

# Scaling Slack

Bing Wei  
Infrastructure@Slack

**Acme Sites** ▼ 🔔

● Matt Kump

---

📁 1 ★ STARRED

# design-work

📁 2 **# events**

👤 Cory, Tina, Dio

---

📁 5

CHANNELS (39) +

📁 3

# accounting-costs

**# brainstorming** 1

# business-ops

**# culture**

🔒 design-chat

# marketing

**# media-and-pr**

🔒 sonic-fanfic

# triage-issues

---

DIRECT MESSAGES (24) +

♥ slackbot

● Brandon Velestuk

● Caroline McCarthy

○ Cory Bujnowicz

○ Fayaz Ashraf

● Graham Hicks

👤 Lane, Pavel

○ Mari Ju

● Matt Hodgins


○ Shannon Tinkley

● Terra Spitzner


**#culture** 19 members | [Add a topic](#)


📞 ⓘ ⚙️ 🔍 Search @ ☆ ⋮

Today


 **Meredith Brown** 12:50 PM  
Really need to give some Kudos to @julie for helping out with the new influx of tweets yesterday. People are really, really excited about yesterday's announcement.


👍 2 ✨ 8 🍌 1 ❤️ 3


 **Kiné Camara** 12:55 PM  
No! It was my pleasure! People are very excited. ⚡

 **Damien Baker** 2:14 PM  
What are our policies in regards to pets in the office? I'm assuming it's a no-go, but thought I would ask here just to make sure that was the case.

My pooch just had surgery and I'd like to be with them if possible.

 **Jake Grimes** 2:18 PM  
Your assumption is correct. Mostly for allergy concerns, but the building also has a restriction on them.


 **Jake Grimes** 2:19 PM  
shared a post ▼

 **Building Policies and Procedures**  
Last edited 2 months ago

---

**SECURITY POLICIES**


- All guests and visitors must sign in
- Guests and visitors must be accompanied throughout the office
- Last to leave is responsible for setting the alarm.

 **Damien Baker** 2:22 PM  
Thanks Jake!


+ Message #culture @ 😊

**About #culture** ✕


---

 Channel Details ▶


---

 Pinned Items ▶


---

 12/19 Members ▶

---

 Shared Files ▶

---

 Notification Preferences ▶

**Acme Sites** ▼ 🔔

● Matt Kump

---

🔖1 **★ STARRED**

# design-work

🔖2 **# events**

📄 Cory, Tina, Dio

---

🔖5

CHANNELS (39) +

🔖3 # accounting-costs

**# brainstorming** 1

# business-ops

**# culture**

🔒 design-chat

# marketing

**# media-and-pr**

🔒 sonic-fanfic

# triage-issues

---

DIRECT MESSAGES (24) +

♥ slackbot

● Brandon Velestuk

● Caroline McCarthy

○ Cory Bujnowicz

○ Fayaz Ashraf

● Graham Hicks

📄 Lane, Pavel

○ Mari Ju

● Matt Hodgins


○ Shannon Tinkley


● Terra Spitzner


**#culture** 19 members | [Add a topic](#)


🔍 Search @ ☆ ⋮


Today


 **Meredith Brown** 12:50 PM  
Really need to give some Kudos to @julie for helping out with the new influx of tweets yesterday. People are really, really excited about yesterday's announcement.  
👍 2 ✨ 8 🍌 1 ❤️ 3

 **Kiné Camara** 12:55 PM  
No! It was my pleasure! People are very excited. ⚡

 **Damien Baker** 2:14 PM  
What are our go, but thoug  
My pooch jus


 **Jake Grimes**  
Your assumpt  
also has a rest

 **Jake Grimes**  
shared a post

 **Build**  
Last e


**SECURITY F**


- All guests
- Guests at
- Last to le

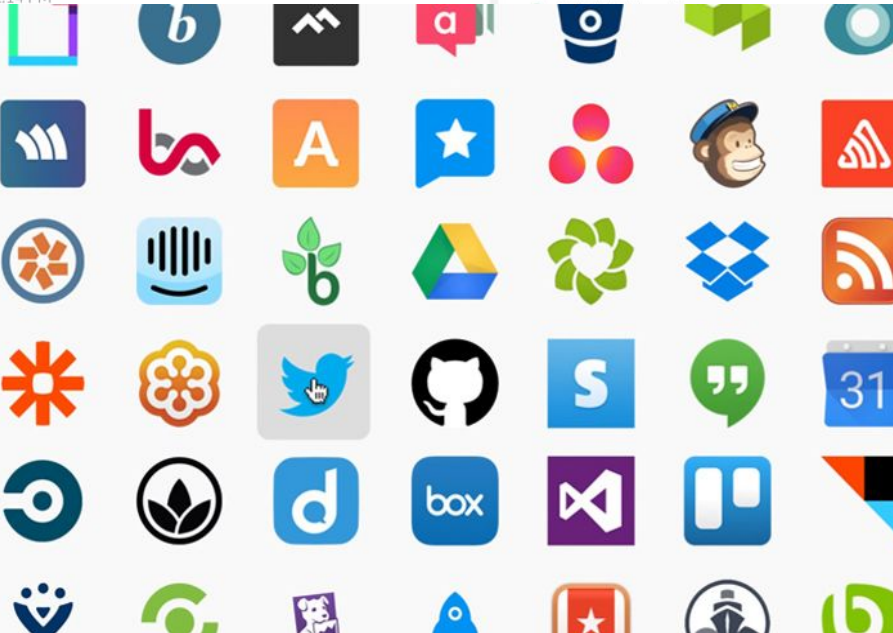
 **Damien Bake**  
Thanks Jake!

+ Message #cu

**About #culture** ✕

 Channel Details ▶

 Pinned Items ▶





Our Mission:

To make people's working  
lives simpler, more pleasant,  
and more productive.





*From* supporting small teams  
*To* serving gigantic organizations of  
hundreds of thousands of users



## Slack Scale

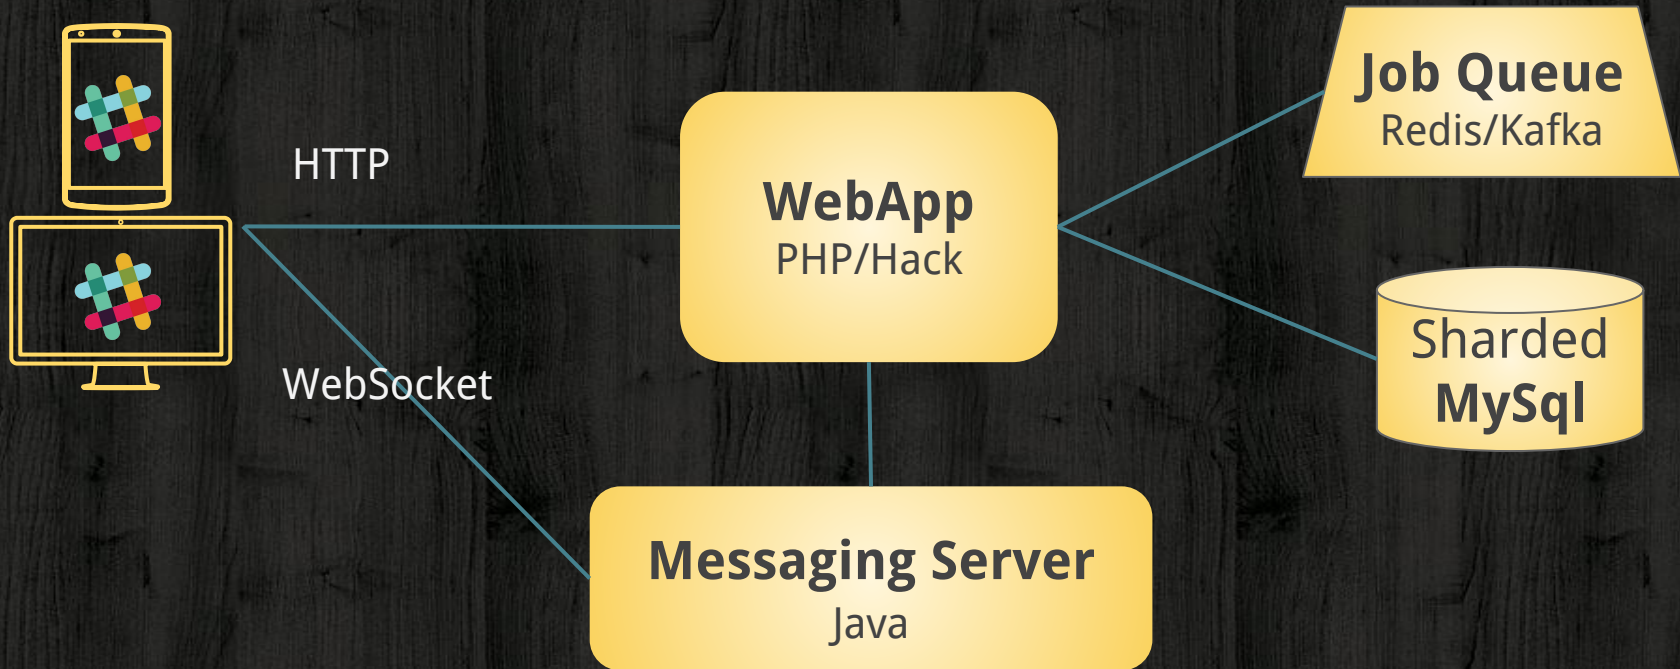
---

- ◆ 6M+ DAU, 9M+ WAU  
5M+ peak simultaneously connected
- ◆ Avg 10+ hrs/weekday connected  
Avg 2+ hrs/weekday in active use
- ◆ 55% of DAU outside of US



# Cartoon Architecture

---



---

# Outline

- ◆ The slowness problem
- ◆ Incremental Improvements
- ◆ Architecture changes
  - Flannel
  - Client Pub/Sub
  - Messaging Server




*Challenge:*

Slowness Connecting to Slack



# Login Flow in 2015

---



```
{
  ok: true,
  - self: {
    id: "U1WKVSEERY",
    name: "binggmail",
    + prefs: {...},
    created: 1469837648,
    manual_presence: "active"
  },
  + team: {...},
  latest_event_ts: "1510364623.000000",
  + channels: [...],
  + groups: [...],
  + ims: [...],
  cache_ts: 1510365223,
  read_only_channels: [ ],
  can_manage_shared_channels: false,
  + mpims: [...],
  + subteams: {...},
  + dnd: {...},
  + users: [...],
  cache_version: "v16-giraffe",
  cache_ts_version: "v2-bunny",
  + bots: [...],
  url: "wss://mpmulti-qdvx.lb.slack-msgs.com/webs
  xxrCpNy2JMxpygOpoWoqLcr6WSY-ALPCzoEVgEYxnseZ6AA
}
```

## Some examples

number of users/channels	response size
30 / 10	200K
500 / 200	2.5M
3K / 7K	20M
30K / 1K	60M

# Login flow in 2015

---



# Real-time Events on WebSocket

---

User



WebSocket:  
100+ types of events

e.g. chat messages,  
typing indicator,  
files uploads,  
files comments,  
threads replies,  
user presence changes,  
user profile changes,  
reactions, pins, stars,  
channel creations,  
app installations,  
etc.

**Messaging Server**

# Login Flow in 2015

---

- ◆ Clients Architecture
  - Download a snapshot of *entire* team
  - Updates trickle in through the WebSocket
  - Eventually consistent snapshot of whole team

# Problems

---

Initial team snapshot takes time

# Problems

---

Initial team snapshot takes time  
Large client memory footprint



# Problems

---

Initial team snapshot takes time  
Large client memory footprint  
Expensive reconnections

# Problems

---

Initial team snapshot takes time

Large client memory footprint

Expensive reconnections

Reconnect storm

---

# Outline

- ◆ The slowness problem
- ◆ Incremental Improvements
- ◆ Architecture changes
  - Flannel
  - Client Pub/Sub
  - Messaging Server



## Improvements

---

- ◆ **Smaller team snapshot**
  - Client local storage + delta
  - Remove objects out + in parallel loading
  - Simplified objects: e.g. canonical Avatars



## Improvements

---

- ◆ Incremental boot
  - Load one channel first



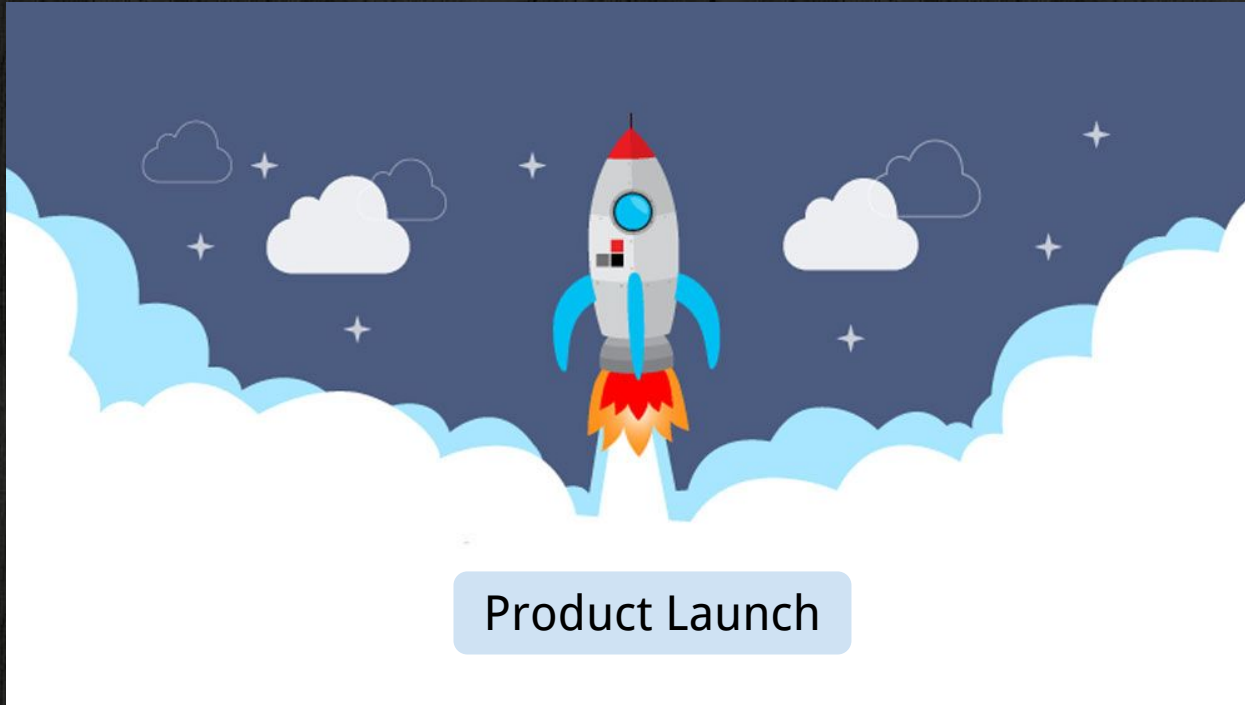
## Improvements

---

- ◆ Rate Limit
- ◆ POPs
- ◆ Load Testing Framework


# Support New Product Features

---



# Cope with New Product Features

---



Product Launch

**Marg, Orla**  
3 members

This is the very beginning of your direct message history with **Marg Hernandez** and **Orla McCourt**. Direct messages are private within this group.

Send me desktop notifications for: [More options](#)

- Activity of any kind
- Only mentions of my name or highlight words

+ Message Marg, Orla

**Aaron Quinn** 2:32 AM

Happy birthday @mark!

11 8 1 1 1



Still...

---

## Limitations

- ◆ What if team sizes keep growing
- ◆ Outages when clients dump their local storage

---

# Outline

- ◆ The slowness problem
- ◆ Incremental Improvements
- ◆ Architecture changes
  - Flannel
  - Client Pub/Sub
  - Messaging Server



# Client Lazy Loading

---

Download less data upfront

- ◆ Fetch more on demand



# Flannel: Edge Cache Service

---

A query engine backed by cache  
on edge locations



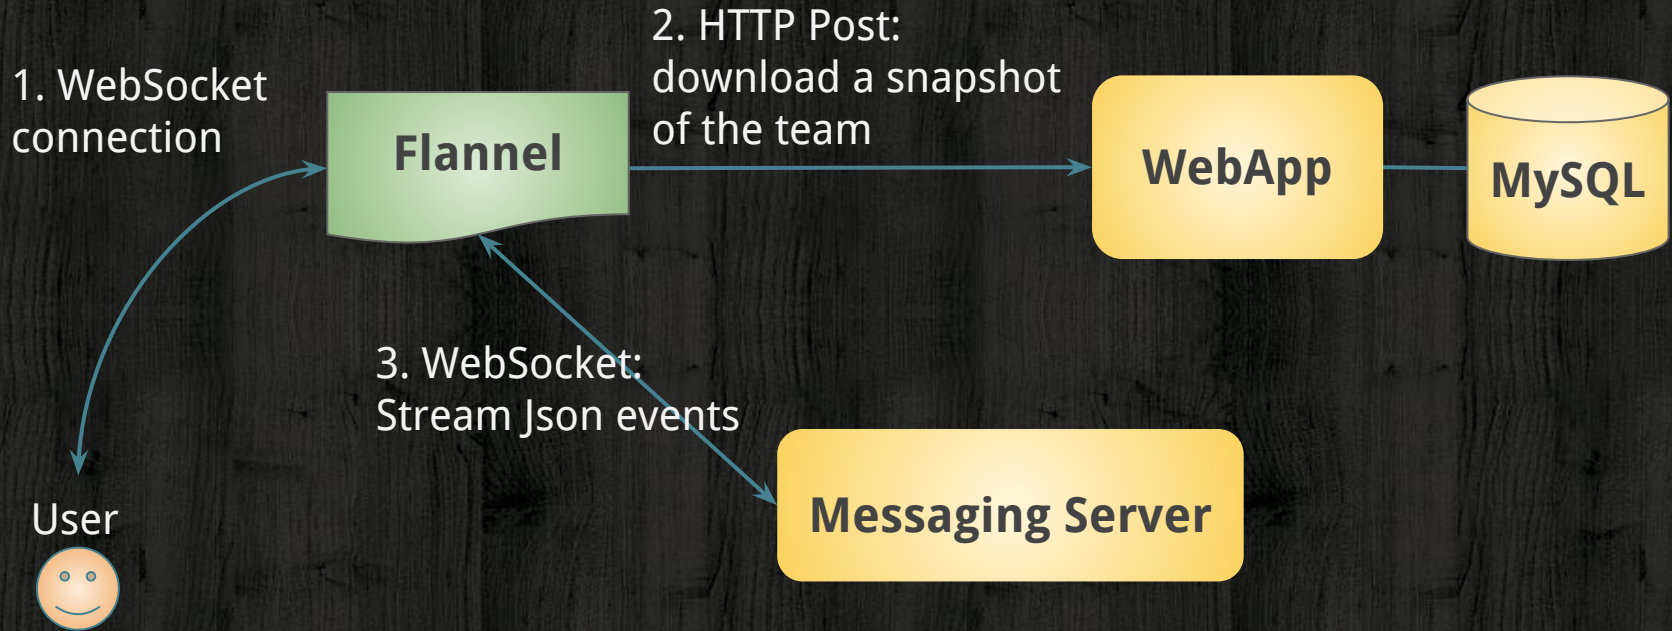
# What are in Flannel's cache

---

- ◆ Support big objects first
  - Users
  - Channels Membership
  - Channels

# Login and Message Flow with Flannel

---



# A Man in the Middle

---



Use real-time events to  
update its cache

E.g. user creation,  
user profile change,  
channel creation,  
user joins a channel,  
channel convert to private

# Edge Locations

---

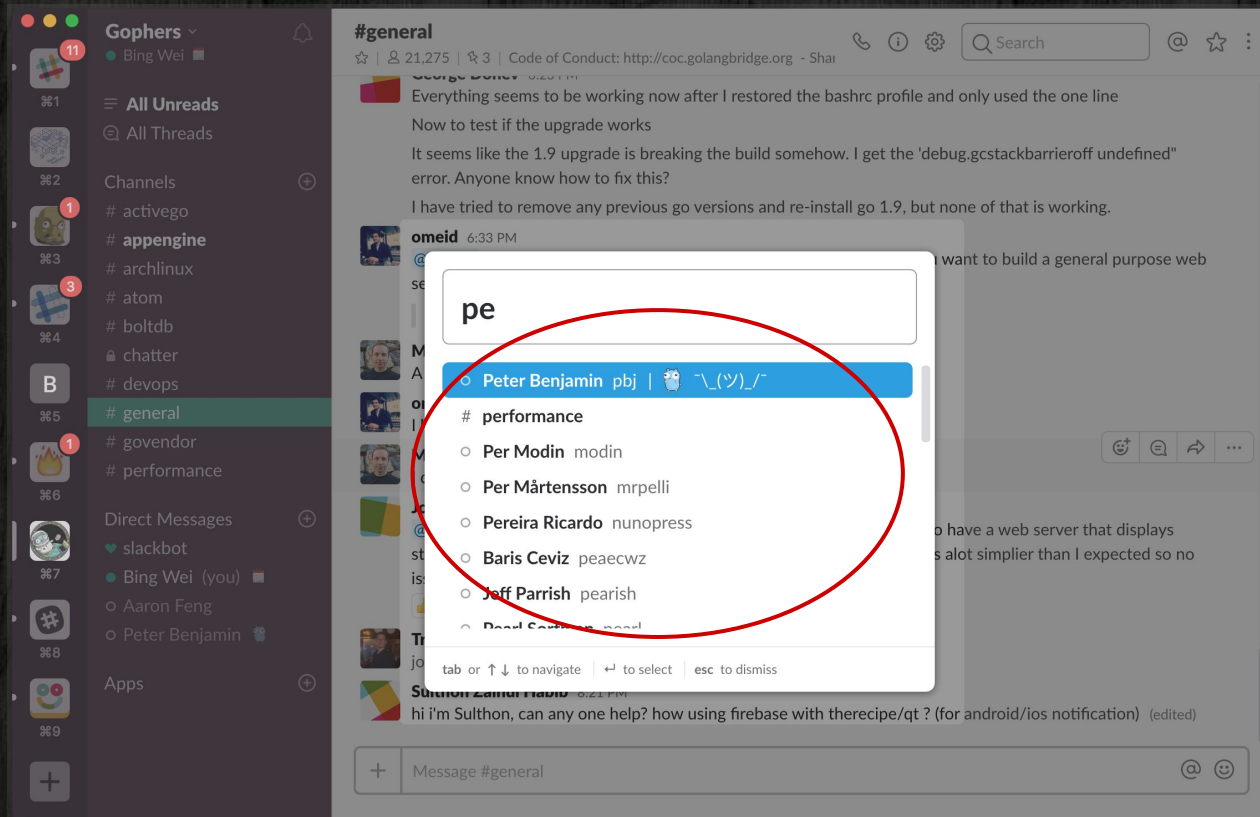


Mix of AWS & Google Cloud



# Examples Powered by Flannel

## Quick Switcher



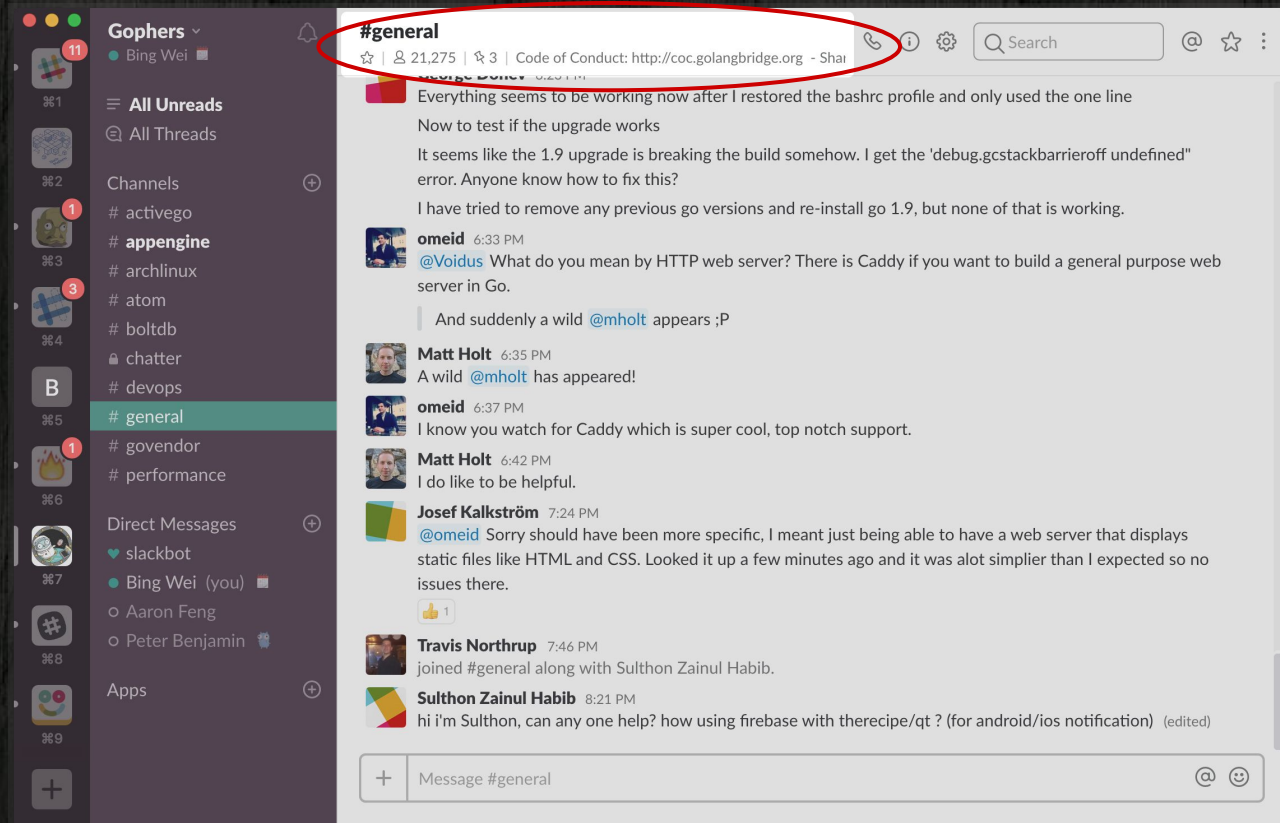
# Examples Powered by Flannel

## Mention Suggestion

The screenshot displays a Slack interface. On the left is a sidebar with a list of channels: #activego, #appengine, #archlinux, #atom, #boltdb, #chatter, #devops, #general, #govendor, and #performance (highlighted). Below the channels are direct messages and apps. The main window shows the #performance channel with a post by Damian Gryski. A mention suggestion dropdown is open over the text '@joeb', listing users: joeb Joe Blubaugh, joeb Joe B, joeblew99 Joe Blew, joa Joa Ebert, and jbarber Joe Barber. The first three items in the list are circled in red. The text '(not in channel)' is visible to the right of each user name in the dropdown.

# Examples Powered by Flannel

## Channel Header



**#general**  
☆ | 👤 21,275 | 📄 3 | Code of Conduct: <http://coc.golangbridge.org> - Shai

Everything seems to be working now after I restored the bashrc profile and only used the one line  
Now to test if the upgrade works  
It seems like the 1.9 upgrade is breaking the build somehow. I get the 'debug.gcstackbarrieroff undefined'  
error. Anyone know how to fix this?  
I have tried to remove any previous go versions and re-install go 1.9, but none of that is working.

**omeid** 6:33 PM  
**@Voidus** What do you mean by HTTP web server? There is Caddy if you want to build a general purpose web server in Go.  
And suddenly a wild **@mholt** appears ;P

**Matt Holt** 6:35 PM  
A wild **@mholt** has appeared!

**omeid** 6:37 PM  
I know you watch for Caddy which is super cool, top notch support.

**Matt Holt** 6:42 PM  
I do like to be helpful.

**Josef Kalkström** 7:24 PM  
**@omeid** Sorry should have been more specific, I meant just being able to have a web server that displays static files like HTML and CSS. Looked it up a few minutes ago and it was alot simpler than I expected so no issues there.

**Travis Northrup** 7:46 PM  
joined #general along with Sulthon Zainul Habib.

**Sulthon Zainul Habib** 8:21 PM (edited)  
hi i'm Sulthon, can any one help? how using firebase with therecipe/qt ? (for android/ios notification)

+ Message #general

# Examples Powered by Flannel

## Channel Sidebar

The screenshot displays a Slack workspace interface. On the left is the 'Channel Sidebar' with a list of channels including #activego, #appengine, #archlinux, #atom, #boltdb, #chatter, #devops, #general, #govendor, and #performance (highlighted). The main channel view shows a message from Chad Kunde (snoproblem) about Arch Linux on Raspberry Pi. The 'About #performance' panel on the right lists 619 members, with a red circle highlighting the list. The member list includes Akshay Kumar, Andrew FigPope, Bing Wei (you), Brent Anderson, Brian Smith, Chris Hines, Damian Gryski, Gopher, huydx, Jason Morgan, Josh Roppo, Kevin Burke, and Konrad Reiche.

**Channel Sidebar**

- Gophers
- All Unreads
- All Threads
- Channels
- # activego
- # appengine
- # archlinux
- # atom
- # boltdb
- # chatter
- # devops
- # general
- # govendor
- # performance
- Direct Messages
- slackbot
- Bing Wei (you)
- Aaron Feng
- Peter Benjamin
- Apps

**#performance**

619 | 1

@dgryski On that note, I have....news

**Chad Kunde (snoproblem)** 9:34 PM  
Arch Linux - aarch64 will boot and run on raspberry-pi 3 unmodified. There's no aarch64 firmware package, so it's not quite full-featured. I haven't had any kernel panics while building docker containers, though, so it's much more stable than my last attempt.

Thursday, November 9th

**Dmitriy Zimnev** 1:16 AM  
@snoproblem could you explain the GOGCCGLAGS environment variable?

**Vincent** 5:18 AM  
joined #performance along with 3 others.

Yesterday

**Vladimir Smirnov** 5:08 AM  
@dgryski I hope there'll be more high performance Aarch64 based boards out there. Cause right now you can find either rpi3-like stuff or something good but for 1k\$+

**Damian Gryski** 9:21 AM  
<https://blog.cloudflare.com/on-the-dangers-of-intels-frequency-scaling/>

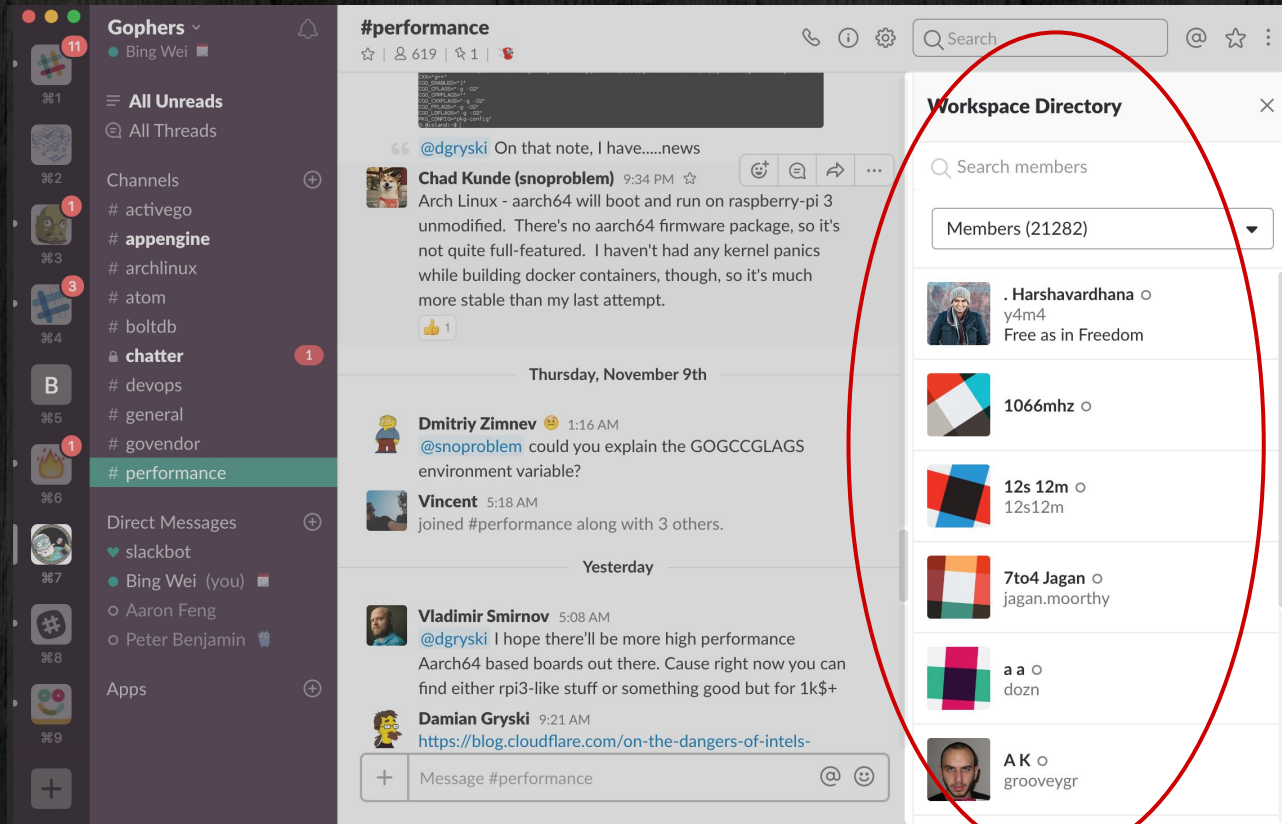
Message #performance

**About #performance**

- Highlights
- 1 Pinned Item
- 619 Members
- Akshay Kumar
- Andrew FigPope
- Bing Wei (you)
- Brent Anderson
- Brian Smith
- Chris Hines
- Damian Gryski
- Gopher
- huydx
- Jason Morgan
- Josh Roppo
- Kevin Burke
- Konrad Reiche

# Examples Powered by Flannel

## Team Directory



The image shows a Slack workspace interface. On the left is a sidebar with a list of channels, including #performance, which is currently selected. The main area displays a conversation in the #performance channel. A red circle highlights a 'Workspace Directory' overlay on the right side of the screen. This overlay contains a search bar for members and a list of team members with their profile pictures, names, and last seen status.

**Workspace Directory**

Search members

Members (21282)

- . Harshavardhana ○  
y4m4  
Free as in Freedom
- 1066mhz ○
- 12s 12m ○  
12s12m
- 7to4 Jagan ○  
jagan.moorthy
- a a ○  
dozn
- AK ○  
grooveygr

# Flannel Results

---

- ◆ Launched Jan 2017
  - Load 200K user team
- ◆ 5M+ simultaneous connections at peak
- ◆ 1M+ clients queries/sec

# Flannel Results

---



---

This is not the end of the story





# Evolution of Flannel

---

# Web Client Iterations

---

## Flannel Just-In-Time Annotation

Right before Web clients are about to access an object, Flannel pushes that object to clients.




## A Closer Look

---

Why does Flannel sit on WebSocket?

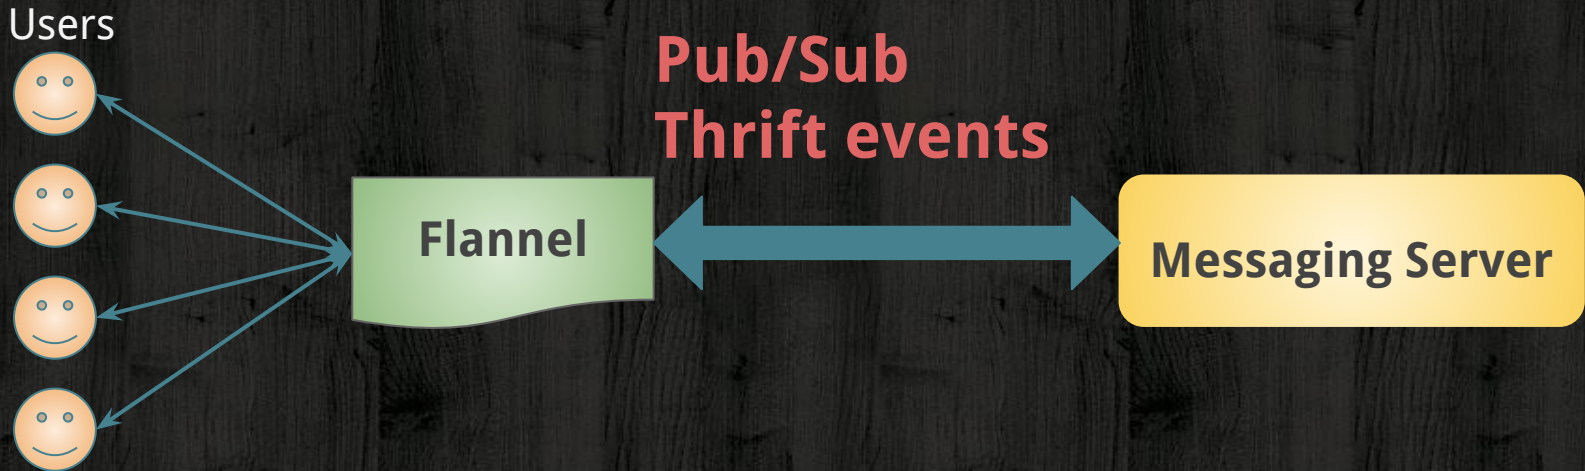
# Old Way of Cache Updates

---



# Publish/Subscribe (Pub/Sub) to Update Cache

---



# Pub/Sub Benefits

---



Less Flannel CPU  
Simpler Flannel code  
Schema data  
Flexibility for cache management

# Flexibility for Cache Management

---

## *Previously*

- ◆ Load when the first user connects
- ◆ Unload when the last user disconnects

# Flexibility for Cache Management

---

## *With Pub/Sub*

- ◆ Isolate received events from user connections





## Another Closer Look

---

- ◆ With Pub/Sub, does Flannel need to be on WebSocket path?

## Next Step

---

Move Flannel out of WebSocket path

## Next Step

---

Move Flannel out of WebSocket path


**Why?**

Separation & Flexibility

# Evolution with Product Requirements

---

## *Grid for Big Enterprise*



Before Grid

---

Team Affinity for Cache Efficiency

Now

---

Team Affinity



Grid Aware

# Grid Awareness Improvements

---

## Flannel Memory



Saves 22G of per host, 1.1TB total

# Grid Awareness Improvements


---

For our biggest customer

DB Shard CPU Idle 25% -> 90%



P99 User Connect Latency 40s -> 4s





# Future

---

Team Affinity



Grid Aware



Scatter & Gather

---

# Outline

- ◆ The slowness problem
- ◆ Incremental Improvements
- ◆ Architecture changes
  - Flannel
  - Client Pub/Sub
  - Messaging Server

# Expand Pub/Sub to Client Side

---

**Client Side Pub/Sub**  
reduces events Clients have to handle

# Presence Events

---



60% of all events

$O(N^2)$

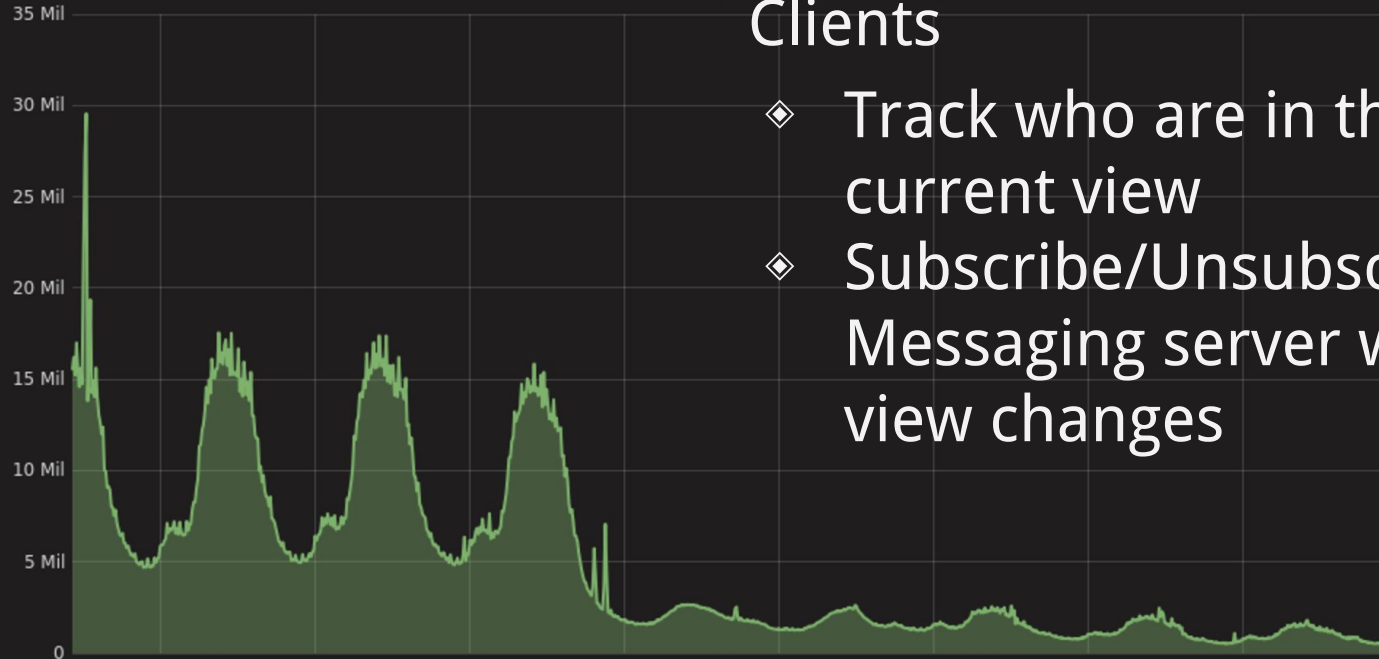
1000 user team  $\Rightarrow$

$1000 * 1000 = 1M$  events



# Presence Pub/Sub

---



## Clients

- ◆ Track who are in the current view
- ◆ Subscribe/Unsubscribe to Messaging server when view changes

---

# Outline

- ◆ The slowness problem
- ◆ Incremental Improvements
- ◆ Architecture changes
  - Flannel
  - Client Pub/Sub
  - Messaging Server

# What is Messaging Server


---

**Messaging Server**



# A Message Router


---





# Events Routing and Fanout

---



# Limitations

---

- ◆ Sharded by Team  
Single point of failure

# Limitations

---

- ◆ Sharded by Team  
Single point of failure
- ◆ Shared Channels  
Shared states among teams



# Topic Sharding

---

## **Everything is a Topic**

public/private channel, DM, group DM,  
user, team, grid

# Topic Sharding

---

 Natural fit for shared channels

# Topic Sharding

---



Natural fit for shared channels



Reduce user perceived failures

# Other Improvements

---

Auto failure recovery



# Other Improvements

---

Auto failure recovery  
Publish/Subscribe

# Other Improvements

---


Auto failure recovery

Publish/Subscribe

Fanout at the edge

# Our Journey

---



# Journey Ahead

A long, straight asphalt road with a yellow center line and white edge lines stretches from the foreground into the distance. The road is flanked by green grassy fields. In the background, there are rolling hills and a range of mountains under a clear sky. Utility poles with power lines are visible on the left side of the road.

Get in touch:  
<https://slack.com/jobs>

---



*Thanks!*

**Any questions?**

@bingw11