

Software Engineering at Google Scale

Marija Mikic
Google Los Angeles



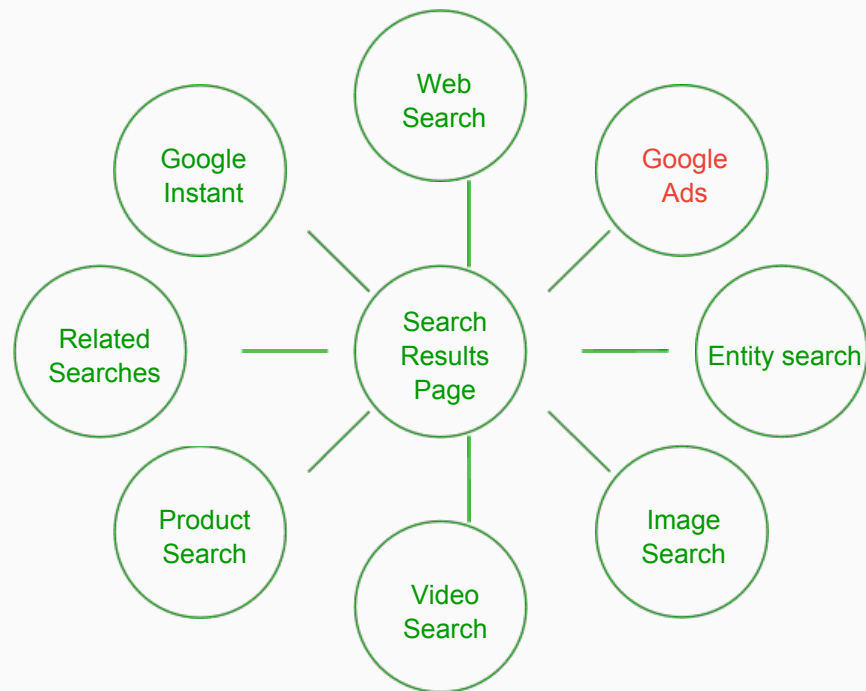
Background

- Ph.D. in Software Engineering in 2004
 - Software architecture
 - System deployment and adaptation
- @ Google for ~12 years
- Worked on several projects within Ads
- Managing since 2008
- My team
 - 27 engineers
 - 3 statisticians
 - 10 linguists
 - PM

Google by the numbers

- **~50,000,000,000** - The number Android app downloads
- **~1,000,000,000** - The number of active Android users
- **~500,000,000** - The number of active gmail accounts
- **~5,000,000** - The amount of miles Google cars have driven to acquire map information for Project Truth
- **~4,000,000** - The amount of queries Google receives per min
- **~1,300,000** - The number of apps in the Google Play Store
- **~840,000** - The amount of miles Google self-driving cars have driven to date
- **~90** - The number of languages supported by Google Translate
- **1** - **The Google user experience**

One Google experience



Many systems work together and produce a single, coherent result within hundreds of milliseconds!

One Google experience

Google

revenant
revenant meaning
revenant showtimes
revenant review
About 58,800,000 results (0.58 seconds)

Showing results for **revenant**
Search instead for reverant

Showtimes for The Revenant

Wed, Apr 13 | Thu, Apr 14 | Fri, Apr 15 | Sat, Apr 16

All times Morning Afternoon Evening Night

ArcLight Santa Monica - Map
4:30pm 6:15 9:30

Valley Plaza 6 - North Hollywood - Map
6:40pm 10:10

Pacific Winnetka All Stadium 21 - Map
4:05pm 7:20

All times are in PT

More showtimes

The Revenant

2015 · Drama film/Thriller · 2h 36m

Play trailer on YouTube

8.1/10 IMDb | 82% Rotten Tomatoes | 76% Metacritic

While exploring the uncharted wilderness in 1823, frontiersman Hugh Glass (Leonardo DiCaprio) sustains life-threatening injuries from a brutal bear attack. When a member (Tom Hardy) of his hunting team kills his young son (Forrest Goodluck) and leaves him for dead, Glass must utilize his survival sk... More

Release date: December 25, 2015 (USA)
Director: Alejandro González Iñárritu
Adapted from: The Revenant
Cinematography: Emmanuel Lubezki
Nominations: Academy Award for Best Picture, More

Available on

- Google Play Movies & TV \$14.99
- YouTube From \$14.99
- iTunes From \$14.99

More services

Critic reviews

View 6+ more

It's an imposing vision, to be sure, but also an inflated and emotionally stunted one, despite an anchoring performance of ferocious 200% commitment from Leonardo DiCaprio. Full review
Justin Chang - Variety

The Revenant (2015) - IMDb

www.imdb.com/title/tt1663202/ · Internet Movie Database ·
★ ★ ★ ★ Rating: 8.1/10 - 321,567 votes
A frontiersman on a fur trading expedition in the 1820s fights for survival after being mauled by a bear and left for dead by members of his own hunting team. ... Leonardo DiCaprio, Tom Hardy, Will Poulter. ... Q: How is Leonardo DiCaprio's role in "The Revenant" different from his usual ...
Full Cast & Crew - Paul Anderson - John Fitzgerald - 2016

The Revenant (2015 film) - Wikipedia, the free encyclopedia

https://en.wikipedia.org/wiki/The_Revenant_(2015_film) · Wikipedia ·
The Revenant is a 2015 American historical drama directed, co-produced, and co-written by Alejandro G. Iñárritu. The screenplay by Iñárritu and Mark L. Smith is ...
Hugh Glass - Kristoffer Joner - Trapping - Man in the Wilderness

Revenant - Wikipedia, the free encyclopedia

https://en.wikipedia.org/wiki/Revenant · Wikipedia ·
A **revenant** is a visible ghost or animated corpse that is believed to have returned from the grave to terrorize the living. The word **revenant** is derived from the ...
Nachzehrer - Revenant (disambiguation) - Wiedergänger - Gjenganger

The Revenant (2015) - Rotten Tomatoes

www.rottentomatoes.com/films/the_revenant_2015/ · Rotten Tomatoes ·

One Google experience

Google

revenant
revenant meaning
revenant showtimes
revenant review

About 58,800,000 results (0.58 seconds)

Showing results for **revenant**
Search instead for reverant

Showtimes for The Revenant

Wed, Apr 13 | Thu, Apr 14 | Fri, Apr 15 | Sat, Apr 16

All times Morning Afternoon Evening Night

ArcLight Santa Monica - Map
4:30pm 6:15 9:30

Valley Plaza 6 - North Hollywood - Map
6:40pm 10:10

Pacific Winnetka All Stadium 21 - Map
4:05pm 7:20

All times are in PT

More showtimes

The Revenant

2015 · Drama film/Thriller · 2h 36m

Play trailer on YouTube

8.1/10 IMDb | 82% Rotten Tomatoes | 76% Metacritic

While exploring the uncharted wilderness in 1823, frontiersman Hugh Glass (Leonardo DiCaprio) sustains life-threatening injuries from a brutal bear attack. When a member (Tom Hardy) of his hunting team kills his young son (Forrest Goodluck) and leaves him for dead, Glass must utilize his survival sk... More

Release date: December 25, 2015 (USA)
Director: Alejandro González Iñárritu
Adapted from: The Revenant
Cinematography: Emmanuel Lubezki
Nominations: Academy Award for Best Picture, More

Available on

Google Play Movies & TV	\$14.99
YouTube	From \$14.99
iTunes	From \$14.99

More services

Critic reviews

View 6+ more

It's an imposing vision, to be sure, but also an inflated and emotionally stunted one, despite an anchoring performance of ferocious 200% commitment from Leonardo DiCaprio. Full review
Justin Chang - Variety

The Revenant (2015) - IMDb

www.imdb.com/title/tt1663202/ · Internet Movie Database ·
★ ★ ★ ★ Rating: 8.1/10 - 321,567 votes

A frontiersman on a fur trading expedition in the 1820s fights for survival after being mauled by a bear and left for dead by members of his own hunting team. ... Leonardo DiCaprio, Tom Hardy, Will Poulter. ... Q: How is Leonardo DiCaprio's role in "The Revenant" different from his usual ...
Full Cast & Crew - Paul Anderson - John Fitzgerald - 2016

The Revenant (2015 film) - Wikipedia, the free encyclopedia

https://en.wikipedia.org/wiki/The_Revenant_(2015_film) · Wikipedia ·
The Revenant is a 2015 American historical drama directed, co-produced, and co-written by Alejandro G. Iñárritu. The screenplay by Iñárritu and Mark L. Smith is ...
Hugh Glass - Kristoffer Joner - Trapping - Man in the Wilderness

Revenant - Wikipedia, the free encyclopedia

https://en.wikipedia.org/wiki/Revenant · Wikipedia ·
A **revenant** is a visible ghost or animated corpse that is believed to have returned from the grave to terrorize the living. The word **revenant** is derived from the ...
Nachzehrer - Revenant (disambiguation) - Wiedergänger - Gjenganger

The Revenant (2015) - Rotten Tomatoes

www.rottentomatoes.com/films/the_revenant_2015/ · Rotten Tomatoes ·

One Google experience

Google

revenant
revenant meaning
revenant showtimes
revenant review

About 58,800,000 results (0.58 seconds)

Showing results for **revenant**
Search instead for reverant

Showtimes for The Revenant

Wed, Apr 13	Thu, Apr 14	Fri, Apr 15	Sat, Apr 16
All times Morning Afternoon Evening Night			
ArcLight Santa Monica - Map			
4:30pm	6:15	9:30	
Valley Plaza 6 - North Hollywood - Map			
6:40pm	10:10		
Pacific Winnetka All Stadium 21 - Map			
4:05pm	7:20		

[More showtimes](#)

The Revenant

R 2015 · Drama film/Thriller · 2h 36m

[Play trailer on YouTube](#)

8.1/10 IMDb	82% Rotten Tomatoes	76% Metacritic
----------------	------------------------	-------------------

While exploring the uncharted wilderness in 1823, frontiersman Hugh Glass (Leonardo DiCaprio) sustains life-threatening injuries from a brutal bear attack. When a member (Tom Hardy) of his hunting team kills his young son (Forrest Goodluck) and leaves him for dead, Glass must utilize his survival sk... [More](#)

Release date: December 25, 2015 (USA)
Director: Alejandro González Iñárritu
Adapted from: The Revenant
Cinematography: Emmanuel Lubezki
Nominations: Academy Award for Best Picture, More

Available on

Google Play Movies & TV	\$14.99
YouTube	From \$14.99
iTunes	From \$14.99

[More services](#)

Critic reviews

[View 6+ more](#)

Its an imposing vision, to be sure, but also an inflated and emotionally stunted one, despite an anchoring performance of ferocious 200% commitment from Leonardo DiCaprio. Full review
Justin Chang - Variety

The Revenant (2015) - IMDb

www.imdb.com/title/tt1663202/ · Internet Movie Database ·

★ ★ ★ ★ ★ Rating: 8.1/10 - 321,567 votes

A frontiersman on a fur trading expedition in the 1820s fights for survival after being mauled by a bear and left for dead by members of his own hunting team. ... Leonardo DiCaprio, Tom Hardy, Will Poulter. ... Q: How is Leonardo DiCaprio's role in "The Revenant" different from his usual ...
Full Cast & Crew - Paul Anderson - John Fitzgerald - 2016

The Revenant (2015 film) - Wikipedia, the free encyclopedia

[https://en.wikipedia.org/wiki/The_Revenant_\(2015_film\)](https://en.wikipedia.org/wiki/The_Revenant_(2015_film)) · Wikipedia ·

The Revenant is a 2015 American historical drama directed, co-produced, and co-written by Alejandro G. Iñárritu. The screenplay by Iñárritu and Mark L. Smith is ... Hugh Glass - Kristoffer Joner - Trapping - Man in the Wilderness

Revenant - Wikipedia, the free encyclopedia



<https://en.wikipedia.org/wiki/Revenant> · Wikipedia ·

A **revenant** is a visible ghost or animated corpse that is believed to have returned from the grave to terrorize the living. The word **revenant** is derived from the ...
Nachzehrer - Revenant (disambiguation) - Wiedergänger - Gjenganger

The Revenant (2015) - Rotten Tomatoes

www.rottentomatoes.com/films/the_revenant_2015/ · Rotten Tomatoes ·

One Google experience

Google **revenant**  

revenant
revenant meaning
revenant showtimes
revenant review

About 58,800,000 results (0.58 seconds)

Showing results for **revenant**
Search instead for reverant

Showtimes for The Revenant

Wed, Apr 13 | Thu, Apr 14 | Fri, Apr 15 | Sat, Apr 16


All times Morning Afternoon Evening Night

ArcLight Santa Monica - Map
4:30pm 6:15 9:30


Valley Plaza 6 - North Hollywood - Map
6:40pm 10:10


Pacific Winnetka All Stadium 21 - Map
4:05pm 7:20

All times are in PT

 More showtimes

The Revenant

 2015 · Drama film/Thriller · 2h 36m





 Play trailer on YouTube

8.1/10 IMDb | 82% Rotten Tomatoes | 76% Metacritic

While exploring the uncharted wilderness in 1823, frontiersman Hugh Glass (Leonardo DiCaprio) sustains life-threatening injuries from a brutal bear attack. When a member (Tom Hardy) of his hunting team kills his young son (Forrest Goodluck) and leaves him for dead, Glass must utilize his survival sk... More

Release date: December 25, 2015 (USA)
Director: Alejandro González Iñárritu
Adapted from: The Revenant
Cinematography: Emmanuel Lubezki
Nominations: Academy Award for Best Picture, More


Available on


 Google Play Movies & TV	\$14.99
 YouTube	From \$14.99
 iTunes	From \$14.99
 More services	


Critic reviews


View 6+ more

It's an imposing vision, to be sure, but also an inflated and emotionally stunted one, despite an anchoring performance of ferocious 200% commitment from Leonardo DiCaprio. Full review
Justin Chang - Variety

The Revenant (2015) - IMDb
www.imdb.com/title/tt1663202/  Internet Movie Database -
★ ★ ★ ★ Rating: 8.1/10 - 321,567 votes
A frontiersman on a fur trading expedition in the 1820s fights for survival after being mauled by a bear and left for dead by members of his own hunting team. ... Leonardo DiCaprio, Tom Hardy, Will Poulter. ... Q: How is Leonardo DiCaprio's role in "The Revenant" different from his usual ...
Full Cast & Crew - Paul Anderson - John Fitzgerald - 2016

The Revenant (2015 film) - Wikipedia, the free encyclopedia
https://en.wikipedia.org/wiki/The_Revenant_(2015_film)  Wikipedia -
The Revenant is a 2015 American historical drama directed, co-produced, and co-written by Alejandro G. Iñárritu. The screenplay by Iñárritu and Mark L. Smith is ...
Hugh Glass - Kristoffer Joner - Trapping - Man in the Wilderness

Revenant - Wikipedia, the free encyclopedia
https://en.wikipedia.org/wiki/Revenant  Wikipedia -
A revenant is a visible ghost or animated corpse that is believed to have returned from the grave to terrorize the living. The word revenant is derived from the ...
Nachzehrer - Revenant (disambiguation) - Wiedergänger - Gjenganger

The Revenant (2015) - Rotten Tomatoes
www.rottentomatoes.com/films/the_revenant_2015/  Rotten Tomatoes -

One Google experience

Google

revenant
revenant meaning
revenant showtimes
revenant review

About 58,800,000 results (0.58 seconds)

Showing results for **revenant**
Search instead for revenant

Showtimes for The Revenant

Wed, Apr 13 | Thu, Apr 14 | Fri, Apr 15 | Sat, Apr 16

All times Morning Afternoon Evening Night

ArcLight Santa Monica - Map
4:30pm 6:15 9:30

Valley Plaza 6 - North Hollywood - Map
6:40pm 10:10

Pacific Winnetka All Stadium 21 - Map
4:05pm 7:20

All times are in PT

More showtimes

The Revenant (2015) - IMDb

www.imdb.com/title/tt1663202/ - Internet Movie Database -
★★★★★ Rating: 8.1/10 - 321,567 votes

A frontiersman on a fur trading expedition in the 1820s fights for survival after being mauled by a bear and left for dead by members of his own hunting team. ... Leonardo DiCaprio, Tom Hardy, Will Poulter. ... Q: How is Leonardo DiCaprio's role in "The Revenant" different from his usual ...
Full Cast & Crew - Paul Anderson - John Fitzgerald - 2016

The Revenant (2015 film) - Wikipedia, the free encyclopedia

[https://en.wikipedia.org/wiki/The_Revenant_\(2015_film\)](https://en.wikipedia.org/wiki/The_Revenant_(2015_film)) - Wikipedia -
The Revenant is a 2015 American historical drama directed, co-produced, and co-written by Alejandro G. Iñárritu. The screenplay by Iñárritu and Mark L. Smith is ...
Hugh Glass - Kristoffer Joner - Trapping - Man in the Wilderness

Revenant - Wikipedia, the free encyclopedia

<https://en.wikipedia.org/wiki/Revenant> - Wikipedia -
A **revenant** is a visible ghost or animated corpse that is believed to have returned from the grave to terrorize the living. The word **revenant** is derived from the ...
Nachzehrer - Revenant (disambiguation) - Wiedergänger - Gjenganger

The Revenant (2015) - Rotten Tomatoes

www.rottentomatoes.com/films/the_revenant_2015/ - Rotten Tomatoes -

The Revenant

2015 - Drama film/Thriller - 2h 36m

Play trailer on YouTube

8.1/10 IMDb | 82% Rotten Tomatoes | 76% Metacritic

While exploring the uncharted wilderness in 1823, frontiersman Hugh Glass (Leonardo DiCaprio) sustains life-threatening injuries from a brutal bear attack. When a member (Tom Hardy) of his hunting team kills his young son (Forrest Goodluck) and leaves him for dead, Glass must utilize his survival sk... More

Release date: December 25, 2015 (USA)
Director: Alejandro González Iñárritu
Adapted from: The Revenant
Cinematography: Emmanuel Lubezki
Nominations: Academy Award for Best Picture, More

Available on

Google Play Movies & TV	\$14.99
YouTube	From \$14.99
iTunes	From \$14.99

More services

Critic reviews

View 6+ more

It's an imposing vision, to be sure, but also an inflated and emotionally stunted one, despite an anchoring performance of ferocious 200% commitment from Leonardo DiCaprio. Full review
Justin Chang - Variety

One Google experience

Google

revenant
revenant meaning
revenant showtimes
revenant review

About 58,800,000 results (0.58 seconds)

Showing results for **revenant**
Search instead for reverant

Showtimes for The Revenant

Wed, Apr 13 | Thu, Apr 14 | Fri, Apr 15 | Sat, Apr 16

All times Morning Afternoon Evening Night

ArcLight Santa Monica - Map
4:30pm 6:15 9:30

Valley Plaza 6 - North Hollywood - Map
6:40pm 10:10

Pacific Winnetka All Stadium 21 - Map
4:05pm 7:20

All times are in PT

More showtimes

The Revenant (2015) - IMDb

www.imdb.com/title/tt1663202/ Internet Movie Database

★★★★★ Rating: 8.1/10 - 321,567 votes

A frontiersman on a fur trading expedition in the 1820s fights for survival after being mauled by a bear and left for dead by members of his own hunting team. ... Leonardo DiCaprio, Tom Hardy, Will Poulter. ... Q: How is Leonardo DiCaprio's role in "The Revenant" different from his usual ...

Full Cast & Crew - Paul Anderson - John Fitzgerald - 2016

The Revenant (2015 film) - Wikipedia, the free encyclopedia

https://en.wikipedia.org/wiki/The_Revenant_(2015_film) Wikipedia

The Revenant is a 2015 American historical drama directed, co-produced, and co-written by Alejandro G. Iñárritu. The screenplay by Iñárritu and Mark L. Smith is ... Hugh Glass - Kristoffer Joner - Trapping - Man in the Wilderness

Revenant - Wikipedia, the free encyclopedia

https://en.wikipedia.org/wiki/Revenant Wikipedia

A **revenant** is a visible ghost or animated corpse that is believed to have returned from the grave to terrorize the living. The word **revenant** is derived from the ...
Nachzehrer - Revenant (disambiguation) - Wiedergänger - Gjenganger

The Revenant (2015) - Rotten Tomatoes

www.rottentomatoes.com/films/the_revenant_2015/ Rotten Tomatoes

The Revenant

2015 · Drama film/Thriller · 2h 36m

Play trailer on YouTube

8.1/10 IMDb | 82% Rotten Tomatoes | 76% Metacritic

While exploring the uncharted wilderness in 1823, frontiersman Hugh Glass (Leonardo DiCaprio) sustains life-threatening injuries from a brutal bear attack. When a member (Tom Hardy) of his hunting team kills his young son (Forrest Goodluck) and leaves him for dead, Glass must utilize his survival sk... More

Release date: December 25, 2015 (USA)
Director: Alejandro González Iñárritu
Adapted from: The Revenant
Cinematography: Emmanuel Lubezki
Nominations: Academy Award for Best Picture, More

Available on

Google Play Movies & TV	\$14.99
YouTube	From \$14.99
iTunes	From \$14.99
More services	


Critic reviews

View 6+ more

Its an imposing vision, to be sure, but also an inflated and emotionally stunted one, despite an anchoring performance of ferocious 200% commitment from Leonardo DiCaprio. Full review
Justin Chang - Variety

Explosive computational requirements

- Growing in many dimensions
 - More users, use more
 - More data
 - web pages, videos, ads, mailboxes, blogs, etc.
 - Better quality
 - more relevant, faster response, “understanding user intent”



Google DC in 1999



Google DC today



How do we build Google systems?

Need to balance:

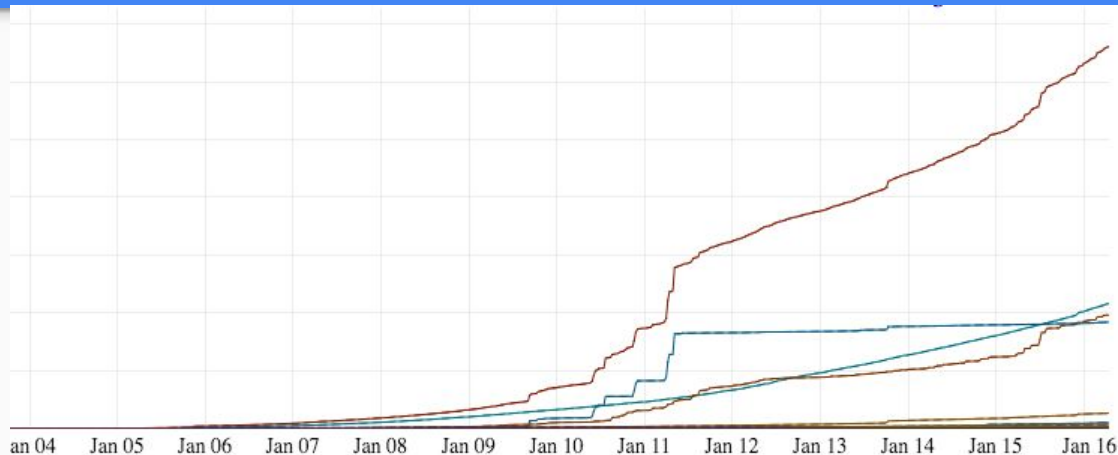
- Simplicity
- Features
- Platform transparency
- Latency
- Scalability
- Fault-tolerance
- Reliability
- Evolvability
- Availability
- ...

The image displays a mobile search interface for the movie 'The Revenant'. At the top, a search bar contains the word 'revenant'. Below it, a dropdown menu lists suggestions: 'revenant', 'revenant meaning', 'revenant filming location', and 'revenant showtimes'. The search results are divided into several sections:


- Showtimes for The Revenant:** A table showing showtimes for Monday, April 25, through Thursday, April 28. It lists theaters like Valley Plaza 6 and Regency Academy Cinemas 6 with showtimes at 7:00pm and 7:10pm.
- The Revenant (2015) - IMDb:** A snippet from the Internet Movie Database showing a rating of 8.1/10 and a brief synopsis of the movie.
- The Revenant Movie Card:** A card featuring a collage of movie stills, the title 'The Revenant', a release date of 2015, a genre of Drama film/Thriller, a duration of 2h 36m, and ratings from IMDb (8.1/10), Rotten Tomatoes (82%), and Metacritic (76%).

At the bottom, there are links to IMDb, Rotten Tomatoes, and Metacritic.

Development scale



Billions of
LOC



~20K active
developers

~100K daily
changes to
code repository

Reuse to the rescue - Google technology stack



Google technology stack (selected subset)

- Common data language
- Common RPC mechanism
- Data storage solutions
 - Bigtable
 - Spanner
 - Colosus (GFS++)
- Processing pipelines
- GUI frameworks
- Logging
- Logs processing
- Testing
- Monitoring/Alerting
- Deployment
- Runtime evolution

Message passing

- Protocol buffers
- Language neutral construct for serializing structured data

```
message SearchRequest {  
  required string query = 1;  
  optional int32 page_number = 2; // Which page number do we want?  
  optional int32 results_per_page = 3; // Number of results to return per page.  
}
```

- Enables exchange of data among Google components or systems - **lingua franca**
- ~50K different message types in Google code tree across ~12K .proto files
- Used in RPC communication and for persistent storage of data

Stubby (gRPC)

- Extends the .proto files with definition of services (<http://www.grpc.io>)
- Services define inputs and outputs in terms of .proto messages
- Enables any two Google components written in any language to communicate
- Forms the “microservice architecture” paradigm

```
// The query service definition.
service QueryService {
  // respond to search
  rpc Search(SearchRequest) returns (SearchReply) {}
}

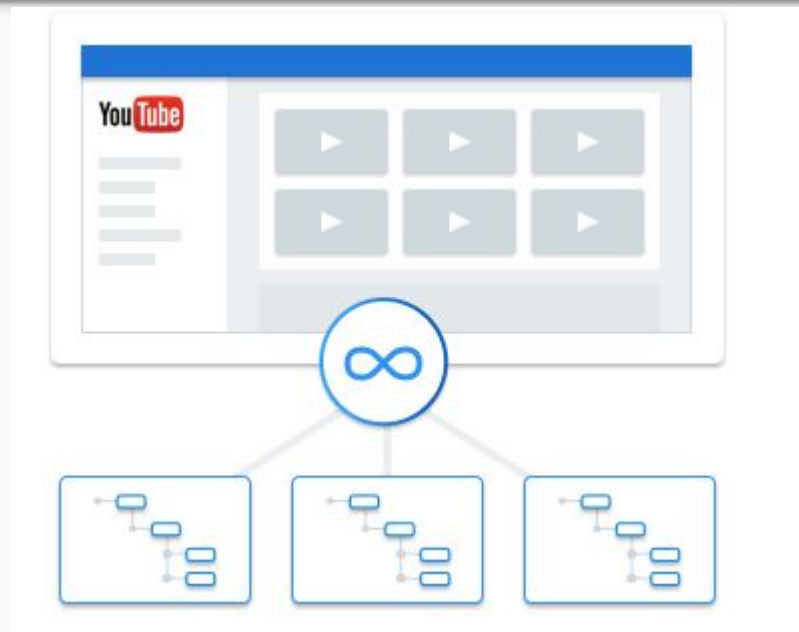
// The request message with the query string.
message SearchRequest {
  required string query = 1;
}

// The reply message with the search and ad results
message SearchReply {
  repeated SearchResult result = 1;

  repeated AdResult ad_result = 2;
}
```

Common storage platform - Bigtable

- Database service powering many Google systems (search, analytics, maps, gmail, ...)
- **Highly scalable**
 - sparsely populated table that can scale to billions of rows and thousands of columns, petabytes of data
- Rows indexed - ideal for single key data and very **low latency**
- **Availability** and **reliability** through redundancy/replication
- High read/write low latency rates
- Ideal for Mapreduce operations



Common logging and log processing

- Reusable constructs for logging system behavior
- Log processing framework for extracting information from the logs
- Common components for data sampling and aggregation
- Data access security

```
123.45.67.89 - 25/Mar/2002 10:15:32 -
```

```
http://www.google.com/search?q=cars -
```

```
Firefox 1.0.7; Windows NT 5.1 -
```

```
740674ce2123e969
```

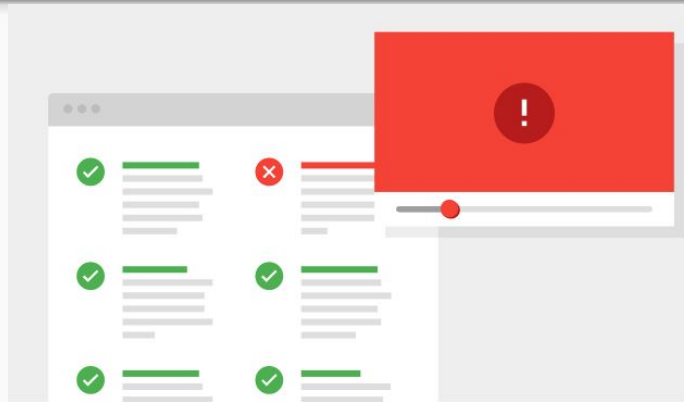


Dataflow systems are commonly needed at Google

- Reusable components for dataflow processing
- Processing organized into stages
- At each stage there is some “worker type” binary that consumes work from previous stages and produces some output data
- Guarantees all work is executed exactly once or reflected in the current system state
- Multiple instances of given worker type can be self-scheduled and grab applicable work
- Resilient to intermittent failures (worker does not “consume” task until it succeeds)

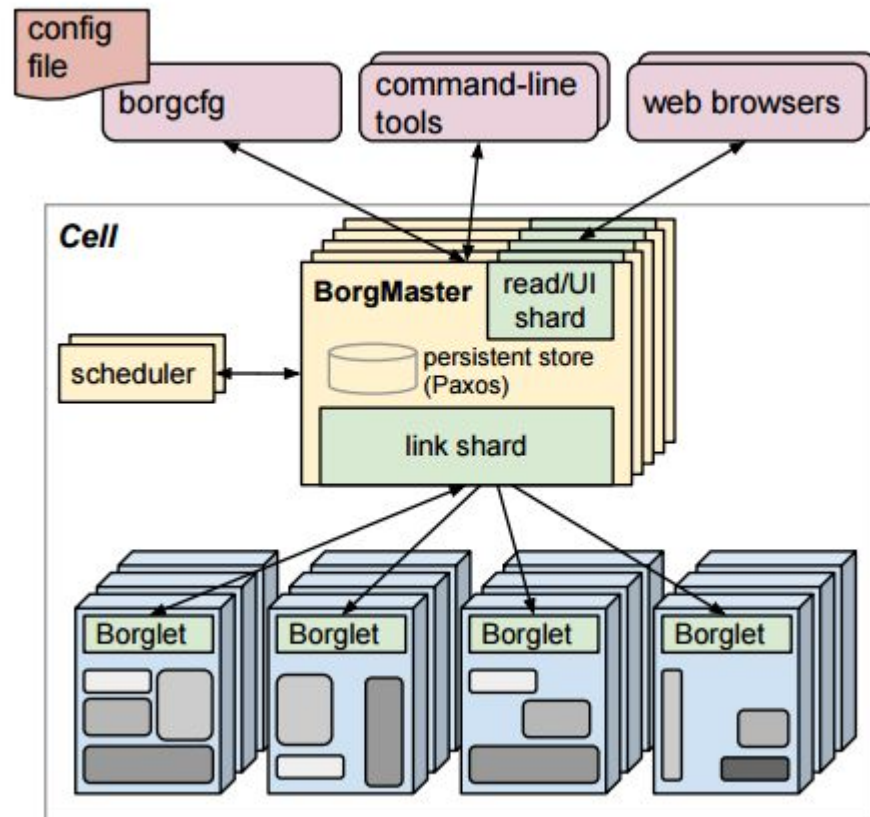
Testing support at Google

- TAP - for fast (unit) tests
 - runs all affected unit tests on every change sent for review
 - reports status on CL
- GUITAR - Google Unified Integration Testing and Releasing
 - longer tests
 - tests that require a lot of setup
 - tests that may require a lot of resources



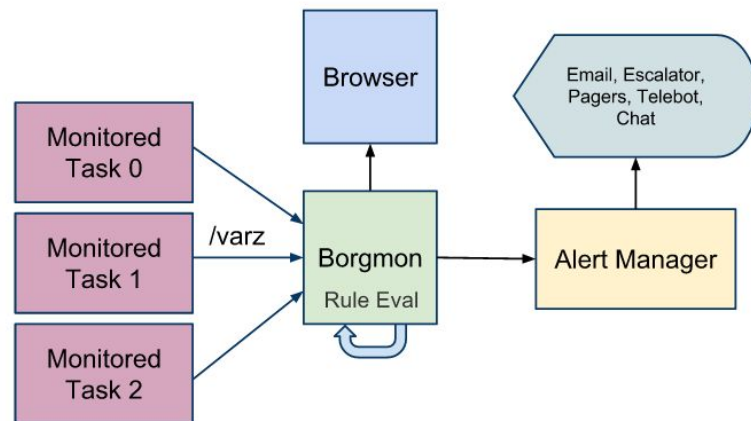
Deployment - Borg

- System for distribution of jobs and tasks within a machine cluster
- Hides the details of resource management and failure handling
- Operates with very high reliability and availability, and supports applications that do the same
- Assigns jobs to machines while satisfying resource constraints/requirements




Common monitoring and alerting

- Varz: common libraries to track the state of the system
 - E.g., number of queries
 - Can monitor state of single task or multiple replicas
- Borgmon
 - System and rules to monitor interesting (combinations of) the variables
 - Automatic dashboard generation
- Alert manager to escalate when defined rules trigger
- Metamon - monitoring borgmon



Controlled system evolution - experiment framework



- Divert small portion of end user traffic to the new version
- Measure effect of a change before it is rolled out to everyone

SWE organization



SWE Organization at Google

- Google systems are organized into Product Areas

Web

- Web Search**
Search billions of web pages
- Toolbar**
Add a search box to your browser
- Google Chrome**
A browser built for speed, simplicity and security
- Bookmarks**
Access your bookmarks and starred items

Mobile

- Mobile**
Get Google products on your mobile phone
- Search for mobile**
Search Google wherever you are
- Maps for mobile**
View maps, your location and get directions on your phone

business

- AdWords**
Attract more customers and only pay for results
- Google Cloud Platform**
Build and host applications and websites, store and analyze data on Google's secure infrastructure
- AdSense**
Create online revenue today
- Analytics**
Know your audience and analyze usage
- Google Apps for Work**
Get email, docs, storage and more, customized for your business
- Google My Business**
Make sure your business looks great on Google Search, Maps and Google+ for free
- AdMob**
Make money from your apps
- Google Domains**
Find a domain and build a website for your business

Media

- YouTube**
Watch, upload and share videos
- Books**
Search the full text of books
- News**
Search thousands of news stories
- Google Photos**
All your photos on all your devices, organized and easy to share
- Google Play**
Your music, movies, books, and Android apps available anywhere
- Image Search**
Search for images on the web
- Video Search**
Search for videos on the web
- Google Cardboard**
Experience virtual reality in a simple, fun, and affordable way

Geo

- Maps**
View maps and directions
- Panoramio**
Explore and share photos of the world
- Earth**
Explore the world from your computer

Specialized Search

- Custom Search**
Create a customized search experience for your community
- Finance**
Business info, news and interactive charts
- Trends**
Explore past and present search trends
- Google Shopping**
Search for stuff to buy
- Scholar**
Search scholarly papers

Home & Office

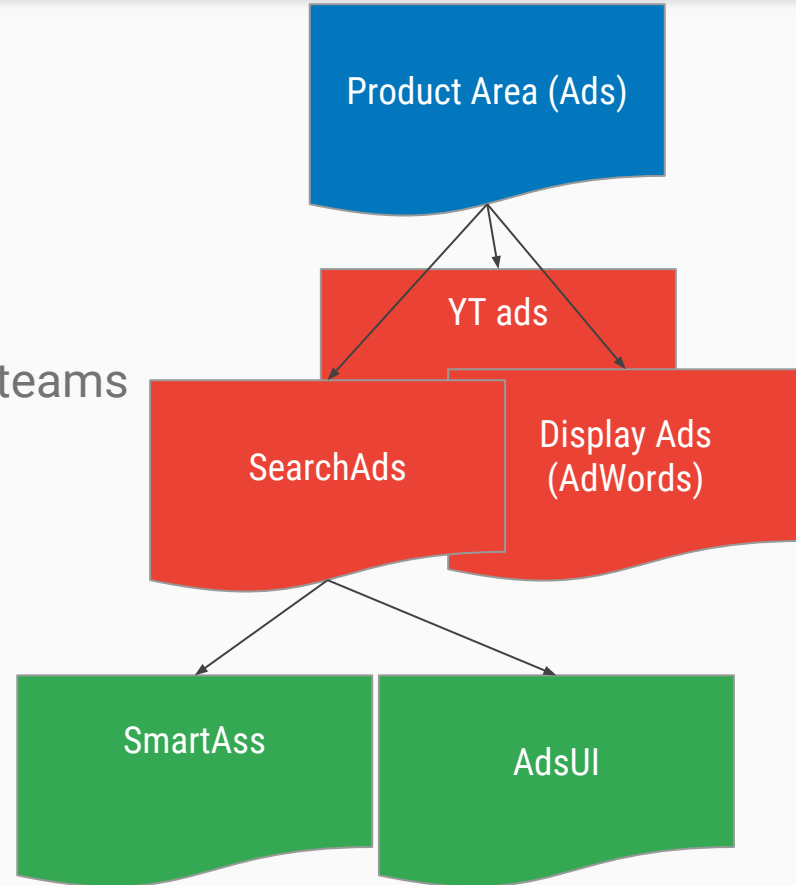
- Gmail**
Fast, searchable email with less spam
- Docs**
Open, edit, and create documents
- Slides**
Open, edit, and create presentations
- Drawings**
Create diagrams and flow charts
- Calendar**
Organize your schedule and share events with friends
- Voice**
One number for all your phones, online voicemail and cheap calling
- Google Cloud Print**
Print anywhere, from any device
- Google Store**
Explore and shop the latest products made with Google
- Drive**
Create, share and keep all your stuff in one place
- Sheets**
Open, edit, and create spreadsheets
- Forms**
Build free surveys
- Sites**
Create websites and secure group wikis
- Translate**
Instantly translate text, web pages, and files between over 50 languages
- Google Wallet**
Make your phone your wallet
- Google Keep**
Save what's on your mind

Social

- Google+**
Real-life sharing, rethought for the web
- Groups**
Create mailing lists and discussion groups
- Blogger**
Share your life online with a blog - it's quick, easy and free
- Hangouts**
Conversations that come to life. Anytime, anywhere, for free

SWE Organization at Google

- Within each PA, there are multiple products
 - Search ads
 - Product listing ads
 - Display ads
 - YT ads
 - ...
- Within each product, there are many projects/teams
 - SmartAss
 - AdsUI
 - ...



Team roles

- **TLs**

- In charge of large project or several medium-size projects
- Plan work for the team (objectives -> design)
- Guide more junior engineers
- Approve designs by more junior engineers

- **TLM**

- In charge of multiple projects, manages 1+ teams (5 in my case)
- Sets project direction and also manages people
- Works across the organization to plan and execute work

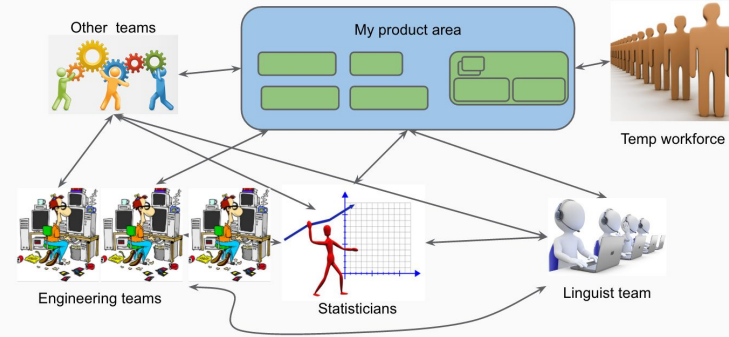
- **PMs**

- Works with TL/Ms from all relevant teams to develop product direction

- **Engineers**

- Junior: works on many tasks designed by more senior team members
- Senior: owns larger parts of project from design to delivery

- **Statisticians, Linguists, ...**



SWE Organization at Google - Team

- Team objectives are planned yearly and quarterly
 - Prioritize (risk/opportunity assessment)
 - Binary vs incremental
- Multiple teams typically share objectives
- Regular iteration schedule
 - Check in on objectives
 - Communicate the status of previous iteration
 - Plan next iteration of work
(design/development/testing/releasing/monitoring)
- Daily stand-up meeting

SWE Organization at Google - Team

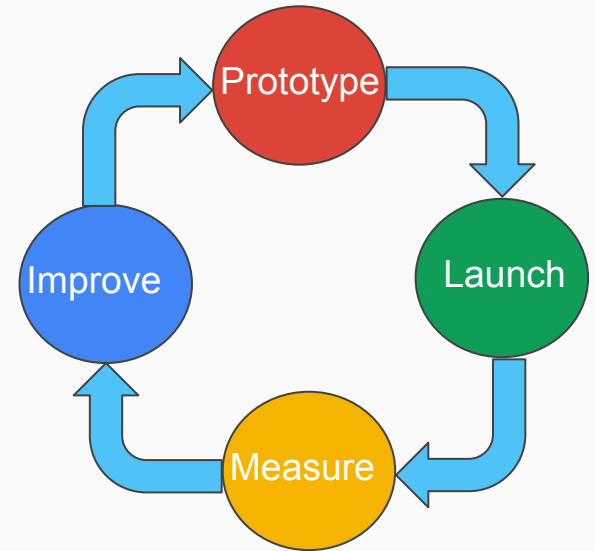
- Team process needs to continuously adapt to
 - Team structure (new eng vs attrition)
 - Project scope/complexity
 - Inter-project dependencies
 - Project deadlines
 - Available resources
- Always re-examine the current state
 - What is working well and what is not
 - What are the risks
- Who participates in what phase of the process
 - Efficiency (can you effectively brainstorm with 10 people)
- Communication
 - Within team and across teams
- Team structure
 - One big team vs. several smaller teams

Development process at Google



Development process at Google

- Prototype-launch-measure-improve
- Requirements
 - Derived from strategic objectives
 - Cascaded to projects and specific objectives within a (set of) project(s)
 - Non-functional requirements
 - Derived from measuring current state
 - Driven by the infrastructure changes
 - E.g., migration to new storage platform



Architecture/Design

- Capture at **different levels** how requirements will be met
- Heavily influenced by Google scale and available infrastructure/components
- Heavily influenced by non-functional requirements
 - Scalability
 - Latency
 - Fault-tolerance
- Frequently involve back-of-the envelope calculations
 - How much data, QPS, CPU,....
- No project should be scheduled for implementation if it does not have a design
 - Schedule design/prototype work first

- High level view
 - Architecture diagram, highlighting new/changed/removed components
- What are the major components, and how do they interact?
- What are the inputs and outputs?
- What changes to the existing system/components are needed
- What is the underlying data model, why is it structured that way
- How much data?
- Other non-functional properties?
- Rollout plan

Production readiness

- Resource requirements
 - Only theoretically unlimited computational resources

Production readiness

- Resource requirements
 - Only theoretically unlimited computational resources
- Latency

Production readiness

- Resource requirements
 - Only theoretically unlimited computational resources
- Latency
- Availability

Production readiness

- Resource requirements
 - Only theoretically unlimited computational resources
- Latency
- Availability
- Scalability

Production readiness

- Resource requirements
 - Only theoretically unlimited computational resources
- Latency
- Availability
- Scalability
- Traffic capacity

Production readiness

- Resource requirements
 - Only theoretically unlimited computational resources
- Latency
- Availability
- Scalability
- Traffic capacity
- Load balancing

Production readiness

- Resource requirements
 - Only theoretically unlimited computational resources
- Latency
- Availability
- Scalability
- Traffic capacity
- Load balancing
- **Monitoring**
- **Alerting**

Design and launch decisions

- Engineers make “smaller” design decisions
- TLs design systems
- Micro-scheduling of design reviews
 - Everyone can schedule them
 - Design review can become design tutoring
 - Within the team for smaller changes
 - With other teams/TL for any bigger/visible change
- What happens during design reviews
 - Flaws and unintended effects (hopefully) identified
 - Alternatives suggested
 - Selection of launch criteria
- Launch reviews
 - Data driven launch decisions
 - Security and privacy sign-off

Some rules of design at Google

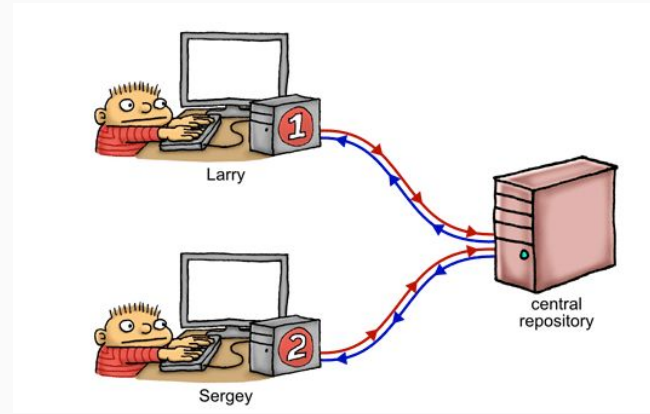
- Design document should be in sync with the implementation
 - New tools make it easier to require
- Feedback should be sought early and often
 - Design needs signoff from relevant stakeholders
- Prototype key elements to evaluate your design
- Understand design choices and limitations of existing components/systems
- Research existing solutions to similar problems
- **Document tradeoffs and alternatives considered**
- **The Pareto principle (20-80)**
 - What is good enough?

Design challenges

- Lots of data, lots of computing power
 - Low-level implementation details matter at scale
- Building products/features from reusable components should be fast
 - Reusable components focus on subset of common developer challenges
 - Designing around the component limitations is often hard
- Right design at X may be very wrong at 10X or 100X
 - Design for ~10X growth, but plan to rewrite before ~100X
- Understanding the impact of proposed changes is critical
- Balancing new features against technical debt

Tools Google engineers use daily

- Source control / changelists / version control

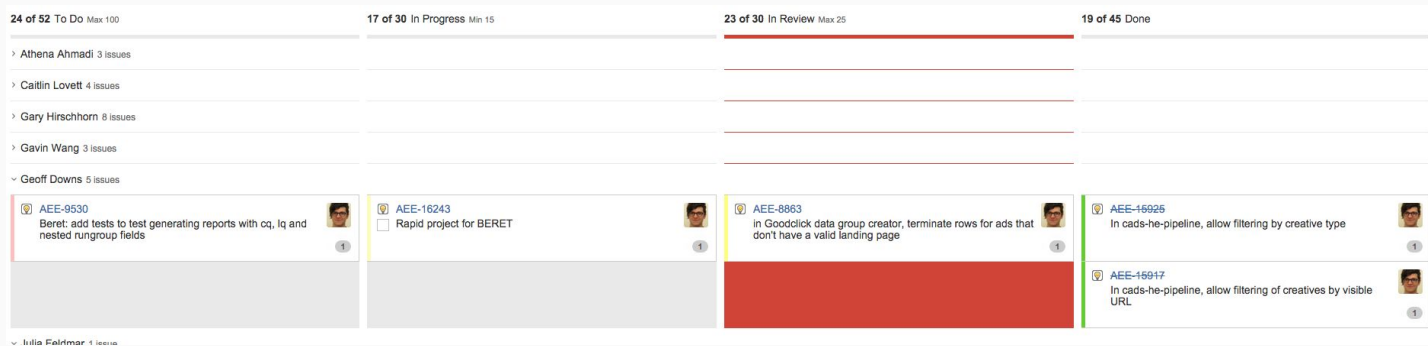


Tools Google engineers use daily

- Source control / changelists / version control
- Bug tracking systems

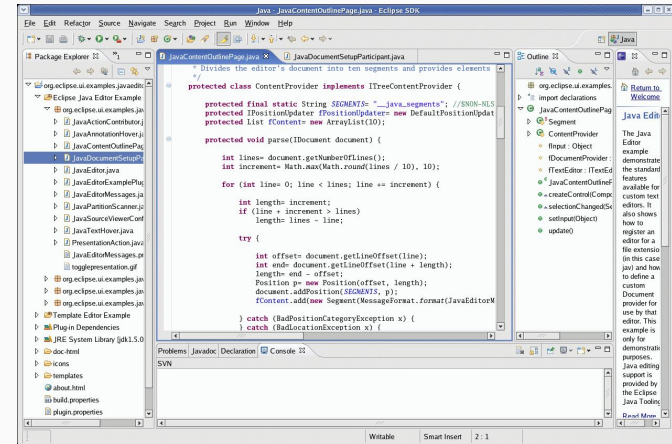
Tools Google engineers use daily

- Source control / changelists / version control
- Bug tracking systems
- Development process tracking tools (jira)



Tools Google engineers use daily

- Source control / changelists / version control
- Bug tracking systems
- Development process tracking tools (jira)
- An editor for programming



Tools Google engineers use daily

- Source control / changelists / version control
- Bug tracking systems
- Development process tracking tools (jira)
- An editor for programming
- One-touch build system
- Deployment tools
- Deployment monitoring tools
- Tons of documentation, codelabs
- Code search
- Test frameworks
- Code reviews, code review tools
- Email, email, email
- IM, hangouts

Implementation support at Google

- Code search
- Single code repository (Piper)
 - Many facets to make sure development continues to scale
- Eclipse plug-ins for Google
- Powerful code review tool
- Automatic code checks
 - Style violations,
 - Testing breakages,
 - Test coverage, ...
- Submit queue
 - Ensures only approved code with all tests passing can be submitted
- Continuous builds/integration tests

Implementation workflow for Google engineers

- Start new jira story/task
- Checkout relevant code from repository
- Make the changes in your favorite editor
- Make it “work”!
- Add adequate logging to measure and report state
- Add/update tests to cover the functionality added/changed/removed
- Add/update release tests
- Run the tests and make sure they pass
- Update relevant documentation
- Create a CL (containing code, tests, and documentation changes)
- Self review the CL, cleanup
- Send for code review

Implementation workflow for Google engineers

- Code review - iterative process
 - change is automatically ran against the main branch
 - at least another pair of eyes, sometimes more
 - address comments
- Are you done?
 - Document release process
 - Follow up CLs
 - Update monitoring and alerting
 - See it through release

Pushmaster is a shared team responsibility (TLs, eng)

- Regular (e.g., bi-weekly) release schedule
- All relevant changes with test plans associated with the release documented
- Staging environment
- Push master builds and pushes to staging
- Push master tests all the new features
- Push master runs test plan
- Full automation of tests, eventually
- Push

- Borg
 - Automatic management jobs/tasks across datacenters and machines
- Release tool - Rapid
 - Packages multiple binaries into single release package
 - Ensures all tests pass
 - Automatic deployment to staging environment
 - Non-automated tests executed in staging
 - Deployment to production

Monitoring and alerting

- Google engineers assume things will fail
 - Processes, machines, DCs
- Provide appropriate alerting in each case
 - Will Borg handle it or someone needs to know
- In case of alerts create playbook entries
 - On call engineer should know how to respond to alert
- Post-mortems for issues that were not anticipated
- On-call - shared responsibility (TLs, eng)
 - Makes sure systems behave
 - Responds to alerts
 - Makes it easier for next on-call
 - Full-time job for 1-2 weeks

Development Scenario #1



Development scenario #1: adding a new field to the proto message

```
message SearchRequest {  
  required string query = 1;  
  optional int32 page_number = 2; // Which page number do we want?  
  optional int32 results_per_page = 3; // Number of results to return per page.  
}
```

Development scenario #1: adding a new field to the proto message

```
message SearchRequest {  
  required string query = 1;  
  optional int32 page_number = 2; // Which page number do we want?  
  optional int32 results_per_page = 3; // Number of results to return per page.  
  required Location query_origin = 4;  
}
```

Development scenario #1: adding a new field to the proto message

```
message SearchRequest {  
  required string query = 1;  
  optional int32 page_number = 2; // Which page number do we want?  
  optional int32 results_per_page = 3; // Number of results to return per page.  
  required Location query_origin = 4;  
}
```

- Any issues with this change?

Backwards/forwards compatibility

- Components generating and reading the data must continue to “understand” each other
- Existing data needs to continue to be valid or needs to be backfilled
- Common language comes with some limitations
- Understanding HOW components interact is required for making changes

Backwards/forwards compatibility

- Components generating and reading the data must continue to “understand” each other
- Existing data needs to continue to be valid or needs to be backfilled
- Common language comes with some limitations
- Understanding HOW components interact is required for making changes

- Proto 3 removes handles backwards compatibility by removing optional/required and making all proto fields optional
- It is up to application logic to enforce/expect exact instance of proto fields

Development Scenario #2



Development scenario #2: debugging a live system

- How often is that 1 line change logged?
- How much information is being logged?
- Where does the logged information live?
 - Temporary logs (expire/get deleted regularly)

Development scenario #2: debugging a live system

- How often is that 1 line change logged?
- How much information is being logged?
- Where does the logged information live?
 - Temporary logs (expire/get deleted regularly)
- Even small changes can have big impact if they happen very frequently

Development Scenario #3



Development scenario #4: scalability, latency, fault tolerance

- Replacing data storage layer
 - Substantial testing in “staging” environment
 - Successful rollout to production
 - 3 days later, lots of slowness, failures
-
- Staging data not big enough to test the behavior
 - Production data was too massive for some (rarely used) RPC call
 - RPC call was triggering cache population which was failing on retries
 - Too many RPC calls to the data store made all other queries to datastore slow

Summary



Summary

- Building Google-scale systems is hard

Summary

- Building Google-scale systems is hard
- Reuse at all levels
 - Domain knowledge
 - Requirements
 - Framework-level solutions
 - Processes
 - Code
 - ...

Summary

- Building Google-scale systems is hard
- Reuse at all levels
 - Domain knowledge
 - Requirements
 - Framework-level solutions
 - Processes
 - Code
 - ...
- Reuse comes at a cost

Summary


- Building Google-scale systems is hard
- Reuse at all levels
- Reuse comes at a cost
- Documenting design decisions is necessary to
 - Communicate proposed changes
 - Understand the impact
 - Understand the tradeoffs

Summary

- Building Google-scale systems is hard
- Reuse at all levels
- Reuse comes at a cost
- Documenting design decisions is necessary to
 - Communicate proposed changes
 - Understand the impact
 - Understand the tradeoffs
- **Developers regularly propose system changes that impact other components and systems**
 - We are still missing tools/processes to ensure these are understood well in advance
 - **Post mortems - learning from mistakes**
 - **No public shaming, goal is success of the project**
 - Every developer must have architecture-level knowledge to be successful

How has software engineering evolved over the past 12 years @Google

- We are still far from developing software by simply connecting components
 - But we are building **way** more complex systems
 - Keeping up with the scale



How has software engineering evolved over the past 12 years @Google

- We are still far from developing software by simply connecting components
 - But we are building **way** more complex systems
 - Keeping up with the scale
- **Moore's law has really enabled Google's growth**

How has software engineering evolved over the past 12 years @Google

- We are still far from developing software by simply connecting components
 - But we are building **way** more complex systems
 - Keeping up with the scale
- Moore's law has really enabled Google's growth
- Visionary investment in
 - Computing power (datacenters)
 - Infrastructure (e.g., Borg, Piper, ...)

How has software engineering evolved over the past 12 years @Google

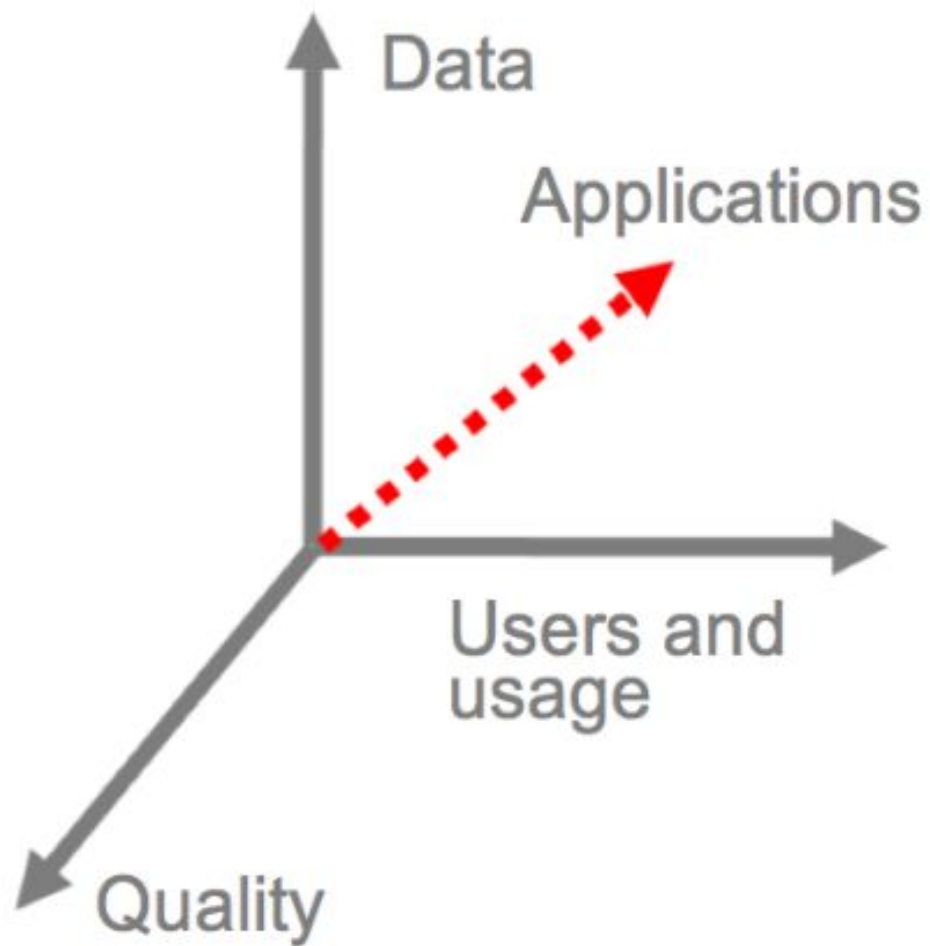
- We are still far from developing software by simply connecting components
 - But we are building **way** more complex systems
 - Keeping up with the scale
- Moore's law has really enabled Google's growth
- Visionary investment in
 - Computing power (datacenters)
 - Infrastructure (e.g., Borg, Piper, ...)
- **Balancing engineering freedom and top-down structure**
 - Significant reduction in effort duplication

How has software engineering evolved over the past 12 years @Google

- We are still far from developing software by simply connecting components
 - But we are building **way** more complex systems
 - Keeping up with the scale
- Moore's law has really enabled Google's growth
- Visionary investment in
 - Computing power (datacenters)
 - Infrastructure (e.g., Borg, Piper, ...)
- Balancing engineering freedom and top-down structure
 - Significant reduction in effort duplication
- **Continuous investment in technical debt**

How has software engineering evolved over the past 12 years @Google

- We are still far from developing software by simply connecting components
 - But we are building **way** more complex systems
 - Keeping up with the scale
- Moore's law has really enabled Google's growth
- Visionary investment in
 - Computing power (datacenters)
 - Infrastructure (e.g., Borg, Piper, ...)
- Balancing engineering freedom and top-down structure
 - Significant reduction in effort duplication
- Continuous investment in technical debt
- Significant investment in developer education
 - Codelabs, classes, training



Thank You!

Q&A

