


The Path to Instant Rebuilds

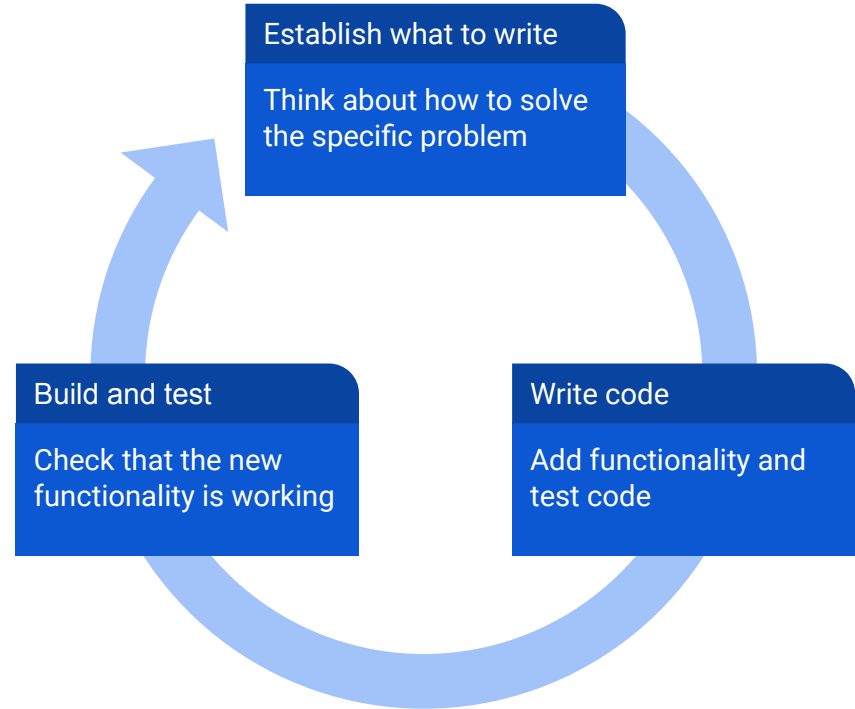

Less Context Switching and
Sub Minute Submit Times

DevLin2022

Software Development Loop



<https://xkcd.com/303/>



Build Startup Time

- Time before the first compile command is running
- Example: Chromium project
 - 30000 source files
 - 10s startup time with Make
 - Created Ninja
 - 1s startup time acceptable
 - https://ninja-build.org/manual.html#_introduction

Startup Time

- Example: Web skeleton loading
 - Don't let the user loose focus or go back
- Example: Chrome browser startup time
 - Regression test
- Solution: Choose your build system wisely

React Loading Skeleton

Loading

A Blog Post




A List



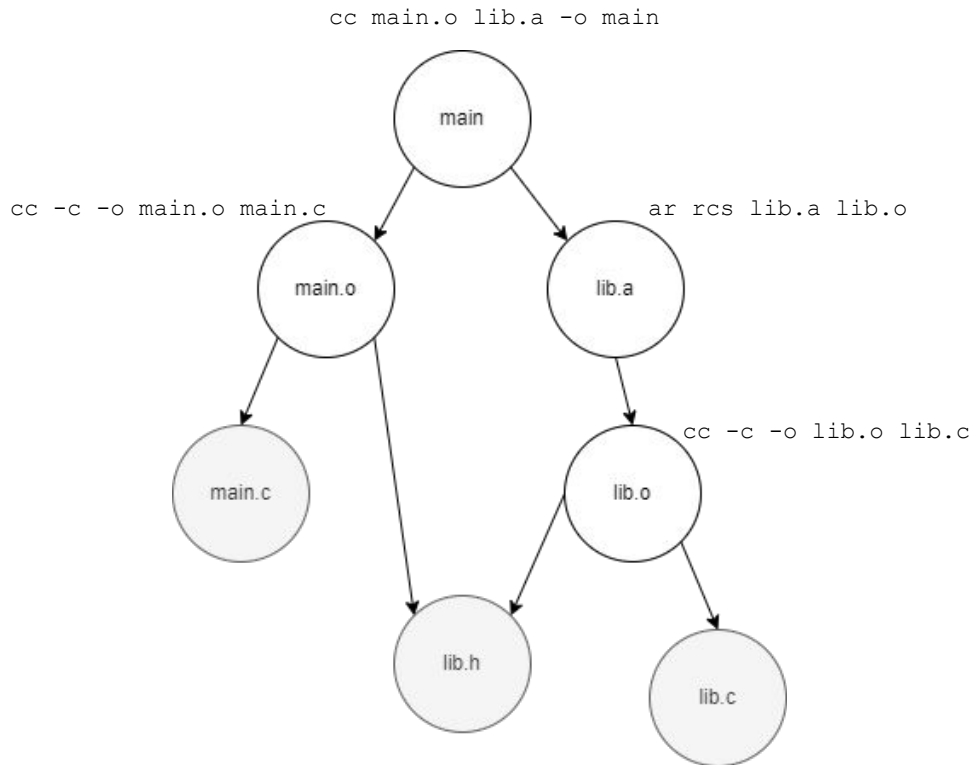
<https://www.npmjs.com/package/react-loading-skeleton>

Build and Test Time



How to Build and Test

- Example C program
 - main.c, lib.c, lib.h
- Build system
 - Describe the dependency graph
 - Examples: Make, Scons, Bazel
- NOTE: Everyone builds the same

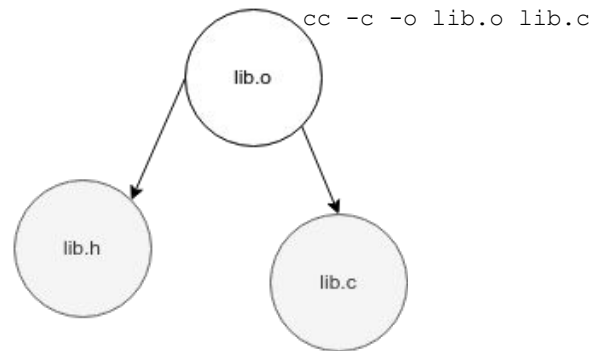


Remote Cache

- Action description
 - Command line
 - Environment variables
 - Input files and content

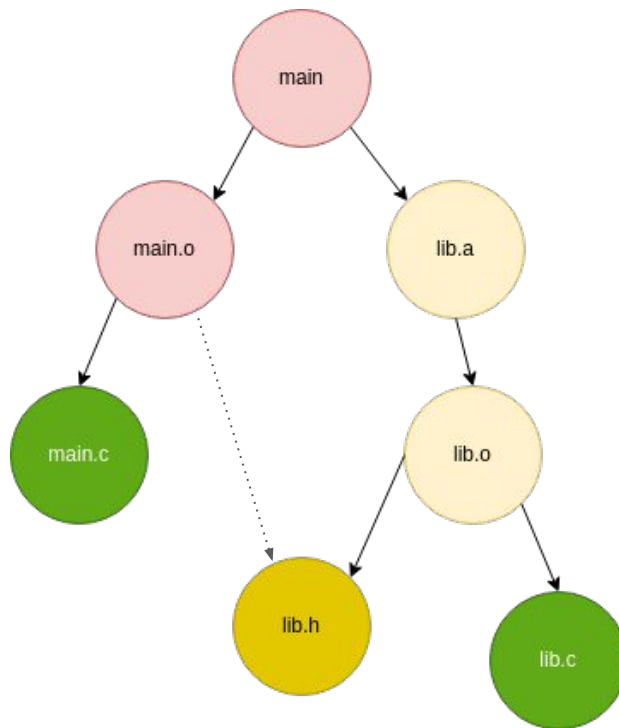
- Store the results on a cache server
 - More than 99% cache hit, built by your colleagues or CI
 - **Reduce build time by 90%**

- Deterministic actions
 - No random numbers
 - No date or time



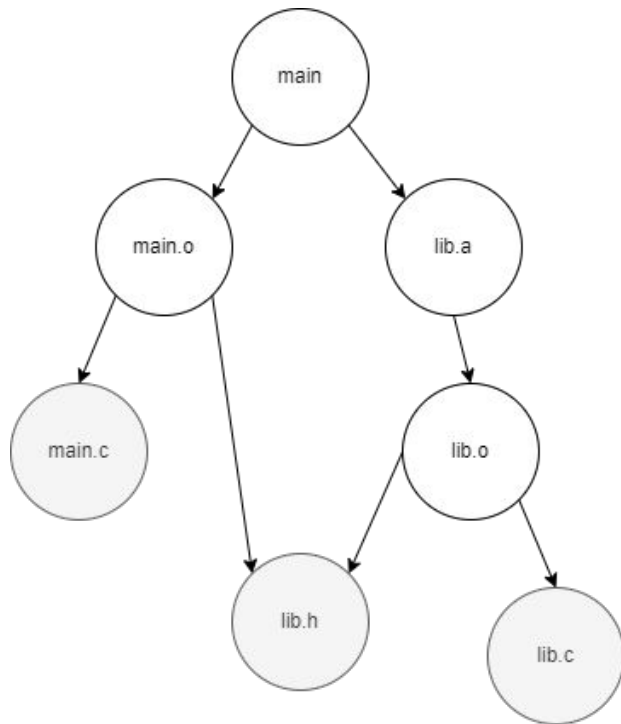
Underdescribed Graph - Cache Poisoning

- What if the graph is underdescribed?
 - This is a bug.
 - Builds might be correct by luck.
 - Workaround: **make clean && make all**
- Avoid cache poisoning
 - Sandbox each action
 - ... or use remote execution



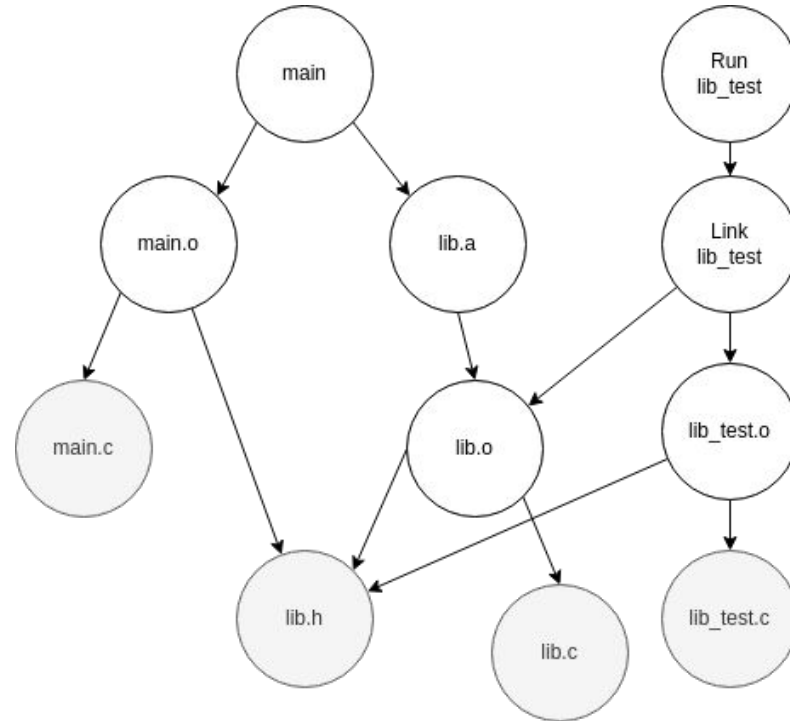
Remote Execution

- Compilations can be done in parallel
 - `make -j1000`
- Underdescribed actions won't build
 - Cannot access local host
 - In practice sandboxed environment
- Move the CPUs to the cluster
 - Small developer machines
 - Large cluster
 - Time share the resources
 - **Reduce build time by another 90%**



No Difference Between Build and Test

- Both are spawning actions
- Run the test in parallel with building main



Remote Execution API

<https://github.com/bazelbuild/remote-apis>

- Clients
 - Bazel - Google (Blaze)
 - Pants - Twitter
 - Buck - Facebook
 - Goma - Chromium (with ninja)
 - ...
- Servers
 - Buildbarn
 - Buildfarm
 - bazel-remote (cache only)
 - ...

Poll - Build System

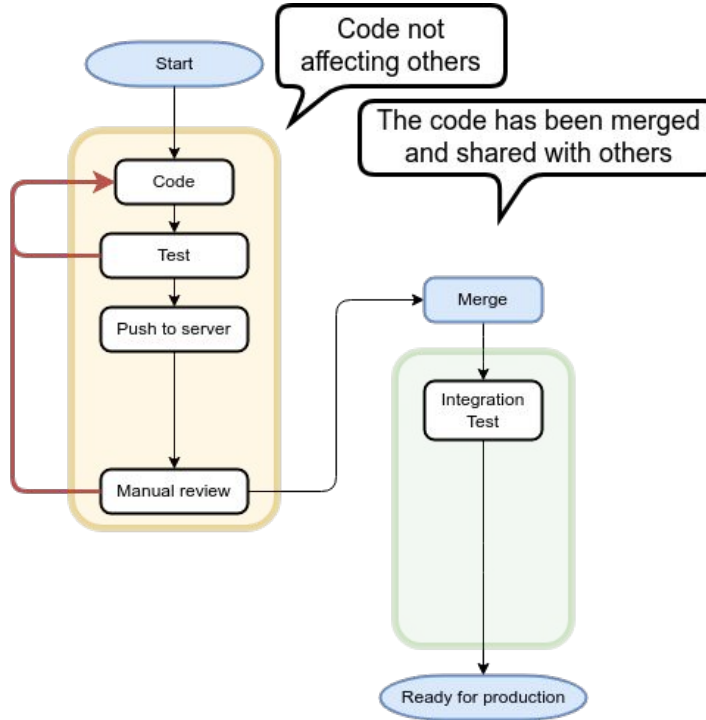
Do you use

- Remote cache
- Remote execution

Continuous Integration

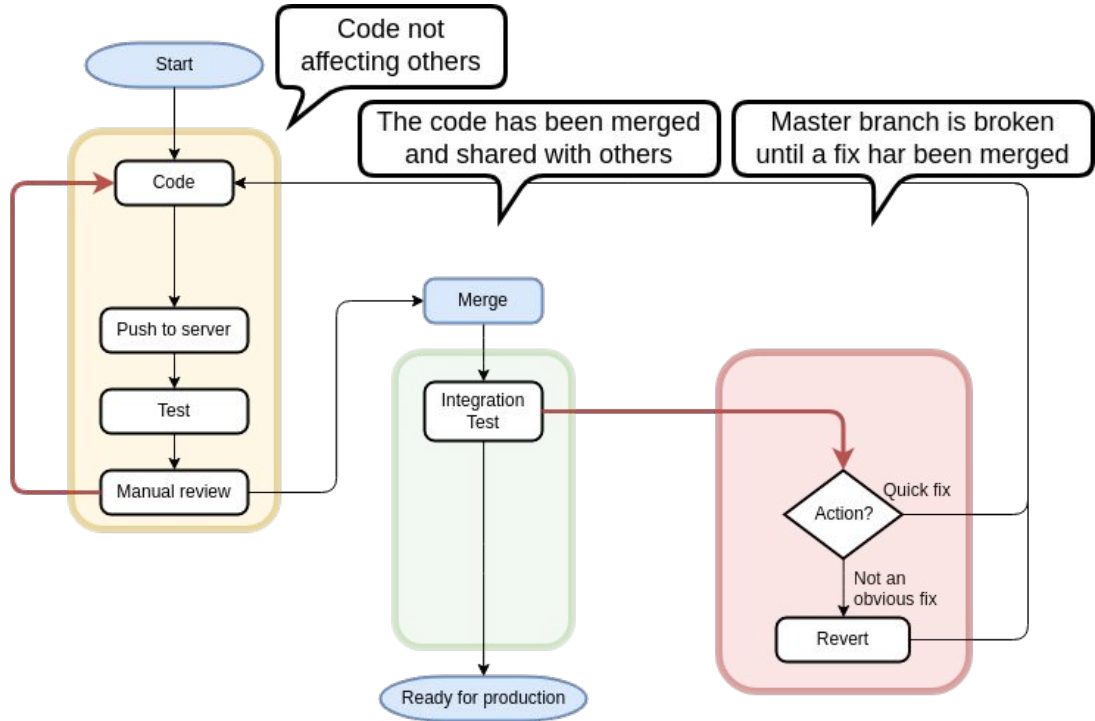
Code Integration

- Two phases
 - Presubmit
 - Assigned to a task
 - Post submit
 - Task done



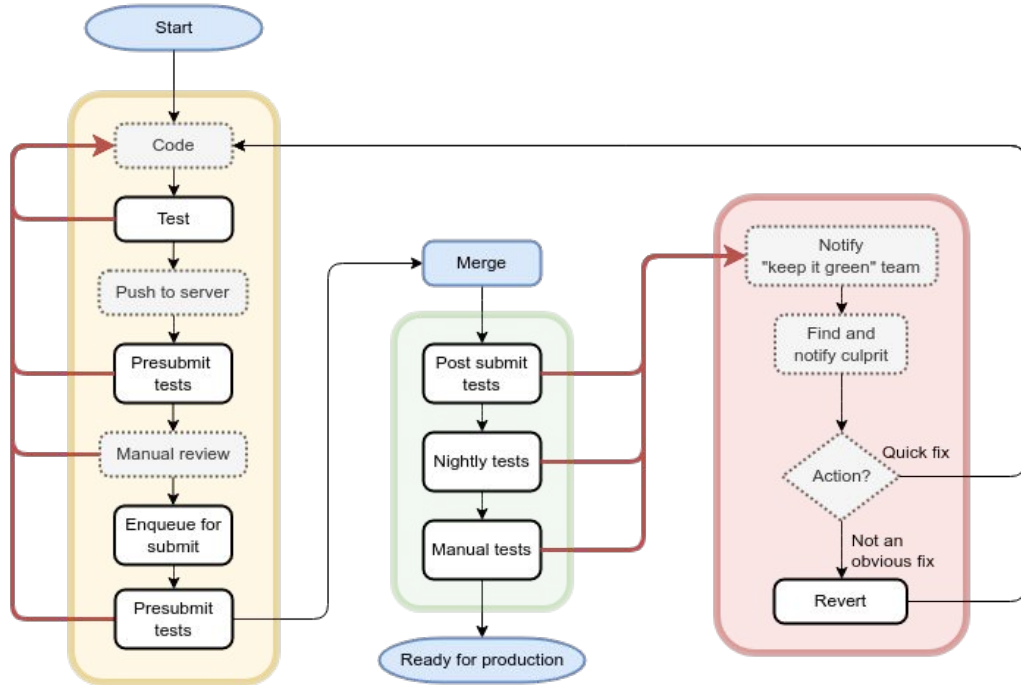
Code Integration

- Three phases
 - Presubmit
 - Post submit
 - Broken state
- Keep repo green
 - “All” tests should always pass



Move Tests Earlier

- Three phases
 - Presubmit
 - Post submit
 - Broken state
- Keep repo green
 - “All” tests should always pass
- Automated tests
 - When to run?
 - Machine cost



Summary - Problems with Latency

- Long startup times
 - Lost focus
- Developers need to juggle multiple tasks in parallel
 - Annoying to wait hours for fixing minor things
- Lower willingness to fix typos
 - Increased technical debt
- Reusing the same commit for multiple things
 - More difficult to review
- Extends the critical path
 - Slower feature growth in product

Poll - Presubmit times


Average time from triggering presubmit checks until merge

- Have at least some tests
- <1h
- <30m
- <10m
- <1m



CI Time - Locally


- Disk and network intensive
 - During startup
- CPU intensive
 - When compiling



- Checkout repository
- Fetch dependencies
- Build and test
- Upload and cleanup

CI Time - Remotely


- CPU intensive tasks minimized



- Checkout repository
- Fetch dependencies
- Build and test
- Upload and cleanup

CI Time - Remotely

- Disk and network intense
 - Not much CPU needed
- Next step
 - Reuse the workspace
 - Requires fully specified build graph



- Checkout repository
- Fetch dependencies
- Build and test
- Upload and cleanup

Resolving Red Repo State

Rebuild Locally

- Reproduce the error locally
- Start debugging
- Fix and test it

Rebuild Locally - Mimic the CI Pipeline

- CI Tools
 - Jenkins pipelines
 - GitHub actions
 - Gitlab pipelines
- Transferring artifacts
 - Deliverables between steps store externally

or...

- Shell scripts
- Defined process or helper scripts
 - ... to reproduce certain CI steps locally



MEROTON
SPEED UP YOUR BUILDS

Visualize Individual Test Results over Time

Results for abseil-cpp@master	1649c03	a87df8e	0064d9d	2b403ec	d819278	ea882fb	f5fd4cc	37a1aa0	1db72eb	fd9f9e7	8279400	1b6a9e9	bc09731	4691321	2fc358d	90184f6
	2022-11-02 19:45:24	2022-11-02 14:12:54	2022-11-01 21:27:28	2022-11-01 19:13:10	2022-10-31 20:32:21	2022-10-31 19:21:32	2022-10-31 17:27:59	2022-10-31 16:24:14	2022-10-31 20:22:03	2022-10-28 05:22:33	2022-10-28 22:45:48	2022-10-28 10:11:42	2022-10-28 19:52:27	2022-10-27 17:15:07	2022-10-25 21:56:36	2022-10-25 20:08:45

bazel test ... --//absl/cleanup/...

```

- //absl/...
+ //absl/algorithm/...
+ //absl/base/...
+ //absl/container/...
+ //absl/debugging/...
+ //absl/flags/...
- //absl/functional/...
  //absl/functional:any_invocable_test
  //absl/functional:bind_front_test
  //absl/functional:function_ref_test
  //absl/functional:function_type_benchmark
+ //absl/hash/...
+ //absl/log/...
+ //absl/memory/...
  //absl/meta:type_traits_test
+ //absl/numeric/...
+ //absl/profiling/...
+ //absl/random/...
+ //absl/status/...
+ //absl/strings/...
+ //absl/synchronization/...
+ //absl/time/...
+ //absl/types/...
  //absl/utility:utility_test

```

220/220	219/220	219/220	219/220	219/220	219/220	220/220	220/220	220/220	220/220	220/220	220/220	220/220	220/220	220/220	220/220	220/220
2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2
20/20	20/20	20/20	20/20	20/20	20/20	20/20	20/20	20/20	20/20	20/20	20/20	20/20	20/20	20/20	20/20	20/20
23/23	23/23	23/23	23/23	23/23	23/23	23/23	23/23	23/23	23/23	23/23	23/23	23/23	23/23	23/23	23/23	23/23
5/5	5/5	5/5	5/5	5/5	5/5	5/5	5/5	5/5	5/5	5/5	5/5	5/5	5/5	5/5	5/5	5/5
11/11	11/11	11/11	11/11	11/11	11/11	11/11	11/11	11/11	11/11	11/11	11/11	11/11	11/11	11/11	11/11	11/11
4/4	3/4	3/4	3/4	3/4	3/4	4/4	4/4	4/4	4/4	4/4	4/4	4/4	4/4	4/4	4/4	4/4
Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass
Pass	Fail	Fail	Fail	Fail	Fail	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass
Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass
Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass
3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3
14/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14	14/14
2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2
Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass
3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3
3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3	3/3
27/27	27/27	27/27	27/27	27/27	27/27	27/27	27/27	27/27	27/27	27/27	27/27	27/27	27/27	27/27	27/27	27/27
2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2	2/2
59/59	59/59	59/59	59/59	59/59	59/59	59/59	59/59	59/59	59/59	59/59	59/59	59/59	59/59	59/59	59/59	59/59
8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8	8/8
7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7	7/7
10/10	10/10	10/10	10/10	10/10	10/10	10/10	10/10	10/10	10/10	10/10	10/10	10/10	10/10	10/10	10/10	10/10
Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass	Pass

bazel test -c opt //absl/cleanup/...

```

- //absl/...
  //absl/cleanup:cleanup_test

```

0/1
Fail


0/1
Fail

1/1
Pass

1/1
Pass

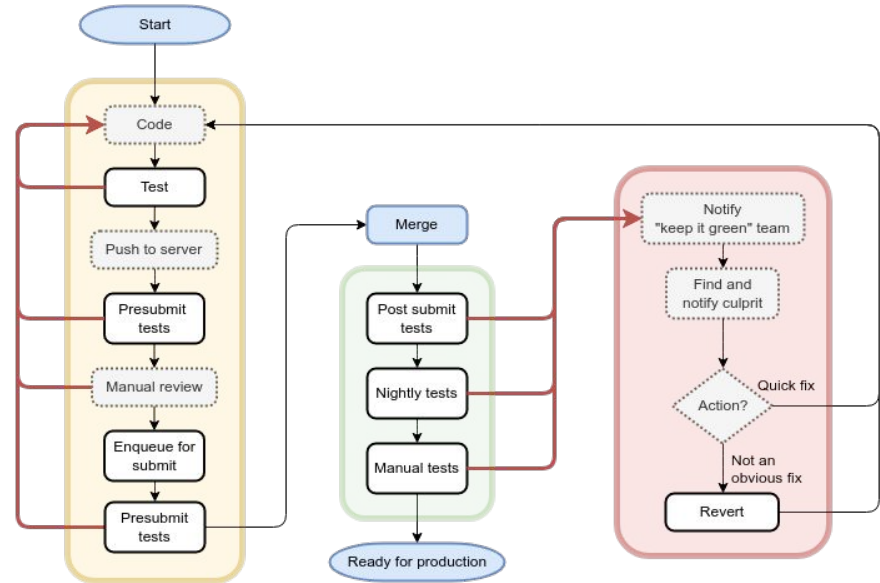
Interval Halving

- Interval halving to find the bad commit
 - `git bisect start HEAD <good-commit> -`
 - `git bisect run ./run_test.sh`
 - `git bisect reset`
- Start automatically
 - ... for nightly tests
- Review suggested revert
 - Fooled by a flaky test?
 - Is there an easy fix instead?
- Manual inspection in case of slow and expensive tests
 - Timeouts
 - Bricking hardware



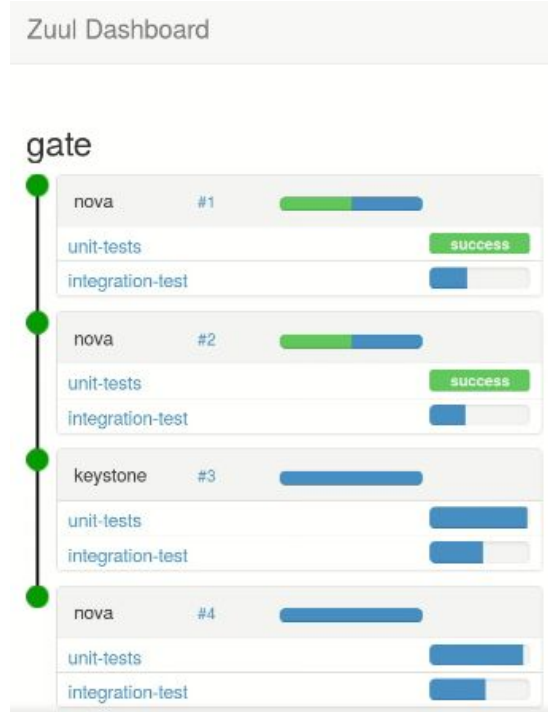
Keep Master Green

- Test what will be the merge result
- Move tests early
 - Includes integration tests



Speculative Gate Queue

- Test the future state after merge
 - Test main+1
 - Test main+1+2
 - Test main+1+2+3
 - Test main+1+2+3+4



<https://zuul-ci.org/>

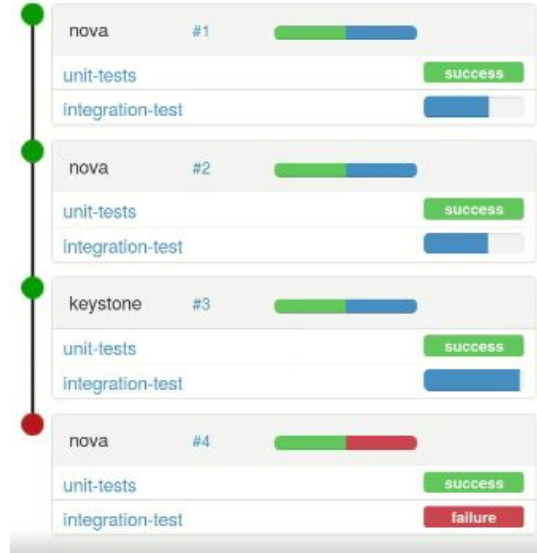
Speculative Gate Queue

- Test the future state after merge
 - Test main+1
 - Test main+1+2
 - Test main+1+2+3
 - Test main+1+2+3+4 - **Failing**

- Is #4 bad?

Zuul Dashboard

gate



<https://zuul-ci.org/>

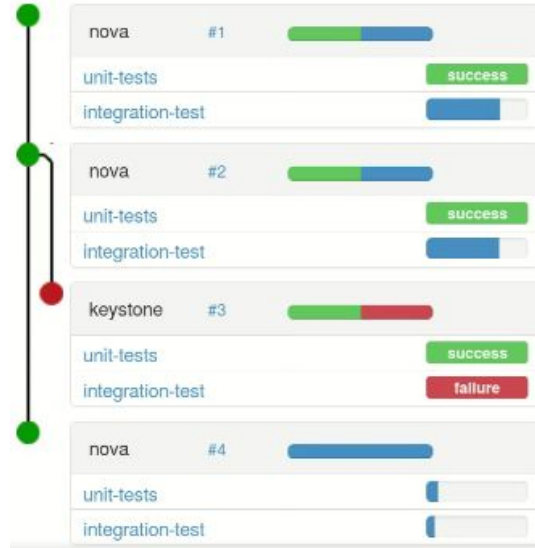
Speculative Gate Queue

- Test the future state after merge
 - Test main+1
 - Test main+1+2
 - Test main+1+2+3 - **Failing**
 - Test main+1+2+4

- Is #3 bad?

Zuul Dashboard

gate

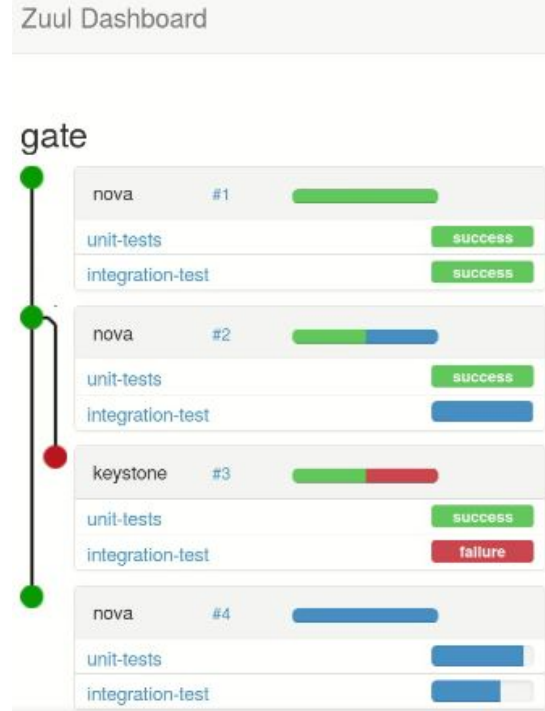


<https://zuul-ci.org/>

Speculative Gate Queue

- Test the future state after merge
 - Test main+1 - **Merging**
 - Test main+1+2
 - Test main+1+2+3 - **Failing**
 - Test main+1+2+4

- Is #3 bad?



<https://zuul-ci.org/>

Speculative Gate Queue

- Test the future state after merge
 - Test main+1 - Merged
 - Test main+1+2 - Merged
 - Test main+1+2+3 - Failing
 - Test main+1+2+4

- #3 is bad



<https://zuul-ci.org/>

Optimal Gate Queue

- Keeping Master Green at Scale - Uber
 - <https://dl.acm.org/doi/pdf/10.1145/3302424.3303970>
- Optimistic gate queue is not optimal
 - Changes have a probability of failing
 - Test 1+2+3
 - Test 1+3+4
 - Test 1+2+4
 - Then try 1+2+3+4
 - Correlated with
 - Size of change
 - Certain files
 - Author...

Poll - Gate Queue

Do you use a gate queue

- Tests on push
- Tests after review
- Speculative gate testing after review

Conclusion

- Quick local incremental build -> Keep mind focused
- Remote cache/execution -> Faster builds, for developers and CI
- Fast CI -> Finish one task at a time
- Fast CI -> Encourage small fixes, less technical debt

- Shell scripts -> Locally reproducible builds, now and in the future
- Shell scripts -> Avoid vendor lock in
- Gate queue -> Keep Master Green

References

Tools of interest:

- Bazel - {Fast, Correct} - Choose two
<https://bazel.build/>
- Buildbarn
<https://github.com/buildbarn/bb-deployments/>
- Zuul CI
<https://zuul-ci.org/>

Previous talks on the subject:

- One Minute Presubmits
<https://docs.google.com/presentation/d/14dxac2omYI5Feaoiw-u09qB1fQgJU7ASNJag7YMg3TI/>
- Selective testing in Bazel
 - BazelCon 2019 - Selective Testing by Benjamin Peterson
https://www.youtube.com/watch?v=9Dk7mtlm7_A
 - <https://github.com/Tinder/bazel-diff>