

Principles of Software Construction: Objects, Design, and Concurrency

DevOps

Charlie Garrod

Chris Timperley



Administrivia


- Homework 6 has been released
 - Sequential implementation due by Tuesday, Nov. 26
 - Parallel implementation due by Wednesday, Dec. 4

Outline

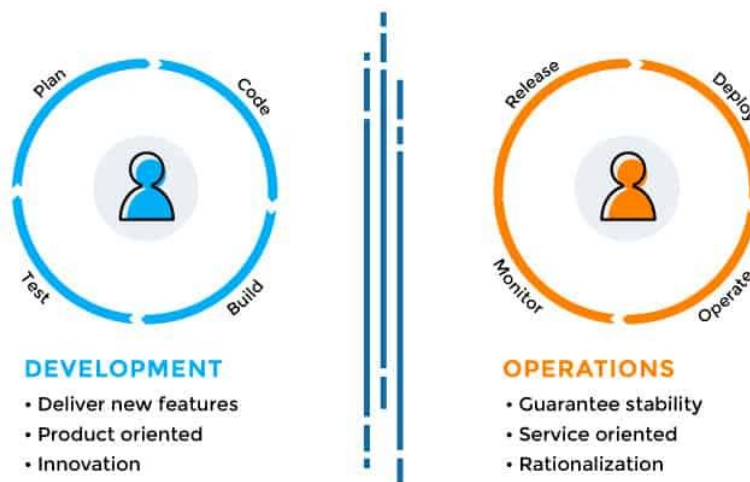
- DevOps and CI/CD
- Large-Scale Version Control
- Release Management

Devs, Ops, and The Wall of Confusion


OLD WAY:



Wall of Confusion





DevOps: Development / Operations



<https://blog.gds.gov.tech/that-ci-cd-thing-principles-implementation-tools-aa8e77f9a350>

Principle: Automation Everywhere

```
----- INSTALL.SH -----  
#!/bin/bash  
  
pip install "$1" &  
easy_install "$1" &  
brew install "$1" &  
npm install "$1" &  
yum install "$1" & dnf install "$1" &  
docker run "$1" &  
pkg install "$1" &  
apt-get install "$1" &  
sudo apt-get install "$1" &  
steamcmd +app_update "$1" validate &  
git clone https://github.com/"$1"/"$1" &  
cd "$1"; ./configure; make; make install &  
curl "$1" | bash &
```



<https://blog.chef.io/automate-all-the-things/>

Principle: Code as Configuration

- Manage configuration files in your version control system
 - Travis, Gradle, Jenkins, ...
- Packaging and installation
 - Docker, package.json, setup.py, pom.xml, ...
- Infrastructure and deployment
 - Docker Compose, Ansible, Puppet, Kubernetes
 - Manage servers and resources
- ...

```
98 lines (85 stoc) | 2.13 KB
1  apply plugin: 'java'
2  apply plugin: 'eclipse'
3  apply plugin: 'checkstyle'
4  apply plugin: 'jacoco'
5
6  test.testLogging {
7      exceptionFormat "full"
8      events "failed", "passed", "skipped"
9  }
10
11 configurations.all {
12     resolutionStrategy {
13         force 'org.ow2.asm:asm:6.2.1'
14         forcedModules = [ 'org.ow2.asm:asm:6.2.1' ]
15     }
16 }
17
18 check.doFirst {
19     List<String> missing = new ArrayList<>();
20     for (name in [ "domain.pdf",
21                 "system_sequence.pdf",
22                 "behavioral_contract.pdf",
23                 "interaction_title_validation.pdf",
24                 "interaction_monastery_scoring.pdf",
25                 "object.pdf",
26                 "rationale.pdf",
27                 "README.md" ]) {
28         String path = "design_documents" + File.separator + name;
29         if (!file(path).exists()) {
30             missing.add(path);
31         }
32     }
33     if (missing.size() != 0) {
34         String message = "The following files were missing:\n\n";
35         message += String.join("\n", missing);
36         message += "\n\nPlease check the expected file names in the handout.";
37         throw new GradleException(message);
38     }
39 }
```

Installation and configuration can be annoying

- Build flags
- Build order
- Static dependencies
- Dynamic dependencies
- Environment variables
- Configuration files
- DLL hell
- ...



Getting the Source Code and Building LLVM

The LLVM Getting Started documentation may be out of date. The [Clang Getting Started](#) page might have more accurate information.

This is an example workflow and configuration to get and build the LLVM source:

1. Checkout LLVM (including related subprojects like Clang):
 - `git clone https://github.com/llvm/llvm-project.git`
 - Or, on windows, `git clone --config core.autocrlf=false https://github.com/llvm/llvm-project.git`
2. Configure and build LLVM and Clang:
 - `cd llvm-project`
 - `mkdir build`
 - `cd build`
 - `cmake -G <generator> [options] ../llvm`

Some common generators are:

- Ninja — for generating Ninja build files. Most LLVM developers use Ninja.
- Unix Makefiles — for generating make-compatible parallel makefiles.
- Visual Studio — for generating Visual Studio projects and solutions.
- Xcode — for generating Xcode projects.

Some Common options:

- `-DLLVM_ENABLE_PROJECTS='...'` — semicolon-separated list of the LLVM subprojects you'd like to additionally build. Can include any of: clang, clang-tools-extra, libcxx, libcxxabi, libunwind, lldb, compiler-rt, lld, polly, or debuginfo-tests.

For example, to build LLVM, Clang, libcxx, and libcxxabi, use `-DLLVM_ENABLE_PROJECTS="clang;libcxx;libcxxabi"`.

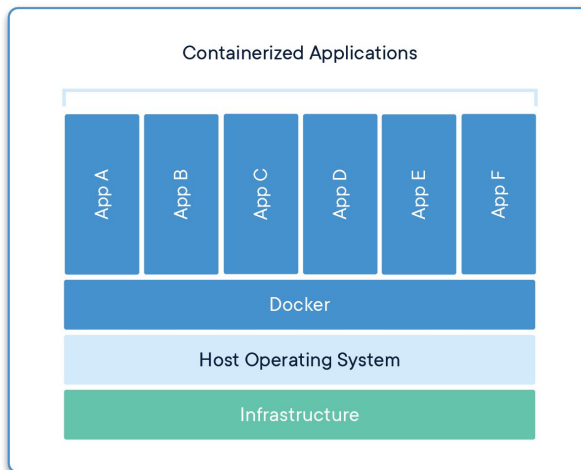
- `-DCMAKE_INSTALL_PREFIX=directory` — Specify for *directory* the full pathname of where you want the LLVM tools and libraries to be installed (default `/usr/local`).
- `-DCMAKE_BUILD_TYPE=type` — Valid options for *type* are Debug, Release, RelWithDebInfo, and MinSizeRel. Default is Debug.
- `-DLLVM_ENABLE_ASSERTIONS=on` — Compile with assertion checks enabled (default is Yes for Debug builds, No for all other build types).

- Run your build tool of choice!
 - The default target (i.e. `ninja` or `make`) will build all of LLVM.
 - The `check-all` target (i.e. `ninja check-all`) will run the regression tests to ensure everything is in working order.
 - CMake will generate build targets for each tool and library, and most LLVM sub-projects generate their own `check-<project>` target.
 - Running a serial build will be *slow*. To improve speed, try running a parallel build. That's done by default in Ninja; for make, use `make -j NNN` (NNN is the number of parallel jobs, use e.g. number of CPUs you have.)
- For more information see [CMake](#)
- If you get an "internal compiler error (ICE)" or test failures, see [below](#).

Consult the [Getting Started with LLVM](#) section for detailed information on configuring and compiling LLVM. Go to [Directory Layout](#) to learn about the layout of the source code tree.



- Uses lightweight containerization
- Full setup including configuration
- Separate container for each service
 - web server, database, logic, ...
 - reduced attack surface
- Used in development and deployment



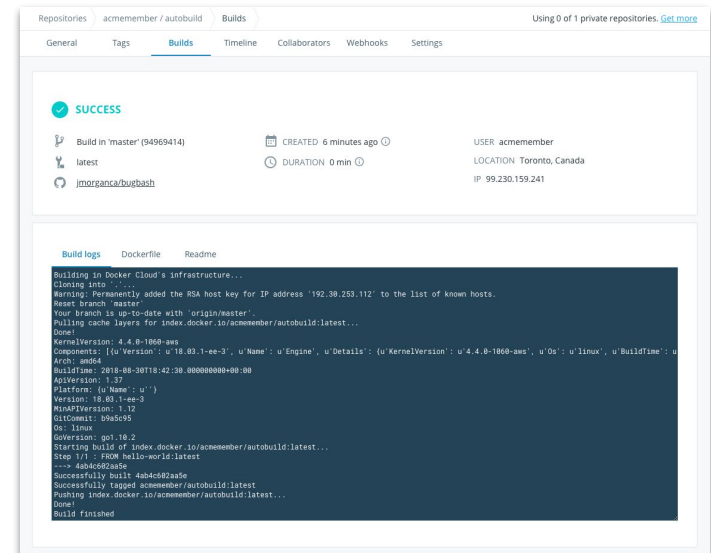

```
FROM ubuntu:18.04
RUN apt-get update \
  && apt-get install -y \
    apt-transport-https \
    ca-certificates \
    curl \
    docker \
    software-properties-common \
    git \
    python \
    python-pip \
    python-dev \
    patchelf \
    python3 \
    python3-pip \
    openjdk-8-jdk \
    locales \
    vim \
  && pip install pipenv \
  && curl -fsSL https://download.docker.com/linux/ubuntu/gpg | apt-key add - \
  && add-apt-repository \
    "deb [arch=amd64] https://download.docker.com/linux/ubuntu \
    $(lsb_release -cs) \
    stable" \
  && apt-get update \
  && apt-get install -y docker-ce \
  && apt-get autoremove -y \
  && apt-get clean \
  && rm -rf /var/lib/apt/lists/* /tmp/* /var/tmp/*
RUN sed -i -e 's/# en_US.UTF-8 UTF-8/en_US.UTF-8 UTF-8/' /etc/locale.gen && \
  locale-gen
ENV LANG en_US.UTF-8
ENV LANGUAGE en_US:en
ENV LC_ALL en_US.UTF-8
```

Docker and DockerHub

- Build an image for each release
- Quickly rollback to stable versions

```
$ docker pull mysql:8.0
```

```
$ docker push christimperley/darjeeling
```



Principle: Rapid Releases and Feedback

- Remove the manual and ceremonial aspects from releases
 - Possibly continuous releases
 - Incremental rollout; quick rollback
- Get feedback on your changes ASAP
 - Continuously measure quality, refine implementation, and rerelease



Travis CI



Principle: Shared Responsibility

- Breakdown the “Wall of Confusion”
- Improve collaboration between dev. and ops. teams
- Reduce “throw it over the fence” syndrome
- Treat failures as a learning experience...



Aside: Postmortems

Example Postmortem

Shakespeare Sonnet++ Postmortem (incident #465)

Date: 2015-10-21

Authors: jennifer, martym, agoogler

Status: Complete, action items in progress

Summary: Shakespeare Search down for 66 minutes during period of very high interest in Shakespeare due to discovery of a new sonnet.

Impact:¹⁶³ Estimated 1.21B queries lost, no revenue impact.

Root Causes:¹⁶⁴ Cascading failure due to combination of exceptionally high load and a resource leak when searches failed due to terms not being in the Shakespeare corpus. The newly discovered sonnet used a word that had never before appeared in one of Shakespeare's works, which happened to be the term users searched for. Under normal circumstances, the rate of task failures due to resource leaks is low enough to be unnoticed.



Trigger: Latent bug triggered by sudden increase in traffic.

Resolution: Directed traffic to sacrificial cluster and added 10x capacity to mitigate cascading failure. Updated index deployed, resolving interaction with latent bug. Maintaining extra capacity until surge in public interest in new sonnet passes. Resource leak identified and fix deployed.

Detection: Borgmon detected high level of HTTP 500s and paged on-call.


Action Items:¹⁶⁵

Action Item	Type	Owner	Bug
Update playbook with instructions for responding to cascading failure	mitigate	jennifer	n/a DONE
Use flux capacitor to balance load between clusters	prevent	martym	Bug 5554823 TODO



<https://blog.codinghorror.com/the-project-postmortem/>
<https://www.developer.com/design/article.php/3637441>
<https://landing.google.com/sre/books/>


Two sides to DevOps



Created by Shocho
from Noun Project

Operations-oriented

- Manage servers automatically
- Easier to identify and fix bugs
- Automatic logging, monitoring, and operations

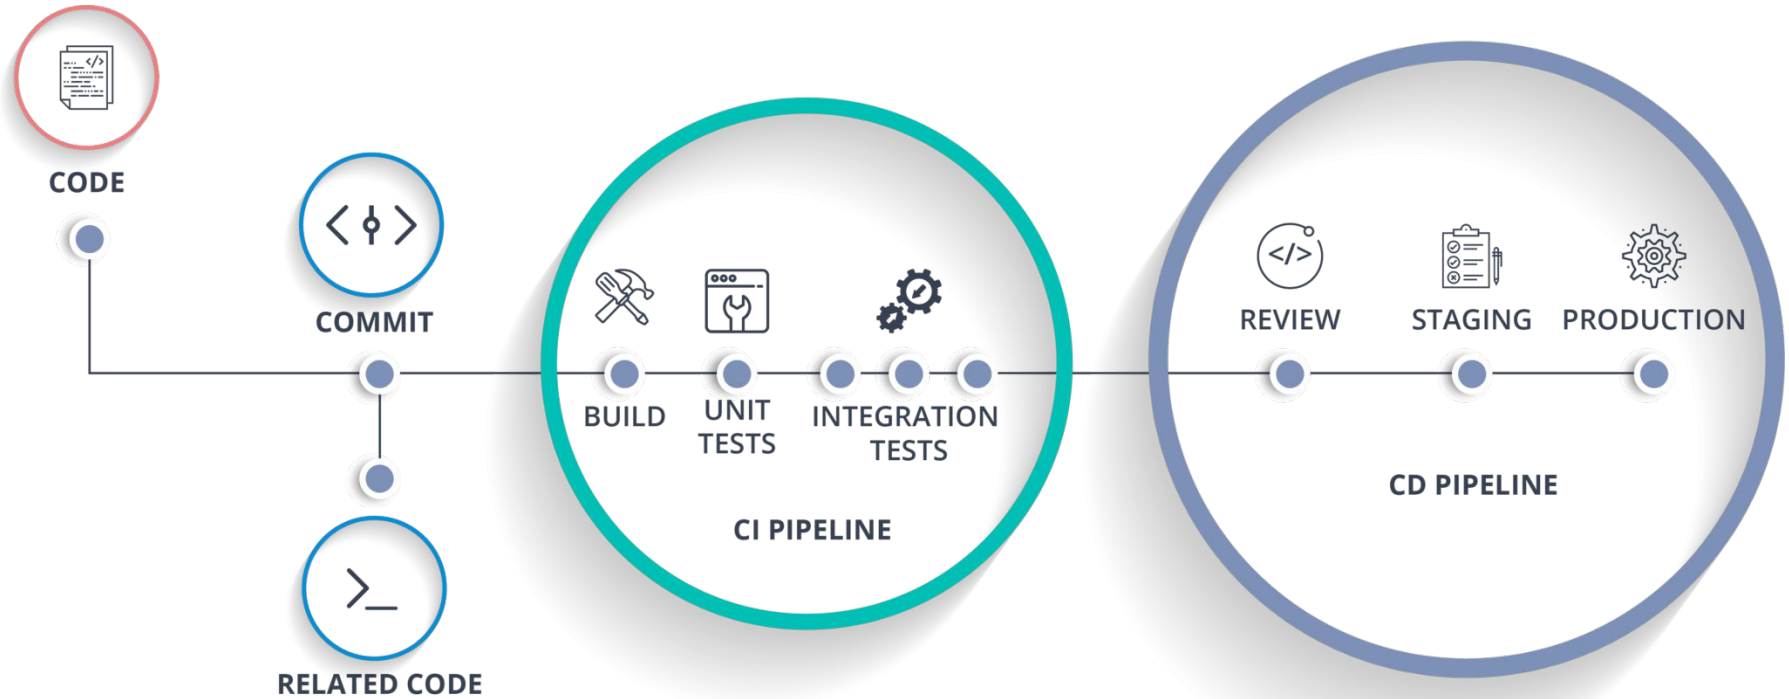


Created by Eucalyp
from Noun Project

Developer-oriented

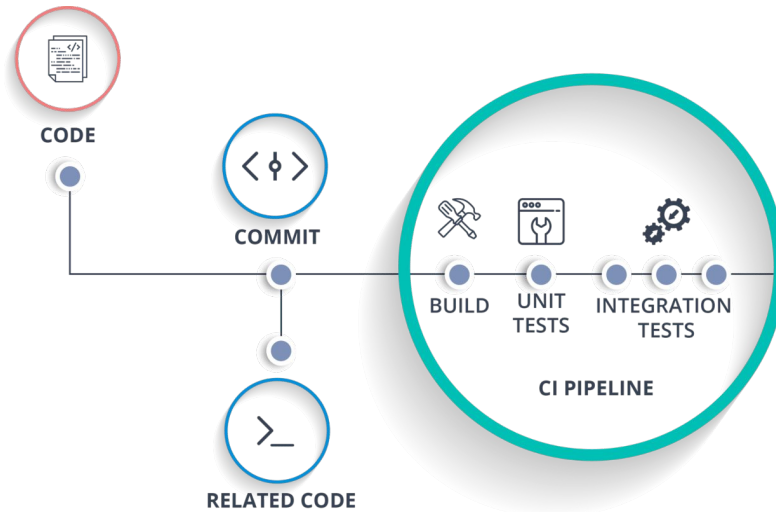
- Agile releases!
- Easier to share and understand code
- Faster onboarding
- Safely push code through CI/CD pipeline

Continuous Integration and Continuous Deployment



<https://dzone.com/articles/learn-how-to-setup-a-cicd-pipeline-from-scratch>

Continuous Integration



Travis CI

JUnit 5



Gradle



SpotBugs

Continuous Integration at Google

Google workflow



- All code is reviewed before commit (by humans and automated tooling)
- Each directory has a set of owners who must approve the change to their area of the repository
- Tests and automated checks are performed before and after commit
- Auto-rollback of a commit may occur in the case of widespread breakage



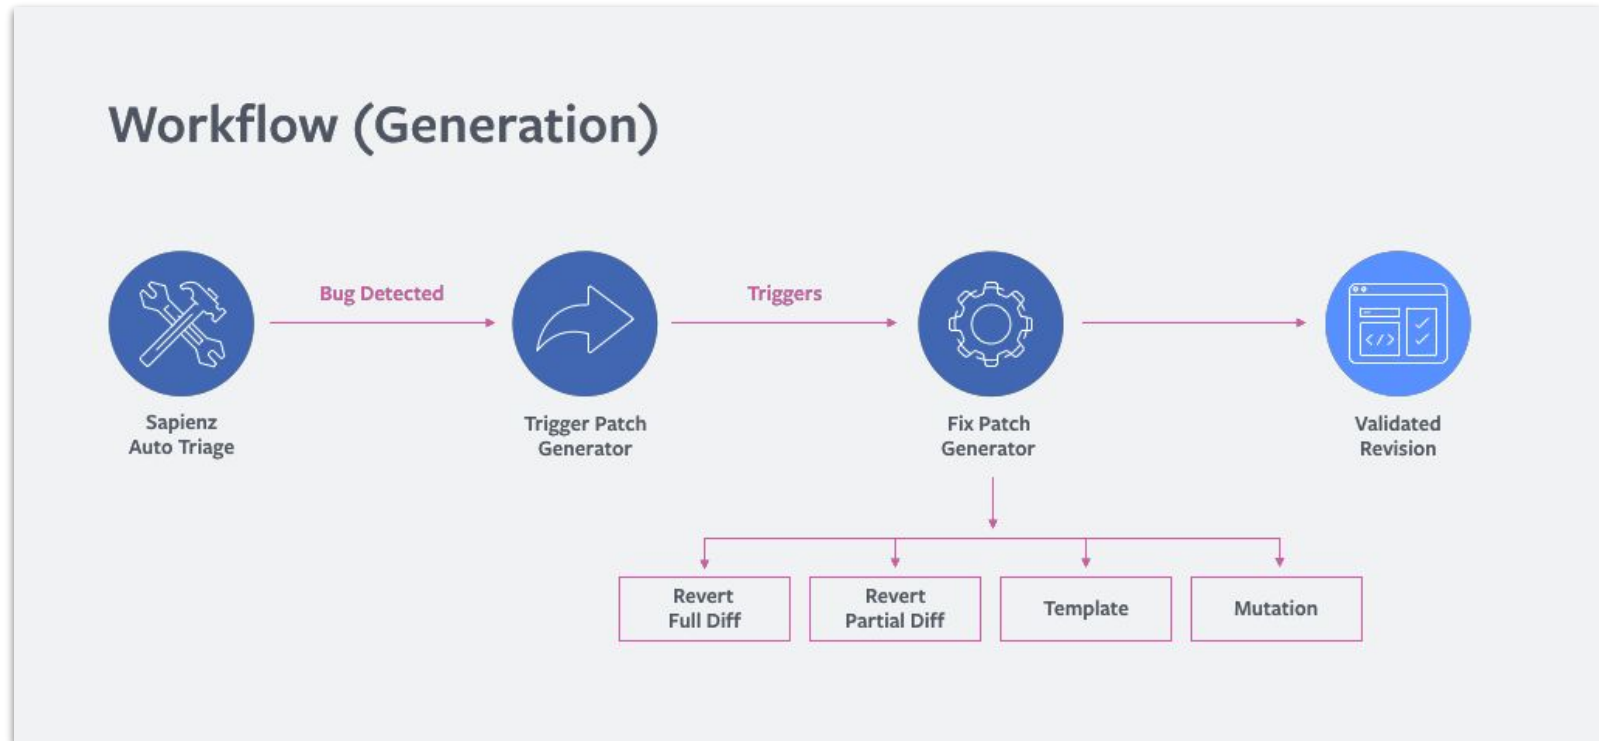
Additional tooling support

Critique	Code review
CodeSearch*	Code browsing, exploration, understanding, and archeology
Tricorder**	Static analysis of code surfaced in Critique, CodeSearch
Presubmits	Customizable checks, testing, can block commit
TAP	Comprehensive testing before and after commit, auto-rollback
Rosie	Large-scale change distribution and management

* See "How Developers Search for Code: A Case Study", In European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering, 2015

** See "Tricorder: Building a program analysis ecosystem". In International Conference on Software Engineering (ICSE), 2015

Aside: Sapienz and SapFix at Facebook



<https://engineering.fb.com/developer-tools/finding-and-fixing-software-bugs-automatically-with-sapfix-and-sapienz/>

Outline

- DevOps and CI/CD
- Large-Scale Version Control
- Release Management

How do you scale to 2 billion lines of code?



Google repository usage

Human users


- 25 thousand Googlers in dozens of offices around the world

On an average workday

- 15 thousand commits by humans
- 30 thousand commits by automated systems
- Billions of file read requests* (800K QPS at daily peak)

*Google recently open-sourced a subset of the internal build system, see www.bazel.io

Millions of changes committed (cumulative)



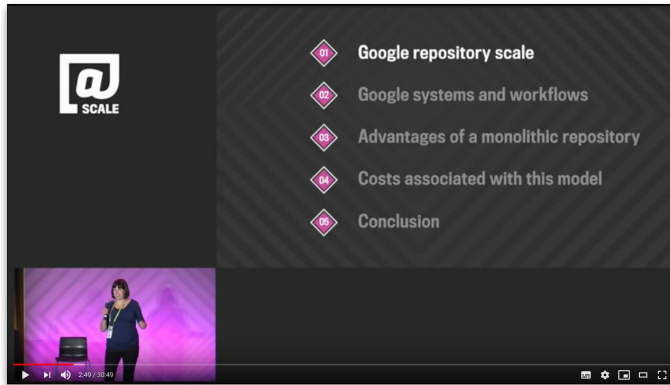
Some perspective

Linux kernel

- 15 million lines of code in 40 thousand files (total)

Google repository

- 15 million lines of code in 250 thousand files *changed per week, by humans*
- 2 billion lines of code, in 9 million source files (total)



Google repository statistics

As of Jan 2015

Total number of files*	1 billion
Number of source files	9 million
Lines of code	2 billion
Depth of history	35 million commits
Size of content	86 terabytes
Commits per workday	45 thousand

*The total number of files includes source files copied into release branches, files that are deleted at the latest revision, configuration files, documentation, and support files.

R. Potvin and J. Levenberg, "The Motivation for a Monolithic Codebase: Why Google stores billions of lines of code in a single repository", in Communications of the ACM, vol. 59, no. 7, 2016.

contributed articles

DOI:10.1145/2854146

Google's monolithic repository provides a common source of truth for tens of thousands of developers around the world.

BY RACHEL POTVIN AND JOSH LEVENBERG

Why Google Stores Billions of Lines of Code in a Single Repository

This article outlines the scale of that codebase and details Google's custom-built monolithic source repository and the reasons the model was chosen. Google uses a homegrown version-control system to host one large codebase visible to, and used by, most of the software developers in the company. This centralized system is the foundation of many of Google's developer workflows. Here, we provide background on the systems and workflows that make feasible managing and working productively with such a large repository. We explain Google's "trunk-based development" strategy and the support systems that structure workflow and keep Google's codebase healthy, including software for static analysis, code cleanup, and streamlined code review.

Google-Scale

Google's monolithic software repository, which is used by 95% of its software developers worldwide, meets the definition of an ultra-large-scale^a system, providing evidence the single-source repository model can be scaled successfully.

The Google codebase includes approximately one billion files and has a history of approximately 35 million commits spanning Google's entire 18-year existence. The repository contains 86TB^a of data, including approximately

^a Total size of uncompressed content, excluding release branches.


» key insights

- Google has shown the monolithic model of source code management can scale to a repository of one billion files, 35 million commits, and tens of thousands of developers.
- Benefits include unified versioning, extensive code sharing, simplified dependency management, atomic changes, large-scale refactoring, collaboration across teams, flexible code ownership, and code visibility.
- Drawbacks include having to create and scale tools for development and execution and maintain code health, as well as potential for codebase complexity (such as unnecessary dependencies).

EARLY GOOGLE EMPLOYEES decided to work with a shared codebase managed through a centralized source control system. This approach has served Google well for more than 16 years, and today the vast majority of Google's software assets continues to be stored in a single, shared repository. Meanwhile, the number of Google software developers has steadily increased, and the size of the Google codebase has grown exponentially (see Figure 1). As a result, the technology used to host the codebase has also evolved significantly.

78 COMMUNICATIONS OF THE ACM | JULY 2016 | VOL. 59 | NO. 7


A recent history of code organization



Monorepo



Single-repo Monolith



Multi-repo



Monolithic repositories (Monorepos)

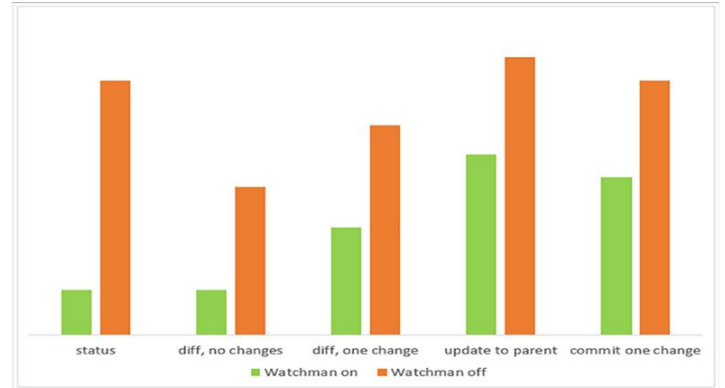
A single version control repository containing multiple:

- Projects
- Applications
- Libraries



POSTED ON JAN 7, 2014 TO CORE DATA, OPEN SOURCE

Scaling Mercurial at Facebook



By Durham Goode Rain



With thousands of commits a week across hundreds of thousands of files, Facebook's main source repository is enormous—many times larger than even the Linux kernel, which checked in at 17 million lines of code and 44,000 files in 2013. Given our size and complexity—and Facebook's practice of shipping code twice a day—improving our source control is one way we help our engineers move fast.

Choosing a source control system


Two years ago, as we saw our repository continue to grow at a staggering rate, we sat down and extrapolated our growth forward a few years. Based on those projections, it appeared likely that our then-current technology, a Subversion server with a Git mirror, would become a productivity bottleneck very soon. We looked at the available options and found none that were both fast and easy to use at scale.

Our code base has grown organically and its internal dependencies are very complex. We could have spent a lot of time making it more modular in a way that would be friendly to a source control tool, but there are a number of benefits to using a single repository. Even at our current scale, we often make large changes throughout our code base, and having a single repository is useful for continuous modernization. Splitting it up would make large, atomic refactorings more difficult. On top of that, the idea that the scaling constraints of our source control system should dictate our code structure just doesn't sit well with us.

We realized that we'd have to solve this ourselves. But instead of building a new system from scratch, we decided to take an existing one and make it scale. Our engineers were comfortable with Git and we

Microsoft | Brian Harry's Blog | Product Blogs | DevOps | Languages | .NET | Platform Development | Data Development | Login

The largest Git repo on the planet


Brian
May 24th, 2017

It's been 3 months since I first wrote about [our efforts to scale Git to extremely large projects and teams](#), with an effort we called "Git Virtual File System". As a reminder, GVFS, together with a set of enhancements to Git, enables Git to scale to VERY large repos by virtualizing both the .git folder and the working directory. Rather than download the entire repo and checkout all the files, it dynamically downloads only the portions you need based on what you use.

A lot has happened and I wanted to give you an update. Three months ago, GVFS was still a dream. I don't mean it didn't exist – we had a concrete implementation, but rather, it was unproven. We had validated on some big repos but we hadn't rolled it out to any meaningful number of engineers so we had only conviction that it was going to work. Now we have proof.

Today, I want to share our results. In addition, we're announcing the next steps in our GVFS journey for customers, including expanded open sourcing to start taking contributions and improving how it works for us at Microsoft, as well as for partners and customers.


Windows is live on Git

Over the past 3 months, we have largely completed the rollout of Git/GVFS to the Windows team at Microsoft.

As a refresher, the Windows code base is approximately 3.5M files and, when checked in to a Git repo, results in a repo of about 300GB. Further, the Windows team is about 4,000 engineers and the engineering system produces 1,760 daily "lab builds" across 440 branches in addition to thousands of pull request validation builds. All 3 of the dimensions (file count, repo size and activity), independently, provide daunting scaling challenges and taken together they make it unbelievably challenging to create a great experience. Before the move to Git, in Source Depot, it was spread across 40+ depots and we had a tool to manage operations that spanned them.

As of my writing 3 months ago, we had all the code in one Git repo, a few hundred engineers using it and a small fraction (<10%) of the daily build load. Since then, we have rolled out in waves across the engineering team.

The first, and largest, jump happened on March 22nd when we rolled out to the Windows OneCore team of about 2,000 engineers. Those 2,000 engineers worked in Source Depot on Friday, went home for the weekend and came back Monday morning to a new experience based on Git. People on my team were holding their breath that whole weekend, praying we weren't going to be pummeled by a mob of angry engineers who showed up Monday unable to get any work done.



Monorepos are also used by open source projects

foursquare / fsqio

Watch 129 Star 199 Fork 43

Code Issues 17 Pull requests 2 Actions Projects 0 Wiki Security Insights

A monorepo that holds all of Foursquare's opensource projects

pants foursquare monorepo mongodb rogue scala

1,096 commits 1 branch 0 packages 23 releases 22 contributors Apache-2.0

Branch: master New pull request Create new file Upload files Find file Clone or download

Eric-Arellano and OniOni Add scripts to modernize to Python 3 Latest commit 332e7a2 on Jun 13

- 3rdparty Update 3rd party map. 3 months ago
- build-support/fsqio pants build committing publish data for push of io.fsq#buildgen-emit... 6 months ago
- scripts/fsqio Add scripts to modernize to Python 3 3 months ago
- src Move "--userconfig" to webpack resolve as it's an npm option. 3 months ago
- test fix buildgen scalac plugin not handling _root_ (#14562) 6 months ago
- .gitignore Add pytest_cache to fsq.io gitignore (#6428) 2 years ago
- .travis.yml Generalize sapling projects under "opensource" and check-in fsqio bui... 7 months ago
- BUILD.opensource Pin scala minor version (#9262) 11 months ago
- BUILD.tools Drop a BUILD.tools in Fsq.io. 3 years ago
- CLA.md Move deployed files to consolidated directory. 4 years ago
- CONTRIBUTING.md Add a sapling quiver build (#14305) 7 months ago
- LICENSE.md Update copyright info in Fsq.io LICENSE. 4 years ago
- README.md Remove pom-resolve code and references (#3927) 2 years ago
- pants Add Fsqio build tooling to Fsq.io (#8728) 11 months ago
- pants-travis-ci.ini Delete the custom JDK install for Fsq.io CI. (#744) 2 years ago
- pants.ini Add publishing config worked out in shading repo and docs (x2) (#15049) 6 months ago
- upkeep Support for WebpackSubsystem and add ivy.xml 7 months ago

README.md

stellar / go

Watch 71 Star 625 Fork 342

Code Issues 205 Pull requests 23 Actions Projects 0 Security Insights

Stellar's public monorepo of go code https://stellar.org/developers

stellar blockchain cryptocurrency horizon

1,743 commits 6 branches 0 packages 78 releases 77 contributors View license

Branch: master New pull request Create new file Upload files Find file Clone or download

leghmculloch services/ticker: set version for release v1.2.0 (#1962) Latest commit 4c28b16 1 hour ago

- .circleci .circleci: remove go mod checks for go 1.12 (#1944) 5 days ago
- .github doc: lower the barrier-to-entry of writing great pr descriptions (#1889) 21 days ago
- .vscode cleanup 2 years ago
- address fixed assert in main_test.go 2 years ago
- amount AllowOverInt64 in amount.IntStringToAmount (#609) last year
- build Remove unnecessary whitespace and fix typo in godoc comments (#1866) last month
- clients Make horizonclient examples fully qualified (#1867) 29 days ago
- crc16 Initial import from go-stellar-base 3 years ago
- docs/reference docs: fix link to stellartoml (#1743) 2 months ago
- exp exp/hubble: write entries to elasticsearch (#1932) 6 hours ago
- handlers Use go vet in CI (#611) 10 months ago
- hash Initial import from go-stellar-base 3 years ago
- keypair keypair: add MustRandom (#1907) 8 days ago
- meta handle bumpseq effects last year
- network Run staticcheck in CI (#1453) 5 months ago
- price Ignore offers which result in overflow errors in path finding search (#... 2 months ago
- protocols services/horizon: Update "Action needed" tags (#1929) 6 days ago
- services services/ticker: set version for release v1.2.0 (#1962) 1 hour ago
- strkey Merge pull request #1579 from stellar/master 3 months ago
- support all: update all http.Servers to use a ReadTimeout 4 hours ago
- tools tools/horizon-ops: Replace requests from FLD access log (#1789) 2 months ago

Why do these companies use monorepos?



Benefits of Monorepos

- Cheaper code reuse
 - Extract reusable code into a new component
 - Easily use that code from elsewhere! No need for more repos.
- Browse, read, and search through the entire codebase
 - Works with grep, IDEs, and special tools out of the box
- Atomic refactorings with a single commit
 - Switch from an old API to a new API in a single commit
- Easier to test, debug, review, and deploy projects that span multiple applications
 - Easier to collaborate across projects and teams.
 - No more internal dependency management!

Drawbacks of Monorepos

- Require collective responsibility for team and developers
- Require trunk-based development
 - More on that later...
- Force you to have only one version of everything
- Scalability requirements for the repository
- Can be hard to deal with updates around things like security issues
- Build and test bloat without very smart build system
- Slow VCS without very smart system
- Permissions?

Outline

- DevOps and CI/CD
- Large-Scale Version Control
- Release Management

How and when should software be released?

Google repository

- 15 million lines of code in 250 thousand files *changed per week, by humans*
- 2 billion lines of code, in 9 million source files (total)

Principle: Quick to Deploy; Slow to Release

*“Get your **** together; fix it in production.”*

Chuck Rossi, former Release Engineering Director at Facebook

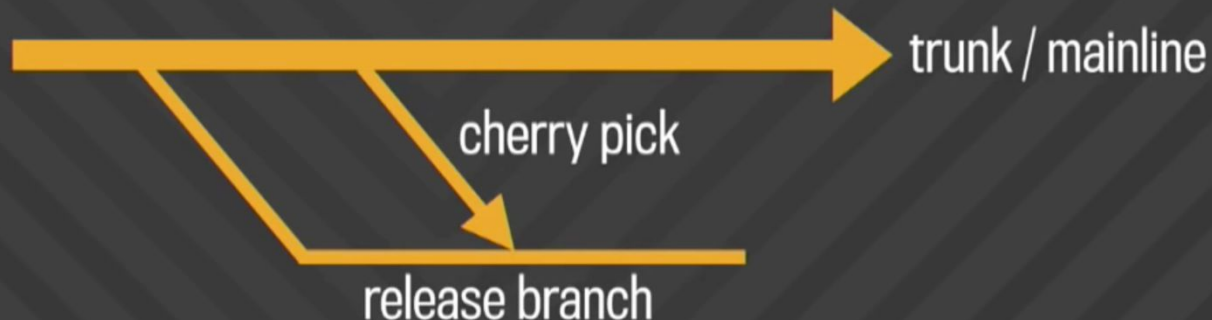


Trunk-based development at Google

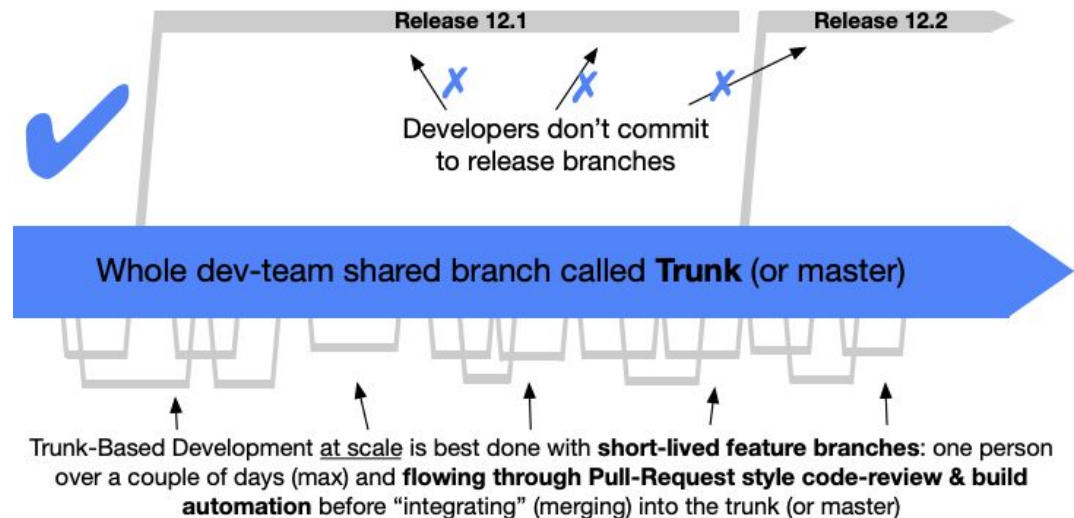

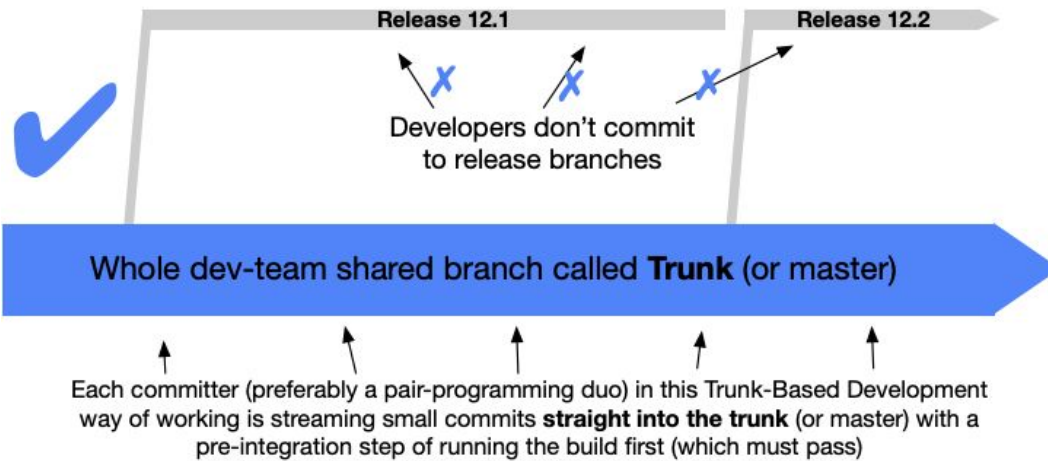
Trunk-based development

Combined with a centralized repository, this defines the monolithic model


- Piper users work at “head”, a consistent view of the codebase
- All changes are made to the repository in a single, serial ordering
- There is no significant use of branching for development
- Release branches are cut from a specific revision of the repository



Trunk-based development

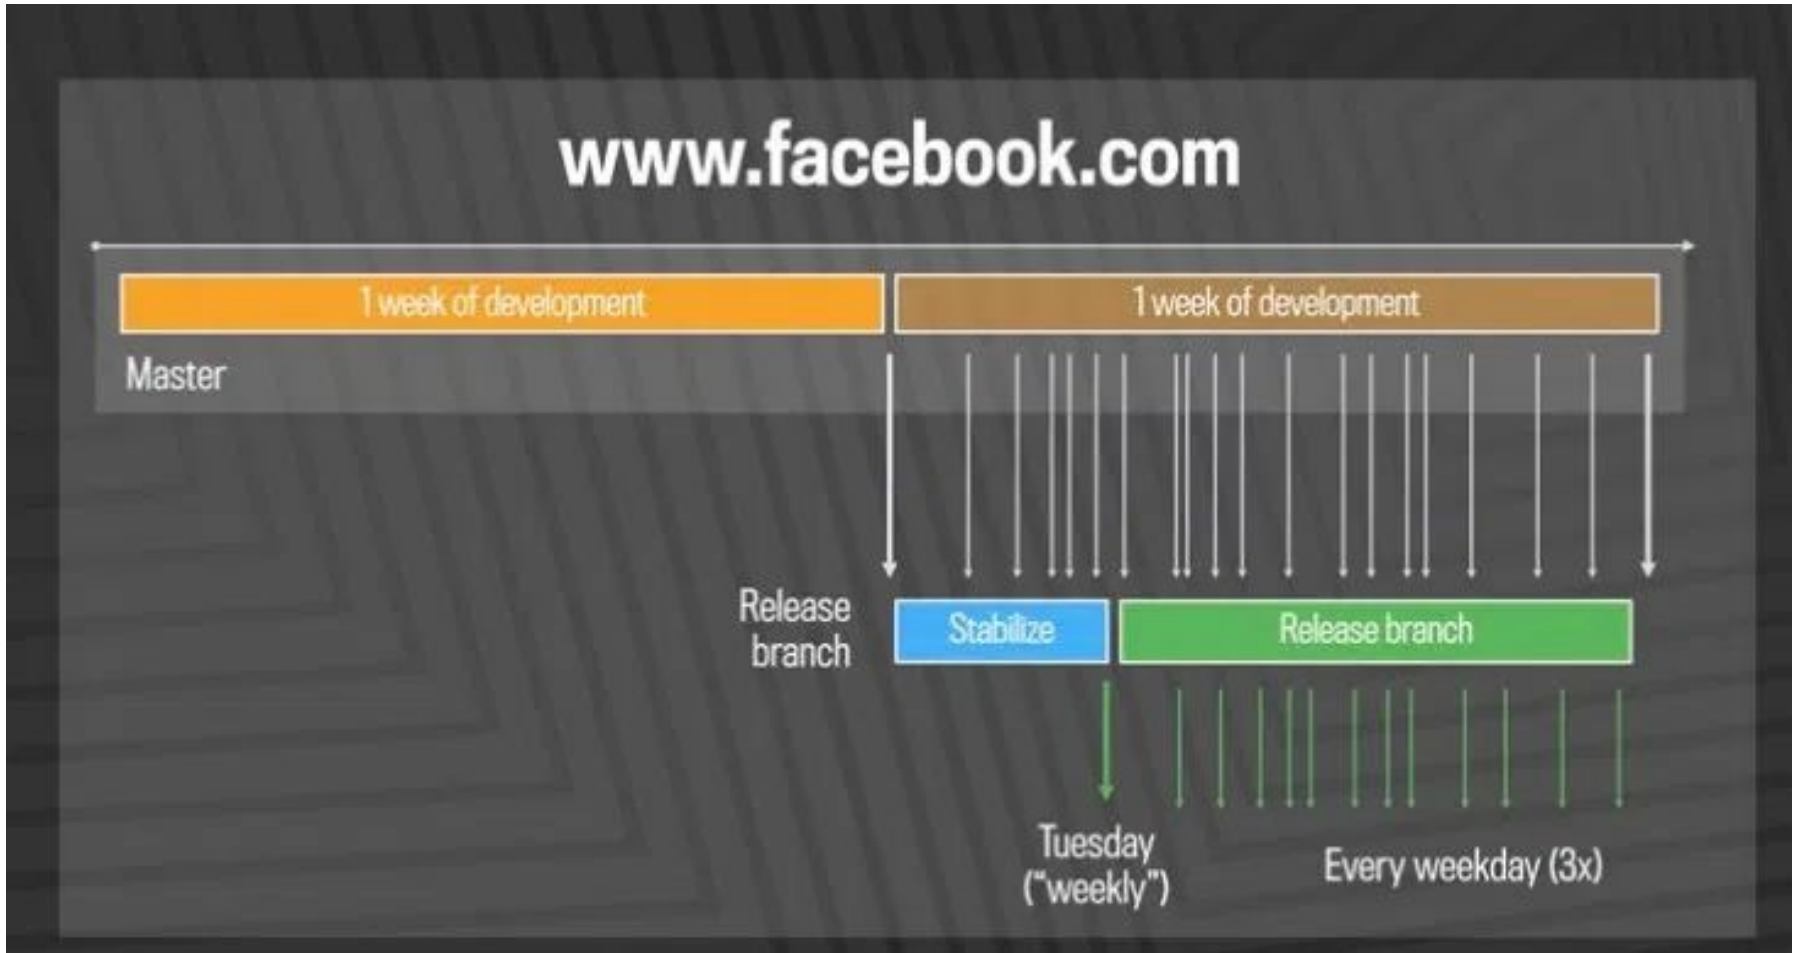


Cherrypicking




<https://www.atlassian.com/blog/git/the-essence-of-branch-based-workflows>

Fresh release branch every week




<https://engineering.fb.com/web/rapid-release-at-massive-scale/>

The number of commits in a branch cut became unsustainable



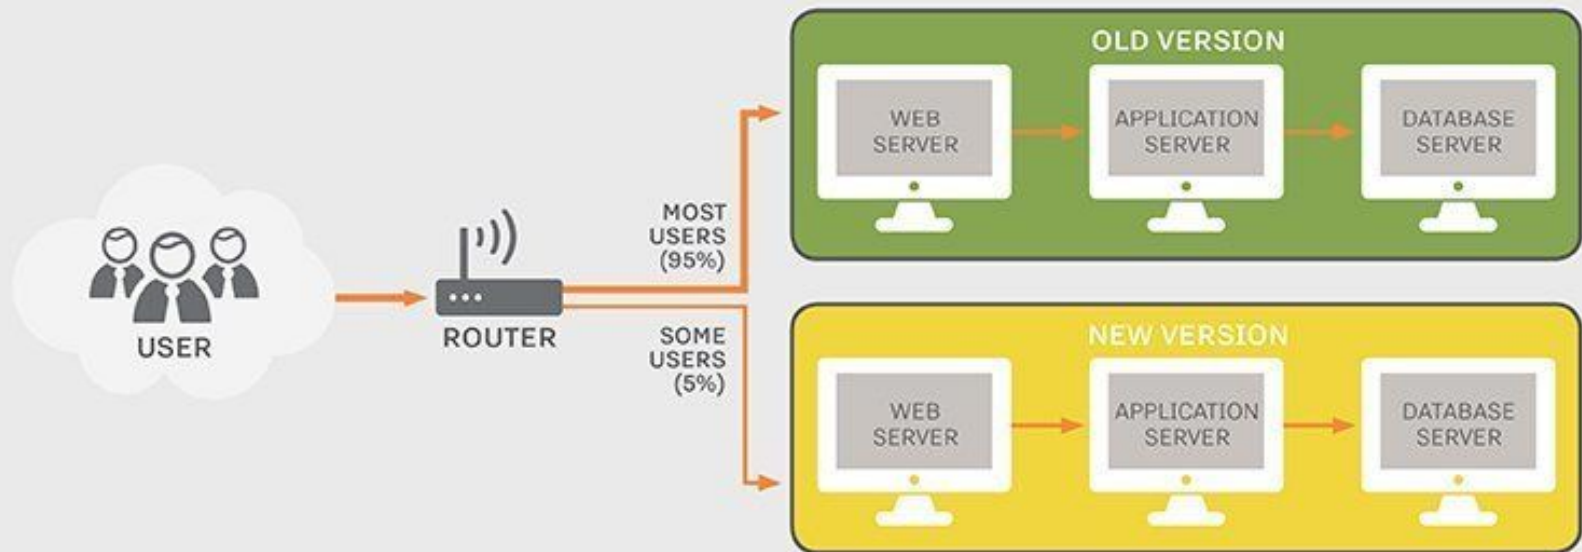
Quasi-continuous push from master (1,000+ devs, 1,000 diffs/day);
10 pushes/day



Principle: Every feature is an experiment

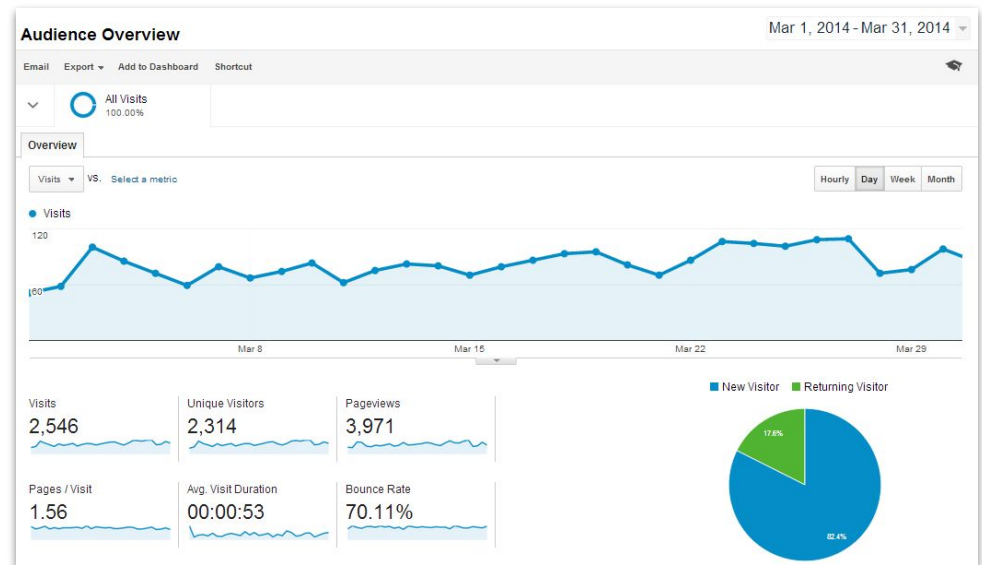
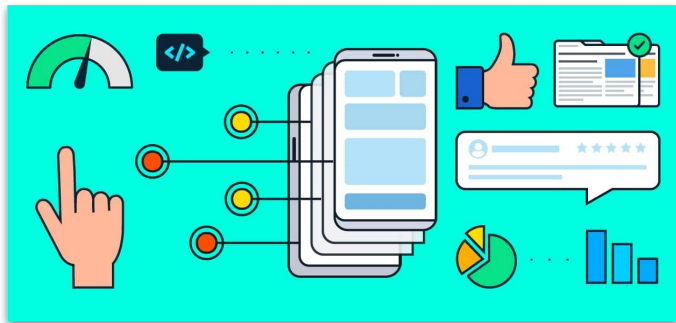


CANARY TESTING

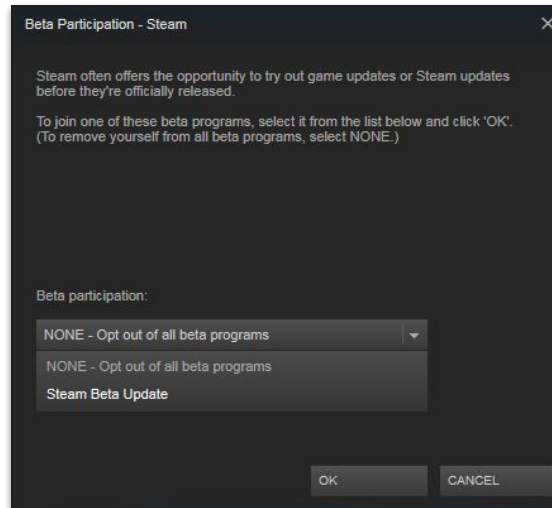


Dark Launching

- Similar to canary testing
- Focuses on user response to frontend changes rather than performance of backend
- Measure user response via *metrics: engagement, adoption*


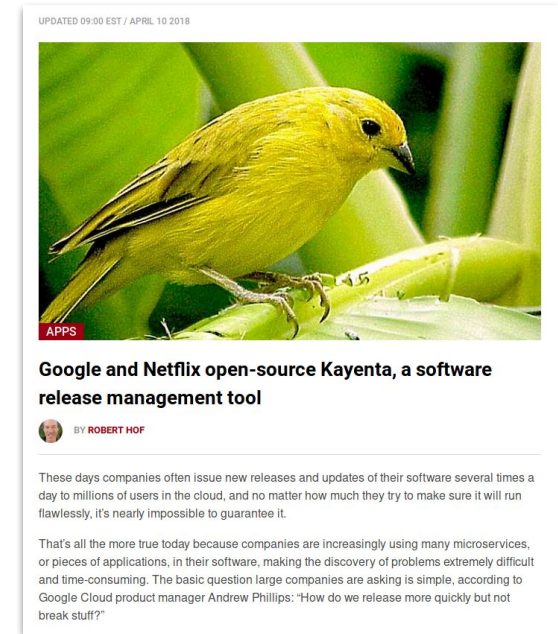
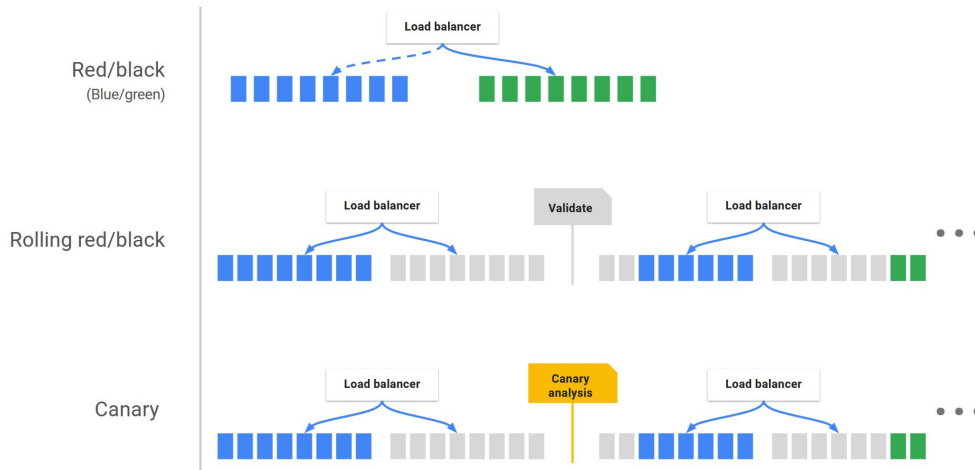


Aside: Opt-In Beta



Automated canary analysis at Netflix

- ~60,000 configuration changes per day, ~4000 commits per day
- Bake an Amazon Machine Image (AMI) for each commit
- Deploy via Spinnaker and Kayenta
- Perform automated canary analysis.
 - If okay, switch to new version.
 - If bad, rollback to old version.




<https://medium.com/netflix-techblog/automated-canary-analysis-at-netflix-with-kayenta-3260bc7acc69>

<https://octopus.com/blog/blue-green-red-black>

<https://siliconangle.com/2018/04/10/google-netflix-open-source-kayenta-software-release-management-tool/>

Control deployments at run-time using feature flags



GateKeeper

Project: 64bit_rollout

Rank	Move	Group	Description
1	▲▼	all users	(delete)

New Restraint

Restraint Type: Age - Older

- Age - Older
- Age - Younger
- Application
- Browser
- Code Location
- Country
- Datacenter
- Is Employee
- Friend Count - Less
- Friend Count - More
- Gatekeeper project ID
- Locale
- Network
- OS
- Remote IP
- Server IP
- Server Time - After
- Server Time - Before

Save Cancel

WHITELIST ME

BLACKLIST ME

On
vuvtxzdzqrp

Alpha n/a

Alpha Def. n/a

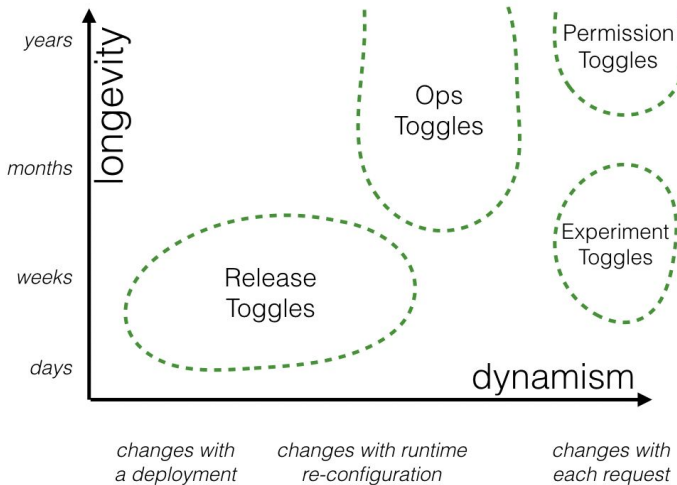
Updated 4/21/09 3:23:04pm

Console none

Name

Description 64 bit rollout

Needs Flush No



<https://martinfowler.com/articles/feature-toggles.html>

<https://docs.microsoft.com/en-us/azure/devops/migrate/phase-features-with-feature-flags?view=azure-devops>

Warning! Feature flags can be dangerous



Knightmare: A DevOps Cautionary Tale

D7 DevOps April 17, 2014 6 Minutes

I was speaking at a conference last year on the topics of DevOps, Configuration as Code, and Continuous Delivery and used the following story to demonstrate the importance making deployments fully automated and repeatable as part of a DevOps/Continuous Delivery initiative. Since that conference I have been asked by several people to share the story through my blog. This story is true – this really happened. This is my telling of the story based on what I have read (I was not involved in this).

This is the story of how a company with nearly \$400 million in assets went bankrupt in 45-minutes because of a failed deployment.

In laymen's terms, Knight Capital Group realized a \$460 million loss in 45-minutes. Remember, Knight only has \$365 million in cash and equivalents. **In 45-minutes Knight went from being the largest trader in US equities and a major market maker in the NYSE and NASDAQ to bankrupt.**

Summary

- DevOps brings development and operations together
 - Automation, Automation, Automation
 - Infrastructure as code
- Release management
 - Versioning and branching strategies
- Continuous deployment is increasingly common
- Exploit opportunities of continuous deployment; perform testing in production and quickly rollback
 - Experiment, measure, and improve