



Enabling Grids for E-science

More details on the gLite IS

Giuseppe Andronico

INFN

EGEE Tutorial

Seoul, 29-30.08.2005

www.eu-egee.org



Information Society



- **Information System**
 - lcg-infosites
 - R-GMA
- **Accounting System**
- **Monitoring System**

lcg-infosites (the present)

If you are a user

Retrieve information of Grid resources and status

Get the information of your jobs status

If you are a middleware developer

Workload Management System:
Matching job requirements and Grid resources

Monitoring Services:
Retrieving information of Grid Resources status and availability

If you are site manager or service

You “generate” the information for example relative to your site or to a given service

```

*****
These are the data for alice: (in terms of CPUs)
*****
#CPU  Free   Total Jobs   Running   Waiting   Computing Element
-----
52     51     0             0          0      ce.prn.hp.com:2119/jobmanager-lcgpbs-long
16     14     3             2          1      lcg06.sinp.msu.ru:2119/jobmanager-lcgpbs-long
[.....]
The total values are:
-----
10347  5565    2717         924        1793
    
```



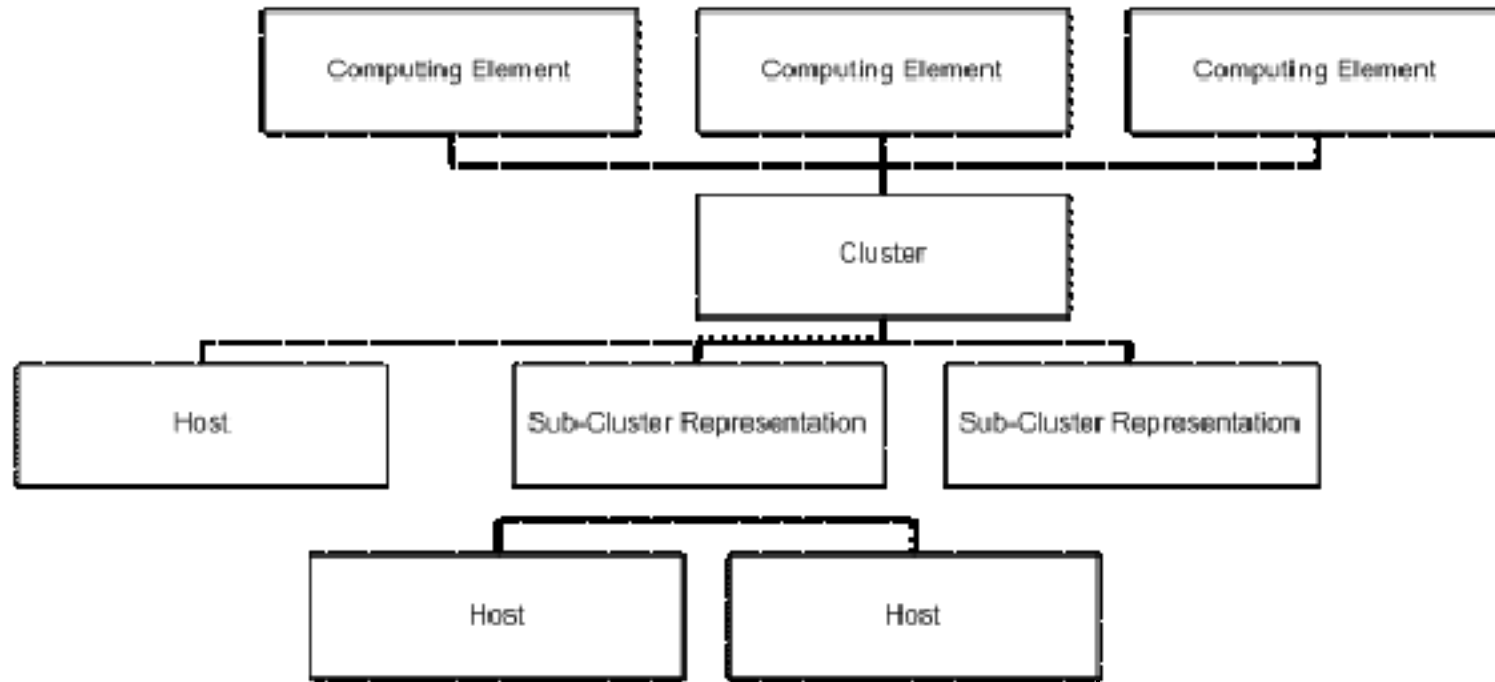
☒ Something has managed this information: (General IS architecture)

☒ Something has provided it: (Providers, Servers)

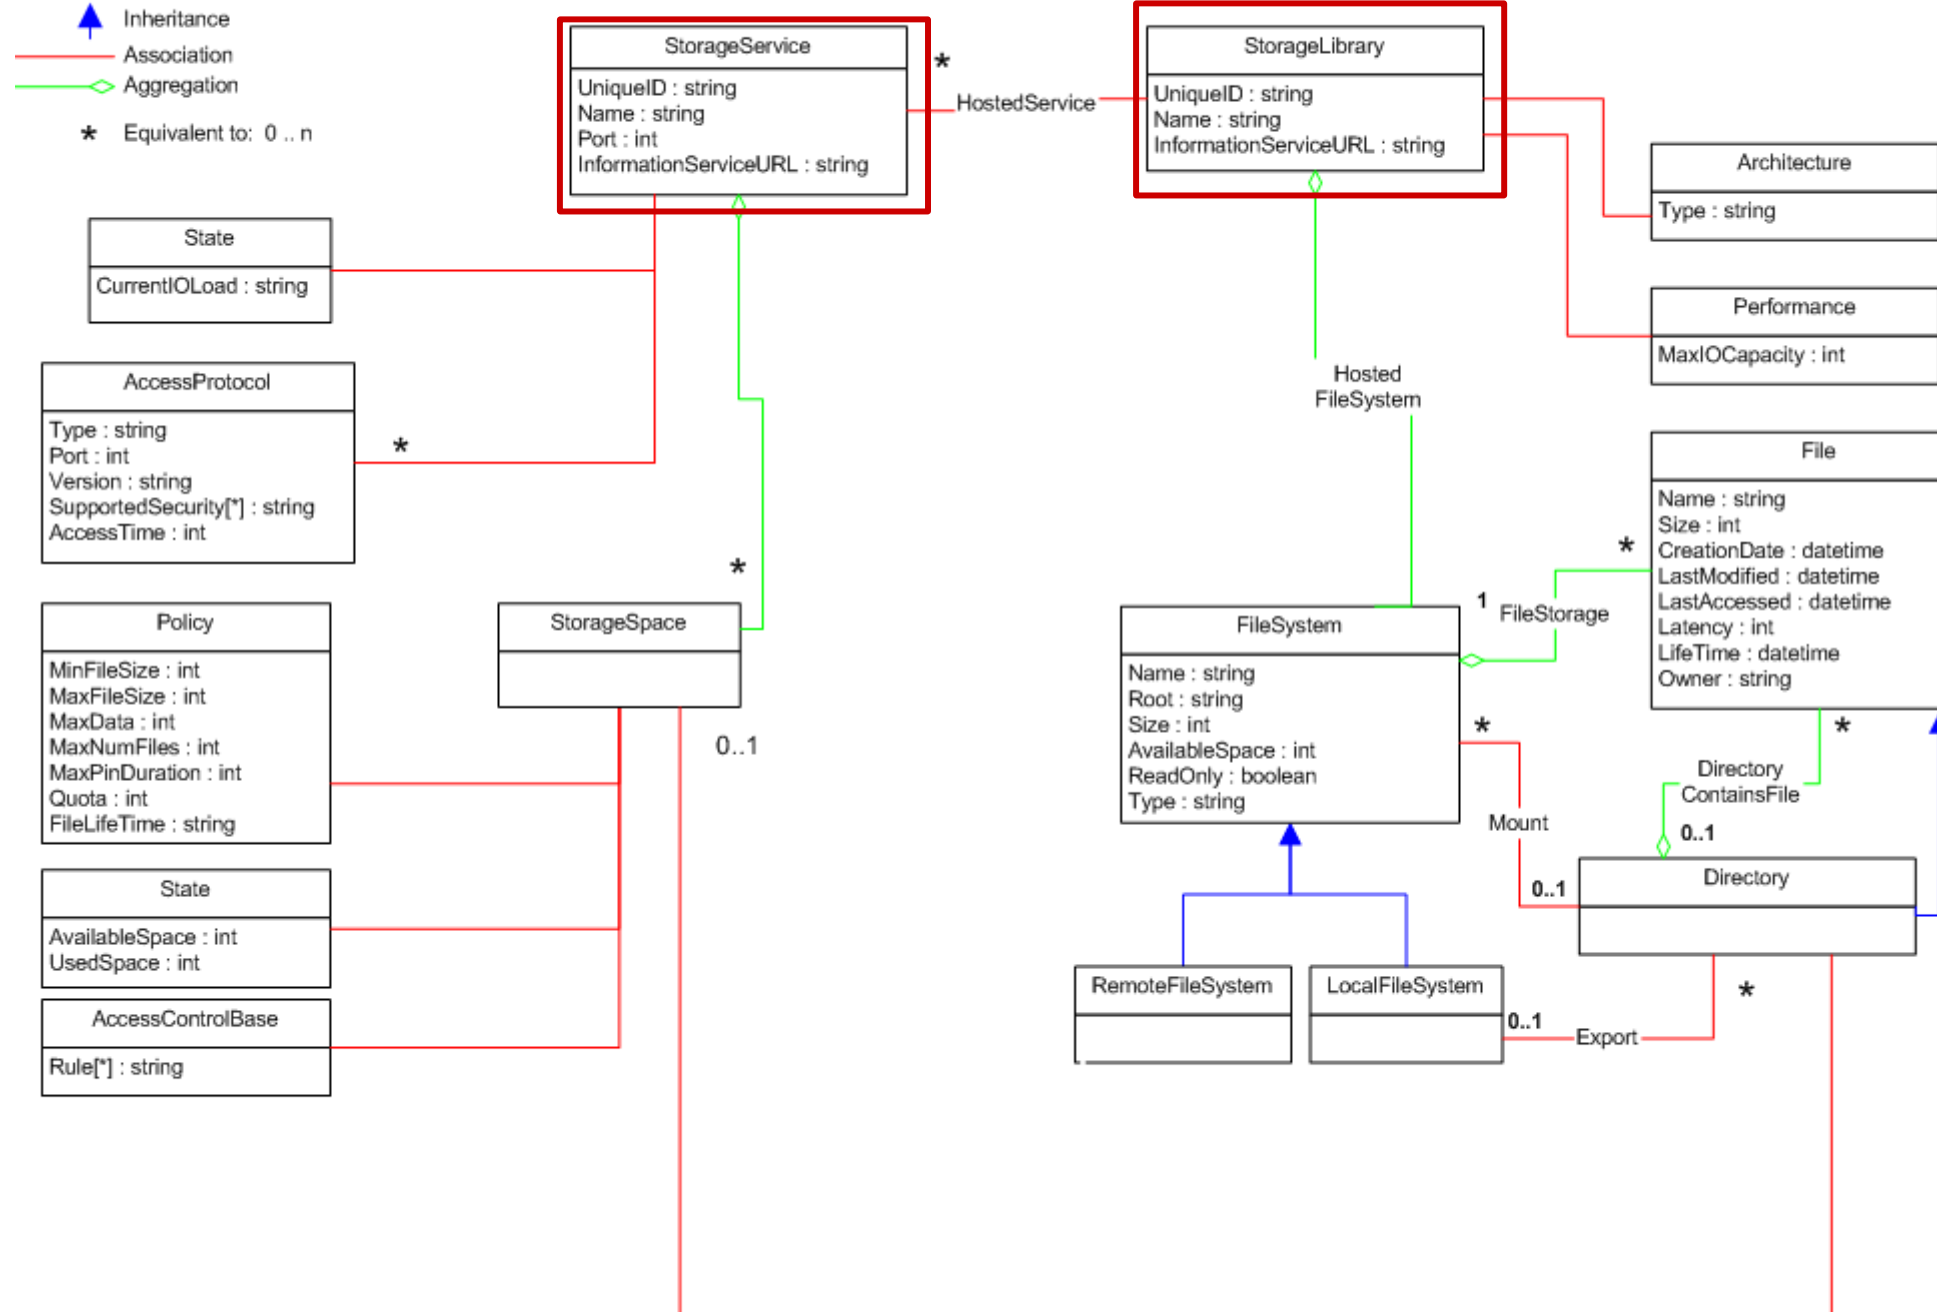
☒ It is following a certain "schema": (GLUE Schema)
 She will use some EGEE/LCG tools and after few moments...

☒ And she has accessed it following a protocol: (Access Protocol: LDAP)

- **Developed within High Energy Physics (HEP) community**
 - DataGrid / EGEE
 - DataTAG
 - Globus
- **Currently defines CEs and SEs**
- **Entire R-GMA Schema (not only GLUE):**
 - For service discovery and monitoring
 - <http://hepunix.rl.ac.uk/egee/jra1-uk/glite-r1/schema/index.html>



Storage Element



MDS: Monitoring and Discovery Service

- ▶ Adopted from Globus
- ▶ It is the general architecture of EGEE/LCG to manage Grid information

General steps:

- 1st. At each site **providers** report static and dynamic service status to **servers**
- 2nd. A **central system** queries these servers and stores the retrieved information in a database
- 3rd. This information will be accessed through a given **access protocol**
- 4th. The central system provides the information in a **given schema**

BDII (a MDS evolution) is the current EGEE/LCG Information System and it is based on LDAP

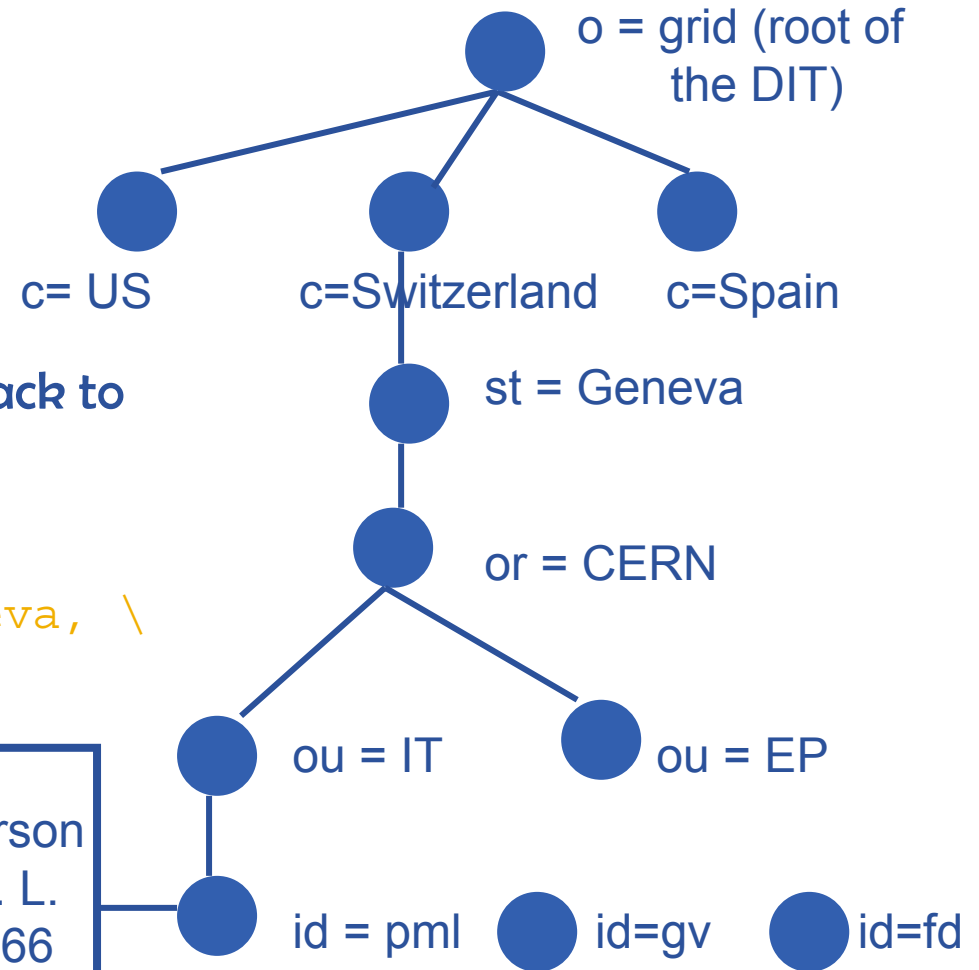
▶ LDAP structures data as a tree

▶ The values of each entry are uniquely named

▶ Following a path from the node back to the root of the DIT, a unique name is built (the DN):

`"id=pml,ou=IT,or=CERN,st=Geneva, \ c=Switzerland,o=grid"`

objectClass:person
 cn: Patricia M. L.
 phone: 5555666
 office: 28-r019



♠ lcg-infosites

- Already deployed in LCG-2 in the last release
- It is intended to be the most complete information retriever for the user:
 - ✓ Once he arrives at the Grid (on UIs)
 - ✓ To be used by the user applications (on WNs)
- Several versions of this script have been included in the software packages of ATLAS and the monitoring services of Alice (MonAlisa)
- You do not need a proxy



This will be tested during
the hands-on session

> `lcg-infosites --vo <your_vo> feature --is <your_bdii>`

- It's mandatory to include the **vo** and the **feature**
- The **-is** option means the BDII you want to query. If not supplied, the BDII defined into the **LCG_GFAL_INFOSYS** will be interrogated

Features and descriptions:

closeSE	Names of the CEs where the user's VO is allowed to run together with their corresponding closest SEs
ce	Number of CPUs, running and waiting jobs and names of the CEs
se	SEs names together with the available and used space
lrc (rnc)	Name of the lrc (rnc) for the user's VO
all	It groups all the features just described
help	Description of the script

> `lcg-infosites --vo alice se --is lxb2006.cern.ch`

```


*****
These are the data for alice: (in terms of SE)
*****
Avail Space (Kb)          Used Space (Kb)          SEs
-----
33948480                  2024792                  se.prd.hp.com
506234244                 62466684                 teras.sara.nl
1576747008                3439903232               gridkap02.fzk.de
1000000000000            5000000000000            castorgrid.cern.ch
304813432                 133280412                gw38.hep.ph.ic.ac.uk
651617160                 205343480                mu2.matrix.sara.nl
1000000000000            1000000000               lcgads01.gridpp.rl.ac.uk
415789676                 242584960                cclcgseli01.in2p3.fr
264925500                 271929024                se-a.ccc.ucl.ac.uk
668247380                 5573396                  seitep.itep.ru
766258312                 681359036                t2-se-02.lnl.infn.it
660325800                 1162928716               tbn17.nikhef.nl
1000000000000            1000000000000            castorftp.cnaf.infn.it
14031532                  58352476                 lcgse01.gridpp.rl.ac.uk
1113085032                1034242456               zeus03.cyf-kr.edu.pl
[... ..]

```

R-GMA (the future)

- **Relational Grid Monitoring Architecture (R-GMA)**
 - Developed as part of the EuropeanDataGrid Project (EDG)
 - Now as part of the EGEE project.
 - Based on the Grid Monitoring Architecture (GMA) from the Global Grid Forum (GGF).
- **Uses a relational data model.**
 - Data is viewed as a table.
 - Data structure defined by the columns.
 - Each entry is a row (tuple).
 - Queried using Structured Query Language (SQL).

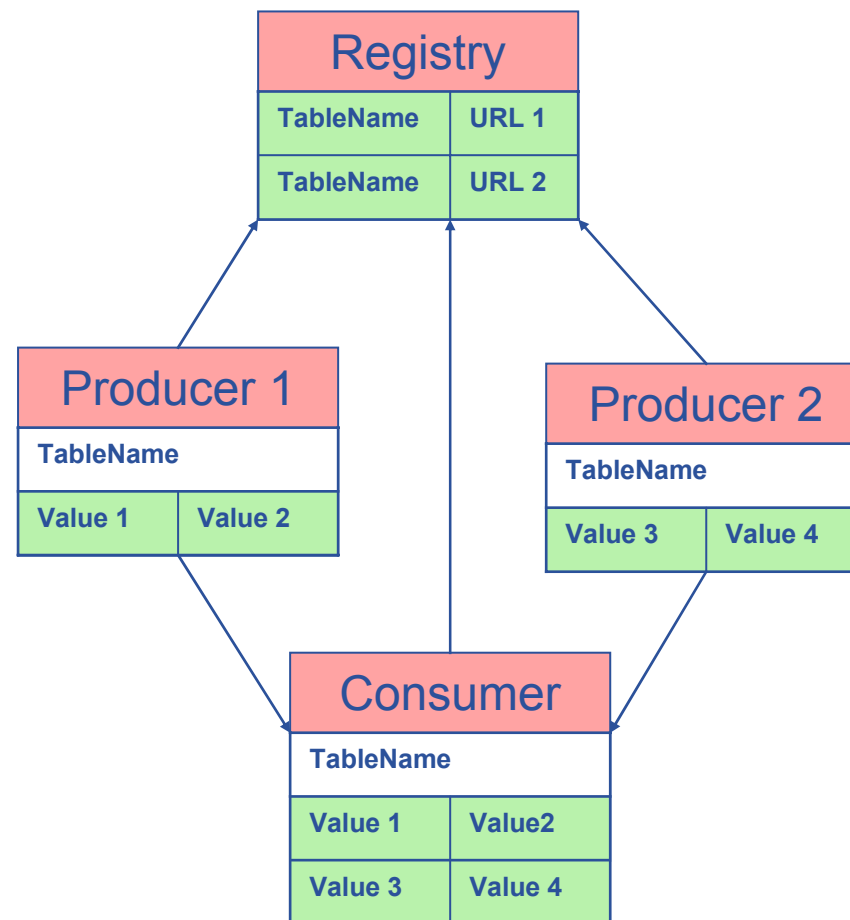
- The Producer stores its location (URL) in the Registry.
- The Consumer looks up producer URLs in the Registry.
- The Consumer contacts the Producer to get all the data.
- Or the Consumer can listen to the Producer for new data.



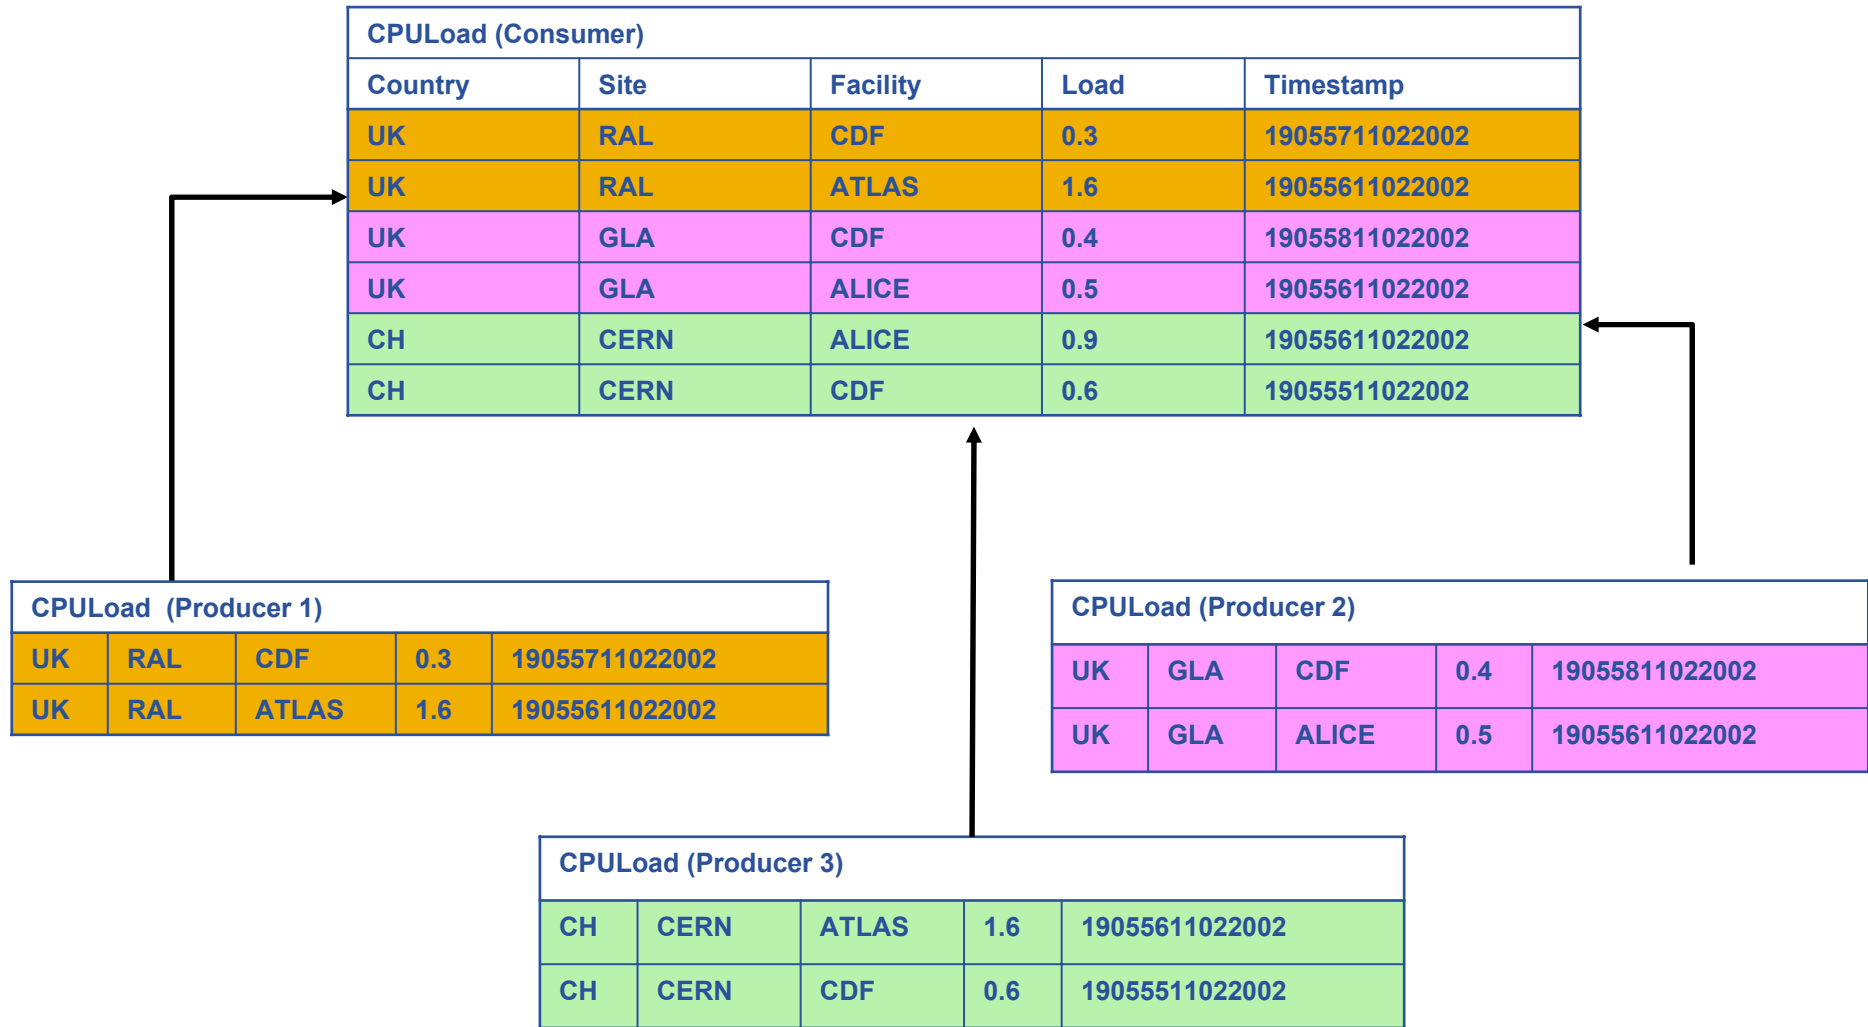
name	ID	birth	Group
Tom	4	1977-08-20	HR

`SELECT * FROM people WHERE group='HR'`

- The Consumer will get all the URLs that could satisfy the query.
- The Consumer will connect to all the Producers.
- Producers that can satisfy the query will send the tuples to the Consumer.
- The Consumer will merge these tuples to form one result set.




Select * from CPUload



Service				
URI	VO	type	emailContact	site
gppse01	alice	SE	sysad@rl.ac.uk	RAL
gppse01	atlas	SE	sysad@rl.ac.uk	RAL
gppse02	cms	SE	sysad@rl.ac.uk	RAL
lxshare0404	alice	SE	sysad@cern.ch	CERN
lxshare0404	atlas	SE	sysad@cern.ch	CERN

ServiceStatus				
URI	VO	type	up	status
gppse01	alice	SE	y	SE is running
gppse01	atlas	SE	y	SE is running
gppse02	cms	SE	n	SE ERROR 101
lxshare0404	alice	SE	y	SE is running
lxshare0404	atlas	SE	y	SE is running



Result Set (Consumer)	
URI	emailContact
gppse02	sysad@rl.ac.uk

SELECT Service.URI Service.emailContact FROM Service S, ServiceStatus SS
WHERE (S.URI= SS.URI and SS.up='n')

- **The easiest way to try out R-GMA.**
 - It is installed on the machine running the Registry and Schema:
<https://rgmasrv.ct.infn.it:8443/R-GMA>
 - You can also install it along with the Producer and Consumer Servlets.
- **Using the Browser you can do the following.**
 - Browse the tables in the schema.
 - Look at the table definitions.
 - See all the available producers for a table.
 - Query a table.
 - Query only selected producers.

R-GMA Browser Home Page - Mozilla

File Edit View Go Bookmarks Tools Window Help

https://rgmasrv.ct.infn.it:8443/R-GMA/ Go Search


Home Bookmarks Webmail Missioni Offerte Ordini FastWeb Mozilla.org

R-GMA Browser

Home

Predefined:

- [Services](#)
- [Site](#)
- [Table Sets](#)



Enabling Grids For E-science

[All tables](#)

- [GLUE Info Providers](#)
- [Network Monitoring](#)
- [Service Discovery](#)
- [CMS](#)
- [GlueSA](#)
- [GlueSAAccessControlBaseRule](#)
- [GlueSE](#)
- [GlueSEAccessProtocol](#)
- [GlueSEAccessProtocolSupportedSec](#)
- [GlueSL](#)
- [GlueService](#)
- [GlueServiceAccessControlRule](#)
- [GlueSubCluster](#)
- [GlueSubClusterSoftwareRunTimeEnv](#)
- [GlueVO](#)
- [JobMonitor](#)
- [NetworkFileTransferThroughput](#)
- [NetworkICMPPacketLoss](#)
- [NetworkOneWayIPDV](#)
- [NetworkRTT](#)
- [NetworkTCPThroughput](#)
- [NetworkUDPPacketLoss](#)
- [NetworkUDPThroughput](#)
- [Service](#)
- [ServiceAssociation](#)
- [ServiceData](#)
- [ServiceStatus](#)
- [Site](#)
- [UserTable](#)

Query: `SELECT Name, Endpoint, Type, MajorVersion, MinorVersion, PatchVersion, Site_Name, WSDL, Semantics, MeasurementDate, MeasurementTime FROM Service`

Name	Endpoint
https://rgmasrv.ct.infn.it:8443/R-GMA/ArchiverServlet	https://rgmasrv.ct.infn.it:8443/R-GM
https://rgmasrv.ct.infn.it:8443/R-GMA/ConsumerServlet	https://rgmasrv.ct.infn.it:8443/R-GM
https://rgmasrv.ct.infn.it:8443/R-GMA/DBProducerServlet	https://rgmasrv.ct.infn.it:8443/R-GM
https://rgmasrv.ct.infn.it:8443/R-GMA/BrowserServlet	https://rgmasrv.ct.infn.it:8443/R-GM
https://rgmasrv.ct.infn.it:8443/R-GMA/SchemaServlet	https://rgmasrv.ct.infn.it:8443/R-GM
https://rgmasrv.ct.infn.it:8443/R-GMA/LatestProducerServlet	https://rgmasrv.ct.infn.it:8443/R-GM
https://rgmasrv.ct.infn.it:8443/R-GMA/CanonicalProducerServlet	https://rgmasrv.ct.infn.it:8443/R-GM
https://rgmasrv.ct.infn.it:8443/R-GMA/StreamProducerServlet	https://rgmasrv.ct.infn.it:8443/R-GM
https://rgmasrv.ct.infn.it:8443/R-GMA/RegistryServlet	https://rgmasrv.ct.infn.it:8443/R-GM
glite-rb.ct.infn.it_Logging_Bookkeeping_Server	http://glite-rb.ct.infn.it/LB/LBServer

Number of rows: 10

- **APIs exist in Java, C, C++, Python.**
 - For clients (servlets contacted behind the scenes)
- **They include methods for...**
 - Creating consumers
 - Creating primary and secondary producers
 - Setting type of queries, type of produces, retention periods, time outs...
 - Retrieving tuples, inserting data
 - ...
- **You can create your own Producer or Consumer.**

- **R-GMA overview page.**
 - <http://www.r-gma.org/>
- **R-GMA in EGEE**
 - <http://hepunx.rl.ac.uk/egee/jra1-uk/>
- **R-GMA Documentation**
 - <http://hepunx.rl.ac.uk/egee/jra1-uk/LCG/doc/>

A generic Grid accounting process involves many subsequent phases that can be divided in:

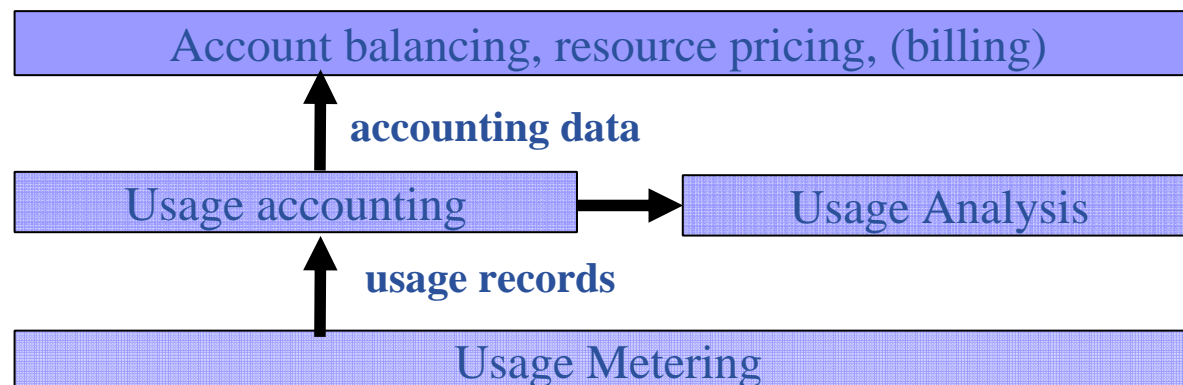
- **Metering:** collection of usage metrics on computational resources.
- **Accounting:** storage of such metrics for further analysis.
- **Usage Analysis:** Production of reports from the available records.
- **Pricing:** Assign and manage prices for computational resources.
- **Billing:** Assign a cost to user operations on the Grid and charge them.

In this presentation we briefly describe these steps and give a quick overview of DGAS, the accounting middleware of the EGEE project.

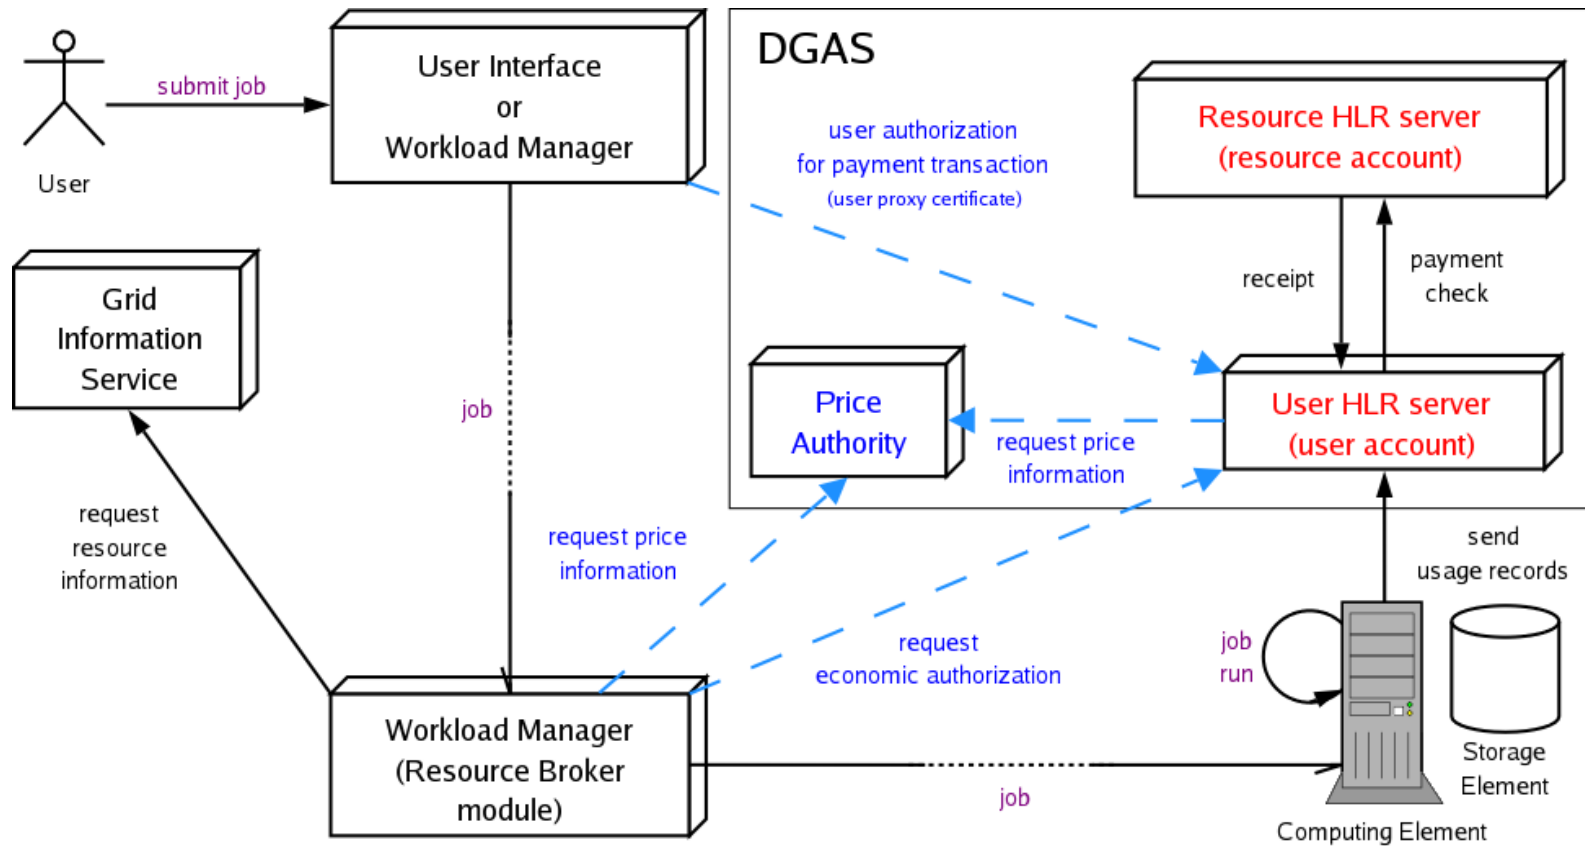
The *Data Grid Accounting System* was originally developed within the EU Datagrid Project and is now being maintained and re-engineered within the EU EGEE Project.

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.



A simplified view of DGAS within the WMS context.



--- Economic accounting (optional)

Usage Metering on Computing Elements is done by lightweight sensors 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.

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).

A process, completely transparent to the Grid User collects the necessary information needed by the Accounting. These, and the corresponding metrics are sent via an *encrypted* channel to the Accounting System *signed with the user credentials*.

The usage of *Grid Resources* by *Grid Users* is registered in appropriate servers, called Home Location Registers (HLRs) where both users and resources are registered.

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 usage metering, but not necessarily resource pricing and billing.

Resource pricing is done by dedicated Price Authorities (PAs) that may use different pricing algorithms: 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.

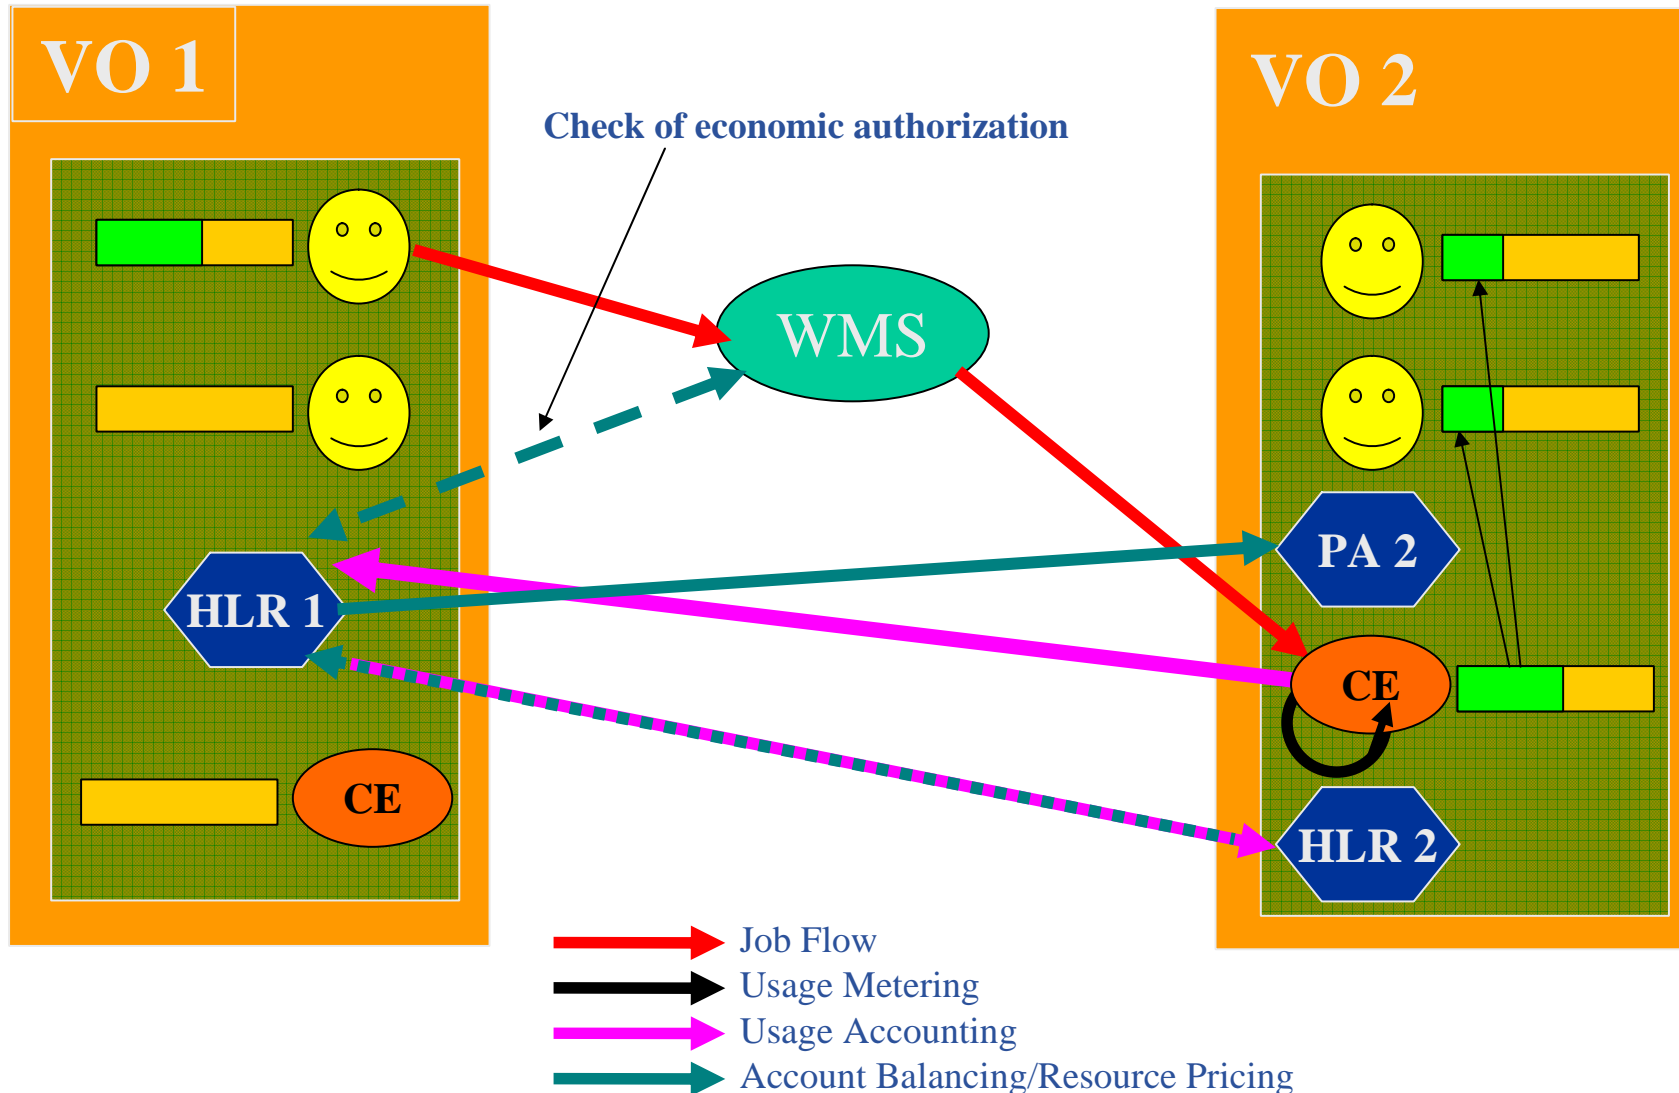
The job cost is determined (by the HLR service) from *resource prices* and *usage records*.

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

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.

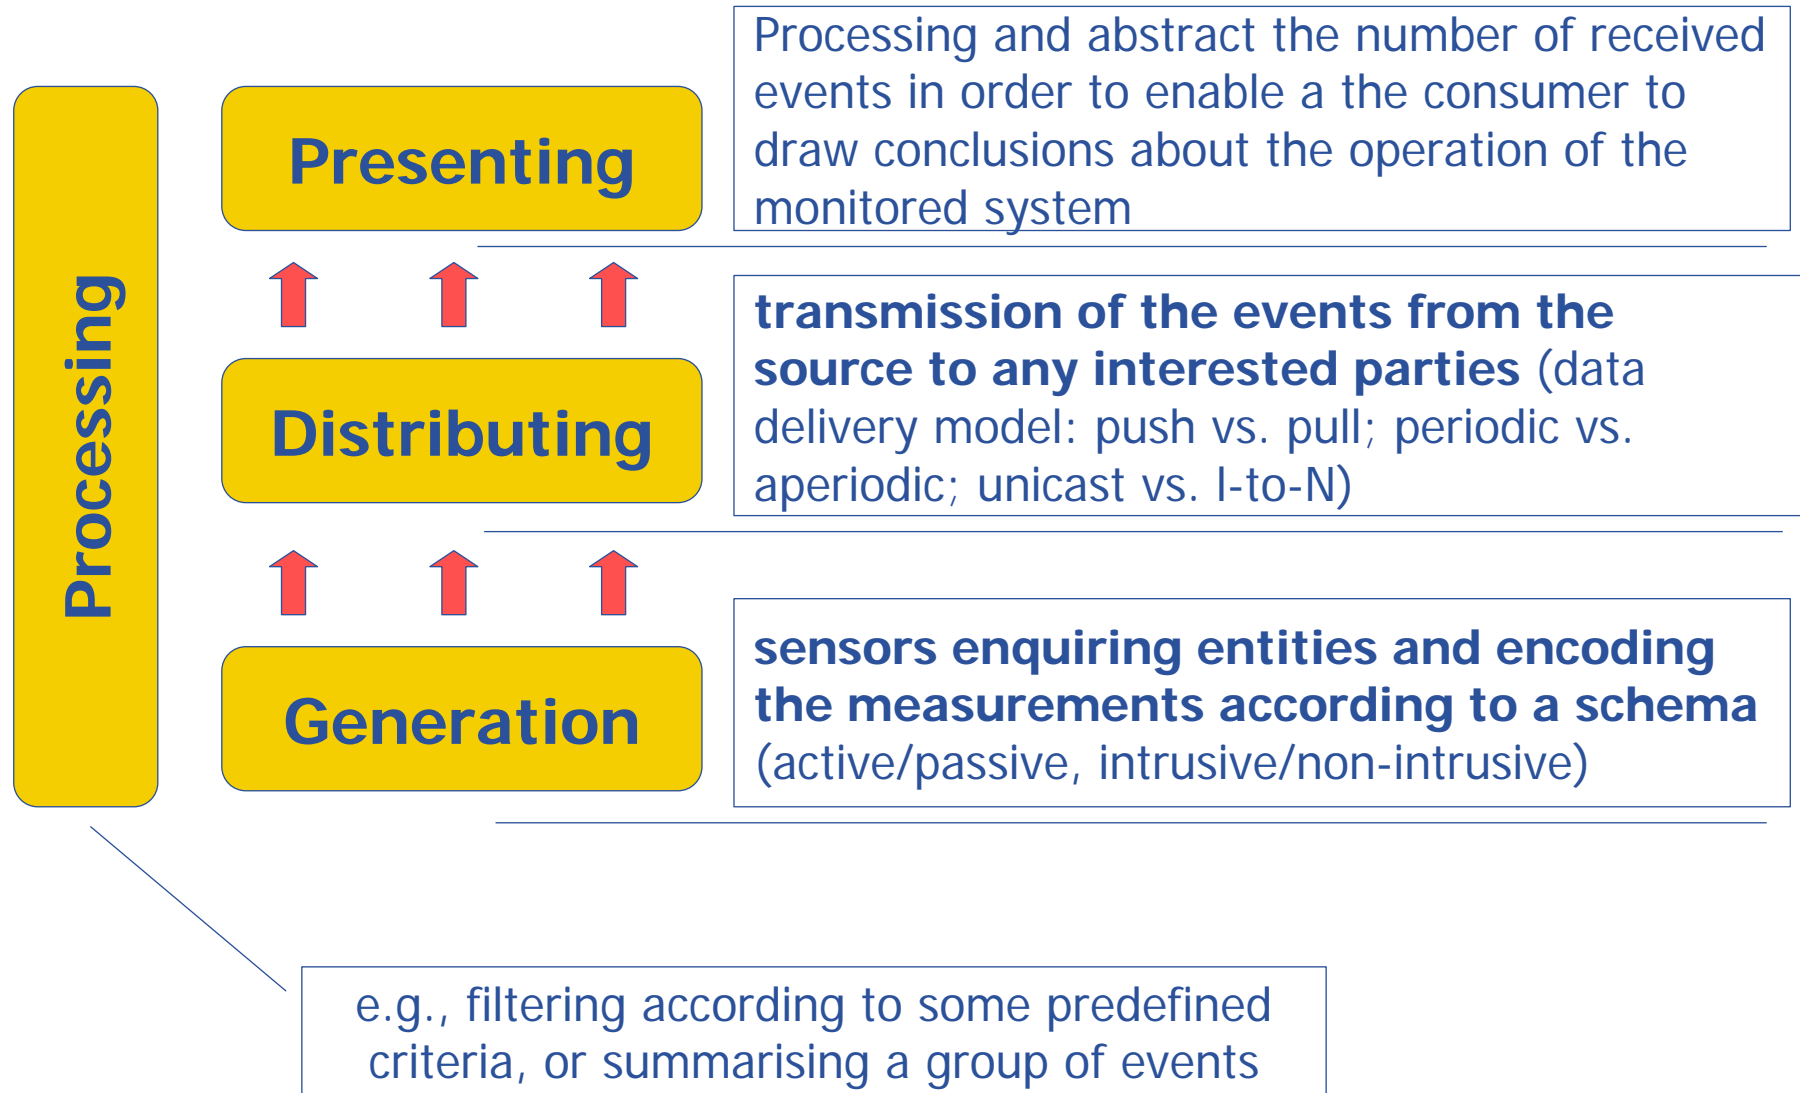


- ***Further information and documentation about DGAS can be found at:***
<http://www.to.infn.it/grid/accounting>

- **Grid Monitoring**
 - the activity of **measuring** significant **grid resources** **related parameters**
 - in order to
 - **analyze usage, behavior and performance of the grid**
 - **detect and notify**
 - *fault situations*
 - *contract violations (SLA)*
 - *user-defined events*

- **Measurement:** the process by which numbers or symbols are assigned to feature of an entity in order to describe them according to clearly defined rules
- **Event:** collection of timestamped data associated with the attribute of an entity [2]
- **Event schema** (or simply schema): defines the typed structure and semantics of all events so that, given an event type, one can find the structure and interpret the semantics of the corresponding event [2]

The four main phases of monitoring



- Virtual Organization:
 1. visualize at various aggregation levels the actual set of resources accessible to its members;
 2. Assess how Grid mapping functionalities from virtual to physical resources and users meet the members' demands
 3. analyze data retrospectively to understand how to improve the effectiveness of VO applications running in a Grid, as the target machine for different executions of the same application can vary over time


- Site Administrator:
 - Visualize the managed Grid services in order to see how they are being used/performing (possibly divided by VO)
- User:
 - Is my job “working” (e.g., consuming CPU?)
- Grid Operation Center:
 - Status of Grid services (e.g., WMS, Service Discovery, CE, SE)
 - Free/busy resources per site/per VO at a given time
 - Timely notification about fault situations

GridICE: architectural insight

- *generation* of events:
 - Sensors: typically perl scripts or c programs
 - Schema:
 - GLUE Schema v.1.1 + GridICE extension
 - *System related (e.g., CPU load, CPU Type, Memory size)*
 - *Grid service related (e.g., CE ID, queued jobs)*
 - *Network related (e.g., Packet loss) [5]*
 - *Job usage (e.g., CPU Time, Wall Time)*
 - All sensors are executed in a periodic fashion

- *distribution* of events:
 - Hierarchical model
 - **Intra-site**: by means of the local monitoring service
 - *default choice, LEMON (<http://www.cern.ch/lemon>)*
 - **Inter-site**: by offering data through the Grid Information Service
 - **Final Consumer**: depending on the client application
 - Mixed data delivery model
 - **Intra-site**: depending on the local monitoring service (push for lemon)
 - **Inter-site**: depending on the GIS (current choice, MDS 2.x, pull)
 - **Final consumer**: pull (browser/application), push (publish/subscribe notification service)


Example deployment in LCG2



Site	Domain	Computing Resources						Storage Resources					
		GK#	Q#	RunJob	WaitJob	SlotLoad	Power	WN#	CPU#	CPUload	Available	Total	%
INFN-PISA2	pi.infn.it	1	7	4	0	100%	13K	2	4	100%	3.5 TB	3.5 TB	0%
INFN-ROMA1	roma1.infn.it	1	2	2	7	5%	235K	21	42	2%	31.1 GB	33.9 GB	8%
INFN-Roma1-CMS	roma1.infn.it	1	2	0	0	0%	48K	5	11	0%	63.2 GB	65.9 GB	4%
INFN-ROMA1-VIRGO	roma1.infn.it	1	2	6	0	86%	45K	7	14	29%	28.4 GB	31.2 GB	9%
INFN-ROMA2	roma2.infn.it	1	6	1	0	10%	86K	5	18	5%	1.1 TB	1.1 TB	3%
INFN-ROMA3	roma3.infn.it	1	3	0	0	0%	38K	4	8	0%	956.7 GB	956.7 GB	0%
INFN-TORINO	to.infn.it	1	8	56	29	100%	297K	28	56	93%	420.6 GB	1.9 TB	79%
NA-ICAR-CNR	dma.unina.it	1	6	3	16	100%	8K	3	3	100%	-	-	-
SNS	sns.it	1	7	1	0	8%	24K	3	6	0%	64.6 GB	67.7 GB	5%
SPACI-LECCE	egee.unile.it	1	6	0	7	0%	6K	1	1	0%	-	-	-
TOKYO-LCG2	icepp.jp	1	2	0	0	0%	-	-	-	-	896.8 GB	1.8 TB	51%
LCG_KNU	knu.ac.kr	1	5	5	484	100%	-	-	-	-	59 GB	61.6 GB	4%
NIKHEF-ELPROD	nikhef.nl	1	6	205	53	94%	-	-	-	-	895.2 GB	1.7 TB	8%
saralcg2	matrix.sara.nl	1	16	42	11	92%	-	-	-	-	90.6 GB	104.4 GB	13%
NCP-LCG2	ncp.edu.pk	1	6	0	0	0%	-	-	-	-	42.3 GB	44.1 GB	4%
PAKGRID-LCG2	pakgrid.org.pk	1	6	0	0	0%	-	-	-	-	59.5 GB	60.3 GB	1%
CYFRONET-LCG2	cyf-kr.edu.pl	1	9	55	179	92%	-	-	-	-	865 GB	2 TB	58%
egee.man.poznan.pl	egee.man.poznan.pl	1	5	0	0	0%	-	-	-	-	255.4 GB	255.6 GB	0%
WARSAW-LCG2	fuw.edu.pl	1	3	0	0	7%	-	-	-	-	348.1 GB	348.1 GB	0%
LIP-LCG2	lip.pt	1	4	12	10	3%	113K	7	23	14%	346 GB	696.9 GB	51%
ROGRID-ICI	grid.ici.ro	1	6	7	7	78%	76K	5	16	54%	138.2 GB	141.6 GB	2%
ITEP	itep.ru	1	7	2	0	5%	-	-	-	-	63.3 GB	68.6 GB	8%
JINR-LCG2	jlnr.ru	1	5	2	0	10%	-	-	-	-	1.7 TB	1.7 TB	1%
RRC-KI	grid.kiae.ru	1	5	1	0	5%	-	-	-	-	762.5 GB	766.3 GB	0%
ru-Moscow-GCRAS-LCG2	wdcb.ru	1	3	0	0	0%	-	-	-	-	-	-	-
RU-Moscow-KIAM-LCG2	keldysh.ru	1	4	1	0	13%	-	-	-	-	97.4 GB	102.5 GB	5%
ru-Moscow-SINP-LCG2	sinp.msu.ru	1	9	33	1	69%	-	-	-	-	87.3 GB	104.9 GB	17%
ru-Novgorod-NOVSU-LCG2	novsu.ac.ru	1	5	0	0	0%	-	-	-	-	23.2 GB	23.4 GB	0%
ru-PNPI-LCG2	pnpi.nw.ru	1	1	58	28	100%	-	-	-	-	-	-	-
ru-PSN-LCG2	psn.ru	1	2	22	29	100%	-	-	-	-	167.8 GB	172.3 GB	3%
HPC2N	hpc2n.umu.se	1	1	50	75	54%	-	-	-	-	929.1 GB	934.9 GB	1%
NSC	nsc.liu.se	1	6	4	1	17%	-	-	-	-	59.4 GB	66.9 GB	11%
GOG-Singapore	ngpp.ngp.org.sg	1	3	0	0	3%	-	-	-	-	-	-	-

Site	Total	All		Broker		BDII		CE		SE		GC		Others					
		✖	⚠	✖	⚠	✖	⚠	✖	⚠	✖	⚠	✖	⚠	✖	⚠				
CERN-CIC	3	-	Disappeared	-	-	-	-	1	-	-	1	-	-	1	-	-	-	-	-
CNAF-T1	29	1	11	-	-	6	1	4	-	-	1	-	-	2	-	-	5	-	-
ESA-ESRIN	9	6	-	-	-	-	-	1	-	-	1	-	-	1	-	-	6	6	-
HPCC-UNILE	2	-	-	-	-	-	-	1	-	-	-	-	-	-	-	-	1	-	-
ifae	83	4	11	1	-	-	-	1	-	-	5	4	-	-	-	-	76	-	11
INAF-TRIESTE	1	-	-	-	-	-	-	-	-	-	1	-	-	-	-	-	-	-	-
INFN-BARI	27	1	-	-	-	-	-	1	-	-	1	-	-	1	-	-	24	1	-
infn-bologna	5	-	-	-	-	-	-	1	-	-	1	-	-	1	-	-	2	-	-
INFN-BOLOGNA-CMS	15	2	1	-	-	-	-	2	1	-	1	-	-	-	-	-	12	1	1
INFN-CAGLIARI	10	-	2	-	-	-	-	2	-	1	1	-	-	1	-	-	6	-	1
INFN-CATANIA	4	-	-	1	-	-	-	1	-	-	1	-	-	1	-	-	-	-	-
INFN-FERRARA	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
INFN-FIRENZE	19	-	-	-	-	-	-	1	-	-	1	-	-	1	-	-	16	-	-
INFN-FRASCATI	9	-	4	-	-	-	-	2	-	1	2	-	2	2	-	1	3	-	-
INFN-LECCE	3	-	-	-	-	-	-	1	-	-	1	-	-	-	-	-	1	-	-
INFN-LNL-2	17	1	-	-	-	-	-	1	1	-	1	-	-	1	-	-	14	-	-
INFN-LNL-LCG	94	-	12	-	-	-	-	1	-	-	1	-	-	-	-	-	92	-	12
INFN-MILANO	33	-	7	-	-	-	-	1	-	-	2	-	1	2	-	1	28	-	5
INFN-NAPOLI-ATLAS	21	-	1	-	-	-	-	2	-	-	1	-	1	1	-	-	17	-	-
INFN-PADOVA	64	-	6	7	-	1	-	1	-	-	4	-	-	1	-	-	50	-	6
INFN-PERUGIA	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
INFN-PISA	17	-	1	-	-	-	-	1	-	-	3	-	-	-	-	-	13	-	1
INFN-PISA2	5	-	1	-	-	-	-	1	-	-	1	-	1	1	-	-	2	-	-
INFN-ROMA1	24	-	-	-	-	-	-	1	-	-	1	-	-	1	-	-	21	-	-
INFN-Roma1-CMS	7	-	-	-	-	-	-	1	-	-	-	-	-	1	-	-	5	-	-
INFN-ROMA1-VIRGO	9	-	-	-	-	-	-	1	-	-	-	-	-	1	-	-	7	-	-
INFN-ROMA2	9	-	2	-	-	-	-	1	-	-	1	-	-	1	-	1	6	-	1
INFN-ROMA3	6	-	-	-	-	-	-	1	-	-	-	-	-	1	-	-	4	-	-
INFN-TORINO	30	-	-	-	-	-	-	1	-	-	-	-	-	1	-	-	28	-	-
INFN-TRIESTE	3	-	3	-	-	-	-	1	-	1	1	-	1	-	-	-	1	-	1
[mi.infn.it]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
NA-ICAR-CNR	5	-	-	-	-	-	-	1	-	-	-	-	-	1	-	-	3	-	-

Running/waiting jobs for a VO



Dissemination: <http://grid.infn.it/gridice>

- [1] S. Andreatta, N. De Bortoli, S. Fantinel, A. Ghiselli, G. L. Rubini, G. Tortone, M. C. Vistoli **GridICE: a monitoring service for Grid systems**, *Future Generation Computer System* 21 (2005) 559–571
- [2] B. Tierney, R. Aydt, D. Gunter, W. Smith, M. Swamy, V. Taylor, R. Wolski, **A Grid Monitoring Architecture**, GFD-I.7
- [3] S. Zaniolas, R. Sakellariou, **A taxonomy of grid monitoring systems**, *Future Generation Computer Systems* 21 (2005) 163–188
- [4] M. Franklin, S. Zdonik, **“Data In Your Face”**: Push Technology in Perspective, ACM SIGMOD '98, Seattle, WA, USA
- [5] S. Andreatta, A. Ciuffoletti, A. Ghiselli, C. Vistoli. **Monitoring the connectivity of a Grid**. Proceedings of the 2nd International Workshop on Middleware for Grid Computing (MGC 2004) in conjunction with the 5th ACM/IFIP/USENIX International Middleware Conference, Toronto, Canada, October 2004.
- [6] S. Andreatta, N. De Bortoli, S. Fantinel, G.L. Rubini, G. Tortone. ***Design and Implementation of a Notification Model for Grid Monitoring Events***. CHEP04, Interlaken (CH), Sep 2004