

facebook

Facebook Cache Invalidation Pipeline

Melita Mihaljevic (melitam@fb.com)
Production Engineer

Agenda


1 Introduction

2 What makes keeping cache consistent challenging?

3 Cache infrastructure and Invalidation Pipeline


4 Q&A

Introduction



What makes keeping cache
consistent challenging ?

Cache (in)consistency



Cache (in)consistency



Challenges

- Rapidly evolving product
- Massive scale
- Social Graph at the center



Evolving product

[thefacebook]
home search global social net invite faq logout

Scott Pearson's Profile (This is you) Scranton

quick search go

My Profile [edit]
My Groups
My Friends
My Messages
My Away Message
My Mobile Info
My Account
My Privacy

Picture [edit]

Information [edit]

Account Info:
Name: Scott Pearson, BS
Member Since: January 12, 2005
Last Update: February 3, 2005 [edit]

Basic Info: [edit]
Email: pearsons2@scranton.edu
Status: Alumnus/Alumna
Sex: Male
Year: 2004
Concentration: Computing Sciences, Mathematics

Phone: 570.499.4818
High School: Dunmore HS '00

Extended Info: [edit]
Screenname: ScottiePP7
Looking For: Friendship, Dating, A Relationship, Random play, Whatever I can get
Interested In: Women
Relationship Status: Single
Political Views: Liberal
Interests: Drinking, Football, Basketball, Tennis, saying you'll have that

Connection
This is you.

Access
Scott is currently logged in from a non-residential location.

Other Schools [edit]

facebook Search for anything on Facebook Janet Aronica Home

Janet Aronica Update Info Activity Log

Marketing Chick at Shareaholic
Studied Communication/Journalism at St. John Fl...
Lives in Boston, Massachusetts
From New York

About Friends 741 Photos 324 Map 26 Likes 424

Status Photo Place Life Event

What's on your mind?

Janet Aronica 2 seconds ago

741 Friends

Activity Recent

Janet listened to Sexy Ch... Guetta on Spotify. Like · Comment

Janet joined 7th Annual Cookout.

Janet became friends with Lunch

Nicole Gregory commented on this.

News Feed

Suggested Item

J. Somerfields Sponsored

With the J. Somerfields App, the latest looks are always at your fingertips.

Shop Now

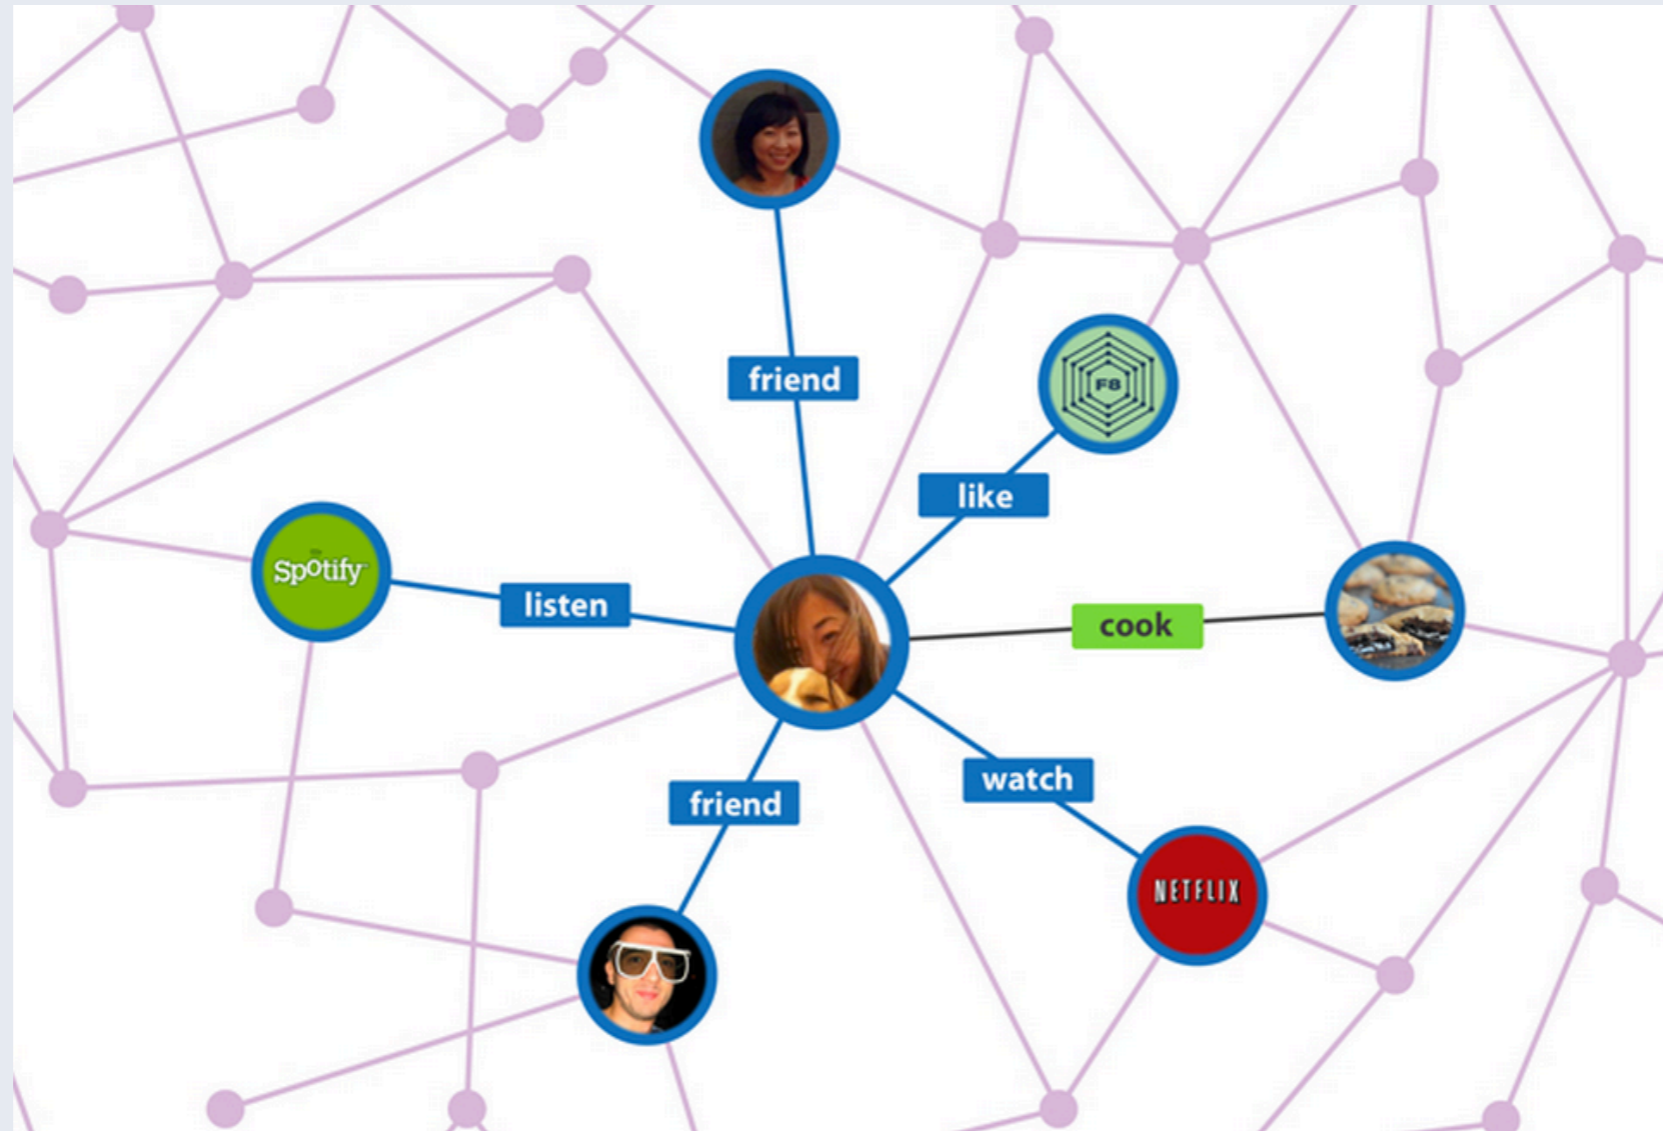
Ch2

Massive scale

- 1.44 billion users
- 829 million daily active users
- 654 million mobile daily active users
- hundreds of thousands of servers



Social graph



People are only one dimension of the social graph

Cache infra and Invalidation Pipeline

TAO

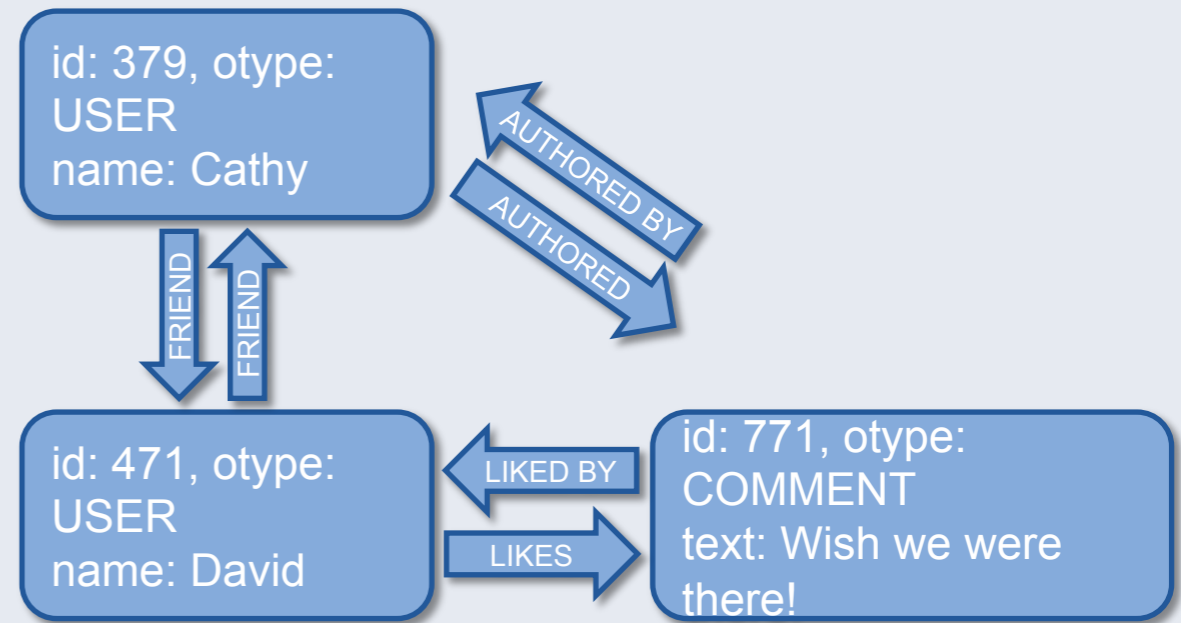
The social graph

Representation for most FB data

- Nodes are “objects”
- Edges are “associations”

Complex queries

- Rendering one page requires multiple accesses to the graph
- Fast access required



TAO

Two level, write-through cache with graph semantics

Server-side graph abstraction

- Simplification of client code

Two level cache

- Leaders coordinate access to DB
- Followers cache most data
- Graceful failover

Cache consistency inside data center

Web servers




Follower cache




Leader cache

refill X

range get



Database

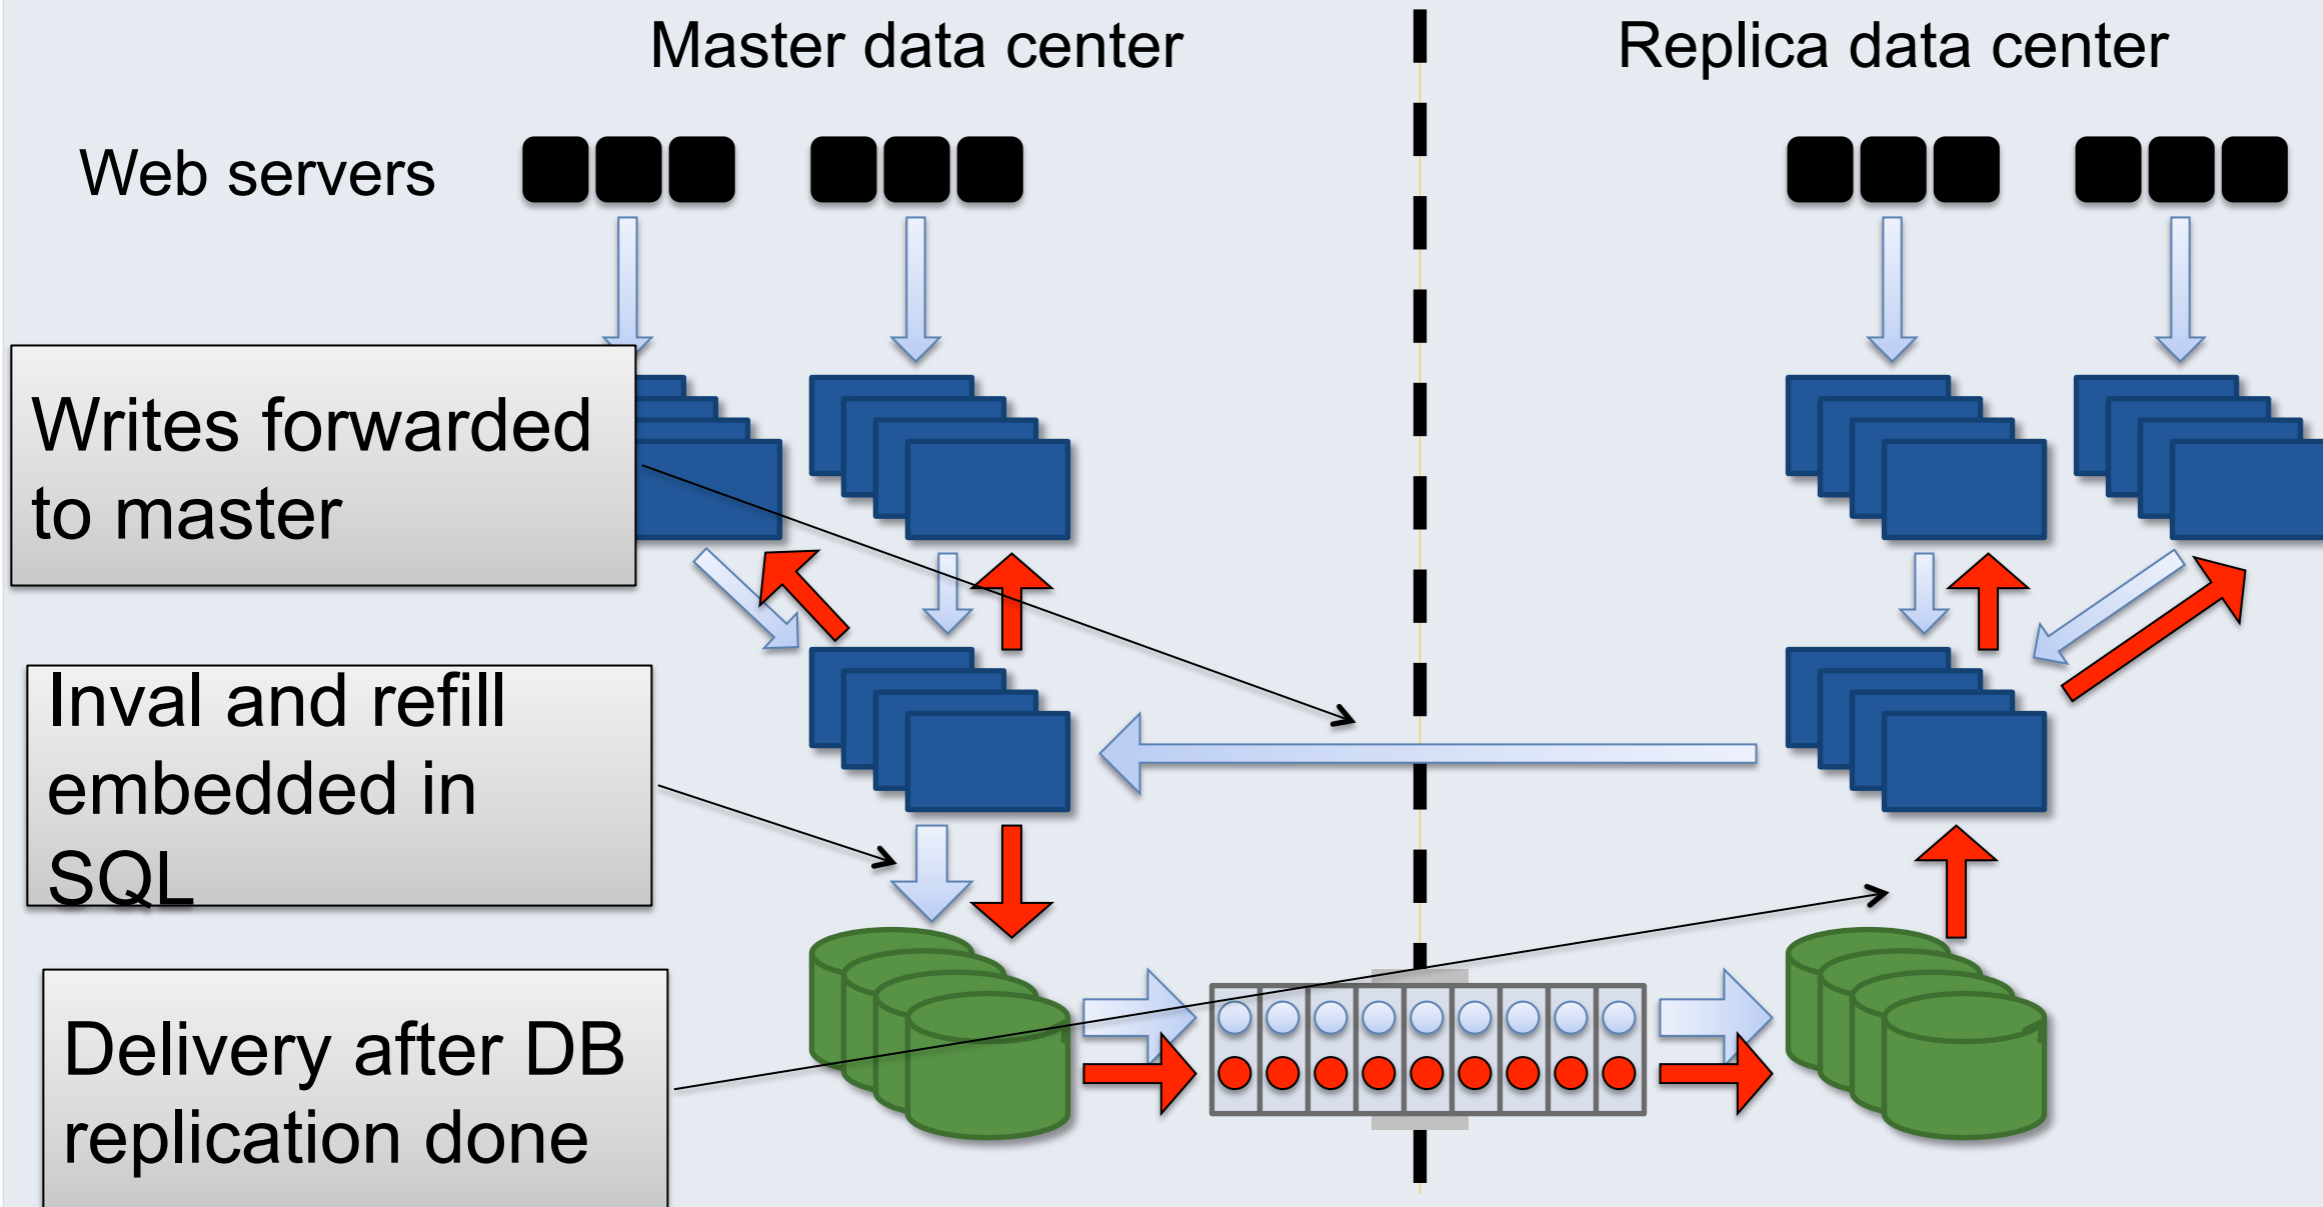


X -> Y ok

X -> Y ok

X -> Y ok

Cache consistency across data centers



Detecting inconsistencies

- Bad:
 - Users can't see comments on Zuck's post
- Better:
 - Measuring consistency while not impacting consistency of the system
 - Detecting data that is inconsistent
 - helps finding the root cause

Fixing inconsistencies

- Fixing the root cause is usually enough
- Invalidation messages will be delivered when the system is healthy
 - Sometimes cleanup is necessary
 - Invalidating shards manually
- “Permanent” inconsistency
- Bugs in invalidation pipeline
 - Cache gets evicted after ttl expires or is replaced by newer value
 - Bad for things that don't change often

Summary



Consistency

is **key**

for keeping
your dog

healthy and happy!

facebook

Q&A