Enabling Grids for E-science in Europe

www.eu-egee.org

4th EGEE Conference
October 24-28, 2005 - Pisa/Italy

Distributed Grid Accounting System

Brief overview & DGAS2APEL Interface

A. Guarise, G. Patania & **R. Piro**

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

gLite

EGEE is a project funded by the European Union under contract INFSO-RI-508833

Accounting &/vs. Monitoring

Enabling Grids for E-science in Europe

Is accounting nothing but monitoring?

- Accounting is about tracing ***single*** jobs, file requests, etc. It is about single transactions (and their logic aggregations) associated to ***accounts***.
- Grid Accounting requires a ***precise mapping*** of usage records to grid job IDs and grid user IDs.
- Monitoring can use accounting information. But not only.

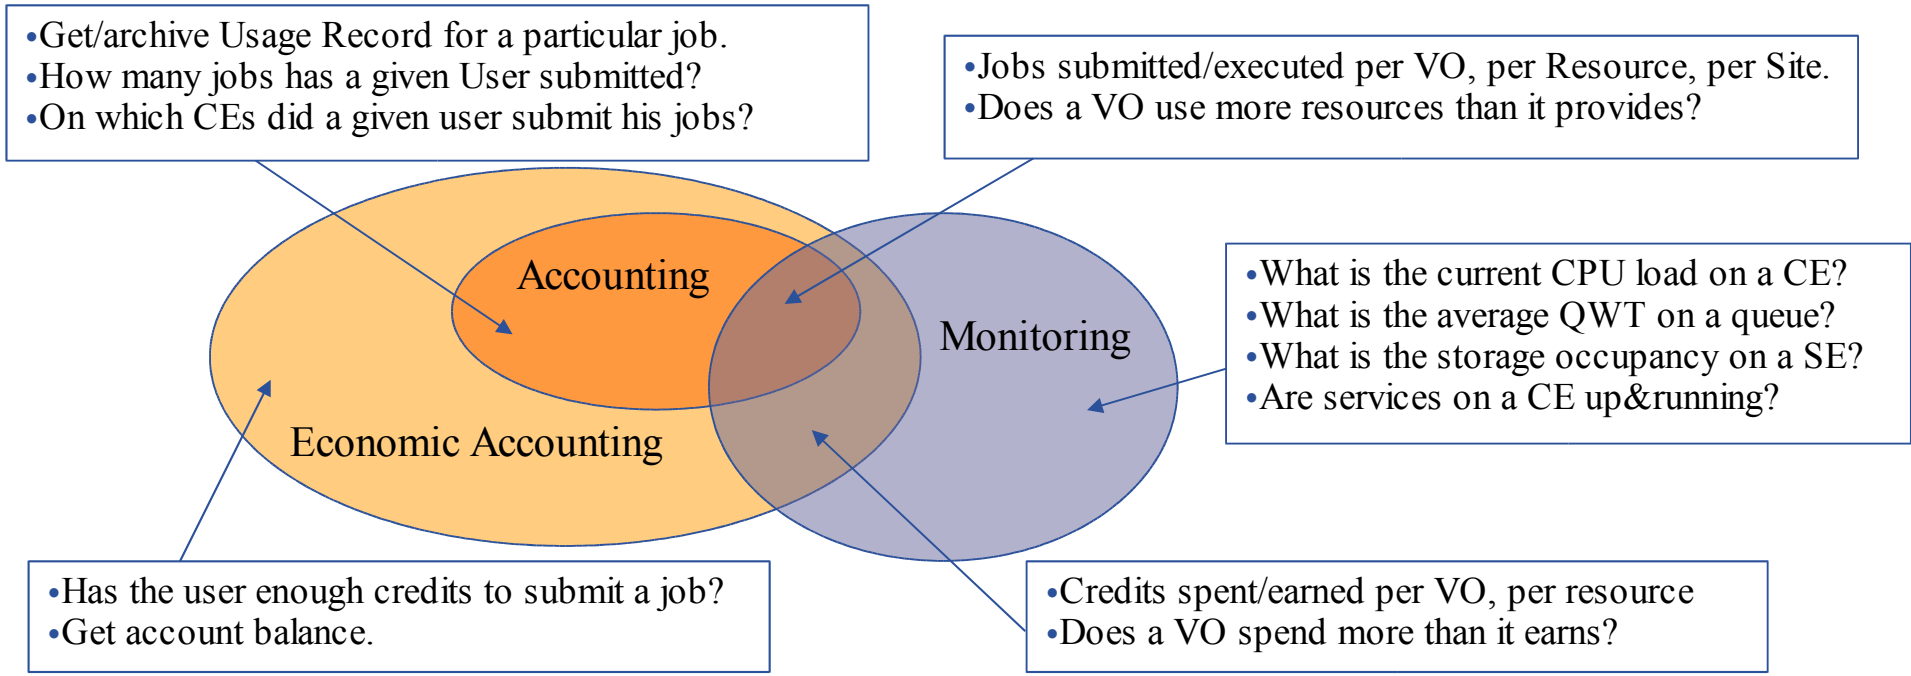

DGAS Architecture (simplified)

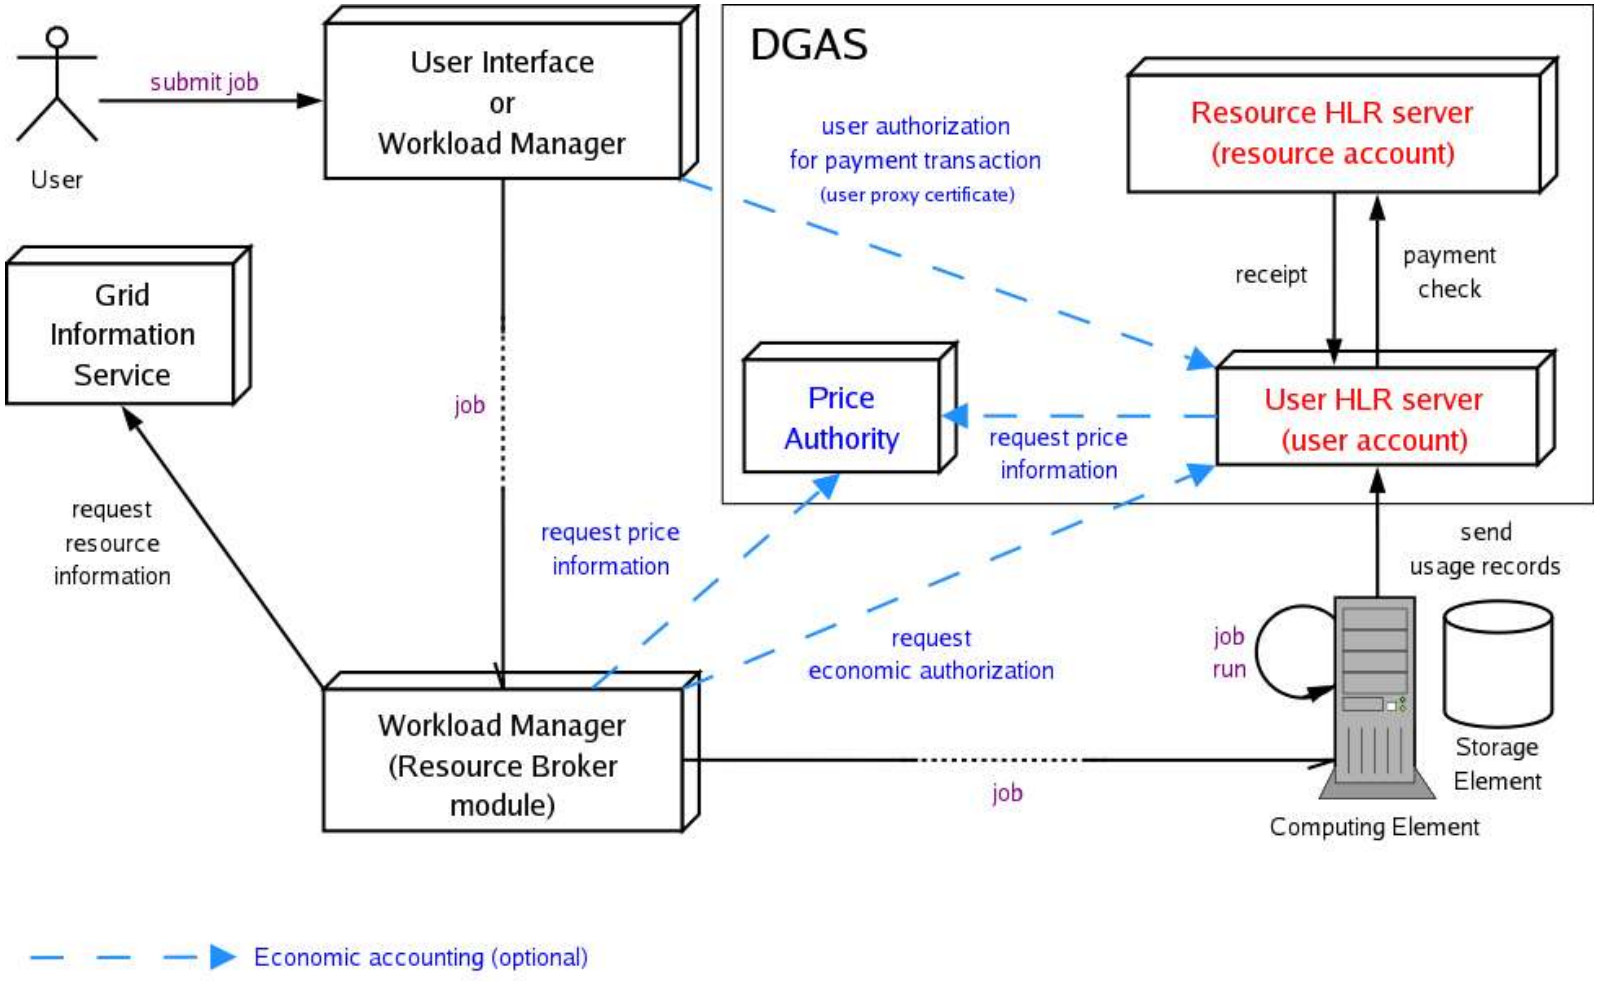

Metering Infrastructure: GIANDUIA

Enabling Grids for E-science in Europe

- GIANDUIA workflow

(**G**ianduia **I**s **A** **N**ice **D**istributed **U**sage-metering **I**nfrastructure for **A**ccounting)

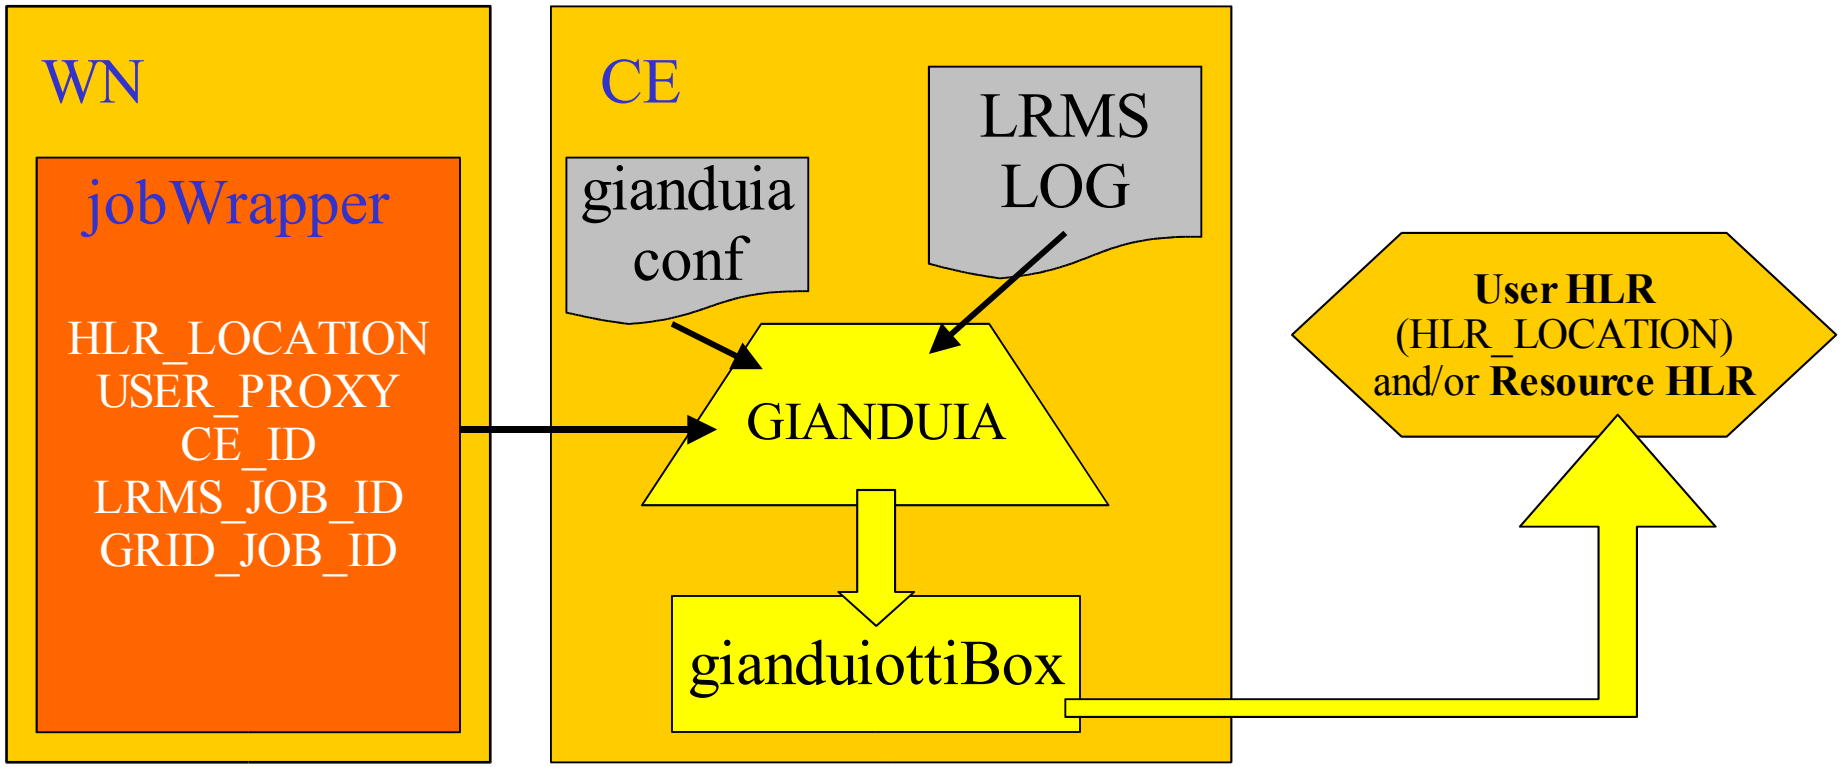

DGAS deployment (1)

Enabling Grids for E-science in Europe

- The DGAS architecture allows *flexible deployment* schemas. The deployment that we propose is to have:
  - *One **User HLR** per VO* (or more for VOs with a large number of users).
    - It stores the accounts and usage records for the registered users of the VO. The user specifies the contact string of his HLR in the job JDL.
  - *One **Resource HLR** per site*, although single HLRs can store information for many sites.
    - These store the accounts and usage records for the registered resources. The contact string is specified on the CE nodes.
  - The ***Gianduia metering system*** installed and configured on the Computing Elements.
  - One or more ***Price Authorities*** (only if Economic Accounting is desired). There should be one PA per Resource HLR.

DGAS deployment (1)

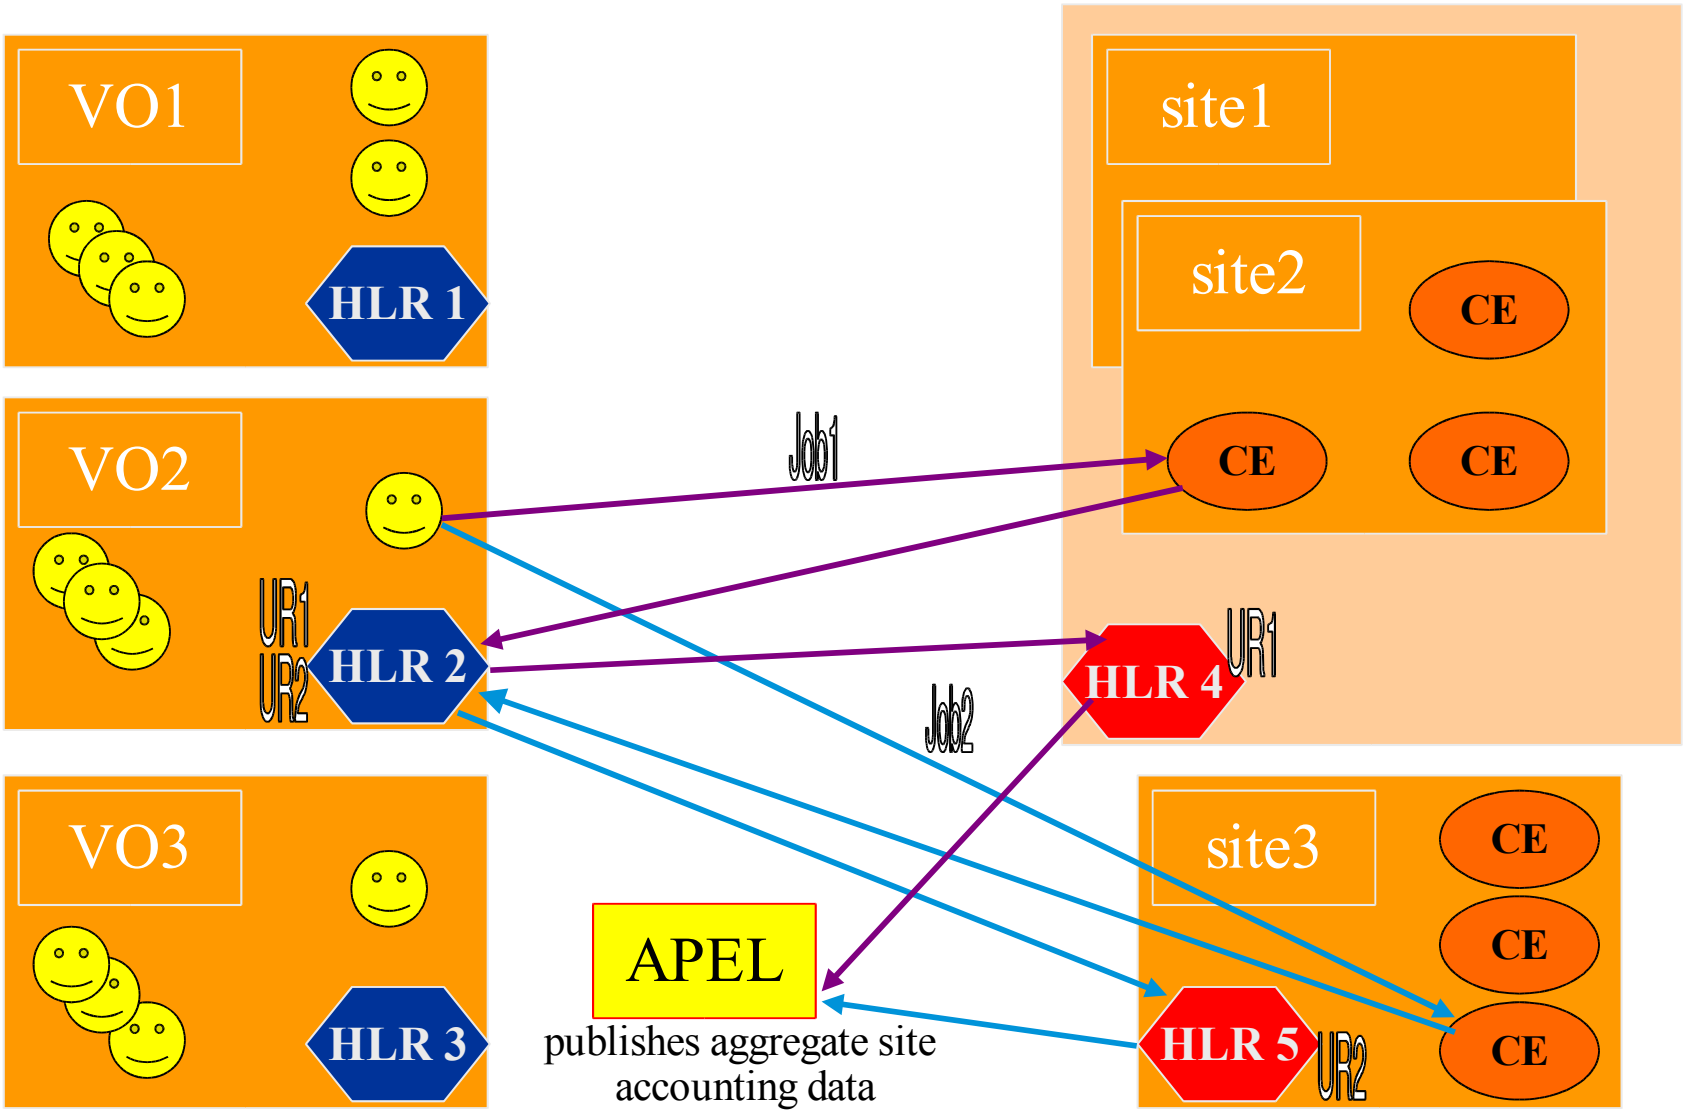

DGAS 2 APEL ...

Enabling Grids for E-science in Europe

- DGAS provides a tool that converts it's accounting records into the format used by APEL (`LcgRecords` table):
  - can be ***periodically executed*** on the Resource HLR;
  - pushes the accounting records to an (either central or local) APEL database;
  - ***keeps track of previously converted accounting records*** (only new records will be processed if re-processing is not forced);
  - ***retries*** to convert accounting records with error conditions from previous executions;
  - for reasons of ***privacy*** the user's certificate subject is not provided (missing authorization mechanism for the access to information published via R-GMA).

DGAS accounting records

- with the DGAS sensors currently installed on INFNgrid (LCG) resources:
  - user's certificate subject
  - user's VO
  - local user ID
  - grid ID of the job
  - CE's grid ID
  - CPU time & wall clock time
  - physical & virtual memory usage
  - accounting timestamp
- with the DGAS sensors of gLite 1.4:
  - user's certificate subject
  - user's FQAN (VOMS certificates)
  - user's VO
  - local user & group ID
  - job ID (both grid and local LRMS ID)
  - CE's grid ID
  - SpecInt2000 & SpecFloat2000
  - number of processors
  - CPU time & wall clock time
  - physical & virtual memory usage
  - accounting timestamp
  - execution start/end timestamps
  - ctime, qtime, etime
  - job's exit status

Important issues adressed by DGAS

Enabling Grids for E-science in Europe

- **Privacy**:
  - Only *authorized (!) access* to accounting data (users, resource admins, VO admins).
- **Security/Reliability**:
  - Usage records transmitted from CE to the User HLR using the *user's proxy certificate*.
  - Usage records stored by both User HLR and Resource HLR (only the Resource HLR for local accounting).
  - User HLRs accept usage records only for registered users.
  - Resource HLRs accept transactions only for registered resources and only from trusted User HLRs (or directly from the resource for local accounting).
  - Usage Record transmissions and transactions between HLRs are *asynchronous* and in case of failures (e.g. temporary network problems) are retried.
- **Scalability**:
  - *Decentralized* infrastructure with an arbitrary number of HLRs/PAs.
- **Accounting for Resources and Users**
  - both site managers and users should be able to access information on their jobs (resource level reporting and user level reporting).

References

- *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:*
  https://edms.cern.ch/file/571271/1/