@0.34.0          py-pycodestyle@2.3.1   tcl@8.6.8
boost@1.66.0           gobject-introspection@1.49.2 libiconv@1.15          lzma@4.32.7              pkgconf@1.4.0          py-pyflakes@1.6.0      texinfo@6.5
boost@1.68.0           gotcha@0.0.2                 libjpeg-turbo@1.5.3    lzo@2.09                 presentproto@1.0       py-pyparsing@2.2.0     tk@8.6.8
bzip2@1.0.6            gotcha@develop               libmng@2.0.3           m4@1.4.18                protobuf@3.5.1.1       py-pytables@3.3.0      trilinos@12.12.1
c-blosc@1.12.1         gperf@3.0.4                  libpciaccess@0.13.5    matio@1.5.9              py-argparse@1.4.0      py-pytest@3.6.0        turbine@1.0.0
cairo@1.14.12          harfbuzz@1.4.6               libpfm4@4.8.0          metis@5.1.0              py-babel@2.4.0         py-pytz@2017.2         turbine@1.0.0
caliper@1.8.0          hdf5@1.8.19                  libpng@1.6.34          mfem@3.3.2               py-bottleneck@1.0.0    py-scipy@1.0.0         umpire@master
cmake@3.11.1           hdf5@1.8.19                  libpthread-stubs@0.4   miniconda2@4.3.30        py-configparser@3.5.0  py-setuptools@39.0.1   unifycr@master
conduit@master         hdf5@1.10.1                  libquo@1.3             miniconda3@4.3.30        py-cycler@0.10.0       py-six@1.11.0          util-macros@1.19.1
curl@7.59.0            hdf5@1.10.1                  libsigsegv@2.11        mpich@3.2.1              py-cython@0.28.1       py-subprocess32@3.2.7  veloc@1.0
damageproto@1.2.1      hdf5@1.10.1                  libsm@1.2.2            mumps@5.1.1              py-dateutil@2.5.2      python@2.7.14          videoproto@2.3.3
darshan-runtime@3.1.6  hdf5@1.10.1                  libtiff@4.0.6          nasm@2.13.03             py-enum34@1.1.6        qhull@2015.2           vtkm@master
darshan-util@3.1.6     help2man@1.47.4              libtiff@4.0.8          ncurses@6.0              py-flake8@3.5.0        qthreads@1.12          vtkm@1.1.0
doxygen@1.8.12         hpctoolkit@2017.06           libtool@2.4            netcdf@4.4.1.1           py-funcsigs@0.4        r@3.4.3                xcb-proto@1.13
dtcmp@1.1.0            hpctoolkit-externals@2017.06 libtool@2.4.2          netlib-scalapack@2.0.2   py-functools32@3.2.3-2 raja@0.5.3             xextproto@7.3.0
er@0.0.3               hwloc@1.11.9                 libtool@2.4.6          nettle@3.3               py-h5py@2.7.1          rankstr@0.0.2          xproto@7.0.31
exmcutils@0.5.3        hwloc@2.0.1                  libunwind@1.1          ninja@1.8.2              py-hypothesis@3.7.0    readline@7.0           xtrans@1.3.5
expat@2.2.2            hypre@2.13.0                 libx11@1.6.5           numactl@2.0.11           py-jinja2@2.9.6        redset@0.0.3           xz@5.2.3
fftw@3.3.7             hypre@2.13.0                 libxau@1.0.8           openblas@0.2.20          py-kiwisolver@1.0.1    ruby@2.2.0             zfp@0.5.0
fixesproto@5.0         icu4c@60.1                   libxcb@1.13            openmpi@3.0.1            py-lit@0.5.0           ruby-ronn@0.7.3        zlib@1.2.11
flex@2.6.4             inputproto@2.3.2             libxdamage@1.1.4       openssl@1.0.2n           py-mako@1.0.4          scr@1.2.2              zsh@5.4.2
font-util@1.3.1        intel-tbb@2018.2             libxdmcp@1.1.2         otf2@2.1                 py-markupsafe@1.0      shuffile@0.0.3         zstd@1.3.0
fontconfig@2.12.3      jdk@8u141-b15                libxext@1.3.3          pango@1.41.0             py-matplotlib@2.2.2    snappy@1.1.7
%
```

E4S Second Release (37+ ST products) exascaleproject/sdk:AHM19 (on Dockerhub)

```
1: adios :            adios@1.13.1  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/adios-1.13.1-v7jyzgyie7n542qppgoz2izthu6xeaj5
2: bolt :             bolt@1.0b1  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/bolt-1.0b1-jenaxkneyprxgq6abwaihlkuuoko4pwv
3: caliper :          caliper@1.8.0  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/caliper-1.8.0-lrmti32xdgyckyhk5vr5okrxtniv2pb5
4: darshan-runtime :          darshan-runtime@3.1.6  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/darshan-runtime-3.1.6-yb2tk7rst4yclkluqaixardes3slhgve
5: gasnet :           gasnet@1.30.0  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/gasnet-1.30.0-hp4d5xsbnhg5gisbkmgopd6pkqmgrczo
6: globalarrays :             globalarrays@5.7  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/globalarrays-5.7-7zbsme3slnsmzkuzgq6ac4ggbdnoakal
7: gotcha :           gotcha@develop  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/gotcha-develop-dcqszr3n36z73pqsm2d745rx5bzvr2hq
8: hdf5 :             hdf5@1.10.1  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/hdf5-1.10.1-jizgfu54nfiqzemokjopdym7l3tov7md
9: hpctoolkit :               hpctoolkit@2017.06  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/hpctoolkit-2017.06-boqjp7bdarhaysswzp6p6w5skt5wa423
10: hypre :           hypre@2.13.0  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/hypre-2.13.0-3kjvfl7rz3e7f6eojvojyfegcwdl6ehb
11: geopm :           geopm@0.4.0  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/geopm-0.4.0-qhho4xnuyymvurjeuqjfml4u42b7a3t6
12: Jupyter :             /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/miniconda3-4.3.30-6hmm62lkf5v6n3fulsw3latyjy2phlba/bin/jupyter
12: kokkos :          kokkos@2.03.00  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/kokkos-2.03.00-a3ksyhg6fflnlufs5sfanqfwxeeeogey
13: legion :          legion@17.10.0  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/legion-17.10.0-cjomljrvcxzbhwlznfc5luw6vwiubnyr
14: libquo :          libquo@1.3  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/libquo-1.3-cdtptdmouswpx5a4nvwxfyld3u3mcj62
15: magma :           magma@2.4.0  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-4.8.5/magma-2.4.0-7cc275vlzmhypm5uuubj4krfsoqshhmr
16: mfem :            mfem@3.3.2  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/mfem-3.3.2-sdrntzuthtzqophdl63b3ujmzy5ytb4g
17: mpich :           mpich@3.2.1  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/mpich-3.2.1-5j57f4j36vhcsxgn2pwndouz27qe4ij4
18: netcdf :          netcdf@4.4.1.1  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/netcdf-4.4.1.1-7vei7dnyaskclsuhpyr6wqdp4xjmdadx
19: openmpi :         openmpi@3.0.1  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/openmpi-3.0.1-hdjeffn2fs3iidk3whvv6smbrnmzqq3e
20: papi :            papi@5.5.1  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/papi-5.5.1-abkudkdhzua3p4lnn7m6ssj3or45fjri
21: papyrus :         papyrus@develop  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/papyrus-develop-77k6v4izzvjx222zbrpiexka7fmjsjgr
22: paraview :        paraview@5.4.1  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/paraview-5.4.1-qxvvvzn5qs435z25jefz2ijlhoivd3f4
23: petsc :           petsc@3.8.4  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/petsc-3.8.4-7naeokjkiniftmkecngpcn736bvnrdhl
24: pdt :             pdt@3.25  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/pdt-3.25-fjjddrbx7lx4hrqmqfwssq4oz46zvj5p
25: qthreads :        qthreads@1.12  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/qthreads-1.12-npkx43id5wewkrbsv6qpr76qisoozbpu
26: raja :            raja@0.5.3  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/raja-0.5.3-zrjr35xwjrfz6wacs4k36ilwc45m6gq6
27: scr :             scr@1.2.2  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/scr-1.2.2-fdqkevg2nf6yedg4qhwersf6ojwikxqz
28: spack :               /usr/local/packages/ecp/spack/bin/spack
28: strumpack :               strumpack@3.1.1  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/strumpack-3.1.1-q4wwcyff7lzrrbwc6np5jxezv6iix7ig
29: sundials :        sundials@3.1.0  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/sundials-3.1.0-xrqsfvumk2jw7aqidjsj7lya4w5kqm3p
30: sz :              sz@1.4.12.3  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/sz-1.4.12.3-dgykqp27gsnnyc2ktm6rnb6bfgxwq7vq
31: tasmanian :               tasmanian@6.0  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/tasmanian-6.0-fv7z3ninw7agbvlw2jhau2hyx5ofyt4k
32: tau :             tau@2.28  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/tau-2.28-2zm23cf4lu74wfp2ufrzo7bu22popu4x
33: trilinos :        trilinos@12.12.1  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/trilinos-12.12.1-kobl2zztgzcukmx5tktvmyradjt6qym7
34: vtkm :            vtkm@1.1.0   /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/vtkm-1.1.0-rze7qodn6y6pbvuivl5hw7pyekuqtiut
35: umpire :          umpire@master  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/umpire-master-4bditlkgpbuznpnnshpf3poxthmadefq
36: unifycr :         unifycr@master  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/unifycr-master-kmxew2he475aeh4jc3edhi4nbsywpekl
37: zfp :             zfp@0.5.0  /usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/gcc-7.3.0/zfp-0.5.0-bqeu73eeiodsoknvtmqakstg3hpx3zav
```

Extreme-scale Scientific Software Stack (E4S)
https://e4s.io

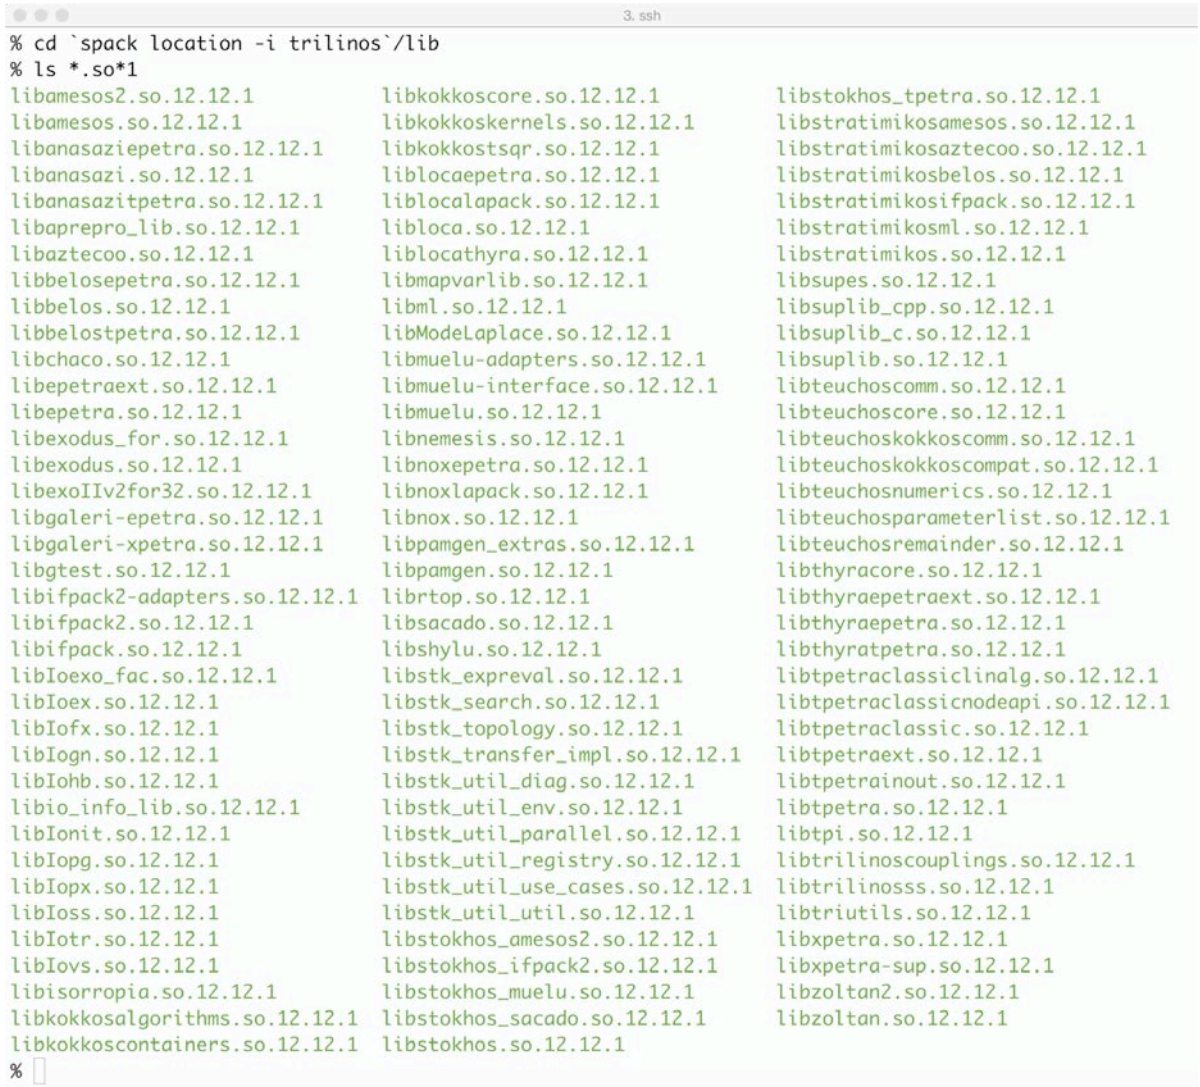

Running MPI applications on other systems

- Applications built with MPI in the E4S container can be replaced by the system MPI!
- This allows fast inter-node communication using the native interconnect.
- Application and data are external to the E4S container.
- Programming models, compilers, runtime libraries, and tools are inside the container.
- We can replace MPI using the MPICH ABI compatibility layer.
- Goal: Build an MPI binary once and run it un-modified on all HPC Linux x86_64 clusters!

Using E4S on NCSA BlueWaters and replacing MPI

Step 1: Allocate a node with the E4S image

- qsub -I -l nodes=2:ppn=32 -l walltime=01:15:00 -l gres=shifter16 -v UDI=exascaleproject/sdk:AHM19
- This allocates a single node for 1:15h
- Specifies the use of Shifter as the container environment
- The image is exascaleproject/sdk:AHM19
- This image was pulled on a compute node previously using:
  - %module load shifter; shifterimg pull exascaleproject/sdk:AHM19
- After this qsub step, we can now launch the job using aprun

Using E4S on NCSA BlueWaters Replacing MPI

Step 2: Create a file called ~/shifter_mpi.sh

```
% cat ~/shifter_mpi.sh
#!/bin/bash
# set up LD_LIBRARY_PATH
for dir in $(echo $CRAY_LD_LIBRARY_PATH:/opt/cray/wlm_detect/default/lib64 | tr
':' ' ')
do
    realpath=$(readlink -f "$dir")
    if [[ -z $LD_LIBRARY_PATH ]]
    then
        eval 'export LD_LIBRARY_PATH=/dsl'$realpath
    else
        eval 'export LD_LIBRARY_PATH=$LD_LIBRARY_PATH:/dsl'$realpath
    fi
done
```

Replacing MPI using cray-mpich-abi package

Step 3: Source this ~/shifter_mpi and setup LD_LIBRARY_PATH

```
% cat run.sh
#!/bin/bash
export CRAY_ROOTFS=SHIFTER
module load shifter
module unload PrgEnv-cray # or any other PrgEnv module currently loaded
module load PrgEnv-gnu # or PrgEnv-intel
module unload cce
module unload cray-mpich
module load cray-mpich-abi
export LD_LIBRARY_PATH=$CRAY_MPICH_DIR/lib:$LD_LIBRARY_PATH
source ~/shifter_mpi.sh
export LD_LIBRARY_PATH=/usr/local/packages/ecp/spack/opt/spack/linux-centos7-x86_64/
gcc-7.3.0/hwloc-1.11.9-7xxgxbg65an7zmrztfcuu3hs73puj6v3/lib:$LD_LIBRARY_PATH
export OMP_NUM_THREADS=2
aprun -b -n 64 -- ./lulesh.host  -i 100
```

Replacing MPI using cray-mpich-abi package

Step 4: run the example

```
% ./run.sh
Running problem size 30^3 per domain until completion
Num processors: 64
Num threads: 2
Total number of elements: 1728000
…
Run completed:
   Problem size        =  30
   MPI tasks           =  64
   Iteration count     =  100
   Final Origin Energy = 8.465100e+07
   Testing Plane 0 of Energy Array on rank 0:
        MaxAbsDiff   = 7.916242e-09
        TotalAbsDiff = 3.030168e-08
        MaxRelDiff   = 1.224484e-13

Elapsed time         =      16.58 (s)
Grid time (us/z/c)   =  6.1409471 (per dom)  (0.095952298 overall)
FOM                  =  10421.845 (z/s)
Elapsed time         =      16.58 (s)
Grid time (us/z/c)   =  6.0131382 (per dom)  (0.22270882 overall)
FOM                  =  4490.1679 (z/s)

Application 81575093 resources: utime ~442s, stime ~20s, Rss ~45404, inblocks ~9110
```

NAS Parallel Benchmark Example

Compare a native build with a container based build

```
% cd ~/scratch/host/NPB3.1
% make suite
% cd bin
% ./run.sh


% cd ~/scratch/demo/NPB3.1/bin/
% ./run.sh
```

E4S VirtualBox OVA image

Contains all four container runtimes and the E4S Singularity image!

- Docker
- Singularity
- Shifter
- Charliecloud

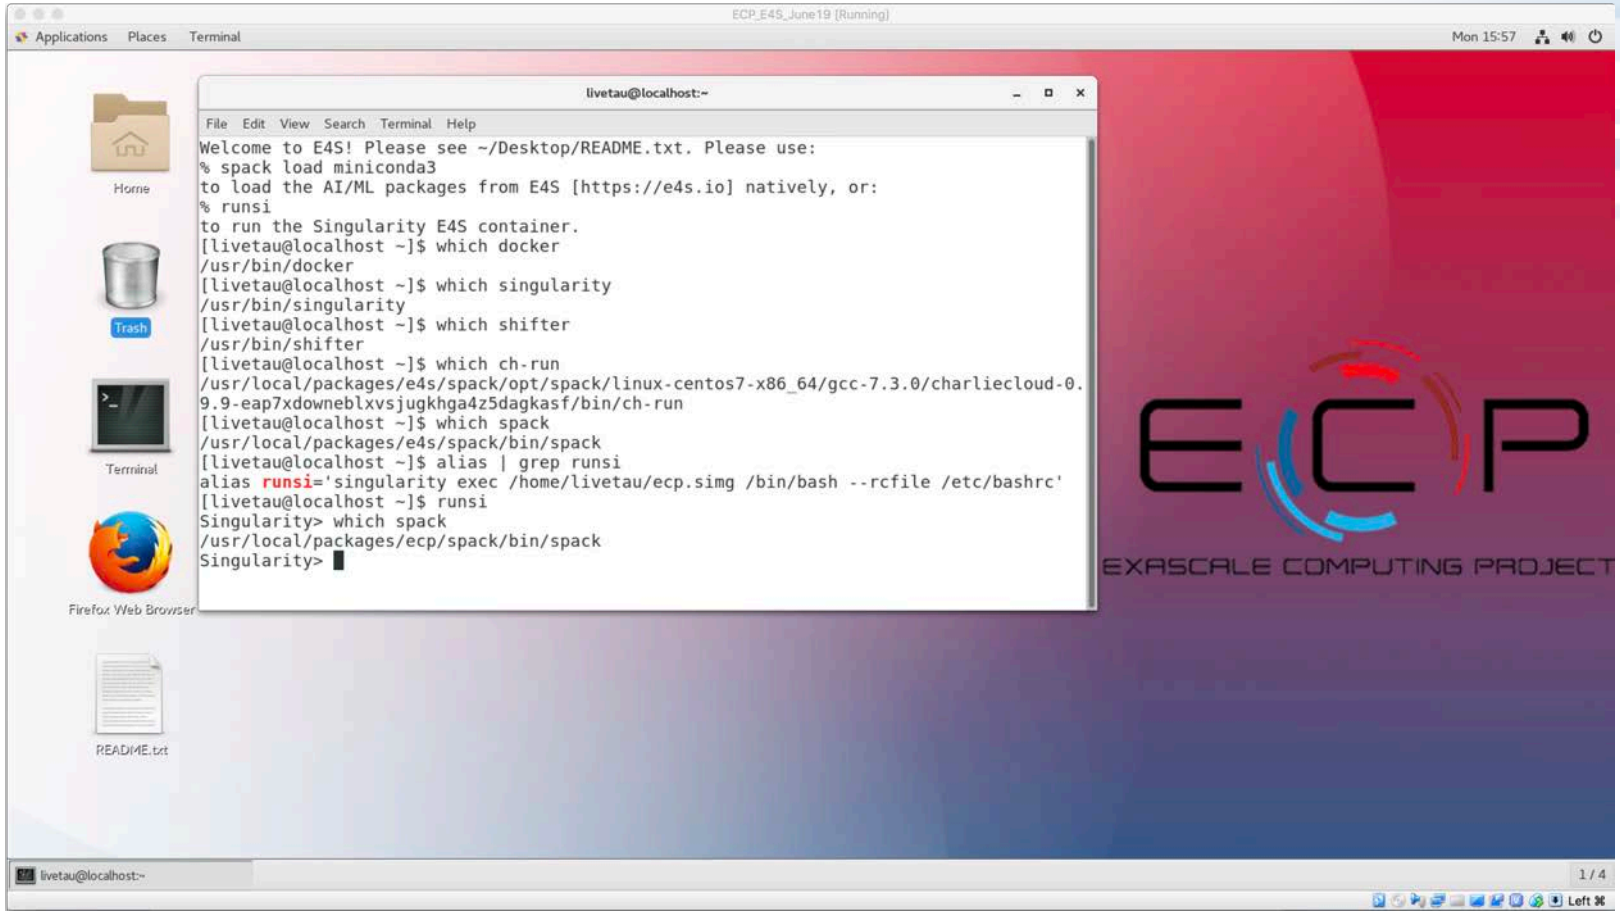

E4S image on Amazon AWS

Contains all four container runtimes and the E4S Singularity image!

- AWS AMI ID (Oregon, us-west-2 region):
  - ami-063e830287b86155c
- Royalty free, public image with HPC, AI, and 4 container runtimes
- Launch EC2 instance with this AMI
  - Login: livetau
  - Password: ****

Future work, issues…

- Increasing the number of ST packages in E4S
- Porting to IBM and ARM platforms
- Support for GPUs and visualization tools
- Addition of CI testing
- Facility deployment
- Scalable startup with full-featured “Supercontainers”
- Improving the launch of MPI applications

E4S: How to get involved

- E4S BoF at SC19
  - Tuesday, Nov. 19, 12:15pm – 1:15pm, Room 405-406-407
- CANOPIE-HPC Workshop at SC19
  - 1st Workshop on Containers and New Orchestration Paradigms for Isolated Environments in HPC
  - Monday, Nov. 18, 2019, 2pm – 5:30pm, Room 704-706
  - https://canopie-hpc.nersc.gov/
- "Container Computing for HPC and Scientific Workflows"
  - Tutorial at SC19, Sunday, Nov. 17, 2019, 1:30pm – 5pm, Room 201

Acknowledgment

“This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of two U.S. Department of Energy organizations (Office of Science and the National Nuclear Security Administration) responsible for the planning and preparation of a capable exascale ecosystem, including software, applications, hardware, advanced system engineering, and early testbed platforms, in support of the nation’s exascale computing imperative.”