



Enabling Grids for E-science

JRA1 status

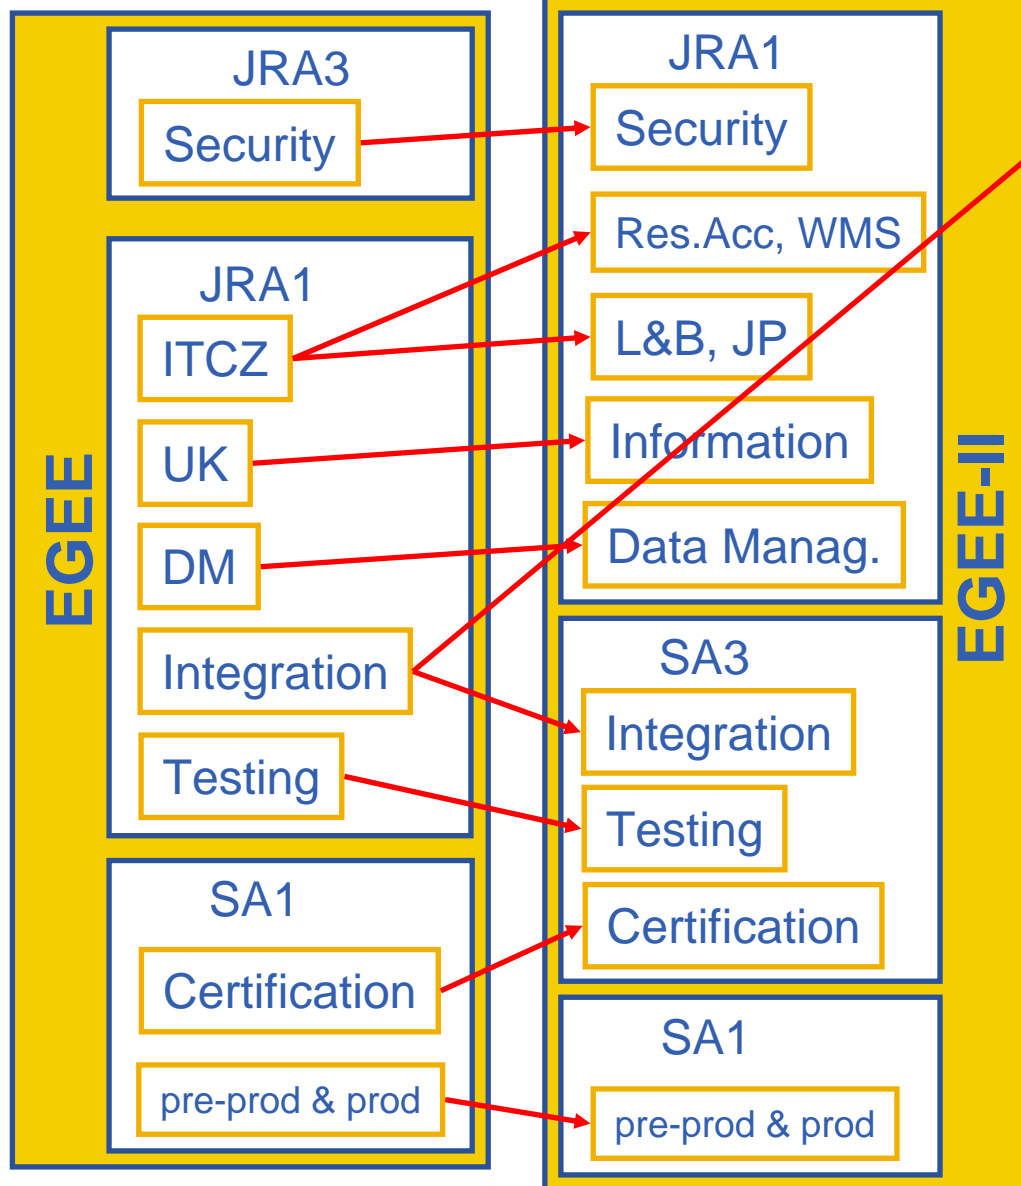
Claudio Grandi - JRA1 Activity Manager - INFN

*EGEE II All Activity Meeting
CERN, 24-25 August 2006*

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

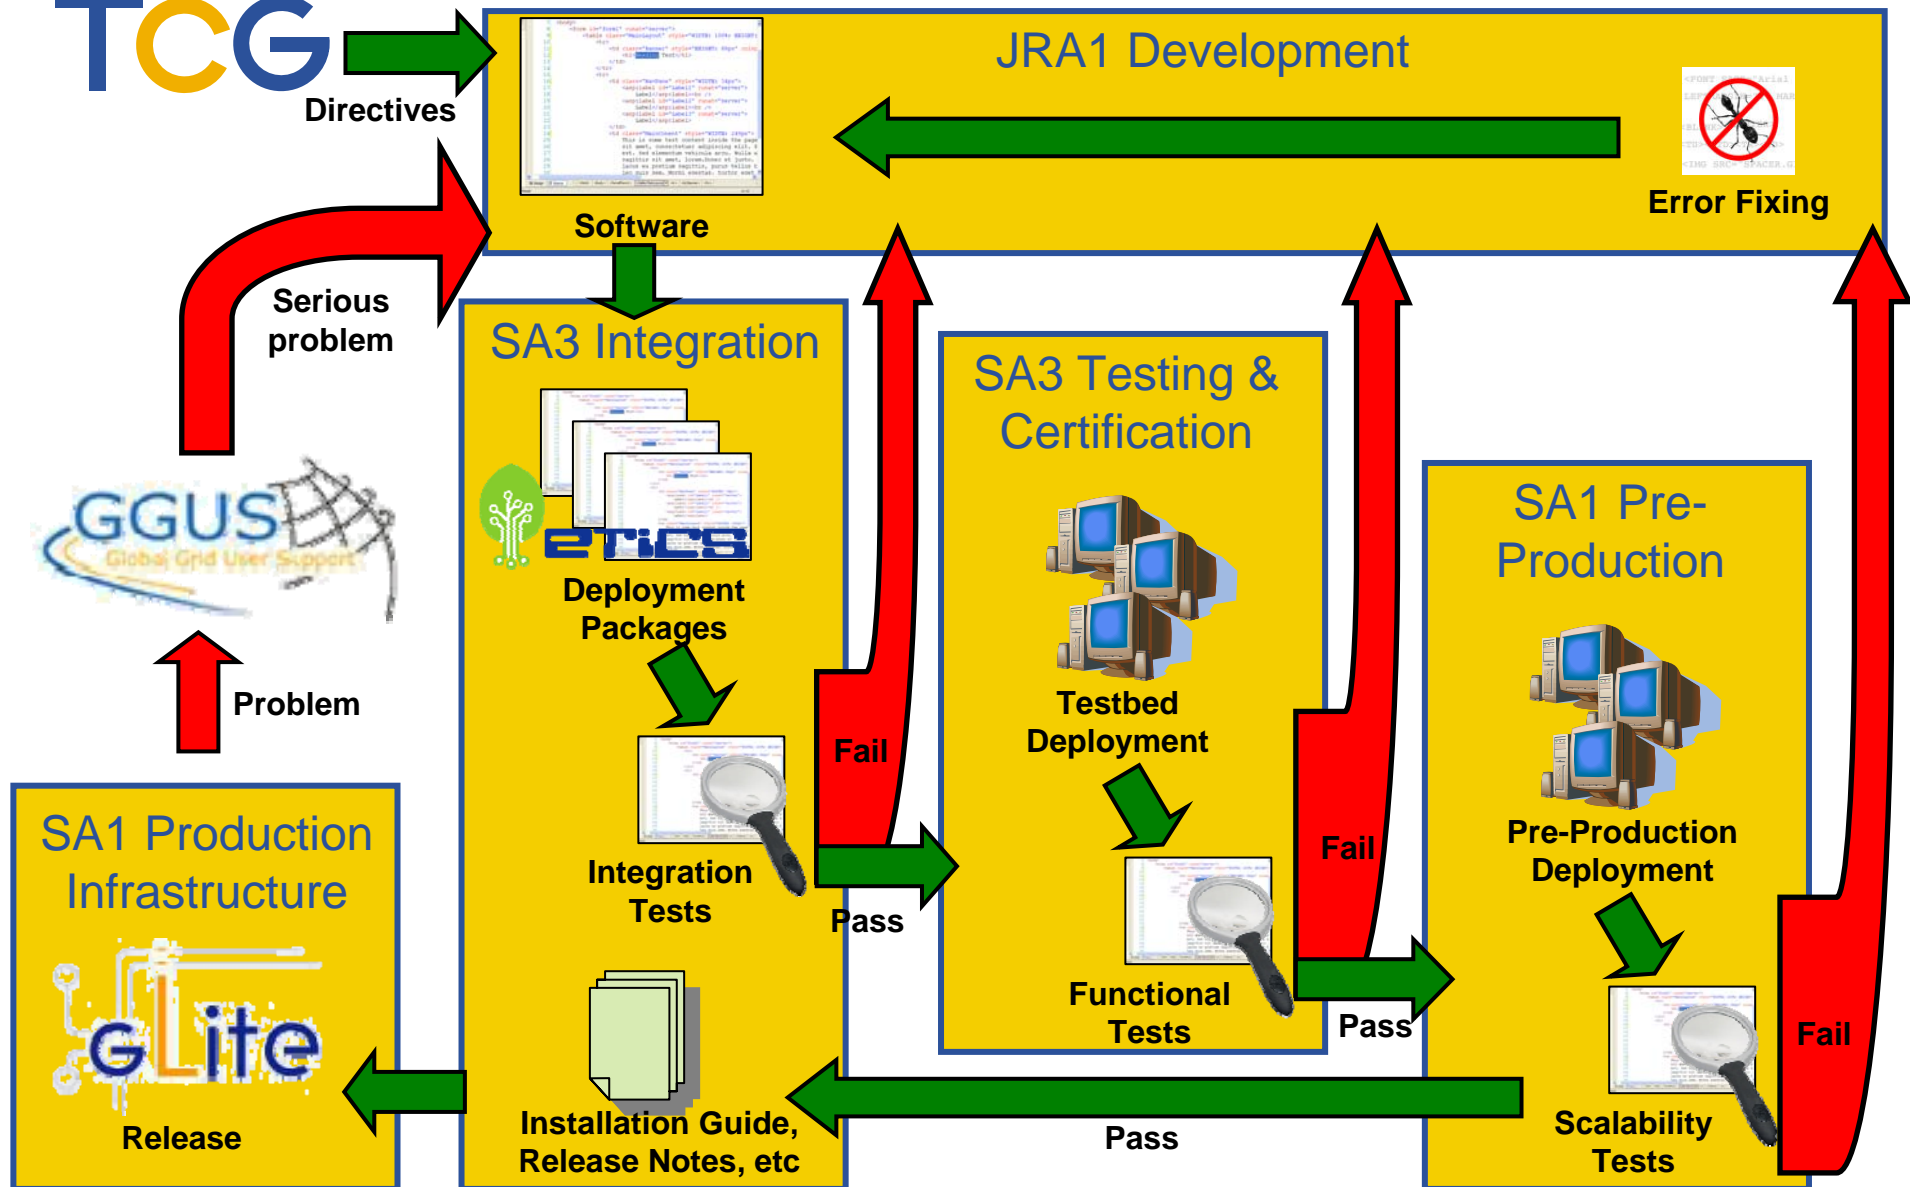


- **From EGEE to EGEE-II; JRA1 role in EGEE-II**
- **Software process; preview test-bed**
- **Milestones and deliverables**
- **Manpower situation**
- **Main achievements and future plans**
- **Reviewers recommendations at the EGEE reviews**



- JRA1 is responsible for developing the middleware
- SA3 is responsible for integration, testing and certification, i.e. to produce the release
- SA1 runs the PPS and PS systems
- ETICS provides the tools for building and testing used by JRA1 and SA3

- **Coordination of work comes from different sources**
 - Was only EMT
 - Now are: new EMT, JRA1 steering, TCG, *LCG-MB?*, *GDB?*, ...
- **New collocation of EMT:**
 - Now in SA3, responsible for the release
 - Joint coordination by SA3, SA1 and JRA1
 - Wider membership (including the Condor team 😊)
- **JRA1 steering group created:**
 - For internal JRA1 coordination
 - All cluster leaders and deputies
- **New software process**



- **Technical Coordination Group (TCG)**
 - gathers & prioritizes user requirements from HEP, Biomed, (industry), sites
 - gLite development is client-driven!
- **Software from EGEE-JRA1 and other projects**
 - JRA1 preview test-bed (currently being set up)
 - early exposure to users of “uncertified” components
- **SA3 Integration Team**
 - Ensures components are deployable and work
 - Deployment Modules implemented high-level gLite node types
 - (WMS, CE, R-GMA Server, VOMS Server, FTS, etc)
 - Build system now spun off into the ETICS project (Jan 2006)
- **SA3 Certification Team**
 - Merge of the JRA1 testing and SA1 certification teams
 - Dedicated test-bed; test release candidates and patches
 - Develop test suites
- **SA1 Pre-Production System**
 - Scale tests by users

- **JRA1 internal:**
 - Steering group (meetings every week):
 - project-eu-egEE-jra1-steering@cern.ch
 - All members (JRA1 All Hands meeting 3-4 times per year):
 - project-eu-egEE-jra1-allmembers@cern.ch
 - Internal cluster mailing lists:
 - Security: project-eu-egEE-jra3-internal@cern.ch
 - Resource access, WMS, L&B, JP: egEE-jra1-itcz@infn.it
 - Information: jra1-uk@physics.gla.ac.uk
 - Data management: it-dep-gm-dm@cern.ch
- **Cross activity:**
 - EMT, with SA3 and SA1 (meetings twice per week):
 - project-eu-egEE-middleware-emt@cern.ch
 - Design Team (meetings almost every month):
 - project-eu-egEE-middleware-design@cern.ch
 - Coordination with users inside the TCG (meetings every second week):
 - project-eu-egEE-tcg@cern.ch
 - gLite discussion list:
 - glite-discuss@cern.ch
 - Task forces

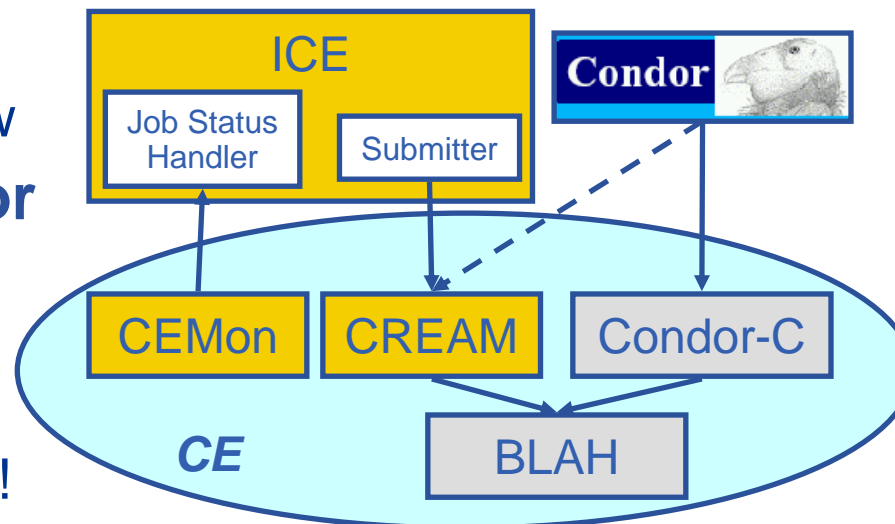
- The SA3 integration and certification teams are focused on providing code for the production infrastructure
 - Strong control over what is accepted, but **slow process for the certification of the new components and of the improvements**
- JRA1 requested a test-bed to expose to users those components not yet considered for certification
 - To get feedback from users and site managers
 - TCG and PEB acknowledged that this is needed, but no resources were foreseen for this activity in the EGEE-II proposal

The JRA1 partners which have also strong commitments in SA1 have been requested to provide resources (machines and manpower) for this activity without compromising their commitment in SA1

→ At present, only INFN and CESNET have committed resources

We need more sites!!!

- **DGAS in certification soon**
 - But will be also on the preview
- **There is now big interest for the CREAM CE:**
 - Test CREAM alone and together with the gLiteCE:
Unique CE with two interfaces!



- **Job Provenance has been ready for a long time**
 - Not only new functionalities, it can also unload the L&B!
- **Need to test glexec running on the WNs**
 - Verify how clients deal with sites with and without it!
- **Last but not least: G-PBOX!**
 - Intense activity in the Job Priorities working group

- **MJRA1.1.1: *Support plan, definition of common components and tools, strategy for multiple platform support***
 - Due PM1 (30/4/06)
 - Delivered on 10/5/06, approved by PEB (13/6/06)
- **MJRA1.2: *Functional Description of Grid Components***
 - due PM3 (30/6/06)
 - Delivered on 7/7/06, pending internal review
- **MJRA1.3: *Grid Components Reengineering Workplan***
 - Work plan for reengineered Grid Foundation and Grid Services
 - Due PM 4 (31/7/06)
 - Delivered on 3/8/06, pending internal review

- **MJRA1.4: *Shibboleth interoperability through dedicated SICS***
 - This includes development and operation of a testbed operated by the SWITCH partner.
 - Due PM 6 (30/9/06)
- **MJRA1.5: *Shibboleth interoperability with attribute retrieval through VOMS***
 - This includes development and operation of a testbed operated by the SWITCH partner.
 - Due PM 9 (31/12/06)
- **DJRA1.1: *Report on Middleware Service Reengineering***
 - Report on progress of reengineering, services delivered to SA3, compliance with TCG requirements, standardisation and cooperation results
 - Due PM10 (31/1/07)

- **For each middleware component we tried to provide:**
 - A description of the services provided
 - Which form (library vs. server etc)
 - How many instances
 - Which other services they communicate with
 - What protocol do they support
 - Who is supposed to be using them
- **Should help in:**
 - identification of the correct component to use for a specific task
 - definition of the number of instances to be deployed for a service
- **...but also:**
 - evaluation of the impact of modifications to components
 - identification of bottlenecks in the architecture

- **Program agreed with TCG**
 - Requirement collection from NA4 (reuse of LCG requirement list)
 - Prioritization of requirements
 - JRA1 and SA3 work plans approved
 - Recently added JSPG and sites requests
- **Partners responsibilities defined**
- **Migration to VDT 1.3.11, support for SL4 and 64-bit and migration to the ETICS build system not explicitly mentioned in the tables**
- **Tried to formalize the idea of “improved usability (e.g. error reporting) and performance” but probably not well represented in the tabular form...**
- **Time scales are subject to modification according to the amount of effort requested for the main JRA1 tasks**
 - Support on the production infrastructure
 - Support to certification and testing (e.g. “CMS WMS” exercise)

JRA1 Manpower situation

Partner	From TA	Actual	Comments
CERN	3	3	CERN DM group spread over many activities
CESNET	5	4.2	Hiring before autumn
CCLRC	7	3	4 people resigned. Hiring still in progress
DATAMAT	5	5.2	
INFN	23	20.5	Proposed reduction of 2.5 due to higher manpower costs. 10 days 'glitch' in June
SWITCH	3	3	
UH.HIP	2	2	John White is working as JRA1 deputy
FOM	2	2	
UvA	2	2	
UiB	1	1	
KTH	1	1	Long term sick leave

Problem in Security: 2 out of 8 people (excluding SWITCH) are not doing development (1 missing in KTH and John White working as JRA1 deputy) → 25% effort missing!!!

Contacts with UH.HIP to get an additional FTE

- **Support on the production infrastructure (GGUS, 2nd line support)**
- **Bug-fixing**
- **Improve robustness and usability (efficiency, error reporting, ...)**
- **Support for SL(C)4 and for x86-64 and IA64**
- **Addressing requests for functionality improvements from users, site administrators, etc... (through the TCG)**
- **Task Forces together with applications and site experts**

- **Delivery of gLite 3.0**
 - Convergence of LCG 2.7.0 and gLite 1.5
 - New WMS/LB/UI/CE
 - Tuning and optimization
 - “CMS-WMS” exercise
 - To speed up bug-fixing and certification
 - Instance of the WMS attached to the PS where patches flow directly from the developers but in a “controlled” way (by SA3 and SA1)
 - Migration to the new version of Condor (6.7.19)
 - Tests by IT-PSS (former LCG-EIS) team (ATLAS & CMS)
- **Porting to VDT 1.3.11 (including GT4 pre-WS)**
 - 89% of the code builds
 - Mandatory step to support Scientific Linux 4 and 64-bit
- **Preview testbed**
 - ICE-CREAM and Job Provenance deployed, G-PBox ready for deployment, glexec on WN’s will follow

- **Security**
 - Enabling glexec on WNs
 - Address user requirements in VOMS, VOMSAdmin
 - Proxy renewal library repackaged without WMS dependencies
 - Design of the Shibboleth-based short-lived credential service
- **Job Management**
 - Improvement in functionality and performance on WMS and LB
 - Development continued on new components
 - ICE-CREAM, G-PBox, including LCAS/LCMAPS plugins, Job Provenance
 - Deployment of DGAS accounting on INFN sites
- **Data Management**
 - Support for SRM v2 in DPM, GFAL and FTS
 - Improvements in Encrypted Data Storage
 - Improvements in LFC distributed service
 - FTS proxy renewal
- **Information**
 - Coding of new R-GMA design

- **Complete migration to VDT 1.3.11 and support for SL(C)4 and 64-bit**
- **Work will continue according to the plans agreed with the TCG and reported in MJRA1.3**
- **In particular:**
 - Continue work on making all services VOMS-aware
 - Improve error reporting and logging of services
 - Improve performances, in particular WMS and LB
 - Support for all LRMSs present on PS in BLAH/gLiteCE
 - Complete support to SRM v2
 - Complete the new Encrypted Data Storage based on GFAL/LFC
 - Complete and test glexec on WNs
 - Activities in the Job Priorities WG (still some confusion there...)
- **Collaboration with EUChinaGrid on IPv6 compliance**

- **Recommendations form the 2nd and the 3rd reviews**
- **Four areas:**
 - Process and releases (2nd review)
 - International collaborations and standards (2nd and 3rd reviews)
 - Industry involvement (2nd and 3rd reviews)
 - Data Management (3rd review)

26^{2-rev}. *“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”*

27^{2-rev}. *“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”*

28^{2-rev}. *“Clarify and advertise a more conservative (in term of time span) and comprehensive release cycle plan for gLite”*

29^{2-rev}. *“Revise the gLite development process to fully integrate the Technical Coordination Group and application developers”*

→ Fulfilled thanks to the new process that includes the TCG and acknowledged by the reviewers

30^{2-rev}. *“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”*

31^{2-rev}. *“Identify in the middleware stack which parts of gLite is “conformant” to standards activities within GGF and where it is currently not“*

→ **The reviewers acknowledged the work done, but:**

21^{3-rev}. *“The EGEE grid infrastructure should continue to evolve, with a balance of application versus technology and/or standards driven evolution”*

Further comment in the text

“[...] the users' understanding of their future requirements should not be taken as incontrovertible; [...] the EGEE engineers should realistically believe that they have a better understanding of new technologies and their effect on application requirements. Thus, there is a need to balance the user requests with both standards and technology driven futures.”

- **Activities in the EGEE Design Team are progressing in parallel to those of the TCG**
- **JRA1 continues to participate to the standardization bodies (e.g. OGF)**
- **The work plans include related activities**
 - e.g. JSDL2JDL translator, UR-compliant CE log file for accounting (together with OSG)
- **The TCG should acknowledge the reviewers comments and accept technical decisions taken by JRA1/SA1/SA3**
 - Not so easy to achieve...

32^{2-rev}. “Make more effective use of the Industry Forum to realize industrial involvement in the development to achieve smoother technology transfer.”

→ More work to be done:

20^{3-rev}. “Fully complete the implementation of recommendation 32 of the second project review”

→ Addressing 20^{3-rev} recommendation

- **Creation of the Industry Task Force**
- **Workshop with HP on *gLite* readiness for industry**
- **Collaboration with the CERN Openlab**

22^{3-rev}. *“EGEE should expand data management capabilities, such as:*

- Tools for the management and curation of data and*
- Tools to monitor storage utilization on a per VO basis.”*

And in the text

- On file transfer: “eventually EGEE will want to support the initiation of very large data transfers by multi-node (e.g., MPI) jobs using parallel data streams along distinct paths [...] the use of global parallel file systems such as IBM GPFS”*
- On Catalog services: “EGEE needs to be open about competing technologies and decide on its best course of action. [...] Well developed integration into excellent Hierarchical Storage Management (HSM) systems should be considered an important requirement.”*
- On Storage Elements: “the Biomed oriented elements with inherent encryption [...] could be another area where investment by the EGEE could help expand its impact both within scientific and more general use.”*

- **The JRA1 Data Management team has 3 FTEs!**
- **Many activities are happening in the National Grid Projects (e.g. GPFS SE with SRM 2.2 interface)**
- **Storage Elements, in particular HSMs, are site choices. EGEE and JRA1 are providing directives about the interface (SRM v2) and tools to interact with SEs through that interface but not the SE implementation**
- **The EDS work is progressing with priority in the DM group**
- **Storage accounting is a new area. JRA1 already started investigating it. The activity is still low priority but may raise according to TCG decisions**

- **JRA1 now working according to the new process**
 - Still a few things to be tuned, but we are on track!
- **Release of gLite 3.0 has been an important milestone**
 - Not painless. Still tuning a few components to make the usable
- **We need a preview test-bed**
 - Initial set up done, but we need more sites
- **Work plans defined and agreed with TCG**
 - This is an iterative work though...
- **Manpower situation under control**
 - But problems with some partners. Still need time before stability
- **No major issues coming from the EU reviewers**
 - Need to understand how to address the DM requests with 3 FTEs