

Nix and LHCb

Chris Burr

University of Manchester

21st March 2018



- PhD student at Manchester working on LHCb
- Supposed to be working on analysis and velo alignment
- Generally interested in computing
- Before starting this work I tried a few things
 - Including packaging ROOT/XRootD with conda

Why this started?

- This started with analysis preservation in mind
 - Post-DaVinci environments can be tricky to share/preserve
 - Docker is great, but can't be used in most places
 - Must be something better
- Looked at various options, settled on Nix
- Nix could be more generally useful everywhere

What is Nix?

- Nix is a “purely functional package manager”
 - Works with Linux and macOS
 - Can be used alongside other package managers
 - There is also a Linux distribution, NixOS



- Nix is a “purely functional package manager”
- Source-based
 - *Binary caches* can be used to avoid compiling everything



- Nix is a “purely functional package manager”
- Source-based
- Packages are built from *Nix expressions*
 - Typically $\mathcal{O}(10)$ lines long
 - Defined using a custom functional language
 - $\sim 14,000$ package definitions available in nixpkgs



- Nix is a “purely functional package manager”
- Source-based
- Packages are built from *Nix expressions*
- Builds aim to be portable, reproducible and deterministic

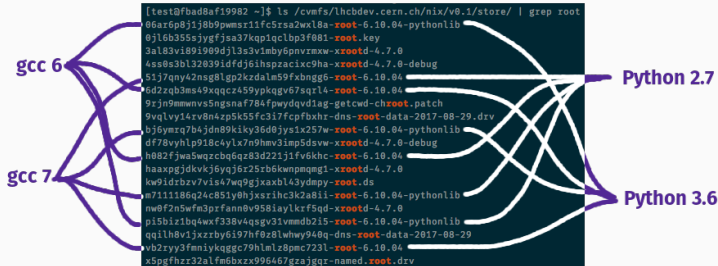


- Nix is a “purely functional package manager”
- Source-based
- Packages are built from *Nix expressions*
- Builds aim to be portable, reproducible and deterministic
- Lots more features available when used fully
 - NixOS
 - Single and multi user modes
 - Transactional approach to updates and configuration



How does Nix work?

- Everything is stored in `/nix/` (by default)
- Packages are kept in `/nix/store`
- Each package lives in a directory named by hash of it's dependencies
 - `gcc6`: `/nix/store/6d2zqb3ms49xqqcz459ypkqgv67sqr14-root-6.10.04/`
 - `gcc7`: `/nix/store/h082fjwa5wqzcbq6qz83d221j1fv6khc-root-6.10.04/`
- Optionally packages can have multiple outputs
 - `bin`, `lib`, `python-lib`, ...



- A collection of nix expressions is known as a *channel*
- Nixpkgs is the most common: <https://github.com/NixOS/nixpkgs>¹

¹There are also release channels at: <https://github.com/NixOS/nixpkgs-channels>

Nix expressions

```
[test@headbf19002 ~]$ readelf -d /cvmfslhcbdev.cern.ch/nix/v0.1/store/6d2zqb3ms49xqcc459ypkgv67sqr14-root-6.10.04/lib/libPyROOT.so

Dynamic section at offset 0xafc00 contains 39 entries:
  Tag             Type                             Name/Value
0x0000000000000001 (NEEDED)     Shared library: [libTree.so]
0x0000000000000001 (NEEDED)     Shared library: [libMathCore.so]
0x0000000000000001 (NEEDED)     Shared library: [libRint.so]
0x0000000000000001 (NEEDED)     Shared library: [libpython3.6m.so.1.0]
0x0000000000000001 (NEEDED)     Shared library: [libTet.so]
0x0000000000000001 (NEEDED)     Shared library: [libRIO.so]
0x0000000000000001 (NEEDED)     Shared library: [libTet.so]
0x0000000000000001 (NEEDED)     Shared library: [libThread.so]
0x0000000000000001 (NEEDED)     Shared library: [libCore.so]
0x0000000000000001 (NEEDED)     Shared library: [libtdc++.so.6]
0x0000000000000001 (NEEDED)     Shared library: [libm.so.6]
0x0000000000000001 (NEEDED)     Shared library: [libgcc_s.so.1]
0x0000000000000001 (NEEDED)     Shared library: [libpthread.so.0]
0x0000000000000001 (NEEDED)     Shared library: [libc.so.6]
0x0000000000000001 (NEEDED)     Shared library: [ld-linux-x86-64.so.2]
0x0000000000000000 (SONAME)      Library soname: [libPyROOT.so]
0x000000000000001d (RUNPATH)     Library runpath: [/cvmfslhcbdev.cern.ch/nix/v0.1/store/6d2zqb3ms49xqcc459ypkgv67sqr14-root-6.10.04/lib
/cvmfslhcbdev.cern.ch/nix/v0.1/store/6d09114wz2bnobgs7J5E2f2symdw-py8m3-3.6.2/lib/cvmfslhcbdev.cern.ch/nix/v0.1/store/hdck6ax230m2605
#f8f93jrf15qica-glibc-2.25/lib/cvmfslhcbdev.cern.ch/nix/v0.1/store/674bxb3fwp3rwrk0hax4zbxwax7ik-gfortran-3.4.8-lib/lib]
```

- A collection of nix expressions is known as a *channel*
- Nixpkgs is the most common: <https://github.com/NixOS/nixpkgs>¹
- Nixpkgs also provides helper functions
 - `buildEnv`: Makes a meta package of symlinks
 - `fetchurl`/`fetchgit`/`fetchpatch`/`fetchchips`/`fetchipfs`
 - `stdenv.mkDerivation`
 - Uses the standard environment to run a `genericBuild`
 - Sets up linker flags and `RUNPATH`
 - Rewrites the interpreter paths of shell scripts to `/nix/store/...`
 - Also uses test suites for many packages

¹There are also release channels at: <https://github.com/NixOS/nixpkgs-channels>

A Nix expression for Gaudi

```
1  { stdenv, fetchurl, fetchpatch, boost, clhep, cmake, cppunit, gperftools
2  , heppdt, jemalloc, libunwind, python, tbb, utillinux, xercesc, zlib
3  , ninja, root, gdb, aida, gsl, libpng }:
4
5  stdenv.mkDerivation rec {
6    name = "gaudi-${version}";
7    version = "v29r0";
8
9    src = fetchurl {
10     url = "https://gitlab.cern.ch/gaudi/Gaudi/repository/${version}/archive.tar.gz";
11     sha256 = "1ijdq1l8rscwij9hgyzrlvga1qg7b0csx76wcd76x3yli8bc766b";
12   };
13
14   buildInputs = [
15     cmake python gdb aida ninja root boost clhep cppunit gperftools heppdt
16     jemalloc libunwind tbb utillinux xercesc zlib gsl libpng
17   ];
18
19   patches = [ ./fix-profiling.patch ];
20
21   cmakeFlags = [
22     "-GNinja"
23   ];
24
25   enableParallelBuilding = true;
26
27   meta = {
28     homepage = https://gaudi.web.cern.ch/gaudi/;
29     description = "A basis for HEP experiment frameworks";
30     platforms = stdenv.lib.platforms.unix;
31     maintainers = with stdenv.lib.maintainers; [ chrisburr ];
32   };
33 }
```

What have I done?

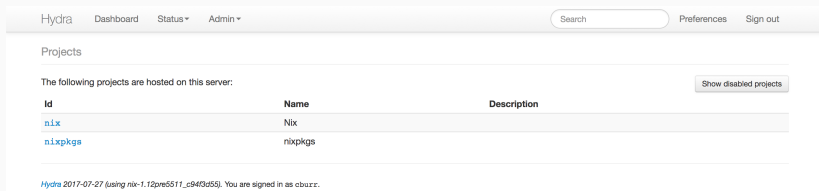
- Installed Nix inside docker without cvmfs mounted
- Built Nix changing `/nix/` to `/cvmfs/lhcbdev.cern.ch/nix/`

- Installed Nix inside docker without cvmfs mounted
- Built Nix changing `/nix/` to `/cvmfs/lhcbdev.cern.ch/nix/`
- Built Nix again....

- Installed Nix inside docker without cvmfs mounted
- Built Nix changing `/nix/` to `/cvmfs/lhcbdev.cern.ch/nix/`
- Built Nix again....
- And it works!!!!
- But the official binary cache can't be used anymore...

- Installed Nix inside docker without cvmfs mounted
- Built Nix changing `/nix/` to `/cvmfs/lhcbdev.cern.ch/nix/`
- Built Nix again....
- And it works!!!!
- But the official binary cache can't be used anymore...
- Have since created a gitlab group: <https://gitlab.cern.ch/lhcb-nix/>
 - **bootstrap**: Use GitLab CI to build nix with a custom store directory
 - Also contains forks of `hydra`, `nix` and `nixpkgs`

Hydra build “farm”



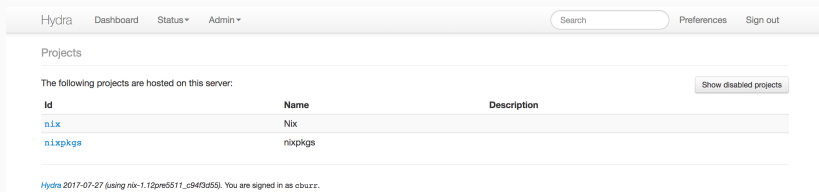
The screenshot shows the Hydra web interface. At the top, there is a navigation bar with 'Hydra', 'Dashboard', 'Status', and 'Admin' menus. A search bar is on the right, along with 'Preferences' and 'Sign out' links. Below the navigation bar, the page title is 'Projects'. A message states 'The following projects are hosted on this server:' followed by a 'Show disabled projects' button. A table lists the projects:

Id	Name	Description
nix	Nix	
nixpkgs	nixpkgs	

At the bottom of the interface, a status message reads: 'Hydra 2017-07-27 (using nix-1.12pre5511_c94f3d55). You are signed in as cbaurr.'

- Set up an instance of Hydra² on openstack: <http://lhcb-hydra.cern.ch:3000/>
- Took less than an hour to get my first build Including setting up PostgreSQL!
 - Uses the local machine for builds
- Since moved to using DBoD and GitLab CI to build a container

²<https://nixos.org/hydra/>



The screenshot shows the Hydra web interface. At the top, there is a navigation bar with 'Hydra', 'Dashboard', 'Status', and 'Admin' menus. A search bar is on the right, along with 'Preferences' and 'Sign out' links. Below the navigation bar, the page title is 'Projects'. A message states 'The following projects are hosted on this server:' followed by a 'Show disabled projects' button. A table lists the projects:

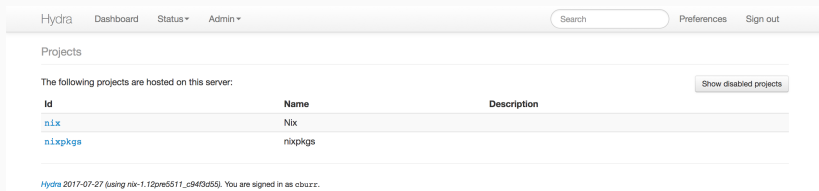
Id	Name	Description
nix	Nix	
nixpkgs	nixpkgs	

At the bottom of the interface, a status message reads: 'Hydra 2017-07-27 (using nix-1.12pre5511_c94f3d55). You are signed in as cbaurr.'

- Set up an instance of Hydra² on openstack: <http://lhcb-hydra.cern.ch:3000/>
- Took less than an hour to get my first build Including setting up PostgreSQL!
 - Uses the local machine for builds
- Since moved to using DBoD and GitLab CI to build a container
- Setting up an extra worker was trivial
 - Just need to be able to SSH to a machine with Nix
 - Docker container on `lb1hcbpr3` with my build of Nix installed

²<https://nixos.org/hydra/>

Hydra build “farm”



The screenshot shows the Hydra web interface. At the top, there is a navigation bar with 'Hydra', 'Dashboard', 'Status', and 'Admin' menus. A search bar is on the right, along with 'Preferences' and 'Sign out' links. Below the navigation bar, the page title is 'Projects'. A message states 'The following projects are hosted on this server:' followed by a table of projects. A 'Show disabled projects' button is located to the right of the table. At the bottom of the page, a status message reads: 'Hydra 2017-07-27 (using nix-1.12pre5511_c94f3d55). You are signed in as cbrucr.'

Id	Name	Description
nix	Nix	
nixpkgs	nixpkgs	

- Set up an instance of Hydra² on openstack: <http://lhcb-hydra.cern.ch:3000/>
- Took less than an hour to get my first build Including setting up PostgreSQL!
 - Uses the local machine for builds
- Since moved to using DBoD and GitLab CI to build a container
- Setting up an extra worker was trivial
 - Just need to be able to SSH to a machine with Nix
 - Docker container on `lb1hcbpr3` with my build of Nix installed
- Support for slaves with different architectures or extra features (AVX?)

²<https://nixos.org/hydra/>

- nixpkgs has a concept of overlays that are applied the main nixpkgs

lncb-software.nix 69 Bytes

```
1 { };
2 self: super:
3
4 {
5   gaudi = super.callPackage ../pkgs/gaudi {};
6 }
```

- nixpkgs has a concept of overlays that are applied the main nixpkgs

lhc-b-software.nix 69 Bytes

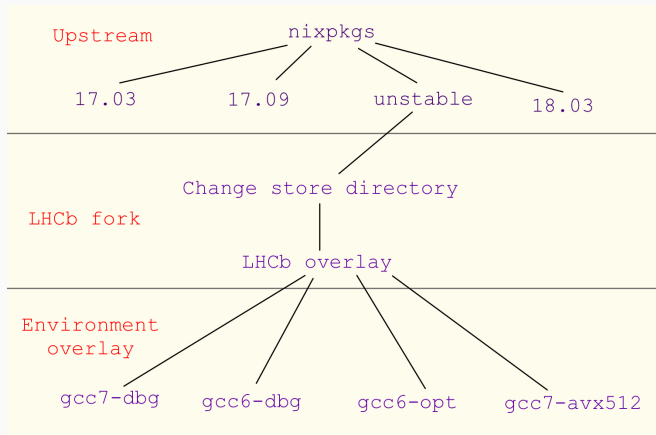
```
1 { };
2 self: super:
3
4 {
5   gaudi = super.callPackage ../pkgs/gaudi {};
6 }
```

- Can also override existing packages or package arguments

gcc-6.nix 235 Bytes

```
1 { };
2 self: super:
3
4 {
5   qt5 = super.qt59;
6   libsForQt5 = super.libsForQt59;
7   gcc = super.gcc6;
8   # Some things really need gcc7
9   aws-sdk-cpp = super.aws-sdk-cpp.override {
10     stdenv = super.overrideCC super.stdenv super.gcc7;
11   };
12 }
```

- Use first overlay to add packages that are unsuitable for upstream
- Second overlay is an argument to nixpkgs to set the environment



Creating environments

- nixpkgs can be used to create environments using `buildEnv`
 - Symlinked to the store directory, similar to an LCG view
- To give a short but comprehensive example:

```
1 { nixpkgs ? builtins.fetchGit { url = https://gitlab.cern.ch/lhcb-nix/nixpkgs.git; ref = "master-lhcb"; }
2   , name ? "user_environment"
3   , extraOverlayPath ? "gcc-7.nix"
4   , extra_packages ? []
5 };
6
7 with import <nixpkgs> { inherit extraOverlayPath; };
8
9 let
10   user_environment = (buildEnv {
11     name = name;
12     paths = (builtins.concatLists [
13       [
14         pkgs.coreutils
15         pkgs.bash
16         pkgs.gcc
17       ]
18       # Python 2
19       (pkgs.root.override { python = python27; }).pythonLib
20       pkgs.python27Full.withPackages(ps: [
21         ps.matplotlib
22         ps.pandas
23       ])
24       # Python 3
25       (pkgs.root.override { python = python36; }).pythonLib
26       pkgs.python36Full.withPackages(ps: [
27         ps.matplotlib
28         ps.pandas
29         ps.snakemake
30       ])
31     ])
32   };
33
34   (builtins.map (s: pkgs.${s}) extra_packages)
35 ];
36 };
37 in user_environment
38
```

Creating environments

- You can then define multiple versions with different arguments
- I've created three as an example:

```
1  {nixpkgs}:
2
3  with import nixpkgs {};
4
5  let
6  ....jobs = {
7  .....# Some example environments for now
8  .....example_environment_gcc6 = callPackage ./make_user_environment.nix {
9  .....  inherit nixpkgs;
10 .....  name = "analysis_environment_gcc6";
11 .....  extraOverlayPath = "gcc-6.nix";
12 .....  };
13 .....example_environment_gcc7 = callPackage ./make_user_environment.nix {
14 .....  inherit nixpkgs;
15 .....  name = "analysis_environment_gcc7";
16 .....  extraOverlayPath = "gcc-7.nix";
17 .....  };
18 .....gaudi_environment_gcc7 = callPackage ./make_user_environment.nix {
19 .....  inherit nixpkgs;
20 .....  name = "gaudi_environment_gcc7";
21 .....  extraOverlayPath = "gcc-7.nix";
22 .....  extra_packages = ["gaudi"];
23 .....  };
24
25 .....# Make a replacement nixpkgs channel
26 .....lhcb_nixpkgs = pkgs.releaseTools.channel {
27 .....  constituents = [];
28 .....  name = "lhcb_nixpkgs";
29 .....  src = pkgs.path;
30 .....  isNix05 = false;
31 .....  };
32 .....};
33 in jobs
34
```

- Full example stored at: <https://gitlab.cern.ch/lhcb-nix/lhcb-environments>
- Built in the lhcb-environments project on hydra

The example environments

- As most of the work is done upstream adding packages is easy
- As these examples are designed to replace PATH entirely they contain:
 - Shells: `bash/zsh/tcsh/dash`
 - Standard utilities: `coreutils/man/grep/tar/findutils/rsync/...`
 - Text editors: `nano/vim/neovim/atom`
 - Version control: `git/svn/hg`

The example environments

- As most of the work is done upstream adding packages is easy
- As these examples are designed to replace PATH entirely they contain:
 - Shells: `bash/zsh/tcsh/dash`
 - Standard utilities: `coreutils/man/grep/tar/findutils/rsync/...`
 - Text editors: `nano/vim/neovim/atom`
 - Version control: `git/svn/hg`
 - Building: `gcc/cmake/ninja/boost/libxml2/tbb/gperftools/...`
 - Debugging: `gdb/lldb/valgrind`
 - TexLive 2017

The example environments

- As most of the work is done upstream adding packages is easy
- As these examples are designed to replace PATH entirely they contain:
 - Shells: `bash/zsh/tcsh/dash`
 - Standard utilities: `coreutils/man/grep/tar/findutils/rsync/...`
 - Text editors: `nano/vim/neovim/atom`
 - Version control: `git/svn/hg`
 - Building: `gcc/cmake/ninja/boost/libxml2/tbb/gperftools/...`
 - Debugging: `gdb/lldb/valgrind`
 - TexLive 2017
 - Python 2.7 with `matplotlib/numpy/pandas/nose/jupyter/...`
 - Python 3.6 with `matplotlib/numpy/pandas/snakemake/...`

The example environments

- As most of the work is done upstream adding packages is easy
- As these examples are designed to replace PATH entirely they contain:
 - Shells: `bash/zsh/tcsh/dash`
 - Standard utilities: `coreutils/man/grep/tar/findutils/rsync/...`
 - Text editors: `nano/vim/neovim/atom`
 - Version control: `git/svn/hg`
 - Building: `gcc/cmake/ninja/boost/libxml2/tbb/gperftools/...`
 - Debugging: `gdb/lldb/valgrind`
 - TexLive 2017
 - Python 2.7 with `matplotlib/numpy/pandas/nose/jupyter/...`
 - Python 3.6 with `matplotlib/numpy/pandas/snakemake/...`
 - XRootD with Python 2.7 and 3.6 bindings
 - ROOT* with Python 2.7 and 3.6 bindings

Try it for yourself in docker!

Try it for yourself in docker! (CERN only due to firewall)

- Install Nix:

```
1 docker run --rm -it centos:7 bash
2 useradd test
3 yum install -y bzip2
4 mkdir -p -m 0755 /cvmfs/lhcbdev.cern.ch/nix
5 chown test /cvmfs/lhcbdev.cern.ch/nix
6 cd /home/test
7 su test bash -c "curl -LO https://chrisburr.me/lhcb-nix-2.0/nix-2.0-2018_03_20-x86_64-linux.tar.bz2"
8 su test bash -c "curl https://chrisburr.me/lhcb-nix-2.0/install | sh"
```

- Install one (or more) of the environments in any directory:

example_environment_gcc6/example_environment_gcc7/gaudi_environment_gcc7

```
1 su test
2 . /home/test/.nix-profile/etc/profile.d/nix.sh
3 export LC_ALL=en_US.utf-8
4 export LANG=en_US.utf-8
5
6 mkdir -p "/cvmfs/lhcbdev.cern.ch/nix/environments/"
7 export LHCB_NIX_ENV_DIR="/cvmfs/lhcbdev.cern.ch/nix/environments/analysis_environment_gcc7"
8 nix-env -ir analysis_environment_gcc7 --profile "${LHCB_NIX_ENV_DIR}" -Q -j8
```

- Set PATH and run!

```
1 su test
2 export LHCB_NIX_ENV_DIR="/cvmfs/lhcbdev.cern.ch/nix/environments/analysis_environment_gcc7"
3 export PATH="${LHCB_NIX_ENV_DIR}/bin"
4 export CMAKE_PREFIX_PATH="${LHCB_NIX_ENV_DIR}"
5 export NIX_SSL_CERT_FILE=/etc/ssl/certs/ca-bundle.crt
6 bash
```


- Downloads are slow:
 - The binary cache is currently compressed on the fly by hydra
 - There is a setting to copy them to a directory/AWS/..
 - This can then be hosted on any web server
 - Plus packages are then signed automatically
- Package signatures aren't checked (see above)
- Some packages have issues being built inside docker containers

My thoughts...

- Software built should be able to run on “any” flavour of Linux
 - Example works with CentOS 6, 7 and Ubuntu
 - Darwin should be fairly easy to add
 - Experimental support for AArch64

- Software built should be able to run on “any” flavour of Linux
- Simpler environments
 - No more (ab)use of `LD_LIBRARY_PATH` or `PYTHON_PATH`
 - Software with conflicting dependencies can be used at the same time

- Software built should be able to run on “any” flavour of Linux
- Simpler environments
- Huge number of packages definitions already written ~14,000
 - Adding new software to an environment is a one line change

- Software built should be able to run on “any” flavour of Linux
- Simpler environments
- Huge number of packages definitions already written ~14,000
- Adding new package definitions is straight forward
 - The standard builder already works with most build systems
 - `RUNPATH` and other paths are set automatically
 - Building Gaudi was trivial
(once I had written definitions for all of it's HEP specific dependencies...)
(and fixed a bug? in the CMake config of the profiling module...)

- Software built should be able to run on “any” flavour of Linux
- Simpler environments
- Huge number of packages definitions already written ~14,000
- Adding new package definitions is straight forward
- Active community, lots of very helpful experts on IRC

- Documentation is lacking some places
 - But it's rapidly improving
 - Figuring things out from the source isn't too difficult

What isn't so good?

- Documentation is lacking some places
- The Nix expression language has a steep learning curve
 - I had never used a functional language like Haskell
 - Might have been easier if I had
 - Doesn't matter simple things like writing packages

- Documentation is lacking some places
- The Nix expression language has a steep learning curve
- Independence from the host system isn't perfect
 - I've read about issues with OpenGL/graphics drivers
 - Kernel
 - Can't be worse than what already exists

What isn't so good?

- Documentation is lacking some places
- The Nix expression language has a steep learning curve
- Independence from the host system isn't perfect
- Sometimes reproducible builds aren't reproducible
 - Only seen this happen due to remote files being removed/changed
 - So long as the original nix store is kept there is always a copy

- Nix is awesome!
- I can see a lot of benefits and potential uses
 - Could avoid issues with missing or conflicting dependencies
 - Defining extra environments is easy (per analysis?/distributable?)
 - Can update old environments where needed (XRootD?)
- Useful resources and some other details in backup

- Nix is awesome!
- I can see a lot of benefits and potential uses
 - Could avoid issues with missing or conflicting dependencies
 - Defining extra environments is easy (per analysis?/distributable?)
 - Can update old environments where needed (XRootD?)
- Useful resources and some other details in backup

Any Questions?

Documentation:

- Introduction to Nix: <https://nixos.org/nixos/nix-pills/>
- Nix manual: <https://nixos.org/nix/manual/>
- Nixpkgs manual: <https://nixos.org/nixpkgs/manual/>

- `PYTHON_PATH` isn't ideal as it is used by all Python versions
- `sitecustomize.py` is aimed for this purpose
 - Uses `$LHCB_NIX_ENV_DIR/lib/pythonX.Y/site-packages/`
- `ROOT` can't be built with simultaneous Python 2 and 3 support
 - Instead make the Python library a separate package
 - Each is then loaded from `lib/pythonX.Y/site-packages`
 - Using `TPython` from the `root` REPL uses Python 2

- Stripped and deleted by default
- `stdenv.mkDerivation` has an option `separateDebugInfo`
- Makes a `-debug` package containing `lib/debug/.build-id/XX/YYYY`
- Can be loaded in GDB by modifying `~/.gdbinit` to contain:
 - `set debug-file-directory ENV_DIR/lib/debug`
 - There are probably other methods available