



01/16/2014

**CASSANDRA AND
RIAK AT
BESTBUY.COM**

WHO WE ARE

- Best Buy is the **world's largest multi-channel consumer electronics retailer**, with stores in the United States, Canada, China and Mexico.
- We are the 10th **largest** online retailer in the United States
- More than **1.6 billion** visitors come to our stores and BestBuy.com each year.
- Our **My Best Buy** loyalty program is among the largest loyalty programs of its kind, **with more than 40 million active members.**
- We provide customers with **outstanding choice, unbiased advice and unmatched support** for the tech needs.



A UNIQUE CUSTOMER PROMISE

- THE **LATEST DEVICES AND SERVICES**, ALL IN ONE PLACE
- **IMPARTIAL & KNOWLEDGEABLE** ADVICE
- **COMPETITIVE PRICES**
- THE ABILITY TO SHOP **WHEN AND WHERE YOU WANT**
- **SUPPORT** FOR THE LIFE OF YOUR PRODUCTS



**JOEL
CRABB**

**Chief Architect,
BestBuy.com**

**KANNAN
SWAMINATHAN**

**Director, Web
Architecture**



NOSQL AT BESTBUY.COM

- **Schema-less design**
 - **Key-value systems**
 - **Sparse column systems**
 - **Non-relational data**
- **Distributed nature**
 - **Eventual consistency**
 - **Active-Active across clouds and datacenters**
 - **High reliability**
 - **Horizontal scaling**



BESTBUY.COM



Search by Keyword, SKU # or Item #



WEEKLY DEALS

Stores

Orders

Help

Credit Cards

Español

PRODUCTS

SERVICES

SHOPS & DEALS

GIFTS

my 

Account

Sign In

Hi, Joel
Sign Out



INTRODUCING

my 

REINVENTING REWARD ZONE®

MORE REWARDS + MORE EXCLUSIVES + MORE MOBILE

[Learn More](#)



HTC One Blue
Only \$49.99
with activation
and 2-year
agreement.
Plus free shipping.
[Shop now >](#)



**Grand Theft
Auto V Midnight
Opening**
Monday night at
select stores.
[Learn more >](#)



LOW PRICE GUARANTEE

We'll match prices on qualifying products. [See details >](#)



FREE SHIPPING

on orders \$25 and up.
[See details >](#) Exclusions apply.



18-MONTH FINANCING

on purchases totaling \$429 or more. [Learn more >](#)




January 16, 2014

©Copyright 2014 - Best Buy, Inc. All rights reserved.

SCALABILITY GOALS

- **Near-Infinite**
- **Bursts**

- **7X traffic spikes**
- **Bursts > 50,000 rps**
- **#4 in eCommerce traffic during 2013 Holiday**



— bestbuy.com | — walmart.com

(from Jan 1, 2013 to Dec 30, 2013) (in millions of visits per day)

(7-day moving averages)



January 16, 2014

©Copyright 2014 - Best Buy, Inc. All rights reserved.

LOW PRICE GUARANTEE


ON SALE

\$599.99

Regular Price: \$799.99

You Save: \$200.00

FREE SHIPPING

 **Add to Cart**

FLEXIBILITY

GOALS

Low cost of change

Fast concepts to site

Daily releases

Multiple versions

One day of work vs. two months

RELIABILITY GOALS

- **100% availability**
- **Zero defects**
- **~ 2s response times**

- **Achieved 100% cloud uptime during Holiday**

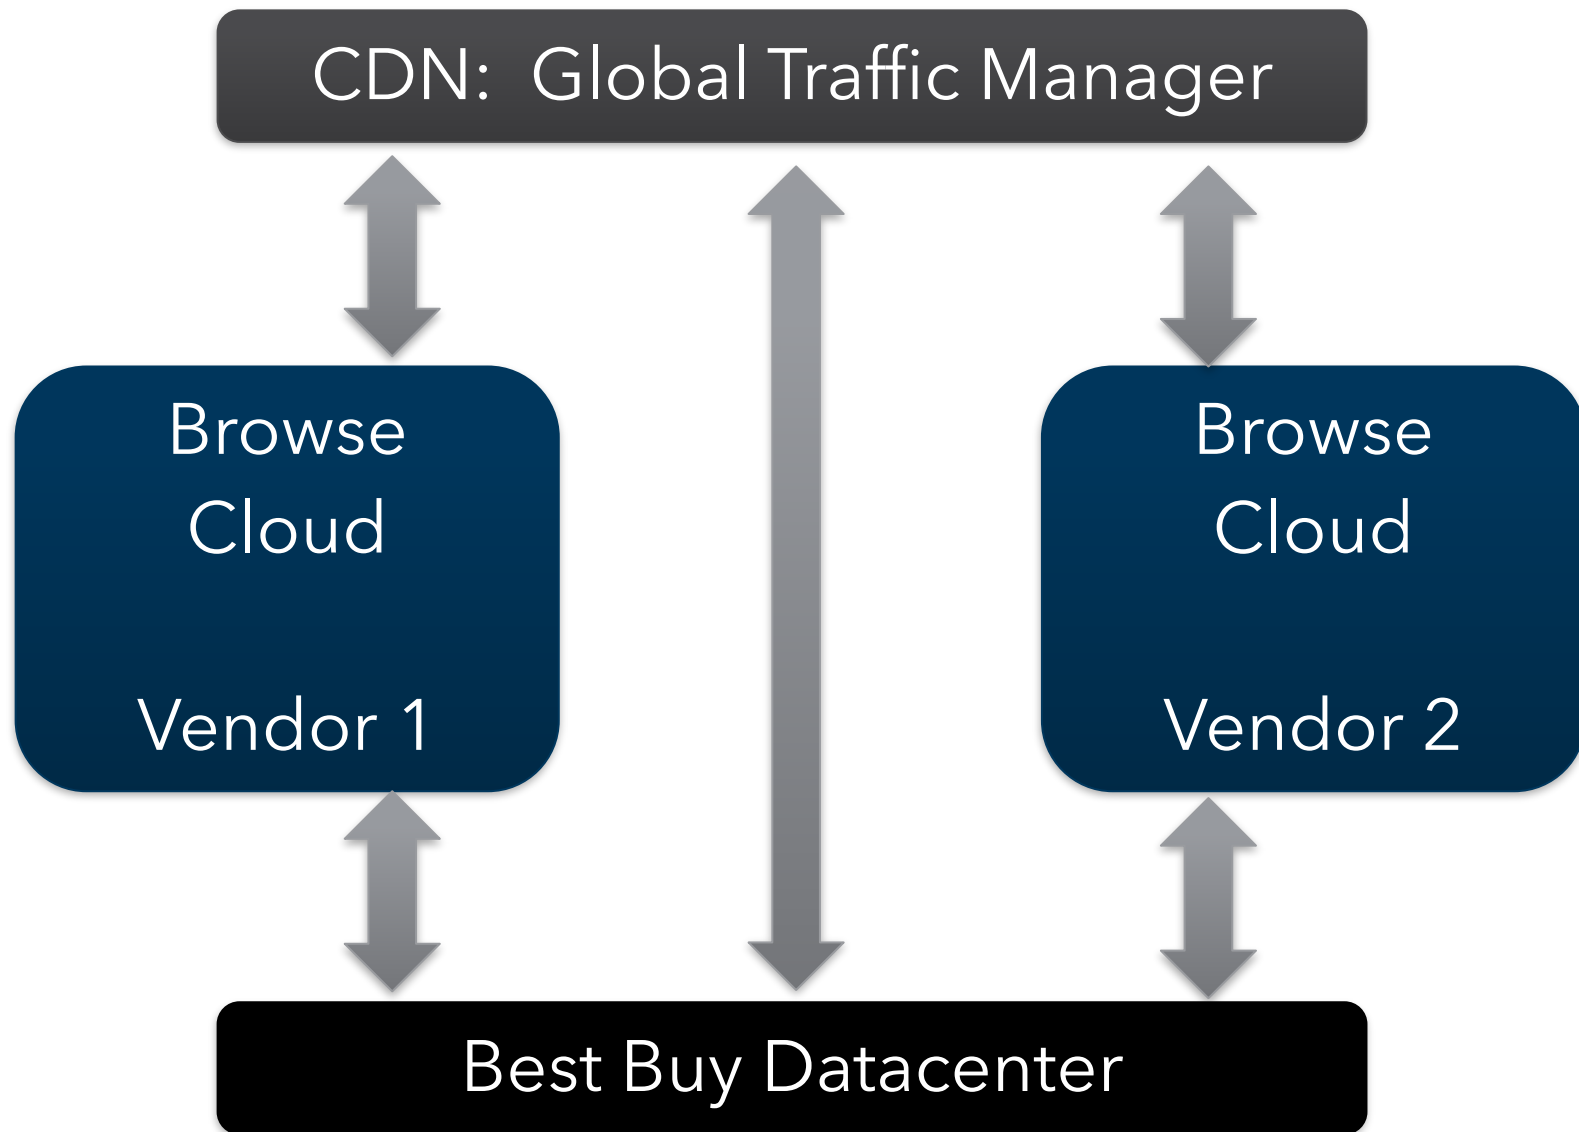


CLOUD ARCHITECTURE CONCEPTS

- **Clouds fail; plan for it**
 - **Multiple availability zones**
 - **Multiple regions**
 - **Multiple vendors**
- **Datacenter connections fail; plan for it**
 - **Serve pages completely from cloud**
 - **Browse-only fallback mode**



CLOUD ARCHITECTURE CONCEPTS





RIAK

PRODUCT CATALOG - REQUIREMENTS

- **Business Requirements**

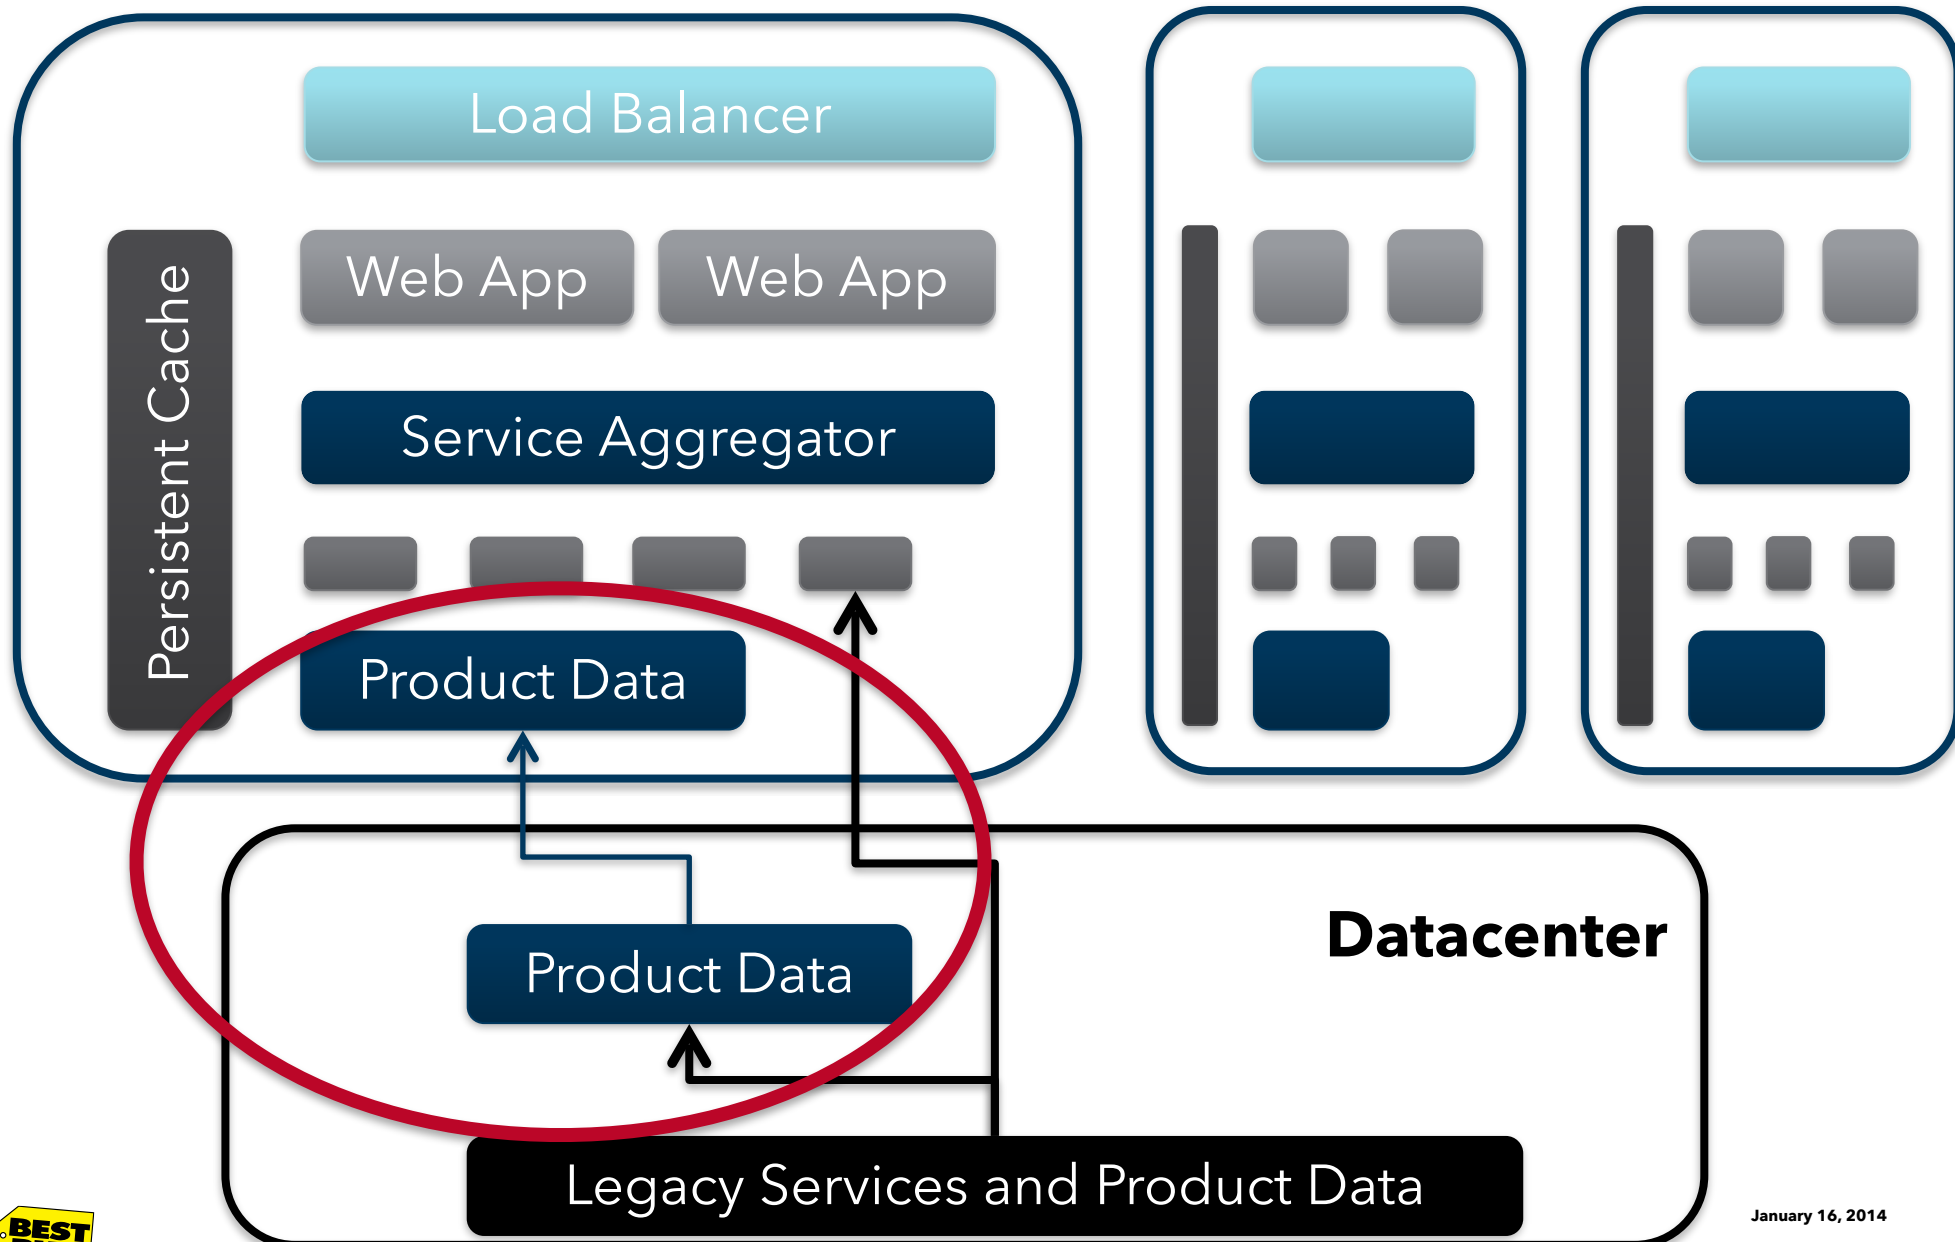
- **Easily add new attributes to products**
- **Use existing product content feeding systems**
- **Provide enhanced product data to all of Best Buy**

- **Technical Requirements**

- **Scale to BestBuy.com's needs**
- **One-way replication from DC to cloud**
- **Provide a Product Catalog API usable by Best Buy applications internally (DC) and externally (cloud)**



BROWSE CLOUD - ARCHITECTURE



WHY RIAK FOR THE PRODUCT CATALOG

- **Key-value Store**
 - Easily add new attributes
 - Different attribute sets for different product categories
- **Hub and Spoke Replication**
 - Multiple cloud instances can be out-of-synch for seconds or minutes; this is OK
 - Connection to Datacenter can be lost
- **Resilience within all tiers**
 - Riak's ring architecture allows single instances to fail with little impact to the system



WHY RIAK FOR THE PRODUCT CATALOG

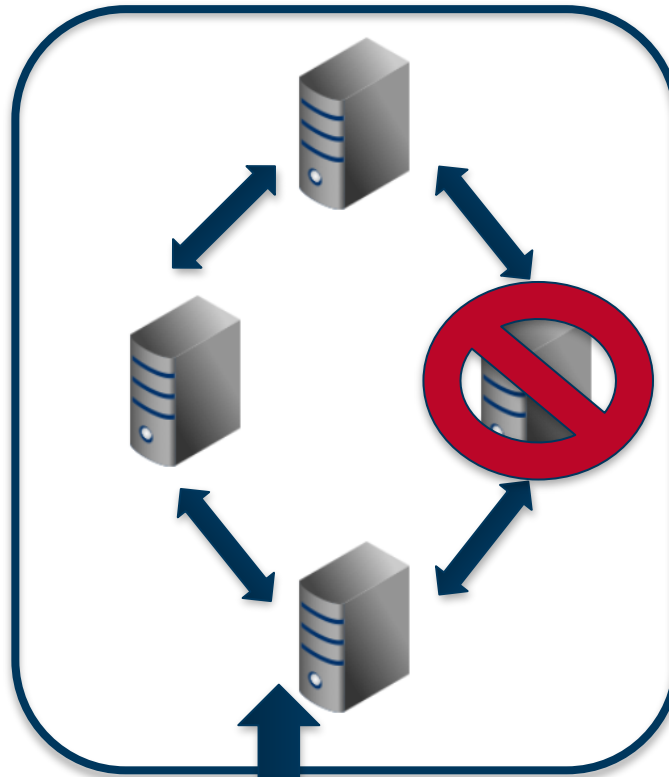

- **Legacy Bridge**
 - **Extract data from legacy system**
 - **Expose to anyone who needs the data**
 - **Extends decision to retire legacy system**
- **Input into Product Evolution**
 - **Best Buy is stretching Riak in multiple areas**
 - **Feedback on Search and Replication**
 - **Direct connection to Riak engineers**




PRODUCT CATALOG DEPLOYMENT

Cloud Replicant

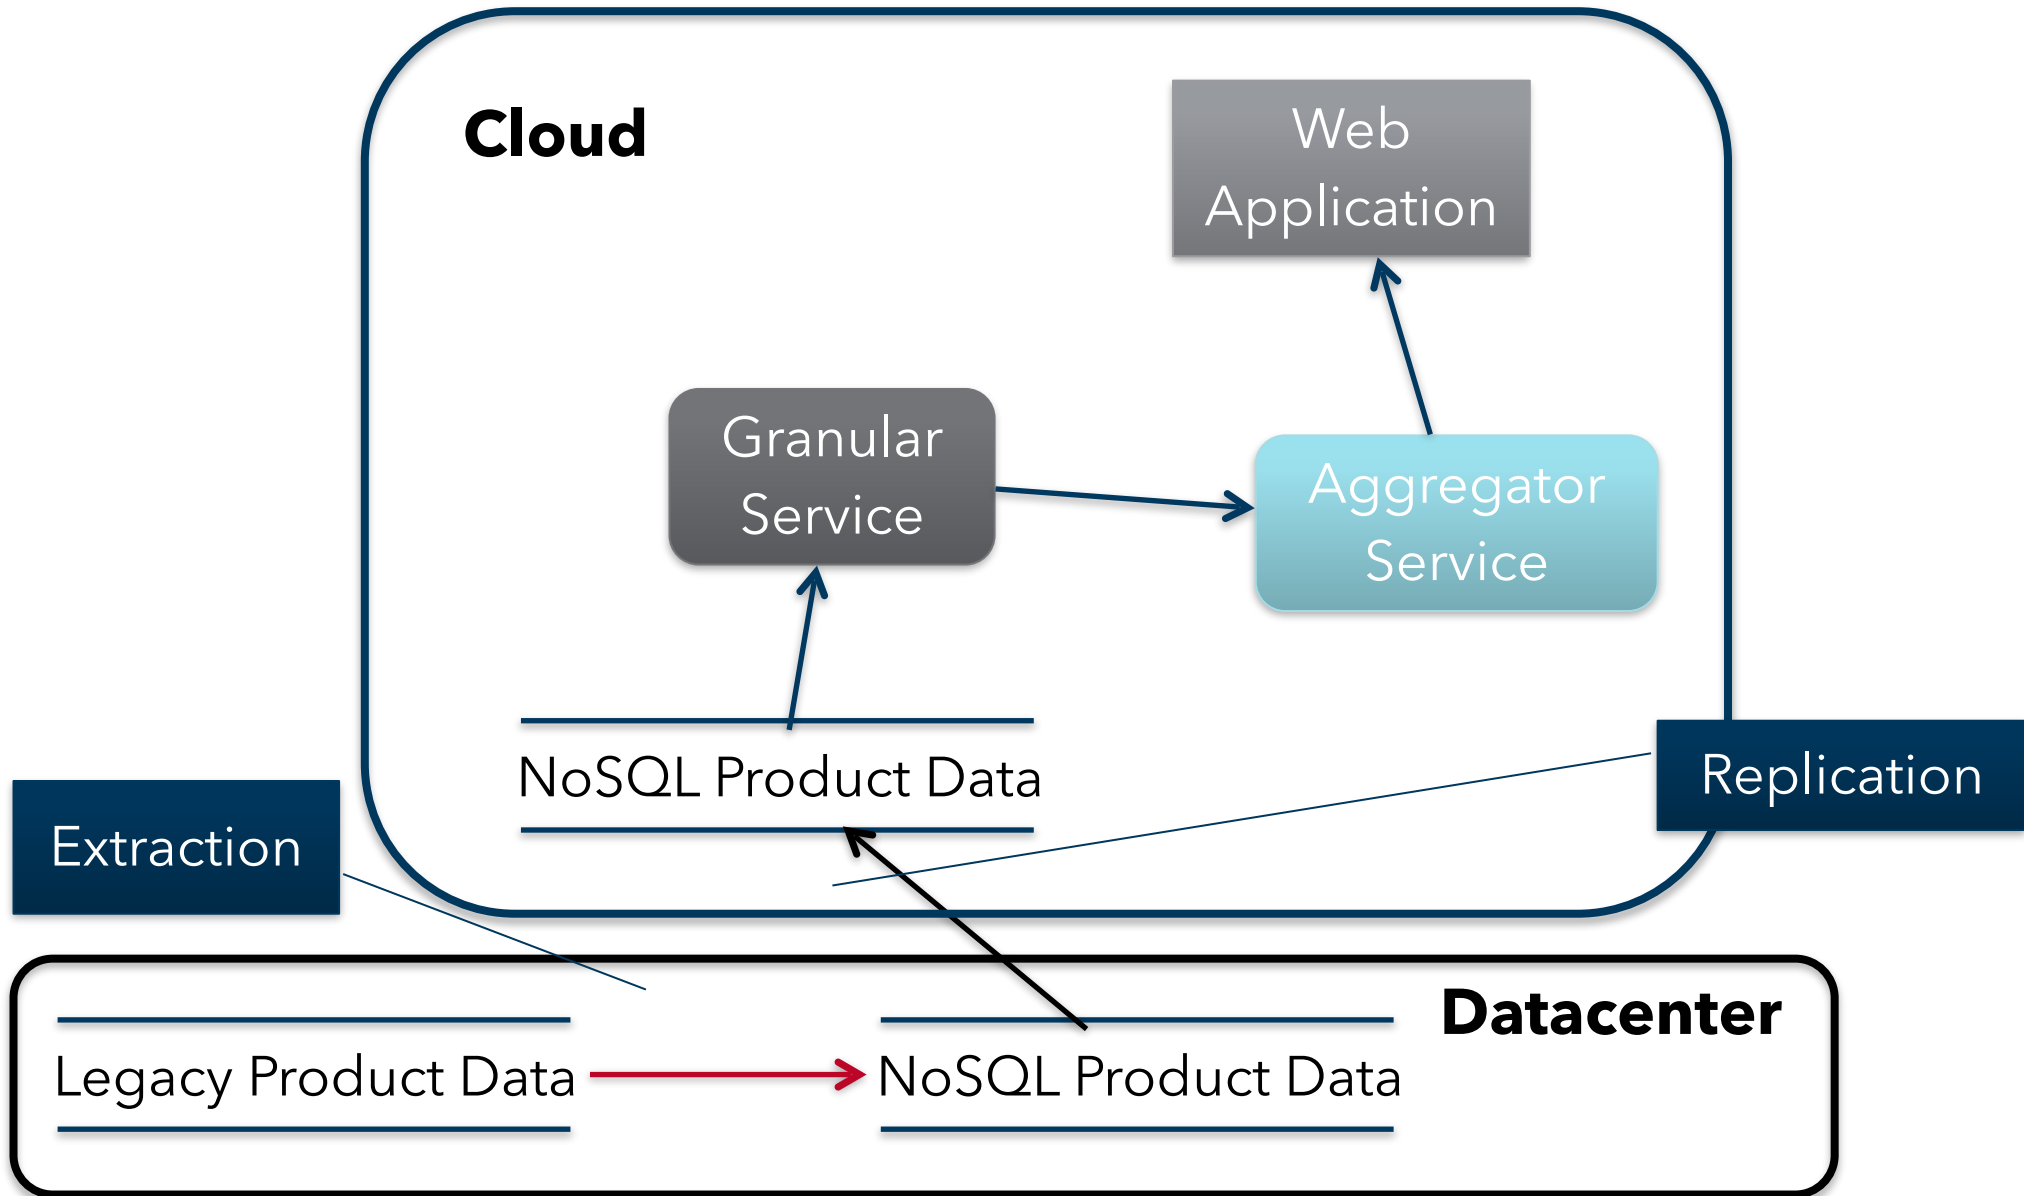
Cloud Replicant



Each Cloud Replicant is a self-contained instance which can intermittently lose connectivity to the Datacenter

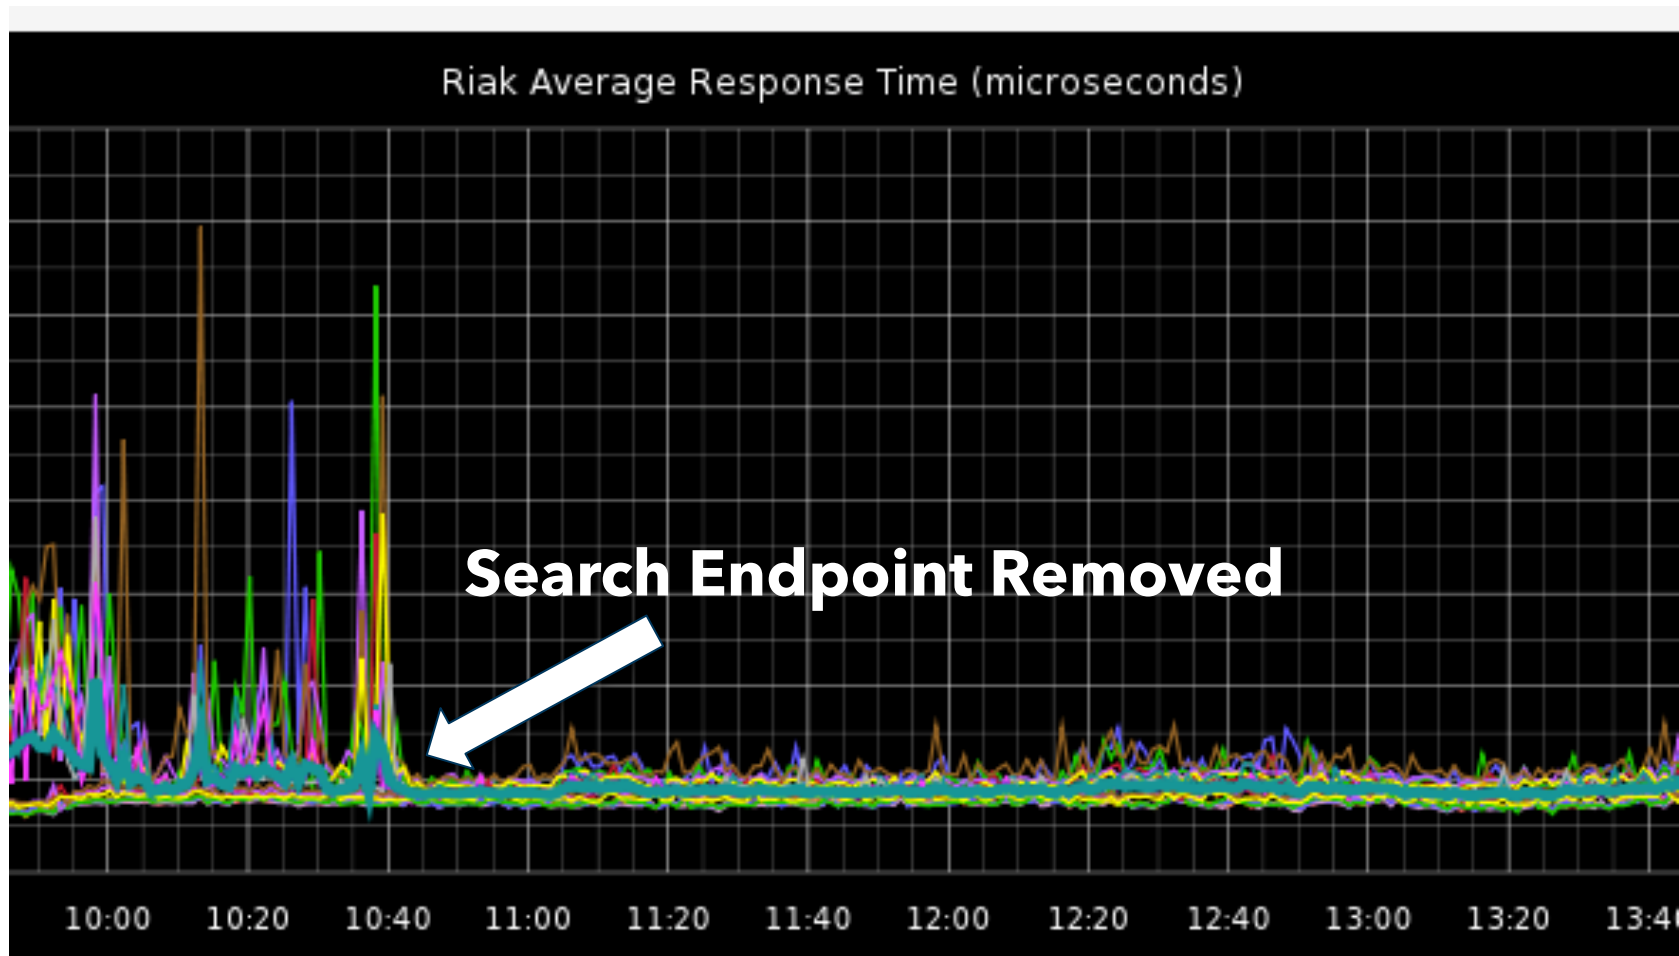


PRODUCT CATALOG - DATA FLOW



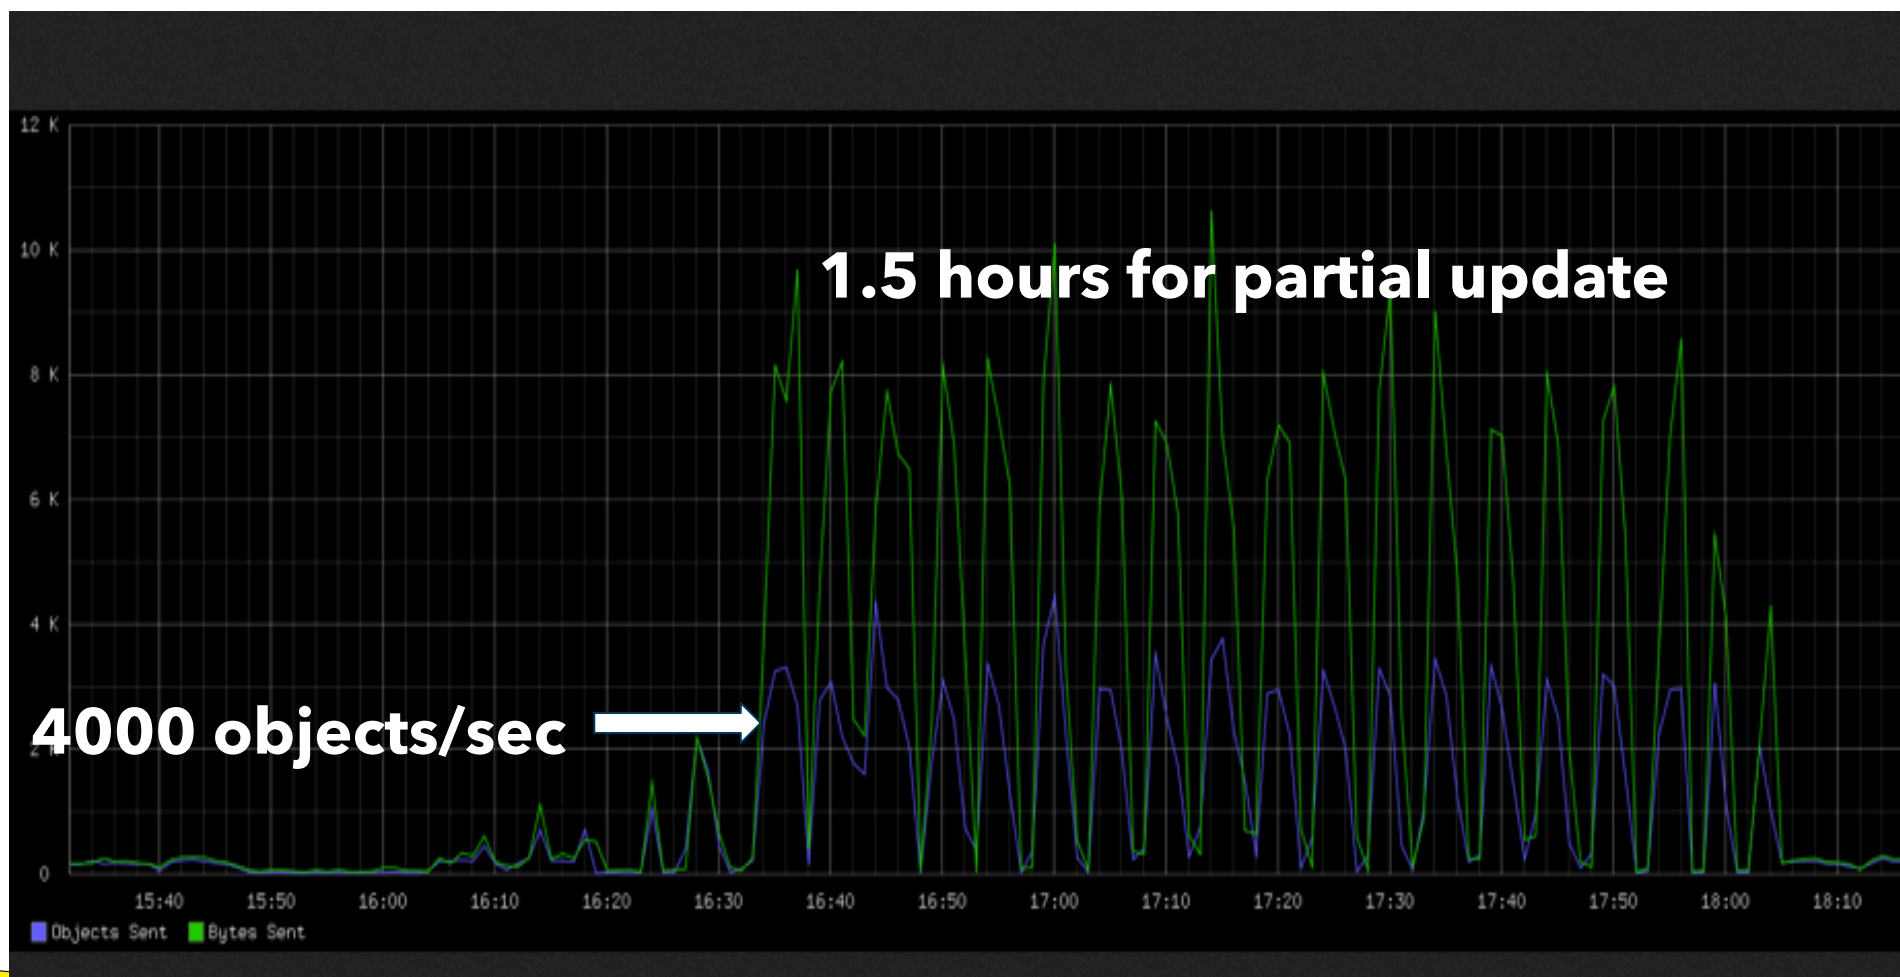
RIAK LESSONS LEARNED

- **Search eats up the CPU (Pre-Yokozuna)**



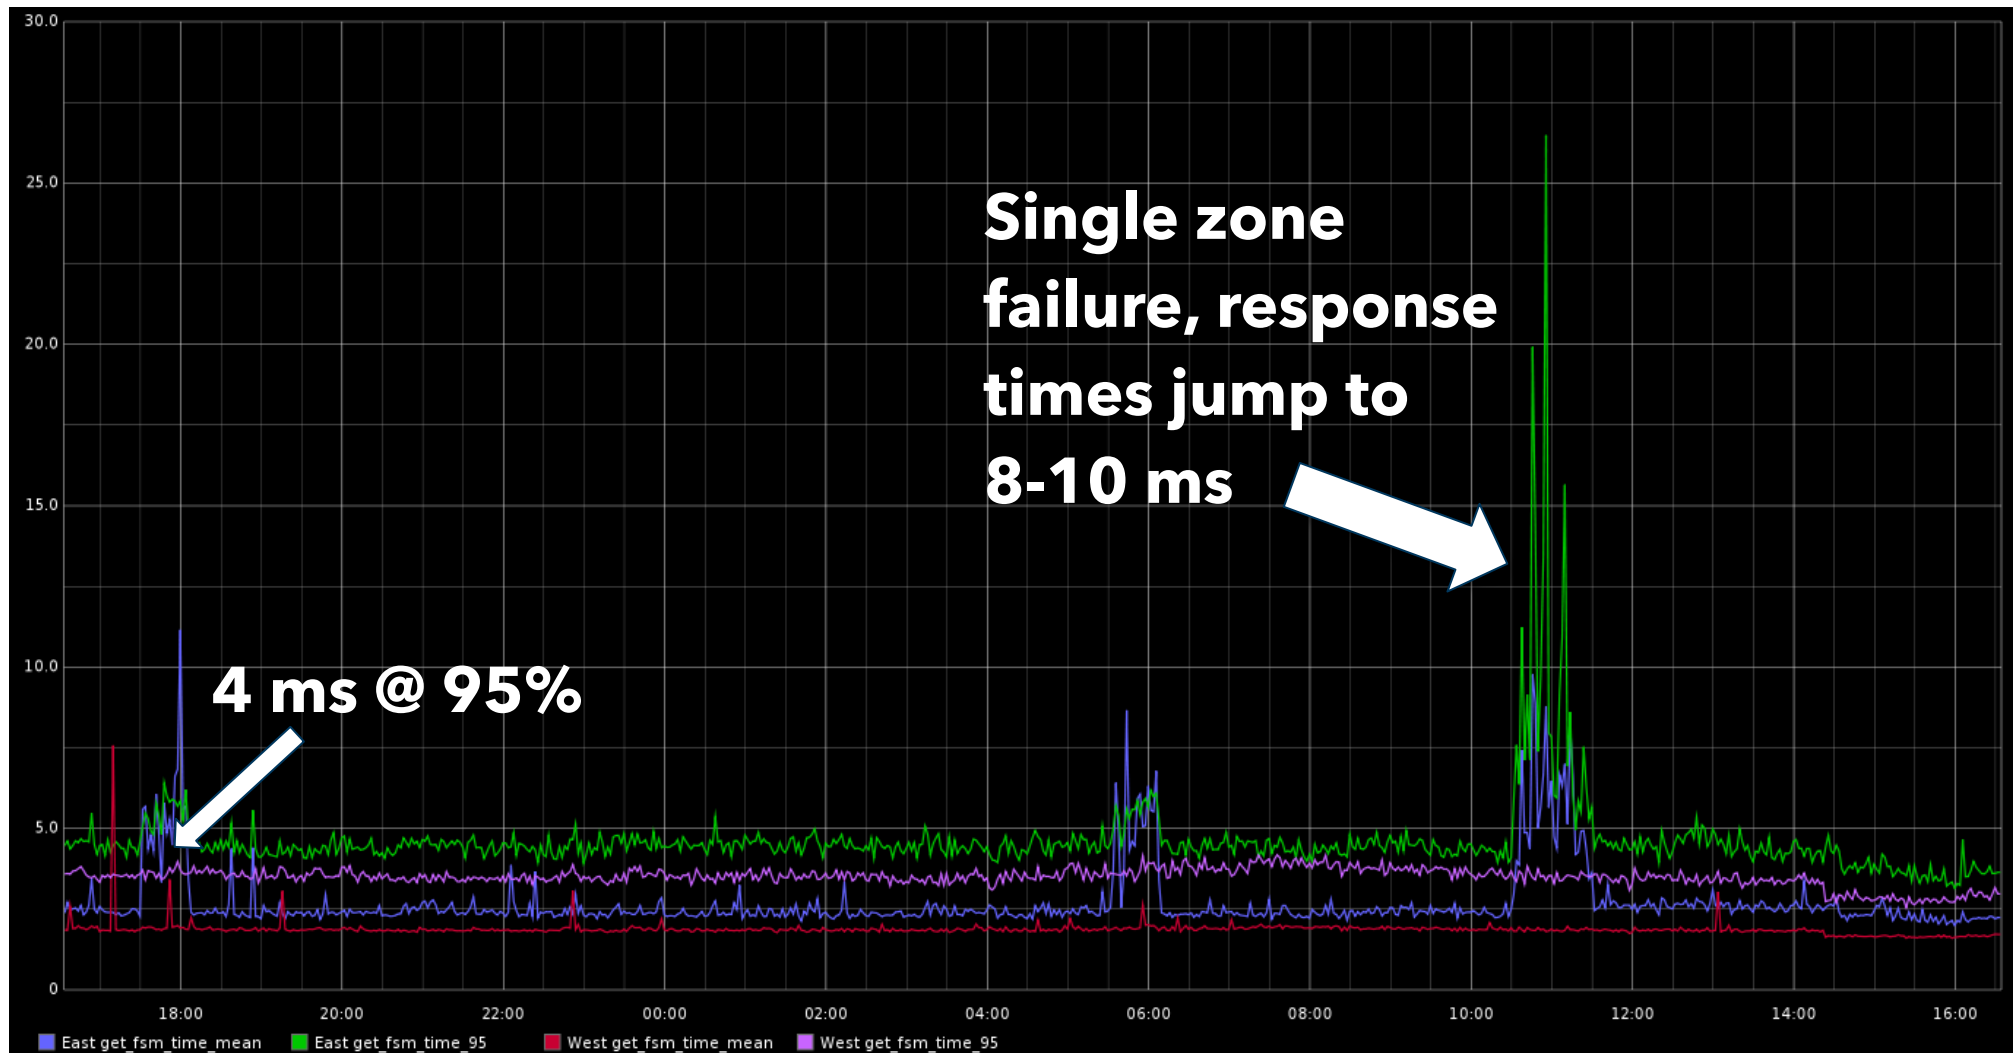
RIAK LESSONS LEARNED

- **Multi-datacenter replication is hard at scale**
- **Object replication fails silently (v1.4.2)**



DID THE PRODUCT CATALOG SUCCEED?

- **Scale** - generally < 5 ms response times in cloud
- **Scope** - provides APIs in cloud and DC





CASSANDRA



January 16, 2014

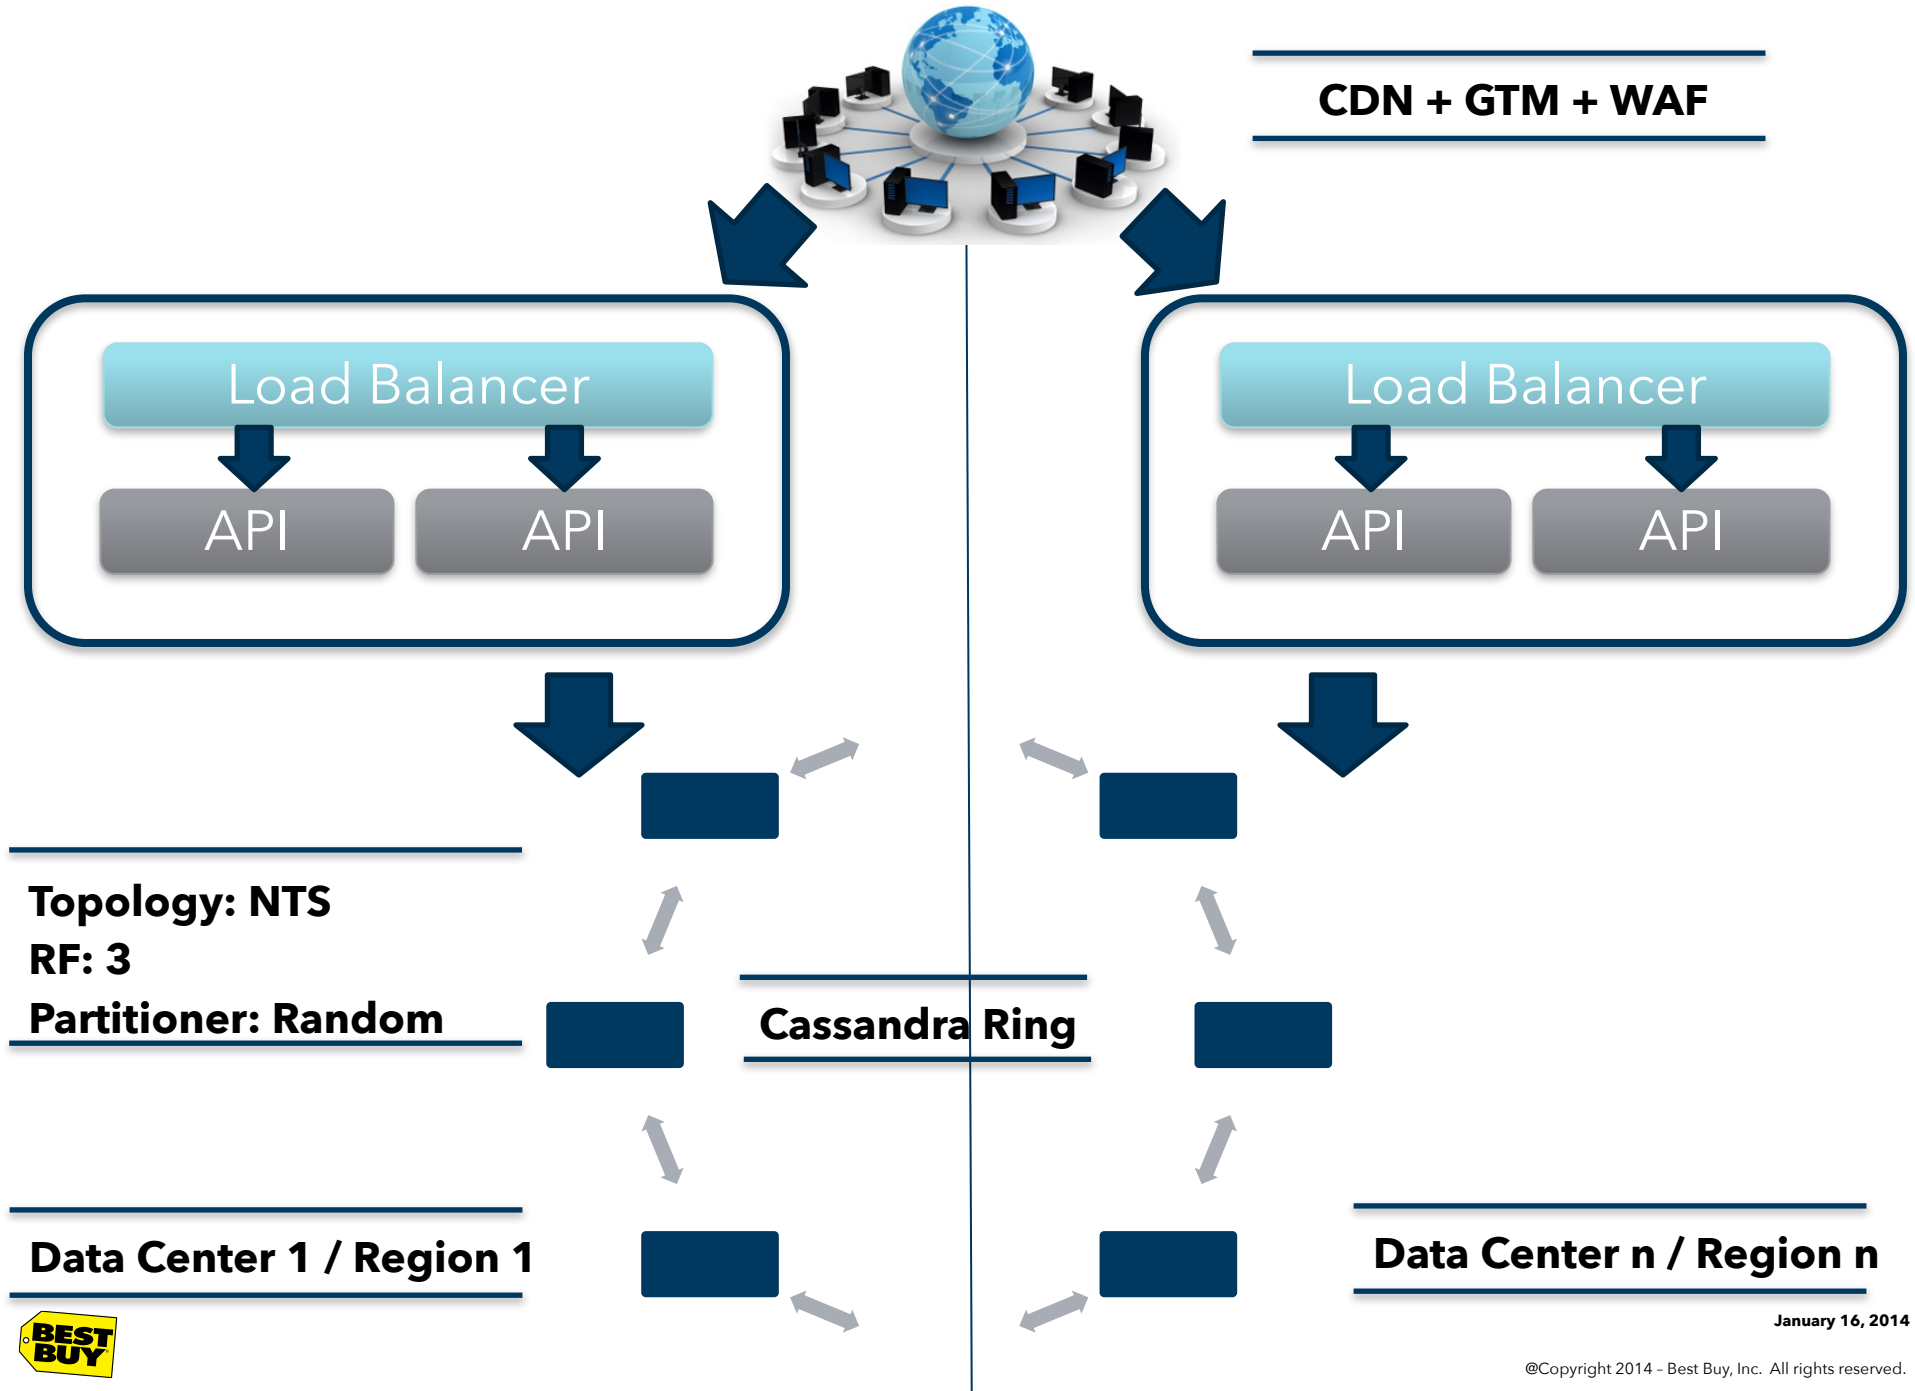
©Copyright 2014 - Best Buy, Inc. All rights reserved.

CUSTOMER GRAPH - REQUIREMENTS

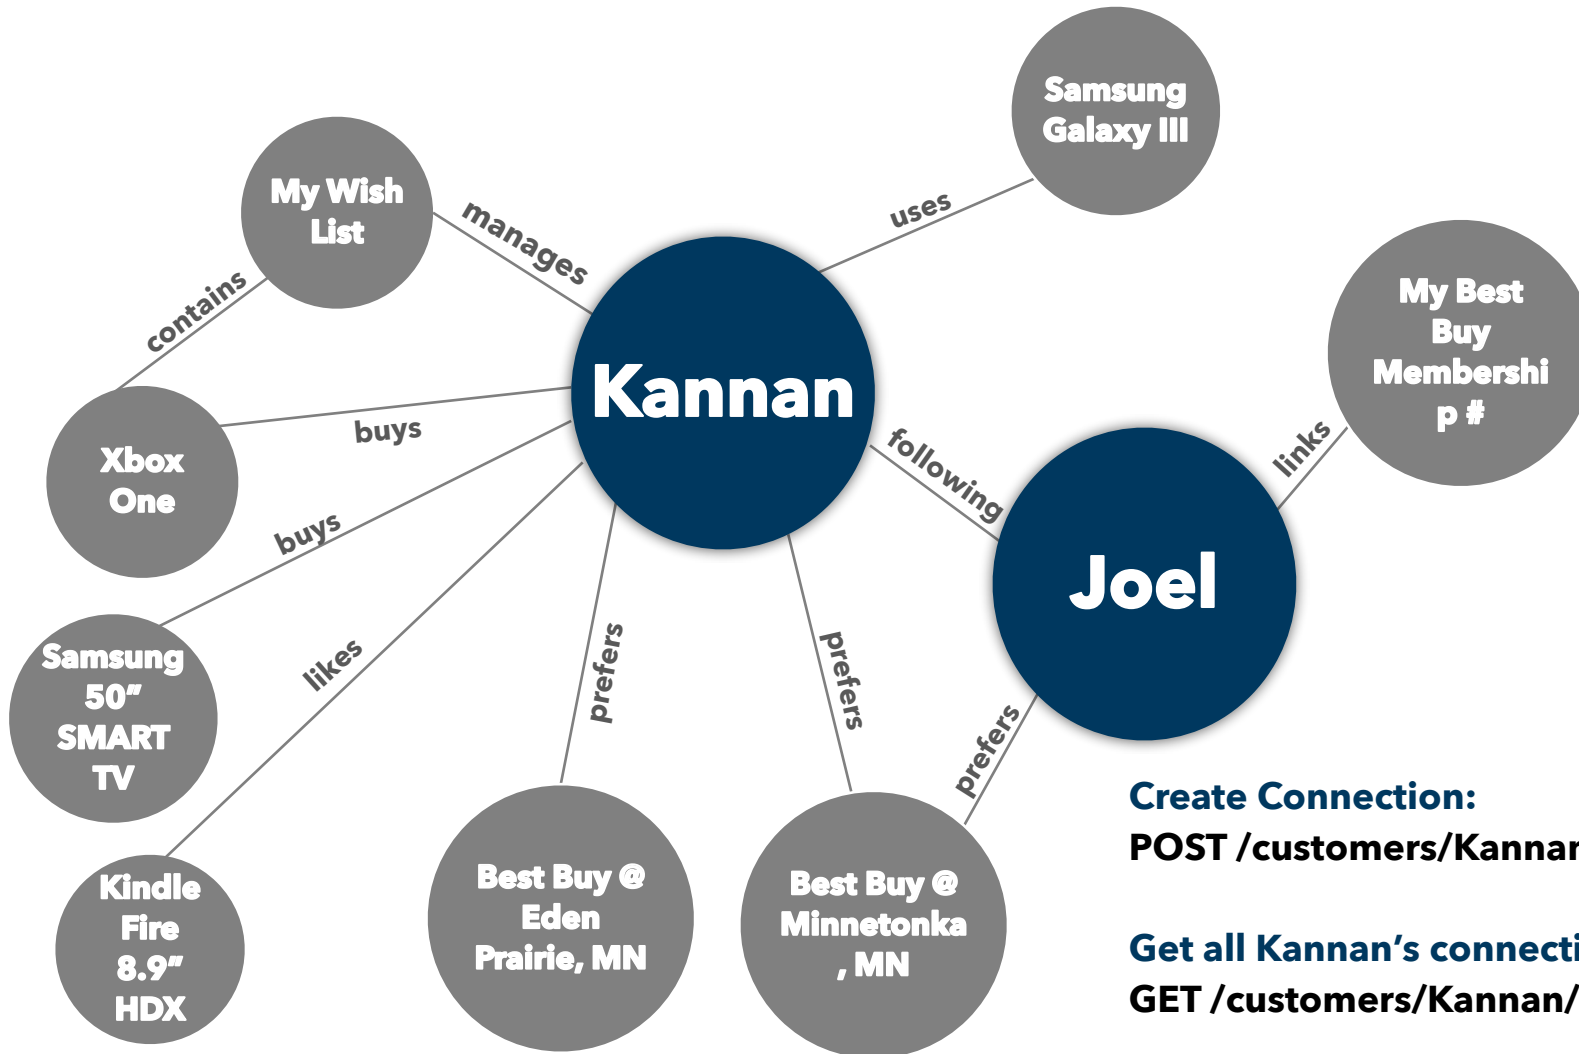
- **Business Requirements**
 - **Create a single identity provider**
 - **Build an adaptive model to support a 360 degree view of the customer, customer segmentation and multi-channel personalization**
 - **Extensible framework to support federated identity interoperability with external providers**
- **Technical Requirements**
 - **Scale to BestBuy.com's needs with a distributed service oriented architecture**
 - **Support risk based authentication and secure service interfaces for consumers**



CUSTOMER GRAPH - ARCHITECTURE



CUSTOMER GRAPH - CONCEPTUAL



Create Connection:

POST /customers/Kannan/following/customers/Joel

Get all Kannan's connections:

GET /customers/Kannan/connections

Get Joel's "followers":

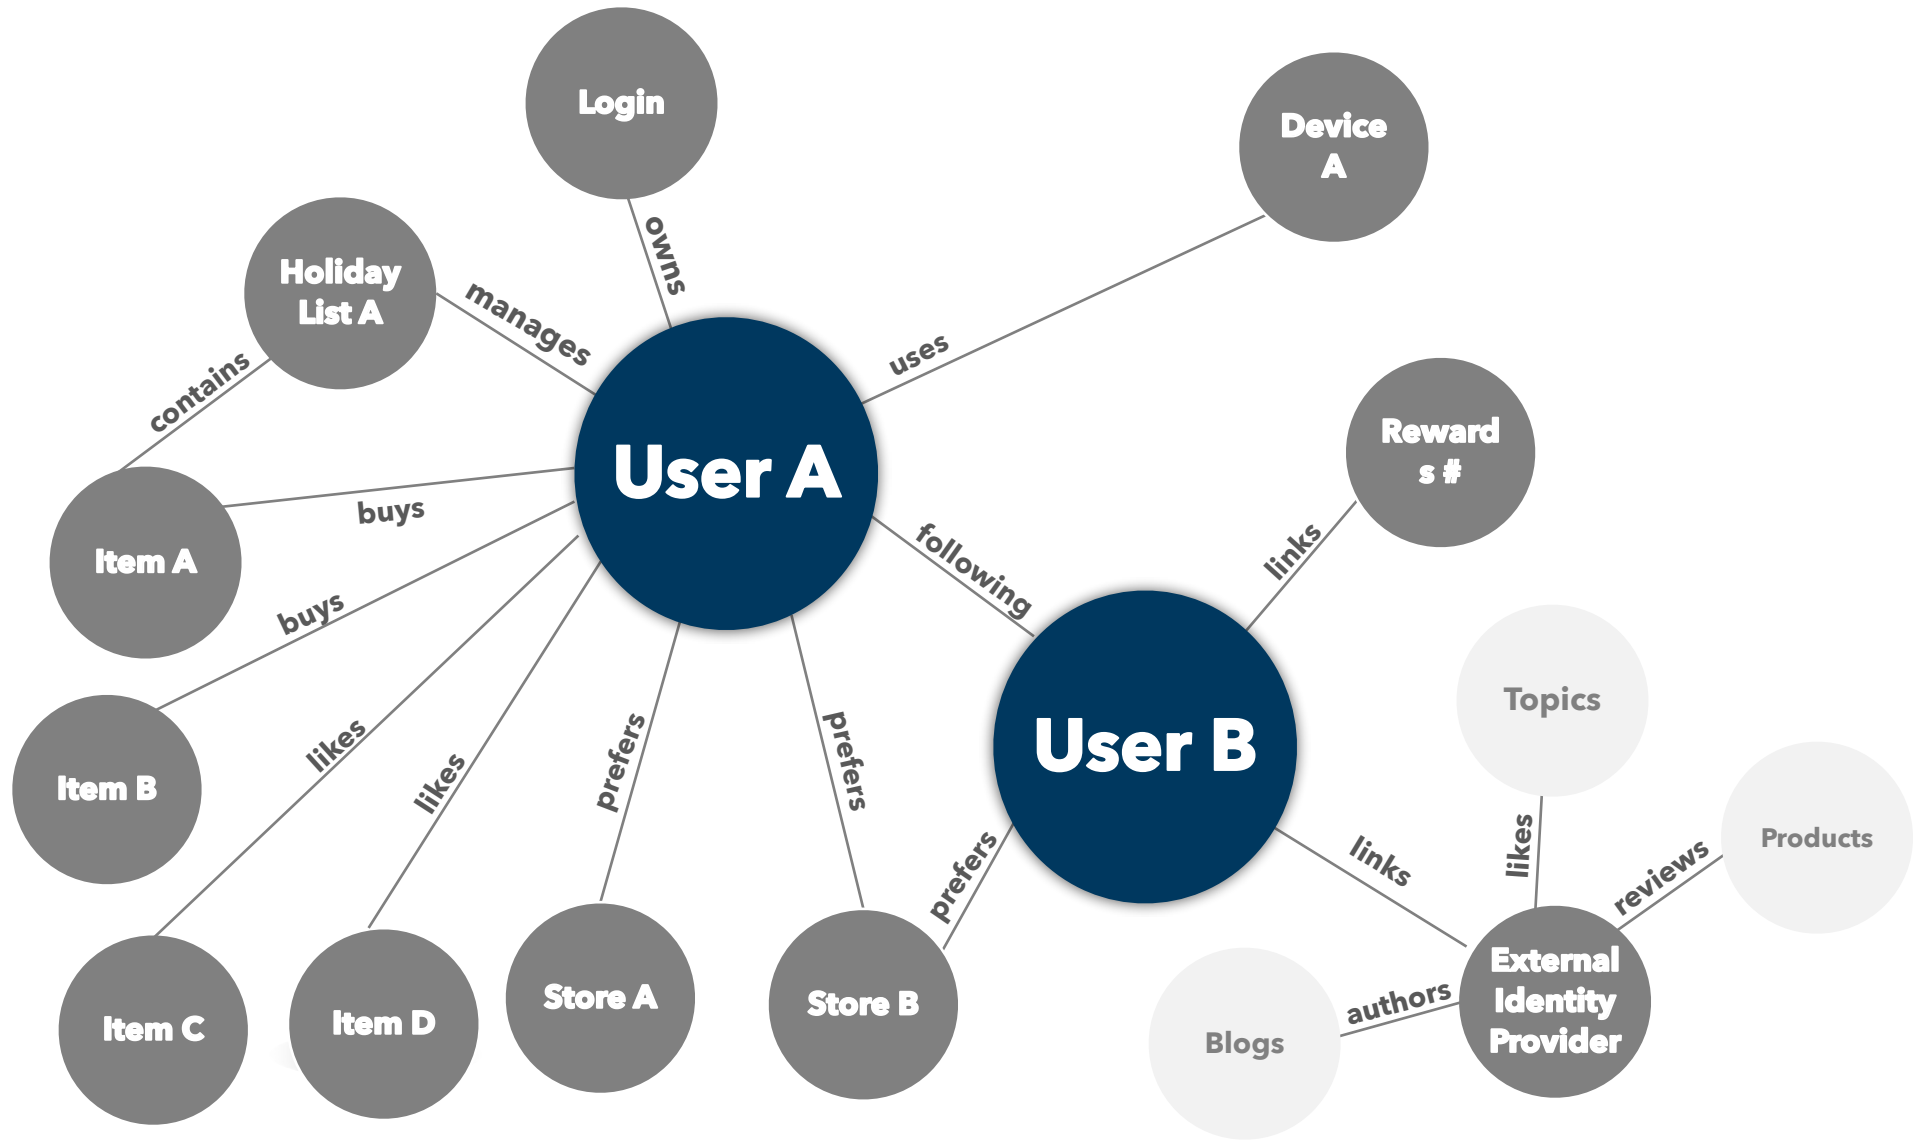
GET customers/Joel/followers

Get customers connecting to a specific store:

GET /stores/{store id}/connecting/customer



CUSTOMER GRAPH - CONCEPTUAL



WHY CASSANDRA?

- **No single point of failure**
- **Delivers on Atomicity, Isolation & Durability**
- **Eventual Consistency**
- **Tunable Consistency for Reads vs. Writes**
- **Linear Scalability**
- **Querying a column slice or a range of row keys**
- **Data can have expiration set**
- **Reliable multiple data center replication**



LESSONS LEARNED - DATA MODELING

- ✓ **Use composite column names; static composite types**
- ✓ **Column names are stored physically sorted and indexed**
 - ✓ **Store "values" as column names**
- ✓ **Wide rows in conjunction with composite columns can be used to build indices, but...**
 - ✓ **For larger data sets, distribute the columns among rows**
 - ✓ **A secondary index is best modeled as a separate column family**



LESSONS LEARNED - DATA MODELING

- ✓ **Do not store an entity as a single column blob**
 - ✓ **Cannot index and query on entity attributes**
 - ✓ **Updates to an entity attribute would require a read and then a write**
- ✓ **Mutate just the required columns on an entity row**
 - ✓ **Do not read and then write**



LESSONS LEARNED - MONITORING

✓ **Heap Size and Use**

✓ **GC**

✓ **Mutation Stage (Writes) & Read Stage (Reads)**

✓ **Active and Pending**

✓ **AE Service, Stream & Message-Streaming-Pool Stages**

✓ **Especially during scheduled repair or rebuild (while adding nodes to a new data center)**

✓ **Streaming requires keep-alive connections (watch out for firewalls terminating idle established connections; update periodic tcp keep-alive ping rate)**

✓ **Compaction Stage & Compaction Count**

✓ **IO Wait, Limits - nofiles, nproc**



IMPLEMENTATION STRATEGY

- **Start with a project team**
 - **Riak - Product Catalog**
 - **Cassandra - Customer Graph**
- **Build expertise in development and operations**
- **As usage grows, create a Platform team**
 - **Combine engineering into one team**
 - **Project team can then focus on business features**
 - **Focus on automation**



CONCLUSIONS

- **Riak and Cassandra have both been successful**
 - **Riak - No complete outages in two years**
 - **Cassandra - Flawless in its first Holiday in 2013**
- **Differences**
 - **Replication patterns are different**
 - **Write and read treatment is different**
 - **Deployment pattern are different**
- **High Availability**
 - **Both have highly resilient architectures**
 - **Both scale linearly**





January 16, 2014

©Copyright 2014 - Best Buy, Inc. All rights reserved.