App Engine Nitty Gritty

Scalability, Fault Tolerance, and Integrating Amazon EC2

Questions: tinyurl.com/AppEngineNittyGritty

In the right place?

FrontSeat.org

Civic Software

Jesse Kocher
Dave Peck
Josh Livni

Obama CTO

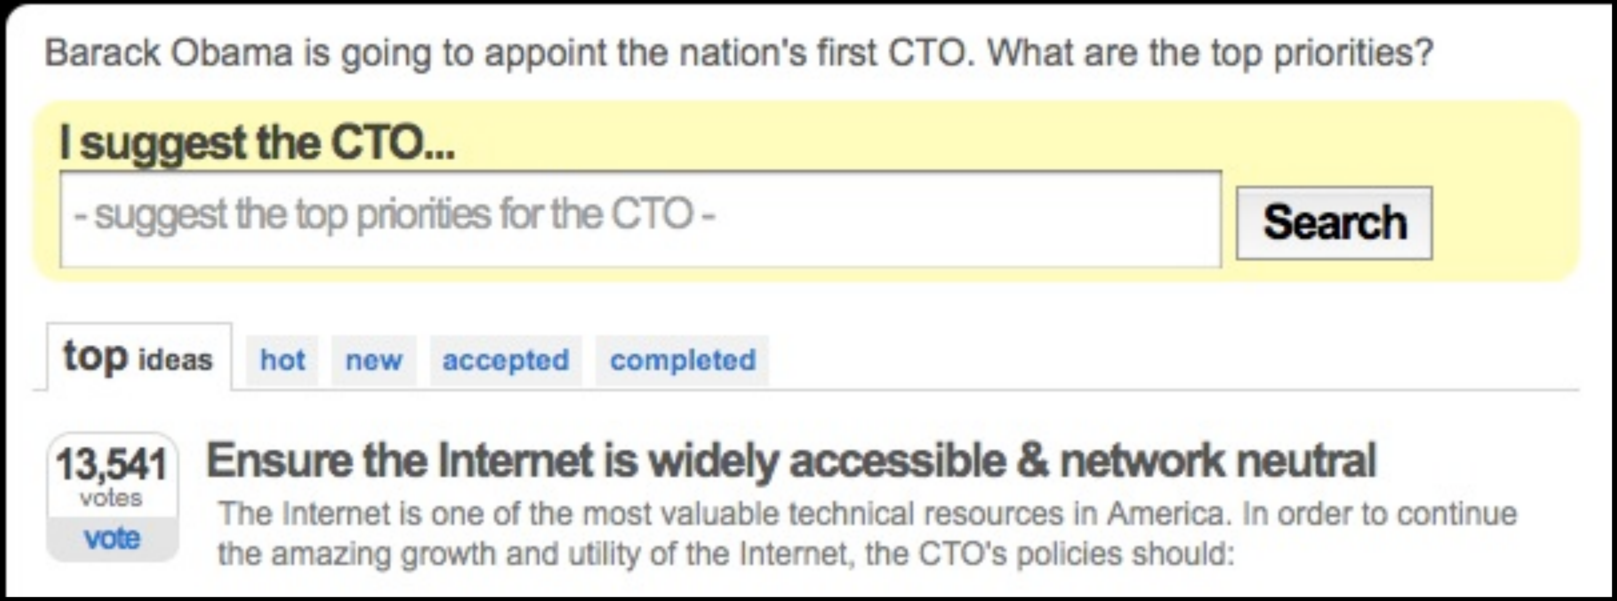

Barack Obama is going to appoint the nation's first CTO. What are the top priorities?

Ensure the Internet is widely accessible & network neutral

The Internet is one of the most valuable technical resources in America. In order to continue the amazing growth and utility of the Internet, the CTO's policies should:

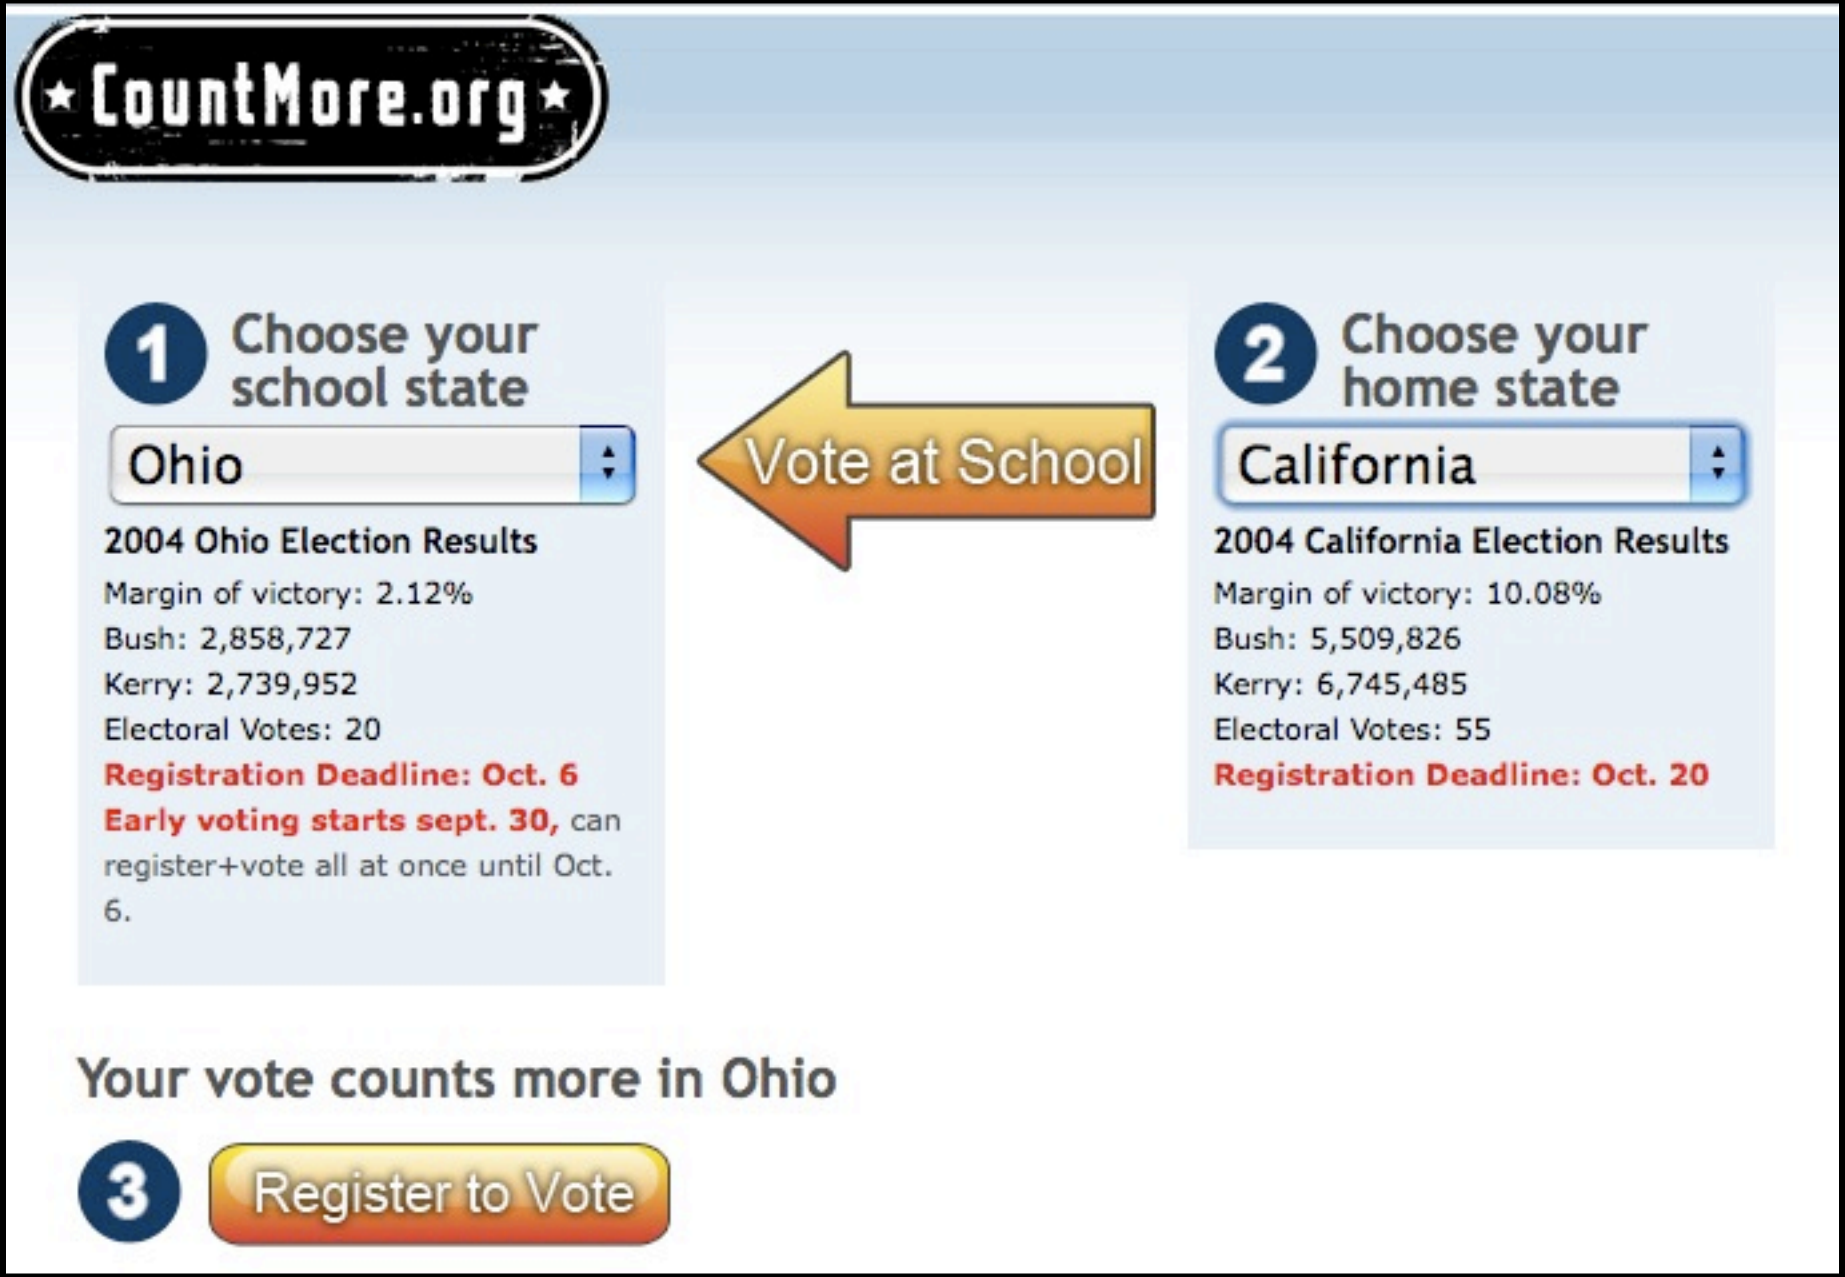

1 Choose your school state

Ohio

2004 Ohio Election Results

Margin of victory: 2.12%

Bush: 2,858,727

Kerry: 2,739,952

Electoral Votes: 20

Registration Deadline: Oct. 6

Early voting starts sept. 30, can register+vote all at once until Oct. 6.

2 Choose your home state

California

2004 California Election Results

Margin of victory: 10.08%

Bush: 5,509,826

Kerry: 6,745,485

Electoral Votes: 55

Registration Deadline: Oct. 20

Your vote counts more in Ohio

3 Register to Vote

Obama CTO

★ CountMore.org ★

Seattle City Light Bill

PLA
Predatory Lending Association

Obama CTO

★ CountMore.org ★

Walk Score™

CITY OF SEATTLE

Seattle City Light Bill

Current Billing

Seattle City Light Bill

Kathy's Climate Kit
The gift of action

PLA
Predatory Lending Association

Walk Score

demo

Why walkability?

- Climate
- Health
- Social capital
- Transit oriented
- Few auto-related deaths

Why walkability?

Home values

Reduced transportation costs

Local businesses

1. Create demand for walkable neighborhoods

2. Provide transparency about the walkability of every property

demo

Bedrooms: 3

Bathrooms: 2

Walk Score: 86

Bedrooms: 3

Bathrooms: 2

Walk Score: 86

Search By Walk Score

Walk Score API

on

Google App Engine

API workflow

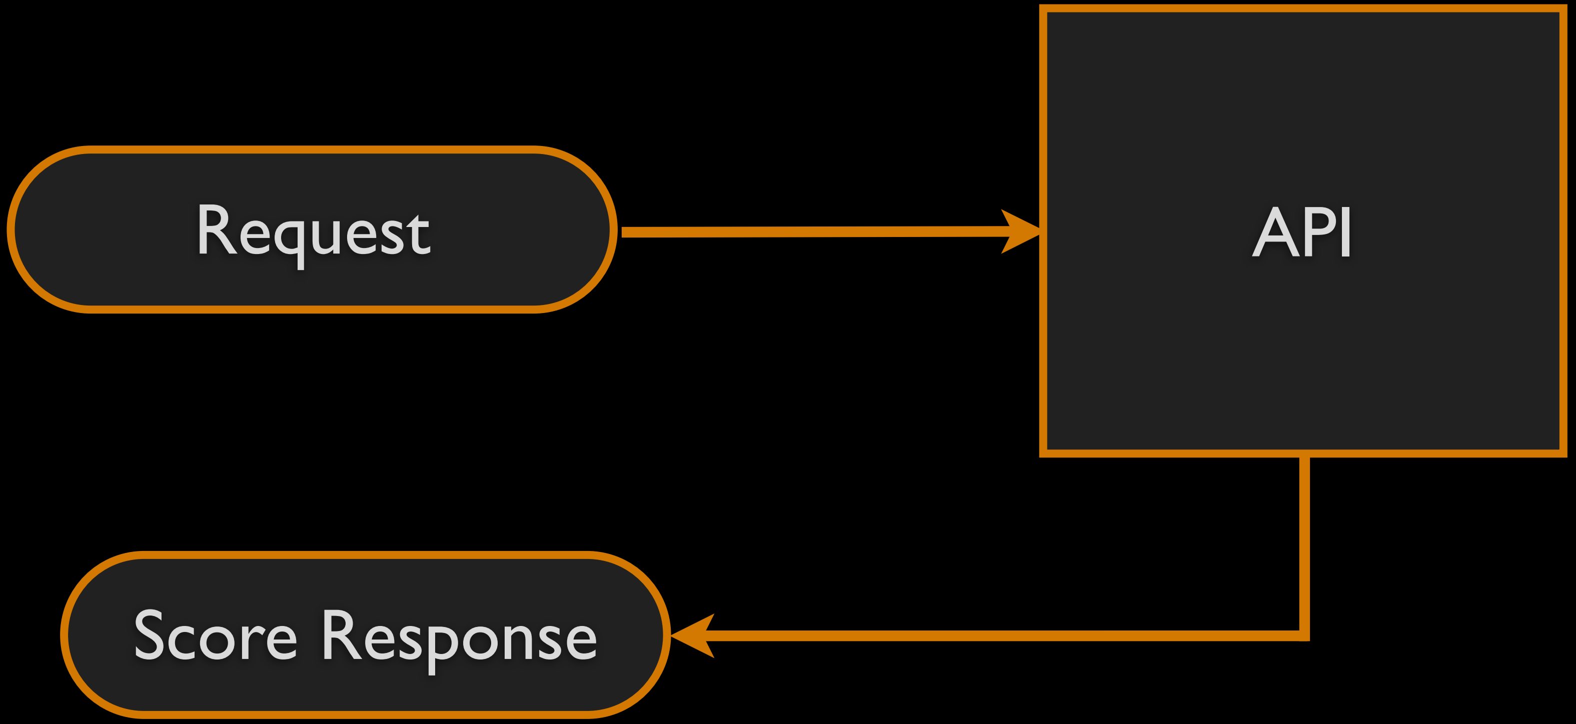

Considerations

Vendor lock in

Cost

AppEngine Failures

Criteria

Core: Respond to a score request

- core functionality should always work.

Secondary: (various)

- let secondary functionality fail by itself

Not a Veg-O-Matic

Great for simple CRUD

Fails in some other areas

Rankings

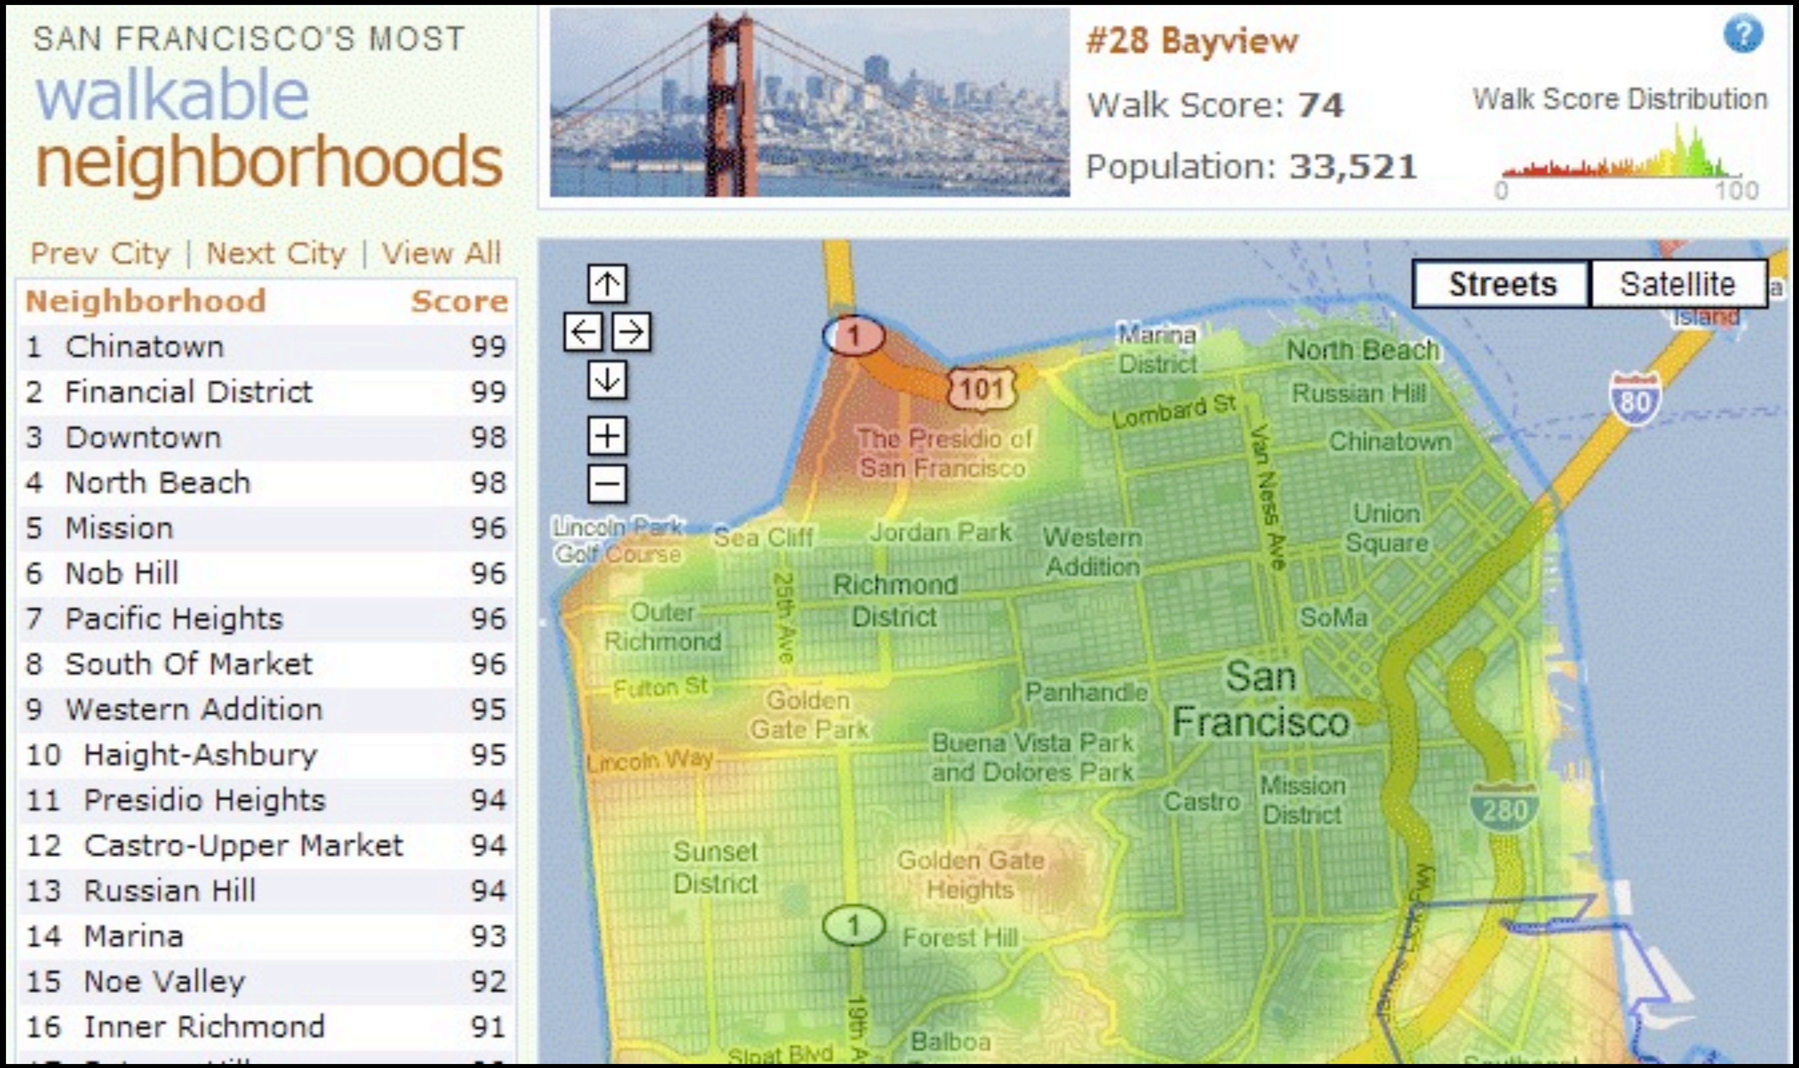

SAN FRANCISCO'S MOST walkable neighborhoods

Prev City | Next City | View All

Neighborhood	Score
1 Chinatown	99
2 Financial District	99
3 Downtown	98
4 North Beach	98
5 Mission	96
6 Nob Hill	96
7 Pacific Heights	96
8 South Of Market	96
9 Western Addition	95
10 Haight-Ashbury	95
11 Presidio Heights	94
12 Castro-Upper Market	94
13 Russian Hill	94
14 Marina	93
15 Noe Valley	92
16 Inner Richmond	91

#28 Bayview

Walk Score: **74**

Population: 33,521

API Usage

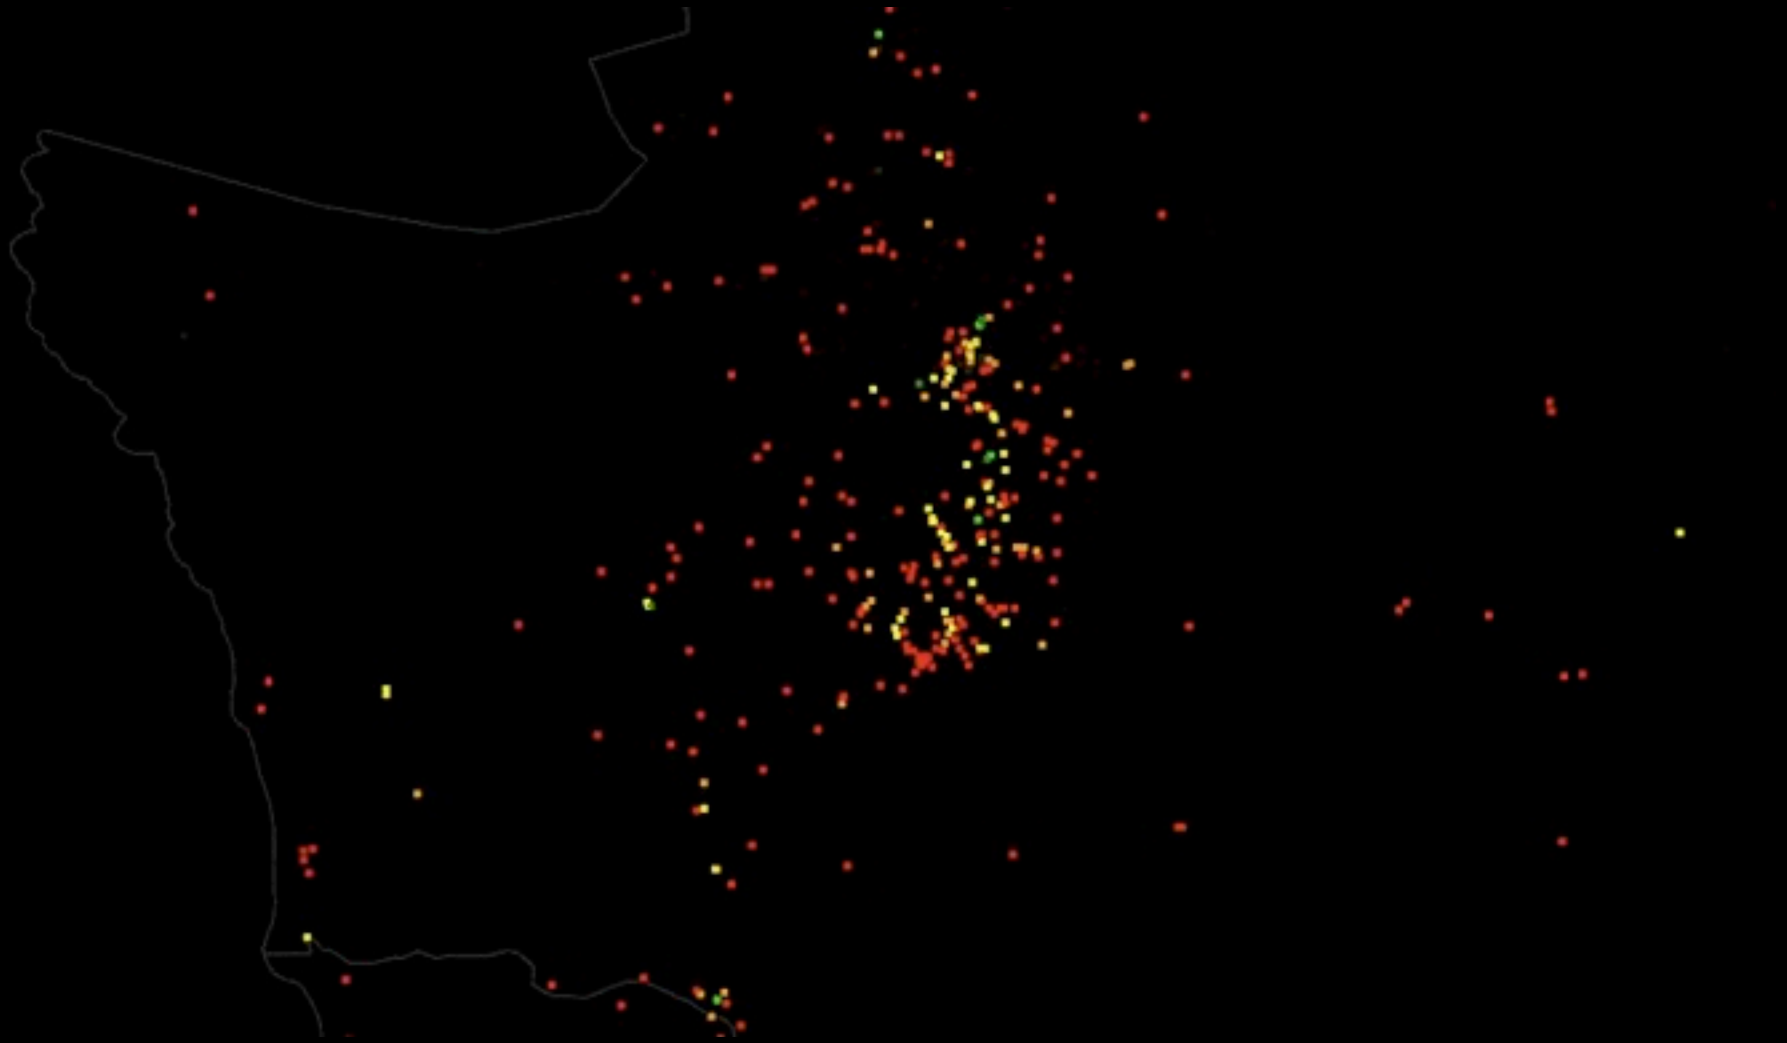

API Usage

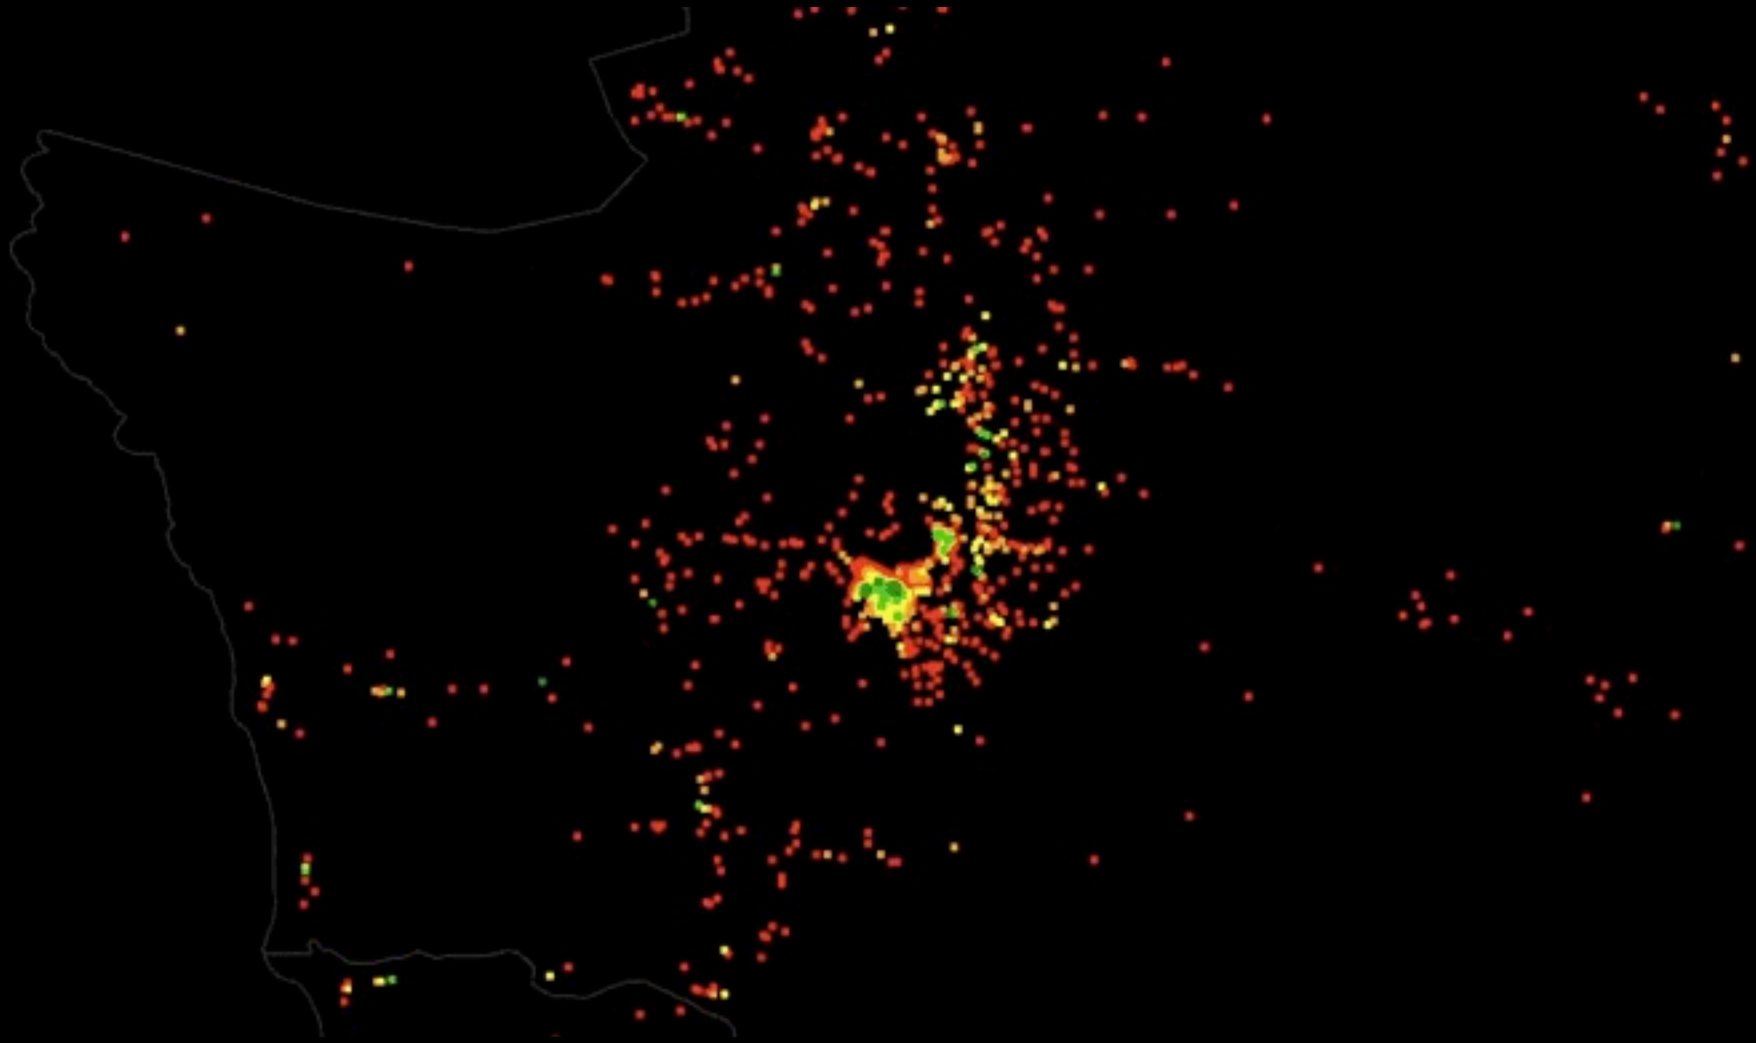

API Usage

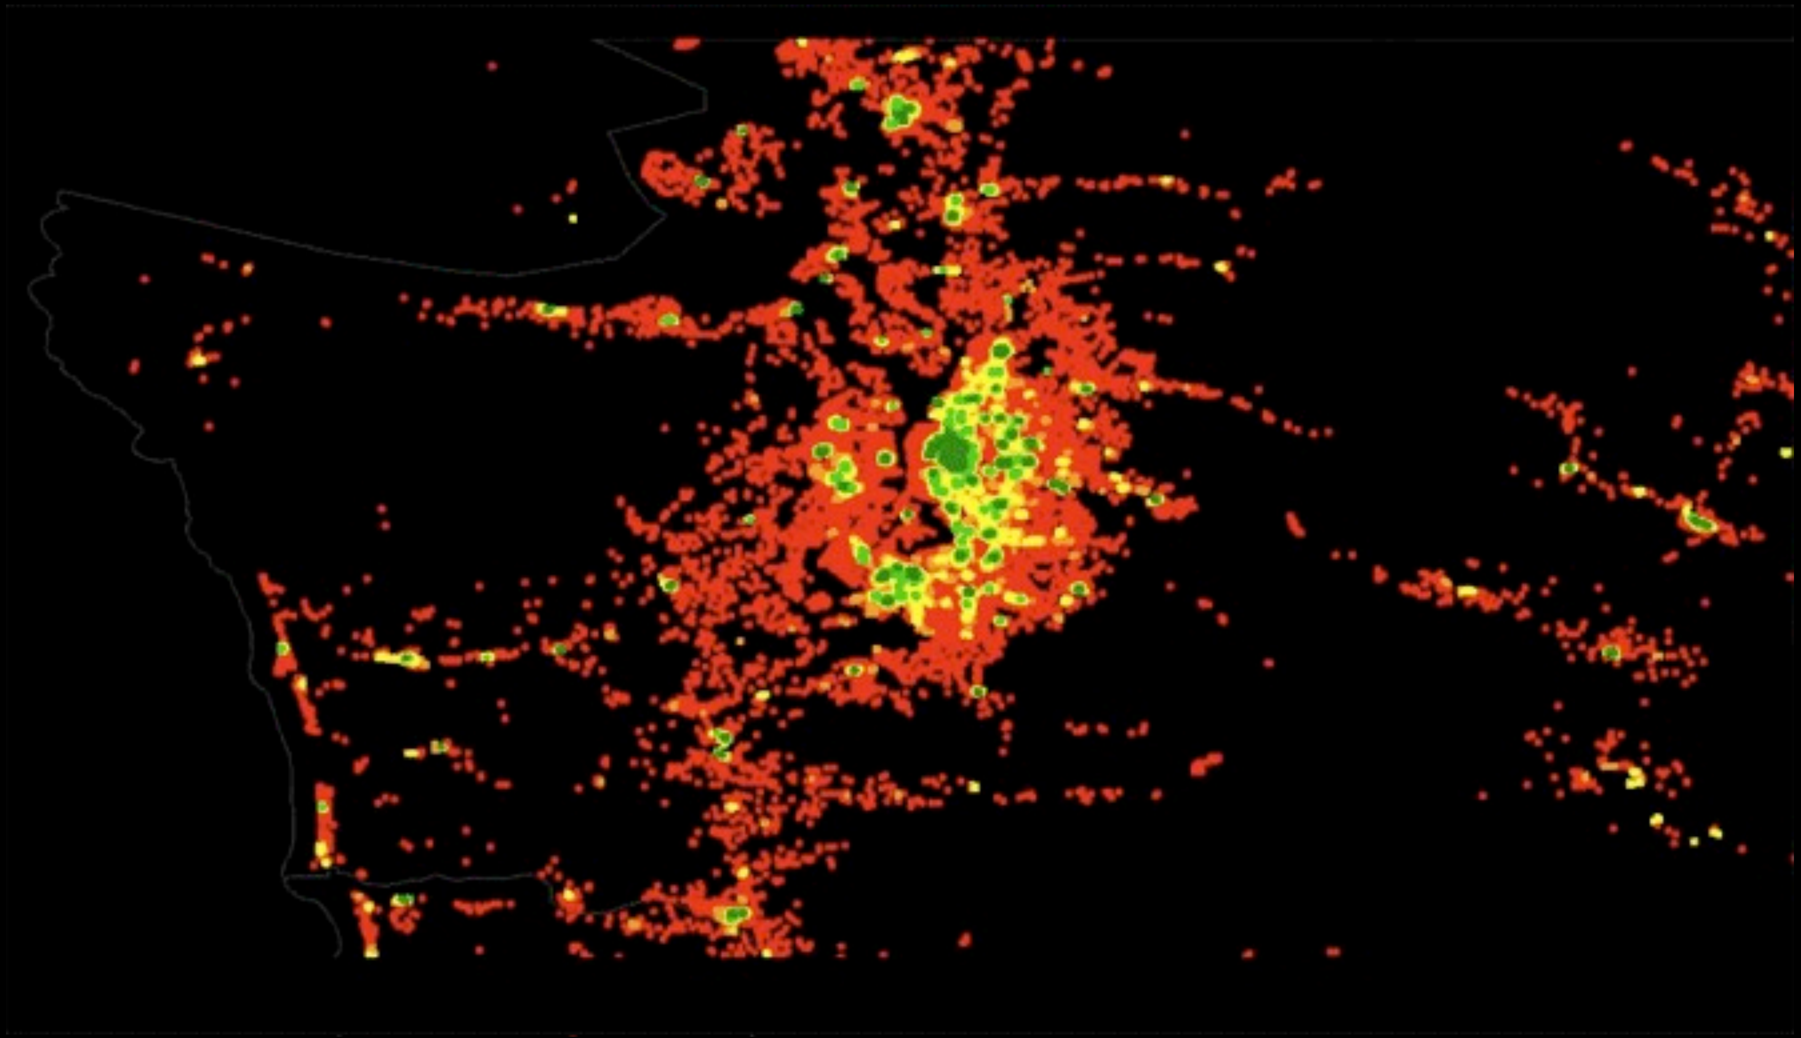

PreCalculation

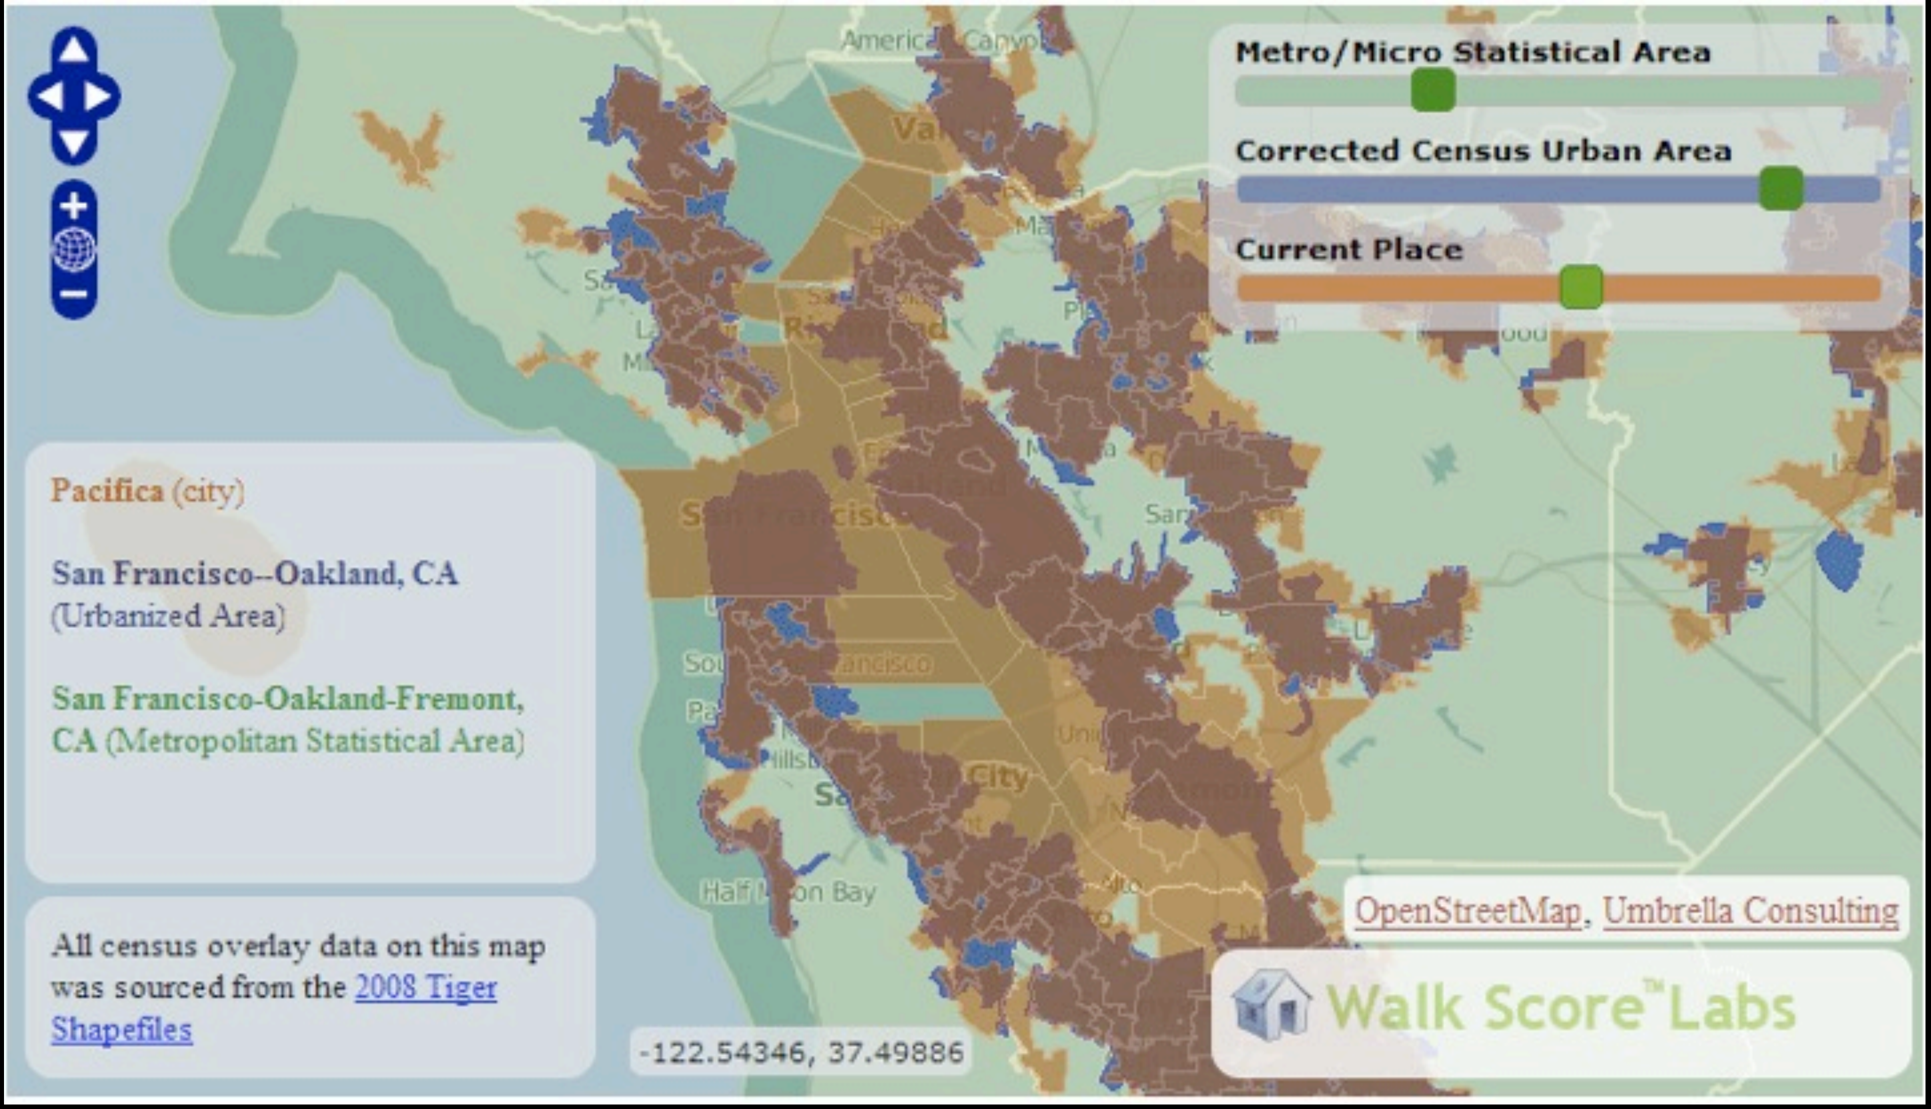

Pacifica (city)

San Francisco--Oakland, CA (Urbanized Area)

San Francisco-Oakland-Fremont, CA (Metropolitan Statistical Area)

All census overlay data on this map was sourced from the 2008 Tiger Shapefiles

Metro/Micro Statistical Area
Corrected Census Urban Area
Current Place
OpenStreetMap, Umbrella Consulting
Walk Score™ Labs
-122.54346, 37.49886
American Canyon
Half Moon Bay

API workflow

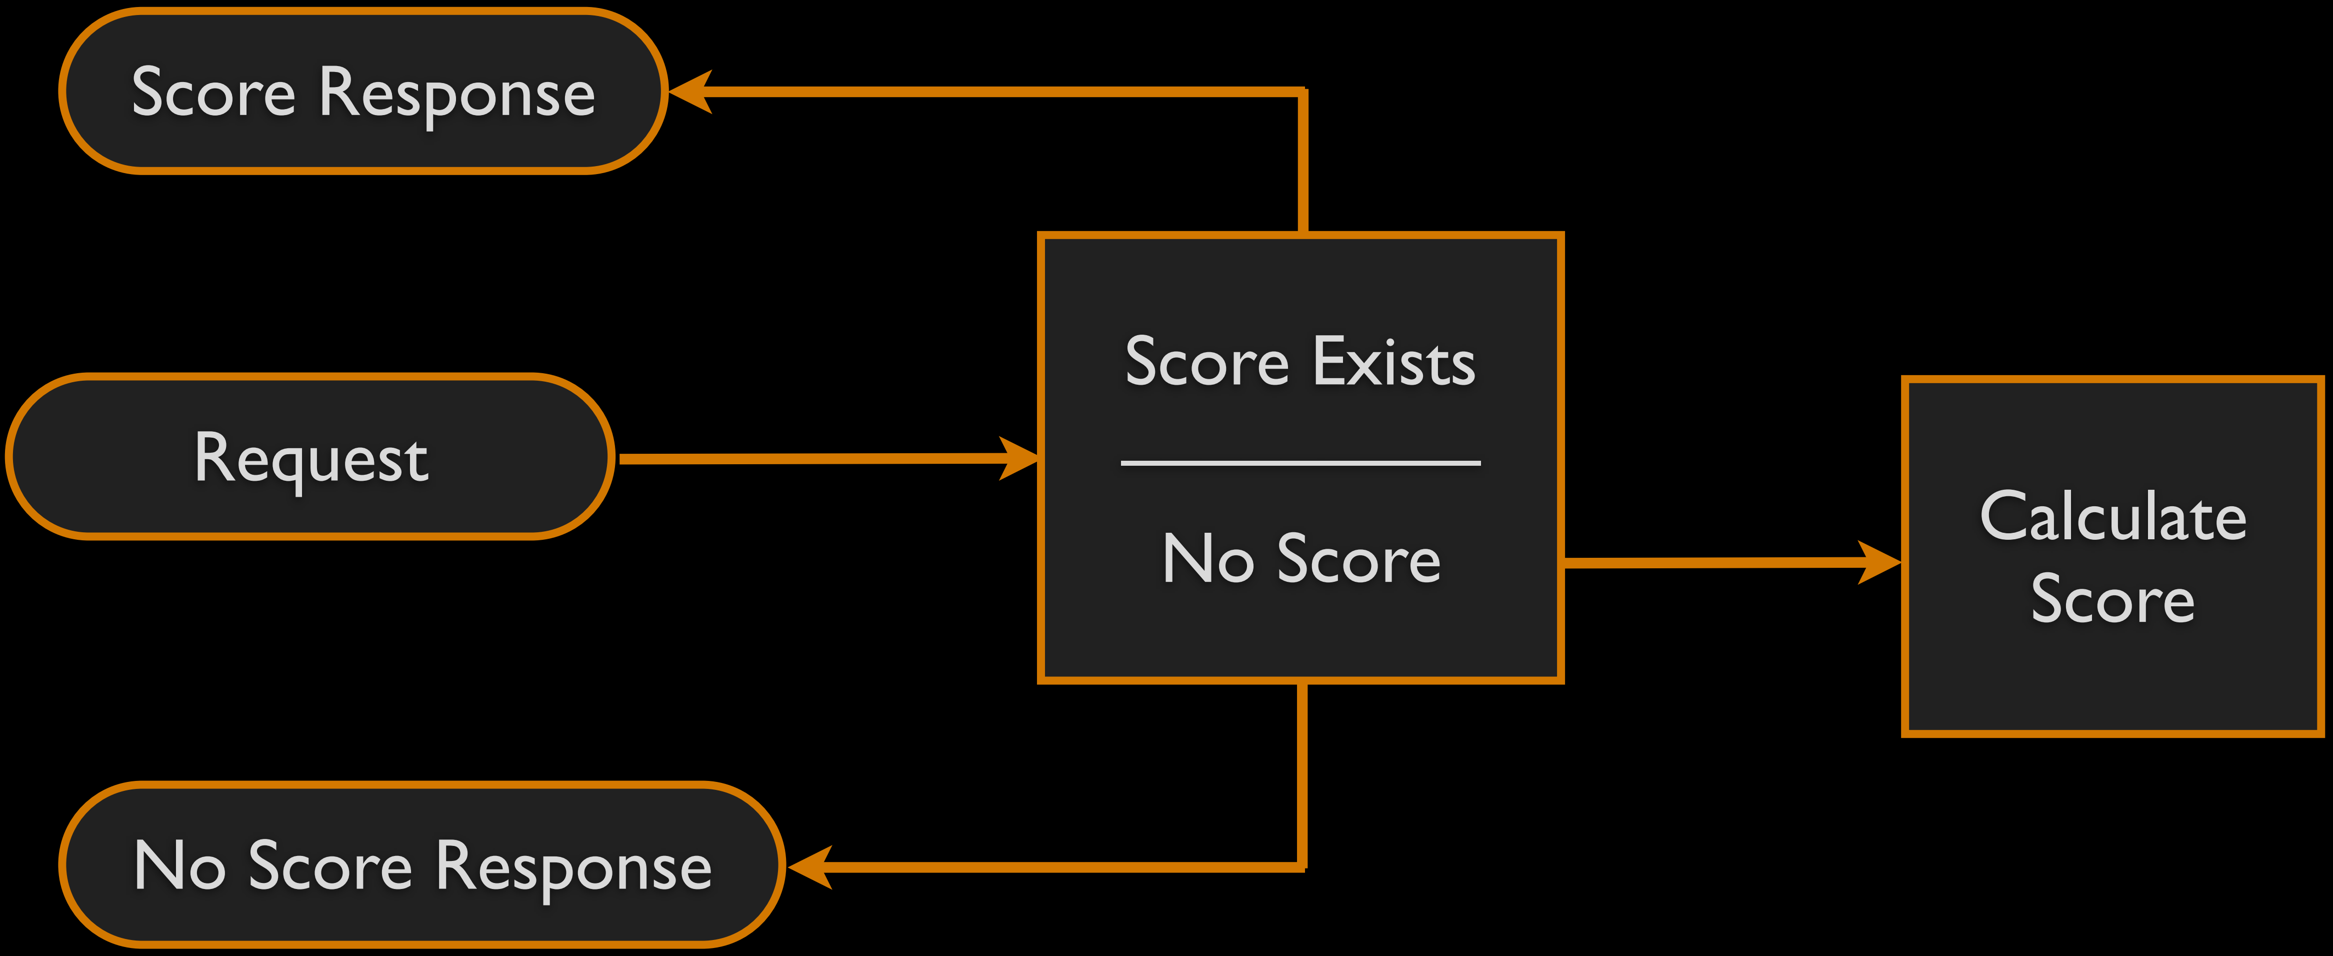

Walk Score Calculation Takes Time!

Give us latitude + longitude

Get a score, if available

Otherwise, queue it up!

!cron jobs

background tasks?

EC2.

Background processing

Background processing

Arbitrary # of URL fetches

Background processing

Arbitrary # of URL fetches

Multi-process execution

App Engine + EC2

Customer

AppEngine

Amazon EC2

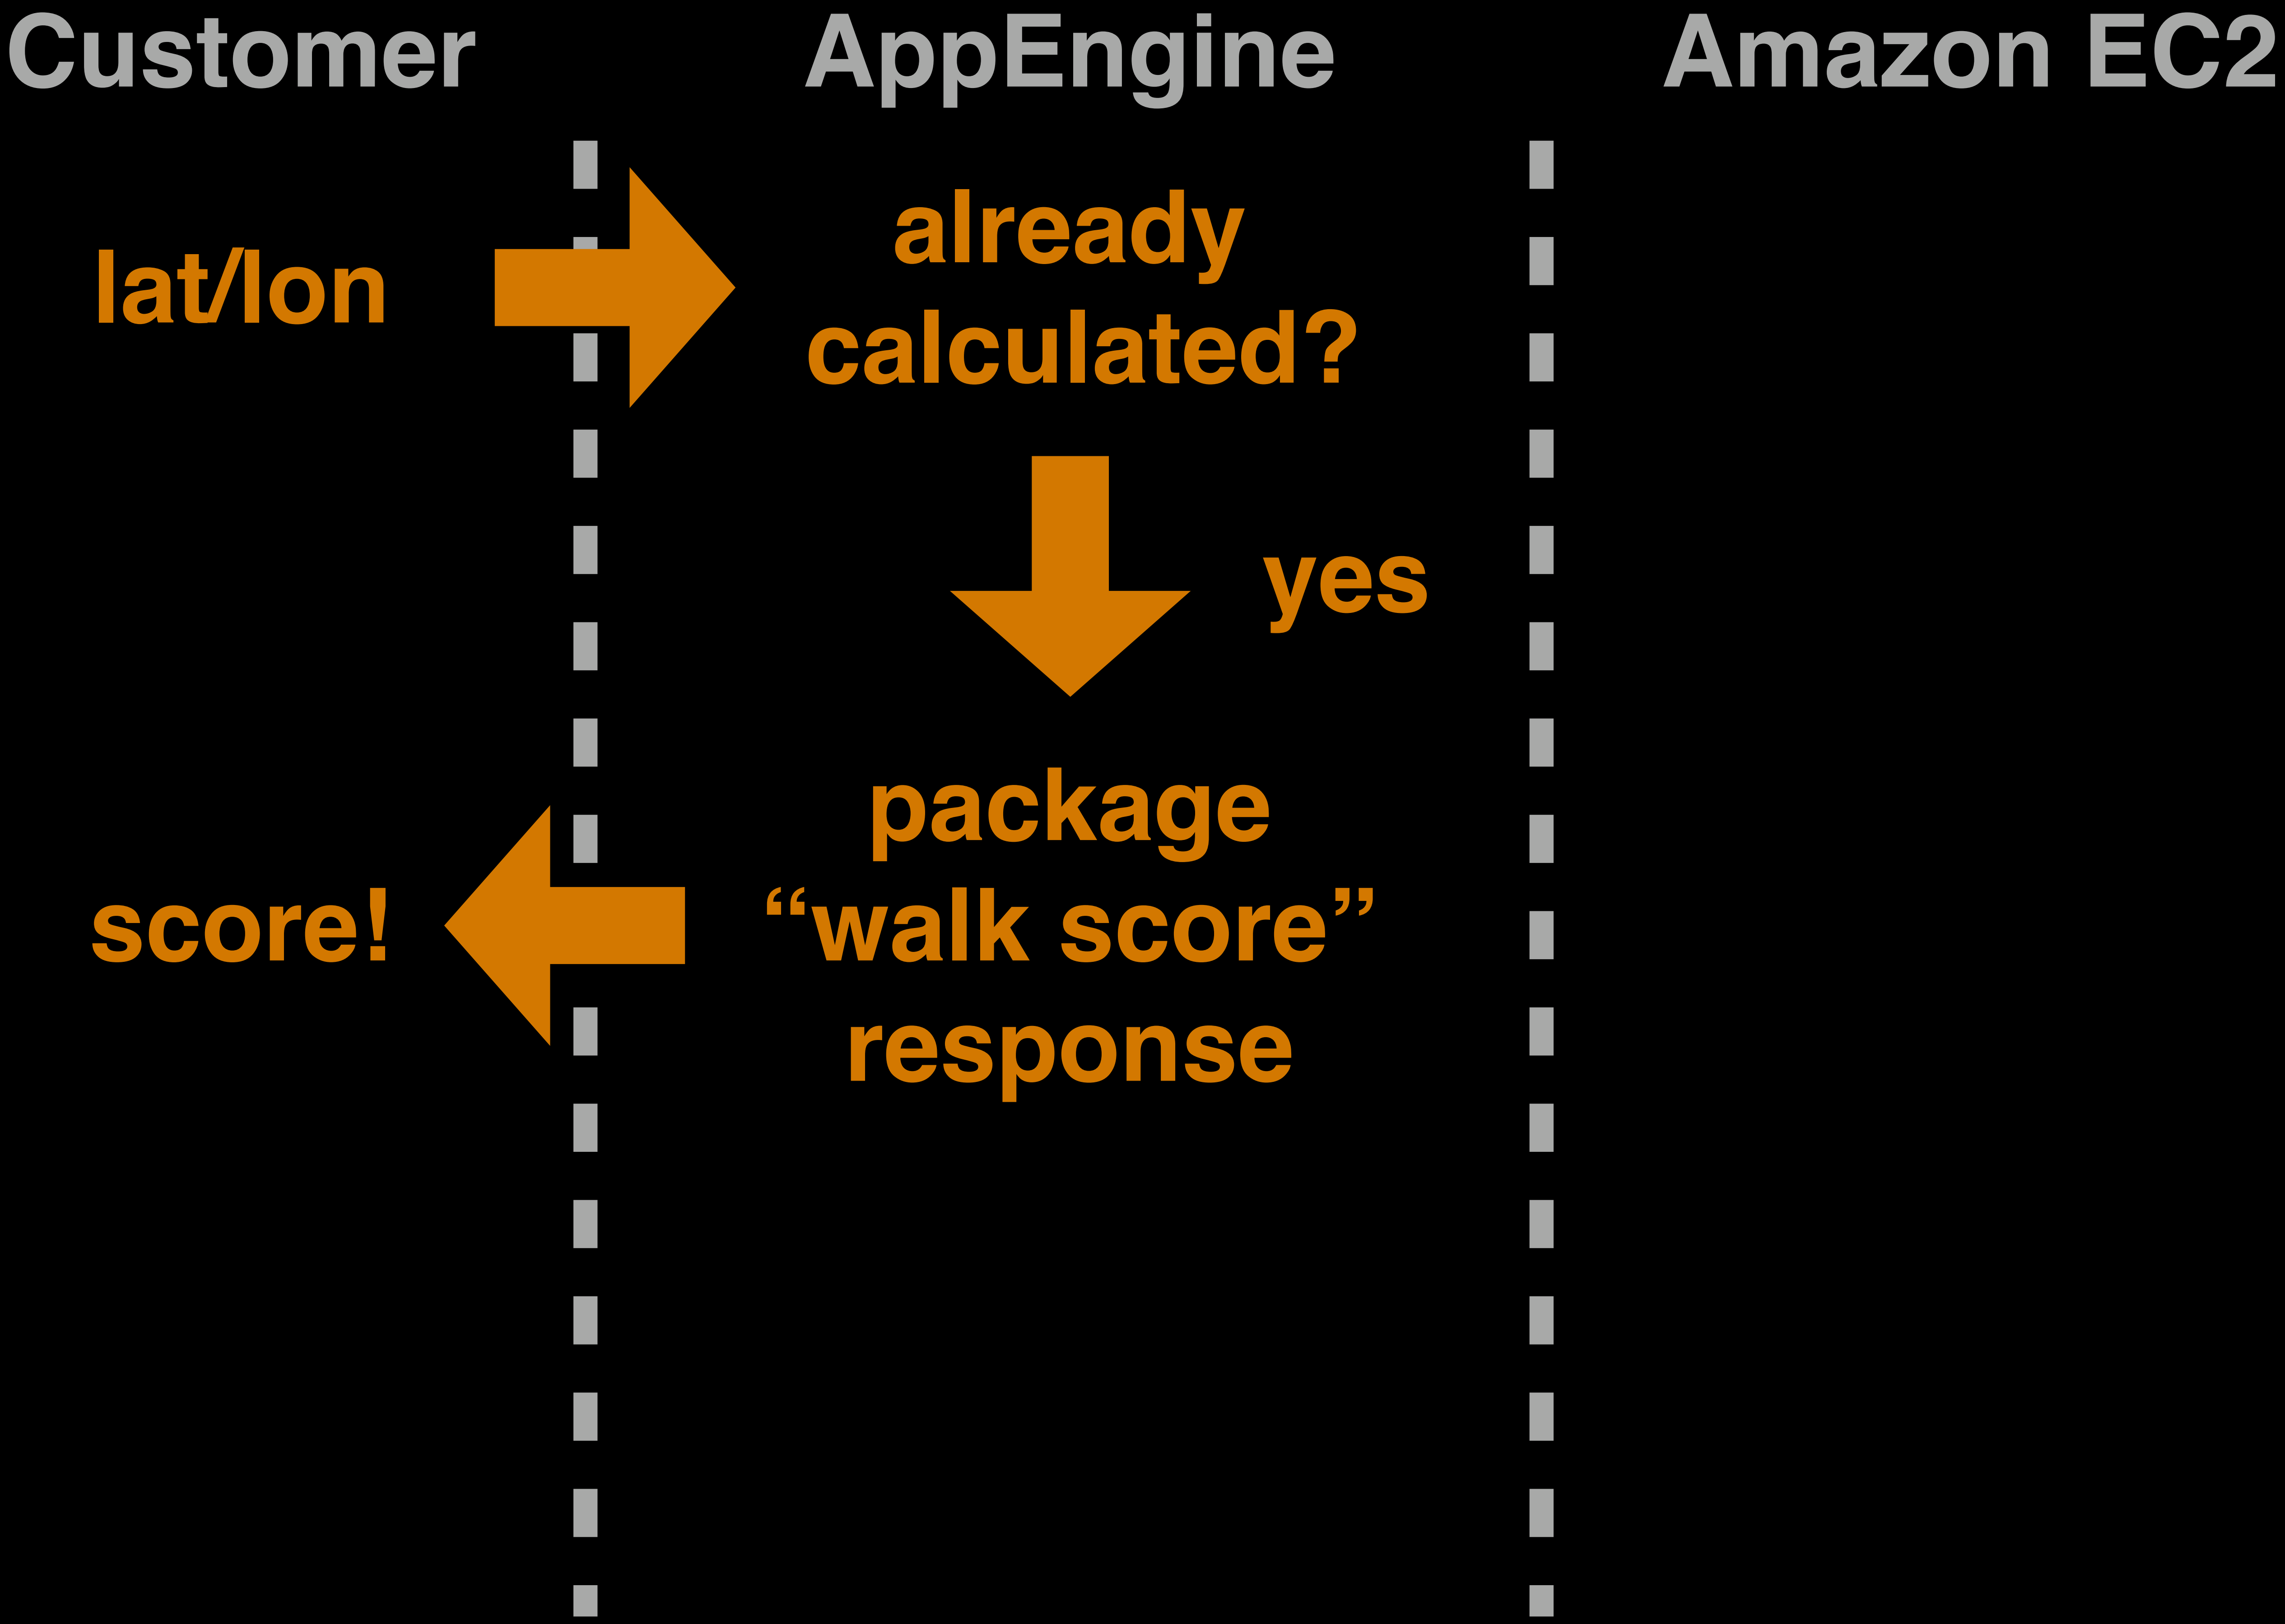

Customer

AppEngine

Amazon EC2

lat/lon

already calculated?

yes

package “walk score” response

score!

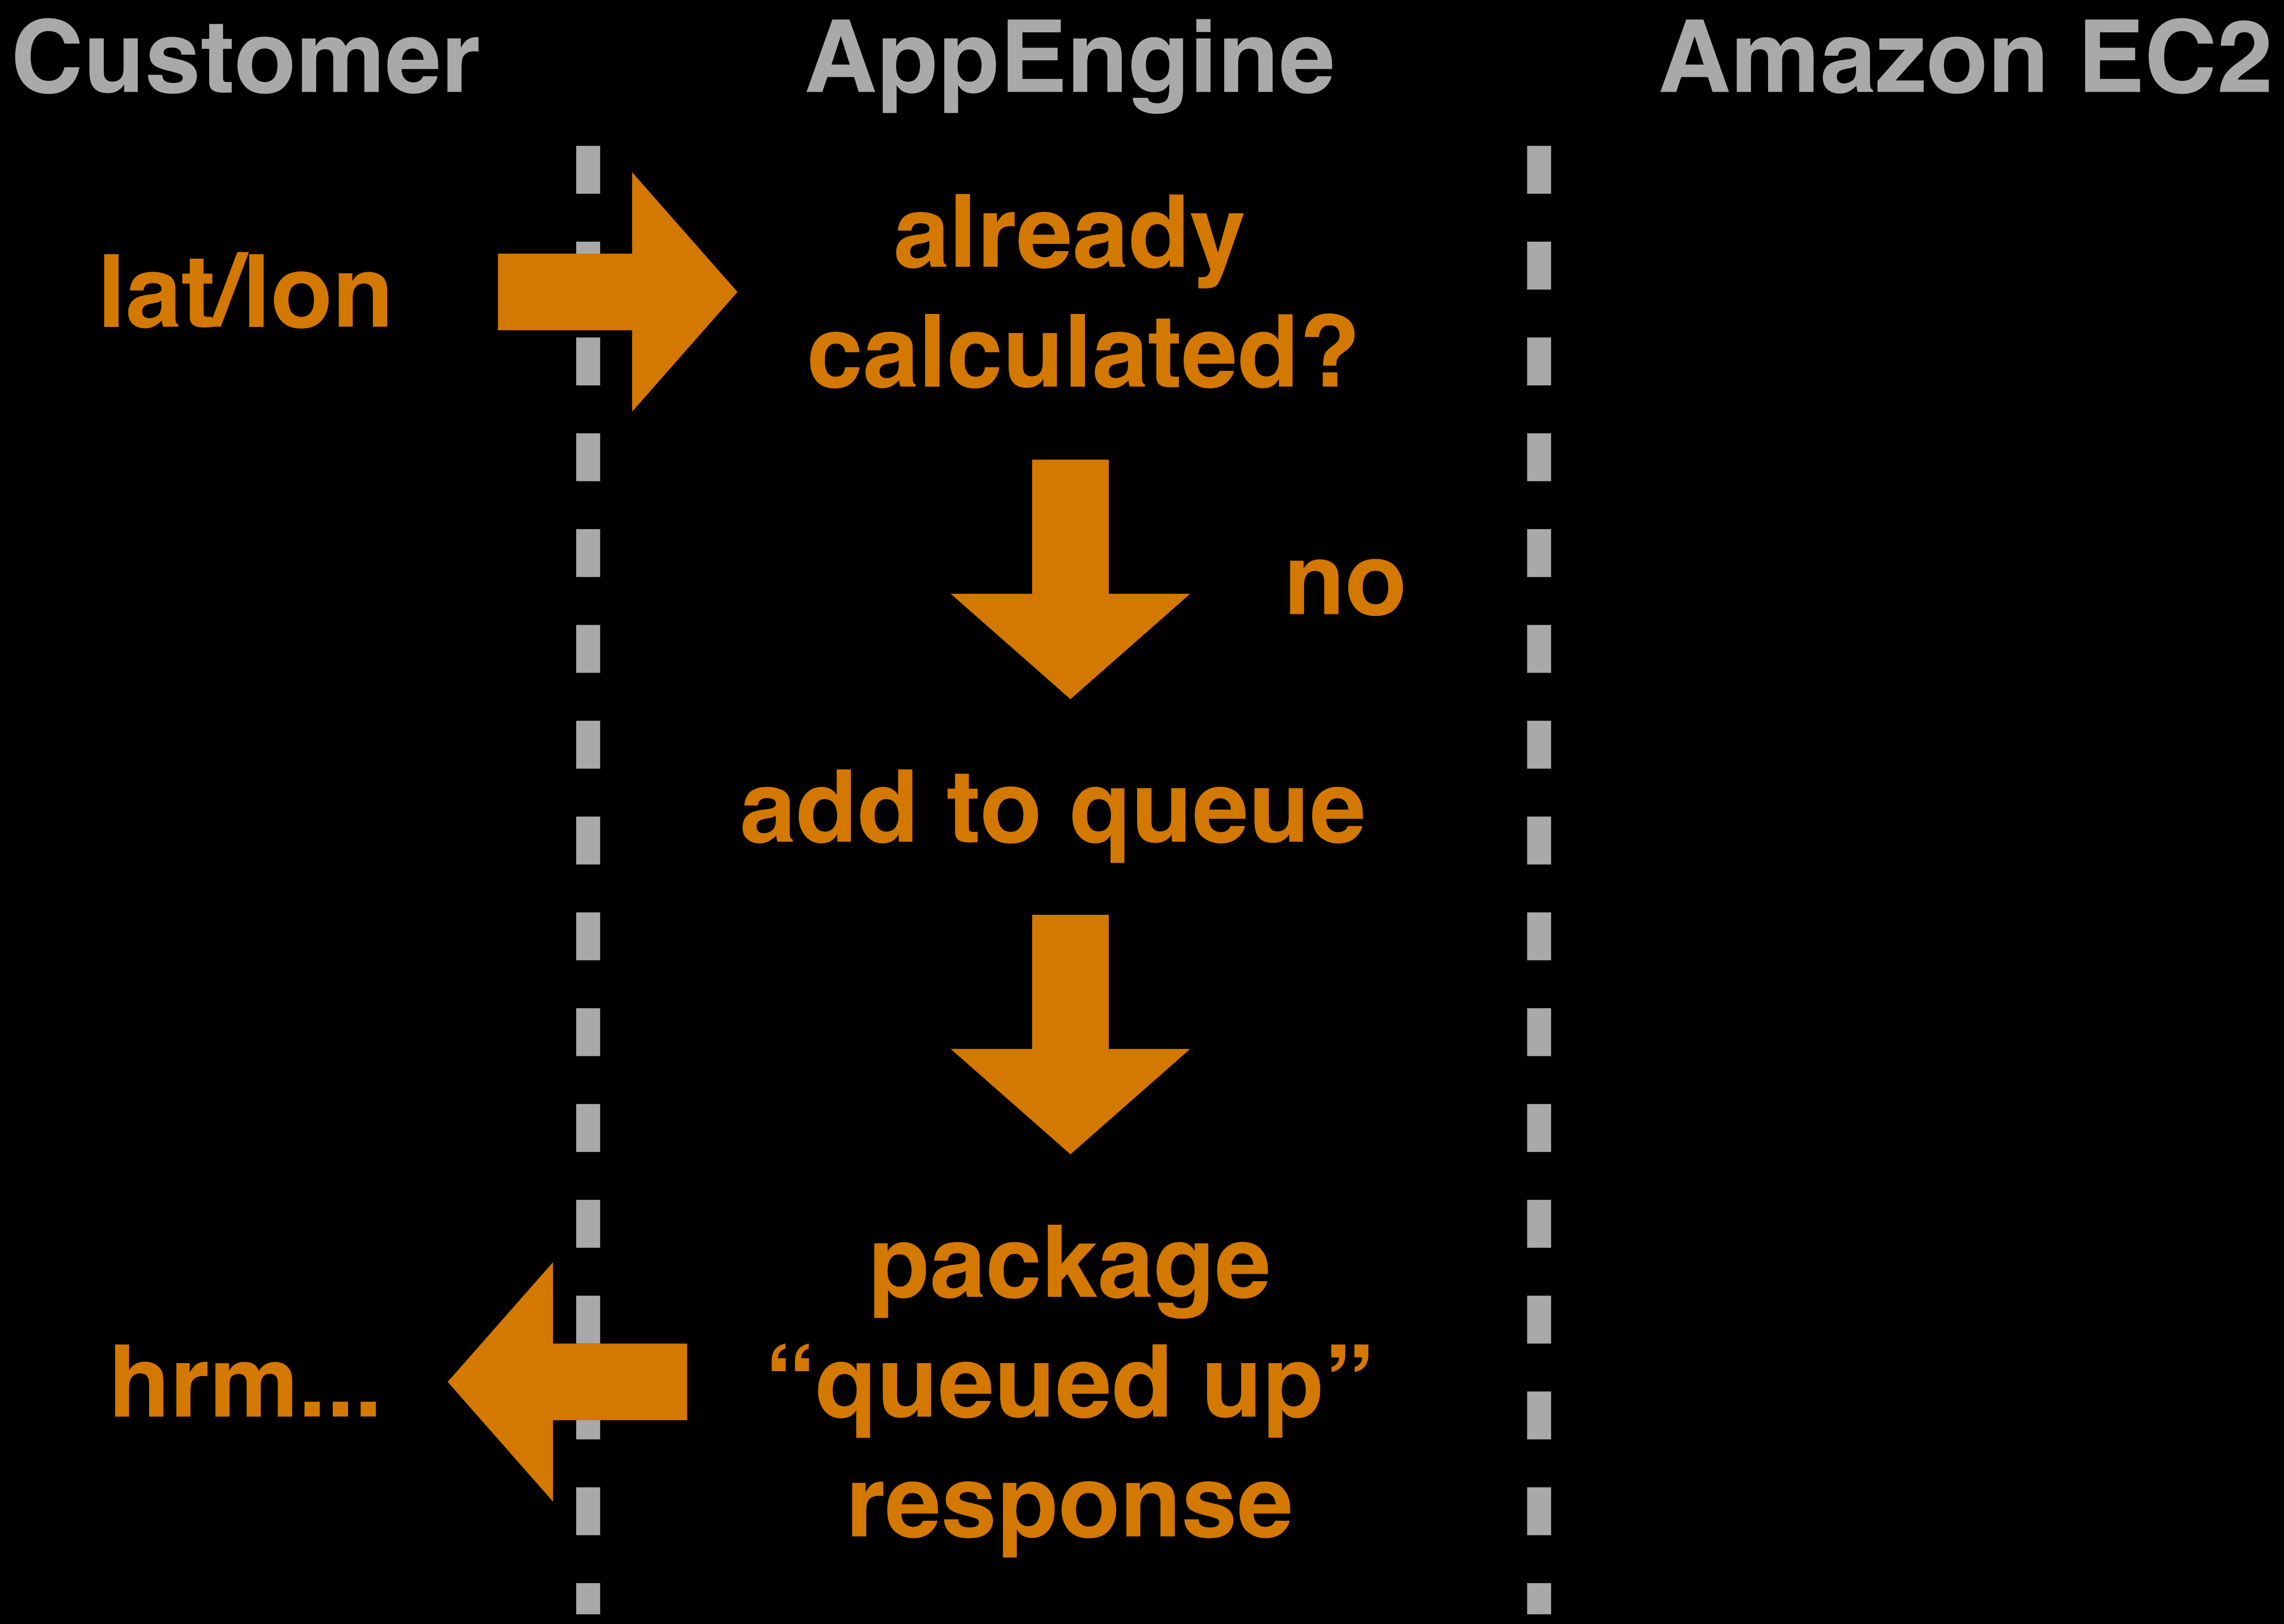

Customer

AppEngine

Amazon EC2

lat/lon

already calculated?

no

add to queue

package “queued up” response

hrm...

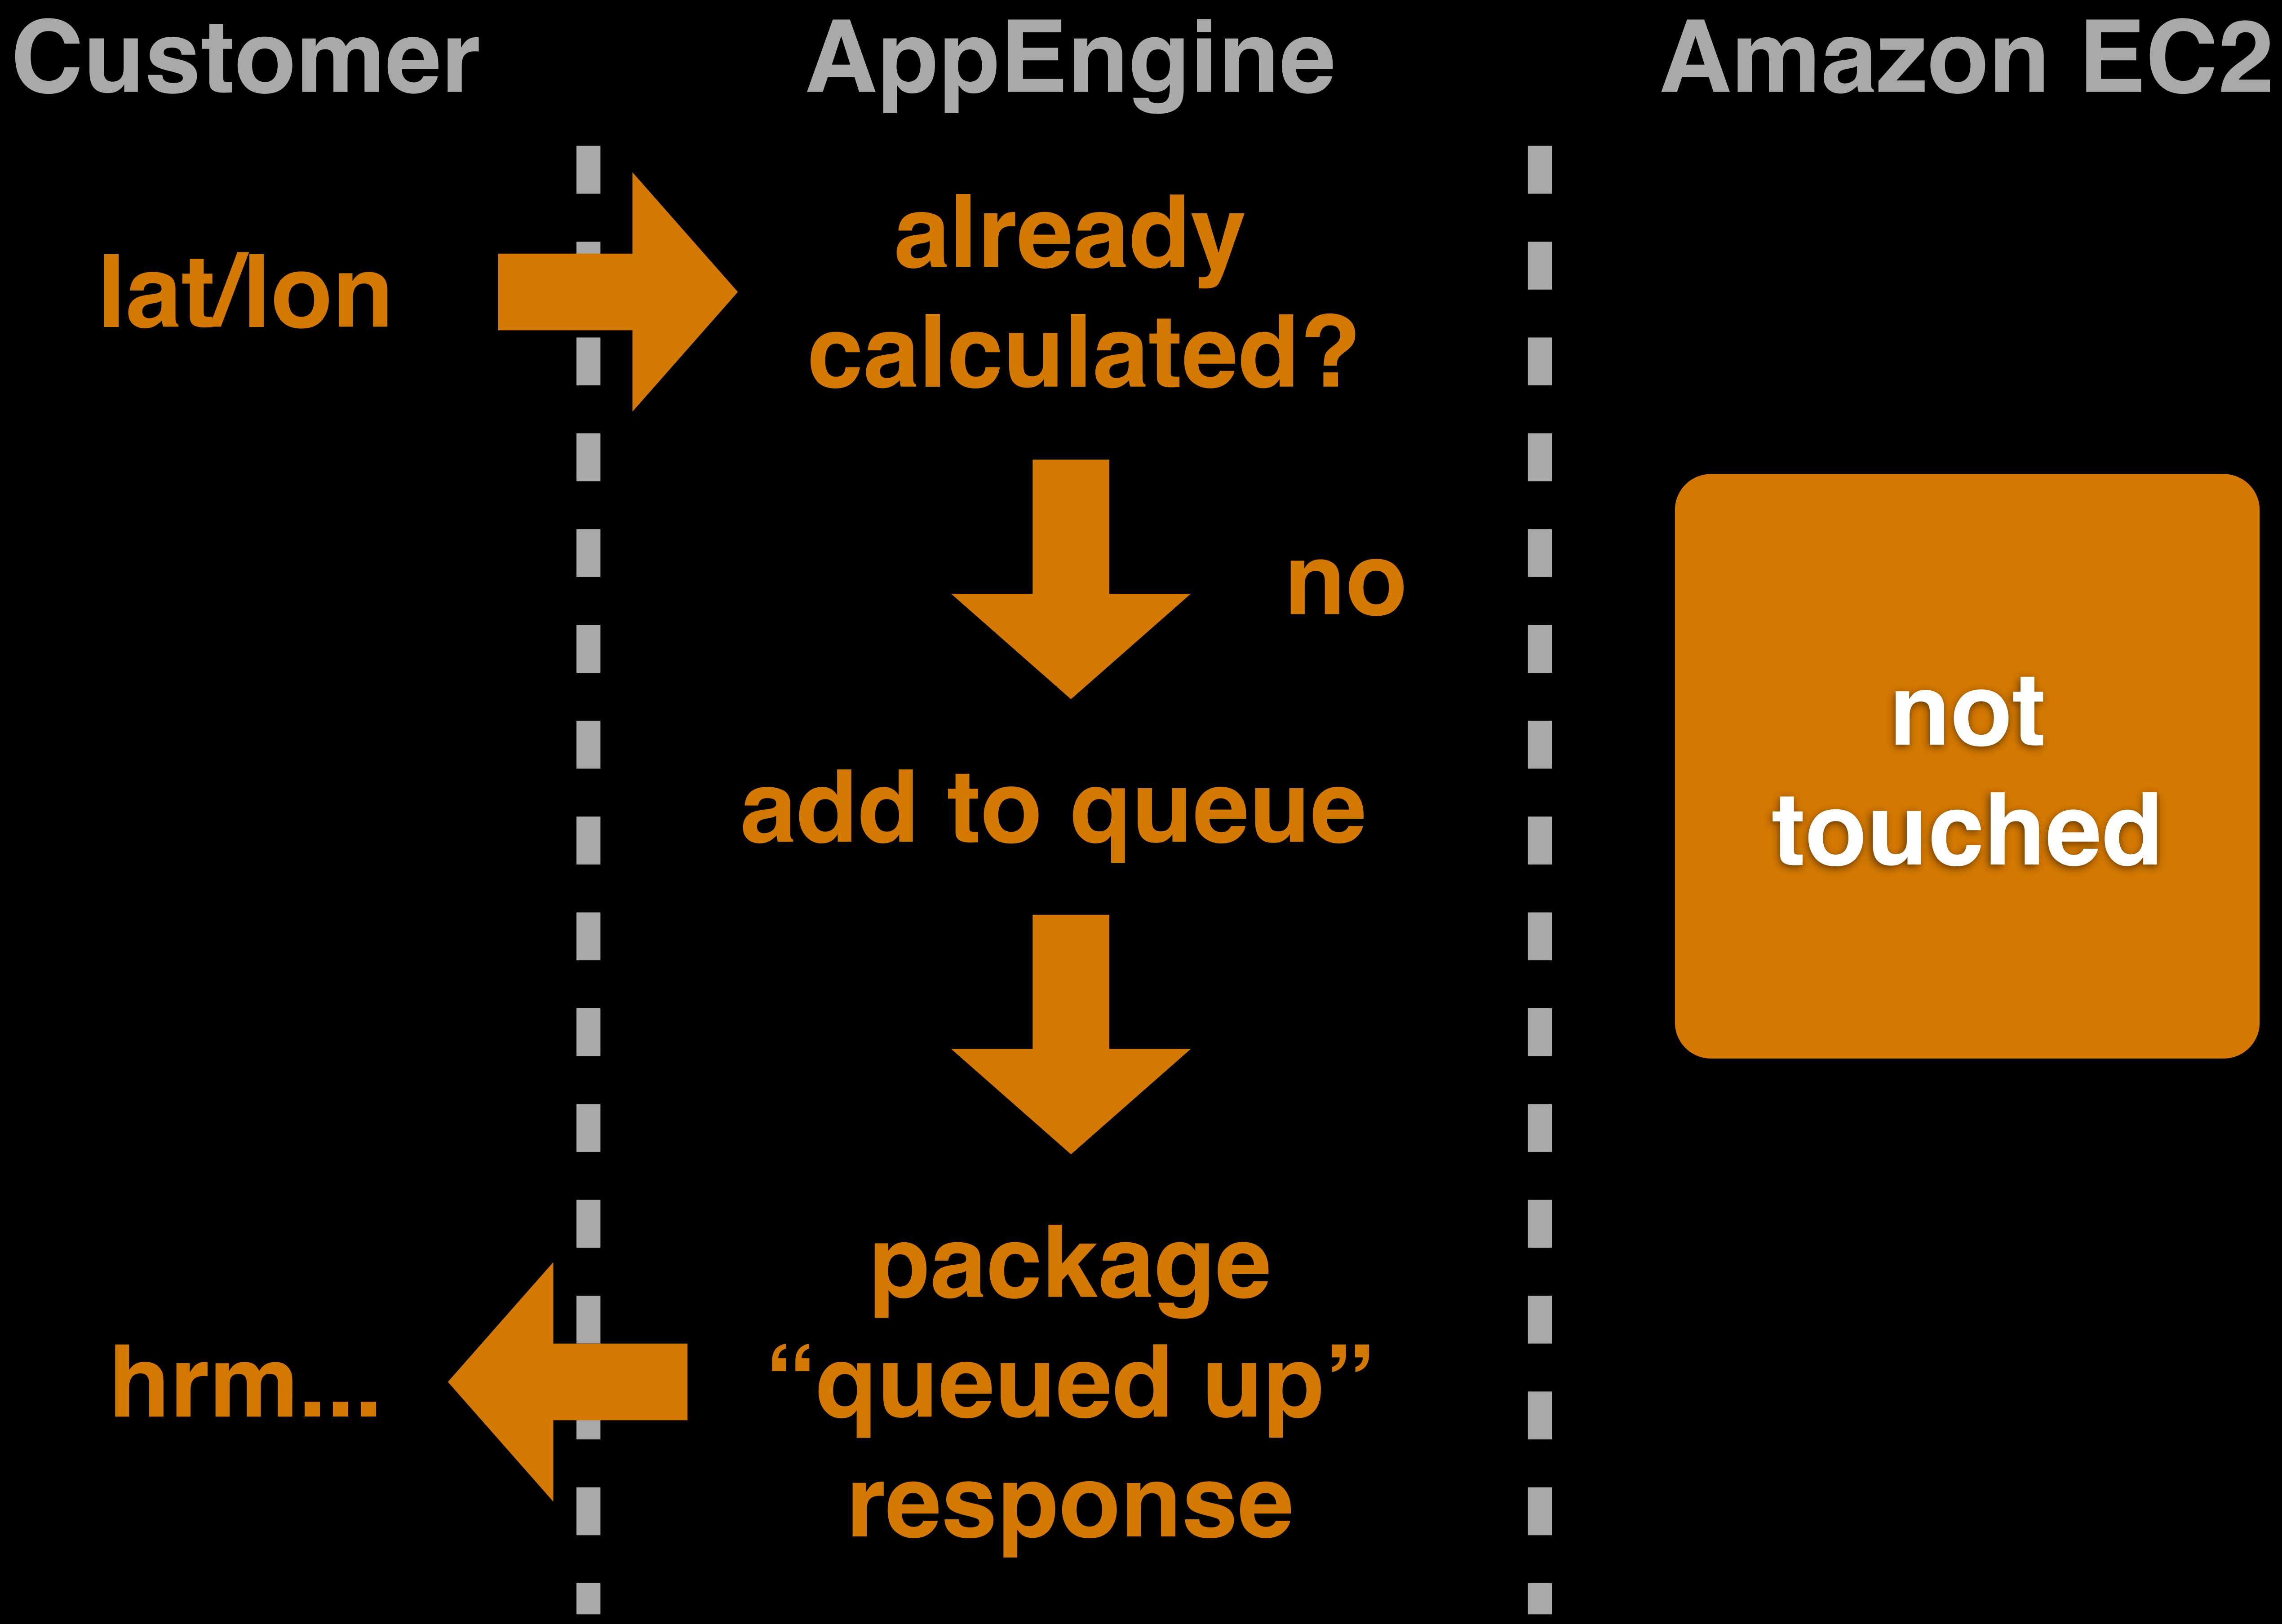

Customer

AppEngine

Amazon EC2

lat/lon

already calculated?

no

add to queue

package “queued up” response

hrm...

not touched

Customer Requests Only Touch AppEngine

Actually running it...

Unpredictable Performance

High Data Store contention

High Data Store contention

Data Store goes offline

High Data Store contention

Data Store goes offline

AppEngine response time increases

High Data Store contention

Data Store goes offline

AppEngine response time increases

AppEngine goes down hard

The Steps To Scalability Hotness

Step 1

Inconsistent model design

Inconsistent model design

Uneven, or no, memcache use

Inconsistent model design

Uneven, or no, memcache use

No data store batching

~5QPS

"Stingingly Scrumptious!"
DAVE'S
HURTIN' HABANERO
HOT SAUCE
TM

Step 2

Memcache on read

Memcache on read

Datastore batching where easy

~10QPS

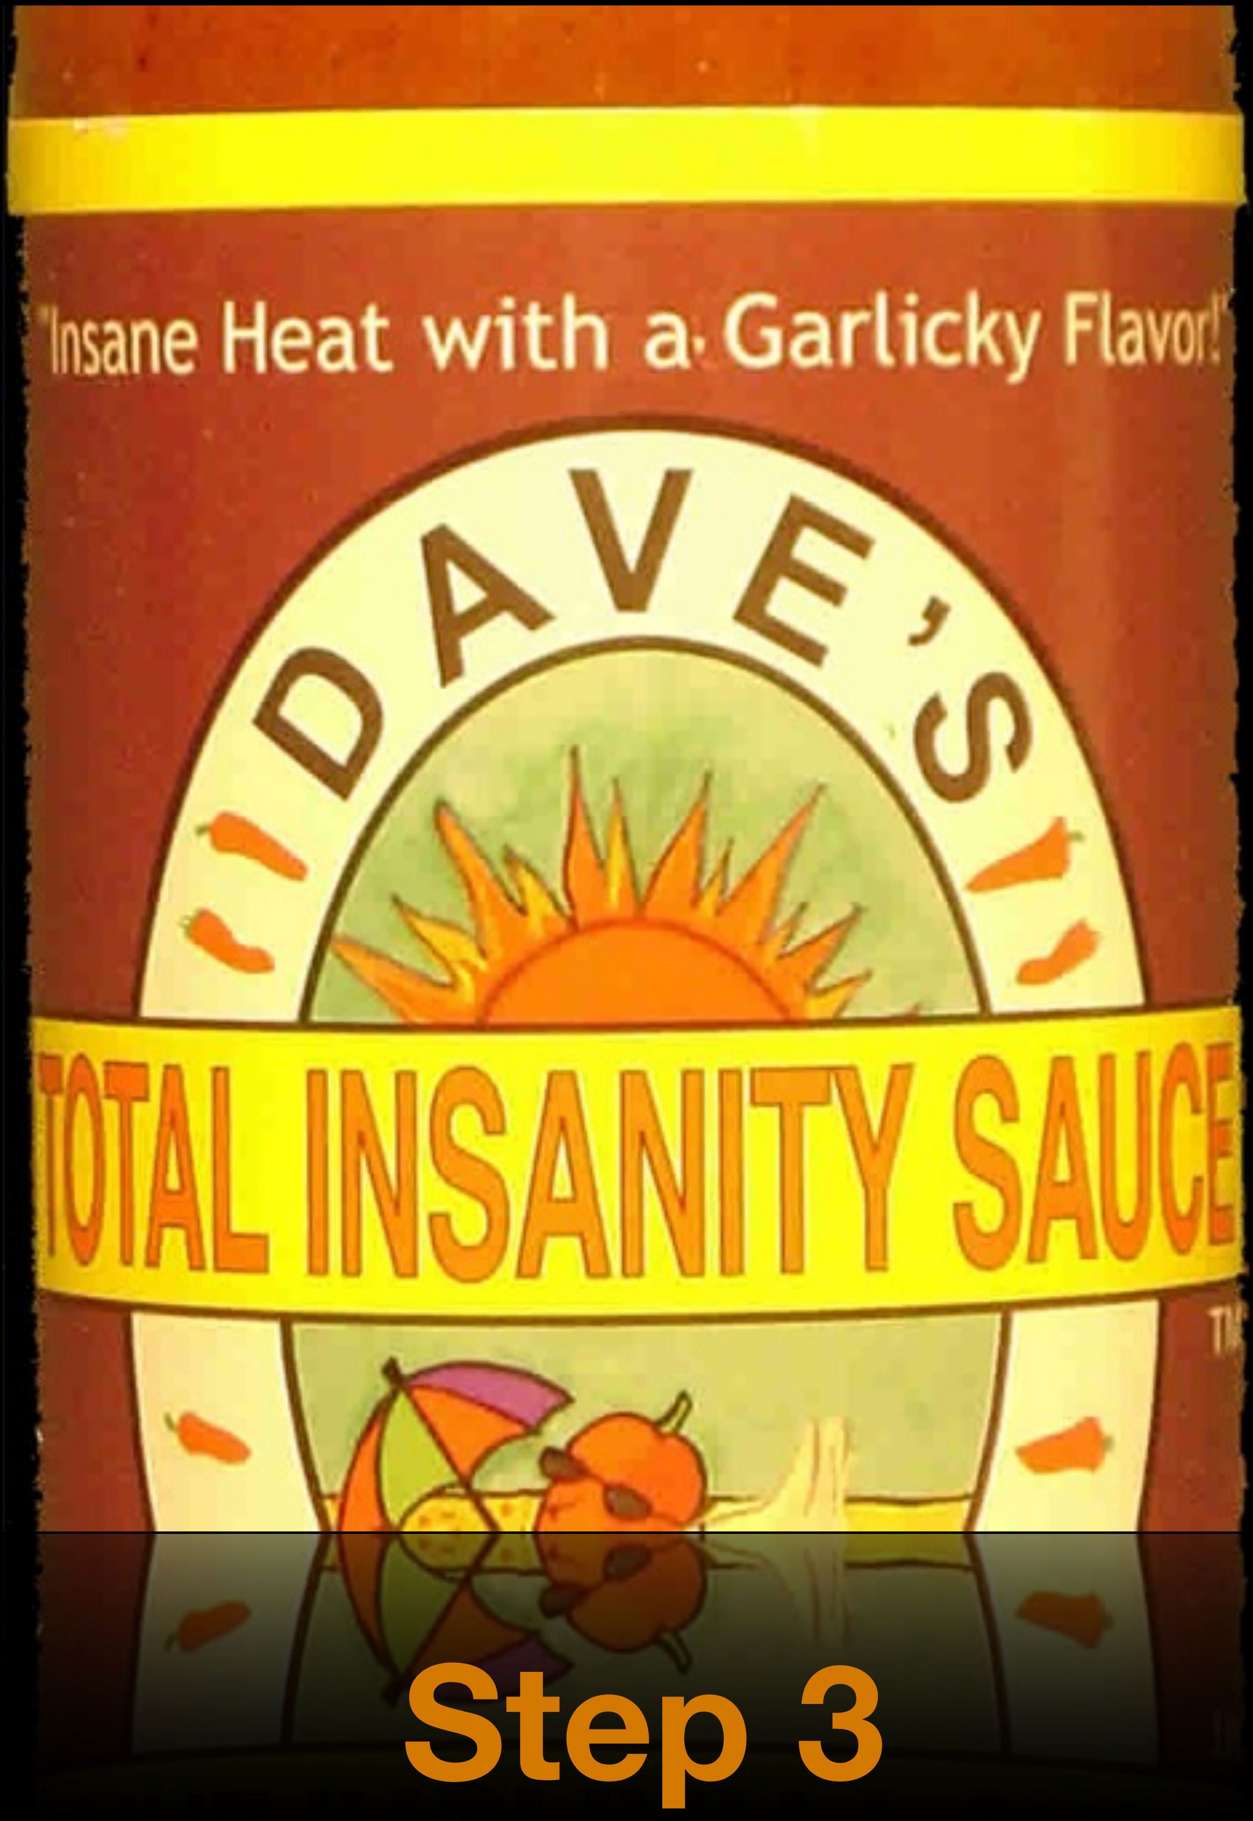
"Insane Heat with a Garlicky Flavor!"
DAVE'S
TOTAL INSANITY SAUCE
TM

Step 3

Memcache usage patterns

Memcache usage patterns

Batch all data store calls

Memcache usage patterns

Batch all data store calls

Decide when not to memcache

Memcache usage patterns

Batch all data store calls

Decide when not to memcache

Decide how long to memcache

Memcache usage patterns

Batch all data store calls

Decide when not to memcache

Decide how long to memcache

Load Test

Memcache usage patterns

Batch all data store calls

Decide when not to memcache

Decide how long to memcache

Load Test

Performance Monitoring

```
try:
    entity.put()
except TimeOut:
    logging.debug()
    entity.put()
```

Applying The Scalability Rules

Counting on App Engine

Keeping track of usage / quotas:

Sharded Counter

Didn't scale...

Counting on App Engine

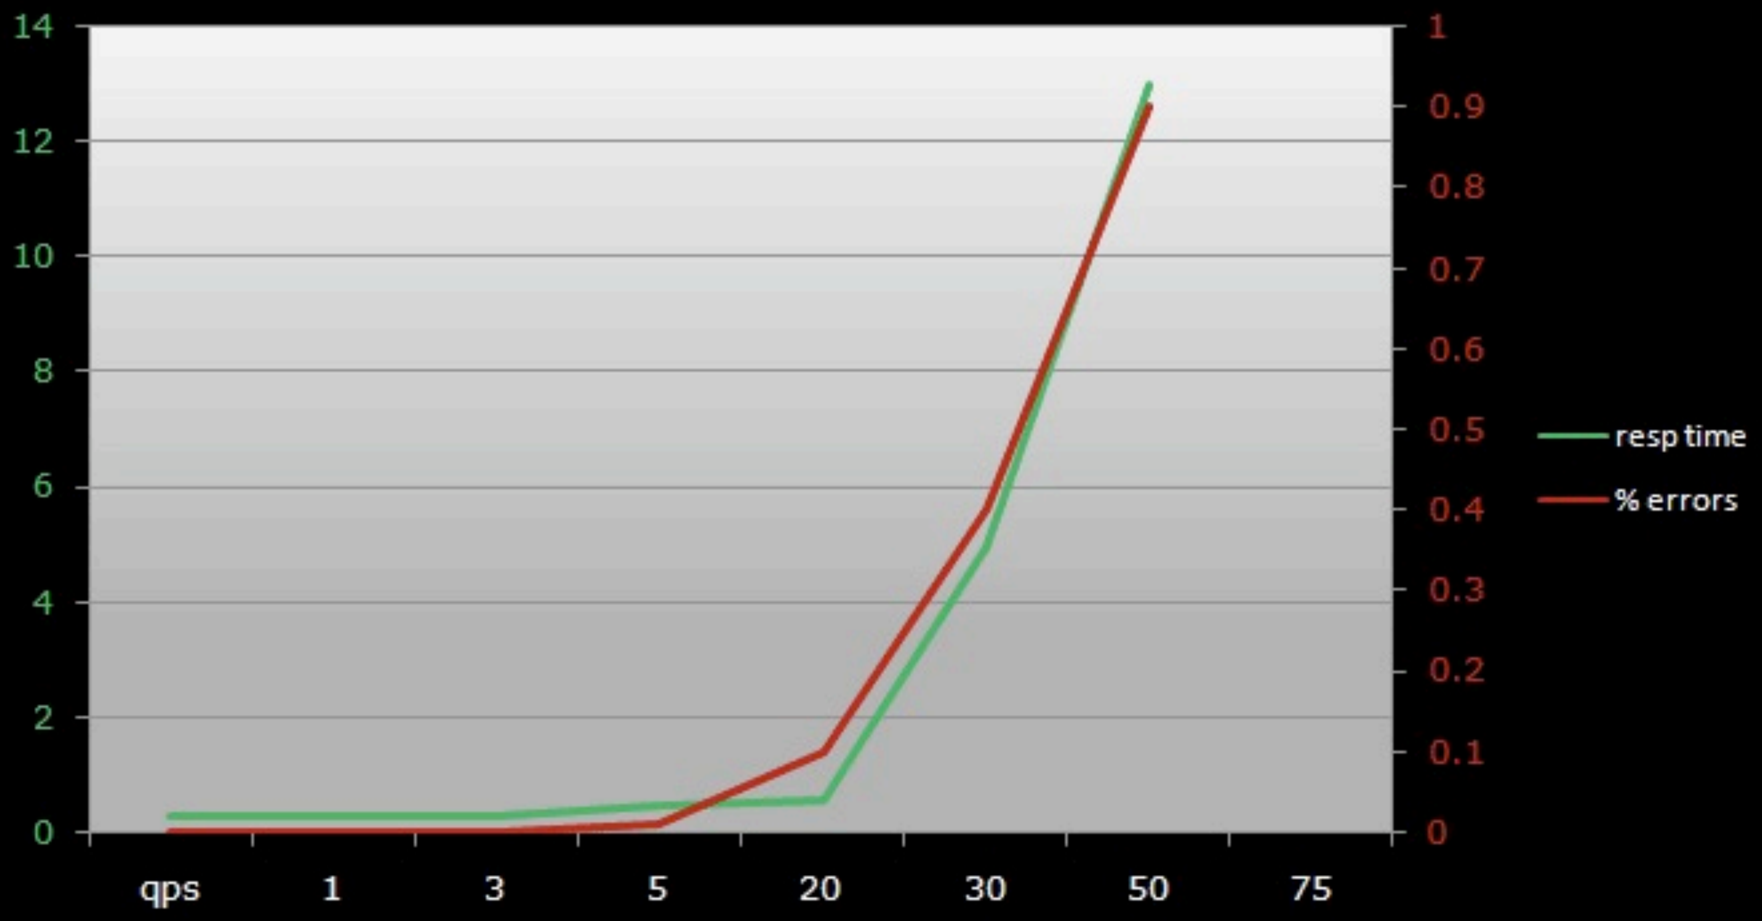

Counting on App Engine

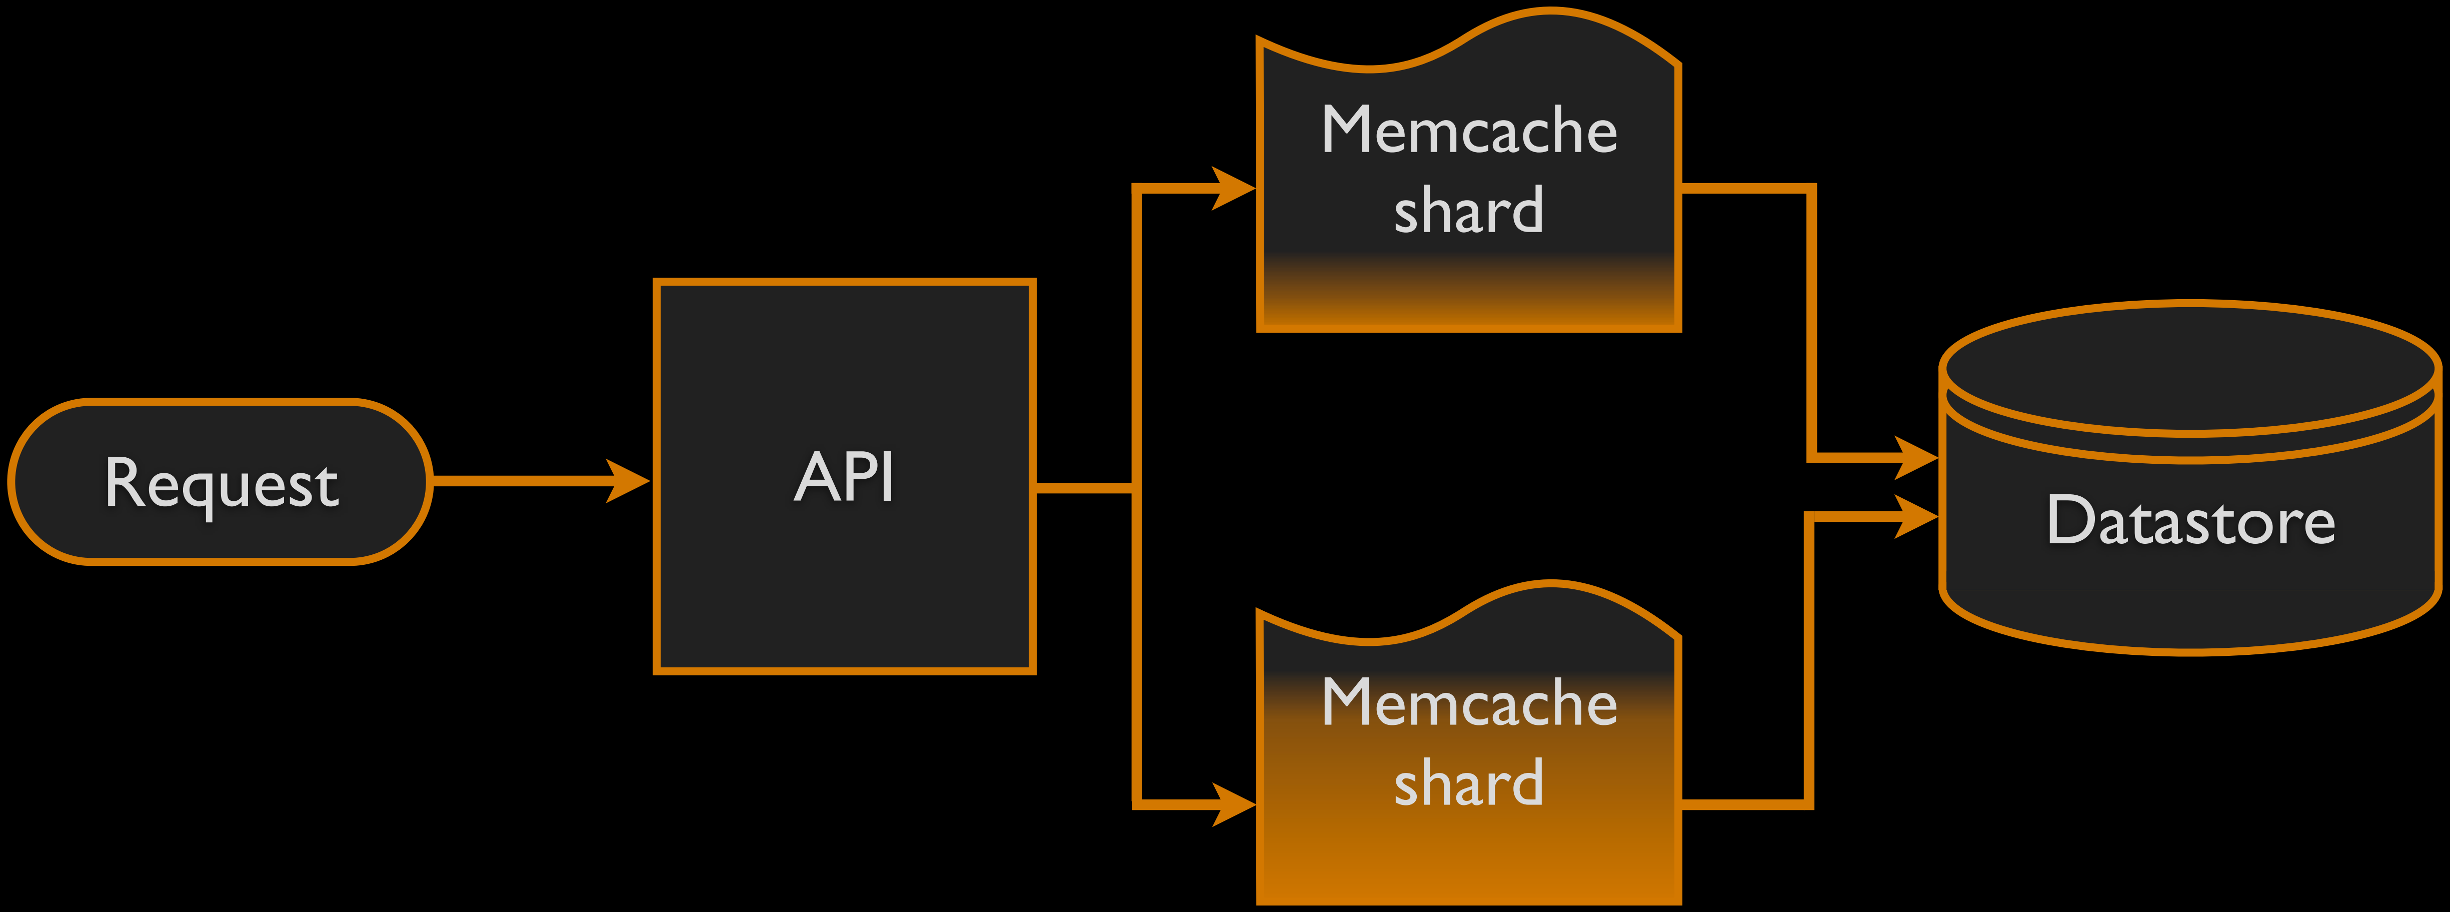

Counting on App Engine

Solution

Memcached (sharded) counter

Advantage

Negligable datastore contention / timeouts

Counting on App Engine

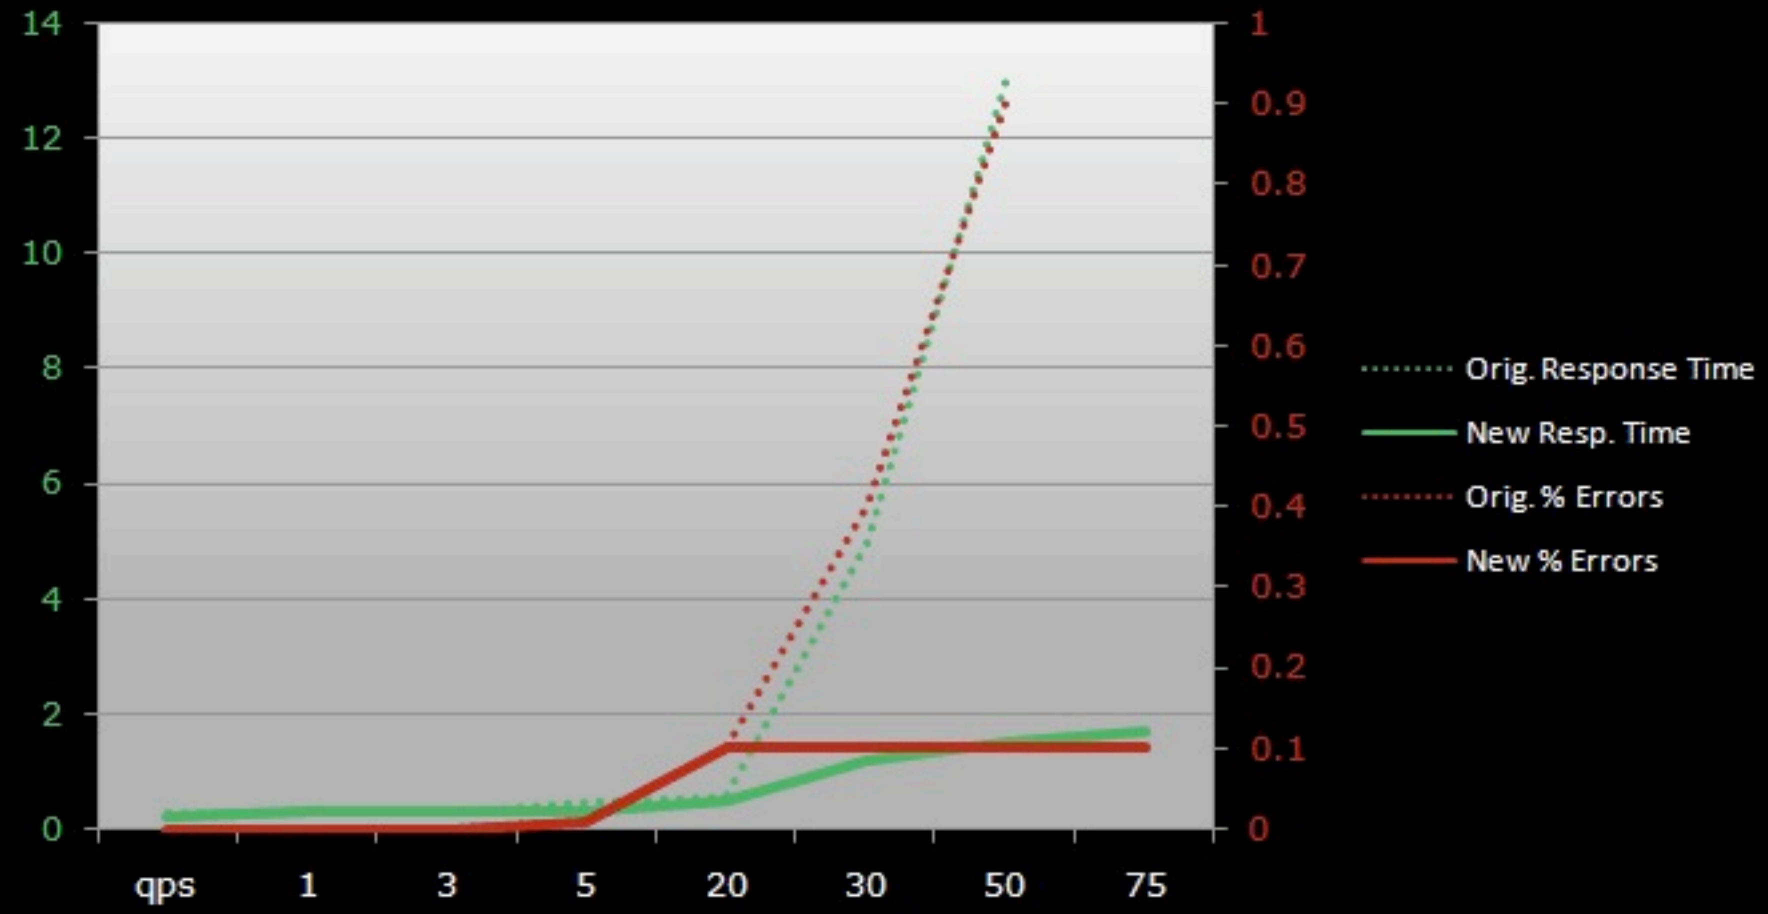

Summary

Use Memcache on write, expecting lossiness

Decide which features fit AppEngine and which don't

Don't let secondary features interfere with core functionality

Results:

1.5 M requests/day

Peaks up to 80 requests/second

A black box

- Well documented
- Community support:
  - irc
  - Google Groups
  - Issue tracker
  - App Engine team

A very good black box

WalkScore.com

A project of FrontSeat.org

jesse@frontseat.org

davepeck@codeorange.com

josh@umbrellaconsulting.com

Q&A

Q&A

Some suggested topics:

Instrumentation

Monitoring

Performance strategies

Calculator details

The cost of indexes

Bug #901

...

feedback? haveasec.com/io

(extra slides if they come up in Q&A)

Instrumentation

Google's Dashboard (insert screenshot)

Google's System Status (insert snapshot here)

Google's Log Viewer (insert snapshot here)

Walk Score's Dashboard (insert snapshot here)

Calculator Details

Performance Strategies

Decouple I/O from computation

Buffer all I/O

Use buffer level marks

Dynamically respond to changing I/O conditions

AppEngine

Amazon EC2

Third Party
Web Services

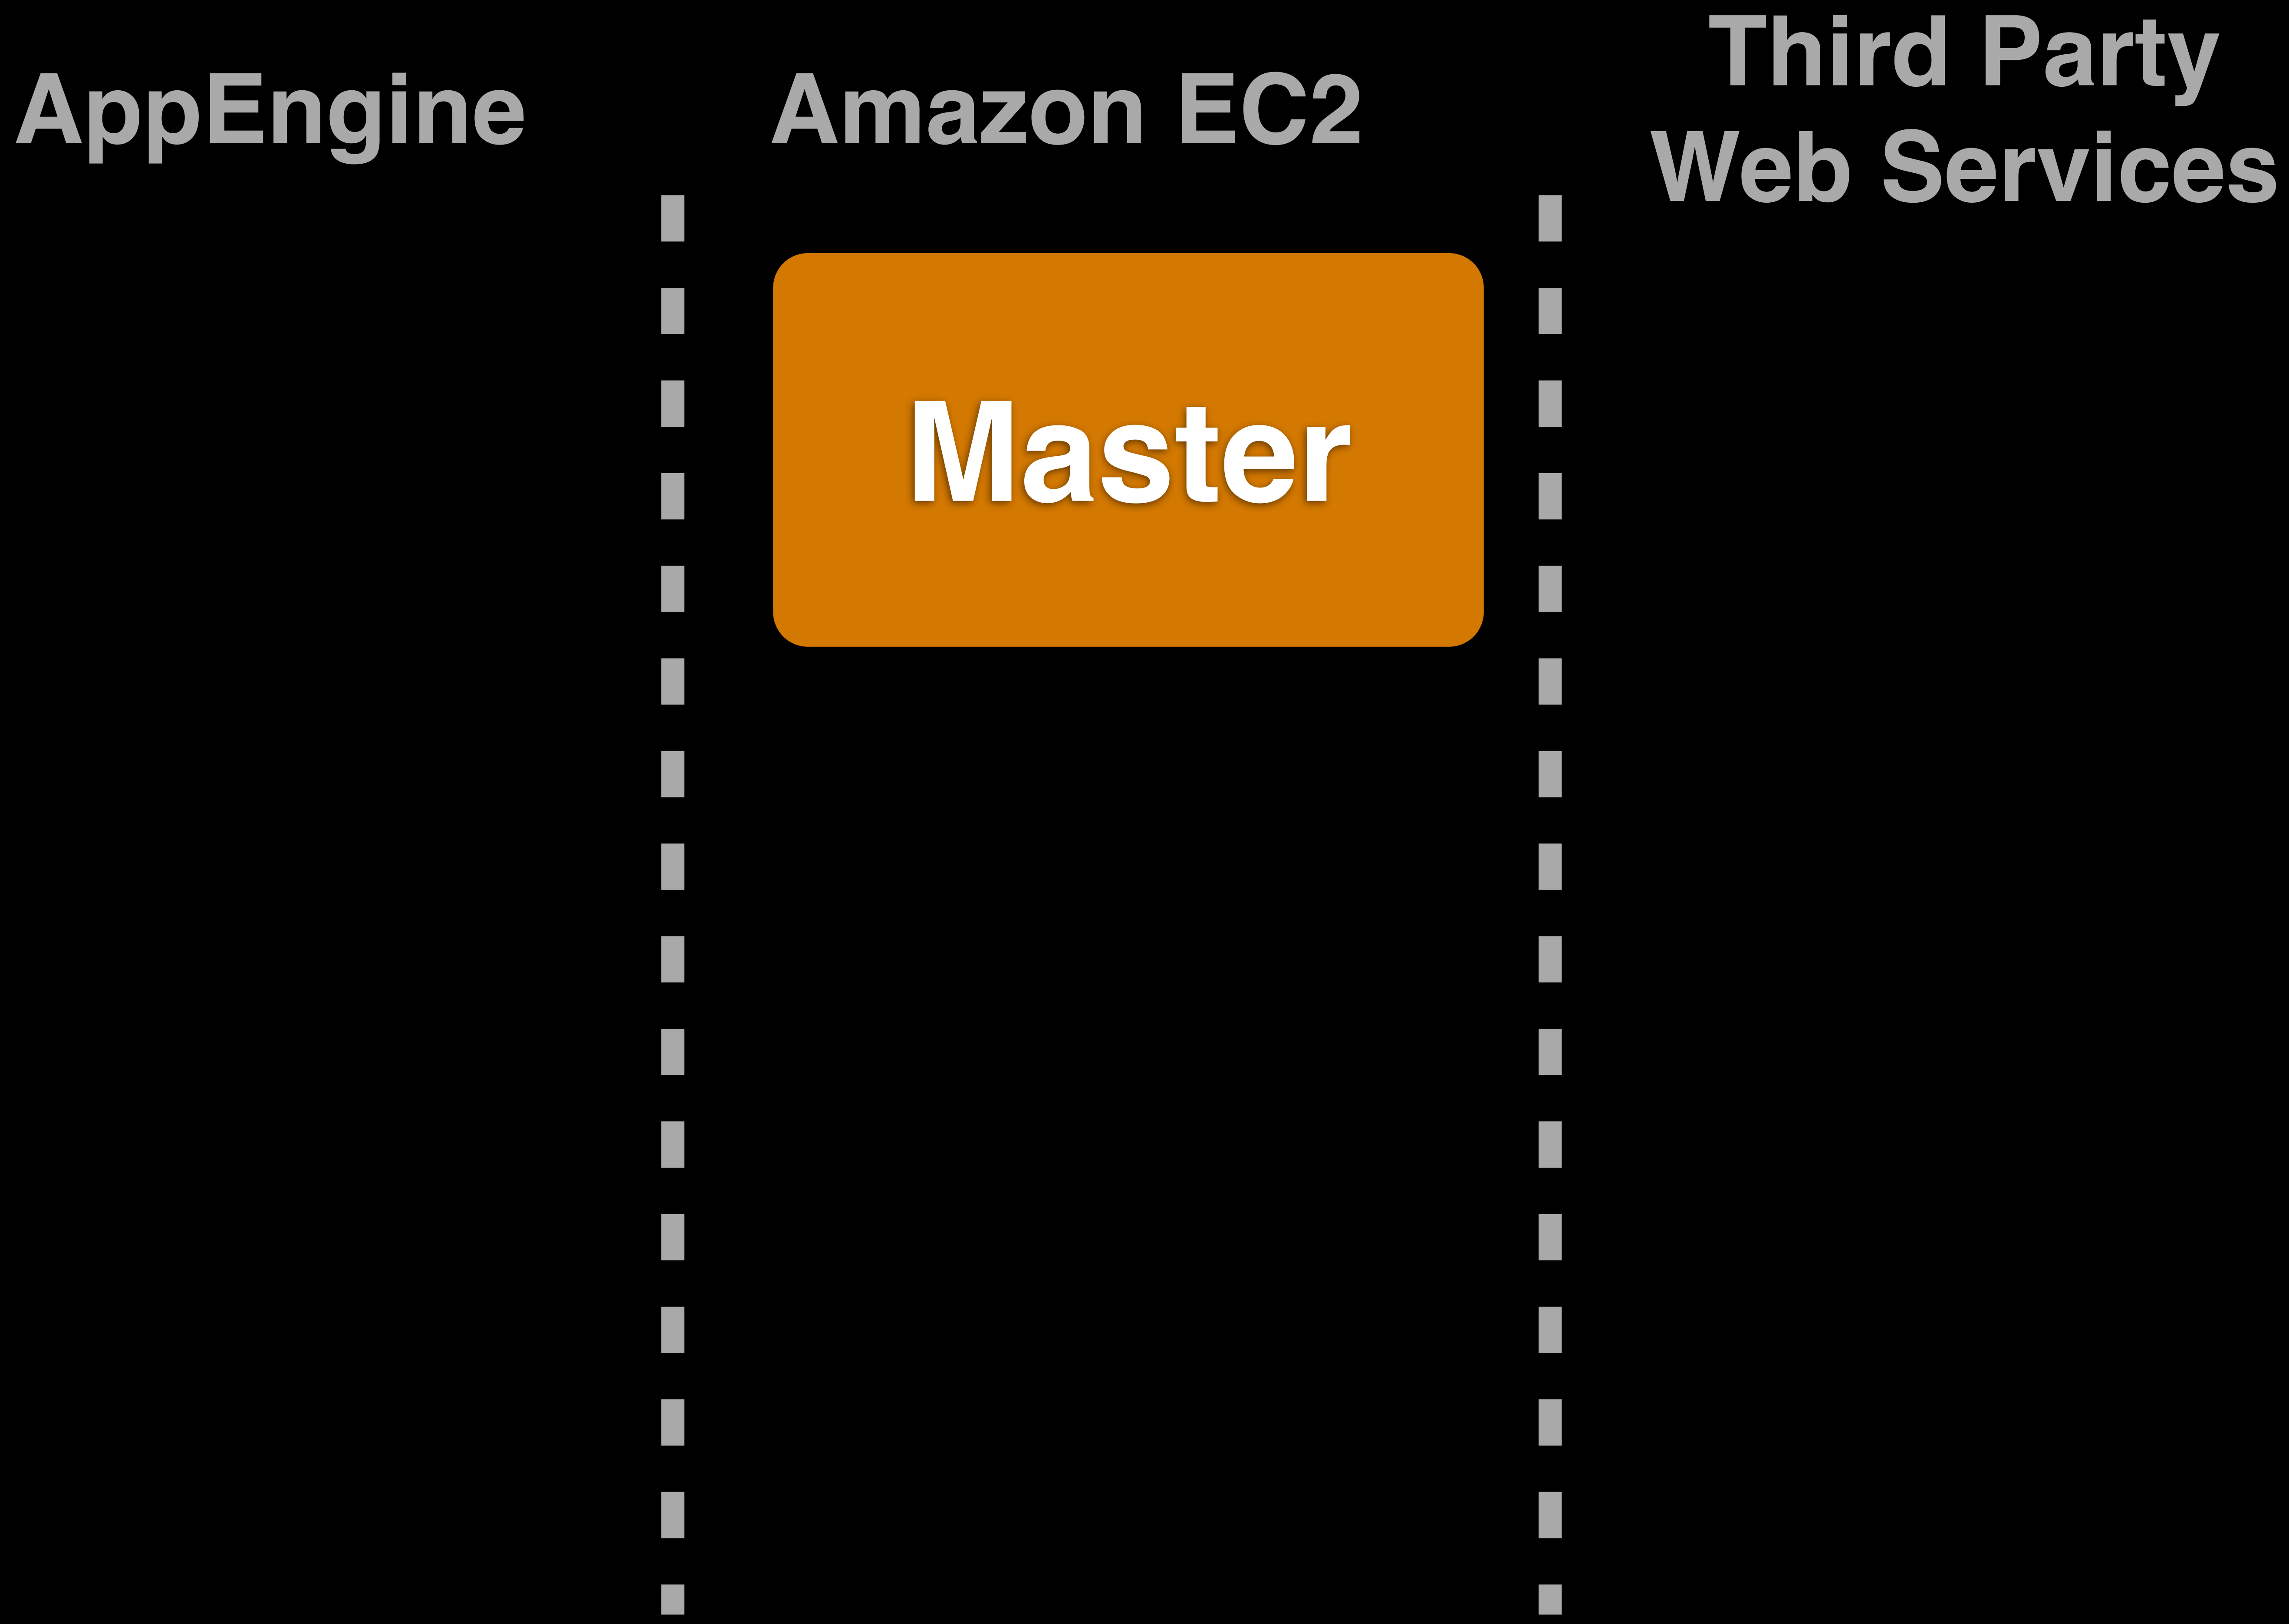

AppEngine

Amazon EC2

Third Party Web Services

Master

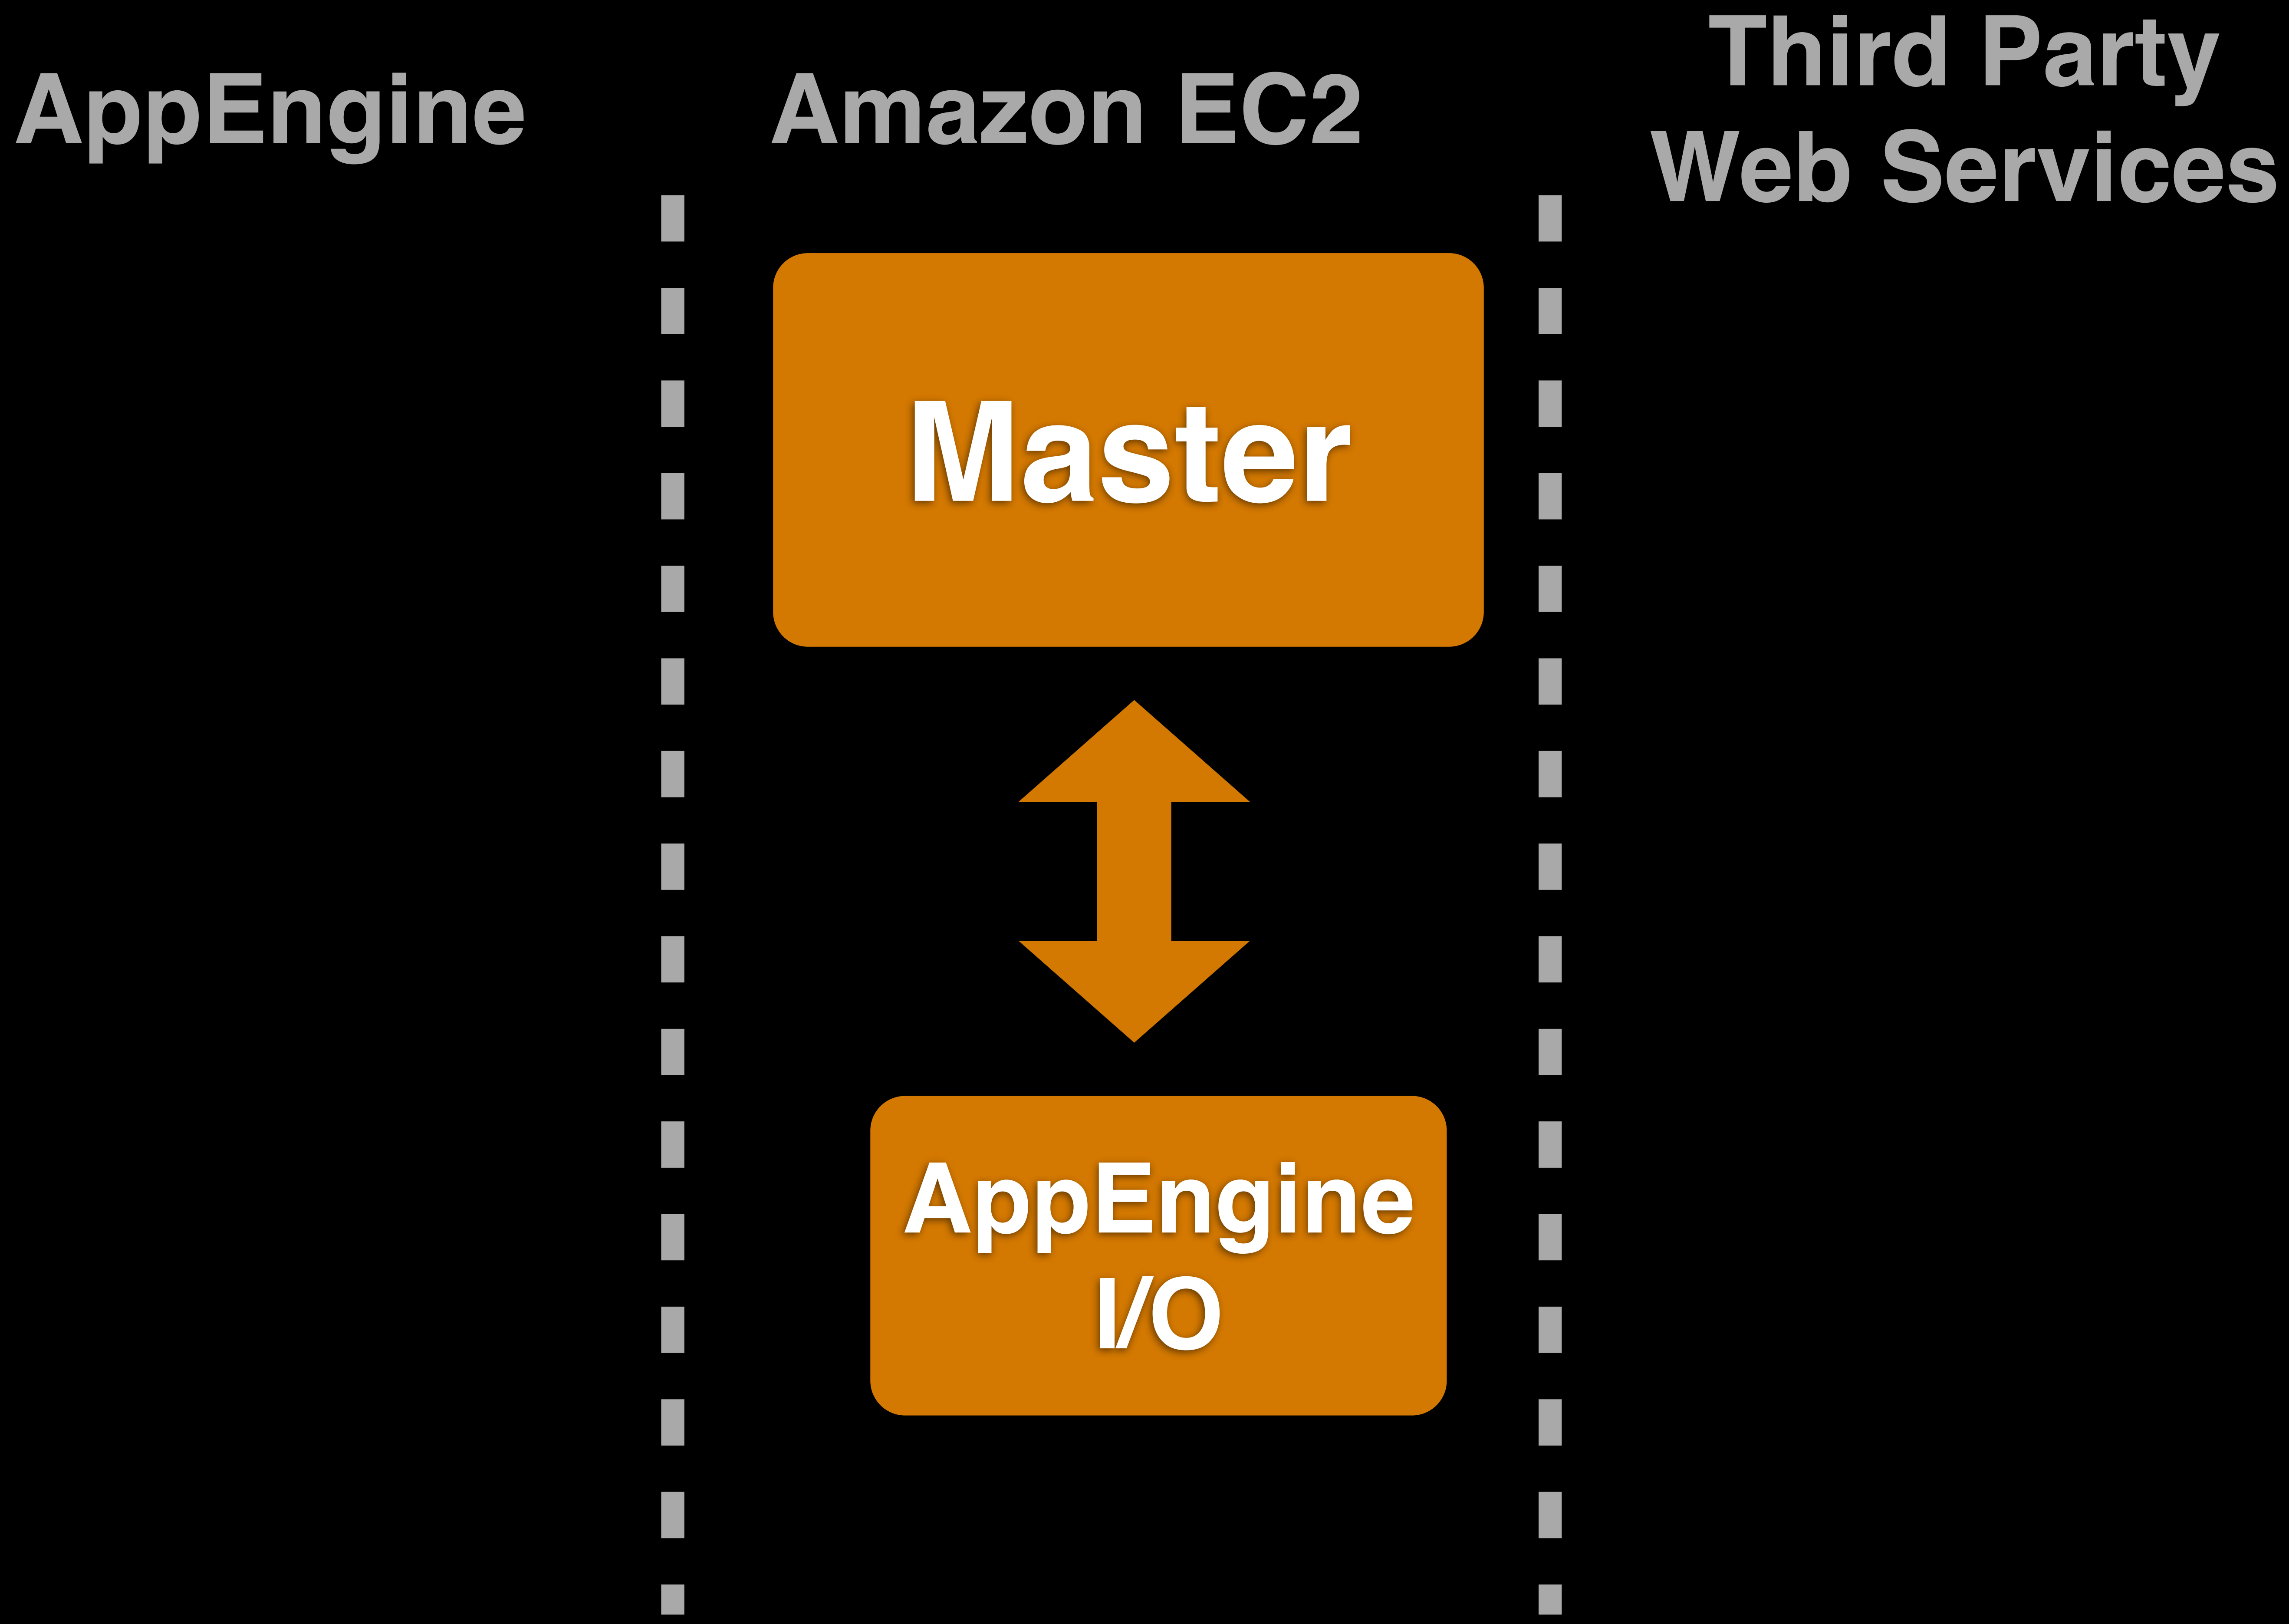

AppEngine

Amazon EC2

Third Party
Web Services

Master

AppEngine
I/O

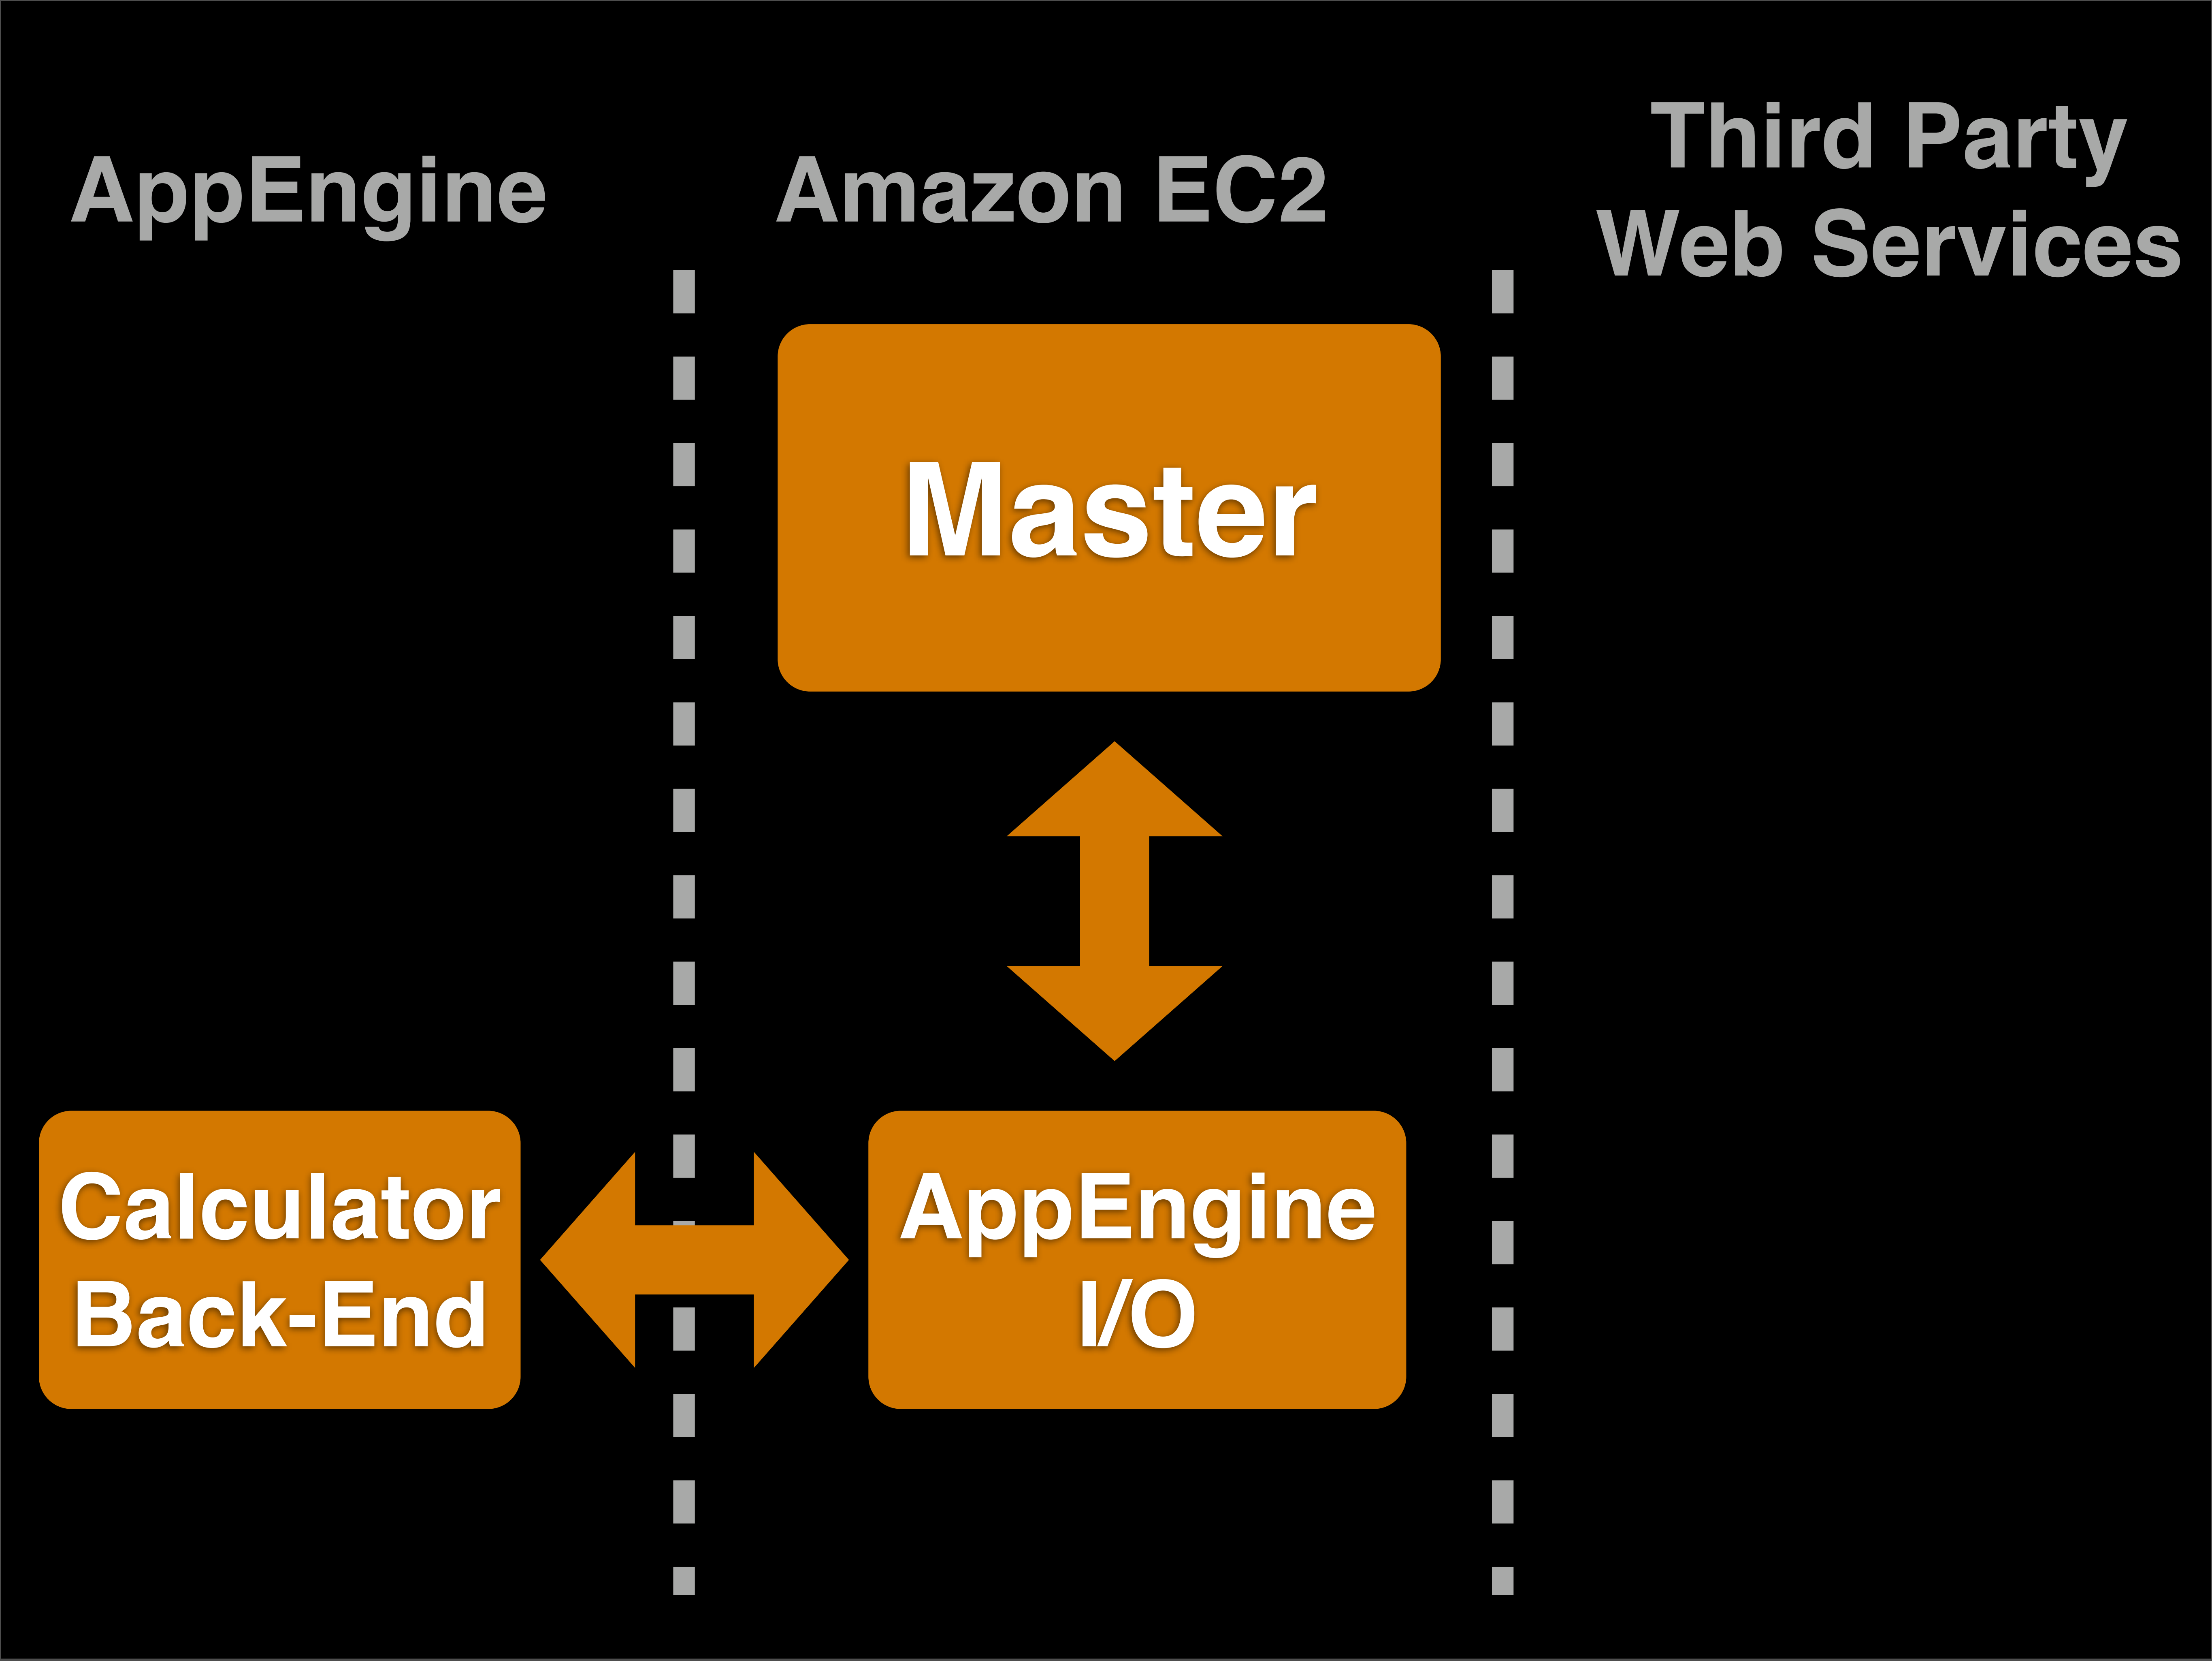
AppEngine
Amazon EC2
Third Party Web Services
Master
Calculator Back-End
AppEngine I/O

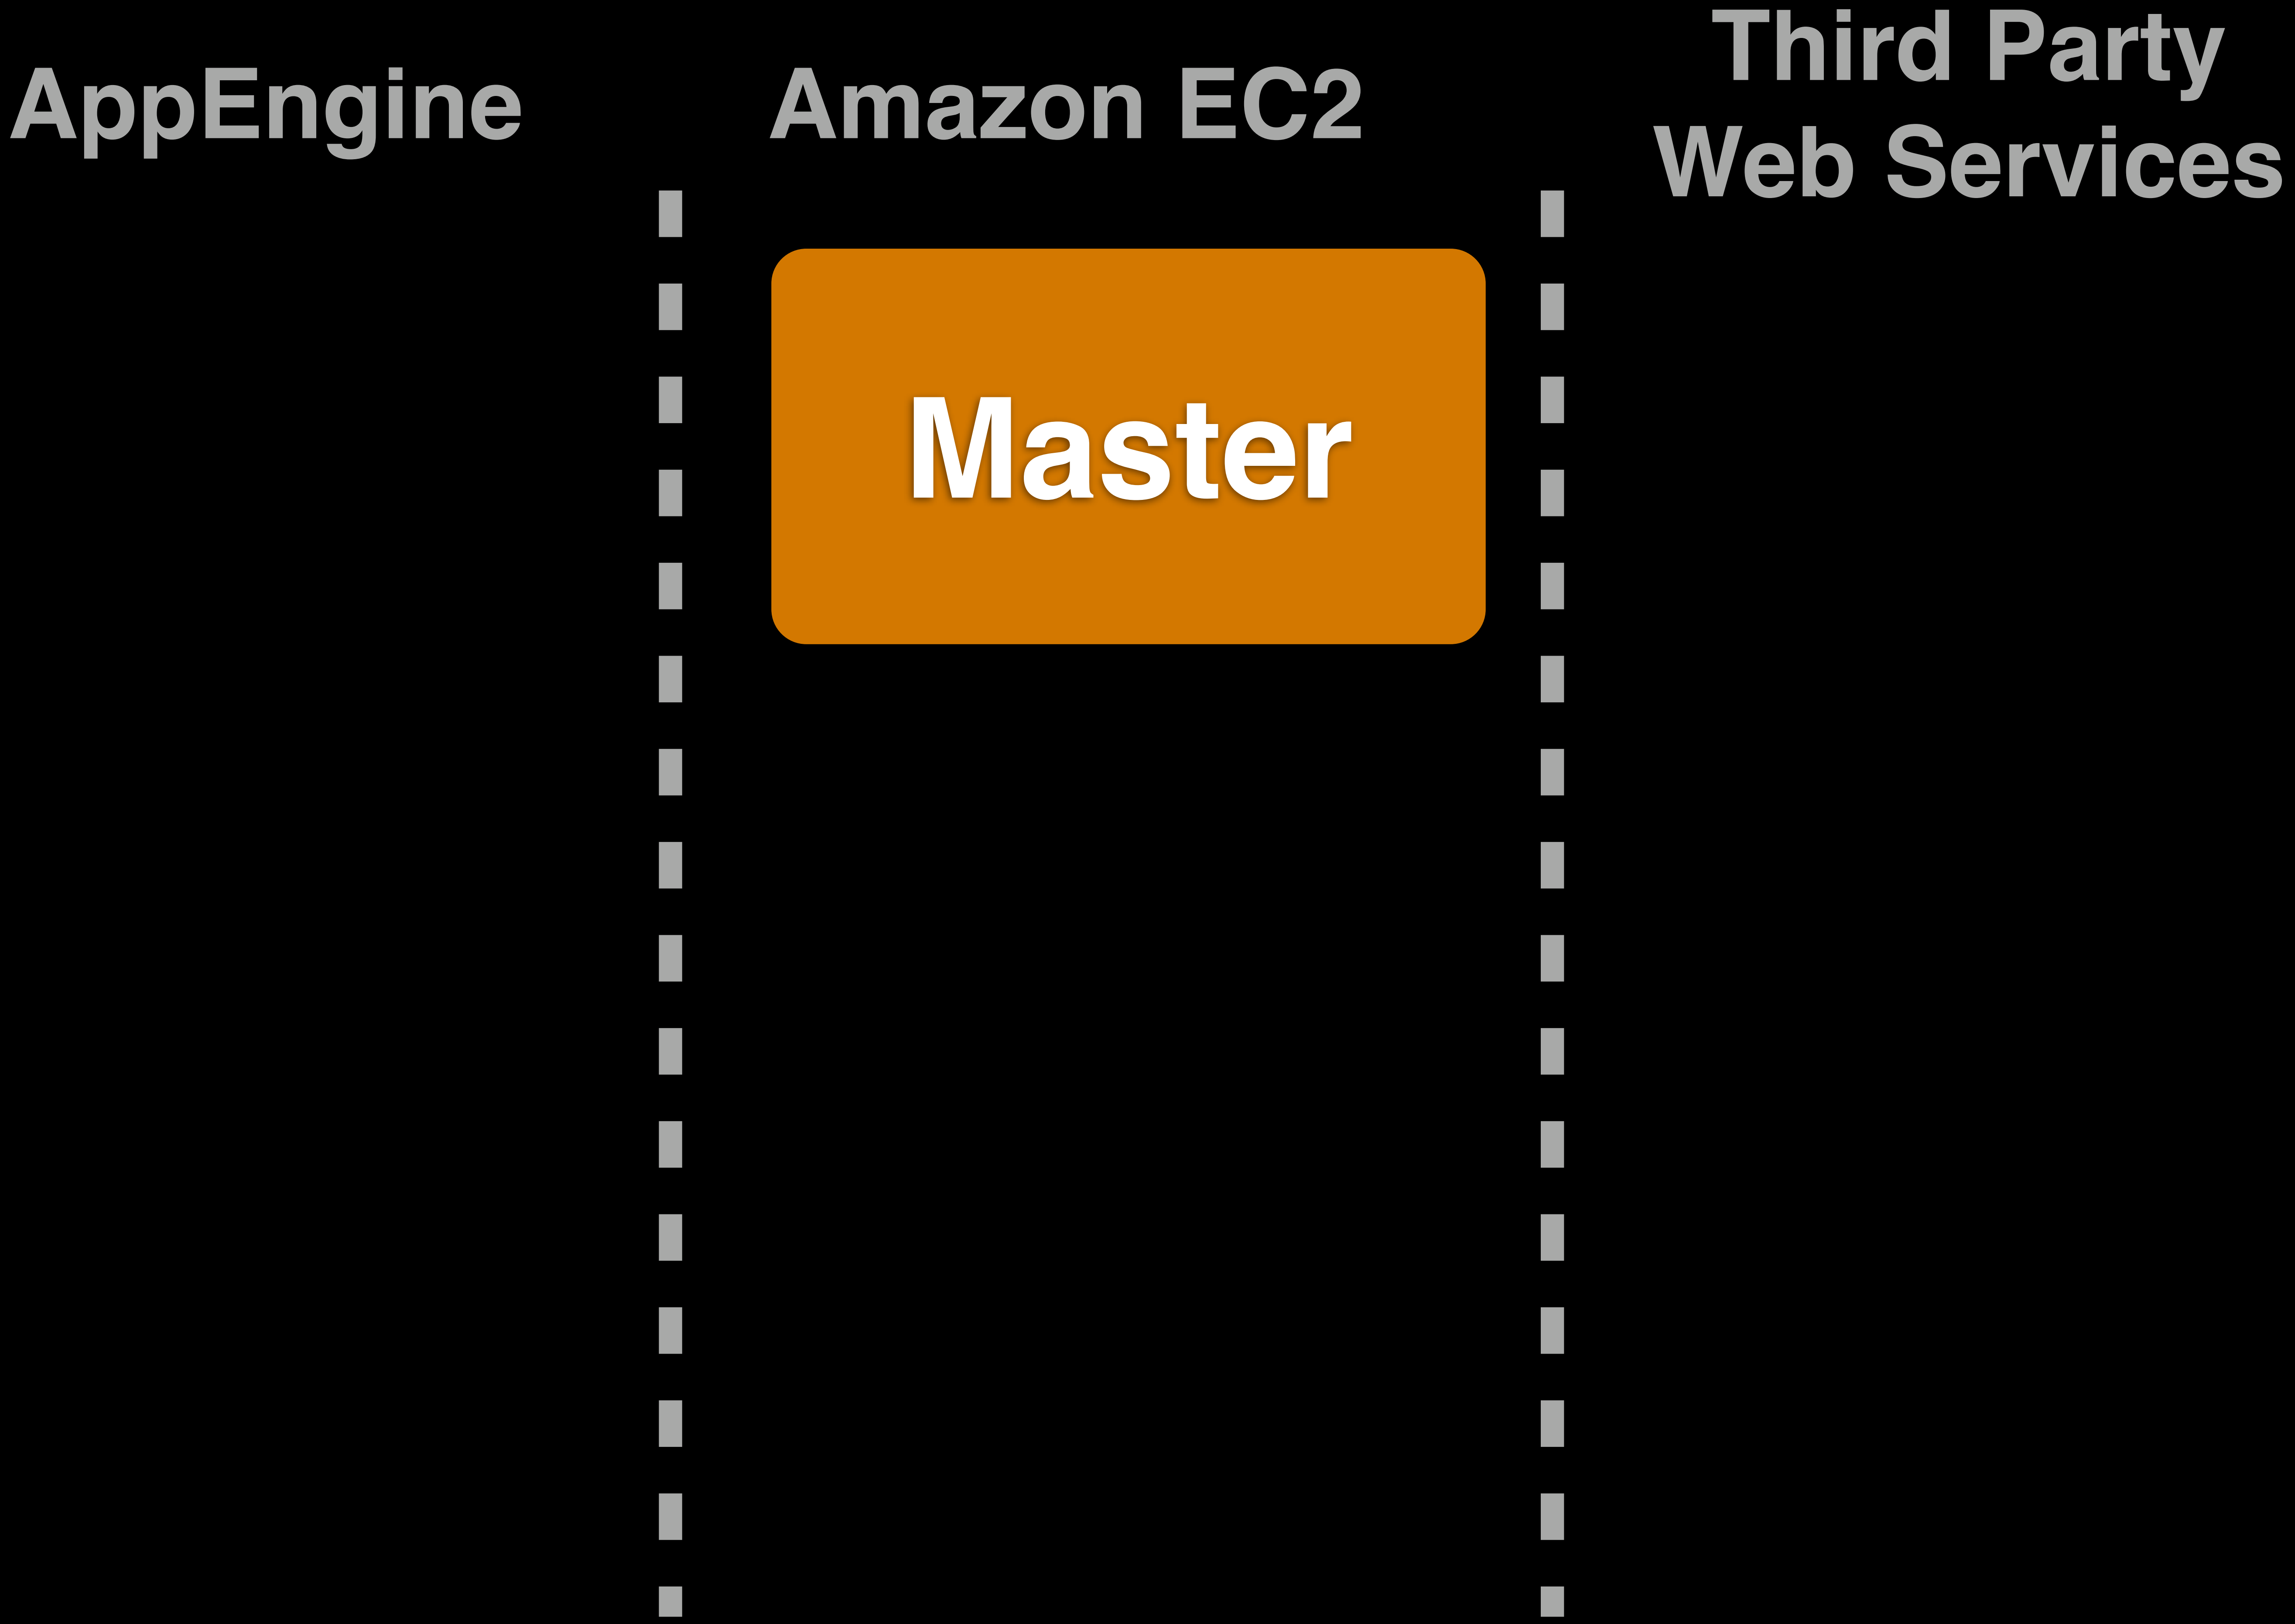

AppEngine

Amazon EC2

Third Party
Web Services

Master

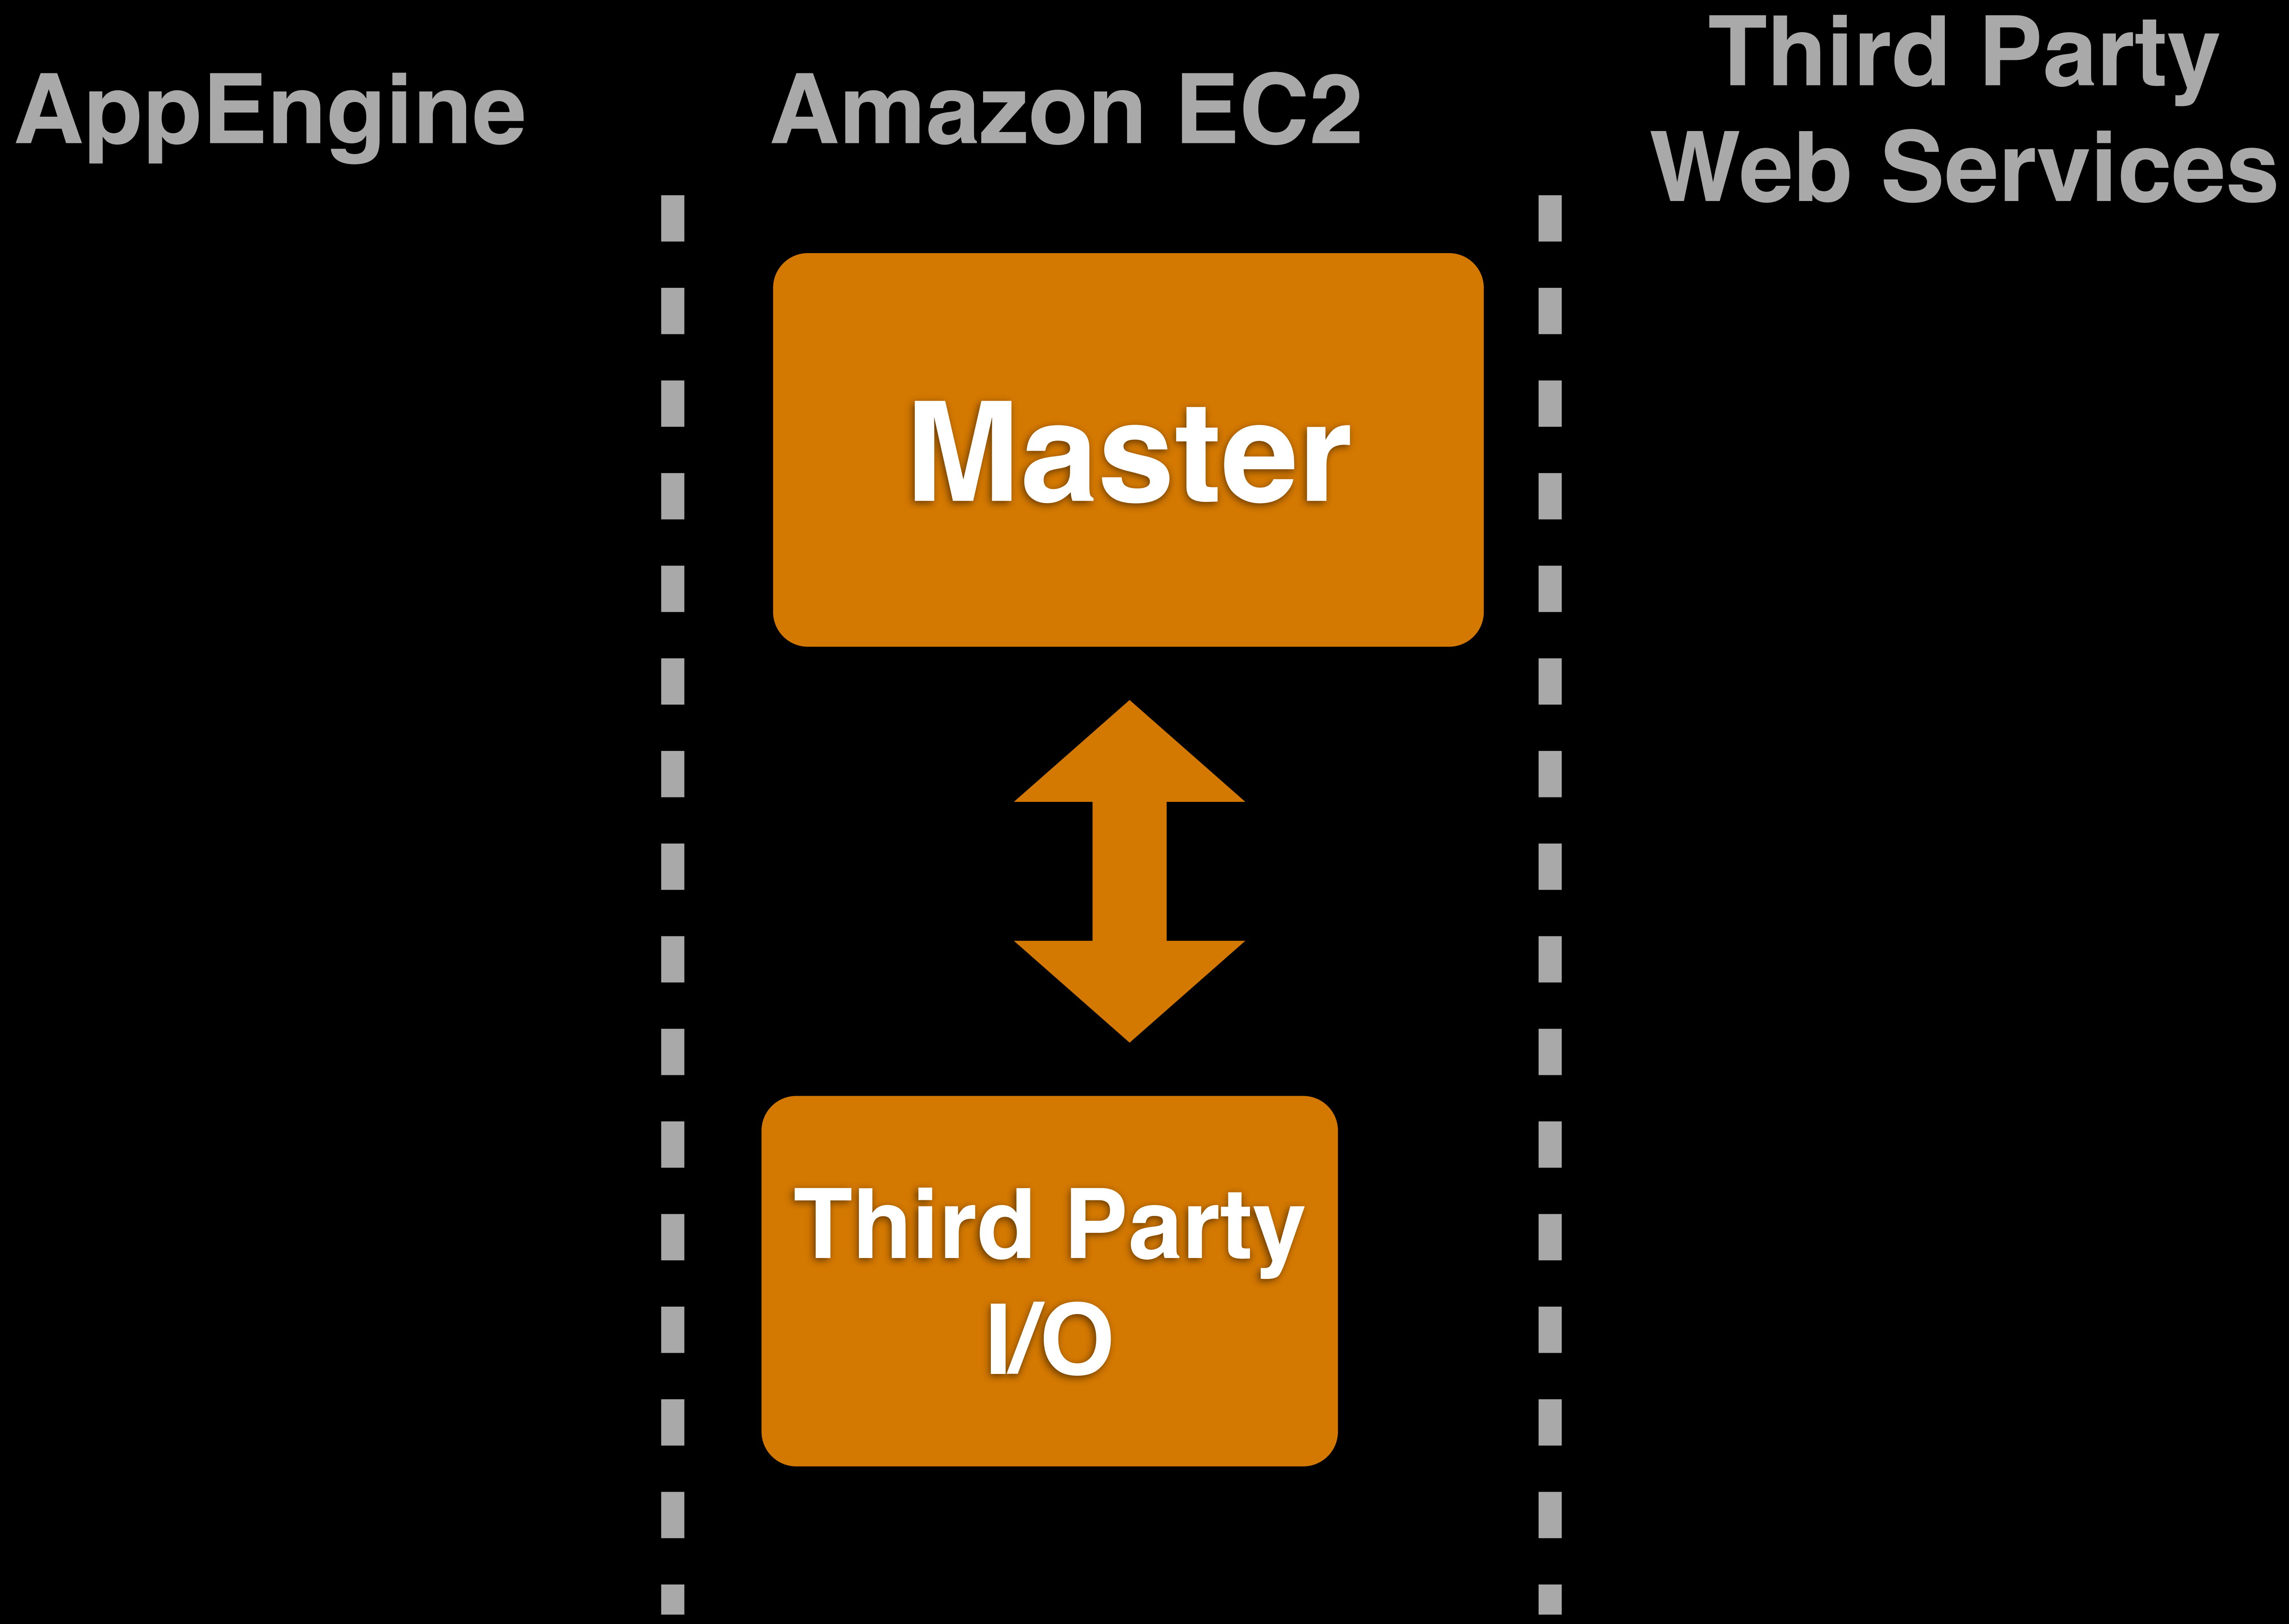

AppEngine

Amazon EC2

Third Party Web Services

Master

Third Party I/O

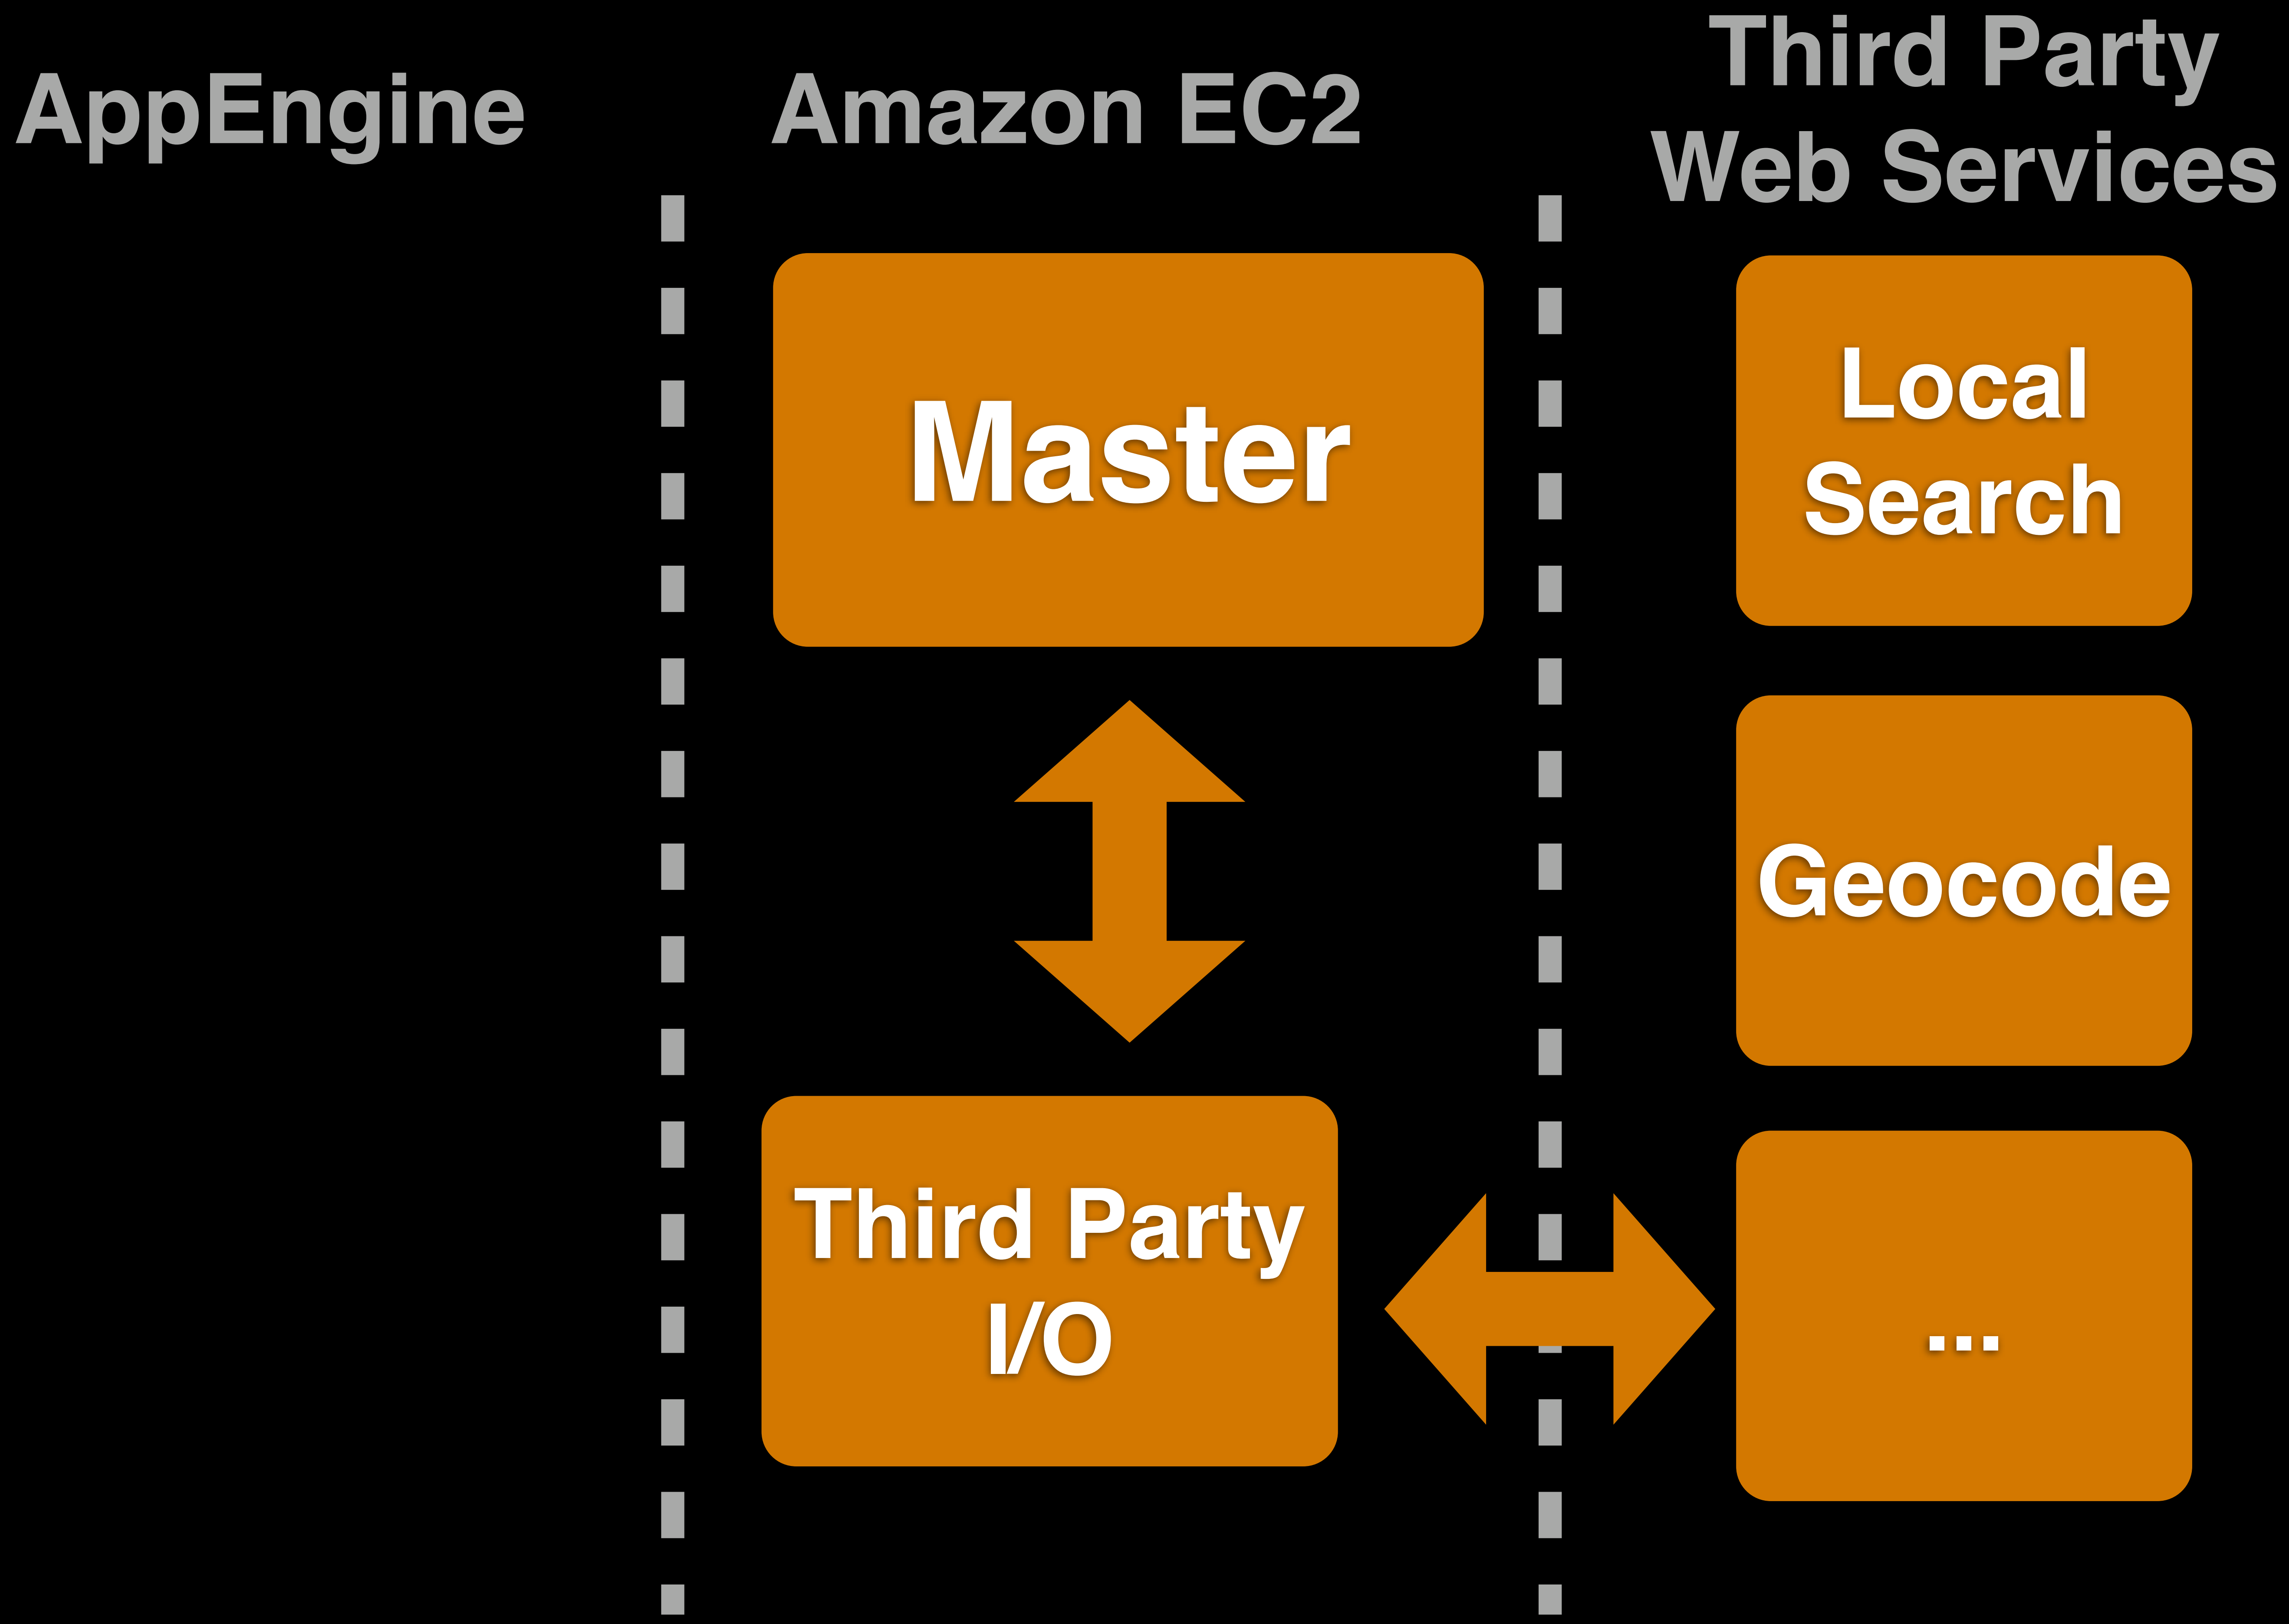

AppEngine

Amazon EC2

Third Party Web Services

Master

Third Party I/O

Local Search

Geocode

...

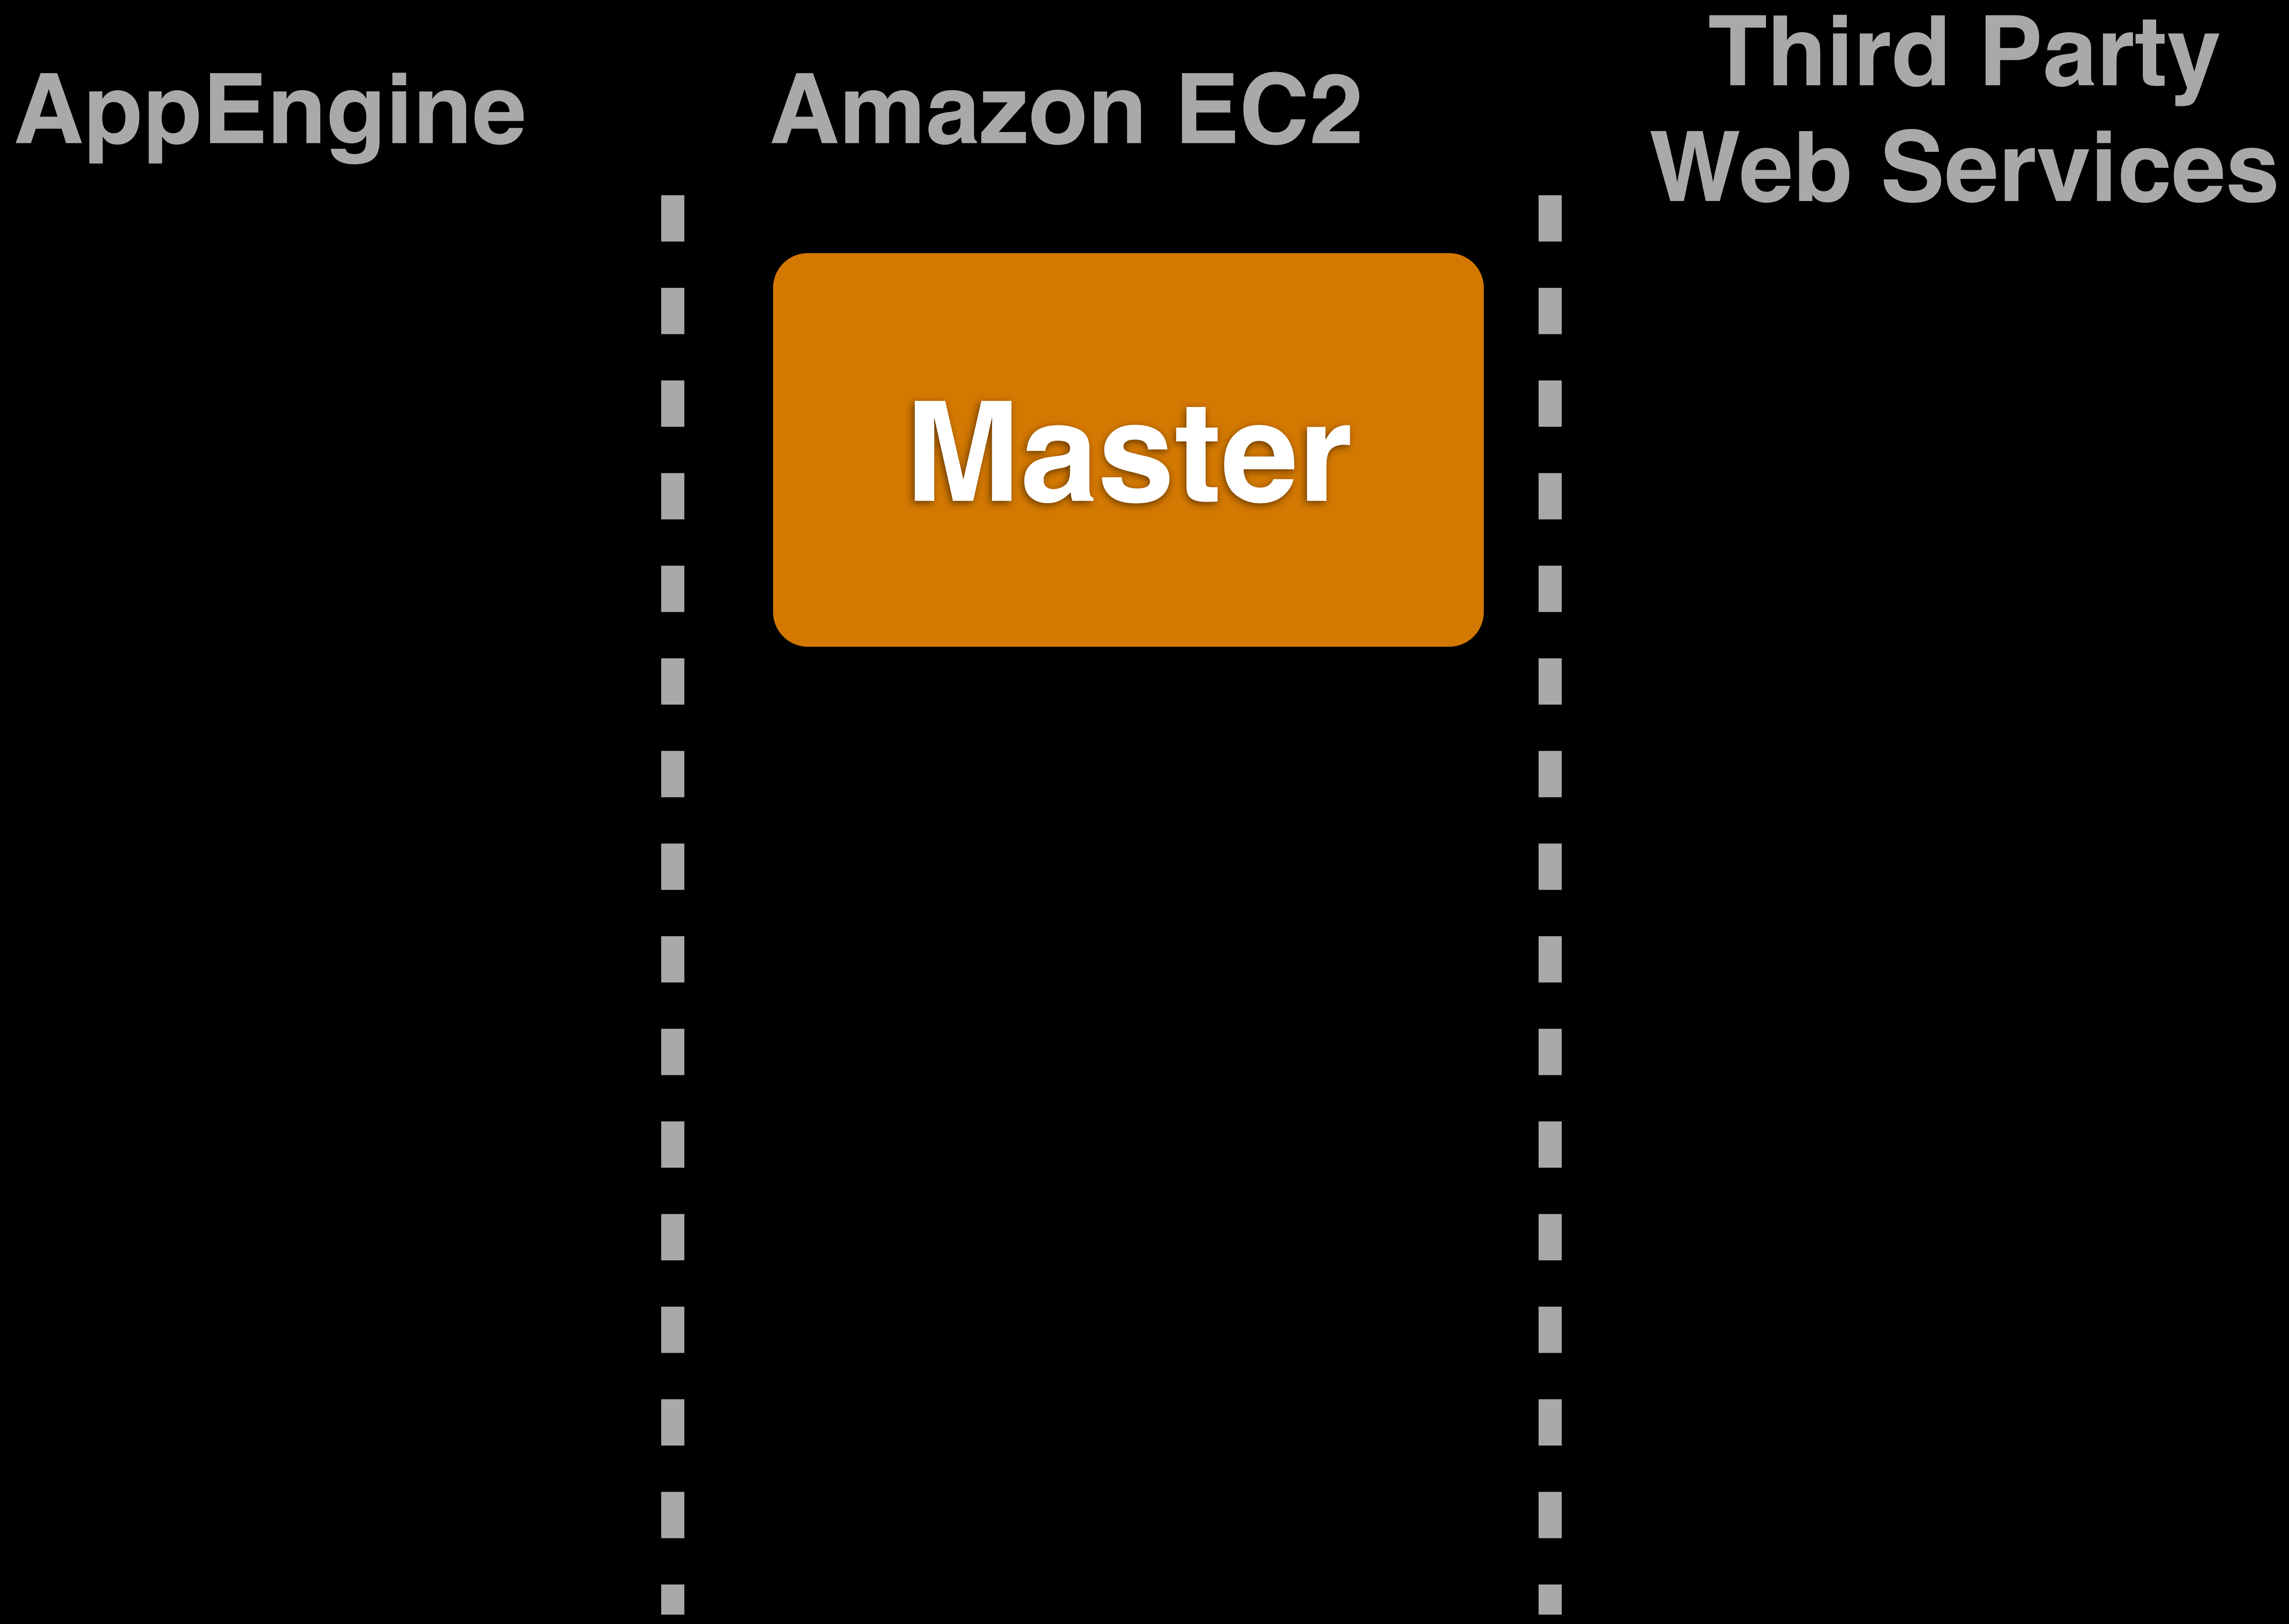

AppEngine

Amazon EC2

Third Party
Web Services

Master

AppEngine

Amazon EC2

Third Party
Web Services

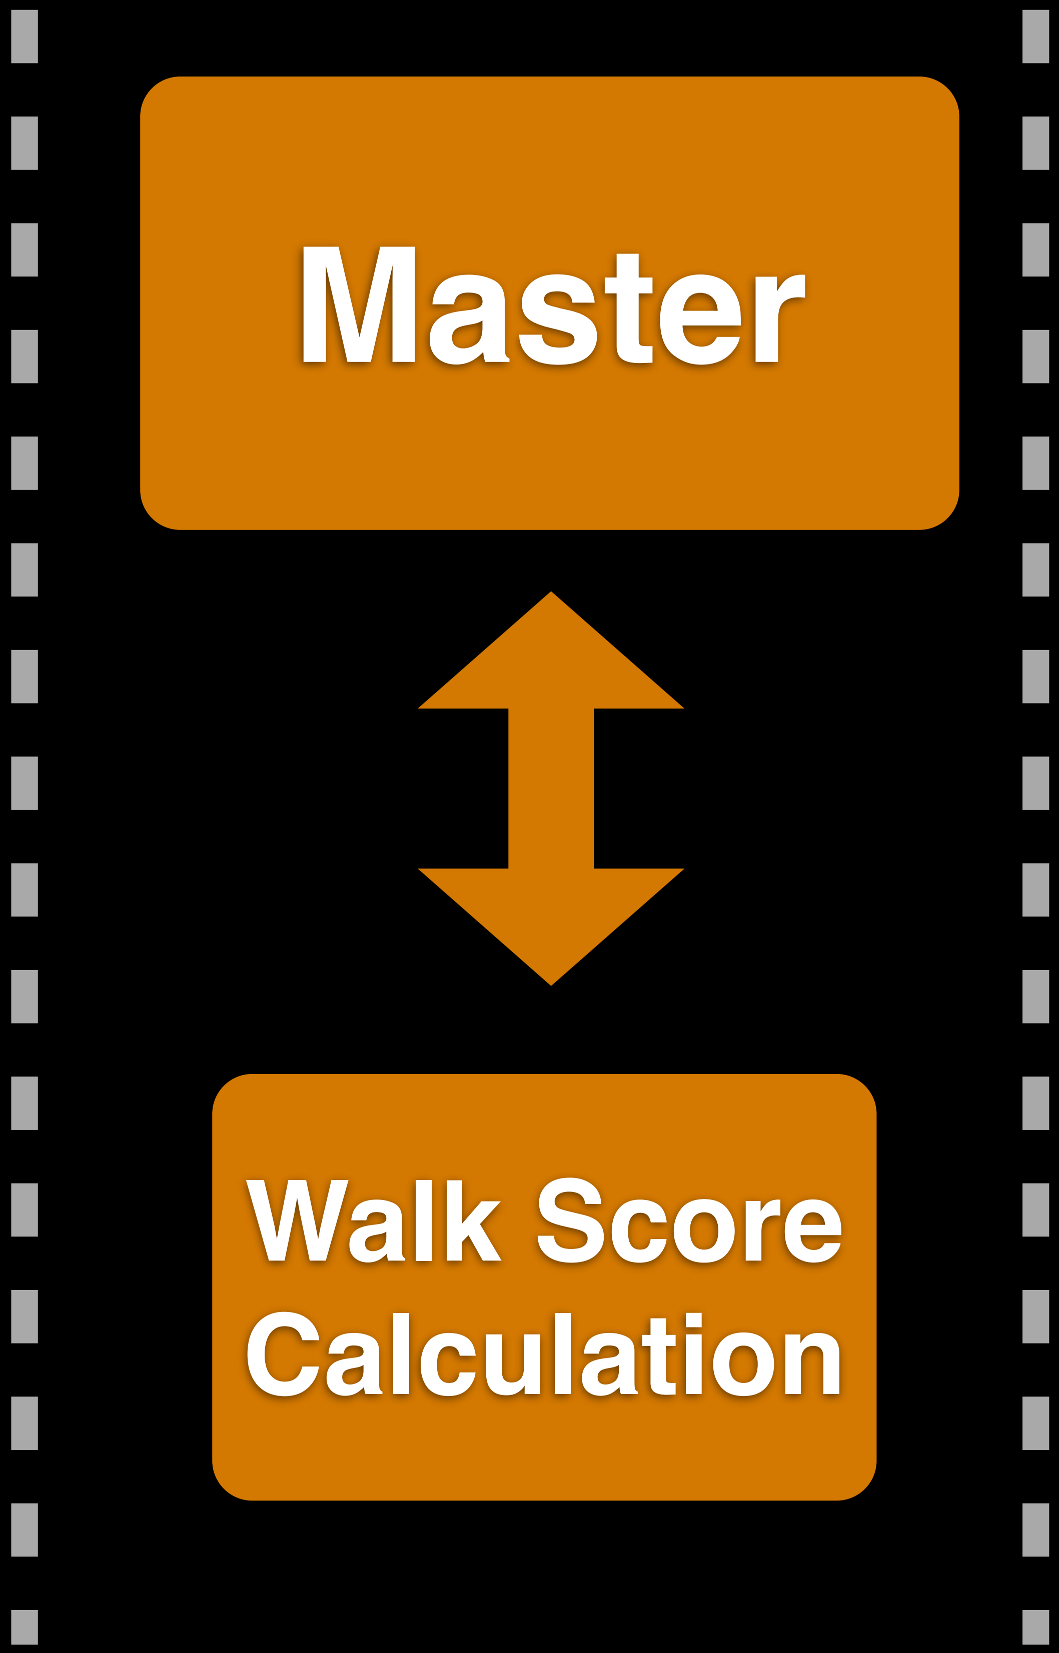

Master

Walk Score
Calculation

AppEngine Quirks & Bugs