


The design and implementation of Grid database services in OGSA-DAI



Mario Antonioletti¹, Malcolm Atkinson², Rob Baxter¹, Andrew Borley³, Neil P. Chue Hong¹, Brian Collins³, Neil Hardman³, Alastair C. Hume¹, Alan Knox³, Mike Jackson¹, Amy Krause¹, Simon Laws³, James Magowan³, Norman W. Paton⁴, Dave Pearson⁵, Tom Sugden¹, Paul Watson^{6,*},[†] and Martin Westhead¹

¹*EPCC, University of Edinburgh, James Clerk Maxwell Building, Mayfield Road, Edinburgh EH9 3JZ, U.K.*

²*National e-Science Centre, 15 South College Street, Edinburgh EH8 9AA, U.K.*

³*IBM United Kingdom Ltd, Hursley Park, Winchester S021 2JN, U.K.*

⁴*Department of Computer Science, University of Manchester, Oxford Road, Manchester M13 9PL, U.K.*

⁵*Oracle UK Ltd, Thames Valley Park, Reading RG6 1RA, U.K.*

⁶*School of Computing Science, University of Newcastle, Newcastle-upon-Tyne NE1 7RU, U.K.*

SUMMARY

Initially, Grid technologies were principally associated with supercomputer centres and large-scale scientific applications in physics and astronomy. They are now increasingly seen as being relevant to many areas of e-Science and e-Business. The emergence of the Open Grid Services Architecture (OGSA), to complement the ongoing activity on Web Services standards, promises to provide a service-based platform that can meet the needs of both business and scientific applications. Early Grid applications focused principally on the storage, replication and movement of file-based data. Now the need for the full integration of database technologies with Grid middleware is widely recognized. Not only do many Grid applications already use databases for managing metadata, but increasingly many are associated with large databases of domain-specific information (e.g. biological or astronomical data). This paper describes the design and implementation of OGSA-DAI, a service-based architecture for database access over the Grid. The approach involves the design of Grid Data Services that allow consumers to discover the properties of structured data stores and to access their contents. The initial focus has been on support for access to Relational and XML data, but the overall architecture has been designed to be extensible to accommodate different storage paradigms. The paper describes and motivates the design decisions that have been taken, and illustrates how the approach supports a range of application scenarios. The OGSA-DAI software is freely available from <http://www.ogsadai.org.uk>. Copyright © 2005 John Wiley & Sons, Ltd.

KEY WORDS: Grid; databases; e-Science; Web Services

*Correspondence to: Paul Watson, School of Computing Science, University of Newcastle, Newcastle-upon-Tyne NE1 7RU, U.K.

[†]E-mail: paul.watson@newcastle.ac.uk



1. INTRODUCTION

This paper examines how databases can be integrated into OGSA Grid environments through the use of the OGSA-DAI (OGSA Data Access and Integration) service-based architecture.

1.1. The need for databases in Grid environments

Early Grid applications, mainly motivated by an attempt to cater to scientific-based computing, were often closely associated with devices or tools that read and/or generated flat files. Consequently, support for files rather than for the management of structured data had the highest profile in the early Grid toolkits (e.g. [1]). However, over time, the file management systems and registries associated with Grid toolkits themselves became complex, and database management systems (DBMSs) were increasingly used to store Grid metadata (e.g. [2]). Contemporaneously, the requirements of the scientific computing community have become more sophisticated with, for example, biological and astronomical communities generating large quantities of data that increasingly use databases for storage and retrieval. Similarly, engineering, medical research, healthcare and many governmental systems can also take advantage of Grids that access and integrate multiple and distributed collections of structured data. These data need to be made available and accessible to distributed sets of users and their applications, which makes these communities ideal candidates to adopt Grid-based infrastructures.

1.2. The OGSA Grid environment

Web Services emerged from the business computing world with the aim of supporting business-to-business relationships in a language-neutral, platform independent way. They are loosely coupled, distributed components that interact with each other through the exchange of messages. To achieve the goal of interoperability, a collection of specifications are currently progressing through standardization bodies such as the W3C and OASIS. The architectural principles on which Web Services are based may be found in [3] and a recent overview is provided in [4].

Starting in 2001, researchers, initially led by Globus (<http://www.globus.org>) and IBM, began developing new Grid standards and technology. The aim was to merge the understanding developed through the design and deployment of early Grid applications (e.g. on Globus Toolkit 2) with the emerging Web Services middleware. Their goal was to allow Grid developers to exploit the huge commercial investment in Web Services infrastructure and tooling. The result was the Open Grid Services Architecture (OGSA) [5], a high-level framework designed to support dynamic virtual organizations that share independently administered data and resources seamlessly across a network of heterogeneous computers. Still under development, this architecture will define the major functional components needed to meet those requirements. OGSA-DAI is one of those high-level functional components.

1.3. Databases and the OGSA Grid environment

It has become increasingly clear that if the OGSA is to support a wide range of communities, then database integration is vital. This paper presents an approach to integrating databases into OGSA so



that communities may publicize and share structured data resources, and applications may access and update data stored in them.

There are two main dimensions of complexity to the problem: reconciling implementation differences between server products within a single database paradigm (IBM, Oracle, Microsoft, etc.), and the variety of database paradigms (object, relational, XML, etc.). Unsurprisingly, none of the currently deployed DBMSs support OGSA Grid integration. Each DBMS is however the result of many hundreds of person-years of effort to provide a wide range of functionality, valuable programming interfaces and tools, and important properties such as security, performance and dependability. As these attributes will be required by applications, we believe that building new Grid-enabled DBMSs from scratch is both unrealistic and uneconomic. We have shown that it is also unnecessary, by developing middleware that integrates existing DBMSs with the OGSA Grid.

Our intention is that, over time, the effort required to achieve this will diminish. As the Grid becomes commercially important, database vendors will embed the middleware functionality directly into their products to provide 'out-of-the-box' support for OGSA Grid integration, by directly supporting the emerging OGSA Grid standards, for example, for data access, notification and data dissemination. This simplifies the structure of applications and yields substantial performance gains. Similarly, it is vital that those designing standards for Grid middleware take into account the special requirements to easily integrate databases across a Grid. One of the motivations for OGSA-DAI is to expose and formulate such requirements. Together, these converging developments will reduce the amount of middleware required to integrate databases into the OGSA Grid.

Therefore, one of the aims of the OGSA-DAI project is to explore the requirements and approaches to integrating databases into the OGSA Grid in order to influence both the emerging standards and established database vendors. OGSA-DAI has designed, developed and released a collection of services for integrating database access and description with the core capabilities of OGSA, thus allowing structured data resources to be seamlessly integrated into OGSA Grid applications. We are now gaining feedback from early adopters.

1.4. Related work

This work focuses on the issues surrounding the design of an efficient, generic, service-based interface to databases. Most previous work on interfaces has been language dependent and not service based (e.g. JDBC). Other work on service-based interfaces to databases is generally more limited in scope or different in context. For example, several relational database vendors support the integration of their products with Web Services (e.g. [6,7]). Such proposals tend to support the calling of Web Services from within SQL queries, the creation of Web Services from stored procedures, and the publication of Web Services based on database requests. This approach contrasts with, and complements, the development of generic database access services such as those provided by OGSA-DAI.

In a Grid setting, the European Data Grid has developed Spitfire [8], a Web Service interface to relational databases for metadata management. Spitfire has developed an infrastructure that allows a client to query a relational database over GSI-enabled HTTP(S). An XML-based protocol is used to represent the query, and its result. The system supports role-based security: a client can specify the role they wish to adopt for a query execution, and a mapping table in the server checks that they are authorized to take on this role. OGSA-DAI has adopted a similar approach to Spitfire for authentication, but differs in supporting multiple access languages, asynchronous and third-party



delivery, as well as an extensible document-based request style interface as opposed to Spitfire's finer-grained operations.

This latter difference is an important aspect of the OGSA-DAI design. Consumers send rich documents that may contain multiple activities (e.g. a set of queries) to OGSA-DAI servers. When compared with database services that only support fine-grained interactions (e.g. single queries), this reduces the number of interactions (important for increasing server and application performance) and offers the opportunity for the server to perform local optimizations (e.g. not materializing intermediate results). Further, unlike other approaches, OGSA-DAI allows developers to define their own activities and make them available to consumers. This feature has been exploited by a number of research groups and could, for example, be used to expose specialist functionality such as local data mining capability to database consumers.

There is clearly also a relationship between OGSA-DAI and other data Grid functionalities. For example, efficient movement [9] and replication [10] of data are core functionalities in a data Grid. OGSA-DAI can make use of existing Grid data movement functionalities, as described in Section 4.3, and a data replication service could use OGSA-DAI to read data from or write data to structured data stores. Higher-level Grid data management systems, such as SRB [11] or Chimera [12], could use OGSA-DAI either to access structured data resources, or to provide access to their own metadata. As such, OGSA-DAI should be seen as one of a range of components that together support access, sharing, management and coordinated use of data on the Grid.

1.5. Terms used in this paper

The term *database* refers to systems used for managing structured data, such as relational databases and collections of semi-structured files. The term *data resource* is used virtually synonymously with database, the only difference being that it could be considered to include devices such as telescopes and scanners that produce or accept structured data, or data integration infrastructures that provide the illusion that a single database is being accessed. The term *data source* denotes that subset of data resources that only produce data, e.g. read only data collections. For these two terms, sometimes the word 'data' is omitted as it is obvious from the context.

A DBMS is responsible for the storage and management of one or more databases. Examples of DBMSs are Oracle 10g, DB2, SQL Server and MySQL. A DBMS will support a particular database *paradigm*, for example relational, object-relational or object. A *database system* (DBS) is created, using a DBMS, to manage a specific database. The DBS includes any associated application software.

The term *OGSA Grid* refers to a system that complies with the Open Grid Service Architecture [13]. When the word *Grid* is used unqualified it will still refer to an OGSA Grid unless explicitly stated otherwise.

The term *metadata* denotes data that describe either data or a service, normally the former. Sometimes the application or system context determines which data are treated as metadata.

1.6. The structure of this paper

The rest of the paper is structured as follows. Section 2 considers the requirements placed upon database integration by OGSA Grid applications. The requirements for, and the architecture of,



OGSA-DAI are described in Section 3, in particular detailing one key aspect of OGSA-DAI—the way in which consumers make requests to an OGSA-DAI Grid Database Service. Section 4 describes this aspect in detail. Finally, the current status and conclusions are presented in Section 5.

The functionalities described in this paper are consistent with those provided by Release 3 of the OGSA-DAI software (July 2003) [14].

2. THE DATABASE REQUIREMENTS OF GRID APPLICATIONS

A typical application that accesses databases may consist of a computation that queries one or more databases and then carries out further analysis on the retrieved data. Therefore, database access should be seen as being only one part of a wider, distributed application. The databases themselves will usually exist independently of any one specific application, in geographically distributed locations. They will be operated under independent and autonomous regimes. If each regime is left to decide independently the access interface with which it presents each database to the OGSA Grid, then this has two immediate negative consequences. Firstly, code will need to be written that provides the OGSA Grid interface to their database. This has cost implications for the service provider making the database available to the OGSA Grid. Secondly, those writing OGSA Grid applications or data integration services will have to write code for each proprietary OGSA Grid database interface to which access is required. This has cost implications for those writing the applications. OGSA-DAI provides a substantial part of those requirements thereby avoiding these costs. The argument is, as with any other Grid infrastructure, that the higher cost of developing a well-engineered and generic solution is amortized across many applications, the net effect being both an overall reduction in costs and an overall gain in infrastructure quality.

OGSA-DAI addresses both of these issues by implementing a standard, service-based, interface through which databases can be presented to OGSA Grid applications. This standard interface is being specified by the GGF Database Access and Integration Services (DAIS) working group [15]. This approach allows those exposing their databases to the OGSA Grid to reuse code to achieve this, whilst consumers of those services can write code in a way that will be, as far as possible, independent of the specific regime under which the database is administered. In order to ensure that the design of OGSA-DAI achieves this aim, it is first necessary to consider the requirements that it must meet.

To integrate databases into OGSA Grid environments, two sets of requirements must be met: firstly, those that deal with the integration of databases into the Grid environment, and secondly those that deal with the functionality offered by databases to applications that must be available via the Grid. These two sets of requirements are considered in turn.

If computational and database components are to be seamlessly combined to create distributed applications, then a set of agreed standards will have to be defined and met by all components. While it is too early in the lifetime of the Grid to state categorically what all the areas of standardization will be, work on existing middleware systems (e.g. CORBA), and emerging research within the Global Grid Forum, suggest that security [16], accounting [17], performance monitoring [18] and scheduling [19] will be important.

Some Grid applications will have extreme performance requirements. In an application that performs CPU-intensive analysis on a huge amount of data accessed by a complex query from a DBS,



then achieving high performance may require utilizing high-performance servers to support the query execution (e.g. a parallel database server) and the computation (e.g. a powerful compute server such as a parallel machine, or a cluster of workstations). However, this may still not produce high performance, unless the communication between the query and analysis components is optimized. Different communication strategies will be appropriate under different circumstances. If all of the query results are required before analysis can begin, then it may be best to transfer all of the results efficiently in a single block from the database server to the compute server. Alternatively, if a significant computation needs to be performed on each element of the result set, then it is likely to be more efficient to stream the results from the DBS to the compute server as they are produced. When streaming, it is important to optimize communication by sending data in blocks, rather than as individual items, and to use flow control to ensure that the consumer is not swamped with or starved of data. Further, it may often be the case that the intended recipient of the result is not the consumer that issued the query. To avoid redundant transmission of data, it should be possible for the query issuer to specify a third-party destination for a result.

We now move beyond considering the requirements that are placed on all Grid middleware by the need to support databases, and consider the requirements that Grid applications will place on the DBMSs themselves. Firstly, there appears to be no reason why Grid applications will not require at least the same functionality, tools and properties as other types of database applications. Consequently, the range of facilities already offered by existing DBMSs will be required. These support both the management of data, and the management of the computational resources used to store and process those data. Specific facilities include the following:

- query and update facilities;
- programming interface;
- indexing;
- high availability;
- recovery;
- replication;
- versioning;
- schema evolution;
- uniform access to data and schema;
- concurrency control;
- transactions;
- bulk loading;
- manageability;
- archiving;
- security;
- integrity constraints;
- change notification (e.g. triggers).

We now consider whether Grid-enabled databases will have requirements beyond those typically found in existing systems. The Grid is intended to support the widescale sharing of large quantities of information. The likely characteristics of such systems may be expected to generate the following set of requirements that Grid-enabled databases will have to meet.



2.1. Metadata-driven access

It is widely recognized that Metadata will be very important for many Grid applications, in the following activities.

- Management or scheduling through provision of system and administrative information (e.g. defining the performance and capacity of a data resource, the charges and policies governing its use, and the current operational state).
- Data discovery or interpretation through provision of data structure and content information, e.g. the schema to which the data conforms and a summary of the content (values) available within that structure. This may be extended with generic information about the interpretation of the schema and content description, e.g. via ontologies [20].
- Resource or access method selection, through indexes or summaries. For example, a database of astronomical objects extracted by processing images may be used as a primary data source in its own right or as a means of selecting parts of images to reanalyse.
- Data selection or evaluation, to inform human judgements about the data. This varies from pervasive provenance data describing the origin of data to specific annotations suggesting interpretations and implications of the data.

Almost all aspects of metadata can have components that are application specific. Many disciplines are developing extensive bodies of practice concerning the organization and description of their data. Many applications involve portals, workflows or bespoke code that first examines metadata according to user requirements and then uses these metadata to locate the data, describe which data are accessed, determine what transformations are necessary, to steer analyses and visualizations, and to carry forward information into automatically generated metadata associated with result sets.

Currently, the use of metadata in Grid applications tends to be relatively simple—it is mainly for mapping the logical names for datasets into the physical locations where they can be accessed [11,21]. As users and developers develop more sophisticated applications, and as the Grid expands into new application areas such as engineering and the life sciences, more sophisticated metadata systems and tools will be required. The result is likely to be a *Semantic Grid* [22] that is analogous to the *Semantic Web* [23].

The use of metadata to locate data has important implications for integrating databases into the Grid because it promotes a two-step access to data [24]. In step one, a search of metadata catalogues is used to locate the databases containing the data required by the application. Those data are then accessed in the second step. A consequence of two-step access is that the application writer does not know the specific resource that will be accessed in the second step. Therefore, the application must be general enough to connect and interface to any of the possible resources returned in step one. This is straightforward if all interfaces are built from the same DBMS, and so offer the same interfaces to the application, but it is more difficult if these interfaces are heterogeneous.

Ideally the two-step approach requires that all resources should provide the same interfaces, but variation in facilities, interfaces and representations is inevitable. Therefore, the variation must be described in metadata that are available via introspection over the resource or service, so that software may be written that adapts to the variations. It also requires that all data are held in a common format, or that the metadata that describe the data are sufficient to allow applications to understand the formats and interpret the data.



A crucial issue is the production of metadata that are reliable. Today most metadata are produced manually. This is both labour intensive and error prone—it will not scale to the envisaged scale of applications and the envisaged number and size of data resources. It is therefore imperative that tools be developed to automatically generate metadata as systems and data collections are built, and to generate data, such as content summaries, by inspecting data resources. OGSA-DAI does not provide such tools but it clearly identifies their necessity and provides a context in which they may be deployed.

Once such metadata are in use, additional infrastructure components are required. Primary among these are registries that have matching algorithms relevant to application requirements that operate over these metadata.

2.2. Multiple database federation

One of the aims of the Grid is to promote the systematic sharing of scientific data. A recent study of the requirements of some early Grid applications concluded that ‘The prospect exists for literally billions of data resources, from large curated databases to individual collections, from major instruments such as the Large Hadron Collider and telescopes to portable instruments, medical scanners, networks of environmental sensors and wearable devices and petabytes of data being accessible in a Grid environment’ [25]. Here, data resources include sensors, instruments and wearable devices. If this prospect is realized then it is expected that many of the advances to flow from the Grid will come from applications that can combine information from multiple data sets. This will allow researchers to combine different types of information on a single entity to gain a more complete picture, and to aggregate the same types of information about different entities.

Achieving this will require support for integrating data from multiple data sources, for example through distributed query and transaction facilities. This has been an active research area for several decades, and needs to be addressed on multiple levels. As was the case for metadata-driven access, the design of federation middleware will be made much more straightforward if databases can be accessed through standard interfaces that hide as much of their heterogeneity as possible.

Even when APIs are standardized, there is a higher-level problem of the semantic integration of multiple databases, which has been the subject of much attention over the past decades [26,27]. In general, the problem complexity increases with the degree of heterogeneity of the set of databases being federated, though the provision of ontologies and metadata can assist. While there is much existing work on the federation on which to build, for example in the area of query processing [28], the Grid should give a renewed impetus to research in this area because there will be clear benefits from utilizing tools that can combine data over the Grid from multiple, distributed repositories. It is also important that the middleware that supports distributed services across federated databases meets the other Grid requirements. For example, distributed queries run across the Grid may process huge amounts of data, and so the performance requirements on the middleware may, in some cases, exceed the requirements on the individual data resources.

The Grid provides mechanisms with which to build trusted contexts across a number of independent members of a virtual organization. It also provides mechanisms for moving code as well as data between computational resources within the virtual organization, either relying on the established trust or using additional safety mechanisms. This means that computation may be moved close to data which often yields savings as code grows much less rapidly than data and the results of many analyses are much smaller than the data examined. Taking advantage of this code mobility poses new planning,



scheduling and optimization challenges which will be aided by a consistent model for computation and data services.

2.3. Summary

In summary, there are sets of requirements that must be met in order to support the construction of Grid applications that access databases. Some are generic across all Grid application components, while others are database specific. Grid applications will require at least the functionality provided by current DBMSs. As these are complex pieces of software, with high development costs, building new, Grid-enabled DBMSs from scratch is not an option. Instead, new facilities must be added by enhancing existing DBMSs, rather than by replacing them. The most commonly used DBMSs are commercial products that are not open-source, and so enhancements have to be achieved by wrapping the DBMSs externally. Later, we hope that DBMS vendors will choose to integrate Grid-enabled data access and integration mechanisms into their code for improved performance, safety and operational convenience.

The OGSA-DAI project is addressing these requirements by designing a standard, service-based interface through which databases are exposed to applications. The rest of this paper describes the work done to date in the project; designing a system to address those requirements that are perceived to be of the highest priority: querying, efficient results delivery, security and federation. Over time, the project aims to move on to addressing the other requirements identified in this section.

3. OGSA-DAI ARCHITECTURE AND FUNCTIONALITY

The goal of OGSA-DAI is to provide a uniform service interface for data access and integration to databases exposed to the Grid, hiding differences such as database driver technology, data formatting techniques and delivery mechanisms. The aim is that through OGSA-DAI, disparate, heterogeneous data sources and resources can be made available as services that are fully integrated with other OGSA services, for example, for data transport and security. OGSA-DAI services will therefore provide the basic operations that can be used by higher-level services to offer greater functionality, such as data federation and distributed queries.

The OGSA-DAI architecture also aims to encourage the design of efficient applications and evolvable services. To address the former, there is support for grouping multiple requests on an OGSA-DAI service into a single message sent to a service. This reduces latency by increasing the granularity of interactions, and reducing both the number of messages exchanged and the quantity of data transferred. To support service evolution, request messages are self-describing in that they identify the specific operation(s) to be called, so allowing a service to evolve to support new operations without requiring existing applications to be modified.

3.1. Architecture

In the rest of this section we give an overview of the OGSA-DAI architecture, and then explain in more detail how these specific requirements are addressed.

Figure 1 illustrates the principal components within the OGSA-DAI architecture. OGSA-DAI Grid Data Services access databases using drivers, and make use of several additional components,

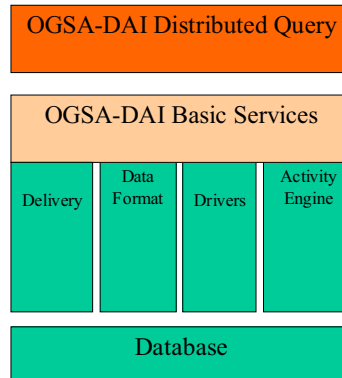


Figure 1. OGSA-DAI components.

for example, for data formatting, data delivery and request handling. Applications can use the core OGSA-DAI GDS components directly to access individual data stores, or can use a distributed query processor, OGSA-DQP [29], to coordinate access to multiple database services.

The important aspects of the architecture are discussed in the following subsections.

3.1.1. Standardized interfaces

A benefit that OGSA-DAI brings over and above paradigm-specific database connectivity models like JDBC is a consistent interface regardless of the underlying technology. For example, OGSA-DAI services support access to multiple relational database systems, XML storage managers and file systems. They provide consistent portType definitions and many shared functionalities independently of the database paradigm or specific system being accessed.

As well as query interfaces, OGSA-DAI will evolve to offer interfaces for other common data operations, including transactions and bulk load.

3.1.2. Metadata

OGSA-DAI services provide metadata about the DBMS, e.g. whether it is an Oracle, DB2 or MySQL, etc. DBMS system that are being exposed to the Grid. Also metadata are provided about the capabilities of the DBMSs that are being exposed to the Grid through the service interfaces as well as any inherent capabilities of the services themselves, such as third-party delivery and transformations that are available to operate on data retrieved from a database. The connection technology that is employed to connect to the databases is also exposed for clients capable of using such information. For relational databases, the database schema may be extracted from the service, which may be helpful to higher-level services such as distributed query processing. The metadata may be provided statically, that is when the service is configured, or dynamically, which may require additional coding. On the whole the static



metadata model is extensible so that communities that employ OGSA-DAI to access databases within a Grid context can provide community-specific metadata for the databases they expose to the Grid. The dynamic model can be extended, but this requires more effort from the developers of applications that want to dynamically extend the OGSA-DAI metadata framework.

3.1.3. Sessions

Sessions are important to database systems, for example to relate the queries sent by a consumer to the appropriate transaction context. Therefore, some way is needed to relate a set of requests sent by a consumer (or consumers) over a period of time, where they are part of the same data access session. A DBMS may concurrently support multiple client sessions. For example, DBMSs provide an interface by which clients can connect to the database and, at this point, a new session is created for that client. This allows, for example, multiple concurrent transactions to be run against the same DBMS. Therefore, OGSA-DAI provides a standard way for consumers to create a new session against a data resource. Whilst it would be possible to support this in a platform-independent manner (for example, by passing a session id with all requests to a database service), a decision was made to meet this requirement using specific mechanisms provided by the Open Grid Services Infrastructure version 1.0 (OGSI 1.0), as described in Section 1.

3.1.4. Security

It is important to be able to authenticate a consumer so that their request can be executed with the appropriate privileges. As yet there is no consensus within the Grid community as to how authentication should be done within virtual organizations. The aim, however, is to adopt whatever solution emerges from the standardization efforts currently being undertaken within the Web Services communities. Currently the status quo is to use slightly modified X509 certificates within the Grid community. Hence OGSA-DAI provides a simple authentication layer that consists of a *RoleMap* file and an implementation class that does the role mapping between the Grid credentials used to access the system, which are based on the X509 certificates, and the credentials required to access the underlying DBMS.

3.1.5. Grouping requests

OGSA-DAI supports a document-oriented interface for database requests, in which a single document may specify a collection of related activities (for example, an update to a database, a query of the updated database, and the delivery of the result of the query to a third party). This increases the granularity of interactions, reducing the number of messages and the quantity of data received. Section 4 describes in detail the structure and capabilities of request documents.

3.1.6. Data delivery options

In traditional client–server systems such as JDBC, a client sends a request (e.g. a query) to the database, and receives the result in return. This synchronous request model will not always meet the extreme performance and capacity requirements of some Grid applications. For these, it is important to support



a range of data delivery options. For example, the ability to directly deliver the data to a third party, that is to consume the result of the query, or to hold the result at the data service until it is required by another service.

OGSA-DAI supports a range of data-delivery options that can be exploited by Grid applications. Basic requests to an OGSA-DAI Grid Data Service are synchronous: the consumer sends a request document (Section 4) and in return receives a message containing the results. OGSA-DAI also offers asynchronous requests in which the result is not returned to the requestor: it may remain in the Grid Data Service until it is accessed by a consumer, or it may be sent directly to another service. A detailed description of the range of delivery options is presented in Section 4.

4. DATABASE REQUESTS

The OGSA-DAI *Grid Data Service* (GDS) supports a document-oriented interface for database requests, in which a single document may specify a collection of related activities (for example, an update to a database, a query of the updated database, and the delivery of the result of the query to a third party). This means that the GDS supports few operations—all accesses and updates to the state of a database are carried out by way of a *perform* operation, the parameter of which is a request document. In OGSA-DAI, a request consists of a collection of linked activities, where each activity represents an operation on the service, for example, to query the database, fetch some data from a remote site, etc. The document-based approach has been followed for several reasons. For example, a requirements analysis suggested that a unit of work with a remote database typically involves some database access, data format translation and delivery, and the document approach reduces the number of round trips required in a distributed setting. The request document thus supports composition of activities. It is now described in more detail.

4.1. Request definition and execution

The root element of a document, generally referred to as a perform document, sent to the *perform* operation is a `gridDataServicePerform` element. This contains the collection of activities to be performed by the GDS.

4.2. Activities

Each perform document may make use of one or more types of activity. Activities dictate the tasks that a GDS instance can perform on behalf of its client via its enactment engine. An activity is mapped to a Java class that is responsible for acting on or processing the content of that activity. This means that it is straightforward to add new kinds of activities to a GDS, either to provide generally useful functions (e.g. format transformations using an XSLT activity) or operations that are relevant to a specific kind of application (e.g. to compute the codon usage of a DNA sequence). We note, however, that the GDS activity model is not intended to develop into a full workflow language, but is intended to remain of limited expressiveness with a view to supporting common data access and transformation tasks.

The sequence of activity execution within a single request is defined by associating the input of one activity with the output of another activity to define a data flow. For example, Figure 2 illustrates

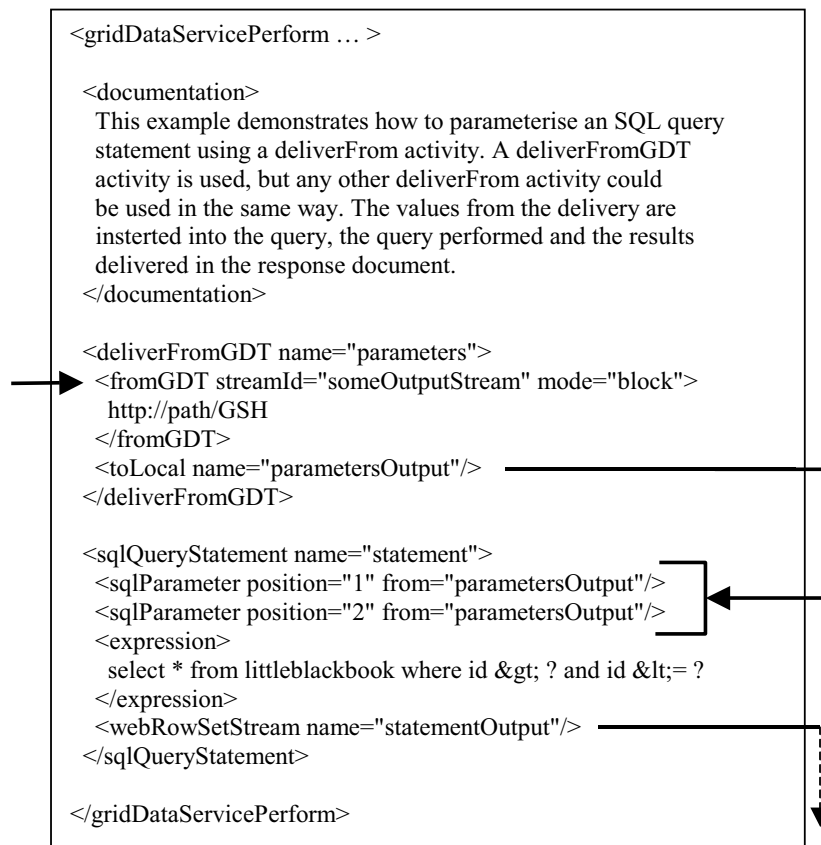


Figure 2. Activity inputs and outputs.

data flows among linked activities using arrows. The *parameter* element gives a name to a literal, which is used as an input to the *sqlQueryStatement*. Where a *parameter* element has no *value*, this must be supplied as an input to the request when it is executed. In this instance it is obtained from a delivery activity, *DeliverToGDT*, which in a GDS is another portType, the Grid Data Transport portType, allowing communication links to be established between a client and a GDS or a GDS and another GDS. The parameter values, in this case, could be supplied by the client or another GDS. The results are returned to the client synchronously in the response.

In support of the configurable and extensible nature of requests, concrete activity types supported by a GDS instance are listed in the *ActivityType* service data element (SDE) [30].

Sections 4.1 and 4.2 have described the framework within which concrete activities are defined. The following sections provide some examples of concrete activities.



4.2.1. Relational database activities

Activities that provide access to SQL queries, updates and stored procedures are supported. At the time of writing, services have been developed for accessing MySQL, Oracle and DB2 databases. An example *sqlQueryStatement* is shown below. The *sqlParameter* elements define values that are to be passed to the SQL expression. The first element references an activity output value named *chromvalue* from elsewhere in the request, while the second explicitly sets the value. The query is followed by the name by which the outputs are to be made available through the *webRowSetStream* element.

```
<sqlQueryStatement name="statement1">
  <sqlParameter position="1" from="chromvalue"/>
  <sqlParameter position="2">frog</sqlParameter>
  <expression>
    select * from chromFrag where chrom=? and organism=?
  </expression>
  <webRowSetStream name="statement1result"/>
</sqlQueryStatement>
```

4.2.2. XML database activities

Activities are provided for access to XML queries and updates. XPath queries and XUpdate are supported for implementation over the Xindice database [31]. The elements of *xPathStatementType* are an optional *collection* element, which identifies the collection over which the XPath query is to be applied; an optional *resourceId* element, which identifies a particular resource in a collection if required; an optional *namespace* element, which defines the namespaces that are used in the XPath expression; *expression* which contains the XPath expression; and *resourceSetStream* which is used to identify the result sequence.

The following shows a sample document for a simple query to an XML data service, retrieving information on chromosome fragments of length less than 20 000.

```
<xPathStatement name="chromStatement">
  <expression>/chromFrag[@len<20000]</expression>
  <resourceSetStream name="chromStatementOutput"/>
</xPathStatement>
```

The result of the query is made available to other activities through the *resourceSetStream* label, which is defined as an activity output. This output can then be used, for example, as input to a delivery, as described in the next section.

4.3. Delivery activities

The ability to deliver data to and/or from a GDS is an intrinsic part of this service. From the client's perspective, the operation of a GDS is either synchronous or asynchronous.

In the synchronous case, the client submits a perform document to *GDS::perform* and expects the results of the complete request in response to this call. Thus data delivery is considered to be part of the




Figure 3. Synchronous delivery of query results.

response to the *perform* operation call. The default delivery is to in-line the results in the response to the client unless instructed otherwise by an explicit delivery activity in the perform document. This is illustrated in Figure 3: a consumer sends a request containing a query to a GDS, and receives the results back directly.

The following *gridDataServicePerform* document to a relational GDS defines a request that queries a single table using SQL, and returns the result to the analyst in the response.

```
<gridDataServicePerform ... >
  <sqlQueryStatement name="statement">
    <expression>
      select * from chromFrag where len &lt; 20000
    </expression>
    <webRowSetStream name="statementresponse"/>
  </sqlQueryStatement>
</gridDataServicePerform>
```

In the asynchronous case, the client submits a perform document to *GDS::perform* and expects the GDS to complete the processing of the request after the *perform* operation has returned control to the client. This is particularly important when long running requests are submitted. Data delivery in this case follows either an active or a passive approach.

In the active case, the perform document must contain enough information for the GDS to interact with delivery protocols or services in order to deliver the data. In the passive case, the GDS waits for external services or protocols to initiate the data delivery. The delivery activities defined within OGSA-DAI are given in Table I. The underlying transport protocol used to support such an activity is provided by the *bind* element of the WSDL.

In Table I, *deliverToGDT* and *deliverFromGDT* involve the use of the *Grid Data Transport* (GDT) portType. This portType acts as an endpoint for data delivery, and must be supported by the Grid Service that is the target of this delivery activity. Thus, for example, a GDS service *S* can deliver to a remote service *R* using *deliverToGDT* only if *R* supports the GDT portType. *R* may or may not be a GDS. The *GridDataTransport* portType defines four operations *getFully*, *getBlock*, *putFully* and *putBlock*.



Table I. OGSA-DAI delivery activities.

Activity	Usage
<i>deliverFromGDT</i> <i>deliverFromGFTP</i> <i>deliverFromURL</i>	Define input configuration depending on the protocol chosen (GDT/GFTP/URL). Also define an output element that an internal activity can reference in order to obtain the data that this activity will retrieve.
<i>deliverToGDT</i> <i>deliverToGFTP</i> <i>deliverToURL</i>	Define an input element that references an output element in some other activity defined in the perform document to which the data will be delivered. The defined output configuration depends on the protocol chosen (GDT/GFTP/URL).
<i>inputStream</i>	Defines an output element that an internal activity can reference in order to obtain data when data are presented to this activity by an external service.
<i>outputStream</i>	Defines an input element that references an output element in one other activity defined in the <i>request</i> document. This provides the data that this activity will send in response to any request by an external service.

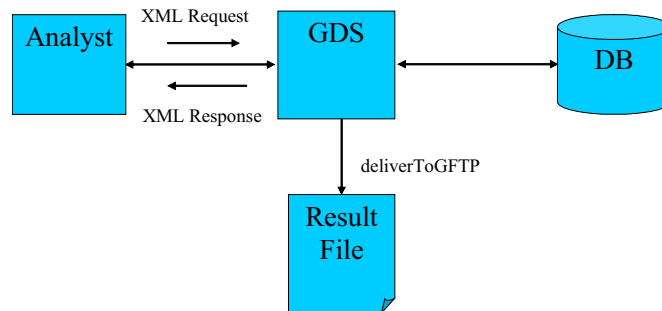


Figure 4. Asynchronous delivery of query results to a file.

GridFTP (GFTP) [32] is the principal data transport service used on the Grid. The activity type *deliverFromGFTP* allows data to be received by a GDS from an external location using GFTP, and *deliverToGFTP* allows data to be delivered by a GDS to an external location using GFTP.

The most straightforward asynchronous scenario for a GDS is that a requester executes a request containing a query to a GDS, the result of which is delivered to a file, in this case using GFTP. This scenario is illustrated in Figure 4.

The following perform document defines a request containing the same query as in Figure 3, but uses GFTP to deliver the result to a (potentially remote) file. The response returned to the analyst indicates that the request has been received, and is queued for evaluation. We note that previous database access middleware, such as JDBC, has tended not to provide this sort of indirect delivery capability.

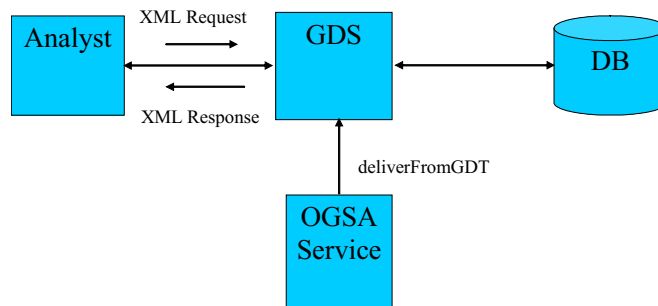


Figure 5. Asynchronous delivery of results from a service.

```
<gridDataServicePerform ... >

  <sqlQueryStatement name="statement">
    <expression>
      select * from chromFrag where len &lt; 20000
    </expression>
    <webRowSetStream name="statementresponse"/>
  </sqlQueryStatement>

  <deliverToGFTP name="d1">
    <fromLocal from="statementresponse"/>
    <toGFTP host="ogsdai.org.uk"
      port="8080" file="path/to/myfile.txt"/>
  </deliverToGFTP>

</gridDataServicePerform>
```

There are many situations in which a GDS may usefully interact directly with another service. For example, it may be useful for the results of a query over one GDS to be conveyed to another GDS by a replication infrastructure. By contrast, Figure 5 illustrates a scenario in which a GDS is used to update a database using data delivered to it from another service. In this example, it is assumed that the other service implements the GDT portType introduced in Section 4.3. The *deliverFromGDT* activity will invoke the *getFully* operation on the related service to obtain the required data for update. These data are then used as the parameter of an *sqlUpdateStatement* activity that inserts the retrieved value into the *chromFrag* table.



```

<gridDataServicePerform ... >

  <deliverFromGDT name="d1">
    <fromGDT streamId="otherrequestasynch/d1"
      mode="full">
      http://ogsadai.org.uk/GDTService/my/GDT/GSH
    </fromGDT>
    <toLocal name="datatoinsert"/>
  </deliverFromGDT>

  <sqlUpdateStatement name="statement">
    <sqlParameter position="1" from="datatoinsert"/>
    <expression>insert into chromFrag values ?</expression>
  </sqlUpdateStatement>

</gridDataServicePerform>

```

5. CONCLUSIONS

With a view to exploring how best to integrate databases and the Grid, the OGSA-DAI project has developed a collection of Grid Database Services, the key features of which are as follows.

- They provide consistent access to multiple database paradigms; this paper has described activities that support access to both Relational and XML DBSs.
- They support flexible data delivery, providing facilities for incremental and bulk delivery, to and from services or files, in a synchronous or an asynchronous manner.
- They are fully integrated with existing Grid functionalities for authentication and data transport.
- They are tracking the ongoing standardization activity within the DAIS Working Group of the Global Grid forum (<http://www.gridforum.org>).

The concepts embodied by the OGSA-DAI design are largely technology independent, assuming only that applications are built from services that communicate by exchanging messages. The first version made available was built on top of the OGSI 1.0 [30]. Experience with several early implementations of OGSI, and a desire to achieve greater synergy between Grid and Web Service developer communities, has resulted in a recent proposal to revise the mapping of OGSI onto Web Services. This is in the form of a set of specifications called WS-Resource Framework (WSRF) [33] which deliver a similar infrastructure, but with an explicit separation between stateful resources and the Web Services which present mechanisms for using those resources. As OGSA-DAI has a relatively simple operational model, it should be fairly straightforward to map it onto whatever emerges as the preferred way of building Grid applications. We are currently working on mappings to WSRF and WS-I.

Many of the advances to flow from the Grid will come from applications that can combine information from multiple data sets. This will allow researchers to combine different types of information on a single entity to gain a more complete picture, and to aggregate the same types



of information about different entities. Achieving this requires support for integrating data from multiple DBSs, for example through distributed query facilities. The design of federation middleware is made much more straightforward by OGSA-DAI providing standard interfaces that hide as much of the heterogeneity as possible. The OGSA-DQP [29] system builds a distributed query processing service on top of a set of OGSA-DAI database services. The query operations (filter, join, etc.) are distributed over nodes acquired from the Grid, so as to distribute the work of query evaluation in a scalable manner—the system can also exploit parallelism in the database operators and other operations invoked from a query, to provide scalable performance. OGSA-DQP exploits the standard OGSA-DAI interfaces to extract metadata about data held behind a service (e.g. its size, the availability of indices, etc.) for optimization. Having determined the best way to evaluate a query across each of the database services, it exploits the standard interfaces to send query requests to all the databases, while processing their responses is simplified by OGSA-DAI's common query result format. Without OGSA-DAI, the DQP service would have to be designed to interface directly with a set of heterogeneous DBMSs—a much harder task. Future development on the OGSA-DAI software is building on this, to produce data integration services which will bring together XML databases and semi-structured data held in files, as well as relational databases.

Overall, we believe that the OGSA-DAI system described in this paper is the most comprehensive attempt to date to provide generic database access capabilities within a Grid setting. The design seeks to provide enhanced functionality compared with other database middleware by supporting access to multiple database paradigms and by providing flexible data delivery both to and from the database services. The OGSA-DAI software (along with OGSA-DQP) is freely available in open-source form from <http://www.ogsadai.org.uk>.

Trademarks

IBM and DB2 are trademarks of International Business Machines Corporation in the United States, other countries, or both. Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both. Microsoft and SQL Server are trademarks of Microsoft Corporation in the United States, other countries, or both. Oracle is a trademark of Oracle Corporation. MySQL is a trademark of MySQL AB in the United States and other countries. Globus Project and Globus Toolkit are trademarks of the University of Chicago.

ACKNOWLEDGEMENTS

This work is supported by the U.K. e-Science Core Programme, the EPSRC and the DTI.

REFERENCES

1. The Globus Toolkit V2.4. <http://www.globus.org/gt2.4/download.html> [July 2004].
2. Dinda P, Plale B. A unified relational approach to Grid information services. *Technical Report GWD-GIS-012-1*, Global Grid Forum, 2001.
3. Booth D, Haas H, McCabe F, Newcomer E, Champion M, Ferris C, Orchard D. Web Services Architecture, W3C Working Group Note, 2004. <http://www.w3.org/TR/2004/NOTE-ws-arch-20040211> [July 2004].



4. Chatterjee S, Webber J. *Developing Enterprise Web Services: An Architects Guide to Best Practices*. Penguin: Harmondsworth, 2003.
5. Foster I, Kesselman C, Nick J, Tuecke S. Grid services for distributed system integration. *IEEE Computer* 2002; **35**(6):37–46.
6. Malaika S, Nelin CJ, Qu R, Reinwald B, Wolfson DC. DB2 and Web Services. *IBM Systems Journal* 2002; **41**(4):666–685.
7. Van Raalte T, Ward R. *Oracle 9i Application Server Web Services Developer's Guide*. Oracle Corporation, 2002.
8. Bell WH, Bosio D, Hoschek W, Kunszt P, McCance G, Silander M. Project spitfire—towards Grid Web Service databases. *Database Access and Integration Services Working Group, 5th Global Grid Forum (GGF5)*, Edinburgh, U.K., 2002. Global Grid Forum, 2002.
9. Madduri RK, Hood CS, Allcock WE. Reliable file transfer in Grid environments. *Proceedings of the 7th Annual IEEE Conference on Local Computer Networks (LCN)*, 2002; 737–738.
10. Stockinger H, Samar A, Holtman K, Allcock B, Foster IT, Tierney B. File and object replication in data Grids. *Proceedings of the 10th IEEE International Symposium on High Performance Distributed Computing (HPDC-10)*, 2001; 76–86.
11. Rajasekar A, Wan M, Moore R. MySRB & SRB: Components of a data Grid. *Proceedings of the 11th IEEE International Symposium on High Performance Distributed Computing (HPDC-11)*. IEEE Computer Society Press: Washington, DC, 2002; 301–310.
12. Foster IT, Vöckler J-S, Wilde M, Zhao Y. Chimera: A virtual data system for representing, querying, and automating data derivation. *Proceedings of the 14th International Conference on Scientific and Statistical Database Management (SSDBM)*, 2002; 37–46.
13. Foster I, Gannon D, Kishimoto H (eds.). *The Open Grid Services Architecture*, GGF OGSA Working Group, 2003. <http://forge.gridforum.org/projects/ogsa-wg/document/draft-ggf-ogsa-spec/en/13> [July 2004].
14. OGSA-DAI V3. <http://www.ogsadai.org/releases/ogsadai.php#R3> [July 2004].
15. Global Grid Forum—Database Access and Integration Services Working Group. <http://forge.gridforum.org/projects/dais-wg/> [July 2004].
16. Global Grid Forum—OGSA Security Working Group. <http://forge.gridforum.org/projects/ogsa-sec-wg/> [July 2004].
17. Global Grid Forum—Accounting Models Research Group. <http://forge.gridforum.org/projects/acct-rg/> [July 2004].
18. Tierney B, Aydt R, Gunter D, Smith W, Swany M, Taylor V, Wolski R. *A Grid Monitoring Architecture*. Global Grid Forum, 2002.
19. Global Grid Forum—Scheduling and Resource Management Area. <http://forge.gridforum.org/projects/srm/> [July 2004].
20. Wroe C, Stevens R, Goble CA, Roberts A, Greenwood M. A suite of Daml+Oil ontologies to describe bioinformatics Web Services and data. *International Journal of Cooperative Information Systems* 2003; **12**(2):197–224.
21. Ripeanu M, Forster I. A decentralized, adaptive, replica location service. *Proceedings of the 11th IEEE International Symposium on High Performance Distributed Computing (HPDC-11)*, Edinburgh, U.K., 24–16 July 2002. IEEE Computer Society Press: Washington, DC, 2002.
22. De Roure D, Jennings N, Shadbolt N. Research agenda for the Semantic Grid: A future e-Science infrastructure. *Report UKeS-2002-02*, U.K. National e-Science Centre, 2002.
23. Berners-Lee T, Hendler J, Lassila O. The Semantic Web. *Scientific American* 2001; (May).
24. Watson P. Databases and the Grid. *Grid Computing*, Berman, Fox, Hey (eds.). Wiley: New York, 2003; 363–384.
25. Pearson D. Data requirements for the Grid. *Scoping Study Report*, U.K. Grid Database Taskforce, 2002.
26. Seth AP, Larson JA. Federated database systems for managing distributed, heterogeneous and autonomous databases. *ACM Computing Surveys* 1990; **22**(3):183–236.
27. Lahiri T, Abiteboul S, Widom J. Ozone: Integrating structured and semistructured data. *Proceedings of the 7th International Workshop on Database Programming Languages*, Kinloch Rannoch, U.K., 1999. Springer: Berlin, 1999.
28. Yu CT, Chang CC. Distributed query processing. *ACM Computing Surveys* 1984; **16**(4):399–433.
29. Alpdemir N, Mukherjee A, Paton NW, Watson P, Fernandes AAA, Gounaris A, Smith J. Service-based distributed query processing on the Grid. *Proceedings of the 1st International Conference on Service-Oriented Computing (ICSOC)*. Springer: Berlin, 2003; 467–482.
30. Tuecke S, Czajkowski K, Foster I, Frey J, Graham S, Kesselman C, Maguire T, Sandholm T, Vanderbilt P, Snelling D. Open Grid Services Infrastructure (OGSI) Version 1.0. *Global Grid Forum Recommendation*, Global Grid Forum, June 2003.
31. Apache Xindice. <http://xml.apache.org/xindice/> [July 2004].
32. Globus GridFTP. <http://www.globus.org/datagrid/gridftp.html> [July 2004].
33. Ferguson D, Foster I, Frey J, Graham S, Leyman F, Nally M, Storey T, Tuecke S, Weerawaranna S. Modeling stateful resources with Web Services. *IBM Systems Journal* 2004; **43**(4).