



# Lambda Architecture

Learn to stop worrying and love human fault tolerance

Jim Scott, Director, Enterprise Strategy and Architecture

Spark Summit 2014



# Fault tolerance

developer

software

hardware

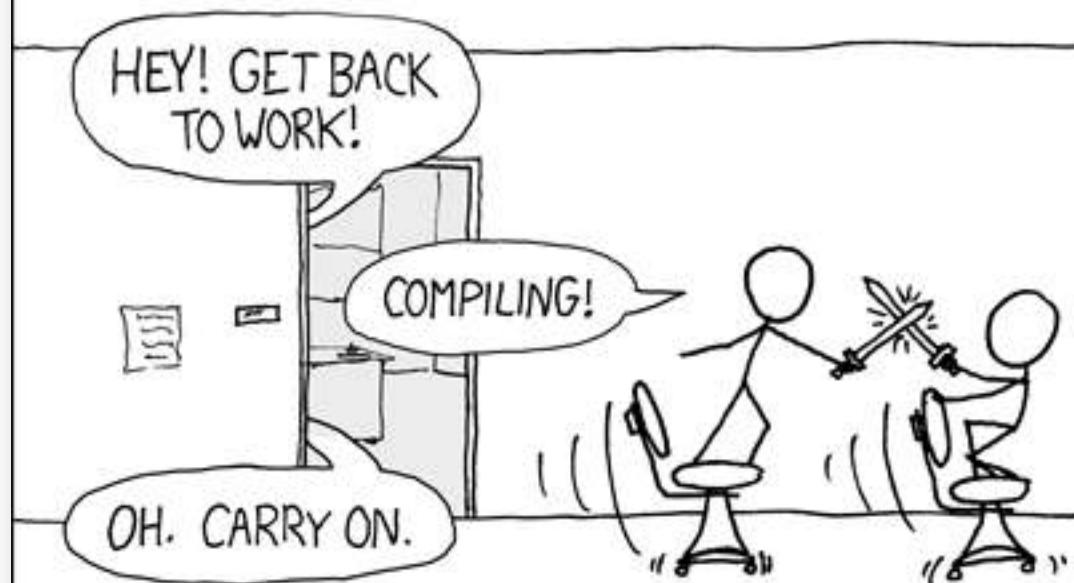




# Let's talk about developers...



THE #1 PROGRAMMER EXCUSE  
FOR LEGITIMATELY SLACKING OFF:  
"MY CODE'S COMPILING."



<http://xkcd.com/303/>

HI, THIS IS  
YOUR SON'S SCHOOL.  
WE'RE HAVING SOME  
COMPUTER TROUBLE.



OH, DEAR - DID HE  
BREAK SOMETHING?  
IN A WAY -



DID YOU REALLY  
NAME YOUR SON  
Robert'); DROP  
TABLE Students;-- ?



OH, YES. LITTLE  
BOBBY TABLES,  
WE CALL HIM.

WELL, WE'VE LOST THIS  
YEAR'S STUDENT RECORDS.  
I HOPE YOU'RE HAPPY.

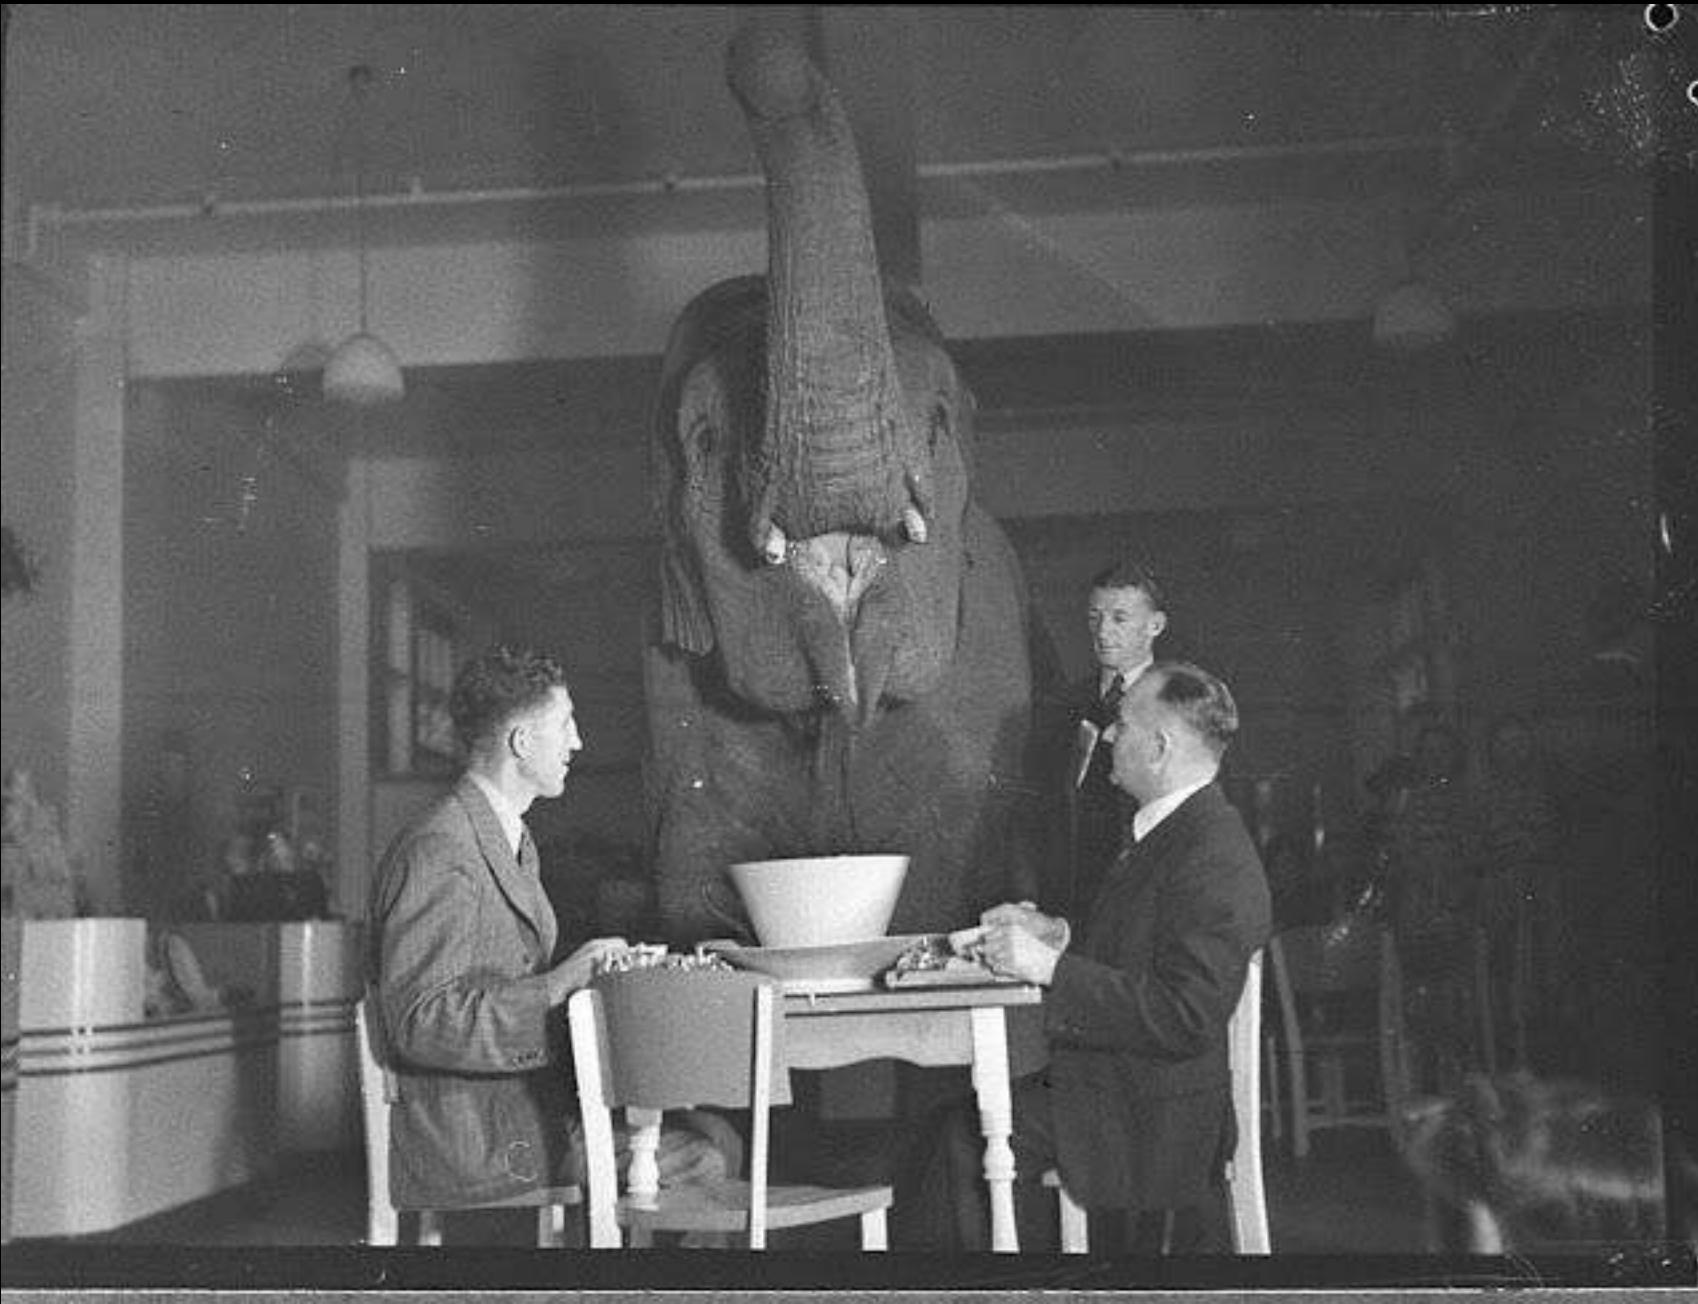


AND I HOPE  
YOU'VE LEARNED  
TO SANITIZE YOUR  
DATABASE INPUTS.

<http://xkcd.com/327/>



MapR







Let's talk about **developer** **first**. tolerance





# When things go wrong ...

## **The Real Reason Facebook Went Down Yesterday: It's Complicated**

Nick O'Neill on September 24, 2010 10:20 AM

Yesterday afternoon Facebook experienced the worst outage that the company has had "in over four years", causing the site to go down for most users for "approximately 2.5 hours". One of the company's engineers followed up with a blog post, explaining exactly what went wrong. The cause of the issue sounds relatively complicated, however the conclusion was that the company had to restart the entire site.

According to Robert Johnson:

The key flaw that caused this outage to be so severe was an unfortunate handling of an error condition. An automated system for verifying configuration values ended up causing much more damage than it fixed.

The intent of the automated system is to check for configuration values that are invalid in the cache and replace them with updated values from the persistent store. This works well for a transient problem with the cache, but it doesn't work when the persistent store is invalid.



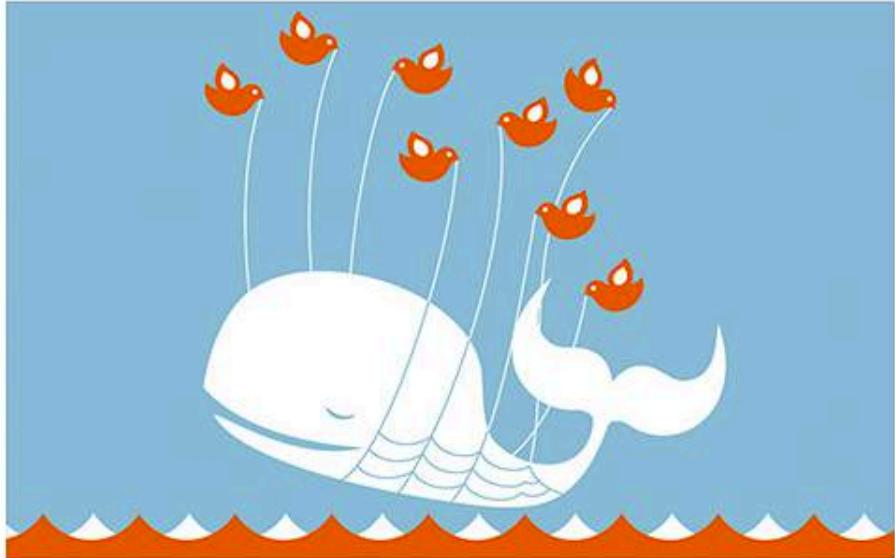
2010  
**unfortunate  
handling of  
error condition**

# When things go wrong ...

**Twitter crashes hard,  
Internet freaks out**

By Julianne Pepitone @CNNMoneyTech June 21, 2012: 3:34 PM ET

**f** Recommend 2k **Twitter** 2,134 **Share** **g+1** 212 Email Print



FROM: TWITTER

Twitter crashed so hard on June 21 that the site didn't even display the famous "Fail Whale." Instead, it simply timed out.

2012  
**cascaded bug**

# When things go wrong ...

## RBS takes £125m hit over IT outage

by Dan Worth 03 Aug 2012

**FOLLOW V3** [Follow @V3\\_co\\_uk](#) [Like](#) [LinkedIn](#) [Follow](#) [Newsletters](#)



The RBS/Natwest technical outage that blighted the bank and its customers in June has cost the firm £125m according to its latest financial reports.

The banking group, which posted overall losses of £1.5bn for the second quarter of its financial year, was hit by a huge outage that affected millions of customers from receiving or making payments and lasted for almost an entire week.

"The immediate software issue was promptly identified and rectified. Despite this, significant manual intervention in a highly automated and complex batch processing environment was required. This resulted in a significant backlog of daily data and information processing," the firm said in its filing.

2012  
**upgrade of batch processing**

# When things go wrong ...

## Google explains reasons behind today's 30-minute service outage

NEWS By Andrew Martonik | Jan 24 2014 | 9:50 pm | 56 COMMENTS



2014  
**bug/bad config**

Proverbial 'software bug' sent a spiral of bad configurations to other systems



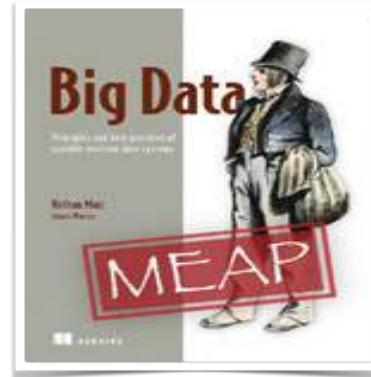


# Lambda Architecture to the rescue!



# Let's step back a bit ...

- Nathan Marz (Backtype, Twitter, stealth startup)
- Creator of ...
  - Storm
  - Cascalog
  - ElephantDB

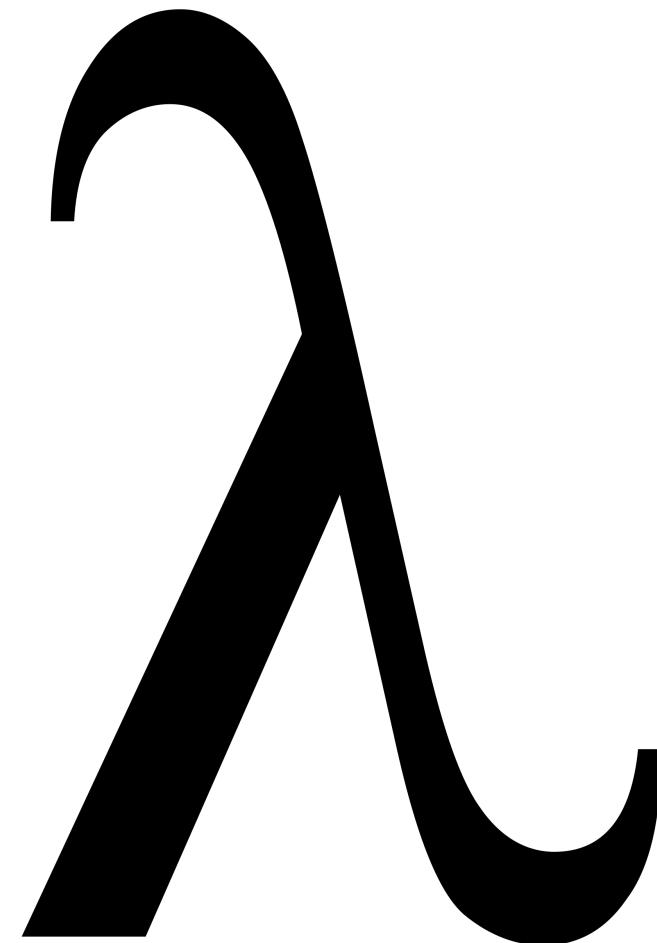


<http://manning.com/marz/>

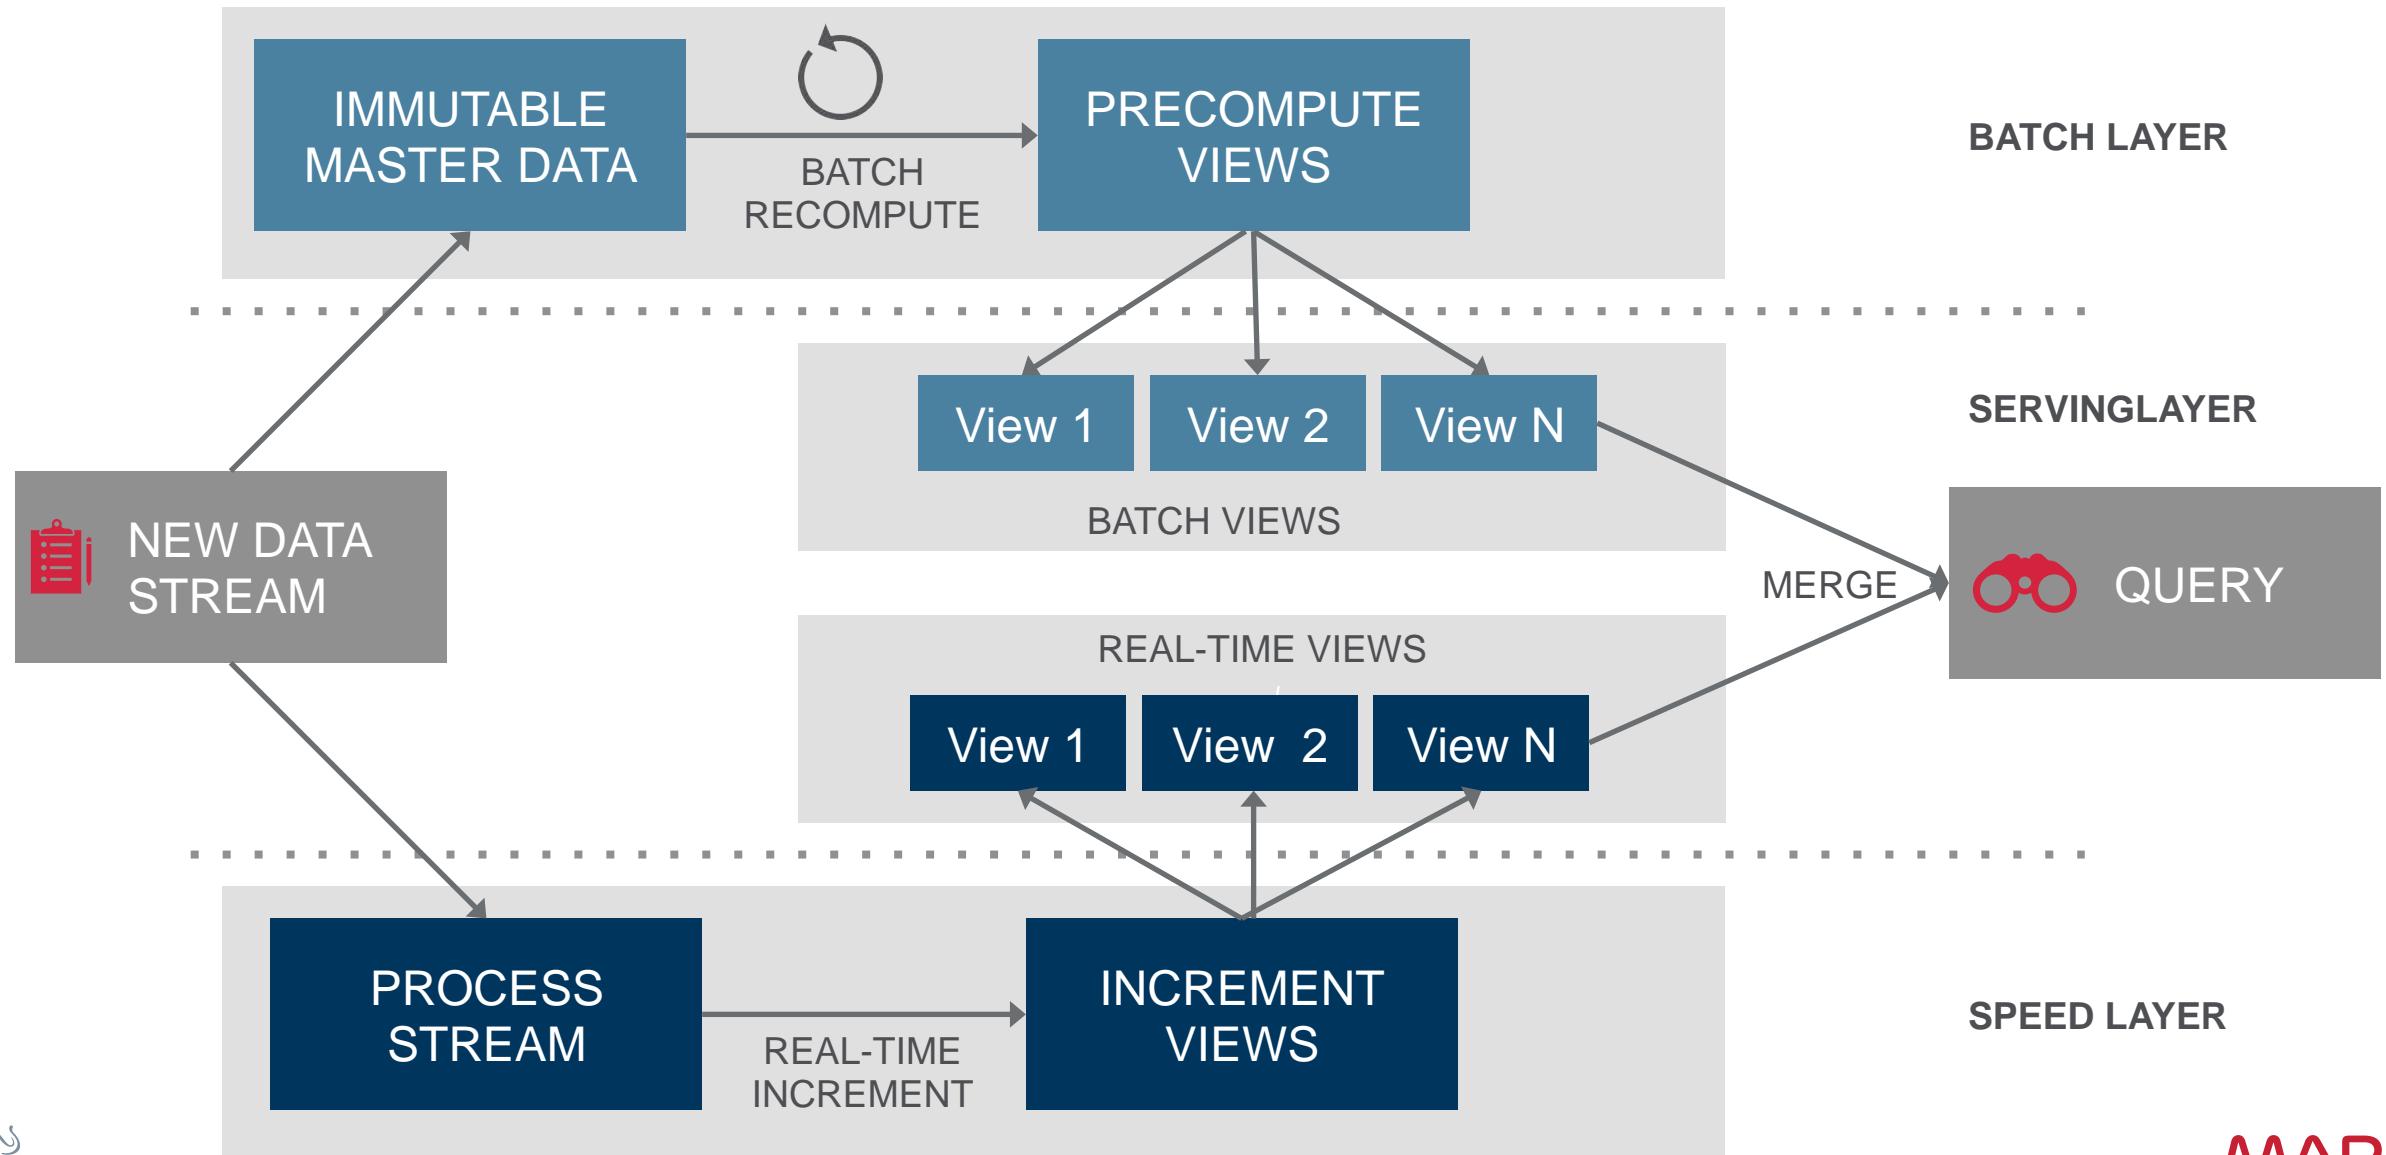
# Lambda Architecture—Requirements

- Fault-tolerant against both hardware failures and human errors
- Support variety of use cases that include low latency querying as well as updates
- Linear scale-out capabilities
- Extensible, so that the system is manageable and can accommodate newer features easily

# Why call it the Lambda Architecture?



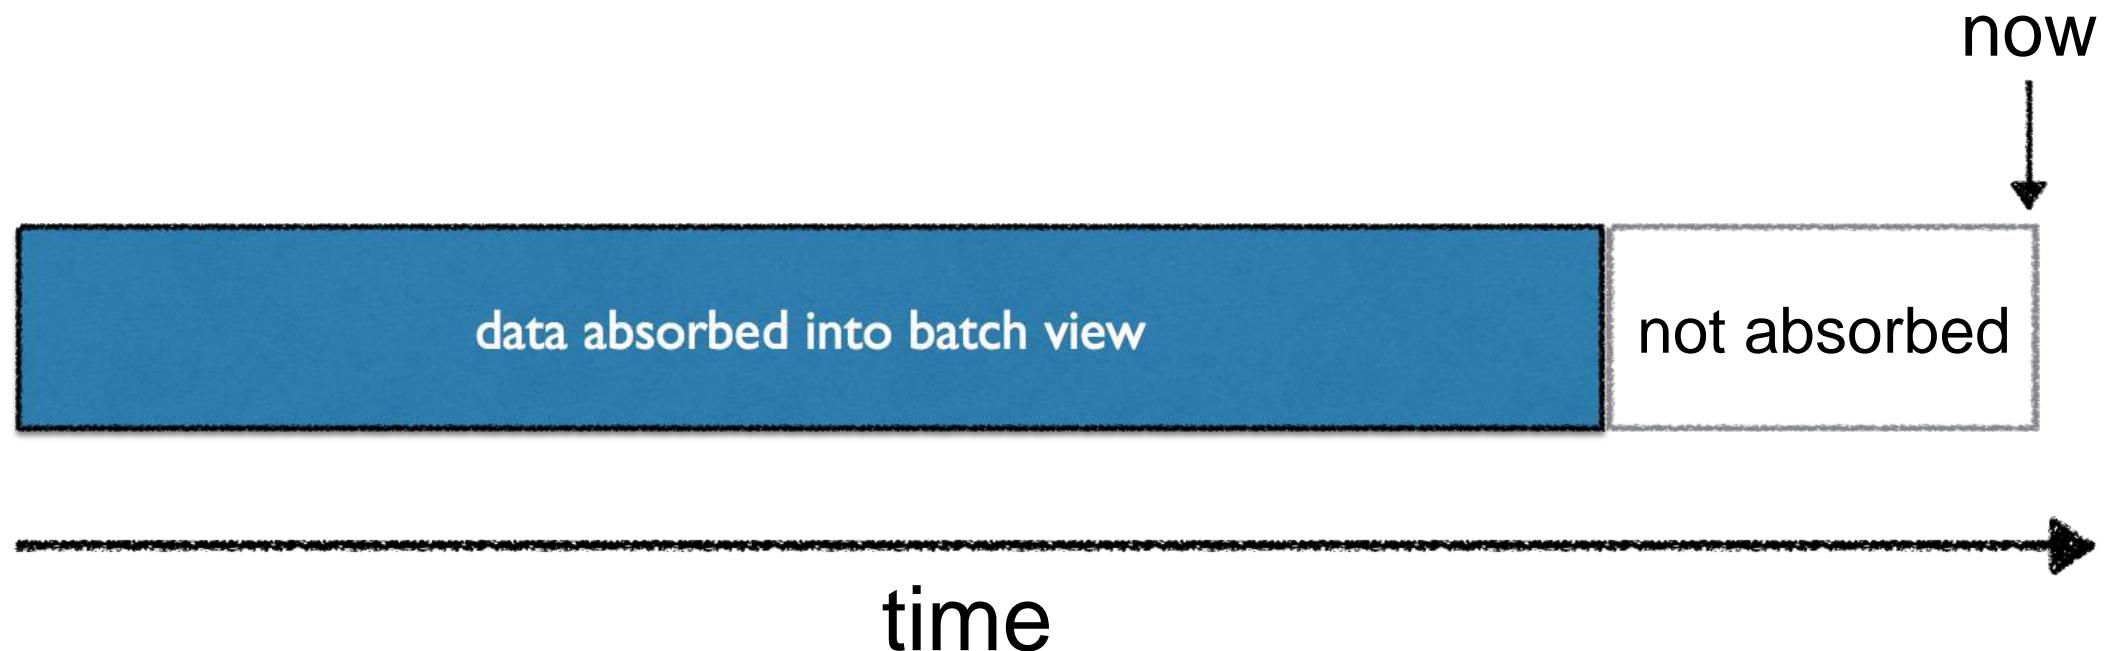
# Lambda Architecture



# Lambda Architecture—Layers

- Batch layer
  - managing the master dataset, an immutable, append-only set of raw data
  - pre-computing arbitrary query functions, called batch views
- Serving layer indexes batch views so that they can be queried in ad hoc with low latency
- Speed layer accommodates all requests that are subject to low latency requirements. Using fast and incremental algorithms, deals with recent data only

# Lambda Architecture—Compensate Batch



# Lambda Architecture—Immutable Data + Views



<http://openflights.org>



MAPR

# Lambda Architecture—Immutable Data + Views

timestamp	airport	flight	action
2014-01-01T10:00:00	DUB	EI123	take-off
2014-01-01T10:05:00	HEL	SAS45	take-off
2014-01-01T10:07:00	AMS	BA99	take-off
2014-01-01T10:09:00	LHR	LH17	landing
2014-01-01T10:10:00	CDG	AF03	landing
2014-01-01T10:10:00	FCO	AZ501	take-off

immutable master dataset

# Lambda Architecture—Immutable Data + Views

timestamp	airport	flight	action
2014-01-01T10:00:00	DUB	EI123	take-off
2014-01-01T10:05:00	HEL	SAS45	take-off
2014-01-01T10:07:00	AMS	BA99	take-off
2014-01-01T10:09:00	LHR	LH17	landing
2014-01-01T10:10:00	CDG	AF03	landing
2014-01-01T10:10:00	FCO	AZ501	take-off

immutable master dataset



air-borne: 2307



air-borne per airline:

airport	planes
AMS	69
CDG	44
DUB	31
FCO	10
HEL	17
LHR	101

air-borne per airline:

airline	planes
AF	59
AZ	23
BA	167
EI	19
LH	201
SAS	28

views



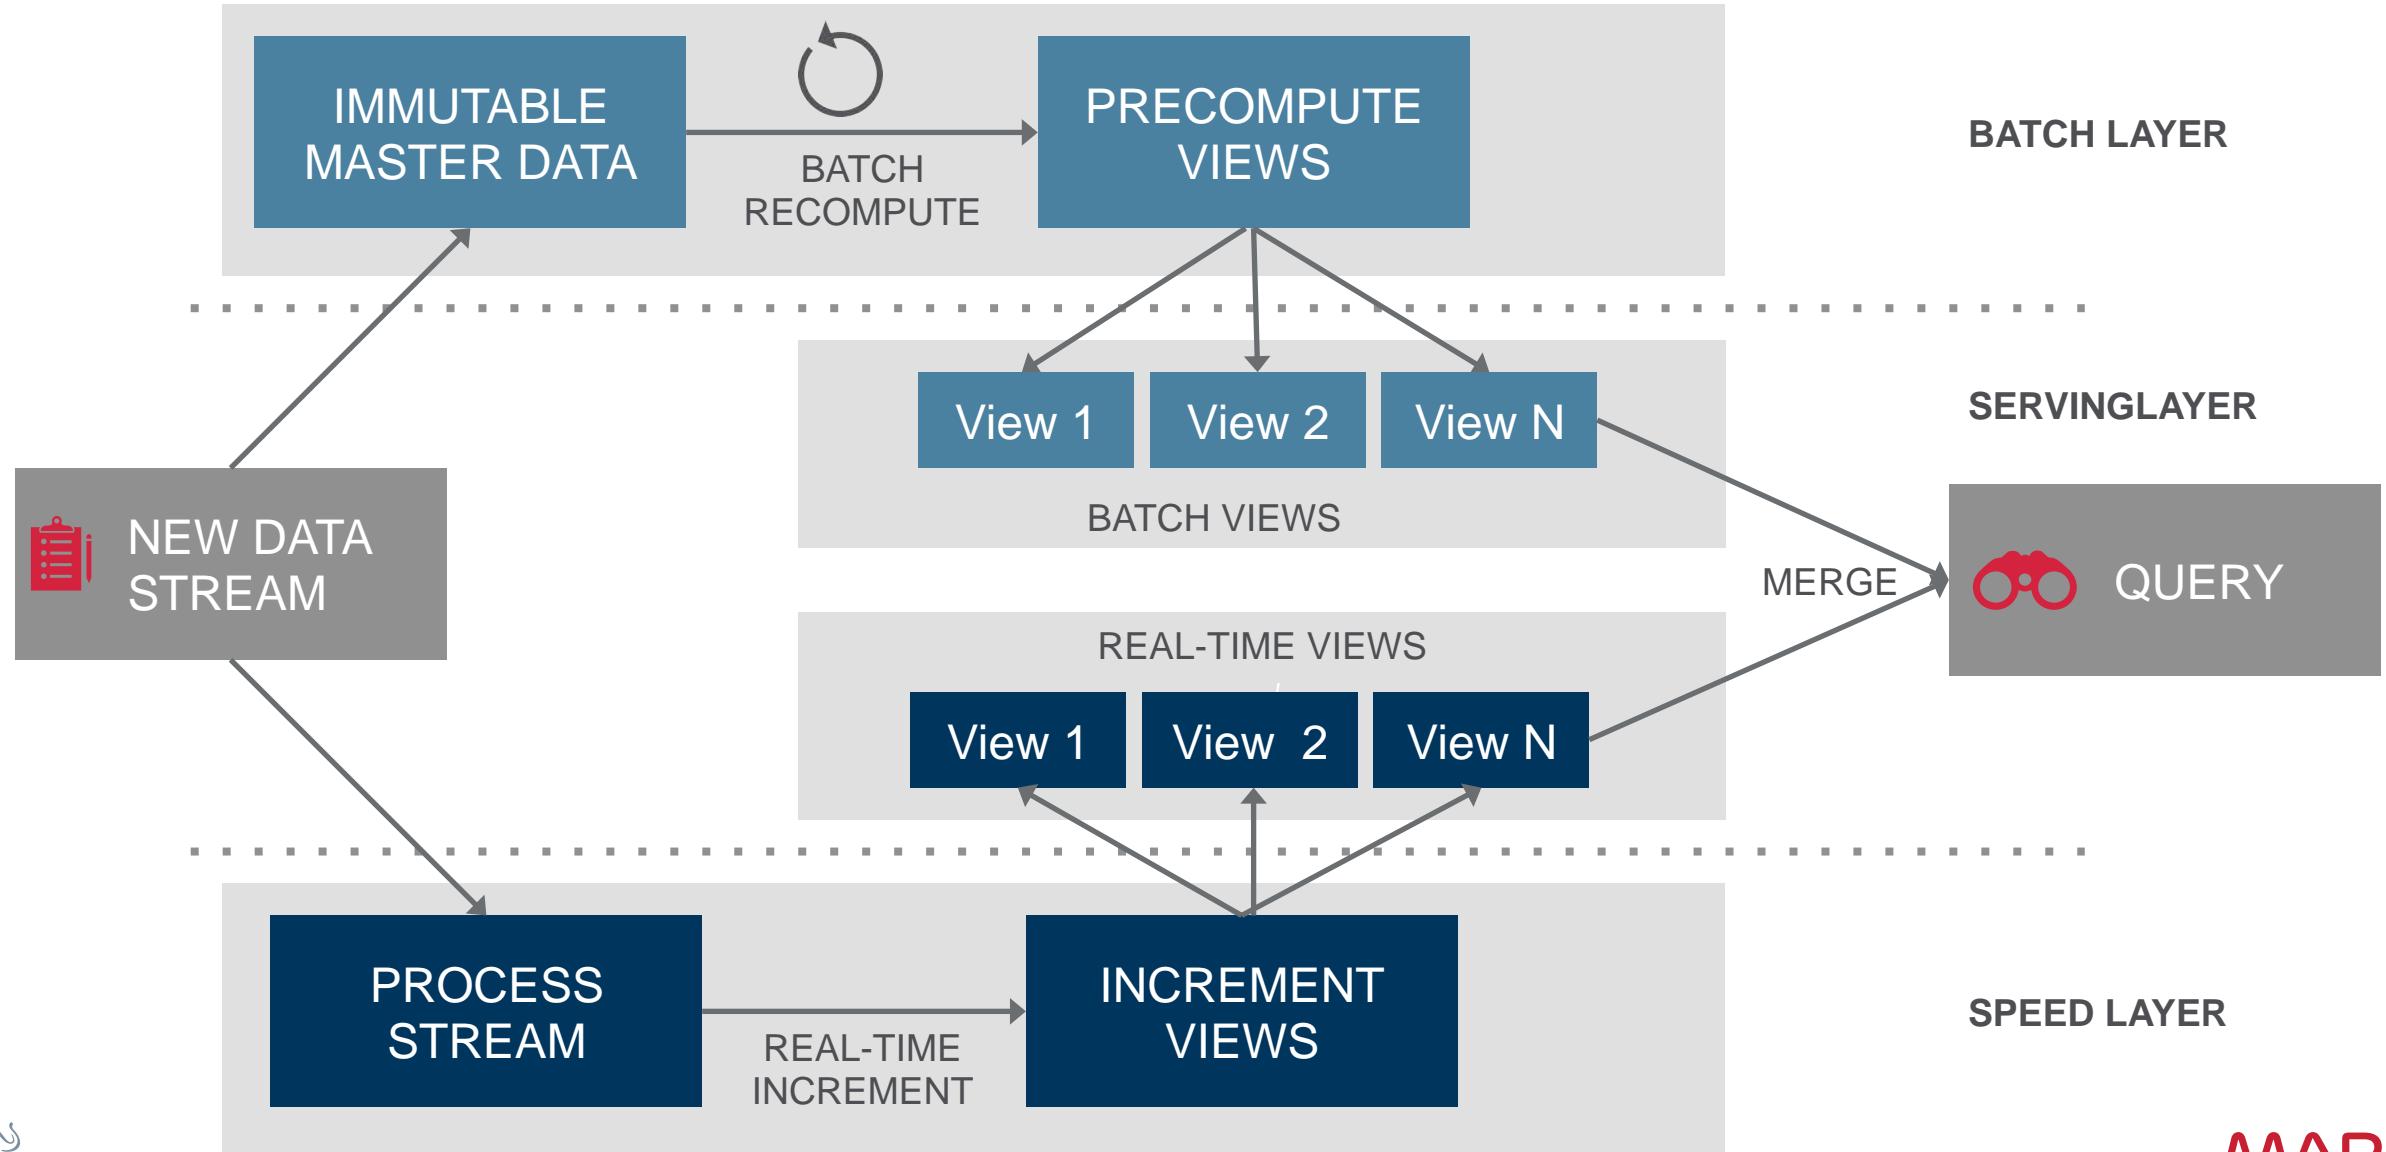
MapR



# Implementing the Lambda Architecture



# Lambda Architecture



# Components

## Processing Frameworks

Technology	Does it fit	Maturity	Easyness of use	Language	Platforms	Comments
Hadoop MapReduce	★★★	★★★	★	Java	Hadoop	
Spark MapReduce	★★★	★	★★	Scala	Spark	
Pig	★★★	★★★	★★★	Pig Latin, Java	Hadoop	Support planned for Spark and Tez
Hive	★★★	★★★	★★★	HiveQL, Java	Hadoop	Support planned for Tez. Spark's clone is "Shark"
Cascading / Scalding	★★★	★★	★			
Cascalog	★★★	★	★			
Crunch / SCrunch	★★★	★★	★			
Pangool	★★★	★	★			

## Batch View Databases

Technology
ElephantDB
SploutSQL
Voldemort (with a ReadOnly backend)
HBase (bulk loading)

## Speed Components

24 December 2013

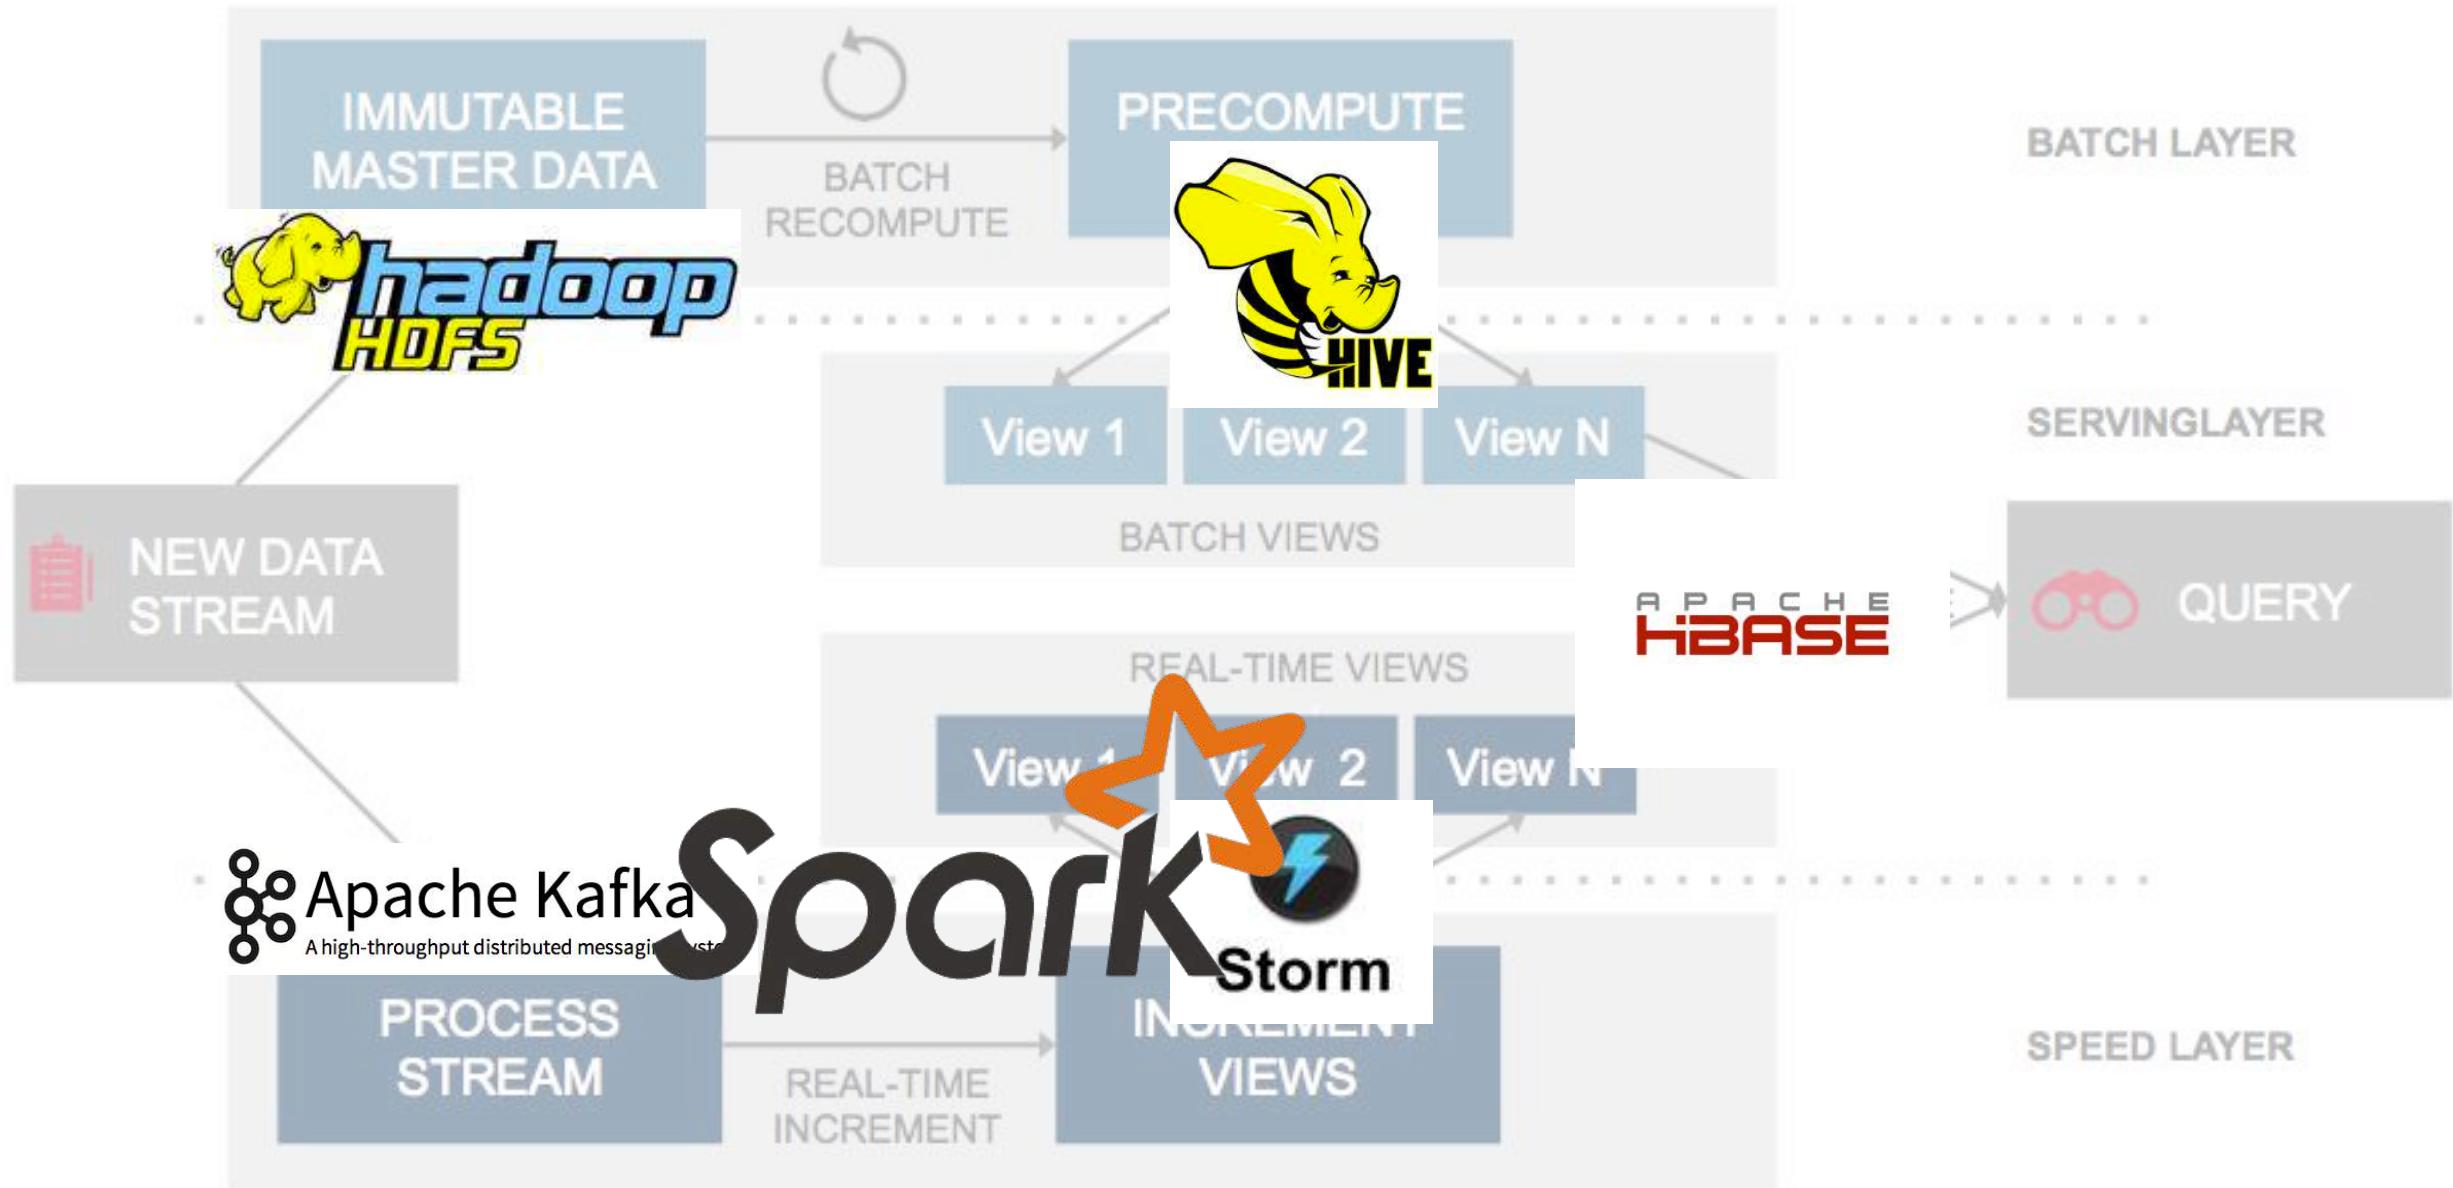
## Stream Processing Frameworks

Technology	Does it fit	Maturity	Complexity	Language	Comments
Apache Storm	★★★	★★★	★	Clojure	originates from Twitter
Apache Samza	★★★	★★	★★	Scala/Java	originates from LinkedIn
Apache S4	★★★	★	★	Java	originates from Yahoo!

## Cloud-based (XaaS) Offerings

Technology	Does it fit	Maturity	Complexity	API	Comments
Kinesis	★★★	★★	★★	Java	introduced in 11/2013



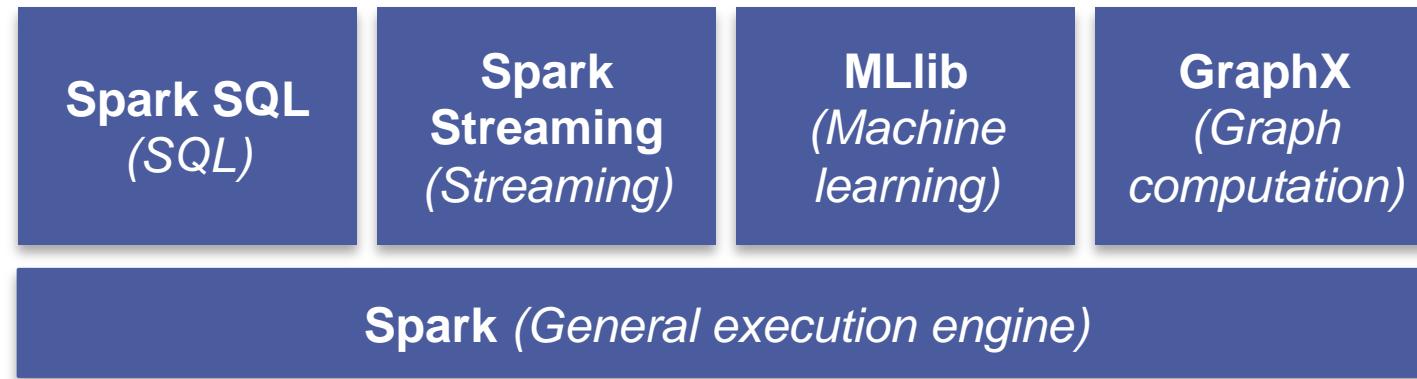


# How about an integrated approach?

- Twitter Summingbird
- Lambdoop
- Apache Spark



# A Unified Platform ...



Continued innovation bringing new functionality, e.g.:

- **BlinkDB** (Approximate Queries)
- **SparkR** (R wrapper for Spark)
- **Java 8** (Closures)



# The Spark Stack from 100,000 ft

4

Spark ecosystem

3

Spark core engine

2

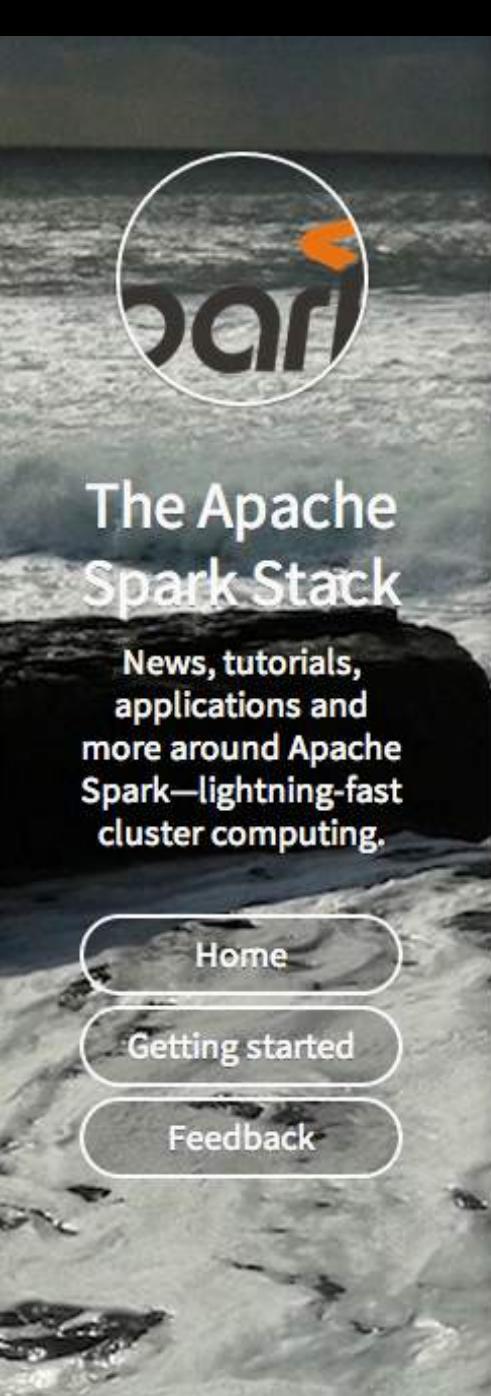
Execution environment

1

Data platform



MAPR



## Spark on fire

30.04.2014

Curt Monash's insightful analysis of the current state of the Spark space incl. ecosystem components, use cases, Hadoop vendor support and upcoming developments: [http://www.dbms2.com/2014/04/30/spark-on-fire/...](http://www.dbms2.com/2014/04/30/spark-on-fire/)

[Read more](#)

## Spark 1.0 and Beyond

28.04.2014

Patrick Wendell at the Bay Area Spark User Meetup....

[Read more](#)

## Latest posts

### Spark on fire

6 days ago

### Spark 1.0 and Beyond

8 days ago

### Adding Native SQL Support to Spark with Catalyst

8 days ago

### Real-time big data processing with Spark Streaming

8 days ago

### Spark at Twitter

9 days ago

<http://spark-stack.org>

# Q & A

Engage with us!

@mapr



maprtech

mapr-technologies



MapR

jscott@mapr.com



maprtech



MAPR