R-GMA: Now With Added Authorization

Tuesday 12 February 2008 15:00 (20 minutes)

R-GMA is currently being used by APEL, the ARDA Dashboard and Service Discovery on the LCG grid and by Grid Ireland. APEL uses a producer at each site to publish accounting data. A consumer is used to accumulate all of the data in a central location where it is migrated off line for later analysis. The ARDA dashboard has a consumer that pulls in monitoring data about the status of grid jobs as published by resource brokers. This enables real time monitoring of the progress of jobs on the grid. The Service Discovery API has an R-GMA plugin. A producer at each site publishes information about available services and their current status. This enables other middleware components to select required services. Grid Ireland are using R-GMA to monitor TCP logs. With the increased robustness of the code and the fine grained authorization we expect that users will find new applications for R-GMA.

3. Impact

The latest version of the R-GMA server improves on the current implementation. It is based around a new design that was derived from our experiences from the initial prototype and subsequent patches to it. Robustness and scalability are at the centre of the new design. Single points of failure have been removed and the servers have been made as autonomous as possible. Reliance on the delivery of individual control messages has also been removed. From a user's perspective, the improvements in functionality are the introduction of fine grained authorization and virtual databases (VDB). The authorization is done using SQL views of tables constructed dynamically from user defined rules and VOMS attributes. VDBs allow for the partitioning of data. We envisage that each VO would have one or more VDBs.

URL for further information:

http://hepunx.rl.ac.uk/egee/jra1-uk/index.html

4. Conclusions / Future plans

From the existing deployment we learned not to rely upon any single message being transmitted successfully. For the new deployment there will no longer be a central registry and schema service. Instead there will be several registry replicas per VDB. For the schema there will be a replica for each VDB at each site supporting that VDB with one defined as the master schema. The design permits alternative databases to be used. Currently we only support MySQL but Oracle will be added in the future.

Provide a set of generic keywords that define your contribution (e.g. Data Management, Workflows, High Energy Physics)

monitoring, information, message, relational

1. Short overview

R-GMA permits users to define their own data structures along with the fine grained authorization rules specifying who can write and read the data. They can then publish data via a producer API without knowledge of potential consumers. A consumer API is used to retrieve the permitted view of information published by the producers. Previous releases of the server have been patches to the original prototype. The new code has been re–engineered for robustness and scalability.

Authors: Dr WILSON, Antony (RAL); Dr FISHER, Steve (RAL)

Co-authors: Mr KUSEJU, Adebiyi (RAL); Mr DUNCAN, Alastair (RAL); Mr CHIDAMBARAM, Arun (RAL); Dr JIANG, Ming (RAL); Ms BHATTI, Parminder (RAL)

Presenter: Dr FISHER, Steve (RAL)

Session Classification: Monitoring, Accounting & Support

Track Classification: Existing or Prospective Grid Services