



Applications report to final project conference

F Harris

(on behalf of WP10, WP9, WP8)



Overview of talk

- ◆ Overview of the applications programme of work until the review
- ◆ Application testing experience in LCG/EDG environments
- ◆ Lessons and planning coming from joint meetings with middleware
- ◆ Lessons and planning coming from joint meeting with security group
- ◆ Planning for application evaluations and the writing of deliverables
- ◆ Output from *AWG* and a forward look
- ◆ Summary



Overview of the applications programme of work until the review, and related milestones

- ◆ **From now till early November** (5-6 working weeks)
 - Evaluate middleware on EDG App TB, (+ LCG Cert TB and LCG-1 service)
- ◆ **From now till ~Nov 24** (WP reports to be finished in 8 weeks time)
 - Write, and agree within WP, the final application evaluations, and **submit to reviewers Nov 24**
- ◆ **Nov 24 to Dec 8**
 - Interaction with reviewers before **submission to PTB members Dec 8**
- ◆ **Dec 8 to Dec 15**
 - Interaction with PTB prior to **PTB approval/disapproval(!) Dec 15**
- ◆ **Now to end March**
 - Planning and development of review and demos **review last week in Feb**
 - Inputs to EGEE (? Mechanisms for this...) from **AWG** and the individual groups



Application testing experience in LCG/EDG

- ◆ Loose Cannon(LC) testing so far of EDG middleware has been largely based on LCG Cert TB (which has MDS-BDII inf sys specially configured). These test have been very positive.(there has also been since some months testing on EDG DEV TB)

- ◆ LCs have recently started tests on LCG-1 service and EDG App TB

- ◆ Experiments about to start on LCG-1 and EDG App TB

- ◆ WP9 and WP10 just about to start on EDG App TB
 - WP9 have been involved in 'data migration' exercise with WP5



Test code sources

LCs July-Sept 03 LCG / EDG testing:

- ◆ Full general purpose generic-HEP-application test suite from Jean-Jacques Blaising: (PERL)
 - Gets info on GRID current status (GRID config file) and creates JDL accordingly
 - Submits jobs, monitor status, retrieve output, report results
 - ◆ Various LC testing JDLs and scripts (/bin/bash, PERL)
 - general and intensive stress tests
 - ◆ Official EDG WP6-SC (edg-site-certification) PERL-OO test suite
 - Basic job submission, output retrieve
 - Basic Data Management tests, Filename registration and file copy
 - My Proxy tests
 - Match-Making specific tests
- (http://marianne.in2p3.fr/datagrid/TestPlan/TESTSTATUS/EDG_20_TEST_STATUS.html)**

30 Sept 2003

F Harris EDG Heidelberg



RB Stress Tests by Massive Job Submission on the LCG-CT

- ◆ RB never crashed
- ◆ ran without problems for several days in a row 20 streams with 100 jobs each (typical error rate ~ 2 % still present)
- ◆ RB stress test in a job storm of 50 streams , 20 jobs each :
 - 50% of the streams ran out of connections between UI and RB. (configuration parameter – but machine constraints)
 - Remaining 50% streams finished normal (2% error rate)
 - Time between job-submit and return of the command (acceptance by the RB) is 3.5 seconds. (independent of number of streams)



Preliminary full simulation and reconstruction tests with ALICE on the CERN LCG-CT

(beginning of Sept 03)

- ◆ Aliroot 3.09.06 (including HBT correl.) fully reconstructed events
- ◆ CPU-intensive, RAM-demanding (up to 600MB ,160MB average) ,long lasting jobs (average 14 hours)
- ◆ Outcome:
 - > 95 % successful job submission, execution and output retrieval in a lightly loaded GRID environment
 - ~ 95 % success (*first estimate*) in a highly job-populated testbed with concurrent job submission and execution (2 streams of 50 AliRoot jobs and concurrent 5 streams of 200 middle-size jobs)
 - My Proxy renewal successfully exploited



Main lessons and planning coming from joint meeting with middleware

(<http://agenda.cern.ch/fullAgenda.php?ida=a035398>)

- ◆ **Broker much more stable- but still need real scalability tests** (? on EDG TB since we have more sites)
 - Many new features (**check pointing, interactive jobs, output data, access cost ranking, accounting, ...**) largely untested by Applications. These will be tested by application groups and LCs.
 - DAGMAN will be evaluated by CMS before the review
- ◆ **New replica manager works well in general**
 - File registration is **very** slow (> 20 secs for 1 12-byte file!) but solutions in hand
 - RLS/RLI needs evaluation (this comes transparently to applications)
 - No 'bulk transfers' (collections) support available - work around that



Lessons from meetings with middleware (2)

◆ Interface to mass storage needs clarification

- Seems that LCG and EDG will be using different software, and that neither is yet fully tested (**mass storage support is essential for all production work**)
- Applications want to see uniform interface (through WP2), and a clear definition of what we can rely on in the coming few weeks (we don't want to be debugging features which have not yet been thoroughly tested by the WPs) - i.e. what modes of usage are absolutely solid (e.g. gridftp.....)

◆ RGMA 'commissioning'

- Good progress has been made but the applications have not yet seriously stressed it
- Some experiments (e.g. CMS/BOSS, LHCb-ATLAS/GANGA) interfacing to RGMA for monitoring
- Applications will naturally perform evaluations of RGMA from now with EDG 2.1. It is essential that we establish a good level of stability on EDG App. TB as soon as possible so as to enable applications to do coherent and reasonable scale tests (that is what we promised reviewers in March)



Other application points

- ◆ 'BrokerInfo' usage
 - Need clarification to users on usage of 'BrokerInfo APIs' and 'WP2 APIs'
- ◆ Requirement by all application groups for outbound IP connectivity from WNs
 - Matter taken on board by LCG GDB
- ◆ Information publishing management
 - Project needs (not in 2.1 before end of EDG) clarification of the semantics of published information so consumers and producers understand each other!
- ◆ Space Management for SEs and WNs
 - Fell off the bottom of priority list - will come back to haunt us(very soon)
- ◆ No facilities for quota management and advance reservation - will be needed 'soon' post-EDG



*Lessons and planning coming from joint meeting with security group
(<http://agenda.cern.ch/fullAgenda.php?ida=a035410>)*

- ◆ **VOMS has been integrated in 2.1** and SCG are confident that it provides a basic security infrastructure
 - WP1 and WP2 have interfaced their services to it
 - No file level ACLs
 - Please note it cannot provide 'quotas and advance reservation' schemes, but provides an infrastructure for such future facilities provided by WHO? in post-EDG developments
- ◆ **Applications view for work till end Feb (and beyond)**
 - If we have any 'serious' problems we will run in non-VOMS mode
 - WP10 are very interested in checking out 'delegation'
 - WP10 need ACLs and encryption
 - WP8 (GENIUS) need proxy delegation
- ◆ **SCG demo at last project review (to be organised by SCG)**
 - implementation of ACLs on a VOMS-aware Storage Element
 - Based on EO testcase (different datasets, different user groups with different access rights)



Planning for the application evaluations

◆ WP8

- Effort shared between tests on LCG-1 service and EDG App TB
 - LCG-1
 - Tests by both LCs and some experiments (ALICE, Atlas, CMS, LHCb...) with experiments tests preparing for data challenges (mass storage support mandatory)
 - EDG App TB
 - Tests by LCs and some experiments (DO, ALICE+?)

◆ WP9 and WP10

- All testing done on EDG App TB (? In longer term there is the issue of their access to LCG-1 after EDG project finishes)



Current thinking about organisation of testing

◆ Experiment work on LCG-1

- will test performance (but without extreme stress testing) doing necessary work preparing data challenges

◆ Testing on EDG App TB

- Will basically be for functionality of new EDG software (checkpointing, MPI, etc..)
- Also will naturally be checking out the stability of RGMA
- Following functionality tests may do some organised stress testing
- We will check areas that caused us problems in previous evaluation
 - E.g. WP9 will check registration of many thousands of files
- Use Common use case template from D8.3
 - All WPs test as much as possible HEPCAL mini use cases
 - WP9 and WP10 test their own specific mini use cases

◆ Organisation of testing

- How are we going to control and monitor testing in next weeks in App TB?
 - Start slowly with basic functionality tests, and plan on an 'awareness' basis special tests such as large scale file handling and registration
 - WP8/9/10 keeping closely in touch and will inform App TB TB of plans



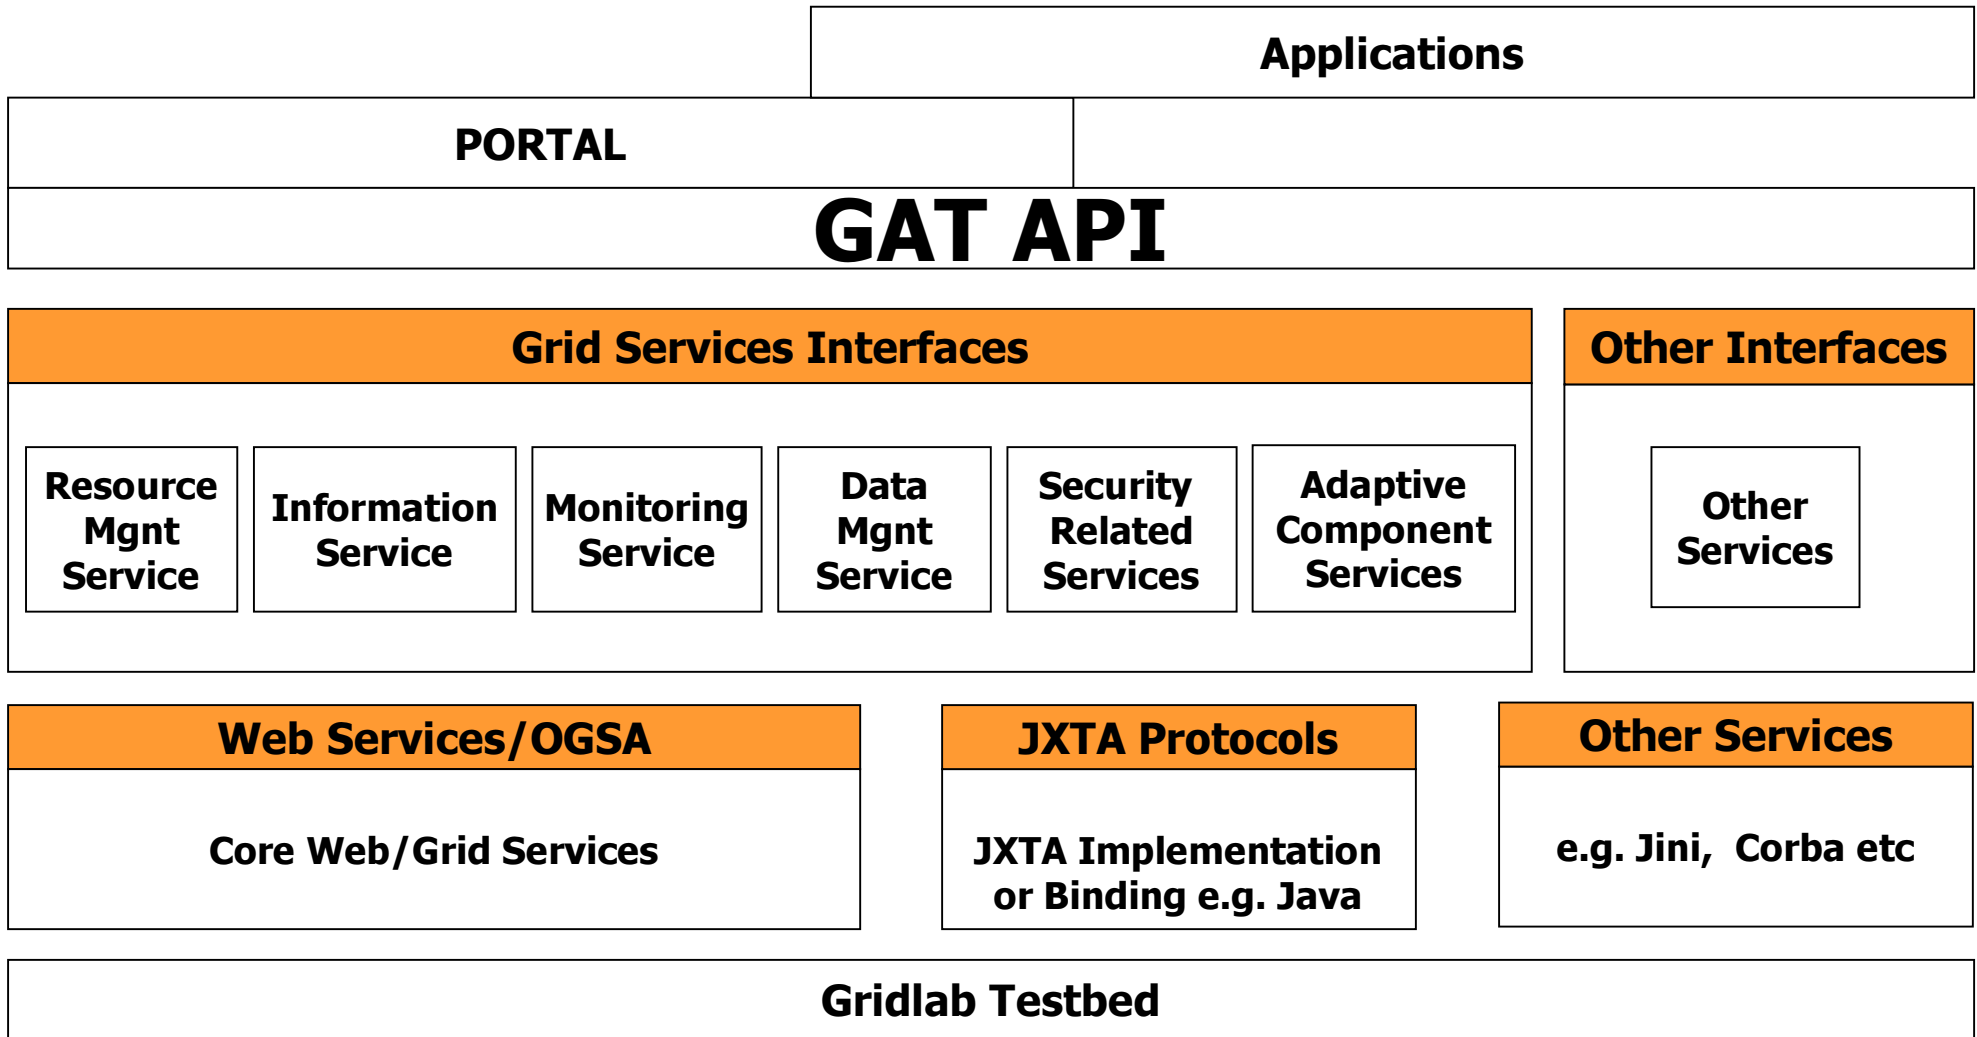
Output from AWG and a forward look

- ◆ **AWG produced** *'Joint Recommendations on middleware'* and *'Joint use cases'*
 - Applications are mapping use cases to 2.0 and 2.1 (summaries in deliverables)
- ◆ **The 'joint use cases' (based on HEPCAL template)** will be used as an input to EGEE applications software development work (common applications interface).
 - Other key inputs will come from *LCG ARDA (Grid Service Architecture)* and from work in other communities (*e.g. Gridlab GAT work*) - essential to have dialogue between these communities
- ◆ **EGEE applications group** includes HEP, Biomedicine and 'other applications'. So the 'AWG flavour' will continue here (see *G Wormser* talk this afternoon)



GAT Overview

I Taylor(Univ of Cardiff)





Follow up on use cases and GAT

- Gridlab looks at AWG joint list of use cases + HEPCAL and their mapping to GAT APIs
- GAT and AWG meets (around February) to discuss use cases and the 'essential subset of APIs' which are necessary
- AWG will compare this to ATF APIs and definitions coming out of LCG/ARDA together with ARDA people (this is essential work prior to EGEE start- Common Application Layer specification due July 1)

HealthGrid 2004

January 29th - 30th 2004,
Clermont-Ferrand, France

<http://clermont2004.healthgrid.org>

The aims of this conference are to reinforce and promote awareness of the possibilities and advantages linked to the deployment of GRID technologies in health. In this context "Health" does not involve only clinical practice but covers the whole range of information from molecular level (genetic and proteomic information) through cells and tissues, to the individual and finally the population level (social healthcare).





Summary

- ◆ *Applications must be totally focussed on evaluations and deliverable writing in the coming weeks, and then preparing the review*
- ◆ *The essential aim of the project should be to fulfill our 'promise' of delivering basic production quality software - (this is definitely higher priority than delivering 'new' features) - I.e. can we really handle many thousands of jobs and millions of files in a multi-site, multi-VO environment, and have high efficiencies?*
- ◆ *Hence our essential requests to the project are:*
 - *That absolute priority should be given to establishing a stable application TB*
 - Bug fixes
 - Maybe switching off features that prove fragile
 - *That we have (extracted from delivery notes) a short (?10 page) summary of the overall view of the 'reliable and fully tested middleware' functionality (and hence a warning off for the use of fancier, untested features, at least for the coming weeks into November)*
 - *That, if we are to have demos for the review, then continued support for the Application TB be provided up to the end of Feb*



◆ *THANKS*

- To all our friends in the middleware, testbed and project office who have put up with the nagging requests of the applications for (stability, stability, stability), and who have survived Datagrid together with us!
- See you soon at the review, and beyond, in other tests of our survival instincts and good humour!