

# Putting together the platform: Riak, Redis, Solr and Spark

*Bryan Hunt*





`$ whoami`

Bryan Hunt

Client Services Engineer

@binarytemple

# Minimum viable product - the ideologically correct doctrine

---

1. Start with what you have
2. Solve problems as they arise
3. Get to the **point where you can't scale**
4. Use a **more scalable** solution





PRODUCTS

SERVICES

DEALS

Search Best Buy



Sign In

Create Account

Appliances

TV & Home Theater

Computers & Tablets

Cell Phones

Cameras & Camcorders

Audio

Car Electronics & GPS

Video Games, Movies & Music

Health, Fitness & Beauty

Home & Office

Wearable Technology

Shop by Brand

# SAVE \$100 ON SELECT ASUS 2- IN-1 LAPTOPS

Powered by Intel®. Plus, save an additional \$50-\$150 with exclusive College Student Deals.

Shop Laptops in the Student Tech Center

Free upgrade to Windows 10.

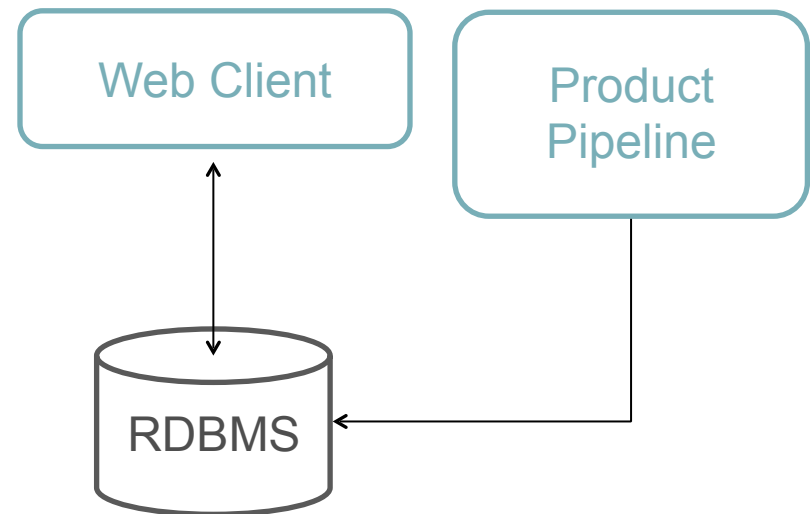
Learn more about Windows 10 ›



EXPERT **SERVICE.** UNBEATABLE **PRICE.**

# Small ecommerce site

- Sessions stored in browser cookie
- As the cart is mutated it is written to the database
- Abandoned carts timeout and items are returned to inventory





**“We’re getting slow”**

that means we’ve succeeded



## Billions of Data Points

Millions of Users, 10 Million Meters in EU  
Only 4 people devops team



**2000**  
**4 DATA POINTS**  
**PER YEAR**  
(traditional gas meter)

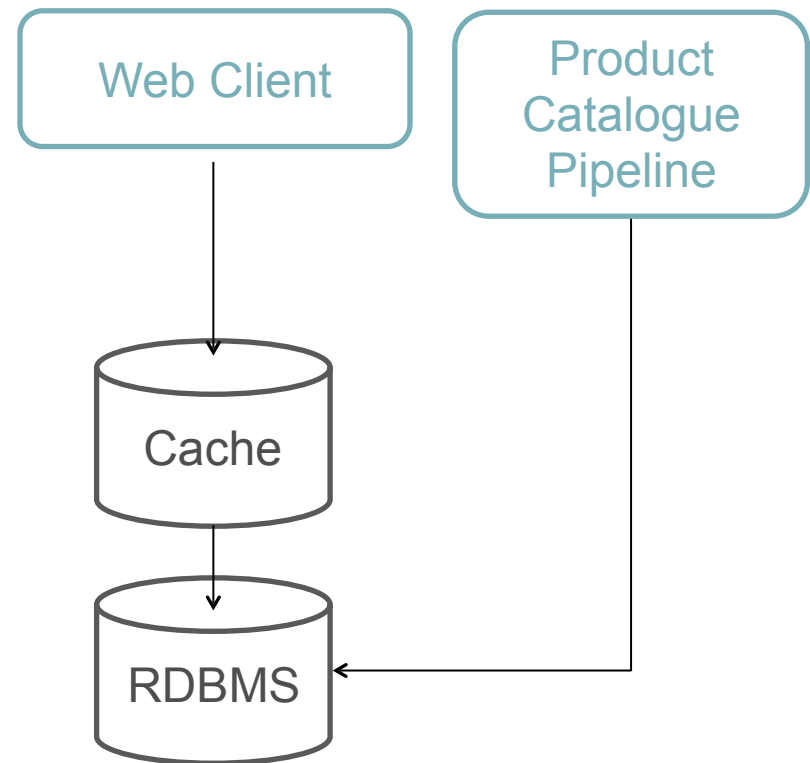


**2013**  
**35,000+ DATA POINTS**  
**PER YEAR**  
(digital meter)

# More users == slow read performance

```
if !cache $item  
  db.fetch($item)
```


- Latency increases when not cached
- Caching layer added with above logic
- Need for cache invalidation logic
- Cache warming....
- Oh my!






# Web store - now support mobile clients

- Mobile client added
- RDBMS write locks increase
- Latency still increases
- Caching layer added capacity but invalidation is a pain
- Write locks are still an issue



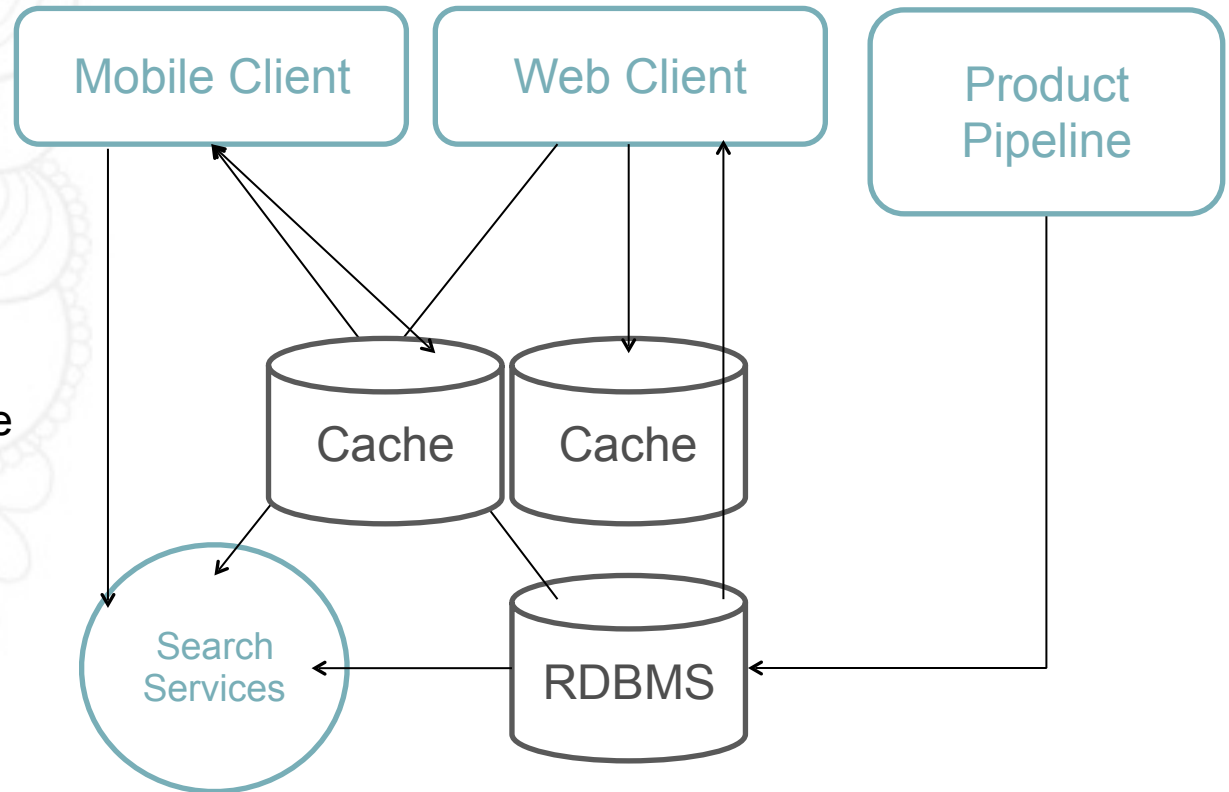
# More and more demands... make it searchable...

How can it take 3 weeks to add a simple feature?




# Web store - growing pains

- SQL profiling/ optimization
- Search is separated
- Don't forget to populate the index
- Never, never run \*that\* query



# Product Catalogue – The Feature Backlog Continues



**Scale and Availability**



**Performance**



**Mobile**



**Analytics**



**Search**

# Enter the Riak


---





# Scalable, Available, Fault Tolerant

Riak has a masterless architecture in which every node in a cluster is capable of serving read and write requests.



# Write Availability and Consistency

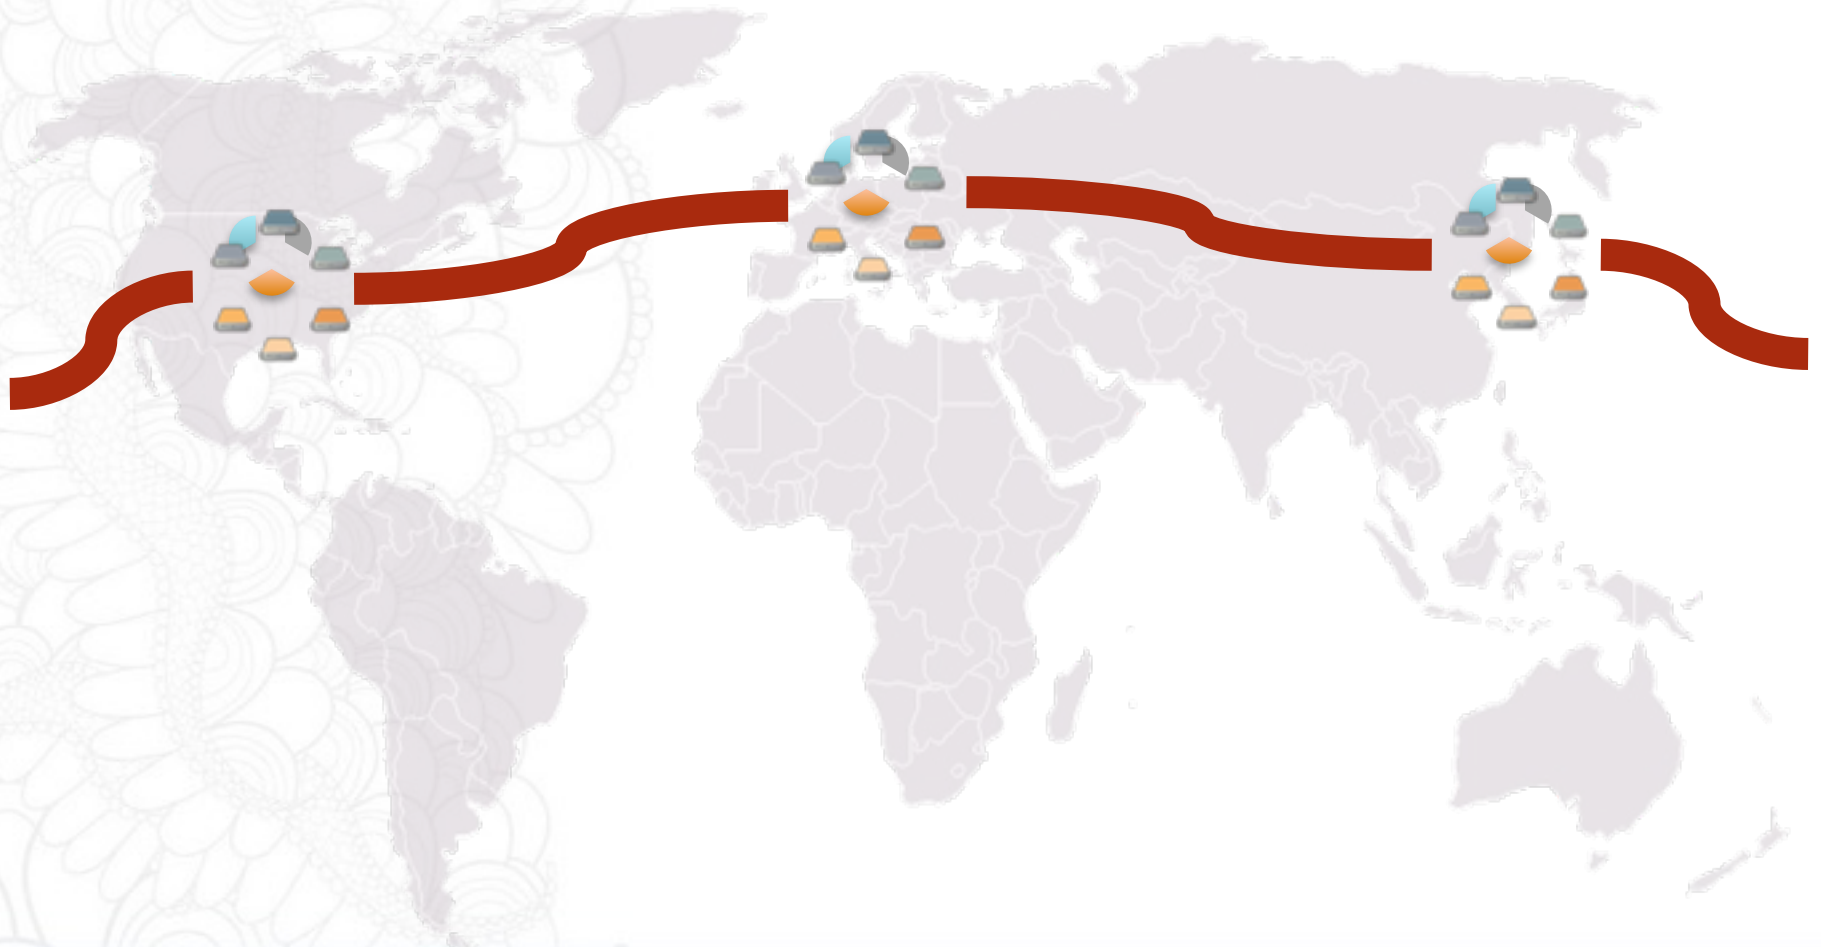
Are they mutually exclusive?

# Conflicts?

**Riak Data Types** (Convergent Replicated Data Types or CRDTs) are a developer-friendly way to keep track of updates in an eventually consistent environment:

- **Counter**  
Keeps tracks of increments and decrements on an integer
- **Flag**  
Values limited to *enable* or *disable*
- **Set**  
A collection of unique binary values that supports add and remove operations on one or more values
- **Map**  
Supports the nesting of Riak Data Types
- **Register**  
A named binary field that can only be used as part of a Map

# Multi-cluster Replication



# Dealing with Search at scale

## **Distributed Full-Text Search**

Standard full-text Solr queries automatically expand into distributed search queries for a complete result set across instances.

## **Index Synchronization**

Data is automatically synchronized between Riak KV and Solr using intelligent monitoring to detect changes, and propagates those to Solr indexes.


## **Auto-Restart**

Monitor Solr OS processes continuously and automatically start or restart them whenever failures are detected.

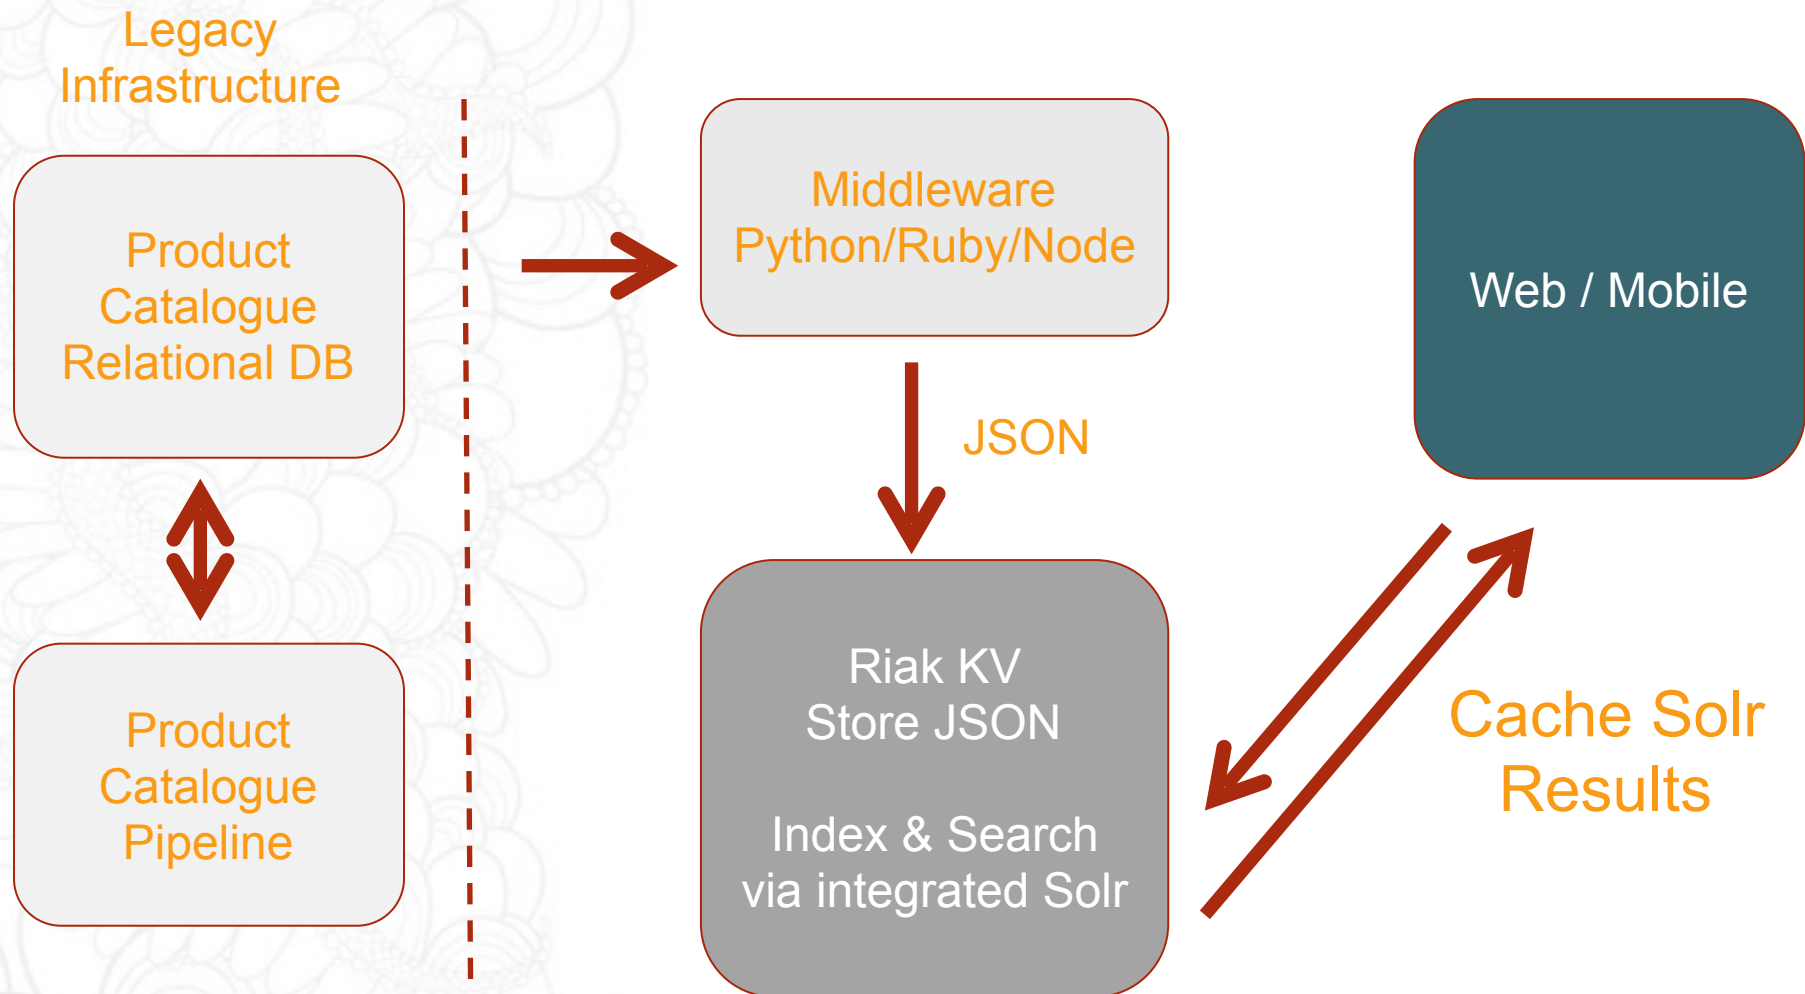


# Make Search Better

we need to get a competitive advantage



# Web store: One step closer to success



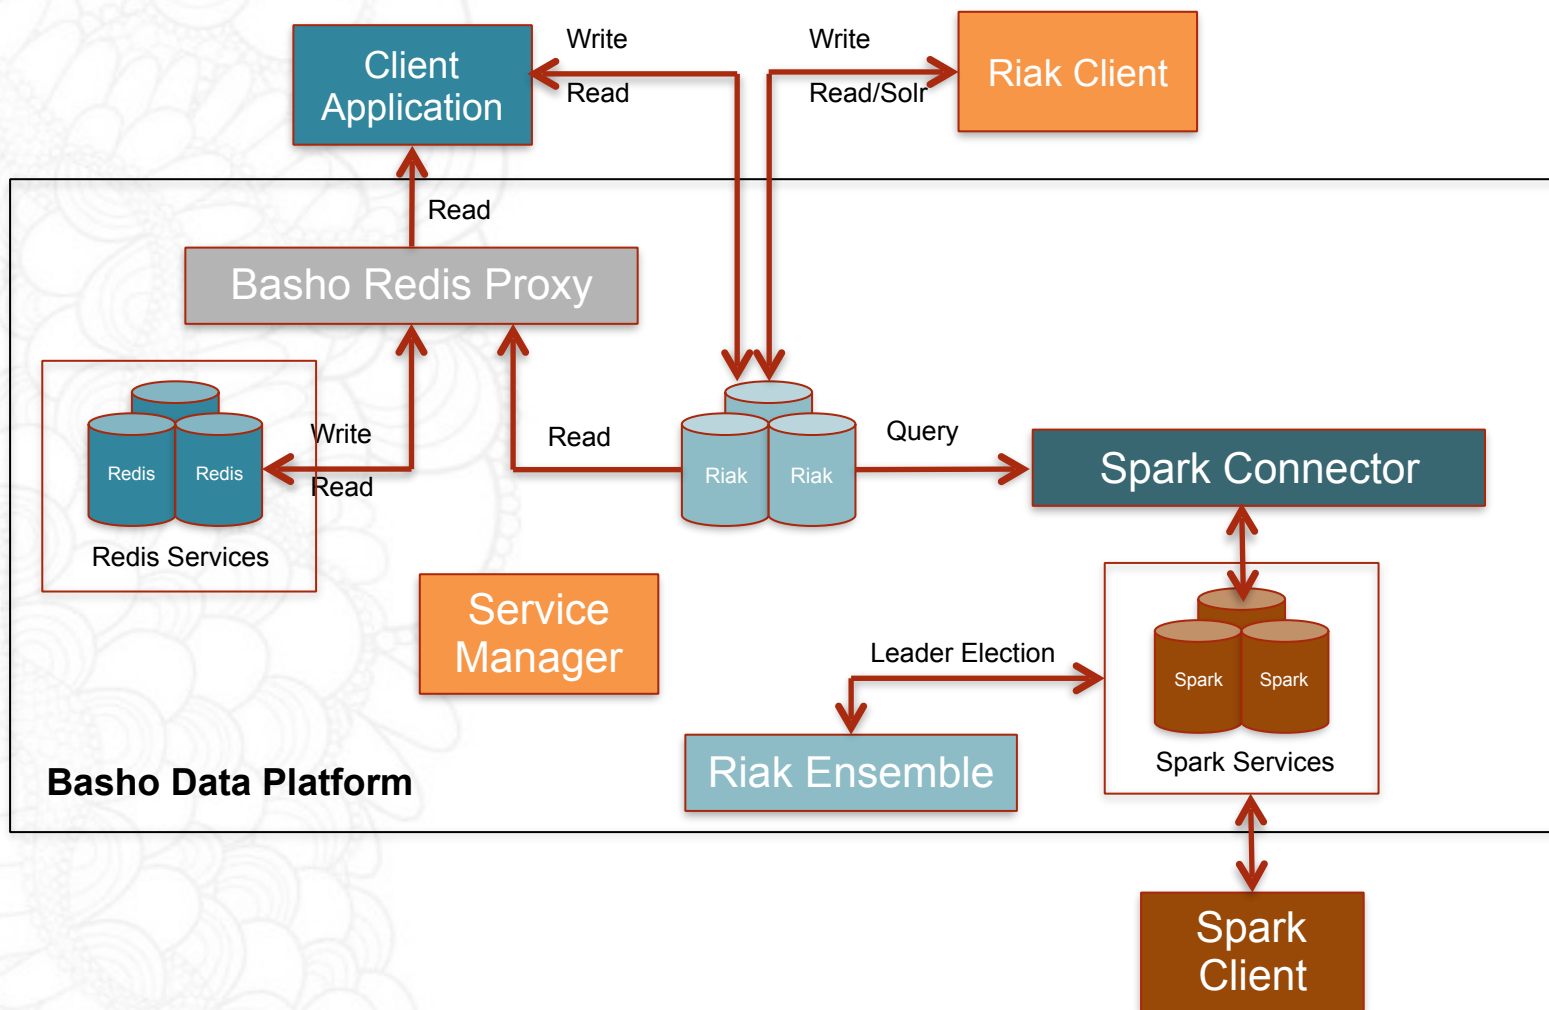
# Basho Data Platform

---



## What it provides..

- Easy to use caching layer
- Automatically indexed search
- Scalable, highly available persistence layer
- Simple pipelining to data analysis





# Caching

---



# Redis - why ?

# Why Use Redis With The Data Platform?

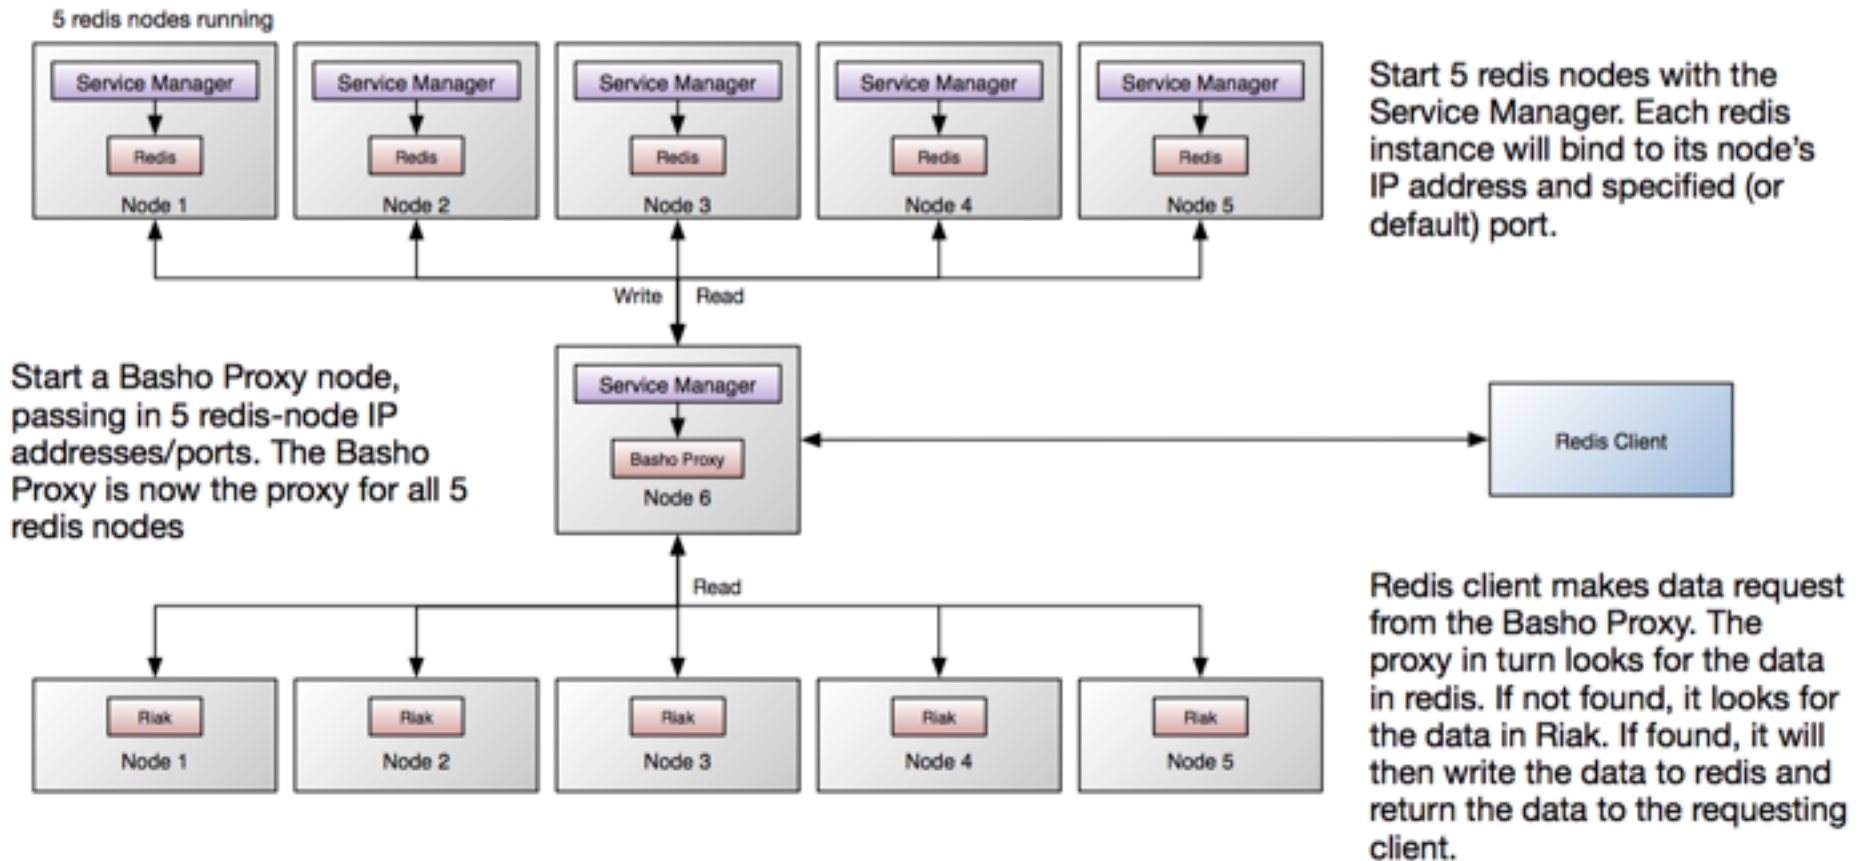
- Redis, by itself, is approximately 6x faster than Riak for reads
- Redis is not built to scale

## With the **Basho Data Platform**:

- Approximately 5x faster read performance compared to Riak for buckets at an 80% cache hit ratio.


# Basho Data Proxy & Redis

## Data Platform Basho Proxy - Redis




# Resilience

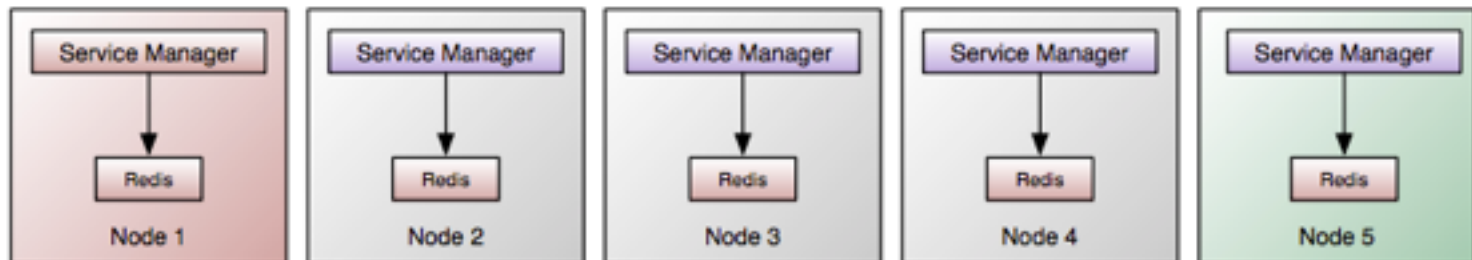
5 redis nodes running



Node 5 goes down, still 4 redis nodes running



Node 1 goes down, node 5 returns to service, 4 redis nodes running





# Basho Data Proxy & Redis

## Redis

- v2.8.21
  - Redis Cluster (3.0+) need not apply
- Cache Proxy using only GET and SET
  - SET uses a TTL, configured as `CACHE_TTL`
- Client may use other commands such as:
  - PEXPIRE to invalidate a cached value
  - INFO to integrate with monitoring solutions
  - MONITOR to watch traffic to Redis hosts

## Basho Data Proxy


- Written in C
- Read-through cache in BDP 1.0
- Write-through cache in a future BDP release
- EE only, until write-through cache is complete and released
- Pre-sharding & connection aggregation built-in

# Analytics

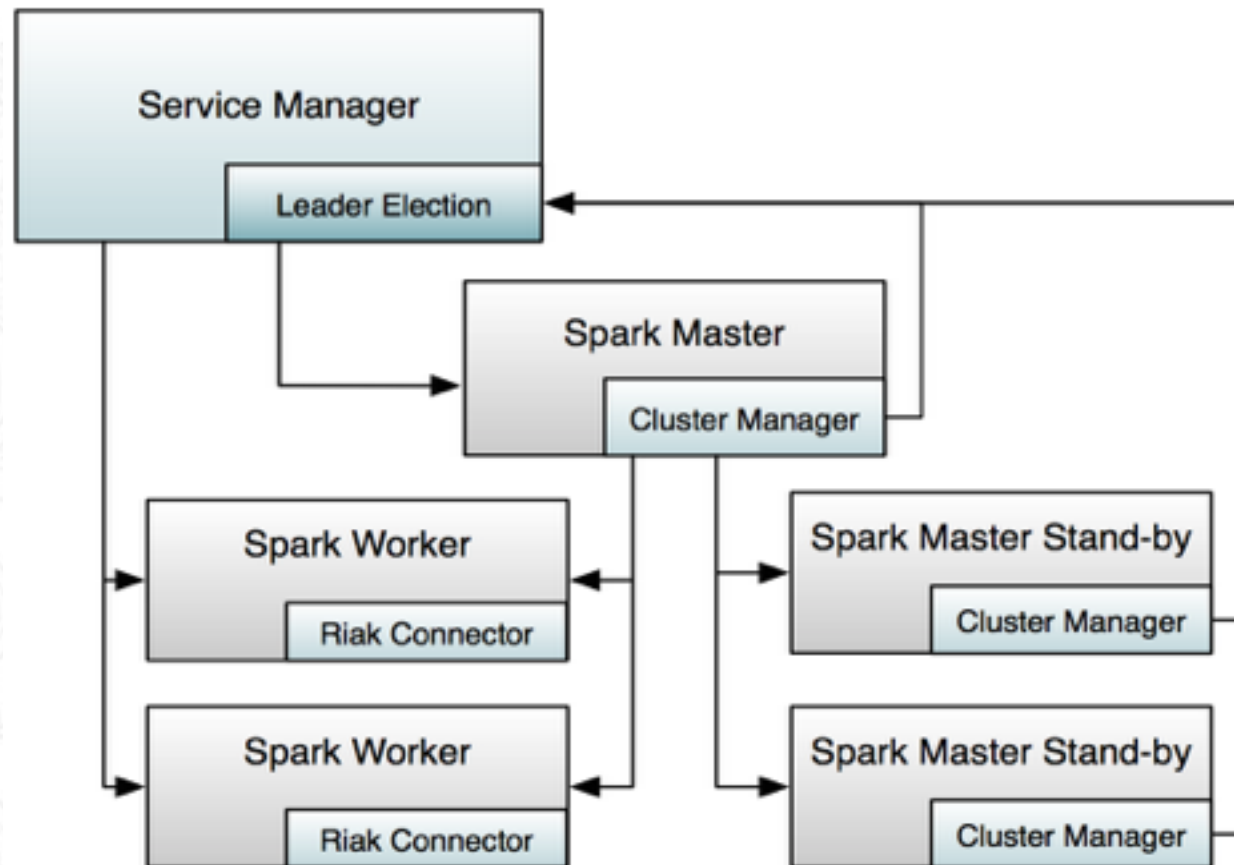
---



# Connecting to BDP with Spark



# Data Platform Spark Cluster



# Spark & Spark Cluster Manager

## Spark

- V1.4.1
- Dependencies
  - Java 8
  - Scala 2.10.5

## Spark Cluster Manager

- Shipped as a JAR in the BDP extras package
- Enables running a HA Spark cluster
  - Spark Master
  - Spark Stand-by Master(s)
- Implemented as a custom Spark recovery mode
- Stores its metadata in a strongly consistent Riak bucket



# Why Basho Data Platform?

---







PRODUCTS

SERVICES

DEALS

Search Best Buy



Sign In | Create Account

Appliances

TV & Home Theater

Computers & Tablets

Cell Phones

Cameras & Camcorders

Audio

Car Electronics & GPS

Video Games, Movies & Music

Health, Fitness & Beauty

Home & Office

Wearable Technology

Shop by Brand

# SAVE \$100 ON SELECT ASUS 2- IN-1 LAPTOPS

Powered by Intel®. Plus, save an additional \$50-\$150 with exclusive College Student Deals.

Shop Laptops in the Student Tech Center

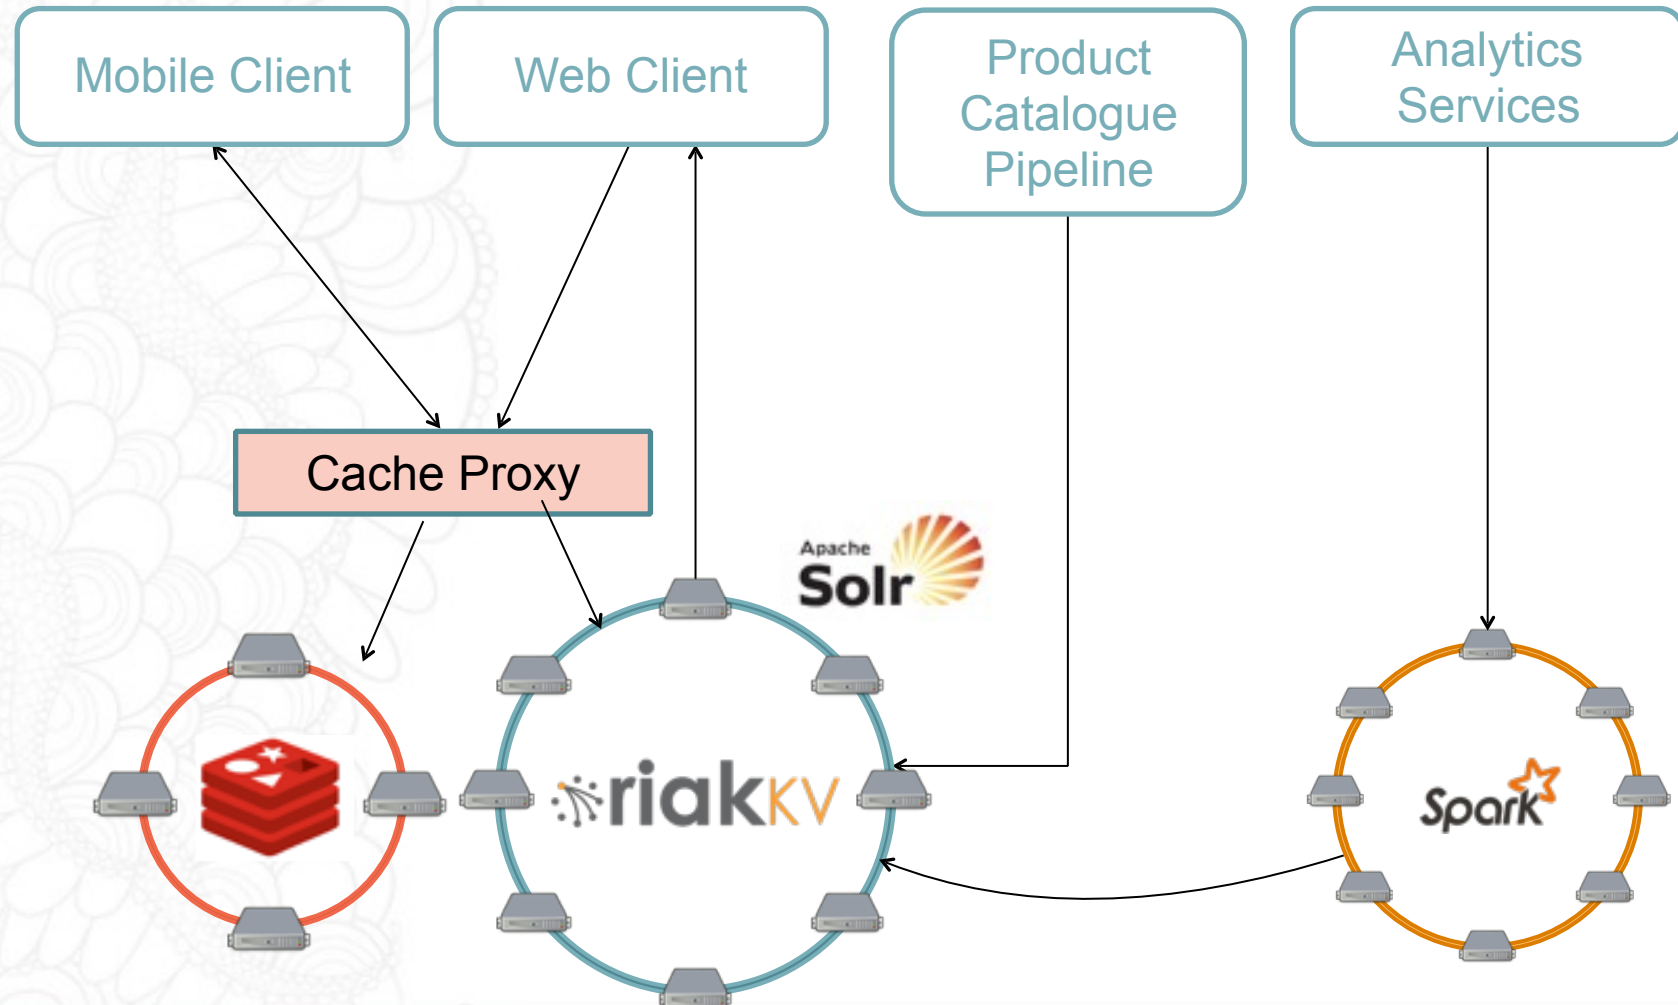
Free upgrade to Windows 10.

Learn more about Windows 10 ›



EXPERT **SERVICE.** UNBEATABLE **PRICE.**

# Product Catalogue with Basho Data Platform




**Faster  
Analysis**



**More  
IOPS/TBs**



**Faster  
Reads**





**QUESTIONS?**

# Spend Time

getting to know us

## OPEN SOURCE

## ENTERPRISE

### **Basho Data Platform ([code](#))**

- **Riak KV with parallel extract**

### **Basho Data Platform, Enterprise**

- **Riak EE with multi-cluster replication**
- **Spark Leader Election Service**

### **Basho Data Platform Add-on's ([code](#))**

- **Spark + Spark Connector**

### **Basho Data Platform Add-on's**

- **Redis + Cache Proxy**
- **Spark Workers + Spark Master**

**[Download a build](#)**

**[Contact us](#) to get started**



**@basho**

**@binarytemple**





*Please*

**Remember to  
rate this session**

*Thank you!*



# Thanks !