


The AMGA Metadata Catalogue

AMGA is the gLite Metadata Catalogue, designed to meet the requirements of the EGEE applications. Main features:


- **Modular back-end** – Supports Oracle, PostgreSQL, MySQL and SQLite
- **Modular front-end** – High-performance TCP Streaming interface and standards compliant Web-Services front-ends.
- **Hierarchical Organisation** – Metadata organised internally in a tree like structure.
- **Dynamic Schemas** – Schemas may be created, modified and deleted by clients at run-time.
- **Integration with Grid Security** – Supports Grid proxy authentication and VOMS authorization
- **Secure client connections using SSL**
- Authorization using **ACLs** with support for **user and group** management in the catalogue.
- **Replication** – Metadata collections can be replicated to improve fault-tolerance, scalability and performance when used by many clients spread across Grid sites.

Performance

For most of the common operations the performance of AMGA is within a few percent of direct database access, allowing applications to profit from Grid Security without a significant overhead.



Architecture




Read Performance*

*Data from LHCb Logging and Bookkeeping (Thanks to C. Cloff)

Metadata on the Grid


On a Grid, metadata is used for two purposes: file metadata and simplified grid-enabled database access.

- **File Metadata** - Applications running on a Data Grid typically generate and use millions of files, often distributed across several Grid sites. A Metadata Catalogue Service provides the means to **describe and discover the data** required by the users and their jobs.



Metadata Catalogue for File Metadata

- **Simplified Grid Enabled Database Access** – Traditional databases do not support the authentication and authorisation mechanisms used on the Grid, like Grid Proxies and VOMS. This makes them unsuitable for being used directly by grid users. A metadata service can provide a database-like service, while offering full integration with Grid Security.



Metadata Catalogue as A Grid-enabled Database Service

Some Applications using AMGA

Grid Jobs Tracking



Ganga is an interface to the Grid for distributed analysis. AMGA is being used for tracking highly dynamic information about job status.

Medical Data Manager

The **Medical Data Manager** was developed by the BioMed application area of EGEE, to demonstrate a secure service for accessing and manipulating medical images and associated metadata from the grid. AMGA is used both for file metadata and for storage of relational data.

Logging and Bookkeeping



AMGA is being evaluated for Logging and Bookkeeping. Currently LHCb is using AMGA to store 15 million log entries of jobs.

File Metadata for GIS applications



UNOSAT is a United Nations programme for providing satellite imagery to field personal working on humanitarian missions. AMGA is used in a prototype developed by UNOSAT that uses the Grid for storing, processing and providing images.

GILDA Applications

- **gMOD** – Grid Movie on Demand. Provides information about movies and streams them to users. AMGA stores detailed information about the movies.
- **gLibrary** – A digital library for electronic documents. AMGA stores and provides access to the information about the documents.

Replication and Distribution

The large size of Data Grids, some of them like the LCG having over 6,000 users and 200 sites, is a challenge that any Grid Service must address to operate successfully in such an environment. Metadata Catalogues, in particular, must serve requests from thousands of users and jobs, originating from geographical dispersed regions.

AMGA supports replication and distribution to provide the required scalability, dependability and fault-tolerance.


Features

- **Asynchronous Replication** – Scale to large number of nodes spread across WANs
- **Master/Slave** – Writes only at the master node, and replicated to all slaves.
- **Catalogue-level Replication** – Replicates metadata commands. Provides **database independence** (works with Oracle, PostgreSQL, MySQL)
- **Partial Replication** – A slave may replicate only part of the metadata tree.
- **Federation** – A node can subscribe to different parts of the metadata tree from different masters.


Performance and current status

Implementation and internal testing is complete. Initial performance results show:


- With one slave: ~100 metadata commands replicated per second
- Close to linear scalability with the increase in the number of slaves.
- Low overhead at the master



Usage Scenarios



Architecture



Performance and Scalability

References

- N. Santos, and B. Koblitz, "Distributed Metadata with the AMGA Metadata Catalog" in Workshop on Next-Generation Distributed Data Management - HPDC-15, Paris, France, June 2006
- N. Santos, and B. Koblitz, "Metadata services on the grid". In Proc. of Advanced Computing and Analysis Techniques (ACAT05), Zeuthen, Berlin, May 2005

Acknowledgments

The authors wish to thank the ARDA team for all of their help and support, and Viktor Pose for the exhaustive testing of AMGA. Nuno Santos would also like to thank the Portuguese Foundation for Science and Technology (FCT) for partially funding his work.

