A new monitoring tool from the site point of view

By Julia Andreeva, Max Boehm, Adrian Casajus, Benjamin Gaidioz, Costin Grigoras, Lukasz Kokoszkiewicz, **Elisa Lanciotti**, Ricardo Rocha, Pablo Saiz, Irina Sidorova, Andrei Tsaregorotsev .

WLCG 2009 Data-Taking Readiness Planning workshop
CERN 13-14 November 2008

Contents

- Introduction
  - Motivation and objective of the activity
- Structure of this new monitoring tool
  - Information work flow from the source to the final display
- Object of the monitoring:
  - Main VO activities at the sites
- The display: Gridmap
  - Structure of the Gridmap
  - Layout of the first prototype
- Current status of the activity and open issues

Monitoring of the activities during the CCRC08

- After the CCRC08 experience, feedback from site administrators:
  - The main monitoring tools during this exercise were experiment specific tools. They are not straightforward to use for a person external to the experiment
  - Sites serving many experiments had to deal with very different tools
  - In many cases the sites could not understand well if they were contributing to the VO activity as expected
- The site administrators would like to be able to compare the experiment's view of the site contribution to the information they get from their own monitoring systems

Objective of this activity

Providing a new tool which should:

- from one unique console, give an overview of the overall status of the services in the site. This should be a tool easy to use, also for persons external to the VO, and which does not require a particular knowledge of each experiment ->AN OVERALL VIEW OF THE SITE FROM ONE UNIQUE CONSOLE
- This tool will extract information from the experiment specific tools (Dashboard, MonALISA, Dirac, Phedex) and will display it in a consistent way, providing links to the source of the information-> A HIGH LEVEL TOOL
- Display the information using the gridmap technology -> FAST DETECTION OF PROBLEMS THANKS TO THE COLOR AND SIZE OF THE MAPS

An additional requirement from sites is to have a clear definition of the targets from the experiments: this information should be displayed by this tool (but it has to be clearly defined and provided by the experiments)

Information work flow

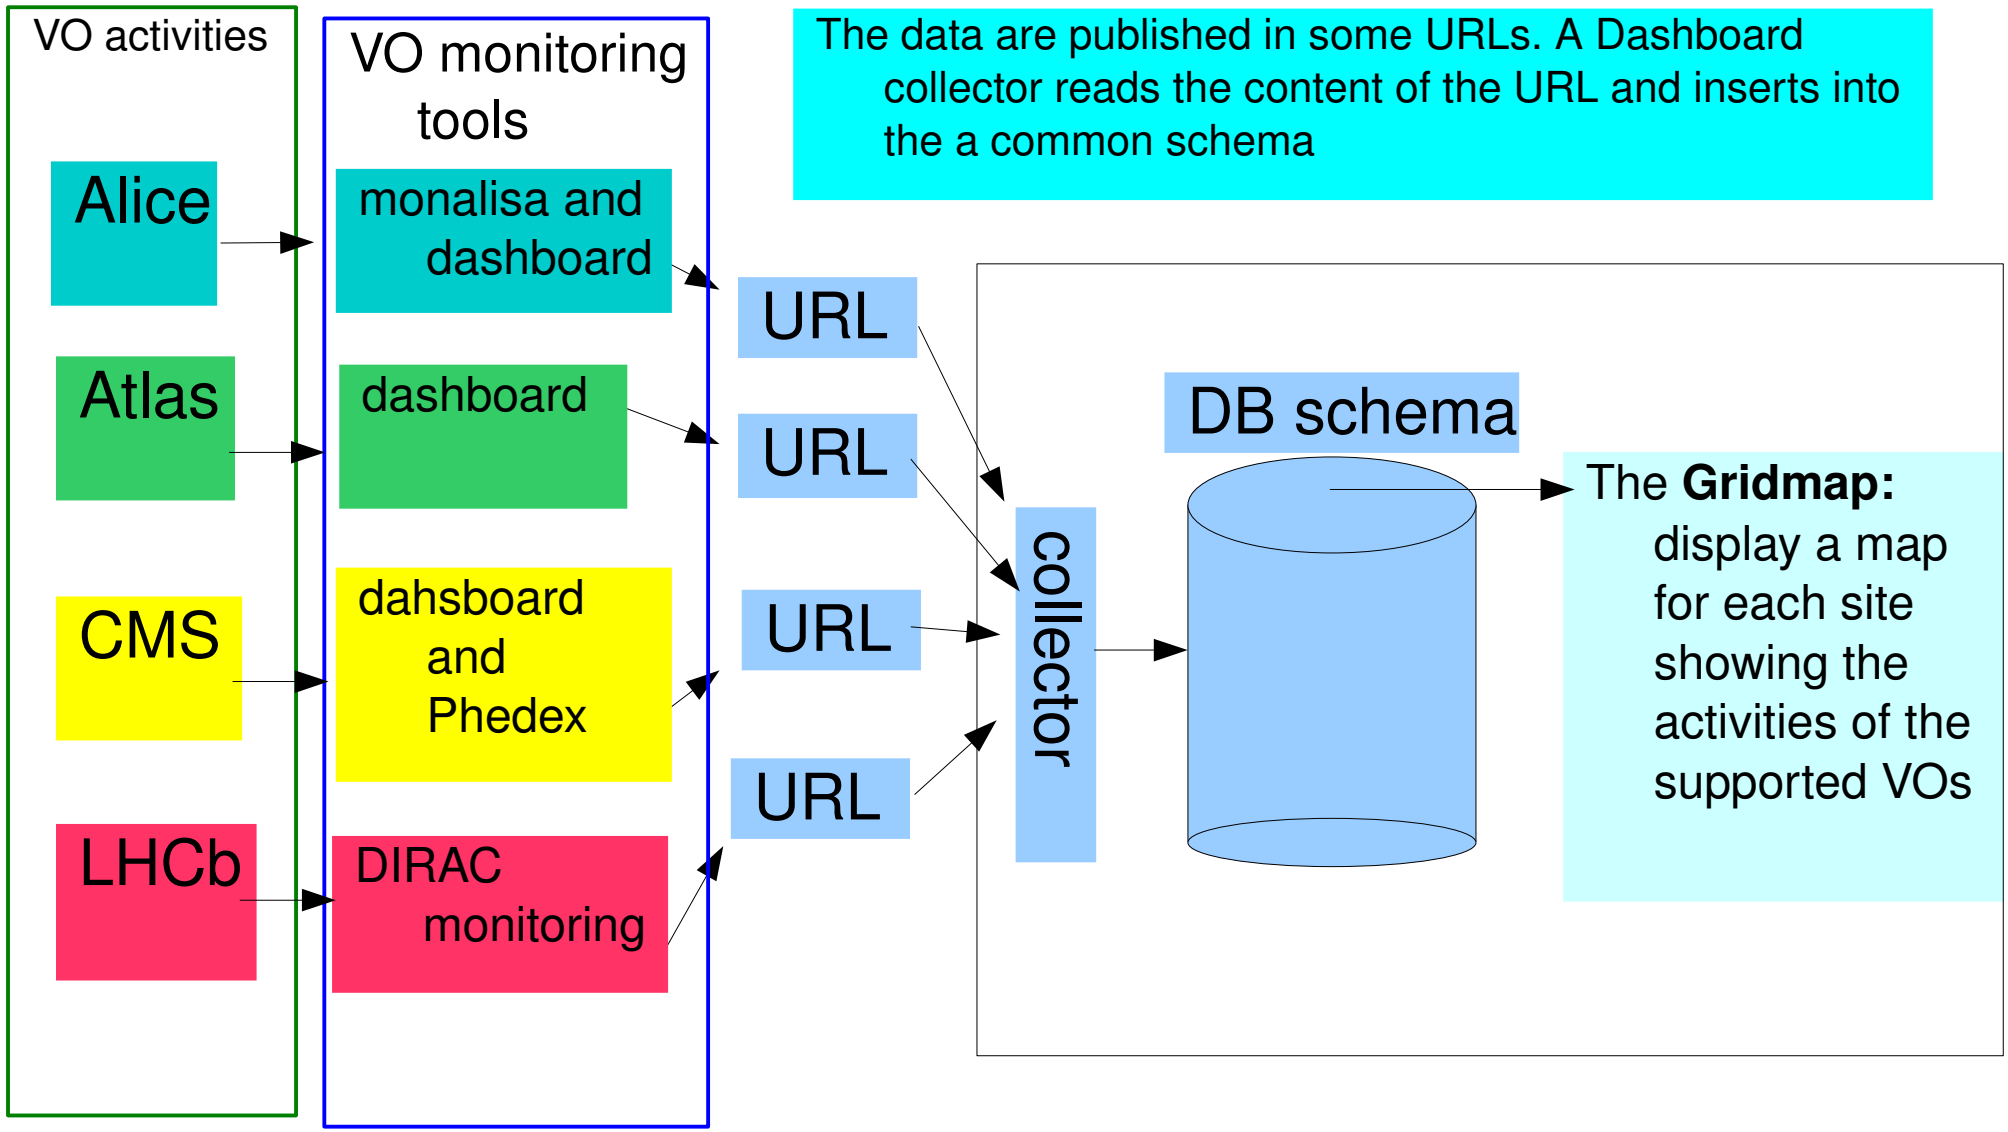

Objects to monitor: provided by the experiment specific monitoring tools

- The information about the experiment activities is retrieved from different tools:
  - Dashboard, MonALISA, DIRAC, Phedex
- We need to
  - Identify some common metrics for all the experiments
  - Agree on the definition of each metric
  - Verify which metrics are interesting from the site point of view and verify whether they are availale in the monitoring systems

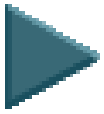

Have a set of metrics which have the same meaning for all the VOs

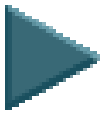

Build a data model and design a schema for the database -> a shared database for the data coming from the different experiments

Activities and metrics to monitor

ACTIVITIES

general site status

job processing

data transfer

...

METRICS.

Number of parallel jobs, Number of completed jobs (in the last hour)

wall time and CPU time (possibly in KSI2K)

...

Success rate (average in the last hour)

Transfer rate (average in the last hour)

...

There are 2 main activities: job processing and data transfer. Their metrics are collected for the 4 experiments

For each main activitiy there are some secondary activities: they are optional and depend on the experiment computing model

Job Processing: MC production, data reconstruction, user analysis, SAM tests...

Data transfer: MC production transfer, transfer t0-t1, transfer t1-t1...

The collector

- A new collector has been developed in the framework of Dashboard, in analogy with the collector of the Site Status Board module
- The collector periodically (every hour) reads the metrics published in the URLs by the experiments and inserts the values into the DB schema

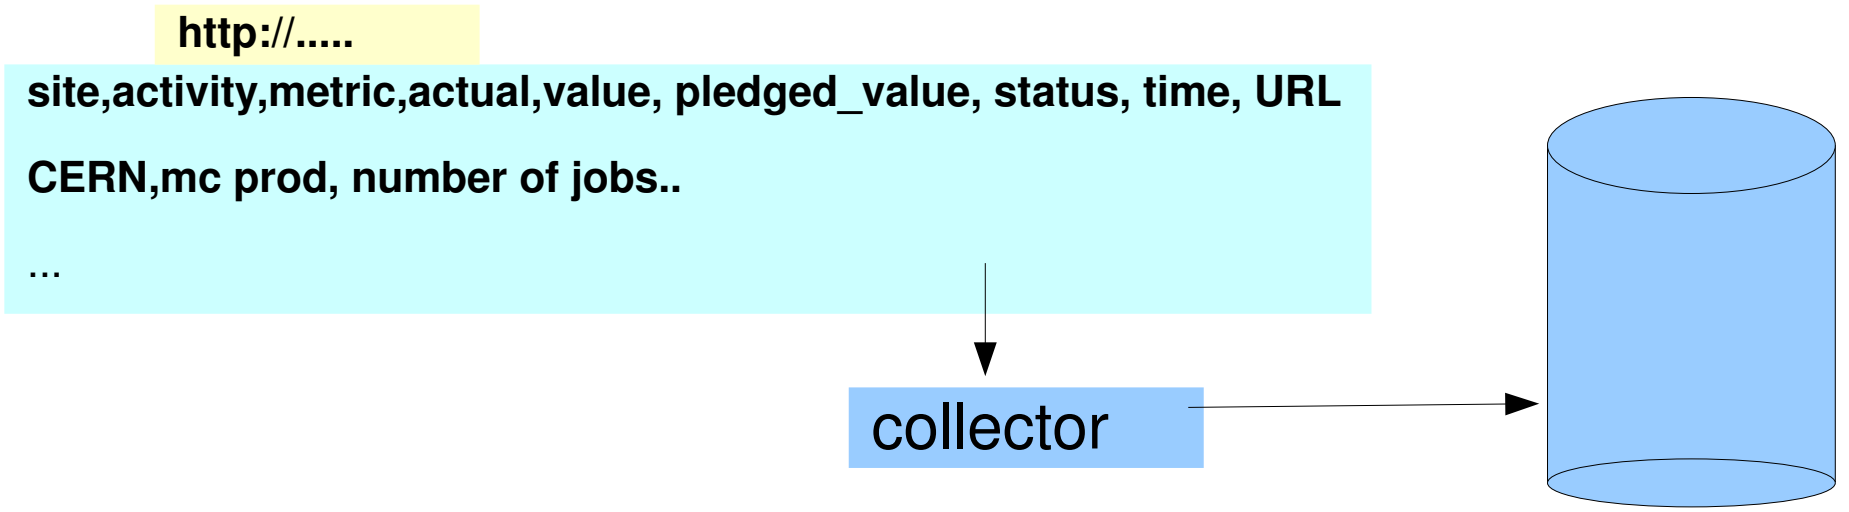

Database schema

Implemented in a development instance of Oracle.

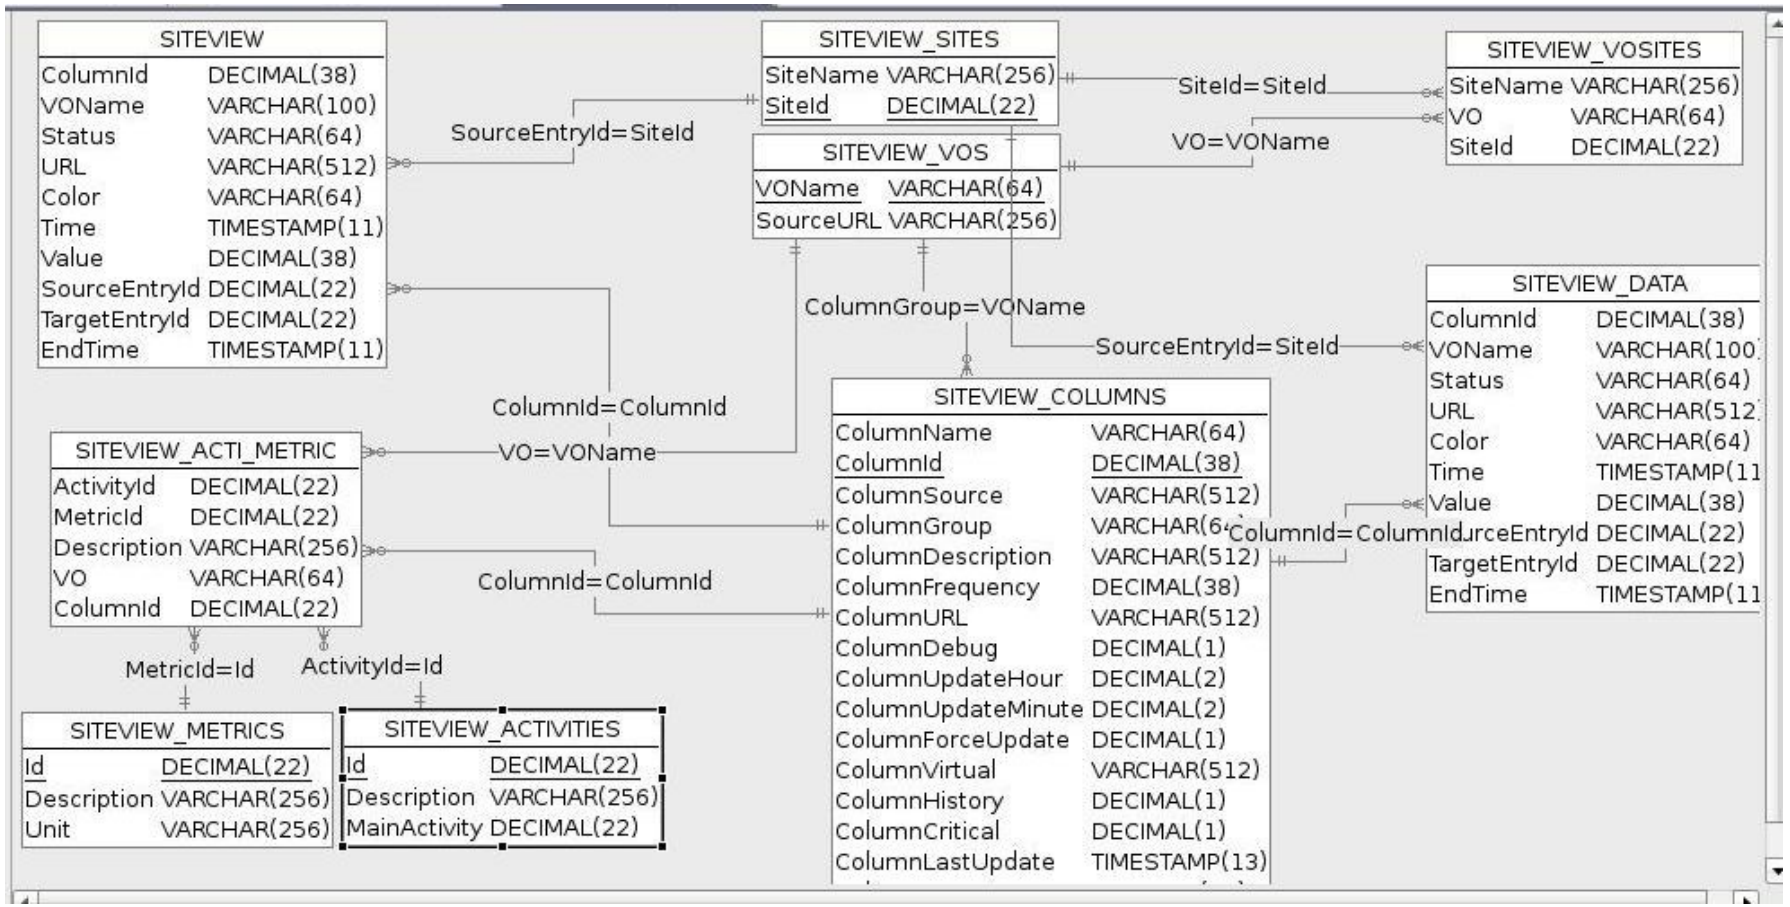

Very flexible: any other activity and metric can be added later, according to the requirements of the VO

The gridmap

- A gridmap for each site
- Inside the site gridmap, the hierarchy is:
  - first level: Activity
  - second level: VO

The main map only shows the main activities, for all the supported Vos

The size of the rectangle should be proportional to some metrics (i.e. number of parallel jobs). A constant offset is added, in order to have all the supported Vos visible, even if they are running fiew (or no) jobs.

The status (=COLOR) should be provided by the VO. Alternatively, it can be computed as a success rate between the measured value and the pledged value. Inthis case the VO has to provide the pledged value

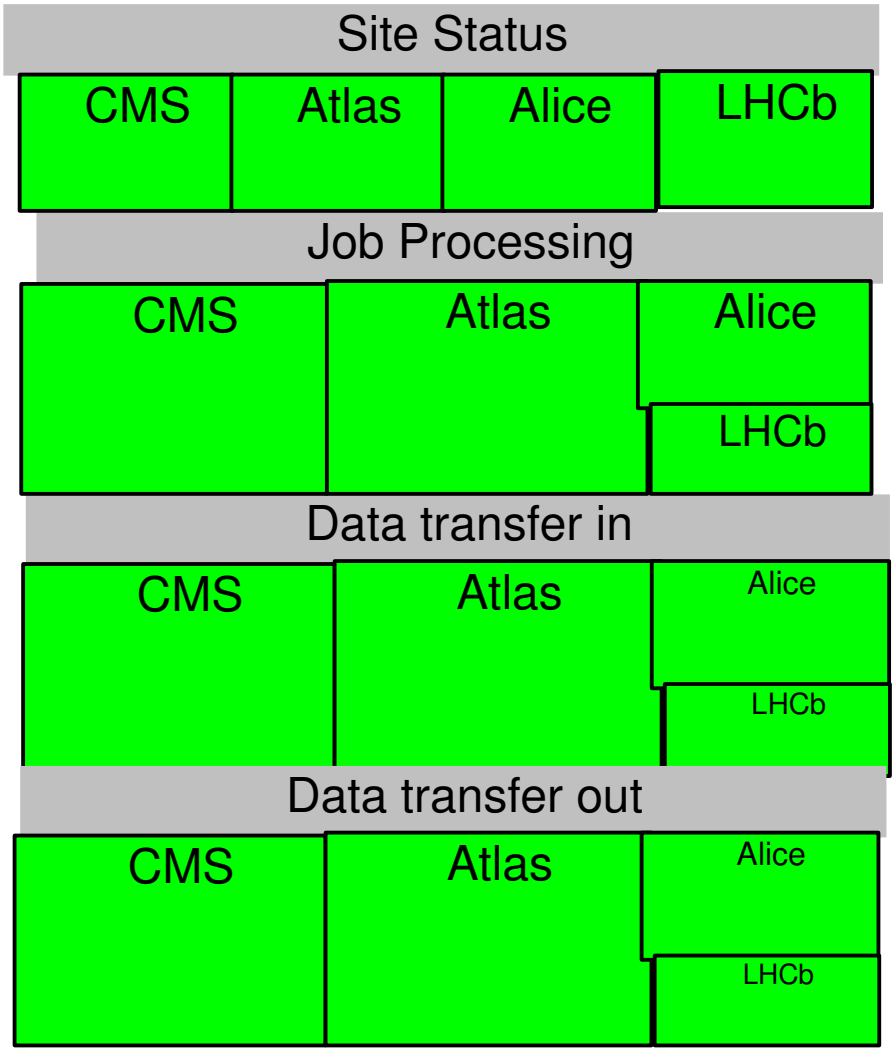

The site status

- It is an overall evaluation of the site status from the VO perspective. If the status is green, then the site is ok for the VO, even if there is no job running
- For LHCb and Atlas it is computed on the basis of SAM VO specific tests
- For CMS it is not only SAM tests, but also production success rate, only taking into account the failures which by exit code probably point to site problems, and wheter the site is visible in BDII
- For Alice it is extracted from Monalisa
- **Important**: if the status is red it does NOT necessarily mean that the responsible is the site. Even if the site status definition is as much as possible related to the site, it can happen sometimes that a site turns red due to problems of the VO. Even in this case it is useful to notify the site about the problem.

A screenshot of the prototype for FZK

- In the main map only the main activities are displayed, for all the supported VOs

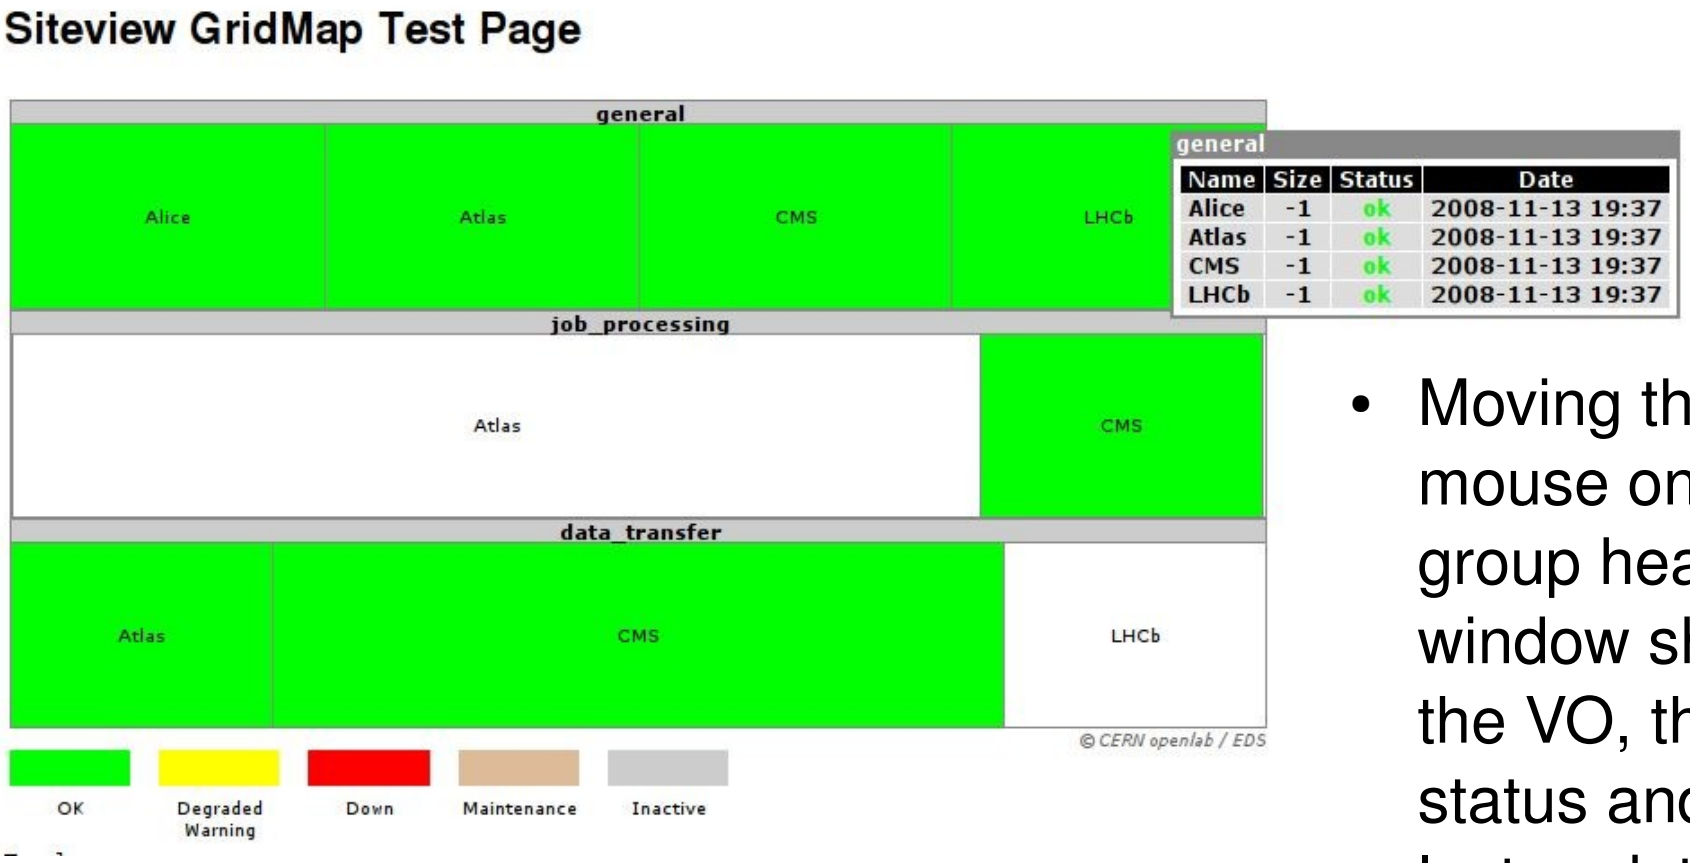

- Moving the mouse on the group header a window shows the VO, the status and the last update

Link to the site status for Alice

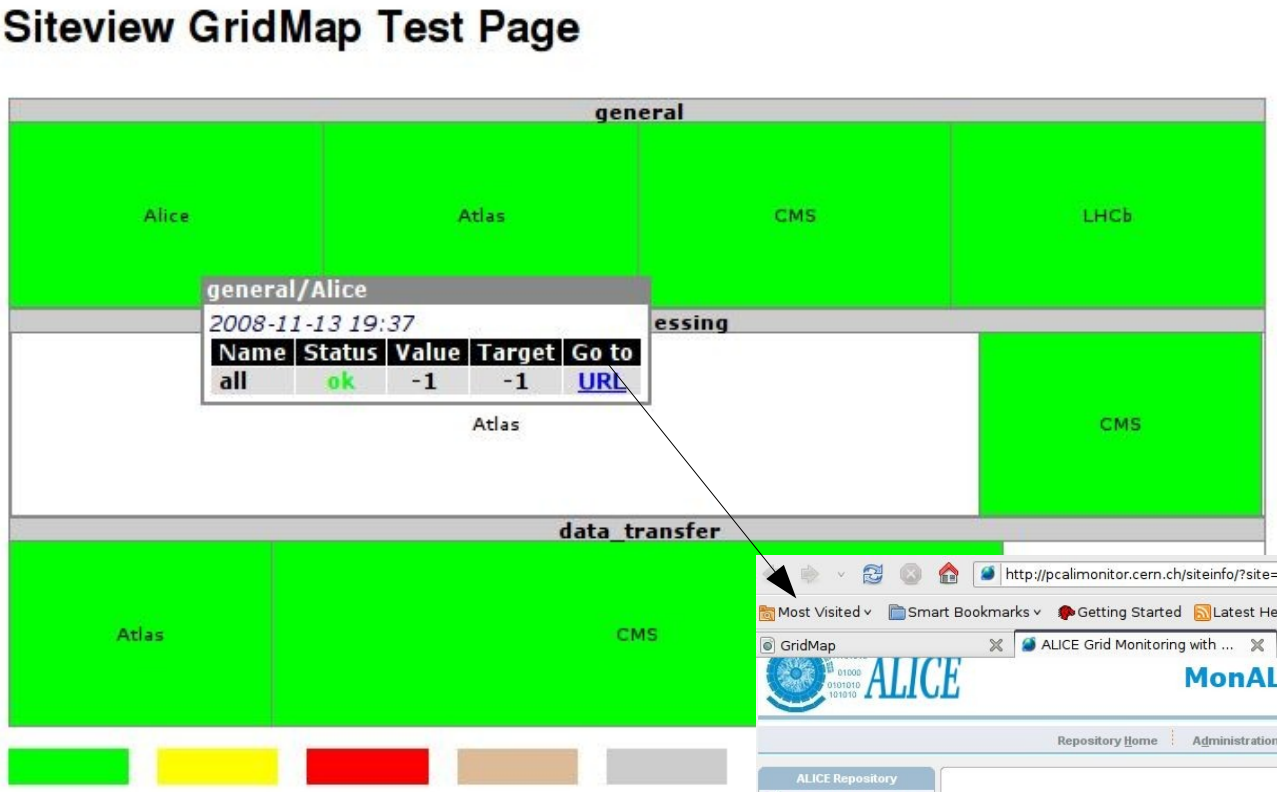

The context help provides a link to the source of information

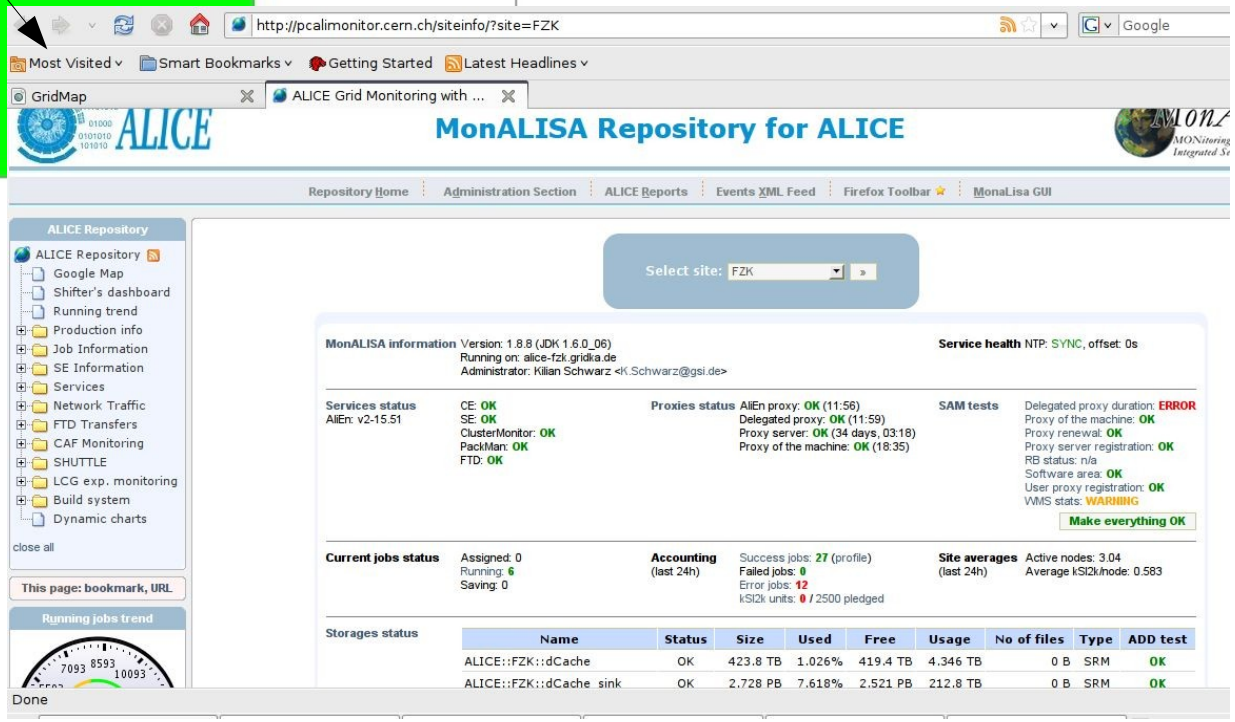

Link to the site status for CMS: Site Status Board

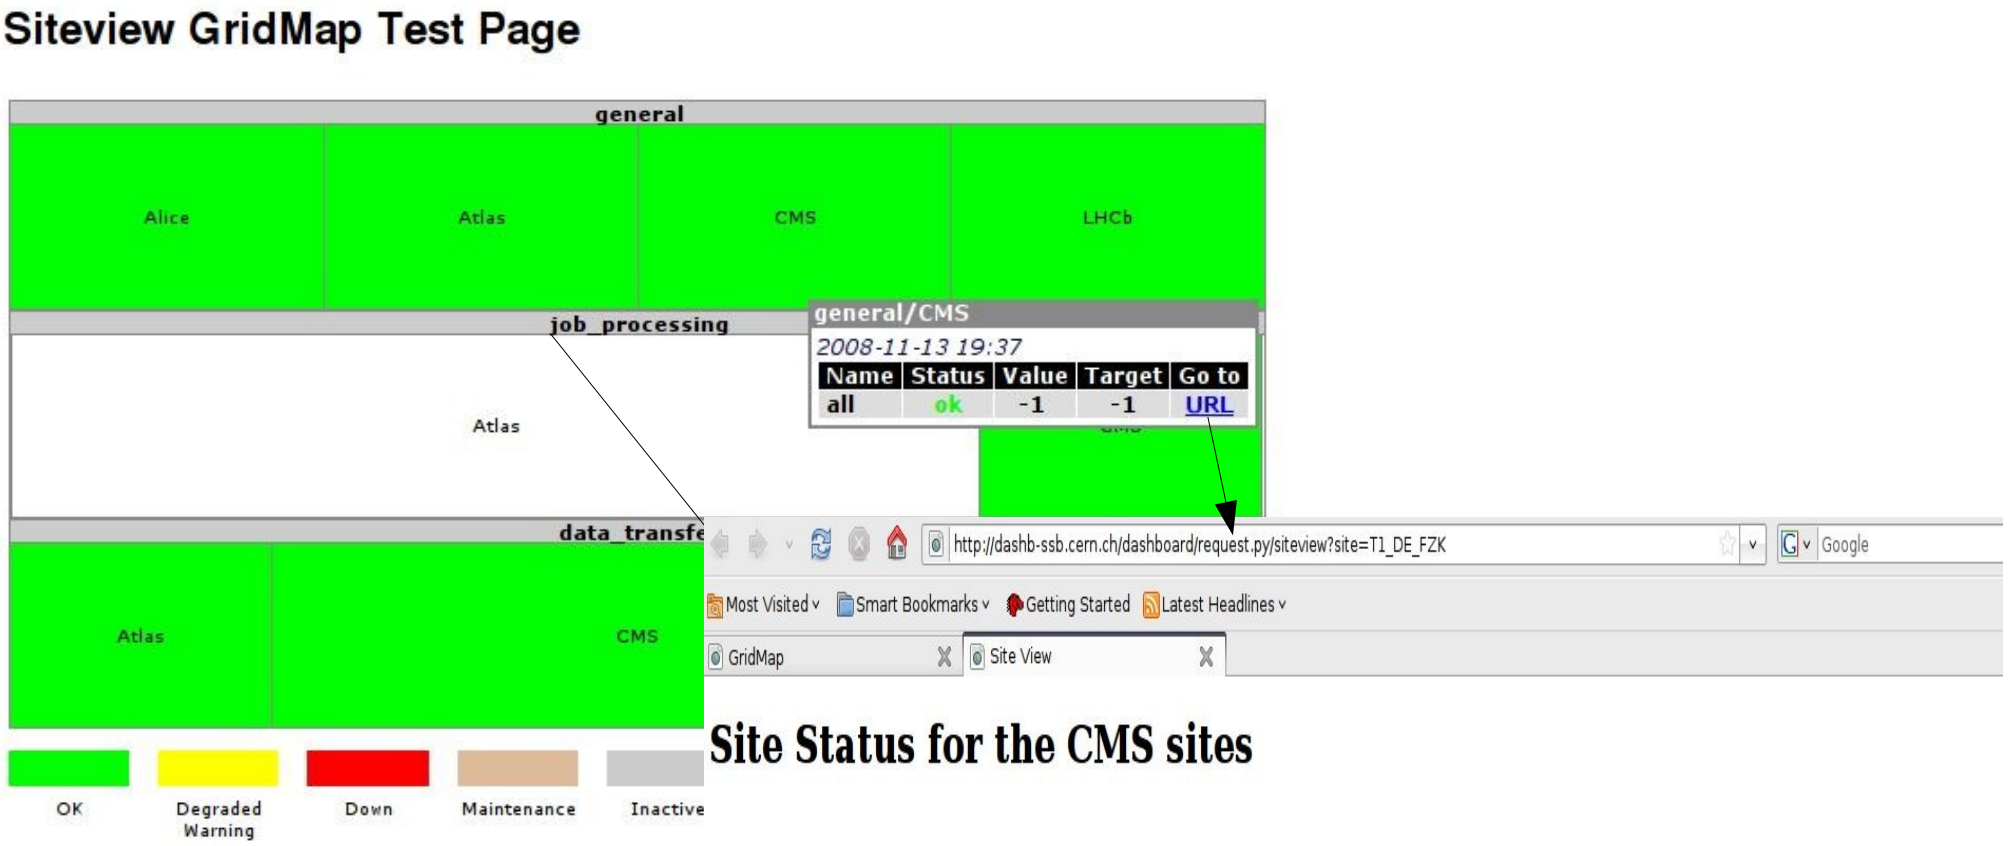

From there all the information and links are available

Site status link for Atlas and LHCb: VO specific SAM tests

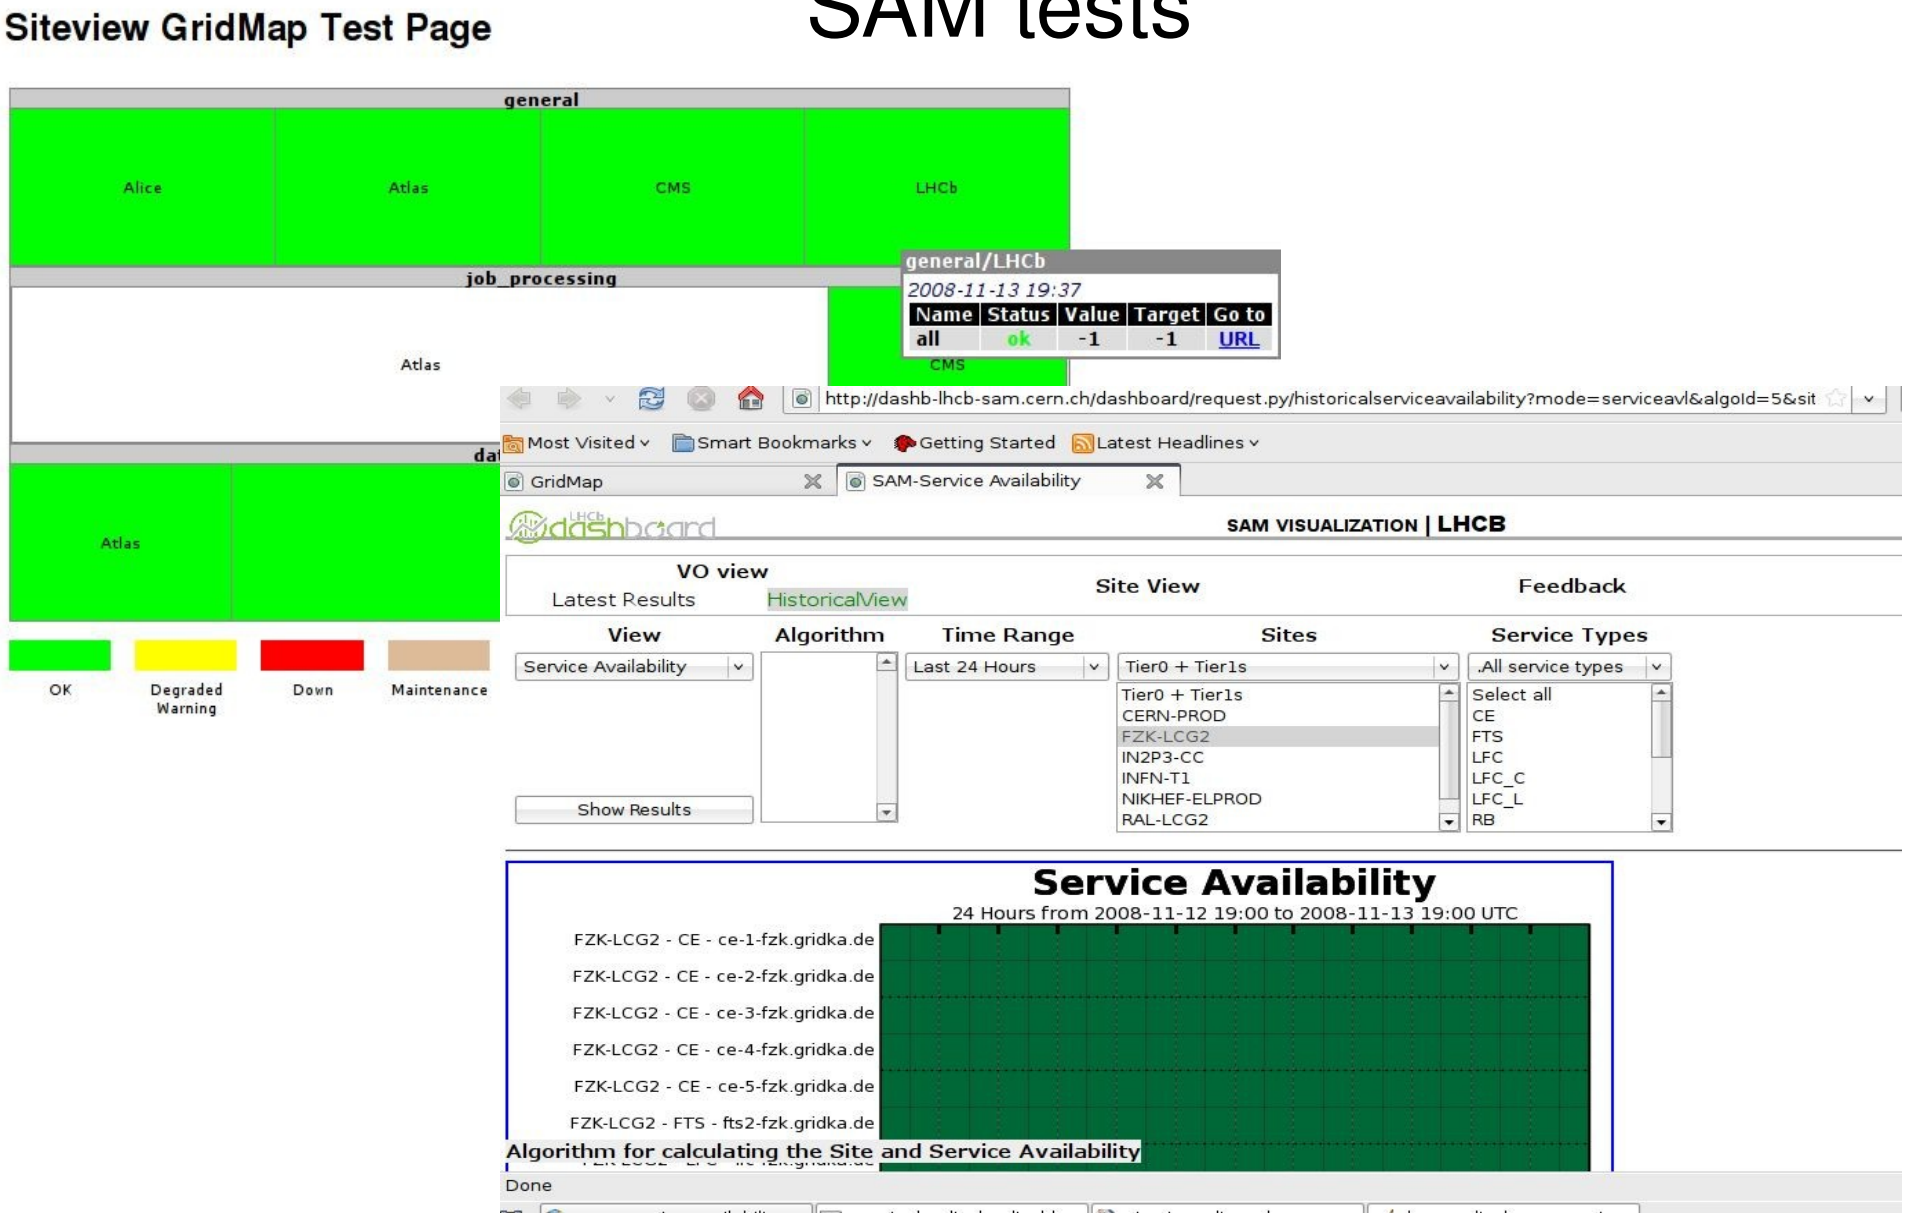

Job processing activity

- Moving the mouse on the header: information about the VO which are supported by the site.
- The number of running jobs is displayed (this is the parameter which determines the size of the rectangle), the status (which determines the color), and the last update time

- **Important**: if the status is red, it doesn't point to a problem of the site. It is a problem of the VO activity

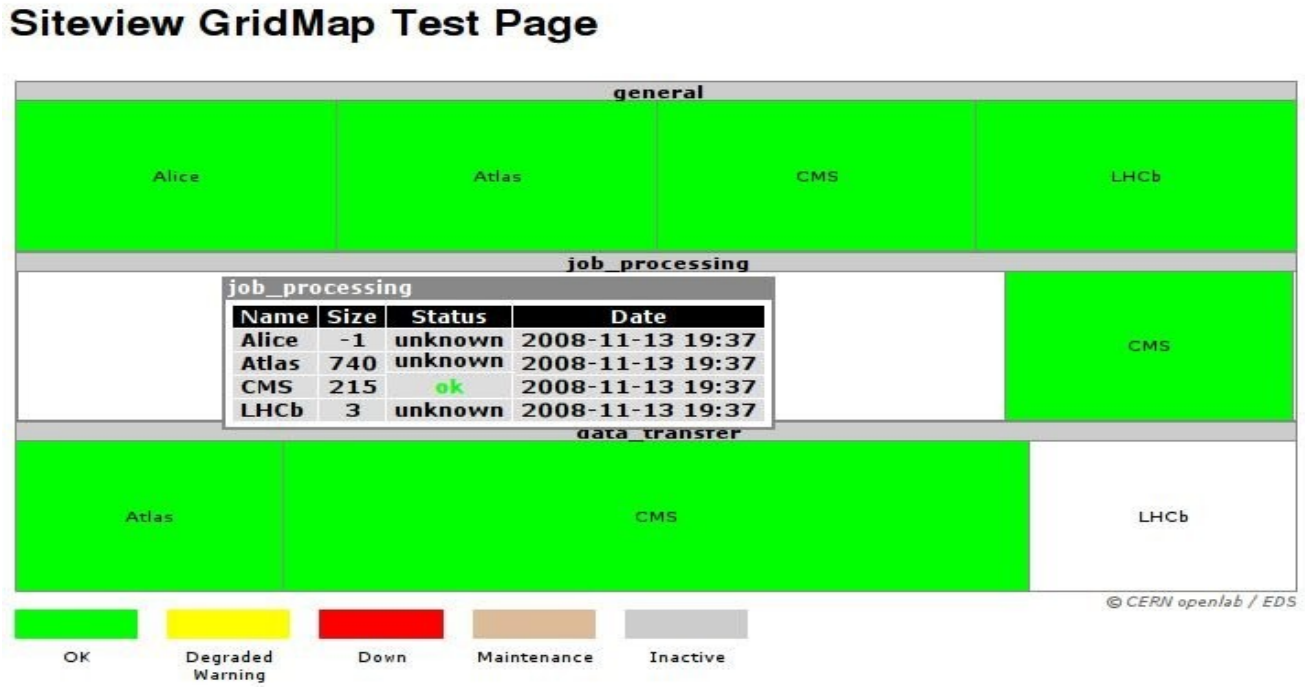

Job processing: Atlas

- All the metrics relative to the job processing activiy are displayed
- A second level gridmap show the different types of jobs (will be shown later)

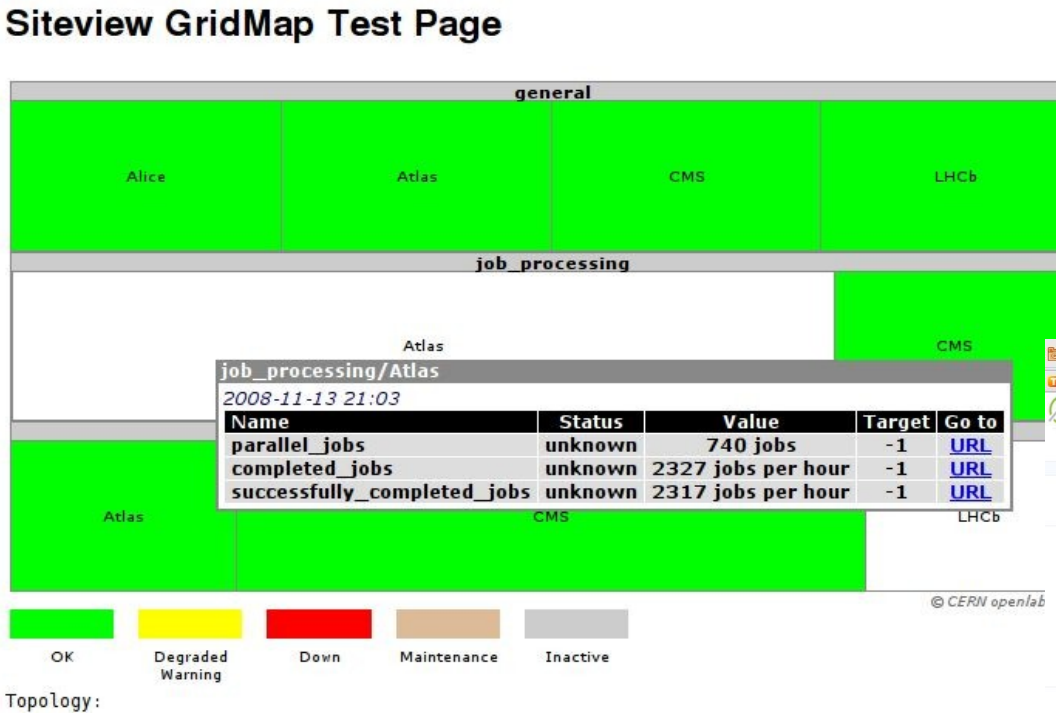

Job Processing: Alice

- All the metrics relative to job processing are given.
- The wall time is given, **normalized in ksi2k** together with the pledged value.

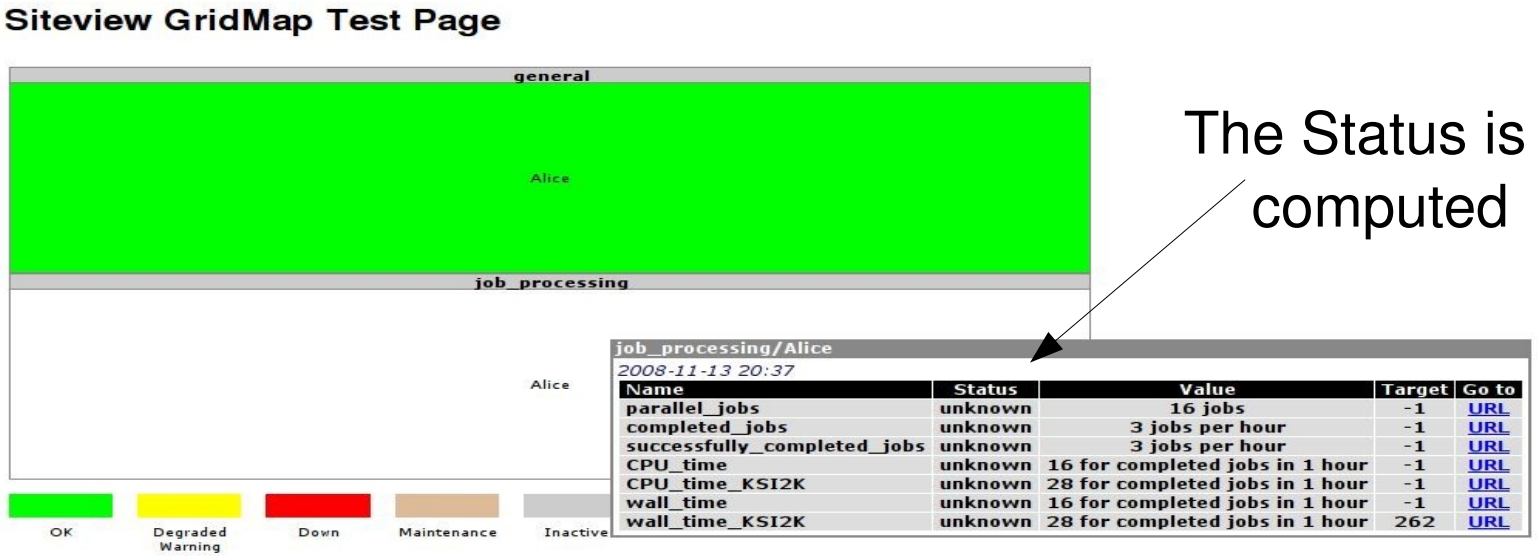

The Status is not computed

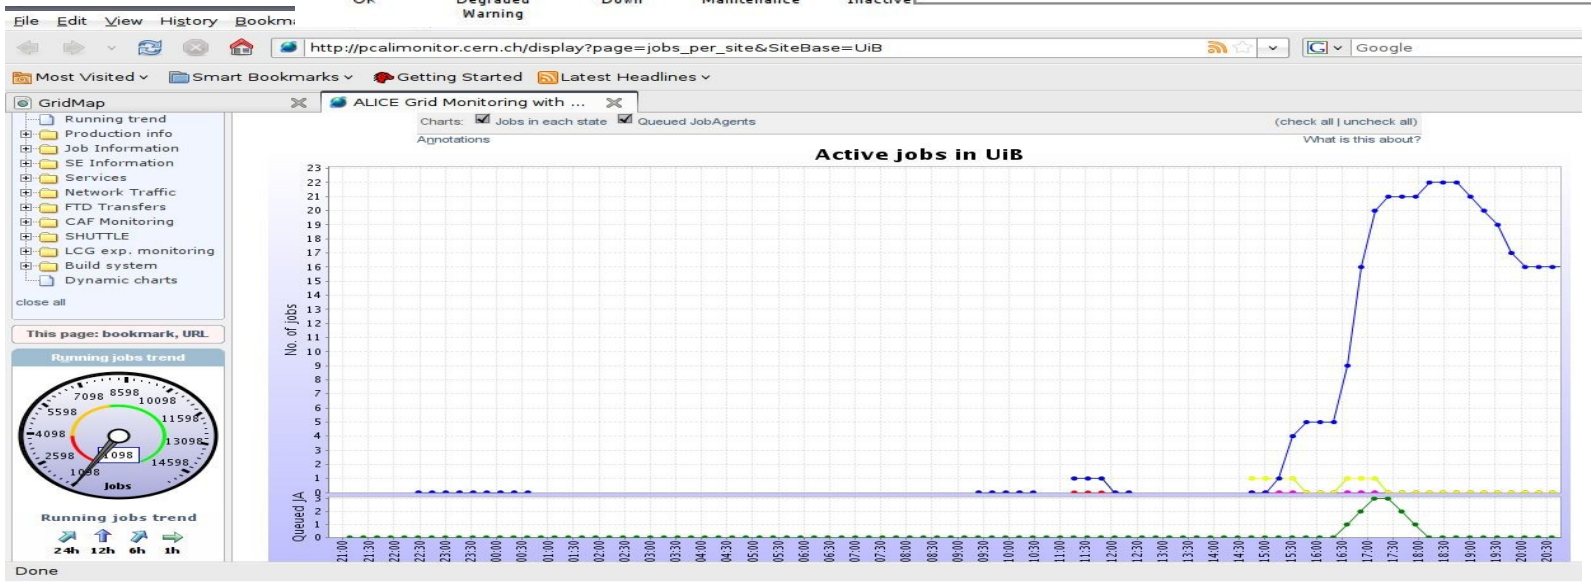

The link in the context help redirects to Monalisa repository for that site and activity

Job processing LHCb

- Job processing LHCb. Link to the Dirac WMS history plot

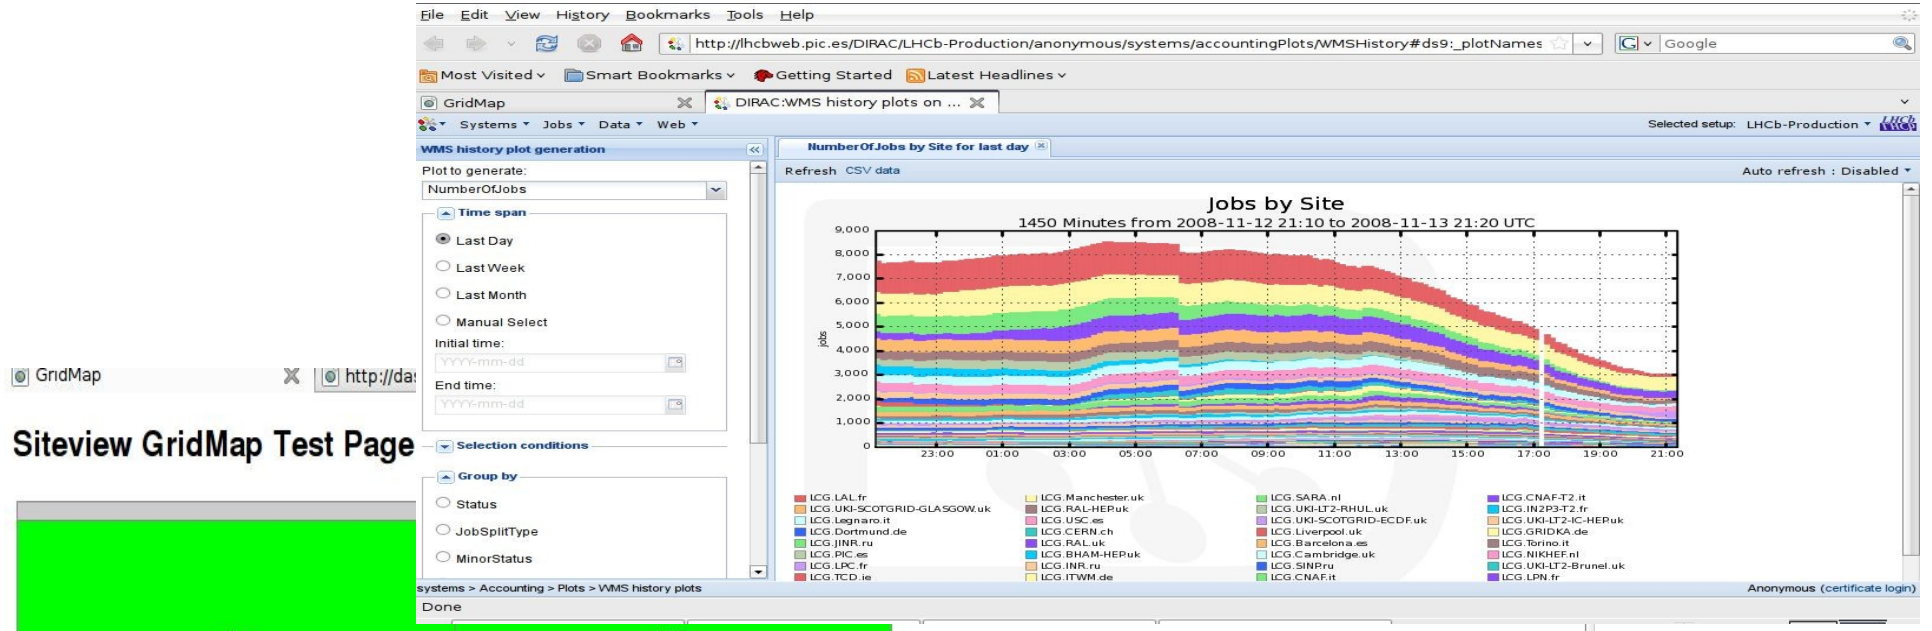

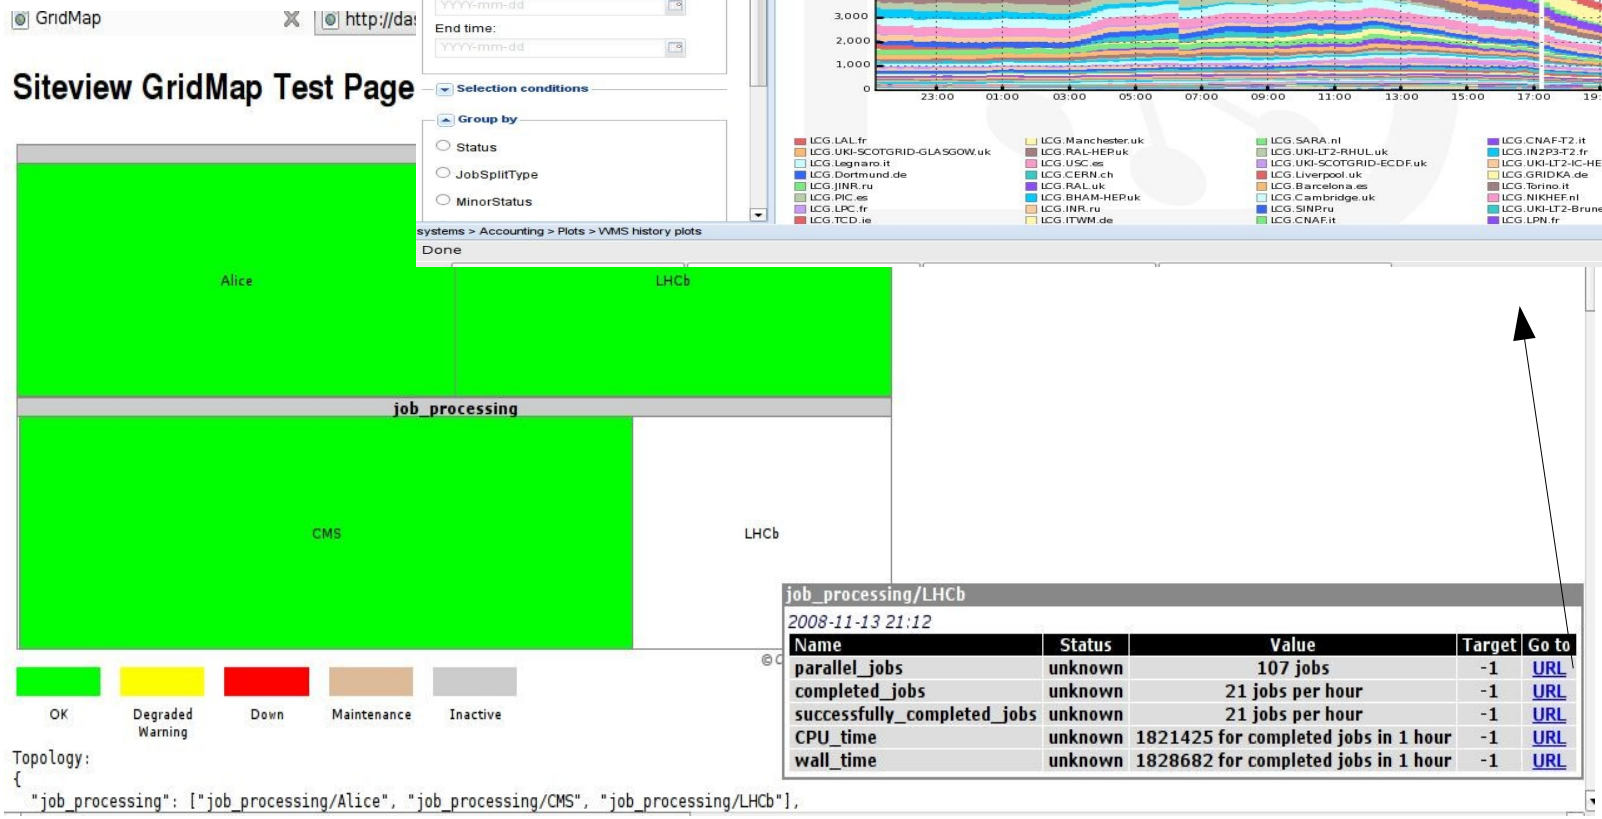

The time is given in second (not normalized)

No pledged value, nor status

The second level map

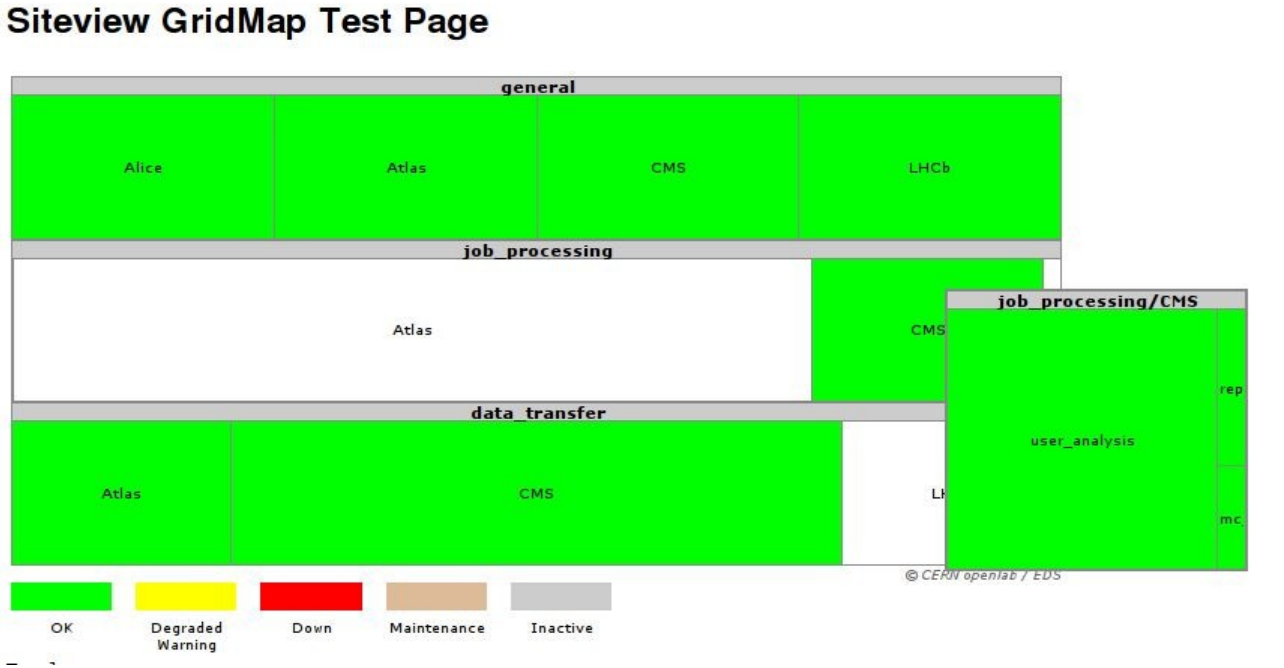

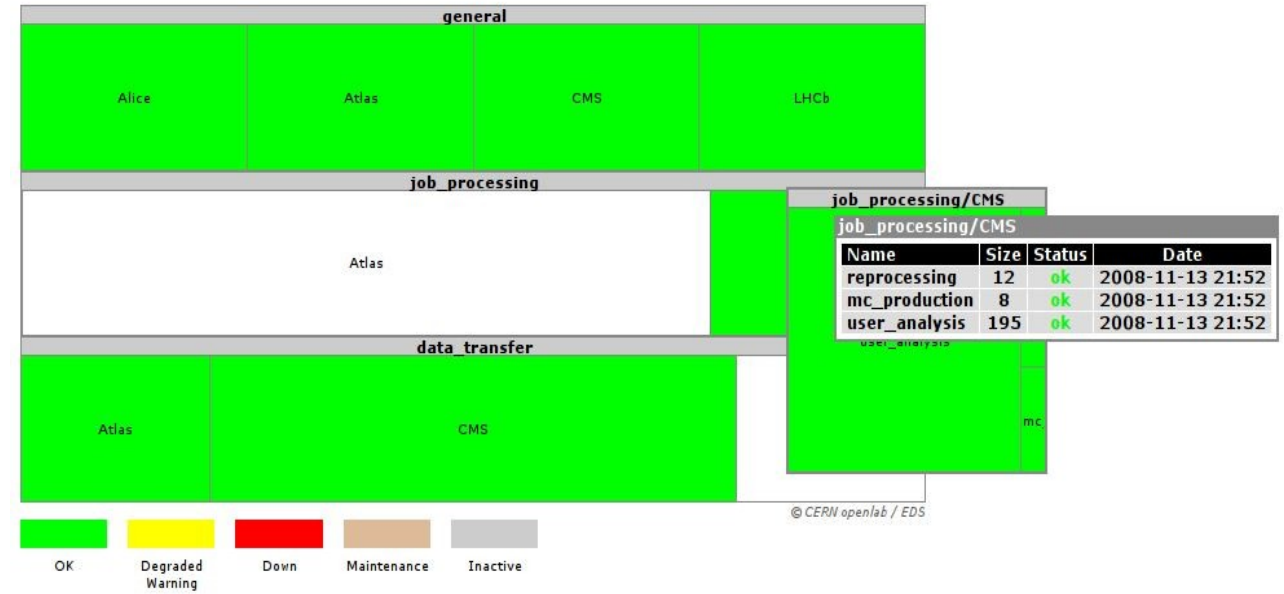

- For ex. Job processing activity for CMS.
- Clicking on the map, the sub activities are show. For ex. CMS has reprocessing, MC production, user analysis
- The context help shows the number of jobs for every sub activity

Data transfer

- It is envisaged to display the total incoming and outcoming rate
- The submap will display the network traffic channel by channel

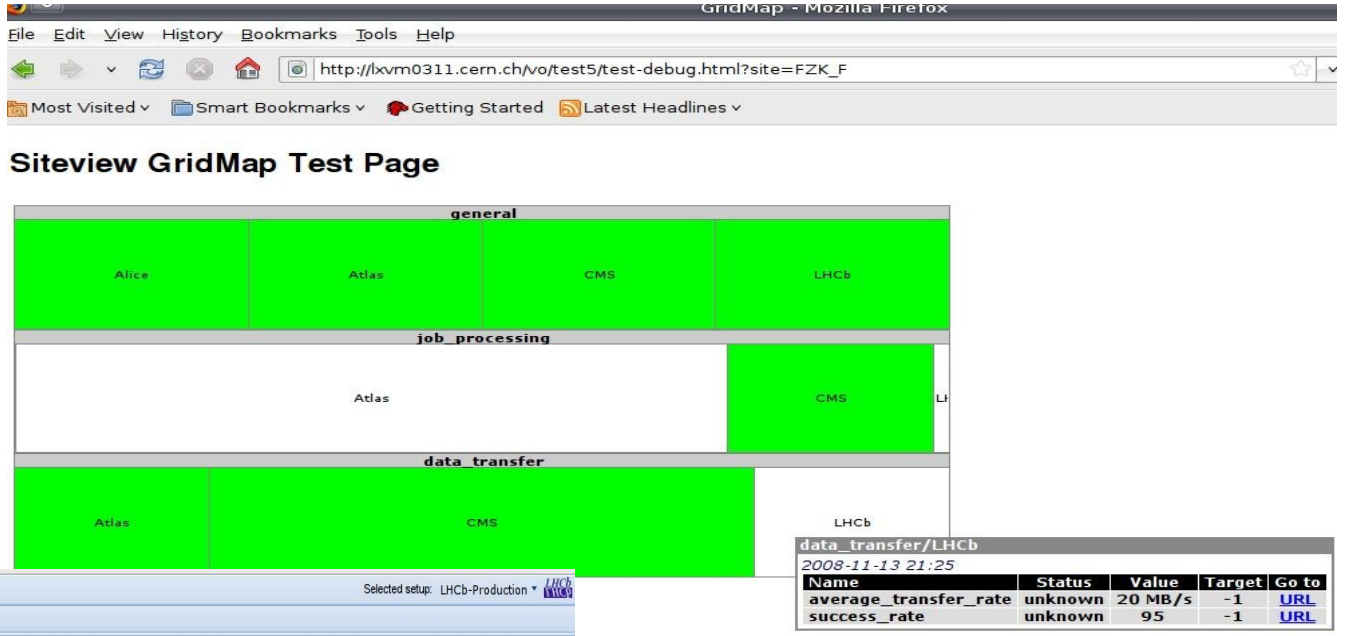

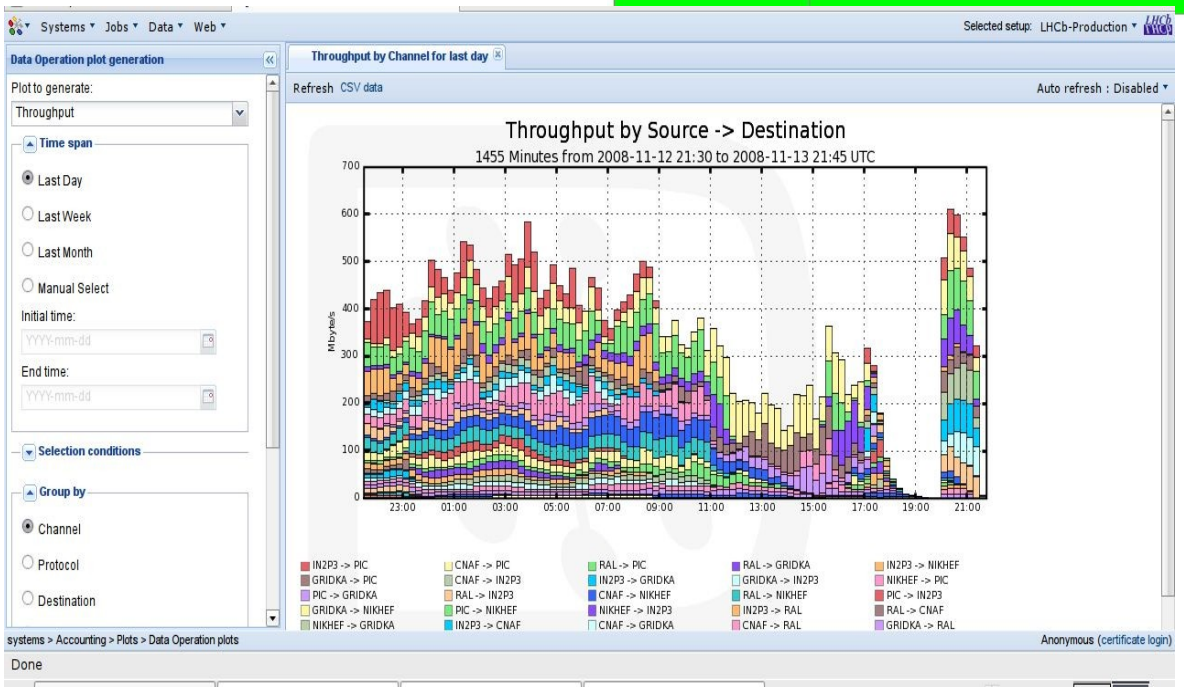

- Implementation still ongoing...

Data transfer for Atlas

- **The status and pledged value are given**
- The status is given for the average transfer rate comparing with the pledged value: it can be ok or warning
- And for the success rate the status is defined following the rule:

- SR<20% : critical
- 20<=SR<50%: warning
- 50<=SR<80%: normal
- SR>80% good

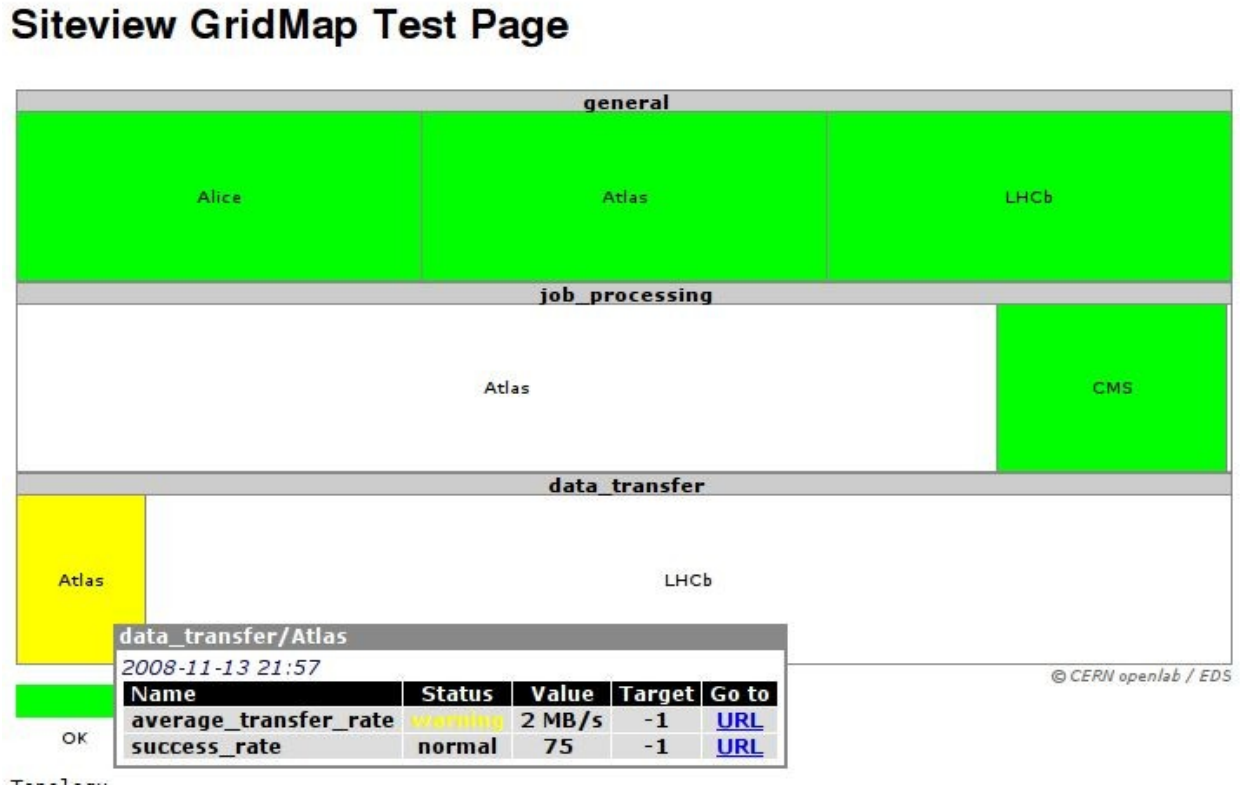

Current status and plans

- A set of activities and metrics has been defined
- The design of the database schema has proved to work fine (still some improvement to do..)
- The collector has been implemented and is currently working and populating the database schema
- The Vos have already provided several metrics. Still some metrics are missing
- In many cases the expected values and the status are not provided by the VOs

Easy part of the job:the tool has a very simple structure

Not that easy...

Difficult, but feasible

Plans: first complete map by the end of the year for pilot use

Open issues

- In many cases it is difficult for the experiments to provide the expected values for the metrics and to provide a status of the activity
- Sometimes the expected value for a given metric is not available, or it is available but it is not in the same database or system, so the observed value cannot be related with the expected one
- In general, experiments are reluctant to assign a status to the activity at a given site.

Some remarks:

If the activity status at a site turns red, it doesn't mean that the site is responsible. It means that the VO activity is having problems.

The site status is given by the 'site status' group (upper rectangles in the Gridmap): if that is green, then the site is ok

This is only an informative tool which should streamline the information work flow and help site administrators and Vos to detect problems in an easy and efficient way

Acknowledgements

- This project can progress only with the collaboration of the experiments which publish data to fill the common database schema
- A fundamental contribution from: P.Saiz (CMS), C.Grigoras (Alice), R.Rocha and B. Gaidioz (Atlas), A. Casajus and A.Tsaregorotsev (LHCb), W.Ollivier (LHCb and Atlas).