Relevance & Personalization with Real Time Big Data Analytics

GROUPON®

Ameya Kanitkar
ameya@groupon.com

GROUPON

About Me

- Lead Engineer working on Real Time Data Infrastructure @ Groupon
- Graduate of Carnegie Mellon University & Pune University

What are Groupon Deals?

Our Relevance Scenario

Users

Deal Personalization Infrastructure Use Cases

Deliver Relevant Experience with High Quality Deals

Deliver Personalized Website, Mobile and Email Experience

Deal Level Performance User Behavior Performance

Deal Personalization Infrastructure Use Cases

Deliver Personalized Emails

Personalize billions of emails for hundreds of millions of users

Offline System

Deliver Personalized Website & Mobile Experience

Personalize one of the most popular e-commerce mobile & web app for hundreds of millions of users & page views

Online System

Earlier System

Email

Offline Personalization Map/Reduce

Online Deal Personalization API

MySQL Store

Data Pipeline (User Logs, Email Records, User History etc)

Scaling: Keeping Up With a Changing Business

Growing Deals

2011

20+

2012

400+

2015

2000+

Growing Users

- 110 Million+ subscribers
- We need to store data like, user click history, email records, service logs etc. This tunes to billions of data points and TB's of data

Changing Business: Shift from Email to Mobile

110 Million+ App Downloads

- Growth in Mobile Business
- Reducing dependence on email marketing
- Change in strategy from daily deal model to deal marketplace

Earlier System

- Scaling MySQL for data such as user click history, email records was painful unless we shard data
- Data Pipeline is not “Real Time”

Ideal System

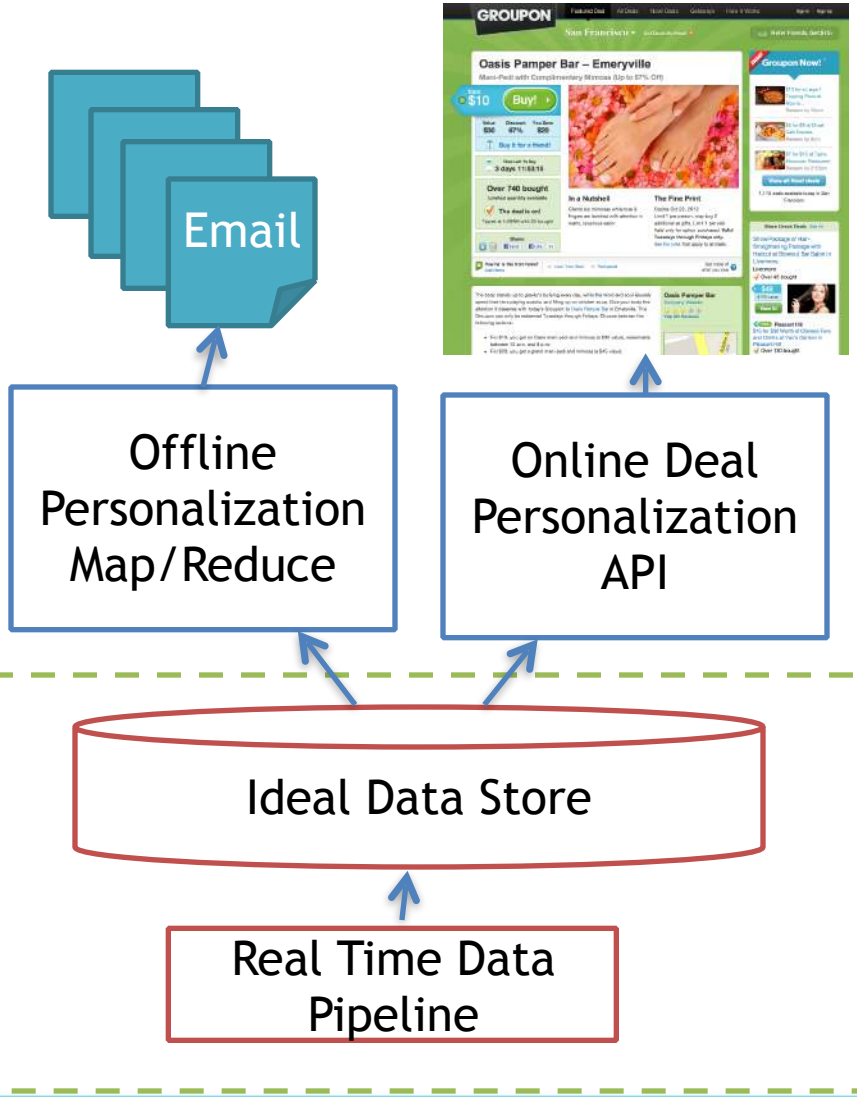

- Common data store that serves data to both online and offline systems
- Data store that scales to hundreds of millions of records
- Data store that plays well with our existing hadoop based systems
- Real Time pipeline that scales and can process about 100,000 messages/ second

Final Design

Email

Web Site Logs

Mobile Logs

Online Deal Personalization API

Offline Personalization Map/Reduce

Kafka Message Broker

Storm

HBase

GROUPON

Two Challenges With HBase

How to scale
100,000
writes/ second?

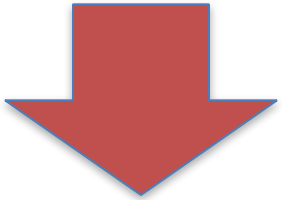

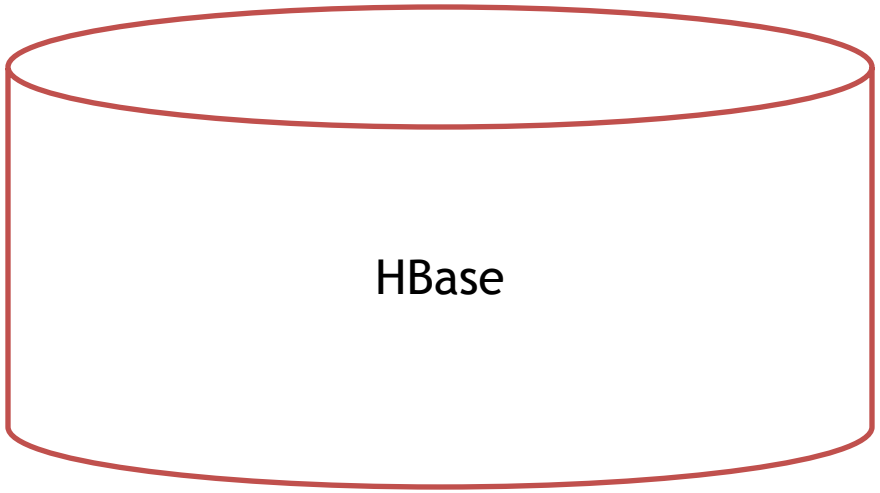

- How to run Map Reduce Programs over HBase without affecting read latency
- How to batch load data in HBase without affecting read latencies

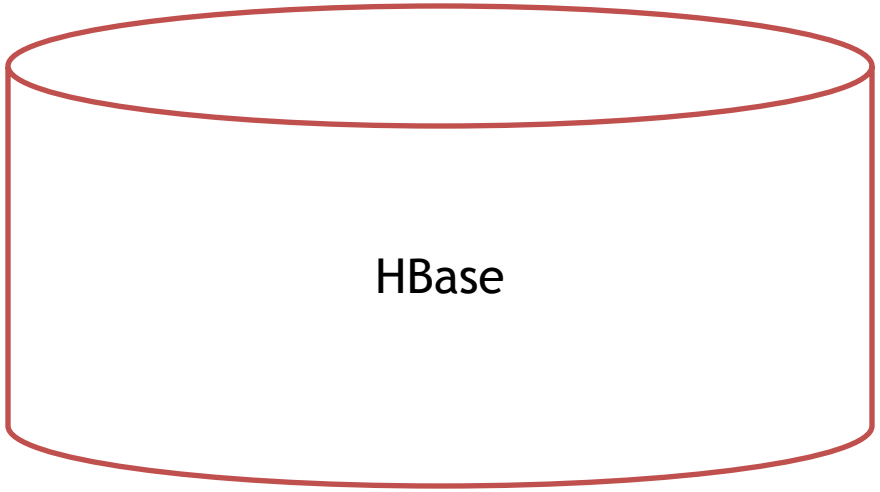

GROUPON

Scaling with HBase

- **Every new write is written as a new column in HBase. This also ensures de-duplication**
- **At the time of writing new event, no read happens**

	Column Family 0
User_id 1	***timestamp_DealId_eventType: {Impression1}***
	timestamp_DealId_eventType: {Impression2}

Final Design

Map Reduce Over HBase

Real Time HBase

Replication

Batch HBase

Bulk Load data via HFiles

Leveraging System for Real Time Analytics

Various requirements from relevance algorithms to pre-compute real time analytics for better targeting

How do women in Berlin convert for Pizza deals?

How women in Berlin are converting for a particular pizza deal?

GROUPON

Leveraging System for Real Time Analytics

More Complex Examples

How do women in Berlin from Mitte area aged 45-50 convert for New York Style Pizza, when deal is located within 2 miles, and when deal is priced between €10-€20?

How do women in Berlin from Mitte area aged 45-50 convert for New York Style Pizza, for this particular deal which happens to be about 2 miles away, and deal is priced for €12

GROUPON

Leveraging System for Real Time Analytics

Even More Complex Examples

How do women in Berlin from Mitte area aged 45-50 convert for New York Style Pizza, when deal is located within 2 miles, and when deal is priced between €10-€20 who also like Activities such as Biking and who have been very active customer of Groupon deals on mobile platform?

How do women in Berlin from Mitte area aged 45-50 convert for New York Style Pizza, for this particular deal which happens to be about 2 miles away, and deal is priced for €12, who also like activities such as biking and who have been very active customer of Groupon deals on mobile platform?

GROUPON

Power of Simple Counting

Turns out all earlier questions can be answered if we could count appropriate events in appropriate bucket

$$\text{Conversion rate for pizza deals for women in Palo Alto} = \frac{\text{No of Purchases by Women in Palo Alto for Pizza Deals}}{\text{No Deal Impressions by Women in Palo Alto for Pizza Deals}}$$

GROUPON

Real Time Analytics Infrastructure

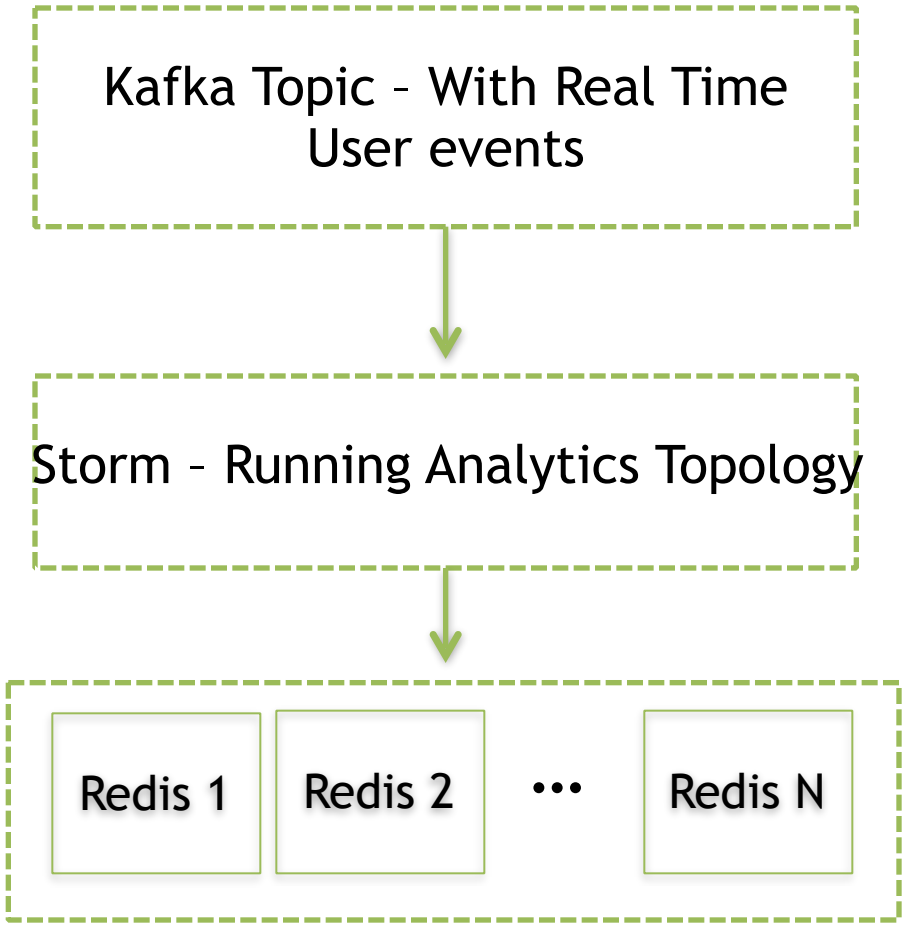

Real Time infrastructure processing 100,000 requests/ second

Storm Topology calculating various dimensions/ buckets and updates appropriate Redis bucket. Redis is sharded from client side

Redis cluster handles over 3 Million events per second. Stores over 14 Billion unique keys

GROUPON

Real Time Analytics Infrastructure - Explained

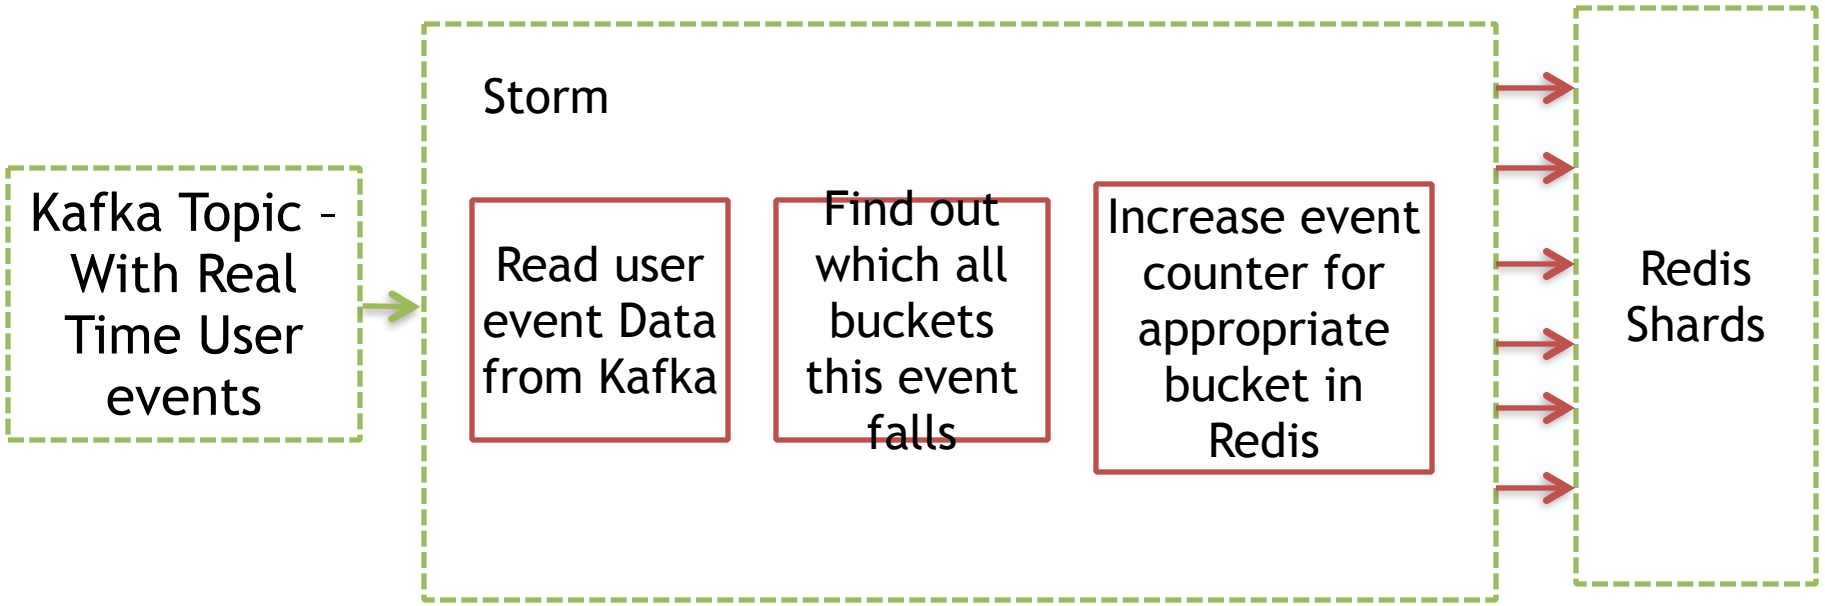

Scaling Challenges - Kafka - Storm

- Massive reads and writes to Kafka cluster created back pressure in Storm topologies that produce data into Kafka
- Storm was hard to scale. We had to try various number of combinations to finalize how many bolts of each type are required for steady state operations and overall how many workers are needed etc. This took more time than we anticipated
- Use “maxSpoutPending” setting in Storm topologies. We found it to be very useful to shield your topologies from sudden increase in traffic
- Build your entire infrastructure – where data duplicates are allowed

GROUPON

Scaling Challenges - Redis

- Reduce memory footprint – use hashes. Very memory efficient compared to normal Redis keys
- In order to support high write operations turned off AOF, turned on RDB backups

Easiest of all other infrastructure pieces – Kafka, Storm, HBase

When Small is Big – Bloom Filters

- Since both Kafka and Storm can send same data twice specially at scale, it was important to build downstream infrastructure that can handle duplicate data.
- However, by very nature Analytics Topology (Counting Topology) cannot handle duplicates
- Storing individual messages for billions of messages is way too expensive and would take lot more memory
- So we used bloom filters. At a very small % error rate, we could effectively de dup data with a very small memory footprint. We store bloom filters in Redis using redis “bit” support

GROUPON

Avoiding Errors – Backups/ Recovery Strategy

With such a high volume system, which also drives so much revenue for the company good backup/ recovery strategy is necessary

Redis

RDB Backups every X hours. RDB backups are stored in HDFS for later use

HBase

HBase Snapshot functionality is used. Snapshot taken every X hours. Can be loaded as necessary

Kafka/ Storm

All input into Kafka topic is stored in HDFS. So any hour/ day can be replayed from HDFS if necessary

GROUPON

We Are Hiring…

GROUPON

Questions?


Questions?
GROUPON
Thanks!
ameya@groupon.com
www.groupon.com/techjobs