



Enabling Grids for E-science

EGEE Middleware reengineering

*Claudio Grandi
(INFN – Bologna)*

*EGEE Final Review
23-24 May 2006
CERN*



www.eu-egee.org
www.glite.org

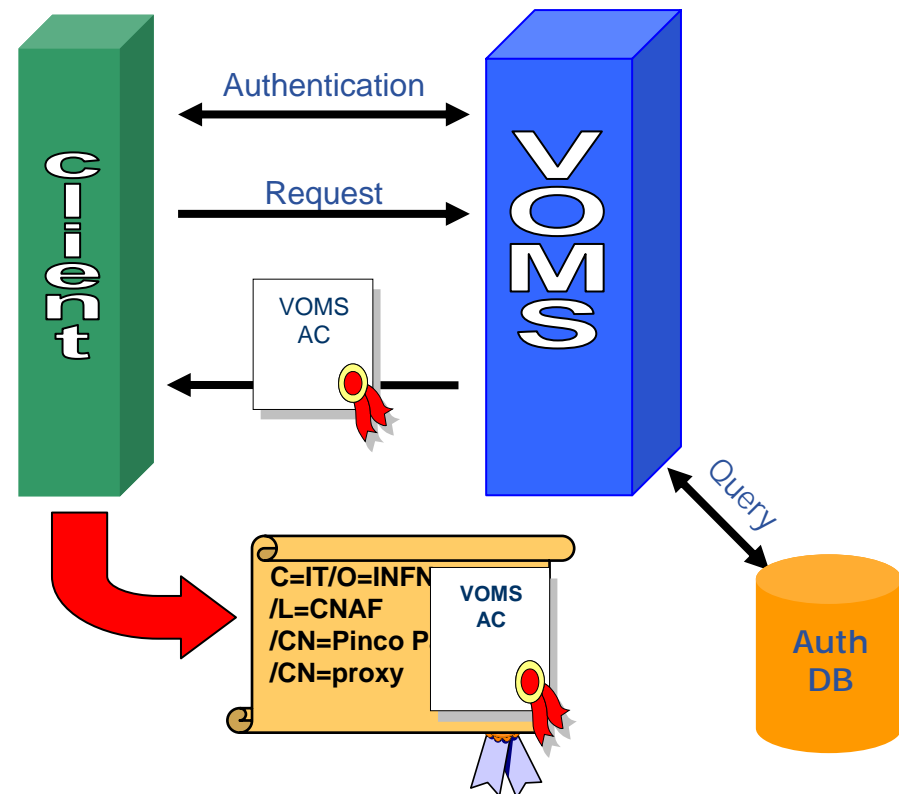


- **Major JRA1 achievements since the 2nd EGEE Review**
- **Answers to recommendations at the 2nd EGEE Review**
- **Summary**

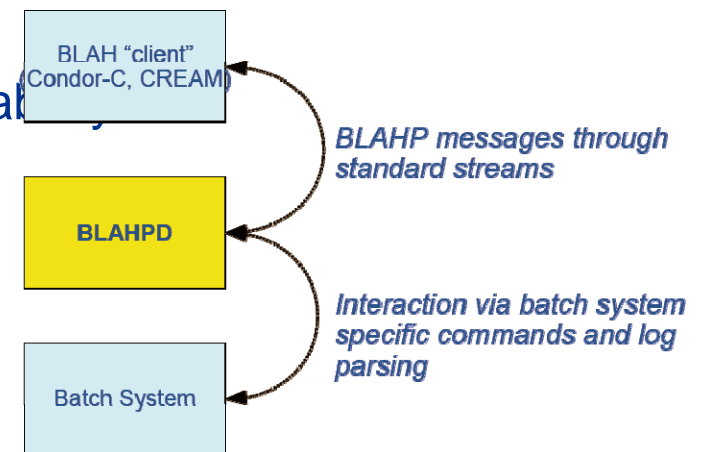
- **Integration & test of gLite 1.5 (released in January)**
 - Including preparation of release notes and documentation
- **Quick-fixes to previous releases**
- **ETICS project started in January**
 - EGEE build system started migration to the new ETICS system
- **Fusion of the JRA1 testing team and the SA1 certification team**
- **Consolidation of gLite 1.5 components for gLite 3.0**
- **Integration & test of the gLite 3.0 release (by new SA3!)**
 - Major effort since LCG 2.7.0 and gLite 1.5 were developed in not completely consistent environments
- **Address feedback received from users on the PPS**
- **Development for next releases (formerly gLite 3.1)**

- **The Relational Grid Monitoring Architecture (R-GMA)** provides a uniform method to access and publish both information and monitoring data.
 - From a user's perspective, an R-GMA installation currently appears similar to a single relational database.
 - Relational implementation of the GGF Grid Monitoring Architecture (GMA)
 - Included in gLite 3.0 and used for job and infrastructure monitoring
- **The gLite Service Discovery** provides a standard set of methods for locating Grid services
 - hides underlying information system
 - plug-ins for R-GMA, BDII and XML files (others could be developed if required)
 - It is in gLite 3.0 and used by WMS and Data Management clients
- **The APEL system uses R-GMA to propagate and display job accounting information for infrastructure monitoring**
 - It is in gLite 3.0
- **The Information System on the production infrastructure is BDII**

- Bare certificates are not enough for defining user capabilities on the Grid
- Users belong to VOs, to groups inside a VO and may have special roles
- **VOMS** provides a way to add attributes to a certificate proxy:
 - mutual authentication of client and server
 - VOMS produces a signed Attribute Certificate (AC)
 - the client produces a new proxy that contains the attributes
- The attributes are used to provide the user with additional capabilities according to the VO policies.
- **VOMS is in gLite 3.0** and is well established on the production system



- **gLite CE (Condor-C GSI enabled)**
 - It is in gLite 3.0 and will be deployed on selected sites together with an LCG-CE
- **CREAM** the new lightweight web service CE
 - is still on development and it is not in gLite 3.0
- **On the production infrastructure all sites will deploy an LCG-CE**
 - based on GT2 GRAM
 - will be dismissed when other CEs prove reliable
- **Batch Local ASCII Helper (BLAH)**
 - Interface between the CE and the Local Resource Manager System
 - e.g. submit, cancel and query
 - It is in gLite 3.0 in the gLite CE
- **CEMon** is a web service to publish the status of a computing resource and of running jobs
 - Supports both synchronous queries or asynchronous notification of events
 - It is in gLite 3.0 but the default is that the WMS queries the BDII



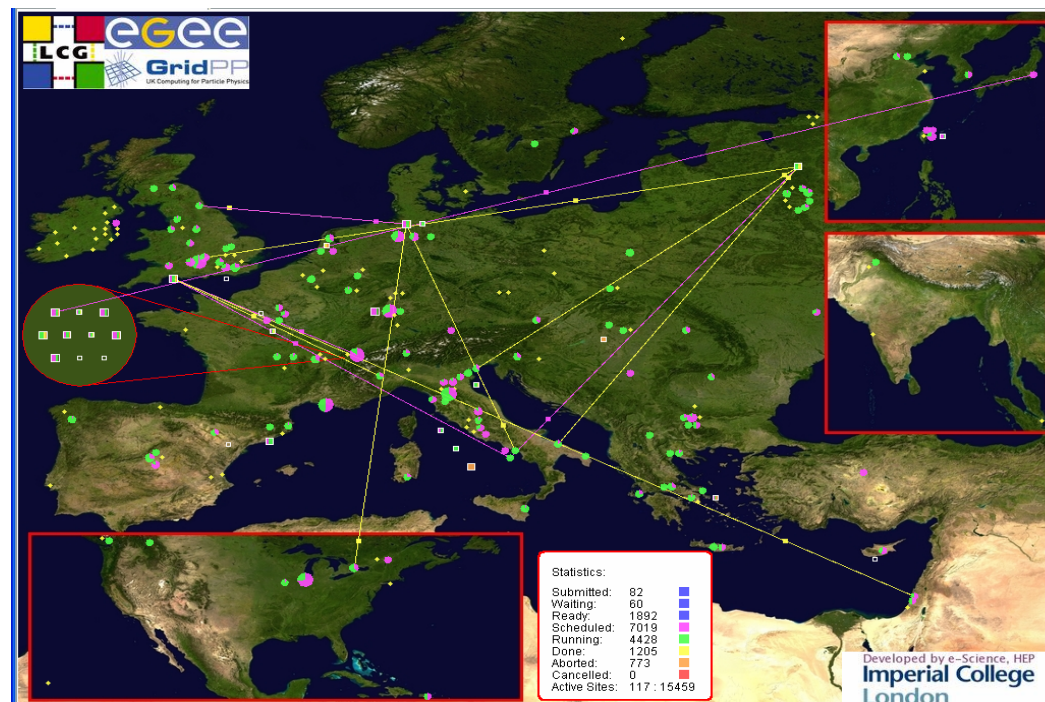
- **Logging and Bookkeeping service**

- Tracks jobs during their lifetime (in terms of events)
- L&B Proxy provides faster, synchronous and more efficient access to L&B services to Workload Management Services
- It is in gLite 3.0

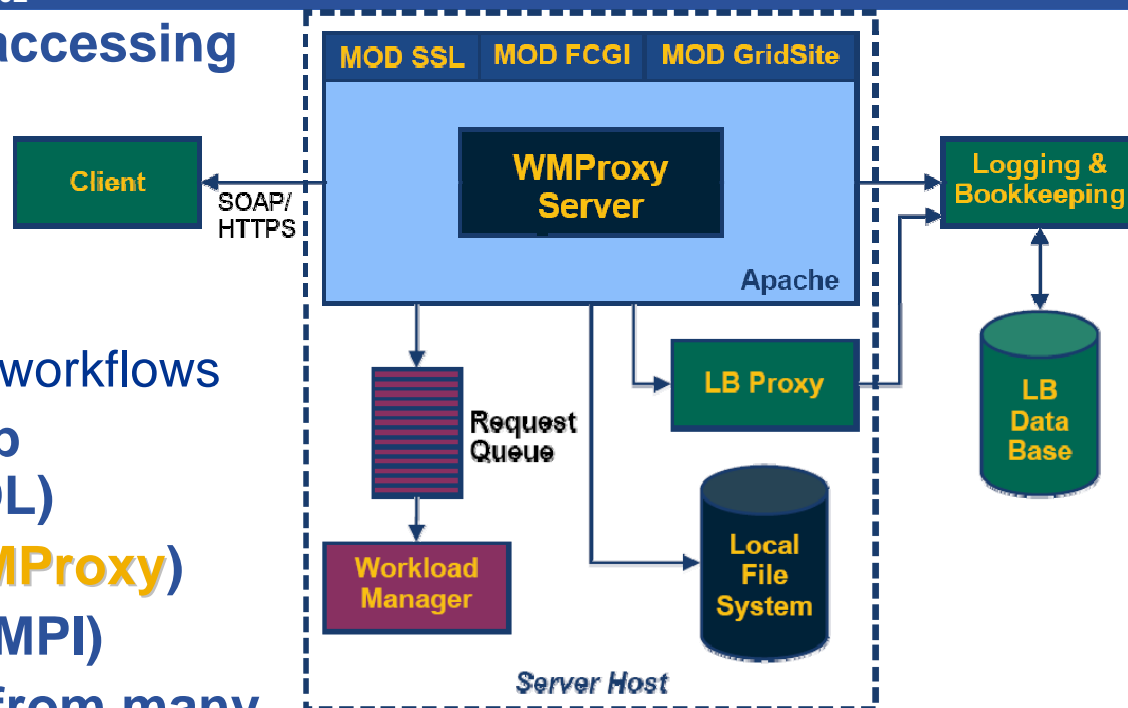
- **Job Provenance**

- Long term job information storage
- Useful for debugging, post-mortem analysis, comparison of job executions in different environments
- Useful for statistical analysis
- It is not in gLite 3.0

- In EGEE II this activity will be carried on by the new CZ cluster



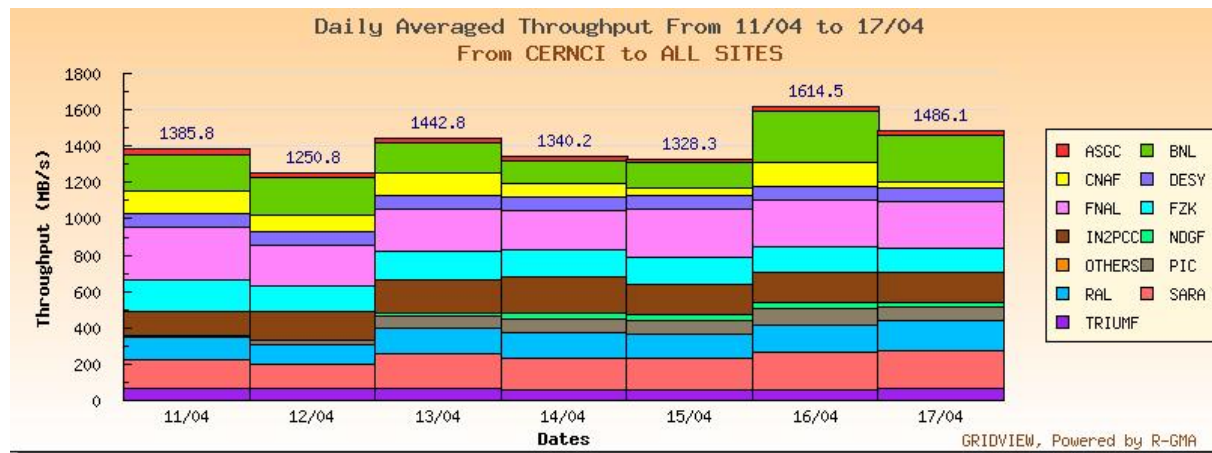
- The **WMS** Helps the user accessing computing resources
 - resource brokering
 - management of job input and output
 - management of complex workflows
- Job specification via a Job Description Language (JDL)
- Web service interface (**WMProxy**)
- Support for parallel jobs (MPI)
- Collection of information from many sources (CEMon, bdll, R-GMA)
- Supports different Data management interfaces (DLI, StorageIndex)
- Job File Perusal: file peeking during job execution
- Submissions via Condor-G
 - in development **ICE**, for submissions to CREAM CE
- The **gLite WMS** is in **gLite 3.0** together with the old LCG-RB



- **DGAS** collects Grid accounting information
 - (User, JobId, user VO, VOMS FQAN, system usage (cpuTime, ...), ...)
 - Secure storage and transfer of accounting records
 - signed and encrypted
 - Compliant with privacy requirements
 - Sites may keep personal information local if requested by local laws
 - It is not in gLite 3.0, is being integrated with APEL
- **GPBOX** is a tool to define, store and propagate fine-grained VO policies
 - based on VOMS groups and roles
 - enforcement of policies done at sites: sites may accept/reject policies
 - May be interfaced to dynamic sources of information
 - e.g. an accounting system to provide fair share
 - It is not in gLite 3.0

- **Storage Elements currently using solutions from other projects**
 - Various implementation from LCG and other external projects
 - disk-based: DPM, dCache
 - tape-based: Castor, dCache
 - Common interface: SRMv1, migrating to SRMv2
 - Support for ACLs in DPM, in future in Castor and dCache
- **Posix-like file access with gLite I/O**
 - support of ACLs from Fireman data catalog
 - It is in gLite 3.0 with limited support
 - Investigating whether it can be replaced by GFAL (by LCG)
 - supports ACLs in the SRM layer
- **Hydra keystore stores keys for data encryption**
 - 3 instances: at least 2 need to be available for decryption
 - Demonstrated at the SRM/DICOM demo at EGEE 4th conference in Pisa (system with gLite I/O, Fireman and AMGA)
 - It is in gLite 3.0 with limited support. Will be fully supported in the future

- **Data Catalog: Fireman**
 - It is in gLite 3.0 with limited support
 - Investigating whether it can be replaced by **LFC** (by LCG)
- **AMGA Metadata Catalog (Joint JRA1-NA4 (ARDA) development)**
 - Support for VOMS-based permissions for data records
 - Support for replication
 - It is in gLite 3.0 with limited support. Will be fully supported in the future
- **Reliable, scalable and customizable File Transfer System (FTS)**
 - Manages transfers through channels
 - mono-directional network pipes between two sites
 - Web service interface
 - Support for different user and administrative roles
 - It is in gLite 3.0 and has been used on the production system for Service Challenges
 - Sustained file transfer record (1.6 GB/s)



“Have more direct pathways and programs where the teams from developers, testers, infrastructural providers, and above all, application developers and users spend focused efforts to identify the usage and concerns with the current LCG-2 and gLite, instead of relying on a fairly long pathway from the application end to the development”

- **Adoption of the process *a-la* EGEE II in fall 2005**
 - Process controlled by the Technical Coordination Group (TCG) keeps the development focused
 - Task Forces with applications, testers and deployment experts
- **Convergence of the LCG and gLite middleware stacks**
 - gLite 3.0 released in May. Deployed on the Production System
 - Middleware deployed on the Pre-Production System (PPS) and exposed to the users in mid-March

“Continuously assess user application feedback, especially in the light of introduction of new services, in order to be able to judge whether continued investment into the R&D of that particular feature would have high return on its value”

- **Achieved with the creation of TCG and Task Forces**
 - Prioritization of the inclusion of gLite 1.5 and LCG 2.7.0 components into gLite 3.0
 - Definition of high priority additions for the next releases
 - Feed-back from the Task Forces about functionality and performance of the components deployed on the PPS
- **Still to be done:**
 - Creation of a *“preview”* system when new components may be exposed to the users at an early stage of development
 - Early feedback before effort is spent in integration and testing

“Clarify and advertise a more conservative (in term of time span) and comprehensive release cycle plan for gLite”

- **After gLite 3.0 adopt a continuous release process**
 - no more *big-bang* releases with fixed deadlines for all
 - develop components as requested by users and sites
 - deploy or upgrade as soon as testing is satisfactory
- **Major releases synchronized with large scale activities of VOs (SCs)**
 - include what is available at that time
- **Support more versions of client code**
- **Keep servers backward compatible as much as possible**
- **This is handled by SA3 in the EGEE II process**
 - Integration process using ETICS tools
 - The new testing team includes the former JRA1 testing and SA1 certification teams

“Revise the gLite development process to fully integrate the Technical Coordination Group and application developers”

- **Started in fall 2005 with TCG and Task Forces**

“Investigate the deliverables of other international grid R&D activities and identify where deliverables could be shared in a mutually collaborative fashion to achieve rapid international interoperations with grids outside of EU”

- **Design Team meetings**
 - Include OSG experts
- **Frequent bi-lateral meetings**
 - OSG, NAREGI, NorduGrid
- **Joint efforts:**
 - GLUE (OSG)
 - glexec (GridPP & Condor)
 - CREAM (OMII-EU)
- **Components developed by other projects, often with EGEE support:**
 - Build system (ETICS)
 - Condor-C / gLite CE (VDT)
 - MyProxy (VDT)
 - VOMRS (VDT)
- **Components developed by EGEE included in other projects systems:**
 - VOMS (VDT)
 - MyProxy/VOMS interface (VDT)
 - BLAHPD (VDT)
 - CEMon (VDT)
 - L&B (will be in VDT)
 - LCMAPS (Globus)
 - AAA toolkit (Globus)
- ***Note that development with external projects that do not have commitments on the development or simply different time scales may be very difficult***

“Identify in the middleware stack which parts of gLite is “conformant” to standards activities within GGF and where it is currently not”

- **Note that some standards are newer than the corresponding gLite middleware**
 - Sometimes it is difficult to adopt a new standard without breaking the existing functionalities
- **Standards adoption:**
 - LCAS, LCMAPS (and glexec) will support the remote authorization callout format based on SAML when specified by OGSA-Authz
 - AAA toolkit supports SAML, XACML, OGSA AuthZ port type and GT4 AuthZ callouts
 - Trustmanager and Util-java compliant to TLS and RFC2252
 - Storage Elements have SRMv1 interface (migrating to SRMv2)
- LFC & Fireman have DLI and StorageIndex interfaces
- CREAM, CEMon, WMPProxy use SOAP, WSDL and WS-I
- building JDSL to JDL translator
- GLUE schema agreed with OSG
- Activity going on with NAREGI to exchange information between GLUE and CIM
- R-GMA developed according to GGF-INFOD; discussions to modify standard
- Service Discovery is compatible (and extends) UDDI interface
- **New proposed standards**
 - VOMS Attribute Certificate (GGF)
 - CREAM contribution to GGF-BES

“Make more effective use of the Industry Forum to realize industrial involvement in the development to achieve smoother technology transfer”

- **Good collaboration with industrial partners:**
 - Datamat is developing important parts of the WMS (WMPProxy) and the User Interface
 - CERN OpenLab, contributing to the porting of gLite on IA-64 architectures
- **First EGEE Industry Day organized in Paris (27 April 2006)**
- **More input needed by the Industry Forum!**

- **EGEE provides a complete middleware stack**
- **Developed according to a well defined process**
 - Controlled by the EGEE Technical Coordination Group
- **gLite 3 will be available on the production infrastructure in June**
- **Development is continuing to provide increased robustness, usability and functionality**



Lightweight Middleware for
Grid Computing

www.glite.org