Enabling Grids for E-sciencE

LHCC Comprehensive Review November 22-23, 2004, CERN

EGEE Middleware

Frédéric Hemmer

LCG/EGEE Middleware Manager

www.eu-egee.org

www.cern.ch/lcg

www.glite.org

Outline

- **What has been done so far**
  - Design Team
  - Architecture & Design Documents
  - Development Testbed (a.k.a. prototype)
  - Development, Integration and Testing Processes
  - … and Software
- **Outcome of last ARDA Workshop**
  - Broaden Scope & Size of Development Testbed
  - Deploy exiting prototype code for ALICE DC'04 Phase III
- **Tackling last review recommendations**
- **Issues & concerns**
- **Next steps**

Work Achieved So Far

Design Team

- **International Design Team formed as of December 2003**
  - Composed of a small number of experienced Middleware providers from AliEn, former EDG, Globus & VDT/NMI
  - Holds roughly monthly face to face meetings
  - Defines and agrees the Middleware building blocks, focusing on existing solutions rather than developing from scratch.
  - Discusses components and how they need to be adapted to fit in the development testbed (prototype)
  - Produced Architecture and Design documents
    - Involving other people as needed

Architecture and Design

- **Project started with staffing essentially complete on April 1st, 2004**
- **Architecture document released in June 2004**
  - Explaining what the building blocks are and the proposed functionalities
  - Based upon the ARDA blueprint
  - Submitted not only to EGEE management but to diverse communities such as the GAG and the experiments for further feedback.
  - Feedback taken into account in delivered version as much as possible
- **Design Document released in August 2004**
  - Explaining what are the proposed external interfaces together with the related WSDL.
  - Submitted not only to EGEE management but to diverse communities such as the GAG and the experiments for further feedback.
  - Feedback taken into account in delivered version as much as possible
- **Those documents have been also used by consortiums such as OSG to prepare their blueprint**
  - And made available to GGF, GridLab, OMII, etc…

Development Testbed

- **A Development Testbed (known as prototype) has been made available as of May 2004**
  - To host prototype middleware as recommended by the ARDA RTAG
  - Many ideas from the ALICE/AliEn system
    - Started with AliEn, adding additional components from other middleware providers
  - Comprises resources at CERN, University of Wisconsin/Madison and INFN
  - Approximately 60 users registered
  - Being expanded with a second VO in Madison
  - Will be further expanded as a result of the ARDA Workshop outcome in October 2004
  - Used by the ARDA Team to try out new middleware
    - And provide feedback on what needs to be changed, corrected, etc…
    - Will use the prototype middleware to deliver 4 end to end prototypes of Analysis
    - ARDA Workshops in June & October 2004 served as (positive) feedback mechanism
  - Bio-Medical community has been invited to use the development testbed
    - But not very active so far

Development, Integration, Testing processes

- **Design discussed, agreed and implemented**
  - Developers introduce their components in prototype
  - Interfaces (WDSL usually) published and made available
    - According to Design Document when possible
  - Components modified to interoperate
    - Tricky language dependent issues
    - Immaturity of code generators (gSoap, A
  - Code checked-in in gLite CVS
- **Prototype updated for direct use by ARDA T**
  - Happened 7 times since May
  - Bug reports and enhancement request received
  - Developers required to provide installation notes
- **Code taken over by Integration for auto-build**
  - Feedback cycle with developers
  - Early difficulties with SCM compliance seems now to be over
  - Usually difficulties with dependencies
    - E.g. X uses classadds 0.9.6, Y uses 0.6.x
    - E.g. X uses GT2.4 from VDT, Y uses GT2.2 from globus
  - Streamline configuration
  - Contributes to documentation
- **Code taken over by Testing**
  - Initial Step is to install Services
    - Usually (especially at the beginning) not suitable docs (installation & tion)
    - Significant time spent iterating with developers
  - Contributes to documentation
  - Build and run a basic suite of test for a given component
- **Code delivered to Operations for certification and validation**
- **Feedback cycle starts**

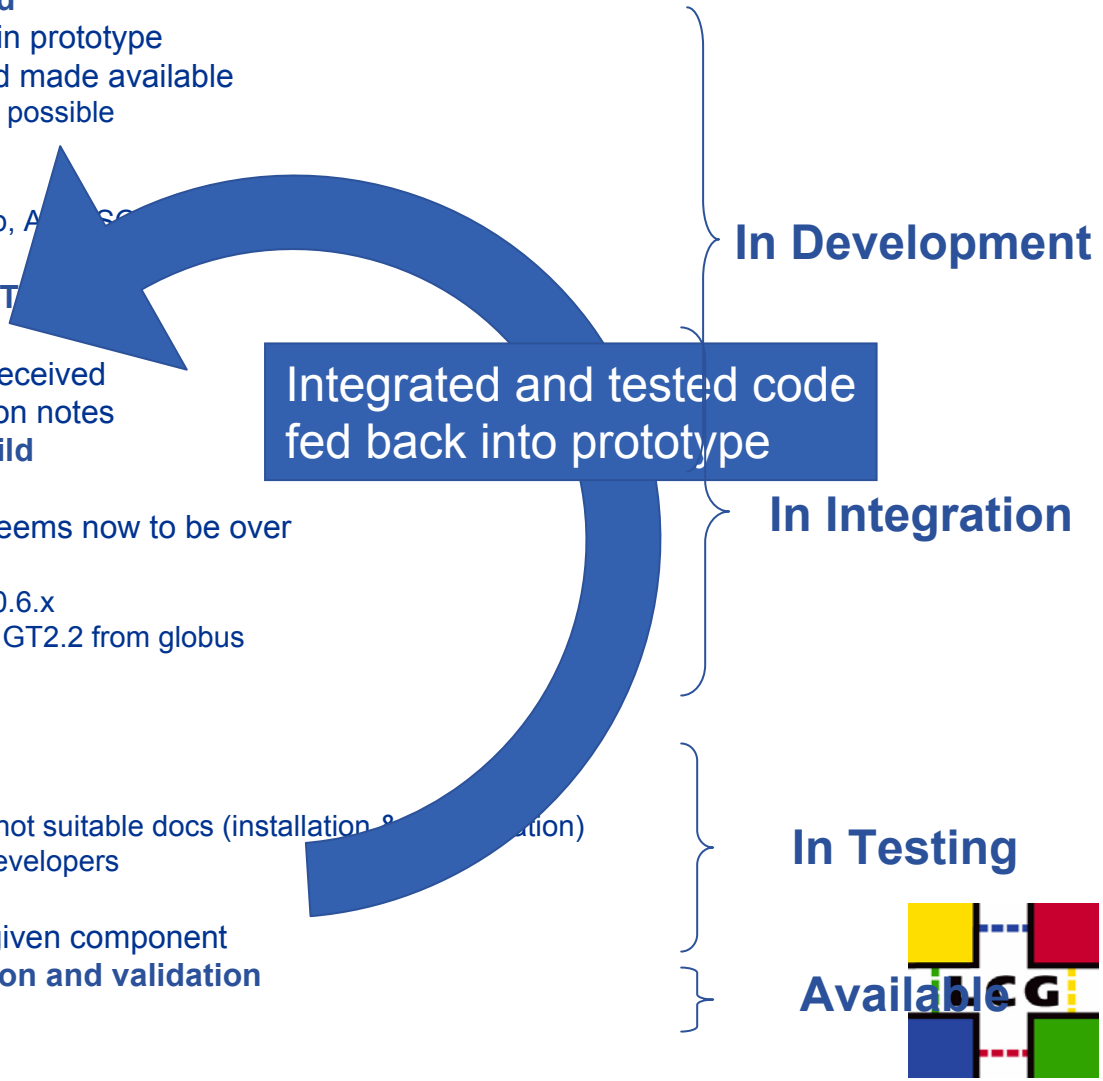

Example (complex) : File Catalog

- **Initially pure AliEn File & Metadata catalog in prototype**
- **Interfacing to RLS discussed, agreed**
  - Perl bindings generated
  - AliEn File catalog interfaced
- **Local Replica Catalog introduced in the Architecture**
  - Discussed many times, finally agreed
  - AliEn refactored to remove Master Replica and introduce SE-ID
  - ServiceIndex Interface exposed from AliEn File Catalog
  - Local Replica Catalog introduced in the prototype
  - Perl bindings generated
  - Waiting to interface to AliEn File Catalog in the prototype
- **At the same time development of biomed Metadata Catalog**
  - Used as an example on how to plug your own catalog
  - Available in the prototype
- **At the same time AliEn refactored to comply to SCM**
  - Now available from the gLite CVS tree
  - And in the build system
- **Not in testing yet**

Integration and Testing

- **Strong Focus of the project**
  - Based on the fact that Middleware so far was insufficiently tested and documented, poor quality, badly configured, released as big bangs etc…
- **Software Configuration and Management plan has been agreed and elaborated**
  - https://edms.cern.ch/document/446241
- **Developers Guide available**
  - https://edms.cern.ch/document/468700
- **Test Plan available**
  - Based upon architecture document, release plan, and Applications requirements
  - https://edms.cern.ch/document/473264/
- **Automated Build system in place**
  - Nightly builds (identify problems as soon as possible) http://glite.web.cern.ch/glite/packages/N20041119/
  - Continuous builds (available for developers to be able to integrate)
  - Weekly Integration builds & package generation (continuous updates available) http://glite.web.cern.ch/glite/packages/I20041029/
- **Testing cluster in place in 3 sites**
  - CERN, NIKHEF and RAL
- **Common logging, configuration, management and error handling being defined with developer groups**
  - Complicated given the different languages http://glite.web.cern.ch/glite/project/pm.asp#sloc

Prototype Middleware Status & Plans (I)

- **Workload Management**
  - AliEn TaskQueue
  - EDG WMS (plus new TaskQueue and Information Supermarket)
  - EDG L&B
- **Computing Element**
  - Globus Gatekeeper + LCAS/LCMAPS
    - Dynamic accounts (from Globus)
  - CondorC
  - Interfaces to LSF/PBS (blahp)
  - "Pull components"
    - AliEn CE
    - gLite CEmon (being configured)

Blue: deployed on development testbed

Red: proposed

Prototype Middleware Status & Plans (II)

- **Storage Element**
  - Existing SRM implementations
    - dCache, Castor, …
    - FNAL & LCG DPM
  - gLite-I/O (re-factored AliEn-I/O)
- **Catalogs**
  - AliEn FileCatalog – global catalog
  - gLite Replica Catalog – local catalog
  - Catalog update (messaging)
  - FiReMan Interface
  - RLS (globus)
- **Data Scheduling**
  - File Transfer Service (Stork+GridFTP)
  - File Placement Service
  - Data Scheduler
- **Metadata Catalog**
  - Simple interface defined (AliEn+BioMed)
- **Information & Monitoring**
  - R-GMA web service version; multi-VO support

Prototype Middleware Status & Plans (III)

- **Security**
  - VOMS as Attribute Authority and VO mgmt
  - myProxy as proxy store
  - GSI security and VOMS attributes as enforcement
    - fine-grained authorization (e.g. ACLs)
    - globus to provide a set-uid service on CE
- **Accounting**
  - EDG DGAS (not used yet)
- **User Interface**
  - AliEn shell
  - CLIs and APIs
  - GAS
    - Catalogs
    - Integrate remaining services
- **Package manager**
  - Prototype based on AliEn backend
  - evolve to final architecture agreed with ARDA team

Schedule for pre-production service

- gLite I/O – Available
- Logging & Bookkeeping, WMS, CE, WN – In testing – end Novemb
- R-GMA – In integration/testing – December
- CE-Notification – In integrati ember
- Replica, File, Combined Catalog – In d ent – December
- File Transfer Service – gration/testing – December
- File Placement Service – ntegration/testing – December
- VOMS – In integration/testing – December
- UI – In integration – December
- AliEn Task Queue & C – In integration/testing – To be deployed on ALICE sites
- Package Man – Discussions w/experiments, deployment – prototype exists
- Grid A – Prototype exists
- Acco ng (DGAS) – In integration – Prototype exists
- Job Provenance – Proof of concept exists

Details to be discussed later this week at the 2nd EGEE conference

LCG Internal Review 2003 Recommendations

www.eu-egee.org

www.cern.ch/lcg

www.glite.org

LCG Internal Review 2003 Recommendations (I)

- **Evident: the M/W is one of the most important components of LCG.**
- **But: the current plan for the LCG M/W area (partly thanks to ARDA) is incomplete.**
  - Also: GTA milestones have been referred to as “mostly reports” (not good).

    Middleware milestones have been defined as development & testing infrastructure, prototype software in a development testbed, released software and associated documentation
- **While the M/W is not under the exclusive control of the LCG project, its milestones are very important and need to be included in the project overview.**
  - They will clearly need to be negotiated between LCG and EGEE

    LCG Milestones are aligned with EGEE ones, within the scope of the ARDA project. Milestones and Deliverables are reviewed by both EGEE and LCG projects. EGEE Deliverable reviews reports are available:

    DJRA1.1 (Architecture): http://edms.cern.ch/document/493614/1

    DJRA1.2 (Design): http://edms.cern.ch/document/487871/0.8
    - Having the same person in charge of both is clearly good

LCG/EGEE Milestones

WBS	EGEE	Comment	Name	Date MS	MS Done
1.5		___	Middleware		
1.3.4			LCG - EGEE Coordination		
1.3.4.1			EGEE senior management appointed	15-07-03	15-07-03
1.3.4.2			Technical design team established	01-09-03	04-12-03
1.3.4.8		i_jan04	EGEE Middleware people hired	29-02-04	23-03-04
1.3.4.9		i_jan04	EGEE Middleware execution plan available	29-02-04	17-03-04
1.3.4.10		i_jan04	EGEE Contract signed	01-04-04	01-04-04
1.5.2.6	None	i_apr04	First version of prototype available for experiments	16-06-04	16-05-04
1.5.2.7	MJRA1.1	i_apr04	Development and integration tools deployed	30-06-04	30-06-04
1.5.2.8	MJRA1.2	i_apr04	Software cluster development & testing infrastructure in place	30-06-04	30-06-04
1.5.2.9	DJRA1.1	i_apr04	Architecture & Planning Document for release candidate 1	30-06-04	30-06-04
1.5.2.10	MJRA1.3	i_apr04	Integration & Testing infrastructure in place; test plan	31-08-04	31-08-04
1.5.2.11	DJRA1.2	i_apr04	Grid Services design document for release candidate 1	31-08-04	31-08-04
1.5.2.12	None	i_apr04	Second version of prototype available for experiments	31-08-04	31-08-04
1.5.2.13	MJRA1.4	i_apr04	Release Candidate 1	31-12-04	

LCG Internal Review 2003 Recommendations (II)

- **It would be useful to have a plan, complete with milestones, manpower (including manpower outside CERN) prepared already by January 2004 (i.e. before EGEE starts in April).**
  - This plan will clearly have to be agreed upon with the EGEE Management, who should arrange for hires of new people to proceed immediately at the project start (or even before).

    A Plan has been elaborated with Responsibilities, Manpower, Milestones and Deliverables withtin the scope of the EGEE Technical Annex (Delivered to EU on January 27th, 2004). An execution plan is available identifying tasks and individuals.

    LCG Milestones have been added to take into account the ARDA project

    People have been hired as of December 2003 (but starting no earlier than April 1st, 2004). For CERN only, this represented 49 interviews and 15 staff selected

    A more detailed technical Release Plan has been elaborated assigning > 100 milestones and responsibilities. This plan is reviewed on a weekly basis.

    https://edms.cern.ch/document/468699

EGEE JRA1 Release Plan

100	PM4 DONE (PM4)	First version based on AliEn Package Manager component exposed and Web Service •*Status August 8: DONE* CVS component: http://jra1mw.cvs.cern.ch:8180/cgi-bin/jra1mw.cgi/org.glite.prototype.alien/AliEn/Service/ Module org.glite.prototype.alien (subdirectory AliEn/Service/PackMan) RPM: AliEn-Client http://alien.cern.ch/dist/1.35-20/i386_linux24/RPMS/AliEn-Client-1.35-20.i386.rpm
101	PM5 DONE (PM7)	Evaluation of alternative package manager solutions that could provide required functionality, scalability and flexibility. Selected candidate for implementation. •*Status September 1: Presentation will be given on September 2nd* •*Status September 15: Presentation given on September 2: http://agenda.cern.ch/fullAgenda.php?ida=a043837; informal meetings with experiments and SA1 ongoing* •*Status September 30: Further discussion scheduled during the ARDA workshop Oct 20-22: http://lcg.web.cern.ch/LCG/peb/arda/workshops/oct04.htm* •*Status October 27: No major objects on plan presented during ARDA workshop – further progress tracked in milestone 97 - DONE*
102	PM6	Prototype of Package Manager Service delivered. •*Status October 19: prototype deployed; documentation missing, needs to be factored out of AliEn.* •*Status October 27: installation documentation missing; factoring out from AliEn not feasible.*
103	PM9	Up to this point iterative bugfix releases and functionality improvements based on user feedback..
	PM9	Functionality freeze
	PM10	Integrated release candidate 1
	PM12	Release 1

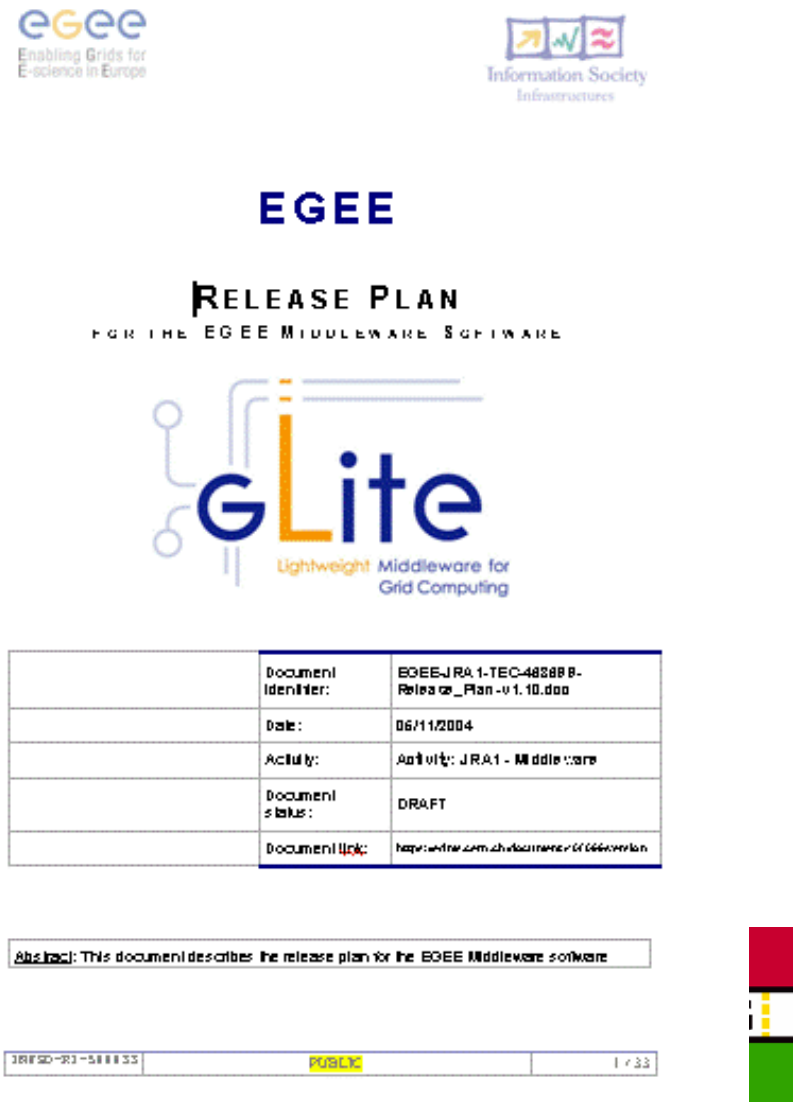

Outcome of last ARDA Workshop

- **Enlarge size and scope of the development testbed**
  - Expose more users than just the ARDA team
  - Made enough computing resources available for realistic analysis
    - Madison installation is being expanded with ~60 CPU's
    - FZK resources will be added
- **Deploy current prototype software on ALICE sites**
  - To handle Phase III of ALICE Data Challenge 2004
  - To provide early feedback to Middleware developers
  - ALICE site managers will provided requested help
    - Still analyzing technical and resources implications

LCG Internal Review 2003 Recommendations (III)

- **M/W planning**
  - ARDA planning should be established by end 2003, involving both the experiments and EGEE m/w experts, as well as AliEn, NorduGrid and US m/w experts.
  - The way to implement the ARDA service Architecture in a fast prototype framework should be proposed by a technical team as soon as possible.
    - The plan should clearly be
      - *(a) consistent with experiment requirements*
      - *(b) complying with the EGEE engagements*
      - *(c) submitted to SC2.*
  - The plan should address (obvious, here for completeness):
    - real timescale and milestones
    - team available for executing the work
    - the m/w components to be taken and their sources
    - the relationship and synchronization with LCG-2,
    - and … OGSI

LCG Internal Review 2003 Recommendations (IV)

- **The six-month timescale for the ARDA prototype should be negotiated with EGEE and the experiments:**
  - Real point: to have a new release for users before end 2004
    - this requires the experiments to be exposed to it earlier (by end summer 2004?)
      - *This timescale does not seem incompatible with EGEE proposal/TA sent to EU → see below*

    Users exposed to prototyping work as of May 2004. Release Candidate 1 expected by the end of the year. Individual components are/will be available on a continuous basis (weekly integration build).

- **M/W development cycle and scope**
  - Component-by-component deployment and avoiding big-bang releases are critical parts of a strategy for avoiding some of the worst problems experienced with EDG.
    - Intermediate (EGEE) releases should be in the milestones of the LCG m/w manager

    Intermediate releases have been provided in the form of the prototype testbed. Seven releases have been produced between May and September 2004 (bug fixes, functionality changes, new components).
  - EGEE has to produce a coherent set of functions and code with clear interfaces allowing multiple implementations of some components

    The Design Document exposes the proposed external interfaces. The EGEE PTF work on WSDL interfaces explicitly to allow multiple implementations. As an example, there are two implementations of the metadata catalog. ATLAS is envisaging to prototype those interfaces in AMI.

LCG Internal Review 2003 Recommendations (V)

- **OGSI and GRAM**
  - ARDA report recommends a fast OGSI-compliant prototype
  - Question is whether OGSI offers a mature solution (on a short timescale) for an ARDA implementation. Numerous issues:
    - Information System and GRAM (critical parts of the GLOBUS kit) used by LCG, but both services have problems of scalability and reliability.
    - In GT3/OGSI GRAM is both slow and not scalable (expected, given it is but the GT2 version wrapped). CondorG/GT3 will be (has been?) demonstrated in SC2003.
      - *Performance evaluation eagerly awaited: the CONDOR/VDT team (should have) incorporated much of the EDG/LCG feedback on the GT2 GRAM.*
    - GT3 IndexService: totally new, looks well designed, still some problems are present.
      - *GLOBUS team has been informed and an improved bi-directional channel has been established with GTA.*

  Globus January 2004 invalidated OGSI strategy. EGEE Middleware activity proposed to stick to Web Services (WS-I) as far as possible waiting for the standardisation efforts to cristallize and provide mature implementation(s).

LCG Internal Review 2003 Recommendations (VI)

- **Federated Grids**
  - Currently: LHC experiments use a number of different Grids
    - Sometimes multiple systems (Grids) are used even within a single experiment
  - Clear that different Grids will coexist (e.g. US Tier2, NorduGrid)
    - numerous reasons (funding included)
  - EGEE alone will also require some “federation” concept (national grids with identical m/w)
  - First priority should be to show that a single Grid can achieve real production quality.
    - Fortunately, this is the LCG

- **Side remark: ARDA may offer a good opportunity for harmonisation of the different efforts**

LCG Internal Review 2003 Recommendations (VII)

- **Realistic Grid analysis tests by experiments before end 2004 are necessary before the Computing TDRs can be written (and their underlying computing model established).**
  - This testing should involve also Tier-2s and possibly even Tier-3s, and requires the s/w from the Application Area to be available in time.
  - These tests will need a stable LCG service
    - but with a significant risk of failure (for the tests)
    - the experiments should chart strategies for failures in specific components – leaving the rest of the system in a "functioning" state

- **Fallback solutions**
  - A fallback solution for Grid m/w is very important
    - especially if LCG-2 evolution does not deliver production-quality m/w in time for the experiment C-TDRs
    - The experiments need (at least) some Grid functionality, available with production quality.
      - *The actual functionalities and how such a fallback solution will be organized should be included in the LCG plan*

Enhanced LCG-2 Middleware is the fallback solution.

Issues and Concerns

Concerns/Issues

- **Almost all components are late compared to original schedule**
- **Difficult to coordinate the middleware developments that are happening in national programmes, EGEE Middleware, LCG Operations, US, etc…**
- **Biomed security & functional requirements are much more stringent than HEP requirements (encryption, anonymity, DICOM servers, metadata handling, …)**
  - Security requirements may complicate considerably overall architecture
- **Building prototypes from diverse building blocks**
  - With requirements of production quality software process
  - Different Security Models
- **Aggressive and diverse timelines may affect quality**
  - Software Engineering and in particular Testing & Integration
- **Multiple reporting lines, coordination meetings, etc…**
  - Generating overheads
  - Many requests for dissemination affecting effectiveness of delivering technical results
- **Convergence with standards and timescales**
  - WSRF not yet approved, LCG needs to take decision in 2005
- **What level of involvement/collaboration with OSG is expected?**

Next steps

Next Steps

- **Deliver scheduled components to pre-production service**
  - Get Operations feedback and adapt as required
- **Deploy prototype Middleware to ALICE sites**
- **Enlarge development testbed and open it to a larger community**
- **Finalize contents of EU RC1 release**
- **Deliver and test all components out of integration builds**
- **Finalize first integrated release as an EU deliverable**

Summary

- **Although EGEE Middleware development activity started in April only**
  - Prototype Middleware has been made available to the ARDA Team
    - With positive feedback
  - Software Engineering process is in place
    - Aimed at delivering production quality software
    - Allowing for rapid development cycles
  - First components are being delivered to LCG pre-production service
- **The challenge is now**
  - To finalize components delivery to LCG pre-production service
    - With the required quality
  - To deploy prototype software to ALICE sites
    - Preferably coming out of Integration and Testing
  - To enlarge the development testbed to more sites and users
    - And to cope with the associated operational workload
  - To deliver first integrated Release as an EU deliverable