# GRID DATABASE MANAGEMENT: ISSUES, REQUIREMENTS AND FUTURE DIRECTIONS

Sandro Fiore, Salvatore Vadacca, Alessandro Negro and Giovanni Aloisio
University of Salento & SPACI Consortium
Euro Mediterranean Centre for Climate Change
viale Gallipoli, 49 – 73100 Lecce - Italy
E-mail: {sandro.fiore, salvatore.vadacca, alessandro.negro, giovanni.aloisio}@unile.it

**KEYWORDS**

Data Grid, Grid Database, Database Access, Database Management, Database Integration.

**ABSTRACT**

Data grids allow to store, manage and share large data collections, huge amount of files, geographically distributed databases, etc. across virtual organizations. Grid data and metadata management is becoming more and more important as the number of involved data sources is continuously increasing and decentralizing. Grid database management services represent a basic and fundamental building block for the next generation petascale production grids. In this paper we present the fundamental concepts related to grid-database access, management and integration, highlighting main requirements and issues, describing research activity in this area and dealing with new Open Grid Forum related standards.

## INTRODUCTION

Grids encourage and promote the publication, sharing and integration of scientific data, distributed across Virtual Organizations (Foster and Kesselman 1998). Grid computing is widely regarded as a new field, distinguished from traditional distributed computing owing to its main focus on large-scale resource sharing and innovative high-performance applications.

It can be considered as an enabling paradigm allowing organizations to easily share, virtualize, integrate and efficiently manage wide spread resources (computational and data sources, sensors and instrumentations, etc.) in a grid environment. Grids link together servers, data sources, tools, services and applications into a single environment by means of a specific glue named *grid middleware*.

In the last years many efforts were devoted to the management (both coarse and fine grained) of data (grid-storage services, storage resource managers, metadata services, replica catalogues, grid-database access and integration services, etc.). Data management represents the real challenge for the next generation petascale grid environments since current production grids are able to produce hundreds of petabytes of data in the next years.

In the last few years, there was an increasing interest in fine grained (database related) grid data management activities and services connected with database access, metadata management, data integration, data transformation, data flow, etc.

Grid Services for database access and integration play a strategic role and provide added value to a production grid environment since they allow to aggregate data, join datasets stored at different sites, infer new knowledge by analyzing structured and distributed data, manage monitoring and accounting information, etc.

In this paper we will talk about the fundamental concepts related to grid-database access, management and integration, highlighting main requirements and issues, describing research activity in this area and dealing with new Open Grid Forum related standards.

The outline of the paper is as follows: in the 2nd section we talk about data grids, whereas in the 3rd section we describe in detail key issues. The 4th section concerns data requirements for grid database services as well as the 5th section shortly introduces new standards and specifications in this area. The 6th section recalls related works whereas the 7th section describes a specific case study (GRelC Project). Finally, we draw our conclusions.

## GRID PARADIGM AND DATA GRIDS

Production grids (i.e. EGEE, Teragrid (Pennington 2002)) produce huge amounts of scientific data that must be available “on the grid” to scientific and engineering applications for further analysis and computations.

Data Grids provide the proper data service framework for Computational Grids, creating virtualized access to widespread data sources (both files and databases). Even if in the last years research activities related to Data Grids have mainly focused on coarse grained data management (files), accessing and integrating legacy/new databases (records) is becoming a fundamental issue.

A grid infrastructure is basically made up of two components: computational and data grids.

While a *computational grid* provides the computing power needed to run applications, a *data grid* provides a robust framework for data management services that enables data access, integration, transfer, replication, virtualization, distribution, etc.

Proceedings of the 2008 High Performance
Computing & Simulation Conference ©ECMS
Waleed W. Smari (Ed.)
ISBN: 978-0-9553018-7-2 / ISBN: 978-0-9553018-6-5 (CD)

## GRID-DBMS: KEY ISSUES

A grid database management system (Aloisio et al. 2004) should provide transparent, secure and efficient *management* (in terms of database access, integration, federation, transformation, etc.) of data sources (relational, hierarchical, object-oriented, etc.) in a grid environment. Since the beginning of the *grid era* many efforts were directed towards *computational access* (e.g. job submission/monitoring, Globus GRAM (Foster 2005), etc.) and *storage management* (e.g. file transfer, GridFTP (Allcock et al. 2005) based storage, SRB (Rajasekar et al. 2003)). Grid database management was addressed starting from the year 2000 (EDG-Spitfire (Bell et al. 2002), GRelC (Aloisio et al. 2005), and OGSA-DAI (Antonioletti et al. 2005)).

In the following we describe some basic elements connected with database management in distributed environments, highlighting how they impact on the application domains (i.e. e-Science) and why they are so relevant for end users (i.e. scientists). In particular, next subsections will be devoted to the discussion of data representation, data organization, data models, query languages, data access, data integration, access control and data flow.

### Data Representation

In order to be domain-independent, data grids must provide support (in terms of access and management) to every type of data format, structure and representation. Data can be both structured and unstructured, characterized by different formats, coding, precision, accuracy and semantics.

Some examples concern bioinformatics (i.e. textual files, relational data sources), astrophysics (i.e. relational DBMS with postGIS extensions), climate scientists (i.e. XML) data banks.

### Data Organization, Data Models and Query Languages

Data can be organized following several data models such as relational and hierarchical. Support in terms of relational or XML engines is widely provided by existing systems: Postgres, MySQL, IBM/DB2, etc. as well as XIndice, eXist, etc. Such DBMSs provide full support in terms of database access and management functionalities, API, SDK, CLI, etc. Different data models adopt different query languages such as SQL (for the relational one) and XPath and XQuery (for the hierarchical one); data grids must provide support to all of them.

### Data Access

Even if DBMSs provide a lot of functionalities for the management of data sources, they are not fully compatible with existing grid middleware (i.e. gLite or Globus (Foster and Kesselman 1997)). They can be accessed in grid by using a "*grid-DBMS*" interface. This grid interface (which has to play a front-end role) must, obviously, provide full support to all of the query languages (SQL, XQuery, XPath, etc.) concerning the target data resources (transparency requirement with regard to the query language).

The specific part of the grid-DBMS that makes a data resource accessible in grid (or "grid enabled") is called Grid Database Access Service (Grid DAS).

It must provide secure, transparent, robust and efficient access to heterogeneous and distributed databases exposing standard interfaces to enable interoperability with other grid components and/or services.

Several research projects exploit the *service-in-the-middle* or *front-end* approach to provide such kind of functionalities, that is, they focus on the development of a transparent, secure and robust grid interface to existing DBMSs. On the contrary, vendor-specific products (i.e. Oracle 11g) generally exploit an *embedded approach* providing within the product, software modules to run on a grid environment (e.g. GSI (Tuecke 2001) support).

### Data Integration

While the Grid DAS is a basic service to expose databases in grid (it provides a first level of virtualization), the Grid Data Integration Service (Grid DIS) is a further necessary building block if we want to provide aggregation capabilities (second level of virtualization). An interesting example is the OGSA-DQP (Alpdemir et al. 2003).

A Grid DIS can be centralized or distributed and in some cases it is integrated into the related Grid DAS providing what we call a Grid DAIS (e.g. GRelC DAIS which will be described later).

Data integration is strongly challenging since it allows both to integrate data within several *application-level* domains (bioinformatics, astrophysics, financial, etc.) and *system-level* distributed environments for monitoring and accounting purposes, etc.

### Data Access Control

Data access control is more important to ensure that the confidentiality of the data is preserved/maintained against unauthorized accesses.

The facilities that the Grid provides to control access must be very flexible in terms of the combinations of restrictions, available policies, etc. User-centric and VO-centric data access control allow managing policies at each level of granularity addressing local site autonomy and user-level policies management (in the first case) and flexibility, scalability and manageability in the VO-level policies management (in the second case).

A combined User-VO data access control allows mixing the benefits related to the two approaches (any combination of insert, update, and delete privileges can be defined with the right level of granularity).

Moreover, the Grid must provide the ability to control access based on user role (as it usually happens for DBMSs). Role based access control is fundamental for collaborative working, when several individuals may perform the same role at the same time and provides a

scalable and manageable way to split users in sub-classes with specific and well-known privileges. Granting and revoking activities must be dynamically performed by administrators and should be easily carried out by using high level interfaces such as data grid portals.
Data access policies should be managed at the Grid-DBMS layer, without entirely relying on the back-end framework. This could (i) enable data access control for trivial data resources such as text files and (ii) prevent the access attempts to the back-end systems for unauthorized users.

## MAIN REQUIREMENTS

In the following we highlight key requirements related to grid-database management, taking into account the most important ones: transparency which is strongly connected with data virtualization, security which is fundamental to protect data, efficiency as a performance index, and, finally, interoperability to ease grid data service composition/interaction.

### Transparency

Transparency is a common requirement for grid services and fundamental to make virtualization a reality. There are various possible types of transparency in a distributed environment. In particular, it relates to:

1. physical data location: the physical location of a database in the grid must be hidden/virtualized by the grid service;
2. naming: an application must be able to access a data source without knowing its name or location. These kind of information must be managed by means of mapping, alias, etc. which conceal data that are not relevant to the end-user, such as connection string for the databases, DBMS port, login and password, etc.;
3. data replication: replication of data improves performance, availability and fault tolerance. The user must not be aware of the existence/management of multiple physical copies of the same data source; she has just to deal with the logical (virtualized) data source name.
4. DBMSs heterogeneity: today many different RDBMSs exist, such as ORACLE, IBM/DB2, PostgreSQL, MySQL, SQLite, etc. Moreover, an increasing number of applications interact with not relational databases such as flat files and XML-based documents in the bioinformatics and climate change domains. This kind of heterogeneity (which includes different APIs, data types, physical support) must be properly handled in order to provide a uniform access interface to different data sources and a grid database access service independent of the back-end systems.

### Efficiency

Performance plays a fundamental role in the data grid environment. High throughput, concurrent accesses, fault tolerance, reduced communication overhead, etc., are important goals that must be achieved by exploiting among the others data localization and query parallelism. Moreover, efficient data delivery mechanisms can reduce the connection time (parallel streams) and the amount of transferred data (data compression).

### Security

Security is crucial for the management of a database in data grid environment. Data security aims at protecting data against unauthorized accesses by (i) preventing unauthorized users from accessing data and (ii) protecting information exchanged in the data grid network. Authentication is strongly required to check user's identity; authorization concerns privileges and read/write permissions. Most important production/research grids adopt the de-facto standard for security Globus Grid Security Infrastructure (GSI). It provides full security support concerning data encryption, data integrity, protection against replay attacks and detection of out of sequence packets. GSI is widely used both in gLite and Globus based grid environments.

### Interoperability

Interoperability can be achieved by standard adoption. Today the adopted paradigm is basically service oriented; more specifically WS-I approach (which means based on SOAP, XML and WSDL W3C standards) is well suited for basic interoperability. OGF specifications issued by the DAIS-WG and discussed in the following section mainly focus grid database access aspects.

## NEW STANDARDS AND SPECIFICATIONS

From a standardization point of view, in 2002 the Global Grid Forum (now Open Grid Forum) established a working group named DAIS (Data Access and Integration) to define a complete and effective set of specifications about these challenging topics. Since the beginning, the DAIS Working Group provides an umbrella under which many efforts are undertaken and people from all around the world are grouping together, giving a strong contribution in this area to the scientific grid community.
Interesting activities about Database Access and Integration recently produced the WS-DAI (Web Service Data Access and Integration) family of specifications (WS-DAI, WS-DAIR and WS-DAIX) (Antonioletti et al. 2006) which defines a set of web service interfaces to relational or XML data resources.
The base interfaces and properties for data access services are described in the Web Services Data Access and Integration (WS-DAI) specification. The WS-DAIR specification extends WS-DAI interfaces to allow access to and provide descriptions of relational data resources. Relational data resources are assumed to be composed of tabular data structures such as relations

and resultsets which are typically accessed either using SQL queries or by row iteration, respectively. The WS-DAIX specification extends WS-DAI interfaces to allow access to and provide descriptions of XML databases. XML data resources are assumed to consist of collections of XML documents that are accessed and modified using XPath, XQuery and/or XUpdate. Interfaces are provided for these languages in this specification.

The keyword highlighted by this standardization activity is interoperability, which can be achieved by different grid middleware providing reference implementations of these specifications.

## MAIN PROJECTS

In the DataGrid area several projects addressed *grid database management*. In particular, the first three were: GRelC (which will be widely described in the next Section), Spitfire and OGSA-DAI.

The *Spitfire Project* was part of the Work Package 2 of the European Data Grid Project and provided a means to access relational databases from the grid. It was a very thin layer on top of an RDBMS (by default MySQL) that provides a JDBC driver. It used Web Service technology (Jakarta Tomcat) to provide SOAP-based RPC (through Apache Axis) to a few user-definable database operations.

The *Open Grid Services Architecture Data Access and Integration* (OGSA-DAI) (Karasavvas et al. 2005) is another project concerned with constructing middleware to assist with access and integration of data from separate data sources via the grid. It is engaged in identifying the requirements, designing solutions and delivering software that will meet this purpose. The project was conceived by the UK Database Task Force and is working closely with the Global Grid Forum DAIS-WG and the Globus team. OGSA-DAI provides a more complex framework with regard to Spitfire and it is currently used in several e-Science projects.

## CASE STUDY: GRELC PROJECT

The Grid Relational Catalog (GRelC) is a research project started at the University of Salento (Italy) and addressing grid-database management issues. It basically aims at providing data grid solutions to access, manage and integrate data sources (e.g. relational databases).

Currently, the top service provided by the GRelC project is GRelC DAIS. It is a general purpose data grid service for database access and integration. This service acts as a standard front-end based on the well-known SOA approach. It provides both basic and advanced primitives to access, query, integrate, manage and interact with different data sources, providing a high level of transparency, concealing the back-end heterogeneity, middleware details related to Globus GSI and VOMS (Alfieri et al. 2003) and other low level issues. It is WS-I based, runs both on Globus and gLite grid middleware/environments and it provides efficient grid-enabled query mechanisms. From a security point of view, it supports both global (by means of VOMS) and local (on the GRelC DAIS side) authorization levels which means (i) fine and coarse grained data policies support and (ii) role-based management.

The GRelC DAIS provides a wide set of functionalities which includes: query submission, grid-database management related to user/VO/ACL, etc. advanced functionalities to transparently and securely integrate heterogeneous, distributed and geographically spread grid data sources (through P2P (Aloisio et al. 2007) connected GRelC DAIS nodes), etc. Moreover, it offers efficient data delivery (query resultsets) exploiting compression and streaming. Support for synchronous and asynchronous queries, is also provided.

The GRelC DAIS is able to transparently and securely integrate heterogeneous, distributed and geographically spread grid data sources, through a connected P2P-based network of GRelC DAIS nodes. The GRelC DAIS is very versatile so it can be used both at VO and site level. It can/is used in both ways depending on VO/user/database constraints and requirements. There is no single point of failure and no centralized management for this service due to the scalable P2P architecture.

Finally, by means of the GRelC Portal, GRelC DAIS nodes can be managed via web. The GRelC Portal eases the access and integration of grid-databases. It completely replaces the Command Line Interface, does not need additional configuration/installation of software and it provides a seamless and ubiquitous way to manage data sources in a grid environment.

Currently the GRelC DAIS is used as core service of the Euro-Mediterranean Centre for Climate Change (CMCC) (Fiore et al. 2008) grid metadata handling framework. Moreover, a wide deployment on the GILDA t-Infrastructure (Andronico et al. 2005) is also available for tutorials and training activities.

## CONCLUSIONS

Data Grids represent the basic framework for next generation petascale grid environments. They provide a set of services to store, access, share, manage, distribute, synchronize and integrate massive amounts of data distributed across heterogeneous and geographically spread grid resources.

In this survey we presented the basic concepts related to grid-database access, management and integration, highlighting main requirements (with particular emphasis on transparency/virtualization) and issues.

We also described research activity in this area, in particular describing the GRelC project as case study.

Finally, we discussed the novel Open Grid Forum WS-DAI family of specifications as a key to address interoperability among grid database access services.

## REFERENCES

Alfieri, R.; R. Cecchini; V. Ciaschini; L. dell'Agnello; A. Frohner; A. Gianoli; K. Lorentey; and Fabio Spataro. 2003. "VOMS, an Authorization System for Virtual

Organizations". In *Proceedings of European Across Grids Conference*. 33-40.
Allcock, W.; J. Bresnahan; R. Kettimuthu; M. Link; C. Dumitrescu; I. Raicu; and I. Foster. 2005. "The Globus Striped GridFTP Framework and Server". *ACM Press*.
Aloisio, G.; M. Cafaro; S. Fiore; and M. Mirto. 2004. "The GRelC Project: Towards GRID-DBMS". In *Proceedings of the IASTED PDCN Conference* (Innsbruck, Austria, Feb. 17-19). IASTED, 1-6.
Aloisio, G.; M. Cafaro; S. Fiore; and M. Mirto. 2005. "The Grid Relational Catalog Project". In *Advances in Parallel Computing - Grid Computing: The New Frontiers of High Performance Computing*. L. Grandinetti. Elsevier, 129-155.
Aloisio, G.; M. Cafaro; S. Fiore; M. Mirto; and S. Vadacca. "GRelC Data Gather Service: a Step Towards P2P Production Grids". 2007. In *Proceedings of 22nd ACM Symposium on Applied Computing* (Seoul, Korea, Mar. 11-15). Vol. I, 561-565.
Alpdemir, M.N.; A. Mukherjee; N.W. Paton; P. Watson; A.A.A. Fernandes; A. Gounaris; and Jim Smith. 2003. "OGSA-DQP: A service-based distributed query processor for the Grid". In *Proceedings of UK e-Science All Hands Meeting Nottingham* (EPSRC, Sep. 24).
Andronico, G.; V. Ardizzone; R. Barbera; R. Catania; A. Carriera; A. Falzone; E. Giorgio; G. La Rocca; S. Manforte; M. Pappalardo; G. Passaro; and G. Platania. 2005. "GILDA: The Grid INFN Virtual Laboratory for Dissemination Activities". *TRIDENTCOM*, 304-305
Antonioletti, M.; M.P. Atkinson; R. Baxter; A. Borley; N.P. Chue Hong; B. Collins; N. Hardman; A. Hume; A. Knox; M. Jackson; A. Krause; S. Laws; J. Magowan; N.W. Paton; D. Pearson; T. Sugden; P. Watson; and M. Westhead. 2005. "The Design and Implementation of Grid Database Services in OGSA-DAI". *Concurrency and Computation: Practice and Experience*, Vol. 17, Issue 2-4, 357-376.
Antonioletti, M.; A. Krause; N.W. Paton; A. Eisenberg; S. Laws; S. Malaika; J. Melton; and D. Pearson. 2006. "The WS-DAI Family of Specifications for Web Service Data Access and Integration". *ACM SIGMOD Record*, Vol. 35, No. 1, 48-55.
Bell, W.H.; D. Bosio; and W. Hoschek. 2002. "Project spitfire - towards Grid Web Service databases". *Database Access and Integration Services Working Group, 5th Global Grid Forum* (Edinburgh, UK).
Fiore, S.; Vadacca S.; Negro A.; and G. Aloisio. 2008. "Euro-Mediterranean Centre for Climate Change Data Grid". Submitted to "DAPSYS 08 (Debrecen, Hungary, Sep 3-5)
Foster, I.; and C. Kesselman. 1997. "Globus: A Metacomputing Infrastructure Toolkit." *International Journal on Supercomputer Applications* Vol. 11, No. 2, 115-128.
Foster, I. and C. Kesselman. 1998. *The Grid: Blueprint for a New Computing Infrastructure*. Morgan Kaufmann.
Foster, I. 2005. "Globus Toolkit Version 4: Software for Service-Oriented Systems". In *Proceedings of IFIP International Conference on Network and Parallel Computing*. 2-13.
Karasavvas, K.; M. Antonioletti; M.P. Atkinson; N.P. Chue Hong; T. Sugden; A.C. Hume; M. Jackson; A. Krause; and C. Palansuriya. 2005. "Introduction to OGSA-DAI Services". *Lecture Notes in Computer Science*, Vol. 3458, 1-12.
Pennington, R. 2002. "Terascale clusters and the teragrid". In *Proceedings of High Performance Computing Asia* (Dec 16-19), 407–413.
Rajasekar, A.; M. Wan; R. Moore; W. Schroeder; G. Kremenek; A. Jagatheesan; C. Cowart; B. Zhu; S.Y. Chen; and R. Olschanowsky. 2003. "Storage Resource Broker - Managing Distributed Data in a Grid". *Computer Society of India Journal, Special Issue on SAN*, Vol. 33, No. 4 (Oct), 42-54.
Tuecke, S. 2001. "Grid Security Infrastructure (GSI) Roadmap". *Internet Draft*.

## AUTHOR BIOGRAPHIES

**SANDRO FIORE** was born in Galatina (LE) in 1976. He received a summa cum laude Laurea degree in Computer Engineering from the University of Lecce (Italy) in 2001, as well as a PhD degree in Informatic Engineering on Innovative Materials and Technologies from the ISUFI-University of Lecce in 2004. Research activities focus on parallel and distributed computing, specifically on advanced grid data management. Since 2004, he is a member of the Center for Advanced Computational Technologies (CACT) of the University of Salento and technical staff member of the SPACI Consortium. Since 2001 he has been the Project Principal Investigator of the Grid Relational Catalog project. Dr. Fiore was involved in the EGEE project (Enabling Grids for E-science) and is currently involved in the EGEE-II project and other national projects (LIBI). Since June 2006, he leads the Data Grid group of the Euro-Mediterranean Centre for Climate Change (CMCC) in Lecce (Italy). He is author and co-author of more than 40 papers in refereed journals/proceedings on parallel and grid computing and holds a patent on advanced data management.
His e-mail address is sandro.fiore@unile.it and his web-page can be found at http://grelc.unile.it/ person.php?surname=Fiore

**SALVATORE VADACCA** was born in Galatina (LE) in 1982. He received summa cum laude bachelor and master degrees in Computer Engineering from the University of Lecce, Italy in 2003 and 2006, respectively. His research interests include data management; distributed, peer-to-peer and grid computing; as well as web design and development. Since 2003, he has been a team member of the GRelC Project. In 2006 he joined the Euro-Mediterranean Centre for Climate Change (CMCC) in Lecce, Italy, where he works in the Data Grid group. Since April 2008 he is a Ph.D. student in Interdisciplinary Science and Technology at ISUFI, Lecce.
His e-mail address is salvatore.vadacca@unile.it and his web-page can be found at http://grelc.unile.it/ person.php?surname=Vadacca

**ALESSANDRO NEGRO** was born in S. Pietro Vernotico (BR) in 1981. He received a summa cum laude bachelor degree in Computer Engineering from the University of Lecce, Italy in February 2004 and a summa cum laude master degree in Computer Engineering from the same University in April 2006. His research interests include GUI Development, Data

Management, Distributed, Grid Computing and Web Services. He is also interested in web and pattern-oriented design. Since 2003 he is a Team Member of the GRelC Project. Since September 2006 he holds a contract position in the LIBI project for the University of Lecce. Since April 2008 he is a Ph.D. student in Interdisciplinary Science and Technology at ISUFI, Lecce.
His e-mail address is alessandro.negro@unile.it and his web-page can be found at http://grelc.unile.it/person.php?surname=Negro

**GIOVANNI ALOISIO** is Full Professor of Information Processing Systems at the Engineering Faculty of the University of Lecce. His research interests are in the area of High Performance Computing, Distributed and Grid Computing and are carried out at the Department of Innovation Engineering of the University of Lecce. He is also the Director of the CACT (Center for Advanced Computational Technologies) of the National Nanotechnology Laboratory (NNL/CNR-INFM). As director of the CACT, which is an international partner of the US National Partnership for Advanced Computational Infrastructure (NPACI), he leads joint research projects on Grid both at a national and international level. He has been a co-founder of the European Grid Forum (Egrid) which then merged into the Global Grid Forum (GGF). He has founded SPACI (Southern Partnership for Advanced Computational Infrastructures), a consortium on ICT and grid computing among the University of Lecce, the University of Calabria and HP Italia. The Consortium is a follow-up of the SPACI project funded by the Italian Ministry of Education, University and Technological Research, to pursue excellence in the field of Computational Science and Engineering. He is also the Vice-Director of the NanoSciences and Grid Computing Section of the "Scuola Superiore ISUFI" and CEO of the SPACI Consortium. He is Director of the Division "ICT and Operations" of the Euro-Mediterranean Centre for Climate Change (CMCC). He is the author of more than 100 papers in refereed journals on parallel & grid computing.
His e-mail address is giovanni.aloisio@unile.it and his web-page can be found at http://grelc.unile.it/person.php?surname=Aloisio