E4S: Extreme-Scale Scientific Software Stack

https://e4s.io

Sameer Shende
Performance Research Laboratory, OACISS, U. Oregon

sameer@cs.uoregon.edu

First Workshop on NSF and DOE HPC Tools

312 Lillis, Universtity of Oregon, Eugene, OR

https://oaciss.uoregon.edu/NSFDOE19

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

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.

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

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

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

- 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
%
```

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

3. ssh

```
% cd `spack location -i trilinos`/lib
% ls *.so*1
libamesos2.so.12.12.1           libkokkoscore.so.12.12.1          libstokhos_tpetra.so.12.12.1
libamesos.so.12.12.1            libkokkoskernels.so.12.12.1       libstratimikosamesos.so.12.12.1
libanasaziepetra.so.12.12.1     libkokkostsqr.so.12.12.1          libstratimikosaztecoo.so.12.12.1
libanasazi.so.12.12.1           liblocaepetra.so.12.12.1          libstratimikosbelos.so.12.12.1
libanasazitpetra.so.12.12.1     liblocalapack.so.12.12.1          libstratimikosifpack.so.12.12.1
libaprepro_lib.so.12.12.1       libloca.so.12.12.1                libstratimikosml.so.12.12.1
libaztecoo.so.12.12.1           liblocathyra.so.12.12.1           libstratimikos.so.12.12.1
libbelosepetra.so.12.12.1       libmapvarlib.so.12.12.1           libsupes.so.12.12.1
libbelos.so.12.12.1             libml.so.12.12.1                  libsuplib_cpp.so.12.12.1
libbelostpetra.so.12.12.1       libModeLaplace.so.12.12.1         libsuplib_c.so.12.12.1
libchaco.so.12.12.1             libmuelu-adapters.so.12.12.1      libsuplib.so.12.12.1
libepetraext.so.12.12.1         libmuelu-interface.so.12.12.1     libteuchoscomm.so.12.12.1
libepetra.so.12.12.1            libmuelu.so.12.12.1               libteuchoscore.so.12.12.1
libexodus_for.so.12.12.1        libnemesis.so.12.12.1             libteuchoskokkoscomm.so.12.12.1
libexodus.so.12.12.1            libnoxepetra.so.12.12.1           libteuchoskokkoscompat.so.12.12.1
libexoIIv2for32.so.12.12.1      libnoxlapack.so.12.12.1           libteuchosnumerics.so.12.12.1
libgaleri-epetra.so.12.12.1     libnox.so.12.12.1                 libteuchosparameterlist.so.12.12.1
libgaleri-xpetra.so.12.12.1     libpamgen_extras.so.12.12.1       libteuchosremainder.so.12.12.1
libgtest.so.12.12.1             libpamgen.so.12.12.1              libthyracore.so.12.12.1
libifpack2-adapters.so.12.12.1  librtop.so.12.12.1                libthyraepetraext.so.12.12.1
libifpack2.so.12.12.1           libsacado.so.12.12.1              libthyraepetra.so.12.12.1
libifpack.so.12.12.1            libshylu.so.12.12.1               libthyratpetra.so.12.12.1
libIoexo_fac.so.12.12.1         libstk_expreval.so.12.12.1        libtpetraclassiclinalg.so.12.12.1
libIoex.so.12.12.1              libstk_search.so.12.12.1          libtpetraclassicnodeapi.so.12.12.1
libIofx.so.12.12.1              libstk_topology.so.12.12.1        libtpetraclassic.so.12.12.1
libIogn.so.12.12.1              libstk_transfer_impl.so.12.12.1   libtpetraext.so.12.12.1
libIohb.so.12.12.1              libstk_util_diag.so.12.12.1       libtpetrainout.so.12.12.1
libio_info_lib.so.12.12.1       libstk_util_env.so.12.12.1        libtpetra.so.12.12.1
libIonit.so.12.12.1             libstk_util_parallel.so.12.12.1   libtpi.so.12.12.1
libIopg.so.12.12.1              libstk_util_registry.so.12.12.1   libtrilinoscouplings.so.12.12.1
libIopx.so.12.12.1              libstk_util_use_cases.so.12.12.1  libtrilinosss.so.12.12.1
libIoss.so.12.12.1              libstk_util_util.so.12.12.1       libtriutils.so.12.12.1
libIotr.so.12.12.1              libstokhos_amesos2.so.12.12.1     libxpetra.so.12.12.1
libIovs.so.12.12.1              libstokhos_ifpack2.so.12.12.1     libxpetra-sup.so.12.12.1
libisorropia.so.12.12.1         libstokhos_muelu.so.12.12.1       libzoltan2.so.12.12.1
libkokkosalgorithms.so.12.12.1  libstokhos_sacado.so.12.12.1      libzoltan.so.12.12.1
libkokkoscontainers.so.12.12.1  libstokhos.so.12.12.1
%
```

Running MPI applications on other systems

- Applications built with MPI in the E4S container can replace the MPI in the container with 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 Frontera at TACC

- Use the training account:

```
% ssh −Y train<id>@frontera.tacc.utexas.edu
% cp −r /home1/00494/tg457572/E4S/Containers/demo .
% cd demo/NPB3.1/
% idev −m 50
% module load tacc-singularity
%cat run_sing.sh
#!/bin/bash
singularity exec /scratch1/00494/tg457572/E4S/Containers/ecp.simg /bin/bash --rcfile /etc/bashrc
% ./run_sing.sh
Singularity> which spack
Singularity> cd NPB3.1; make clean; make suite; ls −l bin/lu.C.64
Singularity> exit
```

Using E4S on Frontera at TACC

- After allocating the node and exiting from Singularity:

  % cd NPB3.1/bin; cat run.sh

  module load tacc-singularity

  export PATH=/scratch1/00494/tg457572/pkgs/mvapich2-231/bin:$PATH

  export MV2_ENABLE_AFFINITY=0

  mpirun -np 64 singularity exec -B /etc/libibverbs.d:/etc/libibverbs.d -B /usr/lib64:/hostlib64 -B /opt:/opt -B /scratch1:/scratch1 /scratch1/00494/tg457572/E4S/Containers/ecp.simg /bin/bash -c ' . /etc/bashrc; spack unload mpich openmpi; spack load gcc; export LD_LIBRARY_PATH=/scratch1/00494/tg457572/pkgs/mvapich2-231/lib:$LD_LIBRARY_PATH:/hostlib64:/hostlib64/libibverbs; ./lu.C.64'

  % ./run.sh

Building on laptop and running on Frontera (TACC)

- Build a Trilinos application on your laptop

```
% docker images
% docker run −v /Users/<login>:/home/<login> -i −t exascaleproject:sdk/AHM19 /bin/bash
% cp −r /usr/local/packages/ecp/apps/demo/trilinos . ; cd trilinos/Zoltan; ./clean.sh; ./compile.sh
% scp Zoltan frontera.tacc.utexas.edu:demo/trilinos/Zoltan
% ssh −Y train<id>@frontera.tacc.utexas.edu
% cd demo/trilinos/Zoltan
% idev −N 2 −m 50
% cat run.sh
% ./run.sh
# Replaces mpi with system MPI on Frontera!
```

E4S VirtualBox OVA image

Contains all four container runtimes and the E4S Singularity image!

- Docker
- Singularity
- Shifter
- Charliecloud

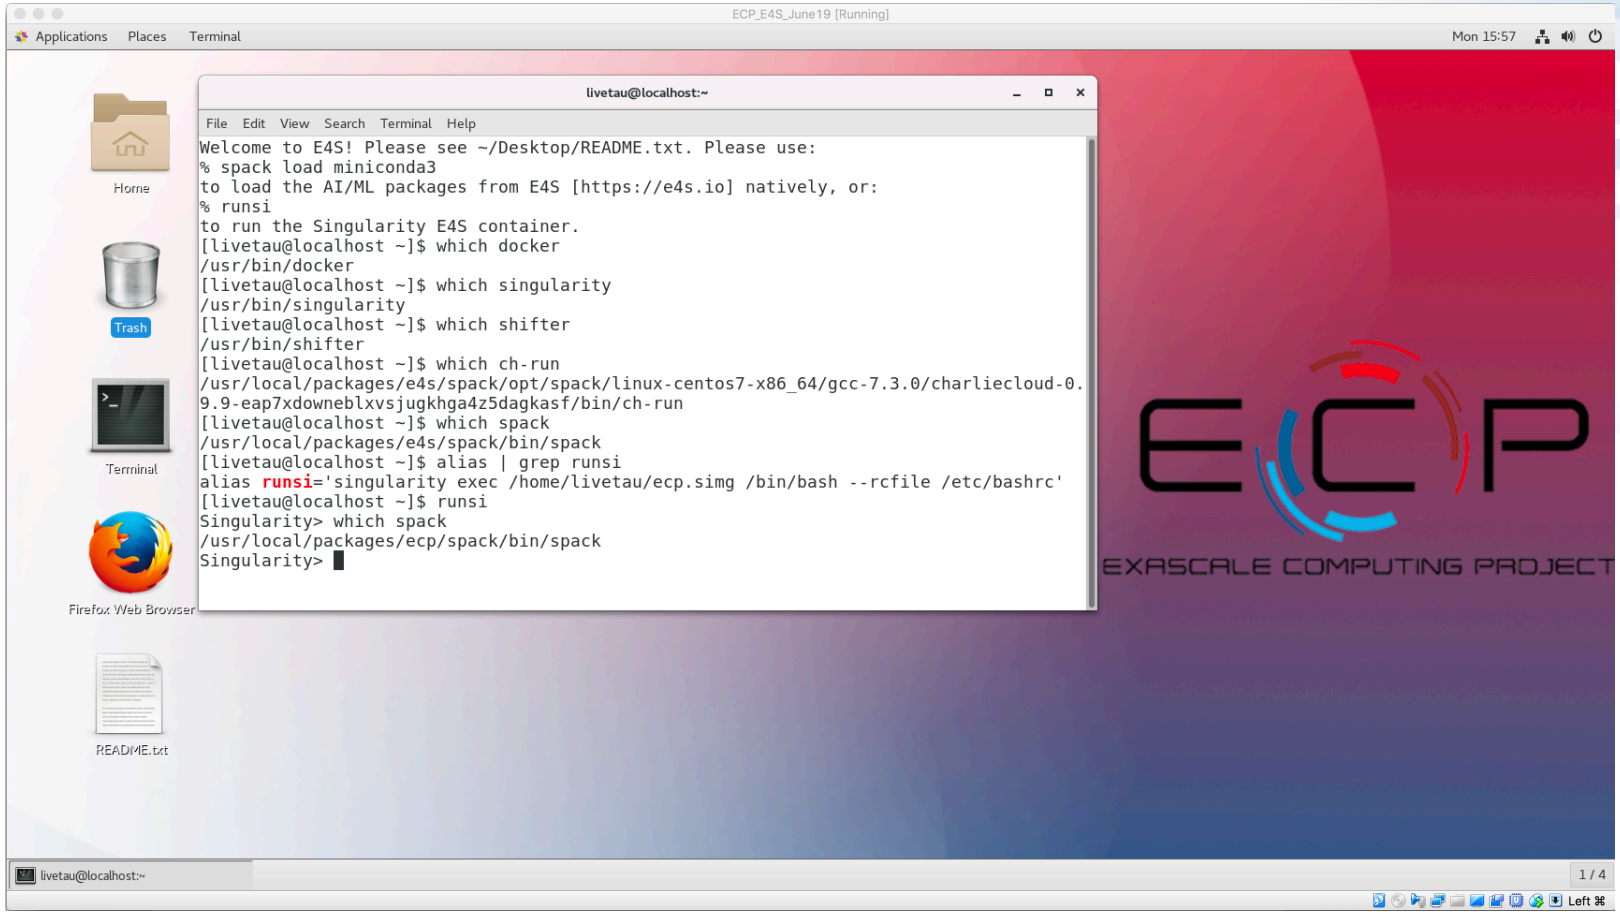

14

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: ***
  - Password: ****
  - Email: sameer@cs.uoregon.edu for login info.

15

Singularity on Theta at ALCF

```
% qsub -A ECP_SDK -t 30 -n 2  -q debug-cache-quad  -I
% /projects/ECP_SDK/tutorial/run_job.sh
module swap PrgEnv-intel PrgEnv-gnu

module swap cray-mpich cray-mpich-abi

export SINGULARITYENV_LIBWLM_DETECT=/opt/cray/wlm_detect/
1.3.2-6.0.6.0_3.8__g388ccd5.ari/lib64

aprun -n 16 -N 8 singularity exec -H $HOME    -B /projects/ECP_SDK:/projects/
ECP_SDK:ro -B /opt:/opt:ro -B /var/opt:/var/opt:ro /projects/ECP_SDK/containers/
singularity/ecp.simg bash -c 'unset CRAYPE_VERSION; source /usr/local/
packages/ecp/misc/bashrc; spack load   trilinos hypre parmetis hdf5 metis
openblas superlu zlib netcdf matio boost@1.66.0 scalapack suite-sparse
tau ;spack unload openmpi mpich ;  export
LD_LIBRARY_PATH=$LIBWLM_DETECT:$CRAY_LD_LIBRARY_PATH:
$CRAYPAT_LD_LIBRARY_PATH:$LD_LIBRARY_PATH   ; /projects/ECP_SDK/
tutorial/demo/trilinos/Zoltan/Zoltan; '
```

Singularity on Quartz at LLNL

MVAPICH2 needs /lib. Mount it as /hostlib64 and add it to LD_LIBRARY_PATH

```
% salloc -N 2
% srun -n 4 -c 2 singularity exec -B /lib64:/hostlib64 -B
$SLURM_SUBMIT_DIR:$SLURM_SUBMIT_DIR -B /usr/tce:/usr/tce ./ecp.simg /
bin/bash -c ' . /etc/bashrc ; spack load trilinos hypre parmetis hdf5 metis
openblas superlu zlib netcdf matio boost@1.66.0 scalapack suite-sparse tau;
spack unload openmpi mpich; export LD_LIBRARY_PATH=/usr/tce/packages/
mvapich2/mvapich2-2.2-intel-18.0.1/lib:$LD_LIBRARY_PATH:/hostlib64; ./Zoltan'
```

Replacing MPI with Shifter on Cori.nersc.gov

% shifterimg images

exascaleproject/sdk:AHM19 …

% To replace MPI with system MPI:

salloc -N 2 -q interactive -t 00:30:00 --image=exascaleproject/sdk:AHM19 -C haswell -L SCRATCH

~sameer/run_shifter.sh

cat ~/run_shifter.sh

srun −n 32 shifter -- /bin/bash -c 'unset CRAYPE_VERSION; . /etc/bashrc ; spack load trilinos hypre parmetis hdf5 metis openblas superlu zlib netcdf matio boost@1.66.0 scalapack suite-sparse tau; spack unload openmpi mpich; ./Zoltan'

18

Future work, issues…

- Increasing the number of ST packages in E4S
- Build cache mirrors for E4S packages
- 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

19

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.”