DISTRIBUTED GRID ACCOUNTING SYSTEM (DGAS)

Basic concepts and proposals for the UR format

Rosario M. Piro – INFN Torino, Italy

INTRODUCTION

- The Data Grid Accounting System (DGAS) was originally developed within the EU Datagrid (EDG) project and is now being maintained and re-engineered within the EU EGEE project and contributes to the OMII-Europe project.
- The purpose of *DGAS* is to implement Resource Usage Metering, Accounting and Account Balancing (through resource pricing) in a fully distributed Grid environment.

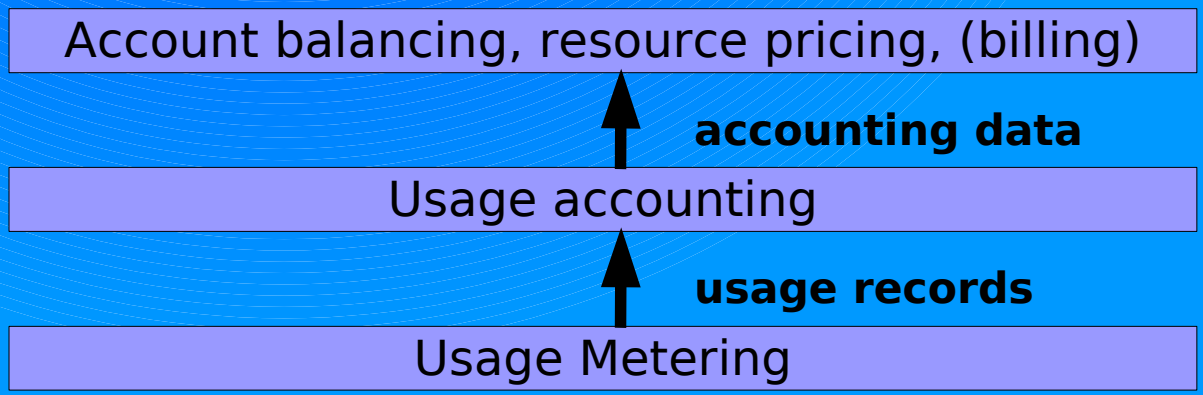

METERING

- Lightweight sensors on Computing Elements:
  - parse PBS/Torque or LSF event logs
    - a common accounting sensor for PBS/Torque, LSF, Condor, SGE is being developed in collaboration with APEL and OSG.
  - pass accounting records to the accounting layer.
- For reliable accounting in a grid environment:
  - collected data must be unequivocally associated to:
    - the grid ID of the user (certificate subject/DN)
    - the grid ID of the job (GlobalJobId)
    - the grid ID of the resource (e.g. Globus contact string).
  - problem: most LRMS logs don't provide grid-related job info. Additional log by LCG, gLite and Condor resources for mapping grid-related info to local info.

ACCOUNTING

- Accounting records are stored on a distributed infrastructure of Home Location Register (HLR) servers (for scalability).
- Accounting records are associated to user accounts and resource accounts identified by:
  - User DN (certificate subject) as user account ID
  - Global resource ID (CE ID) as resource account ID
- Accounting records can be forwarded from Resource/Site HLRs to User HLRs:
  - users/VO admins don't need to query many Site HLRs for accounting information, but only the HLR that managers the respective account (user-level accounting).

ACCOUNT BALANCING

- DGAS Price Autority (PA) servers:
  - responsible for setting resource prices (only CPU for now),
  - prices can be set manually or determined dynamically,
  - pricing algorithms are dynamically linked libraries and can be customized as needed.
- Job cost determined by Resource/Site HLR from resource price (per unit) and resource usage (number of units).
- Account balancing by exchanging virtual credits between user and resource account.
- Resource pricing and account balancing are OPTIONAL in DGAS.

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 VOs 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 information to enable them.

WORK FLOW EXAMPLE

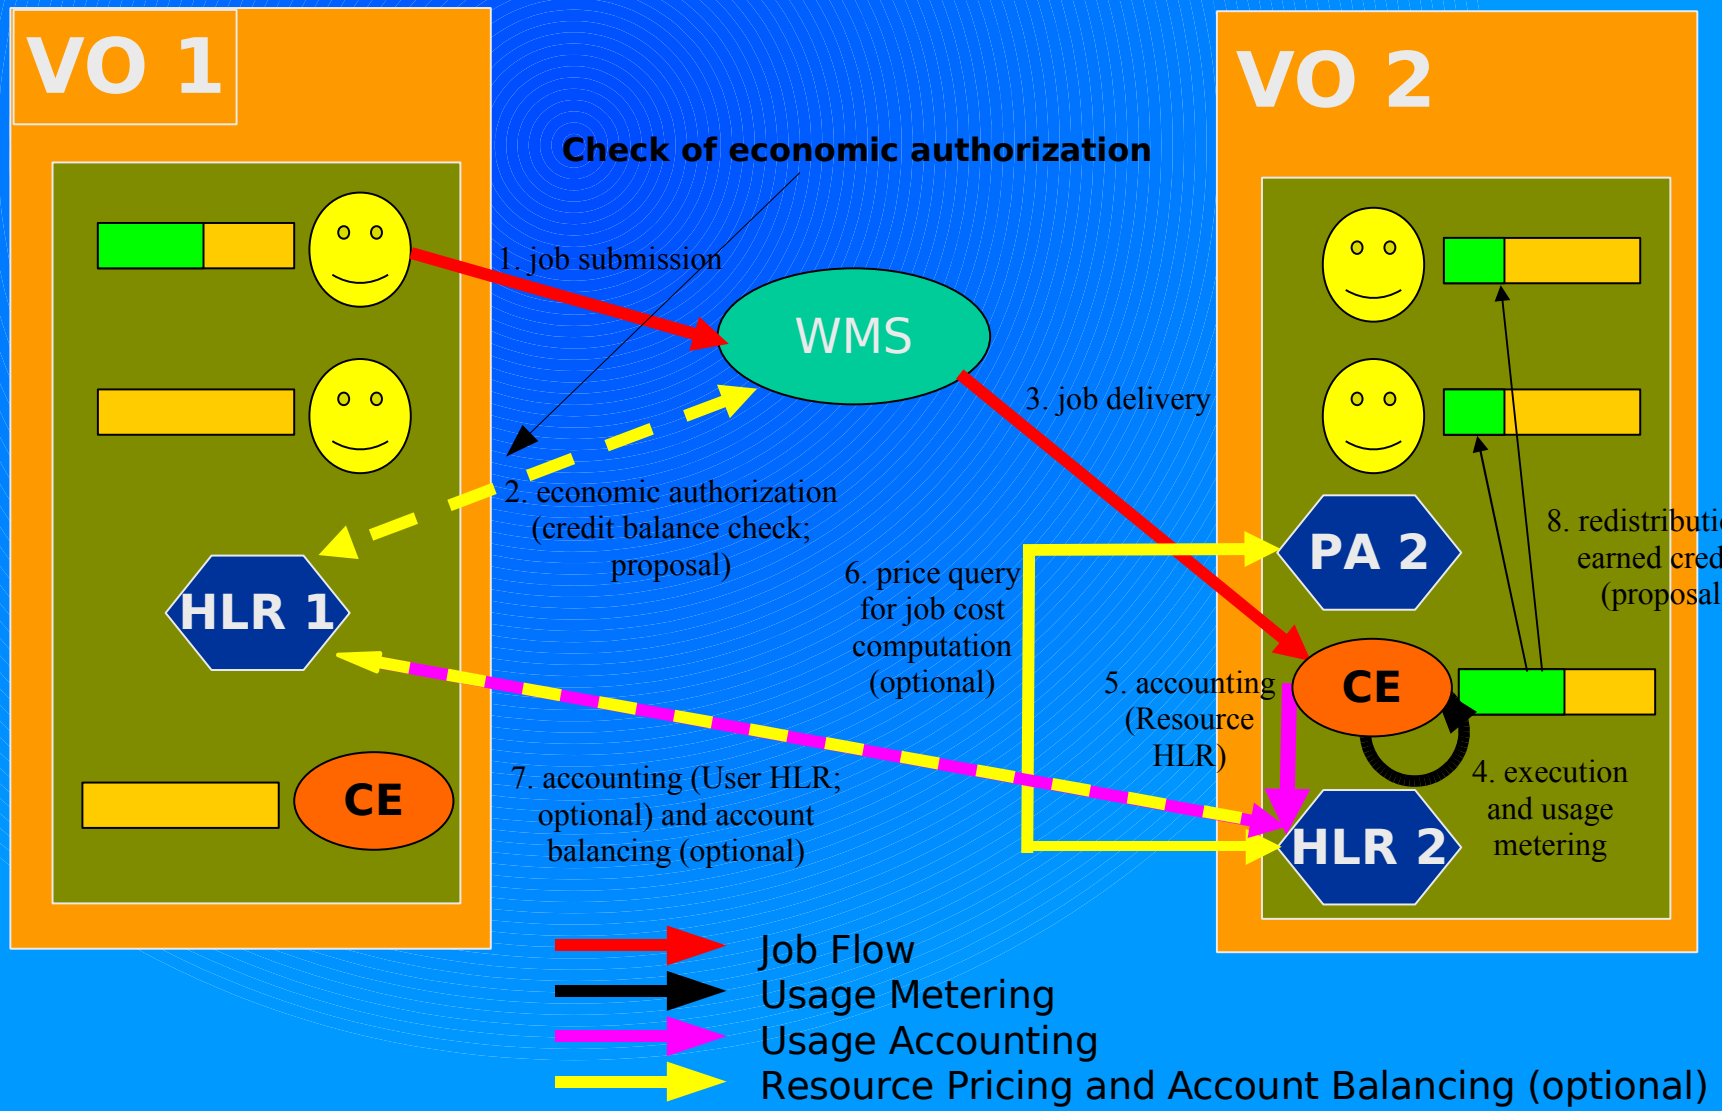

IMPORTANT ISSUES
addressed by DGAS

- Privacy:
  - all communication encrypted (based in Globus GSI)
  - only authorized (!) access to accounting data (users, admins).
- Security/Reliability:
  - Accounting records can be stored by both Resource HLR and User HLR.
  - User HLRs accept accounting records only for registered users and only from trusted Resource HLRs.
  - Resource HLRs accept records only from registered resources.
  - Accounting 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.

UR PROPOSALS (1)

- Many large Grid communities (LCG, OSG, ...) require information on the User VO! ProjectName is semantically different ...
- Several LCG VOs need more detailed information on groups and user roles within the VO (FQAN = Fully Qualified Attribute Name)!
- Proposal: VO and FQAN as parts of a user's identity

```
<UserIdentity>
   <LocalUserId>...</LocalUserId>
   ...
   <VOName>cms</VOName>
   <UserFQAN>/cms/production/Role=Admin,...</UserFQAN>
</UserIdentity>
```

UR PROPOSALS (2)

- Usage records need to be associated to Grid user IDs (DNs), Grid job IDs and Grid resource IDs (e.g. Globus contact string).
  - Grid resources may be logical entities (e.g. LCG CE IDs: *"t2-ce-01.to.infn.it:2119/jobmanager-lcglsf-cert"*). MachineName, Queue and Host are not enough ...
- SiteName and LRMS type would be useful.
- Proposal (in analogy to UserIdentity and JobIdentity):

```
<ResourceIdentity>
   <GlobalResourceId>t2-ce-01.to.infn.it:...</GlobalResourceId>
   <SiteName>INFN-TORINO</SiteName>
   <Queue lrmsType="LSF">cert</Queue>
   <MachineName>Tier2-Cluster-Torino</MachineName>
   <Host>wn17.to.infn.it</Host>
   ...
</ResourceIdentity>
```

UR PROPOSALS (3)

- The measured resource usage of a job depends heavily on the performance of the executing resource.
- How to compare usage records from heterogeneous resources across the Grid? Normalization of resource usage is one of the important issues in multi-organizational environments that are interested in fair share.
- 1 sec CPU time ≠ 1 sec CPU time !!! (for two different processor performances)
- Proposal: add performance information to ResourceIdentity

```
<ResourceIdentity>
   ...
   <Performance units="SpecInt2000">...</Performance>
   <Performance units="SpecFloat2000">...</Performance>
</ResourceIdentity>
```

UR PROPOSALS (4)

- We should go beyond job usage and think about other resource usage!
- Storage accounting, for example, is one of the big upcoming issues ...
- Proposal:
  - JobUsageRecord = UsageRecord + JobIdentity
  - StorageUsageRecord = UsageRecord + FileIdentity

```
<StorageUsageRecord>
   <UserIdentity>...</UserIdentity>
   ...
   <FileIdentity>...</FileIdentity>
   ...
</StorageUsageRecord>
```

UR PROPOSALS (5)

- Users might be charged not only for consumed resources, but also for requested/reserved resources.
- Should it be possible to specify resource requests in the same Usage Record?
- Possible approach: additional attribute?

```
<Disk usage="consumed">...</Disk>
<Disk usage="requested">...</Disk>
<Network usage="consumed">...</Network>
<Network usage="requested">...</Network>

(default should be "consumed" for backward compatibility)
```

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:* *http://jra1mw.cvs.cern.ch:8180/cgi-bin/jra1mw.cgi/org.egee.jra1.deliverables/users-guide/*