

# Status of Packaging HEP Software using Spack

Patrick Gartung  
Scientific Software Infrastructure Department  
Scientific Computing Division  
Fermilab

# Why Packaging?

- Building all of the software needed to run a physics analysis can be complicated and time consuming.
- Some physicists have the background/interest to “DIY”.
- Most physicists want an “Easy Button” so they can get to making histograms.
- The solution: provide packages of prebuilt software or a simple script to build the packages.

# Spack

- originally designed as a “package manager for supercomputers” but applicable to many platforms
- <https://github.com/LLNL/spack> source code
- <http://software.llnl.gov/spack> documentation
- <https://www.computer.org/csdl/proceedings/sc/2015/3723/00/2807623.pdf> official paper
- <https://github.com/LLNL/spack/wiki> Information including info about weekly teleconference
- <https://groups.google.com/d/forum/spack> google group

# Features

- multiple compilers, multiple platforms and multiple variants from one package definition
- package definition declares dependencies on other packages
  - syntax for a range of dependency versions
  - dependencies are concretized into a DAG
  - a hash is generated from the DAG
  - reuse of dependencies that meet the criteria (in the range of versions and variants)
- compiler and linker wrappers that set rpath to declared dependencies
  - your program always find the right library
  - setting rpath circumvents the issue of SIP on OS X 10.11 squashing DYLD\_LIBRARY\_PATH in shell subprocesses
- generates tk environment modules and dotkit files for runtime environment needs
  - can be extended to generate other environment setup scripts — lmod modules in development

# hep-spack

- An add-on repo for Spack with common HEP packages
- <https://github.com/HEP-SF/hep-spack>
- See README.md for directions on adding packages to spack
- clhep, cppunit, fastjet, gccxml, geant4, xerces-c, graphviz, py-py2neo, py-mistune, root6, xrootd
- 376 spack built-in packages and counting
- <https://github.com/HEP-SF/hep-spack/wiki/spack-builtin-packages>

# Platforms I've built on

- OSX10.10 with system clang and bootstrapped gcc 4.9.3
- SL7 with bootstrapped gcc 4.9.3
- SLF6 with bootstrapped gcc 4.9.3
- Ubuntu 14.04 with bootstrapped gcc 4.9.3
- <https://github.com/HEP-SF/hep-spac/wiki/hep-spac-packages> Product Matrix
- <https://github.com/HEP-SF/hep-spac/wiki/hep-spac-packages-built-with-dependencies-by-OS> list of products from spac including hep-spac packages and dependencies built.

# Feedback and Experience

- Spack worked out of the box on SLF6, SL7 and Ubuntu 14.04
  - Digging a little deeper than README.md resolved any initial misunderstandings.
- Spack worked out of the box on OS X with Xcode command line tools installed (clang)
  - Fortran support on OS X requires an Homebrew or MacPorts install of gcc with gfortran, but that's normal.
- Spack developers and community are very helpful.
  - Google group and weekly teleconference (report from HEP community included)
  - Many branches on forked repos and pull requests have W.I.P. implementations of features of interest to HEP
  - Branch that specifically addressed compiling gcc on OS X.
  - Resolved a bug I found compiling gcc on OS X. Spack compiler wrapper generated command "ld -r -rpath ...". On linux the -rpath is ignored, on OS X this errors out.

# W.I.P. Features

- Create tarballs and relocate pre-built binaries
  - <https://github.com/LLNL/spack/pull/445> from Benedikt Hegner
  - <https://github.com/amundson/spack/tree/testmods> from Jim Amundson
  - <https://github.com/LLNL/spack/pull/631> make rpms



# W.I.P. Features

- Development and run-time environment setup
  - <https://github.com/brettviren/spack/tree/feature/views> from Brett Viren
  - <https://github.com/LLNL/spack/pull/721> Recursively generate a list of “module load” commands that could be used to define a runtime environment
  - <https://github.com/LLNL/spack/pull/543> Add the spconfig command to setup up dependencies for a build environment

# W.I.P. Features

- Better platform definitions
  - <https://github.com/LLNL/spack/pull/561> Defines Cray platform but also distinguishes Ubuntu and Centos.
  - <https://github.com/HEP-SF/documents/tree/master/HSF-TN/draft-2015-NAM> HSF naming conventions document
- Parallel builds of independent packages
  - <https://github.com/LLNL/spack/pull/843> from Benedikt Hegner

# Missing Features

- Chaining multiple spack installations together (e.g. cvmfs and home directory) has been requested
  - <https://github.com/LLNL/spack/issues/642>
- What other features are needed to make Spack useable to the HEP community?

# var/spack/repos/fnal/packages/gccxml/package.py

```
from spack import *
```

```
class Gccxml(Package):  
    homepage = "https://github.com/gccxml/gccxml"  
    url      = "https://github.com/gccxml/gccxml"
```

```
    version('2013-07-31', git='https://github.com/gccxml/gccxml.git', commit='3afa')
```

```
    depends_on('cmake @2.8.12.2:')
```

```
    def install(self, spec, prefix):  
        cmake(*std_cmake_args)  
        make()  
        make("install")
```

# var/spack/repos/hep-spack/packages/root/package.py

```
from spack import *
import sys

class Root(Package):
    """ROOT is a data analysis framework."""
    homepage = "https://root.cern.ch"
    url      = "https://root.cern.ch/download/root_v6.07.02.source.tar.gz"

    version('6.07.02', '3fb585bf9fa6ce06ca503173c8bee107')
    version('6.06.02', 'e9b8b86838f65b0a78d8d02c66c2ec55')

    if sys.platform == 'darwin': patch('math_uint.patch')

    depends_on("cmake")
    depends_on("pcre")
    depends_on("fftw")
    depends_on("clhep@2.3.2.2")
    depends_on("geant4")
    depends_on("graphviz")
    depends_on("python")
    depends_on("gsl")
    depends_on("libxml2+python")
    depends_on("jpeg")
    depends_on("openssl@1.0.2f")
    depends_on("freetype")
    if sys.platform != 'darwin':
        depends_on("libpng")

    def install(self, spec, prefix):

        build_directory = join_path(self.stage.path, 'spack-build')
        source_directory = self.stage.source_path
```

# var/spack/repos/hep-spack/packages/root/package.py

```
options=[source_directory]
if '+debug' in spec:
    options.append('-DCMAKE_BUILD_TYPE:STRING=Debug')
else:
    options.append('-DCMAKE_BUILD_TYPE:STRING=Release')
```

```
options.extend(std_cmake_args)
```

```
if sys.platform == 'darwin':
    darwin_options= [
        '-Dcastor=OFF',
        '-Drfio=OFF',
        '-Dbonjour=OFF',
        '-Dcocoa=OFF',
        '-Dx11=ON',
        '-Ddcache=OFF' ]
    options.extend(darwin_options)
```

```
def setup_dependent_environment(self, module, spec, dep_spec):
    """Root wants to set ROOTSYS"""
    os.environ['ROOTSYS'] = self.prefix
```

```
with working_dir(build_directory, create=True):
    cmake(*options)
    make()
    make("install")
```

```
def url_for_version(self, version):
    """Handle ROOT's unusual version string."""
    return "https://root.cern.ch/download/root_v%s.source.tar.gz" % version
```

# var/spack/repos/hep-spack/packages/geant4/package.py

```
from spack import *
```

```
class Geant4(Package):
```

```
    """Geant4 is a toolkit for the simulation of the passage of particles through matter. Its areas of application include high energy, nuclear and accelerator physics, as well as studies in medical and space science."""
```

```
    homepage = "http://geant4.cern.ch/"
```

```
    url      = "http://geant4.cern.ch/support/source/geant4.10.01.p03.tar.gz"
```

```
    version('10.01.p03', '4fb4175cc0dabcd517443fbdccd97439')
```

```
    depends_on("xerces-c")
```

```
    depends_on("cmake")
```

```
    depends_on("clhep@2.3.1.1")
```

```
    def install(self, spec, prefix):
```

```
        cmake_args = list(std_cmake_args)
```

```
        cmake_args.append('-DXERCESC_ROOT_DIR:STRING=%s'%spec['xerces-c'].prefix)
```

```
        cmake_args.append('-DGEANT4_BUILD_CXXSTD=c++11')
```

# Unsolvable DAG

- Made root dependent on one version of clhep and geant4.
- Made geant4 dependent on a different version of clhep.
- Did spack catch this?
- Yes, but the error is a little cryptic:

```
[vagrant@localhost geant4]$ spack install root  
==> Error: Invalid spec: 'clhep@2.3.2.2^cmake@3.2:'. Package clhep requires version 2.3.1.1, but spec asked for 2.3.2.2
```



# Spackathon

- Create spack packages for all of the external packages built with lcgcmake
- <https://gitlab.cern.ch/sft/lcgcmake/blob/master/externals/CMakeLists.txt>
- <http://lcgsoft.web.cern.ch/lcgsoft/release/84/>
- Grab the links to tarfiles from the page above and add Spack package definitions for each to the hep-sf package repo.
- Follow the example from <https://github.com/HEP-SF/hep-spack>
- The target platform is CernCentos7 but you can try on any linux distribution.
- The target compiler will be gcc 5 built with spack.
- For OS X a linux VM or remote linux node is recommended.
- In preparation for the Spackathon, you should try building gcc 5 with spack.

OSX10.10 with clang / gcc 4.2.1

==> 47 installed packages.

```
-- darwin-x86_64 / clang@7.0.2-apple -----  
boost@1.60.0+atomic+chrono~cxx14+date_time~debug+filesystem+graph~icu_support+iostreams+locale+log+math~mpi  
+multithreaded+program_options~python+random+regex+serialization+shared+signals+singlethreaded+system+test  
+thread+wave  
bzip2@1.0.6  
cmake@3.5.2~doc+ncurses+openssl~qt  
fftw@3.3.4+float+long_double~mpi~openmp~quad  
freetype@2.5.3  
ghostscript@9.16  
graphviz@2.38.0  
gsl@2.1  
jpeg@9a  
libpng@1.6.16  
libxml2@2.9.2+python  
ncurses@6.0  
openssl@1.0.2f  
pcre@8.38  
python@2.7.11  
readline@6.3  
root@6.07.02  
sqlite@3.8.5  
swig@3.0.8  
xz@5.2.2  
zlib@1.2.8  
  
-- darwin-x86_64 / gcc@4.2.1 -----  
gcc@4.9.3~binutils~gold gmp@6.1.0 mpc@1.0.3 mpfr@3.1.4
```

OSX10.10 with bootstrapped gcc 4.9.3

```
-- darwin-x86_64 / gcc@4.9.3 -----  
boost@1.60.0+atomic+chrono+date_time~debug+filesystem+graph~icu_support+iostreams+locale+log  
+math~mpi+multithreaded+program_options~python+random+regex+serialization+shared+signals  
+singlethreaded+system+test+thread+wave  
bzip2@1.0.6  
cmake@3.5.2~doc+ncurses+openssl~qt  
environment-modules@3.2.10  
fftw@3.3.4+float+long_double~mpi~openmp~quad  
freetype@2.5.3  
ghostscript@9.16  
graphviz@2.38.0  
gsl@2.1  
jpeg@9a  
libxml2@2.9.2+python  
ncurses@6.0  
openssl@1.0.2f  
pcre@8.38  
python@2.7.11  
readline@6.3  
root@6.07.02  
sqlite@3.8.5  
swig@3.0.2  
tcl@8.6.5  
xz@5.2.2  
zlib@1.2.8
```

SL7 ==> 52 installed packages.

-- linux-x86\_64 / gcc@4.8.5 -----

binutils@2.26+gold~krellpatch~libiberty gcc@4.9.3+binutils+gold gmp@6.1.0 mpc@1.0.3 mpfr@3.1.4

-- linux-x86\_64 / gcc@4.9.3 -----

boost@1.60.0+atomic+chrono~cxx14+date\_time~debug+filesystem+graph~icu\_support+iostreams+locale+log+math~mpi+multithreaded  
+program\_options~python+random+regex+serialization+shared+signals+singlethreaded+system+test+thread+wave

bzip2@1.0.6

cmake@3.5.1~doc+ncurses+openssl~qt

cmake@3.5.2~doc+ncurses+openssl~qt

cppunit@1.12.1

dbus@1.11.2

expat@2.1.0

fastjet@3.1.3

fftw@3.3.4+float+long\_double~mpi~quad

freetype@2.5.3

gccxml@20150423

geant4@10.01.p03

ghostscript@9.16

glib@2.42.1

graphviz@2.38.0

gsl@2.1

jpeg@9a

lcms@2.6

lcov@1.12

libevent@2.0.21+openssl

libffi@3.2.1

libmng@2.0.2

libpng@1.6.16

libtiff@4.0.3

libxcb@1.11.1

libxml2@2.9.2~python

libxml2@2.9.2+python

ncurses@6.0

openssl@1.0.2g

pcre@8.38

py-mistune@0.7.1

py-py2neo@2.0.8

py-setuptools@20.7.0

python@2.7.11

qt@4.8.6~gtk~krellpatch~mesa

readline@6.3

root@6.06.02

root@6.07.02

sqlite@3.8.5

swig@3.0.2

tbb@4.4.3

tmux@2.1

xcb-proto@1.11

xerces-c@3.1.3

xrootd@4.3.0

xz@5.2.2

zlib@1.2.8

SLF6 ==> 41 installed packages.

-- linux-x86\_64 / gcc@4.4.7 -----

binutils@2.26+gold~krellpatch~libiberty gcc@4.9.3+binutils+gold gmp@6.1.0 mpc@1.0.3 mpfr@3.1.4

-- linux-x86\_64 / gcc@4.9.3 -----

boost@1.60.0+atomic+chrono~cxx14+date\_time~debug+filesystem+graph~icu\_support+iostreams+locale+log+math~mpi  
+multithreaded+program\_options~python+random+regex+serialization+shared+signals+singlethreaded+system+test+thread+wave

bzip2@1.0.6

clhep@2.3.1.1+cxx14~debug

cmake@3.5.1~doc+ncurses+openssl~qt

cmake@3.5.2~doc+ncurses+openssl~qt

cppunit@1.12.1

fastjet@3.1.3

fftw@3.3.4+float+long\_double~mpi~quad

freetype@2.5.3

gccxml@20150423

geant4@10.01.p03

ghostscript@9.16

graphviz@2.38.0

gsl@2.1

jpeg@9a

lcov@1.12

libpng@1.6.16

libxml2@2.9.2~python

libxml2@2.9.2+python

ncurses@6.0

openssl@1.0.2g

pcre@8.38

py-mistune@0.7.1

py-py2neo@2.0.8

py-setuptools@20.7.0

python@2.7.11

readline@6.3

root@6.06.02

root@6.07.02

sqlite@3.8.5

swig@3.0.2

tbb@4.4.3

xerces-c@3.1.3

xrootd@4.3.0

xz@5.2.2

zlib@1.2.8

Ubuntu 14.04 ==> 45 installed packages.

```
-- linux-x86_64 / gcc@4.8 -----
binutils@2.26+gold~krellpatch~libiberty flex@2.5.39 gcc@5.3.0+binutils+gold isl@0.14
m4@1.4.17+sigsegv mpfr@3.1.4
bison@3.0.4 gcc@4.9.3+binutils+gold gmp@6.1.0 libsigsegv@2.10 mpc@1.0.3
-- linux-x86_64 / gcc@4.9.3 -----
bison@3.0.4
boost@1.60.0+atomic+chrono~cxx14+date_time~debug+filesystem+graph~icu_support+iostreams+locale+log+math~mpi
+multithreaded+program_options~python+random+regex+serialization+shared+signals+singlethreaded+system+test+thread+wave
bzip2@1.0.6
clhep@2.3.1.1+cxx14~debug
cmake@3.5.2~doc+ncurses+openssl~qt
fastjet@3.1.3
fftw@3.3.4+float+long_double~mpi~quad
flex@2.5.39
freetype@2.5.3
gccxml@20150423
geant4@4.10.01.p03
ghostscript@9.16
gmp@6.1.0
graphviz@2.38.0
jpeg@9a
lcov@1.12
libpng@1.6.16
libsigsegv@2.10
libxml2@2.9.2+python
m4@1.4.17+sigsegv
ncurses@6.0
openssl@1.0.2g
pcre@8.38
py-mistune@0.7.1
py-py2neo@2.0.8
py-setuptools@20.7.0
python@2.7.11
readline@6.3
sqlite@3.8.5
swig@3.0.2
xerces-c@3.1.3
xrootd@4.3.0
xz@5.2.2
zlib@1.2.8
```