

Batch Monitoring and Testing

HEPiX Spring 2011

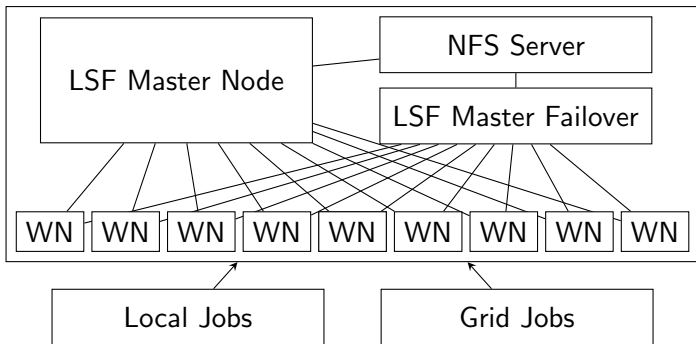
Jérôme Belleman
CERN – IT-PES
May 2011

- 1 Motivations
- 2 What We Have
- 3 What We Want
- 4 Technology Investigations
- 5 Outlook

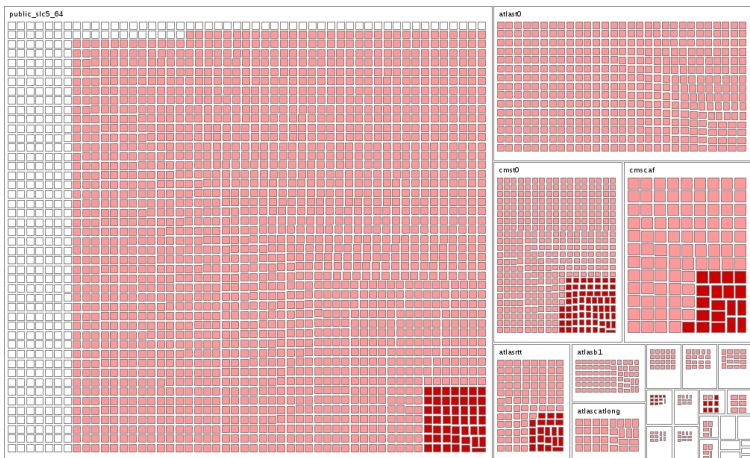
Section 1

Motivations


- Platform LSF 7.0.6
- All resources in one redundant instance
- Different shares for different customers: public, grid and several for CERN experiments



- > 3 700 nodes
- > 30 000 cores



- 180 000 jobs/day
- 25 resource pools
- Supports high-energy physics and other projects in the laboratory






Service Level Status – CERN Batch Service

Section 2


What We Have

- Lemon
 - For Linux-based computer centre machines
 - Used for many services at CERN
 - Client-Server oriented
 - Node-level and alarm-driven monitoring
- Log files must be manually scrutinised
- Commands must be manually issued and their output manually read

By queue, by user, by group:



Node monitoring	→	Job-level monitoring
Limited set of views	→	Add <i>many</i> more (e.g. users, CPU time consumed, hosts, ...)
Static plots	→	More flexible view set
Data collected once for plots	→	Keep raw data
Missing correlations	→	Allow for correlations (e.g. jobs/user, CPU time/queue, hosts/host partition, ...)



Section 3

What We Want

We need to answer questions such as:

- Why is a job not running?
- Does fairshare behave as intended?
- Is the use unreasonable?
- Are we low on resources?

Monitoring:

- Determine correct setup of fairshare
- Investigate more advanced LSF features
- Decide about needs for restructuring queues
- Reduce problem identification time

Profiling:

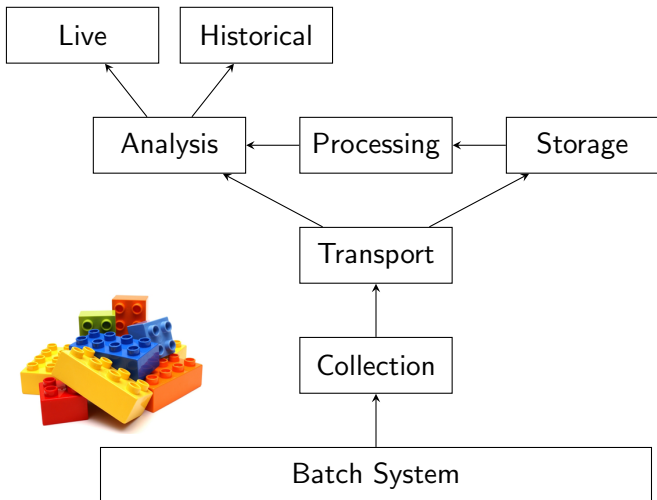
- How resources are used and understand under-utilisation
- Identify job types and define application profiles
- Capacity planning

- For the service managers
- For the service desk
- For the end user
- Pluggable: “Write your own data miner against our data”

We need an integration instance:

- To test significant changes
- To test new queues, fine-tune fairshare, . . .
- Which may constantly be submitted jobs if needs be


- Survey of existing tools
- Integrate existing monitoring data sources
- Batch system independent
- IT-wide effort
- Lego-like, interchangeable building blocks



Section 4

Technology Investigations

- Relational databases:
Oracle, MySQL,
PostgreSQL
- NoSQL databases:
Cassandra, MongoDB,
CouchDB, Riak
- Time series databases:
OpenTSDB, RRDtool,
Whisper
- Transport: ActiveMQ,
rsyslog
- User interface and
dashboards: Django,
Graphite, Flot



From To Autoreload WxH:

active +


Metric: Rate Right Axis
 Aggregator:

Tags Downsample
 x

Axes **Key**

Y Y2
 Label
 Format
 Range
 Log scale

1438 points retrieved, 29 points plotted in 45ms.



Section 5

Outlook

Current monitoring focused on nodes and alarms.

But we also want:

- Job-level monitoring
- To perform correlations
- Heavy analytics

How?

- Survey existing tools
- Batch-system independent
- Lego-like, interchangeable building blocks

When?

- Currently scouting existing tools
- First prototype in 6 months

We are interested in other people's experiences!

Questions & Discussion