EGEE User Forum



Contribution ID: 71

Type: Oral contribution

User Applications of R-GMA

Thursday 2 March 2006 17:00 (30 minutes)

The Relational Grid Monitoring Architecture (R-GMA) provides a uniform method to access and publish both information and monitoring data. It has been designed to be easy for individuals to publish and retrieve data. It provides information about the grid, mainly for the middleware packages, and information about grid applications for users. From a user's perspective, an R-GMA installation appears as a single virtual database. R-GMA provides a flexible infrastructure in which producers of information can be dynamically created and deleted and tables can be dynamically added and removed from a schema. All of the data that is published has a timestamp, enabling its use for monitoring. R-GMA is currently being used for job monitoring, application monitoring, network monitoring, grid FTP monitoring and the site functional tests (SFT).

R-GMA is a relational implementation of the Global Grid Forum's (GGF) Grid Monitoring Architecture (GMA). GMA defines producers and consumers of information and a registry that knows the location of all consumers and producers. R-GMA provides Consumer, Producer, Registry and Schema services.

The consumer service allows the user to issue a number of different types of query: history, latest and continuous. History queries are queries over time sequenced data and latest queries correspond to the intuitive idea of current information. For a continuous query, new data are broadcast to all subscribed consumers as soon as those data are published via a producer. Consumers are automatically matched with producers of the appropriate type that will satisfy their query.

Data published by application code is stored by a producer service. R-GMA provides a producer service that includes primary and secondary producers. Primary producers are the initial source of data within an R-GMA system. Secondary producers can be used to republish data in order to co-locate information to speed up queries (and allow multi-table queries), to reduce network traffic and to offer different producer properties. It is envisaged that there will be numerous primary producers and one or two secondary producers for each subset of data. Both primary and secondary producers may use memory or a database to store the data and may specify retention periods. Memory producers give the best performance for continuous queries, whereas database producers give the best performance where joins are required.

It is not necessary for users to know where other producers and consumers are: this is managed by the local producer and consumer services on behalf of the user. In most cases it is not even necessary to know the location of the local producer and consumer services, as worker nodes and user interface nodes are already configured to point to their local R-GMA producer and consumer services.

There are already a number of applications using R-GMA. The first example is job monitoring. There was a requirement to allow grid users to monitor the progress of their jobs and for VO administrators to get an overview of what was happening on the grid. The problems were that the location in which a grid job would end up was not known in advance, and that worker nodes were behind firewalls so they were not accessible remotely.

SA1 has adopted the job wrapper approach, as this did not require any changes to the application code. Every job is put in a wrapper that periodically publishes information about the state of the process running the job and its environment. These data are currently being published via the SA1 JobMonitoring table within R-GMA. A second application has been written to run on the resource broker nodes. This application examines the logging and bookkeeping logs and publishes data about the changes in state of grid jobs. These data are made available via the SA1 JobStatusRaw table.

Both the producer in the job wrapper and the producers on the resource broker nodes make use of R-GMA memory primary producers. A database secondary producer is used to aggregate the data.

Other uses of R-GMA include application monitoring, network monitoring and gridFTP monitoring. There are a number of different ways to implement application monitoring including the wrapper approach, as the job monitoring, and instrumentation of the application code. Instrumentation of the code can mean using a logging service, e.g. log4j, which publishes data via R-GMA, or calling R-GMA API methods directly from the application code.

The network monitoring group, NA4, have been using R-GMA to publish a number of network metrics. They used memory primary producers in the network sensors to publish the data and a database secondary producer to aggregate the data.

SA1 have made use of the consumer service for monitoring grid FTP metrics. They have written a memory primary producer that sits on the gridFTP server nodes and publishes statistics about the file transfers. A continuous consumer is used to pull in all the data to a central location, from where it is written to an Oracle database for analysis. This was used for Service Challenge 3.

Two patterns have emerged from the use made of R-GMA for monitoring. In both patterns data is initially published using memory primary producers. These may be short lived and only make the data available for a limited time, e.g. the lifetime of a grid job. In one pattern data are made persistent by using a consumer to populate an external database which applications query directly. In the other pattern, an R-GMA secondary producer is used to make the data persistent and also make it available for querying through R-GMA.

In the coming months we plan to add support for multiple Virtual Data Bases, authorization within the context of a Virtual Data Base using VOMS attributes, registry replication, load balancing over multiple R-GMA servers and support for Oracle.

R-GMA is an information and monitoring system that has been specifically designed for the grid environment. It can be used by systems, VOs and individuals and is already in use in production.

Authors: Dr DJAOUI, Abdeslem (RAL); Mr DUNCAN, Alastair (RAL); Mr WILSON, Antony (RAL); Mr WALK, John (RAL); Dr CORNWALL, Linda (RAL); Dr CRAIG, Martin (RAL); Mr BYROM, Rob (RAL); Dr MIDDLETON, Robin (RAL); Dr FISHER, Steve (RAL); Mr HICKS, Steve (RAL)

Presenter: Dr FISHER, Steve (RAL)

Session Classification: 2c: Special type of jobs (MPI, SDJ, interactive jobs, ...) - Information systems

Track Classification: Special type of jobs (MPI, SDJ, interactive jobs, ...) and information systems