

A P A C H E  
**HBASE**

**All that is new...**

**Michael Stack**  
***Berlin Buzzwords 2013***

**Ingredients: Overview**

**(30%), Current (10%),**

**Next (40%), Future (20%)**

**Me (100%)**

**"...scalable,  
distributed  
datastore."**


**Who uses it?**



*Pinterest*

ebay®

facebook.



Dropbox



StumbleUpon

OP@WER



**How is it  
used?**

# *How is it used?*


- **Multifarious, multitenant**
  - *a.k.a dumping ground*
- **Reconciliation Store**
- **Time Series/Metrics**
- **Lots-o-entities Store**

*How is it used?*

**OLAP & OLTP**



**Who runs the  
project?**



facebook

cloudera



淘宝网  
Taobao.com

# Currently



**Last one was 0.94.8, this week...**



- **Release every month**
- **Each more stable**
- **& more performant**
- **Some features...**





**> 1000**

**So far...**





# CONSISTENCY

IS THE LAST REFUGE OF THE UNIMAGINATIVE.



OSCAR WILDE





## Client


- Client Request to RS
- ... pause...
- ... pause...
- ... pause...
- ... pause...
- ... pause...
- ... pause...
- ... pause...
- ... pause...
- ... pause...
- ...Response!

## Region Server

- **A RS dies**
- → Log splitting
- → Regions available
- ... idle...
- ... idle...
- → Client connected to new RS
- Request Served...

## Master

- ... idle...
- ... detects Failed RS ...
- ... starts the recovery ...
- ... starts log splitting ...
- ... assign regions ...
- ... replay complete ...





- **Hadoop HDFS Fixes**
- **Faster Recovery**
  - **Detection**
  - **Replay**
  - **Assign**

96



96





protobuf

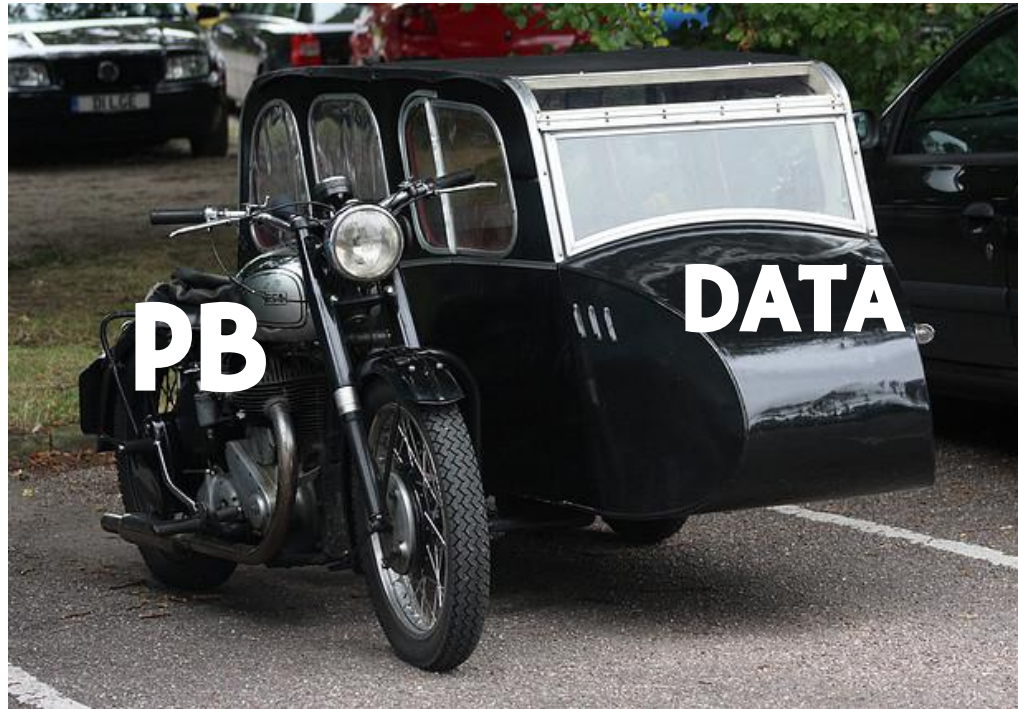
Protocol Buffers - Google's data interchange format

- **System tables**
- **Filesystem**
- **Up in zookeeper**
- **Over the wire**

# RPC

96

- **Implements Protobuf Service**
- **Specification!**
- **Data on the side**
  - *Encoding*
  - *Compression*



# 96 Snapshots

- **By Table**
  - **Snapshot, clone, restore, export**
- **Inexpensive**
  - **Just metadata**
- **Good for...**
  - **Backups**
  - **Replication**
  - **Offline processing**





# Compactions

- **Pluggable**
  - **Tiered**
  - **Striped**
- **Smarter Triggers**





96


# Tests

- **Cluster test module**
  - Standalone or cluster
  - Sizeable
    - x data
    - x runtime
- **Borrows test types/mechanisms**
  - Netflix "ChaosMonkey"
  - Accumulo linked-list dataloss checker

A yellow pill with the number 96 embossed on it, set against a dark background.

# Miscellaneous

- **Smarter load balancer**
- **Revamped**
  - **Metrics**
  - **UI**
  - **etc.**
- **Hadoop 1.x and 2.x**



# HDFS

- **Short Circuit Reads**
- **Checksums**
  - **Inline**
  - **Native**
- **Favored node**



# HBase Ecosystem











### My Energy Use

Show my  by year



■ You  
■ All similar homes  
■ Efficient similar homes  
[What homes are compared?](#)

### Base 0.97.0-SNAPSHOT API

<a href="#">Result</a>	<a href="#">getRowOrBefore</a> (byte[] row, ...)	Return the row that matches
<a href="#">ResultScanner</a>	<a href="#">getScanner</a> (byte[] family)	Gets a scanner on the current
<a href="#">ResultScanner</a>	<a href="#">getScanner</a> (byte[] family, ...)	Gets a scanner on the current
<a href="#">Scanner</a>	<a href="#">getScanner</a> (Scan scan)	Returns a scanner on the current



# kiji.org

- **Entity-centric, simple model**
- **AVRO**
  - **Types & compound types**
- **Each cell is schema versioned**
  - **Schema migration**
- **Works across MR & REST, etc.**
- **Open-source**
- **Production deploys**



S

Q

L





# SQL





# SQL

# Phoenix



- "A SQL skin over HBase"
- Coprocessors, custom filters, jdbc driver
- <https://github.com/forcedotcom/phoenix>

next

next

**1.0?**

next

**Faster scans  
(OLAP)**

# next

## **Latency Resilience/"Latency Tolerance"\***

- **Bring home the outliers**



***Related: QoS***

# next

## **Flexibility-Usability trade-off**

- **Ergonomics**
- **Forgiveness**

**ORDER  
& PAY**



**TAKE-AWAY**





**Thank You!**

**[stack@apache.org](mailto:stack@apache.org)**