



Docker Spacklove

or: How I learned to stop worrying and love the package manager

Hadrien Grasland

LAL – Orsay

2018-09-19

Introduction

- **Good comment on a CHEP talk:**
 - « *Verrou looks like a really nice tool, but I wouldn't build a patched Valgrind just to try it out* »
- **Encouraged me to work on a long-time pain point:**
 - Software packaging & portability challenges in HEP
 - How toolchain improvements could address them

(Conflicting) Requirements



Two visions of portability

- **Most of CERN & WLCG:**
 - Software is portable if it runs on all SLC6 Grid nodes
 - Can assume CVMFS availability, slowly deprecating AFS
- **Rest of the world:**
 - Software is portable if it runs well on *my computer*
 - Using many Linux distros, some other Unices
 - Install in one command, integrate well with my system
 - Always-on & very fast Internet is not a given

Two user profiles

- **« Testers »: Experimenting around**
 - Want to get something running as quickly as possible
 - Can tolerate low integration with the host OS
- **« Developers »: Putting oneself at ease**
 - Want to leverage full system (dev tools, GUIs, drivers...)
 - Can tolerate an occasional large build if automated
- **Can assume some computing & Unix proficiency**

Figuring it out



Serving the « testers »

- **Containers seem best for this use case**
 - Lighter than a VM
 - Simpler to build and maintain
 - Download & run with 1-2 commands
 - Easier to attach system resources (e.g. filesystem)
- **Which implementation?***
 - Singularity is a good idea for scientific computing
 - But Docker wins *massively* on usability for now

Scaling up to « developers »

- **Need a custom build for full system integration**
 - Use system compiler/libc/tools, link with HW drivers...
 - Avoid too much isolation from host system
- **Integration vs reproducibility trade-off**
 - Remember: no manpower to test many configurations!
 - Best provide our dependencies + opt-in to system libs
 - Need a package manager to sort the dependencies out

Picking package managers

- **Various forms of flexibility needed:**
 - Many programming languages (C++, Python, Fortran...)
 - Good portability across Linuxes, ideally Unices
 - Minimize tampering with user environment
 - Allow opt-in dependencies on host system
 - Coexistence of « cousin » builds (e.g. ROOT C++14 & 17)
- **Spack seems most appropriate* for this use case**

What package managers bring to containers

- **Simplified maintenance**
 - Turn big Dockerfile into a few spack commands
 - Reusing work across containers is much simpler
- **Easier for « testers » to become « developers »**
 - Dockerfiles are specific to a base Linux distribution
 - But Spack commands are portable across systems
 - They help reproducing a container's setup locally



In practice

Stress-testing the approach

- **Previously moved my build recipes to Dockerfiles**
 - Simplifies automated testing (easy rollback on failure)
 - Easier to move my dev environment to another machine
 - Quickly share a dev environment with coworkers
- **Here, tried to move most shell work to Spack**
 - Test projects: Verrou, Templight, ACTS, Gaudi
 - Aiming for full builds (~all optional dependencies)

Building and testing Gaudi* (before)

```
1 # === DOCKER-SPECIFIC HACKERY ===
2
3 FROM hgrasland/root-tests
4 LABEL Description="opensUSE Tumbleweed environment for Gaudi" Version="0.1"
5 CMD bash
6
7
8 # === SYSTEM SETUP ===
9
10 # Update the host system
11 RUN zypper ref && zypper dup -y
12
13 # Install non-ROOT requirements
14 RUN zypper in -y doxygen graphviz cppunit-devel gdb unzip libxerces-c-devel
15     uuid-devel libunwind-devel gperf-tools gperf-tools-devel
16     jemalloc-devel ncurses5-devel ninja wget python2-nose
17     python2-networkx which curl libuuid-devel
18
19 # === INSTALL (OLDER) BOOST ===
20
21 # Download Boost v1.66 (Gaudi does not support v1.67+ yet)
22 RUN git clone --recursive -j8 --branch=boost-1.66.0 --depth=1
23     https://github.com/boostorg/boost.git
24
25 # Build and install Boost
26 RUN cd boost
27     && ./bootstrap.sh --with-python=python2.7
28     && ./b2 -j8
29     && ./b2 install
30
31 # Work around Boost's brain damaged build system
32 RUN cp -rf boost/libs/program_options/include/boost/*
33     /usr/local/include/boost/
34     && cp -rf boost/libs/utility/include/boost/*
35     /usr/local/include/boost/
36     && cp -rf boost/libs/circular_buffer/include/boost/*
37     /usr/local/include/boost/
38     && cp -rf boost/libs/ptr_container/include/boost/*
39     /usr/local/include/boost/
40     && cp -rf boost/libs/assign/include/boost/*
41     /usr/local/include/boost/
42
43 # Get rid of the Boost build directory
44 RUN rm -rf boost
45
46 # === INSTALL C++ GUIDELINE SUPPORT LIBRARY ===
47
48 # Download the GSL
49 RUN git clone --depth=1 https://github.com/Microsoft/GSL.git
50
51 # Build the GSL
52 RUN cd GSL && mkdir build && cd build
53     && cmake -GNinja -DCMAKE_BUILD_TYPE=RelWithDebInfo
54     -DGSL_CXX_STANDARD=14 ..
55     && ninja
56
57
58 # Check that the GSL build is working properly
59 RUN cd GSL/build && ctest -j8
60
61 # Install the GSL
62 RUN cd GSL/build && ninja install
63
64 # Get rid of the GSL build directory
65 RUN rm -rf GSL
66
67 # === INSTALL RANGE-V3 ===
68
69 # Download the range-v3 library (v0.3.5)
70 RUN git clone --branch=0.3.6 --depth=1
71     https://github.com/ericniebler/range-v3.git
72
73 # Build range-v3
74 RUN cd range-v3 && mkdir build && cd build
75     && cmake -GNinja -DRANGES_CXX_STD=14 .. && ninja
76
77 # Check that the range-v3 build is working properly
78 RUN cd range-v3/build && ctest -j8
79
80 # Install range-v3
81 RUN cd range-v3/build && ninja install
82
83 # Get rid of the range-v3 build directory
84 RUN rm -rf range-v3
85
86 # === INSTALL AIDA ===
87
88 # Download, extract and delete the AIDA package
89 RUN mkdir AIDA && cd AIDA
90     && wget
91     ftp://ftp.slac.stanford.edu/software/freehep/AIDA/v3.2.1/aida-3.2.1.zip
92     && unzip -q aida-3.2.1.zip
93     && rm aida-3.2.1.zip
94
95 # Install the AIDA headers
96 RUN cp -r AIDA/src/cpp/AIDA /usr/include/
97
98 # Get rid of the rest of the package, we do not need it
99 RUN rm -rf AIDA
100
101 # === INSTALL CLHEP ===
102
103 # Download CLHEP
104 RUN git clone --branch=CLHEP_2_4_1_0 --depth=1
105     https://gitlab.cern.ch/CLHEP/CLHEP.git
106
107 # Build CLHEP
108 RUN cd CLHEP && mkdir build && cd build
109     && cmake -GNinja .. && ninja
110
111 # Test our CLHEP build
112 RUN cd CLHEP/build && ctest -j8
113
114 # Install CLHEP
115 RUN cd CLHEP/build && ninja install
116
117
118 # Get rid of the CLHEP build directory
119 RUN rm -rf CLHEP
120
121 # === INSTALL HEPDC V2 ===
122
123 # Download and extract HepDC v2
124 RUN curl
125     http://lcgapp.cern.ch/project/simu/HepPDT/download/HepPDT-2.06.01.tar.gz \
126     | tar -xz
127
128 # Build and install HepDC
129 RUN cd HepPDT-2.06.01 && mkdir build && cd build
130     && ../configure && make -j8 && make install
131
132 # Get rid of the HepDC build directory
133 RUN rm -rf HepPDT-2.06.01
134
135 # === INSTALL HEPMC V3 ===
136
137 # Download HepMC v3
138 RUN git clone --depth=1 https://gitlab.cern.ch/hepmc/HepMC3.git
139
140 # Build and install HepMC
141 RUN cd HepMC3 && mkdir build && cd build
142     && cmake -DCMAKE_BUILD_TYPE=RelWithDebInfo ..
143     && make -j8 && make install
144
145 # Get rid of the HepMC build directory
146 RUN rm -rf HepMC3
147
148 # === INSTALL HEPMC V2 ===
149
150 # NOTE: Why are we overwriting our HepMC3 install with a HepMC2 one, you may
151 # wonder? The answer has to do with RELAX being hopelessly broken, and
152 # expecting the CMake files of HepMC3 together with the headers of HepMC2
153
154 # Download HepMC v2
155 RUN git clone --depth=1 https://gitlab.cern.ch/hepmc/HepMC.git
156
157 # Build HepMC
158 RUN cd HepMC && mkdir build && cd build
159     && cmake -Dmomentum=GEV -Dlength=MM .. && make -j8
160
161 # Test our build of HepMC
162 RUN cd HepMC/build && make test -j8
163
164 # Install HepMC and remove bits of HepMC3 (ugh...)
165 RUN cd HepMC/build && make install
166     && rm /usr/local/lib64/libHepMC.so /usr/local/lib64/libHepMC.a
167
168 # Get rid of the HepMC build directory
169 RUN rm -rf HepMC
170
171 # === INSTALL RELAX ===
172
173 # Download and extract RELAX (yes, this file is not actually gzipped)
174 RUN curl http://lcgpackages.web.cern.ch/lcgpackages/tarFiles/sources/RELAX-root6.tar.gz \
175     | tar -x
176
177 # Build and install RELAX (wow, such legacy, much hacks!)
178 RUN cd RELAX && mkdir build && cd build
179     && ln -s 'which genreflex' /genreflex
180     && export CXXFLAGS="-Iusr/local/include/root/"
181     && cmake .. -DCMAKE_BUILD_TYPE=RelWithDebInfo
182     && make -j8 && make install
183     && rm /genreflex && unset CXXFLAGS
184
185 # Get rid of the RELAX build directory
186 RUN rm -rf RELAX
187
188 # === ATTEMPT A GAUDI TEST BUILD ===
189
190 # Clone the Gaudi repository
191 RUN git clone --origin upstream https://gitlab.cern.ch/gaudi/Gaudi/
192
193 # Configure Gaudi
194 RUN cd Gaudi && mkdir build && cd build
195     && cmake -DGAUDI_DIAGNOSTICS_COLOR=ON -GNinja ..
196
197 # Configure the run-time linker
198 # NOTE: I am not sure why this is needed for this build specifically, but the
199 # Gaudi build will fail to find CLHEP if we don't do it.
200 RUN ldconfig
201
202 # Build Gaudi
203 RUN cd Gaudi/build && ninja
204
205 # Test the Gaudi build
206 # NOTE: Some Gaudi tests do ptrace system calls, which are not allowed in
207 # unprivileged docker containers because they leak too much information
208 # about the host. You can allow the container to run these tests
209 # by passing the "--security-opt=seccomp:unconfined" flag to docker run,
210 # but for some strange reason this flag cannot be passed to docker build.
211 # Therefore, we disable these tests during the docker image build.
212 RUN cd Gaudi/build && ctest -j8 -E "(google_auditors\heapchecker|event_timeout_abort)"
213
214 # Remove build byproducts to keep image light
215 RUN cd Gaudi/build && ninja clean
216
217 # === FINAL CLEAN UP ===
218
219 # Discard the system package cache to save up space
220 RUN zypper clean
```

Building and testing Gaudi* (after)

```
1 FROM hgrasland/root-tests:latest-cxx17
2 LABEL Description="openSUSE Tumbleweed environment for Gaudi" Version="0.1"
3 CMD bash
4
5 # Use my Gaudi package development branch
6 #
7 # TODO: Remove this once it's integrated in Spack
8 #
9 RUN cd /opt/spack && git fetch HadrienG2 && git checkout gaudi-package
10
11 # Build a spack spec for Gaudi
12 RUN echo "export GAUDI_SPACK_SPEC=\"\`gaudi@develop +tests +optional
13         ^ \`${ROOT_SPACK_SPEC}\`\"" >> \`${SETUP_ENV}
14
15 # Build Gaudi and its dependencies using Spack
16 RUN spack build \`${GAUDI_SPACK_SPEC}
17
18 # Test the Gaudi build
19 #
20 # NOTE: Some Gaudi tests do ptrace system calls, which are not allowed in
21 #       unprivileged docker containers because they leak too much information
22 #       about the host. You can allow the container to run these tests
23 #       by passing the "--security-opt=seccomp:unconfined" flag to docker run,
24 #       but for some strange reason this flag cannot be passed to docker build.
25 #       Therefore, we disable these tests during the docker image build.
26 #
27 RUN spack cd --build-dir \`${GAUDI_SPACK_SPEC} \
28     && cd spack-build \
29     && spack env gaudi+tests+optional \
30     ctest -j8 -E "(google_auditors\|.heapchecker|event_timeout_abort)"
31
32 # Drop the build to save space in the final Docker image
33 #
34 # (This will preserve dependencies, therefore reinstalling should be quick)
35 #
36 RUN spack clean \`${GAUDI_SPACK_SPEC}
```

- Does the same thing...
 - In 36 LOCs instead of 234
 - 2/3 of which are comments
- Where did the logic go?
 - Much is handled by Spack
 - Remainder is in packages
 - More easily reusable!

Conclusion

- **Containers and package managers are friends**
 - Share recipes across Dockerfiles & with others
 - Delegate repetitive churn to the package manager
 - Smooth user transition from virtualized to local build
- **State of my experiment:**
 - Verrou & Templight are merged in official Spack repo
 - Gaudi & ACTS* are waiting for ROOT in [WIP] PRs

* Only the DD4hep part is submitted at the moment, to avoid stacking [WIP] PRs.

Questions?



VMs vs Containers

- **Advantages of VM:**
 - Better portability to non-Linux / ancient kernels
 - Tigher isolation of untrusted code from host
 - Much better for GUI and other GPU work
- **Advantages of containers:**
 - Easier for end users
 - Lighter system resource usage
 - Can document & reproduce the build process

Docker vs Singularity

- **Singularity is a better idea for scientific compute**
 - No root-equivalent permission needed
 - No unnecessary isolation from host system
 - Self-contained images reduce need for registries
- **Docker is a more mature product**
 - More reliable toolchain
 - More straightforward workflows
 - Much better documentation on the Web

Comparing package managers

- Options which I studied most closely:

	Multi-language	Multi-OS	Env-preserving	Vendored deps	Native deps	Easy	Fast	Cousin builds	Non-HEP users
LCGCMake	Y	N	Y	Y	N	N	N	N	N
APT / RPM	Y	~	N	N	Y	Y	Y	N	Y
Spack	Y	Y	Y	Y	~	Y	N	Y	Y
Nix	Y	~	Y	Y	N	N	N	Y	Y

- Important points that did not matter for *this* study:
 - Institutional support from WLCG & CERN
 - Relocation & distribution of pre-built binaries