

# Tests too slow? Rethink your CI!

---

Simon Stewart

Ex Selenium Project Lead

Build Nerd

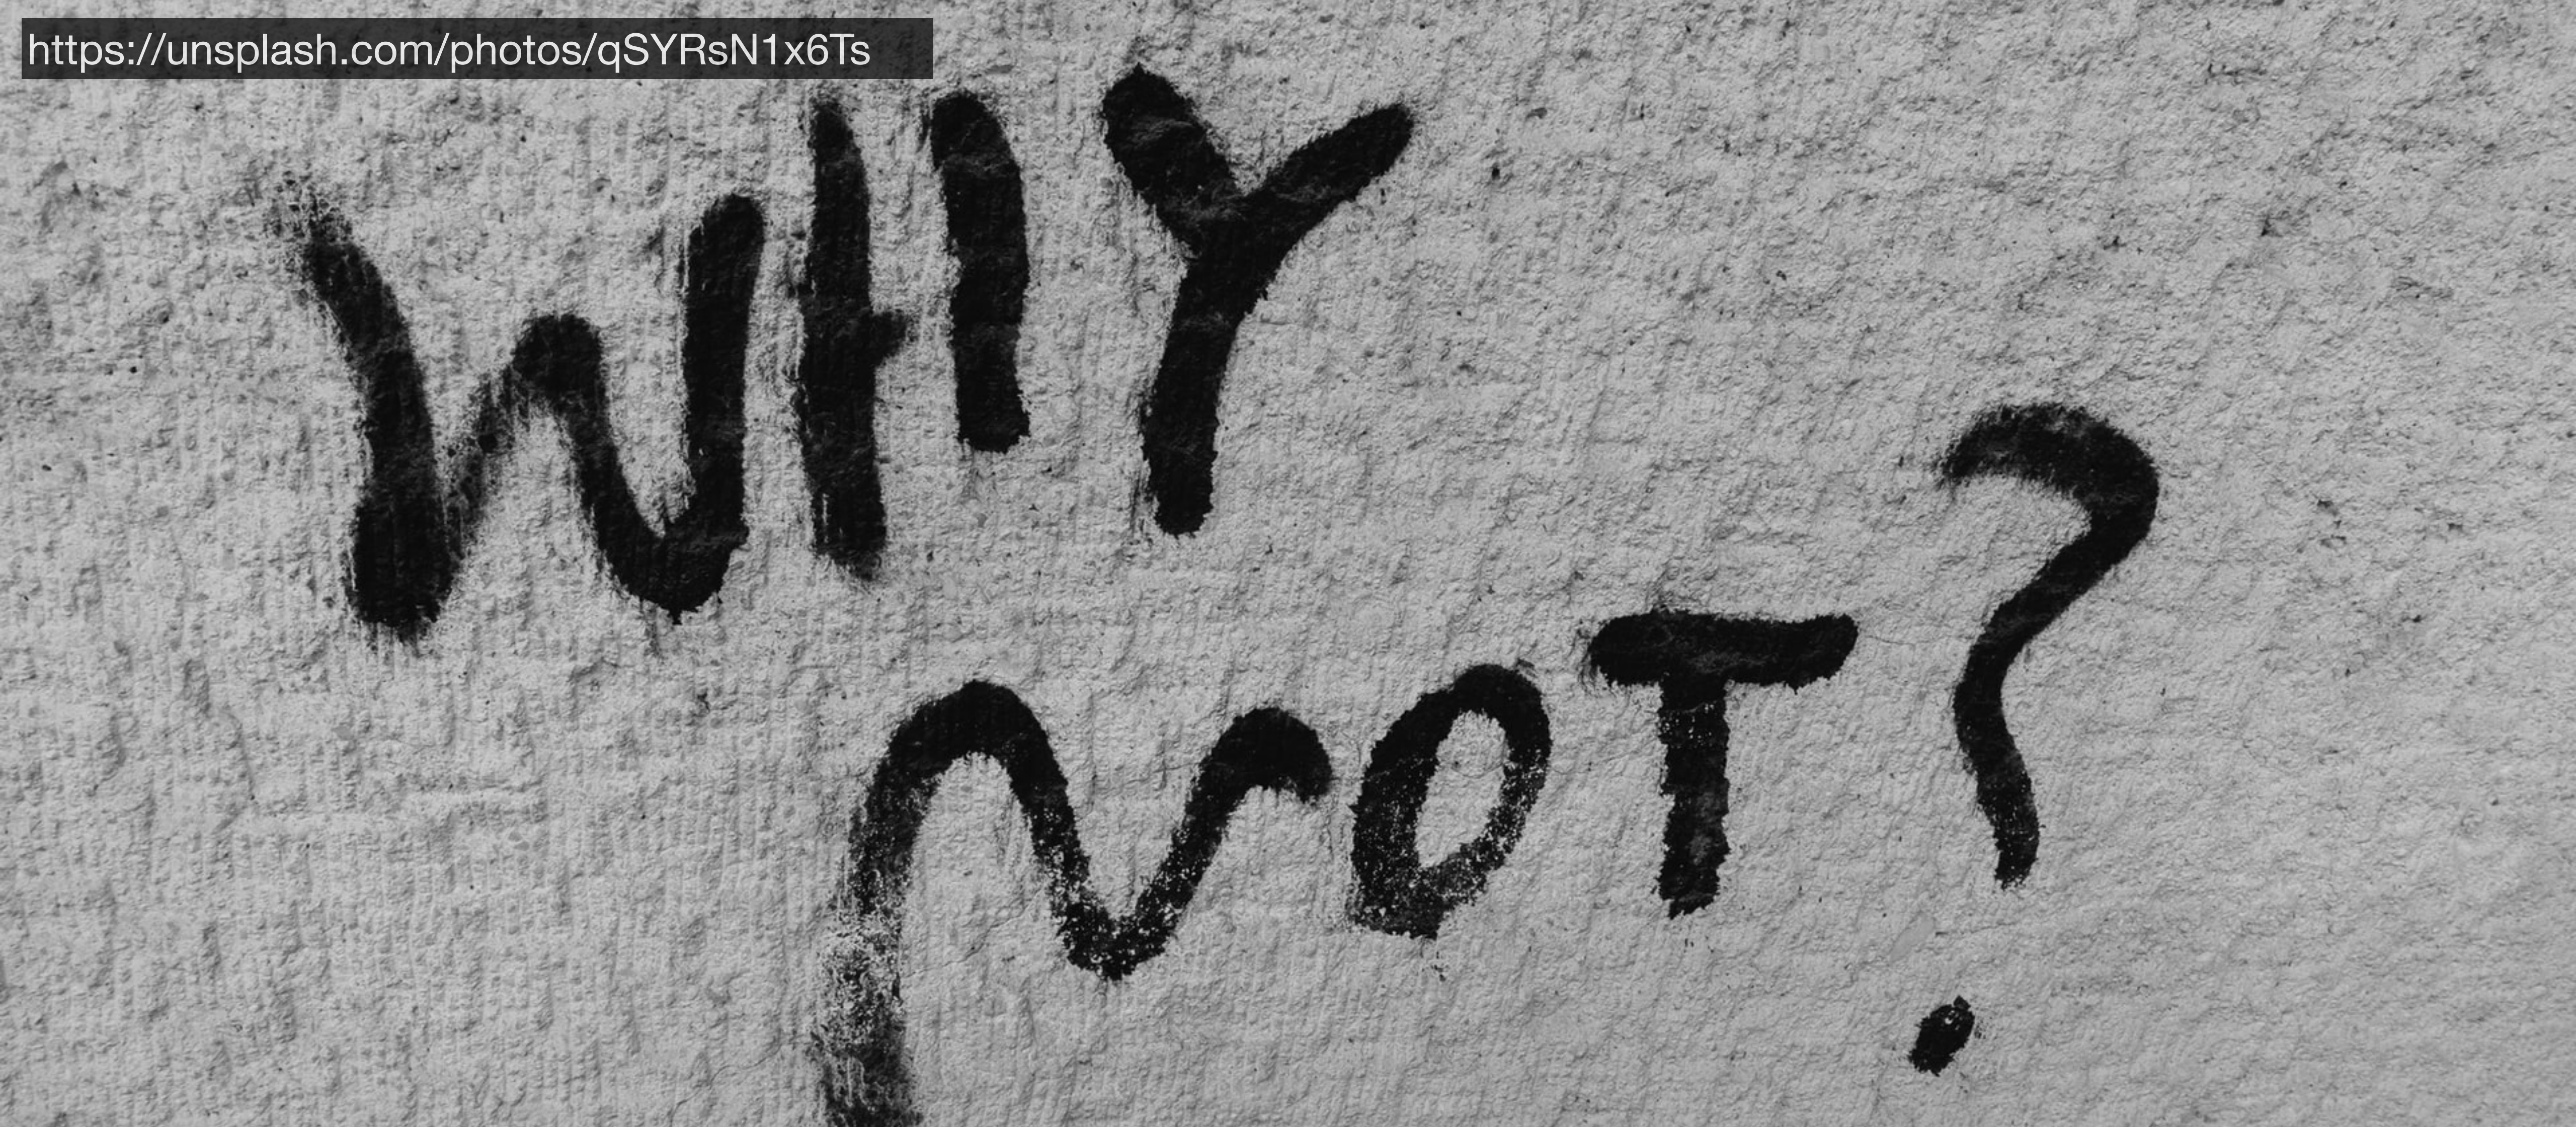
@shs96c on the socials

# General ideas, specific example

---

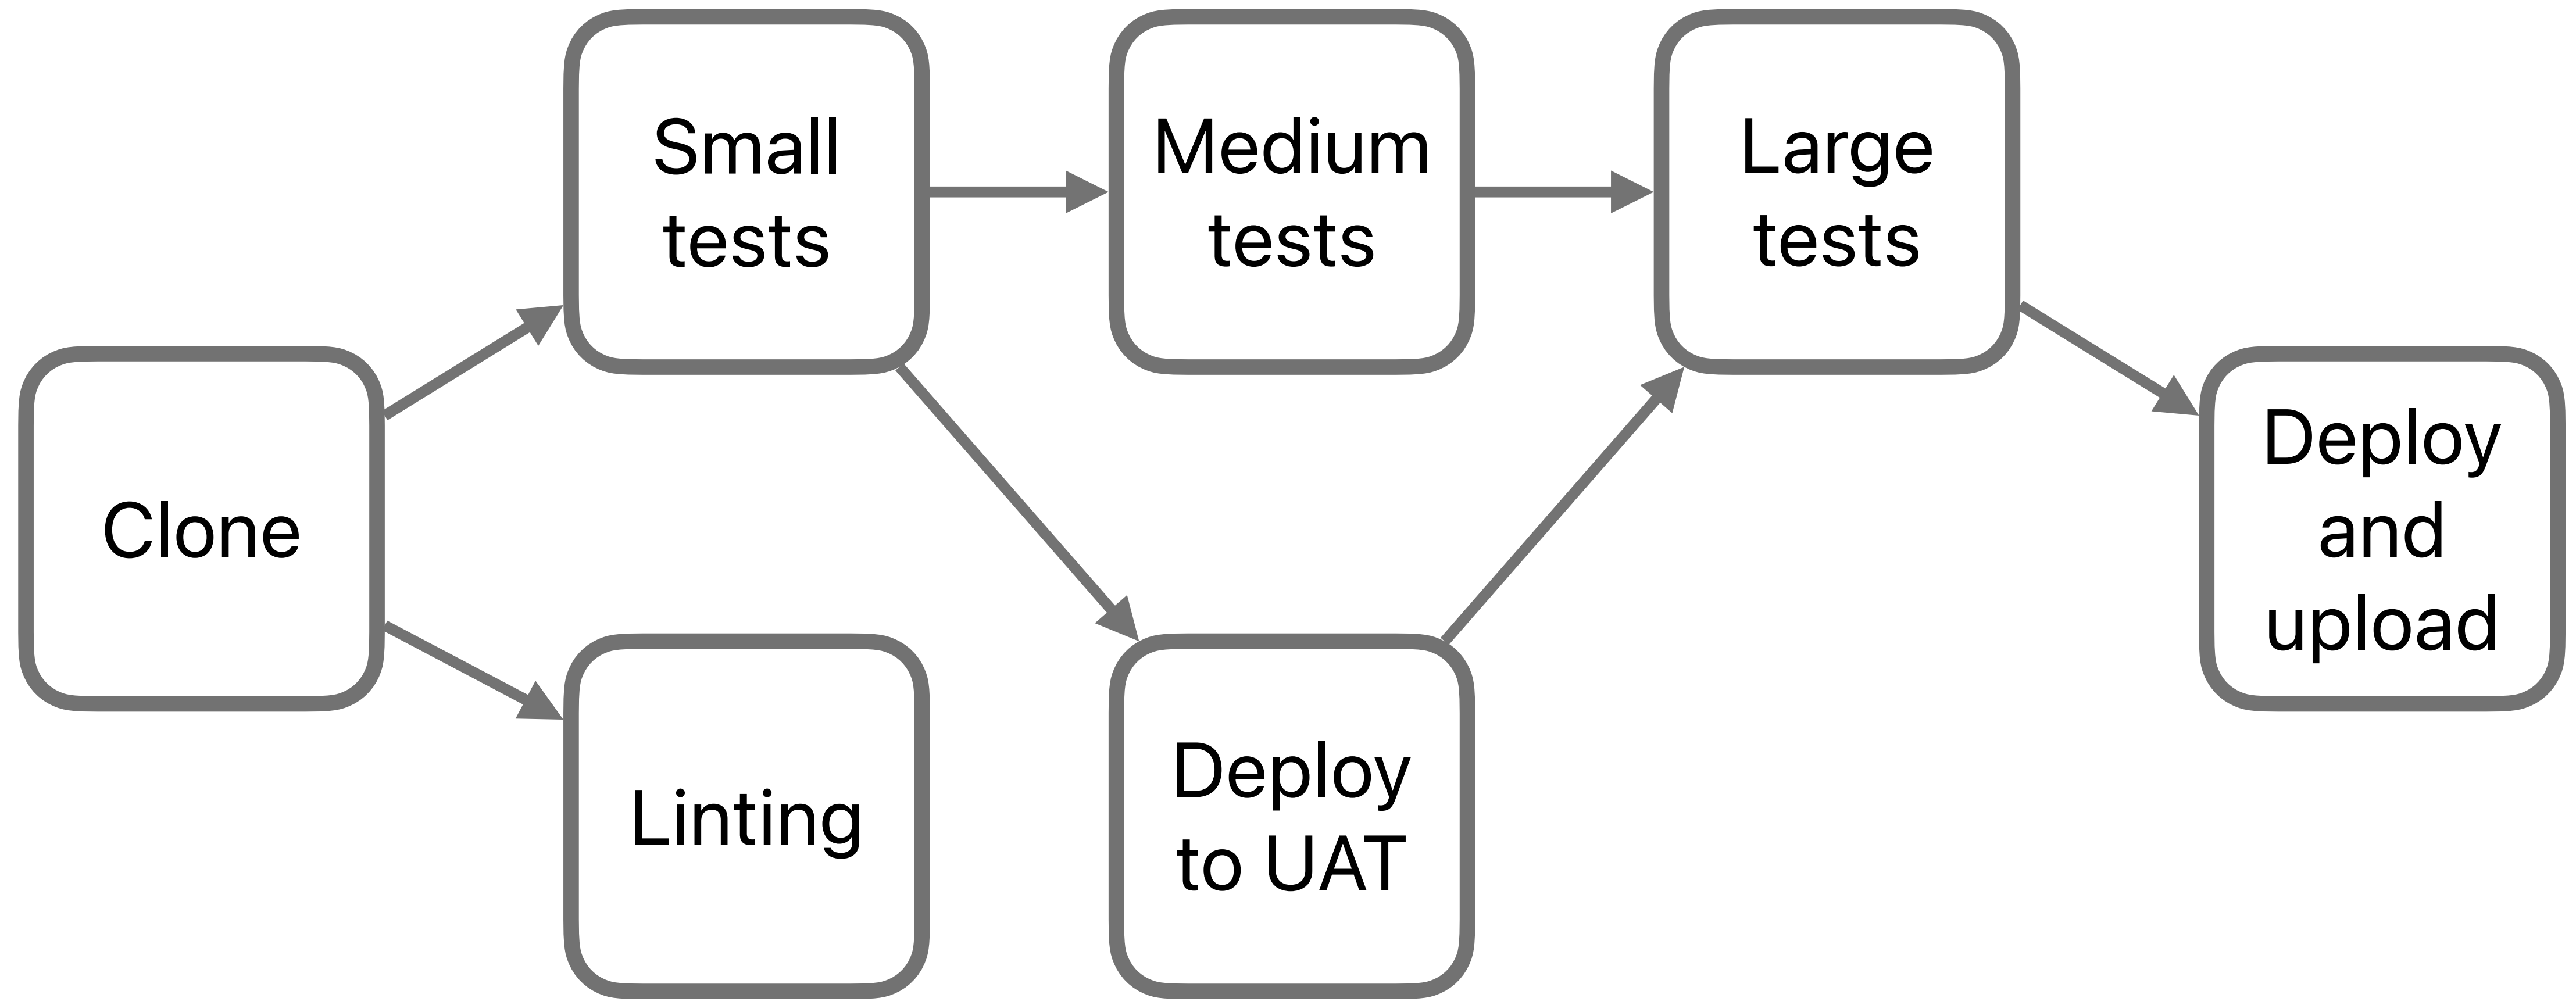
- The ideas in this talk can be applied on any code base...
- ... but I'm going to use one project as an example...
- ... and that project is Selenium
- Just remember, when you hear "Selenium" in this talk, think "my own project" 😊





Why use Continuous Integration?

Is this change safe?



A regular CI pipeline

“BUILD ALL THE THINGS!”



Why do we do this?

We don't know what's changed

How our CI build used to work

github.com/AlDanial/cloc v 1.98 T=13.33 s (174.8 files/s, 27369.3 lines/s)

Language	files	blank	comment	code
Java	1186	23873	27707	108445
C#	398	7657	16705	37157
JavaScript	240	6794	17787	32334
Rust	38	4103	886	25734
Ruby	274	4783	7075	17035
Python	193	5503	7399	13779
SUM:	2329	52713	77559	234484

By the numbers

A polyglot repo



Build times were climbing. Why?

10 months ago: ~45 minutes


9 months ago: ~2 hours

7 months ago: ~5 hours



Something had to change

How we build Selenium



# Introducing Bazel

A polyglot, graph-based build tool




Hermetic builds

Far less “works on my machine”




Purely functional builds

Same inputs, same outputs, amazing caching




Build graphs

Not always visible, always there



Changing a file


If we change this...



Changing a file


... we rebuild these






Changing another file

If we change this...




Changing another file

... we rebuild these




Adding a file

If we add this...




Adding a file

... we rebuild these




Removing a file

If we remove this...



Removing a file

... we rebuild these



Fast local builds

Graphs and pure functions == parallelism



Distributed builds

Cacheable! Distributable!

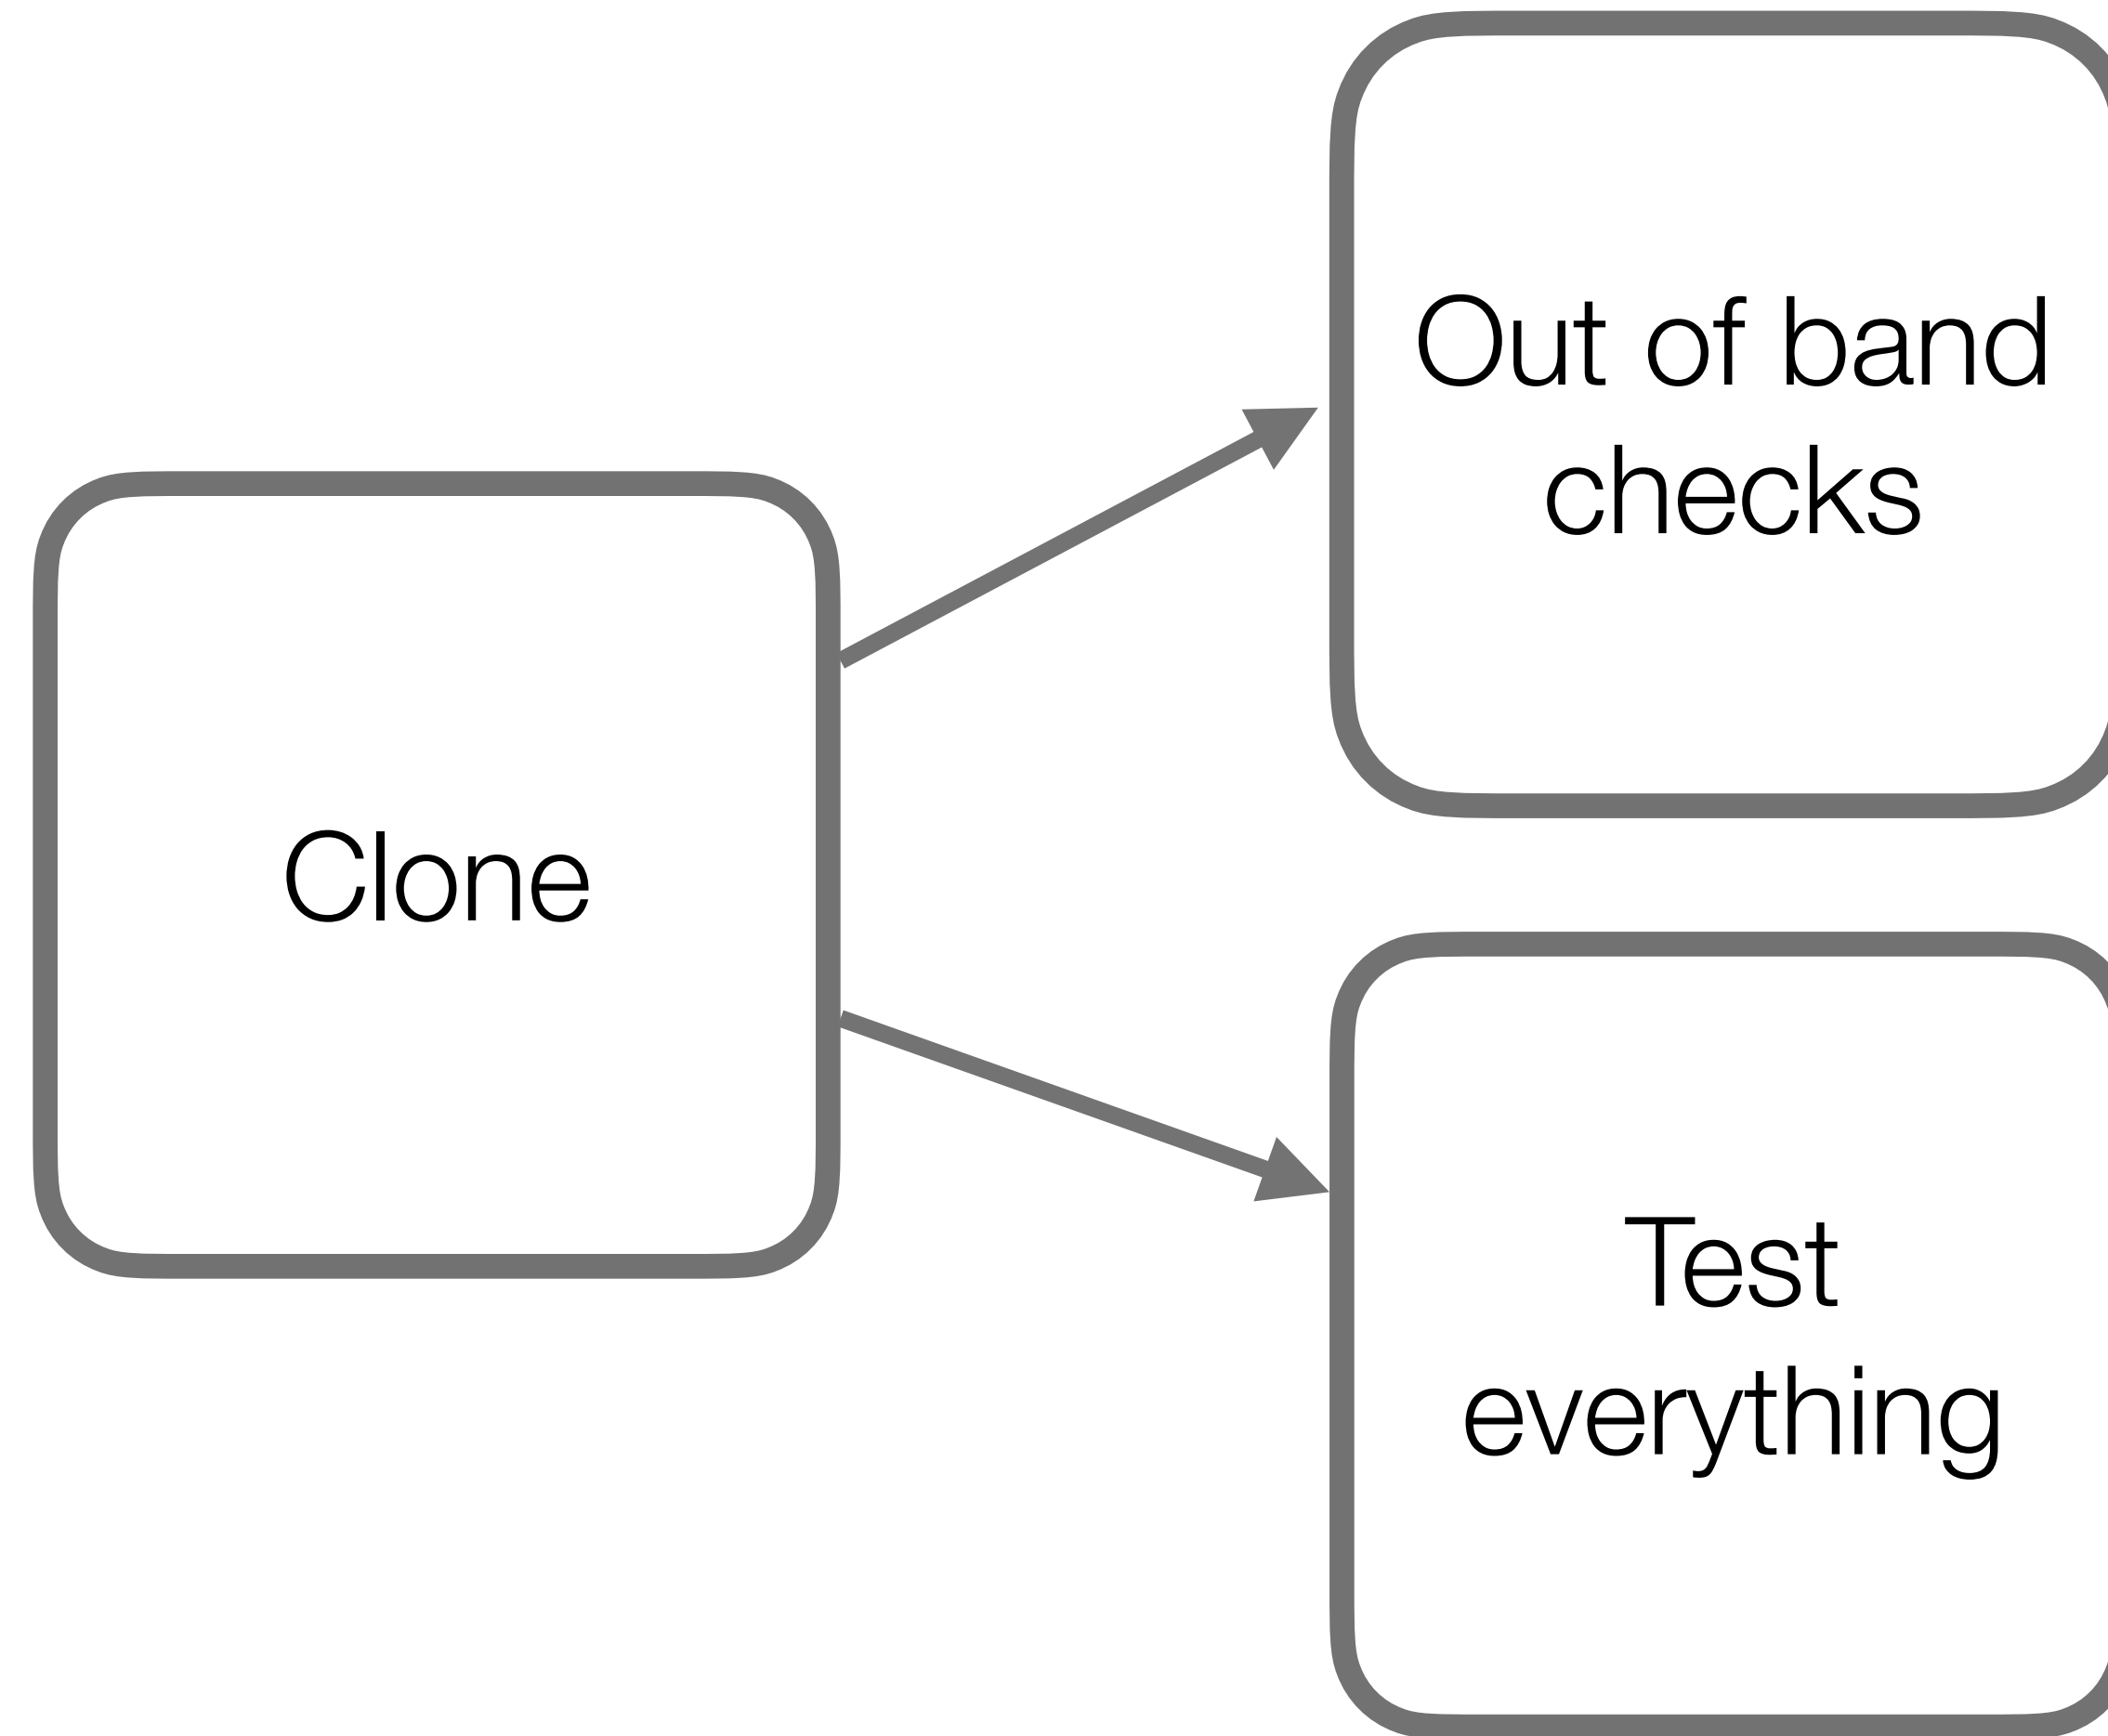




“Bazel as a Service”

Lots of ways to run an RBE

With this in mind, a better CI



Our new CI pipeline

---

# Out of band checks

---


- For “whole repo” operations
- Runs a code formatter
  - Java
  - Rust
  - Build files
- Any changes, and the build fails



# Test everything

---

- Runs every test we can on EngFlow
  - In parallel
- Runs tests that used to be in their own workflows
- Maintain a list of tests that are known to be troublesome
  - Retry tests at most twice
- Final step ensures all deployable targets build cleanly



267 workflow run results

Event ▾



















Status ▾

Branch ▾

Actor ▾

This workflow has a `workflow_dispatch` event trigger.

Run workflow ▾

 <b>[java] keep the streams alive</b> CI - RBE #1469: Commit <a href="#">0c86525</a> pushed by joerg1985	trunk	 1 hour ago  38m 40s	...
 <b>[rb] mark minimizing test as flaky</b> CI - RBE #1462: Commit <a href="#">afa349a</a> pushed by titusfortner	trunk	 6 hours ago  9m 33s	...
 <b>[rb] fix http proxy configuration for chrome (#13093)</b> CI - RBE #1460: Commit <a href="#">194b060</a> pushed by titusfortner	trunk	 10 hours ago  18m 50s	...
 <b>Fix logging levels in http.js and webdriver.js (#13098)</b> CI - RBE #1458: Commit <a href="#">76ace92</a> pushed by diemol	trunk	 12 hours ago  25m 39s	...
 <b>Allow external uri to be configurable for components that support ser...</b> CI - RBE #1455: Commit <a href="#">945e4f4</a> pushed by diemol	trunk	 13 hours ago  37m 5s	...
 <b>add rust files to update version script</b> CI - RBE #1451: Commit <a href="#">aaec17e</a> pushed by titusfortner	trunk	 yesterday  18m 13s	...

Recent test runs

So. Much. Faster

🏠 Summary

Jobs

✓ Format

✓ Test

Run details

🕒 Usage

📄 Workflow file

Triggered via push 1 hour ago

Status

Total duration

Artifacts

👤 joerg1985 pushed -> 0c86525 `trunk`

**Success**

**38m 40s**

—

**ci-rbe.yml**

on: push

✓ F.../ Check format scri... 1m 59s

✓ Test / All RBE tests 38m 32s



Here's what we do now

A fast formatting check, then run all the tests

### Cache Hits and Execution ?



#### Actions

<b>● Cached</b>	<b>73%</b>
● Fetched from disk cache	0%
● Fetched from remote cache	73%
<b>● Executed</b>	<b>27%</b>
● Executed locally	17%
● Executed remotely	10%

Cache hit rates may not include all expected data points

Bazel uses an internal cache to speed up incremental builds. Such cache hits are not reported by Bazel, so the data shown here does not include them.

[Learn more](#)

# Logs in EngFlow

The RBE has extensive caching









412 workflow run results

Event ▾ Status ▾ Branch ▾ Actor ▾

This workflow has a `workflow_dispatch` event trigger.

Run workflow ▾

 <b>[rb] mark minimizing test as flaky</b> CI #4386: Commit <a href="#">afa349a</a> pushed by titusfortner	trunk	📅 6 hours ago 🕒 43m 0s	...
 <b>Fix logging levels in http.js and webdriver.js (#13098)</b> CI #4381: Commit <a href="#">76ace92</a> pushed by diemol	trunk	📅 13 hours ago 🕒 1h 1m 55s	...
 <b>add rust files to update version script</b> CI #4373: Commit <a href="#">aaec17e</a> pushed by titusfortner	trunk	📅 yesterday 🕒 21m 57s	...
 <b>[py] add more logging to driver startup</b> CI #4372: Commit <a href="#">2621f82</a> pushed by titusfortner	trunk	📅 yesterday 🕒 11m 54s	...
 <b>[py] fix incorrect version changes</b> CI #4331: Commit <a href="#">1c9e702</a> pushed by titusfortner	trunk	📅 4 days ago 🕒 2h 0m 30s	...
 <b>Update mirror info (Thu Nov 2 12:06:36 UTC 2023)</b> CI #4327: Commit <a href="#">86df0ad</a> pushed by selenium-ci	trunk	📅 4 days ago 🕒 1m 54s	...

Not everything builds remotely

Less work, faster builds

# Summary

---

- The secret sauce is **Bazel**
- **Bazel** allows us...
  - ... to **query** our build graph
  - ... use **remote execution** and caching
- **Pervasive caching** allows us to **share a repo** with more people.
- **Remote builds** allow us to **test everything and expect it to be fast**
- This allows us to **dramatically decrease CI time**
- We could go faster!



# Challenges

---

- Bazel
  - Not everyone is familiar with Bazel
  - Not everyone likes Bazel
  - Not everything builds with Bazel
- As we add more tests, RBE build times will get worse.
  - We'll have to carefully structure our repo to avoid this
  - Target determination may help



# Q & A

---

Simon Stewart  
Ex Selenium Project Lead  
Build Nerd  
@shs96c on the socials