



Enabling Grids for
E-science in Europe

www.eu-egee.org

*Workshop on e-Infrastructures
May 13, 2005 - Amsterdam*

DGAS

DataGrid Accounting System

Basic concepts and current status

R. Piro

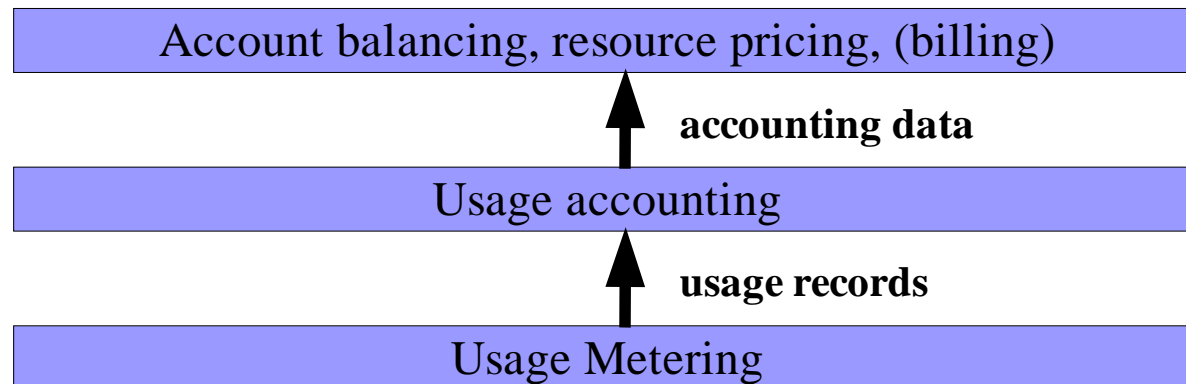
Istituto Nazionale di Fisica Nucleare (INFN) - Torino, Italy



The *Data Grid Accounting System* was originally developed within the EU Datagrid Project and is now being maintained and re-engineered within the EU EGEE Project.

The Purpose of *DGAS* is to implement *Resource Usage Metering, Accounting and Account Balancing* (through resource pricing) in a fully distributed Grid environment. It is conceived to be distributed, secure and extensible.

The system is designed in order for Usage Metering, Accounting and Account Balancing (through resource pricing) to be independent layers.



Usage Metering is done by lightweight sensors (DGAS Gianduia) installed on the Computing Elements. These sensors parse PBS/LSF/Torque event logs to build *Usage Records* that can be passed to the accounting layer (DGAS Pushd).

For a reliable accounting of resource usage (essential for billing) it is important that the collected data is *unequivocally* associated to the unique grid ID of the user (certificate subject/DN), the resource (CE ID) as well as the job (global job ID).

The usage of *Grid Resources* by *Grid Users* is registered in appropriate servers, called Home Location Registers (HLRs) that manage both user and resource accounts.

In order to achieve scalability, accounting records can be stored on an arbitrary number of independent HLRs. At least one HLR per VO is foreseen, although a finer granularity is possible.

Each HLR keeps the records of all grid jobs submitted or executed by each of its registered users or resources, thus being able to furnish usage information with many granularity levels:

Per user or resource,
per group of users or resources,
per VO.

Accounting requires accurate usage metering, but not necessarily resource pricing and billing.

Balancing and Resource Pricing

Resource pricing is done by dedicated *Price Authorities* (PAs) that may use different pricing algorithms such as: manual setting of fixed prices, dynamical determination of prices according to the state of a resource.

In order to achieve scalability, prices can be established by an arbitrary number of independent PAs. At least one PA per VO is foreseen (VOs will want to retain control on the pricing of their resources).

Price algorithms are dynamically linked by the PA server and can be re-implemented according to the resource owners' needs.

The *job cost* is determined (by the HLR service) from *resource prices* and *usage records*.

Account balancing is done by exchanging virtual credits between the *User HLR* and the *Resource HLR*.

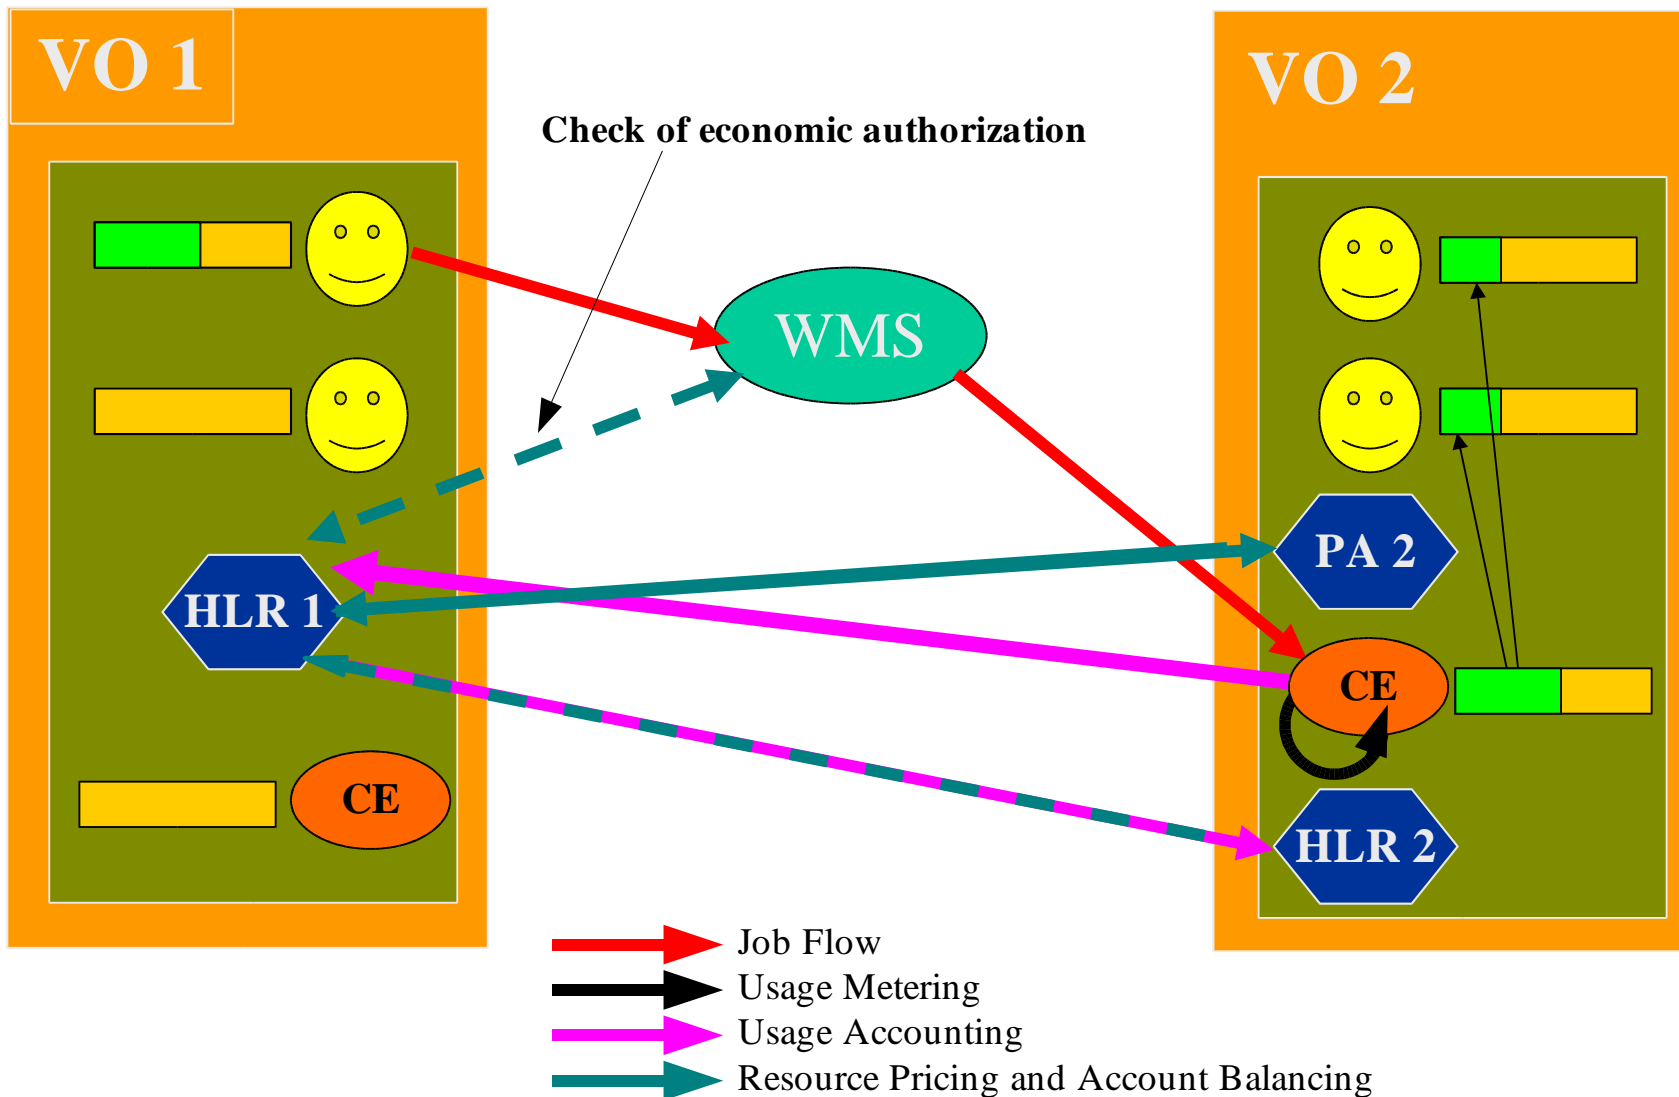
What about billing/charging?

The *Account Balancing* provided by DGAS is intentionally generic. It may be used for different use cases, such as:

- > *Monitoring* of overall resource consumption by users and resource contribution by owners.
- > *Redistribution of credits* earned by a VO's resources to the VO's users (for balanced resource sharing between VOs).
- > *Billing/charging* of users **after** resource usage.
- > *Credit/quota acquisition* by users **before** resource usage.

The purpose of DGAS is not to define (and hence limit) the economic interactions between users and resource owners, but to provide the necessary means to enable them.

Example of economic accounting



- **Privacy:**

- All communication has to be encrypted.

- Only authorized (!) access to accounting data (users, resource admins, VO admins).

- **Security/Reliability:**

- Authorization of the accounting/payment transaction (indirectly) by the user.

- Usage records transmitted from CE to HLR using the user's proxy certificate.

- Usage records stored by both User HLR and Resource HLR.

- User HLRs accept usage records only for registered users.

- Resource HLRs accept transactions only for registered resources and only from trusted User HLRs.

- Usage Record transmissions and transactions between HLR are asynchronous and in case of failures (e.g. temporary network problems) are retried.

- **Scalability:**

- Decentralized infrastructure with an arbitrary number of HLRs/PAs.

- gLite prototype testbed:
 - DGAS sensors installed on 3 Computing Elements (PBS, LSF)
 - one HLR for both user and resource accounts
 - over 7000 jobs accounted for 3 test users since March
- INFN-grid (national grid of the Italian National Institute for Nuclear Physics):
 - DGAS included in release 2.4.0 of INFN-grid (based on LCG 2.4.0)
 - one HLR for user accounts and one HLR for resource accounts
 - DGAS sensors currently being deployed on all CEs with release 2.4.0 (first jobs from some of the CEs accounted)
- For now only Usage Metering and Accounting; Pricing and Account Balancing tested, but not yet used for production.

- Standardization of service interfaces:

Transmission of usage records from DGAS sensors to User HLR with the GGF Usage Records format (the old DGAS protocol is still maintained for backward compatibility).

Ongoing development to replace the DGAS protocol for HLR and PA queries with SOAP/Web Services (possibly the GGF Resource Usage Service format?).

- *Further information and documentation about DGAS can be found at:*
<http://www.to.infn.it/grid/accounting>
- *EGEE/gLite User's Guides for DGAS components:*
<http://jra1mw.cvs.cern.ch:8180/cgi-bin/jra1mw.cgi/org.egee.jra1.deliverables/users-guide/>