Timelines @ Twitter

QCon London 2012

Arya Asemanfar

Poll-based Timeline

Tweets

Sam Pullara @sampullara 1m
@kevino it is mirroring -- i don't think there is a way to do as if it was a second monitor ... yet.
In reply to kevin oliver

news.yc Popular @newsycombinator 10m
Flowers regenerated from 30,000-year-old frozen fruit tissue j.mp/zPe7EH

The Onion @TheOnion 10m
SPORTSWIRE: Groans Abound As Tim Duncan Raises Hand Once Again At City Council Meeting onion.com/yuXGvC

news.yc Popular @newsycombinator 10m
How I built a Hacker News mobile web app j.mp/wUvRy5

Maria Popova @brainpicker 14m
YES "What we need is seriously engaged art that can teach us again that we're smart." David Foster Wallace on art vs TV j.mp/wVHltY

Poll-based Timeline

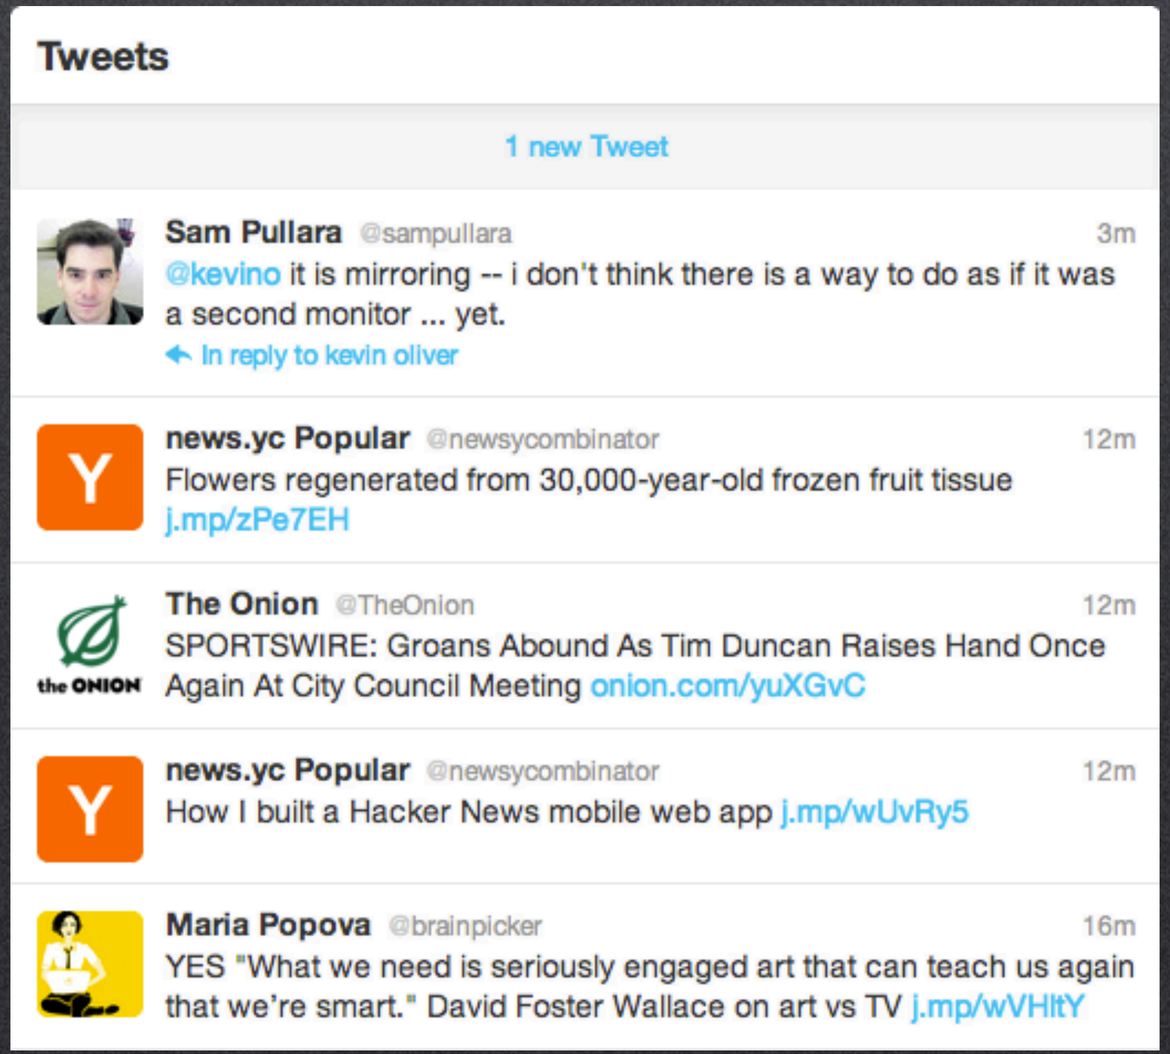

Search


QCon London
Search
Tweets Top / All
QCon London @qconlondon 1h
Vote on the most important software development trends for 2012 t.co/Zaio6Rge, this will be announced in the evening keynote tonight
Open Source @opensource_ 2h
Sold out QCon kicks off in London: big data, mobile, cloud, HTML 5 #opensource t.co/csHookQv
Krzysztof Wilczynski @kwilczynski 2h
New slide from QCon 2012 by @patrickdebois t.co/d9RbHru7
View media
Dušan Omerčević @dusano 3h
#QConLondon 2012 tutorial on Continuous delivery aka. Making deployment a non-event t.co/HXYMVS6H
View media
Liisi Toom @liisi 3h
QCon London is ON! ZeroTurnaround is in the house! #fb pic.twitter.com/xLWTrDMC
View photo

Streaming (aka Push)

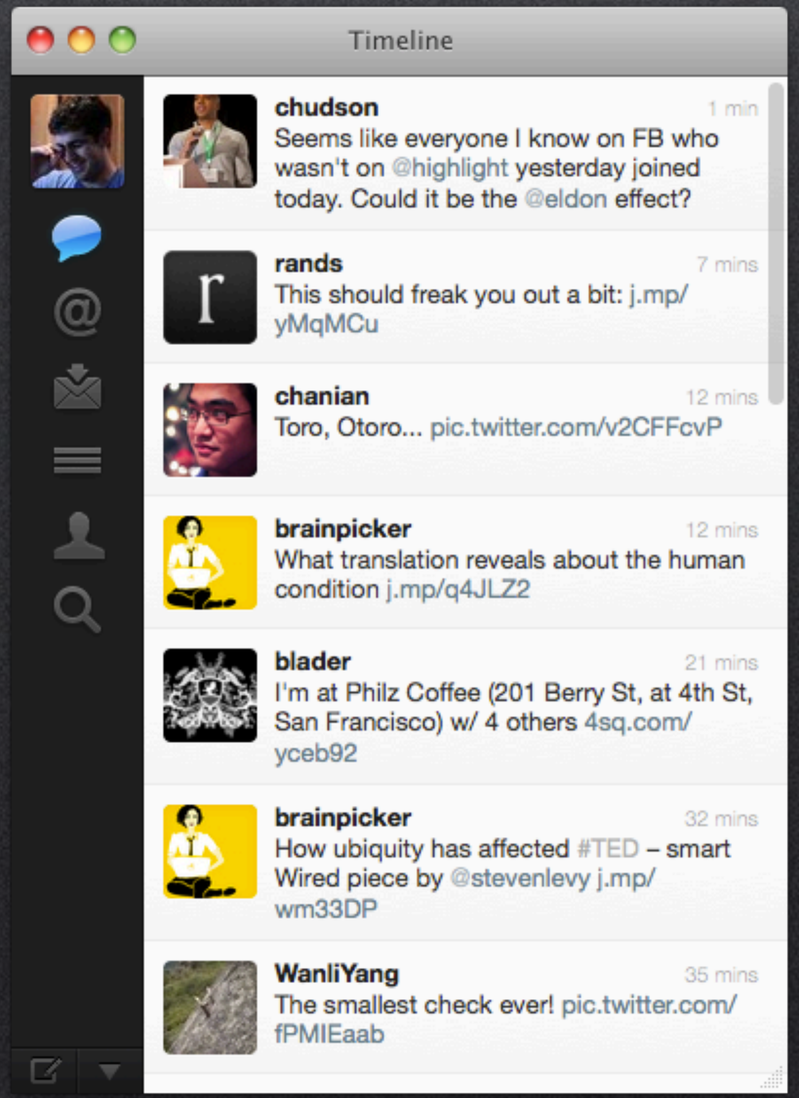

Timeline Delivery

100M+ active users worldwide

Poll-based

- 200K QPS
- latency* 1ms median, 4ms @ p99

Search

- 30K QPS

* to determine which tweets you should see

Timeline Delivery

250M New Tweets Per Day

- 3K/sec - daily average
- 5k/sec - daily peak
- 10K+/sec - several large events in February

users' home timelines

new tweet

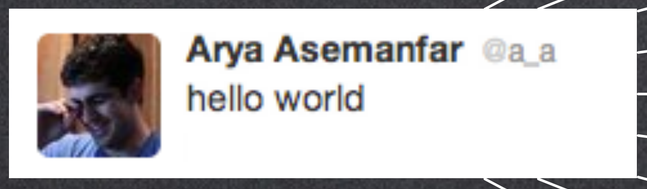

Timeline Delivery

26 billion deliveries every day

- 18 million per minute

Latency to deliver to 1 million followers

- 3.5 seconds average
- or 285K per second

Architecture

What's happening?

hello world

129 Tweet


HTTP
Proxy


HTTP
Proxy
tweet
API


HTTP
Proxy
tweet
API
queue

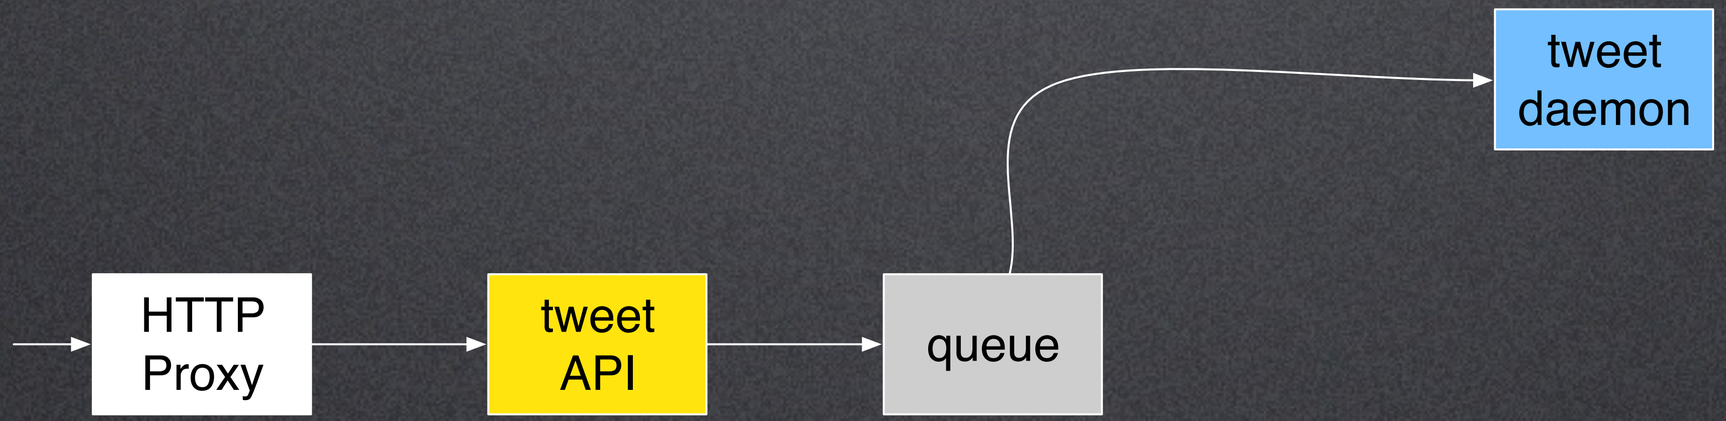
HTTP
Proxy
tweet
API
queue
tweet
daemon

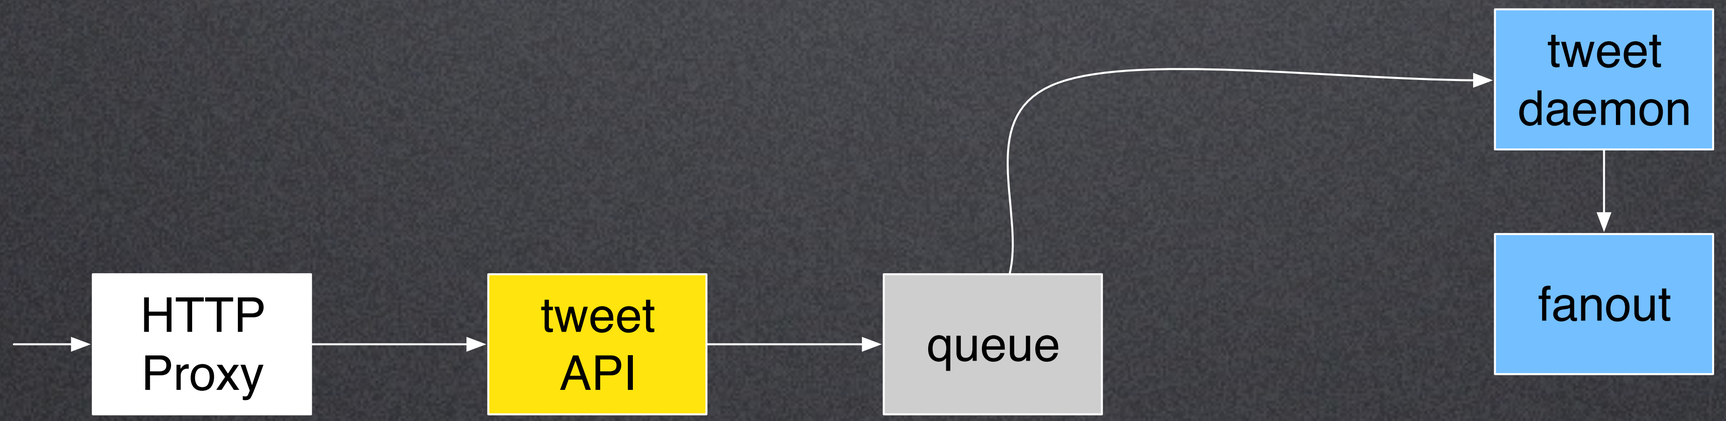
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout

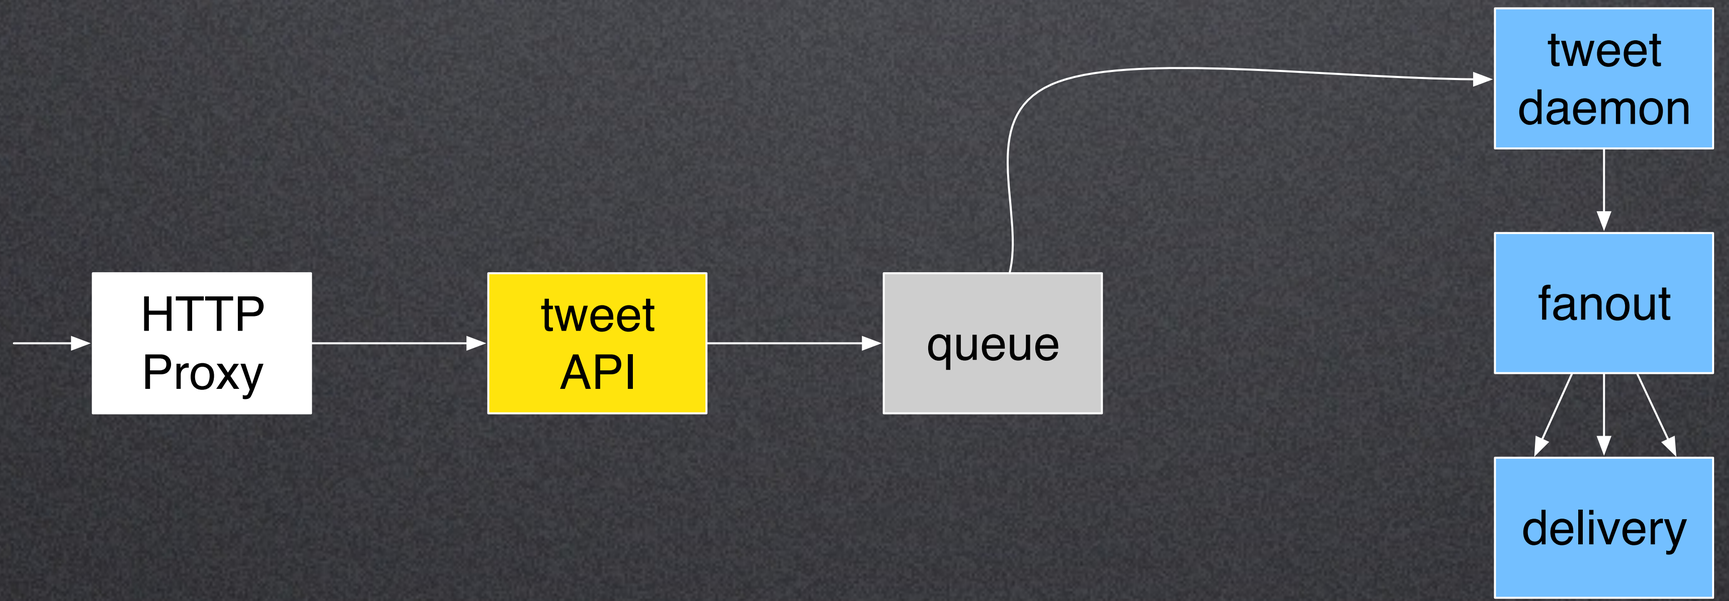
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery

Pipelined Delivery

fanout

Pipelined Delivery

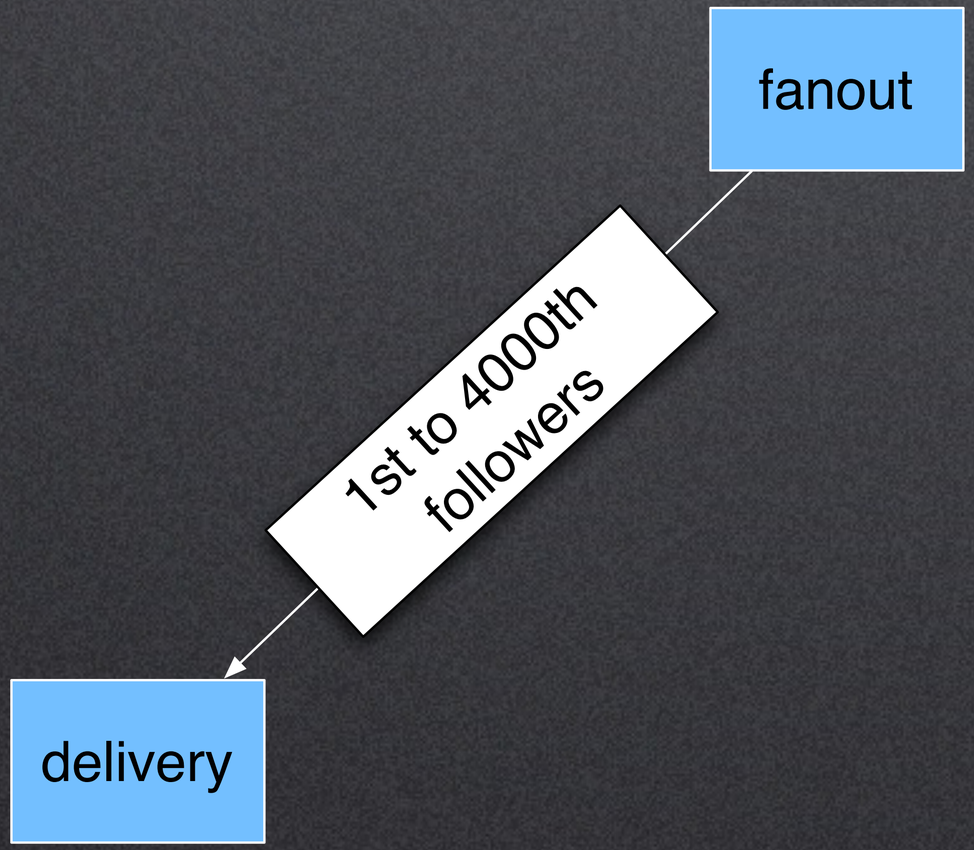

Pipelined Delivery

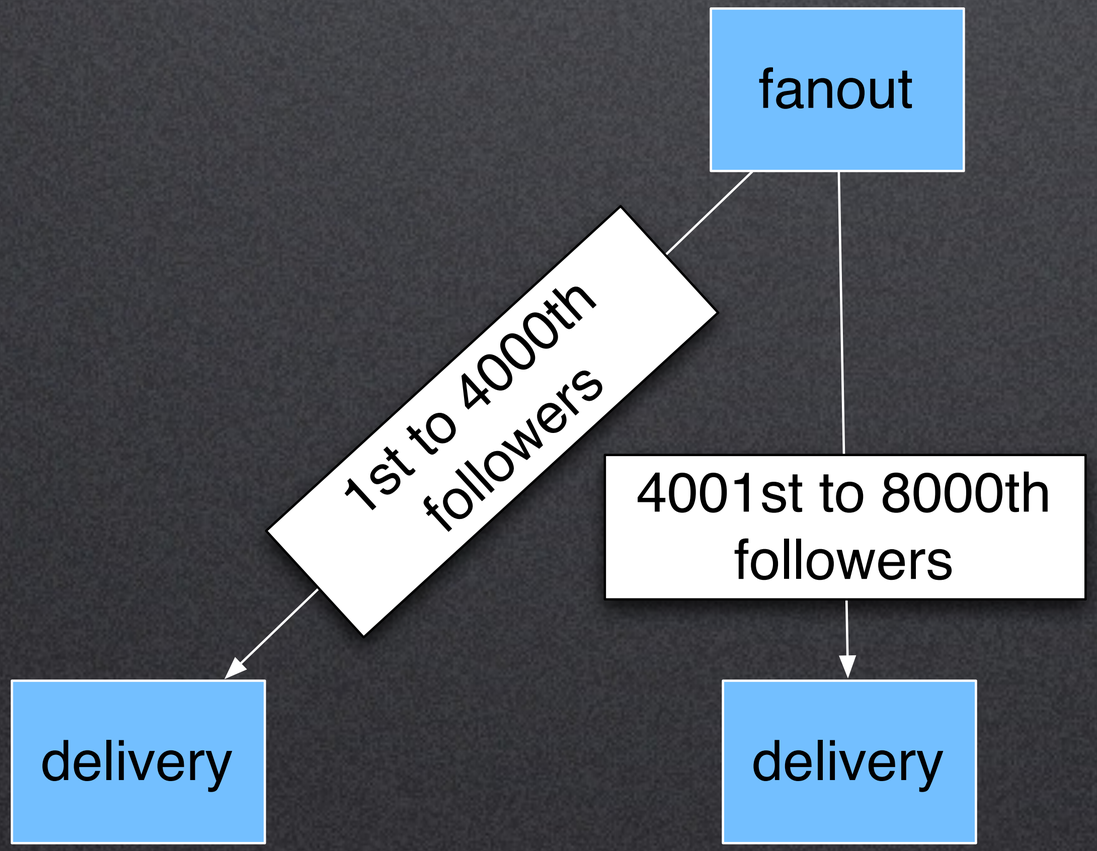

Pipelined Delivery

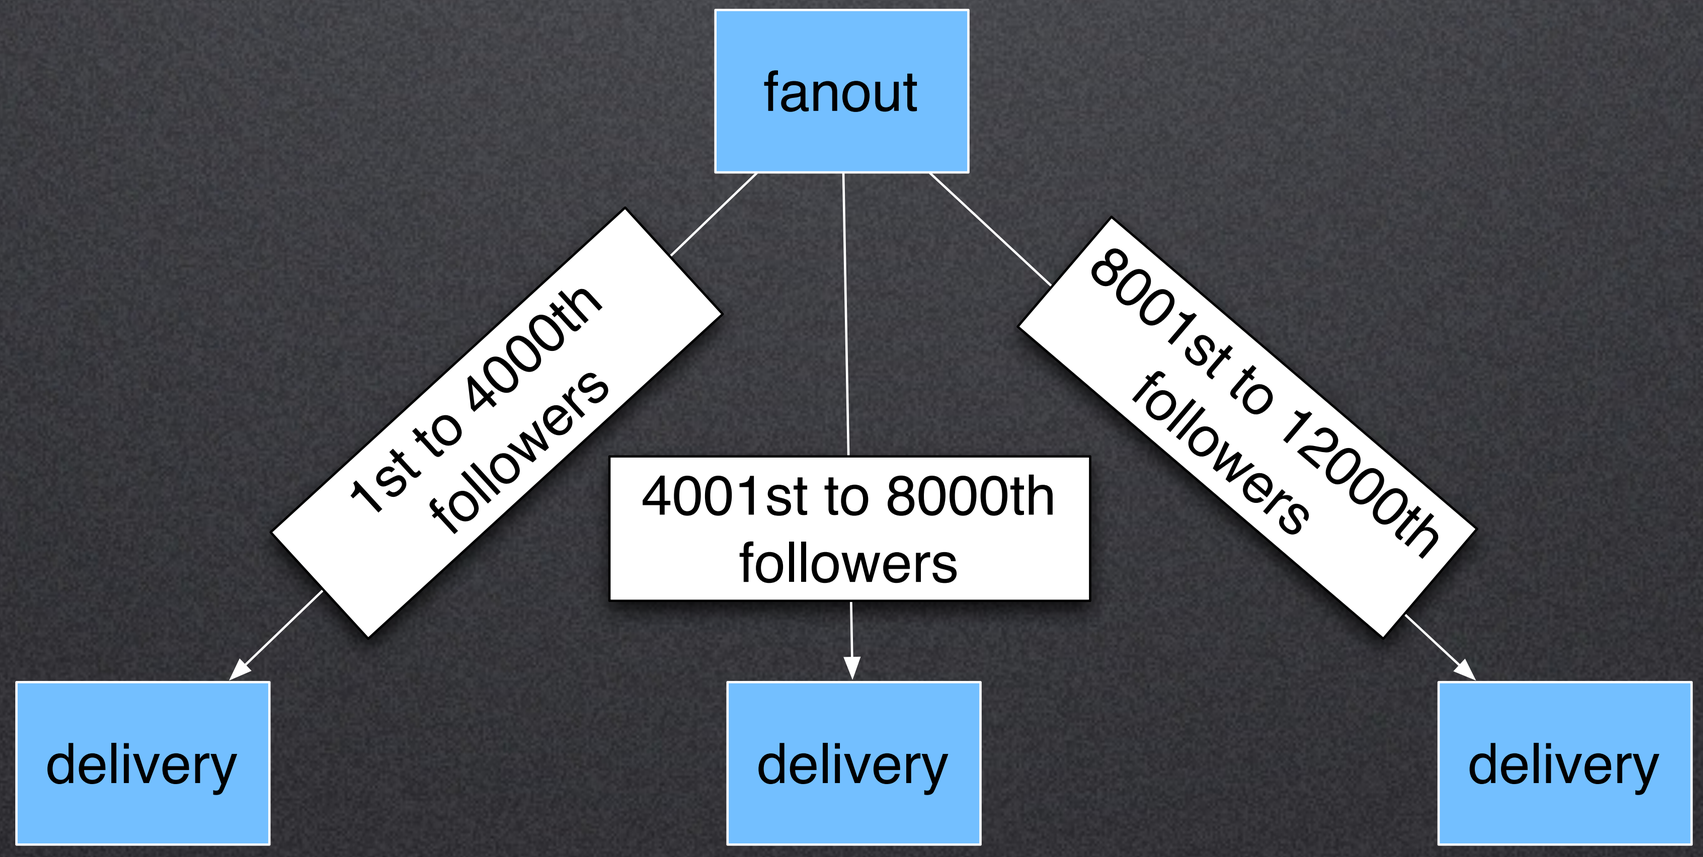

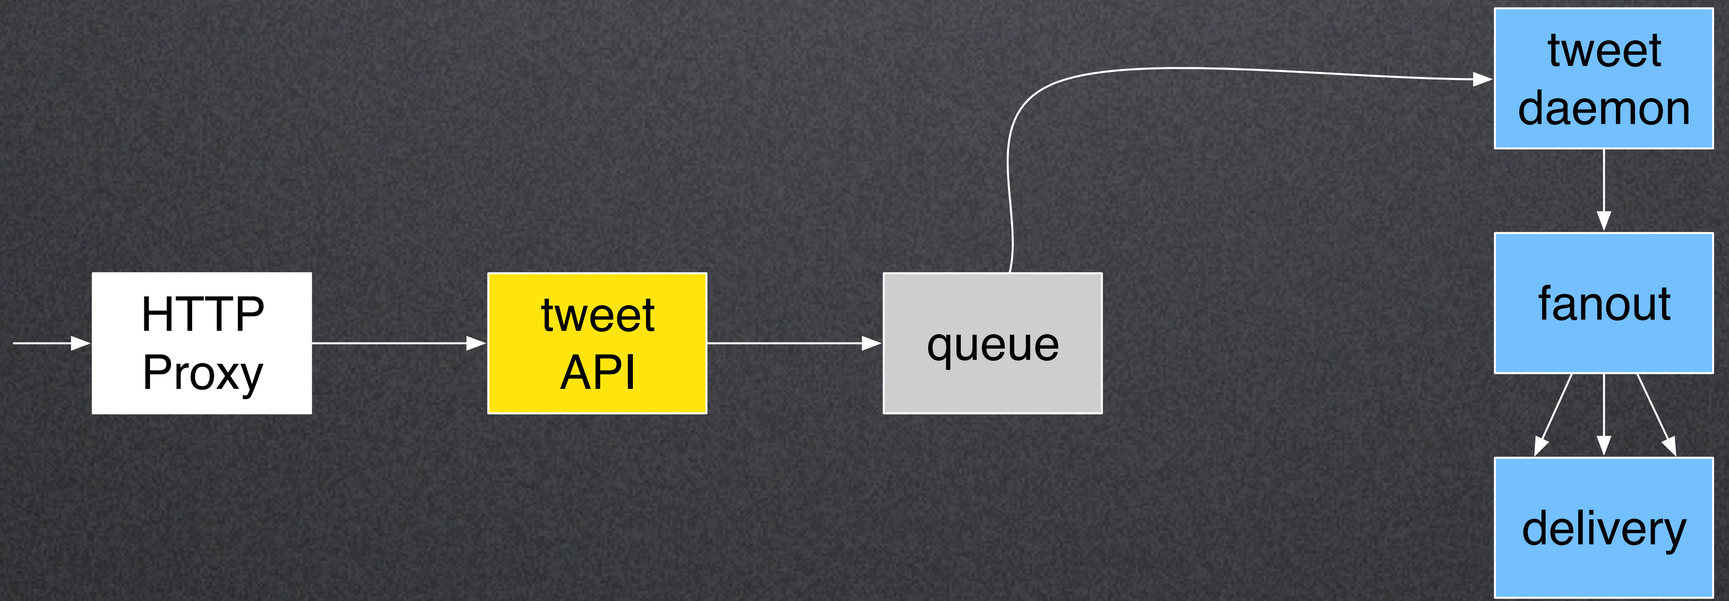
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery

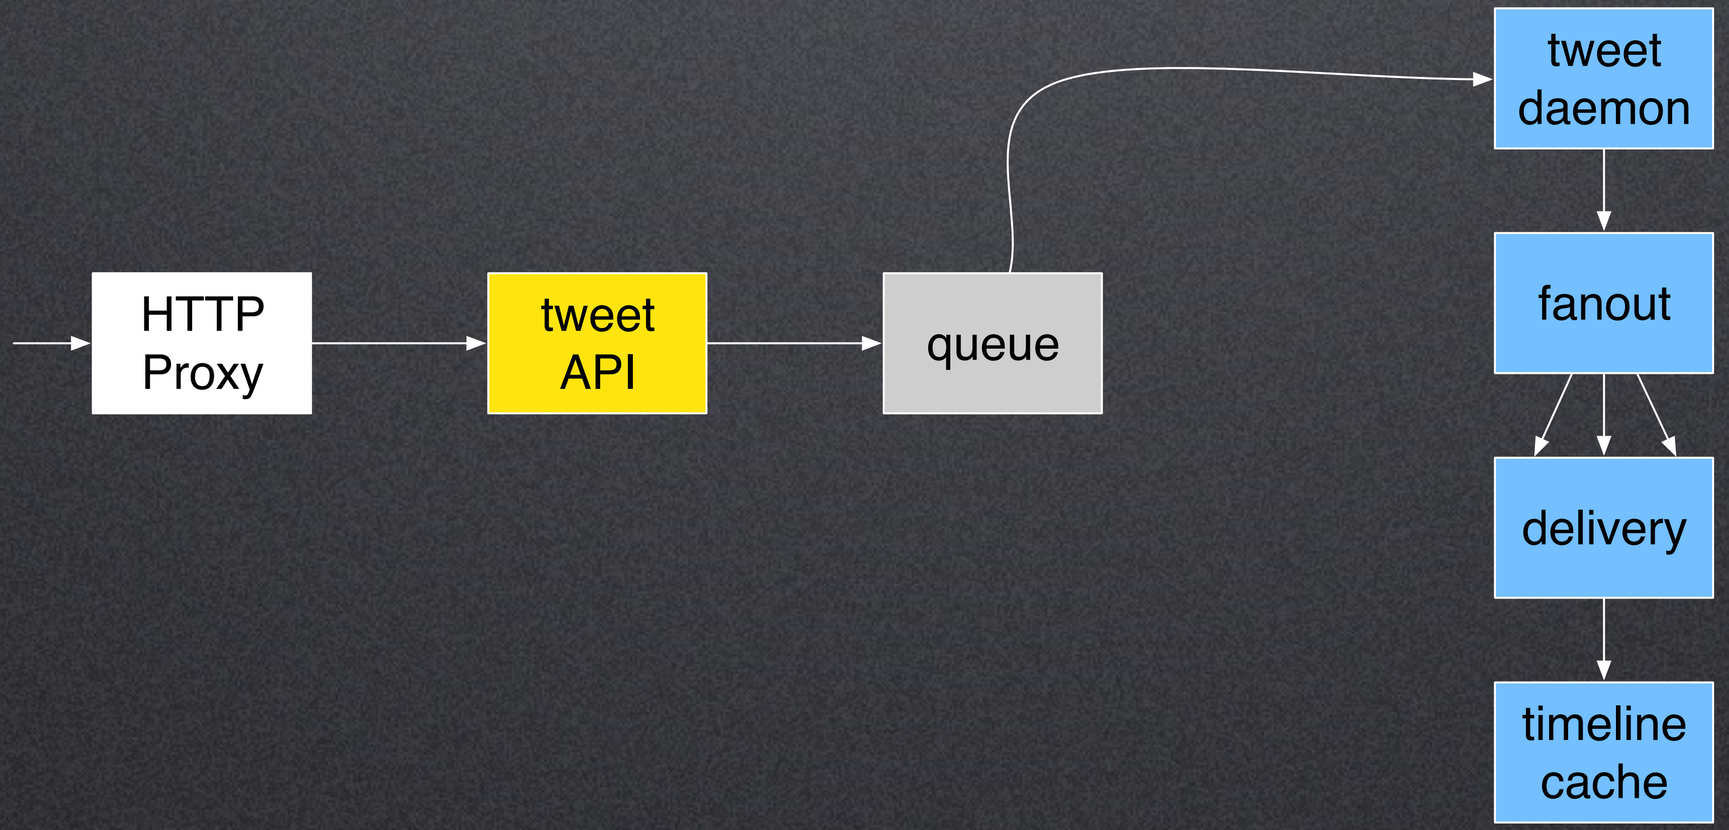
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
cache

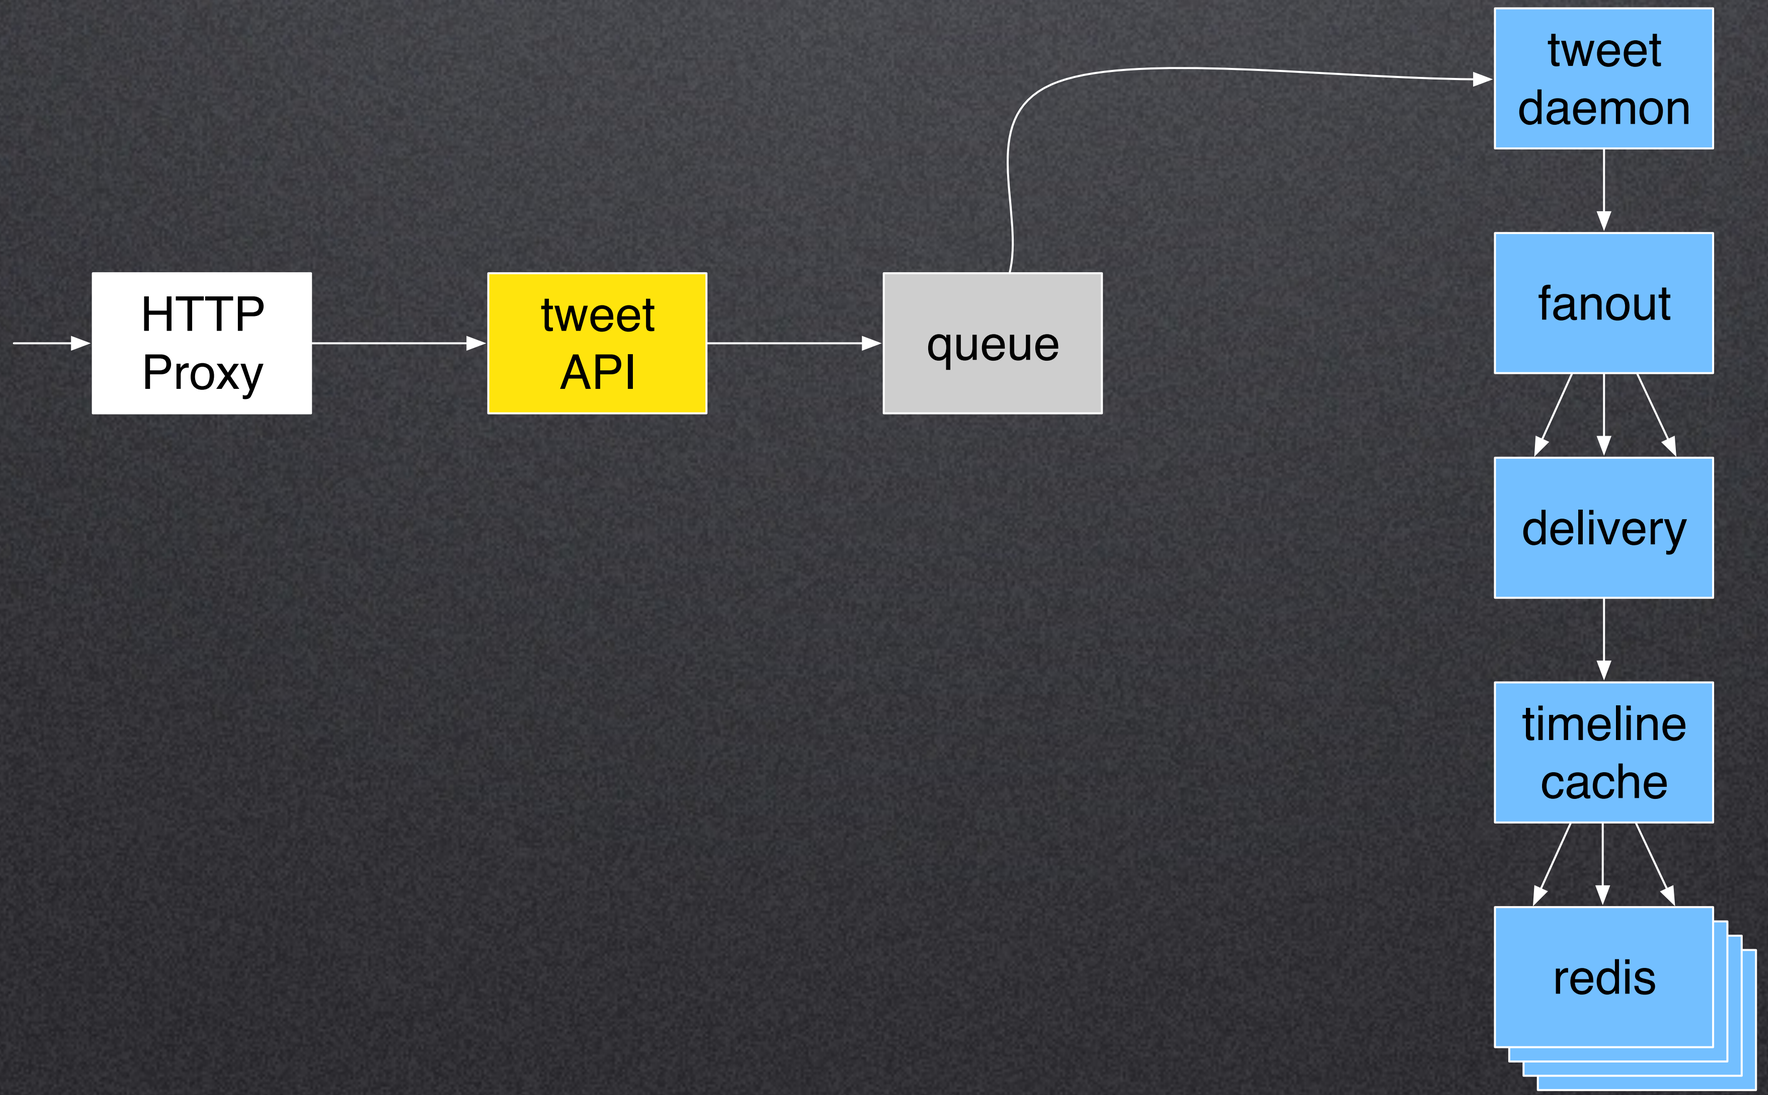
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
cache
redis

RPUSHX to timeline ID

Redis list structure:

tweet ID	user ID	bitfield
tweet ID	user ID	bitfield
tweet ID	user ID	bitfield
tweet ID	user ID	bitfield
tweet ID	user ID	bitfield
tweet ID	user ID	bitfield

Redis list structure:

tweet ID	user ID	bitfield	
tweet ID	user ID	bitfield	
tweet ID	user ID	bitfield	retweet ID
tweet ID	user ID	bitfield	
tweet ID	user ID	bitfield	retweet ID
tweet ID	user ID	bitfield	

Redis list structure:

tweet ID	user ID	bitfield	
tweet ID	user ID	bitfield	
tweet ID	user ID	bitfield	retweet ID
tweet ID	user ID	bitfield	
tweet ID	user ID	bitfield	retweet ID
tweet ID	user ID	bitfield	

retweet bit set

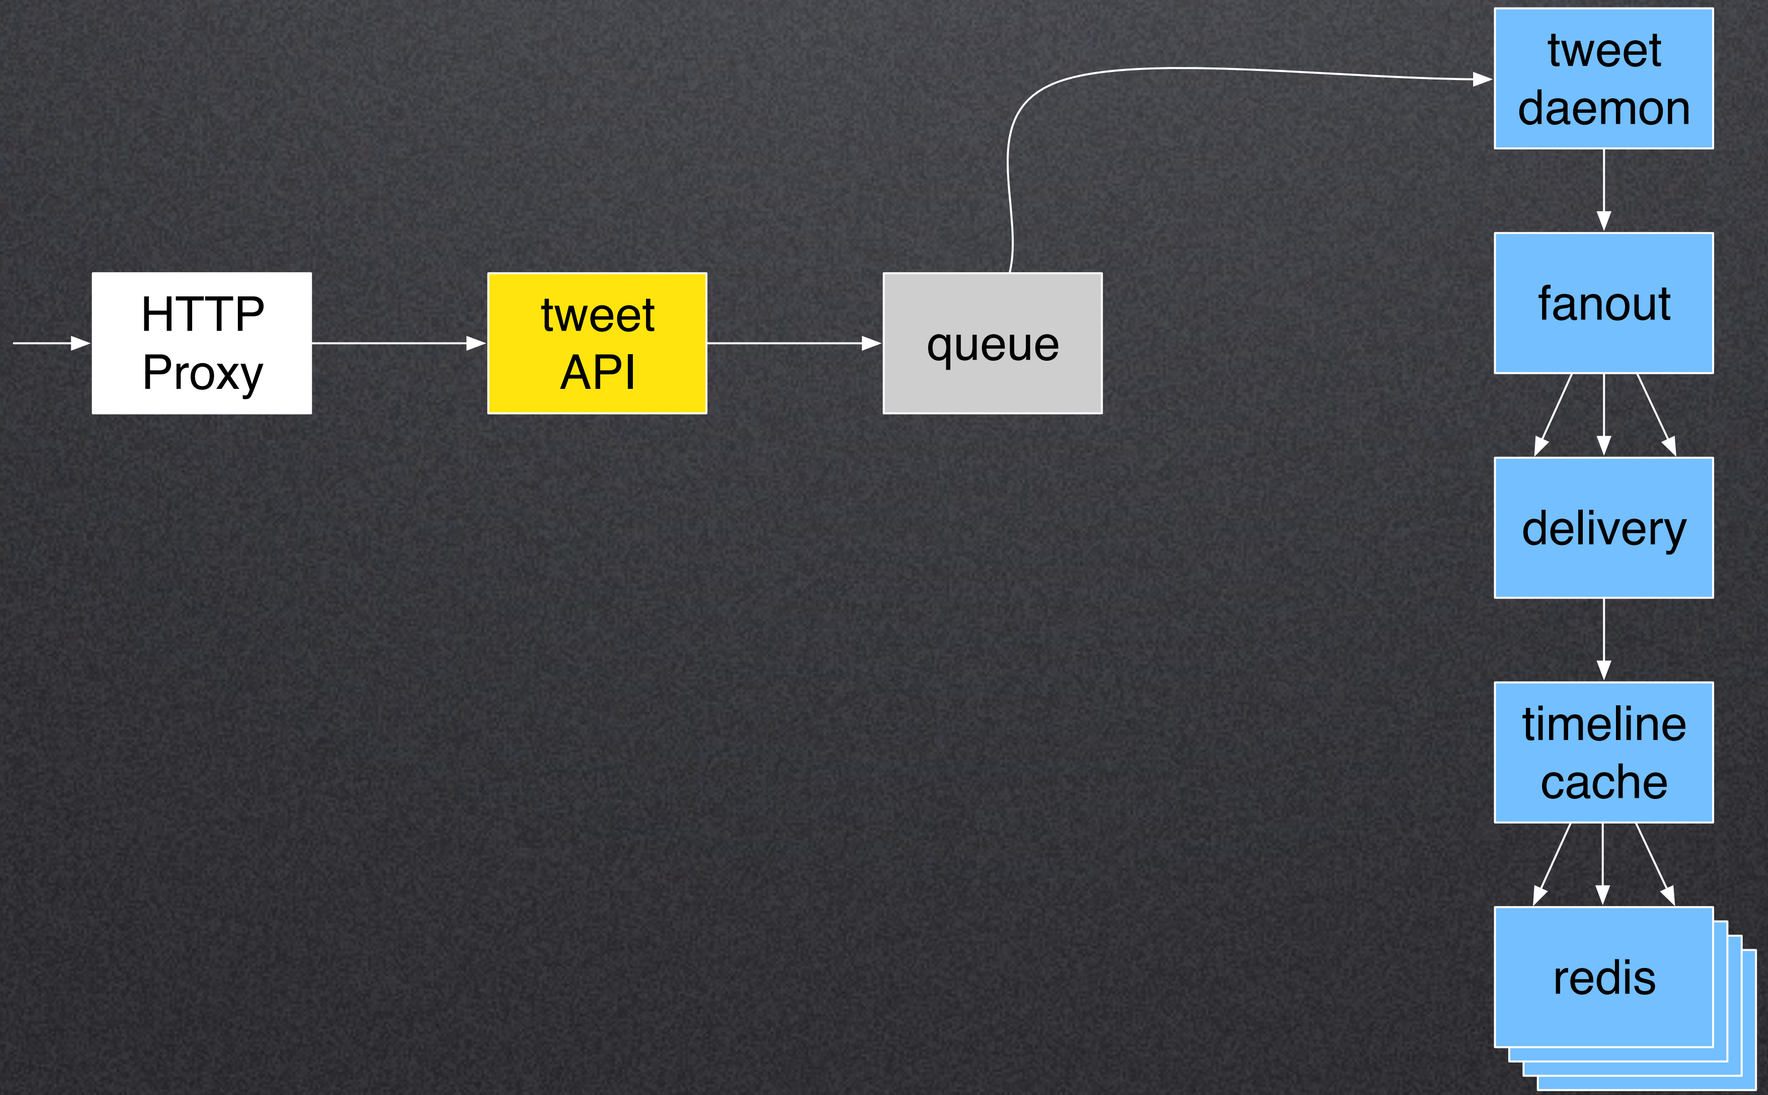
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
cache
redis

Poll-based Timelines

Tweets

Sam Pullara @sampullara 1m
@kevino it is mirroring -- i don't think there is a way to do as if it was a second monitor ... yet.
In reply to kevin oliver

news.yc Popular @newsycombinator 10m
Flowers regenerated from 30,000-year-old frozen fruit tissue j.mp/zPe7EH

The Onion @TheOnion 10m
SPORTSWIRE: Groans Abound As Tim Duncan Raises Hand Once Again At City Council Meeting onion.com/yuXGvC

news.yc Popular @newsycombinator 10m
How I built a Hacker News mobile web app j.mp/wUvRy5

Maria Popova @brainpicker 14m
YES "What we need is seriously engaged art that can teach us again that we're smart." David Foster Wallace on art vs TV j.mp/wVHltY

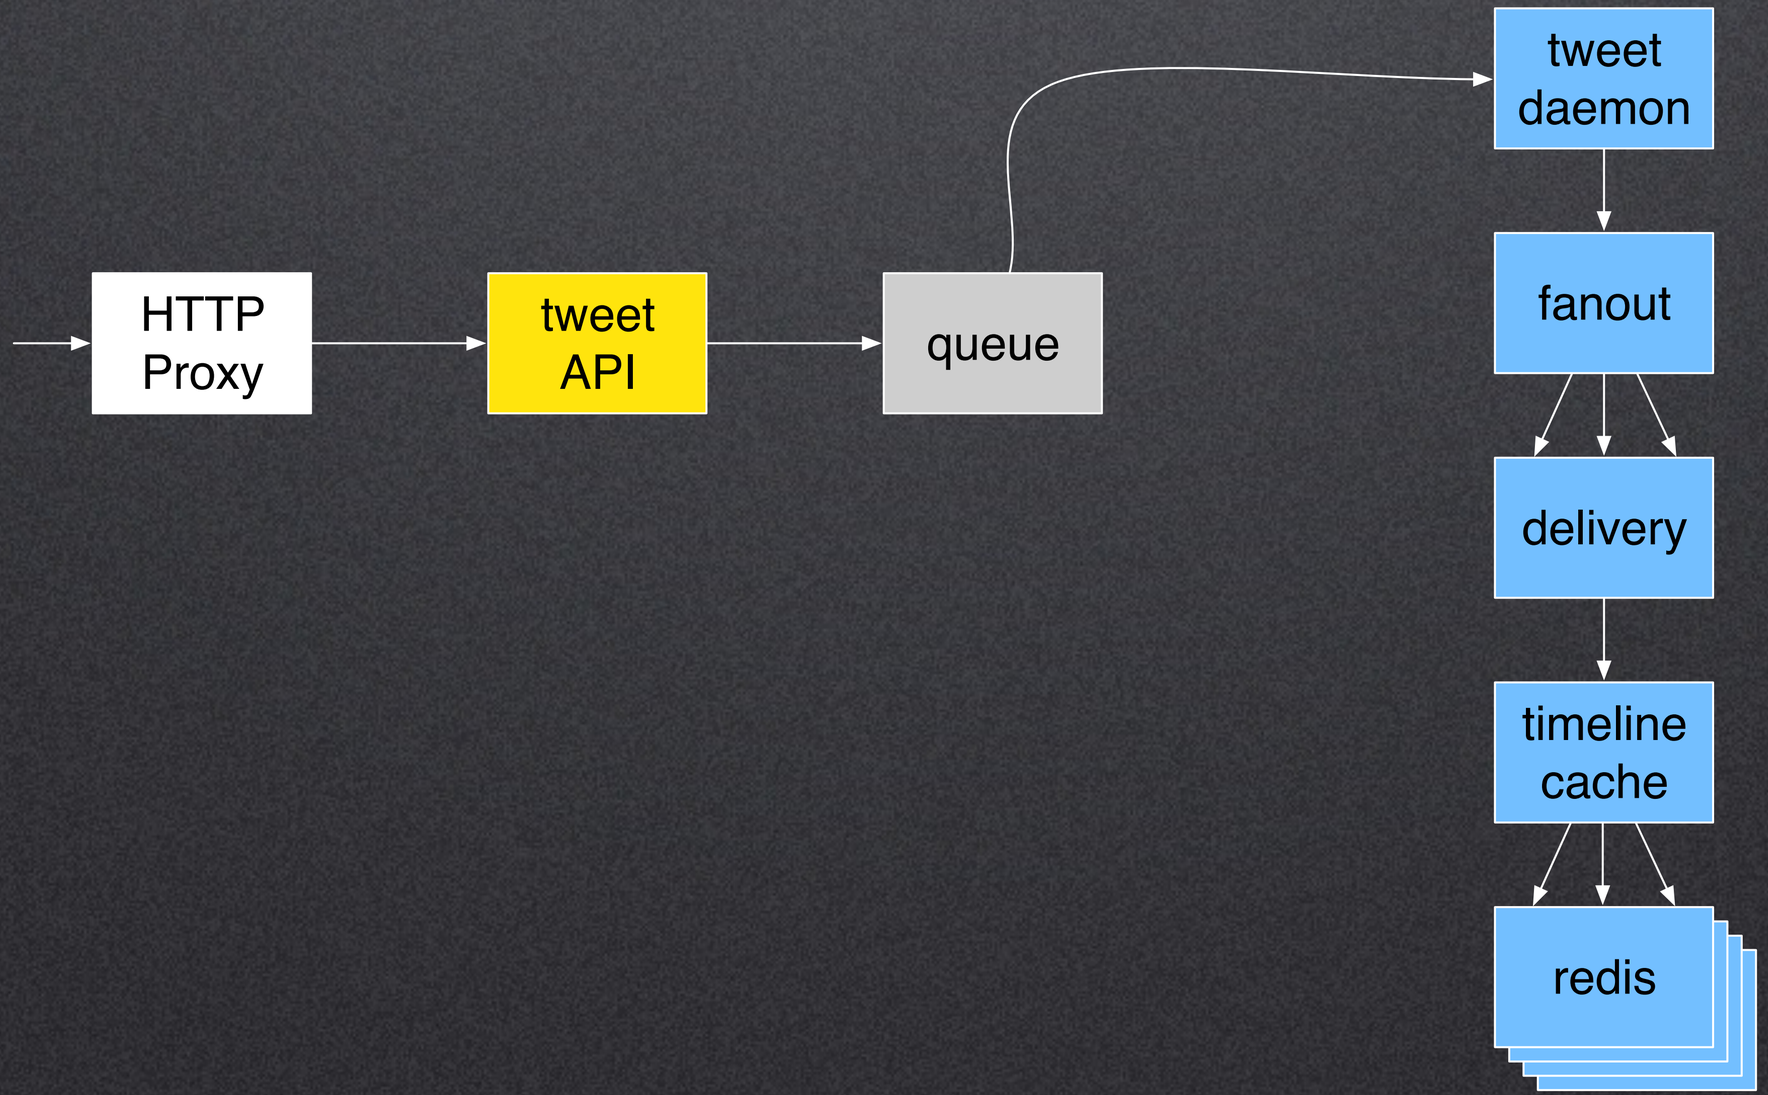
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
cache
redis

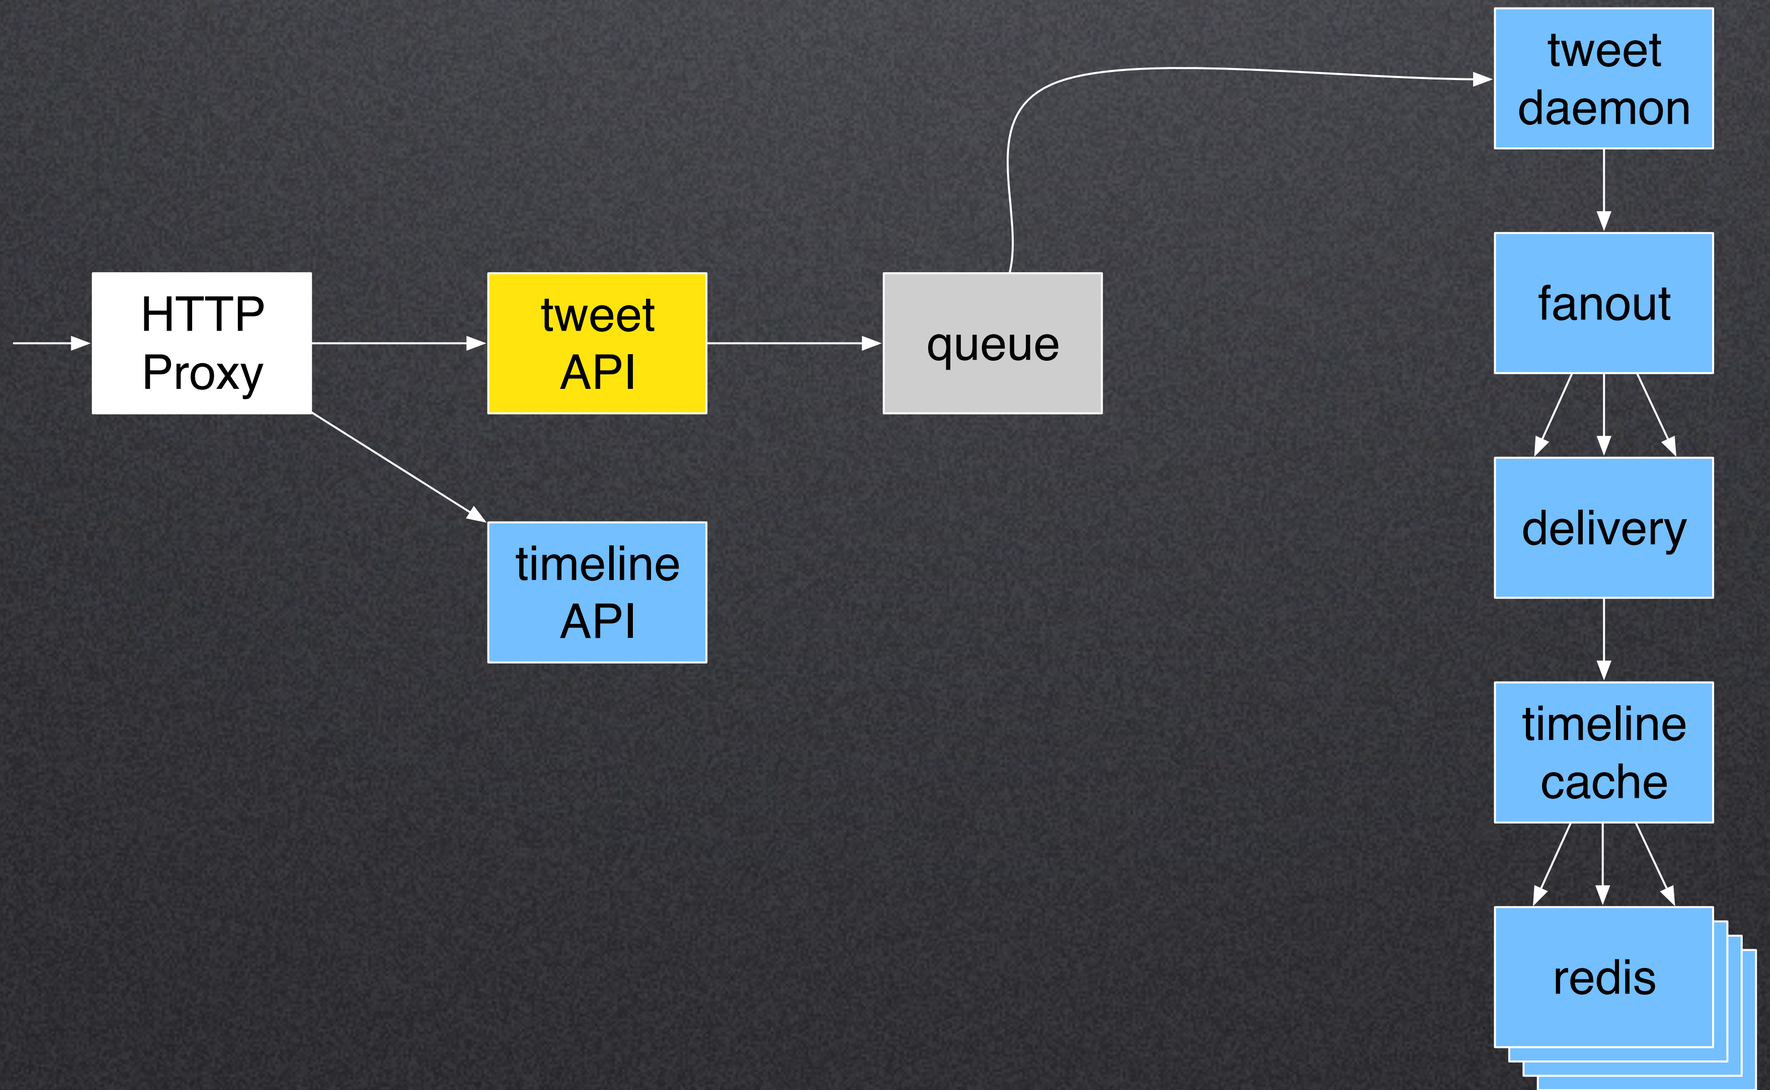
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
cache
redis
timeline
API

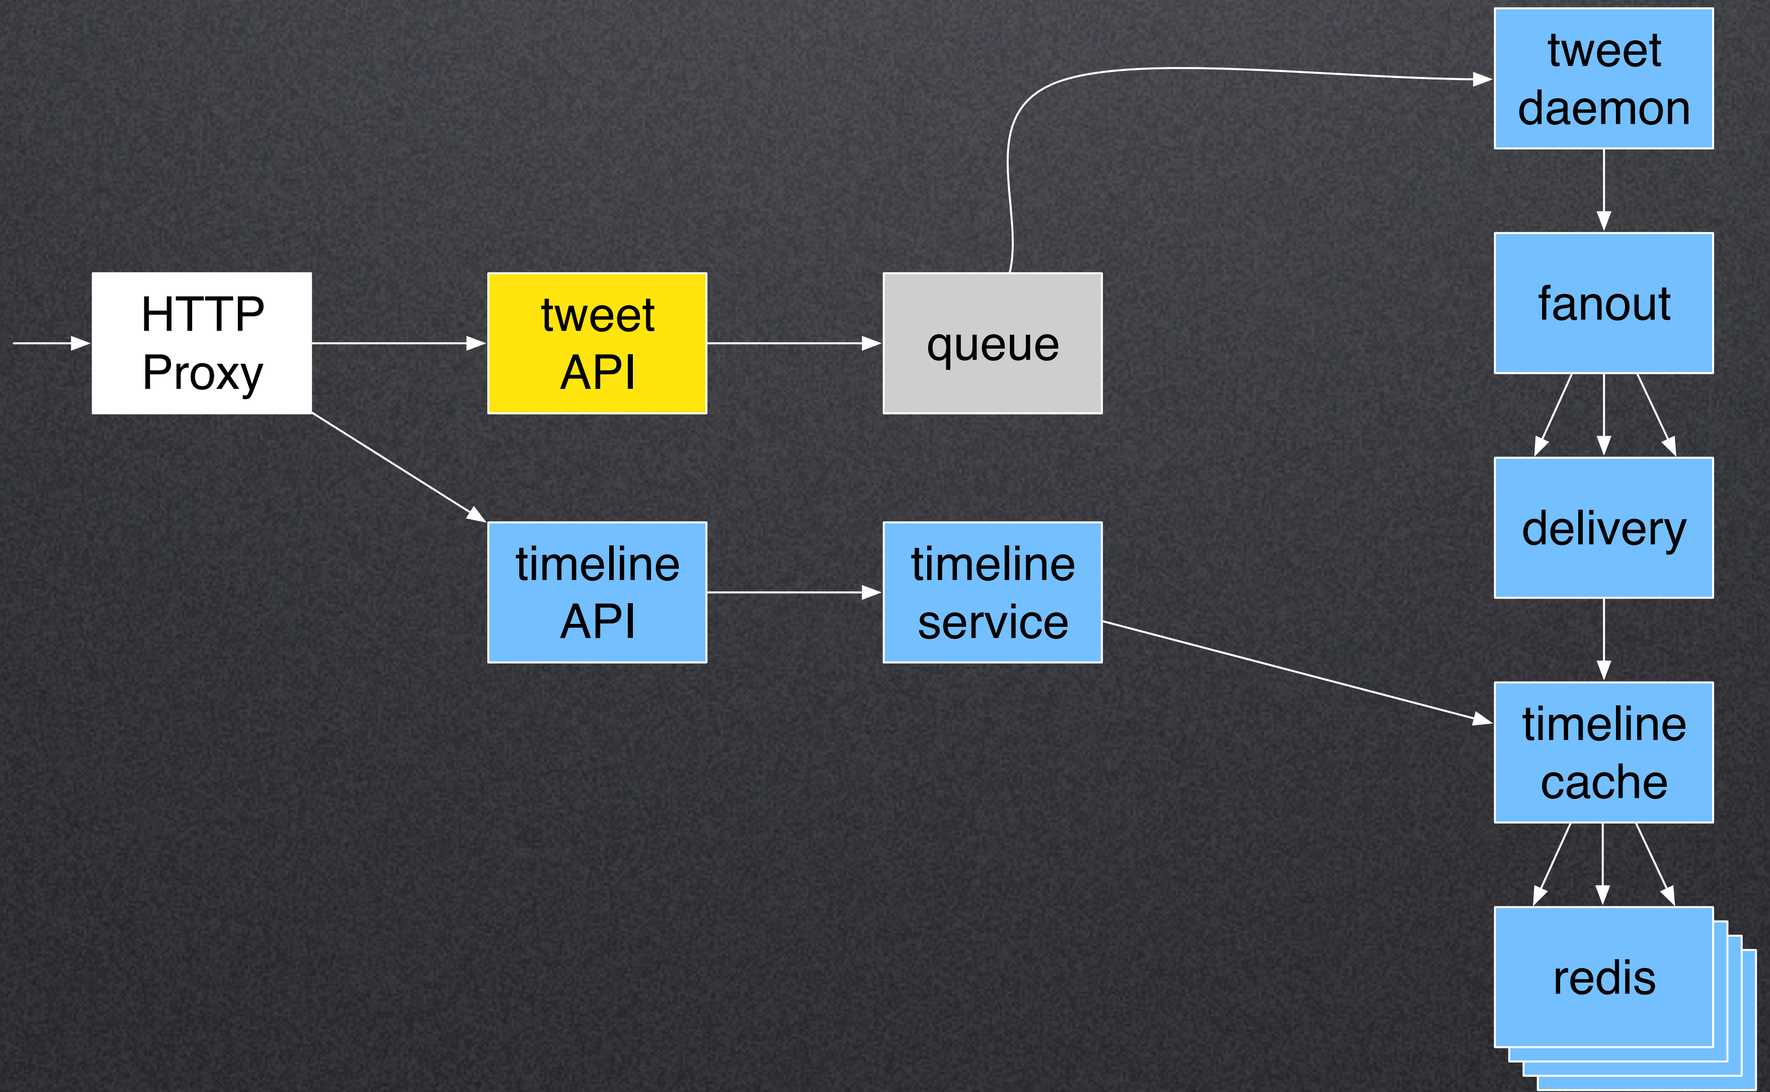
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
API
timeline
service
timeline
cache
redis

Loading your home timeline

three pieces of data

- timeline
- tweets
- users

Loading your home timeline

three pieces of data

- timeline
- tweets
- users

Loading your home timeline

three pieces of data

- timeline
- tweets
- users

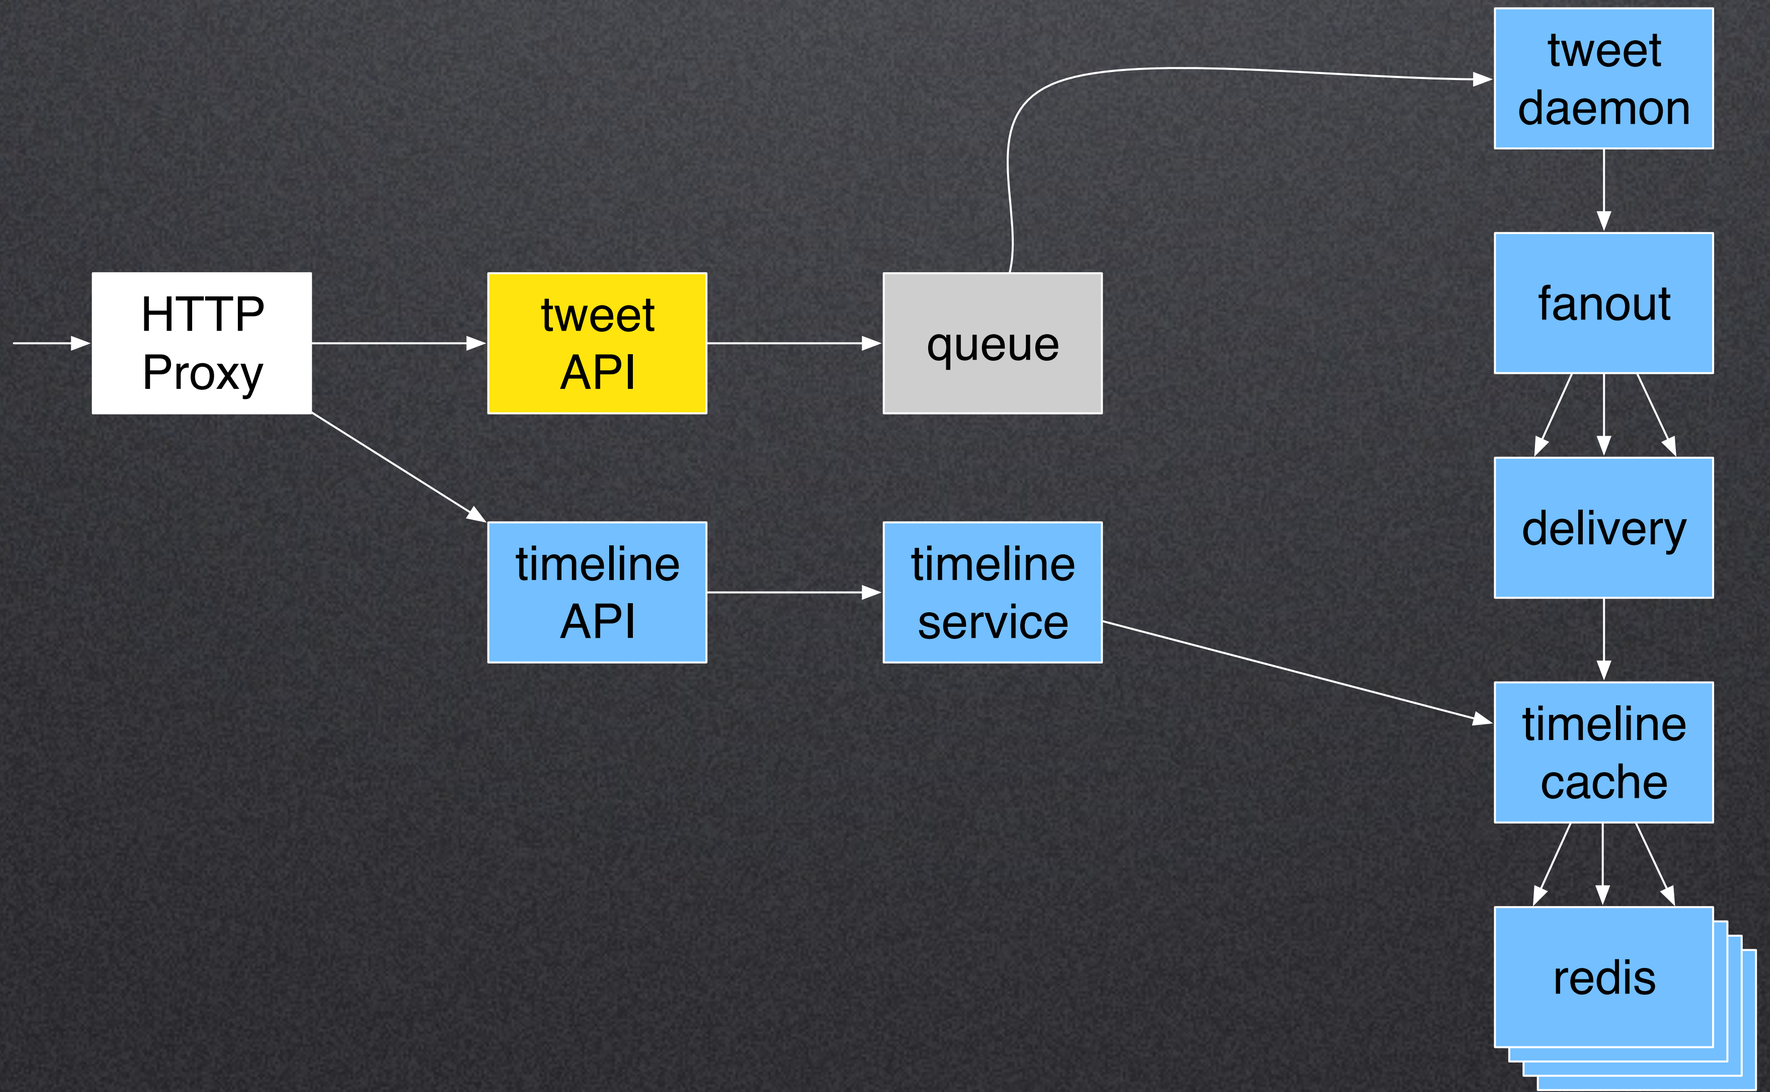
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
API
timeline
service
timeline
cache
redis

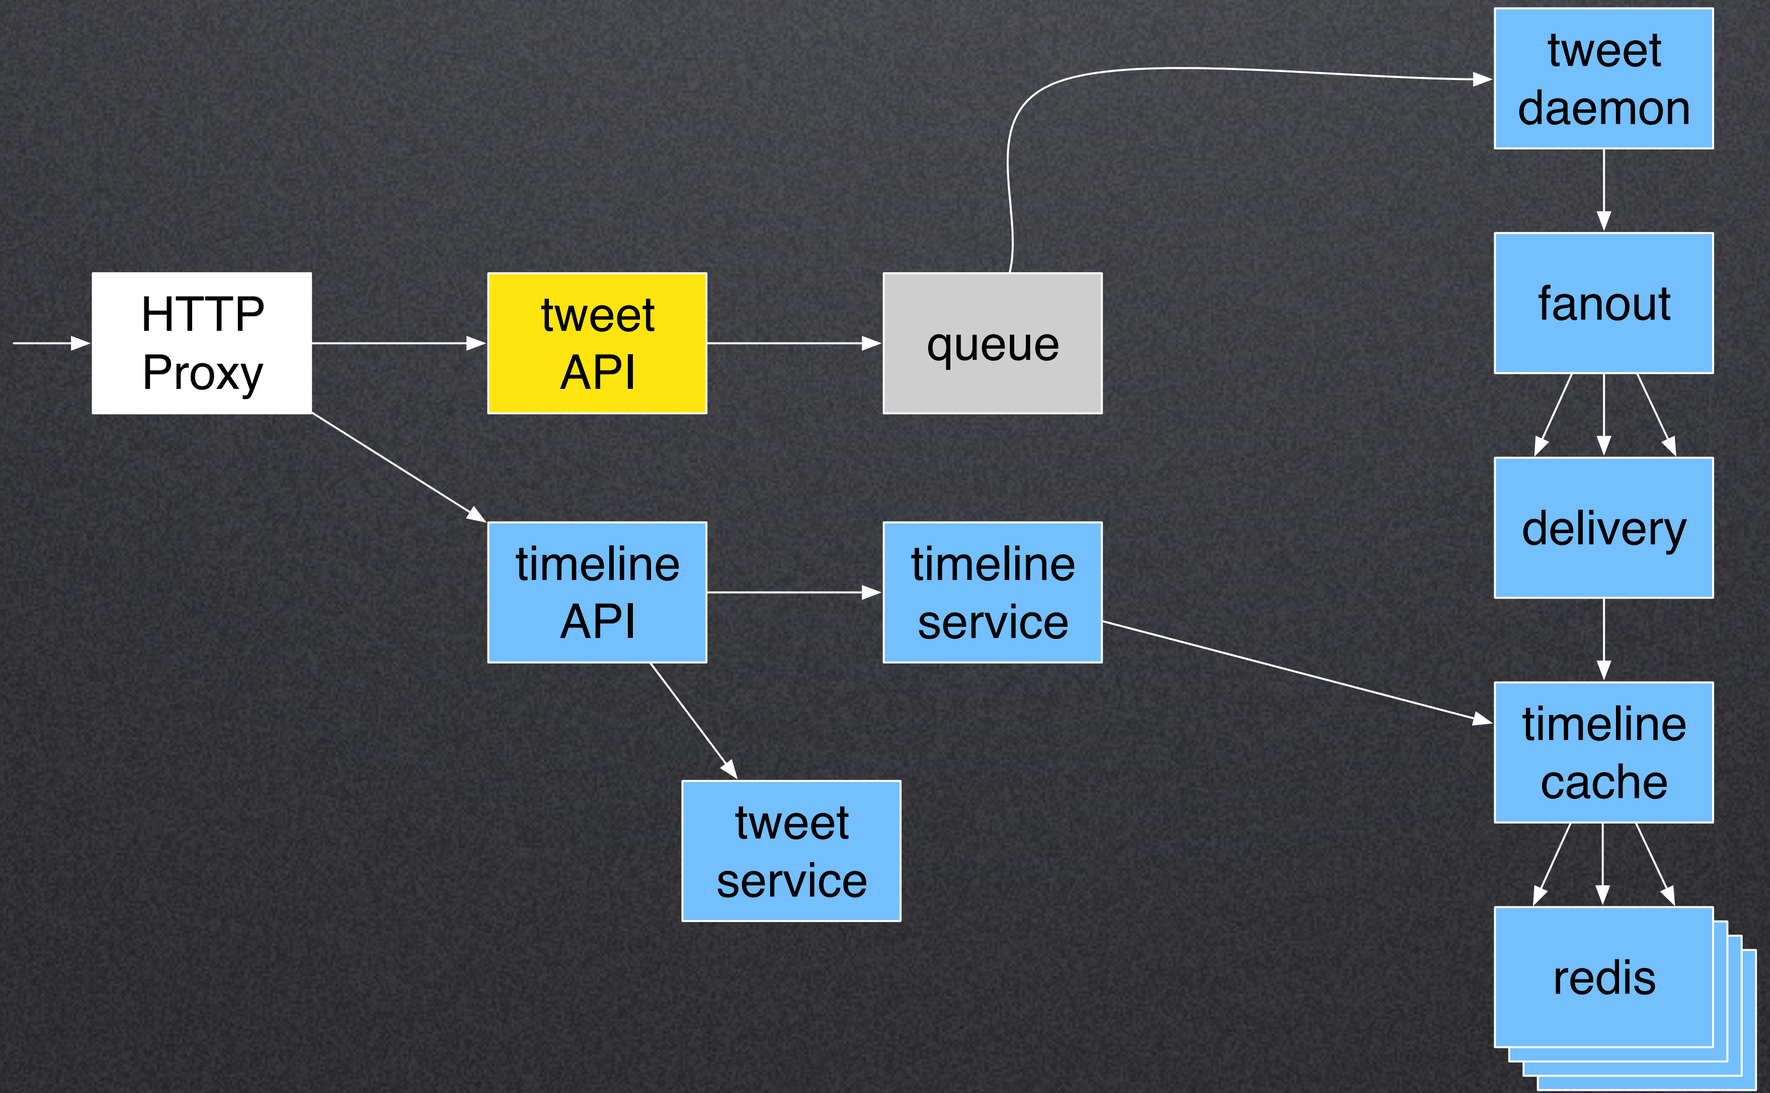
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
API
timeline
service
timeline
cache
tweet
service
redis

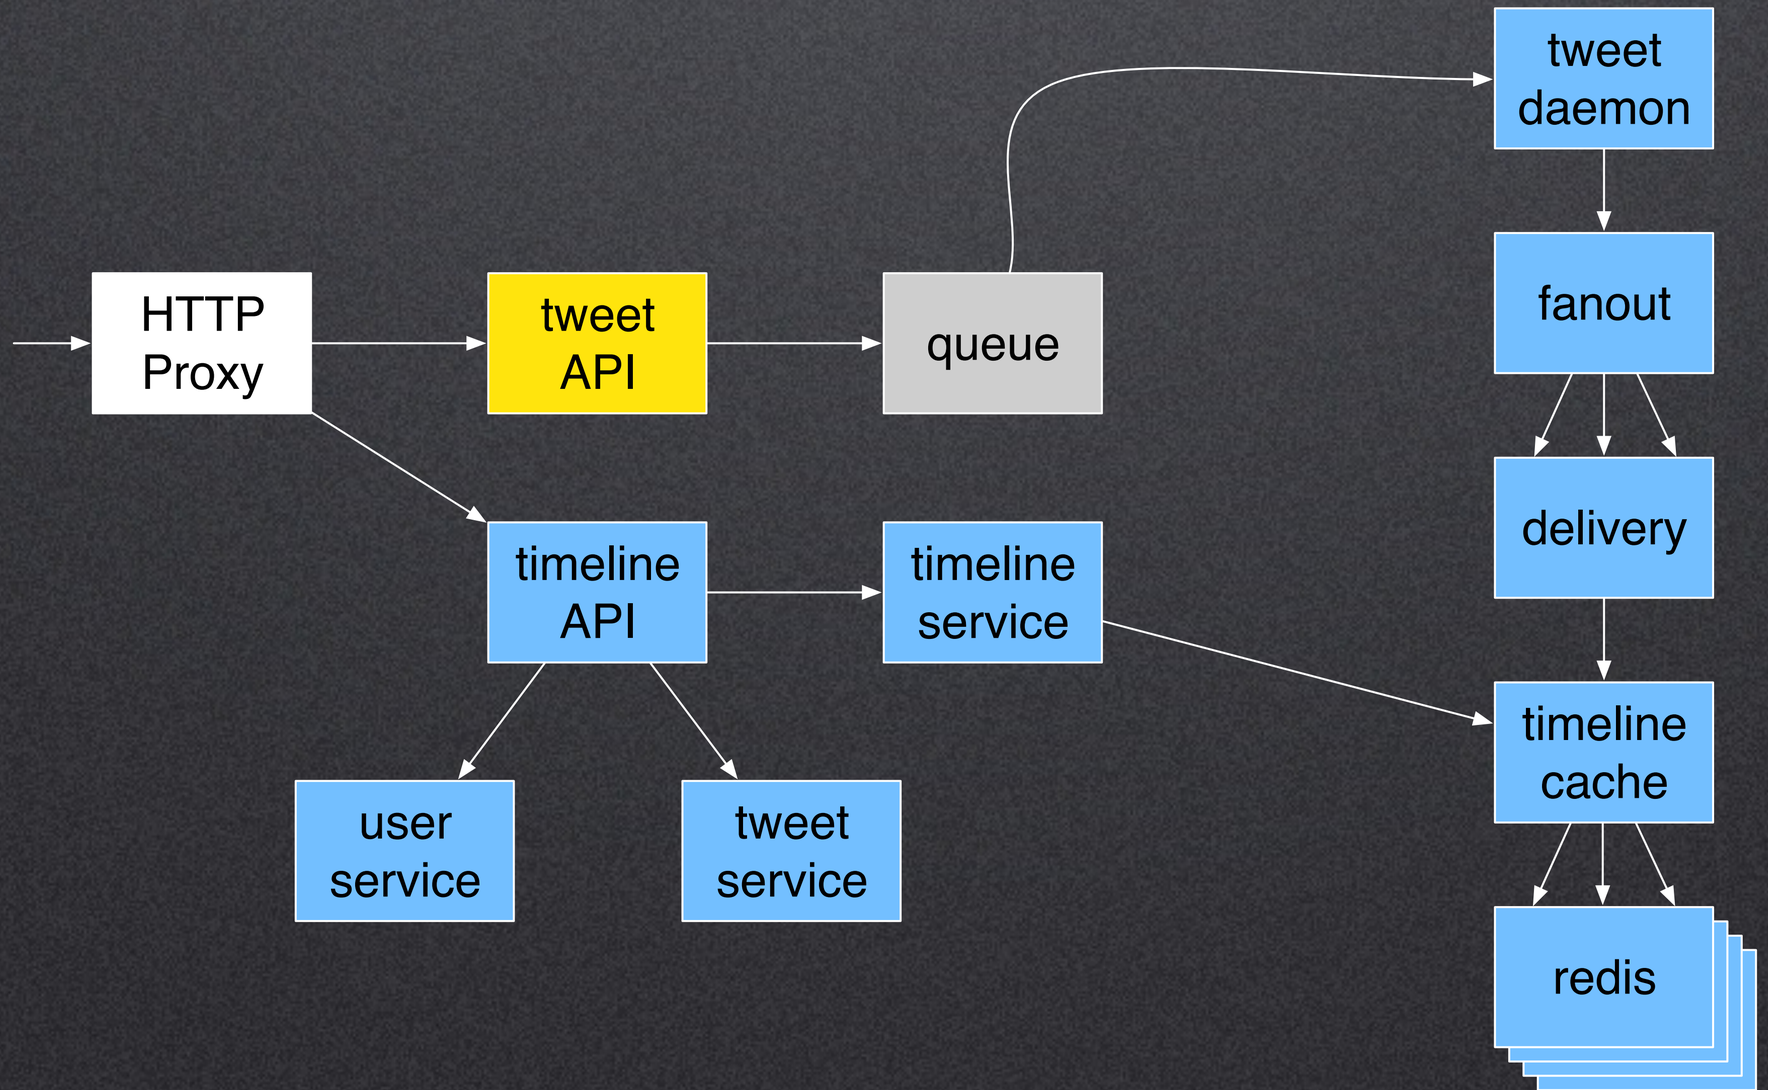
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
API
timeline
service
timeline
cache
user
service
tweet
service
redis

Hot Spots

users

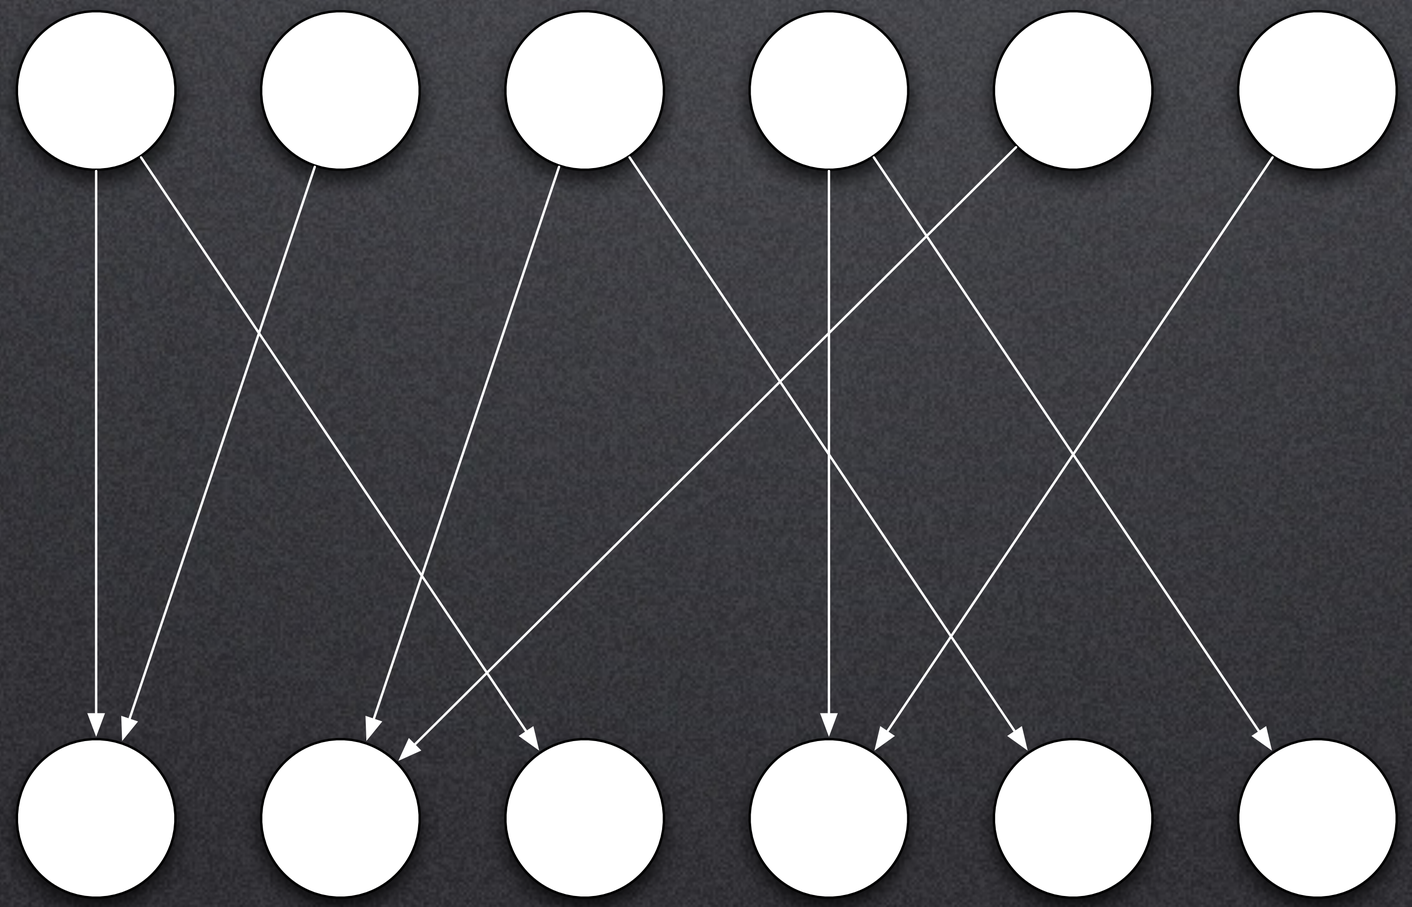

timelines

users

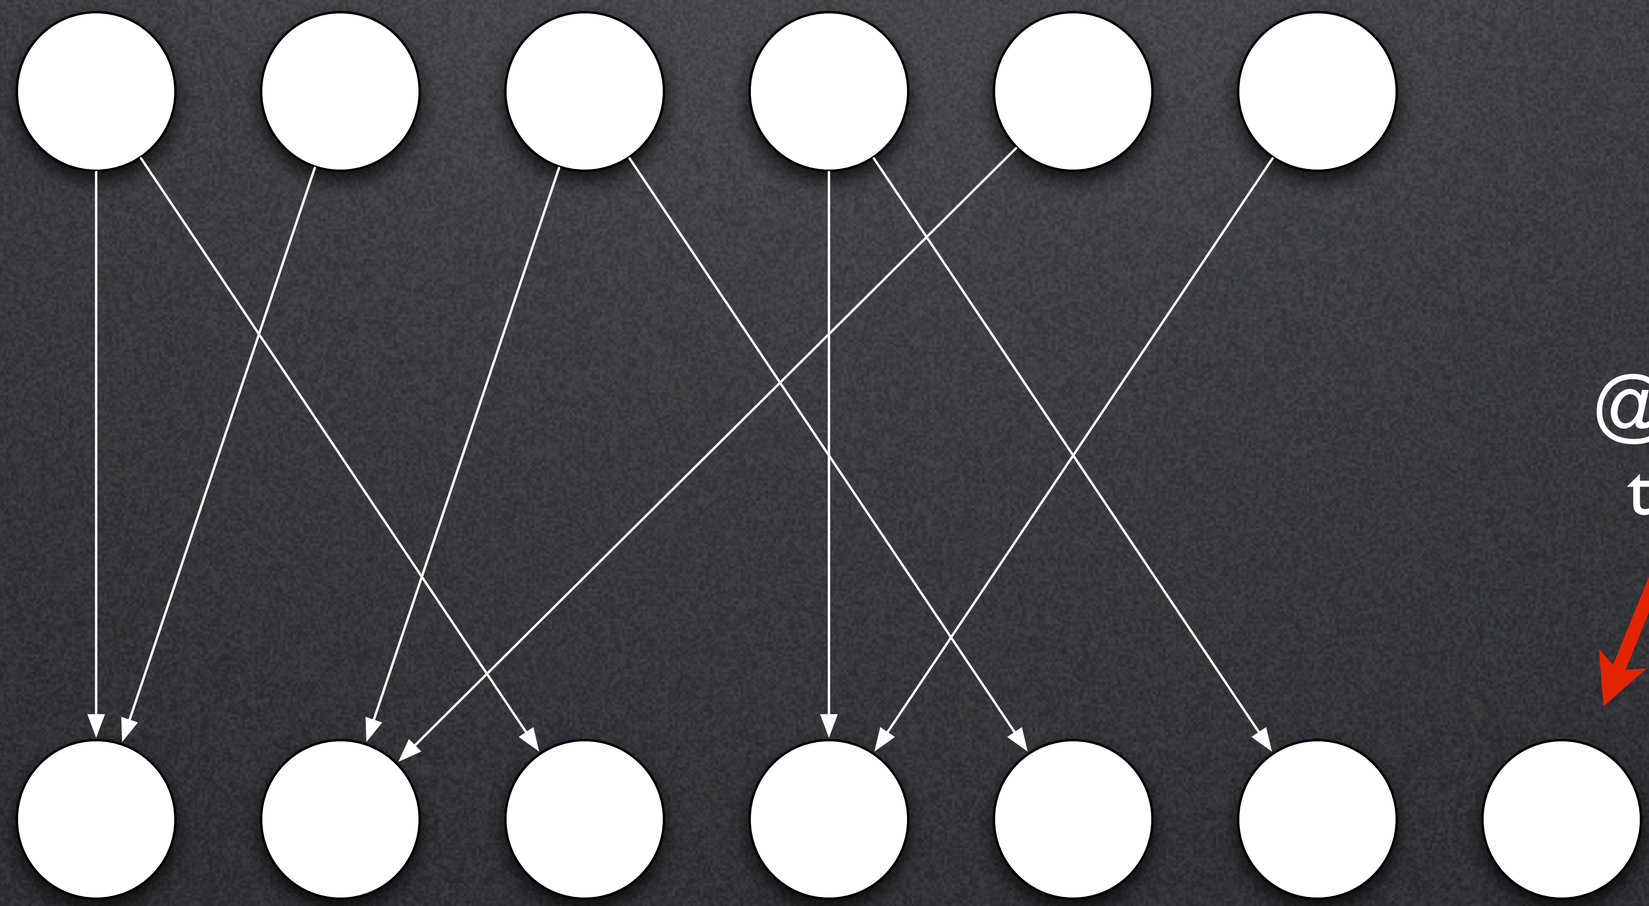

@diarioas
timeline

timelines

users

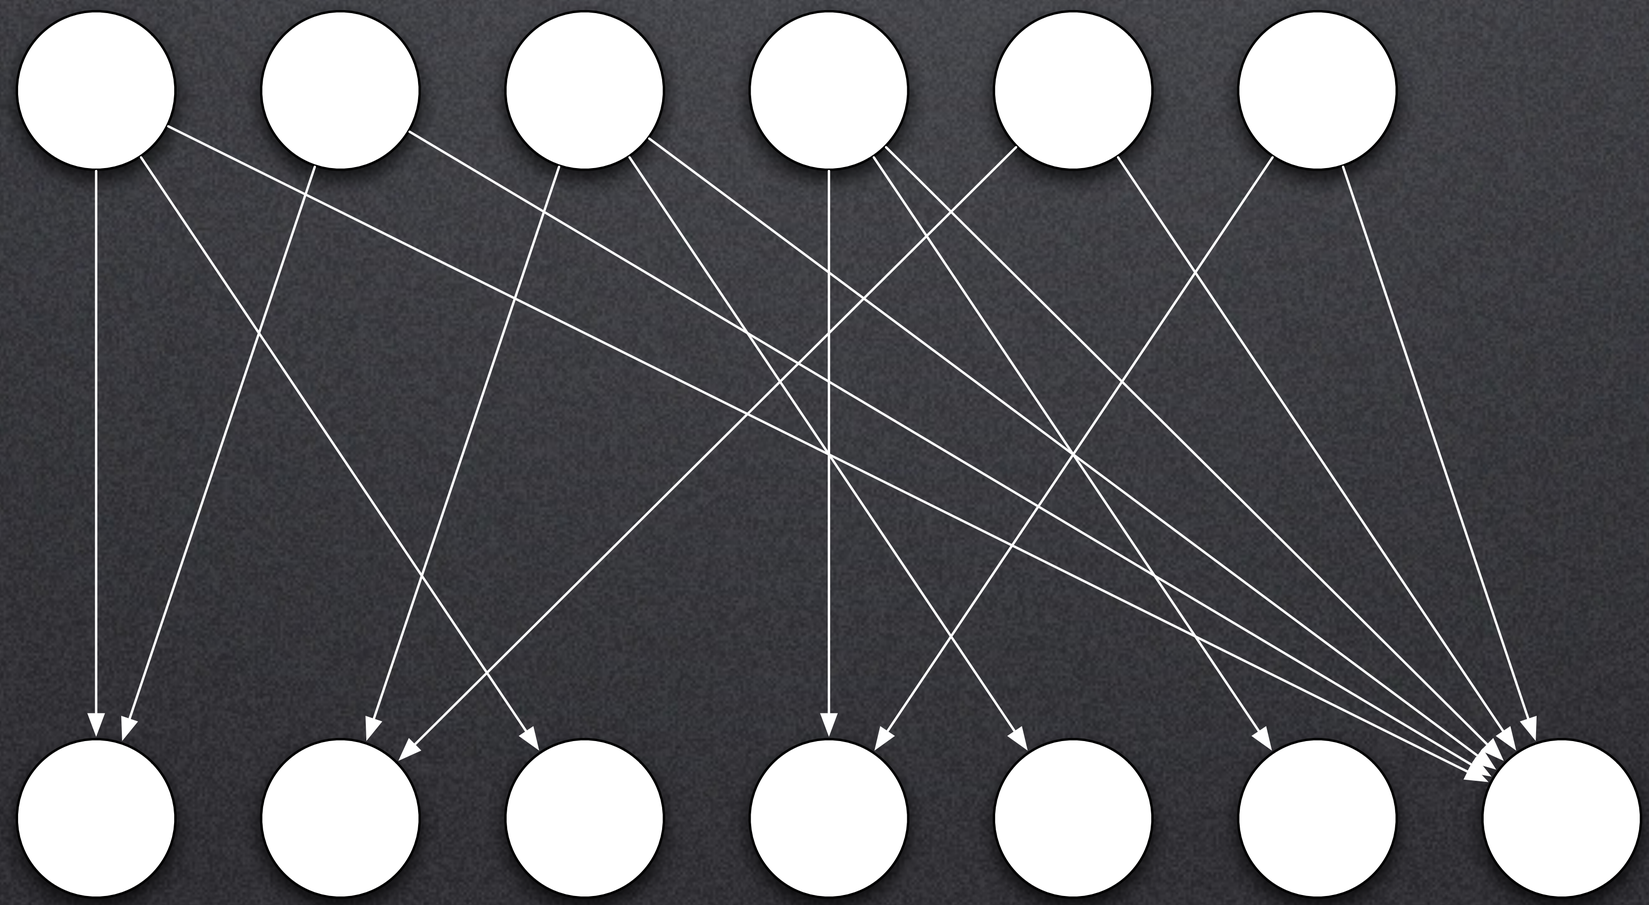

timelines

users

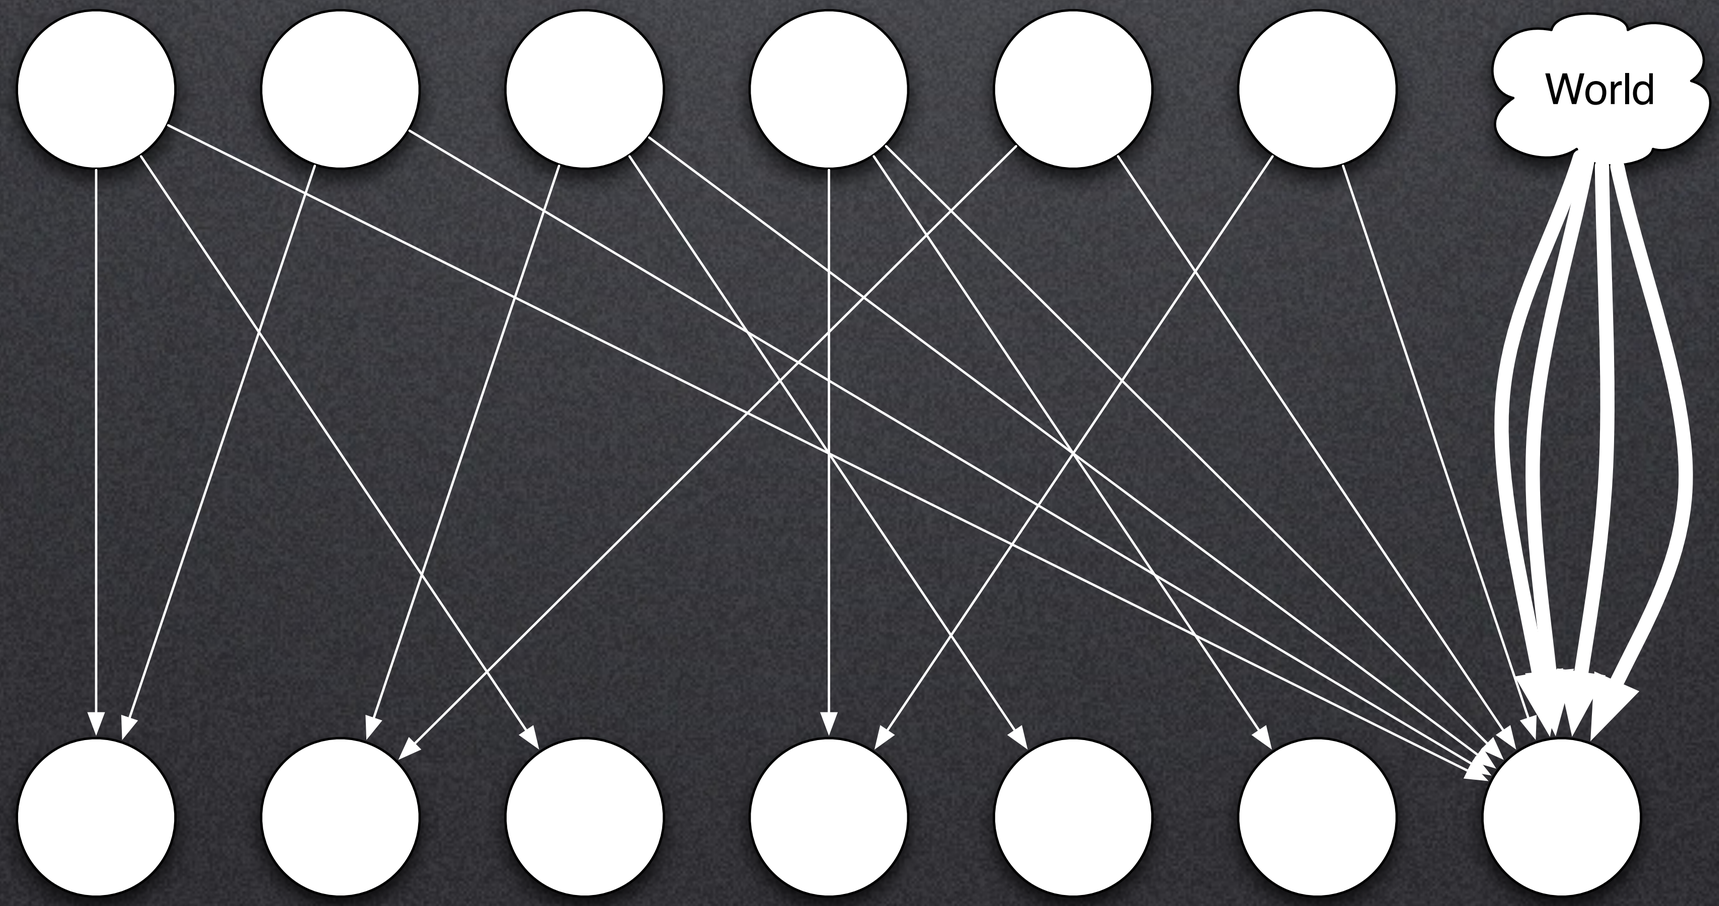

timelines

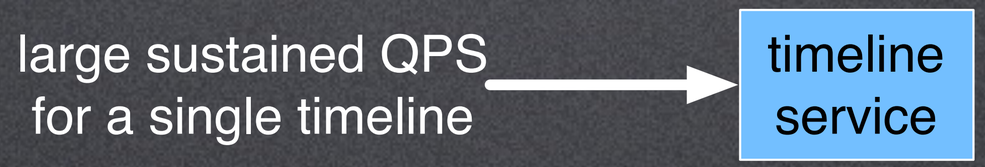
large sustained QPS
for a single timeline
timeline
service

timeline service

timeline service

timeline service

timeline service

timeline service

timeline service

timeline service

timeline service

timeline cache

timeline service

timeline service

timeline service

timeline service

timeline cache

timeline cache

timeline cache

timeline cache

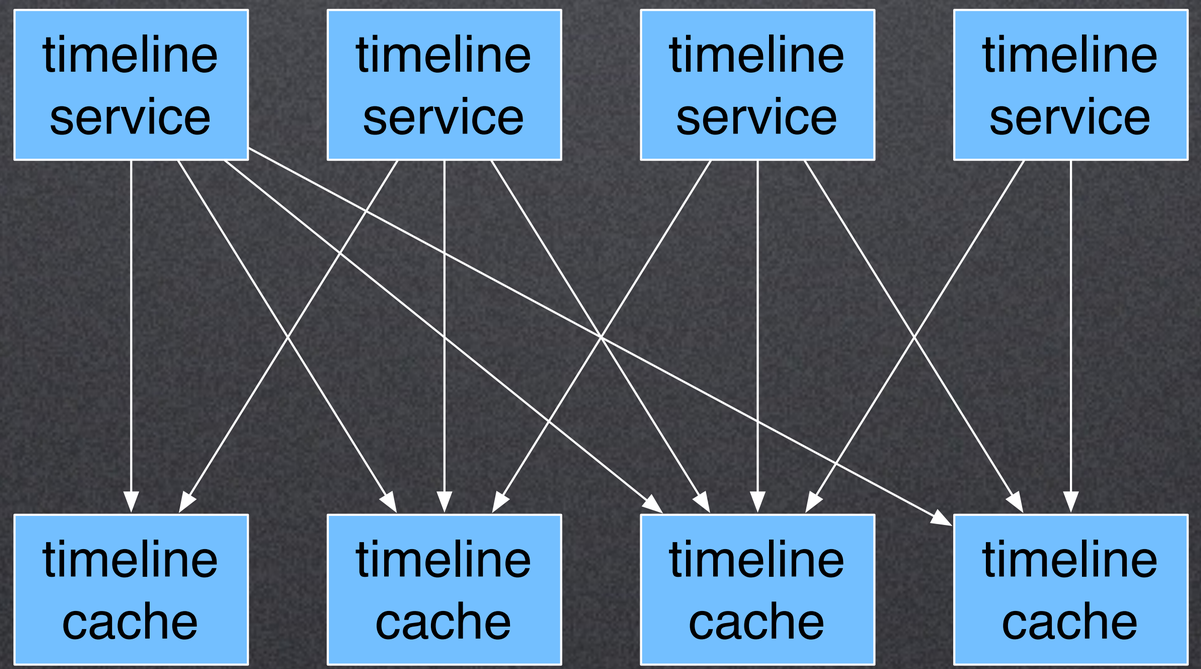
timeline service
timeline service
timeline service
timeline service
timeline cache
timeline cache
timeline cache
timeline cache

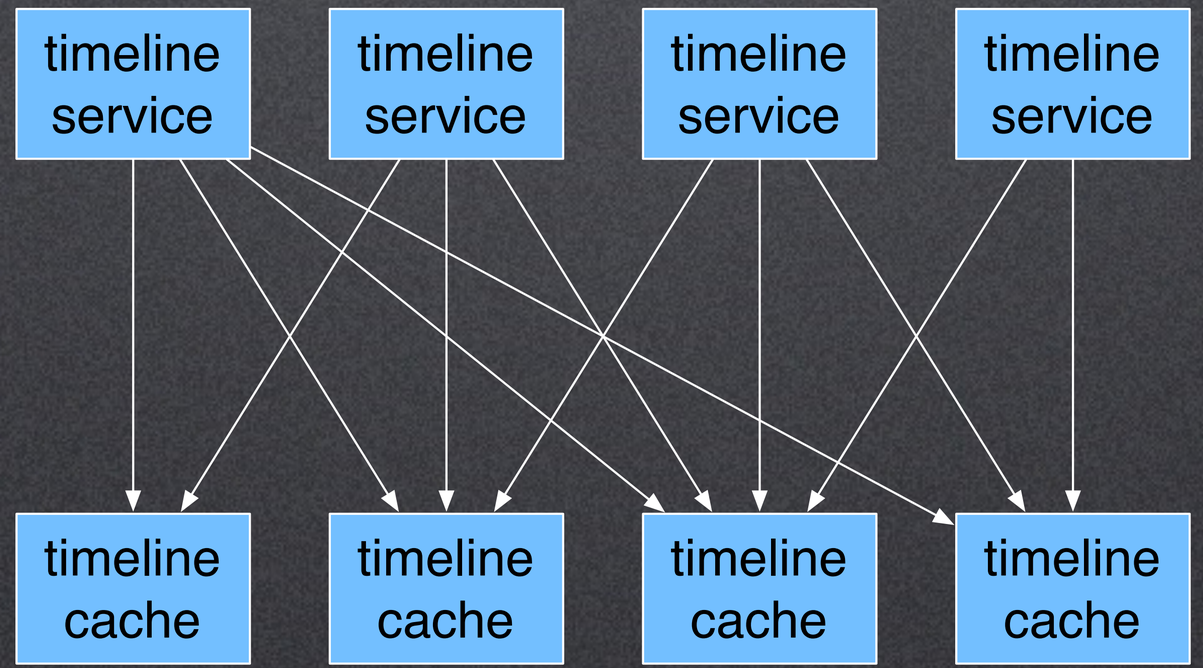

partitioning layer

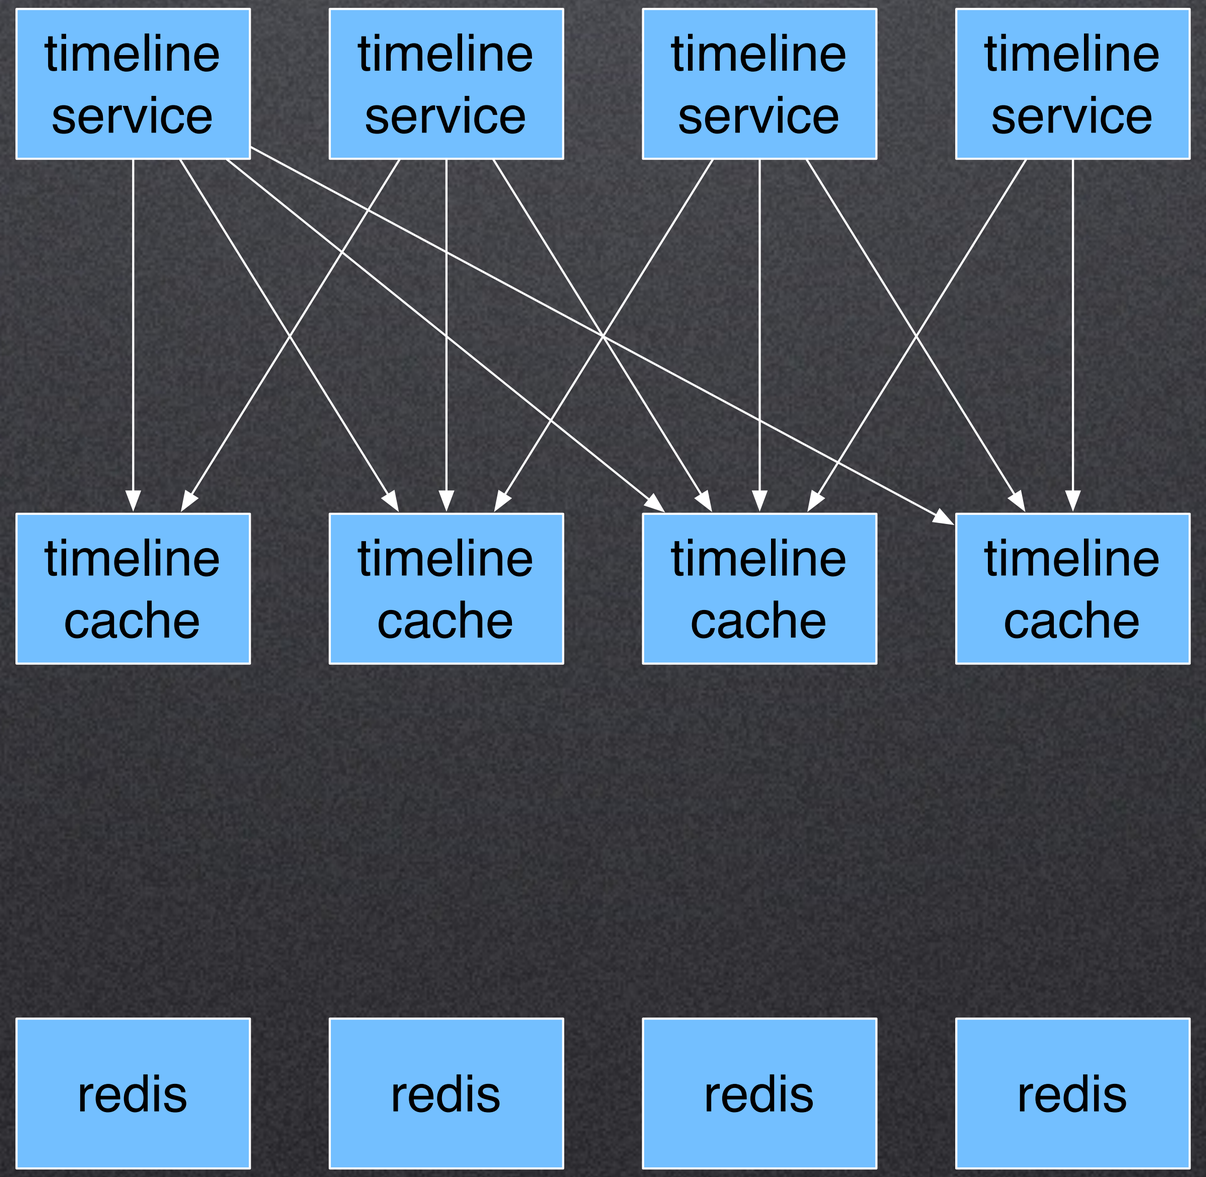

partitioning layer

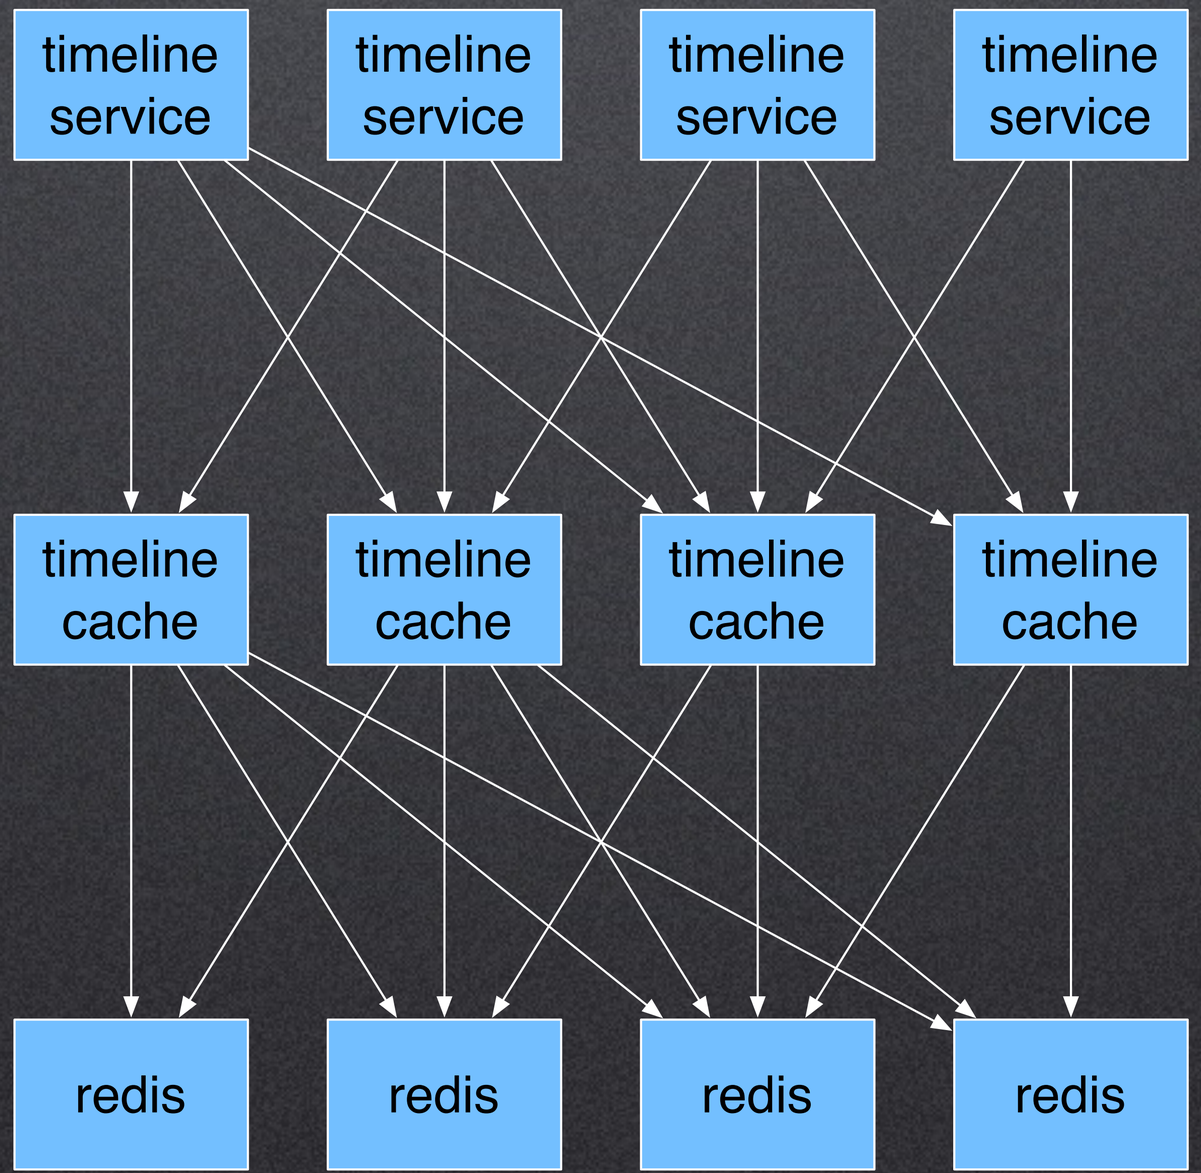
timeline service
timeline service
timeline service
timeline service
timeline cache
timeline cache
timeline cache
timeline cache
redis
redis
redis
redis

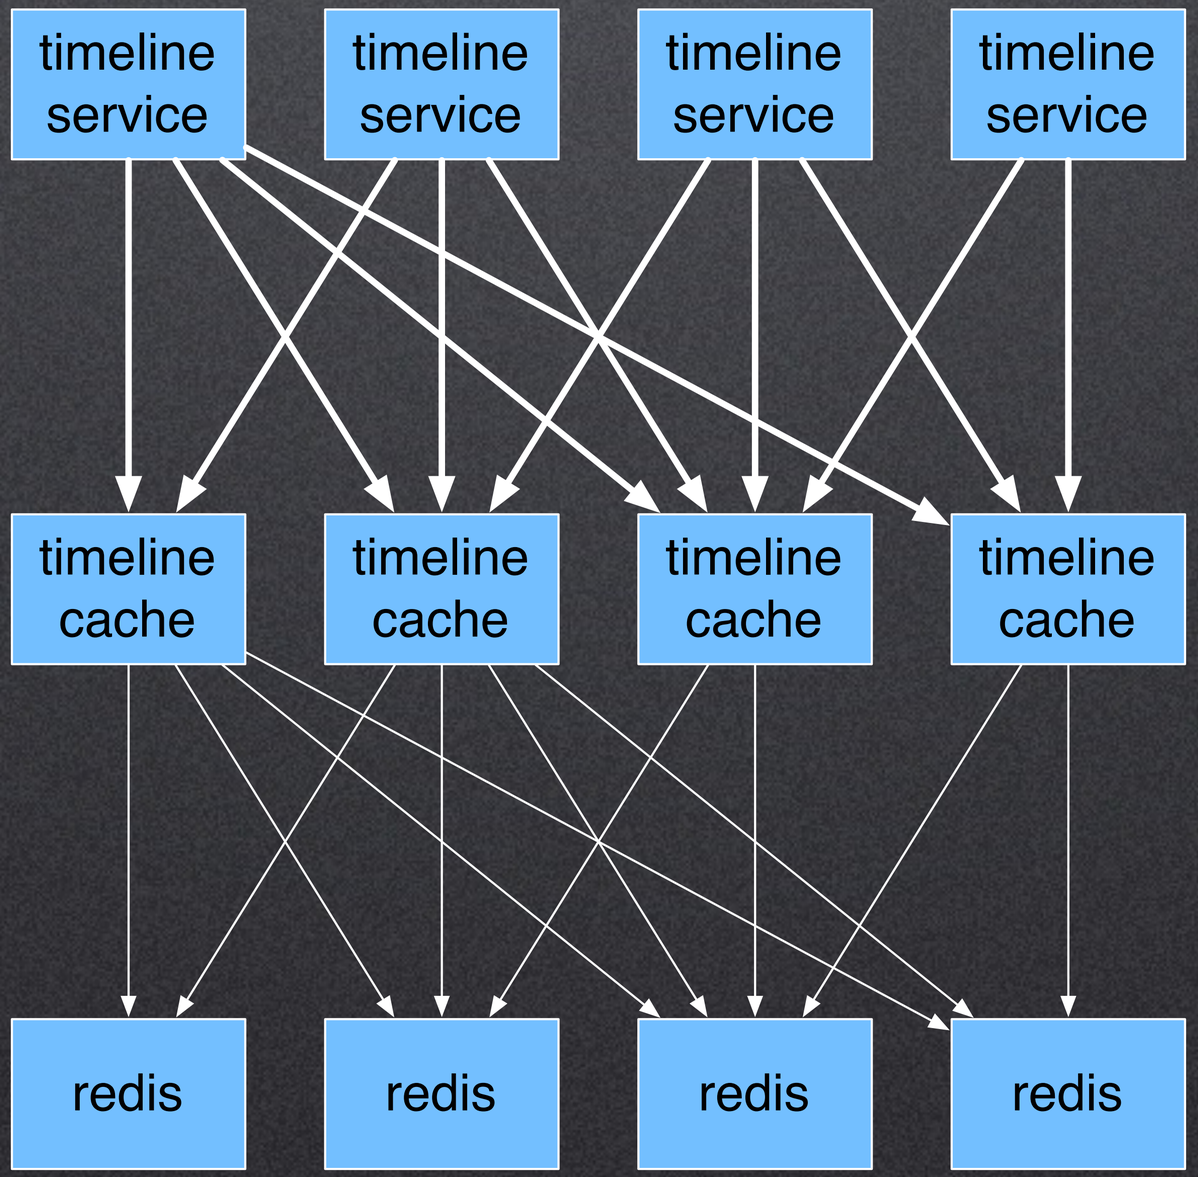
timeline service
timeline service
timeline service
timeline service
timeline cache
timeline cache
timeline cache
timeline cache
redis
redis
redis
redis

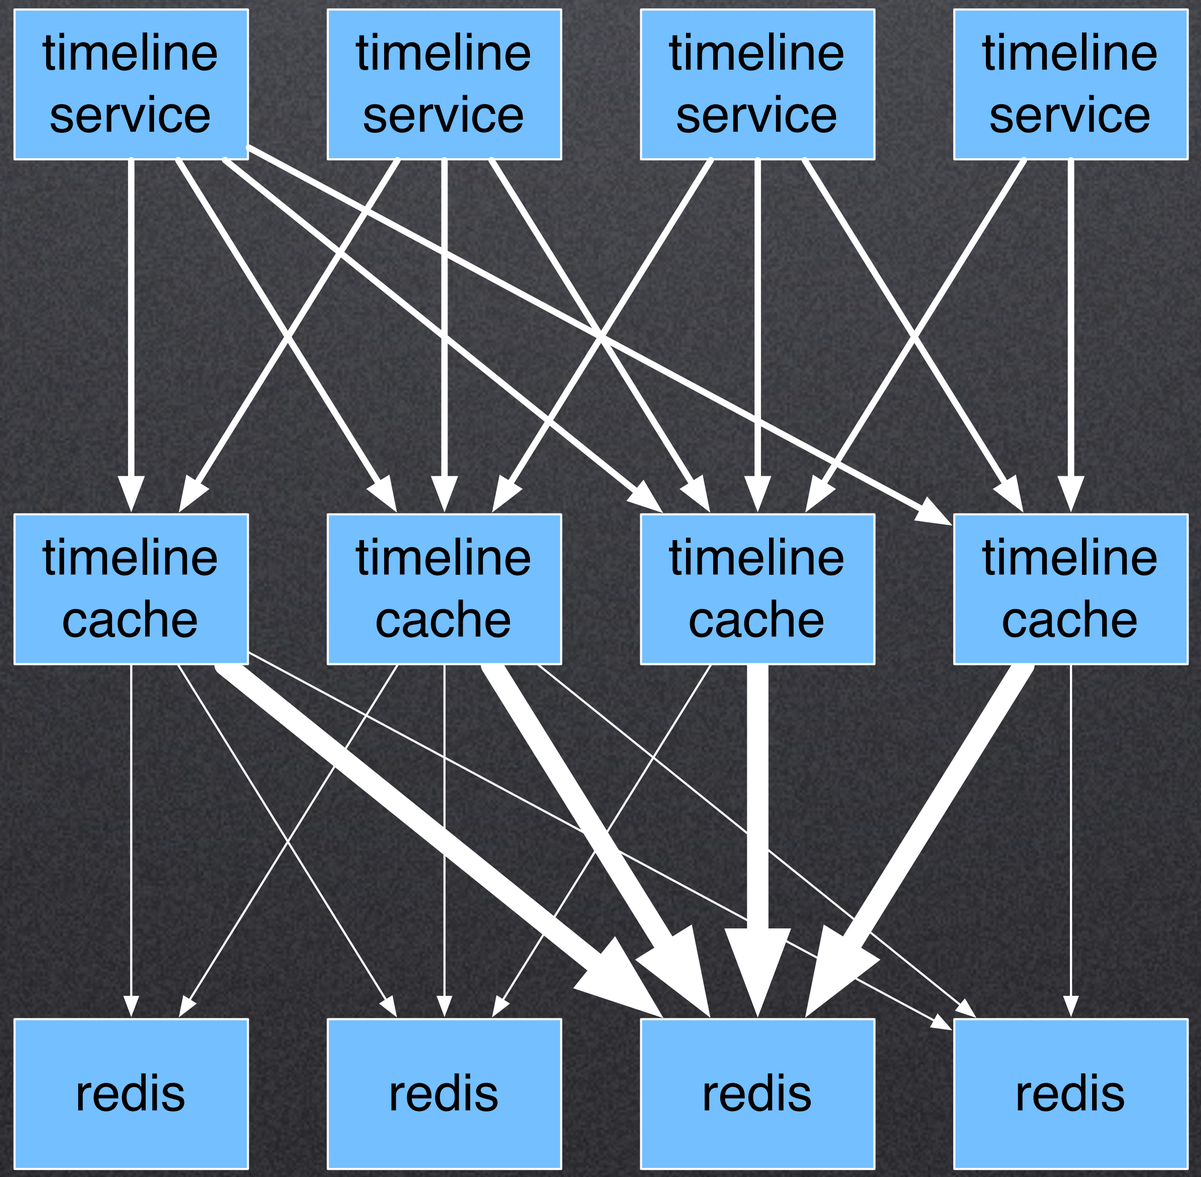
timeline service
timeline service
timeline service
timeline service
timeline cache
timeline cache
timeline cache
timeline cache
redis
redis
redis
redis

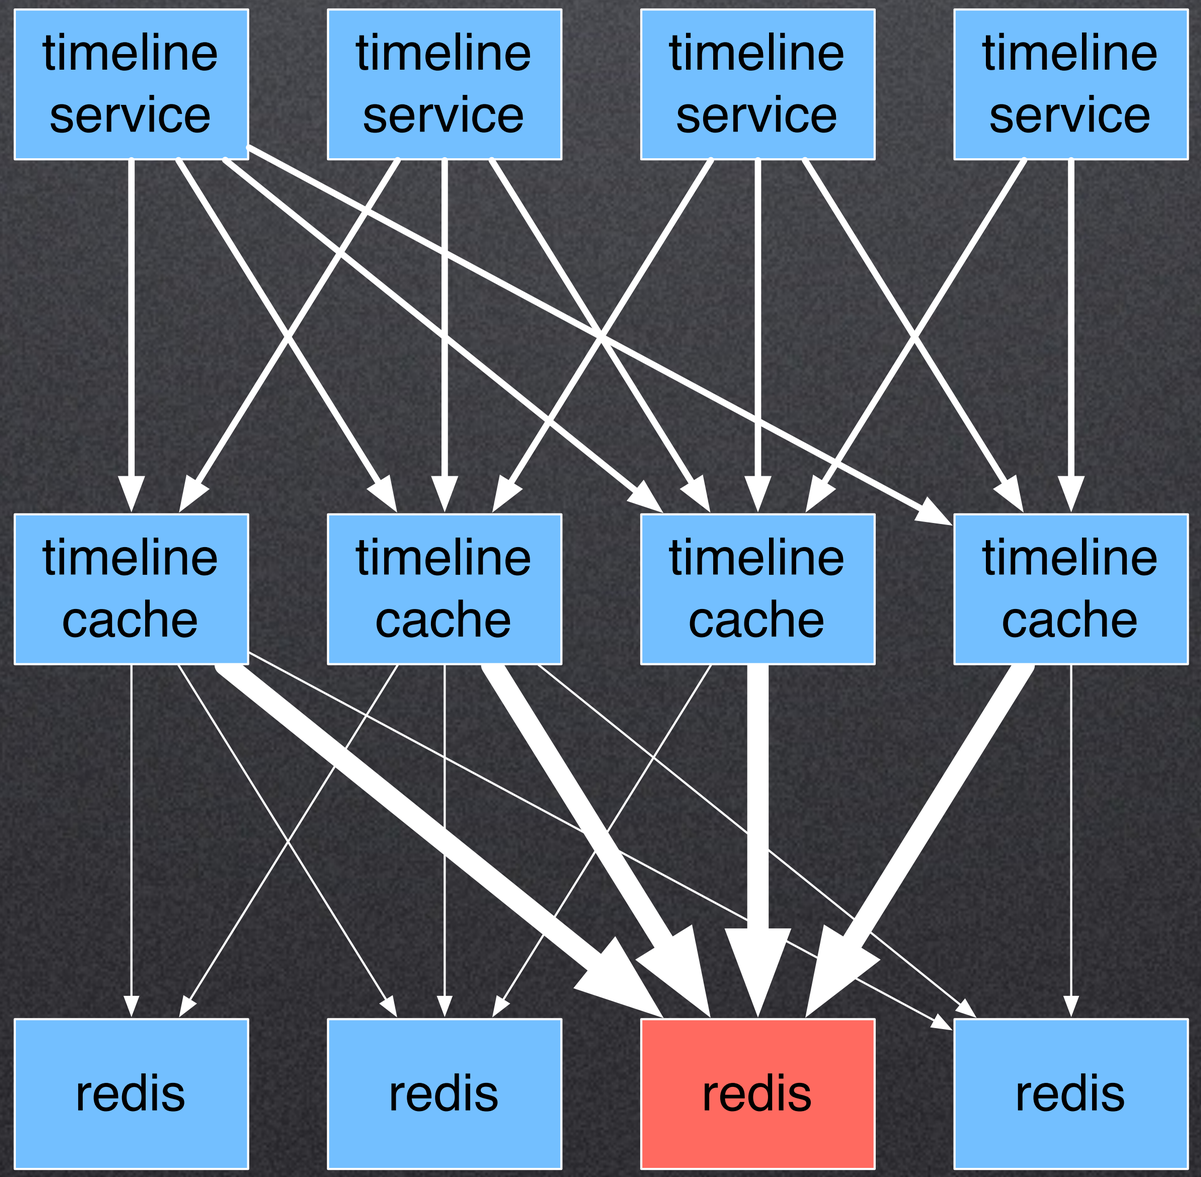
timeline service
timeline service
timeline service
timeline service
timeline cache
timeline cache
timeline cache
timeline cache
redis
redis
redis
redis

in-process cache

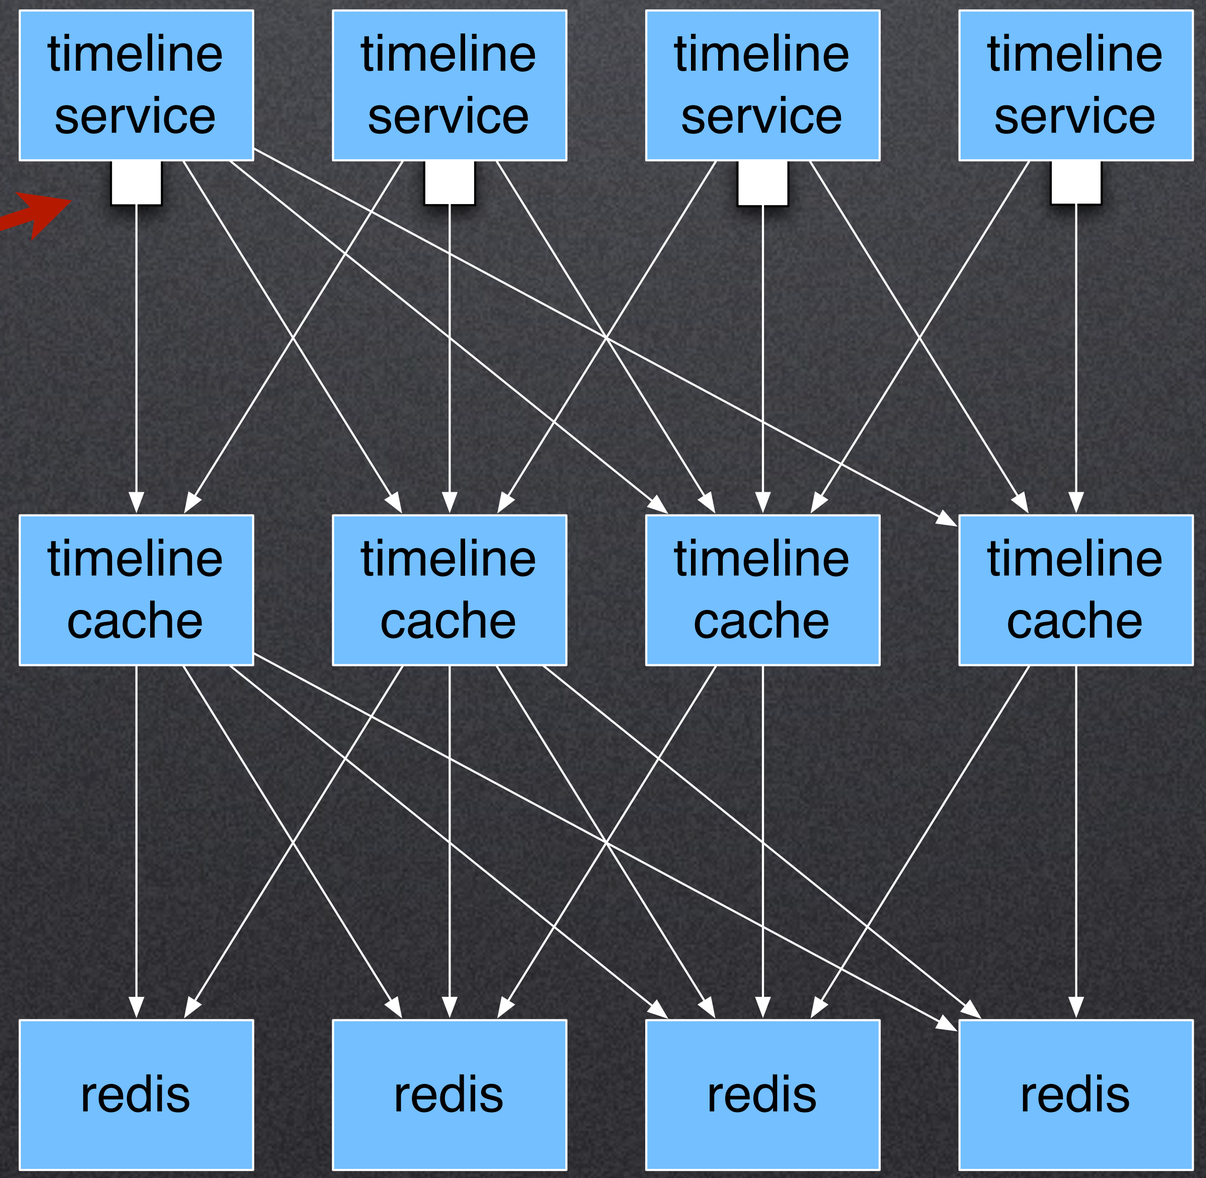

In-Process Cache


timeline
service


timeline
service
timeline
cache


timeline
service
timeline
cache


timeline
service
timeline
cache

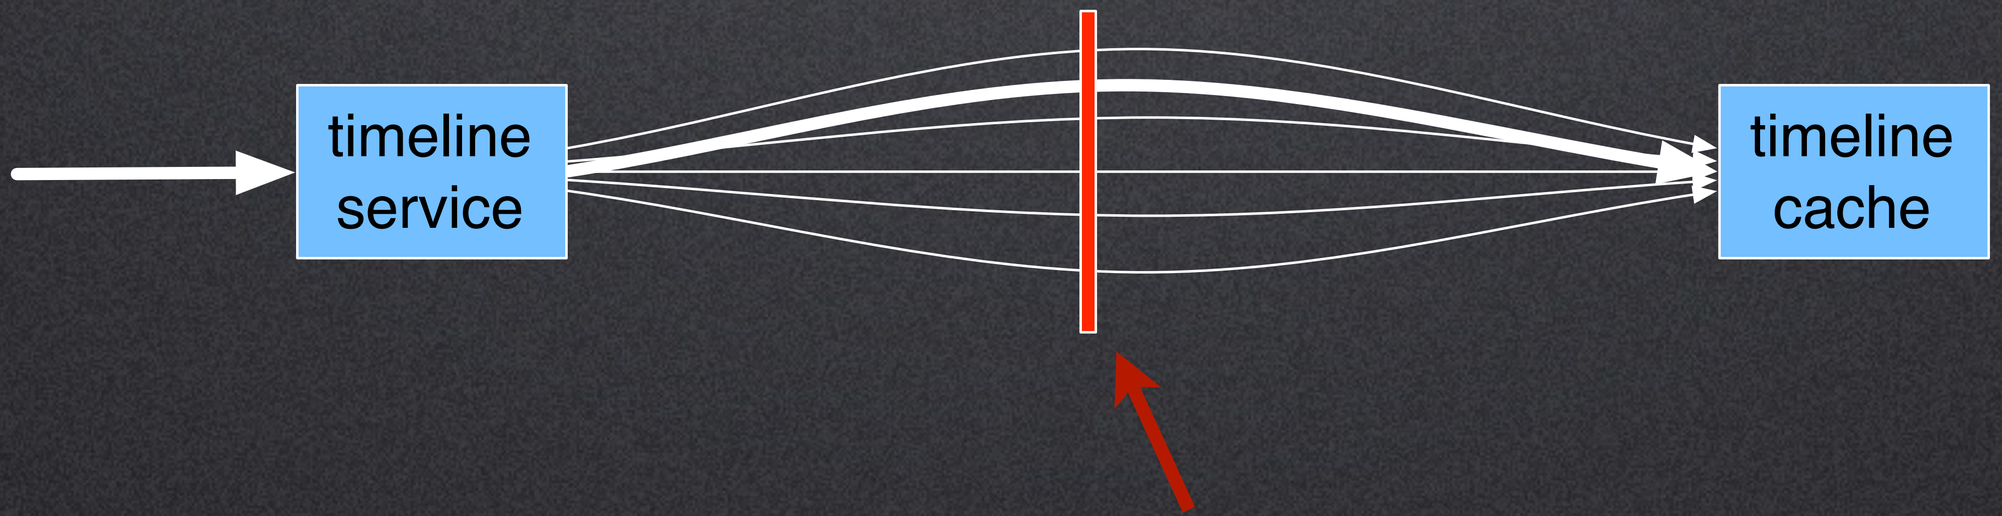
timeline service
timeline cache
detects high-frequency requests

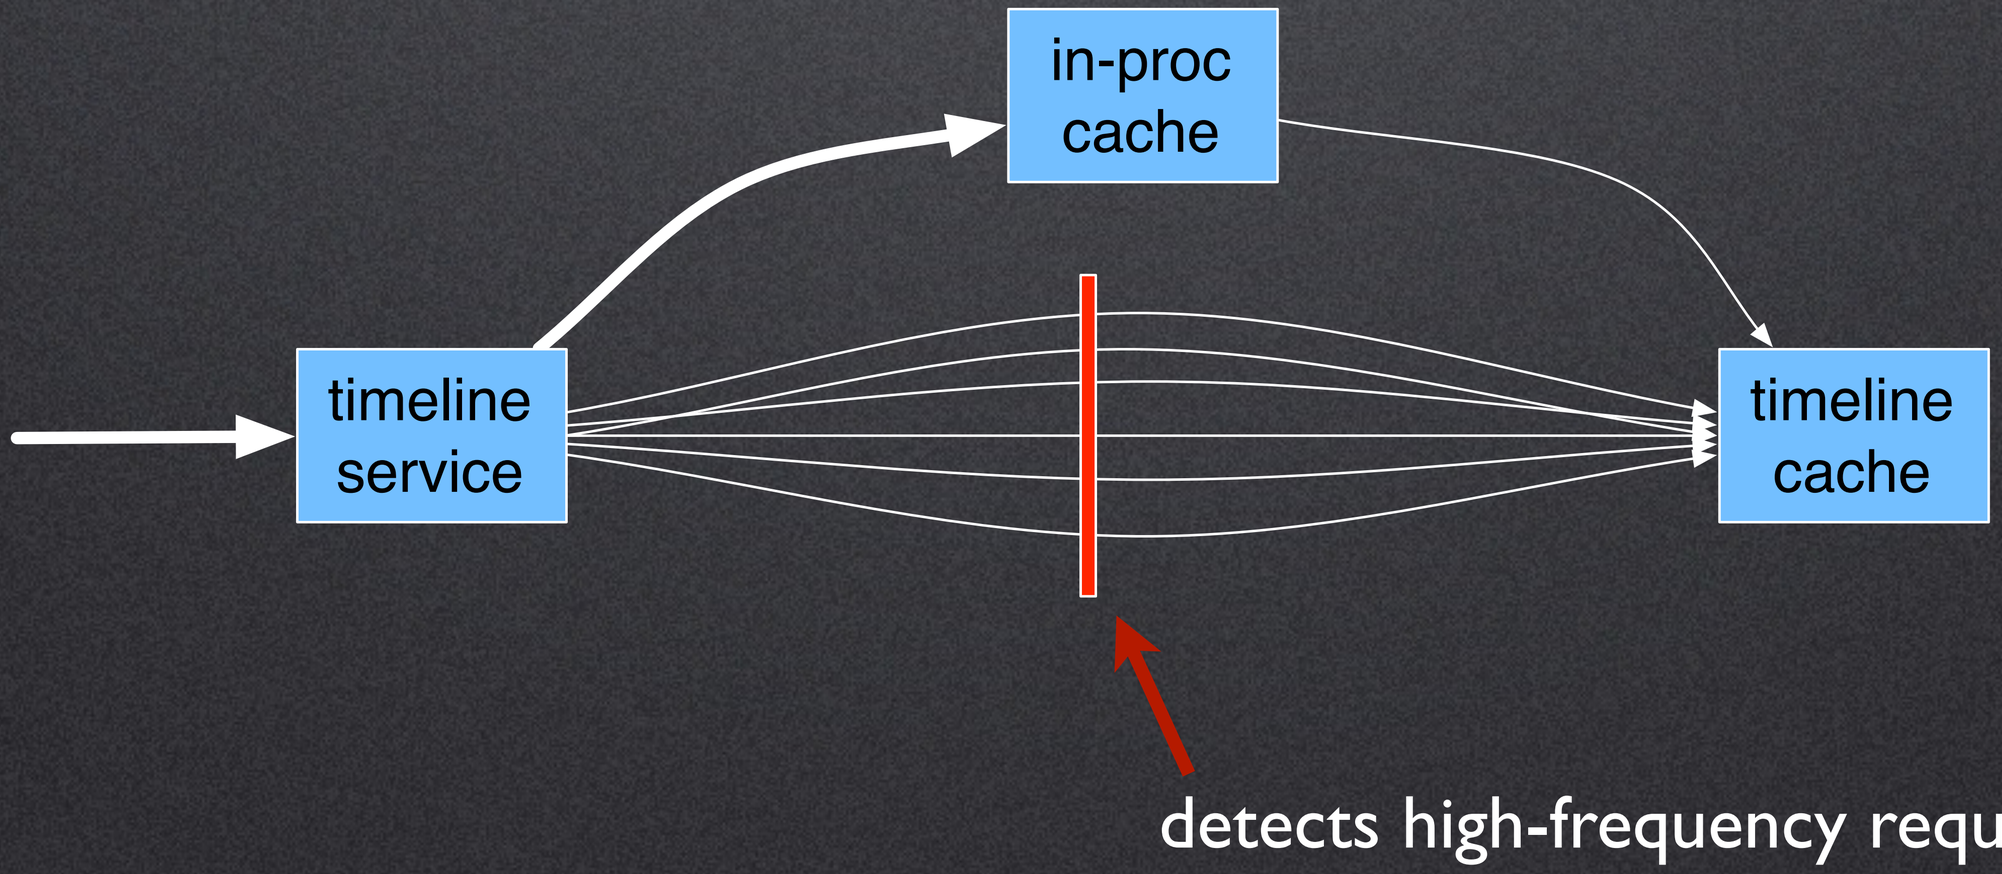
in-proc
cache
timeline
service
timeline
cache
detects high-frequency requests

Search


QCon London
Search
Tweets Top / All
QCon London @qconlondon 1h
Vote on the most important software development trends for 2012 t.co/Zaio6Rge, this will be announced in the evening keynote tonight
Open Source @opensource_ 2h
Sold out QCon kicks off in London: big data, mobile, cloud, HTML 5 #opensource t.co/csHookQv
Krzysztof Wilczynski @kwilczynski 2h
New slide from QCon 2012 by @patrickdebois t.co/d9RbHru7
View media
Dušan Omerčević @dusano 3h
#QConLondon 2012 tutorial on Continuous delivery aka. Making deployment a non-event t.co/HXYMVS6H
View media
Liisi Toom @liisi 3h
QCon London is ON! ZeroTurnaround is in the house! #fb pic.twitter.com/xLWTrDMC
View photo

What's happening?

hello world

129 Tweet

HTTP
Proxy

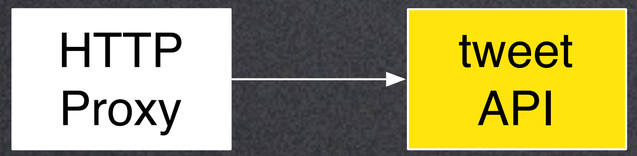
HTTP
Proxy
tweet
API


HTTP
Proxy
tweet
API
queue


HTTP
Proxy
tweet
API
queue
ingester

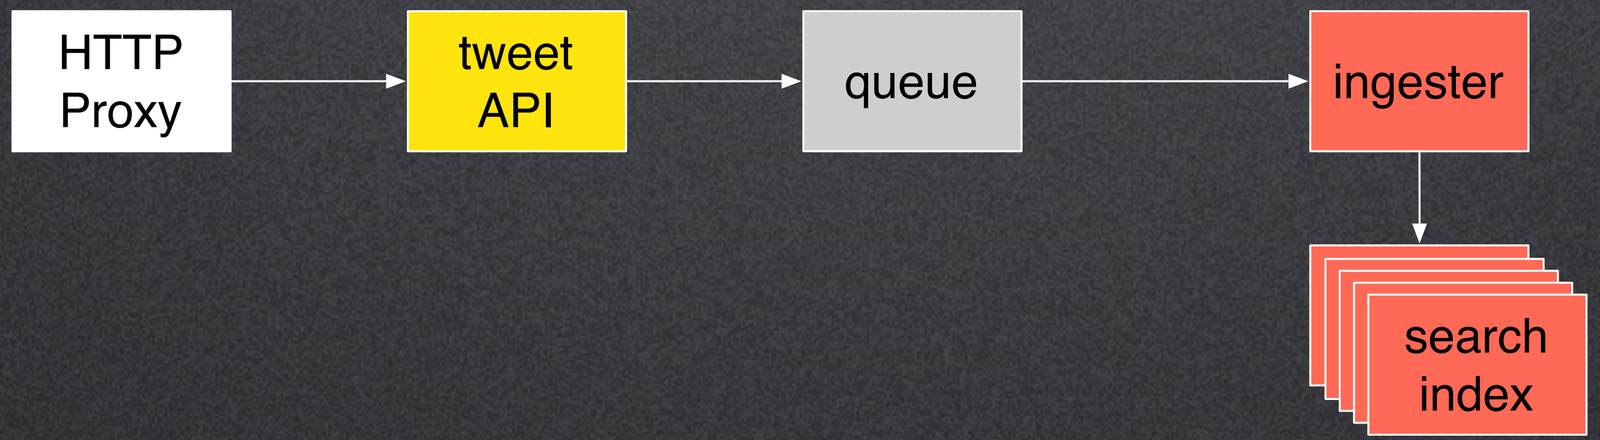
HTTP
Proxy
tweet
API
queue
ingester
search
index

Search


QCon London
Search
Tweets Top / All
QCon London @qconlondon 1h
Vote on the most important software development trends for 2012 t.co/Zaio6Rge, this will be announced in the evening keynote tonight
Open Source @opensource_ 2h
Sold out QCon kicks off in London: big data, mobile, cloud, HTML 5 #opensource t.co/csHookQv
Krzysztof Wilczynski @kwilczynski 2h
New slide from QCon 2012 by @patrickdebois t.co/d9RbHru7
View media
Dušan Omerčević @dusano 3h
#QConLondon 2012 tutorial on Continuous delivery aka. Making deployment a non-event t.co/HXYMVS6H
View media
Liisi Toom @liisi 3h
QCon London is ON! ZeroTurnaround is in the house! #fb pic.twitter.com/xLWTrDMC
View photo

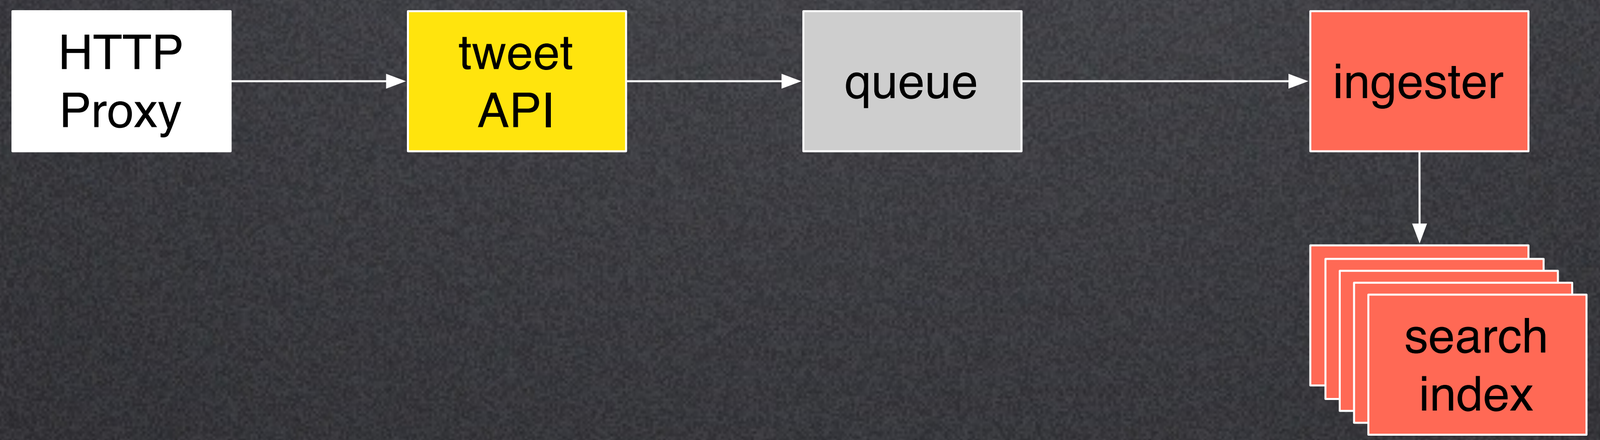
HTTP
Proxy
tweet
API
queue
ingester
search
index

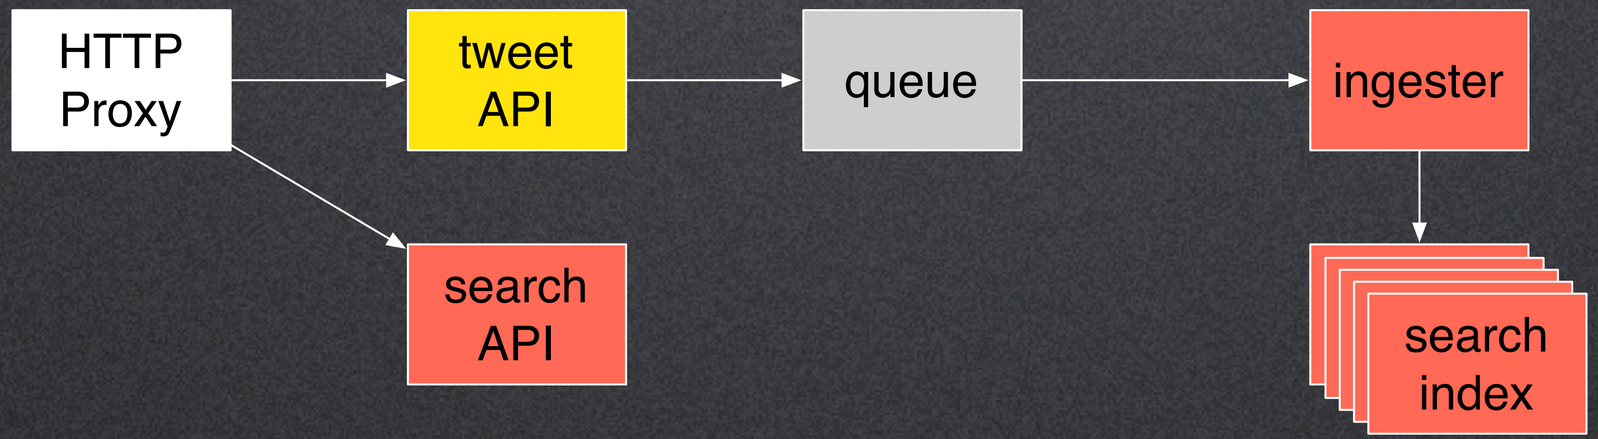
HTTP
Proxy
tweet
API
queue
ingester
search
API
search
index

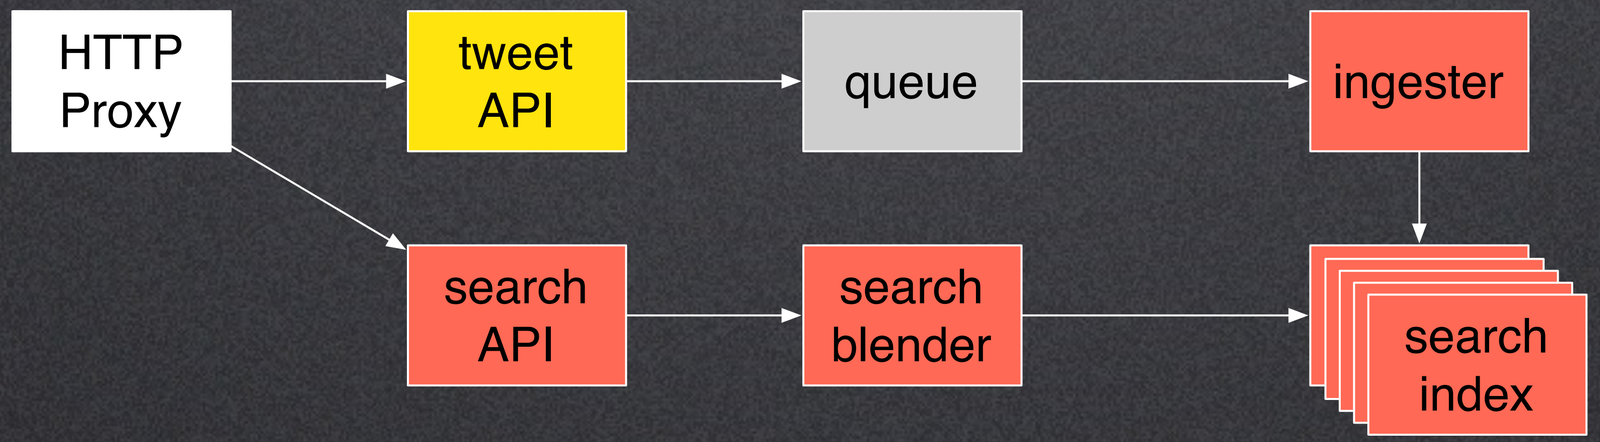
HTTP
Proxy
tweet
API
queue
ingester
search
API
search
blender
search
index

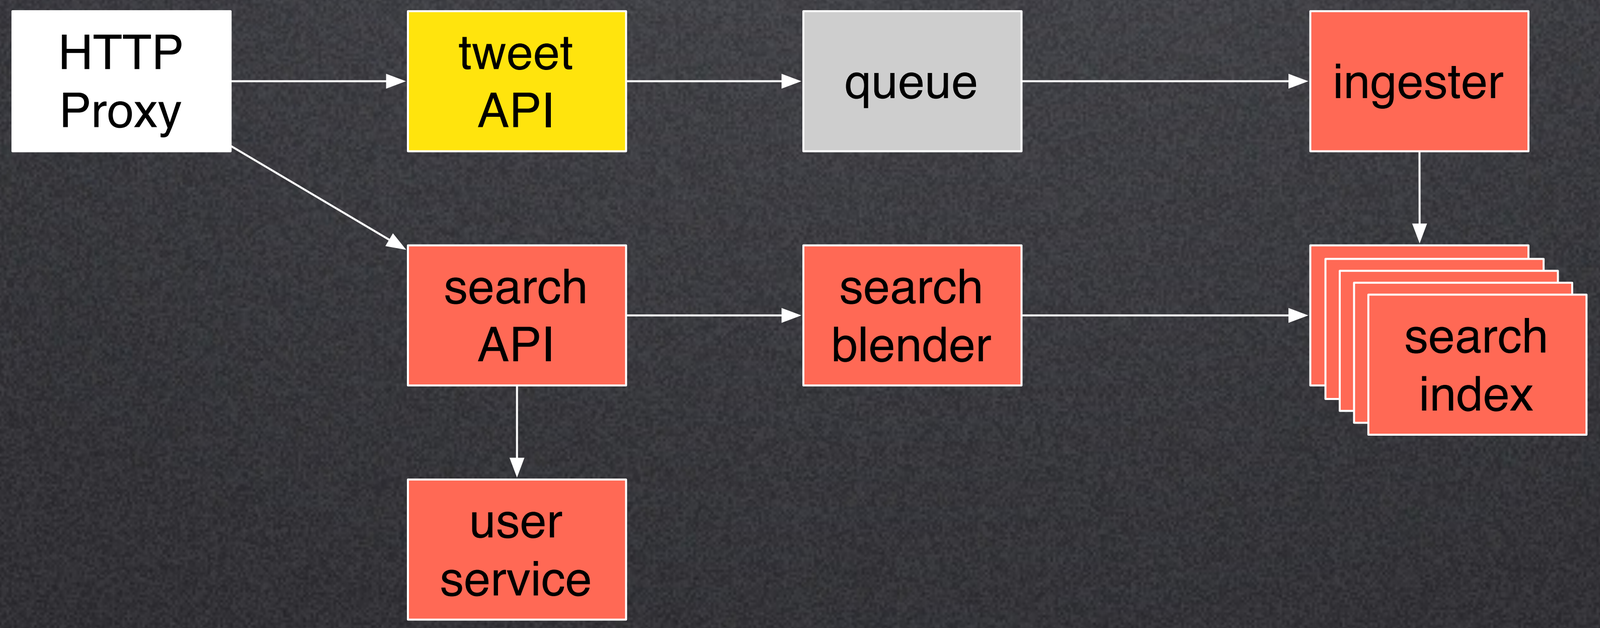
HTTP
Proxy
tweet
API
queue
ingester
search
API
search
blender
search
index
user
service


search
index


search index
search index
search index
search index

New Tweet

search index

search index

search index

search index

search
blender

search
index

search
index

search
index

search
index

search blender

search blender

search blender

search blender

search index

search index

search index

search index

search blender

search blender

search blender

search blender

search index

search index

search index

search index

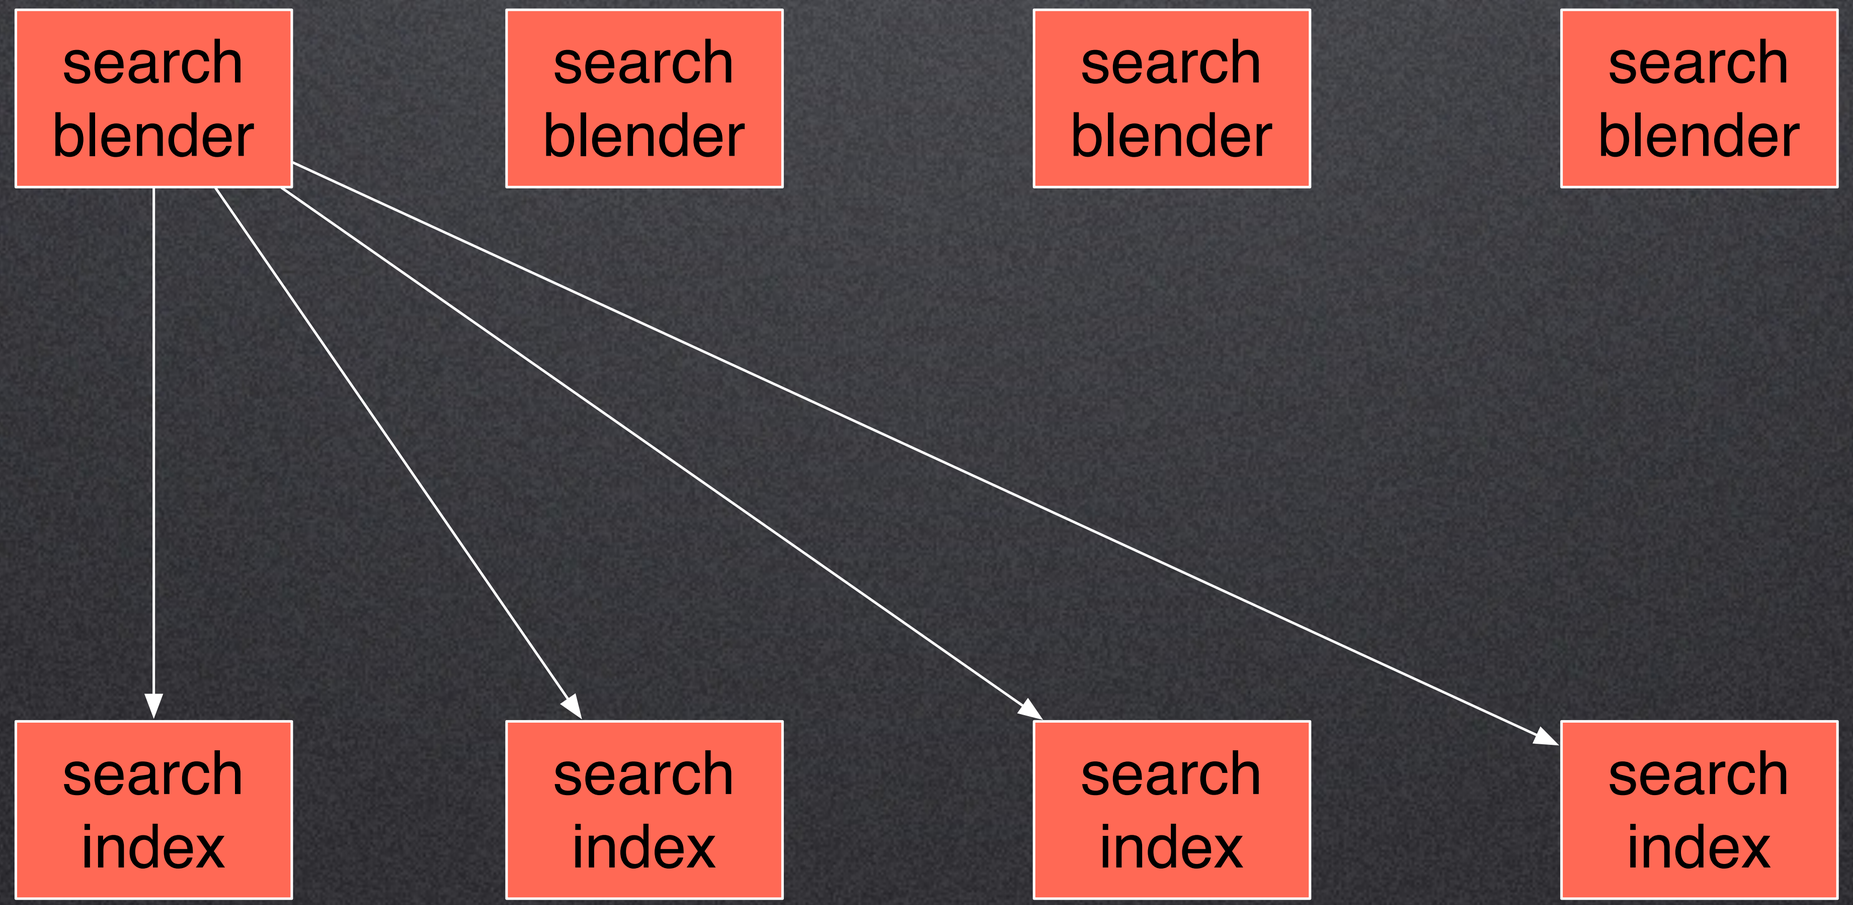
search blender
search blender
search blender
search blender
search index
search index
search index
search index

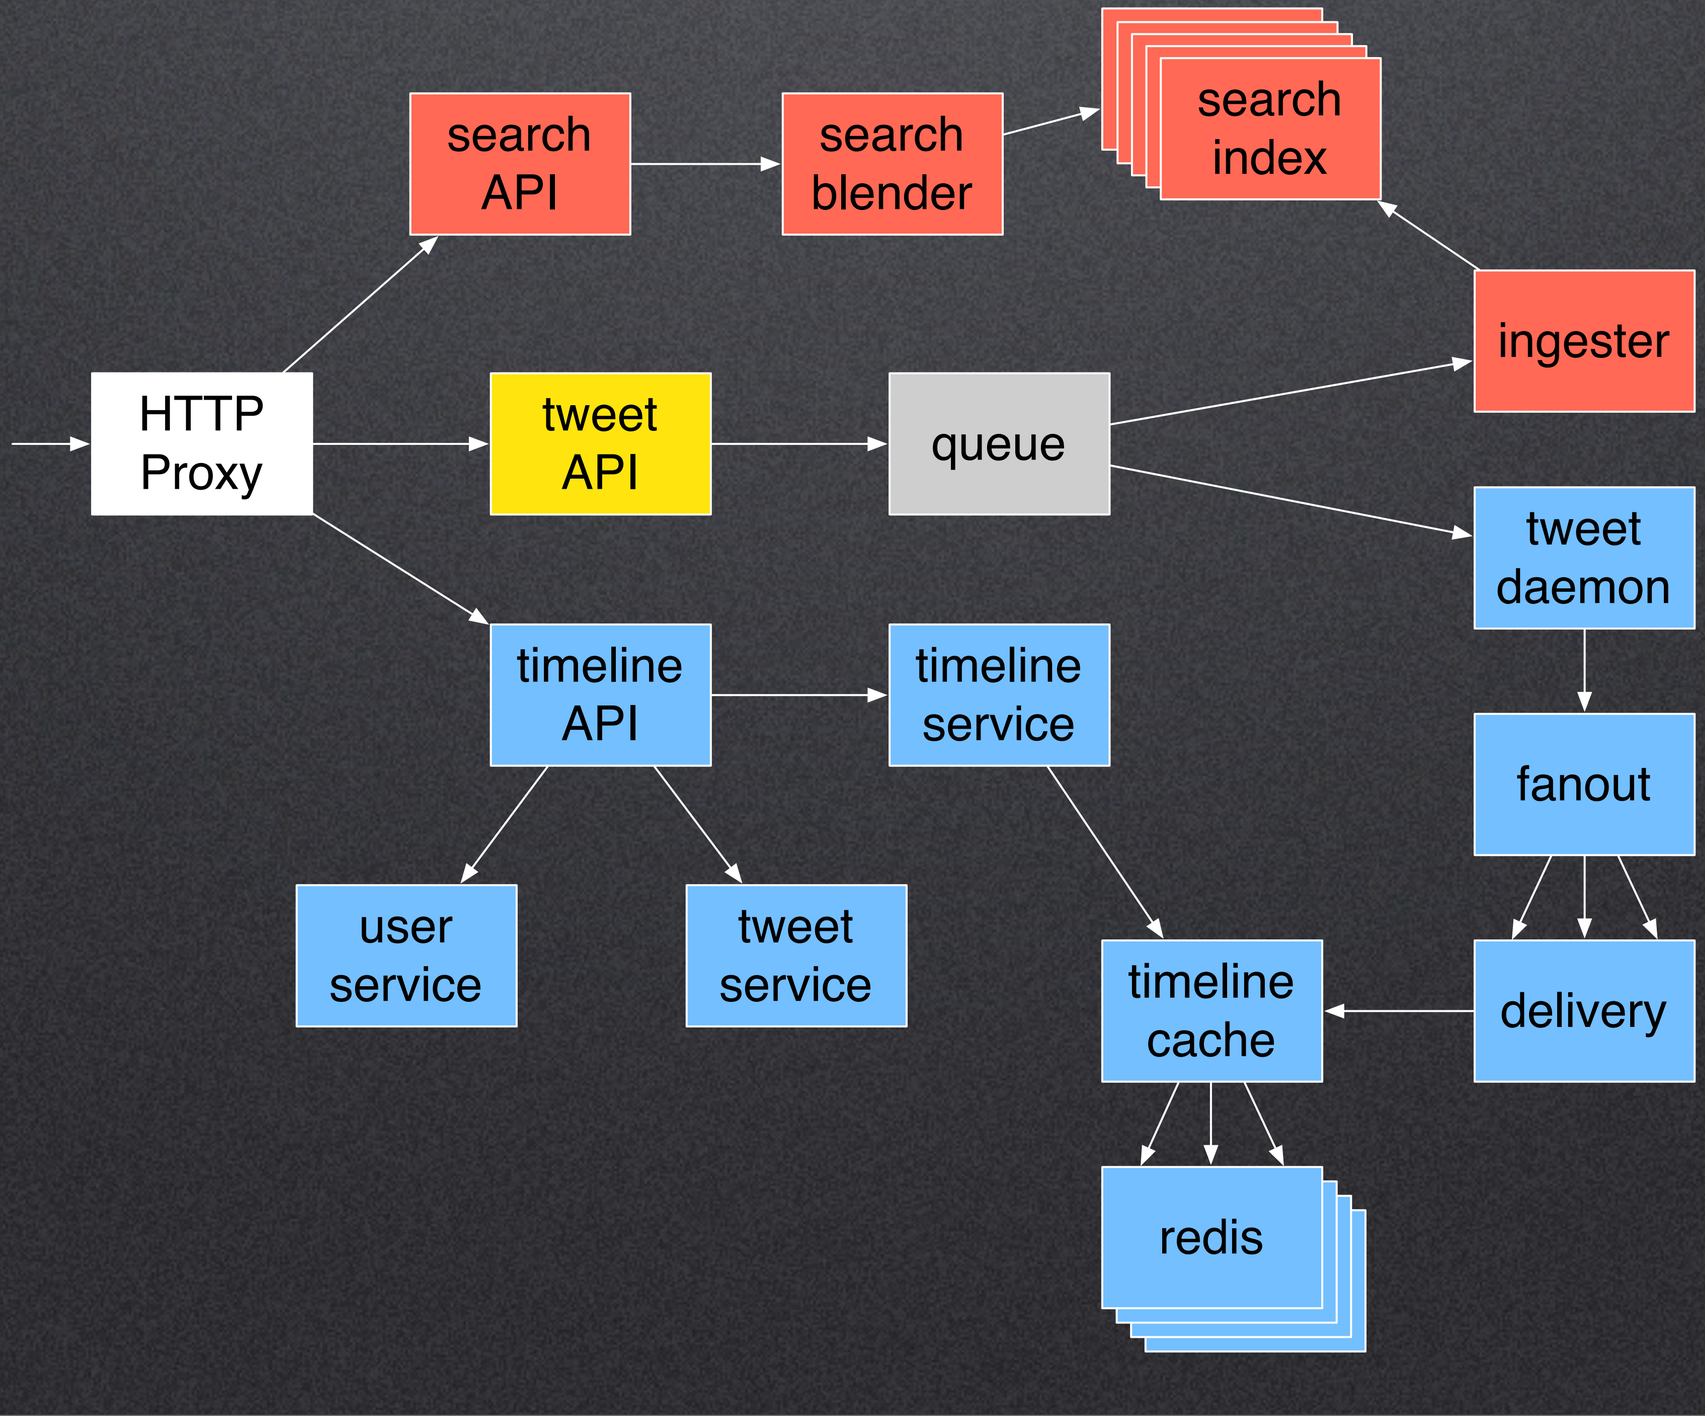
search API
search blender
search index
HTTP Proxy
tweet API
queue
ingester
tweet daemon
timeline API
timeline service
fanout
user service
tweet service
timeline cache
delivery
redis

search API

search blender

search index

ingester

HTTP Proxy

tweet API

queue

Durable Stores

tweet daemon

timeline API

timeline service

fanout

user service

tweet service

timeline cache

delivery

redis

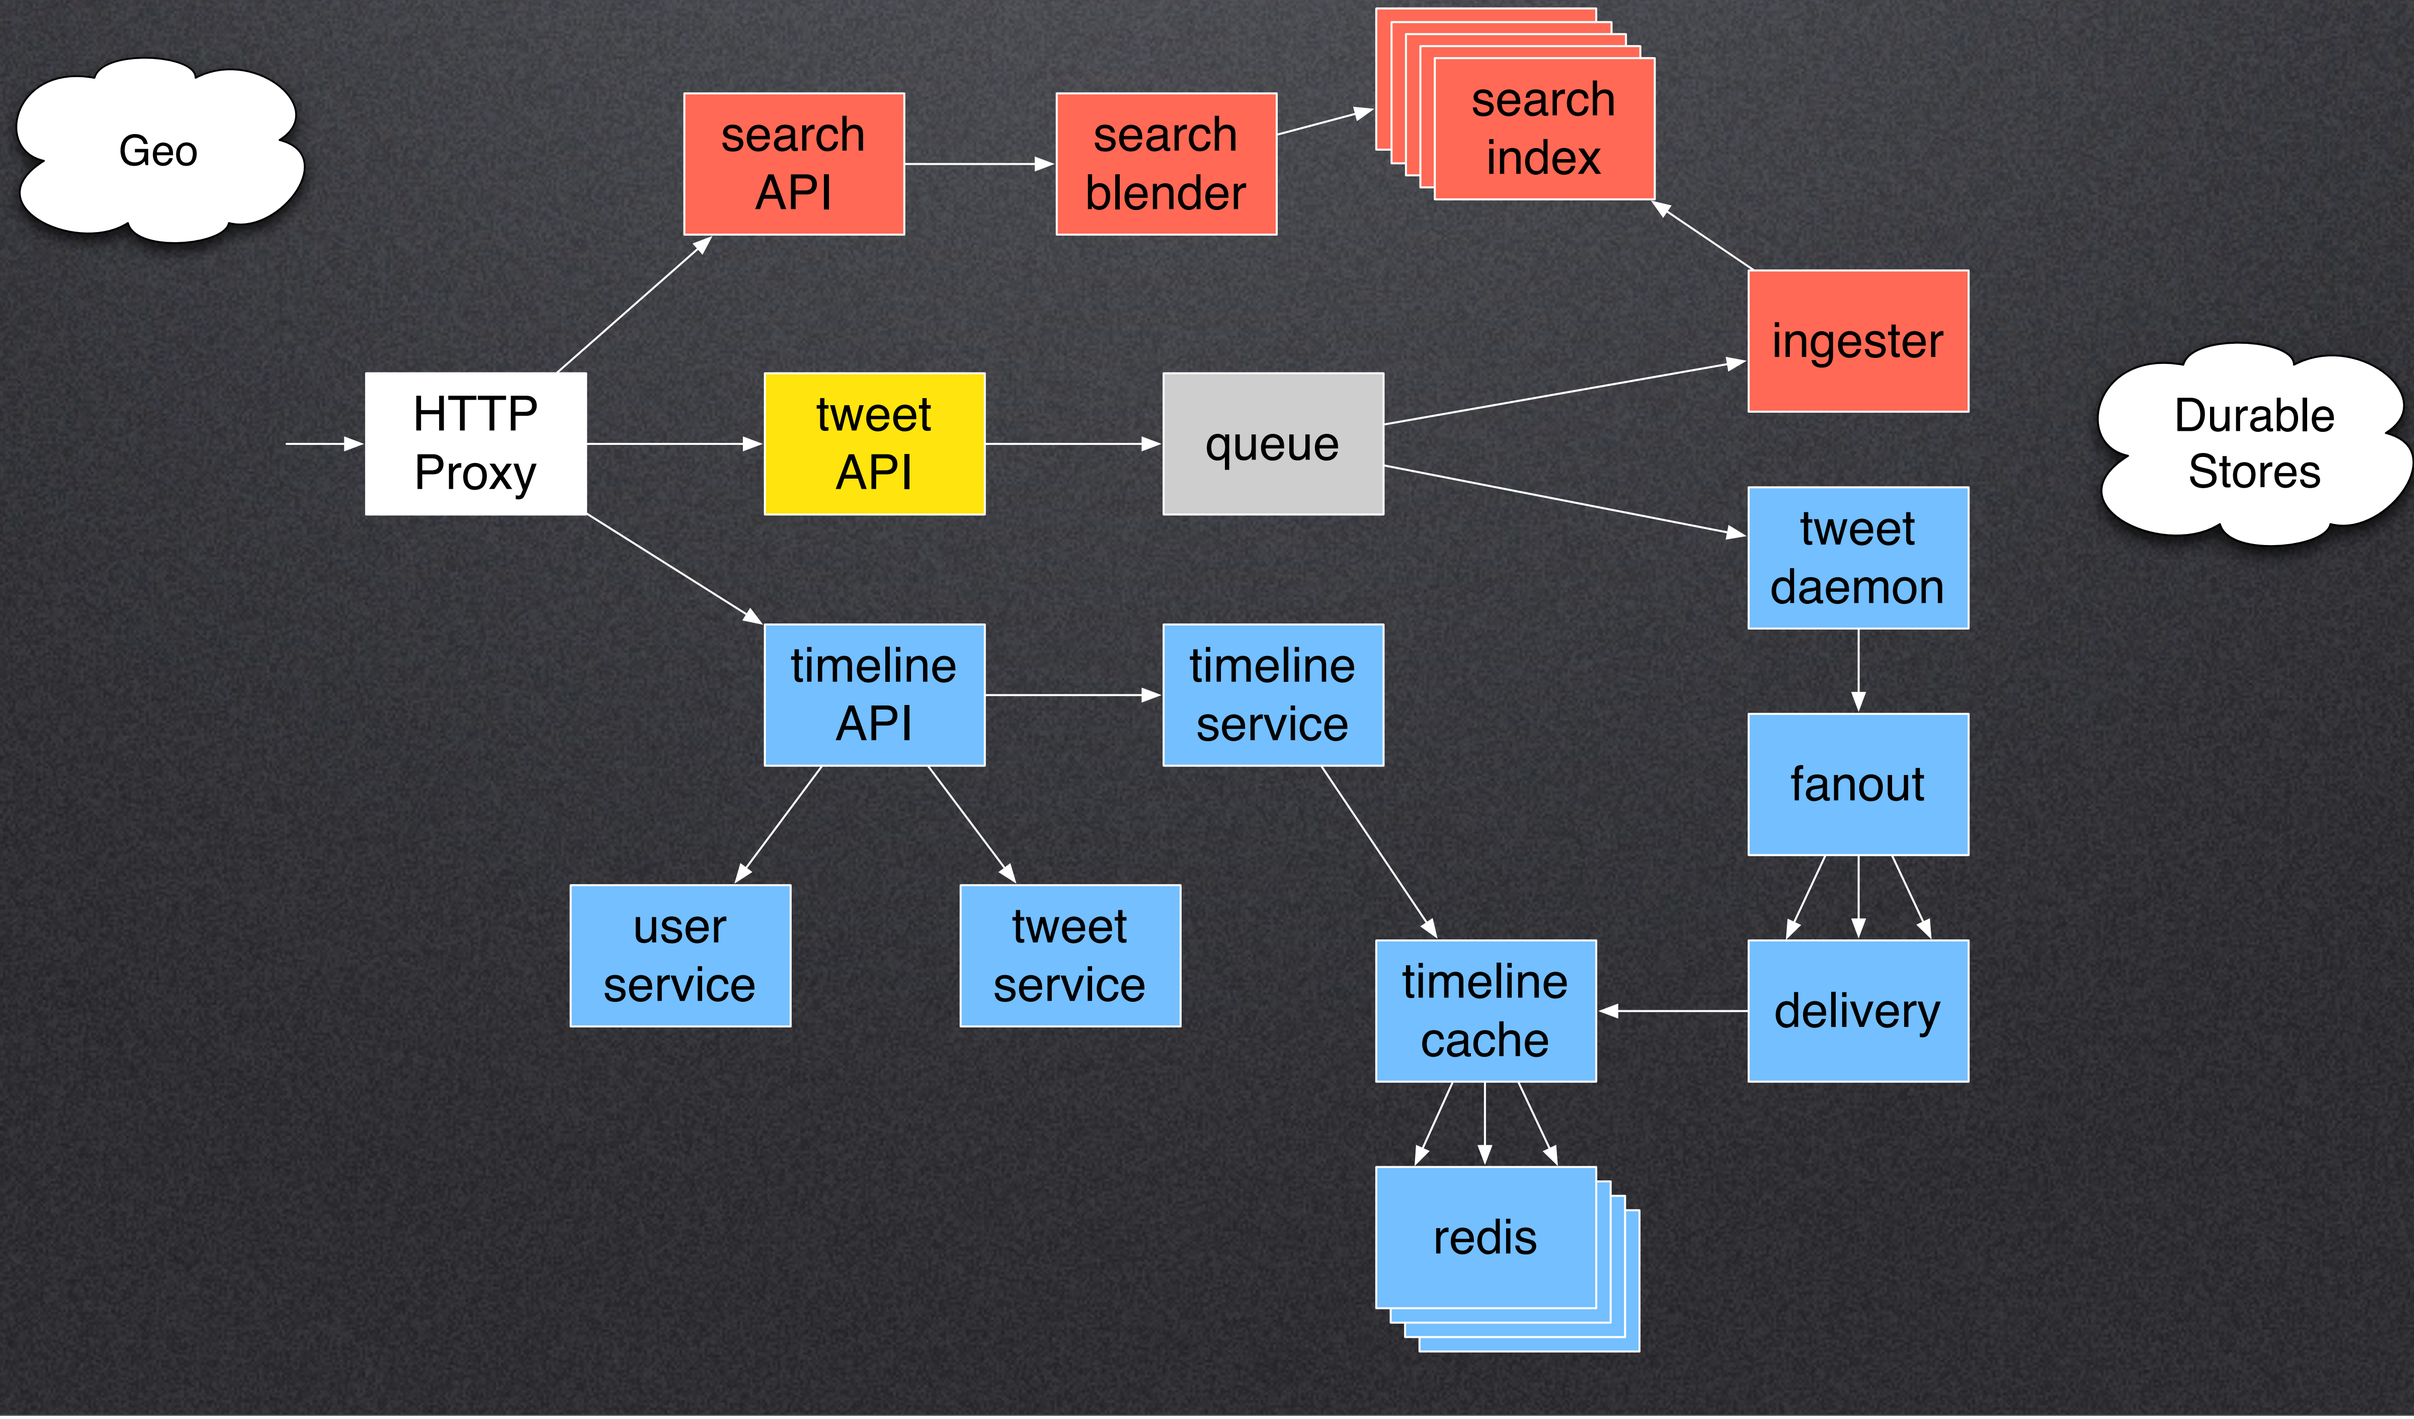
Geo
search API
search blender
search index
ingester
Durable Stores
HTTP Proxy
tweet API
queue
tweet daemon
timeline API
timeline service
fanout
user service
tweet service
timeline cache
delivery
redis

Geo

search API

search blender

search index

HTTP Proxy

tweet API

queue

ingester

Durable Stores

tweet daemon

timeline API

timeline service

fanout

user service

tweet service

timeline cache

delivery

redis

URL expansion

Geo

search API

search blender

search index

HTTP Proxy

tweet API

queue

ingester

Durable Stores

tweet daemon

timeline API

timeline service

fanout

user service

tweet service

timeline cache

delivery

redis

URL expansion

Geo

search API

search blender

search index

HTTP Proxy

tweet API

queue

ingester

Durable Stores

tweet daemon

timeline API

timeline service

fanout

user service

tweet service

timeline cache

delivery

redis

URL expansion

Geo

search API

search blender

search index

HTTP Proxy

tweet API

queue

ingester

Durable Stores

tweet daemon

timeline API

timeline service

fanout

user service

tweet service

timeline cache

delivery

redis

URL expansion

Geo

search API

search blender

search index

HTTP Proxy

tweet API

queue

ingester

Durable Stores

tweet daemon

timeline API

timeline service

fanout

user service

tweet service

timeline cache

delivery

redis

URL expansion

HTTP Proxy

timeline API

Finagle - JVM RPC Library

open source:
http://github.com/twitter/finagle

Finagle - RPC Common Library

Components

- connection management
- protocol codecs
- transient error handling
- distributed tracing
- service discovery
- observability

Finagle – RPC Common Library

Components

- connection management
- protocol codecs
- transient error handling
- distributed tracing
- service discovery
- observability

Building a client

```
ClientBuilder()
  .cluster(TimelineServiceCluster)
  .hostConnectionCoresize(5)
  .hostConnectionLimit(10)
  .hostConnectionIdleTime(5.seconds)
  .failureAccrualParams(5, 10.seconds)
  .retries(3)
  .timeout(500.milliseconds)
```

Building a client

```
ClientBuilder()
  .cluster(TimelineServiceCluster)
  .hostConnectionCoresize(5)
  .hostConnectionLimit(10)
  .hostConnectionIdleTime(5.seconds)
  .failureAccrualParams(5, 10.seconds)
  .retries(3)
  .timeout(500.milliseconds)
```

Building a client

```
ClientBuilder()
  .cluster(TimelineServiceCluster)
  .hostConnectionCoresize(5)
  .hostConnectionLimit(10)
  .hostConnectionIdleTime(5.seconds)
  .failureAccrualParams(5, 10.seconds)
  .retries(3)
  .timeout(500.milliseconds)
```

Building a client

```
ClientBuilder()
  .cluster(TimelineServiceCluster)
  .hostConnectionCoresize(5)
  .hostConnectionLimit(10)
  .hostConnectionIdleTime(5.seconds)
  .failureAccrualParams(5, 10.seconds)
  .retries(3)
  .timeout(500.milliseconds)
```

Building a client

```
ClientBuilder()
  .cluster(TimelineServiceCluster)
  .hostConnectionCoresize(5)
  .hostConnectionLimit(10)
  .hostConnectionIdleTime(5.seconds)
  .failureAccrualParams(5, 10.seconds)
  .retries(3)
  .timeout(500.milliseconds)
```

Building a client

```
ClientBuilder()
  .cluster(TimelineServiceCluster)
  .hostConnectionCoresize(5)
  .hostConnectionLimit(10)
  .hostConnectionIdleTime(5.seconds)
  .failureAccrualParams(5, 10.seconds)
  .retries(3)
  .timeout(500.milliseconds)
```

Building a client

```
ClientBuilder()
  .cluster(TimelineServiceCluster)
  .hostConnectionCoresize(5)
  .hostConnectionLimit(10)
  .hostConnectionIdleTime(5.seconds)
  .failureAccrualParams(5, 10.seconds)
  .retries(3)
  .timeout(500.milliseconds)
```

Building a client

```
ClientBuilder()
  .cluster(TimelineServiceCluster)
  .hostConnectionCoresize(5)
  .hostConnectionLimit(10)
  .hostConnectionIdleTime(5.seconds)
  .failureAccrualParams(5, 10.seconds)
  .retries(3)
  .timeout(500.milliseconds)
```

Finagle - RPC Common Library

Components

- connection management
- protocol codecs
- transient error handling
- distributed tracing
- service discovery
- observability

Finagle – RPC Common Library

Components

- connection management
- protocol codecs
- transient error handling
- distributed tracing
- service discovery
- observability

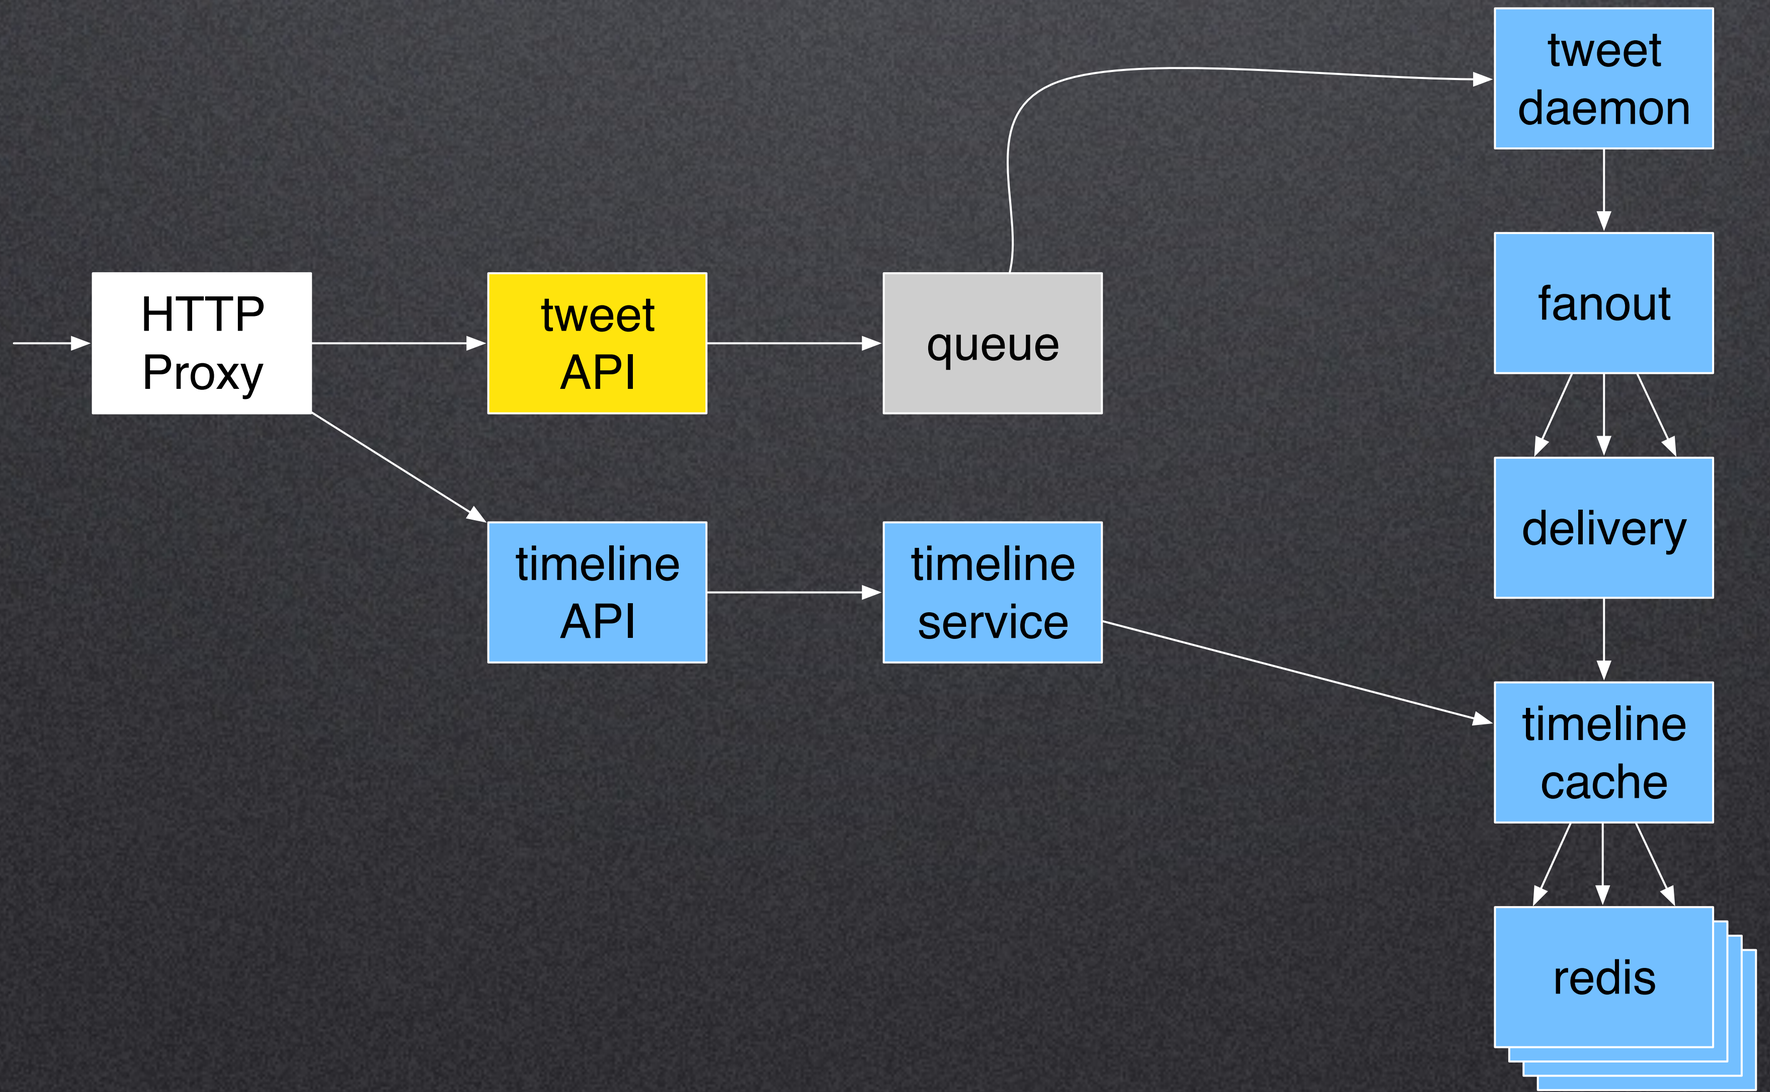
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
API
timeline
service
timeline
cache
redis

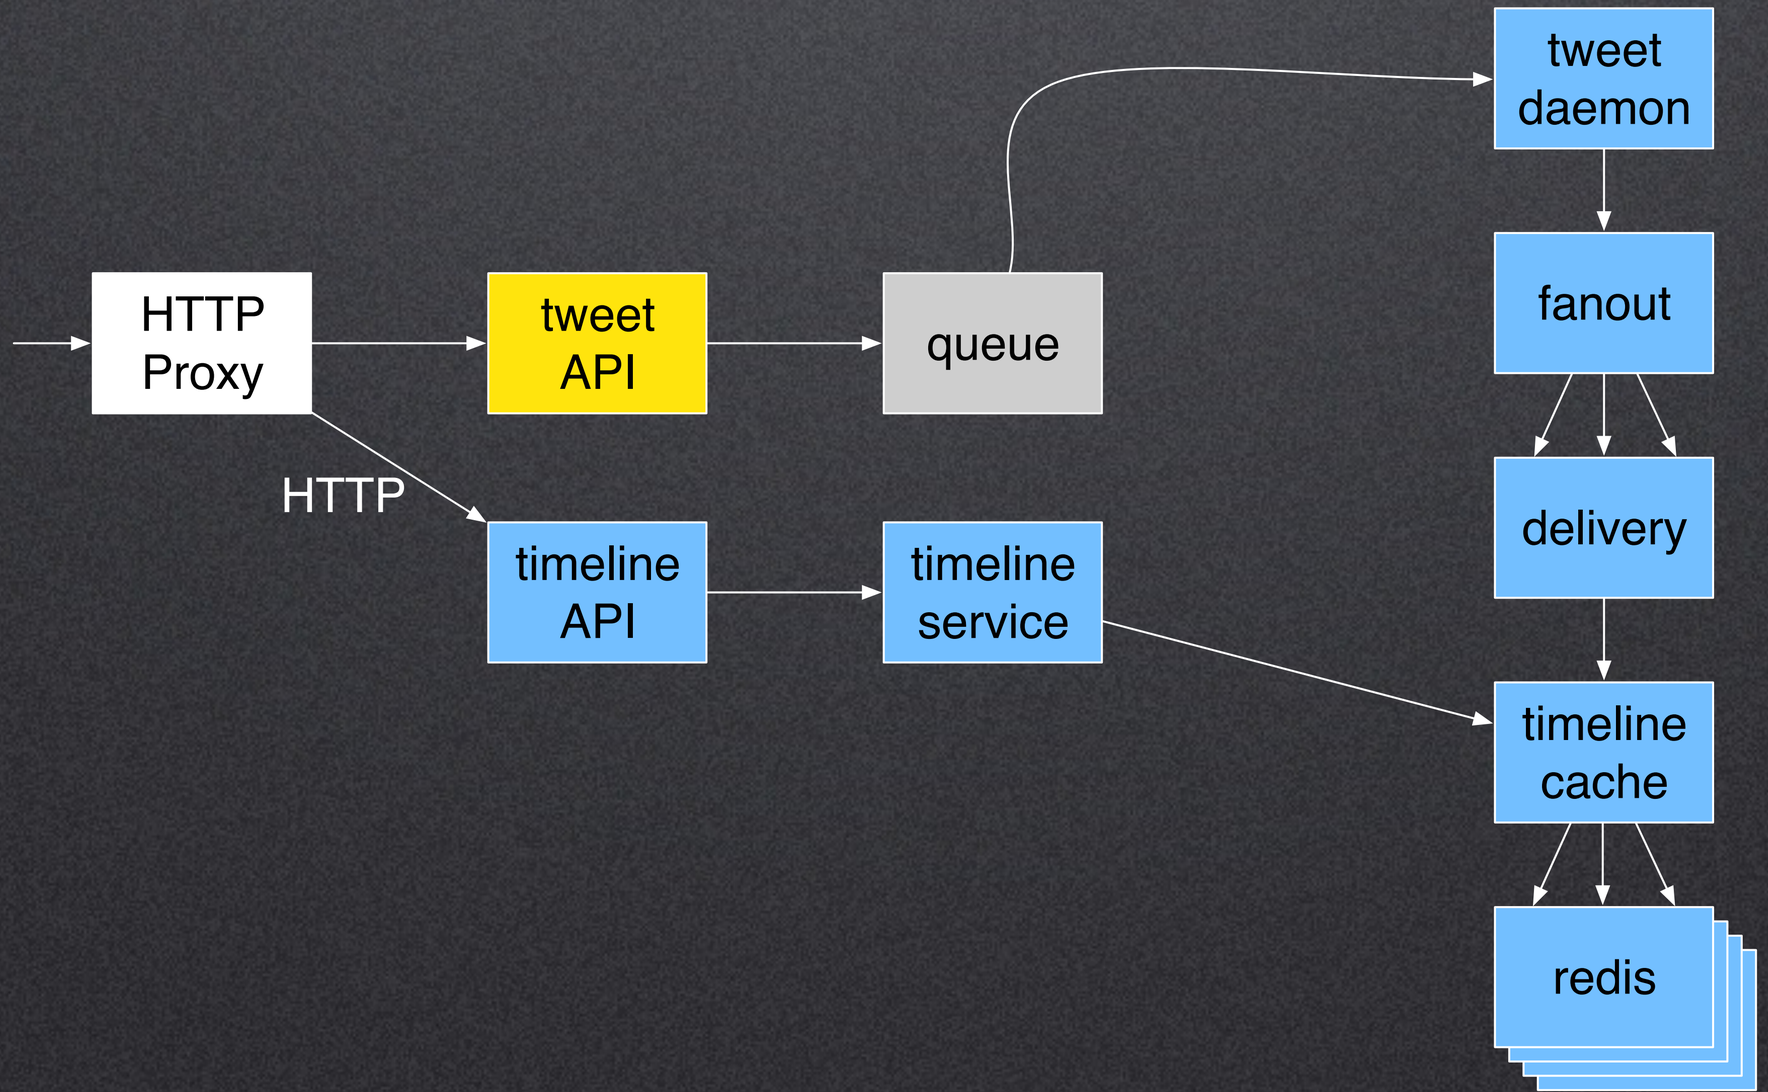
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
HTTP
timeline
API
timeline
service
timeline
cache
redis

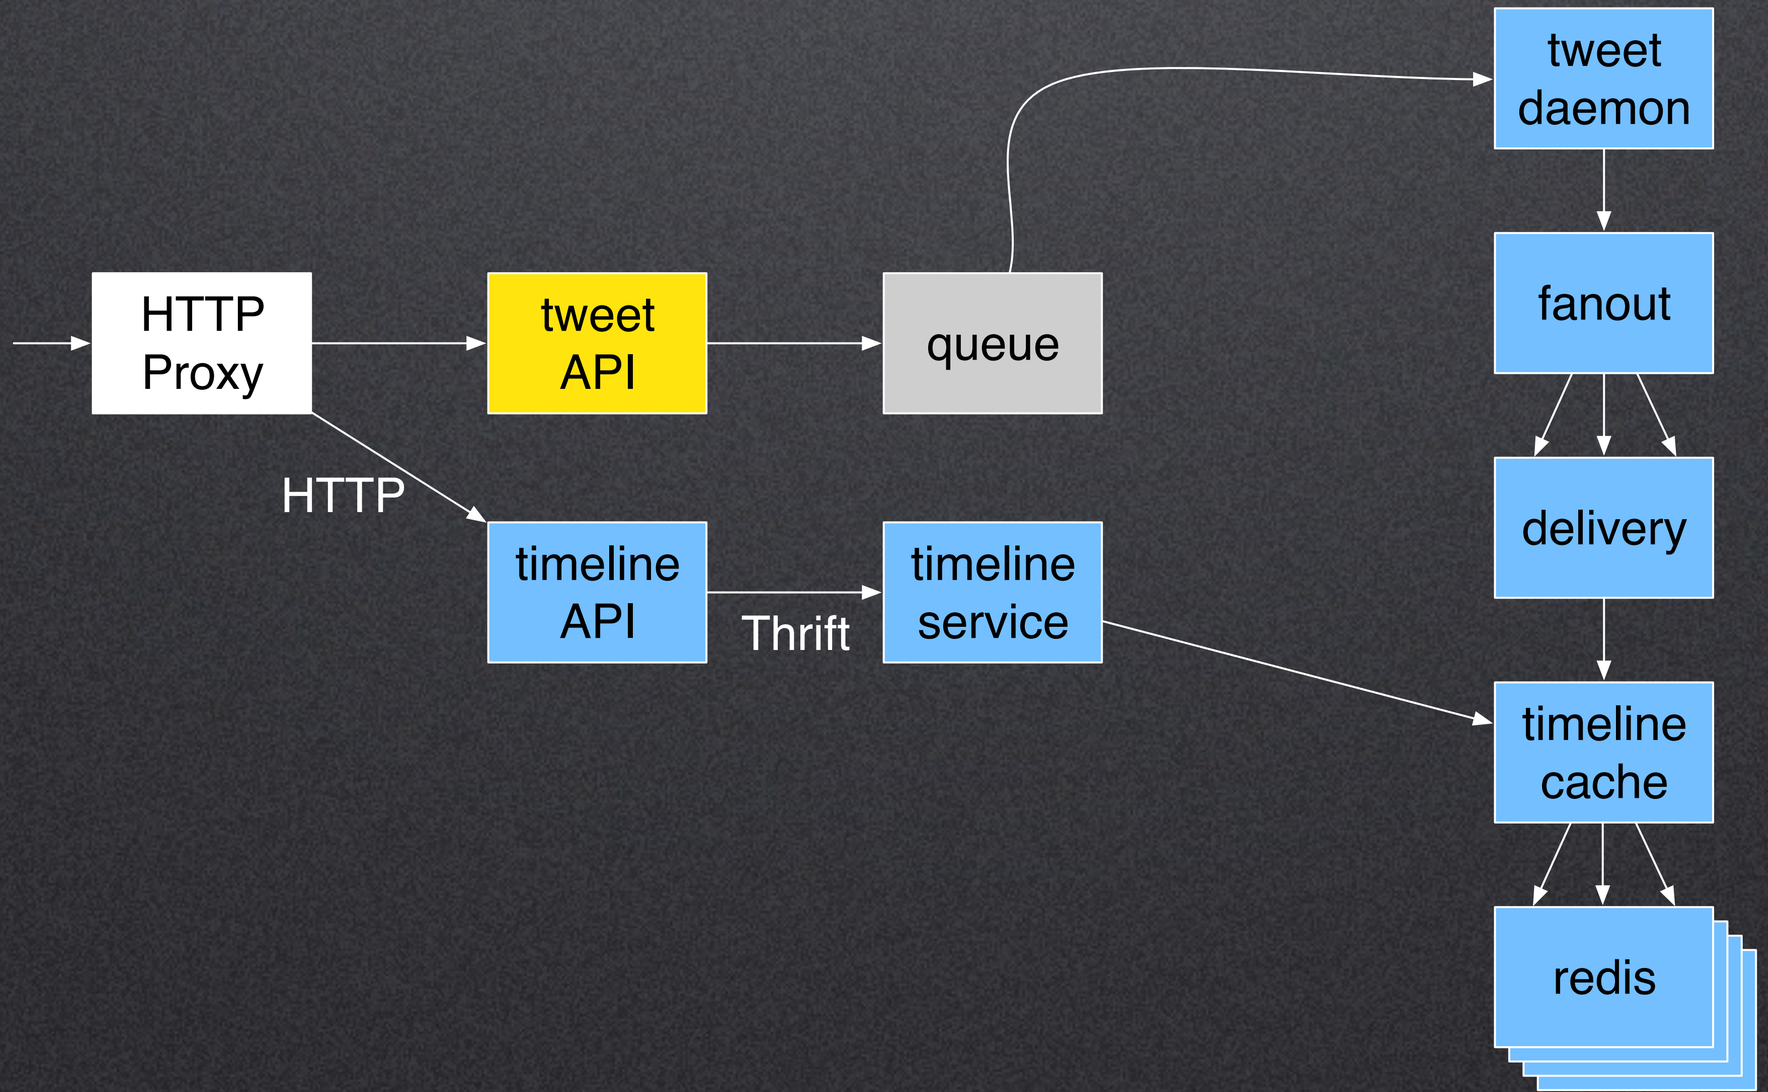
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
HTTP
timeline
API
Thrift
timeline
service
timeline
cache
redis

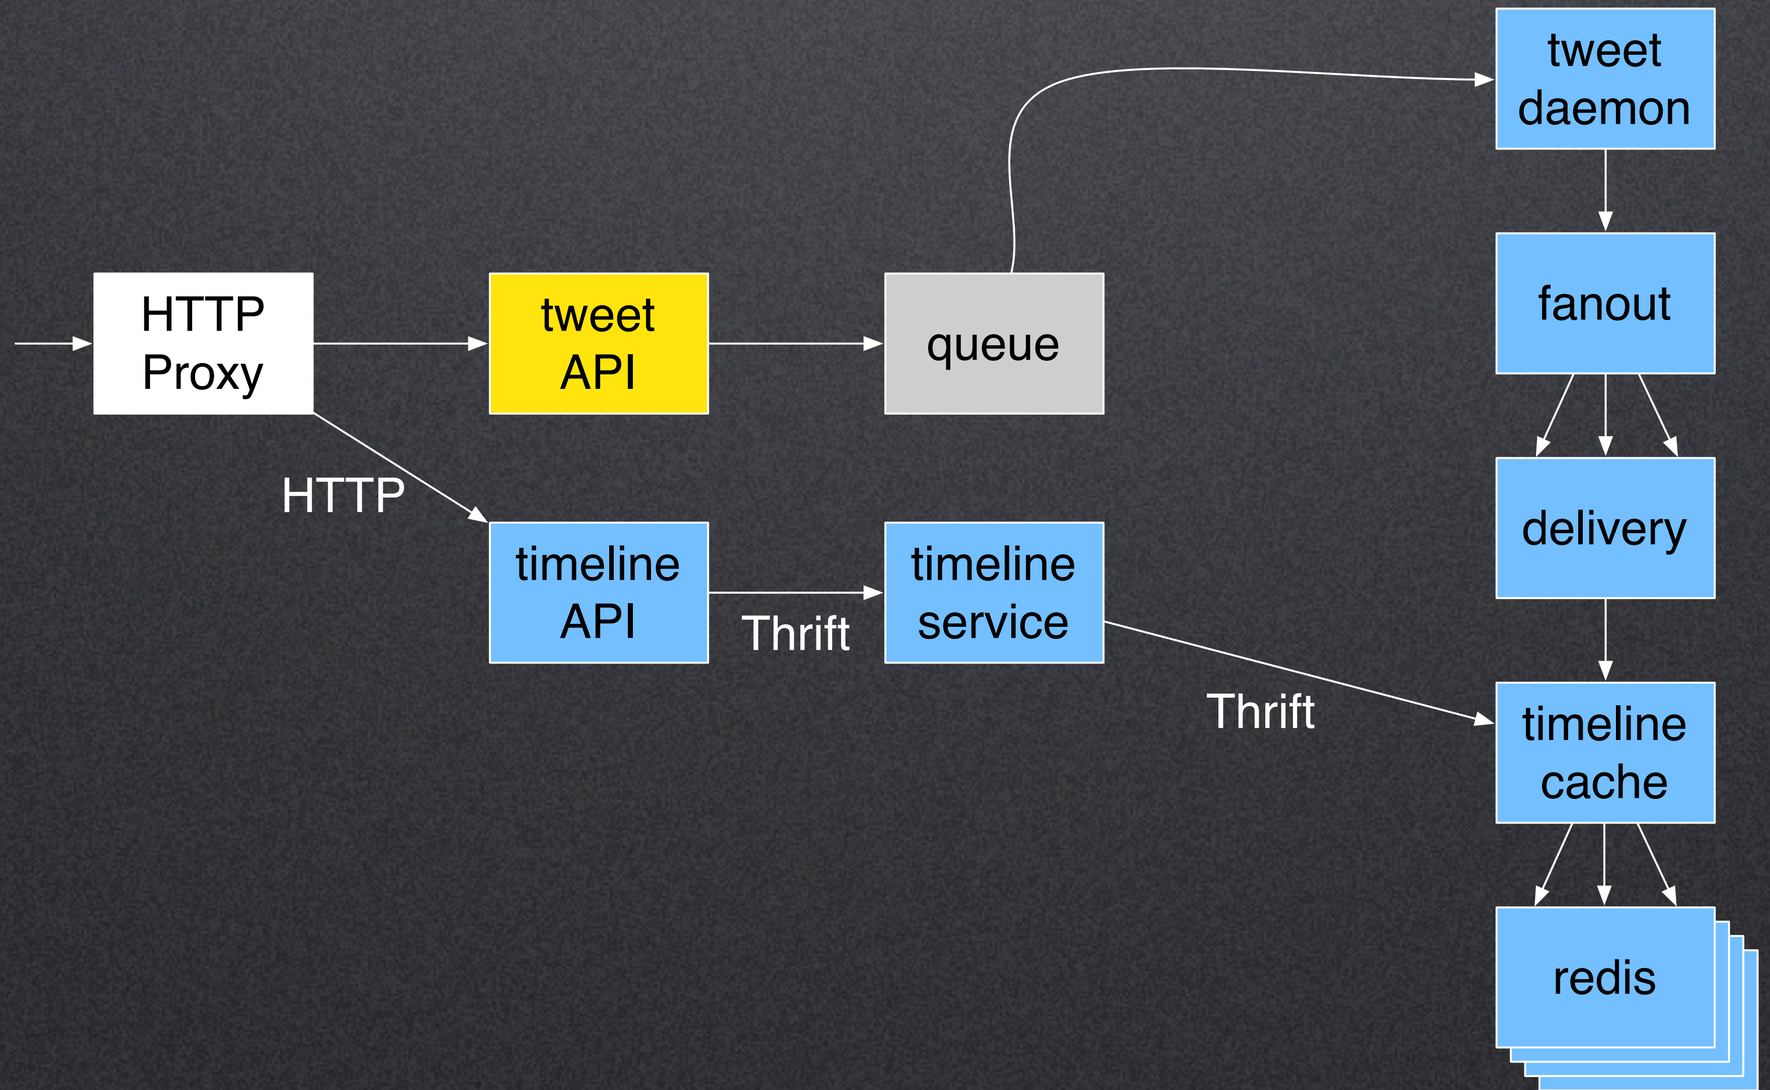
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
HTTP
timeline
API
Thrift
timeline
service
Thrift
timeline
cache
redis

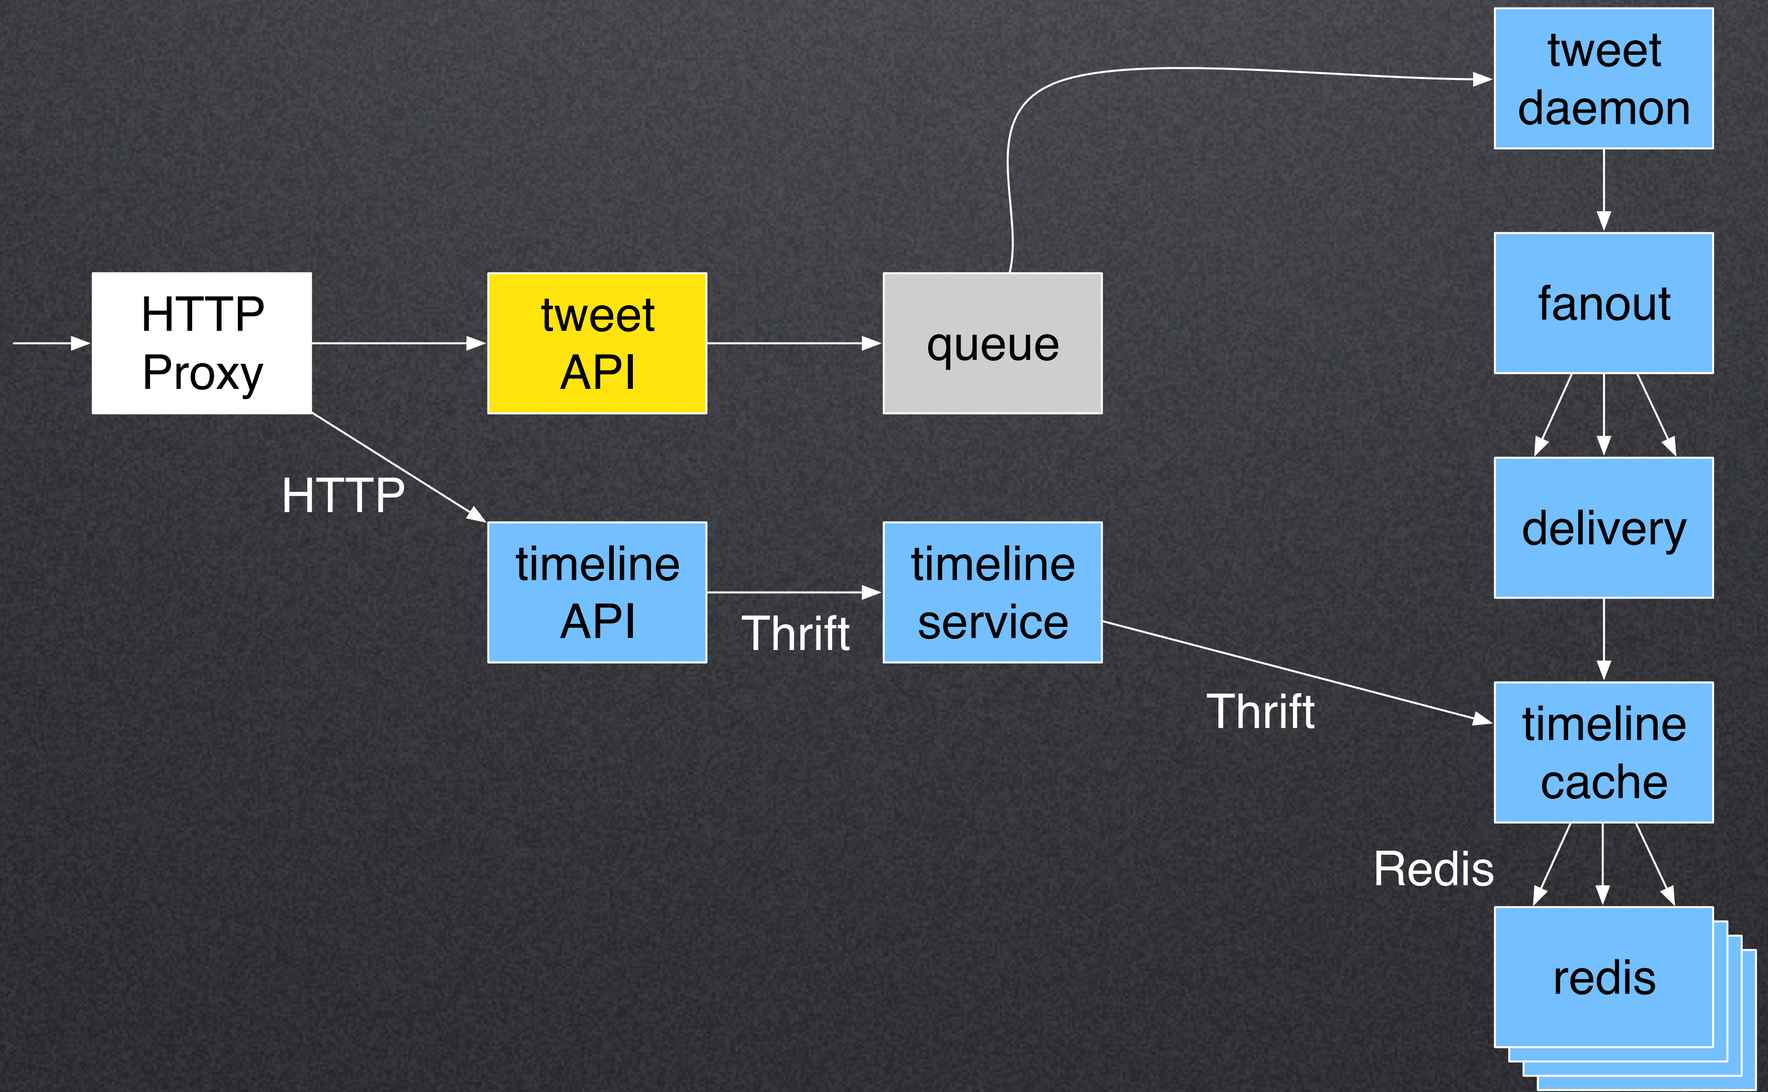
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
HTTP
timeline
API
Thrift
timeline
service
Thrift
timeline
cache
Redis
redis

Finagle - RPC Common Library

Components

- connection management
- protocol codecs
- transient error handling
- distributed tracing
- service discovery
- observability

Looking Forward

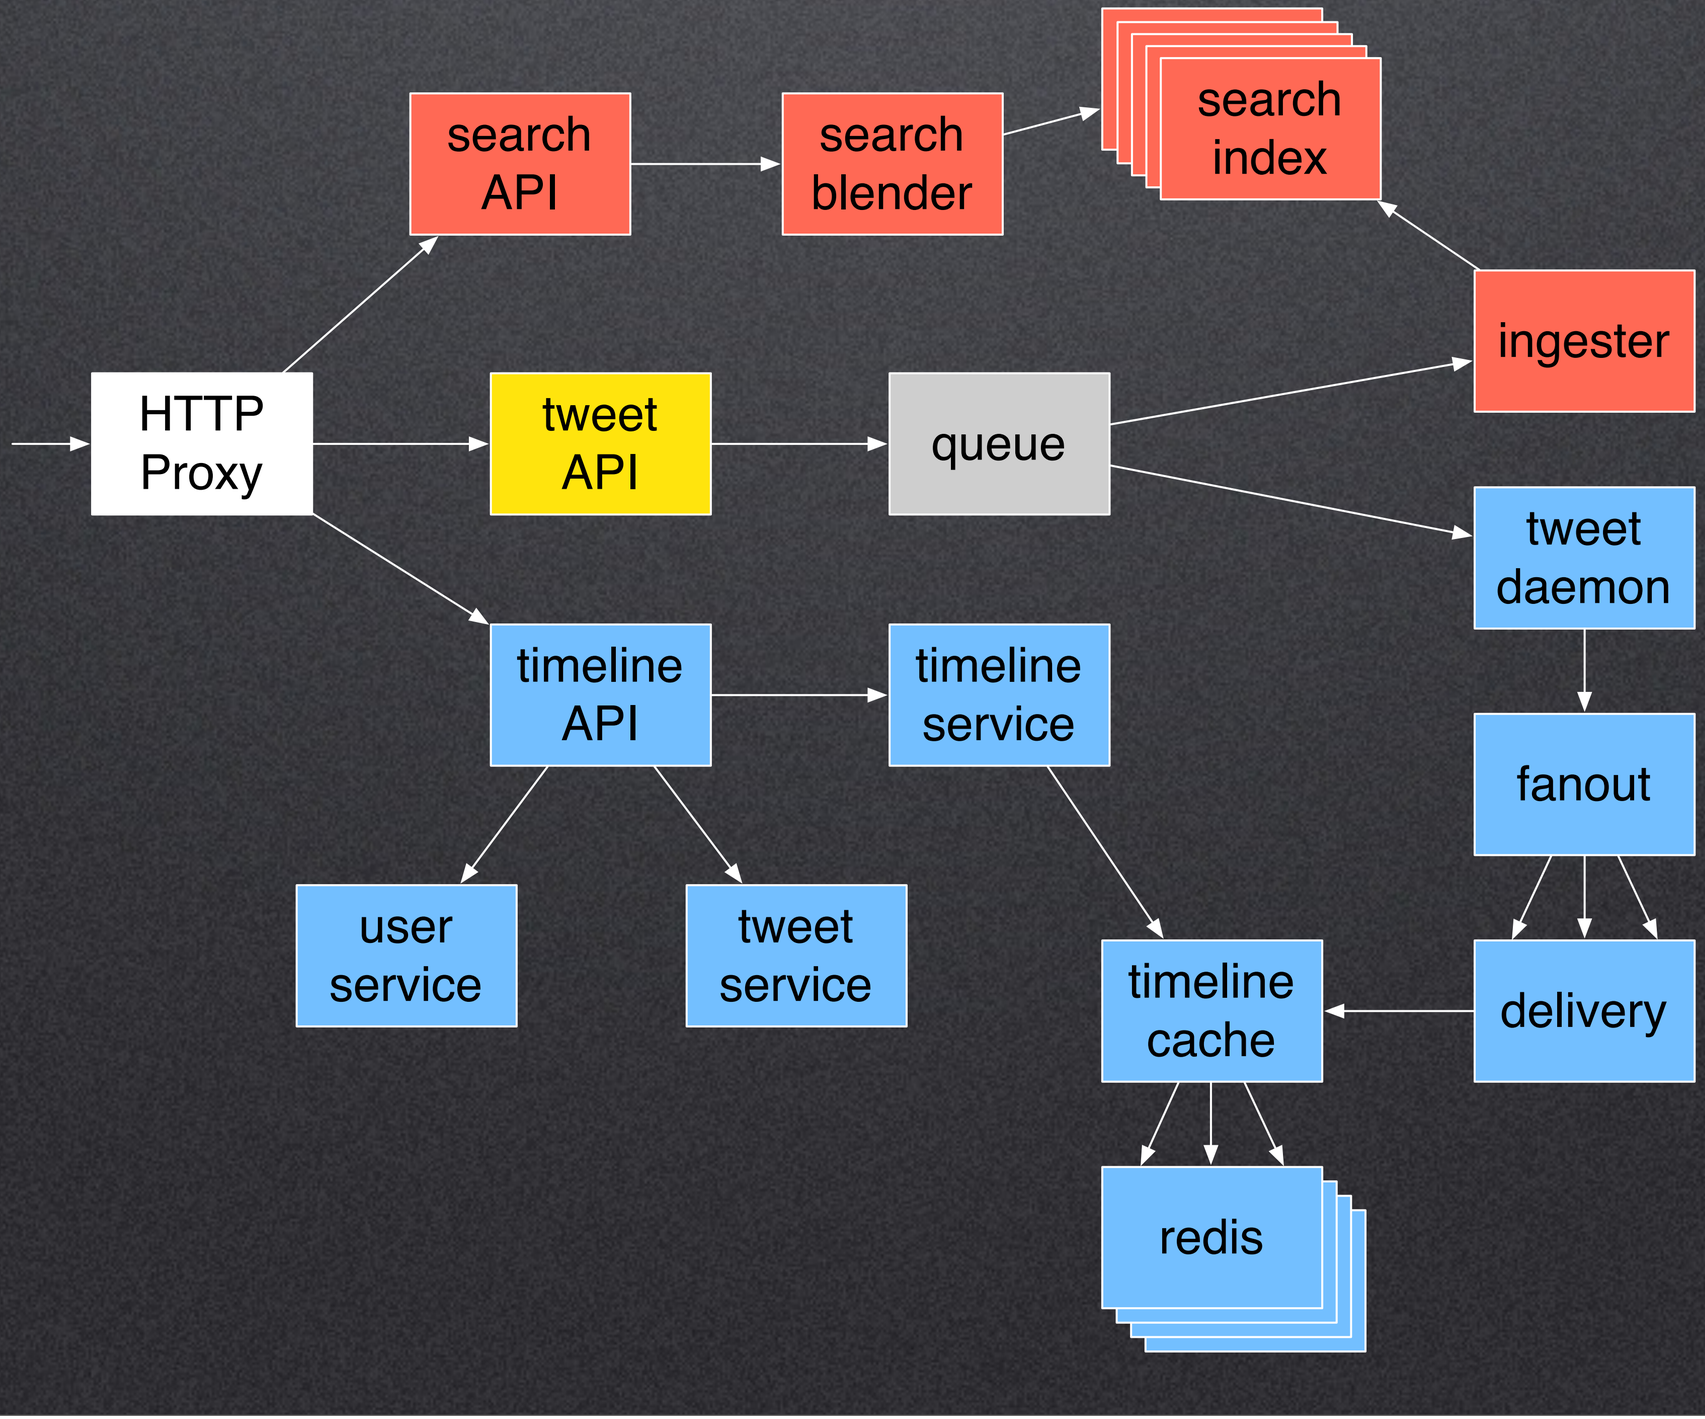
search API
search blender
search index
HTTP Proxy
tweet API
queue
ingester
tweet daemon
timeline API
timeline service
fanout
user service
tweet service
timeline cache
delivery
redis

search API

search blender

search index

HTTP Proxy

tweet API

queue

ingester

tweet daemon

timeline API

timeline service

fanout

user service

tweet service

timeline cache

delivery

redis

Case Study

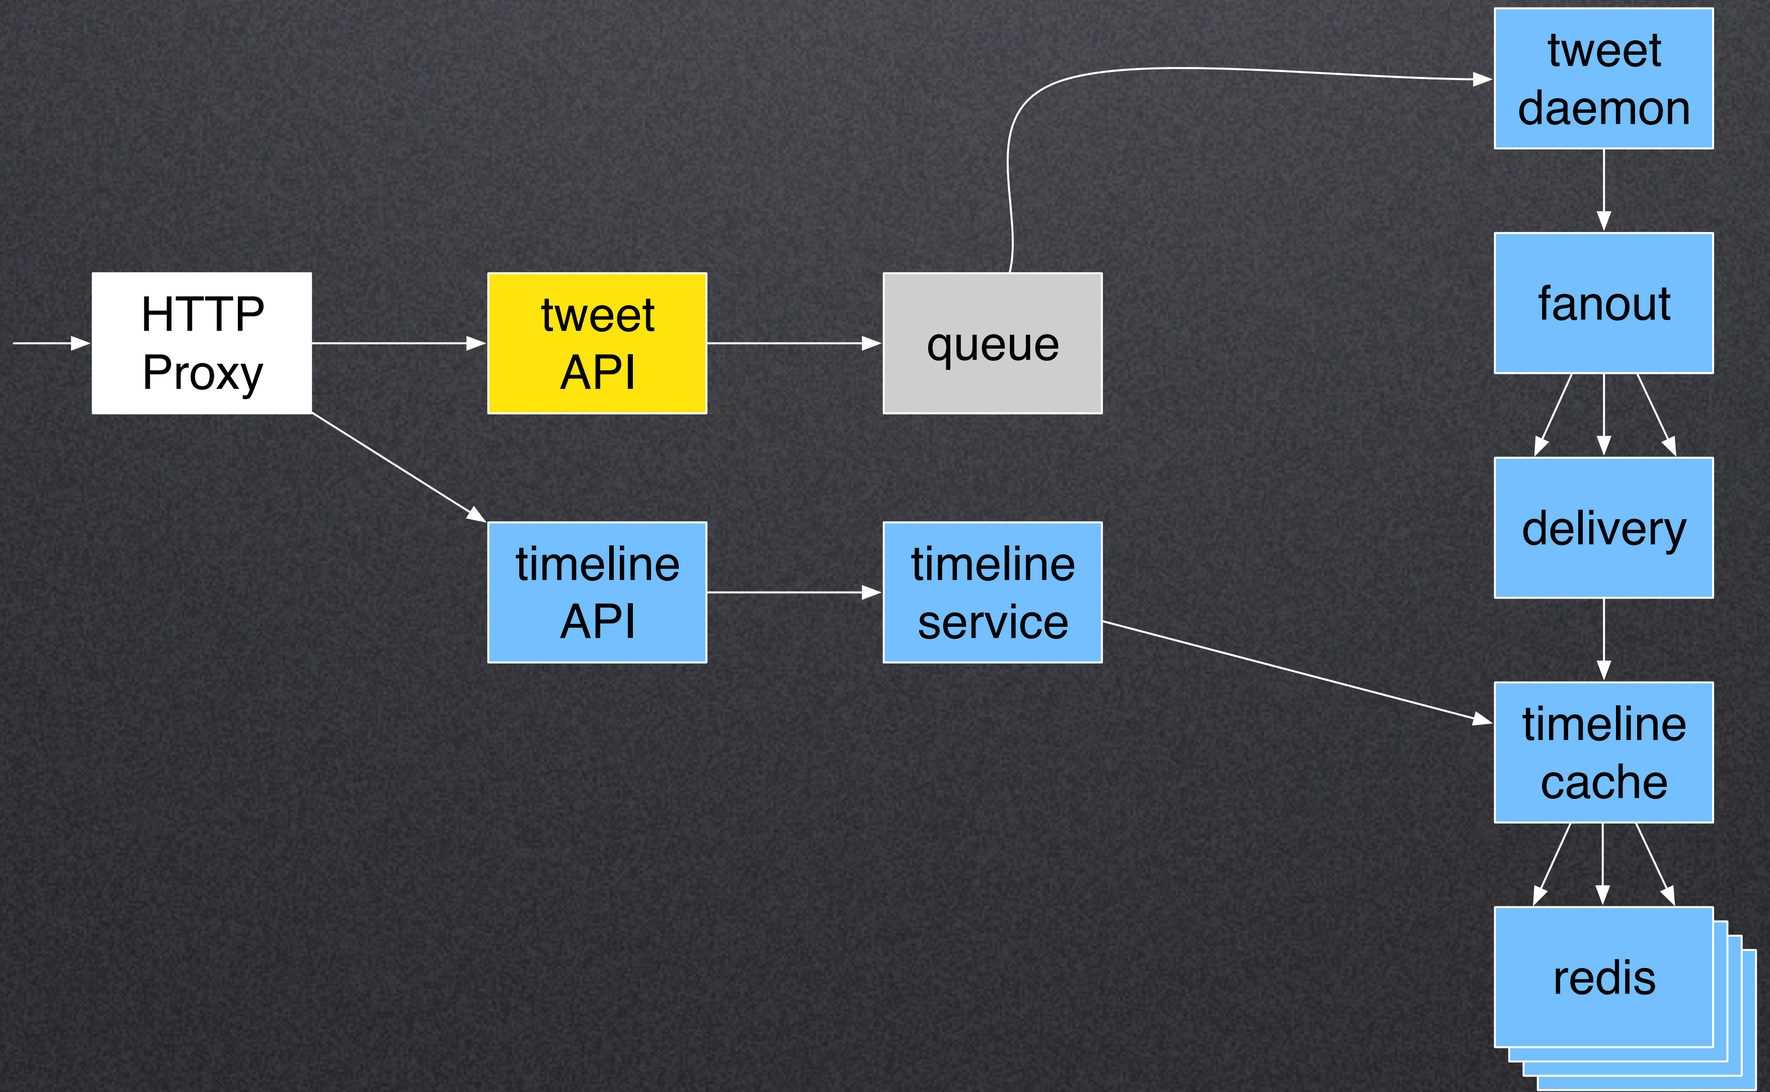
HTTP
Proxy
tweet
API
queue
tweet
daemon
fanout
delivery
timeline
API
timeline
service
timeline
cache
redis

Create
Tweet


Create Tweet
Load Followers

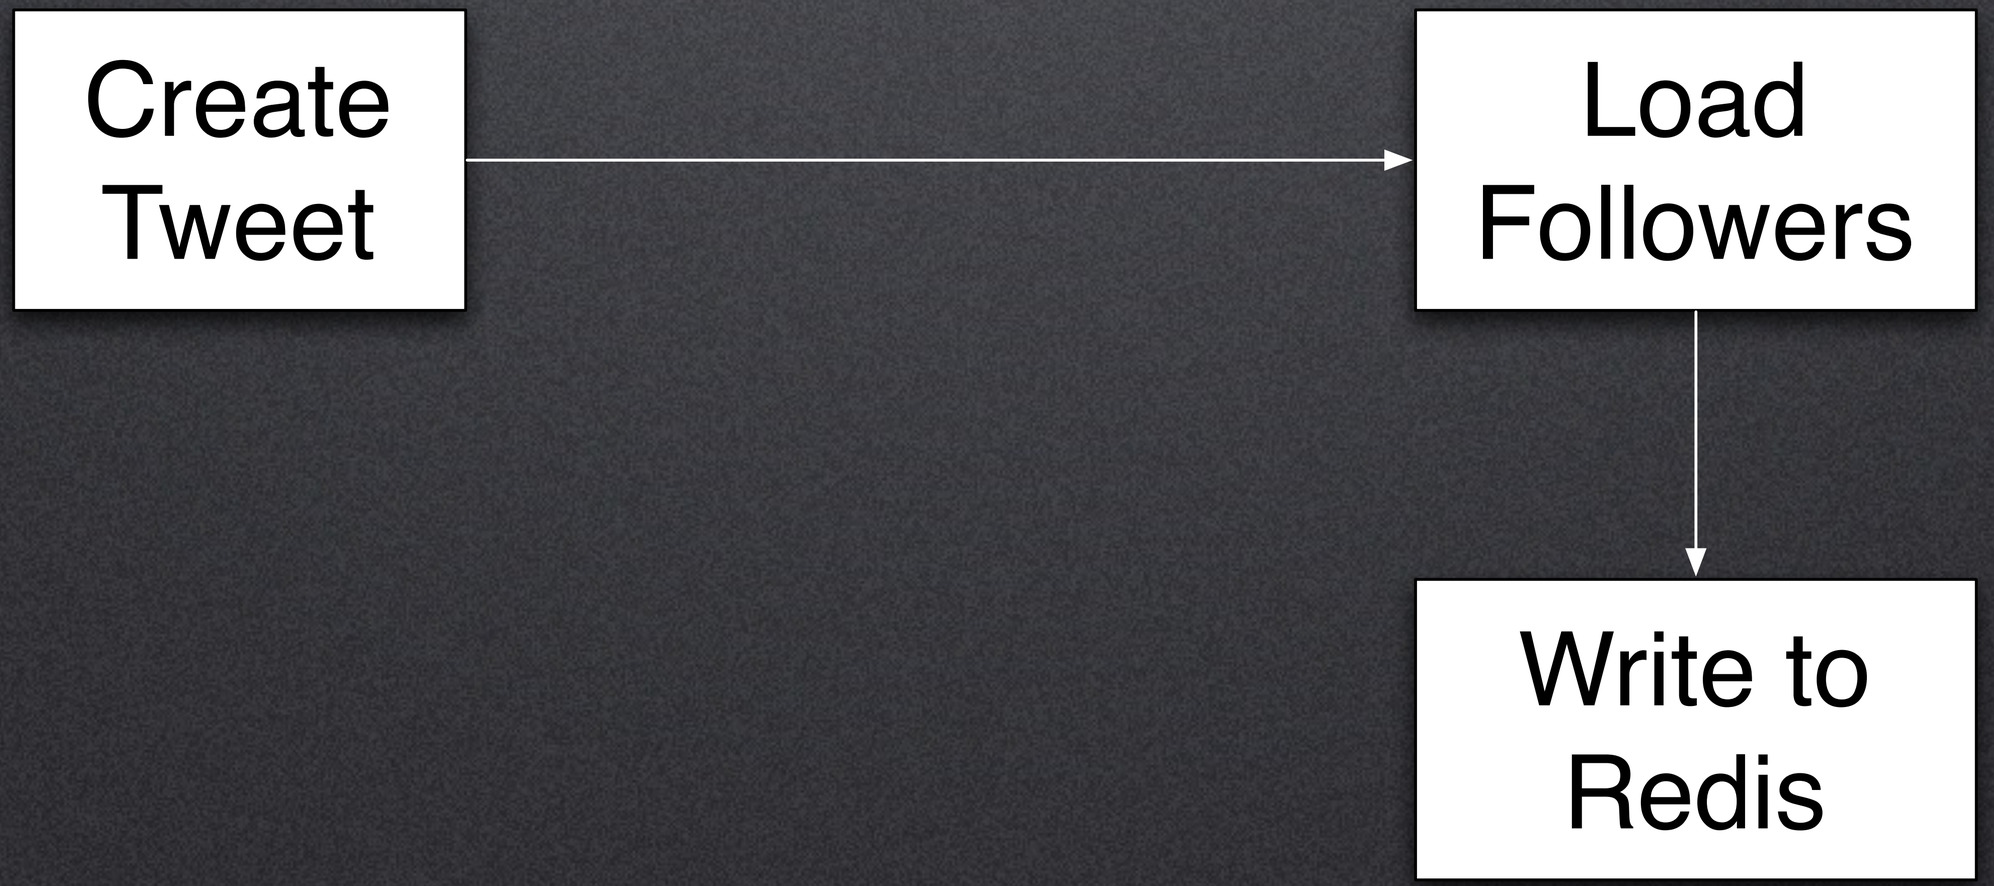
Create Tweet
Load Followers
Write to Redis

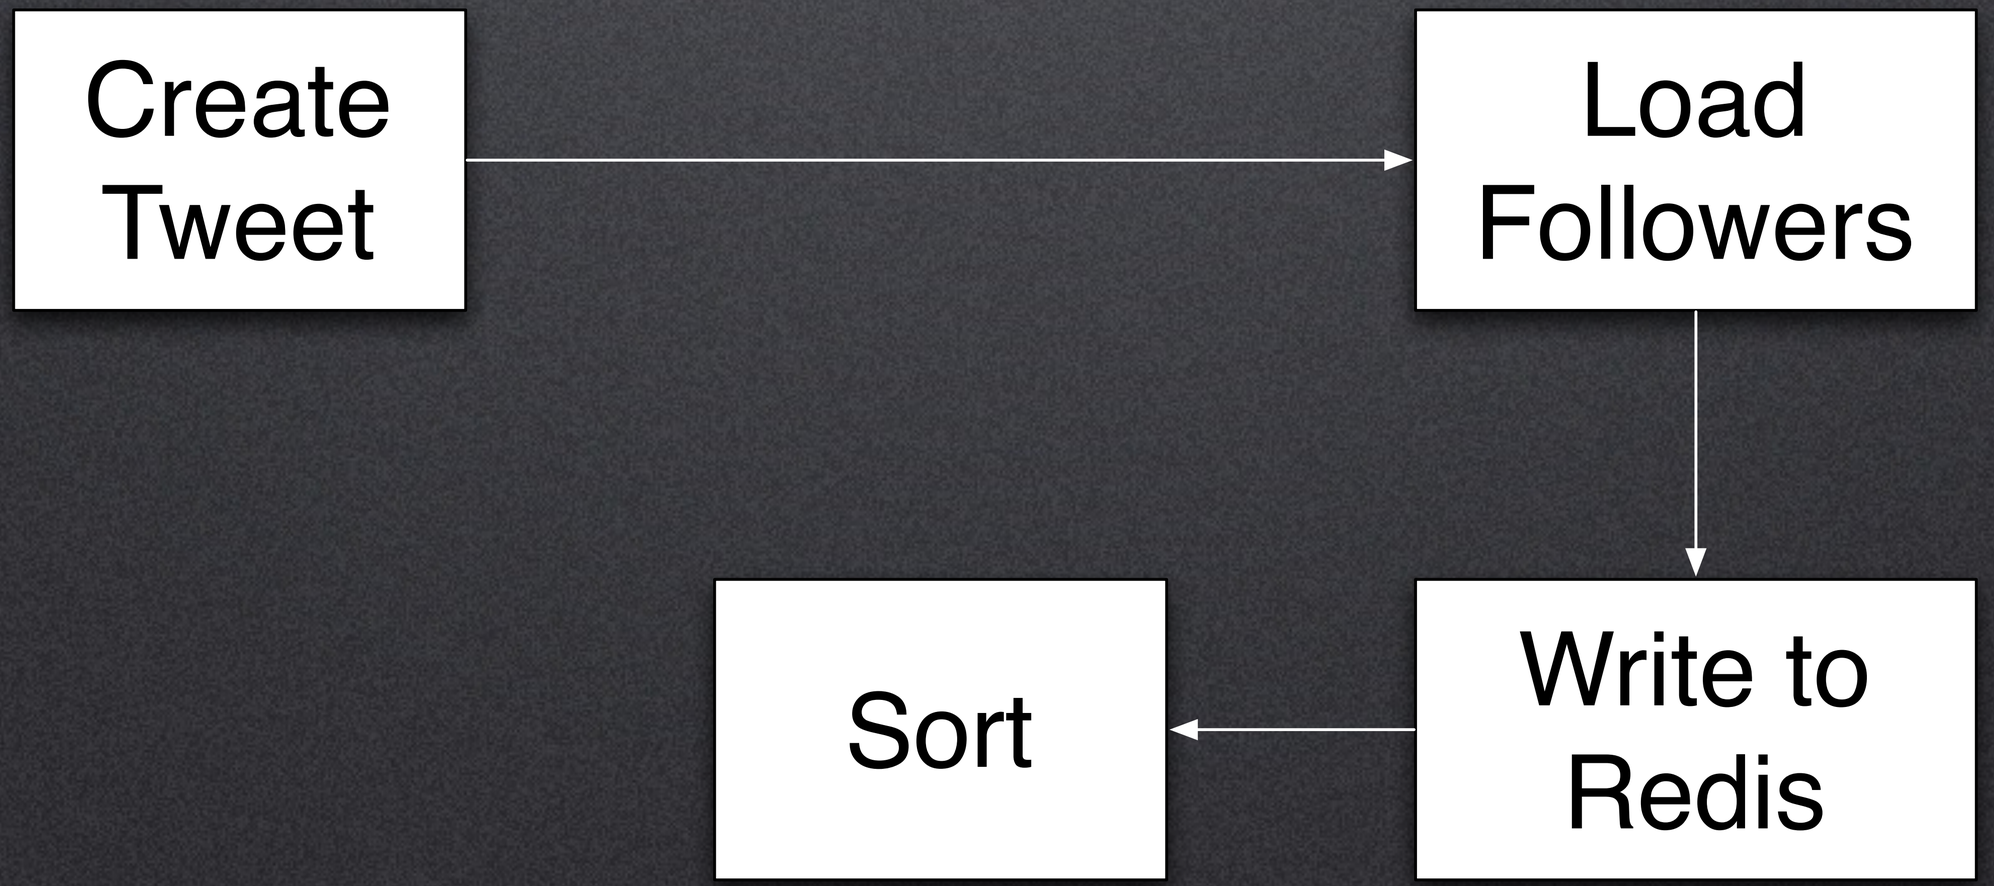
Create Tweet
Load Followers
Write to Redis
Sort

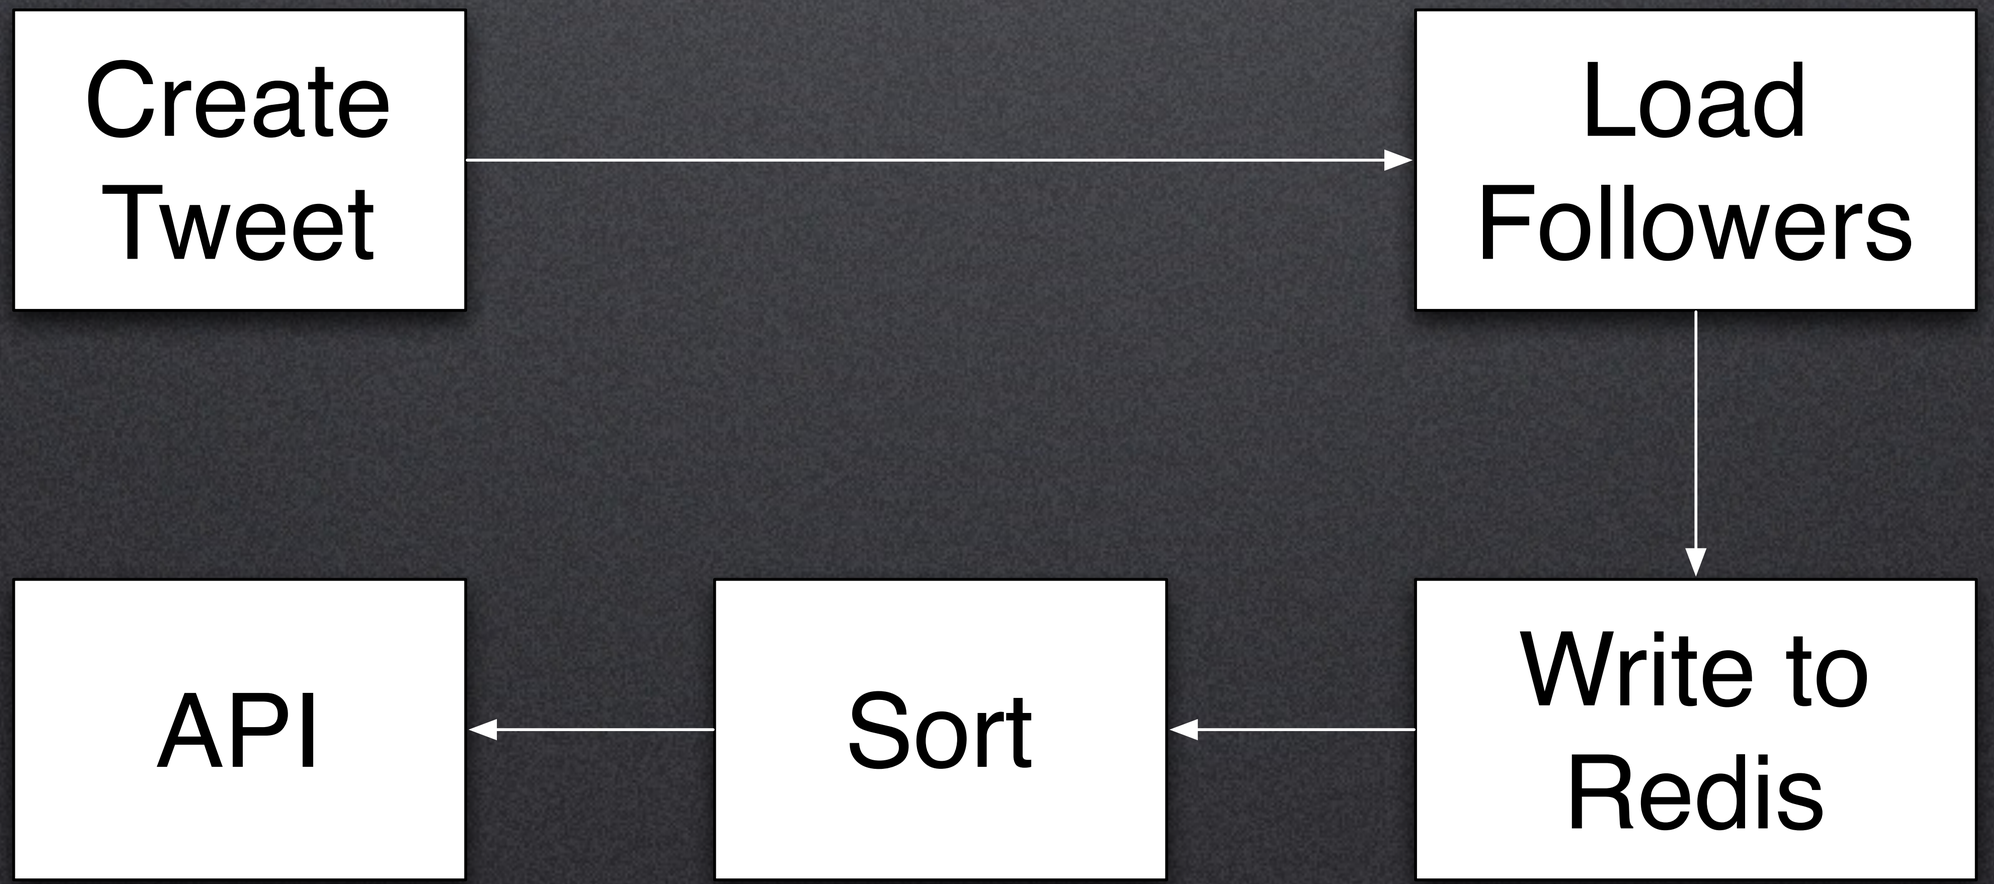
Create Tweet
Load Followers
Write to Redis
Sort
API

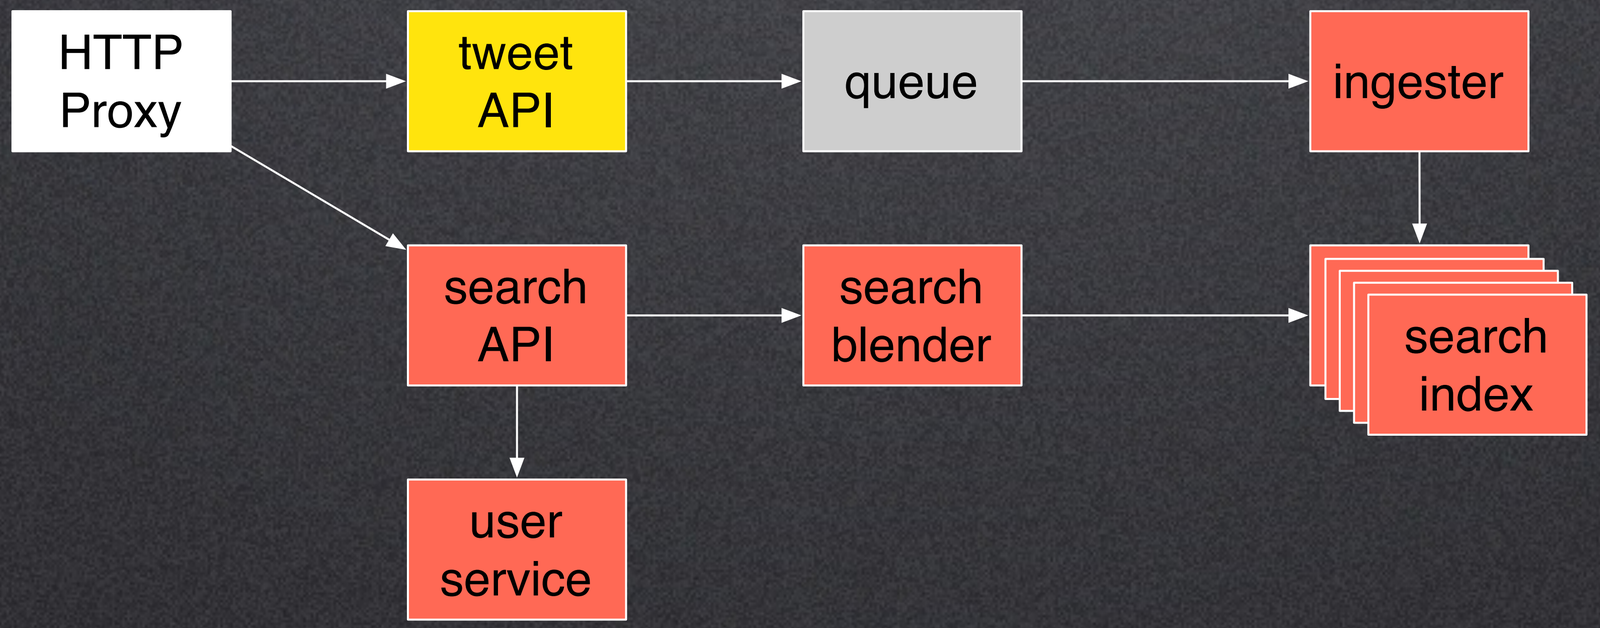
HTTP
Proxy
tweet
API
queue
ingester
search
API
search
blender
search
index
user
service

Create
Tweet


Create Tweet
Tokenize

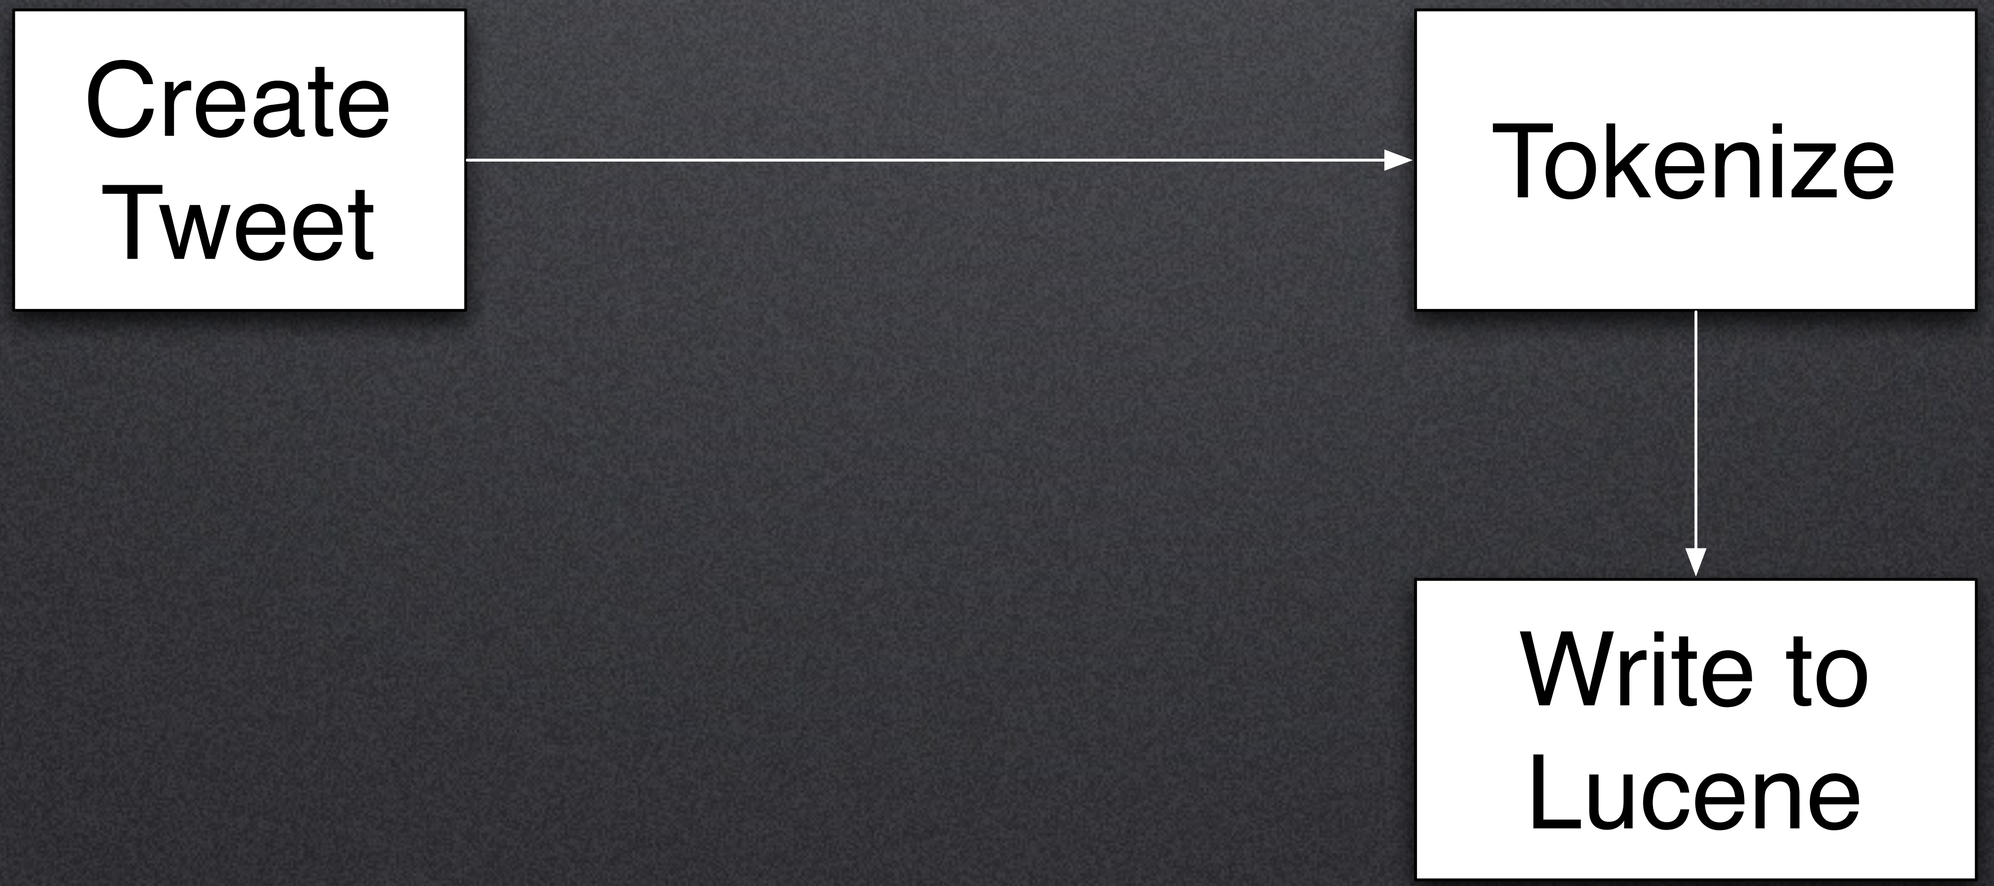
Create Tweet
Tokenize
Write to Lucene

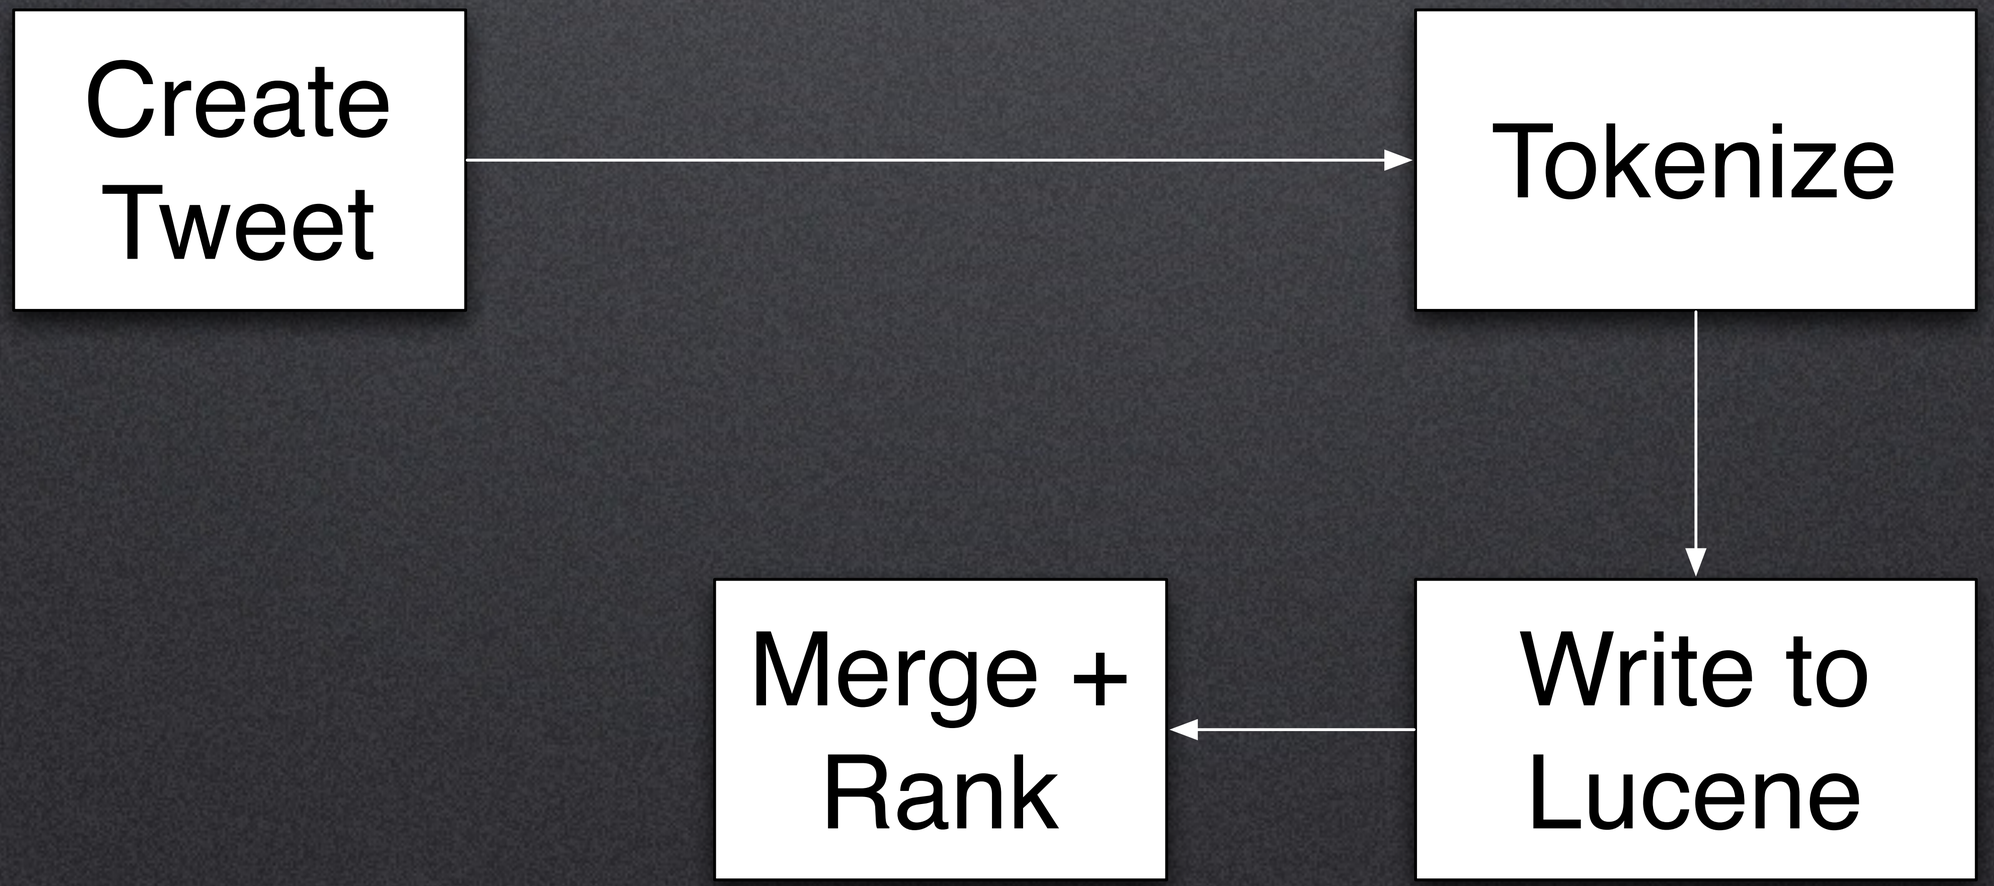
Create Tweet
Tokenize
Write to Lucene
Merge + Rank

Create Tweet

Tokenize

Write to Lucene

Merge + Rank

Search Query

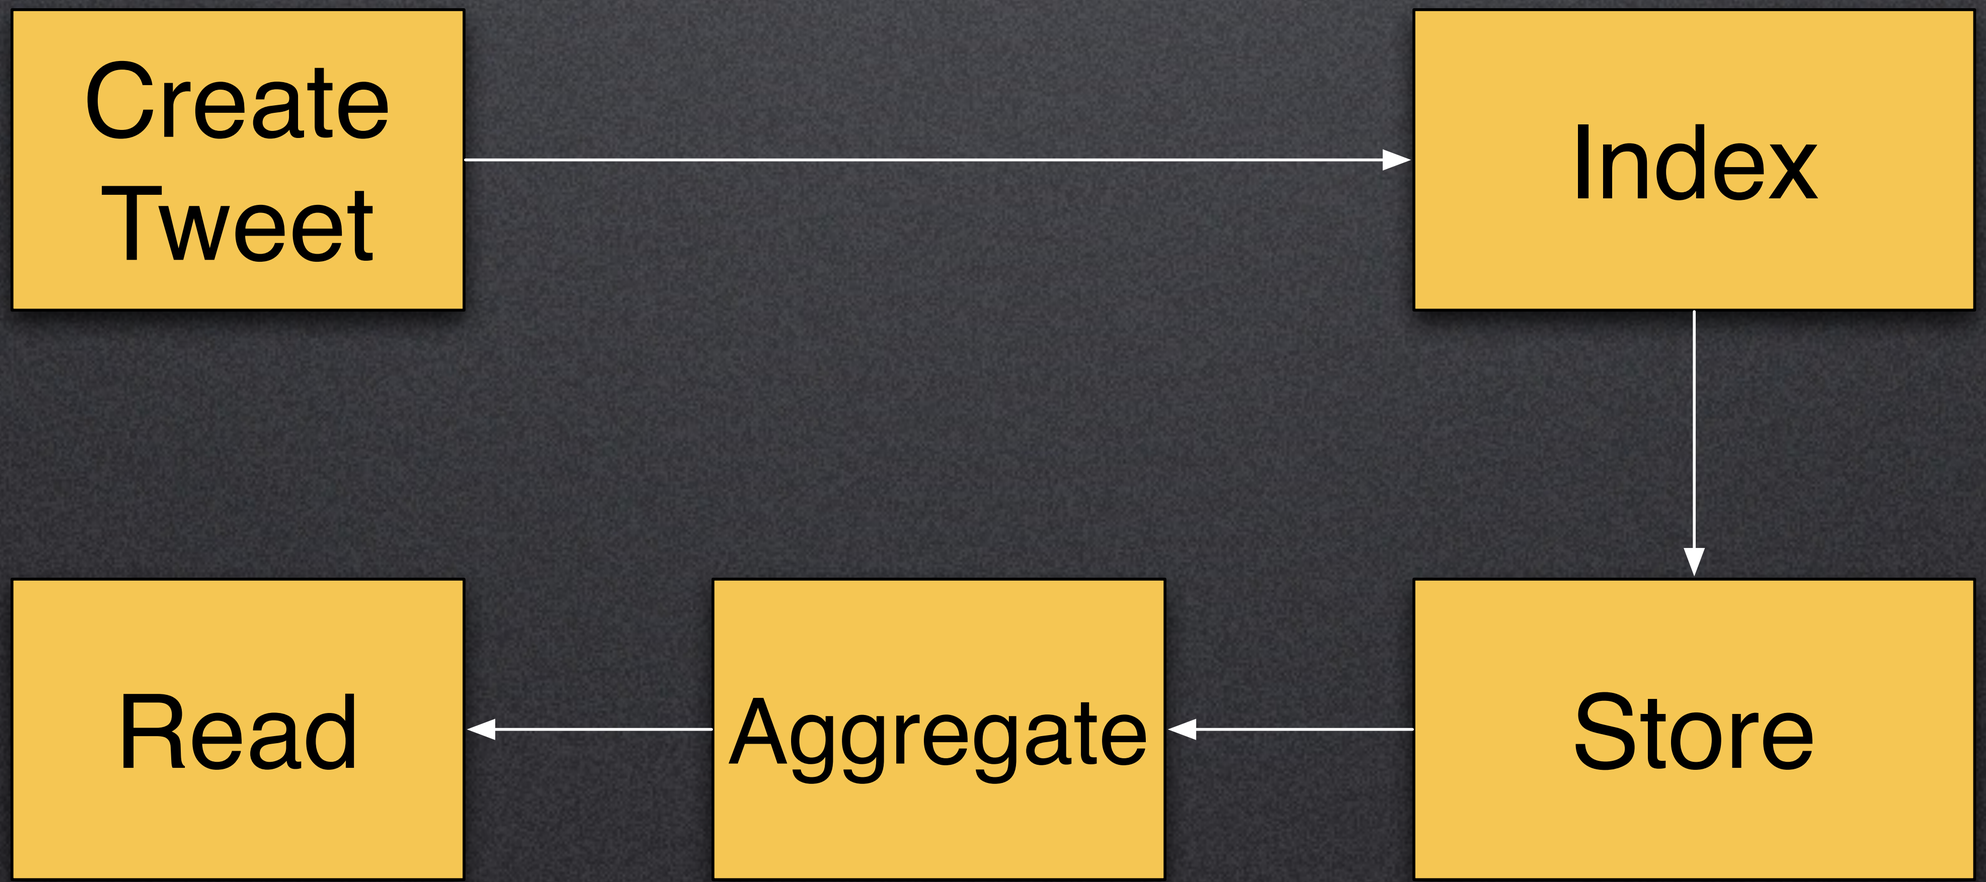
Create Tweet
Index
Store
Aggregate
Read

@joe tweets “hello world”
@joe followed by @sam and @frank

search index -> [“hello”,“world”]
fanout index -> [@sam, @frank]

@joe: hello world

@sam

@frank

Redis

world

hello

Lucene


@sam
Redis
1726371
1415044
1056831
98124

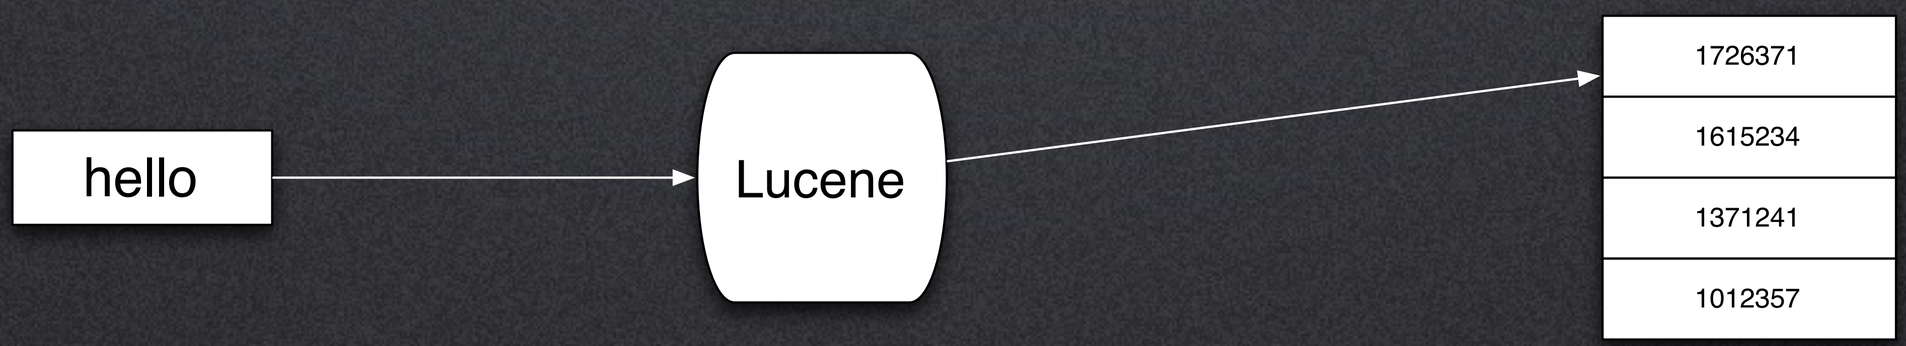
hello
Lucene
1726371
1615234
1371241
1012357

User Intent

"QCon London"

@sam's
home timeline

User Intent

Query Expansion

"QCon London"

"QCon"
AND
"London"

@sam's
home timeline

User Intent

Query Expansion

"QCon London"

"QCon"
AND
"London"

@sam's
home timeline

"home_timeline:@sam"

User Intent

Query Expansion

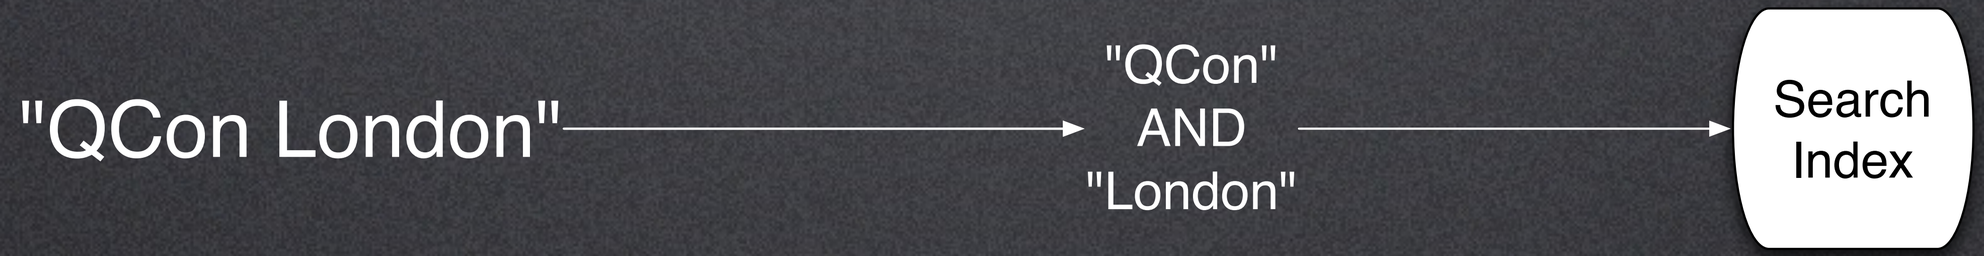

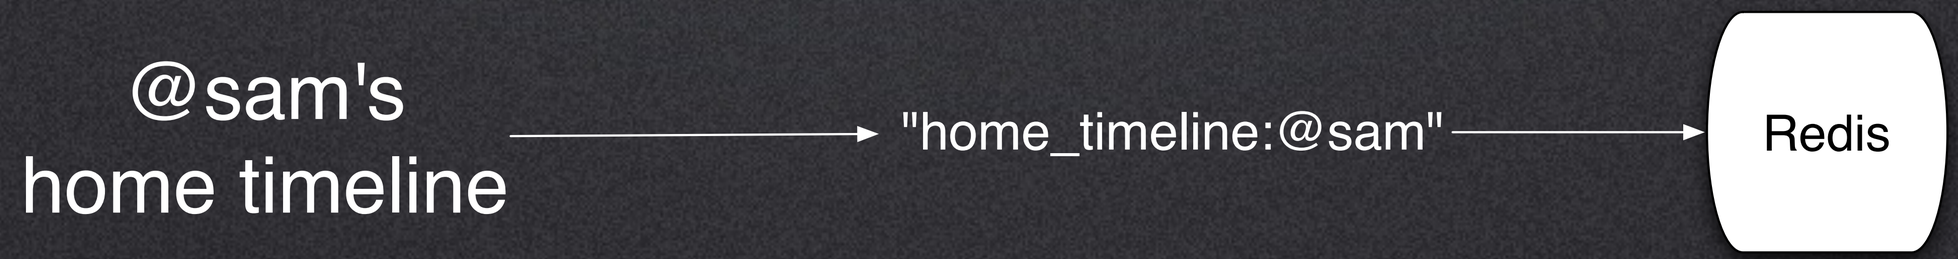

User Intent

Query Expansion

"QCon London"

"QCon"
AND
"London"

Search
Index

union of timelines

@sam's
home timeline

"user_timeline:@joe"
OR
"user_timeline:@frank"

Redis

Fan-out vs Fan-in

Delivery Type	Read-time Cost	Write-time Cost
Fan-out		
Fan-in		
???		

Delivery Type	Read-time Cost	Write-time Cost
Fan-out	Low O(1)	
Fan-in		
???		

Delivery Type	Read-time Cost	Write-time Cost
Fan-out	Low O(1)	Low to High O(N)
Fan-in		
???		

Delivery Type	Read-time Cost	Write-time Cost
Fan-out	Low O(1)	Low to High O(N)
Fan-in	Low to High O(N)	
???		

Delivery Type	Read-time Cost	Write-time Cost
Fan-out	Low O(1)	Low to High O(N)
Fan-in	Low to High O(N)	Low O(1)
???		

Delivery Type	Read-time Cost	Write-time Cost
Fan-out	Low O(1)	Low to High O(N)
Fan-in	Low to High O(N)	Low O(1)
Hybrid		

Delivery Type	Read-time Cost	Write-time Cost
Fan-out	Low O(1)	Low to High O(N)
Fan-in	Low to High O(N)	Low O(1)
Hybrid	Low	

Delivery Type	Read-time Cost	Write-time Cost
Fan-out	Low O(1)	Low to High O(N)
Fan-in	Low to High O(N)	Low O(1)
Hybrid	Low	Low to Medium

@sam's
home timeline

"home_timeline:@sam"

Redis

@sam's
home timeline

"home_timeline:@sam"
OR
"user_timeline:@barackobama"

Redis

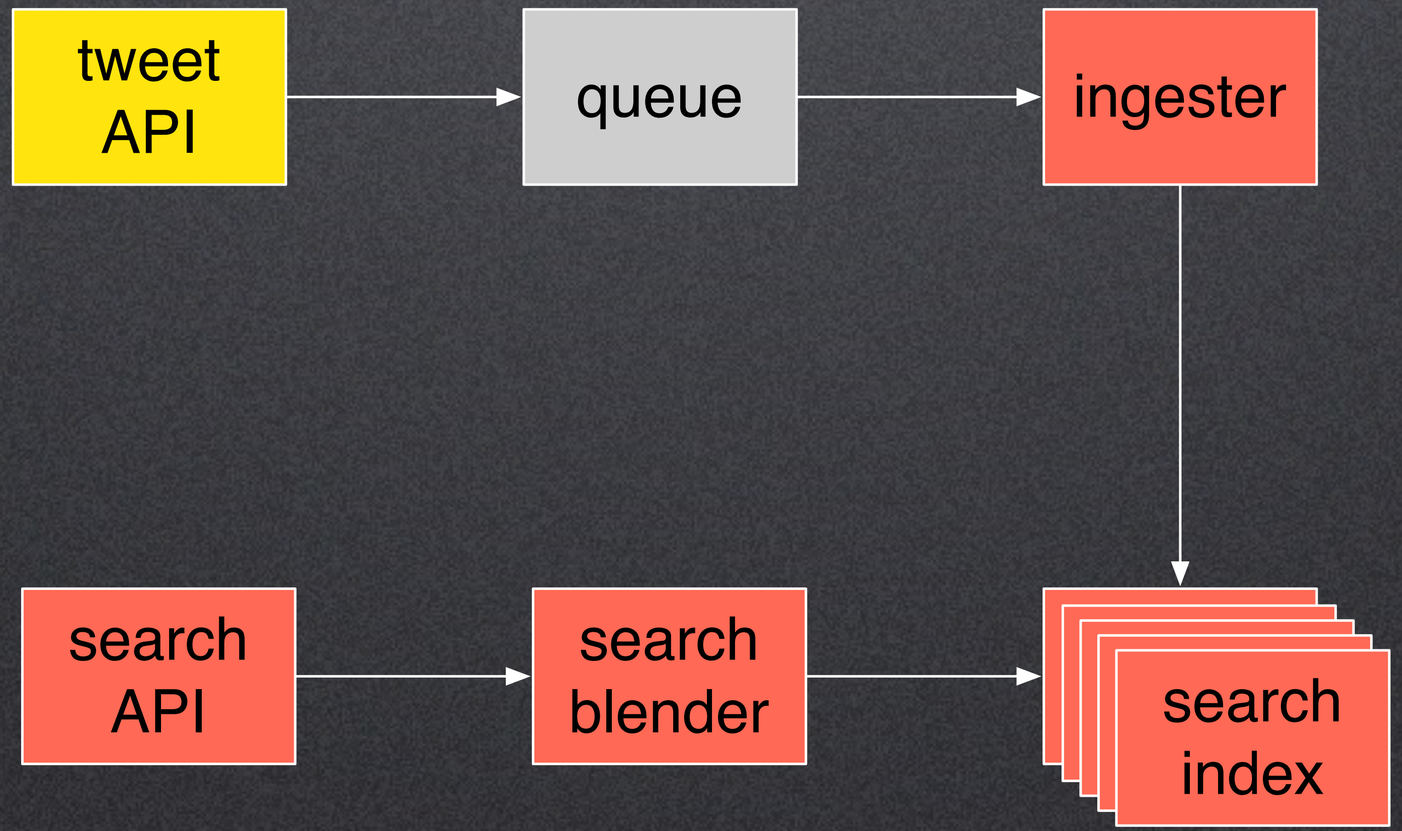
tweet
API
queue
ingester
search
API
search
blender
search
index

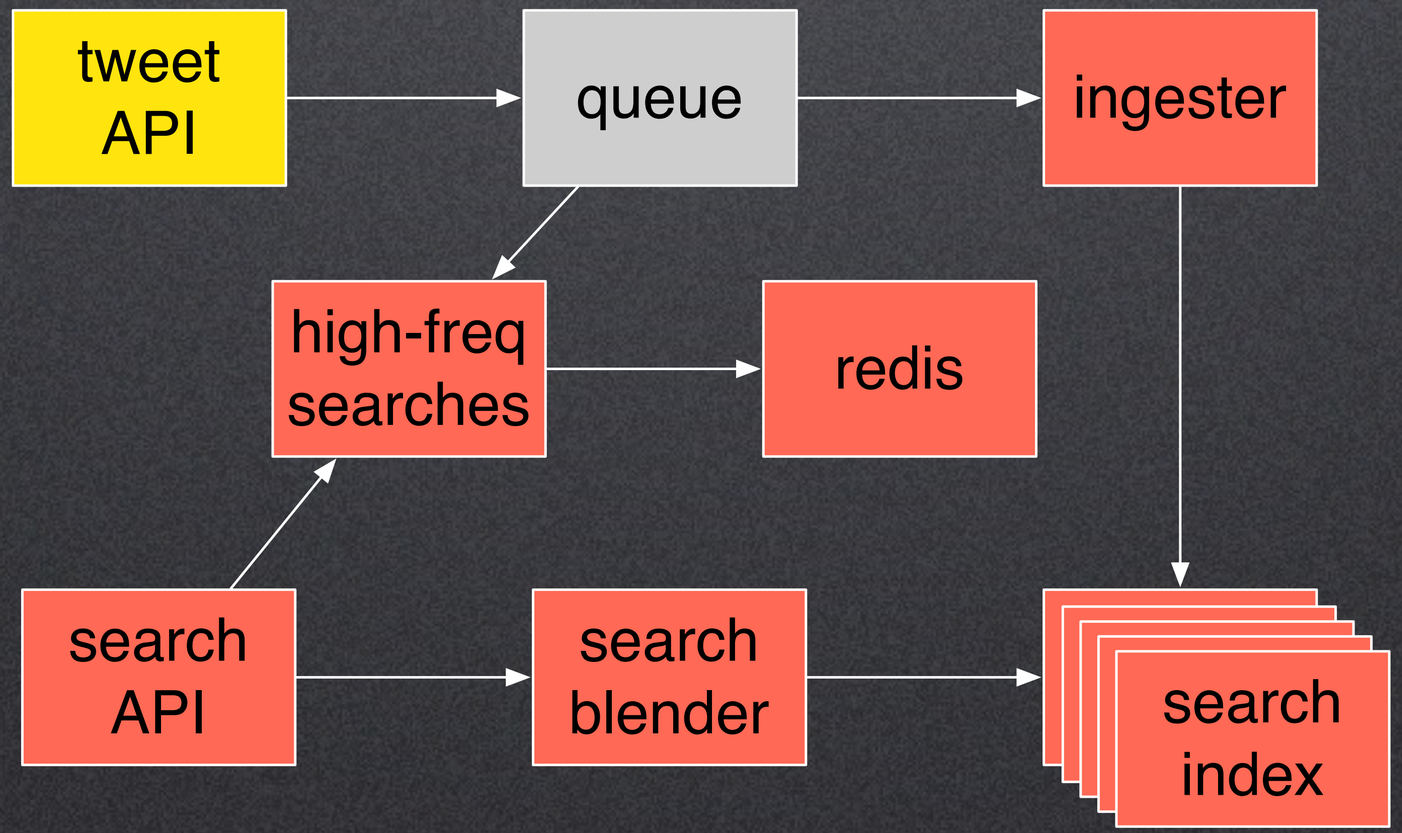
tweet
API
queue
ingester
high-freq
searches
redis
search
API
search
blender
search
index

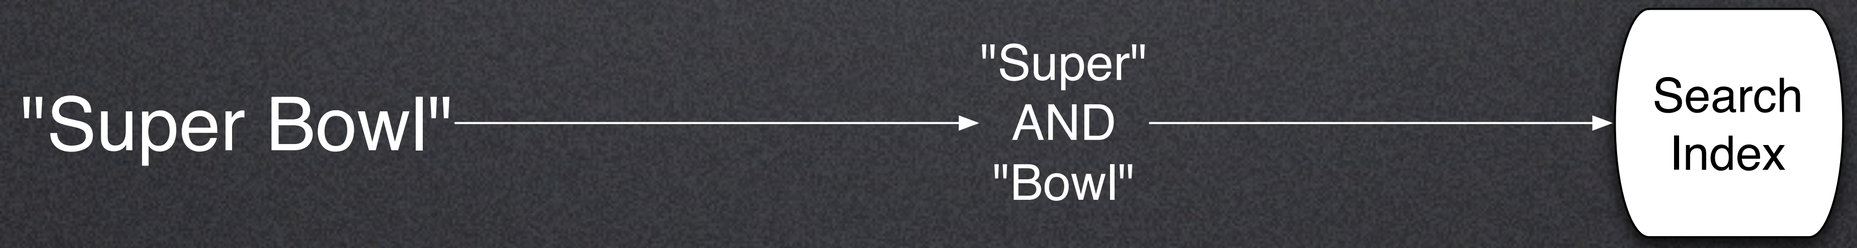
"Super Bowl"
"Super"
AND
"Bowl"
Search
Index

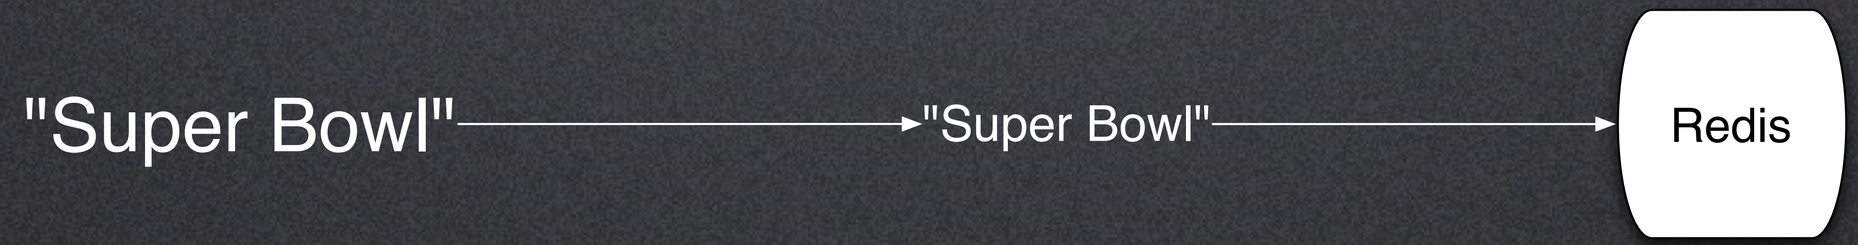
"Super Bowl"
"Super Bowl"
Redis

Redis vs Search Index

Redis Eviction Model

home_timeline:14653230

home_tiimeline:6014086

home_timeline:12512351

home_timeline:15123512

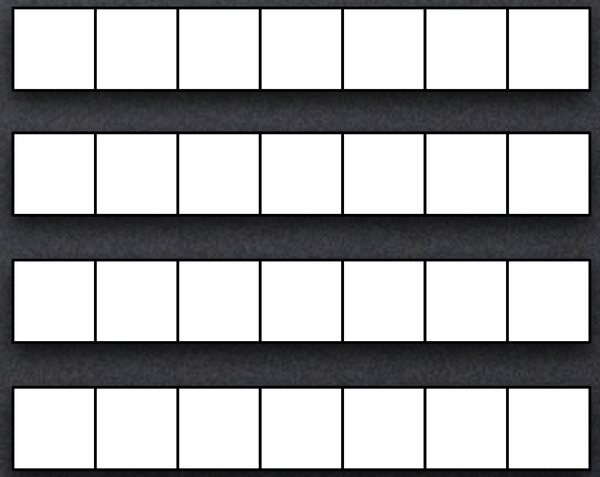

Redis Eviction Model

home_timeline:14653230

home_tiimeline:6014086

home_timeline:12512351

home_timeline:15123512

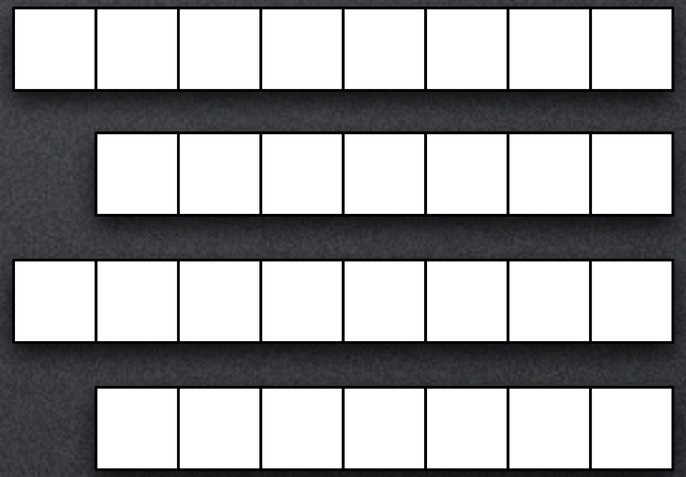

Redis Eviction Model

home_timeline:14653230

home_tiimeline:6014086

home_timeline:12512351

home_timeline:15123512

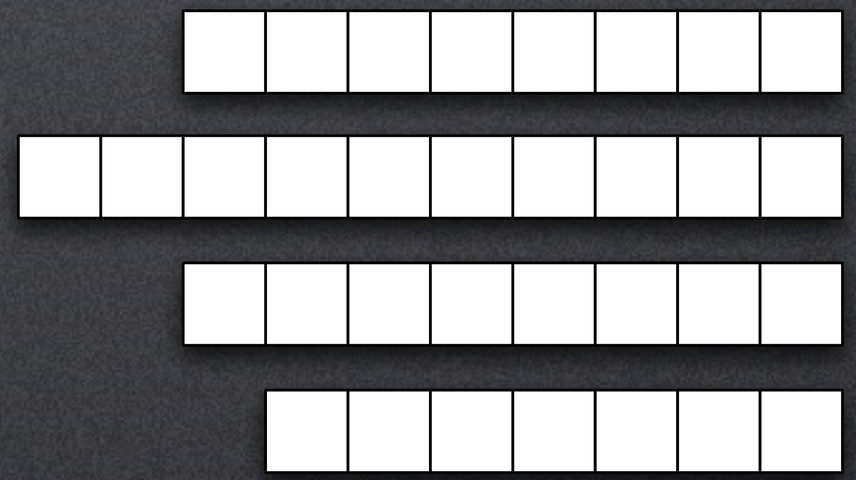

Redis Eviction Model

home_timeline:14653230

home_tiimeline:6014086

home_timeline:12512351

home_timeline:15123512

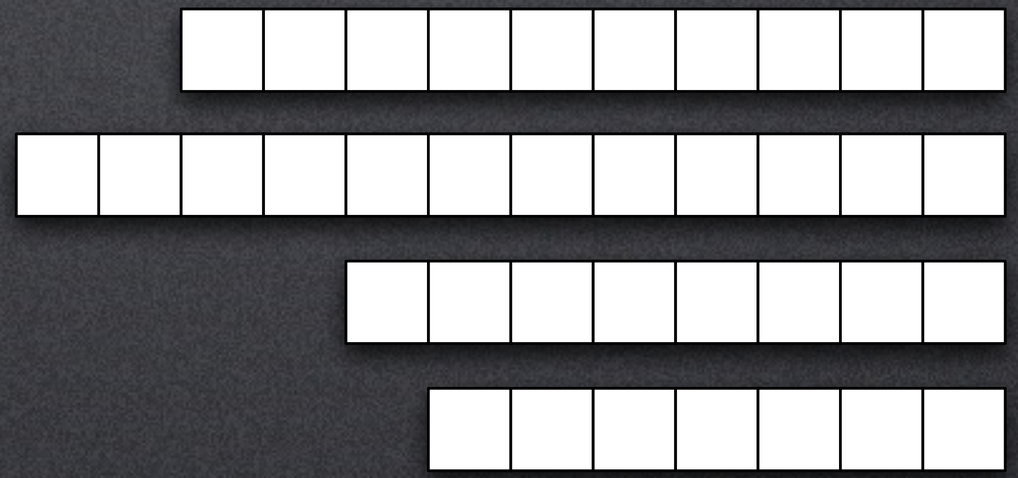

Redis Eviction Model

home_timeline:14653230

home_tiimeline:6014086

~~home_timeline:12512351~~

home_timeline:15123512

Redis Eviction Model

home_timeline:14653230

home_tiimeline:6014086

home_timeline:15123512

Search Index Eviction Model

QCon

London

hello

world

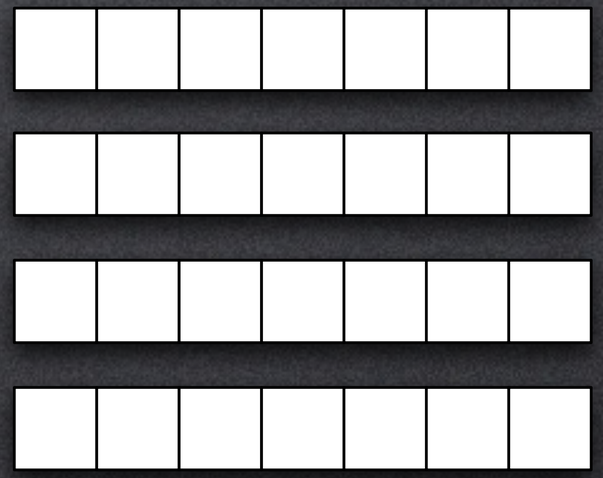

Search Index Eviction Model

QCon

London

hello

world

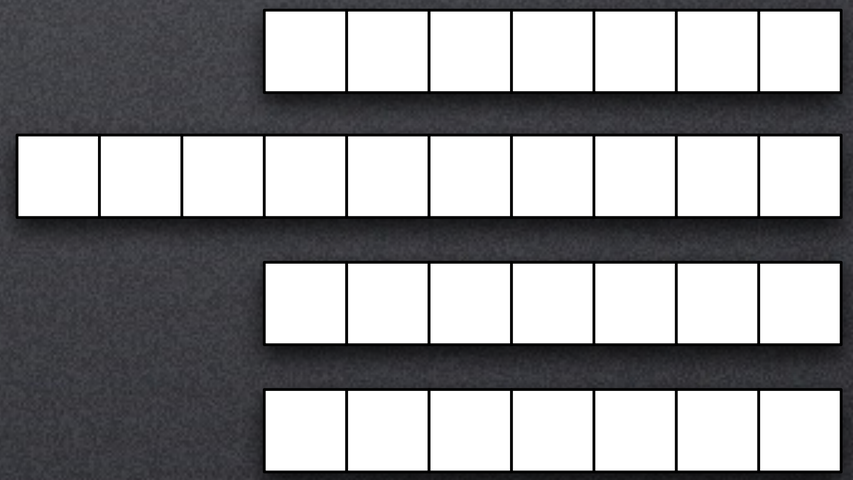

Search Index Eviction Model

QCon

London

hello

world

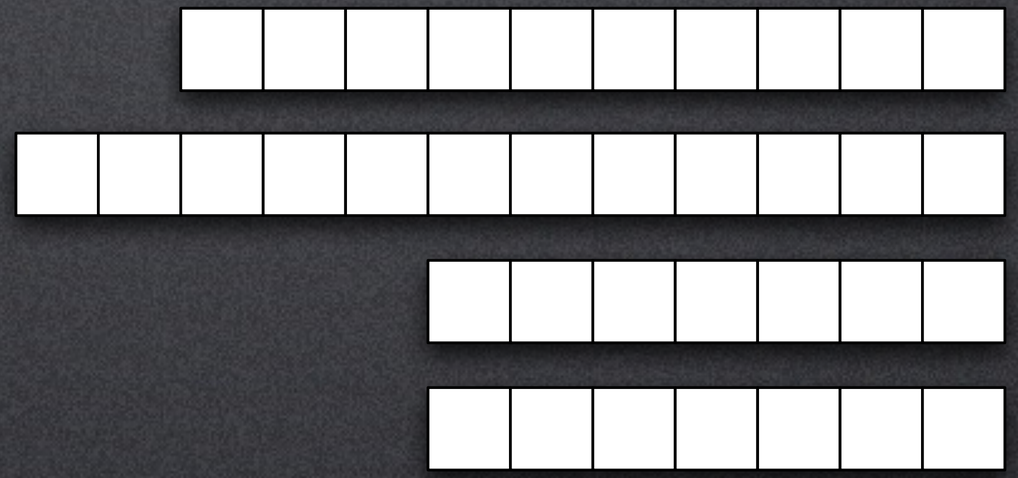

Search Index Eviction Model

QCon

London

hello

world

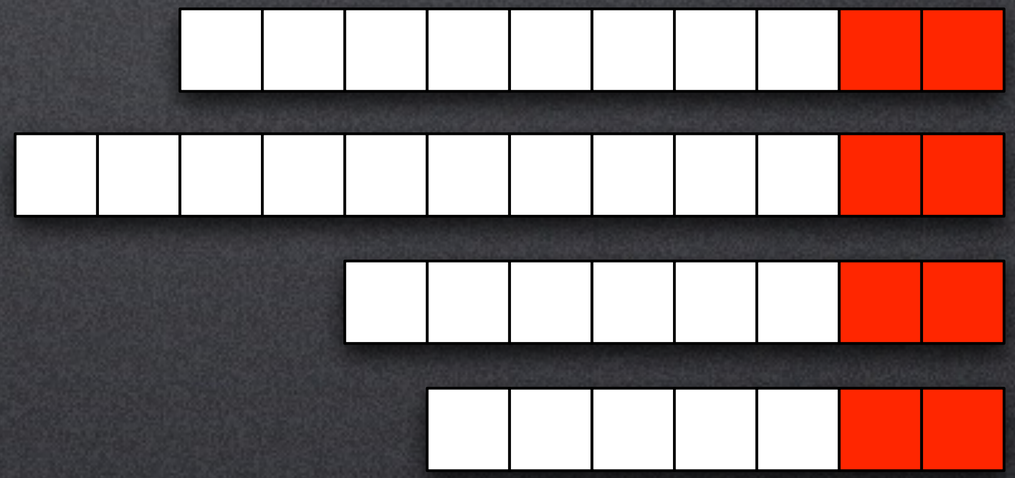

Search Index Eviction Model

QCon

London

hello

world

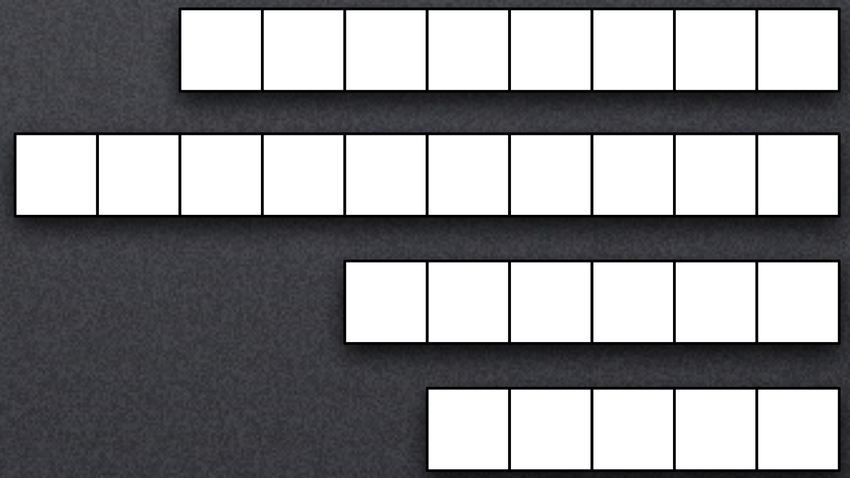

Conclusion

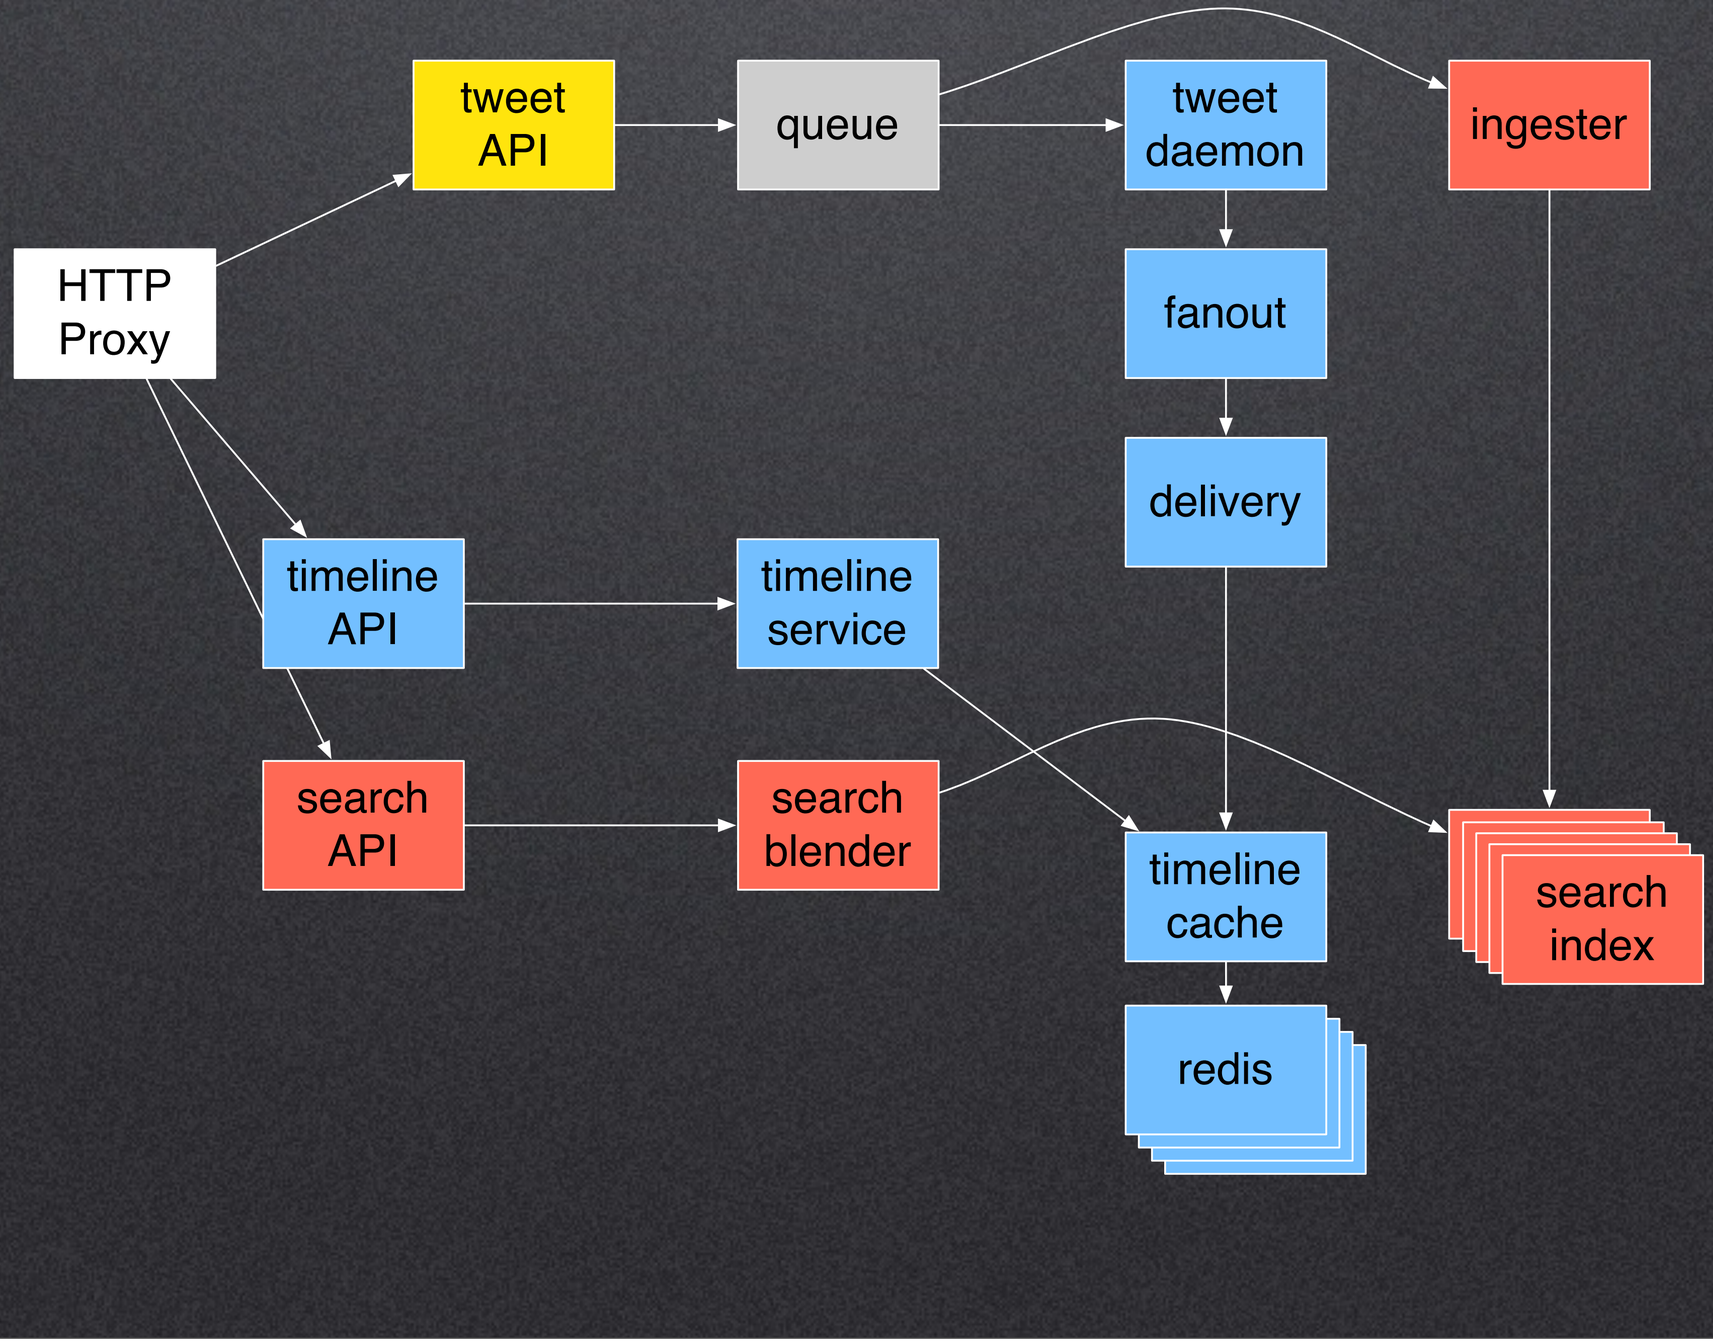
tweet API
queue
tweet daemon
ingester
HTTP Proxy
fanout
delivery
timeline API
timeline service
search API
search blender
timeline cache
search index
redis

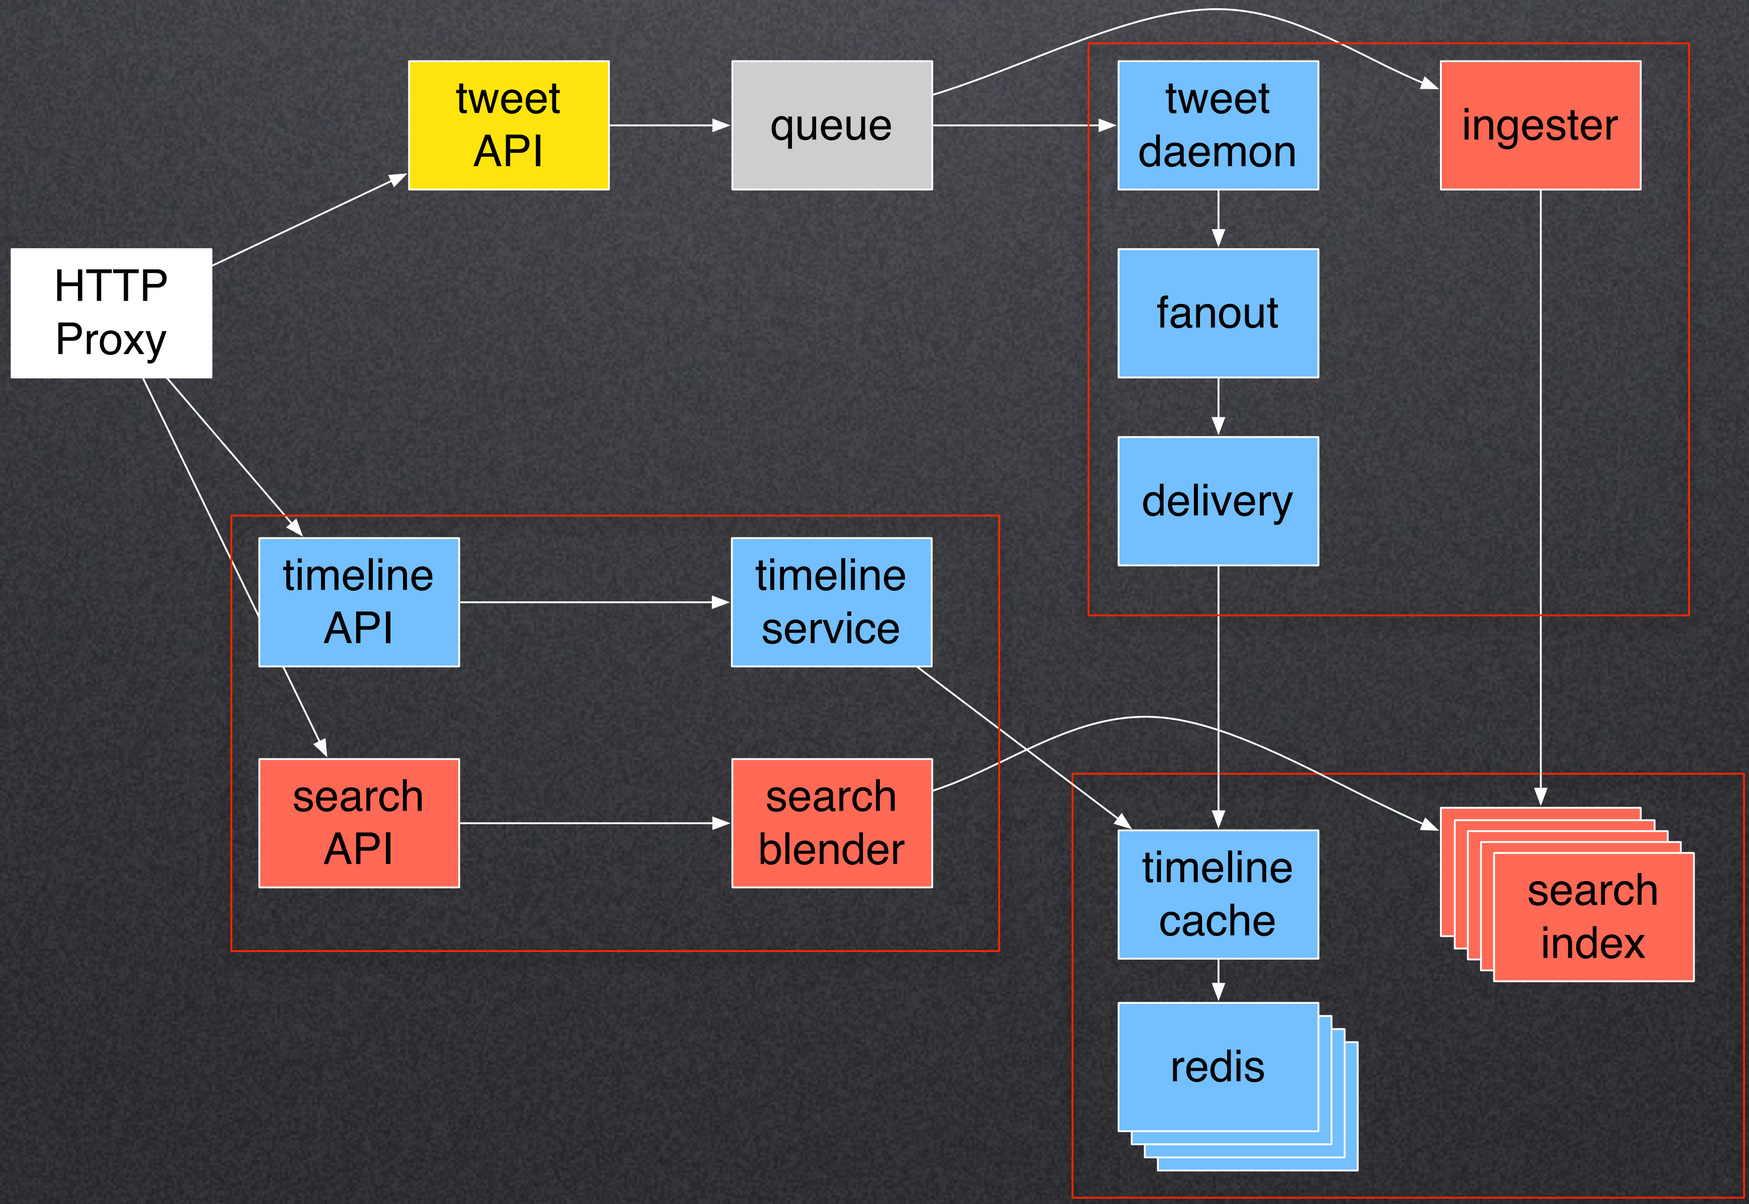
tweet API
queue
tweet daemon
ingester
HTTP Proxy
fanout
delivery
timeline API
timeline service
search API
search blender
timeline cache
search index
redis

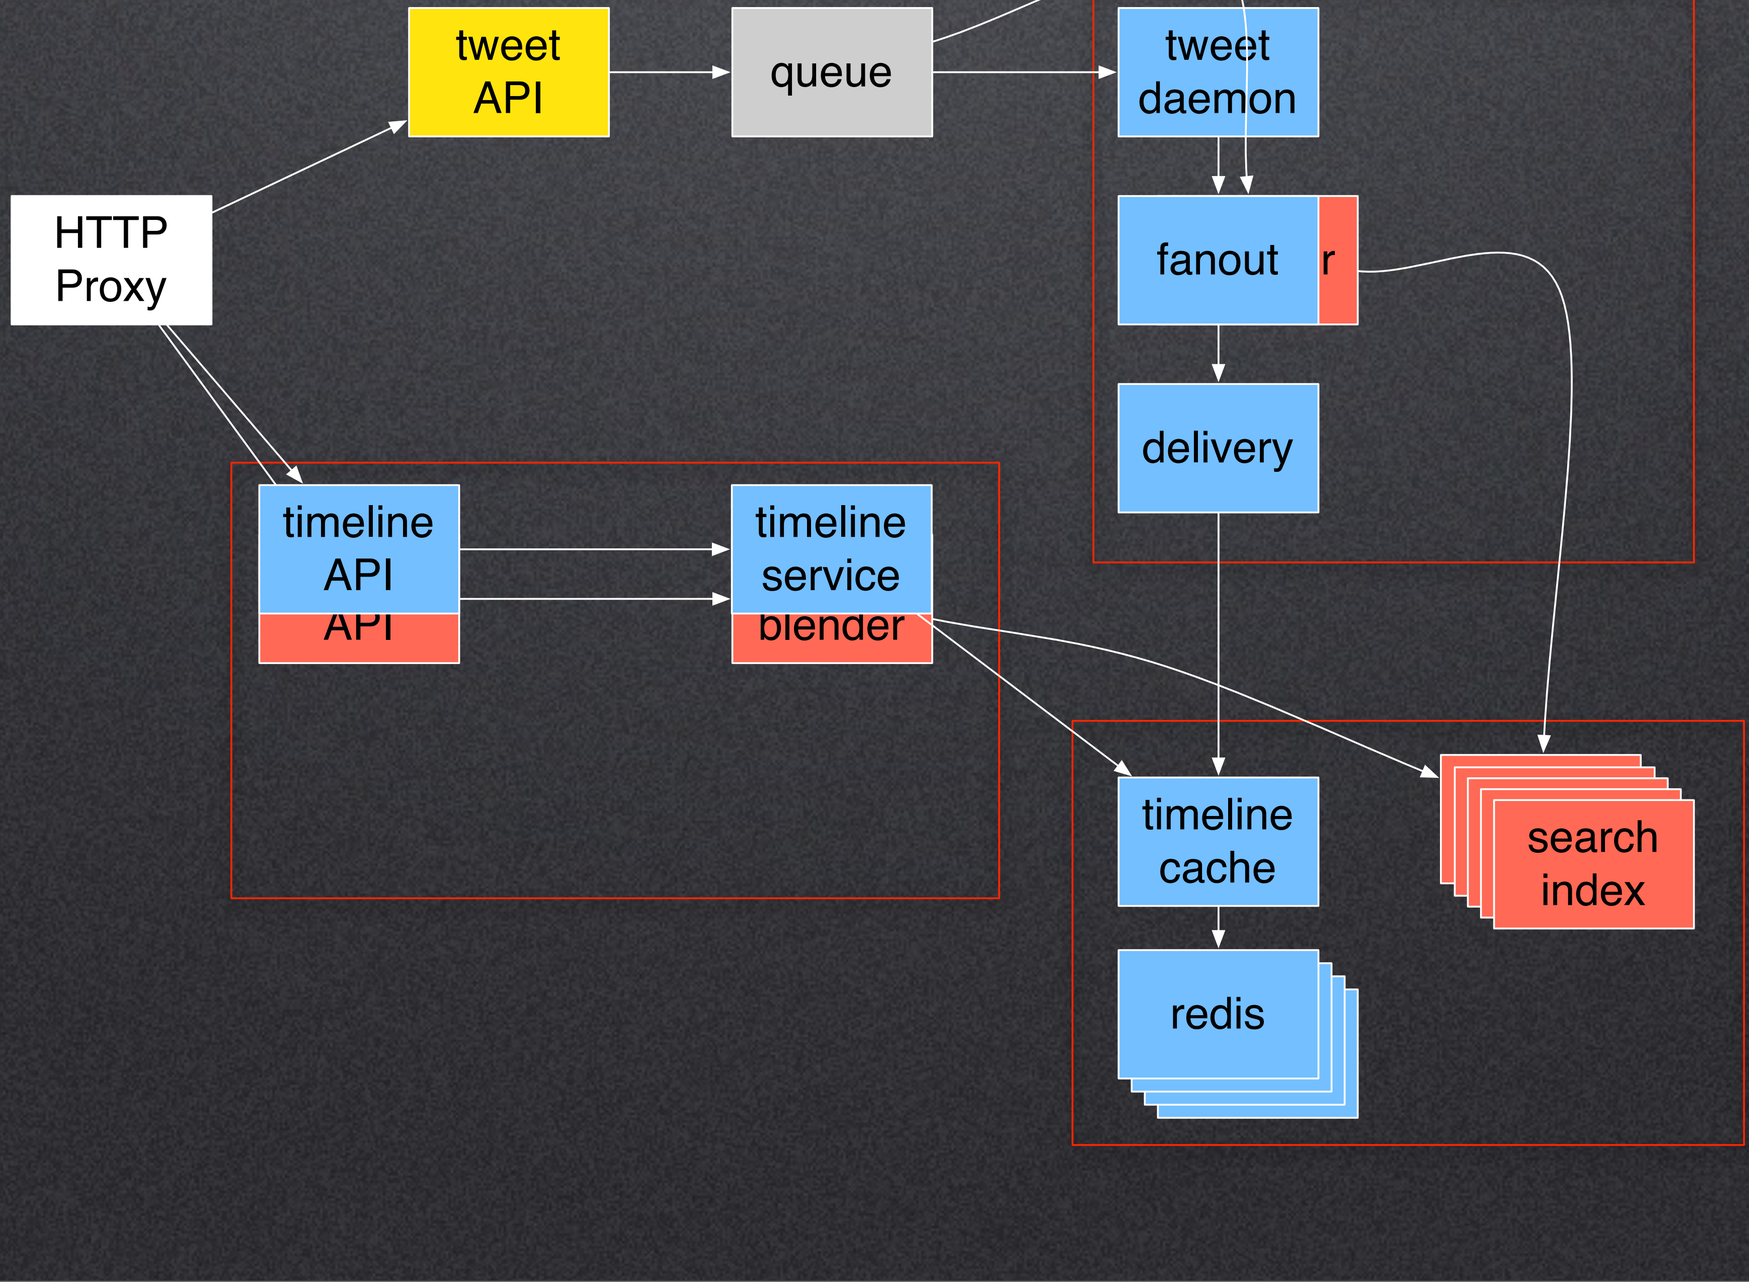
tweet API
queue
tweet daemon
HTTP Proxy
fanout
r
delivery
timeline API
API
timeline service
blender
timeline cache
search index
redis

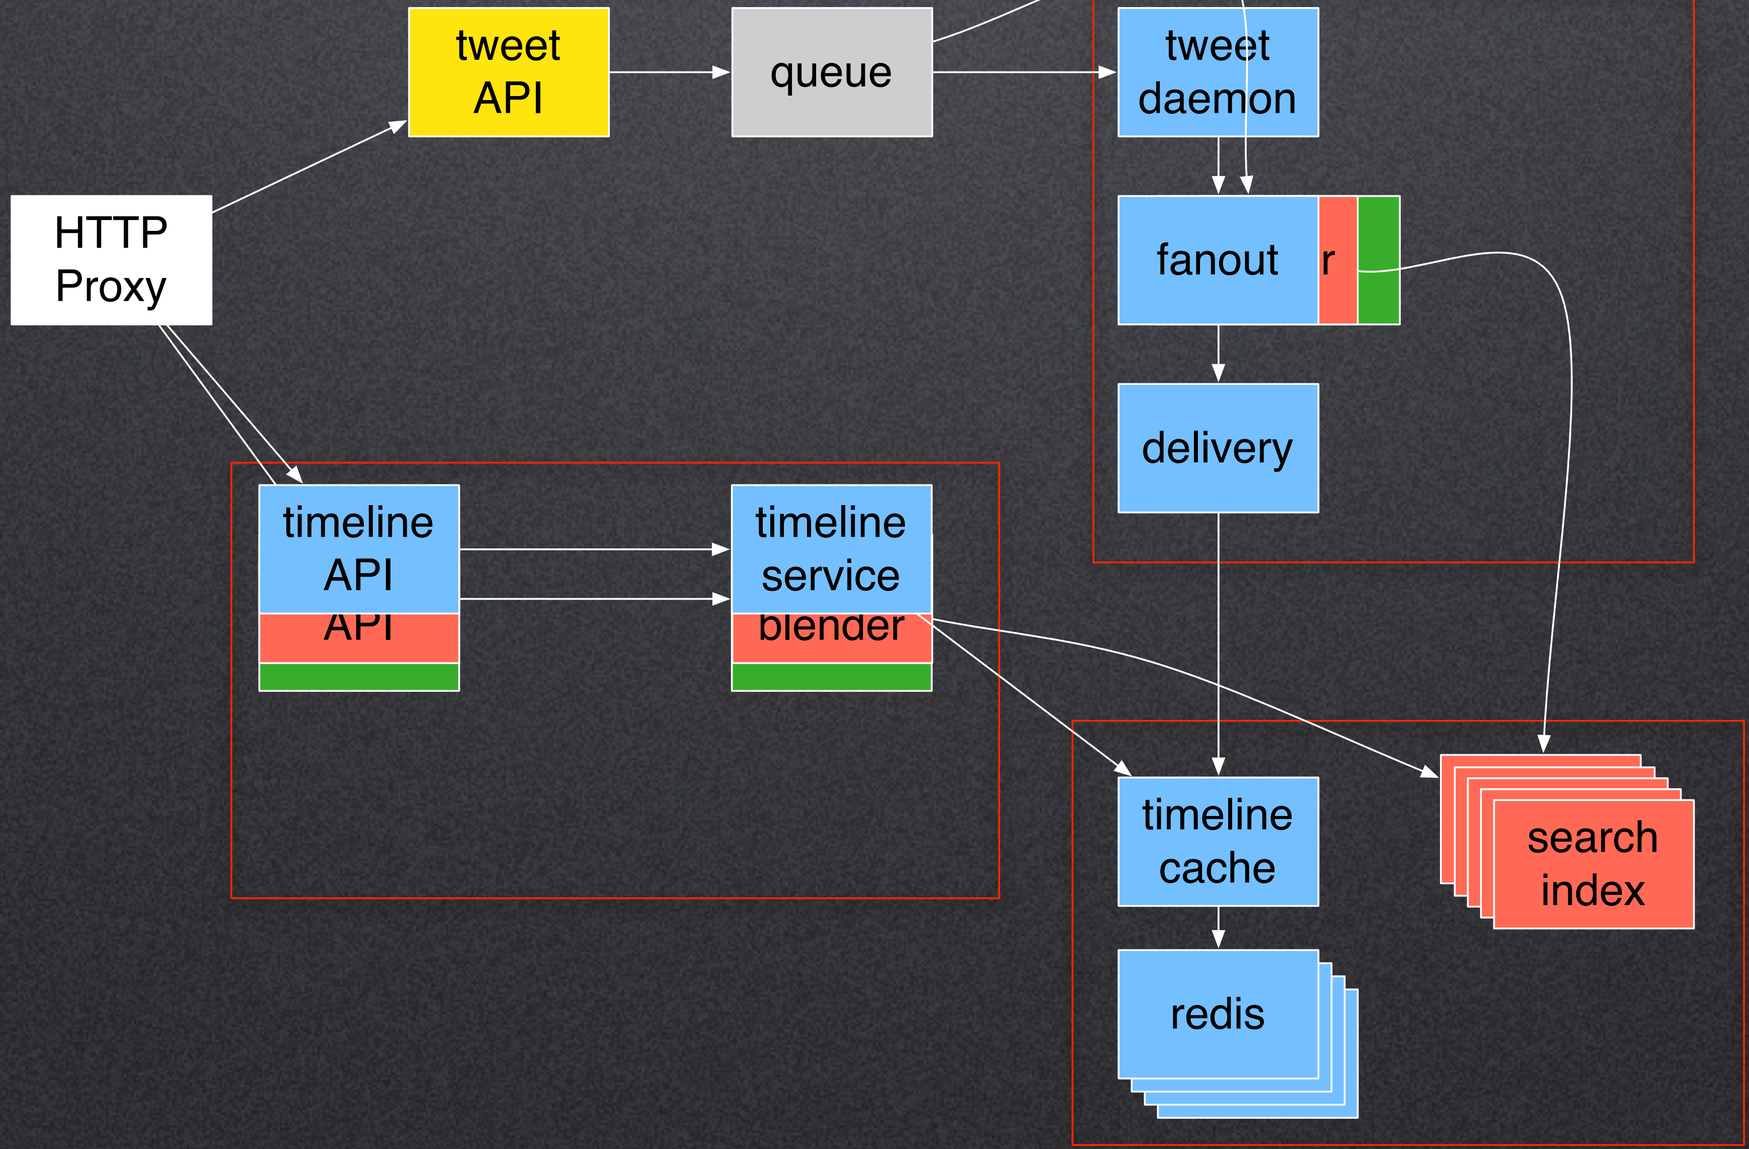
tweet API
queue
tweet daemon
HTTP Proxy
fanout
r
delivery
timeline API
API
timeline service
blender
timeline cache
search index
redis

Questions?

Arya Asemanfar
@a_a