



WP3 Status

Steve Fisher / RAL

30/9/2003

[<s.m.fisher@rl.ac.uk>](mailto:s.m.fisher@rl.ac.uk)



Plan of Talk

- EDG 2.0 Release
- Latest tag (2.1 candidate)
- Wrapping up EDG

R-GMA:

- A general Grid Information and Monitoring System answering SQL queries

- Offers global view of distributed Grid monitoring information

Status



- Release 2.0 on Application testbed showed a number of scalability problems
- These have all been fixed and tested on the WP3 testbed.
- Should be deployed ASAP

EDG 2.0 Release



- Revised APIs
 - The C API code totally rewritten
 - Removed deprecated methods throughout
- Improved tools
 - Command Line Interface (edg-rgma)
 - Browser
- New validation suite
 - edg-rgma-check
 - edg-rgma-examples
- Schema
 - GLUE 1.1
 - Service and ServiceStatus

Latest Tag

Latest Tag

- Robustness
- Performance
- Nagios Integration
- Security
- GRM/PROVE

With exception of GRM/PROVE integration all work related to making the system more robust

Performance



- Optimizelt revealed a number of problems
 - Slow Java I/O – now use NIO for streaming
 - SQL parsing – now use hand coded parser for the inserts
 - XML parsing too slow – DOM replaced by SAX
 - Too many threads – now a big reduction

R-GMA metrics



- A range of measurements:
 - Freshness
 - Completeness of information
 - Response time
- Application, development and WP3 testbeds
- Scripts interpret the measurements
 - Is the absence of information correct?
 - If it is lost – where is it lost?
- Results linked from WP3 web page

Nagios integration



- We are using Nagios to monitor the status of EDG services
 - including R-GMA services
- Nagios provides:
 - Graphical displays
 - Alarm mechanism
- Will also use it to display R-GMA metrics

Security



- We now have authentication in place
- If switched on:
 - you must have a proxy
 - Tomcat machines must have a certificate
 - Otherwise invisible to the user

GRM/PROVE



- To understand the performance of parallel applications
 - Mainly of interest to WP9 and WP10
- GRM now uses the Mercury monitor (from GridLab) to deal with high data rates – but currently no security.
- Uses R-GMA to locate the application and Mercury contact point

Final months

Final months of EDG



- Most effort will be needed in
 - Support
 - Documentation
 - Attention to any new bugs
- And “demos” (as time permits) of
 - Multiple VO support
 - Registry & Schema replication
 - Authorisation
 - Ranglia (Ganglia integration)
 - Can probably be deployed independently
 - Grid Services

Support



- Help experiments and other users to make best use of R-GMA
 - RB / R-GMA interface
 - BOSS
 - GANGA
 - D0
 - BaBar
 - Network Monitoring
 - Logging and Bookkeeping
 - UK – e-Science

Ganglia Integration



- Ganglia is popular cluster monitoring system
- Completing a component (Ranglia) to make Ganglia information available as R-GMA tables
- Uses CanonicalProducer – i.e. on demand generation.
- Must also interface to LEMON

Grid Services



- Already have schema and registry as OGSIGrid services
- We will continue the transition to Grid services
 - Will provide wrappers for compatibility with our current APIs.

Summary



- All R-GMA problems seen on application testbed are now believed to be eliminated
- Will continue to focus on robustness and reliability to meet the needs of LCG for both Information **and** Monitoring.