



Berlin Buzzwords, May 2014

Fault tolerance

developer

software

hardware



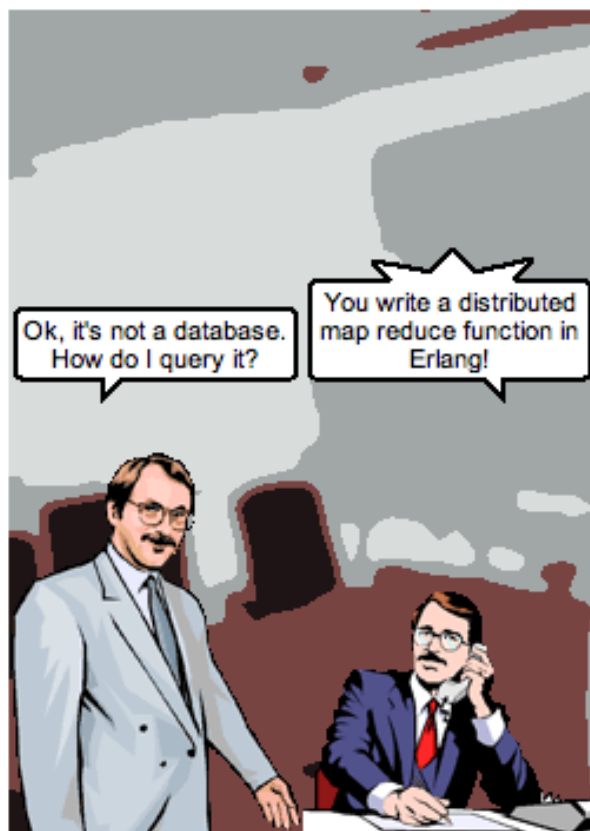
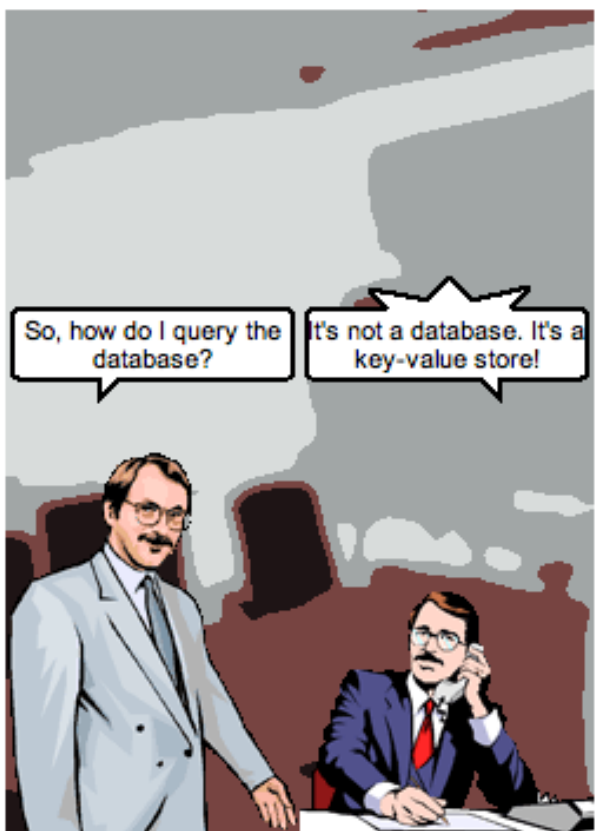
Let's talk about developers ...

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





Fault-tolerance





Let's talk about **developmental 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.



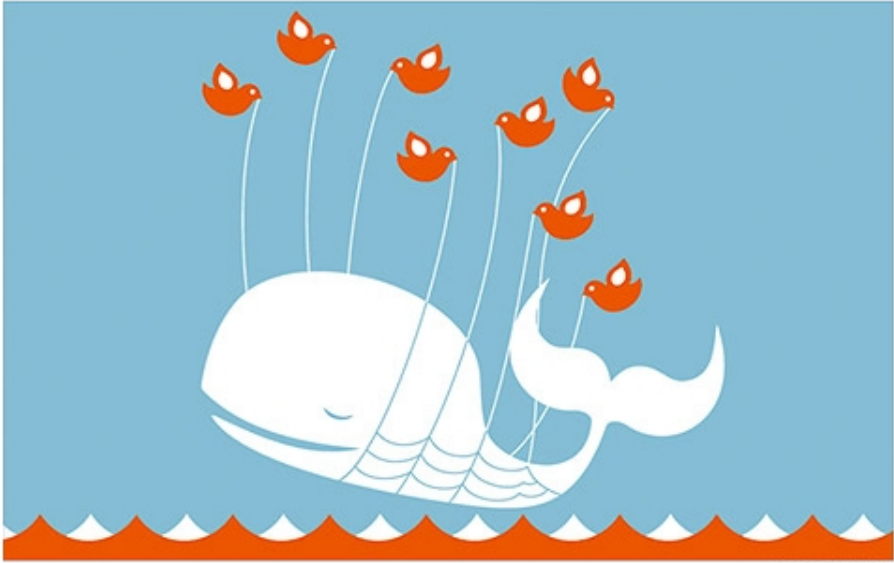
2010
**unfortunate
handling of
error condition**

When things go wrong ...

Twitter crashes hard, Internet freaks out CNNMoney

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

Recommend 2k Tweet 2,134 Share g+ 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



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 affecting 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



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

2014
bug/bad config

Lambda Architecture to the rescue!

Let's step back a bit ...

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

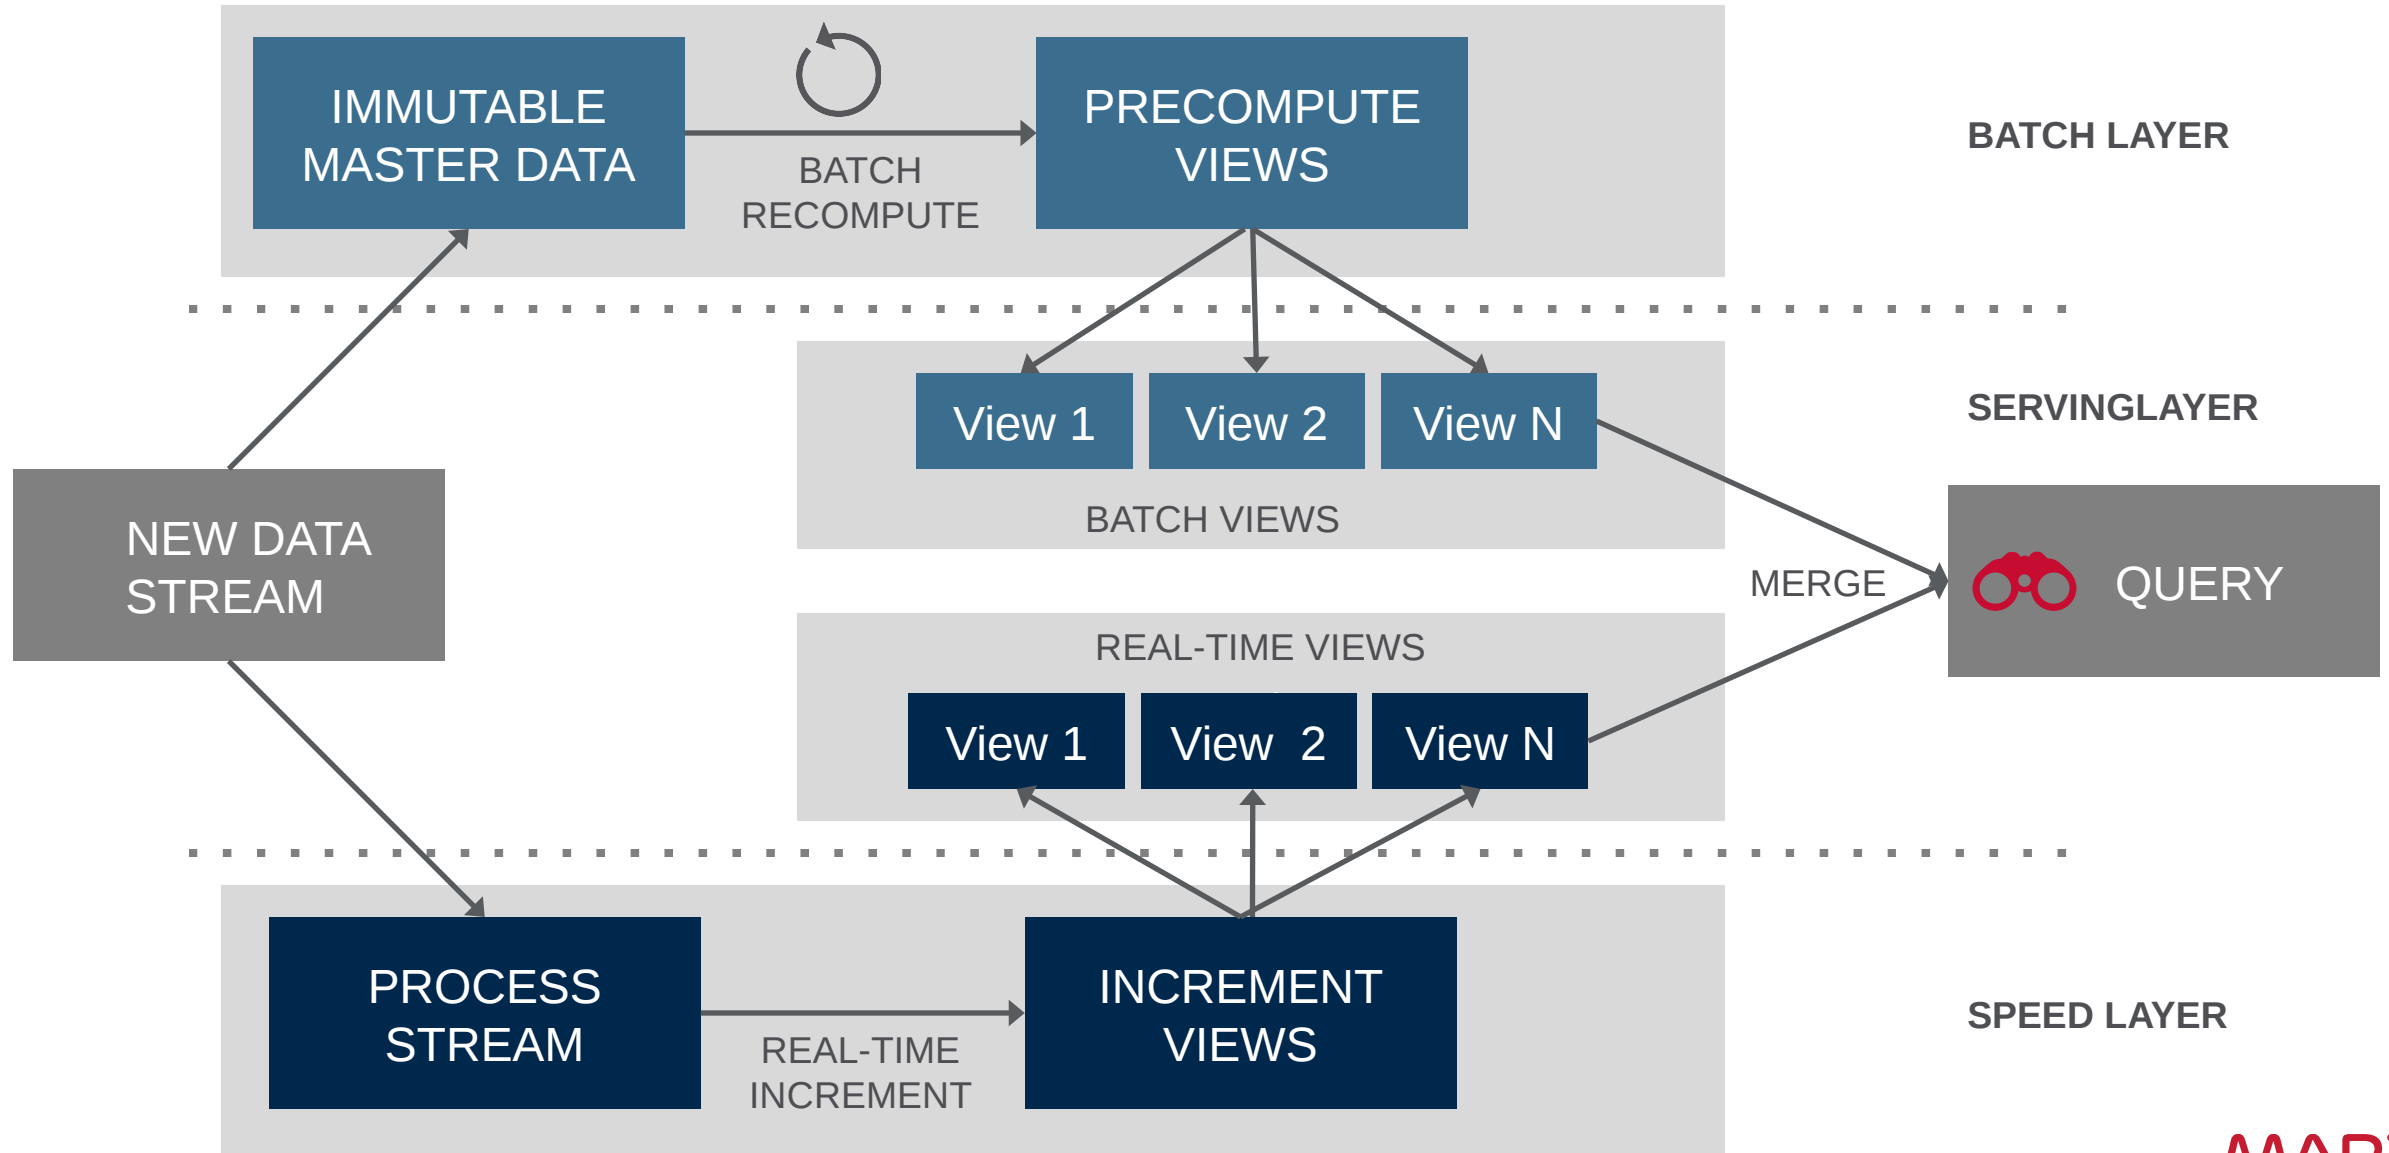


<http://manning.com/marz/>

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

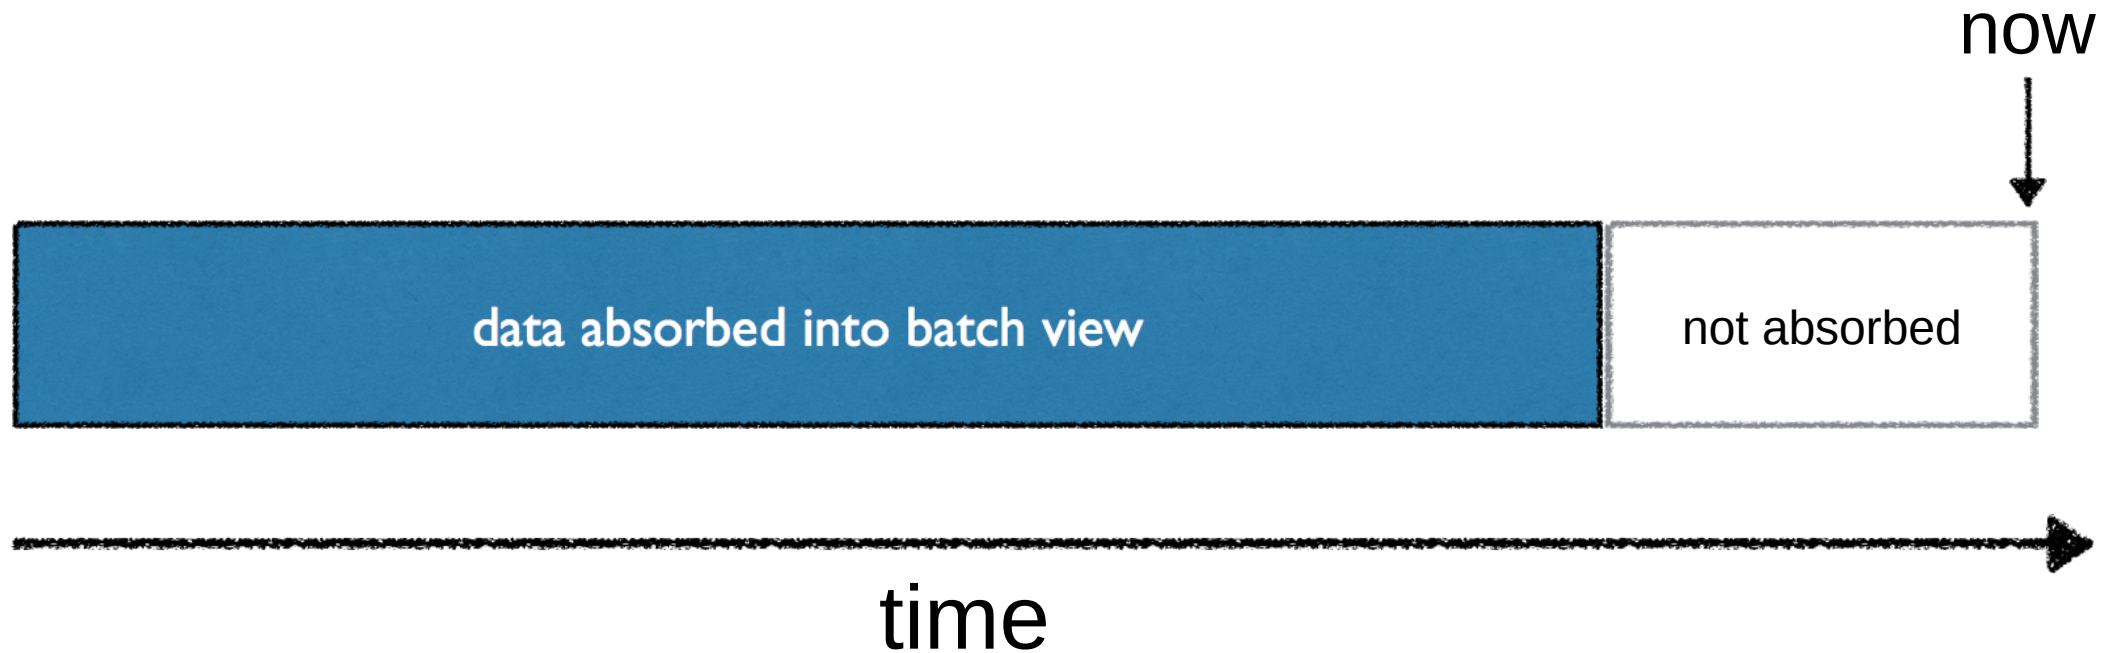
Lambda Architecture



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



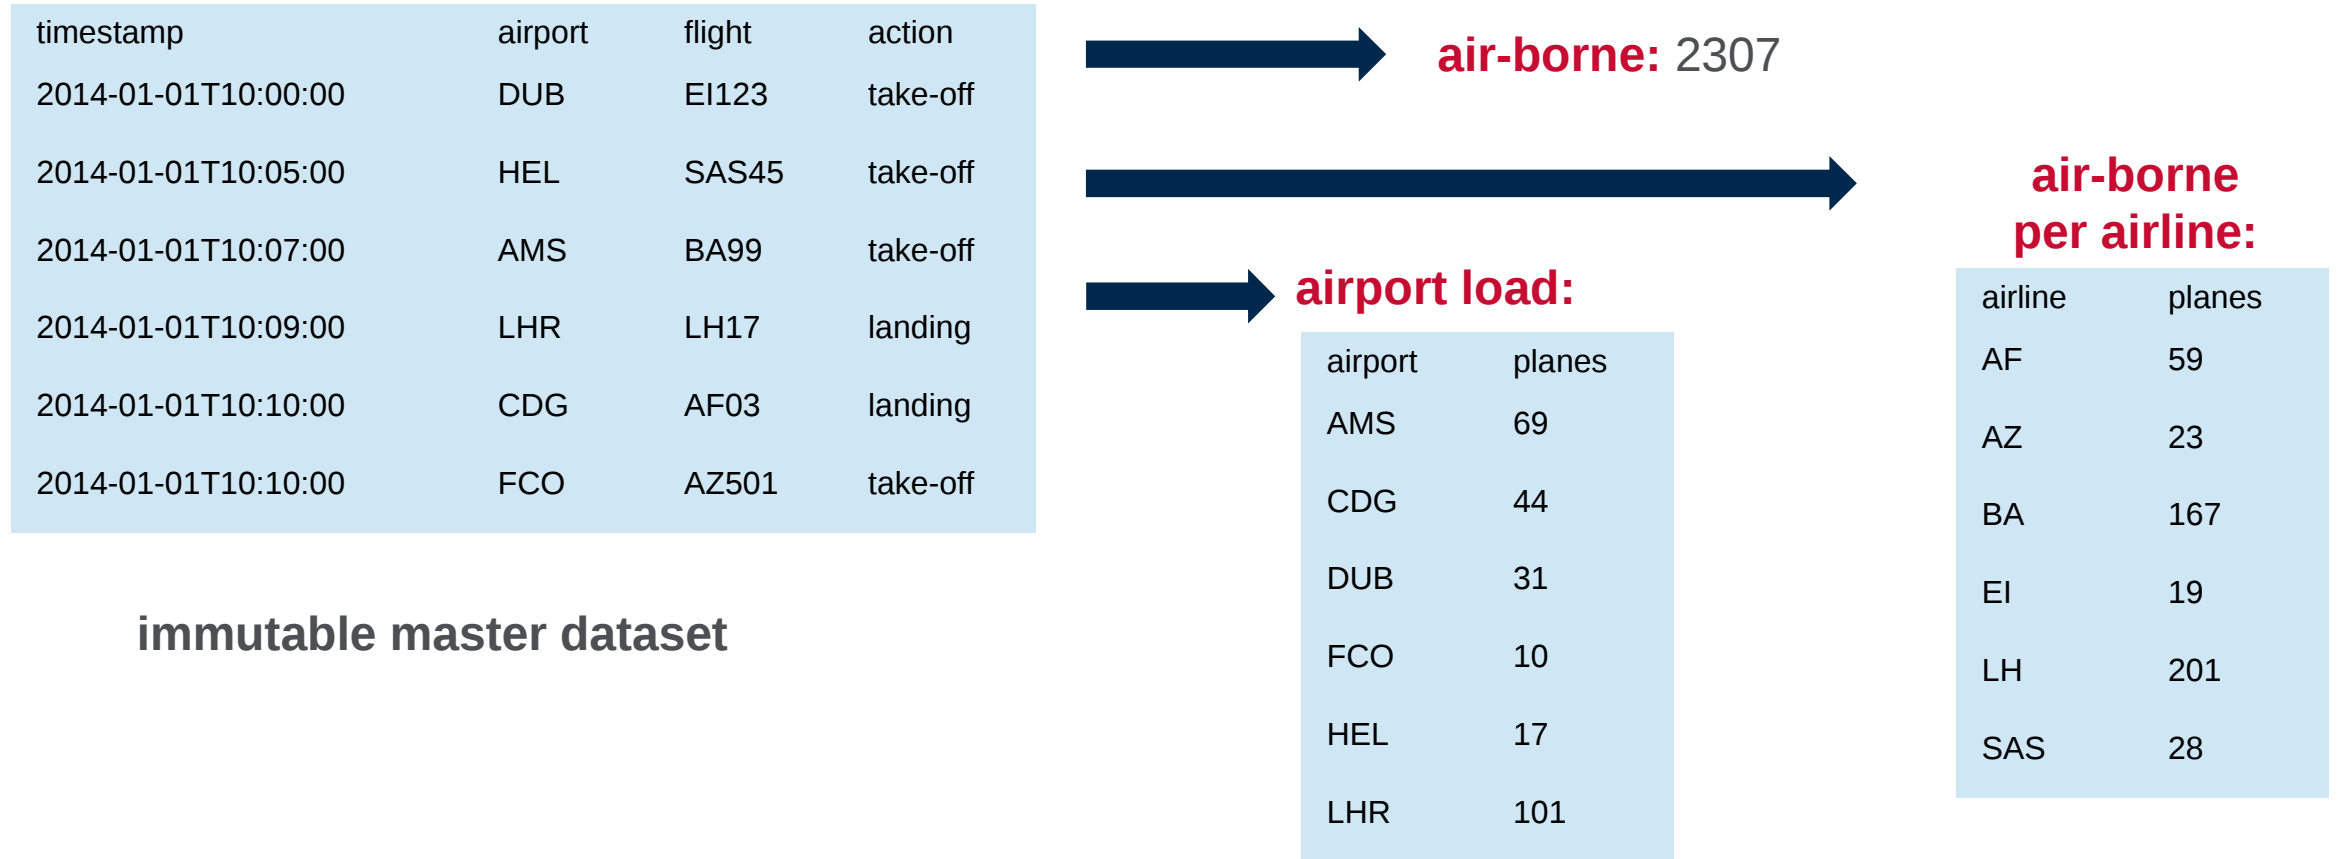
<http://openflights.org>

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



views

Lambda Architecture

A repository dedicated to the Lambda Architecture (LA). We collect and publish examples and good practices around the LA.

Updates

- 20 Jan 2014 » [Lambda Architecture: A state-of-the-art](#) by [Pere Ferrera](#)
- 19 Jan 2014 » [An example Lambda Architecture for real-time analysis of hashtags using Trident, Hadoop and Splout SQL](#) by [Pere Ferrera](#)
- 25 Dec 2013 » [Twitter Summingbird](#) by [Michael Hausenblas](#)
- 25 Dec 2013 » [Lambdoop](#) by [Michael Hausenblas](#)
- 25 Dec 2013 » [Issues in Combined Static and Dynamic Data Management](#) by [Michael Hausenblas](#)
- 24 Dec 2013 » [Where Polyglot Persistence meets the Lambda Architecture](#) by [Michael Hausenblas](#)
- 24 Dec 2013 » [Speed Components](#) by [Michael Hausenblas](#)
- 12 Dec 2013 » [Batch Components](#) by [Nathan Bijns](#)
- 11 Dec 2013 » [A real-time architecture using Hadoop and Storm](#) by [Nathan Bijns](#)
- 10 Dec 2013 » [Why are we doing this and why are we doing this now?](#) by [Michael Hausenblas](#)

What is the Lambda Architecture?

[Nathan Marz](#) came up with the term Lambda Architecture (LA) for a generic, scalable and fault-tolerant data processing architecture, based on his experience working on distributed data processing systems at Backtype and Twitter.

The LA aims to satisfy the needs for a robust system that is fault-tolerant, both against hardware failures and human mistakes, being able to serve a wide range of workloads and use cases, and in which low-latency reads and updates are required. The resulting system should be linearly scalable, and it should scale out rather than up.

Tweets



Michael Hausenblas @mhausenblas 8 May
I've made the [@ApacheSpark](#) slide deck from today's #DSDay available online now:

speakerdeck.com/mhausenblas/ap...

[#bigdata](#) [@LambdaArch](#)

Show Summary



Nathan Bijns @nathan_gs 7 May

RT [@prokopp](#): "Lambda Architecture in a nutshell. How to achieve Velocity and Volume with Big Data. j.mp/QdWO2B" [#LambdaArch](#)

Show Summary



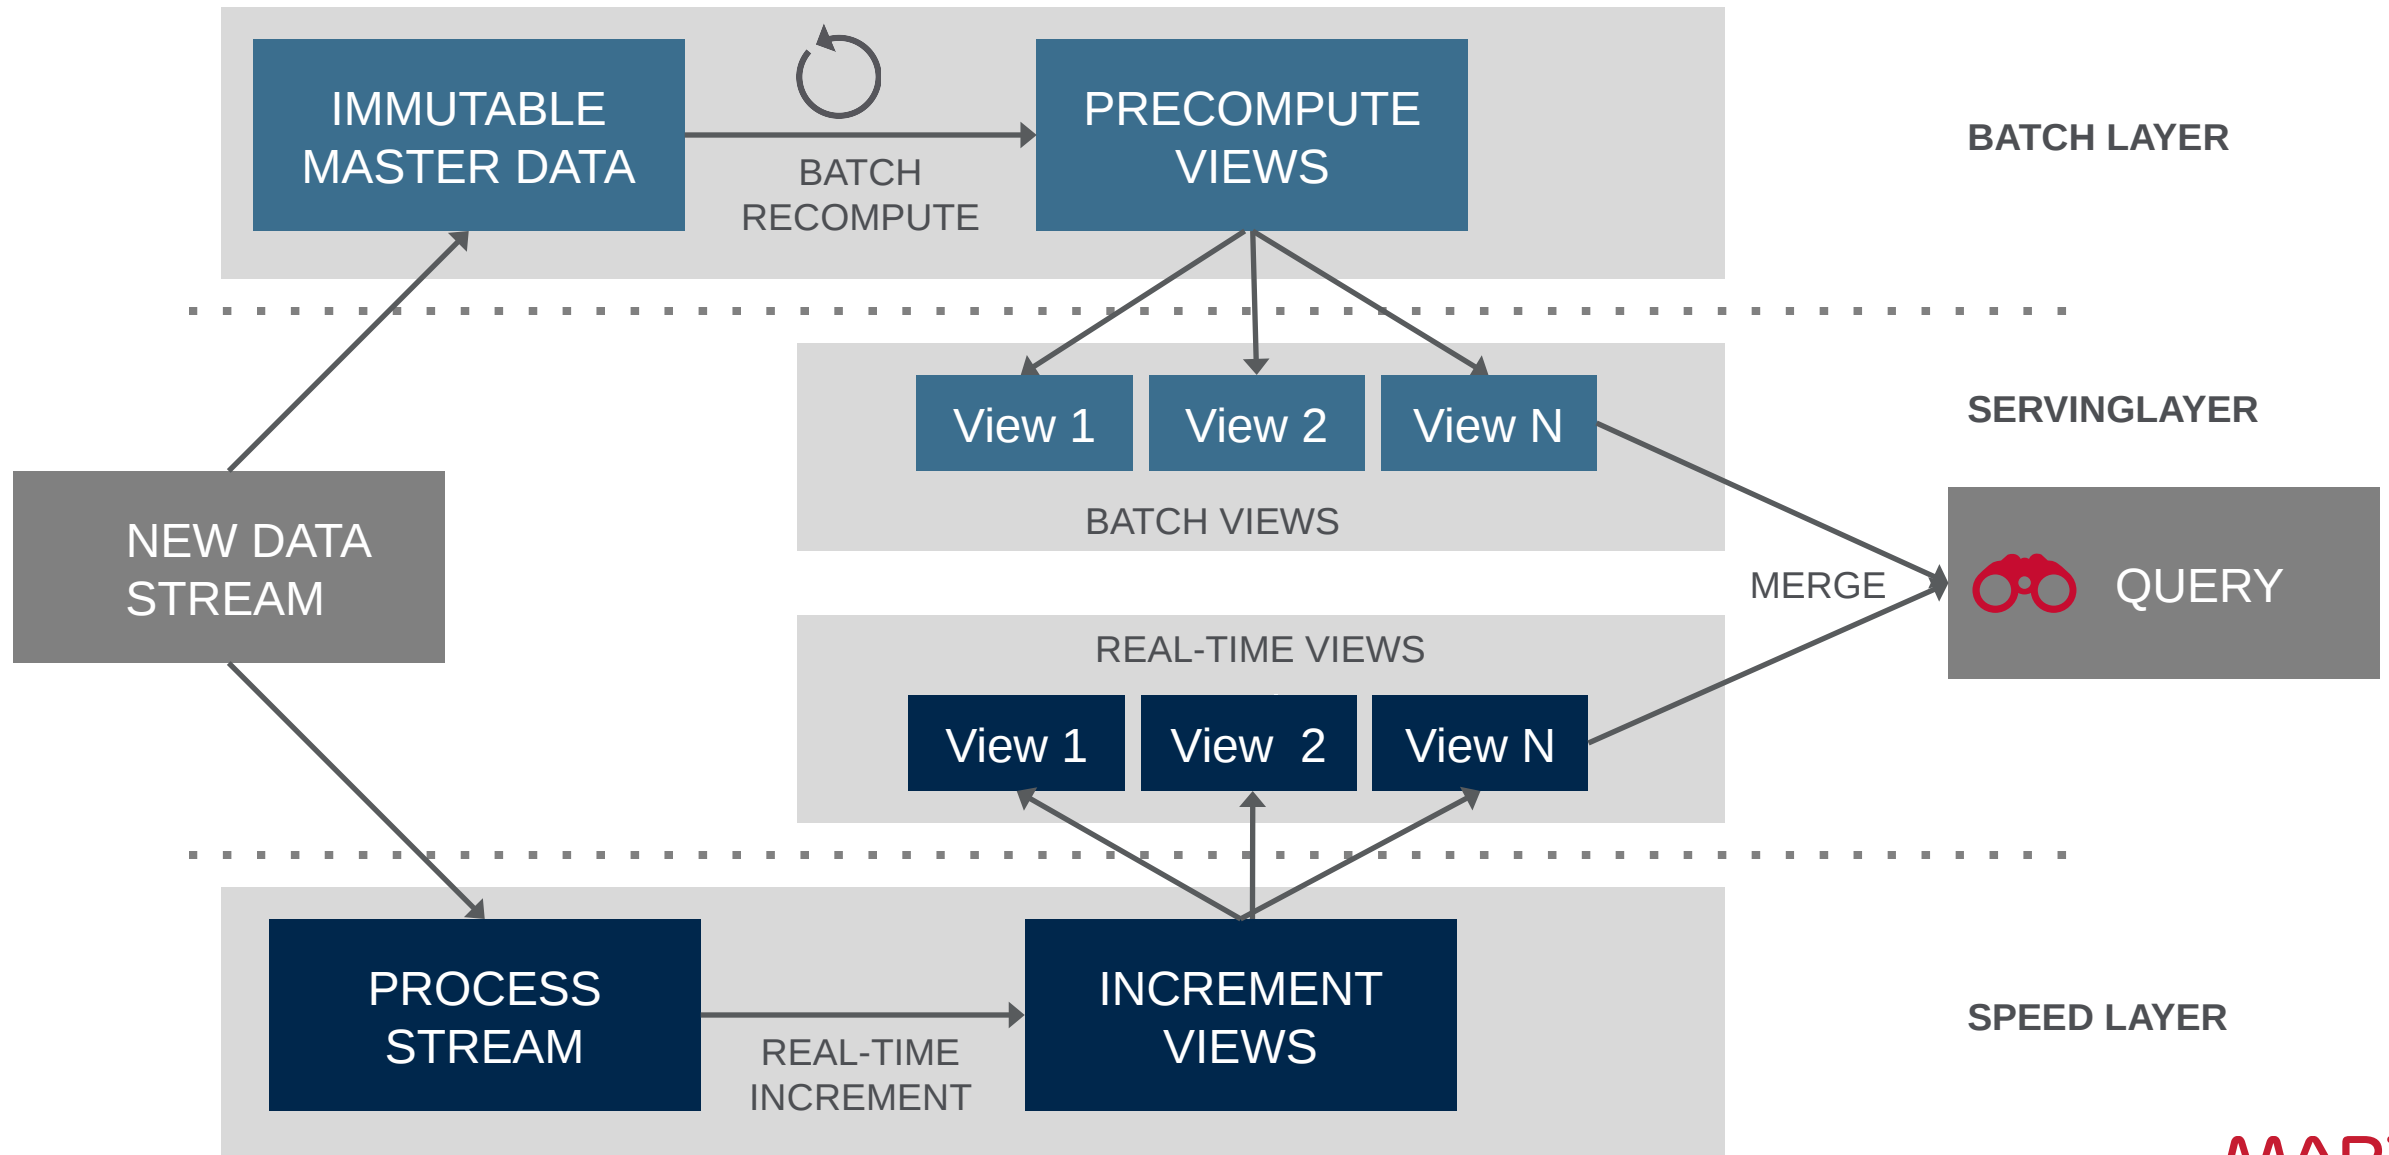
Nathan Bijns @nathan_gs 29 Apr

My slides for A real-time ([#LambdaArch](#)) architecture using [#Hadoop](#) & [#Storm](#), presented during [#nosql14](#) are online: slideshare.net/nathan_gs/a-re...

Show Media

Implementing the Lambda Architecture

Lambda Architecture



Components

<http://lambda-architecture.net/components/2013-12-12-batch-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	★★★	★	★			

Speed Components

24 December 2013

Batch View Databases

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

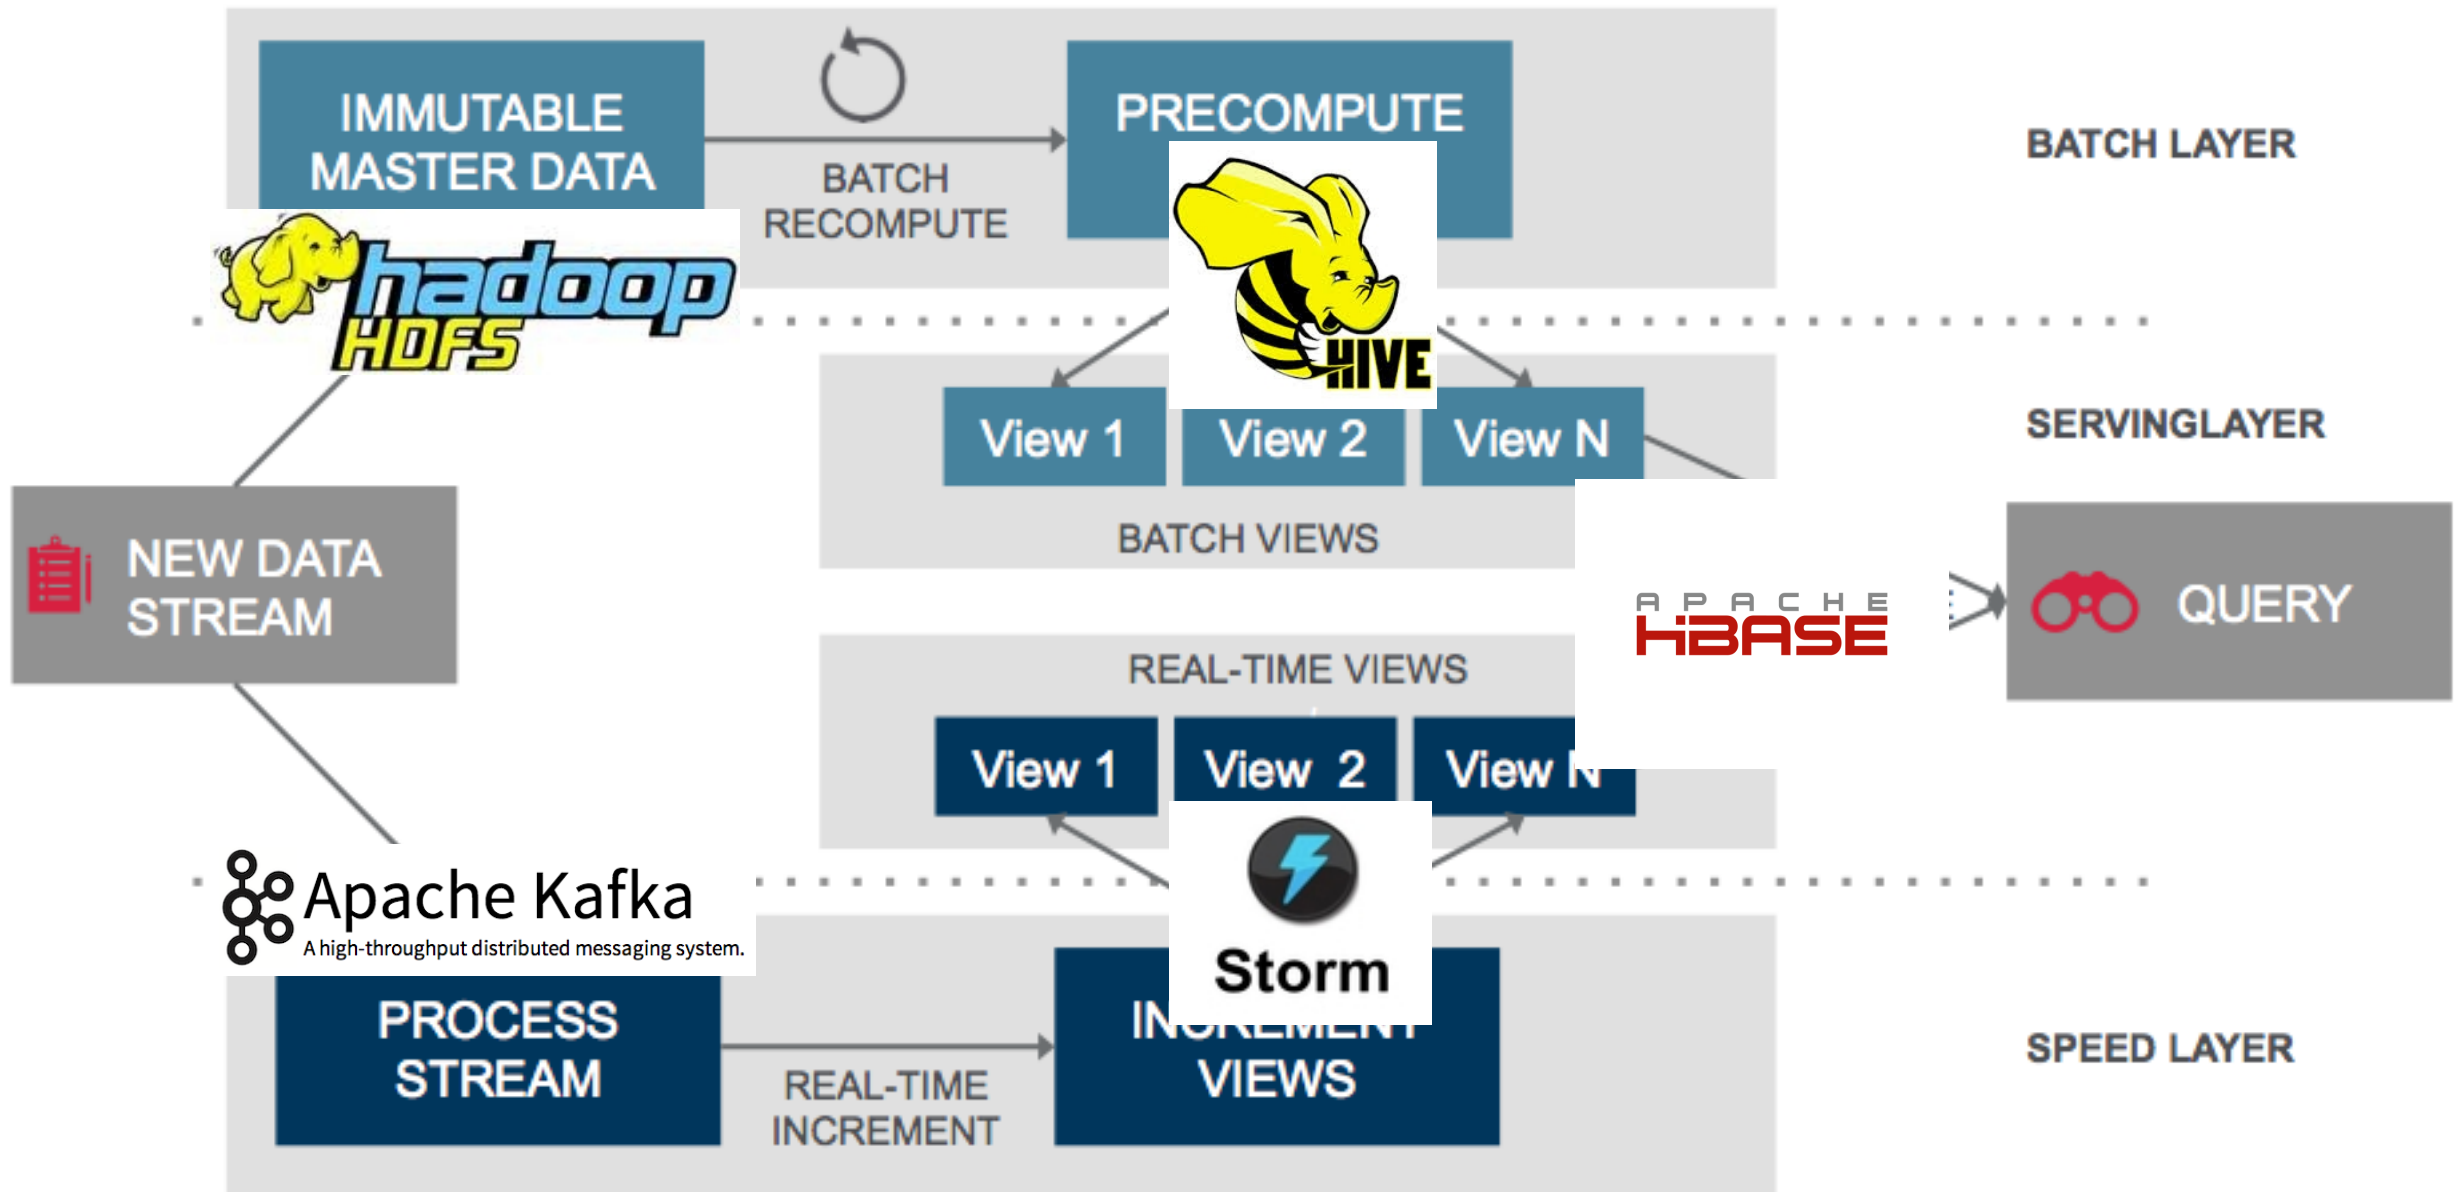
Stream Processing Frameworks

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

<http://lambda-architecture.net/components/2013-12-24-speed-components/>



How about an integrated approach?

- [Twitter Summingbird](#)
- [Lambdoop](#)
- Apache Spark



Apache Spark 101

- Originally developed in 2009 in UC Berkeley's AMP Lab
- A top-level Apache project as of 2014
- Databricks are commercial shepherds
- Enterprise support from Hadoop distributions

The screenshot shows the Apache Spark website homepage. At the top is the Spark logo with the tagline "Lightning-fast cluster computing". Below the logo is a navigation bar with links for "Download", "Related Projects", "Documentation", "Community", and "FAQ". A central banner states: "Apache Spark™ is a fast and general engine for large-scale data processing." To the right, a "Latest News" section lists recent updates: "Spark 0.9.1 released (Apr 09, 2014)", "Submissions and registration open for Spark Summit 2014 (Mar 20, 2014)", "Spark becomes top-level Apache project (Feb 27, 2014)", and "Spark 0.9.0 released (Feb 02, 2014)". Below this is a green "Download Spark" button. On the left, a "Speed" section features a bar chart comparing Hadoop and Spark running times for logistic regression. The chart shows Hadoop at 110 seconds and Spark at 0.9 seconds. Below the chart, text states: "Run programs up to 100x faster than Hadoop MapReduce in memory, or 10x faster on disk." and "Spark has an advanced DAG execution engine that supports cyclic data flow and in-memory computing." To the right of the chart is a code snippet for reading a file and performing a flatMap operation.

Speed
Run programs up to 100x faster than Hadoop MapReduce in memory, or 10x faster on disk.
Spark has an advanced DAG execution engine that supports cyclic data flow and in-memory computing.

Engine	Running time (s)
Hadoop	110
Spark	0.9

Logistic regression in Hadoop and Spark

```
file = spark.textFile("hdfs://...")  
file.flatMap(lambda line: line.split())  
    .map(lambda word: (word, 1))
```

Ease of Use
Write applications quickly in Java, Scala or Python.

Latest News
Spark 0.9.1 released (Apr 09, 2014)
Submissions and registration open for Spark Summit 2014 (Mar 20, 2014)
Spark becomes top-level Apache project (Feb 27, 2014)
Spark 0.9.0 released (Feb 02, 2014)
[Archive](#)

Download Spark

Related Projects:


- [Shark \(SQL\)](#)
- [Spark Streaming](#)
- [MLlib \(machine learning\)](#)
- [GraphX \(graph\)](#)

<https://spark.apache.org/>

The Spark Community

March 27th 2010 - February 15th 2014
Commits to master, excluding merge commits

Contribution Type: **Commits** ▾



AdMobius

YAHOO!

guavus
Go Decisively

amplab
UC BERKELEY

DATABRICKS

TAGGED

阿里巴巴
Alibaba.com

amazon
web services

intel

wandisco

CLOUD
PHYSICS

CONVIVA

Adobe

MAPR

IBM

Hortonworks

4B

webtrends

stratio

ClearStory
DATA
Now You See It

cloudera

ADATAO
BIG DATA INTELLIGENCE

sharethrough

HHMI
janelia farm
research campus

OOYALA

quantiFind

bizo

tuplejump

MAPR

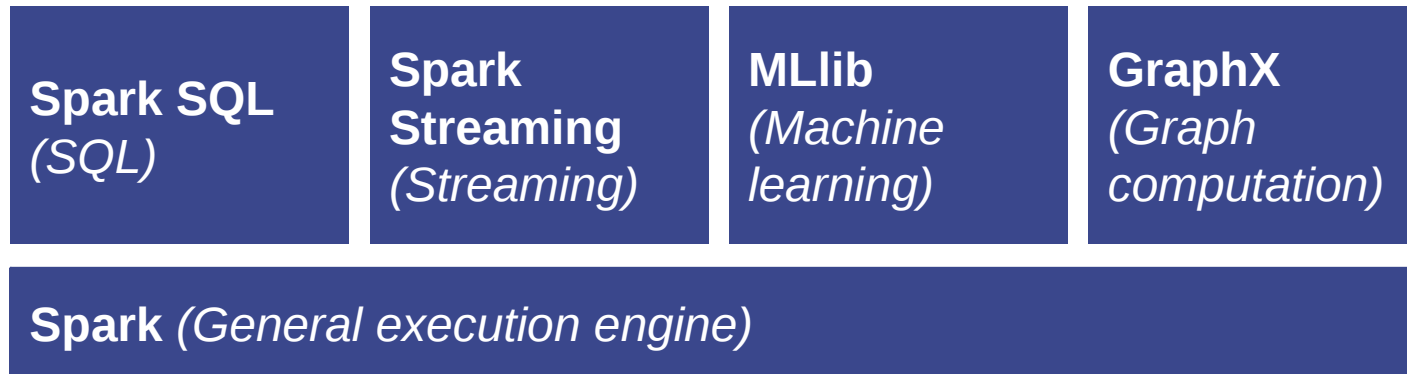
The Spark Stack from 100,000 ft

Spark ecosystem

Spark core engine

ment

A Unified Platform ...



Continued innovation bringing new functionality, e.g.:

- **BlinkDB** (Approximate Queries)
- **SparkR** (R wrapper for Spark)
- **Tachyon** (off-heap RDD caching)

Get Started Immediately

- Multi-language support
- Interactive Shell
- Cloud-native

Python

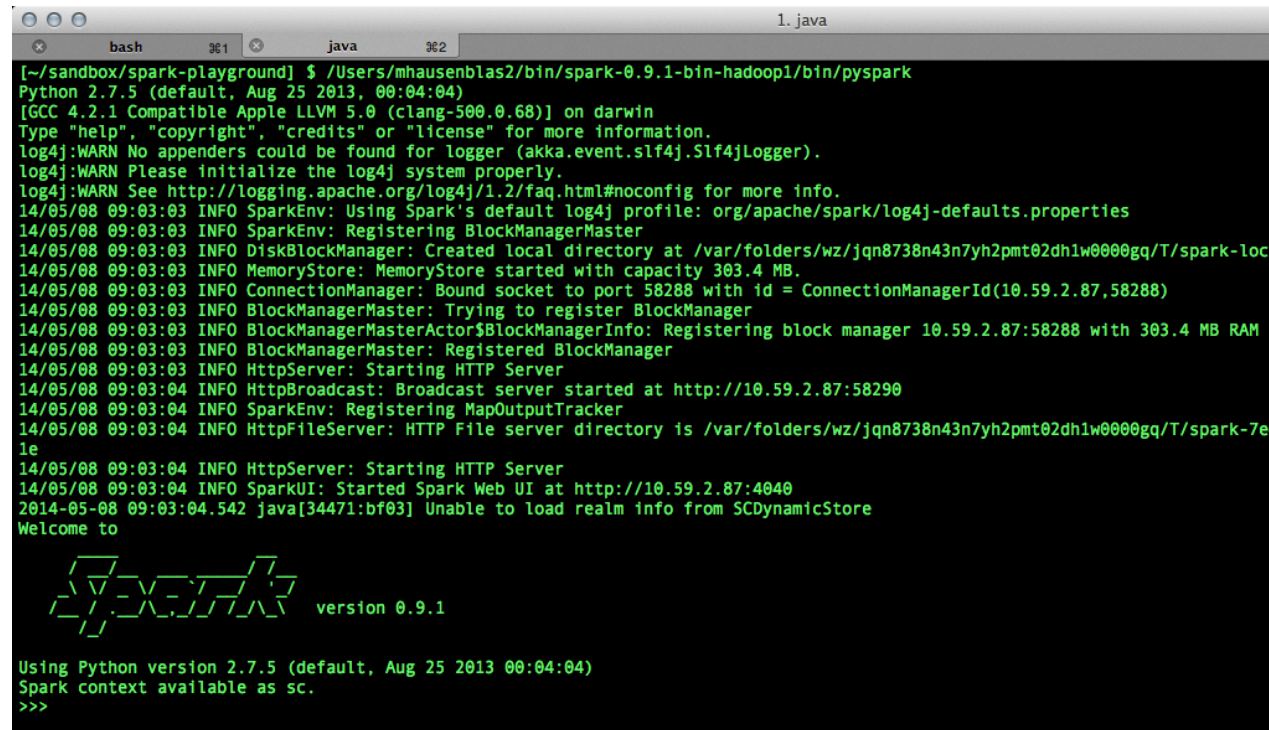
```
lines = sc.textFile(...)
lines.filter(lambda s: "ERROR" in s).count()
```

Scala

```
val lines = sc.textFile(...)
lines.filter(x => x.contains("ERROR")).count()
```

Java

```
JavaRDD<String> lines = sc.textFile(...);
lines.filter(new Function<String, Boolean> () {
    Boolean call(String s) {
        return s.contains("error");
    }
}).count();
```



```
bash 361 java 362 1. java
[~/sandbox/spark-playground] $ /Users/mhausenblas2/bin/spark-0.9.1-bin-hadoop1/bin/pyspark
Python 2.7.5 (default, Aug 25 2013, 00:04:04)
[GCC 4.2.1 Compatible Apple LLVM 5.0 (clang-500.0.68)] on darwin
Type "help", "copyright", "credits" or "license" for more information.
log4j:WARN No appenders could be found for logger (akka.event.slf4j.Slf4jLogger).
log4j:WARN Please initialize the log4j system properly.
log4j:WARN See http://logging.apache.org/log4j/1.2/faq.html#noconfig for more info.
14/05/08 09:03:03 INFO SparkEnv: Using Spark's default log4j profile: org/apache/spark/log4j-defaults.properties
14/05/08 09:03:03 INFO SparkEnv: Registering BlockManagerMaster
14/05/08 09:03:03 INFO DiskBlockManager: Created local directory at /var/folders/wz/jqn8738n43n7yh2pmt02dh1w000ggq/T/spark-loc
14/05/08 09:03:03 INFO MemoryStore: MemoryStore started with capacity 303.4 MB.
14/05/08 09:03:03 INFO ConnectionManager: Bound socket to port 58288 with id = ConnectionManagerId(10.59.2.87,58288)
14/05/08 09:03:03 INFO BlockManagerMaster: Trying to register BlockManager
14/05/08 09:03:03 INFO BlockManagerMasterActor$BlockManagerInfo: Registering block manager 10.59.2.87:58288 with 303.4 MB RAM
14/05/08 09:03:03 INFO BlockManagerMaster: Registered BlockManager
14/05/08 09:03:03 INFO HttpServer: Starting HTTP Server
14/05/08 09:03:04 INFO HttpBroadcast: Broadcast server started at http://10.59.2.87:58290
14/05/08 09:03:04 INFO SparkEnv: Registering MapOutputTracker
14/05/08 09:03:04 INFO HttpFileServer: HTTP File server directory is /var/folders/wz/jqn8738n43n7yh2pmt02dh1w000ggq/T/spark-7e
1e
14/05/08 09:03:04 INFO HttpServer: Starting HTTP Server
14/05/08 09:03:04 INFO SparkUI: Started Spark Web UI at http://10.59.2.87:4040
2014-05-08 09:03:04.542 java[34471:bf03] Unable to load realm info from SCDynamicStore
Welcome to

DEXTRA version 0.9.1

Using Python version 2.7.5 (default, Aug 25 2013 00:04:04)
Spark context available as sc.
>>>
```

Expressive API

map

reduce

Expressive API

map
filter
groupBy
sort
union
join
leftOuterJoin
rightOuterJoin

reduce
count
fold
reduceByKey
groupByKey
cogroup
cross
zip

sample
take
first
partitionBy
mapWith
pipe
save ...



The Apache Spark Stack

News, tutorials, applications and more around Apache Spark—lightning-fast cluster computing.

Home

Getting started

Feedback

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_EMEA



maprtech

MapR Technologies



+MaprTechnologies

mhausenblas@mapr.com



maprtech