NoSQL Databases and Monitoring

Database Futures Workshop

Jérôme Belleman Christos Margiolas
CERN – IT-PES

June 2011

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Outline

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Section 1

Monitoring Context

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

CERN Batch System Monitoring

- Investigate more advanced features of the batch system
- Reduce problem identification time
- Lego-like, interchangeable set of building blocks

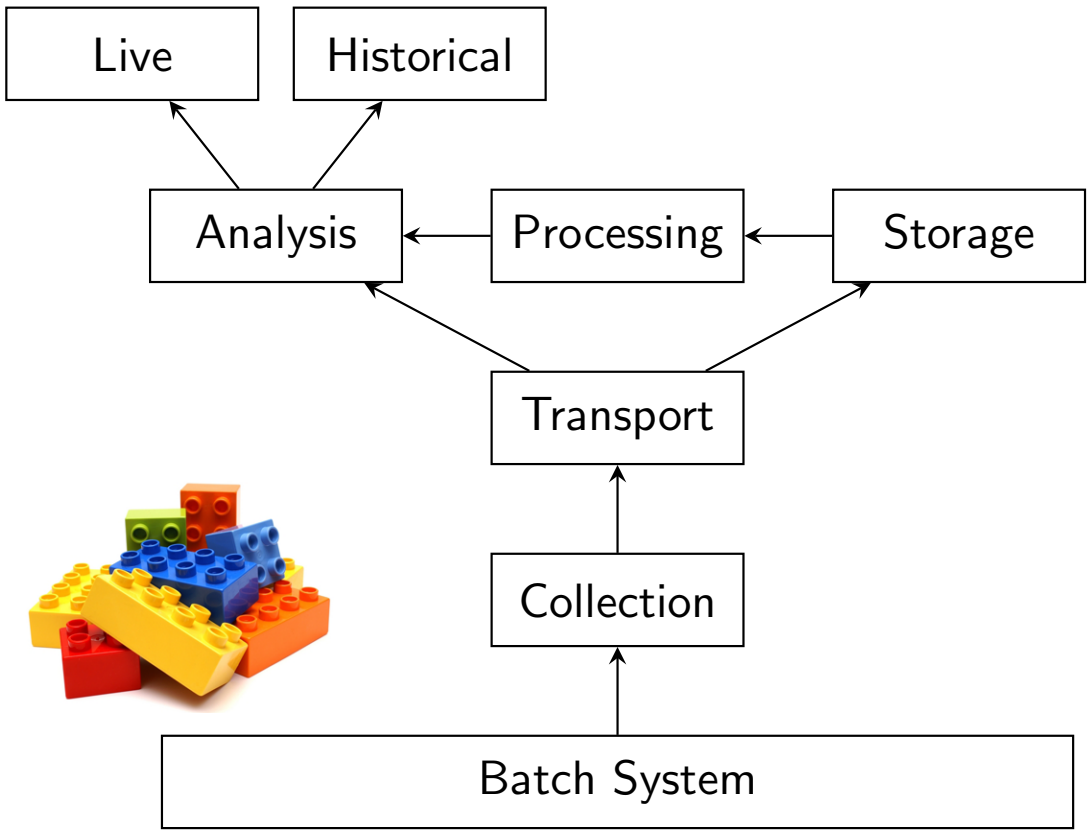

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

What to Monitor?

- Number of jobs
- Users
- Hosts
- Slots
- Queues
- ...

→ Correlations & heavy analytics

- Live representations
- Historical representations

→ Browse metrics over time

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

How Much to Monitor?

- > 200 000 jobs/day
- > 3 700 nodes
- > 30 000 cores

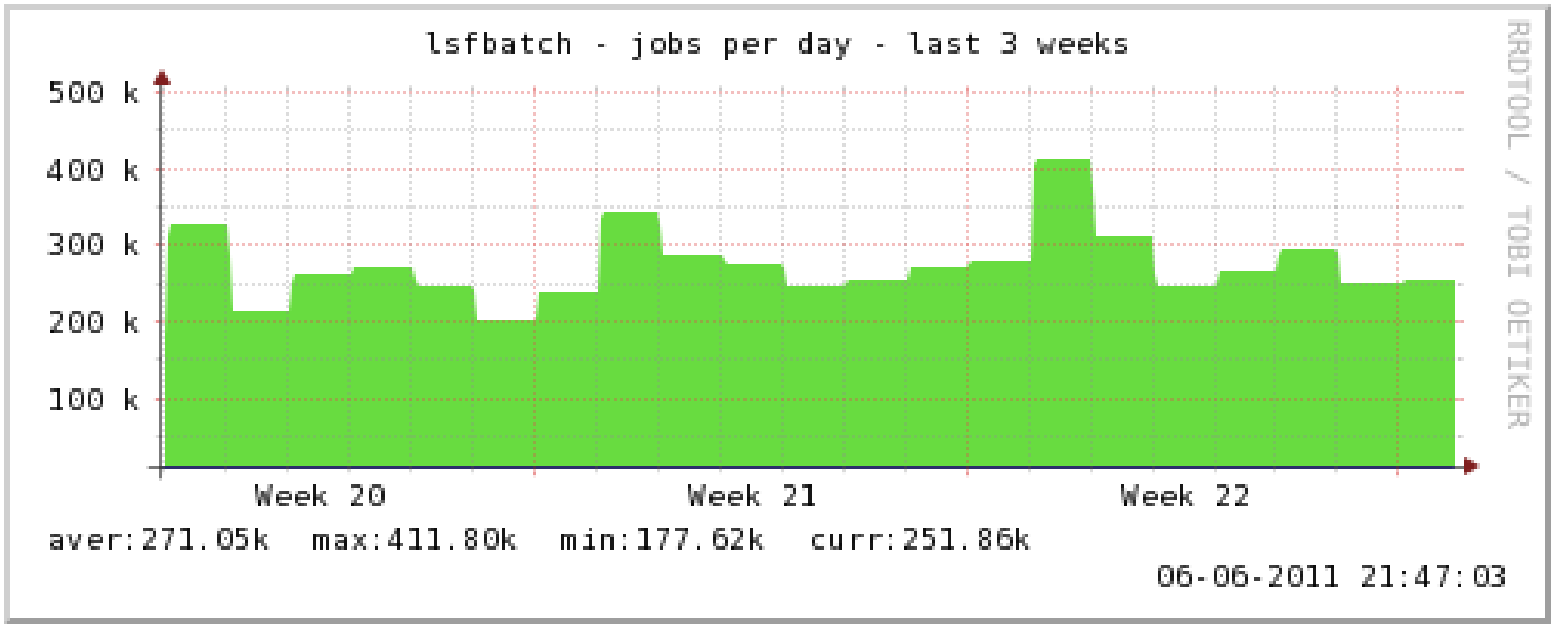

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

What Do We Need?

- Flexible schemas
- Scalability
- Data collection → Many small writes, very often
- Heavy analytics → Massive reads, quite often
- Horizontal & vertical aggregations
- Time views

→ Let's investigate the use of a NoSQL database

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Section 2

NoSQL Databases for Monitoring Purposes

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

General Concepts

- NoSQL ⇝ Non Relational
- A broader set of structured storage
- Different consistency models
- Different classes of NoSQL databases

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Different Classes of NoSQL Databases

- Column-oriented databases (e.g. Google BigTable, Apache HBase, Apache Cassandra)
- Key-value stores (e.g. Amazon Dynamo, Project Voldemort, MongoDB, Riak, Google BigTable, Apache Cassandra)
- Document stores (e.g. Apache CouchDB, MongoDB)

We've been looking at Cassandra.

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Key-Value Stores such as Cassandra

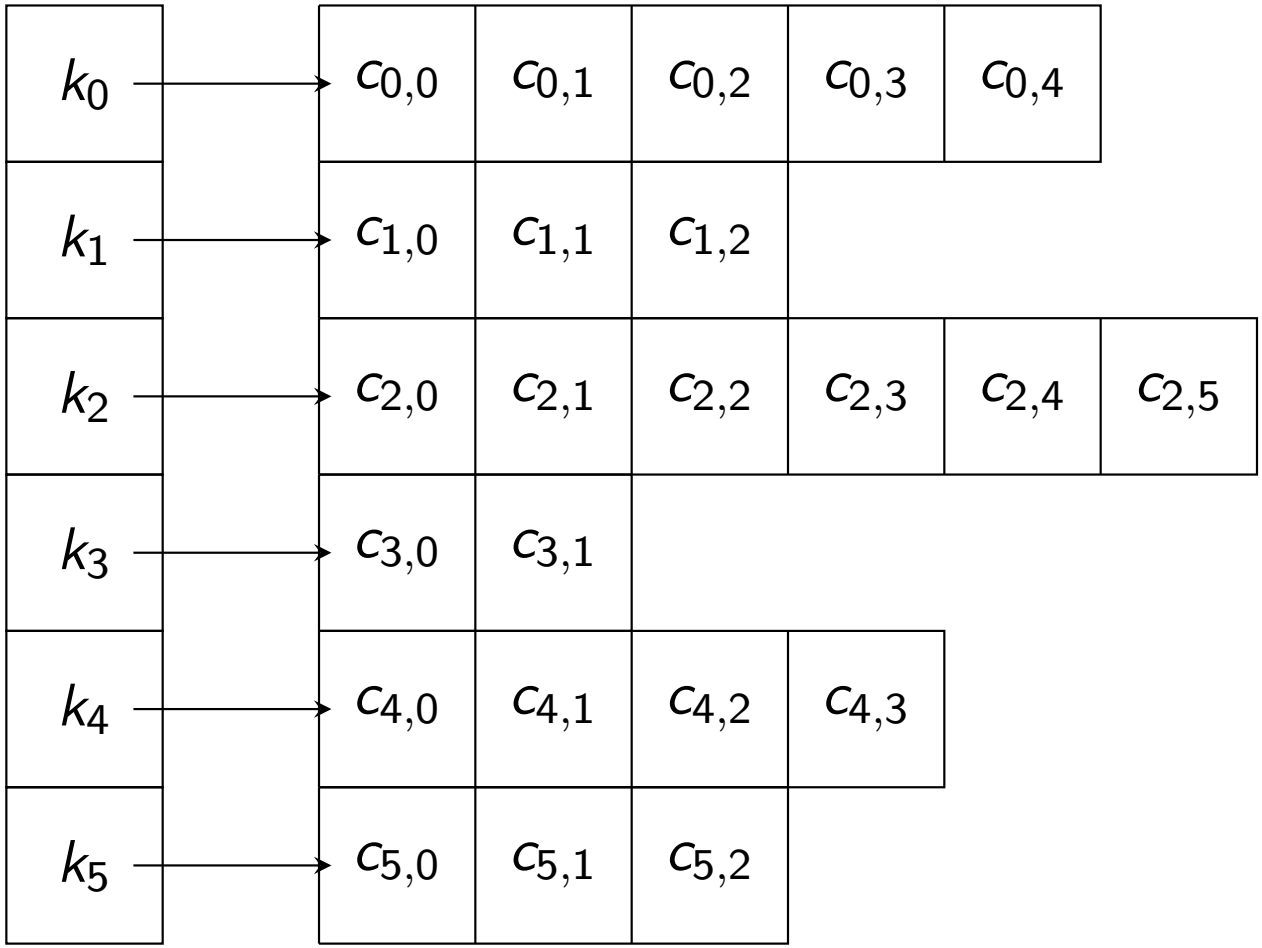

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Key-Value Store Schema Design

- No relationships between multiple tables
- Multiple tables with different keys

  ↓

- New query → New table

  ↓

- Redundancy
- Constraints on the application level

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

How Fast?

- (Quite) Simple-minded; e.g. relaxed about consistency
- High throughput
- Horizontal scalability
- No need for object-relational mapping

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Suitable for Monitoring?

SQL	NoSQL	For Monitoring?
Immediate consistency	Eventual consistency	
High throughput	Higher throughput	
Rigid schemas	Flexible schemas	
H. Scalability	H. Scalability	
Indexes	Another table	
DB-Level constraints	Application constraints	

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Suitable for Monitoring?

SQL	NoSQL	For Monitoring?
Immediate consistency	Eventual consistency	✓
High throughput	Higher throughput	
Rigid schemas	Flexible schemas	
H. Scalability	H. Scalability	
Indexes	Another table	
DB-Level constraints	Application constraints	

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Suitable for Monitoring?

SQL	NoSQL	For Monitoring?
Immediate consistency	Eventual consistency	✓
High throughput	Higher throughput	✓
Rigid schemas	Flexible schemas	
H. Scalability	H. Scalability	
Indexes	Another table	
DB-Level constraints	Application constraints	

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Suitable for Monitoring?

SQL	NoSQL	For Monitoring?
Immediate consistency	Eventual consistency	✓
High throughput	Higher throughput	✓
Rigid schemas	Flexible schemas	✓
H. Scalability	H. Scalability	
Indexes	Another table	
DB-Level constraints	Application constraints	

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Suitable for Monitoring?

SQL	NoSQL	For Monitoring?
Immediate consistency	Eventual consistency	✓
High throughput	Higher throughput	✓
Rigid schemas	Flexible schemas	✓
H. Scalability	H. Scalability	✓
Indexes	Another table	
DB-Level constraints	Application constraints	

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Suitable for Monitoring?

SQL	NoSQL	For Monitoring?
Immediate consistency	Eventual consistency	✓
High throughput	Higher throughput	✓
Rigid schemas	Flexible schemas	✓
H. Scalability	H. Scalability	✓
Indexes	Another table	✓
DB-Level constraints	Application constraints	

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Suitable for Monitoring?

SQL	NoSQL	For Monitoring?
Immediate consistency	Eventual consistency	✓
High throughput	Higher throughput	✓
Rigid schemas	Flexible schemas	✓
H. Scalability	H. Scalability	✓
Indexes	Another table	✓
DB-Level constraints	Application constraints	~

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Suitable for Monitoring?

SQL	NoSQL	For Monitoring?
Immediate consistency	Eventual consistency	✓
High throughput	Higher throughput	✓
Rigid schemas	Flexible schemas	✓
H. Scalability	H. Scalability	✓
Indexes	Another table	✓
DB-Level constraints	Application constraints	~

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

What Do We Need?

- Flexible schemas
- Scalability
- Data collection → Many small writes, very often
- Heavy analytics → Massive reads, quite often
- Horizontal & vertical aggregations
- Time views

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

What Do We Need?

- ☑ Flexible schemas
- ☐ Scalability
- ☐ Data collection → Many small writes, very often
- ☐ Heavy analytics → Massive reads, quite often
- ☐ Horizontal & vertical aggregations
- ☐ Time views

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

What Do We Need?

- ☑ Flexible schemas
- ☑ Scalability
- ☐ Data collection → Many small writes, very often
- ☐ Heavy analytics → Massive reads, quite often
- ☐ Horizontal & vertical aggregations
- ☐ Time views

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

What Do We Need?

- ☑ Flexible schemas
- ☑ Scalability
- ☑ Data collection → Many small writes, very often
- ☐ Heavy analytics → Massive reads, quite often
- ☐ Horizontal & vertical aggregations
- ☐ Time views

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

What Do We Need?

- ☑ Flexible schemas
- ☑ Scalability
- ☑ Data collection → Many small writes, very often
- ☑ Heavy analytics → Massive reads, quite often
- ☐ Horizontal & vertical aggregations
- ☐ Time views

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

What Do We Need?

- ☑ Flexible schemas
- ☑ Scalability
- ☑ Data collection → Many small writes, very often
- ☑ Heavy analytics → Massive reads, quite often
- ☑ Horizontal & vertical aggregations
- ☐ Time views

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

What Do We Need?

- ☑ Flexible schemas
- ☑ Scalability
- ☑ Data collection → Many small writes, very often
- ☑ Heavy analytics → Massive reads, quite often
- ☑ Horizontal & vertical aggregations
- ☑ Time views

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Section 3

Examples

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

A Time Series Database: OpenTSDB

- On top of Apache HBase
- Horizontal aggregation → Free time browsing
- Vertical aggregation → Drill down details

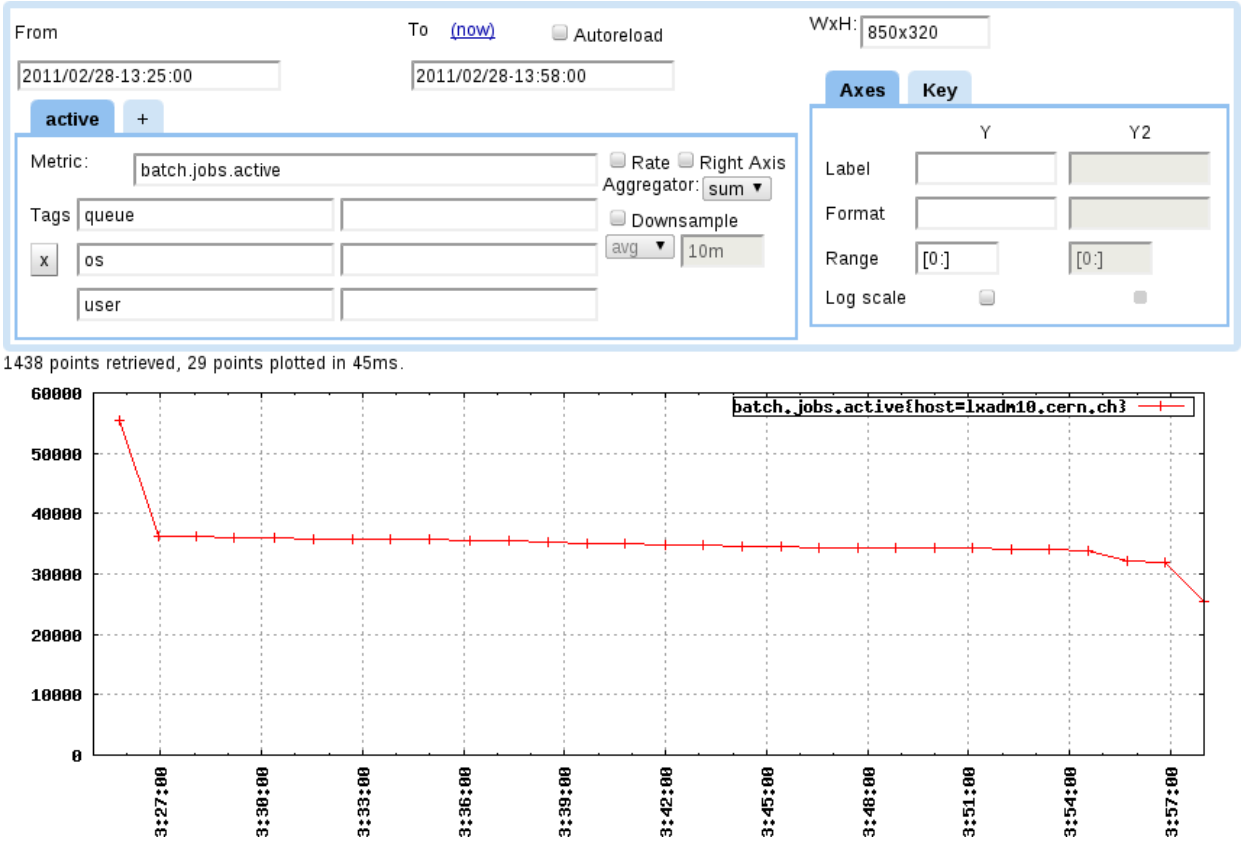

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Batch Accounting with Cassandra

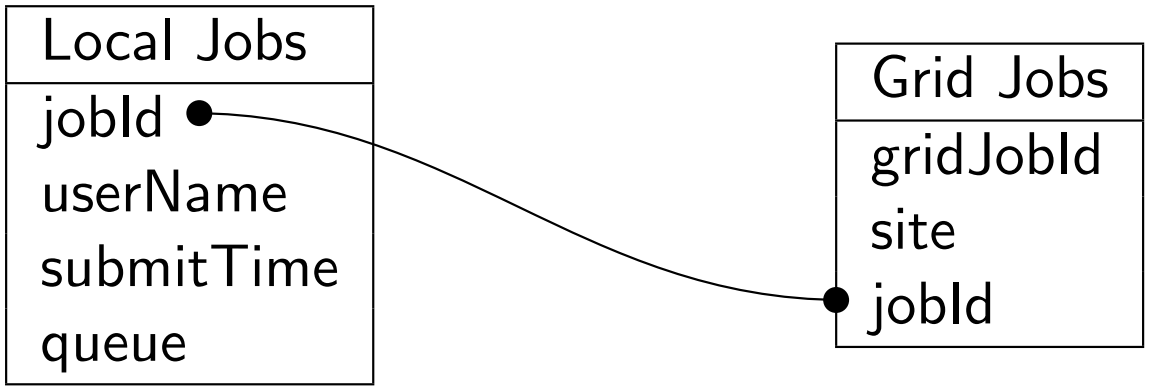

16	fred	3:04	1nh	42	cern
21	bob	6:10	8nh	70	ral
28	alice	9:49	1nh		
37	bob	12:08	1nh	83	ral

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Batch Accounting with Cassandra

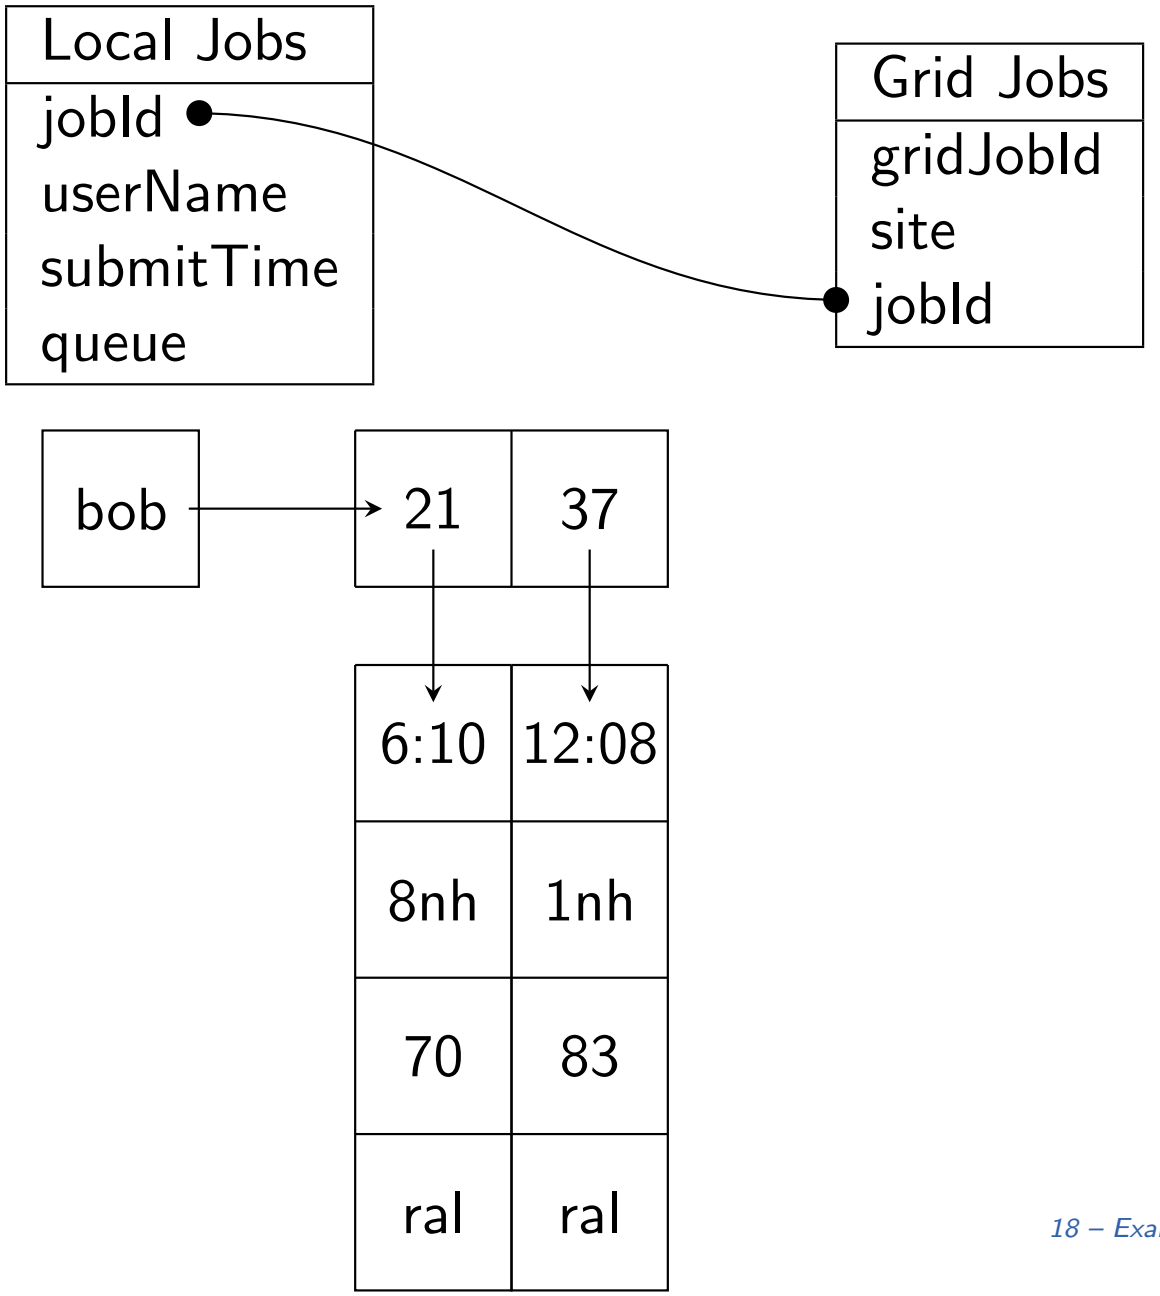

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Section 4

Outlook

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Outlook

- Scalability and throughput through simplicity
- Complacency regarding consistency acceptable for monitoring
- Flexibility
- Many available open-source implementations
- New schema design approach

→ We're investigating the use of a NoSQL database

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it

Thanks!

Questions?

CERN IT Department
CH-1211 Genève 23
Switzerland
www.cern.ch/it