Spack: Package Management for HPC

ATPESC 2025 Software Sustainability Track

Todd Gamblin

Spack Project Lead
Distinguished Member of Technical Staff
Lawrence Livermore National Laboratory

THE LINUX FOUNDATION

We build codes from hundreds of small, complex pieces

Just when we're starting to solve the problem of how to create software using reusable parts, it founders on the nuts-and-bolts problems outside the software itself.

P. DuBois & T. Epperly. ***Why Johnny Can't Build***. Scientific Programming. Sep/Oct 2003.

- **Pros are well known:**
  - Teams can and must reuse each others' work
  - Teams write less code, meet deliverables faster
- **Cons:**
  - Teams must ensure that components work together
  - Integration burden increases with each additional library
  - Integration must be repeated with each update to components
  - **Components must be vetted!**
- **Managing changes over time is becoming intractable**

Build-time incompatibility; fail fast

Appears to work; subtle errors later

Modern scientific codes rely on icebergs of dependency libraries

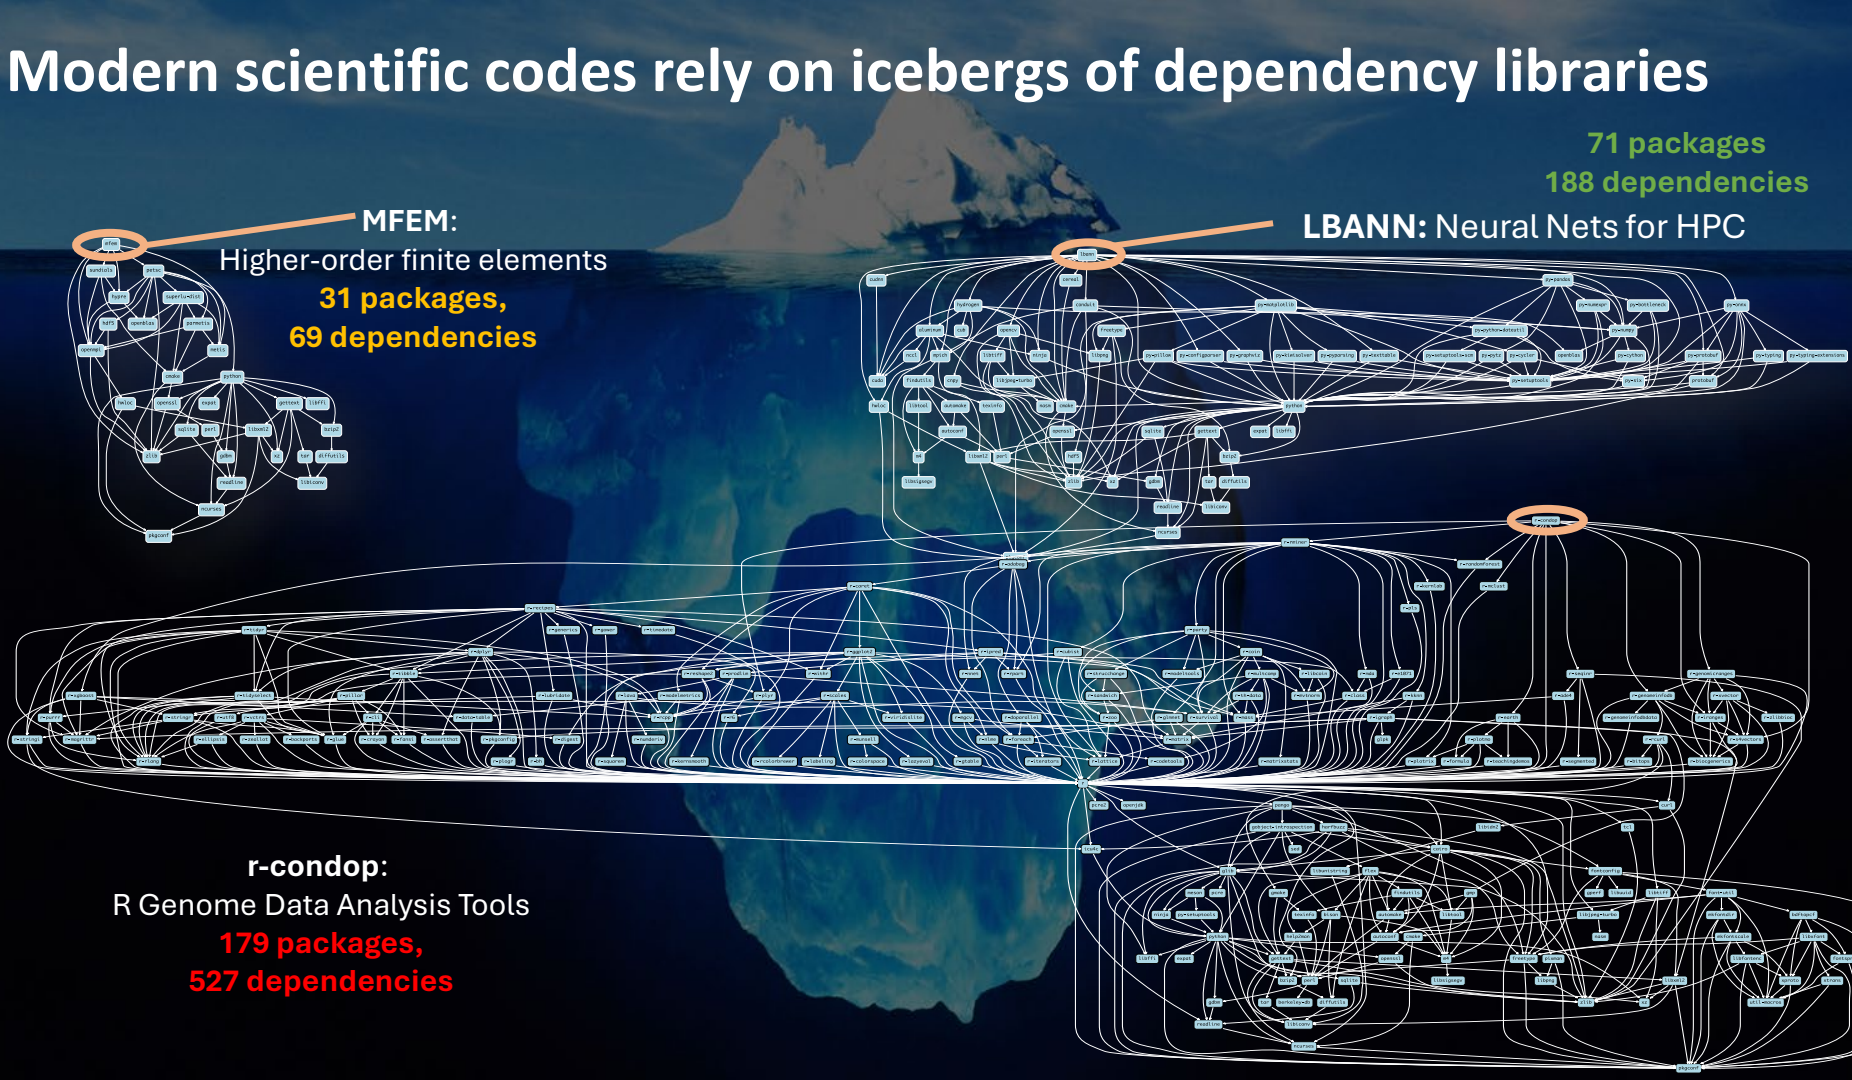

Even proprietary software builds on top of open source

LLNL ARES

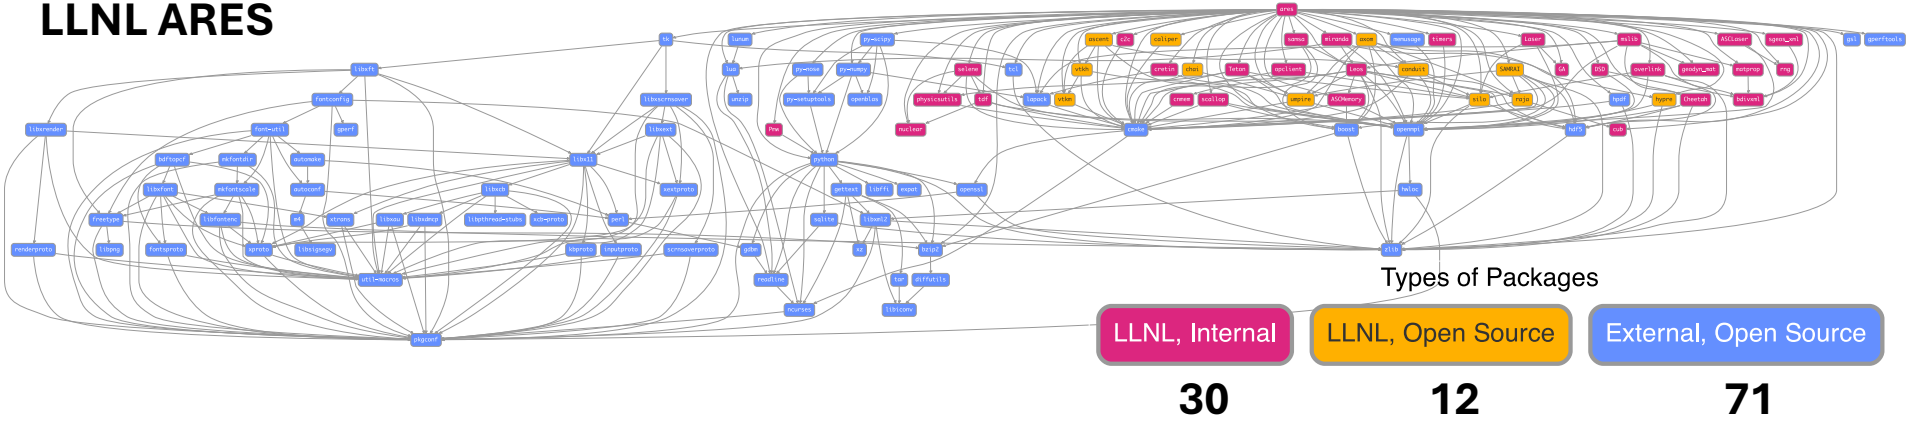

- Open source is critical for nearly every application
- We *cannot* replace all these OSS components with our own
  - How do we put them all together effectively?
  - Do you *have* to integrate this stuff by hand?

ECP’s E4S stack is even larger than these codes

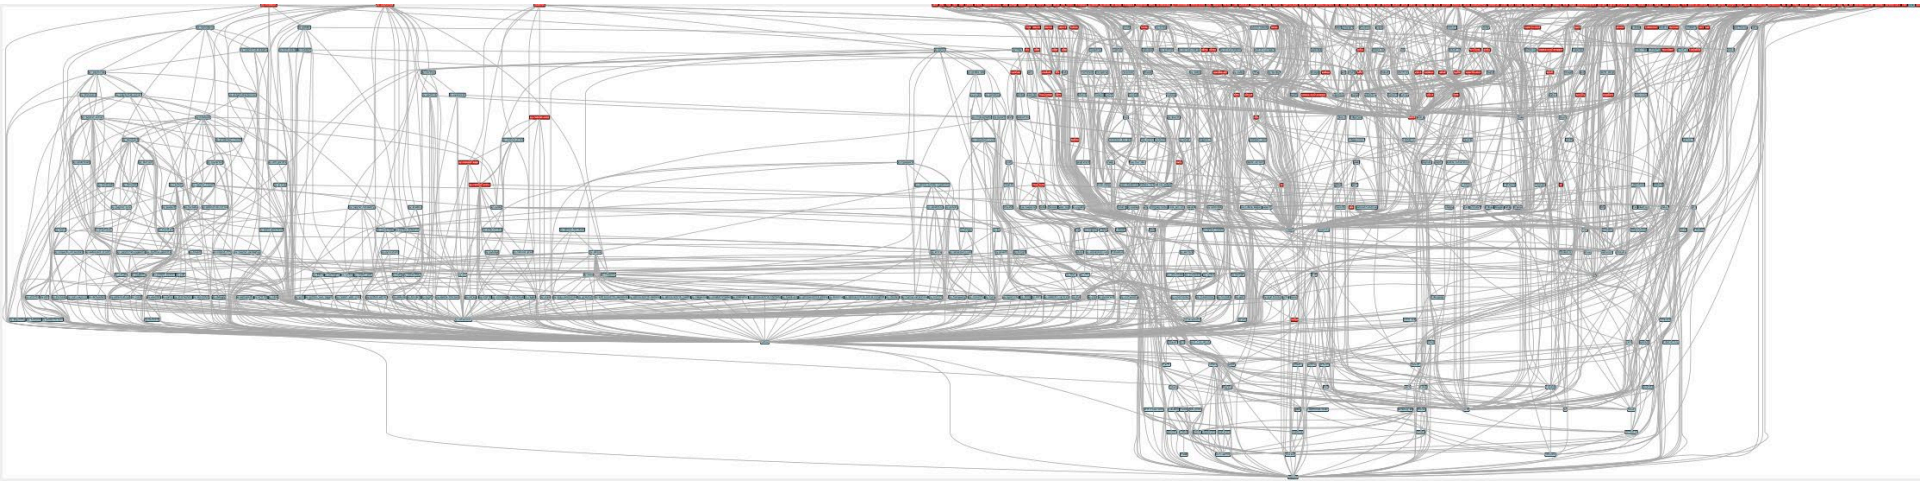

- Red boxes are the packages in it (about 100)
- Blue boxes are what *else* you need to build it (about 600)
- It’s infeasible to build and integrate all of this manually

Some history:
How to install software on a supercomputer, circa 2013

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

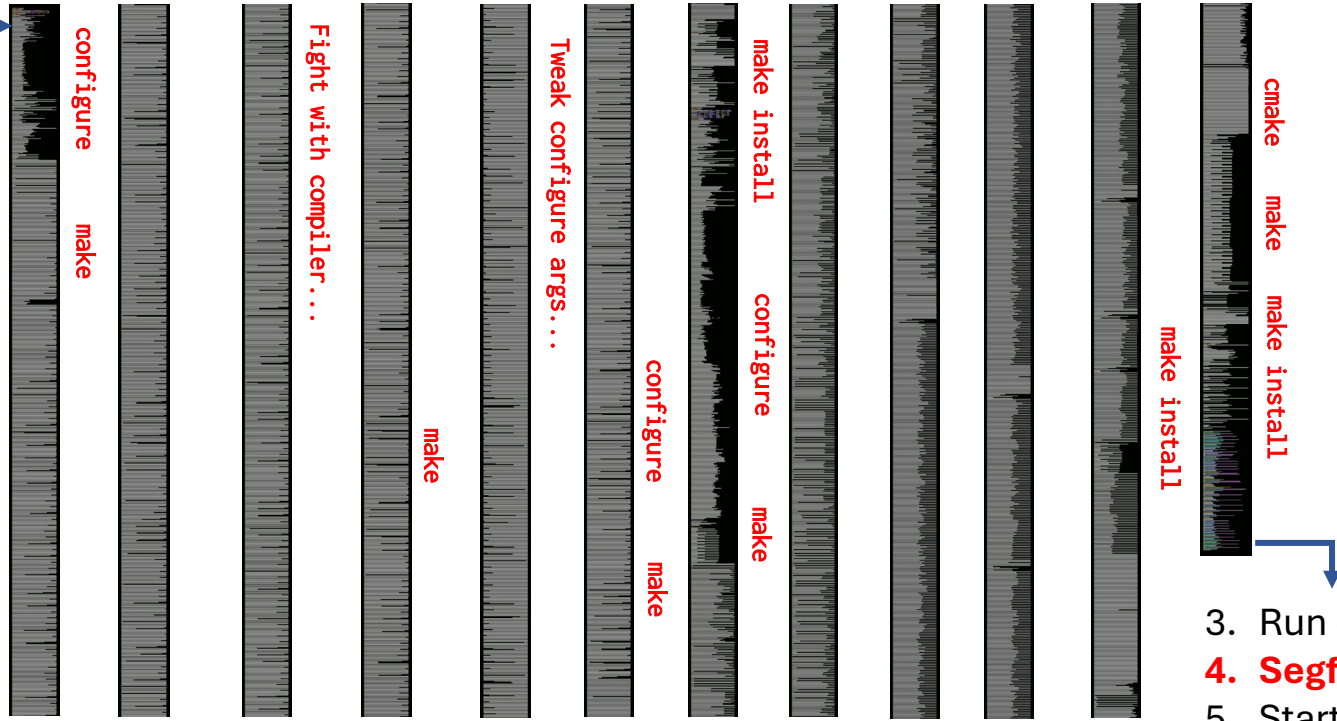

3. Run code
4. **Segfault!?**
5. Start over...

What does a package managers do to help?

- **Does not** replace Cmake/Autotools
- Manages ***dependencies***
  - Drives package-level build systems
  - Ensures **consistency** and **compatibility** among builds of packages in the ecosystem
- Stores **community knowledge**
  - Cache of package build recipes
  - Determining magic configure lines takes time

Package Manager

- Package installation
- Dependency relationships, conflicts
- May drive package-level build systems

High Level Build System

- Cmake, Autotools
- Handle library abstractions
- Generate Makefiles, etc.

Low Level Build System

- Make, Ninja
- Handles dependencies among *commands* in a single build

Many package managers (conda, pip, apt, etc.) make simplifying assumptions about the software ecosystem

- **1:1 relationship between source code and binary (per platform)**
  - Good for reproducibility (e.g., Debian)
  - Bad for performance optimization
- **Binaries should be as portable as possible**
  - What most distributions do
  - Again, bad for performance
- **Toolchain is the same across the ecosystem**
  - One compiler, one set of runtime libraries
  - Or, no compiler (for interpreted languages) and *just one language*

Outside these boundaries, users are typically on their own

High Performance Computing (HPC) violates many of these assumptions

- **Code is typically distributed as source**
  - With exception of vendor libraries, compilers
- **Often build many variants of the same package**
  - Developers' builds may be very different
  - Many first-time builds when machines are new
- **Code is optimized for the processor and GPU**
  - Must make effective use of the hardware
  - Can make 10-100x perf difference
- **Rely heavily on system packages**
  - Need to use optimized libraries that come with machines
  - Need to use host GPU libraries and network
- **Multi-language**
  - C, C++, Fortran, Python, others all in the same ecosystem

Some Supercomputers

Oak Ridge National Lab
Power9 / NVIDIA

RIKEN
Fujitsu/ARM a64fx

Lawrence Berkeley National Lab
AMD Zen / NVIDIA

Argonne National Lab
Intel Xeon / Xe

Oak Ridge National Lab
AMD Zen / Radeon

Lawrence Livermore National Lab
AMD Zen / Radeon

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
- **HPC containers are often optimized per-system**
  - Not clear that we can ever build one container for all facilities

We need something more flexible to **build** versions of containers

Overview & Community

THE LINUX FOUNDATION

Spack enables Software distribution for HPC

- Spack automates the build and installation of scientific software
- Packages are *parameterized,* so that users can easily tweak and tune configuration

No installation required: clone and go

```
$ git clone https://github.com/spack/spack
$ spack install hdf5
```

Simple syntax enables complex installs

```
$ spack install hdf5@1.10.5                    $ spack install hdf5@1.10.5 cppflags="-O3 –g3"
$ spack install hdf5@1.10.5 %clang@6.0         $ spack install hdf5@1.10.5 target=haswell
$ spack install hdf5@1.10.5 +threadssafe       $ spack install hdf5@1.10.5 +mpi ^mpich@3.2
```

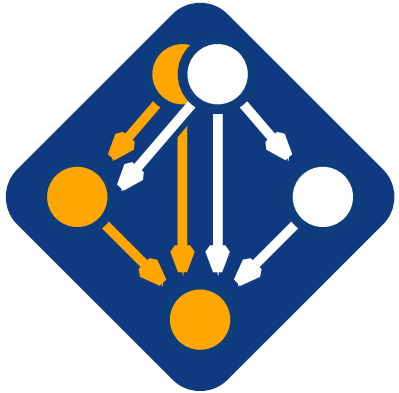

github.com/spack/spack

- Ease of use of mainstream tools, with flexibility needed for HPC
- In addition to CLI, Spack also:
  - Generates (but does **not** require) *modules*
  - Allows conda/virtualenv-like *environments*
  - Provides many devops features (CI, container generation, more)
  - Supports *binary "buildcaches"* so that you don't have to build everything from source

Anyone can use Spack!

- **End Users of HPC Software**
  - Install and run HPC applications and tools
- **HPC Application Teams**
  - Manage third-party dependency libraries
- **Package Developers**
  - People who want to package their own software for distribution
- **User support teams at HPC Centers**
  - People who deploy software for users at large HPC sites

Spack was critical for ECP's mission to create a robust, capable exascale software ecosystem.

EXASCALE COMPUTING PROJECT

https://e4s.io

- Used for building software on the three U.S. exascale systems
- ECP built the Extreme Scale Scientific Software Stack (E4S) with Spack – more at https://e4s.io
- Project continues on ASC and ASCR funding

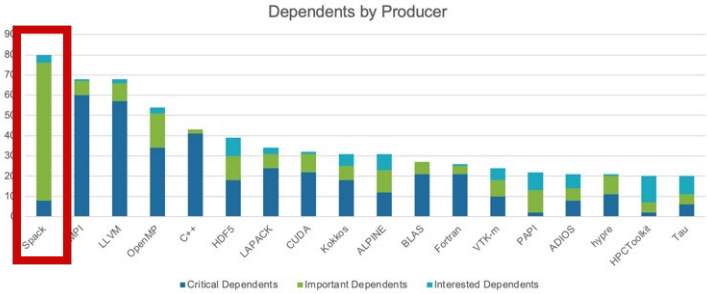

Spack was the most depended-upon project in ECP

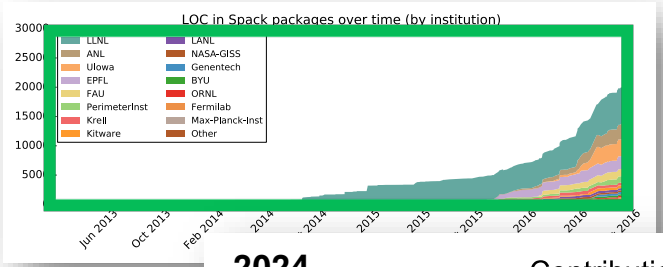

2016

Fall 2020

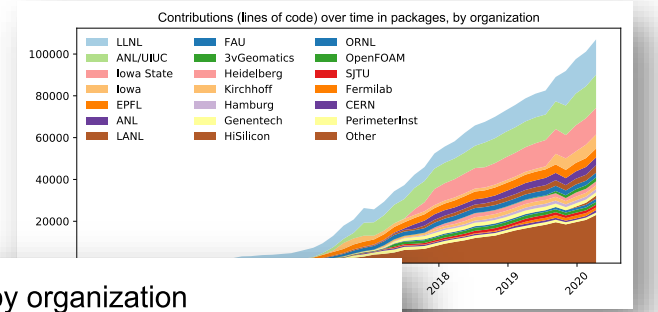

2024

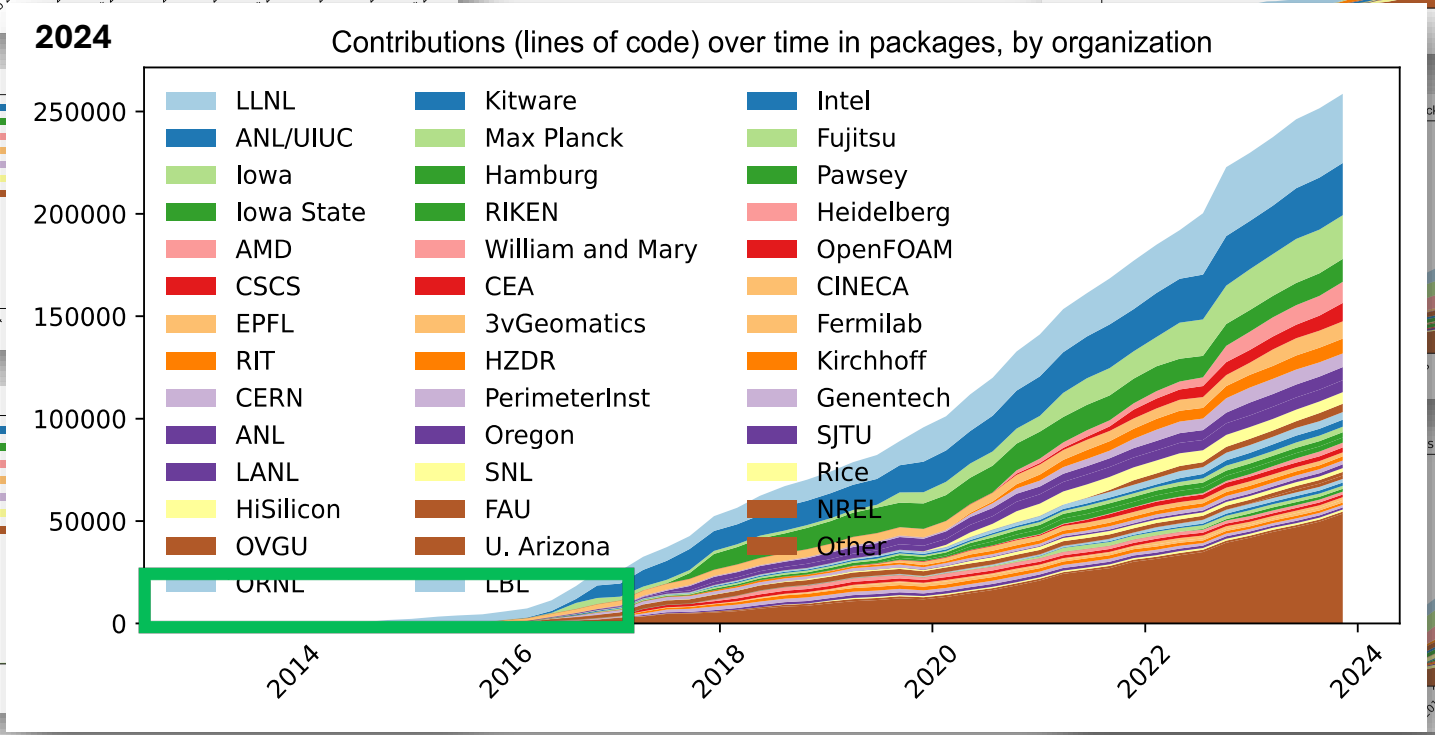

Spack sustains the HPC software ecosystem with the help of many contributors

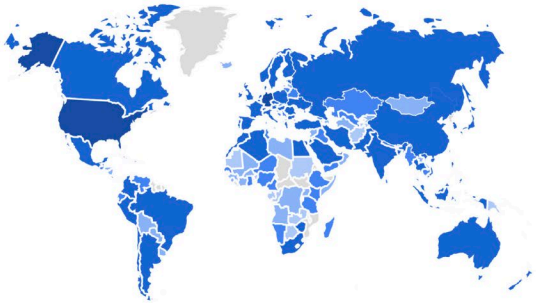

COUNTRY	USERS
United States	23K
Germany	5.3K
China	4.6K
India	4.5K
United Kingdom	3.3K
France	3K
Japan	2.4K

2023 aggregate documentation user counts from GA4
(note: yearly user counts are almost certainly too large)

Over 8,000 software packages
Over 1,300 contributors

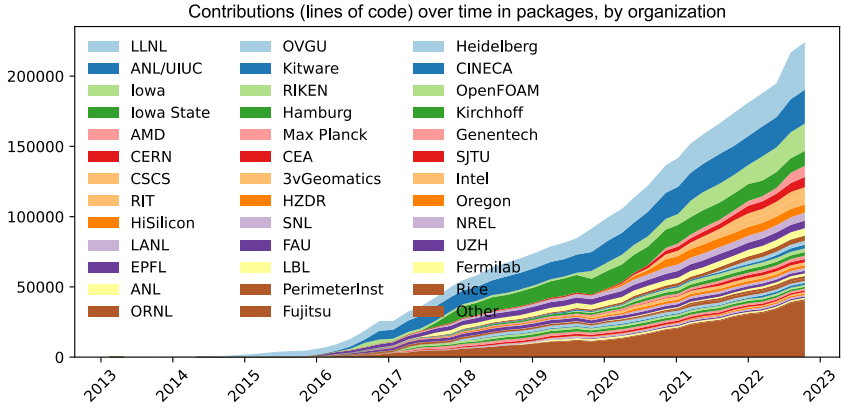

Contributors continue to grow worldwide!

Spack users have diverse roles across many types of institutions

What type of user are you?

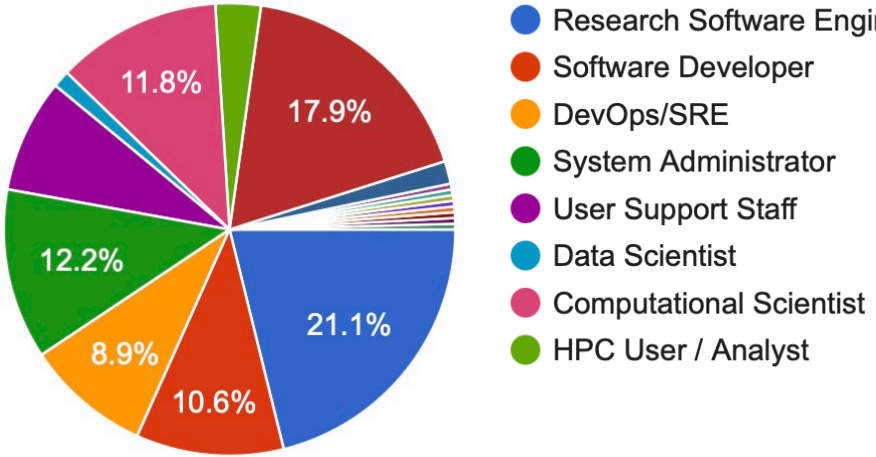

Where do you work?

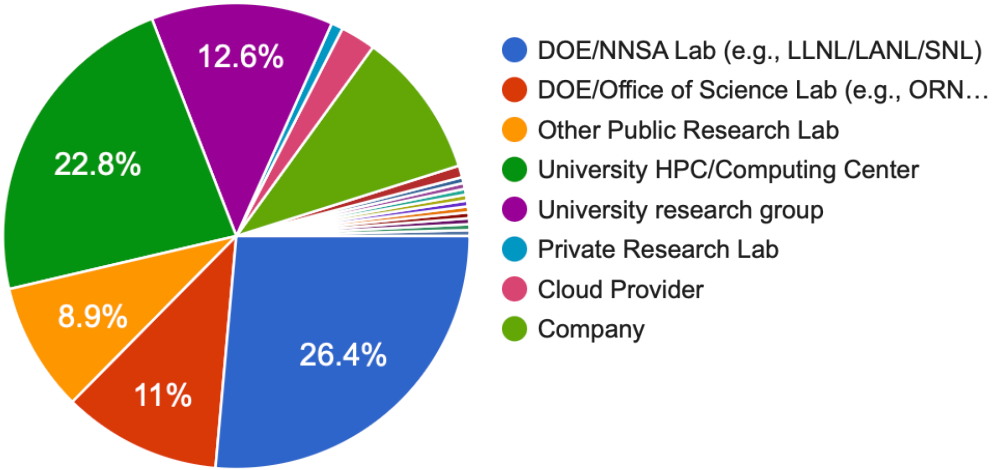

What does the Spack project look like now?

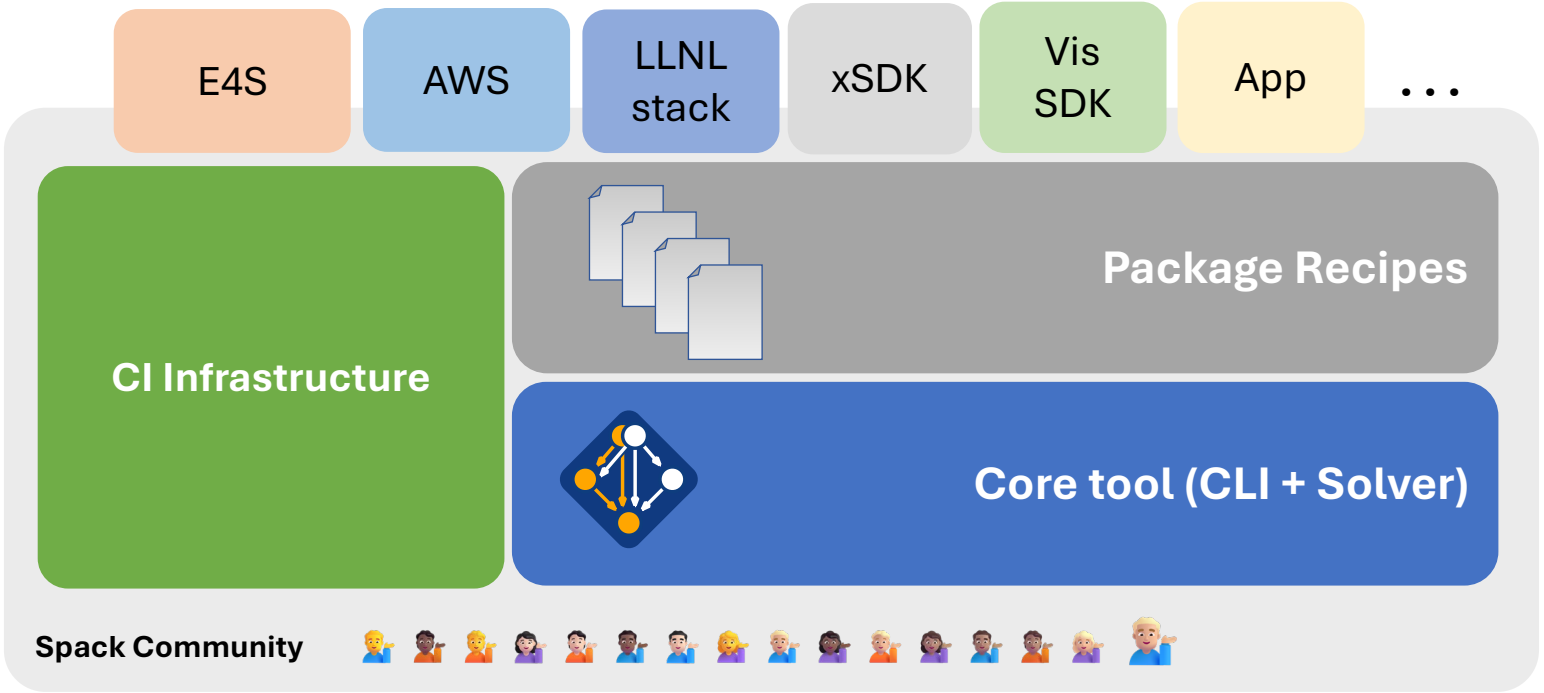

We started conversations with Linux Foundation in December 2021, and talked through mid-2022

- We wanted:
  - A neutral project home
    - To encourage more participation in the project
  - A way to fund project activities:
    - More continuous integration resources
    - User meetings, Slack, etc.
- Talked to LF onboarding team
- Learned about LF's basic requirements:
  - Technical charter
  - Open Governance
  - Trademark assignment

We joined forces with Kokkos to start a larger umbrella, which eventually became HPSF

- Spack and Kokkos were two of the most adopted projects during ECP
- Enable performance portability at different levels
  - Spack: build level
  - Kokkos: application/runtime level
- Goals:
  - Leverage proven track record of community building
  - Leverage industry and labs' familiarity
  - Get more projects on board to build an umbrella organization

With HPSF, we've formalized our governance with the Technical Steering Committee

Todd Gamblin, LLNL
TSC Chair

Greg Becker
LLNL

Massimiliano Culpo
n.p. complete s.r.l

Tammy Dahlgren
LLNL

Wouter Deconinck
U. Manitoba

Ryan Krattiger
Kitware

Mark Krentel
Rice University

John Parent
Kitware

Marc Paterno
Fermilab

Luke Peyralans
U. Oregon

Phil Sakievich
Sandia

Peter Scheibel
LLNL

Adam Stewart
TU Munich

Harmen Stoppels
Stoppels Consulting

Response to LF/HPSF seems positive

Has Spack's transition into Linux Foundation / HPSF given you more confidence in the project?

242 responses

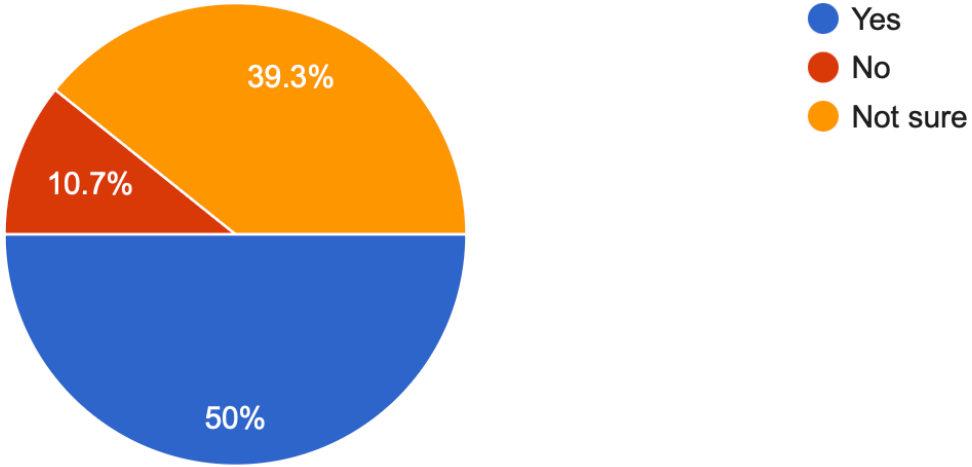

Spack Usage

Spack provides a *spec* syntax to describe customized installations

```
$ 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 cppflags="-O3 –g3"    set compiler flags
$ spack install mpileaks@3.3 target=zen2      set target microarchitecture
$ 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 packages are *templates*
They use a simple Python DSL to define how to build

Not shown: patches, **resources**, **conflicts**, other directives.

```
from spack import *

class Kripke(CMakePackage):
    """Kripke is a simple, scalable, 3D Sn deterministic particle
       transport proxy/mini app.
    """

    homepage = "https://computation.llnl.gov/projects/co-design/kripke"
    url      = "https://computation.llnl.gov/projects/co-design/download/kripke-openmp-1.1.tar.gz"

    version('1.2.3', sha256='3f7f2eef0d1ba5825780d626741eb0b3f026a096048d7ec4794d2a7dfbe2b8a6')
    version('1.2.2', sha256='eaf9ddf562416974157b34d00c3a1c880fc5296fce2aa2efa039a86e0976f3a3')
    version('1.1', sha256='232d74072fc7b848fa2adc8a1bc839ae8fb5f96d50224186601f55554a25f64a')

    variant('mpi',    default=True, description='Build with MPI.')
    variant('openmp', default=True, description='Build with OpenMP enabled.')

    depends_on('mpi', when='+mpi')
    depends_on('cmake@3.0:', type='build')

    def cmake_args(self):
        return [
            '-DENABLE_OPENMP=%s' % ('+openmp' in self.spec),
            '-DENABLE_MPI=%s' % ('+mpi' in self.spec),
        ]

    def install(self, spec, prefix):
        # Kripke does not provide install target, so we have to copy
        # things into place.
        mkdirp(prefix.bin)
        install('../spack-build/kripke', prefix.bin)
```

Base package
(CMake support)

Metadata at the class level

Versions

Variants (build options)

Dependencies
(same spec syntax)

Install logic
in instance methods

Don't typically need install() for CMakePackage, but we can work around codes that don't have it.

Spack Specs can constrain versions of dependencies

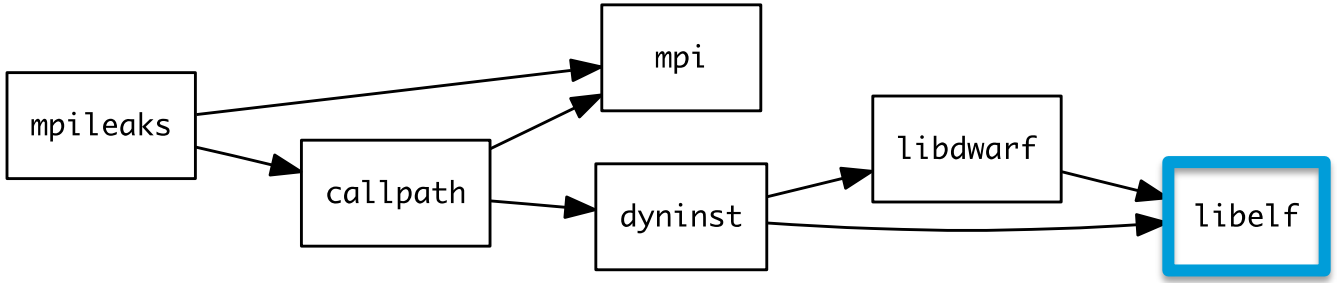


```
$ spack install mpileaks %intel@12.1 ^libelf@0.8.12
```

- Spack ensures *one* configuration of each library per DAG
  - Ensures ABI consistency.
  - User does not need to know DAG structure; only the dependency *names.*
- Spack can ensure that builds use the same compiler, or you can mix
  - Working on ensuring ABI compatibility when compilers are mixed.

Spack handles ABI-incompatible, versioned interfaces like MPI

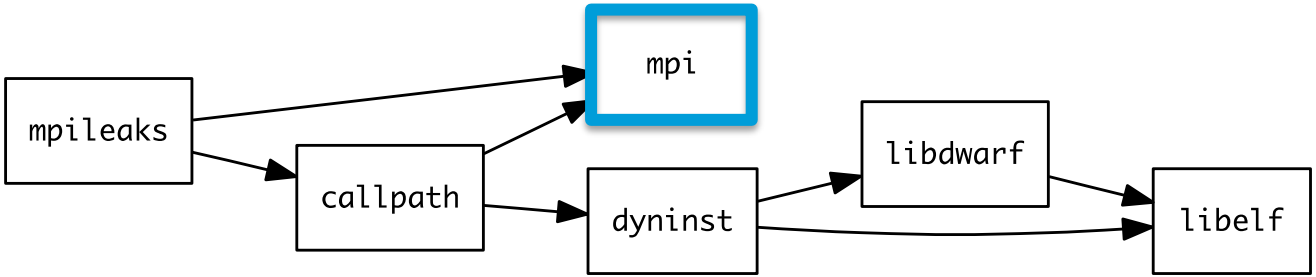

- mpi is a *virtual dependency*
- Install the same package built with two different MPI implementations:

```
$ spack install mpileaks ^mvapich@1.9
```

```
$ spack install mpileaks ^openmpi@1.4:
```

- Let Spack choose MPI implementation, as long as it provides MPI 2 interface:

```
$ spack install mpileaks ^mpi@2
```

Concretization fills in missing configuration details when the user is not explicit.

mpileaks ^callpath@1.0+debug ^libelf@0.8.11

User input: *abstract* spec with some constraints

Normalize

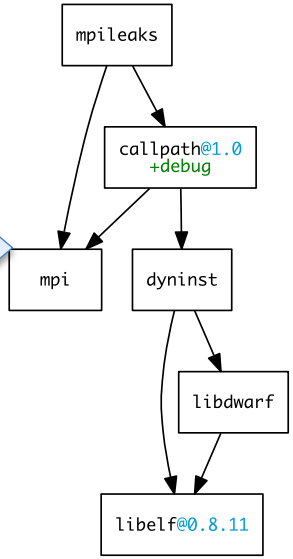

Abstract, normalized spec with some dependencies

Concretize

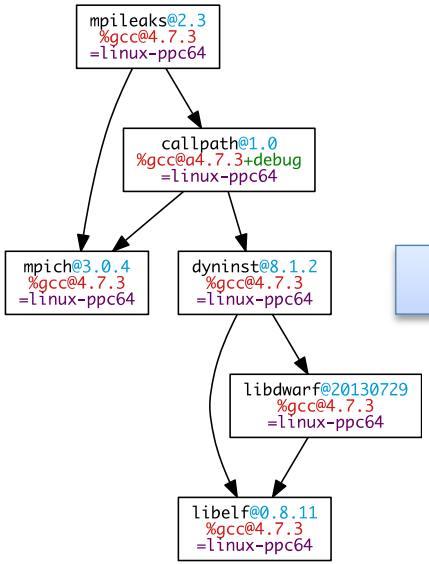

Concrete spec is fully constrained and can be passed to install

Store

spec.json


```
{
  "spec": {
    "_meta": {
      "version": 4
    },
    "nodes": [
      {
        "name": "mpileaks",
        "version": "1.0",
        "arch": {
          "platform": "linux",
          "platform_os": "rhel8",
          "target": {
            "name": "cascadelake"
          }
        },
        "compiler": {
          "name": "gcc",
          "version": "10.3.1"
        },
        "namespace": "builtin",
        "parameters": {
          "build_system": "autotools",
          "stackstart": "0",
        },
        "hash": "o7awlh5q6wccrraon4yd2mfmkdtvvnxe"

        "dependencies": [
          {
            "name": "adept-utils",
            "hash": "r3s7ywbhvtixgc3bknpqovl3dn2adce2",
            "parameters": {
              "deptypes": [
                "build",
                "link"
              ],
              "virtuals": []
            }
          },
          {
```


Detailed provenance stored with installed package

Spack handles combinatorial software complexity

Dependency DAG

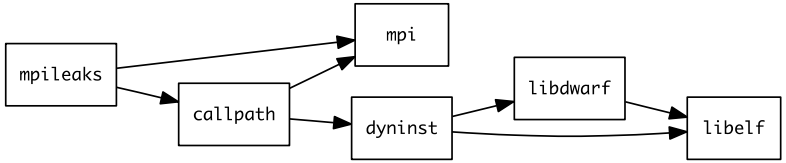

Installation Layout

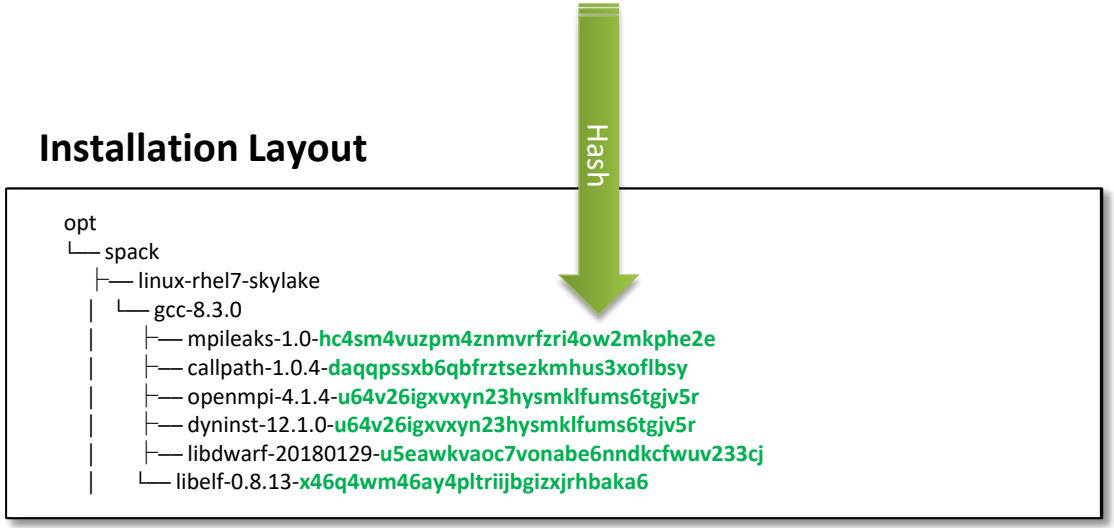

- Each unique dependency graph is a unique ***configuration***.
- Each configuration in a unique directory.
  - Multiple configurations of the same package can coexist.
- **Hash** of entire directed acyclic graph (DAG) is appended to each prefix.
- Installed packages automatically find dependencies
  - Spack embeds RPATHs in binaries.
  - No need to use modules or set LD_LIBRARY_PATH
  - Things work *the way you built them*

Spack environments enable users to build customized stacks from an abstract description

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
```

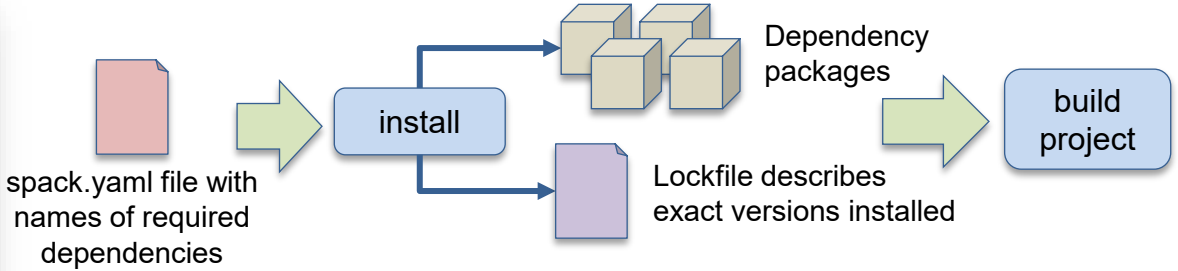

Concrete spack.lock file (generated)

```
{
 "concrete_specs": {
   "6s63so2kstp3zyvjezglndmavy6l3nul": {
     "hdf5": {
         "version": "1.10.5",
         "arch": {
             "platform": "darwin",
             "platform_os": "mojave",
             "target": "x86_64"
         },
         "compiler": {
             "name": "clang",
             "version": "10.0.0-apple"
         },
         "namespace": "builti
         "parameters":
```

- spack.yaml describes project requirements
- spack.lock describes exactly what versions/configurations were installed, allows them to be reproduced.
- Can also be used to maintain configuration together with Spack packages.
  - E.g., versioning your own local software stack with consistent compilers/MPI implementations
  - Allows developers and site support engineers to easily version Spack configurations in a repository

Environments have enabled us to add build many features to support developer workflows

package.py

spack.yaml configuration

spack external find

Automatically find and configure external packages on the system

spack test

Packages know how to run their own test suites

package.py

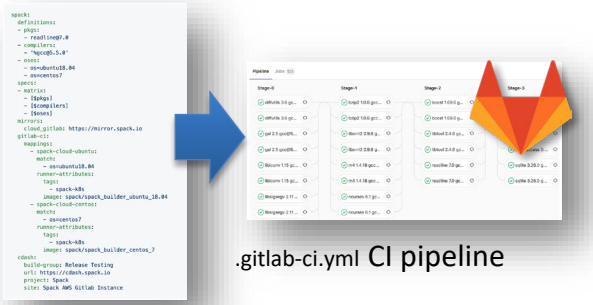

spack.yaml

.gitlab-ci.yml CI pipeline

spack ci

Automatically generate parallel build pipelines
(more on this later)

spack containerize

Turn environments into container build recipes

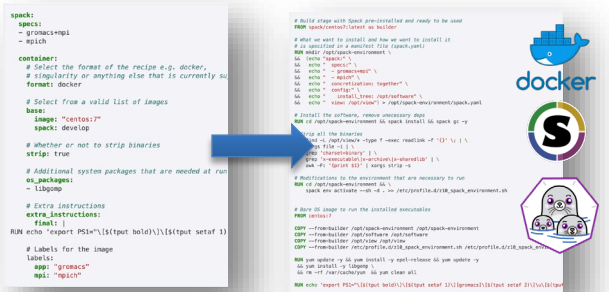

spack develop lets developers work on many packages at once

- Developer features so far have focused on single packages (spack dev-build, etc.)
- New spack develop feature enables development environments
  - Work on a code
  - Develop multiple packages from its dependencies
  - Easily rebuild with changes
- Builds on spack environments
  - Required changes to the installation model for dev packages
  - dev packages don't change paths with configuration changes
  - Allows devs to iterate on builds quickly

```
$ 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
```

Teams are building development front-ends top of on Spack

- Many different approaches:
  - Thin **mack** wrapper for MARBL team at LLNL
  - **spack-manager** at Sandia
  - **SpackDev** at Fermilab
  - **spack-organizer** at CEA
  - **spack_cmake** at LANL

- Workflow seems to be converging around environments + spack develop

- We are trying to bring many of the features from these scripts into core
  - Most of the front-ends are opinionated in one way or another
  - We want spack to support but not force these workflows

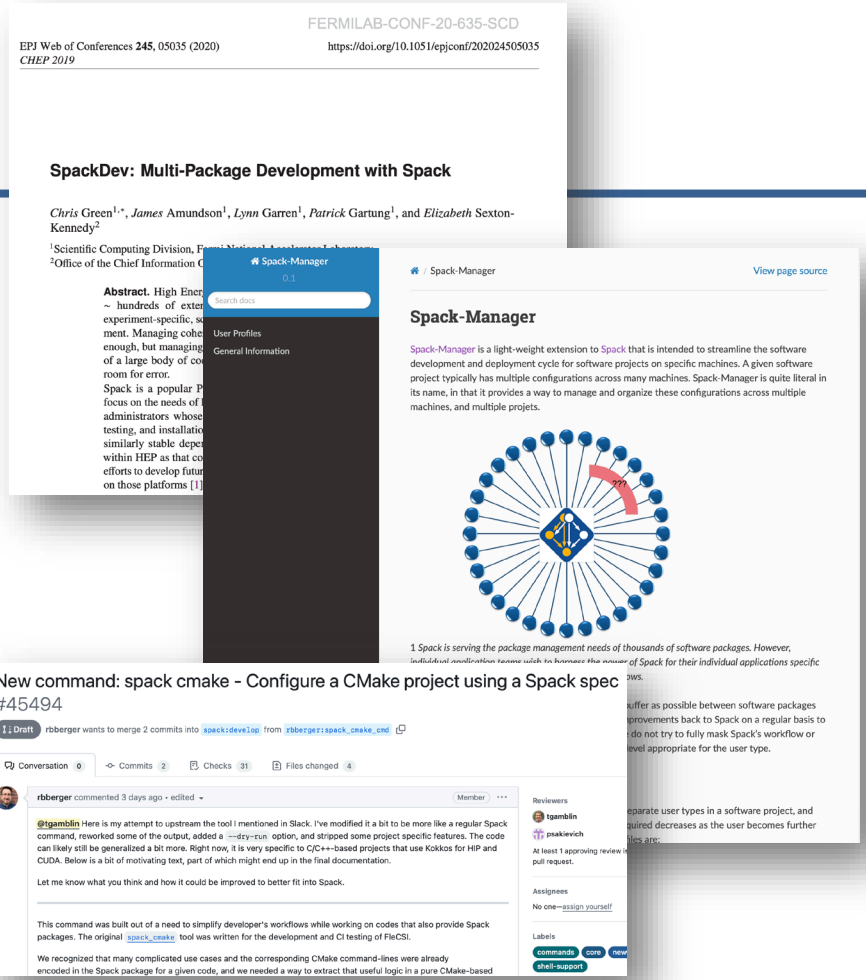

The Road to Spack v1.0

THE LINUX FOUNDATION

The road to v1.0 has been long

- We wanted:
  - ✅ 2020 New ASP-based concretizer
  - ✅ 2021 Reuse of existing installations
  - ✅ 2022 Stable production CI
  - ✅ 2022 Stable binary cache
  - ✅ 2025 Compiler dependencies
  - ✅ 2025 Separate builtin repo
  - ✅ 2025 Stable package API
- v1.0:
  - Changes the dependency model for compilers
  - Enables users to use entirely custom packages
  - Improves reproducibility
  - Improves stability 🤞
- This is the largest change to Spack... ever.

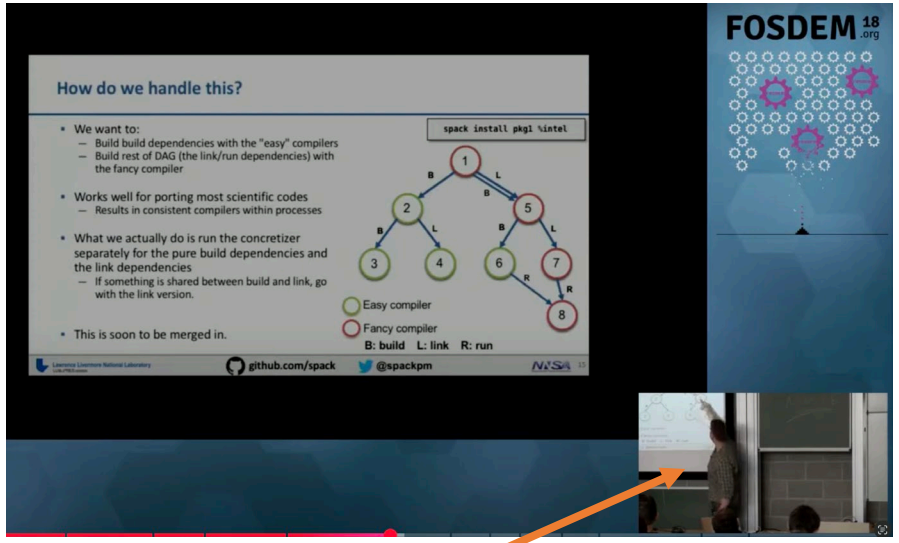

Todd, presenting how simple all this would be at FOSDEM in 2018

Spack packages use a *lot* of (declarative) conditional logic

CudaPackage: a mix-in for packages that use CUDA

```
class CudaPackage(PackageBase):
    variant('cuda', default=False,
            description='Build with CUDA')

    variant('cuda_arch',
            description='CUDA architecture',
            values=any_combination_of(cuda_arch_values),
            when='+cuda')

    depends_on('cuda', when='+cuda')

    depends_on('cuda@9.0:',     when='cuda_arch=70')
    depends_on('cuda@9.0:',     when='cuda_arch=72')
    depends_on('cuda@10.0:',    when='cuda_arch=75')

    conflicts('%gcc@9:', when='+cuda ^cuda@:10.2.89 target=x86_64:')
    conflicts('%gcc@9:', when='+cuda ^cuda@:10.1.243 target=ppc64le:')
```

cuda is a variant (build option)

cuda_arch is only present if cuda is enabled

dependency on cuda, but only if cuda is enabled

constraints on cuda version

compiler support for x86_64 and ppc64le

There is a lot of expressive power in the Spack package DSL.

First challenge: we needed a new concretizer to model the expressiveness of the DSL

Contributors

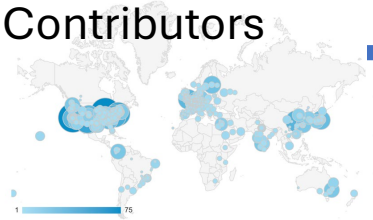

- new versions
- new dependencies
- new constraints

package.py repository

This part is NP-hard!

concretizer

spack developers

yaml default config packages.yaml

admins, users

yaml local preferences config packages.yaml

users

yaml l environment config spack.yaml

users

Command line constraints

spack install hdf5@1.12.0 +debug

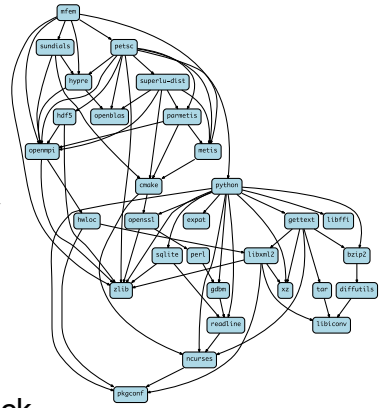

Concrete spec is fully constrained and can be built.

Is stored in spack.lock file after solve.

We reimplemented Spack's concretizer using *Answer Set Programming (ASP)*

- Originally a greedy, custom Python algorithm
- ASP is a *declarative* programming paradigm
  - Looks like Prolog
  - Built around modern CDCL SAT solver techniques
- ASP program has 2 parts:
  1. Large list of facts generated from recipes (problem instance)
  2. Small logic program (~700 lines of ASP code)
- Algorithm is conceptually simpler:
  - Generate facts for all possible dependencies
  - Send facts and our logic program to the solver
  - Read results and rebuild the resolved DAG
- Using **Clingo**, the Potassco grounder/solver package

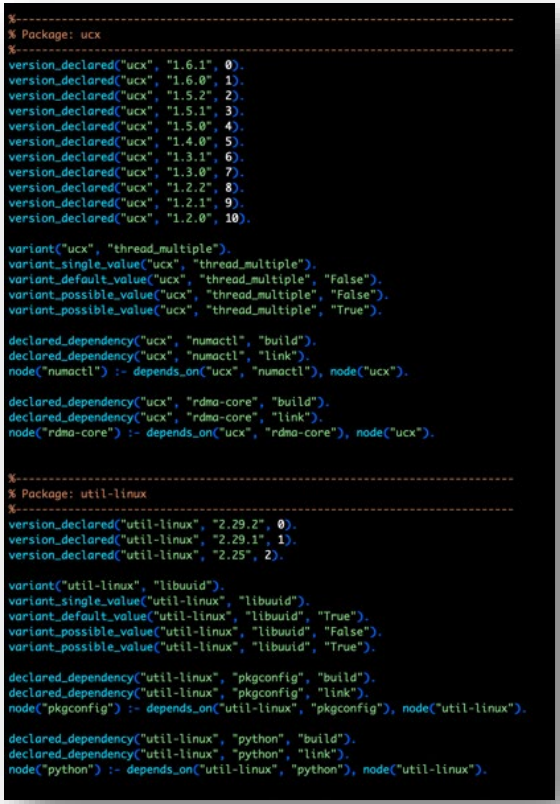

```
%-----------------------------------------------------------------------------
% Package: ucx
%-----------------------------------------------------------------------------
version_declared("ucx", "1.6.1", 0).
version_declared("ucx", "1.6.0", 1).
version_declared("ucx", "1.5.2", 2).
version_declared("ucx", "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 HDF5 package

Spack's concretizer is implemented using Answer Set Programming (ASP)

ASP looks like Prolog but is converted to SAT with optimization

Facts describe the graph

```
node("lammps").
node("cuda").
variant_value("lammps", "cuda", "True").
depends_on("lammps", "cuda").
```

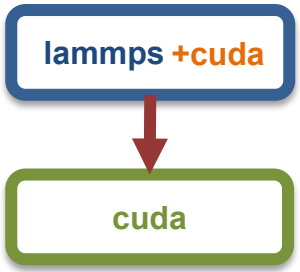

First-order rules (with variables) describe how to resolve nodes and metadata

```
node(Dependency) :- node(Package), depends_on(Package, Dependency).
```

```
node("mpi")
```

```
node("hdf5").
depends_on("hdf5", "mpi").
```

Ground Rule

Grounding converts a first-order logic program into a propositional logic program, which can be solved.

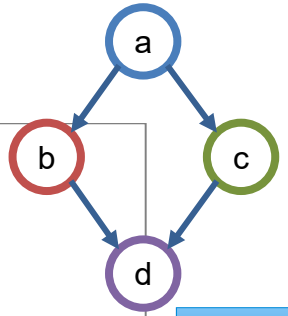


```
depends_on(a, b).
depends_on(a, c).
depends_on(b, d).
depends_on(c, d).

node(Dep)
  :- node(Pkg),
     depends_on(Pkg, Dep).

% at least one is true
1 { node(a); node(b) }.
```

First-order Logic Program

ound

```
depends_on(a, b).
depends_on(a, c).
depends_on(b, d).
depends_on(c, d).

node(b) :- node(a).
node(c) :- node(a).
node(d) :- node(c).
node(d) :- node(b).

% at least one is true
1 { node(a); node(b) }.
```

Propositional Program

```
Answer 1:
node(b)
node(d)

Answer 2:
node(a)
node(b)
node(c)
node(d)
```

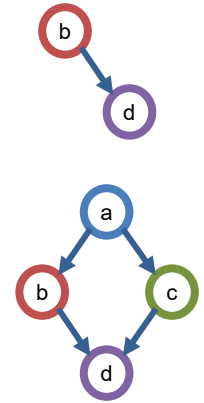

Stable Models (Answer Sets)

Answer 1: Only node(b) is true
Answer 2: Both node(a) and node(b) are true

ASP searches for *stable models* of the input program

- Stable models are also called ***answer sets***
- A ***stable model*** (loosely) is a set of true atoms that can be deduced from the inputs, where every rule is idempotent.
  - Similar to fixpoints
  - Put more simply: a *set of atoms where all your rules are true!*
- Unlike Prolog:
  - Stable models contain everything that can be derived (vs. just querying values)
  - Good ways to do optimization to select the “best” stable model
  - ASP is guaranteed to complete!

Second challenge: Spack's original concretizer did not reuse existing installations

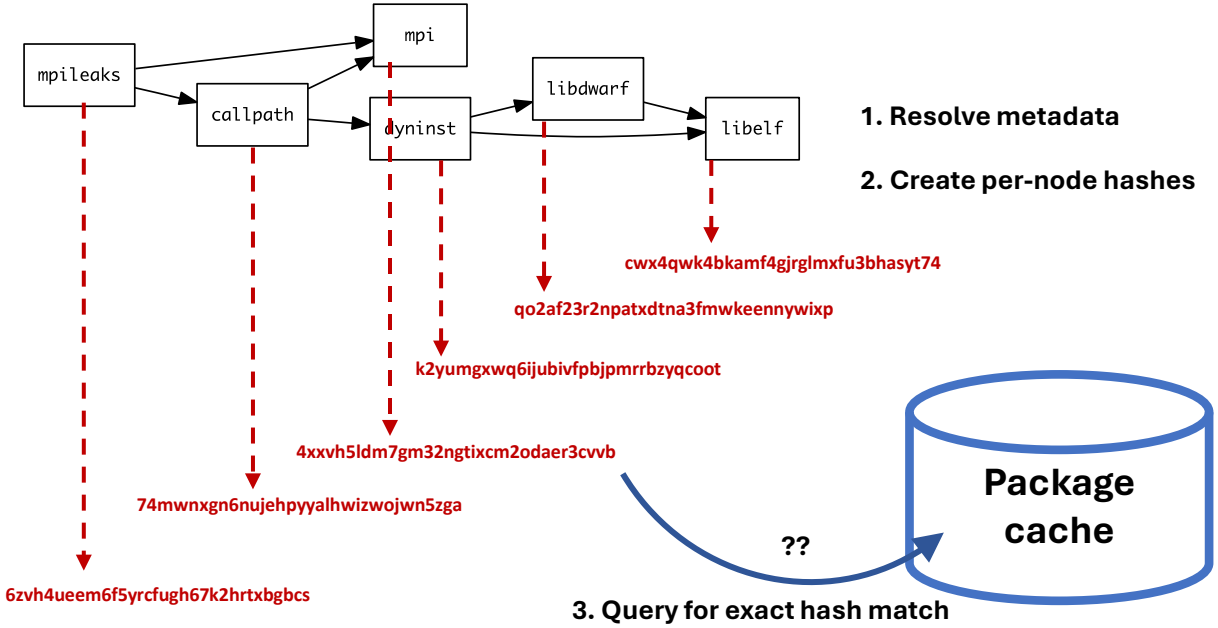

- Hash matches are very sensitive to small changes
- In many cases, a satisfying cached or already installed spec can be missed
- Nix, Spack, Guix, Conan, and others reuse this way

--reuse (now the default) was enabled by ASP

- --reuse tells the solver about all the installed packages!
- Add constraints for all installed packages, with their hash as the associated ID:

```
installed_hash("openssl","lwatuuysmwkhuahrncywvn77icdhs6mn").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","node","openssl").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","version","openssl","1.1.1g").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","node_platform_set","openssl","darwin").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","node_os_set","openssl","catalina").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","node_target_set","openssl","x86_64").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","variant_set","openssl","systemcerts","True").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","node_compiler_set","openssl","apple-clang").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","node_compiler_version_set","openssl","apple-clang","12.0.0").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","concrete","openssl").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","depends_on","openssl","zlib","build").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","depends_on","openssl","zlib","link").
imposed_constraint("lwatuuysmwkhuahrncywvn77icdhs6mn","hash","zlib","x2anksgssxsxa7pcnhzg5k3dhgacglze").
```

Minimizing builds is surprisingly simple in ASP

1. Allow the solver to *choose* a hash for any package:

```
{ hash(Package, Hash) : installed_hash(Package, Hash) } 1 :- node(Package).
```

2. Choosing a hash means we impose its constraints:

```
impose(Hash) :- hash(Package, Hash).
```

3. Define a build as something *without* a hash:

```
build(Package) :- not hash(Package, _), node(Package).
```

4. Minimize builds!

```
#minimize { 1@100,Package : build(Package) }.
```

There's more to it than this, but you get the idea…

With and without --reuse optimization

```
(spackle):solver> spack solve -Il hdf5
==> Best of 9 considered solutions.
==> Optimization Criteria:
  Priority  Criterion                                                  Installed  ToBuild
  1         number of packages to build (vs. reuse)                            -       20
  2         deprecated versions used                                           0        0
  3         version weight                                                     0        0
  4         number of non-default variants (roots)                             0        0
  5         preferred providers for roots                                      0        0
  6         default values of variants not being used (roots)                  0        0
  7         number of non-default variants (non-roots)                         0        0
  8         preferred providers (non-roots)                                    0        0
  9         compiler mismatches                                                0        0
  10        OS mismatches                                                      0        0
  11        non-preferred OS's                                                 0        0
  12        version badness                                                    0        2
  13        default values of variants not being used (non-roots)              0        0
  14        non-preferred compilers                                            0        0
  15        target mismatches                                                  0        0
  16        non-preferred targets                                              0        0

  -   zzngfs3  hdf5@1.10.7%apple-clang@13.0.0~cxx~fortran~hl~ipo~java+mpi+shared~szip~threadsafe+tools api=default b
  -   nsylovq      ^cmake@3.21.4%apple-clang@13.0.0~doc+ncurses+openssl+ownlibs~qt build_type=Release arch=darwin-bi
  -   xdbaqeo          ^ncurses@6.2%apple-clang@13.0.0~symlinks+termlib abi=none arch=darwin-bigsur-skylake
  -   kfureok              ^pkgconf@1.8.0%apple-clang@13.0.0 arch=darwin-bigsur-skylake
  -   5ekd4ap          ^openssl@1.1.1l%apple-clang@13.0.0~docs certs=system arch=darwin-bigsur-skylake
  -   xz6a265              ^perl@5.34.0%apple-clang@13.0.0+cpanm+shared+threads arch=darwin-bigsur-skylake
  -   xgt3tls                  ^berkeley-db@18.1.40%apple-clang@13.0.0+cxx~docs+stl patches=b231fcc4d5cff05e5c3a4814
  -   65edjf6                  ^bzip2@1.0.8%apple-clang@13.0.0~debug~pic+shared arch=darwin-bigsur-skylake
  -   662adoo                      ^diffutils@3.8%apple-clang@13.0.0 arch=darwin-bigsur-skylake
  -   fu7tfsr                          ^libiconv@1.16%apple-clang@13.0.0 libs=shared,static arch=darwin-bigsur-skyla
  -   vjg67nd                  ^gdbm@1.19%apple-clang@13.0.0 arch=darwin-bigsur-skylake
  -   tjceldr                      ^readline@8.1%apple-clang@13.0.0 arch=darwin-bigsur-skylake
  -   xevvljj                  ^zlib@1.2.11%apple-clang@13.0.0+optimize+pic+shared arch=darwin-bigsur-skylake
  -   xelfobh      ^openmpi@4.1.1%apple-clang@13.0.0~atomics~cuda~cxx~cxx_exceptions+gpfs~internal-hwloc~java~legacy
  -   zruns75          ^hwloc@2.6.0%apple-clang@13.0.0~cairo~cuda~gl~libudev+libxml2~netloc~nvml~opencl~pci~rocm+sha
  -   ib4fnkf              ^libxml2@2.9.12%apple-clang@13.0.0~python arch=darwin-bigsur-skylake
  -   dwiv2ys                  ^xz@5.2.5%apple-clang@13.0.0~pic libs=shared,static arch=darwin-bigsur-skylake
  -   blitnbl          ^libevent@2.1.12%apple-clang@13.0.0+openssl arch=darwin-bigsur-skylake
  -   h7jalyu          ^openssh@8.7p1%apple-clang@13.0.0 arch=darwin-bigsur-skylake
  -   7v7bqx2              ^libedit@3.1-20210216%apple-clang@13.0.0 arch=darwin-bigsur-skylake
```

Pure hash-based reuse: all misses

```
(spackle):spack> spack solve --reuse -Il hdf5
==> Best of 10 considered solutions.
==> Optimization Criteria:
  Priority  Criterion                                                  Installed  ToBuild
  1         number of packages to build (vs. reuse)                            -        4
  2         deprecated versions used                                           0        0
  3         version weight                                                     0        0
  4         number of non-default variants (roots)                             0        0
  5         preferred providers for roots                                      0        0
  6         default values of variants not being used (roots)                  0        0
  7         number of non-default variants (non-roots)                         2        0
  8         preferred providers (non-roots)                                    0        0
  9         compiler mismatches                                                0        0
  10        OS mismatches                                                      0        0
  11        non-preferred OS's                                                 0        0
  12        version badness                                                    6        0
  13        default values of variants not being used (non-roots)              1        0
  14        non-preferred compilers                                           15        4
  15        target mismatches                                                  0        0
  16        non-preferred targets                                              0        0

  -   yfkfnsp  hdf5@1.10.7%apple-clang@12.0.5~cxx~fortran~hl~ipo~java+mpi+shared~szip~threadsafe+tools api=defaul
  [+] zd4m26e      ^cmake@3.21.1%apple-clang@12.0.5~doc+ncurses+openssl+ownlibs~qt build_type=Release arch=darwin
  [+] 53i52xr          ^ncurses@6.2%apple-clang@12.0.5~symlinks+termlib abi=none arch=darwin-bigsur-skylake
  [+] us36bwr          ^openssl@1.1.1l%apple-clang@12.0.5~docs+systemcerts arch=darwin-bigsur-skylake
  [+] 74mwnxg              ^zlib@1.2.11%apple-clang@12.0.5+optimize+pic+shared arch=darwin-bigsur-skylake
  -   3ijfnel      ^openmpi@4.1.1%apple-clang@12.0.5~atomics~cuda~cxx~cxx_exceptions+gpfs~internal-hwloc~java~leg
  -   jxexyb7          ^hwloc@2.6.0%apple-clang@12.0.5~cairo~cuda~gl~libudev+libxml2~netloc~nvml~opencl~pci~rocm+
  [+] ckdn5zf              ^libxml2@2.9.12%apple-clang@12.0.5~python arch=darwin-bigsur-skylake
  [+] k7auat3                  ^libiconv@1.16%apple-clang@12.0.5 libs=shared,static arch=darwin-bigsur-skylake
  [+] kzyumgx                  ^xz@5.2.5%apple-clang@12.0.5~pic libs=shared,static arch=darwin-bigsur-skylake
  [+] grgtlcd              ^pkgconf@1.8.0%apple-clang@12.0.5 arch=darwin-bigsur-skylake
  -   nnc66ug          ^libevent@2.1.12%apple-clang@12.0.5+openssl arch=darwin-bigsur-skylake
  [+] 63xbksk          ^openssh@8.6p1%apple-clang@12.0.5 arch=darwin-bigsur-skylake
  [+] snhgldt              ^libedit@3.1-20210216%apple-clang@12.0.5 arch=darwin-bigsur-skylake
  [+] qbkmtdd          ^perl@5.34.0%apple-clang@12.0.5+cpanm+shared+threads arch=darwin-bigsur-skylake
  [+] tnvkifs              ^berkeley-db@18.1.40%apple-clang@12.0.5+cxx~docs+stl patches=b231fcc4d5cff05e5c3a4814f
  [+] 7d5woqt              ^bzip2@1.0.8%apple-clang@12.0.5~debug~pic+shared arch=darwin-bigsur-skylake
  [+] vh6di3i              ^gdbm@1.19%apple-clang@12.0.5 arch=darwin-bigsur-skylake
  [+] qgy3v4l                  ^readline@8.1%apple-clang@12.0.5 arch=darwin-bigsur-skylake
```

With reuse: 16 packages were reusable

Third challenge: we needed to allow multiple versions of build dependencies in the DAG

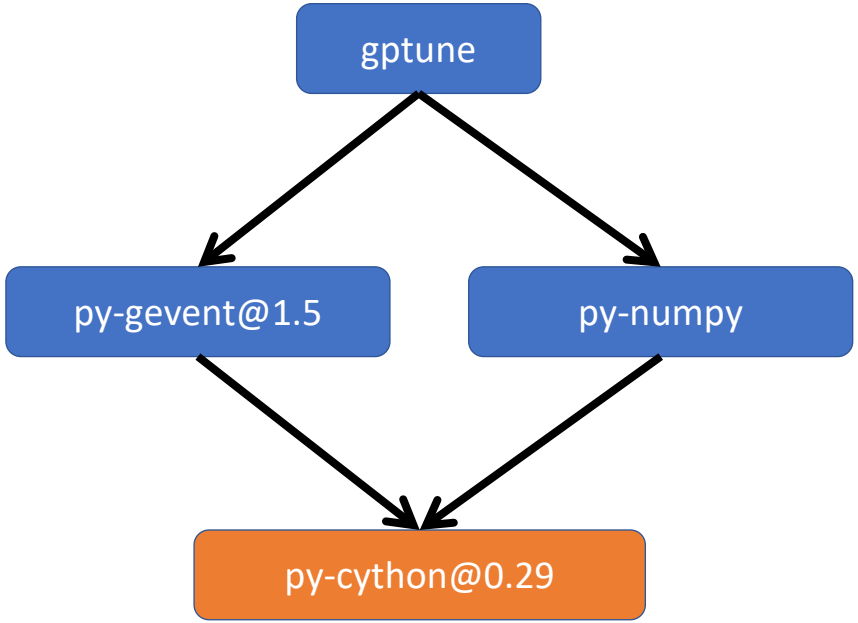

- Only one configuration per package allowed in the DAG
  - Ensures ABI compatibility but is too restrictive
- Needed to enable compiler mixing with compiler dependencies
- Also needed for Python ecosystem
  - In the example py-numpy needs to use py-cython@0.29 as a build tool
  - That enforces using an old py-gevent, because newer versions depend on py-cython@3.0 or greater

Objective: dependency splitting

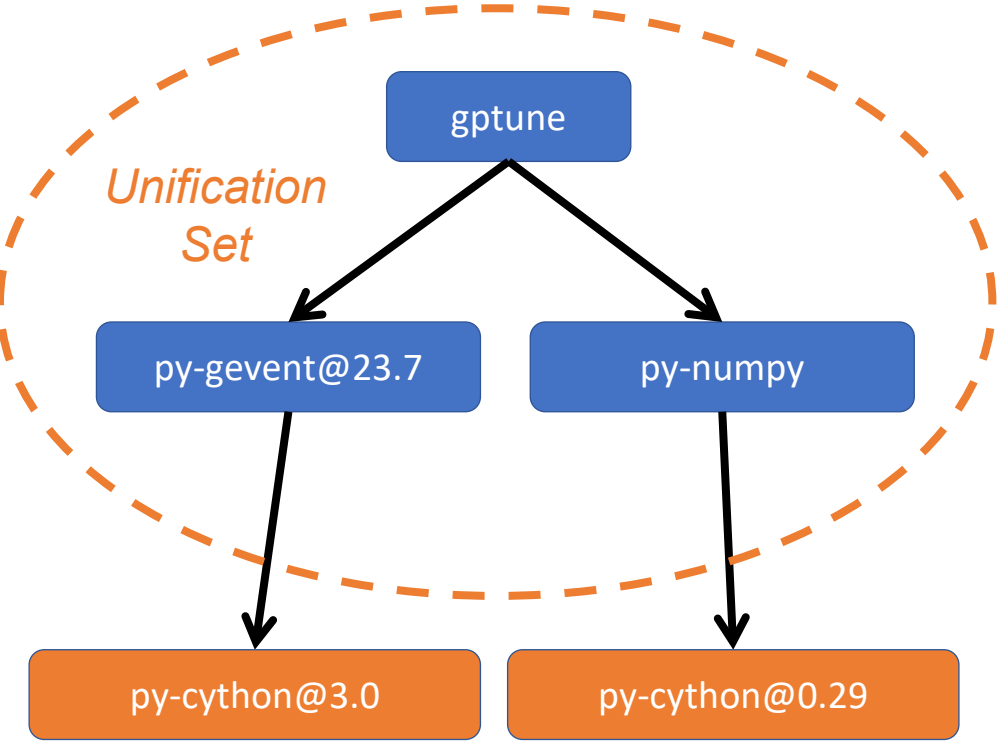

- The constraint on build dependencies can be relaxed, without compromising the ABI compatibility
- Having a single configuration of a package is now enforced on unification sets
- These are the set of nodes used together at runtime (the one shown is for gptune)
- This allows us to use the latest version of py-gevent, because now we can have two versions of py-cython

We want to dynamically “split” nodes when needed

1. Start with deducing single dependency nodes:

```
node(DependencyName)
   :- dependency_holds(PkgName, DependencyName)
```

2. Allow solver to **choose** to duplicate a node:

Converted node identifier from **name** to (**name, id)**

```
1 {
   depends_on(PkgNode, node(0..Y-1, DepNode), Type)
   : max_dupes(DepNode, Y)
} 1
   :- dependency_holds(PkgNode, DepNode).
```

3. Re-encode package metadata so that it can be associated with duplicates

First try at allowing duplicates in a single solve

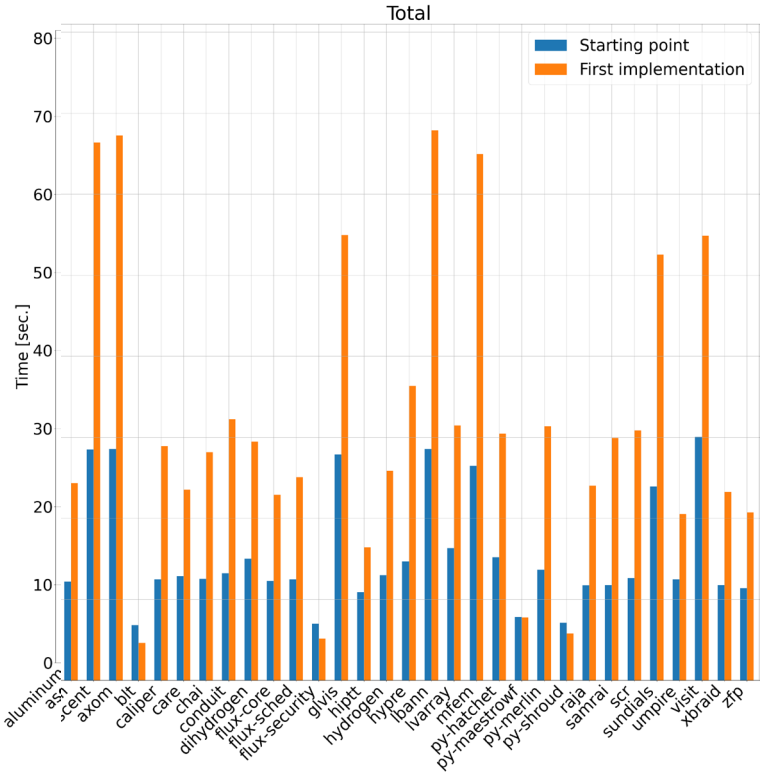

Increased runtimes by >> 2x in some cases

Cycle detection in the solver is *expensive*

```
path(A, B) :- depends_on(A, B).
path(A, C) :- path(A, B), depends_on(B, C).

% this constraint says "no cycles"
:- path(A, B), path(B, A).
```

- Has to maintain path() predicate representing paths between nodes
- Cycles are actually rare in solutions
  - Switched to post-processing for cycle detection
  - Only do expensive solve if a cycle is detected in a solution
- Eventually moved this calculation *into* the solver using some custom directives from the developers

50%+ improvement in solve time

Unification sets can be expensive too

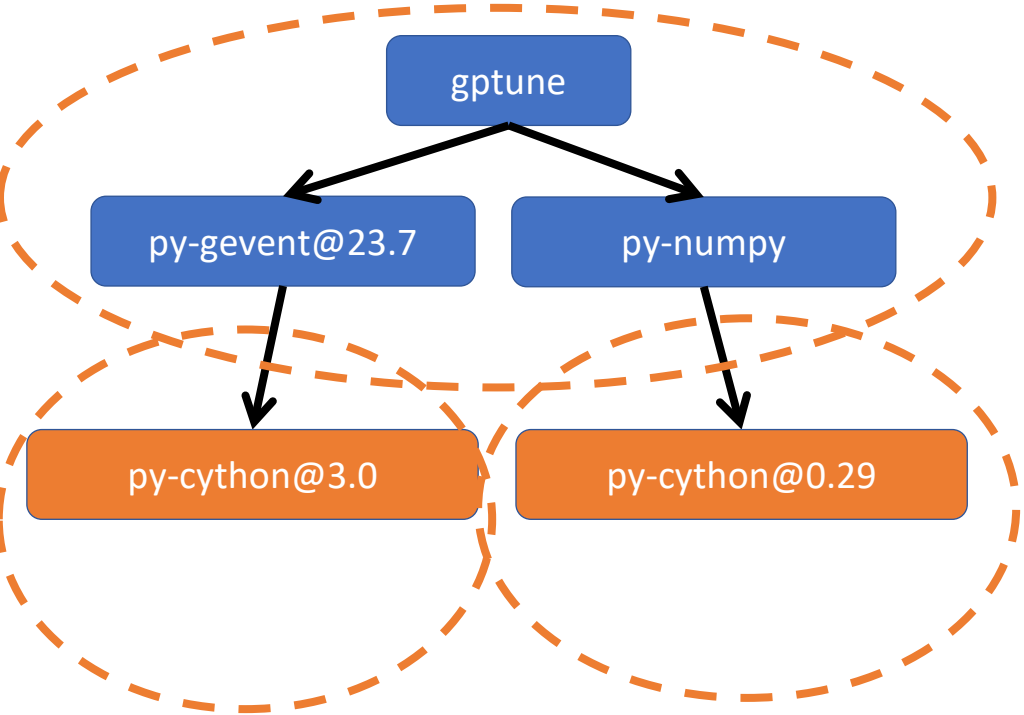

- Unification set creation was originally recursive for *any* build dependencies
  - Ends up blowing up grounding
- Mitigation:
  - Only create new sets for explicitly *marked* build tools
  - Transitive build dependencies that are not from marked build tools go into a *common* unification set
- Need better heuristics to split when necessary for full generality

Through many different optimizations, we were able to reclaim enough performance to make duplicate build dependencies tractable

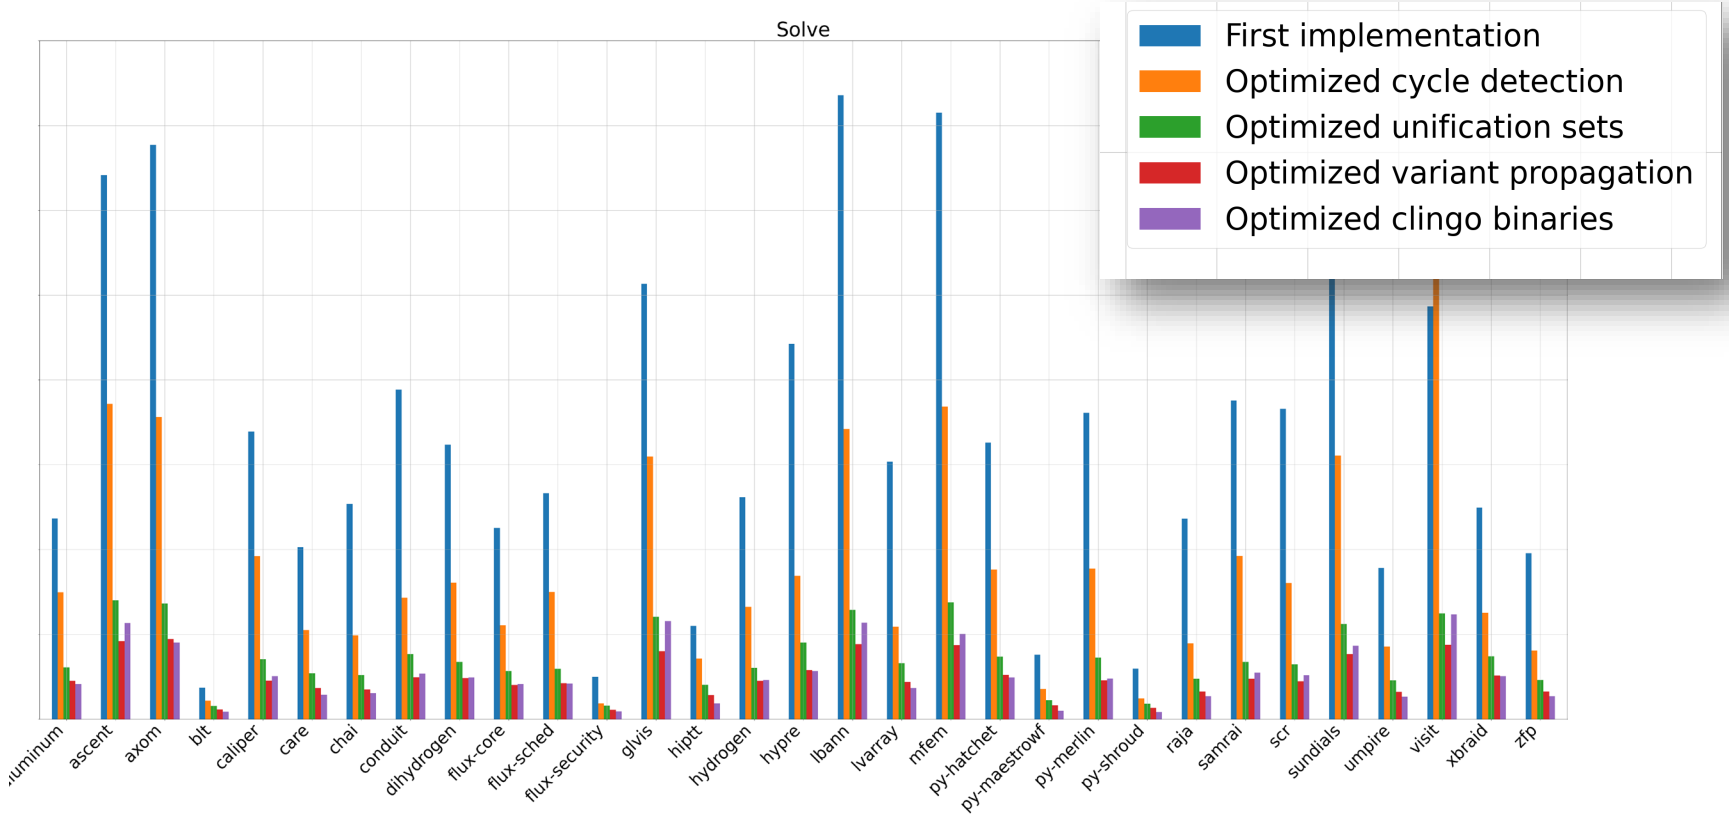

It was not trivial to find a model that was both performant *and* tightly coupled

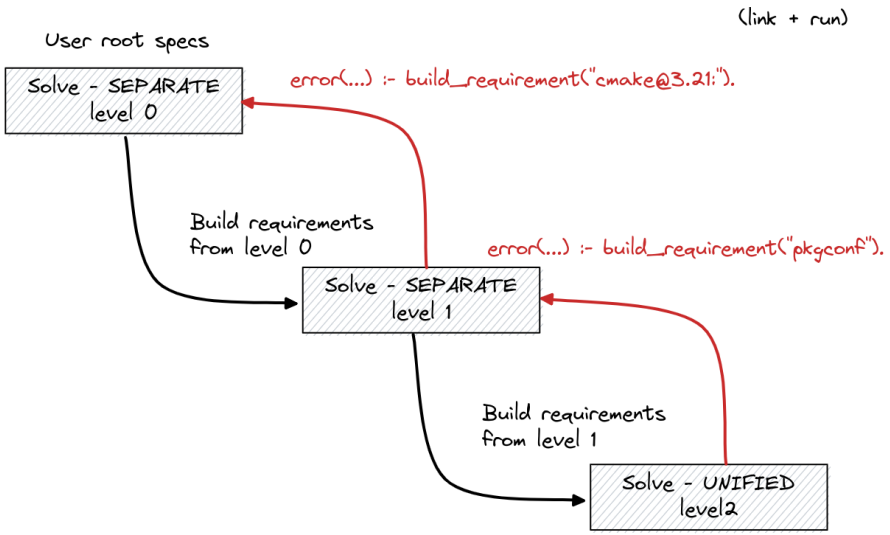

- We tried an iterative version with multiple solves
- Multiple solves had some disadvantages:
  - Slower due to overhead of multiple solves
  - Not coupled, so feedback from build to run environment (and back) was awkward
  - Packagers needed to “help” the solver
- Requiring packagers to provide solve hints in packages isn’t practical

Fourth Challenge: v1.0 adds language dependencies

```
depends_on("c", type="build")
depends_on("cxx", type="build")
depends_on("fortran", type="build")
```

- Spack has historically made these compilers available to every package
  - A compiler was actually “something that supports c + cxx + fortran + f77”
  - Made for a lot of special cases
  - Also makes for duplication of purely interpreted packages (e.g. python)
- Required in 1.0 if you want to use c, cxx, or fortran
  - No-op in v0.23 and prior as we prepared for this feature

Compiler Dependencies

- Compilers are now build dependencies
- Runtime libraries modeled as packages
  - gcc-runtime is injected as link dependency by gcc
  - packages depend on c, cxx,
    fortran virtuals, which are satisfied by gcc node
- glibc is an automatically detected external
  - Injected as a `libc` virtual dependency
  - Does not require user configuration
- Will eventually be able to choose implementations (e.g., musl)

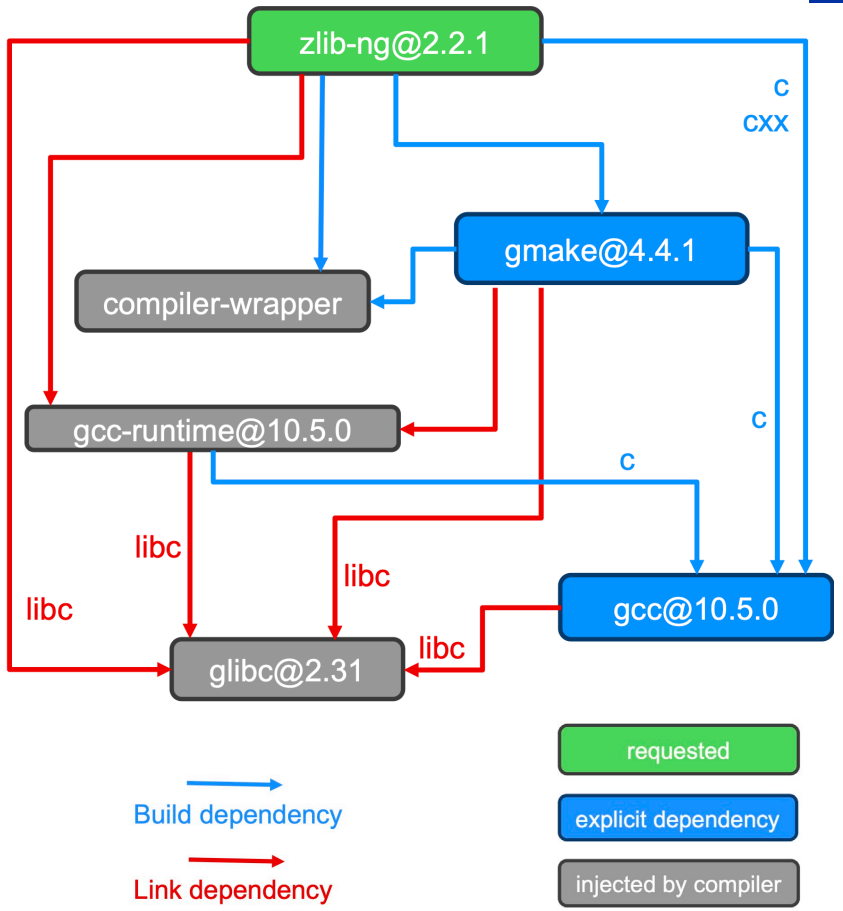

Spack 1.x introduces *toolchains*

toolchains.yaml

```yaml
toolchains:
  clang_gfortran:
    - spec: %c=clang
      when: %c
    - spec: %cxx=clang
      when: %cxx
    - spec: %fortran=gcc
      when: %fortran
    - spec: cflags="-O3 -g"
    - spec: cxxflags="-O3 -g"
    - spec: fflags="-O3 -g"
```

```
spack install foo %clang_gfortran
```

```yaml
toolchains:
  intel_mvapich2:
    - spec: %c=intel-oneapi-compilers @2025.1.1
      when: %c
    - spec: %cxx=intel-oneapi-compilers @2025.1.1
      when: %cxx
    - spec: %fortran=intel-oneapi-compilers @2025.1.1
      when: %fortran
    - spec: %mpi=mvapich2 @2.3.7-1 +cuda
      when: %mpi
```

```
spack install foo %intel_mvapich2
```

- Can lump many dependencies, flags together and use them with a single name
- Any spec in a toolchain can be *conditional*
  - Only apply when needed

Configuring compilers in Spack v1.*

Spack v0.x

compilers.yaml

```yaml
compilers:
  - compiler:
      spec: gcc@12.3.1
      paths:
         c: /usr/bin/gcc
         cxx: /usr/bin/g++
         fc: /usr/bin/gfortran
      modules: [...]
```

Spack v1.x

packages.yaml

```yaml
packages:
  gcc:
    externals:
    - spec: gcc@12.3.1+binutils
      prefix: /usr
      extra_attributes:
        compilers:
          c: /usr/bin/gcc
          cxx: /usr/bin/g++
          fc: /usr/bin/gfortran
        modules: [...]
```

- We automatically convert compilers.yaml, when no compiler is configured
- We will still support *reading* the old configuration until at *least* v1.1
- All fields from `compilers.yaml` are supported in `extra_attributes`

Final challenge: Splitting the package repository

- Spack is two things:
  - Command line tool `spack`
  - Package repository with 8,500+ recipes
- Community wanted
  - package updates without tool changes (e.g. new bugs)
  - tool updates without package changes (reproducibility)
- But coupling between tool and packages was tight
  1. Package classes are in core: `CMakePackage`, `AutotoolsBuilder`, etc.
  2. Compiler wrapper was not a package until recently
  3. Packages *live* in Spack's GitHub repository with a long (git) history

Spack now has a Stable Package API

- Repositories define API version used
  - Versioned *per commit*
- Spack defines API version(s) supported
  - Will complain if a repo is too new
- Packages can only import from:
  - spack.package
  - Core Python
- Any 1.x Spack will support the same package API as all prior 1.x versions
  - Won't break package API unless we bump the major version

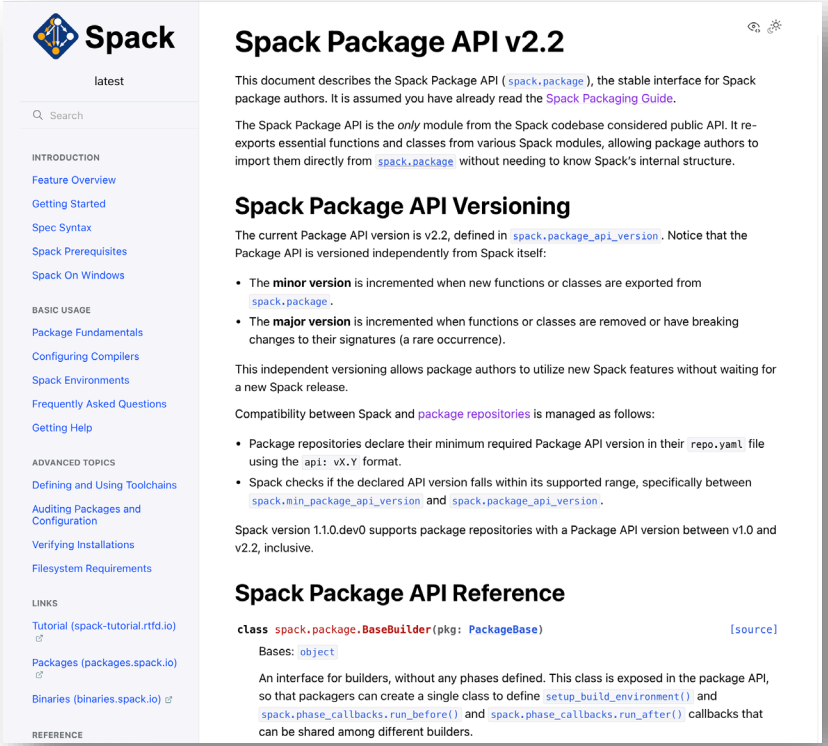

https://spack.rtfd.io/en/latest/package_api.html

Package split process

- Sync packages to **`spack/spack-packages`**
  - Git history is preserved 😌
- Turn package repositories into Python namespace packages
  - **`spack.pkg.builtin`** is now **`spack_repo.builtin`**
- Move build systems to **`spack_repo.builtin.build_systems`**
- Update packages to use fewer Spack internals
- Enable CI on **`spack/spack-packages`**
- Make Spack support Git-based package repositories

You can now specify the package repo version in an environment or config

Pin a commit

```
spack:
  repos:
    builtin:
      git: https://github.com/spack/spack-packages.git
      commit: aec1e3051c0e9fc7ef8feadf766435d6f8921490
```

Work on a branch

```
spack:
  repos:
    builtin:
      git: https://github.com/spack/spack-packages.git
      destination: /path/to/clone/of/spack-packages
      branch: develop
```

Useful commands after repo split

1. `spack repo migrate`: fixes imports in custom repos for you
2. `spack repo set --destination ~/spack-pkgs builtin`: put packages in your favorite location
3. (`spack repo update`: update & pin package repos 

™)

New docs: **https://spack.readthedocs.io/en/latest/repositories.html**

Bonus feature:
Spack now supports concurrent builds!

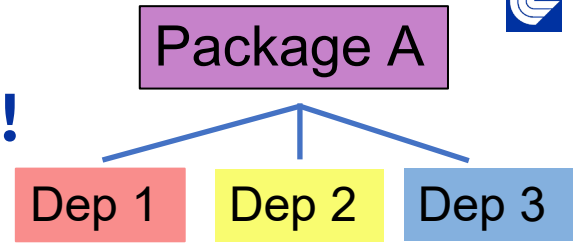

- We *sort of* supported this already
  - But the user had to launch multiple spack processes
  - e.g., `srun -N 4 -n 16 spack install hdf5`

- Now spack handles on-node parallelism itself!
  - Spack now has a scheduler loop
  - Monitors dependencies, starts multiple processes, polls for completion
  - User can control max concurrent processes with '-p'

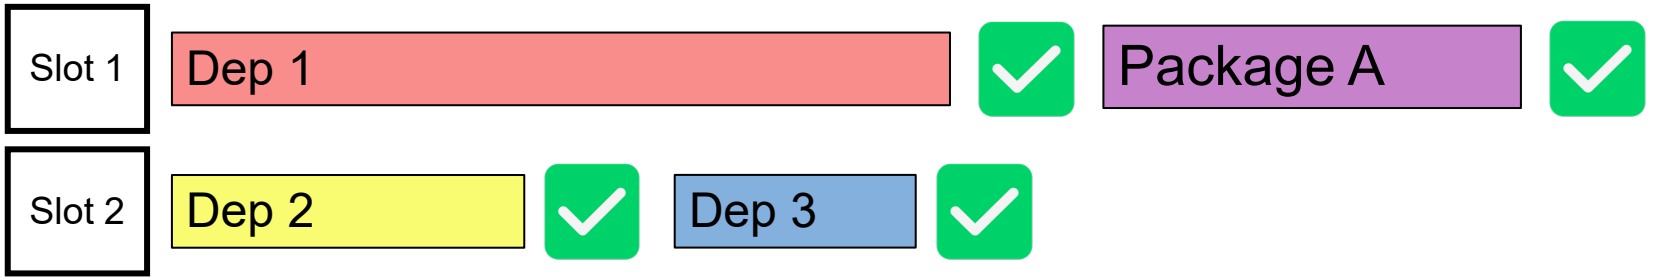

But wait! There's more!

Spack is a core project in the
High Performance Software Foundation

Join us at the Spack User Meeting at
HPSFCon 2026 next year!

@hpsf.bsky.social

hpsf.io

Join us after ISC!

- Join us and 3,800+ others on Spack slack
- Contribute packages, docs, and features on GitHub
- Continue the tutorial at spack-tutorial.rtfd.io

slack.spack.io

★ Star us on GitHub!
github.com/spack/spack

@spackpm.bsky.social

@spack@hpc.social

@spackpm

spack.io

We hope to make distributing & using HPC software easy!

Hands-on time!

THE LINUX FOUNDATION

Tutorial Materials

Find these slides and associated scripts here:

spack-tutorial.rtfd.io

We also have a `#tutorial` chat room on Spack slack. Join at:

slack.spack.io

You can ask questions here any time!

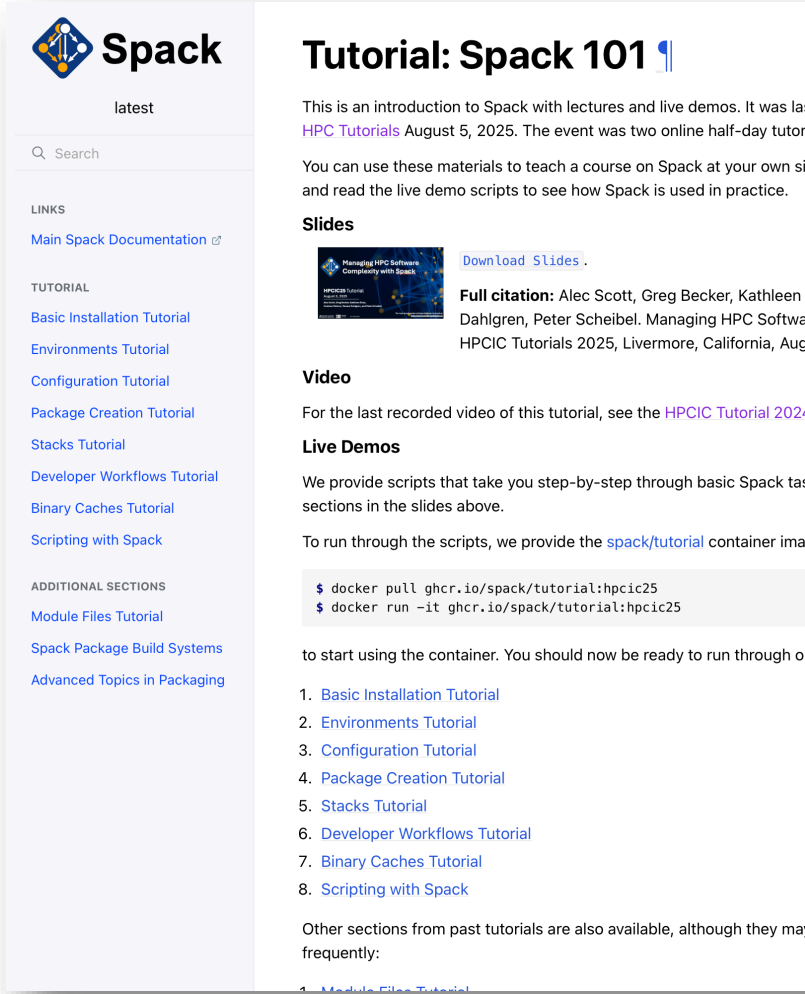

Claim a VM instance at: bit.ly/spack-vms

	A	B	C	D	E
1	**Spack Tutorial VM Instances**				
2	Instructions:	1. Put your name in a box below to claim an account on a VM insta			
3		2. Log in to your VM:			
4		ssh <IP address>			
5		Login/password are both the username from your column below			
6					
7		**Login / Password**			
8	**IP Address**	spack1	spack2	spack3	spack4
9	35.90.43.21				
10	35.91.36.120		Your Name		
11	34.217.149.171				

If you're in the **spack2** column, your login and password are *both* **spack2**

ssh spack2@3.73.129.196

Claim a login by putting your name in the Google Sheet

Build systems moved to **spack/spack-packages**

```python
from spack_repo.builtin.build_systems.autotools import AutotoolsPackage
from spack_repo.builtin.build_systems.cmake import CMakePackage

from spack.package import *


class ZlibNg(AutotoolsPackage, CMakePackage):
    ...
```

More on direct dependencies with %

- You could previously write:

```
pkg %gcc +foo      # +foo would associate with pkg, not gcc - will error in 1.0
```

- Now you'll need to write:

```
pkg +foo %gcc      # +foo associates with pkg
```

- We want these to be symmetric:

```
pkg +foo %dep +bar  # `pkg +foo` depends on `dep +bar` directly
pkg +foo ^dep +bar  # `pkg +foo` depends on `dep +bar` directly or transitively
```

- spack style --spec-strings --fix can remedy this automatically
  - Fixes YAML files, scripts, package.py files
  - Alternative was to have a very hard-to-explain syntax – we surveyed users and they decided it was better to break a bit than to explaining the subtleties of the first 10 years of Spack forever

Breaking changes ⚠️

1. It is no longer safe to assume every node has a compiler.
   a. The tokens `{compiler}`, `{compiler.version}`, and `{compiler.name}` in `Spec.format` expand to none if a Spec does not depend on C, C++, or Fortran.
   b. `spec.compiler` will default to the c compiler if present, else cxx, else fortran for backwards compatibility.
   c. The new default install tree projection is `{architecture.platform}/{architecture.target}/{name}-{version}-{hash}`

2. The syntax `spec["name"]` will only search link/run dependencies and *direct* build dependencies.
   - Previously, this would find deep, transitive deps, which was almost always the wrong behavior.
   - You can still hop around in the graph, e.g. spec["cmake"][”bzip2"] will find cmake’s link dependency

3. The % sigil in specs means “direct dependency”.
   - Can now say: `foo %cmake@3.26 ^bar %cmake@3.31`
   - `^` dependencies are unified, `%` dependencies are not

Art vs. science: reusing builds is not quite enough

- We get strange behavior when we have to build new packages
  - E.g.: Cmake depends on openssl for https
  - Minimizing builds will toggle this feature *off* to avoid a dependency
- **We want to prioritize reusing a package *if* the user already installed it**
  - Has to be more important than defaults, or we would never reuse
- **We want to prioritize package defaults *if* the package had to be built anyway**
  - Make *new* builds follow defaults

How can we do both?

minimize builds > package defaults

package defaults > minimize builds

We devised a two-level optimization scheme

```
build_priority(P, 200) :- build(P),    node(P).
build_priority(P, 0)   :- not build(P), node(P).

% priority + 200 IF we are building
#minimize{
   W@2+Priority,P
   : version_weight(P, W), build_priority(P, Priority)
}.
```

- Minimize builds, *unless we have to build*
  - Prioritize *defaults* for specs we *have* to build
- Last trick to get this to work:
  - All criteria must be formulated as minimizations
  - No built configuration can be “better” than a reused configuration

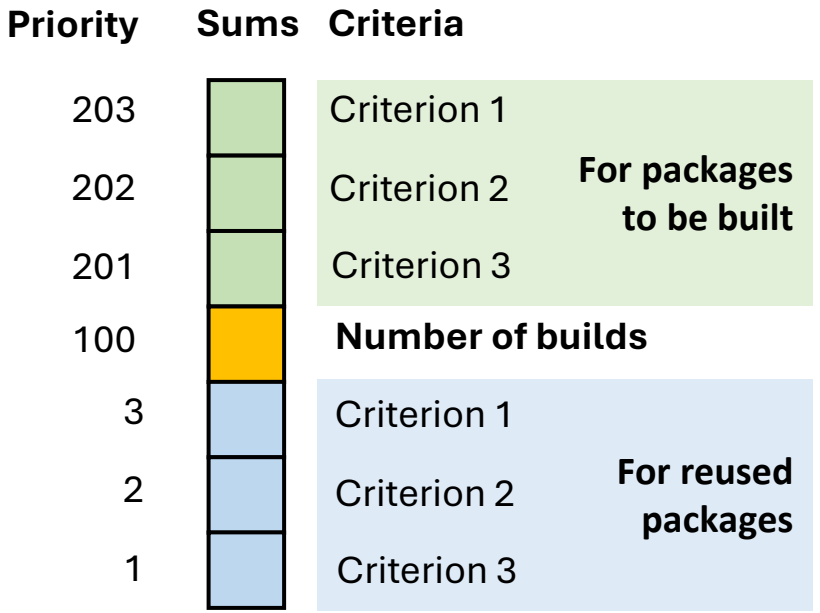

Objective vectors of sums are compared lexicographically from highest to lowest priority

About 2/3 of respondents are from the US

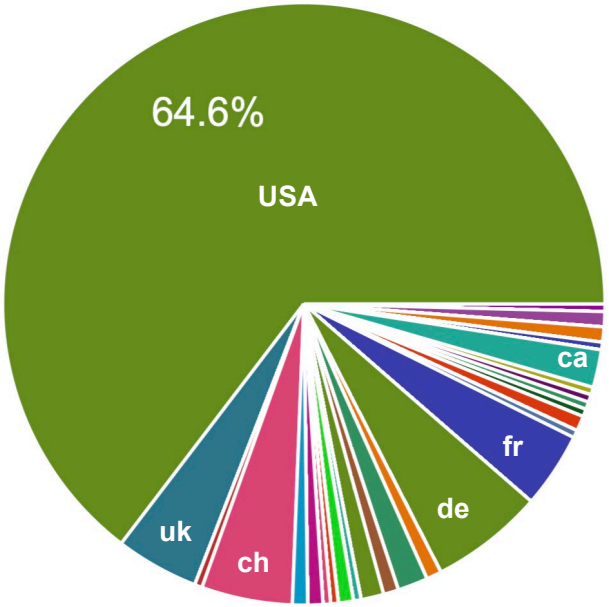

2025 Survey

Comparing to our documentation data, some countries are under-represented.

1 response

4 responses

3 responses

	Total	2,784 100% of total
1	United States	1,114 (40.01%)
2	Germany	214 (7.69%)
3	China	203 (7.29%)
4	United Kingdom	176 (6.32%)
5	India	116 (4.17%)
6	France	115 (4.13%)
7	Japan	88 (3.16%)
8	Italy	85 (3.05%)
9	Switzerland	81 (2.91%)
10	Hong Kong	60 (2.16%)
11	Canada	46 (1.65%)
12	Australia	45 (1.62%)

Past month: active users at spack.readthedocs.io

82% of users are doing HPC; 27% AI

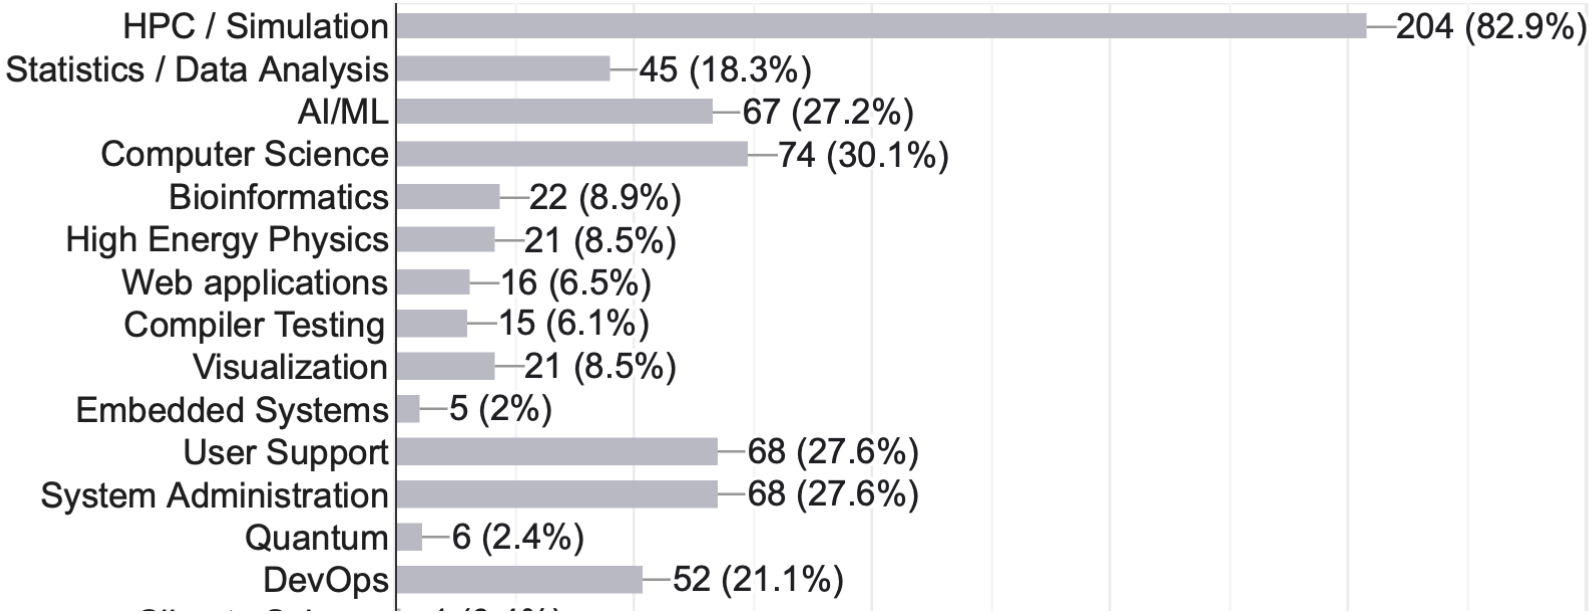

Most users are also contributors!

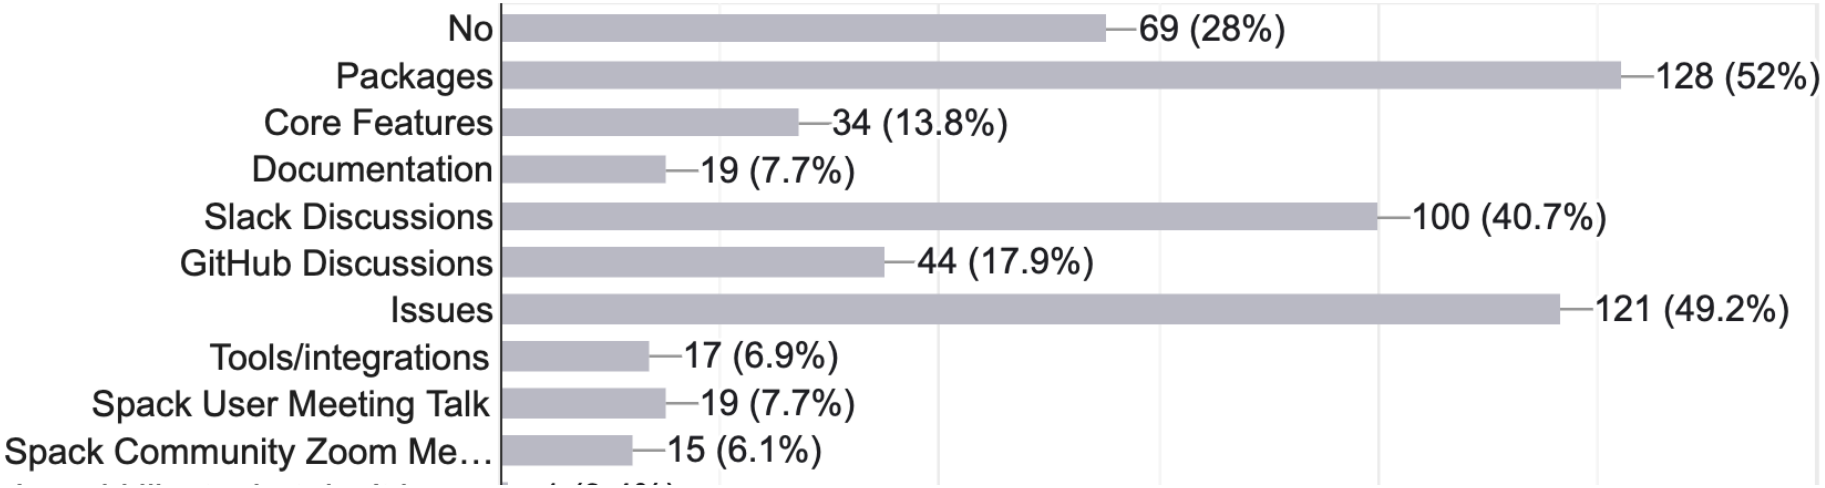

Most are within 2 releases of latest

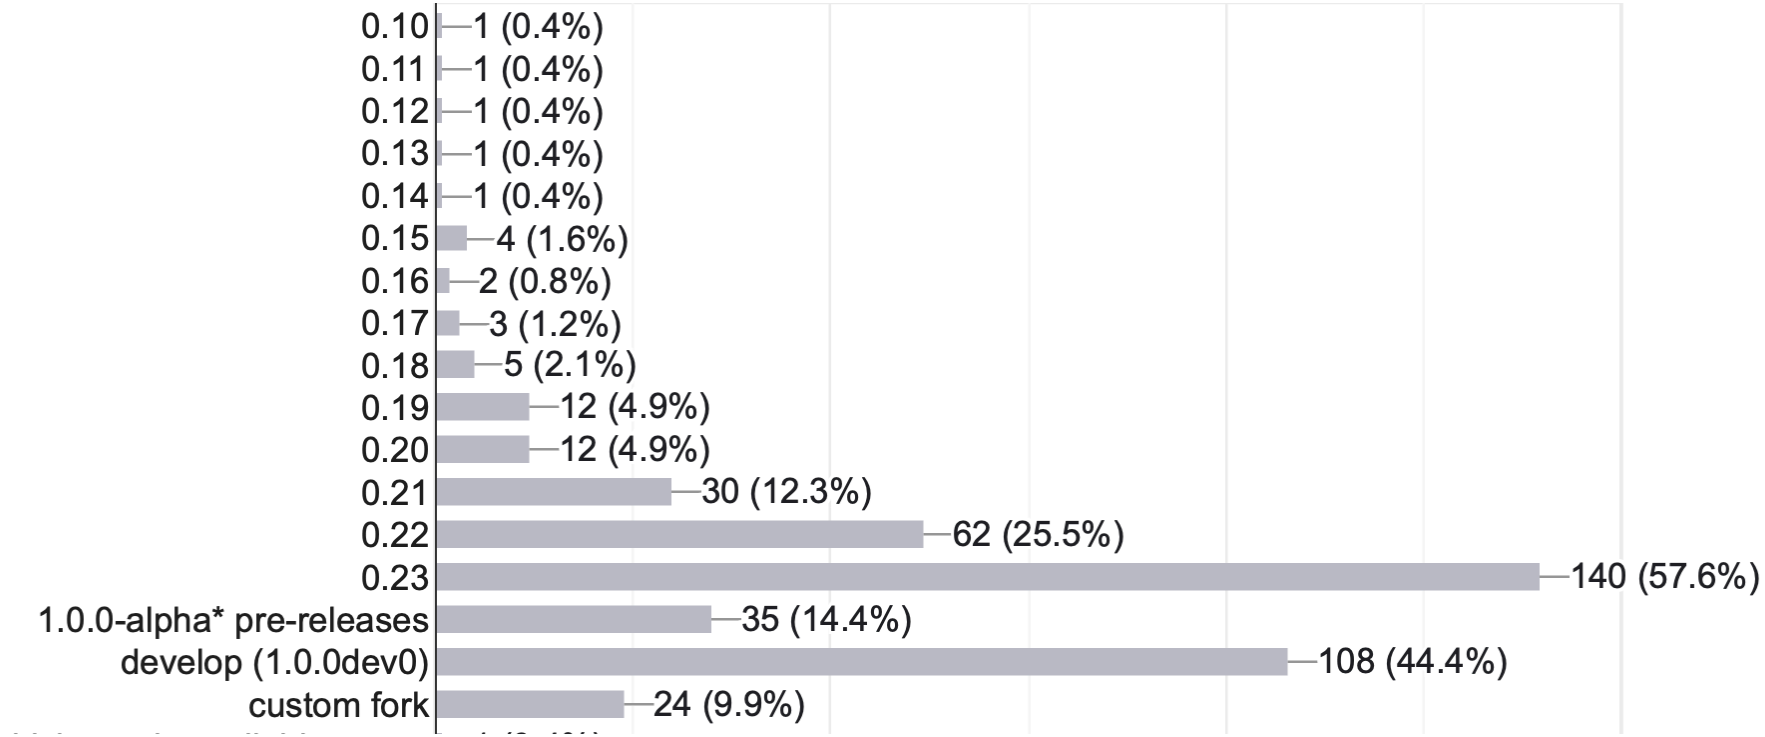

More users are on stable releases now than on develop

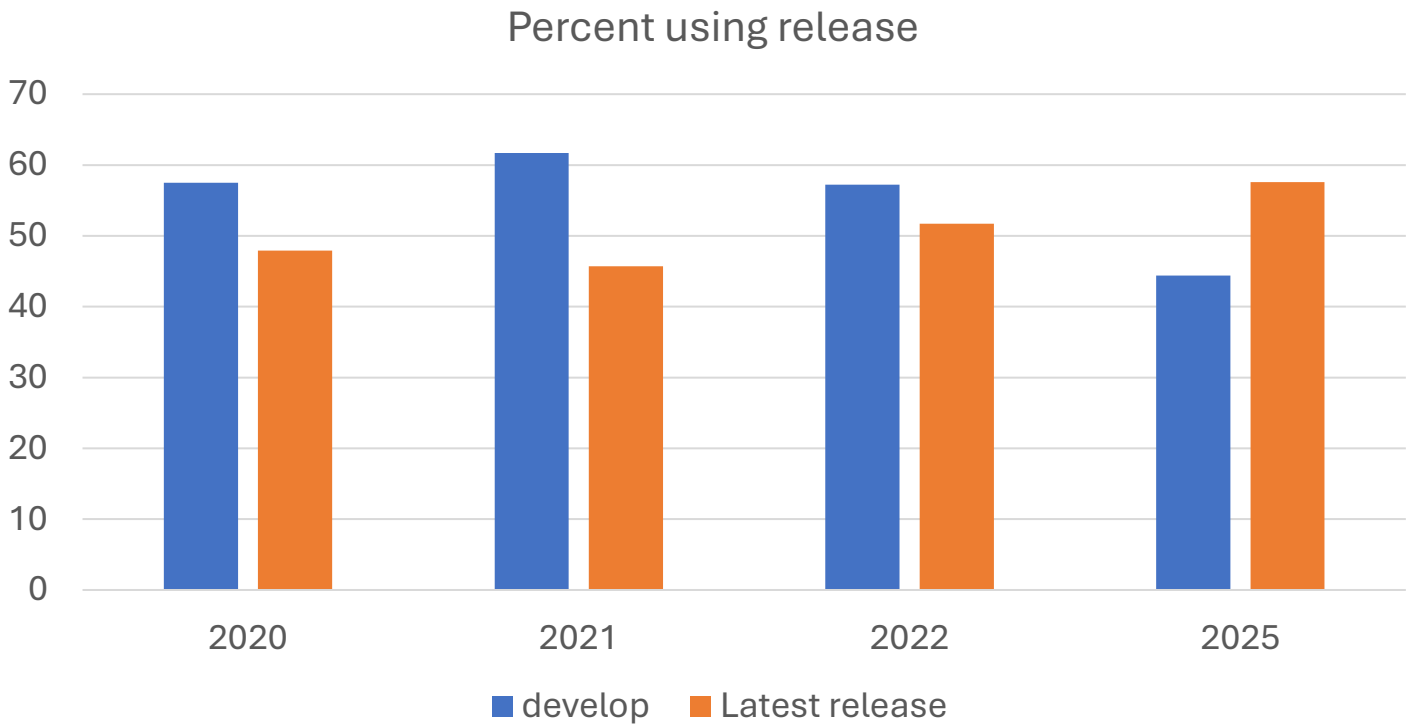

Environments have become the preferred way to load packages

How do you get installed Spack packages into your environment?

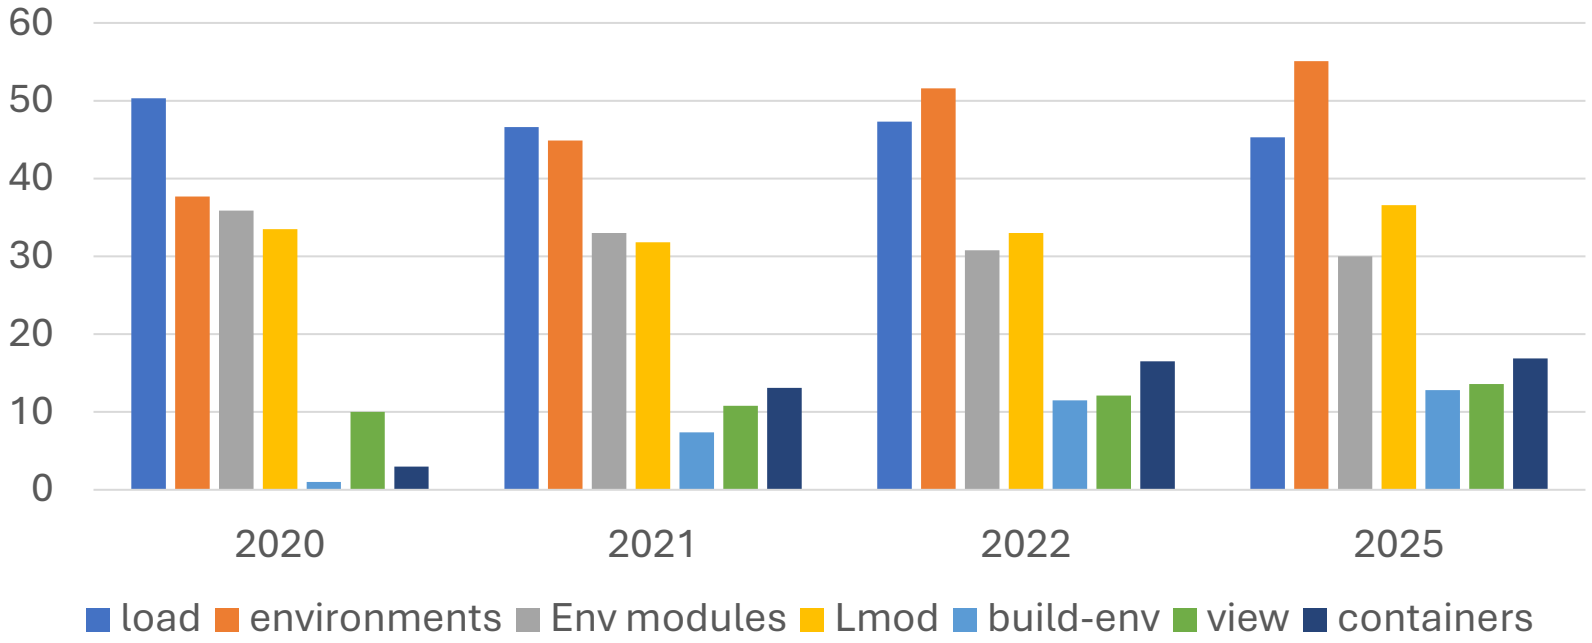

There are many packages outside of Spack's builtin repo

Do you have your own local Spack package repositories?

244 responses

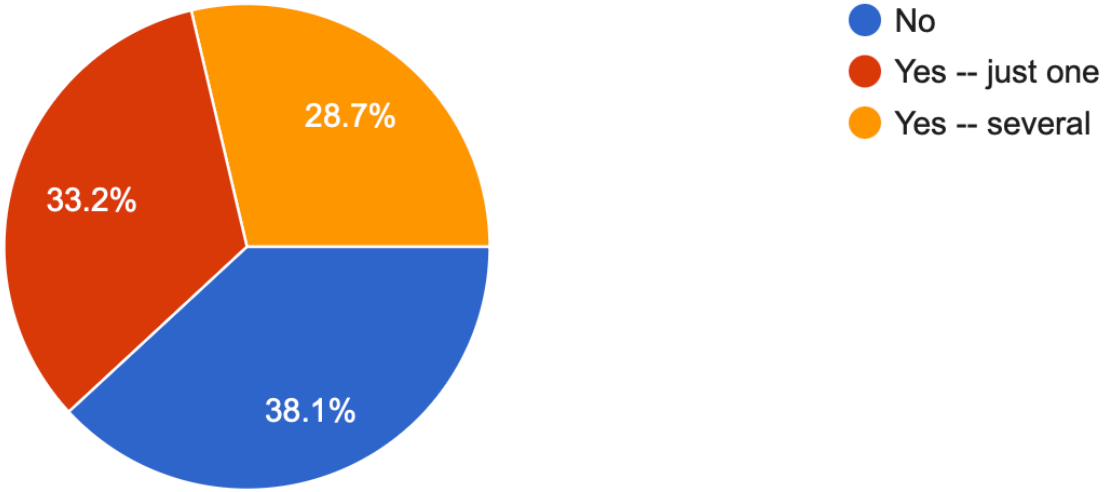

Users want better error messages, more performance, and a separate package repo

Rank these TBD Spack features by importance

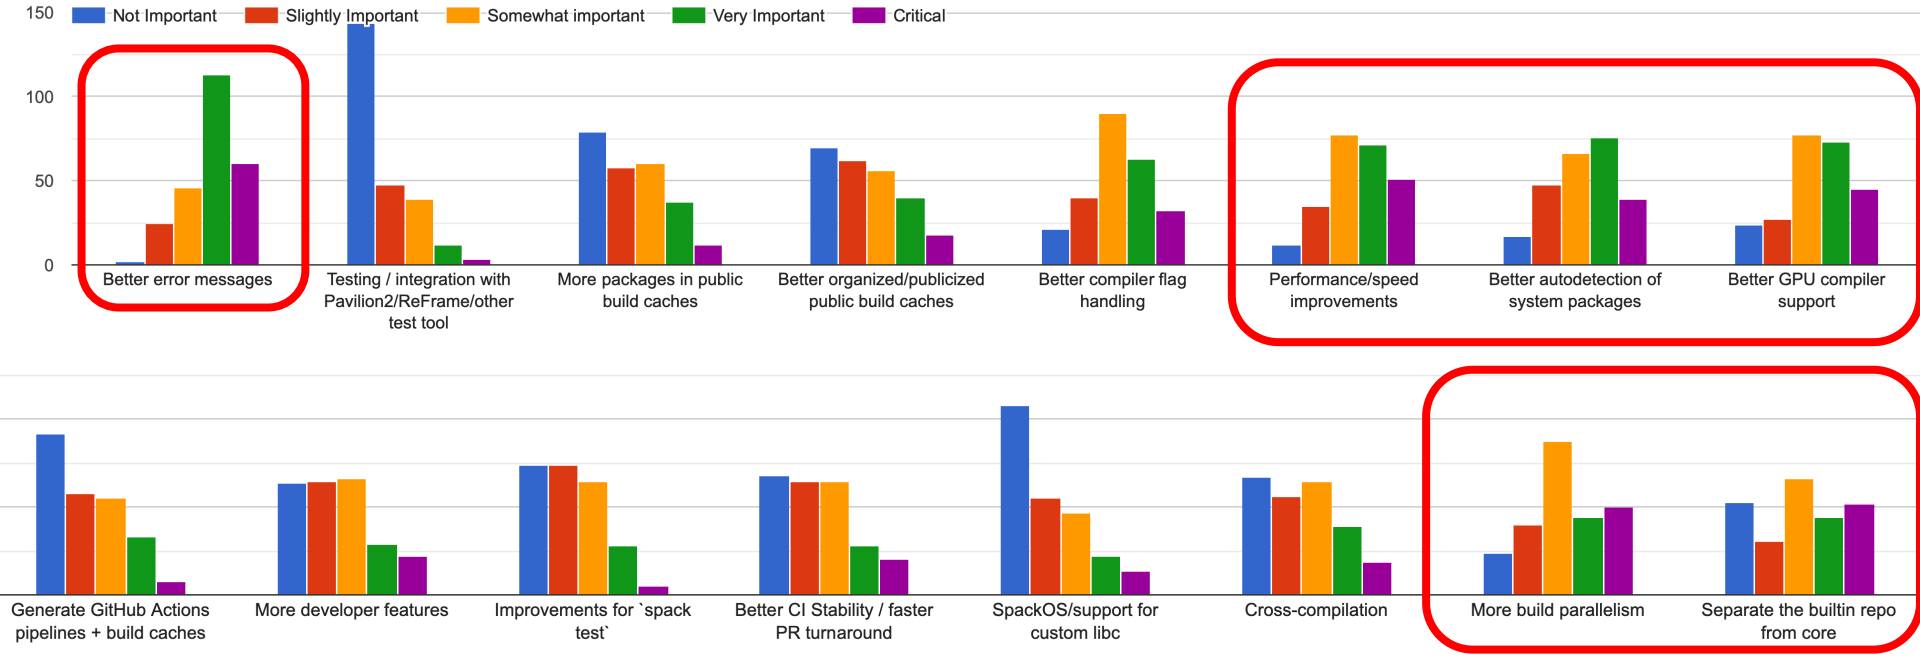

An isolated compilation environment allows Spack to easily swap compilers

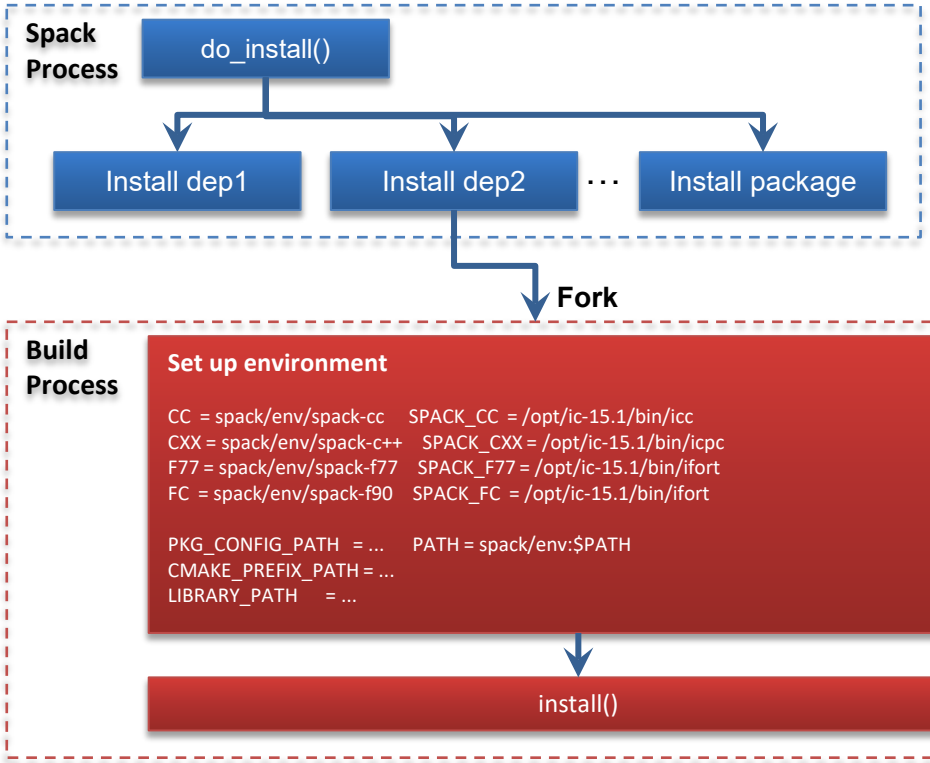

- **Forked build process isolates environment for each build.**
  Uses compiler wrappers to:
  - Add include, lib, and RPATH flags
  - Ensure that dependencies are found automatically
  - Load Cray modules (use right compiler/system deps)

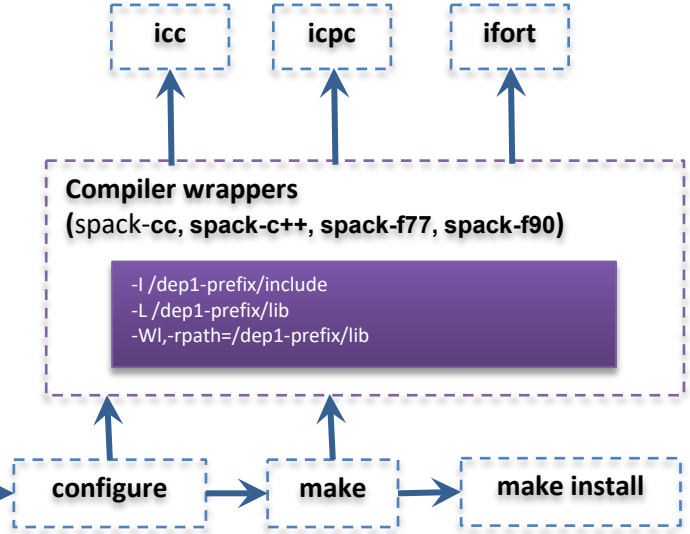

Core contributions are less diverse, but still comprise many organizations