


Scaling: From 0 to 25 million





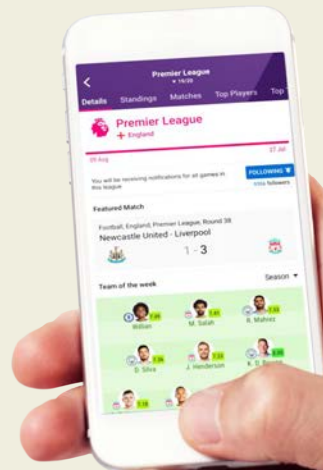
Monthly

28+ million users

1.4 PB+ of traffic

400+ billion requests


1.6M req/s peak





Full page cache

Cache Stampede problem



Proactive caching



+ - 2h



Full page cache





**Know how your
users interact
with your app
(system)**





SofaScore

EL CLÁSICO

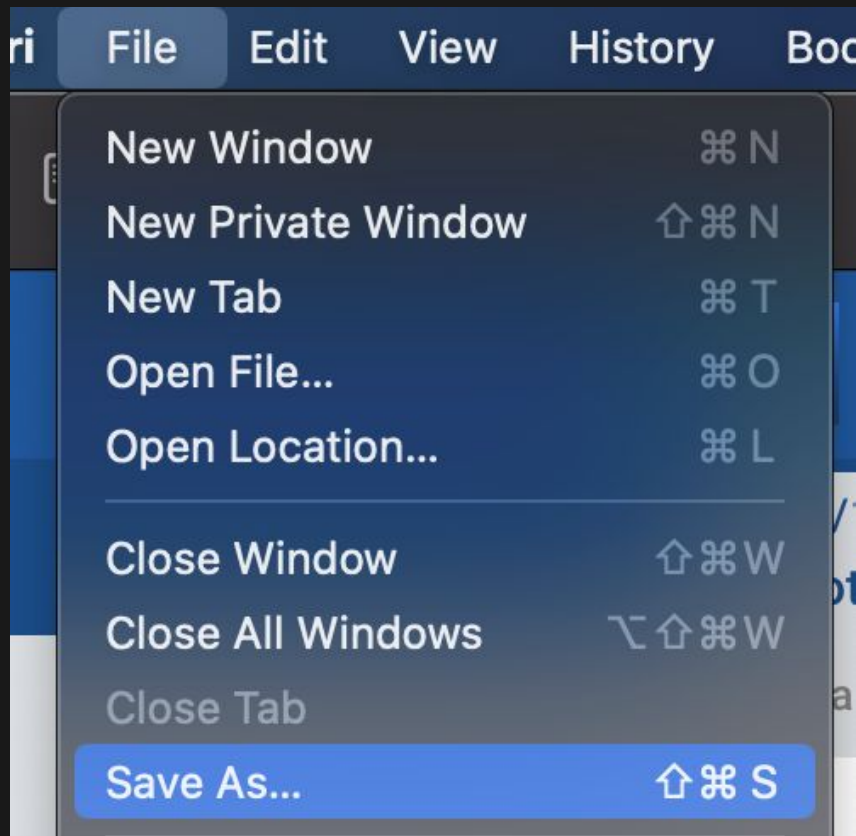


VS















Varnish

HTTP accelerator

The cache is often the
most critical piece of
software for a web-based
business



Request coalescing



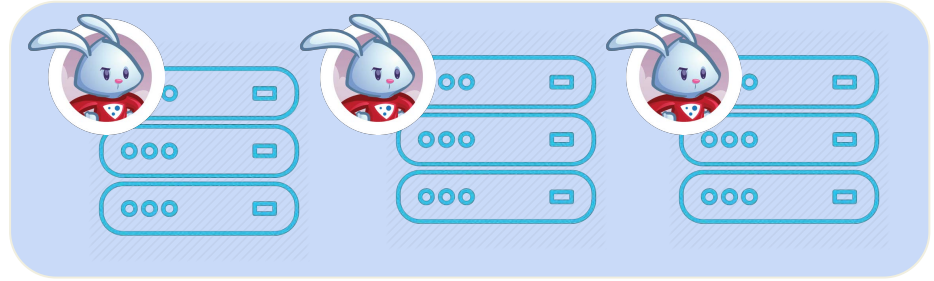
Stateless!



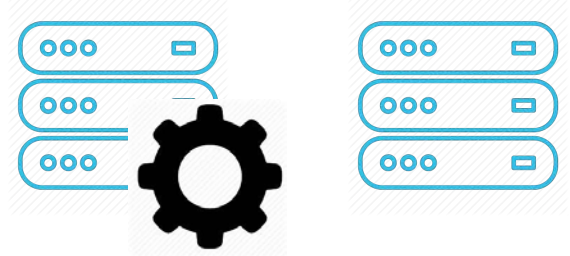
Amazon Machine Image (AMI)



Elastic Load Balancer



Auto Scaling Group



Sofascore





**Be stateless
where possible**





mongoDB®

webscale



mongoDB®

- Replication issues
- Data type errors
- Difficult to analyse data
- Locking
- Foreign keys





Why postgres?

- Open source
- Advanced SQL capabilities
- Minimal locking! (MVCC)
- Great performance!



**~ 1 million
queries / s**

**~ 90k
transactions / s**





Pros and Cons



- Simple Replication
- Simple configuration
- Leader election



- Stability
- Foreign keys, consistency
- Analytical power





**Recognize the
right tool for the
job and change if
necessary**





Microservices?



microservices

serverless

Multi AZ

auto scaling

**YOUR APP
WITH 0 USERS**


imgflip.com



Async communication



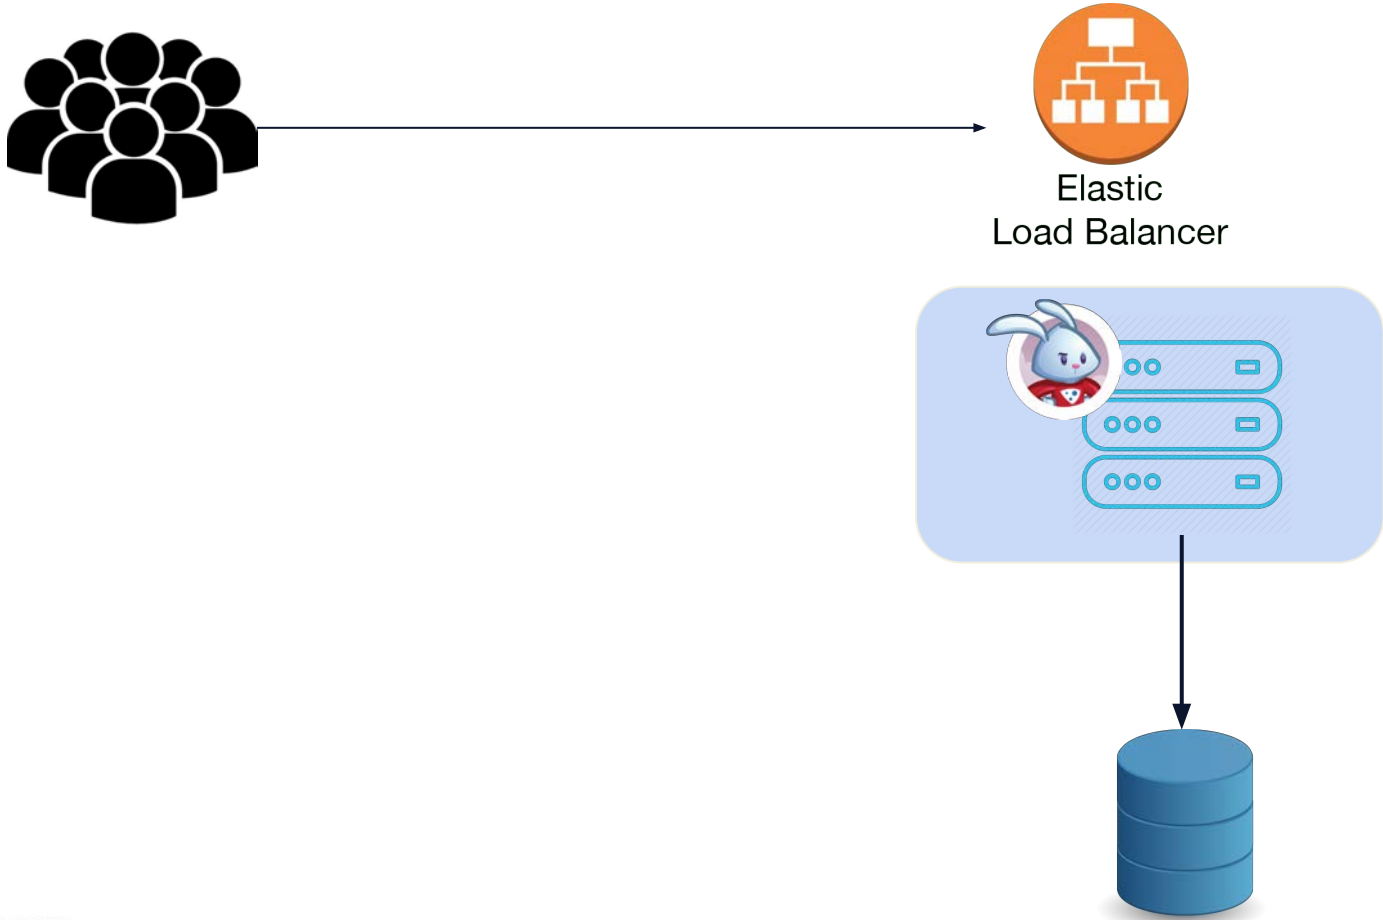
Queues

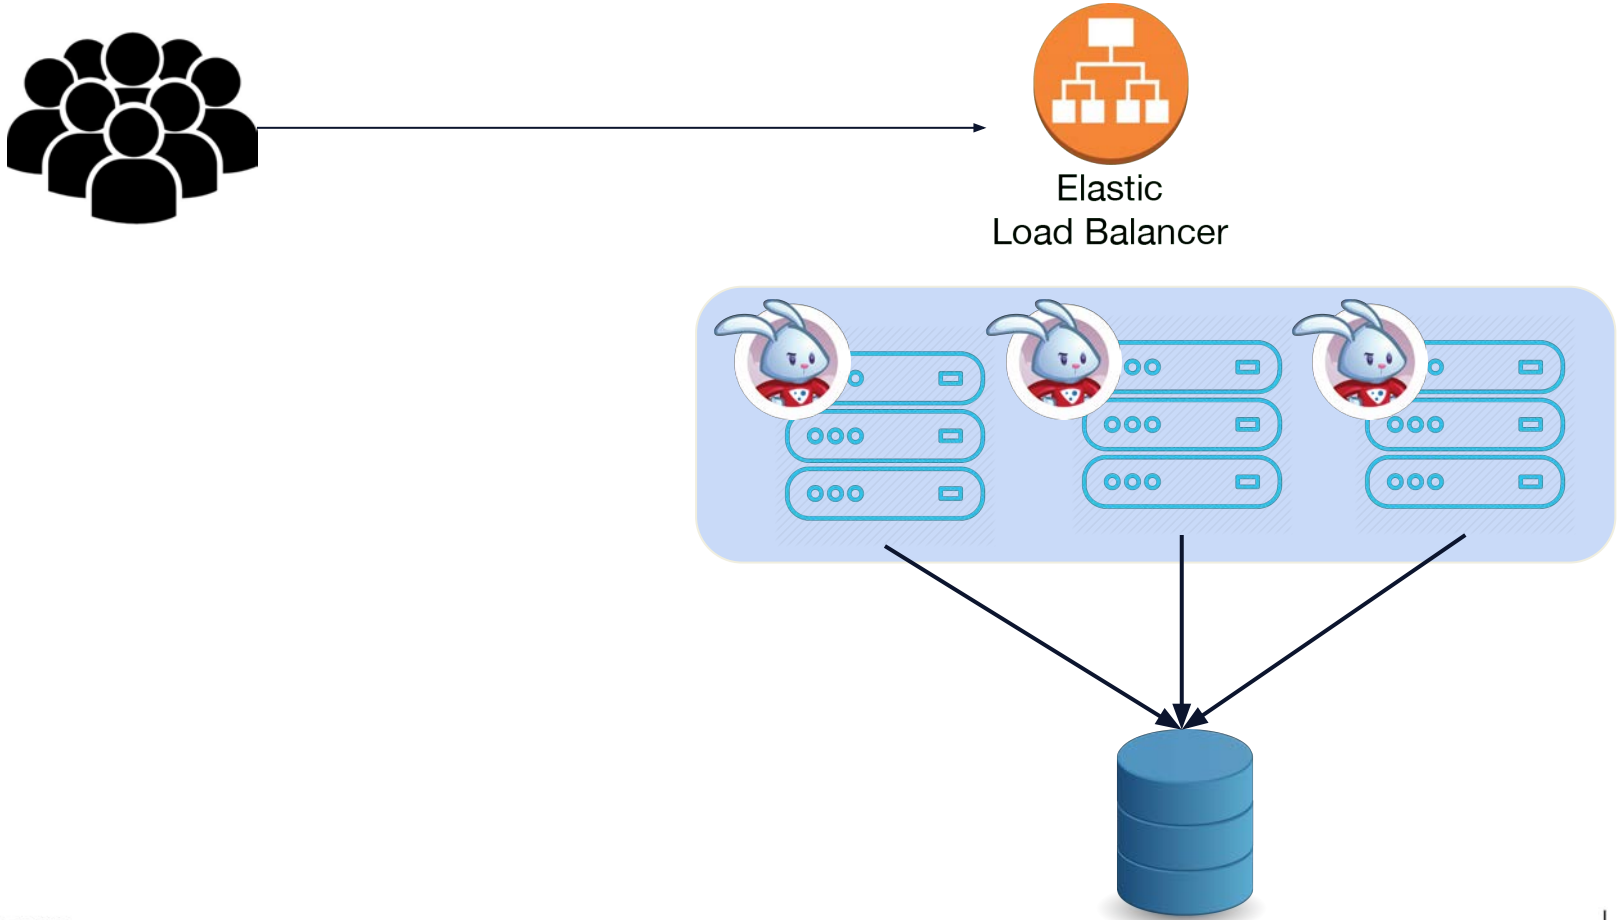





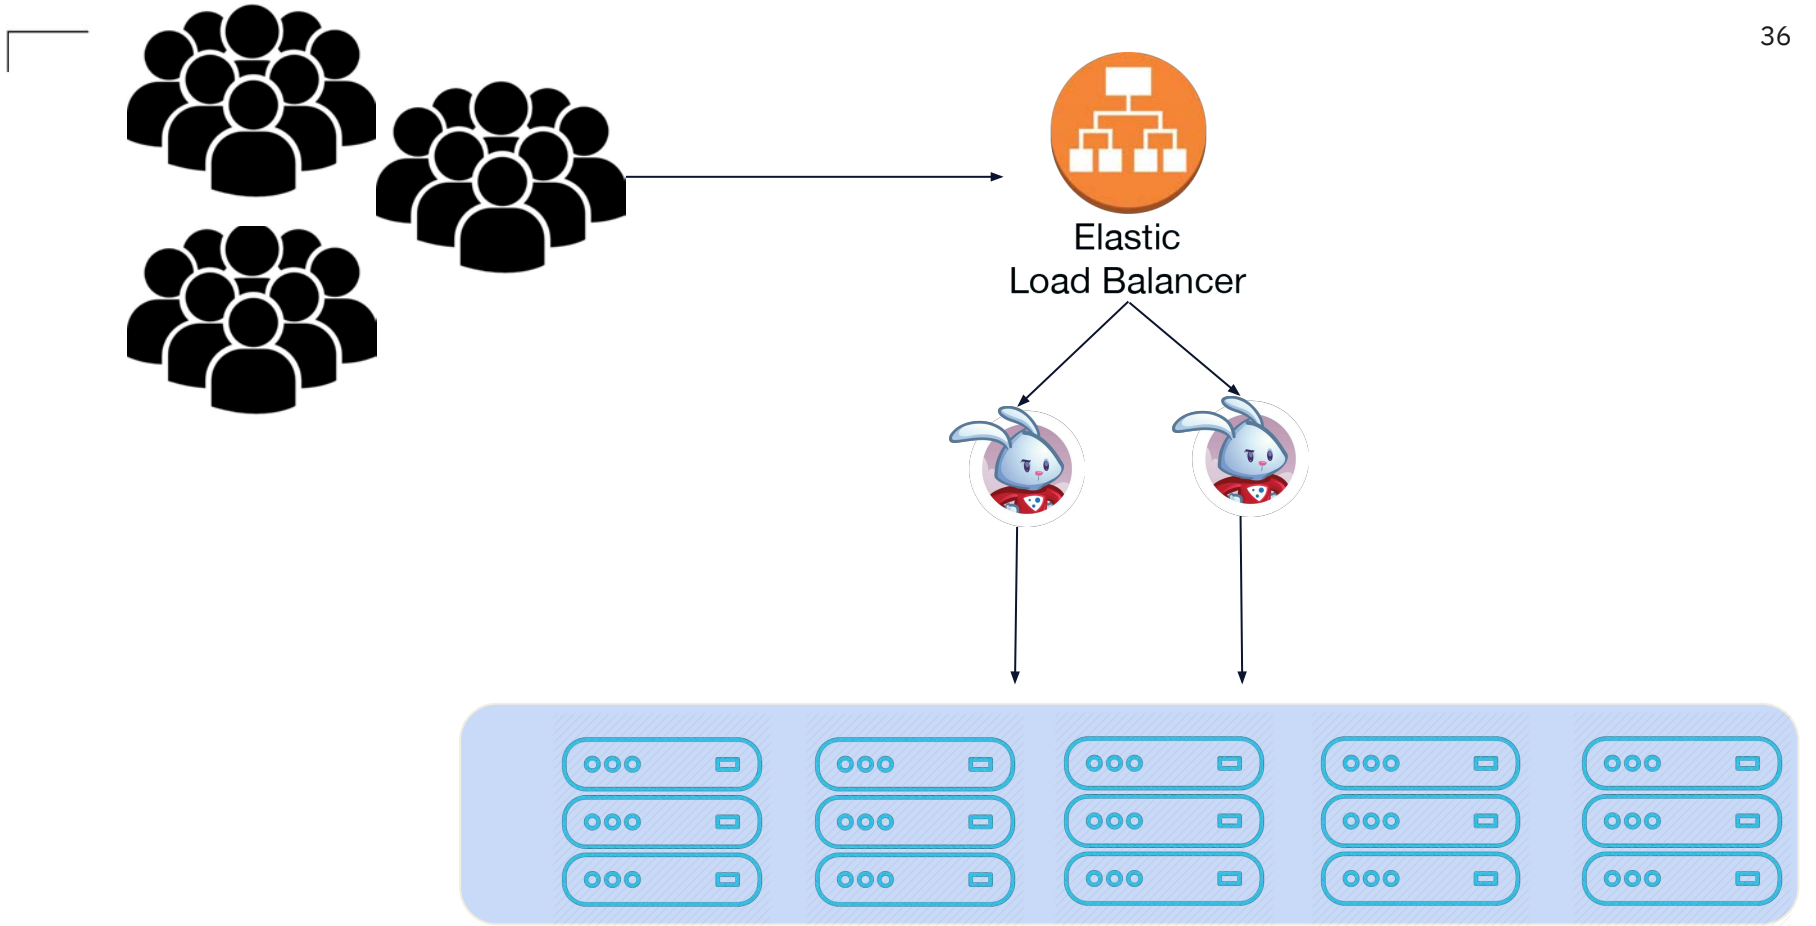
Queue everything
that is slow

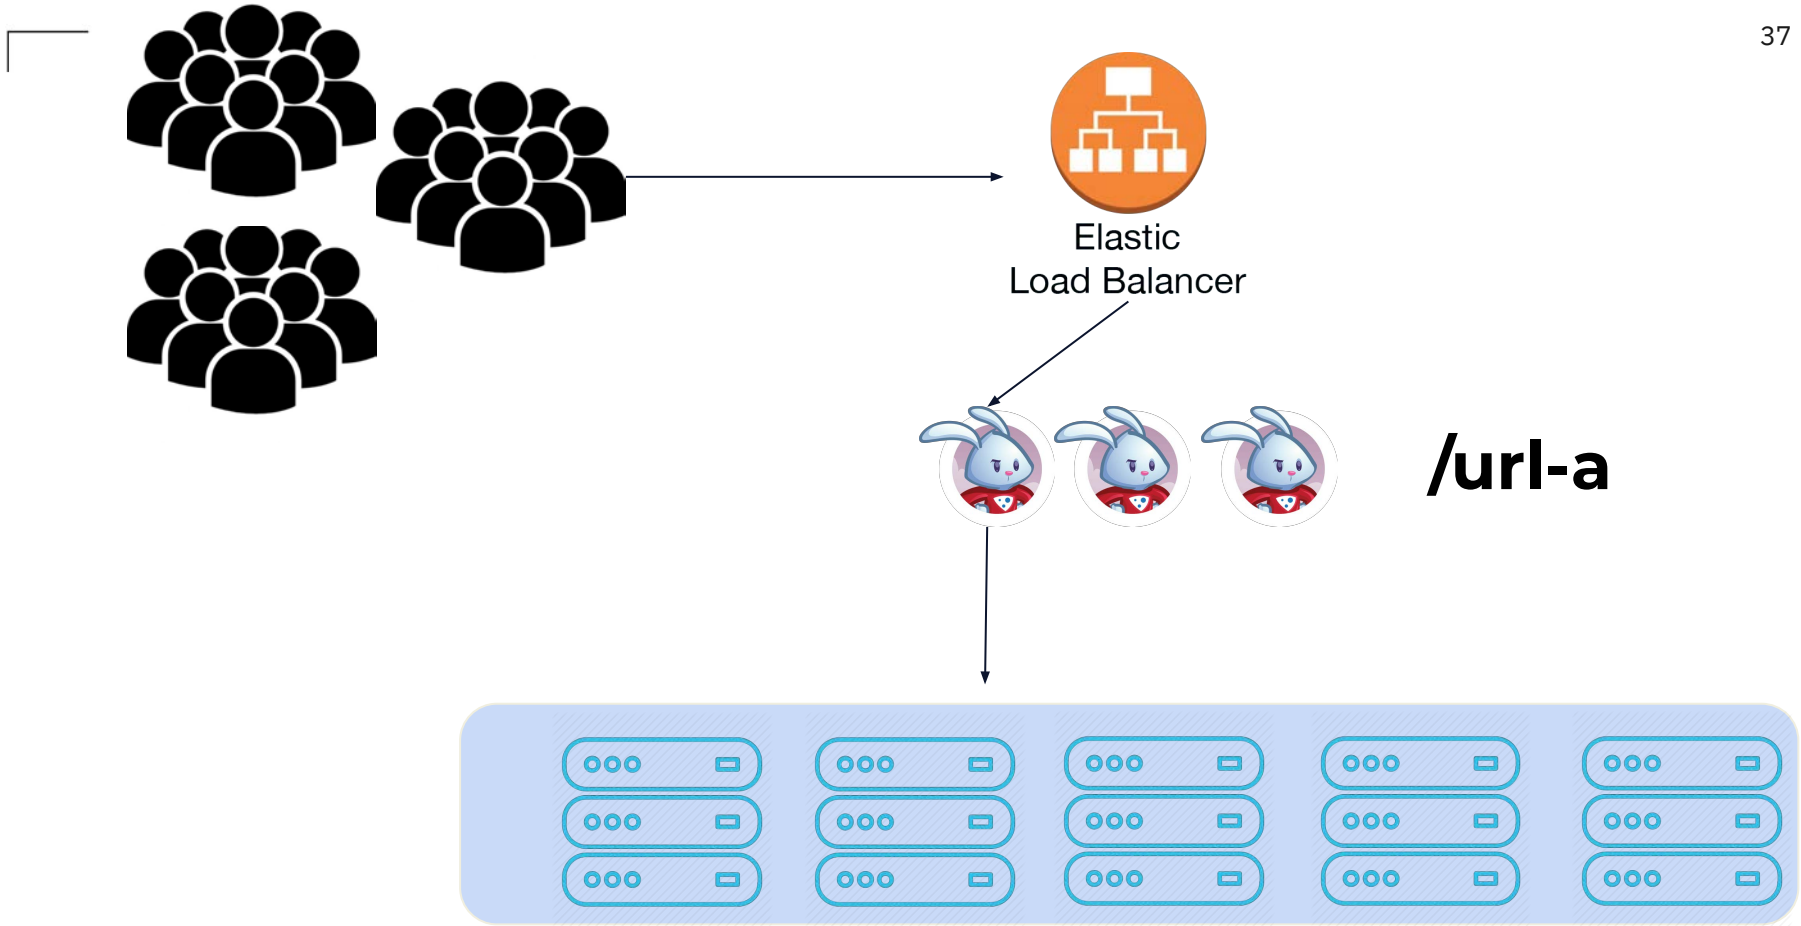


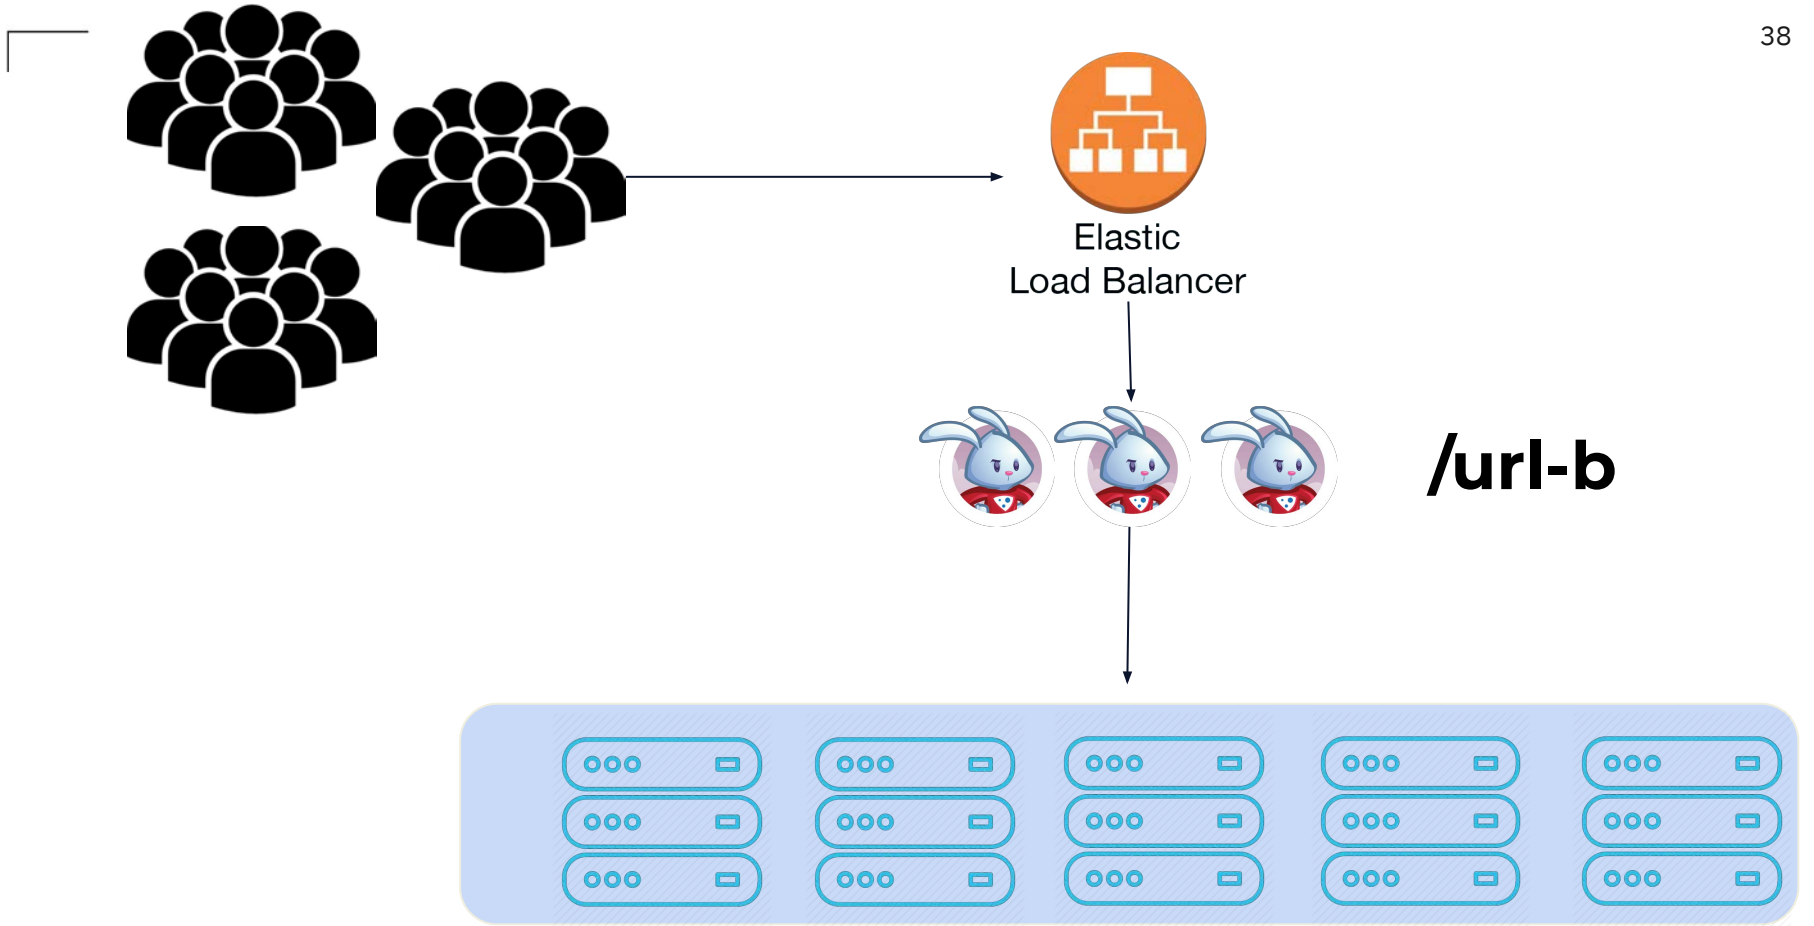


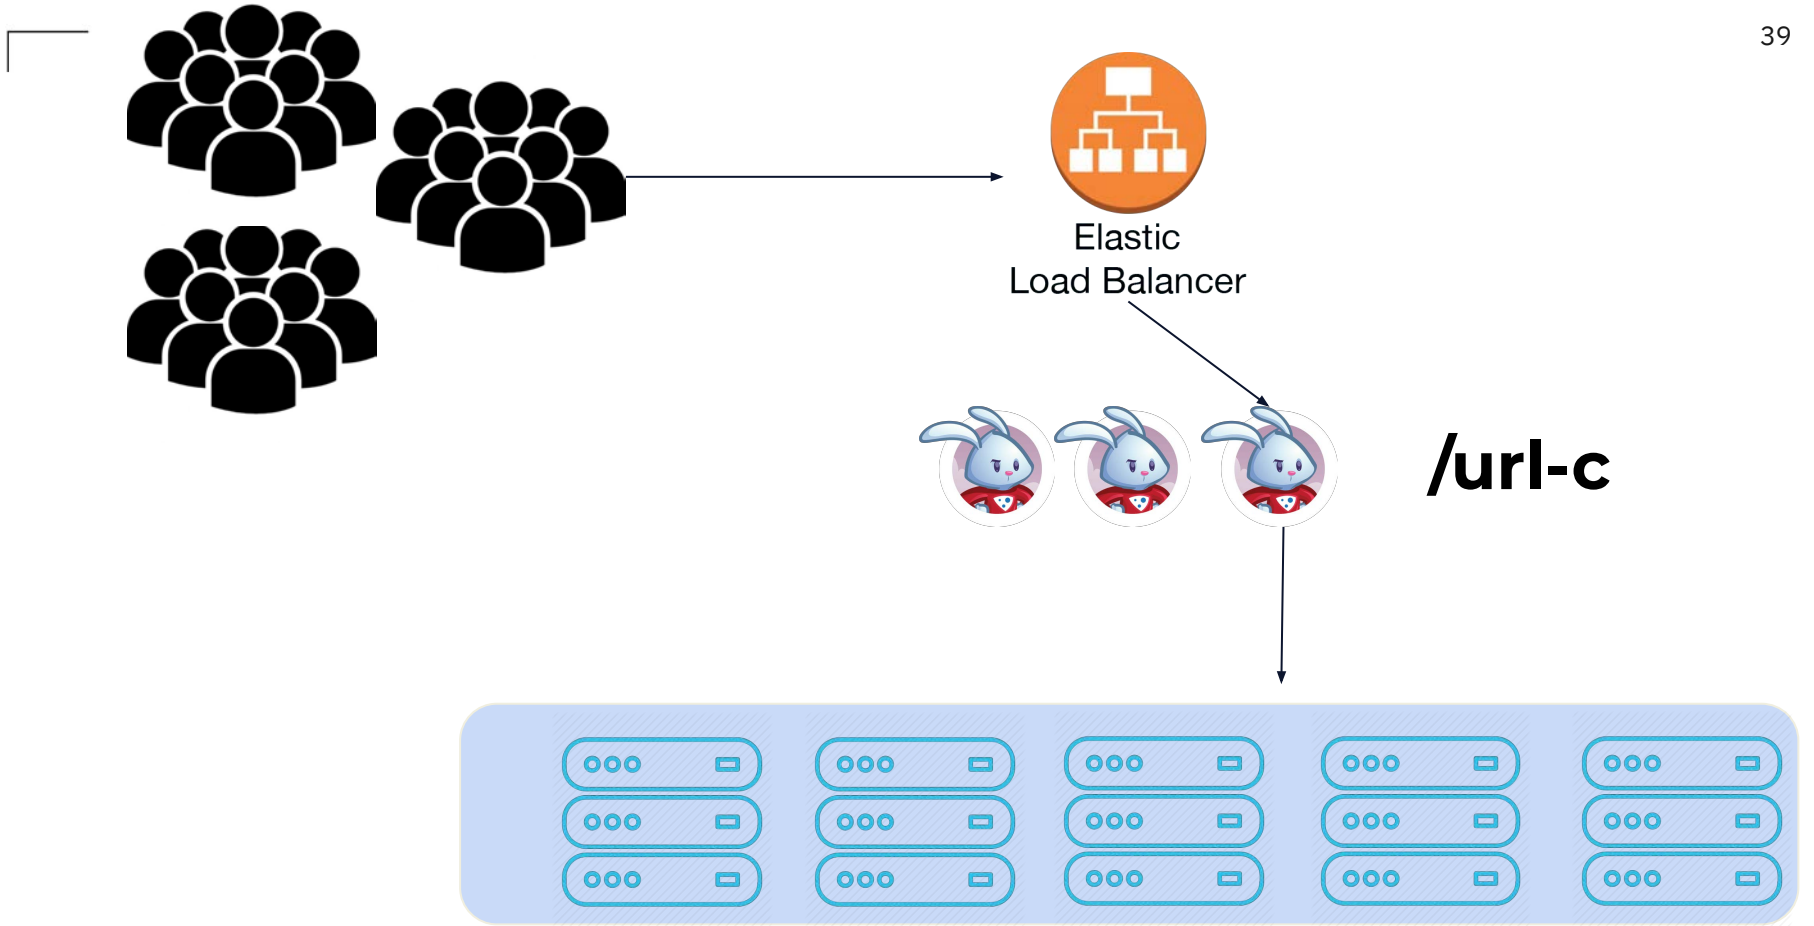


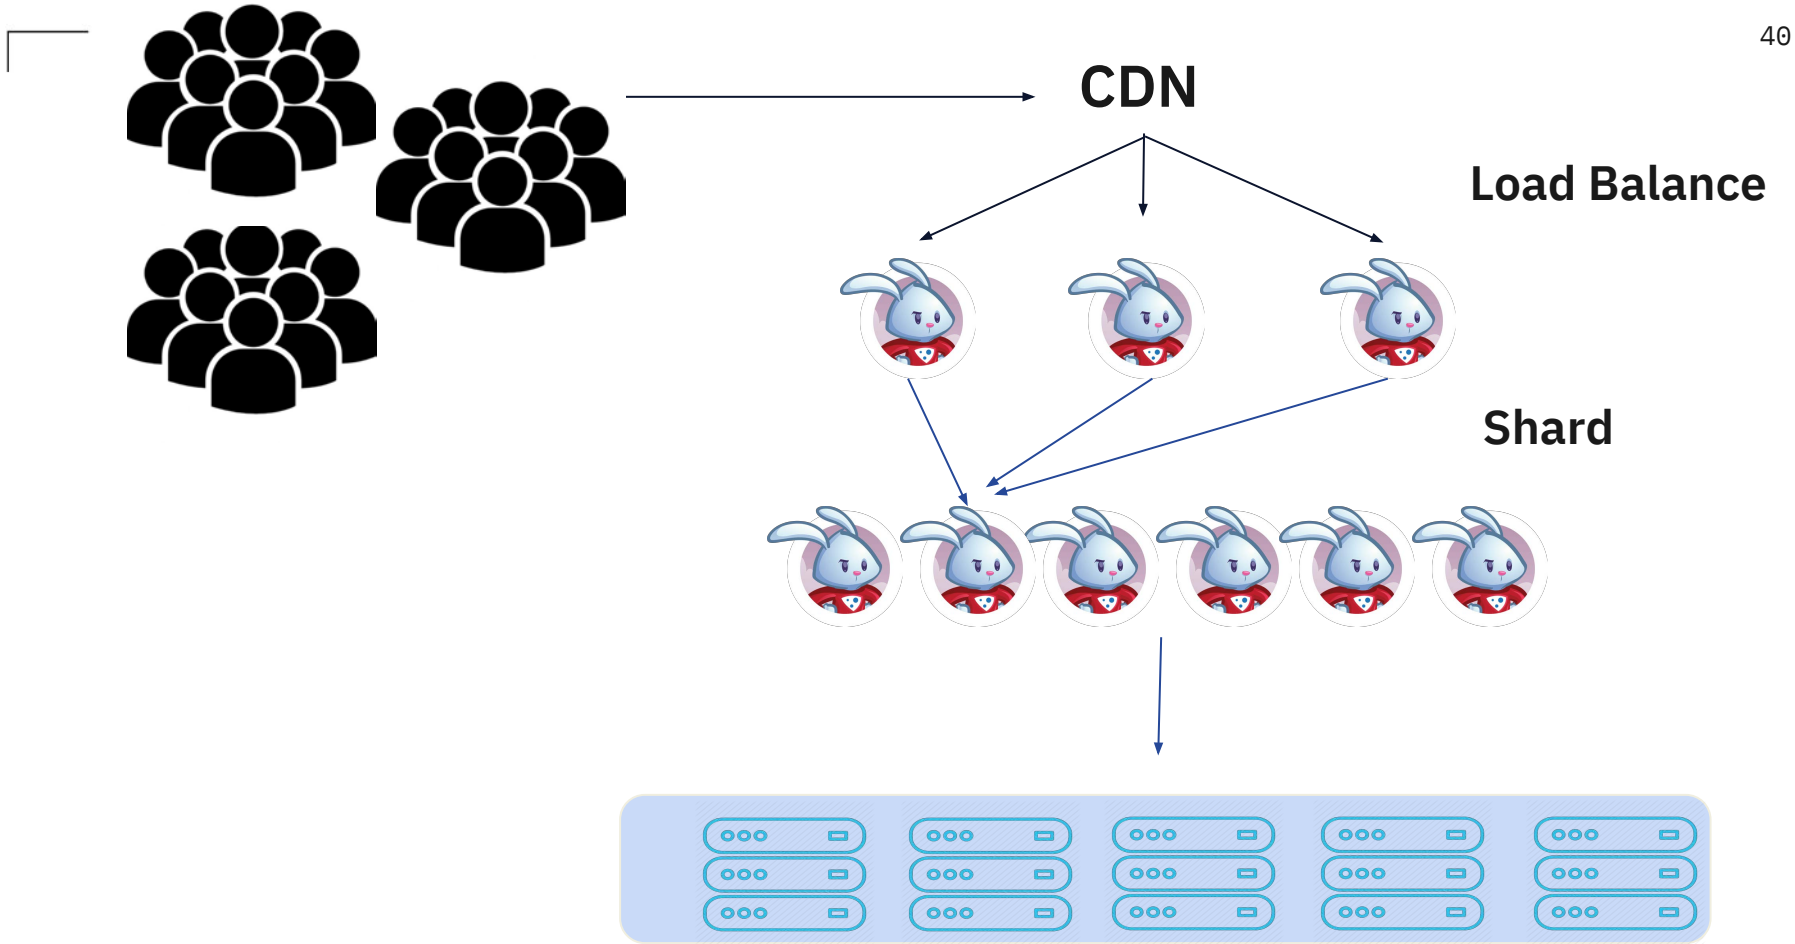













CACHE





**But note: cache
invalidation is hard**



Cache-Control header

max-age: 60

s-maxage: 300

stale-while-revalidate: 3600


stale-if-error: 86400






MB = \$\$\$\$








10x







Migrating everything from cloud to DC









:: Cost and usage [Info](#)

Current month costs

\$1,684,974.20

Forecasted month end costs

\$2,002,283.67 Up 446% over last month

Last month costs

\$366,573.53



Lot's of spare RAM

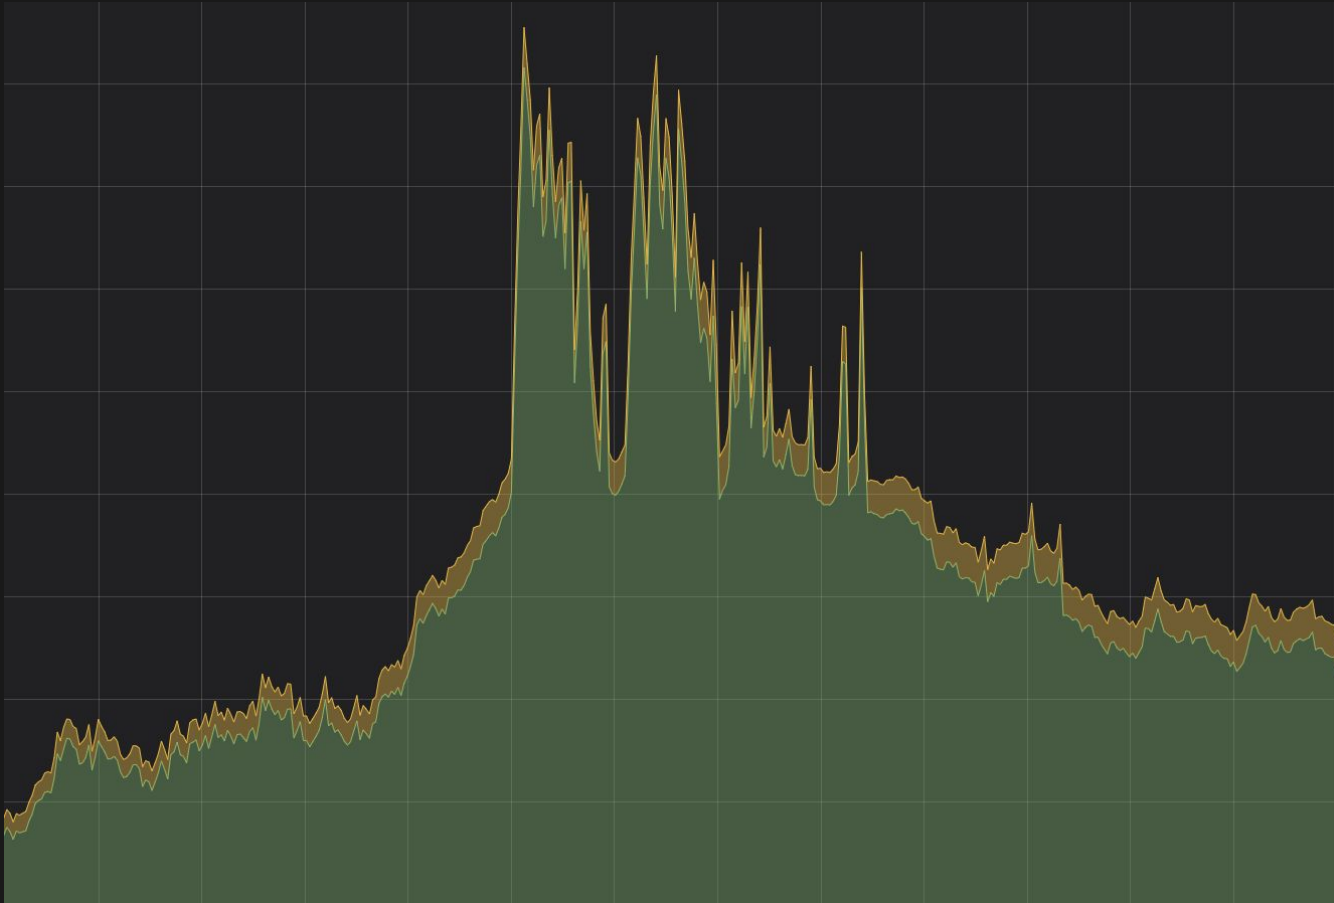
Cache!





Cost of people?






APM

Type


56

/sofa_web_event_list_json	5.9%
/sofa_web_event_details_json	5.64%
/app_api_v1_user_events	4.65%
/app_api_v1_event_votes	3.84%
/sofascore_mobile_mobile4_event_head2head	3.79%
/sofa_web_team	3.48%
/app_api_v1_player_events	3.05%
/sofa_web_event_details	2.84%
/sofa_web_event_details_lineups_json	2.62%
/sofascore_mobile_mobile4_event_vote	2.55%
/sofa_web_event_details_amp	2.45%
/sofascore_mobile_mobile4_team_events	1.92%
/sofascore_mobile_mobile4_event_tweets	1.57%

Optimizations



Longer cache





**Monitor everything
optimize something**

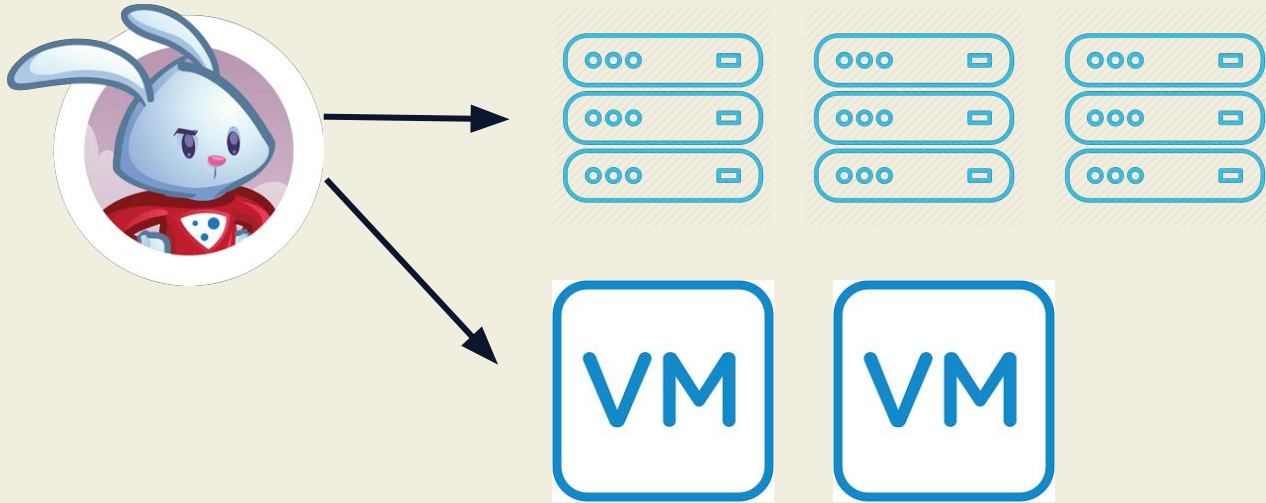


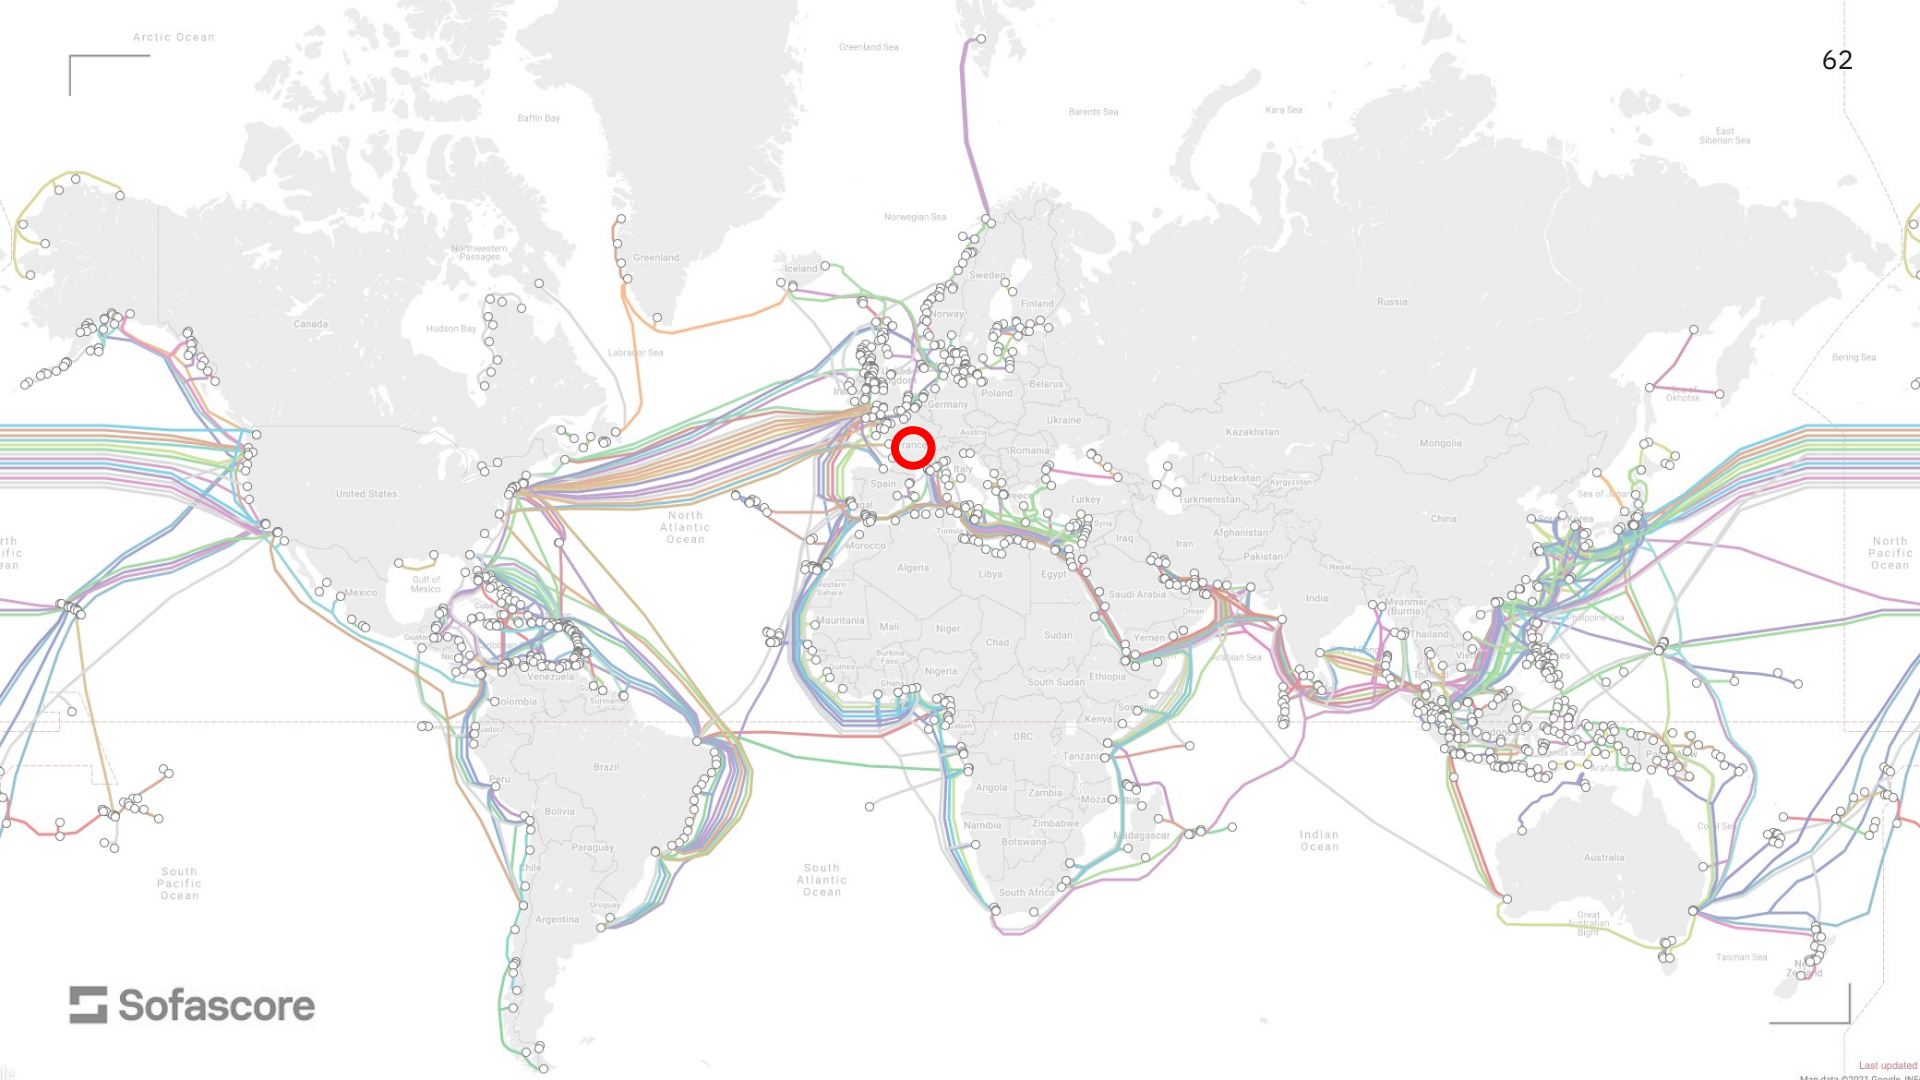


**What if the
number of users
exceeds current
capacity?**



Dedicated + VPS







Josip Stuhli

Nov 10, 2020, 10:25 PM



to [REDACTED]

Hi,

I hope this email finds you well.

I have been informed by other mods that you are from New Zealand. I'm interested in seeing what kind of user experience mobile users from countries outside Europe get using SofaScore.

Could you record your screen while using SofaScore app so I can see what kind of latency are you getting? You can open the app, go through live matches, open details, maybe some tabs, such as lineups and statistics. Maybe open up match list for Saturday/Sunday, team/player info and such.

I'm hoping to use that info to speed up the loading of the data in the app itself.

Thanks,

Josip Stuhli




SofaScore CTO

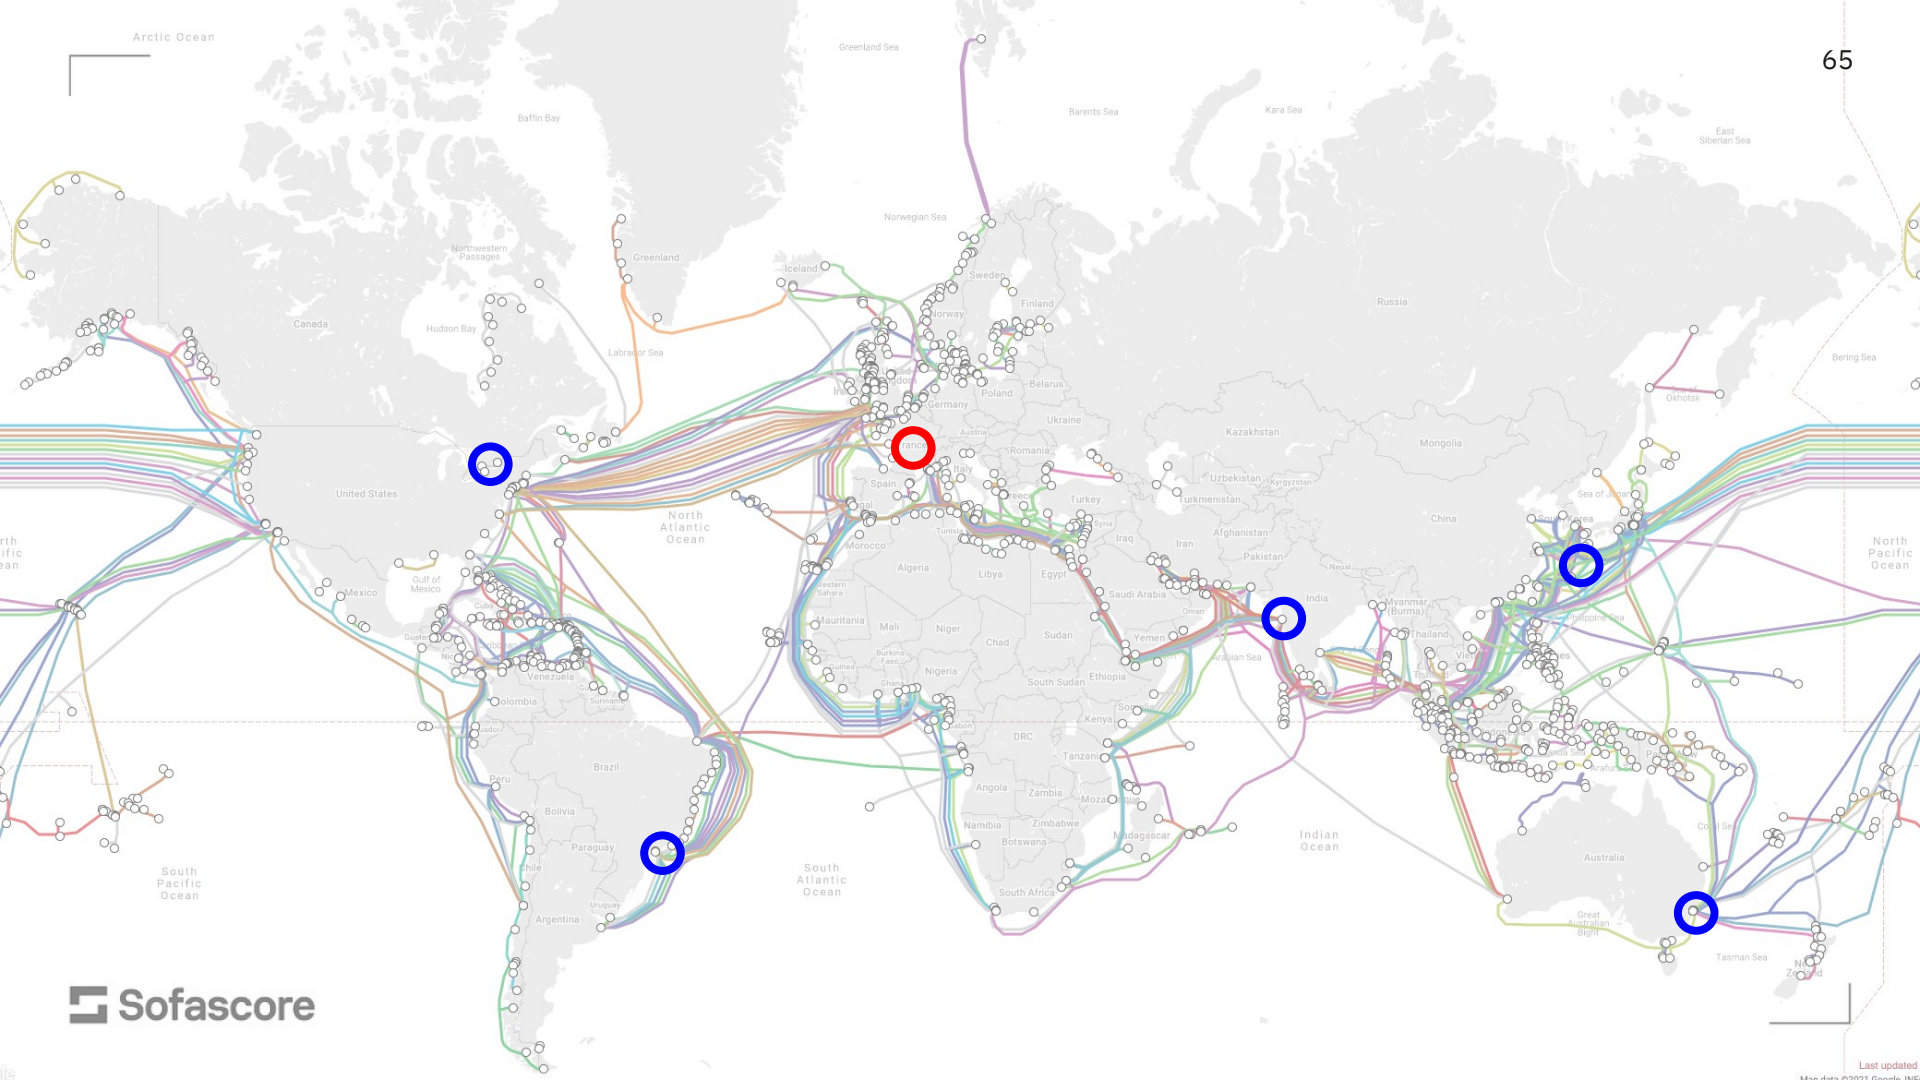



Christopher Karkau [@chris_karkau](#)

to me ▾

Yes mate I'm from New Zealand but living in Australia and everything seems to be working fine







Australia




Shouldn't CDN do
that?







kubernetes



LONGHORN



Real time updates



NATS Messaging



NATS is an open-source messaging system. The NATS server is written in the Go programming language. Client libraries to interface with the server are available for dozens of major programming languages. The core design principles of NATS are performance, scalability, and ease of use. [Wikipedia](#)

Developer(s): [Synadia Communications, Inc.](#)

Initial release: 2011

License: Apache 2.0

Stable release: 2.8.2 / 4 May 2022; 6 days ago

Written in: [Go](#)



500k+ msgs/sec



What to do with 2 PB of data?





1.75 PB -> 46.8TB
1.4 Trillion rows

Insert:

rows_per_sec

2.15 Mil

ClientIP	count()	uniq_methods	uniq_user_agent	uniq_routes	min(EdgeStartTimes	max(EdgeStartTime:
203.161.33.185	9382	1	1	14	2024-11-19 17:31:13	2024-11-19 18:30:12
185.51.61.123	5151	1	1	3	2024-11-19 18:10:17	2024-11-19 18:27:37
192.64.113.199	3706	1	1	3	2024-11-19 17:31:00	2024-11-19 18:29:19
223.19.209.2	2725	1	2	31	2024-11-19 17:31:08	2024-11-19 18:29:32
46.134.139.172	2643	1	1	21	2024-11-19 17:31:00	2024-11-19 17:42:41
2a02:8070:6482:7...	2335	2	3	6	2024-11-19 18:26:16	2024-11-19 18:29:33
41.113.128.118	2221	1	1	40	2024-11-19 17:32:57	2024-11-19 18:29:29
203.175.9.145	1840	1	234	10	2024-11-19 17:31:11	2024-11-19 18:29:25
36.134.84.234	1591	1	1	6	2024-11-19 17:31:03	2024-11-19 18:29:04
98.201.93.87	1494	2	1	43	2024-11-19 17:38:28	2024-11-19 18:29:16

Clickhouse + Superset


MAU YoY

October 2024

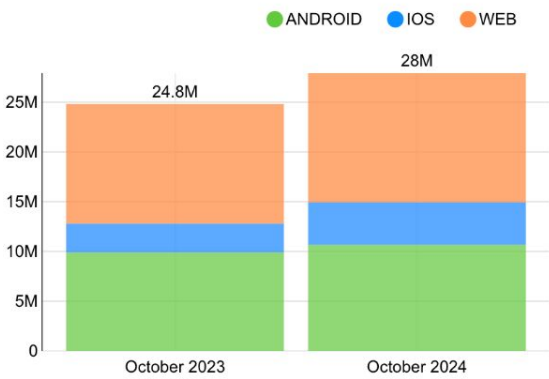
28M

+12.5% compared to previous year


MAU Over Time



MAU YoY Composition



MAU YoY Change Over Time



- hot
- atives
- leads overview
- s
- quisition
- age
- quisition cohorts
- ce purchases
- ns
- journey
- journey
- ns
- nd awareness ove...

All Users Add comparison +

Realtime overview


ACTIVE USERS IN LAST 30 MINUTES

1,914,892

ACTIVE USERS IN LAST 5 MINUTES

509,336

ACTIVE USERS PER MINUTE



View user snapshot



Active users by First user source

#1 (direct)

119,891

66.38%



FIRST USER SOURCE

ACTIVE USERS

FIRST USER SOURCE	ACTIVE USERS
(direct)	120K
google	37K
google-play	17K
Apple	5K

Active users by Audience

#1 All Users

15,382

11.44%



AUDIENCE

ACTIVE USERS

AUDIENCE	ACTIVE USERS
All Users	15K
Growth - Android use...	15K
ads_click custom les...	14K
Open event (Do) GA4	11K

Views by Page title and screen name

#1 Football - scor...ds | Sofascore

10,080

18%



PAGE TITLE AND S...

VIEWS

PAGE TITLE AND S...	VIEWS
Football - scores, sch...	10K
Crystal Palace vs Ma...	2.6K
Calcio - risultati, cale...	2.1K
logos de Futebol de	1.9K

Event count by Event name

#1 screen_view_custom

41,157,034

32.72%



EVENT NAME

EVENT COUNT

EVENT NAME	EVENT COUNT
screen_view_custom	41M
ads_impression_cust...	15M
ad_impression	12M
ads_value	12M

Infra cost is 0.6% of revenue

Take home

- Cache all the things!
- Be stateless where possible
- Know how users interact with the app
- Recognize the right tool for the job and change if necessary
- Cloud should be the default, but keep your options open
- Queue everything that is slow
- Monitor everything - optimize something
- Keep it simple for as long as you can

Thank you!



<https://www.linkedin.com/in/jstuhli/>