JRA1 Middleware

Frédéric Hemmer – JRA1 Manager

All Activities Meeting
March 10-11, 2005, CERN, Switzerland

Outline

- EU Review Q&A
- EU Review Recommendations

- Plans for the next 15 months
- Deliverables/Milestones

- gLite Release 1.0 status

Part I: EU Review follow-up

EU Review Q&A

EU Review Recommendations

EU Review Q&A (I)

- **Workflow support is seen as being important for future expansion of support for new applications.**
  - Possibility could be to work with GridLab to port their Triana s/w on top of gLite. Also the future link with myGrid could help with a work-flow facility. Similarly, OMII have commissioned work on a BPEL service that should be available in 2006.
  - Also have the possibility of myGrid with TAVERNA (UK funding to work with EGEE has been accepted). NA4 is also involved in EMBRACE which would also like to use TAVERNA.
  - ***JRA1 does not consider this as part of its mandate. This point was actually raised as part of NA4 Q&A***

EU Review Q&A (II)

- **How will JRA1 achieve the quality levels specified in its quality plan?**
  - We also have the opportunity to update these metrics as part of DJRA2.2.
  - **See next three slides**

How will JRA1 achieve the quality levels specified in its quality plan?

- **JRA1 Quality Plan and Development Process mandates a number of tests, checks and verifications**
  - Developers' code reviews: done by the development cluster
  - Coding guidelines checks: implemented in the build system
  - Unit testing: implemented by developers, executed by the build system (although this can be improved)
  - Coverage testing: implemented in the build system
  - Integration and system testing: the build system enforces strict controls on internal and external dependencies and a common configuration model facilitate installation and configuration
  - Functional testing: done by the test teams
  - External packages: we make any possible effort to use only mainstream, widely used versions of these packages to minimize the chances of hitting undiscovered problems

➔ The development process includes all verifications recommended by the ISO 10007 SCM guidelines and many other sources

How will JRA1 achieve the quality levels specified in its quality plan?

- **Software Quality Metric 1: Defects/SLOC < 10% (Injection Rate)**
  - This is a standard software quality metric used in the evaluation of pre- and post-release software quality. It takes into account:
    - Requirements, architectural and design defects
    - Defects found and fixed by the developers during development (errors), unit testing and internal code reviews (these bugs are not reported in the bug tracking system)
    - Defects found by integrators/testers during tests
    - Defects found by users after the release
  - It's based on statistics[1] and indicates how many bugs a normal project should expect to inject in the code during its development lifetime

- **Software Quality Metric 2: Defects/KSLOC ~ 10 (Post-Release Defect Density)**
  - This is a standard software quality metric used in the evaluation of post-release software quality. It takes into account:
    - Defects found by users after the release
    - Undetected defects
  - It's based on statistics[2] and indicates how good is a software for users

- **Statistically, between 75% - 95% of the latent defects are removed during individual reviews and by the use automated integration processes[1] (Rayleigh model), the rest are removed by direct testing. In projects of more than 100 KSLOCs typically about 30% to 50% of the remaining defects are found during the first year**
- **The gLite development effort is not a typical project and present challenges not found by other project (distributed development, multiple languages, etc), so it is realistic to expect at least one order of magnitude more issues**

[1] SEI statistics, QSM Rayleigh modeling, various defect removal models – Ferguson, Rico (see for example http://software.isixsigma.com/library/content/c030910a.asp)

[2] A study conducted by Reasoning (http://www.reasoning.com) over 200 open source projects using automated code inspection techniques found an average defects density of 0.57 defects/KSLOC on released software. The real defect density is about one order of magnitude greater (includes functional and user perception defects)

How will JRA1 achieve the quality levels specified in its quality plan?

- **The theory applied to gLite**
  - 630000 SLOCs * 10% = 63000 defects
  - Worst case: 63000 * 25% * 30% = 4725 defects to be found in year 1
    Best case: 63000 * 5% * 50% = 1575 defects to be found in year 1

    (the rest are found later or never found – not all code is actually exercised in all possible conditions)
  - However about 50% of the code was developed and used before EGEE, so we should only consider half of the SLOCs and about half of the defects ➔ 2350
  - So far 983 defects have been found in about 6 months of testing

    ➔ up to 1350 defects can be expected to be found in the next 6 months

  - This means that our post-release defect density (DD) will be:

    2.5-17.5 defects/KSLOC

  - As a comparison: Shuttle Navigation Systems → 0.1 defects/KSLOCs
    WinXP → 2.6 defects/KSLOCs
    MySQL 5 → 0.9 defects/KSLOCs

    Average CMM[1] 2 → 3.2 defects/KSLOCs

  - In order to improve the quality and reduce the defect density we have to increase either the process efficiency (more code reviews and unit testing) and/or increase the test efficiency (more functional and regression testing)
  - The main problem is that we have no previous data to understand how efficient really is the development process

[1] CMM: Capability Maturity Model

EU Review Q&A (III)

- **Why do we want to replace LCG-2 with gLite – You should not only tell us the differences in features from LCG-2 but also the reasons for the change**
  - Slides were added by Frédéric for an Atlas presentation, addressing these points.

Job Management Services Improvements to LCG-2

- **Essentially addressing shortcomings of LCG-2**
  - Robustness
  - Efficiency
- **Read-only Information System cache (ISM)**
  - Reduce reliance on the Information System
  - Updated by
    - Active polling of resources (CE in push mode)
    - Notification of available resources (CE in pull mode)
    - Combination of both
- **Task Queue**
  - Holding job requests when matchmaking is not possible
  - Persistent queue
- **Condor-C**
  - Reliable job submission between the WM and the CE
- **Support for DAGs**
- **Interfaces to catalogs**
  - RLS, StorageIndex, DLI
- **Bulk job submission**
  - Currently through DAGs without dependencies

Data Management Services Improvements to LCG-2

- **Addressing shortcomings of LCG-2 data management**
  - RLS poor performance
  - A non distributed catalog
  - Lack of consistent grid-storage interfaces
  - Unreliable data transfer layer

- **Fireman Catalog**
  - Hierarchical Name Space
  - Bulk Operations
  - ACLs
  - Web Services Interface
  - Performance/scalability

- **Metadata Catalog**
  - Consistent interfaces agreed by Applications
- **gLite I/O**
  - Support of ACL's
  - Support of Fireman catalog in addition to RLS
- **File Transfer Service**
  - Did not exist on LCG-2

Information Systems Improvements to LCG-2

- Authentication
- Service Discovery
- Registry Replication

EU Review Q&A (III)

- **Clarify what are the economic models that are to be supported by the accounting service.**
  - Better explain DGAS from a user and application point of view.
  - **The aim is to support Economic Models that use information about the load of a Computing Element to compute a fair price for its usage.**
  - **DGAS implements a Price Authority service that keeps the records of the prices assigned to resources. The prices can be assigned manually by site managers, or via an 'economic model' that does this process automatically.**
  - **Actually DGAS uses a plugin structure to implement economic models, that is, different economic models can be implemented and simply plugged in the Price Authority.**
  - **So an 'economic model' is the algorithm used by the Price Authority to compute the price for a Grid Resource.**
  - **Actually a couple of Economic Models algorithms are provided, that compute resource prices according to the load of the resources. Other algorithms can be designed and implemented and easily plugged in the structure, by us, or third parties using furnished APIs.**

EU Recommendations

- **Make the migration from LCG-2 to gLite 1.0 more application focused**
  - **So we should work with applications directly**
    - **Not clear how we organize this**
  - **But need the Pre Production Service with *significant* resources**
- **Current gLite should be a stepping stone towards more robust standard-based Grid infrastructure**
  - Collaborations with other International grid middleware R&D initiatives are strongly desirable to select additional components for integration and deployment
  - **Will continue involvement in GGF**
  - **Will continue collaboration with NGS, NorduGrid, OMII, OSG, …**
  - **But clearly labour intensive and stretching resources**
- **Presuming that the decision is made that a co-scheduling facility is essential**
  - A timetable for its implementation should be announced
  - The implementation should not be considered until the end of the contract
  - **Co-scheduling requires the ability to claim (reserve) multiple classes of resource.**
  - **Existing frameworks for (advance) reservations have been explored since the beginning of the EU DataGrid project, and have so far stopped short for lack of reservable resoucres. For computing resources, initial tests were made with the MAUI scheduler (that is not currently integrated deployed in the infrastructure). Advance reservation support for storage resources (via the recent SRM v2 interface) is being prototyped by the IT/CZ cluster, with the purpose of evaluating the interaction with existing implementations of the matchmaking engine.**
  - **Too early for publishing a timetable**

CRM Workshop Rome, 17/18 February, 2005

- **First attempt to define common standard interface to compute resources (inspired by SRM effort)**
- **Attendees from Condor, EGEE, Globus, INFN, LCG, NAREGI, nextgrid, NorduGrid, unigrids**
- **Main Outcome:**
  - Evaluation of JSDL as common language for job submission is ongoing in all projects
  - A Computing Element white paper will be written based on an early draft by Ian Foster and Miron Livny
  - Follow up meeting during EGEE3 conference in Athens

Future Work

- **Testing and software packaging. Reinforce these even further**
  - **JRA1 heartily endorses this recommendation!**

Part II: Plans for the next Period

Plans for the next 15 months

Deliverables/Milestones

Plan for next period

- **Incremental releases (bug fixes) are foreseen on a regular basis for the whole period**
- **Interim releases (functionality enhancements) are foreseen on a regular basis for the whole period**
  - Based on a work plan being elaborated now
    - Based on feed back from applications and operations
    - Will seek priorities from SA1 and NA4
    - E.g. second version of FTS being tested now
- **PM10-PM12: Integration & Testing of Release 1 whose functionality has been frozen in December'04. A common configuration mechanism will be proposed for the services to adopt by PM15.**
- **PM14: Revision of the architecture document taking into account feedback from applications and operations, pending issues such as the site proxy and network element. Progress of standards such as JSDL, WS-Addressing, WS-Notification,… will be addressed**
- **PM15: Revision of the Design document addressing changes in service decomposition, interfaces and review of Web Services technologies.**

Plan for next period

- PM18: Test plan for Release 2
- PM19: Improvements of services based on Applications and Operations feedback. Functionality of Release 2 will be frozen.
- PM20: Integration of Release 2
- PM21: Release2 delivered to SA1
- PM24: Final report including assessment of work completed and outstanding issues

Milestones & Deliverables

Milestone	Month		Description
DJRA1.3	M12	03-2005	Software and associated documentation (Release 1)
DJRA1.4	M14	05-2005	Architecture and Planning (Release 2)
DJRA1.5	M15	06-2005	Design of grid services (Release 2)
LCG 1.5.2.17	M15	06-2005	Interim Release for LCG (1.3 ?)
DJRA1.5	M15	06-2005	Design of grid services (Release 2)
MJRA1.6	M18	09-2005	Test plan for core Grid components and overall Integration (Rel 2)
MJRA1.7	M19	10-2005	Software for the second release candidate available
MJRA1.8	M20	11-2005	Release Candidate 2 enters testing and validation period (Rel 2)
DJRA1.6	M21	12-2006	Software and associated documentation (Release 2)
DJRA1.7	M24	03-2006	Final Report

Part II: gLite Middleware Status

Status

Modules	Availability for ARDA	Availability for SA1 February 7, 2005 (1.0.6) February 21, 2005 (1.0.8) March 9, 2005 (1.0.9)
Complete WMS with Task Queue, Pull mode and Interface to Data Management	February 2, 2005 February 2, but problems. Fixed on February 11. CE only in pull mode.	February 18, 2005 WMS problems in 1.0.6 Functional and basic stress tests available.
Single File Catalog (FiReMan)	February 2, 2005. February 2, 2005. Weekly upgrades provided.	February 7, 2005 Functional tests available.
gLite I/O	Available.	Available. Functional tests and basic stress tests available. 19 bugs reported, 4 open.
File Transfer/Placement Service	February 2, 2005. February 2. Frozen version. New version being finalized for Service Challenges.	February 7, 2005 Testing has stopped.
User Interfaces	February 2, 2005 for WMS and RGMA. Lcg_utils needs to be ported. February 17 (WMS, Data Management, R-GMA, VOMS) in AFS.	February 25, 2005 In 1.0.9 (March 9, 2005) No formal testing. Bugs reported through daily usage.
R-GMA	Available. Tests cases being defined and implemented by UK.	Available. Deployment problems also observed by RAL/NIKHEF.
VOMS	Available. Testing will be carried by INFN.	February 7, 2005 Problems using old server solved (1.0.8)

WMS with Data Management

- **JDL**

```
Executable = "/bin/sh";
Arguments = "job4.sh";
StdOutput = "stdout.log";
StdError = "stderr.log";
StorageIndex =  "http://lxb2028.cern.ch:8080/EGEE/glite-data-catalog-service-fr/services/SEIndex" ;
InputData = { "lfn:/tmp/hemmer/test/InputFile" };
DataAccessProtocol= "gridftp";
InputSandBox = {"job4.sh", "sqrt" };
OutputSandbox = {"stdout.log", "stderr.log"};
```

- **WMS Matching**

```
[lxplus.cern.ch] glite-job-list-match job4.jdl

Selected Virtual Organisation name (from proxy certificate extension): EGEE
Connecting to host gundam.cnaf.infn.it, port 7772

**********************************************************************
                    COMPUTING ELEMENT IDs LIST
 The following CE(s) matching your job requirements have been found:

                *CEId*
 lxb2022.cern.ch:2119/blah-lsf-jra1_high
 lxb2022.cern.ch:2119/blah-lsf-jra1_low
**********************************************************************
```

Command Line Tools

Workload Management

- glite-job-attach
- glite-job-cancel
- glite-job-get-chkpt
- glite-job-list-match
- glite-job-logging-info
- glite-job-output
- glite-job-status
- glite-job-submit
- glite-lb-dump
- glite-lb-load
- glite-lb-logevent
- glite-lb-purge

R-GMA

- rgma (interpreter)

VOMS

- voms-proxy-destroy
- voms-proxy-info
- voms-proxy-init

StorageIndex

- glite-seindex-list

Catalog

- glite-catalog-chkperm
- glite-catalog-chmod
- glite-catalog-create
- glite-catalog-getacl
- glite-catalog-getattr
- glite-catalog-getdefacl
- glite-catalog-getguid
- glite-catalog-getreplica
- glite-catalog-ls
- glite-catalog-mkdir
- glite-catalog-mv
- glite-catalog-rm
- glite-catalog-rmdir
- glite-catalog-setacl
- glite-catalog-setattr
- glite-catalog-setdefacl
- glite-catalog-setdefperm
- glite-catalog-setreplica
- glite-catalog-setschema
- glite-catalog-stat
- glite-catalog-symlink
- glite-catalog-touch

gLite I/O (Catalog & SRM)

- glite-get
- glite-put
- glite-rm

Summary

- **WMS**
  - Task Queue, Pull mode, Data management interface
    - Available in the prototype
    - Used in the testing testbed
      - *Basic functionality and stress test available*
    - Delivered to SA1
      - *Now working on the certification testbed*
      - *Submission to LCG-2 demonstrated*
- **Catalog**
  - MySQL and Oracle
    - Available in the prototype
    - Used in the testing testbed
      - *Basic test available*
    - Delivered to SA1
      - *But not used yet*
- **gLite I/O**
  - Available in the prototype
  - Used in the testing testbed
    - Basic functionality and stress test available
  - Delivered to SA1
      - *But not used yet*
- **FTS**
  - Available in the prototype
    - But should not be used anymore
  - JRA1 testing has stopped
  - FTS is being evolved with LCG
    - Milestone on March 15, 2005
- **UI**
  - Available in the prototype
    - Incudes data management
    - Delivered to SA1
    - Not yet formally tested
- **R-GMA**
  - Available in the prototype
  - Testing has shown deployment problems
- **VOMS**
  - Available in the prototype
  - No tests available
  - Delivered to SA1
    - Interoperability problems observed but now over

Open Issues (high level)

- **Condor-C Proxy renewal**
  - Promised for a long time
  - Only received on March 3, 2005 (part of Condor 6.7.6 to be released March 11, 2005)
- **File Transfer Service**
  - The version in Release 1 is *not* the one that is going to be used for service Challenges
  - FTS is being evolved as required by SA1 service Challenges
  - Weekly monitoring
  - Status at:
    http://egee-jra1-dm.web.cern.ch/egee-jra1-dm/transfer/index.htm
  - Milestone to have a stable “demo” running by March 15, 2005
- **VOMS**
  - Incompatible with previous version
- **RGMA**
  - Difficult to install and make it work
- **Service Discovery API**
  - Changes requested by developers
  - Will not be used by other services in Release 1

FTS stability

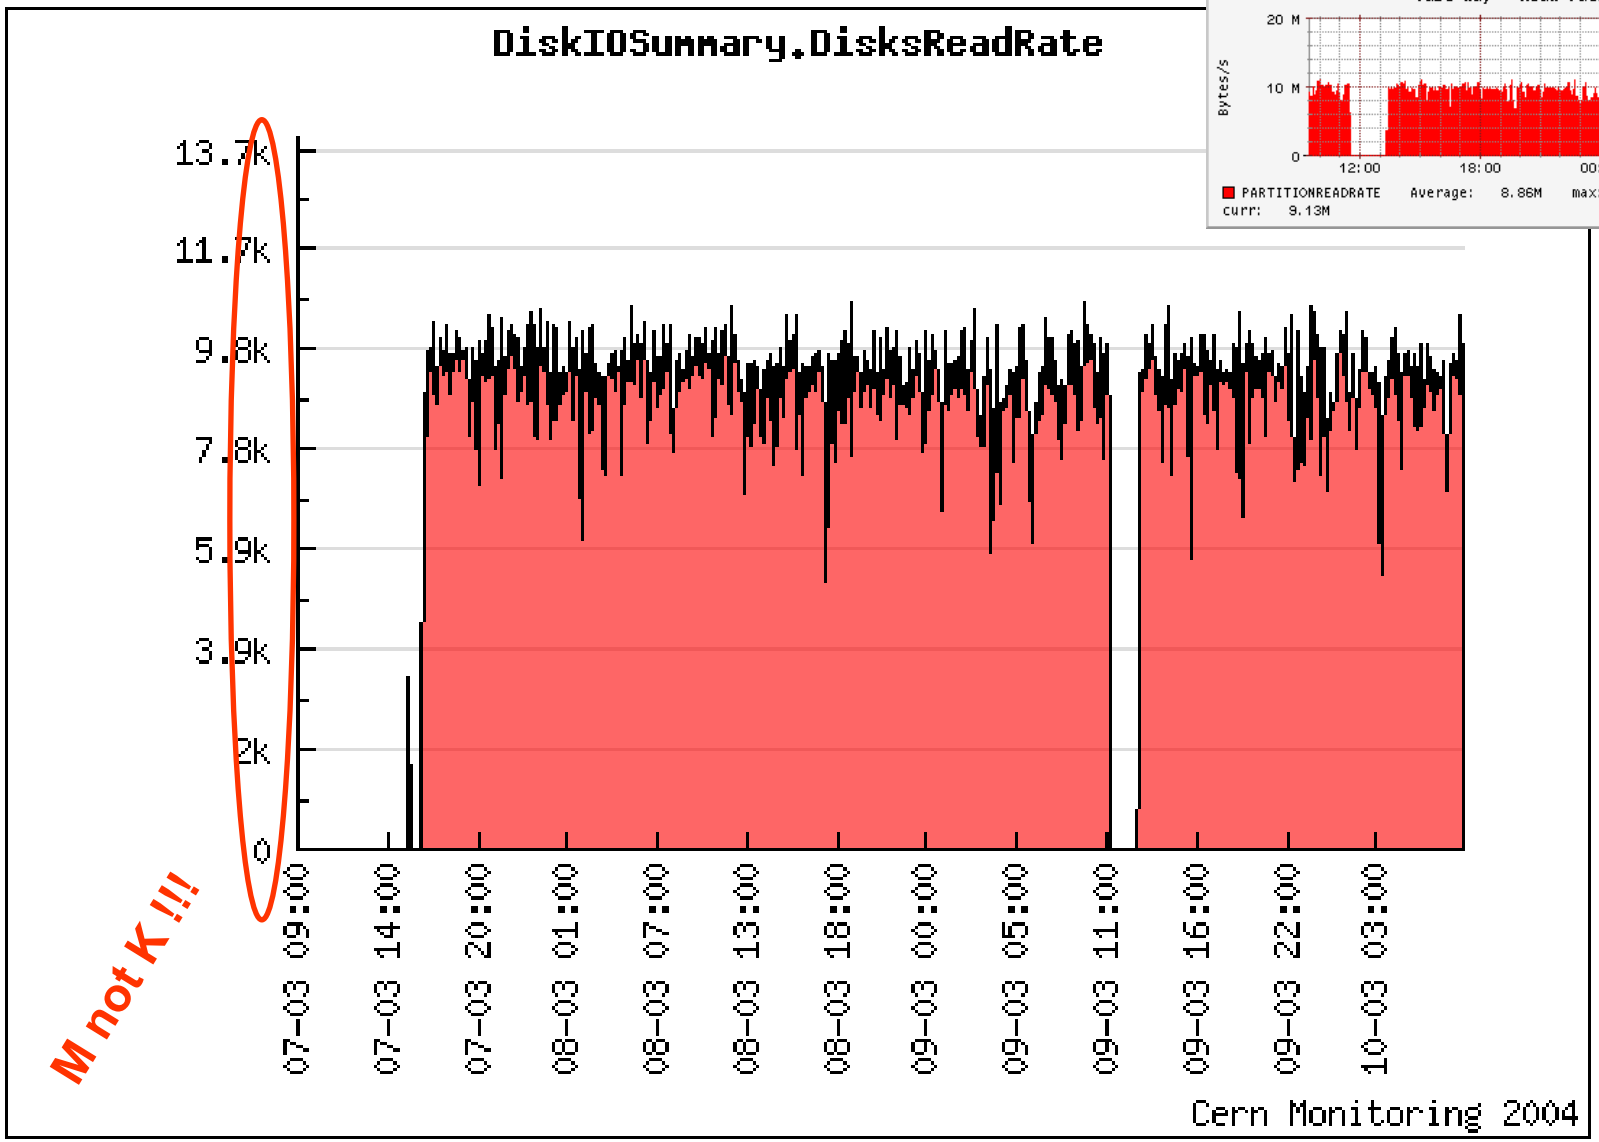

Testing reports

- **JRA1**
  - WMS: https://edms.cern.ch/document/567478
  - Catalog: https://edms.cern.ch/document/568048
  - gLite I/O: https://edms.cern.ch/document/571060
- **SA1**
  - Weekly reports: http://egee-docs.web.cern.ch/egee-docs/list.php?dir=.\Testing\Testing%20Reports\&
  - Daily reports: http://egee-docs.web.cern.ch/egee-docs/list.php?dir=.\Testing\Meeting%20minutes\&
- **NA4**
  - ARDA: http://lcg.web.cern.ch/lcg/PEB/arda/LCG_ARDA_Glite.htm
- **gLite Progress**
  - http://cern.ch/egee-jra1/gLite/Status.htm

Documentation[1]

- **Workload Management**
  - LB: https://edms.cern.ch/file/571273/1/LB-guide.pdf
  - JDL: https://edms.cern.ch/document/555796/
  - WMS: https://edms.cern.ch/file/572489/1/WMS-guide.pdf
- **Data Management**
  - Overview: https://edms.cern.ch/file/570643/1/EGEE-TECH-570643-v1.0.pdf
  - Glite I/O: https://edms.cern.ch/file/570771/1/EGEE-TECH-570771-v1.0.pdf
  - FiReMan: https://edms.cern.ch/file/570780/1/EGEE-TECH-570780-v1.0.pdf
  - Lcg-utils comparison: http://egee-jra1-dm.web.cern.ch/egee-jra1-dm/lcg-utils-glite.htm
  - Metadata (Java): https://edms.cern.ch/file/570779/1/EGEE-TECH-570779-java-v1.0.pdf
- **R-GMA**
  - User Guides (Java, C, C++, Python), API's, Service Discovery: http://hepunx.rl.ac.uk/egee/jra1-uk/glite/doc/index.html
- **VOMS**
  - User Guide: https://edms.cern.ch/file/571991/1/voms-guide.pdf
  - Administrator Guide: https://edms.cern.ch/document/572406/1
- **DGAS**
  - https://edms.cern.ch/file/571271/1/EGEE-DGAS-HLR-Guide.pdf
  - https://edms.cern.ch/file/571271/1/EGEE-DGAS-Gianduia-guide.pdf
  - https://edms.cern.ch/file/571271/1/EGEE-DGAS-PA-Guide.pdf

[1] Links may be temporary; documentation has not been reviewed yet

Other points

- **Before February 18, 2005, SA1 received weekly integration releases *at the same time* than JRA1 Testing team**
  - Addressing the wish of SA1 to see the code as early as possible
  - Caused confusion and/or duplication of efforts, as the same bugs were discovered by the two teams
  - It was decided to release weekly integration builds one week *after* the have been used by JRA1 testing
  - Naming changed
    - Example: I20050225_RC vs. I20050225
- **Weekly releases are too frequent for SA1 testing**
  - No decision yet
  - JRA1 will for now still produce a weekly release
- **Testing efforts have been merged between JRA1/SA1 and NA4**
  - High priority goal to port LCG-2 Test Suite to gLite
  - Coordinated by a daily meeting
  - But communication still needs to be improved

gLite Middleware Summary

- **All modules are available in the prototype**
- **gLite distribution has been cleaned up containing only modules as defined at the review**
- **We are in the middle of the testing**
  - There is an agreement between JRA1/SA1 and NA4 on how to proceed with testing
  - Additional efforts have been requested from CCLRC and INFN
- **Developers priority is to fix bugs**
- **Too early to say that the middleware will be ready for end users at the end of the month**
  - But LCG certification has started
- **Documentation is on the critical path**
  - For deployment
  - For end users
  - And will need to be reviewed