

# What's New in Logging and Bookkeeping

How does the instance presented here differ from your site's L&B server...

## Sneak Preview of L&B for Monte Bianco

The L&B server used in this demonstration at the 2012 EGI Technical Forum is a **preliminary version of L&B** that will be released with EMI-3 Monte Bianco in late 2012. It is an evolution of the previous version, L&B 3.2, fully backward compatible on protocol level. There are several new and interesting features, introduced below. In most cases they build on features that have already been available earlier as a preview or limited functionality.

## Virtual Machine States

In addition to previously supported job types including gLite, native CREAM, PBS/Torque or sandbox transfers, a new state machine has been implemented to follow the states of virtual machines. It is loosely based on the Open Nebula state diagram and assumes that events will be sent to it from different sources on different levels such as the virtualization resource manager, virtualization hypervisor or the virtual image itself. Information from those various sources is combined to acquire a most up-to-date and detailed picture of the virtual machine's current state.

## Inter-Job Relationships

Another extension of a principle already present in previous versions of L&B, there is now support for logging mutual relationship between *any* two jobs of any type. Thus, on top of linking collection or DAG subjobs to their parents, or sandbox transfers to their respective compute jobs, any two jobs in general may now be linked. The nature of the relationship is discernable from the types of the two jobs. For instance – as shown in the demonstration – compute jobs can be linked to virtual resources used to run them.

## Job Queries through URLs

Although the ability to use a Web browser to contact L&B and peruse its information has been there for a long time, users could only list all their jobs, unable to filter the output. Now L&B recognizes specific query strings in the URL, allowing you to specify job query conditions and have the L&B server list only Job IDs matching your conditions. For instance, a query string such as this can be used:

```
?query=jobtype=simple&status=running|scheduled
```

It will list only jobs of *simple* (single gLite) type, whose state is either *running* or *scheduled*.

Similar to query conditions, flags recognized by the API are also applicable to URL queries. For instance the following query string will request that child state histograms are recalculated for all collections included in the output, and that job states are populated with classad data (job descriptions):

```
?flags=childhist_thorough+classad
```

## Branded HTML Output

Sites may now decide to style their L&B's HTML output to fit their institutes' visual styles. An HTML header file can be passed to the L&B server. It will then be included between the `<style>` and `</style>` tags of any HTML page generated by L&B. Several style classes are recognized and contents are formatted accordingly.

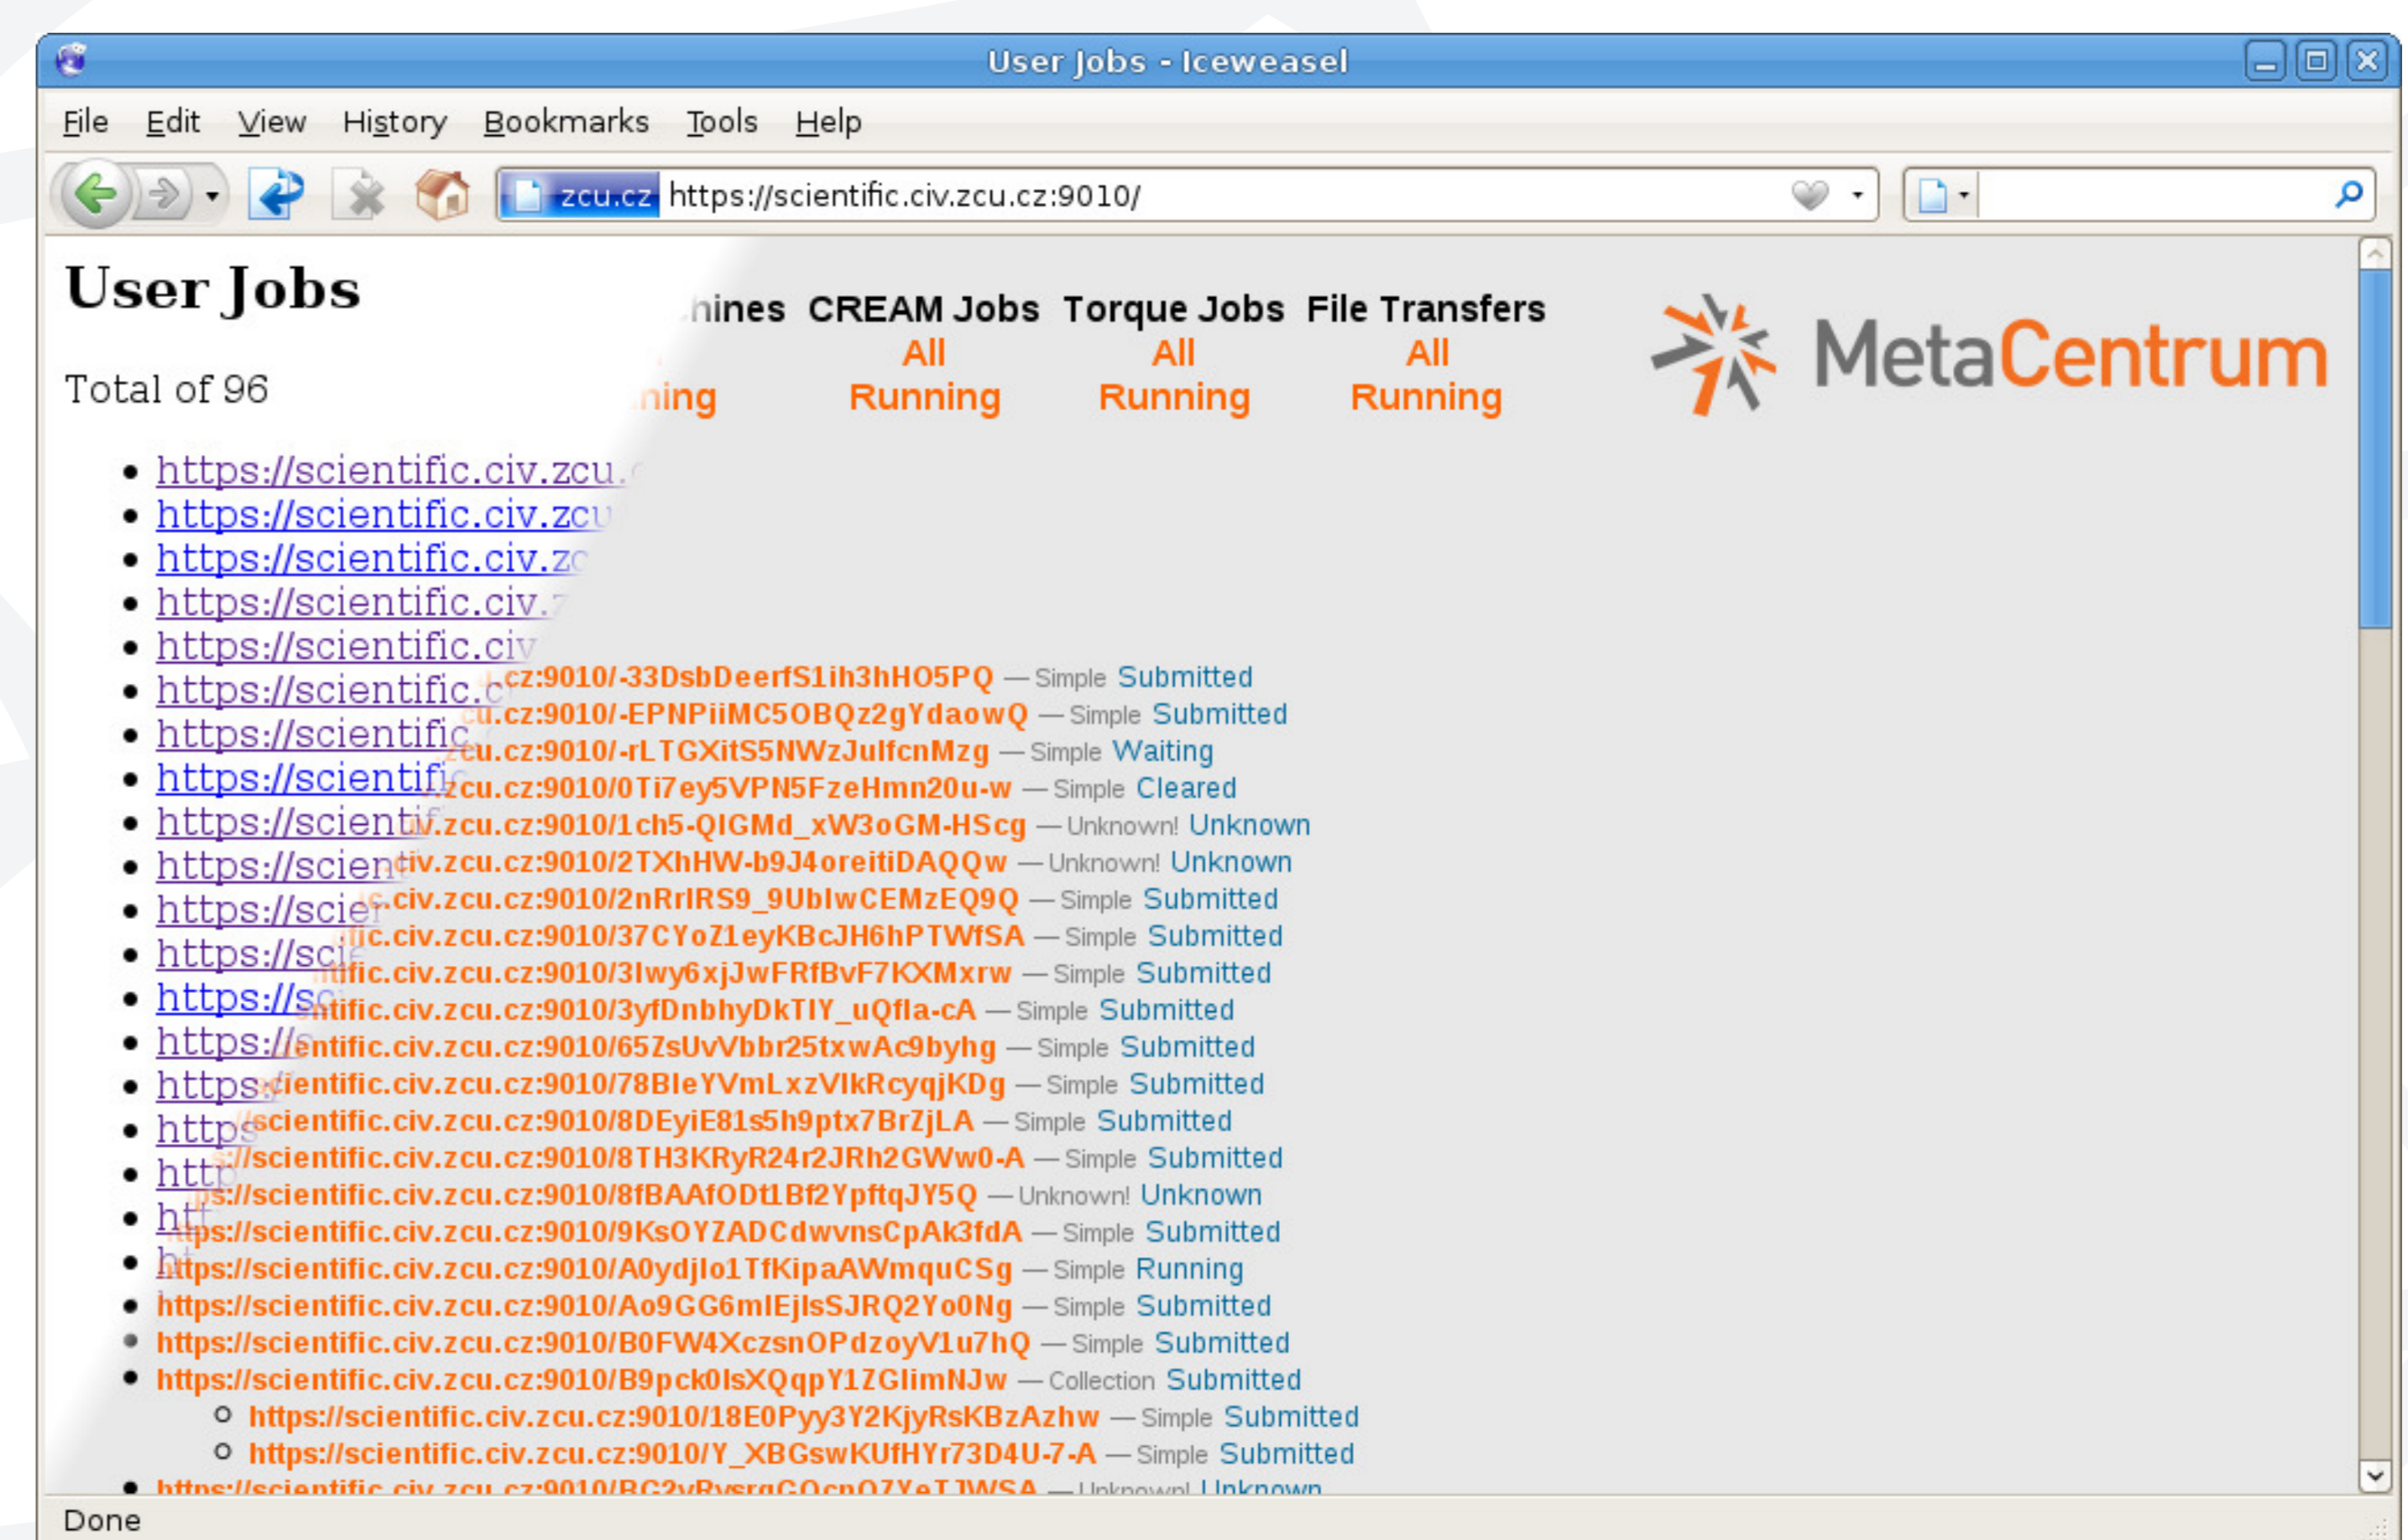


Fig. 1: Traditional vs. branded visual style in L&B

This feature is useful not only for adjusting colors and adding logos, but also for adding page headers with pre-determined common links to the L&B server or elsewhere.

## Alternative Authentication

Relying partly on new functionality to be included in EMI's Common Authentication Library (caNI), L&B supports multiple authentication mechanisms simultaneously: namely traditional proxy certificate-based GSI authentication and Kerberos GSS. Since that allows users to access L&B with multiple identities, a Globus-style grid map file is used to tie those together. This allows for instance Czech national grid users, who rely on Kerberos authentication, to log events for their jobs and then use certificates in their browsers to authenticate before viewing job details.