

Gentoo Prefix as a Physics Software Manager

Benda Xu (Tsinghua University), Guilherme Amadio (CERN),
Fabian Groffen (Gentoo), Michael Haubenwallner (Gentoo)




<https://www.gentoo.org>

Nov. 5, 2019 CHEP'19 Adelaide

Sophisticated Analysis and Simulation

- Data analysis and detector simulation advance fast.
- Physics software stack is getting deeper.

Example (Pythia physics event generator)



Brace the Complexity

- Bundling dependencies all together is not elegant.
 - ▶ Modularity is crucial for maintainability.
- Package manager is the industry standard to solve this.
 - ▶ It is invented by GNU/Linux distributions.
 - ▶ Debian founder Ian Murdock (1973–2015):
the single biggest advancement Linux has brought to the industry.

Brace the Complexity

- Bundling dependencies all together is not elegant.
 - ▶ Modularity is crucial for maintainability.
- Package manager is the industry standard to solve this.
 - ▶ It is invented by GNU/Linux distributions.
 - ▶ Debian founder Ian Murdock (1973–2015):
the single biggest advancement Linux has brought to the industry.
- Every GNU/Linux distribution has its package manager.

Example (Install the Z Shell)

```
apt install zsh           #Debian
dnf install zsh          #Redhat
pacman -S zsh            #Archlinux
zypper install zsh      #SUSE
emerge zsh               #Gentoo
```

Characteristics of Gentoo

- Full GNU/Linux distribution (2000–)
 - ▶ General purpose: from daily to scientific use cases.
 - ▶ Portability: x86, amd64, sparc, arm/64, alpha, mips, riscv, ...
 - ▶ Meta-distribution: ultimate flexibility for specific needs.

Characteristics of Gentoo

- Full GNU/Linux distribution (2000–)
 - ▶ General purpose: from daily to scientific use cases.
 - ▶ Portability: x86, amd64, sparc, arm/64, alpha, mips, riscv, ...
 - ▶ Meta-distribution: ultimate flexibility for specific needs.
- Community driven: spontaneous and distributed development.
 - ▶ Packages, like the build recipe of *Pythia*, are shared pieces of wisdom of the community.

Characteristics of Gentoo

- Full GNU/Linux distribution (2000–)
 - ▶ General purpose: from daily to scientific use cases.
 - ▶ Portability: x86, amd64, sparc, arm/64, alpha, mips, riscv, ...
 - ▶ Meta-distribution: ultimate flexibility for specific needs.
- Community driven: spontaneous and distributed development.
 - ▶ Packages, like the build recipe of Pythia, are shared pieces of wisdom of the community.

Example (`emerge -t pythia`)

```
[ebuild N ] sci-physics/pythia-8.2.26:8::gentoo
[ebuild N ] sci-physics/hepmc-2.06.09-r1::gentoo
[ebuild N ] sci-physics/lhapdf-6.2.3::gentoo
[ebuild N ] dev-libs/boost-1.71.0:0/1.71.0::gentoo
[ebuild U ] dev-util/boost-build-1.71.0::gentoo
[ebuild N ] sci-physics/fastjet-3.0.6-r1::gentoo
[ebuild N ] sci-physics/siscone-3.0.3::gentoo
```

Gentoo Package Manager for Physics


Gentoo Prefix installs a complete Gentoo userspace into a directory, called $\${EPREFIX}$.

- Try it: download and run the bootstrap script¹

```
wget https://goo.gl/czXbjP -O bootstrap.sh
chmod +x bootstrap.sh
./bootstrap.sh
```

- Bootstrap has 3 Stages:

- 1 Compile python and portage in $\${EPREFIX}/tmp$.
- 2 Install gcc in $\${EPREFIX}/tmp$ by portage.
- 3 Build Gentoo in $\${EPREFIX}$ by portage and Stage 2 gcc.



¹<https://wiki.gentoo.org/wiki/Project:Prefix>

Full GNU Userspace of Gentoo Prefix:

```
$ tree -L 1 -d
```

```
/cvmfs/gentoo
```

```
|-- bin  
|-- etc  
|-- lib  
|-- lib64  
|-- run  
|-- sbin  
|-- tmp  
|-- usr  
`-- var
```

- ELF Program Headers INTERP points to dynamic loader in `${EPREFIX}`. Dynamic loader use libraries from `${EPREFIX}/lib`, etc.

```
$ ldd `which cp`  
linux-vdso.so.1 (0x00007ffe245e2000)  
libacl.so.1 => /cvmfs/gentoo/usr/lib64/libacl.so.1 (0x00007f0f7ec29000)  
libc.so.6 => /cvmfs/gentoo/lib64/libc.so.6 (0x00007f0f7ea59000)  
/cvmfs/gentoo/lib64/ld-linux-x86-64.so.2 (0x00007f0f7ec59000)
```

Full GNU Userspace of Gentoo Prefix:

```
$ tree -L 1 -d
```

```
/cvmfs/gentoo
```

```
|-- bin  
|-- etc  
|-- lib  
|-- lib64  
|-- run  
|-- sbin  
|-- tmp  
|-- usr  
`-- var
```

- ELF Program Headers INTERP points to dynamic loader in $\${EPREFIX}$. Dynamic loader use libraries from $\${EPREFIX}/lib$, etc.

```
$ ldd `which cp`  
linux-vdso.so.1 (0x00007ffe245e2000)  
libacl.so.1 => /cvmfs/gentoo/usr/lib64/libacl.so.1 (0x00007f0f7ec29000)  
libc.so.6 => /cvmfs/gentoo/lib64/libc.so.6 (0x00007f0f7ea59000)  
/cvmfs/gentoo/lib64/ld-linux-x86-64.so.2 (0x00007f0f7ec59000)
```

- Gentoo Prefix is a self-contained userspace,
 - ▶ Neither LD_LIBRARY_PATH nor ELF DT_RPATH is needed.
 - ▶ One Prefix runs on all: best distributed via cvmfs to heterogeneous GNU/Linux hosts, like CentOS, Ubuntu.

Full GNU Userspace of Gentoo Prefix:

```
$ tree -L 1 -d
```

```
/cvmfs/gentoo
```

```
|-- bin  
|-- etc  
|-- lib  
|-- lib64  
|-- run  
|-- sbin  
|-- tmp  
|-- usr  
`-- var
```

- ELF Program Headers INTERP points to dynamic loader in $\${EPREFIX}$. Dynamic loader use libraries from $\${EPREFIX}/lib$, etc.

```
$ ldd `which cp`  
linux-vdso.so.1 (0x00007ffe245e2000)  
libacl.so.1 => /cvmfs/gentoo/usr/lib64/libacl.so.1 (0x00007f0f7ec29000)  
libc.so.6 => /cvmfs/gentoo/lib64/libc.so.6 (0x00007f0f7ea59000)  
/cvmfs/gentoo/lib64/ld-linux-x86-64.so.2 (0x00007f0f7ec59000)
```

- Gentoo Prefix is a self-contained userspace,
 - ▶ Neither `LD_LIBRARY_PATH` nor `ELF DT_RPATH` is needed.
 - ▶ One Prefix runs on all: best distributed via `cvmfs` to heterogeneous GNU/Linux hosts, like CentOS, Ubuntu.
- At the same time, it shares the global filesystem view, as opposed to filesystem namespaces/containers.

- autotools-based build systems

```
./configure --prefix="${EPREFIX}" ...
```

Inside Gentoo Prefix

- autotools-based build systems

```
./configure --prefix="${EPREFIX}" ...
```

- CMake-based build systems

```
cmake -DCMAKE_INSTALL_PREFIX="${EPREFIX}" ...
```

Inside Gentoo Prefix

- autotools-based build systems

```
./configure --prefix="${EPREFIX}" ...
```

- CMake-based build systems

```
cmake -DCMAKE_INSTALL_PREFIX="${EPREFIX}" ...
```

- Python, Perl, R, Haskell packages
inherit language interpreter directory prefix.

- autotools-based build systems

```
./configure --prefix="${EPREFIX}" ...
```

- CMake-based build systems

```
cmake -DCMAKE_INSTALL_PREFIX="${EPREFIX}" ...
```

- Python, Perl, R, Haskell packages
inherit language interpreter directory prefix.
- Toolchain

`gcc`, `binutils` `sysroot=${EPREFIX}`, search for headers and libraries in `EPREFIX`.

`gcc` inject `EPREFIX` dynamic loader.

`glibc` look for configurations in `EPREFIX/etc`.

Reference: <https://goo.gl/wQEkhE>

Interoperability with Other Package Managers

Thanks to the flexibility of portage, other package formats can be straightforwardly converted into ebuilds. A few examples:

R CRAN Gentoo R-Overlay

Python PyPI Gentoo PyPI ebuild generator

Emacs ELPA Gentoo ELPA ebuild generator

Octave Forge Gentoo Octave overlay

Java Maven Gentoo Maven overlay (on-going)

Rust Cargo Gentoo Cargo ebuild generator (on-going)

Do not exist, you are welcomed to try :)

Julia Pkg Gentoo Julia overlay...

cf. <https://github.com/JuliaLang/julia/issues/30528>

Packages in Gentoo

- Big official package repository.
cf. <https://packages.gentoo.org>

```
/var/db/pkg $ ls -1 | grep ^sci
```

```
sci-astronomy
```

```
sci-biology
```

```
sci-calculators
```

```
sci-chemistry
```

```
sci-electronics
```

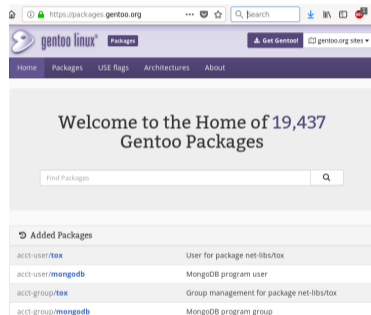
```
sci-geosciences
```

```
sci-mathematics
```

```
sci-physics
```

```
sci-visualization
```

- R Overlay: 18993 packages, e.g. `emerge ggplot2`
- You can do anything on supercomputers:
 - ▶ Install the latest KDE (K Desktop Environment) suite.
 - ▶ Install your own `slurm` by Gentoo to distribute jobs.
- Life on the supercomputers made easier.



g gentoo.org Packages Get Gentoo gentoo.org sites

Home Packages USE flags Architectures About

Welcome to the Home of 19,437 Gentoo Packages

Find Packages

Added Packages

acct-user/tox	User for package net-libs/tox
acct-user/mongodb	MongoDB program user
acct-group/tox	Group management for package net-libs/tox
acct-group/mongodb	MongoDB program group

Case Study 1: Stone-Aged Computing Cluster

- Default environment is RHEL (Redhat Enterprise Linux) 5.

RHEL	Release	End of Security Updates	Lastest Version
5	2007-03-15	2017-03-31	5.11
6	2010-11-10	2020-11-30	6.10
7	2014-06-10	2024-06-30	7.7
8	2019-05-07	2029-05-??	8.0

Case Study 1: Stone-Aged Computing Cluster

- Default environment is RHEL (Redhat Enterprise Linux) 5.

RHEL	Release	End of Security Updates	Lastest Version
5	2007-03-15	2017-03-31	5.11
6	2010-11-10	2020-11-30	6.10
7	2014-06-10	2024-06-30	7.7
8	2019-05-07	2029-05-??	8.0

```
$ cat /etc/redhat-release
Red Hat Enterprise Linux Server release 5.5 (Tikanga)
```

```
$ python -c "import sys; print sys.version"
2.4.3 (#1, Jun 11 2009, 14:09:37)
[GCC 4.1.2 20080704 (Red Hat 4.1.2-44)]
```

```
$ ld --version
GNU ld version 2.17.50.0.6-14.el5 20061020
Copyright 2005 Free Software Foundation, Inc.
```

```
$ ldd --version
ldd (GNU libc) 2.5
Copyright (C) 2006 Free Software Foundation, Inc.
```

Case Study 1: Short Cut to Modern Age

- Gentoo Prefix: no restriction by the host OS:

- ▶ Python

```
$ python -c "import sys;print(sys.version)"  
3.6.6 (default, Aug 31 2018, 03:10:49)  
[GCC 8.2.0]
```

- ▶ Linker

```
$ ld --version  
GNU ld (Gentoo 2.30 p1) 2.30.0  
Copyright (C) 2018 Free Software Foundation, Inc.
```

- ▶ Dynamic loader

```
$ ldd --version  
ldd (Gentoo 2.19-r1 p3) 2.19  
Copyright (C) 2014 Free Software Foundation, Inc.
```

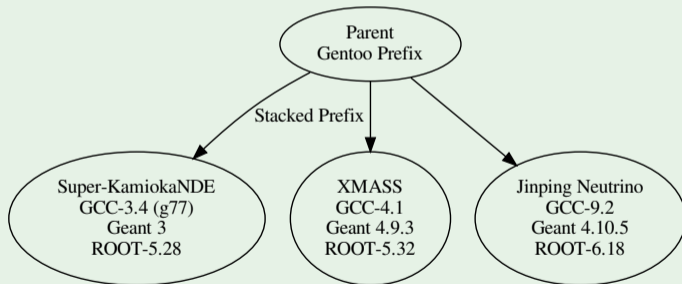
glibc-2.19 is the last glibc supporting linux-2.6.18 of RHEL 5.

Case Study 2: Simultaneous Multiple Versions

- Stacked-Prefix can make thin instances based on a common one.

Example (Legacy of experimental physics)

- Coexistence of software suites to satisfy diverse requirements of varied scenarios.
- Quick deployment of arbitrary testing environments.
- Keeping Filesystem Hierarchy Standard (FHS) without needing containers.

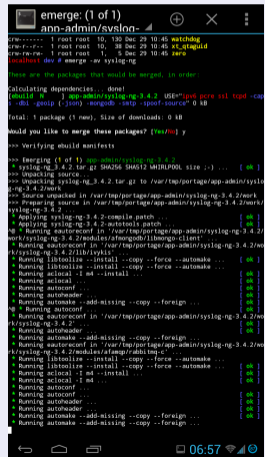


Case Study Bonus

Gentoo on Android

- Be an early adopter of arm64 supercomputing.
- Works on all known rooted android devices.
<http://wiki.gentoo.org/wiki/Project:Android>

Vendor	Device	Android Variant	Version	Tarball	Works?	Reported by
Huawei	P8 Lite 2017 PRA_LX1	stock (EMUI)	5.0.1	20180226	✓ Yes	kavra
Amazon	Kindle Fire (5th Gen)	stock (Fire OS)	5.3.1	20180226	✓ Yes	Tkdestroyer2
Samsung	Galaxy J7 SM-J710F	LineageOS	15.1	20180226	✓ Yes	RomainL
HTC	Nexus 9 LTE (flounder_lte)	LineageOS	14.1	20180225	✓ Yes	hoefling
Xiaomi	Redmi Note	LineageOS	14.1	20180225	✓ Yes	Pieracuti



```
emerge: (1 of 1)
ann-admin/syslog-ng
emerge: [1] root root 10. 130 Dec 29 10:45 watchdog
emerge: [2] root root 10. 38 Dec 29 10:45 kt_qtauid
emerge: [3] root root 1. 5 Dec 29 10:45 zero
emerge: [4] emerge -av syslog-ng

These are the packages that would be merged, in order:

Calculating dependencies... done!
[ebuild] * app-admin/syslog-ng-3.4.2 USE="apk core ssl tclap -cap
+ -dbi -gsoap (-jison) -mongodb -netp -spoof -source" 0 kB
Total: 1 package (1 new), Size of downloads: 0 kB
Would you like to merge these packages? [Y/n] y

*** Verifying ebuild manifests

*** Emerging (1 of 1) app-admin/syslog-ng-3.4.2
* syslog-ng-3.4.2 tar gz SHA256 SHA512 WHIRLPOOL size (-) ... [ ok ]
*** unpacking source ...
*** Unpacking syslog-ng-3.4.2 tar.gz to /var/tmp/portage/app-admin/syslog-ng-3.4.2/work
*** Source unpacked in /var/tmp/portage/app-admin/syslog-ng-3.4.2/work
*** Preparing source in /var/tmp/portage/app-admin/syslog-ng-3.4.2/work/app-admin/syslog-ng-3.4.2
*** Applying syslog-ng-3.4.2-compile patch ... [ ok ]
*** Applying syslog-ng-3.4.2-subtools patch ... [ ok ]
* Running autoreconf in /var/tmp/portage/app-admin/syslog-ng-3.4.2/work/syslog-ng-3.4.2/modules/atmospod/libmospod-client
* Running autoreconf in /var/tmp/portage/app-admin/syslog-ng-3.4.2/work/syslog-ng-3.4.2/lib/eykis
* Running libtoolize --install --copy --force --automake ... [ ok ]
* Running libtoolize --install --copy --force --automake ... [ ok ]
* Running aclocal -I md --install ... [ ok ]
* Running automake ... [ ok ]
* Running automake --add-missing --copy --foreign ... [ ok ]
* Running autoconf in /var/tmp/portage/app-admin/syslog-ng-3.4.2/work/syslog-ng-3.4.2
* Running autoreconf ... [ ok ]
* Running automake --add-missing --copy --foreign ... [ ok ]
* Running autoreconf in /var/tmp/portage/app-admin/syslog-ng-3.4.2/work/syslog-ng-3.4.2/modules/atmospod/libmospod-client
* Running libtoolize --install --copy --force --automake ... [ ok ]
* Running libtoolize --install --copy --force --automake ... [ ok ]
* Running aclocal -I md --install ... [ ok ]
* Running automake ... [ ok ]
* Running autoconf ... [ ok ]
* Running automake --add-missing --copy --foreign ... [ ok ]
* Running automake --add-missing --copy --foreign ... [ ok ]
```

Gentoo on Android

- Be an early adopter of arm64 supercomputing.

Example (CERN)

- Managed as LCG releases

```
$ /cvmfs/sft.cern.ch/lcg/contrib/gentoo/startprefix
```

Example (JUNO@IHEP)

- On lxslc7 login nodes

```
$ /workfs/juno/xubd/gentoo/startprefix
```

- inode quota is not enough, and it is mounted by squashfs, unionfs and FUSE.

Conclusion

- Gentoo (1999–) has been the paradise for geek users and developers for decades.
 - ▶ From the community, but share the mindset of scientists.
- Gentoo portage is a feasible package manager for high energy physics.
 - ▶ It meets all the needs of us, and is especially good on cvmfs.
- The physics use case of Gentoo is a **natural consequence** of its flexibility and expertise in building operating systems.
 - ▶ Sustainable free ride on the developments outside science.

More information

- Homepage: <http://hep.tsinghua.edu.cn/~orv>
<heroxbd AT gentoo.com>
<orv AT tsinghua.edu.cn>
- ① G. Amadio and Benda Xu, *Portage: Bringing Hackers' Wisdom to Science*, <https://arxiv.org/abs/1610.02742>
- ② <https://wiki.gentoo.org/wiki/Project:Prefix>



Ebuild Example: ROOT

- Gentoo Ebuilds are written in carefully organized bash.
- <https://packages.gentoo.org/packages/sci-physics/root>
- <https://gitweb.gentoo.org/repo/gentoo.git/tree/sci-physics/root/root-6.12.06-r11.ebuild>

```
# Copyright 1999-2019 Gentoo Authors
# Distributed under the terms of the GNU General Public License v2

EAPI=6

# ninja does not work due to fortran
CMAKE_MAKEFILE_GENERATOR=emake
FORTRAN_NEEDED="fortran"
PYTHON_COMPAT=( python2_7 python3_{5,6} )

inherit cmake-utils eapi7-ver elisp-common eutils fortran-2 \
        prefix python-single-r1 toolchain-funcs

DESCRIPTION="C++ data analysis framework and interpreter from CERN"
HOMEPAGE="https://root.cern"
SRC_URI="https://root.cern/download/${PN}_v${PV}.source.tar.gz"
```