




# real-time delivery architecture



designing twitter

# what are the goals?




→ evolve from being solely a web stack

**ROUTING**

**PRESENTATION**

**LOGIC**

**STORAGE &  
RETRIEVAL**





# what are the goals?

- evolve from being solely a web stack
- isolate responsibilities and concerns
- site speed and reliability
- developer innovation speed

# A Pattern Language

Towns · Buildings · Construction



Christopher Alexander

Sara Ishikawa · Murray Silverstein

WITH

Max Jacobson · Ingrid Fiksdahl-King

Shlomo Angel

## Pull

## Push

---

### Targeted

twitter.com

User / Site Streams

home\_timeline API


Mobile Push (SMS, etc.)

---

### Queried

Search API

Track / Follow Streams

 **Nelson Minar** @nelson  
That part of the hip hop record where  
a recording of someone urinating.  
Expand

 **Nelson Minar** @nelson  
@kvox awesome photo, dapper man  
vintage airplane?  
View conversation


 **Nelson Minar** @nelson  
@bookscout oh bullshit, merlot is a  
Ch. Petrus, for instance.  
View conversation

 **Nelson Minar** @nelson  
Yesterday was South Park, today is  
Area startup tour? Go kart track? Still  
from Woodside, CA

 **Nelson Minar** @nelson  
@skatterbean Harris Ranch is a fun  
narrow. Don't count on the fuel work  
View conversation

 **Nelson Minar** @nelson  
@bradfordcross no worries, thanks!  
stopped using Techmeme. Just make scrolling work consistently! :-)  
View conversation

36m

 **Tim O'Reilly** @timoreilly  
"online messages that evoke feelings  
unintentionally joy!)...a poetry of the  
Expand

 **Tim O'Reilly** @timoreilly  
What's the craziest thing your dad  
results from @make's homage to Facebook  
Expand

 **Tim O'Reilly** @timoreilly  
Potentially wonderful, Also very scary  
banker in Spain [bit.ly/KExVdk](http://bit.ly/KExVdk) Fund  
View media

 **Tim O'Reilly** @timoreilly  
Old friend selling beautiful cottage on  
Queens, NYC: [bit.ly/KEvKqf](http://bit.ly/KEvKqf)  
Expand

 **Ivan Tarradellas** @itarradellas  
@deepakmohoni @timoreilly @quixote  
"Go to the edge of the cliff and jump  
down"  
Retweeted by Tim O'Reilly  
View conversation Reply Retw

 **Chris Matts** @PapaChrisMatts  
Looking forward to #QCONNewYork on Monday. Talking about Lazy  
Learner. The other the "Agile Individuals track" speakers are most  
#awesome!  
Retweeted by QCon New York  
Expand

 **QCon New York** @qconnewyork  
"Spring Data - NoSQL, NoProblems" talk by @PeterBell at QCon  
New York. #springframework #nosql #mongodb #neo4j  
[qconnewyork.com/qcon-nyc-2012/...](http://qconnewyork.com/qcon-nyc-2012/)  
Expand

 **QCon New York** @qconnewyork  
Want to learn how #APM can help you get the most out of your  
#bigdata application? See @mikopp's session at #qconnewyork  
[qconnewyork.com/qcon-nyc-2012/...](http://qconnewyork.com/qcon-nyc-2012/)  
Expand

 **QCon New York** @qconnewyork  
Creating maintainable, sustainable software can be hard but not  
impossible-Alexander von Zitzewitz #qconnewyork [goo.gl/o3Dy7](http://goo.gl/o3Dy7)  
Expand

 **QCon New York** @qconnewyork  
don't miss your chance to save \$200 for #qconnewyork! early bird  
expires Friday 6/15/2012! [secure.trifork.com/qcon-nyc-2012/...](http://secure.trifork.com/qcon-nyc-2012/)  
Expand

15 Jun

14 Jun

13 Jun

12 Jun

11 Jun



**Nelson Minar** @nelson

That part of the hip hop record where  
a recording of someone urinating.

Expand

36m



**Tim O'Reilly** @timoreilly

"online messages that evoke feelings  
unintentionally joy!)...a poetry of the

Expand



**Chris Matts** @PapaChrisMatts

Looking forward to #QCONNewYork on Monday. Talking about Lazy  
Learner. The other the "Agile Individuals track" speakers are most  
#awesome!

Retweeted by QCon New York

Expand

15 Jun



**Tim O'Reilly** @timoreilly

What's the craziest thing your dad e  
results from @make's homage to Fa

Expand



**QCon New York** @qconnewyork

"Spring Data - NoSQL, NoProblems" talk by @PeterBell at QCon  
New York. #springframework #nosql #mongodb #neo4j  
qconnewyork.com/qcon-nyc-2012/...

Expand

14 Jun



**Tim O'Reilly** @timoreilly

Potentially wonderful, Also very sca  
banker in Spain bit.ly/KExVdk Funde

View media



**QCon New York** @qconnewyork

Want to learn how #APM can help you get the most out of your  
#bigdata application? See @mikopp's session at #qconnewyork  
qconnewyork.com/qcon-nyc-2012/...

Expand

13 Jun



**Nelson Minar** @nelson

Yesterday was South Park, today is  
Area startup tour? Go kart track? Su

from Woodside, CA



**Tim O'Reilly** @timoreilly

Old friend selling beautiful cottage on huge garden lot in Dayside,  
Queens, NYC: bit.ly/KEvKqf

Expand



**QCon New York** @qconnewyork

Creating maintainable, sustainable software can be hard but not  
impossible-Alexander von Zitzewitz #qconnewyork goo.gl/o3Dy7

Expand

12 Jun



**Ivan Tarradellas** @itarradellas

@deepakmohoni @timoreilly @quixote Spain Yes, Ray Bradbury  
"Go to the edge of the cliff and jum  
down"

Retweeted by Tim O'Reilly

View conversation Reply Retw



**QCon New York** @qconnewyork

don't miss your chance to save \$200 for #qconnewyork! early bird  
expires Friday 6/15/2012! secure.trifork.com/qcon-nyc-2012/...

Expand

11 Jun





**Tim O'Reilly** @timoreilly

31m

"online messages that evoke feelings of despair (and, often and unintentionally joy!)...a poetry of the New Aesthetic" [bit.ly/NlePmd](http://bit.ly/NlePmd)

Expand



**Nelson Minar** @nelson

36m

That part of the hip hop record where they think it's funny to include a recording of someone urinating.

Expand



**Tim O'Reilly** @timoreilly

2h

What's the craziest thing your dad ever made? [bit.ly/KEIn4D](http://bit.ly/KEIn4D) Some results from @make's homage to Father's Day

Expand



**Tim O'Reilly** @timoreilly

4h

Potentially wonderful, Also very scary: crowdfunded lawsuit against banker in Spain [bit.ly/KExVdk](http://bit.ly/KExVdk) Funded 24 hrs [bit.ly/KExSOG](http://bit.ly/KExSOG)

[View media](#)



**Tim O'Reilly** @timoreilly

4h

Old friend selling beautiful cottage on huge garden lot in Bayside, Queens, NYC: [bit.ly/KEvKqf](http://bit.ly/KEvKqf)

Expand



**Nelson Minar** @nelson

23h

Yesterday was South Park, today is Buck's. What's next on my Bay Area startup tour? Go kart track? Suburban garages?

[from Woodside, CA](#)



**Ivan Tarradellas** @itarradellas

16 Jun

@deepakmohoni @timoreilly @quixotic @dpatil Yes, Ray Bradbury "Go to the edge of the cliff and jump off. Build your wings on the way down"

[Retweeted by Tim O'Reilly](#)

[View conversation](#) [Reply](#) [Retweet](#) [Favorite](#)



**Chris Matts** @PapaChrisMatts

15 Jun

Looking forward to #QCONNewYork on Monday. Talking about Lazy Learner. The other the "Agile Individuals track" speakers are most #awesome!

[Retweeted by QCon New York](#)

What's happening?




i'm speaking at @cal! → "real-time delivery architecture at @twitter"



71


Tweet



Write API




Fanout





# insert

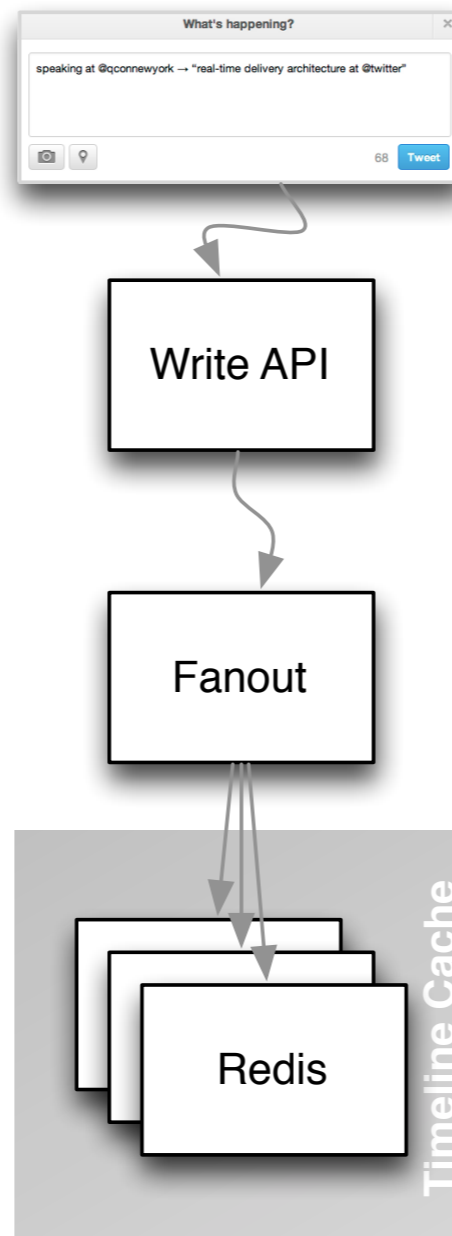
- keyed off “recipient”
- pipelined 4k “destinations” at a time
- replicated



# using redis

→ native list structure

→ RPUSSHX to only add to cached timelines




Tweet ID	User ID	Bits
8 bytes	8 bytes	4 bytes


# using redis

→ native list structure

→ RPU SHX to only add to cached timelines

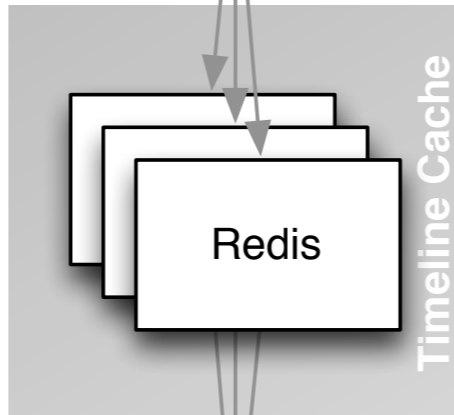


Tweet ID	User ID	Bits	
Tweet ID	User ID	Bits	Tweet ID
Tweet ID	User ID	Bits	
Tweet ID	User ID	Bits	
Tweet ID	User ID	Bits	Tweet ID
Tweet ID	User ID	Bits	
Tweet ID	User ID	Bits	
Tweet ID	User ID	Bits	Tweet ID
Tweet ID	User ID	Bits	
Tweet ID	User ID	Bits	

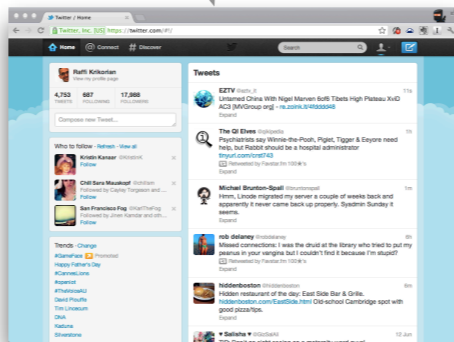


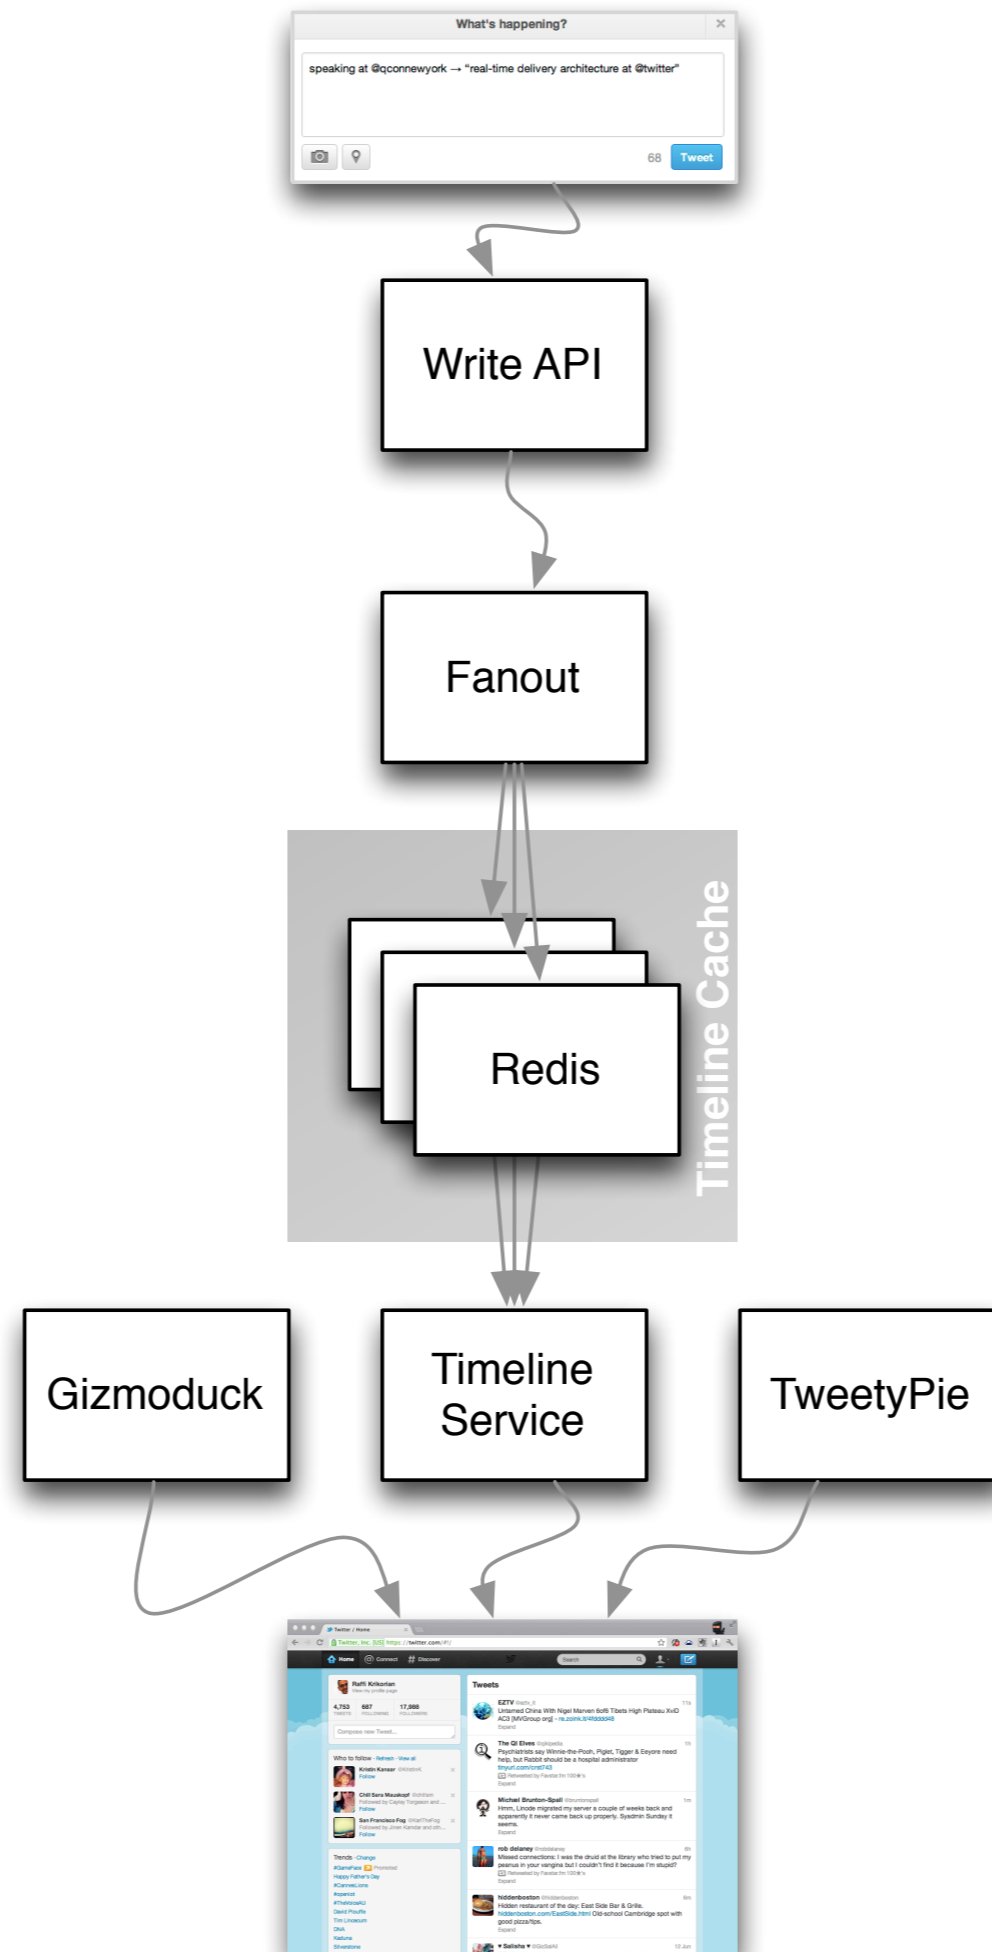
Write API

Fanout




Timeline Service






	<b>Pull</b>	<b>Push</b>
<b>Targeted</b>	twitter.com home_timeline API	User / Site Streams Mobile Push (SMS, etc.)
<b>Queried</b>	Search API	Track / Follow Streams

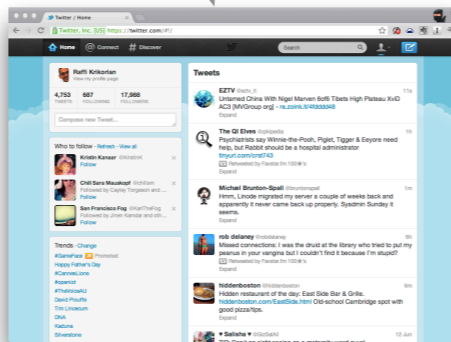


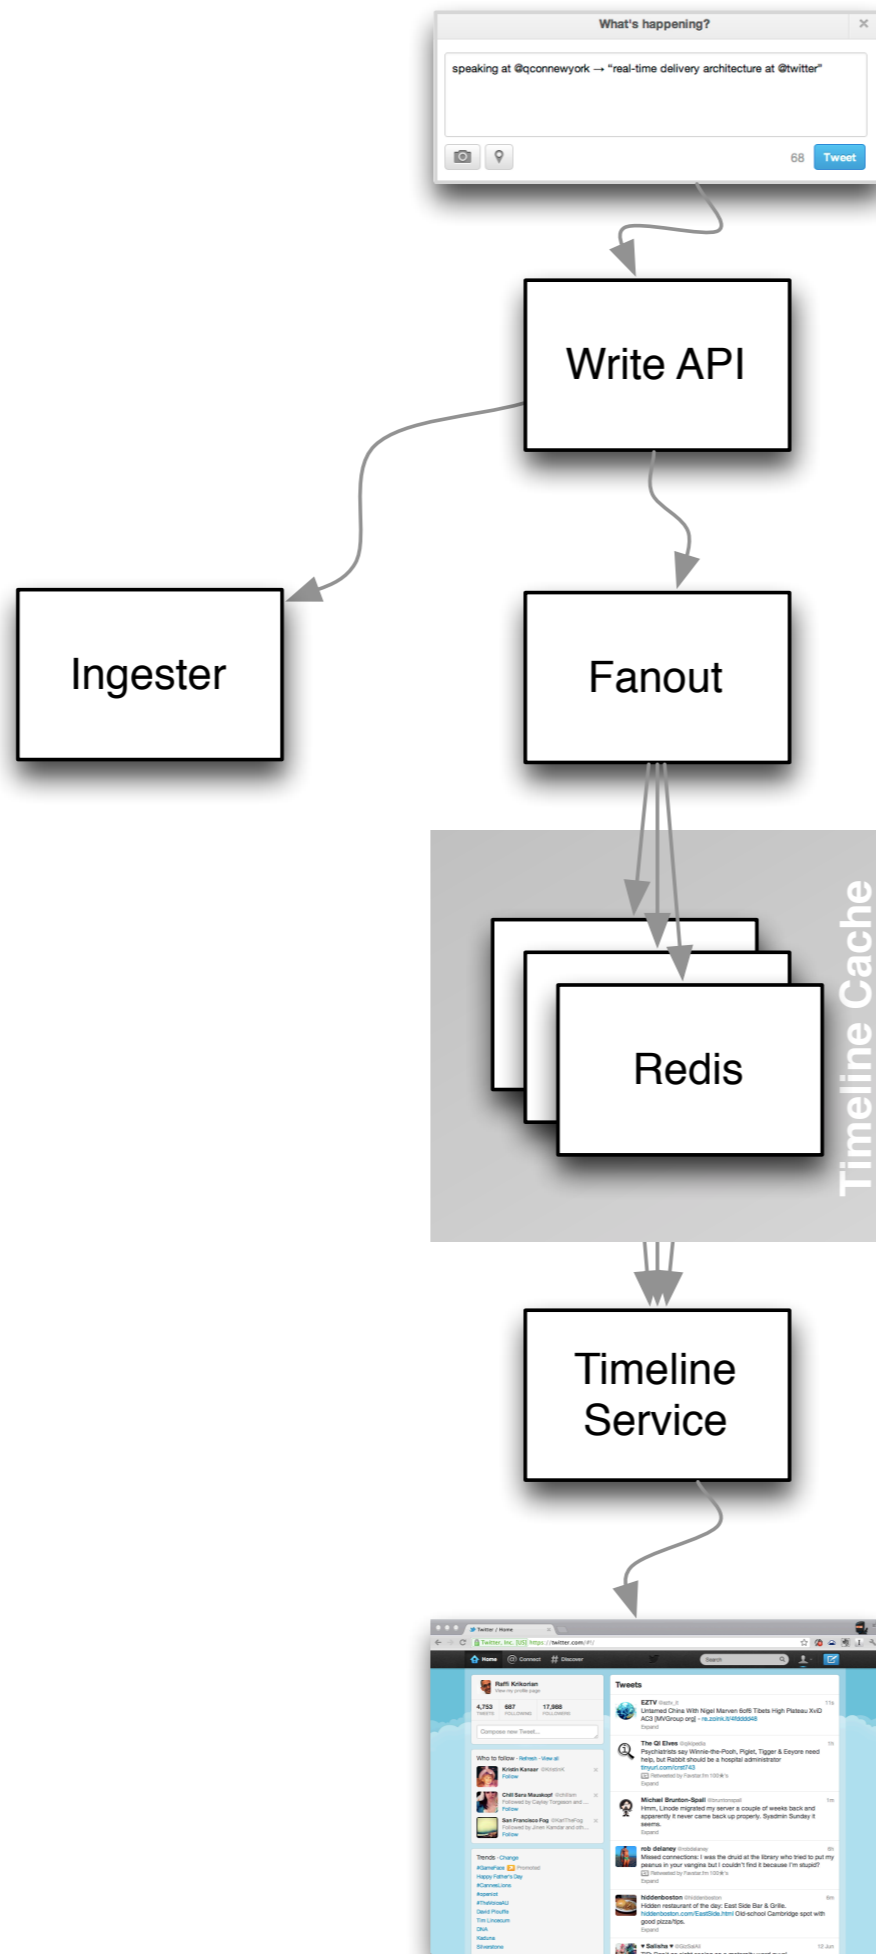
Write API

Fanout

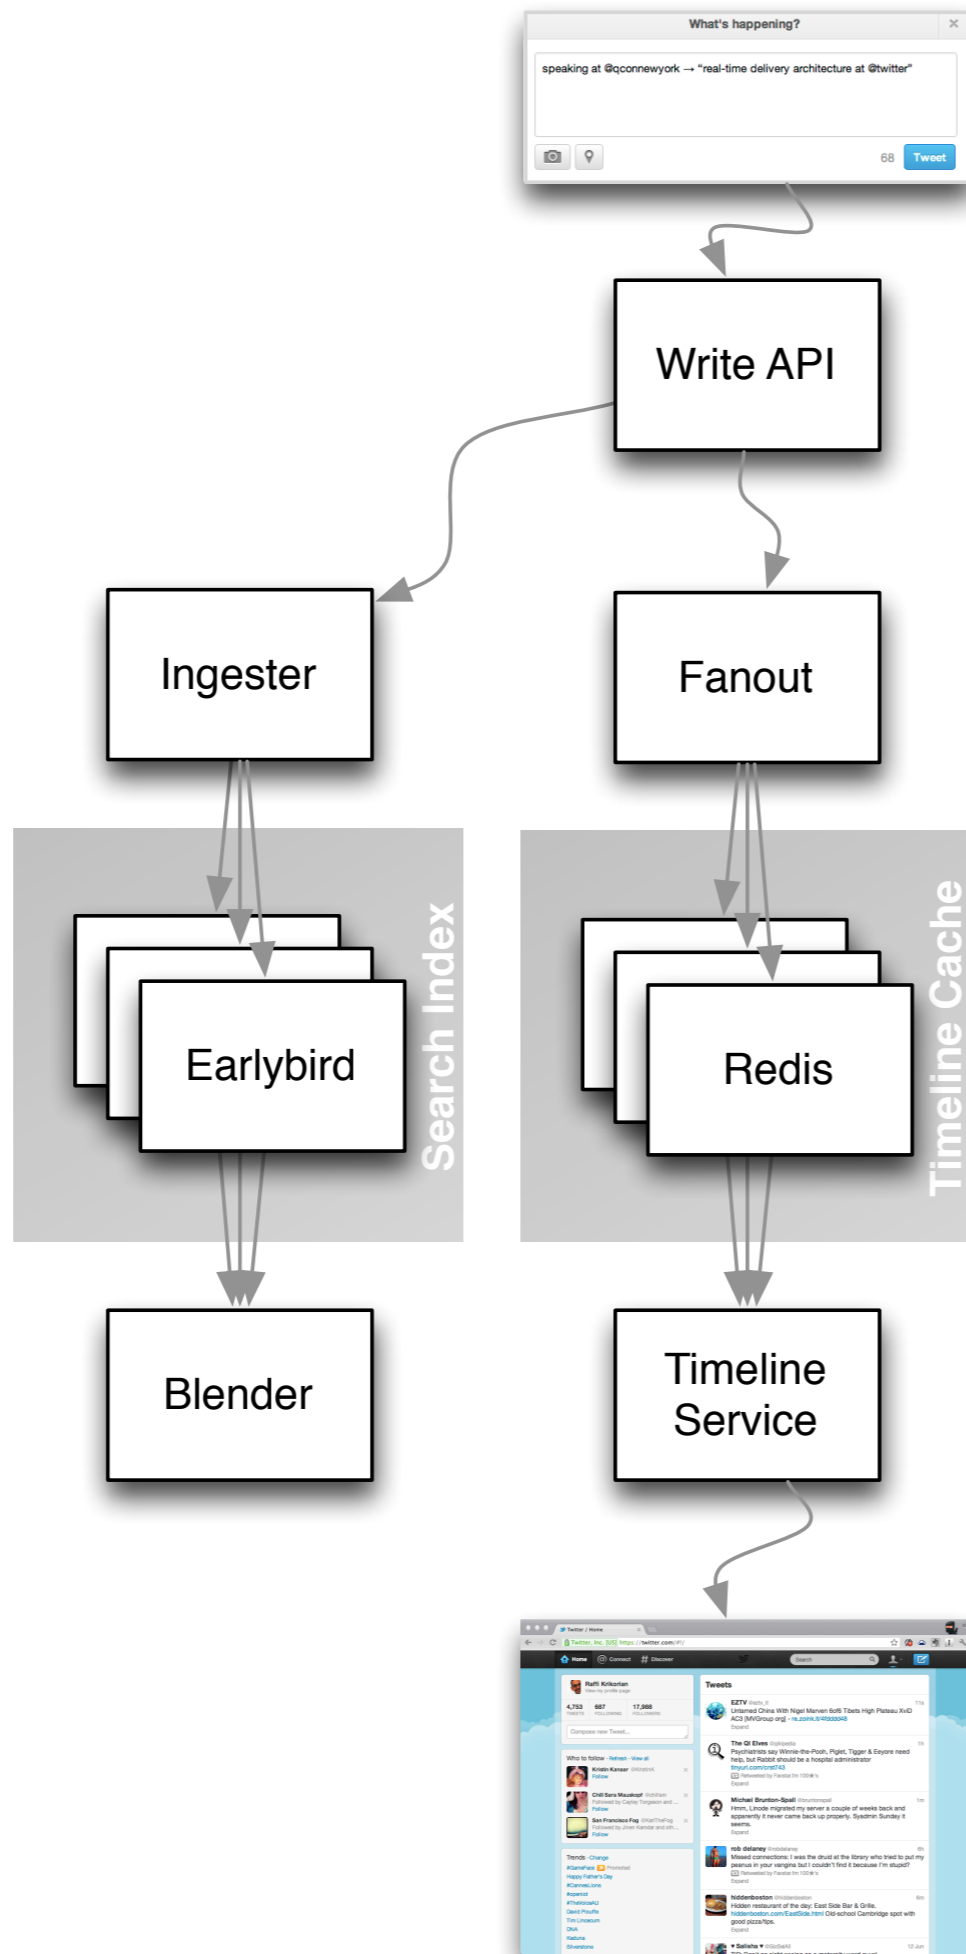


Timeline Service










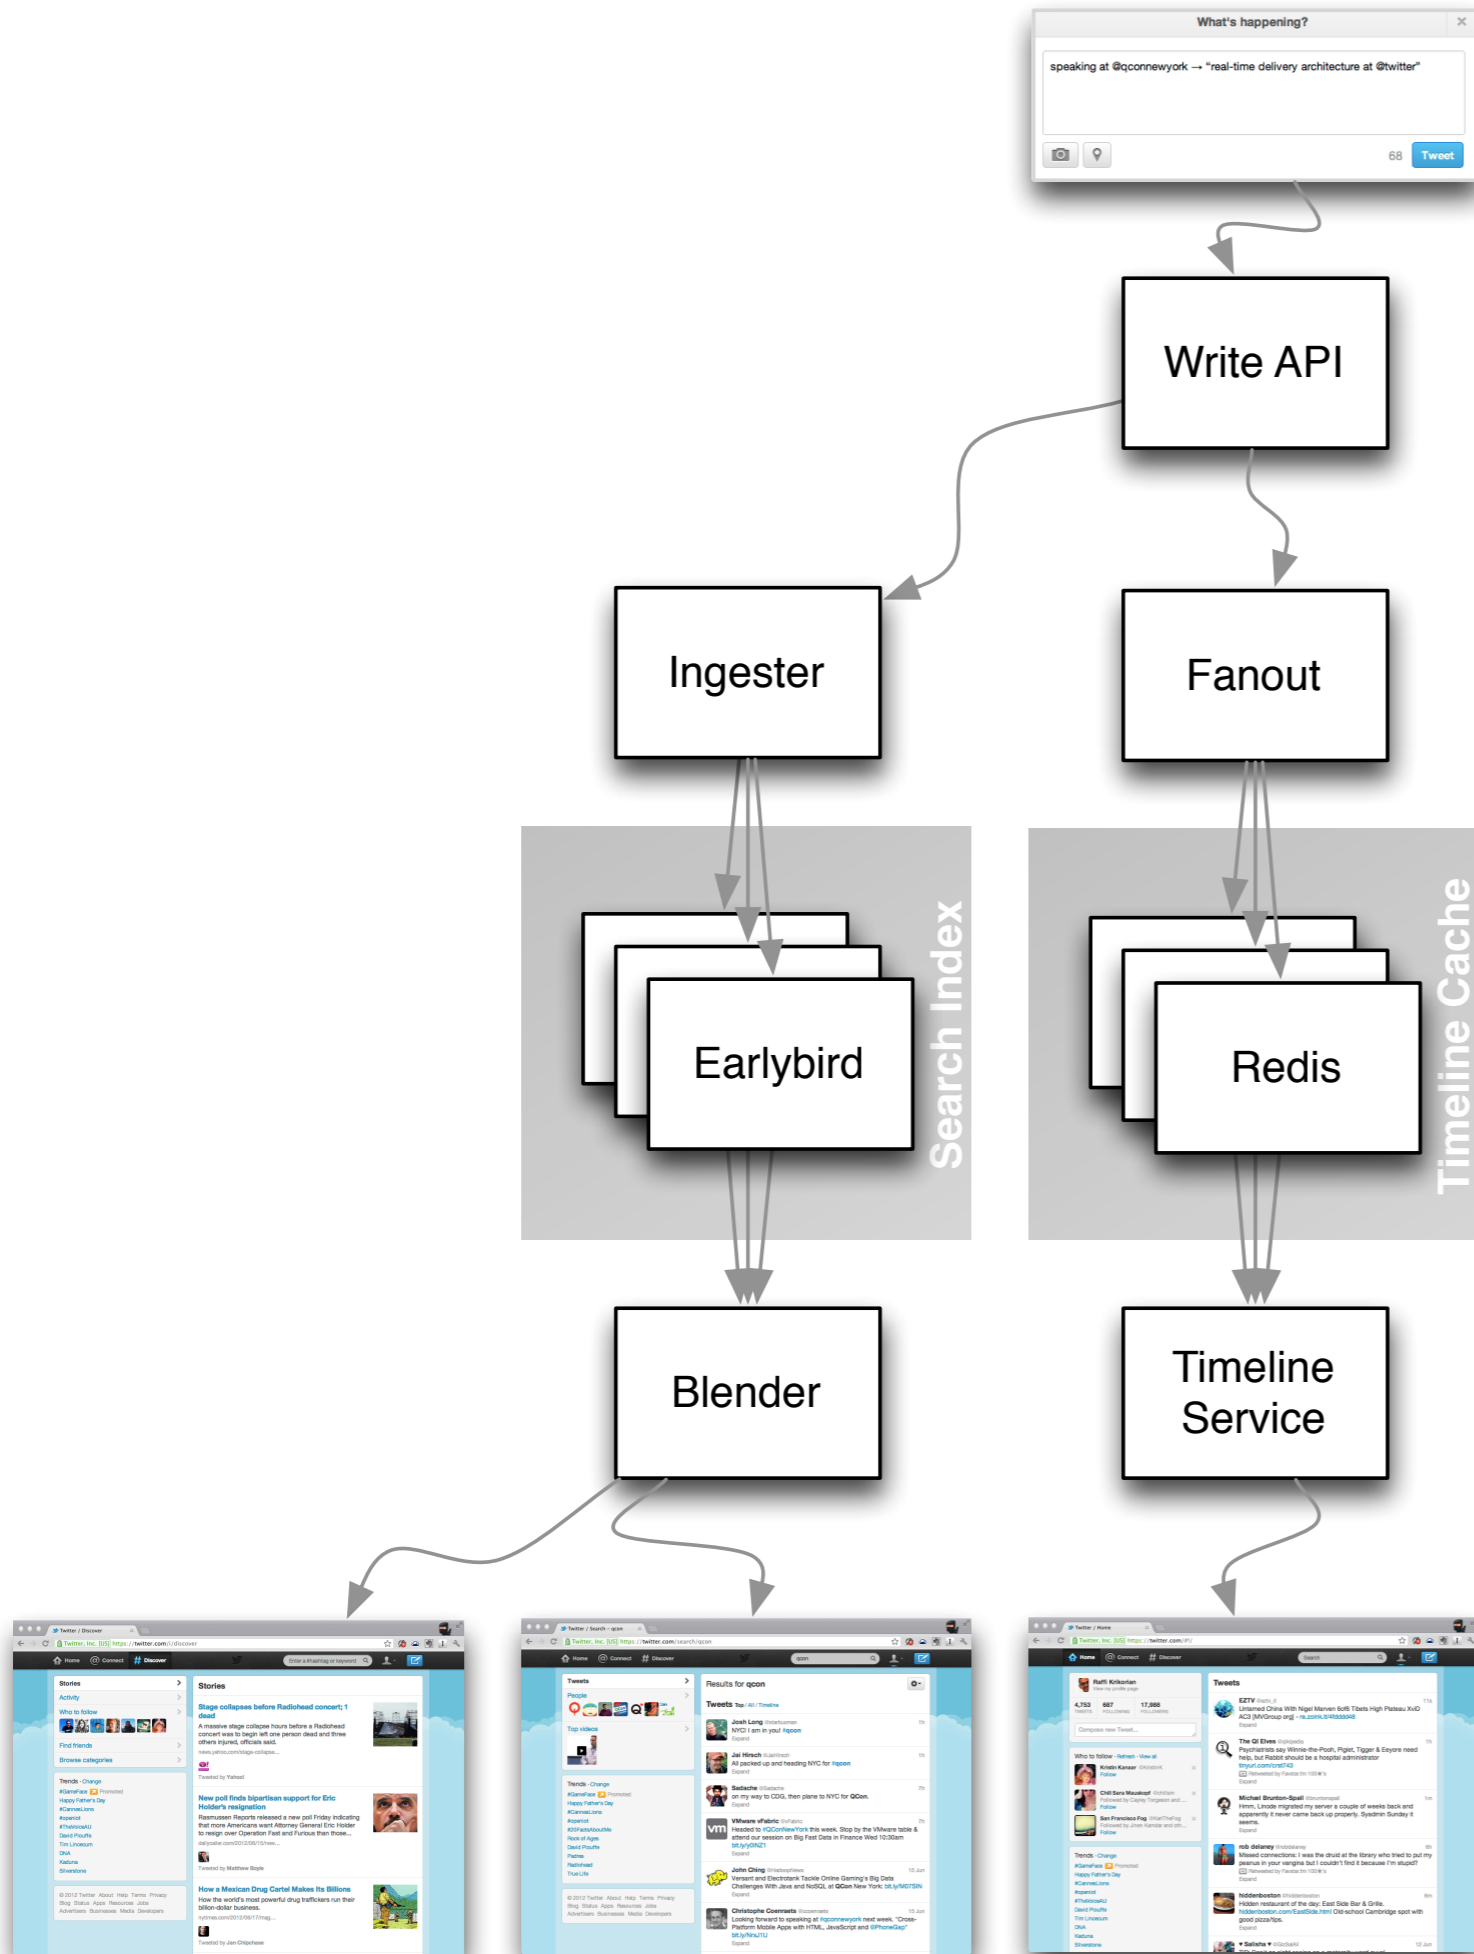
# blender

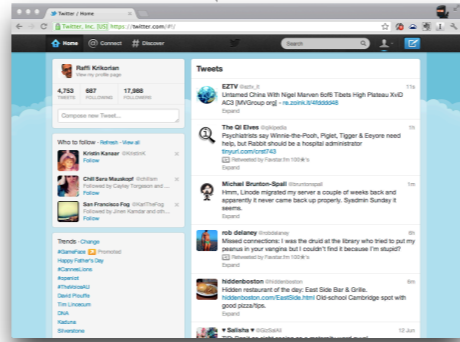
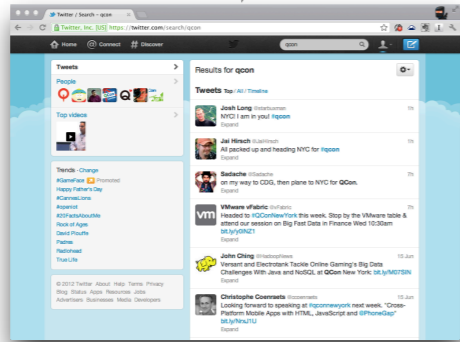
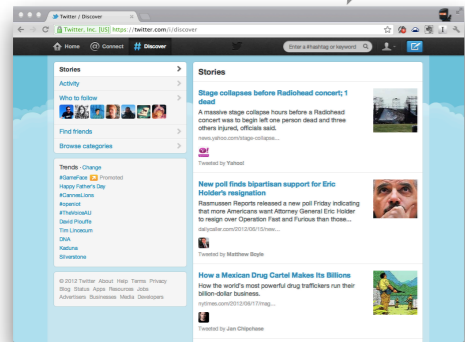
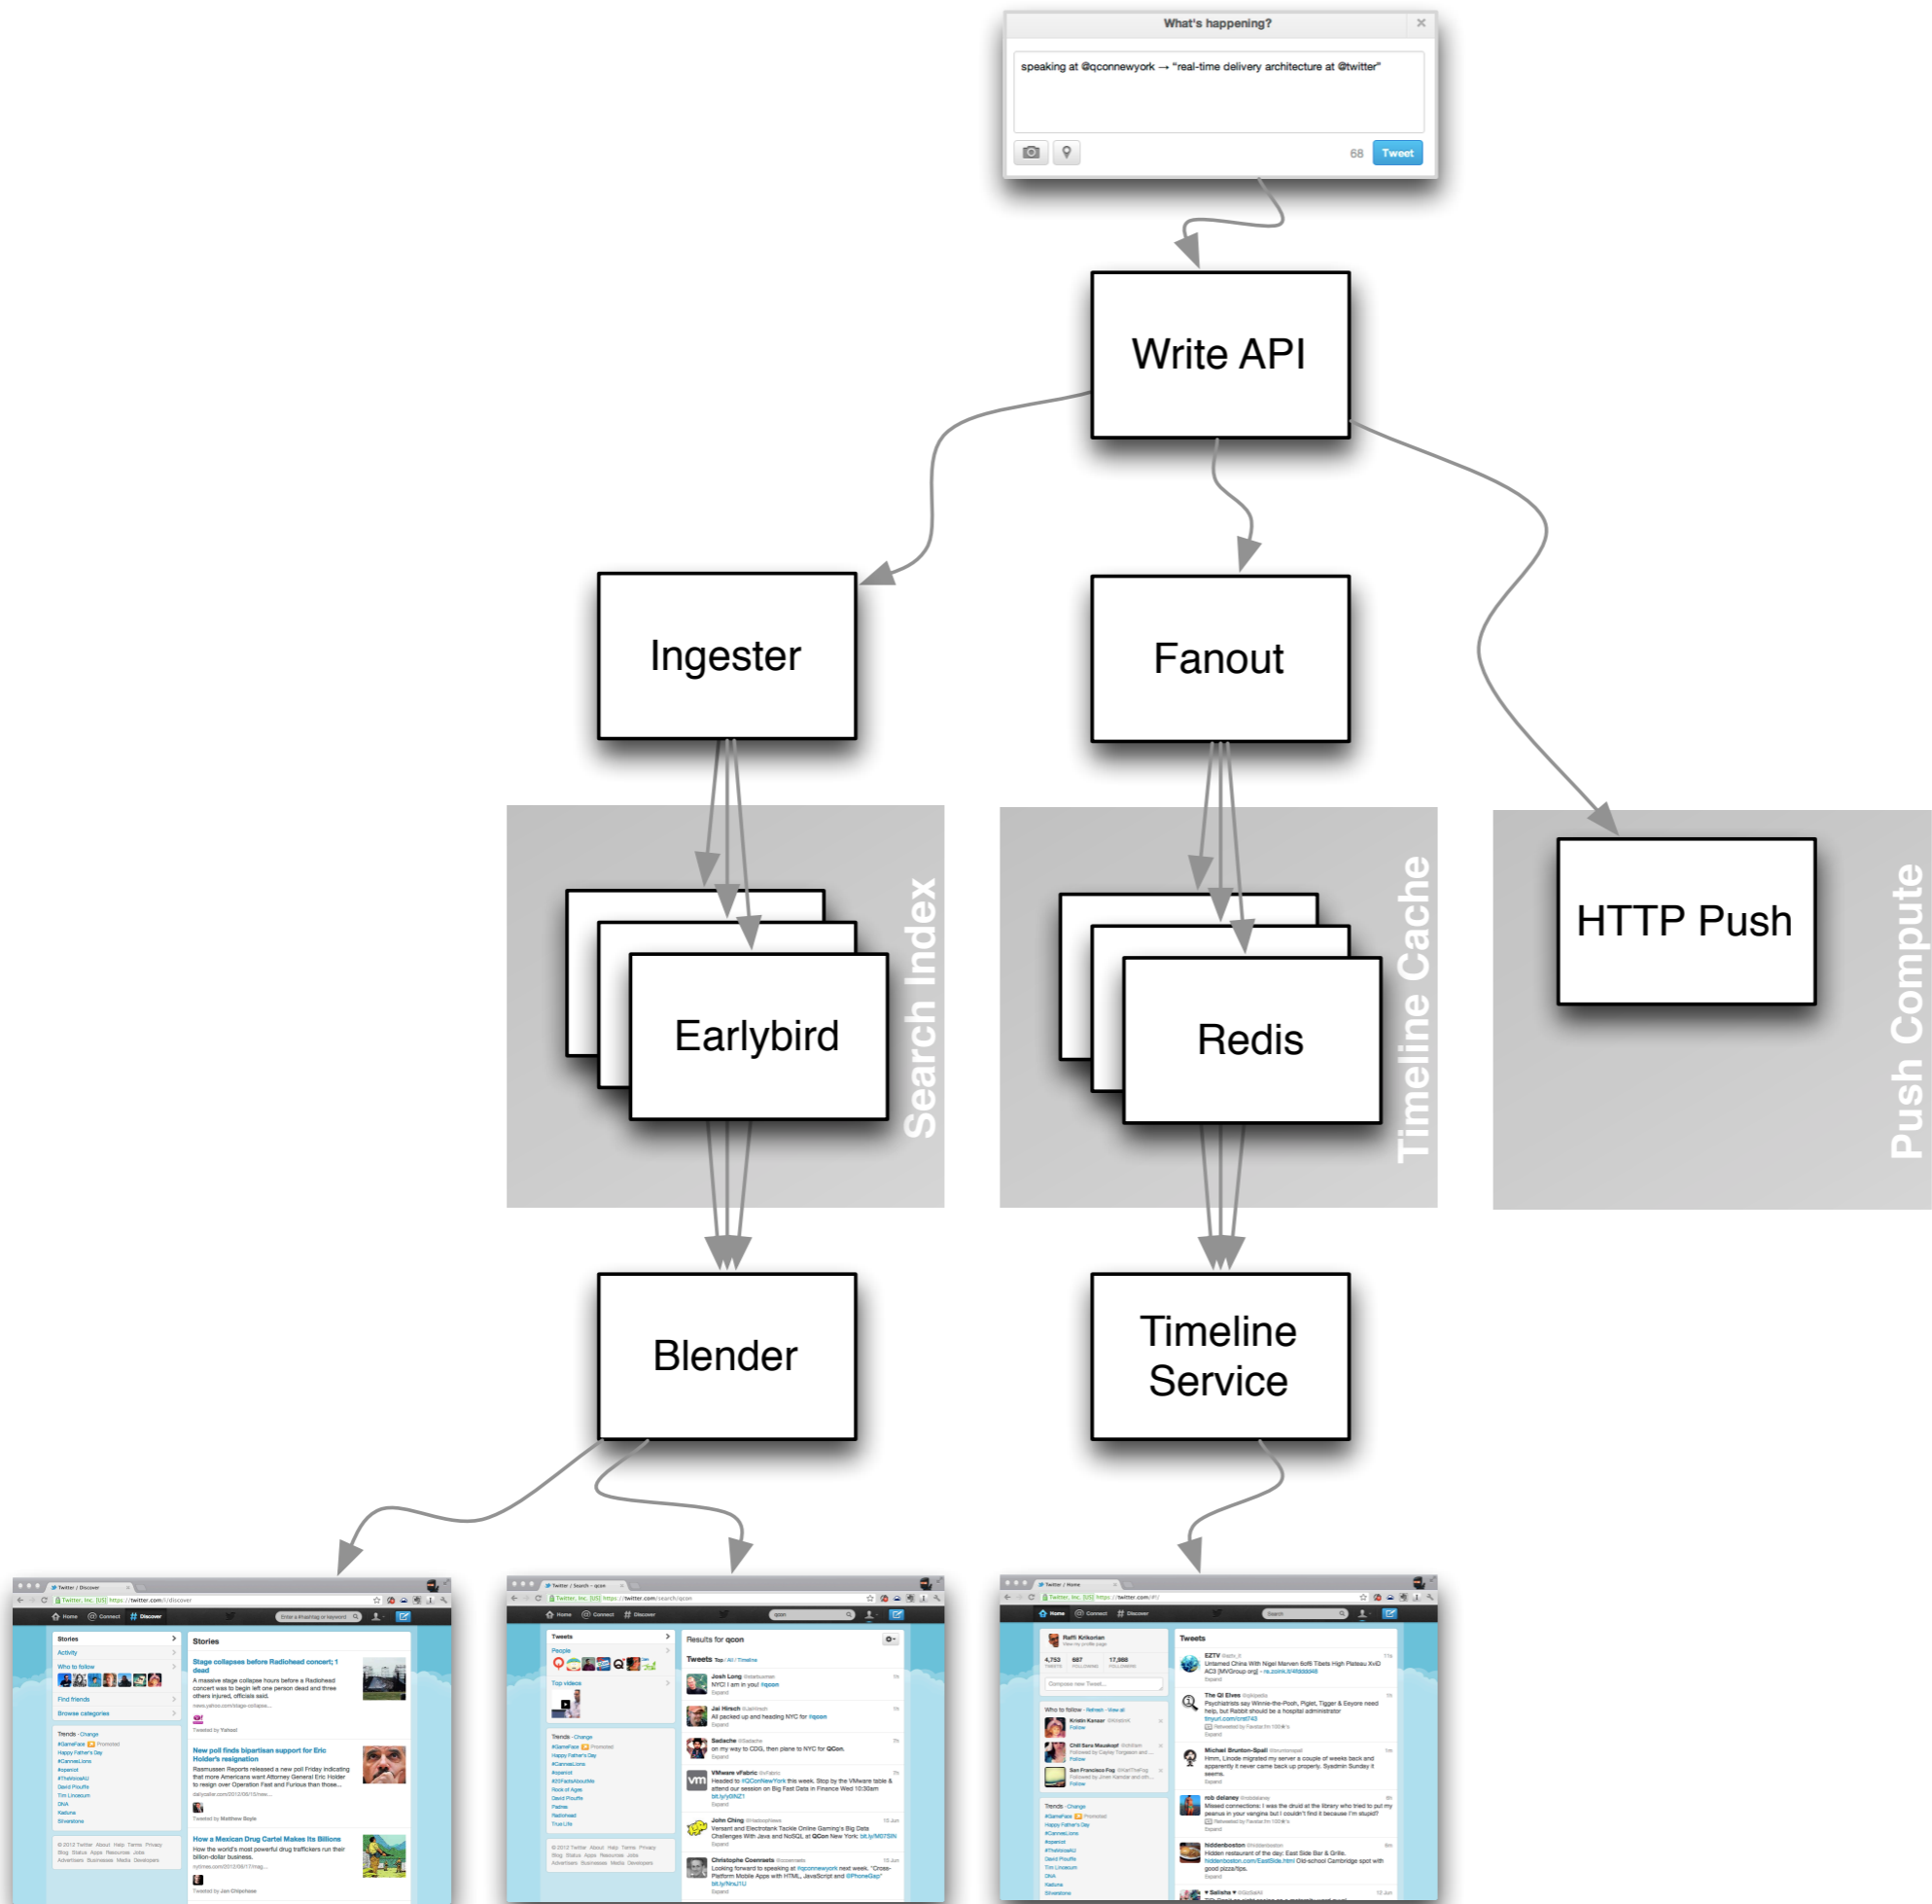
...queries one replica of all indexes

...merges & ranks results



	<b>Pull</b>	<b>Push</b>
<b>Targeted</b>	twitter.com home_timeline API	User / Site Streams Mobile Push (SMS, etc.)
<b>Queried</b>	Search API	Track / Follow Streams





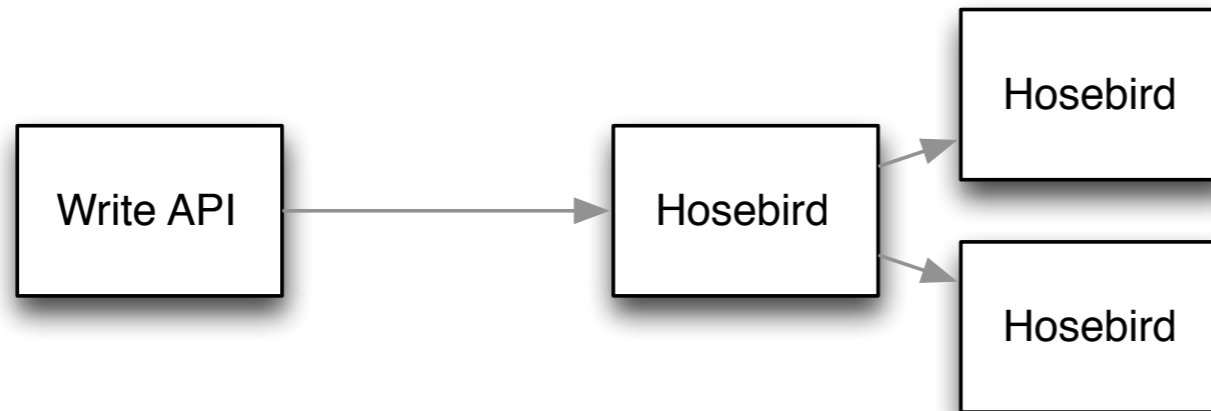
# http push / hosebird

- maintains persistent connections with end clients
- processes tweet & social graph events
- event-based “router”



# event propagation


- write API sends all events into hosebird; sees content creation events, social graph changes, etc.
- different queues for public tweets, protected tweets, social events, etc.



# event cascading


- bandwidth management
- simultaneous connection management  
(~1m long lived & open connections to this cluster)






# firehose

- edge machine simply outputs the public tweet queue
- only allow a limited number of firehoses per hosebird box for bandwidth management





# track / follow

- simple query based on tweet content
- keeps list of terms / users of interest
- parses public tweets at the edge, and if term matches a token, or user is of interest, then route



## user streams

- replicate home timeline experience
- upon login, obtain “following” list
- keep cached following list coherent by seeing social graph updates
- route tweet if from a followed user



Write API

Ingester

Fanout

Search Index

Timeline Cache

Push Compute

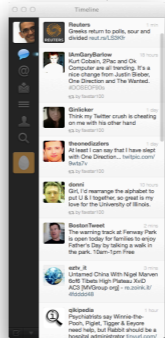
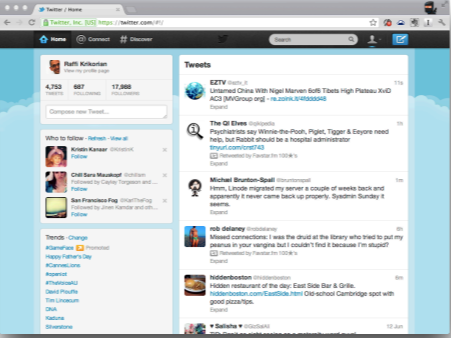
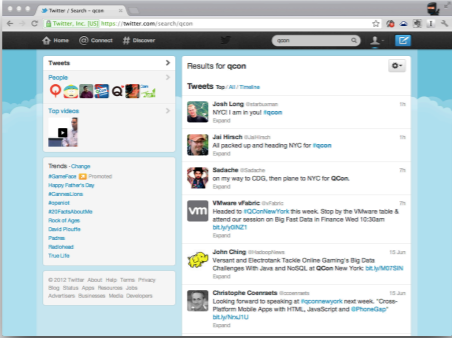

Earlybird

Redis

HTTP Push

Blender

Timeline Service



## Pull

## Push

---

### Targeted

twitter.com

home\_timeline API

User / Site Streams


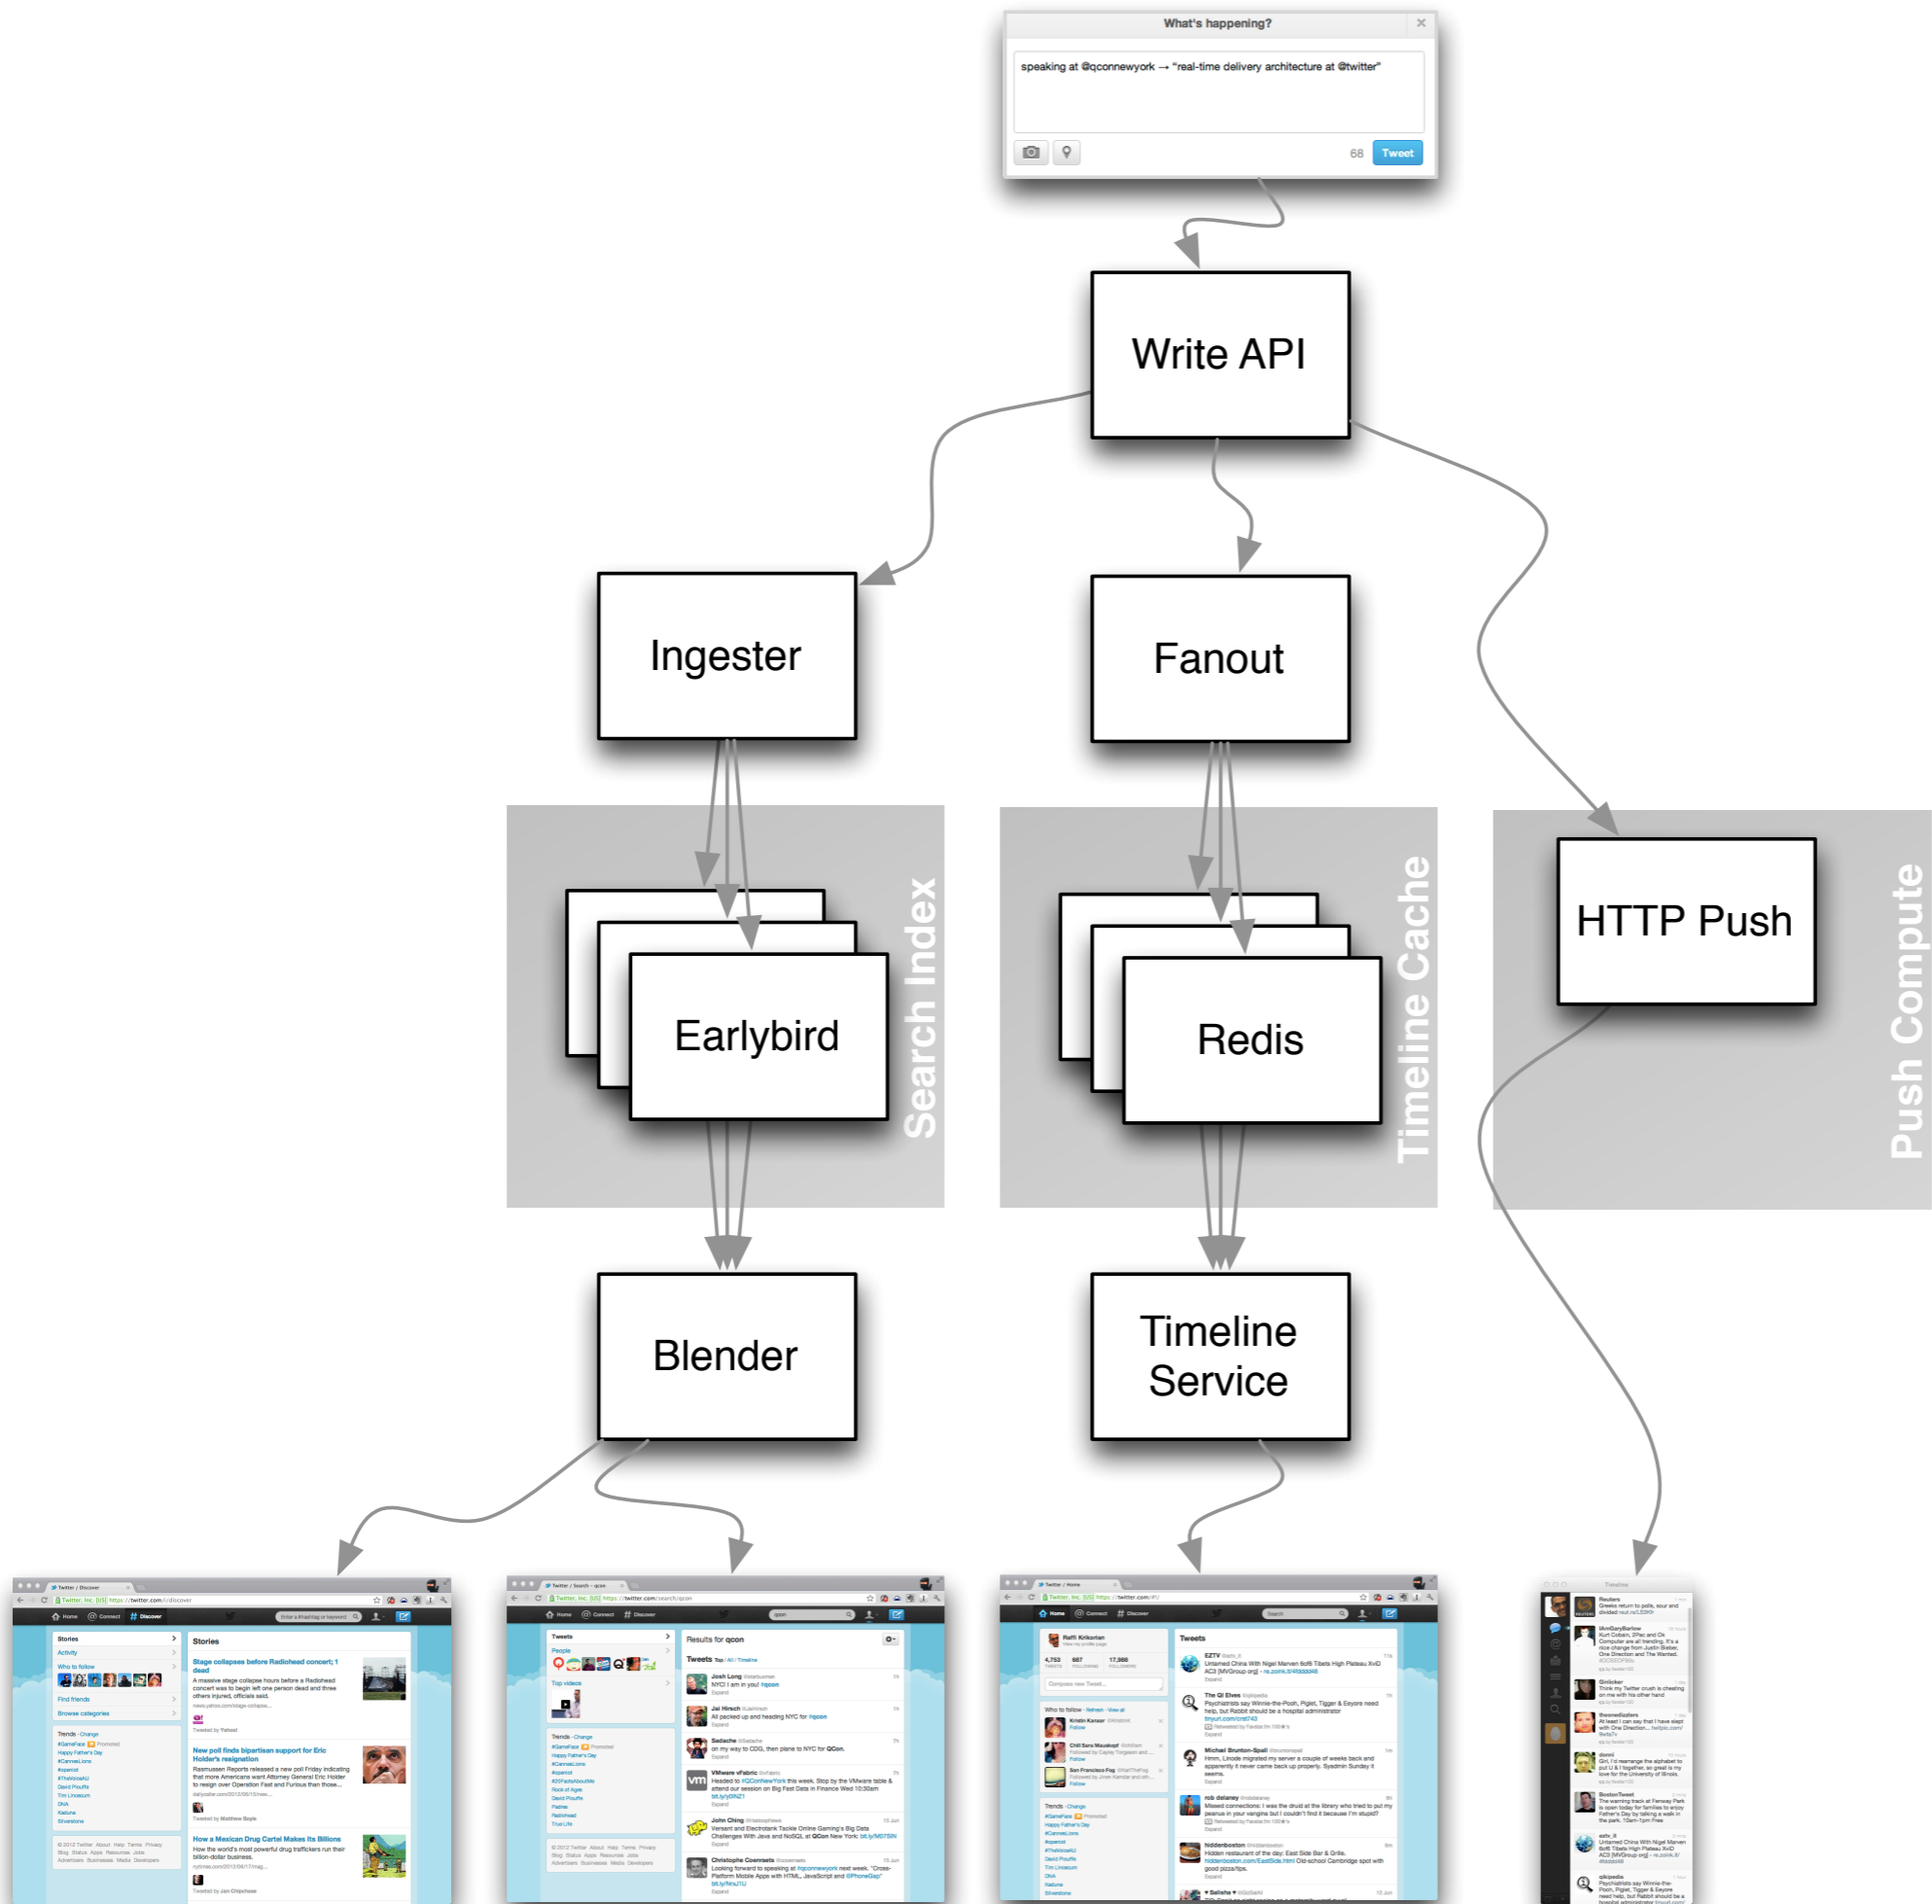
Mobile Push (GCM, etc)

---

### Queried

Search API

Track / Follow Streams



Write API

Ingester

Fanout

Search Index

Earlybird

Timeline Cache

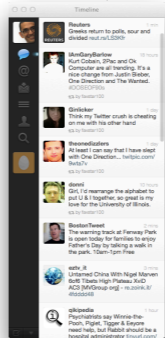


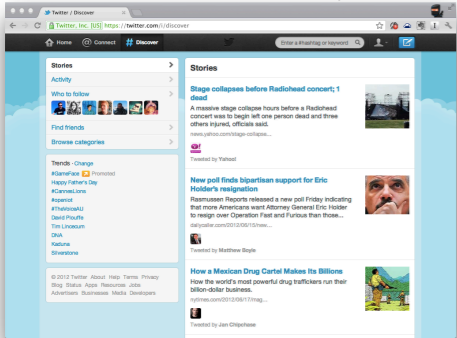
Redis

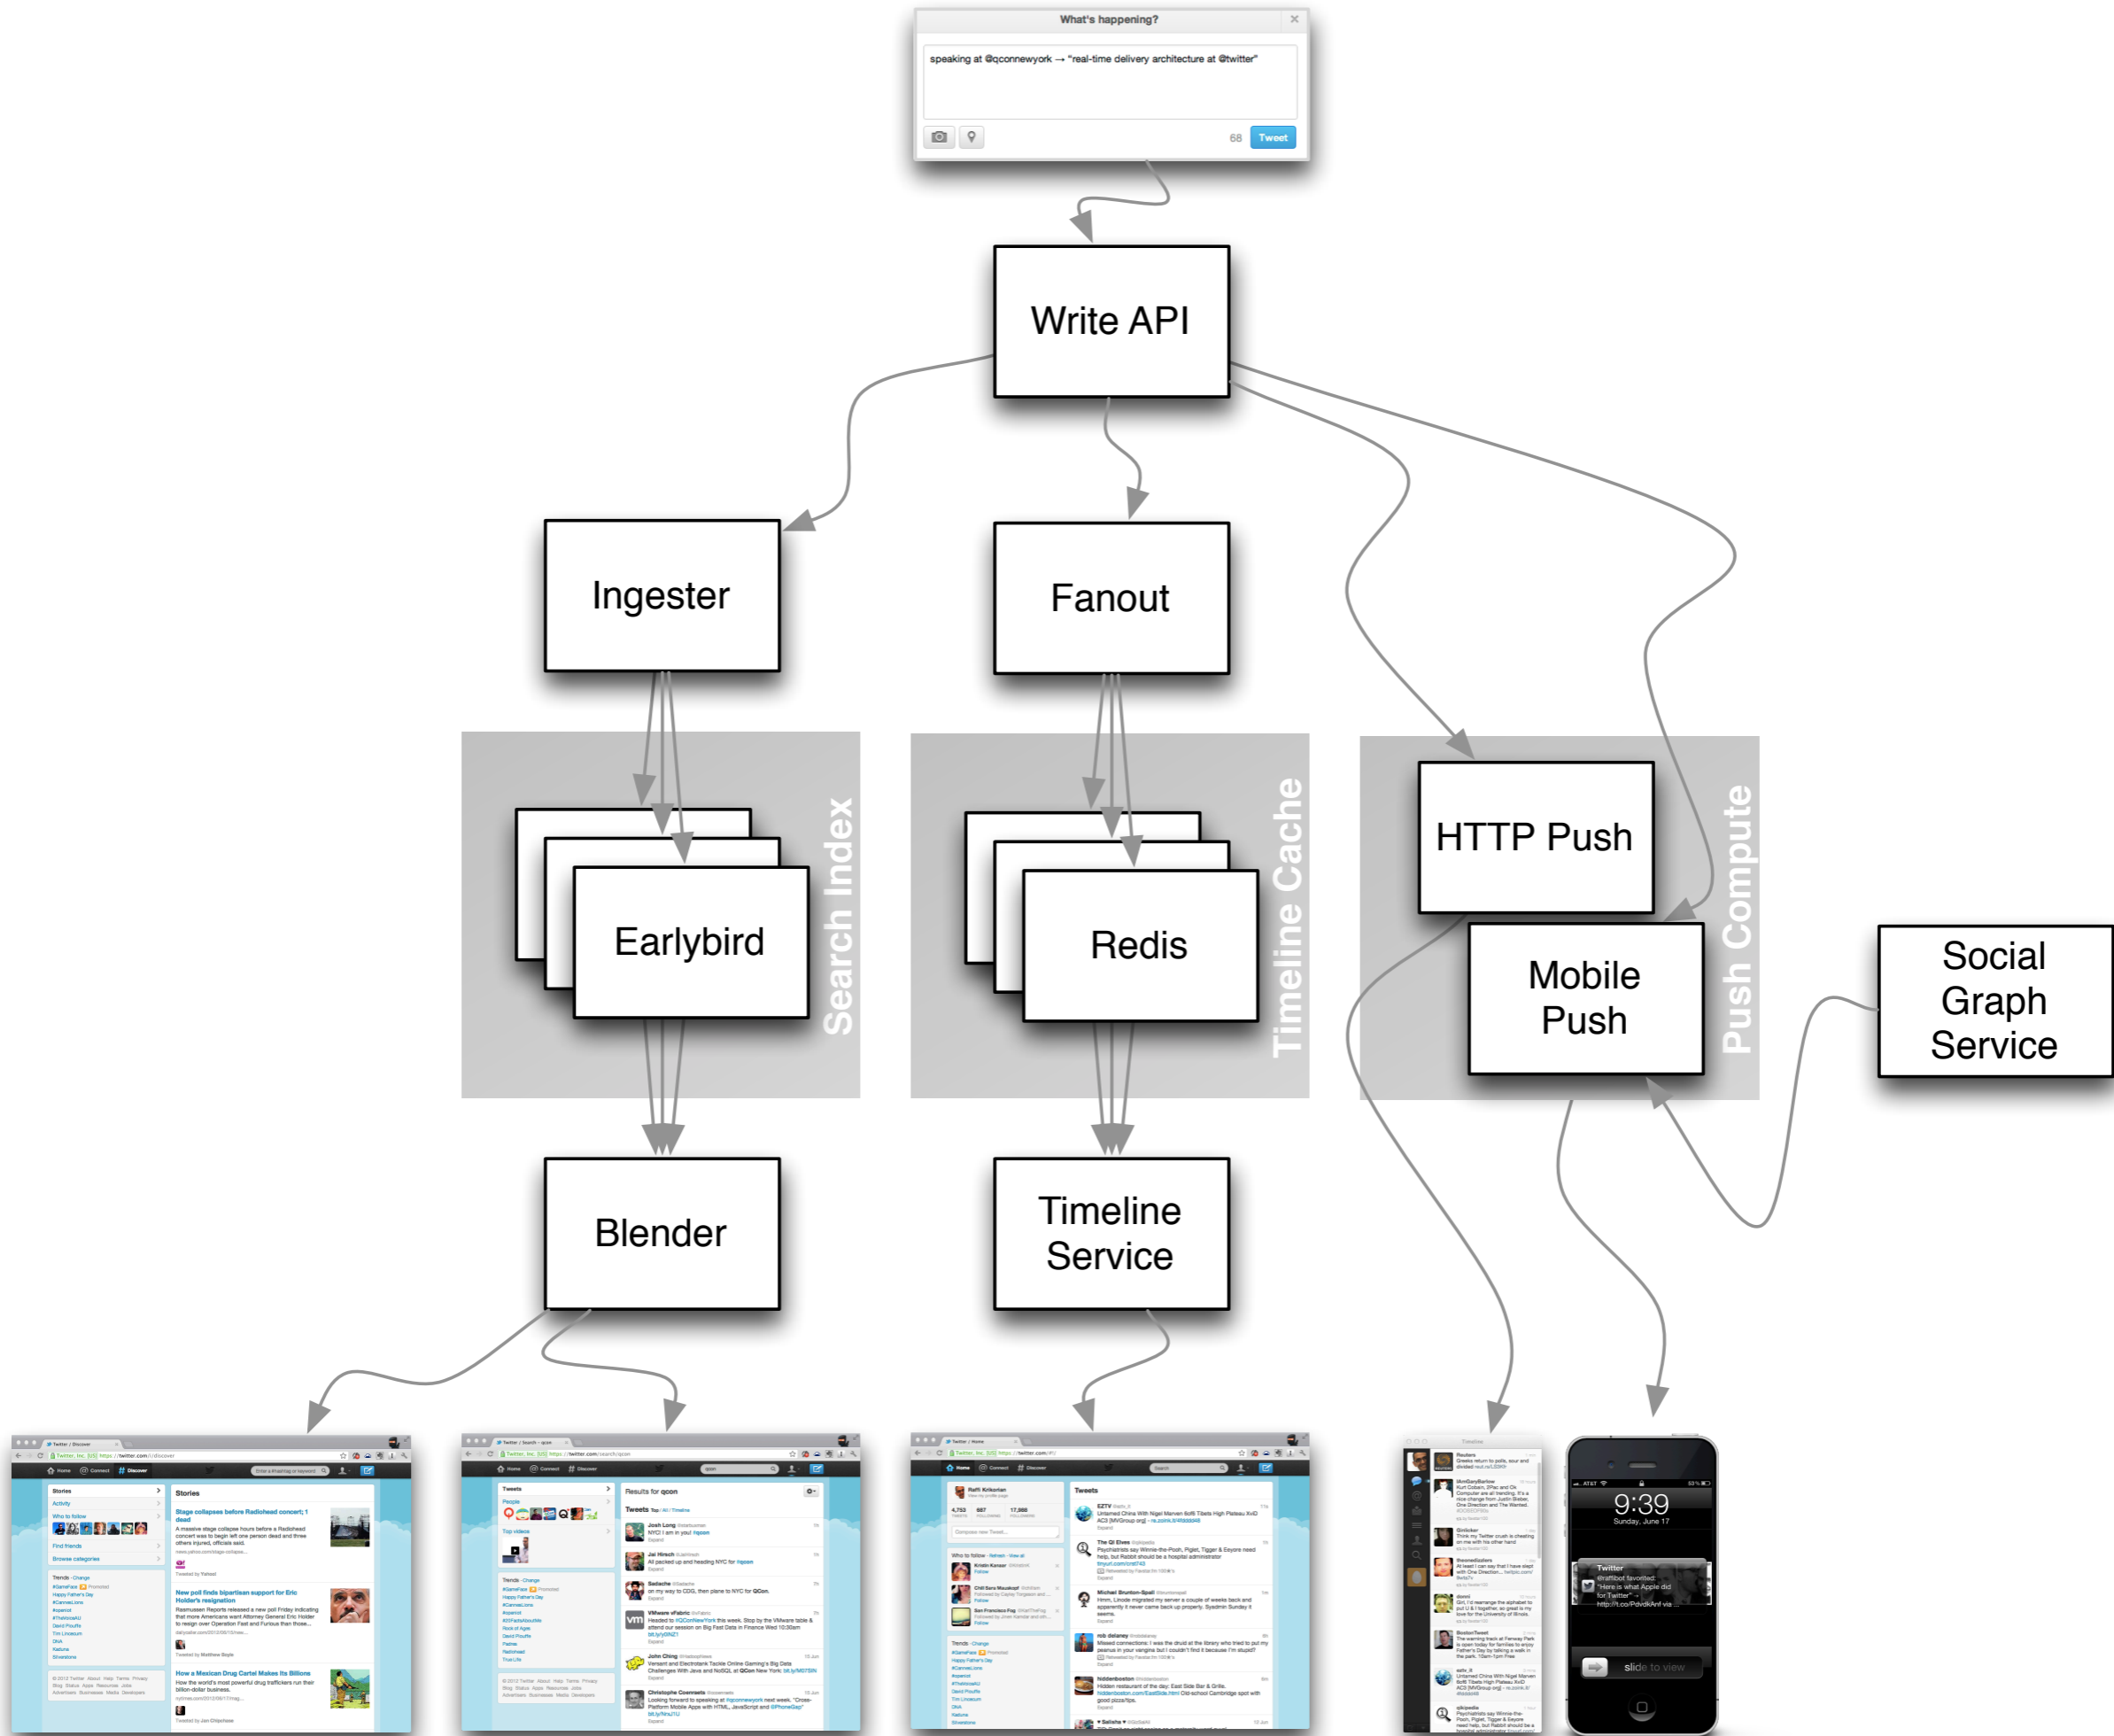
Push Compute


HTTP Push

Blender

Timeline Service




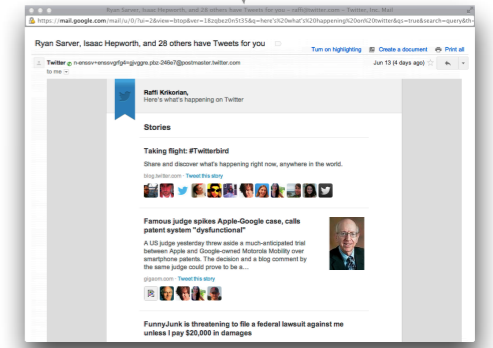
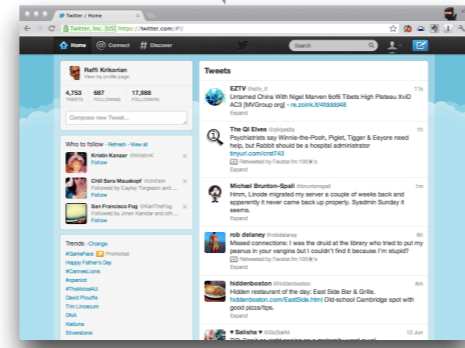
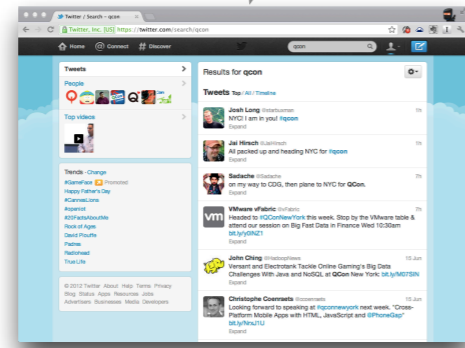

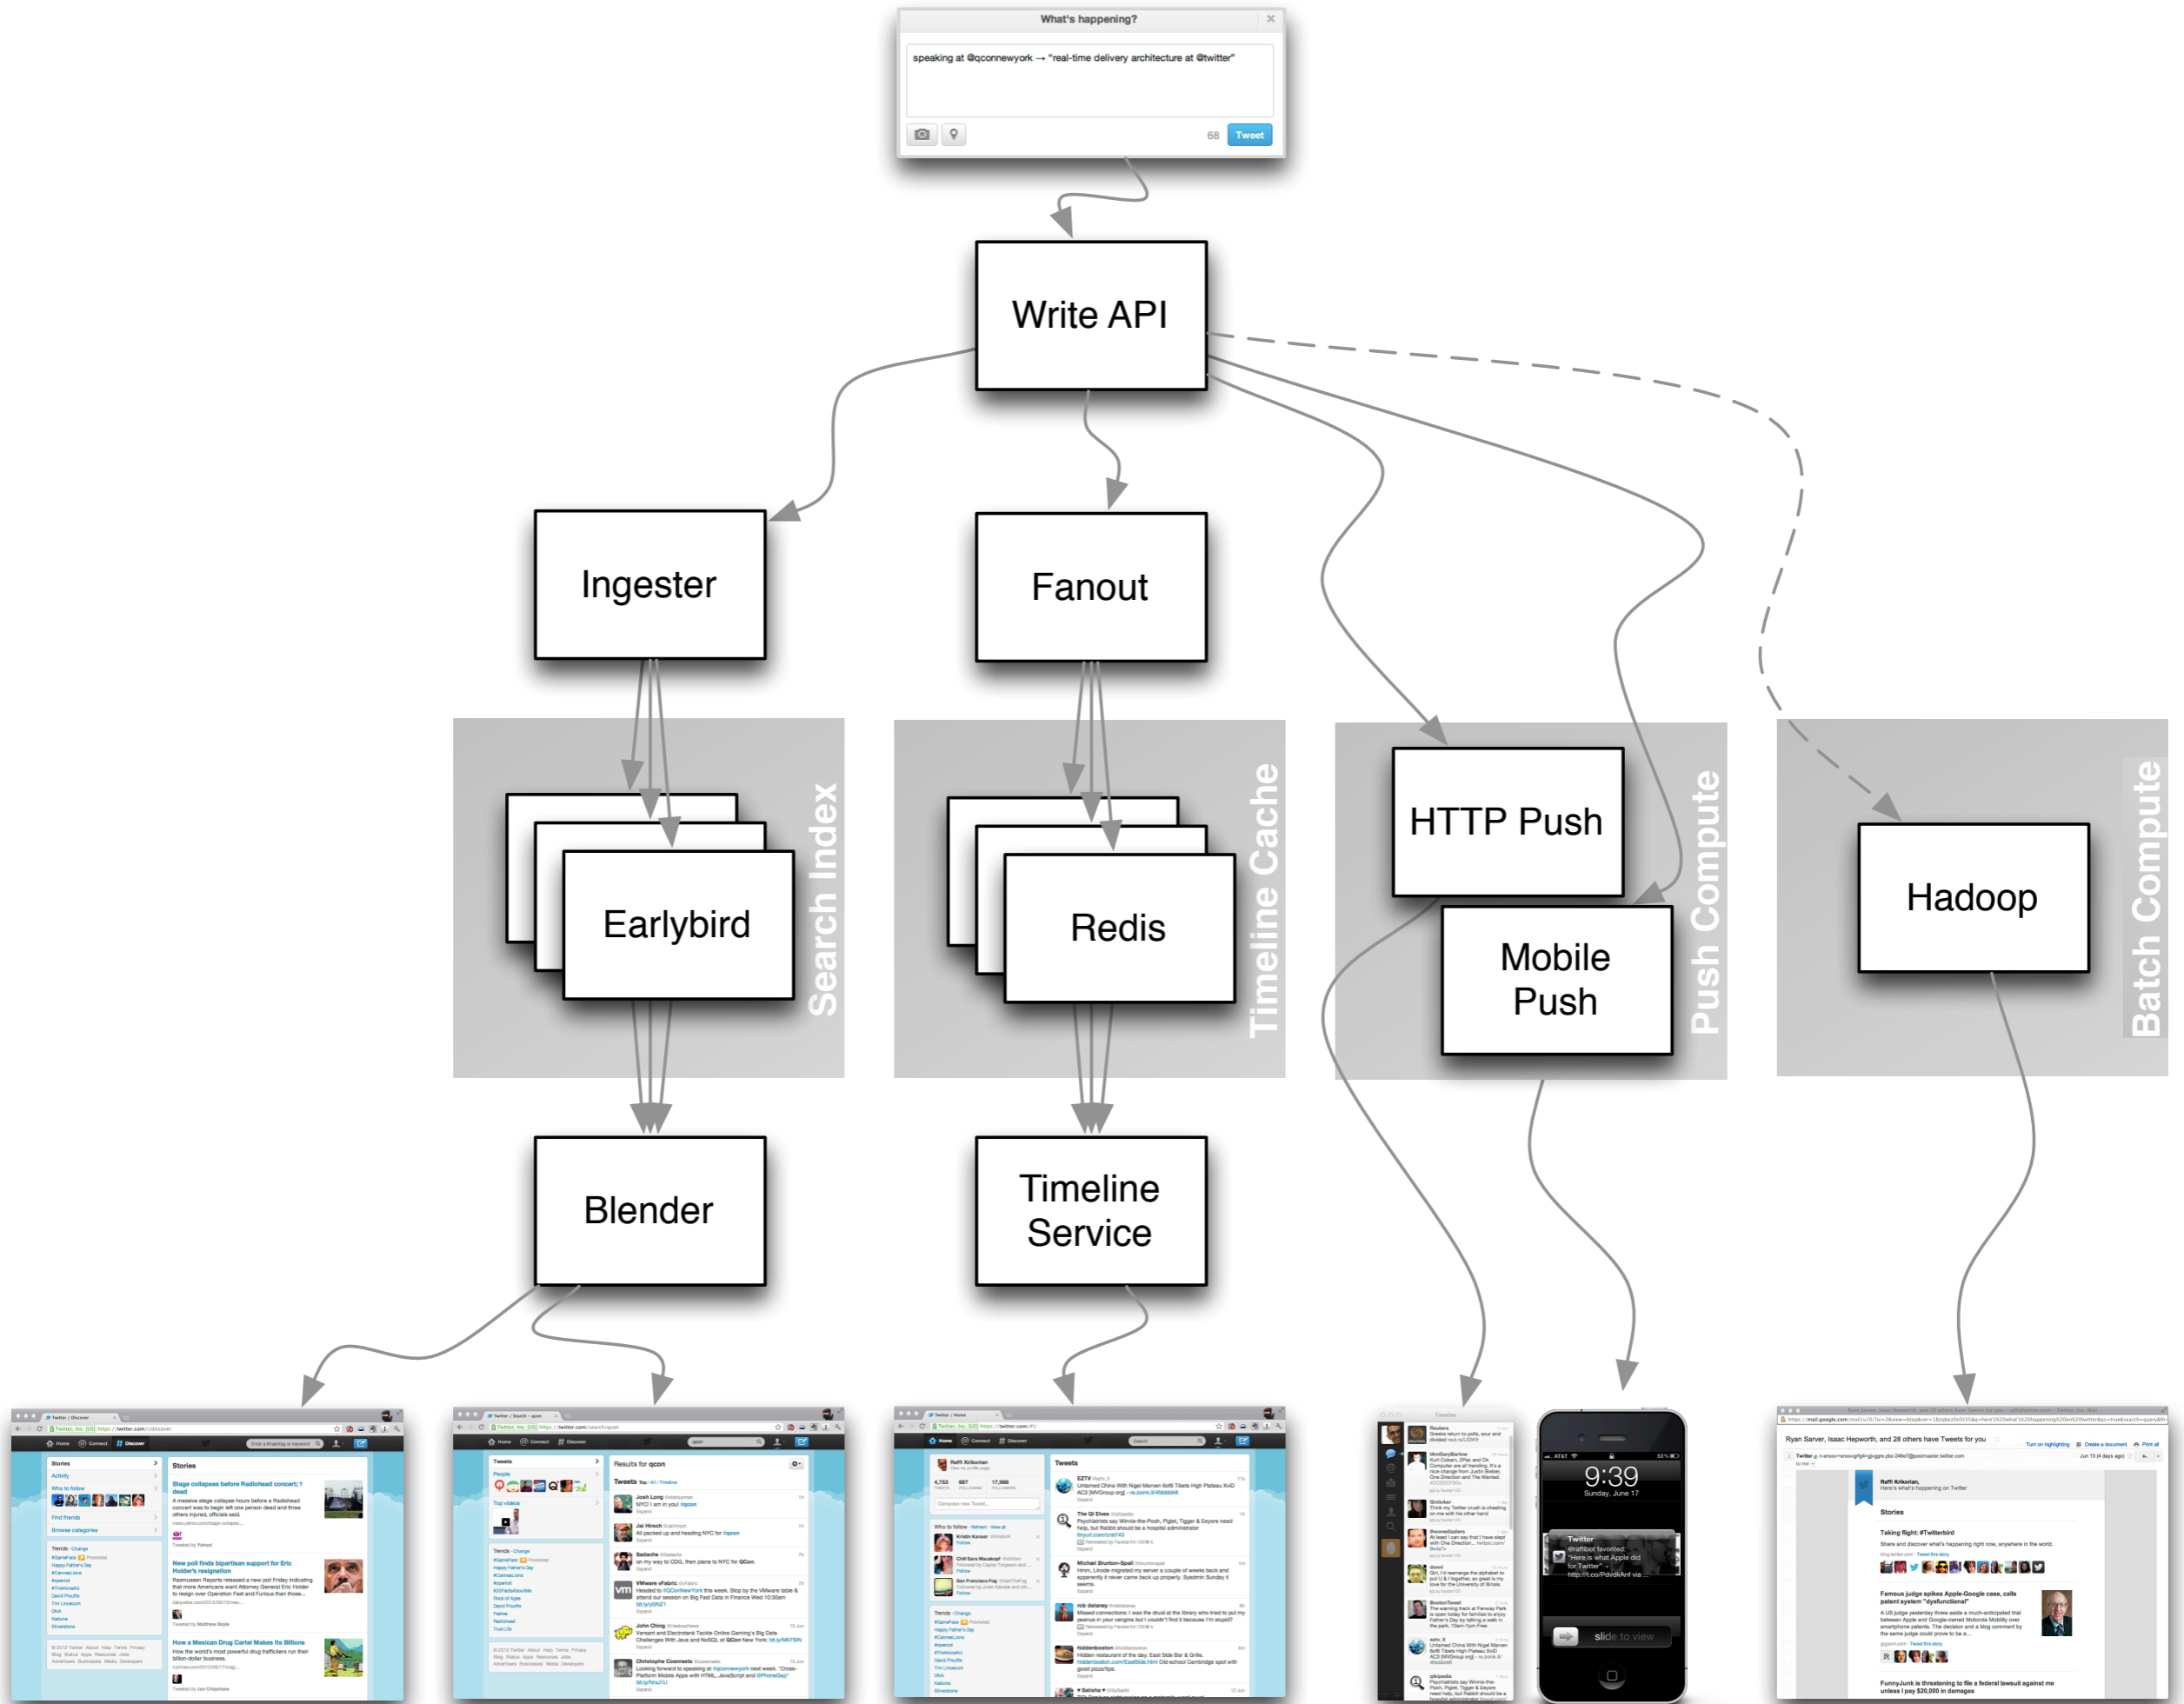


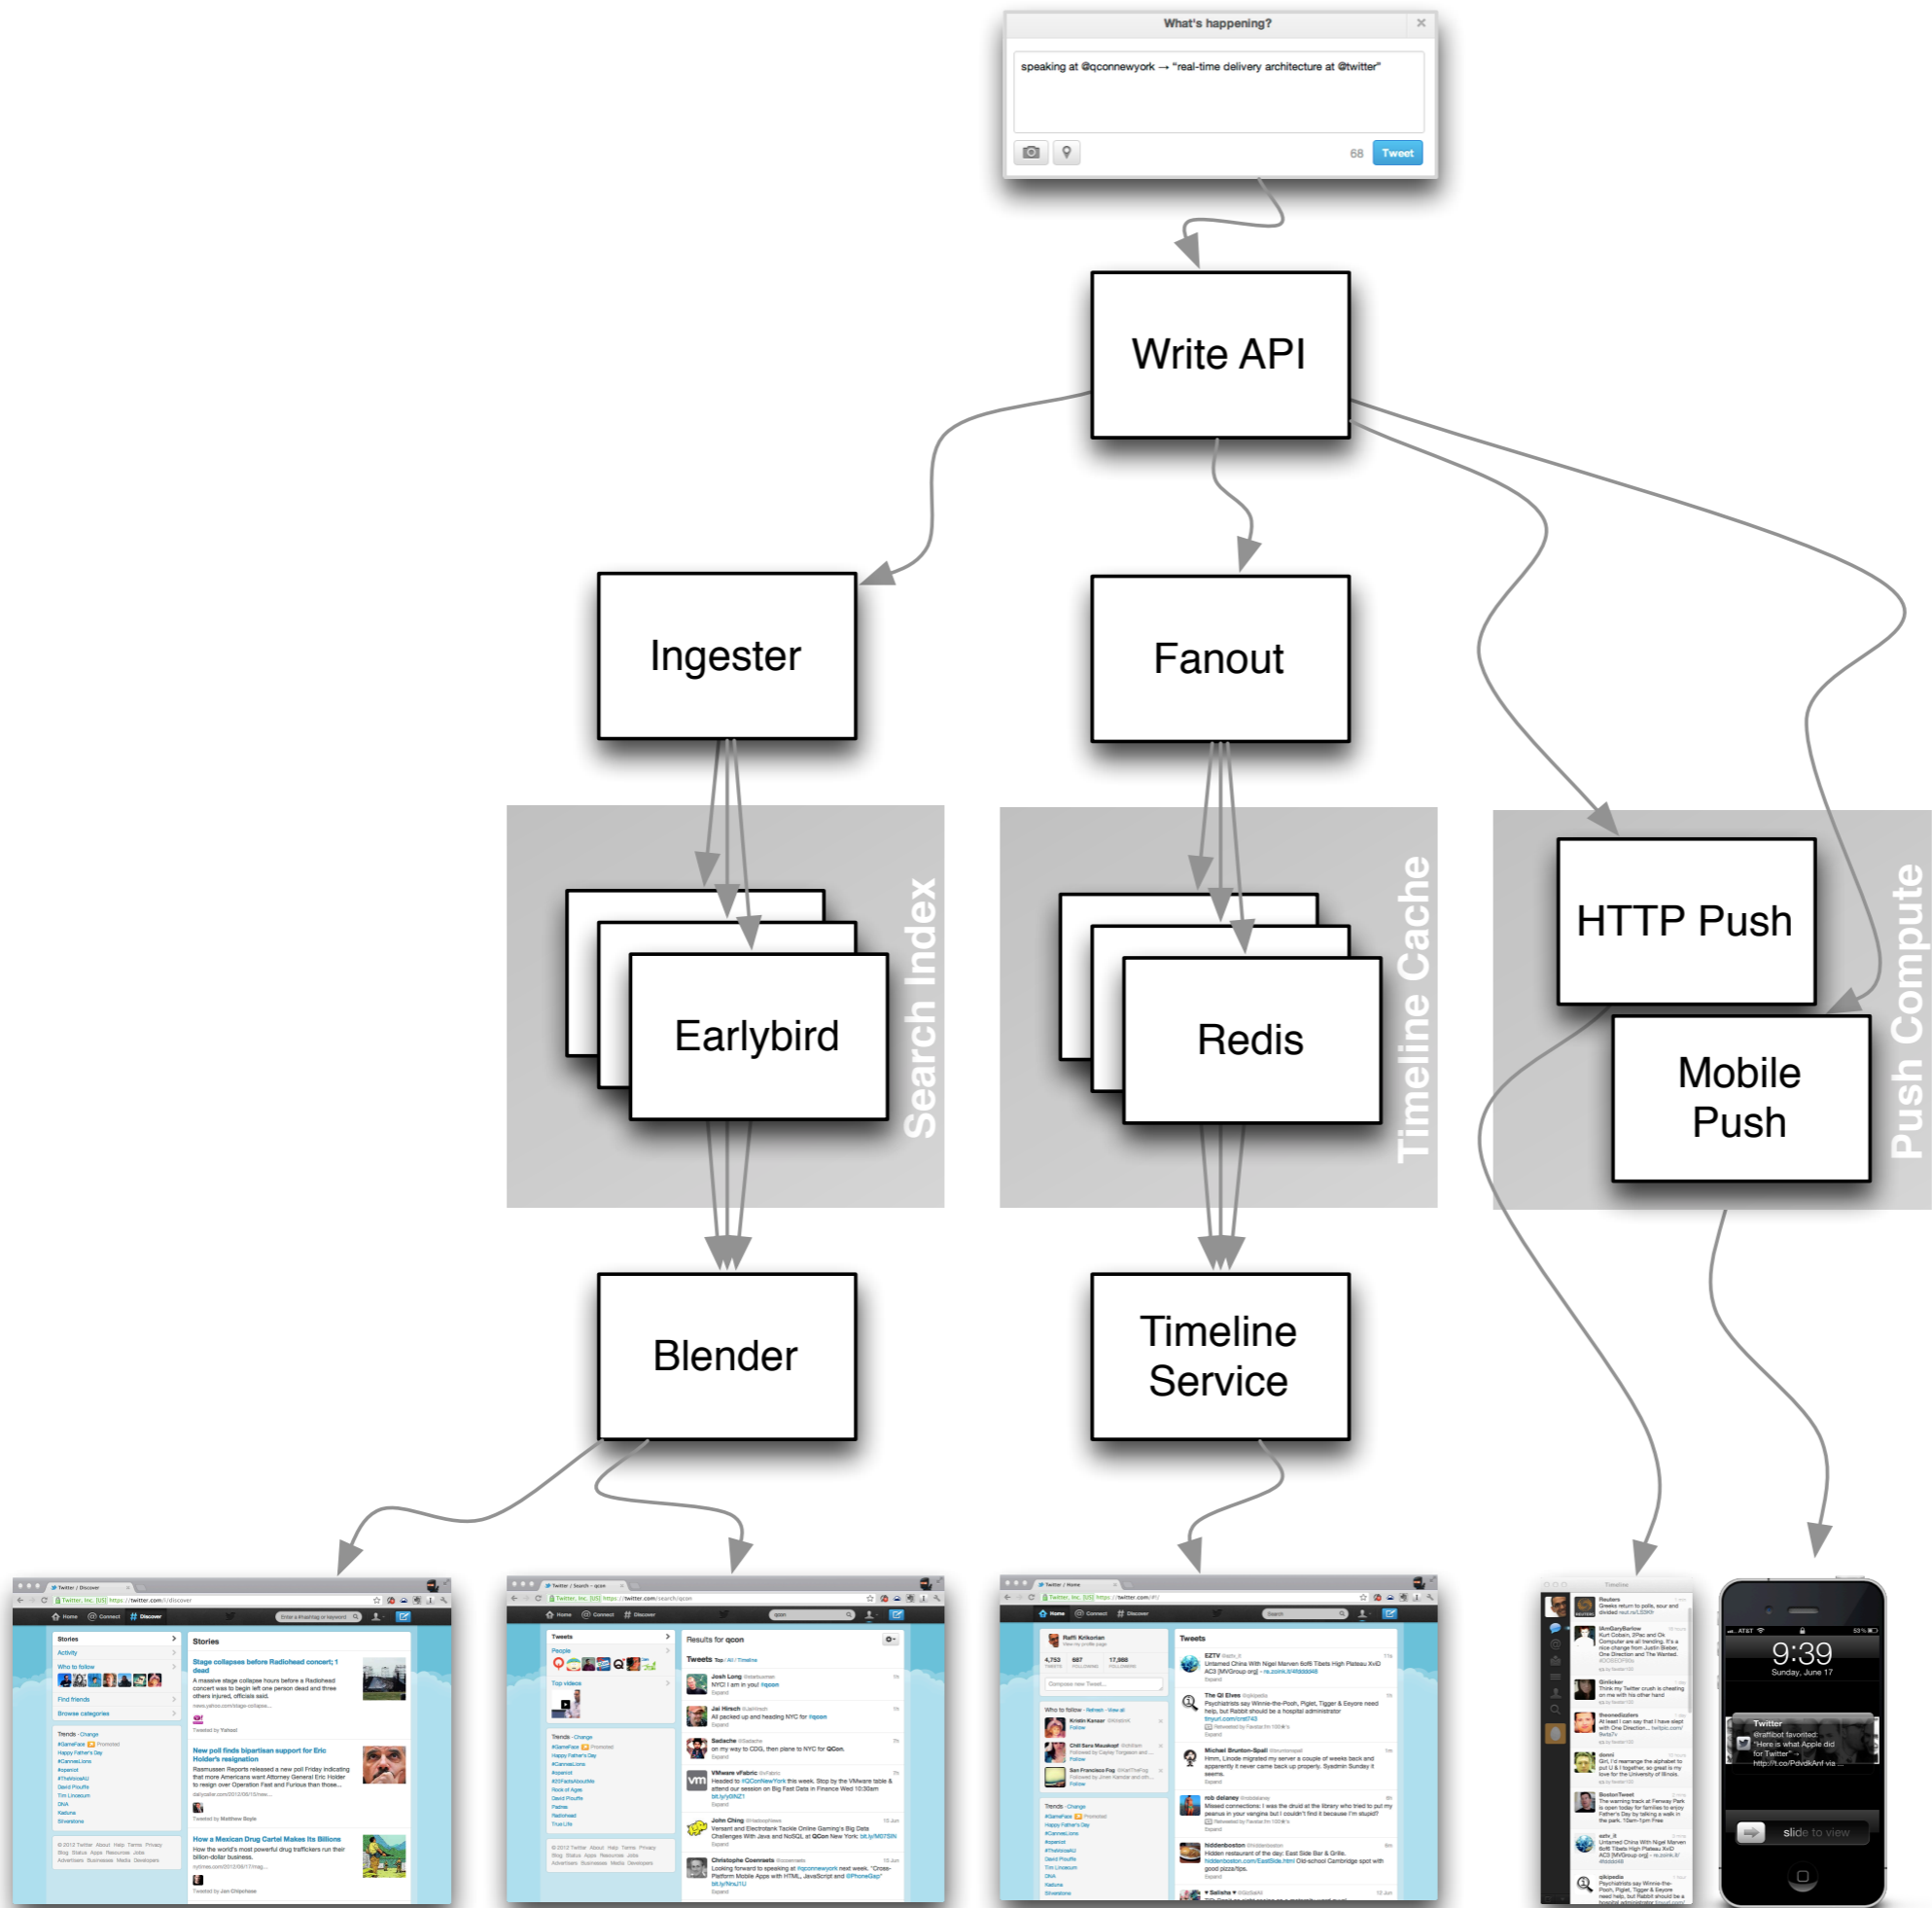




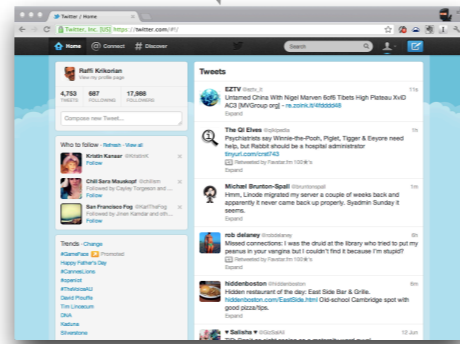
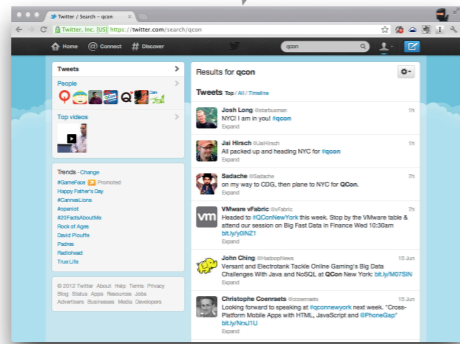
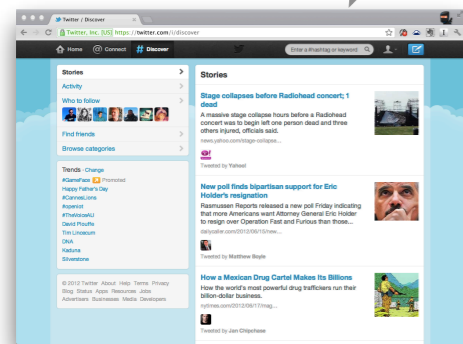
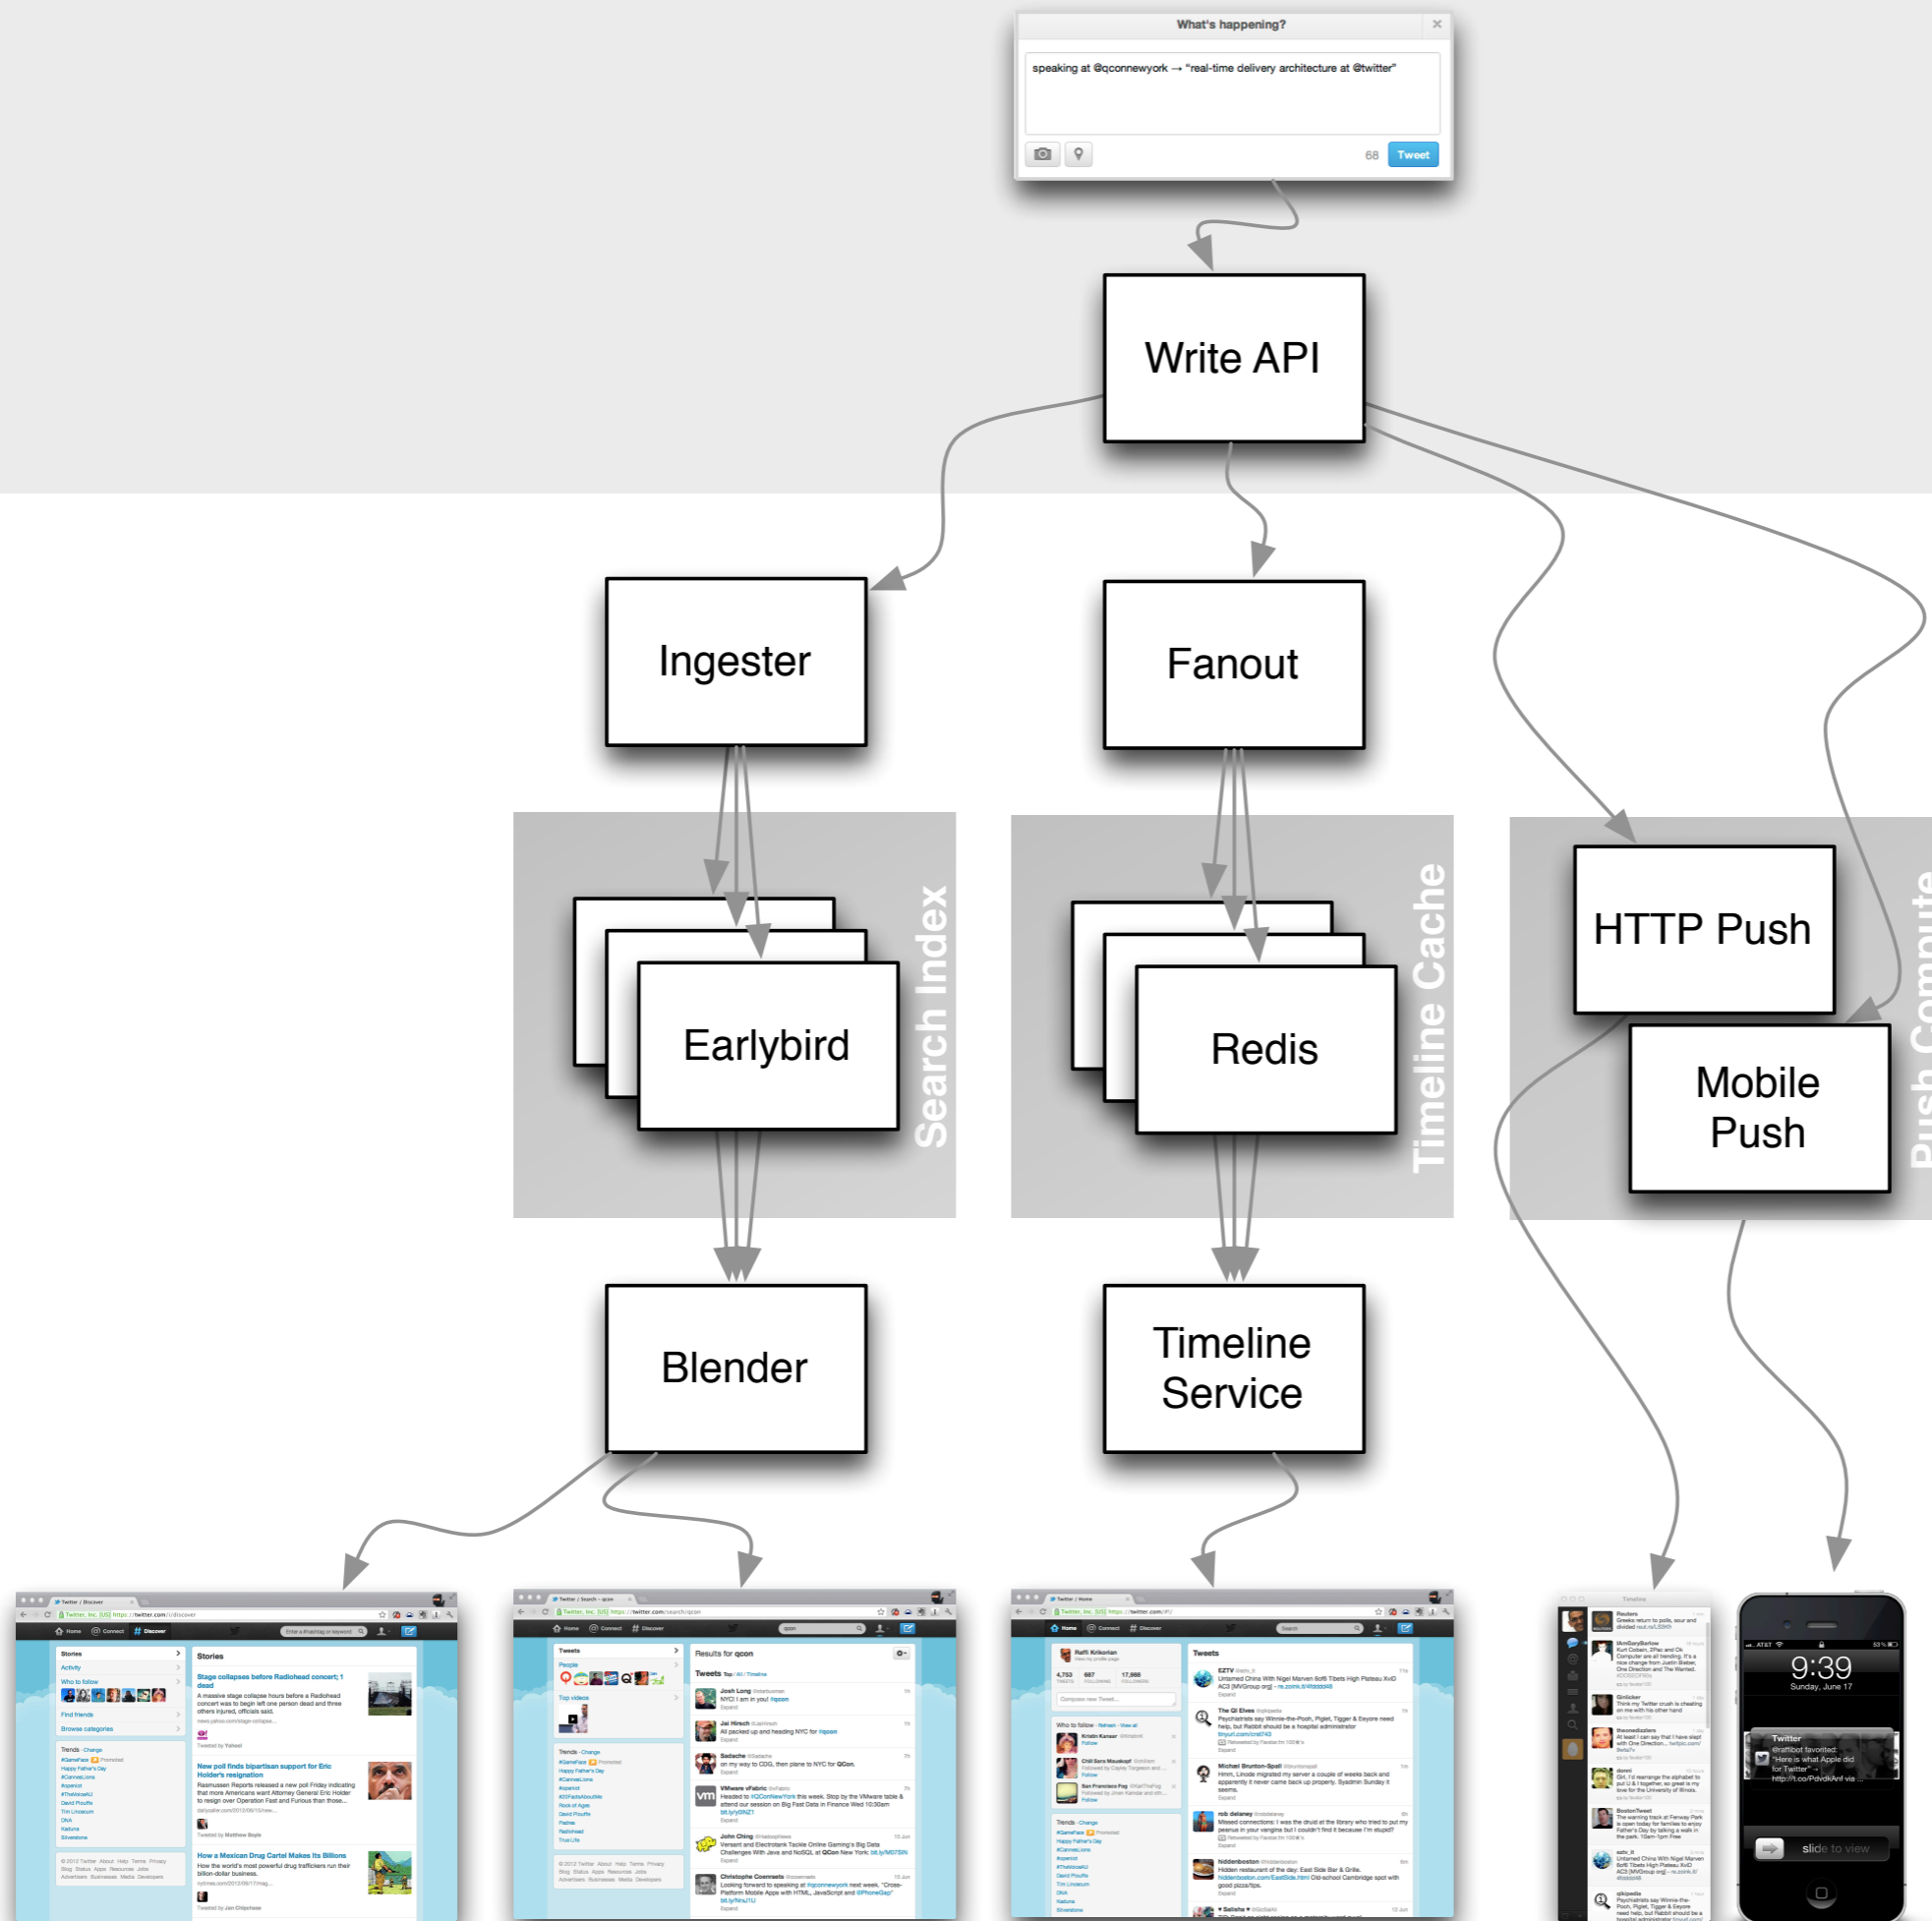
	<b>Pull</b>	<b>Push</b>
<b>Targeted</b>	twitter.com home_timeline API	User / Site Streams Mobile Push (SMS, etc.)
<b>Queried</b>	Search API	Track / Follow Streams








# Synchronous Path





Write API

Ingester

Fanout

Search Index

Earlybird

Timeline Cache

Redis

HTTP Push

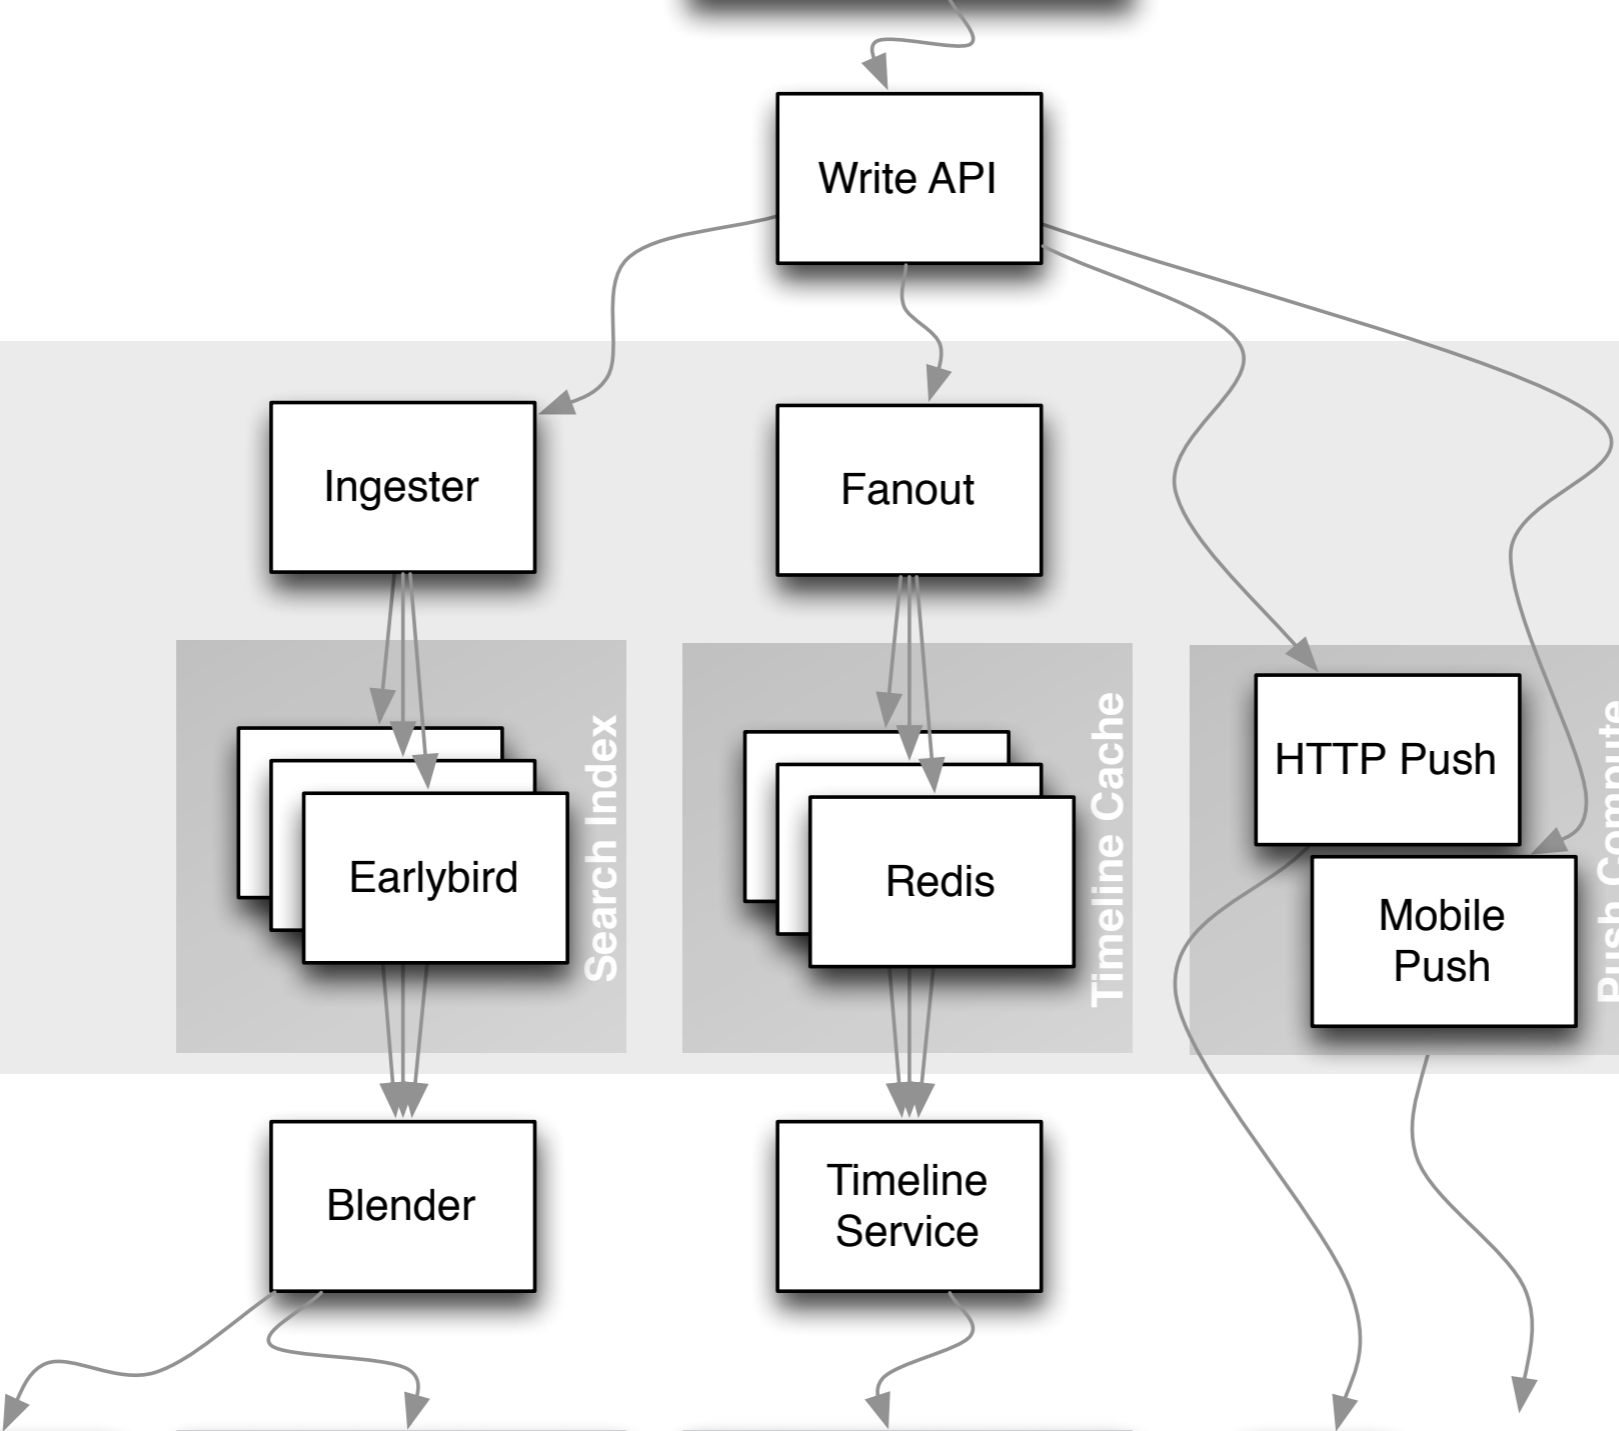
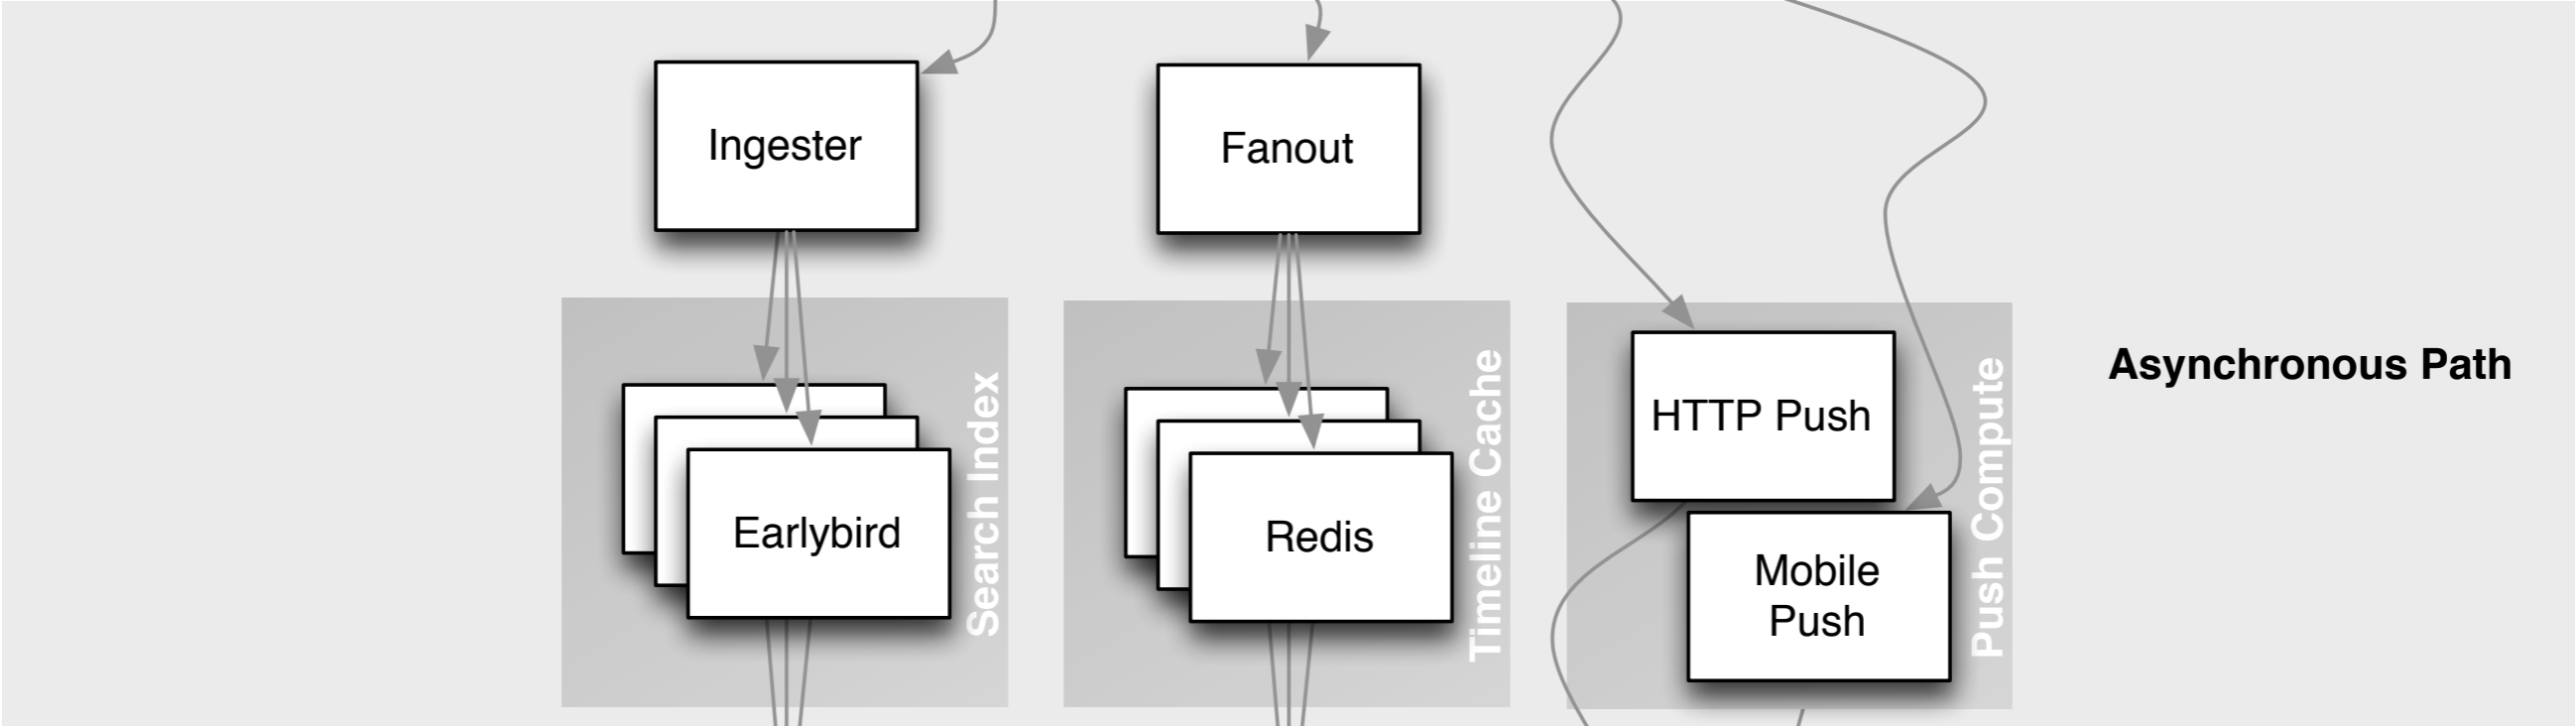
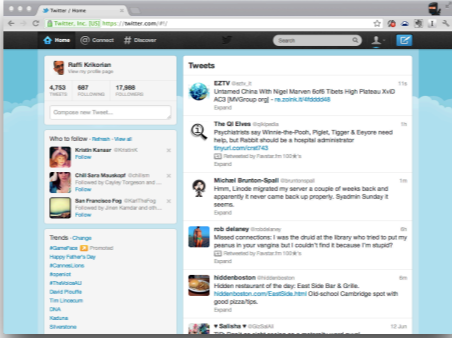

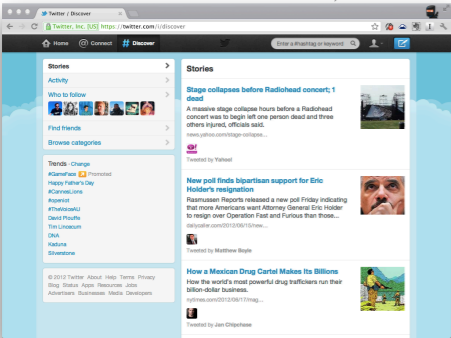
Mobile Push


Push Compute

Asynchronous Path

Blender

Timeline Service





Write API

Ingester

Fanout

Search Index

Earlybird

Timeline Cache

Redis

Push Compute


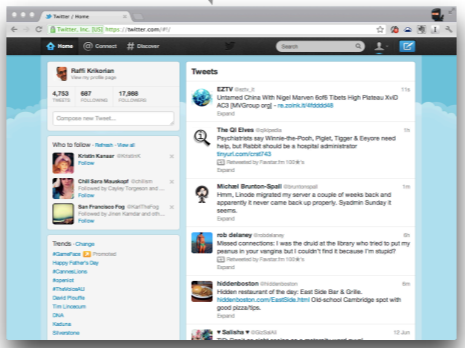
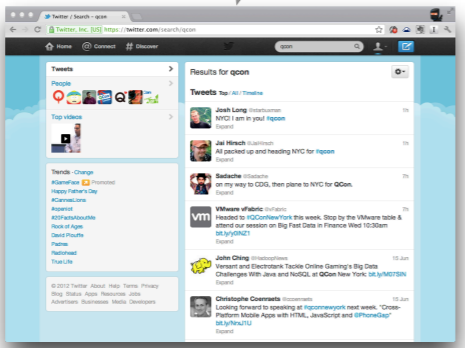
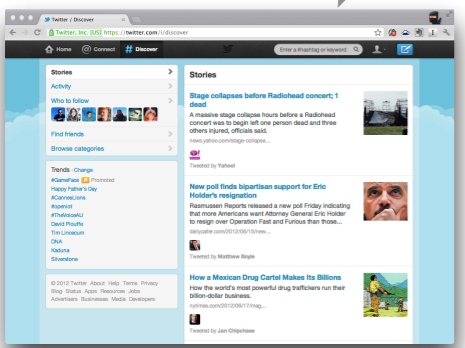
HTTP Push

Mobile Push


Blender

Timeline Service

Query Path







Write API

Ingester

Fanout

Search Index  
Earlybird



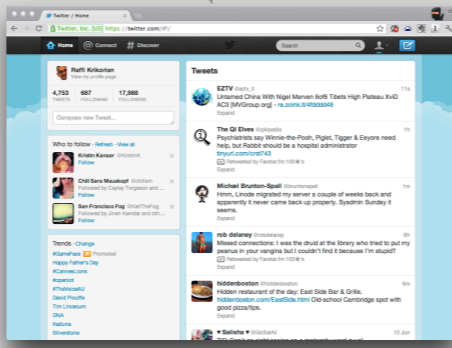
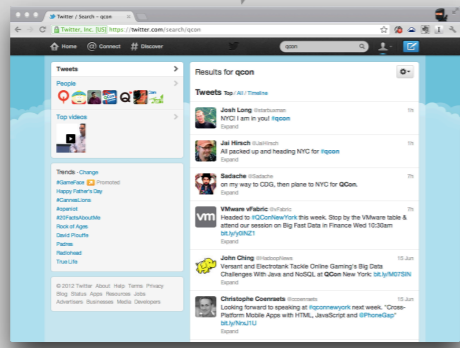
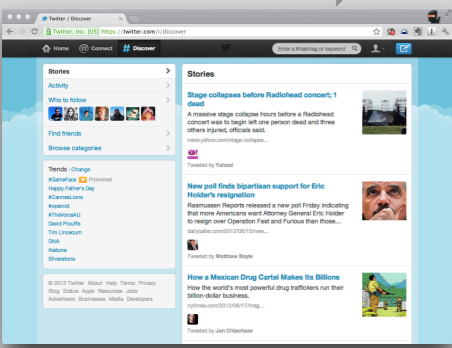
Timeline Cache  
Redis

Push Compute  
HTTP Push  
Mobile Push


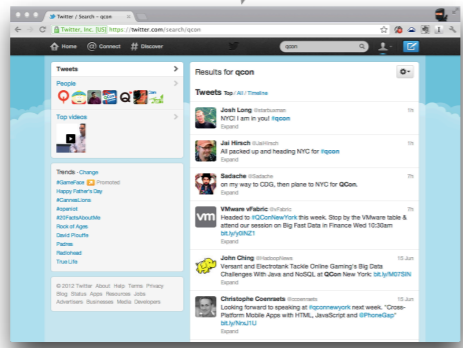

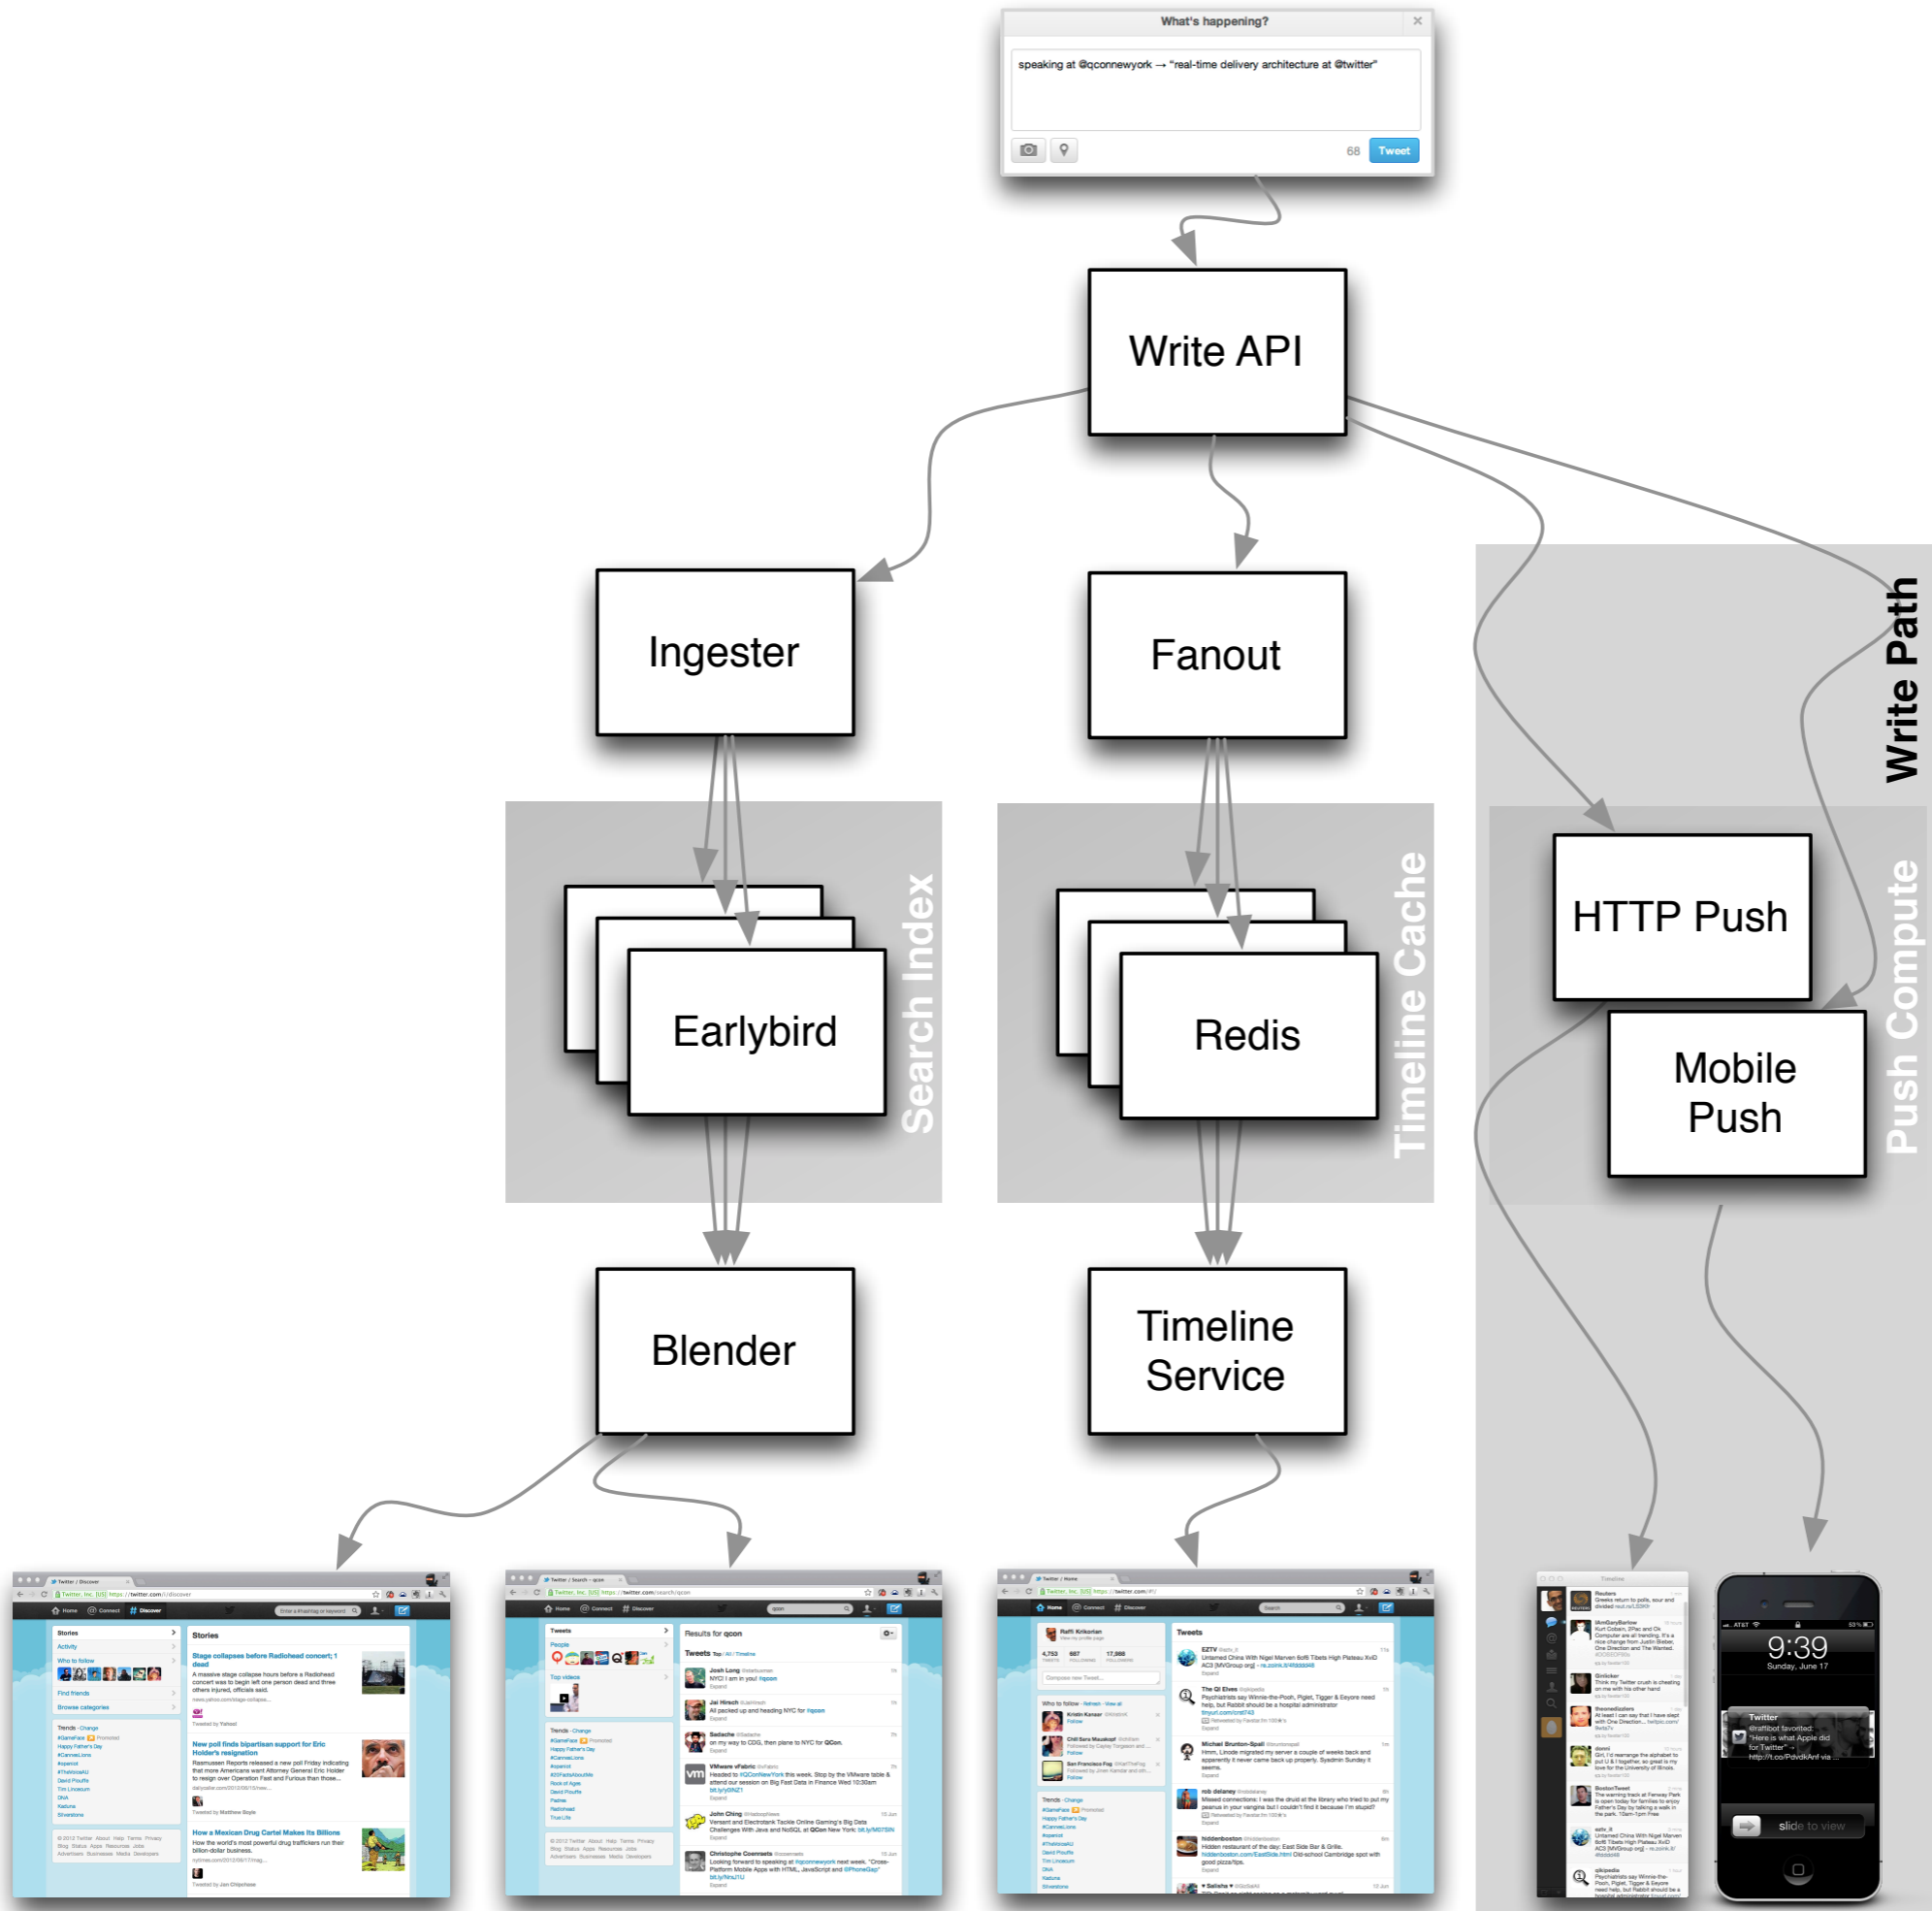
Blender

Timeline Service

Read Path









things we're trying...

What's happening?

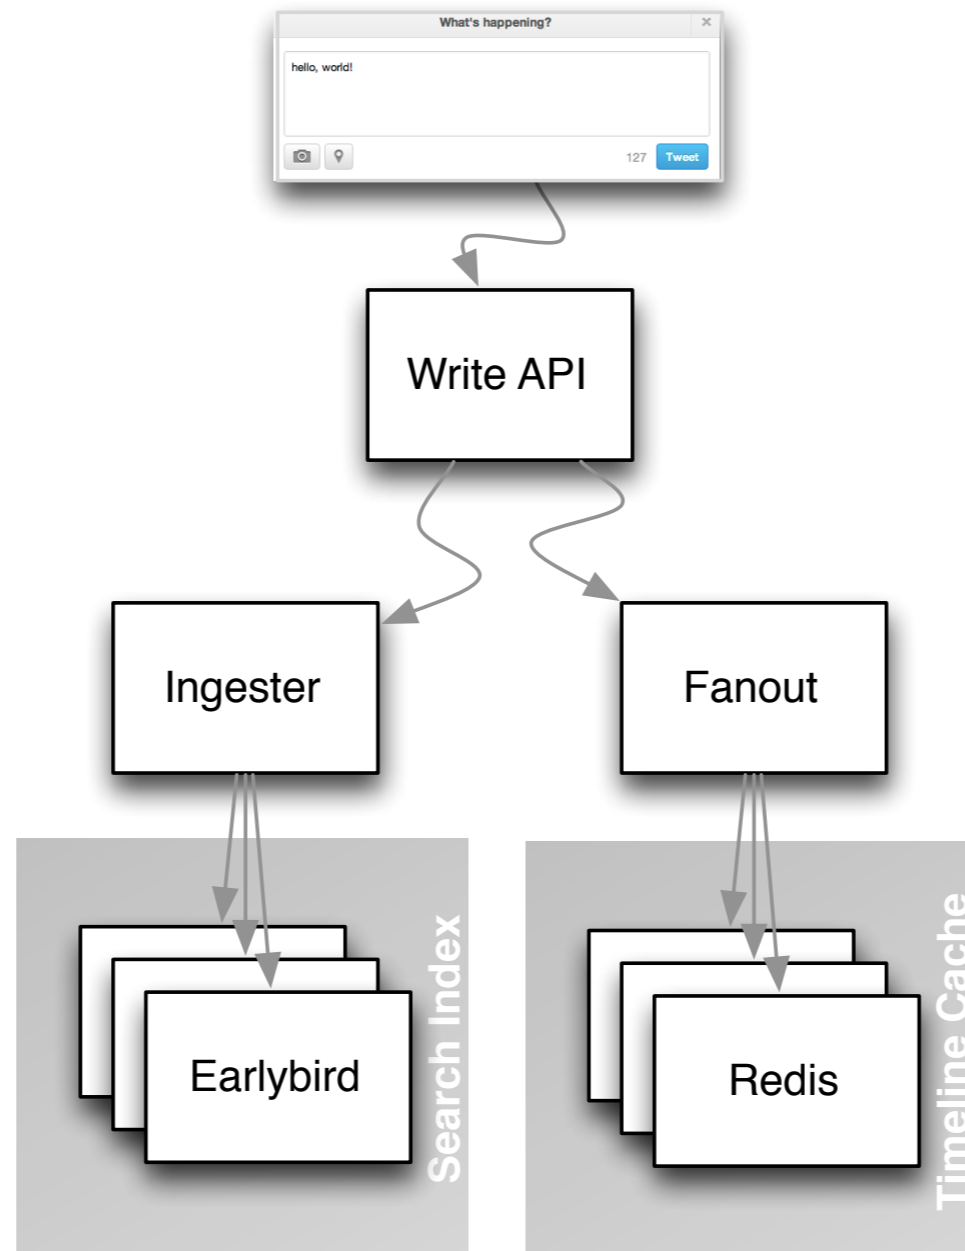


hello, world!



127

Tweet



search index

→ ['hello', 'world']

fanout index

→ [@danadanger, ...]

## User Intent

---

“Hello, world”

@raffi’s home timeline

## Query Expansion

---

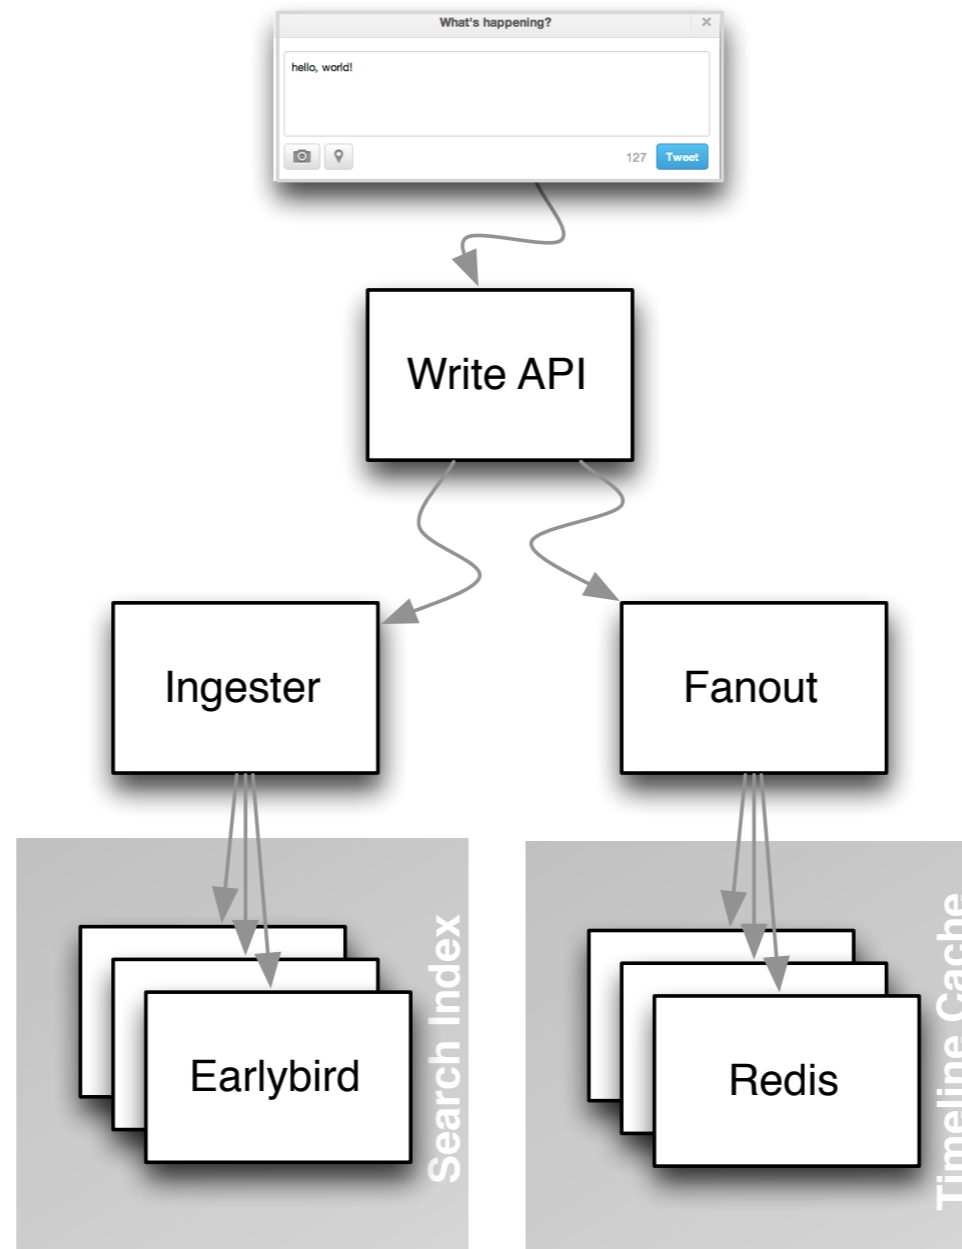
“Hello” AND “world”

user\_timeline:nelson  
OR  
user\_timeline:danadanger

# fan-in

→  $O(1)$  write

→  $O(n)$  read



# fan-out

→  $O(n)$  write

→  $O(1)$  read

## **User Intent**

---

“Hello, world”

## **Query Expansion**

---

“Hello” AND “world”

---

@raffi's home timeline

home\_timeline:raffi

## User Intent

---

“Hello, world”

@raffi’s home timeline

## Query Expansion

---

“Hello” AND “world”

home\_timeline:raffi

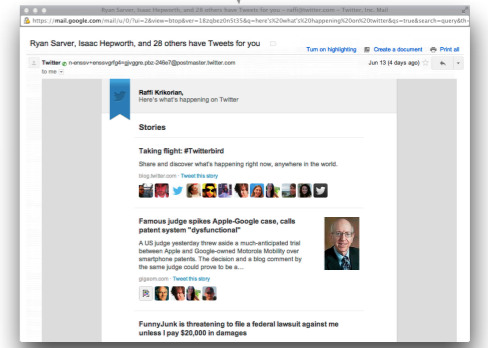

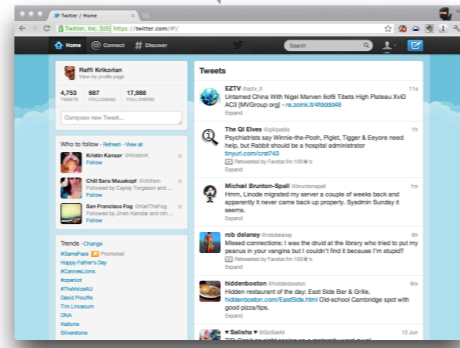
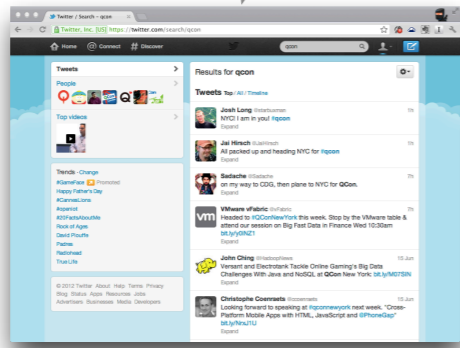


OR

user\_timeline:taylorswift13



# streaming compute

- continuous computation
- driven by the events that come into twitter
- generalizing the push mechanism



# timeline query statistics

- >150m active users worldwide
- 300k qps poll-based timelines  
@ 1ms p50 / 4ms p99
- 30k qps search-based timelines


# tweet input

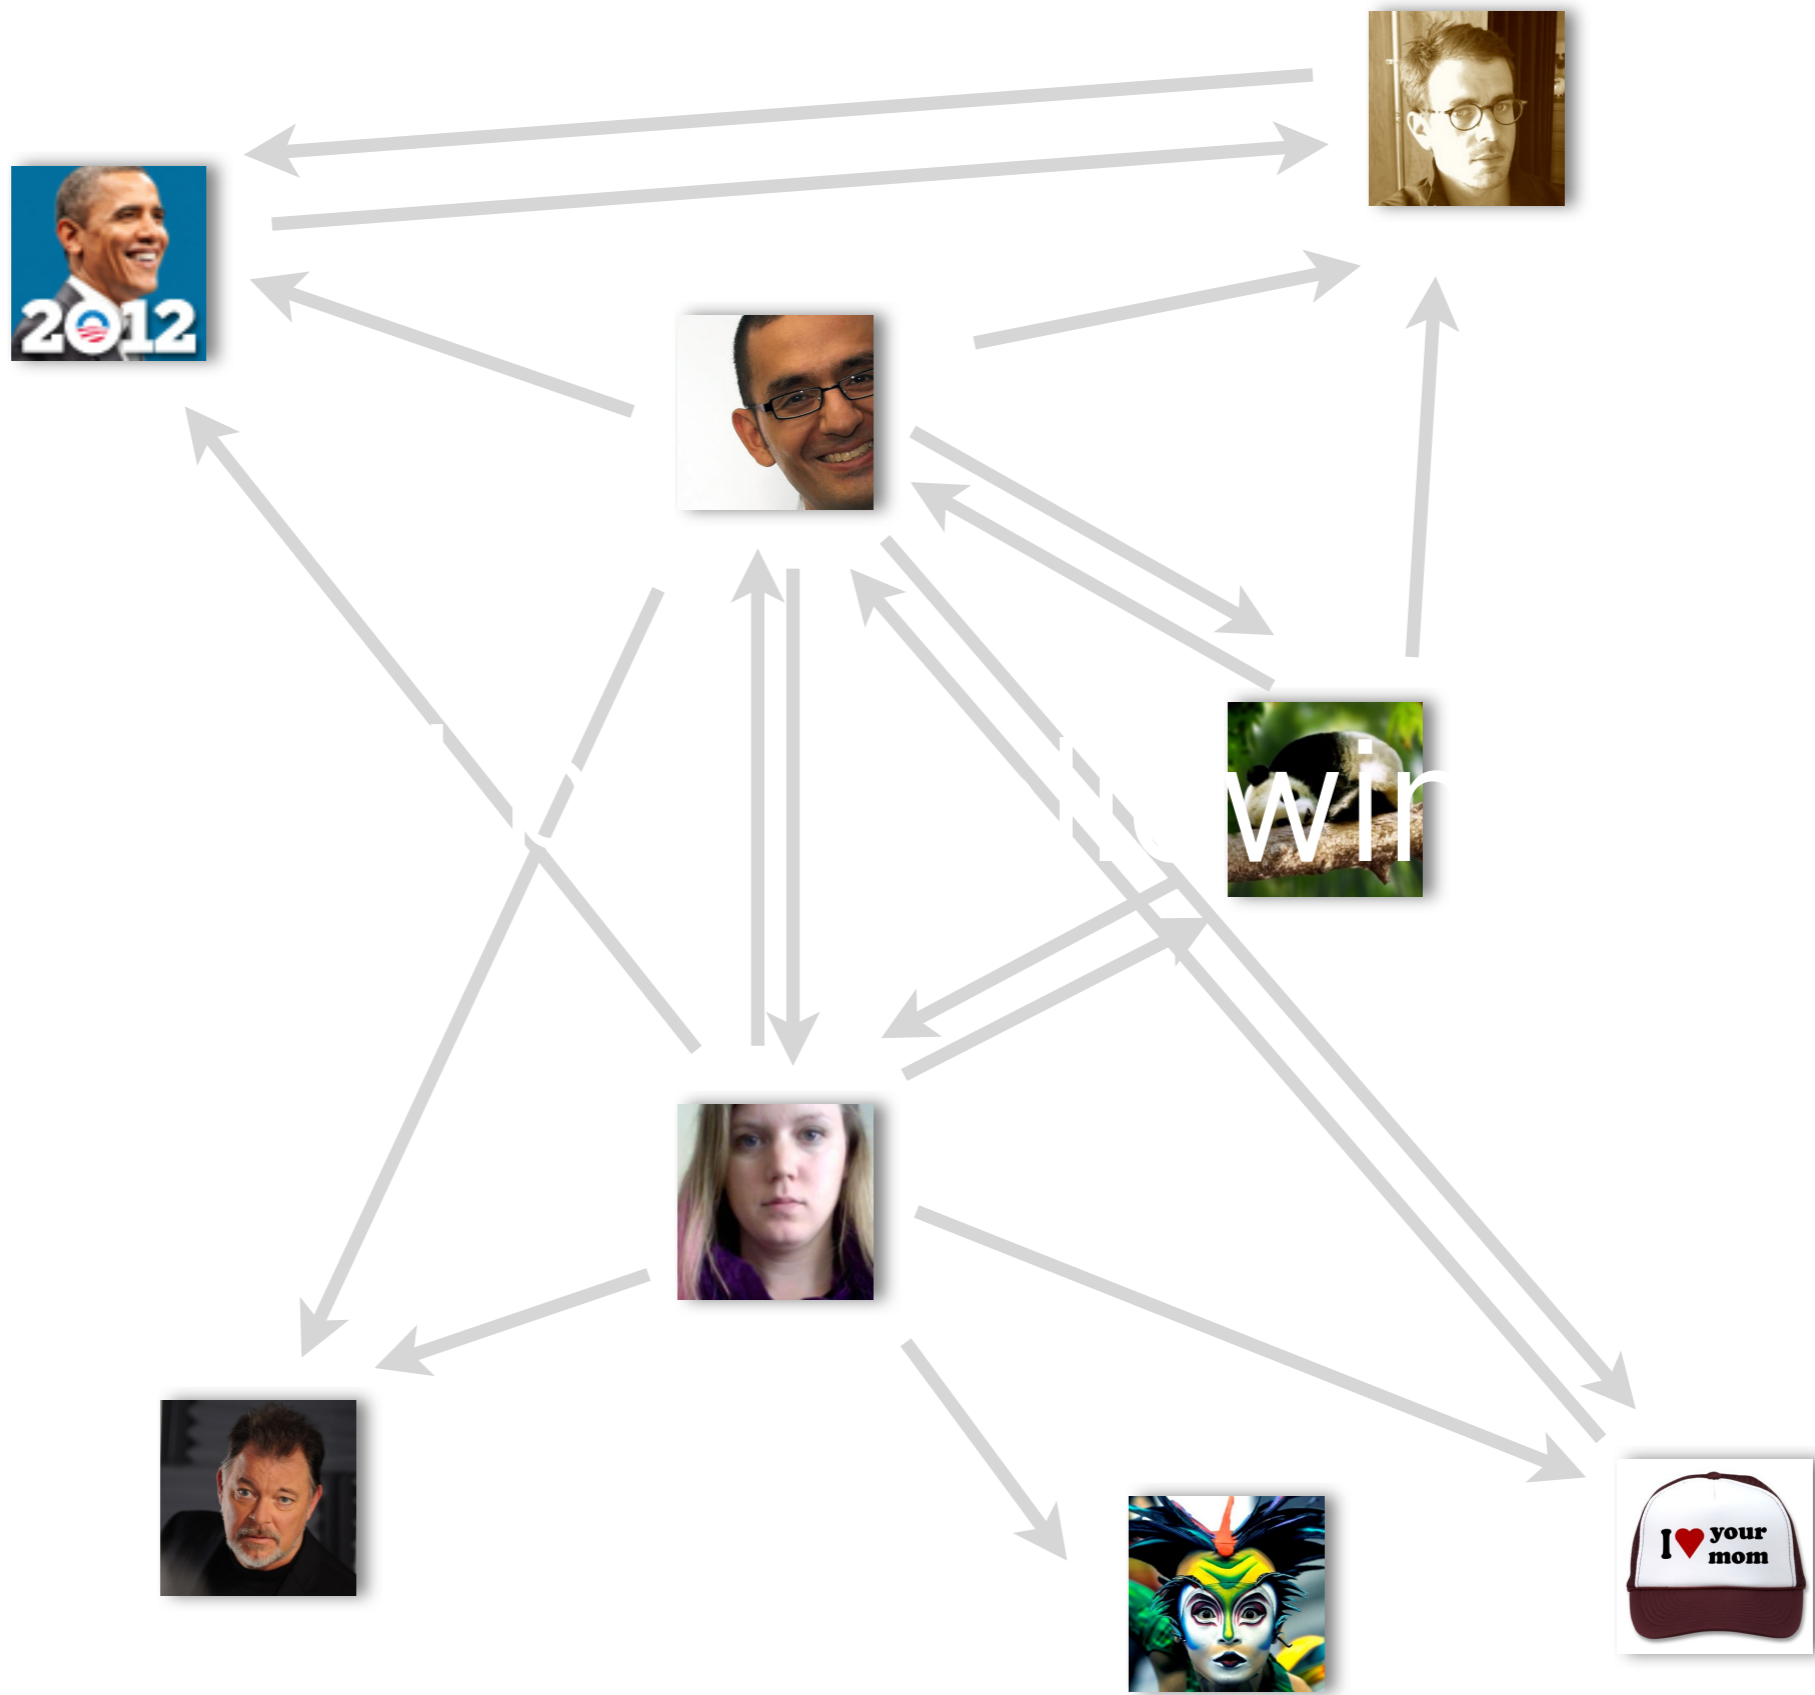
→ ~340m tweets per day

→ ~4K/sec daily average

→ ~6K/sec daily peak

→ >10K/sec during large events






# timeline delivery statistics

→ 26b deliveries / day (~18m / min)

→ 3.5 seconds @ p50 to deliver to 1m

→ ~300k deliveries / sec



thanks!