

Incremental Package Builds

Guillaume Maudoux
@layus

NixCon 2017



Louvain-la-Neuve hold a huge bike event



@layus

laïus (🇫🇷) a long, vague and emphatic speech

@layus


laius (🇫🇷) a long, vague and emphatic speech

layus a nixpkgs contributor

@layus

laius (🇫🇷) a long, vague and emphatic speech

layus a nixpkgs contributor



I contributed to build a castle in France



So I...

- ▶ like building stuffs

So I...

- ▶ like building stuffs
- ▶ started a PhD on incremental builds

So I...

- ▶ like building stuffs
- ▶ started a PhD on incremental builds
- ▶ want your feedback on this presentation

Incremental package builds

Incremental

adjective [...] occurring in **especially small increments**.

Incremental

adjective [...] occurring in **especially small increments**.

noun the amount or degree by which something changes;

Incremental build systems

- ▶ Works better with small steps
- ▶ Can reuse older build products
- ▶ Can detect what needs to be built

Plan

Firefox – Works better with small steps


i3 – Can reuse older build products

Nix store – Can detect what needs to be built

Firefox – Small steps

Building Firefox takes very long

50s	unpack	1 min
1s	patch	
156s	configure	2 min 30 s
4006s	build	1 hour 6 min
31s	install	
8s	fixup	



We need checkpoints!



We need checkpoints!



We need checkpoints!



We need checkpoints!



Trivial idea

We could split each package phase in a different derivation...

Trivial idea

We could split each package phase in a different derivation...

- ▶ Only a small, local fix
- ▶ Not clean for `/nix/store`
- ▶ Already done by external wrappers (`firefox`)

Incremental builds is about small steps

- ▶ Nix is incremental at the package level.
- ▶ Nix is a package manager

Incremental builds is about small steps

- ▶ Nix is incremental at the package level.
- ▶ Nix is a package manager

- ▶ Build systems are incremental at command level.
- ▶ let's use that!

Build systems details – make

1. pros:

- ▶ well known

2. cons:

- ▶ requires previous builds as an input
- ▶ uses timestamps to detect changes

Build systems details – ccache

`ccache` memoizes compiler invocations

`sccache` can share it's cache on the network

`nix` is much like `sccache`

Build systems details – nix-make

edolstra / **nix-make**

<> Code

Issues 0

Pull requests 0

Projects 0

Insights

Experimental Nix build management stuff


15 commits


1 branch


0 releases

Branch: **master** ▾

New pull request

 **edolstra** Produce more sensible store path names

 **examples** *`dependencyClosure` now allows a search path, e.g.,

 **lib** Produce more sensible store path names

Build systems details – nix-make

```
let {  
  inherit (import ../../lib) compileC link;  
  
  hello = link {  
    objects = compileC {  
      main = ./hello.c;  
    };  
  };  
  
  body = [hello];  
}
```

Build systems details – nix-make

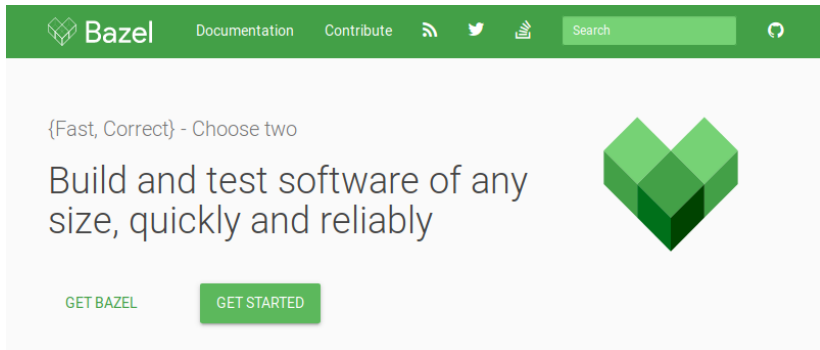
1. pros:

- ▶ very small steps
- ▶ compatible with nix






2. cons:

- ▶ every intermediate .o file ends up in the store
- ▶ requires to port projects to nix-make

Build systems details – bazel



The image shows the top portion of the Bazel website. It features a green navigation bar with the Bazel logo, links for 'Documentation' and 'Contribute', social media icons for RSS, Twitter, and GitHub, a search bar, and a help icon. Below the navigation bar, the main content area has a light gray background. On the left, the text reads '{Fast, Correct} - Choose two' followed by 'Build and test software of any size, quickly and reliably'. On the right is a large 3D Bazel logo. At the bottom left are two buttons: 'GET BAZEL' and 'GET STARTED'.

 **Bazel** Documentation Contribute    Search 

{Fast, Correct} - Choose two

Build and test software of any size, quickly and reliably

[GET BAZEL](#) [GET STARTED](#)

Build systems details – bazel

1. pros:

- ▶ caches arbitrary commands
- ▶ uses sandboxing to guarantee correctness

2. cons:

- ▶ conflicts with nix

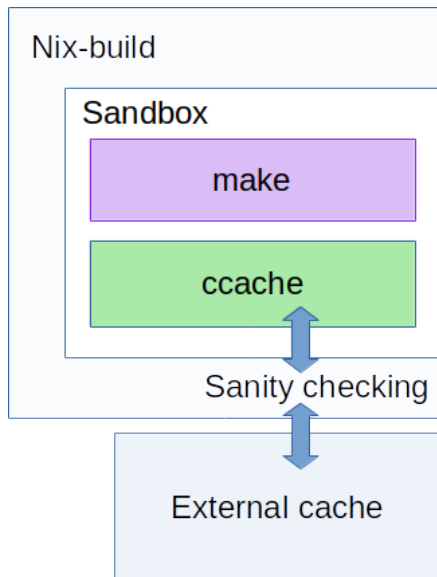
`bazel = nix + build system`

Intermediate solution

- ▶ Use caching in the build system
- ▶ Control caching from nix-build

Easy to implement, if we trust the build system.

Allow caching inside nix-build



TO DO LIST:

1. WAIT FOR TONIGHT
2. TRY TO TAKE OVER THE WORLD!



still a work in progress

i3 – Reusing old builds

The issue

Requirements

- ▶ Need to patch i3 for an annoying bug
- ▶ Can only be tested on this machine
- ▶ Needs to be included in NixOS config

Typical debug session

- ▶ build a custom version (`nix-shell`)
- ▶ write an overlay for that package (`nix-build`)
- ▶ modify nixos to use it, and use debug flags (`nixos-rebuild`)
- ▶ test and restart

Current options

We need an easy solution to achieve simple, local package development:

Current options

We need an easy solution to achieve simple, local package development:

1. Mount the nix store RW

Current options

We need an easy solution to achieve simple, local package development:

1. Mount the nix store RW
2. Insert a symlink in the store

Current options

We need an easy solution to achieve simple, local package development:

1. Mount the nix store RW
2. Insert a symlink in the store
3. Configure services with paths outside the store

Current options

We need an easy solution to achieve simple, local package development:

1. Mount the nix store RW
2. Insert a symlink in the store
3. Configure services with paths outside the store

Current options

We need an easy solution to achieve simple, local package development:

1. Mount the nix store RW
2. Insert a symlink in the store
3. Configure services with paths outside the store

Hacky !!

Existing tools

`git rebase --interactive` brings you right at the conflicting commit, in a fake environment

`nix-shell -A` setups the right environment to build a package. But you *can not* write to the store

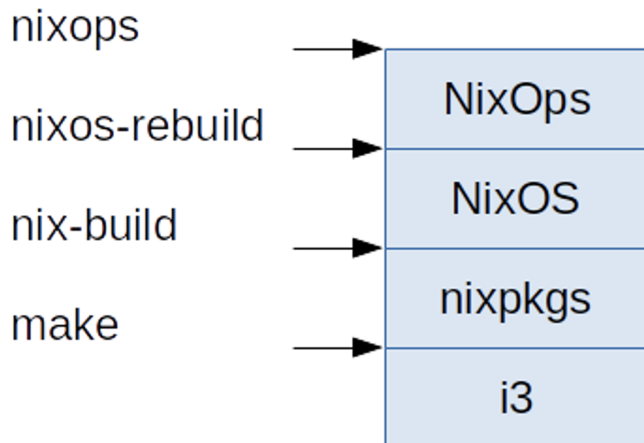
Solving the nix-shell issue

To make nix-shell more handy, it could generate a random \$out on each invocation, and allow writes to that store location.

We want a trade-off between correctness and ease of use.

```
nix build --hack i3
```


Nix* “OSI” model



nix-shell is a zipper

nix-shell

make



Going further

```
nixos-rebuild test --hack i3
```

Drop me in a shell where I can patch i3, then use that for this nixos version.

And even further

We can even use caching when implemented! A “clean” `nix-build` could reuse results cached by a “hacky” build.

And even further

We can even use caching when implemented! A “clean” `nix-build` could reuse results cached by a “hacky” build.

TL;DR: Never compile the same build step twice.

And even further

We can even use caching when implemented! A “clean” `nix-build` could reuse results cached by a “hacky” build.

TL;DR: Never compile the same build step twice.

Well, *if* implemented.

Nix store – Track changes


/nix/store has many similar packages

```
diffoscope /nix/store/4w573p7v5jjjkph5ndzmbwg55v2ids1y-poppler-data-0.4.7 /nix/s
--- /nix/store/4w573p7v5jjjkph5ndzmbwg55v2ids1y-poppler-data-0.4.7
+++ /nix/store/ccc9py7hfgd5v4q7hk2fv3l4rjllh12i-poppler-data-0.4.7
lib
  pkgconfig
    poppler-data.pc
    @@ -1,8 +1,8 @@
  -poppler_datadir=/nix/store/4w573p7v5jjjkph5ndzmbwg55v2ids1y-poppler-data-0.
  +poppler_datadir=/nix/store/ccc9py7hfgd5v4q7hk2fv3l4rjllh12i-poppler-data-0.

  Name: poppler-data
  Description: Encoding files for use with poppler
  Version: 0.4.7

  -Cflags: -DPOPPLER_DATADIR=/nix/store/4w573p7v5jjjkph5ndzmbwg55v2ids1y-poppl
  +Cflags: -DPOPPLER_DATADIR=/nix/store/ccc9py7hfgd5v4q7hk2fv3l4rjllh12i-poppl
```


/nix/store has many similar packages



Nix store --optimize

```
$ du -shc /nix/store/{X,Y,Z}-poppler-data-0.4.7
12M  /nix/store/4w573p7v5jjjkph5ndzmbwg55v2ids1y-poppler-data-0
60K  /nix/store/ccc9py7hfgd5v4q7hk2fv3l4rjllh12i-poppler-data-0
60K  /nix/store/pshrgbbmvkxp6lf4hzwn04560brf52lp-poppler-data-0
13M  total
```

Nix store --optimize



Mass rebuilds use a lot of network.


Every week, unstable branch is merged into master.
Every week, nixpkgs master receives mass-rebuild commits.

Updating once per month makes you download a full new distro each time.

Can we do better than that ?

Binary diffs of substitutes

By storing and sharing diffs of binary packages, we could save bandwidth and hydra space.



Content addressed storage

We can even go further for derivations whose only difference is their \$out.

NixOS / rfcs Watch 35

[Code](#) [Pull requests 11](#) [Projects 0](#) [Insights](#)

Intensional Store #17

[Open](#) **wmertens** wants to merge 2 commits into `nixos:master` from `wmertens:master`

[Conversation 12](#) [Commits 2](#) [Files changed 1](#)



wmertens commented on 11 Aug

A practical proposal to store build products under their output hash instead of their input hash.

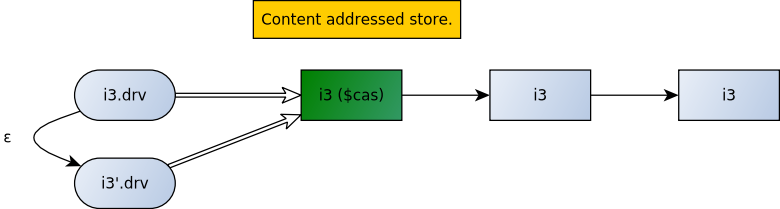


4





1

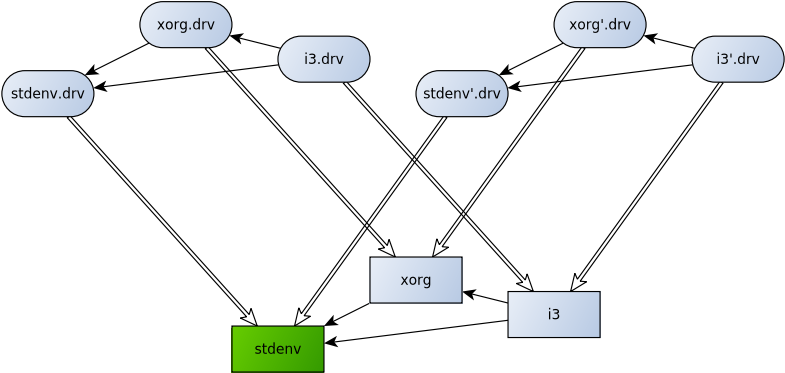
Content addressed storage



CAS & change propagation



CAS & change propagation



Content addressed storage

1. Pros:

- ▶ Does not propagate changes to dependencies
- ▶ Less compiling
- ▶ Faster updates
- ▶ Not too difficult to implement
- ▶ Outsourcing reproducibility is easy

2. Cons:

- ▶ Changes Nix
- ▶ Real impact is unknown

CAS & refactorings

This screenshot shows a GitHub pull request page for the repository `NixOS/nixpkgs`. The page title is "Just strip everything by default, fully tested version :-)" with the PR number `#15339`. The pull request is from branch `layus:fix-strip` to `nixos:staging`. It shows 4 commits and 11 files changed, with a net change of +125 lines and -37 lines. A comment from `layus` dated 10 May 2016 explains that this PR is a fully tested and modified version of PR `#15087`, taking into account discovered bugs. The comment notes that `Stdenv` now supports the `dontStripPath` attribute, which is a black-list version of the former `stripDebugList` and `stripAllList` (maintained for compatibility). The right sidebar shows no reviews and one assignee, `vcunat`.

This repository Search Pull requests Issues Marketplace Explore

NixOS / nixpkgs Unwatch 141 Star 1,967 Fork 2,306

Code Issues 1,987 Pull requests 498 Projects 9 Insights

Just strip everything by default, fully tested version :-)

#15339

Open layus wants to merge 4 commits into nixos:staging from layus:fix-strip

Conversation 35 Commits 4 Files changed 11 +125 -37

layus commented on 10 May 2016 • edited Contributor

This is the same PR as #15087, but fully tested and modified to take into account the discovered bugs.

Stdenv now supports the `dontStripPath` attribute which is the black-list version of the former `stripDebugList` and `stripAllList` (maintained for compatibility).

Reviewers
No reviews

Assignees
vcunat

CAS & Cached builds

Caching builds can make them less stable
(reproducible)

Content addressed store would help to catch
unstable builds

Conclusion

There are still a lot of possible improvements to nix.

I have started to work on cached builds, and CAS is under RFC.

Comments welcome.

Questions ?