



Large objects in the Cloud

# Riak Cloud Storage

- Cloud Storage software backed by Riak
- Simple API
- Multi-tenant, Per-tenant Reporting
- Pluggable Authentication
- Multi Data Center Replication (Enterprise)
- DTrace Support, Detailed Stats, etc
- Preliminary CloudStack integration

# Simple Storage Service (S3) Protocol

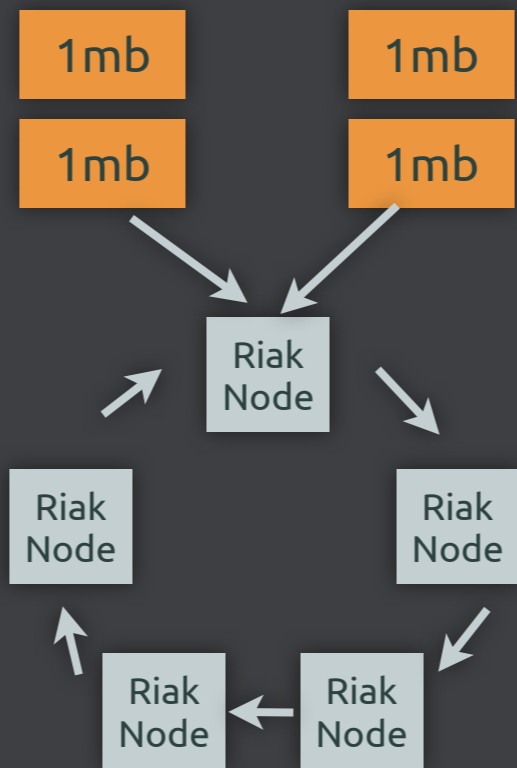
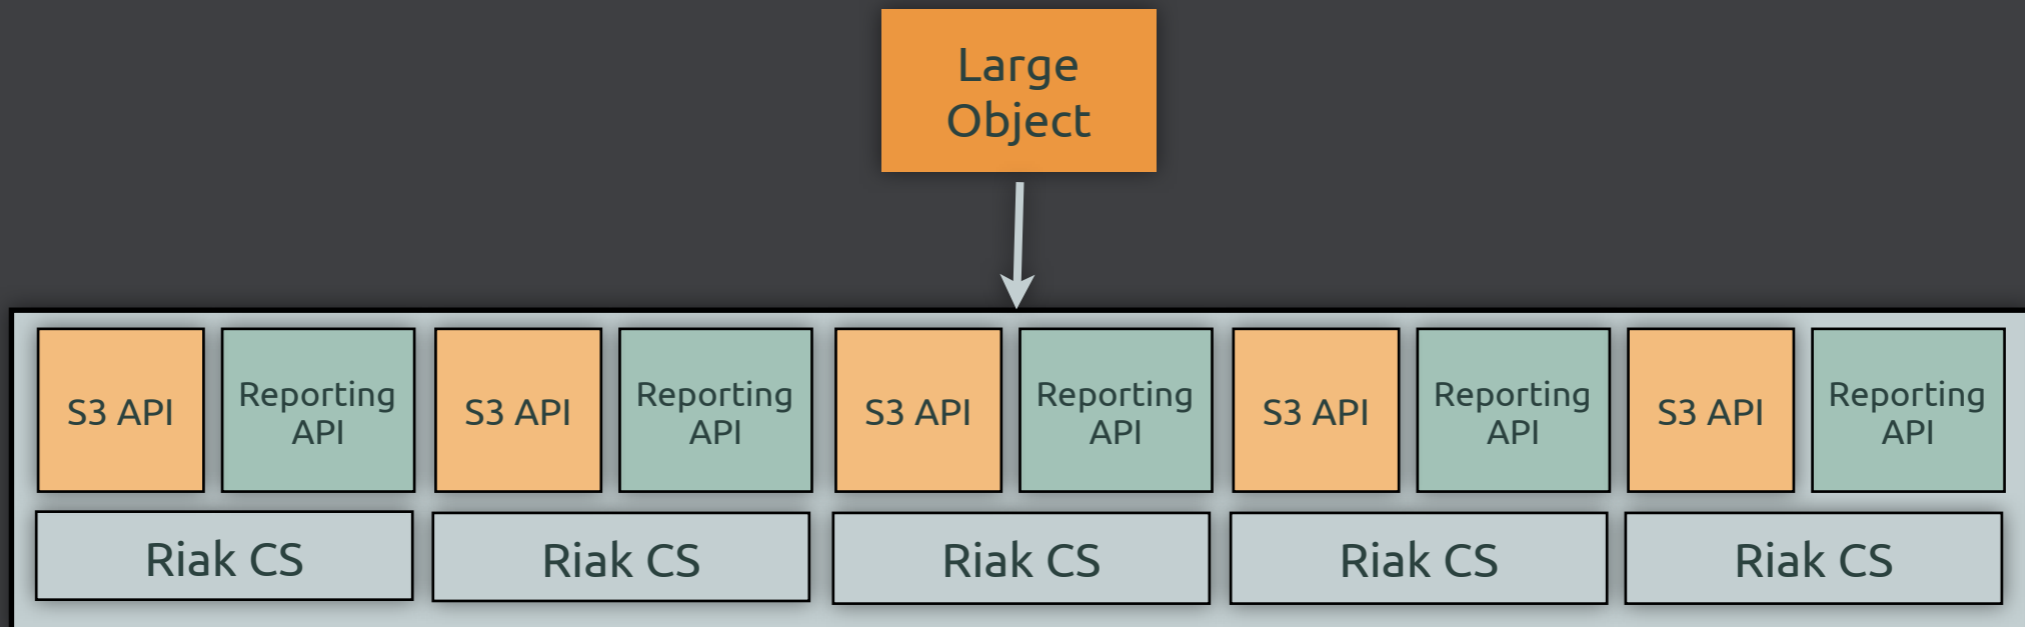
- Straight forward API
  - Make buckets, list buckets, etc
  - GET / PUT / DELETE - operations
- Use any existing Amazon S3 client library ;)

e.g.

```
s3cmd put test-file s3://test-bucket
```

# Riak

- Key-Value Store + Extras
- Distributed, horizontally scalable
- Fault-tolerant
- Highly-available
- Built for the Web
- Inspired by Amazon's Dynamo



# Coming Soon

- Riak CS 1.4
  - Swift API
  - Keystone Integration
  - COPY Object
  - Object Versioning
  - Additional exotic S3 features

**On March 20, 2013**

**Riak CS**

**became open source**



# RICON 2013

*A Distributed Systems Conference for Developers*



## RICON | WEST

*San Francisco*


October ...




## RICON | EAST

*New York City*

May 13th-14th at New World Stages in world-famous Midtown Manhattan.

 Tickets

 More Info



## RICON | EUROPE

*London*

Provisionally  
scheduled for  
November 2013



GOTO Zurich 2013

# APACHE HBASE: INTRODUCTION TO A COLUMN- ORIENTED DATA STORE

Christian Gügi, Lead Architect, Sentric



HBase is an  
open source, distributed,  
column-oriented data store  
modeled after Google's BigTable

- Sorted map data store
- Table consists of rows, each has a row key (primary key)
- Each row may have any number of columns ( $\text{Map}\langle\text{byte}[], \text{byte} []\rangle$ )
- Rows are sorted lexicographically based on row key

Different types of data separated into different "column families"

Data is all byte[]

Row key	Data
amuller	info: { 'height': '2.0m', 'state': 'ZH' } roles: { 'IBM': 'Sales Manager' }
cguegi	info: { 'height': '1.85m', 'state': 'BE' } roles: { 'Sentric': 'Architect'@ts=2011, 'Sentric': 'Mentor'@ts=2012, 'SBDUG': 'Founder' }

Different rows may have different sets of columns (table is sparse)

Single cell may have different values at different timestamps

### info Column Family

Row key	Column key	Timestamp	Value
amuller	info:height	1333883187	2.0m
amuller	info:state	1273871824	ZH
cguegi	info:height	1325755229	1.85m
cguegi	info:state	1325751049	TG

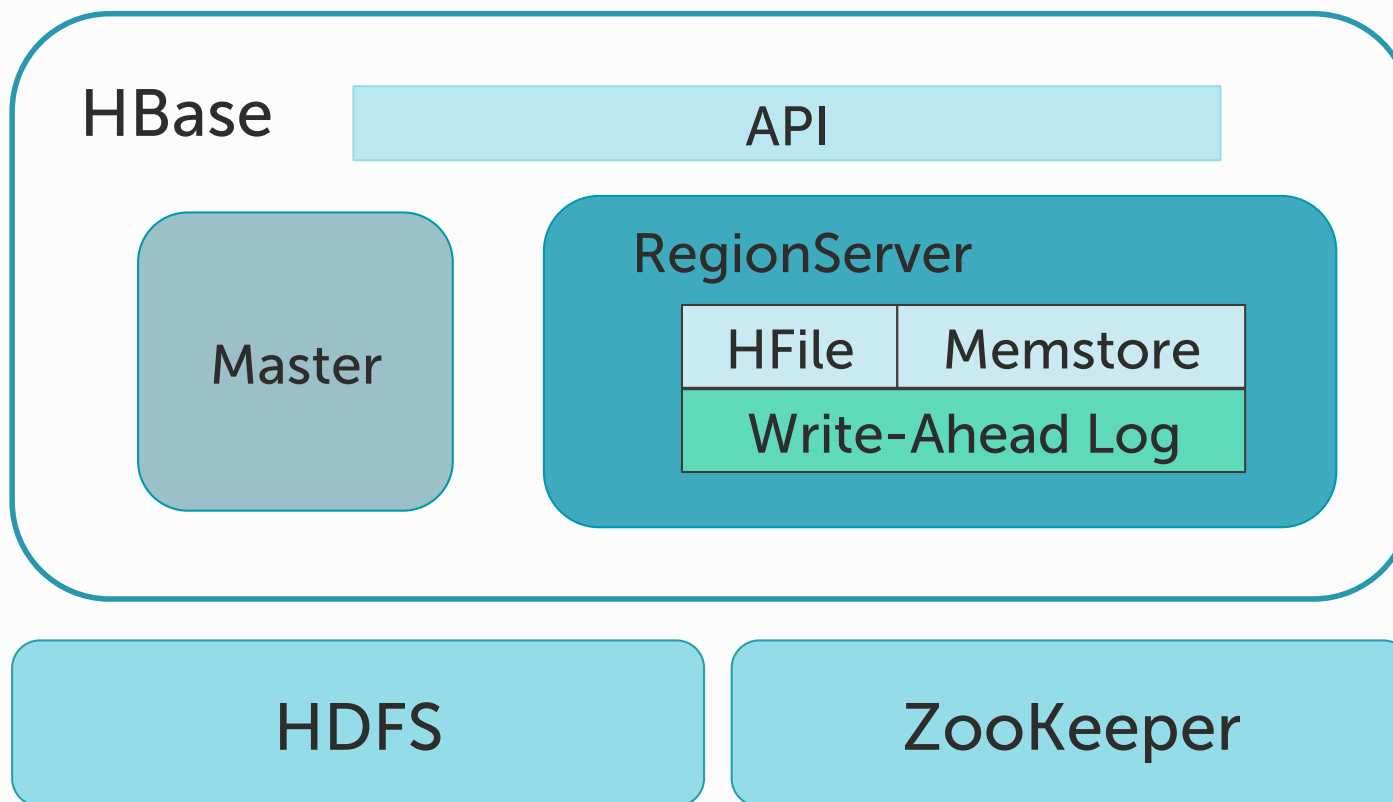
### roles Column Family

Row key	Column key	Timestamp	Value
amuller	roles:IBM	1320105636	Developer
cguegi	roles:SBDUG	1330561785	Founder
cguegi	roles:Sentric	1325376723	Mentor
cguegi	roles:Sentric	1293840959	Architect

Sorted on disk by row key,  
column key, descending ts



Unix timestamp



[HBase: The Definitive Guide]

- Favors Consistency over Availability
- Great Hadoop integration
- Ordered range partitions
- Automatically shards/scales
- Sparse column storage



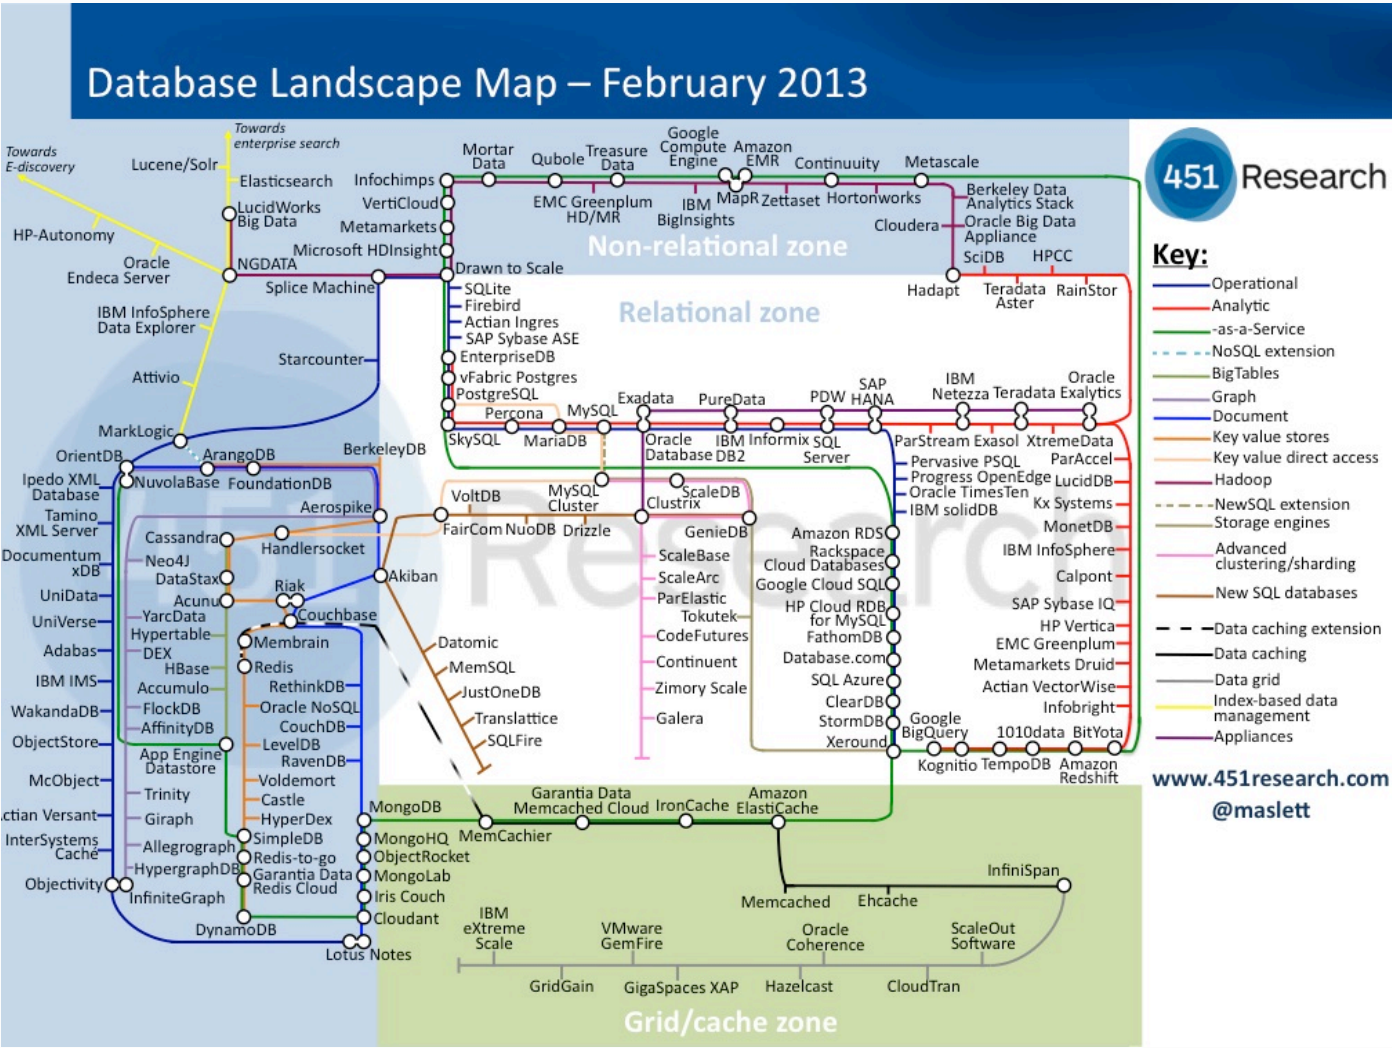
Q&A

THANK YOU!

@chrisgugi



- <http://hbase.apache.org>
- <http://www.sentric.ch>
- <http://bigdata-usergroup.ch>
- <http://about.me/cguegi>



Source: [http://blogs.the451group.com/information\\_management/2013/02/04/updated-database-landscape-map-february-2013/](http://blogs.the451group.com/information_management/2013/02/04/updated-database-landscape-map-february-2013/)