

NoSQL Databases and Monitoring

IT Monitoring Working Group

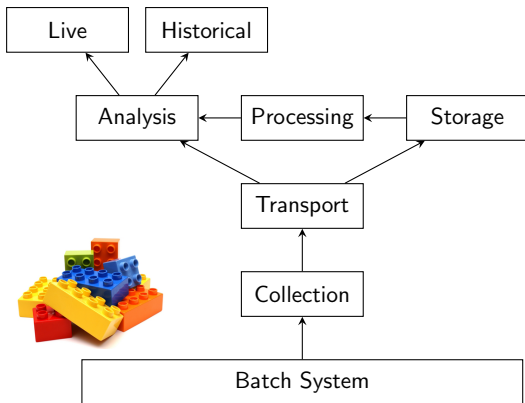
Jérôme Belleman Christos Margiolas
CERN – IT-PES
19 September 2011

- 1 Monitoring Context
- 2 NoSQL Databases for Monitoring Purposes
- 3 Examples
- 4 Outlook

Section 1

Monitoring Context

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



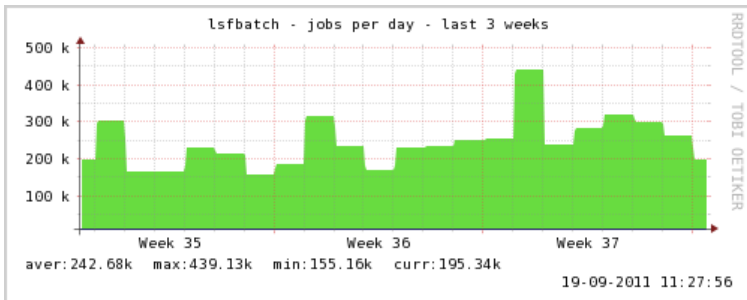
- Jobs (number, CPU usage, I/O, ...)
- Users
- Hosts
- Slots
- Queues
- ...

→ Correlations & heavy analytics

- Live representations
- Historical representations

→ Browse metrics over time

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



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

→ Let's investigate the use of a NoSQL database

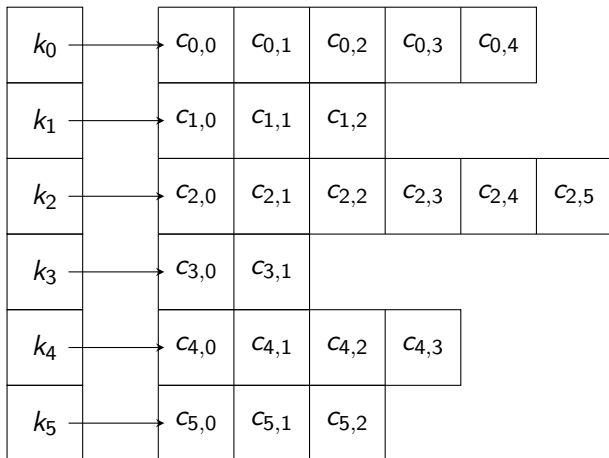
Section 2

NoSQL Databases for Monitoring Purposes

- NoSQL \rightsquigarrow Non Relational
- A broader set of structured storage
- Different consistency levels
- 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.



- (Quite) Simple-minded
- High throughput
- Horizontal scalability
- No need for object-relational mapping

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

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

Section 3

Examples

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

From To Autoreload WxH:

active +


Metric: Rate Right Axis
 Aggregator:

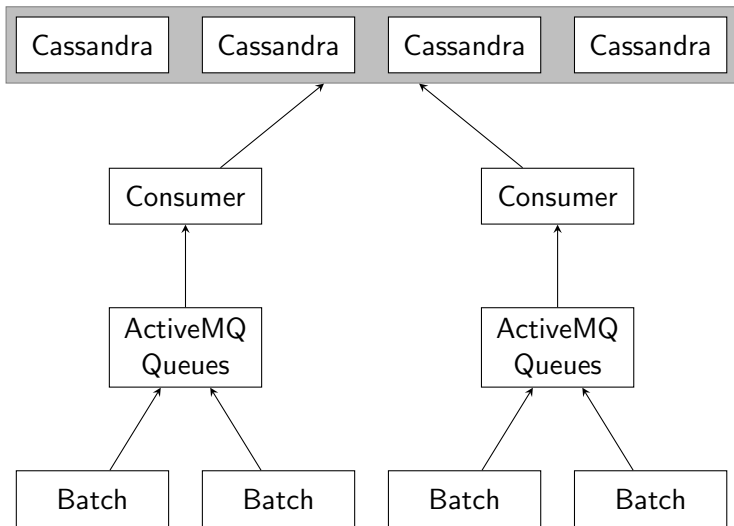
Tags: Downsample
 Avg

Axes **Key**


Y Y2
 Label
 Format
 Range
 Log scale

1438 points retrieved, 29 points plotted in 45ms.





- Reading from bhpert output
- > 30 000 groups/users every 5 minutes



Section 4

Outlook

- Throughput through simplicity
- Relaxed consistency acceptable for monitoring
- Horizontal scalability
- Flexibility
- Many available open-source implementations
- New schema design approach

→ NoSQL databases seem a good fit for monitoring



Questions?