

Cassandra **SF** 2011

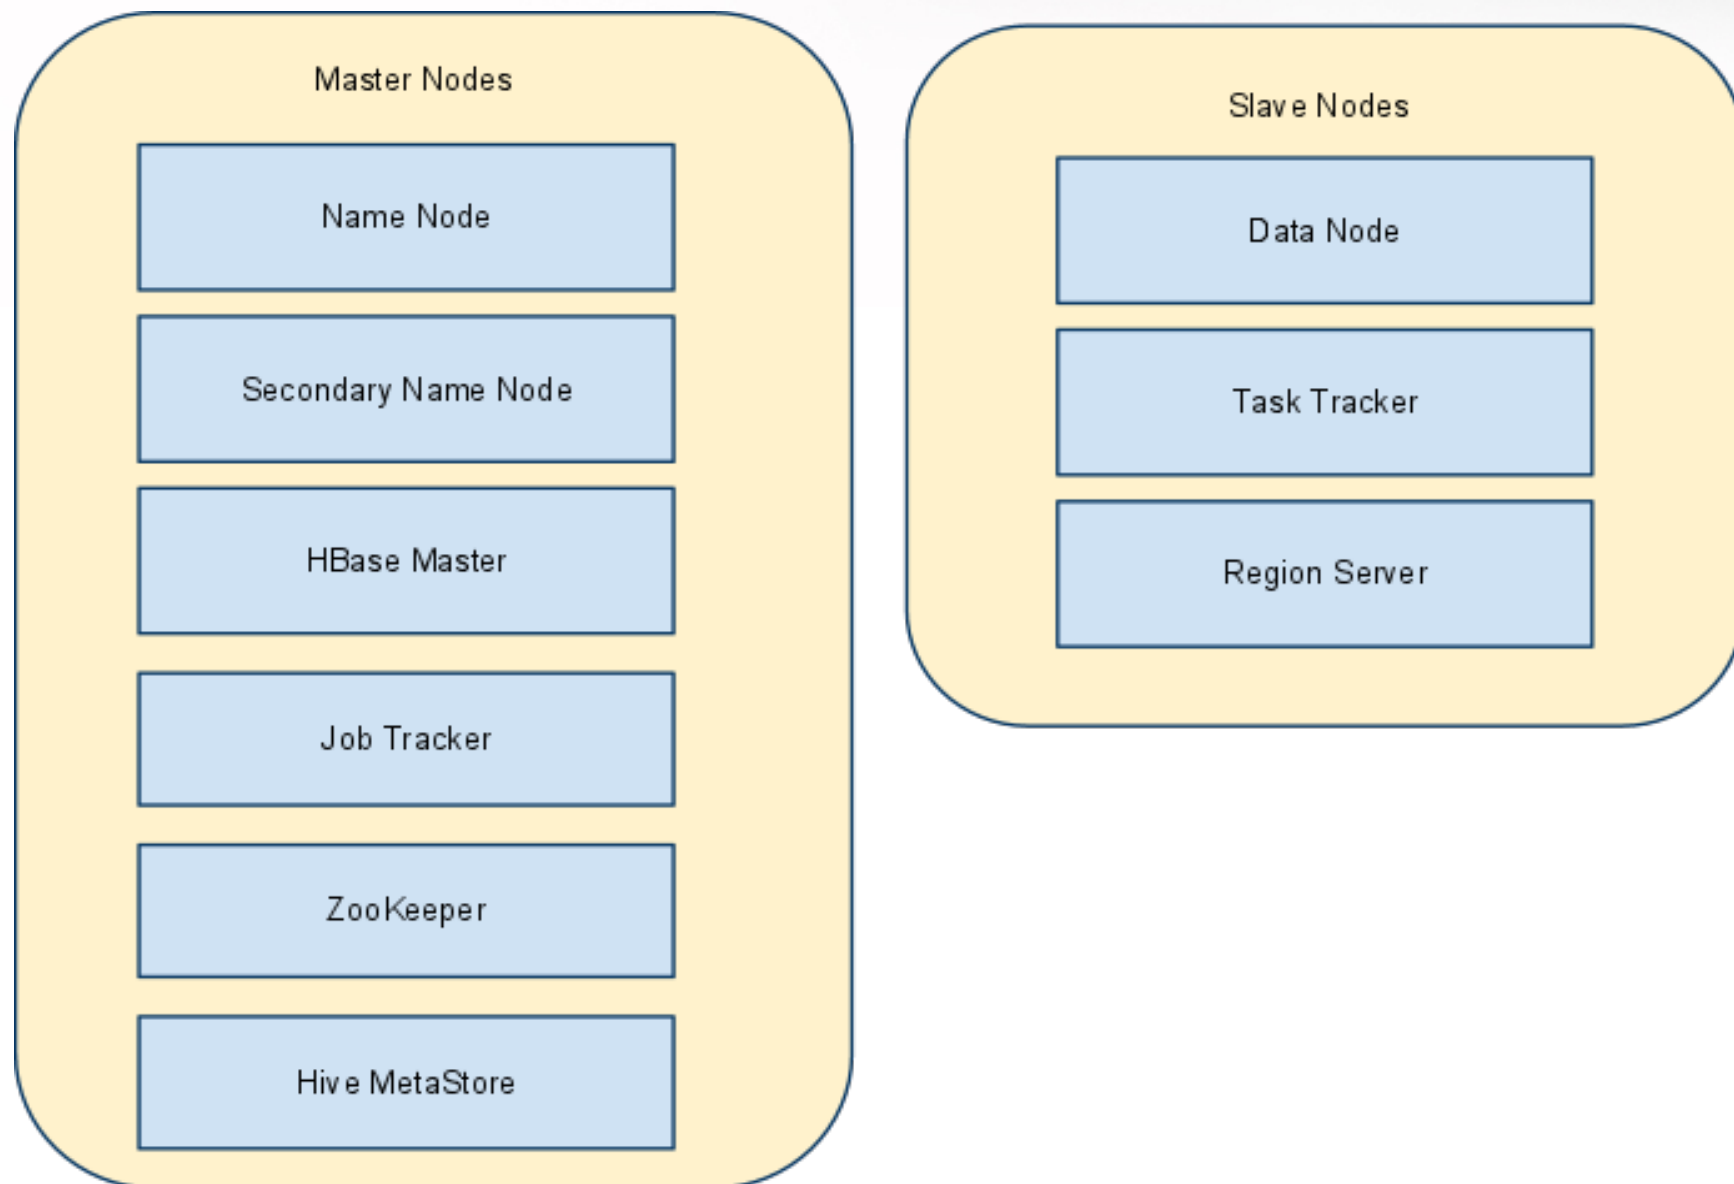
Brisk

Better Hadoop with Cassandra

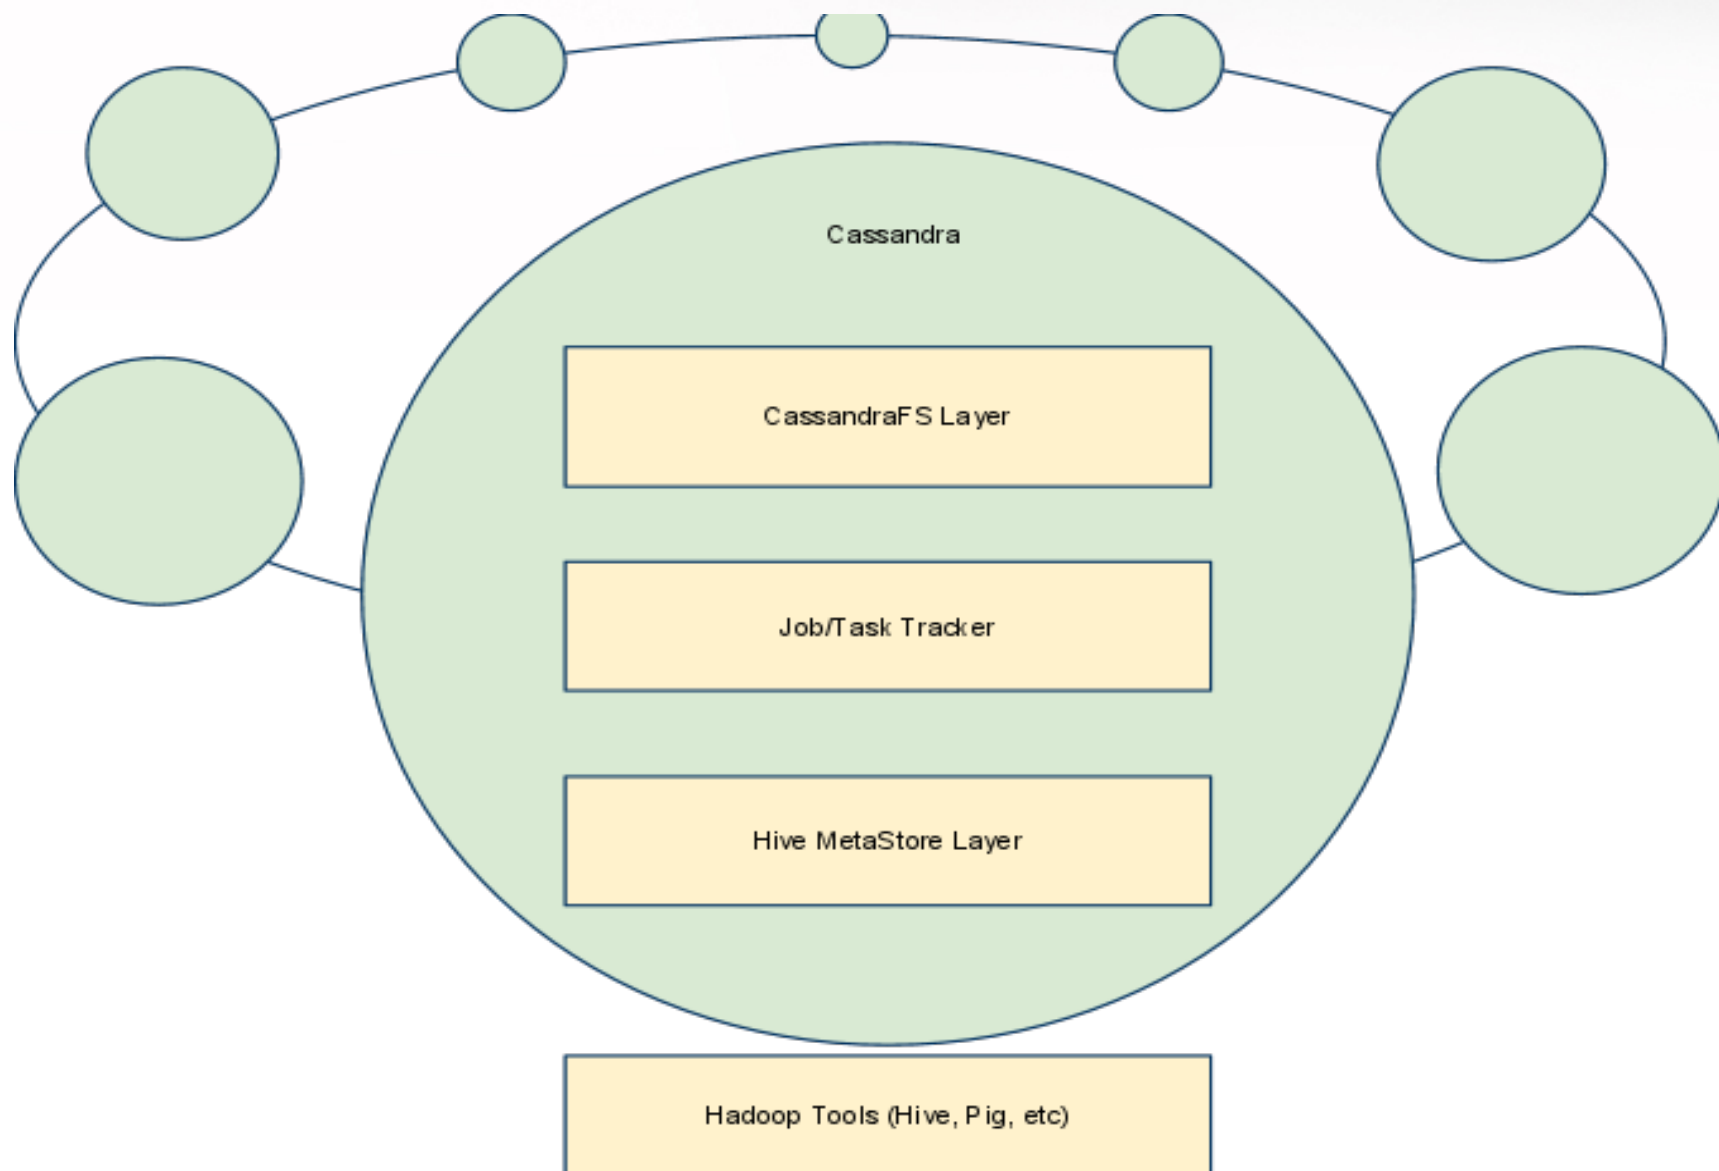
Jake Luciani (@tjake)



The Traditional Hadoop Stack



The Brisk Hadoop Ring

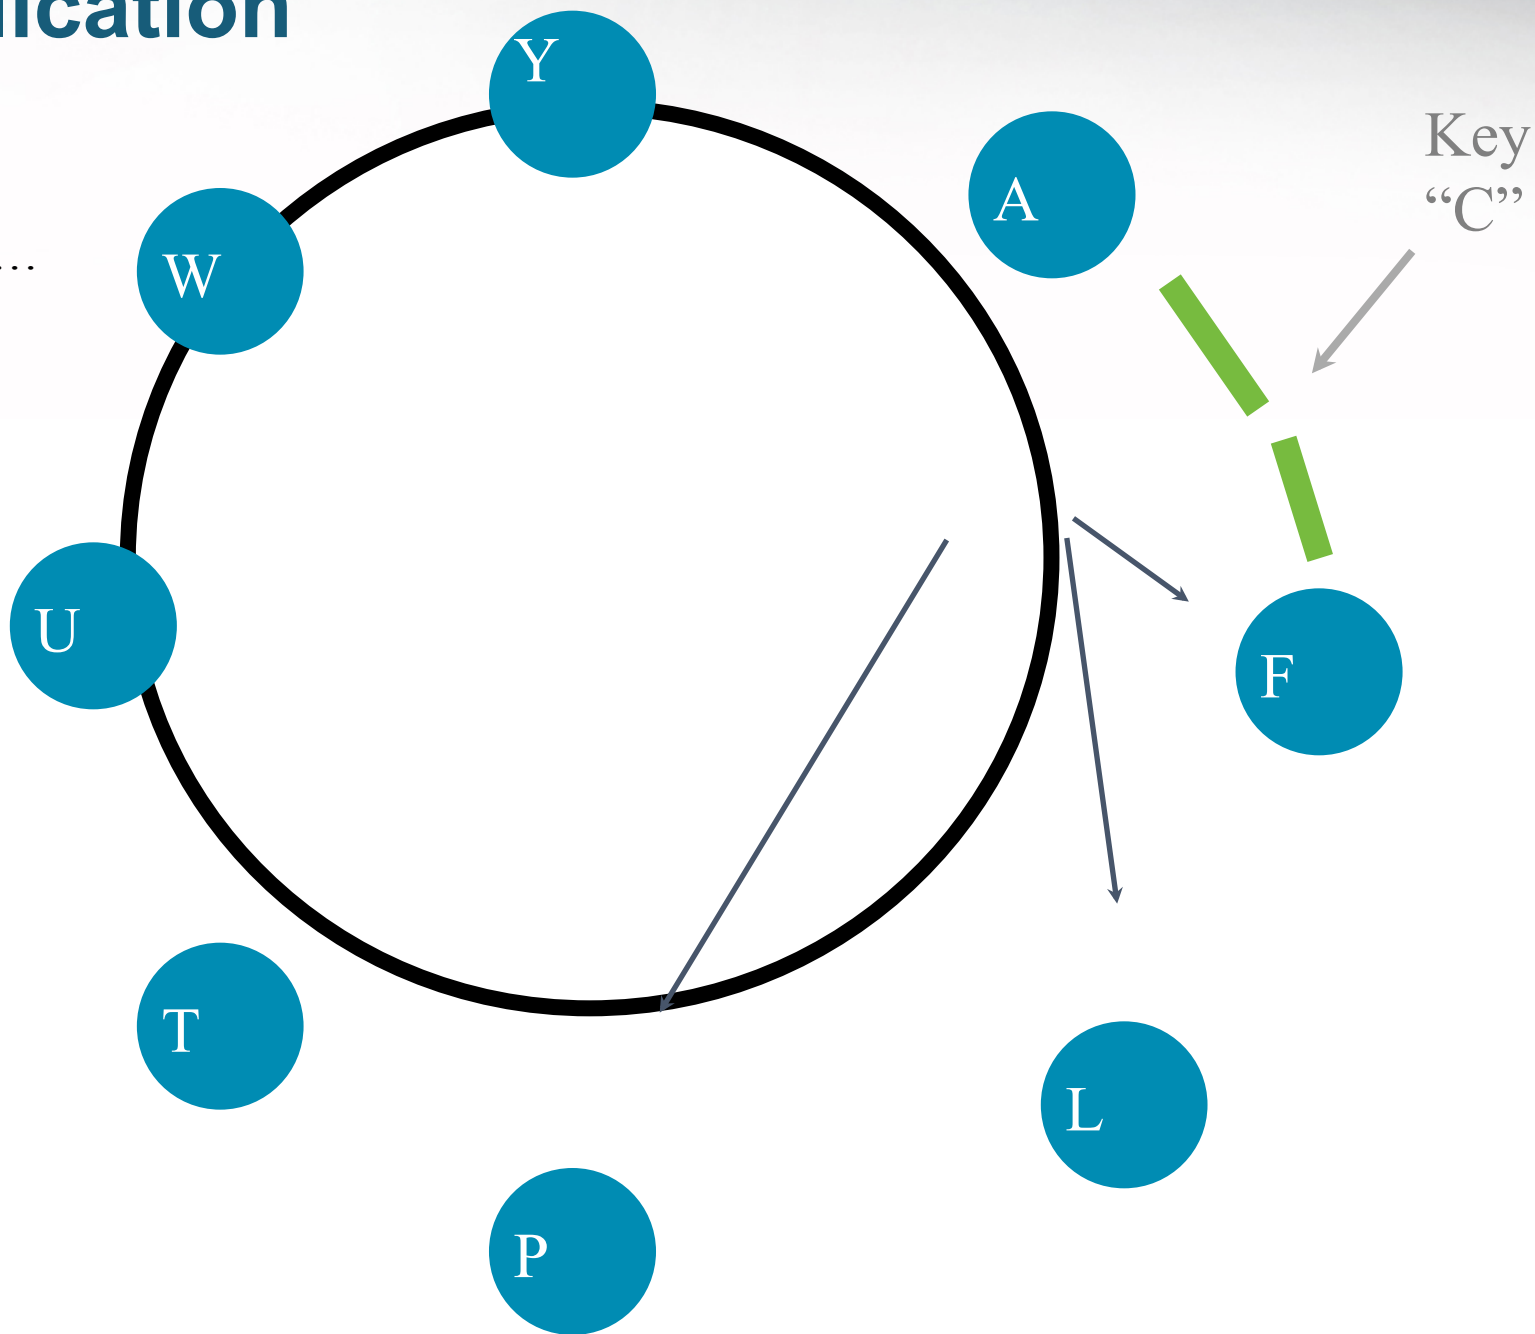


Brisk

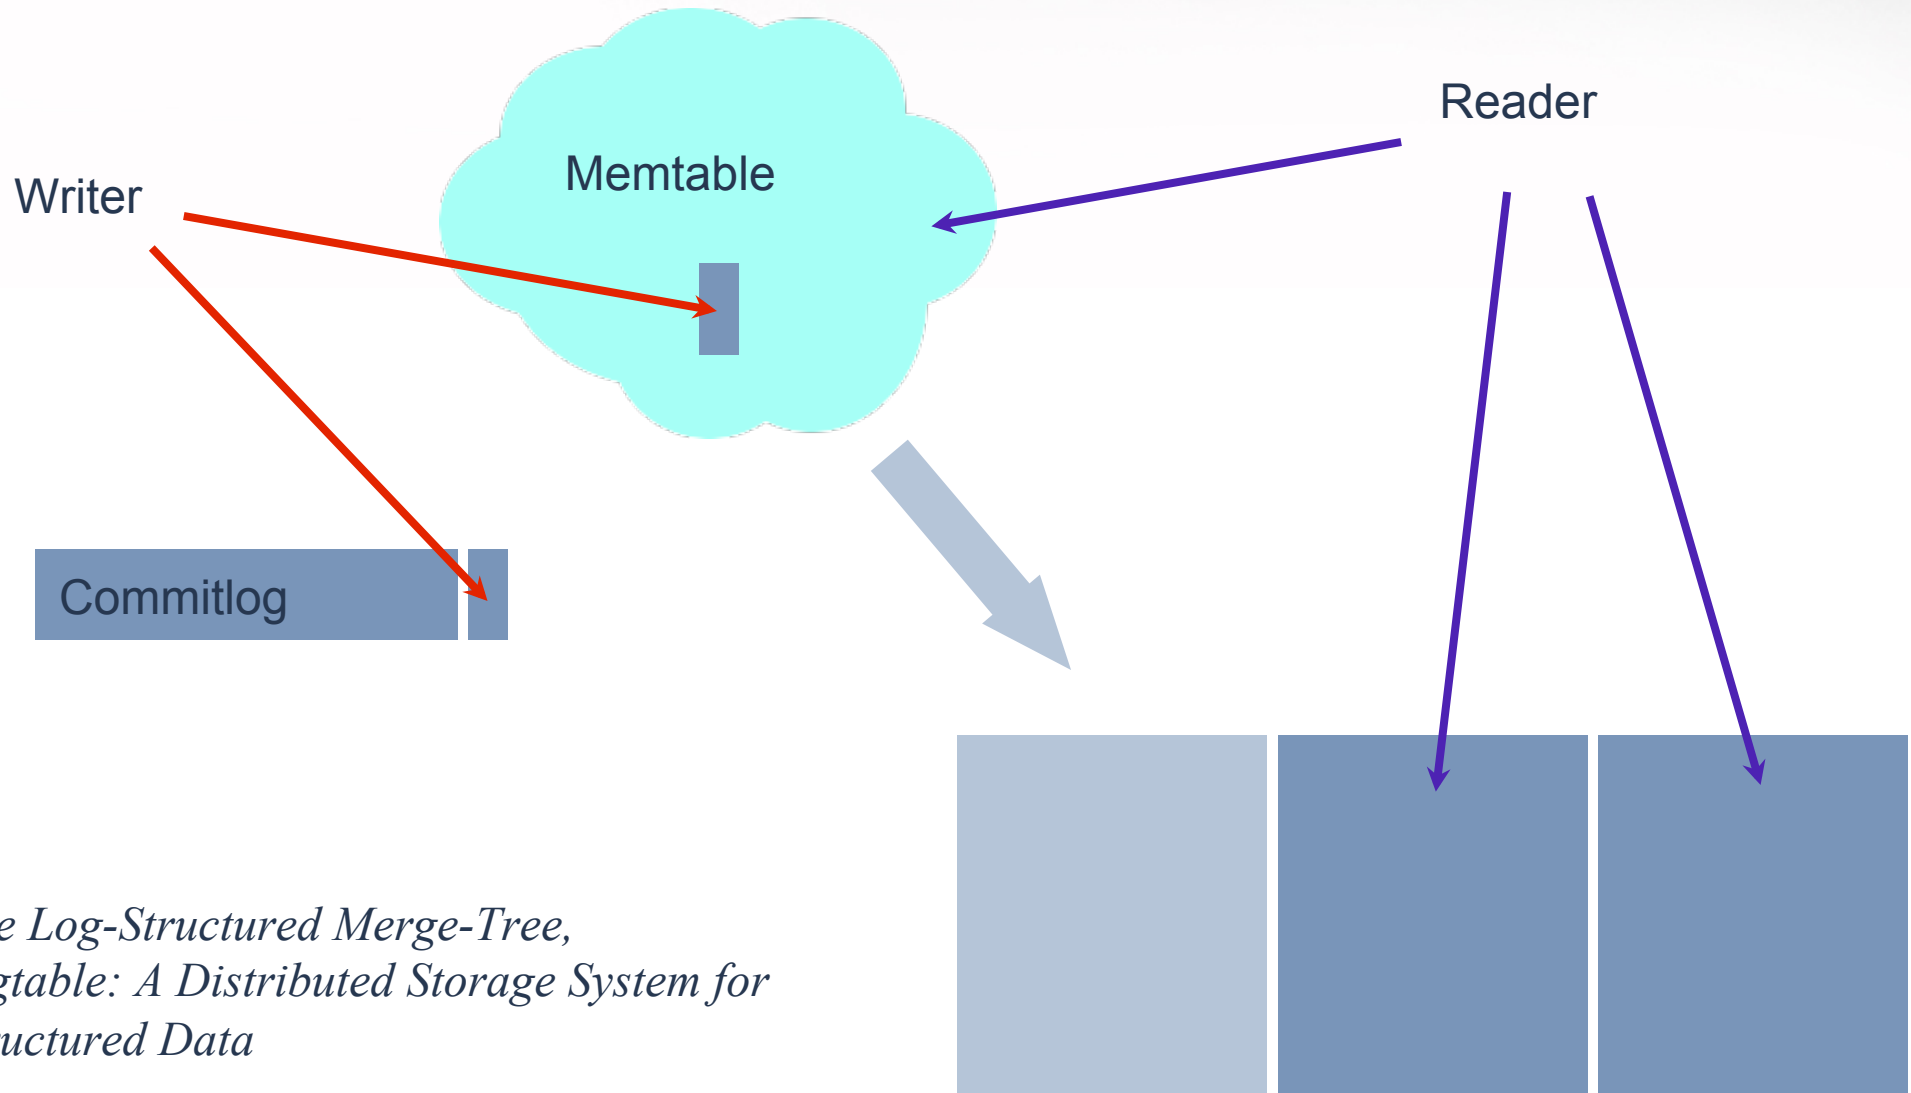
- Easy to deploy and operate
- No single points of failure
- Scale and change nodes with no downtime
- Cross-DC, multi-master clusters
- Allocate resources for OLAP vs OLTP
 - With no ETL
- Based on Cassandra 0.8 and Hadoop 0.20.203

P2P replication

inode
c45caed-5bd7-...



Locally-managed data



*The Log-Structured Merge-Tree,
Bigtable: A Distributed Storage System for
Structured Data*

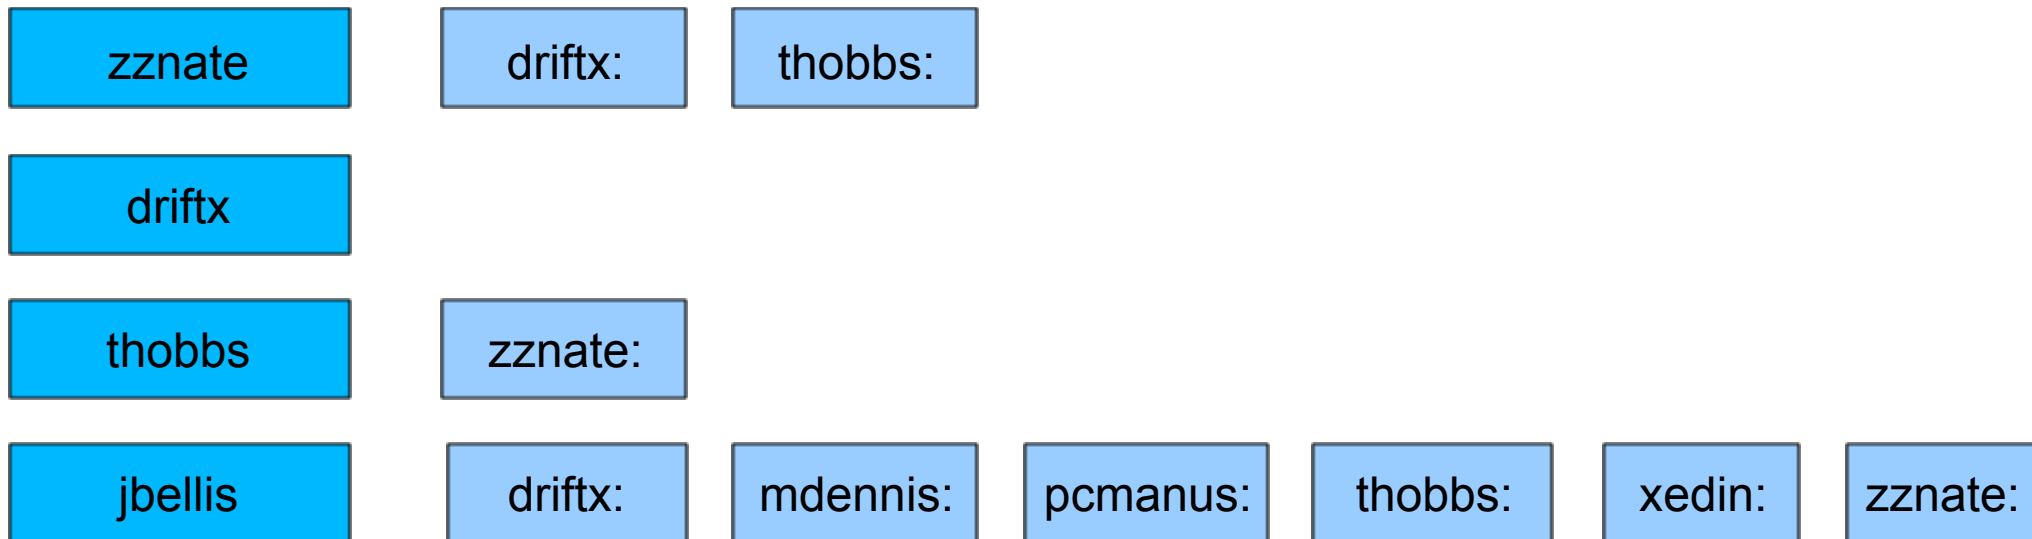
Data Model

- ColumnFamilies contain rows + columns
- Sparse model: no rewriting for ALTER
- (Not really schemaless for a while now)

zznate	Password: *	Name: Nate	
driftx	Password: *	Name: Brandon	
thobbs	Password: *	Name: Tyler	
jbellis	Password: *	Name: Jonathan	Site: datastax.com

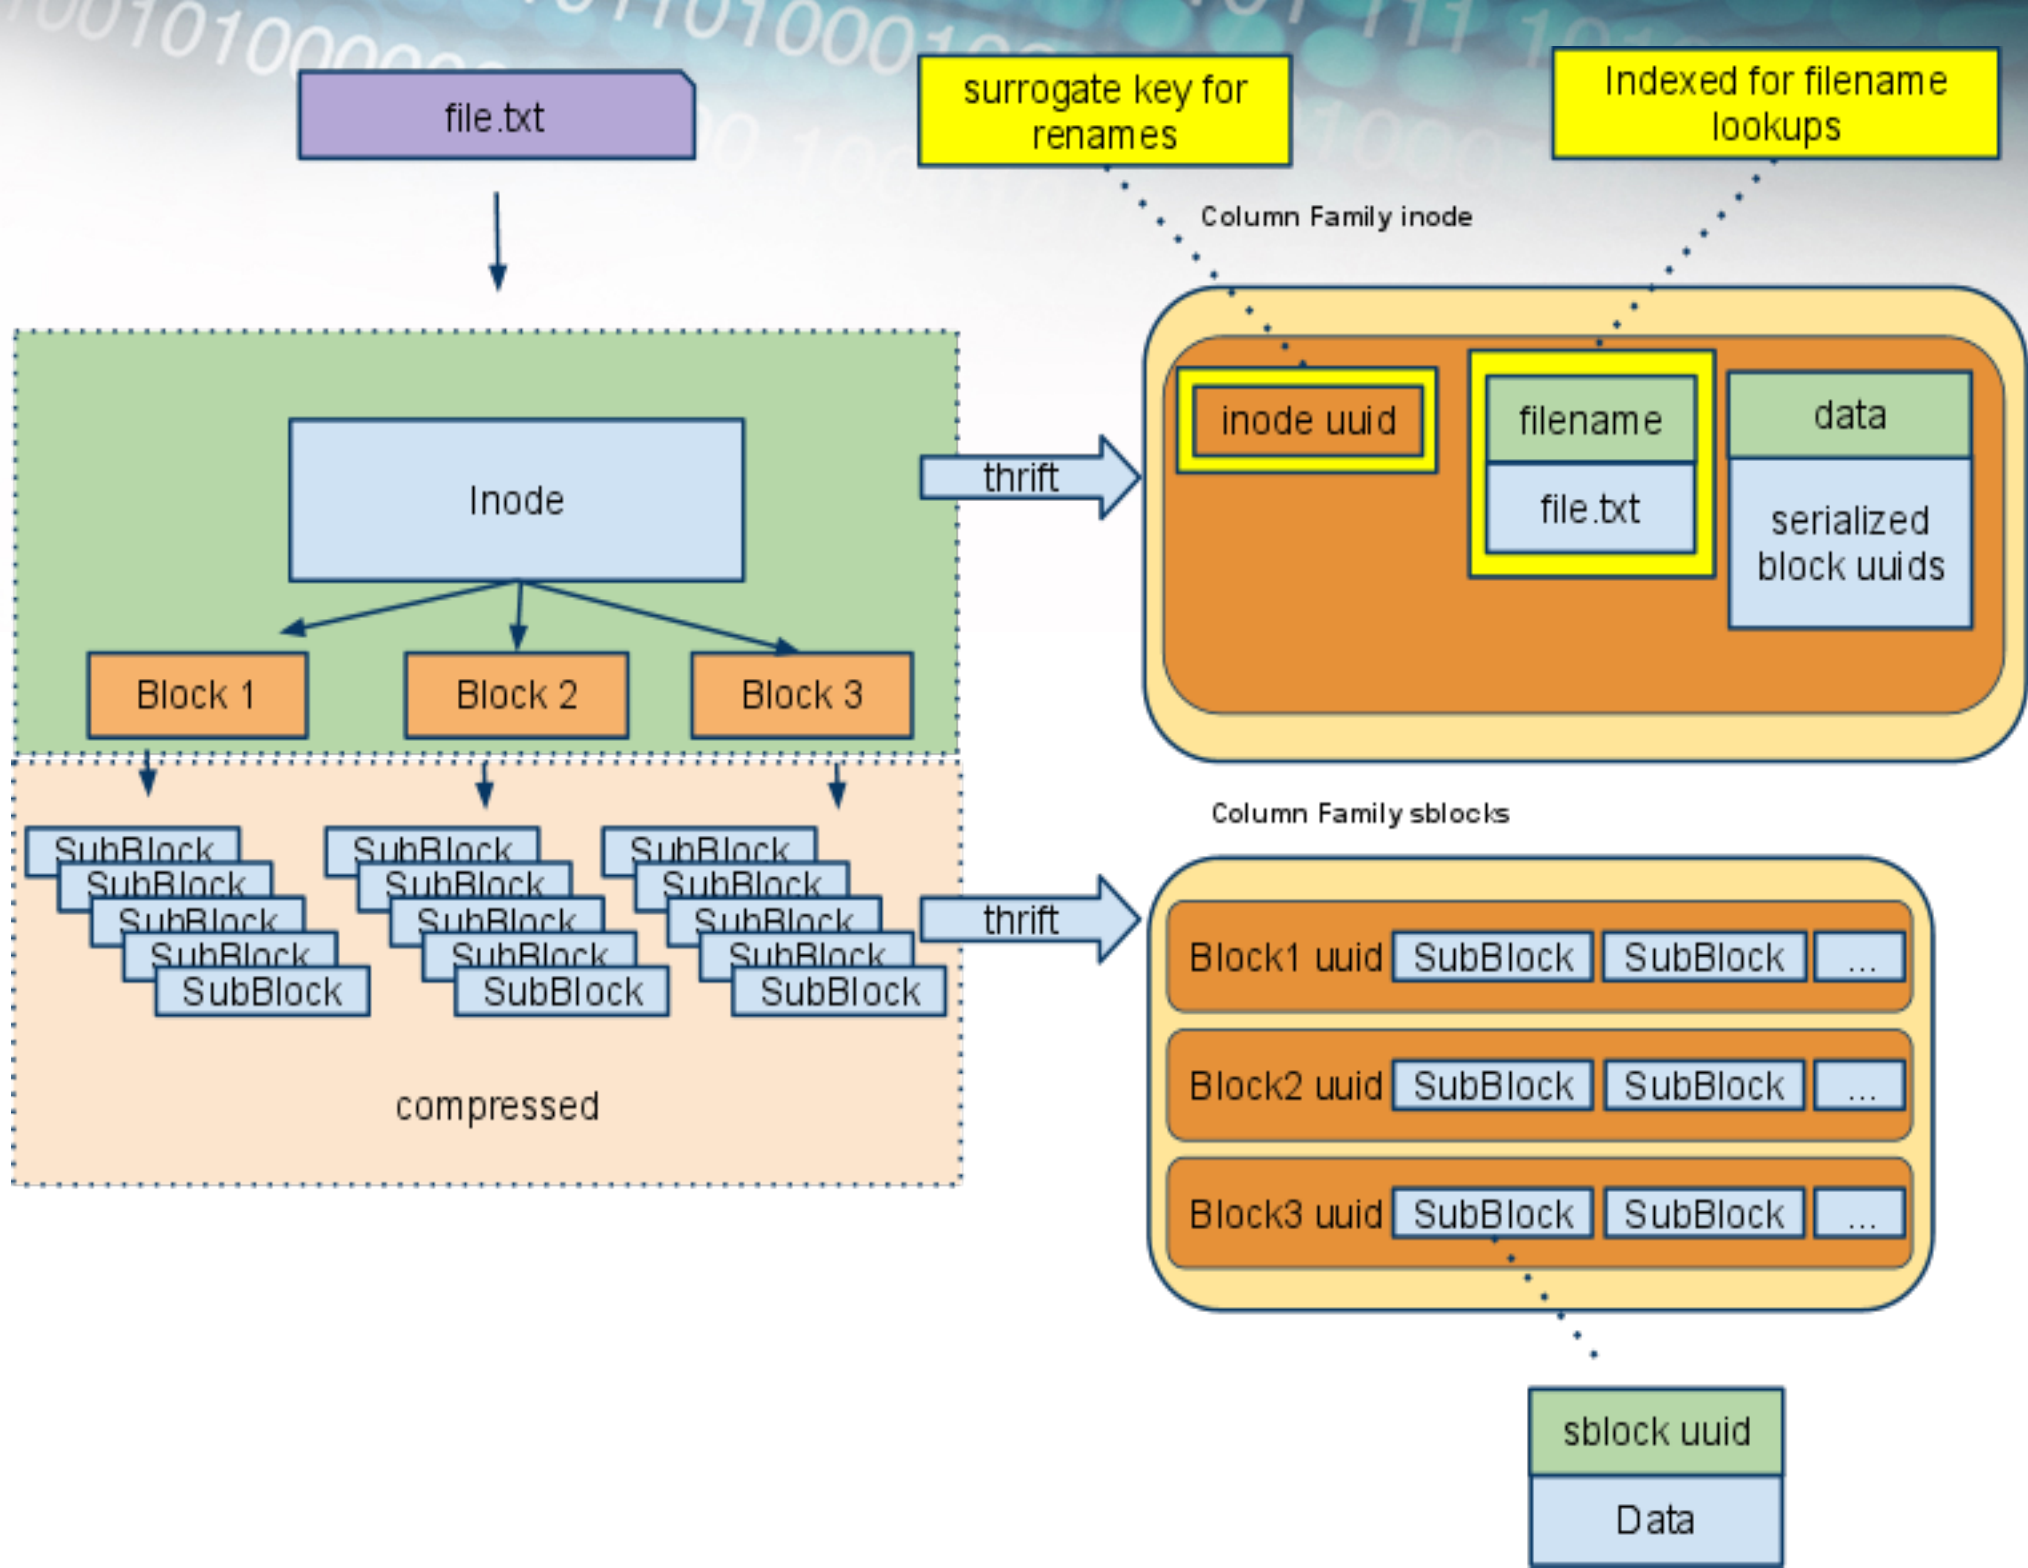
Large rows

- Sparse model + sstable storage means we can have arbitrarily large rows
- Columns are sorted by comparator

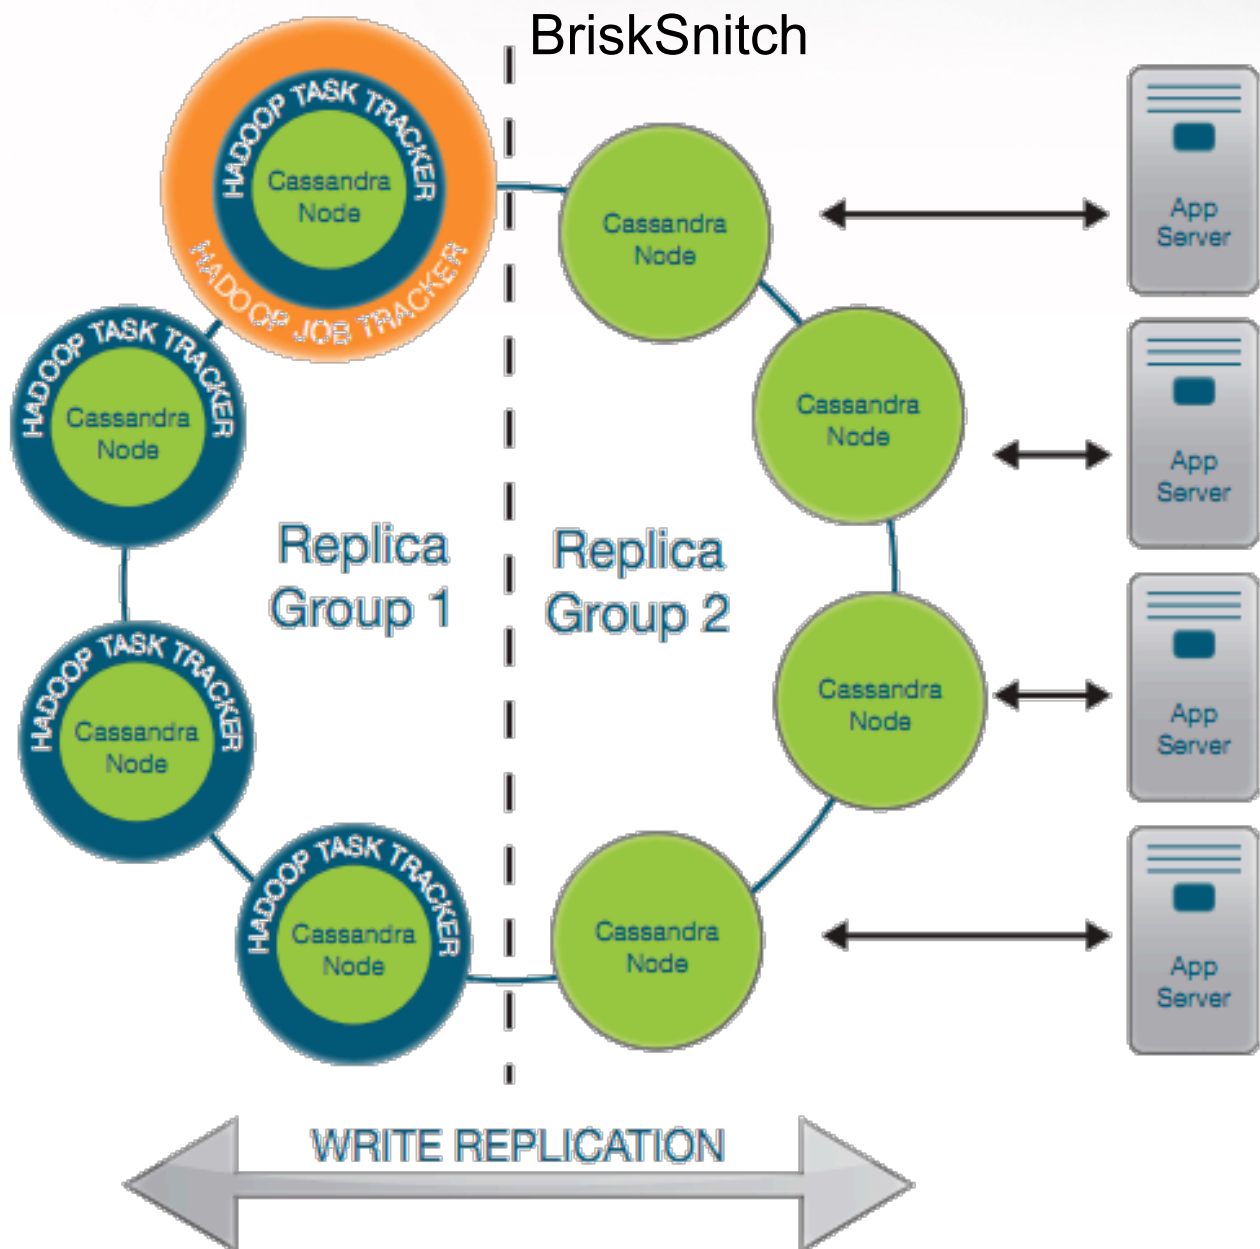


CassandraFS

- data stored as ByteBuffer internally — excellent fit for blocks
- inode, sblocks columnfamilies
 - (blocks split into subblocks to smooth write impact)
- local reads mmap data directly (no rpc)
- blocks are compressed with google snappy
- `hadoop distcp hdfs:///mydata cfs:///mydata`



Hybrid Workloads



Hive support

- Hive MetaStore in Cassandra
 - Unified schema view from any node, with no external systems and no SPOF
 - Automatically maps Cassandra column families to Hive tables
- Supports static and dynamic column families (and supercolumns)

Hive: CFS and ColumnFamilies

CFS:

```
CREATE TABLE users (name STRING, zip INT);  
LOAD DATA LOCAL INPATH 'kv2.txt' OVERWRITE INTO TABLE users;  
select count(*), zip from invites group by bar;
```

Column Family (fixed):

```
CREATE EXTERNAL TABLE Keyspace1.Users(name STRING, zip INT)  
STORED BY 'org.apache.hadoop.hive.cassandra.CassandraStorageHandler';
```

Column Family (dynamic):

```
CREATE EXTERNAL TABLE Keyspace1.Users  
(row_key STRING, column_name STRING, value string)  
STORED BY 'org.apache.hadoop.hive.cassandra.CassandraStorageHandler';
```

Pig Support

with standard Cassandra:

```
$ export PIG_HOME=/path/to/pig
$ export PIG_INITIAL_ADDRESS=localhost
$ export PIG_RPC_PORT=9160
$ export
PIG_PARTITIONER=org.apache.cassandra.dht.RandomPartitioner
$ contrib/pig/bin/pig_cassandra
grunt>
```

with Brisk:

```
$ bin/brisk pig
grunt>
```

Pig: CFS and ColumnFamilies

CFS:

```
grunt> data = LOAD 'cfs:///example.txt' using PigStorage() as  
(name:chararray, value:long);
```

ColumnFamily:

```
data = LOAD 'cassandra://Demo1/Scores' using  
CassandraStorage() AS (key, columns: {T: tuple(name, value)});
```

ColumnFamily Slices:

```
data = LOAD 'cassandra://Demo1/  
Scores&slice_start=M&slice_end=S' using CassandraStorage()  
AS (key, columns: {T: tuple(name, value)});
```

Portfolio Manager Demo

- Low-Latency Side (OLTP)
 - Live tick prices for NASDAQ stocks
 - Thousands of Portfolios
- Batch Analytics Side (OLAP)
 - Historical Stock prices
 - Historical VAR (Value At Risk) calculation for each portfolio using Hive
 - VAR number is written back to Cassandra from Hive

Portfolio Demo ColumnFamilies

Portfolios

Portfolio1	GOOG	LNKD	P	AMZN	AAPL
	5	7	50	100	4

LiveStocks

AAPL	LAST
	\$950.00

StockHist

GOOG	2011-01-01	2011-01-02	2011-01-03	2011-01-04
	\$2.99	\$1.00	\$9.29	\$29.10

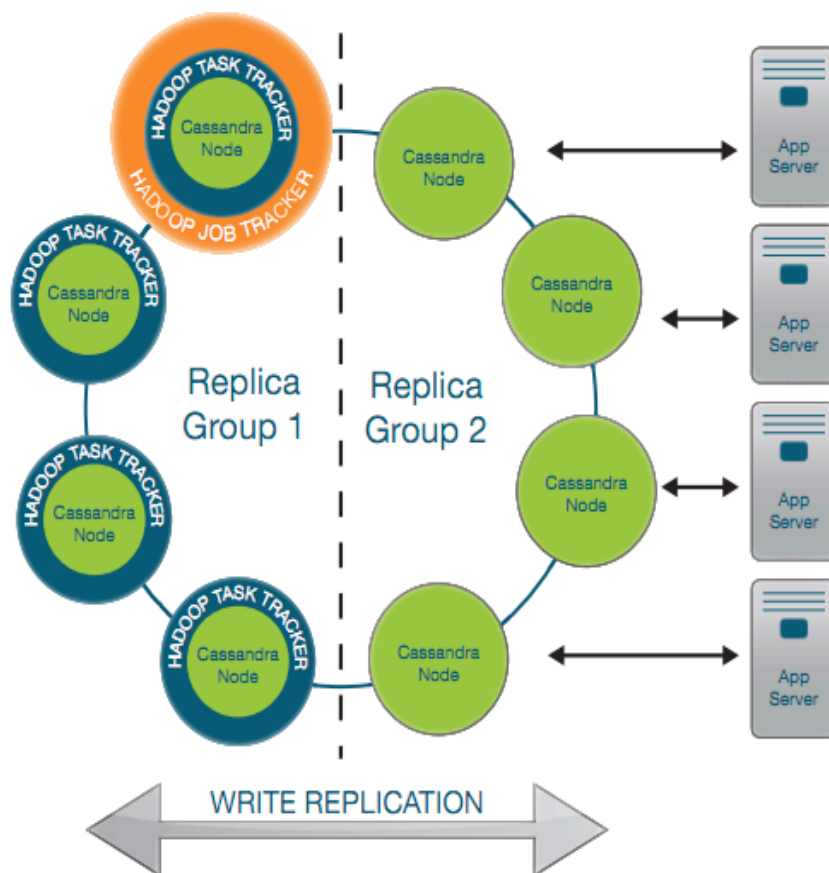
Portfolio Demo Workloads

OLAP:

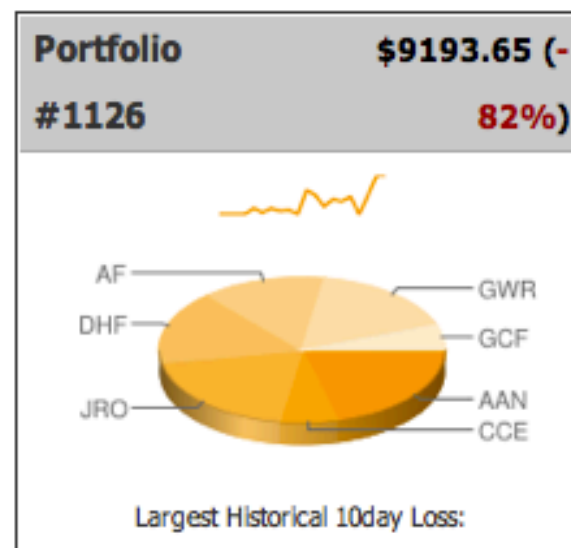


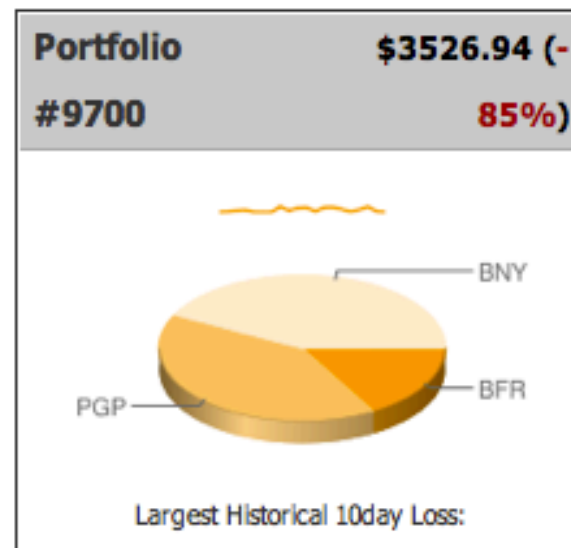
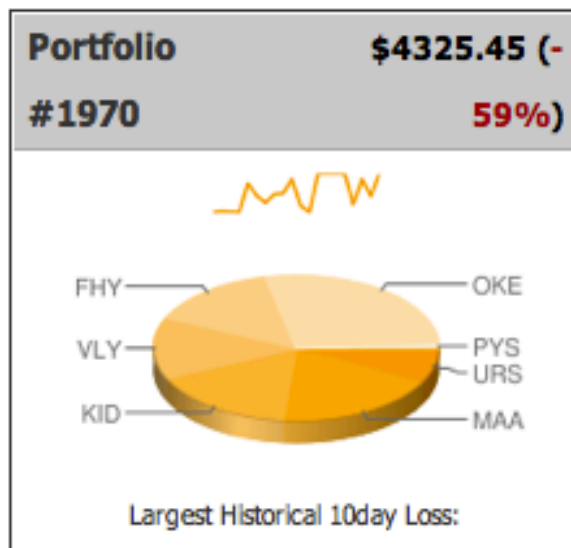
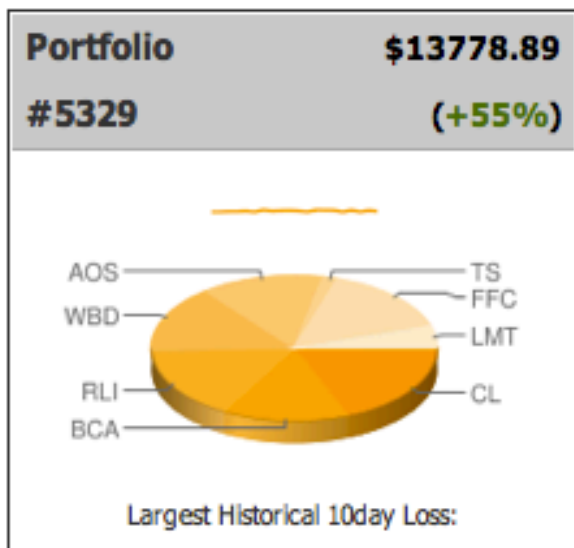
- Hive <- Cassandra
 - Portfolios
 - Historical Prices
- Intermediate Results
- Hive -> Cassandra
- Every N minutes

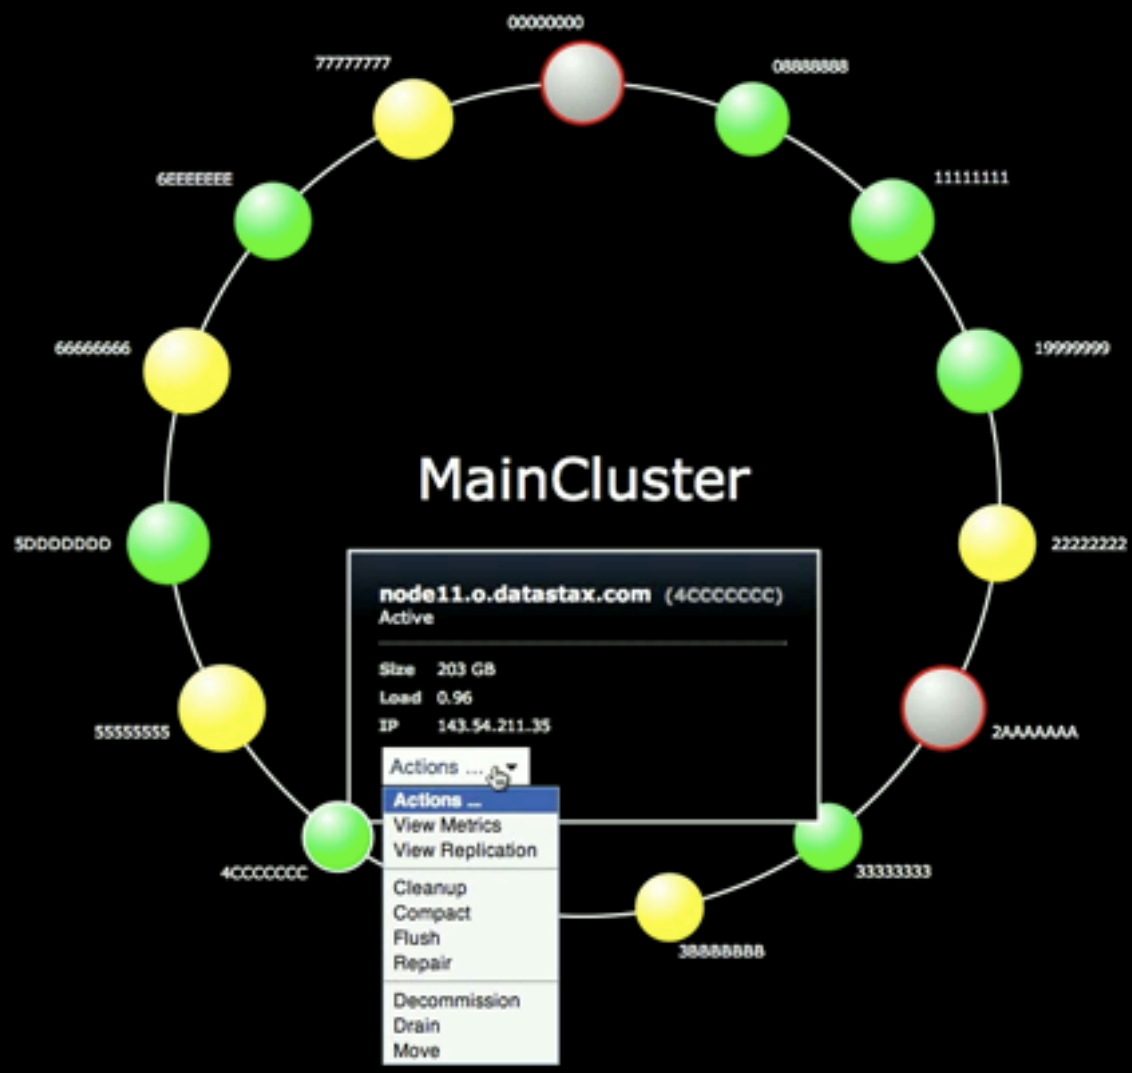
OLTP:

- WebBased Portfolios
- Live Prices for today
- Historical Prices



Portfolio Manager Demo





Dashboard

Cluster Management

Events & Alerts

Performance

Data Modeling

Data Explorer

Hadoop Jobs

[Refresh](#)

Auto-refreshes every minute.

[View Full Details](#)

	Job	Progress	Started	Duration	User
	TeraSort	<div style="width: 33%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 1/2 Reduces: 0/1	5/12/11 12:27 PM	40s	root
	TeraSort	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 2/2 Reduces: 1/1	5/12/11 12:26 PM	34s	root
	TeraGen	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 2/2 Reduces: 0/0	5/12/11 12:25 PM	44s	root
	PiEstimator	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 10/10 Reduces: 1/1	5/12/11 12:19 PM	49s	riptano
	select count (*) from server_request(Stage-1)	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 15/15 Reduces: 1/1	5/12/11 10:56 AM	1m 31s	riptano
	select count(*) size, method from myT...size(Stage-2)	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 10/10 Reduces: 10/10	5/11/11 8:30 PM	57s	riptano
	select count(*) size, method from myT...size(Stage-1)	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 2/2 Reduces: 10/10	5/11/11 8:29 PM	50s	riptano
	select count(*), method from myTest...method(Stage-1)	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 2/2 Reduces: 10/10	5/11/11 8:24 PM	46s	riptano
	select count(*), method from myTest...method(Stage-1)	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 2/2 Reduces: 10/10	5/11/11 8:12 PM	49s	riptano
	select count(*), method from myTest...method(Stage-1)	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 2/2 Reduces: 1/1	5/11/11 8:04 PM	21s	riptano
	select min(duration) from myTest(Stage-1)	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 2/2 Reduces: 1/1	5/11/11 8:02 PM	24s	riptano
	select max(duration) from myTest(Stage-1)	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 2/2 Reduces: 1/1	5/11/11 8:02 PM	24s	riptano
	select count(*) from myTest(Stage-1)	<div style="width: 100%;"><div style="background-color: #0070C0; height: 10px;"></div></div> Maps: 2/2 Reduces: 1/1	5/11/11 8:00 PM	25s	riptano

Open Source

- <https://github.com/riptano/brisk>
- https://github.com/riptano/brisk/tree/master/demos/portfolio_manager
- Apache License

