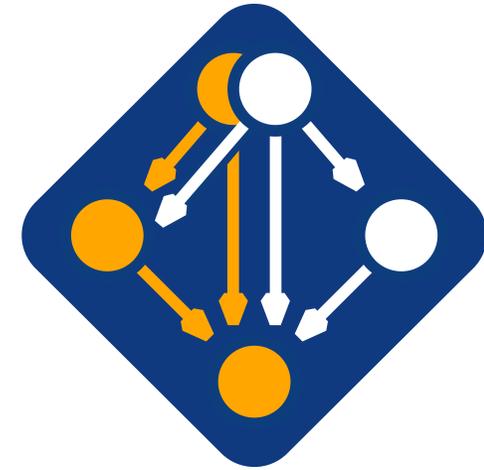


# The Future of Spack

HEP Software Developer Tools & Packaging Working Group  
October 20, 2021



Todd Gamblin  
*Advanced Technology Office*  
*Livermore Computing*

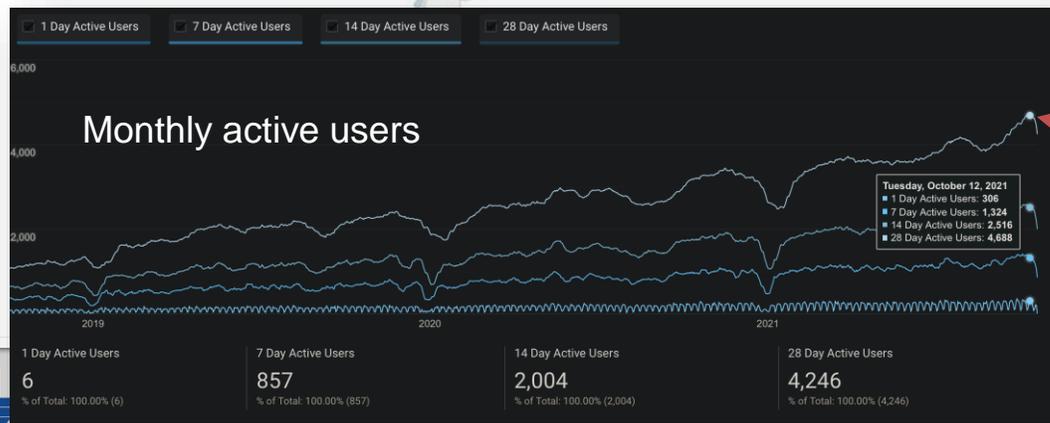
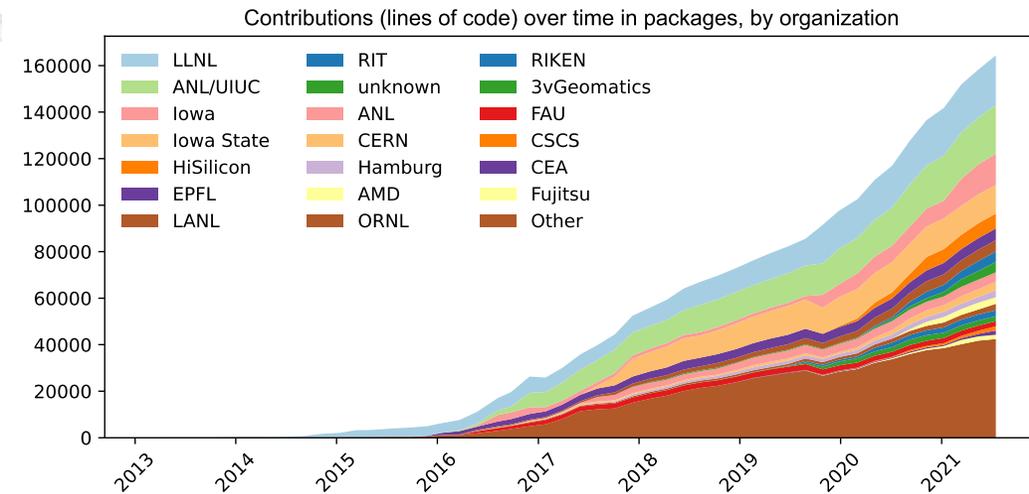


# The Spack community continues to grow!

**5,900+** software packages  
**900+** contributors



Package contribution rate increased in 2020



Broke 4,600 monthly active users on docs site in October 2021

# Spack core development is sustained by a number of key DOE programs and partners



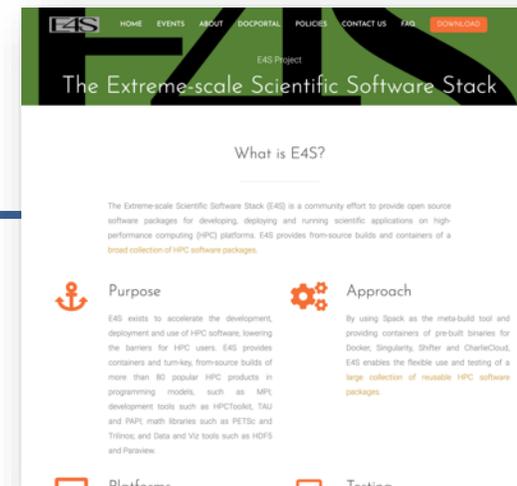
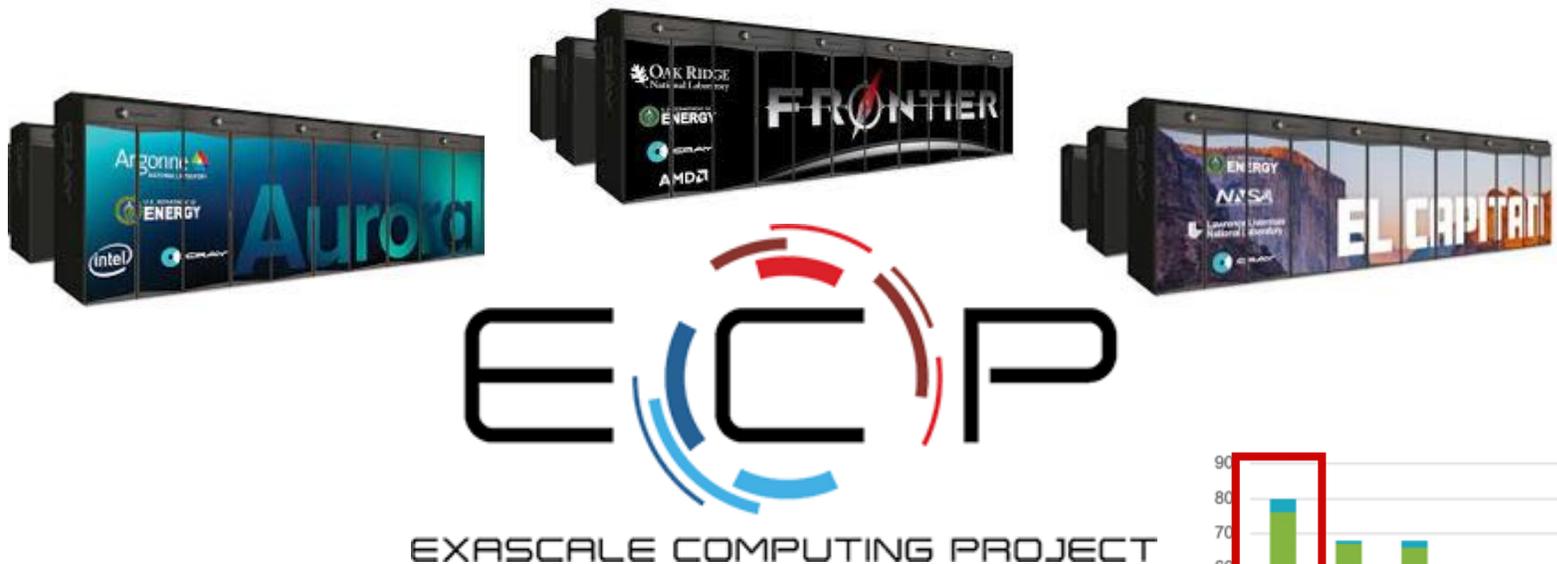
- Common Compute Environment
- Computation Systems & Software Environment
- Facility Operations & User Support



- Packaging Technologies Project
- Facility Integration
- Testing Task Force

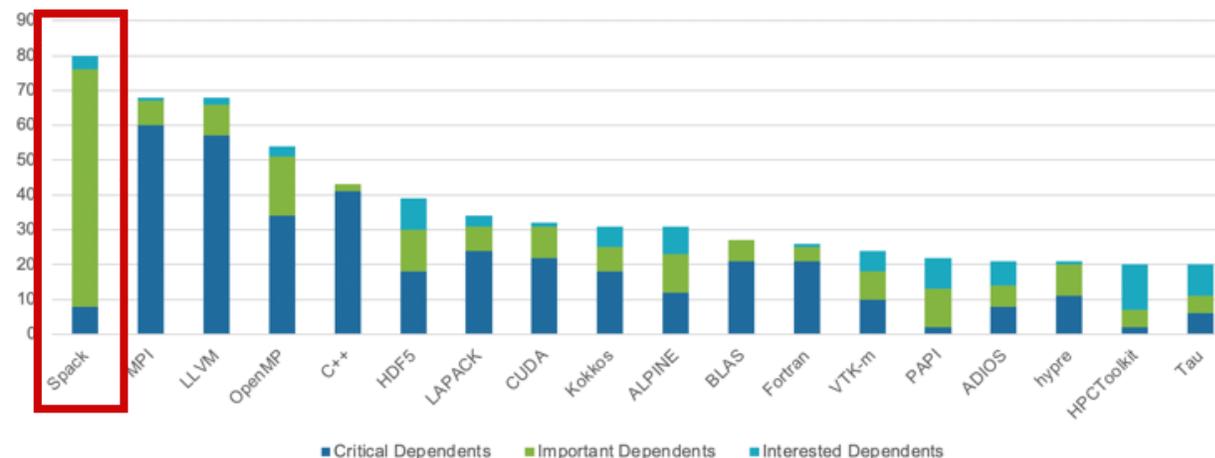


# Spack is a critical part of ECP



<https://e4s.io>

Dependents by Producer



Spack is the most depended-upon project in ECP

- Spack is used to integrate the software ecosystem for the three upcoming U.S. exascale systems
- ECP has built the Extreme Scale Scientific Software Stack (E4S) with Spack – more at <https://e4s.io>

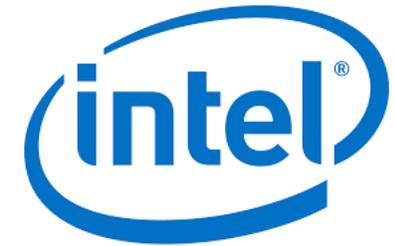
# We have worked with HPE to improve Spack support in the Cray Programming Environment

- Using Cray Programming Environment's MPI, libsci, etc. currently requires a fair amount of configuration
  - Users have to register externals and go through modules
  - PrgEnvs make it hard to be precise about dependencies
- PE team has worked with us to develop a JSON format to describe PE contents 
- PE team has also been responsive with other requests:
  - Standalone compiler executables (no modules)
  - Standalone MPI wrappers (like normal MPICH)
- Things are getting better!

```
{
  "name": "cray-netcdf",
  "version": "4.7.4.0",
  "arch": {
    "platform": "cray",
    "platform_os": "sles15",
    "target": {
      "name": "x86_64"
    }
  },
  "compiler": {
    "name": "cce",
    "version": "9.0.0"
  },
  "parameters": {
    "shared": true,
    "dap": false,
    "hdf4": false,
    "parallel-netcdf": false,
    "mpi": true,
    "jna": false,
    "doxygen": false,
    "doc": false
  },
  "provides": "netcdf-fortran",
  "dependencies": {
    "cray-mpich": {
      "hash": "mvl4uwf63n4l7tuwfrdyqfxmmit7yu54",
      "type": [
        "link"
      ]
    },
    "cray-hdf5": {
      "hash": "tdma2n3hxn25lh1xn7dbkvt23jo2f3mc",
      "type": [
        "link"
      ]
    },
    "zlib": {
      "version": "1.2.11",
      "type": [
        "link"
      ]
    }
  },
  "prefix": "/opt/cray/pe/netcdf-hdf5parallel/4.7.4.0/cce/90",
  "rpm": "cray-netcdf-4.7.4.0-crayclang90-202007092040.ac3e2015515ab-0.sles15.x86_64.rpm",
  "hash": "hsm5hatcfepa6hbfynpu2iv6cxfi0d7i"
},
```

# Spack's widespread adoption has made it a de facto standard, drawing contribution and collaboration from many vendors

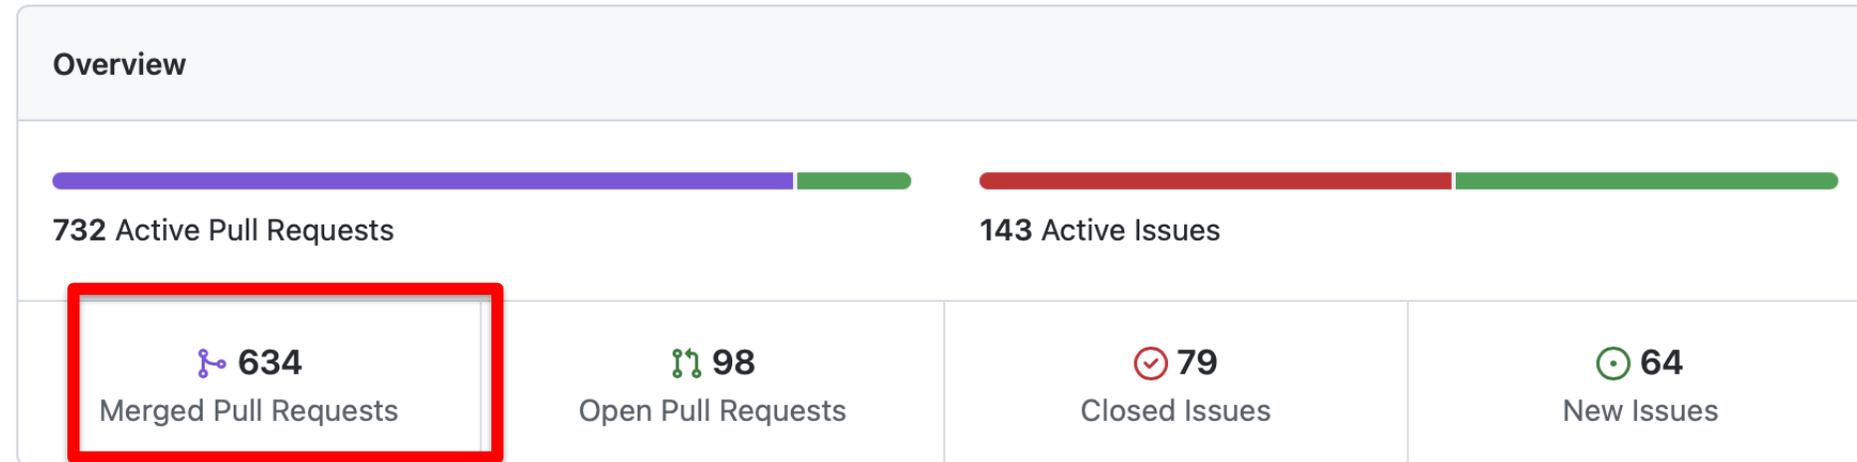
- **AWS** invests in cloud credits for Spack build farm
  - Joint Spack tutorial in July with AWS had 125+ participants
  - Joint AWS/AHUG Spack Hackathon drew 60+ participants
- **AMD** has contributed ROCm packages and compiler support
  - 55+ PRs mostly from AMD, also others
  - ROCm, HIP, aocc packages are all in Spack now
- **Intel** contributing OneApi support and licenses for our build farm
- **NVIDIA** contributing NVHPC compiler support and other features
- **Fujitsu and RIKEN** have contributed a **huge** number of packages for ARM/a64fx support on Fugaku
- **ARM** and **Linaro** members contributing ARM support
  - 400+ pull requests for ARM support from various companies



# Spack GitHub remains a very busy place

September 20, 2021 – October 20, 2021

Period: 1 month ▾

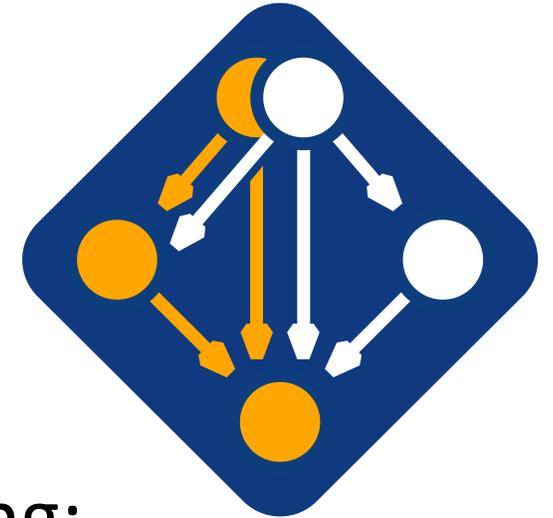


Excluding merges, **179 authors** have pushed **641 commits** to develop and **696 commits** to all branches. On develop, **1,273 files** have changed and there have been **22,569 additions** and **5,572 deletions**.



# What are the priorities going forward?

1. Improving support for GPUs and exascale machines (Cray PE)
2. Developer workflows
3. CI, Testing, and Binary Distribution
4. Deeper modeling of the software ecosystem
5. Broader reach into non-traditional HPC, including:
  - C++ community
  - Windows
  - Cloud environments

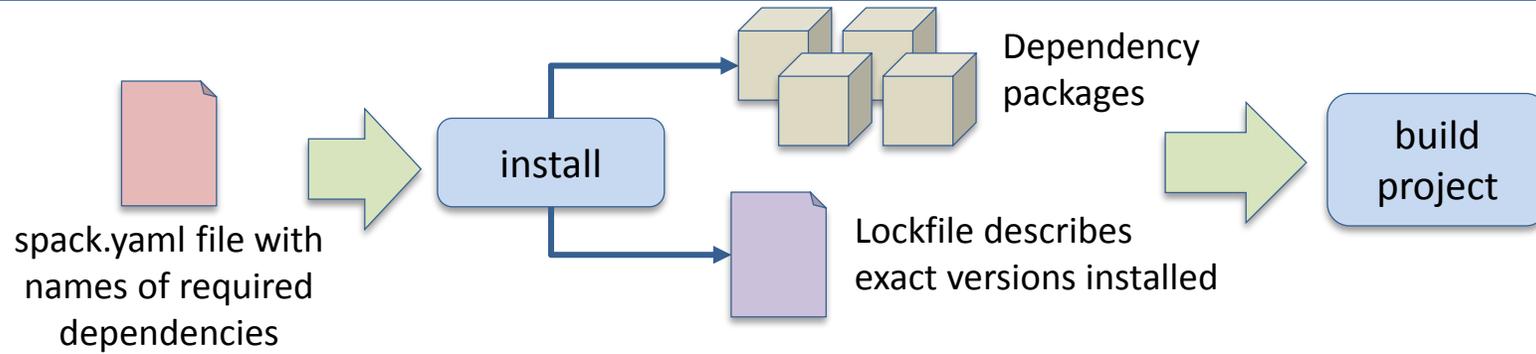


# Our strategy going forward is to focus on building and testing for the compiler/runtime stacks on EA systems



- **We added much-needed support for new vendor compilers**
  - oneapi: Intel, nvhpc: NVIDIA, aocc: AMD
- **Working to add tests so that we can use new spack test capability on EA systems**
  - Community can integrate smoke tests into Spack packages
  - We will be running continuous smoke tests in CI for the ECP stack
- **GPU integration across the stack will be an ongoing focus**
  - Focus on CUDA and HIP packages and common usage
  - Focus on interoperability for mixed compilers, continue to work with vendors
- **Aggressive tests of relocatable binary builds on all EA systems**
  - Keep things working with constant pull request testing

# Spack's environment abstraction allows users to work with a subset of packages



## Two files:

- spack.yaml describes project requirements and configuration
- spack.lock describes exactly what was installed
  - Allows you to reproduce an installation
- An environment can be versioned in a repository
  - Technique used by uberenv, serac, other projects to manage dependencies
- MAPP repository structure maps fairly well to a Spack environment

## Simple spack.yaml file

```
spack:
# include external configuration
include:
- ../special-config-directory/
- ./config-file.yaml

# add package specs to the `specs` list
specs:
- hdf5
- libelf
- openmpi
```

## Concrete spack.lock file (generated)

```
{
  "concrete_specs": {
    "6s63so2kstp3zyvjezglndmavy6l3nui": {
      "hdf5": {
        "version": "1.10.5",
        "arch": {
          "platform": "darwin",
          "platform_os": "mojave",
          "target": "x86_64"
        },
      },
      "compiler": {
        "name": "clang",
        "version": "10.0.0-apple"
      },
      "namespace": "builtin",
      "parameters": {
        "cxx": false,
        "debug": false,
        "fortran": false,
        "hl": false,
        "mpi": true,
      }
    }
  }
}
```

# We have built a number of developer workflow features on top of Spack environments (similarities with SpackDev)

```
class Cmake(Package):
    executables = ['cmake']

    @classmethod
    def determine_spec_details(cls, prefix, exes_in_prefix):
        exe_to_path = dict(
            (os.path.basename(p), p) for p in exes_in_prefix
        )
        if 'cmake' not in exe_to_path:
            return None

        cmake = spack.util.executable.Executable(exe_to_path['cmake'])
        output = cmake('--version', output=str)
        if output:
            match = re.search(r'cmake.*version\s+(\S+)', output)
            if match:
                version_str = match.group(1)
                return Spec('cmake@{0}'.format(version_str))
```

package.py

## spack external find

Automatically find and configure external packages on the system

```
packages:
  cmake:
    externals:
      - spec: cmake@3.15.1
        prefix: /usr/local
```

spack.yaml configuration

## spack test

Packages know how to run their own test suites

```
class Libsigsegv(AutotoolsPackage, GNUMirrorPackage):
    """GNU libsigsegv is a library for handling page faults in user mode..."""

    # - spack package contents...

    extra_install_tests = 'tests/libs/'

    def test(self):
        data_dir = self.test_suite.current_test_data_dir
        smoke_test_c = data_dir.join('smoke_test.c')

        self.run_test(
            'cc',
            ['%s' % self.prefix.include,
             '%s' % self.prefix.lib, 'libsigsegv',
             smoke_test_c,
             '-o', 'smoke_test']
        )
        purpose = check_linking()

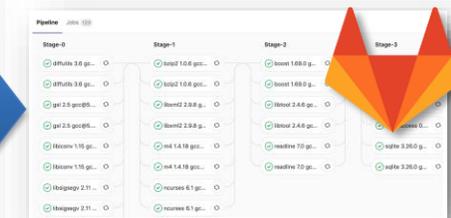
        self.run_test(
            'smoke_test', [data_dir.join('smoke_test.out')],
            purpose='run built smoke test')

        self.run_test('sigsegv1', [TestPassed], purpose='check sigsegv1 output')
        self.run_test('sigsegv2', [TestPassed], purpose='check sigsegv2 output')
```

package.py

```
spack:
  definitions:
    - pkgst:
      - readLine@7.0
      - compilers:
        - "gcc@5.0"
    - ossl:
      - openssl@1.0.2k
      - openssl@1.0.2k
  specs:
    - matrix:
      - [pkgst]
      - [compilers]
      - [ossl]
  mirrors:
    cloud_gitlab: https://mirror.spack.io
  gitlab-ci:
    mappings:
      - spack-cloud-ubuntu:
        match:
          - os=ubuntu18.04
        runner-attributes:
          tags:
            - spack-k8s
        image: spack/spack_builder_ubuntu_18_04
      - spack-cloud-centos:
        match:
          - os=centos7
        runner-attributes:
          tags:
            - spack-k8s
        image: spack/spack_builder_centos_7
  cdash:
    build-group: Release Testing
    url: https://cdash.spack.io
    project: Spack
    site: Spack AWS Gitlab Instance
```

spack.yaml



gitlab-ci.yml CI pipeline

## spack ci

Automatically generate parallel build pipelines (more on this later)

## spack containerize

Turn environments into container build recipes

```
spack:
  specs:
    - gromacs+mpi
    - mpich

  container:
    # Select the format of the recipe e.g. docker,
    # singularity or anything else that is currently supported
    format: docker

    # Select from a valid list of images
    bases:
      image: "centos:7"
      spack: develop

    # Whether or not to strip binaries
    strip: true

    # Additional system packages that are needed at run time
    os_packages:
      - libgomp

    # Extra instructions
    extra_instructions:
      final: |
        RUN echo "PS1='${S}(${S} ${S}) ${S}(${S} ${S}) ${S}(${S} ${S})"

    # Labels for the image
    labels:
      app: "gromacs"
      mpi: "mpich"
```

```
# Build stage with Spack pre-installed and ready to be used
FROM spack/centos:latest as builder

# What we want to install and how we want to install it
# It is specified in a manifest file (spack.yaml)
RUN mkdir /opt/spack-environment/
    && echo "spack" > /opt/spack-environment/spack.yaml
    && echo "spack" > /opt/spack-environment/spack.yaml
    && echo "mpich" > /opt/spack-environment/spack.yaml
    && echo "gromacs" > /opt/spack-environment/spack.yaml
    && echo "concentrations: together" > /opt/spack-environment/spack.yaml
    && echo "config" > /opt/spack-environment/spack.yaml
    && echo "install_tree: /opt/spack" > /opt/spack-environment/spack.yaml

# Install the software, remove unnecessary deps
RUN cd /opt/spack-environment && spack install && spack gc -y

# Strip all the binaries
RUN find -L /opt/spack-environment -type f -exec readlink -f {} \; | \
  xargs -r -n 1 -I {} sed -i 's|{}|$(readlink -f {})|g'
    && echo "strip" > /opt/spack-environment/spack.yaml

# Modifications to the environment that are necessary to run
RUN cd /opt/spack-environment && \
  spack env activate --sh -e -> /etc/profile.d/218_spack_environment.sh

# Bare OS image to run the installed executables
FROM centos:7

COPY --from=builder /opt/spack-environment /opt/spack-environment
COPY --from=builder /opt/spack-environment /opt/spack-environment
COPY --from=builder /etc/profile.d/218_spack_environment.sh /etc/profile.d/218_spack_environment.sh
RUN yum update -y && yum install -y epel-release && yum update -y
    && yum install -y libgomp && \
    && rm -f /var/cache/yum && yum clean all

RUN echo "export PS1='${S}(${S} ${S}) ${S}(${S} ${S}) ${S}(${S} ${S})" > /etc/profile.d/218_spack_environment.sh
```



# spack develop lets developers work on many packages at once

- Workflow:
  - Set up an environment for a code
  - Develop multiple packages from its dependencies
  - Easily rebuild with changes
- You need to supply a version to `spack develop`
  - Spack needs to know what version to check out.
  - Can be a git version, tarball, etc.
- Develop installations are install in the regular Spack install tree, with a unique hash

```
$ spack env activate .  
$ spack add myapplication  
$ spack develop axom@0.4.0  
$ spack develop mfem@4.2.0
```

```
$ ls  
spack.yaml  axom/  mfem/
```

```
$ cat spack.yaml  
spack:  
  specs:  
    - myapplication # depends on axom, mfem  
  
  develop:  
    - axom @0.4.0  
    - mfem @develop
```

# We have added git versioning to Spack

- Users can now specify a full, 40-char git commit as a version
  - Works in environments or on the command line

```
$ spack install zlib @53ce2713117ef2a8ed682d77b944df991c499252
```

- This was tricky because we needed a way to compare a commit to a version
  - MBS only needs to be able to fetch by commit, not compare
  - Packages have conditional logic with versions
  - We can compare versions to commits based on tags in a repository
- We developed an internal representation for commit versions
  - Lexicographic tuple comparison:  
  
(<version>, "", <commits since prior tag>)  
  
– Comes before any <version>.x
  - Allows commits to be compared by distance between versions.
- We can iterate on these semantics
  - there are certain types of versioning schemes for which they won't work



# The AML team has used Spack environments to accelerate their workflow for PyTorch

- **LLNL Applied ML team needed to deploy**
  - PyTorch + Kull development environment
  - On ppc64le with system MPI
- **Before Spack**
  - Everybody built from scratch
  - People wrote scripts and passed them around
  - **Days were spent trying to debug build differences**
- **After spack**
  - Versioned reproducible spack environments in a git repo
  - Standard environments in a shared team directory
  - **Team members can set up a reproducible, customizable environment in ~20 minutes.**
    - Change python version, PyTorch version on the fly
    - Leverage binary caches to avoid redundant builds.

```

spack:
  specs:
    - py-horovod
    - py-torch
    - python
    - py-h5py

  packages:
    all:
      providers:
        mpi:
          - mvapich2@2.3
        lapack:
          - openblas threads=openmp
        blas:
          - openblas threads=openmp
      buildable:true
      variants: [+cuda cuda_arch=37]
      compiler:[gcc@7.3.0]
    ...
  python:
    version: [3.8.6]
  cudnn:
    version:
      - 8.0.4.30-11.1-linux-x64
  py-torch:
    buildable:true
    variants: +cuda +distributed
  mvapich2:
    externals:
      - spec: mvapich2@2.3.1%gcc@7.3.0
        prefix:/usr/tce/packages/mvapich2/mvapich2-2.3-gcc-7.3.0
  compilers:
    - compiler:
        operating_system: rhel7
      paths:
        cc:/usr/tce/packages/gcc/gcc-7.3.0/bin/gcc
        cxx:/usr/tce/packages/gcc/gcc-7.3.0/bin/g++

```

spack.yaml file

This stack is not as complex as MARBL, but seemed like a good basis.

# Under this new workflow, users can describe their environment in a spack.yaml

- First section is familiar
  - List of packages with hashes
- spack.yaml contains:
  - Hashes
  - Per-package options
  - Compiler options
  - Flags
  - External packages
  - Local package recipe repositories
  - Etc.
- If this is too long, some of this can be moved to external includes

```

spack:
  specs:
    - marbl @develop build_type=Release
    - miranda @develop
    - blast @wktexports
    - exo @wktexports
    - adiak @950e3bf91519ecb7b7ee7fa3063bfab23c0e2c9
    - ascent ~fortran~openmp @587f6cf9503ef6176e59a046f6331baed5e36ce6
    - axom ~lua~openmp @587f6cf9503ef6176e59a046f6331baed5e36ce6
    - blt @43022da4dfed5a50a02fbd0355defd03f12157cd
    - caliper~libdw @85601f48e7f883fb87dec85e92c849eec2bb61f7
    - camp @85601f48e7f883fb87dec85e92c849eec2bb61f7
    - care @7f43ed9ed8400f6173b8434b6471142a8ff4d882
    - chai @d3282bc95c533efb90ec0a06085e455daa97df6b
    - conduit @f54f834eb8aaff4fc97613e04cfdb360997867be
    - dray ~test~utils~openmp @c0bee76f2dce29139bde1084bf085d7d1c1b01b4
    - e14 @aded490988f1d0a11ff74f9be7135d95e25e90ca
    - glvis @20aeb2c03ce70f445232dba74179e03c94da0c2c
    - gotcha @e0455990e57e5b74e16343816cd0d2d4f38d65de
    - irep @5d4d2893b25c4dfe4ad05dd6d8110179980c2a6b
    - leilak @1886056c398a6919bf8cce4216732fcd1d8643954
    - mfem +shared @9d8043b9e78dcdcd86639bbb28d3bd7b514fb5e2
    - raja ~openmp @9cb6370bb2868a3ebba23cdce927f5f7f9da530
    - ransbox @edf072bfa7b3f6e0fd6eb106abbe65ae5f677abe
    - samrai @39017121bda44fff713fe3b01cb1e063be93023b
    - selene @6f9b15713c738d70b125bc08aef72925d961a02e
    - spheral @8cc54824c2937405203c3803ab44960fc26d506d
    - tribol @b9185d317bf14d87462ca345086931580c591eb4
    - umpire ~openmp @5201a47a35e3844160dcbecd0916f8c96aa7dd07
    - vtkh @cd6004c94b083b096fda5f994b491b8229dacd79
    - hdf5 @1.8 +cxx+fortran-mpi
    - netcdf-c ~mpi @3.7.2
    - python @1.76.0
    - boost @8.3.4.1
  view: false
  concretization: together

  repos:
    - ~/src/llnl.wci.mapp
    - $spack/var/spack/repos/builtin
    - ~/src/llnl.wci

  compilers:
    - compiler:
      spec: intel@18.0.2
      paths:
        cc: /usr/tce/bin/icc-18.0.2
        cxx: /usr/tce/bin/icpc-18.0.2
        f77: /usr/tce/bin/fort-18.0.2
        fc: /usr/tce/bin/fort-18.0.2
      flags: {}
      operating_system: rhel7
      target: x86_64
      modules: [gcc/4.9.3, intel/18.0.2]
    
```

options,  
versions/hashes

package repos

compiler info

```

packages:
  all:
    compiler: [intel@18.0.2]
    providers:
      mpi: [mvapich2]
      blas: [netlib-lapack]
      lapack: [netlib-lapack]
  hypre:
    variants: +shared
  mpi:
    buildable: false
    externals:
      - spec: mvapich2@2.3%intel@18.0.2 process_managers=slurm arch=linux-rhel7-ivybridge
        prefix: /usr/tce/packages/mvapich2/mvapich2-2.3-intel-18.0.2
  blas:
    buildable: false
  lapack:
    buildable: false
  netlib-lapack:
    buildable: false
    externals:
      - spec: netlib-lapack@3.6.1+shared
        prefix: /usr
  cuda:
    buildable: false
    externals:
      - spec: cuda@10.2
        prefix: /opt/cudatoolkit/10.2
  # Basic build deps
  autoconf:
    buildable: false
    externals:
      - spec: autoconf@2.69
        prefix: /usr
  automake:
    buildable: false
    externals:
      - spec: automake@1.13.4
        prefix: /usr
  bzip2:
    buildable: false
    externals:
      - spec: bzip2@1.0.6
        prefix: /usr
  cmake:
    version: [3.14.5]
    buildable: false
    externals:
      - spec: cmake@3.14.5
        prefix: /usr/tce/packages/cmake/cmake-3.14.5
  gettext:
    buildable: false
    externals:
      - spec: gettext@0.19.8.1
        prefix: /usr
  libtool:
    buildable: false
    externals:
      - spec: libtool@2.4.2
        prefix: /usr
  m4:
    buildable: false
    externals:
      - spec: m4@1.4.16
        prefix: /usr
  perl:
    buildable: false
    externals:
      - spec: perl@5.16.3
        prefix: /usr
  pkg-config:
    buildable: false
    externals:
      - spec: pkg-config@0.27.1
        prefix: /usr
  tar:
    buildable: false
    externals:
      - spec: tar@1.26
        prefix: /usr
    
```

external  
package prefs

MPI

BLAS/LAPACK

build  
dependencies

# Spack generates a spack.lock file that enables you to reproduce the environment

- Users specify their constraints in spack.yaml
  - The rest of configuration is automated by the *concretizer*
  - The concretizer is a constraint solver that reconciles package requirements with yours
  - Details are beyond the scope of this presentation
- If you modify spack.yaml, you can either:
  - Run `spack install` again (this concretizes before installing)
  - Run `spack concretize --force` to see the concretized environment before installing (shown at right)
- spack.lock contains all the decisions the concretizer made:
  - Versions
  - Compilers, compiler versions
  - Variant values
  - Optional dependencies
  - Target architecture
- Open question: how best to manage spack.lock files

```
(rzenies)~> spack -e ~/envs/marbl concretize -f
Concretized marbl build_type=Release
marbl@2021.9.2%intel@18.0.2-ipo build_type=Release arch=linux-rhel7-ivybridge
Ablast@kwtexports%intel@18.0.2+comp-glvls-ipo+lua+opacity+physcsutils build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Abit@0.4.1%intel@18.0.2 arch=linux-rhel7-ivybridge
Acmake@3.14.5%intel@18.0.2+doc+ncurses+openssl+ownlibs-qt build_type=Release patches=ic540040c7e203
Acamp@0.1.0%intel@18.0.2-cuda-ipo-roc-rcm-tests amdgpu_target-none build_type=RelWithDebInfo cuda_arch=none
Aex@kwtexports%intel@18.0.2+EXO_ENABLE_RZ_SRC-ipo build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Afd@tk@0.2.1%intel@18.0.2-ipo+mpi+shared build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Amvapich@2.3%intel@18.0.2-alloc-cuda-debug+regcache+wrapperpath_ch3_rnk_bits32_fabrics+my
Ascent@0.7.1%intel@18.0.2+adios-adios2+babelflow-cuda+doc+dray-fides+fortran-ipo+mfem+mpi+openmp+
Aconduit@0.7.2%intel@18.0.2+adios+doc+doxygen+fortran+hd5+hd5f_compat-ipo+mpi-python+shared-5
Ahd@5@1.8.22%intel@18.0.2+cxx+fortran+hl-ipo+java+mpi+shared+szip-build+threadsafetools api-def
Apkgconf@1.8.0%intel@18.0.2 arch=linux-rhel7-ivybridge
Azlib@1.2.11%intel@18.0.2+optimize+pic+shared arch=linux-rhel7-ivybridge
Adray@0.1.6%intel@18.0.2-cuda-logging+mpi+openmp+shared+stats-test+utils cuda_arch=none patches
Aappcomp@0.3%intel@18.0.2+mpi+openmp+shared arch=linux-rhel7-ivybridge
Amem@1.3.0%intel@18.0.2+amg+axom+camp+conduit-cuda-debug+examples-gnutls-glib-lapack-ll
Amp@superlu-dist+threadsafe+mpire+zlib amdapu_target-none cuda_arch=none time-auto arch=linux-rhel7-ivybridge
Aaxom@0.5.0%intel@18.0.2+cpp14-cuda-debug+examples+fortran+hd5-ipo+lua+mfem+
Araja@0.12.1%intel@18.0.2-cuda-examples+exercises-ipo+openmp-roc-rcm+shared-tests+amd
Aumpire@5.0.1%intel@18.0.2+c-cuda-devicestest+examples-fortran-ipo+numa+openmp+roc
9b493dbf81f tests=none arch=linux-rhel7-ivybridge
Ahypr@2.22.0%intel@18.0.2+complex-cuda-debug+int64-internal-superlu-mixed+int+mpi+open
Anetlib-lapack@3.6.1%intel@18.0.2-external-blas-ipo+lapack+shared-xblas build_type
Ametis@5.1.0%intel@18.0.2-gdb-int64+real64+shared build_type=Release patches=4991da9358
Anetcdf-c@4.8.0%intel@18.0.2-dap-fsync+hd44-ipo+mpi-parallel-netcdf+pic+shared arch=lin
Amd@1.4.2%intel@18.0.2+Zslib+axom arch=linux-rhel7-ivybridge
Avtk-h@0.7.1%intel@18.0.2+contourtree-cuda-logging+mpi+openmp+serial+shared cuda_arch=none and
Avtk-m@1.6.0%intel@18.0.2+64bit+ascent_types-cuda+double+precision+hip-ipo+logging+mpi+o
Acaliper@2.6.0%intel@18.0.2+adiak-cuda+fortran+gotcho-ipo+libdw+libnum+nd+mpi+papi+scamper+
Alibunwind@1.5.0%intel@18.0.2-pic+xx-zlib arch=linux-rhel7-ivybridge
Aappi@6.0.1%intel@18.0.2-cuda-example-infiniband+msensors-nvml+powercap+rapl+sde+shared+sta
Apython@3.7.2%intel@18.0.2+bz2+cctypes+dbm+debug+libxml2+lzma-nis-optimizations+pic+pyexpat+pyt
b843_c129e34472ee39b1083c38a49a06a8573d8b308743a120ca4c0818cdf1801 arch=linux-rhel7-ivybridge
Aabzip@2.1.0%intel@18.0.2-debug+pic+shared arch=linux-rhel7-ivybridge
Aexpat@2.4.1%intel@18.0.2+libss arch=linux-rhel7-ivybridge
Alibffi@3.3%intel@18.0.2 arch=linux-rhel7-ivybridge
Alibmd@1.3%intel@18.0.2 arch=linux-rhel7-ivybridge
Agdbm@1.21%intel@18.0.2 arch=linux-rhel7-ivybridge
Areadline@8.1%intel@18.0.2 arch=linux-rhel7-ivybridge
Ancurses@6.2%intel@18.0.2+symlinks+termlib+abi=none arch=linux-rhel7-ivybridge
Agettext@0.19.8.1%intel@18.0.2+bzip2+curses+git-libunistring+libxml2+tar+xxz patches=9acdb4
Alibffi@3.3%intel@18.0.2 patches=26f26c6f29a7ce9b370a3ab2610f99365b4bd7b82e7c31df41a33370
Aopenssl@1.1.1%intel@18.0.2+docs+systemcerts arch=linux-rhel7-ivybridge
Aperl@5.16.3%intel@18.0.2+cpanm+shared+threads patches=0e0c10ed90aeb0459cd851f88081d4
Asqlite@3.35.0%intel@18.0.2+column-metadata+fts+functions+rtree arch=linux-rhel7-ivybridge
Autil-linux-wvd@2.36.2%intel@18.0.2 arch=linux-rhel7-ivybridge
Axz@5.2.5%intel@18.0.2-pic+libs+shared+static arch=linux-rhel7-ivybridge
Ae14@develop%intel@18.0.2 arch=linux-rhel7-ivybridge
Alua@5.1.5%intel@18.0.2+pcfiles+shared arch=linux-rhel7-ivybridge
Anzip@6.0%intel@18.0.2 arch=linux-rhel7-ivybridge
Agotcha@1.0.3%intel@18.0.2-ipo-test build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Airep@2021.06.22%intel@18.0.2-ipo build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Alua-luapost@35.0%intel@18.0.2 arch=linux-rhel7-ivybridge
Apy-lua@1.0%intel@18.0.2 arch=linux-rhel7-ivybridge
Apy-cython@0.29.24%intel@18.0.2 arch=linux-rhel7-ivybridge
Apy-setuptools@57.4.0%intel@18.0.2 arch=linux-rhel7-ivybridge
Apy-mpi4py@3.0.3%intel@18.0.2 arch=linux-rhel7-ivybridge
Ateila@1.0%intel@18.0.2-ipo build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Aleos@8.3.4.1%intel@18.0.2-cuda-debug+filters+raja+siloum+umpire arch=linux-rhel7-ivybridge
Aboost@1.76.0%intel@18.0.2+atomic+chrono+clanglibcpp+container+context+coroutine+date_time+debug+e
alization+shared+signals+singlethreaded+system+taggedlayout+test+thread+timer+versionedlayout+wave+cxstd+98+visibility+hid
Asilo@4.10.2%intel@18.0.2+fortran+fpzip+hd5+hzipp+mpi+shared-silex patches=7b5a1dc2a0e358e670
Aopacity@2.1.3.1%intel@18.0.2-ipo build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Aphyscsutils@2.4.201111.04%intel@18.0.2-ipo build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Aransbox@0.2.2%intel@18.0.2-cuda-debug+thr arch=linux-rhel7-ivybridge
Aselene@2.4.0%intel@18.0.2-ipo+mpi build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Anuclear@184%intel@18.0.2 arch=linux-rhel7-ivybridge
Arng@3.0%intel@18.0.2-debug arch=linux-rhel7-ivybridge
Atdf@2.3.60%intel@18.0.2-cuda-ipo+mpi build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Atribol@2021.7.21%intel@18.0.2-ipo build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Amranda@2021.9.2%intel@18.0.2-debug-ipo build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Ascmr@2021.2.16%intel@18.0.2-debug-ipo-silo build_type=RelWithDebInfo arch=linux-rhel7-ivybridge
Aoverlink@2.1.2%intel@18.0.2-cuda-debug+thr arch=linux-rhel7-ivybridge
```

Fully concretized MARBL environment

# Spack workflow for developer environment

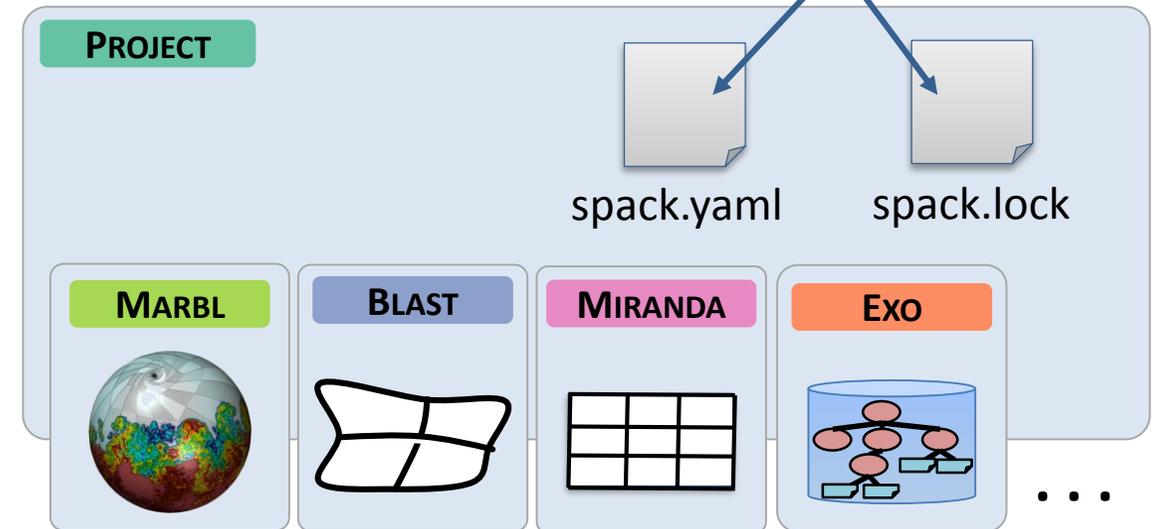
## CLI workflow

```
$ git clone ssh://git@rzgitlab.llnl.gov:7999/mapp/mapp  
$ cd mapp  
$ spack env activate .  
$ spack develop blast@develop  
$ spack develop miranda@develop  
$ srun -N 2 -n 16 --exclusive spack install
```

} Components to work on

spack can do multi-node builds

Top-level YAML manifest and lock file(s)



Subdirectories with components

# E4S is ECP's curated, Spack-based software distribution

- E4S is just a set of Spack packages
  - 60+ packages (297 including dependencies)
  - Growing to include all of ST and more
- Users can install E4S packages:
  - In their home directory
  - In a container
- Facilities can install E4S packages:
  - On bare metal
  - In a container
- Users and facilities can choose parts they want
  - `spack install` only the packages you want
  - Or just edit the list of packages (and configurations) you want in a `spack.yaml` file

```
spack:
  specs:
    - openpmc-api
    - py-libensemble^python@3.7.3
    - hypre
    - mfem
    - trilinos@12.14.1+dtk+intrepid2+shards
    - sundials
    - strumpack
    - superlu-dist
    - superlu
    - tasmanian
    - mercury
    - hdf5
    - adios2
    - dyninst
    - pdt
    - tau
    - hpctoolkit
    - adios
    - darshan-runtime
    - darshan-util
    - veloc
    - scr
    - parallel-netcdf
    - qthreads
    - papyrus@develop
    - bolt
    - raja
    - upcxx
    - kokkos+openmp
    - openmpi
    - umpire
    - libquo
    - globalarrays
    - gotcha
    - caliper
    - papi
    - py-jupyterhub
    - zfp
    - sz
    - libnrm
    - rempi
    - ninja
    - kokkos-kernels
    #- turbine
    #- aml
    #- unifyfs
    #- flecsi+cinch
    #- petsc
    #- faodel
  packages:
    all:
      providers:
        mpi: [spectrum-mpi]
      target: [ppc64le]
  cuda:
    buildable: false
    version: [10.1.243]
    modules:
      cuda@10.1.243: cuda/10.1.243
  spectrum-mpi:
    buildable: false
    version:
      - 10.3.1.2
    modules:
      spectrum-mpi@10.3.1.2: spectrum-mpi/10.3.1.2-20200121
  config:
    misc_cache: $spack/cache
    build_stage: $spack/build-stage
    install_tree: $spack/$padding:512
  view: false
  concretization: separately
```

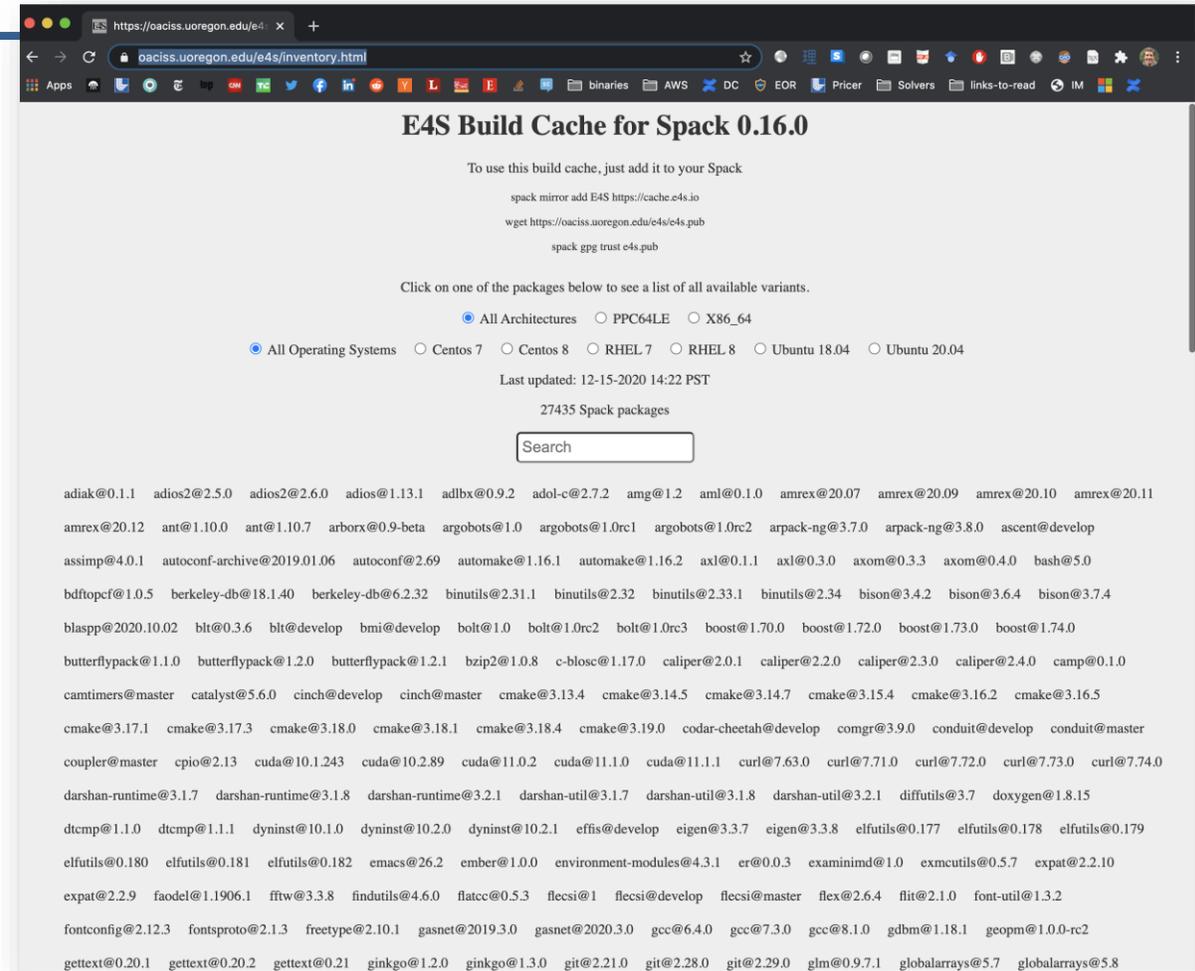


Actual E4S manifest (`spack.yaml`) for OLCF Ascent

More on E4S at <https://e4s.io>

# E4S team has built a binary cache with over 50,000+ Spack binary packages

- Built for multiple OS's, architectures
- E4S team is working with ECP projects to accelerate their build pipelines
- Improved performance of cloud CI for one project by 10-100x
  - Previously, builds took too long for free cloud CI
  - Project can now iterate faster using Spack/E4S binaries
- We are rapidly building out binary build capabilities for Spack
  - Aim to have optimized binaries for most platforms in Frontier/El Capitan timeframe

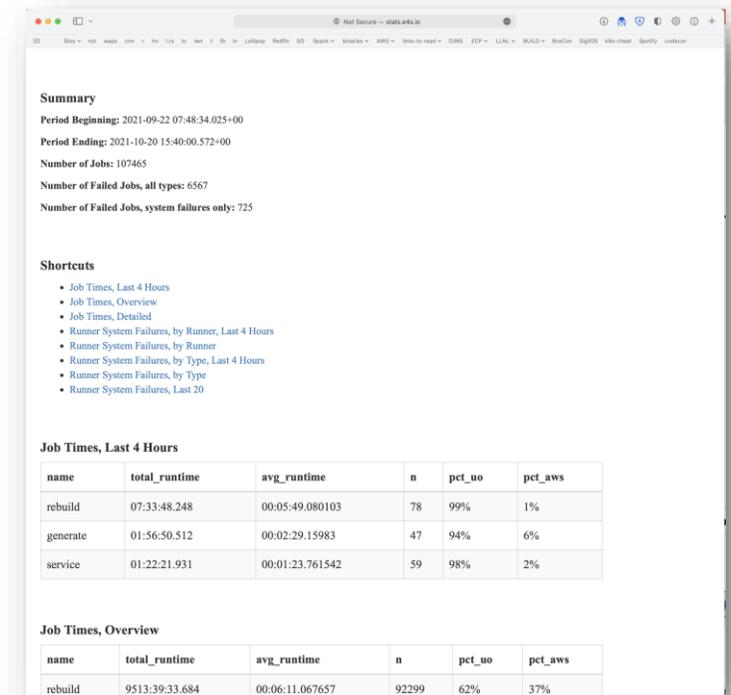


<https://oaciss.uoregon.edu/e4s/inventory.html>



# Future CI directions focus on scalability and testing

- Scaling tests up to handle every PR has been very difficult
  - Driven by GitLab
  - Using Kubernetes builders
  - Using a cluster at U. Oregon
- Concretization of large environments was slowing turnaround
  - 55 min to concretize E4S environment (each spec separately)
  - Brought this down to 2.5 min this past weekend
- Amazon and E4S/UO team helping to pinpoint errors
- Once we have a stable, rolling release of spack develop branch, we'll make the build cache public
  - Rolling binaries for develop
  - Long-lived snapshots for each release



Summary

Period Beginning: 2021-09-22 07:48:34.025+00  
Period Ending: 2021-10-20 15:40:00.572+00  
Number of Jobs: 107465  
Number of Failed Jobs, all types: 6567  
Number of Failed Jobs, system failures only: 725

Shortcuts

- Job Times, Last 4 Hours
- Job Times, Overview
- Job Times, Detailed
- Runner System Failures, by Runner, Last 4 Hours
- Runner System Failures, by Runner
- Runner System Failures, by Type, Last 4 Hours
- Runner System Failures, by Type
- Runner System Failures, Last 20

Job Times, Last 4 Hours

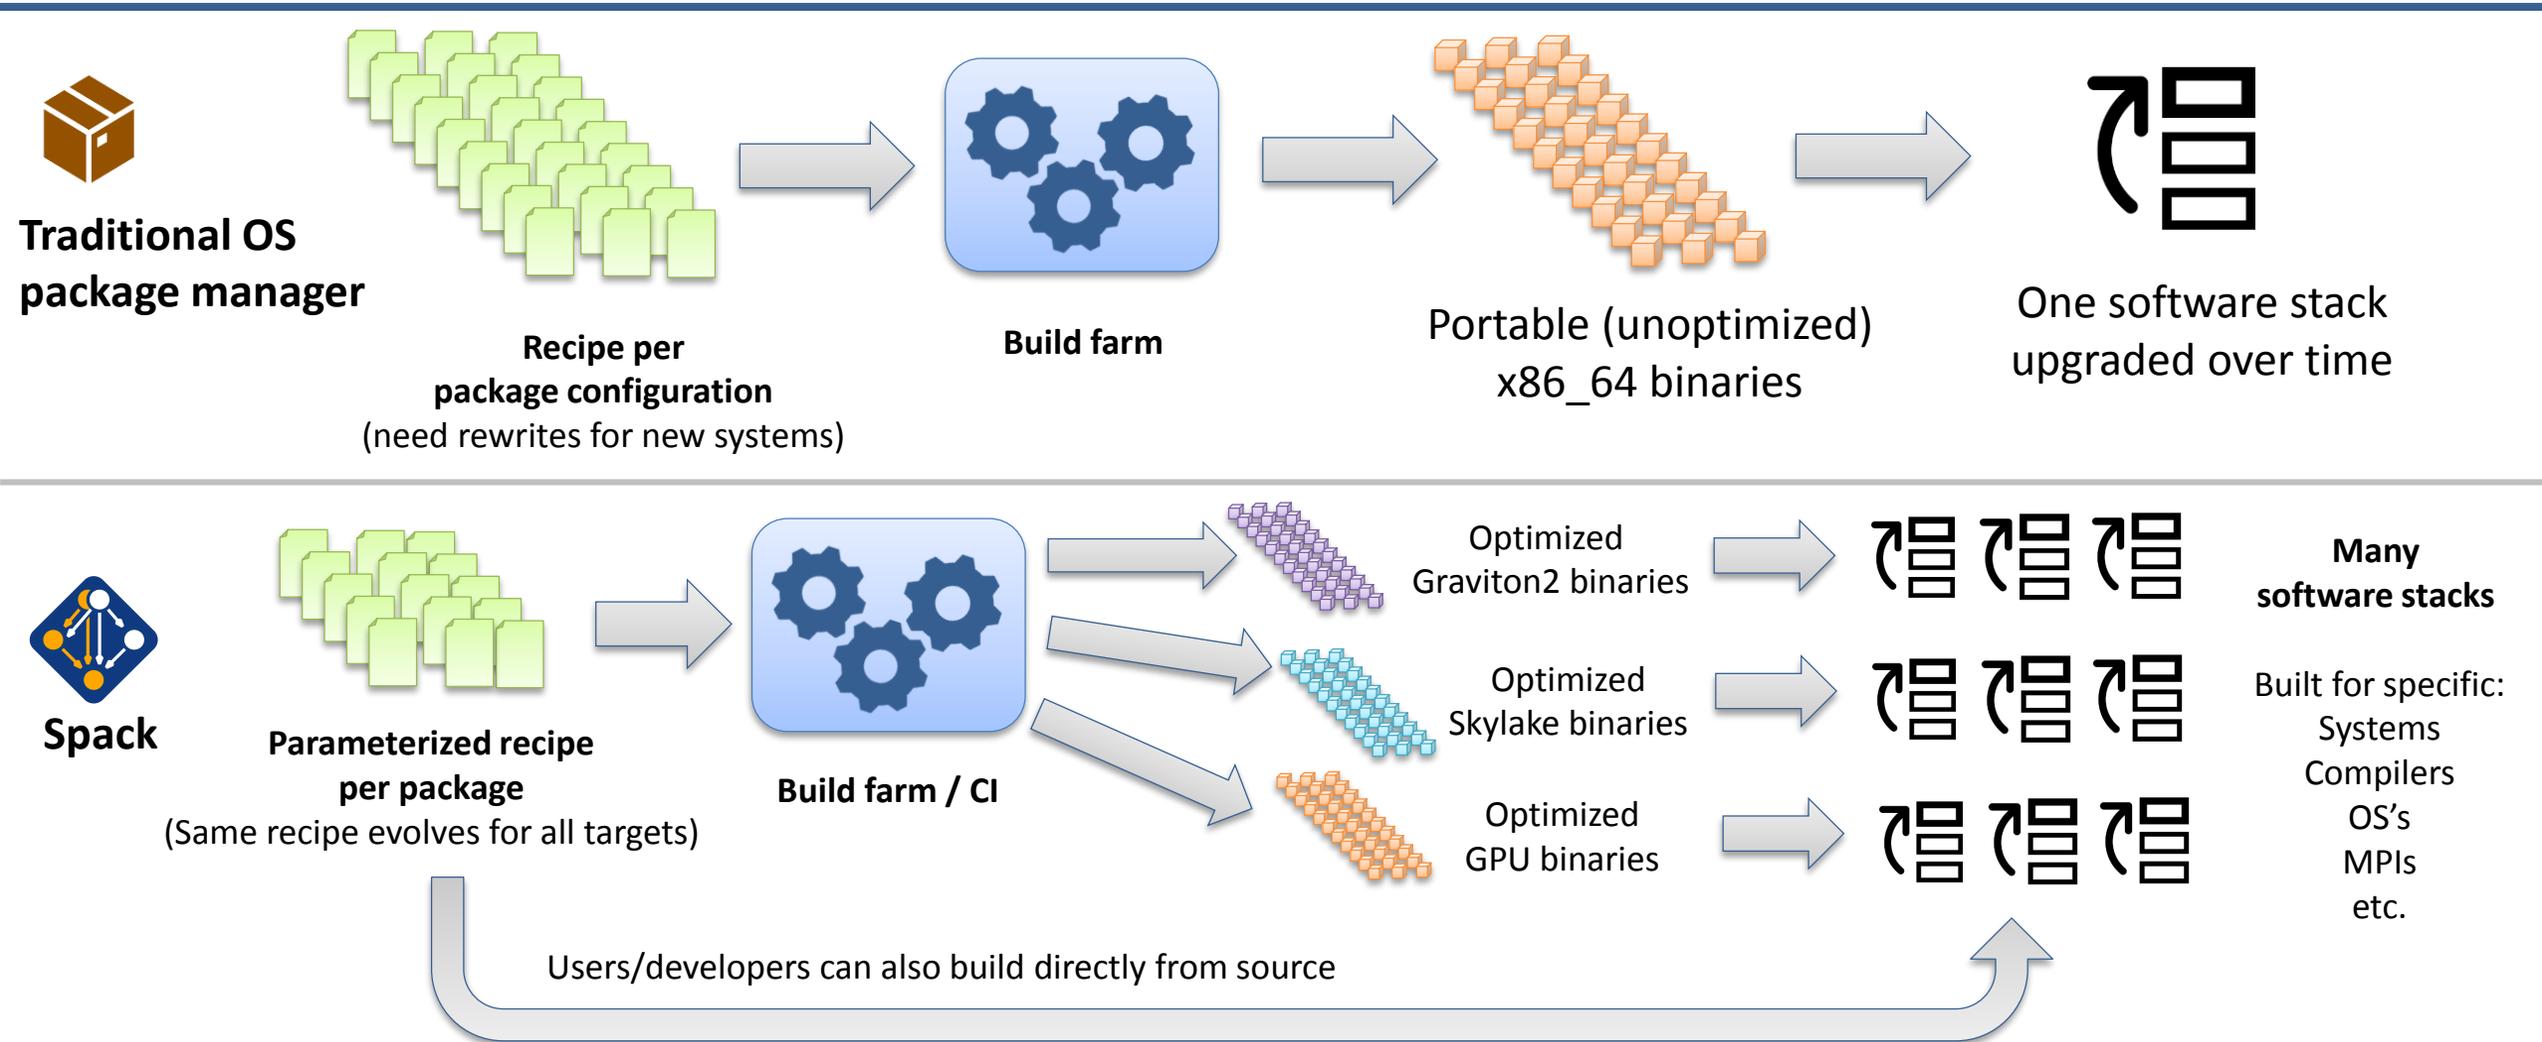
name	total_runtime	avg_runtime	n	pct_uo	pct_aws
rebuild	07:33:48.248	00:05:49.080103	78	99%	1%
generate	01:56:50.512	00:02:29.15983	47	94%	6%
service	01:22:21.931	00:01:23.761542	59	98%	2%

Job Times, Overview

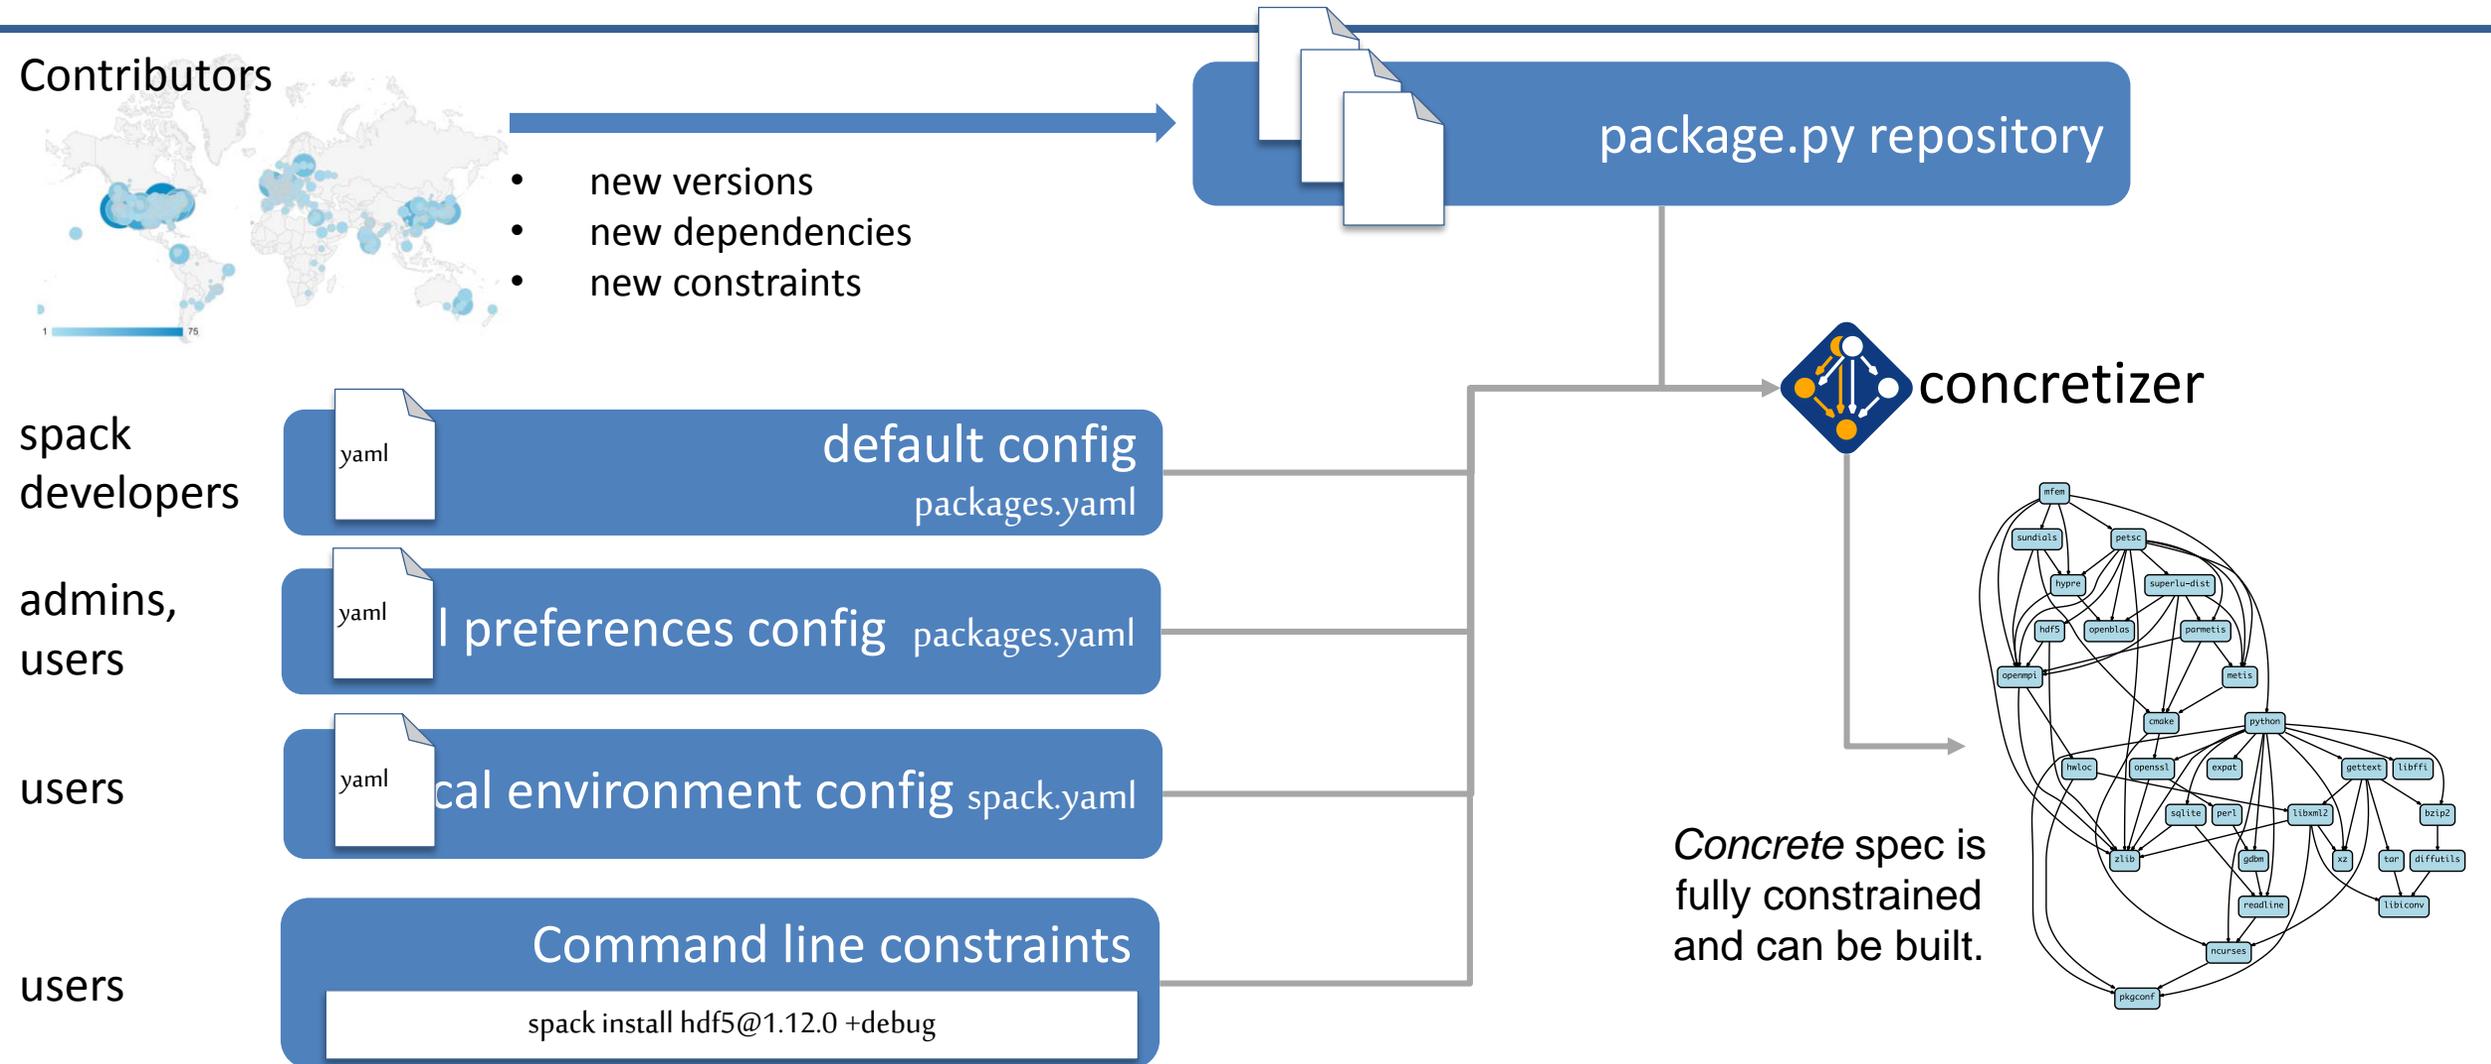
name	total_runtime	avg_runtime	n	pct_uo	pct_aws
rebuild	9513:39:33.684	00:06:11.067657	92299	62%	37%

<http://stats.e4s.io>

# Spack's model lowers the maintenance burden of optimized software stacks

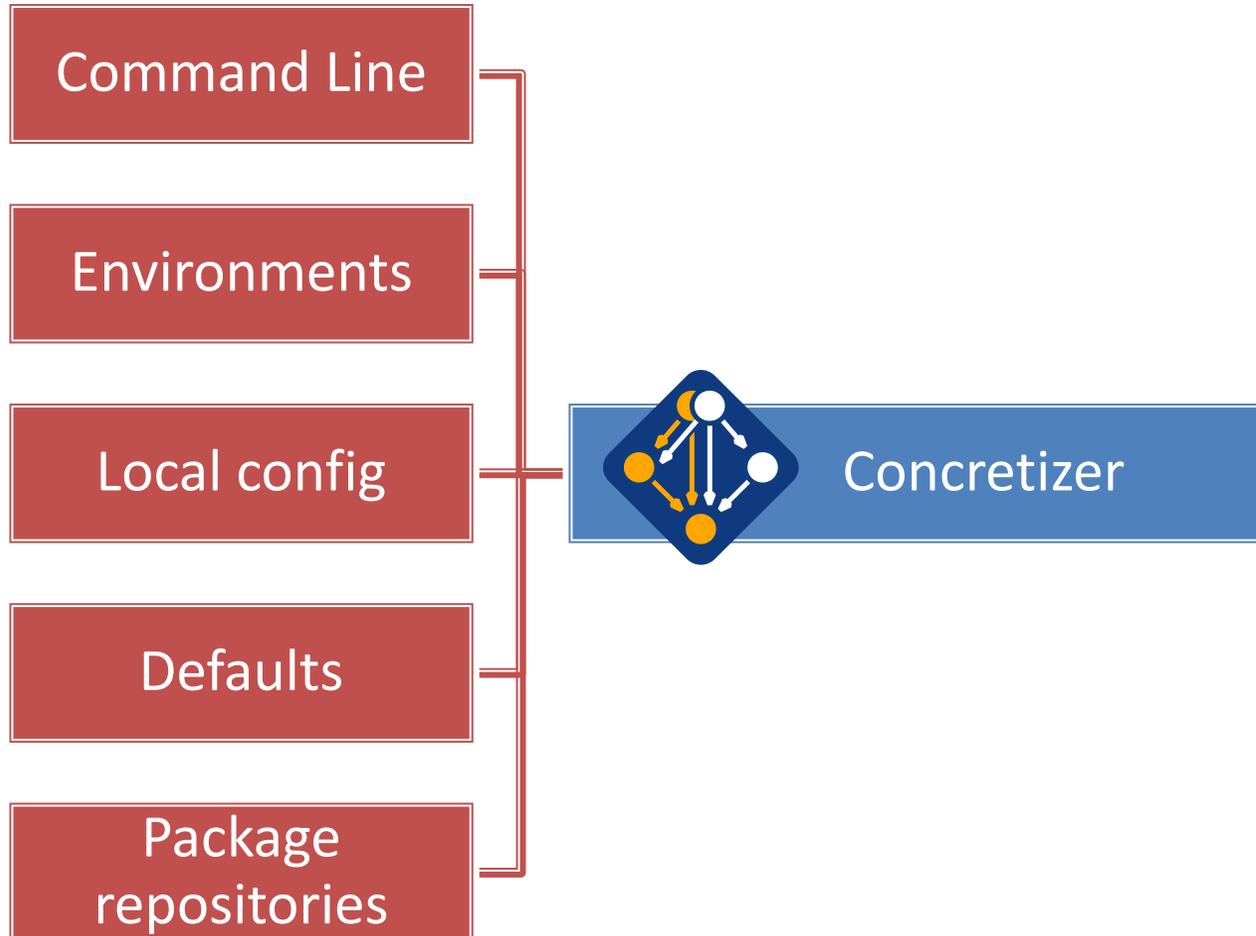


# High level view of a Spack package build



# Spack's concretizer has gotten pretty complicated

## Sources for constraints



- Current implementation is ad-hoc:
  - Traverse the DAG
  - Evaluate conditions, add dependencies
  - Fill in defaults from many sources
  - Repeat until DAG doesn't change
- Issues:
  - Limited support for backtracking causes some graphs to resolve incorrectly
  - Some constraints are strictly ordered
  - Lots of conditional complexity
- Design doesn't scale to all the criteria
  - Hard to add new features/logic
  - Can be slow

# The new concretizer will be default in 0.17

- New concretizer leverages Clingo (see [potassco.org](http://potassco.org))
- Clingo is an Answer Set Programming (ASP) solver
  - ASP looks like Prolog; leverages SAT solvers for speed/correctness
  - ASP program has 2 parts:
    1. Large list of facts generated from our package repositories and config
      - 20,000 – 30,000 facts is typical – includes dependencies, options, etc.
    2. Small logic program (~800 lines), including constraints and optimization criteria
- New algorithm on the Spack side is conceptually simpler:
  - Generate facts for all possible dependencies, send to logic program
  - Optimization criteria express preferences more clearly
  - Build a DAG from the results
- New concretizer solves many specs that current concretizer can't
  - Backtracking is a huge win – many issues resolved
  - Currently requires user to install clingo with Spack
  - Solver will be automatically installed from public binaries in 0.17.0

```
-----  
% Package: ucx  
-----  
%  
version_declared("ucx", "1.6.1", 0).  
version_declared("ucx", "1.6.0", 1).  
version_declared("ucx", "1.5.2", 2).  
version_declared("ucx", "1.5.1", 3).  
version_declared("ucx", "1.5.0", 4).  
version_declared("ucx", "1.4.0", 5).  
version_declared("ucx", "1.3.1", 6).  
version_declared("ucx", "1.3.0", 7).  
version_declared("ucx", "1.2.2", 8).  
version_declared("ucx", "1.2.1", 9).  
version_declared("ucx", "1.2.0", 10).  
  
variant("ucx", "thread_multiple").  
variant_single_value("ucx", "thread_multiple").  
variant_default_value("ucx", "thread_multiple", "False").  
variant_possible_value("ucx", "thread_multiple", "False").  
variant_possible_value("ucx", "thread_multiple", "True").  
  
declared_dependency("ucx", "numactl", "build").  
declared_dependency("ucx", "numactl", "link").  
node("numactl") :- depends_on("ucx", "numactl"), node("ucx").  
  
declared_dependency("ucx", "rdma-core", "build").  
declared_dependency("ucx", "rdma-core", "link").  
node("rdma-core") :- depends_on("ucx", "rdma-core"), node("ucx").  
  
-----  
% Package: util-linux  
-----  
%  
version_declared("util-linux", "2.29.2", 0).  
version_declared("util-linux", "2.29.1", 1).  
version_declared("util-linux", "2.25", 2).  
  
variant("util-linux", "libuuid").  
variant_single_value("util-linux", "libuuid").  
variant_default_value("util-linux", "libuuid", "True").  
variant_possible_value("util-linux", "libuuid", "False").  
variant_possible_value("util-linux", "libuuid", "True").  
  
declared_dependency("util-linux", "pkgconfig", "build").  
declared_dependency("util-linux", "pkgconfig", "link").  
node("pkgconfig") :- depends_on("util-linux", "pkgconfig"), node("util-linux").  
  
declared_dependency("util-linux", "python", "build").  
declared_dependency("util-linux", "python", "link").  
node("python") :- depends_on("util-linux", "python"), node("util-linux").
```

Some facts for the HDF5 package

# The new concretizer enables significant simplifications to packages, particularly complex constraints in SDKs

- Dependencies and other constraints within SDKs could get very messy
- The new concretizer removes the need for some of the more painful constructs
  - When conditions are now much more general
  - Can be solved together with other constraints.
- Also allows for new constructs, like specializing dependencies

**In some cases we needed cross-products of dependency options:**

*Before*

```
depends_on('foo+A+B', when='+a+b')
depends_on('foo+A~B', when='+a~b')
depends_on('foo~A+B', when='~a+b')
depends_on('foo~A~B', when='~a~b')
```

*After*

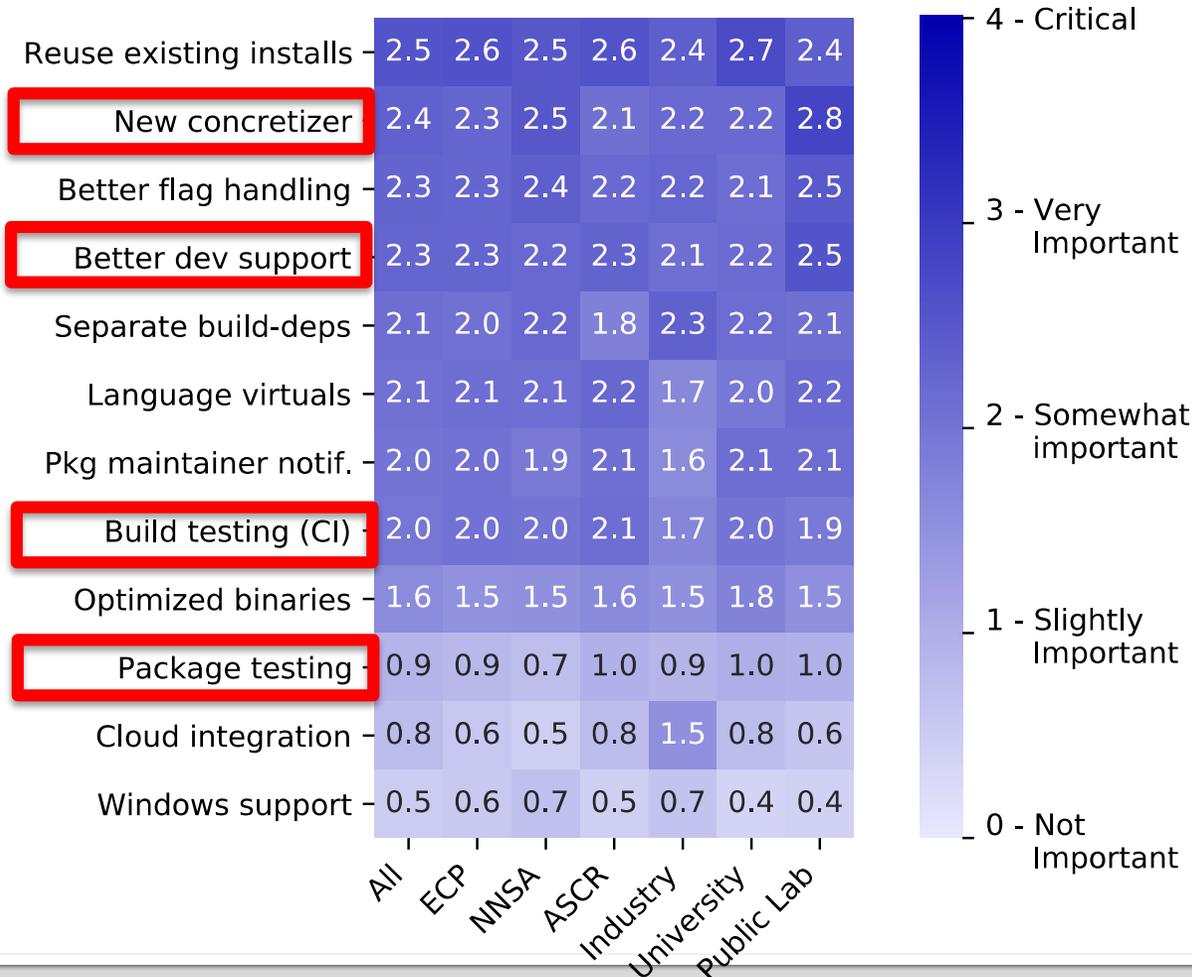
```
depends_on('foo')
depends_on('foo+A', when='+a')
depends_on('foo+B', when='+b')
```

**Specializing a virtual did not previously work:**

```
depends_on('blas')
depends_on(
    'openblas threads=openmp', when='^openblas'
)
```

# Four of the top six most wanted features in Spack were tied to the new concretizer

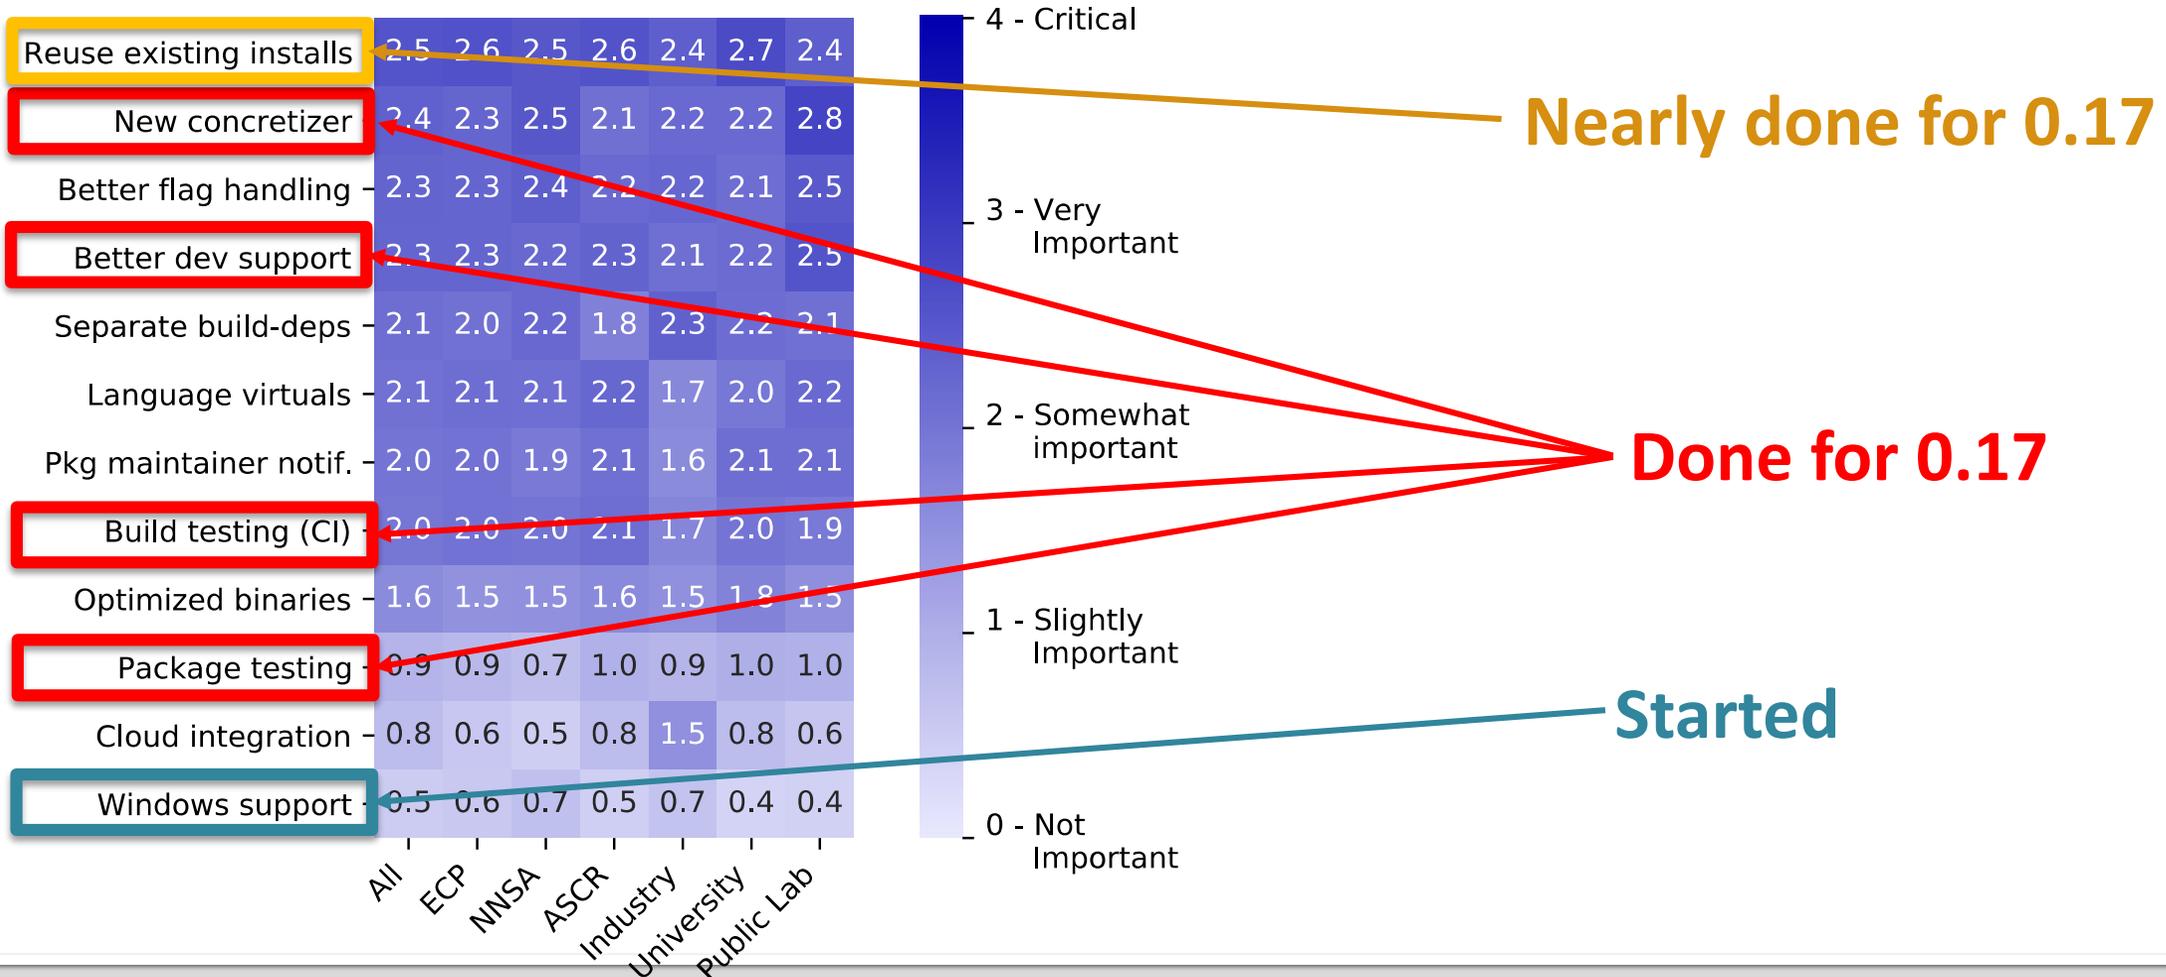
Average feature importance by workplace



- Complexity of packages in Spack is increasing
  - many more package solves require backtracking than a year ago
  - Many variants, conditional dependencies, special compiler requirements
- More aggressive reuse of existing installs requires better dependency resolution
  - Need to be able to analyze how to configure the build to work with installed packages
- Separate resolution of build dependencies also requires a more sophisticated solver
  - Makes the solve even more combinatorial
  - Needed to support mixed compilers, version conflicts between different package's build requirements

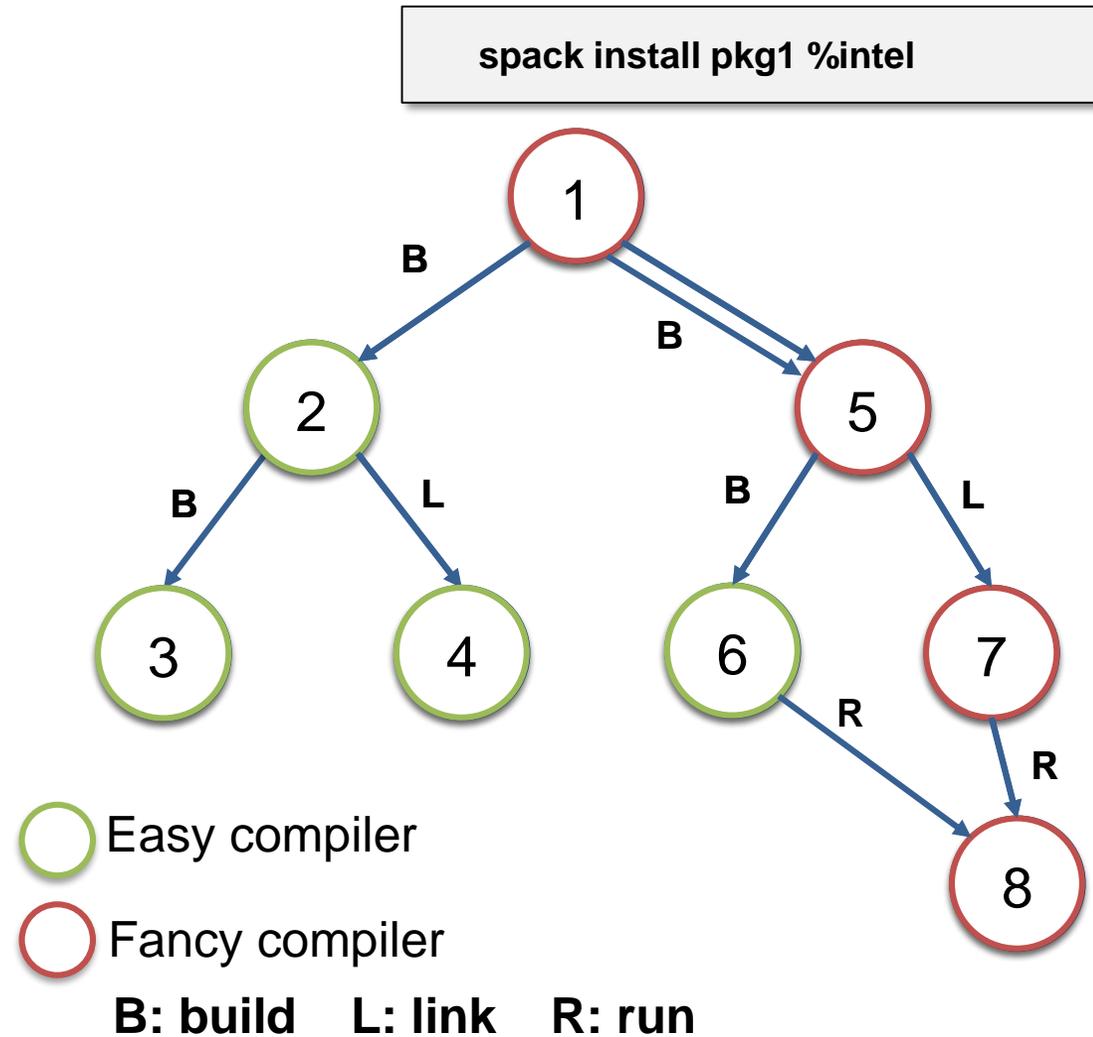
# Four of the top six most wanted features in Spack are tied to the new concretizer

Average feature importance by workplace



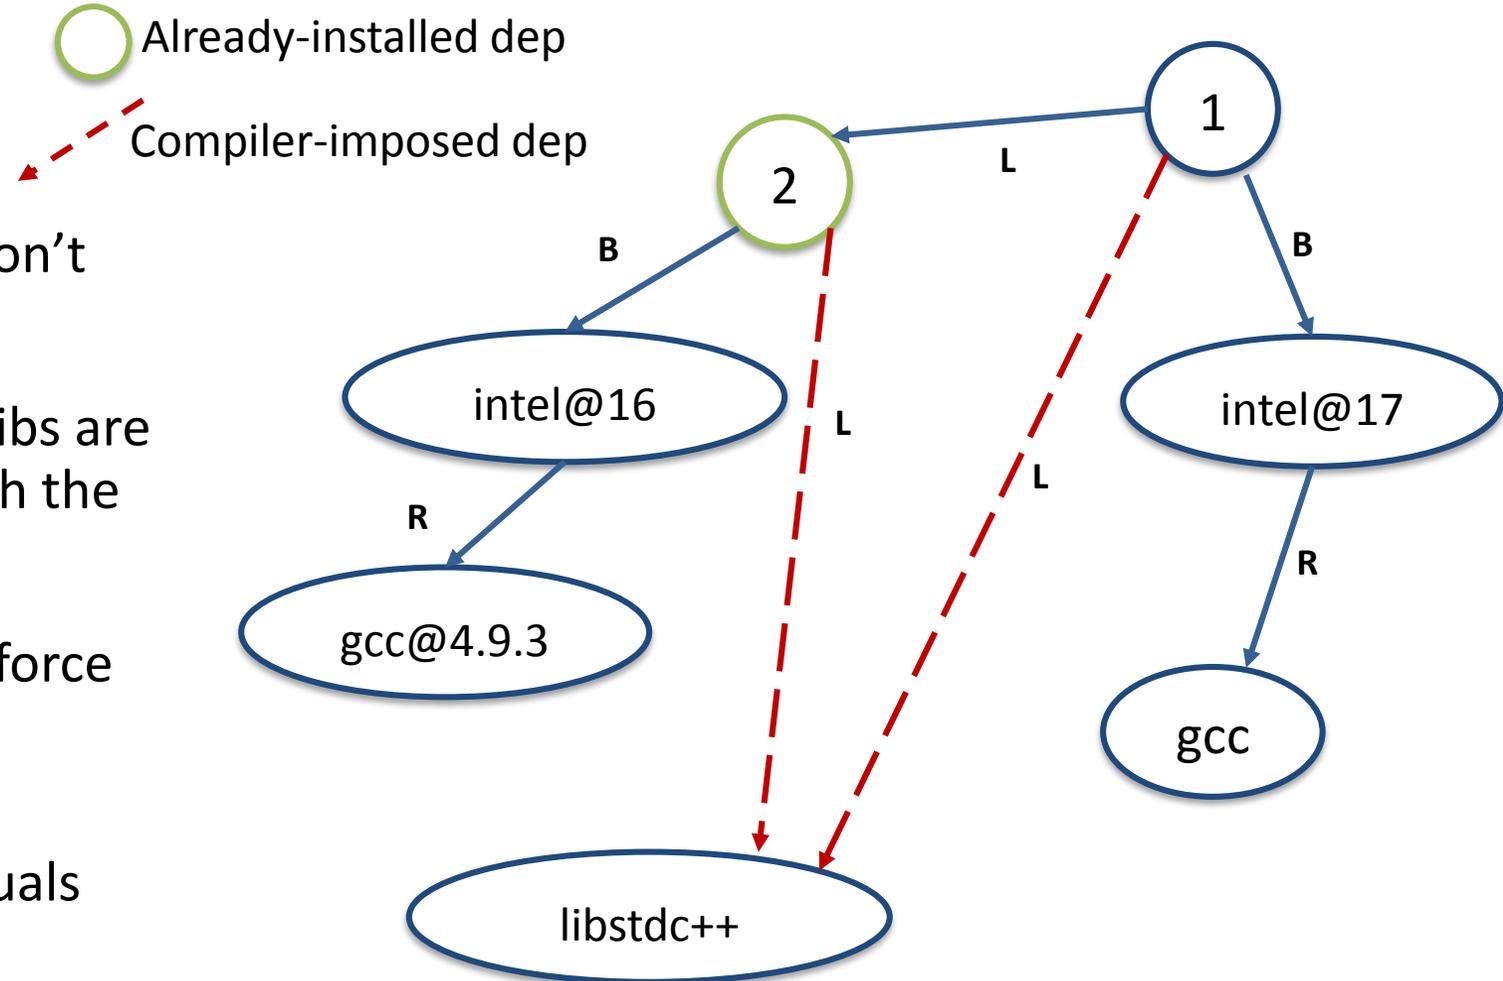
# Separate concretization of build dependencies

- We want to:
  - Build build dependencies with the "easy" compilers
  - Build rest of DAG (the link/run dependencies) with the fancy compiler
- This required significant concretizer modifications
- Gets into issues like bootstrapping



# New compiler dependency model

- Now consistency is enforced via link dependencies from 1 and 2
- If the libstdc++ versions from 1 and 2 don't match, then this won't resolve
- If they do, then we know that the C++ libs are compatible and can build this, even with the old dependency.
- We currently use some heuristics to enforce
  - Moving to SAT makes a lot of sense, again
- Will enable us to support language virtuals



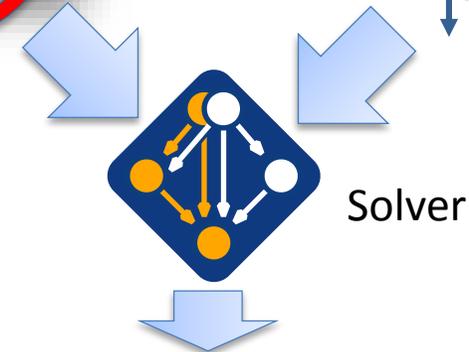
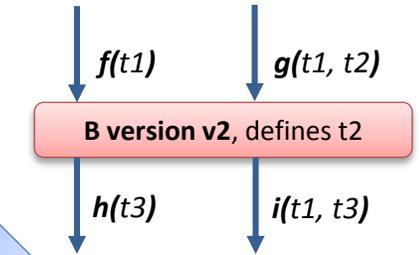
# Ongoing research: BUILD is a 3-year research project, started at LLNL in 2020

- Basic premise: humans can't generate all the compatibility constraints
  - Version ranges, conflicts, in Spack packages not precise
  - rely on maintainers to get right.
- BUILD aims to understand software compatibility at the binary level
  - Develop ABI compatibility models
  - Enable *automatic* and ABI-compatible reuse of system binaries, foreign binary packages
- WIP: better dependency solvers can enable users to solve *around* system dependencies**
  - find "closest" match to a prior build, using new packages
  - Reproduce a prior build with new requirements

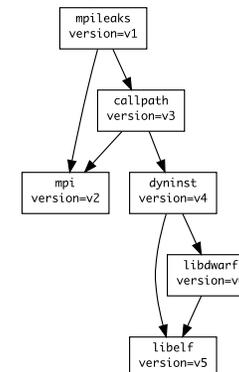
Human-generated constraints



Compatibility Models



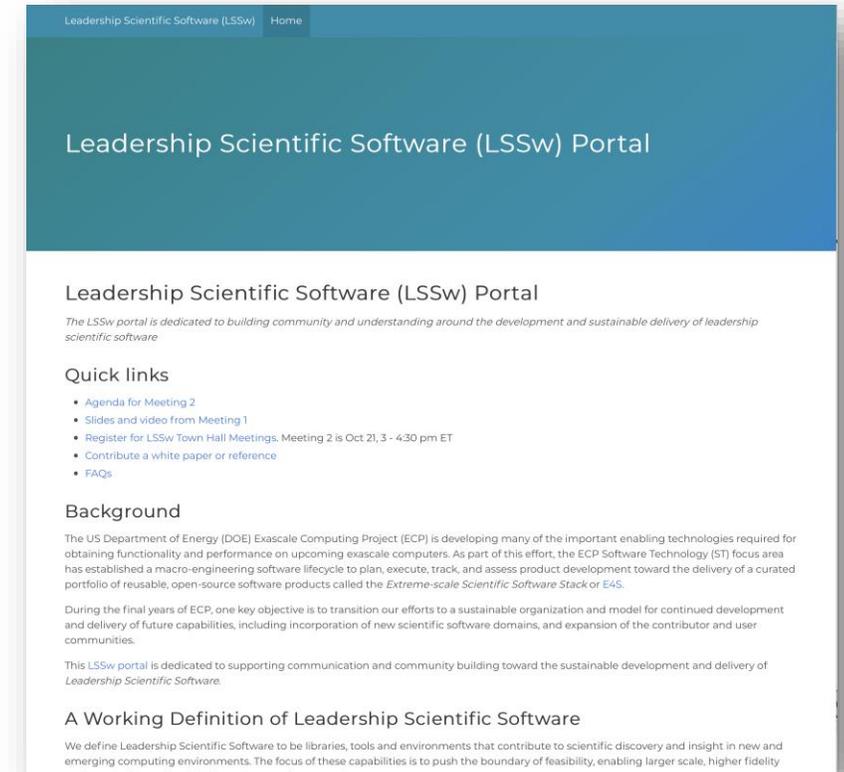
Solver



Resolved  
ABI-compatible  
Graph

# After ECP

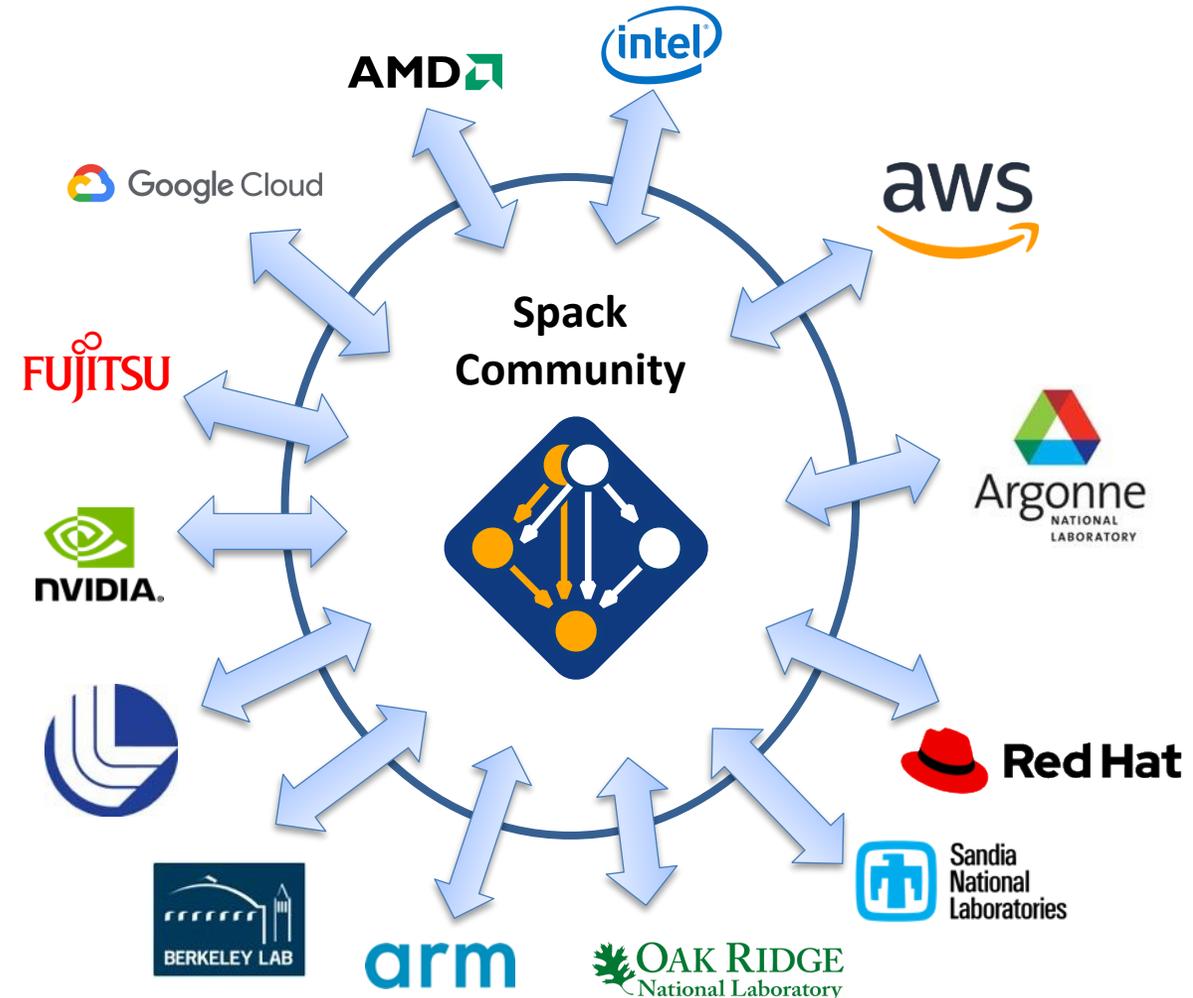
- We are looking at longer-term sustainability directions after ECP
- Opportunities (everything is in flux at this point):
  - ASCR Workshop on the Science of Scientific-Software Development and Use
    - just came out
  - Leadership Scientific Software Meeting series
    - <https://lssw.io>
- We want to be part of any post-ECP sustainability effort!
  - Likely some type of work in conjunction with E4S



The screenshot shows the Leadership Scientific Software (LSSw) Portal website. The header is a dark blue bar with the text "Leadership Scientific Software (LSSw) Home" on the left and "Leadership Scientific Software (LSSw) Portal" on the right. Below the header, the main content area is white. It features a title "Leadership Scientific Software (LSSw) Portal" followed by a subtitle: "The LSSw portal is dedicated to building community and understanding around the development and sustainable delivery of leadership scientific software". There is a "Quick links" section with a bulleted list: "Agenda for Meeting 2", "Slides and video from Meeting 1", "Register for LSSw Town Hall Meetings, Meeting 2 is Oct 21, 3 - 4:30 pm ET", "Contribute a white paper or reference", and "FAQs". Below this is a "Background" section with two paragraphs of text. The first paragraph discusses the US Department of Energy (DOE) Exascale Computing Project (ECP) and the ECP Software Technology (ST) focus area. The second paragraph discusses the transition to a sustainable organization and model. The third paragraph states the portal's dedication to supporting communication and community building. The final section is titled "A Working Definition of Leadership Scientific Software" and defines it as libraries, tools, and environments that contribute to scientific discovery and insight in new and emerging computing environments.

# Spack's long-term strategy is based around broad adoption and collaboration

- **Not sustainable without a community**
  - Broad adoption incentivizes contributors
  - Cloud resources and automation absolutely necessary
- **Spack preserves build knowledge in a cross-platform, reusable way**
  - Minimize rewriting recipes when porting
- **CI ensures builds continue to work as packages evolve**
  - Keep packages flexible but verify key configurations
- **Any suggestions on sustainability models would be appreciated!**





**Disclaimer**

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.