



# Continuous Integration at Google Scale

By John Micco

Developer Infrastructure

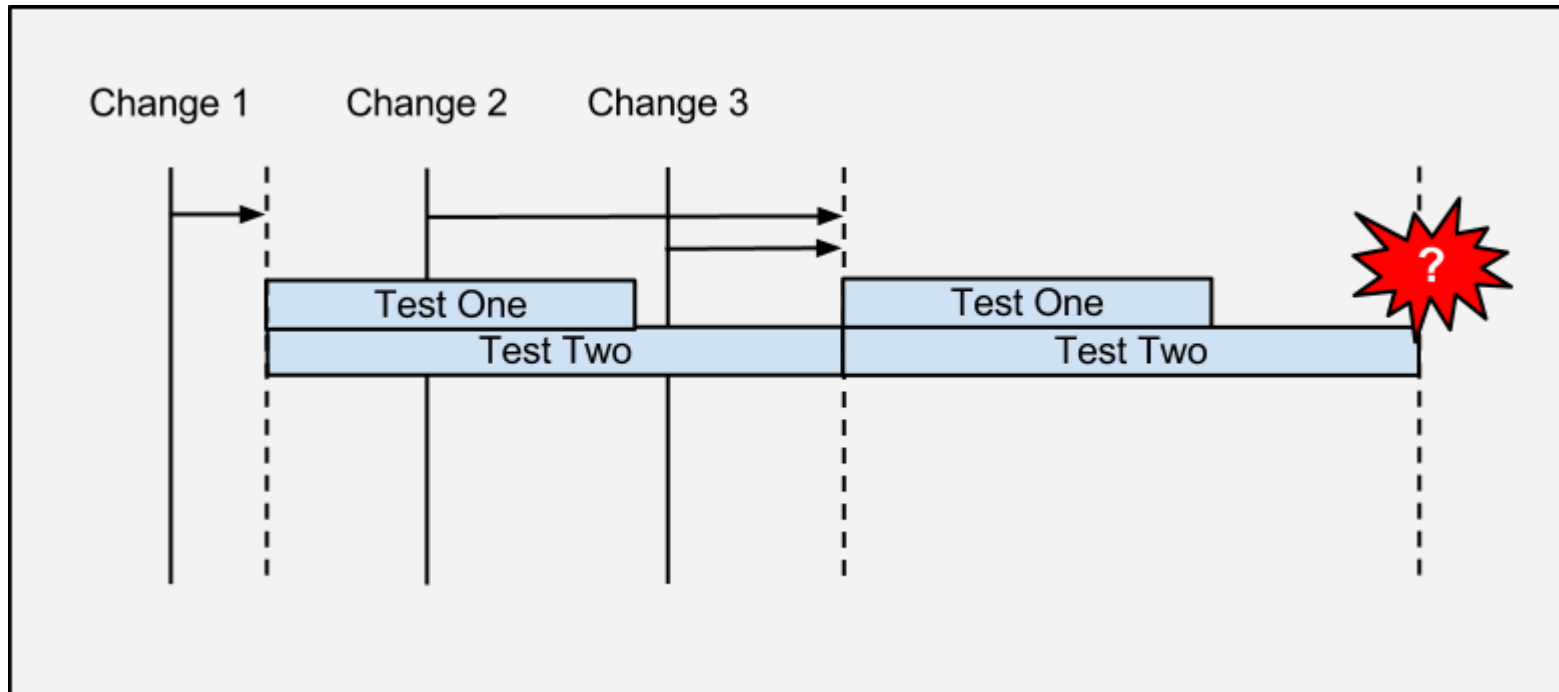
- >10,000 developers in 40+ offices
- 5000+ projects under active development
- 17k submissions per day (1 every 5 seconds)
- Single monolithic code tree with mixed language code
- Development on one branch - submissions at head
- All builds from source
- 20+ sustained code changes per minute with 60+ peaks
- 50% of code changes monthly
- 100+ million test cases run per day

1. Continuous Integration Goals
2. Continuous Integration at Google
3. Practical Matters

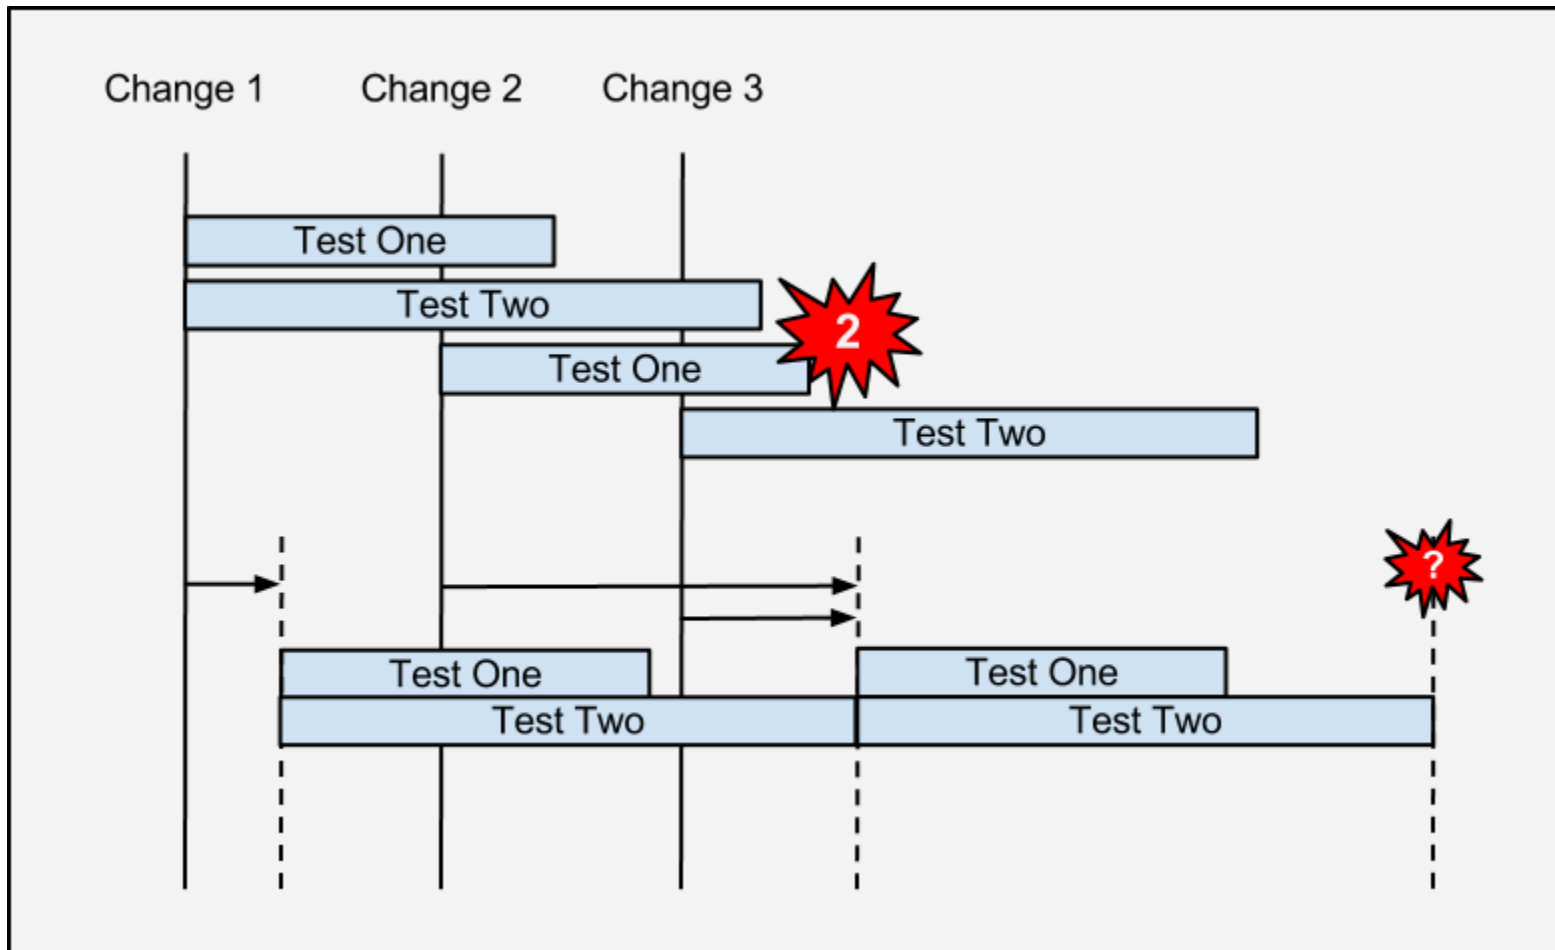
- Provide real-time information to build monitors
  - Identify failures fast
  - Identify culprit Changes
  - Handle flaky tests
- Provide frequent green builds for cutting releases
  - Identify recent green builds
  - Show results of all testing together
  - Allow release tooling to choose a green build
  - Handle flaky tests

- Develop Safely
  - Sync to last green changelist
  - Identify whether change will break the build before submit
  - Submit with confidence
  - Handle flaky tests

- Triggers builds in continuous cycle
- Cycle time = longest build + test cycle
- Tests many changes together
- Which change broke the build?



- Triggers tests on every change
- Uses fine-grained dependencies
- Change 2 broke test 1



# Continuous Integration Display

Current Status

Grid

Test Log

Coverage

Project Maintenance

Project Health (beta)

History

Failed / Broken

Target name

Search target

<< Head

< Newer

CLS 30805794

- 30804822

Older >

Showing 12 of 1166 targets: Failed / Broken [Remove all filters](#)

Changelist and submit time:	30805794 13 hrs	30805729 14 hrs	14 hrs	30805717 14 hrs	30805645 14 hrs	30805578 14 hrs	30805555 14 hrs	30805504 14 hrs	30805495 14 hrs	30805465 14 hrs	30805343 14 hrs	30805322 14 hrs	30805308 14 hrs	30805298 14 hrs	30805279 14 hrs	30805270 14 hrs	30805264 14 hrs	30805233 14 hrs	30805119 14 hrs	30805108 14 hrs	30805099 14 hrs	30805021 14 hrs	30804936 14 hrs	30804921 14 hrs	30804890 14 hrs	...	Prev. finished		
Project Status:	failing	failing	failing	failing	failing	failing	failing	failing	failing	failing	failing	failing	failing	failing	failing	failing	failing	failing	failing	passing	passing	passing	passing	passing	passing	passing	failing		
Affected targets:	Failed: 1 Passed: 766	Failed: 4 Passed: 528	Failed: 2 Passed: 471	Failed: 2 Passed: 183	Failed: 2 Passed: 740	Failed: 1 Passed: 226	Passed: 1	Failed: 1 Passed: 474	Failed: 2 Passed: 113	Failed: 2 Timed out: 1 Passed: 163	Failed: 1 Passed: 30	Failed: 1 Passed: 175	Failed: 1 Passed: 193	Failed: 2 Passed: 254	Failed: 1 Passed: 293	Failed: 1 Passed: 254	Passed: 14	Failed: 3 Passed: 249	Passed: 266	Passed: 1	Passed: 250	Passed: 265	Passed: 250	Passed: 166	Failed: 1 Passed: 273				

Tests:

✓	✓	✓	✓	✓					✓		F		✓										✓			...	✓
✓	✓	✓	✓	✓	✓			✓	F	✓		✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	...	✓
F	✓	✓	✓	✓	✓							✓	✓	✓	✓	✓	✓	F	✓			✓	✓	✓	✓	...	✓
✓	✓	✓	F	✓	✓			✓	✓	F			✓	✓	✓	✓		✓	✓			✓	✓	✓	✓	...	✓
✓	✓	✓	✓	✓	✓			✓	✓		✓		✓	✓	✓	✓		✓	✓			✓	✓	✓	✓	...	✓
✓	✓	✓	✓	✓	✓			✓	✓				✓	✓	F	✓	✓	✓	✓			✓	✓	✓	✓	...	✓
✓	✓	✓	✓	✓	✓			✓	✓	⏸			✓	✓	✓	✓		F	✓			✓	✓	✓	✓	...	✓
✓	✓	F	✓	✓	F	✓		✓	✓				✓	✓	✓	✓		✓	✓			✓	✓	✓	✓	...	✓
✓	✓	✓	✓	F	✓	✓		✓	✓				✓	✓	✓	✓		✓	✓			✓	✓	✓	✓	...	✓
✓	✓	✓	✓	F	F	F		F	F	F		F	F	F	F	F		F				✓	✓	✓	✓	...	✓
✓	✓	F	✓	✓	✓	✓		✓	✓			✓	✓	✓	✓	✓		✓	✓			✓	✓	✓	F	...	✓



- Identifies failures sooner
- Identifies culprit change precisely
  - Avoids divide-and-conquer and tribal knowledge
- Lowers compute costs using fine grained dependencies
- Keeps the build green by reducing time to fix breaks
- Accepted enthusiastically by product teams
- Enables teams to ship with fast iteration times
  - Supports submit-to-production times of less than 36 hours for some projects

- Requires enormous investment in compute resources (it helps to be at Google) grows in proportion to:
  - Submission rate
  - Average build + test time
  - Variants (debug, opt, valgrind, etc.)
  - Increasing dependencies on core libraries
  - Branches
- Requires updating dependencies on each change
  - Takes time to update - delays start of testing

- Makes testing available before submit
- Uses fine-grained dependencies
  - Recalculate any dependency changes
- Uses same pool of compute resources at high priority
- Avoids breaking the build
- Captures contents of a change and tests in isolation
  - Tests against head
  - Identifies problems with missing files
- Integrates with
  - submission tool - submit iff testing is green
  - Code Review Tool - results are posted to the review thread

## Pending CL 30795386 : Presubmit Still Running

### ▼ Still Running (1)



//javatests/com/google/payments/testing/malbec/scenarios/fromconsole/sellersignup:LargeTapTests [\[Details & Test History\]](#)

### ▼ Newly Failing (1)



//javatests/com/google/moneta/storedvalue/service:LargeTests [\[Details & Test History\]](#)

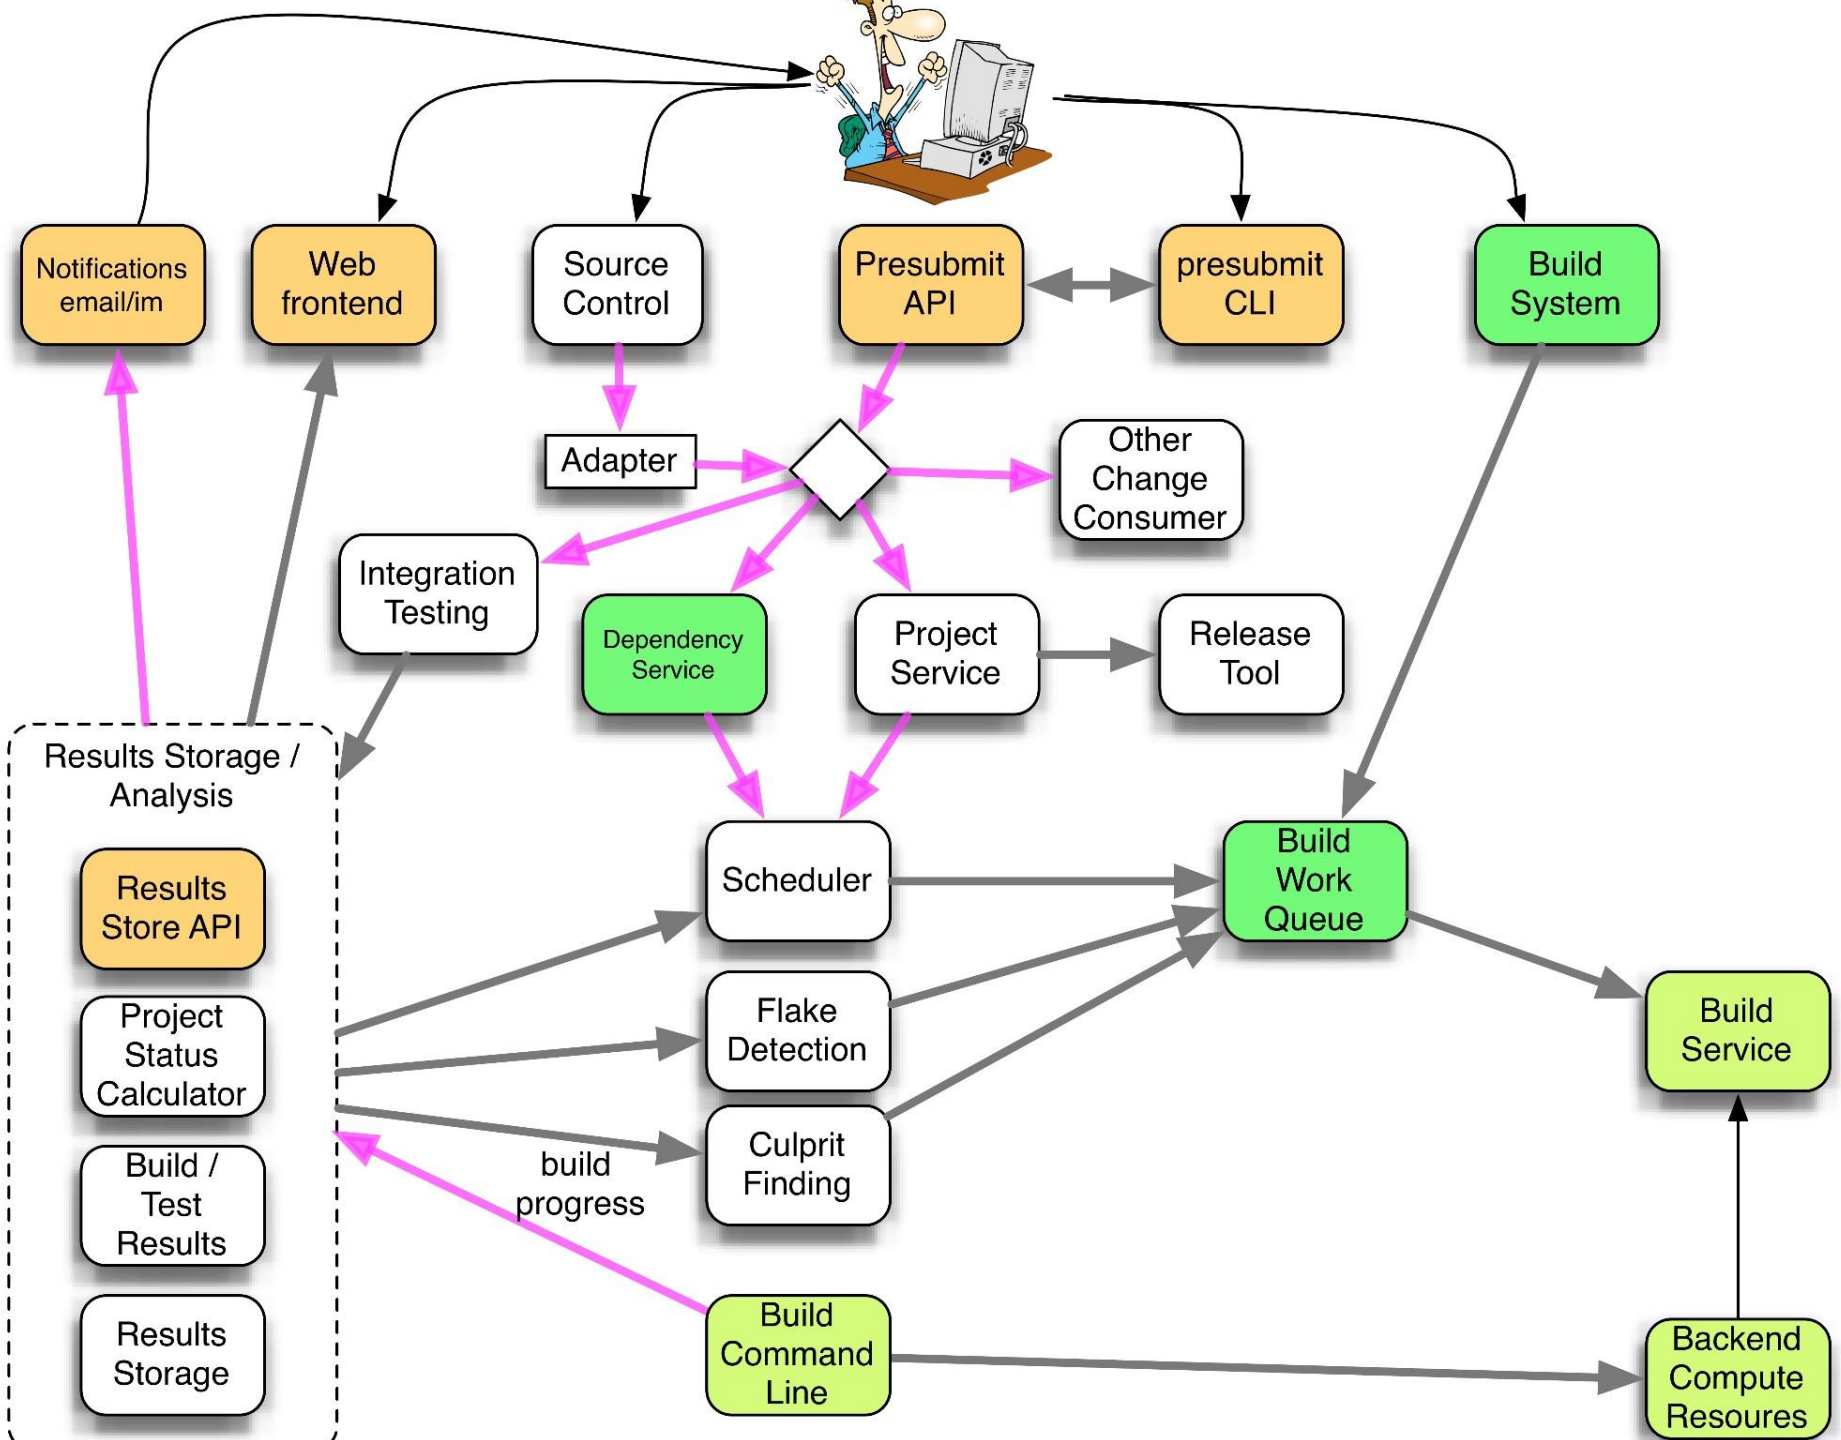
### ▼ Newly Passing (1)



//javatests/com/google/checkout/external/virtualproxycard/servers:RpcFunctionalTests [\[Details & Test History\]](#)

### ▶ Still Passing (1366)

### ▶ Skipped (223)



- System assumes tests pass or fail reliably given code
  - Tests that don't have this property are "flaky"
- Sources of test flakiness:
  - Infrastructure
    - machine failure
    - environment / setup problems
    - leakage - one test impacting another
    - Overloading resources
  - Tests
    - race conditions
    - external dependencies
    - timeouts
  - Code-under-test
    - memory problems
    - order dependence (e.g. hash tables)

- Causes

- Inability to find changes breaking the build - false positives
- Inability to identify green builds for releases
- Wasted work for build monitors
- Wasted compute resources
- Inappropriately failing presubmits - wasting developer time


- Solutions (Google does all of these):

- Fix them!!!
  - Difficult - requires developer time
- Hide them
  - Retry causing delays
  - Identify infrastructure flakes
  - Use metrics to ignore
- Track them
  - Provide metrics to prioritize fix / hide

- Sources of growth in test execution time
  - More developers = increased submission rate
  - More tests
  - Longer running tests
  - Tests consuming more resources (threading)
- Examine the growth trends
  - Predict compute needs
  - Look for any build system features required



# Build / Test Compute Resources



- Problems
  - Quadratic execution time growth
  - Ultimately cannot run every affected test @ every change
  - Low latency results still top requirement
- Solution: Just in time scheduling (JIT)



## Continuous Integration:

- Run every test affected at every ~~changelist~~.

JIT

as often <sup>^</sup> as possible

## In Production:

- Build and run tests concurrently on Google's distributed build and test backend.

# JIT Scheduling



Schedule tests to run only when system has capacity.



Produce project-wide results at periodic changelists.

## Same User Experience; Lower Cost



### Culprit finding

- Failures / breaks between changes may be more difficult to localize to the offending change.
- **Short-term:** Command-line tool to find culprits
- **Longer Term:** Integrated automatic culprit finding

## Same User Experience; Lower Cost



### Flaky Tests

- Tests which only pass some of the time could cause fewer green statuses for projects.
- **Short Term:** Optionally retry failed tests
- **Longer Term:** Tightly integrated flake mitigation and automatic / manual re-running of suspected flakes

## Q &amp; A

## For more information:

- <http://google-engtools.blogspot.com/2011/06/testing-at-speed-and-scale-of-google.html>
- <http://www.youtube.com/watch?v=b52aXZ2yi08>
- <http://www.infoq.com/presentations/Development-at-Google>
- <http://google-engtools.blogspot.com/>
- <http://misko.hevery.com/2008/11/11/clean-code-talks-dependency-injection/>
- [https://www.youtube.com/watch?v=KH2\\_sB1A6IA&feature=youtube\\_gdata\\_player](https://www.youtube.com/watch?v=KH2_sB1A6IA&feature=youtube_gdata_player)