

The GLite logo consists of the word 'GLite' in a blue, sans-serif font. The 'L' is significantly larger and has a yellow vertical bar on its right side. To the left of the 'L' is a stylized circuit diagram with a vertical line and a horizontal line at the top. Below 'GLite' is the text 'Lightweight Middleware for Grid Computing' in a smaller, blue font.

# GLite middleware status and plans

*Claudio Grandi (INFN and CERN)*

*John White (HIP)*

- **Introduction**

-  **GLite Process**

The GLite logo consists of a stylized 'G' with a circuit-like pattern of lines and dots to its left. Below the 'G' is the text 'GLite' in a bold, sans-serif font, followed by 'Process' in a smaller font. Underneath 'GLite' is the tagline 'Lightweight Middleware for Grid Computing' in a very small font.

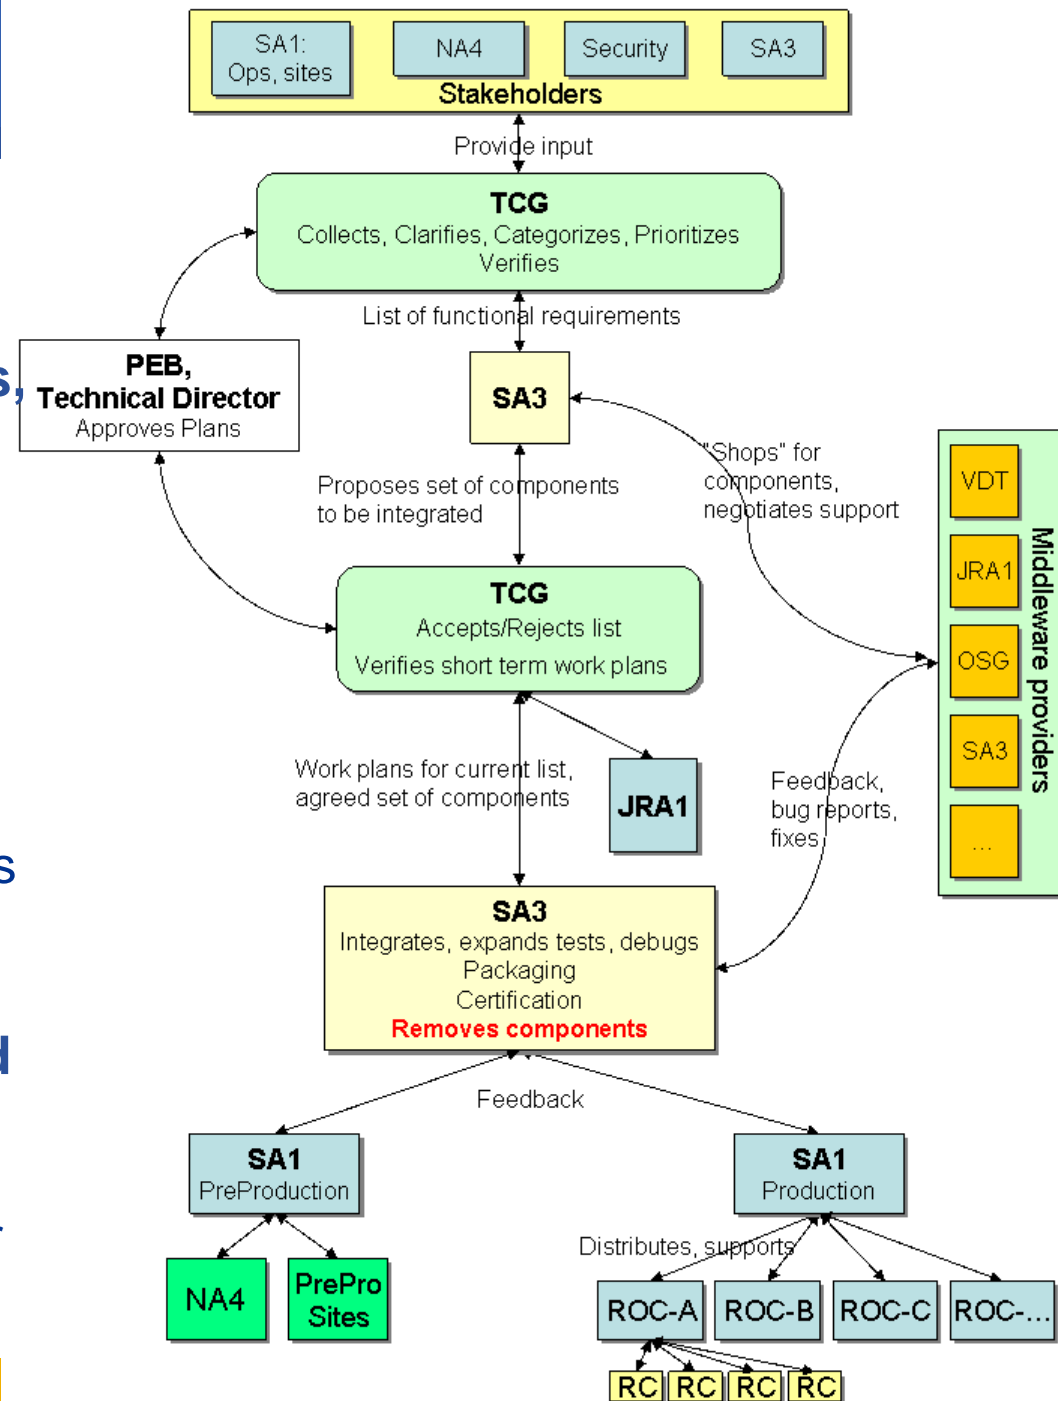
-  **GLite Status and plans**

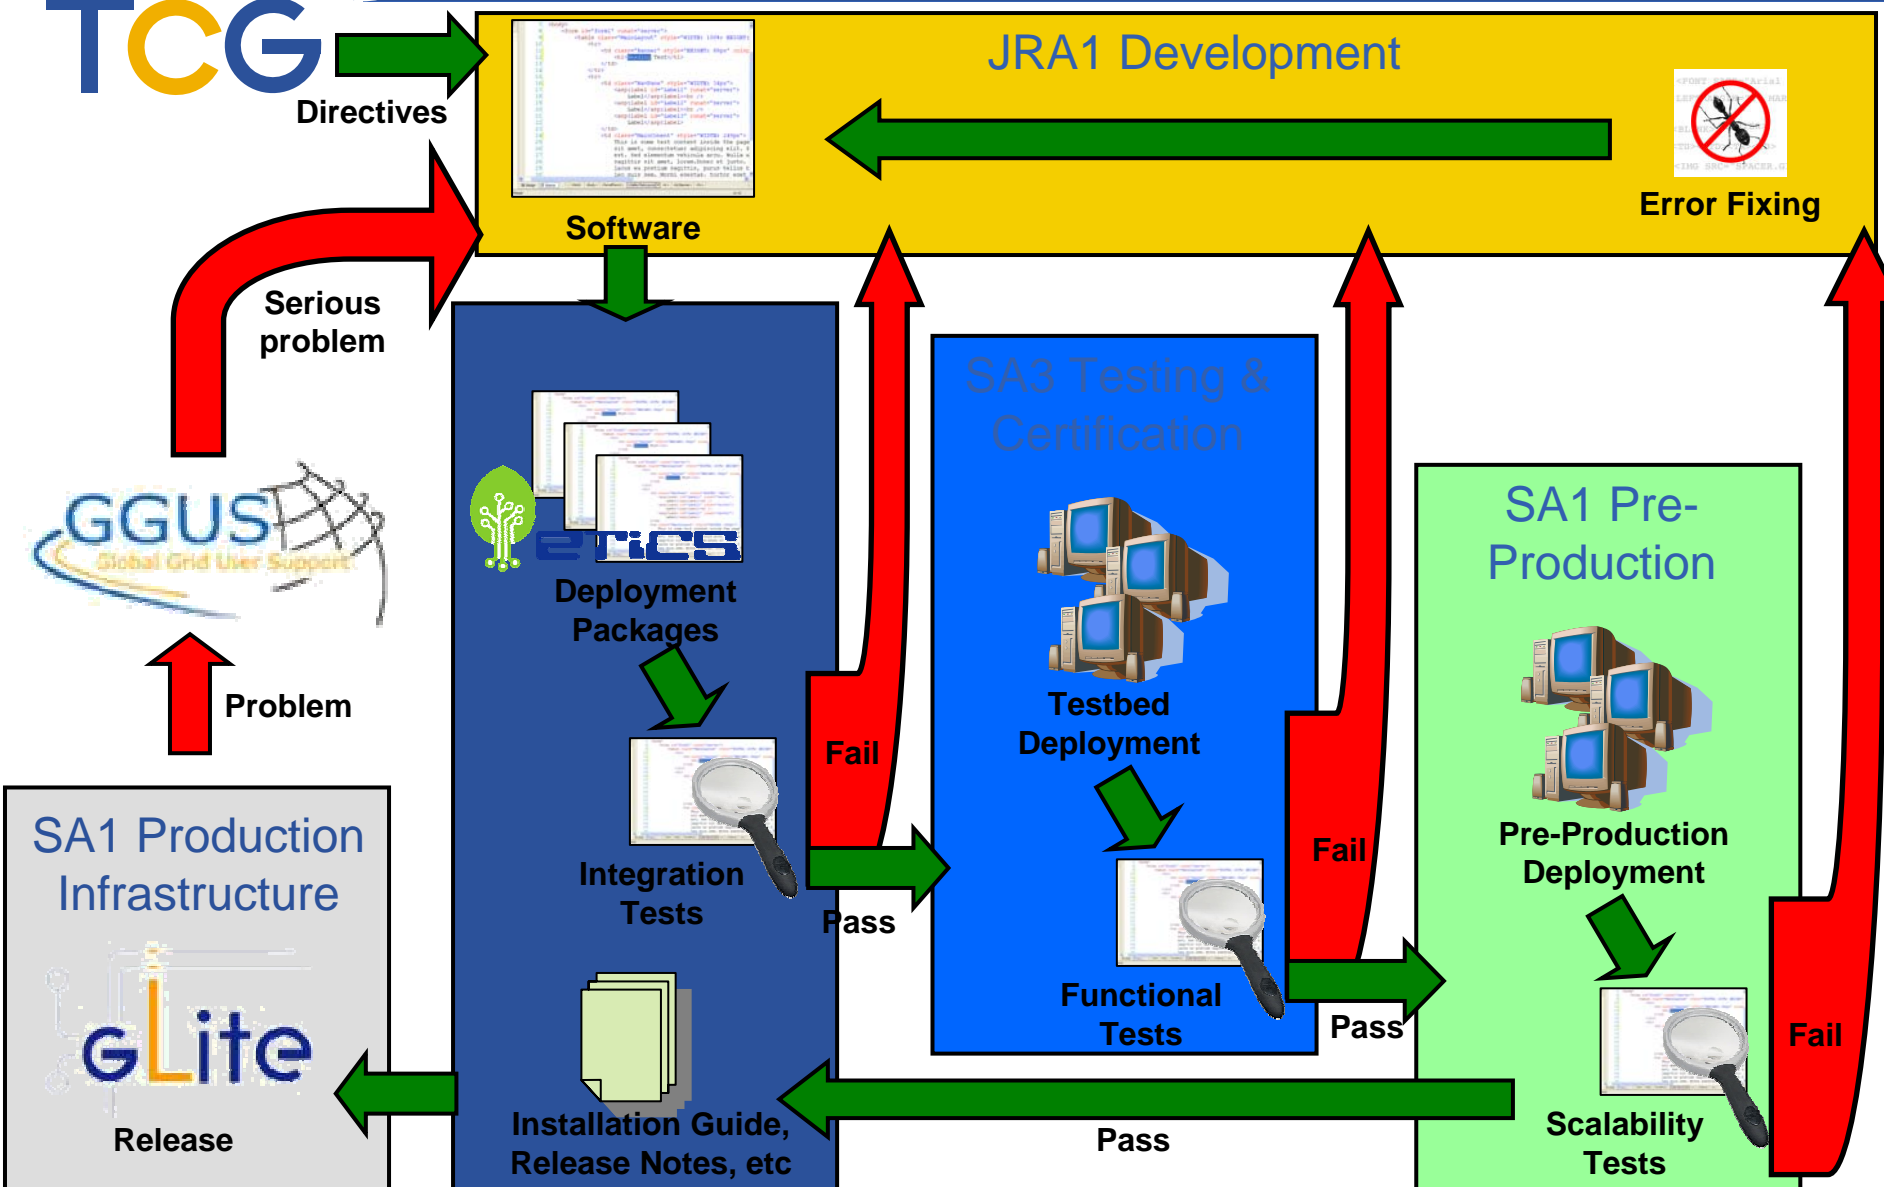
The GLite logo is identical to the one in the previous block, featuring a stylized 'G' with circuit-like patterns and the text 'GLite Status and plans'.

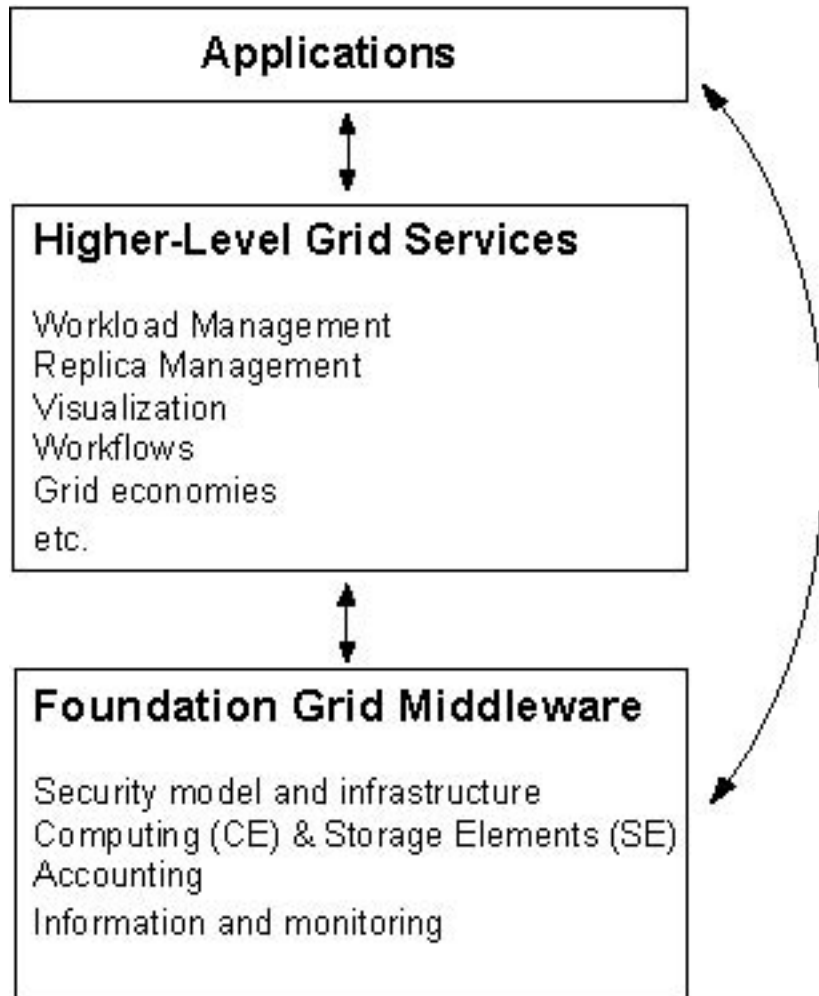
- **Summary**

- This presentation describes the components in gLite 3.0 and the components in development by EGEE JRA1
- The definition of gLite 3.0 and of the priorities for the further development has been controlled by the EGEE Technical Coordination Group (TCG)
  - Includes representatives of the LHC experiments, of the EGEE Biomed community, of the SA1, SA3, NA4, JRA1 activities, of security and of EGEE sites
- The TCG developed a requirement list used as a base for the work of JRA1, SA3 and SA1:  
<https://uimon.cern.ch/twiki/bin/view/EGEE/PriorityList>
- In this presentation the requirements references are given with respect to the list used by LCG:  
<https://uimon.cern.ch/twiki/bin/view/LCG/SummaryOpenIssuesTF> (that is a part of the TCG list)

- **Process controlled by the Technical Coordination Group**
- **Task Forces** with developers, applications, testers and deployment experts
- **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)**
  - Next major release foreseen for the autumn







- Applications have access both to Higher-level Grid Services and to Foundation Grid Middleware
- Higher-Level Grid Services are supposed to help the users building their computing infrastructure but should not be mandatory
- Foundation Grid Middleware will be deployed on the EGEE infrastructure
  - Must be complete and robust
  - Should allow interoperation with other major grid infrastructures
  - Should not assume the use of Higher-Level Grid Services
- In line with BS group results

- In the following the current status and the plans for each component are given
- For all components the main priorities for the developers are:
  - support on the production infrastructure (GGUS, 2<sup>nd</sup> 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

- **VOMS**: provides a way to add attributes to a certificate proxy. Enables VO policies. In gLite 3.0 (1.a)
- **LCAS/LCMAPS**: maps grid to local users. In gLite 3.0
- **VOMS-aware proxy renewal**: used by WMS and available as separate library in one of the next updates of gLite 3.0 (1.d)
- **glexec**: change user identity on a service
  - available on the CE. A service may be made available on the WNs if accepted by sites (support for pilot jobs, 5.k,l)
- **Delegation**: delegate some subset of user privileges to another entity. In gLite 3.0
- **JobRepository**: Collects the user, job and user-mapping information from jobs handled by LCMAPS
  - Not in gLite 3.0, work for later and with OSG



- **BDII**: based on an LDAP, by LCG, is used as Information system in gLite 3.0
- **R-GMA**: provides a uniform method to access and publish distributed information and monitoring data
  - used for job and infrastructure monitoring in gLite 3.0 (6.c,d)
  - adding authorization
- **Service Discovery**: provides a standard set of methods for locating Grid services
  - R-GMA, BDII and XML files backends; will add local caches
  - Used by some DM and WMS components in gLite 3.0
- **CEMon**: Web service to publish status of a computing resource to clients
  - Supports synchronous queries and asynchronous notification (6.d)
  - Uses the same information (GIP) used by BDII
  - In gLite 3 CEMon will be available to the users but the baseline is that the WMS queries the bdII

- **LCG-CE: based on GT2 GRAM**
  - To be replaced when other CEs prove reliability
- **gLite-CE: based on GSI enabled Condor-C**
  - supported by Condor. More efficient. Uses BLAH (see below)
  - deployed for the first time on the PS in gLite 3.0
- **CREAM: new lightweight web service CE**
  - Not in gLite 3 release. Will need exposure to users on dedicated system.
  - WSDL interface (5.j)
  - Will support bulk submission of jobs from WMS and optimization of input/output file transfer (5.d). Uses BLAH
- **BLAH: interfaces the CE and the local batch system**
  - May handle arbitrary information passing from CE to LRMS
    - patches to support this and logging for accounting being added now
  - Used by gLite-CE and CREAM

- **APEL: Uses R-GMA to propagate and display job accounting information for infrastructure monitoring**
  - Reads LRMS log files provided by LCG-CE and BLAH
  - Preparing an update for gLite 3.0 to use the files form BLAH
- **DGAS: Collects, stores and transfers accounting data. Compliant with privacy requirements**
  - Reads LRMS log files provided by LCG-CE and BLAH.
  - Stores information in a site database (HLR) and optionally in a central HLR. Access granted to user, site and VO administrators
  - Not yet certified in gLite 3.0. Deployment plan:
    - certify and activate local sensors and site HLR in parallel with APEL
    - replace APEL sensors with DGAS (DGAS2APEL)
    - certify and activate central HLR. Test scalability to the PS scale
  - (1.b, 6.a,b)

- **Storage Element**

- Common interface: SRMv1, migrating to SRMv2 (3.a)
- Various implementation from LCG and other external projects
  - disk-based: DPM, dCache; tape-based: Castor, dCache
- Support for ACLs in DPM, in future in Castor and dCache (3.c)
  - After the summer: synchronization of ACLs between SEs (4.4.d)
- Common rfiio library for Castor and DPM by end of May

- **Posix-like file access:**

- Grid File Access Layer (GFAL) by LCG
  - Support for ACL in the SRM layer (currently in DPM only)
  - Support for SRMv2 being added now (May). In the summer add thread safety and interface to the information system (3.b,e)
- gLite I/O
  - Support for ACLs from the file catalog and interfaced to Hydra for data encryption
  - Not certified in gLite 3.0. To be dismissed when all functionalities are in GFAL.

- **File Catalogs**
  - LFC from LCG (4.3.a)
    - In June: interface to POOL (4.3.c)
    - In the summer: LFC replication and backup
  - Fireman
    - Not certified in gLite 3.0. To be dismissed when all functionalities are in LFC
- **Hydra: stores keys for data encryption**
  - Being interfaced to GFAL (finish by July)
  - Currently only one instance, but in future there will be 3 instances: at least 2 need to be available for decryption
  - Not yet certified in gLite 3.0. Certification will start soon
- **AMGA Metadata Catalog: generic metadata catalogue**
  - Joint JRA1-NA4 (ARDA) development. Used mainly by Biomed
  - Not yet certified in gLite 3.0. Certification will start soon

- **FTS: Reliable, scalable and customizable file transfer (4.1.a,b,c,d,h,i)**
  - Manages transfers through channels
    - mono-directional network pipes between two sites
  - Web service interface
  - Automatic discovery of services (addresses 4.2.e)
  - Support for different user and administrative roles (1.b)
  - Adding support for pre-staging and new proxy renewal schema (4.1.f). In the medium term add support for SRMv2 (3.b,4.1.g), delegation, VOMS-aware proxy renewal (1.d)
  - Status at <https://twiki.cern.ch/twiki/bin/view/EGEE/DMFtsWorkPlan>

- **WMS helps the user accessing computing resources**
  - resource brokering, management of job input/output, ... (5.c)
- **LCG-RB: GT2 + Condor-G**
  - To be replaced when the gLite WMS proves reliability
- **gLite WMS: Web service (WMPProxy) + Condor-G**
  - Management of complex workflows (DAGs) and compound jobs
    - bulk submission and shared input sandboxes (5.d)
    - support for input files on different servers (scattered sandboxes)
  - Support for *shallow resubmission* of jobs
  - Job File Perusal: file peeking during job execution (5.i)
  - Supports collection of information from CEMon, BDII, R-GMA and from DLI and StorageIndex data management interfaces
  - Support for parallel jobs (MPI) when the home dir is not shared
  - Deployed for the first time on the PS with gLite 3.0

- Short term updates: migration to GLUE 1.2 (see job priorities), bug fixing (shallow resubmission of compound jobs) (2)
- Medium term updates: new Condor (increase efficiency), fixes to increase performance (5.b), job prologue/epilogue, short jobs (SDJ) , improve interactive access to jobs (e.g. *top*, *ls*), UI round-robin (5.i)
- Long term updates: bulk match making, high availability RB (5.a)
- **Logging and Bookkeeping service**
  - Tracks jobs during their lifetime (in terms of events)
  - LBProxy for fast access
  - L&B API and CLI to query jobs
  - Support for “CE reputability ranking“: maintains recent statistics of job failures at CE’s and feeds back to WMS to aid planning
- **Job Provenance: long term job information (6.d)**
  - If deployed will also help unloading the L&B
  - Not yet certified in gLite 3.0. Certification will start when requested by the TCG



- **GPBOX: Interface to define, store and propagate fine-grained VO policies (1.b, 5.f,h)**
  - based on VOMS groups and roles
  - enforcement of policies at sites: sites may accept/reject policies
  - Not yet certified in gLite 3.0. Certification will start when requested by the TCG.
- **Current plans: test job prioritization without GPBOX:**
  1. Mapping of VOMS groups to batch system shares (via GIDs?)
  2. Two queues (long/short) for ATLAS & CMS
  3. Publish info on the share in the CE GLUE 1.2 schema (VOView)
    - The gLite WMS is being modified to support GLUE 1.2
    - Testing with GPBOX if the “Service Class” is published
  4. WMS match-making depending on submitter VOMS certificate
    - But no ranking of resources based on priority offered yet
  5. Settings are not dynamic (via e-mail or CE updates)
- **If GPBOX is needed for LHC, tests must start now!**
  - 12 months are needed to bring it to production quality

- **New gLite components in production**
  - address requirements in terms of functionality and scalability
  - components deployed for the first time need extensive testing
- **New gLite components ready for evaluation**
- **Almost all Critical and High priority items in the Open Issues Twiki page are addressed or being addressed**
- **Need to uniform the “requirement lists”: too many!**
  - unique list together with non-HEP applications
  - uniform requirements and make them consistent
  - propose to use the TCG list