

Accounting Training: Regional Accounting Portal

Álvaro Simón
Iván Díaz

CESGA





Outline

The Accounting Portal

Regionalization

Portal Structure

Portal Deployment

Conclusions

- The Accounting Portal is the central tool for viewing EGI accounting information and reports.
- It also houses the WLCG and InterNGI usage reports.
- The CPU and User accounting is received daily from APEL.
- Other sources of information provide information about VOs, user roles, Grid topology, etc.
- It is developed, maintained and hosted in the CESGA.

- Accounting of CPU time (Normalized or not), Wall Time (Normalized or not), Number of jobs, and efficiency.
- Grouping by Country/Region, date, VO, usage by country.
- Selection of LHC VOs, TOP10, custom.
- Local and grid job separation.
- Tier1, Tier2, Country, OSG, Unregistered, VO discipline, VO Metrics and Custom Views.
- VO Manager, VO Member, Site Admin, User views, powered by UserDN data.
- WLCG and InterNGI reports.



Outline

The Accounting Portal

Regionalization

Portal Structure

Portal Deployment

Conclusions

- The regionalization of most Operational Tools was one of the tasks defined on the EGI DOW.
- Part of the transition of the infrastructure to a Federated organization.
- Several of the operational tools were successfully regionalized.
- The complexity and interdependences of the Accounting infrastructure has slowed down the regionalization progress.

- The regionalization of the Portal means having local slimmed down versions of the APEL server and Portal for NGLs that wish it.
- Those instances don't need functionality like country/region reporting or WLCG/InterNGI usage reports.
- They would also only show a subset of the accounting information.
- In the future specific regional functionality could be developed, but that is currently out of scope.

- The regional version of the Accounting Potyal shows much less data and it is easier to interpret those data.
- There is no possibility of seeing foreign data by error when reporting, and there is no need of filtering them.
- Load can be diverted from the central instance on some scenarios, and responsivity can be improved.
- The regional instance can validate the data from the central instance, and local staff can cross check those with their APEL installation.



Outline

The Accounting Portal

Regionalization

Portal Structure

Portal Deployment

Conclusions

- The Portal is a web application with a strong legacy component and extensive and complex functionality.
- It is a LAMP application, that is, with PHP and MySQL servers on a CentOS server.
- Both servers are on the same machine, for the moment the user load is low enough.
- It is currently implemented as a virtual machine, so it can migrate seamlessly to better hardware.

- PHP 5 is used to serve the pages and create the graphs (using internal requests).
- The code was modularized and simplified, so all the views share common code with small changes.
- PHP also generates CSV and XML listings for external operational tools and users.
- There is no framework use or ORM adaptation for the moment.

- For the relational database back-end we use MySQL 5.1.
- Many of the table schemas are determined by APEL and other operational tools, but there is extensive index optimizations, since writes are very sporadic.
- Apart from the reporting views, which have to process millions of lines and have appreciable delay, most views are generated quickly.
- Since the number of entries grows steadily, MySQL could be migrated to a dedicated server in the future.



Outline

The Accounting Portal

Regionalization

Portal Structure

Portal Deployment

Conclusions

- Since the initial deployment is very complex, the Regional Portal is provided as a pre-configured VM
- This VM hosts all the needed software and is tested against the regional server.
- Regional staffs only need to deploy the VM and integrate it on their regional network and naming schemes.
- The central staff would still offer support for bug fixing and would backport changes from the main portal.

- There are several configuration files to be modified:
- The Apache server must be configured with the correct IPs, DNS addresses, host certificates and local policies for web servers.
- MySQL is already configured, but depending on disk space and CPU resources it can be fine-tuned.
- The Portal cron scripts configuration must point to the correct Regional APEL server.
- The regional APEL server must accept the portal credentials.

- The regional staff should guarantee the operation of their VM instance.
- They would be required to push down hot code fixes, particularly those fixing inconsistencies or incorrect accounting.
- They should start/stop the pertinent services on downtimes/code updates.
- They would be also responsible of replication, network connectivity, certificate chain updates and renewal and so on.

- The code is maintained in a Git repository (git-scm.com), a distributed version control source repository.
- Git is used for the normal repository and there is a special branch for the regional version.
- Usually hot fixes and improvements are pushed down on the regional instances.
- Nevertheless, there is the possibility to push up and distribute local modifications by second parties.



Outline

The Accounting Portal

Regionalization

Portal Structure

Portal Deployment

Conclusions

- The Regional Portal will be distributed as a pre-configured VM
- Regional staff will have to do a initial configuration effort and the usual admin effort.
- Code updates and bug fixes will be circulated using Git.
- The Portal is based on proven and easy to use components that should insure minimal problems.