EGEE Middleware reengineering

Claudio Grandi – JRA1 Activity Manager - INFN

EGEE-II Final EU Review (CERN)
8-9 July 2008

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

Outline

- Activity goals and description
- Achievements and description of work done
- Issues
- Future plans
- Summary

JRA1 activity goals

- **Continue to support and evolve the gLite open source implementation of application-independent grid middleware**
  - Application-independent ***foundation services***
  - Set of ***higher-level services*** working on top of the foundation
  - Follow a Service Oriented Architecture
- **Activity targeted to the support of the Production System**
  - Gradually deploy new components on the PS, support and maintain them
    - Prompt fixing of bugs and support to the Global Grid User Support (**GGUS**)
    - Stability, scalability, manageability
  - Further evolve the middleware stack
    - Address applications needs
    - Interoperability with other infrastructures
    - Attention to emerging standards

JRA1 in Numbers

JRA1 Partners

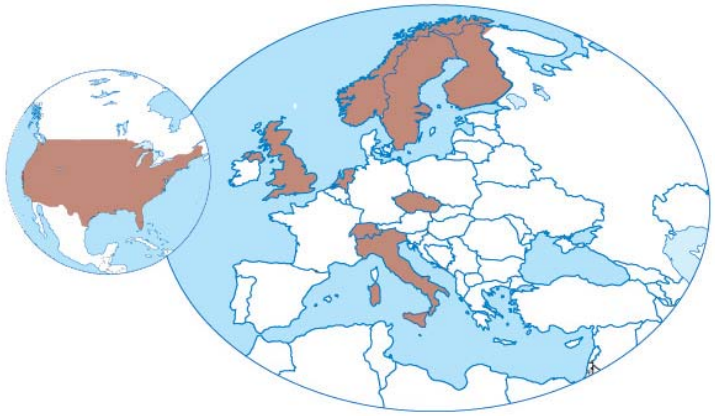

EGEE-II Budget

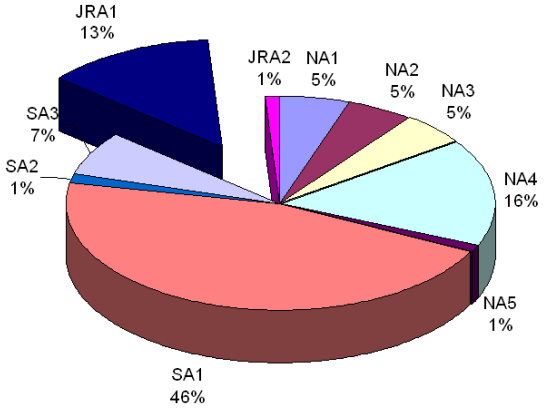

Manpower: 11 partners, 9 countries, 51.5 FTE

Partner	Country	FTEs	People
CERN	Switzerland	3	3
CESNET	Czech Republic	5	12
CCLRC	UK	7	7
Elsag Datamat	Italy	5	9
INFN	Italy	20,5	25
SWITCH	Switzerland	3	4
UH HIP	Finland	3	3
FOM	Netherlands	2	3
UvA	Netherlands	1	3
UiB	Norway	1	1
KTH	Sweden	1	1
U.Chicago	USA		
U.South California	USA		
U.Wisconsin	USA		
Total		**51,5**	**71**

Achievements in the first year

- **gLite 3.0 release**
  - Convergence of the LCG-2 and EGEE software
  - First production release of many EGEE components
- **Started using the new software process**
  - Including the continuous release process
- **Preview Test Bed**
  - New functionalities exposed to the users in an early stage of development
- **Experimental services**
  - Production-scale tests of services and fast development cycles
- **Shibboleth integration in gLite**
  - Short Lived Credential Service (SLCS) and VOMS Attributes to Shibboleth (VASH) service
- **Significant improvement in the performance of the gLite Workload Management System**

Achievements in the second year

- **gLite 3.1 release**
  - SL4 support
  - New version of VDT (including Globus Toolkit 4)
  - Full adoption of ETICS
  - Full adoption of the new software process
- **Long term sustainability: *gLite restructuring***
- **Strategy for the Compute Element revisited (CREAM)**
- **Changes in authorization to support *pilot jobs***
- **Review of gLite authorization – framework re-design**
- **Stability and scalability of WMS/LB and CE addressed**
- **Improvements in interoperability (especially OSG, ARC)**
- **Support of SRMv2.2 in Data Management components**
- **Encrypted Data Storage prototype delivered**
  - Used by EGEE Biomed community and also by EUIndiaGrid
- **Re-designed R-GMA implemented**

gLite restructuring

- **The goals of the gLite restructuring have been presented in detail during the previous review**
  - Address the long term sustainability of gLite, with particular attention to the portability to other platforms
- **Started on 28/5/07 and continued to the end of 2007**
  - Coexistence with other high priority activities and the support to the production infrastructure continued
- **The dependencies of all packages have been critically reviewed by gLite experts (other JRA1 developers, SA3 and SA1 experts, etc…)**
  - Rationalization of packages and dependencies

gLite restructuring outcome

- **Internal code reorganization and clean-up**
  - Developers are often pushed to release quick fixes but a more organic approach is needed
- **Removal of obsolete components**
  - Many functionalities were not used by applications but were still supported
- **Dependencies versions unified**
  - Use well supported external distributions (e.g. Jpackage for Java libraries) whenever possible
- **Client-server separation improved**
  - Lighter and more portable clients to be installed on User Interfaces and Worker Nodes
- **"Culture of maintainability" improved in the developers' community**

service	SL4/SLC4/i386
glite-WN	Released
glite-UI	Released
glite-AMGA_postgres	Released
glite-AMGA_oracle	Released
glite-BDII	Released
lcg-CE	Released
glite-FTM	Released
glite-LB	Released
glite-LFC_mysql	Released
glite-LFC_oracle	Released
glite-MON	Released
glite-PX	Released
glite-SE_dcache_*	Released
glite-SE_dpm_disk	Released
glite-SE_dpm_mysql	Released
glite-TORQUE_utils	Released
glite-TORQUE_client	Released
glite-TORQUE_server	Released
glite-VOMS_oracle	Released
glite-VOMS_mysql	Released
glite-VOBOX	Released
glite-WMS	Released
glite-CREAM	Certification
glite-FTA_oracle	Certification
glite-FTS_oracle	Certification
glite-SE_classic	Certification

Compute Element strategy

- **Legacy service in production on the EGEE infrastructure**
  - **LCG-CE**: based on Globus GRAM
- **Two development streams in EGEE:**
  - **gLite CE**: based on Condor-C, allows for VO-schedulers
  - **CREAM**: lightweight Web Service, compatible with BES standard
  - Common components: interface to local resource managers (BLAH) and security infrastructure (glexec, LCAS/LCMAPS)
  - Both CEs showed to be adequate to the infrastructure needs
- **The project decided that it was not possible to support both systems and decided to focus on the CREAM CE**
  - Approaches to interoperability solutions do not change
    - Gateways built using BLAH as a translator to other infrastructures
    - Condor-G submissions to CREAM being finalized
  - Most of the development of the gLite-CE is reused in CREAM
  - The LCG-CE guarantees the operations until CREAM reaches the needed maturity

CREAM status

- **Stress test on CREAM shows promising performance**

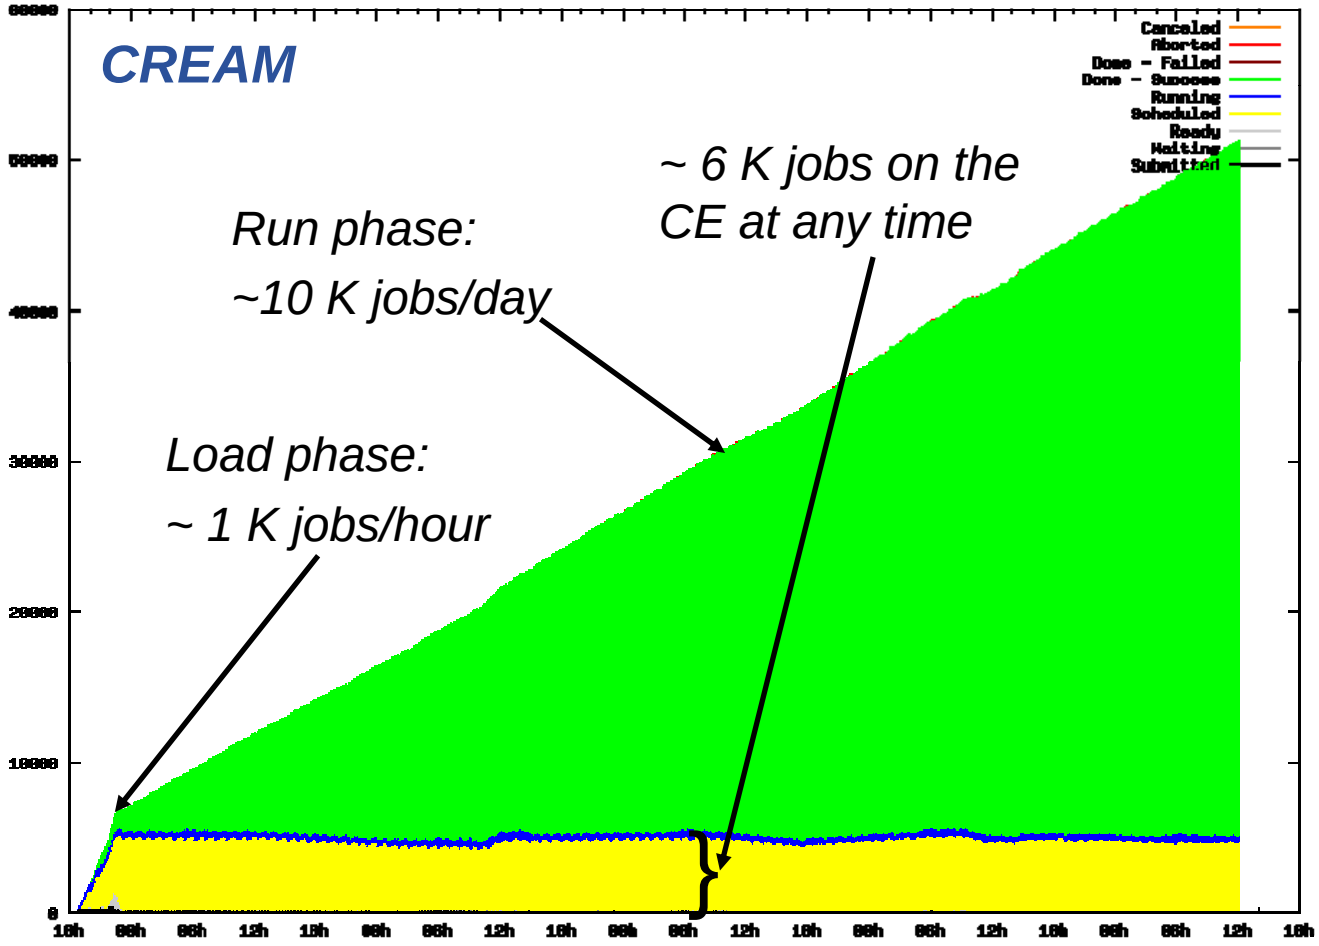

- Test system with 1000 batch slots
- > 90000 jobs by 50 users in 8 days
- Aggressive load phase
- Only a few failures observed

- **CREAM is now in certification and all issues related to its use in a production environment are being addressed**

Authorization framework

- **C. Witzig (SWITCH) nominated EGEE Security Architect**

- **VOMS-based authorization sometimes inconsistent in gLite services**
- **The Milestone document MJRA1.7 has been re-scoped:**
  - "Authorization mechanisms in gLite"
  - Critical analysis of gLite authorization – inconsistencies highlighted
- **Medium term solution (6 months) proposed**
  - Already discussed in the Technical Coordination Group (TCG)
  - The work-plan is being finalized in the context of EGEE-III
- **Longer term solution (1.5 years) being re-designed**
  - The high level architecture has already been discussed in the TCG
  - The detailed design and the work-plan are being finalized

JRA1 All-Hands meetings

Regular JRA1 plenary meetings
Contributions by external experts

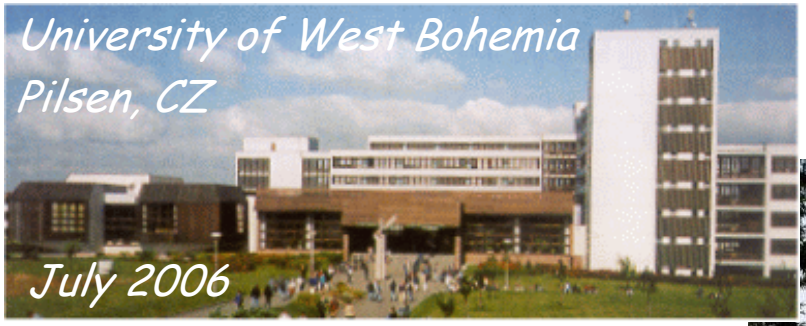

Community-building **events for a group distributed over different countries**

Issues

- **The adoption of ETICS took much more than expected**
  - ETICS helped significantly in understanding the code but
    - Has been adopted when it was in an early stage of development
    - Its focus initially has been mainly on integration and packaging
    - Now addressing also developers' needs but it is still not completely satisfactory in terms of performance and functionality
  - Managing a large number of components, configurations and dependencies is still time consuming
- **The software process is weak in the phase of passing the code from JRA1 to SA3**
  - *Clusters of competence* created in EGEE-III
- **Limited use of the preview test bed**
  - The effort has been concentrated on the experimental services
- **Inconsistencies in the gLite authorization mechanisms**
  - Will be addressed in EGEE-III, as described above

Future plans

- **JRA1 continues in EGEE-III**
  - Manpower reduction (-62% w.r.t. EGEE-II)
  - Concentrate on software maintenance (stability, scalability)
  - Address the limitations of the authorization framework
  - Address interoperability issues mainly through adoption of established standards
- **Continue the collaboration with other partners**
  - Several gLite components are included in VDT and used on OSG
  - Software of the gLite CE included in Condor and part of Red Hat
  - gLite used on infrastructures other than EGEE
- **Creation of a gLite Consortium**
  - Address the needs of middleware development in the **EGI** era

Summary

- **The main achievement of JRA1 in EGEE-II is the consolidation of gLite, addressing the long term sustainability of the middleware**
  - Adoption of a software process and a build infrastructure capable of assuring the long term sustainability of the software
  - Reduction of the complexity of the software that did grow beyond a sustainable level, ensuring in future the possibility to port the software to new platforms
  - Consolidation of the key services to reach the stability and scalability needed by the infrastructure
  - Attention to the needs of the applications needing to inter-operate with other grid infrastructures, both by adopting international standards and by developing specific solutions
  - Consolidation of security practices in particular in the field of authorization

Security

VOMS, VOMS-Admin	Tool for Attribute Certificate management, including its WS-based administrative interface
LCAS/LCMAPS/SCAS	Framework for authorization and mapping to local user accounts; includes now a service (SCAS) offering the same functionality on a central site service; will converge in the revised AuthZ framework
TrustManager Util-Java	Certificate validation tool and security utilities
gJAF	Framework for policy-based authorization in Java; will converge in the revised AuthZ framework
Delegation	WS-based framework for proxy certificate delegation
Job Repository	Auditing tool for Computing Elements
Glexec	Tool for local identity switching based on proxy and attribute certificates. Used on the Compute Elements and also on the Worker Nodes (for pilot jobs)

Test-utils	Test suite to generate certificates used to test middleware
CGSI-gSOAP	Library that allows gSOAP to use GSI authentication. Implements the https and httpg protocols
Proxy-renewal	Tool for proxy and Attribute Certificate renewal. Talks to MyProxy and VOMS
G-PBOX	Framework for XACML-based policies management. Includes a Policy Decision Point, interface libraries and GUI; will converge in the revised AuthZ framework
SLCS	Service that issues Short-lived Credentials based on a successful authentication at a Shibboleth Identity Provider
VASH service	Service that manages attributes flow from Shibboleth to VOMS
Resource Access	
gLite CE	Computing Element based on GSI-enabled Condor-C

LCG CE	Computing Element based on the Globus Toolkit 4 GRAM (by SA3)
CREAM/CEMon	WS-based Computing Element. Compliant with the current definition of OGF-BES and JSDL
BLAHP	Layer that interfaces a Computing Element to the local batch system
	Job Management Services
gLite Workload Management / ICE	Service responsible for the distribution and management of computing tasks on available Computing Elements. Includes a WS-based interface and multi language client tools
LCG RB	Service responsible for thedistribution and management of computing tasks on LCG-CEs (by SA3)
Logging & Bookkeeping	Service to track jobs being processed by middleware components. Provides a query interface to the users
Job Provenance	Service to archive job information gathered from the Logging and Bookkeeping. Provides provenance information about the jobs and data mining capabilities.

Data Management Services

DPM	Lightweight Storage Element for disk pools offering a standard SRM interface
GFAL lcg_utils	Library to offer POSIX-like interface to SRM-based Storage Element Collection of tools that offer UNIX-like file management on Storage Elements
gLiteIO / FiReMan	Service providing secure remote access to files stored on Storage Elements / File and Replica Catalogue that offers a hierarchical UNIX-like view of files stored on grid Storage Elements
LFC	File Catalogue that offers a hierarchical UNIX-like view of files stored on grid Storage Elements
FTS	Service to manage file transfers on network channels between Storage Elements. Includes a WS-based interface
Encrypted Data Storage	Framework for access to encrypted data, including the Hydra secure key-store and the possibility to split keys over multiple servers with SSSS
AMGA	General purpose metadata catalogue (by NA4)

Information Services, Monitoring and Accounting

BDII	LDAP database populated with information of grid resources. Information is stored in standard GLUE format (by SA3)
R-GMA	Relational implementation of the OGF Grid Monitoring Architecture. May store any kind of information produced by grid components or applications
Service Discovery	Library providing a standard interface for services location. Uses BDII, R-GMA or flat files as back-ends
DGAS	Accounting service providing sensors to collect information on Computing Elements, databases for storage of usage records and a system for securely move them
APEL	Accounting services providing sensors to collect information on Computing Elements, and a framework based on R-GMA to pass usage records to the Grid Operation Centre (by SA1)

www.glite.org